

Survey paper

When meaning meets purpose: A survey on semantic and task-oriented communications

Quang Tuan Do^a, Tung Son Do^a, Thanh Phung Truong^a, Dongwook Won^a, Nhu-Ngoc Dao^c,
Wonjong Noh^{b,*}, Sungrae Cho^{a,*}

^a Department of Computer Science and Engineering, Chung-Ang University, Seoul, 06974, South Korea

^b School of Software, Hallym University, Chuncheon, 24252, South Korea

^c Research Institute for Data Science and Artificial Intelligence, Hallym University, Chuncheon, 24252, South Korea

ARTICLE INFO

Keywords:

Semantic communication
Task-oriented communication
Deep learning
Joint source-channel coding
Large language models
Sixth-generation wireless networks

ABSTRACT

As wireless systems evolve toward sixth-generation (6G) networks, traditional bit-centric approaches cannot support intelligent applications that require meaning-aware and effectiveness-aware transmission. This survey examines semantic communication (SC) and its specialized subset, task-oriented semantic communication (ToSC), which prioritizes task-specific utility over general semantic completeness. Unlike prior surveys that treat this distinction descriptively, we synthesize the information bottleneck framework with empirical evidence from the surveyed literature to characterize the qualitative regimes under which SC and ToSC diverge or converge, and to distill design heuristics. An architectural taxonomy along functional axes (encoding strategy, transmission mode, task-relevance mechanism, and adaptivity) captures the recurring mechanisms shared across implementations. We further analyze the deep learning model families (autoencoders, convolutional and graph neural networks, Transformers, and large language and multimodal models) that implement semantic extraction, compression, and task reasoning. A structured evaluation framework addresses metric definitions and the conflicts and trade-offs between semantic fidelity, task effectiveness, and communication efficiency. Across ten application domains, we identify domain-specific risk factors, unresolved limitations, and four cross-domain weakness patterns: knowledge dependency as the most pervasive bottleneck, weak-baseline comparisons that inflate reported gains, brittleness in safety-critical deployments, and inconsistent evaluation methodology that prevents meaningful cross-domain comparison. Open challenges are consolidated into a dependency-ordered research roadmap that prioritizes foundational theory, standardized benchmarks, and scalable knowledge base management. This survey positions SC as the overarching paradigm for 6G information-directed communications, with ToSC as its action-driven specialization, and offers practical criteria for deciding when each applies and how to evaluate it.

1. Introduction

1.1. Background and motivation

Sixth-generation (6G) wireless systems are expected to support applications that impose stringent requirements on latency, bandwidth, and energy efficiency, ranging from extended reality and autonomous systems to large-scale Internet of Things (IoT) and smart healthcare (Dang et al., 2020; Saad et al., 2019; Popovski et al., 2021). Traditional communication systems, designed for accurate bit-level transfer, are ill-suited to these demands: optimizing for symbol fidelity while disregarding the meaning, context, or task-relevance of information leads to redundant transmissions and inefficient resource usage in intelligent,

context-aware systems (Xin et al., 2024; Shen et al., 2023; S. Zhang et al., 2023). This qualitative gap, where lowering bit error rates does not improve semantic relevance or downstream effectiveness, motivates rethinking the design objective: communication systems should be aware of semantics and their pragmatic use (Strinati and Barbarossa, 2021; Yang et al., 2022).

1.2. Semantic communication and task-oriented specialization

Semantic communication (SC) redefines the core objective of communication, moving from the accurate replication of symbols from

* Corresponding authors.

E-mail addresses: dqtuan@uclab.re.kr (Q.T. Do), tsdo@uclab.re.kr (T.S. Do), tptruong@uclab.re.kr (T.P. Truong), dwwon@uclab.re.kr (D. Won), nddao@hallym.ac.kr (N.-N. Dao), wonjong.noh@hallym.ac.kr (W. Noh), srcho@cau.ac.kr (S. Cho).

<https://doi.org/10.1016/j.engappai.2026.116063>

Received 23 March 2026; Received in revised form 18 July 2026; Accepted 18 August 2026

Available online 31 August 2026

0952-1976/© 2026 Elsevier Ltd. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

source to destination toward conveying the sender's intended meaning or semantic content. Although the concept has been explored for decades, recent developments in artificial intelligence (AI) and machine learning offer advanced tools required to extract, represent, and interpret meaning from complex and heterogeneous data, reigniting interest in SC and accelerating its progress. This approach draws inspiration from the efficiency of human communication, which emphasizes intent and context (Niu et al., 2022; Wang et al., 2025). SC shifts the design question from "how to transmit data" to "what to transmit": by selectively conveying only semantically relevant information, SC reduces redundant communication, conserves bandwidth, and lowers latency (Lu et al., 2022; Yang et al., 2022).

Within this broader paradigm, task-oriented semantic communication (ToSC) is a specialization of SC in which the communication objective is to enable a specific task outcome at the receiver. Success is therefore measured by task-level performance metrics, such as classification accuracy, detection probability, control reward, decision latency, or mission success rate rather than symbol- or bit-level fidelity. ToSC extracts and transmits only the semantic information that is causally useful for the task under resource constraints, while suppressing irrelevant semantics. Early formulations include the multi-user visual question answering (VQA) system of Xie et al. (2021a) and the personalized saliency encoder of Kang et al. (2022); broader treatments of the ToSC design space appear in Getu et al. (2024), Gündüz et al. (2023).

Consider wireless image transmission. Under general reconstruction-oriented SC, the receiver aims to reconstruct a perceptually faithful image for human viewing or broad reuse; background and fine-grained texture must be preserved, and success is measured by perceptual fidelity (e.g., PSNR/SSIM) and quality-of-experience. Under action-oriented ToSC, the receiver's goal may be classification or detection; only class- or target-discriminative features are essential, permitting the encoder to discard details that do not affect task performance. Communication success is measured by task metrics (e.g., top-1/top-5 accuracy, F1-score, false-alarm/miss rates) and end-to-end decision latency under bandwidth/energy constraints (Getu et al., 2024; Gündüz et al., 2023). This contrast clarifies that general SC prioritizes reusable meaning, whereas ToSC prioritizes pragmatic effectiveness for a specific action.

This meaning- and action-centric perspective calls for theoretical frameworks and metrics that extend beyond bit error rate to include semantic fidelity for general SC and pragmatic success for ToSC. In this survey, we treat SC as the overarching paradigm guiding information-directed communications for 6G and beyond, and ToSC as a key specialization tailored to action-driven systems.

1.3. Contributions and organization

This survey provides a systematic examination of SC with specialized emphasis on ToSC, presenting both theoretical foundations and practical implementations. The key contributions of this work are as follows:

1. **Theoretical Foundations Traced to Their Design Consequences:** We trace the theoretical lineage showing that knowledge-base mismatch degrades semantic channel capacity independently of physical channel quality, synthesizing the progression from classical information theory through semantic and strongly semantic information theories, Bao et al.'s semantic communication theory, and goal-oriented communication, and connecting each phase to its practical implications for modern ToSC system design.
2. **A Unified SC/ToSC Framework Stating When the Distinction Matters:** Unlike prior surveys that treat the SC/ToSC distinction as a labeling convention, we synthesize the information bottleneck (IB) framework (originating with Tishby et al., 2000,

mapped to SC by Beck et al., 2023) with empirical evidence from the surveyed literature to characterize the qualitative regimes under which SC and ToSC diverge or converge. Concrete examples demonstrate the conditions under which SC-optimized systems underperform at tasks and vice versa, and three design principles translate these conditions into practical engineering criteria.

3. **Evaluation Metrics and Their Conflicts:** Metrics are organized by Weaver level (Level B fidelity for SC, Level C effectiveness for ToSC, plus cross-level efficiency metrics). We analyze the conflicts and trade-offs between metric categories, including Level B vs. Level C, effectiveness vs. efficiency, and efficiency vs. fidelity, and provide practitioner guidance for metric selection under competing objectives.
4. **Ten Domains Distilled into Four Weakness Patterns:** From ten application areas we distill four recurring weakness patterns not systematically documented in prior surveys (developed in Section 5 and mapped in Fig. 7), each domain carrying a critical assessment of limitations, negative results, and unresolved challenges.
5. **Research Trends and Prioritized Research Directions:** Five research priorities, ranked by the number of downstream barriers each unblocks, distill open challenges from research trends spanning system intelligence (large language model/large multimodal model (LLM/LMM) integration, intent-based and explainable systems), system architecture (multi-user, cooperative agents, edge computing), and physical infrastructure (task-aware design, reconfigurable intelligent surface (RIS)-assisted systems).

These contributions differentiate this survey from prior SC/ToSC surveys along the axes made explicit in Table 1. Rather than cataloging schemes by application domain or by neural-network backbone (the organizations that dominate earlier surveys), we organize implementations along four *functional axes* (encoding strategy, transmission mode, task-relevance mechanism, and adaptivity) and overlay a complementary AI-model-family lens (Section 3.4), so that recurring design mechanisms become visible across otherwise dissimilar systems. Moreover, we ground the SC/ToSC distinction in the information-bottleneck framework and the surveyed empirical evidence to state *when it matters*.

Among the most closely related recent surveys, the survey of semantic radio access networks (Meng et al., 2026) offers the most systematic treatment of how semantic communication integrates into the radio access network and protocol stack, but confines its analysis to that layer and so neither derives when the SC/ToSC distinction matters nor analyzes inter-metric conflict. The survey of generative diffusion for wireless (Fan et al., 2026) catalogs diffusion-model techniques across the wireless protocol stack more comprehensively than any other, yet treats diffusion as a model family rather than as a decoding strategy spanning Weaver Levels B and C (Section 6.1). The dedicated survey of generative-AI-enabled SC by Liang and Li (2026) likewise indexes the field by generative architecture, communication modality, and application task, not by the communication objective the generated output serves. The contemporary-SC survey of Nguyen et al. (2026) comes closest by aggregate coverage in Table 1 and spans the broadest range of emerging technique families, from theory of mind and generative AI to joint source-channel coding (JSCC); its organizing axis, however, remains the technique itself, leaving open the condition-based synthesis of *when* SC and ToSC diverge and the inter-metric conflict analysis developed here. This combination of a structured cross-survey coverage matrix, a mechanism-level architectural taxonomy, and a critical, condition-based synthesis is, to our knowledge, not jointly offered by existing surveys.

The remainder of this paper is organized as follows (see Fig. 1 for the overall structure). Section 2 presents the theoretical evolution from classical information theory to goal-oriented communication. Section 3 introduces the ToSC framework, including a critical analysis




Fig. 1. Overall structure of this survey.

of when the SC/ToSC distinction matters. Section 4 establishes the evaluation framework with metric conflict analysis. Section 5 examines ten application domains with per-domain critical assessments and a cross-domain synthesis. Section 6 surveys research trends. Section 7 discusses challenges and presents prioritized research directions. Section 8 concludes with cross-cutting findings. To contextualize these contributions, Table 1 identifies the specific analytical gaps in prior surveys that this work addresses, focusing on the *types of analysis* that remain missing, not topical coverage breadth. We also distinguish our scope from the complementary view of communication as part of integrated sensing, computation, and communication: whereas recent surveys of edge perception organize the field around how the network edge senses and processes its physical environment (Lyu et al., 2026), this survey centers on the communication objective itself (what meaning or task outcome the transmitted representation must preserve), treating sensing and computation as sources of semantics rather than the organizing axis. Our review covers the semantic- and task-oriented-communication literature available through mid-2026.

Notation. Throughout this paper, X denotes the source data, Z denotes the compressed semantic representation, T denotes the task variable (e.g., a class label), and Y denotes the channel output. $I(\cdot;\cdot)$ denotes mutual information, $H(\cdot)$ denotes Shannon entropy, and $H(\cdot|\cdot)$ denotes conditional entropy, all in the standard information-theoretic sense; $\bar{H}_s(\cdot)$ denotes Bao et al.'s model-theoretic *semantic* entropy (the average logical information of a message under a given knowledge base and inference procedure), distinct from the Shannon entropy $H(\cdot)$. In the IB formulation (Section 3.6.1), $\beta > 0$ controls the compression-relevance trade-off. K_s and K_d denote the source and destination

knowledge bases, respectively. $\langle W, K, I, M \rangle$ refers to Bao et al.'s 4-tuple system model (world model, background knowledge, inference engine, message interpreter) introduced in Section 2; the bare I in this tuple is Bao et al.'s inference component, distinct from the mutual-information operator $I(\cdot;\cdot)$ above.

2. Theoretical evolution: From classical information theory to task-oriented principles

The development of ToSC builds upon a theoretical lineage that progressively extended communication theory beyond symbol transmission toward meaning and purpose. The semiotic foundations of Morris (1938) distinguish three dimensions of signs: syntactics (formal structure), semantics (meaning conveyed), and pragmatics (effect on interpreters), which map directly onto Weaver's three communication levels (Shannon and Weaver, 1998), illustrated in Fig. 2:

- **Level A (Technical):** Accuracy of symbol transmission, the domain of Classical Information Theory (CIT) (Shannon, 1948).
- **Level B (Semantic):** Precision with which transmitted symbols convey intended meaning.
- **Level C (Effectiveness):** Impact of received meaning on achieving desired behavior (Weaver, 1953).

CIT provides the mathematical foundation for Level A but explicitly abstracts away Levels B and C. This abstraction is adequate for bit-pipe communication but creates a fundamental mismatch for intelligent 6G applications where the *relevance* and *utility* of transmitted information

Table 1
Positioning of this survey relative to prior work.

Survey	Theory	Architecture	Metrics	Applications	Critical synthesis	Primary focus
Y. Liu et al. (2024)	◦	◦	◦	•	×	Broad SC overview
Lu et al. (2023)	•	•	◦	◦	×	SC enabling-technique taxonomy
Wheeler and Natarajan (2023)	◦	•	×	◦	×	Context-based design framework
Getu et al. (2023)	◦	×	•	◦	×	SC/GOC performance metrics
Guo et al. (2024a)	◦	•	×	◦	◦	SC security threats and countermeasures
Xin et al. (2024)	•	×	◦	×	×	Mathematical foundations
Nguyen et al. (2026)	•	•	◦	•	◦	Contemporary SC (Theory of Mind, generative AI, JSCC)
Meng et al. (2025)	×	◦	×	◦	◦	Secure SC across the life cycle
Meng et al. (2026)	◦	•	×	◦	×	Semantic radio access networks
Fan et al. (2026)	◦	•	×	•	×	Generative diffusion for wireless
Liang and Li (2026)	◦	•	◦	•	◦	Generative-AI-enabled SC
Getu et al. (2024)	•	◦	◦	◦	◦	Goal-oriented SC: theory, techniques, and challenges
Gündüz et al. (2023)	•	•	◦	◦	◦	Context/semantics/task-oriented foundations (tutorial)
This work	•	•	•	•	•	SC/ToSC distinction conditions, metric-conflict analysis, cross-domain synthesis, KB stress-test, prioritized roadmap

Coverage (operational scoring, per aspect): • = a dedicated treatment with comparative synthesis across multiple works; ◦ = the aspect is addressed but only descriptively or for isolated schemes; × = absent or mentioned only in passing. “Critical synthesis” denotes the presence of explicit cross-work comparative analysis organized around stated trade-off axes (rather than per-topic description), independent of which specific trade-offs are examined; a • rating additionally requires that this comparative synthesis span the majority of the four content dimensions (theory, architecture, metrics, and applications) rather than a single sub-area. The • cross-work analyses scored for This work are located in the theoretical-lineage synthesis (Section 2), the architectural taxonomy along stated functional axes (Section 3.3, Table 3), the metric-conflict analysis (Section 4), the cross-domain weakness synthesis (Section 5, Fig. 7), and the prioritized roadmap (Section 7); all scores reflect the authors’ assessment of each comparator’s cited version.




Fig. 2. Weaver’s three levels of communication (Level A: Technical, Level B: Semantic, and Level C: Effectiveness). The mapping of these levels to the performance metrics is discussed in Section 4; ToSC primarily targets Level C.

matter more than its bit-level fidelity. Bridging this gap has been the central theoretical challenge driving SC research, progressing through three distinct phases: quantifying semantic content (Level B theories), modeling semantic exchange in communication systems, and formalizing goal-directed communication (Level C theories). We summarize these phases below and compare them in Supplementary Table S1; readers seeking detailed mathematical treatments are referred to Xin et al. (2024), Lu et al. (2023). Fig. 3 traces this progression across the three conceptual phases, mapping each theoretical milestone to its enabling contribution in modern ToSC practice.

2.1. Phase 1: Quantifying semantic content (Level B)

The earliest attempts to formalize semantic information (SI) emerged from the Semantic Information Theory (SIT) of Bar-Hillel and Carnap (1953), which defined SI as the logical content implied by a statement within a formal language model. A statement that narrows the set of possible states carries more SI than a vague one, analogous to Shannon entropy but operating over logical rather than probabilistic spaces. However, SIT suffered from the Bar-Hillel–Carnap paradox: contradictions, being maximally “informative” in a logical sense, received the highest SI values despite being false (Bar-Hillel and Carnap, 1953).




Fig. 3. Theoretical evolution of semantic communication foundations. Five milestones trace the progression from Shannon’s classical information theory (Level A, 1948) through semantic information quantification (SIT, SSIT; Level B, 1953–2004) and system-level semantic exchange modeling (Bao et al.; Level B, 2011), to goal-oriented communication theories (Level C, 2008–2021). Each forward-link box indicates how the historical contribution enables modern ToSC practice. Phase braces correspond to the three subsections of Section 2. The horizontal axis orders milestones by theoretical progression and is *not* drawn to a linear time scale; the year labels denote approximate periods, and goal-oriented communication theories (2008–2021) overlap the 2011 system-level work.

The Strongly Semantic Information Theory (SSIT) of Floridi (2004) resolved this by incorporating truth values directly into SI quantification. SSIT defines a discrepancy function $f(s)$ measuring how far a statement s deviates from the actual state of affairs, and derives informativeness as $g(s) = 1 - f(s)^2$. This formulation ensures that false statements carry negative informational value, aligning SI with the practical intuition that only truthful content supports effective decision-making. Subsequent extensions incorporating “truthlikeness” (D’Alfonso, 2011; Niiniluoto, 2012) further refined the granularity of this analysis. Despite these advances, both SIT and SSIT remain limited to logical propositions and do not address how SI propagates through a communication system, a gap that motivated the next theoretical phase. However, the core insight, that information content should be measured by its *semantic* rather than statistical properties, directly underpins modern SC system design: the sentence-similarity metrics used to train and evaluate systems such as DeepSC (Xie et al., 2021b) are practical descendants of SIT’s attempt to quantify meaning, replacing formal logic with learned neural embeddings (see Section 4).

2.2. Phase 2: Modeling end-to-end semantic exchange

Bao et al. (2011) moved beyond individual-statement semantics to model SC as a system. Their framework characterizes both source and destination by a 4-tuple $\langle W, K, I, M \rangle$ (world model, background knowledge, inference engine, and message interpreter) and defines semantic channel capacity as follows:

$$C_{sem} = \sup_{P(X|W)} [I(X; Y) - H(W|X) + \bar{H}_s(Y)] \quad (1)$$

where $I(X; Y)$ is the Shannon mutual information between the sent and received messages (the syntactic channel term, as in CIT); $H(W|X)$ is the semantic-encoder *equivocation*, the residual ambiguity of the world model W given the transmitted message under the sender’s knowledge base and inference procedure; and $\bar{H}_s(Y)$ is the average model-theoretic *semantic* information the receiver can recover from Y under its own knowledge base and inference procedure, distinct from the Shannon output entropy $H(Y)$. Because this interpretive term enters additively, Bao et al.’s Semantic Channel Coding Theorem yields a capacity that may lie *above* or *below* the classical capacity $\sup_{P(X|W)} I(X; Y)$, according to whether the receiver’s semantic-interpretation gain $\bar{H}_s(Y)$ or the encoder equivocation $H(W|X)$ dominates. This capacity depends on the alignment between the source’s and destination’s knowledge bases (K_s vs. K_d): when $K_s \neq K_d$, semantic noise arises that has no counterpart in classical channel noise. The insight that *knowledge mismatch degrades semantic capacity independently of physical channel*

quality is foundational for understanding the knowledge-base dependency that pervades modern ToSC systems (discussed in Section 3). However, the framework remained conceptual, relying on idealized assumptions (e.g., $K_s = K_d$) that are rarely satisfied in practice; moreover, the world model W is assumed to be finite and enumerable, which is unrealistic for continuous sensor data in applications such as autonomous driving or medical imaging. Güler et al. (2018) subsequently formulated meaning recovery over a noisy channel as a Bayesian game of incomplete information, in which an external agent of unknown (adversarial or helpful) intent shapes the receiver’s interpretation, establishing game-theoretic foundations for robust semantic exchange. This Phase 2 insight, that knowledge-base alignment is a system-level design parameter as important as channel quality, manifests directly in modern ToSC architectures: the knowledge-base synchronization mechanisms surveyed in Section 3 (e.g., the curriculum learning of Farshbafan et al. (2023) and the knowledge graph mapping of Zhou et al. (2022)) are practical attempts to satisfy the $K_s = K_d$ assumption that Bao et al.’s capacity formula requires.

2.3. Phase 3: The pragmatic turn: Goal-Oriented Communication (Level C)

The transition to Level C required a fundamental reconceptualization: communication as a means to achieve goals rather than to transfer meaning. Juba and Sudan (2008) demonstrated that two agents without a shared language can still communicate successfully if the receiver can verify whether its goal has been achieved, a result limited to problems in the PSPACE complexity class but conceptually significant. Their framework was subsequently generalized into Goal-Oriented Communication (GOC) theory (Juba and Sudan, 2009; Goldreich et al., 2012), which distinguishes *meta-goals* (internal intentions) from *syntactic goals* (observable outcomes) and proves that goal achievement is possible even without shared linguistic structure, provided agents are cooperative and can sense progress toward their objectives. A key theoretical result established that constructing universal communication strategies under this framework is equivalent to designing online learning algorithms (Juba and Vempala, 2011), forging a direct link between communication theory and machine learning.

Kountouris and Pappas (2021) further bridged these theoretical phases by introducing a semantics-empowered framework for networked intelligent systems that integrates timing, context, and goal-awareness into communication design. This progression, from SIT’s logical semantics through system-level SC models to GOC’s task-driven formulation, provides the theoretical justification for ToSC as a design philosophy. Where classical SC optimizes for semantic fidelity at Level B, ToSC inherits GOC’s principle that communication success is measured solely by whether the receiver’s task is accomplished.

However, these foundational works largely relied on trial-and-error communication paradigms and offered guidance for designing protocols or strategies for simple communication systems, such as server-printer interactions; their focus remained on a mathematical theory of GOC for traditional computer models. Despite their limited scope, these early GOC systems marked a milestone: the ultimate measure of communication success is task achievement, not symbol accuracy. Modern ToSC systems operationalize this principle through learned goal verification. For example, the personalized saliency framework of Kang et al. (2022) lets the receiver's task define feature relevance, while the IB-based edge inference of Shao et al. (2021) jointly trains encoder and task head, directly realizing Juba's theoretical link between learning algorithms and communication strategies.

3. ToSC framework: A specialization of SC

3.1. Defining ToSC

Traditional communication systems, grounded in Shannon's theory, prioritize reliable bit transmission while often neglecting semantic content or task-specific data utility, an approach increasingly insufficient for modern applications that prioritize informed machine decisions over mere data relay, such as autonomous driving, healthcare monitoring, and industrial automation. ToSC addresses these shortcomings by transmitting only task-relevant semantic content rather than all available data, acknowledging the varying value of information across applications and tailoring data to tasks for more effective resource allocation and communication efficiency (Gündüz et al., 2023).

ToSC has two primary objectives: effectiveness and resource conservation. Effectiveness ensures task success, defined by task-specific outcomes such as high classification accuracy or reliable robotic control, with emphasis on the receiver's functional outcome. The second, often the primary driver of ToSC adoption, is communication efficiency: reducing transmitted data by conveying only task-relevant semantics lowers bandwidth, latency, and potentially energy consumption, important for resource-constrained environments such as IoT and mobile settings (Sagduyu et al., 2024).

Distinguishing ToSC from general SC clarifies its unique focus, despite their close relationship: Fig. 4 and Table 2 illustrate the workflow and paradigm differences. General SC aims to transmit meaning accurately and efficiently so the receiver's interpretation aligns with the sender's intent, whereas ToSC is a pragmatic specialization emphasizing task-specific utility: to enhance efficiency, it may intentionally discard semantically valid but task-irrelevant meaning that general SC, designed to preserve meaning faithfully, would retain (Zhou et al., 2024). For example, an image classification ToSC system may transmit only features discriminative for the target classes, ignoring background or texture details that a general semantic system focused on reconstruction may preserve. Goal-oriented communication, central to ToSC, inherently involves both understanding meaning and achieving a specific objective, so the receiver's primary concern is effectiveness in accomplishing the predefined goal rather than semantic completeness alone (Getu et al., 2024).

A defining characteristic of ToSC is its intrinsic reliance on shared knowledge accessible to both the transmitter and receiver (Bao et al., 2011; Strinati and Barbarossa, 2021). The term "knowledge base" has been applied loosely in the literature to everything from ontologies to pre-trained model weights, which obscures important distinctions. We therefore separate three layers of shared knowledge and reserve the term *knowledge base* for the first two:

- **Static knowledge base (KB):** fixed, explicitly shared, and addressable knowledge agreed upon a priori, such as ontologies, codebooks, label sets, and shared datasets or embeddings.
- **Dynamic context:** session- or runtime-state synchronized between endpoints during operation, such as recent observations, dialogue history, or an evolving task state.

- **Implicit (parametric) knowledge:** knowledge embedded in model weights (pre-trained or fine-tuned encoders, and large language and multimodal models). We treat this as a *carrier* of knowledge rather than a knowledge base in itself: a model is a source of implicit knowledge that can *back* or *substitute for* an explicit KB, for instance through retrieval-augmented generation, but it is not directly addressable or jointly editable in the way a shared KB is.

Across these layers, shared knowledge (1) establishes the context that makes "task-relevant semantics" well-defined, (2) reduces transmission volume by letting communication focus on information not already inferable from shared knowledge, and (3) enables intelligent encoding and decoding at both ends of the link (Guo et al., 2024a). We use this three-layer terminology consistently throughout the survey; the interplay between implicit parametric knowledge and explicit KBs in large-model-based systems is examined in Section 6.

3.2. Functional components and system building blocks

ToSC systems require a distinct set of functional blocks compared to traditional architectures, reflecting their emphasis on meaning and task execution (Gündüz et al., 2023; Barbarossa et al., 2023). Fig. 5 presents a consolidated system model with IB notation. While specific implementations may vary, several core components are typically present.

3.2.1. Task-tailored semantic encoders and decoders

At the core of a ToSC framework are the semantic encoder and decoder, which replace or augment the conventional source and channel coding blocks. The semantic encoder analyzes the source data, often using contextual information or a KB, to extract task-relevant features and prepare them for transmission, typically as a compressed representation such as a latent vector.

The semantic decoder receives the (potentially noisy) transmitted information and interprets it, often using shared knowledge or contextual cues; its function depends heavily on the task, ranging from directly performing the task (e.g., classifying the input based on the received semantic features) to reconstructing a representation of the source data that contains the information required by a downstream task execution unit.

These encoders and decoders are task-tailored, with design and training optimized specifically for the intended goal. Deep neural networks (DNNs) commonly implement these components for their ability to learn complex mappings and extract salient features from data, for example convolutional autoencoders (CAEs) for image-related tasks (Bourtsoulatze et al., 2019), convolutional neural networks (CNNs) for feature extraction and classification (Lyu et al., 2024; Eldeeb et al., 2024), and Transformer architectures for sequential or relational data such as text or multimodal inputs (Xie et al., 2021b, 2022). A key feature of several modern ToSC systems is JSCC and end-to-end training, where the semantic encoder, channel transmission (or its model), semantic decoder, and task execution module are jointly optimized as a single system to maximize task performance under given communication constraints (Bourtsoulatze et al., 2019; J. Xu et al., 2023); this end-to-end deep JSCC formulation is examined in detail in Section 3.3.

3.2.2. Role of KBs and context acquisition

Unlike traditional, content-agnostic systems, ToSC systems frequently rely on shared knowledge or context between transmitter and receiver: this understanding, often formalized as a KB, is essential for both extracting relevant semantics at the transmitter and correctly interpreting them at the receiver (Xin et al., 2024). The KB can take various forms depending on the application, including background information, ontologies defining terms and relationships, shared datasets used to train the semantic codecs, base models for rendering (such as an




Fig. 4. Comparative study of traditional models, SC systems, and ToSC systems: workflow of (a) traditional communication, (b) SC, and (c) ToSC. In (a) every bit must be recovered, so channel errors corrupt the delivered message. In (b) the semantic extractor and reconstructor exchange a learned semantic representation, and the receiving system recovers the intended meaning despite bit errors. In (c) the ToSC pipeline conveys only the semantics the receiving system needs to execute the task (switching the light on), rather than the full semantic content. A shared knowledge base supports semantic interpretation in both paradigms; as marked in the figure, it is optional, but often required, for general SC (dashed) and essential for ToSC's task relevance and interpretation (solid; Table 2).

Table 2
Comparison of communication paradigms.

Feature	Traditional (Shannon)	General Semantic Communication	ToSC
Primary Goal	Accurate symbol/bit recovery	Faithful conveyance of meaning (semantics)	Effective task completion/Goal achievement
Focus Level	Technical (Level A)	Semantic (Level B)	Effectiveness (Level C)
Key Metrics	BER, SER, Capacity, Throughput	Semantic Similarity, Semantic Error	Task Success Rate, Task Error, Efficiency
Info. Value	All bits treated equally	Based on contribution to overall meaning	Based on relevance to specific task/goal
Role of Task	Agnostic	Implicitly related to meaning extraction	Central driver of communication design
Knowledge Base	Generally absent	Often required for semantic interpretation	Essential for task relevance & interpretation

avatar's initial pose and mesh in AR), or task models trained in the cloud and broadcast to all users, as in multi-user VQA transceivers (Xie et al., 2021a), and SC effectiveness depends heavily on this shared knowledge's quality and consistency.

In addition to static KBs, dynamic context acquisition modules can also play an important role (Strinati et al., 2024; Kountouris and Pappas, 2021): by sensing the environment, tracking user intent, or maintaining a history of interactions (potentially via memory modules), they let the system dynamically adapt its semantic processing to the current situation. For example, recent architectures that integrate large models let the foundation model's embedded knowledge stand in for a dedicated shared knowledge base, segmenting content by task prompt and reconstructing it with a conditional generative model, thereby operating without explicit KB synchronization (Chen et al., 2024b).

However, the reliance on KBs is itself a liability: a KB must be built, kept up to date, and kept consistent across parties, and each requirement becomes harder in dynamic multi-agent systems. Semantic mismatch, where the KBs at the transmitter and receiver differ or become outdated, can lead to misinterpretation and ultimately task failure (Bao et al., 2011). Recent approaches address this through curriculum learning frameworks that actively synchronize a common language across communication links (Farshbafan et al., 2023), and through mapping abstract KBs to computable knowledge graphs that provide reproducible mechanisms for resolving semantic ambiguity (Zhou et al., 2022). Addressing these KB management issues, including modeling knowledge evolution, is therefore critical for the widespread deployment of reliable ToSC systems (Guo et al., 2024a).




Fig. 5. Consolidated ToSC system model with IB notation. The semantic encoder E , conditioned on source knowledge base KB_s , compresses source X into representation Z . The decoder D , conditioned on KB_d , produces the task output $\hat{T}$. The IB objective governs the compression–task-relevance trade-off: SC corresponds to $T=X$ (full reconstruction), while ToSC targets a downstream task $T \neq X$, discarding task-irrelevant information. The red link highlights that KB mismatch ($K_s \neq K_d$) introduces semantic noise independent of channel quality.

3.2.3. Task execution and evaluation units

Located on the receiver side, the task execution unit uses the decoded SI to perform the intended goal (Xie et al., 2021a; Lu et al., 2022); its nature is entirely application-dependent: (1) a classifier that outputs a label, (2) a controller that generates actuation commands for a robot, (3) a decision-making algorithm, (4) a rendering engine that reconstructs a scene, or (5) an inference engine that answers a question.

Complementing the execution unit, an evaluation mechanism, either implicit or explicit, assesses the task execution unit’s performance against predefined metrics, important for system optimization during training (e.g., calculating the loss function based on task error) and potentially for adapting system behavior during operation. The performance metrics employed are detailed in Section 4.

To illustrate the ToSC framework, we consider a VQA system where the task is answering a textual question about an image (the source data). Here, the task-tailored semantic encoder typically employs a CNN for visual feature extraction and a Transformer for interpreting the question’s semantics, fusing both into a joint, task-relevant semantic representation. This encoding and decoding process draws on shared knowledge: in VQA, a static KB of common-sense ontologies about objects and relationships, and contextual data from training sets, complemented by the implicit knowledge carried in the pre-trained encoder weights. If the VQA system components are distributed (edge processing versus a central server), the semantic representation is transmitted over a communication channel; the semantic decoder receives it and, often using multimodal reasoning techniques, interprets the fused semantics to extract the information needed to answer the question; the task execution unit renders this information as the final textual answer. An evaluation unit then assesses success using metrics such as VQA accuracy, comparing the generated answer to the ground truth.

3.2.4. Interaction between components

This end-to-end flow, traced concretely in the VQA example above, is common to ToSC systems regardless of application: guided by the task requirements, the KBs, and contextual information, the encoder’s compressed representation crosses the channel to a semantic decoder that either directly produces the task output or hands an intermediate representation to the task execution unit for execution and evaluation.

ToSC frameworks are defined by tight integration and often joint optimization of these functional blocks (Gündüz et al., 2023), which differs from the traditional modular approach, where source coding, channel coding, and modulation are optimized independently; this separation’s optimality guarantees do not extend to the finite-length regimes typical of ToSC (Kostina and Verdú, 2013). Architectural variations exist, such as split learning, where the encoder resides on an edge device while the decoder and task execution unit operate on a more powerful server for distributed resource usage.

3.3. Implementation architectures

The conceptual framework of ToSC has been realized through diverse architectural designs. We structure the architectural taxonomy along four functionally meaningful axes that directly inform engineering decisions: (1) encoding strategy (end-to-end JSCC vs. modular separate source-channel coding), (2) transmission mode (analog vs. digital), (3) task-relevance mechanism (how the system identifies and prioritizes task-relevant information), and (4) adaptivity (static vs. channel-adaptive vs. task-adaptive). These axes capture the semantic communication *design logic* that the neural network backbone alone does not; the AI model family is complementary, determining *how* each axis is realized, and is analyzed separately in Section 3.4. Table 3 maps representative systems along the four axes. These dimensions are not individually new: the joint-vs-separate and analog-vs-digital distinctions are established JSCC design choices (cf. the overview of Gündüz et al. (2023)), and task-relevance and adaptivity recur throughout the ToSC literature; our contribution is consolidating them into a single comparative frame overlaid with the model-family lens of Section 3.4, not introducing each axis individually.

3.3.1. Encoding strategy: End-to-end JSCC vs. Modular

The most fundamental architectural distinction in SC/ToSC is whether source and channel coding are performed jointly or separately.

End-to-end JSCC systems train a single DNN that maps source data directly to channel input symbols, bypassing the traditional separation of source coding, channel coding, and modulation. Bourtsoulatzis et al. (2019) demonstrated that such deep JSCC encoders achieve better rate–distortion performance (measured by PSNR and SSIM under noisy channels) than separated source-channel coding, establishing the foundational architecture for modern systems. Subsequent work expanded this: Kurka and Gündüz (2021) introduced bandwidth-adaptive Deep-JSCC with successive refinement, transmitting images progressively in layers that receivers aggregate to trade bandwidth for reconstruction quality with graceful degradation. Dai et al. (2022) proposed NTSCC, which bridges neural transform coding and JSCC by learning a latent entropy model that guides joint rate–distortion–channel optimization, achieving gains over both purely learned compression and conventional JSCC. Kurka and Gündüz (2020) further improved wireless image transmission by exploiting channel output feedback to enable iterative refinement at the transmitter (DeepJSCC-f); for a detailed treatment of deep JSCC, see the overview by Gündüz et al. (2023). This approach offers graceful degradation under varying channel conditions: performance degrades smoothly rather than exhibiting the cliff effect of digital systems, and naturally supports variable-rate transmission when

combined with bandwidth-adaptive mechanisms. However, end-to-end JSCC requires differentiable channel models for gradient-based training, which limits applicability to channels with tractable mathematical descriptions.

Modular (separate coding) systems retain the classical separation principle, applying a learned semantic encoder for source compression followed by conventional or learned channel coding. This approach offers compatibility with existing communication infrastructure and allows independent optimization of each module, exemplified by systems such as the VQA framework of Xie et al. (2021a), which extracts visual semantic features with a frozen ImageNet-pre-trained ResNet-101 encoder before standard transmission. The drawback is that the separation theorem's optimality guarantees, which hold for infinite block lengths, do not extend to the finite-length, delay-constrained regimes typical of ToSC applications, where JSCC can offer substantial gains (Gündüz et al., 2023).

3.3.2. Transmission mode: Analog vs. digital

Within JSCC, a secondary but consequential distinction is whether the encoder output is transmitted as continuous-valued (analog) channel symbols or quantized to a finite constellation (digital).

Analog JSCC transmits continuous-valued encoder outputs directly over the channel, preserving the smooth mapping between source semantics and channel symbols. This enables graceful degradation: as the signal-to-noise ratio (SNR) varies, reconstruction or task quality degrades proportionally rather than catastrophically. Most deep JSCC systems, including Boursoulatz et al. (2019) and the multi-task framework of Lyu et al. (2024), operate in this mode. Analog operation, however, is incompatible with digital communication infrastructure (e.g., existing modulation and coding schemes, ARQ protocols).

Digital JSCC quantizes encoder outputs to discrete symbols compatible with standard digital modulation, enabling integration with existing protocol stacks and standard error correction. Tung et al. (2022) proposed DeepJSCC-Q, which learns to map continuous encoder outputs to a finite constellation while preserving the graceful degradation property through a soft-to-hard quantization strategy, demonstrating that digital JSCC can approach the performance of its analog counterpart. Quantization noise is the cost, and hybrid approaches using learned quantization to minimize the quantization-task performance gap remain an active research direction.

3.3.3. Task-relevance mechanism

A defining design choice in ToSC is how the system identifies which information is relevant to the downstream task, an axis along which three approaches are prevalent.

Implicit (loss-function-driven): The most common approach defines task relevance implicitly through end-to-end training with a task-specific loss function (e.g., cross-entropy for classification, mean average precision (mAP) for detection). The encoder learns to retain task-relevant features and suppress irrelevant ones as a byproduct of optimizing the loss. Encoder-decoder pairs trained end-to-end for a classification task (Sagduyu et al., 2023) and autoencoder-based systems trained with reconstruction loss exemplify this approach. Its appeal is simplicity, but relevance is fixed at training time and cannot adapt to new tasks without retraining.

Explicit (ranking/filtering): Some systems decouple feature extraction from relevance assessment by applying a separate ranking or filtering stage. The personalized encoder of Kang et al. (2022) scores scene-graph triplets by their saliency to a user's stated interest and transmits only the matching subgraph, and the goal-oriented avatar framework of Wang et al. (2024) transmits skeleton keypoints ranked by their importance to avatar pose recovery. This approach supports dynamic relevance adaptation but requires an explicit relevance model.

IB-principled: A theoretically grounded alternative uses the information bottleneck framework (Eq. (2)) to derive representations that are maximally informative about the task while maximally compressing

the source. The deep variational information bottleneck (Alemi et al., 2017) provides tractable bounds for DNN-based optimization. The IB approach offers a principled compression-relevance trade-off controlled by a single parameter β , but is computationally intensive and sensitive to the accuracy of variational approximations.

3.3.4. Adaptivity

ToSC systems vary in their ability to adapt to changing conditions during deployment.

Static: The encoder and decoder are fixed after training, with performance optimal for the training distribution but degrading under distribution shift (channel variation, task drift, KB mismatch). Most early ToSC systems operate in this mode.

Channel-adaptive: The encoder adjusts its output based on real-time channel state information (CSI). The channel-adaptive gated network of Lyu et al. (2024) prunes output features based on channel SNR, dynamically adjusting the data rate while maintaining task performance. X. Chen et al. (2025) use a masked auto-encoder with a channel-aware extractor that selects patches for transmission based on CSI.

Task-adaptive: The encoder adjusts its behavior based on changing task requirements. Multi-task learning architectures (Eldeeb et al., 2024) learn shared representations that serve multiple tasks via task-specific decoder heads, though keeping the tasks' features separable within one backbone requires dedicated remedies such as the domain-adaptation design of Zhang et al. (2022), which projects each task's features into a task-specific domain so that only those features need to be transmitted. More flexible approaches use foundation models (LLMs/LMMs) to dynamically decompose and adapt to new tasks (Xie et al., 2024), though at substantial computational cost.

3.3.5. Advanced integration approaches

Beyond the core design axes, recent architectures extend ToSC through integration with emerging technologies: O-RAN integration embeds semantic intelligence within the radio access network via a semantic plane and semantic RIC (Strinati et al., 2024), and the broader design space for embedding such semantic processing across the radio-access plane is surveyed by Meng et al. (2026). Hybrid approaches also combine ToSC with complementary paradigms such as Integrated Sensing and Communications (ISAC), where communication signals serve dual sensing purposes (Sagduyu et al., 2024); these are discussed in detail as research trends in Section 6.

3.4. The role of AI model families in ToSC

The functional axes above describe *what* architectural choices a ToSC system makes; the AI model family determines *how* those choices are realized, and is where the engineering of semantic extraction and task inference actually happens. Because these methods are central to the artificial-intelligence perspective on ToSC, we analyze the model families that recur across the surveyed systems along four roles (semantic extraction, semantic compression, task reasoning, and knowledge representation) and summarize their suitability and limitations relative to the architectures of Section 3.3 in Table 4.

Autoencoders and variational autoencoders form the backbone of deep JSCC: a learned encoder-bottleneck-decoder maps the source directly to channel symbols, performing extraction and compression jointly (Boursoulatz et al., 2019). Disentangled variants such as the β -VAE expose interpretable latent factors that can be selectively transmitted (Ma et al., 2023). They excel at analog end-to-end JSCC but offer little explicit task reasoning, and their latent space is largely opaque, which complicates structured knowledge representation and interpretability.

Convolutional neural networks are the dominant choice for spatial and visual semantic extraction, encoding images into compact, task-discriminative features for classification- or detection-oriented ToSC (Lyu et al., 2024). Their inductive bias and hardware efficiency suit channel-adaptive, loss-driven (implicit-relevance) designs,

Table 3
Taxonomy of representative ToSC architectures along functional design Axes.

Representative system	Encoding strategy	Tx mode	Task-relevance	Adaptivity	Backbone	Domain/Task
Bourtsoulatze et al. (2019)	JSCC (end-to-end)	Analog	Implicit (recon. loss)	Static	CAE	Image reconstruction
Lyu et al. (2024)	JSCC (end-to-end)	Analog	Implicit (unified loss)	Channel-adaptive	CAE + gated net	Multi-task (recon. + classif.)
Eldeeb et al. (2024)	JSCC (end-to-end)	Analog	Implicit (MTL loss)	Task-adaptive	CAE/CNN + MTL heads	Autonomous vehicles
Xie et al. (2021a) (MU-DeepSC)	Modular	Digital	Implicit (task loss)	Static	ResNet-101 + MAC	VQA (multi-user)
Wang et al. (2023) (TSAR)	Modular	Digital	Explicit (base-knowledge selection)	Static	SANet	AR avatar rendering
X. Chen et al. (2025)	JSCC (end-to-end)	Analog	Explicit (patch scoring)	Channel-adaptive	Masked AE	Flexible multi-task
Wang et al. (2024b) (I-JSCC)	JSCC (end-to-end)	Analog	Explicit (feature importance)	Channel-adaptive	CNN	Multi-user image
Ma et al. (2023)	Modular	Digital	IB-principled (β -VAE)	Static	β -VAE	Explainable SC
Du et al. (2025)	Modular (split)	Digital	Implicit (LMM reasoning)	Task-adaptive	LMM + lightweight enc.	Vehicular VQA
Strinati et al. (2024)	Modular (O-RAN)	Digital	Explicit (semantic plane)	Task-adaptive	Protocol-level	Network-wide mgmt.

but they capture local structure better than long-range or relational dependencies and do not natively represent structured knowledge.

Transformers and Vision Transformers model long-range dependencies and support cross-modal fusion through attention, underpinning text SC (e.g., DeepSC Xie et al., 2021b) and multimodal, multi-user systems (Xie et al., 2022). Attention also provides a natural mechanism for explicit task-relevance ranking. The cost is computational: quadratic attention and large parameter counts raise latency and energy use, which constrains deployment on resource-limited edge devices.

Graph neural networks (GNNs) operate directly on graph-structured semantics such as scene graphs and knowledge graphs, making them the natural substrate for knowledge representation and relational task reasoning; they underpin explainable scene-graph SC (Sun et al., 2023) and cross-modal graph-based SC (Chen et al., 2024a). The dependence on an explicit graph-construction step is the main weakness: when the extracted relations are noisy, incomplete, or absent, the reasoning built on them is degraded accordingly.

Large language and multimodal models act simultaneously as cross-modal semantic codecs, task decomposers, and carriers of the implicit (parametric) knowledge identified in our three-layer knowledge definition; the KG-LLM scheme of Salehi et al. (2025) illustrates the coupling of an explicit graph with parametric knowledge. They offer the strongest task reasoning and zero-shot adaptivity of the families surveyed, but heavy inference cost, hallucination, and inconsistency across model versions and fine-tunes restrict real-time and edge use, as discussed further in Section 6. The moderate semantic-compression rating in Table 4 reflects the model family's intrinsic role as a within-modality codec, no denser than a tuned autoencoder bottleneck; the order-of-magnitude reductions associated with large-model ToSC (e.g., the extreme image-to-text compression of Ribouh et al. in Section 6.1.1) arise instead from the cross-modal transform the model family enables (representing an image by a short text description), a distinct mechanism from the bottleneck compression scored in that column.

No single family dominates all four roles, and practical ToSC systems compose them, for example a CNN or ViT extractor feeding a GNN- or LLM-based reasoner, followed by an autoencoder bottleneck for transmission. The functional-axis taxonomy and this model-family

lens are therefore complementary: the axes fix the communication design logic, while the model family determines the achievable semantic granularity, reasoning depth, and deployment cost.

3.5. Semantic processing mechanisms

ToSC systems process information through three stages that collectively determine what is transmitted: representation, extraction, and filtering. Task-relevant semantics are represented as latent vectors (learned by AE/DNN encoders) (Choi et al., 2024), knowledge graphs (subject-predicate-object triples for structured data) (Sun et al., 2023; Zhou et al., 2022), or explicit feature vectors (bounding boxes, skeleton joints, classifier inputs) (S. Liu et al., 2025; Wang et al., 2024). Relevance is defined either implicitly (Sagduyu et al., 2023) or explicitly (S. Liu et al., 2025; Wang et al., 2024), per the task-relevance axis of Section 3.3.

Extraction techniques include (1) DNN-based feature learning (CNNs, Transformers, AEs), (2) classical CV algorithms (object detectors, key-point estimators) (Yang et al., 2022), (3) NLP methods (end-to-end text semantic encoders and knowledge-graph construction) (Zhou et al., 2021, 2022), and (4) the IB framework (Tishby et al., 2000), which derives compressed representations maximizing task-relevant mutual information via variational approximations (Alemi et al., 2017; Barbarossa et al., 2023). After extraction, semantic ranking prioritizes elements by task contribution, adaptive transmission adjusts fidelity to channel conditions and target reconstruction quality (W. Zhang et al., 2023), and semantic filtering removes task-irrelevant content (Park and Yoon, 2025).

This pipeline entails a fundamental trade-off. Blau and Michaeli (2019) proved that rate-distortion bounds conflict with perceptual quality, motivating perceptual/generative over pixel-fidelity compression, a distinction within semantic fidelity (Level B). The case for task-oriented (Level C) compression rests instead on the information-bottleneck argument: Shao et al. (2021) used a variational information-bottleneck formulation to characterize an achievable rate-accuracy trade-off between communication rate and task performance. Recent work refines this IB-based rate control in several directions: rate-optimized IB with SNR-adaptive networks for resilient image SC (Wang

Table 4
AI model families in ToSC: Suitability across semantic roles.

Model family	Semantic extraction	Semantic compression	Task reasoning	Knowledge Repr.	Representative use/Key limitation in ToSC
Autoencoder/VAE	•	•	×	×	Backbone of deep JSCC (Boursoulatz et al., 2019); β -VAE for interpretable selective transmission (Ma et al., 2023). Opaque latent space; little task reasoning.
CNN	•	◦	×	×	Spatial/visual feature extraction for classification-oriented ToSC (Lyu et al., 2024; Eldeeb et al., 2024). Weak at long-range and relational reasoning.
Transformer/ViT	•	◦	•	◦	Long-range dependencies and cross-modal fusion; text SC (Xie et al., 2021b) and multimodal multi-user SC (Xie et al., 2022). High compute and latency.
GNN	◦	◦	•	•	Graph-structured semantics (scene/knowledge graphs); explainable scene-graph SC (Sun et al., 2023) and cross-modal graph SC (Chen et al., 2024a). Needs explicit graph construction; brittle to noisy relations.
LLM/LMM	•	◦	•	•	Cross-modal codec, task decomposition, and carrier of implicit knowledge (Salehi et al., 2025). Inference cost, hallucination, version/fine-tune drift.

Suitability per role (operational scoring): • = a primary, purpose-built mechanism for the role in the surveyed ToSC literature; ◦ = capable of the role but as a secondary function or with known limitations; × = not an inherent strength, typically delegated to another family. Generative and diffusion models are not listed as a separate family because they act as a *decoding* mode rather than an extraction/compression backbone; they are treated in Section 6.1.4. *Mapping to the functional axes of Table 3*: autoencoders/VAEs realize analog end-to-end JSCC and, in disentangled β -VAE form, the information-bottleneck-principled modular design of Ma et al. (2023); CNNs suit channel-adaptive, implicit-relevance designs; Transformers/ViTs enable explicit task-relevance ranking; GNNs underpin knowledge-structured, relational modular designs such as the structured scene-graph SC of Sun et al. (2023); and LLMs/LMMs enable task-adaptive modular reasoning. The Backbone column of Table 3 gives the realized per-system pairing. These per-role suitability ratings reflect the authors' assessment of the surveyed literature rather than a numerical benchmark, and a single rating necessarily abstracts over variation within each broad family.

et al., 2026), feature disentanglement into common and private components to prune redundancy in multi-device cooperative ToSC (M. Sun et al., 2026), and a conditional rate–distortion formulation with task prompts and feedback for task-adaptive coding (He et al., 2026a). Defining relevance too narrowly risks brittleness; finding the optimal balance between compression and task effectiveness remains a central challenge.

3.5.1. Cross-mechanism synthesis

Some of the surveyed mechanisms recur far beyond the domains that introduced them. *Information-theoretic extraction* (IB-based or variational) and *saliency-based prioritization* appear in nearly every application domain, from text and image transmission to autonomous driving and extended reality, and form the foundational layer of any ToSC pipeline. *Adaptive transmission* and *feedback coupling* recur for a different reason: real-world channels and tasks are non-stationary, and static semantic compression breaks down once the environment or the task objective evolves. *Robustness control* mechanisms (adversarial training, channel-adaptive encoding, redundancy injection) appear wherever deployment conditions are uncertain, reflecting the tension between aggressive semantic compression and system reliability. Supplementary Table S2 maps each recurring mechanism to representative systems, application domains, and the underlying reason for its recurrence.

3.6. When does the SC/ToSC distinction matter? A critical analysis

The preceding sections established that SC optimizes for semantic fidelity (Level B) while ToSC optimizes for task effectiveness (Level C). However, this distinction risks being merely definitional unless we can identify *when it leads to materially different system behavior*. This subsection synthesizes the IB framework, originating with Tishby et al. and mapped to SC by Beck et al., with empirical evidence from the surveyed literature to characterize the qualitative regimes under which SC and ToSC diverge or converge, together with the design principles that follow.

3.6.1. Formal grounding via the information bottleneck

The SC/ToSC distinction can be formalized through the IB framework (Tishby et al., 2000; Goldfeld and Polyanskiy, 2020). Beck et al. (2023) cast this IB formulation in a wireless physical-layer semantic-communication setting, using a mutual-information objective to trade off source-information preservation against task-relevance. Given source data X , a compressed representation Z , and a task variable T , the IB objective minimizes:

$$\mathcal{L}_{\text{IB}} = I(X; Z) - \beta \cdot I(Z; T) \quad (2)$$

where $I(\cdot; \cdot)$ denotes mutual information and β controls the trade-off between compression and task-relevance. General SC corresponds to the regime where $T = X$ (the task is full reconstruction), so $I(Z; T) = I(Z; X)$, and the system preserves as much source information as possible at a given rate. ToSC corresponds to $T \neq X$ (the task is a downstream function of the source, e.g., a class label), and the system discards source information that is irrelevant to T .

This formalization reveals that **SC and ToSC diverge whenever $I(X; T) < H(X)$** , that is, whenever the task requires strictly less information than the source contains. The gap $H(X) - I(X; T)$ quantifies the “task-irrelevant” information that ToSC can discard but SC must preserve. For high-dimensional sources (images, video, point clouds) with low-dimensional tasks (classification, detection), this gap is typically large, making the distinction practically significant. The strict inequality itself holds for essentially any lossy task (since $H(X) - I(X; T) = H(X|T) \geq 0$, with equality only when X is a deterministic function of T); it is therefore the *magnitude* of this conditional entropy, not the inequality alone, that decides whether the SC/ToSC distinction is of engineering significance.

3.6.2. Concrete evidence of divergence

The surveyed literature, complemented by a representative vehicular design pattern, shows that SC-optimized and ToSC-optimized systems produce materially different outcomes:

Case 1: Personalized task-oriented selection vs. full reconstruction. Kang et al. (2022) design a personalized, saliency-based task-oriented encoder for UAV image sensing that represents each image as a scene graph of subject-relation-object triplets and transmits only the triplets matching a user's stated interest (e.g., "man wearing a jacket"), weighted by a personalized attention map, reporting on the order of a 64% reduction in communication cost. A reconstruction-oriented SC system would instead preserve the entire image; the two systems thus retain materially different content from the same source, illustrating how a task objective and a reconstruction objective *diverge in what they preserve*. Xiao et al. (2022) formalize a related divergence through rate-distortion theory for strategic semantic communication, in which transmitter and receiver hold different distortion measures and choose their encoding and decoding strategies rationally, showing that the achievable rate-distortion performance depends on the equilibrium the two sides reach rather than on signal fidelity alone.

Case 2: Autonomous driving: detection vs. scene reconstruction. As a representative vehicular-network design pattern, a ToSC system transmitting only detected object features (bounding boxes, class labels, trajectories) would achieve high mAP for downstream planning using a fraction of the bandwidth required for full scene reconstruction, but such systems fail when the downstream task changes unexpectedly, for example when a previously unmodeled obstacle type appears, or when the planning module requires contextual information (road surface conditions, weather cues) discarded by the task-specific encoder. This reveals a fundamental **brittleness trade-off**: ToSC's efficiency gains come at the cost of reduced adaptability to task drift.

Case 3: Multi-task systems as the boundary case. Lyu et al. (2024) design a multi-task SC system that jointly optimizes for image reconstruction (Level B) and classification (Level C). By sweeping the loss-balance parameter between the two objectives, they show that the two cannot be maximized simultaneously: as the weighting shifts toward classification, the classification accuracy rises while the reconstruction PSNR falls (their Fig. 10), tracing a non-trivial, qualitatively reported accuracy-PSNR trade-off between Level B and Level C performance. This trade-off is consistent with the information-bottleneck picture above. Because the classification target is a coarsening of the source (T a deterministic function of X), the data-processing inequality bounds task-relevant by reconstruction-relevant information, $I(Z;T) \leq I(Z;X)$. We caution that this is an information-content bound, not a guarantee that task-relevant *features* form a literal subset of reconstruction-prioritized ones. As Case 4 below notes, low-variance yet high-value features can be task-critical while a reconstruction objective deprioritizes them. The observed trade-off is therefore suggestive evidence for this bounded-information relationship in their multi-task setting rather than a proof of feature-set containment.

Case 4: Low-variance, high-value features. Divergence can also arise at the opposite extreme of compression, when task-critical information is low-variance yet high-value: a small lesion occupying a few pixels of a medical image, or a distant pedestrian in a driving scene. Variance-driven compression preserves the dominant background while discarding precisely the subtle, low-energy features on which the task depends; the target then occupies a small region within a large task-irrelevant background, so the effective task-information gap $H(X) - I(X;T)$ is *large* despite the target's low variance. Given a relevance signal, a task-aware encoder can protect these high-value features that a reconstruction-oriented objective would deprioritize, so SC and ToSC *diverge*. Without such a signal (as in rare-pathology cases where the salient feature is not flagged), both objectives discard the feature and task outcomes degrade in tandem, a *co-failure* rather than a benign convergence.

3.6.3. When the distinction collapses

The SC/ToSC distinction becomes negligible in two important scenarios:

(1) **Low-rate regimes with energy-dominant tasks.** When the channel bandwidth is severely limited and the task's discriminative information is carried by the highest-variance (energy-dominant) components of the source, both SC and ToSC systems are forced onto the same small set of most-salient features. Under this condition, and only under it, the features necessary for reconstruction and those necessary for the task converge, because both must prioritize the highest-variance components of the data. This is an expected consequence of the rate-constrained regime: at very low SNR or extreme compression ratios, both objectives are forced onto the same high-variance components, so SC- and ToSC-optimized systems should perform similarly. This *success-convergence* (both objectives retain the same features and the task still succeeds) must be distinguished from the *co-failure* of Case 4 above, in which both objectives discard a low-variance but task-critical feature and degrade in tandem: the former makes ToSC unnecessary, the latter makes it ineffective without a relevance signal. The convergence argument here therefore applies only when the task's discriminative information is energy-dominant, and should not be assumed for fine-grained or safety-critical perception.

(2) **Tasks that require full source fidelity.** Some tasks (open-ended diagnostic reading, where a clinician may inspect the entire image for any abnormality so that almost all of the source is potentially task-relevant, or speech communication where prosodic nuances carry meaning) inherently require high source fidelity. For these, the "task-irrelevant" information gap $H(X) - I(X;T)$ is small, and ToSC offers minimal advantage over SC. This observation has practical implications: *not all applications benefit from the ToSC approach*, and engineers should evaluate the information gap before investing in task-specific encoder design.

3.6.4. Implications for system design

This analysis yields three practical design principles for practitioners:

1. **Estimate the task-information gap before committing to ToSC.** If the downstream task is complex (e.g., multi-step reasoning, open-ended generation) or if task requirements may change, the efficiency gains of ToSC may not justify the loss of generality. A multi-task or progressive encoding approach may be more appropriate.
2. **KB alignment is plausibly more critical for ToSC than for SC.** We argue that, because ToSC systems discard "redundant" source information a receiver could otherwise exploit to recover from KB mismatch, they should be *more sensitive* to KB degradation than general SC systems, which retain that information (cf. Section 3). We present this as a reasoned design conjecture rather than an empirically settled result: the surveyed literature establishes that KB mismatch degrades semantic capacity for SC broadly (Bao et al., 2011) but does not isolate a strictly greater task-performance degradation for ToSC than SC under matched KB drift. Quantifying it rigorously requires the KB-dependent rate-distortion bounds we flag as Priority 1 (Section 7). This conjecture nonetheless has direct implications for deployment in dynamic environments where KBs evolve independently.
3. **Design for graceful degradation.** ToSC encoders should include mechanisms for fallback to higher-fidelity transmission when task confidence is low or when the system detects distribution shift, rather than operating at a fixed compression point optimized for a single task scenario.

Each of these principles presupposes that semantic fidelity and task effectiveness can actually be measured; Section 4 examines the metrics available for doing so.

4. Performance evaluation metrics

Evaluating SC and ToSC systems requires metrics that go beyond classical bit error rate to capture meaning preservation and task success. Following the three-level framework established in Section 2 (cf. Fig. 2), this section organizes metrics by their target level: Level B (semantic fidelity) for general SC, Level C (task effectiveness) for ToSC, and cross-level efficiency metrics. We conclude with an analysis of metric conflicts and practitioner guidance for metric selection. Supplementary Table S3 provides a consolidated overview of metrics across all levels.

4.1. Metrics for semantic fidelity and quality (Level B)

Level B metrics assess how faithfully meaning and content survive transmission and reconstruction; they are highly modality-dependent.

- **Image/Video Quality:** Common metrics evaluate image and video quality. Peak signal-to-noise ratio (PSNR) quantifies reconstruction quality by measuring the ratio of maximum signal power to corrupting noise power, assessing pixel-level similarity. Some JSCC frameworks use it as a primary metric for image recovery. In contrast, the structural similarity index measure (SSIM) is a perceptual metric quantifying quality degradation by comparing luminance, contrast, and structure. SSIM can be incorporated into the loss function for training deep JSCC frameworks to directly optimize perceptual quality (Lyu et al., 2024). Beyond pixel-level and structural metrics, *learned perceptual metrics* are increasingly adopted in SC evaluation. The Learned Perceptual Image Patch Similarity (LPIPS) computes distance in the feature space of a pre-trained deep network, capturing perceptual differences that PSNR and SSIM miss; it is particularly relevant for generative SC systems (Section 6.1) where reconstructions may be perceptually plausible but pixel-misaligned. The Fréchet Inception Distance (FID) measures the distributional similarity between generated and real images, providing a population-level quality assessment. Both metrics connect directly to the rate-distortion-perception trade-off formalized by Blau and Michaeli (2019): optimizing for PSNR forces divergence from perceptual quality (as measured by LPIPS/FID), making the choice of metric a design decision with architectural consequences.
- **Text Semantic Similarity:** Text semantic similarity metrics evaluate meaning preservation after transmission. Traditional metrics like Bilingual Evaluation Understudy (BLEU) assess quality by comparing n-gram overlaps between original and received sentences, but may not fully capture meaning if synonyms are used. A more advanced approach uses deep learning models, like sentence-BERT, to map sentences into high-dimensional vector embeddings (Yan et al., 2022b). Semantic similarity is then the cosine similarity between the original and received sentence embeddings. The resulting score (0 to 1) offers a precise measure of meaning preservation and can serve as a key performance indicator or quality constraint.
- **Speech Quality:** Evaluation of speech SC systems depends on the goal: accurate transcription (speech-to-text) or high-fidelity reconstruction (speech-to-speech). For speech-to-text, word error rate (WER) is a standard metric measuring accuracy by counting word substitutions, deletions, and insertions against a ground-truth text (Han et al., 2022). To better capture meaning, semantic similarity scores from sentence embeddings (e.g., BERT) can be used. For speech-to-speech, both objective and subjective metrics are used. Objective metrics include the signal-to-distortion ratio (SDR), the L2 error between the original and recovered signals (Weng and Qin, 2021), and Mel cepstral distortion (MCD), a distance between speech spectra (Han et al., 2022). Perceptual metrics approximate human judgment, like the

ITU-T standard perceptual evaluation of speech quality (PESQ) (Rix et al., 2001). Subjective listening tests yield a mean opinion score (MOS), where human raters score the recovered audio's naturalness (Ribeiro et al., 2011).

4.2. Metrics for task effectiveness and goal achievement (Level C)

For task-oriented specializations of SC, metrics at the effectiveness level are central, as they directly quantify the successful completion of a given task. These metrics are inherently goal-oriented and task-specific, with success being defined by the application's ultimate objective. This subsection details the primary metrics that operate at this level, starting with the Value of Information (VoI).

4.2.1. Value of Information (VoI)

VoI is a goal-oriented semantic metric: it quantifies what a piece of information is worth for achieving a specific task. Unlike content-agnostic metrics, VoI is inherently content-aware and task-dependent. It is formally defined as the expected improvement in task performance gained from receiving new information, often expressed as the difference in the expected value of a utility function before and after the information is acquired. Because it directly measures task-relevance, VoI is well suited to optimizing resource allocation and decision-making. Examples of its formulation include:

- **Control Systems:** In a robotic control task, such as a UAV navigating to a waypoint, the VoI of a command can be directly tied to its impact on goal completion. For instance, it could be formulated as the reduction in the predicted distance to the target (W. Wu et al., 2024). A command that corrects a large deviation has a higher VoI than a minor adjustment made while already on course.
- **Classification Tasks:** In a remote visual classification task (e.g., medical diagnosis), the VoI can be linked to model uncertainty. If a low-resolution image results in a low-confidence prediction (high entropy), the VoI of receiving a higher-resolution version is high, as it is expected to significantly reduce classification error. If the model is already highly confident, the VoI of additional data is low.
- **Object Detection:** For an autonomous vehicle, the VoI of detecting an object can be weighted by its class and kinematics. The VoI of correctly identifying a pedestrian crossing the road is far greater than that of a distant, stationary building. This allows the system to prioritize processing and bandwidth for safety-critical information.

As it captures task-relevance, VoI can be combined with other metrics like the Age of Information for intelligent decision-making, such as re-ordering a command queue at a robot receiver to prioritize high-value, fresh information (W. Wu et al., 2024).

4.2.2. Metrics for task-specific success

For many applications, task success is quantified by direct, objective performance indicators. These metrics are highly specific to the task at hand.

- **Classification/Recognition Tasks:** A primary metric for classification and recognition tasks is classification accuracy, the proportion of correctly identified samples. Performance is often optimized by minimizing a loss function, such as cross-entropy. For instance, recent work on multi-task SC uses classification accuracy to evaluate image classification, comparing methods based on coding rate reduction against cross-entropy minimization, especially under label corruption (Lyu et al., 2024).

- **Object Detection Tasks:** Object detection evaluation assesses both classification and localization accuracy. Localization is measured by Intersection over Union (IoU), the overlap between predicted and ground-truth bounding boxes. The primary performance metric is mAP, the average of Average Precision (AP) scores across all object classes; AP summarizes a class's precision-recall curve. mAP is typically calculated at a specific IoU threshold (e.g., mAP@0.5) or averaged over a range of thresholds, and often reported separately by object scale. Optimization involves a multi-task loss combining classification and bounding box regression losses. For example, a UAV object detection study might use MS COCO standards, with mAP at a 0.5 IoU threshold (AP_{50}) as the primary metric, while also considering performance on small (AP_S), medium (AP_M), and large (AP_L) objects to show robustness (Song et al., 2025).
- **Reconstruction/Estimation Tasks:** For 3D reconstruction tasks, evaluation focuses on geometric accuracy by measuring positional deviation between original and reconstructed models. For articulated objects like human avatars, assessments calculate the error of individual joints or keypoints, while overall scene distortion is measured using 3D comparison metrics. For instance, a study on AR avatar pose recovery evaluates accuracy at two levels: using mean per joint position error (MPJPE) for skeleton joint errors, and a point-to-point (P2Point) error metric for the overall geometric difference in the final rendered scene (Wang et al., 2024).

4.2.3. Metrics for timeliness and information freshness

For dynamic, real-time scenarios, the timeliness of information is critical to task effectiveness.

- **Age of Information (AoI):** AoI quantifies information freshness at the destination, defined as the time elapsed since the generation of the most recently received status update. Its main shortcoming is that it is content-agnostic; it penalizes outdated information even if it remains correct.
- **Age of Incorrect Information (AoII):** To address the limitation of AoI, AoII was introduced as a semantics-empowered metric combining age and correctness. AoII is a penalty computed as the product of an information-penalty function (capturing state error) and a non-decreasing time function (increasing the penalty as the mismatch persists) (Maatouk et al., 2022). It effectively measures the time for which the destination has an incorrect understanding of the source's state.

4.2.4. Metrics for user-centric and system-wide performance

Effectiveness can also be measured from a user-centric perspective or at the scale of an entire network.

- **Quality-of-Experience (QoE):** QoE is a user-centric metric quantifying subjective satisfaction. In SC, it is often a composite metric that balances trade-offs between task effectiveness and communication efficiency, reflecting diverse user priorities. The final QoE score is often calculated using functions (e.g., logistic) that map objective metrics to a satisfaction score (0 to 1). For instance, one model for SC networks bases QoE on two components: semantic accuracy (how well the task was performed) and semantic rate (SI transmitted per second, in suts/s) (Yan et al., 2022a). A weighted sum of satisfaction scores for these components calculates the final QoE, capturing user preferences for either high accuracy or a higher semantic rate.
- **System Throughput in Message (STM):** Traditional throughput in bits is unsuitable for semantic networks as it does not capture the successful delivery of meaning. Instead, network-level metrics are needed to measure total semantic throughput. STM has been proposed to measure the aggregate message rate successfully received by all users in a network (Xia et al., 2023). It uses a

bit-rate-to-message-rate (B2M) transformation function to convert a link's achievable bit rate into an achievable message rate. This B2M transformation incorporates a semantics-relevant term that accounts for the background knowledge and inference capability of the semantic coding models. By maximizing STM, network resources like user association and bandwidth can be allocated to optimize overall network efficiency from a semantic perspective.

4.3. Metrics for communication efficiency and resource utilization

A key motivation for SC is improving communication efficiency. Quantifying these resource savings is essential for demonstrating the superiority of semantic approaches over traditional systems.

- **Bandwidth/Data Rate:** Bandwidth and data rate are fundamental efficiency measures, quantified as bits per pixel (bpp) for images, bandwidth reduction percentage, or compression ratio (Huang et al., 2022). To measure the transmission of meaning, specialized metrics based on the semantic unit (sut) have been proposed. These include the semantic transmission rate (S-R), which measures effective semantic information transmitted per second (suts/s), and semantic spectral efficiency (S-SE), the rate of successful semantic transmission per unit bandwidth (suts/s/Hz), derived by normalizing S-R by channel bandwidth (Yan et al., 2022b).
- **Latency:** For time-sensitive systems like real-time tasks, end-to-end latency, the total time from source data generation to final task execution, is a critical metric. It comprises processing latency (e.g., semantic encoding/decoding) and transmission latency (influenced by data volume, coding rates, retransmissions). The objective is to minimize this latency while meeting reliability and accuracy requirements. For systems with parallel streams, minimizing the maximum latency among all streams is key to avoiding straggler effects. For example, a generative SC framework might transmit semantic modalities, like a textual prompt and other signals, in parallel (Qiao et al., 2024). Optimization could involve a reliable retransmission scheme for the critical prompt and adaptive modulation for other signals, with a latency-aware resource allocation scheme assigning power to minimize the maximum delay under semantic quality constraints.

4.4. Metric conflicts, trade-offs, and practitioner guidance

The metrics defined above are not independent; in practice, optimizing one metric frequently degrades another. Understanding these conflicts is essential for practitioners designing SC/ToSC systems. We identify three categories of metric tension, visualized in Fig. 6, and provide guidance for resolving them.

Level B vs. Level C conflict. Semantic fidelity and task effectiveness can directly conflict: PSNR/SSIM-optimized systems retain details irrelevant to classifiers, while accuracy-optimized systems may produce reconstructions unintelligible to human operators. *Practitioner guidance:* When the system must serve both human inspection and automated tasks, adopt a multi-objective loss with an explicit Pareto trade-off parameter, and report both Level B and Level C metrics. Single-metric optimization is appropriate only when the use case is purely automated. As a concrete illustration, Lyu et al. (2024) demonstrate this trade-off in a multi-task image transmission system: sweeping the loss-balance parameter traces a qualitative accuracy-PSNR trade-off curve (their Fig. 10) on which gains in classification accuracy come at the cost of reconstruction PSNR and vice versa, so the two objectives cannot be maximized at a single operating point and are not freely substitutable. Choosing the operating point on this frontier is a design decision rather than a property the system can optimize away.

Task effectiveness vs. efficiency (timeliness) conflict. Maximizing VoI or accuracy may require additional data or complex inference,




Fig. 6. Metric conflict and trade-off map for SC/ToSC evaluation. The three vertices are the primary metric families: semantic fidelity (Level B), task effectiveness (Level C), and communication efficiency (cross-level). Each edge denotes a pairwise conflict with its governing mechanism: fidelity vs. effectiveness (the Level B and Level C objectives trace a Pareto frontier and cannot be jointly maximized), effectiveness vs. efficiency (richer inference raises value-of-information but costs bandwidth and increases age-of-information), and efficiency vs. fidelity (aggressive compression removes the redundancy, i.e. the robustness margin, that buffers channel errors and knowledge-base mismatch). As the central node states, no single operating point optimizes all three simultaneously, so system design reduces to selecting a justified point on this surface.

increasing latency and degrading AoI. For example, in UAV command-and-control, a stale but high-value command is counterproductive: value-of-information and age-of-information objectives must be traded off so that timely, high-value commands take precedence (W. Wu et al., 2024). *Practitioner guidance:* Define a composite utility function that explicitly penalizes both task error and staleness, such as AoI-weighted accuracy. Prioritize AoI minimization for safety-critical decisions and accuracy maximization for non-time-critical analytics within the same system using a priority queue based on VoI.

Efficiency vs. fidelity (robustness) conflict. Aggressive semantic compression reduces bandwidth but also reduces the redundancy that protects against channel errors and KB mismatch. A system operating at maximum semantic spectral efficiency (S-SE) has no margin for error recovery if the receiver's KB diverges from the sender's. *Practitioner guidance:* Design compression rates with explicit margin for the expected KB mismatch level and channel variability. In dynamic environments, adaptive compression that increases redundancy when confidence is low (measured by receiver feedback) is preferable to fixed-rate schemes.

Limitations of current metrics. Two gaps remain in current evaluation practice. First, VoI is theoretically compelling but rarely computable in closed form for realistic scenarios; most implementations use proxy metrics (e.g., prediction confidence) that may not faithfully represent true task value. Second, no existing metric captures *cross-task generalizability*, the degree to which a semantic representation supports not just the intended task but also related tasks that may arise during deployment. Developing such a metric is an open research challenge that directly affects the practical viability of ToSC systems in multi-task environments.

The need for standardized evaluation. Underlying all of these tensions is a methodological obstacle: there is no standardized way

to evaluate SC/ToSC systems, so reported gains are frequently incomparable across papers. As our cross-domain analysis shows (Pattern 4, Section 5), works differ in datasets, channel models, compression baselines, and even the definition of “task success”, which makes it impossible to determine whether a new system genuinely advances the state of the art or merely benefits from a weaker baseline or an easier channel. Addressing this requires three coordinated efforts. First, a *unified benchmark* (public, multi-modal datasets carrying both reconstruction ground truth and task annotations) so that Level B fidelity and Level C effectiveness can be measured on the same footing. Second, a *standardized test protocol* that fixes the channel models, the compression baselines (including strong neural codecs, not only JPEG/BPG), the SNR operating points, and the reporting requirements, so that results are reproducible and directly comparable. Third, an *open evaluation platform*, namely a public leaderboard with reference implementations and held-out test channels, that additionally reports *deployment-oriented* measures (inference latency, model size, energy per inference, and robustness under KB mismatch) rather than task accuracy alone, so that practical deployability, not benchmark accuracy in isolation, becomes a first-class evaluation criterion. We develop this as a prioritized engineering direction in Section 7 (Priority 2).

5. Use cases in ToSC

The principles of ToSC address key challenges in high-impact technological domains. In areas from mobile robotics to industrial automation, the core challenges are similar: maintaining reliable, low-latency performance in complex, resource-constrained environments. ToSC reduces communication overhead and energy consumption not through

Table 5

Comparison of representative ToSC works across ten application domains. Reported gains are each work's own comparison, predominantly measured against traditional codecs (e.g., JPEG/BPG) or raw/full-state transmission rather than a state-of-the-art learned codec; see the weak-baseline caveat (Pattern 2, Section 5).

Domain	Reference	Core technique	Key result
Vehicular	Eldeeb et al. (2024)	Multi-task SC for traffic sign recognition/classification	89% transmitted bit reduction
	Zhang et al. (2022)	Unified multi-task encoder backbone	Comparable accuracy to single-task models with much less transmission overhead
	Feng et al. (2025)	Edge-enabled multimodal semantic fusion	95%+ image processing runtime reduction
	Lv et al. (2024)	Importance-aware SC (VIS-SemCom)	~70% data reduction, high segmentation quality
	C. Liu et al. (2023)	Adaptable-ratio SC with resource allocation	Jointly adapts compression ratio and wireless resources to multi-user delay constraints
Sat/UAV	Hassan et al. (2024)	Joint semantic-channel encoding for LEO satellites	Reduced latency while preserving data quality
	W. Chen et al. (2025)	FSO-SC with VQ-VAE for remote sensing images	3 dB power gain; 60% overhead reduction
	Z. Xu et al. (2025)	Task-oriented transmission with UoI metric	Minimized UoI and power consumption
	Song et al. (2024)	Cognitive SC with knowledge graph for UAV detection	Only 0.6% accuracy drop under noisy channels
	Xu et al. (2024)	Hierarchical federated learning for UAV swarm SC	Higher PSNR and lower MSE than baseline
	Y. Xu et al. (2023)	TOSA framework with DRL for C&C packets	Near-constant transmit count vs. linear growth; matched effectiveness
Surveillance	Shah et al. (2021)	PeR-VIS: semantic description-based person retrieval	High IoU and detection rate, surpassed prior best
	Shao et al. (2023)	TOCOM-TEM: task-oriented edge video analytics	62%–85% comm. cost reduction at matched accuracy
	Nandakishore et al. (2024)	Background-as-shared-knowledge, dynamic ROI only	86%+ bit savings with improved PSNR
	A.Y. Huang et al. (2023)	SPE: privacy-preserving SC via semantic scene graphs	Parameter-only exchange (~50 MB vs. up to 50 GB raw)
IoT/IIoT	Yang et al. (2021)	Task-relevant semantic extraction for AI-task classification	>40% accuracy gain; much lower bandwidth
	Xie and Qin (2020)	L-DeepSC: pruned/quantized SC for IoT	Model 90% smaller with minimal perf. loss
	Hu et al. (2023)	Robust SC under semantic noise	Masked VQ-VAE codebook robust to adversarial semantic noise
	H. Chen et al. (2024)	Edge-assisted optimal model selection per device	Near-optimal solution in simulations
	Moqurrah et al. (2022)	Instant Anonymity: perturbation-based privacy	Lower query error vs. k-anonymity
	Willner et al. (2017)	oneM2M + Semantic Web for Industry 4.0	Integrates existing formal information models
	Zhao et al. (2024)	Distributed RIS with probabilistic SC	Up to ~2× semantic-aware rate vs. centralized RIS
	Zhang et al. (2024)	ISC with semantic flags for tool wear monitoring	12.9% delay reduction
Healthcare	Y. Yang et al. (2024)	ISCSC: joint beamforming + semantic extraction	70% transmission rate improvement
	Bura and Jamalpur (2025)	DeepSemECG for robust biomedical signal SC	Outperformed baseline on CDI under noise
	M. Yang et al. (2024)	Superimposed SC for real-time ECG monitoring	98.8% accuracy at 50× compression
Robotics	Li and Aijaz (2025)	ToSC + GenAI agent for multi-robot anomaly detection	Data reduced to 2.5% of raw images
	Wu et al. (2025)	H2R Bridge: few-shot human intent via a fine-tuned vision-language model	92%+ accuracy with 5 examples per action
	Suresh et al. (2024)	XR-integrated SC for human–multi-robot exploration	Underlying SC operates at a rate-normalized SNR up to 20 dB lower than classical
Digital Twin	S. Chen et al. (2025)	Goal-oriented SC with feature & temporal selection	59.5%–80% comm. load reduction (sim.)
	Tang et al. (2024)	DRL-based compact SC for UAV-assisted DT sync	8.23% sync delay; 57% energy savings
	B. Li et al. (2025)	Joint optimization of sync, power, and resources	13.2% sync latency reduction
	Kreuzer et al. (2024)	Review of AI in digital twins (incl. predictive maintenance)	Surveys proactive-maintenance techniques
Metaverse/XR	Y. Huang et al. (2023)	ISCom: interest-aware SC for point cloud streaming	10–22 FPS improvement on mobile
	W. Yang et al. (2024)	Streamlined SC + GenAI for multi-modal XR data	Efficient heterogeneous data transmission
	Chen et al. (2024a)	Cross-modal SC with GNNs and cross-modal attention	High cross-modal semantic similarity across SNRs
	Wang et al. (2024)	Goal-oriented SC for avatar synchronization	95.6% latency reduction
	G. Liu et al. (2023)	Vision-based SC encoding skeleton coordinates	51 bytes vs. 8.243 MB per frame
	Y. Zhang et al. (2026)	LLM-integrated SC for context-aware 6G XR	Goal-directed bandwidth allocation
	Li et al. (2025a)	Federated cloud–edge-end SC for privacy-preserving XR	96% delay reduction; 44% quality gain
Text SC	Xie et al. (2021b)	DeepSC: Transformer-based end-to-end text SC	Superior similarity vs. Huffman+LDPC
	Xie and Qin (2020)	L-DeepSC: lightweight SC via pruning/quantization	12.3 MB to 1.28 MB model size
	Xie et al. (2021a)	MU-DeepSC: multi-user SC for text + image (VQA)	Joint multi-modal multi-user optimization
	Zhou et al. (2022)	Knowledge graph-based SC with S-P-O triples	Structured semantic transmission + reasoning
	C. Liu et al. (2024)	Triplet-based explainable SC with filtering	Reduced data, maintained interpretability
	Shi et al. (2023)	RIS-assisted JSCC for text	Improved BLEU at low SNR
Speech SC	Weng and Qin (2021)	DeepSC-S: attention-based end-to-end speech SC	Higher SDR and PESQ than separate coding at low SNR
	Han et al. (2022)	SC preserving semantic content + speaker identity	Competitive PESQ/MOS, reduced data

simple compression, but by extracting and transmitting only task-essential information. Table 5 provides a consolidated comparison of representative works, and the subsections below examine how ToSC enhances system performance in each domain, with explicit analysis of task-specific requirements, semantic information extraction, and performance metrics. Each subsection concludes with a critical assessment that identifies domain-specific limitations, negative results, and unresolved challenges.

5.1. Autonomous vehicular and vehicular communication networks

Autonomous transportation is a widely studied use case for ToSC. An autonomous vehicle must process sensor data in real time for navigation, obstacle avoidance, and traffic compliance, and transmitting complete raw sensor streams introduces latency and bandwidth overhead that these functions cannot tolerate. What actually feeds a driving decision is far sparser: recognized traffic signs, detected obstacles, and planned trajectories. ToSC extracts and transmits those

elements and discards the rest; success is judged at the driving level, by safety (accident rates), navigation accuracy, and real-time response under varying traffic conditions.

Recent work demonstrates ToSC's potential in driving scenarios, though important caveats apply. Eldeeb et al. (2024) proposed a multi-task SC framework for autonomous vehicles that extracts semantic features from traffic signs for both reconstruction and classification tasks, achieving up to an 89% reduction in transmitted bits while maintaining high task performance. However, serving several driving subtasks from one encoder requires keeping their features separable: the unified multi-task system of Zhang et al. (2022) employs domain adaptation precisely to project each task's features into its own domain and transmit only the task-specific ones, a design burden that grows when additional driving subtasks (lane detection, depth estimation) are added. To handle multi-user and multimodal data, Feng et al. (2025) developed a system for edge-enabled autonomous driving that fuses semantic features from multiple vehicles and data types at an edge server, reducing image processing runtime by over 95% compared to traditional JPEG methods. Lv et al. (2024) focused on efficient communication by developing an importance-aware SC system for vehicular image segmentation. Their VIS-SemCom system prioritizes critical objects like vehicles and pedestrians, cutting transmitted data by about 70% at no cost to segmentation quality. However, all three systems define "task-relevant" features at training time, a definition that can be brittle when unfamiliar obstacle types or unexpected road conditions fall outside the training distribution; the adaptable-compression framework of C. Liu et al. (2023) addresses a complementary axis, jointly adapting the semantic compression ratio and wireless resources to multi-user delay constraints.

Critical assessment. Despite encouraging results, the cited studies share a common limitation: evaluation under controlled or simulated channel conditions. No work has demonstrated ToSC reliability under real-world vehicular channel dynamics (fast fading, Doppler spread, handover). All systems also assume a fixed task taxonomy (predefined sign classes, known obstacle types); performance under open-set conditions (where previously unseen objects appear that were not in the training distribution) remains untested. The safety-critical nature of autonomous driving demands formal guarantees on worst-case performance that no existing ToSC system provides.

5.2. Satellite and UAV-aided communication networks

Satellite and UAV links are where bandwidth scarcity is most clearly a physical limit: a satellite downlink pass delivers a fraction of what the onboard sensors capture, and a UAV's energy budget caps both its computation and its transmissions. The tasks these platforms serve (Earth observation, remote sensing, surveillance missions, real-time control of aerial platforms) therefore force a choice about what crosses the link: geospatial features, detected anomalies, mission-critical status updates, or compressed observations that keep only what analysis needs. Evaluation centers on mission completion rates, data quality retained under the bandwidth budget, and latency for time-critical operations. The surveyed systems split by platform: satellite work concentrates on downlink compression under extreme bandwidth constraints, UAV work on uplink efficiency with energy-constrained onboard processing. Both rely on task-driven feature selection, but differ in their handling of KB synchronization: satellite systems assume ground-station KBs are static, whereas UAV swarms require dynamic KB updates across mobile agents.

Hassan et al. (2024) explore SC for resource-constrained 6G LEO satellite communication in Earth observation missions. They propose a framework that uses joint semantic-channel encoding and optimization algorithms to minimize latency in imagery downlinks while meeting quality-of-service constraints. Results show their scheme effectively reduces latency and compression ratios compared to baselines, while consistently preserving required data quality. For FSO-based satellite

communication, W. Chen et al. (2025) propose an FSO-SC scheme for remote sensing image transmission. Their method uses a VQ-VAE to extract key semantic features. Experiments demonstrated a power gain of over 3 dB and a 60% reduction in communication overhead compared to traditional systems, all while maintaining image quality. To improve efficiency in Satellite-Integrated Internet, Z. Xu et al. (2025) introduce a task-oriented transmission (TUT) framework. Their "utility loss of information (UoI)" metric, which captures freshness and task status, guides the semantic coding. Simulations showed their approach effectively minimizes long-term UoI and power consumption, outperforming several benchmarks.

On the UAV side, where agile, low-latency control dominates the design requirements, Song et al. (2025) proposed a cognitive SC system for UAV-based object detection that uses a knowledge graph (KG) to improve performance with limited onboard resources. A semantic extractor on the UAV sends compressed data, and the KG at the ground server enhances semantic interpretation; across low-SNR AWGN and Rayleigh channels their evaluation reports consistent mean-average-precision gains over an attention-based deep JSCC baseline. An earlier version of this work (Song et al., 2024) reported superior resilience under noisy channel conditions, with a minimal 0.6% accuracy drop compared to over 3% for baseline systems. To train semantic models in a UAV swarm, Xu et al. (2024) designed an SC architecture using hierarchical federated learning (HFL). This design improves reliability by organizing UAVs into clusters for decentralized training. Experiments showed the HFL-powered model significantly outperformed the baseline under dynamic conditions, achieving higher PSNR and lower mean squared error (MSE). Y. Xu et al. (2023) developed a task-oriented semantics-aware (TOSA) framework to improve efficiency when sending command and control (C&C) signals to UAVs. A deep reinforcement learning (DRL) agent learns to drop redundant or less timely packets based on their value. The learned policy cuts resource consumption sharply: it transmits only the most valuable command-and-control signals, keeping the number of transmissions nearly constant as packet repetition grows while achieving nearly the same effective transmission rate as bit-oriented transmission of every packet.

Critical assessment. Satellite and UAV ToSC systems face a unique challenge: the extreme asymmetry between uplink and downlink capabilities makes KB synchronization difficult. Most frameworks assume the ground station's KB is perfectly aligned with the onboard encoder, but orbital dynamics, firmware update cycles, and mission-specific reconfiguration introduce persistent KB drift. The energy constraints of small satellites and UAVs impose hard limits on encoder complexity, yet most proposed systems rely on computationally intensive DNNs. How the reported laboratory performance translates to space-qualified hardware is a question the surveyed work has yet to address.

5.3. Video surveillance

Large-scale video surveillance generates prohibitive volumes of raw video, yet its downstream tasks (person identification, activity recognition, anomaly detection, threat assessment) depend not on pixel-level fidelity but on sparse cues: human pose, facial characteristics, motion patterns, behavioral descriptors. ToSC transmits the cues instead of the pixels, cutting data volume without sacrificing detection accuracy, identification precision, or response time.

For person retrieval, Shah et al. proposed PeR-ViS, a system that uses a semantic description (e.g., clothing, gender) as a query instead of an image (Shah et al., 2021). Their deep learning approach acts as a cascade filter to identify the target individual, achieving a high average IoU and detection rate that surpassed the best prior result. Shao et al. developed TOCOM-TEM, a task-oriented framework for low-latency edge video analytics (Shao et al., 2023). By removing spatial and temporal redundancy, the framework cut the communication cost from roughly 130 KB to 20–50 KB per frame (a 62%–85% reduction) at matched accuracy (~85% multiple-object-detection accuracy) for a

multi-camera pedestrian occupancy prediction task. In static surveillance environments, Nandakishore et al. introduced a technique that uses the unchanging background as shared knowledge (Nandakishore et al., 2024). The system transmits only the dynamic Region of Interest (e.g., a person). This method achieved average bit savings of over 86% and simultaneously improved PSNR compared to a baseline. To address privacy, Huang et al. proposed SPE, a decentralized, semantic privacy-preserving framework for edge nodes (A.Y. Huang et al., 2023). The system generates a semantic scene graph instead of transmitting raw video and trains models by only exchanging parameters. Communication overhead falls accordingly: nodes exchange only model parameters, about 50 MB per node versus local datasets of up to 50 GB, roughly 0.1% of the volume that sharing raw data would require.

Critical assessment. Video surveillance ToSC faces a tension between compression and accountability: discarding “task-irrelevant” visual information at the encoder means the original scene cannot be reconstructed for forensic review or legal evidence. Systems like PeR-ViS that operate on semantic descriptions rather than images inherently lose the ability to verify false positives after the fact. The privacy-preserving approaches (e.g., scene graph transmission) trade one risk for another: while raw video is not transmitted, the semantic graph itself may leak sensitive information about activities and social interactions. The ethical and regulatory implications of semantic-level surveillance remain largely unexplored.

5.4. IoT and industrial IoT environments

IoT deployments face a different constraint: enormous device counts with tiny per-device budgets. An IoT network monitors environments, controls devices, aggregates data, and automates decisions, and for each of these the payload that matters is small: an abstracted sensor reading, a status indicator, an anomaly alert, a control command. ToSC strips the raw-data redundancy around those payloads, and is measured on system responsiveness, energy efficiency, network utilization, and task completion accuracy across the deployment. We consider general IoT first and the specialized Industrial IoT (IIoT) setting next.

To address challenges in the AIoT, Yang et al. (2021) proposed SC-AIT, a semantic communication architecture that transmits only the task-relevant semantic information and discards content irrelevant to the target task. Implemented for image classification, it achieved much lower bandwidth requirements and over 40% classification-accuracy gains when compared to conventional bit-level (technical-level) communication. Xie and Qin (2020) developed L-DeepSC, a lite distributed semantic communication system for resource-constrained IoT devices. A model trained in the cloud is compressed via pruning and quantization for execution on devices, reducing the model size by nearly 90% (from 12.3 MB to 1.28 MB) with minimal impact on performance. However, Hu et al. (2023) show that conventional semantic codebooks are vulnerable to *semantic noise* (adversarial, sample-dependent and sample-independent perturbations that distort the intended meaning) and propose a masked VQ-VAE codebook to restore robustness, a concern compounded in IoT deployments where heterogeneous devices and time-varying inputs stress the shared codebook. To manage the complexity-accuracy trade-off in edge-assisted SC, H. Chen et al. (2024) proposed a framework where an edge server selects the optimal semantic model for each IoT device’s task, achieving a near-optimal solution in simulations. A common limitation across these IoT systems is that they assume stationary data distributions; in practice, sensor drift and firmware updates continuously alter the statistical properties of the data, yet none of the proposed frameworks include mechanisms for online model adaptation.

Smart factories intensify those demands: decisions must be real-time, devices from different vendors must interoperate, and process data is commercially sensitive. IIoT tasks (predictive maintenance, quality control, production optimization, safety monitoring) consume machine health indicators, production quality metrics, and operational

status updates, and the payoff is realized in production: reduced downtime, improved quality, lower cost.

For data privacy in 5G-enabled IIoT, Moqurab et al. proposed Instant Anonymity, a lightweight semantic privacy framework (Moqurab et al., 2022). The perturbation-based technique is suitable for resource-constrained devices and demonstrated superior data utility with much lower query error compared to the traditional k-anonymity model. Focusing on interoperability, Willner et al. argued for semantic-level communication in smart factories (Willner et al., 2017). They proposed using the oneM2M standard, which supports Semantic Web technologies, to implement the Industrie 4.0 Asset Administration Shell concept, enabling the reuse of hundreds of existing formal information models and their linkage to thousands of data sources. To enhance spectral efficiency in IIoT, Zhao et al. investigated a framework combining distributed RIS with probabilistic semantic communication (PSC) (Zhao et al., 2024). By jointly optimizing system parameters, their distributed RIS schemes achieved up to roughly twice the sum semantic-aware transmission rate of a centralized RIS deployment with the same total number of elements. For tool wear monitoring, Zhang et al. (2024) developed an Industrial Semantic Communications (ISC) framework. The system classifies the machine’s state and transmits only a semantic flag during irrelevant “idling” periods. This intelligent filtering reduced end-to-end transmission delay by 12.9% compared to a standard scheme.

Critical assessment. IIoT environments present perhaps the hardest deployment scenario for ToSC: extreme device heterogeneity, limited compute on edge nodes, and the need for long-term (multi-year) model stability. Most cited systems are evaluated on homogeneous device populations with stationary data distributions. In practice, sensor drift, firmware updates, and evolving production processes continuously alter the data distribution, requiring mechanisms for online model adaptation that most current ToSC frameworks do not address, though few-shot adaptive schemes such as that of Han et al. (2025) begin to provide such online adaptation, albeit for shifting channel conditions rather than input-data drift. The interoperability challenge is also underestimated: while semantic standards like oneM2M exist, achieving semantic alignment across thousands of heterogeneous devices from different manufacturers remains an open engineering problem.

5.5. Medical and healthcare systems

Healthcare shows both the promise and the limits of ToSC. A clinician monitoring a patient in real time, reading a scan, planning treatment, or responding to an emergency works from diagnostic features (waveform morphology, biomarker thresholds, abnormality indicators), not from full-fidelity signal reconstruction, so delivering only those features cuts transmission overhead. The difficulty is that in medicine the boundary between “task-relevant” and “task-irrelevant” is less clear-cut than elsewhere: subtle signal features may carry diagnostic significance that is not apparent from the training data.

Y. Yang et al. (2024) introduce an Integrated Sensing, Computing, and Semantic Communication (ISCSC) framework for e-Healthcare. By jointly optimizing beamforming and semantic extraction, their approach aims to maximize the Semantic Secrecy Rate (SSR) for privacy and minimize sensing error, improving the transmission rate by up to 70% compared to traditional benchmarks. Bura and Jamalpur (2025) developed DeepSemComm, a framework for robustly transmitting biomedical signals whose DeepSemECG model preserves the diagnostic integrity of ECG signals in noisy environments, achieving lower Clinical Distortion Index (CDI) than a CNN autoencoder baseline under additive noise. For IIoT-based healthcare, M. Yang et al. (2024) propose a framework for real-time ECG monitoring using superimposed semantic communication, achieving a classification accuracy of 98.8% with a compression rate of 50x. However, healthcare spans both regimes identified in Section 3.6. Open-ended diagnostic reading,

where a clinician inspects the entire image for any abnormality, approaches the small-gap convergence case ($H(X) - I(X;T)$ small), so ToSC offers little advantage over high-fidelity SC. Targeted detection of a subtle, low-variance abnormality is instead the low-variance, high-value case: the diagnostic target sits within a large task-irrelevant background, so a task-aware encoder *can* outperform reconstruction, but only if it is given a relevance signal identifying the target. Absent such a signal, aggressive task-oriented compression risks discarding clinically relevant information, particularly for rare pathologies absent from the training distribution, where the semantic encoder has no basis for judging relevance.

Critical assessment. The two-regime split above has direct deployment consequences. The reported 50x compression ratio for ECG monitoring, while impressive, has not been validated against the full spectrum of cardiac pathologies; rare arrhythmias that were absent from the training set may be filtered out as “task-irrelevant.” Regulatory approval for diagnostic systems that operate on lossy semantic representations (rather than full-fidelity data) remains an unresolved barrier. The security frameworks also focus on data confidentiality but do not address the risk of adversarial semantic attacks that could alter diagnostic outcomes.

5.6. Robot collaboration and human-robot interaction

Multi-robot collaboration and human-robot interaction demand communication that scales with team size. A teammate needs the semantics of the shared goal, not another robot’s raw sensor feed: recognized objects, spatial relationships, motion parameters, gesture interpretations. ToSC transmits exactly that, which keeps coordination within the bandwidth budget.

Li and Aijaz (2025) introduce a framework combining ToSC with a Generative AI (GenAI) agent to enhance multi-robot systems. In an anomaly detection task, robots use a semantic encoder before transmitting images. This reduced data transmission volumes to as little as 2.5% of that required for raw images without giving up classification accuracy. To address data scarcity in human-robot collaboration, Wu et al. (2025) present the H2R Bridge, a framework for understanding human intent from few examples. Their system uses a fine-tuned vision-language model to interpret human actions and translate them into robot commands. With only five training examples per action, it reached over 92% accuracy. For exploration in hazardous environments, Suresh et al. (2024) present a framework integrating human operators with a multi-agent robotic team using extended reality and SC. To avoid overwhelming operators, the system transmits only the essential meaning of exploration data. The chosen SC system was previously shown to operate effectively at a rate-normalized SNR up to 20 dB lower than classical communication, ensuring timely decision-making.

Critical assessment. Multi-robot ToSC systems are evaluated primarily on static task definitions with cooperative agents. Real-world multi-robot deployments face adversarial conditions (competing objectives, communication jamming, sensor spoofing) that are not addressed. The H2R Bridge’s few-shot intent recognition, while encouraging, relies on a pre-trained vision-language model whose failure modes under domain shift (e.g., from laboratory to industrial environments) are not characterized. The reliance on shared KBs among robots creates a single point of failure: if the KB becomes corrupted or desynchronized, the entire multi-agent system may fail simultaneously rather than degrading gracefully.

5.7. Digital Twin synchronization and predictive maintenance

Digital Twin (DT) systems require real-time synchronization between physical assets and their virtual counterparts for predictive maintenance and performance optimization (S. Chen et al., 2024). ToSC addresses the bandwidth saturation inherent in continuous full-state transmission by extracting only task-relevant semantics: state change indicators, degradation signals, and anomaly features. For robotic arm reconstruction in DT environments, S. Chen et al. (2024) proposed a goal-oriented framework that pairs a feature-selection algorithm with deep

reinforcement learning-based temporal selection, transmitting only the most informative features and time steps rather than periodic full-state updates; the extended version (S. Chen et al., 2025) reduces communication load by at least 59.5% (strict reconstruction-error constraint) and 80% (relaxed constraint) in simulation while maintaining reconstruction accuracy.

Tang et al. (2024) proposed a framework using compact SC systems to minimize latency in UAV-assisted DT synchronization. By transmitting only essential information, their DRL-based optimization algorithm reduced DT sync delay by 8.23%, data dropping rates by 15.31%, and energy consumption by up to 57.14% compared to benchmarks. For mobile edge networks, advanced semantic communication frameworks address the mobility of user devices and dynamic channel conditions, achieving up to 13.2% reduction in synchronization latency through joint optimization of synchronization strategy, transmission power, and computational resource allocation (B. Li et al., 2025). Okegbile et al. (2026) instead target the security of this synchronization, integrating split federated learning with task-oriented secure semantic communication and a token-based semantic-defense mechanism that distinguishes adversarial from authentic semantic data during physical-to-virtual updates.

For predictive maintenance, AI-driven digital-twin systems already base proactive decisions on degradation signatures and failure precursors (vibration anomalies, temperature deviations, wear metrics) (Kreuzer et al., 2024; Lv and Xie, 2024; S. Ma et al., 2024); a ToSC formulation would transmit only these precursors rather than complete sensor streams, keeping communication overhead minimal.

Critical assessment. DT synchronization presents a unique temporal challenge: the “ground truth” physical state evolves continuously, but semantic updates are discrete and lossy. Current systems report average synchronization improvements but do not analyze worst-case desynchronization events, precisely the scenarios where predictive maintenance decisions are most consequential. The feature- and temporal-selection approach, while effective for the repetitive motion of robotic-arm reconstruction, remains untested on irregular or previously unseen motion patterns. A deeper gap remains: no formal definition of “semantic synchronization fidelity” exists that captures not just state accuracy but the impact of synchronization errors on downstream maintenance decisions.

5.8. Metaverse and Extended Reality applications

Metaverse and Extended Reality (XR) applications stream visual, audio, haptic, and spatial content in real time, yet only a small fraction of it matters to the user at any instant. ToSC exploits exactly that, extracting user attention, interaction intentions, avatar motion, and environmental context instead of complete sensory streams (Wang et al., 2024).

The works in this domain span three architectural strategies: (1) region-of-interest selection that compresses spatial data by transmitting only user-attended regions (ISCom, skeleton-based systems); (2) cross-modal fusion that jointly encodes multiple modalities through shared semantic representations (graph neural network and Transformer-based approaches); and (3) foundation model-assisted systems that exploit LLM reasoning for context-aware bandwidth allocation. These strategies trade off compression aggressiveness against reconstruction controllability, with the most aggressive approaches (skeleton-only transmission at 51 bytes) achieving orders-of-magnitude compression but sacrificing the ability to reconstruct visual detail for verification.

For XR device optimization, the interest-aware ISCom framework improves point cloud video streaming, raising the minimum rendering frame rate by 10 FPS and up to 22 FPS on mobile devices through region-of-interest selection and lightweight encoder-decoder architectures tailored to resource-constrained XR devices (Y. Huang et al., 2023). Yang et al. integrate semantic awareness with generative AI

to achieve streamlined transmission, using multi-modal fusion and superposition schemes for heterogeneous data (W. Yang et al., 2024).

Cross-modal semantic communication suits Metaverse applications because it integrates text, audio, image, and haptic modalities through graph neural networks and Transformer-based encoders. These systems compress semantic features through cross-modal attention mechanisms while using generative AI for multimodal data reconstruction in virtual scenarios (Chen et al., 2024a). Total transmitted data decreases, and cross-modal semantic similarity stays high under the authors' own similarity measure across the reported channel conditions.

For avatar synchronization, goal-oriented frameworks achieve up to 95.6% latency reduction (Wang et al., 2024), while skeleton-based systems compress avatar data to just 51 bytes (G. Liu et al., 2023). LLM-integrated approaches enable context-aware bandwidth allocation for immersive experiences (Y. Zhang et al., 2026), and federated cloud-edge-end architectures reduce transmission delay by 96% while addressing privacy concerns through distributed semantic processing and 3D scene reconstruction (Li et al., 2025a). The skeleton-based and LLM-integrated systems instantiate strategies (1) and (3) above, while the avatar-synchronization and federated cloud-edge-end frameworks apply the same task-relevance lever to latency and privacy objectives.

Critical assessment. Metaverse/XR applications demand extremely low latency (e.g., sub-20 ms transmission budgets for augmented reality Wang et al., 2024) while simultaneously requiring high visual fidelity for user comfort. Current ToSC systems report impressive compression ratios (e.g., 51 bytes for skeleton data) but do not evaluate perceptual quality from the user's perspective; metrics like MPJPE may be acceptable numerically while producing visually jarring artifacts that break immersion. The reliance on pre-shared base knowledge (e.g., an initial avatar mesh) creates a cold-start problem for new users or environments. The LLM-assisted approaches introduce inference latency that may conflict with the stringent real-time requirements of XR applications.

Text and speech SC, the two modality-level domains, are the last of the ten; their techniques underpin many of the application-oriented systems above.

5.9. Text-based semantic communication

Text-based SC, brought to prominence by the foundational DeepSC system, targets sentence-level meaning preservation across noisy channels for applications including machine translation, dialogue, and information retrieval. Performance is evaluated through Level B metrics (BLEU, sentence similarity, WER), though the gap between n-gram overlap and true meaning preservation motivates embedding-based measures (Section 4).

Building on earlier deep-learning JSCC for text (Farsad et al., 2018), Xie et al. (2021b) introduced DeepSC, the first Transformer-based end-to-end SC system for text, jointly optimizing semantic encoding and channel coding. DeepSC demonstrated that transmitting learned semantic features rather than source-coded bits achieves superior sentence similarity, particularly at low SNR, compared to conventional source-channel separation schemes such as Huffman coding combined with low-density parity-check (LDPC) coding. Subsequent work extended this design in several directions. Xie and Qin (2020) developed L-DeepSC, a lightweight version using pruning and quantization to reduce the model from 12.3 MB to 1.28 MB for resource-constrained IoT deployment. For multi-user scenarios, MU-DeepSC (Xie et al., 2021a) enables simultaneous text and image transmission for tasks such as VQA, with jointly optimized transceivers. Zhou et al. (2022) proposed representing text semantics via knowledge graphs, specifically subject-predicate-object triples, enabling explicit, structured semantic transmission that supports reasoning at the receiver. The triplet-based explainable SC scheme of C. Liu et al. (2024) further developed this direction by filtering redundant triplets to reduce data volume while maintaining interpretability.

A distinct architectural branch applies JSCC to text by mapping sentences directly to channel symbols without intermediate bit representations. The RIS-assisted text SC system of Shi et al. (2023) demonstrated that jointly optimizing semantic encoding with RIS phase shifts improves BLEU scores in low-SNR conditions, illustrating how physical-layer enhancements complement semantic-level design.

Critical assessment. Text SC benefits from the maturity of NLP models, but several limitations persist. Most systems are evaluated on short, well-formed sentences from curated datasets; performance on noisy, informal, or domain-specific text (medical reports, legal documents) is largely uncharacterized. The reliance on pre-trained language model weights as implicit (parametric) knowledge creates the same KB alignment weakness identified across other domains: if the transmitter and receiver use different model versions or fine-tuning, semantic decoding degrades. Evaluation is dominated by BLEU and sentence similarity, neither of which captures pragmatic correctness, i.e., whether the received text leads to the correct downstream action. No benchmark yet evaluates text SC under realistic multi-turn dialogue conditions with evolving context.

5.10. Speech and audio semantic communication

Speech SC targets both speech-to-text (transcription) and speech-to-speech (reconstruction) scenarios, transmitting phonemic content, prosodic features, and speaker identity. Performance spans Level B (WER, PESQ, MOS) and Level C (downstream task accuracy for classifiers or command interpreters).

Weng and Qin (2021) introduced DeepSC-S, among the first deep learning-based SC systems for speech, using an attention-based encoder-decoder trained end-to-end to extract and transmit speech semantics. DeepSC-S demonstrated that semantic-level transmission achieves higher signal-to-distortion ratio and perceptual speech quality than conventional separate source-channel coding, especially at low SNR, by learning squeeze-and-excitation attention weights that emphasize the essential speech information. Han et al. (2022) further improved speech SC efficiency by proposing a system that preserves both semantic content and speaker characteristics, achieving competitive PESQ and MOS scores while reducing transmitted data volume. The system exploits the observation that much of the speech signal is acoustically redundant for meaning extraction (pauses, filler sounds, and predictable phonetic transitions), enabling aggressive compression without semantic loss.

A recent direction integrates neural audio codecs (e.g., EnCodec Défossez et al., 2023, which uses residual vector quantization) into the SC pipeline. These codecs learn discrete token representations of audio that are suited to standard digital transmission while preserving perceptual quality, bridging the analog JSCC and digital modular paradigms identified in Section 3.3.

Critical assessment. Speech SC shares the knowledge dependency challenge: encoder-decoder pairs must be trained on matched data distributions, and performance degrades when the speaker population, language, or acoustic environment at deployment diverges from the training set. We found no work addressing multilingual or code-switching scenarios, though both are common in real-world communication. Evaluation under realistic acoustic conditions (background noise, reverberation, multi-speaker overlap) is sparse. Most critically, speech SC for task-oriented applications (voice commands for robotics, diagnostic interviews in telemedicine) has received minimal attention; the majority of work focuses on reconstruction fidelity rather than downstream task performance, leaving the SC/ToSC distinction for speech largely unexplored.

5.11. Cross-domain synthesis: Recurring patterns and lessons

The critical assessments above reveal several cross-cutting patterns that recur across individual application domains:

Pattern 1: knowledge dependency is the dominant bottleneck. Nearly every domain depends on aligned knowledge (an explicit shared KB where one is maintained, or, for the parametric text- and speech-SC systems, the implicit knowledge embedded in model weights), yet none has solved knowledge synchronization for dynamic conditions (Fig. 7, column 1). Partial solutions exist (Zhou et al., 2022; Farshbafan et al., 2023; J. Zhang et al., 2026), but severity scales with how distributed the system is, suggesting knowledge management should be treated as a first-class system component.

Pattern 2: Reported gains are benchmarked against weak baselines. Across the domains surveyed above, most studies compare ToSC systems against traditional schemes (JPEG, raw transmission) rather than against state-of-the-art learned compression (e.g., neural codecs, VVC). This inflates the apparent advantage of ToSC. A more honest comparison would benchmark against the best available compression method and isolate the additional gain attributable to task-awareness; the severity of this effect nonetheless varies by domain (Fig. 7), being most pronounced where no strong learned baseline is yet established.

Pattern 3: Safety-critical domains expose the brittleness trade-off. Autonomous driving, healthcare, robotics, and industrial monitoring all reveal the same tension: ToSC's efficiency gains come from discarding information deemed task-irrelevant, but the definition of "relevant" is fixed at training time. In domains where the cost of missing information is catastrophic (misdiagnosed pathology, undetected pedestrian, unsafe robot actuation near a human collaborator, unrecognized equipment failure), this brittleness is unacceptable without formal guarantees or fallback mechanisms.

Pattern 4: Evaluation methodology is inconsistent across domains. There is no standardized protocol for evaluating ToSC systems. Different domains use different metrics, channel models, compression baselines, and definitions of "task success". This makes cross-domain comparison impossible and hinders the identification of generalizable design principles. The community would benefit from standardized benchmark suites (cf. the role of ImageNet in computer vision or GLUE in NLP) that enable fair comparison across ToSC implementations. Fig. 7 maps the severity of these four patterns across all ten application domains, summarizing that knowledge dependency is the most pervasive weakness across the surveyed domains.

6. Research trends in ToSC

The trends below respond to weaknesses the preceding sections exposed: the surveyed systems largely assume a fixed task, an aligned and static knowledge base, and a benign channel, and each assumption fails somewhere in deployment. Supplementary Table S4 summarizes these research areas.

6.1. Model integration and semantic inference

6.1.1. Integration of ToSC systems with LLMs and LMMs

The integration of LLMs and LMMs enables ToSC systems to dynamically reason about task requirements, overcoming the fixed task definitions that constrain traditional approaches (Chen et al., 2024b). Foundation models serve three roles: as semantic codecs that transform data across modalities (e.g., image-to-text), as a source of implicit knowledge for task decomposition, and as intelligence engines in distributed systems. W. Yang et al. (2025) proposed M-GSC, where an LLM decomposes complex tasks and specifies standardized semantic representations, enabling scalable multi-user operation. Li et al. (2026) applied goal-oriented keyframe selection for wireless video, reporting PSNR gains of 33%–38% over deep-JSCC video baselines (DeepWiVe

and DVST). Ribouh and Saleem (2025) demonstrated extreme compression (a compression ratio of about 4250×, reducing one image to a 1.4Kb textual description) by converting images to textual descriptions via a vision-language model, with a generative model synthesizing images at the receiver, trading pixel-level fidelity for the large reduction in transmitted data. Du et al. (2025) explored a split architecture with a lightweight vehicular encoder and cloud-based LMM for VQA. Salehi et al. (2025) instead coupled knowledge-graph extraction with LLM-based encoding and decoding for text, with the explicitly transmitted graph serving as the shared knowledge base and the model's parametric knowledge backing it; this reduced the transmitted text volume by 30% while reaching 84% semantic similarity, against 63% for a deep-learning DeepSC baseline. More recent LMM-based designs push further on multimodality and real-world validation: Yao et al. (2026) unify generation and comprehension within a single transmission by replacing the visual decoder with a multimodal language model that captions the received content without a separate text stream, while Ding et al. (2026) report a foundation-model-aided video SC system evaluated over 3GPP-standard channels in simulation and validated over the air on a software-defined-radio testbed, one of the few prototype-level rather than simulation-only evaluations.

This trend loosens ToSC's dependence on task-specific models in favor of knowledge-driven designs (Wang et al., 2023). However, LLM-based codecs introduce new failure modes: hallucinated descriptions that diverge from source content, computational costs that conflict with low-latency requirements, and poorly characterized conditions for when cross-modal transformation preserves task-relevant information. These prominently reported large-model and generative figures warrant the same methodological caution as the domain results of Section 5: each is typically a single-study, single-dataset result, and absolute compression ratios (e.g., the ~4250× of Ribouh et al.) are measured against raw or reconstruction baselines rather than a state-of-the-art learned codec (the weak-baseline caveat, Pattern 2).

The role a foundation model plays in ToSC is clarified by the three-layer knowledge taxonomy introduced earlier: a model's parameters constitute *implicit (parametric) knowledge* rather than an addressable knowledge base, so they can substitute for or back an explicit KB but cannot be queried or synchronized like one. For large-model ToSC, this framing carries practical consequences. Keeping parametric knowledge up to date requires either costly retraining or RAG; RAG couples the model to an external, updatable store of dynamic context, gaining weight-free knowledge refresh at the cost of retrieval latency and a synchronization dependency. The nearest surveyed instantiation is the knowledge-graph-plus-LLM scheme of Salehi et al. (2025), which transmits an explicit graph as the shared KB while the model's parametric knowledge backs it; it is not RAG in the strict sense, since it carries structured knowledge in the message rather than querying an external index at inference. A full retrieval-augmented SC pipeline, in which the encoder or decoder queries an updatable external store in real time, is to our knowledge anticipated rather than empirically demonstrated for semantic links, mirroring the prompt-injection status noted in Section 6.4. Because parametric knowledge is implicit, the transmitter and receiver must also hold *consistent* models: differing model versions or fine-tuning yield divergent decodings of the same semantic representation, a silent failure mode analogous to knowledge-base misalignment. Most limiting for deployment, foundation models impose severe edge resource overhead: frontier foundation models now commonly contain on the order of trillions of parameters, far exceeding edge-node GPU memory, with prohibitive inference latency if invoked on every sample. To address this overhead, H. Zhang et al. (2026) propose a cloud-edge model-evolution scheme that incrementally updates a small edge feature extractor via knowledge distillation from a cloud-resident foundation model, preserving low inference latency at a modest accuracy cost. Y. Sun et al. (2026) likewise target edge deployment with a large-AI-model agent that aligns multimodal data and plans communication policies from natural-language intent, reporting a

Application Domain	KB Dependency	Weak Baselines	Safety Brittleness	Eval Gap
Vehicular	■	■	■	□
Satellite / UAV	■	□	□	□
Surveillance	■	□	□	■
IoT / IIoT	■	□	■	□
Healthcare	■	□	■	■
Robotics	■	□	■	□
Digital Twin	■	□	□	□
Metaverse / XR	■	□	□	□
Text SC	■	■	□	□
Speech SC	■	□	□	□

Legend: ■ Severe □ Moderate □ Minimal

Fig. 7. Cross-domain weakness matrix mapping four recurring weakness patterns across ten SC/ToSC application domains. Columns represent KB dependency, weak baselines, safety-critical brittleness, and evaluation methodology gaps. Cell glyphs indicate severity (filled square = high/severe, half-filled = moderate, empty = low/minimal). Severities are the authors' qualitative synthesis of the per-domain critical assessments in Section 5, not an independent measurement; rows follow the domain order of Table 5. The uniformly severe first column reflects that knowledge dependency (an explicit shared knowledge base where one is maintained, or the implicit parametric knowledge of text- and speech-SC models) is the most pervasive weakness across the surveyed domains.

concrete compute budget on the order of 9.4 TFLOPs (dominated by the language model and driven chiefly by input length, and met within a single consumer-grade GPU such as an RTX 4090), together with a 33% reduction in communication delay, while flagging the limited planning capacity of resource-constrained edge models. Islam and Chang (2026) make the cloud-versus-edge trade-off explicit, pairing a cloud-hosted, API-based large model that attains higher semantic fidelity with a compact on-device model that removes the external API dependency at a measurable fidelity cost.

Together these results make the central deployment tension concrete for latency-sensitive settings such as vehicular networks and industrial IoT: the broad generalization that motivates foundation models strains the compute, memory, and latency budgets of edge devices, though the schemes above (cloud–edge distillation, input-length-aware edge agents) show the tension is real but tractable. These edge-generic budgets must be read against the far tighter envelopes of the latency-critical domains: vehicular V2X messaging operates under millisecond-scale end-to-end latency budgets on embedded-class automotive compute, and industrial IoT control loops run within tight low-power, limited-memory envelopes. In these regimes, even a distilled edge agent must be aggressively compressed, so the cloud–edge distillation and input-length-aware offloading schemes above are prerequisites rather than refinements for these settings. The mechanisms for closing this gap are the same ones that already shrank earlier SC models toward IoT-class budgets: pruning and quantization (L-DeepSC compressed a Transformer text encoder from 12.3 to 1.28 MB Xie and Qin, 2020), knowledge distillation to a lightweight student (C. Liu et al., 2025) including the cloud–edge distillation of H. Zhang et al. (2026), and lightweight-encoder designs with fast on-device training for resource-constrained terminals (Peng et al., 2024a). Applied to a trillion-parameter foundation model, reaching the megabyte-scale memory and millisecond-scale latency envelope of automotive and industrial-IoT edge nodes would demand reductions of several orders of magnitude.

However, deployment-oriented quantities are reported only sporadically across the large-model ToSC systems surveyed here: Sun et al. give a compute budget and Zhang et al. an order-of-magnitude parameter

count, while on-device latency, memory footprint, and energy per inference go unreported for the keyframe-selection, extreme-compression, split-VQA, and knowledge-graph schemes. Moreover, the budgets that do exist are edge-generic: the surveyed large-model ToSC literature reports no vehicular-network- or industrial-IoT-specific inference-latency, compute, or storage figures. Even the one explicitly vehicular scheme, the split-VQA encoder of Du et al. (2025), reports only workstation-GPU response times and FLOPs rather than vehicle-grade deployment budgets, so per-domain deployment envelopes for these latency-critical settings remain an open measurement gap rather than a settled quantity, precisely the gap the standardized-evaluation argument of Section 4 targets.

6.1.2. Intent-based ToSC systems

Intent-based systems convey the underlying purpose of a communication exchange itself, one abstraction above the task-relevant features that earlier ToSC transmits (Guo et al., 2025). Thomas and Saad (2022) integrated neuro-symbolic AI (GFlowNets) to learn causal data structures, enabling nodes to reason about intent and redefining reliability based on meaning preservation. C. Liu et al. (2025) addressed deployment constraints via knowledge distillation, transferring large-model capabilities to lightweight student models with channel-adaptive modules. Aceta et al. (2022) applied this approach to industrial human-machine interaction, using ontology-driven intent interpretation to generate executable commands without large training datasets.

Intent-based communication pushes ToSC toward true Level C operation by transmitting purpose rather than content (Gündüz et al., 2023). The practical limitation is that intent inference requires causal reasoning capabilities that current neuro-symbolic systems provide only for narrow domains; scaling to open-ended, multi-party interactions remains an open challenge.

6.1.3. Explainable ToSC systems

Explainable ToSC replaces opaque feature vectors with structured, human-interpretable representations (C. Liu et al., 2024). Three approaches span different modalities: Ma et al. (2023) use a β -variational

autoencoder to disentangle features into semantically interpretable components with selective transmission; C. Liu et al. (2024) represent text semantics as subject-relation-object triplets with relevance-based filtering; and Sun et al. (2023) extract scene graphs of objects and relationships for visual tasks (DeepSC-SG).

Explainability is essential for trust in safety-critical domains (Kang et al., 2022), but a fundamental tension persists: structured representations (scene graphs, triplets) impose a discrete bottleneck that may discard continuous semantic nuances captured by dense latent vectors, precisely the nuances that determine performance in fine-grained tasks such as medical image analysis.

6.1.4. Generative and diffusion-model-based semantic communication

A shift is underway in semantic communication from *reconstruction-based* to *generation-based* decoding: rather than recovering a faithful replica of the transmitted signal, the receiver uses a generative model to synthesize an output that is perceptually convincing and task-appropriate, conditioned on a compact semantic description transmitted over the channel. This fundamentally changes what is preserved across the link, from source-level fidelity to semantic-level consistency.

Grassucci et al. (2023) initiated this line of research by introducing a diffusion-model-based semantic communication framework in which the transmitter extracts a compact semantic map and the receiver employs a conditional diffusion process to generate high-quality reconstructions from minimal transmitted information. M. Yang et al. (2025) proposed a diffusion-based JSCC scheme with perceptual optimization, demonstrating that diffusion decoders can achieve superior perceptual quality compared to conventional JSCC approaches, particularly at very low channel bandwidth ratios. T. Wu et al. (2024) took a complementary physical-layer approach, introducing a channel denoising diffusion model that operates after channel equalization to learn the channel-input distribution and iteratively remove residual channel noise; combined with a JSCC encoder, it lowers reconstruction error and improves robustness across both Gaussian-noise and fading channels at varying signal-to-noise ratios. More recent generative designs couple the decoder to explicit knowledge and to new modalities: Fang et al. (2026) pair a vision-language large model that extracts a concise cross-modal description with a diffusion reconstructor guided by a shared semantic knowledge base, while Ma et al. (2026) extend generative decoding to arbitrary-length video through a multimodal-large-model keyframe extractor, reaching channel-bandwidth ratios on the order of 10^{-4} – 10^{-3} .

From a theoretical standpoint, generative decoding is closely connected to the rate–distortion–perception trade-off formalized by Blau and Michaeli (2019), which establishes that perfect perceptual quality (i.e., the reconstructed distribution matching the source distribution) requires a higher distortion than the classical rate–distortion bound. Diffusion-based SC systems operate precisely in this regime: they sacrifice point-wise reconstruction accuracy in exchange for outputs that are distributionally faithful to the source, achieving high perceptual quality at extreme compression ratios. This reshapes the boundary between SC and ToSC: since the generated output preserves meaning and utility without preserving the source signal itself, we position generative SC not as a paradigm separate from either but as a *decoding strategy* (generation-based rather than reconstruction-based) that spans both. Where the output is judged by perceptual fidelity it serves general SC (Level B), and where it is judged by downstream task success it serves ToSC (Level C), with the rate–distortion–perception trade-off governing the balance. A broader catalog of generative diffusion models across wireless sensing, transmission, and applications is given by Fan et al. (2026), and a dedicated survey of generative-AI-enabled SC (organized by generative architecture, communication modality, and application task) by Liang and Li (2026).

Practical deployment nonetheless lags. Diffusion models impose substantial computational cost through their iterative sampling process,

which conflicts with the low-latency requirements of many ToSC applications. Controllability is a second obstacle: ensuring that the generated output faithfully reflects the transmitted semantic description, rather than hallucinating plausible but incorrect content, requires careful conditioning mechanisms that are not yet fully understood. Formal information-theoretic guarantees are largely absent as well: unlike classical rate–distortion theory, no established framework bounds the task-performance loss incurred by generation-based decoding under arbitrary channel conditions.

6.2. System-level architecture

6.2.1. Multi-user ToSC systems

Multi-user ToSC introduces challenges in coordinating diverse objectives, classified as single-modal (managing interference across separate tasks) or multimodal (fusing complementary data for shared objectives) (H. Zhang et al., 2023). Xie et al. (2021a) introduced MU-DeepSC for multimodal VQA, fusing image and text features via a MAC network with jointly optimized transceivers, later extending this to a unified Transformer framework supporting both modes (Xie et al., 2022). A related thread serves multiple *tasks* rather than multiple users from a shared encoder: Lyu et al. (2024) unified coding-rate-reduction maximization, which yields discriminative features for classification, with mean-square-error minimization for image recovery in a single deep JSCC objective, and added a gated network that adaptively prunes transmitted features according to channel conditions. For resource-constrained deployment, Peng et al. (2024a) combined federated learning, knowledge distillation, and curriculum learning (Fed-CL-KD). Federated learning more broadly bridges single-node ToSC and distributed network AI, enabling collaborative semantic model training without aggregating raw data (Xiang et al., 2025; S. Zhang et al., 2023).

6.2.2. Cooperative agents in ToSC

Cooperative multi-agent ToSC enables agents to share distributed observations while exploiting task correlations (Razlighi et al., 2024, 2025). Razlighi et al. (2024) proposed a split semantic encoder with a common unit for shared features and specific units for task-specific messages, later extending this to distributed sources with partial observations (Razlighi et al., 2025). He et al. (2023) used deep reinforcement learning to jointly learn what, when, and to whom to transmit, co-optimizing communication protocols with agent policies. A fundamental limitation is that cooperation can be counterproductive when tasks are weakly correlated: the coordination overhead may exceed collaboration gains, particularly under scarce communication resources.

6.2.3. ToSC in edge computing

Integrating ToSC with mobile edge computing (MEC) transforms semantic transmission into a joint communication-computation optimization problem, where edge devices transmit only task-essential semantic information (Y. Xu et al., 2023). Cang et al. (2023) jointly optimized the semantic extraction factor, transmit power, and computation resources to minimize task execution delay. Guo et al. (2024) introduced DTCN, where edge servers act as intelligent semantic relays that fuse device-transmitted features with their own background knowledge before forwarding. Shao et al. (2023) provided theoretical grounding by deriving deterministic information-bottleneck bounds on how much video data must be delivered over the wireless channel to satisfy edge inference accuracy.

6.3. Physical layer and infrastructure

6.3.1. Task-aware encoder and decoder design

A growing body of work designs adaptive encoder–decoders that jointly optimize semantic prioritization with channel state awareness. Wang et al. (2024b) proposed I-JSCC, where the encoder scores each feature's task importance and a higher-level mechanism determines per-user transmission ratios based on both importance and channel conditions. Lyu et al.'s channel-adaptive gated network (Section 6.2.1) applies the same principle, dynamically adjusting data rate to maintain task performance (Lyu et al., 2024). X. Chen et al. (2025) employ a masked auto-encoder that scores data patches by task importance and then selects patches for transmission based on real-time CSI. At the physical layer, Guo and Yang (2022) designed MIMO precoders directly from semantic gradients rather than traditional channel metrics. A complementary robustness-oriented design pairs semantic-aware masking with a discrete codebook to harden the task-oriented link (Li et al., 2025b).

6.3.2. Semantic-aware ToSC systems

Semantic-aware ToSC systems explicitly evaluate the importance of each semantic feature for the downstream task and allocate wireless resources accordingly, via three complementary approaches to quantifying feature importance (C. Xu et al., 2025). *Objective importance scoring* estimates each feature's contribution to task performance and transmits only the most critical ones; Wang et al. (2024b) introduced a metric for task-oriented semantic spectral efficiency (TOSSE) and used reinforcement learning to maximize it. *Application-specific ranking* maps importance to domain constraints: the GSAR framework (Wang et al., 2024) ranks avatar skeleton nodes by their importance to pose recovery and assigns high-priority nodes to more reliable subchannels, while Gholami et al. (2025) allow flexible semantic similarity ranges that adapt compression per user. *Subjective utility models* incorporate human decision biases; Vaidanis et al. (2025) used cumulative prospect theory to link feature importance to an agent's risk perception. A complementary line empirically models the semantic-loss function itself, quantifying how compression and channel conditions degrade downstream task accuracy (Nguyen et al., 2025).

6.3.3. Reconfigurable intelligent surface assisted ToSC systems

Reconfigurable intelligent surfaces (RIS) address a fundamental limitation of ToSC: semantic-level efficiency gains can be negated by poor physical channel quality, and joint designs that optimize source beamforming and RIS phase shifts to form favorable channels for the semantic link directly target this failure mode (J. Ma et al., 2024). However, the key innovation for ToSC is optimizing RIS phase shifts based on the semantic importance of data streams, not merely channel state information (Guo et al., 2024b). Shi et al. (2023) demonstrated that RIS-assisted text SC achieves higher BLEU scores in low-SNR conditions and provides robustness to channel estimation errors. Jiang et al. (2024) introduced an RL-based controller that jointly optimizes subchannel allocation and RIS phase shifts based on both user needs and semantic priorities, achieving superior task performance compared to channel-only optimization. Related surface-assisted designs jointly optimize the surface configuration with semantic objectives, including stacked intelligent metasurfaces for task-oriented links (Huang et al., 2024) and a RIS-aided optimization strategy for semantic communication (Hu et al., 2024). Mao et al. (2026) broaden this surface-assisted design to joint sensing, communication, computing, and control, using RIS and semantic awareness to minimize closed-loop execution latency under a minimum-information-entropy (control-quality) constraint for industrial IoT. This dual optimization of physical channel characteristics and semantic priorities distinguishes ToSC applications of RIS from conventional signal-strength-oriented approaches (M. Liu et al., 2025).

6.4. Security and privacy against semantic attacks

Because ToSC systems are built on learned semantic encoders, decoders, and shared knowledge bases, their security surface differs fundamentally from that of bit-level systems: the attack target is no longer the raw payload but the AI models and the semantic representations they exchange. Transmitting only task-relevant features reduces, but does not eliminate, the exposure of source information: discarding task-irrelevant content shrinks the attack surface, yet the model-inversion attacks examined below show that recognizable source content can still be reconstructed from the transmitted features alone. Moreover, the same learned components expose AI-native vulnerabilities at every stage of the pipeline: inference-time eavesdropping and perturbation, training-time poisoning of the shared model, and large-model-specific manipulation (Guo et al., 2025). We group the most extensively studied threats into four classes with their corresponding defenses, and then flag further AI-native risks that remain largely unexamined for semantic links.

Model inversion and privacy leakage (inference time). An eavesdropper who intercepts transmitted semantic symbols can invert the encoder to reconstruct the source. Chen et al. (2023) formalized this as the *model inversion eavesdropping attack*, reconstructing recognizable images under both white-box (encoder known) and black-box (an inverse network trained from intercepted signal–image pairs) settings and across channel SNRs, and proposed a lightweight defense based on random permutation and substitution of transmitted features. However, their ablation shows that random permutation alone is insufficient: the black-box attacker, whose inverse network is trained on intercepted signal–image pairs, still reconstructs visually recognizable images under permutation-only obfuscation, and it is the combination with random substitution that renders the eavesdropped images unrecognizable under both attack types, so the protection ultimately rests on the secrecy and reliable delivery of the per-transmission permutation–substitution parameters rather than on either transformation alone.

Defenses against inversion broadly follow two routes. The first is representation disentanglement: Wang et al. (2024a) proposed IBAL, which uses the information-bottleneck principle (Eq. (2)) with adversarial training to extract minimal task-sufficient representations that maximize adversary reconstruction distortion; Erak et al. (2026) introduced CLAD, using contrastive and adversarial training to enforce statistical independence between task-relevant and private features; and Erdemir et al. (2022) studied privacy-aware communication over a wiretap channel using a variational-autoencoder-based JSCC, showing through simulation that the learned encoder preserves reconstruction quality at the legitimate receiver while confusing a passive eavesdropper about a sensitive source attribute. The second is feature obfuscation: J. Xu et al. (2025) proposed FSSC, which dynamically encrypts features during semantic extraction and transmission to block reconstructions from intercepted signals, while a prototype adversarial alignment module simultaneously hardens image tasks against adversarial perturbations. Yet these representation- and obfuscation-based defenses are bounded by limits made explicit below. IBAL and Erdemir et al.'s encoder both degrade only the eavesdropper they were trained against: a surrogate whose reconstruction distortion IBAL maximizes, and a passive listener whose sensitive-attribute leakage the variational scheme minimizes. A stronger or differently structured real eavesdropper can reduce the protected margin; adversarial robustness (below) carries the same threat-model dependence. Moreover, CLAD's enforced statistical independence between task-relevant and private features is attainable only insofar as the private attribute is uncorrelated with the task variable: any task-predictive component of a sensitive attribute cannot be removed without conceding task performance, leaving an irreducible privacy–utility floor. And the feature encryption of FSSC, like the permutation–substitution scheme above, ultimately reduces to the secrecy and reliable delivery of its per-transmission keys rather than to the obfuscation itself.

Data poisoning and backdoor attacks (training time). Because the shared semantic model is learned, often collaboratively, an adversary who corrupts the training data or model updates can implant lasting malicious behavior. In the task-oriented multi-user setting, Peng et al. (2024b) showed that malicious participants can degrade a collaboratively trained semantic decoder through label-flipping, clean-label, and semantic-noise-injection poisoning. They further proposed DPAD-MUSC, an adversarial-reinforcement-learning attack–defense game that detects and excludes poisoned samples, recovering on the order of 15%–25% task accuracy relative to no defense, depending on the attack type (over 15% under clean-label, over 20% under label-flipping, and over 25% under semantic-noise-injection attacks). Its scope is the limitation: a detection-and-exclusion scheme confined to data poisoning, it shares the adaptive-evasion coupling flagged below for the Wei et al. detector and does not address the model-update or backdoor attacks treated next.

Backdoor attacks are subtler still: Yang et al. (2026b) demonstrated a *polymorphic* backdoor (SemBugger) in which intensity-graded triggers, injected by poisoning, make a single trigger family produce multiple distinct malicious outputs while preserving benign fidelity (attack success rate above 90% at under 2 dB PSNR degradation); critically, they also derived a certified *semantic-smoothing* defense (which Yang et al. report as the first certified-robustness result for SC backdoors) that suppresses triggers with a provable lower bound on efficacy. However, as with any randomized-smoothing scheme, this certified radius is governed by the injected noise level. In the reported evaluation this clean-transmission cost is small: the defense drives attack success below 1% at only about 0.2–0.4 dB of added PSNR distortion. The operative limitation is therefore not fidelity loss but adaptivity: the guarantee weakens once an adaptive adversary drives the trigger strength beyond the certified region, the very regime such defenses bound but cannot eliminate. On the detection side, Z. Wei et al. (2025) flagged poisoned samples by their deviation in semantic feature space using a threshold-based similarity test that, unlike neuron pruning or pair-constrained schemes, neither modifies the model nor restricts the data format. That test presumes the poison induces a detectable feature-space deviation; an adaptive poisoner optimizing for feature-space stealth could hold the deviation below the threshold, so detection and evasion remain locked in an escalating cycle.

Adversarial examples and robustness (inference time). Small, deliberately crafted perturbations to the input or the transmitted signal can flip a downstream decision without perceptibly altering the source. Such perturbations span two attack surfaces: *input-space* changes applied before encoding, and *over-the-air* signal-domain perturbations crafted to survive the wireless channel, the latter studied extensively for deep-learning-based wireless and NextG links (Sagduyu et al., 2023; Sadeghi and Larsson, 2019; Flowers et al., 2020; Kim et al., 2022). The dominant defense is adversarial training (exposing the model to worst-case perturbations during learning), which Wang et al. (2024a) instantiate as adversarial learning that trains the encoder to fool a model-inversion adversary by maximizing its reconstruction distortion, and J. Xu et al. (2025) fold into their obfuscation framework. The same primitive doubles as a robustness mechanism against non-adversarial channel distribution shift: K. Wei et al. (2025) used projected-gradient adversarial training to mimic diverse channel distortions, improving classification accuracy by more than 10% under unseen channel states without increasing model complexity. These gains are nonetheless bounded by the threat model seen during training: adversarial training hardens the model only against perturbations resembling those it was optimized against, so an adaptive attacker using a larger budget or an unmodeled perturbation geometry can still induce errors. The defense raises the attacker's cost rather than eliminating the vulnerability. This convergence, with adversarial training serving both security and channel robustness, illustrates how AI-native defenses can be co-designed with the communication objective rather than added on afterward.

Large-model-specific threats. As Section 6.1.1 discusses, foundation models increasingly serve as semantic codecs and reasoning engines, importing the security vulnerabilities of large models into the communication link. Hoang and Le (2025) showed that a backdoor embedded during pre-training of an LLM-based (BERT/ROBERTa) semantic transceiver *persists through downstream fine-tuning* and resists model pruning, by mapping triggered inputs to a predefined output representation, a threat with no analogue in fixed-codebook SC. Beyond backdoors, prompt-injection and instruction-manipulation attacks, in which transmitted content is parsed by the receiver model as instructions rather than as data, are a recognized large-model risk still largely unexamined in the SC setting. They are most acute for the agentic ToSC systems of Section 6.1.1: an agent that plans communication policies from natural-language intent (Y. Sun et al., 2026), together with any tool use it triggers, exposes an instruction channel an attacker can hijack. Candidate mitigations (isolating instruction from data, authenticating message provenance, and policy-filtering agent actions) are established for general large-model systems but not yet adapted to semantic links; this threat is therefore anticipated rather than empirically demonstrated, and remains an urgent open problem. These large-model threats are especially consequential because the same parametric knowledge is reused across tasks and users, so a single compromised model can propagate errors system-wide.

Further AI-native threats (largely unexamined). The four classes above are those with demonstrated SC/ToSC instantiations and defenses; several AI-native risks documented for learned models more broadly remain open for semantic links and should not be read as out of scope. We use *model inversion* above in the communication-specific sense of inverting the encoder to reconstruct the transmitted *input*; distinct from it are the classical attribute model-inversion attack, which recovers training-data attributes from model access (Fredrikson et al., 2015), and membership-inference attacks, which test whether a record was used in training (Shokri et al., 2017), both pertinent once a shared semantic model or knowledge base is exposed. Model-extraction (stealing) attacks can clone the encoder or KB from query access (Tramèr et al., 2016), against which credential-based access control that renders a stolen semantic model low-fidelity for unauthorized users is an early defense (Yang et al., 2026a); gradient-inversion attacks recover inputs from shared gradients (Zhu et al., 2019), and the closely related embedding-inversion attacks recover them from the intermediate representations that semantic systems transmit; and availability attacks (denial of service, and sponge inputs that inflate inference energy and latency (Shumailov et al., 2021)) target the real-time guarantees ToSC depends on. Most SC-specific of all, the shared knowledge base or retrieval-augmented store that Section 6.1.1 couples to large-model ToSC is itself an *integrity* target distinct from training-time model poisoning: an adversary who poisons or tampers with entries in the shared KB or the retrieval index can steer decoding system-wide at inference time without ever modifying the model weights, a threat with no analogue in bit-level systems and directly implied by this survey's own finding that knowledge dependency is the dominant cross-domain bottleneck (Section 5). Beyond integrity, the store's confidentiality is also at stake: Hu et al. (2026) demonstrate a reconstruction attack in which an adversarial large-language-model API provider recovers nearly 90% of a retrieval-augmented knowledge base through adaptive embedding-space queries. Authentication and provenance for KB entries, together with integrity verification of the retrieval index, are the natural defenses but remain unaddressed for semantic links. Watermarking and provenance for transmitted semantic models, needed to attribute a leaked or tampered model, are still at an early stage. These are named open problems (Section 7), not problems already addressed.

These vulnerabilities are most damaging precisely in the safety-critical domains of Pattern 3 (Section 5), where ToSC is most attractive and where a reconstructed private attribute, a poisoned task decision, or a backdoored large model can cause physical harm rather than mere

data loss. Two lessons cut across all four threat classes. Security in ToSC is inseparable from the learning pipeline: defenses must protect the model and its shared knowledge, not merely encrypt the channel. And every defense surveyed trades against task utility or computational cost, with formal, verifiable guarantees still rare; certified defenses such as that of Yang et al. (2026b) are the exception rather than the rule. Because several of these threats can co-occur in practice, recent designs also compose multiple specialized defenses within a single system, for instance a mixture-of-experts that activates different security experts according to each node's declared security requirements (He et al., 2026b). Standardized adversarial benchmarks and verifiable, certifiable defense mechanisms are therefore a priority research direction (Section 7).

7. Challenges and future research opportunities

The research trends in Section 6 and the cross-domain analysis in Section 5 reveal a field that has made rapid progress in demonstrating ToSC's potential but has not yet resolved the foundational issues preventing deployment at scale. The remaining obstacles are of three kinds, and they are not equally mature: the theory has gaps, the algorithms face known but unsolved barriers, and real-world deployment has barely been attempted. For each, we identify the specific research trends it affects and propose concrete research directions, prioritized by their potential to unblock multiple downstream problems simultaneously.

7.1. Challenges in ToSC

7.1.1. Challenges in theory

The theoretical foundations of ToSC remain incomplete in ways that directly constrain the research trends discussed in Section 6. We identify four foundational gaps, ordered by their downstream impact:

- **Definition and Quantification of SI:** the classical framework of Shannon (1948) treats all bits as equally valuable, but the IB formulation (Eq. (2)) reveals that the *relevant* information content of a source X with respect to a task T is $I(X; T)$, which can be far smaller than $H(X)$. No universally accepted measure of this task-relevant semantic content exists. Existing proposals, including logical informativeness from SIT (Bar-Hillel and Carnap, 1953), truthfulness-weighted informativeness from SSIT (Floridi, 2004), and mutual-information-based relevance from the IB framework (Tishby et al., 2000), each capture different facets but have not been unified into a single, computable metric that applies across modalities and tasks. A viable semantic entropy measure must satisfy three requirements: (1) it must reduce to Shannon entropy when the task is full reconstruction ($T = X$), (2) it must decrease monotonically as task-irrelevant information is discarded, and (3) it must be estimable from finite samples via variational bounds (Alemi et al., 2017). Developing and validating such a measure is the deepest open problem here: until one exists, task-relevant information can be bounded (Eq. (2)) but not directly measured.
- **Modeling Semantic Noise:** Semantic noise arises not from channel impairments but from mismatches in interpretation between sender and receiver. Bao et al.'s semantic channel model (Section 2) formalized one source of this noise: when the transmitter's and receiver's KBs diverge ($K_s \neq K_d$), the receiver's semantic-interpretation term $\bar{H}_s(Y)$ falls (the receiver recovers less model-theoretic meaning from the same channel output under its mismatched knowledge base), so the semantic channel capacity C_{sem} degrades independently of the physical channel. Güler et al. (2018) further showed that under asymmetric information, the achievable meaning recovery is bounded by the Bayesian game equilibrium, which may be far below the semantic capacity achievable with aligned KBs. Yet current theoretical models treat

semantic noise either as a fixed KB offset or as an additive perturbation analogous to channel noise; neither captures the *compositional* nature of real-world semantic mismatches, where a single misaligned concept (e.g., confusing "high vibration from tool wear" with "vibration from an adjacent press" in an IIoT setting) can cascade into task failure. Recent work begins to move in this direction: Fayaz et al. (2026) decompose semantic noise into source-, channel-, and receiver-side components and show that source- and receiver-side noise can dominate task performance even when channel-induced semantic errors are negligible. New models must represent semantic noise as a structured, concept-level perturbation that interacts with the task variable T , not merely the source X .

- **Semantic Rate-Distortion Theory:** Classical rate-distortion theory characterizes the minimum rate needed to reconstruct a source within a fidelity constraint, but ToSC requires characterizing the minimum rate needed to *perform a task* within a performance constraint, a problem of a different kind when the task requires strictly less information than the source (i.e., $I(X; T) < H(X)$, the regime where SC and ToSC diverge, as analyzed in Section 3.6). This problem has deep roots in information theory. The remote source coding framework, originating with Dobrushin and Tsybakov (1962) and Witsenhausen (1980), established rate-distortion bounds for compressing a source observed through a noisy mapping, precisely the setting of task-oriented compression where the encoder observes raw data X but cares only about a downstream variable T . The CEO problem (Courtade and Weissman, 2014) extends this to distributed settings directly applicable to multi-agent ToSC. Stavrou and Kountouris (2023) recently formalized a rate-distortion approach to goal-oriented communication, providing the closest existing theoretical treatment of what the field needs. Liu et al. (2022) further formulated rate-distortion bounds when the transmitter observes only an indirect version of the semantic source, proving divergence from Shannon limits. Blau and Michaeli (2019) demonstrated an additional complication: rate-distortion bounds conflict with perceptual quality metrics, meaning that minimizing reconstruction error forces divergence from true semantic perception. The remaining gap is therefore narrower than sometimes claimed: it is specifically a *KB-conditioned* extension of these known frameworks that incorporates KB alignment quality as a parameter. When KBs are perfectly aligned, the encoder can exploit shared knowledge to achieve lower rates; when KBs degrade, the effective rate-distortion bound shifts, requiring higher rates to compensate for semantic ambiguity. Developing such KB-dependent bounds (cf. Priority 1 in the research roadmap below) would provide the first principled benchmarks for ToSC system design.
- **Linking Semantics to Task Performance:** The IB framework provides the formal structure for this link, since $I(Z; T)$ measures how much task-relevant information the representation Z preserves, but translating this into closed-form performance predictions for practical systems remains open. The concrete divergence evidence from Section 3.6 (Cases 1–3) demonstrates that the relationship between semantic representation quality and task performance is non-trivial. Kang et al. (2022) showed that task-oriented and reconstruction-oriented encoders *diverge in what they preserve*, with the personalized task-oriented encoder transmitting only the user-relevant scene-graph subgraph (on the order of a 64% communication-cost reduction) while a reconstruction objective must retain the whole image. Lyu et al. (2024) reported a qualitative accuracy-PSNR trade-off (their Fig. 10) on which classification accuracy and reconstruction PSNR are not simultaneously maximized. KB mismatch introduces an additional confounding factor that current models ignore: even a representation with high $I(Z; T)$ under aligned KBs may yield poor task performance when $K_s \neq K_d$, because the receiver misinterprets

the semantic content; analogously, a representation with high $I(Z;T)$ on the training distribution can degrade sharply under the domain and semantic shifts encountered at deployment (H. Li et al., 2025). Establishing rigorous, KB-aware performance bounds that predict task accuracy as a function of representation quality, channel conditions, and KB alignment is essential for moving ToSC from empirical demonstration to principled engineering.

7.1.2. Challenges in algorithm development

The research trends in Section 6 demonstrate advanced algorithmic solutions, but several fundamental barriers remain that cut across all trend areas:

- **Joint Source-Channel Coding:** ToSC requires a joint approach where semantic content is encoded and protected against channel impairments simultaneously, often through end-to-end learning (Gündüz et al., 2023). The resulting optimization problem is non-convex: it involves tuning the high-dimensional, non-linear parameters of neural networks for both the source and channel coding aspects. Kostina and Verdú (2013) showed that JSCC can strictly outperform separate coding in the finite blocklength regime, yet the optimal joint mapping remains intractable for DNN-parameterized systems. A specific technical gap is the difficulty of backpropagating gradients through non-differentiable wireless channel layers: current solutions often rely on simplified, differentiable channel approximations or reinforcement learning, which may not capture the full complexity of real-world effects like deep fades or non-Gaussian interference, limiting the robustness of the learned solution.
- **Robust Optimization Under Uncertainty:** ToSC systems must operate reliably in real-world situations involving inherent uncertainty, including variations in wireless channel conditions, environmental interference, and, uniquely, *semantic noise*. Hu et al. (2023), motivating their robust masked VQ-VAE codebook, report that conventional semantic systems are degraded by such noise: adversarial, sample-dependent and sample-independent perturbations that distort the intended meaning, arising from ambiguities in language or data representation, mismatches in contextual understanding, or discrepancies in interpretation between sender and recipient. This requires optimization techniques that maintain performance under these stochastic variations while guaranteeing the intended meaning is accurately conveyed and used despite potential semantic ambiguities.
- **Achieving Robust and Resource-Aware Multi-Task Generalization:** Developing ToSC systems that perform a wide range of tasks, potentially simultaneously, and still generalize to new, unknown tasks or changing environmental/network conditions requires algorithms that learn representations beneficial for multiple goals while concurrently optimizing often-conflicting objectives. Transfer across heterogeneous tasks is the hard case; for example, a semantic representation optimized for a classification task may be entirely unsuitable for a generative one, such as image reconstruction. The resource-aware dimension is partially addressed by the adaptable-compression framework of C. Liu et al. (2023), which jointly adapts the semantic compression ratio and wireless resources across users and channel conditions, though transferring a single semantic representation across heterogeneous tasks remains largely open; versatile deployment depends on balancing these competing trade-offs.
- **Lack of Unified Semantic-Related Performance Evaluation Metrics:** Algorithmic development is also slowed by the lack of broadly acknowledged, unified performance metrics for reliably assessing semantic-related performance (Barbarossa et al., 2023). Unlike traditional communication systems, which rely on measurements such as bit error rate, ToSC needs measures that both

reliably correlate with actual task success across applications and are mathematically well-founded. This absence of uniform criteria makes it difficult to objectively compare algorithmic approaches and steer the optimization process.

- **Lack of Interpretability and Explainability for DL-based ToSC Systems:** Many modern ToSC algorithms, especially those based on deep learning, operate as “black boxes” (Ma et al., 2023; C. Liu et al., 2024): this opacity makes it difficult to grasp how semantic features are detected, retrieved, and processed and how these learned representations contribute to task outcomes, in turn making it difficult to diagnose errors, improve algorithm design, or establish trust in these systems, especially in critical applications. Future research should emphasize algorithms that incorporate explainable AI methodologies to clarify semantic feature selection and decision-making.
- **Computational Complexity and Scalability:** Modern semantic processing frequently relies on deep DNNs and complex optimization techniques that require significant compute power and memory, making implementation on resource-constrained devices (IoT sensors, mobile devices), and scaling to vast, heterogeneous networks, hard on both the algorithmic and the engineering side (Xie and Qin, 2020; C. Liu et al., 2025). Still needed are lightweight algorithms, model compression techniques, and distributed learning frameworks that reduce computing cost while maintaining semantic fidelity and task performance.

7.1.3. Challenges for real-world deployment

The gap between the laboratory results reported in Sections 5–6 and operational deployment is substantial. The cross-domain patterns identified in Section 5 (knowledge dependency and brittleness most directly, with the two evaluation-side patterns addressed under Priority 2 below) manifest as concrete engineering challenges:

- **Hardware and Computational Constraints:** Many advanced semantic processing algorithms, particularly those based on DNNs, are computationally intensive, demanding substantial processing power and memory. Implementing them on edge devices or under limited power budgets demands hardware-friendly implementations and efficient algorithmic optimizations that cut latency and energy consumption while preserving semantic fidelity.
- **Interoperability with Legacy Systems and Standardization:** Existing communication infrastructures are largely built upon traditional, bit-centric protocols, so integrating a new semantic layer requires ensuring interoperability and coexistence with existing standards and hardware, via interfaces and transition strategies that allow a gradual shift toward semantic communication. The lack of standardization for SI representation, protocols, and KBs further complicates interoperability between systems from different vendors or designed for different applications; future deployment should rely on the development of common standards and dynamic resource management frameworks.
- **Environmental and Channel Variability:** Real-world wireless networks are characterized by highly dynamic and unpredictable conditions, including fluctuating channel quality, varying interference levels, and user mobility, so ToSC systems that consistently extract, transmit, and interpret semantic content despite these variations typically need real-time adaptive mechanisms that continuously adjust semantic models and communication strategies to the current operational environment.
- **KB Alignment and Management Problem:** Deployment also faces the KB alignment and management problem: complete and accurate KBs are expensive to build, and even once established, independent updates or differing training experiences pull distributed transmitters and receivers apart, degrading performance; dynamically updating KBs across distributed systems compounds the difficulty. A related issue is the inherent inconsistency between




Fig. 8. Prioritized research roadmap for semantic communication, structured as a top-down dependency hierarchy. Priority 1 (semantic rate-distortion theory with KB-dependent bounds) is the enabling root at the top, unblocking standardized benchmarks (Priority 2) and, via them, KB management design (Priority 3). Priority 3 in turn unblocks both robust JSCC for safety-critical deployment (Priority 4) and foundation model-based universal codecs (Priority 5). Arrows indicate enabling dependencies: solving lower-layer priorities removes barriers for higher-layer ones. The ordering reflects the analysis in Sections 5 and 7, where each priority is ranked by the number of downstream challenges it unblocks.

source and destination KBs, even assuming real-time sharing: goal-oriented semantic-language frameworks assume the shared language and KB track local, time-varying goals (Kalfa et al., 2021), while Farshbafan et al. (2023) proposed curriculum learning to actively synchronize a common language but at significant alignment overhead. As such, ToSC systems, particularly multi-user ones, must be designed to operate despite KB discrepancies, a requirement none of the surveyed designs yet meets robustly.

- **Security and Privacy Concerns:** Transmitting semantically rich information introduces new security vulnerabilities and privacy risks: malicious actors can manipulate the semantic content to cause system failure or inject harmful instructions, while even unmanipulated semantic features can inadvertently reveal sensitive personal or operational details. This requires privacy-preserving semantic compression that sanitizes data while retaining task-relevant utility, together with authentication and integrity protection for the semantic content itself and for the shared knowledge base and retrieval store against poisoning or tampering.

7.2. Prioritized research directions

Rather than enumerating all possible future directions, we distill the challenges above into five high-priority research problems, ranked by the number of downstream barriers each unblocks and informed by the cross-domain patterns identified in Section 5 and the limitations observed across the research trends in Section 6.

Priority 1 [Theoretical]: Semantic rate-distortion theory with KB-dependent bounds. This is the theoretical contribution the field most needs. A formal framework linking transmission rate, task-specific distortion (not signal-level distortion), and KB alignment quality would provide: (a) fundamental limits against which all ToSC systems can be benchmarked, (b) principled guidance for encoder design rather than empirical trial-and-error, and (c) a formal basis for the SC/ToSC trade-off analyzed in Section 3.6. Viable directions include extending

the information bottleneck framework (Eq. (2)) to incorporate KB mismatch as a noise source (H. Li et al., 2025), and using variational methods (Alemi et al., 2017) to approximate the resulting bounds for practical DNN-based systems.

Priority 2 [Engineering]: Standardized benchmarks and evaluation protocols. The cross-domain analysis in Section 5 identified inconsistent evaluation as a field-wide problem. The community urgently needs: (a) public, multi-modal datasets with both reconstruction and task annotations, enabling fair SC-vs-ToSC comparison; (b) a standardized evaluation protocol specifying channel models, compression baselines (including state-of-the-art neural codecs, not just JPEG), and reporting requirements; (c) metrics that capture cross-task generalizability, not just single-task performance; and (d) an open, hosted evaluation platform (a public leaderboard with reference implementations and held-out test channels) that reports deployment-oriented measures (inference latency, model size, energy per inference, and robustness to KB mismatch) alongside task accuracy, so that comparisons reflect practical deployability rather than benchmark accuracy in isolation. A template for this exists in adjacent fields (ImageNet, GLUE, LibriSpeech). Establishing such benchmarks and the platform to run them would accelerate reproducible research and enable the field to move beyond inflated comparisons against weak baselines.

Priority 3 [Algorithmic/Engineering]: Scalable KB management with formal consistency guarantees. Knowledge dependency emerged as the most pervasive bottleneck across the surveyed application domains. Deployment at scale requires distributed KB management systems that provide: (a) provable consistency guarantees under network partitions and asynchronous updates, drawing on techniques from distributed databases and consensus algorithms, specifically conflict-free replicated data types (CRDTs) for eventual consistency without coordination overhead, Raft-based consensus for strong consistency in small clusters, and vector clock-based versioning for detecting and resolving KB divergence across asynchronous agents; (b) graceful degradation protocols that detect KB drift (e.g., via Bayesian KB alignment monitoring) and increase transmission redundancy or fall back to SC

mode; (c) lightweight KB synchronization mechanisms suitable for resource-constrained devices (IoT sensors, UAVs). Engagement with standardization bodies (ETSI ISG ZSM, O-RAN Alliance, the IEEE P1876 working group on semantic communication standards, and the 3GPP SA1 Study Item on AI/ML-assisted communication) will be essential for ensuring that KB management solutions achieve cross-vendor interoperability.

Priority 4 [Engineering]: Robust JSCC with formal safety guarantees for critical applications. The brittleness trade-off identified in safety-critical domains (Pattern 3) demands JSCC algorithms that provide: (a) formal worst-case performance bounds, not just average-case improvements; (b) graceful fallback to higher-fidelity transmission when task confidence drops below a threshold; (c) robustness to non-differentiable channel effects (deep fades, non-Gaussian interference) that current gradient-based training cannot handle; and (d) certified robustness to adversarial, poisoning, and backdoor attacks, supported by standardized adversarial benchmarks, so that the AI-native defenses surveyed in Section 6.4 can be compared and verified rather than assessed ad hoc. Hybrid model-driven and data-driven approaches offer a path toward interpretable, certifiable JSCC systems: deep unfolding (Monga et al., 2021) maps iterative optimization algorithms onto trainable neural network layers, preserving convergence guarantees while learning channel-adapted parameters; neuro-symbolic methods (Thomas and Saad, 2022) combine learned feature extraction with rule-based reasoning to provide verifiable decision paths.

Priority 5 [Algorithmic]: Foundation models as universal semantic codecs. Large-scale foundation models (LLMs and LMMs) offer a potential solution to the KB alignment problem by substituting for an explicit, shared knowledge base. The key research challenges are: (a) efficient adaptation and distillation of foundation models for edge deployment, achieving the compression benefits demonstrated in Section 6 without the computational overhead; (b) formal analysis of when a foundation model's parametric knowledge provides sufficient alignment for ToSC tasks vs. when domain-specific fine-tuning is required; (c) security implications of relying on centralized foundation models as a substitute for shared KBs (single point of attack, model and knowledge-base/retrieval-store poisoning). A closely related challenge is the emergence of generative and diffusion-model-based decoding (Section 6.1), which fundamentally shifts the SC/ToSC boundary by replacing source-level reconstruction with semantic-level generation; extending the rate-distortion-perception trade-off (Blau and Michaeli, 2019) to formally bound task-performance loss under generative decoding remains an open theoretical problem that intersects with Priorities 1 and 4. Fig. 8 visualizes the dependency ordering among the five priorities.

8. Conclusion

This survey aimed to answer the questions that catalog-style surveys of SC leave open: when the distinction between SC and ToSC actually matters, where current systems fall short, and what the field must prioritize to move from laboratory demonstrations to deployable technology. Three cross-cutting findings emerge from our analysis, synthesizing the conditional SC/ToSC distinction of Section 3 with the four cross-domain weakness patterns of Section 5: the third finding consolidates the weak-baseline and evaluation-inconsistency patterns, while the brittleness pattern is carried forward as research Priority 4 (robust JSCC with formal safety guarantees).

First, the SC/ToSC distinction is substantive but conditional. Synthesizing the information bottleneck framework (Section 3.6) with empirical evidence from the surveyed literature, we identified that SC and ToSC diverge whenever the task requires strictly less information than the source contains, a condition that holds for most practical applications involving high-dimensional data and low-dimensional tasks. However, the distinction collapses in low-rate regimes and for

tasks requiring full source fidelity (e.g., open-ended diagnostic reading of medical images), whereas for low-variance, high-value features (e.g., detecting a tiny lesion within a large irrelevant background) extreme compression can instead drive the two apart; the choice between SC and ToSC should therefore be an explicit engineering decision based on the task-information gap, not a default assumption.

Second, knowledge dependency is the field's most underestimated weakness. Nearly every application domain we surveyed depends on aligned knowledge, whether an explicit shared knowledge base or the implicit knowledge embedded in model weights, yet none has solved knowledge management for real-world conditions. The theoretical framework of Bao et al. shows that KB mismatch introduces semantic noise that degrades capacity independently of physical channel quality, and our cross-domain analysis confirms this as the most pervasive bottleneck. We further argue that ToSC systems, by discarding "redundant" information an SC receiver could otherwise use to absorb KB errors, are *plausibly more* sensitive to this problem than general SC systems (a conjecture grounded in the information-bottleneck view rather than a direct empirical comparison, whose rigorous test depends on the KB-dependent bounds and a controlled SC-vs-ToSC robustness benchmark we identify as open problems in Priorities 1–2), with direct implications for deployment strategy.

Third, the field's evaluation methodology systematically overstates ToSC's advantages: as the weak-baseline and evaluation-inconsistency patterns of Section 5 document, reported gains are typically measured against weak baselines (traditional codecs or raw transmission) under controlled channel conditions, with no standardized protocol for cross-system comparison. Establishing rigorous, standardized benchmarks is not merely optional but a prerequisite for the field's scientific credibility.

Based on these findings, we proposed five research priorities ordered by the number of downstream barriers each unblocks: semantic rate-distortion theory with KB-dependent bounds, standardized benchmarks, scalable KB management with consistency guarantees, robust JSCC with formal safety guarantees, and foundation models as universal semantic codecs. These priorities are dependency-ordered, not a wish list. SC and ToSC change what 6G communication systems optimize for: meaning and task effectiveness instead of bit fidelity. The efficiency gains documented across ten application domains suggest the practical value of this shift, although their magnitude is likely overstated by the prevalent benchmarking against weak baselines (Section 5) and remains to be confirmed against state-of-the-art neural codecs. However, realizing this potential requires the field to move beyond demonstrating what ToSC *can* do under favorable conditions and instead establish what it *guarantees* under realistic ones. The analytical framework, critical assessments, and prioritized research directions presented in this survey are intended to accelerate that transition.

CRedit authorship contribution statement

Quang Tuan Do: Writing – review & editing, Writing – original draft, Visualization, Validation, Methodology, Investigation, Formal analysis, Conceptualization. **Tung Son Do:** Writing – original draft, Investigation, Formal analysis, Conceptualization. **Thanh Phung Truong:** Writing – original draft, Validation, Methodology. **Dongwook Won:** Writing – original draft, Investigation, Conceptualization. **Nhu-Ngoc Dao:** Validation, Supervision, Methodology. **Wonjong Noh:** Writing – review & editing, Validation, Supervision. **Sungrae Cho:** Validation, Supervision, Resources, Project administration.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work the authors used Claude (Anthropic) in order to improve the language and readability of the manuscript and to assist in editing during the revision. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported in part by the IITP (Institute of Information & Communications Technology Planning & Evaluation) - ITRC (Information Technology Research Center) (IITP-2026-RS-2022-00156353, 40%) grant funded by the Korea government (MSIT) and in part by the NRF (National Research Foundation of Korea) (Grant No. RS-2024-00453301, MASC: Convergence of Multimodal AI and Semantic Communications for Hyper-Connected Hyper-Intelligent Public Safety Systems, 30% / No. RS-2026-25471520) grant funded by the Korea government (MSIT) in 2026.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.engappai.2026.116063>.

Data availability

No data was used for the research described in the article.

References

- Aceta, C., Fernández, I., Soroa, A., 2022. KIDE4I: A generic semantics-based task-oriented dialogue system for human-machine interaction in Industry 5.0. *Appl. Sci.* 12 (3), 1192.
- Alemi, A.A., Fischer, I., Dillon, J.V., Murphy, K., 2017. Deep variational information bottleneck. In: *Proc. Int. Conf. Learn. Represent.*. ICLR.
- Bao, J., Basu, P., Dean, M., Partridge, C., Swami, A., Leland, W., Hendler, J.A., 2011. Towards a theory of semantic communication. In: *Proc. 2011 IEEE Netw. Sci. Workshop*. IEEE, pp. 110–117.
- Bar-Hillel, Y., Carnap, R., 1953. Semantic information. *Br. J. Philos. Sci.* 4 (14), 147–157.
- Barbarossa, S., Communiello, D., Grassucci, E., Pezone, F., Sardellitti, S., Di Lorenzo, P., 2023. Semantic communications based on adaptive generative models and information bottleneck. *IEEE Commun. Mag.* 61 (11), 36–41.
- Beck, E., Bockelmann, C., Dekorsy, A., 2023. Semantic information recovery in wireless networks. *Sensors* 23 (14), 6347.
- Blau, Y., Michaeli, T., 2019. Rethinking lossy compression: The rate-distortion-perception tradeoff. In: *Proc. Int. Conf. Mach. Learn.*. PMLR, pp. 675–685.
- Boursoulatte, E., Kurka, D.B., Gündüz, D., 2019. Deep joint source-channel coding for wireless image transmission. *IEEE Trans. Cogn. Commun. Netw.* 5 (3), 567–579.
- Bura, V.K., Jamalpur, B., 2025. DeepSemComm: a deep learning-based semantic communication framework with case studies in text and biomedical signal transmission. *Eur. Phys. J. Plus* 140 (6), 597.
- Cang, Y., Chen, M., Yang, Z., Hu, Y., Wang, Y., Zhang, Z., Wong, K.-K., 2023. Resource allocation for semantic-aware mobile edge computing systems. In: *Proc. 2023 IEEE Globecom Workshops*. GC Wkshps, IEEE, pp. 1585–1590.
- Chen, H., Fang, F., Wang, X., 2024. Semantic extraction model selection for IoT devices in edge-assisted semantic communications. *IEEE Commun. Lett.* 28 (7), 1733–1737.
- Chen, X., Gan, S., Feng, C., Wang, X., Quek, T.Q., 2025. Take what you need: Flexible multi-task semantic communications with channel adaptation. *arXiv preprint arXiv:2502.08221*.
- Chen, W., Ju, C., Yuan, T., Zhan, Y., Zhang, M., Wang, D., 2025. Free space optical semantic communication for satellite remote sensing image transmission. *IEEE Trans. Commun.* 73 (10), 9168–9183.
- Chen, M., Liu, M., Wang, C., Song, X., Zhang, Z., Xie, Y., Wang, L., 2024a. Cross-modal graph semantic communication assisted by generative AI in the metaverse for 6G. *Research* 7, 0342.
- Chen, M., Liu, M., Zhang, Z., Xu, Z., Wang, L., 2024b. Task-oriented semantic communication with foundation models. *China Commun.* 21 (7), 65–77.
- Chen, S., Spyros-Papastavridis, E., Jin, Y., Deng, Y., 2024. Goal-oriented semantic communication for robotic arm reconstruction in digital twin. In: *Proc. IEEE Global Commun. Conf. (GLOBECOM)*. IEEE, pp. 4830–4835.
- Chen, S., Spyros-Papastavridis, E., Jin, Y., Deng, Y., 2025. Goal-oriented semantic communication for robot arm reconstruction in digital twin: Feature and temporal selections. *IEEE J. Sel. Areas Commun.* 43 (9), 3072–3087.
- Chen, Y., Yang, Q., Shi, Z., Chen, J., 2023. The model inversion eavesdropping attack in semantic communication systems. In: *Proc. IEEE Global Commun. Conf. GLOBECOM*.
- Choi, J., Park, J., Grassucci, E., Communiello, D., 2024. Semantic communication challenges: Understanding dos and avoiding don'ts. In: *Proc. 2024 IEEE 99th Veh. Technol. Conf. VTC2024-Spring*, IEEE, pp. 1–5.
- Courtade, T.A., Weissman, T., 2014. Multiterminal source coding under logarithmic loss. *IEEE Trans. Inform. Theory* 60 (1), 740–761.
- Dai, J., Wang, S., Tan, K., Si, Z., Qin, X., Niu, K., Zhang, P., 2022. Nonlinear transform source-channel coding for semantic communications. *IEEE J. Sel. Areas Commun.* 40 (8), 2300–2316.
- D'Alfonso, S., 2011. On quantifying semantic information. *Information* 2 (1), 61–101.
- Dang, S., Amin, O., Shihada, B., Alouini, M.-S., 2020. What should 6G be? *Nat. Electron.* 3 (1), 20–29.
- Défossez, A., Copet, J., Synnaeve, G., Adi, Y., 2023. High fidelity neural audio compression. *Trans. Mach. Learn. Res.*
- Ding, J., Jiang, P., Wen, C.-K., Li, X., Jin, S., 2026. Foundation model-aided channel-adaptive video semantic communication and prototype validation. *IEEE Trans. Wirel. Commun.* 25, 2892–2906.
- Dobrushin, R.L., Tsybakov, B.S., 1962. Information transmission with additional noise. *IEEE Trans. Inform. Theory* 8 (5), 293–304.
- Du, B., Du, H., Niyato, D., Li, R., 2025. Task-oriented semantic communication in large multimodal models-based vehicle networks. *IEEE Trans. Mob. Comput.* 24 (10), 9822–9836.
- Eldeeb, E., Shehab, M., Alves, H., 2024. A multi-task oriented semantic communication framework for autonomous vehicles. *IEEE Wirel. Commun. Lett.* 13 (12), 3469–3473.
- Erak, O., Alhussein, O., Tong, W., 2026. Contrastive learning and adversarial disentanglement for privacy-aware task-oriented semantic communication. *IEEE Internet Things J.* 13 (12), 25948–25965.
- Erdemir, E., Dragotti, P.L., Gündüz, D., 2022. Privacy-aware communication over a wiretap channel with generative networks. In: *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*. ICASSP, IEEE, pp. 2989–2993.
- Fan, D., Meng, R., Xu, X., Liu, Y., Nan, G., Feng, C., Han, S., Gao, S., Xu, B., Niyato, D., Quek, T.Q.S., Zhang, P., 2026. Generative diffusion models for wireless networks: Fundamental, architecture, and state-of-the-art. *IEEE Commun. Surv. Tutor.* 28, 5632–5677, [arXiv:2507.16733](https://arxiv.org/abs/2507.16733).
- Fang, Z., Sun, M., Wang, S., Xu, X., Gao, H., Huang, J., Han, S., Zhang, P., 2026. Cross-modal generative semantic communications powered by semantic knowledge base. *IEEE Trans. Netw. Sci. Eng.* 13, 5568–5585.
- Farsad, N., Rao, M., Goldsmith, A., 2018. Deep learning for joint source-channel coding of text. In: *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*. ICASSP, IEEE, pp. 2326–2330.
- Farshbafan, M.K., Saad, W., Debbah, M., 2023. Curriculum learning for goal-oriented semantic communications with a common language. *IEEE Trans. Commun.* 71 (3), 1430–1446.
- Fayaz, N., Mallik, R.K., Lall, B., 2026. Semantic noise decomposition and saliency-weighted mutual information in task-oriented communication. *IEEE Commun. Lett.* 30, 2148–2152.
- Feng, Y., Shen, H., Shan, Z., Yang, Q., Shi, X., 2025. Semantic communication for edge intelligence enabled autonomous driving system. *IEEE Netw.* 39 (2), 149–157, [arXiv:2406.10606](https://arxiv.org/abs/2406.10606).
- Floridi, L., 2004. Outline of a theory of strongly semantic information. *Minds Mach.* 14, 197–221.
- Flowers, B., Buehrer, R.M., Headley, W.C., 2020. Evaluating adversarial evasion attacks in the context of wireless communications. *IEEE Trans. Inf. Forensics Secur.* 15, 1102–1113.
- Fredrikson, M., Jha, S., Ristenpart, T., 2015. Model inversion attacks that exploit confidence information and basic countermeasures. In: *Proc. 22nd ACM SIGSAC Conf. Comput. Commun. Secur.*. pp. 1322–1333.
- Getu, T.M., Kaddoum, G., Bennis, M., 2023. Making sense of meaning: A survey on metrics for semantic and goal-oriented communication. *IEEE Access* 11, 45456–45492.
- Getu, T.M., Kaddoum, G., Bennis, M., 2024. A survey on goal-oriented semantic communication: Techniques, challenges, and future directions. *IEEE Access* 12, 51223–51274.
- Gholami, N., Moghim, N., Ghahfarokhi, B.S., Salavati, P., Thomas, C.K., Shetty, S., Rahmati, T., 2025. Flexible semantic-aware resource allocation: Serving more users through similarity range constraints. *arXiv preprint arXiv:2504.20939*.
- Goldfeld, Z., Polyanskiy, Y., 2020. The information bottleneck problem and its applications in machine learning. *IEEE J. Sel. Areas Inf. Theory* 1 (1), 19–38.
- Goldreich, O., Juba, B., Sudan, M., 2012. A theory of goal-oriented communication. *J. ACM* 59 (2), 1–65.
- Grassucci, E., Venturini, F., Communiello, D., 2023. Generative semantic communication: Diffusion models beyond bit recovery. In: *Proc. IEEE Global Commun. Conf. GLOBECOM*, IEEE, pp. 3603–3608.
- Güler, B., Yener, A., Swami, A., 2018. The semantic communication game. *IEEE Trans. Cogn. Commun. Netw.* 4 (4), 787–802.
- Gündüz, D., Qin, Z., Aguerri, I.E., Dhillon, H.S., Yang, Z., Yener, A., Wong, K.K., Chae, C.-B., 2023. Beyond transmitting bits: Context, semantics, and task-oriented communications. *IEEE J. Sel. Areas Commun.* 41 (1), 5–41.

- Guo, J., Chen, H., Song, B., Chi, Y., Yuen, C., Yu, F.R., Li, G.Y., Niyato, D., 2024. Distributed task-oriented communication networks with multimodal semantic relay and edge intelligence. *IEEE Commun. Mag.* 62 (6), 82–89.
- Guo, S., Wang, Y., Zhang, N., Su, Z., Luan, T.H., Tian, Z., Shen, X., 2024a. A survey on semantic communication networks: Architecture, security, and privacy. *IEEE Commun. Surv. Tutor.* 27 (5), 2860–2894.
- Guo, J., Yang, C., 2022. Learning precoding for semantic communications. In: *Proc. 2022 IEEE Int. Conf. Commun. Workshops (ICC Workshops)*. IEEE, pp. 163–168.
- Guo, S., Zhang, A., Wang, Y., Feng, C., Quek, T.Q., 2025. Semantic-enabled 6G communication: A task-oriented and privacy-preserving perspective. *IEEE Netw.* 40 (2), 113–121.
- Guo, S., Zhou, J., Li, Z., Zhang, P., Zhang, S., 2024b. Reconfigurable intelligent surfaces-assisted task-oriented communications for AI-driven vertical applications. In: *Proc. 2024 22nd Int. Symp. Model. Optim. Mobile, Ad Hoc, Wireless Netw.. WiOpt, IEEE*, pp. 99–104.
- Han, T., Yang, Q., Shi, Z., He, S., Zhang, Z., 2022. Semantic-preserved communication system for highly efficient speech transmission. *IEEE J. Sel. Areas Commun.* 41 (1), 245–259.
- Han, W., Zhang, A., Feng, C., Chen, R., Zhang, S., Guo, S., 2025. Few-shot adaptive learning for robust task-oriented semantic communication. In: *Proc. AAAI 2025 Workshop Artif. Intell. Wireless Commun. Netw.. AI4WCN*.
- Hassan, S.S., Nguyen, L.X., Tun, Y.K., Han, Z., Hong, C.S., 2024. Semantic enabled 6G LEO satellite communication for earth observation: A resource-constrained network optimization. In: *Proc. IEEE Global Commun. Conf. (GLOBECOM)*. IEEE, pp. 4472–4478.
- He, J., Deng, Y., Wu, S., Cao, B., Meng, S., Zhang, W., Zhang, Q., 2026a. Task-adaptive semantic communication with feedback: A conditional rate-distortion approach. *IEEE Trans. Cogn. Commun. Netw.* 12, 7157–7171.
- He, G., Feng, M., Zhang, Y., Liu, G., Dai, Y., Jiang, T., 2023. Deep reinforcement learning based task-oriented communication in multi-agent systems. *IEEE Wirel. Commun.* 30 (3), 112–119.
- He, J., Luo, X., Kang, J., Xiong, Z., Li, X., Chen, C., Niyato, D., 2026b. Mixture of experts-based semantic communication for secure Artificial Intelligence of Things. *IEEE Trans. Veh. Technol.* 75 (2), 3324–3329.
- Hoang, D.T.M., Le, L.B., 2025. Backdoor attacks on large language model based semantic communication systems. In: *Proc. IEEE Int. Conf. Commun. Workshops. ICC Workshops*, pp. 518–523.
- Hu, X., Tian, Y., Li, Q., Kho, Y.H., Wang, X., Xiao, B., Yang, Z., Li, W., 2024. A novel RIS-aided optimization strategy for semantic communication system. *IEEE Wirel. Commun. Lett.* 13 (6), 1655–1659.
- Hu, L., Wang, H., Xiang, T., Chen, Z., He, J., Guo, S., 2026. Vector-guided embezzler: A reconstruction attack against retrieval-augmented generation in agentic AI communication. *IEEE Trans. Cogn. Commun. Netw.* 12, 8303–8316.
- Hu, Q., Zhang, G., Qin, Z., Cai, Y., Yu, G., Li, G.Y., 2023. Robust semantic communications with masked VQ-VAE enabled codebook. *IEEE Trans. Wirel. Commun.* 22 (12), 8707–8722.
- Huang, G., An, J., Yang, Z., Gan, L., Bennis, M., Debbah, M., 2024. Stacked intelligent metasurfaces for task-oriented semantic communications. *IEEE Wirel. Commun. Lett.* 14 (2), 310–314.
- Huang, Y., Bai, B., Zhu, Y., Qiao, X., Su, X., Yang, L., Zhang, P., 2023. ISCom: Interest-aware semantic communication scheme for point cloud video streaming on metaverse XR devices. *IEEE J. Sel. Areas Commun.* 42 (4), 1003–1021.
- Huang, A.Y., Chen, Y., Huang, D., Zhao, M., 2023. Semantic privacy-preserving for video surveillance services on the edge. In: *Proc. 8th ACM/IEEE Symp. Edge Comput.*, pp. 300–305.
- Huang, D., Gao, F., Tao, X., Du, Q., Lu, J., 2022. Toward semantic communications: Deep learning-based image semantic coding. *IEEE J. Sel. Areas Commun.* 41 (1), 55–71.
- Islam, A., Chang, K., 2026. Beyond bit fidelity: Dual-mode domain-aware knowledge-enhanced semantic text communication framework. *IEEE Trans. Netw. Sci. Eng.* 13, 7964–7986.
- Jiang, P., Wen, C.-K., Jin, S., Li, G.Y., 2024. RIS-enhanced semantic communications adaptive to user requirements. *IEEE Trans. Commun.* 72 (7), 4134–4148.
- Juba, B., Sudan, M., 2008. Universal semantic communication I. In: *Proc. 40th Annu. ACM Symp. Theory Comput.*, pp. 123–132.
- Juba, B., Sudan, M., 2009. Universal Semantic Communication II: A Theory of Goal-Oriented Communication. Tech. Rep. TR08-095, MIT CSAIL, URL <https://people.csail.mit.edu/madhu/papers/2008/usc2.pdf>.
- Juba, B., Vempala, S., 2011. Semantic communication for simple goals is equivalent to on-line learning. In: *Algorithmic Learning Theory: 22nd International Conference, ALT 2011, Espoo, Finland, October 5-7, 2011. Proceedings 22*. Springer, pp. 277–291.
- Kalfa, M., Gok, M., Atalik, A., Tegin, B., Duman, T.M., Arıkan, O., 2021. Towards goal-oriented semantic signal processing: Applications and future challenges. *Digit. Signal Process.* 119, 103134.
- Kang, J., Du, H., Li, Z., Xiong, Z., Ma, S., Niyato, D., Li, Y., 2022. Personalized saliency in task-oriented semantic communications: Image transmission and performance analysis. *IEEE J. Sel. Areas Commun.* 41 (1), 186–201.
- Kim, B., Sagduyu, Y.E., Davaslioglu, K., Erpek, T., Ulukus, S., 2022. Channel-aware adversarial attacks against deep learning-based wireless signal classifiers. *IEEE Trans. Wirel. Commun.* 21 (6), 3868–3880.
- Kostina, V., Verdú, S., 2013. Lossy joint source-channel coding in the finite blocklength regime. *IEEE Trans. Inform. Theory* 59 (5), 2545–2575.
- Kountouris, M., Pappas, N., 2021. Semantics-empowered communication for networked intelligent systems. *IEEE Commun. Mag.* 59 (6), 96–102.
- Kreuzer, T., Papapetrou, P., Zdravkovic, J., 2024. Artificial intelligence in digital twins—A systematic literature review. *Data Knowl. Eng.* 151, 102304.
- Kurka, D.B., Gündüz, D., 2020. DeepJSCC-f: Deep joint source-channel coding of images with feedback. *IEEE J. Sel. Areas Inf. Theory* 1 (1), 178–193.
- Kurka, D.B., Gündüz, D., 2021. Bandwidth-agile image transmission with deep joint source-channel coding. *IEEE Trans. Wirel. Commun.* 20 (12), 8081–8095.
- Li, P., Aijaz, A., 2025. Task-oriented connectivity for networked robotics with generative AI and semantic communications. In: *Proc. IEEE INFOCOM 2025 - IEEE Conf. Comput. Commun. Workshops. INFOCOM WKSHPS*, IEEE, pp. 1–6, arXiv: 2503.06771.
- Li, B., Cai, H., Liu, L., Fei, Z., 2025. Delay-aware digital twin synchronization in mobile edge networks with semantic communications. *IEEE Trans. Veh. Technol.* 74 (7), 10974–10983.
- Li, N., Deng, Y., Niyato, D., 2026. Goal-oriented semantic communication for wireless video transmission via generative AI. *IEEE Trans. Wirel. Commun.* 25, 10841–10854.
- Li, Y., Jiang, S., Liu, B., Wang, B., Wang, L., Rao, M., 2025a. Semantic communication-enabled cloud-edge-end-collaborative metaverse services architecture. In: *Proc. IEEE Int. Workshop Signal Process. Adv. Wireless Commun. (SPAWC)*. IEEE, pp. 1–5.
- Li, Y., Jin, Z., Song, T., Song, X., Hu, J., 2025b. Robust task-oriented communication with semantic-aware masking and discrete codebook. In: *Proc. 2025 IEEE Wireless Commun. Netw. Conf.. WCNC*, IEEE, pp. 1–6.
- Li, H., Shao, J., He, H., Song, S., Zhang, J., Letaief, K.B., 2025. Tackling distribution shifts in task-oriented communication with information bottleneck. *IEEE J. Sel. Areas Commun.* 43 (7), 2667–2683.
- Liang, C., Li, D., 2026. Generative AI-enabled semantic communication: State-of-the-art, applications, and the way ahead. *IEEE Commun. Surv. Tutor.* 28, 3976–4015.
- Liu, M., An, J., Huang, C., Yuen, C., 2025. Over-the-air ode-inspired neural network for dual task-oriented semantic communications. *IEEE Trans. Cogn. Commun. Netw.* 12, 805–819.
- Liu, G., Du, H., Niyato, D., Kang, J., Xiong, Z., Soong, B.H., 2023. Vision-based semantic communications for metaverse services: A contest theoretic approach. In: *Proc. IEEE Global Commun. Conf. (GLOBECOM)*. IEEE, pp. 2426–2432.
- Liu, C., Guo, C., Yang, Y., Chen, M., Quek, T.Q., 2025. Lightweight task-oriented semantic communication empowered by large-scale AI models. *IEEE Trans. Veh. Technol.* 74 (9), 14823–14827.
- Liu, C., Guo, C., Yang, Y., Jiang, N., 2023. Adaptable semantic compression and resource allocation for task-oriented communications. *IEEE Trans. Cogn. Commun. Netw.* 10 (3), 769–782.
- Liu, C., Guo, C., Yang, Y., Ni, W., Zhou, Y., Li, L., Quek, T.Q., 2024. Explainable semantic communication for text tasks. *IEEE Internet Things J.* 11 (24), 39820–39833.
- Liu, S., Li, N., Deng, Y., Quek, T.Q., 2025. Goal-oriented semantic communication for wireless visual question answering. *IEEE J. Sel. Areas Commun.* 43 (12), 4247–4261, arXiv:2411.02452.
- Liu, J., Shao, S., Zhang, W., Poor, H.V., 2022. An indirect rate-distortion characterization for semantic sources: General model and the case of gaussian observation. *IEEE Trans. Commun.* 70 (9), 5946–5959.
- Liu, Y., Wang, X., Ning, Z., Zhou, M., Guo, L., Jedari, B., 2024. A survey on semantic communications: technologies, solutions, applications and challenges. *Digit. Commun. Netw.* 10 (3), 528–545.
- Lu, Z., Li, R., Lu, K., Chen, X., Hossain, E., Zhao, Z., Zhang, H., 2023. Semantics-empowered communications: A tutorial-cum-survey. *IEEE Commun. Surv. Tutor.* 26 (1), 41–79.
- Lu, K., Zhou, Q., Li, R., Zhao, Z., Chen, X., Wu, J., Zhang, H., 2022. Rethinking modern communication from semantic coding to semantic communication. *IEEE Wirel. Commun.* 30 (1), 158–164.
- Lv, J., Tong, H., Pan, Q., Zhang, Z., He, X., Luo, T., Yin, C., 2024. Importance-aware image segmentation-based semantic communication for autonomous driving. *arXiv preprint arXiv:2401.10153*.
- Lv, Z., Xie, S., 2024. Artificial intelligence in the digital twins: State of the art, challenges, and future research topics. *Digit. Twin* 1 (1), 12.
- Lyu, Z., Cao, X., Song, X., Li, Y., Wang, J., Wang, S., Cui, Y., Yuan, W., Yu, X., Zhu, G., Liu, H., Xu, J., Ng, D.W.K., Cui, S., 2026. Sense smarter, think better: A survey on edge perception for next-generation networks. *arXiv preprint arXiv:2605.18457*.
- Lyu, Z., Zhu, G., Xu, J., Ai, B., Cui, S., 2024. Semantic communications for image recovery and classification via deep joint source and channel coding. *IEEE Trans. Wirel. Commun.* 23 (8), 8388–8404.
- Ma, S., Flanigan, K.A., Bergés, M., 2024. State-of-the-art review: The use of digital twins to support artificial intelligence-guided predictive maintenance. *arXiv preprint arXiv:2406.13117*.
- Ma, J., Li, Q., Liu, R., Pandharipande, A., Ge, X., 2024. Enhanced semantic information transfer on RIS-assisted communication systems. *IEEE Wirel. Commun. Lett.* 13 (8), 2225–2229.
- Ma, S., Qiao, W., Wu, Y., Li, H., Shi, G., Gao, D., Shi, Y., Li, S., Al-Dhahir, N., 2023. Task-oriented explainable semantic communications. *IEEE Trans. Wirel. Commun.* 22 (12), 9248–9262.

- Ma, Y., Yin, H., Qiao, L., Sun, S., Gao, Z., Xu, Y., Zhang, W., 2026. LGVSC: A large-model-driven generative video semantic communication framework. *IEEE Trans. Veh. Technol.* 1–6.
- Maatouk, A., Assaad, M., Ephremides, A., 2022. The age of incorrect information: An enabler of semantics-empowered communication. *IEEE Trans. Wirel. Commun.* 22 (4), 2621–2635.
- Mao, S., Yuen, C., Liu, L., Xiao, M., Yu, S., Zhang, N., 2026. RIS-enhanced semantic-aware sensing, communication, computation, and control for Internet of Things. *IEEE Trans. Wirel. Commun.* 25, 2231–2246.
- Meng, R., Gao, S., Fan, D., Gao, H., Wang, Y., Xu, X., Wang, B., Lv, S., Zhang, Z., Sun, M., et al., 2025. A survey of secure semantic communications. *J. Netw. Comput. Appl.* 104181.
- Meng, R., Huang, Z., Yan, J., Sun, M., Liu, Y., Feng, C., Xu, X., Zhang, Z., Gao, S., Zhang, P., Quek, T.Q.S., 2026. Semantic radio access networks: Architecture, state-of-the-art, and future directions. *IEEE Trans. Cogn. Commun. Netw.* 12, 7076–7097, arXiv:2512.20917.
- Monga, V., Li, Y., Eldar, Y.C., 2021. Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing. *IEEE Signal Process. Mag.* 38 (2), 18–44.
- Moqurrah, S.A., Anjum, A., Tariq, N., Srivastava, G., 2022. InstantAnonymity: A lightweight semantic privacy guarantee for 5G-enabled IIoT. *IEEE Trans. Ind. Inf.* 19 (1), 951–959.
- Morris, C.W., 1938. Foundations of the theory of signs. In: *International Encyclopedia of Unified Science*. Chicago University Press, pp. 1–59.
- Nandakishore, K., Rithvik, R.V.S., Khan, M.Z.A., 2024. Efficient image transmission using semantic communication for static environment surveillance. In: *Proc. 2024 IEEE Int. Conf. Adv. Netw. Telecommun. Syst. ANTS, IEEE*, pp. 1–6.
- Nguyen, T.T., Le, T.-D., Ha, V.N., Chou, H.-f., Eappen, G., Tran, D.-D., Nguyen-Kha, H., Thiruvassagam, P., Garces-Socarras, L.M., Gonzalez-Rios, J.L., et al., 2025. A semantic-loss function modeling framework with task-oriented machine learning perspectives. In: *Proc. 2025 IEEE Int. Conf. Mach. Learn. Commun. Netw. ICMLCN, IEEE*, pp. 1–6, arXiv:2503.09903.
- Nguyen, L.X., Raha, A.D., Aung, P.S., Niyato, D., Han, Z., Hong, C.S., 2026. A contemporary survey on semantic communications: Theory of mind, generative AI, and deep joint source-channel coding. *IEEE Commun. Surv. Tutor.* 28, 2377–2417, arXiv:2502.16468.
- Niiniluoto, I., 2012. Truthlikeness, vol. 185, Springer Science & Business Media.
- Niu, K., Dai, J., Yao, S., Wang, S., Si, Z., Qin, X., Zhang, P., 2022. A paradigm shift toward semantic communications. *IEEE Commun. Mag.* 60 (11), 113–119.
- Okegbile, S.D., Gao, H., Cai, J., 2026. A novel secure split federated semantic learning framework and its optimization for digital twin network evolution. *IEEE Trans. Mob. Comput.* 25 (1), 1302–1319.
- Park, J., Yoon, S.W., 2025. Transmit what you need: task-adaptive semantic communications for visual information. *IEEE J. Sel. Areas Commun.* 43 (12), 4182–4197, arXiv:2412.13646.
- Peng, J., Xing, H., Li, Y., Feng, L., Xu, L., Lei, X., 2024a. Task-oriented multi-user semantic communication with lightweight semantic encoder and fast training for resource-constrained terminal devices. *IEEE Wirel. Commun. Lett.* 13 (9), 2427–2431.
- Peng, J., Xing, H., Xu, L., Luo, S., Dai, P., Feng, L., Song, J., Zhao, B., Xiao, Z., 2024b. Adversarial reinforcement learning based data poisoning attacks defense for task-oriented multi-user semantic communication. *IEEE Trans. Mob. Comput.* 23 (12), 14834–14851.
- Popovski, P., Chiariotti, F., Croisfelt, V., Kalør, A.E., Leyva-Mayorga, I., Marchegiani, L., Pandey, S.R., Soret, B., 2021. Internet of things (IIoT) connectivity in 6G: An interplay of time, space, intelligence, and value. *arXiv preprint arXiv:2111.05811*.
- Qiao, L., Mashhadi, M.B., Gao, Z., Foh, C.H., Xiao, P., Bennis, M., 2024. Latency-aware generative semantic communications with pre-trained diffusion models. *IEEE Wirel. Commun. Lett.* 13 (10), 2652–2656.
- Razlighi, A.H., Bockelmann, C., Dekorsy, A., 2024. Semantic communication for cooperative multi-task processing over wireless networks. *IEEE Wirel. Commun. Lett.* 13 (10), 2867–2871.
- Razlighi, A.H., Tillmann, M.H., Beck, E., Bockelmann, C., Dekorsy, A., 2025. Cooperative and collaborative multi-task semantic communication for distributed sources. In: *Proc. ICC 2025-IEEE Int. Conf. Commun.* IEEE, pp. 3966–3971.
- Ribeiro, F., Florêncio, D., Zhang, C., Seltzer, M., 2011. Crowdmos: An approach for crowdsourcing mean opinion score studies. In: *Proc. 2011 IEEE Int. Conf. Acoust., Speech Signal Process. ICASSP, IEEE*, pp. 2416–2419.
- Ribouh, S., Saleem, O., 2025. Large language model-based semantic communication system for image transmission. *arXiv preprint arXiv:2501.12988*.
- Rix, A.W., Beerends, J.G., Hollier, M.P., Hekstra, A.P., 2001. Perceptual evaluation of speech quality (PESQ)-a new method for speech quality assessment of telephone networks and codecs. *ICASSP, In: Proc. 2001 IEEE Int. Conf. Acoust., Speech, Signal Process.*, vol. 2, IEEE, pp. 749–752.
- Saad, W., Bennis, M., Chen, M., 2019. A vision of 6G wireless systems: Applications, trends, technologies, and open research problems. *IEEE Netw.* 34 (3), 134–142.
- Sadeghi, M., Larsson, E.G., 2019. Adversarial attacks on deep-learning based radio signal classification. *IEEE Wirel. Commun. Lett.* 8 (1), 213–216.
- Sagduyu, Y.E., Erpek, T., Yener, A., Ulukus, S., 2024. Will 6G be semantic communications? opportunities and challenges from task oriented and secure communications to integrated sensing. *IEEE Netw.* 38 (6), 72–80.
- Sagduyu, Y.E., Ulukus, S., Yener, A., 2023. Task-oriented communications for NextG: End-to-end deep learning and AI security aspects. *IEEE Wirel. Commun.* 30 (3), 52–60.
- Salehi, S., Erol-Kantarci, M., Niyato, D., 2025. LLM-enabled data transmission in end-to-end semantic communication. In: *Proc. 2025 IEEE Symp. Comput. Commun. ISCC*, pp. 1–6.
- Shah, P., Garg, A., Gajjar, V., 2021. Per-vis: Person retrieval in video surveillance using semantic description. In: *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis.* pp. 41–50.
- Shannon, C.E., 1948. A mathematical theory of communication. *Bell Syst. Tech. J.* 27 (3), 379–423.
- Shannon, C.E., Weaver, W., 1998. *The Mathematical Theory of Communication*. University of Illinois Press.
- Shao, J., Mao, Y., Zhang, J., 2021. Learning task-oriented communication for edge inference: An information bottleneck approach. *IEEE J. Sel. Areas Commun.* 40 (1), 197–211.
- Shao, J., Zhang, X., Zhang, J., 2023. Task-oriented communication for edge video analytics. *IEEE Trans. Wirel. Commun.* 23 (5), 4141–4154.
- Shen, L.-H., Feng, K.-T., Hanzo, L., 2023. Five facets of 6G: Research challenges and opportunities. *ACM Comput. Surv.* 55 (11), 1–39.
- Shi, J., Chan, T.-T., Pan, H., Lok, T.-M., 2023. Reconfigurable intelligent surface assisted semantic communication systems. In: *Proc. 2023 IEEE Int. Conf. Signal Process., Commun. Comput. ICSPCC, IEEE*, pp. 1–6.
- Shokri, R., Strinati, M., Song, C., Shmatikov, V., 2017. Membership inference attacks against machine learning models. In: *Proc. 2017 IEEE Symp. Secur. Priv. SP, IEEE*, pp. 3–18.
- Shumailov, I., Zhao, Y., Bates, D., Papernot, N., Mullins, R., Anderson, R., 2021. Sponge examples: Energy-latency attacks on neural networks. In: *Proc. 2021 IEEE Eur. Symp. Secur. Priv. (EuroS&P)*, IEEE, pp. 212–231.
- Song, X., Yuan, L., Qu, Z., Zhou, F., Wu, Q., Quek, T.Q., Hu, R.Q., 2024. Knowledge graph driven UAV cognitive semantic communication systems for efficient object detection. In: *Proc. ICC 2024 - IEEE Int. Conf. Commun.* IEEE, pp. 1685–1690, arXiv:2401.13995.
- Song, X., Zhou, F., Ding, R., Qu, Z., Li, Y., Wu, Q., Al-Dhahir, N., 2025. UAV cognitive semantic communications enabled by knowledge graph for robust object detection. *IEEE Trans. Commun.* 73 (8), 6052–6067.
- Stavrou, P.A., Kountouris, M., 2023. The role of fidelity in goal-oriented semantic communication: A rate distortion approach. *IEEE Trans. Commun.* 71 (7), 3881–3896.
- Strinati, E.C., Barbarossa, S., 2021. 6G networks: Beyond Shannon towards semantic and goal-oriented communications. *Comput. Netw.* 190, 107930.
- Strinati, E.C., Di Lorenzo, P., Sciancalepore, V., Aijaz, A., Kountouris, M., Gündüz, D., Popovski, P., Sana, M., Stavrou, P.A., Soret, B., et al., 2024. Goal-oriented and semantic communication in 6G AI-native networks: The 6G-GOALS approach. In: *Proc. 2024 Joint Eur. Conf. Netw. Commun. & 6G Summit (EuCNC/6G Summit)*. IEEE, pp. 1–6.
- Sun, Y., Liu, Y., Guo, S., Qiu, X., Chen, J., Hao, J., Niyato, D., 2026. Edge large AI model agent-empowered cognitive multimodal semantic communication. *IEEE Trans. Mob. Comput.* 25 (1).
- Sun, S., Qin, Z., Xie, H., Tao, X., 2023. Task-oriented explainable semantic communications based on structured scene graphs. In: *Proc. IEEE Global Commun. Conf. (GLOBECOM)*. IEEE, pp. 3222–3227.
- Sun, M., Wu, D., Zhang, P., Wang, R., 2026. Disentangled information bottleneck guided multidevice cooperative task-oriented semantic communication. *IEEE Trans. Commun.* 74, 6259–6274.
- Suresh, A., Beck, E., Dekorsy, A., Rückert, P., Tracht, K., 2024. Human-integrated multi-agent exploration using semantic communication and extended reality simulation. In: *Proc. 2024 10th Int. Conf. Autom. Robot. Appl. (ICARA)*, Athens, Greece. pp. 22–24.
- Tang, J., Nie, J., Bai, J., Xu, J., Li, S., Zhang, Y., Yuan, Y., 2024. UAV-assisted digital twin synchronization with tiny machine learning-based semantic communications. *IEEE Internet Things J.* 11 (17), 28437–28451.
- Thomas, C.K., Saad, W., 2022. Neuro-symbolic artificial intelligence (AI) for intent-based semantic communication. In: *Proc. IEEE Global Commun. Conf. (GLOBECOM)*. IEEE, pp. 2698–2703.
- Tishby, N., Pereira, F.C., Bialek, W., 2000. The information bottleneck method. *arXiv Preprint Physics/0004057*.
- Tramèr, F., Zhang, F., Juels, A., Reiter, M.K., Ristenpart, T., 2016. Stealing machine learning models via prediction APIs. In: *Proc. 25th USENIX Secur. Symp. USENIX Security* 16, pp. 601–618.
- Tung, T.-Y., Kurka, D.B., Jankowski, M., Gündüz, D., 2022. DeepJSCC-Q: Channel input constrained deep joint source-channel coding. In: *Proc. IEEE Int. Conf. Commun. ICC, IEEE*, pp. 3880–3885.
- Vaidanis, S., Stavrou, P.A., Kountouris, M., 2025. Optimization for semantic-aware resource allocation under CPT-based utilities. In: *Proc. 2025 IEEE 26th Int. Workshop Signal Process. Artif. Intell. Wireless Commun. SPAWC, IEEE*, pp. 1–5, arXiv:2506.04952.
- Wang, Z., Deng, Y., Aghvami, A.H., 2023. Task-oriented and semantics-aware communications for augmented reality. In: *Proc. IEEE Global Commun. Conf. (GLOBECOM)*. IEEE, pp. 2215–2220.

- Wang, Z., Deng, Y., Aghvami, A.H., 2024. Goal-oriented semantic communications for avatar-centric augmented reality. *IEEE Trans. Commun.* 72 (12), 7982–7995.
- Wang, Y., Guo, S., Deng, Y., Zhang, H., Fang, Y., 2024a. Privacy-preserving task-oriented semantic communications against model inversion attacks. *IEEE Trans. Wirel. Commun.* 23 (8), 10150–10165.
- Wang, Y., Han, H., Feng, Y., Zheng, J., Zhang, B., 2025. Semantic communication empowered 6G networks: Techniques, applications, and challenges. *IEEE Access* 13, 28293–28314.
- Wang, Y., Han, S., Xu, X., Liang, H., Meng, R., Dong, C., Zhang, P., 2024b. Feature importance-aware task-oriented semantic transmission and optimization. *IEEE Trans. Cogn. Commun. Netw.* 10 (4), 1175–1189.
- Wang, W., Wang, C., Zhang, Z., Xu, W., Zhang, P., Jiang, F., Xu, J., Cai, Y., Pang, M., Xu, L., Quek, T.Q.S., 2026. Resilient image semantic communication based on rate-optimized information bottleneck theory. *IEEE Trans. Netw. Sci. Eng.* 13, 3127–3143.
- Weaver, W., 1953. Recent contributions to the mathematical theory of communication. *ETC: A Rev. Gen. Semant.* 261–281.
- Wei, Z., Jiang, Y., Huang, J., Zhong, F., Gyawali, S., 2025. Detecting backdoor attacks via similarity in semantic communication systems. In: *Lecture Notes in Networks and Systems*. Springer, pp. 315–329, arXiv:2502.03721.
- Wei, K., Xie, R., Xu, W., Lu, Z., Xiao, H., 2025. Robust semantic communication via adversarial training. *IEEE Trans. Veh. Technol.* 74 (12), 19849–19853.
- Weng, Z., Qin, Z., 2021. Semantic communication systems for speech transmission. *IEEE J. Sel. Areas Commun.* 39 (8), 2434–2444.
- Wheeler, D., Natarajan, B., 2023. Engineering semantic communication: A survey. *IEEE Access* 11, 13965–13995.
- Willner, A., Diedrich, C., Younes, R.B., Hohmann, S., Kraft, A., 2017. Semantic communication between components for smart factories based on oneM2M. In: *Proc. 2017 22nd IEEE Int. Conf. Emerg. Technol. Factory Autom. ETFA, IEEE*, pp. 1–8.
- Witsenhausen, H.S., 1980. Indirect rate distortion problems. *IEEE Trans. Inform. Theory* 26 (5), 518–521.
- Wu, T., Chen, Z., He, D., Qian, L., Xu, Y., Tao, M., Zhang, W., 2024. CDDM: Channel denoising diffusion models for wireless semantic communications. *IEEE Trans. Wirel. Commun.* 23 (9), 11168–11183.
- Wu, W., Yang, Y., Deng, Y., Aghvami, A.H., 2024. Goal-oriented semantic communications for robotic waypoint transmission: The value and age of information approach. *IEEE Trans. Wirel. Commun.* 23 (12), 18903–18915.
- Wu, D., Zhao, Q., Fan, J., Qi, J., Zheng, P., Hu, J., 2025. H2R bridge: Transferring vision-language models to few-shot intention meta-perception in human robot collaboration. *J. Manuf. Syst.* 80, 524–535.
- Xia, L., Sun, Y., Niyato, D., Li, X., Imran, M.A., 2023. Joint user association and bandwidth allocation in semantic communication networks. *IEEE Trans. Veh. Technol.* 73 (2), 2699–2711.
- Xiang, Z., Li, Y., Yu, F., Wang, Y., Li, Y., Rao, Z., 2025. Task-oriented semantic communication via federated learning. In: *Proc. 2025 IEEE Int. Perform., Comput., Commun. Conf. IPCCC, IEEE*, pp. 1–7.
- Xiao, Y., Zhang, X., Li, Y., Shi, G., Başar, T., 2022. Rate-distortion theory for strategic semantic communication. In: *Proc. 2022 IEEE Inf. Theory Workshop. ITW, IEEE*, pp. 279–284.
- Xie, H., Qin, Z., 2020. A lite distributed semantic communication system for Internet of Things. *IEEE J. Sel. Areas Commun.* 39 (1), 142–153.
- Xie, H., Qin, Z., Li, G.Y., 2021a. Task-oriented multi-user semantic communications for VQA. *IEEE Wirel. Commun. Lett.* 11 (3), 553–557.
- Xie, H., Qin, Z., Li, G.Y., Juang, B.-H., 2021b. Deep learning enabled semantic communication systems. *IEEE Trans. Signal Process.* 69, 2663–2675.
- Xie, H., Qin, Z., Tao, X., Han, Z., 2024. Toward intelligent communications: Large model empowered semantic communications. *IEEE Commun. Mag.* 63 (1), 69–75.
- Xie, H., Qin, Z., Tao, X., Letaief, K.B., 2022. Task-oriented multi-user semantic communications. *IEEE J. Sel. Areas Commun.* 40 (9), 2584–2597.
- Xin, G., Fan, P., Letaief, K.B., 2024. Semantic communication: A survey of its theoretical development. *Entropy* 26 (2), 102.
- Xu, Z., Jiao, J., Yuan, G., Jiang, S., Wang, W., Wang, Y., Zhang, Q., 2025. Task-oriented semantic coding and utility optimal transmission in satellite-integrated internet. *IEEE Trans. Geosci. Remote Sens.* 63, 1–16.
- Xu, C., Ma, Y., Tafazolli, R., Wang, J., 2025. Data-importance-aware power allocation for adaptive real-time communication in computer vision applications. *IEEE J. Sel. Areas Commun.* 43 (12), 4015–4026, arXiv:2504.08922.
- Xu, J., Tung, T.-Y., Ai, B., Chen, W., Sun, Y., Gündüz, D., 2023. Deep joint source-channel coding for semantic communications. *IEEE Commun. Mag.* 61 (11), 42–48.
- Xu, J., Wang, X., Chen, L., Lin, L., Hu, J., Khawaja, S.A., Dev, K., 2025. Fortifying secure semantic communication: A next-generation defense framework against model inversion and adversarial threats. *IEEE Wirel. Commun.* 32 (5), 64–71.
- Xu, J., Yao, H., Zhang, R., Mai, T., Huang, S., Guo, S., 2024. Federated learning powered semantic communication for UAV swarm cooperation. *IEEE Wirel. Commun.* 31 (4), 140–146.
- Xu, Y., Zhou, H., Deng, Y., 2023. Task-oriented semantics-aware communication for wireless UAV control and command transmission. *IEEE Commun. Lett.* 27 (8), 2232–2236.
- Yan, L., Qin, Z., Zhang, R., Li, Y., Li, G.Y., 2022a. QoE-aware resource allocation for semantic communication networks. In: *Proc. IEEE Global Commun. Conf. (GLOBECOM)*. IEEE, pp. 3272–3277.
- Yan, L., Qin, Z., Zhang, R., Li, Y., Li, G.Y., 2022b. Resource allocation for text semantic communications. *IEEE Wirel. Commun. Lett.* 11 (7), 1394–1398.
- Yang, W., Du, H., Liew, Z.Q., Lim, W.Y.B., Xiong, Z., Niyato, D., Chi, X., Shen, X., Miao, C., 2022. Semantic communications for future internet: Fundamentals, applications, and challenges. *IEEE Commun. Surv. Tutor.* 25 (1), 213–250.
- Yang, M., Gao, D., Li, J., Xu, W., Shi, G., 2024. Superimposed semantic communication for IoT-based real-time ECG monitoring. *IEEE J. Biomed. Health Inf.* 28 (7), 3819–3830.
- Yang, Y., Guo, C., Liu, F., Liu, C., Sun, L., Sun, Q., Chen, J., 2021. Semantic communications with AI tasks. *arXiv preprint arXiv:2109.14170*.
- Yang, X., Lai, Y., Li, G., Wu, J., Zhou, K., Chen, M., 2026a. SemanAegis: Toward credential-aware semantic communication against knowledge leakage threats. *IEEE Trans. Mob. Comput.* 25 (4), 5032–5049.
- Yang, X., Lai, Y., Li, G., Wu, J., Zhou, K., Li, J., Chen, M., 2026b. Toward polymorphic backdoor against semantic communication via intensity-based poisoning. *IEEE Trans. Inf. Forensics Secur.* 21, 4505–4520.
- Yang, M., Liu, B., Wang, B., Kim, H.-S., 2025. Diffusion-aided joint source channel coding for high realism wireless image transmission. *IEEE Trans. Mach. Learn. Commun. Netw.* 3, 1227–1243.
- Yang, W., Xiong, Z., Mao, S., Quek, T.Q., Zhang, P., Debbah, M., Tafazolli, R., 2025. Rethinking generative semantic communication for multi-user systems with large language models. *IEEE Wirel. Commun.* 32 (5), 170–178.
- Yang, W., Xiong, Z., Quek, T.Q., Shen, X., 2024. Streamlined transmission: A semantic-aware XR deployment framework enhanced by generative AI. *IEEE Netw.* 38 (6), 29–38.
- Yang, Y., Yang, Z., Yuan, W., Liu, F., Cao, X., Huang, C., Zhang, Z., Shikh-Bahaei, M., 2024. Integrated sensing, computing and secure semantic communication for smart healthcare systems. In: *Proc. IEEE Globecom Workshops (GC Wkshps)*. IEEE, pp. 1–6.
- Yao, L., Wei, Y., Zhou, Y., Li, Z., Chi, N., Wong, K.-K., 2026. You only transmit once: Unified generation and comprehension for efficient semantic communication. *IEEE J. Sel. Areas Commun.* 44, 3349–3364.
- Zhang, H., Di, B., Zhang, H., Niyato, D., Song, L., 2026. Feature-foundation model evolution for low-latency semantic communication. *IEEE Trans. Wirel. Commun.* 25, 11915–11929.
- Zhang, G., Hu, Q., Qin, Z., Cai, Y., Yu, G., 2022. A unified multi-task semantic communication system with domain adaptation. In: *Proc. IEEE Global Commun. Conf. (GLOBECOM)*. IEEE, pp. 3971–3976.
- Zhang, J., Li, S., Sun, Y., Cai, K., Chen, H., Xu, X., Cui, S., 2026. Cooperative semantic knowledge base update for semantic communication networks. *IEEE Trans. Commun.* 74, 1391–1405.
- Zhang, S., Lim, W.Y.B., Ng, W.C., Xiong, Z., Niyato, D., Shen, X.S., Miao, C., 2023. Toward green metaverse networking: Technologies, advancements, and future directions. *IEEE Netw.* 37 (5), 223–232.
- Zhang, Y., Sun, Y., Guo, L., Chen, W., Ai, B., Gunduz, D., 2026. Multimodal LLM integrated semantic communications for 6G immersive experiences. *IEEE Commun. Mag.* 64 (5), 68–75, arXiv:2507.04621.
- Zhang, H., Wang, H., Li, Y., Long, K., Nallanathan, A., 2023. DRL-driven dynamic resource allocation for task-oriented semantic communication. *IEEE Trans. Commun.* 71 (7), 3992–4004.
- Zhang, M., Yang, Y., Song, B., Chen, Z., Ma, X., Li, D., Li, N., Cao, F., Zhang, Y., Zhang, H., 2024. Industrial semantic communications for tool wear monitoring in industrial IoT networks. In: *Proc. 2024 IEEE/CIC Int. Conf. Commun. China (ICCC Workshops)*. IEEE, pp. 248–253.
- Zhang, W., Zhang, H., Ma, H., Shao, H., Wang, N., Leung, V.C., 2023. Predictive and adaptive deep coding for wireless image transmission in semantic communication. *IEEE Trans. Wirel. Commun.* 22 (8), 5486–5501.
- Zhao, Z., Yang, Z., Huang, C., Wei, L., Yang, Q., Zhong, C., Xu, W., Zhang, Z., 2024. A joint communication and computation design for distributed RISs assisted probabilistic semantic communication in IIoT. *IEEE Internet Things J.* 11 (16), 26568–26579.
- Zhou, H., Deng, Y., Liu, X., Pappas, N., Nallanathan, A., 2024. Goal-oriented semantic communications for 6G networks. *IEEE Internet Things Mag.* 7 (5), 104–110.
- Zhou, F., Li, Y., Zhang, X., Wu, Q., Lei, X., Hu, R.Q., 2022. Cognitive semantic communication systems driven by knowledge graph. In: *Proc. ICC 2022-IEEE Int. Conf. Commun.*. IEEE, pp. 4860–4865.
- Zhou, Q., Li, R., Zhao, Z., Peng, C., Zhang, H., 2021. Semantic communication with adaptive universal transformer. *IEEE Wirel. Commun. Lett.* 11 (3), 453–457.
- Zhu, L., Liu, Z., Han, S., 2019. Deep leakage from gradients. In: *Adv. Neural Inf. Process. Syst.*, vol. 32.