RATE-DISTORTION OPTIMIZED COMMUNICATION FOR COLLABORATIVE PERCEPTION

Anonymous authorsPaper under double-blind review

000

001

002003004

006

008 009

010 011

012

013

014

015

016

017

018

019

021

024

025

026

027 028 029

031

032

033

034

037

040

041

042

043

044

046

047

048

051

052

ABSTRACT

Collaborative perception emphasizes enhancing environmental understanding by enabling multiple agents to share visual information with limited bandwidth resources. While prior work has explored the empirical trade-off between task performance and communication volume, a significant gap remains in the theoretical foundation. To fill this gap, we draw on information theory and introduce a pragmatic rate-distortion theory for multi-agent collaboration, specifically formulated to analyze performance-communication trade-off in goal-oriented multiagent systems. This theory concretizes two key conditions for designing optimal communication strategies: supplying pragmatically relevant information and transmitting redundancy-less messages. Guided by these two conditions, we propose RDcomm, a communication-efficient collaborative perception framework that introduces two key innovations: i) task entropy discrete coding, which assigns features with task-relevant codeword-lengths to maximize the efficiency in supplying pragmatic information; ii) mutual-information-driven message selection, which utilizes mutual information neural estimation to approach the optimal redundancy-less condition. Experiments on 3D object detection and BEV segmentation demonstrate that RDcomm achieves state-of-the-art accuracy on DAIR-V2X and OPV2V, while reducing communication volume by up to 108×. The code will be released.

1 Introduction

Multi-agent collaborative perception enhances environmental understanding by enabling agents to jointly perceive and share information. This paradigm has shown clear advantages over single-agent sensing, particularly in overcoming occlusions and limited fields of view, and has been widely adopted in tasks such as 3D object detection Wang et al. (2020) and BEV segmentation Xu et al. (2022a).

In this field emphasizing multi-agent collaboration, a fundamental challenge remains: the trade-off between task performance and communication volume Hu et al. (2022). While sharing richer information tends to preserve collaboration quality, it introduces significant communication overhead; conversely, aggressively limiting communication may result in the loss of task-critical information, ultimately degrading overall performance. Recent works tackle this trade-off by extracting informative and compact representations from visual observations to serve as collaborative messages. One representative line of work focuses on spatial selection, aiming to transmit only task-relevant regions, such as those with high detection confidence Hu et al. (2022) or sparse observation coverage Xu et al. (2025). Another direction leverages neural compression techniques Ballé et al. (2018); Van Den Oord et al. (2017), reducing the size of transmitted features through autoencoders Shao et al. (2024); Hu et al. (2024) or channel reduction Li et al. (2021); Lu et al. (2024). Despite some empirical gains, previous approaches are heuristic in nature, relying on manually designed communication strategies or intuitive criteria. Such approaches lack theoretical grounding and provide no principled guidance on what to communicate or how to encode it under bandwidth constraints.

To fill this gap, we take an information-theoretic perspective and propose *pragmatic rate-distortion* theory for multi-agent collaboration, which explicitly models the trade-off between communication bit-rate and task-specific pragmatic distortion. Our theoretic analysis extends Shannon's classical rate-distortion framework Shannon et al. (1959) in two key aspects. First, we introduces pragmatic distortion, a task-driven metric that reflects the impact of message degradation on downstream task performance, distinct from reconstruction-based distortions Blau & Michaeli (2019); Cover (1999).

Second, our theory generalizes to distributed communication among multiple agents, where both message senders and receivers observe the environment. We thus account for inter-agent redundancy, a factor typically neglected in traditional rate—distortion analysis. Building upon these extensions, our theory ultimately characterizes the minimal communication cost required to meet a specified distortion threshold, and derives two key conditions that an optimal communication strategy should satisfy: *pragmatic-relevant* and *redundancy-less*. We envision this theoretical framework as a foundation for analyzing communication efficiency in broader multi-agent tasks.

Inspired by the two theoretical conditions, we propose **RDcomm** (**R**ate-**D**istortion guided pragmatic **comm**unication), a novel communication-efficient collaborative perception system, which optimizes both message selection and coding to reduce communication overhead while preserving collaborative complementarity and task effectiveness. Specifically, we design the two core components of RDcomm based on the two derived conditions: i) Based on the pragmatic-relevant condition, we propose a novel task entropy discrete coding module. It first utilizes learned codebooks to quantize feature vectors, and then applies variable-length coding guided by task relevance, assigning shorter codewords to more informative features. ii) Based on the redundancy-less condition, we propose a novel feature selection module leveraging mutual information neural estimation Belghazi et al. (2018). This module enables agents to perform an inter-agent handshake process to assess message redundancy by quantifying the mutual information between shared and locally observed features. We validate RDcomm on two representative perception tasks: 3D object detection and BEV semantic segmentation, using both real-world dataset DAIR-V2X Yu et al. (2022) and simulation dataset OPV2V Xu et al. (2021). Experimental results show that RDcomm reduces communication volume by up to 108 times compared against existing methods.

Our main contributions are summarized as follows:

- We introduce a pragmatic rate-distortion theory for multi-agent collaboration, which characterizes the performance-communication trade-off, and concretize two optimal conditions: i) supply pragmatic-relevant information; ii) avoid inter-agent redundancy.
- We propose RDcomm, a communication-efficient collaborative perception framework that
 is designed to approach the two optimal conditions with two innovations: i) task entropy
 discrete coding; ii) mutual-information-driven message selection. Experiments on detection and segmentation tasks demonstrate that RDcomm achieves dual superiority in both
 performance and communication efficiency.

2 Related works

2.1 COMMUNICATION-EFFICIENT COLLABORATIVE PERCEPTION

In multi-agent collaboration, a key challenge is to balance task performance and communication cost Hu et al. (2022). Early collaboration transmits raw sensor data Han et al. (2023) and achieves high accuracy but suffers from heavy bandwidth usage. Late collaboration reduces bandwidth by sending final predictions, but degrades performance under noise Lu et al. (2023); Hu et al. (2022). To address this, intermediate collaboration transmits feature maps to strike a balance between performance and efficiency. Prior works mainly improve efficiency via: i) spatial selection, which transmits features at critical regions; and ii) feature compression. For spatial selection, Where2comm Hu et al. (2022) selects high-confidence regions, CodeFilling Hu et al. (2024) removes redundant collaborators, and CoSDH Xu et al. (2025) targets unobserved areas. For compression, techniques include value quantization Wang et al. (2020); Shao et al. (2024), vector quantization Hu et al. (2024), and channel reduction Li et al. (2021); Lu et al. (2024). However, these methods are primarily heuristic and lack theoretical guarantees for communication efficiency. In our work, we provide a theoretical framework grounded in rate-distortion analysis, offering explicit conditions for optimal communication.

2.2 Rate-distortion theory background

Rate-distortion theory Shannon et al. (1959) provides a fundamental framework for lossy compression by characterizing the minimum bits required to represent a signal X as a compressed representation Z under a specified distortion constraint $D[X,Z] \leq \delta$. The goal is to find a probabilistic encoding map p(Z|X) that minimizes the mutual information I(X;Z), as formulated in (1):

$$\operatorname{Rate}(\delta) = \min_{p(Z|X)} I(X;Z) \quad \text{s.t. } D[X,Z] \le \delta$$
 (1)

In lossy compression, Z is an approximate reconstruction of X. A typical example is a Gaussian source $X \sim \mathcal{N}(\mu, \sigma^2)$ with mean squared error (MSE) distortion, where the optimal rate under distortion level δ has a closed-form solution $R(\delta) = h(X) - \frac{1}{2}\log(2\pi e\delta)$, which reflects the total information in X minus the portion tolerable under the distortion budget. Despite its general application in visual compression Ballé et al. (2018), classical rate-distortion analysis mostly caters to fidelity-based distortion metrics Blau & Michaeli (2019) and single-source settings. In general, any distortion measure $d: \mathcal{X} \times \mathcal{Z} \to \mathbb{R}_{\geq 0}$ of the form $D[X,Z] = \mathbb{E}_{p(x,z)}[d(x,z)]$ is valid, as long as there exists a $z \in \mathcal{Z}$ such that D[X,Z] is finite Dubois et al. (2021). Our work extends this framework to multi-agent collaborative perception by incorporating task-specific distortion and inter-agent redundancy.

3 Pragmatic rate-distortion theory for collaboration

We introduce the pragmatic rate-distortion theory for multi-agent collaboration. Our theoretical analysis follows three high-level steps: i) We reformulate the collaboration task objective in a rate-distortion formulation (Sec. 3.1); ii) Defining the pragmatic distortion for collaboration (Sec. 3.2); iii) We derive the minimal transmission rate under constrained distortion, and present the conditions for an optimal communication strategy (Sec. 3.3). See detailed proofs in the Appendix A.3.

3.1 PROBLEM FORMULATION

We first introduce the problem we target to solve in rate-distortion formulation. In collaborative tasks, our goal is to optimize model parameters and the message generation strategy, in order to achieve the minimal transmission bits under constrained task loss, that is (2):

$$\min_{Z_{s},\theta} \sum_{r=1}^{N} |\{Z_{s\to r}\}_{s\neq r}| \quad \text{s.t. } \sum_{r=1}^{N} L_{Y}\left(Y_{r}, \Phi_{\theta}\left(X_{r}, \{Z_{s\to r}\}_{s\neq r}\right)\right) \leq L_{max}$$
 (2)

where N denotes the number of agents; $\{Z_{s \to r}\}_{s \neq r}$ represents the messages sent from agent s to other agents; $|\cdot|$ measures the information volume; $\Phi_{\theta}(\cdot)$ is the task model parameterized by θ ; X_r and Y_r denote the local observation and ground-truth label of agent r, respectively; $L_Y(\cdot)$ is the loss function associated with task Y, and L_{\max} specifies the maximum tolerable task loss. We then present (3) as a formal rate-distortion optimization version of (2), which defines the minimal communication bits $\mathrm{Rate}(\delta)$ and serves as the foundational objective.

$$\operatorname{Rate}(\delta) = \min_{p(Z_{s \to r}|X_s)} I(X_s; Z_{s \to r}) \quad \text{s.t. } \operatorname{D}_Y[X_s, Z_{s \to r}|X_r] \le \delta. \tag{3}$$

As shown in (3), the communication volume is captured by $I(X_s; Z_{s \to r})$, which quantifies the amount of information from the original observation X_s that is preserved in the transmitted message $Z_{s \to r}$. We denote $D_Y[X_s, Z_{s \to r} \mid X_r]$ as the pragmatic distortion for the collaborative task Y, which measures the degradation in task performance when transmitting $Z_{s \to r}$ instead of X_s , given the local observation X_r . See detailed discussions on our problem formulation in Appendix A.3.1.

3.2 Pragmatic distortion for collaborative perception

In this section we make the pragmatic distortion $D_Y[X_s, Z_{s \to r} \mid X_r]$ in (3) explicit for collaborative perception. Let X_s , X_r denote the sender's and receiver's local observations, respectively, Y denotes the perception task target. $Z_{s \to r}$ is compressed from X_s . We define $D_Y[X_s, Z_{s \to r} | X_r]$ in (4):

$$D_Y\left[X_s, Z_{s \to r} | X_r\right] = B_{risk}\left[Y | Z_{s \to r}, X_r\right] - B_{risk}\left[Y | X_s, X_r\right] \tag{4}$$

where $B_{risk}[Y|X_r,X_s]$ denotes the Bayes Risk Dubois et al. (2021), which measures the minimum achievable prediction error for the target Y given the joint inputs X_r and X_s . We define the distortion $D_Y[X_s,Z_{s\to r}|X_r]$ as the increase in Bayes Risk when predicting Y with the compressed representation $Z_{s\to r}$ instead of the original signal X_s , while conditioning on the existing local information X_r . Formally, the Bayes Risk is $B_{risk}[Y|X] = \inf_f \mathbb{E}_{p(X,Y)} \left[L(Y,f(X))\right]$, where f is any predictor. For perception tasks, the loss is typically computed independently at each BEV location. Accordingly, we define the overall Bayes Risk as the average of pixel-wise Bayes Risks over all locations, as in (5), where $i \in \mathcal{S}$ denotes a BEV coordinate within the perception range \mathcal{S} , and $Y_{(i)}$ and $f(X)_{(i)}$ denote the corresponding ground-truth label and model prediction.

$$B_{risk}[Y|X] = \inf_{f} \mathbb{E}_{p(X,Y)} \left[L(Y, f(X)) \right] = \inf_{f} \frac{1}{|Y|} \sum_{i \in S} \mathbb{E}_{p(X,Y_{(i)})} \left[L(Y_{(i)}, f(X)_{(i)}) \right]$$
 (5)

We instantiate the pragmatic distortion for two representative perception tasks: BEV segmentation and 3D object detection. Specifically, for BEV segmentation, we adopt the per-pixel cross-entropy (CE) loss; for 3D object detection, we adopt the widely used CenterPoint loss Yin et al. (2021). We directly present the final derivations of the pragmatic distortion $D_Y[X_s, Z_{s \to r} | X_r]$ in Tab. 1. Note that we express pragmatic distortions in terms of entropy, which is naturally adaptive for information theory analysis.

Table 1: Pragmatic distortions differ from classical reconstruction distortion by considering task entropy $H(Y|\cdot)$ and local redundancy X_r . Proofs are provided in Appendix A.3.4.

Task (loss function)	Distortion $D_Y[X_s, Z_{s \rightarrow r} X_r]$	$Z_{s \to r} X_r $	
Lossy reconstruction Ballé et al. (2018) (MSE)	$rac{1}{ X_s } \left\ X_s - Z_{s ightarrow r} ight\ _2^2, ext{no } X_r$	(6)	
BEV segmentation (CE)	$\frac{1}{ Y } \sum_{s \in S} [\operatorname{H}(Y_{(i)} \mid Z_{s \to r}, X_r) - \operatorname{H}(Y_{(i)} \mid X_s, X_r)]$	(7)	
3D detection. (CenterPoint)	$\frac{1}{ Y } \sum_{i \in \mathcal{S}} [H(Y_{(i,c)} \mid Z_{s \to r}, X_r) - H(Y_{(i,c)} \mid X_s, X_r) + \frac{1}{2} \sum_{k \in \mathcal{K}} \left(e^{H(Y_{(i,k)} \mid Z_{s \to r}, X_r) - 1} - e^{H(Y_{(i,k)} \mid Z_{s \to r}, X_r) - 1} \right) $	$ X_s,X_r)-1$	
	NO.	(8)	

3.3 MINIMAL BIT-RATE OF COLLABORATIVE MESSAGE AND OPTIMAL CONDITIONS

In this section, we present the trade-off between communication bit-rate and distortion by incorporating the pragmatic distortions (7) and (8) into objective (3).

Theorem 1 (Minimal bit-rate $\operatorname{Rate}(\delta)$ of collaborative message under distortion constraint, see proof in A.3.5). Consider a message sender agent a_s and a message receiver agent a_r and their observation denoted as X_s, X_r , where the sender compresses X_s as $Z_{s \to r}$ and transmits it to the receiver to collaborates in achieving task target Y. Then, the minimal transmission bit-rate $\operatorname{Rate}(\delta) = \min_{p(Z_{s \to r}|X_s)} I(X_s; Z_{s \to r})$ s.t. $\operatorname{D}_Y[X_s, Z_{s \to r}] \leq \delta$ is given in (10).

$$\operatorname{Rate}(\delta) = \min_{\substack{p(Z_{s \to r} \mid X_s) \text{ s.t. } \operatorname{D}_Y[X_s, Z_{s \to r}] \leq \delta}} \operatorname{I}(X_s; Z_{s \to r}) = \operatorname{I}(Y; X_s \mid X_r) - \delta$$

$$= \operatorname{H}(X_s) - \underbrace{\operatorname{H}(X_s \mid Y)}_{\text{information in } X_s \text{ irrelevant to } Y} - \underbrace{\operatorname{I}(Y; X_s; X_r)}_{\text{information in } X_s \text{ redundant with } X_r \text{ about } Y} - \delta$$

$$(9)$$

The minimal bit-rate $\mathrm{Rate}(\delta)$ can be achieved only if the following two conditions are satisfied. **Pragmatic-relevant.** The transmitted message $Z_{s\to r}$ should contain only information relevant to the receiver's task Y, as formalized in (11):

$$H(Z_{s \to r}|Y) = 0 \tag{11}$$

Equation (11) implies that the uncertainty of $Z_{s \to r}$ is eliminated given a specific task target Y, indicating a unique mapping from Y to $Z_{s \to r}$. Consequently, $Z_{s \to r}$ should exclude any information unrelated to the task Y, and the task-relevant messages should be prioritized during coding.

Redundancy-less. The transmitted message $Z_{s\to r}$ should avoid maintaining information that is already contained in the receiver's observation X_r , which is expressed as (12):

$$I(Z_{s\to r}; X_r) = 0 \tag{12}$$

Equation (12) shows that the mutual information between the transmitted message $Z_{s\to r}$ and the observation of receiver X_r should be eliminated. In other words, $Z_{s\to r}$ should avoid containing the redundant information, but exactly supply the information missed in X_r .

4 RDCOMM: EFFICIENT COMMUNICATION FOR CO-PERCEPTION

Inspired by the theoretical analysis, we introduce **RDcomm**, a novel communication-efficient collaborative perception framework with three main components (Fig. 1): i) a perception pipeline that provides the basic functionalities for perception tasks; ii) a task entropy discrete coding module (Sec. 4.1), following the pragmatic-relevant condition (11), which adopts a novel variable-length coding guided by task relevance; iii) a mutual-information-driven selection module (Sec. 4.2), following the redundancy-less condition (12), which selects complementary messages for transmission.

Perception pipeline. The perception pipeline couples a BEV encoder with task-specific decoders. The BEV encoder accepts either LiDAR or camera inputs and maps sensor data into a unified bird's-eye-view (BEV) representation, enabling consistent spatial alignment across agents. Denoting the

217

218 219

220

221

222 223

224

225 226 227

228 229

230

231

232

233

234

235 236

237 238

239

240

241

242

243

244

245

246

247 248

249 250

251

253

254

255

256

257

258 259

260

261

266

267

268

Figure 1: RDcomm features two key components: i) task entropy discrete coding for improving the pragmatic relevance of message, which assigns short codewords to the codes with high confidence frequency; ii) mutualinformation-driven message selection, which measures message redundancy by mutual information estimation.

observation of the sender agent by X_s , the BEV encoder $\Phi_{\text{bev}}(\cdot)$ produces a BEV feature map by $F_s = \Phi_{\text{bev}}(X_s) \in \mathbb{R}^{h \times w \times c}$, followed by task-specific decoders Φ_{task} for downstream tasks such as 3D object detection and BEV segmentation. Backbone details are provided in Appendix A.2.1. We then focus on compressing the feature F_s into a collaboration message $Z_{s\to r}$ for any receiver agent.

4.1 TASK ENTROPY DISCRETE CODING

Our first objective is to approach the pragmatic-relevant condition $H(Z_{s\to r}|Y)=0$ in (11). We approximate it by minimizing the task-conditioned entropy $H(Z_{s\to r}|Y)$.

Layered vector quantization. We begin by constraining $H(Z_{s\to r})$ via vector quantization inspired by Zhu et al. (2022), where the core idea is to represent each vector in F_s with the nearest embedding \mathbf{e}_i in a codebook $\mathbf{B} = [\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n] \in \mathbb{R}^{n \times d}$ with n learnable embeddings. We further implement a layered discrete auto-encoder $\Phi_{vq}(\cdot)$ to quantize the BEV feature F_s by $F_s^q = \Phi_{vq}(F_s, \mathbf{B}_{base}, \mathbf{B}_{res})$. ${f B}_{base}$ is used to approximate the basic coarse-grained information of F_s with small codebook volume n, and the residual error $F_s - \Phi_{vq}(F_s, \mathbf{B}_{base})$ is further approximated by a fine-grained codebook ${\bf B}_{res}$ with larger volume. This layered quantization is described in (13)(14):

$$Z_{base}^{q} = \arg\min_{i} \|\mathbf{B}_{base}[i] - f_{in}(F_{s})\|_{2}, \quad Z_{res} = F_{s} - Z_{base}^{q}$$

$$Z_{res}^{q} = \arg\min_{i} \|\mathbf{B}_{res}[i] - f_{in}(Z_{res})\|_{2}, \quad F_{s}^{q} = f_{out}(Z_{res}^{q} + Z_{base}^{q})$$
(13)

$$Z_{res}^q = \arg\min_{i} \|\mathbf{B}_{res}[i] - f_{in}(Z_{res})\|_2, \quad F_s^q = f_{out}(Z_{res}^q + Z_{base}^q)$$
 (14)

where $f_{in}(\cdot)$, $f_{out}(\cdot)$ are MLP projectors to bridge the distribution gap between continuous vectors and codebook embeddings, and the input feature map F_s is flattened before quantization.

Task-aware priority and encoding. While quantization reduces representation length by restricting the vector-space volume, we further improve coding efficiency by introducing task bias. Specifically, we prioritize task-relevant messages for selection and encode them with shorter code lengths. Recall that our objective is to minimize the task-conditioned entropy $H(Z_{s\to r}|Y)\to 0$, which can be expanded as (15).

$$\min_{Z_{s \to r}} \mathbb{E}_Y \sum_{Z_{s \to r}} \left[-p(Z_{s \to r}|Y) \log p(Z_{s \to r}|Y) \right] \tag{15}$$

Note that maximizing $p(Z_{s\to r}|Y)$ to 1 provides a sufficient solution for minimizing $H(Z_{s\to r}|Y)$, and $p(Z_{s\to r}|Y) \propto p(Y|Z_{s\to r})p(Z_{s\to r})$ for a given target distribution p(Y). Therefore we priority the messages with high task confidence $p(Y|Z_{s\rightarrow r})$ for transmission. We implement this with a confidence generator $\Phi_{\mathrm{conf}}(\cdot)$ producing scores $C_s = \Phi_{\mathrm{conf}}(F_s) \in \mathbb{R}^{h \times w}$. The confidence mask is $M_c = \mathbf{1}[C_s > \tau_c]$, where τ_c is a confidence threshold, and quantized features are selected as $F_{sc}^q = M_c \odot F_s^q$. In practice, we instantiate $\Phi_{\mathrm{conf}}(\cdot)$ by reusing the task decoder $\Phi_{\mathrm{task}}(\cdot)$.

We further reduce the average coding length of the quantized F_{sc}^q , where we encode each embedding e_i in $\{B_{base}, B_{res}\}$ considering the joint effect of task confidence $p(Y|e_i)$ and occurrence frequency $p(\mathbf{e}_i)$. Specifically, we define the confidence frequency $p_c(\cdot)$ for each embedding \mathbf{e}_i as (16):

$$p_c(\mathbf{e}_i) = \sum_{F_s} \sum_{\{(u,v): \mathbf{e}_i \in \Phi_{vq}(F_s)[u,v]\}} \Phi_{conf}(F_s)[u,v]$$

$$\tag{16}$$

where $p_c(\mathbf{e}_i)$ represents the total task confidence predicted from \mathbf{e}_i across the entire dataset. We compute it by accumulating the confidence scores $\Phi_{\mathrm{conf}}(F_s)[u,v]$ at spatial locations (u,v) where $F_s[u,v]$ are quantized with embedding \mathbf{e}_i . To improve coding efficiency for the task-relevant embeddings, we propose to assign shorter code lengths for the embeddings with higher confidence frequency. For implementation, this work provides a straightforward yet effective solution by applying Huffman coding Huffman (2007), where we set the Huffman weight of \mathbf{e}_i to be its confidence frequency $p_c(\mathbf{e}_i)$. Our task-entropy coding ultimately produces an index map $D_s \in (\{0,1\}^{l_{u,v}})^{h \times w}$ to represent the input feature $F_s \in \mathbb{R}^{h \times w \times c}$, where $\{0,1\}^{l_{u,v}}$ denotes binary strings of variable length $l_{u,v}$. $D_s[u,v] = [D_s^{base}[u,v] \parallel D_s^{res}[u,v]] \in \{0,1\}^{l_{u,v}}$ denotes the code at location (u,v), which consists of coarse-grained semantic information $D_s^{base}[u,v]$ and fine-grained $D_s^{res}[u,v]$.

Discussion. i) The relation between coding method and the information optimization target (15): the high-confidence selection module $\Phi_{\operatorname{conf}(\cdot)}$ reduces the task-conditioned entropy $\operatorname{H}(Z_{s \to r}|Y)$ governed by $p(Z_{s \to r}|Y) \propto p(Y|Z_{s \to r})p(Z_{s \to r})$. We instantiate this by defining the confidence frequency $p_c(\cdot)$, and use it as the weight for entropy coding. ii) The intuition behind is simple: with limited bandwidth, high-confidence messages are selected more often, making their embeddings frequent. Accordingly, we weight entropy coding using the confidence frequency $p_c(\mathbf{e}_i)$. iii) We use the accumulation of confidence rather than the average to distinguish embeddings with similar task relevance but different occurrence frequencies. iv) The confidence-frequency coding is applied post-training and remains lossless for task performance, as it only reassigns embedding indices.

4.2 Complementary Selection with Mutual Information Estimation

In this section we focus on reducing inter-agent message redundancy according to the redundancy-less condition $I(Z_{s\to r}; X_r) = 0$ (12). We approximate this target by (17). Ω denotes the selection region.

$$\min_{\Omega} I(\hat{F}_{sc}^q[\Omega]; F_r) \tag{17}$$

To obtain the redundancy-less region Ω , we perform feature selection via mutual information neural estimation: we approximate $I(\hat{F}_{sc}^q; F_r)$ by with a learnable estimator $\Phi_{\mathrm{MI}}(\hat{F}_{sc}^q, F_r)$ and select features accordingly. Note that we use the coarse-grained compression $\hat{F}_{sc}^q = \mathbf{B}_{base}[D_s^{base}]$ as an abstract of F_s , which is pre-handed to the receiver to estimate its redundancy with the receiver's information F_r . This abstraction \hat{F}_{sc}^q helps estimate semantic redundancy by transmitting D_s^{base} with much smaller communication volume, which is 10 times smaller than the lossless message D_s .

Mutual information neural estimation. Consider the variable pair $\mathbf{s}, \mathbf{r} \in \mathbb{R}^c$ that are included in $\hat{F}^q_{sc}, F_r \in \mathbb{R}^{h \times w \times c}$ respectively, the mutual information $I(\mathbf{s}, \mathbf{r})$ can be estimated using (18):

$$I(\mathbf{s}, \mathbf{r}) := D_{KL}(\mathbb{P}_{\mathbf{s}, \mathbf{r}} \parallel \mathbb{P}_{\mathbf{s}} \otimes \mathbb{P}_{\mathbf{r}}) \ge \sup_{T: \mathbb{R}^c \times \mathbb{R}^c \to \mathbb{R}} \left\{ \mathbb{E}_{\mathbb{P}_{\mathbf{s}, \mathbf{r}}}[T(\mathbf{s}, \mathbf{r})] - \log(\mathbb{E}_{\mathbb{P}_{\mathbf{s}} \otimes \mathbb{P}_{\mathbf{r}}}[e^{T(\mathbf{s}, \mathbf{r})}]) \right\}$$
(18)

where $\mathbb{P}_{\mathbf{s},\mathbf{r}}$ denotes the joint distribution of \mathbf{s} and \mathbf{r} , and $\mathbb{P}_{\mathbf{s}},\mathbb{P}_{\mathbf{r}}$ denote marginals; D_{KL} is KL divergence; $T:\mathbb{R}^c\times\mathbb{R}^c\to\mathbb{R}$ is any function that projects the (\mathbf{s},\mathbf{r}) pair into a real number. Formulation (18) shows that mutual information $I(\mathbf{s},\mathbf{r})$ actually measures the divergence between the joint distribution $\mathbb{P}_{\mathbf{s},\mathbf{r}}$ and marginal distribution $\mathbb{P}_{\mathbf{s}}\otimes\mathbb{P}_{\mathbf{r}}$. While (18) employs the KL divergence as a standard estimation, we instead adopt the GAN-style divergence Nowozin et al. (2016) following Li et al. (2020) to facilitate optimization, and use the estimation target in (19). $\sigma(\cdot)$ denotes sigmoid.

$$\hat{\mathbf{I}}(\mathbf{s}, \mathbf{r}) \ge \sup_{T: \mathbb{R}^{c} \times \mathbb{R}^{c} \to \mathbb{R}} \left\{ \mathbb{E}_{\mathbb{P}_{\mathbf{s}, \mathbf{r}}} [\log \sigma(T(\mathbf{s}, \mathbf{r}))] + \mathbb{E}_{\mathbb{P}_{\mathbf{s}} \otimes \mathbb{P}_{\mathbf{r}}} [\log (1 - \sigma(T(\mathbf{s}, \mathbf{r})))] \right\}$$
(19)

The inequality results from approximating $T(\cdot)$ with limited representation ability, which is a subclass of projection $\mathbb{R}^c \times \mathbb{R}^c \to \mathbb{R}$. Our objective is to optimize $T(\cdot)$ to maximize the estimation bound in (18), where $T(\cdot)$ is implemented by a learnable estimator $\Phi_{\mathrm{MI}}(\cdot)$, with the loss defined in (20):

$$\mathcal{L}_{MI} = -\frac{1}{|\mathcal{P}_{\mathbf{s},\mathbf{r}}|} \sum_{(\mathbf{s},\mathbf{r})\in\mathcal{P}_{\mathbf{s},\mathbf{r}}} \log \sigma \left(\Phi_{\mathrm{MI}}(\mathbf{s},\mathbf{r})\right) - \frac{1}{|\mathcal{P}_{\mathbf{s}}||\mathcal{P}_{\mathbf{r}}|} \sum_{\mathbf{s}\in\mathcal{P}_{\mathbf{s}},\mathbf{r}\in\mathcal{P}_{\mathbf{r}}} \log \left(1 - \sigma \left(\Phi_{\mathrm{MI}}(\mathbf{s},\mathbf{r})\right)\right)$$
(20)

where $\mathcal{P}_{s,r}$ denotes the set sampled from $\mathbb{P}_{s,r}$ where s, r are visual features from two agents at the same location. $\mathcal{P}_{s}, \mathcal{P}_{r}$ are sampled from the marginals where s, r are randomly combined. We see

that the mutual information estimator $\Phi_{\mathrm{MI}}(\cdot)$ actually serves as a discriminator. It predicts whether a feature pair (\mathbf{s},\mathbf{r}) is drawn from the joint distribution $\mathbb{P}_{\mathbf{s},\mathbf{r}}$, i.e., two agents' observations at the same location that exhibit similar patterns, or from the product of marginals $\mathbb{P}_{\mathbf{s}}\otimes\mathbb{P}_{\mathbf{r}}$, in which observations are randomly paired and likely to differ in pattern due to non-corresponding locations.

Redundancy-less feature selection. To reduce redundancy in transmission, we prioritize the message s with a low mutual information score $\Phi_{\mathrm{MI}}(\mathbf{s},\mathbf{r})$. Specifically, we obtain the redundancy map $R_{s \to r} = \Phi_{\mathrm{MI}}(\hat{F}_{sc}^q, F_r) \in \mathbb{R}^{h \times w}$ and derive the redundancy-less selection mask $M_{MI} = 1[R_{s \to r} < \tau_{MI}]$, where τ_{MI} is a redundancy threshold. The feature finally sent to the receiver is $Z_{s \to r} = M_{MI} \odot \hat{F}_{sc}^q$. This selection filters out the messages in $Z_{s \to r}$ that are already covered by F_r . The total communication volume of RDcomm is computed as $|D_s \odot M_c \odot M_{MI}| + |D_s^{base} \odot M_c|$. The first term measures the volume of the selected lossless information $Z_{s \to r}$, and the second term measures the cost for identifying redundancy and is much smaller than the first term by setting \mathbf{B}_{base} with small codebook size n and dimension d with limited segments.

Message smoothing and fusion. Note that $Z_{s \to r}$ is obtained under sparse masks M_C and M_{MI} . Although it preserves salient information, the sparsity may degrade semantic content. We mitigate this by applying a UNet Ronneberger et al. (2015) $\Phi_{\rm smth}(\cdot)$ to smooth and dilate $Z_{s \to r}$, propagating sparse signals to neighboring regions. The receiver then obtains the enhanced perception results as $\bar{Y}_r = \Phi_{\rm task}(\Phi_{\rm fusion}(F_r, \Phi_{\rm smth}(Z_{s \to r}))$, where $\Phi_{\rm fusion}(\cdot)$ is instantiated using the effective maxfusion operation following Hu et al. (2024).

4.3 Training

We train RDcomm in three stages. First, we train the BEV encoder Φ_{bev} and the task decoder Φ_{task} with task loss \mathcal{L}_{task} , which corresponds to the CenterPoint loss for 3D detection and the per-pixel cross-entropy loss for BEV segmentation. Second, we train the vector quantization module Φ_{vq} with both task loss \mathcal{L}_{task} and feature reconstruction loss $\mathcal{L}_{recon} = \|F_s^q - F_s\|_2^2$. After that, the confidence frequency $p_c(\cdot)$ is updated. Finally, we train the mutual-information estimator Φ_{MI} using \mathcal{L}_{MI} . In the later stages of training, the thresholds τ_c, τ_{MI} are randomly varied to facilitate bandwidth adaptation.

5 EXPERIMENTS

To evaluate RDcomm, we conduct experiments on two representative collaborative perception tasks: 3D object detection and BEV semantic segmentation. To validate the effectiveness and generality of the proposed RDcomm, our evaluation spans both LiDAR and camera modalities, up to 5 collaborating agents, and varying bandwidth constraints.

Collaborative 3D detection. We evaluate collaborative 3D detection on two datasets: the real-world DAIR-V2X Yu et al. (2022) and the simulated OPV2V Xu et al. (2021). DAIR-V2X is a vehicle-to-infrastructure dataset with 9K frames of 2-agent collaboration between a vehicle and a roadside unit (RSU). Each agent is equipped with a LiDAR and a 1920×1080 camera, where the RSU uses a 300-channel LiDAR and the vehicle a 40-channel LiDAR. OPV2V is a vehicle-to-vehicle dataset simulated with CARLA Dosovitskiy et al. (2017), containing 12K frames. Our experiments involve up to 3 agents, each equipped with a 64-channel LiDAR and four RGB cameras at 800×600 resolution. We evaluate both LiDAR and camera modalities and report Average Precision (AP) at IoU thresholds of 30%, 50%, and 70%. Following Lu et al. (2024), the perception range is set to 204.8m×102.4m. Collaborative BEV segmentation. We evaluate BEV semantic segmentation on OPV2V dataset following CoBEVT Xu et al. (2022a). Each agent predicts a BEV semantic occupancy map with camera inputs, ground truth classes include dynamic vehicles, drivable area, and lane. We involve collaboration among up to 5 agents. Performance is measured by Intersection-over-Union (IoU) between predictions and ground-truth BEV labels, with the perception range fixed at 100m×100m.

5.1 QUANTITATIVE ANALYSIS

Benchmark comparison. Fig. 2 and Fig. 3 compare RDcomm with previous collaborative perception methods in terms of the trade-off between perception performance and communication volume on 3D detection and BEV segmentation. The baselines include no collaboration, CodeFilling Hu et al. (2024), V2X-ViT Xu et al. (2022b), CoBEVT Xu et al. (2022a), DiscoNet Li et al. (2021), AttFuse Xu et al. (2021), V2VNet Wang et al. (2020), Fcooper Chen et al. (2019), HEAL Lu et al. (2024), and late Fusion (directly share the final perception results). We see that: i) RDcomm consistently delivers a superior perception—communication trade-off across all bandwidth settings for both detection and segmentation; the gains persist across LiDAR/camera modalities and multiple semantic classes.

Figure 2: RDcomm achieves the best performance–communication trade-off in 3D detection, across DAIR-V2X/OPV2V datasets with LiDAR/camera input modalities.

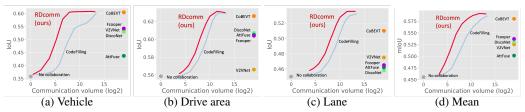
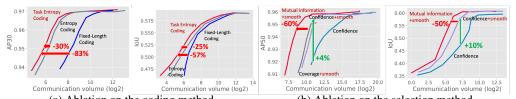


Figure 3: RDcomm achieves the best performance-communication trade-off in BEV segmentation.

ii) Under extreme bandwidth constraints, RDcomm achieves larger gains than prior methods: for detection, +11.49/+19.82% (LiDAR/camera) on DAIR-V2X and +12.01/+22.92% on OPV2V with a 50K times reduction relative to uncompressed features; for segmentation, +5.69% mIoU at a 1K times reduction. iii) RDcomm outperforms previous communication-efficient SOTA, CodeFilling, with significantly reduced communication cost in detection: 15/13 times less (LiDAR/camera) on DAIR-V2X and 54/108 times less on OPV2V; and 8 times less for segmentation.

Ablation on coding method. Fig. 4a compares the proposed task entropy coding in RDcomm against: i) classic entropy coding weighted by occurrence frequency Huffman (2007); ii) fixed-length coding Hu et al. (2024), where code length is $\log 2$ of the codebook volume. We see that task entropy coding saves 83/57%(detection/segmentation) in communication volume compared to fixed-length coding, and 30/25% compared to occurrence-driven entropy coding. The reason is that our task entropy coding prioritizes pragmatic-rich codes by assigning them shorter codewords, whereas classic entropy coding may waste short codewords on high-frequency but pragmatically weak codes. Theoretically, we use confidence mask M_C to reduce the entropy $\mathrm{H}(Z_{s \to r}|Y)$, and entropy coding weighted by confidence frequency $p_c(\cdot)$ shortens the expected code length toward the entropy limit implied by this entropy.

Ablation on selection method. Fig. 4b compares our mutual-information—driven (MI) redundancy selection against other redundancy selection: i) confidence-based Hu et al. (2024); ii) LiDAR-coverage—based Xu et al. (2025). Our MI selection reduces communication volume by 60/50% (detection/segmentation) relative to these baselines. The gain arises since MI leverages a pragmatic yet lightweight abstraction to identify redundancy, providing richer cues than one-dimensional confidence or coverage signals. Theoretically, the mutual information estimator encourages agents to transmit $Z_{s \to r}$ with low $I(Z_{s \to r}, F_r)$ relative to the receiver's feature F_r . Meanwhile, we observe that the smoothing module improves AP50 by 4% and IoU by 10%, demonstrating its effectiveness in mitigating sparsity under high selection rates.



(a) Ablation on the coding method (b) Ablation on the selection method Figure 4: Ablation study on OPV2V detection/segmentation, evaluating coding and selection methods.

Lossless bit-rate. We compare the bit-rate of RDcomm with the optimal bit-rate $\mathrm{Rate}(\delta)$ in (10) under lossless pragmatic compression, i.e., $\delta=0$. To reduce the error in estimating $\mathrm{Rate}(0)$, we exclude the receiver's information X_r in the experiment; the optimal rate is then $\mathrm{I}(Y;X_s)$, which is tightly upper-bounded by $\mathrm{H}(Y)\approx\log_2(4)$, i.e., 2 bits-per-pixel (bpp). Here we define "lossless" as a perfor-

Table 2: Segmentation performance with compression. bpp: bits-per-pixel.

Method	mIoU↑	bpp↓
VQVAE-32	0.38 (82%)	10
VQVAE-128	0.40 (87%)	14
RDcomm-128	0.44 (95%)	4
No compression	0.46 (100%)	4096

mance drop of less than 5%. Tab. 2 compares RDcomm's

BEV segmentation performance with the no-compression scheme. RDcomm attains 95% of mIoU with 4 bpp, close to the 2 bpp upper bound. Here bpp describes BEV feature. We also report VQVAE Van Den Oord et al. (2017) results (codebook sizes 32/128, segment number 2), where RDcomm uses a codebook size of 128. RDcomm delivers higher accuracy at a lower rate, indicating that effective pragmatic compression cannot be achieved by merely tuning codebook size, requires selective transmission and shorter codewords for task-relevant codes.

Cost for transmitting abstract. Tab. 3 reports the share of bandwidth consumed by transmitting the pragmatic abstraction \hat{F}_{sc}^q on DAIR-V2X detection. We observe that abstraction transmission accounts for only 9%–11% of the total communication volume, yet is effective to identify redundancy as showed in Fig. 4b.

Table 3: Allocation of communication volume.

AP30	0.82	0.79	0.76
total bits	9054	449	166
abstract bits	882(9%)	49 (11%)	19 (11%)

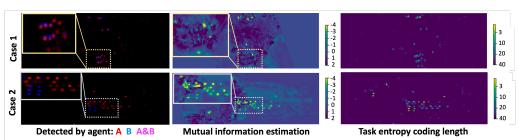


Figure 5: Visualization of mutual information estimation and task entropy coding length on DAIR-V2X.

5.2 QUALITATIVE ANALYSIS

Fig. 5 visualizes mutual-information estimates and task-entropy code lengths in two cases. In case 1, both agents A and B detect the same vehicles, forming high-mutual-information regions. In case 2, A and B detect different vehicles, yielding low-mutual-information regions that are prioritized for sharing. We also observe that task-entropy coding assigns short codewords to task-relevant regions. Long codewords are assigned to background areas, which are omitted when bandwidth is limited.

6 CONCLUSIONS

This work investigates the trade-off between task performance and communication volume from an information-theoretic perspective. We formulate a pragmatic rate—distortion theory for collaborative perception, deriving the optimal bit rate for message transmission and two necessary conditions for optimal compression: pragmatic-relevant and redundancy-less. Guided by these two conditions, we propose RDcomm, a communication-efficient collaborative perception method with two novel components: i) task entropy discrete coding and ii) mutual-information-driven message selection. Experiments covering both detection and segmentation show that RDcomm achieves state-of-the-art perception—communication trade-offs across both LiDAR and camera modalities.

Limitations and future work. Our study focuses on perception tasks. Future directions include extending the framework to broader tasks such as navigation, manipulation, and scene captioning, and incorporating additional modalities such as motion and language.

Reproducibility Statement. We have taken multiple steps to ensure the reproducibility of our work. The detailed formulations of our theoretical results, including all assumptions and proofs, are provided in Appendix A.3. The design of RDcomm and its components is described in Sec. 4, with training details in Sec. 4.3 and backbone details reported in Appendix. A.2.1. The datasets used in our experiments (DAIR-V2X and OPV2V) are publicly available. We will release our source code and configuration files in the camera-ready version to further facilitate reproducibility.

REFERENCES

- Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. Variational image compression with a scale hyperprior. *arXiv preprint arXiv:1802.01436*, 2018.
- Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeshwar, Sherjil Ozair, Yoshua Bengio, Aaron Courville, and Devon Hjelm. Mutual information neural estimation. In *International conference on machine learning*, pp. 531–540. PMLR, 2018.
- Yochai Blau and Tomer Michaeli. Rethinking lossy compression: The rate-distortion-perception tradeoff. In *International Conference on Machine Learning*, pp. 675–685. PMLR, 2019.
- Qi Chen, Xu Ma, Sihai Tang, Jingda Guo, Qing Yang, and Song Fu. F-cooper: Feature based cooperative perception for autonomous vehicle edge computing system using 3d point clouds. In *Proceedings of the 4th ACM/IEEE Symposium on Edge Computing*, pp. 88–100, 2019.
- Thomas M Cover. Elements of information theory. John Wiley & Sons, 1999.
- Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator. In *Conference on robot learning*, pp. 1–16. PMLR, 2017.
- Yann Dubois, Benjamin Bloem-Reddy, Karen Ullrich, and Chris J Maddison. Lossy compression for lossless prediction. *Advances in Neural Information Processing Systems*, 34:14014–14028, 2021.
- Yushan Han, Hui Zhang, Huifang Li, Yi Jin, Congyan Lang, and Yidong Li. Collaborative perception in autonomous driving: Methods, datasets, and challenges. *IEEE Intelligent Transportation Systems Magazine*, 15(6):131–151, 2023.
- Yue Hu, Shaoheng Fang, Zixing Lei, Yiqi Zhong, and Siheng Chen. Where2comm: Communication-efficient collaborative perception via spatial confidence maps. *Advances in neural information processing systems*, 35:4874–4886, 2022.
- Yue Hu, Juntong Peng, Sifei Liu, Junhao Ge, Si Liu, and Siheng Chen. Communication-efficient collaborative perception via information filling with codebook. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15481–15490, 2024.
- David A Huffman. A method for the construction of minimum-redundancy codes. *Proceedings of the IRE*, 40(9):1098–1101, 2007.
- Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Point-pillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 12697–12705, 2019.
- Maosen Li, Siheng Chen, Ya Zhang, and Ivor Tsang. Graph cross networks with vertex infomax pooling. *Advances in neural information processing systems*, 33:14093–14105, 2020.
- Yiming Li, Shunli Ren, Pengxiang Wu, Siheng Chen, Chen Feng, and Wenjun Zhang. Learning distilled collaboration graph for multi-agent perception. *Advances in Neural Information Processing Systems*, 34:29541–29552, 2021.
- Yifan Lu, Quanhao Li, Baoan Liu, Mehrdad Dianati, Chen Feng, Siheng Chen, and Yanfeng Wang. Robust collaborative 3d object detection in presence of pose errors. In 2023 IEEE International Conference on Robotics and Automation (ICRA), pp. 4812–4818. IEEE, 2023.
- Yifan Lu, Yue Hu, Yiqi Zhong, Dequan Wang, Yanfeng Wang, and Siheng Chen. An extensible framework for open heterogeneous collaborative perception. *arXiv preprint arXiv:2401.13964*, 2024.
- Sebastian Nowozin, Botond Cseke, and Ryota Tomioka. f-gan: Training generative neural samplers using variational divergence minimization. *Advances in neural information processing systems*, 29, 2016.
- Jonah Philion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16*, pp. 194–210. Springer, 2020.

- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pp. 234–241. Springer, 2015.
 - Claude E Shannon et al. Coding theorems for a discrete source with a fidelity criterion. *IRE Nat. Conv. Rec*, 4(142-163):1, 1959.
 - Jiawei Shao, Teng Li, and Jun Zhang. Task-oriented communication for vehicle-to-infrastructure cooperative perception. In 2024 IEEE 34th International Workshop on Machine Learning for Signal Processing (MLSP), pp. 1–6. IEEE, 2024.
 - Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.
 - Tsun-Hsuan Wang, Sivabalan Manivasagam, Ming Liang, Bin Yang, Wenyuan Zeng, and Raquel Urtasun. V2vnet: Vehicle-to-vehicle communication for joint perception and prediction. In Computer vision–ECCV 2020: 16th European conference, Glasgow, UK, August 23–28, 2020, proceedings, part II 16, pp. 605–621. Springer, 2020.
 - Junhao Xu, Yanan Zhang, Zhi Cai, and Di Huang. Cosdh: Communication-efficient collaborative perception via supply-demand awareness and intermediate-late hybridization. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 6834–6843, 2025.
 - Runsheng Xu, Hao Xiang, Xin Xia, Xu Han, Jinlong Liu, and Jiaqi Ma. Opv2v: An open benchmark dataset and fusion pipeline for perception with vehicle-to-vehicle communication. 2022 International Conference on Robotics and Automation (ICRA), pp. 2583–2589, 2021.
 - Runsheng Xu, Zhengzhong Tu, Hao Xiang, Wei Shao, Bolei Zhou, and Jiaqi Ma. CoBEVT: Cooperative bird's eye view semantic segmentation with sparse transformers. *CoRL*, 2022a.
 - Runsheng Xu, Hao Xiang, Zhengzhong Tu, Xin Xia, Ming-Hsuan Yang, and Jiaqi Ma. V2x-vit: Vehicle-to-everything cooperative perception with vision transformer. In *European conference on computer vision*, pp. 107–124. Springer, 2022b.
 - Tianwei Yin, Xingyi Zhou, and Philipp Krahenbuhl. Center-based 3d object detection and tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 11784–11793, 2021.
 - Haibao Yu, Yizhen Luo, Mao Shu, Yiyi Huo, Zebang Yang, Yifeng Shi, Zhenglong Guo, Hanyu Li, Xing Hu, Jirui Yuan, et al. DAIR-V2X: A large-scale dataset for vehicle-infrastructure cooperative 3d object detection. *In Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition (CVPR)*, 2022.
 - Xiaosu Zhu, Jingkuan Song, Lianli Gao, Feng Zheng, and Heng Tao Shen. Unified multivariate gaussian mixture for efficient neural image compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 17612–17621, 2022.

A APPENDIX

A.1 STATEMENTS

- **LLM Usage.** We used an LLM (ChatGPT) solely for language refinement, such as improving grammar, clarity, and readability of sentences. The research ideas, methodology, experiments, and overall writing structure were entirely developed by the authors.
- A.2 MODEL
- A.2.1 SINGLE-AGENT PERCEPTION PIPELINE
- The perception pipeline comprises two components: a BEV encoder and a task-specific decoder. The BEV encoder projects sensor inputs into bird's-eye-view (BEV) representations, enabling consistent

spatial alignment and collaboration across different views. Task-specific decoders are then applied for downstream detection or segmentation tasks.

BEV Encoder. Our framework supports either LiDAR or camera inputs, where we commonly denote the observation of the ith agent as X_i . We obtain BEV feature as $F_i = \Phi_{enc}(X_i) \in \mathbb{R}^{h \times w \times c}$, where $\Phi_{enc}(\cdot)$ denotes the complete BEV encoder. F_i then serves as the information source for selection and compression. For LiDAR inputs, we adopt the widely used PointPillars encoder Lang et al. (2019) to extract BEV features from point clouds. For camera inputs, we employ the Lift-Splat-Shoot Philion & Fidler (2020) module following Lu et al. (2024), which lifts image features into 3D frustums and aggregates them into the BEV plane via learned depth distributions. For both LiDAR and camera modality, the extracted BEV features are further processed by a 2D convolutional ResNeXt-based backbone Lu et al. (2024).

Decoder. Based on the BEV feature F_i , we incorporate task-specific decoders. For 3D object detection, the decoder consists of a classification head, a box regression head, and a direction estimation head to predict object bounding boxes, following Lu et al. (2024); Hu et al. (2022). For BEV semantic segmentation, we employ a MLP as decoder to produce dense, per-pixel semantic predictions, following Xu et al. (2022a).

A.3 THEORY

In this section, we provide: i) further discussion on our problem formulation; ii) proof of the proposed propositions and theories.

A.3.1 DISCUSSION ON PROBLEM FORMULATION

Note that our theory formulation equation 2 equation 3 is not an ad hoc assumption, but consistent with the learning objects in pragmatic compression. Table 4 reveals that the optimization objective equation 3 is dual-equivalent to several prior approaches Hu et al. (2022); Ballé et al. (2018), as they all share the same La-

Table 4: Problem formulations of bandwidth-constrained collaboration. \mathcal{R} : bit-rate, \mathcal{D} : distortion/loss.

Formulation Type	Optimization target	Distortion/loss \mathcal{D}
Constrained task loss Hu et al. (2022) Task-compression joint loss Ballé et al. (2018) Pragmatic rate-distortion (ours)	$\begin{aligned} & \min \mathcal{D} \text{s.t. } \mathcal{R} \leq \delta \\ & \min \mathcal{D} + \lambda \mathcal{R} \\ & \min \mathcal{R} \text{s.t. } \mathcal{D} \leq \delta \end{aligned}$	task loss function task loss function pragmatic distortion

grangian target min $\mathcal{D} + \lambda \mathcal{R}$, which is a weighted sum of distortion \mathcal{D} and communication rate \mathcal{R} with weight λ . In the Section 3.2 we make the task distortion $D_Y[X_s, Z_s \mid X_T]$ explicit.

As an extreme case, consider early collaboration Han et al. (2023), where agents directly transmit raw sensor data (i.e., $Z_s = X_s$). In this case, the communication volume becomes $I(X_s; Z_s) = H(X_s)$, corresponding to the full information content of X_s , and the distortion $D_Y[X_s, X_s|X_T]$ is zero.

A.3.2 DISCUSSION ON PRAGMATIC DISTORTION

From Tab. 8 we see that: i) for collaborative BEV segmentation, pragmatic distortion is expressed as the gap in conditional entropy, based on the Bayes risk $B_{risk}[Y|X] = \frac{1}{|Y|} \sum_{i \in \mathcal{S}} H\left(Y_{(i)}|Z_s,X_r\right)$; ii) for collaborative 3D detection, the distortion further incorporates an exponential term of conditional entropy, where $\mathcal{K} = \{loc, size, ori\}$ denotes the set of regression losses (location, size, and orientation), which contribute more uncertainty than classification; iii) compared to the widely used MSE distortion in image reconstruction Ballé et al. (2018), the pragmatic distortion defined in our theory differs in two aspects: first, it emphasizes task uncertainty rather than fidelity, second, it accounts for the receiver's information (e.g., X_r) to analyze redundancy.

A.3.3 Proof of Proposition 1: Bayes risk R[Y|X] for perception tasks

Proposition 1 (Bayes risk R[Y|X] for perception tasks, see proof in A.3.3). Given an observation input X and a perception task target Y, we focus on the Bayes risk R[Y|X] to measure the difficulty of predicting Y from X. In object detection task, the detection results and corresponding target label are denoted as $\hat{Y}, Y \in \mathbb{R}^{h \times w \times (8+K)}$, where the (8+K) channels stand for classification heatmap $\hat{Y}_c, Y_c \in \mathbb{R}^{h \times w \times K}$, offset $\hat{Y}_o, Y_o \in \mathbb{R}^{h \times w \times 3}$, size $\hat{Y}_s, Y_s \in \mathbb{R}^{h \times w \times 3}$, rotation $\hat{Y}_r, Y_r \in \mathbb{R}^{h \times w \times 2}$. Here [h, w] denotes the BEV perception range. The total loss is shown in equation 21:

$$L_{total} = L_{heatmap} + \lambda_2 L_{offset} + \lambda_3 L_{size} + \lambda_4 L_{rotation}$$
 (21)

where $\lambda_2, \lambda_3, \lambda_4$ are the loss weights. Consider N objects involved in ground truth, the heatmap loss optimizes the model to classify the foreground object from background, where we adopt the focal loss $L_{focal}(y,\hat{y}) = -\sum_{k=1}^K \alpha_k \left(1-\hat{y}_k\right)^\gamma y_k \log \hat{y}_k$, where α_k, γ are hyper-parameters in focal loss, here we consider a simplified situation that $\alpha_k = 1, \gamma = 0$ and the loss degenerates into cross entropy loss L_{ce} . L_{offset}, L_{size} , and $L_{rotation}$ are LI loss. Specifically, we consider the situation that the elements in Y_o follow Gaussian distribution $Y_{o(i,j)}|X \sim \mathcal{N}(\mu_o(X), \sigma_o^2(X))$, $Y_{s(i,j)}|X \sim \mathcal{N}(\mu_s(X), \sigma_s^2(X)), Y_{r(i,j)}|X \sim \mathcal{N}(\mu_r(X), \sigma_r^2(X))$, and the number of objects is $\overline{N_{obj}}$. The Bayes risk of object detection with centerpoint detection loss is given as equation 22:

$$R[Y|X] = \sum_{i \le h, j \le w} H(Y_{c(i,j)}|X) + \overline{N_{obj}} \sqrt{2/\pi} \left(\lambda_2 \sigma_o(X) + \lambda_3 \sigma_s(X) + \lambda_4 \sigma_r(X)\right)$$
(22)

When the task Y refers to occupancy prediction, the regression terms are put off and the Bayes risk R[Y|X] solely consists of the terms of conditional entropy, as shown in equation 23.

$$R[Y|X] = \sum_{i \le h, j \le w} H(Y_{c(i,j)}|X) = \sum_{i \le h, j \le w} H(Y_{(i,j)}|X)$$
(23)

In each communication round, messages are transmitted between connected agents as shown in equation 2, where the connection is established by the pre-defined collaboration principle. We denote the observation of message sender/reciever as X_s, X_r , the perception target as Y. The message $\mathcal{P}_{s \to r}$ is obtained via $\mathcal{P}_{s \to r} = C(X_s)$, where $C(\cdot)$ is a compressor that reduces the transmission bit-rate. The pragmatic distortion is defined in equation 24, where Y is the perception target, R[Y|X] denotes the Bayes risk when predicting Y from X, $R[Y|X_r, X_s]$ denotes the Bayes risk when predicting Y from the fused information of X_r, X_s .

$$D_{Y}[X_{s}, Z_{s}] = R[Y \mid X_{r}, Z_{s}] - R[Y \mid X_{r}, X_{s}]$$
(24)

To analyze this distortion in perception tasks, we need to:

- Give the specific formulation of Bayes risk $R[Y|X_r]$ (single perception) and $R[Y|X_r,X_s]$ (collaborative perception) in detection 3D task with centerpoint loss(for example).
- Reformulate the distortion D_Y by introducing the task related Bayes risk $R[Y|X_r,X_s]$.
- Reformulate the distortion D_V by introducing the supply-request information.

Definition 1 (Bayes risk). Let $X \in \mathcal{X}$ be the input variable (features, observed data), $Y \in \mathcal{Y}$ be the target variable (labels), P(X,Y) denote the joint probability distribution of X and Y, $L(Y,\hat{Y})$ be the loss function quantifying the discrepancy between a prediction $\hat{Y} = f(X)$ and the true value Y, and $f: \mathcal{X} \to \mathcal{Y}$ be a predictive model. The Bayes risk is defined as the infimum of the expected loss over all possible decision functions, as shown in equation 25:

$$R = \inf_{f} \mathbb{E}_{X,Y} \left[L(Y, f(X)) \right]$$
 (25)

Bayes risk is the minimum achievable loss by an ideally trained model. It captures unavoidable uncertainty in the data, such as the ambiguity due to overlapping classes in classification tasks or stochastic noise in target variables for regression tasks. For any model f, the expected loss satisfies $\mathbb{E}[L(Y, f(X))] \geq R_{Bayes}$. With a given loss function, the Bayes risk completely depends on the data distribution P(X, Y), it indicates the "difficulty" of learning the projection $f: \mathcal{X} \to \mathcal{Y}$. Due to its property to characterize data distribution, we utilize the differences of Bayes risk to measure the pragmatic distortion.

Bayes risk for perception tasks In this section, we derive the Bayes risk of perception tasks with specific loss functions.

First, we review the formulation of centerpoint loss. Suppose that the observation from camera or LiDAR can be represented by 3D voxel feature $X \in \mathbb{R}^{D \times h \times w \times C}$, the detection results and corresponding target label are denoted as $\hat{Y}, Y \in \mathbb{R}^{h \times w \times (8+K)}$, where the (8+K) channels stand for classification heatmap $\hat{Y}_c, Y_c \in \mathbb{R}^{h \times w \times K}$, offset $\hat{Y}_o, Y_o \in \mathbb{R}^{h \times w \times 3}$, size $\hat{Y}_s, Y_s \in \mathbb{R}^{h \times w \times 3}$, rotation $\hat{Y}_T, Y_T \in \mathbb{R}^{h \times w \times 2}$. The total loss is:

$$L_{\text{total}} = L_{\text{heatmap}} + \lambda_2 L_{\text{offset}} + \lambda_3 L_{\text{size}} + \lambda_4 L_{\text{rotation}}$$
 (26)

where $,\lambda_2,\lambda_3,\lambda_4$ are the loss weights. Consider N objects in ground truth, the heatmap loss optimizes the model to classify the foreground object from background, we utilize focal loss $L_{\rm focal}\left(y,\hat{y}\right)=-\sum_{k=1}^K\alpha_k\left(1-\hat{y}_k\right)^\gamma y_k\log\hat{y}_k$, where α_k,γ are hyper-parameters in focal loss, here we consider a simplified situation that $\alpha_k=1,\gamma=0$ and the loss degenerates into cross entropy loss $L_{ce}.$ $L_{\rm offset}$, $L_{\rm size}$, and $L_{\rm rotation}$ are L1 loss.

Now we derive the Bayes risk in 3D object detection with centerpooint detection loss. To simplify the formulation, we approximately decompose the total Bayes risk into the sum of Bayes risk on each location as shown in equation 27:

$$R[Y|X] = \inf_{f} \mathbb{E}_{X,Y} [L(Y, f(X))] = \inf_{f} \sum_{i < h, j < w} \mathbb{E}_{X,Y_{(i,j)}} [L(Y_{(i,j)}, f(X)_{(i,j)})]$$
(27)

We regard the perception task at each region as independent tasks, and we define the located Bayes risk for perception tasks in equation 28:

$$R[Y_{(i,j)}|X] = \inf_{f} \mathbb{E}_{X,Y_{(i,j)}} \left[L(Y_{(i,j)}, f(X)_{(i,j)}) \right]$$
(28)

We derive the Bayes risk for two primarily used loss function involved in perception tasks: focal loss, MSE loss, and L1 loss.

Focal loss. The true distribution of $P(Y_c, X)$ satisfies $p(Y_{c(i,j,k)} = 1|X) = p_{i,j,k}$, and we have $P(Y_c|X) = \prod_{i,j,k} p(Y_{c(i,j,k)}|X)$ since different classes and locations are independent. The Bayes optimal prediction is the true conditional distribution: $\hat{Y}_{c(i,j,k)}^* = p_{(i,j,k)}$, then we have:

$$R[Y_{c(i,j)}|X] = \mathbb{E}_{X,Y_{c(i,j)}} L_{ce}(Y_{c(i,j)}, p_{(i,j)})$$
(29)

$$= \mathbb{E}_X \sum_{Y_{c(i,j)}} p(Y_{c(i,j)}|X) L_{ce}(Y_{c(i,j)}, p_{(i,j)})$$
(30)

$$= \mathbb{E}_X \sum_{k=1}^K -p(Y_{c(i,j,k)} = 1|X) \log p(Y_{c(i,j,k)} = 1|X)$$
(31)

$$= H(Y_{c(i,j)}|X) \tag{32}$$

MSE loss. Given a specific X, we assume that the elements in offset target Y_o follow Gaussian distribution $Y_{o(i,j)}|X \sim \mathcal{N}(\mu_o(X), \sigma_o^2(X))$, and the number of objects is $\overline{N_{obj}}$. This assumption is reasonable, since minimizing MSE loss can be regarded as MLE(Maximum likelihood estimation) when $Y_{o(i,j)}|X \sim \mathcal{N}(\mu_o(X), \sigma_o^2(X))$. The Bayes optimal prediction is $f(X) = \hat{Y}_{o(i,j)}^* = \mu_o(X)$. Put this into equation 27, the Bayes risk is derived as $\overline{N_{obj}}\sqrt{2/\pi}\sigma_o(X)$. Similarly, we can derive the Bayes risk for the size and rotation targets by assuming their distributions follow a Gaussian distribution $Y_{s(i,j)}|X \sim \mathcal{N}(\mu_s(X), \sigma_s^2(X)), Y_{r(i,j)}|X \sim \mathcal{N}(\mu_r(X), \sigma_r^2(X))$. Combining the Bayes risk of the individual loss function described in equation 26, we obtain the Bayes risk of object detection with centerpoint detection loss as equation 33:

$$R_{centerpoint}[Y|X] = \sum_{i \le h, j \le w} H(Y_{c(i,j)}|X) + \overline{N_{obj}} \sqrt{2/\pi} \left(\lambda_2 \sigma_o(X) + \lambda_3 \sigma_s(X) + \lambda_4 \sigma_r(X)\right)$$
(33)

L1 loss. Given a specific X, we assume that the elements in target Y_o follow Laplace distribution $p(Y_o \mid X) = \frac{1}{2b_{o\mid X}} \exp\left(-\frac{|Y-\mu_o(X)|}{b_{o\mid X}}\right)$. This assumption is reasonable, since minimizing L1 loss can be regarded as MLE(Maximum likelihood estimation) when $Y_{o(i,j)}|X \sim \frac{1}{2b_{o\mid X}} \exp\left(-\frac{|Y-\mu_o(X)|}{b_{o\mid X}}\right)$. The Bayes optimal prediction is $f(X) = \hat{Y}^*_{o(i,j)} = \operatorname{median}(Y_{o(i,j)} \mid X) = \mu_o(X)$. Put this into equation 27, the Bayes risk is derived as:

$$= \mathbb{E}_X \frac{1}{2b_{o|X}} 2b_{o|X}^2 \tag{37}$$

$$= \mathbb{E}_X b_{o|X}$$

$$= b_{o|X}$$

$$b_{o|X} \text{ is a constant}$$

$$(39)$$

On the other hand, when $p(Y_{o(i,j)} = y|X = x) = \frac{1}{2b_{o|X}} \exp\left(-\frac{|y-\mu_o(x)|}{b_{o|X}}\right)$, we can formulate the conditional entropy $H(Y_{o(i,j)}|X)$ as:

$$H(Y_{o(i,j)}|X) = \mathbb{E}_{x \sim X} H(Y_{o(i,j)}|X = x)$$

$$= \mathbb{E}_{x \sim X} \int_{-\infty}^{\infty} -p(y) \log p(y) dy$$

$$(41)$$

$$= \mathbb{E}_{x \sim X} \int_{-\infty}^{\infty} -\frac{1}{2b_{o|X}} \exp\left(-\frac{|y - \mu_o(x)|}{b_{o|X}}\right) \left(\log \frac{1}{2b_{o|X}} - \frac{|y - \mu_o(x)|}{b_{o|X}}\right) dy \qquad (42)$$

$$= \mathbb{E}_{x \sim X} \log(2b_{o|X}) + 1$$

$$= \log(2b_{o|X}) + 1$$

$$b_{o|X} \text{ is a constant}$$

$$(43)$$

Combining equation 39 with equation 44, we have:

$$R[Y_{o(i,j)}|X] = \frac{1}{2}e^{H(Y_{o(i,j)}|X)-1}$$
(45)

(44)

A.3.4 Proof of Tab. 1: Pragmatic distortion for collaborative perception.

In this section, we derive the pragmatic distortion in collaborative perception. To achieve this, we start from decomposing the contribution of ego agent and other agents.

Now we derive the pragmatic distortion in collaborative perception task. Consider a simple scenario with 2 collaborators, and the observations/features of message sender and receiver are X_s and X_r , and the sender compresses X_s into Z_s to transmit, we define the pragmatic distortion as shown in equation 46, which measures the increase of Bayes risk after replacing the collaboration message X_s with Z_s :

$$D_{Y}[X_{s}, Z_{s}] = R[Y \mid X_{r}, Z_{s}] - R[Y \mid X_{r}, X_{s}]$$
(46)

We give a specific formulation by replacing the Bayes risk in equation 46 with the Bayes risk of centerpoint loss in equation 33, as shown in equation 47:

$$D_{Y_{(i,j)}}[X_s, Z_s] = H(Y_{c(i,j)}|X_r, Z_s) - H(Y_{c(i,j)}|X_r, X_s) +$$
(47)

$$\frac{1}{2}\lambda_2 \left(e^{H(Y_{o(i,j)}|X_r,Z_s)-1} - e^{H(Y_{o(i,j)}|X_r,X_s)-1} \right) + \tag{48}$$

$$\frac{1}{2}\lambda_3(e^{H(Y_{s(i,j)}|X_r,Z_s)-1} - e^{H(Y_{s(i,j)}|X_r,X_s)-1}) +$$
(49)

$$\frac{1}{2}\lambda_4(e^{H(Y_{r(i,j)}|X_r,Z_s)-1} - e^{H(Y_{r(i,j)}|X_r,X_s)-1})$$
(50)

We consider a degraded version by ignoring the regression loss, which is suitable for semantic occupancy prediction task, as shown in equation 51:

$$D_{Y_{(i,j)}}[X_s, Z_s] = H(Y_{(i,j)}|X_r, Z_s) - H(Y_{(i,j)}|X_r, X_s)$$
(51)

A.3.5 PROOF OF THEOREM 1: OPTIMAL BIT-RATE OF COLLABORATIVE MESSAGE

In this section, we derive the optimal transmission bit-rate in collaborative perception task. Consider the same collaboration situation described in Tab. 1 with 2 collaborators, and the observations/features of message sender and receiver are X_s and X_r , and the sender compresses X_s into Z_s to transmit. Our goal is to derive the minimum bit-rate needed to transmit Z_s while guaranteeing a limited pragmatic distortion, as shown in equation 52:

$$\operatorname{Rate}(\delta) = \min_{p(Z_s|X_s)} I(X_s; Z_s) \quad \text{s.t. } \operatorname{D}_Y[X_s, Z_s] \le \delta.$$
 (52)

For occupancy prediction, put pragmatic distortion equation 51 into the constraint in equation 52, we have equation 55:

$$D_{Y}[X_{s}, Z_{s}] = H(Y|X_{r}, Z_{s}) - H(Y|X_{r}, X_{s})$$
(53)

$$= [H(Y|X_r) - I(Y; Z_s|X_r)] - [H(Y|X_r) - I(Y; X_s|X_r)]$$

$$= I(Y; X_s|X_r) - I(Y; Z_s|X_r) < \delta$$
(54)

This inequality condition also satisfies for object detection task, which corresponds to the distortion defined in equation 50 by considering two approximation:

• First-order approximation.

$$e^{H(Y_{o(i,j)}|X_r,Z_s)-1} - e^{H(Y_{o(i,j)}|X_r,X_s)-1} \ge \frac{1}{e} (H(Y_{o(i,j)}|X_r,Z_s) - H(Y_{o(i,j)}|X_r,X_s))$$
(56)

• Decomposition of joint entropy, with the assumption that the existing of $\operatorname{object}(Y_c)$ is independent with the location (Y_o) , size (Y_s) , and heading (Y_r) .

$$H(Y_c, Y_o, Y_s, Y_r) = H(Y_c) + H(Y_o) + H(Y_s) + H(Y_r)$$
 (57)

Given that, We reformulate equation 52 as shown in equation 64:

$$\operatorname{Rate}(\delta) = \min_{p(Z_s|X_s) \text{ s.t. D}_Y[X_s, Z_s] \le \delta} I(X_s; Z_s)$$
(58)

$$\geq \min_{p(Z_s|X_s)} \inf_{\text{s.t. Dy}} I(X_s; Z_s|X_r)$$

$$(59)$$

$$\geq \min_{\substack{p(Z_s|X_s) \text{ s.t. D}_Y[X_s,Z_s] \leq \delta \\ p(Z_s|X_s) \text{ s.t. D}_Y[X_s,Z_s] \leq \delta}} I(X_s;Z_s|X_r)$$

$$\geq \min_{\substack{p(Z_s|X_s) \text{ s.t. D}_Y[X_s,Z_s] \leq \delta \\ p(Z_s|X_s) \text{ s.t. D}_Y[X_s,Z_s] \leq \delta}} I(Y;X_s|X_r)$$

$$\geq \min_{\substack{p(Z_s|X_s) \text{ s.t. D}_Y[X_s,Z_s] \leq \delta \\ p(Z_s|X_s) \text{ s.t. D}_Y[X_s,Z_s] \leq \delta}} I(Y;X_s|X_r) - \delta$$
(61)

$$\geq \min_{p(Z_s|X_s)} \min_{s \in D_Y[X_s|Z_s] \le \delta} I(Y; X_s|X_r) - \delta \tag{61}$$

$$=I(Y; X_s \mid X_r) - \delta \quad (\text{no } Z_s)$$

$$\tag{62}$$

$$= H(X_s) - \left[H(X_s) - I(Y; X_s) \right] - \left[I(Y; X_s) - I(Y; X_s \mid X_r) \right] - \delta$$

$$(63)$$

$$= H(X_s) - \left[H(X_s) - I(Y; X_s)\right] - \left[I(Y; X_s) - I(Y; X_s \mid X_r)\right] - \delta$$

$$= H(X_s) - \underbrace{H(X_s \mid Y)}_{\text{information in } X_s \text{ irrelevant to } Y} - \underbrace{I(Y; X_s; X_r)}_{\text{information in } X_s \text{ redundant with } X_r \text{ about } Y$$

$$(63)$$

We make assumption that the variables follow the Markov chain $Y \leftrightarrow X_s \leftrightarrow Z_s$ and $X_r \leftrightarrow X_s \leftrightarrow Z_s$. Next, we will explain the reasoning behind each inequality and the conditions for these inequality to achieve equality.

The first inequality equation 59 is satisfied when Markov chain $X_r \leftrightarrow X_s \leftrightarrow Z_s$ holds. This is because equation 65:

$$I(Z_s; X_s, X_r) = I(Z_s; X_r) + I(Z_s; X_s | X_r) = I(Z_s; X_s) + I(Z_s; X_r | X_s)$$
(65)

The Markov chain $X_r \leftrightarrow X_s \leftrightarrow Z_s$ leads to $I(Z_s; X_r | X_s) = 0$. Then we have equation 66:

$$I(Z_s; X_s) = I(Z_s; X_r) + I(Z_s; X_s | X_r) \ge I(Z_s; X_s | X_r)$$
(66)

Here we can see the equality condition for the first inequality equation 59 is that, $I(Z_s; X_r) = 0$, which means Z_s , the compressed version of X_s , should not have redundant information in X_r .

The second inequality equation 60 is satisfied since due to DPI(Data Processing Inequality) given the Markov chain $Y \leftrightarrow X_s \leftrightarrow Z_s$. This is because equation 67:

$$I(Z_s; X_s, Y) = I(Z_s; Y) + I(Z_s; X_s | Y) = I(Z_s; X_s) + I(Z_s; Y | X_s)$$
(67)

The Markov chain $Y \leftrightarrow X_s \leftrightarrow Z_s$ leads to $I(Z_s; Y|X_s) = 0$. Then we have equation 68:

$$I(Z_s; X_s) = I(Z_s; Y) + I(Z_s; X_s | Y) > I(Z_s; Y)$$
(68)

We can see that the equality condition for the second inequality equation 60 is that, $I(Z_s; X_s | Y) = 0$. We can derive that equation 69:

$$I(Z_s; X_s | Y) = H(Z_s | Y) - H(Z_s | X_s, Y) = 0$$
(69)

We can see from equation 69 that $H(Z_s|Y) = H(Z_s|X_s,Y)$, since Z_s is a compressed version of X_s , the uncertainty $H(Z_s|X_s,Y)$ is 0, therefore $H(Z_s|Y)=0$. This implies that Z_s is completely task-relative, it does not contains information unrelated to the task Y.

The third inequality is derived from equation 55, and the equality condition is achieved when the distortion budget is sufficiently utilized.