Task-Oriented Source Coding Reduces to Standard Source Coding for Designed Sources

Joss Armstrong

***Abstract*—When the source is designed by a coordination mechanism with smooth utility, the task-oriented source coding programme reduces to standard source coding on the oracle allocation. Taylor's theorem with a zero linear term converts the task loss to MSE on the oracle allocation, from which the task-oriented rate-distortion function, the optimal quantizer, and the rate saving all follow by substitution into standard results. The optimal quantizer is Lloyd-Max on the oracle allocation rather than a task-weighted variant, and the rate saving from task-aware coding is dimensionality reduction rather than a task-specific structural insight. The resource-allocation and power-scheduling examples used to motivate goal-oriented quantization in the recent literature fall within this designed-source subclass, where the general machinery reduces to standard tools.**

***Index Terms*—task-oriented source coding, goal-oriented quantization, rate-distortion, mechanism design, Lloyd-Max**

I. Introduction

The task-oriented source coding programme asks how compression should change when the receiver cares about a downstream task rather than reconstruction fidelity. The answer from the recent literature is that both the rate-distortion function and the quantizer design should be recomputed with respect to the task loss. Zou et al. [1] formalize this as goal-oriented quantization (GOQ), deriving high-resolution optimality conditions that depend on the Jacobian of the decision function and the Hessian of the goal (Propositions 1–2, Algorithm 1). Sun et al. [2] extend the framework to joint precoding and quantization design, with a modified Lloyd iteration whose encoder and decoder conditions are driven by the task loss (Algorithm 2, Eqs. 26–28). The broader programme includes task-based quantization for hardware-limited inference [3], information-bottleneck formulations [4], and the semantic/goal-oriented communication survey of Gündüz et al. [5], which frames the field.

The examples used to motivate the GOQ framework, resource allocation in [1] and power scheduling in [2], share a structural property that the general theory does not exploit. In both cases a system designer chooses the objective, so the task-relevant quantity is a known deterministic function of the source rather than a latent variable to be inferred. This letter shows that under this *designed-source* condition, combined with smooth concave utility (the standard assumption for resource allocation [6]), the entire task-oriented programme reduces to standard tools. The task loss reduces to MSE on the designer's optimal action via Taylor's theorem, and the rate-distortion function, the optimal quantizer, and the rate saving all follow from standard results. In particular, the optimal quantizer is ordinary Lloyd-Max on the optimal action, not a task-weighted variant, and the rate saving is dimensionality reduction rather than a task-specific structural insight.

J. Armstrong is with Ericsson Ireland. ORCID: 0009-0009-3462-9679.

II. Mechanism Class and Variable Mapping

A. The MISES Mechanism Class

A coordinator manages a resource pool on behalf of agents with demand types $t \sim F$ over $\mathcal{T} \subseteq \mathbb{R}^d$. The utility $U : \mathcal{T} \times \mathbb{R}^r \to \mathbb{R}$ satisfies:

A1. β-smoothness: $\nabla^2_{rr} U(t, r) \succeq -\beta I$ for all (t, r).

A2. α-strong concavity: $\nabla^2_{rr} U(t, r) \preceq -\alpha I$ for all (t, r).

Smooth concavity is the standard modelling assumption for resource allocation in mechanism design [6]. The mechanism class and the three theorems that follow from it are established in [7]. The oracle allocation is $\phi^*(t) = \arg\max_r U(t, r)$. A category signal $c : \mathcal{T} \to \{1, \ldots, K\}$ induces the category allocation $\phi(k) = \mathbb{E}[\phi^*(t) \mid c(t) = k]$ and within-category variance $\varepsilon = \mathbb{E}[\|\phi^*(t) - \phi(c(t))\|^2]$.

B. The Task-Oriented Source Coding Framework

Source T with distribution F over $\mathcal{T} \subseteq \mathbb{R}^d$. Action space $\mathcal{A} \subseteq \mathbb{R}^r$. Task loss $\ell : \mathcal{T} \times \mathcal{A} \to \mathbb{R}_{\geq 0}$. Decision rule $g : \hat{\mathcal{T}} \to \mathcal{A}$ maps the compressed representation to an action. The task-oriented rate-distortion function is

$$R_{\text{task}}(D) = \min_{\substack{p(\hat{c}|t): \\ \mathbb{E}[\ell(t, g(\hat{c}))] \leq D}} I(T; \hat{C}). \tag{1}$$

The GOQ framework [1] considers an M-level quantizer $Q_M : \mathcal{T} \to \{z_1, \ldots, z_M\}$ and defines the optimality loss

$$L(Q; f) = \alpha_f \int_{\mathcal{T}} \left[f(\chi(Q(g)); g) - f(\chi(g); g) \right] \varphi(g)\, dg \tag{2}$$

where $\chi(g) = \arg\min_x f(x; g)$ is the optimal decision function, f is the goal function, and φ is the density of g.

C. Variable Mapping

Table I summarizes the correspondence. The structural specialization is that the task loss depends on t only through $\phi^*(t)$, a known, deterministic, dimension-reducing function of the source. This is natural when the mechanism designer chooses ϕ^*.

TABLE I
VARIABLE MAPPING BETWEEN THE TASK-ORIENTED AND MISES FRAMEWORKS.

Task-oriented	MISES	Role
T (source)	t (type)	Input
$\hat{C}$ (compressed)	C (category)	Compressed
$g(\hat{C})$ (action)	$\phi(C)$ (allocation)	Decoded action
$\ell(t, g(\hat{c}))$	$U(t,\phi^*(t)) - U(t,\phi(c))$	Task/welfare loss
$\chi(g)$ (oracle)	$\phi^*(t)$	Full-info action
$R_{\text{task}}(D)$	$\log K_{\min}(D)$	Min. rate

III. MAIN RESULTS

Proposition 1 (Task loss is MSE)**.** *Under A1–A2, for any category assignment c with category allocation $\phi(k) = \mathbb{E}[\phi^*(t) \mid c(t) = k]$,*

$$\frac{\alpha}{2}\,\varepsilon \;\le\; \mathbb{E}\big[\ell(t,\phi(c(t)))\big] \;\le\; \frac{\beta}{2}\,\varepsilon. \tag{3}$$

For quadratic utility ($\alpha = \beta$), the sandwich is tight: $\mathbb{E}[\ell] = \frac{\alpha}{2}\,\varepsilon$.

Proof: Expand $U(t,\cdot)$ around $\phi^*(t)$. The linear term vanishes by the first-order condition $\nabla_r U(t,\phi^*(t)) = 0$. By A1–A2, $-\beta I \preceq \nabla^2_{rr} U \preceq -\alpha I$, so

$$\tfrac{\alpha}{2}\|\phi^*(t)-\phi(k)\|^2 \le U(t,\phi^*(t))-U(t,\phi(k)) \le \tfrac{\beta}{2}\|\phi^*(t)-\phi(k)\|^2.$$

Take expectations over $t \sim F$. ■

Proposition 2 (Task-oriented R-D is standard R-D)**.** *Under A1–A2,*

$$R^{\phi^*}_{\text{MSE}}\!\left(\frac{2D}{\beta}\right) \;\le\; R_{\text{task}}(D) \;\le\; R^{\phi^*}_{\text{MSE}}\!\left(\frac{2D}{\alpha}\right) \tag{4}$$

where $R^{\phi^*}_{\text{MSE}}(\delta)$ is the standard MSE rate-distortion function for the source $\phi^*(T)$. For quadratic utility, $R_{\text{task}}(D) = R^{\phi^*}_{\text{MSE}}(2D/\alpha)$.

Proof: By Proposition 1, $\mathbb{E}[\ell] \le D$ implies $\varepsilon \le 2D/\alpha$ (from the lower bound) and $\varepsilon \ge 2D/\beta$ (from the upper bound). Substitute into the definition of $R^{\phi^*}_{\text{MSE}}$. ■

Proposition 3 (Optimal quantizer is Lloyd-Max)**.**

(a) Decoder: *The category allocation $\phi(k) = \mathbb{E}[\phi^*(t) \mid c(t) = k]$ minimizes the task distortion upper bound among all decoders $\hat{\phi} : \{1,\ldots,K\} \to \mathbb{R}^r$.*

(b) Encoder: *Among all K-partitions of $\mathcal{T}$, the Lloyd-Max quantizer applied to $\phi^*(t)$ minimizes ε and hence the task distortion.*

Proof: (a) The conditional mean minimizes MSE by the MMSE property. By (3), minimizing ε minimizes the task loss bounds. (b) Minimizing ε over K-partitions with conditional-mean centroids is the Lloyd-Max problem on $\phi^*(t)$ by definition. ■

Remark 1 (Reduction of GOQ to Lloyd-Max). Zou et al. [1] derive the GOQ algorithm (Algorithm 1) with a task-weighted nearest-neighbour condition using the matrix $E_{f,\chi}(g) = B_{f,\chi}(g) + A_{f,\chi}(g)$, where $A_{f,\chi} = J_\chi^\top H_f J_\chi$ (Proposition 2), and a gradient-based representative update. They note (p. 48) that when the first-order optimality condition holds, $B_{f,\chi} = 0$. Sun et al. [2] state the conditions most cleanly, with the encoder assigning by task loss (Eq. 26) and the decoder minimizing task loss (Eq. 28).

Under A1–A2, Proposition 1 implies that the task loss is MSE on $\phi^*(t)$, so (i) Zou's task-weighted distance reduces to Euclidean distance on $\phi^*(t)$; (ii) Sun's encoder condition becomes standard nearest-neighbour on $\phi^*(t)$; (iii) Sun's decoder condition becomes the conditional mean. Both algorithms reduce to standard Lloyd-Max on $\phi^*(t)$.

Proposition 4 (Rate saving is dimensionality reduction)**.** *When $r < d$ (allocation dimension $<$ type dimension), $R_{\text{task}}(D) < R_{\text{recon}}(D)$ for all $D > 0$. For Gaussian T with linear oracle allocation $\phi^*(t) = At$, $A \in \mathbb{R}^{r\times d}$, the rate saving at high resolution is*

$$R_{\text{recon}}(D) - R_{\text{task}}(D) \;\approx\; \frac{d-r}{2}\ \textit{nats/symbol.} \tag{5}$$

Proof: By Proposition 2, R_{task} is rate-distortion on $\phi^*(T) \in \mathbb{R}^r$, while R_{recon} is rate-distortion on $T \in \mathbb{R}^d$. For $r < d$, ϕ^* discards $d-r$ irrelevant dimensions. The Gaussian case follows from $R(\delta) = (n/2)\log(\sigma^2/\delta)$. ■

IV. EXAMPLES

A. Gaussian Type, Linear Oracle

Let $T \sim \mathcal{N}(0,\Sigma_T)$ with $\Sigma_T \in \mathbb{R}^{d\times d}$, and $\phi^*(t) = At$ where $A \in \mathbb{R}^{r\times d}$, $r < d$. Utility: $U(t,a) = -(a-At)^\top(a-At)$, so $\alpha = \beta = 2$ (quadratic).

Standard derivation. Characterize the distortion-rate function for the task loss $\ell(t,a) = \|At - a\|^2$, optimizing over encoders and decoders with a non-standard distortion measure. The solution requires eigendecomposition of $A\Sigma_T A^\top$ and reverse water-filling in the eigenspace.

MISES reduction.

1) $\phi^*(T) = AT \sim \mathcal{N}(0, A\Sigma_T A^\top)$, an r-dimensional Gaussian.
2) By Proposition 2 (tight for quadratic): $R_{\text{task}}(D) = R^{\phi^*}_{\text{MSE}}(D)$ = reverse water-filling on the eigenvalues of $A\Sigma_T A^\top$.
3) The optimal K-level quantizer (Proposition 3), standard vector Lloyd-Max on AT.

The rate saving is $(d-r)/2$ nats per symbol, the bits wasted by reconstruction-optimal coding on the $d-r$ dimensions that the task ignores.

B. Nonlinear Oracle, Non-Gaussian Source

Let $T = (T_1, T_2)$ with $T_i \overset{\text{iid}}{\sim} \text{Uniform}[0,1]$. A coordinator allocates a shared resource proportionally to the product of two demand attributes: $U(t,r) = -(r - t_1t_2)^2$, so $\phi^*(t) = t_1t_2$. This is quadratic ($\alpha = \beta = 2$), nonlinear, and dimension-reducing ($d = 2 \to r = 1$).

The source $\phi^*(T) = T_1T_2$ has density $f_{\phi^*}(x) = -\ln x$ on $[0,1]$, which concentrates near zero.

Standard derivation. Optimize a K-level quantizer for the task loss $\ell(t,a) = (t_1t_2 - a)^2$ over the joint distribution of (T_1,T_2), a two-dimensional integral at each Lloyd iteration.

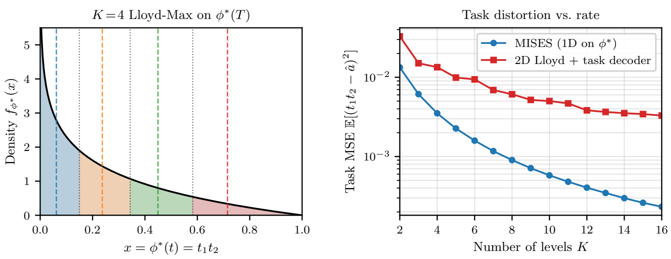

Fig. 1. Nonlinear example: $\phi^*(t) = t_1 t_2$, $T_i \sim \text{Uniform}[0,1]$. Left: density $f_{\phi^*}(x) = -\ln x$ with $K = 4$ Lloyd-Max boundaries (dotted) and representatives (dashed). Right: task MSE vs. number of quantizer levels.

MISES reduction. Run scalar Lloyd-Max on $\phi^*(T) = T_1 T_2$ with density $f(x) = -\ln x$. The two-dimensional task-loss optimization becomes a one-dimensional MSE quantizer design. Figure 1 (left) shows the $K = 4$ Lloyd-Max partition on f_{ϕ^*}, with boundaries bunching near zero where the density concentrates. Figure 1 (right) compares the task MSE of the MISES quantizer (1D Lloyd-Max on $\phi^*(T)$) against a 2D k-means quantizer on T with a task-optimal decoder ($\hat{a}_k = \mathbb{E}[\phi^*(T) \mid \text{cell } k]$). The MISES quantizer achieves lower task distortion at every K because it allocates all codewords to the one-dimensional task-relevant variable, while the 2D quantizer wastes codewords distinguishing types that map to the same oracle allocation.

The example illustrates two effects absent from the Gaussian case: the oracle allocation is nonlinear (ϕ^* is multiplicative, not affine), and the induced source distribution is non-Gaussian (the $-\ln x$ density has no closed-form rate-distortion function). Both effects are irrelevant to the reduction, which requires only A1–A2.

V. Discussion

Propositions 1 to 4 replace the following components of the task-oriented source coding programme:

- The GOQ high-resolution analysis of Zou et al. [1] (Propositions 1–2, ∼5 pp. including appendix proofs) and the task-weighted quantization algorithm (Algorithm 1, ∼2 pp.), which reduce to standard Lloyd-Max on $\phi^*(t)$ (Proposition 3, 4 lines).
- The joint precoding-quantization design of Sun et al. [2] (Algorithm 2, Eqs. 26–28), whose modified encoder/decoder conditions become standard nearest-neighbour and conditional mean (Remark 1).
- The task-oriented rate-distortion characterization, which requires random coding with task-distortion typicality and task-specific Fano arguments in general, and reduces to standard rate-distortion on $\phi^*(T)$ (Proposition 2, 5 lines).
- The finding that $R_{\text{task}} < R_{\text{recon}}$, which reduces to a dimension count (Proposition 4, 3 lines).

The reduction holds because the MISES mechanism class specializes the task-oriented setting in one structural way: the task-relevant quantity $\phi^*(t)$ is a known deterministic function chosen by the designer, not a latent variable to be inferred. Combined with strong concavity, the task loss becomes MSE on $\phi^*(t)$ via Taylor's theorem with a zero linear term, and the entire downstream apparatus inherits the simplification.

The reduction also applies to the information bottleneck (IB) [4]. The IB minimizes $I(T;C)$ subject to $I(C;Y) \geq I_0$, where Y is a relevant variable. Under A1–A2, the relevant variable is $Y = \phi^*(T)$, a deterministic function of T chosen by the designer. Since ϕ^* is deterministic, $I(C;\phi^*(T)) \leq I(C;T)$ with equality when C depends on T only through $\phi^*(T)$. The optimal encoder therefore compresses $\phi^*(T)$ directly, and the three self-consistent IB equations [4] reduce to standard rate-distortion equations on $\phi^*(T)$ with MSE distortion (by Proposition 1). The Gaussian IB is a known special case of this reduction [8]; the MISES result shows it holds for the entire mechanism-design class.

The scope of the reduction is the smooth-utility mechanism class (A1–A2, deterministic oracle allocation). The GOQ framework and its machinery remain necessary for the general stochastic-source setting, where the task-relevant quantity is not designed but given by nature [1, 5]. The contribution here is that a practically important subclass exists, resource allocation, power scheduling, and network coordination, where the task-relevant quantity is a deterministic function chosen by the system designer and the utility is smooth and concave by construction [6]. For this designed-source subclass, a substantially shorter path is available.

References

[1] H. Zou, C. Zhang, S. Lasaulce, L. Saludjian, and H. V. Poor, "Goal-oriented quantization: Analysis, design, and application to resource allocation," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 1, pp. 42–54, 2023.

[2] Y. Sun, H. Zou, C. Zhang, S. Lasaulce, and M. Kieffer, "Goal-oriented compression for L_p-norm-type goal functions: Application to power consumption scheduling," *J. Franklin Inst.*, vol. 361, no. 10, p. 106926, 2024.

[3] N. Shlezinger, Y. C. Eldar, and M. R. D. Rodrigues, "Hardware-limited task-based quantization," *IEEE Trans. Signal Process.*, vol. 67, no. 20, pp. 5223–5238, 2019.

[4] N. Tishby, F. C. Pereira, and W. Bialek, "The information bottleneck method," in *Proc. 37th Annu. Allerton Conf. Commun. Control Comput.*, 1999, pp. 368–377.

[5] D. Gündüz, Z. Qin, I. E. Aguerri, H. S. Dhillon, Z. Yang, A. Yener, K. K. Wong, and C.-B. Chae, "Beyond transmitting bits: Context, semantics, and task-oriented communications," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 1, pp. 5–30, 2023.

[6] A. Mas-Colell, M. D. Whinston, and J. R. Green, *Microeconomic Theory*. Oxford University Press, 1995.

[7] J. Armstrong, "MISES: Mechanism-informed source encoding for structured environments," 2026, zenodo, DOI: 10.5281/zenodo.19819209.

[8] G. Chechik, A. Globerson, N. Tishby, and Y. Weiss, "Information bottleneck for Gaussian variables," *J. Mach. Learn. Res.*, vol. 6, pp. 165–188, 2005.