INNOCOMM: LEARNING DECISION-RELEVANT LATENT MESSAGES FOR MULTI-AGENT LANGUAGE MODELS

Anonymous authors
Paper under double-blind review

ABSTRACT

When evidence is distributed across large language model agents, a communication interface must preserve the information needed for receiver decisions under a real byte budget. We introduce INNOCOMM, a sender-local latent codec that maps locally contextualized evidence into 128-dimensional slots, serializes them as INT8 packets, and delivers them to a set-consistent language-model receiver. The codec is trained to improve receiver decisions through task supervision, frozen-base behavior matching, and packet–receiver co-adaptation rather than full activation reconstruction. On two annotation-derived candidate-decision tasks, the uncompressed interface exhibits stable content utility across three receiver seeds on Development, with a separate component confirmation on a held-out FEVEROUS population. In single-receiver-seed Development experiments, a fixed-reader four-slot packet compresses the full FP16 hidden-bank payload by approximately $1{,}054\times$ to a mean of 1,428.872 application-layer bytes per request while lowering NLL by 0.044406 relative to the full base. Under a tighter 146-byte-per-sender budget, co-adapting the shared codec and memory adapter reduces NLL by 0.042625 (95% CI $[0.019817, 0.063856]$) relative to a matched codec-only continuation, while preserving content effects, base utility, and tested permutation consistency. A strong adapted text interface remains ahead by 0.044275 NLL at the same mean byte budget. Together, these results show that co-adapting the shared packet and receiver interface improves compact latent communication under an actual wire budget.

1 INTRODUCTION

When language-model agents hold different pieces of evidence, communication must turn their local observations into information that a receiver can use. We study a *sender-local* setting in which each agent encodes only the public query and its own report, while a language-model receiver uses the resulting message multiset to choose among public candidates.

Compactness alone does not ensure a useful message. Quantization and a short latent representation change the memory presented to the receiver, coupling the problems of retaining evidence and learning how to read it. INNOCOMM addresses both by learning messages against receiver task risk and co-adapting their receiver-side interface, with budget-adapted text and candidate scores as communication references.

We introduce INNOCOMM, a decision-oriented latent communication interface. For each sender, a frozen contextual encoder processes the public query and one local report. A learned packetizer maps the resulting variable-length hidden states to a small number of 128-dimensional slots. The slots pass through a real INT8 serialize–deserialize round trip, deterministic sender-set canonicalization, and a shared decoder before entering an LLM receiver through an existing memory adapter. Training optimizes the receiver's candidate distribution with task supervision and frozen-base behavior matching; it does not require reconstructing the sender's full activation bank.

To adapt the readout to compressed messages, we compare fixed-reader training with packet–receiver co-adaptation. The fixed-reader codec learns to use the existing memory adapter, whereas the aligned model jointly updates the packetizer, shared decoder, and a copy of that adapter while keeping

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

the sender encoder and 7B LLM backbone frozen. A matched codec-only continuation uses the same initialization, data, objective, training budget, selection rule, and wire format, isolating the contribution of receiver adaptation.

Across two annotation-derived decision tasks, the uncompressed contextual interface shows report-specific utility on development data across three receiver seeds, with a separate FEVEROUS component confirmation. In subsequent single-receiver-seed experiments, the fixed-reader four-slot codec reduces the full FP16 hidden-bank payload by approximately $1{,}054\times$ while lowering NLL. At 146 bytes per sender, co-adapting the shared codec and memory adapter improves on matched codec-only continuation.

Our contributions are:

- We introduce a sender-local latent codec that learns decision-relevant messages, transmits them through an explicit INT8 wire format, and receives them through a set-consistent language-model interface.
- We co-adapt the codec and the receiver's existing memory adapter without updating the language-model backbone or increasing message bytes, and establish a direct gain over matched shared-codec continuation on development data.
- We place latent packets, budget-adapted text, candidate scores, and activation interfaces on a common risk–byte frontier using content controls and realized application-layer byte budgets.

2 Related Work

Learned communication and latent LLM collaboration. Continuous messages have long been learned from downstream objectives in multi-agent systems. DIAL propagates task gradients through a noisy channel, while CommNet jointly learns continuous communication and cooperative policies (Foerster et al., 2016; Sukhbaatar et al., 2016). Recent LLM-agent work uses internal representations to bypass repeated language generation. CIPHER communicates an expectation over output embeddings (Pham et al., 2024), and LatentMAS transfers last-layer latent thoughts through shared working memory (Zou et al., 2026). Other work transforms KV caches across heterogeneous agents (Chen et al., 2026), or analyzes a frozen latent channel with post-hoc sparse coding (Wu & Zhu, 2026). InnoComm focuses on a complementary question: how to learn a finite-capacity packet from each sender's local evidence, execute its actual wire representation, and measure the receiver decision induced by its content.

Compressing context, memory, and KV representations. Latent bottlenecks can turn large inputs into a small set of representations. Perceiver uses asymmetric cross-attention to map high-dimensional inputs into a compact latent array (Jaegle et al., 2021). ICAE learns memory slots through autoencoding, language modeling, and instruction tuning objectives (Ge et al., 2024). BTR combines downstream task learning with query-aware distillation and representation recovery to construct low-precision token representations for retrieval-augmented models (Cao et al., 2024). KIVI instead quantizes the KV cache used inside a single model (Liu et al., 2024). For communication, KVComm selects model-native KV pairs according to attention importance (Shi et al., 2026), whereas KVCOMM reuses KV caches across contexts through online offset correction (Ye et al., 2025). InnoComm learns new fixed-width messages from contextual sender states and jointly learns their decoder and receiver-side memory interface under a fixed transmission budget.

Set-consistent reception and behavior alignment. Messages from multiple senders form a multiset. Deep Sets characterizes permutation-invariant set functions, and Set Transformer models interactions among set elements with attention (Zaheer et al., 2017; Lee et al., 2019). Our protocol realizes this constraint by canonicalizing serialized sender packets, decoding them with shared modules, and testing the final probabilities under independent sender and candidate permutations. The training objective also uses a frozen base distribution in the spirit of knowledge distillation (Hinton et al., 2015); the receiver-alignment contribution is identified separately through matched training arms.

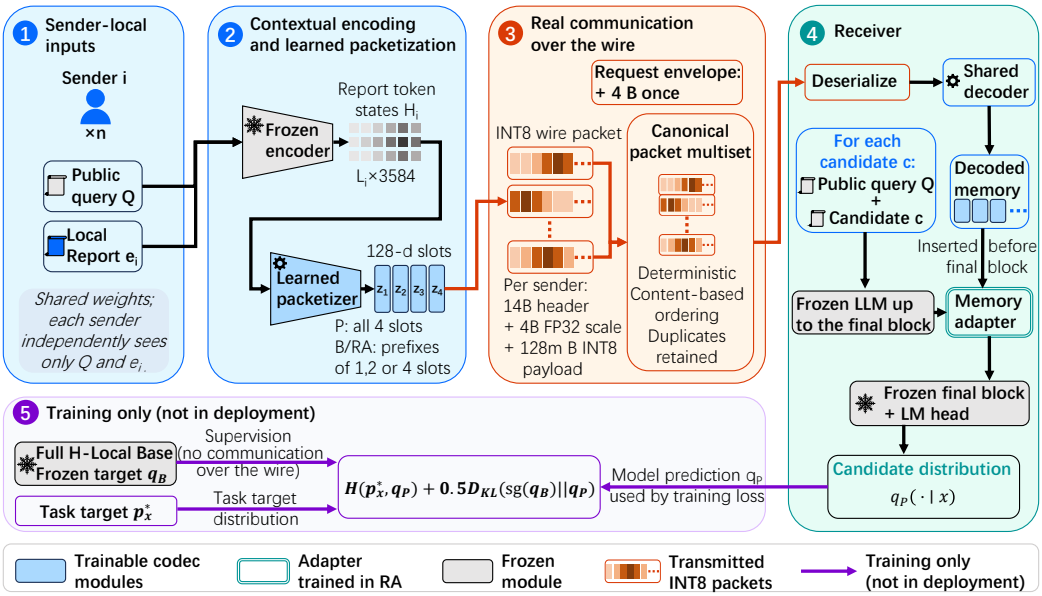


Figure 1: INNOCOMM encodes each sender's local query–report states as an INT8 packet, then canonically receives the packet multiset through a frozen language-model backbone. Receiver-aligned training also adapts the memory interface; the frozen full H-LOCAL base supplies the distributional target q_B during training and is absent at deployment.

Task-oriented multirate coding. Task-oriented communication optimizes remote inference rather than input reconstruction. ARTOVeQ, for example, uses nested codebooks and progressive learning across link rates (Fishel et al., 2025). The successive packetizer in INNOCOMM supports one checkpoint at one, two, or four slots. Each prefix is independently requantized and transmitted in full, so the interface is multirate but not an append-only refinement bitstream. Our study centers on whether these quantized messages preserve report-specific decision value and how their receiver changes their position on a measured risk–byte frontier.

3 PROBLEM SETUP

3.1 SENDER-LOCAL CANDIDATE DECISIONS

For request x, the public input contains a query Q_x and a candidate set $\mathcal{C}_x = \{c_1, \ldots, c_{K_x}\}$. Sender i holds only its local report $e_{x,i}$ and produces the transmitted latent message

$$w_{x,i} = f_\phi(Q_x, e_{x,i}). \tag{1}$$

The receiver observes the public input and the message multiset and returns

$$q_P(\cdot \mid x) = g_{\theta,P}(Q_x, \mathcal{C}_x, \{\{w_{x,i}\}\}_{i=1}^{n_x}), \tag{2}$$

where protocol P specifies message construction, wire representation, and reception. Models and protocol are shared in advance, but a sender cannot read other reports, the full report coalition, gold supports, source identifiers, or evaluation targets. Candidate semantics are public, although the implemented sender encoder reads only Q_x and $e_{x,i}$; candidate text enters when the receiver scores each option.

The task construction supplies the target distribution p_x^*. Messages preserve multiplicity, and their arrival order must not affect the output. Synchronously permuting candidates and candidate indices should induce the same permutation of the output distribution. We test both properties with separate original and transformed forward passes.

3.2 RECEIVED RISK AND WIRE COST

Our primary metric is equal-task/source negative log-likelihood. Let $\mathcal{S}_t$ be the sources for task t, and let $\ell_{s,t}(P)$ aggregate the pre-specified row weights within source s. Then

$$R(P) = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \frac{1}{|\mathcal{S}_t|} \sum_{s \in \mathcal{S}_t} \ell_{s,t}(P). \tag{3}$$

We write $\Delta(A, P) = R(A) - R(P)$, so a positive value favors P. Confidence intervals are computed by paired resampling of sources.

Let N_{fixed} be the trained, frozen, query-conditioned N-MEM receiver, B_{full} the full H-LOCAL base, and T_{full} the corresponding full-text receiver. We measure preservation of the base's utility relative to N_{fixed} using

$$C_{90,B}(P) = \big[R(N_{\text{fixed}}) - R(P)\big] - 0.9\big[R(N_{\text{fixed}}) - R(B_{\text{full}})\big]. \tag{4}$$

$C_{90,\text{Text}}$ replaces B_{full} with T_{full}. The mean-byte-matched text mixture T_{mix} is used only for the communication comparison.

Communication cost is the realized application-layer payload

$$W(P; x) = b_{\text{request}} + \sum_i (b_{\text{header},i} + b_{\text{payload},i} + b_{\text{metadata},i}) + b_{\text{control}}. \tag{5}$$

We first average request bytes within each task and then average FEVEROUS and MuSiQue equally. An average-budget mixture may randomly choose a complete endpoint independently of content and labels, but its mode byte is charged. A per-sender hard cap instead constrains every message in every request.

We test report-specific utility, improvement over the full-bank receiver, and risk at matched communication cost. The direct alignment estimand is

$$\Delta_{\text{align}} = R(P_{\text{codec-only}}) - R(P_{\text{aligned}}), \tag{6}$$

for arms matched in initialization, data, objective, training budget, selection, and wire bytes.

4 METHOD

4.1 OVERVIEW

Figure 1 summarizes the sender-local codec and its receiver interface. The full H-LOCAL base provides a frozen training target; deployment uses only the transmitted packets and the receiver's public inputs.

4.2 SENDER-LOCAL SLOT ENCODING

The frozen encoder $E_{\bar{\theta}}$ independently processes each query–report pair and retains the final-layer states of report tokens,

$$H_{x,i} = E_{\bar{\theta}}(Q_x, e_{x,i}) \in \mathbb{R}^{L_{x,i} \times 3584}. \tag{7}$$

Before any sender exchange, a LayerNorm and bias-free projection map each state to the slot dimension:

$$X_{x,i} = \text{LN}(H_{x,i}) W_{\text{in}}^{\top} \in \mathbb{R}^{L_{x,i} \times d}, \quad W_{\text{in}} \in \mathbb{R}^{d \times 3584},\ d = 128. \tag{8}$$

The fixed-rate INNOCOMM-P packetizer uses four learned queries U. With four-head cross-attention A and a residual feed-forward network,

$$Z_{x,i} = A(U, X_{x,i}) + \text{FFN}(A(U, X_{x,i})), \qquad Z_{x,i} \in \mathbb{R}^{4 \times 128}. \tag{9}$$

The multirate INNOCOMM-B/RA packetizer generates successive slots, with each z_j written as a column vector. Query j conditions on the mean of earlier slots,

$$\bar{z}_{<j} = \begin{cases} 0, & j = 1, \\ \frac{1}{j-1} \sum_{k<j} z_k, & j > 1, \end{cases} \qquad Q_j = q_j + W_{\text{prefix}} \bar{z}_{<j}. \tag{10}$$

One checkpoint can therefore run with $m \in \{1, 2, 4\}$ prefix slots. A shared linear decoder maps every slot into one 3,584-dimensional memory vector. No loss reconstructs $H_{x,i}$ element by element.

4.3 EXECUTABLE PACKETS AND SET-CONSISTENT RECEPTION

Each sender quantizes its selected slots using one symmetric FP32 scale. The packet contains a 14-byte header, a 4-byte scale, and $128m$ signed INT8 payload bytes; a request adds a 4-byte envelope. The resulting costs are 146, 274, and 530 bytes per sender for one, two, and four slots, respectively. Their equal-task means are 396.512, 740.632, and 1,428.872 bytes per request. INNOCOMM-P selects an FP16 endpoint and then evaluates a real INT8 round trip; B/RA train with a straight-through estimator and always evaluate deserialized bytes.

The receiver canonically orders the serialized packets using a deterministic content-based key and retains duplicates. It decodes and concatenates the resulting memory vectors without sender identifiers, padding, source metadata, or uncharged length side channels; Appendix C specifies the exact key and wire fields. For candidate c, it forms a fixed Yes/No prompt, computes the final-position logit difference $s_c = \ell_{\texttt{Yes}} - \ell_{\texttt{No}}$, and returns

$$q_P(c \mid x) = \mathrm{softmax}_{c \in \mathcal{C}_x}(s_c). \tag{11}$$

4.4 DECISION TRAINING AND PACKET–RECEIVER ALIGNMENT

Behavior-matched training combines the candidate target with the frozen full-base distribution q_B:

$$\mathcal{L}_{\mathrm{match}} = H(p_x^*, q_P) + 0.5 D_{\mathrm{KL}}\left(\mathrm{sg}(q_B) \| q_P\right). \tag{12}$$

Task-only arms remove the KL term. We cache the frozen base distributions and stop gradients through these targets; deployment does not call the base.

The aligned and codec-only arms start from the same previously trained B codec and H-LOCAL memory adapter. Codec-only continuation updates the packetizer and decoder; aligned training also updates a copy of the existing adapter. Both keep the language-model backbone frozen and use the same message format and byte budgets. Shared training samples one of the one-, two-, or four-slot budgets per example visit, whereas fixed-1 training uses only one slot. Each selected prefix is requantized and transmitted in full. Appendix D gives the initialization, parameter counts, and selection rules.

5 EXPERIMENTAL SETUP

5.1 TASKS, POPULATIONS, AND EVIDENCE SCOPE

We construct multi-sender candidate-decision tasks from FEVEROUS (Aly et al., 2021) and MuSiQue (Trivedi et al., 2022) annotations. FEVEROUS candidates are SUPPORTS, REFUTES, and INSUFFICIENT; MuSiQue candidates are ANSWERABLE and INSUFFICIENT. Each original claim or question defines one source. Deterministic construction rules derive evaluation instances (rows) with two to six messages, and rows from one source never cross data splits. Appendix A gives the complete annotation-to-report construction, soft targets, and within-source weights.

The task also contains relation-counterfactual instances in which the visible query, reports, and candidates are unchanged but the target differs through an unobserved relation variable. They receive one quarter of the within-source weight for affected sources; the reported NLL therefore includes this partially observable component (Appendix A).

The training set contains 768 sources per task; the source-disjoint development set contains 192 per task. A separate component-confirmation set contains 512 FEVEROUS sources. Row counts and population roles are reported in Appendix A. All primary systems use the Qwen2.5-7B-Instruct receiver lineage (Qwen Team, 2024). The uncompressed study evaluates receiver seeds 20260822, 20260823, and 20260824; subsequent packet studies use seed 20260824, selected by a pre-specified median rule. Thus the final packet results are single-receiver-seed development evidence, while the earlier H-LOCAL component has three-seed development and a separate FEVEROUS confirmation.

5.2 TRAINING AND CONTROLS

Learned modules train on the training set and select checkpoints on the development set under a common AdamW schedule and epoch budget. LLM forward passes use bfloat16, cached banks use

Table 1: Content utility of the contextual interface. Values are paired NLL contrasts with 95% CIs. Positive values favor H-LOCAL in the first three rows; positive H-LOCAL−T-MEM values favor text. IC is the independently held FEVEROUS component confirmation.

Contrast	3-seed Development	FEVEROUS IC
N-MEM − H-LOCAL	**0.045836 [0.020387, 0.070470]**	**0.122030 [0.082980, 0.160038]**
Replacement − H-LOCAL	**0.133957 [0.100647, 0.165682]**	**0.188653 [0.139954, 0.237256]**
E-Local − H-LOCAL	**0.041891 [0.019417, 0.063558]**	**0.115697 [0.078790, 0.151875]**
H-LOCAL − T-MEM	0.139852 [0.090568, 0.184275]	0.203214 [0.149241, 0.255122]
$C_{90,\text{Text}}$	−0.121283 [−0.161463, −0.075931]	−0.170690 [−0.218288, −0.122371]

FP16, and wire packets use INT8. Training and evaluation run on 48 GiB NVIDIA RTX A6000 GPUs. Appendix D gives the optimization and selection details.

The comparisons span uncompressed hidden-state and text receivers, low-byte text and candidate-score references, and fixed-reader, codec-only, and receiver-aligned packets. Report-specificity is tested with fixed no-evidence states, same-codec null input, natural report replacement, and—for B/RA—a separately optimized no-evidence codec. Null RA arms share the B codec and H-LOCAL adapter initialization of the real-message arms before optimization on N-MEM. Natural replacements draw full reports from another source while matching sender count and report-length quartile without reading the current label; Appendix E gives the complete stage-specific controls.

5.3 METRICS, UNCERTAINTY, AND BYTE MATCHING

The primary estimand is the paired source NLL difference from Eq. 3. Sources are equally weighted within tasks and the two tasks are equally weighted. We use 5,000 paired source-cluster percentile bootstrap replicates; three-seed effects first average paired model differences for the same source. Packet-stage intervals condition on the trained and selected endpoint and quantify source-resampling uncertainty. For alignment, we report the pooled paired interval together with task-specific point effects. Accuracy and Brier score are secondary; source-weighted Brier from P and raw-row Brier from B/RA are not compared as a trend.

All communication baselines execute their actual representation path. Text uses UTF-8 payload caps of 132, 260, and 516 bytes; with the common 14-byte header, these match the packet's total 146-, 274-, and 530-byte sender budgets after accounting for its 4-byte scale. Candidate scores transmit one FP16 value per candidate per sender. The average-budget text mixture chooses a full endpoint with public, content-independent randomness and pays a one-byte control field; it does not average predictions or satisfy a per-request hard cap. Exact reference construction appears in Appendix C.

6 MAIN RESULTS

6.1 SENDER-LOCAL CONTEXTUAL STATES CARRY DECISION-RELEVANT CONTENT

The uncompressed H-LOCAL interface consistently benefits from the correct local reports. Across three receiver seeds, contextual report states outperform the no-evidence, natural-replacement, and embedding-only controls (Table 1).

The seed-20260824 endpoint also shows these content and contextualization effects on 512 separate FEVEROUS sources. Both candidate and sender reversal produce zero maximum probability change and perfect argmax agreement on Development and confirmation data.

6.2 A FOUR-SLOT INT8 PACKET PRESERVES AND IMPROVES FULL-BASE UTILITY

The fixed-reader INNOCOMM-P compresses each hidden bank into four 128-dimensional slots. Table 2 shows that the real INT8 packet uses a mean of 1,428.872 bytes per request, compared with 1,506,300.164 for the full FP16 bank: a $1{,}054\times$ reduction. Its mean wire cost is also about $151\times$ lower than that of the FP16 readout K/V reference.

Table 2: Development risk and application-layer wire cost for the four-slot fixed-reader packet (receiver seed 20260824). Brier is source-weighted.

Endpoint	Mean bytes/request ↓	NLL ↓	Brier ↓
Full H-LOCAL FP16	1,506,300.164	0.788889	0.501977
Readout K/V FP16	215,221.428	0.789261	0.501885
INNOCOMM-P FP16	2,794.598	**0.743531**	**0.469562**
INNOCOMM-P INT8	**1,428.872**	0.744483	0.470272

Table 3: Packet–receiver co-adaptation at an identical one-slot wire budget. The paired effect is codec-only NLL minus aligned NLL; positive values favor alignment.

Endpoint	Adapter	FEVEROUS	MuSiQue	Pooled	Paired effect [95% CI]
Codec-only shared	fixed	0.907535	0.570583	0.739059	reference
INNOCOMM-RA shared	**train**	**0.894500**	**0.498370**	**0.696435**	**0.042625 [0.019817, 0.063856]**
Codec-only fixed-1	fixed	0.917467	0.557777	0.737622	reference
INNOCOMM-RA fixed-1	train	0.918610	0.502473	0.710541	0.027081 $[-0.002876, 0.055971]$

Relative to full H-LOCAL, the INT8 packet improves FEVEROUS and MuSiQue NLL by 0.064145 and 0.024668, for a pooled improvement of 0.044406 $[0.012625, 0.075203]$. Its $C_{90,B}$ is 0.047665 $[0.017897, 0.077661]$, retaining at least 90% of the base's measured utility relative to the fixed no-evidence receiver. The FP16-to-INT8 NLL point change is 0.000952 while bytes fall by another 48.87%. Neither uniform nor segment projections retain the learned codec's report-specific utility; Appendix F reports the complete matrix.

6.3 PACKET–RECEIVER CO-ADAPTATION YIELDS A MATCHED LOW-BUDGET GAIN

We next compare aligned and codec-only continuations under the same 146-byte per-sender wire contract. Table 3 isolates whether the memory adapter is trainable. Shared alignment reduces pooled NLL from 0.739059 to 0.696435, a paired improvement of 0.042625 $[0.019817, 0.063856]$. The FEVEROUS and MuSiQue point effects are 0.013036 and 0.072214.

The corresponding fixed-1 interval crosses zero, so the direct alignment result is established by the shared construction. Aligned training adds 3,674,368 adapter parameters but no transmitted bytes and does not update the 7B backbone. The selected aligned packet retains report-specific utility, base utility, and the tested permutation behavior (Section 7).

6.4 COMMUNICATION PERFORMANCE AT REALIZED BYTE COST

Figure 2 places the endpoints in one risk–byte coordinate system. At 396.512 mean bytes per request, the aligned one-slot packet reaches NLL 0.696435, while the mean-byte-matched text reference reaches 0.652159, an advantage of 0.044275 $[0.015197, 0.074273]$. Head-truncated text also satisfies the 146-byte-per-sender cap, using 366.904 mean bytes at NLL 0.653758; receiver adaptation lowers this endpoint from 0.825488. Complete values appear in Appendix G.

Among the evaluated latent endpoints, aligned one-slot improves on unaligned four-slot NLL (0.696435 versus 0.727777) with 72% fewer bytes.

7 ANALYSIS AND ABLATIONS

7.1 THE ALIGNED PACKET REMAINS REPORT-SPECIFIC

The aligned receiver continues to benefit from the correct reports. Figure 3 distinguishes a separately optimized no-evidence codec, an input-null intervention on the same codec, the frozen N-MEM reference, and natural report replacement. All four paired intervals favor the correct packet, and the full-base comparison remains positive. The improvement in received risk is therefore accompanied by report-specific utility rather than sensitivity to arbitrary memory alone.

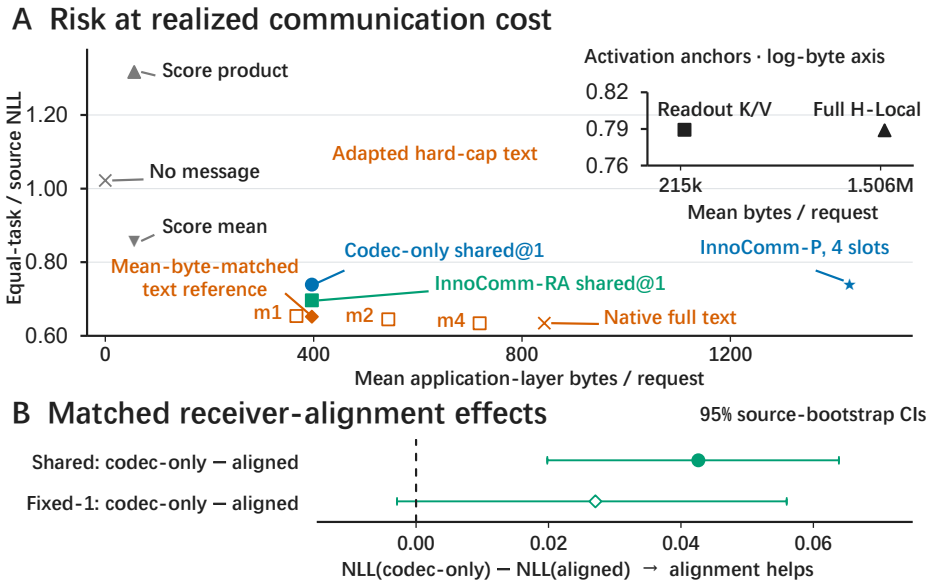


Figure 2: Communication risk and receiver alignment on single-receiver-seed development data. (A) Equal-task/source NLL versus mean application-layer bytes; open squares are hard-cap text, the diamond is the mean-byte-matched mixture including its one-byte control, and activation anchors use a log-byte inset. (B) Codec-only-minus-aligned NLL with 95% source-bootstrap intervals; positive values favor alignment.

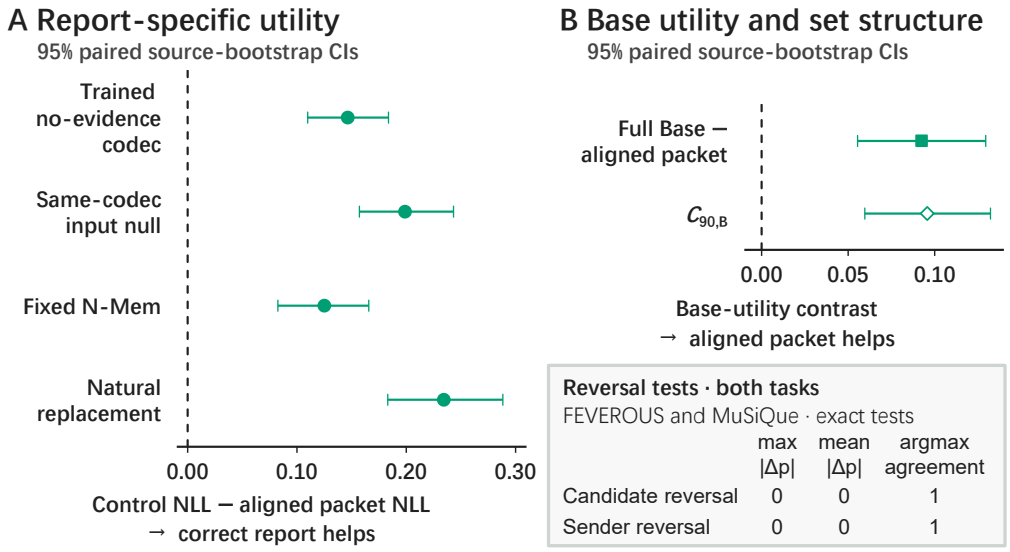


Figure 3: Controls for the one-slot aligned shared codec on single-receiver-seed development data. (A) Control-minus-packet NLL; positive values favor the correct packet. (B) Base-utility and exact reversal tests on both tasks. Error bars are 95% source-bootstrap intervals.

Independent sender and candidate reversal tests return zero maximum and mean probability change with perfect argmax agreement for the evaluated H-LOCAL, P, B@1, and RA@1 endpoints. The tests use independent forward passes for each interface. For packet endpoints, those passes include actual INT8 serialization, deserialization, and canonical reception.

7.2 MULTIRATE, TRAINING-SCOPE, AND QUANTIZATION ABLATIONS

Shared checkpoints are selected by mean development NLL across all three budgets; among aligned candidates that retain report-specific utility, one-slot NLL selects the representative arm. The paired shared-versus-fixed intervals nevertheless cross zero (Appendix F, Figure 4).

Alignment changes the evaluated rate curve. Codec-only risk decreases modestly as more slots are added, whereas the aligned shared codec reaches its lowest risk at one slot and does not improve with additional slots.

Frozen-base matching lowers pooled point risk in the tested P and B configurations, but the pre-alignment one-slot paired interval crosses zero and the task-level effects differ.

Appendix F reports the complete intervals, rate curves, and quantization results.

8 LIMITATIONS

The experiments evaluate FEVEROUS- and MuSiQue-derived candidate-decision tasks with one Qwen2.5-7B receiver lineage. The uncompressed interface is evaluated across three receiver seeds and a separate FEVEROUS population, whereas the packet and alignment results currently use one receiver seed. Evaluation across additional receiver lineages and populations remains open.

9 CONCLUSION

INNOCOMM learns sender-local INT8 messages for a set-consistent language-model receiver. The fixed-reader codec preserves report-specific utility under substantial hidden-bank compression, and co-adapting the shared codec with the existing memory adapter improves NLL over matched codec-only continuation without increasing transmitted bytes or updating the backbone. These results support receiver co-adaptation as an effective training step for compact latent messages.

AI USE STATEMENT

Generative AI tools were used for grammar and spelling checks, language polishing, related-work retrieval, and programming assistance for a subset of the experiments. All AI-assisted work was reviewed by the authors, who take responsibility for the final content of this work.

REPRODUCIBILITY STATEMENT

Appendices A–D specify the task construction, statistical estimands, model interface, serialization rules, and training and selection procedures. Appendices E–I report the controls, ablations, communication comparisons, and measured cost scope. The anonymous supplementary artifact contains trainable packet states, checkpoint, data, and environment manifests, anonymized aggregate data underlying Figures 2–5, figure-to-source mappings, and entry points for validating the released aggregate quantities.

REFERENCES

Rami Aly, Zhijiang Guo, Michael Sejr Schlichtkrull, James Thorne, Andreas Vlachos, Christos Christodoulopoulos, Oana Cocarascu, and Arpit Mittal. FEVEROUS: Fact extraction and verification over unstructured and structured information. In *Advances in Neural Information Processing Systems Datasets and Benchmarks Track*, 2021. URL https://arxiv.org/abs/2106.05707.

Qingqing Cao, Sewon Min, Yizhong Wang, and Hannaneh Hajishirzi. BTR: Binary token representations for efficient retrieval augmented language models. In *International Conference on Learning Representations*, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/hash/c0f721d329c1a10546869c783e866fb7-Abstract-Conference.html.

Siyi Chen, Xiaoyan Zhang, Meng Wu, Jonathan Tremblay, Valts Blukis, Stan Birchfield, Rene Vidal, Alvaro Velasquez, Sijia Liu, and Qing Qu. See what I see, know what I think: Dense latent communication across heterogeneous agents. *arXiv preprint arXiv:2606.13594*, 2026. URL https://arxiv.org/abs/2606.13594.

Eyal Fishel, May Malka, Shai Ginzach, and Nir Shlezinger. Remote inference over dynamic links via adaptive rate deep task-oriented vector quantization. *arXiv preprint arXiv:2501.02521*, 2025. URL https://arxiv.org/abs/2501.02521.

Jakob N. Foerster, Yannis M. Assael, Nando de Freitas, and Shimon Whiteson. Learning to communicate with deep multi-agent reinforcement learning. In *Advances in Neural Information Processing Systems*, 2016. URL https://arxiv.org/abs/1605.06676.

Tao Ge, Jing Hu, Lei Wang, Xun Wang, Si-Qing Chen, and Furu Wei. In-context autoencoder for context compression in a large language model. In *International Conference on Learning Representations*, 2024. URL https://arxiv.org/abs/2307.06945.

Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. URL https://arxiv.org/abs/1503.02531.

Andrew Jaegle, Felix Gimeno, Andrew Brock, Oriol Vinyals, Andrew Zisserman, and Joao Carreira. Perceiver: General perception with iterative attention. In *International Conference on Machine Learning*, 2021. URL https://arxiv.org/abs/2103.03206.

Juho Lee, Yoonho Lee, Jungtaek Kim, Adam Kosiorek, Seungjin Choi, and Yee Whye Teh. Set transformer: A framework for attention-based permutation-invariant neural networks. In *International Conference on Machine Learning*, 2019.

Zirui Liu, Jiayi Yuan, Hongye Jin, Shaochen Zhong, Zhaozhuo Xu, Vladimir Braverman, Beidi Chen, and Xia Hu. KIVI: A tuning-free asymmetric 2bit quantization for KV cache. In *International Conference on Machine Learning*, 2024. URL https://arxiv.org/abs/2402.02750.

Chau Pham, Boyi Liu, Yingxiang Yang, Zhengyu Chen, Tianyi Liu, Jianbo Yuan, Bryan A. Plummer, Zhaoran Wang, and Hongxia Yang. Let models speak ciphers: Multiagent debate through embeddings. In *International Conference on Learning Representations*, 2024. URL https://arxiv.org/abs/2310.06272.

Qwen Team. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*, 2024. URL https://arxiv.org/abs/2412.15115.

Xiangyu Shi, Marco Chiesa, Gerald Q. Maguire Jr., and Dejan Kostić. KVComm: Enabling efficient LLM communication through selective KV sharing. In *International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=F7rUng23nw.

Sainbayar Sukhbaatar, Arthur Szlam, and Rob Fergus. Learning multiagent communication with backpropagation. In *Advances in Neural Information Processing Systems*, 2016. URL https://arxiv.org/abs/1605.07736.

Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. MuSiQue: Multihop questions via single-hop question composition. *Transactions of the Association for Computational Linguistics*, 10:539–554, 2022. doi: 10.1162/tacl_a_00475. URL https://aclanthology.org/2022.tacl-1.31/.

Di Wu and Xiaohui Zhu. Post-hoc sparse coding of latent communication between vision-language model agents. *arXiv preprint arXiv:2608.10198*, 2026. URL https://arxiv.org/abs/2608.10198.

Hancheng Ye, Zhengqi Gao, Mingyuan Ma, Qinsi Wang, Yuzhe Fu, Ming-Yu Chung, Yueqian Lin, Zhijian Liu, Jianyi Zhang, Danyang Zhuo, and Yiran Chen. KVCOMM: Online cross-context KV-cache communication for efficient LLM-based multi-agent systems. *arXiv preprint arXiv:2510.12872*, 2025. URL https://arxiv.org/abs/2510.12872.

Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabás Póczos, Ruslan Salakhutdinov, and Alexander J. Smola. Deep sets. In *Advances in Neural Information Processing Systems*, 2017.

Jiaru Zou, Ruizhong Qiu, Gaotang Li, Xiyuan Yang, Katherine Tieu, Pan Lu, Ke Shen, Hanghang Tong, Yejin Choi, Jingrui He, James Zou, Mengdi Wang, and Ling Yang. Latent collaboration in multi-agent systems. In *International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=syG9I9ofd8.

A Data Construction and Statistical Estimation

Unless stated otherwise, packet, budget, and alignment results use Qwen2.5-7B, receiver seed 20260824, and the development set. The three-seed H-Local experiment and its separate FEVEROUS confirmation are explicitly identified. Every contrast is written as left-system risk minus right-system risk, so a positive value favors the system on the right.

A.1 Annotation-derived candidate decisions

The experiments construct candidate-decision tasks from annotations rather than reporting the original open-benchmark metrics. Each original claim or question defines a source; each derived evaluation instance specifies one message multiset and target from that source.

Table 4: Frozen construction of the two candidate-decision tasks.

Step	FEVEROUS	MuSiQue
Original source	A non-insufficient claim and one valid evidence set	A question with two to four supporting paragraphs aligned to its decomposition
Candidate semantics	Supports, Refutes, Insufficient	Answerable, Insufficient
Sender report	One report per evidence element, preserving page context and joining the formal evidence with an [EVIDENCE] delimiter	One report per decomposition hop, formatted as title [EVIDENCE] paragraph
Natural complete row	All required reports; target is the original label	All decomposition hops; target is Answerable
Incomplete natural rows	For multi-report sources, remove the last report or retain only the first, targeting Insufficient; for a one-report source, duplicate the report and retain its label	The same completion-family rule, plus an atomic family built from the first hop
Constructor-only relation counterfactual	Keep report text and sender multiplicity fixed; collapse the unobserved role_id to one role and target Insufficient	Same
Evaluation rows	Canonical candidate order and two to six messages	Same

The evaluated sender and receiver do not observe role_id. A natural-complete row and its role-collapse counterpart therefore expose the same query, report multiset, and candidates while assigning different targets. These rows model uncertainty from an unobserved relation variable and introduce irreducible target ambiguity under the sender-local contract; they receive 0.25 total within-source weight. We include them identically for all systems and do not treat their performance as evidence that the packet identifies role changes.

Discrete labels use fixed smoothing with $\epsilon = 10^{-4}$:

$$p_x^*(c) = \begin{cases} 1-\epsilon, & c = y_x, \\ \epsilon/(K_x - 1), & c \neq y_x. \end{cases} \tag{13}$$

Candidate-reversed rows are used only for equivariance tests. Reports are deterministically resolved from the original annotations and are not regenerated from model predictions. The sender reads only the public query and its own report. Candidate text is used by receiver-side scoring, and training additionally accesses the frozen target and base distribution.

A.2 SOURCE WEIGHTS AND BOOTSTRAP UNCERTAINTY

For derived evaluation instances $\mathcal{X}_s$ belonging to source s,

$$\ell_s(P) = \sum_{x \in \mathcal{X}_s} a_{x|s} \left[- \sum_{c \in \mathcal{C}_x} p_x^*(c) \log \max(q_P(c \mid x), 10^{-6}) \right]. \tag{14}$$

If a source has relation-counterfactual rows, they receive 0.25 total weight and natural rows receive 0.75; each group divides its weight equally. Without a counterfactual, all rows are equally weighted. Sources are equal within a task, and FEVEROUS and MuSiQue are equal in pooled risk. Three-seed H-LOCAL contrasts first average the three paired model differences for each source.

Intervals use 5,000 paired source-cluster bootstrap draws. P uses source-weighted Brier; B and RA use a task-equal raw-row Brier score. The separate FEVEROUS confirmation is used only for the H-LOCAL component and not for P/B/RA.

Table 5: Experimental populations, shown as sources / derived instances.

Population	FEVEROUS	MuSiQue	Role
Training	768 / 1,762	768 / 2,545	training
Development	192 / 444	192 / 626	selection and evaluation
FEVEROUS confirmation	512 / 1,139	—	H-LOCAL component confirmation

B MODEL AND IMPLEMENTATION DETAILS

B.1 BACKBONE AND CONTEXTUAL BANK

All systems use the same Qwen2.5-7B-Instruct-based receiver state, fixed before the contextual-bank and packet experiments. Its language-model backbone, input embeddings, and LM head remain frozen in packet experiments. Runs use Python 3.10.20, PyTorch 2.11.0+cu128, Transformers 4.53.3, and 48 GiB NVIDIA RTX A6000 GPUs. The receiver context limit is 8,192 tokens and the H-LOCAL bank limit is 3,072 tokens; the supplementary environment manifest records the CUDA, cuDNN, driver, and remaining package versions.

For each report, the encoder input is:

```
Encode evidence for candidate assessment.
query: {public query}
one sender report: {local report}
```

No additional special tokens are inserted. An empty report is represented by EOS. The Qwen decoder runs with an all-one attention mask and default zero-based positions; the interface keeps only the final-layer span aligned to the report body. Public-prefix and header states are excluded. Cached banks are FP16, and a row exceeding the frozen bank limit raises an error rather than being silently truncated.

H-LOCAL uses these states; E-LOCAL uses input embeddings. N-MEM contextualizes the query with "no evidence supplied," keeps at most 64 public states, repeats them to the corresponding report-token length, and restores the sender-block partition. T-MEM jointly contextualizes all reports on the receiver side. The fixed no-evidence reference is therefore a trained, query-conditioned system that can preserve report-block lengths; it is not one global constant packet.

B.2 PACKETIZER, MEMORY ADAPTER, AND SCORING

The P packetizer has four learned queries, four attention heads, and an FFN of width $2d$ for $d \in \{128, 256\}$. Including the shared decoder, the P codecs have 1,053,568 and 2,365,696 trainable parameters, respectively. The B/RA successive 128-dimensional codec has 1,069,952 parameters. The memory adapter is inserted before Qwen's final Transformer block. It layer-normalizes hidden states and memory, applies width-256 four-head cross-attention, and returns a residual update; it has

3,674,368 parameters. The complete aligned system trains 4,744,320 parameters while freezing the LLM.

Each candidate receives an `Answer Yes or No only` prompt. The leading space strings `" Yes"` and `" No"` must tokenize as distinct single tokens. Candidate reversal maps output probabilities back to the original semantic labels before comparison.

C Serialization, Bytes, and Communication References

C.1 Exact binary contract

Each sender packet begins with a 14-byte big-endian header with layout `>4sBBHHI`: four-byte magic `IKP1`, one-byte version, one-byte dtype code, two-byte slot rows, two-byte dimension, and four-byte payload length. A four-byte big-endian FP32 scale and row-major signed INT8 payload follow. The request adds a four-byte envelope. For sender slots Z,

$$s = \begin{cases} \max|Z|/127, & \max|Z| > 0, \\ 1, & \text{otherwise}, \end{cases} \qquad q = \mathrm{clip}(\mathrm{round}(Z/s), -127, 127), \tag{15}$$

and decoding returns $\hat{Z} = sq$.

The receiver sorts serialized sender blocks by `(SHA256 digest, raw bytes)` and retains duplicate blocks. The digest provides a deterministic content key; the raw bytes break digest ties without introducing sender metadata.

Table 6: Exact INT8 packet cost.

Slots	Header	Scale	Payload	Bytes/sender
1	14	4	128	146
2	14	4	256	274
4	14	4	512	530

Task	1 slot	2 slots	4 slots
FEVEROUS	466.005	871.050	1,681.140
MuSiQue	327.019	610.214	1,176.604
Equal-task	**396.512**	**740.632**	**1,428.872**

For n senders and m slots, $b_i(m) = 14 + 4 + 128m$ and $B_{\text{request}} = 4 + \sum_{i=1}^{n} b_i(m)$. Text uses the same 14-byte sender header without a scale. Its 132/260/516-byte payload caps become 146/274/530 bytes per sender after the common header, matching the corresponding packet budgets after accounting for the packet scale. Truncation discards an incomplete final UTF-8 character. The pre-specified mean-byte-matched text reference chooses head m1 with probability 0.9186871599 and head m4 with probability 0.0813128401, plus a one-byte mode field, for 396.511837 expected bytes per request.

C.2 Text and candidate-score construction

For candidate-score communication, sender i evaluates every candidate with the same Yes/No score used by the receiver and computes $r_i(c) = \mathrm{softmax}_{c \in \mathcal{C}_x}(s_{i,c})$. The mean baseline averages and renormalizes sender distributions. The product baseline computes

$$r_{\text{prod}}(c) = \mathrm{softmax}_{c \in \mathcal{C}_x}\left(\sum_i \log \max(r_i(c), 10^{-6})\right). \tag{16}$$

FP16 scores cost $4 + n(14 + 2K_x)$ bytes per request.

Text `head` uses the UTF-8 prefix. `query-overlap` extracts lowercase `[a-z0-9]+` terms, splits sentences after `.?!`, ranks sentences by unique term overlap with the query, preserves original order for ties, and packs them under the byte cap with spaces charged. Selection is sender-local; the receiver jointly contextualizes received text. Readout K/V is an internal activation-interface reference in the same receiver lineage, not a reproduction of external KVComm systems.

D TRAINING AND SELECTION DETAILS

All learned arms train for eight epochs with AdamW, learning rate 10^{-4}, weight decay 0.01, gradient clipping 1.0, 5% linear warm-up, effective batch size 16, and gradient accumulation 16. Shared arms run one prefix per example visit, sample the three budgets uniformly, and select on equal-task/equal-budget NLL. Fixed arms train and select only at their specified budget. Exact ties choose the earliest epoch. Null arms replace report inputs with N-MEM while retaining the same procedure and behavior target.

B codecs are independently initialized and use the slot dimension and interface structure selected in P; they do not continue a P checkpoint. Every RA packet arm starts from B `successive_match`, seed 20260824, epoch 7, and the H-LOCAL adapter. Shared checkpoints are selected by the equally weighted mean NLL over the one-, two-, and four-slot budgets; fixed arms use their specified budget. Among aligned candidates that satisfy the content criteria, the representative arm is selected by one-slot NLL. The B adapted-text receiver selects epoch 4 by query-overlap m1 NLL (0.664098); the RA continued-text arm selects epoch 0 with the same value and reproduces its adapted warm start.

Table 7: Training and selection details for the principal learned systems.

Stage/arm	Trainable scope	Params	Epoch	Selection NLL
P slot128-task	codec, fixed reader	1,053,568	6	0.768853
P slot128-match	codec, fixed reader	1,053,568	6	0.743531
P slot256-task	codec, fixed reader	2,365,696	7	0.753771
P slot256-match	codec, fixed reader	2,365,696	4	0.746119
B fixed1-match	fixed-1 codec	1,069,952	6	0.743223
B fixed2-match	fixed-2 codec	1,069,952	2	0.743325
B successive-match	shared codec	1,069,952	7	0.732464
B successive-task	shared codec	1,069,952	6	0.756604
B successive-null	shared N-MEM codec	1,069,952	2	0.851061
RA codec-only shared	codec, adapter fixed	1,069,952	0	0.732464
RA codec-only fixed1	codec, adapter fixed	1,069,952	6	0.737622
RA shared	codec and adapter	4,744,320	7	0.697771
RA fixed1	codec and adapter	4,744,320	8	0.710541
RA null shared	N-MEM codec and adapter	4,744,320	2	0.852516
RA null fixed1	N-MEM codec and adapter	4,744,320	4	0.852486
B adapted text	text memory adapter	3,674,368	4	0.664098
RA continued text	text memory adapter	3,674,368	0	0.664098

E COMPLETE CONTENT AND STRUCTURE CONTROLS

E.1 FIXED-READER FOUR-SLOT PACKET

Table 8: Complete content and utility contrasts for the selected four-slot INNOCOMM-P INT8 packet.

Contrast	FEVEROUS	MuSiQue	Pooled	95% CI
Same-codec N-input − P	0.276887	0.155663	0.216275	[0.160866, 0.282498]
Fixed N − P	0.100486	0.053502	0.076994	[0.043608, 0.109887]
No message − P	0.342703	0.212579	0.277641	[0.237211, 0.318787]
Natural replacement − P	0.323279	0.062387	0.192833	[0.136537, 0.253053]
P − base	−0.064145	−0.024668	−0.044406	[−0.075203, −0.012625]
P − native text	0.127469	0.092680	0.110074	[0.061216, 0.150990]
$C_{90,B}$	0.067779	0.027551	0.047665	[0.017897, 0.077661]
$C_{90,\mathrm{Text}}$	−0.104673	−0.078062	−0.091368	[−0.127920, −0.048050]

P uses input-null and replacement controls; separately trained no-evidence codecs are additionally evaluated for B and RA. For the selected P endpoint, NLL is 0.926588 on FEVEROUS and 0.562378 on MuSiQue (pooled: 0.744483); the corresponding source-weighted Brier scores are 0.559684 and 0.380860 (pooled: 0.470272).

E.2 CODEC-ONLY AND RECEIVER-ALIGNED PACKETS

Table 9: Complete one-slot content and base-utility controls. P denotes the real-report packet endpoint in each row group.

Endpoint	Contrast	FEVEROUS	MuSiQue	Pooled	95% CI
B shared	Trained null − P	0.174995	0.043213	0.109104	[0.078399, 0.143714]
	Same-codec N-input − P	0.303082	0.081779	0.192431	[0.135781, 0.260346]
	Fixed N − P	0.119538	0.045296	0.082417	[0.053046, 0.112355]
	Replacement − P	0.278730	0.088000	0.183365	[0.133747, 0.235226]
	Base − P	0.083197	0.016462	0.049830	[0.021322, 0.080587]
	P − native text	0.108417	0.100886	0.104651	[0.056071, 0.145867]
RA shared	Trained null − P	0.177342	0.115547	0.146445	[0.109825, 0.183726]
	Same-codec N-input − P	0.264603	0.133101	0.198852	[0.157109, 0.243226]
	Fixed N − P	0.132574	0.117510	0.125042	[0.082551, 0.165691]
	Replacement − P	0.282036	0.186815	0.234426	[0.183062, 0.288349]
	P − base	−0.096233	−0.088676	−0.092454	[−0.129500, −0.055449]
	$C_{90,B}$	0.099867	0.091559	0.095713	[0.059634, 0.132244]
	P − native text	0.095381	0.028672	0.062026	[0.019355, 0.099346]
RA fixed-1	Trained null − P	0.171491	0.112398	0.141945	[0.093743, 0.187513]
	Same-codec N-input − P	0.334698	0.142721	0.238710	[0.193971, 0.287567]
	Fixed N − P	0.108463	0.113407	0.110935	[0.060085, 0.159459]
	Replacement − P	0.297522	0.180974	0.239248	[0.182102, 0.297890]
	P − base	−0.072123	−0.084572	−0.078348	[−0.121267, −0.034412]
	$C_{90,B}$	0.075757	0.087456	0.081606	[0.037915, 0.124241]

The aligned shared and fixed-1 packets satisfy all four content controls and base preservation. Natural replacement preserves report-like form while changing source content, and is interpreted as a content effect rather than a cross-model ranking.

E.3 CANDIDATE AND SENDER PERMUTATIONS

Table 10: Independent permutation tests on evaluated endpoints. Agreement is argmax agreement after restoring semantic candidate order.

Evidence object	Receiver seeds	Tasks	Candidate reversal	Sender reversal
H-LOCAL Development	20260822/23/24	F/M	max/mean $\Delta = 0$, agr.=1	max/mean $\Delta = 0$, agr.=1
H-LOCAL confirmation	20260824	F	max $\Delta = 0$, agr.=1	max $\Delta = 0$, agr.=1
P four-slot INT8	20260824	F/M	max $\Delta = 0$, agr.=1	max $\Delta = 0$, agr.=1
B one-slot arms	20260824	F/M	max/mean $\Delta = 0$, agr.=1	max/mean $\Delta = 0$, agr.=1
RA shared/fixed-1 @1	20260824	F/M	max/mean $\Delta = 0$, agr.=1	max/mean $\Delta = 0$, agr.=1

The prespecified acceptance criterion requires a maximum absolute probability change of at most 10^{-4} and identical argmax predictions. For every endpoint in Table 10, the observed maximum change was zero in both the candidate- and sender-reversal tests. Each result compares independent forward passes on the original and reversed inputs. The reported RA results cover the evaluated one-slot endpoints.

F PACKETIZATION AND TRAINING-SCOPE ABLATIONS

F.1 FIXED-READER PACKETIZERS AND QUANTIZATION

Neither deterministic segment nor uniform projection passes the content criteria. The selected slot128-match endpoint has INT8 NLL 0.744483 versus 0.743531 in FP16, and its base-preservation contrast remains positive after quantization.

F.2 MULTIRATE AND TRAINING-SCOPE COMPARISONS

The shared-versus-fixed intervals cross zero for B at one and two slots (0.004163 $[-0.011718, 0.020254]$ and 0.012769 $[-0.004516, 0.031008]$) and for RA at one slot (0.014107

Table 11: Four-slot FP16 packetizer sweep. Brier is source-weighted; content indicates whether all specified report-content criteria are satisfied.

Arm	NLL	Brier	Mean bytes/request	Content
segment128	0.944752	0.611305	2,794.598	no
uniform128	0.941096	0.609740	2,794.598	no
slot128-task	0.768853	0.473615	2,794.598	yes
slot128-match	**0.743531**	**0.469562**	2,794.598	**yes**
segment256	0.941977	0.622581	5,547.558	no
uniform256	0.958065	0.633616	5,547.558	no
slot256-task	0.753770	0.471625	5,547.558	yes
slot256-match	0.746119	0.470519	5,547.558	yes

Table 12: Base- and text-utility contrasts for the same FP16 packetizer sweep. Entries are point estimates with 95% source-bootstrap intervals.

Arm	$C_{90,B}$	$C_{90,\mathrm{Text}}$
segment128	-0.152604 [-0.187469,-0.117567]	-0.291637 [-0.337386,-0.243224]
uniform128	-0.148948 [-0.182828,-0.114878]	-0.287981 [-0.334735,-0.238724]
slot128-task	0.023295 [-0.016100,0.060701]	-0.115738 [-0.156757,-0.071937]
slot128-match	**0.048617 [0.019160,0.078964]**	-0.090416 [-0.126663,-0.048450]
segment256	-0.149829 [-0.186069,-0.112645]	-0.288862 [-0.338338,-0.235010]
uniform256	-0.165917 [-0.210019,-0.122798]	-0.304950 [-0.364337,-0.245476]
slot256-task	0.038377 [0.003496,0.072337]	-0.100655 [-0.141765,-0.056895]
slot256-match	0.046029 [0.017351,0.074410]	-0.093004 [-0.128357,-0.050955]

$[-0.006354, 0.035744]$). Shared alignment has a positive interval, 0.042625 $[0.019817, 0.063856]$, whereas fixed-1 alignment remains inconclusive at 0.027081 $[-0.002876, 0.055971]$.

At one slot, RA shared has lower NLL and raw-row Brier than RA fixed-1, while raw-row accuracy is 0.583330 versus 0.586931. Representative-arm selection uses one-slot equal-task/source NLL.

A descriptive comparison using the frozen pre-alignment B predictions gives a task-only-minus-behavior-matched one-slot NLL difference of 0.012845 $[-0.008348, 0.038961]$. The point estimate favors behavior matching, but the interval crosses zero and the task-level effects differ.

Figure 4 summarizes the prefix and training-scope comparisons.

Table 13: Selected codec-only B and receiver-aligned RA endpoints.

Stage/arm	Slots	Mean bytes/request	FEVEROUS	MuSiQue	Pooled NLL
B fixed1-match	1	396.512	0.923851	0.562594	0.743223
B fixed2-match	2	740.632	0.906067	0.580583	0.743325
B successive-match	1	396.512	0.907535	0.570583	0.739059
B successive-match	2	740.632	0.896287	0.564825	0.730556
B successive-match	4	1,428.872	0.890584	0.564970	0.727777
B successive-task	1	396.512	0.945090	0.558720	0.751905
B successive-task	2	740.632	0.958994	0.558361	0.758677
B successive-task	4	1,428.872	0.960508	0.557954	0.759231
B successive-null	1	396.512	1.082531	0.613797	0.848164
B successive-null	2	740.632	1.092487	0.614072	0.853280
B successive-null	4	1,428.872	1.088897	0.614579	0.851738
RA codec-only fixed1	1	396.512	0.917467	0.557777	0.737622
RA codec-only shared	1	396.512	0.907535	0.570583	0.739059
RA codec-only shared	2	740.632	0.896287	0.564825	0.730556
RA codec-only shared	4	1,428.872	0.890584	0.564970	0.727777
RA null shared	1	396.512	1.071842	0.613917	0.842879
RA null shared	2	740.632	1.102596	0.613580	0.858088
RA null shared	4	1,428.872	1.099638	0.613524	0.856581
RA fixed1	1	396.512	0.918610	0.502473	0.710541
RA shared	**1**	**396.512**	**0.894500**	**0.498370**	**0.696435**
RA shared	2	740.632	0.896738	0.499497	0.698117
RA shared	4	1,428.872	0.897582	0.499940	0.698761

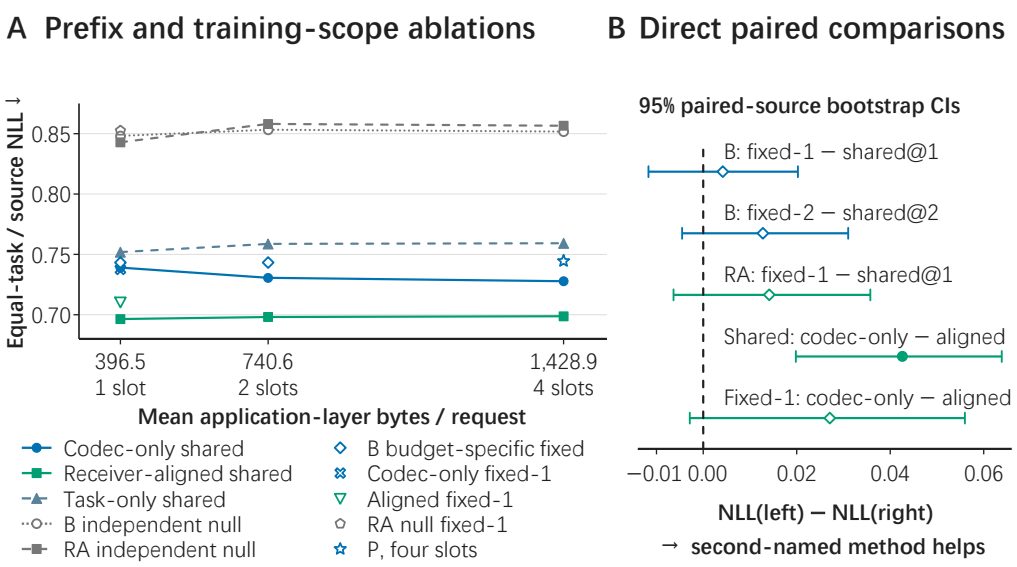


Figure 4: Prefix and training-scope ablations on single-receiver-seed development data. (A) Lines connect budgets evaluated from one shared checkpoint; isolated markers denote separately trained fixed-budget codecs and the four-slot P codec. The B shared and RA codec-only shared endpoints coincide because the latter selects its warm start. (B) Paired NLL differences with 95% source-bootstrap intervals; positive values favor the second-named method. The aligned curve does not improve with additional slots, and the shared-versus-fixed intervals cross zero.

G Full Communication Frontier

G.1 Text endpoints

Table 14: Text communication endpoints. The continued models select epoch 0 and therefore exactly reproduce the adapted warm-start predictions.

Receiver	Selection	Budget	Mean bytes/request	Pooled NLL
Frozen parent	head	m1	366.904	0.825488
Frozen parent	head	m2	543.219	0.679721
Frozen parent	head	m4	718.727	0.642594
Frozen parent	query-overlap	m1	366.902	0.833875
Frozen parent	query-overlap	m2	543.214	0.682201
Frozen parent	query-overlap	m4	718.726	0.647401
Adapted/continued	head	m1	366.904	0.653758
Adapted/continued	head	m2	543.219	0.644965
Adapted/continued	head	m4	718.727	0.634101
Adapted/continued	query-overlap	m1	366.902	0.664098
Adapted/continued	query-overlap	m2	543.214	0.649167
Adapted/continued	query-overlap	m4	718.726	0.636496

Adapting the receiver reduces head-m1 NLL from 0.825488 to 0.653758. This change is larger than the differences among frozen text selection heuristics, which motivates treating receiver adaptation as part of a strong communication baseline.

G.2 Latent, text, score, and activation references

The mean-byte-matched text reference's advantage over RA shared is 0.044275 $[0.015197, 0.074273]$, with FEVEROUS/MuSiQue point effects 0.066620/0.021931. The hard-cap head-m1 endpoint also

Table 15: Complete Development risk–byte frontier for the principal communication identities.

System	Mean bytes/request	Pooled NLL	Payload identity
No communication	0.000	1.022123	No message
Candidate-score mean	55.556	0.857570	Per-sender FP16 candidate scores, mean aggregation
Candidate-score product	55.556	1.316921	Same scores, product aggregation
Adapted text head m1	366.904	0.653758	146-byte/sender hard-cap text
Codec-only shared @1	396.512	0.739059	One-slot INT8 latent packet
INNOCOMM-RA shared @1	**396.512**	**0.696435**	Receiver-aligned one-slot INT8 packet
Mean-byte-matched text reference	396.512	0.652159	Random endpoint mixture, including one-byte control
Adapted text head m2	543.219	0.644965	Mid-budget hard-cap text
Adapted text head m4	718.727	0.634101	High-budget hard-cap text
Native full text	842.754	0.634408	Complete sender-local text
INNOCOMM-P four-slot INT8	1,428.872	0.744483	Fixed-reader latent packet
Readout K/V FP16	215,221.428	0.789261	Internal activation readout
Full H-LOCAL FP16	1,506,300.164	0.788889	Complete contextual activation bank

uses fewer mean bytes and has lower point risk than RA shared, so the text result does not depend on occasionally invoking the m4 endpoint. Alignment reduces the earlier B packet's measured gap to its mean-byte-matched text reference from 0.086956 to 0.044275, a descriptive reduction of 49.1%; the two mixtures use slightly different contracts, so the matched alignment contrast in Table 3 remains the direct method test.

H CACHED-HIDDEN SYSTEM MEASUREMENTS

The profiled endpoint is B `successive_match@1`, receiver seed 20260824, on an RTX A6000. Sender contextual states are cached before timing. For each task, 32 rows receive two warm-up passes followed by five measured passes, producing 160 timing samples; repeated passes are not treated as independent sources.

Table 16: Cached-hidden latency in milliseconds (mean / p50 / p95).

Component	FEVEROUS	MuSiQue
Packet encode + INT8 round trip + reconstruction	1.678 / 1.663 / 2.351	1.056 / 1.009 / 1.340
Receiver candidate scoring	41.696 / 37.271 / 58.023	29.073 / 28.962 / 29.676
Cached-hidden total	43.373 / 39.241 / 59.592	30.130 / 29.965 / 30.874

Packet preparation accounts for 3.87% and 3.51% of mean cached-hidden time. The measurement excludes cold sender encoding, physical network transfer, cross-machine scheduling, FLOPs, energy, and full-inference peak memory. These measurements characterize the pre-alignment one-slot B endpoint; no receiver-aligned timing is reported.

I PRE-ALIGNMENT DIAGNOSTICS

These descriptive diagnostics use the frozen one-slot B shared codec and its pre-alignment mean-byte-matched text-mixture reference on 384 development sources. Table 17 reports the mean packet-minus-reference NLL gap; positive values favor text. The diagnostics characterize the pre-alignment system rather than the trained RA model.

The group means do not follow a consistent monotone pattern across length, uncertainty, margin, or sender count. The row-level overlap includes examples correctly predicted only by the packet and only by text, despite the text reference's lower aggregate risk.

Figure 5 visualizes these source-group gaps and row-level error overlaps.

Table 17: Packet–text gap across source quartiles.

Grouping variable	Q1	Q2	Q3	Q4
Full-bank length	0.132825	0.014094	0.114356	0.086548
Base entropy	0.097376	0.089955	0.079833	0.080660
Base margin	0.083138	0.045765	0.099933	0.118988
Packet margin	0.013577	0.115605	0.126707	0.091934
Text margin	0.062136	0.023713	0.087251	0.174724
Report bytes	0.130341	0.048315	0.069263	0.100228

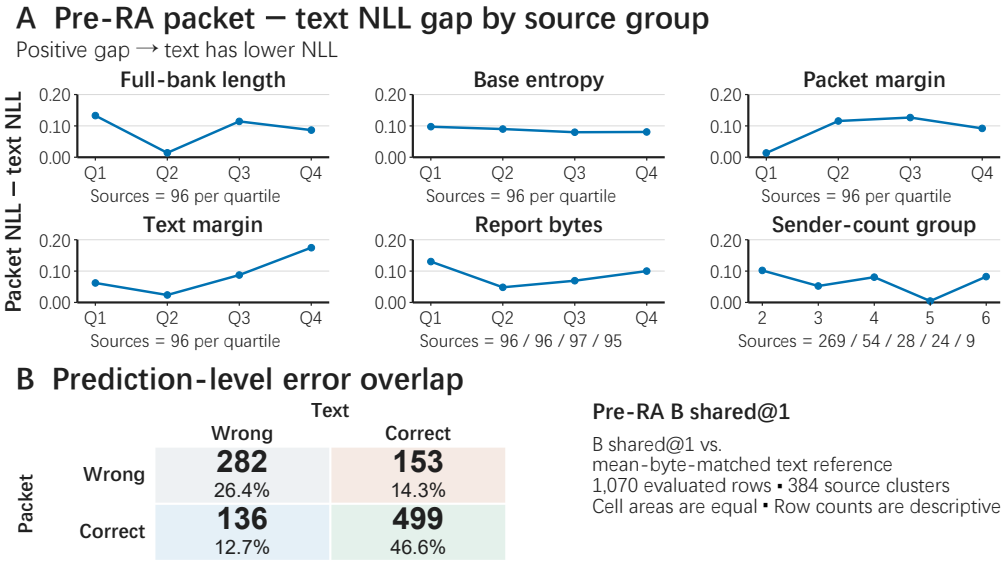


Figure 5: Pre-alignment development diagnostics from the frozen B codec. **A:** Descriptive packet-minus-text-reference NLL gaps by source group, with group sizes shown; positive values favor text. **B:** Packet/text correctness for 1,070 evaluation rows from 384 sources. Equal-sized cells show counts and percentages. Group means have no confidence intervals; row-level counts are descriptive and are not treated as independent source replicates.