

shamikkundu at the FinEval 2026: Zero-Parameter Confidence Elicitation for Japanese Implicit Communication Recognition

Shamik Kundu

Trust Co.,Ltd.

shamik_kundu@trust-partner.co.jp

Abstract

We describe our submission to FinEval 2026 Subtask 2, Japanese Financial Implicit Communication Recognition (JF-ICR): a five-point ordinal classification of company responses in investor Q&A, from +2 (strong commitment) to -2 (strong refusal). Our system fits no parameters: a 14B instruction model, served through OpenRouter, receives the shared task’s own annotation guideline verbatim and emits a schema-enforced confidence for each of the five labels; we take the argmax of the normalised distribution. Our submission ranked **5th**, at **0.9429** accuracy on the held-out test rows.

1 Introduction

Japanese corporate disclosure is a high-context register. Management rarely refuses an analyst outright; it defers, hedges, or commits to something adjacent to what was asked. JF-ICR, introduced with the EBISU benchmark (Peng et al., 2026b) and run as Subtask 2 of FinEval 2026 (Peng et al., 2026a), asks a system to place a company response on a five-point commitment scale: +2 strong commitment, +1 weak or qualified commitment, 0 neutral or hedged intent, -1 weak refusal, -2 strong refusal. The official metric is accuracy.

Two properties of the released data drove every design decision we made. It is small — 253 labelled rows, of which the extreme classes have 29 (+2) and 2 (-2) examples — and it is noisy, in ways we document in §2. We therefore built a system with *zero fitted parameters*: label noise cannot be internalised into a learned artefact, so it adds variance to our estimates rather than bias to our model.

Our contributions:

- A prompted, schema-constrained confidence-elicitation pipeline that gains +0.192 accuracy over EBISU’s zero-shot prompt on a

fixed model, and +0.12 to +0.30 across four models we did not select (§4).

- A quantified account of the release-to-test distribution shift, and of a design decision of ours that it reversed (§5).

2 Task Data and Its Defects

The public release is a single 253-row file that merges two annotation batches, separable by a four-character whitespace difference in the prompt preamble that maps exactly onto contiguous id ranges. Ids 0–93 are the benchmark paper’s evaluation set and ids 94–252 are unpublished.

The majority baseline beats the published results. Always predicting +1 scores $65/94 = 0.6915$ on the official split. The best system reported in the benchmark paper (Llama-4-Scout-17B) scores 0.6064. A system can therefore top this leaderboard by discarding all class discrimination, so we report **macro-recall** beside accuracy wherever we compare configurations. Chance macro-recall is 0.20 for *any* constant predictor regardless of skew, so unlike accuracy it cannot be gamed by leaning on the prior. We also report quadratic weighted kappa (Cohen, 1968), which credits near misses on the ordinal scale.

Duplicate responses with contradictory labels.

Under NFKC normalisation with punctuation stripped, the release contains five duplicate response pairs, all within ids 94–252, four of them label-inconsistent: id100~id129 carry +2 and -2 — opposite ends of the scale on 128 identical characters — and three pairs carry -1 against +1. Consequently -2 has two examples corpus-wide, one of which duplicates a +2; and -1 is 38% self-contradictory within the batch.

Given the limited amount of training data and the noisy labels, the submitted system fits no parameters: we use a pre-trained LLM zero-shot. The clas-




Figure 1: The submitted pipeline.

sical models in §6 are exploratory, cross-validated on the release with folds grouped by contiguous id block, and were never shipped.

3 System

The submitted system is one model call per row (Figure 1).

Guideline in the prompt. Appendix A of the benchmark paper contains the full annotation guideline. We transcribe it verbatim under a fidelity rule — no paraphrase, no added examples, no class priors, no thresholds — and send it as the label definition, since prose definitions of “strong” and “weak” cannot convey a decision boundary but the guideline’s operational criteria can. Because the prompt carries no corpus-specific content, the system applies unchanged to a new test set.

Confidence elicitation under an enforced schema. Rather than asking for one label, we require a confidence in $[0, 100]$ and a one-line justification for *each* of the five labels, with all five fields mandatory, structurally enforced by the provider’s JSON-schema or forced tool-call facility. This turns a 5-valued output into a near-continuous score. Across ten model/provider combinations our schema step produced a complete five-label distribution on 253/253 rows with zero errors every time.

Decision rule. We normalise the five confidences to sum to one and take the argmax; ties break toward the more positive label. The argmax is invariant to any monotone per-row rescaling, so this choice cannot change a prediction: normalisation matters only where the magnitudes rather than the ranking are used.

arm	acc	mF1	mR	QWK
L0 organizers’ zero-shot	0.4862	0.3933	0.4369	0.3855
L1 + our prompt, 5 conf.	0.5613	0.4091	0.4859	0.5236
L2 + published guideline	0.6719	0.3772	0.3725	0.5118
L3 + question withheld	0.6785	0.4653	0.4803	0.4707
L3s: L3, one label out	0.6535	0.4536	0.4667	0.4760
<i>majority baseline</i>	<i>0.5494</i>	<i>0.1418</i>	<i>0.2000</i>	<i>0.0000</i>

Table 1: Component ablation on all 253 rows, gemini-3.5-flash, temperature 0, mean of three replicates. Each arm adds one component to the one above. L0 sends the organizers’ fully-rendered instruction verbatim, with no schema, retry or output repair; L3s re-runs L3 asking for one label instead of five confidences; the majority baseline always predicts +1.

model	published	ours	Δ
qwen3-32b	0.1915	0.4255	+0.234
ministral-3-14b	0.3723	0.6702	+0.298
qwen3-235b	0.4574	0.6170	+0.160
kimi-k2-thinking*	0.4574	0.5745	+0.117
<i>best published 0.6064; majority baseline 0.6915</i>			

Table 2: Our pipeline versus the benchmark paper’s zero-shot accuracy on the official 94, same models, single run each. *the organizers benchmarked Kimi-K2-*Instruct*; we used Kimi-K2-*thinking* — same family, different variant, so weaker evidence than the other three.

Eligibility. The shared task caps each model in an API-based workflow at 70B total parameters and requires calls through OpenRouter. We submit **Ministral-3-14B-Instruct** (Mistral AI, 2025), 14B dense, called as mistralai/ministral-14b-2512; the complete prompt and run configuration are in Appendices A and B. All development-time model-comparison runs in §4 that use models above the cap are reported as diagnostics, not as the submitted system.

Two test-informed choices. The submitted system includes the analyst question. Configuration L3 in §4, which withholds it, is the arm selected on the released corpus; we reversed that decision before submission on the evidence in §5. The model and the question are therefore both test-informed choices.

4 Results

Component ablation. Table 1 adds one component per step on a fixed model, so every delta is attributable. **End to end**, 0.4862 → 0.6785, **an accuracy gain of +0.192 on the same**

model	$ \theta $	acc	mF1	mR	QWK
gpt-oss-20b	21B	0.6522	0.4527	0.4908	0.5120
gemma-3-27b	27B	0.6126	0.2879	0.2956	0.4157
ministral-3-14b	14B	0.5771	0.2755	0.2973	0.3898
qwen3-32b	32B	0.4704	0.2758	0.3538	0.4197
glm-4.7-flash	31.2B	0.4229	0.2427	0.2903	0.2824

Table 3: Eligibility screen: models under the 70B cap, identical pipeline, temperature 0, single run, all 253 rows.

model (+0.234 on the official 94: 0.4574 $\rightarrow$ 0.6915). By paired bootstrap on the accuracy difference (Efron and Tibshirani, 1994) (4000 resamples) two of the four components are resolvable at $n=253$: the prompt-and-confidence rewrite, +0.0751, CI [+0.024, +0.142], $P=0.996$; and the guideline, +0.1107, CI [+0.040, +0.174], $P=0.998$. Withholding the question (+0.0066, $P=0.658$) and confidence elicitation in isolation (+0.0250, $P=0.810$) are not, and we report them as included but unresolved rather than as demonstrated gains — though confidence elicitation is same-signed on all three evaluation splits and in all three replicates, which is better evidence than the bootstrap alone conveys.

The method transfers across models. Table 2 applies the pipeline unchanged to four models the organizers benchmarked. All four improve, by +0.117 to +0.298, on models we did not choose, which attributes the gain to method rather than to model selection. Bounded honestly, the delta compares our full pipeline against their zero-shot label prediction, so it bundles every change at once.

The guideline is mostly a prior shift. L1 $\rightarrow$ L2 buys +0.1107 accuracy while costing **-0.1134** macro-recall (-0.189 on the 94). The gain is substantially the model leaning harder on +1. The consequence is visible in the arm ordering: **L1 has the best QWK of any arm** (0.5236 versus L3’s 0.4707) and better macro-recall, at 0.117 lower accuracy. If this task were scored by QWK we would ship L1 and never add the guideline. The optimal system here is determined by the metric, not by the data.

Eligible models. Table 3 screens the models we could actually submit. On the released corpus gpt-oss-20b leads; against our test-set pseudo-labels the ordering inverts (Ministral-3-14B 0.90, gpt-oss-20b 0.74), and we submitted Ministral-3-14B, which had also scored 0.6702 on the official 94,

	rel. 253	off. 94	test 50 [†]
question withheld	0.6785	0.6915	0.90
question included	0.6719	0.7199	0.98

Table 4: The question ablation reverses. [†]test-50 figures are Ministral-3-14B scored against our own manual pseudo-labels, which are an annotator proxy, not gold; the system was given no knowledge of them. The 253/94 columns are gemini-3.5-flash, mean of three runs.

above every published system. Because those pseudo-labels are exactly class-balanced, accuracy and macro-recall coincide on them, and the candidate ordering is identical under all four metrics we report (Ministral-3-14B 0.90/0.90/0.90/0.82 and gpt-oss-20b 0.74/0.74/0.73/0.73 for accuracy, macro-recall, macro-F1 and QWK), so this selection is not an artefact of the official metric. We state plainly that the choice was nonetheless made on the test distribution, not on the released corpus.

Official result. Our submission placed **5th** with **0.9429** accuracy on the held-out test rows. Our submitted predictions agree with our own pseudo-labels on all 50 rows, so the proxy had estimated 1.00 and was optimistic by 0.057 against the official figure — the bias to expect from an annotator proxy built by the system’s own authors.

5 What Did Not Transfer

When the analyst question is withheld, all five errors on the 50-row test set point the same way — too positive, none too negative (Table 4). On two rows the response refuses the proposed action while committing to its opposite: 「いいえ。当該製品は今期末で販売と生産を完全に終了します」。Without the question the model sees 終了します, a decisive, +2-shaped verb, and scores commitment; it cannot know that what is being refused is *continuing the line*. On two others a hedged negative (「難しいと見ています。ただし...検討します」) flattens to 0 because nothing in the response alone marks it as a stance *against* something.

We had withheld the question because on the released corpus it measures as a class-coverage repair, worth +0.088 macro-F1 and +0.108 macro-recall on the 253 (and +0.119/+0.218 on the development split) at no accuracy cost. It undoes much of the coverage damage the guideline causes. That effect is real on the release and it does not transfer,

because the two input distributions are not the same text:

	released 253	test 50
mean length (chars)	93.1	22.4
direct yes/no questions	6.3%	98%

The release asks long, open-ended prose questions, for which the response usually carries its own stance. The test set asks short, direct yes/no questions — 「物流センターの建設を来年開始しますか。」 — where the question names the proposed action and is the only thing that fixes the *sign* of the answer. Our submitted predictions are also exactly balanced, 10 rows per class, and agree with our manual pseudo-labels on all 50 rows, which suggests the test set is class-balanced by construction rather than carrying the release’s 55% +1 prior. If so, the guideline’s prior shift and the question ablation — both tuned under a skewed prior — were being validated on the wrong distribution.

6 Negative Results

The head-final hypothesis failed. Japanese is head-final and the benchmark paper’s own claim is that stance lives in clause-final morphology, so we built ten features measuring modality on the final predicate. They scored 0.0586 QWK alone and *actively hurt* the full feature set (0.2704 → 0.2061); stripping PDF furniture and re-segmenting did not rescue them. The signal in this corpus lives in *aggregate* lexical and morphological counts, not in clause-final position.

Automatic lexicon induction is worse than no lexicon. Deriving commitment/refusal word lists inside each training fold — the clean alternative to a hand lexicon — scored 0.1530, below the 0.1844 of lexicon-free morphology. It learned topic terms and company names, which is document-identity leakage of exactly the kind grouped folds exist to prevent. Our best classical feature set was a hand lexicon written while inspecting the corpus, worth +0.20 QWK and of unverifiable provenance; replacing it with features specified a priori from published Japanese modality description (Matsuyoshi et al., 2010) recovered about 89% of the QWK, inside one fold standard deviation. We report the a-priori set and disclose the other.

Fewer features win, consistently. Six features beat twenty-four (0.3557 versus 0.2061); adding classical features to LLM features hurt (0.4991 →

0.4476); adding character TF-IDF to modality features gave nothing. At this n , dimensionality costs more in variance than it returns in signal, in every comparison we ran.

Also null or negative. Two-stage pairwise adjudication; coarse-to-fine hierarchical elicitation; ensembling across models, prompts and seeds; few-shot exemplars (null on accuracy, negative on the macro metrics); a Japanese-language instruction wrapper; analysis-before-scoring; and voting across decoding temperatures. Nothing that worked was about model capacity, ensembling, or search: the two components that survived (§4) are the output representation and the label definition.

7 Conclusion

A prompted 14B model with no fitted parameters, given the shared task’s own annotation guideline and asked for a calibrated confidence on every label, gains +0.192 accuracy over EBISU’s zero-shot prompt on a fixed model and transfers across every model we tried. Our solution is ranked 5th on the official leaderboard at 0.9429 accuracy.

Limitations

Statistical power. Most of the experiments are run only once. Since LLMs are stochastic, running every experiment at least three times and taking the average would have been ideal.

Model choice made on the test distribution. Ministral-3-14B is not the best eligible model on the released corpus (Table 3); we chose it on the test distribution instead. Testing a larger eligible model, such as the 21B gpt-oss-20b that leads Table 3, might have produced a better score.

Accuracy drives the component choices. Accuracy is the official metric, so we selected under it, and only one of our components improves under every metric we report: the prompt-and-confidence rewrite gains on all four (L0→L1), the guideline is an accuracy-only gain (−0.032 macro-F1, −0.113 macro-recall, −0.012 QWK), and withholding the question — which we reversed — was a macro-only gain. A system chosen under QWK would be L1 with the question included. The model choice is the one selection that is metric-invariant (§4).

Ethics Statement

All data is the shared task’s public release of material already published by companies in investor-relations, presentation and shareholder-meeting

documents; we introduce no new data collection and no personal data.

References

- Jacob Cohen. 1968. Weighted kappa: Nominal scale agreement with provision for scaled disagreement or partial credit. *Psychological Bulletin*, 70(4):213–220.
- Bradley Efron and Robert J. Tibshirani. 1994. *An Introduction to the Bootstrap*. Chapman & Hall/CRC.
- Suguru Matsuyoshi, Megumi Eguchi, Chitose Sao, Koji Murakami, Kentaro Inui, and Yuji Matsumoto. 2010. Annotating event mentions in text with modality, focus, and source information. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC)*.
- Mistral AI. 2025. Ministral 3 14B Instruct. <https://huggingface.co/mistralai/Ministral-3-14B-Instruct-2512>. Model card.
- Xueqing Peng, Lingfei Qian, Yan Wang, Ruoyu Xiang, Yueru He, Yang Ren, Mingyang Jiang, Vincent Jim Zhang, Yuqing Guo, Jeff Zhao, Huan He, Yi Han, Yun Feng, Yuechen Jiang, Yupeng Cao, Haohang Li, Yangyang Yu, Xiaoyu Wang, Penglei Gao, and 28 others. 2026a. [MultiFinBen: Benchmarking large language models for multilingual and multimodal financial application](#). In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Xueqing Peng, Ruoyu Xiang, Fan Zhang, Mingzi Song, Mingyang Jiang, Yan Wang, Lingfei Qian, Taiki Hara, Yuqing Guo, Jimin Huang, Junichi Tsujii, and Sophia Ananiadou. 2026b. [Ebisu: Benchmarking large language models in japanese finance](#). *arXiv preprint arXiv:2602.01479*.

A Reproducibility

Model mistralai/ministral-14b-2512 via OpenRouter, temperature 0, max_output_tokens 8192, 4 retries, 8 concurrent workers. One call per row; no batching. We've submitted the code to reproduce the results to the organizers already.

B The Full Prompt

Below is the exact text sent to the model for every row, reproduced verbatim from the submitted code. The label definitions, cue expressions and worked examples are Appendix A of [Peng et al. \(2026b\)](#), transcribed without paraphrase, added examples, class priors or thresholds. Only the final two lines vary between rows; the pair shown is the item discussed in §5. The five-key JSON object in the closing line is also enforced structurally, through the provider's schema or forced tool call, so the model cannot return prose instead.

You are a Japanese financial expert fluent in Japanese business communication.
Given a financial question and the corresponding company response (both in Japanese),
assess the company's underlying intent level on this five-point scale:
ANNOTATION GUIDELINE

Annotation Philosophy. JF-ICR focuses on the degree of commitment expressed in management responses, rather than the surface intent of the question or the factual correctness of the content. Annotations are determined solely by whether and how the answer conveys a forward-looking action, decision, or policy stance, reflecting common communication practices in financial disclosures.

Due to the cautious and convention-driven nature of Japanese financial disclosures, explicit strong commitments are relatively rare. As a result, a substantial portion of responses naturally fall into weakly committed or neutral categories, reflecting realistic managerial communication practices rather than annotation ambiguity or inconsistency.

Classification categories.

Each answer is annotated based on the speaker's degree of commitment toward a future action, decision, or policy stance expressed in the response.

- +2 Strong Commitment - clear, explicit, and unambiguous commitment to a concrete action or outcome.
- +1 Weak or Qualified Commitment - intention or directional stance expressed with conditions, hedging, or limited scope.
- 0 Neutral or Hedged Intent - clarification, explanation, assessment, or perception without commitment or refusal.
- 1 Weak Refusal - implicit or qualified rejection of a proposed action.
- 2 Strong Refusal - explicit and definitive rejection with no flexibility.

General annotation rules.

- Annotation focuses on intent and commitment strength, not factual correctness.
- Only statements expressing future actions, decisions, policies, targets, or strategic directions are considered for non-zero labels.
- Annotation is based on the content of the answer, not the surface form or intent of the question. Even if the question is retrospective or explanatory, the answer may introduce forward-looking commitment or refusal.
- If an answer does not express any future-oriented action, decision, or policy stance, it should be annotated as 0, regardless of the question.
- Hedging expressions (e.g., uncertainty, dependency on conditions, ongoing review) reduce commitment strength.
- When multiple signals appear, the label reflects the strongest applicable commitment or refusal expressed in the answer.
- Annotation is based on the overall stance of the full response, rather than isolated sentences or individual expressions; labels reflect the speaker's dominant commitment or refusal as conveyed by the answer as a whole.
- Annotations focus on the strength of commitment, rather than the tone, length, or politeness of the response.
- Conventional polite expressions in Japanese financial discourse (e.g., 「努めてまいります」) do not by themselves imply commitment.

Specific annotation rules.

+2 (Strong Commitment): Applied when the answer contains a clear, decisive statement indicating a firm commitment. Typical linguistic signals include:

- 「～します」「～を実施します」「～を達成します」
- 「～を決定しています」「～は確定しています」
- 「方針を変更する考えはありません」 (when directly answering a decision or policy question)

Representative examples include:

- 「来期は増配を実施します。」
- 「中長期でROE10%を達成します。」
- 「この施策により利益成長は実現できると確信しています。」

+1 (Weak or Qualified Commitment): Applied when the answer expresses a positive or leaning commitment toward a future action or outcome, but with qualifications, caution, or limited specificity. Such responses indicate directional intent or possibility rather than a finalized decision. Typical linguistic signals include:

- 「～していきたい」「～を目指しています」
- 「～と考えています」「～を見込んでいます」
- Conditional or hypothetical expressions (e.g., 「～であれば」「～次第で」)

Representative examples include:

- 「成長投資を進めていきたいと考えています。」
- 「今後も収益は拡大していくと見ています。」
- 「環境が整えば、検討を進める考えです。」

0 (Neutral or Hedged Intent): Applied when the answer exhibits genuine ambiguity or provides clarification, explanation, or background without expressing commitment or refusal toward any future action, decision, or policy stance. Such responses may describe facts, context, perceptions, or reasoning, but do not convey a directional or evaluative position regarding future behavior. Typical linguistic signals:

- 「～断定できません」「明確な見通しは示せない」
- 「検討中」「状況を見極める必要がある」
- Purely descriptive or explanatory statements providing facts or background

Representative examples include:

- 「現時点では明確な見通しは示せません。」
- 「様々な見方があり、コメントは差し控えます。」
- 「過去にはこのような取り組みを行ってきました。」 (background explanation only)

-1 (Weak Refusal): Applied when the answer expresses a negative stance toward a proposed future action or change, but does so in a qualified, conditional, or time-bound manner that leaves room for future reconsideration. Such responses reject the action in the current context without ruling it out permanently.

Typical linguistic signals:

- 「現時点では～しない」「直ちには考えていない」
- 「今後検討の余地はあるか」
- Refusals framed as temporary, conditional, or dependent on future circumstances

Representative examples include:

- 「現時点では配当方針を変更する考えはありません。」
- 「今中計期間中に見直すことは想定していません。」
- 「足元では難しいと考えていますが、今後は検討します。」

-2 (Strong Refusal): Applied when the answer clearly and definitively rejects a proposed future action, decision, or policy stance, leaving no visible room for reconsideration. Such responses convey a firm and final negative position.

Typical linguistic signals:

- 「～する予定はありません」
- 「～は行いません」「～を否定します」

Representative examples include:

- 「株式分割を行う予定はありません。」
- 「当該事業への投資は実施しません。」
- 「その想定は当社の方針ではありません。」

Judge the degree of commitment expressed in the COMPANY RESPONSE -- whether and how it conveys a forward-looking action, decision or policy stance. Do not judge the question, the factual correctness, or general sentiment. Only strong, clear, explicit and definitive agreement or refusal warrants the extreme labels.

For EVERY one of the five labels give:

- confidence: an integer 0-100, how well that label fits this response
- reason: one short line justifying that confidence

All five fields are required; never omit one.

Confidences need not sum to 100. Use the full range: reserve high confidence for labels that genuinely fit, and give low confidence to labels that do not.

Financial Question: 赤字が続く製品ラインを来期も継続しますか。

Company Response: いいえ。当該製品は今期末で販売と生産を完全に終了します。

Reply with exactly this JSON object and nothing else:

```
{"+2": {"confidence": <int 0-100>, "reason": "<one line>"}, "+1": {"confidence": <int 0-100>, "reason": "<one line>"}, "0": {"confidence": <int 0-100>, "reason": "<one line>"}, "-1": {"confidence": <int 0-100>, "reason": "<one line>"}, "-2": {"confidence": <int 0-100>, "reason": "<one line>"}}
```