

EVA-Bench: A New End-to-end Framework for Evaluating Voice Agents

Tara Bogavelli, Gabrielle Gauthier Melançon, Katrina Stankiewicz,
Oluwanifemi Bamgbose, Fanny Riols, Hoang Nguyen, Raghav Mehndiratta,
Lindsay Brin, Hari Subramani, Joseph Marinier,
Sridhar Krishna Nemala, Anil Madamala, Srinivas Sunkara

ServiceNow AI Research

Correspondence: tara.bogavelli@servicenow.com

Abstract

Voice agents are increasingly deployed across enterprise applications. However, no existing benchmark jointly addresses realistic conversation simulation and comprehensive voice-specific evaluation. We present **EVA-Bench**, an end-to-end evaluation framework that addresses both. On the simulation side, EVA-Bench orchestrates dynamic bot-to-bot audio conversations with automatic simulation validation that detects user simulator error and appropriately regenerates conversations before scoring. On the measurement side, EVA-Bench introduces two composite metrics: **EVA-A** (Accuracy) and **EVA-X** (Experience). EVA-Bench includes 213 scenarios across three enterprise domains, a controlled perturbation suite for accent and noise robustness, and multi-trial measurements that distinguish peak from reliable capability. Across 12 systems spanning three architectures, we find: (1) no system simultaneously exceeds 0.5 on both **EVA-A_{pass@1}** and **EVA-X_{pass@1}**; (2) peak and reliable performance diverge substantially (median $\text{pass@k} - \text{pass}^k$ gap of 0.44 on EVA-A); and (3) accent and noise perturbations expose substantial robustness gaps, with effects varying across architectures, systems, and metrics (mean Δ up to 0.314). We release EVA-Bench under an open-source license ¹.

1 Introduction

Voice agents are Artificial Intelligence (AI) systems designed to carry out tasks through spoken conversations, and their deployment across a wide range of applications is rapidly growing (Moore, 2025). Voice agents operate under constraints that are fundamentally distinct from text: speech is ephemeral and linear, real-time timing shapes the naturalness of interaction, and acoustic conditions vary widely across callers. These properties give rise to failure modes with no direct text analog (Ray et al., 2026;

Chen et al., 2026), and render evaluation frameworks designed for text-based agents (Yao et al., 2025; Qin et al., 2024; Patil et al., 2025) insufficient for assessing voice agent quality. Rigorous evaluation of voice agents must therefore address two distinct challenges: how conversations are simulated, and how quality is measured.

The simulation challenge concerns constructing interactions that are valid proxies for real deployment conditions. This requires complete multi-turn interactions as only full conversations expose how an agent recovers from misunderstandings and resolves tasks end-to-end. Conversations must reflect the task-oriented nature of real voice agent deployments, user behavior must reflect natural human spoken dialogue, and acoustic conditions must reflect real-world environments, including variation in accents and background noise. Simulated users must themselves be validated: a simulator that drifts from its scenario, behaves unrealistically, or varies across runs undermines any downstream evaluation, conflating agent variance with simulator variance.

The measurement challenge concerns capturing the full scope of voice agent quality once valid simulations are in place. Task completion and turn-taking dynamics, while necessary, leave critical failure modes undetected (Cao et al., 2026; Mehta, 2025; Andres et al., 2026). On the accuracy side, an agent may call the correct tools yet violate system policy or produce spoken outputs containing incorrect entities (e.g. wrong confirmation codes, or monetary amounts) that are catastrophic in production yet undetectable from transcript-level evaluation alone. On the user experience side, an agent may achieve low response latency yet fail to make meaningful progress across turns, repeat prior questions, or present users with an excessive number of spoken options that would overwhelm a user’s working memory. Addressing the measurement challenge requires evaluation across a broader set

¹<https://github.com/ServiceNow/eva>

Table 1: **Feature comparison of contemporary voice agent evaluation frameworks.** **Multi-turn:** supports multi-turn conversations. **Tool Use:** supports tool calling by the voice agent. **Realistic Audio:** evaluates under realistic audio conditions with perturbations of user acoustic or behavioral conditions (e.g., accent, background noise, disfluencies). **Simulator Validation:** applies runtime quality control to the conversation simulation before scoring. **Multi-Trial:** reports performance aggregated across multiple trials per scenario, rather than a single trial. $\sim$ denotes partial support, and — for Simulator Validation means the criterion does not apply due to the absence of a simulator. EVA-Bench is the only framework combining live multi-turn simulation across both speech-to-speech (S2S) and cascade architectures with a realistic audio environment, automated simulator validation, and pass@k measurement via multi-trial consistency measurement.

Framework	Interaction	Supported Architectures	Multi-turn	Tool Use	Realistic Audio	Simulator Validation	Multi Trial
EVA-Bench	Live bot-to-bot	S2S, Cascade	✓	✓	✓	✓	✓
τ -Voice (Ray et al., 2026)	Live bot-to-bot	S2S	✓	✓	✓	✗	✗
FDB-v3 (Lin et al., 2026a)	Real human audio	S2S, Cascade	✓	✓	$\sim$	—	✗
VoiceAgentBench (Jain et al., 2025)	Static	S2S, Cascade	✓	✓	✗	—	✗
CAVA (Held et al., 2025)	Partial simulation	S2S, Cascade	✓	✓	✗	✗	✗
FDB-v2 (Lin et al., 2026b)	Live, auto examiner	S2S	✓	✗	✗	✗	✗
FD-Bench (Lin et al., 2025)	Live, simulated	S2S	✓	✗	✗	✗	✗

of dimensions than existing benchmarks provide. Additionally, voice agents are not architecturally uniform: cascade systems chain separate speech-to-text (STT), large language model (LLM), and text-to-speech (TTS) components, while audio-native systems process audio inputs directly, either end-to-end via speech-to-speech (S2S) models, or via hybrid systems that pair a large audio language model (LALM) with a TTS model (full definitions in Appendix A). These architectures have fundamentally different mechanisms, yet must be evaluated on equal footing for benchmarks to meaningfully compare them.

We present EVA-Bench, a benchmark designed to solve both of these challenges. On the simulation side, EVA-Bench conducts fully automated bot-to-bot audio simulation over dynamic multi-turn dialogues, with validation-gated quality control ensuring consistency across repeated trials. On the measurement side, EVA-Bench introduces two primary composite scores: EVA-A (Accuracy) and EVA-X (Experience). EVA-A captures task completion, faithfulness to policy and context, and speech fidelity of key entities. EVA-X captures conversation progression, conciseness for spoken delivery, and turn-taking timing. Diagnostic metrics are also available for further analysis. Across 12 evaluated systems, EVA-Bench reveals that accuracy and experience remain jointly unsatisfied, that peak capability consistently overstates reliable performance, and that robustness to acoustic perturbations varies substantially across systems and

metrics. Our contributions:

- We introduce EVA-Bench: an end-to-end evaluation framework for voice agents that generates realistic bot-to-bot conversations through a user simulator with validation-gated quality control, and supports controlled acoustic perturbations across independent trials.
- We define EVA-A and EVA-X, joint accuracy and experience metrics that surface failure modes invisible to existing benchmarks and enable direct comparison between audio-native and cascade voice agents.
- We create three enterprise benchmark datasets with a total of 213 scenarios focused on surfacing voice-specific failure modes.
- We show empirical findings on cascade vs. audio-native tradeoffs, perturbation sensitivity, and behavioral consistency across trials.

2 Related Work

Many existing voice benchmarks focus on individual components such as audio understanding (Yang et al., 2021; Ao et al., 2024; Cui et al., 2025), TTS quality (Huang et al., 2022; Manku et al., 2025; Huang et al., 2025), or conversational dynamics (Peng et al., 2025; Arora et al., 2025), rather than end-to-end voice agent behavior. We summarize the key gaps below.

Conversation Simulation. Effective voice agent evaluation requires live, task-oriented multi-turn conversations, yet most benchmarks fall short: FDB-v1 and FDB-v1.5 (Lin et al., 2025,




Figure 1: **EVA-Bench framework overview.** The simulation orchestrates parallel per-scenario bot-to-bot audio sessions over WebSocket in which the User Simulator — configured with a scenario-specific goal, persona, and conversational TTS voice — interacts with the Voice Agent under test. The Tool Executor handles all agent tool calls deterministically. Completed conversations pass through **Simulator Validation** that trigger automatic regeneration on failure before entering the **Quality Measurements** phase, which produces **EVA-A** and **EVA-X** pass@1, pass@k, and pass^k scores in addition to **Diagnostic** metrics.

2026c) are heavily scripted without tool use, VoiceAgentBench (Jain et al., 2025) uses static TTS-synthesized queries, and FDB-v3 (Lin et al., 2026a) remains single-turn despite improved realism. The closest prior work, τ -Voice (Ray et al., 2026) and FDB-v2 (Lin et al., 2026b), conduct live bot-to-bot multi-turn simulation but lack automated validation of simulator behavior across trials, leaving open whether scores reflect agent quality or simulator issues. Additionally, τ -Voice couples accent variation with persona changes, confounding acoustic and behavioral effects. EVA-Bench addresses these gaps with multi-trial bot-to-bot simulation, automatic simulator validation, and a controlled perturbation suite that isolates each factor independently. Table 1 provides a feature-level comparison.

Voice Agent Quality Measurement. Existing benchmarks converge on a narrow set of metrics. VoiceAgentBench (Jain et al., 2025) assesses tool use accuracy but no conversational quality; τ -Voice (Ray et al., 2026) adds turn-taking measures but not faithfulness; FDB-v3 (Lin et al., 2026a) introduces transcript-level response quality but omits policy faithfulness and audio-level entity accuracy. No existing framework measures whether the agent makes efficient conversational progress, avoids excessive cognitive load, or speaks correct information, all of which are consequential failure modes for enterprise deployment.

3 Methodology

3.1 Conversation Simulation

Data Design. Constructing a benchmark dataset well-suited to voice agent evaluation requires careful attention to both domain relevance and sce-

nario specificity. EVA-Bench comprises three domains reflecting real-world enterprise voice agent deployments: Airline Customer Service Management (CSM), Healthcare Human Resources Service Delivery (HRSD), and Enterprise Information Technology Service Management (ITSM). Scenarios within each domain are designed to reflect the task-oriented nature of real voice agent interactions, focusing on high-contact cases where users are most likely to call an agent, such as flight rebooking rather than initial booking. Each scenario consists of a *user goal* specifying the user’s intended outcome with explicit constraints (e.g., departure before 10pm, fare below a specified amount), a *user persona* defining speaking style, patience, and personality, a *scenario database* containing the data the agent’s tools query and modify, and a *ground truth* specifying the expected database state after successful task resolution. User goals are accompanied by a decision tree that eliminates ambiguity about intended outcomes and user choices throughout the conversation, enabling repeatable evaluation. Scenarios are further designed to surface voice-specific failure modes by requiring agents to correctly handle key entities (e.g. confirmation codes, names, and domain-specific identifiers) that are frequently misheard in spoken interactions. Scenarios were constructed through a joint generation pipeline and validated through human reviewers and frontier-model stress testing in text-only mode; more details on data domains, scenario examples, and dataset construction and validation can be found in Appendix C.

Multi-Turn Conversations. EVA-Bench evaluates agents through fully automated bot-to-bot conversations. A user simulator, built on a high-

quality cascade pipeline, receives the user goal, decision tree, and persona as input and communicates with the agent over a live audio WebSocket (Figure 1). Both sides of the interaction operate over audio, enabling evaluation of cascade and audio-native architectures under identical conditions. To establish that the user simulator produces realistic and correct behaviors, we validated it directly with seven human reviewers. They read the transcript and listened to the audio for 50 conversations in the Airline CSM domain, rating three dimensions: content realism (whether a real user would plausibly say this), audio naturalness (whether the voice sounds human), and goal-correctness (whether the user remained faithful to its assigned goal). All 50 conversations passed on all three dimensions. See Appendix D for full simulator details.

Controlled Perturbations. EVA-Bench introduces a perturbation suite that varies user acoustic and behavioral conditions independently. Acoustic perturbations include accent variations, background noises, and connection degradation. Behavioral perturbations model caller variation in personality and speaking style. Each perturbation axis is independently controlled, enabling conditions to be applied in isolation or combination to disentangle each factor’s effect on performance. See Appendix G for more details.

Runtime Simulation Validation. Before any evaluation metrics are computed, each simulated conversation passes through automated validation checks. User Behavioral Fidelity (LLM-as-Judge (Zheng et al., 2023)) checks whether the user simulator faithfully executed its assigned goal without deviations that would corrupt agent evaluation. The judge prompt contains specific corruption types to check for. User Speech Fidelity uses a LALM-as-Judge to verify that the simulator’s spoken audio accurately conveyed its intended content, using a nearly identical prompt to the Speech Fidelity judge explained in 3.2.1. Conversations failing any check are automatically regenerated, ensuring that evaluation scores reflect agent behavior rather than simulator artifacts. Across four systems evaluated on all domains, 12.0% of trials required regeneration due to user simulator error (almost exclusively due to user behavioral drift), with speech fidelity accounting for less than 4% of reruns. Full validation details, including judge selection methodology and per-check rerun breakdowns, are provided in Appendix D.

3.2 Voice Agent Quality Measurement

EVA-Bench evaluates each conversation across three layered metric categories: **Accuracy (EVA-A)**, **Experience (EVA-X)**, and **Diagnostic Metrics**. These are described in the following subsections, and a table summarizing all metrics is provided in Appendix E.

Architecture-Aware Implementations. Because cascade, hybrid, and S2S systems divide responsibility for speech understanding and generation differently, we score the reasoning-oriented metrics (Faithfulness, Conciseness, Conversation Progression) against whichever component is responsible for a given capability in each architecture. In cascade systems, the LLM alone is responsible for reasoning: it never sees or produces raw audio, so it is scored on the STT transcript it read and the text it wrote before TTS. In S2S and hybrid systems, the model is responsible for audio understanding directly, so the assistant is judged against the user’s intended text rather than a transcript, since these systems have no separable intermediate text on the input side. For S2S, the agent’s output is likewise scored from a post-hoc transcript of the model’s own audio, since there is no separable intermediate text on the output side either.

LLM Judge Development. Judges were developed through a five-stage process: (1) metric definition and rating scale construction, (2) judge prompt construction, (3) development set construction via synthetic generation and multi-model consensus labeling with human adjudication on disagreements, (4) prompt refinement guided by disagreement analysis, and (5) judge model selection based on accuracy and macro-F1 against the development set (see Appendix E.2 for more details). We further validated our judges through a human agreement study detailed in Section 4.2.

3.2.1 EVA-A: Accuracy Metrics

Task completion alone is a necessary but insufficient measure of accuracy. An agent can reach the correct end state while fabricating a policy detail, misreading a confirmation code aloud, or proceeding without required confirmations. Below are the metrics we propose to measure Accuracy.

Task Completion. A deterministic binary metric comparing the SHA-256 hash of the scenario database’s final state against the ground-truth state. A score of 1 indicates the agent made exactly the correct tool calls with correct parameters; 0 indicates any deviation, i.e. wrong, missing, or extra

changes. Because the user simulator produces repeatable outcomes, failures are unambiguously attributable to agent error. See Appendix E.3.1 for more details.

Faithfulness. An LLM-as-Judge metric evaluating whether the agent’s actions remain grounded in the instructions, policies, tool results, and user inputs. This complements Task Completion: high Task Completion with low Faithfulness indicates the task was completed but with material errors along the way (e.g., misrepresenting fees). See Appendix E.3.2 for more details.

Speech Fidelity. A LALM-as-Judge metric evaluating whether the agent’s spoken audio accurately reproduces the intended text, with particular attention to high-stakes named entities (e.g. confirmation codes, dates, dollar amounts). For S2S systems where no intended text exists, the metric instead verifies that key entities from user turns and tool responses are correctly spoken. To our knowledge, this is the only metric in any end-to-end voice agent benchmark that evaluates the quality of the agent’s spoken output at the audio level. See Appendix E.3.3 for more details.

3.2.2 EVA-X: Experience Metrics

The quality of a conversational experience with a voice agent is shaped by several key factors: whether the conversation moves purposefully toward resolution, whether responses are concise enough to follow without replay, and whether the timing of the agent’s replies feels natural.

Conversation Progression. An LLM-as-Judge metric that evaluates whether the agent efficiently moves the conversation forward by avoiding repetition, retaining context across turns, and driving toward task resolution without stalling or backtracking. See Appendix E.4.1 for more details.

Conciseness. An LLM-as-Judge metric that evaluates whether the agent’s responses are appropriately brief for spoken delivery. Phone callers cannot skim or re-read long responses; verbose agents fail users when they impose cognitive overload by providing too many details or questions. See Appendix E.4.2 for more details.

Turn-Taking. A deterministic timestamp-based metric measuring whether the agent spoke at the right time, neither interrupting the user nor introducing excessive silence. Each turn is routed to a semantically appropriate scoring function: agent-interrupted turns are scored on overlap duration, barge-in count, and post-interrupt recovery latency;

user-interruption turns on agent yield latency; and uninterrupted turns on a piecewise-linear latency curve. Turns involving tool calls receive a more lenient latency threshold, reflecting a longer expected duration than a purely conversational turn. This metric also takes into account when an agent fails to respond to a user turn and the conversation hangs. See Appendix E.4.3 for more details.

3.2.3 Diagnostic Metrics

Diagnostic metrics are not included in EVA-A or EVA-X scores. Their purpose is to make main metric failures actionable by providing more granular information on key failure areas. For example, Transcription Accuracy (Key Entities) is an LLM-as-Judge diagnostic metric that identifies domain-specific key entities in user speech (confirmation codes, names, dates, IDs) and verifies whether each was correctly transcribed in cascade systems using semantic rather than exact match. This surfaces failures that word error rate (WER) misses entirely: a confirmation code off by one character scores near-perfect on WER but is functionally unusable. Additional diagnostic metrics cover authentication outcomes, response latency, etc. (complete list provided in Appendix E.5).

3.2.4 Aggregate Metrics

Metrics for each dimension are aggregated into per-dimension scores (EVA-A, EVA-X), designed to capture both average and consistent performance. Measuring consistency requires a binary notion of success per conversation, so that we can assess how often a system succeeds across repeated trials of the same scenario. Simple averaging is problematic as it can mask major failures on individual metrics.

Pass Criteria. We define a *pass threshold* τ_m for each metric m , calibrated to the point at which performance is acceptable given the metric scale and implementation (Appendix E). A conversation passes on a dimension if *every* metric meets its threshold. Concretely, a conversation passes on accuracy if $(\text{Task Completion} = 1.0) \wedge (\text{Faithfulness} \geq 0.5) \wedge (\text{Speech Fidelity} \geq 0.95)$ and passes on experience if $(\text{Turn-Taking} \geq 0.8) \wedge (\text{Conversation Progression} \geq 0.5) \wedge (\text{Conciseness} \geq 0.5)$.

Pass Metric Definitions. The binary pass/fail scoring yields three aggregate statistics, each reported as EVA-A and EVA-X variants. **pass@1** is the fraction of $T = Nk$ trials (N scenarios, k trials each) that pass, measuring average performance. **pass@k** is the fraction of scenarios where

Table 2: **Human-judge inter-annotator agreement (κ , Spearman ρ) for the judge (IAA-J) versus the human linguist inter-annotator baseline (IAA-H).** Quadratic-weighted κ is used for the 1–3 ordinal metrics; unweighted κ for the binary Speech Fidelity metric. 95% CIs from 10,000 conversation-level bootstrap resamples. To verify that the pooled pairing design does not distort agreement estimates, we compute κ using a single randomly selected annotator per record; results fall within 0.007 of the pooled estimates across all metrics. IAA-J κ ranges from 0.777 to 0.845 across the 4 metrics; Spearman ρ agrees with κ within 0.008 for every metric, indicating no systematic calibration bias. Note that IAA-H is a practical rather than strict ceiling: when two human annotators disagree and the judge agrees with one, IAA-J is inflated relative to IAA-H.

Metric	IAA-J κ	95% CI	IAA-J ρ	IAA-J _{rand} κ	95% CI	IAA-H κ	95% CI	IAA-H ρ
Faithfulness	0.836	[0.729, 0.915]	0.844	0.832	[0.697, 0.932]	0.740	[0.566, 0.870]	0.756
Conv. Progression	0.845	[0.753, 0.911]	0.843	0.841	[0.724, 0.931]	0.769	[0.627, 0.875]	0.782
Conciseness	0.823	[0.754, 0.874]	0.826	0.823	[0.745, 0.883]	0.825	[0.749, 0.881]	0.825
Speech Fidelity	0.777	[0.704, 0.835]	0.781	0.770	[0.683, 0.846]	0.754	[0.685, 0.817]	0.754

at least one of k trials passes, measuring ceiling performance. **pass^k** measures reliability, by raising each scenario’s pass rate $\hat{p}_i$ to the k -th power and averaging across all N scenarios. This captures the probability that the system passes all k independent trials in a given scenario. The difference between **pass@k** and **pass^k** quantifies the gap between ceiling (peak) and consistent (reliable) performance. Formal definitions are provided in Appendix A.

4 Empirical Analysis

4.1 Experiment Setup

We evaluated 12 systems in total: seven cascade, two hybrid, and three S2S. Configuration details are provided in Appendix B. Under the clean (unperturbed) condition, systems were evaluated on all 213 scenarios with $k = 5$ trials per scenario. We additionally evaluated under three perturbation conditions: French-accented user speech, coffee shop background noise, and both combined. To maintain feasibility across 12 systems and three conditions, perturbed evaluations used a randomly sampled subset of 90 scenarios (30 per domain) with $k = 3$ trials per scenario. The same subset was used across all systems. GPT-Realtime and Gemini Live models were evaluated using their native SDK (OpenAI, 2024; Google DeepMind, 2024), Scribe-v2.2-Realtime + Gemini-3-Flash + TTS-Conversational-v3-Alpha was evaluated using ElevenAgents (ElevenLabs, 2023), and the remaining systems were evaluated using Pipecat (Daily, 2024). All systems were evaluated using the respective framework’s default turn detection configuration.

4.2 Evaluation Reliability

EVA-Bench scores reflect agent behavior, not measurement noise. EVA-Bench scores reflect genuine differences in agent behavior rather than evaluation artifacts: judge stochasticity contributes minimally to observed variance, and score differences between systems exceed measurement noise.

We substantiated this with two sets of analyses: human-judge agreement to establish judge validity, and variance decomposition to complement bootstrap confidence intervals (CIs) in characterizing and bounding relative measurement noise.

Judge agreement meets human IAA across metrics. We measured inter-annotator agreement (IAA) between the LLM judge and/or human annotators using quadratic-weighted Cohen’s κ . Our reference baseline was the IAA of two independent annotations from human linguist labelers (IAA-H). For judge inter-annotator agreement (IAA-J), each human annotation was paired with the judge score and pooled, with 95% bootstrap confidence intervals (10,000 resamples) computed at the conversation level. The judge met the human IAA on all four metrics, with IAA-J κ ranging from 0.777 to 0.845 across metrics. For every metric, Spearman’s ρ and κ (both computed on the same judge–human ratings) differed by at most 0.008, supporting the absence of systematic calibration bias. Agreement numbers are in Table 2 and more details are shown in Appendix E.2.

Pass/fail thresholds are validated by binary IAA and threshold sensitivity analyses. We measured IAA on the binary pass/fail scores in addition to the full ordinal scale. Binarizing is meaningful here because annotators were told upfront which ratings mapped to pass versus fail. On this dichotomization, IAA-J was consistently close to the IAA-H baseline, differing by at most 0.11 across metrics. IAA-J agreement reached "substantial" or "near perfect agreement" by standard conventions for Cohen’s κ (Landis and Koch, 1977) on every metric; IAA-H did the same on all but Conversation Progression, where two human labelers agree moderately (0.553) (agreement scores in Table 3). This supports the finding that the chosen pass/fail thresholds for each metric are well-calibrated against human judgment. To further validate our pass/fail thresholds, we examined score distribu-

Table 3: **Human–judge inter-annotator agreement (IAA) on the pass/fail thresholds, compared against the human–human agreement baseline.** 95% CIs from 10,000 record-level bootstrap resamples.

Metric	IAA-J κ	95% CI	IAA-H κ
Faithfulness	0.860	[0.754, 0.947]	0.731
Conv. Progression	0.664	[0.492, 0.805]	0.553
Conciseness	0.738	[0.651, 0.809]	0.741
Speech Fidelity	0.777	[0.704, 0.835]	0.754

tions and conduct threshold sensitivity sweeps (Appendix E), which further confirm that pass rates are not artifacts of scores clustering near the cutoff.

Score variance reflects trial behavior, not judge noise. Observed metric scores for a given model within a domain reflect variance from three sources: scenario difficulty, trial stochasticity (conversation trajectories), and LLM judge stochasticity. We used several complementary analyses on a subset of systems to characterize the contribution of each source to total variance. These analyses demonstrated that our main findings reflect genuine differences in model behavior rather than measurement noise.

In linear mixed-effects models fit per metric and model, trial stochasticity was the dominant source of variance, consistently exceeding scenario-level variance: it accounted for 40.4–80.4% of total variance on judge-graded metrics and 52.7–100.0% on deterministic metrics, versus only 5.5–37.5% for judge stochasticity and 0.0–47.3 % for scenario variance (Table 35). A separate permutation test on trial and judge standard deviations confirmed that trial-level variability significantly exceeds judge-level variability on every judge-graded model $\times$ metric combination tested ($p < 0.0001$ for all 16 combinations; Table 37). Low judge variance indicates low measurement noise and validates measurement reliability; in contrast, within-scenario trial variance reflects genuine stochastic model behavior, and is the motivation for including multiple trials per scenario. Full variance decomposition results are reported in Appendix H.

4.3 Main Findings

Accuracy–Experience Frontier. No evaluated system cleared 0.5 on both EVA-A and EVA-X pass@1 jointly, and only GPT-Realtime-1.5 (0.47, 0.57) cleared 0.4 on both. This is reflected in the pass@1 Pareto frontier (Figure 2), which contains only four systems spanning two disjoint regions.

Within EVA-X, the gap was almost entirely driven by Turn-Taking, where mean scores stratify cleanly by architecture (cascade: 0.28–0.58;




Figure 2: **Accuracy vs Experience.** Accuracy and Experience scores for pass@1, $\pm$ 95% CIs across domains. Four systems are on the Pareto frontier, two S2S and two cascade: Gemini-3.1-Flash-Live, GPT-Realtime-1.5, Scribe + Gemini-3-Flash + Conversational v3, and Nova + GPT-5.4 + Sonic (left to right).

S2S: 0.82–0.83) (Table 4). Conciseness and Conversation Progression showed no comparable separation. The two hybrid systems both fell within the cascade EVA-X range (0.000, 0.029), suggesting that hybrid systems may not inherit the latency advantages of fully speech-to-speech architectures, though more systems would be needed to explore this.

Among the evaluated cascade systems, we observed a consistent accuracy–experience trade-off, pointing to a potential underlying capability–latency tension. The three cascade systems that performed best on accuracy (Nova + GPT-5.4 + Sonic, Scribe + Gemini-3-Flash + Conversational v3, Parakeet + Gemma-4-31B + Kokoro) struggled on experience, with mean latencies on tool-call turns above 5 s. In contrast, the two cascade systems with better experience (Whisper + Qwen3.5-27B + Voxtral, Cohere + Gemma-4-26B + Voxtral) achieved tool-call turn latencies below 2.7 s but also lower accuracy. No cascade system exceeded 0.25 on both dimensions, with no overlapping CIs. Full latency breakdowns appear in Table 27.

Consistency Analysis. Across all 12 evaluated systems, peak performance (pass@k) substantially exceeded reliable performance (pass^k) on both axes: the median gap was 0.44 on EVA-A and 0.24 on EVA-X. Under the pass^k interpretation (the probability of passing all five trials on a given scenario) even the strongest systems fell well below

Table 4: **Accuracy and Experience metrics** for all evaluated systems under clean-audio conditions across 5 trials ($k = 5$), pooled equal-weighted across the three EVA-Bench domains. Each cell shows the pooled point estimate $\pm$ the percentile bootstrap CI half-width ($\alpha = 0.05$). The three pass-rate columns share a single shading scale so they can be visually compared; each submetric column is scaled independently. Darker = higher point estimate.

Arch.	System	EVA-A			Task Completion	Faithfulness	Speech Fidelity
		pass@1	pass@k	pass^k	Mean	Mean	Mean
Cascade	Cohere + Gemma-4-26B + Voxtral	0.207 $\pm$ 0.041	0.416 $\pm$ 0.070	0.060 $\pm$ 0.028	0.338 $\pm$ 0.049	0.375 $\pm$ 0.036	0.983 $\pm$ 0.003
	Scribe + Gemini-3-Flash + Conversational v3	0.490 $\pm$ 0.052	0.730 $\pm$ 0.058	0.269 $\pm$ 0.055	0.736 $\pm$ 0.043	0.457 $\pm$ 0.042	0.977 $\pm$ 0.006
	Ink-whisper + Haiku-4.5 + Sonic 3	0.234 $\pm$ 0.041	0.516 $\pm$ 0.069	0.057 $\pm$ 0.028	0.374 $\pm$ 0.044	0.518 $\pm$ 0.033	0.983 $\pm$ 0.003
	Nova-3 + GPT-5.4 + Sonic 3	0.504 $\pm$ 0.044	0.809 $\pm$ 0.048	0.217 $\pm$ 0.048	0.609 $\pm$ 0.043	0.754 $\pm$ 0.027	0.989 $\pm$ 0.003
	Nova-3 + GPT-5.4-mini + Aura-2	0.210 $\pm$ 0.045	0.448 $\pm$ 0.069	0.062 $\pm$ 0.032	0.465 $\pm$ 0.050	0.270 $\pm$ 0.033	0.974 $\pm$ 0.005
	Parakeet-1.1 + Gemma-4-31B + Kokoro	0.403 $\pm$ 0.045	0.748 $\pm$ 0.055	0.169 $\pm$ 0.046	0.637 $\pm$ 0.051	0.466 $\pm$ 0.035	0.954 $\pm$ 0.009
	Whisper + Qwen3.5-27B + Voxtral	0.205 $\pm$ 0.033	0.518 $\pm$ 0.066	0.033 $\pm$ 0.019	0.417 $\pm$ 0.051	0.546 $\pm$ 0.033	0.913 $\pm$ 0.010
Hybrid	Gemini-3-Flash + Gemini-3.1-Flash-TTS	0.431 $\pm$ 0.047	0.812 $\pm$ 0.055	0.158 $\pm$ 0.043	0.674 $\pm$ 0.041	0.443 $\pm$ 0.036	0.969 $\pm$ 0.006
	Ultravox-Realtime	0.270 $\pm$ 0.047	0.503 $\pm$ 0.072	0.108 $\pm$ 0.037	0.473 $\pm$ 0.055	0.292 $\pm$ 0.035	0.971 $\pm$ 0.007
S2S	Gemini-3.1-Flash-Live	0.292 $\pm$ 0.048	0.552 $\pm$ 0.069	0.132 $\pm$ 0.043	0.473 $\pm$ 0.052	0.238 $\pm$ 0.035	0.995 $\pm$ 0.003
	GPT-Realtime-1.5	0.467 $\pm$ 0.052	0.710 $\pm$ 0.061	0.283 $\pm$ 0.056	0.739 $\pm$ 0.046	0.360 $\pm$ 0.041	0.996 $\pm$ 0.002
	GPT-Realtime-mini	0.163 $\pm$ 0.041	0.318 $\pm$ 0.063	0.059 $\pm$ 0.030	0.345 $\pm$ 0.054	0.125 $\pm$ 0.031	0.977 $\pm$ 0.012
Arch.	System	EVA-X			Turn-Taking	Conciseness	Conv. Progression
		pass@1	pass@k	pass^k	Mean	Mean	Mean
Cascade	Cohere + Gemma-4-26B + Voxtral	0.209 $\pm$ 0.027	0.647 $\pm$ 0.069	0.015 $\pm$ 0.011	0.567 $\pm$ 0.024	0.809 $\pm$ 0.007	0.598 $\pm$ 0.032
	Scribe + Gemini-3-Flash + Conversational v3	0.024 $\pm$ 0.018	0.061 $\pm$ 0.035	0.004 $\pm$ 0.006	0.451 $\pm$ 0.019	0.774 $\pm$ 0.007	0.804 $\pm$ 0.023
	Ink-whisper + Haiku-4.5 + Sonic 3	0.009 $\pm$ 0.006	0.042 $\pm$ 0.031	0.000 $\pm$ 0.000	0.312 $\pm$ 0.020	0.784 $\pm$ 0.007	0.710 $\pm$ 0.023
	Nova-3 + GPT-5.4 + Sonic 3	0.007 $\pm$ 0.006	0.031 $\pm$ 0.024	0.000 $\pm$ 0.000	0.283 $\pm$ 0.019	0.835 $\pm$ 0.007	0.737 $\pm$ 0.020
	Nova-3 + GPT-5.4-mini + Aura-2	0.113 $\pm$ 0.023	0.416 $\pm$ 0.070	0.005 $\pm$ 0.004	0.583 $\pm$ 0.019	0.835 $\pm$ 0.008	0.428 $\pm$ 0.025
	Parakeet-1.1 + Gemma-4-31B + Kokoro	0.010 $\pm$ 0.009	0.035 $\pm$ 0.032	0.000 $\pm$ 0.000	0.308 $\pm$ 0.015	0.829 $\pm$ 0.007	0.774 $\pm$ 0.024
	Whisper + Qwen3.5-27B + Voxtral	0.273 $\pm$ 0.034	0.684 $\pm$ 0.065	0.051 $\pm$ 0.021	0.561 $\pm$ 0.029	0.685 $\pm$ 0.010	0.612 $\pm$ 0.026
Hybrid	Gemini-3-Flash + Gemini-3.1-Flash-TTS	0.000 $\pm$ 0.000	0.000 $\pm$ 0.000	0.000 $\pm$ 0.000	0.019 $\pm$ 0.003	0.801 $\pm$ 0.007	0.618 $\pm$ 0.029
	Ultravox-Realtime	0.029 $\pm$ 0.020	0.081 $\pm$ 0.039	0.006 $\pm$ 0.007	0.417 $\pm$ 0.020	0.750 $\pm$ 0.010	0.429 $\pm$ 0.030
S2S	Gemini-3.1-Flash-Live	0.589 $\pm$ 0.035	0.979 $\pm$ 0.021	0.240 $\pm$ 0.045	0.830 $\pm$ 0.017	0.801 $\pm$ 0.009	0.636 $\pm$ 0.029
	GPT-Realtime-1.5	0.566 $\pm$ 0.039	0.939 $\pm$ 0.034	0.216 $\pm$ 0.040	0.815 $\pm$ 0.013	0.801 $\pm$ 0.008	0.679 $\pm$ 0.024
	GPT-Realtime-mini	0.406 $\pm$ 0.036	0.893 $\pm$ 0.043	0.099 $\pm$ 0.032	0.818 $\pm$ 0.015	0.722 $\pm$ 0.009	0.388 $\pm$ 0.033

peak, suggesting that single-trial scores systematically overstate deployment-grade reliability.

Robustness Analysis. We evaluated each system under three perturbation conditions (**accented speech**, **background noise**, and **both combined**) measuring performance degradation against a clean baseline across the 90 subsampled scenarios (see subsection 4.1) via paired sign-flip permutation tests ($n=90$, Holm–Bonferroni corrected). The architectures diverged in their failure modes: cascade systems were vulnerable on both accuracy and experience metrics, while S2S systems suffered primarily on experience metrics. **Accented speech** drove the largest accuracy failures: cascade Task Completion dropped by a mean 10 points (worst system: 17 points), while no S2S model showed significant accuracy degradation (0/27 model-metric pairs). **Background noise** exposed S2S experience failures (EVA-X $\bar{\Delta}=-0.16$, Turn-Taking $\bar{\Delta}=-0.076$); cascade EVA-X was largely floored (baselines 0.007–0.273) and showed little further room to drop, but cascade Turn-Taking did degrade substantially under noise ($\bar{\Delta}=-0.12$). Cascade accuracy also suffered under noise (Task Completion $\bar{\Delta}=-0.10$ vs. S2S $\bar{\Delta}=-0.04$). The **combined** condition revealed the full spread: cascade Task Completion dropped a mean 19 points, while the S2S model mean re-

mained within 5 points. Within cascades, however, robustness varied considerably, ranging from 11–87% of per-system metric-perturbation combinations showing significant degradation. The two most robust cascade systems degraded primarily on experience metrics, more closely resembling S2S systems than other cascades. Turn-Taking was the most perturbation-sensitive metric, with 81% of measurements across perturbations and systems showing significant degradation. See Figure 3 and Appendix G.

4.4 Failure Mode Analysis

EVA-Bench’s metrics enable detailed error analysis. We report four targeted analyses below, with additional analyses in Appendix F.

Named entity transcription as an accuracy bottleneck. EVA-Bench’s diagnostic metrics revealed a consistent correlation between Transcription Accuracy (Key Entities) and Task Completion across all three domains, pointing to transcription as a candidate bottleneck for cascade performance. Across all cascade systems, mean transcription accuracy on key entities was strongly correlated with mean Task Completion (Pearson $r = 0.93$, $p = 0.002$) and the relationship held within each domain ($r = 0.88$ – 0.93 , all $p < 0.01$). Cascade systems with transcription accuracy below




Figure 3: Perturbation effect on **Turn-Taking** pooled across domains. Bars show the mean delta from clean trials; 95% bootstrap CIs on the per-scenario delta. Bar colors encode the perturbation condition: ■ accent, ■ background noise, ■ accent + background noise. Asterisks mark cells significant after Holm-Bonferroni correction ($* p < 0.05$, $** p < 0.01$, $*** p < 0.001$). Models, left to right: **Cascade**: Cohere + Gemma-4-26B + Voxtral, ElevenAgents (Scribe + Gemini-3-Flash + Conversational v3), Ink + Haiku-4.5 + Sonic, Nova + GPT-5.4 + Sonic, Nova + GPT-5.4-mini + Aura, Parakeet + Gemma-4-31B + Kokoro, Whisper + Qwen3.5-27B + Voxtral; **Hybrid**: Gemini-3-Flash + Gemini-3.1-Flash, Ultravox; **S2S**: Gemini-3.1-Flash-Live, GPT-Realtime-1.5, GPT-Realtime-mini.

70% on key entities showed Task Completion rates 39% lower than systems above this threshold (0.37 vs 0.60). See Figure 5 for STT Transcription Accuracy (Key Entities), Figure 6 for correlation with Task Completion, and Appendix F.1 for more details.

Evaluated S2S systems show contrasting patterns relative to other systems. The higher EVA-X scores for S2S systems reflect a clear Turn-Taking lead with on-time rate and conversation completion up +27.9 pp and +15.2 pp on average respectively, though the strongest cascade systems came within a few percentage points. The conversation completion gap is partly explained by many cascade systems often failing to respond to short utterances and spelled content (Table 25). S2S systems were flagged with policy violations in more conversations on average (+24.6 pp), though this comparison is potentially confounded as systems that fail to reach completion have fewer occasions to violate a policy. (see Appendix F.2 for details).

Speech fidelity failures concentrate on alphanumeric entities. Entity errors (character substitutions, omissions, and phonetic confusions) were the dominant failure mode across all evaluated models, motivating Speech Fidelity as an audio-level metric: a caller who hears an incorrect confirmation code cannot detect the error from context alone. See Appendix F.3 for examples.

5 Conclusion

We presented EVA-Bench, an end-to-end evaluation framework for voice agents that jointly addresses simulation fidelity and measurement comprehensiveness. Together, validation-gated bot-to-bot simulation, architecture-agnostic composite metrics (EVA-A and EVA-X), and a multi-trial

consistency framework (pass@1, pass@k, pass^k) enable comparison across cascade, hybrid, and S2S voice agents under comparable conditions.

Our evaluation of 12 systems produced three central findings:

(1) While the best-performing cascade and S2S systems achieved comparable accuracy, experience quality diverges sharply along architecture lines, with the S2S–cascade gap on EVA-X driven almost entirely by turn-taking.

(2) Peak and reliable performance diverged substantially across all systems: the median pass@k–pass^k gap is 0.44 on EVA-A, indicating that single-trial evaluation scores systematically overstate deployment-grade quality regardless of architecture.

(3) We observed that cascade and S2S systems degrade differently under acoustic perturbation.

These observations highlight dimensions of voice agent quality that existing benchmarks leave unmeasured: joint accuracy–experience measurement, peak-versus-reliable performance gaps, and perturbation-specific degradation patterns across voice agent architectures.

EVA-Bench is released as a fully open-source, extensible framework. The metric suite is modular, and pass criteria and perturbation conditions are configurable, allowing practitioners to adapt the benchmark to their deployment requirements. New domains can be added by configuring agents and scenarios. We invite the community to extend the benchmark with additional domains, languages, and evaluation dimensions as the field evolves.

Limitations

While EVA-Bench provides a rigorous, extensible evaluation framework for conversational voice

agents, several limitations should be acknowledged.

Metrics LLM-based judges may exhibit stylistic biases, systematically favoring outputs resembling their training distribution. This risk may be higher when evaluated systems share a model family with the judge. Moreover, LALM-based judges remain an emerging paradigm and may exhibit lower reliability than text-only judges under equivalent conditions. In addition, the current scoring scheme applies binary task completion judgments with no partial credit for agents that satisfy all but one sub-goal, potentially suppressing fine-grained discriminability between systems differing primarily in terminal sub-goal robustness.

Framework EVA-Bench assumes that bot-to-bot simulation with a validated user simulator constitutes a valid proxy for real caller interactions. If the simulator systematically diverges from real callers in ways our behavioral fidelity validator does not detect, scores may not transfer to production conditions. PCM-to- μ -law audio conversion introduces quality degradation, and bot-to-bot interface timing may not fully represent production deployments. Imprecise pipeline event timestamps (VAD events, log reconciliation across systems) may lead to inaccurate response speed values. Full reproduction requires access to commercial model APIs. Generating $k=5$ validated trials per scenario carries non-trivial cost; a full run across all 213 scenarios and 12 systems requires several hundred model calls per trial. All tools use a declarative mock executor rather than live APIs; while schemas reflect real-world APIs, mocks cannot capture failure modes, latency variance, partial responses, or schema drift. Production latency-reduction strategies (speculative execution, response pre-fetching, streaming-aware scheduling) are not used; reported latency metrics may not reflect achievable production lower bounds.

Simulation All scenarios are currently English-only. The user simulator is built on a high-quality STT-LLM-TTS pipeline (scribe-v2.2-realtime + GPT-5.1 + ElevenLabs v3 Conversational) which may therefore under-represent natural disfluencies. The simulator also does not systematically generate interruptions, meaning EVA-Bench under-stresses turn-taking robustness and cannot fully differentiate systems that handle barge-in gracefully. As the simulator relies on a commercial system, its behavior may change across versions. Lastly, while

we employ validators to detect off-policy behavior, perfect adherence cannot be guaranteed.

Experiments Our evaluation covered 12 systems across three architectural classes, but this sample size precludes broad inferential claims about model classes. All frameworks used default turn detection settings, and system prompts were not optimized. Both were deliberate choices to preserve fair baselines. Targeted tuning would likely yield higher scores. Different latency optimizations could also be explored for cascade pipelines, such as pre-tool speech, streamed LLM responses, and speculative decoding; our framework supports configuring these, but they were not included in these results. Robustness experiments include only one accent (French) and one background noise condition (coffee shop). The provider’s voice library does not consistently offer the same voice across a range of accents, and applying a custom accent modification to an existing voice not trained for it carries uncertain reliability. We therefore selected a dedicated voice ID chosen to represent the French accent authentically, rather than modifying the baseline voice, and informally validated the chosen voice with native French speakers for realism. As a result, the accent condition necessarily confounds accent with voice identity, and observed differences may partially reflect properties of this particular voice or noise condition rather than accent or noise variation broadly.

Our plans for future work and addressing certain limitations are iteratively outlined and addressed on our website ².

Ethics Statement

Our work focuses on responsible development and evaluation of voice agent systems. All evaluation scenarios in EVA-Bench are fully synthetic — no real caller data, recordings, or personally identifiable information were used in dataset construction or evaluation.

All simulation tools respect the copyright and privacy guidelines. Although EVA-Bench enables large-scale evaluation of voice agents, we cannot guarantee that models evaluated through it will not generate harmful or biased output during evaluation runs. Researchers and practitioners are strongly encouraged to implement appropriate content filtering and bias detection before deploying evaluated systems in production environments.

²<https://servicenow.github.io/eva/#limitations>

Additionally, we acknowledge that our current domain coverage is limited, and the simulation is conducted exclusively in English, which may inadvertently reinforce existing representational biases in audio AI systems. We encourage the community to expand EVA-Bench with more diverse languages, accents, cultural contexts, and domains. We are actively working on extending it to other languages.

Reproducibility Statement

We are committed to full reproducibility of our evaluation framework and experimental results. All EVA-Bench code, configuration files, evaluation scripts, scenario data, and documentation are publicly released as open-source under a repository which includes setup instructions, environment specifications, and scripts to reproduce all reported evaluations (<https://github.com/ServiceNow/eva>).

We provide comprehensive implementation details including all model configurations (Appendix B), data distributions across domains (Appendix C), judge prompts (Appendix K), and metric definitions and thresholds (Appendix E). All LLM-as-judge configurations are fully specified to enable result replication across research groups. We note that full reproduction of the reported scores requires access to commercial model APIs; results may vary across API versions and deployment configurations, as discussed in limitations.

Acknowledgments

We thank the ServiceNow Machine Learning Data Linguists for their careful work reviewing the data scenarios, providing feedback on domain design, and annotating conversation samples: Tiffany Do, Ryan Dux, Maria Kossenko, Keerthana Gopinathan, Anne Heaton-Dunlap, Nidhi Kumari, Ranjani Iyer. We are also grateful to Sai Rajeswar, Masoud Hashemi, Vikas Yadav, and Chris Pal for their feedback and comments on this work. Finally, we thank Joyce Li and Nitin Aggarwal for their product insights and guidance on enterprise use case requirements.

References

Miguel E. Andres, Vadim Fedorov, Rida Sadek, Enric Spagnolo-Arrizabalaga, and Nadescha Trudel. 2026. [Testing the Testers: Human-Driven Quality Assessment of Voice AI Testing Platforms](#).
Junyi Ao, Yuancheng Wang, Xiaohai Tian, Dekun Chen,

Jun Zhang, Lu Lu, Yuxuan Wang, Haizhou Li, and Zhizheng Wu. 2024. [SD-Eval: A Benchmark Dataset for Spoken Dialogue Understanding Beyond Words](#). *Advances in Neural Information Processing Systems*, 37:56898–56918.

Siddhant Arora, Zhiyun Lu, Chung-Cheng Chiu, Ruoming Pang, and Shinji Watanabe. 2025. [Talking Turns: Benchmarking Audio Foundation Models on Turn-Taking Dynamics](#). In *International Conference on Learning Representations*, volume 2025, pages 52754–52781.

Hongliu Cao, Ilias Driouich, and Eoin Thomas. 2026. [Beyond Task Completion: Revealing Corrupt Success in LLM Agents through Procedure-Aware Evaluation](#).

Yiming Chen, Xianghu Yue, Chen Zhang, Xiaoxue Gao, Robby T Tan, and Haizhou Li. 2026. [VoiceBench: Benchmarking LLM-Based Voice Assistants](#). *Transactions of the Association for Computational Linguistics*, 14:378–398.

Wenqian Cui, Xiaoqi Jiao, Ziqiao Meng, and Irwin King. 2025. [VoxEval: Benchmarking the knowledge understanding capabilities of end-to-end spoken language models](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16735–16753, Vienna, Austria. Association for Computational Linguistics.

Daily. 2024. Pipecat: Open Source Framework for Voice and Multimodal Conversational AI. <https://github.com/pipecat-ai/pipecat>. Accessed: 2026.

ElevenLabs. 2023. ElevenLabs: Documentation. <https://elevenlabs.io/docs/eleven-agents/overview>. Accessed: 2026.

Google DeepMind. 2024. Gemini Live API: Low-Latency Bidirectional Voice and Video Interactions. <https://ai.google.dev/gemini-api/docs/live>. Accessed: 2026.

Will Held, Michael J. Ryan, Aditya Shrivastava, Ali Sartaz Khan, Caleb Ziems, Ella Li, Martijn Bartelds, Michael Sun, Tan Li, Woody Gan, and Diyi Yang. 2025. [CAVA: Comprehensive Assessment of Voice Assistants](#). <https://github.com/SALT-NLP/CAVA>. A benchmark for evaluating large audio models (LAMs) capabilities across six domains: turn taking, instruction following, function calling, tone awareness, safety, and latency.

Mattias Heldner and Jens Edlund. 2010. Pauses, gaps and overlaps in conversations. *Journal of Phonetics*, 38(4):555–568.

Wen-Chin Huang, Erica Cooper, and Tomoki Toda. 2025. [SHEET: A Multi-purpose Open-source Speech Human Evaluation Estimation Toolkit](#). In *Inter-speech 2025*, pages 2355–2359.

- Wen-Chin Huang, Erica Cooper, Yu Tsao, Hsin-Min Wang, Tomoki Toda, and Junichi Yamagishi. 2022. [The VoiceMOS Challenge 2022](#). In *Interspeech 2022*, pages 4536–4540.
- Dhruv Jain, Harshit Shukla, Gautam Rajeev, Ashish Kulkarni, Chandra Khatri, and Shubham Agarwal. 2025. [VoiceAgentBench: Are Voice Assistants ready for agentic tasks?](#)
- J Richard Landis and Gary G Koch. 1977. The measurement of observer agreement for categorical data. *biometrics*, pages 159–174.
- Guan-Ting Lin, Chen Chen, Zhehuai Chen, and Hung yi Lee. 2026a. [Full-Duplex-Bench-v3: Benchmarking Tool Use for Full-Duplex Voice Agents Under Real-World Disfluency](#).
- Guan-Ting Lin, Shih-Yun Shan Kuan, Jiatong Shi, Kai-Wei Chang, Siddhant Arora, Shinji Watanabe, and Hung-yi Lee. 2026b. [Full-duplex-bench-v2: A multi-turn evaluation framework for duplex dialogue systems with an automated examiner](#). In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 27–36, San Diego, California, United States. Association for Computational Linguistics.
- Guan-Ting Lin, Shih-Yun Shan Kuan, Qirui Wang, Jiachen Lian, Tingle Li, Shinji Watanabe, and Hung-yi Lee. 2026c. [Full-Duplex-Bench v1.5: Evaluating Overlap Handling for Full-Duplex Speech Models](#). In *ICASSP 2026-2026 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 19447–19451. IEEE.
- Guan-Ting Lin, Jiachen Lian, Tingle Li, Qirui Wang, Gopala Anumanchipalli, Alexander H Liu, and Hung-yi Lee. 2025. [Full-duplex-bench: A benchmark to evaluate full-duplex spoken dialogue models on turn-taking capabilities](#). In *2025 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 1–8. IEEE.
- Ruskin Raj Manku, Yuzhi Tang, Xingjian Shi, Mu Li, and Alexander Smola. 2025. [EmergentTTS-Eval: Evaluating TTS Models on Complex Prosodic, Expressiveness, and Linguistic Challenges Using Model-as-a-Judge](#). In *Advances in Neural Information Processing Systems*, volume 38. Curran Associates, Inc.
- Sushant Mehta. 2025. [Beyond Accuracy: A Multi-Dimensional Framework for Evaluating Enterprise Agentic AI Systems](#). *arXiv preprint arXiv:2511.14136*.
- Olivia Moore. 2025. [AI Voice Agents: 2025 Update](#). <https://a16z.com/ai-voice-agents-2025-update/>. Andreessen Horowitz.
- OpenAI. 2024. [Realtime API Documentation](#). <https://platform.openai.com/docs/guides/realtime>. Accessed: 2026.
- Shishir G Patil, Huanzhi Mao, Fanjia Yan, Charlie Cheng-Jie Ji, Vishnu Suresh, Ion Stoica, and Joseph E Gonzalez. 2025. [The Berkeley Function Calling Leaderboard \(BFCL\): From Tool Use to Agentic Evaluation of Large Language Models](#). In *Forty-second International Conference on Machine Learning*.
- Yizhou Peng, Yi-Wen Chao, Dianwen Ng, Yukun Ma, Chongjia Ni, Bin Ma, and Eng Siong Chng. 2025. [FD-Bench: A Full-Duplex Benchmarking Pipeline Designed for Full Duplex Spoken Dialogue Systems](#). In *Interspeech 2025*, pages 176–180.
- Bidyapati Pradhan, Surajit Dasgupta, Amit Kumar Saha, Omkar Anustoop, Sriram Puttagunta, Vipul Mittal, and Gopal Sarda. 2025. [SyGra: A Unified Graph-Based Framework for Scalable Generation, Quality Tagging, and Management of Synthetic Data](#).
- Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, and 1 others. 2024. [ToolLLM: Facilitating large language models to master 16000+ real-world APIs](#). In *International Conference on Learning Representations*, volume 2024, pages 9695–9717.
- Soham Ray, Keshav Dhandhanian, Victor Barres, and Karthik R Narasimhan. 2026. [\$\tau\$ -Voice: Benchmarking Full-Duplex Voice Agents on Real-World Domains](#). In *Forty-third International Conference on Machine Learning*.
- Gabriel Skantze. 2021. [Turn-Taking in Conversational Systems and Human-Robot Interaction: A Review](#). *Computer Speech & Language*, 67:101178.
- Tanya Stivers, Nicholas J Enfield, Penelope Brown, Christina Englert, Makoto Hayashi, Trine Heine-mann, Gertie Hoymann, Federico Rossano, Jan Peter De Ruiter, Kyung-Eun Yoon, and 1 others. 2009. Universals and cultural variation in turn-taking in conversation. *Proceedings of the National Academy of Sciences*, 106(26):10587–10592.
- Shu-wen Yang, Po-Han Chi, Yung-Sung Chuang, Cheng-I Jeff Lai, Kushal Lakhotia, Yist Y. Lin, Andy T. Liu, Jiatong Shi, Xuankai Chang, Guan-Ting Lin, Tzu-Hsien Huang, Wei-Cheng Tseng, Ko tik Lee, Da-Rong Liu, Zili Huang, Shuyan Dong, Shang-Wen Li, Shinji Watanabe, Abdelrahman Mohamed, and Hung yi Lee. 2021. [SUPERB: Speech Processing Universal PERFORMANCE Benchmark](#). In *Interspeech 2021*, pages 1194–1198.
- Shunyu Yao, Noah Shinn, Pedram Razavi, and Karthik R Narasimhan. 2025. [\$\tau\$ -bench: A Benchmark for Tool-Agent-User Interaction in Real-World Domains](#). In *The Thirteenth International Conference on Learning Representations*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, and 1 others. 2023. [Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena](#). *Advances in Neural Information Processing Systems*, 36:46595–46623.

Xuhui Zhou, Weiwei Sun, Qianou Ma, Yiqing Xie, Jiarui Liu, Weihua Du, Sean Welleck, Yiming Yang, Graham Neubig, Sherry Tongshuang Wu, and Maarten Sap. 2026. Mind the Sim2Real Gap in User Simulation for Agentic Tasks. *arXiv preprint arXiv:2603.11245*.

A Definitions & Key Terms

STT (*Speech-to-Text*.) A model or service that transcribes spoken audio into text. In voice agent pipelines, STT serves as the input stage, converting user utterances into a form consumable by a downstream language model.

TTS (*Text-to-Speech*.) A model or service that synthesizes natural-sounding speech from text. In voice agent pipelines, TTS serves as the output stage, rendering language model responses as audio delivered to the user.

S2S (*Speech-to-Speech*.) An end-to-end model architecture that accepts raw audio as input and produces audio as output, bypassing discrete STT and TTS stages entirely. S2S systems may also emit text output, used to invoke tool calls, alongside or prior to synthesizing a spoken response. S2S systems minimize modality-crossing latency and preserve paralinguistic signals throughout the pipeline.

LALM. A large language model that accepts audio as a direct input modality—without prior transcription—and produces text output. AudioLLMs internalize acoustic features (prosody, speaker characteristics, noise) that would otherwise be lost in an STT transcription step. Also referred to as an **AudioLLM** or a **SpeechLLM**.

Hybrid. A voice agent architecture that combines an AudioLLM (for audio-aware language understanding) with a discrete TTS module (for speech synthesis). Hybrid systems are Audio-Native on the input side while retaining a text-to-speech output stage, and occupy a middle ground between fully cascaded (STT+LLM+TTS) and fully end-to-end (S2S) pipelines.

Audio-Native. An umbrella term for voice agent architectures that process audio directly at one or more stages, rather than relying solely on text-based STT→LLM→TTS cascades. Audio-Native systems include both S2S and *Hybrid* configurations.

Multi-turn. A conversational interaction comprising more than one exchange between user and agent—i.e., at least one agent response followed by a subsequent user utterance within the same session. Multi-turn evaluation tests an agent’s ability to maintain task state, handle clarifications, and recover from errors across dialogue turns.

Bot-to-Bot. An automated evaluation protocol in which a *user simulator* (a separate agent or model)

plays the role of the human caller, conducting full spoken conversations with the system under evaluation.

pass@1. The fraction of all trials that pass, measuring average performance: $\text{pass}@1 = \frac{1}{T} \sum_{t=1}^T \mathbf{1}[\text{pass}_t^{(d)}]$, where $T = Nk$ is the total number of trials across N scenarios and k trials each.

pass@k. The fraction of scenarios where at least one of k trials passes, measuring ceiling performance: $\text{pass}@k = \frac{1}{N} \sum_{i=1}^N \mathbf{1} \left[\sum_{j=1}^k \mathbf{1}[\text{pass}_{i,j}^{(d)}] \geq 1 \right]$. See also: *Peak*.

pass^k The expected probability that a system passes all k independent future trials on a given scenario, measuring reliable performance: $\text{pass}^k = \frac{1}{N} \sum_{i=1}^N \hat{p}_i^k$, where $\hat{p}_i = \frac{1}{k} \sum_{j=1}^k \mathbf{1}[\text{pass}_{i,j}^{(d)}]$ is the per-scenario pass rate across k trials. See also: *Reliable*.

Peak. Generally refers to $\text{pass}@k$ scores. A peak score characterizes the best-case performance of a system over k independent attempts—i.e., the probability that at least one of k runs produces a correct outcome. Peak scores are informative about a system’s ceiling capability but do not reflect consistency.

Reliable. Generally refers to pass^k scores. A reliability score characterizes the probability that a system produces a correct outcome on *every one* of k independent attempts. Reliability scores penalize variance and are informative about how consistently a system can be expected to succeed across repeated deployments.

Table 5: Self-hosted model configurations.

Model	Model Abrv	Model ID	Type	GPU	CPU	Precision	Deployment
Gemma-4-26B	Gemma-4-26B	google/gemma-4-26B-A4B-it	LLM	2× H100	–	BF16	vLLM
Gemma-4-31B	Gemma-4-31B	google/gemma-4-31B-it	LLM	2× H100	–	BF16	vLLM
Qwen3.5-27B	Qwen3.5-27B	Qwen/Qwen3.5-27B	LLM	4× H100	8× 8GB	BF16	vLLM
Parakeet-1.1	Parakeet	nvdiia/parakeet-ctc-1.1b	STT	1× H100	–	BF16	Nvidia NIM
Whisper-Large-v3	Whisper	openai/whisper-large-v3	STT	1× H100	4× 64GB	FP16	vLLM
Cohere-transcribe	Cohere	CohereLabs/cohere-transcribe-03-2026	STT	1× H100	8× 64GB	BF16	vLLM
Kokoro	Kokoro	hexgrad/Kokoro-82M	TTS	1× H100	8× 32GB	FP32	Remsky Kokoro
Voxtral-4B-TTS	Voxtral	mistralai/Voxtral-4B-TTS-2603	TTS	1× H100	–	BF16	vLLM

B Experiment Parameters

Below are the turn detection and model configurations for all evaluated and judge models; the user simulator is detailed separately in Appendix D. The agent prompts are provided in Appendix J. Finally, all the third-party software packages used are shared in Appendix L.

B.1 Self-Hosted Models

All self-hosted models were served on NVIDIA H100 GPUs. Models served via vLLM used vllm-openai v0.19.0 version, with specific configurations shown in Table 6. Table 5 lists the hardware and serving configurations.

Parameter	Gemma	Qwen
temperature	1.0	1.0
top_p	0.95	0.95
top_k	64	20
max_tokens	12000	—
min_p	—	0.0
presence_penalty	—	1.5
repetition_penalty	—	1.0
enable_thinking	false	false
skip_special_tokens	false	—

Table 6: Configurations for self-hosted models: **Gemma** (Gemma-4-26B, Gemma-4-31B) and **Qwen** (Qwen3.5-27B).

B.2 API-Hosted Models

For ElevenLabs, we used ElevenAgents with the following models: Scribe-v2.2-Realtime, Gemini-3-Flash, and TTS-Conversational-v3-Alpha. We used the default agent parameters, listed in Table 7.

Table 8 lists all the other API-hosted models.

B.3 Turn Detection Configurations

We used the default turn detection configurations for each framework in our experiments. Each framework offers varying levels of configurability, making it difficult to standardize exact parameters and turn strategies across evaluations.

Pipecat. The default start strategy used VAD (voice activity detection) or transcription receipt to determine when the user begins speaking, and the stop strategy uses AI-powered turn detection via

Table 7: ElevenLabs ElevenAgents parameters.

Component	Parameter	Value
STT	Filter background speech	disabled
TTS	Expressive mode	enabled
	Voice	Lauren B — Friendly & Engaging Customer Care Agent
LLM	Voice ID	3liN8q8YoeB9Hk6AboKe
	Temperature	0
	Reasoning effort	minimal
	Limit token usage	—1
	Parallel tool calling	disabled
Tools	Cascade timeout	8 s
	Wait for response	enabled
	Pre-tool speech	auto
	Execution mode	immediate
	Tool call sound	none
Agent	Response timeout	20 s
	Eagerness	normal
	Spelling patience	auto
	Speculative turn	enabled
	Re-transcribe on timeout	disabled
	Take turn after silence	15 s
	End call after silence	disabled
	Max conversation duration	600 s

LocalSmartTurnAnalyzerV3 to determine when the user finishes speaking.³

OpenAI Realtime. We used the default server VAD, which uses periods of silence to detect turn boundaries. Default values were used for threshold, prefix_padding_ms, and silence_duration.⁴

ElevenAgents. The turn “eagerness” parameter was left at its default setting (normal).⁵

Gemini Live. We used the default automatic VAD provided.⁶

EVA-Bench makes turn detection parameters and options configurable via the CLI, so practitioners can run experiments using the turn detection settings available to their chosen framework. The only exception is ElevenAgents, where users must

³<https://docs.pipecat.ai/api-reference/server/utilities/turn-management/user-turn-strategies#base-parameters-2>

⁴<https://developers.openai.com/api/docs/guides/realtime-vad>

⁵<https://elevenlabs.io/docs/eleven-agents/customization/conversation-flow#turn-eagerness>

⁶<https://docs.cloud.google.com/vertex-ai/generative-ai/docs/model-reference/multimodal-live>

Table 8: API-hosted model configurations. Note that Ultravox-Realtime is offered as a realtime model service in Pipecat.

Model	Provider	Type	Model ID	Parameters
GPT-5.2	OpenAI	LLM	gpt-5.2	reasoning: default
GPT-5.4	OpenAI	LLM	gpt-5.4	reasoning: default
GPT-5.4-mini	OpenAI	LLM	gpt-5.4-mini	reasoning: default
GPT-Realtime-1.5	OpenAI	S2S	gpt-realtime-1.5	voice: Marin
GPT-Realtime-mini	OpenAI	S2S	gpt-realtime-mini	voice: Marin
Gemini-3-Flash	Google	LLM	gemini-3-flash-preview	reasoning: default
Gemini-3.1-Flash-Live	Google	LALM	gemini-3.1-flash-live-preview	voice: Leda
Gemini-3.1-Flash-TTS	Google	TTS	gemini-3.1-flash-tts-preview	voice: provider default
Claude Opus 4.6	AWS Bedrock	LLM	us.anthropic.claude-opus-4-6-v1	reasoning: default
Haiku-4.5	AWS Bedrock	LLM	us.anthropic.claude-haiku-4-5-20251001-v1:0	–
Ultravox-Realtime	Ultravox	LALM	–	–
Ink-whisper	Cartesia	STT	ink-whisper	–
Sonic 3	Cartesia	TTS	sonic-3	–
Nova-3	Deepgram	STT	nova-3	–
Aura-2	Deepgram	TTS	aura-2-helena-en	–

register and configure their agents separately prior to evaluation.

B.4 Experiments Compute Resources

Running EVA-Bench experiments involves costs across three distinct components of the pipeline. First, the **user simulator** is powered by Eleven-Labs’ ElevenAgents (Appendix D), a fully hosted proprietary conversational AI system. Second, the **agent under evaluation** introduces inference costs whenever the model is proprietary. Closed-source models are billed per token for LLM inference, and per second of audio for STT and per character for TTS when the agent operates as a full voice pipeline. Note that a conversation will on average last four to five minutes. Third, the **LLM-as-judge** component adds a separate layer of cost, as four of the main metrics rely on model-based scoring. We evaluated 213 scenarios over five trials on clean data and 90 scenarios over three trials per perturbation across three perturbation types, yielding $213 \times 5 + 90 \times 3 \times 3 = 1,875$ evaluation samples per model.

Researchers should therefore treat full EVA-Bench runs as non-trivial compute investments, particularly when benchmarking multiple models simultaneously.

C Data

We created three datasets on different enterprise domains, each selected to target a distinct axis of difficulty for voice agents. All three require accurate transcription of structured named entities over voice (e.g., confirmation codes and employee identifiers), but differ in their primary challenge. **Airline Customer Service Management (CSM)** tests temporal reasoning and complex policy adherence in high-stakes flight rebooking scenarios. **Healthcare Human Resources Service Delivery (HRSD)** stresses entity density, requiring callers to communicate multiple registration and license numbers across clinical and administrative HR workflows. **Enterprise Information Technology Service Management (ITSM)** introduces branching conversational flows (e.g., incident resolution attempts must fail before ticket escalation is permitted) and tiered authentication reflecting the access sensitivity of different workflows. Dataset statistics are summarized in Table 9.

Within each domain, scenarios span three dimensions: **Single-Intent** (one workflow per call), **Multi-Intent** (one to four concurrent workflows, testing compositional task completion without context loss), and **Adversarial** (hard policy constraints under social pressure, e.g., refusing compensation to an ineligible caller).

Section C.1 introduces the workflows for the three data domains and Section C.2 describes the data generation pipeline and validation process. For sample conversations and scoring please see our website.⁷

C.1 Workflows

The **Airline CSM** domain covers 50 scenarios across seven workflow categories backed by 15 tools. It is high-stakes and time-pressured, with heavy dependence on accurate transcription of named entities such as confirmation codes, flight numbers, and passenger names.

The **Healthcare HRSD** domain covers 83 scenarios across 12 single-intent workflows backed by 47 tools, extended with dual-intent, triple-intent, and adversarial variants. It has the highest per-workflow complexity in EVA-Bench, averaging 8.7 expected tool calls per scenario. Its defining challenge is the density of named entities communicated over voice — NPI numbers, DEA registration numbers, state license numbers, and OTP codes — where a single transcription error can cascade into

authentication or policy failures.

The **Enterprise ITSM** domain covers 80 scenarios across 21 workflows backed by 59 tools, spanning single- to quadruple-intent and adversarial variants. Its defining characteristic is a branching flow structure, where incident workflows gate escalation on failed resolution attempts. Authentication is tiered across three levels — standard, OTP-elevated, and manager-level — reflecting the sensitivity of different workflows.

Across all three domains, scenarios span 1–4 required workflows per task (single-, multi-, and adversarial-intent), giving each domain a graded difficulty distribution rather than a uniform one.

Table 9: Comparison of EVA-Bench data domains. Multi-intent scenarios present one to four concurrent workflows within a single call. Auth tiers refers to distinct levels of caller authentication required across flows.

	Airline CSM	Healthcare HRSD	Enterprise ITSM
Scenarios	50	83	80
Workflows	7	12	21
Tools	15	47	59
Avg. tool calls	3.14	8.7	8.3
Min / max tool calls	1 / 6	1 / 18	1 / 18
Auth tiers	1	2	3

C.2 Data Generation Pipeline

Synthetic Data Generation with SyGra

Scenarios were generated using SyGra (Pradhan et al., 2025), a graph-based synthetic data generation pipeline. Each scenario requires three jointly consistent components: a user goal (including a decision tree that constrains the user simulator to a deterministic outcome), a scenario database (the backend state the agent’s tools query and modify), and an expected final database state (the ground truth against which task completion is evaluated). Joint generation is essential: the expected final state must be consistent with both the user goal and the initial database. Independent generation would introduce silent inconsistencies; for example, a flight number referenced in the user goal that does not exist in the scenario database would corrupt the evaluation signal.

Generation proceeded in the following stages:

1. **Policy specification.** Domain policies and workflow constraints were defined and reviewed prior to generation.
2. **Joint generation.** SyGra generated user goals, initial databases, and expected final states jointly from a workflow graph, using GPT-5.2 as the generative backbone.

⁷<https://servicenow.github.io/eva/#demo>

3. **Multi-intent composition.** Multi-intent scenarios were constructed by combining single-intent records into coherent multi-workflow user goals, with expected final states merged accordingly.
4. **Adversarial scenario design.** Adversarial scenarios were hand-designed around specific policy boundary conditions, then verified against tool executor behavior to confirm that the policy violation is achievable but detectable by a correctly behaving agent.

Human Review

Following generation, all scenarios went through multiple rounds of manual review by domain-specific reviewer teams (Airline CSM: 7 reviewers; Healthcare HRSD: 11 reviewers; Enterprise ITSM: 11 reviewers). Reviewers verified that: (1) policies were applied consistently across scenarios within a domain; (2) user goals were specific enough to admit exactly one correct resolution; (3) expected final states were internally consistent with both the user goal and the initial database; (4) adversarial scenarios were correctly specified, with a clearly identifiable policy violation; (5) scenarios within a reviewer’s assigned section were not duplicates or semantic equivalents of one another; and (6) scenarios were realistic representations of the domain, with unrealistic scenarios flagged for revision or removal. Because each domain’s scenarios were partitioned into non-overlapping workflow sections and reviewed within-section, this process checks for duplication within a section but does not include a cross-section or cross-domain duplication check. Records identified as ambiguous, inconsistent, unrealistic, or duplications were corrected or discarded.

Frontier Model Stress Testing

As a final validation step, we ran three frontier models — OpenAI/Gpt-5.4, Google/Gemini 3.1 Pro, and Anthropic/Claude Opus 4.6 — on a text-only version of each scenario, bypassing the audio pipeline and providing conversation transcripts directly. For every scenario on which any model scored zero on task completion, we manually investigated whether the failure reflected genuine model error or a dataset issue: an ambiguous policy, an under-specified user goal, a bug in the tool executor, or an inconsistency between the initial and expected database states. Records with identified dataset issues were corrected or removed. All se-

lected samples had a task completion of 1 on at least one of the frontier models showing that all tasks are solvable.

This process provides high confidence that task completion failures in the full audio evaluation reflect real agent errors rather than evaluation artifacts.

D User Simulator Details

We used ElevenLabs ElevenAgents (ElevenLabs, 2023) as the user simulator with the following cascade system: Scribe v2.2 Realtime + GPT-5.1 + Eleven V3 Conversational. We selected these models for their high transcription accuracy, User Behavioral Fidelity, user realism for GPT-5.1 (Zhou et al., 2026), and for their naturalness and realism for Eleven v3 Conversational. ElevenLabs also provides a large voice library, enabling testing of a wide variety of user accents, languages, speaking styles, etc.

We created four ElevenLabs agents for the user simulator, covering two accents (English and French) and two genders each. When creating a new agent, select *Blank Agent* as the starting template, then apply the configuration described in Table 10. The four voice names are listed in Table 11. All parameters not listed are set to their default values provided by ElevenLabs at agent creation.

Table 10: ElevenLabs User Simulator Configuration

Parameter	Value
TTS model family	V3 Conversational
Expressive mode	Enabled (no tags selected)
Language	English
LLM	GPT-5.1
System prompt	{{prompt}}
Default personality	Disabled
First message	None
Interruptible	Disabled
Advanced > Input audio	μ -law telephony, 8000 Hz
Advanced > Eagerness	Eager
Advanced > Take turn after silence	15 s
Advanced > Max conversation duration	600 s
Tools > System tools	End conversation (enabled)

When enabling the "End Conversation" system tool, the name must be `end_call`, and the description to provide is shown below. This allows the simulator to hang up programmatically.

"End Conversation" Tool Description

Use this to end the phone call and hang up.
Call this function when its time to end the call and one of the following is true:

1. The agent has confirmed your request is resolved (all steps are completed) and you have said goodbye
2. The agent has initiated a transfer to a live agent
3. The agent has been unable to make progress for at least 5 consecutive turns

4. The agent says goodbye or indicates the conversation is over
5. The agent indicates that the remainder of your request cannot be fulfilled.
6. If the assistant says something along the lines of "I'm sorry I encountered an error processing your request."

IMPORTANT: never call this tool in the same turn that you provide the agent with data, an identifier, a request to transfer to a live agent, an approval to proceed, or any kind of additional information.
Before calling this tool, always say a brief goodbye first.

Once the agent is configured, click Publish in the top-right corner. The agent-id can be retrieved from the Widget tab of the agent dashboard, under Embed code.

The simulator is prompted in EVA-Bench with a specific user goal and is instructed to stay on task, communicate all required named entities clearly, and terminate the conversation when the goal is accomplished, or the task is clearly unlikely to succeed. The system prompts are provided in Appendix I.

D.1 User Simulation Validation

Table 12 lists the three validation metrics used to verify user-simulator behavior and conversation integrity before any EVA-Bench evaluation is performed.

D.1.1 Conversation Valid End.

Before invoking any LLM judges, we run a deterministic check on the conversation logs to verify that each simulation terminated correctly. A valid end state is one in which either the agent failed to respond to the user, or the user invoked the end-call tool. Conversations meeting this criterion proceed to the User Speech and Behavioral Fidelity judges; all others are rerun.

This gate primarily catches infrastructure-level failures: WebSocket connections that closed unexpectedly, conversations that failed to start, or user simulator timeouts. These are pipeline errors rather than agent errors, and filtering them deterministically before any judge invocation avoids wasting LLM calls on malformed simulations.

Table 11: ElevenLabs Voice Name per Agent

Accent	Gender	Voice Name	Voice ID
English	Female	Natalee Champlin	KpTQ5yzwazQWLkvnK59A
English	Male	Eric - Smooth, Trustworthy	cjVigY5qzO86Huf0OWal
French	Female	Mariva Viva Muse - Warm and Energetic	1hIScOW98xkqE5ttC10C
French	Male	Jamie - French Accent & Charismatic	K8nDX2f6wjv6bCh5UeZi

Table 12: EVA-Bench validation metrics. These metrics ensure the quality of the user simulator and the integrity of each conversation prior to evaluation.

Metric	Type	Scale	Pass Thresholds
Conversation Finished	Deterministic	{0, 1}	1.0
User Behavioral Fidelity	LLM-as-Judge	{0, 1}	1.0
User Speech Fidelity	LALM-as-Judge	1-3 $\rightarrow$ [0, 1]	-

D.1.2 User Behavioral Fidelity.

Motivation. A key assumption underlying our evaluation is that the user simulator faithfully follows its assigned goal and decision logic, since the ground truth end database state is derived from this assumption. If the user deviates, the agent may fail to reach the ground truth state through no fault of its own. We therefore define a validation gate that detects user behavior errors capable of corrupting the evaluation, organized into five *corruption types*:

1. **Extra modifications.** The user makes requests beyond its stated goal that invoke modification tools writing to the scenario database. The user simulator prompt explicitly instructs the user to decline any such offers from the agent, but we check for violations regardless.
2. **Premature ending.** In our simulations, the user is responsible for ending the call once its goal is complete. If the user hangs up prematurely—for example, providing actionable information and ending the call in the same turn—the agent has no opportunity to execute the required tool calls. We therefore verify that the user does not terminate the conversation in the same turn it provides critical information to the agent.
3. **Missing information.** If the user fails to provide information the agent needs to complete the task, the evaluation is corrupted since task success cannot reasonably be expected.
4. **Duplicate modifications.** Occasionally, the user simulator (particularly when using non-primary models) enters a loop and repeats requests the agent has already fulfilled. The agent then acts on the duplicate request, causing redundant writes to the scenario database that cause the final state comparison to fail.

5. **Decision tree violations.** Each user is given a structured decision tree governing how to navigate choices during the interaction (e.g., “*accept the earlier flight if the price difference is under \$200, otherwise decline*”). We verify that the user adheres to this logic, since deviations would cause the agent to reach a final state inconsistent with the ground truth.

Implementation. We prompt a judge model with a description of the five corruption types, the full conversation trace (including agent tool calls), and the list of available agent tools. The judge first produces a *corruption analysis* assessing each corruption type in turn, then outputs a binary flag per corruption type, and finally an overall binary rating for the conversation. Conversations receiving a rating of 0 are rerun. See [subsection K.7](#) for the full prompt.

Validation. To validate the judge, we constructed a human-annotated dataset of real user simulator failures from earlier evaluation artifacts alongside correct behavior examples, labeled with both overall ratings and per-corruption-type annotations. GPT-5.2 with medium reasoning achieved 100% accuracy across three independent runs, and was therefore selected as the judge model.

Error Categorization. Across 714 trials flagged by the User Behavioral Fidelity judge—spanning four evaluated systems across all domains—premature ending was the most frequent corruption type (63.9%), followed by decision tree violations (52.9%), missing information (28.6%), extra modifications (1.3%), and duplicate modifications (0%). Notably, 42.4% of flagged trials exhibited two or more corruption types simultaneously. These numbers are summarized in [Table 13](#).

D.1.3 User Speech Fidelity.

Motivation. As noted in [subsection 3.1](#), speech fidelity flags are rare in practice, which may invite questions about why this validation gate was included at all. The answer lies in one of the primary failure modes we observe in voice agents: failure to transcribe and understand key entities. To attribute

such failures confidently to the agent rather than the user simulator, we must verify that the user’s synthesized speech correctly conveys all critical entities. This is particularly consequential because most flows begin with an authentication step requiring the agent to correctly capture entities such as names, confirmation codes, and account IDs—if the user’s speech corrupts these, the agent cannot proceed regardless of its own capability.

Implementation. We adapt the same prompt used to evaluate agent speech fidelity (Section K.3), with one key modification to the rating scale. Unlike for the agent, we do not require the user’s speech to precisely mirror every word of the user-side LLM output; only key entities and major informational content must be conveyed accurately. The judge therefore rates each turn on a 3-point scale: 3 indicates full fidelity, 2 indicates minor errors that do not affect the agent’s ability to progress (e.g., slight disfluencies, but all entities intact), and 1 indicates an entity error or significant omission or addition that would prevent the conversation from proceeding sensibly. Any conversation in which any turn receives a rating of 1 is rerun. The full prompt can be found in [subsection K.6](#).

Validation. Because this prompt is closely derived from the agent speech fidelity judge, we inherit its validation. That judge achieved high inter-annotator agreement with human linguists ($\kappa = 0.777$, 95% CI [0.704, 0.835]), and the core capability it requires—accurately parsing audio and detecting entity-level errors—is shared. The rating scale and its interpretation are sufficiently well-defined that additional annotation studies were not deemed necessary. See [table 2](#) for more details on human-judge agreement.

Table 13: Distribution of corruption types among trials flagged by the User Behavioral Fidelity judge. Percentages sum to more than 100% because a single trial may exhibit multiple corruption types.

Corruption Type	Cases	% of Flagged
Premature ending	456	63.9%
Decision tree violation	378	52.9%
Missing information	204	28.6%
Extra modifications	9	1.3%
Duplicate modifications	0	0%

E Metric Details

All EVA-Bench metrics are summarized in Table 14.

E.1 Log Processing and Variable Extraction

Available Logs Every conversation typically produces three independent log streams, which we merge and replay deterministically to recover the variables needed by our metrics. The streams are:

- **Audit log.** The agent’s own record: user-turn transcripts, the assistant’s LLM output, including tool calls and their responses, each with parameters and timestamps. It does not contain detailed speech framework events like turn detection or TTS events.
- **Framework events.** Detailed events with wall-clock timestamps, extracted from speech frameworks (Pipecat, S2S framework, ...). The *intended* assistant text comes from the pipeline’s LLM output or TTS input, whichever is logged. This can be shorter than the assistant’s full LLM output in the audit log, since an interruption may cancel text before it reaches the TTS model.
- **User simulator events.** Events from the user simulator: per-speaker start/end markers for each audio segment, the simulator’s own outgoing text (the user’s *intended* utterance, and our source of truth for scoring the agent-side user transcript), and the transcript of the assistant’s audio channel. This transcript can differ from the assistant’s LLM output because of speech or transcription errors, but for S2S models it is sometimes the only record of the assistant’s output.

Extracted variables. For each turn we extract per-role text and audio variables and a linearised conversation trace that interleaves user and assis-

tant turns with tool calls. Table 15 gives the default mapping from log source to variable. The conversation trace feeds judge metrics needing a chronological transcript; the per-turn fields feed metrics needing intermediate states, such as Transcription Accuracy Key Entities.

Pipeline-specific implementations. Pipeline type overrides the default mapping in two ways, reflecting which signals are available and trustworthy for each architecture. These differences affect the metric definitions (main paper, Section 3.2), and most judge prompts have fragments swapped in dynamically to match the evaluated architecture.

- **Cascade.** All three streams are unmodified.
- **S2S.** No separable STT or TTS steps, so no user STT transcript and no separate text intent for the output. Both roles in `conversation_trace` are sourced from the user simulator events: intended text for user turns, transcribed text for assistant turns.
- **Hybrid.** User turns sourced as in S2S (simulator intended text, no STT); assistant turns sourced as in cascade (TTS entries from framework events).

Log consolidation. We merge the streams into one timeline sorted by timestamp, indexing turns so that assistant turn i responds to user turn i . Because the streams come from independent components, they can drift in edge cases, especially around interruptions and rapid turn switches: events can be missing or delayed, and one user turn can be detected as two. We handle several such cases during processing. To validate this, we manually inspected a stratified sample of conversations across all three pipeline types; residual errors were rare and did not systematically bias any reported metric.

Table 14: EVA-Bench metrics organized by category. All EVA-A and EVA-X scores are normalized to $[0, 1]$ prior to aggregation. Thresholds are used for `pass@k` and `pass^k` computation: a run is considered successful if all EVA-A and EVA-X metrics meet their respective thresholds simultaneously.

Category	Metric	Type	Scale	Pass Thresholds
EVA-A (Accuracy)	Task Completion	Deterministic	$\{0, 1\}$	1.0
	Faithfulness	LLM-as-Judge	$1-3 \rightarrow [0, 1]$	0.50
	Speech Fidelity	LALM-as-Judge	$[0, 1]$	0.95
EVA-X (Experience)	Conciseness	LLM-as-Judge	$1-3 \rightarrow [0, 1]$	0.50
	Conversation Progression	LLM-as-Judge	$1-3 \rightarrow [0, 1]$	0.50
	Turn-Taking	Deterministic	$[0, 1]$	0.80
Diagnostic	Authentication Success	Deterministic	$\{0, 1\}$	—
	Response Latency	Deterministic	seconds	—
	Speakability	LLM-as-Judge	$\{0, 1\}$	—
	STT Word Error Rate	Deterministic	$[0, \infty)$	—
	Tool Call Validity	Deterministic	$[0, 1]$	—
	Transcription Key Entities	LLM-as-Judge	$[0, 1]$	—

Table 15: Default mapping from log source to extracted variable. The upper block lists the per-turn text and audio fields; the lower block lists the entries that compose the linear `conversation_trace`. Pipeline-specific overrides are listed in the text. * Empty if the framework emits neither record. † Falls back to `user_simulator_events / assistant_speech` for S2S. ‡ Uses `user_simulator_events / user_speech` for audio-native pipelines (S2S/Hybrid).

Variable	Source
<code>intended_user_turns[i]</code>	<code>user_simulator_events / user_speech</code>
<code>transcribed_user_turns[i]</code>	<code>audit_log / user</code>
<code>intended_assistant_turns[i]</code>	<code>framework_logs / tts_text, llm_response*</code>
<code>transcribed_assistant_turns[i]</code>	<code>user_simulator_events / assistant_speech</code>
<code>audio_timestamps_{role}_turns[i]</code>	<code>user_simulator_events / audio_start, audio_end</code>
<code>conversation_trace (assistant)</code>	<code>audit_log / assistant</code> , truncated to framework-log prefix†
<code>conversation_trace (user)</code>	<code>audit_log / user‡</code>
<code>conversation_trace (tools)</code>	<code>audit_log / tool_call, tool_response</code>

Interruptions We flag a likely interruption where one speaker’s `audio_start` fires while the other’s audio session is still open, labeling the entries `[user/assistant interrupts]`, `[likely cut off by user/assistant]`. We say *likely* because the logs record text sent to TTS, not audio played, so the exact truncation point is unknown. Judges use the label to avoid penalizing the agent for intended words that were never spoken. For example, Speech Fidelity should not penalize the agent for intended words that were never spoken, and Conversation Progression should not penalize the agent for repeating interrupted sentences.

Audio recordings. Every session also writes three 16-bit PCM mono WAV files captured from the audio bus: `audio_user.wav` (user simulator output only), `audio_assistant.wav` (assistant output only), and `audio_mixed.wav` (both speakers, for human review, used by no metric). The role-specific files feed the speech-fidelity metrics, so the judge need not attribute speech to the correct speaker.

E.2 Judge Development and Validation

LLM-as-judge evaluations are only as reliable as the judges themselves. We developed each judge through a structured pipeline: metric definition, prompt construction, development dataset construction, prompt improvements, model selection, and final validation against human annotation.

Metric definition and prompt construction. For each judge metric, we first defined the rating scale and the failure modes the judge is expected to catch, specifying how a judge should detect and categorize every issue. This categorical and granular framing is intentional: it preserves actionable signal about *how* a voice agent fails, rather than collapsing everything into a single score. It also helps

the judge accurately score by offering detailed categories. From this exercise, initial judge prompts emerged.

Development dataset construction. We constructed a development dataset for each metric using synthetic data generation followed by multi-model consensus labeling. For data generation, we sampled naturally occurring conversations and also prompted frontier models to produce conversations exhibiting specific failure modes within each metric. This targeted generation was key to achieving class balance.

For labeling, we called three frontier models, Gemini 3.1 Pro, Claude Opus 4.6, and GPT-5.2, as judges on each generated sample. When all three models agreed on a rating, that rating was assigned as ground truth. When they disagreed, a human reviewer examined the sample and selected the correct label.

Prompt improvements and model selection. We used the development sets to refine the prompts, focusing on samples where judges disagreed and analyzing their explanations. This process revealed ambiguities in the judge instructions and informed targeted prompt improvements. We then formally evaluated each of the three frontier models as judge candidates on the development datasets, reporting accuracy alongside macro-averaged F1, and selected the ones achieving the highest combined score (see Table 16). For Speech Fidelity, Gemini 3.1 Pro achieved the highest performance, but we selected Gemini 3 Flash for deployment: it achieved nearly identical performance at substantially lower inference cost. Model parameters used for each judge are provided in Table 17.

Human agreement validation. To assess the validity of the finalized judge prompts and selected models, we separately constructed a held-out test

Table 16: Judge validation datasets and results per candidate model. Size = number of annotated examples. Acc. = judge accuracy against adjudicated human labels. F1 = macro F1-score. **Bold** indicates the selected judge model for each metric.

Metric	Dataset	Claude Opus 4.6		GPT-5.2		Gemini 3 Flash		Gemini 3.1 Pro	
	Size	Acc.	F1	Acc.	F1	Acc.	F1	Acc.	F1
Faithfulness	137	83.9%	80.7%	70.8%	68.9%	–	–	81.0%	78.4%
Conciseness	100	90.8%	73.4%	92.3%	84.0%	–	–	91.2%	74.4%
Conv. Progression	136	74.3%	73.1%	78.7%	76.7%	–	–	69.1%	66.6%
Speech Fidelity	147	–	–	–	–	89.1%	83.2%	91.2%	85.9%

Table 17: Judge model identifiers and inference parameters used during evaluation. Parameters not listed were left at provider defaults.

Parameter	Value
GPT-5.2 gpt-5.2	
max_tokens	100,000
Claude Opus 4.6 us.anthropic.claude-opus-4-6-v1	
provider defaults	–
Gemini 3 Flash gemini-3-flash-preview	
temperature	0.0
max_tokens	40,000
reasoning_effort	minimal

set of 63 conversation samples per metric, never used during prompt development. Each sample was labeled by two expert linguists. For text-based metrics (Faithfulness, Conversation Progression, Conciseness), linguists were shown the conversation traces and the judge outputs from all three frontier model candidates, presented blind to which model produced which output. We provided judge outputs to the linguists because voice agent transcripts can be long and complex; without the judge’s surfaced evidence and analysis, annotators found it difficult to reliably identify subtle failure modes in long traces. Linguists verified the judge’s reasoning against the transcript directly rather than accepting it uncritically. For Speech Fidelity, linguists listened to the raw audio only along with the intended text, with no judge output, since this metric requires direct perceptual evaluation of synthesized speech. We report inter-annotator agreement in the main paper using Cohen’s κ in Table 2.

E.3 Accuracy Metrics

E.3.1 Task Completion

What it measures. Task Completion is the bottom-line accuracy check: it asks whether the changes the agent committed to the scenario database match the expected end state specified in the dataset’s ground truth. Unlike the judge metrics, this is a deterministic code-based check; the same conversation always yields the same binary verdict.

Method. Each scenario in the dataset specifies the database state we expect after a successful run.

During execution, every state-mutating tool call writes through to a per-record copy of the scenario database, and the final state is captured at the end of the conversation. The metric canonically serialises both states, computes their SHA-256 hashes, and reports a pass if the two hashes match. When they do not match, a structured diff is returned.

Inputs.

- `expected_scenario_db`: the dataset’s ground-truth final state for this record;
- `final_scenario_db`: the actual final state captured at end-of-run

Authentication gate. Authentication state lives in a dedicated `session` field of the scenario database, and is verified separately from the hash comparison rather than folded into it. The reason is that some scenarios can be satisfied by different valid authentication paths. We therefore verify authentication separately via a superset check: every key-value pair in the expected session must be present in the actual session.

Rating scale. Task completion is binary: 1.0 when the hashes match, 0.0 otherwise.

Pass/fail thresholding. For pass-related aggregations such as `pass@k`, the pass threshold is set to 1.0 (i.e. exact match).

E.3.2 Faithfulness

What it measures. Faithfulness is a conversation-level judge that asks whether the assistant remained grounded in the information, policies, and instructions available to it throughout the run. Unlike Task Completion which only checks whether the goal was achieved, this metric also penalizes the path: a conversation that concludes successfully but hallucinates a policy or skips a required verbal confirmation will receive a low faithfulness score.

Judge and prompt template. The judge uses *Claude Opus 4.6*. As for other text-judge metrics, specific prompt fragments were written to account for differences in the conversation trace for different pipeline architectures. The full prompt is

available in Appendix K.2.

Inputs.

- the `conversation_trace` (with tool calls and responses inline);
- the agent’s configuration (instructions/tools);
- the simulated `current_date_time`, used to resolve temporal references;

Failure modes. The judge scores five disjoint dimensions, each defined to be non-overlapping so that any given issue maps to exactly one.

- `fabricating_tool_parameters`: the assistant called a tool with a parameter value that cannot be traced to any user statement, prior tool result, policy entitlement, simple arithmetic, or standard domain mapping.
- `misrepresenting_tool_result`: the assistant inaccurately conveyed something a tool returned: wrong field value, contradicted status, omitted material caveats, or arithmetic errors when computing values from tool data.
- `violating_policies`: the assistant contradicted the agent instructions: skipped a required verification step, executed an irreversible write without the disclosure or confirmation the policy requires, or stated a policy incorrectly.
- `failing_to_disambiguate`: the assistant proceeded on ambiguous or contradictory user input without clarification.
- `hallucination`: a residual category for information stated to the user that has no source in any tool response, user utterance, agent instruction, or system context, and that is not already captured by the four preceding dimensions.

Rating scale and aggregation. Each failure mode is rated on a 3-point integer scale:

- **3**: no issue on this dimension.
- **2**: minor or ambiguous issue with low user impact (e.g. a fabricated parameter that reaches a read-only tool, is caught quickly, and never surfaces to the user; a borderline policy deviation; a small phrasing embellishment that does not alter any decision).
- **1**: clear violation with material impact (e.g. a fabricated parameter impacting the database state; an irreversible action without the policy-required disclosure; misstating a fee, balance, or eligibility rule the user could act on).

The ratings are aggregated by *minimum* into a single overall rating, so a single rating-1 dimension produces an overall rating of 1. The overall rating is normalized to $[0, 1]$.

Pass/fail thresholding. For pass-related aggregations, such as `pass@k`, we count a conversation as a faithfulness pass if its overall rating is ≥ 2 . This threshold sits exactly at the rubric’s own load-bearing boundary: rating 1 is reserved for violations with *material* impact on the user, while rating 2 is explicitly defined as covering minor or ambiguous issues. As such, rating-2 conversations are still treated as acceptable: at the current capability level of voice-LLM systems, demanding strict rating-3 perfection on every run would be unrealistic and would compress meaningful differences between systems into a uniformly low pass-rate. This cutoff point is also validated by the pass/fail human-agreement detailed in Section 4.2 and Table 3. Across all 12 systems, ratings are distributed 43.5% at rating 1, 32.0% at rating 2, and 24.4% at rating 3, giving a pass rate of 56.4% under the reported ≥ 2 criterion and 24.4% under the stricter $= 3$ cut.

E.3.3 Speech Fidelity

What it measures. Speech Fidelity asks whether the assistant’s spoken audio actually matches what the system intended (or was expected) to communicate. It is the audio-side complement to Faithfulness: faithfulness scores what the LLM *decided* to say, speech fidelity scores whether the user could correctly hear the entities the LLM (or the upstream tool responses) returned. Errors here would result in the user receiving the wrong information, such as a garbled flight number, a confirmation code with one substituted character, or a dollar amount whose digit is dropped.

Judge and prompt template. Speech Fidelity is an audio-judge metric: it sends a multimodal request (audio + textual context) to *Gemini 3 Flash*. The audio is encoded as `base64 WAV` and accompanied by the per-turn entity context the judge needs to verify. To reduce upstream noise and cost, the assistant audio is silence-trimmed before being sent to the judge. The full prompt is available in Appendix K.3.

Inputs.

- the audio recording of the assistant speech;
- the intended assistant turns for cascade and hybrid systems, the user turns and tool outputs for S2S systems;

Two pipeline-specific variants. Unlike the text judge metrics, speech fidelity does not have a single prompt with pipeline-aware fragments: the cas-

cade/hybrid case and the S2S case have qualitatively different inputs and use different prompts.

- **Cascade and Hybrid.** Both architectures expose an *intended* text-side reference for the assistant. The judge task is a direct word-for-word comparison: did the audio reproduce the intended text, with particular attention to key entities.
- **S2S.** S2S systems do not typically expose any text-side intent, so there is nothing to compare the audio against in the cascade sense. We instead reformulate the question as an *entity articulation* check: does the assistant clearly and correctly speak the entities it was supposed to convey? The judge receives a *redacted conversation trace* with only user utterances and tool responses. These are the entity sources the assistant can say. The judge transcribes the assistant audio itself and checks, per turn, that any entities it speaks that originate from the trace are clearly audible.

Rating scale and aggregation. Each rated turn is given a binary rating in $\{0, 1\}$, where 1 means the audio is correct, and 0 means there is an issue. The conversation-level score is the mean across rated turns; for S2S, turns with no entities are dropped from both the numerator and the denominator.

Pass/fail thresholding. For pass-related aggregations such as pass@k, the pass threshold is set to 0.95, i.e. a conversation passes if at most a small fraction of rated turns are flagged with an entity error. We use a stricter threshold than the rubric-grounded 0.5 used for the 3-point judge metrics for two reasons. First, the rating scale is already binary: there is no rating-2 “minor issue” band whose inclusion this threshold has to negotiate, so the rubric-grounding argument used elsewhere does not apply. Second, the underlying error is high-stakes: a single garbled confirmation code or dollar amount can render a whole conversation operationally wrong, so tolerating more than the occasional turn-level slip would mask exactly the kind of failure this metric exists to catch. A threshold sweep from 0.50 to 1.00 confirms this cutoff is not a fragile choice (Table 18). Pass rates move little below 0.80 (e.g. cascade: 0.372 $\rightarrow$ 0.355; s2s: 0.314 $\rightarrow$ 0.310), and hybrid remains the clear low performer throughout. Movement is concentrated in the 0.80–0.95 band, most visibly for hybrid (0.298 $\rightarrow$ 0.263). The ordering of cascade and s2s is close enough to be threshold-sensitive: cascade leads through 0.90, but the two are within 0.004 of each other by 0.95–1.00.

Threshold	Cascade	Hybrid	S2S
0.50	0.372	0.309	0.314
0.60	0.370	0.309	0.312
0.70	0.363	0.308	0.311
0.80	0.355	0.298	0.310
0.85	0.340	0.283	0.309
0.90	0.320	0.268	0.309
0.95	0.305	0.263	0.309
1.00	0.305	0.263	0.309

Table 18: Speech Fidelity threshold sweep: pooled EVA-A pass@1 by architecture across thresholds from 0.50 to 1.00.

Limitations. Manual inspection of a stratified sample of rated turns surfaced a few false positives (the judge rates 0 on a turn whose audio actually reproduced the intended content). They are typically caused by upstream log-processing artifacts rather than by judge mistakes. The metric compares the audio against *intended_assistant_turns*, and that reference is only as good as the merge of the framework log and the user simulator audio events described in Appendix E.1. These false positives are concentrated on turns adjacent to interruptions and rapid turn switches.

E.4 Experience Metrics

E.4.1 Conversation Progression

What it measures. Conversation Progression is a conversation-level experience judge that asks whether the assistant moved the conversation forward without redundancy: no repeated tool calls with identical parameters, no restating of information already communicated, retention of established facts, and well-formed clarification questions. Unlike Faithfulness, which scores the *correctness* of the assistant’s choices, conversation progression scores the *efficiency*.

Judge and prompt template. The judge is *GPT-5.2*. As for other text judge metrics, specific prompt fragments were written to account for differences in the conversation trace for different pipeline architectures. The prompt template is available in Appendix K.4.

Inputs. The `conversation_trace` (with tool calls and responses inline).

Failure modes. The judge scores four disjoint dimensions:

- `unnecessary_tool_calls`: the assistant called a tool without justification: same tool with the same parameters after a successful prior re-

sponse, a tool with empty/missing required parameters that produced a predictable error, or a tool whose result was already available from a previous response.

- `information_loss`: the assistant failed to retain or act on an established fact: re-asked the user for information already provided, ignored a constraint the user explicitly stated, or failed to use a required value returned by a prior tool.
- `redundant_statements`: the assistant restated information it had already communicated to the user (repeated explanations, repeated status updates, multiple recaps across turns). The exception is a single recap at the very end of the conversation.
- `question_quality`: the assistant's questions were poorly formed or missing: overly broad/vague questions when enough information was on hand to act, multiple questions bundled in a turn that a single tool call could resolve, or missing clarifications when input was genuinely ambiguous.

Rating scale and aggregation. Each failure mode is rated on a 3-point integer scale:

- **3**: no issue on this dimension.
- **2**: a single isolated issue that does not significantly impact conversation flow (e.g. one unnecessary tool call that didn't slow things down, a single redundant restatement, one vague question), or a borderline case where it is unclear whether the issue constitutes a real progression problem.
- **1**: multiple instances of the same type of issue in this dimension, or a single severe issue that clearly derailed or stalled the conversation.

An aggregated score is computed on the per-dimension scores, based on the following rules. If no failure mode is flagged, the overall rating is 3. If any dimension is rated 1, *or* if three or more dimensions are flagged, the overall score is 1. Anything in between is 2. The overall rating is normalized to $[0, 1]$.

Pass/fail thresholding. We count a conversation as a conversation-progression pass if its overall rating is ≥ 2 . The threshold sits at the rubric's own load-bearing boundary: rating 1 is reserved for runs where progression *materially* broke down, while rating 2 covers a small number of isolated minor inefficiencies that did not impede the outcome; conversations at rating 2 are still treated as acceptable. This cutoff point is also validated by the pass/fail

human-agreement detailed in Section 4.2 and Table 3. Across all 12 systems, ratings are distributed 18.8% at rating 1, 42.4% at rating 2, and 38.8% at rating 3, giving a pass rate of 81.2% under the reported ≥ 2 criterion and 38.8% under the stricter $= 3$ cut.

E.4.2 Conciseness

What it measures. Conciseness is an experience metric that asks whether each assistant turn is appropriately brief and voice-appropriate: a listener consuming the response in real time should be able to absorb its content in a single pass.

Judge and prompt template. The judge is run as a single LLM-as-judge call per conversation (one call returning a per-turn array), using *GPT-5.2*. The prompt template is available in Appendix K.5.

Inputs. The `conversation_trace`.

Failure modes. When a turn is rated below 3, the judge tags it with one or more of the following (a turn may carry multiple tags):

- `verbosity_or_filler`: unnecessary wording, hedging, or repetition within the same turn beyond what the context requires.
- `excess_information_density`: too many distinct facts, options, numbers, steps, or requests packed into one turn for a listener to retain in real time.
- `over_enumeration_or_list_exhaustion`: reading out long lists exhaustively rather than summarising, or presenting multiple options with excessive per-option detail rather than inviting follow-up.
- `contextually_disproportionate_detail`: more background, clarification, or explanation than the situation actually warrants.

Rating scale and aggregation. Each assistant turn is rated on a 3-point integer scale:

- **3** (highly concise): no issue.
- **2** (adequate but not optimally concise): exactly one minor failure mode is present, but the response remains processable and does not meaningfully overwhelm the listener.
- **1** (not concise): one or more significant failure modes are present that would materially increase cognitive load and hinder comprehension when spoken.

Each per-turn rating is normalized to $[0, 1]$, and the conversation-level Conciseness score is the *mean* of these per-turn normalized ratings across all rated turns. Unlike Faithfulness (minimum across di-

Bucket	Conciseness
[0.0, 0.1)	0.0
[0.1, 0.2)	0.0
[0.2, 0.3)	0.0
[0.3, 0.4)	0.3
[0.4, 0.5)	0.7
[0.5, 0.6)	6.9
[0.6, 0.7)	15.6
[0.7, 0.8)	27.0
[0.8, 0.9)	35.0
[0.9, 1.0]	14.5
Mean (std)	0.78 (0.11)
Threshold	≥ 0.50

Table 19: Distribution of conversation-level Conciseness scores across all 12 systems, by 0.1-width bucket.

mensions) and Conversation Progression (count-aware aggregation across dimensions), Conciseness is therefore a continuous score in $[0, 1]$.

Pass/fail thresholding. Following the same convention as Faithfulness and Conversation Progression, we count a conversation as a Conciseness pass if its conversation-level normalized score is ≥ 0.5 , i.e. if its mean per-turn rating is ≥ 2 . Because Conciseness aggregates per-turn scores, this threshold has a slightly different rubric-level interpretation: a pass means the assistant was on average at least “adequate but not optimally concise” across the conversation, with isolated rating-1 turns tolerated as long as they are offset by rating-3 turns elsewhere. The threshold is intentionally lenient: verbosity degrades the listening experience but does not cause any material harm. Across all 12 systems, per-turn conciseness scores average 0.78 (std 0.11), and under 1% of conversations fall below 0.5 (Table 19), confirming that in practice the threshold screens only extreme verbosity rather than typical performance. A stricter cutoff would separate systems more sharply on this metric; we expose the threshold as a configurable parameter so downstream users can raise it if finer-grained discrimination on conciseness is desired.

E.4.3 Turn-Taking

What it measures. Turn-Taking is a deterministic experience metric that scores the temporal quality of each exchange: whether the agent responds in a natural conversational window, avoids interrupting the user, and yields promptly when interrupted. Unlike prior works (Ray et al., 2026; Lin et al., 2026a) that score turn-taking as flat per-turn la-

tency scalars, we rate each turn differently based on if interruptions occurred, and we relax latency thresholds on turns involving tool calls.

Inputs. Per-turn event timestamps from the user simulator events.

Per-turn scoring. Each turn is first classified by its interrupt condition and routed to the corresponding scoring function (Table 20): s_{latency} for uninterrupted turns, s_{agent} or s_{yield} for agent- or user-initiated interruptions, and their minimum when both occur.

- **Uninterrupted** (s_{latency}): response latency ℓ (end of user utterance to agent response onset) is mapped through a piecewise-linear curve: hard-zero for premature responses, a rising ramp toward the sweet spot as the response settles into the natural window, a flat sweet spot over the natural conversational window, and a falling ramp back to 0 once silence grows long enough to disrupt the exchange. Region boundaries and tool-call-aware variants are detailed in Table 21. Boundaries are set deliberately permissive to reflect current voice agents capabilities, as explained below.
- **Agent interrupts user** (s_{agent}): scored as the minimum of three sub-scores (Table 22), each capped at 0.5 so an interruption is never cost-free. *Overlap* penalizes total simultaneous speech duration o up to $o_{\text{max}} = 2,000$ ms; *count* penalizes the number of distinct agent segments n overlapping user speech (>1 ms intersection) up to $N_{\text{max}} = 3$; *post-interrupt* measures how cleanly the agent resumes after interrupting: the silent gap before its next settled response is scored on the same latency curve as uninterrupted turns.
- **User interrupts agent** (s_{yield}): linearly penalizes the agent’s yield latency Δt up to $\Delta t_{\text{max}} = 2,000$ ms, rewarding immediate yield with a full score of 1 (Table 20).

The conversation-level score is the mean of per-turn scores, giving a continuous value in $[0, 1]$.

Region boundaries calibration. Human turn-taking is fast; inter-turn gaps center near 200 ms (Stivers et al., 2009; Heldner and Edlund, 2010), a target no evaluated voice agent can meet. Even the more lenient 700–1,000 ms silence thresholds targeted by some deployed dialogue systems (Skantze, 2021) are too strict for our setting: the fastest systems in our pool respond in ≈ 1 s on average on non-tool turns, and most cascade systems in 2–5 s (Table 27). A sweet-spot ceiling in that range

Table 20: **Per-turn scoring regime.** Each turn is routed to the appropriate scoring function conditioned on its interrupt condition. The latency curve s_{latency} (Table 21) and the agent-interruption sub-scores s_{agent} (Table 22) are defined separately. For the yield score, Δt is the agent’s yield latency across the turn boundary and $\Delta t_{\text{max}} = 2,000$ ms.

Turn Condition	Score Function
Uninterrupted	s_{latency}
Agent interrupted user	$s_{\text{agent}} = \min(s_{\text{overlap}}, s_{\text{count}}, s_{\text{post}})$
User interrupted agent	$s_{\text{yield}} = \max(0, 1 - \frac{\Delta t}{\Delta t_{\text{max}}})$
Both interruptions	$\min(s_{\text{agent}}, s_{\text{yield}})$

would place the majority of turns of nearly every system in the late-penalty region, compressing the score range and leaving the metric unable to discriminate between systems. We therefore set the boundaries beyond the levels prior work suggests, calibrated to the latencies voice agents actually achieve. The boundary values (Table 21) can be modified based on the use case.

Pass/fail thresholding. A conversation passes if its Turn-Taking score is ≥ 0.8 . Most other metrics inherit their cutoff from an ordinal rubric, but turn-taking scores are continuous and have no such intrinsic threshold, so we justify the threshold τ_{tt} at 0.8 on four grounds.

(i) *Conversational interpretation.* The per-record score is the mean of per-turn scores in $[0, 1]$, where a turn earns 1 when its latency is acceptable. A threshold of 0.8 admits at most one off-bracket turn in five; an interpretable, conversation-level reading of the cutoff (“4/5 turns on-time”).

(ii) *The latency region boundaries are already lenient.* Given that the boundaries are deliberately permissive, a low τ_{tt} would compound this leniency, making the metric trivially easy to pass and reducing its ability to discriminate systems. We therefore aim to select a stricter τ_{tt} .

(iii) *Calibration to current capability.* Among the 12 benchmarked systems (Table 27), a wide gap separates systems achieving real-time turn-taking

Table 21: **Latency curve boundaries by turn type.** Tool-call turns receive more lenient upper thresholds ($\ell_{\text{sweet-high}}$, $\ell_{\text{hard-late}}$) to absorb tool execution latency; lower breakpoints are unchanged.

Breakpoint	Std / Tool	Description
$\ell_{\text{hard-early}}$	−500 / −500	Likely interruption; score = 0; score ramps up above.
$\ell_{\text{sweet-low}}$	500 / 500	Lower edge of optimal window; score = 1 here.
$\ell_{\text{sweet-high}}$	2,000 / 3,000	Upper edge of optimal window; score ramps down beyond.
$\ell_{\text{hard-late}}$	3,500 / 5,000	Excessively delayed; score = 0 at or beyond.

Table 22: **Agent interruption sub-scores**, each capped at $M = 0.5$. The minimum of the three is taken, so the weakest dimension dominates. $o_{\text{max}} = 2,000$ ms, $N_{\text{max}} = 3$.

Sub-score	Formula
Overlap	$s_{\text{overlap}} = \max(0, M(1 - \frac{o}{o_{\text{max}}}))$
Count	$s_{\text{count}} = \max(0, M(1 - \frac{n-1}{N_{\text{max}}-1}))$, $n \geq 1$
Post-interrupt	$s_{\text{post}} = s_{\text{latency}}(\Delta t)$

from those that do not: the three S2S systems score 0.815–0.830, while the next-highest scores 0.583, a gap exceeding 0.23. Any $\tau_{\text{tt}} \in [0.6, 0.8]$ produces the same pass/fail partition of all 12 systems, so the precise value within this range does not change conclusions. A higher threshold (e.g., 0.9) is miscalibrated to current capability: no system in our pool would reliably qualify.

(iv) *Rankings are stable across the threshold.* We sweep $\tau_{\text{tt}} \in \{0.50, 0.55, \dots, 0.95\}$ and recompute EVA-X pass@1 at each value (Figure 4). Rankings are preserved across the full range: Spearman $\rho = 0.968$ between endpoints, and 61/66 (92.4%) of pairwise orderings hold strictly across all ten thresholds (63/66, 95.5%, counting statistically tied pairs as preserved). The three pairs that flip are systems within 0.15 pass@1 of each other. The per-system pass@1 vector is likewise linearly stable around the anchor: Pearson $r = 0.998$ at $\tau_{\text{tt}} = 0.75$ and $r = 0.995$ at $\tau_{\text{tt}} = 0.85$, never falling below 0.910 across $[0.50, 0.95]$ (Table 23).

Table 23: Pearson correlation between per-system EVA-X pass@1 vectors at $\tau_{\text{tt}} = 0.8$ and other candidate thresholds ($n = 12$ systems).

τ_{tt}	Pearson r vs. $\tau_{\text{tt}} = 0.80$
0.50	0.910
0.55	0.947
0.60	0.967
0.65	0.979
0.70	0.991
0.75	0.998
0.80	1.000 (anchor)
0.85	0.995
0.90	0.980
0.95	0.950

Between the two grounds that narrow the choice, the “4/5 on-time” reading and the already-lenient bracket, we select the top of the $[0.6, 0.8]$ range. This value reflects current model capability, not a fixed property of the metric: EVA-Bench exposes `pass_at_k_threshold` as a configurable parameter, and we expect the default to rise as model latencies improve.

E.5 Diagnostic Metrics

Diagnostic metrics include:

- *Authentication Success Rate*, surfacing if agents




Figure 4: Sensitivity of EVA-X pass@1 to the turn-taking pass-threshold τ_{tt} . We sweep τ_{tt} from 0.50 to 0.95 in 0.05 increments and recompute EVA-X pass@1 at each value, holding the Conversation Progression and Conciseness thresholds fixed at 0.5. The dashed vertical line marks $\tau_{tt} = 0.8$. (a) Architecture-aggregated. Each curve is the mean per-scenario EVA-X pass@1 (averaged over $k = 5$ trials per scenario), pooled with equal weight across all (system, scenario) pairs within an architecture class (cascade, hybrid, S2S); shaded bands are percentile bootstrap 95% CIs from 1,000 resamples of (system, scenario) pairs. (b) Per-system. One curve per benchmarked system, colored by architecture; shaded bands are bootstrap 95% CIs from 1,000 resamples of scenarios within the system. Architecture ordering (S2S > cascade > hybrid) is preserved at every threshold; per-system rankings flip only between close competitors.

- were able to pass this critical gate (deterministic);
 - *Response Latency*, in seconds decomposed into turns with and without tool calls (deterministic);
 - *Speakability*, evaluating whether the agent text is voice-friendly prior to synthesis (LLM-as-Judge);
 - *STT Word Error Rate*, computed via `jiwer` (deterministic);
 - *Tool Call Validity*, the fraction of tool calls with correctly formatted parameters (deterministic);
 - *Transcription Accuracy for Key Entities*, STT accuracy specifically on key entities such as names, dates, and alphanumeric codes (LLM-as-Judge)
- Together, these metrics let practitioners dig deeper, and understand why failures occur. For example, a failed task can be traced to a transcription error via *Transcription Accuracy for Key Entities* on a confirmation code, rather than to correct transcription undermined by faulty reasoning.




Figure 5: **Mean transcription accuracy on key entities per cascade STT model, pooled across the three EVA-Bench domains.** Each bar is the mean of per-scenario accuracies (one value per scenario, averaged over $k=5$ trials); whiskers are 95% normal-approximation CIs on that mean. Models are sorted ascending. Nova-3 aggregates the two cascades that share it (Nova-3 + GPT-5.4 + Sonic 3 and Nova-3 + GPT-5.4-mini + Aura-2), giving it twice as many scenario observations and a correspondingly tighter interval. Scribe-v2.2-Realtime saturates near 1.0 on this benchmark; the gap to the next-best STT (Parakeet-1.1) exceeds 0.19.

F Metrics Analysis

F.1 Key Entity Transcription Accuracy & Task Completion

The domains in EVA-Bench are entity-dense; agents must extract confirmation codes, names, authentication codes, employee IDs, and similar information. These entities are needed not only to authenticate users, but also to complete downstream task steps (e.g., looking up a ticket, updating a record, or routing a request). Figure 5 breaks down key entity transcription accuracy by cascade STT model, pooled across the three EVA-Bench domains; Scribe-v2.2-Realtime leads, with the remaining models spanning 0.59 to 0.78. As shown in Figure 6, on these domains, transcription accuracy on key entities and task completion are tightly coupled for all the evaluated cascade systems. However, transcription accuracy on key entities alone does not predict task completion. Two systems use Nova-3 for STT and score similarly on transcription, yet the GPT-5.4 system completes meaningfully more tasks than the GPT-5.4-mini one, showing that how well the LLM reasons over transcribed entities also matters.

F.2 Patterns Across Pipeline Architectures

Among the systems we evaluated, the S2S systems scored highest on the responsiveness facets, though only narrowly ahead of the strongest cascades. They scored lowest on policy adherence. We focus on four facets that surface these patterns most clearly: on-time turns, conversation completion, authentication rate, and policy violations (from the faithfulness judge).

For each facet, a per-system mean averages that




Figure 6: Correlation between mean key entity transcription accuracy and mean task completion across the 7 evaluated cascade systems. Pearson $r = 0.930$, $p = 0.002$.

system’s per-domain values. The pipeline-type mean is the unweighted mean of its constituent system means, with a 95% percentile bootstrap interval from resampling those means with replacement (10,000 iterations, fixed seed). Pairwise effect sizes $\Delta = \mu_A - \mu_B$ are reported with a percentile bootstrap interval from independently resampling each group’s per-system means. Differences are in percentage points ($\Delta \times 100$); metrics are bounded in $[0, 1]$. For pipeline types with few systems (hybrid, $n=2$; S2S, $n=3$), the percentile bootstrap over system means is degenerate; its interval cannot extend beyond the observed system values, so we report these intervals as descriptive spread, not sampling-based confidence intervals. All numbers use the clean (unperturbed) runs of the systems in the main paper, across all three domains, yielding 1,065 clean conversations per system. Tab. 24 reports all values; Figure 7 visualizes the means with per-system points overlaid.

Given the small number of systems per type, we read these as observations about the systems evaluated, not pipeline-type effects. The observations below cover cascade and S2S; we report hybrid in the table but set it aside, since two divergent systems give only a range.

F.2.1 Observations

On-time rate (*higher is better*): a sub-metric of the Turn-Taking score (App. E.4.3). We report it directly because its definition is unambiguous: the share of agent turns whose response latency falls in $[200, 2750)$ ms, or $[200, 4000)$ ms when the turn involves a tool call. Note that the boundaries are not exactly the same as in the Turn-Taking score definition, given that here we don’t have the piecewise-linear curves.

Table 24: **Architecture comparison on the four analyzed facets**, computed on clean runs only. On-time turns, conversation completion (Finishes), authentication success (Auth.), and policy violations (Violates). Higher is better for the first three; lower is better for violations.

(a) Per-pipeline-type means with 95% bootstrap intervals over systems-within-type, and the S2S–Cascade pairwise difference (Δ) with its 95% bootstrap interval. n is the number of systems per row. Mean rows use the $[0, 1]$ metric scale; the Δ row is in percentage points ($\Delta \times 100$), with its interval from independently resampling each group’s per-system means.

Pipeline type	n	Facet (mean [95% interval])			
		On-time $\uparrow$	Finishes $\uparrow$	Auth. $\uparrow$	Violates policy $\downarrow$
Cascade	7	0.55 [0.44, 0.66]	0.82 [0.74, 0.93]	0.77 [0.66, 0.86]	0.53 [0.41, 0.65]
Hybrid	2	0.23 [0.01, 0.46]	0.89 [0.87, 0.90]	0.88 [0.84, 0.92]	0.63 [0.56, 0.70]
Speech-to-Speech	3	0.83 [0.80, 0.89]	0.97 [0.95, 0.99]	0.88 [0.83, 0.93]	0.78 [0.69, 0.88]
Δ S2S–Cascade (pp)	—	+27.9 [+16.4, +39.6]	+15.2 [+6.0, +24.5]	+11.3 [+2.4, +20.9]	+24.6 [+11.7, +38.4]

(b) Per-system means, averaged across the three domains. Pipeline-type means and bootstrap intervals in (a) and Fig. 7 are computed from these per-system values.

Pipeline type	System	On-time $\uparrow$	Finishes $\uparrow$	Auth. $\uparrow$	Violates $\downarrow$
Cascade	Nova-3 + GPT-5.4 + Sonic 3	0.48	0.66	0.85	0.28
	Ink-whisper + Haiku-4.5 + Sonic 3	0.44	0.67	0.65	0.48
	Whisper-Large-v3 + Qwen3.5-27B + Voxtral-4B-TTS	0.76	0.75	0.60	0.42
	Cohere-transcribe + Gemma-4-26B + Voxtral-4B-TTS	0.70	0.79	0.68	0.60
	Nova-3 + GPT-5.4-mini + Aura-2	0.66	0.90	0.83	0.73
	Scribe-v2.2-Realtime + Gemini-3-Flash + TTS-Conversational-v3-Alpha	0.49	0.97	0.89	0.61
	Parakeet-1.1 + Gemma-4-31B + Kokoro	0.33	0.97	0.87	0.61
Hybrid	Ultravox-Realtime	0.46	0.87	0.84	0.70
	Gemini-3-Flash + Gemini-3.1-Flash-TTS	0.01	0.90	0.92	0.56
Speech-to-Speech	GPT-Realtime-1.5	0.80	0.99	0.93	0.69
	GPT-Realtime-mini	0.80	0.97	0.83	0.88
	Gemini-3.1-Flash-Live	0.89	0.95	0.88	0.78

All three S2S systems we evaluated sit above every cascade in our sample, but the gap is narrow for the strongest cascades: *Whisper-Large-v3 + Qwen3.5-27B + Voxtral-4B-TTS* (0.76) and *Cohere-transcribe + Gemma-4-26B + Voxtral-4B-TTS* (0.70) come within 4–10 pp of the lowest S2S system (0.80). The remaining five cascades trail more substantially (0.33–0.66).

Conversation Completion (*higher is better*): captures the failure mode where the agent fails to respond to a user turn, leading the conversation to time out. None of the S2S pipelines we evaluated struggled with completing conversations, whereas cascade is more variable. Only two cascade systems reached S2S-comparable values, *Scribe-v2.2-Realtime + Gemini-3-Flash + TTS-Conversational-v3-Alpha* and *Parakeet-1.1 + Gemma-4-31B + Kokoro*.

To understand why conversations do not finish correctly, we examined every conversation ending in an agent-side inactivity timeout, classifying each by the preceding user turn: *short turn* (under 5 words), *short $\wedge$ confirm* (short turn opening with *yes/no/sure/ok*), and *spelled content* (NATO phonetic words or alphanumeric codes). Table 25 breaks timeouts down per system.

Five of seven cascades show short-turn shares

of 52–68% of their timeouts, often beginning with a confirmation token (“Yes, that’s correct”) after which the agent produces no follow-up. Four systems avoid this pattern: *Scribe-v2.2-Realtime + Gemini-3-Flash + TTS-Conversational-v3-Alpha*, *Parakeet-1.1 + Gemma-4-31B + Kokoro*, *Ultravox-Realtime*, and all three S2S systems. All but *Ultravox-Realtime* keep overall failure below 5%. Spelled-content failures persist across all pipelines and dominate timeouts among systems that handle short turns well. These hangs occur predominantly during authentication.

Authentication success (*higher is better*): the proportion of conversations in which the agent successfully authenticates the user before proceeding. This probes how well each architecture handles entity-dense input without transcript-level evaluation.

Raw scores mix two phenomena: how often the agent completes the authentication exchange without hanging, and how often it authenticates correctly when it does. Restricting to conversations that finished correctly separates these (Table 26). S2S systems are essentially unchanged ($\Delta \leq 0.8$ pp), finishing 95–99% of clean conversations. Several cascades move substantially: *Nova-3 + GPT-5.4 + Sonic 3* rises from 0.85 to 0.94 (matching the




Figure 7: **Per-pipeline-type means on the four facets, computed on clean runs only.** Error bars are 95% percentile bootstrap intervals over systems-within-type, shown as descriptive uncertainty bands; faint dots are per-system means. Hybrid ($n = 2$) is shown as individual system points only.

Table 25: **Per-system view on clean runs (1,065 conversations per system).** *Conv. fail rate* = $1 - \text{conversation_correctly_finished}$ (any failure type, conversation-pooled). *Timeouts* = number of those failures that ended with the agent going silent. The three right-most columns are the share of *this system's* timeouts whose preceding user turn matched each pattern. Systems are ordered within each pipeline class by timeout count (descending).

Pipeline	System	Conv. fail rate	Timeouts	Short turn	Short $\wedge$ confirm	Spelled content
Cascade	Ink-whisper + Haiku-4.5 + Sonic 3	36.5%	395	63%	43%	17%
	Nova-3 + GPT-5.4 + Sonic 3	35.8%	381	68%	39%	12%
	Whisper-Large-v3 + Qwen3.5-27B + Voxtral-4B-TTS	28.2%	300	62%	52%	14%
	Cohere-transcribe + Gemma-4-26B + Voxtral-4B-TTS	22.6%	241	56%	37%	21%
	Nova-3 + GPT-5.4-mini + Aura-2	10.7%	117	52%	20%	24%
	Scribe-v2.2-Realtime + Gemini-3-Flash + TTS-Conversational-v3-Alpha	2.6%	28	0%	0%	68%
	Parakeet-1.1 + Gemma-4-31B + Kokoro	2.5%	27	19%	19%	48%
Hybrid	Ultravox-Realtime	14.0%	149	6%	1%	63%
	Gemini-3-Flash + Gemini-3.1-Flash-TTS	10.5%	112	30%	14%	46%
S2S	Gemini-3.1-Flash-Live	4.5%	52	2%	0%	58%
	GPT-Realtime-mini	2.9%	32	0%	0%	66%
	GPT-Realtime-1.5	1.4%	15	0%	0%	73%

best S2S), *Ink-whisper + Haiku-4.5 + Sonic 3* from 0.65 to 0.86, and *Whisper-Large-v3 + Qwen3.5-27B + Voxtral-4B-TTS* from 0.60 to 0.73. The cascade authentication deficit is therefore in large part a conversation completion artifact. *GPT-Realtime-1.5* leads on both axes: it finishes 99% of clean conversations and authenticates correctly 93% of the time.

Policy violations (*lower is better*): the dominant failure mode contributing to low faithfulness scores, since policy adherence requires strict instruction-following. Across all pipelines, 61% of conversations are flagged with at least one policy violation. Common cases include the agent performing write actions without explicit user authorization, or fabricating policies.

GPT-Realtime-1.5 (0.69, the lowest S2S violation rate) lands in the same range as several cas-

cades; the other two S2S systems sit higher. The lowest raw violation rate comes from *Nova-3 + GPT-5.4 + Sonic 3* (0.28). These per-conversation rates are confounded, however: a system that stalls or fails early has fewer occasions to violate a policy in the first place. This shows up in cascades that hang often; *Ink-whisper + Haiku-4.5 + Sonic 3* rises from 0.48 to 0.63 (+15 pp) once restricted to completed conversations, so its low raw rate partly reflects fewer opportunities rather than stronger instruction-following. Because S2S systems progress further than most cascades, their higher rates should be read as a per-conversation difference, not as evidence of weaker instruction-following per se.

Table 26: **Per-system authentication success on clean runs, raw vs. conditional on conversation completion.** *Auth (raw)* is computed over all 1,065 clean conversations; *Auth | finished* is computed only on conversations where *conversation_correctly_finished* = 1. Δ (pp) is their difference. Systems with low finish rates move the most.

Pipeline	System	Finish rate	Auth (raw)	Auth finished	Δ (pp)
Cascade	Nova-3 + GPT-5.4 + Sonic 3	0.66	0.85	0.94	+8.8
	Ink-whisper + Haiku-4.5 + Sonic 3	0.67	0.65	0.86	+20.5
	Whisper-Large-v3 + Qwen3.5-27B + Voxtral-4B-TTS	0.75	0.60	0.73	+12.6
	Cohere-transcribe + Gemma-4-26B + Voxtral-4B-TTS	0.79	0.68	0.73	+4.9
	Nova-3 + GPT-5.4-mini + Aura-2	0.90	0.83	0.86	+2.9
	Scribe-v2.2-Realtime + Gemini-3-Flash + TTS-Conversational-v3-Alpha	0.97	0.89	0.90	+1.1
Hybrid	Parakeet-1.1 + Gemma-4-31B + Kokoro	0.97	0.87	0.88	+0.3
	Ultravox-Realtime	0.87	0.84	0.86	+2.0
	Gemini-3-Flash + Gemini-3.1-Flash-TTS	0.90	0.92	0.94	+2.4
Speech-to-Speech	GPT-Realtime-1.5	0.99	0.93	0.94	+0.5
	GPT-Realtime-mini	0.97	0.83	0.84	+0.8
	Gemini-3.1-Flash-Live	0.95	0.88	0.88	+0.7

F.3 Speech Fidelity

Across all models, **entity mispronunciation** (letter substitutions, digit omissions, spurious insertions, and phonetic confusions between similar-sounding characters) is by far the dominant failure class, accounting for the majority of flagged turns.

Substitutions. In a single alphanumeric code, one character is replaced by a visually or acoustically unrelated one. *Sonic 3* consistently rendered the airport code LAX as LEX (CSM, record 1.2.3) and substituted S C 3 for A C 3 in an ITSM request prefix across multiple turns (record 16). *Aura-2* substituted 1 for I in a medical license number, saying L 1 C instead of L I C (HRSD, record A6). *Ultravox-Realtime* misread the confirmation code ZKLX8E as ZKLXIE, substituting 8 for I (CSM, record 7.2.9), and rendered the hotel voucher prefix GQSIHM as DQSIHM, substituting D for G (record 7.3.1).

Omissions. One or more characters are silently dropped from an entity. *GPT-Realtime-1.5* read the 10-digit NPI 3342331444 as only 9 digits, dropping the trailing 4 (HRSD, record 1.1), and omitted the 6 from registration ID 358607, producing 35807 (record 9.1). *Kokoro* dropped the A from the spelled-out code MEAL across multiple trials, producing MEL (CSM, record 2.2.5), and omitted the letter K from flight number SK915, rendering it as S 915 consistently across three consecutive turns (record 2.2.2). *Aura-2* omitted one of the five repeated digits, saying EMP-0-5-5-5-5 instead of EMP-0-5-5-5-5-5 (ITSM, record 22).

Insertions. Spurious characters or words are injected into an entity that should be reproduced verbatim. *Gemini-3.1-Flash-Live* inserted the name Pavel into the middle of the visa petition number EH23328710672, breaking the code entirely

(HRSD, record 11.2). *Ultravox-Realtime* added a spurious E to the meal voucher prefix in multiple trials, saying MEAL-7EMMHTS-PAX0 instead of MEAL-7MMHTS-PAX0 (CSM, record 2.3.2). *Voxtral-4B-TTS* appended a trailing 3 to the request identifier REQ-SW-9befdac7c2e6, producing REQ-SW-9befdac7c2e63 (ITSM, record 76).

Phonetic confusions. A character is replaced by one that sounds similar when spoken aloud. *GPT-Realtime-1.5* repeatedly confused C and P, rendering the facility code IXC as IXP across two consecutive turns (HRSD, record D3.3), and swapped D for B in the policy number PDZP6L $\rightarrow$ PBZP6L (record T2.1). *Gemini-3.1-Flash-Live* confused Z and V, articulating the DEA prefix ZS as VS (HRSD, record A2). *Kokoro* persistently mispronounced SK130 as S-Cone 130 across two turns (record 4.2.4).

F.4 Turn-Taking & Response Speed

Table 27 reports per-system latency and turn-taking sub-metrics, ultimately contributing to EVA-X pass@1. Three columns give the share of agent turns that arrive early, on-time, or late (summing to 1.0 per system); the response-speed columns give mean per-turn latency in seconds, split by whether the turn involved a tool call.

The timing buckets (early/on-time/late) are hard-thresholded diagnostics, not the piecewise-linear curve used to compute the scored metric, so their exact boundaries differ from the Turn-Taking scoring definition. A turn counts as early when latency < 200 ms, late when latency ≥ 2.75 s (or ≥ 4 s on tool-call turns), and on-time otherwise.

Across systems, the turn-taking score is driven mostly by the on-time and late rates; the early rate is always small and barely varies (largest 0.148, *GPT-Realtime-mini*; most below 0.06), so it separates systems little. Comparing latency with and

Table 27: **Per-system turn-taking and response-speed detail under clean-audio conditions, pooled equal-weighted across the three domains.** Each cell shows the pooled point estimate $\pm$ the percentile bootstrap CI half-width ($\alpha = 0.05$). Turn-taking rates are in $[0, 1]$; response speed is mean per-turn latency in seconds. Cells are shaded per column with darker = better (lower for latency and error-rate columns; higher otherwise).

Arch.	System	EVA-X			Turn-Taking			Response Speed (s)		
		pass@1	Score	Early	On-time	Late	Interrupt.	Overall	w/ tools	w/o tools
Cascade	Cohere + Gemma-4-26B + Voxtral	0.209 ±0.027	0.567 ±0.024	0.113 ±0.009	0.703 ±0.015	0.184 ±0.014	0.110 ±0.009	2.415 ±0.107	2.660 ±0.253	2.171 ±0.077
	Scribe + Gemini-3-Flash + Conversational v3	0.024 ±0.018	0.451 ±0.019	0.008 ±0.002	0.489 ±0.022	0.504 ±0.022	0.006 ±0.002	4.158 ±0.118	5.294 ±0.166	2.420 ±0.119
	Ink-whisper + Haiku-4.5 + Sonic 3	0.009 ±0.006	0.312 ±0.019	0.024 ±0.004	0.440 ±0.015	0.536 ±0.016	0.023 ±0.004	3.399 ±0.045	3.535 ±0.081	3.354 ±0.060
	Nova-3 + GPT-5.4 + Sonic 3	0.007 ±0.006	0.283 ±0.021	0.029 ±0.005	0.484 ±0.017	0.487 ±0.018	0.028 ±0.005	3.943 ±0.112	5.191 ±0.189	2.914 ±0.078
	Nova-3 + GPT-5.4-mini + Aura-2	0.113 ±0.021	0.583 ±0.020	0.042 ±0.006	0.658 ±0.016	0.299 ±0.016	0.041 ±0.006	3.148 ±0.074	3.648 ±0.095	2.377 ±0.113
	Parakeet-1.1 + Gemma-4-31B + Kokoro	0.010 ±0.009	0.308 ±0.015	0.030 ±0.005	0.327 ±0.017	0.643 ±0.018	0.030 ±0.005	4.832 ±0.114	5.373 ±0.128	4.291 ±0.216
	Whisper + Qwen3.5-27B + Voxtral	0.273 ±0.035	0.561 ±0.029	0.056 ±0.008	0.759 ±0.015	0.186 ±0.012	0.054 ±0.008	2.251 ±0.050	2.309 ±0.069	2.164 ±0.067
Hybrid	Gemini-3-Flash + Gemini-3.1-Flash-TTS	0.000 ±0.000	0.019 ±0.003	0.017 ±0.004	0.010 ±0.003	0.973 ±0.005	0.017 ±0.004	7.466 ±0.257	9.100 ±0.413	5.244 ±0.189
	Ultravox-Realtime	0.029 ±0.019	0.417 ±0.021	0.040 ±0.007	0.457 ±0.019	0.503 ±0.020	0.040 ±0.007	4.838 ±0.225	6.994 ±0.301	1.393 ±0.063
S2S	Gemini-3.1-Flash-Live	0.589 ±0.034	0.830 ±0.017	0.011 ±0.005	0.886 ±0.012	0.104 ±0.012	0.011 ±0.005	2.288 ±0.070	2.846 ±0.088	1.288 ±0.086
	GPT-Realtime-1.5	0.566 ±0.040	0.815 ±0.013	0.121 ±0.012	0.799 ±0.016	0.079 ±0.009	0.072 ±0.011	1.798 ±0.061	2.395 ±0.089	1.011 ±0.031
	GPT-Realtime-mini	0.406 ±0.035	0.818 ±0.015	0.148 ±0.013	0.805 ±0.014	0.047 ±0.006	0.106 ±0.011	1.481 ±0.045	1.868 ±0.063	1.067 ±0.040

without tool calls, every system slows on tool-call turns (no overlapping CIs) so the effect is consistent, but its magnitude is system-specific: from roughly 0.1–0.2 s for some cascades (*Whisper-Large-v3* + *Qwen3.5-27B* + *Voxtral-4B-TTS*; *Ink-whisper* + *Haiku-4.5* + *Sonic 3*) up to several seconds for *Ultravox-Realtime* and *Gemini-3-Flash* + *Gemini-3.1-Flash-TTS*. Together these sub-metrics explain much of the EVA-X pass@1 gap between cascade and S2S systems.

F.5 Cross Domain Variability

Tables 29 and 31 report, for every (system, metric) pair, the sample standard deviation (ddof=1) of the per-domain point estimates across the three EVA domains (CSM, ITSM, HR). These complement Tables 30 (per-domain values) and the pooled tables 4, by quantifying how much of the pooled headline number is hiding domain dispersion.

F.5.1 Accuracy and Experience results per domain

Accuracy results are provided for all systems per domain: CSM in Table 28a, HR in Table 28b, and ITSM in Table 28c. Cross-domain variability of accuracy metrics is provided in Table 29.

Experience results are provided for all systems per domain: CSM, HR and ITSM in Table 30a, 30b, and 30c respectively. Cross-domain variability of experience metrics is provided in Table 31.

Table 28: **Accuracy metrics for evaluated systems under clean-audio conditions across CSM, HR, and ITSM domains.** Each subtable shows results for $k = 5$ trials. Each cell shows the point estimate $\pm$ the percentile bootstrap CI half-width ($\alpha = 0.05$). The three pass-rate columns share a single shading scale (so pass@1 vs. pass@k vs. pass^k are visually comparable); each submetric column is scaled independently. Darker = higher point estimate.

(a) CSM Domain

Arch.	System	EVA-A			Task Completion	Faithfulness	Speech Fidelity
		pass@1	pass@k	pass^k	Mean	Mean	Mean
Cascade	Cohere-transcribe + Gemma-4-26B + Voxtral-4B-TTS	0.246 $\pm$ 0.084	0.500 $\pm$ 0.140	0.066 $\pm$ 0.061	0.368 $\pm$ 0.092	0.398 $\pm$ 0.070	0.989 $\pm$ 0.006
	Scribe-v2.2-Realtime + Gemini-3-Flash + TTS-Conversational-v3-Alpha	0.656 $\pm$ 0.092	0.900 $\pm$ 0.080	0.385 $\pm$ 0.122	0.808 $\pm$ 0.080	0.666 $\pm$ 0.088	0.981 $\pm$ 0.012
	Ink-whisper + Haiku-4.5 + Sonic 3	0.272 $\pm$ 0.076	0.640 $\pm$ 0.140	0.065 $\pm$ 0.064	0.440 $\pm$ 0.092	0.374 $\pm$ 0.068	0.989 $\pm$ 0.004
	Nova-3 + GPT-5.4 + Sonic 3	0.628 $\pm$ 0.076	0.940 $\pm$ 0.060	0.278 $\pm$ 0.108	0.732 $\pm$ 0.068	0.784 $\pm$ 0.050	0.996 $\pm$ 0.003
	Nova-3 + GPT-5.4-mini + Aura-2	0.216 $\pm$ 0.096	0.460 $\pm$ 0.140	0.098 $\pm$ 0.081	0.456 $\pm$ 0.092	0.298 $\pm$ 0.078	0.979 $\pm$ 0.009
	Parakeet-1.1 + Gemma-4-31B + Kokoro	0.540 $\pm$ 0.096	0.880 $\pm$ 0.100	0.294 $\pm$ 0.111	0.672 $\pm$ 0.096	0.600 $\pm$ 0.076	0.966 $\pm$ 0.019
	Whisper-Large-v3 + Qwen3.5-27B + Voxtral-4B-TTS	0.272 $\pm$ 0.071	0.680 $\pm$ 0.140	0.045 $\pm$ 0.049	0.470 $\pm$ 0.072	0.470 $\pm$ 0.072	0.918 $\pm$ 0.017
Hybrid	Gemini-3-Flash + Gemini-3.1-Flash-TTS	0.488 $\pm$ 0.092	0.840 $\pm$ 0.100	0.188 $\pm$ 0.089	0.696 $\pm$ 0.088	0.478 $\pm$ 0.080	0.965 $\pm$ 0.012
	Ultravox-Realtime	0.324 $\pm$ 0.100	0.540 $\pm$ 0.140	0.131 $\pm$ 0.081	0.428 $\pm$ 0.104	0.334 $\pm$ 0.076	0.971 $\pm$ 0.015
S2S	Gemini-3.1-Flash-Live	0.356 $\pm$ 0.104	0.660 $\pm$ 0.140	0.170 $\pm$ 0.102	0.504 $\pm$ 0.096	0.342 $\pm$ 0.084	1.000 $\pm$ 0.000
	GPT-Realtime-1.5	0.424 $\pm$ 0.116	0.640 $\pm$ 0.140	0.271 $\pm$ 0.120	0.540 $\pm$ 0.112	0.424 $\pm$ 0.092	0.998 $\pm$ 0.003
	GPT-Realtime-mini	0.176 $\pm$ 0.092	0.300 $\pm$ 0.120	0.085 $\pm$ 0.075	0.288 $\pm$ 0.104	0.164 $\pm$ 0.076	0.971 $\pm$ 0.034

(b) HR Domain

Arch.	System	EVA-A			Task Completion	Faithfulness	Speech Fidelity
		pass@1	pass@k	pass^k	Mean	Mean	Mean
Cascade	Cohere-transcribe + Gemma-4-26B + Voxtral-4B-TTS	0.229 $\pm$ 0.070	0.386 $\pm$ 0.108	0.087 $\pm$ 0.049	0.318 $\pm$ 0.089	0.408 $\pm$ 0.055	0.985 $\pm$ 0.005
	Scribe-v2.2-Realtime + Gemini-3-Flash + TTS-Conversational-v3-Alpha	0.373 $\pm$ 0.082	0.614 $\pm$ 0.108	0.199 $\pm$ 0.074	0.655 $\pm$ 0.080	0.320 $\pm$ 0.066	0.973 $\pm$ 0.009
	Ink-whisper + Haiku-4.5 + Sonic 3	0.239 $\pm$ 0.063	0.482 $\pm$ 0.108	0.052 $\pm$ 0.033	0.318 $\pm$ 0.077	0.700 $\pm$ 0.052	0.982 $\pm$ 0.006
	Nova-3 + GPT-5.4 + Sonic 3	0.422 $\pm$ 0.080	0.651 $\pm$ 0.108	0.213 $\pm$ 0.073	0.496 $\pm$ 0.087	0.740 $\pm$ 0.033	0.986 $\pm$ 0.005
	Nova-3 + GPT-5.4-mini + Aura-2	0.217 $\pm$ 0.060	0.434 $\pm$ 0.108	0.042 $\pm$ 0.032	0.455 $\pm$ 0.084	0.298 $\pm$ 0.051	0.979 $\pm$ 0.006
	Parakeet-1.1 + Gemma-4-31B + Kokoro	0.318 $\pm$ 0.065	0.651 $\pm$ 0.096	0.083 $\pm$ 0.051	0.590 $\pm$ 0.082	0.386 $\pm$ 0.057	0.946 $\pm$ 0.016
	Whisper-Large-v3 + Qwen3.5-27B + Voxtral-4B-TTS	0.167 $\pm$ 0.058	0.337 $\pm$ 0.096	0.042 $\pm$ 0.032	0.335 $\pm$ 0.082	0.651 $\pm$ 0.048	0.899 $\pm$ 0.021
Hybrid	Gemini-3-Flash + Gemini-3.1-Flash-TTS	0.381 $\pm$ 0.070	0.783 $\pm$ 0.096	0.135 $\pm$ 0.062	0.660 $\pm$ 0.063	0.380 $\pm$ 0.052	0.982 $\pm$ 0.005
	Ultravox-Realtime	0.272 $\pm$ 0.072	0.506 $\pm$ 0.108	0.125 $\pm$ 0.069	0.523 $\pm$ 0.092	0.229 $\pm$ 0.046	0.973 $\pm$ 0.008
S2S	Gemini-3.1-Flash-Live	0.248 $\pm$ 0.075	0.434 $\pm$ 0.108	0.111 $\pm$ 0.060	0.501 $\pm$ 0.082	0.164 $\pm$ 0.051	0.994 $\pm$ 0.005
	GPT-Realtime-1.5	0.349 $\pm$ 0.084	0.614 $\pm$ 0.108	0.176 $\pm$ 0.078	0.812 $\pm$ 0.060	0.230 $\pm$ 0.055	0.994 $\pm$ 0.004
	GPT-Realtime-mini	0.106 $\pm$ 0.051	0.229 $\pm$ 0.096	0.029 $\pm$ 0.029	0.369 $\pm$ 0.080	0.080 $\pm$ 0.031	0.978 $\pm$ 0.008

(c) ITSM Domain

Arch.	System	EVA-A			Task Completion	Faithfulness	Speech Fidelity
		pass@1	pass@k	pass^k	Mean	Mean	Mean
Cascade	Cohere-transcribe + Gemma-4-26B + Voxtral-4B-TTS	0.147 $\pm$ 0.055	0.362 $\pm$ 0.100	0.028 $\pm$ 0.030	0.328 $\pm$ 0.073	0.319 $\pm$ 0.063	0.974 $\pm$ 0.007
	Scribe-v2.2-Realtime + Gemini-3-Flash + TTS-Conversational-v3-Alpha	0.440 $\pm$ 0.082	0.675 $\pm$ 0.100	0.222 $\pm$ 0.086	0.745 $\pm$ 0.067	0.385 $\pm$ 0.064	0.977 $\pm$ 0.008
	Ink-whisper + Haiku-4.5 + Sonic 3	0.190 $\pm$ 0.067	0.425 $\pm$ 0.112	0.054 $\pm$ 0.050	0.362 $\pm$ 0.065	0.480 $\pm$ 0.054	0.977 $\pm$ 0.006
	Nova-3 + GPT-5.4 + Sonic 3	0.463 $\pm$ 0.068	0.838 $\pm$ 0.088	0.162 $\pm$ 0.066	0.597 $\pm$ 0.073	0.738 $\pm$ 0.054	0.985 $\pm$ 0.005
	Nova-3 + GPT-5.4-mini + Aura-2	0.198 $\pm$ 0.060	0.450 $\pm$ 0.112	0.045 $\pm$ 0.042	0.482 $\pm$ 0.082	0.215 $\pm$ 0.044	0.962 $\pm$ 0.008
	Parakeet-1.1 + Gemma-4-31B + Kokoro	0.350 $\pm$ 0.075	0.713 $\pm$ 0.100	0.131 $\pm$ 0.068	0.650 $\pm$ 0.088	0.414 $\pm$ 0.057	0.951 $\pm$ 0.014
	Whisper-Large-v3 + Qwen3.5-27B + Voxtral-4B-TTS	0.175 $\pm$ 0.045	0.537 $\pm$ 0.113	0.012 $\pm$ 0.010	0.412 $\pm$ 0.085	0.516 $\pm$ 0.060	0.923 $\pm$ 0.015
Hybrid	Gemini-3-Flash + Gemini-3.1-Flash-TTS	0.425 $\pm$ 0.073	0.812 $\pm$ 0.088	0.150 $\pm$ 0.069	0.665 $\pm$ 0.067	0.470 $\pm$ 0.049	0.960 $\pm$ 0.011
	Ultravox-Realtime	0.215 $\pm$ 0.065	0.463 $\pm$ 0.113	0.068 $\pm$ 0.049	0.467 $\pm$ 0.092	0.312 $\pm$ 0.061	0.969 $\pm$ 0.008
S2S	Gemini-3.1-Flash-Live	0.273 $\pm$ 0.075	0.562 $\pm$ 0.112	0.115 $\pm$ 0.075	0.412 $\pm$ 0.082	0.209 $\pm$ 0.051	0.991 $\pm$ 0.008
	GPT-Realtime-1.5	0.628 $\pm$ 0.080	0.875 $\pm$ 0.075	0.402 $\pm$ 0.098	0.865 $\pm$ 0.053	0.426 $\pm$ 0.055	0.996 $\pm$ 0.004
	GPT-Realtime-mini	0.207 $\pm$ 0.065	0.425 $\pm$ 0.112	0.063 $\pm$ 0.050	0.378 $\pm$ 0.090	0.133 $\pm$ 0.040	0.981 $\pm$ 0.013

Table 29: **Cross-domain variability of accuracy metrics:** each cell is the sample standard deviation (ddof=1) of the per-domain point estimates across the three EVA domains (CSM, ITSM, HR). Larger values indicate the system’s performance depends more strongly on domain. K = 5.

Arch.	System	EVA-A			Task Completion	Faithfulness	Speech Fidelity
		pass@1	pass@k	pass^k	SD	SD	SD
Cascade	Cohere-transcribe + Gemma-4-26B + Voxtral-4B-TTS	0.053	0.074	0.030	0.027	0.049	0.008
	Scribe-v2.2-Realtime + Gemini-3-Flash + TTS-Conversational-v3-Alpha	0.148	0.150	0.101	0.077	0.184	0.004
	Ink-whisper + Haiku-4.5 + Sonic 3	0.041	0.111	0.007	0.062	0.166	0.006
	Nova-3 + GPT-5.4 + Sonic 3	0.109	0.147	0.058	0.118	0.026	0.006
	Nova-3 + GPT-5.4-mini + Aura-2	0.011	0.013	0.032	0.015	0.048	0.010
	Parakeet-1.1 + Gemma-4-31B + Kokoro	0.120	0.119	0.111	0.042	0.117	0.010
	Whisper-Large-v3 + Qwen3.5-27B + Voxtral-4B-TTS	0.058	0.172	0.019	0.085	0.094	0.013
Hybrid	Gemini-3-Flash + Gemini-3.1-Flash-TTS	0.054	0.028	0.027	0.019	0.055	0.011
	Ultravox-Realtime	0.055	0.039	0.035	0.048	0.056	0.002
S2S	Gemini-3.1-Flash-Live	0.057	0.113	0.033	0.052	0.093	0.005
	GPT-Realtime-1.5	0.144	0.144	0.113	0.174	0.113	0.002
	GPT-Realtime-mini	0.052	0.099	0.028	0.049	0.043	0.005

Table 30: **Experience metrics for evaluated systems under clean-audio conditions across CSM, HR, and ITSM domains.** Each subtable shows results for k=5 trials. Each cell shows the point estimate $\pm$ the percentile bootstrap CI half-width ($\alpha = 0.05$). The three pass-rate columns share a single shading scale (so pass@1 vs. pass@k vs. pass^k are visually comparable); each submetric column is scaled independently. Darker = higher point estimate.

(a) CSM Domain

Arch.	System	EVA-X			Turn-Taking	Conciseness	Conv. Progression
		pass@1	pass@k	pass^k	Mean	Mean	Mean
Cascade	Cohere-transcribe + Gemma-4-26B + Voxtral-4B-TTS	0.220 $\pm$ 0.052	0.680 $\pm$ 0.120	0.014 $\pm$ 0.014	0.664 $\pm$ 0.035	0.790 $\pm$ 0.012	0.572 $\pm$ 0.068
	Scribe-v2.2-Realtime + Gemini-3-Flash + TTS-Conversational-v3-Alpha	0.052 $\pm$ 0.048	0.120 $\pm$ 0.100	0.009 $\pm$ 0.015	0.469 $\pm$ 0.041	0.792 $\pm$ 0.013	0.832 $\pm$ 0.048
	Ink-whisper + Haiku-4.5 + Sonic 3	0.008 $\pm$ 0.012	0.040 $\pm$ 0.060	0.000 $\pm$ 0.000	0.391 $\pm$ 0.028	0.755 $\pm$ 0.015	0.682 $\pm$ 0.048
	Nova-3 + GPT-5.4 + Sonic 3	0.004 $\pm$ 0.008	0.020 $\pm$ 0.040	0.000 $\pm$ 0.000	0.291 $\pm$ 0.034	0.825 $\pm$ 0.014	0.796 $\pm$ 0.042
	Nova-3 + GPT-5.4-mini + Aura-2	0.108 $\pm$ 0.044	0.440 $\pm$ 0.140	0.003 $\pm$ 0.005	0.577 $\pm$ 0.036	0.812 $\pm$ 0.018	0.440 $\pm$ 0.056
	Parakeet-1.1 + Gemma-4-31B + Kokoro	0.024 $\pm$ 0.028	0.080 $\pm$ 0.080	0.000 $\pm$ 0.001	0.274 $\pm$ 0.034	0.842 $\pm$ 0.014	0.832 $\pm$ 0.040
	Whisper-Large-v3 + Qwen3.5-27B + Voxtral-4B-TTS	0.224 $\pm$ 0.064	0.700 $\pm$ 0.120	0.021 $\pm$ 0.021	0.661 $\pm$ 0.031	0.654 $\pm$ 0.019	0.558 $\pm$ 0.048
Hybrid	Gemini-3-Flash + Gemini-3.1-Flash-TTS	0.000 $\pm$ 0.000	0.000 $\pm$ 0.000	0.000 $\pm$ 0.000	0.031 $\pm$ 0.009	0.805 $\pm$ 0.013	0.634 $\pm$ 0.062
	Ultravox-Realtime	0.048 $\pm$ 0.052	0.120 $\pm$ 0.100	0.013 $\pm$ 0.020	0.483 $\pm$ 0.036	0.711 $\pm$ 0.020	0.540 $\pm$ 0.054
S2S	Gemini-3.1-Flash-Live	0.504 $\pm$ 0.064	1.000 $\pm$ 0.000	0.119 $\pm$ 0.067	0.788 $\pm$ 0.039	0.806 $\pm$ 0.020	0.702 $\pm$ 0.056
	GPT-Realtime-1.5	0.560 $\pm$ 0.080	0.940 $\pm$ 0.080	0.194 $\pm$ 0.083	0.821 $\pm$ 0.029	0.783 $\pm$ 0.017	0.634 $\pm$ 0.052
	GPT-Realtime-mini	0.376 $\pm$ 0.076	0.900 $\pm$ 0.100	0.089 $\pm$ 0.064	0.801 $\pm$ 0.036	0.705 $\pm$ 0.016	0.406 $\pm$ 0.062

(b) HR Domain

Arch.	System	EVA-X			Turn-Taking	Conciseness	Conv. Progression
		pass@1	pass@k	pass^k	Mean	Mean	Mean
Cascade	Cohere-transcribe + Gemma-4-26B + Voxtral-4B-TTS	0.241 $\pm$ 0.043	0.687 $\pm$ 0.096	0.014 $\pm$ 0.010	0.549 $\pm$ 0.046	0.818 $\pm$ 0.011	0.584 $\pm$ 0.052
	Scribe-v2.2-Realtime + Gemini-3-Flash + TTS-Conversational-v3-Alpha	0.012 $\pm$ 0.024	0.024 $\pm$ 0.036	0.004 $\pm$ 0.008	0.469 $\pm$ 0.023	0.774 $\pm$ 0.010	0.788 $\pm$ 0.030
	Ink-whisper + Haiku-4.5 + Sonic 3	0.002 $\pm$ 0.005	0.012 $\pm$ 0.024	0.000 $\pm$ 0.000	0.218 $\pm$ 0.037	0.793 $\pm$ 0.012	0.741 $\pm$ 0.033
	Nova-3 + GPT-5.4 + Sonic 3	0.012 $\pm$ 0.014	0.048 $\pm$ 0.048	0.000 $\pm$ 0.000	0.273 $\pm$ 0.039	0.859 $\pm$ 0.011	0.729 $\pm$ 0.031
	Nova-3 + GPT-5.4-mini + Aura-2	0.130 $\pm$ 0.036	0.458 $\pm$ 0.108	0.004 $\pm$ 0.003	0.594 $\pm$ 0.032	0.861 $\pm$ 0.008	0.405 $\pm$ 0.035
	Parakeet-1.1 + Gemma-4-31B + Kokoro	0.002 $\pm$ 0.005	0.012 $\pm$ 0.024	0.000 $\pm$ 0.000	0.344 $\pm$ 0.020	0.825 $\pm$ 0.013	0.701 $\pm$ 0.046
	Whisper-Large-v3 + Qwen3.5-27B + Voxtral-4B-TTS	0.258 $\pm$ 0.058	0.614 $\pm$ 0.108	0.053 $\pm$ 0.033	0.481 $\pm$ 0.060	0.673 $\pm$ 0.019	0.625 $\pm$ 0.045
Hybrid	Gemini-3-Flash + Gemini-3.1-Flash-TTS	0.000 $\pm$ 0.000	0.000 $\pm$ 0.000	0.000 $\pm$ 0.000	0.014 $\pm$ 0.004	0.795 $\pm$ 0.012	0.561 $\pm$ 0.049
	Ultravox-Realtime	0.022 $\pm$ 0.024	0.072 $\pm$ 0.060	0.004 $\pm$ 0.008	0.432 $\pm$ 0.029	0.742 $\pm$ 0.019	0.453 $\pm$ 0.051
S2S	Gemini-3.1-Flash-Live	0.622 $\pm$ 0.055	1.000 $\pm$ 0.000	0.261 $\pm$ 0.078	0.861 $\pm$ 0.021	0.798 $\pm$ 0.012	0.570 $\pm$ 0.042
	GPT-Realtime-1.5	0.593 $\pm$ 0.060	0.964 $\pm$ 0.048	0.244 $\pm$ 0.072	0.825 $\pm$ 0.016	0.807 $\pm$ 0.010	0.696 $\pm$ 0.041
	GPT-Realtime-mini	0.354 $\pm$ 0.055	0.867 $\pm$ 0.084	0.075 $\pm$ 0.048	0.827 $\pm$ 0.019	0.725 $\pm$ 0.014	0.313 $\pm$ 0.042

(c) ITSM Domain

Arch.	System	EVA-X			Turn-Taking	Conciseness	Conv. Progression
		pass@1	pass@k	pass^k	Mean	Mean	Mean
Cascade	Cohere-transcribe + Gemma-4-26B + Voxtral-4B-TTS	0.168 $\pm$ 0.042	0.575 $\pm$ 0.113	0.016 $\pm$ 0.026	0.487 $\pm$ 0.044	0.821 $\pm$ 0.010	0.639 $\pm$ 0.041
	Scribe-v2.2-Realtime + Gemini-3-Flash + TTS-Conversational-v3-Alpha	0.007 $\pm$ 0.010	0.037 $\pm$ 0.050	0.000 $\pm$ 0.000	0.415 $\pm$ 0.029	0.755 $\pm$ 0.012	0.793 $\pm$ 0.035
	Ink-whisper + Haiku-4.5 + Sonic 3	0.018 $\pm$ 0.015	0.075 $\pm$ 0.063	0.000 $\pm$ 0.000	0.328 $\pm$ 0.035	0.804 $\pm$ 0.008	0.708 $\pm$ 0.038
	Nova-3 + GPT-5.4 + Sonic 3	0.005 $\pm$ 0.008	0.025 $\pm$ 0.037	0.000 $\pm$ 0.000	0.285 $\pm$ 0.029	0.821 $\pm$ 0.011	0.688 $\pm$ 0.031
	Nova-3 + GPT-5.4-mini + Aura-2	0.100 $\pm$ 0.038	0.350 $\pm$ 0.113	0.007 $\pm$ 0.009	0.577 $\pm$ 0.030	0.831 $\pm$ 0.011	0.439 $\pm$ 0.038
	Parakeet-1.1 + Gemma-4-31B + Kokoro	0.003 $\pm$ 0.005	0.013 $\pm$ 0.025	0.000 $\pm$ 0.000	0.306 $\pm$ 0.021	0.821 $\pm$ 0.011	0.787 $\pm$ 0.039
	Whisper-Large-v3 + Qwen3.5-27B + Voxtral-4B-TTS	0.338 $\pm$ 0.062	0.738 $\pm$ 0.100	0.079 $\pm$ 0.050	0.539 $\pm$ 0.058	0.727 $\pm$ 0.015	0.651 $\pm$ 0.036
Hybrid	Gemini-3-Flash + Gemini-3.1-Flash-TTS	0.000 $\pm$ 0.000	0.000 $\pm$ 0.000	0.000 $\pm$ 0.000	0.013 $\pm$ 0.003	0.803 $\pm$ 0.012	0.659 $\pm$ 0.039
	Ultravox-Realtime	0.018 $\pm$ 0.020	0.050 $\pm$ 0.050	0.001 $\pm$ 0.002	0.336 $\pm$ 0.042	0.797 $\pm$ 0.016	0.294 $\pm$ 0.045
S2S	Gemini-3.1-Flash-Live	0.643 $\pm$ 0.068	0.938 $\pm$ 0.062	0.340 $\pm$ 0.088	0.840 $\pm$ 0.026	0.799 $\pm$ 0.014	0.637 $\pm$ 0.047
	GPT-Realtime-1.5	0.545 $\pm$ 0.065	0.912 $\pm$ 0.063	0.208 $\pm$ 0.067	0.798 $\pm$ 0.023	0.812 $\pm$ 0.014	0.708 $\pm$ 0.042
	GPT-Realtime-mini	0.487 $\pm$ 0.057	0.912 $\pm$ 0.062	0.134 $\pm$ 0.052	0.825 $\pm$ 0.020	0.736 $\pm$ 0.017	0.445 $\pm$ 0.050

Table 31: **Cross-domain variability of experience metrics:** each cell is the sample standard deviation (ddof=1) of the per-domain point estimates across the three EVA domains (CSM, ITSM, HR). Larger values indicate the system’s performance depends more strongly on domain. K = 5.

Arch.	System	EVA-X			Turn-Taking	Conciseness	Conv. Progression
		pass@1	pass@k	pass^k	SD	SD	SD
Cascade	Cohere-transcribe + Gemma-4-26B + Voxtral-4B-TTS	0.038	0.063	0.001	0.090	0.017	0.036
	Scribe-v2.2-Realtime + Gemini-3-Flash + TTS-Conversational-v3-Alpha	0.024	0.052	0.004	0.031	0.018	0.024
	Ink-whisper + Haiku-4.5 + Sonic 3	0.008	0.032	0.000	0.088	0.026	0.030
	Nova-3 + GPT-5.4 + Sonic 3	0.004	0.015	0.000	0.009	0.021	0.055
	Nova-3 + GPT-5.4-mini + Aura-2	0.016	0.058	0.002	0.010	0.024	0.020
	Parakeet-1.1 + Gemma-4-31B + Kokoro	0.012	0.039	0.000	0.035	0.011	0.067
	Whisper-Large-v3 + Qwen3.5-27B + Voxtral-4B-TTS	0.058	0.063	0.029	0.092	0.038	0.048
Hybrid	Gemini-3-Flash + Gemini-3.1-Flash-TTS	0.000	0.000	0.000	0.010	0.005	0.051
	Ultravox-Realtime	0.017	0.036	0.006	0.075	0.043	0.125
S2S	Gemini-3.1-Flash-Live	0.075	0.036	0.112	0.037	0.004	0.066
	GPT-Realtime-1.5	0.024	0.026	0.026	0.015	0.015	0.040
	GPT-Realtime-mini	0.071	0.023	0.031	0.014	0.016	0.068

G Perturbation Analysis

Beyond standard evaluation conditions, EVA-Bench supports robustness testing through a structured perturbation system applied to the user simulator. Perturbations operate along three axes (behavior, accent, and audio degradation) and can be composed to simulate realistic deployment conditions that clean-audio benchmarks do not capture. All perturbations are applied exclusively to the user simulator; the agent under evaluation receives no special configuration, ensuring that robustness scores reflect genuine system sensitivity rather than evaluation artifacts.

Below we describe the robustness testing options available in EVA-Bench as well as additional results of the perturbation testing experiment described in the main body of the paper.

G.1 Perturbation Types

Behavioral Perturbations. The three behavioral personas (*aggressive_impatient*, *elderly_slow*, and *forgetful_disorganized*) constitute the behavioral perturbation axis. Each modifies both the simulator’s conversational prompt and its voice to produce acoustically and behaviorally consistent speech. Behavioral perturbations stress-test the agent’s ability to handle non-canonical caller patterns: interruptions and rapid speech, slow delivery with long silences, and disfluency-laden turns with mid-utterance corrections.

Accent Perturbations. EVA-Bench supports four non-native English accent variants for the user simulator: French, Indian, Spanish, and Chinese. Each accent uses a dedicated voice to ensure phonetic authenticity. Accent perturbations are mutually exclusive with behavioral perturbations, as both require a specific voice id. Together, the four accents allow systematic evaluation of speech recognition robustness across a representative sample of global English speaker populations commonly encountered in enterprise deployments.

Audio Degradation. Two audio degradation mechanisms are supported and can be composed with any behavioral or accent perturbation. First, *background noise* mixes an ambient audio track into the user simulator’s speech at a configurable signal-to-noise ratio (default: 15 dB). Eight noise environments are included: airport gate, baby crying, background music, bad connection static, coffee shop, loud construction, NYC street, and road noise — spanning the range of acoustic environ-

ments in which enterprise voice agents are commonly exposed to. Second, *connection degradation* applies a stack of VoIP artifacts — codec compression, simulated packet loss, and volume fluctuation — on top of any other active perturbation, reproducing the degradation characteristic of real telephony infrastructure.

G.2 Perturbation Experiment

Methods. For each model, metric, and domain (or pooled across domains), we test whether each perturbation condition - accent, background noise, and the combination (accent + background noise) - alters performance as compared to the clean baseline. We analyze scenario-level mean deltas between a perturbation run and a paired baseline: within each scenario, performance is averaged across trials (3 trials per perturbation condition, 5 trials per clean baseline) and the paired difference is computed as $\delta = \bar{x}_{\text{perturbation}} - \bar{x}_{\text{baseline}}$, yielding one δ per scenario.

We assess each perturbation with a paired sign-flip permutation test on scenario-level deltas. This approach makes no distributional assumptions and is appropriate for the bounded, ordinal nature of our metrics. Under the null hypothesis of no perturbation effect ($\mathbb{E}[\delta] = 0$), the sign of each paired difference is exchangeable; we independently flip the sign of each scenario delta with probability 0.5 and recompute the mean across 10,000 permutations to construct the null distribution. We report two-sided p -values, defined as the fraction of permutations whose absolute mean is at least as large as the observed $|\bar{\delta}|$. Each estimate is accompanied by a 95% percentile bootstrap confidence interval on the mean delta, computed by resampling scenarios with replacement (1,000 bootstrap samples).

To control for multiple comparisons across the three perturbation conditions within each model $\times$ metric $\times$ domain combination, raw p -values are Holm–Bonferroni corrected. Effects are considered significant when the corrected p -value is below $\alpha = 0.05$.

Results. Within the perturbation effect plotted in Figure 8, models are always listed in the following order (left to right): **Cascade:** Cohere-transcribe + Gemma-4-26B + Voxtral-4B-TTS, Scribe + Gemini-3-Flash + Conversational v3, Ink-whisper + Haiku-4.5 + Sonic 3, Nova-3 + GPT-5.4 + Sonic 3, Nova-3 + GPT-5.4-mini + Aura-2, Parakeet-1.1 + Gemma-4-31B + Kokoro, Whisper-Large-v3 + Qwen3.5-27B + Voxtral-4B-TTS; **Hy-**

brid: Gemini-3-Flash + Gemini-3.1-Flash-TTS, Ultravox-Realtime; **S2S:** Gemini-3.1-Flash-Live, GPT-Realtime-1.5, GPT-Realtime-mini. The complete results are shared in Table 32a, Table 32b, and Table 32c.

train of thought and need a moment to remember what you were saying. You also make mistakes when you speak and have to repeat yourself (ex. "hmm yeah one second... let me find that... ok its A E 2 B oh wait sorry actually its A F 2 B")

Default Persona

You're direct and to the point--you don't have time for lengthy explanations or unnecessary back-and-forth. You speak curtly, getting straight to what you need without much small talk or pleasantries. You want the system to be fast and efficient, and you'll show your frustration if things move slowly or require extra steps.

Elderly & Slow

You are elderly and have difficulty understanding fast speech. You speak extremely slowly, with frequent deliberate pauses. You occasionally ask the agent to repeat themselves slowly. You do not rush. You frequently use ellipses (...) in your output to indicate pauses. Ex. ("Ok yes... my confirmation code is... W... K... 2... E... X... B...")

Aggressive & Impatient

You are impatient and easily frustrated when the agent does not resolve your requests immediately. You speak very quickly and often interrupt the agent mid-sentence when they are talking for too long to make your frustration clear and ask them to hurry it up. Or you ask why they are taking so long if there is a long silence. Express your frustration whenever progress is not being made and remember to interrupt often. You frequently output words in all caps to indicate your frustration and add emphasis.

Forgetful & Disorganized

You are forgetful and prone to disfluencies (um..., uh..., huh..., let me think..., hold on a second..., let me find that piece of information..., etc). You frequently use ellipses (...) in your output to indicate pauses. You often forget the information you need and have to search for it in the middle of your turn. Simple things like your name and date of birth you remember easily, but for any specific codes and IDs you need a couple of seconds to find it. You often lose your

Table 32: **Perturbation effects across acoustic conditions: pooled mean Δ (perturbed – clean) per model and metric.** Cells are shaded by magnitude (red = degradation, green = improvement). Significance from sign-flip permutation tests with Holm–Bonferroni correction within each (model, metric) family of three conditions: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

(a) Accent


Model	EVA-A	EVA-X	TaskComp	Faith	SpeechFidelity	ConvProg	TurnTake	Concise
<i>Cascade</i>								
Cohere-transcribe + Gemma-4-26B + Voxtral-4B-TTS	−0.090**	−0.073*	−0.176***	+0.106**	+0.006	+0.020	−0.165***	−0.015*
Scribe + Gemini-3-Flash + Conversational v3	−0.058	+0.001	−0.001	−0.024	−0.012	−0.000	+0.006	+0.006
Ink-whisper + Haiku-4.5 + Sonic 3	−0.103***	−0.013	−0.109***	+0.035	−0.009	−0.027	−0.124***	−0.031**
Nova-3 + GPT-5.4 + Sonic 3	−0.135***	−0.004	−0.144***	−0.007	−0.003	+0.019	−0.160***	+0.003
Nova-3 + GPT-5.4-mini + Aura-2	−0.054*	−0.046*	−0.121***	+0.060*	−0.016*	−0.029	−0.164***	−0.018*
Parakeet-1.1 + Gemma-4-31B + Kokoro	−0.040	+0.011	−0.033	+0.030	+0.003	−0.055	+0.121***	−0.003
Whisper-Large-v3 + Qwen3.5-27B + Voxtral-4B-TTS	−0.074**	−0.161***	−0.128***	+0.085*	+0.015	−0.027	−0.165***	−0.004
<i>Hybrid</i>								
Ultravox-Realtime	+0.008	−0.008	−0.019	+0.023	+0.004	+0.097**	−0.015	−0.002
Gemini-3-Flash + Gemini-3.1-Flash-TTS	−0.116**	+0.000	−0.152**	−0.139***	+0.008	−0.150***	+0.096***	−0.014
<i>S2S</i>								
Gemini-3.1-Flash-Live	−0.014	+0.033	−0.007	−0.023	+0.005	−0.028	+0.028	−0.001
GPT-Realtime-1.5	+0.021	+0.033	+0.041	+0.013	−0.002	+0.030	+0.014	−0.000
GPT-Realtime-mini	+0.016	+0.035	−0.039	+0.019	+0.018	−0.027	+0.025	+0.004

(b) Background Noise


Model	EVA-A	EVA-X	TaskComp	Faith	SpeechFidelity	ConvProg	TurnTake	Concise
<i>Cascade</i>								
Cohere-transcribe + Gemma-4-26B + Voxtral-4B-TTS	−0.046	+0.079*	−0.072*	+0.006	+0.032*	−0.070*	+0.097***	−0.023***
Scribe + Gemini-3-Flash + Conversational v3	−0.043	−0.014	−0.016	−0.032	−0.003	−0.007	−0.065***	−0.005
Ink-whisper + Haiku-4.5 + Sonic 3	−0.051	−0.010	+0.002	+0.005	−0.007	−0.071*	−0.056**	−0.022**
Nova-3 + GPT-5.4 + Sonic 3	−0.187***	−0.004	−0.199***	+0.010	−0.005	−0.007	−0.162***	−0.005
Nova-3 + GPT-5.4-mini + Aura-2	−0.095**	−0.050*	−0.187***	+0.118***	−0.002	+0.036	−0.227***	+0.002
Parakeet-1.1 + Gemma-4-31B + Kokoro	−0.066	−0.004	−0.059	+0.010	+0.004	−0.021	−0.163***	+0.001
Whisper-Large-v3 + Qwen3.5-27B + Voxtral-4B-TTS	−0.104***	−0.206***	−0.195***	−0.013	+0.024	−0.175***	−0.281***	−0.033**
<i>Hybrid</i>								
Ultravox-Realtime	+0.041	−0.016	−0.004	+0.020	+0.008	+0.077*	−0.044**	−0.011
Gemini-3-Flash + Gemini-3.1-Flash-TTS	−0.056	+0.000	−0.004	−0.092**	−0.001	−0.032	+0.060***	−0.006
<i>S2S</i>								
Gemini-3.1-Flash-Live	−0.062*	−0.115**	−0.033	−0.018	−0.001	−0.043	−0.057*	−0.006
GPT-Realtime-1.5	−0.082*	−0.230***	−0.055	−0.050	−0.011	−0.105***	−0.102***	−0.018
GPT-Realtime-mini	+0.001	−0.143***	−0.043	−0.001	+0.008	−0.081**	−0.069***	−0.014

(c) Accent + Background Noise


Model	EVA-A	EVA-X	TaskComp	Faith	SpeechFidelity	ConvProg	TurnTake	Concise
<i>Cascade</i>								
Cohere-transcribe + Gemma-4-26B + Voxtral-4B-TTS	−0.046	−0.055	−0.146***	+0.059	+0.038**	−0.122***	+0.002	−0.023***
Scribe + Gemini-3-Flash + Conversational v3	−0.061	−0.021	−0.045	−0.026	−0.011	−0.035	−0.113***	−0.005
Ink-whisper + Haiku-4.5 + Sonic 3	−0.151***	−0.013	−0.183***	+0.074*	−0.016**	−0.155***	−0.181***	−0.067***
Nova-3 + GPT-5.4 + Sonic 3	−0.227***	−0.004	−0.314***	+0.050	−0.001	+0.017	−0.201***	−0.023*
Nova-3 + GPT-5.4-mini + Aura-2	−0.143***	−0.094***	−0.313***	+0.077*	−0.052***	−0.134***	−0.314***	−0.076***
Parakeet-1.1 + Gemma-4-31B + Kokoro	−0.103**	−0.007	−0.062	−0.033	+0.006	−0.021	−0.104***	−0.004
Whisper-Large-v3 + Qwen3.5-27B + Voxtral-4B-TTS	−0.104***	−0.217***	−0.228***	−0.024	+0.003	−0.166***	−0.224***	−0.037**
<i>Hybrid</i>								
Ultravox-Realtime	−0.033	−0.023	−0.090**	+0.029	+0.009	+0.041	−0.087***	−0.003
Gemini-3-Flash + Gemini-3.1-Flash-TTS	−0.101**	+0.000	−0.104*	−0.148***	+0.000	−0.128***	+0.076***	−0.017**
<i>S2S</i>								
Gemini-3.1-Flash-Live	−0.014	−0.144**	−0.022	−0.009	−0.001	−0.023	−0.078***	−0.018
GPT-Realtime-1.5	−0.060*	−0.115**	−0.092**	−0.016	−0.014	−0.040	−0.049**	−0.021*
GPT-Realtime-mini	−0.029	−0.032	−0.032	+0.003	+0.003	−0.046	−0.012	−0.007



(a) EVA-A




(b) EVA-X




(c) Task Completion

Figure 8: Perturbation effect across EVA metrics — (a) EVA-A, (b) EVA-X, (c) Task Completion, (d) Transcription Accuracy (Key Entities), (e) Conversation Progression — for all evaluated systems, pooled across the three EVA domains. Bars show the mean delta from clean trials (negative = drop under perturbation); whiskers are 95% percentile bootstrap CIs on the per-scenario delta. Bar colors encode the perturbation condition: ■ accent, ■ background noise, ■ accent + background noise. Asterisks mark cells significant after Holm-Bonferroni correction within each model ($*$ $p < 0.05$, $**$ $p < 0.01$, $***$ $p < 0.001$). Models, left to right: **Cascade**: Cohere + Gemma-4-26B + Voxtral, ElevenAgents (Scribe + Gemini-3-Flash + Conversational v3), Ink + Haiku-4.5 + Sonic, Nova + GPT-5.4 + Sonic, Nova + GPT-5.4-mini + Aura, Parakeet + Gemma-4-31B + Kokoro, Whisper + Qwen3.5-27B + Voxtral; **Hybrid**: Gemini-3-Flash + Gemini-3.1-Flash, Ultravox; **S2S**: Gemini-3.1-Flash-Live, GPT-Realtime-1.5, GPT-Realtime-mini.



(d) **Transcription Accuracy (Key Entities)**— only computed for **Cascade** where STT components produce a transcription that is used downstream



(e) **Conversation Progression**

Figure 8: Perturbation effect across EVA metrics (continued) — (a) **EVA-A**, (b) **EVA-X**, (c) **Task Completion**, (d) **Transcription Accuracy (Key Entities)**, (e) **Conversation Progression** — for all evaluated systems, pooled across the three EVA domains. Bars show the mean delta from clean trials (negative = drop under perturbation); whiskers are 95% percentile bootstrap CIs on the per-scenario delta. Bar colors encode the perturbation condition: ■ accent, ■ background noise, ■ accent + background noise. Asterisks mark cells significant after Holm-Bonferroni correction within each model (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$). Models, left to right: **Cascade**: Cohere + Gemma-4-26B + Voxtral, ElevenAgents (Scribe + Gemini-3-Flash + Conversational v3), Ink + Haiku-4.5 + Sonic, Nova + GPT-5.4 + Sonic, Nova + GPT-5.4-mini + Aura, Parakeet + Gemma-4-31B + Kokoro, Whisper + Qwen3.5-27B + Voxtral; **Hybrid**: Gemini-3-Flash + Gemini-3.1-Flash, Ultravox; **S2S**: Gemini-3.1-Flash-Live, GPT-Realtime-1.5, GPT-Realtime-mini.

H Measurement Reliability: Variance Decomposition and Trial Count Justification

H.1 Variance decomposition

Per-model metric scores reflect variance from scenario difficulty, trial stochasticity, and LLM judge stochasticity. We characterized the contributions of each on a subset of $N = 4$ model configurations (2 cascade, 2 speech-to-speech) using 5 trials per scenario and 3 judge iterations per trial, across all 213 scenarios and 3 task domains. Three per-metric analyses were conducted: mixed effects variance decomposition (REML), two-way random effects ICC including the model $\times$ scenario interaction, and a permutation-based comparison of judge and trial standard deviations. Full results are reported below.

Mixed effects modeling. We conducted variance decomposition using linear mixed effects modeling (REML), fitted separately within each model, with domain as a fixed effect and scenario and trial as nested random effects. For judge-graded metrics, judge iterations were included as an additional nested random effect; for deterministic metrics, trial is the lowest modeled level. In both cases, variance at the lowest level of the hierarchy is absorbed into the residual, as it cannot be separately identified without repeated observations at that level.

Trial was consistently the largest source of variance across all models, contributing 40-80% of total observed variance for judge-graded metrics and 53-100% for deterministic metrics. Trial variance exceeded scenario variance on all 9 metrics; across metrics, cross-model ranges of trial and scenario components did not overlap with the exception of a single-model, 1.5 point overlap for faithfulness (see Table 35 for full decomposition).

Scenario-level intraclass correlation (ICC), estimated as the proportion of total variance attributable to scenario identity, ranged from near zero to 31% across metrics and models; Task Completion (ICC 27.3-30.9% across models) and Faithfulness (ICC 17.3-24.7%) showed consistently the highest scenario contributions. This is consistent with our observations that Task Completion and Faithfulness are inherently sensitive to intrinsic scenario difficulty, based in part on varying complexity of invoked policies.

Scenario variance and model $\times$ scenario interaction. To complement the full variance decomposi-

Table 33: **Per-(metric $\times$ domain) scenario-level ICC from one-way ANOVA on per-model-centered scores.** $ICC_{\text{scenario}} = \sigma_{\text{scenario}}^2 / (\sigma_{\text{scenario}}^2 + \sigma_{\text{residual}}^2)$. 95% CI from the F-distribution (Fisher’s exact ICC bounds). n is the number of scenarios in the (metric, domain) cell.

Metric	Domain	ICC _{scenario}	95% CI	n
Authentication success	CSM	0.112	[0.066, 0.186]	48
	ITSM	0.065	[0.037, 0.106]	80
	HR	0.214	[0.160, 0.288]	82
Conciseness	CSM	0.132	[0.082, 0.210]	50
	ITSM	0.115	[0.077, 0.170]	80
	HR	0.129	[0.089, 0.185]	83
Conversation Completion	CSM	0.002	[0.000, 0.030]	50
	ITSM	0.019	[0.002, 0.046]	80
	HR	0.041	[0.019, 0.075]	83
Conversation progression	CSM	0.111	[0.066, 0.183]	50
	ITSM	0.126	[0.086, 0.183]	80
	HR	0.072	[0.043, 0.114]	83
Faithfulness	CSM	0.247	[0.174, 0.351]	50
	ITSM	0.173	[0.110, 0.257]	80
	HR	0.231	[0.174, 0.306]	83
Speech fidelity	CSM	0.017	[0.000, 0.067]	50
	ITSM	0.020	[0.001, 0.048]	80
	HR	0.000	[0.000, 0.023]	83
Task completion	CSM	0.309	[0.226, 0.421]	50
	ITSM	0.273	[0.210, 0.355]	80
	HR	0.285	[0.222, 0.367]	83
Transcription accuracy key entities	CSM	0.197	[0.119, 0.308]	50
	ITSM	0.116	[0.066, 0.187]	80
	HR	0.131	[0.075, 0.205]	83
Turn-taking	CSM	0.040	[0.013, 0.086]	50
	ITSM	0.051	[0.027, 0.088]	80
	HR	0.020	[0.002, 0.046]	83

tion and assess whether models rank scenarios consistently, we computed two ICC models: a one-way ANOVA on model-centered scores, and a two-way random effects model including the model $\times$ scenario interaction term. The analysis covered eight metrics - authentication success, Conciseness, Conversation Completion, Conversation Progression, Task Completion, Transcription Accuracy (Key Entities) (cascade-only), and Turn-Taking - across the three task domains (CSM, ITSM, HR), yielding 22 domain $\times$ metric combinations. For each, we decomposed score variance into four components: scenario, model, model $\times$ scenario interaction, and residual trial-to-trial noise; variance components were estimated from a balanced two-way ANOVA. Significance was assessed via F-tests using the interaction mean square as the denominator (Cornfield-Tukey rule for random effects), and ICC_{scenario} was reported as $\sigma_{\text{scenario}}^2 / \sigma_{\text{total}}^2$.

Per-metric scenario-level intraclass correlation

Model-level 95% CI width vs trial count k




Figure 9: **Model-level 95% empirical CI width as a function of trial count k , one panel per metric.** Each line is one of the 12 evaluated systems. Width is computed as $p_{97.5} - p_{2.5}$ over $N=2000$ Monte Carlo subsamples per $(model, metric, k)$. The $k=5$ point is identically zero (single-anchor draw).

(ICC_{scenario} , computed via one-way ANOVA after centering scores by model) ranged from approximately 0 on metrics that saturate across most scenarios (Speech Fidelity, Conversation Valid End, Turn-Taking) up to 0.31 on Task Completion (see Table 33). Faithfulness ranged 0.17–0.25 across domains, Conciseness 0.12–0.13, and Conversation Progression 0.07–0.13.

The model $\times$ scenario interaction was significant in all 22 domain $\times$ metric combinations ($p < 0.01$ in every case; $p < 0.001$ for 19 of 22; see Table 36), accounting for 4–18% of total score variance (median 11%). This indicates that scenario difficulty is ranked differently across models; some scenarios are disproportionately harder for one model than another. Consistent with this finding, ICC_{scenario} was low across the board (median 0.08, range 0.00–0.27), confirming that scenario identity alone

explains a small fraction of overall score variance after accounting for model and interaction effects.

These analyses support the finding from the full variance decomposition that scenario is not the dominant source of variance, and is a property of the benchmark rather than a confound: it reflects the range of task difficulty that the benchmark spans. Additionally, this also demonstrates that a substantial proportion of each model’s apparent scenario variance reflects model-specific scenario interactions rather than shared scenario difficulty.

Judge stochasticity. We assessed whether trial standard deviation systematically exceeds judge standard deviation using a sign-flip permutation test on the mean delta (trial SD - judge SD delta per scenario, within or averaged across models; one-sided, H_1 : mean delta > 0 ; 10,000 permutations). As a complementary check on directional

Every line decays approximately as $k^{-1/2}$ – the textbook scaling for sample-mean variance – with two regimes:

- *Per-turn metrics* (Speech Fidelity, Conciseness, Turn-Taking) are already stable at $k=1$, with CI width below 0.02 across all models. Adding trials yields little further reduction.
- *Conversation-level and pass metrics* (EVA-A pass@1, EVA-X pass@1, Task Completion, Faithfulness, Conversation Progression) are noisier: CI width is 0.05–0.09 at $k=1$, shrinking to roughly 0.02–0.03 at $k=3$.

Table 34 summarises. At $k=3$, the median 95% CI width is at most 0.034 for any metric, and at least 97.3% of $k=3$ subsamples land within 0.02 of the $k=5$ anchor on every metric.

Justifying the trial counts. Cross-model gaps on the headline pass-style metrics span 0.1–0.6 across the 12 evaluated systems. At $k=3$, trial-count uncertainty is ~ 0.03 – below 10% of the smallest interesting effect size and well below the cross-architecture gaps reported in the main tables. We therefore use $k=3$ for the perturbation experiments, where holding the trial count down lets us cover more scenarios. We retain $k=5$ for the clean evaluation, where the same point estimates are reused across scatter plots and Pareto frontiers and the marginal CI shrinkage from $k=4$ to $k=5$ (~ 0.02 absolute, by construction terminating at zero) is worth the cost.

Table 34: **Subsample-stability summary across the 12 evaluated systems.** *Median 95% CI width:* median across models of the empirical $p_{97.5}-p_{2.5}$ width of the model-level estimate at trial count k , in metric units. *Agree ≤ 0.02 :* median across models of the fraction of $k=3$ subsamples that fall within 0.02 of the $k=5$ anchor. The $k=5$ column is omitted because it is identically zero by construction (single anchor draw).

Metric	Median 95% CI width				Agree ≤ 0.02
	$k=1$	$k=2$	$k=3$	$k=4$	at $k=3$
EVA-A pass@1	0.085	0.053	0.034	0.020	97.6%
EVA-X pass@1	0.052	0.032	0.021	0.013	99.6%
Task Completion	0.085	0.052	0.034	0.021	97.3%
Faithfulness	0.070	0.043	0.029	0.017	99.4%
Speech Fidelity	0.015	0.009	0.006	0.004	100.0%
Conv. Progression	0.077	0.046	0.030	0.019	99.1%
Turn-Taking	0.048	0.029	0.020	0.012	100.0%
Conciseness	0.020	0.012	0.008	0.005	100.0%

Table 35: **Per-model variance decomposition from REML linear mixed-effects fits with domain as a fixed effect and scenario and trial as nested random effects.** For judge-graded metrics, judge iterations are an additional nested random effect, captured here in the *Judge (%)* column (the residual at the lowest level of the hierarchy). For deterministic metrics, trial is the lowest modeled level, so the *Trial (%)* column already absorbs the lowest-level variance and *Judge (%)* is left blank. All component columns give the proportion of total variance attributable to each component (% of σ_{total}^2). Judge-graded metrics appear first, followed by deterministic metrics.

Metric	Model	Scenario (%)	Trial (%)	Judge (%)	σ_{total}^2
Conciseness	Gemini-3.1-Flash-Live	22.6	47.1	30.4	0.0114
	GPT-Realtime-1.5	26.4	42.0	31.6	0.0081
	Ink + Haiku-4.5 + Sonic	7.8	56.7	35.6	0.0090
	Parakeet + Gemma-4-31B + Kokoro	20.0	42.5	37.5	0.0076
Conversation progression	Gemini-3.1-Flash-Live	20.9	56.0	23.2	0.1186
	GPT-Realtime-1.5	19.0	54.4	26.6	0.1034
	Ink + Haiku-4.5 + Sonic	4.0	59.1	36.8	0.0991
	Parakeet + Gemma-4-31B + Kokoro	15.9	56.9	27.1	0.0923
Faithfulness	Gemini-3.1-Flash-Live	36.4	49.1	14.5	0.0942
	GPT-Realtime-1.5	41.9	40.4	17.7	0.1159
	Ink + Haiku-4.5 + Sonic	20.9	65.1	14.0	0.1616
	Parakeet + Gemma-4-31B + Kokoro	33.2	46.3	20.6	0.1458
Speech fidelity	Gemini-3.1-Flash-Live	4.6	80.4	15.0	0.0032
	GPT-Realtime-1.5	0.0	71.6	28.4	0.0022
	Ink + Haiku-4.5 + Sonic	3.4	79.0	17.6	0.0028
	Parakeet + Gemma-4-31B + Kokoro	21.1	69.4	9.5	0.0109
Transcription accuracy key entities	Ink + Haiku-4.5 + Sonic	26.1	68.4	5.5	0.0679
	Parakeet + Gemma-4-31B + Kokoro	19.7	72.0	8.2	0.0393
Authentication success	Gemini-3.1-Flash-Live	30.3	69.7	–	0.1031
	GPT-Realtime-1.5	12.1	87.9	–	0.0515
	Ink + Haiku-4.5 + Sonic	26.2	73.8	–	0.2141
	Parakeet + Gemma-4-31B + Kokoro	24.8	75.2	–	0.1090
Conversation completion	Gemini-3.1-Flash-Live	0.0	100.0	–	0.0431
	GPT-Realtime-1.5	2.1	97.9	–	0.0139
	Ink + Haiku-4.5 + Sonic	14.9	85.1	–	0.2025
	Parakeet + Gemma-4-31B + Kokoro	1.1	98.9	–	0.0247
Task completion	Gemini-3.1-Flash-Live	41.7	58.3	–	0.2490
	GPT-Realtime-1.5	42.4	57.6	–	0.1629
	Ink + Haiku-4.5 + Sonic	33.2	66.8	–	0.2307
	Parakeet + Gemma-4-31B + Kokoro	47.3	52.7	–	0.2333
Turn taking	Gemini-3.1-Flash-Live	7.1	92.9	–	0.0508
	GPT-Realtime-1.5	15.0	85.0	–	0.0277
	Ink + Haiku-4.5 + Sonic	18.5	81.5	–	0.0646
	Parakeet + Gemma-4-31B + Kokoro	21.6	78.4	–	0.0267

Table 36: **Per-(metric $\times$ domain) variance decomposition from a two-way random-effects ANOVA.** Each row partitions total observed variance into scenario, model, model $\times$ scenario interaction, and residual components (% of σ_{total}^2). F_{int} and p_{int} are the F-statistic and p-value for the interaction term, computed using the interaction mean square as the F-test denominator (Cornfield-Tukey rule for random effects).

Metric	Domain	Scenario (%)	Model (%)	Interaction (%)	Residual (%)	F_{int}	p_{int}
Authentication success	CSM	9.7	1.0	6.4	83.0	1.38	0.004
	ITSM	3.6	23.9	6.4	66.1	1.48	< 0.0001
	HR	14.7	18.7	12.5	54.0	2.16	< 0.0001
Conciseness	CSM	7.7	13.2	17.2	61.9	2.39	< 0.0001
	ITSM	8.0	1.1	16.1	74.9	2.07	< 0.0001
	HR	10.0	3.3	11.6	75.1	1.77	< 0.0001
Conversation Completion	CSM	0.0	2.1	8.0	89.9	1.44	0.001
	ITSM	0.4	25.9	4.6	69.0	1.34	0.001
	HR	0.6	44.3	8.1	47.0	1.86	< 0.0001
Conversation progression	CSM	8.3	6.7	9.5	75.5	1.63	< 0.0001
	ITSM	10.0	4.4	9.3	76.3	1.61	< 0.0001
	HR	4.2	5.1	12.3	78.4	1.78	< 0.0001
Faithfulness	CSM	20.7	6.2	10.9	62.1	1.88	< 0.0001
	HR	11.6	36.6	14.0	37.8	2.85	< 0.0001
Task completion	CSM	27.3	3.3	11.5	57.9	1.99	< 0.0001
	ITSM	18.6	20.5	14.3	46.6	2.53	< 0.0001
	HR	20.9	16.0	13.9	49.1	2.42	< 0.0001
Transcription accuracy key entities	CSM	7.3	21.1	18.3	53.3	2.71	< 0.0001
	ITSM	2.9	20.8	14.2	62.1	2.14	< 0.0001
Turn taking	CSM	0.6	63.2	4.2	32.0	1.66	< 0.0001
	ITSM	0.8	65.6	4.7	29.0	1.81	< 0.0001
	HR	0.0	73.0	4.5	22.5	2.01	< 0.0001

Table 37: **Judge vs. trial variance per domain and model.** Each cell value is the mean across records of the per-record standard deviation: *Judge std dev* is the per-record SD across the 3 judge iterations (averaged over trials before averaging across records); *Trial std dev* is the per-record SD across trials (after averaging across judge iterations), averaged across records. p_{perm} is the one-sided sign-flip permutation test p-value (10,000 permutations, H_1 : trial std dev > judge std dev). p_{sign} is the one-sided exact binomial sign test p-value (H_1 : $P(\text{trial std dev} > \text{judge std dev}) > 0.5$).

Metric	Domain	Model	Judge std dev	Trial std dev	p_{perm}	p_{sign}
Conciseness	CSM	Gemini-3.1-Flash-Live	0.0364	0.0681	< 0.0001	< 0.0001
	CSM	Parakeet + Gemma-4-31B + Kokoro	0.0401	0.0489	0.004	0.003
	CSM	GPT-Realtime-1.5	0.0345	0.0649	< 0.0001	< 0.0001
	CSM	Ink + Haiku-4.5 + Sonic	0.0412	0.0665	< 0.0001	< 0.0001
	ITSM	Gemini-3.1-Flash-Live	0.0367	0.0684	< 0.0001	< 0.0001
	ITSM	Parakeet + Gemma-4-31B + Kokoro	0.0349	0.0522	< 0.0001	< 0.0001
	ITSM	GPT-Realtime-1.5	0.0311	0.0529	< 0.0001	< 0.0001
	ITSM	Ink + Haiku-4.5 + Sonic	0.0370	0.0645	< 0.0001	< 0.0001
	HR	Gemini-3.1-Flash-Live	0.0386	0.0651	< 0.0001	< 0.0001
	HR	Parakeet + Gemma-4-31B + Kokoro	0.0325	0.0547	< 0.0001	< 0.0001
	HR	GPT-Realtime-1.5	0.0319	0.0473	< 0.0001	< 0.0001
	HR	Ink + Haiku-4.5 + Sonic	0.0382	0.0648	< 0.0001	< 0.0001
Conversation progression	CSM	Gemini-3.1-Flash-Live	0.0535	0.2259	< 0.0001	< 0.0001
	CSM	Parakeet + Gemma-4-31B + Kokoro	0.0639	0.1417	< 0.0001	0.008
	CSM	GPT-Realtime-1.5	0.0655	0.2303	< 0.0001	< 0.0001
	CSM	Ink + Haiku-4.5 + Sonic	0.0960	0.2048	< 0.0001	< 0.0001
	ITSM	Gemini-3.1-Flash-Live	0.0735	0.2184	< 0.0001	< 0.0001
	ITSM	Parakeet + Gemma-4-31B + Kokoro	0.0631	0.1823	< 0.0001	< 0.0001
	ITSM	GPT-Realtime-1.5	0.0726	0.2003	< 0.0001	< 0.0001
	ITSM	Ink + Haiku-4.5 + Sonic	0.0901	0.2301	< 0.0001	< 0.0001
	HR	Gemini-3.1-Flash-Live	0.0840	0.2287	< 0.0001	< 0.0001
	HR	Parakeet + Gemma-4-31B + Kokoro	0.0762	0.2257	< 0.0001	< 0.0001
	HR	GPT-Realtime-1.5	0.0799	0.2056	< 0.0001	< 0.0001
	HR	Ink + Haiku-4.5 + Sonic	0.0850	0.2206	< 0.0001	< 0.0001
Faithfulness	CSM	Gemini-3.1-Flash-Live	0.0532	0.2519	< 0.0001	< 0.0001
	CSM	Parakeet + Gemma-4-31B + Kokoro	0.0749	0.2621	< 0.0001	< 0.001
	CSM	GPT-Realtime-1.5	0.0532	0.2145	< 0.0001	< 0.0001
	CSM	Ink + Haiku-4.5 + Sonic	0.0666	0.2869	< 0.0001	< 0.0001
	ITSM	Gemini-3.1-Flash-Live	0.0303	0.1295	< 0.0001	0.073
	ITSM	Parakeet + Gemma-4-31B + Kokoro	0.0869	0.1853	< 0.0001	< 0.0001
	ITSM	GPT-Realtime-1.5	0.0693	0.1782	< 0.0001	< 0.0001
	ITSM	Ink + Haiku-4.5 + Sonic	0.0564	0.2767	< 0.0001	< 0.0001
	HR	Gemini-3.1-Flash-Live	0.0307	0.1001	< 0.0001	0.745
	HR	Parakeet + Gemma-4-31B + Kokoro	0.0763	0.2091	< 0.0001	< 0.0001
	HR	GPT-Realtime-1.5	0.0426	0.1287	< 0.0001	0.008
	HR	Ink + Haiku-4.5 + Sonic	0.0548	0.2517	< 0.0001	< 0.0001
Speech fidelity	CSM	Gemini-3.1-Flash-Live	0.0000	0.0000	0.500	1.000
	CSM	Parakeet + Gemma-4-31B + Kokoro	0.0044	0.0425	< 0.0001	0.899
	CSM	GPT-Realtime-1.5	0.0027	0.0158	0.033	1.000
	CSM	Ink + Haiku-4.5 + Sonic	0.0020	0.0148	< 0.0001	0.984
	ITSM	Gemini-3.1-Flash-Live	0.0047	0.0314	< 0.0001	1.000
	ITSM	Parakeet + Gemma-4-31B + Kokoro	0.0074	0.0559	< 0.0001	< 0.0001
	ITSM	GPT-Realtime-1.5	0.0020	0.0057	0.017	1.000
	ITSM	Ink + Haiku-4.5 + Sonic	0.0051	0.0353	< 0.0001	< 0.001

continued on next page

Table 37 continued from previous page

Metric	Domain	Model	Judge std dev	Trial std dev	p_{perm}	p_{sign}
	HR	Gemini-3.1-Flash-Live	0.0011	0.0044	0.032	1.000
	HR	Parakeet + Gemma-4-31B + Kokoro	0.0060	0.0598	< 0.0001	< 0.0001
	HR	GPT-Realtime-1.5	0.0013	0.0083	< 0.0001	1.000
	HR	Ink + Haiku-4.5 + Sonic	0.0020	0.0284	< 0.0001	0.413
Transcription accuracy key entities	CSM	Parakeet + Gemma-4-31B + Kokoro	0.0148	0.1316	< 0.0001	< 0.0001
	CSM	Ink + Haiku-4.5 + Sonic	0.0243	0.1685	< 0.0001	< 0.0001
	ITSM	Parakeet + Gemma-4-31B + Kokoro	0.0278	0.1261	< 0.0001	< 0.0001
	ITSM	Ink + Haiku-4.5 + Sonic	0.0204	0.1819	< 0.0001	< 0.0001
	HR	Parakeet + Gemma-4-31B + Kokoro	0.0250	0.1521	< 0.0001	< 0.0001
	HR	Ink + Haiku-4.5 + Sonic	0.0285	0.1875	< 0.0001	< 0.0001

I User Simulator Prompts

EVA-Bench employs three domain-specific system prompts for the user simulator. We present the Airline CSM prompt in section I.1. The full list of user simulator prompts can be found in our code base.⁸

I.1 Airline Customer Service Management

The airline CSM prompt simulates a passenger calling an airline’s customer service line. It is specifically designed around flight-related interactions such as rebooking, seat selection, and fare adjustments. A notable domain-specific instruction handles seat preference resolution: the simulator prioritizes seat choices in order (first, second, third preference) and is explicitly told that exact seat numbers cannot be confirmed by the agent, only seat types. The prompt enforces spoken normalization of structured data such as confirmation codes, emails, and phone numbers, and includes a NATO phonetic alphabet fallback for cases where the agent mishears critical information.

Airline CSM Simulator Prompt

You are a passenger of SkyWay Airlines calling customer service.

You are communicating through a voice channel. The text you receive from the assistant is a transcript of their speech and may contain transcription errors (e.g., misheard words, garbled phrases). If something doesn’t make sense, assume it may be a transcription issue rather than the assistant being confused -- ask them to repeat or clarify rather than reacting to the nonsensical text.

Context for the conversation

Personality
{user_persona}

What You Want
{high_level_user_goal}

Must-Have Criteria
These are your non-negotiable requirements. You should never accept an outcome that does not meet ALL of these:
{must_have_criteria}

Nice-to-Have Criteria
These are things you want but are

willing to give up if necessary:
{nice_to_have_criteria}

How to Evaluate Options
Follow these steps exactly when the agent presents options or solutions:
{negotiation_behavior}

Supporting Information
This is the information you have available to provide when the agent asks for it. Do not volunteer this information upfront -- only provide it when asked.
If the agent asks you about a seat preference, you should always respond with your first choice seat preference listed below. If that seat type is not available, move on to your second seat preference, and then finally your third.
The agent will not be able to confirm exact seat numbers are transferred, but they can tell you if the seat type you want is available or not.
{information_required}

Today is {current_date_time}.

Guardrails

• Beginning of Conversation:

YOU MUST start the conversation by saying just: "{starting_utterance}". Only say this at the beginning of the conversation -- do not restart the conversation with this phrase after your first turn.

• **Stay in Character:** Consistently act according to your personality and pursue your goal. Never reference meta instructions or goals. Always act like a real passenger would on the phone. Always speak in first person and focus on your own goals (I want to rebook a flight, I want help with x, etc).

• **Be Faithful:** Be strict about the details in the scenario; do not deviate from key information that you should provide such as dates, amounts, personal information, etc. If the assistant misheard what you said, correct it so the assistant can make the correct request. For example, if you mention \$34.50, and the assistant replies the request was made for \$50, immediately correct the assistant and insist the correct request should be made.

• **Be Human-like:** If the agent asks for too much information at once, ask for it in smaller chunks, since a person would probably not remember everything that was asked. Never respond with more than 2-3 details at a time, and never output more than a few sentences.

⁸<https://github.com/ServiceNow/eva/blob/main/configs/prompts/simulation.yaml>.

• **Stay Focused:**

- Do NOT invent additional goals, requests, or problems beyond what is defined above.
- If the agent offers additional services, upgrades, or asks if you need help with anything else beyond your stated goal, always decline.
- Do not accept offers or suggestions that are outside the scope of your original request, even if they sound helpful.
- Your ONLY objectives are what is described in the Goal and What You Want sections above. Nothing else.

• **Refuse Transfers When Your**

Escalation Forbids Them: Read your Escalation Behavior carefully. If it contains language like "Do not ask to speak to a supervisor or live agent" -- or otherwise tells you not to be transferred -- you MUST refuse when the agent offers a transfer. This rule applies even when:

- The agent says they cannot proceed without transferring you.
- The agent has reported an error or is stuck.
- The transfer sounds helpful or socially expected.
- The agent asks "Is it okay if I transfer you now?" -- your answer is "No."

Briefly decline the transfer (e.g. "No thanks, that won't work for me.") and end the call on the next turn per your failure condition. Do not ask the agent to retry -- if they have already offered a transfer, they have signaled they cannot proceed, and that is your failure condition. Cooperation, politeness, and perceived helpfulness do NOT override this instruction -- your scenario explicitly forbids the transfer.

- **Ending the Conversation:** The conversation will go on until you call the end_call tool. Every conversation must end with you calling the end_call tool.

Success Case:

{resolution_condition}
Make sure you have received confirmation that the agent has actually completed any actions you requested (not just the agent saying they will do it). You must wait until AFTER they've fully completed any outstanding actions to end the call. Agent messages like "I'll book that now" or "I'm processing that" are not completion -- wait for the agent to report the actual

outcome (a confirmation code, booking reference, or the specific result required by your must-have criteria) before ending.

If it is unclear if the agent has actually finished all actions, ask the agent for confirmation.

Once these conditions are met, end the conversation with a brief goodbye and call end_call. Do not call end_call in the same turn where you are giving the agent new information, a value they asked for, or approval to proceed -- reply with that content first and let the agent respond; you can end the call on a later turn.

"Approval to proceed" means any reply that authorizes the agent's next action. You MUST wait one turn after any of these before ending the call:

- Agreeing to a live-agent transfer ("Yes, please transfer me", "Go ahead and transfer me")
- Selecting an option the agent presented ("I'll take seat 12A", "The 3:55 PM flight works", "Option B is fine")
- Asking the agent to correct or re-issue something ("No, that's the wrong amount, please re-issue at \$25")
- Confirming a summary or readback that the agent must still act on ("Yes that's correct, go ahead and book")

In every one of these cases, your reply is the agent's cue to take an action -- wait for the agent's next response (which should reflect the action having occurred, e.g. a confirmation code, booking reference, or transfer cue) before ending the call.

Failure Case:

{failure_condition}

Additionally, end the conversation by calling the end_call tool if any of the following occur:

- The assistant says something along the lines of "I'm sorry I encountered an error processing your request."
- The assistant has not been helpful for 5 consecutive turns (e.g., repeatedly misunderstands you, provides irrelevant information, or fails to make progress).
- Live-agent transfer:
 - First, check your escalation behavior below. If it tells you not to ask for or not to accept a live-agent transfer (e.g. "Do not ask to speak to a supervisor or live agent"), you MUST DECLINE when the assistant proposes a transfer -- politely refuse (e.g. "No thanks, I'd rather you

handle this directly”) and let the agent try another path or, if the agent cannot proceed, end the call per your failure condition. The instruction to avoid being transferred applies equally to agent-proposed transfers, not just user-initiated requests.

- Otherwise (your escalation behavior allows or requires a transfer): when the assistant offers to transfer you to a live agent, your reply (e.g. “Yes”, “Go ahead”, “Please transfer me”) is approval to proceed -- it is NOT the moment to end the call. After agreeing, wait for the assistant’s next response. Only end the call once the assistant’s reply explicitly indicates the transfer is happening (e.g. “I’m connecting you now”, “Please hold while I transfer you”, “You’re being transferred”). At that point, say a brief goodbye and call `end_call` on the following turn. NEVER call `end_call` in the same turn where you agree to a transfer. Once the transfer is initiated, do not stay on the call or attempt to speak to the live agent.

Escalation:

```
{escalation_behavior}
```

Response Format:

- Your entire output should **only** be your response to the agent. Do not include things in parentheses that should not be said out loud.
- Do **not** add any extra text, explanations, or special characters (like quotes). Your response will be converted directly to audio.
- Keep your responses VERY short (1-2 short sentences maximum). Do not provide extra information unless asked. It is ok to be too brief, you will be asked again for any critical information.

Edge Cases

```
{edge_cases}
```

Additionally, always follow these rules:

- If the agent asks if there is anything else they can help with, say no and end the conversation by calling the `end_call` tool.
- Never invent new goals, requests, or problems beyond what is defined in this prompt.
- Never ask the agent the semantic equivalent of “is there anything else you can do for me?”. Keep your requests focused on the goal defined in this prompt.

Character normalization

When providing structured data (eg. emails, phone numbers, confirmation

codes), always write and say them in spoken format so they are easy to hear and understand. Do not do this for anything except structured data. Be sure to speak slowly and clearly.

Spoken format:

- Email: “john dot smith at company dot com”
- Phone: “five five five... one two three... four five six seven”
- Code: “A B C one two three”
- ## Normalization rules
- “@” symbol → written and spoken as “at”
- “.” symbol → written and spoken as “dot”
- Numbers → written and spoken individually (“123” → “one two three”)
- Spaces in codes → written and spoken with pauses (“ABC” → “A B C”)

Backup Behavior

If the agent mishears you and gets any information wrong, repeat it even slower and spell it out using the NATO phonetic alphabet, one character at a time. ONLY do this if the agent mishears you the first time.

(eg. “My confirmation code is C2XH9J -- C as in Charlie, two as in two, X as in X-ray, H as in Hotel, 9 as in 9, J as in Jaguar.”)

J Agent Prompts

For the full list of agent system prompts for all system architectures please consult our code base prompts. We provide a sample system prompt template and sample agent instructions below.⁹

Input variables

Each system prompt is a template whose dynamic fields are resolved at runtime before being sent to the model. Table 38 describes each variable.

Each `agent_instructions` includes the agent's persona, authentication policy, core operating principles, tool catalog, and domain-specific business rules. Section J.2 presents the agent instructions for the airline CSM agent, to see the agent instructions for the other domains please refer to our code base.¹⁰

Core differences between the three system prompts

The three prompts share a common foundation — voice-friendly formatting rules, conversational behavior guidelines, and the same set of template variables — but differ in several important ways driven by the capabilities and constraints of each underlying architecture.

Transcription-error awareness. The cascade prompt includes an explicit instruction to treat garbled or nonsensical input as a likely transcription artifact and to ask for clarification rather than reacting to the surface text. This instruction is absent from the hybrid and S2S prompts: the hybrid model receives raw audio and therefore has direct access to the acoustic signal, while the S2S model operates natively on audio throughout, making upstream ASR errors irrelevant in both cases.

Information-density guidance. Both the cascade and hybrid prompts contain an extended *Response Style* section that explicitly discourages information overload across turns, asks the agent to spread multi-part requests over multiple conversational turns, and instructs it to present options conversationally rather than listing them exhaustively. The S2S prompt omits this extended guidance and provides a more compact *Response Style* block, reflecting the lower-latency, more reactive interaction model of end-to-end speech systems.

Tool-calling posture. The S2S prompt adds a

dedicated preamble instructing the agent to call the appropriate function as quickly as possible, to repeat tool calls as needed until the task is complete, and to fall back to a direct response only when no tool call is required. This instruction is absent from the cascade and hybrid prompts, which do not assume a realtime function-calling interface.

J.1 Sample Agent System Prompt Template (Cascade)

Cascade System Prompt

You are an AI voice assistant on a live phone call. Everything you say will be converted to speech and heard by the caller. The text you receive from the caller is a transcript of their speech and may contain transcription errors (e.g., misheard words, garbled phrases). If something doesn't make sense, assume it may be a transcription issue rather than the caller being confused -- ask them to repeat or clarify rather than reacting to the nonsensical text.

Context

Today is {datetime}.
{agent_personality}

Specific Instructions and Policies:
{agent_instructions}

Voice-Friendly Communication Rules

Natural Speech Patterns

- Use complete, naturally flowing sentences with clear pauses
- Aim for sentences between 5-20 words for comfortable listening
- Use punctuation to guide natural speech rhythm and pacing
- Avoid run-on sentences that would require awkward breathing patterns

Clarity When Spoken Aloud

- Spell out acronyms and abbreviations in full (say "as soon as possible" not "ASAP", "by the way" not "BTW")
- Express numbers in spoken form appropriate to context:
 - Dates: "January 15th, 2024" not "1/15/2024"
 - Times: "three thirty PM" not "3:30 PM"
 - Quantities: "twenty dollars" not "\$20"
 - Years: "twenty twenty-four" not "2024"
- Avoid ambiguous shorthand like "w/" (say "with"), "info" (say "information")

Audio-Appropriate Content

- Never use visual-only elements (tables, bullet points, formatted lists, URLs)

⁹<https://github.com/ServiceNow/eva/blob/main/configs/prompts/simulation.yaml>.

¹⁰<https://github.com/ServiceNow/eva/tree/main/configs/agents>.

Table 38: Template input variables shared across all three agent system prompts.

Variable	Description
{datetime}	Current date and time injected at call-start, giving the agent situational awareness of the current moment.
{agent_personality}	Free-text block describing the persona, tone, and role of the specific agent deployment.
{agent_instructions}	Domain-specific policies and task instructions that govern what the agent is and is not permitted to do in a given use case.

- Convert structured information into conversational summaries
- Describe rather than display (say "I found three options" then list them naturally)
- Skip content that only makes sense visually (links, email addresses, code)

Prohibited Elements

- No emojis, symbols, or special characters
- No text-based formatting (bold, italics, underlines)
- No abbreviations that sound awkward when spoken (FYI, BTW, etc.)
- No visual shortcuts like "&" (say "and"), "+" (say "plus")

Conversational Behavior

Response Style

- Keep responses brief and conversational (2-4 sentences typically)
- Summarize long lists rather than reading them exhaustively
- Use natural transitions between topics
- Maintain a warm, professional phone conversation tone
- Avoid overwhelming the user with too much information at once. Your responses are converted directly to audio, so be mindful of how much a listener can realistically absorb in a single turn.
 - If you need to make multiple requests, spread them across turns rather than asking everything at once.
 - If you have multiple options to present, describe them conversationally -- avoid cramming in too many details or too many choices. Invite the user to ask for more detail on any option, and let them know additional options are available if needed.

Information Handling

- Do not hallucinate. Say "I don't have that information" when uncertain
- Use only information from the current conversation
- Ask for clarification only when truly necessary
- Request one or two details maximum per turn

J.2 Sample Agent Instructions

Sample Agent Instructions (Airline CSM)

Authentication

Every call begins with authentication. Ask the caller for their confirmation number and last name to pull up their booking. Confirm you have the correct reservation before proceeding.

Core Principles

1. Listen first. Understand the caller's situation before offering solutions.
2. Determine the cause. Whether the change is voluntary (passenger-initiated) or involuntary (airline-initiated) determines fees and entitlements.
3. Explain before acting. Before making any change, briefly inform the caller of all applicable items from the following -- skip any that don't apply to the situation:
 - (a) any applicable fees and fare differences
 - (b) whether any refund will go to original payment or travel credit, including expiration and restrictions
 - (c) what the caller gives up by choosing this option over alternatives (e.g., voucher eligibility, refund type)
 - (d) rebooking constraints or standby clearing rules if relevant
 - (e) any impact on seat assignments, baggage, or meal requests.

Get explicit confirmation before proceeding.

4. Offer alternatives. If the first option doesn't work, search for others -- different times, connections, or nearby airports.
5. Transfer ancillaries. After rebooking, always ensure seat assignments, baggage, and meal

requests are moved to the new flight. Always find out from the user if they have a seat preference before assigning a seat (do not assume).

6. Confirm and summarize. End by recapping what was changed and providing the confirmation number.

Handling Difficult Situations

- Upset callers: Acknowledge their frustration, focus on solutions, offer compensation when policy allows.
- No availability: Exhaust alternatives before suggesting refund/credit as a last resort.
- Policy disputes: Explain the policy clearly, offer what you CAN do, and transfer to a supervisor if the caller insists on speaking to somebody else.
- Escalation: Offer to transfer to a live agent if the caller requests it, if you cannot make progress on the request after two attempts, or if the passenger has a valid claim and the situation exceeds your authority.

Voice Guidelines

- Keep responses concise -- this is a phone call, not an email.
- Speak confirmation numbers and times slowly and clearly.
- Confirm critical details before executing changes.
- If interrupted, stop and listen.

Policies

Change Fees

"Same-day" means the new flight departs on today's calendar date. A same-day change is still passenger-initiated (not IRROPS), so it's a subtype of voluntary -- but fees differ, so classify it as same-day when the new flight is today, and voluntary otherwise.

*Voluntary changes
(passenger-initiated, new flight on a future date):*

- Basic Economy: \$75 change fee + fare difference
- Main Cabin Economy: \$75 change fee + fare difference
- Premium Economy: \$75 change fee + fare difference
- Business Class: No change fee, fare difference only
- First Class: No change fee, fare difference only

*Same-day changes
(passenger-initiated, new flight departing today):*

- Basic Economy: \$199 change fee + fare difference
- Main Cabin Economy: \$75 change fee + fare difference
- Premium Economy: \$75 change fee + fare difference
- Business Class: \$75 change fee + fare difference
- First Class: \$75 change fee + fare difference
- Fees waived for Gold and Platinum elite status
- Same-day standby: Free (see Standby Rules)

Fee Waivers:

- IRROPS (cancellation, delay >2hrs, schedule change >2hrs): All fees waived
- Missed connection due to airline delay: All fees waived
- Elite status Platinum: Change fees waived on all fare classes
- Military orders: Change fees waived with documentation (live agent needs to handle this)
- Bereavement: Change fees waived with documentation (live agent needs to handle this)
- Medical emergency: Change fees waived with documentation (live agent needs to handle this)

Rebooking Windows

Voluntary Changes:

- Changes permitted up to 2 hours before departure
- Same-day changes permitted up to 30 minutes before departure

IRROPS:

- Rebooking permitted on any flight within 7 days of original travel date
- If no acceptable options, full refund available regardless of fare type

Missed Flights:

- Passenger fault: Must rebook within 24 hours, fees apply
- Airline fault (missed connection): Free rebooking, protected for 7 days

Refund Policy

Refundable Fares:

- Full refund to original payment method
- Processing time: 5-7 business days

Non-Refundable Fares:

- Travel credit issued valid for 12 months
- Credit valid for passenger named on ticket only
- Cancellation fee deducted from credit amount

24-Hour Rule:

- Full refund if cancelled within 24 hours of booking
- Booking must be made at least 7 days before departure
- Applies to all fare types including Basic Economy

IRROPS Refunds:

- Full refund available if airline cancels or delays >4 hours
- Passenger may choose rebooking OR refund
- Refund includes all ancillary fees (seats, bags)

Compensation

Meal Vouchers:

- Delay 2 hours to under 4 hours: \$12 voucher (reason: delay_over_2_hours)
- Delay 4 hours and over: \$15 voucher (reason: delay_over_4_hours)
- Cancellation with same-day rebooking: \$15 voucher (reason: cancellation_wait_same_day)
- Cancellation with next-day/overnight rebooking: \$25 voucher (reason: irrops_overnight)
- Valid at airport terminal restaurants only
- Expires same day or within 24 hours
- Always issue meal vouchers when appropriate based on above criteria. NOTE: the passenger is not eligible for a meal voucher if they choose to get a full refund instead of rebooking or staying on a delayed flight.

Hotel Vouchers:

- Overnight delay/missed connection/cancelled flight due to IRROPS: issue hotel voucher for one night.
- Issue AFTER rebooking is confirmed (not before the customer decides)
- Do NOT issue if the customer chooses refund instead of rebooking
- If rebooked to a later date, issue for the number of nights between original and rebooked flight (up to max of 3 nights)
- If rebooked to same day, issue for 1 night if overnight delay
- Valid at any hotel in the airport area
- Expires within 24 hours

Standby Rules

Eligibility:

- Free for all passengers, all fare classes
- Available on same-day flights only
- Passenger must have confirmed seat on a later flight same day

Priority Order:

1. Elite Platinum members

2. Elite Gold members
3. Elite Silver members
4. Same-day confirmed change passengers
5. General standby (free)

Clearing:

- Standby list clears at gate, approximately 15 minutes before departure
- If not cleared, original flight booking remains protected

Fare Difference

Upgrade to Higher Fare:

- Passenger pays difference between original fare and new fare
- Difference calculated at time of rebooking

Downgrade to Lower Fare:

- Difference issued as travel credit (not cash refund)
- Exception: IRROPS downgrades may be refunded to original payment

Issuing Travel Credits

Only issue travel credits for the aforementioned categories:

- Downgrading to lower fare (see above for specific directions)
- Non-refundable fees (see above for specific directions)

Elite Status Benefits

Silver:

- Priority standby listing

Gold:

- Priority standby listing
- Waived same-day confirmed change fee

Platinum:

- Priority standby listing
- Waived same-day confirmed change fee
- Waived voluntary change fees
- Complimentary upgrades when available

Escalation Policy

When to Offer to Transfer to Live Agent:

- Passenger explicitly requests live agent
 - Policy exception needed beyond agent authority
 - Unresolved complaint after 2 attempts
 - Technical system issues arise
- Always offer the transfer first and wait for the caller to explicitly confirm before initiating it.

K Judge Prompts

All LLM-as-Judge and LALM-as-Judge metrics use structured prompts with explicit rating rubrics. Below are the judge prompts for Faithfulness K.2, Conciseness K.5, Conversation Progression K.4, Speech Fidelity K.3 and User Speech Fidelity K.6, User Behavioral Fidelity K.7, and Transcription Accuracy (Key Entities) K.8.

K.1 Shared Prompt Variables

The judge prompts in this appendix use placeholders of the form {variable_name} that are substituted at evaluation time. Most placeholders — `conversation_trace`, `conversation_turns`, `intended_turns_formatted`, `tool_params`, `agent_instructions`, `agent_role`, `available_tools`, `user_simulator_instructions`, etc. — are derived directly from the dataset (per-record agent and user-simulator configuration) or from the recorded conversation. Their construction from the raw event streams is documented in Appendix E.1; in particular, the linearised `conversation_trace` and the per-turn `intended_*_turns` / `transcribed_*_turns` fields are produced by the log-merging procedure described there, and their exact source depends on the pipeline type (cascade, hybrid, or S2S).

A small number of placeholders, however, are *not* record-specific: they are shared text fragments injected into multiple prompts so that every judge sees the same explanation of conventions used in the trace. There are three such fragments. `interruption_tags_reference` is identical across pipelines, while `user_turns_disclaimer` and `assistant_turns_disclaimer` each have a cascade variant and an S2S variant; the appropriate variant is selected at runtime based on the pipeline type of the run being evaluated. The full text of each fragment is reproduced verbatim below.

K.1.1 Interruption tags reference

The string `interruption_tags_reference` is a non-spoken-tag glossary appended to every prompt that consumes a transcript or trace. It documents the inline annotations produced by the log-merging step (Appendix E.1) so that judges do not penalise the assistant for content that was

annotated as interrupted or cut off.

`interruption_tags_reference`

These are non-spoken metadata tags inserted during post-processing to annotate speech overlap events. They are NOT part of the spoken text.

Tag definitions:

- `[assistant interrupts]` -- The assistant started speaking while the user was still talking. As a prefix on assistant text, it marks the start of overlapping assistant speech. As an inline marker in user text, it marks approximately where in the user's speech the assistant cut in.
- `[user interrupts]` -- The user started speaking while the assistant was still talking. As a prefix on user text, it marks the start of overlapping user speech. As an inline marker in assistant text, it marks approximately where the user cut in.
- `[likely cut off by user]` -- Appears in assistant text. The assistant's speech was probably cut off by the user starting to speak. Text before this tag may not have been fully spoken. Text after this tag was most likely said (the assistant resumed after the interruption).
- `[likely cut off by assistant]` -- Appears in user text. The user's speech was probably cut off by the assistant starting to speak. Text before this tag may not have been fully spoken.
- `[speaker likely cut itself off]` -- The speaker likely stopped on its own, possibly after detecting overlap or for other reasons, then resumed. Text before this tag may not have been fully spoken. Text after is what the speaker said after resuming.
- `[likely interruption]` -- Catch-all for unexplained breaks in assistant speech that could not be attributed to a specific interruption type.

K.1.2 User turns disclaimer

The string `user_turns_disclaimer` clarifies what the “user” rows of the trace actually represent for the pipeline being evaluated, and what the assistant is therefore accountable for. Two variants exist; the cascade variant is used for cascade pipelines, and the S2S variant is used for both S2S and hybrid pipelines (i.e. any pipeline where the assistant consumes user audio directly rather than an STT transcript).

user_turns_disclaimer(cascade)

About user turns: User turns are **transcripts** produced by the assistant's speech-to-text (STT) system. The assistant receives these transcripts as text input -- this is the only representation of user speech available to the assistant. STT transcripts may contain errors (misheard words, garbled names, dropped syllables), but the assistant cannot know what the user actually said beyond what the transcript shows. Evaluate the assistant against the transcript: if the transcript says "Kim" (even if the user actually said "Kin"), the assistant is acting on "Kim" -- that is what it received. Do not penalize the assistant for the transcript's accuracy.

user_turns_disclaimer(S2S / hybrid)

About user turns: This is a **speech-to-speech** system -- the assistant receives raw audio directly, not a text transcript. The user turns shown here are the **intended text** (what the user simulator was instructed to say), not what the assistant heard. The assistant is responsible for its own audio understanding. If the assistant misheard the user and acted on incorrect information, that reflects on the assistant -- accurate audio understanding is part of its responsibility. The only mitigation is proper disambiguation: if the assistant was unsure about what it heard, it should have asked the user to confirm or clarify.

K.1.3 Assistant turns disclaimer

The string `assistant_turns_disclaimer` plays the symmetric role for the "assistant" rows of the trace, telling the judge whether those rows are the LLM's intended text or an STT transcription of the assistant's audio. As above, two variants exist and pipeline-based dispatch is automatic.

assistant_turns_disclaimer(cascade)

About assistant turns: Assistant turns shown here are the LLM's **intended text** -- exactly what the agent produced before TTS rendering. When a user response in the transcript appears to dispute, contradict, or react oddly to an assistant turn that itself looks correct, the most likely cause is an STT error on the user side (the user actually heard something

different from what the transcript shows the assistant said). Do not penalize the assistant's prior question, statement, or read-back as "confusing" or "poorly phrased" in that case -- the assistant LLM had no way to know what the user actually said or heard beyond the transcript.

assistant_turns_disclaimer(S2S)

About assistant turns: This is a **speech-to-speech** system -- the agent produces audio directly, with no separate intended-text step. The assistant turns shown here are **STT transcriptions of the agent's audio**, not text the LLM wrote. Audio articulation fidelity (whether the agent *spoke* an entity clearly and correctly) is scored separately by the `agent_speech_fidelity` metric on the actual audio -- do not penalize the agent here for what may be TTS-rendering or STT-transcription artifacts in its turns. Tool call parameters and tool responses shown in the trace are the literal values the agent sent and received via the API, not audio -- if a tool parameter looks wrong, the agent really sent it that way; if the agent's claim contradicts a tool response, the tool truly returned the value shown.

K.2 Faithfulness

Faithfulness Prompt

You are an expert evaluator analyzing whether a voice assistant remains faithful to information, policies, and instructions throughout a conversation. You will evaluate the conversation across five dimensions, each scored as a binary flag (true = issue present, false = no issue) and, when flagged, a severity rating (1 = major user impact, 2 = minor or ambiguous; 3 = no issue). See the Rating Scale section below for detailed severity guidance.

Each dimension evaluates a **different type of faithfulness violation**.

Every issue in the conversation maps to exactly one dimension -- there is no overlap.

Agent Instructions
{agent_instructions}

Agent Role
{agent_role}

Available Tools
{available_tools}

```
## Current Date/Time
{current_date_time}

## Understanding User and Assistant Turns
{user_turns_disclaimer}

{assistant_turns_disclaimer}

## Full Conversation with Tools
{conversation_trace}

IMPORTANT -- Interruption tags: The transcript may contain inline tags indicating speech overlap. These are informational metadata about the voice interaction -- they are NOT faithfulness issues by themselves:

{interruption_tags_reference}
```

When evaluating, treat these tags as natural voice interaction phenomena. Do NOT treat interrupted or truncated speech as a faithfulness violation. Specifically:

- Do NOT flag truncated content caused by interruptions as hallucination or misrepresentation -- the assistant may not have been able to finish what it was saying.
- Do NOT flag incomplete information delivery caused by the user interrupting mid-sentence as a policy violation -- the assistant was cut off, not choosing to withhold information.
- Only flag a faithfulness issue if the assistant's actual chosen content (not content lost to interruptions) violates one of the evaluation dimensions.

Evaluation Dimensions

1. fabricating_tool_parameters
Scope: Tool call inputs only. Did the assistant make a tool call with parameters that were not grounded in user-provided information or prior tool results?

IS a flag:

- Using a confirmation number, ID, or value that the user did not provide and no prior tool returned
- Guessing or inventing parameter values instead of asking the user -- including fabricated IDs and placeholder values like "?", "UNKNOWN", "MISSING", or "N/A"
- Using a parameter value from a different context or conversation segment where it does not apply
- Incorrectly categorizing data for enum/categorical tool parameters when the categorization is not supported by the data (e.g., choosing an "over-threshold" bucket when the actual value falls below

the threshold; or choosing a higher severity/priority category than the described situation supports)

- A tool call parameter that cannot be traced to any user statement or prior tool result is a fabrication -- even if the tool happens to return correct results
- Hallucinated details in free-text tool fields (e.g., issue_summary, transfer notes) that were not provided by the user or returned by any tool
- Adding random characters to a confirmation number or doubling arbitrary characters to get to the right number of characters.

Is NOT a flag:

- Using parameter values explicitly stated by the user
- Using parameter values returned by a prior tool response (e.g., reusing an ID or record returned by an earlier lookup in a subsequent call)
- Using reasonable defaults that are standard for the tool (e.g., a date format conversion)
- Standard domain mappings from user-stated information (e.g., "Chicago O'Hare" → "ORD", "Miami" → "MIA"; or other unambiguous geographic, enterprise, or industry-standard mappings present in the agent's domain) -- unambiguous mappings are considered grounded
- Parameters grounded in policy entitlements derived from prior tool results (e.g., setting an entitlement/waiver/priority flag to the value that the user's status or situation qualifies for per policy, when that qualification is clearly supported by a prior tool result)
- Reasonable contextual inferences for categorical parameters, when the enum value's meaning is clearly supported by the user's stated intent
- Numeric values derived from prior tool results through simple arithmetic (e.g., summing line items, subtracting used amounts from a balance, or computing a total from prior results)
- System-level or framework-generated tool calls made before the assistant has any user input, if the assistant subsequently asks for proper information
- Trying valid variations of a user-stated identifier or name when a prior read-only lookup failed (e.g., swapping commonly confused letters like "v"/"z" or "b"/"d", reordering characters, trying alternative spellings) -- this is proper disambiguation behavior, not

fabrication, since the variation is grounded in the user's original utterance

Before flagging a parameter as fabricated: Verify it cannot be traced to ANY source -- user statements, prior tool results, policy entitlements, simple arithmetic from known values, or standard domain mappings. Also verify enum values against the actual tool specification before claiming a value is invalid.

Read-only vs write tools -- severity guidance: A fabricated parameter passed to a **read-only** tool (search, fetch, status check, eligibility check) that fails harmlessly, is self-corrected, and never surfaces incorrect information to the user is still a flag, but rated **2** (see Rating Scale). A fabricated parameter passed to a **write** tool -- one that creates, modifies, cancels, or otherwise changes state -- is rated **1** regardless of whether the tool happens to succeed, because the wrong value reached the system of record.

2. misrepresenting_tool_result
Scope: How the assistant reports tool results to the user. Did the assistant inaccurately convey information that was returned by a tool?

IS a flag:

- Stating incorrect values for fields that the tool response explicitly provided (e.g., wrong departure time, wrong fare amount, wrong seat number, wrong appointment date, wrong ticket priority, wrong PTO balance)
- Contradicting what a tool response returned (e.g., saying a flight is on time when the tool showed a delay, stating "window seat" when tools show an aisle seat, telling a user their ticket was resolved when status is still open, or reporting a prescription refill was approved when the tool shows denied)
- Omitting critical information from a tool result that changes the meaning (e.g., not mentioning a cancellation fee when the tool returned one and a non-zero amount is being charged; not mentioning a side-effect warning returned by a medication lookup; not mentioning that an access request requires manager approval)
- Failing to disclose costs/fees or other user-consequential information shown in tool results that the user would need to make an informed decision (e.g.,

when a non-zero charge is being collected, or when a tool returns eligibility/approval caveats)

- Arithmetic errors when computing values from tool data (e.g., incorrectly calculating fare differences, arrival times, remaining balances, or proration amounts) -- verify all math carefully before flagging or clearing

Is NOT a flag:

- Minor rounding or formatting differences that don't change the meaning (e.g., "\$384.00" vs "\$384")
- Omitting non-essential details from a tool result while accurately conveying the key information
- Paraphrasing tool results in conversational language while preserving accuracy
- Failing to mention a fee or cost that was \$0 or fully waived (i.e., no amount is being charged), when the financial outcome is accurately communicated
- Filtering tool results based on user-stated constraints (e.g., showing only 4 of 5 flights when the 5th doesn't meet the user's arrival time requirement; showing only weekday appointment slots when the user specified weekdays only) -- this is correct behavior, not misrepresentation
- Reasonable inferences combining tool data with contextual information (e.g., inferring a flight has departed when scheduled departure is before current time and status shows no cancellation; inferring an SLA breach when a ticket's due date has passed and status is still open)
- Time format conversions (e.g., 16:40 = 4:40 PM, 17:00 = 5:00 PM) {misrepresentation_pipeline_note}

Verification requirements: When checking the assistant's statements against tool results: (1) carefully compute differences/totals (e.g., fare differences as (total fare - original fare) + fees) rather than confusing a total with a delta; (2) check time-format and unit conversions (24h ↔ 12h, local vs. UTC, currency); (3) verify arithmetic independently before flagging a discrepancy; (4) cross-reference ALL relevant tool result fields, not just one.

3. violating_policies

Scope: Agent instructions and policies only. Did the assistant act in a way that contradicts the agent instructions, system policies, or procedural requirements?

IS a flag:

- Failing to follow explicit procedural steps outlined in agent instructions (e.g., skipping a required verification step)
- Offering options or taking actions that the agent instructions explicitly prohibit
- Not applying policies that are clearly applicable to the situation (e.g., not offering an entitled benefit, not following a required disclosure)
- Stating a policy incorrectly, or significantly changing a policy's meaning
- **Temporal sequencing for consequential actions:** When instructions require "explain before acting" or "get explicit confirmation before proceeding," the assistant must pause for user confirmation BETWEEN read operations and write operations that have financial consequences or are irreversible. Executing such read and write operations in the same turn without intermediate user confirmation violates these instructions. Summarizing results TO the caller after the fact does NOT satisfy a requirement to get confirmation FROM the caller before acting.
- **Irreversible write operations** (cancellations, rebookings, refunds, account/record changes, ticket submissions, access grants, etc.) executed (a) without disclosing what the agent's policies (see Agent Instructions above) require to be surfaced before that write (e.g., a fee, balance impact, eligibility caveat, or approval requirement -- varies by policy), or (b) without a clear signal of user intent for that specific action. Do not invent disclosure requirements the policy does not impose.

Is NOT a flag:

- Following reasonable interpretations of ambiguous instructions
- Minor stylistic deviations from instructions that don't affect the outcome (e.g., slightly different wording for a required disclosure)
- Actions not covered by any explicit policy or instruction
- Proactive issuance of no-cost benefits the user is clearly entitled to (e.g., goodwill compensation during a service disruption, waiving a fee the user is explicitly entitled to) without explicit confirmation -- these are beneficial actions with

no negative consequence, and the user's entitlement or explicit request serves as sufficient basis

- When a user explicitly requests a specific action AND the general cost structure has been communicated, proceeding without re-stating exact amounts (if not yet knowable) is not a clear violation
- Proceeding with a write tool based on implicit intent (the user described the action they want and provided the necessary information, but did not give a verbatim "yes, do it now"), when the agent's policies do not require explicit confirmation for that action
- "Predict-then-correct": stating an action should have no fee/cost based on policy *before* calling the tool, then proceeding to call the tool -- no violation **when the tool result confirms the prediction** (e.g., "there should be no fee based on policy" → tool call → fee is waived as predicted). If the tool result contradicts the prediction (e.g., a fee is charged anyway), that is a faithfulness issue: the user was committed to the action under incorrect cost expectations, and an after-the-fact correction doesn't undo it.

Evaluating policy application:

When two policy paths could apply (e.g., same-day change vs. missed flight; in-warranty repair vs. out-of-warranty replacement), consider timeline and eligibility carefully. If the user is within the more favorable policy window (e.g., a flight hasn't departed yet, a request is within SLA, a return is within the return window), applying the more favorable applicable policy is not a violation. Also, if two policy paths produce the identical fee/outcome, choosing one over the other is not a material violation.

4. failing_to_disambiguate

Scope: Handling of ambiguous or contradictory information.

Did the assistant make assumptions or proceed without clarification when the user's input was ambiguous or contradictory? {disambiguation_context}

IS a flag:

- Proceeding with an action when the user's request could reasonably refer to multiple options and the assistant did not ask which one
- Making assumptions about user intent when the user provided contradictory information (e.g., user says two different dates)
- Choosing between conflicting pieces

of information without asking the user to clarify

- Not clarifying ambiguous input that has an impact on the downstream conversation. For example, "after noon" and "afternoon" could refer to different times of day and should not be silently inferred. The agent should not make a decision that excludes available options without validating the user's intent.
- When unable to retrieve some information, if the conversation contains multiple differing versions of a confirmation code or name, the assistant should actively disambiguate rather than silently defaulting to one version or the latest one. Making look-up tool calls is inexpensive and should be done to resolve any ambiguity.
- Failing to consider possible input errors when a lookup fails for an uncommon name or alphanumeric code (e.g., not asking the user to spell it out or verify)
- Not leveraging required information, such as specific confirmation number or names, that could be reasonably inferred from the conversation.

Is NOT a flag:

- Proceeding when the user's intent is clear and unambiguous
- Asking a clarifying question when the user's request is ambiguous (this is correct behavior)
- Making a reasonable inference when the context makes the intent obvious (e.g., user says "my flight" when they only have one flight)
- Retrying a lookup with a corrected spelling after the user confirms or spells out the information -- this is proper disambiguation behavior
- Trying valid different combinations of names and confirmation codes when a lookup fails (e.g., swapping commonly confused letters like "v"/"z" or "b"/"d", reordering characters)

5. hallucination

Scope: Information the assistant states to the user that has no source -- not already covered by the preceding dimensions. Did the assistant present information that was not provided by the user, not returned by any tool response, and not stated in the agent instructions or system context?

IS a flag:

- Stating facts, details, or numbers that do not appear in any tool response, user utterance, agent instruction, or system context

(e.g., inventing a gate number, adding a benefit the user doesn't have, fabricating an ID or reference number, inventing an amount or deadline)

- Presenting fabricated policies, timelines, or conditions not found in any available source
- Claiming the system can perform lookups or actions using identifiers not supported by any available tool (e.g., offering to look up a record by an identifier the tools don't accept, or offering a capability not present in the available tools)
- Misidentifying the brand, company, or agent role (e.g., using a different airline name, naming the wrong organization)

Is NOT a flag:

- Stating information that is directly inferable from tool results and/or system context (e.g., computing an arrival time from departure + duration, calculating an expiration date from current date + valid_months, or computing a remaining amount from a limit minus used amount)
- Referencing the current date/time from the system context -- this is grounded information, NOT hallucination
- Providing general conversational courtesies that don't assert factual claims
- Hedged, commonsense caveats (e.g., "you may want to verify at the counter") that don't contradict tool results or policy -- only flag fabricated information presented as definitive fact
- General domain knowledge (e.g., standard check-in windows, typical appointment lead times, standard password-reset flows) that is reasonable and not contradicted by tool results

Critical verification step: Before flagging hallucination, check ALL available sources: (1) all tool responses in the conversation, (2) user utterances, (3) agent instructions, (4) the Current Date/Time field and other system context metadata -- do NOT assume these fields are empty without verifying. Information derived from system context (e.g., current date) is grounded, not hallucinated.

Disambiguation from other dimensions:

- If the assistant misquotes, distorts or embellishes a tool result → flag under `misrepresenting_tool_result` (the source exists but was reported incorrectly)

- If an unsupported capability is offered in passing → flag here; if actually attempted via fabricated tool call → flag under fabricating_tool_parameters
- If the assistant states something with NO source at all → flag here

You will focus only on the above dimensions. You will NOT consider conversation flow, task completion, or other criteria outside of faithfulness.

Rating Scale

For all five dimensions, determine if there is evidence that one or more issues should be flagged and rate that dimension based on the following guidelines. Severity hinges on the **impact on the user** -- both inside the conversation and beyond it.

- **3** (No faithfulness issues):
 - No issue with this dimension
- **2** (Minor or ambiguous faithfulness issues -- low user impact):
 - A single isolated issue that does not materially affect the outcome, such as:
 - * Calling a **read-only** tool (e.g., a lookup, search, or status tool) with wrong parameters, when the error is caught quickly and no incorrect information reaches the user
 - * A small hallucination or misstatement with no consequence inside the conversation and no downstream effect (e.g., a minor phrasing embellishment that does not alter any decision or action)
 - * A minor ungrounded tool parameter that doesn't affect results
 - * A minor policy deviation that doesn't affect the user's decision-making or their understanding of the policy
 - * Skipping a read-back-and-confirm step on values (dates, names, identifiers) the user stated in preceding turns, when the values written to the tool call match what appears in the user's transcribed utterance. **This carve-out does NOT apply when the transcript shows signals of a likely transcription/STT divergence** -- e.g., the user explicitly said an earlier readback was wrong, OR the transcript contains two or more different versions of the same value across turns. In those cases the assistant had clear signal that a readback was needed to surface the discrepancy, and skipping it likely committed a wrong value

- to the system -- rate **1**, not 2.
- Minor instruction-following deviations that do not materially affect the outcome (e.g., slight formatting differences, omitting low-importance optional steps)
- Borderline cases where it is unclear whether a faithfulness violation occurred due to ambiguous instructions, incomplete context, or reasonable interpretation differences
- Adopting incorrect terminology from the user (e.g., wrong brand or product name) while processing the correct record, when it doesn't cause confusion or incorrect actions
- If something appears as being borderline an issue, it should probably be rated 2.
- **1** (Clear faithfulness violations -- major user impact):
 - Any issue that materially affects the user, either during the conversation or afterward. Such as:
 - * Calling a **write** tool with wrong or fabricated parameters -- regardless of whether the call succeeds.
 - * Committing a value to a write tool when the transcript shows clear signals of a likely transcription/STT divergence (e.g., the user explicitly said an earlier readback was wrong, OR the transcript contains two or more different versions of the same value across turns) and the assistant did not resolve the ambiguity before committing -- **rate 1 even if the written value happens to match the most recent transcribed utterance**, because the divergence signal made a readback necessary and the assistant likely committed a wrong value.
 - * Executing an **irreversible action without the explicit user confirmation that the agent's policies (see Agent Instructions) require** for that specific action (e.g., a "summarize and confirm before submitting" rule). If the agent's policy does not require explicit confirmation for that action, implicit intent from the conversation flow is sufficient -- do not flag the absence of a verbatim "yes, do it now."
 - * Executing an **irreversible action without first disclosing what the agent's policies (see Agent Instructions) require to be surfaced** before that write (e.g., balance impact,

fee, eligibility caveat, approval requirement -- varies by policy). Do not invent disclosure requirements the policy does not impose.

- * **Financial or reputational impact** on the user or the company -- e.g., communicating incorrect charges, fees, refunds, balances, or amounts; making commitments the company cannot honor; misstating obligations in a way that could harm trust
- * Leaving the user with an **incorrect understanding of company policy** that they could act on in the future (e.g., misstating eligibility rules, entitlement conditions, or process requirements), even if the current conversation ends up fine
- * Hallucinating information not present in tool results, especially consequential figures (costs, balances, dates, approvals) communicated to the user as fact
- Any faithfulness issue that repeatedly prevents the conversation from progressing is also rated 1.

For the final rating of the conversation, use the minimum rating across all dimensions as the overall faithfulness rating (i.e., if any dimension is rated 1, overall rating is 1; if all dimensions are 3, overall rating is 3; if there are no 1s but at least one 2, overall rating is 2).

Response Format

Respond in JSON format:

```
{
  "dimensions": {
    "fabricating_tool_parameters": {
      "evidence": "<string>",
      "flagged": <bool: true if issue is present, false otherwise>,
      "rating": <int: 1, 2, or 3>
    },
    "misrepresenting_tool_result": {
      "evidence": "<string>",
      "flagged": <bool: true if issue is present, false otherwise>,
      "rating": <int: 1, 2, or 3>
    },
    "violating_policies": {
      "evidence": "<string>",
      "flagged": <bool: true if issue is present, false otherwise>,
      "rating": <int: 1, 2, or 3>
    },
    "failing_to_disambiguate": {
      "evidence": "<string>",
      "flagged": <bool: true if issue is present, false otherwise>,
      "rating": <int: 1, 2, or 3>
    }
  }
}
```

```
}},
"hallucination": {
  "evidence": "<string: 1-2 sentences citing specific examples from the transcript, or 'None' if not flagged>",
  "flagged": <bool: true if issue is present, false otherwise>,
  "rating": <int: 1, 2, or 3>
}},
"rating": <int: 1, 2, or 3 -- minimum rating across all dimensions>
}}
```

disambiguation_context (Cascade)

Since the assistant is working from a speech-to-text transcript, it should account for potential transcription errors, and clarify any ambiguity in the user's intent, especially when they lead to write/irreversible operations. It's not needed to clarify if the tools called are simple lookups, but if the lookups fail, the assistant is expected to clarify the user's intent.

disambiguation_context (S2S / hybrid)

Since the assistant processes raw audio directly (speech-to-speech), it should account for potential audio perception errors -- mishearing letters, numbers, names, or codes is common with spoken input. The assistant should clarify any ambiguity, especially for alphanumeric codes, names, and values that lead to write/irreversible operations. It's not needed to clarify if the tools called are simple lookups, but if the lookups fail, the assistant is expected to clarify the user's intent. The bar for disambiguation is higher than for a text-based system because the assistant knows it is working from audio and should anticipate mishearings.

misrepresentation_pipeline_note (S2S)

****Speech-to-speech scoping for this dimension.**** Because assistant turns in the trace are STT-transcribed audio (see **About assistant turns** above), token-level discrepancies between an assistant utterance and a tool result -- dropped/added dashes, single-character substitutions, missing or extra digits within long alphanumeric IDs, altered spacing -- typically reflect TTS-rendering or STT-transcription artifacts and are scored by 'agent_speech_fidelity', not here. Only flag

```
'misrepresenting_tool_result'
when the discrepancy is
structural/semantic (wrong field,
wrong order of magnitude, wrong
category) or when downstream signals
-- subsequent tool calls, follow-up
actions, user objections -- show the
agent was internally operating on a
wrong value.
```

K.3 Speech Fidelity

Speech Fidelity Prompt

You are an expert evaluator judging the fidelity of this audio file against the intended text. You will listen to one audio clip and verify that the spoken content faithfully reproduces the intended text, with special attention to TTS-critical entities. The audio provided is a recording of the agent's side of a conversation, and contains only the agent responses, not the user.

```
## Intended Turns
{intended_turns_formatted}
```

```
## IMPORTANT: Comparison Rules
```

Your task is to compare the **exact intended text** word-for-word against what you hear in the audio. The TTS-critical entities highlight which parts are most important to verify, but they do NOT replace or override the intended text.

```
## Understanding the Intended Text
```

The intended text may contain non-spoken tags and markers. You must understand these to evaluate fairly.

```
### Audio-Direction Tags
Tags like [slow], [firm], [annoyed]
describe how the words were meant to
be spoken. They are NOT spoken aloud
and should never be expected in the
audio.
```

```
### Interruption Tags
{interruption_tags_reference}
```

The tags tell you that certain portions of the intended text were likely never spoken, because the speaker was interrupted or cut themselves off. Do NOT penalize for missing words that fall in a region the tags indicate was not spoken.

Key principle: If a tag indicates that a section of text was likely not spoken aloud (due to interruption or cut-off), do NOT penalize for those words being missing from the audio.

Only evaluate fidelity for words that were reasonably expected to have been spoken.

```
## Evaluation Criteria
```

For each intended turn, compare what you hear in the audio against the intended text. Focus especially on **TTS-critical entities** listed for each turn.

Entity categories to watch:

- Confirmation codes (e.g., ZK3FFW, FAR0UM, 8JVSDF)
- Domain-specific identifiers (e.g., flight numbers like "SkyWay 410", ticket or incident numbers, order numbers, case IDs)
- Dollar amounts (e.g., \$15, \$1,285.00)
- Short alphanumeric codes (e.g., seat numbers like "21C", room numbers, extension numbers)
- Spelled-out codes (e.g., "Z K three F F W") -- verify EVERY letter and digit individually; "K O L T S F" vs "K O L T S S F" is an error
- Reference IDs with segments (e.g., REF-8JVSDF-001, MEAL-FAR0UM-PAX0) -- verify each segment; "M E L" vs "M E A L" is an error
- Times (e.g., 3:55 PM, 10:30 AM)
- Dates (e.g., March 25th, February 3rd)
- Names (e.g., Mr. Rivera, Rodriguez)

What constitutes an error (rating = 0):

- Any entity spoken incorrectly (wrong digits, letters, amounts, numbers)
- Missing words that change the meaning or omit an entity
- Added words that introduce a factually incorrect entity
- Substituted words that alter an entity value

What to ignore (does NOT cause rating = 0):

- Minor pronunciation variations that do not change the identity of an entity (e.g., "Ms." vs "Miss" is acceptable)
- Filler words ("um", "uh", "so") added or omitted
- End-of-audio cut-off: if the audio cuts off at the very END of the last turn, missing trailing words is acceptable as long as all entities in that turn were spoken correctly before the cut-off
- Slight pacing or prosody differences
- Non-spoken tags: [slow], [firm], [annoyed], and all interruption tags listed above
- Words in regions flagged by interruption tags as likely not spoken

```

## Rating Scale (per turn)
• 1 (High Fidelity): All entities are spoken correctly. Non-entity words are faithfully reproduced with no meaningful omissions or additions.
• 0 (Low Fidelity): One or more entity errors, OR significant non-entity word errors that change the meaning of the turn.
## Response Format
Respond with a JSON object. Each turn entry must include the turn_id matching the turn number shown in the Intended Turns above:
{
  "turns": [
    {
      "turn_id": <int: the turn number from the Intended Turns>,
      "transcript": <string: your transcription of the audio for this turn, use only the audio for this not the intended text>
      "explanation": "<string: 1-3 sentence analysis of fidelity for this turn, citing specific intended vs actual mismatches, noting any regions skipped due to interruption flags>",
      "rating": <0 or 1>
    }
  ]
}

```

K.4 Conversation Progression

Conversation Progression Prompt

You are an expert evaluator analyzing whether a voice assistant effectively moved a conversation forward. You will evaluate the conversation across four dimensions, each scored as a binary flag (true = issue present, false = no issue).

Each dimension evaluates a **different type of action**. Every issue in the conversation maps to exactly one dimension -- there is no overlap. Ensure to consider both the assistant agent instructions and the following agent dimensions when evaluating the conversation.

IMPORTANT -- Scope boundary with faithfulness: This metric evaluates whether the conversation moved forward efficiently. It does NOT evaluate whether the assistant followed policies, complied with user constraints, or acted faithfully to its instructions -- those are faithfulness concerns. If an issue is primarily about the assistant violating a policy or acting against the user's explicit instructions (e.g., taking an action the user

said not to, not disclosing fees), do NOT flag it here even if it also affected conversation flow. Only flag issues where the assistant's conversational choices (questions asked, information repeated, tools called) were themselves inefficient or counterproductive.

IMPORTANT -- Voice conversation context: This is a voice (spoken) conversation, which means speech recognition errors are common. When the assistant repeats a request because the previous attempt was misheard or garbled, this is expected behavior in a voice interface, not a progression issue.

IMPORTANT -- Interruption tags: The transcript may contain inline tags indicating speech overlap. These are informational metadata about the voice interaction -- they are NOT conversation progression issues by themselves:

```
{interruption_tags_reference}
```

When evaluating, treat these tags as natural voice interaction phenomena. Do NOT penalize interruptions themselves. Only flag an issue if the interruption caused observable consequences (e.g., information loss because the agent's cut-off speech contained critical details that were never restated, or unnecessary repetition because the agent repeated already-heard information after being interrupted).

```

## Understanding the Conversation Trace
{user_turns_disclaimer}

```

```
{assistant_turns_disclaimer}
```

```

## Full Conversation with Tools
{conversation_trace}

```

```
## Evaluation Dimensions
```

```

### 1. unnecessary_tool_calls
Scope: Tool call actions only.
Were any tool calls unjustified -- repeated without reason, made without required information, or made for data already available?

```

IS a flag:

- Calling the same tool with the same parameters after a prior successful response (no new user input or error in between)
- Calling a tool with empty or missing required parameters, causing a predictable error (e.g., calling a lookup tool with empty

strings before asking the user for the required identifier)

- Calling a tool when the needed information was already returned by a previous tool response
- Calling a tool to verify something a prior tool response already confirmed

Is NOT a flag:

- Retrying a tool call after a tool error with corrected parameters
- Calling the same tool with different parameters (e.g., different IDs or search criteria)
- Sequential tool calls that each return new, necessary information (e.g., a record lookup followed by a status check followed by a related-details fetch)
- A tool call that fails unexpectedly (the assistant could not have predicted the failure)
- Tool calls that are necessary for the task but were executed prematurely (e.g., before the user confirmed) -- premature execution is a faithfulness/policy compliance issue, not a conversation progression issue
- Tool calls that follow standard agent instructions (e.g., automatically carrying over related attributes or defaults when taking an action) even if the user did not explicitly request those specific actions

CAVEAT: If the model makes 3 or more unnecessary tool calls, this dimension should be rated 1.

2. information_loss

Scope: The assistant's memory of established facts. Did the assistant fail to retain or act on information already established in the conversation -- whether from the user's statements or from prior tool responses?

This dimension is about the assistant **forgetting or ignoring known facts**, regardless of how that failure manifests (re-asking, wrong assumptions, ignoring constraints).

{information_loss_pipeline_note}

IS a flag:

- Re-asking the user for information they already provided (e.g., asking for the confirmation number or reference ID after the user stated it and it was used successfully).
- Ignoring a constraint the user explicitly stated (e.g., the user ruled out a particular action but the assistant still asks about it or asks for the details that would only be needed to take that action)

- Failing to use relevant data from a prior tool response when it was needed for the next step (e.g., not using an identifier returned by an earlier lookup when making a follow-up tool call that requires it)

Is NOT a flag:

- Asking for information the user has not yet provided
- Asking a clarifying question about genuinely ambiguous information
- Asking for authentication or identification details required by the agent instructions (e.g., a confirmation number, reference ID, or user's name) at the start of the conversation
- The assistant acting on information that contradicts what the user said, when the contradiction is due to a faithfulness or policy violation -- flag that under faithfulness, not here. Only flag here if the assistant demonstrably forgot or ignored previously established facts within the conversation flow.

Disambiguation from other dimensions:

- If the assistant re-asks for user-provided info → flag here (not redundant_statements)
- If the assistant makes an unnecessary tool call because it forgot a prior result → flag under unnecessary_tool_calls (the tool action is the observable problem)
- If the assistant proceeds with an action that contradicts the user's stated preference (e.g., choosing a different option than the one the user requested) → this is a faithfulness violation, not information_loss. Only flag here if the assistant clearly forgot the user's input, not if it chose to override it.

3. redundant_statements

Scope: The assistant repeating its own previous output. Did the assistant restate information it had already communicated to the user?

This dimension **ONLY** covers the assistant repeating **its own prior utterances** -- not forgetting user input (that is information_loss) and not tool call issues (that is unnecessary_tool_calls).

IS a flag:

- Restating details, times, amounts, or status information the assistant already told the user in an earlier turn (outside of a final recap) when the user did not ask for it
- Repeating the same explanation or instruction in multiple turns when the user has acknowledged and moved on

Is NOT a flag:

- A brief recap or summary at the end of the conversation (this is helpful, not redundant). However, if the assistant provides multiple recaps across different turns, only the final one is exempt -- earlier recaps that restate already-communicated information are still flagged.
- Confirming back details to the user once for verification (e.g., reading back a confirmation number the user just provided)
- Stating information for the first time, even if it was available from a tool response earlier
- Repeating information in direct response to the user explicitly requesting confirmation or asking to hear it again (the user must clearly ask -- simply continuing the conversation is not a request for repetition)
- Re-explaining a policy or constraint when the user continues to challenge, dispute, or insist against it -- the assistant must reiterate its position in these cases and should not be penalized for doing so. However, if the assistant repeats the exact same explanation verbatim across multiple turns, flag it -- the assistant should vary its phrasing.
- Repeating a request for information (e.g., confirmation code, spelling) when speech recognition or transcription errors clearly caused the previous attempt to fail (e.g., garbled text, partial characters, obvious mishearing visible in the transcript). Do NOT apply this exception when the transcript shows no evidence of ASR failure -- the assistant re-asking without cause is still a flag.

4. question_quality

Scope: The quality and appropriateness of the assistant's questions, where the issue is NOT caused by forgetting information (that is information_loss). Did the assistant ask poorly formed questions or fail to ask a necessary clarifying question?

IS a flag:

- Asking an overly broad or vague question when the assistant had enough information to take action (e.g., "What would you like to do?" when the user already stated a clear goal that the assistant remembers but chose not to act on)
- Asking multiple questions at once when a single tool call could have resolved the need
- Failing to ask for clarification

when the user's request was genuinely ambiguous, and instead proceeding with assumptions

- Failing to ask for clarification when there are multiple options that meet the user's requirements
- Failing to ask for required information before taking an action (e.g., not asking for required details for a tool call before making the tool call, when those details have not been made available through a previous tool call, or inputs from the user)
- Failing to provide necessary information for the user to make a decision (e.g., not providing clear information about the details of the options available to the user)
- Taking an irreversible action (e.g., cancellation, rebooking, ticket submission, access grant, account change) without first confirming when user input is ambiguous or contradicts system data (e.g., user claims a 4-hour delay but system shows 45 minutes -- assistant should clarify before acting)

Is NOT a flag:

- Asking for required authentication or identification information required by the agent instructions (e.g., confirmation number, reference ID, user's name)
- Asking a clarifying question when the user's intent is genuinely ambiguous
- Asking a follow-up question based on new information from a tool response
- Asking the user to confirm an error-prone value (alphanumeric ID, code, date, dollar amount, spelled-out name) just provided -- read-backs of error-prone values are standard voice-agent practice required by typical agent policies and are NOT a question quality issue. This holds even when the transcript shows the assistant's read-back as matching what the user said: STT may render different spoken audio as identical text, so the read-back can still be catching a real audio-level mismatch.
- Not disclosing fees, costs, or other policy-required details before taking an action -- policy compliance (e.g., whether the assistant explained consequences before an irreversible action) is a faithfulness concern, not a conversation progression issue. This dimension only evaluates whether the assistant's questions and information-sharing effectively moved the conversation forward.
- Referencing information that exists

in the agent instructions (e.g., standard fees, policies) without verifying it via a tool call -- the agent is expected to know its own instructions. Only flag if the information was genuinely unknown and required a tool call or user input.

Disambiguation from information_loss:

- If the assistant asks "What would you like to do?" because it FORGOT the user already stated their goal → flag under information_loss
- If the assistant asks "What would you like to do?" when the user's goal is clear and remembered but the assistant chose a vague question over taking action → flag here

Rating Scale

For all four dimensions, determine if there is evidence that one or more issues should be flagged and rate that dimension based on the following guidelines:

- **3** (No progression issue):
 - No issue with this dimension
- **2** (Minor progression issue):
 - A single isolated issue that does not significantly impact the conversation flow (e.g., one unnecessary tool call that didn't slow things down, a single redundant restatement, one vague question)
 - A borderline case where it is unclear whether the issue constitutes a real progression problem
- **1** (Clear progression issue):
 - Multiple instances of the same type of issue in this dimension
 - A single severe issue that clearly derailed or stalled the conversation (e.g., ignoring a stated constraint or user requirement before carrying out a write operation, failing to ask for required information before taking action, asking an overly vague question when the user's goal was clear, making an overly vague assumption not supported by user inputs/conversation history when multiple options exist)

Overall Rating

The final rating considers BOTH the severity within each dimension AND the total number of flagged dimensions:

- **3:** No dimension is flagged (all dimensions rated 3)
- **2:** One or two dimensions are flagged at rating 2 (minor), AND no dimension is rated 1
- **1:** Any of the following:
 - Any dimension is rated 1 (clear issue within a single dimension)
 - Three or more dimensions are

flagged (even if each is individually minor, widespread issues across many areas constitute a clear overall progression problem)

Response Format

Respond in JSON format. The "evidence" field must ALWAYS contain 1-2 sentences referencing specific parts of the transcript, even when flagged is false. When not flagged, briefly explain why no issue was found.

```
{
  "dimensions": {
    "unnecessary_tool_calls": {
      "evidence": "<string: REQUIRED -- cite transcript examples if flagged, or explain why clean if not>",
      "flagged": "<bool: true if issue is present, false otherwise>",
      "rating": "<int: 1, 2, or 3>"
    },
    "information_loss": {
      "evidence": "<string: REQUIRED>",
      "flagged": "<bool: true if issue is present, false otherwise>",
      "rating": "<int: 1, 2, or 3>"
    },
    "redundant_statements": {
      "evidence": "<string: REQUIRED>",
      "flagged": "<bool: true if issue is present, false otherwise>",
      "rating": "<int: 1, 2, or 3>"
    },
    "question_quality": {
      "evidence": "<string: REQUIRED>",
      "flagged": "<bool: true if issue is present, false otherwise>",
      "rating": "<int: 1, 2, or 3>"
    }
  },
  "rating": "<int: 1, 2, or 3>"
}
```

information_loss_pipeline_note (S2S)

****Speech-to-speech scoping for this dimension.**** Because assistant turns in the trace are STT-transcribed audio (see *About assistant turns* above), variant token-level readings of the same alphanumeric identifier across nearby assistant turns -- dropped/added dashes, single-character substitutions, missing or extra digits within long IDs, altered spacing or capitalization -- typically reflect TTS-rendering or STT-transcription artifacts on a value the agent is reading consistently in audio. These are scored by 'agent_speech_fidelity', not here. Only flag 'information_loss' when the discrepancy is structural/semantic (different entity, wrong field, wrong category -- e.g., addressing the user by an entirely different

first name, or referencing a different person/record than the tool returned), or when downstream signals -- subsequent tool calls made with a wrong value, follow-up actions taken on stale data, user objections that the agent then fails to incorporate -- show the agent was internally operating on a wrong value or had genuinely lost track of the established fact."

K.5 Conciseness

Conciseness Prompt

You are an expert evaluator judging the conciseness and voice-appropriateness of assistant responses in a voice conversation.

Conversation
{conversation_turns}

Understanding the Conversation Format

The conversation is grouped by turn_id. Each turn may contain:

- **user:** What the user said
- **assistant:** What the assistant said (there may be multiple assistant entries within a single turn -- e.g., the assistant speaks, calls a tool, then speaks again)
- **tool_call:** A tool invocation made by the assistant
- **tool_response:** The result returned by the tool

When a turn contains multiple assistant entries, evaluate them **together as a single unit** -- they represent the assistant's complete response within that turn. Tool calls and responses between assistant entries explain why the assistant spoke in multiple parts (it was waiting for data). It could also be due to interruptions from the user.

Understanding Interruption Tags

{interruption_tags_reference}

Key principle: When interruption tags are present, the assistant may not have been able to finish what it was saying. Do NOT penalize for truncated or fragmented content caused by interruptions. Only evaluate the conciseness of content the assistant chose to say, not content that might have been cut off.

Instructions

The conversation includes user, assistant, tool_call, and

tool_response entries. Rate only the assistant's spoken content. User turns, tool calls, and tool responses are provided for context only.

For each turn that contains assistant content, evaluate whether the assistant's response is appropriately concise and easy to digest when spoken aloud to a human.

The assistant is expected to follow conversational voice guidelines:

- Keep responses brief and conversational (typically 2-4 sentences)
- Summarize long lists rather than reading them exhaustively
- Avoid overwhelming the listener with too much information at once
- Spread multiple requests across turns when possible
- Present options conversationally and avoid cramming excessive detail into one turn

Evaluation Criteria

When evaluating each turn, consider:

- Does the response get to the point without filler, rambling, or unnecessary content?
- Is all the information relevant and necessary given the conversation context?
- Is the amount of detail reasonable for someone listening to -- not reading -- the response?
- If the response enumerates options or items (e.g., "Option one is... Option two is..."), does the structure help the user? The volume should not be overwhelming.
- Is the provided information justified by context (e.g., confirming a detail the user may have misheard)? Or is it inappropriate (e.g., excessive itemization or explanation when the user may only care about the end result)?
- Within turns, is repetition avoided? Across turns there may be valid reasons for repetition, but it should usually not occur within a single turn.
- Essential information -- such as confirmation codes, reference IDs, ticket numbers, or other specific details the user needs to note down -- should never be penalized, regardless of length.

Allowed Exceptions (Voice Interaction Realities)

The assistant may occasionally produce longer turns when the context requires precise information transfer. The following cases should NOT be penalized for verbosity or information density. The turn itself may still be penalized for other

reasons.

1. Phonetic Confirmation of Codes

- When confirming a confirmation code, reference number, ticket number, or similar identifier, the assistant may spell characters using the NATO phonetic alphabet (e.g., "B as in Bravo, F as in Foxtrot").
- This is especially appropriate when the user previously misheard or asked for clarification.

2. Reference or Identifier Delivery

- When providing an identifier the user needs to note down (e.g., a ticket number, reference code, or voucher code), the assistant may read the whole code out loud.
- This information is essential and should not be penalized regardless of length.

3. End-of-Call Wrap-Up

- The final assistant turn in a conversation may include a slightly longer recap or confirmation of next steps (e.g., summarizing the action taken, confirming what will be sent or followed up on, thanking the user).
- Minor additional detail in this final wrap-up should not be penalized unless it becomes excessively long or introduces unrelated information.

Important principle: Information given in assistant turns must be short enough for an average person to easily follow in real-time conversation and retain in working memory.

Failure Modes

When a response is not optimally concise, identify which of the following failure modes are present. A turn may have multiple failure modes.

verbosity_or_filler

Contains unnecessary wording, repetition within the same turn, hedging, or explanation beyond what the context requires.

excess_information_density

Presents too many distinct facts, options, numbers, steps, or requests at once, making it difficult for a listener to process in real time. Note: bundling closely related transactional details that the user

needs to act on or remember together (e.g., confirming a reference number, date, and one or two key details in a single turn) is expected behavior -- only flag this when the volume of information genuinely exceeds what a listener can comfortably retain.

over_enumeration_or_list_exhaustion

Reads out long lists instead of summarizing, or presents multiple options with excessive detail rather than inviting follow-up.

contextually_disproportionate_detail

Provides more background, clarification, or explanation than the situation warrants.

Contextual Leniency and Failure Mode Priority

Conciseness should be evaluated with respect to the conversational context. If additional wording or detail is clearly necessary for the user to understand or act on the information, a modest increase in verbosity should be considered acceptable and should NOT be penalized.

If none of the above are present, return an empty list for `failure_modes`.

Rating Scale For Each Turn With Assistant Content

- **3** (Highly Concise / No Cognitive Overload) -- The response is clear, appropriately scoped for voice, and comfortably digestible in real time. No failure modes are present. A turn that delivers a few closely related facts as part of a single transactional step (e.g., confirming the key details of a request or incident) still qualifies as 3 if the listener can comfortably absorb it in one pass.
- **2** (Adequate but Not Optimally Concise) -- One minor failure mode is present, but the response remains reasonably processable in a voice setting and does not meaningfully overwhelm the listener. Reserve this rating for turns where you can identify specific content that should have been omitted or deferred to a later turn -- not merely for turns that happen to contain several necessary details.
- **1** (Not Concise / Causes Cognitive Overload) -- One or more significant failure modes are present that materially increase cognitive load and would hinder comprehension in a voice conversation.

Provide one entry per `turn_id` in the

conversation.

Response Format

Provide your response as a valid JSON array, one entry per turn. Each entry must include the `turn_id` matching the turn number shown in the conversation above.

- If the turn contains assistant content, rate it with 1, 2, or 3.
- If the turn does not contain assistant content (e.g., user-only turn), set rating to null.

```
[
  {
    "turn_id": <int: the turn number
from the conversation>,
    "explanation": "<Detailed analysis
referencing the evaluation criteria
and explicitly linking identified
weaknesses to the listed failure
modes to justify the selected rating
(1-3). Empty string if rating is
null.>",
    "failure_modes":
    ["<failure_mode_1>",
    "<failure_mode_2>", ...],
    "rating": <int: 1, 2, or 3, or null
if no assistant content>
  }
]
```

If the turn is rated 3 or null, `failure_modes` must be an empty list: `[]`.

K.6 User Speech Fidelity

User Speech Fidelity Prompt

You are an expert evaluator judging the fidelity of text-to-speech (TTS) audio against the intended text. You will listen to one audio clip and verify that the spoken content faithfully reproduces the intended text, with special attention to TTS-critical entities.

Evaluation Mode: User

Intended Turns
{intended_turns_formatted}

Understanding the Intended Text

The intended text may contain non-spoken tags and markers. You must understand these to evaluate fairly.

Audio-Direction Tags

Tags like [slow], [firm], [annoyed] describe how the words were meant to be spoken. They are NOT spoken aloud and should never be expected in the audio.

Interruption Tags {interruption_tags_reference}

The tags tell you that certain portions of the intended text were likely never spoken, because the speaker was interrupted or cut themselves off. Do NOT penalize for missing words that fall in a region the tags indicate was not spoken.

Key principle: If a tag indicates that a section of text was likely not spoken aloud (due to interruption or cut-off), do NOT penalize for those words being missing from the audio. Only evaluate fidelity for words that were reasonably expected to have been spoken.

Evaluation Criteria

TTS-Critical Entities (check these carefully)

- **Personal names:** "John Smith" vs "Jim Smith"
- **Dates and times:** "December 15th" vs "December 50th", "3:45 PM" vs "3:15 PM"
- **Reference codes:** Confirmation numbers, incident numbers, booking IDs (e.g., "QWMN62" vs "QWN62")
- **Numeric values:** Dollar amounts, quantities, percentages (e.g., "\$150" vs "\$115")
- **Addresses:** Street numbers, street names, cities (e.g., "123 Main Street" vs "124 Main Street")
- **Contact information:** Phone numbers, email addresses (e.g., "tom_cobb@gmail.com")
- **Flight/route numbers:** "UA204" vs "UA240"
- **Serial numbers and other identifiers**

Error Types

- **Missing words:** Words in the intended text that were not spoken AND were reasonably expected to have been spoken (i.e., not in a region flagged by interruption tags)
- **Added words:** Extra words spoken that are not in the intended text
- **Wrong words:** Words spoken incorrectly or substituted with different words
- **Entity errors:** Any of the TTS-critical entities above spoken incorrectly

What to Ignore

- Non-spoken tags: [slow], [firm], [annoyed], and all interruption tags listed above
- Words in regions flagged by interruption tags as likely not spoken
- Minor pronunciation variations that do not change meaning (accent

differences)

- Natural filler words (um, uh) if they do not affect core content
- Missing words at the END of the LAST turn only (audio recordings are often cut off before the final utterance completes). However, missing words in the middle of the last turn, or missing words in any earlier turn, should still be penalized.

Rating Scale (per turn)

• **3 (High Fidelity):**

- All expected entities spoken correctly (names, dates, destinations, codes, etc)
- All words reasonably expected to have been spoken are present and accurate.
- Minor pronunciation variations acceptable.
- No audio tags spoken out loud.

• **2 (Medium Fidelity):**

- All entities spoken correctly (names, dates, destinations, codes, etc)
- Part of a turn may be missing (often in the first turn, the first few words are missing)
- Some words that were reasonably expected may be missing or spoken slightly incorrectly, but they are not critical and the conversation is able to progress even with this issue.
- Potential issues with audio tags being said out loud

• **1 (Low Fidelity):**

- One or more entity errors (missing entities, incorrect entities, etc) OR
- Some other major error that prevents the conversation from continuing in a sensible manner.

Response Format

Respond with a JSON object. Each turn entry must include the turn_id matching the turn number shown in the Intended Turns above:

```
{
  "turns": [
    {
      "turn_id": <int: the turn number
        from the Intended Turns>,
      "transcript": <string: your
        transcription of the audio for this
        turn, use only the audio for this not
        the intended turns>,
      "explanation": "<succinct analysis;
        for score 1 or 2, quote the specific
        issue with intended vs actual;
        note any regions skipped due to
        interruption tags>",
      "rating": <1, 2, 3>
    }
  ]
}
```

K.7 User Behavioral Fidelity

User Behavioral Fidelity Prompt

You are an expert evaluator determining whether a simulated user's behavior has corrupted the voice agent evaluation.

Your job is to determine whether the user's behavior caused the agent to be evaluated unfairly -- specifically, whether the user's actions led to the database being in a different state than it should be, or prevented the agent from completing actions it otherwise would have.

Conversation Evidence
{conversation_evidence}

How the conversation ended
The conversation ended due to:
{conversation_end}

User Simulator Instructions
The following is the full system prompt the user simulator was given for this conversation, including the user's persona, goal, decision tree, must-have criteria, and end-of-call rules. This is the source of truth for what the user-sim was told to do. Evaluate user behavior against these instructions, not against generic notions of "what a real user would do."
```\n{user\_simulator\_instructions}\n```\n

## Modification Tools  
The following are the tools that modify database state. These are the only tools relevant to corruption analysis -- read-only tools are not a concern.  
{modification\_tools}

## Evaluation Criteria

### Guiding principle: judge the user from the user's perspective  
The user is a simulated caller who only hears the agent's spoken words. They cannot see tool calls, tool responses, or the agent's internal state. Judge the user's behavior -- what they said, what they refused to say, when they ended the call -- based on what the agent said and asked them, not on what the tool-call trace reveals. Use the tool-call trace asymmetrically:

- **To exonerate the user (always allowed):** confirm that an agent failure occurred that the user couldn't have known about -- e.g., the agent claimed "I'm submitting

your request" but no modification tool was called (hallucinated tool call), or the agent's tool returned an error the user couldn't have prevented.

- **To convict the user (use carefully):** only when the user's *visible behavior* (their words, their refusals, their choice to end) plausibly caused the modification problem. A missing or wrong modification by itself is not evidence against the user; the user's behavior must be the proximate cause.

Analyze the conversation for the following corruption scenarios:

### Corruption Type 1: User invented requests that caused extra modifications

The user made requests OUTSIDE of their assigned goal that caused the agent to call one or more modification tools listed above.

- Only flag this if the user's off-script request directly led to a modification tool being called.
- If the user went off-script but the agent only called read-only tools (e.g., searching, looking up information), this is NOT corruption.

### Corruption Type 2: User ended the conversation prematurely  
The user ended the conversation before the agent had the opportunity to complete the necessary modification tools to fulfill the user's goal.

Applicability:

- This type only applies when the conversation ended due to the user calling the end call tool.
- If the conversation ended due to the agent failing to respond to the user's final turn, this corruption type should never be flagged since the user wasn't the one to end the call.

Decision procedure: flag this ONLY if a reasonable user following their decision tree could have enabled the missing modification by staying on the call. The burden is on showing the user blocked an action the agent was on track to complete.

Do NOT flag (agent failures that exonerate the user, even if a required modification did not happen):

- **Agent error or retry offer.** The agent reported a tool error or technical issue and offered to "try again" / "look into it." The user is not obligated to wait through indefinite retries; ending after an error is consistent with their

failure condition.

- **Hallucinated tool call.** The agent claimed to perform an action ("I'm submitting your request now", "I've assigned the SLA") but the modification tool was never actually called in the trace. Staying on the call longer would not have made the agent call the tool -- the agent already believed the work was done.
- **Off-script demand.** The agent asked the user to do or provide something that is not in the user's goal or persona (run a diagnostic command, provide a value the agent should already have, perform a task outside the scenario). The user cannot comply with invented prerequisites, and ending after such a demand is not corruption.
- **Skipped wrap-up before the agent's own goodbye.** The user expressed they were done or that the issue was resolved, and the agent failed to call its required closing modification (e.g., mark\_resolved, close\_ticket) before the goodbye exchange. The agent should call wrap-up tools before closing rapport, not after -- the user is not responsible for prompting the agent to call its own wrap-up tool.
- **Agent stuck or unhelpful** for multiple consecutive turns, or the agent said it could not help.
- **Resolution condition met.** All necessary modifications were completed and the user's goal was satisfied.

DO flag:

- The agent was actively and competently progressing (asking for a specific piece of information that exists in the user's goal, about to call a tool) and the user cut it off without cause.
- The user's decision tree explicitly required waiting for a confirmation / request ID, and the user ended without one despite the agent making normal progress (no error, no stuckness, no off-script demand, no hallucinated tool call).
- Special case -- transfer to a live agent:
  - The user is instructed to always end the call AFTER the agent transfers to a live agent. This should not be penalized.
  - However the user must give the agent a chance to actually call the transfer\_to\_agent tool. If the user says "Ok go ahead and transfer me. Goodbye" and ends before the tool call lands, this is a premature end and MUST be flagged.

### Corruption Type 3: User failed to provide required information  
The user failed to provide information from their goal that the agent explicitly asked for, preventing the agent from completing a necessary modification tool call.

- Only flag this if the agent clearly asked for specific information that was available in the user's goal, the user failed to provide it, and this directly prevented a modification tool from being called.
- Do NOT flag this if the agent never asked for the information.

### Corruption Type 4: User looping caused duplicate modifications  
The user repeatedly made the same request in a loop, causing the agent to call the same modification tool multiple times when it should have only been called once.

- Only flag this if the looping directly caused duplicate or extra modification tool calls.
- If the user looped but the agent handled it correctly (did not call extra modification tools), this is NOT corruption.

### Corruption Type 5: User violated decision tree instructions causing a wrong modification  
The user explicitly violated a specific instruction in their decision tree (negotiation behavior, edge cases, escalation behavior, resolution condition, or failure condition) AND this violation directly caused a modification tool to be called with different parameters than it would have been if the user had followed their instructions correctly.

- Examples: the user accepted an option that did not meet their must-have criteria when they should have rejected it; the user ignored an edge case instruction (e.g., accepted a standby flight when told to reject standby) and this led to a modification; the user failed to follow their failure condition and instead accepted an unsuitable resolution.
- Only flag this if the violation directly caused a modification tool to be called incorrectly. If the user deviated from instructions but no modification tool was affected, this is NOT corruption.
- Do NOT flag this if the agent only presented options that failed to meet the user's must-have criteria AND the user had no correct option to choose -- in that case the agent failed, not the user. Only flag this if the user had a correct action available

(e.g., rejecting all options, asking for alternatives, triggering the failure condition) but chose incorrectly instead.

## Rating

#### Binary Rating:

- **1 (Clean):** The user's behavior did not corrupt the agent evaluation. None of the corruption types above occurred. Minor deviations from the user's instructions that did not affect database state are acceptable.
- **0 (Corrupted):** One or more corruption types occurred -- the user's behavior caused the agent to be evaluated against an incorrect database state.

Respond in JSON format:

```
{
 "corruption_analysis": {
 "extra_modifications": {
 "analysis": "<reasoning about whether the user made off-script requests that caused modification tool calls>",
 "detected": <bool>
 },
 "premature_ending": {
 "analysis": "<reasoning about whether the user ended the call before the agent could complete necessary modifications>",
 "detected": <bool>
 },
 "missing_information": {
 "analysis": "<reasoning about whether the user failed to provide requested information that blocked a modification>",
 "detected": <bool>
 },
 "duplicate_modifications": {
 "analysis": "<reasoning about whether user looping caused duplicate modification tool calls>",
 "detected": <bool>
 },
 "decision_tree_violation": {
 "analysis": "<reasoning about whether the user violated a specific instruction and whether a correct action was available, and whether this caused an incorrect modification>",
 "detected": <bool>
 }
 },
 "rating": <int: 0 or 1>
}
```

## K.8 Transcription Accuracy (Key Entities)

### Transcription Accuracy (Key Entities) Prompt

You are an expert evaluator analyzing Speech-to-Text (STT) transcription accuracy for key entities across an entire conversation.

Your task:

1. For EACH user turn, identify all key entities in the EXPECTED text
2. Check if each entity appears CORRECTLY in the TRANSCRIBED text

3. Mark each entity as correct or incorrect
4. For entities in regions that were likely never spoken aloud (as indicated by interruption tags), still include them in the output but mark them as skipped

#### ## What Counts as an Entity

An entity must have a **specific, concrete value** -- something that could be passed as an input to a program or tool (not an AI, but a script or database lookup). Ask yourself: could this value be stored in a variable and used programmatically?

- Names (people, places, organizations): e.g. "John Smith", "Austin", "Delta Airlines"
- Specific dates and times: e.g. "December 15th", "3:45 PM" -- NOT vague references like "tomorrow morning" or "later today"
- Confirmation codes / reference numbers: e.g. "ABC123", "ZK3FFW"
- Flight numbers: e.g. "UA 204"
- Amounts and prices: use the specific value only, e.g. "\$120" -- for qualifier phrases like "under \$120", only use the specific value
- Addresses: e.g. "123 Main Street"
- Phone numbers: e.g. "555-867-5309"
- Email addresses: e.g. "john@example.com"
- Other specific identifiers: seat numbers, loyalty numbers, booking IDs, etc.

**Not an entity:** vague temporal words ("tomorrow", "next week", "morning"), general descriptors ("the cheap flight", "a long trip"), or open-ended qualifiers ("less than an hour", "around noon").

#### ## Understanding Tags in the Expected Text

The expected text may contain non-spoken tags and markers. These are metadata -- they were never said aloud and must not be treated as entities or evaluated.

#### ### Audio-Direction Tags

Tags like [slow], [firm], [annoyed] describe how the words were meant to be spoken. Ignore them entirely.

#### ### Interruption Tags

{interruption\_tags\_reference}

These markers indicate that parts of the expected text may never have been spoken aloud, because the user was interrupted or talked over. An entity that was never spoken cannot be correctly transcribed, so you must

NOT penalize for entities in regions that were likely not said, instead mark them as skipped.

**Key principle:** Only evaluate entities that were reasonably expected to have been spoken aloud. If a tag indicates the user was interrupted or talked over before or during an entity, still include the entity in your output but set `skipped: true` and explain why in the analysis. The correct field should reflect your best assessment of whether the transcription matched, but skipped entities will be excluded from accuracy metrics downstream.

#### ## User Turns to Evaluate

{user\_turns}

#### ## Correctness Criteria

- Entity must be present (not missing) -- unless in a region flagged by interruption tags
- Entity value must match (minor formatting variations OK)
- Numbers: "150" and "one hundred fifty" are equivalent
- Dates: "December 15th" and "Dec 15" are equivalent
- Names: Case-insensitive exact match required

Important note: The expected text will often feature things formatted like "one two three" instead of "123". Your goal is to evaluate the semantic equivalence, meaning these are considered equivalent if they were heard in audio.

#### ## Examples

##### Example Input:

Turn 1:

Expected: My confirmation is A B C one two three on December 15th.

Transcribed: My confirmation is ABC123 on December 15th.

Turn 2:

Expected: Transfer one hundred fifty to account 1 2 3 4 5.

Transcribed: Transfer \$115 to account 12345.

Turn 3:

Expected: [slow] The code is X X F six O H, with the letter O, [assistant interrupts] not zero.

Transcribed: The code is X... X F 6 O H with the letter O.

Turn 4:

Expected: My phone number is four zero four five five five [assistant interrupts] zero eight five six.

Transcribed: My phone number is 404-555.

**Example Response:**

```
[
 {{
 "turn_id": 1,
 "entities": [
 {{
 "type": "confirmation_code",
 "value": "A B C one two three",
 "transcribed_value": "ABC123",
 "analysis": "Matches exactly",
 "correct": true,
 "skipped": false
 }},
 {{
 "type": "date",
 "value": "December 15th",
 "transcribed_value": "December 15th",
 "analysis": "Matches exactly",
 "correct": true,
 "skipped": false
 }}
],
 "summary": "All 2 key entities transcribed correctly."
 }},
 {{
 "turn_id": 2,
 "entities": [
 {{
 "type": "amount",
 "value": "one hundred fifty",
 "transcribed_value": "$115",
 "analysis": "Amount wrong: $150 vs $115",
 "correct": false,
 "skipped": false
 }}
],
 {{
 "type": "account_number",
 "value": "1 2 3 4 5",
 "transcribed_value": "12345",
 "analysis": "Matches exactly",
 "correct": true,
 "skipped": false
 }}
],
 "summary": "1 out of 2 entities correct. Amount error."
}},
 {{
 "turn_id": 3,
 "entities": [
 {{
 "type": "confirmation_code",
 "value": "X X F six O H",
 "transcribed_value": "X F 6 O H",
 "analysis": "Missing one X -- transcribed 5 characters instead of 6. The code appears before the [assistant interrupts] tag so it is evaluated normally.",
 "correct": false,
 "skipped": false
 }}
],
 "summary": "1 entity found before interruption, partially incorrect
```

```
(missing one X). No entities after [assistant interrupts] tag to skip."
}},
 {{
 "turn_id": 4,
 "entities": [
 {{
 "type": "phone_number",
 "value": "four zero four five five five zero eight five six",
 "transcribed_value": "404-555",
 "analysis": "The full number is 404-555-0856. The [assistant interrupts] tag appears after 'five five five', meaning the last four digits ('zero eight five six') were likely drowned out by the agent speaking over the user. The transcription captured the portion before the interruption. Skipping because the entity spans into the interrupted region and cannot be fully evaluated.",
 "correct": false,
 "skipped": true
 }}
],
 "summary": "1 entity found. Phone number spans into interrupted region -- skipped. Partial transcription (404-555) matches the portion before the interruption."
 }}
]
```

**## Response Format**

Respond with a JSON object. Each turn entry must include the turn\_id matching the turn number shown in the User Turns to Evaluate section above:

```
[
 {{
 "turn_id": <int: the turn number from the User Turns to Evaluate section>,
 "entities": [
 {{
 "type": "<name|date|flight_number, etc ...>",
 "value": "<entity value from expected text>",
 "transcribed_value": "<how it appeared or 'missing'>",
 "analysis": "<brief reason; if skipped, explain why the entity falls in an interrupted region>",
 "correct": <true|false>,
 "skipped": <true|false>
 }}
],
 "summary": "<1-2 sentence summary for this turn>"
 }}
]
```

Table 39: Third-party dependencies and their licenses.

License	Packages
MIT	pydantic, elevenlabs, litellm, deepgram-sdk, onnxruntime, azure-cognitiveservices-speech, cartesia, assemblyai, setuptools, fastapi, pyyaml, pydub, jaconv, more-itertools, pytest, pytest-cov, ruff, mypy, inflect
Apache-2.0	openai, aioboto3, google-generativeai, google-genai, google-cloud-speech, google-cloud-texttospeech, aiofiles, jiwer, streamlit, pytest-asyncio, regex
BSD-2-Clause	pipecat-ai
BSD-3-Clause	uvicorn, websockets, httpx, pandas, numpy, python-dotenv
MIT / Apache-2.0	structlog
MIT / MPL-2.0	tqdm

## L Licenses & Usage

### L.1 Third Party Licenses

EVA-Bench depends on the third-party software packages listed in Table 39. Full license texts are available in the code repository `THIRD_PARTY_NOTICES` file. All existing artifacts are used in a manner consistent with the license type and intended use.

### L.2 EVA-Bench License

EVA-Bench is released under the MIT license and is intended to be used in the following ways:

- **Academic research:** Benchmarking and evaluating conversational voice agents across task accuracy and experiential quality dimensions in research settings.
- **Enterprise evaluation:** Organizations may use EVA-Bench to evaluate and compare voice agent deployments across enterprise domains.
- **Framework extension:** Researchers and practitioners are encouraged to extend EVA-Bench with new data domains and metrics.

We discourage use of EVA-Bench for purposes that misrepresent evaluation results, such as selective reporting of metrics or use in misleading competitive benchmarking claims. Evaluation outputs should be interpreted in conjunction with the statistical methodology described in Section 4.3 and with awareness of the limitations section.

Additionally, the EVA-Bench dataset and evaluation scenarios should not be used as training data for voice agents or language models. Training on benchmark data undermines the integrity of future evaluations and may lead to artificially inflated performance scores that do not reflect real-world capability.

## M AI Usage

We used large language model (LLM) assistants to support writing and phrasing in select sections of this paper. AI coding agents were additionally used to assist with code scaffolding and implementation during the development of the framework. All research design, experimental analysis, result interpretation, and scientific conclusions were conducted entirely by the human authors.

## N Human Annotations

**Use of Human Annotators.** We employed human annotators to measure judge-human agreement on evaluation scores. All annotators are full-time AI data linguists employed by the same company as the authors; annotation is a core component of their regular job responsibilities. A total of seven linguists participated, with each sample independently annotated by two linguists.

**Annotation Instructions.** Annotators were given the same prompts and scoring rubrics used by the LLM judge (see Section K). No additional guidelines or supplementary instructions were provided beyond the judge prompts, in order to ensure a fair comparison between human and automated judgments.

**Recruitment and Compensation.** Annotators were recruited from an internal team of full-time AI data linguists. As salaried employees, their compensation is determined by standard employment agreements. No external crowdworkers or research subjects were recruited for this work.

**Consent and Data Use.** The annotation task involved evaluating synthetically generated conversational data; no personal data from the annotators was collected or used. All annotators were informed of the purpose of the annotations and how they would be used in the research prior to beginning the task.

**Risks.** The annotation tasks posed no foreseeable risks to participants. The data consisted of synthetic enterprise voice agent conversations and did not contain sensitive, offensive, or personally identifiable content.