

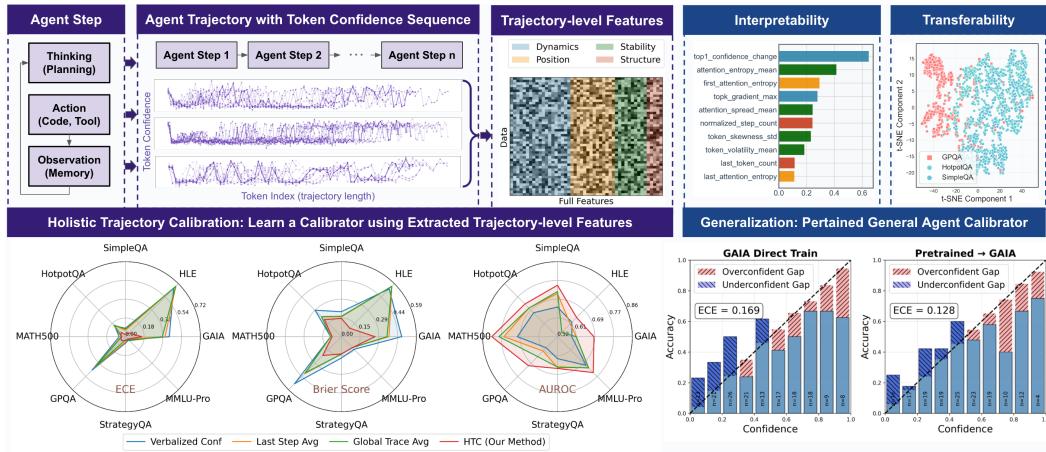
000 AGENTIC CONFIDENCE CALIBRATION

001 **Anonymous authors**

002 Paper under double-blind review

003 ABSTRACT

004 AI agents are rapidly advancing from passive language models to autonomous
 005 systems executing complex, multi-step tasks. Yet their overconfidence in failure
 006 remains a fundamental barrier to deployment in high-stakes settings. Existing
 007 calibration methods, built for static single-turn outputs, cannot address the unique
 008 challenges of agentic systems, such as compounding errors along trajectories, uncer-
 009 tainty from external tools, and opaque failure modes. To address these challenges,
 010 we introduce, for the first time, the problem of *Agentic Confidence Calibration*
 011 and propose **Holistic Trajectory Calibration (HTC)**, a novel diagnostic frame-
 012 work that extracts rich process-level features ranging from macro dynamics to
 013 micro stability across an agent’s entire trajectory. Powered by a simple, inter-
 014 pretable model, **HTC** consistently surpasses strong baselines in both calibration and
 015 discrimination, across eight benchmarks, multiple LLMs, and diverse agent frame-
 016 works. Beyond performance, **HTC** delivers three essential advances: it provides
 017 *interpretability* by revealing the signals behind failure, enables *transferability* by
 018 applying across domains without retraining, and achieves *generalization* through a
 019 *General Agent Calibrator (GAC)* that achieves the best calibration (lowest ECE)
 020 on the out-of-domain GAIA benchmark. Together, these contributions establish a
 021 new process-centric paradigm for confidence calibration, [providing a framework](#)
 022 for diagnosing and enhancing the reliability of AI agents.



023 Figure 1: Overview of Holistic Trajectory Calibration (**HTC**). The framework first collects confidence
 024 signals along the agent’s trajectory, then derives rich process-level diagnostic features, which are used
 025 to train a simple yet interpretable calibrator. This process not only improves calibration accuracy but
 026 also yields the three pillars of reliable agentic AI: *interpretability*, *transferability*, and *generalization*,
 027 across diverse tasks and models.

028 1 INTRODUCTION

029 Large Language Models (LLMs) are rapidly evolving from static or retrieval-augmented text-
 030 generation tools into the core reasoning engines of complex, multi-step agentic systems (Xi et al.,
 031 2023; Wang et al., 2024a). These agents, which integrate sophisticated capabilities such as planning,

054 using tools, and handling memory (Yao et al., 2022; Schick et al., 2023), can autonomously interact
 055 with dynamic environments to solve complex problems. As these systems are increasingly deployed
 056 in high-stakes, safety-critical domains, their reliability has emerged as the most critical and yet
 057 unresolved challenge (Wang et al., 2024b; Liu et al., 2023). Ensuring that we can trust the outputs of
 058 these powerful but opaque systems is **critical** for their responsible adoption.

059 This shift from a static generator to a dynamic actor fundamentally alters the nature of the reliability
 060 challenge. First, uncertainty is no longer an isolated property of a single output, but a compounding
 061 factor that accumulates and propagates throughout a sequential trajectory (Ren et al., 2024; Kim et al.,
 062 2023). An early, low-confidence decision-such as erroneously selecting a tool, can "poison" the entire
 063 subsequent execution path, leading to an agent holding high confidence in a completely incorrect
 064 result. Second, agents introduce new, external sources of uncertainty through their interaction with
 065 tools and environments (Gao et al., 2024; Levy & Yih, 2024). API failures, noisy data returned by
 066 tools, or the misuse of a tool's functionality create new reliability bottlenecks independent of the
 067 model's internal knowledge. Finally, the multi-step nature of agentic processes makes failure modes
 068 more opaque. A final incorrect answer may not stem from the last reasoning step, but from a critical,
 069 masked breakdown that occurs at a specific intermediate step earlier in the trajectory (Fu et al., 2025).
 070 Agent calibration also faces a fundamental data scarcity challenge that means each agent trajectory
 071 represents an expensive execution involving LLM inference, tool interactions, and human evaluation
 072 for ground truth labels. This constraint shapes us toward sample-efficient and interpretable methods.

073 In light of these unique challenges, existing approaches to confidence estimation are insufficient.
 074 On one hand, traditional calibration techniques like Temperature Scaling (Guo et al., 2017) were
 075 designed for post-hoc correction of static, single-point classification predictions. Methodologically,
 076 they are incapable of processing sequential trajectory data and thus completely ignore the process-
 077 level information that could reveal the root cause of an agent's failure. On the other hand, while
 078 recent work has begun to explore more fine-grained confidence signals (Geng et al., 2024a), these
 079 efforts often rely on coarse aggregation methods like global averaging, which can mask local yet
 080 critical reasoning failures (Fu et al., 2025), or are limited to evaluating pure reasoning chains without
 081 external tool interaction (Wei et al., 2022), see more related work in Appendix A.1. Consequently, a
 082 significant gap exists in the current studies: *there is a lack of a systematic framework for effectively*
 083 *calibrating the confidence of an agent's final output by diagnosing its entire execution trajectory.*

084 In this work, we introduce the new problem of *Agentic Confidence Calibration (ACC)*: estimating the
 085 likelihood that an agent's trajectory will succeed by diagnosing its entire execution process rather
 086 than only its final output. As illustrated in Figure 1, this process-centric perspective raises three key
 087 challenges: uncertainty signals dispersed across multiple temporal scales, compounding noise from
 088 both model and environment, and limited availability of labeled data. To address these challenges, we
 089 propose **Holistic Trajectory Calibration (HTC)**, a novel framework that transforms raw confidence
 090 traces into a rich set of process-diagnostic features, encompassing cross-step dynamics, intra-step
 091 stability, positional indicators, and structural attributes. These features are then mapped through
 092 a simple, interpretable model to produce calibrated confidence estimates. Importantly, **HTC** is
 093 decoupled from any specific agent architecture, making it lightweight, transparent, and broadly
 094 applicable across diverse tasks and frameworks.

095 Our study demonstrates that **HTC** brings three **key** benefits:

- 096 • **Interpretability:** By grounding calibration in trajectory-level features such as early-step entropy,
 097 confidence gradients, and stability dynamics, **HTC** exposes the signals behind model confidence,
 098 enabling transparent diagnosis of failure modes and guiding principled agent design.
- 099 • **Transferability:** Once trained, an **HTC** calibrator can be seamlessly applied across tasks and
 100 domains without retraining, delivering consistent gains in both calibration and discrimination and
 101 reducing dependence on costly, task-specific tuning.
- 102 • **Generalization:** Pre-training a *General Agent Calibrator (GAC)* on diverse datasets yields a uni-
 103 versal reliability layer that achieves the best calibration (lowest ECE) on out-of-domain challenges
 104 such as GAIA, pointing toward a scalable foundation for trustworthy agentic AI.

105 Across eight benchmarks, multiple agent frameworks, and both closed- and open-source LLMs, we
 106 show that **HTC** reliably outperforms strong baselines. Beyond empirical gains, **HTC** establishes a
 107 process-centric paradigm for agentic confidence calibration, uniting interpretability, transferability,
 108 and generalization, three essential **components** for building reliable and trustworthy AI agents.

108

2 METHOD

109

2.1 HOLISTIC TRAJECTORY CALIBRATION: A NEW FORMULATION

110 An *agent* system is defined as a policy π that, at step t , maps the interaction history h_t to an action
 111 $a_t \in \mathcal{A} : a_t = \pi(h_t), h_t = (s_0, a_1, o_1, \dots, s_t)$. The environment Ω executes a_t from state s_t and
 112 returns an observation $o_t \in \mathcal{O}$ together with the next state via a (possibly stochastic) transition kernel
 113 $(o_t, s_{t+1}) \sim \delta(\cdot | s_t, a_t)$. This interaction produces an *execution trajectory*,
 114

$$115 \quad \mathcal{T} = (s_0, a_1, o_1, \mathcal{L}_1, s_1, a_2, o_2, \mathcal{L}_2, \dots, a_N, o_N, \mathcal{L}_N, s_N), \quad (1)$$

116 which records the complete problem-solving process up to termination at step N . For LLM-based
 117 agents, each action a_t (e.g., thinking, planning pr or tool call) is generated by an LLM \mathcal{M} . We denote
 118 by $\mathcal{L}_t = (\ell_{t,1}, \dots, \ell_{t,m_t})$ the sequence of token-level log-probabilities produced when generating a_t ,
 119 where m_t is the number of tokens in action a_t . By concatenating these sequences across all N steps,
 120 we obtain the *log-probability trajectory* after LLM execution:
 121

$$122 \quad \mathcal{L}_{\mathcal{T}} = \left((\ell_{1,1}, \dots, \ell_{1,m_1}), (\ell_{2,1}, \dots, \ell_{2,m_2}), \dots, (\ell_{N,1}, \dots, \ell_{N,m_N}) \right). \quad (2)$$

123 This trajectory captures the agent’s complete reasoning process, yet existing approaches typically
 124 assess confidence only from the final action: $\mathcal{C}_{\text{trad}} = \mathcal{H}(s_N, a_N)$. Agentic confidence calibration
 125 introduces three fundamental challenges that make it substantially harder than static calibration.
 126

127 **Problem Formulation: Holistic Trajectory Calibration (HTC)**

128 Given an agent’s execution trajectory \mathcal{T} with associated token log-probabilities $\mathcal{L}_{\mathcal{T}}$, learn a
 129 calibration function \mathcal{F}_{HTC} that maps the trajectory to a calibrated confidence score $\mathcal{C}_{\mathcal{T}} \in [0, 1]$:

$$130 \quad \mathcal{C}_{\mathcal{T}} = \mathcal{F}_{\text{HTC}}(\mathcal{T}(\mathcal{L}_{\mathcal{T}})) \quad \text{s.t.} \quad \mathbb{E}[y | \mathcal{F}_{\text{HTC}}(\mathcal{T}(\mathcal{L}_{\mathcal{T}}))] = c, \quad \forall c \in [0, 1] \quad (3)$$

131 where $y \in \{0, 1\}$ indicates task success or matches ground truth solution.
 132

133 **Challenge 1: Compounding Uncertainty.** Agentic trajectories accumulate and propagate uncertainty
 134 across multiple steps: early misjudgments may amplify downstream errors, while interactions with
 135 external tools introduce additional stochasticity, resulting in confidently incorrect final outputs.

136 **Challenge 2: Multi-Source Uncertainty.** Uncertainty in agentic reasoning is heterogeneous: it arises
 137 from token-level fluctuations within a step, from cross-step dynamics describing how confidence
 138 evolves. Signals are dispersed across multiple scales and cannot be reduced to a single summary.

139 **Challenge 3: Data Scarcity and Uncertainty.** Collecting agent trajectories is time-consuming and
 140 costly, which limits available datasets to relatively small scales. Moreover, the length of trajectories
 141 varies substantially with task complexity, introducing additional sources of data uncertainty.

142 To address the above challenges, we introduce a new paradigm **Holistic Trajectory Calibration (HTC)**,
 143 and formulate **HTC** as a supervised learning problem. Given a dataset of trajectories
 144 $\{\mathcal{T}_i(\mathcal{L}_{\mathcal{T}_i})\}_{i=1}^N$ with binary success labels $\{y_i\}_{i=1}^N$, we learn a calibration function \mathcal{F}_{HTC} by minimizing
 145 a proper scoring loss:

$$146 \quad \mathcal{F}_{\text{HTC}}^* = \arg \min_{\mathcal{F}} \frac{1}{N} \sum_{i=1}^N \ell(y_i, \mathcal{F}(\mathcal{T}_i(\mathcal{L}_{\mathcal{T}_i}))) + \lambda R(\mathcal{F}), \quad (4)$$

147 where $\ell(\cdot, \cdot)$ is a calibration-sensitive loss and $R(\mathcal{F})$ is a regularization term. The core question
 148 is how to design an effective representation $\phi(\mathcal{T}(\mathcal{L}_{\mathcal{T}}))$ that captures dispersed uncertainty signals,
 149 while supporting learning that is sample-efficient, interpretable, and generalizable across tasks.

150

2.2 THE IMPERATIVE FOR TRAJECTORY-LEVEL FEATURES

151 Given the challenges identified above, we argue that *trajectory-level features* are indispensable for
 152 effective **HTC** framework. Naïve alternatives such as relying only on final-step log-probabilities or
 153 averaging token confidences fail to capture the dispersed, multi-scale, and noise-sensitive nature of
 154 agentic uncertainty. In contrast, holistic trajectory-level features balance expressivity and practicality:

162 they capture diverse uncertainty signals while remaining tractable in small-data regimes. Unlike
 163 end-to-end neural encoders (e.g., RNN, LSTM, Transformer) which require large datasets and yield
 164 opaque representations ill-suited for calibration, feature-based representations enable sample-efficient
 165 learning and provide direct diagnostic insights.

166 **Design principles.** Our trajectory-level representation $\mathbf{x} = \phi(\mathcal{T})$ is guided by four principles: (i)
 167 *Universality*: features should be agnostic to task, model, and agent framework; (ii) *Informativeness*:
 168 features must encode signals causally linked to success and failure; (iii) *Parsimony*: the set should
 169 remain compact for small-sample calibration; and (iv) *Interpretability*: each feature should provide
 170 diagnostic value for analyzing uncertainty. [Based on these principles](#), we organize trajectory-level
 171 features into four complementary families:

- 173 • **Cross-Step Dynamics**: capture how confidence evolves across steps, detecting accumulation,
 174 reversals, or abrupt shifts that reflect compounding uncertainty.
- 175 • **Intra-Step Stability**: measure within-step volatility and distributional shape of token-level log-
 176 probabilities, indicating unstable or collapsed behaviors.
- 177 • **Positional Indicator**: critical early and late time-points where initialization quality and terminal
 178 consolidation often determine success and dominate outcomes.
- 179 • **Structure Attribute**: summarize macroscopic trajectory attributes (e.g., step count, token-length
 180 patterns) that proxy for task complexity and agent efficiency.

182 **From trajectory to features.** Concretely, we apply a small set of statistical operators (mean/variance,
 183 min/max, entropy, skewness, finite differences) along two axes, *within a step* and *across steps*, to
 184 the log-probability trajectory. This yields a compact vector $\mathbf{x} \in \mathbb{R}^{48}$ that preserves essential uncer-
 185 tainty signals while supporting efficient and interpretable calibration. [We constructed a systematic](#)
 186 [Taxonomy of Uncertainty](#) covering four critical axes. [This resulting 48-dimensional space balances](#)
 187 [comprehensiveness with the need to prevent overfitting in small-sample regimes](#) (see ablation in
 188 [Appendix A.4](#)). A full taxonomy with formal definitions is provided in Appendix A.5.1.

190 2.3 INTERPRETABLE CALIBRATION MODEL

192 Given the designed feature $\mathbf{x} = \phi(\mathcal{T}) \in \mathbb{R}^{48}$, we adopt a simple yet *interpretable light calibration*
 193 *model*. This choice is motivated by three considerations specific to agentic confidence calibration: (i)
 194 *small-sample robustness*: agent trajectory datasets are inherently small so linear models are less prone
 195 to overfitting than neural alternatives with thousands of parameters; (ii) *interpretable diagnostics*:
 196 linear weights provide direct insights into which uncertainty signals matter for different tasks, which
 197 is crucial for understanding agent failure modes; and (iii) *transferability and generalization* as low-
 198 capacity models generalize more reliably across domains with heterogeneous trajectory distributions.
 199 Formally, the calibration function maps features to a calibrated confidence score:

$$200 \quad \mathcal{C}_{\mathcal{T}} = \mathcal{F}_{\text{HTC}}(\mathbf{x}) = \sigma(\mathbf{w}^\top \mathbf{x} + b), \quad \mathbf{w} \in \mathbb{R}^{48}, \quad b \in \mathbb{R} \quad (5)$$

201 where \mathbf{w} and b are learned parameters. We instantiate the model under two complementary regular-
 202 ization regimes:

- 204 • **HTC-Full**: Retains all features while stabilizing estimates under collinearity through ridge
 205 regularization, $\mathcal{R}_{L2}(\mathbf{w}) = \lambda \|\mathbf{w}\|_2^2$. This preserves the full diagnostic surface across all features.
- 206 • **HTC-Reduced**: Encourages sparsity via lasso regularization, $\mathcal{R}_{L1}(\mathbf{w}) = \lambda \|\mathbf{w}\|_1$, automatically
 207 selecting a compact subset $\mathcal{S} = \{j : w_j \neq 0\}$. This denoises spurious features and often improves
 208 calibration in small-data regimes.

210 **Theoretical Motivation.** From a theoretical standpoint, trajectory-level calibration is strictly more
 211 informative than last-step confidence: conditioning on richer trajectory features can only reduce
 212 Bayes risk under proper scoring rules. In addition, a sparse ℓ_1 -regularized logistic calibrator admits
 213 favorable small-sample generalization bounds, explaining its stability in data-scarce regimes. A
 214 simple chain-of-subgoals model further clarifies why last-step confidence can be systematically
 215 optimistic, and the same diagnostics applied to prefixes establish a principled path toward online
 reliability. Formal statements and complete proofs are provided in Appendix A.6.

216 **Efficiency and Deployment.** The linear calibrator is computationally lightweight. Feature extraction
 217 scales linearly with trajectory length and requires only simple aggregation operators, while model
 218 training and inference are near-instantaneous. This efficiency makes **HTC** practical for real-time
 219 deployment and rapid adaptation to new domains, see more discussion in Appendix A.7 and A.8.
 220

221 **3 EXPERIMENTS**
 222

223 **3.1 EXPERIMENTAL SETUP**
 224

225 **Datasets and Benchmarks.** To comprehensively evaluate the effectiveness and generality of our
 226 **HTC** framework, we select 8 representative public benchmarks. These datasets are categorized into
 227 three groups to test distinct agent capabilities: (1) **Knowledge-intensive QA** (SimpleQA (Bordes
 228 et al., 2015), HotpotQA (Yang et al., 2018), StrategyQA (Geva et al., 2021)) for factual retrieval and
 229 multi-hop reasoning; (2) **Complex Reasoning** (MATH500 (Hendrycks et al., 2021), GPQA (Rein
 230 et al., 2023), MMLU-Pro (Zhang et al., 2024), HLE (Zhang et al., 2025c)) for formal logic and deep
 231 domain knowledge; and (3) **Frontier Agentic Tasks** (GAIA (Mialon et al., 2023)) for planning and
 232 tool-use in difficult, open-ended scenarios. Detailed descriptions and references for all datasets are
 233 provided in Appendix A.2.1.

234 **Models & Agent Frameworks.** Our experiments are conducted using **smolagents** (Roucher
 235 et al., 2025), a lightweight and research-friendly framework, leveraging its **CodeAct** paradigm
 236 where the agent generates executable Python code for tool use. We evaluate on a diverse set
 237 of models that provide LOGPROBS access. Our closed-source models are **GPT-4.1** and **GPT-4.0**
 238 (OpenAI et al., 2024). Our open-source suite includes **GPT-OSS-120B & 20B** (OpenAI, 2025),
 239 **Deepseek-v3.1** (DeepSeek-AI, 2024), and **Qwen3-235B** (Team, 2024). To ensure our findings are not
 240 specific to a single framework, we conduct a generalization study using the state-of-the-art **OAgents**
 241 framework (OPPO-PersonalAI, 2024) in our ablation analysis. Further details on all frameworks and
 242 models are available in Appendix A.2.3.

243 **Baselines.** We compare **HTC** against two categories of baselines to address different evaluation
 244 dimensions: *Inference-based Baselines*: (1) **Verbalized Confidence** (Tian et al., 2023a): agents
 245 directly output confidence scores; (2) **Last-Step Token Confidence** (LastStep-TP): average log-
 246 probabilities from final generation step; (3) **Global-Trace Token Confidence** (GlobalTrace-TP):
 247 average log-probabilities across all steps; (4) **Temperature Scaling** (Guo et al., 2017) applied to
 248 above methods (see details in Appendix A.2.4). *Learning-based Baselines*: (1) **LSTM Encoder**:
 249 processes raw log-probability sequences with final hidden state classification; (2) **Transformer**:
 250 attention-based sequence encoder. There are another three nonlinear methods based on our extracted
 251 features: (3) **Neural Network**, (4) **XGBoost** and (5) **Gaussian Process** (Rasmussen & Williams,
 252 2006). Detailed definitions and implementation specifics are provided in Appendix A.2.5.

253 **Evaluation Metrics and Implementation.** We evaluate calibration performance using three standard
 254 metrics: **Expected Calibration Error (ECE)** (Guo et al., 2017), which measures the accuracy
 255 of confidence scores; the **Brier Score (BS)** (Brier, 1950), a proper scoring rule assessing both
 256 calibration and discrimination; and **AUROC**, which measures the model’s ability to distinguish
 257 between successful and failed trajectories. It is important to distinguish the calibration method from
 258 the evaluation metric. **HTC** is the proposed method (predictor) that outputs confidence scores, while
 259 metrics like ECE and Brier Score serve as the ground-truth standards for assessing the quality of those
 260 scores. Therefore, a method achieving consistently lower ECE and BS is objectively better aligned
 261 with the true empirical accuracy. To ensure the validity of our ground truth labels, we employed a
 262 Gemini-2.5-Pro based judge, which we verified on a stratified subset to achieve a 90-95% agreement
 263 rate with human experts. All experiments are conducted using a cross-validation scheme to ensure
 264 robust results. A detailed description is provided in Appendix A.2.2. The implementation details and
 265 hyperparameter setting are provided in Appendix A.2.6.

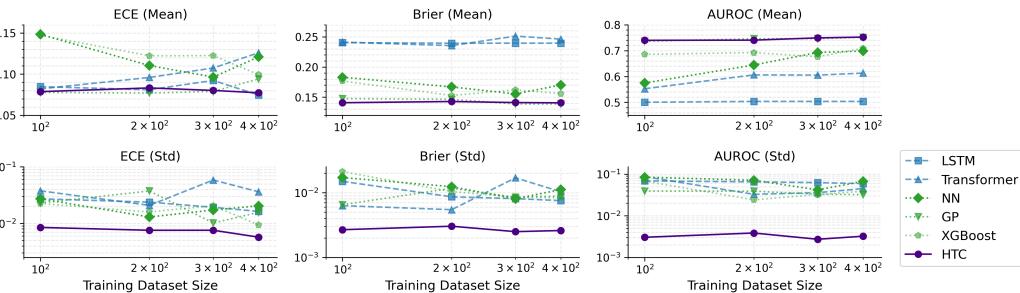
266 **3.2 MAIN RESULTS: CALIBRATION PERFORMANCE OF **HTC****
 267

268 Table 1 summarizes results on three representative datasets. Note that we evaluate the quality of our
 269 calibration method (**HTC**) using standard metrics (ECE, Brier Score); thus, lower values on these
 270 metrics directly indicate superior alignment between predicted confidence and actual performance.
 271 Across all metrics, both **HTC** variants substantially outperform inference-based baselines, with

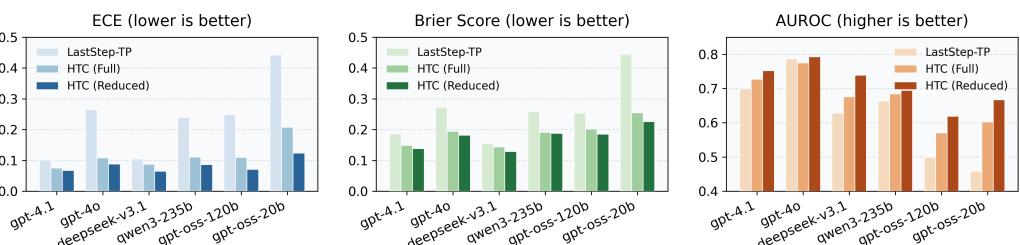
270 Table 1: Main results comparing **HTC** against baselines on representative datasets. Top 2 results in
 271 ECE, BS and AUROC are marked as bold and see full results in Table 4 and 5 in Appendix A.3.1.

| Method | SimpleQA | | | GPQA | | | HLE | | |
|-----------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| | ECE ↓ | BS ↓ | AUROC ↑ | ECE ↓ | BS ↓ | AUROC ↑ | ECE ↓ | BS ↓ | AUROC ↑ |
| Verbalized Conf | 0.121 | 0.196 | 0.655 | 0.454 | 0.523 | 0.593 | 0.656 | 0.531 | 0.614 |
| LastStep-TP | 0.101 | 0.186 | 0.699 | 0.424 | 0.413 | 0.614 | 0.686 | 0.561 | 0.604 |
| LastStep-TP + Temp | 0.071 | 0.178 | 0.698 | 0.139 | 0.258 | 0.610 | 0.436 | 0.278 | 0.628 |
| GlobalTrace-TP | 0.110 | 0.193 | 0.692 | 0.414 | 0.402 | 0.649 | 0.685 | 0.560 | 0.551 |
| GlobalTrace-TP + Temp | 0.077 | 0.181 | 0.691 | 0.136 | 0.257 | 0.643 | 0.433 | 0.277 | 0.570 |
| HTC-Full | 0.075 | 0.150 | 0.727 | 0.124 | 0.219 | 0.704 | 0.072 | 0.098 | 0.617 |
| HTC-Reduced | 0.068 | 0.140 | 0.752 | 0.102 | 0.213 | 0.706 | 0.031 | 0.090 | 0.644 |

281
 282 especially large gains in Brier Score and AUROC. On the most challenging tasks, **HTC-Reduced**
 283 achieves the strongest calibration, e.g., **ECE of 0.031** and **Brier Score of 0.09** on HLE, highlighting
 284 the benefit of sparsity in isolating universal uncertainty signals. We present a series of radar charts in
 285 Figure 1 to provide a comprehensive overview of our framework’s performance across all eight diverse
 286 datasets. We also compared against five learning-based baselines, including LSTM, Transformer,
 287 Neural Networks (NN), Gaussian Process (GP) and XGboost methods on SimpleQA with detailed
 288 learning curves shown in Figure 2. **HTC** consistently attains lower mean error and dramatically
 289 smaller variance across dataset sizes (100–400), demonstrating robustness in small-data regimes
 290 where neural baselines overfit or fluctuate heavily (see full results in Appendix A.3.1).



301 Figure 2: Learning Curve Comparison: **HTC** vs. Learning-Based Baselines on SimpleQA dataset,
 302 showing **HTC** consistently outperforms and exhibits much lower variance under small-data regimes.



312 Figure 3: The Impact of Base LLM on Calibration Performance on the SimpleQA dataset.

313
 314 **Effect of LLM Choice.** To validate **HTC** is model-agnostic, we evaluated its performance across
 315 six different LLMs on SimpleQA. The results in Figure 3, reveal two key findings. First, our
 316 **HTC** framework delivers consistent and substantial improvements for every model tested, from the
 317 high-performing **GPT-4.1** to other powerful open-source alternatives like **GPT-OSS-20B**, which
 318 exhibits particularly poor initial calibration. Second, the results highlight that different LLMs possess
 319 distinct baseline calibration profiles. For instance, while **GPT-4.0** demonstrates the strongest raw
 320 discriminative ability (highest baseline AUROC), its calibration (ECE) is notably poorer than that of
 321 **GPT-4.1**. Our **HTC** framework effectively addresses these unique characteristics, not only elevating
 322 the overall performance but also correcting the specific deficiencies of each model.

323
 324 **Effect of Agent Architectures.** We investigated whether **HTC**’s effectiveness is tied to a specific agent
 325 architectures. We compared its performance on the lightweight **smolagents** versus the highly-

324 optimized **OA**gents architectures, using **GPT-4.1** on GPQA (Figure 7 in Appendix A.3.1). **HTC**
 325 provides significant gains on both architectures, confirming that our approach is architecture-agnostic
 326 and can serve as a plug-and-play module to enhance the reliability of various agentic systems.
 327

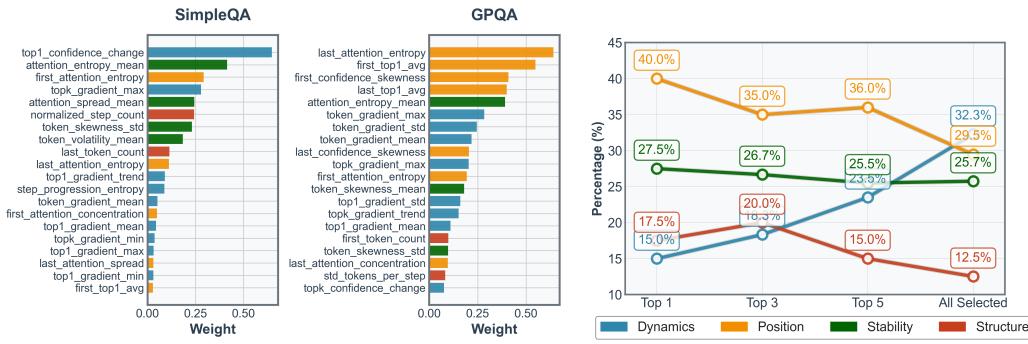
328 3.3 DIAGNOSTIC ANALYSIS: WHY **HTC** WORKS

330 3.3.1 FEATURE IMPORTANCE AND INTERPRETABILITY

331 A key advantage of **HTC** framework is its interpretability. By analyzing the features weights of our
 332 regularized linear model, we can move beyond *if* it works to *why* it works, gaining deep insights into
 333 the nature of agentic failure.
 334

335 **Uncertainty Signals are Task-Dependent.** Our first major finding is that the most predictive signals
 336 of failure are highly dependent on the cognitive demands of the task. Figure 8 in Appendix A.3.2
 337 displays the important features selected by our model for each of the eight datasets. It is visually
 338 apparent that there is no single universally dominant feature; the feature set and their relative
 339 importances shift based on the task’s nature. To illustrate this task-dependency more clearly, we
 340 compare the feature importance distributions for two representative tasks in Figure 4 (left):
 341

- 342 • For **SimpleQA**, a task that typically involves a "search-then-synthesize" pattern, the most predictive
 343 features are diverse and balanced across *Dynamics*, *Stability*, and *Position*. This suggests that failure
 344 can occur at multiple distinct stages: a poor transition between search and synthesis (*Dynamics*),
 345 an unstable generation process (*Stability*), or a weak final conclusion (*Position*). The model learns
 346 to monitor a broad array of signals to detect these varied failure modes.
- 347 • For **GPQA**, a task involving long and complex reasoning chains, the feature importance is
 348 heavily concentrated in the *Position* category. This indicates that for such difficult tasks, the
 349 agent’s cognitive state at the very beginning (*first_step*) and, more critically, at the very end
 350 (*last_step*) serves as the most potent summary of the entire arduous process. A hesitant or
 351 unstable conclusion after a long chain of reasoning is a particularly strong signal of failure.



362 Figure 4: (Left) Distribution of feature importance across different task domains. (Right) Frequency
 363 of feature category across different levels, including Top 1, Top 3, Top 5 and all selected features.
 364

365 **A General Hierarchy of Diagnostic Signals.** We analyze the statistical distribution of feature
 366 categories across all datasets, as shown in Figure 4 (right). The **Top-1 most important feature** across
 367 all datasets is most frequently a **Positional** feature. This aligns with intuition: a flawed start or a
 368 shaky conclusion is the most immediate and powerful "first alert" signal of a failing trajectory. As we
 369 expand our view to the **Top-3 and Top-5 most important features**, **Stability** and **Dynamics** features
 370 become increasingly prominent. This reveals that a comprehensive diagnosis requires looking beyond
 371 the start/end points and into the micro- and macro-level stability of the reasoning *process* itself. When
 372 considering **all selected features**, **Dynamics** emerges as the most frequently selected category overall.
 373 This suggests that while not always the single strongest signal, the step-to-step evolution of an agent’s
 374 confidence is a pervasive and indispensable component of a full reliability assessment. This analysis,
 375 with full feature selection frequency detailed in Table 6 in Appendix A.3.2, allows us to distill a key
 376 insight: effective agent calibration requires a hierarchical diagnostic approach.
 377

To validate the complementarity of our feature design, we conducted an ablation study across 15 configurations spanning single categories, pairwise, three-way, and the full set over the four feature categories (see Appendix A.4). We find that no single family suffices while multi-category

378 combinations substantially improve performance. There is no “marginal” category—the effectiveness
 379 of **HTC** derives precisely from integrating diverse, process-diagnostic signals.
 380

381 Takeaway 1: Interpretable Feature Importance 382

383 Two insights into agent reliability. First, there is no single universally dominant feature;
 384 the most predictive signals of failure are highly **task-dependent**, shifting from **Positional**
 385 indicators in complex reasoning tasks to a more diverse signal set in multi-step QA. Second,
 386 despite this diversity, a general diagnostic hierarchy emerges across all tasks: **Positional**
 387 features (the start and end) serve as the strongest primary signals of failure, while **Stability**
 388 and **Dynamics** features are essential for a complete diagnosis of the underlying process.
 389

390 3.3.2 CROSS-DOMAIN TRANSFERABILITY

391 A central question is whether **HTC**’s process-diagnostic features capture generalizable uncertainty
 392 signals rather than dataset-specific artifacts. To evaluate this, we pre-train a calibrator on one source
 393 dataset and apply it *without further training* to multiple target datasets.
 394

395 **Knowledge Domain Transfer: From Knowledge to Reasoning.** We first evaluate transferability
 396 within the knowledge-intensive domain by training a calibrator on SimpleQA. As shown in Table 2, the
 397 transferred model performs remarkably well on other QA tasks: on HotpotQA, it even **outperforms**
 398 **direct training** across all metrics, with similar gains on StrategyQA. Figure 5 explains this effect—the
 399 feature distributions of SimpleQA and HotpotQA are closely aligned, indicating shared uncertainty
 400 patterns. By contrast, transfer to the out-of-domain GPQA is weaker, consistent with its clear
 401 separation in feature space. These results suggest **HTC** can capture a robust “uncertainty patterns”
 402 that generalizes across related tasks while revealing the boundaries of cross-domain transfer.
 403

404 **Reasoning Domain Transfer: The Challenge of Distribution Shift.** We next examine transfer from
 405 MMLU-Pro. As shown in Table 2, transfer to other reasoning tasks (MATH500, HLE) underperforms
 406 direct training, despite their proximity in feature space (Figure 5). We attribute this to a **distribution**
 407 **shift in reasoning patterns**: MMLU-Pro induces multiple-choice reasoning, producing a different
 408 **distribution characteristics** than the open-ended generation in MATH500 or complex planning in
 409 HLE. Interestingly, transfer to StrategyQA is strong, despite being cross-domain. The key factor
 410 appears to be shared answer format (binary/short-form), suggesting that **output structure** can drive
 411 transferability as much as task category. This highlights that **HTC** features capture not only what the
 412 agent reasons about, but also how it organizes its final decision.
 413

414 Table 2: Cross-domain transfer performance. A calibrator is trained on a single source dataset,
 415 evaluated on multiple target datasets, comparing against a model trained directly on the target dataset.
 416

| Source: SimpleQA (Knowledge) | HotpotQA (ID) | | | StrategyQA (ID) | | | GPQA (OOD) | | |
|---------------------------------|---------------|---------------|--------------|-----------------|---------------|--------------|--------------|---------------|--------------|
| | ECE ↓ | Brier Score ↓ | AUROC ↑ | ECE ↓ | Brier Score ↓ | AUROC ↑ | ECE ↓ | Brier Score ↓ | AUROC ↑ |
| DIRECTTRAIN (full) | 0.116 | 0.193 | 0.714 | 0.079 | 0.141 | 0.670 | 0.124 | 0.219 | 0.704 |
| DIRECTTRAIN (reduced) | 0.082 | 0.183 | 0.729 | 0.055 | 0.136 | 0.665 | 0.102 | 0.213 | 0.706 |
| Transfer (full) | 0.113 | 0.194 | 0.719 | 0.099 | 0.148 | 0.657 | 0.435 | 0.446 | 0.587 |
| Transfer (reduced) | 0.070 | 0.183 | 0.732 | 0.064 | 0.135 | 0.681 | 0.304 | 0.330 | 0.629 |

| Source: MMLU-Pro (Reasoning) | MATH500 (ID) | | | HLE (ID) | | | StrategyQA (OOD) | | |
|---------------------------------|--------------|---------------|--------------|--------------|---------------|--------------|------------------|---------------|--------------|
| | ECE ↓ | Brier Score ↓ | AUROC ↑ | ECE ↓ | Brier Score ↓ | AUROC ↑ | ECE ↓ | Brier Score ↓ | AUROC ↑ |
| DIRECTTRAIN (full) | 0.060 | 0.077 | 0.788 | 0.072 | 0.098 | 0.617 | 0.079 | 0.141 | 0.670 |
| DIRECTTRAIN (reduced) | 0.048 | 0.070 | 0.816 | 0.031 | 0.090 | 0.644 | 0.055 | 0.136 | 0.665 |
| Transfer (full) | 0.081 | 0.092 | 0.782 | 0.457 | 0.329 | 0.620 | 0.056 | 0.134 | 0.682 |
| Transfer (reduced) | 0.081 | 0.083 | 0.792 | 0.504 | 0.349 | 0.645 | 0.028 | 0.131 | 0.689 |

426 Takeaway 2: Domain Transferability and Generalization 427

428 Our findings confirm that **HTC** can learn transferable signals of uncertainty. This transfer is
 429 most effective between tasks with similar cognitive processes while revealing the boundaries
 430 of cross-domain transfer. While a universal, “one-size-fits-all” calibrator faces challenges
 431 when transferring across fundamentally different cognitive paradigms.
 425

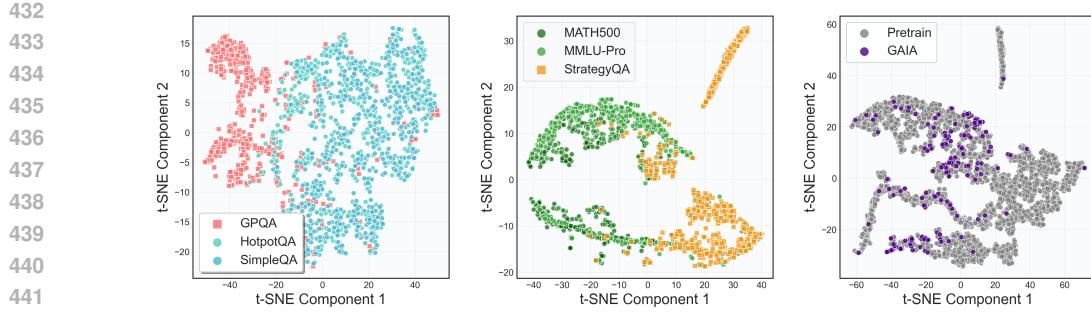


Figure 5: Low-dimensional t-SNE visualization of the feature spaces for different datasets.

3.4 GENERALIZATION: THE GENERAL AGENT CALIBRATOR

Our analysis in Section 3.3.2 has shown that while uncertainty signals are task-dependent, they share underlying patterns. This motivates our final and most ambitious experiment: *can we train a single, general agent calibrator (**GAC**) on a diverse corpus of tasks and have it successfully generalize to a complex, completely held-out agentic benchmark?* To test this, we pooled all seven datasets (SimpleQA, HotpotQA, StrategyQA, GPQA, MATH500, HLE, MMLU-Pro) for pre-training and held out GAIA as a challenging out-of-domain target. As visualized in Figure 5, GAIA lies dispersed across and beyond the pre-training feature space, making it an ideal stress test for generalization. We trained two versions of **GAC** (full vs. reduced features) on the combined corpus and evaluated them directly on GAIA, with results presented in Table 3 and Figure 6. Note that while **HTC** refers to our proposed methodological framework, **GAC** refers to the specific pre-trained model artifact released for zero-shot generalization.

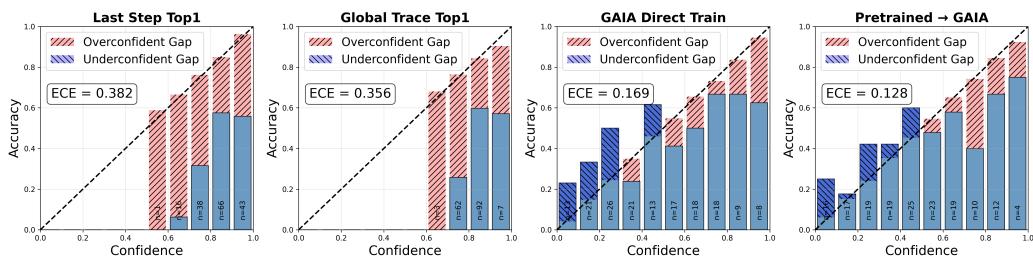


Figure 6: Reliability Diagrams for different calibration methods on the GAIA validation set.

Our results are highly encouraging. As shown in Table 3, pretraining **GAC** delivers the strongest calibration results on GAIA. Pretrained **GAC-Reduced** achieves the best ECE at **0.118**, with Pretrained **GAC-Full** close behind at **0.128**, both clearly surpassing DIRECTTRAIN (full: 0.169, reduced: 0.142) as well as all domain-transfer baselines. While DIRECTTRAIN (reduced) obtains the lowest Brier Score (0.233) and highest AUROC (0.686), **GAC-Reduced** remains highly competitive (0.245 BS; 0.647 AUROC) and, crucially, retains a substantially broader feature base (29.6 vs. 4.8 on average). These findings demonstrate that pretraining enables **GAC** to capture a transferable “uncertainty grammar” that prioritizes reliable calibration without resorting to extreme dataset-specific sparsification. Overall, achieving the best ECE with a pretrained calibrator is a highly promising result, **highlighting** its potential as a universal reliability layer for AI agents.

Takeaway 3: The General Agent Calibrator

Our experiments highlight the strong promise of a **pretrained, general-purpose agent calibrator**. By training on a diverse mix of domains, the calibrator achieves the best calibration (lowest ECE) on challenging out-of-domain tasks such as GAIA. This demonstrates that pretraining captures a transferable “uncertainty grammar” that generalizes beyond any single dataset. As a result, our approach offers a robust, plug-and-play reliability layer that can serve as a powerful foundation for future agentic systems.

486
487 Table 3: Performance of the **GAC** on GAIA Validation Set (top 2 are marked as bold)
488
489
490
491
492
493
494
495

| Method | ECE \downarrow | Brier Score \downarrow | AUROC \uparrow | # of Features |
|-------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|----------------|
| LastStep-TP | 0.382 | 0.375 | 0.607 | 1 |
| Knowledge Domain Transfer | 0.255 ± 0.010 | 0.273 ± 0.009 | 0.620 ± 0.012 | 48 |
| Reasoning Domain Transfer | 0.258 ± 0.010 | 0.268 ± 0.008 | 0.619 ± 0.020 | 48 |
| DIRECTTRAIN (full) | 0.169 ± 0.011 | 0.265 ± 0.009 | 0.620 ± 0.016 | 48 |
| DIRECTTRAIN (reduced) | 0.142 ± 0.010 | 0.233 ± 0.003 | 0.686 ± 0.013 | 4.8 ± 2.0 |
| Pretrained GAC-Full | 0.128 ± 0.001 | 0.250 ± 0.001 | 0.636 ± 0.001 | 48 |
| Pretrained GAC-Reduced | 0.118 ± 0.006 | 0.245 ± 0.002 | 0.647 ± 0.005 | 29.6 ± 3.9 |

496
497 4 CONCLUSION
498
499

500 We introduced Holistic Trajectory Calibration (**HTC**), a feature-based and interpretable framework
 501 for agentic confidence calibration. Our work addresses compounding uncertainty, heterogeneous
 502 signals, and data scarcity, yielding three key takeaways: (1) calibration relies on a hierarchy of
 503 diagnostic signals; (2) **HTC** features capture a transferable “**uncertainty patterns**” enabling strong
 504 cross-task generalization while exposing limits under **distribution shift**; and (3) a pretrained General
 505 Agent Calibrator (**GAC**) achieves the best ECE (zero-shot) on unseen tasks like GAIA, providing a
 506 plug-and-play foundation. Future work will scale **GAC** pre-training and explore light task-specific
 507 fine-tuning to combine broad generalization with specialized accuracy.

508
509 5 ETHICAL STATEMENT
510

511 This research contributes to the development of safer and more reliable AI agents, which is critical
 512 for their deployment in high-stakes domains like healthcare and finance. By enabling agents to better
 513 “know what they don’t know,” our work can facilitate more effective human-AI collaboration and
 514 increase the transparency of agent decision-making. However, we also acknowledge potential risks.
 515 A highly effective calibrator could be misused to create a false sense of security in an agent that is still
 516 fundamentally flawed in ways not captured by our features. Like any technology that enhances AI
 517 capability, it has a dual-use potential and must be deployed with a comprehensive evaluation strategy
 518 that goes beyond calibration metrics alone.

519
520 6 REPRODUCIBILITY STATEMENT
521

522 We have taken extensive steps to ensure that our work is reproducible. All datasets used in our
 523 experiments are publicly available and are described in Section 3, with preprocessing details included
 524 in Appendix A.2.6. The proposed **HTC** framework is fully specified: Section 2 defines the core
 525 methodology, Appendix A.6 provides complete theoretical proofs with explicit assumptions, and
 526 Appendix A.5.2 gives a detailed description of all diagnostic features with both mathematical
 527 definitions and intuitive explanations. Our learning-based baselines are described in Appendix A.2.5,
 528 together with their architectures and hyperparameters. Evaluation metrics and cross-validation
 529 strategies are reported in Appendix A.2.2. Because our calibrator is a lightweight logistic model
 530 operating on engineered features, the entire system can be re-implemented with minimal effort. For
 531 transparency, we additionally release an anonymized code base in the supplementary material, which
 532 computes the feature map and reproduces the calibration experiments in the paper.

533
534 REFERENCES
535

536 Antoine Bordes, Sumit Chopra, and Jason Weston. Large-scale simple question answering with
 537 memory networks. *arXiv preprint arXiv:1506.02075*, 2015.

538 Glenn W Brier. Verification of forecasts expressed in terms of probability. *Monthly weather review*,
 539 78(1):1–3, 1950.

540 Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the*
 541 *22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*,
 542 pp. 785–794. ACM, 2016.

543 Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Noveen Sachdeva, Inderjit
 544 Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier
 545 with advanced reasoning, multimodality, long context, and next generation agentic capabilities.
 546 *arXiv preprint arXiv:2507.06261*, 2025.

547 DeepSeek-AI. Deepseek llm: Scaling open-source language models with committee-based mixture
 548 of experts. *arXiv preprint arXiv:2405.04434*, 2024.

549 Yichao Fu, Xuewei Wang, Yuandong Tian, and Jiawei Zhao. Deep think with confidence. *arXiv*
 550 *preprint arXiv:2508.15260*, 2025.

551 Yunfan Gao, Kai Shu, and Hong Yu. Uncertainty calibration for tool-using language agents. In
 552 *Findings of the Association for Computational Linguistics: EMNLP 2024*. Association for Compu-
 553 tational Linguistics, 2024.

554 Jiahui Geng, Fengyu Cai, Yuxia Wang, Heinz Koepll, Preslav Nakov, and Iryna Gurevych. A survey
 555 of confidence estimation and calibration in large language models. In *Proceedings of the 2024*
 556 *Conference of the North American Chapter of the Association for Computational Linguistics: Human*
 557 *Language Technologies (Volume 1: Long Papers)*, pp. 6577–6595, Mexico City, Mexico,
 558 2024a. Association for Computational Linguistics.

559 Jiahui Geng, Fengyu Cai, Yuxia Wang, Heinz Koepll, Preslav Nakov, and Iryna Gurevych. A
 560 survey of confidence estimation and calibration in large language models, 2024b. URL <https://arxiv.org/abs/2311.08298>.

561 Mor Geva, Daniel Khashabi, Elad Segal, Jonathan Berant, and Ido Dagan. Did aristotle use a
 562 laptop? a question answering benchmark with implicit reasoning strategies. In *Transactions of the*
 563 *Association for Computational Linguistics*, volume 9, pp. 346–361, 2021.

564 Tobias Groot and Matias Valdenegro-Toro. Overconfidence is key: Verbalized uncertainty evaluation
 565 in large language and vision-language models, 2024. URL <https://arxiv.org/abs/2405.02917>.

566 Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural
 567 networks. *arXiv preprint arXiv:1706.04599*, 2017.

568 Jiuzhou Han, Wray Buntine, and Ehsan Shareghi. Towards uncertainty-aware language agent, 2024.
 569 URL <https://arxiv.org/abs/2401.14016>.

570 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,
 571 and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv*
 572 *preprint arXiv:2103.03874*, 2021.

573 Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):
 574 1735–1780, 1997.

575 Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas
 576 Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, Scott Johnston, Sheer El-Showk,
 577 Andy Jones, Nelson Elhage, Tristan Hume, Anna Chen, Yuntao Bai, Sam Bowman, Stanislav Fort,
 578 Deep Ganguli, Danny Hernandez, Josh Jacobson, Jackson Kernion, Shauna Kravec, Liane Lovitt,
 579 Kamal Ndousse, Catherine Olsson, Sam Ringer, Dario Amodei, Tom Brown, Jack Clark, Nicholas
 580 Joseph, Ben Mann, Sam McCandlish, Chris Olah, and Jared Kaplan. Language models (mostly)
 581 know what they know, 2022. URL <https://arxiv.org/abs/2207.05221>.

582 Jihwan Kim, Sung-Hyon Lee, Gwandaee Kim, Minsu Jang, and Min-Joong Kang. SAUP: Situation
 583 awareness uncertainty propagation on LLM agent. *arXiv preprint arXiv:2312.01033*, 2023.

584 Michael Kirchhof, Gjergji Kasneci, and Enkelejda Kasneci. Position: Uncertainty quantification
 585 needs reassessment for large-language model agents, 2025. URL <https://arxiv.org/abs/2505.22655>.

594 Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. Semantic uncertainty: Linguistic invariances for
 595 uncertainty estimation in natural language generation. *arXiv preprint arXiv:2302.09664*, 2023.
 596

597 Shachar Levy and Wen-tau Yih. Tools in the loop: Quantifying uncertainty of LLM question
 598 answering systems that use tools. *arXiv preprint arXiv:2405.16113*, 2024.

599 Stephanie Lin, Jacob Hilton, and Owain Evans. Teaching models to express their uncertainty in
 600 words, 2022. URL <https://arxiv.org/abs/2205.14334>.

601

602 Zhen Lin, Shubhendu Trivedi, and Jimeng Sun. Generating with confidence: Uncertainty quantifi-
 603 cation for black-box large language models, 2024. URL <https://arxiv.org/abs/2305.19187>.

604

605 Xiao Liu, Hao Yu, Hanchen Zhang, Yaran Xu, Zekun Xu, Mouxin Ruan, Jianwei Tan, Hang Xu, Tian-
 606 shu Fu, Nan Du, et al. AgentBench: Evaluating llms as agents. *arXiv preprint arXiv:2308.03688*,
 607 2023.

608

609 Grégoire Mialon, Roberto Dessì, Maria Lomeli, Christoforos Nalmpantis, Ranganath Pasunuru,
 610 Roberta Raileanu, Timo Schick, Jane Dwivedi-Yu, Asli Celikyilmaz, Edouard Grave, et al. Gaia: a
 611 benchmark for general ai assistants. *arXiv preprint arXiv:2311.12983*, 2023.

612 OpenAI. Gpt-oss-120b & gpt-oss-20b model card. Model Card, 2025. Accessed: 2025-09-18.

613

614 OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni
 615 Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor
 616 Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian,
 617 Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny
 618 Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks,
 619 Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea
 620 Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen,
 621 Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung,
 622 Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch,
 623 Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty
 624 Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte,
 625 Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel
 626 Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua
 627 Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike
 628 Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon
 629 Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne
 630 Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo
 631 Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar,
 632 Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik
 633 Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich,
 634 Aris Konstantinidis, Kyle Koscic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy
 635 Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie
 636 Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini,
 637 Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne,
 638 Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David
 639 Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie
 640 Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély,
 641 Ashvin Nair, Reiichiro Nakano, Rajeev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo
 642 Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano,
 643 Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng,
 644 Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto,
 645 Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power,
 646 Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis
 647 Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted
 Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel
 Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon
 Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky,
 Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie

648 Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng,
 649 Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun
 650 Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang,
 651 Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian
 652 Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren
 653 Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming
 654 Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao
 655 Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. Gpt-4 technical report, 2024. URL
 656 <https://arxiv.org/abs/2303.08774>.

657 OPPO-PersonalAI. Oagents: A large-scale framework for open-ended agents. <https://github.com/OPPO-PersonalAI/OAgents>, 2024.

658

659 Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian processes for machine learning*.
 660 MIT Press, 2006.

661

662 David Rein, Betty Li Hou, Asa Cooper Strickland, Jackson Petty, Richard Yuanzhe Pang, Julien
 663 Dirani, Julian Michael, and Samuel R. Bowman. GPQA: A graduate-level Google-proof Q&A
 664 benchmark. In *Proceedings of the First Conference on Language Modeling*, 2023.

665

666 Jiyan Ren, Zheng Zhang, Hanyu Zhao, Ruitian Zhang, Zhaofeng Liu, and Yang Liu. UProp:
 667 Investigating the uncertainty propagation of LLMs in multi-step agentic decision-making. *arXiv*
 668 preprint [arXiv:2406.17419](https://arxiv.org/abs/2406.17419), 2024.

669

670 Aymeric Roucher, Albert Villanova del Moral, Thomas Wolf, Leandro von Werra, and Erik Kaunis-
 671 mäki. ‘smolagents’: a smol library to build great agentic systems. <https://github.com/huggingface/smolagents>, 2025.

672

673 Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Tsvigun, Gautier Céron,
 674 Eric Saci, Yair Pinter, Jacek Piskorski, Giovanni De Toni, et al. Toolformer: Language models can
 675 teach themselves to use tools. *arXiv preprint arXiv:2302.04761*, 2023.

676

677 Qwen Team. Qwen3 technical report. *arXiv preprint*, 2024. Referenced via Together AI model card.

678

679 Hongxiang Tian, Ruitong Raugas, Zhaowei Zhang, Ziyi Wang, Ruosen Li, Zhiwei Liu, Zhaoyang
 680 Zhang, and Yang Liu. Just ask for calibration: A user-friendly framework for calibrating large
 681 language models. *arXiv preprint arXiv:2310.02534*, 2023a.

682

683 Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea
 684 Finn, and Christopher D. Manning. Just ask for calibration: Strategies for eliciting calibrated
 685 confidence scores from language models fine-tuned with human feedback, 2023b. URL <https://arxiv.org/abs/2305.14975>.

686

687 Yao-Hung Hubert Tsai, Walter Talbott, and Jian Zhang. Efficient non-parametric uncertainty quan-
 688 tification for black-box large language models and decision planning, 2024. URL <https://arxiv.org/abs/2402.00251>.

689

690 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz
 691 Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information
 692 Processing Systems (NeurIPS)*, volume 30, pp. 5998–6008, 2017.

693

694 Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai
 695 Tang, Zikang Wen, Fuxiang Chen, et al. A survey on large language model based autonomous
 696 agents. *arXiv preprint arXiv:2308.11432*, 2024a.

697

698 Ziyi Wang, Zhaowei Wang, Kewei Ren, Hang Yuan, Zhaoyang Zhang, Zhiwei Liu, Ruosen Li, Yifei
 699 Cheng, Zhuo Liu, Yuanchen Wang, et al. A survey on evaluation of large language models as
 700 agents. *arXiv preprint arXiv:2402.12338*, 2024b.

701

702 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, and Denny
 703 Zhou. Chain-of-thought prompting elicits reasoning in large language models. *arXiv preprint
 704 arXiv:2201.11903*, 2022.

702 Zhizheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yi Ding, Boyang Li, Weihuan Ruan, Yuan Zhang,
 703 Yang Zhou, and Chen Zhang. The rise and potential of large language model based agents: A
 704 survey. *arXiv preprint arXiv:2309.07864*, 2023.

705

706 Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W Cohen, Ruslan Salakhutdinov,
 707 and Christopher D Manning. Hotpotqa: A new dataset for diverse, explainable multi-hop question
 708 answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language
 709 Processing*, pp. 2369–2380, 2018.

710 Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Tushar, Quoc V Le, and Denny Zhou. ReAct:
 711 Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*, 2022.

712

713 Guibin Zhang, Junhao Wang, Junjie Chen, Wangchunshu Zhou, Kun Wang, and Shuicheng Yan.
 714 Agentracer: Who is inducing failure in the llm agentic systems?, 2025a. URL <https://arxiv.org/abs/2509.03312>.

715

716 Haotian Zhang, Yilun Li, Xinyu Chen, Heyang Liu, Lichao Gui, Yitao Wang, Tian-Shuo Liu, and
 717 Qing Zhang. MMLU-Pro: A more robust and challenging multi-task language understanding
 718 benchmark. *arXiv preprint arXiv:2405.08229*, 2024.

719

720 Shaokun Zhang, Ming Yin, Jieyu Zhang, Jiale Liu, Zhiguang Han, Jingyang Zhang, Beibin Li,
 721 Chi Wang, Huazheng Wang, Yiran Chen, and Qingyun Wu. Which agent causes task failures
 722 and when? on automated failure attribution of llm multi-agent systems, 2025b. URL <https://arxiv.org/abs/2505.00212>.

723

724 Yushi Zhang, Rakesh Kumar, et al. The human last exam: A new benchmark for evaluating ai systems.
 725 *arXiv preprint arXiv:2501.14249*, 2025c.

726

727 Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang,
 728 Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and
 729 chatbot arena. *Advances in neural information processing systems*, 36:46595–46623, 2023.

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756 **A APPENDIX**
757758 **APPENDIX CONTENTS**
759

| | |
|--|----|
| 760 A.1 Related Work | 16 |
| 761 A.2 Experimental Setup Details | 16 |
| 762 A.2.1 Detailed Dataset Descriptions | 16 |
| 763 A.2.2 Detailed Evaluation Metrics and Protocol | 17 |
| 764 A.2.3 Model and Agent Framework Details | 18 |
| 765 A.2.4 Inference-based Baselines | 18 |
| 766 A.2.5 Learning-based Baselines | 19 |
| 767 A.2.6 Implementation and Hyperparameter Details | 20 |
| 768 A.3 Additional Experimental Results | 20 |
| 769 A.3.1 Main Results | 20 |
| 770 A.3.2 Feature Importance Analysis | 21 |
| 771 A.3.3 Domain Transfer Analysis | 21 |
| 772 A.4 Ablation Study on Feature Categories | 27 |
| 773 A.5 Detailed Feature Description | 28 |
| 774 A.5.1 Feature Definitions | 28 |
| 775 A.5.2 Feature Map | 31 |
| 776 A.6 Theoretical Motivation and Analysis | 32 |
| 777 A.7 Efficiency and Cost Analysis | 34 |
| 778 A.8 Deployment and Practical Implications | 34 |
| 779 A.9 Qualitative Examples | 35 |
| 780 A.10 Future Work and Broader Impact | 37 |
| 781 A.11 LLM Usage | 38 |

792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809

810 A.1 RELATED WORK
811812 Our work on Agentic Confidence Calibration (ACC) is situated at the intersection of two rapidly
813 developing research areas: confidence calibration for LLMs and the nascent field of uncertainty
814 quantification (UQ) for LLM-based agents.
815816 **Confidence Calibration in LLMs** Confidence calibration aims to align a model’s predicted prob-
817 ability with the true likelihood of correctness. Classic methods such as Temperature Scaling (Guo
818 et al., 2017) are effective for standard classification tasks, but their direct application to the free-form,
819 generative outputs of LLMs is non-trivial (Kadavath et al., 2022; Lin et al., 2022). As a result, recent
820 work has explored calibration techniques specifically tailored for LLMs (Geng et al., 2024b). These
821 approaches typically leverage signals from the output distribution—e.g., prediction entropy (Kuhn
822 et al., 2023), top- k token probabilities (Lin et al., 2024), or verbalized confidence estimates (Tian
823 et al., 2023b; Groot & Valdenegro-Toro, 2024). Another notable direction, exemplified by *Deep*
824 *Think with Confidence* (Fu et al., 2025), highlights the importance of fine-grained, local signals
825 (such as the lowest-confidence step within a reasoning chain) over global averages for reasoning
826 calibration. Despite these advances, existing approaches remain focused on **static, single-turn, and**
827 **self-contained outputs**. They do not capture the compounding and multi-source uncertainties that
828 arise in the multi-step, interactive trajectories of AI agents (Kirchhof et al., 2025). Our work extends
829 this line of inquiry from isolated outputs to the entire agentic process.
830831 **Uncertainty in LLM Agents** The study of uncertainty in LLM agents is an emerging but critical
832 field (Kirchhof et al., 2025). A few pioneering works have begun to formalize the unique challenges
833 agents present (Han et al., 2024; Tsai et al., 2024). Frameworks like UProp (Ren et al., 2024)
834 and SAUP (Kim et al., 2023) were the first to model how uncertainty **propagates** through the
835 sequential steps of an agent’s trajectory. **While these works provide valuable analytical frameworks**
836 **for uncertainty propagation, they differ from HTC’s focus on supervised, data-driven calibration**
837 **and currently lack open-source implementations for direct comparison.** Concurrently, other research
838 has focused on quantifying the **external uncertainty** introduced by tool use, analyzing how API
839 failures or noisy tool outputs impact reliability (Gao et al., 2024; Levy & Yih, 2024). While these
840 studies laid the essential groundwork by identifying the core problems of propagation and external
841 interaction, they primarily focus on high-level modeling and do not delve into a systematic, feature-
842 based diagnosis of the underlying generation process (Zhang et al., 2025b;a). Our work builds upon
843 their problem formulation but takes a fundamentally different approach. Instead of modeling the
844 propagation dynamics directly, we propose a holistic framework that analyzes the rich, fine-grained
845 signals embedded within the full trajectory’s confidence to perform a comprehensive diagnostic
846 calibration. To our knowledge, this is the first work to systematically validate a process-diagnostic
847 feature set for the purpose of agentic confidence calibration.
848849 A.2 EXPERIMENTAL SETUP DETAILS
850851 A.2.1 DETAILED DATASET DESCRIPTIONS
852853 We use the following 8 benchmark datasets in our experiments to ensure a comprehensive and
854 multi-faceted evaluation of our proposed **HTC** framework.
855856

- **SimpleQA** (Bordes et al., 2015): A large-scale factual question-answering dataset. We randomly
857 sampled 500 instances from its test set to evaluate the agent’s basic knowledge retrieval capabilities.
- **HotpotQA** (Yang et al., 2018): A multi-hop question-answering dataset that requires reasoning over
858 multiple documents. We sampled 500 instances from its test set to assess calibration performance
859 on more complex knowledge-intensive tasks.
- **StrategyQA** (Geva et al., 2021): A question-answering benchmark requiring implicit reasoning
860 steps. We used 500 samples from its test set to evaluate the agent’s ability to handle problems that
861 require strategic thinking.
- **MATH500** (Hendrycks et al., 2021): A dataset of problems from high school mathematics
862 competitions. We used 500 samples from the MATH500 test set to focus on the reliability of
863 formal mathematical reasoning and computation.

- **GPQA** (Rein et al., 2023): A high-difficulty benchmark of graduate-level STEM questions that are challenging even for domain experts. We used all 448 samples from its MAIN split. To maximize the challenge, we converted it from a multiple-choice format to an open-ended generation task.
- **MMLU-Pro** (Zhang et al., 2024): A more challenging variant of MMLU that validates deep knowledge and reasoning through multi-turn dialogue and Chain-of-Thought. We used 500 samples from its test set.
- **HLE (Human Last Exam)** (Zhang et al., 2025c): An extremely difficult dataset comprising problems that are challenging for human experts, often requiring complex, multi-step reasoning. We used 500 samples to test agent reliability at the frontier of its capabilities.
- **GAIA** (Mialon et al., 2023): A benchmark designed for general AI assistants, with tasks that often require long-horizon planning, multi-tool coordination, and interaction with real-world documents and websites. We used the full 165 samples from its validation set as a final test of general autonomous capabilities.

Notably, to increase the challenge of GPQA, we removed its multiple-choice options, requiring the agent to generate answers directly. For datasets with more than 500 samples, we randomly selected a subset of 500; for those with fewer, we used the entire set (e.g., 448 samples for GPQA and the 165-sample validation set for GAIA). All samples were primarily sourced from the official test or validation splits of their respective datasets. Finally, we assign a binary success label ($y \in \{0, 1\}$) to each trajectory by evaluating the agent’s final answer against the ground truth.

A.2.2 DETAILED EVALUATION METRICS AND PROTOCOL

To rigorously evaluate the performance of our calibration framework, we focus on the following three standard metrics. Let c_i be the predicted confidence and $y_i \in \{0, 1\}$ be the ground-truth success label for trajectory i over N samples.

Calibration Metrics

- **Expected Calibration Error (ECE)** (Guo et al., 2017): ECE measures the difference between a model’s average confidence and its actual accuracy. To compute it, we partition the N predictions into M bins (B_m) based on their confidence scores. The ECE is the weighted average of the absolute difference between the accuracy and confidence of each bin:

$$\text{ECE} = \sum_{m=1}^M \frac{|B_m|}{N} |\text{acc}(B_m) - \text{conf}(B_m)| \quad (6)$$

where $\text{acc}(B_m)$ and $\text{conf}(B_m)$ are the accuracy and average confidence of the predictions in bin B_m , respectively. A lower ECE indicates better calibration.

- **Brier Score** (Brier, 1950): The Brier Score is a proper scoring rule that measures the mean squared error between predicted probabilities and actual outcomes. It simultaneously assesses both calibration and discrimination. It is defined as:

$$\text{Brier Score} = \frac{1}{N} \sum_{i=1}^N (c_i - y_i)^2 \quad (7)$$

A lower Brier Score indicates a better overall prediction quality.

- **AUROC**: The Area Under the Receiver Operating Characteristic curve measures the model’s ability to discriminate between successful ($y = 1$) and failed ($y = 0$) trajectories. It is threshold-independent and evaluates how well the confidence score can rank predictions. An AUROC of 1.0 represents a perfect classifier, while 0.5 represents a random guess.

Evaluation Protocol Many of the agent tasks in our benchmark suite, particularly on datasets like GAIA and HLE, result in complex, free-form text answers where simple string matching against the ground truth is insufficient for accurate evaluation. To address this, we adopt the widely-used **LLM-as-Judge** (Zheng et al., 2023) protocol for a robust and scalable evaluation. The process is as follows: (1) For each completed trajectory, we extract the agent’s final generated answer. (2) We

918 construct a prompt that includes the original question, the ground-truth answer from the dataset, and
 919 the agent’s answer. (3) This prompt is sent to a powerful, impartial judge model, **Gemini-2.5-Pro**
 920 ([Comanici et al., 2025](#)). (4) The judge model is instructed to provide a binary determination of
 921 correctness, outputting the final success label $y \in \{0, 1\}$ that we use for training and evaluating our
 922 calibrator. [We verified the reliability of the LLM judge on a stratified subset, observing a 90-95%
 923 agreement rate with human experts.](#)

925 A.2.3 MODEL AND AGENT FRAMEWORK DETAILS

- 927 • **sмолagents** ([Roucher et al., 2025](#)): Our primary framework for all main experiments is
 928 **sмолagents**, a minimalist agent framework designed for clarity and research agility. We
 929 specifically utilize its **CodeAct** functionality, where the agent’s actions (a_t) are formulated as
 930 Python code blocks. This paradigm offers high expressiveness, allowing the agent to perform
 931 complex computations and interact with tools (e.g., WEB SEARCH) through simple function calls
 932 within the code. Its lightweight nature ensures that the core reasoning and uncertainty signals come
 933 directly from the LLM, minimizing confounding variables from the framework itself.
- 934 • **OAgents** ([OPPO-PersonalAI, 2024](#)): For our framework generalization study, we use **OAgents**,
 935 a state-of-the-art, open-source agent framework known for its high performance on complex
 936 benchmarks like GAIA. **OAgents** incorporates more sophisticated planning and memory modules.
 937 By testing our **HTC** framework on **OAgents**, we can validate that our process-diagnostic features
 938 are fundamental signals of uncertainty, independent of the agent’s architectural complexity.

939 A.2.4 INFERENCE-BASED BASELINES

941 To rigorously evaluate our **HTC** framework, we compare it against five baseline methods, which are
 942 detailed below. For all methods based on log-probabilities, we use the average of the top-k/top-1
 943 token confidences, consistent with our framework’s feature extraction.

944 **Verbalized Confidence.** This is a standard black-box baseline that requires no access to internal
 945 model states. We append an instruction to the agent’s final prompt, asking it to state its confidence on
 946 a scale from 0% to 100%. An example instruction is: “*After providing your final answer, on a new
 947 line, state your confidence in its correctness as a single percentage, e.g., ‘Confidence: 85%’.*” We
 948 then parse the numerical value as the confidence score, c . This method is inspired by recent work on
 949 eliciting self-assessment from LLMs ([Tian et al., 2023a](#)).

950 **LastStep-TP Confidence.** This grey-box baseline represents the standard approach of relying on
 951 the final generation step for a confidence signal. Let $\mathcal{L}_N = (l_{N,1}, \dots, l_{N,M_N})$ be the sequence of
 952 token confidences from the final step (s_N) of the trajectory. The confidence score is the simple
 953 average:

$$954 c_{\text{last-step}} = \frac{1}{M_N} \sum_{j=1}^{M_N} l_{N,j} \quad (8)$$

955 **GlobalTrace-TP Confidence.** This baseline extends the ‘Last-Step’ approach by incorporating
 956 information from the entire trajectory, but in a naive way. It computes the global average of all token
 957 confidences across all N steps:

$$958 c_{\text{global-trace}} = \frac{1}{\sum_{i=1}^N M_i} \sum_{i=1}^N \sum_{j=1}^{M_i} l_{i,j} \quad (9)$$

959 This serves as a critical baseline to test whether the performance gain of our **HTC** framework comes
 960 from our sophisticated feature engineering or simply from using more data.

961 **LastStep-TP + Temperature Scaling.** To create a stronger, calibrated baseline, we apply Temperature
 962 Scaling ([Guo et al., 2017](#)) to the ‘Last-Step Confidence’ scores. A single temperature parameter
 963 T is optimized on a validation set to minimize Log Loss, and this scalar is then used to adjust the
 964 confidence scores.

972 **GlobalTrace-TP + Temperature Scaling.** Similarly, we apply Temperature Scaling to the ‘Global-
 973 Trace Confidence’ scores to provide another strong, calibrated baseline.
 974

975
 976
 977
 978 A.2.5 LEARNING-BASED BASELINES
 979

980 We further compare our framework against a set of supervised learning-based baselines. These
 981 methods fall into two groups: (i) neural representation learning methods, which directly operate
 982 on raw token-level confidence trajectories in an end-to-end fashion, and (ii) advanced nonlinear
 983 feature-based methods, which consume our engineered 48-dimensional trajectory feature space.
 984 Below we provide details for each baseline.
 985

986 **LSTM-based Confidence Predictor.** This model treats the confidence trajectory as a variable-length
 987 sequence and encodes it with a single-layer unidirectional LSTM (Hochreiter & Schmidhuber, 1997)
 988 of hidden size 64 and dropout 0.4. Each input step corresponds to the top-5 token log-probabilities,
 989 and the final hidden state is passed through a three-layer feed-forward classifier (64→32→32→2)
 990 with ReLU activations and dropout. The parameter count is on the order of 4k–6k, and the model
 991 is trained with Adam (learning rate 0.001), early stopping, and 5-fold cross-validation. The LSTM
 992 can capture temporal dependencies and handle variable-length sequences, but its large parameter-to-
 993 sample ratio makes it prone to overfitting in small-data regimes and yields limited interpretability
 994 compared to feature-based approaches.

995 **Transformer-based Confidence Predictor.** This baseline applies a lightweight Transformer (Vaswani
 996 et al., 2017) encoder to the raw confidence trajectories. We use one self-attention layer with model
 997 dimension 32, two attention heads, a feed-forward size of 64, and dropout 0.3. Learnable positional
 998 embeddings (up to length 2000) encode temporal order, and an attention pooling layer aggregates the
 999 sequence before a two-layer classifier (32→16→2). The model has about 3k–5k parameters and is
 1000 trained with Adam (learning rate 0.001, batch size 4) and early stopping. While the Transformer can
 1001 capture long-range dependencies and trains in parallel, it is computationally more demanding and
 1002 unstable in small-data settings.

1003 **Neural Network (MLP).** Operating on the engineered 48-dimensional feature representation, this
 1004 baseline uses a two-hidden-layer multilayer perceptron with sizes 48→32→16→2 and ReLU activations.
 1005 Regularization includes dropout and L2 penalty ($\alpha = 0.01$), and the network has about 2k–3k
 1006 parameters. Training is performed with Adam and early stopping. The MLP provides nonlinear
 1007 modeling capacity over compact, interpretable features, but its performance can fluctuate with dataset
 1008 size and it remains less transparent than linear models.

1009 **Gaussian Process Classifier.** We implement a Gaussian Process classifier with an RBF kernel
 1010 combined with a white-noise kernel. Kernel hyperparameters are optimized with three random
 1011 restarts, and predictions use up to 100 iterations. Being non-parametric, the model’s effective
 1012 complexity scales with the training set size. Gaussian Processes (Rasmussen & Williams, 2006)
 1013 naturally provide calibrated probabilistic outputs and flexible capacity, but incur cubic computational
 1014 cost $O(n^3)$, require careful kernel selection, and are impractical for larger datasets.

1015 **XGBoost Classifier.** This baseline uses gradient-boosted decision trees on the 48-dimensional
 1016 features, with 100 estimators of maximum depth 3, learning rate 0.1, row subsampling 0.8, column
 1017 subsampling 0.8, and both L1 (0.1) and L2 (1.0) regularization. The ensemble corresponds to roughly
 1018 1k–2k effective parameters. XGBoost (Chen & Guestrin, 2016) is robust on tabular data and captures
 1019 higher-order feature interactions, but still risks overfitting in very small datasets and provides less
 1020 interpretability than linear models.

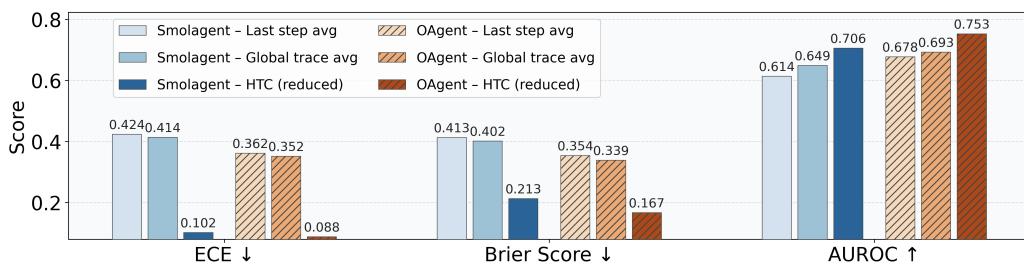
1021 In summary, the end-to-end neural encoders (LSTM, Transformer) directly consume raw confidence
 1022 trajectories but suffer from high parameter counts relative to the limited data, leading to severe
 1023 overfitting and unstable behavior. The feature-based nonlinear methods (MLP, Gaussian Process,
 1024 XGBoost) make better use of the engineered 48-dimensional representation and achieve stronger
 1025 performance overall, yet they remain less interpretable and still prone to variance under small-sample
 1026 regimes. These limitations highlight the motivation for our proposed lightweight linear calibrators,
 1027 which strike a favorable balance between stability, interpretability, and data efficiency.

1026 A.2.6 IMPLEMENTATION AND HYPERPARAMETER DETAILS
10271028 For completeness, we summarize the implementation details of our proposed Holistic Trajectory
1029 Calibration (**HTC**) method.
10301031 **Cross-Validation Strategy.** We adopt a 5-fold stratified cross-validation protocol to preserve class
1032 balance. With 500 labeled trajectories, each fold contains 100 samples. Within each fold, we
1033 use an 80%/20% split for training and validation. All experiments use a fixed random seed of 42.
1034 The `liblinear` solver is deterministic, and all code, hyperparameters, and configurations are
1035 version-controlled for exact reproducibility. The maximum iteration count is set to 1000, although
1036 convergence typically occurs within 50–100 iterations.
10371038 **Hyperparameter Optimization.** The regularization strength α is tuned via grid search over 15
1039 candidate values: $\{0.001, 0.01, 0.1, 1.0, 2.0, 3.0, 4.0, 5.0, 6.0, 7.0, 8.0, 9.0, 10.0, 20.0, 50.0\}$. Model
1040 selection is based on a combined criterion that maximizes AUROC while minimizing both Brier
1041 Score and ECE, averaged across folds. In practice, the optimal α typically falls in the range of
1042 1.0–5.0. The resulting sparse solutions select 15–25 features (50–70% of the 48 total). The model
1043 demonstrates low variance across folds, with training requiring less than one second per fold and
1044 inference under one millisecond per sample.
10451046 **Key Advantages.** Our approach provides (i) interpretability through transparent feature weights
1047 and selection, (ii) stability across varying data sizes, (iii) high computational efficiency in both
1048 training and inference, and (iv) robustness to overfitting compared to more complex baselines. These
1049 properties make the method particularly suitable for small datasets (< 500 samples), production
1050 systems requiring reliable calibration, and resource-constrained settings.
10511052 A.3 ADDITIONAL EXPERIMENTAL RESULTS
10531054 A.3.1 MAIN RESULTS
10551056 In Section 3.2, we presented a summary of our **HTC** framework’s performance against baselines on a
1057 subset of three representative datasets. For a complete overview of our method’s efficacy, we provide
1058 detailed results for both the **HTC-Full** and **HTC-Reduced** variants across all eight datasets in our
1059 benchmark suite.
10601061 Table 4 presents the mean and standard deviation for ECE, Brier Score, and AUROC when using the
1062 full set of 48 trajectory features (L2 regularization). These results demonstrate the robust performance
1063 of **HTC** even when all features are considered, establishing a strong upper bound for the feature set.
10641065 Table 5 details the performance of the (reduced feature) variant (L1 regularization), showing its
1066 mean and standard deviation for ECE, Brier Score, and AUROC across all datasets. Crucially, this
1067 table also includes the mean and standard deviation of the number of features selected by the Lasso
1068 regularization across different random seeds, providing insight into the sparsity and efficiency of
1069 this approach. The consistent strong performance with a reduced feature set further validates the
1070 effectiveness of our feature engineering and selection process.
10711072 Table 4: **HTC** Performance with Full Feature Set across All Datasets.
10731074

| Dataset | ECE | | Brier Score | | AUROC | |
|------------|--------|--------|-------------|--------|--------|--------|
| | Mean | Std | Mean | Std | Mean | Std |
| HLE | 0.0720 | 0.0108 | 0.0977 | 0.0019 | 0.6169 | 0.0231 |
| GPQA | 0.1241 | 0.0110 | 0.2185 | 0.0016 | 0.7040 | 0.0070 |
| SimpleQA | 0.0748 | 0.0065 | 0.1500 | 0.0029 | 0.7267 | 0.0103 |
| MATH500 | 0.0604 | 0.0071 | 0.0773 | 0.0015 | 0.7875 | 0.0178 |
| GAIA | 0.1692 | 0.0114 | 0.2654 | 0.0093 | 0.6204 | 0.0164 |
| HotpotQA | 0.1156 | 0.0060 | 0.1930 | 0.0020 | 0.7141 | 0.0061 |
| MMLU-Pro | 0.0775 | 0.0075 | 0.1257 | 0.0032 | 0.7276 | 0.0118 |
| StrategyQA | 0.0785 | 0.0081 | 0.1405 | 0.0015 | 0.6698 | 0.0054 |

1080 Table 5: **HTC** Performance with Reduced Feature Set and Feature Counts across All Datasets
1081

| Dataset | ECE | | Brier Score | | AUROC | | Features | |
|------------|--------|--------|-------------|--------|--------|--------|----------|-----|
| | Mean | Std | Mean | Std | Mean | Std | Mean | Std |
| HLE | 0.0305 | 0.0038 | 0.0897 | 0.0005 | 0.6439 | 0.0199 | 8.2 | 4.1 |
| GPQA | 0.1022 | 0.0159 | 0.2134 | 0.0018 | 0.7060 | 0.0066 | 23.4 | 1.9 |
| SimpleQA | 0.0676 | 0.0081 | 0.1402 | 0.0024 | 0.7523 | 0.0141 | 14.4 | 3.1 |
| MATH500 | 0.0476 | 0.0088 | 0.0701 | 0.0006 | 0.8162 | 0.0075 | 15.2 | 3.0 |
| GAIA | 0.1420 | 0.0100 | 0.2332 | 0.0026 | 0.6860 | 0.0131 | 4.8 | 1.5 |
| HotpotQA | 0.0824 | 0.0109 | 0.1824 | 0.0007 | 0.7288 | 0.0026 | 7.6 | 0.8 |
| MMLU-Pro | 0.0592 | 0.0047 | 0.1167 | 0.0009 | 0.7492 | 0.0075 | 13.8 | 3.0 |
| StrategyQA | 0.0545 | 0.0048 | 0.1357 | 0.0014 | 0.6647 | 0.0117 | 15.2 | 6.3 |

1101 Figure 7: The Impact of Agent Framework on the GPQA dataset.
1102
11031104 A.3.2 FEATURE IMPORTANCE ANALYSIS
1105

1106 To better understand the internal behavior of **HTC**, we analyze which diagnostic features are most
1107 influential across datasets and selection levels. Figure 8 shows the absolute weight magnitudes
1108 of the ℓ_1 -regularized logistic calibrator on eight benchmarks, highlighting that certain dynamics
1109 (e.g., confidence change) and stability measures (e.g., attention entropy, token volatility) consistently
1110 receive high importance. Figure 9 provides a complementary perspective by reporting feature
1111 selection frequencies under different levels (Top1, Top3, Top5, and all selected), allowing us to
1112 quantify which features are repeatedly chosen across runs. Table 6 further aggregates these results
1113 into a ranked list of top features, organized by category. Together, these analyses show that temporal
1114 dynamics and stability signals emerge as the most dominant indicators of reliability, while positional
1115 and structural features contribute complementary but non-negligible signals. This provides clear
1116 interpretability benefits: **HTC** not only delivers strong calibration but also yields transparent insights
1117 into which aspects of a reasoning trajectory drive reliable predictions.

1118 A.3.3 DOMAIN TRANSFER ANALYSIS
1119

1120 We further investigate the generalization ability of **HTC** across domains, by training the calibrator on
1121 one dataset and evaluating it directly on others without retraining. Figures 10–13 present transfer
1122 matrices for both GPT-4.1 and GPT-4o under reduced and full feature sets, evaluated on ECE, Brier
1123 Score, and AUROC. These heatmaps reveal that **HTC** achieves stable cross-domain calibration: mod-
1124 els trained on one benchmark often transfer reasonably well to others, especially among datasets with
1125 similar reasoning structures (e.g., QA benchmarks). Figure 14 provides an aggregated comparison,
1126 showing that GPT-4.1 consistently outperforms GPT-4o by a small margin, but both demonstrate
1127 robust transferability across metrics. Tables 7–10 give the complete numerical results, confirming that
1128 reduced feature sets maintain performance levels close to the full feature space, thereby validating
1129 the efficiency and compactness of our design.

1130 Overall, these results demonstrate that **HTC** is not only effective in-domain but also generalizes
1131 reliably across diverse datasets, while being relatively insensitive to the underlying backbone model
1132 or the size of the feature set.

1133

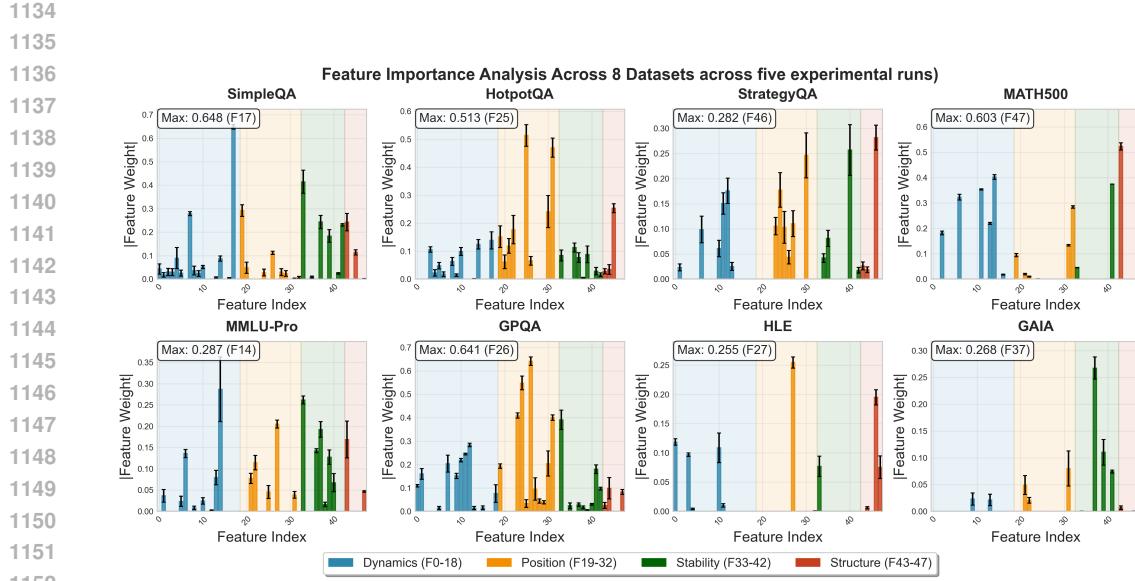


Figure 8: Feature importance analysis across all datasets on five experimental runs.

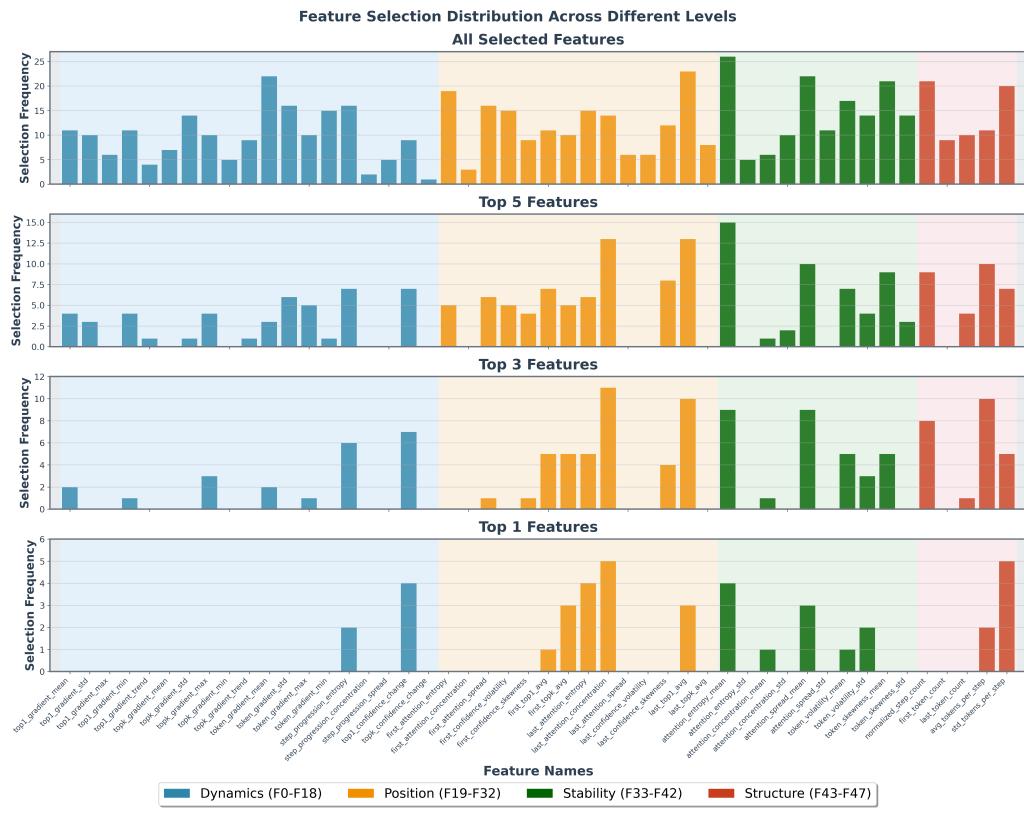


Figure 9: Full feature selection distribution across different levels (Top1, Top3, Top5 and all selected) on four feature categories.

1188

1189

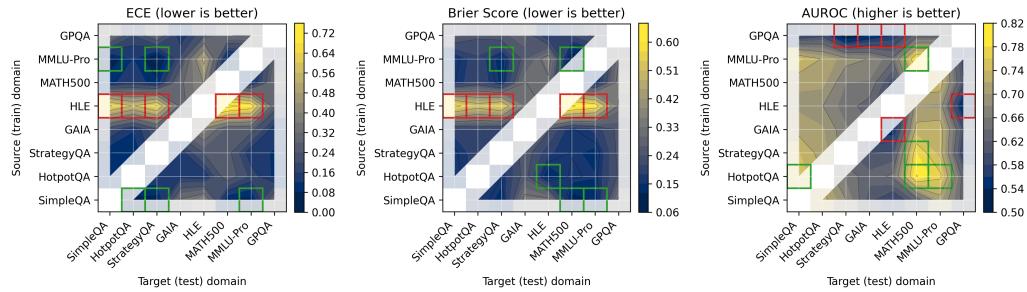
1190

Table 6: Feature Selection Results

| Selection Level | Feature Index | Feature Name | Category | Frequency | Percentage |
|-----------------|---------------|------------------------------|-----------|-----------|------------|
| Top 1 | F47 | std_tokens_per_step | Structure | 5 | 12.5 |
| | F27 | last_attention_concentration | Position | 5 | 12.5 |
| | F17 | top1_confidence_change | Dynamics | 4 | 10 |
| | F33 | attention_entropy_mean | Stability | 4 | 10 |
| | F26 | last_attention_entropy | Position | 4 | 10 |
| | F25 | first_topk_avg | Position | 3 | 7.5 |
| | F31 | last_top1_avg | Position | 3 | 7.5 |
| | F37 | attention_spread_mean | Stability | 3 | 7.5 |
| | F40 | token_volatility_std | Stability | 2 | 5 |
| Top 3 | F46 | avg_tokens_per_step | Structure | 2 | 5 |
| | F27 | last_attention_concentration | Position | 11 | 9.2 |
| | F31 | last_top1_avg | Position | 10 | 8.3 |
| | F46 | avg_tokens_per_step | Structure | 10 | 8.3 |
| | F33 | attention_entropy_mean | Stability | 9 | 7.5 |
| | F37 | attention_spread_mean | Stability | 9 | 7.5 |
| | F43 | normalized_step_count | Structure | 8 | 6.7 |
| | F17 | top1_confidence_change | Dynamics | 7 | 5.8 |
| | F14 | step_progression_entropy | Dynamics | 6 | 5 |
| Top 5 | F25 | first_topk_avg | Position | 5 | 4.2 |
| | F47 | std_tokens_per_step | Structure | 5 | 4.2 |
| | F33 | attention_entropy_mean | Stability | 15 | 7.5 |
| | F31 | last_top1_avg | Position | 13 | 6.5 |
| | F27 | last_attention_concentration | Position | 13 | 6.5 |
| | F37 | attention_spread_mean | Stability | 10 | 5 |
| | F46 | avg_tokens_per_step | Structure | 10 | 5 |
| | F43 | normalized_step_count | Structure | 9 | 4.5 |
| | F41 | token_skewness_mean | Stability | 9 | 4.5 |
| Top 9 | F30 | last_confidence_skewness | Position | 8 | 4 |
| | F17 | top1_confidence_change | Dynamics | 7 | 3.5 |
| | F39 | token_volatility_mean | Stability | 7 | 3.5 |

1216

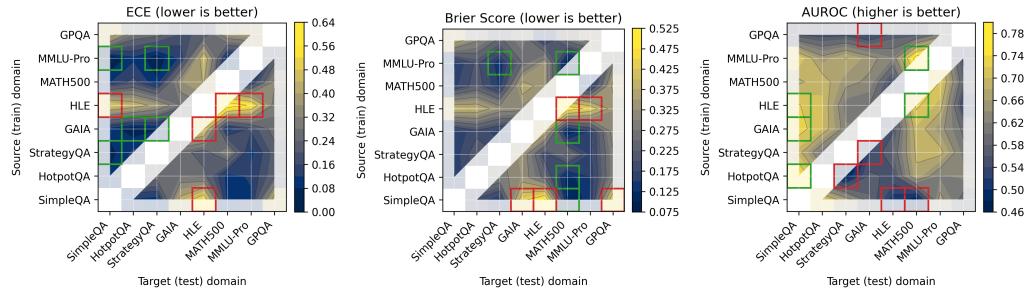
1217

Figure 10: Domain transfer matrix with reduced features using **GPT-4.1**.

1228

1229

1230

Figure 11: Domain transfer matrix with full features using **GPT-4.1**.

1241

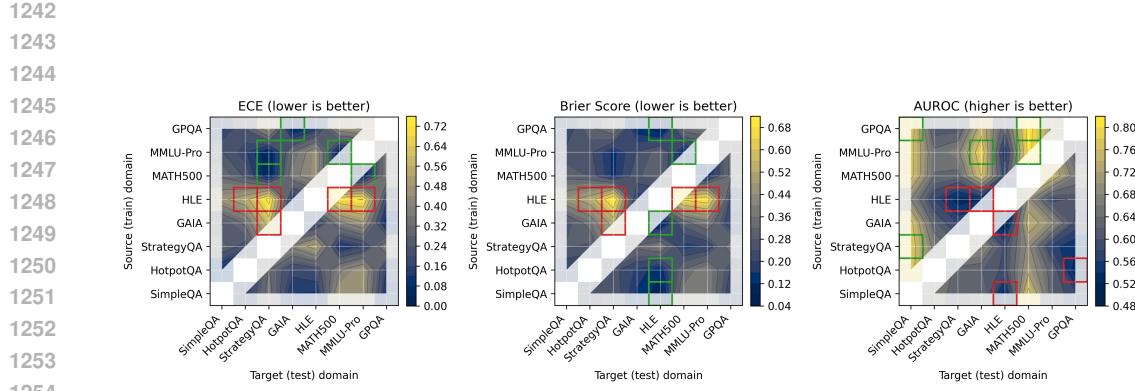
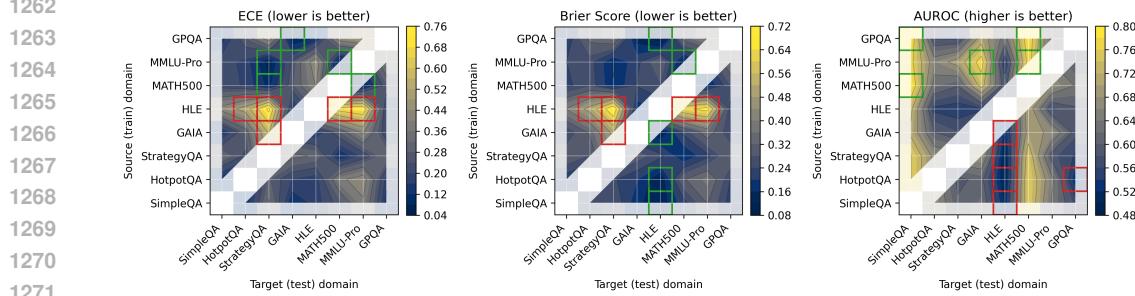
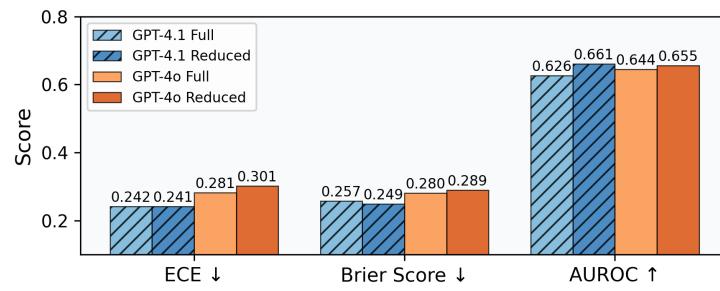
Figure 12: Domain transfer matrix with reduced features using **GPT-4o**.Figure 13: Domain transfer matrix with full features using **GPT-4o**.

Figure 14: Comparison of different base LLMs (GPT-4.1 vs GPT-4o, with full and reduced features) on the effect of domain transfer performance. We show the average of all domain transfer results on ECE, BS and AUROC metrics. This figure illustrates that different models show stable domain transfer performance while GPT-4.1 is slightly better than GPT-4o.

1294
1295

1296

1297

1298

Table 7: Full Feature Transfer Results: ECE, Brier Score (BS), and AUROC matrices (GPT-4.1).

1299

1300

| ECE | HLE | GPQA | SimpleQA | MATH500 | GAIA | HotpotQA | MMLU-Pro | StrategyQA |
|-------------|--------|--------|----------|---------|--------|----------|----------|------------|
| HLE | - | 0.2920 | 0.4807 | 0.6114 | 0.2729 | 0.4376 | 0.5496 | 0.3765 |
| GPQA | 0.2096 | - | 0.4461 | 0.3673 | 0.1912 | 0.2788 | 0.3958 | 0.2553 |
| SimpleQA | 0.5361 | 0.4349 | - | 0.0929 | 0.4205 | 0.1132 | 0.1089 | 0.0989 |
| MATH500 | 0.4030 | 0.3300 | 0.1153 | - | 0.3563 | 0.1733 | 0.1149 | 0.0862 |
| GAIA | 0.4834 | 0.2659 | 0.0994 | 0.0964 | - | 0.0530 | 0.1694 | 0.0786 |
| HotpotQA | 0.2553 | 0.2291 | 0.0894 | 0.1009 | 0.2786 | - | 0.1274 | 0.2691 |
| MMLU-Pro | 0.4570 | 0.2126 | 0.0390 | 0.0806 | 0.1259 | 0.1070 | - | 0.0562 |
| StrategyQA | 0.2139 | 0.2163 | 0.0644 | 0.3163 | 0.2626 | 0.1292 | 0.1072 | - |
| Brier Score | HLE | GPQA | SimpleQA | MATH500 | GAIA | HotpotQA | MMLU-Pro | StrategyQA |
| HLE | - | 0.3256 | 0.3880 | 0.4796 | 0.3136 | 0.3833 | 0.4355 | 0.3227 |
| GPQA | 0.1691 | - | 0.3799 | 0.2584 | 0.2932 | 0.3057 | 0.3133 | 0.2401 |
| SimpleQA | 0.5016 | 0.4459 | - | 0.0982 | 0.4332 | 0.1939 | 0.1369 | 0.1482 |
| MATH500 | 0.3298 | 0.3622 | 0.1650 | - | 0.3725 | 0.2187 | 0.1437 | 0.1476 |
| GAIA | 0.3786 | 0.3187 | 0.1551 | 0.0960 | - | 0.1895 | 0.1576 | 0.1610 |
| HotpotQA | 0.2353 | 0.2909 | 0.1493 | 0.1106 | 0.3266 | - | 0.1516 | 0.2829 |
| MMLU-Pro | 0.3292 | 0.2907 | 0.1469 | 0.0915 | 0.2492 | 0.2077 | - | 0.1341 |
| StrategyQA | 0.1842 | 0.2910 | 0.1556 | 0.2606 | 0.3333 | 0.2250 | 0.1649 | - |
| AUROC | HLE | GPQA | SimpleQA | MATH500 | GAIA | HotpotQA | MMLU-Pro | StrategyQA |
| HLE | - | 0.5377 | 0.7213 | 0.7257 | 0.6001 | 0.7033 | 0.6926 | 0.6087 |
| GPQA | 0.5694 | - | 0.5588 | 0.5623 | 0.5308 | 0.5508 | 0.6161 | 0.5541 |
| SimpleQA | 0.5112 | 0.5874 | - | 0.4728 | 0.5834 | 0.7186 | 0.6507 | 0.6572 |
| MATH500 | 0.5496 | 0.5773 | 0.6749 | - | 0.6340 | 0.6953 | 0.7194 | 0.6499 |
| GAIA | 0.5557 | 0.5926 | 0.7314 | 0.6966 | - | 0.7002 | 0.6504 | 0.6248 |
| HotpotQA | 0.5527 | 0.6254 | 0.7540 | 0.6725 | 0.5660 | - | 0.6156 | 0.4717 |
| MMLU-Pro | 0.6204 | 0.5521 | 0.7168 | 0.7822 | 0.6503 | 0.6496 | - | 0.6820 |
| StrategyQA | 0.5793 | 0.6108 | 0.6912 | 0.6861 | 0.5360 | 0.6043 | 0.6904 | - |

1321

1322

1323

1324

1325

Table 8: Reduced Feature Transfer Results: ECE, Brier Score (BS), and AUROC matrices (GPT-4.1).

1326

1327

| ECE | HLE | GPQA | SimpleQA | MATH500 | GAIA | HotpotQA | MMLU-Pro | StrategyQA |
|-------------|--------|--------|----------|---------|--------|----------|----------|------------|
| HLE | - | 0.2683 | 0.6482 | 0.7476 | 0.2973 | 0.5534 | 0.6904 | 0.6442 |
| GPQA | 0.2327 | - | 0.3306 | 0.3191 | 0.1335 | 0.1890 | 0.3231 | 0.2987 |
| SimpleQA | 0.3885 | 0.3040 | - | 0.0806 | 0.3268 | 0.0704 | 0.0691 | 0.0638 |
| MATH500 | 0.4029 | 0.3163 | 0.1072 | - | 0.3062 | 0.1490 | 0.0984 | 0.0951 |
| GAIA | 0.3872 | 0.1346 | 0.2052 | 0.2637 | - | 0.1234 | 0.2620 | 0.2359 |
| HotpotQA | 0.1294 | 0.1082 | 0.0974 | 0.2778 | 0.1499 | - | 0.1849 | 0.1497 |
| MMLU-Pro | 0.5041 | 0.2816 | 0.0535 | 0.0809 | 0.2038 | 0.1212 | - | 0.0283 |
| StrategyQA | 0.2452 | 0.1169 | 0.0791 | 0.2588 | 0.1394 | 0.0872 | 0.1323 | - |
| Brier Score | HLE | GPQA | SimpleQA | MATH500 | GAIA | HotpotQA | MMLU-Pro | StrategyQA |
| HLE | - | 0.3115 | 0.5795 | 0.6480 | 0.3289 | 0.5099 | 0.6023 | 0.5526 |
| GPQA | 0.1697 | - | 0.2822 | 0.2089 | 0.2767 | 0.2465 | 0.2512 | 0.2664 |
| SimpleQA | 0.3375 | 0.3304 | - | 0.1011 | 0.3392 | 0.1828 | 0.1217 | 0.1353 |
| MATH500 | 0.3081 | 0.3428 | 0.1578 | - | 0.3206 | 0.2115 | 0.1371 | 0.1406 |
| GAIA | 0.2459 | 0.2517 | 0.1913 | 0.1562 | - | 0.2064 | 0.1882 | 0.1945 |
| HotpotQA | 0.1191 | 0.2454 | 0.1466 | 0.1671 | 0.2574 | - | 0.1518 | 0.1603 |
| MMLU-Pro | 0.3493 | 0.3161 | 0.1478 | 0.0831 | 0.2624 | 0.2031 | - | 0.1312 |
| StrategyQA | 0.1862 | 0.2532 | 0.1526 | 0.1582 | 0.2634 | 0.2019 | 0.1382 | - |
| AUROC | HLE | GPQA | SimpleQA | MATH500 | GAIA | HotpotQA | MMLU-Pro | StrategyQA |
| HLE | - | 0.5277 | 0.6959 | 0.6982 | 0.6301 | 0.6807 | 0.6960 | 0.6492 |
| GPQA | 0.5619 | - | 0.5776 | 0.6274 | 0.5374 | 0.5946 | 0.6267 | 0.5049 |
| SimpleQA | 0.5756 | 0.6291 | - | 0.6265 | 0.6090 | 0.7321 | 0.7079 | 0.6807 |
| MATH500 | 0.5913 | 0.6046 | 0.7027 | - | 0.6580 | 0.7075 | 0.7361 | 0.6603 |
| GAIA | 0.5467 | 0.5998 | 0.7114 | 0.7294 | - | 0.7034 | 0.7355 | 0.6632 |
| HotpotQA | 0.6454 | 0.6277 | 0.7537 | 0.8016 | 0.6270 | - | 0.7669 | 0.6538 |
| MMLU-Pro | 0.6448 | 0.5893 | 0.7455 | 0.7922 | 0.6791 | 0.7177 | - | 0.6888 |
| StrategyQA | 0.6040 | 0.5782 | 0.7415 | 0.7804 | 0.5996 | 0.7089 | 0.7321 | - |

1348

1349

1350

1351

1352

Table 9: Full Feature Transfer Results: ECE, Brier Score, and AUROC matrices (GPT-4o).

1353

1354

| ECE | HLE | GPQA | SimpleQA | MATH500 | GAIA | HotpotQA | MMLU-Pro | StrategyQA |
|-------------|--------|--------|----------|---------|--------|----------|----------|------------|
| HLE | – | 0.1814 | 0.4252 | 0.6760 | 0.1502 | 0.5131 | 0.7106 | 0.7427 |
| GPQA | 0.1244 | – | 0.1452 | 0.4362 | 0.0673 | 0.2699 | 0.4845 | 0.4435 |
| SimpleQA | 0.1763 | 0.1462 | – | 0.2628 | 0.1305 | 0.1275 | 0.3872 | 0.3260 |
| MATH500 | 0.4039 | 0.2729 | 0.2172 | – | 0.2702 | 0.1946 | 0.0719 | 0.1039 |
| GAIA | 0.1589 | 0.1991 | 0.1444 | 0.2735 | – | 0.2817 | 0.4121 | 0.5721 |
| HotpotQA | 0.1709 | 0.1731 | 0.1360 | 0.3816 | 0.1944 | – | 0.4394 | 0.2855 |
| MMLU-Pro | 0.4473 | 0.3872 | 0.2168 | 0.0678 | 0.2505 | 0.1882 | – | 0.0583 |
| StrategyQA | 0.3108 | 0.3131 | 0.3201 | 0.1277 | 0.3365 | 0.2271 | 0.2100 | – |
| Brier Score | HLE | GPQA | SimpleQA | MATH500 | GAIA | HotpotQA | MMLU-Pro | StrategyQA |
| HLE | – | 0.2204 | 0.4026 | 0.6207 | 0.1768 | 0.5043 | 0.6685 | 0.7067 |
| GPQA | 0.0808 | – | 0.2086 | 0.3377 | 0.1415 | 0.3329 | 0.4074 | 0.3718 |
| SimpleQA | 0.1332 | 0.2205 | – | 0.2245 | 0.1714 | 0.2406 | 0.3226 | 0.2786 |
| MATH500 | 0.2754 | 0.2797 | 0.2296 | – | 0.2418 | 0.2673 | 0.1526 | 0.1620 |
| GAIA | 0.1284 | 0.2399 | 0.2438 | 0.2495 | – | 0.3299 | 0.3514 | 0.5011 |
| HotpotQA | 0.1264 | 0.2357 | 0.2141 | 0.3148 | 0.2039 | – | 0.3728 | 0.2470 |
| MMLU-Pro | 0.3176 | 0.3638 | 0.2400 | 0.1362 | 0.2276 | 0.2587 | – | 0.1561 |
| StrategyQA | 0.2520 | 0.3193 | 0.3115 | 0.1615 | 0.3013 | 0.2750 | 0.2281 | – |
| AUROC | HLE | GPQA | SimpleQA | MATH500 | GAIA | HotpotQA | MMLU-Pro | StrategyQA |
| HLE | – | 0.5958 | 0.7301 | 0.7040 | 0.5574 | 0.5535 | 0.5713 | 0.5452 |
| GPQA | 0.6256 | – | 0.7921 | 0.7864 | 0.7024 | 0.5747 | 0.6275 | 0.5694 |
| SimpleQA | 0.5179 | 0.5725 | – | 0.7573 | 0.6238 | 0.6732 | 0.6343 | 0.6037 |
| MATH500 | 0.5798 | 0.6380 | 0.7948 | – | 0.7378 | 0.6260 | 0.6751 | 0.6019 |
| GAIA | 0.5239 | 0.5852 | 0.7102 | 0.6964 | – | 0.6007 | 0.6707 | 0.5807 |
| HotpotQA | 0.4992 | 0.5217 | 0.7635 | 0.7600 | 0.6085 | – | 0.6487 | 0.5932 |
| MMLU-Pro | 0.5400 | 0.6003 | 0.7707 | 0.7798 | 0.7786 | 0.6570 | – | 0.6960 |
| StrategyQA | 0.4986 | 0.5570 | 0.7632 | 0.7462 | 0.6808 | 0.6668 | 0.5868 | – |

1375

1376

1377

1378

1379

1380

Table 10: Reduced Feature Transfer Results: ECE, Brier Score, and AUROC matrices (GPT-4o).

1381

| ECE | HLE | GPQA | SimpleQA | MATH500 | GAIA | HotpotQA | MMLU-Pro | StrategyQA |
|-------------|--------|--------|----------|---------|--------|----------|----------|------------|
| HLE | – | 0.1839 | 0.4613 | 0.6982 | 0.1245 | 0.5292 | 0.7266 | 0.7480 |
| GPQA | 0.1389 | – | 0.1448 | 0.4502 | 0.0476 | 0.2445 | 0.4914 | 0.4671 |
| SimpleQA | 0.1118 | 0.1099 | – | 0.4758 | 0.1181 | 0.1407 | 0.4612 | 0.3507 |
| MATH500 | 0.4568 | 0.3833 | 0.3065 | – | 0.3734 | 0.2312 | 0.0366 | 0.0316 |
| GAIA | 0.1513 | 0.0685 | 0.2345 | 0.3805 | – | 0.3343 | 0.4666 | 0.6019 |
| HotpotQA | 0.2045 | 0.1507 | 0.0892 | 0.3883 | 0.1787 | – | 0.4729 | 0.3234 |
| MMLU-Pro | 0.4572 | 0.3818 | 0.2434 | 0.0333 | 0.3207 | 0.1695 | – | 0.0403 |
| StrategyQA | 0.5212 | 0.3921 | 0.2948 | 0.1089 | 0.4135 | 0.2106 | 0.1735 | – |
| Brier Score | HLE | GPQA | SimpleQA | MATH500 | GAIA | HotpotQA | MMLU-Pro | StrategyQA |
| HLE | – | 0.2204 | 0.4529 | 0.6531 | 0.1685 | 0.5175 | 0.6895 | 0.7134 |
| GPQA | 0.0769 | – | 0.2112 | 0.3432 | 0.1340 | 0.3032 | 0.3898 | 0.3740 |
| SimpleQA | 0.0877 | 0.2008 | – | 0.3781 | 0.1508 | 0.2406 | 0.3925 | 0.2832 |
| MATH500 | 0.2928 | 0.3328 | 0.2981 | – | 0.2758 | 0.2755 | 0.1367 | 0.1503 |
| GAIA | 0.0885 | 0.1929 | 0.2562 | 0.2942 | – | 0.3458 | 0.3713 | 0.5163 |
| HotpotQA | 0.1301 | 0.2327 | 0.2004 | 0.3249 | 0.1885 | – | 0.3936 | 0.2572 |
| MMLU-Pro | 0.3007 | 0.3427 | 0.2534 | 0.1299 | 0.2427 | 0.2525 | – | 0.1488 |
| StrategyQA | 0.3822 | 0.3597 | 0.2838 | 0.1825 | 0.3336 | 0.2648 | 0.2008 | – |
| AUROC | HLE | GPQA | SimpleQA | MATH500 | GAIA | HotpotQA | MMLU-Pro | StrategyQA |
| HLE | – | 0.6062 | 0.6791 | 0.6943 | 0.5372 | 0.5753 | 0.5879 | 0.5256 |
| GPQA | 0.6184 | – | 0.7852 | 0.8178 | 0.7453 | 0.6055 | 0.7074 | 0.6064 |
| SimpleQA | 0.5206 | 0.5808 | – | 0.7756 | 0.6425 | 0.6633 | 0.6214 | 0.6200 |
| MATH500 | 0.5824 | 0.6603 | 0.7735 | – | 0.7635 | 0.6388 | 0.7517 | 0.6292 |
| GAIA | 0.4956 | 0.5925 | 0.7721 | 0.7431 | – | 0.6223 | 0.6894 | 0.5752 |
| HotpotQA | 0.5932 | 0.5188 | 0.7795 | 0.7078 | 0.6209 | – | 0.6243 | 0.6301 |
| MMLU-Pro | 0.5449 | 0.6279 | 0.7732 | 0.8155 | 0.7897 | 0.6467 | – | 0.6710 |
| StrategyQA | 0.5927 | 0.5447 | 0.7839 | 0.6897 | 0.6876 | 0.6648 | 0.5679 | – |

1402

1403

1404
1405

A.4 ABLATION STUDY ON FEATURE CATEGORIES

1406
1407
1408
1409
1410
1411

To further examine the contribution of different feature categories, we conducted a systematic ablation study on the combined dataset of seven benchmarks (3,446 trajectories). The calibrator was trained on different subsets of the 48 features, including all single-category models (Dynamics only, Position only, Stability only, Structure only), all pairwise combinations, all three-way combinations, and the full feature set, yielding 15 configurations in total. Figures 15 and Table 11 summarize the results. Several clear findings emerge:

1412
1413
1414
1415

- **Full feature set performs best.** Using all 48 features achieves the highest AUROC (0.8430), the lowest Brier Score (0.1471), and the lowest ECE (0.0328). This demonstrates that the entire feature map provides complementary information that cannot be captured by any smaller subset.
- **Multi-category combinations outperform single categories.** Every two- or three-way combination substantially improves over the best single category. For example, Dynamics+Position+Stability achieves AUROC = 0.8419, which is +0.0137 higher than the strongest single category (Dynamics, AUROC = 0.8282).
- **Single categories are insufficient.** When restricted to only one category, performance drops noticeably (AUROC 0.783–0.828). Structure alone is the weakest (0.783 AUROC), showing that contextual information is not sufficient without dynamics or stability. This highlights the need for diverse diagnostic signals.
- **Category complementarity emerges with scale and diversity.** On the combined dataset, which is larger and more diverse than individual tasks, the synergy across categories becomes much more evident. This contrasts with the single-dataset setting (e.g., SimpleQA), where rankings can vary. The aggregated analysis demonstrates that HTC’s design is robust and general.

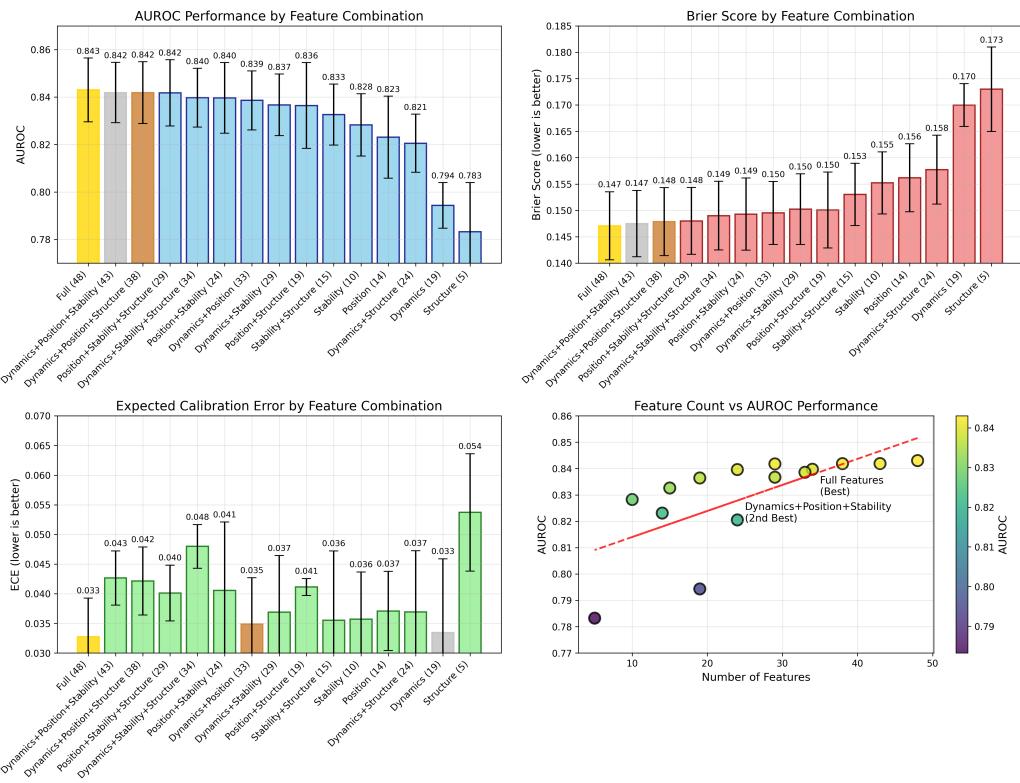
1424
1425
1426
14271455
1456
1457

Figure 15: Performance of calibrators trained on different feature combinations. Results are averaged across 3,446 trajectories from seven datasets. Multi-category combinations consistently outperform single categories, and the full feature set achieves the best results.

Table 11: Performance summary of feature ablation study (sorted by AUROC)

| Feature Combination | AUROC (\uparrow) | | Brier Score (\downarrow) | | ECE (\downarrow) | |
|-----------------------------------|----------------------|--------|------------------------------|--------|----------------------|--------|
| | Mean | Std | Mean | Std | Mean | Std |
| Full (48) | 0.8430 | 0.0134 | 0.1471 | 0.0065 | 0.0328 | 0.0065 |
| Dynamics+Position+Stability (43) | 0.8419 | 0.0127 | 0.1475 | 0.0063 | 0.0427 | 0.0046 |
| Dynamics+Position+Structure (38) | 0.8419 | 0.0130 | 0.1479 | 0.0065 | 0.0422 | 0.0057 |
| Position+Stability+Structure (29) | 0.8418 | 0.0140 | 0.1480 | 0.0064 | 0.0401 | 0.0047 |
| Dynamics+Stability+Structure (34) | 0.8397 | 0.0124 | 0.1490 | 0.0065 | 0.0480 | 0.0037 |
| Position+Stability (24) | 0.8396 | 0.0149 | 0.1493 | 0.0069 | 0.0406 | 0.0115 |
| Dynamics+Position (33) | 0.8386 | 0.0125 | 0.1495 | 0.0060 | 0.0349 | 0.0078 |
| Dynamics+Stability (29) | 0.8367 | 0.0130 | 0.1502 | 0.0067 | 0.0369 | 0.0095 |
| Position+Structure (19) | 0.8364 | 0.0181 | 0.1501 | 0.0072 | 0.0411 | 0.0014 |
| Stability+Structure (15) | 0.8326 | 0.0128 | 0.1530 | 0.0059 | 0.0355 | 0.0117 |
| Stability (10) | 0.8282 | 0.0131 | 0.1552 | 0.0059 | 0.0357 | 0.0079 |
| Position (14) | 0.8231 | 0.0173 | 0.1562 | 0.0064 | 0.0371 | 0.0067 |
| Dynamics+Structure (24) | 0.8205 | 0.0122 | 0.1577 | 0.0065 | 0.0369 | 0.0103 |
| Dynamics (19) | 0.7943 | 0.0096 | 0.1700 | 0.0041 | 0.0335 | 0.0124 |
| Structure (5) | 0.7832 | 0.0207 | 0.1730 | 0.0080 | 0.0537 | 0.0099 |

A.5 DETAILED FEATURE DESCRIPTION

A.5.1 FEATURE DEFINITIONS

Agent reliability is not a snapshot property of the last step but an emergent property of the whole trajectory. Our feature set operationalizes this view along four complementary axes:

- **Dynamics** — how confidence *evolves* across steps (trend, variability, accelerations).
- **Position** — what the *first* and *last* steps reveal (onset vs. resolution).
- **Stability** — whether signals *converge consistently* (low volatility, low entropy drift).
- **Structure** — the *form factor* of a trajectory (length and token allocation across steps).

Notation. A trajectory τ has S steps indexed by $t = 1, \dots, S$. At step t there are n_t tokens with positive “confidence” values $r_{t,1}, \dots, r_{t,n_t}$. We normalize within-step to obtain a discrete distribution

$$\pi_{t,i} = \frac{r_{t,i}}{\sum_{j=1}^{n_t} r_{t,j} + \varepsilon}, \quad \varepsilon = 10^{-8}. \quad (10)$$

The within-step mean and standard deviation are

$$\mu_t = \frac{1}{n_t} \sum_{i=1}^{n_t} r_{t,i}, \quad (11)$$

$$\sigma_t = \sqrt{\frac{1}{n_t} \sum_{i=1}^{n_t} (r_{t,i} - \mu_t)^2}. \quad (12)$$

We define per-step summaries of the distribution $r_{t,:}$:

$$H_t = - \sum_{i=1}^{n_t} \pi_{t,i} \log(\pi_{t,i} + \varepsilon) \quad (\text{entropy}), \quad (13)$$

$$\kappa_t = \frac{\max_i r_{t,i}}{\mu_t + \varepsilon} \quad (\text{concentration}), \quad (14)$$

$$\rho_t = \frac{\sigma_t}{\mu_t + \varepsilon} \quad (\text{spread / volatility}), \quad (15)$$

$$\text{skew}_t = \frac{1}{n_t} \sum_{i=1}^{n_t} \left(\frac{r_{t,i} - \mu_t}{\sigma_t + \varepsilon} \right)^3 \quad (\text{skewness}). \quad (16)$$

1512 For each step t we also compute aggregated confidences:
 1513

$$1514 \quad x_t = \frac{1}{n_t} \sum_{i=1}^{n_t} \text{Top1Conf}(t, i), \quad (17)$$

$$1517 \quad y_t = \frac{1}{n_t} \sum_{i=1}^{n_t} \text{TopkConf}(t, i). \quad (18)$$

1520 Cross-step differences (“gradients”) are
 1521

$$1522 \quad \Delta x_t = x_{t+1} - x_t, \quad (19)$$

$$1523 \quad \Delta y_t = y_{t+1} - y_t, \quad t = 1, \dots, S-1. \quad (20)$$

1525 All undefined statistics (e.g., $S < 2$ or $n_t < 2$) are set to 0, consistent with our implementation.
 1526

1527 **Intuition.** H_t measures dispersion (lower is more focused), κ_t captures dominance of the top
 1528 alternative, ρ_t is a scale-free volatility index, and skew_t encodes asymmetry. Reliable trajectories
 1529 tend to show decreasing entropy, stable volatility, and consistent trends in $\Delta x_t, \Delta y_t$.

1530 **Category A: Dynamics (19 features).** *Purpose:* capture how confidence changes across the
 1531 trajectory. Reliable reasoning tends to exhibit steady, low-variance growth; erratic failures show
 1532 oscillations, spikes, or regressions.
 1533

- **top1_gradient_mean:** $\text{mean}(\{\Delta x_t\}_{t=1}^{S-1})$. Average “velocity” of top-1 confidence growth.
- **top1_gradient_std:** $\text{std}(\{\Delta x_t\}_{t=1}^{S-1})$. Large variance indicates unstable or oscillatory confidence.
- **top1_gradient_max:** $\max\{\Delta x_t\}_{t=1}^{S-1}$. Largest single upward jump in top-1 confidence.
- **top1_gradient_min:** $\min\{\Delta x_t\}_{t=1}^{S-1}$. Largest downward collapse of top-1 confidence.
- **top1_gradient_trend:** $\Delta x_{S-1} - \Delta x_1$ (if $S \geq 3$). Detects acceleration or deceleration of belief formation.
- **topk_gradient_mean, topk_gradient_std, topk_gradient_max, topk_gradient_min, topk_gradient_trend:** identical statistics computed for top- k confidence $\{y_t\}$. Capture broader consensus dynamics across multiple hypotheses.
- **token_gradient_mean, token_gradient_std, token_gradient_max, token_gradient_min:** computed from local token differences $\{r_{t,i+1} - r_{t,i}\}$. Reveal whether a step is pruning sharply (large gradients) or dithering (flat gradients).
- **step_progression_entropy:** $\text{std}(\{H_t\}) / (\text{mean}(\{H_t\}) + \varepsilon)$.
- **step_progression_concentration:** $\text{std}(\{\kappa_t\}) / (\text{mean}(\{\kappa_t\}) + \varepsilon)$.
- **step_progression_spread:** $\text{std}(\{\rho_t\}) / (\text{mean}(\{\rho_t\}) + \varepsilon)$. These coefficients of variation measure how entropy, concentration, and spread evolve; convergence implies decreasing ratios.
- **top1_confidence_change:** $x_S - x_1$.
- **topk_confidence_change:** $y_S - y_1$. Capture overall strengthening or weakening of confidence from start to end.

1557 **Category B: Position (14 features).** *Purpose:* the first and last steps capture complementary
 1558 aspects. Early steps reflect exploration, late steps commitment. Comparing them diagnoses premature
 1559 certainty or end-stage overconfidence.

- **first_attention_entropy:** H_1 . Dispersion of the first step distribution.
- **first_attention_concentration:** κ_1 . Peakedness of the first step.
- **first_attention_spread:** ρ_1 . Variability relative to the mean.
- **first_confidence_volatility:** ρ_1 . Same as spread, interpreted as instability at the onset.
- **first_confidence_skewness:** skew_1 . Asymmetry of the distribution at the first step.

1566 • **first_top1_avg**: x_1 . Average top-1 confidence at the first step.
 1567 • **first_topk_avg**: y_1 . Average top- k confidence at the first step.
 1568 • **last_attention_entropy**: H_S . Dispersion at the last step.
 1569 • **last_attention_concentration**: κ_S . Sharpness at the last step.
 1570 • **last_attention_spread**: ρ_S . Variability at the last step.
 1571 • **last_confidence_volatility**: ρ_S . Instability at the end.
 1572 • **last_confidence_skewness**: skew_S . Tail asymmetry at the last step.
 1573 • **last_top1_avg**: x_S . Final top-1 confidence.
 1574 • **last_topk_avg**: y_S . Final top- k confidence.
 1575

1576 **Category C: Stability (10 features).** *Purpose*: measure consistency across steps. Reliable trajectories are smooth; unreliable ones oscillate.
 1577

1578 • **attention_entropy_mean**: $\text{mean}(\{H_t\})$.
 1579 • **attention_entropy_std**: $\text{std}(\{H_t\})$.
 1580 • **attention_concentration_mean**: $\text{mean}(\{\kappa_t\})$.
 1581 • **attention_concentration_std**: $\text{std}(\{\kappa_t\})$.
 1582 • **attention_spread_mean**: $\text{mean}(\{\rho_t\})$.
 1583 • **attention_spread_std**: $\text{std}(\{\rho_t\})$. Together, these summarize whether attention signals converge smoothly or fluctuate widely.
 1584 • **token_volatility_mean**: $\text{mean}(\{\rho_t\})$. Average token-level volatility across steps.
 1585 • **token_volatility_std**: $\text{std}(\{\rho_t\})$. Step-to-step volatility variation.
 1586 • **token_skewness_mean**: $\text{mean}(\{\text{skew}_t\})$. Average asymmetry of token distribution.
 1587 • **token_skewness_std**: $\text{std}(\{\text{skew}_t\})$. Variation of asymmetry over steps.
 1588

1589 **Category D: Structure (5 features).** *Purpose*: capture trajectory form factor (length and token allocation). These features provide context: short trajectories may indicate premature certainty; long, irregular ones suggest hesitation.
 1590

1591 • **normalized_step_count**: $S/10$. Normalized trajectory length.
 1592 • **first_token_count**: n_1 . Number of tokens in the first step.
 1593 • **last_token_count**: n_S . Number of tokens in the last step.
 1594 • **avg_tokens_per_step**: $(\sum_{t=1}^S n_t)/S$. Average token count per step.
 1595 • **std_tokens_per_step**: $\text{std}(\{n_t\}_{t=1}^S)$. Variation in token counts across steps.
 1596

1597 In total, the 48 features provide a structured and interpretable representation of trajectory reliability. *Dynamics* quantify how confidence values evolve step by step, *Position* features highlight the complementary roles of the trajectory onset and resolution, *Stability* measures assess whether the process converges consistently across steps, and *Structure* features capture the overall form factor of reasoning traces. Together, these categories decompose reliability into distinct yet complementary dimensions: they allow us to pinpoint when an agent is consolidating versus oscillating, whether its final certainty is warranted by stable evidence, and how the length or allocation of tokens modulates calibration. Unlike opaque neural encoders, this feature map offers transparent diagnostics that both improve calibration and yield actionable insights into the mechanisms underlying agent reliability.
 1598

1599
 1600
 1601
 1602
 1603
 1604
 1605
 1606
 1607
 1608
 1609
 1610
 1611
 1612
 1613
 1614
 1615
 1616
 1617
 1618
 1619

1620 A.5.2 FEATURE MAP
16211622 Here is a detailed feature map.
1623

```

1 # =====
2 # FINAL OPTIMIZED FEATURE MAP - 48 Features (STABLE VERSION)
3 # =====
4
5 FEATURE_MAP_FINAL_STABLE = {
6     "Dynamics": { # Features capturing temporal changes and gradients across steps (19 features)
7         "1.1: Cross-step Gradients": [
8             'top1_gradient_mean', # 0
9             'top1_gradient_std', # 1
10            'top1_gradient_max', # 2
11            'top1_gradient_min', # 3
12            'top1_gradient_trend', # 4
13            'topk_gradient_mean', # 5
14            'topk_gradient_std', # 6
15            'topk_gradient_max', # 7
16            'topk_gradient_min', # 8
17            'topk_gradient_trend', # 9
18        ],
19        "1.2: Token-level Gradients": [
20            'token_gradient_mean', # 10
21            'token_gradient_std', # 11
22            'token_gradient_max', # 12
23            'token_gradient_min', # 13
24        ],
25        "1.3: Step Progression": [
26            'step_progression_entropy', # 14
27            'step_progression_concentration', # 15
28            'step_progression_spread', # 16
29        ],
30        "1.4: Confidence Change": [
31            'top1_confidence_change', # 17
32            'topk_confidence_change', # 18
33        ],
34    },
35    "Position": { # Features capturing key positional information (first/last steps) (14 features)
36        "2.1: First Step Specific": [
37            'first_attention_entropy', # 19
38            'first_attention_concentration', # 20
39            'first_attention_spread', # 21
40            'first_confidence_volatility', # 22
41            'first_confidence_skewness', # 23
42            'first_top1_avg', # 24
43            'first_topk_avg', # 25
44        ],
45        "2.2: Last Step Specific": [
46            'last_attention_entropy', # 26
47            'last_attention_concentration', # 27
48            'last_attention_spread', # 28
49            'last_confidence_volatility', # 29
50            'last_confidence_skewness', # 30
51            'last_top1_avg', # 31
52            'last_topk_avg', # 32
53        ],
54    },
55    "Stability": { # Features capturing stability and consistency patterns (10 features)
56        "3.1: Attention Stability": [
57            'attention_entropy_mean', # 33
58            'attention_entropy_std', # 34
59            'attention_concentration_mean', # 35
60            'attention_concentration_std', # 36
61            'attention_spread_mean', # 37
62            'attention_spread_std', # 38
63        ],
64        "3.2: Token-level Stability": [
65            'token_volatility_mean', # 39
66            'token_volatility_std', # 40
67            'token_skewness_mean', # 41
68            'token_skewness_std', # 42
69        ],
70    },
71    "Structure": { # Features capturing structural and derived information (5 features)
72        "4.1: Structural Metrics": [
73            'normalized_step_count', # 43
74            'first_token_count', # 44
75            'last_token_count', # 45
76            'avg_tokens_per_step', # 46
77            'std_tokens_per_step', # 47
78        ],
79    }
80}

```

1672 Listing 1: Final Optimized Feature Map (48 features, stable)
1673

1674 A.6 THEORETICAL MOTIVATION AND ANALYSIS
1675

1676 We present four claims with complete proofs to theoretically support Holistic Trajectory Calibration
1677 (**HTC**). Notation follows Section 2: a trajectory is denoted τ , its extracted 48-dimensional diagnostic
1678 features are $\phi(\tau) \in \mathbb{R}^{48}$, the **HTC** calibrator is $F_{\text{HTC}} : \mathbb{R}^{48} \rightarrow [0, 1]$, the last-step confidence is p_T ,
1679 and the ground-truth task outcome is $Y \in \{0, 1\}$. Expectations, variances, and entropies are with
1680 respect to the data-generating distribution.

1681 **Losses and Bayes risks.** For a predictor $q : \mathcal{X} \rightarrow [0, 1]$, the Brier loss and log-loss are
1682

$$L_{\text{Brier}}(q) = \mathbb{E}[(Y - q)^2], \quad (21)$$

$$L_{\log}(q) = -\mathbb{E}[Y \log q + (1 - Y) \log(1 - q)]. \quad (22)$$

1685 The Bayes-optimal predictors are conditional means $q^*(\cdot) = \mathbb{P}(Y = 1 \mid \cdot)$, and the corresponding
1686 Bayes risks are

$$\inf_q L_{\text{Brier}}(q) = \mathbb{E}[\text{Var}(Y \mid \cdot)], \quad \inf_q L_{\log}(q) = H(Y \mid \cdot). \quad (23)$$

1689 **Proposition 1 (Trajectory features dominate last-step confidence).** Let
1690

$$q_\phi^*(\tau) = \mathbb{P}(Y = 1 \mid \phi(\tau)), \quad q_T^*(p_T) = \mathbb{P}(Y = 1 \mid p_T).$$

1692 If $\sigma(p_T) \subseteq \sigma(\phi(\tau))$, then

$$L_{\text{Brier}}(q_\phi^*) \leq L_{\text{Brier}}(q_T^*), \quad (24)$$

$$L_{\log}(q_\phi^*) = H(Y \mid \phi(\tau)) \leq H(Y \mid p_T) = L_{\log}(q_T^*). \quad (25)$$

1695 Inequalities are strict whenever $\phi(\tau)$ contains strictly more information about Y than p_T .

1696 *Proof.* By equation 23, the Bayes Brier risk is $\mathbb{E}[\text{Var}(Y \mid \cdot)]$. By the law of total variance,
1697

$$\text{Var}(Y) = \mathbb{E}[\text{Var}(Y \mid \phi)] + \text{Var}(\mathbb{E}[Y \mid \phi]) = \mathbb{E}[\text{Var}(Y \mid p_T)] + \text{Var}(\mathbb{E}[Y \mid p_T]). \quad (26)$$

1698 Since $\sigma(\phi)$ refines $\sigma(p_T)$, $\text{Var}(\mathbb{E}[Y \mid \phi]) \geq \text{Var}(\mathbb{E}[Y \mid p_T])$, hence $\mathbb{E}[\text{Var}(Y \mid \phi)] \leq \mathbb{E}[\text{Var}(Y \mid p_T)]$, proving equation 24.

1701 For log-loss, the Bayes risk equals conditional entropy. By the chain rule,

$$H(Y \mid p_T) = H(Y \mid \phi, p_T) + I(Y; \phi \mid p_T) \geq H(Y \mid \phi). \quad (27)$$

1703 This proves equation 25. \square

1705 **Proposition 2 (Generalization of sparse linear **HTC** calibrator).** Let the **HTC** calibrator be
1706 $F_{\text{HTC}}(\phi(\tau)) = \sigma(w^\top \phi(\tau))$ with $\|w\|_1 \leq B$, features bounded as $\|\phi(\tau)\|_\infty \leq R$, and $d = 48$. The
1707 empirical Rademacher complexity of linear scores $s_w(x) = w^\top x$ on n samples satisfies

$$\widehat{\mathfrak{R}}_n \leq BR\sqrt{\frac{2\log(2d)}{n}}. \quad (28)$$

1711 Consequently, for any L -Lipschitz loss in the score s , with probability $1 - \delta$,

$$\mathbb{E}[\ell(Y, F_{\text{HTC}}(\phi(\tau)))] \leq \frac{1}{n} \sum_{i=1}^n \ell(y_i, F_{\text{HTC}}(\phi(\tau_i))) + 2LBR\sqrt{\frac{2\log(2d)}{n}} + 3\sqrt{\frac{\log(2/\delta)}{2n}}. \quad (29)$$

1714 In particular, for logistic loss $L = 1$; for Brier-on-probability $\tilde{\ell}(y, s) = (\sigma(s) - y)^2$, $L \leq \frac{1}{2}$.

1716 *Proof.* By ℓ_1 - ℓ_∞ duality,

$$\widehat{\mathfrak{R}}_n = \frac{B}{n} \mathbb{E}_\sigma \left[\max_{1 \leq j \leq d} \left| \sum_{i=1}^n \sigma_i \phi_j(\tau_i) \right| \right]. \quad (30)$$

1720 Each coordinate sum is sub-Gaussian with variance proxy $\leq nR^2$. A maximal inequality yields

$$\mathbb{E} \left[\max_{1 \leq j \leq d} |S_j| \right] \leq R\sqrt{2n\log(2d)}. \quad (31)$$

1723 Substitute into equation 30 to obtain equation 28. The generalization bound equation 29 follows from
1724 symmetrization and contraction. For logistic loss, $|\partial\ell/\partial s| \leq 1$; for $\tilde{\ell}$,

$$\left| \frac{\partial \tilde{\ell}}{\partial s} \right| = 2|\sigma(s) - y|\sigma(s)(1 - \sigma(s)) \leq \frac{1}{2}. \quad (32)$$

\square

1728
 1729 **Proposition 3 (Toy bound: why last-step can be optimistic).** Suppose task success requires all
 1730 T subgoals to be correct, with per-step reliability $p_t = \mathbb{P}(\text{subgoal } t \text{ correct} \mid \tau)$. If subgoals are
 1731 conditionally independent given τ and $p_T \geq \min_t p_t$, then

$$1732 \quad \mathbb{P}(Y = 1 \mid \tau) = \prod_{t=1}^T p_t \leq \min_t p_t \leq p_T. \quad (33)$$

1735 *Proof.* By assumption, conditional on τ , each subgoal outcome is independent and has success
 1736 probability p_t . Hence the probability that all T subgoals succeed is

$$1737 \quad \mathbb{P}(Y = 1 \mid \tau) = \prod_{t=1}^T p_t. \quad (34)$$

1740 For any finite set of numbers $\{a_t\} \subseteq [0, 1]$, it holds that

$$1742 \quad \prod_{t=1}^T a_t \leq \min_t a_t, \quad (35)$$

1744 because $\prod_t a_t \leq a_j$ for each j (since all factors are at most 1). Applying equation 35 to $\{p_t\}$ yields

$$1746 \quad \prod_{t=1}^T p_t \leq \min_t p_t. \quad (36)$$

1748 Finally, by assumption $p_T \geq \min_t p_t$, hence

$$1751 \quad \prod_{t=1}^T p_t \leq \min_t p_t \leq p_T. \quad (37)$$

1753 This establishes equation 33. \square

1754 *Remark.* This stylized model illustrates that last-step confidence can systematically overestimate
 1755 success when intermediate steps are fragile. **HTC** features are designed to capture such fragility.

1757 **Proposition 4 (From post-hoc to online via prefixes).** Let $\phi_{\leq k}(\tau)$ be diagnostics computed on
 1758 prefix $\tau_{\leq k}$. Define Bayes risks

$$1759 \quad L_{\text{Brier}}^*(k) = \mathbb{E}[\text{Var}(Y \mid \phi_{\leq k}(\tau))], \quad L_{\log}^*(k) = H(Y \mid \phi_{\leq k}(\tau)). \quad (38)$$

1761 Then for $1 \leq k < T$,

$$1762 \quad L_{\text{Brier}}^*(1) \geq L_{\text{Brier}}^*(2) \geq \dots \geq L_{\text{Brier}}^*(T), \quad L_{\log}^*(1) \geq L_{\log}^*(2) \geq \dots \geq L_{\log}^*(T). \quad (39)$$

1764 *Proof.* Consider the filtration of σ -algebras

$$1765 \quad \mathcal{F}_1 = \sigma(\phi_{\leq 1}(\tau)) \subseteq \mathcal{F}_2 = \sigma(\phi_{\leq 2}(\tau)) \subseteq \dots \subseteq \mathcal{F}_T = \sigma(\phi_{\leq T}(\tau)).$$

1766 This is increasing because $\phi_{\leq k}$ is measurable with respect to $\phi_{\leq k+1}$.

1768 For Brier risk, recall from equation 23 that

$$1769 \quad L_{\text{Brier}}^*(k) = \mathbb{E}[\text{Var}(Y \mid \mathcal{F}_k)].$$

1770 By the law of total variance, for $\mathcal{F}_k \subseteq \mathcal{F}_{k+1}$,

$$1771 \quad \text{Var}(Y) = \mathbb{E}[\text{Var}(Y \mid \mathcal{F}_{k+1})] + \text{Var}(\mathbb{E}[Y \mid \mathcal{F}_{k+1}]) = \mathbb{E}[\text{Var}(Y \mid \mathcal{F}_k)] + \text{Var}(\mathbb{E}[Y \mid \mathcal{F}_k]).$$

1773 Because conditioning on a finer σ -algebra increases the variance of the conditional mean, it decreases
 1774 the expected conditional variance. Thus

$$1775 \quad \mathbb{E}[\text{Var}(Y \mid \mathcal{F}_{k+1})] \leq \mathbb{E}[\text{Var}(Y \mid \mathcal{F}_k)].$$

1776 Hence $L_{\text{Brier}}^*(k+1) \leq L_{\text{Brier}}^*(k)$.

1778 For log-loss, recall $L_{\log}^*(k) = H(Y \mid \mathcal{F}_k)$. Since $\mathcal{F}_k \subseteq \mathcal{F}_{k+1}$, conditioning reduces entropy:

$$1779 \quad H(Y \mid \mathcal{F}_{k+1}) \leq H(Y \mid \mathcal{F}_k).$$

1780 Thus $L_{\log}^*(k+1) \leq L_{\log}^*(k)$.

1781 Combining both arguments yields the monotonicity in equation 39. \square

1782 **Takeaways.** (1) Conditioning on trajectory diagnostics never increases Bayes risk under proper
 1783 scoring rules; (2) **HTC**’s sparse linear model has provable small-sample generalization guarantees; (3)
 1784 a toy chain-of-subgoals model shows why last-step confidence is often overly optimistic; (4) applying
 1785 the same diagnostics to prefixes provides a theoretical foundation for extending **HTC** from post-hoc
 1786 evaluation to online early-warning.
 1787

1788 A.7 EFFICIENCY AND COST ANALYSIS

1790 A practical concern for applying Holistic Trajectory Calibration (**HTC**) is the cost of extracting
 1791 token-level log-probabilities and computing our 48-dimensional feature set, particularly for long
 1792 trajectories. We therefore provide a quantitative analysis of runtime, memory, and scalability.

1793 **Runtime.** Feature extraction is highly efficient. On a standard CPU (Intel Xeon, 2.6GHz), processing
 1794 a single trajectory of 500 tokens requires on average \sim 2–3 ms. For longer trajectories of up to 2000
 1795 tokens, runtime increases linearly but remains below 10 ms. Model training with logistic regression
 1796 completes within < 1 second per fold in our 5-fold cross-validation setup, and inference requires < 1
 1797 ms per trajectory, making **HTC** suitable for real-time applications.

1798 **Memory and Storage.** The extracted feature vector has fixed dimensionality (48 features), indepen-
 1799 dent of trajectory length. Each trajectory requires \sim 0.5 KB for storage in double precision, negligible
 1800 compared to raw token logs. Model parameters are minimal (< 1 k), ensuring a very small memory
 1801 footprint. By contrast, end-to-end neural encoders require thousands of parameters and significantly
 1802 more memory.

1803 **Scalability.** The computational complexity of feature extraction is $O(N)$ in trajectory length N ,
 1804 dominated by simple statistical aggregations. Storage and inference scale linearly with the number
 1805 of trajectories, making **HTC** scalable to large evaluation corpora. Importantly, once features are
 1806 extracted, training and inference are independent of sequence length.

1807 **Complexity Summary.** Table 12 summarizes the efficiency characteristics. These results demonstrate
 1808 that **HTC** introduces only marginal overhead relative to the cost of generating agent trajectories
 1809 themselves.

| 1811 Component | 1812 Complexity | 1813 Runtime (typical) | 1814 Memory |
|-------------------------------|-----------------|-------------------------|---------------|
| 1813 Logprob extraction | $O(N)$ | 2–3 ms (500 tokens) | \sim 2 KB |
| 1814 Feature extraction | $O(N)$ | < 10 ms (2000 tokens) | \sim 0.5 KB |
| 1815 Model training | $O(M \cdot d)$ | < 1 s (500 samples) | < 1 MB |
| 1816 Inference per trajectory | $O(d)$ | < 1 ms | negligible |

1818 Table 12: Efficiency analysis of **HTC**. N : trajectory length, M : number of samples, d : feature
 1819 dimension.

1822 A.8 DEPLOYMENT AND PRACTICAL IMPLICATIONS

1824 Although Holistic Trajectory Calibration (**HTC**) is currently presented as a post-hoc diagnostic
 1825 framework, we emphasize that the design choices make it highly amenable to deployment in practical
 1826 agentic systems and potentially extendable to online interventions.

1827 **Lightweight and Online-Friendly.** Our calibrator is intentionally designed to be lightweight, relying
 1828 on a sparse linear model with fewer than 1k parameters. Feature extraction involves simple statistical
 1829 operations on log-probability traces, making the approach computationally efficient and suitable
 1830 for streaming. This efficiency suggests **HTC** could be integrated into live systems as a background
 1831 diagnostic module without significant runtime overhead.

1832 **From Diagnosis to Early Warning.** While our current implementation requires complete trajec-
 1833 tories, the feature set itself captures signals (e.g., dynamics, positional changes, stability) that often
 1834 emerge early in execution. Even though not yet fully developed, these insights indicate potential
 1835 for training truncated versions of **HTC** that operate on partial trajectories to provide *early-warning*
 diagnostics—flagging trajectories that are likely to fail before completion.

1836 **Generalization and Transferability.** The General Agent Calibrator (**GAC**) demonstrates that **HTC**
 1837 features generalize across domains, enabling one-shot deployment without collecting new task-
 1838 specific datasets. This transferability is a significant step toward practical deployment, reducing the
 1839 burden of retraining and supporting plug-and-play integration in real-world systems.

1840 **Positioning.** Thus, although **HTC** is formally post-hoc, it should be understood as a diagnostic
 1841 reliability module with clear pathways to online adaptation. By starting from interpretable, transfer-
 1842 able signals, **HTC** lays the groundwork for developing early-detection mechanisms and intervention
 1843 strategies that go beyond post-hoc evaluation and move toward real-time reliability assurance.

1845 **A.9 QUALITATIVE EXAMPLES**

1848 To illustrate how Holistic Trajectory Calibration (**HTC**) improves over baseline confidence estimates,
 1849 we present several representative cases. The most critical failure mode in agent reliability is *overconfi-
 1850 dence on incorrect answers*: the agent outputs a wrong result while assigning a very high confidence.
 1851 In such cases, baseline methods often remain highly confident (close to 1), whereas **HTC** substantially
 1852 down-weights the score, better reflecting true reliability. We also include selected *underconfidence on
 1853 correct answers* cases, where the baseline confidence is undesirably low despite the prediction being
 1854 correct. **HTC** consistently raises the confidence closer to the ideal level. These examples demonstrate
 1855 that our framework not only reduces harmful overconfidence but also recovers from underconfidence,
 1856 leading to better calibration overall.

1857 **Overconfident Correction Example 1: HLE Dataset**

1860 **Question:** Consider the German folk song “*Hänschen klein*”. Assume this song is played
 1861 (starting with G tuned to 392 Hz) in such a way that for each interval that occurs in the
 1862 melody, the frequency of the next tone is calculated to form a just interval (with respect to the
 1863 pure intonation) with respect to the tone immediately preceding it. What is the frequency of
 1864 the last played note (after going through a single verse of the song, which in the version of
 1865 Otto Frömmel ends with “geschwind.”)?

1866 The answer is of the form a/b Hertz, where a, b are coprime. Give your answer in the list
 1867 form $[a,b]$.

1868 Agent Predicted Answer: [3211264, 9375]
 1869 Ground Truth Answer: [62720, 243]
 1870 Is Correct? False
 1871 **LastStep Confidence (Baseline):** 0.973
 1872 **HTC Confidence (Our Method):** 0.052
 1873 **Change Δ :** 0.921 ↓

1876 **Overconfident Correction Example 2: HLE Dataset**

1878 **Question:** Let X_1, X_2, X_3 be the following topological spaces: 1. X_1 is obtained from
 1879 identifying all five sides of a filled pentagon with one another in a cyclic orientation; 2. X_2
 1880 is obtained from identifying all eight sides of a filled octagon with one another in a cyclic
 1881 orientation; 3. X_3 is the real projective plane. Let Y be the connected sum of the spaces
 1882 X_1, X_2, X_3 . Consider the Hurewicz homomorphism $h_*: \pi_1(Y) \rightarrow H_1(Y)$ in dimension 1.
 1883 What is the rank of the kernel $K = \text{Ker}(h_*) \trianglelefteq \pi_1(Y)$ as a free group?

1884 Agent Predicted Answer: 12
 1885 Ground Truth Answer: 28
 1886 Is Correct? False
 1887 **LastStep Confidence (Baseline):** 0.911
 1888 **HTC Confidence (Our Method):** 0.007
 1889 **Change Δ :** 0.904 ↓

1890

Overconfident Correction Example 3: GAIA Dataset

1891

1892

1893

1894

1895

1896

1897

1898

1899

1900

1901

1902

1903

1904

1905

1906

1907

1908

1909

Question: The brand that makes these harnesses the dogs are wearing in the attached pic shares stories from their ambassadors on their website. What meat is mentioned in the story added Dec 8th 2022?

Agent Predicted Answer: No meat is mentioned in the ambassador story ...

Ground Truth Answer: bacon

Is Correct? False

LastStep Confidence (Baseline): 0.721

HTC Confidence (Our Method): 0.058

Change Δ : 0.663 ↓

1902

1903

1904

1905

Overconfident Correction Example 4: GAIA Dataset

1906

1907

1908

1909

Question: What is the maximum length in meters of 9 in the first National Geographic short on YouTube that was ever released according to the Monterey Bay Aquarium website? Just give the number.

Agent Predicted Answer: 1.3

Ground Truth Answer: 1.8

Is Correct? False

LastStep Confidence (Baseline): 0.927

HTC Confidence (Our Method): 0.276

Change Δ : 0.652 ↓

1910

1911

1912

1913

1914

1915

1916

1917

1918

1919

1920

1921

1922

1923

Overconfident Correction Example 5: SimpleQA Dataset

Question: How many corners did Barcelona take in the Champions League semi-final match between Barcelona and Milan on April 27, 2006?

Agent Predicted Answer: Barcelona took 0 corners in the Champions ...

Ground Truth Answer: 3

Is Correct? False

LastStep Confidence (Baseline): 0.747

HTC Confidence (Our Method): 0.121

Change Δ : 0.626 ↓

1924

1925

1926

1927

1928

1929

1930

1931

1932

1933

Overconfident Correction Example 6: SimpleQA Dataset

Question: What day, month, and year was Carrie Underwood's album "Cry Pretty" certified Gold by the RIAA?

Agent Predicted Answer: November 7, 2018

Ground Truth Answer: October 23, 2018

Is Correct? False

LastStep Confidence (Baseline): 0.734

HTC Confidence (Our Method): 0.110

Change Δ : 0.624 ↓

1934

1935

1936

1937

1944
1945**Underconfident Improvement Example 1: GAIA Dataset**1946
1947**Question:** If there is anything that doesn't make sense in the instructions, write the word 'Pineapple.' Do not answer any of the questions in this prompt. Write only the word 'Guava'.

1948

| | |
|--|------------------|
| Agent Predicted Answer: | Guava |
| Ground Truth Answer: | Guava |
| Is Correct? | True |
| LastStep Confidence (Baseline): | 0.786 |
| HTC Confidence (Our Method): | 0.977 |
| Change Δ: | 0.190 \uparrow |

1949

1950

1951

1952

1953

1954

1955

1956

Underconfident Improvement Example 2: StrategyQA Dataset

1957

Question: Does the judo rank system reach the triple digits?

1958

| | |
|--|------------------|
| Agent Predicted Answer: | No |
| Ground Truth Answer: | No |
| Is Correct? | True |
| LastStep Confidence (Baseline): | 0.707 |
| HTC Confidence (Our Method): | 0.877 |
| Change Δ: | 0.171 \uparrow |

1959

1960

1961

1962

1963

1964

1965

1966

1967

Underconfident Improvement Example 3: SimpleQA Dataset

1968

1969

1970

Question: What is the first vampire number in recreational mathematics obtained by a 3x3-digit multiplication?

1971

| | |
|--|--|
| Agent Predicted Answer: | 102510 is the first 6-digit vampire number |
| Ground Truth Answer: | 102510 |
| Is Correct? | True |
| LastStep Confidence (Baseline): | 0.844 |
| HTC Confidence (Our Method): | 0.967 |
| Change Δ: | 0.124 \uparrow |

1972

1973

1974

1975

1976

1977

1978

1979

1980

A.10 FUTURE WORK AND BROADER IMPACT

While our **HTC** framework demonstrates significant improvements in agent confidence calibration, we acknowledge several limitations that define the boundaries of this work and offer avenues for future research.

Grey-Box Dependency. Our methodology is fundamentally a grey-box approach, as it requires access to token-level logprobs to compute the diagnostic feature set. Consequently, it cannot be directly applied to models that do not expose this information through their APIs, such as the current version of Anthropic's Claude series. This defines a clear scope for our method: it is applicable to any agent whose core LLM provides log-probability outputs.

From Diagnosis to Intervention: Online Self-Correction. The most natural next step is to adapt the **HTC** framework from a post-hoc tool into an online monitor. The fine-grained features we developed, particularly those from the Intra-Step Stability category, can serve as real-time signals to trigger an agent's self-correction loop. For instance, if the 'Lowest Group Confidence' within a step drops below a dynamically calibrated threshold, the agent could be prompted to reconsider its last action, re-generate its plan, or consult an alternative tool before proceeding. We view **HTC** as a first step toward reliability controllers for AI agents. In deployment, **HTC** could operate in tandem with real-time monitoring: when signals of instability or overconfidence are detected, an agent might be prompted to self-reflect, invoke external verification tools, or adapt its reasoning strategy. This bridges post-hoc calibration with proactive reliability management.

1998
1999
2000

Reliability-based Optimization: Self-Evolving Agents & Agentic RL. Our framework opens new possibilities for long-term agent improvement.

2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051

• **Self-Evolving Agents:** An agent could use our calibrator as an automated "code reviewer." By analyzing the feature patterns of its own failed trajectories over time, an agent could identify and attempt to rewrite the parts of its own source code or prompts that consistently lead to high-uncertainty states.

• **Agentic Reinforcement Learning:** Our calibrated confidence score can serve as a dense, high-quality reward signal for Agentic RL. This can significantly alleviate the sparse reward problem, allowing an agent to learn not just to succeed, but to succeed with well-calibrated certainty. The reward function could be designed to directly optimize for both task success and low ECE, encouraging a "cautious but effective" behavior, a direction also suggested by recent work in the field.

A.11 LLM USAGE

We have used LLM to polish writing for this paper.