

# GAICo: A Deployed and Extensible Framework for Evaluating Diverse and Multimodal Generative AI Outputs

Nitin Gupta<sup>1\*</sup>, Pallav Koppiseti<sup>1</sup>, Kausik Lakkaraju<sup>1</sup>, Biplav Srivastava<sup>1</sup>

<sup>1</sup>University of South Carolina  
Columbia, South Carolina, USA

## Abstract

The rapid proliferation of Generative AI (GenAI) into diverse, high-stakes domains necessitates robust and reproducible evaluation methods. However, practitioners often resort to ad-hoc, non-standardized scripts, as common metrics are often unsuitable for specialized, structured outputs (e.g., automated plans, time-series) or holistic comparison across modalities (e.g., text, audio, and image). This fragmentation hinders comparability and slows AI system development. To address this challenge, we present GAICo (Generative AI Comparator): a deployed, open-source Python library that streamlines and standardizes GenAI output comparison. GAICo provides a unified, extensible framework supporting a comprehensive suite of reference-based metrics for unstructured text, specialized structured data formats, and multimedia (images, audio). Its architecture features a high-level API for rapid, end-to-end analysis, from multi-model comparison to visualization and reporting, alongside direct metric access for granular control. We demonstrate GAICo’s utility through a detailed case study evaluating and debugging complex, multi-modal AI Travel Assistant pipelines. GAICo empowers AI researchers and developers to efficiently assess system performance, make evaluation reproducible, improve development velocity, and ultimately build more trustworthy AI systems, aligning with the goal of moving faster and safer in AI deployment. **Since its release on PyPI in Jun 2025, the tool has been downloaded over 16K times, across versions, by Dec 2025, demonstrating growing community interest.**

**Deployed Library** — [pypi.org/project/GAICo](https://pypi.org/project/GAICo)

**Code** — [github.com/ai4society/GenAIResultsComparator](https://github.com/ai4society/GenAIResultsComparator)

**Documentation** —

[ai4society.github.io/projects/GenAIResultsComparators](https://ai4society.github.io/projects/GenAIResultsComparators)

**Extended Version** — [doi.org/10.48550/arXiv.2508.16753](https://doi.org/10.48550/arXiv.2508.16753)

## Introduction

The adoption of Generative AI (GenAI), particularly Large Language Models (LLMs), has been transformative across industries. However, this rapid integration has outpaced the development of standardized, accessible, and domain-appropriate evaluation tools. While standard Natural Language Processing (NLP) metrics like BLEU (Papineni et al.

\*Contact authors: [ningt@email.sc.edu](mailto:ningt@email.sc.edu) & [biplav.s@sc.edu](mailto:biplav.s@sc.edu).  
Copyright © 2026, Association for the Advancement of Artificial Intelligence ([www.aaai.org](http://www.aaai.org)). All rights reserved.

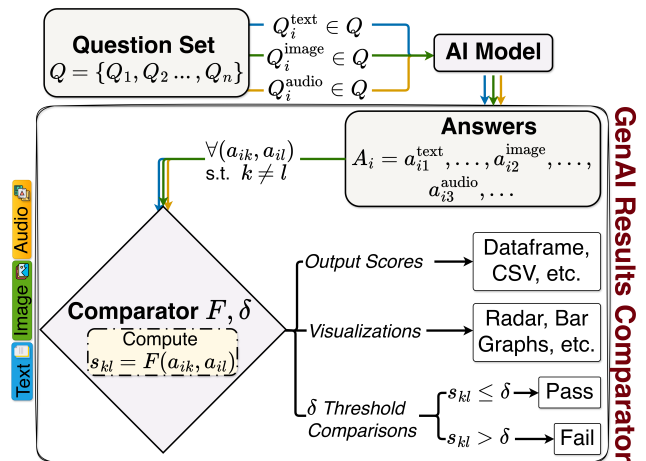


Figure 1: The multi-modal GAICo workflow. The framework processes answers from multi-modal (text, image, audio) AI models, computes pairwise similarity scores ( $s_{kl}$ ), and constructs several outputs: raw data reports, visualizations, and pass/fail assessments against a threshold  $\delta$ . (any distance function, or conversely,  $1 - \text{similarity metric}$ ).

2002) and ROUGE (Lin 2004) are useful, they are often insufficient for structured outputs like AI plans or time-series forecasts. In the absence of a unified tool, developers create bespoke evaluation scripts, leading to a fragmented and non-reproducible landscape that hinders the comparison and iterative improvement of AI systems.

In response, we introduce GAICo (Generative AI Comparator), a deployed and extensible open-source Python library built to address these challenges. GAICo provides a unifying framework that makes the application of diverse, reference-based metrics simple and reproducible across different data types (modalities). By providing a consistent interface, GAICo directly supports the iterative workflows common in applied AI development. According to a PyPI package statistics source (PyPI Stats 2025), the library has been downloaded over 16,000 times in just over 6 months (Jun-Dec 2025), demonstrating its widespread and growing community interest.

Common scenarios for usage of GAICo are evaluating

an LLM’s output along a range of metrics (Figure 1) and output of multiple LLMs across a single metric (notebooks ‘example-1’ and ‘example-2’ in Table 2). Consider the practical challenge of building a significantly more complex composite “AI Travel Agent” that uses an orchestrator LLM to generate an itinerary and specialist models to create corresponding images and audio summaries (System Pipelines in Figure 2). Evaluating different combinations of these components requires writing separate, ad-hoc scripts for each data modality. This process is slow, error-prone, and makes it difficult to pinpoint the root cause of a poor output: was it a flawed prompt from the orchestrator or poor execution by a specialist model? This exemplifies the need for a unified tool that can efficiently handle complex, multi-modal comparisons.

Our contributions, aligned with the “Tools and Methodologies for Moving Faster and Safer” track, are: (1) **A Unified and Extensible Framework:** A general framework for comparing GenAI outputs, with an object-oriented design that allows seamless integration of metrics for text, structured data (AI plans, time-series), and multimedia. (2) **A Streamlined Practitioner Workflow:** The `Experiment` class, which simplifies the end-to-end evaluation of complex AI systems by abstracting score calculation, visualization, and reporting into a few lines of code. (3) **Demonstrated Utility for Debugging Composite AI:** A case study on AI Travel Assistants (Figure 2) showcasing GAICo’s effectiveness in performing nuanced, multi-stage evaluations to isolate performance issues. (4) **A Deployed and Maintained Tool:** A public tool on PyPI with comprehensive documentation, optional dependencies to manage installation size, and a robust testing suite to lower the barrier to adoption.

## Related Works

### General-Purpose AI/NLP Evaluation Libraries

General-purpose AI and Natural Language Processing (NLP) evaluation libraries provide foundational metrics for various tasks. For instance, the Hugging Face `evaluate` library (Wolf et al. 2019) offers a wide array of metrics for NLP and related tasks, providing a standardized `compute` method for individual metrics and integrating well within the Hugging Face ecosystem. Similarly, Scikit-learn (Pedregosa et al. 2011) provides fundamental metrics for classification, regression, and clustering, adaptable for certain aspects of text evaluation (e.g., cosine similarity on TF-IDF vectors). Libraries like NLTK (Bird, Klein, and Loper 2009) and SpaCy (Honnibal et al. 2020) offer foundational text processing capabilities and some basic metrics (e.g., BLEU in NLTK).

While indispensable, these libraries primarily offer individual text metrics and lack a unified framework for evaluating diverse, multi-modal GenAI outputs or streamlining multi-model comparisons. This fragmentation forces developers to manually integrate disparate tools. GAICo directly addresses this issue by providing a single, extensible API that unifies evaluation across diverse data types, significantly reducing development overhead and improving reproducibility.

A key differentiator for GAICo is its design choice to *decouple evaluation from LLM inference*, positioning it as a post-hoc, reproducible comparison framework. This contrasts with general-purpose libraries like Hugging Face `evaluate` by offering a more integrated workflow and specialized metrics, and with end-to-end frameworks, e.g., Ragas (Es et al. 2024), DeepEval (Ip and Vongthongsri 2025), that are tightly coupled with LLM APIs for “LLM-as-a-judge” evaluations. While these integrated frameworks are powerful, their coupling can introduce complexities like API costs, rate-limits, and nondeterminism. GAICo’s focus on pre-generated outputs simplifies the evaluation pipeline, making it faster and more reliable for developers focused on comparative analysis. The `Experiment` class, discussed in Section 3, further streamlines this by offering a high-level abstraction for multi-model comparison, plotting, and reporting. A detailed comparison of these approaches is provided in the extended version (Gupta et al. 2025).

### Domain-Specific Evaluation Approaches

Evaluation in specialized AI domains often relies on custom tools, leading to fragmented practices. GAICo bridges this gap by integrating relevant metrics into a single, consistent framework.

**Automated Planning** Evaluation in automated planning typically focuses on properties like plan length, cost, or state-space coverage, often requiring specialized planners or validators (e.g., PDDL parsers). However, when LLMs generate planning sequences, the output might not strictly adhere to formal planning languages, necessitating metrics that compare the *content and order of actions* rather than strict plan validity. GAICo addresses this with `PlanningLCS` and `PlanningJaccard` metrics, which are designed to robustly compare action sequences, including concurrent actions, directly from LLM outputs. These are inspired by foundational work on measuring plan diversity and similarity (Srivastava et al. 2006, 2007) and sequence analysis (Hirschberg 1977), adapted specifically for GenAI evaluation contexts.

**Time-Series Analysis** Standard time-series evaluation metrics include Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), or R-squared for forecasting accuracy; however, when comparing generated time-series *shapes or sequences of values* from LLMs, metrics like Dynamic Time Warping (DTW) (Berndt and Clifford 1994) become crucial for handling variations in speed or phase. GAICo integrates `TimeSeriesDTW` and `TimeSeriesElementDiff` to provide robust, normalized comparisons of time-series data, regardless of their textual representation.

**Multimedia (Image/Audio)** Evaluation of generated images and audio often involves specialized libraries and metrics. For images, common metrics include Structural Similarity Index (SSIM) (Wang et al. 2004), Peak Signal-to-Noise Ratio (PSNR) (Sara, Akter, and Uddin 2019), or perceptual hashes (Farid 2021). For audio, metrics like Signal-to-Noise Ratio (SNR) (Yuan et al. 2019), Perceptual Evalu-

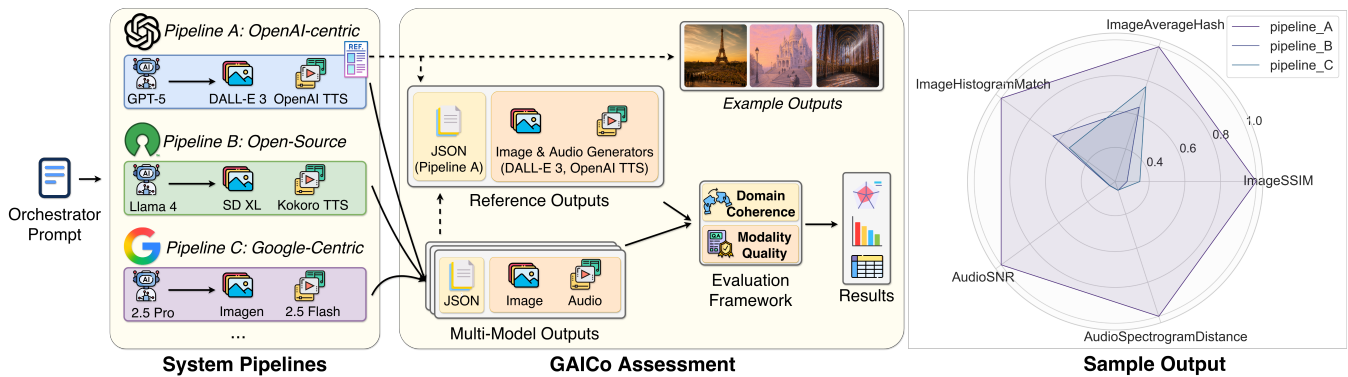


Figure 2: Illustration of GAICo for a composite AI system case study. (Right) A user uses LLMs to get results tailored to their needs and then analyzes them with GAICo. They start the process in three parallel pipelines that generate multi-modal outputs (JSON plan, image, audio). The user then performs a two-part evaluation: (a) **Plan Coherence**, comparing the JSON outputs of all pipelines against a single *Baseline Plan Reference* (derived from Pipeline A’s output) and (b) **Modality Quality**, comparing each pipeline’s generated image and audio against *per-pipeline references*. These per-pipeline references are generated by feeding the *same prompts/scripts* from the respective pipeline’s JSON output into baseline specialist generators (from Pipeline A). (Left) The output radar plot shows image and audio fidelity relative to references derived from each pipeline’s prompts. For more examples, see Table 2.

ation of Speech Quality (PESQ) (Rix et al. 2001), or Short-Time Objective Intelligibility (STOI) (Taal et al. 2010) are used. While underlying libraries (e.g., scikit-image (Van der Walt et al. 2014), Pillow (Clark 2015), librosa (McFee et al. 2015), scipy (Virtanen et al. 2020)) provide these functionalities, GAICo unifies their application within its framework. The rise of multi-modal generative models, such as Meta’s ImageBind (Girdhar et al. 2023), further underscores the need for unified evaluation tools capable of assessing outputs across diverse modalities. This allows for consistent, normalized comparison of multimedia outputs alongside text and structured data, bridging the gap between general LLM evaluation and highly specialized domain-specific analysis.

GAICo’s contribution lies in integrating these diverse, domain-specific metrics into a single, extensible library, accelerating development and improving the reliability of composite AI systems.

## The GAICo Framework

The GAICo framework is architected to provide a robust, extensible, and user-friendly solution for evaluating GenAI outputs. Its design balances a high-level, streamlined workflow with the flexibility required for granular analysis, centered on three components: an extensible `BaseMetric` class, a comprehensive metric library, and a high-level `Experiment` class. Detailed descriptions of each metric and a UML class diagram for GAICo are available in the extended version (Gupta et al. 2025).

### Core Concepts: BaseMetric and Extensibility

At the heart of GAICo’s architecture is the `BaseMetric` abstract class, which establishes a consistent foundation for all metrics. It mandates that every metric, regardless of data modality, implements a sin-

gle core method: `calculate(generated_texts, reference_texts)`. This simple interface is key to the framework’s flexibility, transparently handling various input formats (single items, lists, NumPy arrays) for efficient batch processing. This architecture makes GAICo highly extensible: a developer can introduce a new metric by simply inheriting from `BaseMetric` and implementing the `calculate()` method, immediately integrating it into the GAICo ecosystem.

### Comprehensive Metric Library

GAICo provides a carefully designed library of metrics spanning textual analysis, structured data, and multimedia (audio and image). This comprehensive approach addresses the reality that modern GenAI applications produce complex, multi-modal outputs that traditional text-only tools cannot adequately assess. The library includes n-gram-based (BLEU, ROUGE), text similarity (Jaccard, Cosine), and semantic (BERTScore) metrics. For structured data, it offers specialized metrics for automated planning sequences (`PlanningLCS`) and time-series (`TimeSeriesDTW`). For multimedia, it integrates established metrics for image (SSIM, PSNR) and audio (SNR, Spectrogram Distance) quality.

### The Streamlined Experiment Class

The `Experiment` class is a high-level API designed to streamline the entire evaluation workflow. It orchestrates the common task of comparing multiple model outputs against a single reference, encapsulating score calculation, visualization, and reporting into a few lines of code. Its `compare()` method automates the generation of scores, publication-ready plots (bar or radar), application of quality thresholds, and export of comprehensive CSV reports. By abstracting away boilerplate code, the `Experiment` class directly

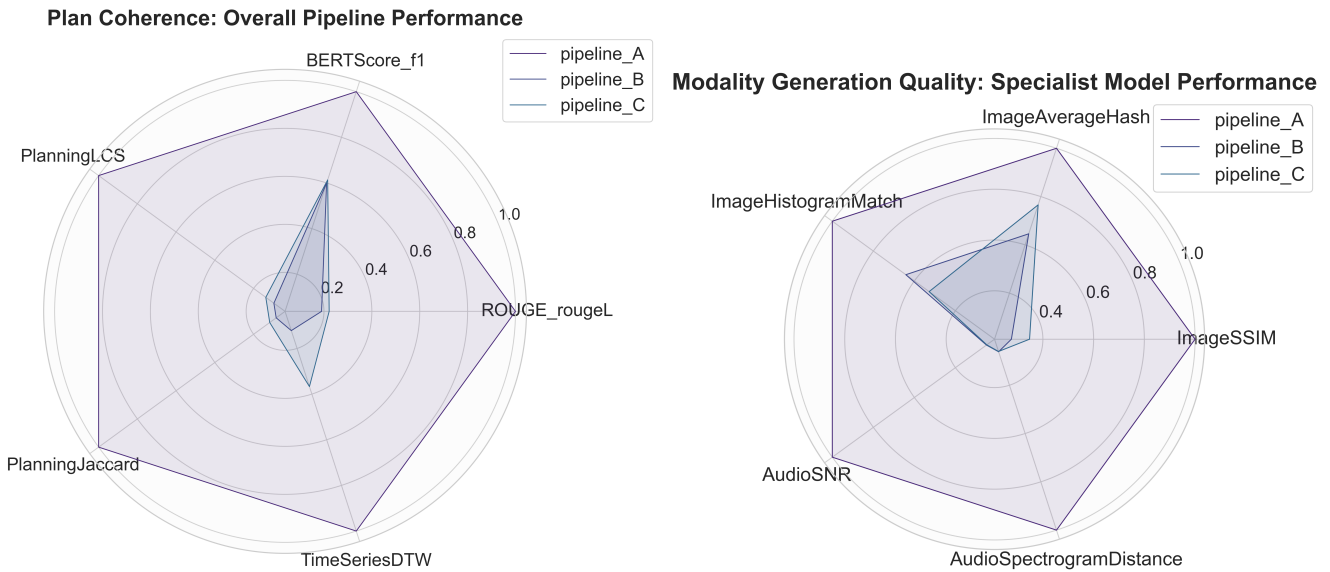


Figure 3: Radar plots generated by GAICo comparing pipeline performance across various metrics. (Right) Modality Generation Quality, assessing the specialist models’ fidelity against references generated from *their own pipeline’s prompts*. Each axis is a metric, and each line represents a pipeline’s averaged score. (Left) Plan Coherence, showing the relative strengths of orchestrator LLMs against a universal human-curated reference.

improves developer productivity and enables rapid, reproducible benchmarking.

### Deployment and Maintainability

GAICo is a professionally engineered tool deployed on the Python Package Index (PyPI) for easy installation via `pip`. To balance functionality with accessibility, it uses optional dependencies (e.g., `pip install 'gaico[bertscore]'`), minimizing the installation footprint. The project’s quality is ensured through a comprehensive test suite (pytest (Krekel et al. 2004)), continuous integration (CI) for automated testing and deployment, and pre-commit hooks (Sottile 2014) for code consistency. Extensive documentation (MkDocs (Christie 2014)), including runnable Jupyter notebooks and a full API reference, lowers the barrier to adoption for the applied AI community.

### Case Study: Evaluating Composite AI Travel Assistants

To demonstrate GAICo’s utility in a realistic scenario, we evaluated three distinct, composite AI systems tasked with creating a multi-modal travel itinerary. This study highlights how GAICo enables a nuanced evaluation crucial for debugging and improving such systems. An overview of the case study is demonstrated in Figure 2. The code for the case study can be found in notebook number 7, ‘case\_study.ipynb’, in Table 2.

### Setup and Methodology

Our goal was to assess AI systems built from different combinations of state-of-the-art generative models. Each system,

or “pipeline”, consists of an orchestrator LLM that generates a travel plan in JSON format, which in turn provides prompts for downstream, specialized models for image and audio generation.

**Master Orchestrator Prompt** All orchestrator LLMs received the same detailed prompt, designed to elicit a 3-day Paris itinerary in a structured JSON format, including fields for planning sequences, time-series, image prompts, and audio prompts (latter two for the specialist models). The full prompt is included in the extended version (Gupta et al. 2025).

We constructed three pipelines using different combinations of orchestrator LLMs and specialist image/audio models (Pipeline A: OpenAI-centric, Pipeline B: Open-Source, Pipeline C: Google-centric). Each pipeline received an identical master prompt to generate a 3-day Paris itinerary in a structured JSON format, which included prompts for the downstream models.

A key challenge in composite systems is attributing failure. To address this, we designed a two-part evaluation strategy, seamlessly executed using GAICo: **(1) Plan Coherence:** We evaluated the core travel planning ability of the orchestrator LLMs by comparing the generated JSON content (text, plans, budgets) from all pipelines against a single high-quality reference plan (from Pipeline A). **(2) Modality Generation Quality:** We evaluated the fidelity of the specialist models by comparing each pipeline’s generated image and audio against a reference created using that pipeline’s *own* generated prompts. This isolates the specialist model’s execution quality from the orchestrator’s prompt-generation quality.

Pipeline	Plan Coherence Metrics					Modality Generation Metrics				
	ROUGE-L <sub>T</sub>	BERTScore-F1 <sub>T</sub>	LCS <sub>Pl</sub>	Jaccard <sub>Pl</sub>	DTW <sub>TS</sub>	SSIM <sub>I</sub>	AvgHash <sub>I</sub>	HistMatch <sub>I</sub>	SNR <sub>A</sub>	SpecDist <sub>A</sub>
Pipeline A	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
Pipeline B	0.190	0.599	0.095	0.083	0.122	0.276	0.646	<b>0.642</b>	<b>0.249</b>	<b>0.261</b>
Pipeline C	<b>0.222</b>	<b>0.613</b>	<b>0.137</b>	<b>0.117</b>	<b>0.367</b>	<b>0.347</b>	<b>0.766</b>	0.528	0.247	0.260

Table 1: Combined evaluation results. Plan coherence metrics (*left*) and modality generation metrics (*right*) are reported for each pipeline. Subscripts denote evaluation domains:  $T$  = text,  $Pl$  = planning,  $TS$  = time series,  $I$  = image, and  $A$  = audio. Scores are averaged over three days; higher scores between Pipeline B and Pipeline C are highlighted.

## Results and Analysis

The quantitative results, summarized in Table 1 and visualized in Figure 3, demonstrate GAICo’s value in accelerating AI development and improving system safety.

**Orchestrator Performance (Plan Coherence)** As shown in the left side of Table 1, Pipeline A’s orchestrator (GPT-5), serving as the high-quality baseline, expectedly achieved perfect scores (1.000) across all plan coherence metrics. In contrast, both Pipeline B (Llama 4) and Pipeline C (Gemini 2.5 Pro) showed significantly lower performance, indicating substantial deviations from the baseline plan. As shown in the radar plot in Figure 3 (left), Pipeline C generally outperformed Pipeline B across these metrics, suggesting a better ability to approximate the reference budget forecast.

**Specialist Model Performance (Modality Generation Quality)** The modality generation quality results (right of Table 1) reveal the execution fidelity of the specialist models against their respective high-quality references. Pipeline A’s specialist models achieved perfect scores (1.000) as their outputs were used to define the modality-specific references. For image generation, Pipeline C generally outperformed Pipeline B across Image SSIM and Image AverageHash, indicating stronger structural and perceptual similarity to the baseline. However, Pipeline B showed a slightly higher Image HistogramMatch score, suggesting its color distribution was closer to the baseline. This highlights the value of multi-metric evaluation for a nuanced understanding.

Focusing on the audio metrics, both Pipeline B and Pipeline C produced audio summaries with notably lower Audio SNR and Audio SpectrogramDist compared to Pipeline A. This indicates that while their orchestrators might provide valid scripts, their specialist audio models are not yet capable of matching the fidelity of the baseline OpenAI TTS. Figure 3 (right) visually represents these performance differences.

**Analysis** This case study vividly demonstrates how GAICo facilitates “moving faster and safer” in the development of complex AI systems.

**Faster Development Cycles** Without GAICo, this multi-modal evaluation would require numerous, disparate scripts for JSON parsing, text analysis, plan validation, and multimedia comparison, followed by manual aggregation. GAICo streamlined this entire process into a single, unified workflow. This enables rapid iteration, allowing developers to

swap components (e.g., an orchestrator LLM) and immediately receive quantitative feedback on the impact across all modalities, **drastically reducing the time and maintenance costs** associated with benchmarking.

**Improved Reliability and Safer Deployment** The two-part evaluation, enabled by GAICo, is critical for enhancing system **reliability**. By separating “Plan Coherence” from “Modality Generation Quality”, we can precisely diagnose failures. For instance, the results show that Pipeline B’s poor performance stems from both a weak orchestrator (low scores in Table 1 left) and less-than-optimal specialist models (Table 1 right). This granular insight allows for targeted optimization, improving the orchestrator’s planning logic or replacing a specific specialist model, which is essential for building trustworthy systems and preventing incidents caused by underperforming components. This directly contributes to the goal of producing **better AI solutions** by enabling a more rigorous and efficient development process.

## Demonstrated Use Cases and Documentation

To maximize ease of use and adoption, key factors in a tool’s real-world impact, GAICo is supported by extensive documentation and a comprehensive suite of example notebooks. These resources demonstrate the library’s versatility across a range of applications, from simple quick-start guides to advanced, domain-specific evaluations in finance, planning, and multimedia. Table 2 summarizes the available examples, which are all publicly accessible in our code repository.

## Conclusion

In this section, we highlight the unique role of GAICo in evaluating LLM outputs in the context of prior work. Then, we discuss how it fits into the role of “moving faster and safer” in AI system development.

## GAICo and Its Novelty

GAICo addresses a critical gap in GenAI evaluation by providing the first unified, extensible, and deployed framework for **reproducible** assessment across diverse data modalities. Recent studies have found that even with access to the original source code, many ML experiments cannot be reproduced, mainly due to missing details in documentation and inconsistent evaluation procedures (Lopresti and Nagy 2021). GAICo replaces this current landscape of ad-hoc, non-comparable scripts with a standardized approach, integrating metrics for text, structured data (planning, time-series), and multimedia. This standardization is GAICo’s

Notebook Name	Description
<b>Quickstart Examples</b>	
01. <i>case_study.ipynb</i>	Full code for the composite AI travel assistant evaluation presented in this paper.
02. <i>example-1.ipynb</i>	Evaluating multiple models using a single metric via the <code>calculate()</code> method.
03. <i>example-2.ipynb</i>	Evaluating a single model across multiple metrics using their <code>calculate()</code> methods.
04. <i>example-audio.ipynb</i>	Evaluating AI-generated audio using specialized audio metrics.
05. <i>example-image.ipynb</i>	Evaluating AI-generated images using specialized image metrics.
06. <i>example-structured_data.ipynb</i>	Comparing text-based metrics with specialized metrics for time-series and planning.
07. <i>quickstart.ipynb</i>	A simple, quickstart workflow using GAICo’s <code>Experiment</code> module.
<b>Intermediate Examples</b>	
08. <i>example-DeepSeek.ipynb</i>	Evaluating DeepSeek for Point of View (POV) document analysis.
09. <i>example-audio_data.ipynb</i>	Advanced audio analysis for TTS and music generation tasks.
10. <i>example-election.ipynb</i>	Evaluating models on sensitive election-related questions.
11. <i>example-finance.ipynb</i>	Evaluating models on finance domain questions by iterating over a dataset.
12. <i>example-planning.ipynb</i>	Evaluation of various travel plans using planning-specific metrics.
13. <i>example-recipes.ipynb</i>	Evaluating models on recipe generation with parallelization via <code>JOBLIB</code> .
14. <i>example-timeseries.ipynb</i>	In-depth evaluation and perturbation of time-series data.
<b>Advanced Examples</b>	
15. <i>example-llm_faq.ipynb</i>	Comparing various LLM responses (Phi, Mixtral, etc.) on a FAQ dataset.
16. <i>example-threshold.ipynb</i>	Exploring default and custom thresholding techniques for LLM evaluation.
17. <i>example-viz.ipynb</i>	Demonstrations of creating custom visualizations for evaluation results.

Table 2: Summary of GAICo Example Notebooks. Those referenced in this paper are *emphasized*. All available at: <https://github.com/ai4society/GenAIResultsComparator/tree/main/examples>.

core novelty; it ensures that once a metric is chosen, its application and reporting are **consistent**, directly improving the **quality** and **productivity** of the evaluation process. As demonstrated in our case study, this unified approach enables nuanced debugging, such as distinguishing orchestrator faults from specialist model failures, which is crucial for enhancing the **reliability** of complex AI systems.

### GAICo as a Deployed Tool

GAICo is engineered as a practical utility to help the community move **faster and safer**. Its design prioritizes **ease of use** and **adoption** through a simple PyPI installation, optional dependencies to manage installation size, comprehensive documentation, and a suite of runnable examples. This accessibility allows teams to **accelerate development** cycles; the streamlined `Experiment` class replaces disparate scripts with a single, consistent API, improving developer **productivity**. The framework enables a **safer deployment** of AI by facilitating precise failure analysis. By distinguishing a poor plan from a poor execution, GAICo improves system **reliability**, reduces long-term maintenance costs, and allows teams to mitigate weaknesses before they become incidents. Ultimately, GAICo is a deployed tool designed to have a significant **impact** by enabling the applied AI community to produce **better AI solutions** through rigorous, standardized, and efficient evaluation.

### Limitations and Future Work

Despite its strengths, GAICo has some limitations in its current scope. First, the current metric suite, while broad, primarily focuses on reference-based comparisons and does not yet include metrics for fairness, bias, toxicity, or operational aspects like latency and computational cost. Furthermore, current visualizations are static plots (bar, radar). There is no built-in interactive dashboard or web-based UI for dynamic exploration of results. Finally, while supporting structured data like planning and time-series, the library does not yet offer general-purpose comparison for arbitrary complex structured data (e.g., nested JSONs, knowledge graphs) or multi-turn conversational outputs.

We now discuss how one may extend the limitations in the future. A key priority is to incorporate metrics for fairness, bias, toxicity, and operational statistics, providing developers with a more holistic toolkit for building safer and more responsible AI systems. GAICo’s extensible design makes this easy. Future work also includes exploring developing interactive dashboards or deeper integration with MLOps platforms (like MLflow (Zaharia et al. 2018)) to enable dynamic exploration and tracking of evaluation results. Lastly, further work would also focus on adding general-purpose comparison capabilities for complex structured data (e.g., JSON diffing, graph similarity) and extending metrics to handle multi-turn conversational outputs.

## Acknowledgments

This work is partially supported by NSF Awards #2454027, NAIIR250014, and Faculty Award by JP Morgan Research.

## References

- Berndt, D. J.; and Clifford, J. 1994. Using dynamic time warping to find patterns in time series. In *Proceedings of the 3rd international conference on knowledge discovery and data mining*, 359–370.
- Bird, S.; Klein, E.; and Loper, E. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. ” O’Reilly Media, Inc.”.
- Christie, T. 2014. MkDocs. Accessed Aug 15, 2025.
- Clark, A. 2015. Pillow (PIL Fork) Documentation.
- Es, S.; James, J.; Anke, L. E.; and Schockaert, S. 2024. Ragas: Automated evaluation of retrieval augmented generation. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, 150–158.
- Farid, H. 2021. An overview of perceptual hashing. *Journal of Online Trust and Safety*, 1(1).
- Girdhar, R.; El-Nouby, A.; Liu, Z.; Singh, M.; Alwala, K. V.; Joulin, A.; and Misra, I. 2023. Imagebind: One embedding space to bind them all. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 15180–15190.
- Gupta, N.; Koppiseti, P.; Lakkaraju, K.; and Srivastava, B. 2025. GAICo: A Deployed and Extensible Framework for Evaluating Diverse and Multimodal Generative AI Outputs. *arXiv preprint arXiv:2508.16753*.
- Hirschberg, D. S. 1977. Algorithms for the longest common subsequence problem. *Journal of the ACM*, 24(4): 664–675.
- Honnibal, M.; Montani, I.; Van Landeghem, S.; and Boyd, A. 2020. spaCy: Industrial-strength Natural Language Processing in Python.
- Ip, J.; and Vongthongsri, K. 2025. deepeval.
- Krekel, H.; Oliveira, B.; Pfanschmidt, R.; Bruynooghe, F.; Laughner, B.; and Bruhin, F. 2004. pytest x.y. <https://github.com/pytest-dev/pytest>. Version x.y. Contributors include Holger Krekel, Bruno Oliveira, Ronny Pfanschmidt, Floris Bruynooghe, Brianna Laughner, Florian Bruhin, and others.
- Lin, C.-Y. 2004. ROUGE: A Package for Automatic Evaluation of Summaries. In *Text Summarization Branches Out*, 74–81. Barcelona, Spain: Association for Computational Linguistics.
- Lopresti, D.; and Nagy, G. 2021. Reproducibility: evaluating the Evaluations. In *International Workshop on Reproducible Research in Pattern Recognition*, 12–23. Springer.
- McFee, B.; Raffel, C.; Liang, D.; Ellis, D. P.; McVicar, M.; Battenberg, E.; and Nieto, O. 2015. librosa: Audio and music signal analysis in python. *SciPy*, 2015: 18–24.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In Isabelle, P.; Charniak, E.; and Lin, D., eds., *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, 311–318. Philadelphia, Pennsylvania, USA: Association for Computational Linguistics.
- Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; Vanderplas, J.; Passos, A.; Cournapeau, D.; Brucher, M.; Perrot, M.; and Duchesnay, E. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12: 2825–2830.
- PyPI Stats. 2025. GAICo Download Statistics. <https://pepy.tech/projects/gaico?timeRange=threeMonths&category=version&includeCIDDownloads=true&granularity=daily&viewType=line&versions=0.3.0,0.2.0,0.1.5,0.1.4>. Accessed Aug 18, 2025.
- Rix, A. W.; Beerends, J. G.; Hollier, M. P.; and Hekstra, A. P. 2001. Perceptual evaluation of speech quality (PESQ)—a new method for speech quality assessment of telephone networks and codecs. In *2001 IEEE international conference on acoustics, speech, and signal processing. Proceedings (Cat. No. 01CH37221)*, volume 2, 749–752. IEEE.
- Sara, U.; Akter, M.; and Uddin, M. S. 2019. Image quality assessment through FSIM, SSIM, MSE and PSNR—a comparative study. *Journal of Computer and Communications*, 7(3): 8–18.
- Sottile, A. 2014. pre-commit. Accessed Aug 15, 2025.
- Srivastava, B.; Kambhampati, S.; Do, M. B.; and Nguyen, T. A. 2006. Finding inter-related plans. *ICAPS 2006*, 18.
- Srivastava, B.; Nguyen, T. A.; Gerevini, A.; Kambhampati, S.; Do, M. B.; and Serina, I. 2007. Domain Independent Approaches for Finding Diverse Plans. In *IJCAI*, 2016–2022.
- Taal, C. H.; Hendriks, R. C.; Heusdens, R.; and Jensen, J. 2010. A short-time objective intelligibility measure for time-frequency weighted noisy speech. In *2010 IEEE international conference on acoustics, speech and signal processing*, 4214–4217. IEEE.
- Van der Walt, S.; Schönberger, J. L.; Nunez-Iglesias, J.; Boulogne, F.; Warner, J. D.; Yager, N.; Gouillart, E.; and Yu, T. 2014. scikit-image: image processing in Python. *PeerJ*, 2: e453.
- Virtanen, P.; Gommers, R.; Oliphant, T. E.; Haberland, M.; Reddy, T.; Cournapeau, D.; Burovski, E.; Peterson, P.; Weckesser, W.; Bright, J.; van der Walt, S. J.; Brett, M.; Wilson, J.; Millman, K. J.; Mayorov, N.; Nelson, A. R. J.; Jones, E.; Kern, R.; Larson, E.; Carey, C. J.; Polat, İ.; Feng, Y.; Moore, E. W.; VanderPlas, J.; Laxalde, D.; Perktold, J.; Cimrman, R.; Henriksen, I.; Quintero, E. A.; Harris, C. R.; Archibald, A. M.; Ribeiro, A. H.; Pedregosa, F.; van Mulbregt, P.; and SciPy 1.0 Contributors. 2020. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17: 261–272.
- Wang, Z.; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4): 600–612.
- Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.

Yuan, T.; Deng, W.; Tang, J.; Tang, Y.; and Chen, B. 2019. Signal-to-noise ratio: A robust distance metric for deep metric learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4815–4824.

Zaharia, M. A.; Chen, A.; Davidson, A.; Ghodsi, A.; Hong, S. A.; Konwinski, A.; Murching, S.; Nykodym, T.; Ogilvie, P.; Parkhe, M.; Xie, F.; and Zumar, C. 2018. Accelerating the Machine Learning Lifecycle with MLflow. *IEEE Data Eng. Bull.*, 41: 39–45.