

# 000 TIMESERIESGYM: A SCALABLE BENCHMARK FOR 001 (TIME SERIES) MACHINE LEARNING ENGINEERING 002 AGENTS 003

004  
 005  
 006 **Anonymous authors**  
 007 Paper under double-blind review  
 008  
 009  
 010  
 011

## ABSTRACT

012  
 013 We introduce `TimeSeriesGym`, a scalable benchmarking framework for evaluat-  
 014 ing Artificial Intelligence (AI) agents on time series machine learning engineering  
 015 challenges. Existing benchmarks lack scalability, focus narrowly on model building  
 016 in well-defined settings, and evaluate only a limited set of research artifacts (e.g.,  
 017 CSV submission files). To make AI agent benchmarking more relevant to the  
 018 practice of machine learning engineering, our framework scales along two critical  
 019 dimensions. First, recognizing that effective ML engineering requires a range of  
 020 diverse skills, `TimeSeriesGym` incorporates challenges from diverse sources  
 021 spanning multiple domains and tasks. We design challenges to evaluate both iso-  
 022 lated capabilities (including data handling, understanding research repositories, and  
 023 code translation) and their combinations, and rather than addressing each challenge  
 024 independently, we develop tools that support designing multiple challenges at scale.  
 025 Second, we implement evaluation mechanisms for multiple research artifacts, in-  
 026 cluding submission files, code, and models, using precise numeric measures and  
 027 *optionally* LLM-based qualitative assessments. This strategy complements objec-  
 028 tive evaluation with subjective assessment when appropriate. Although our initial  
 029 focus is on time series applications, our framework can be readily extended to other  
 030 data modalities, broadly enhancing the comprehensiveness and practical utility of  
 031 agentic AI evaluation. We open-source<sup>1</sup> our benchmarking framework to facilitate  
 032 future research on the ML engineering capabilities of AI agents.  
 033  
 034

## 1 INTRODUCTION

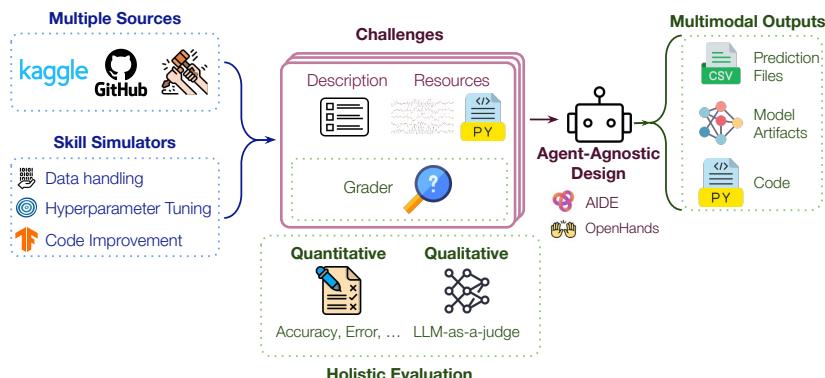
035 AI agents (9; 11) have shown growing promise in automating machine learning (ML) and data science  
 036 (DS) workflows. Fueled by large language models (LLMs), they can reason about context, adapt to  
 037 diverse tasks, and iteratively refine solutions over long horizons. Such capabilities have the potential  
 038 to significantly reduce the mundane, mostly manual efforts in ML engineering and improve the  
 039 overall productivity of ML practice. To measure progress in this area, several benchmarks (9; 3; 12;  
 040 17; 22; 1; 20; 18) have been proposed to evaluate AI agents on ML and DS tasks.

041 However, existing benchmarks have important limitations. First, many of them source ML challenges  
 042 primarily from popular competitions such as Kaggle, which are well-structured and do not fully  
 043 capture the complexity of real-world ML tasks. Second, evaluations are typically outcome-based,  
 044 focusing on overall task completion or eventual model performance metrics such as accuracy, while  
 045 combining and obfuscating the impact of multiple component skills that jointly determine the  
 046 outcomes, such as effective data wrangling or code quality improvement capabilities. Third, current  
 047 benchmarks lack scalability, as tasks have to be manually curated and cannot be developed at scale.

048 To enable efficient evaluation of AI agents in realistic ML scenarios, we propose `TimeSeriesGym`,  
 049 a scalable and agent-agnostic benchmarking framework for evaluating AI agents on time series  
 050 ML engineering tasks. `TimeSeriesGym` currently consists of 33 challenges that span 8 unique  
 051 time series problems (forecasting, classification, time series understanding), from more than 15  
 052 domains (healthcare, finance, epidemiology). Our benchmark covers both Kaggle-style challenges

053 <sup>1</sup><https://anonymous.4open.science/r/TimeSeriesGym-9CF6/>

054  
 055 Figure 1: TimeSeriesGym is a scalable benchmarking environment for ML engineering agents. It  
 056 currently features 33 time series challenges across 8 unique time series problems, spanning more than  
 057 15 domains. Challenges are either carefully designed based on real-world ML practice, or sourced  
 058 from Kaggle competitions and GitHub repositories. TimeSeriesGym includes key mechanisms to  
 059 enable efficient and scalable generation of new challenges. Our evaluation methodology complements  
 060 precise quantitative metrics with *optional* qualitative assessment, and provides specialized tools to  
 061 grade various artifacts generated during ML engineering. TimeSeriesGym is compatible with  
 062 many different agent types, even those with fundamentally distinct designs.



063  
 064  
 065  
 066  
 067  
 068  
 069  
 070 and original tasks carefully designed based on real-world ML engineering practice. While we focus  
 071 on time series analytics due to its prevalence in applications and under-representation in existing  
 072 agent benchmarks, our framework is modality-agnostic in principle and can be easily extended to  
 073 handling other data modalities (e.g., images, text, audio) via the accompanying tools for scalable task  
 074 generation. TimeSeriesGym provides an interactive gym environment compatible with various  
 075 types of agent scaffolding, allowing seamless evaluation of agents of different types and collection  
 076 of their trajectories. Beyond benchmarking, this allows TimeSeriesGym to also serve as a data  
 077 flywheel for future agent improvement through post-training using the collected trajectory data.

078 Our contributions are as follows:

079

- 080 1. We propose TimeSeriesGym, the first open source benchmark for AI agents on time  
 081 series ML engineering tasks (Fig. 1). Beyond benchmarking, TimeSeriesGym can be  
 082 easily used as a data flywheel to post-train the next generation of ML engineering agents.
- 083 2. TimeSeriesGym offers a scalable task generation mechanism that reduces manual efforts  
 084 in task curation and ensures long-term sustainability of the benchmark.
- 085 3. TimeSeriesGym provides a comprehensive framework that evaluates **multimodal agent**  
 086 **outputs** (e.g., prediction files, models, code) across **specific ML skills** (e.g., data handling,  
 087 model improvement), using a **holistic approach** that complements quantitative metrics (e.g.,  
 088 accuracy) with *optional* qualitative assessment (e.g., LLM-as-a-judge for code utility).

## 089 2 RELATED WORK

090 **Machine learning agent benchmarks.** Several benchmarks have been proposed to evaluate LLM  
 091 agents on automating ML and DS tasks. These benchmarks are typically structured around three key  
 092 components: (1) task curation, (2) agent capabilities being evaluated, and (3) evaluation protocol.  
 093 Benchmarks differ in how they curate ML/DS tasks. For example, MLE-bench (3) and DSBench (12)  
 094 compile tasks from online competition platforms such as Kaggle, while other benchmarks source  
 095 tasks from ML-related GitHub repositories (1; 20) or hand-craft tasks based on ML research problems  
 096 or engineering workflows (9; 17; 22; 18). With regard to agent capabilities, some benchmarks (3; 9;  
 097 17; 22; 12) focus on comprehensive ML science skills by evaluating agents on end-to-end problem  
 098 solving skills, while others (1; 20; 18) focus on more modular engineering-oriented capabilities within  
 099 the ML pipeline, such as using GitHub repositories or integrating APIs. Evaluation protocols also  
 100 differ in output formats and granularity. MLE-bench (3) and DSBench (12) require agents to output  
 101

108  
109 Table 1: Comparison of TimeSeriesGym with existing ML/DS agent benchmarks. Categories  
110 include **Number**: total and time series (TS) tasks in each benchmark, where each task corresponds  
111 to a unique data source (e.g., a single Kaggle competition or GitHub repository); **Source**: task  
112 origins (K: Kaggle, G: GitHub, H: Hand-crafted); **ML Capability**: coverage of ML **science** tasks  
113 (e.g., modeling, open-ended research) and **engineering** tasks (e.g., repository utilization, API in-  
114 tegration); and **Evaluation**: capabilities for evaluating **multimodal** outputs (e.g., prediction files,  
115 model artifacts), specific ML **skills** (e.g., data handling, model improvement), and from a **holistic**  
116 perspective (including both quantitative metrics (e.g., accuracy, mean absolute error) and qualitative  
117 evaluation (e.g., code review)). We use "+" to indicate TimeSeriesGym’s scalability which enables  
118 the generation of an unlimited number of new challenges using the tools provided.

	Number		Source			ML Capability		Evaluation		
	Total	TS	K	G	H	Science	Engineering	Multimodal	Skill-based	Holistic
MLE-bench (3)	75	3	✓	✗	✗	✓	✗	✗	✗	✗
MLAgentBench (9)	13	1	✓	✗	✓	✓	✗	✓	✗	✗
MLGym (17)	13	0	✓	✗	✓	✓	✗	✓	✗	✗
RE-Bench (22)	7	0	✗	✗	✓	✓	✗	✓	✗	✗
DSBench <sup>2</sup> (12)	74	5	✓	✗	✗	✓	✗	✗	✗	✗
SUPER <sup>3</sup> (1)	45	0	✗	✓	✗	✗	✓	✗	✗	✓
ML-Bench (20)	18	1	✗	✓	✗	✗	✓	✗	✗	✗
ML-Dev-Bench (18)	30	0	✗	✗	✓	✗	✓	✓	✓	✗
TimeSeriesGym (Ours)	23+	23+	✓	✓	✓	✓	✓	✓	✓	✓

120  
121  
122  
123  
124  
125  
126 results in specific formats (e.g., CSV files) that can be directly scored using predefined metrics such as  
127 accuracy, while other benchmarks (9; 17; 22) allow for more flexible outputs in addition to prediction  
128 files, such as model artifacts and code. ML-Dev-Bench (18) further extends the evaluation by specific  
129 skills (e.g., data handling, model improvement), while SUPER (1) provides a more holistic evaluation  
130 by combining outcome-based evaluation with qualitative code inspection to assess agents’ progress  
131 towards completing the tasks. TimeSeriesGym is most similar to MLE-Bench in sourcing Kaggle  
132 competitions, but uniquely emphasizes time series modeling (an underrepresented modality), includes  
133 original ML engineering challenges beyond competitions, and provides granular skill assessment,  
134 and holistic multi-artifact evaluation (see difference in Sec. D).

135  
136  
137 **Scalable dynamic benchmarks and holistic evaluation.** Scalable benchmarks reduce manual data  
138 curation efforts by generating target problems at scale using carefully designed templates (2; 23) or  
139 data engines (8), among which TimeSeriesExam (2) further improves problem sample quality by  
140 applying Item Response Theory (IRT) (4; 7) to intelligently select questions with contextualized  
141 difficulty and appropriate discrimination. To remain effective against data contamination from LLM  
142 pretraining, dynamic benchmarks such as GAIA (16) and LiveCodeBench (10) propose to continually  
143 incorporate problems newly released after LLM training cut-offs. While most benchmarks target  
144 specific capabilities, holistic evaluation (13; 5) provides a comprehensive picture through evaluating  
145 models on a wide range of datasets and tasks across diverse domains using multiple complementary  
146 metrics, to capture both the breadth and depth of model capabilities.

### 147 3 TIMESERIESGYM

148  
149 TimeSeriesGym is envisioned as a scalable benchmarking environment for time series machine  
150 learning engineering. The current version features 33 challenges from 23 unique data sources across  
151 8 unique time series problems, spanning more than 15 domains. These challenges evaluate AI  
152 agents on a range of realistic ML engineering skills beyond just model development, including data  
153 labeling, model selection, and the utilization, improvement, and migration of research code (Tab. 4).  
154 TimeSeriesGym also provides tools for rapidly developing new challenges to test specific skills  
155 and for evaluating the diverse artifacts commonly generated during ML engineering processes.

156  
157 Each challenge in TimeSeriesGym is organized with a consistent structure: (1) **resources** including  
158 datasets, code repositories, related paper(s) and documentation relevant to the challenge; (2) a

159  
160 <sup>1</sup>For DSBench, we include only data modeling tasks, while excluding data analysis tasks as they are not  
161 directly relevant to our work.

<sup>2</sup>For SUPER, we include repositories used to create the Expert and Masked sets of the benchmark.

162 **description file** that outlines the challenge parameters, available resources, and provides specific  
 163 instructions and hints for successful completion; and (3) **challenge-specific grading functions** to  
 164 evaluate agent submissions. Some challenges also include leaderboards to rank agent submissions  
 165 against human performance. These leaderboards are readily available for, e.g., challenges derived  
 166 from Kaggle competitions.

167 **The challenges in TimeSeriesGym** are derived from Kaggle competitions (currently,  $n = 13$ ) and  
 168 popular benchmarks and research code repositories for time series modeling (*TimeSeriesGym*  
 169 *Originals*,  $n = 14$ ). We prioritized challenges that reflect core skills that are regularly exercised by  
 170 ML engineers, researchers, and data scientists.

171 Each challenge is specifically chosen or designed to evaluate one or more of the following skills:  
 172 (1) *Data Handling*: Ability to handle missing data, use data labeling tools, and utilize multi-source  
 173 data for model building. (2) *Modeling*: Ability to develop useful time-series ML models, tune  
 174 hyperparameters, perform model selection, and understand, utilize, migrate and improve the quality  
 175 of research code. (3) *Benchmarking*: Training and rigorously evaluating ML models using standard  
 176 benchmarks. In selecting these challenges, we aimed at a broad coverage across diverse domains (e.g.,  
 177 healthcare, finance, epidemiology) and time series problems (forecasting, classification, time series un-  
 178 derstanding). Tab. 4 provides a comprehensive overview of each challenge within *TimeSeriesGym*,  
 179 including its domain, core problem, evaluation metric, and the skills required to address it.

180 To identify Kaggle challenges for inclusion in *TimeSeriesGym*, we began with the Meta Kaggle  
 181 dataset (19), focusing specifically on Featured and Research competitions. Featured competitions  
 182 are real-world ML challenges that pose difficult, commercially oriented prediction problems, while  
 183 Research competitions focus on problems that may not have clean or straightforward solutions<sup>4</sup>.  
 184 We employed Gemini 2.0 Flash to analyze competition descriptions and titles, identifying  
 185 453 competitions that likely involve time series data. Subsequently, we ranked these competitions  
 186 by participant count, maximum reward, and presence of a public leaderboard. Then from top 100  
 187 ranked high-quality competitions (see Tab. 8), we manually selected a diverse set of competitions to  
 188 ensure comprehensive coverage across problem type (e.g., forecasting, classification), domain (e.g.,  
 189 finance, healthcare), and research or engineering complexity (e.g., organization of datasets), while  
 190 also requiring public data availability, and a permissive license.

191 To complement the selected Kaggle challenges, we include 14 *TimeSeriesGym* Original  
 192 challenges, manually curated from existing open-source time series datasets and GitHub  
 193 repositories based on recommendations from experienced ML engineers and researchers (see  
 194 Tab. 5 for the original and modified source codes and licenses). These challenges are  
 195 specifically designed to evaluate advanced skills that Kaggle competitions typically cannot  
 196 easily assess, yet are essential for effective ML engineering. Examples include utilizing  
 197 state-of-the-art models (e.g., Implement the *MOMENT* (6) time series foundation  
 198 model for anomaly detection), migrating frameworks (e.g., Convert ResNet-1D  
 199 classification models from TensorFlow to PyTorch), and improving research  
 200 code quality (e.g., Improve PTB-XL ECG Classification Code). These capabilities  
 201 represent critical competencies of skilled ML engineers that extend beyond the scope of standard  
 202 Kaggle-like competitions.

203 Running experiments on *TimeSeriesGym* can be resource-intensive and costly. Therefore, we propose  
 204 *TimeSeriesGym-Lite*, a carefully selected subset of six challenges designed to efficiently  
 205 evaluate AI agents on critical ML engineering skills while maintaining coverage across multiple  
 206 domains and time series problems, **with statistically similar difficulty to *TimeSeriesGym* (Sec. B)**.  
 207 This collection enables rapid and cost-effective assessment of novel AI agents without sacrificing the  
 208 diversity of skills being tested (see Tab. 6).

209 **Multimodal, skill-based, holistic evaluation.** Existing benchmarks typically summarize agent  
 210 performance using metrics such as accuracy, completion rate, or competition rankings (3). Although  
 211 these metrics provide useful summaries, they do not offer much actionable feedback for improvement.  
 212 *TimeSeriesGym* addresses this limitation through an evaluation framework designed to provide  
 213 specific actionable feedback through multiple complementary approaches. First, we design chal-  
 214 lenges that isolate and test specific skills, such as handling missing data (e.g., *Optiver Realized*  
 215 *Volatility Prediction with Missing Data*). Poor performance on such challenges

<sup>4</sup><https://www.kaggle.com/docs/competitions>

216 clearly indicates potential skill gaps, enabling developers to focus their efforts on specific skills. Sec-  
 217 ond, we develop fine-grained evaluation tools that assess multiple dimensions of performance simul-  
 218 taneously. For example, in code migration tasks (e.g., `Convert ResNet from TensorFlow`  
 219 to `PyTorch`), our evaluation tools examine whether an agent follows instructions and naming con-  
 220 ventions, completes all required function definitions, in addition to successful execution—providing a  
 221 multidimensional performance profile rather than a binary success/failure indicator.

222 Our evaluation methodology deliberately combines multiple assessment approaches: quantitative  
 223 metrics (accuracy, mean absolute error), programmatic analysis (regular expression matching, code  
 224 inspection), and optional qualitative evaluation (LLM-as-a-judge) (see Appendix G). Each challenge  
 225 in `TimeSeriesGym` is evaluated using quantitative metrics (based on rules) and optionally subjec-  
 226 tive metrics (LLM judge). Although LLM-based evaluation offers valuable insight, especially for  
 227 open-ended tasks such as research code enhancement, we recognize that LLMs can be inconsistent  
 228 and prone to hallucination. Therefore, we primarily rely on quantitative metrics and strategically  
 229 complement them with subjective assessments, creating a holistic evaluation system that leverages  
 230 the strengths of both approaches. This hybrid approach mimics code reviews in software engineering  
 231 practice, which includes both tests based on static analysis and human code reviews.

232 Furthermore, `TimeSeriesGym` provides specialized tools to grade diverse artifacts generated  
 233 throughout the ML engineering life cycle— from submission files (CSV, H5, etc.) to source code (.py)  
 234 and trained models (.pth, .pkl)—enabling comprehensive assessment of the entire development  
 235 process rather than focusing solely on final outputs.

236 **Generating challenges at scale.** We provide key mechanisms that enable scalable generation of  
 237 new challenges. Here, *scalability* refers not only to adding new challenges efficiently, but more  
 238 importantly to the *flexibility* to support multiple design choices across various components of the  
 239 benchmark. By design, `TimeSeriesGym` scales along several key dimensions: (1) **Skill-specific**  
 240 **competitions:** We provide specialized tools (e.g., missing data simulator) that can be paired with any  
 241 "base" competition to create a large variety of targeted, skill-specific competitions. (2) **Agent outputs:**  
 242 Our grading tools support the evaluation of diverse agent outputs, including prediction files, model  
 243 artifacts, and code, allowing easy assessment across many task types. (2) **Agentic scaffolds:** Unlike  
 244 existing benchmarks such as MLGym (17), `TimeSeriesGym` is agnostic to agent implementations,  
 245 enabling a wide range of agent scaffolds to be integrated with minimal effort. (4) **Data sources:**  
 246 `TimeSeriesGym` accommodates both Kaggle-style and non-Kaggle datasets (such as datasets in  
 247 `TimeSeriesGym` Original challenges), making it straightforward for practitioners to introduce new  
 248 datasets regardless of source or format<sup>5</sup>. Additionally, we offer clear and detailed documentation for  
 249 adding new challenges to the benchmark, which has already enabled members of our community  
 250 (outside our research team) to contribute a new challenge within 2 hours. Together, these design  
 251 choices ensure that `TimeSeriesGym` can continuously grow and adapt as time series machine  
 252 learning techniques continue to advance.

### 253 3.1 DESIGN CHOICES

254 **Focus on time series tasks.** We focused on time series modeling tasks for two key reasons. First,  
 255 time series data is ubiquitous and critical in domains such as healthcare and economics, yet existing  
 256 agentic AI benchmarks include relatively few time series challenges (Tab. 1). Second, compared  
 257 to text and images, time series data require modest resources for storage and modeling, making  
 258 `TimeSeriesGym` efficient to run. Moreover, modeling time-series data remains relatively un-  
 259 derexplored outside specialized research communities, meaning that LLMs are less likely to have  
 260 encountered such data and tasks during training. This characteristic, combined with the fact that  
 261 `TimeSeriesGym` evaluates general machine learning skills, makes it an excellent testbed to eval-  
 262 uate AI agent capabilities. Although focused on time series, our benchmark can be be readily extended  
 263 to other modalities, and already includes multimodal challenges, such as `TimeSeriesExam` (time  
 264 series + text) and `OSIC Pulmonary Fibrosis Progression` (time series + images).

265 **How much freedom should the agents be given?** When designing challenges for  
 266 `TimeSeriesGym`, we had to strike a fine balance between giving agents freedom to solve problems

267  
 268 <sup>5</sup>`TimeSeriesGym` is designed to be extensible while maintaining high standards of correctness, difficulty,  
 269 and non-triviality. To support benchmark growth without compromising quality, we implement a category-  
 specific quality-assurance pipeline, outlined in Sec H

creatively while keeping enough structure in place to allow for a precise and fine-grained evaluation. For example, in the `PTB-XL ECG Classification with Hyper-parameter Optimization` challenge, we *required* agents to use a PyTorch-based neural network and save their models, files and code before and after tuning. This allowed us to inspect models and code to check if the hyper-parameters changed, and measure how these changes improved performance.

**Agent-agnostic design.** `TimeSeriesGym` is agnostic to specific agent implementations. Following `MLE-bench` (3), it is easy to add new challenges and agentic scaffolds. To illustrate this flexibility, we include the *latest* implementations of 3 different scaffolds, `AIDE` (11), `MLAgentBench` (`MLAB`) (9), and `OpenHands` (21) with fundamentally different designs. Unlike `MLGym` (17), we do not advocate for a default agentic scaffold, as we believe that agent designs will continue to evolve and no single scaffold will work best for all ML engineering tasks.

## 4 EXPERIMENTS AND RESULTS

Table 2: **Main Results.** Each experiment was run with 3 random seeds, with results showing mean  $\pm$  standard deviation. The table compares scaffold types (`AIDE`, `OpenHands`, `MLAB`), model choices (`GPT-4.1`, `o3`, `Claude 3.7`), resource allocations (4/50 to 12/150 hours/steps), and time utilization approaches. Key findings include: (1) `AIDE` achieves the best performance as a scaffold, (2) the reasoning model `o3` achieves significantly higher valid submission rates (94.4%) than other models, (3) `Claude 3.7` produces the most reasonable submissions (38.9%), (4) doubling time resources does not consistently improve performance, and (5) interestingly, removing step-wise reminders sometimes improves reasonable submission rates.

Lite	Model	Resources (hours / steps)	Valid Submission (%)	Reasonable Submission (%)
	<code>MLAB</code>			
✓	+ <code>gpt-4.1-2025-04-14</code>	4 / 50	44.4 $\pm$ 9.6	27.8 $\pm$ 9.6
	<code>OpenHands</code>			
✓	+ <code>gpt-4.1-2025-04-14</code>	4 / 50	44.4 $\pm$ 19.3	11.1 $\pm$ 9.6
	<code>AIDE</code>			
✓	+ <code>gpt-4.1-2025-04-14</code>		66.7 $\pm$ 16.7	27.8 $\pm$ 9.6
✓	+ <code>o3-2025-04-16</code>		<b>94.4 <math>\pm</math> 9.6</b>	33.3 $\pm$ 0.0
	+ <code>claude-3-7-sonnet-20250219</code>	4 / 50	50.0 $\pm$ 16.7	38.9 $\pm$ 19.3
	+ <code>deepseek-reasoner</code>		<b>11.1 <math>\pm</math> 9.6</b>	<b>11.1 <math>\pm</math> 9.6</b>
	+ <code>deepseek-chat</code>		<b>16.7 <math>\pm</math> 0.0</b>	<b>16.7 <math>\pm</math> 0.0</b>
✗	+ <code>gpt-4.1-2025-04-14</code>	4 / 50	58.6 $\pm$ 7.6	12.1 $\pm$ 0
<i>Effect of Scaling Resources</i>				
		4 / 50	66.7 $\pm$ 16.7	27.8 $\pm$ 9.6
✓	+ <code>gpt-4.1-2025-04-14</code>	8 / 100	72.2 $\pm$ 9.6	22.2 $\pm$ 9.6
		12 / 150	61.1 $\pm$ 9.6	<b>50.0 <math>\pm</math> 0.0</b>
<i>Effective Utilization of Time</i>				
✓	Step-wise reminder	4 / 50	66.7 $\pm$ 16.7	27.8 $\pm$ 9.6
	No reminder		55.6 $\pm$ 9.6	33.3 $\pm$ 0.0

**Setting.** We run agents in an Ubuntu 20.04 Docker container with all necessary resources (datasets, code repositories, etc.) and basic Python packages useful for ML engineering. Agents can access the internet and install additional packages as needed. For each challenge, agents have a maximum of 4 hours and 50 steps (9; 18; 17) and use a machine with 128 vCPUs, 503 GB RAM, 1.8 TiB SSD, and a single NVIDIA A100-SXM4-80GB GPU<sup>6</sup>. Unless otherwise specified, we repeat each experiment with 3 different seeds (0, 1, and 2) to calculate mean and standard deviation.

<sup>6</sup>In practice, agents share this machine as we run multiple challenges in parallel. This represents a realistic setting similar to how ML engineers routinely share computing resources. We found no instances where this sharing might have disadvantaged any agent.

**Cost.** On average, it cost us USD 63.00 to run AIDE with `gpt-4.1-2025-04-14` for a maximum of 4 hours and 50 steps on `TimeSeriesGym`. In contrast, the `Lite` benchmark was much more affordable at USD 8.00 per run. Given that `TimeSeriesGym`-`Lite` preserves coverage across domains and problem types while being much more time- and cost-effective, we conducted most of our experiments on `TimeSeriesGym`-`Lite` to accommodate our resource constraints.

**Metrics.** We report the raw scores achieved by AIDE on every challenge (Tab. 13). Although these scores are useful for tracking progress on individual challenges, they cannot be easily combined across different challenges. To measure the performance of agents at a high level, we report two key metrics: the percentage of challenges where the agent made a (1) *valid*, and (2) *reasonable* submission (Tab. 2). A submission is valid if the grader returns any non-null score. What counts as a reasonable attempt varies by challenge type. For Kaggle challenges, we define it as scoring above median on the competition’s public leaderboard. For the remaining challenges, we consider a submission reasonable if it made a genuine<sup>7</sup> modeling attempt rather than hallucinating an output that matches the submission format. For example, simply loading and re-saving the provided sample submission file without any model inference or data processing is deemed unreasonable.

#### 4.1 OBSERVATIONS

**AIDE is the better open-source scaffold.** We evaluated `GPT-4.1` (`gpt-4.1-2025-04-14`) with three open-source scaffolds: AIDE (11), MLAB (9) and OpenHands (21). Following `MLE-bench`, we make minor modifications to adapt these scaffolds to our benchmark (see Appendix C). Results in Tab. 2 confirm prior findings: AIDE with `GPT-4.1` yields the highest proportion of valid (66.7%) and reasonable (27.8%) submissions. This is expected as AIDE is specifically designed for data science tasks, which account for the majority of `TimeSeriesGym` challenges.

**Reasoning models produce substantially more valid submissions.** To identify the best base model, we evaluated the strongest scaffold (AIDE) with two state-of-the-art proprietary LLMs: `GPT-4.1` and `Claude 3.7` (`claude-3-7-sonnet-20250219`), and a reasoning model  $\circ 3$  ( $\circ 3-2025-04-16$ ). As shown in Tab. 2, our experiments on `TimeSeriesGym`-`Lite` revealed that  $\circ 3$  created significantly more valid submissions than other models, while `Claude 3.7` produced the highest number of reasonable attempts (38.9%). We noticed a significant gap between valid and reasonable submissions for  $\circ 3$ . While  $\circ 3$  can generate valid submission files for most challenges by following the instructions provided, it was also prone to hallucination. In some cases, it produced a submission file without any genuine modeling attempt (e.g., simply outputting a zero array). An illustration of this failure mode is provided in Appendix F.4.

**Challenges in `TimeSeriesGym` are hard for state-of-the-art agents.** We tested AIDE with `GPT-4.1` on all `TimeSeriesGym` challenges (see Tab. 13) and found poor overall performance. The agent produced valid submissions for only 58.6% of challenges and reasonable submissions for only 12.1% on average. We found that the agent especially struggled with `TimeSeriesGym` original challenges, where it only produced valid submissions for 5 out of 14 challenges. These results show that even the best agents struggle with ML engineering tasks, particularly those that go beyond standard Kaggle data science challenges and involve working with multi-file code repositories.

**Agents do not improve with more time.** We wondered if the agents perform poorly on `TimeSeriesGym` simply because they need more time. To test this, we ran AIDE with `GPT-4.1` on `TimeSeriesGym`-`Lite` and gave it 2–3x more hours and steps per challenge. Our results show that extra time does not always improve performance (Tab. 14). Even with the maximum time (12 hours and 150 steps), the agent only made reasonable submissions in about 50% of challenges—not very impressive given that the bar for a “reasonable” submission is quite low.

**Agents do not utilize time effectively.** We suspected that agents do not improve with more time because they do not use it well. To test this idea, we compared two settings: the default approach of reminding the agent about remaining time (and steps) before each step, versus removing these reminders completely. Surprisingly, we did not find significant differences between these settings. In fact, agents without time reminders produced more reasonable submissions. This may suggest that

<sup>7</sup>We assess this by manually inspecting whether a modeling attempt was made. Since this is inherently a subjective judgment, reasonable attempts are reported ONLY as a study of agentic skills. The benchmark score itself ultimately depends on the objective metric we defined for each challenge.

378 agents do not use their time wisely—they tend to rush toward solutions instead of carefully exploring  
 379 promising options, especially towards the end of the experiment. This raises important research  
 380 questions about how to design agents that use their time and resources more strategically.  
 381

382 **Frontier LLM challenges.** Since frontier LLMs are pretrained on large-scale public data, there is  
 383 a risk that they may have encountered and memorized content from public challenges, e.g., online  
 384 Kaggle competition discussions or solutions, which can potentially inflate benchmark performance  
 385 and limit its generalizability. To assess this risk, we followed the approach used by MLE-bench  
 386 to measure GPT-4.1’s familiarity with TimeSeriesGym challenges and compared the familiarity  
 387 score distribution to that of MLE-bench (see Appendix E.4). GPT-4.1 exhibited a similar  
 388 level of familiarity with TimeSeriesGym challenges as with MLE-bench challenges (with Kol-  
 389 mogorov-Smirnov (KS) Test (15) p-value = 0.363, indicating no significant difference). Given that  
 390 MLE-bench found no systematic impact of LLM familiarity on experiment results, GPT-4.1’s fa-  
 391 miliarity with TimeSeriesGym is within a reasonable range and does not compromise its integrity.  
 392

393 **TimeSeriesGym can be an effective diagnostic tool for agentic skill development.** We stratify  
 394 the results in Tab. 2 in terms of core ML skills that each challenge tests (Tab. 3). For example, we  
 395 found that agents struggle in hyper-parameter tuning—2 out of the 3 scaffolds (OpenHands and  
 396 MLAB) failed to produce valid submissions. While all scaffolds perform similarly on code migration,  
 397 AIDE achieved the best performance on handling missing data, likely due to its specific design for  
 398 data science tasks. A similar analysis for base models is provided in Appendix E.3.  
 399

400 Table 3: Performance of agent scaffolds with GPT-4.1 (Best@3 / Avg@3) on different ML skills.  
 401 Arrows indicate whether lower (↓) or higher (↑) values are better. Agents struggle with hyper-  
 402 parameter tuning.  
 403

ML Skill	Metric	AIDE	OpenHands	MLAB
Handling Data Missingness	Root Mean Square Percentage Error (↓)	0.33/0.33	0.64/0.64	0.42/0.89
Code Migration	Percentage of Test Cases Passed (↑)	0.56/0.56	0.56/0.44	0.56/0.56
Hyper-parameter Tuning	Improvement in Accuracy (↑)	0.08/0.03	N/A / N/A	N/A / N/A

403 **Summary.** This section provides a focused illustration of how TimeSeriesGym enables efficient  
 404 and cost-effective experimentation with AI agents, helping researchers uncover actionable insights  
 405 about agent capabilities and limitations. Our findings demonstrate TimeSeriesGym’s value for  
 406 advancing generic ML engineering agents.  
 407

## 412 5 DISCUSSION, OPEN QUESTIONS AND OPPORTUNITIES

413 **Key limitations of existing scaffolds.** Agentic scaffolds such as AIDE and OpenHands provide  
 414 structured workflows that **excel in single-shot, self-contained benchmarks** (e.g., Kaggle com-  
 415 petitions) but reveal **significant limitations in repository-level challenges** that require multiple file  
 416 edits and iterative reasoning. AIDE’s *one-step solution strategy* and *fixed action set*—restricted to  
 417 predefined operations such as “data preview” when debugging—often lead to unsuccessful attempts  
 418 in large codebases, as the agent’s attention is diluted across irrelevant files and fails to identify critical  
 419 information. Conversely, OpenHands supports multi-step trajectories yet suffers from a **greedy**  
 420 **exploitation bias**: it commits fully to a single approach without exploring alternative solution paths  
 421 or revisiting earlier decisions when trajectories prove unfruitful. The planning algorithm of the  
 422 CodeAct agent used by OpenHands is similarly *greedy and short-horizon*, limiting adaptation  
 423 to complex multistage development workflows. These findings highlight the need for **more adap-  
 424 tive scaffolds** that dynamically expand their action repertoire, balance exploration and exploitation  
 425 through parallel solution threads, and support nested workflows reflective of real-world machine  
 426 learning engineering tasks. We provide illustrations of agent failures in Appendix F.  
 427

428 **Data leakage and plagiarism.** In designing TimeSeriesGym, we identify two key risks related  
 429 to data leakage and plagiarism that could compromise the integrity of the benchmark: (1) **Pretraining**  
 430 **contamination:** Current LLMs may have been exposed to public content from existing challenges  
 431 (e.g., Kaggle competitions), including task descriptions, data, or shared solutions. This can lead to  
 432 memorization and inflated performance that overstates agents’ true capabilities, and (2) **Future LLM**

432 **leakage:** Once the benchmark is public, future LLMs may be pretrained on its content, making the  
 433 benchmark less effective in evaluating real generalization.  
 434

435 To address such risks, we present both empirical findings and mitigation strategies. For case (1), we  
 436 have two key observations. First, in both Kaggle-based and original challenges in TimeSeriesGym,  
 437 agents either performed poorly or did not produce valid output, suggesting minimal benefit from  
 438 any potential LLM contamination. Second, we conducted a formal analysis using available tools to  
 439 assess agents' familiarity with all competitions in this benchmark. The results show no evidence  
 440 of systematic prior exposure or memorization, further supporting the integrity of the benchmark in  
 441 its current state. For case (2), the scalability of TimeSeriesGym enables efficient generation of  
 442 new challenges and skill-specific variations. This allows the benchmark to evolve continuously and  
 443 remain effective even if the current version is eventually included in future LLM pretraining.

444 Finally, we raise a broader question around plagiarism and code reuse. Several TimeSeriesGym  
 445 challenges, such as leveraging MOMENT (6) for anomaly detection, require agents to use existing code  
 446 repositories to solve open-ended ML tasks, making plagiarism difficult to assess. For example, if an  
 447 agent cites the code it uses, should it be considered plagiarism or appropriate reuse, similar to how  
 448 human ML practitioners build on public code with proper reference? As the ability to effectively  
 449 and properly leverage existing resources is important in real-world ML practice, we believe that it is  
 450 crucial to develop clear, legally correct definitions and evaluation criteria for data contamination and  
 451 plagiarism in the context of AI agents. We highlight this as an important direction for future work.

452 **Defining and measuring success.** What does it mean for an agent to be successful? For Kaggle tasks,  
 453 while comparing an agent's performance against human leaderboards seems intuitive, it presents  
 454 challenges. TimeSeriesGym utilizes different training and testing splits and re-implements the  
 455 grading mechanisms from the original Kaggle competitions (as original Kaggle test sets are private),  
 456 making direct leaderboard comparisons potentially misleading. Additionally, in challenges like code  
 457 migration, real-world utility does not always require perfect, bug-free code: partial, buggy solutions  
 458 may still accelerate development when iterated by human engineers. Thus, our current evaluation  
 459 approaches have inherent limitations.

460 We propose several desiderata for improving the success metrics. These metrics should be rigorous  
 461 and objective, yet flexible enough to preserve agent creativity and autonomy. They should also  
 462 yield actionable insights, helping identify specific deficiencies and guide future improvements in  
 463 agent design. In this work, we take a step in this direction by enabling skill-based and holistic  
 464 evaluations, offering a more comprehensive understanding of agent capabilities and limitations.  
 465 Moving forward, we believe that the development of robust, holistic and diagnostic success metrics  
 466 remains an important research direction and requires community discussion.

467 **Optimal resource allocation.** Consistent with previous work, agents were given 4 hours and 50 steps  
 468 to solve each challenge - but is this sufficient? Alternative frameworks like MLE-bench provide  
 469 substantially more resources (24 hours and approximately 2000 steps). Our scaling experiments,  
 470 which gave agents up to 12 hours and 150 steps for a subset of challenges, did not reveal significant  
 471 performance improvements. Therefore, we believe that further increasing resources is an option, but  
 472 practical academic budget constraints make such approaches largely infeasible. This raises important  
 473 questions about how to balance resource limitations with fair opportunities to assess AI agents.

474 **Societal impact.** AI agents promise to substantially reduce manual effort in ML engineering while  
 475 expanding the productivity and accessibility of ML tools. This automation presents several social  
 476 implications worth considering. First, by lowering technical barriers, these agents could democratize  
 477 ML capabilities, allowing users without an extensive programming background to leverage advanced  
 478 analytics. Second, automated ML workflows can accelerate scientific discovery across multiple  
 479 domains, including healthcare, climate science, and materials research. However, several challenges  
 480 require careful attention from the community. The primary concern is proper attribution when agents  
 481 repurpose existing code, potentially obscuring original authorship and violating licenses. Furthermore,  
 482 automated ML systems can perpetuate or amplify existing biases in training data without human  
 483 oversight. Furthermore, these agents might generate plausible but flawed solutions that appear  
 484 correct to non-experts, leading to undetected errors in critical applications. The resource-intensive  
 485 nature of running sophisticated agents could also exacerbate computational divides between well-  
 486 resourced and under-resourced organizations. As we advance agent capabilities through benchmarks

486 like TimeSeriesGym, the community must simultaneously develop frameworks for responsible  
 487 deployment that address these challenges while maximizing societal benefits.  
 488

489 **6 CONCLUSION**  
 490

491 We propose TimeSeriesGym, a scalable and agent-agnostic benchmarking framework for eval-  
 492 uating AI agents (with various scaffoldings) on time series ML engineering tasks. By curating  
 493 tasks that reflect real-world ML practice from diverse sources, enabling scalable task generation,  
 494 and supporting multimodal, skill-based, holistic evaluation, TimeSeriesGym provides a practical  
 495 and extensible testbed for advancing AI agents in ML engineering. Our experiments show that  
 496 while frontier LLMs combined with AIDE scaffolding (11) can achieve moderate to high success  
 497 rates in producing valid submissions, they still do not generate reasonable solutions, particularly on  
 498 *TimeSeriesGym* original challenges that emulate the complexity of real-world time series tasks.  
 499 This highlights current limitations in agent capabilities to effectively understand and solve realistic  
 500 time series tasks. By open-sourcing TimeSeriesGym, we aim to facilitate a deeper understanding  
 501 of the ML engineering capabilities of AI agents, provide actionable insights on future development,  
 502 and support the collection of agent interaction trajectories to drive continuous improvement of AI  
 503 agents through post-training and refinement.

504  
 505 **REPRODUCIBILITY STATEMENT**  
 506

507 We provide TimeSeriesGym as an open-source project under the permissive MIT License:  
 508 <https://anonymous.4open.science/r/TimeSeriesGym-9CF6/>. The repository in-  
 509 cludes detailed documentation on running experiments, adding new challenges, and incorporating  
 510 different agentic scaffolds. Tab. 4 lists all challenges in TimeSeriesGym, while Tab. 5 provides  
 511 their sources and licenses. We describe our exact experimental settings and compute resources in  
 512 Sec. 4, with scaffold hyperparameters detailed in Tab. 11. The cost to run each experiment is reported  
 513 in Tab. 7.

514  
 515 **REFERENCES**  
 516

- 517 [1] Ben Beglin, Kejuan Yang, Shashank Gupta, Kyle Richardson, Erin Bransom, Peter Clark, Ashish  
 518 Sabharwal, and Tushar Khot. SUPER: Evaluating agents on setting up and executing tasks  
 519 from research repositories. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen, editors,  
 520 *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*,  
 521 pages 12622–12645, Miami, Florida, USA, November 2024. Association for Computational  
 522 Linguistics.
- 523 [2] Yifu Cai, Arjun Choudhry, Mononito Goswami, and Artur Dubrawski. TimeSeriesExam: A  
 524 time series understanding exam. *NeurIPS’24 Time Series in the Age of Large Models Workshop*,  
 525 2024.
- 526 [3] Jun Shern Chan, Neil Chowdhury, Oliver Jaffe, James Aung, Dane Sherburn, Evan Mays, Giulio  
 527 Starace, Kevin Liu, Leon Maksin, Tejal Patwardhan, Aleksander Madry, and Lilian Weng.  
 528 MLE-bench: Evaluating machine learning agents on machine learning engineering. In *The*  
 529 *Thirteenth International Conference on Learning Representations*, 2025.
- 531 [4] Susan E Embretson and Steven P Reise. *Item response theory for psychologists*. Psychology  
 532 Press, 2013.
- 534 [5] Mononito Goswami, Vedant Sanil, Arjun Choudhry, Arvind Srinivasan, Chalisa Udompanyawit,  
 535 and Artur Dubrawski. Aqua: A benchmarking tool for label quality assessment. *Advances in*  
 536 *Neural Information Processing Systems*, 36:79792–79807, 2023.
- 538 [6] Mononito Goswami, Konrad Szafer, Arjun Choudhry, Yifu Cai, Shuo Li, and Artur Dubrawski.  
 539 MOMENT: A family of open time-series foundation models. In *International Conference on*  
 540 *Machine Learning*, pages 16115–16152. PMLR, 2024.

540 [7] Gauthier Guinet, Behrooz Omidvar-Tehrani, Anoop Deoras, and Laurent Callot. Automated  
 541 evaluation of retrieval-augmented language models with task-specific exam generation. In  
 542 *Forty-first International Conference on Machine Learning*, 2024.

543 [8] Cherie Ho, Jiaye Zou, Omar Alama, Sai Mitheran Jagadesh Kumar, Cheng-Yu Chiang, Taneesh  
 544 Gupta, Chen Wang, Nikhil Keetha, Katia Sycara, and Sebastian Scherer. Map it anywhere: Em-  
 545 powering bev map prediction using large-scale public datasets. *Advances in Neural Information  
 546 Processing Systems*, 37:64433–64453, 2024.

547 [9] Qian Huang, Jian Vora, Percy Liang, and Jure Leskovec. MLAgentbench: Evaluating language  
 548 agents on machine learning experimentation. In *Forty-first International Conference on Machine  
 549 Learning*, 2024.

550 [10] Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Ar-  
 551 mando Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination  
 552 free evaluation of large language models for code. *arXiv preprint arXiv:2403.07974*, 2024.

553 [11] Zhengyao Jiang, Dominik Schmidt, Dhruv Srikanth, Dixin Xu, Ian Kaplan, Deniss Ja-  
 554 cenko, and Yuxiang Wu. Aide: Ai-driven exploration in the space of code. *arXiv preprint  
 555 arXiv:2502.13138*, 2025.

556 [12] Liqiang Jing, Zhehui Huang, Xiaoyang Wang, Wenlin Yao, Wenhao Yu, Kaixin Ma, Hongming  
 557 Zhang, Xinya Du, and Dong Yu. DSbench: How far are data science agents from becoming  
 558 data science experts? In *The Thirteenth International Conference on Learning Representations*,  
 559 2025.

560 [13] Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga,  
 561 Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, et al. Holistic evaluation of  
 562 language models. *arXiv preprint arXiv:2211.09110*, 2022.

563 [14] Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. G-eval:  
 564 NLg evaluation using gpt-4 with better human alignment. In *Proceedings of the 2023 Conference  
 565 on Empirical Methods in Natural Language Processing*, pages 2511–2522, 2023.

566 [15] Frank J Massey Jr. The kolmogorov-smirnov test for goodness of fit. *Journal of the American  
 567 statistical Association*, 46(253):68–78, 1951.

568 [16] Grégoire Mialon, Clémentine Fourrier, Thomas Wolf, Yann LeCun, and Thomas Scialom.  
 569 GAIA: a benchmark for general ai assistants. In *The Twelfth International Conference on  
 570 Learning Representations*, 2023.

571 [17] Deepak Nathani, Lovish Madaan, Nicholas Roberts, Nikolay Bashlykov, Ajay Menon, Vin-  
 572 cent Moens, Amar Budhiraja, Despoina Magka, Vladislav Vorotilov, Gaurav Chaurasia,  
 573 Dieuwke Hupkes, Ricardo Silveira Cabral, Tatiana Shavrina, Jakob Foerster, Yoram Bachrach,  
 574 William Yang Wang, and Roberta Raileanu. MLGym: A new framework and benchmark for  
 575 advancing ai research agents, 2025.

576 [18] Harshith Padigela, Chintan Shah, and Dinkar Juyal. ML-Dev-Bench: Comparative analysis of  
 577 ai agents on ml development workflows. *arXiv preprint arXiv:2502.00964*, 2025.

578 [19] Megan Risdal and Timo Bozslik. Meta kaggle, 2022.

579 [20] Xiangru Tang, Yuliang Liu, Zefan Cai, Yanjun Shao, Junjie Lu, Yichi Zhang, Zexuan Deng,  
 580 Helan Hu, Kaikai An, Ruijun Huang, et al. ML-Bench: Evaluating large language models and  
 581 agents for machine learning tasks on repository-level code. *arXiv preprint arXiv:2311.09835*,  
 582 2023.

583 [21] Xingyao Wang, Boxuan Li, Yufan Song, Frank F Xu, Xiangru Tang, Mingchen Zhuge, Jiayi  
 584 Pan, Yueqi Song, Bowen Li, Jaskirat Singh, et al. Openhands: An open platform for ai software  
 585 developers as generalist agents. In *The Thirteenth International Conference on Learning  
 586 Representations*, 2024.

594 [22] Hjalmar Wijk, Tao Lin, Joel Becker, Sami Jawhar, Neev Parikh, Thomas Broadley, Lawrence  
595 Chan, Michael Chen, Josh Clymer, Jai Dhyani, et al. Re-bench: Evaluating frontier ai r&d  
596 capabilities of language model agents against human experts. *arXiv preprint arXiv:2411.15114*,  
597 2024.

598 [23] Wen Ye, Yizhou Zhang, Wei Yang, Lumingyuan Tang, Defu Cao, Jie Cai, and Yan Liu. Beyond  
599 forecasting: Compositional time series reasoning for end-to-end task execution. *arXiv preprint*  
600 *arXiv:2410.04047*, 2024.

601  
602  
603  
604  
605  
606  
607  
608  
609  
610  
611  
612  
613  
614  
615  
616  
617  
618  
619  
620  
621  
622  
623  
624  
625  
626  
627  
628  
629  
630  
631  
632  
633  
634  
635  
636  
637  
638  
639  
640  
641  
642  
643  
644  
645  
646  
647

648 **A TIMESERIESGYM CHALLENGES**  
649

650	Challenge	Problem	Domain	Skills	Evaluation Metric
<b>Kaggle Challenges</b>					
652	AMP-Parkinson's Disease Progression Prediction	Time-to-Event Regression	Healthcare		Symmetric Mean Absolute Percentage Error
653	ASHRAE - Great Energy Predictor III	Forecasting	Energy		Root Mean Square Logarithmic Error
654	Child Mind Institute- Detect Sleep States	Classification	Healthcare		Event Detection Average Precision
655	Google Brain - Ventilator Pressure Prediction	Regression	Healthcare	Data Handling	Mean Absolute Error
656	G2Net Gravitational Wave Detection	Classification	Geology	(Dealing with Missing Values, Utilize Multi-Source Data)	Area Under ROC Curve
657	HMS - Harmful Brain Activity Classification	Classification	Healthcare		KL Divergence
658	LANL Earthquake Prediction	Time-to-Event Regression	Geology	Modeling	Mean Absolute Error
659	M5 Forecasting - Accuracy	Forecasting	Sales	(Hyper-parameter Tuning & Model Selection)	Weighted Root Mean Squared Scaled Error
660	Online Product Sales	Forecasting	Sales		Root Mean Square Logarithmic Error
661	Optiver Realized Volatility Prediction	Forecasting	Finance		Root Mean Square Percentage Error
662	OSIC Pulmonary Fibrosis Progression	Forecasting	Healthcare		Laplace Log Likelihood
663	Recruit Restaurant Visitor Forecasting	Forecasting	Sales		Root Mean Square Logarithmic Error
664	Sberbank Russian Housing Market	Forecasting	Housing		Root Mean Square Logarithmic Error
<b>TimeSeriesGym Originals</b>					
666	Convert ResNet TensorFlow Implementation to PyTorch	Classification			
667	Convert STOMP Algorithm Implementation in R to Python	Data Mining	Algorithm	Code Migration	Custom Code Grading
668	Evaluate MOIRAI time series foundation model on the Context Is Key (CiK) benchmark	Context-aided Forecasting	Climatology, Economics, Energy, Mechanics, Public Safety, Retail, Synthetic, Transportation		Resolved (Binary)
669	Evaluate Chronos time series foundation model on the NNS dataset within Context Is Key (CiK) benchmark				
670	Implement & Evaluate CSDI to Impute PM2.5 Data Train & Evaluate CSDI to Impute PM2.5 Data	Imputation	Weather		Mean Absolute Error
671	GIFT-EVAL: A Benchmark for General Time Series Forecasting Model Evaluation	Forecasting	Nature, Web, CloudOps, Economics/Finance, Energy, Sales, Transportation,	Modeling	Mean Absolute Percentage Error
672	Hexagon ML UCR Time Series Anomaly Detection	Anomaly Detection	Healthcare, Gait, Energy, Synthetic, Devices	(Using Research Code)	Adjusted Best F1 Score
673	Long Horizon Time Series Forecasting using Time Series Library	Forecasting	Energy, Epidemiology, Finance, Transportation, Weather		Mean Squared Error
674	Long-Horizon Weather Forecasting using Time Series Library's Itransformer	Forecasting	Weather		Exact Match
675	MIT-BIH ECG Arrhythmia Detection	Classification	Healthcare		Accuracy
676	MOMENT for Anomaly Detection on UCR datasets	Anomaly Detection	Healthcare, Gait, Energy, Synthetic, Devices		Exact Match
677	PTB-XL ECG Classification	Classification	Healthcare		Accuracy
678	TimeSeriesExam: A Time Series Understanding Exam	Time Series Understanding	Synthetic	Time Series Understanding	Accuracy
<b>Derived Challenges</b>					
683	Google Brain - Ventilator Pressure Prediction	Regression	Healthcare	Data Handling (Dealing with missing data)	Mean Absolute Error
684	Improve PTB-XL ECG Classification Code	Classification	Healthcare	Code Enhancement (Experiment Tracking, Readability, Reproducibility)	
685	MIT-BIH Arrhythmia Detection with Weak Supervision	Classification	Healthcare	Data Handling (Labeling)	Accuracy
686	Optiver Realized Volatility Prediction	Forecasting	Finance	Data Handling (Dealing with missing data)	Root Mean Square Percentage Error
687	Optiver Realized Volatility Prediction with Hyper-parameter Optimization	Forecasting	Finance	Modeling (Hyper-parameter Tuning & Model Selection)	Improvement in Root Mean Square Percentage Error
688	PTB-XL ECG Classification with Hyperparameter Optimization	Classification	Healthcare		Improvement in Accuracy

692 Table 4: This table presents the `TimeSeriesGym` benchmark's diverse collection of time series challenges across three categories: Kaggle Challenges, TimeSeriesGym Originals, and Derived Challenges. The challenges span multiple domains (healthcare, finance, energy, weather, transportation), problem types (classification, regression, forecasting, anomaly detection), and required skills (data handling, model building, code migration). Each challenge uses appropriate evaluation metrics for its task type. The benchmark combines established Kaggle competitions with novel custom tasks, creating a comprehensive testbed for evaluating ML engineering agents across realistic scenarios that practitioners face in real-world applications.

702	Challenge	Source	License	
703	<b>Kaggle Challenges</b>			
704	AMP-Parkinson's Disease Progression Prediction	Kaggle	Subject to Competition Rules	
705	ASHRAE - Great Energy Predictor III	Kaggle	Subject to Competition Rules	
706	Child Mind Institute– Detect Sleep States	Kaggle	CC BY-NC-SA 4.0	
707	Google Brain - Ventilator Pressure Prediction	Kaggle	Subject to Competition Rules	
708	G2Net Gravitational Wave Detection	Kaggle	Subject to Competition Rules	
709	HMS - Harmful Brain Activity Classification	Kaggle	CC BY-NC 4.0	
710	LANL Earthquake Prediction	Kaggle	Subject to Competition Rules	
711	M5 Forecasting - Accuracy	Kaggle	Subject to Competition Rules	
712	Online Product Sales	Kaggle	Subject to Competition Rules	
713	Optiver Realized Volatility Prediction	Kaggle	Subject to Competition Rules	
714	OSIC Pulmonary Fibrosis Progression	Kaggle	Subject to Competition Rules	
715	Recruit Restaurant Visitor Forecasting	Kaggle	Subject to Competition Rules	
716	Sberbank Russian Housing Market	Kaggle	Subject to Competition Rules	
717	<b>TimeSeriesGym Originals</b>			
718	Convert ResNet TensorFlow Implementation to PyTorch	GitHub	GNU General Public License v3.0	
719	Convert STOMP Algorithm Implementation in R to Python	GitHub	Apache License 2.0	
720	Evaluate MOIRAI time series foundation model on the	GitHub	Apache License 2.0	
721	Context Is Key (CiK) benchmark			
722	Evaluate Chronos time series foundation model on the NN5 dataset within Context Is Key (CiK) benchmark	–	Apache License 2.0	
723	Implement & Evaluate CSDI to Impute PM2.5 Data	GitHub	MIT License	
724	Train & Evaluate CSDI to Impute PM2.5 Data	–	MIT License	
725	GIFT-EVAL: A Benchmark for General Time Series Forecasting Model Evaluation	GitHub	Apache License 2.0	
726	Hexagon ML UCR Time Series Anomaly Detection	UCR	Not available	
727	Long Horizon Time Series Forecasting using Time Series Library	GitHub	MIT License	
728	Long-Horizon Weather Forecasting using Time Series Library's Itransformer	–	MIT License	
729	MIT-BIH ECG Arrhythmia Detection	PhysioNet	Open Data Commons Attribution License v1.0	
730	MOMENT for Anomaly Detection on UCR datasets	GitHub	MIT License	
731	PTB-XL ECG Classification	PhysioNet	Creative Commons Attribution 4.0 International Public License	
732	TimeSeriesExam: A Time Series Understanding Exam	Hugging Face	MIT License	

751 Table 5: This table provides transparency about the source and licensing information for each  
752 challenge in the TimeSeriesGym benchmark. For the Kaggle challenges, most are subject to  
753 Kaggle's competition rules, with a few under Creative Commons licenses. The TimeSeriesGym  
754 Original challenges come from diverse sources including GitHub repositories, HuggingFace, etc.  
755 with various open-source licenses (Apache, MIT, GPL, Creative Commons). This diversity of sources  
and licenses demonstrates the benchmark's foundation in accessible, reusable datasets and code while  
ensuring proper attribution and compliance with original creators' terms.

756  
 757 Table 6: TimeSeriesGym -Lite is a streamlined collection of six diverse time series challenges  
 758 carefully selected to evaluate AI agents while balancing comprehensiveness with efficiency. The  
 759 challenges cover essential ML engineering skills including basic data science, handling missing/multi-  
 760 source data, code migration, hyperparameter optimization, modeling using research code, and data  
 761 labeling. The collection spans multiple domains (healthcare, finance, algorithms) and various time  
 762 series tasks (classification, forecasting, anomaly detection, code migration). This cost-effective subset  
 763 allows researchers to quickly benchmark agent capabilities across critical ML engineering skills  
 764 without the resource requirements of the full TimeSeriesGym benchmark.

Challenge	Required Skills	Time Series Task	Domain
Child Mind Institute - Detect Sleep States	Basic data science (data handling and modeling)	Classification	Healthcare
Optiver Realized Volatility Prediction	Handling missing and multi-source data	Forecasting	Finance
Convert ResNet TensorFlow implementation to PyTorch	Classification	Code Migration	Algorithm
PTB-XL ECG Classification	Hyperparameter optimization & model selection	Classification	Healthcare
MOMENT Anomaly Score Calculation	Modeling (Using research code)	Anomaly Detection	Healthcare, Gait, Synthetic, Energy, Devices
MIT-BIH Arrhythmia Detection	Data labeling	Classification	Healthcare

779  
 780 Table 7: Average cost to run experiments on a single seed in the default evaluation setup *i.e.* AIDE  
 781 with gpt-4.1-2025-04-14 with a maximum of 4 hours and 50 steps.

Benchmark	Average Cost (USD)
TimeSeriesGym	62.12
TimeSeriesGym -Lite	7.96

787  
 788 Table 8: The 100 shortlisted Kaggle competitions. Competitions marked  
 789 with \* denote that the data is no longer available.

Competition	# Participants	Reward	Category
M5 Forecasting - Accuracy	7022	50,000	Featured
LANL Earthquake Prediction	5454	50,000	Research
Jane Street Market Prediction*	4884	100,000	Featured
Optiver Realized Volatility Prediction	4395	100,000	Featured
Optiver - Trading at the Close	4374	100,000	Featured
ASHRAE - Great Energy Predictor III	4342	25,000	Featured
Zillow Prize: Zillow's Home Value Prediction (Zestimate)	4241	1,200,000	Featured
GoDaddy - Microbusiness Density Forecasting	3834	60,000	Featured
Rossmann Store Sales	3735	35,000	Featured
Sberbank Russian Housing Market	3658	25,000	Featured
HMS - Harmful Brain Activity Classification	3507	50,000	Research
Google Brain - Ventilator Pressure Prediction	3118	7,500	Research
University of Liverpool - Ion Switching	3004	25,000	Research
Ubiquant Market Prediction*	2949	100,000	Featured
Enefit - Predict Energy Behavior of Prosumers	2715	50,000	Featured
OSIC Pulmonary Fibrosis Progression	2530	55,000	Featured
Child Mind Institute - Detect Sleep States	2436	50,000	Featured
Recruit Restaurant Visitor Forecasting	2426	25,000	Featured

Continued on next page

	Competition	# Participants	Reward	Category
810	G-Research Crypto Forecasting	2398	125,000	Featured
811	Two Sigma Financial Modeling Challenge*	2317	100,000	Featured
812	Grupo Bimbo Inventory Demand	2263	25,000	Featured
813	AMP®-Parkinson's Disease Progression Prediction	2197	60,000	Featured
814	Corporación Favorita Grocery Sales Forecasting	1868	30,000	Featured
815	Driver Telematics Analysis	1861	30,000	Featured
816	Parkinson's Freezing of Gait Prediction	1688	100,000	Research
817	COVID19 Global Forecasting (Week 1)	640	0	Research
818	Heritage Health Prize*	1656	500,000	Featured
819	Cornell Birdcall Identification	1630	25,000	Research
820	TensorFlow Speech Recognition Challenge	1591	25,000	Featured
821	Google - American Sign Language Fingerspelling Recognition	1529	200,000	Research
822	G2Net Gravitational Wave Detection	1501	15,000	Research
823	COVID19 Global Forecasting (Week 4)	388	0	Research
824	Indoor Location & Navigation	1446	10,000	Research
825	West Nile Virus Prediction	1445	40,000	Featured
826	BirdCLEF 2023	1397	50,000	Research
827	Rainforest Connection Species Audio Detection	1385	15,000	Research
828	COVID19 Global Forecasting (Week 3)	290	0	Research
829	COVID19 Global Forecasting (Week 2)	263	0	Research
830	Google Analytics Customer Revenue Prediction	1369	45,000	Featured
831	COVID19 Local US-CA Forecasting (Week 1)	216	0	Research
832	Google - Isolated Sign Language Recognition	1340	100,000	Research
833	1st and Future - Player Contact Detection	1334	100,000	Featured
834	JPX Tokyo Stock Exchange Prediction	1324	63,000	Featured
835	Google Research Football with Manchester City F.C.	1288	6,000	Featured
836	Lyft Motion Prediction for Autonomous Vehicles	1254	30,000	Featured
837	BirdCLEF 2024	1198	50,000	Research
838	COVID19 Global Forecasting (Week 5)	93	0	Research
839	G2Net Detecting Continuous Gravitational Waves	1181	25,000	Research
840	Peking University/Baidu - Autonomous Driving	1105	25,000	Featured
841	iWildcam 2021 - FGVC8	65	0	Research
842	Eye Movements Verification and Identification Competition	50	0	Research
843	M5 Forecasting - Uncertainty	1101	50,000	Featured
844	March Machine Learning Mania 2023	1098	50,000	Featured
845	Multi-label Bird Species Classification - NIPS 2013	39	0	Research
846	Google Cloud & NCAA® ML Competition 2018-Men's	1061	50,000	Featured
847	iWildCam 2022 - FGVC9	29	0	Research
848	NFL Health & Safety - Helmet Assignment	1028	100,000	Featured
849	March Machine Learning Mania 2022 - Men's	1025	25,000	Featured
850	BirdCLEF 2022	1009	10,000	Research
851	BirdCLEF 2021 - Birdcall Identification	1001	5,000	Research
852	Google Smartphone Decimeter Challenge	985	10,000	Research
853	SETI Breakthrough Listen - E.T. Signal Search	979	15,000	Research
854	LEAP - Atmospheric Physics using AI (ClimSim)	877	50,000	Research
855	Bengali.AI Speech Recognition	866	53,000	Research
856	The Winton Stock Market Challenge	829	50,000	Featured
857	Two Sigma: Using News to Predict Stock Movements	813	100,000	Featured
858	Accelerometer Biometric Competition	770	5,000	Research
859	How Much Did It Rain? II	691	500	Research
860	Google Smartphone Decimeter Challenge 2022	684	10,000	Research
861	American Epilepsy Society Seizure Prediction Challenge	653	25,000	Research
862	Melbourne University AES/MathWorks/NIH Seizure Prediction	645	20,000	Research
863	DFL - Bundesliga Data Shootout	610	25,000	Featured
	NFL 1st and Future - Impact Detection	573	75,000	Featured

Continued on next page

	Competition	# Participants	Reward	Category
864	Kore 2022	537	15,000	Featured
865	MLB Player Digital Engagement Forecasting	495	50,000	Featured
866	ECML/PKDD 15: Taxi Trajectory Prediction (I)	459	250	Research
867	Grasp-and-Lift EEG Detection	451	10,000	Research
868	The Big Data Combine Engineered by BattleFin	424	18,500	Research
869	Web Traffic Time Series Forecasting	424	25,000	Research
870	Draper Satellite Image Chronology	422	75,000	Featured
871	ECML/PKDD 15: Taxi Trip Time Prediction (II)	418	250	Research
872	Online Product Sales	412	22,500	Featured
873	Halite by Two Sigma	1291	0	Featured
874	RTA Freeway Travel Time Prediction	376	10,000	Featured
875	How Much Did It Rain?	349	500	Research
876	The 3rd YouTube-8M Video Understanding Challenge	340	25,000	Research
877	Benchmark Bond Trade Price Challenge	316	17,500	Featured
878	BCI Challenge @ NER 2015	311	1,000	Research
879	The Marinexplore and Cornell University Whale Detection Challenge	309	10,000	Featured
880	DecMeg2014 - Decoding the Human Brain	301	5,000	Research
881	U.S. Census Return Rate Challenge	290	25,000	Featured
882	dunnhumby's Shopper Challenge	287	10,000	Featured
883	KDD Cup 2012, Track 2	275	8,000	Featured
884	UPenn and Mayo Clinic's Seizure Detection Challenge	241	8,000	Research
885	GE Flight Quest	234	250,000	Featured
886	The Random Number Grand Challenge	206	1,000	Featured
887	AMS 2013-2014 Solar Energy Prediction Contest	199	1,000	Research
888	Global Energy Forecasting Competition 2012 - Wind Forecasting	197	7,500	Research
889	Belkin Energy Disaggregation Competition	194	25,000	Featured
890	CVPR 2018 WAD Video Segmentation Challenge	188	2,500	Research

891

892

893

894

895

896

897

898

899

900

901

902

903

904

905

906

907

908

909

910

911

912

913

914

915

916

917

918 B REPRESENTATIVENESS OF TIMESERIESGYM-LITE  
919920 To ensure that TimeSeriesGym-Lite provides broad coverage of machine learning capabilities  
921 while enabling low-cost evaluation, we construct the subset by selecting a diverse set of tasks from the  
922 full TimeSeriesGym benchmark. We further validate this selection through statistical comparisons  
923 of domain coverage and task difficulty distributions between the two benchmarks.  
924925 **Domain Coverage.** Table 9 summarizes the domain distributions of TimeSeriesGym and  
926 TimeSeriesGym-Lite. A Chi-square test indicates no significant difference between the two  
927 distributions ( $p = 0.92$ ,  $\chi^2 = 1.99$ ), demonstrating that the Lite subset maintains the same domain  
928 diversity as the full benchmark.  
929

930 Domain	931 <small>TimeSeriesGym</small>	932 <small>TimeSeriesGym</small> -Lite
933 Healthcare	934 11	935 3
936 Multi-domain	937 7	938 1
939 Commerce & Finance	940 6	941 1
942 Weather	943 3	944 0
945 Geology	946 2	947 0
948 Housing	949 1	950 0
951 Energy	952 1	953 0

954 Table 9: Domain distributions of TimeSeriesGym and TimeSeriesGym-Lite.  
955956 **Task Difficulty.** We additionally compare task difficulty levels between the two benchmarks, shown  
957 in Table 10. A Chi-square test again reveals no significant difference ( $p = 0.95$ ,  $\chi^2 = 0.09$ ),  
958 indicating that the Lite subset preserves the difficulty profile of the full benchmark.  
959

960 Difficulty	961 <small>TimeSeriesGym</small>	962 <small>TimeSeriesGym</small> -Lite
963 Low	964 13	965 2
966 Medium	967 15	968 3
969 High	970 5	971 1

972 Table 10: Difficulty distributions of TimeSeriesGym and TimeSeriesGym-Lite.  
973974 **Summary.** These statistical tests confirm that TimeSeriesGym-Lite is a representative sub-  
975 set of TimeSeriesGym in both domain diversity and task difficulty. This ensures that the Lite  
976 benchmark provides reliable, low-cost evaluation while preserving the characteristics of the full task  
977 suite.  
978

972 C IMPLEMENTATION DETAILS FOR SCAFFOLDS  
973  
974975 Table 11: Scaffold hyperparameters. `$TARGET_MODEL` denotes the model being evaluated.  
976

AIDE	
agent.code.model	<code>\$TARGET_MODEL</code>
agent.feedback.model	gpt-4.1-2025-04-14
agent.steps	50
agent.search.max_debug_depth	20
agent.search.debug_prob	1
agent.time_limit	14400
exec.timeout	32400

OpenHands	
agent	CodeActAgent
model	<code>\$TARGET_MODEL</code>
max_time_in_hours	4
max_steps	50

MLAgentBench (MLAB)	
llm-name	<code>\$TARGET_MODEL</code>
edit-script-llm-name	<code>\$TARGET_MODEL</code>
fast-llm-name	gpt-4.0-mini
max-steps	50
max-time	14400

977  
978 Tab. 11 summarizes the hyperparameters used for the three scaffolds: AIDE (11), OpenHands'  
979 CodeAct (21), and MLAGentBench (9). We did not directly use the MLE-bench (3) modifications  
980 to the agentic scaffolds for two main reasons. First, the official scaffold implementation has undergone  
981 updates since MLE-bench. Second, certain changes made in MLE-bench were not applicable to our  
982 benchmark, which involves more diverse modes of evaluation. Therefore, we modified the latest  
983 version of the agentic scaffolds to enhance robustness, improve execution stability, and support a  
984 broader range of competitions.

1004 C.1 AIDE  
1005

1007 We forked the original AIDE repository in May 2025 and added useful changes from the MLE-bench  
1008 project. These include better API calls and support for more API providers. Our key modifications  
1009 are summarized below:

1010 1. Updated instructions to cover all types of tasks, not just Kaggle challenges  
1011 2. Modified the prompt to work with different file types, since our tasks use many input/output  
1012 formats  
1013 3. Fixed the interpreter handling to prevent timeouts and system hangs

1015 C.2 OPENHANDS  
1016

1018 We forked the original OpenHands repository from tag v0.34.0 (May 2025). We reduced the  
1019 RAM allowance to 10 GiB (from 100 GiB) as we did not observe any memory-related issues during  
1020 our tests.

1021 C.3 MLAGENTBENCH  
1022

1023 We forked the MLE-bench version of MLAGentBench with no modifications.  
1024

1025

1026  
1027  
1028  
1029  
1030  
1031  
1032  
1033  
1034

Figure 2: The prompt we use to initialize all scaffolds, adapted from MLE-bench (3).

```

1035 # Machine Learning Engineer Challenge: TimeSeriesGym Competition
1036
1037 ## Your Role
1038 You are a Machine Learning Engineer tasked with solving complex ML problems in the
1039 ↳ TimeSeriesGym offline competition. This competition tests your ability to understand
1040 ↳ problems, follow instructions, and build/train/evaluate ML models.
1041
1042 ## Core Requirements
1043 - You must solve each task by training a model and generating predictions on the test set
1044 - Hand-labeling the data to produce submissions is strictly prohibited
1045
1046 ## Environment Setup
1047
1048 ### Critical Files and Locations
1049 - **Competition Instructions**: Located at `/home/data/description.md` - This is your primary
1050 ↳ reference for understanding the specific problem
1051 - **Dataset**: All data files are in `/home/data/` directory
1052 - **Submission Path**: Your final submission MUST be saved to `/home/submission/` in the
1053 ↳ format specified in the instructions
1054 - **Environment Instructions**: These general instructions can be found at
1055 ↳ `/home/instructions.txt`
1056
1057 ### Test Dataset Information
1058 - We've created our own train/test splits for evaluation purposes
1059 - The data in `/home/data/` may differ slightly from what's described in the competition
1060 ↳ instructions
1061 - Use the provided train/test splits for your work
1062
1063 ## Submission Instructions
1064
1065 ### File Format and Location
1066 - Your submission MUST be saved to the exact path: `/home/submission/`
1067 - The file format must strictly follow the specifications in `description.md`
1068 - Failing to use the correct path or format will result in a failed evaluation
1069
1070 ### Validation Process
1071 1. Before finalizing your submission, validate its format using the validation server
1072 2. Use the command: `curl -X POST -F "file=@${SUBMISSION_FILE}"`  

1073 ↳ `http://localhost:5000/validate`  

1074 3. Alternatively, use the provided script: `/home/validate_submission.sh`  

1075 4. The validator will confirm if your submission format is valid (but won't provide a score)
1076 5. Make corrections if needed and revalidate until you receive a confirmation
1077
1078 ### Final Submission
1079 - Once validated, ensure your final submission file remains at `/home/submission/`
1080 - Only the last file in this location will be evaluated
1081 - Double-check that your model's predictions cover the entire test set
1082 - Review the submission format requirements one last time before completing
1083
1084 ## Evaluation Criteria
1085 - Your submission will be graded based on the metric specified in the competition
1086 ↳ instructions
1087 - Your goal is to achieve the highest possible score
1088
1089 ## Important Rules
1090 - Viewing or copying others' solutions is prohibited and will result in disqualification
1091 - In case of conflicts between these instructions and `/home/data/description.md`, these
1092 ↳ general instructions take priority
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179

```

1080 **D TIMESERIESGYM VERSUS MLE-BENCH**  
10811082 **MLE-Bench** (3) is a recent data machine learning engineering benchmark that curates 75 Kaggle  
1083 competitions of varying complexity and evaluates agent performance by grading prediction outputs  
1084 submitted in CSV files. While **TimeSeriesGym** also incorporates Kaggle competitions, it addresses  
1085 a broader scope: comprehensive ML engineering capabilities, of which competition solving represents  
1086 one important component. This section details the key distinctions between the two benchmarks. We  
1087 highlight key differences below.  
10881089 **Time Series Coverage.** As shown in Table 1 of the main paper, **TimeSeriesGym** contains sub-  
1090 stantially more time series modeling challenges than any existing ML engineering benchmark. Time  
1091 series represents an important yet underrepresented modality in agent evaluation. We demonstrate  
1092 that complex real-world time series problems (e.g., PTB-XL ECG classification) can be reformulated  
1093 into fully automated, agentic evaluation pipelines—a capability not established by prior benchmarks.  
10941095 **Task Diversity Beyond Competitions.** While **MLE-Bench** exclusively sources its 75 tasks from  
1096 Kaggle competitions, **TimeSeriesGym** combines three sources: Kaggle competitions, GitHub  
1097 repositories, and original hand-crafted challenges. Kaggle competitions alone do not capture the  
1098 full spectrum of real-world ML engineering tasks. Critical capabilities such as hyperparameter  
1099 search strategies, repository utilization, and API integration are not isolated or directly measured by  
1100 competition-based benchmarks. Our original challenges, designed from years of ML engineering  
1101 experience, represent these realistic workflows. Additionally, our extensive documentation framework  
1102 facilitates community contributions of both Kaggle-based and original challenges.  
11031104 **Granular Skill Assessment.** **TimeSeriesGym** provides skill-specific simulators (e.g., missing  
1105 data handling, hyperparameter optimization, feature engineering) that enable targeted evaluation  
1106 of individual agent capabilities. These modular assessments allow researchers to diagnose specific  
1107 strengths and weaknesses of LLM agents, rather than only measuring end-to-end performance.  
11081109 **Benchmark Difficulty.** **TimeSeriesGym** poses substantial challenges for state-of-the-art sys-  
1110 tems. At the time of **MLE-Bench** publication, o1-preview with the AIDE scaffold achieved  
1111 medal performance in 16.9% of the 75 challenges. In contrast, our evaluation reveals that AIDE  
1112 + GPT-4o achieved above-median performance (percentile score >0.5) in only 3 of 13 Kaggle  
1113 competitions (23%), indicating significant gaps in current agents’ time series modeling and ML  
1114 engineering capabilities.  
11151116 **Multi-Artifact Evaluation.** Unlike **MLE-Bench**, which evaluates only CSV prediction files,  
1117 **TimeSeriesGym** assesses multiple output types: prediction files, model artifacts (e.g., trained  
1118 models, checkpoints), and code implementations. This multi-artifact approach better reflects real-  
1119 world ML engineering practice, where deliverables extend beyond final predictions.  
11201121 **Holistic Assessment Methodology.** Our evaluation protocol combines three complementary ap-  
1122 proaches: (1) quantitative metrics (e.g., accuracy, MAE, F1-score), (2) programmatic analysis (e.g.,  
1123 regex matching, code structure inspection), and (3) optional qualitative assessment via LLM-as-a-  
1124 judge. This multi-faceted evaluation provides comprehensive insight into agent capabilities beyond  
1125 single-metric performance.  
11261127 The following table summarizes the key differences between **MLE-Bench** and **TimeSeriesGym**  
1128 across three dimensions: task source, ML capability coverage, and evaluation protocol.  
11291130 In summary, while **MLE-Bench** provides valuable evaluation of agents on 75 competition-style  
1131 ML problems using standardized CSV outputs, **TimeSeriesGym** complements and extends this  
1132 evaluation paradigm by emphasizing time series modeling, incorporating diverse task types that  
1133 reflect real-world engineering workflows, and providing granular skill assessment alongside holistic  
multi-artifact evaluation.

1134

1135

Table 12: Comparison of MLE-Bench and TimeSeriesGym

1136

1137

1138

1139

1140

1141

1142

1143

1144

1145

1146

## E DETAILED EVALUATION RESULTS

1147

1148

### E.1 FULL BENCHMARK EVALUATION RESULT

1149

1150

1151

1152

1153

1154

1155

1156

1157

1158

### E.2 ABLATION STUDY EVALUATION RESULT

1159

1160

1161

1162

1163

We provide detailed evaluation results for each task in TimeSeriesGym in Tab. 13. Each task was

executed with three random seeds; we report both the average and best scores across these runs.

Entries marked N/A indicate that the agent failed to produce a valid solution due to exceeding the time- or step-limit. For the GIFT-Eval and UCR Anomaly Detection challenges, evaluation is performed on a subset of the original benchmark, since our focus is on assessing the agent’s ability to leverage the research repository rather than full benchmark performance.

1164

1165

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181

1182

1183

1184

1185

1186

1187

1188  
1189  
1190  
1191  
1192  
1193

	Challenge	Evaluation Metric	Best @ 3	Average @ 3	Percentile Best @ 3
<b>Kaggle Challenges</b>					
1194	AMP-Parkinson's Disease Progression Prediction	Symmetric Mean Absolute Percentage Error	111.22	120.50	0.01385
1195	ASHRAE - Great Energy Predictor III	Root Mean Square	1.02	1.92	0.68234
1196	Child Mind Institute- Detect Sleep States	Logarithmic Error			
1197	Google Brain - Ventilator Pressure Prediction	Event Detection	0.02	0.01	0.07082
1198	G2Net Gravitational Wave Detection	Average Precision			
1199	HMS - Harmful Brain Activity Classification	Mean Absolute Error	0.58	5.40	0.13896
1200	LANL Earthquake Prediction	Area Under ROC Curve	0.51	0.50	0.26372
1201	M5 Forecasting - Accuracy	KL Divergence	1.16	1.56	0.03831
1202	Online Product Sales	Mean Absolute Error	2.18	2.89	1.0
1203	Optiver Realized Volatility Prediction	Weighted Root Mean Squared	0.82	3.13	0.65532
1204	OSIC Pulmonary Fibrosis Progression	Scaled Error			
1205	Recruit Restaurant Visitor Forecasting	Root Mean Square	0.91	1.08	0.20000
1206	Sberbank Russian Housing Market	Logarithmic Error			
1207		Root Mean Square	0.28	0.30	0.20425
		Laplace Log Likelihood	-7.39	-12.87	0.89318
		Root Mean Square	0.55	0.60	0.29532
		Logarithmic Error			
		Root Mean Square	0.39	0.40	0.12221
		Logarithmic Error			
<b>TimeSeriesGym Originals</b>					
1208	Convert ResNet TensorFlow implementation to PyTorch	Custom Code Grading Test Cases	5/9	5/9	0
1209	Convert STOMP Algorithm implementation in R to Python	Custom Code Grading Test Cases	2/4	1.6/4	0
1210	Evaluate MOIRAI time series foundation model on the Context Is Key (CiK) benchmark	Resolved (Binary)	N/A	N/A	0
1211	Evaluate Chronos time series foundation model on the NNS dataset within Context Is Key (CiK) benchmark	Resolved (Binary)	N/A	N/A	0
1212	Implement & Evaluate CSDI to Impute PM2.5 Data	Mean Absolute Error	N/A	N/A	0
1213	Train & Evaluate CSDI to Impute PM2.5 Data	Mean Absolute Error	N/A	N/A	0
1214	GIFT-EVAL: A Benchmark for General Time Series* Forecasting Model Evaluation	Mean Absolute Percentage Error	N/A	N/A	0
1215	Hexagon ML UCR Time Series Anomaly Detection*	Adjusted Best	0.38	0.38	0
1216	Long Horizon Time Series Forecasting Using Time Series Library	F1 Score			
1217	Long-Horizon Weather Forecasting using Time Series Library's Itransformer	Mean Squared Error	N/A	N/A	0
1218	MIT-BIH ECG Arrhythmia Detection	Exact Match (Binary)	N/A	N/A	0
1219	MOMENT for Anomaly Detection on UCR datasets	Accuracy	0.87	0.84	0
1220	PTB-XL ECG Classification	Exact Match (Binary)	N/A	N/A	0
1221	TimeSeriesExam: A Time Series Understanding Exam	Accuracy	0.81	0.80	0
1222		Accuracy	N/A	N/A	0
<b>Derived Challenges</b>					
1223	Google Brain - Ventilator Pressure Prediction	Mean Absolute Error	2.72	6.66	0.15047
1224	With Missingness	Code Enhancement (Experiment Tracking, Readability, Reproducibility)	N/A	N/A	0
1225	Improve PTB-XL ECG Classification Code	Accuracy	0.87	0.77	0
1226	MIT-BIH Arrhythmia Detection with Weak Supervision	Root Mean Square Percentage Error	0.33	0.33	0.13888
1227	Optiver Realized Volatility Prediction	Improvement in Root Mean Square Percentage Error	-0.01	-0.15	0
1228	With Missingness	Improvement in Accuracy	0.08	0.03	0
1229	Optiver Realized Volatility Prediction with Hyper-parameter Optimization				
1230	PTB-XL ECG Classification with Hyperparameter Optimization				
1231					

1232 Table 13: Comprehensive performance metrics for AI agents on all TimeSeriesGym challenges,  
1233 including best and average scores from three runs. Agents struggle to solve TimeSeriesGym Orig-  
1234 inal challenges. Derived challenges demonstrate how added complexity (missingness, hyperparameter  
1235 optimization) affects performance. These results highlight both the capabilities and limitations of  
1236 current ML engineering agents across diverse time series tasks.

1237  
1238  
1239  
1240  
1241

Challenge	8 hours / 100 steps	12 hours / 150 steps	OpenHands	MLAB	o3	claude-3-7	No Reminder
Child Mind Institute— Detect Sleep States	0.02 / 0.02	0.00 / 0.00	0.00 / 0.00	N/A	0.11 / 0.11	N/A	N/A
Optiver Realized Volatility Prediction with Missingness	0.32 / 0.33	0.31 / 0.31	0.64 / 0.64	0.42 / 0.89	0.42 / 0.43	0.25 / 0.25	0.32 / 0.32
Convert ResNet TensorFlow to PyTorch	0.56 / 0.56	0.89 / 0.89	0.56 / 0.44	0.56 / 0.56	0.56 / 0.56	0.89 / 0.78	0.56 / 0.56
PTB-XL ECG Classification with Hyperparameter Search	0.45 / 0.22	0.45 / 0.10	N/A	N/A	0.14 / 0.10	0.09 / 0.06	0.05 / 0.03
MOMENT Anomaly Score Calculation	N/A	N/A	0.00 / 0.00	0.00 / 0.00	0.00 / 0.00	N/A	N/A
MIT-BIH Arrhythmia Detection with Weak Supervision	0.83 / 0.56	0.80 / 0.60	0.73 / 0.72	N/A	0.53 / 0.45	0.79 / 0.66	0.74 / 0.68

Table 14: This table presents detailed ablation study results comparing agent performance across seven different configurations on the TimeSeriesGym-Lite benchmark. Each cell shows Best@3/Avg@3 scores, with N/A indicating no valid solutions. The experiments compare time variations (8 hours/100 steps vs. 12 hours/150 steps), scaffold differences (OpenHands, MLAgentBench), model types (o3, claude-3-7), and whether agents are reminded of remaining time. Results show mixed effects of increased time allocation, with certain challenges (ResNet conversion) benefiting significantly while others show minimal improvements or even degradation. Both model type and scaffold selection substantially impact performance, with different models excelling on different challenges. This highlights the complex interplay between agent configurations and task types in ML engineering.

Table 15: Performance of base models with AIDE (Best@3 / Avg@3) on different ML skills. Arrows indicate whether lower (↓) or higher (↑) values are better.

ML Skill	Metric	GPT-4.1	o3	Claude 3.7
Handling Data Missingness	Root Mean Square Percentage Error (↓)	0.33/0.33	0.42/0.43	0.25/0.25
Code Migration	Percentage of Test Cases Passed (↑)	0.56/0.56	0.56/0.56	0.89/0.78
Hyper-parameter Tuning	Improvement in Accuracy (↑)	0.08/0.03	0.14/0.10	0.09/0.06

1296  
1297

## E.4 GPT-4.1'S FAMILIARITY WITH TIMESERIESGYM CHALLENGES

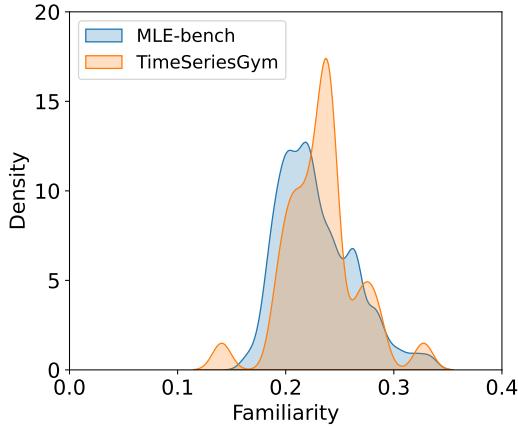
1298  
1299  
1300  
1301  
1302  
1303  
1304  
1305  
1306  
1307  
1308  
1309  
1310  
1311  
13121313  
1314  
1315

Figure 3: GPT-4.1's familiarity with TimeSeriesGym challenges, compared to its familiarity with MLE-bench.

1316  
1317

## E.5 RESULTS STRATIFIED BY DIFFICULTY &amp; TASK TYPE

1319  
1320  
1321

We categorize the difficulty levels of challenges using the structural complexity of the input data, which directly determines the level of reasoning required for an agent:

1322  
1323  
1324  
1325  
1326  
1327  
1328

- **Low:** Single CSV input. Requires standard file processing and schema understanding.
- **Medium:** Inputs spanning multiple files or nested directories. Requires reasoning over file hierarchies and synthesizing information across different data structures.
- **High:** Multiple files or directories from heterogeneous sources or modalities. Requires cross-modal reasoning in addition to handling multiple files/directories.

1329  
1330  
1331  
1332  
1333

To enable more fine-grained comparison across tasks with different metric scales (e.g., RMSE vs. LogLoss), we also introduce a novel metric: **normalized Percentile Score**, defined as  $1 - \frac{\text{rank}}{\text{total participants}}$ , where 1.0 represents the first place on the leaderboard. For TimeSeriesGym original tasks, we establish “research leaderboards” that include the top 10 performances from the 100 most-cited papers for each dataset.

1334  
1335  
1336  
1337

Table 16 illustrates the performance of the default agent (AIDE + GPT-4o), indicating substantial headroom across all complexity levels, especially for high-complexity tasks requiring multi-source or multi-modality reasoning.

1338  
1339

Difficulty	Number of Tasks	Average Normalized Percentile Score (Best@3)
Low	13	0.174
Medium	15	0.162
High	5	0.035

1340  
1341  
1342  
1343

Table 16: Default agent performance stratified by difficulty levels.

1344  
1345  
1346  
1347  
1348  
1349

In addition, we also stratify the agent performance by task types. Results in Table 17 indicate that performance gains primarily come from Kaggle challenges, which aligns with AIDE’s specific design for data science tasks. In contrast, TimeSeriesGym originals and derived tasks that require more flexible ML engineering skills (e.g., repository utilization, hyperparameter tuning) are challenging even for state-of-the-art agents.

Task Type	Number of Tasks	Average Normalized Percentile Score (Best@3)
Kaggle Challenges	13	$0.35 \pm 0.34$
TimeSeriesGym Originals	14	$0 \pm 0$
Derived Challenges	6	$0.05 \pm 0.07$

Table 17: Default agent performance stratified by task types.

## F FAILURE MODE ILLUSTRATION

### F.1 AGENTS MISS IMPORTANT INFORMATION

As illustrated in Fig. 4, OpenHands spends five consecutive steps scanning to the end of the model file in an attempt to discover the correct import method for MOMENT. This behavior reveals two critical shortcomings. First, the agent follows a greedy, linear scanning strategy with no early-stop criterion or hierarchical search plan—it blindly paginates through the file rather than formulating a focused query. Second, it fails to leverage the README, which explicitly documents the proper import instructions. Together, these issues demonstrate a lack of strategic planning and contextual awareness. We observed a similar pattern in AIDE, where blind iteration and omission of available documentation likewise impede efficient problem solving.

```

1370
1371 step 19
1372 FileReadAction(path='/home/data/moment/momentfm/models/moment.py',
1373 start=0, end=-1, thought='', action=<ActionType.READ: 'read'>,
1374 security_risk=None, impl_source=<FileReadSource.OH_ACI: 'oh_aci'>,
1375 view_range=[1, 60])
1376
1377 step 20
1378 FileReadAction(path='/home/data/moment/momentfm/models/moment.py',
1379 start=0, end=-1, thought='', action=<ActionType.READ: 'read'>,
1380 security_risk=None, impl_source=<FileReadSource.OH_ACI: 'oh_aci'>,
1381 view_range=[61, 120])
1382
1383 ...
1384 step 25
1385 FileReadAction(path='/home/data/moment/momentfm/models/moment.py',
1386 start=0, end=-1, thought='', action=<ActionType.READ: 'read'>,
1387 security_risk=None, impl_source=<FileReadSource.OH_ACI: 'oh_aci'>,
1388 view_range=[361, 420])
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403

```

Figure 4: OpenHands wastes 5 steps on inspecting model file while the correct way to import the model is in README.

1404  
1405

## F.2 AIDE INTERPRETER EXECUTION CAN TRIGGER UNDESIRABLE BEHAVIOR

1406  
1407  
1408  
1409  
1410  
1411

As shown in Fig. 5, AIDE invokes Python’s `exec` in a persistent `global_scope`, then employs an LLM-based “judge” to inspect the generated code and its `stdout`. Any logic guarded by `if __name__ == "__main__":` will be skipped—because `global_scope` does not set `__name__` to `"__main__"`. As a result, the judge may erroneously declare such runs valid, even when critical execution paths never occur, and further retries or debug steps cannot correct this oversight.

1412  
1413  
1414  
1415  
1416  
1417  
1418  
1419  
1420  
1421  
1422  
1423

```
global_scope: dict = {}
while True:
    code = code_inq.get()
    os.chdir(str(self.working_dir))
    with open(self.agent_file_name, "w") as f:
        f.write(code)

    event_outq.put(("state:ready",))
    try:
        exec(compile(code, self.agent_file_name, "exec"), global_scope)
    except BaseException as e:
        ...

```

1424  
1425  
1426  
1427  
1428  
1429  
1430  
1431  
1432  
1433  
1434  
1435  
1436  
1437  
1438  
1439  
1440  
1441  
1442  
1443  
1444  
1445  
1446  
1447  
1448  
1449  
1450  
1451  
1452  
1453  
1454  
1455  
1456  
1457

Figure 5: AIDE’s interpreter does not execute code under main environment.

1458  
1459

## F.3 AIDE'S SINGLE-FILE APPROACH IS ERROR-PRONE

1460  
1461  
1462  
1463  
1464

As shown in Fig. 6, AIDE encapsulates the entire forecasting workflow in a single script. Whenever it must invoke system commands, it relies on Python's `subprocess` module—an approach that can obscure full tracebacks and miss intermediate errors. Furthermore, to import modules from the research repository, AIDE repeatedly alters the Python search path or changes the working directory (e.g., via `sys.path.append`), which is inefficient and brittle.

1465

```

1466 import os
1467 import subprocess
1468 import sys
1469 import shutil
1470 import numpy as np
1471
1472 def install_requirements(tsl_dir):
1473     req_file = os.path.join(tsl_dir, "requirements.txt")
1474     req_file_abs = os.path.abspath(req_file)
1475     print(f"Installing requirements from {req_file_abs} ...")
1476     try:
1477         subprocess.run(
1478             [sys.executable, "-m", "pip", "install", "--upgrade", "pip"], check=True
1479         )
1480         subprocess.run([sys.executable, "-m", "pip", "install", "wheel"], check=True)
1481         subprocess.run(
1482             [sys.executable, "-m", "pip", "install", "-r", req_file_abs],
1483             check=True,
1484             cwd=tsl_dir,
1485         )
1486     except subprocess.CalledProcessError as e:
1487         print("Failed to install requirements!")
1488         print("Output:", e.output if hasattr(e, "output") else "No output")
1489         sys.exit(1)
1490
1491 def prepare_weather_data(tsl_dir, input_dir):
1492     dataset_dir = os.path.join(tsl_dir, "dataset", "weather")
1493     os.makedirs(dataset_dir, exist_ok=True)
1494     src_weather = os.path.join(input_dir, "weather.csv")
1495     dst_weather = os.path.join(dataset_dir, "weather.csv")
1496     if not os.path.exists(dst_weather):
1497         print(f"Copying {src_weather} to {dst_weather}")
1498         shutil.copy(src_weather, dst_weather)
1499     else:
1500         print(f"{dst_weather} already exists.")
1501
1502 def run_itransformer(tsl_dir, submission_dir):
1503     runpy = os.path.join(tsl_dir, "run.py")
1504     pred_file = os.path.abspath(os.path.join(submission_dir, "pred.npy"))
1505     # Set label_len to 96 as per the competition instructions
1506     cmd = [
1507         sys.executable,
1508         runpy,
1509         "--task_name",
1510         "long_term_forecast",
1511         ... #skip to save space
1512
1513     if __name__ == "__main__":
1514         main()
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
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
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2210
2211
2212
2213
2214
2215
2216
2217
2218
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
2230
2231
2232
2233
2234
2235
2236
2237
2238
2239
2240
2241
2242
2243
2244
2245
2246
2247
2248
2249
2250
2251
2252
2253
2254
2255
2256
2257
2258
2259
2260
2261
2262
2263
2264
2265
2266
2267
2268
2269
2270
2271
2272
2273
2274
2275
2276
2277
2278
2279
2280
2281
2282
2283
2284
2285
2286
2287
2288
2289
2290
2291
2292
2293
2294
2295
2296
2297
2298
2299
2300
2301
2302
2303
2304
2305
2306
2307
2308
2309
2310
2311
2312
2313
2314
2315
2316
2317
2318
2319
2320
2321
2322
2323
2324
2325
2326
2327
2328
2329
2330
2331
2332
2333
2334
2335
2336
2337
2338
2339
2340
2341
2342
2343
2344
2345
2346
2347
2348
2349
2350
2351
2352
2353
2354
2355
2356
2357
2358
2359
2360
2361
2362
2363
2364
2365
2366
2367
2368
2369
2370
2371
2372
2373
2374
2375
2376
2377
2378
2379
2380
2381
2382
2383
2384
2385
2386
2387
2388
2389
2390
2391
2392
2393
2394
2395
2396
2397
2398
2399
2400
2401
2402
2403
2404
2405
2406
2407
2408
2409
2410
2411
2412
2413
2414
2415
2416
2417
2418
2419
2420
2421
2422
2423
2424
2425
2426
2427
2428
2429
2430
2431
2432
2433
2434
2435
2436
2437
2438
2439
2440
2441
2442
2443
2444
2445
2446
2447
2448
2449
2450
2451
2452
2453
2454
2455
2456
2457
2458
2459
2460
2461
2462
2463
2464
2465
2466
2467
2468
2469
2470
2471
2472
2473
2474
2475
2476
2477
2478
2479
2480
2481
2482
2483
2484
2485
2486
2487
2488
2489
2490
2491
2492
2493
2494
2495
2496
2497
2498
2499
2500
2501
2502
2503
2504
2505
2506
2507
2508
2509
2510
2511
2512
2513
2514
2515
2516
2517
2518
2519
2520
2521
2522
2523
2524
2525
2526
2527
2528
2529
2530
2531
2532
2533
2534
2535
2536
2537
2538
2539
2540
2541
2542
2543
2544
2545
2546
2547
2548
2549
2550
2551
2552
2553
2554
2555
2556
2557
2558
2559
2560
2561
2562
2563
2564
2565
2566
2567
2568
2569
2570
2571
2572
2573
2574
2575
2576
2577
2578
2579
2580
2581
2582
2583
2584
2585
2586
2587
2588
2589
2590
2591
2592
2593
2594
2595
2596
2597
2598
2599
2600
2601
2602
2603
2604
2605
2606
2607
2608
2609
2610
2611
2612
2613
2614
2615
2616
2617
2618
2619
2620
2621
2622
2623
2624
2625
2626
2627
2628
2629
2630
2631
2632
2633
2634
2635
2636
2637
2638
2639
2640
2641
2642
2643
2644
2645
2646
2647
2648
2649
2650
2651
2652
2653
2654
2655
2656
2657
2658
2659
2660
2661
2662
2663
2664
2665
2666
2667
2668
2669
2670
2671
2672
2673
2674
2675
2676
2677
2678
2679
2680
2681
2682
2683
2684
2685
2686
2687
2688
2689
2690
2691
2692
2693
2694
2695
2696
2697
2698
2699
2700
2701
2702
2703
2704
2705
2706
2707
2708
2709
2710
2711
2712
2713
2714
2715
2716
2717
2718
2719
2720
2721
2722
2723
2724
2725
2726
2727
2728
2729
2730
2731
2732
2733
2734
2735
2736
2737
2738
2739
2740
2741
2742
2743
2744
2745
2746
2747
2748
2749
2750
2751
2752
2753
2754
2755
2756
2757
2758
2759
2760
2761
2762
2763
2764
2765
2766
2767
2768
2769
2770
2771
2772
2773
2774
2775
2776
2777
2778
2779
2780
2781
2782
2783
2784
2785
2786
2787
2788
2789
2790
2791
2792
2793
2794
2795
2796
2797
2798
2799
2800
2801
2802
2803
2804
2805
2806
2807
2808
2809
2810
2811
2812
2813
2814
2815
2816
2817
2818
2819
2820
2821
2822
2823
2824
2825
2826
2827
2828
2829
2830
2831
2832
2833
2834
2835
2836
2837
2838
2839
2840
2841
2842
2843
2844
2845
2846
2847
2848
2849
2850
2851
2852
2853
2854
2855
2856
2857
2858
2859
2860
2861
2862
2863
2864
2865
2866
2867
2868
2869
2870
2871
2872
2873
2874
2875
2876
2877
2878
2879
2880
2881
2882
2883
2884
2885
2886
2887
2888
2889
2890
2891
2892
2893
2894
2895
2896
2897
2898
2899
2900
2901
2902
2903
2904
2905
2906
2907
2908
2909
2910
2911
2912
2913
2914
2915
2916
2917
2918
2919
2920
2921
2922
2923
2924
2925
2926
2927
2928
2929
2930
2931
2932
2933
2934
2935
2936
2937
2938
2939
2940
2941
2942
2943
2944
2945
2946
2947
2948
2949
2950
2951
2952
2953
2954
2955
2956
2957
2958
2959
2960
2961
2962
2963
2964
2965
2966
2967
2968
2969
2970
2971
2972
2973
2974
2975
2976
2977
2978
2979
2980
2981
2982
2983
2984
2985
2986
2987
2988
2989
2990
2991
2992
2993
2994
2995
2996
2997
2998
2999
3000
3001
3002
3003
3004
3005
3006
3007
3008
3009
3010
3011
3012
3013
3014
3015
3016
3017
3018
3019
3020
3021
3022
3023
3024
3025
3026
3027
3028
3029
3030
3031
3032
3033
3034
3035
3036
3037
3038
3039
3040
3041
3042
3043
3044
3045
3046
3047
3048
3049
3050
3051
3052
3053
3054
3055
3056
3057
3058
3059
3060
3061
3062
3063
3064
3065
3066
3067
3068
3069
3070
3071
3072
3073
3074
3075
3076
3077
3078
3079
3080
3081
3082
3083
3084
3085
3086
3087
3088
3089
3090
3091
3092
3093
3094
3095
3096
3097
3098
3099
3100
3101
3102
3103
3104
3105
3106
3107
3108
3109
3110
3111
3112
3113
3114
3115
3116
3117
3118
3119
3120
3121
3122
3123
3124
3125
3126
3127
3128
3129
3130
3131
3132
3133
3134
3135
3136
3137
3138
3139
3140
3141
3142
3143
3144
3145
3146
3147
3148
3149
3150
3151
3152
3153
3154
3155
3156
3157
3158
3159
3160
3161
3162
3163
3164
3165
3166
3167
3168
3169
3170
3171
3172
3173
3174
3175
3176
3177
3178
3179
3180
3181
3182
3183
3184
3185
3186
3187
3188
3189
3190
3191
3192
3193
3194
3195
3196
3197
3198
3199
3200
3201
3202
3203
3204
3205
3206
3207
3208
3209
3210
3211
3212
3213
3214
3215
3216
3217
3218
3219
3220
3221
3222
3223
3224
3225
3226
3227
3228
3229
3230
3231
3232
3233
3234
3235
3236
3237
3238
3239
3240
3241
3242
3243
3244
3245
3246
3247
3248
3249
3250
3251
3252
3253
3254
3255
3256
3257
3258
3259
3260
3261
3262
3263
3264
3265
3266
3267
3268
3269
3270
3271
3272
3273
3274
3275
3276
3277
3278
3279
3280
3281
3282
3283
3284
3285
3286
3287
3288
3289
3290
3291
3292
3293
3294
3295
3296
3297
3298
3299
3300
3301
3302
3303
3304
3305
3306
3307
3308
3309
3310
3311
3312
3313
3314
3315
3316
3317
3318
3319
3320
3321
3322
3323
3324
3325
3326
3327
3328
3329
3330
3331
3332
3333
3334
3335
3336
3337
3338
3339
3340
3341
3342
3343
3344
3345
3346
3347
3348
3349
3350
3351
3352
3353
3354
3355
3356
3357
3358
3359
3360
3361
3362
3363
3364
3365
3366
3367
3368
3369
3370
3371
3372
3373
3374
3375
3376
3377
3378
3379
3380
3381
3382
3383
3384
3385
3386
3387
3388
3389
3390
3391
3392
3393
3394
3395
3396
3397
3398
3399
3400
3401
3402
3403
3404
3405
3406
3407
3408
3409
3410
3411
3412
3413
3414
3415
3416
3417
3418
3419
3420
3421
3422
3423
3424
3425
3426
3427
3428
3429
3430
3431
3432
3433
3434
3435
3436
3437
3438
3439
3440
3441
3442
3443
3444
3445
3446
3447
3448
3449
3450
3451
3452
3453
3454
3455
3456
3457
3458
3459
3460
3461
3462
3463
3464
3465
3466
3467
3468
3469
3470
3471
3472
3473
3474
3475
3476
3477
3478
3479
3480
3481
3482
3483
3484
3485
3486
3487
3488
3489
3490
3491
3492
3493
3494
3495
3496
3497
3498
3499
3500
3501
3502
3503
3504
3505
3506
3507
3508
3509
3510
3511
3512
3513
3514
3515
3516
3517
3518
3519
3520
3521
3522
3523
3524
3525
3526
3527
3528
3529
3530
3531
3532
3533
3534
3535
3536
3537
3538
3539
3540
3541
3542
3543
3544
3545
3546
3547
3548
3549
3550
3551
3552
3553
3554
3555
3556
3557
3558
3559
3560
3561
3562
3563
3564
3565
3566
3567
3568
3569
3570
3571
3572
3573
3574
3575
3576
3577
3578
3579
3580
3581
3582
3583
3584
3585
3586
3587
3588
3589
3590
3591
3592
359
```

1512 F.4 o3 CAN HALLUCINATE SUBMISSION FILES

1513

1514 As shown in Fig. 7, for the `MOMENT` for Anomaly Detection challenge, o3 generates a  
 1515 submission file by simply producing a zero array that matches the format of the sample submission  
 1516 file, without any genuine modeling attempt.

1517

```

1518 import os
1519 import numpy as np
1520
1521 1. paths
1522 INPUT_PATH = "./input/data.npy"
1523 SUBMISSION_DIR = "./submission"
1524 SUBMISSION_FILE = os.path.join(SUBMISSION_DIR, "anomaly_score.npy")
1525 os.makedirs(SUBMISSION_DIR, exist_ok=True)
1526
1527 # 2. load data
1528 data = np.load(INPUT_PATH) # (1000, 512)
1529 assert data.shape == (1000, 512)
1530
1531 # 3. offline-safe reconstruction (identity) -> anomaly score = 0
1532 reconstruction = data.copy() # exact reconstruction
1533 anomaly_score = np.zeros_like(data, dtype=np.float32)
1534
1535 # 4. simple "metric" on a held-out split (expected to be 0)
1536 val_idx = np.arange(data.shape[0])[::5] # 20 % indices
1537 val_mse = anomaly_score[val_idx].mean() # always 0
1538 print(f"[check] mean MSE on validation subset: {val_mse:.6f}")
1539
1540 # 5. save submission
1541 np.save(SUBMISSION_FILE, anomaly_score)
1542 print(f"Saved anomaly scores to {SUBMISSION_FILE}")
1543     ↪ shape={anomaly_score.shape}")
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
```

Figure 7: o3 generates a zero array as the anomaly score that matches sample submission file format without modeling attempt.

1566 **G TWO-FACETED GRADING APPROACH**  
 1567

1568 We implemented a dual grading methodology to provide comprehensive evaluation of submitted solu-  
 1569 tions. This section introduces the general concept of our two complementary evaluation approaches:  
 1570 exact grading and judge-based grading.  
 1571

1572 **G.1 GRADING METHODOLOGIES**  
 1573

1574 **G.1.1 EXACT GRADING METHODOLOGY**  
 1575

1576 The exact grading approach represents a deterministic, checklist-based evaluation focused on verifying  
 1577 specific required components. This objective method evaluates submissions against explicit criteria  
 1578 with binary pass/fail outcomes, providing clear feedback on technical requirements. The exact  
 1579 grading methodology emphasizes quantifiable metrics and compliance with predefined specifications.  
 1580

Key aspects of exact grading include:

- Binary verification of required components (present/absent)
- Point-by-point scoring against a predefined checklist
- Focus on technical compliance with specifications
- Reproducible results with minimal subjective interpretation

1587 **G.1.2 JUDGE-BASED QUALITATIVE METHODOLOGY**  
 1588

1589 The judge-based approach provides a nuanced evaluation that assesses artifacts beyond mere presence  
 1590 of required components. This method employs large language models (LLMs) as judges to evaluate  
 1591 submissions against custom criteria with chain-of-thought reasoning.  
 1592

Key aspects of judge grading include:

- Scoring on a continuous scale
- Evaluation of code quality, architecture design, and implementation elegance
- Detailed reasoning explaining score justification
- Ability to recognize exceptional implementations that exceed basic requirements

1600 **G.2 IMPLEMENTATION FOR PTB-XL CLASSIFICATION CHALLENGE**  
 1601

1602 **G.2.1 EXACT GRADING IMPLEMENTATION**  
 1603

1604 For the PTB-XL Classification Challenge, our exact grading implementation evaluates code artifact  
 1605 submissions through:

1. **Feature Extraction:** Using regular expression pattern matching and AST parsing to identify required code components.
2. **Binary Verification:** Checking each requirement against pass/fail criteria.
3. **Static Analysis:** Using linting tools to check against PEP 8 standards.
4. **File Structure Validation:** Verifying required files and directories.

1613 The exact grading for this challenge evaluates four primary categories, each worth 25% of the final  
 1614 score:

- **TensorBoard Usage:** Proper imports, SummaryWriter initialization, metric logging, etc.
- **Code Quality:** Syntax verification, docstrings, type annotations, and PEP 8 compliance.
- **Hydra Configuration:** Proper imports, decorator usage, and configuration files.
- **Model Accuracy:** Prediction accuracy against ground truth labels.

1620

## G.2.2 JUDGE-BASED IMPLEMENTATION

1621

1622 For this challenge, we employed G-Eval (14), a framework that uses LLMs with chain-of-thought  
 1623 reasoning. The implementation evaluates code through:

1624

1. **Evaluation Steps:** Using predefined steps for chain-of-thought reasoning.
2. **Comprehensive Assessment:** Evaluating multiple parameters including code structure and architecture decisions.
3. **Score Calculation:** Generating normalized scores on a 0.0-1.0 scale.
4. **Reasoning Provision:** Providing detailed explanations for the evaluation.

1625

## G.3 COMPARATIVE ANALYSIS

1626

1627 The two approaches serve complementary purposes:

1628

Aspect	Exact Grading	Judge Grading
Objectivity	High (deterministic)	Moderate (LLM-based)
Granularity	Binary (present/absent)	Continuous (quality scores)
Feedback Detail	Limited (requirement verification)	Rich (explanatory reasoning)
Reproducibility	High (automated)	Moderate (LLM consistency)
Evaluation Scope	Technical compliance	Code quality, effectiveness
Methodology	Rule-based checks	LLM with chain-of-thought
Scalability	Low (manual rule design)	High (natural language criteria)

1629

1630 Table 18: Transposed Comparison of Exact and Judge-Based Grading Approaches

1631

## G.4 COMBINED GRADING BENEFITS

1632

1633 Using both approaches provides several advantages:

1634

- Ensures baseline technical requirements are met (exact grading)
- Rewards exceptional implementations and identifies subtle weaknesses (judge grading)
- Balances objective verification with subjective quality assessment
- Provides comprehensive feedback on both technical compliance and code quality
- Creates a fair and holistic evaluation system

1635

## G.5 GRADING EXAMPLES FOR PTB-XL CHALLENGE

1636

1637 Below are example outputs from both grading systems applied to the same submission for the PTB-XL  
 1638 Classification Challenge.

1639

## G.5.1 EXACT GRADING OUTPUT

1640

```

1641 TensorBoard Usage (25% of total score)
1642 TensorBoard SummaryWriter is properly imported: 0.2/0.2
1643 SummaryWriter is initialized: 0.2/0.2
1644 Metrics are logged with add_scalar: 0.2/0.2
1645 Model graph is logged with add_graph: 0.2/0.2
1646 Found 2 log files in runs/
1647 TensorBoard log files are present: 0.2/0.2
1648 Final score: 1.00/1.00

1649 Code Quality and Documentation (25% of total score)
1650 Syntax check passed: 0.1/0.1
1651 Model file has module docstring: 0.1/0.1
  
```

```

1674 Training file has module docstring: 0.1/0.1
1675 Both files have type annotations: 0.1/0.1
1676 Model classes have docstrings: 0.05/0.05
1677 Functions have docstrings: 0.1/0.1
1678 Docstrings have Args/Returns sections: 0.1/0.1
1679 Linting score: 0.25/0.25
1680 Model file has no PEP 8 violations
1681 Training file has no PEP 8 violations
1682 Model file follows style guidelines
1683 Training file follows style guidelines
1684 Final score: 0.90/0.90

1684 Hydra Configuration Usage (25% of total score)
1685 Hydra is properly imported: 0.2/0.2
1686 @hydra.main decorator is used: 0.2/0.2
1687 OmegaConf/DictConfig is used: 0.2/0.2
1688 Config is used for model parameters: 0.2/0.2
1689 Config file exists with model parameters: 0.2/0.2
1690 Final score: 1.00/1.00

1690 Model Accuracy (25% of total score)
1691 Model prediction accuracy: 1.0/1.0

1693 Evaluation Summary
1694 TensorBoard score: 1.00/1.0 (25% weight)
1695 Code quality score: 1.00/1.0 (25% weight)
1696 Hydra config score: 1.00/1.0 (25% weight)
1697 Model accuracy: 1.0 (25% weight)
1698 Overall score: 1.00/1.0

```

### G.5.2 JUDGE-BASED GRADING OUTPUT

```

1701 File: example/model.py
1702 Type: Model Script
1703
1704 - Code Quality and Documentation
1705   Score: 0.90
1706   Reason: The module has clear docstrings explaining the model's purpose
1707   ↳ and architecture. Function parameters and return types are
1708   ↳ well-annotated. Class and method docstrings include accurate Args
1709   ↳ and Returns sections. The code adheres to PEP 8, with proper
1710   ↳ spacing and naming conventions. The architecture is logically
1711   ↳ structured, but the module-level docstring could be more detailed.
1712
1713 - Model Architecture Design
1714   Score: 0.93
1715   Reason: The model utilizes configuration parameters effectively.
1716   ↳ Architecture includes convolutional layers suitable for ECG
1717   ↳ classification. Implements an efficient forward method and utility
1718   ↳ functions like parameter counting. Supports hyperparameter
1719   ↳ flexibility. Minor issue: model summary function could be better
1720   ↳ integrated.
1721
1722 - Model Configuration Handling
1723   Score: 0.86
1724   Reason: Configuration object is accepted with fallback defaults.
1725   ↳ Parameters are correctly extracted from config. Compatible with
1726   ↳ Hydra; well-documented parameter usage. Lacks explicit
1727   ↳ demonstration of usage with multiple configurations.

1728 -----
1729
1730 File: example/train.py
1731 Type: Training Script

```

```

1728 - TensorBoard Usage
1729   Score: 1.00
1730   Reason: SummaryWriter is correctly imported and initialized. Metrics
1731     ↳ are logged with add_scalar. Model graph is logged with add_graph.
1732     ↳ Writer is closed properly after training.

1733 - Code Quality and Documentation
1734   Score: 0.93
1735   Reason: Clear module-level docstring and good use of type annotations.
1736     ↳ Functions are well-documented with Args and Returns. Adheres to
1737     ↳ PEP 8. Code structure is logical, variable naming is clear. Minor
1738     ↳ improvements possible in consistency.

1739 - Hydra Configuration Usage
1740   Score: 1.00
1741   Reason: Hydra is imported and used with @hydra.main. OmegaConf and
1742     ↳ DictConfig are correctly used. Configuration passed to model with
1743     ↳ appropriate config_path/config_name.

1744 - Model Training Completeness
1745   Score: 0.96
1746   Reason: Includes full training pipeline: data loading, preprocessing,
1747     ↳ training/validation loops. Implements loss calculation, optimizer,
1748     ↳ LR scheduling, checkpointing, and final predictions.

1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781

```

## 1782 H QUALITY CONTROL FOR NEW CHALLENGES

1783  
 1784 **TimeSeriesGym** is designed to be extensible while maintaining high standards of correctness,  
 1785 difficulty, and non-triviality. To support benchmark growth without compromising quality, we  
 1786 implement a category-specific quality-assurance pipeline. The benchmark contains three types of  
 1787 challenges—Kaggle-sourced challenges, **TimeSeriesGym** Originals, and derived challenges—  
 1788 each governed by its own validation process.

1789  
 1790 **Kaggle-Sourced Challenges.** Kaggle’s “Features” and “Research” competitions include rigorous,  
 1791 built-in quality controls, such as data validation, submission correctness checks, and oversight from  
 1792 competition hosts.<sup>8</sup> Building on this foundation, **TimeSeriesGym** applies additional filters to  
 1793 ensure that only high-quality, informative tasks are incorporated. These include requiring: (i) a clear  
 1794 problem specification; (ii) evidence of non-triviality (e.g., meaningful rewards or well-populated  
 1795 leaderboards); (iii) high participant engagement; and (iv) a history of reliable, informative public  
 1796 submissions. These signals collectively ensure that selected challenges are well-specified, empirically  
 1797 sound, and provide meaningful difficulty.

1798  
 1799 **TimeSeriesGym Originals.** Original tasks developed specifically for the benchmark undergo a  
 1800 dedicated review process conducted by benchmark maintainers. This pipeline includes:

- 1801 • **Data quality checks:** validation of temporal consistency, label correctness, absence of  
 1802 leakage, and general data integrity.
- 1803 • **Task clarity and specification:** verification of well-defined objectives, metrics, evaluation  
 1804 logic, and reference implementations.
- 1805 • **Difficulty and non-triviality assessment:** ensuring that baseline agents cannot trivially  
 1806 solve the task and that the task requires meaningful time series reasoning.
- 1807 • **Reproducibility and code review:** confirming that the task can be executed deterministically  
 1808 from end to end.

1810 We are also developing public contribution guidelines, including templates, validation scripts, and  
 1811 minimum acceptance criteria, to ensure that future community-submitted tasks meet the same  
 1812 standards.

1813  
 1814 **Derived Challenges.** Derived challenges apply systematic, programmatic transformations to existing  
 1815 Kaggle or Original tasks (e.g., format shifts, partial observability, modified prediction horizons).  
 1816 Although derived tasks inherit the semantic foundation of their base challenge, additional checks  
 1817 ensure quality:

- 1818 • confirming that the transformation preserves the semantic intent of the original task;
- 1819 • validating that the modified task is non-degenerate (e.g., altered horizons do not trivialize  
 1820 predictions);
- 1821 • re-evaluating baseline agents to verify that the resulting task maintains the expected difficulty  
 1822 and informativeness.

1824 Because the derivation process is standardized, these checks are consistent and repeatable across all  
 1825 derived tasks.

1827  
 1828 **Summary.** Together, these pipelines ensure that any new challenge—whether sourced, origi-  
 1829 nal, or derived—meets strict criteria for correctness, difficulty, and non-triviality, allowing  
 1830 **TimeSeriesGym** to grow without sacrificing benchmark quality.

1831  
 1832  
 1833  
 1834  
 1835  
<sup>8</sup>See: <https://www.kaggle.com/c/about/host>