

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 ZEPHYRUS: AN AGENTIC FRAMEWORK FOR WEATHER SCIENCE

Anonymous authors

Paper under double-blind review

ABSTRACT

Foundation models for weather science are pre-trained on vast amounts of structured numerical data and outperform traditional weather forecasting systems. However, these models lack language-based reasoning capabilities, limiting their utility in interactive scientific workflows. Large language models (LLMs) excel at understanding and generating text but cannot reason about high-dimensional meteorological datasets. We bridge this gap by building a novel agentic framework for weather science. Our framework includes a Python code-based environment for agents (*ZEPHYRUSWORLD*) to interact with weather data, featuring tools like an interface to WeatherBench 2 dataset, geoquerying for geographical masks from natural language, weather forecasting, and climate simulation capabilities. We design *ZEPHYRUS*, a multi-turn LLM-based weather agent that iteratively analyzes weather datasets, observes results, and refines its approach through conversational feedback loops. We accompany the agent with a new benchmark, *ZEPHYRUSBENCH*, with a scalable data generation pipeline that constructs diverse question-answer pairs across weather-related tasks, from basic lookups to advanced forecasting, extreme event detection, and counterfactual reasoning. Experiments on this benchmark demonstrate the strong performance of *ZEPHYRUS* agents over text-only baselines, outperforming them by up to 35 percentage points in correctness. However, on harder tasks, *ZEPHYRUS* performs similarly to text-only baselines, highlighting the challenging nature of our benchmark and suggesting promising directions for future work.

1 INTRODUCTION

Large language models (LLMs) have demonstrated remarkable capabilities across diverse scientific domains (Birhane et al., 2023), revolutionizing fields from drug discovery (Zheng et al., 2024; Wu et al., 2024b) and materials science (Lei et al., 2024; Jablonka et al., 2023) to network biology (Theodoris et al., 2023). These models excel at processing textual content such as scientific literature, source code (Jiang et al., 2024), and structured data tables (Zhang et al., 2024). However, their application to domains requiring reasoning over high-dimensional numerical data remains limited (Wang et al., 2024).

Meteorology offers a compelling yet challenging case study, as combining natural language reasoning with complex atmospheric data has the potential to greatly advance weather research. Weather prediction is a critical scientific challenge, with profound implications spanning agriculture, disaster preparedness, transportation, and energy management (Alley et al., 2019). The field has witnessed remarkable progress through machine learning approaches, with foundation models (Nguyen et al., 2023; Kurth et al., 2023; Lam et al., 2023; Bi et al., 2023; Nguyen et al., 2024) now achieving state-of-the-art performance in medium-range forecasting, often surpassing traditional physics-based numerical simulations (Molteni et al., 1996; Bauer et al., 2015). However, current weather models operate exclusively on structured numerical datasets such as reanalysis data, cannot incorporate valuable alternative modalities like textual weather bulletins or field station reports, and crucially, lack interactive natural language interfaces for querying or reasoning.

Weather science workflows require substantial technical expertise to orchestrate complex ecosystems of tools, datasets, and models. Researchers must navigate disparate data sources, integrate outputs from multiple forecasting systems, combine observational datasets with model predictions, and




Figure 1: **Overview:** We develop ZEPHYRUS, an agentic framework for weather science. Given a query, the LLM-based agent ZEPHYRUS writes a code block which is sent to the code execution server. The server orchestrates several tools to execute the code block and returns the execution results to the agent. The agent either decides to execute more code to refine its output or respond back to the user. Refer to Appendix A.8 for the full prompt.

coordinate between different computational environments and APIs. This dependency on extensive technical knowledge creates barriers for domain experts, limiting broader participation in weather science. Traditional meteorological workflows therefore require expert interpretation to translate computational outputs into actionable insights, increasing costs and limiting their utility in human-in-the-loop decision-support systems.

Multimodal LLMs can handle data from diverse modalities and offer a potential pathway to address these challenges. Models capable of jointly processing text with images (Wang et al., 2022; Alayrac et al., 2022; Li et al., 2022; Liu et al., 2023c), video (Zhao et al., 2022; Zhang et al., 2023; Cheng et al., 2024; Lin et al., 2024; Zhang et al., 2025), and audio (Chu et al., 2023; Défossez et al., 2024; Wu et al., 2024a; 2025; Doh et al., 2025; Ghosh et al., 2025) have shown impressive cross-modal reasoning abilities. Yet atmospheric data poses unique challenges: its spatiotemporal, multi-channel structure is fundamentally different from conventional modalities, requiring specialized approaches for effective integration with language models. Initial attempts to bridge this gap have shown promise but remain limited in scope. Early vision-language approaches to meteorology (Chen et al., 2024a; Li et al., 2024; Ma et al., 2024) have focused on narrow applications like extreme weather prediction using restricted variable subsets, falling short of general-purpose meteorological reasoning. More recent multimodal weather-language models (Varambally et al., 2025) demonstrate the potential of this direction but still fail to match established baselines across many important meteorological tasks. This persistent gap highlights a fundamental challenge: despite significant progress in both weather foundation models and LLMs, no existing system successfully unifies meteorological data with natural language reasoning for broad, interactive scientific applications.

We address this challenge by first introducing an agentic environment that enables LLMs to interact programmatically with meteorological data and models. We setup ZEPHYRUSWORLD, a comprehensive execution environment that exposes weather-focused capabilities through easy-to-use Python APIs. The system includes interfaces to the WeatherBench 2 dataset (Rasp et al., 2024), geo-query functionality for translating between coordinates and named locations, state-of-the-art forecasting models (Nguyen et al., 2024), and physics-based simulators. A FastAPI backend parallelizes code execution from LLM-generated queries.

We then develop two code-generating systems of increasing sophistication within this agentic framework. ZEPHYRUS-DIRECT generates Python code in a single step to solve weather problems directly (Gao et al., 2023). ZEPHYRUS-REFLECTIVE employs an iterative execution-refinement (Yao et al., 2023b): it executes code to manipulate weather data, analyzes the results, and refines both code and output before providing a final answer. Both approaches can automatically detect and correct errors produced during code execution. Figure 1 gives an overview of our entire agentic pipeline.

108 To systematically evaluate these approaches, we construct `ZEPHYRUSBENCH`, a comprehensive
 109 benchmark built on ERA5 reanalysis data (Hersbach et al., 2020) from WeatherBench 2 (Rasp
 110 et al., 2024). The benchmark combines human-authored and semi-synthetic tasks spanning 2158
 111 question-answer pairs across 46 distinct tasks. Tasks range from basic data lookups and forecasting
 112 to challenging research problems involving extreme event detection, forecast report generation, and
 113 prediction and counterfactual analysis. We also implement robust evaluation schemes to assess
 114 the scientific accuracy of all generated answers across diverse meteorological reasoning tasks. We
 115 summarize our key contributions below.

116 • We develop `ZEPHYRUSWORLD`, an agentic environment providing unified Python APIs for meteo-
 117 rological data, forecasting models, and climate simulation tools.
 118 • We introduce two code-generating systems that leverage `ZEPHYRUSWORLD`: `ZEPHYRUS-DIRECT`
 119 for single-step code generation and `ZEPHYRUS-REFLECTIVE` for iterative execution-refinement
 120 workflows to solve open-ended meteorological problems.
 121 • We curate `ZEPHYRUSBENCH`, a challenging weather reasoning benchmark with 2062 question-
 122 answer pairs across 46 meteorological task types.
 123 • Our evaluation shows that LLM agents achieve encouraging results on the benchmark, suggesting
 124 that they can be effective assistants to weather scientists.

2 RELATED WORK

125 **Weather Foundation Models.** Neural network-based weather forecasting systems (Lam et al., 2023;
 126 Price et al., 2025; Bi et al., 2023; Pathak et al., 2022; Nguyen et al., 2023; Bodnar et al., 2024; Nguyen
 127 et al., 2024) have revolutionized meteorological prediction by demonstrating superior performance
 128 compared to conventional physics-based approaches (Molteni et al., 1996) while being significantly
 129 more computationally efficient. Nevertheless, these architectures are predominantly trained for
 130 forecasting. In particular, they do not support conversational interfaces or cross-domain reasoning
 131 capabilities.

132 **Agentic frameworks for scientific discovery** Agentic frameworks implement the per-
 133 ceive–reason–plan–act loop by pairing LLMs with tools, memory, and feedback to pursue
 134 long-horizon goals. Core patterns include interleaving reasoning with tool calls (ReAct (Yao et al.,
 135 2023a)), self-critique with episodic memory (Reflexion Shinn et al. (2023)), and self-supervised
 136 learning of API use (Toolformer (Schick et al., 2023)). General-purpose libraries such as AutoGen
 137 provide a standard interface for multi-agent conversation and tool invocation, making these patterns
 138 reusable across tasks (Wu et al., 2024c).

139 In many scientific applications, these frameworks appear as domain agents and self-driving labs. In
 140 chemistry, ChemCrow couples an LLM controller with a curated set of expert tools for synthesis and
 141 analysis (Bran et al., 2024), while Coscientist integrates retrieval, code execution, and laboratory APIs
 142 to plan and run experiments end-to-end (Boiko et al., 2023). Biomedical agents extend the approach
 143 across literature, databases, and analysis workflows (e.g., Biomni (Huang et al., 2025)). Despite these
 144 advances across multiple scientific domains, weather science remains largely unexplored territory for
 145 agentic approaches.

146 **General-Purpose Vision-Language Models.** Multi-modal vision language models (Li et al., 2021;
 147 Alayrac et al., 2022; Li et al., 2022; 2023; Liu et al., 2023c;b;a; 2024) demonstrate strong visual
 148 reasoning capabilities on general-purpose evaluation benchmarks. However, adapting these models
 149 for applications in weather science presents considerable difficulties. Standard VLM architectures
 150 assume RGB visual inputs and exhibit weaknesses in quantitative analytical tasks. Meteorological
 151 data presents fundamentally different challenges through high-dimensional, structured atmospheric
 152 measurements requiring specialized integration approaches for language model compatibility. While
 153 weather-language hybrid models (Varambally et al., 2025) seem promising, they underperform
 154 relative to domain-specific baselines across critical meteorological applications.

155 **Multimodal Weather Datasets.** Recent research has developed several multimodal frameworks
 156 that combine weather observations with textual information. These include the Terra collection
 157 (Chen et al., 2024b), which integrates geographical imagery with descriptive text for general earth
 158 observation, and ClimateIQA (Chen et al., 2024a), which focuses on extreme weather detection
 159 through wind measurement analysis. Similarly, WeatherQA (Ma et al., 2024) specializes in severe
 160 weather interpretation using remote sensing data and expert commentary, while CLLMate (Li et al.,

162 2024) connects media reports with ERA5 observations for weather event classification. Despite
 163 these valuable contributions, existing frameworks are narrow in scope. They concentrate on narrow
 164 applications or utilize only small subsets of atmospheric variables. This approach overlooks a
 165 fundamental characteristic of atmospheric dynamics: weather systems involve complex multi-scale
 166 interactions across numerous meteorological parameters. To address these limitations, our benchmark
 167 incorporates diverse weather reasoning tasks, both human-implemented and semi-synthetically
 168 generated, that span across most WeatherBench2 data channels.

169 3 ZEPHYRUS: AN AGENTIC FRAMEWORK FOR WEATHER SCIENCE

171 3.1 ZEPHYRUSWORLD: AN AGENTIC ENVIRONMENT FOR WEATHER SCIENCE

173 The fragmented nature of weather science tools makes it challenging for LLMs to effectively leverage
 174 them for scientific tasks. To address this, we introduce ZEPHYRUSWORLD, a comprehensive agentic
 175 environment that unifies weather science capabilities from diverse tools through a clean Pythonic
 176 interface. Given a question, we leverage LLMs’ ability (Gao et al., 2023; Jimenez et al.) to generate
 177 Python code and execute it in a sandboxed environment. The output is then fed back to the model,
 178 along with any execution errors. We design high-level APIs for the tools for ease of use, and include
 179 documentation extracted from the docstrings in the models context at inference time.

180 The environment encompasses several essential weather science tools:

- 182 **WeatherBench 2 Data Indexer.** The environment provides the model access to the data through
 183 the `xarray` dataset interface.
- 184 **Geolocator.** This tool provides comprehensive geospatial functionality for weather data analysis.
 185 It handles forward geocoding (place names to coordinates) and reverse geocoding (coordinates
 186 to location names) using the Natural Earth dataset (Natural Earth, 2024). Key operations include
 187 finding geographic features at specific coordinates, retrieving boolean masks and area-weighted
 188 maps for regions, listing sublocations, and calculating geodistances. Built using `geopandas` and
 189 `shapely`, it maintains precomputed spatial caches for fast lookups.
- 190 **Forecaster.** We incorporate the Stormer model (Nguyen et al., 2024), a transformer-based neural
 191 weather prediction system trained on WeatherBench 2. We chose it for its strong performance at
 192 short to medium range forecasts while being orders of magnitude more efficient than traditional
 193 numerical models. Our implementation abstracts checkpoint loading and preprocessing, providing a
 194 simple interface to run forecasts from arbitrary atmospheric initial conditions and return outputs as
 195 `xarray` datasets.
- 196 **Simulator.** Our JAX-GCM simulator is an intermediate complexity atmospheric model built on
 197 NeuralGCM’s dynamical core (Kochkov et al., 2024). It incorporates physical parameterizations
 198 from the SPEEDY Fortran model (Molteni et al., 1996), including radiation, moist physics (clouds
 199 and convection), and vertical and horizontal diffusion. We use the default T32 configuration
 200 (approximately 3.5° resolution) with 8 vertical layers. Built on JAX, we can run 5-day simulations
 201 in only $\approx 25s$ on an A100 GPU.

202 **Code Execution Server.** ZEPHYRUSWORLD requires a system capable of handling multiple weather
 203 analysis tasks simultaneously without resource conflicts. We implement a FastAPI-based server-client
 204 architecture that processes multiple weather analysis requests in parallel using resource pools to
 205 prevent conflicts between simultaneous executions. Each execution follows a strict protocol of
 206 acquiring resources, loading datasets, injecting tools, and executing code with timeout protection
 207 before returning outputs and errors to the client. More details are presented in Appendix A.1

208 3.2 THE ZEPHYRUS FAMILY OF WEATHER AGENTS

209 We design agentic systems that leverage ZEPHYRUSWORLD to solve complex meteorological
 210 tasks. Our approach constructs prompts containing comprehensive documentation of ZEPHYRUS-
 211 WORLD tools, variable descriptions, units, and coordinate systems. The models generate Python
 212 functions using these tools to solve the given questions, which execute on ZEPHYRUSWORLD’s code
 213 execution server. Any execution errors or timeouts are returned to the models, which regenerate code
 214 until the error is resolved. We implement two distinct systems that differ in their execution strategy
 215 and refinement approach. Both systems intentionally maintain simple designs to isolate and measure
 the agentic capabilities of LLMs for solving weather science problems.

216 **ZEPHYRUS-DIRECT** generates a complete Python solution in one attempt and reports the execution
 217 output as the final answer. This model runs the error-correction loop for a maximum of 5 times.
 218

219 **ZEPHYRUS-REFLECTIVE** implements a multi-turn workflow that alternates between code gener-
 220 ation and execution phases. The agent executes individual code blocks and receives the output as
 221 observations. The execution results are fed back to the LLM, which analyzes the observations and
 222 decides on the next step. This iterative process enables the model to assess the scientific plausibility
 223 of outputs, identify anomalies or mistakes in results, and refine subsequent code blocks to address
 224 logical errors. We run the interaction loop for a maximum of 20 times per question.
 225

226 The complete prompts for both systems are presented in Appendix A.8.
 227

228 4 ZEPHYRUSBENCH: A COMPREHENSIVE WEATHER BENCHMARK

229 Weather science problems require analyzing complex atmospheric patterns, modeling trends, and
 230 combining data from multiple sources. We introduce ZEPHYRUSBENCH, a comprehensive bench-
 231 mark that evaluates how effectively LLMs can assist in real-world meteorological workflows. The
 232 benchmark comprises 46 distinct meteorological tasks with answers derived from curated weather
 233 reports and human-generated or verified code.
 234

235 4.1 DATASET CURATION

236 We base our tasks around the ERA5 reanalysis dataset (Hersbach et al., 2020), specifically from
 237 WeatherBench 2 (Rasp et al., 2024). The dataset provides global atmospheric data from 1979 to 2022.
 238 We use 1.5° spatial resolution with 6-hourly temporal resolution.
 239

240 The capabilities measured by our curated tasks range from basic data lookups and computations to
 241 more advanced problems involving forecasting, challenging research problems including extreme
 242 event detection, forecast report generation, prediction analysis, and counterfactual reasoning. We
 243 design tasks with increasing difficulty levels (Easy, Medium, Hard) based on the complexity of
 244 tool usage required to answer them, from simple single-step data queries to multi-step analytical
 245 workflows. Table A.2 provide an overview of the task types we implement as part of our benchmark.
 246

247 For each task-type, we define natural language templates with placeholders such as location, variable,
 248 and time window. To create task-specific examples, these placeholders are filled by randomly
 249 sampling inputs, and the corresponding ground truth is computed deterministically using human-
 250 written or human-verified synthetic code applied to the raw ERA5 data. Figure 7 shows an example
 251 template, and a sample generated from it.
 252

253 Using our framework, we construct a benchmark dataset comprising 2158 test samples spread across
 254 46 tasks. For a detailed breakdown of dataset statistics, please refer to Appendix A.2. We provide
 255 more details about how the tasks are implemented in the subsequent sections.
 256

257 4.1.1 HUMAN-GENERATED TASKS

258 The human-generated tasks span across the Easy and Hard difficulty levels and represent realistic
 259 meteorological queries curated in conjunction with a domain expert. For each task, a graduate
 260 student created a question template and wrote Python code to answer the query. Easy tasks focus
 261 on basic data retrieval operations like finding extrema, querying specific values, and identifying
 262 locations with particular weather conditions. Medium-difficulty tasks introduce forecasting elements,
 263 asking for future weather predictions at specific locations and times, and/or implementing complex
 264 data analysis pipelines. Hard tasks incorporate more complex analytical concepts such as anomaly
 265 detection relative to baselines and counterfactual scenario analysis. They demand comprehensive
 266 meteorological expertise and mirror real-world operational workflows. These include extreme weather
 267 event detection, comprehensive weather assessments, and generation of detailed forecast discussions
 268 that span regional to global scales. For instance, ENSO outlook reports require synthesizing complex
 269 interactions between multiple atmospheric and oceanic variables to produce coherent, scientifically
 270 grounded forecasts. We source the expert-generated weather discussion reports from several online
 271




Figure 2: **Semi-synthetic task generation pipeline**: Weather-related texts are processed by a claim extraction agent to identify scientifically meaningful observational claims. Verified claims are transformed into reusable templates and manually reviewed. Code is generated by an LLM and verified by humans to validate each sample from a template against ERA5 meteorological data. We combine the verifier code with the templates and WeatherBench data to produce novel samples.

sources, such as the NOAA website¹ and IRI Seasonal Climate Forecasts/Outlooks². For extreme weather event tasks, we use records from the EM-DAT international disaster database (Delforge et al., 2025), matching event entries by date and location to the ERA5 data.

4.1.2 SEMI-SYNTHETIC TASK GENERATION

To increase task diversity, we implement a semi-synthetic pipeline that transforms unstructured weather-related text into verifiable benchmark tasks. Figure 2 provides an overview of the procedure. The process begins with a claim extraction agent that analyzes weather texts from various sources, using an LLM to identify scientifically meaningful observational claims about weather phenomena. The agent focuses on quantifiable changes, trends, extremes, and relationships between variables.

These claims are then converted into question templates where we can substitute different locations, time periods, and weather variables to generate multiple benchmark examples from each original claim. For each template, an LLM writes a verification code block that can validate any instance generated from that template against the ERA5 data. This verification step ensures that the generated questions are not only linguistically coherent but also scientifically accurate when tested against actual meteorological observations. We generate multiple candidate instances from each template through this approach. Finally, we manually review them for scientific interest and code correctness. In this way, we generate 30 distinct human-validated synthetic task types.

We also include a semi-synthetic meteorological claim verification task to test whether models are capable of validating claims extracted from meteorological reports against the weather data. From pre-processed NOAA meteorological reports, we select individual claims and pair them with the 24-hour slice of WeatherBench2 data corresponding to the report’s date. Negative instances are generated by systematically negating claims using an LLM. All examples are human-verified to ensure clarity, verifiability, and correctness of negation. More details on this task are included in Appendix A.2.2.

Difficulty	Human-Gen. Tasks	Human-Gen. Samples	Synthetic Tasks	Synthetic Samples	Total Samples
Easy	5	699	0	0	699
Medium	5	574	30	290	864
Hard	5	595	1	0	595
Total	15	1,868	31	290	2,158

Table 1: **ZEPHYRUSBENCH Statistics**: Number of unique tasks and samples, grouped by difficulty and generation form.

¹<https://www.wpc.ncep.noaa.gov/discussions/hpcdiscussions.php?disc=pmdepd>

²<https://iri.columbia.edu/our-expertise/climate/forecasts/seasonal-climate-forecasts/>

324 4.2 EVALUATION METRICS
325326 Since all our tasks are designed around weather tasks with objectively correct answers, we design an
327 evaluation pipeline that can assess the scientific correctness of the answers produced by the models.
328 The model answers fall into five primary categories: **numeric**, **temporal**, **spatial (location-based)**
329 and **descriptive**. Given that model outputs are in natural language, we evaluate them through a
330 multi-stage process:
331332 1. **Verification:** Determine whether the models response contains a relevant and valid answer. At this
333 stage, we merely assess whether or not the response has an appropriate answer to the given question,
334 and not its correctness. We use `gpt-4.1-mini` for this purpose.
335 2. **Extraction:** Extract the specific answer from the model response using another LLM prompt.
336 3. **Scoring:** Apply scoring methods specific to the type of question, which are detailed below.
337338 **Numerical Answers.** For numerical responses, we record the Standardized Median Absolute Error
339 between the predicted and reference values. In addition, we also report the 25%, 75% and 99%
340 quantiles of the standarized absolute error to provide a more complete picture of the error distribution.
341 We use quantiles rather than means because large outliers can significantly skew mean values,
342 obscuring typical model performance patterns. To compare across variables with different scales
343 and units, we divide the absolute error by the standard deviation of the corresponding variable in the
344 dataset.
345346 **Time-based Answers.** We evaluate tasks with time values as responses using Median Absolute Error.
347 We omit the standarization step, since all the answers are in the same units (that is, hours). Like the
numerical answers case, we also report the 25%, 75% and 99% quantiles.
348349 **Location-based Answers.** For questions whose answers are geographic locations, we first match the
350 extracted location name to one of the expected entries from the NaturalEarth dataset (e.g., mapping
351 “USA” to “United States of America”). For countries, we use the `country_converter` library
352 (Stadler, 2017). For other geographic entities such as continents and water bodies, we apply fuzzy
353 string matching (Bachmann et al., 2023), accepting matches above a predefined similarity threshold.
354355 To quantitatively assess the geographic deviation between predicted and reference locations, we
356 employ the Earth Mover’s Distance (EMD) (Monge, 1781) as a primary evaluation metric. We begin
357 by generating surface area-weighted masks over a latitude–longitude grid for both the predicted and
358 reference locations. These masks are normalized to form probability distributions. To account for the
359 curvature of the Earth, we compute pairwise distances between grid points using geodesic distance.
360 The EMD is then calculated using the `POT` library (Flamary et al., 2021). As a complementary metric,
we also report Location Accuracy, which simply measures whether the predicted and reference
location strings are an exact match.
361362 **Descriptive Answers.** To evaluate descriptive answers, we extract individual discussion points from
363 both the model’s response and the reference answer. We then classify each extracted claim from
364 the model’s response as either SUPPORTED, REFUTED, or NEUTRAL against the reference answer,
365 obtaining logit scores from the language model and applying softmax normalization. Similarly, we
366 perform the same procedure for claims from the reference text compared against the model response.
367368 We then define two complementary metrics: precision measures the validity of the model’s claims
369 by computing the proportion that are supported rather than refuted by the reference answer, excluding
370 neutral classifications: Precision = $\frac{\sum_{i \in S} P_{\text{model} \rightarrow \text{ref}}(\text{Supported}_i)}{\sum_{i \in S} P_{\text{model} \rightarrow \text{ref}}(\text{Supported}_i) + \sum_{i \in S} P_{\text{model} \rightarrow \text{ref}}(\text{Refuted}_i)}$
371 where $S = \{i : P_{\text{model} \rightarrow \text{ref}}(\text{Neutral}_i) < 0.5\}$ and $P_{\text{model} \rightarrow \text{ref}}(\text{Supported}_i)$ denotes the probability that
372 model claim i is supported by the reference answer.
373374 Recall measures coverage by evaluating how well the model response addresses the ref-
375 erence claims, computed as the average support probability across all reference points:
376377
$$\text{Recall} = \frac{1}{N} \sum_{i=1}^N P_{\text{ref} \rightarrow \text{model}}(\text{Supported}_i)$$
 where N is the number of reference claims and
 $P_{\text{ref} \rightarrow \text{model}}(\text{Supported}_i)$ denotes the probability that reference claim i is supported by the model
answer.
378




Figure 3: Percentage of questions in the complete dataset answered correctly by each LLM and model type. Definitions of correctness for each question type are detailed in Appendix A.4.




Figure 4: Plots showing (top) error distribution on numerical tasks (bottom-left) location accuracy (bottom-right) discussion scores for GPT-5-Mini and Gemini-2.5-Flash.

Finally, we define the **discussion score** as the F1 score = $\frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$.

Extreme Weather Tasks. In order to evaluate the extreme-weather tasks, we report two metrics: (1) F1 score, which only assesses whether the model correctly predicts the *occurrence* of an extreme event anywhere in the world, without considering event type or exact location. (2) Earthmover’s Distance, which measures the agreement between the reference and predicted list of countries.

Correctness. For ease of presentation, we define correctness criteria that vary depending on the task type. Rather than requiring exact matches, we consider an answer correct if it falls within an acceptable range of the target response. The precise criteria for determining correctness for each task type are detailed in Appendix A.4.

5 EXPERIMENTAL RESULTS

We evaluate model performance across all task types from Section 4 using four backend models: OpenAI GPT-5-Mini, GPT-5-Nano, Google Gemini 2.5 Flash, and OpenAI gpt-oss-120b. We compare three experimental settings: (1) a text-only baseline that attempts to answer weather reasoning questions using only natural language metadata without access to structured weather data or numerical inputs, (2) ZEPHYRUS-DIRECT, and (3) ZEPHYRUS-REFLECTIVE. The text-only baseline measures the extent to which models can utilize their prior meteorological knowledge.

432 The correctness results across all models and settings are presented in Figure 3. We observe that the
 433 ZEPHYRUS agents significantly outperform the text-only baseline across all models, demonstrating
 434 the agentic framework’s ability to effectively ground answers by leveraging meteorological data
 435 from WeatherBench. For GPT-5-Mini, ZEPHYRUS-DIRECT and ZEPHYRUS-REFLECTIVE achieve
 436 48.5% and 54.7% correctness respectively, compared to only 19.9% for the text-only baseline. This
 437 substantial improvement holds consistently across other LLMs, with ZEPHYRUS agents achieving
 438 28.6-35.4% higher correctness than their text-only counterparts. The multi-turn execute-observe-
 439 solution framework implemented in ZEPHYRUS-REFLECTIVE enables it to outperform ZEPHYRUS-
 440 DIRECT by 3.5-6.2% across most models. The exception is GPT-5-Nano, where ZEPHYRUS-
 441 REFLECTIVE performs slightly worse than ZEPHYRUS-DIRECT, likely due to the smaller model’s
 442 limited reasoning capabilities affecting the more complex multi-turn approach.
 443

444 Figure 4 enables fine-grained analysis of error distributions for numerical tasks, location accuracy, and
 445 discussion scores for descriptive answers using GPT-5-Mini and Gemini 2.5 Flash as backend models.
 446 The agents particularly excel at numerical and location prediction tasks, achieving substantially lower
 447 Standardized Absolute Errors and higher location accuracies compared to text-only baselines. For
 448 location prediction, ZEPHYRUS-REFLECTIVE with GPT-5-Mini achieves strong performance with
 449 86.6% accuracy. Once again, the reflective variant enjoys a small benefit in performance over the
 450 Direct approach. The difference between ZEPHYRUS-DIRECT and ZEPHYRUS-REFLECTIVE is
 451 pronounced on numerical tasks for GPT-5-Mini, while both variants perform similarly with Gemini
 452 2.5 Flash.

453 However, all models struggle with the challenging task of generating textual weather reports. The
 454 best performing model (ZEPHYRUS-REFLECTIVE with GPT-5-Mini) only achieves a discussion
 455 score of 0.177. Nevertheless, ZEPHYRUS-REFLECTIVE demonstrates significant advantages over
 456 both ZEPHYRUS-DIRECT and text-only variants for these descriptive tasks. While the text-only
 457 variant lacks access to meteorological information, ZEPHYRUS-DIRECT produces rigid answers
 458 by directly outputting program results, making it ill-suited for nuanced textual generation. The
 459 execute-observe-solution framework in ZEPHYRUS-REFLECTIVE proves more effective.

460 Performance breakdown by difficulty level reveals interesting patterns (detailed results in Appendices
 461 A.5 and A.7). On easy tasks, which primarily involve data analysis questions, ZEPHYRUS agents
 462 perform well with 78.7-88.1% correctness. Medium difficulty tasks show moderate performance with
 463 39.9-50.5% correctness. However, on hard tasks, all models struggle significantly, with ZEPHYRUS
 464 agents achieving similar performance to text-only baselines. This suggests that while current LLMs
 465 can effectively solve simple data analysis problems that pop up in meteorology, they do not yet
 466 possess the capability to reason about abstract weather phenomena even when provided with tools.

467 Task-wise analysis of “Hard” tasks reveals nuanced insights. For generating meteorological dis-
 468 cussions and forecasts for the continental United States, models show promise with ZEPHYRUS-
 469 REFLECTIVE + GPT-5-Mini achieving an average discussion score of 0.31. This contrasts sharply
 470 with global climate forecasting tasks spanning three months, where all models fail completely,
 471 highlighting the current limitations in long-term, large-scale weather reasoning.

472 6 CONCLUSION

473 We tackled the challenging problem of enabling LLMs to reason over high-dimensional weather data
 474 by developing, to our knowledge, the first agentic model for meteorology. Our contributions include:
 475 (1) ZEPHYRUSWORLD, an agentic environment with comprehensive meteorological tools, (2) the
 476 ZEPHYRUS family of agents that leverage these tools, and (3) a scalable data pipeline producing
 477 a large, diverse benchmark dataset (ZEPHYRUSBENCH). Our empirical evaluation shows that the
 478 agentic framework enables effective reasoning about meteorological data, significantly outperforming
 479 text-only baselines. The agents excel at most tasks but struggle with complex challenges like forecast
 480 report generation. Beyond advancing weather science, our work provides a sandbox for developing
 481 more effective agentic workflows. Future work could explore using larger datasets to train agents that
 482 produce more scientifically accurate responses.

483 6.1 ETHICS STATEMENT

484 The authors have followed the ICLR Code of Ethics and have no conflicts of interest to declare.

486 6.2 REPRODUCIBILITY STATEMENT
487

488 The `ZEPHYRUSBENCH` dataset will be made public for use as a benchmark dataset and further
489 research on AI for weather science. All code used to create the dataset and the Zephyrus models will
490 be open-sourced. Detailed results have been included in the paper and its appendices, along with the
491 LLM models used to produce all results.

492 6.3 LLM USAGE STATEMENT
493

494 All LLM usage for the creation of `ZEPHYRUSBENCH` and the evaluation of different models
495 (`ZEPHYRUS-REFLECTIVE`, `ZEPHYRUS-DIRECT`) has been carefully described in the paper. We
496 acknowledge the routine use of LLMs for coding assistance and refinement of writing. LLMs were
497 not used for ideation, conceptual development, literature review, or other substantial contributions of
498 this work.

500 REFERENCES
501

502 Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel
503 Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language
504 model for few-shot learning. *Advances in neural information processing systems*, 35:23716–23736,
505 2022.

506 Richard B Alley, Kerry A Emanuel, and Fuqing Zhang. Advances in weather prediction. *Science*,
507 363(6425):342–344, 2019.

509 Max Bachmann, layday, Georgia Kokkinou, Jeppe Fihl-Pearson, dj, Henry Schreiner, Moshe Sherman,
510 Michał Górný, pekkarr, Delfini, Dan Hess, Guy Rosin, Hugo Le Moine, Khuang Tang, Nicolas
511 Renkamp, Trenton H, glynn, and odidev. `rapidfuzz/rapidfuzz`: Release 3.6.1, December 2023.
512 URL <https://doi.org/10.5281/zenodo.10440201>.

513 Peter Bauer, Alan Thorpe, and Gilbert Brunet. The quiet revolution of numerical weather prediction.
514 *Nature*, 525(7567):47–55, 2015.

516 Kaifeng Bi, Lingxi Xie, Hengheng Zhang, Xin Chen, Xiaotao Gu, and Qi Tian. Accurate medium-
517 range global weather forecasting with 3d neural networks. *Nature*, 619(7970):533–538, 2023.

518 Abeba Birhane, Atoosa Kasirzadeh, David Leslie, and Sandra Wachter. Science in the age of large
519 language models. *Nature Reviews Physics*, 5(5):277–280, 2023.

521 Cristian Bodnar, Wessel P Bruinsma, Ana Lucic, Megan Stanley, Johannes Brandstetter, Patrick
522 Garvan, Maik Riechert, Jonathan Weyn, Haiyu Dong, Anna Vaughan, et al. Aurora: A foundation
523 model of the atmosphere. *arXiv preprint arXiv:2405.13063*, 2024.

525 Daniil A Boiko, Robert MacKnight, Ben Kline, Gabe Gomes, et al. Autonomous chemical research
526 with large language models. *Nature*, 624:570–578, 2023.

527 Andres M Bran, Sam Cox, Oliver Schilter, Carlo Baldassari, Andrew D White, Philippe Schwaller,
528 et al. Augmenting large language models with chemistry tools. *Nature Machine Intelligence*, 6(5):
529 525–535, 2024.

531 Jian Chen, Peilin Zhou, Yining Hua, Dading Chong, Meng Cao, Yaowei Li, Zixuan Yuan, Bing Zhu,
532 and Junwei Liang. Vision-language models meet meteorology: Developing models for extreme
533 weather events detection with heatmaps. *arXiv preprint arXiv:2406.09838*, 2024a.

534 Wei Chen, Xixuan Hao, Yuankai Wu, and Yuxuan Liang. Terra: A multimodal spatio-temporal
535 dataset spanning the earth. *Advances in Neural Information Processing Systems*, 37:66329–66356,
536 2024b.

538 Zesen Cheng, Sicong Leng, Hang Zhang, Yifei Xin, Xin Li, Guanzheng Chen, Yongxin Zhu, Wenqi
539 Zhang, Ziyang Luo, Deli Zhao, et al. Videollama 2: Advancing spatial-temporal modeling and
audio understanding in video-llms. *arXiv preprint arXiv:2406.07476*, 2024.

540 Yunfei Chu, Jin Xu, Xiaohuan Zhou, Qian Yang, Shiliang Zhang, Zhijie Yan, Chang Zhou, and
 541 Jingren Zhou. Qwen-audio: Advancing universal audio understanding via unified large-scale
 542 audio-language models. *arXiv preprint arXiv:2311.07919*, 2023.

543

544 Alexandre Défossez, Laurent Mazaré, Manu Orsini, Amélie Royer, Patrick Pérez, Hervé Jégou,
 545 Edouard Grave, and Neil Zeghidour. Moshi: a speech-text foundation model for real-time dialogue.
 546 *arXiv preprint arXiv:2410.00037*, 2024.

547

548 D. Delforge, V. Wathelet, R. Below, C. Lanfredi Sofia, M. Tonnelier, J. A. F. van Loenhout, and
 549 N. Speybroeck. EM-DAT: The Emergency Events Database. *International Journal of Disaster Risk
 Reduction*, pp. 105509, 2025. doi: 10.1016/j.ijdrr.2025.105509. URL <https://doi.org/10.1016/j.ijdrr.2025.105509>.

550

551

552 Seungheon Doh, Keunwoo Choi, and Juhan Nam. Talkplay: Multimodal music recommendation
 553 with large language models. *arXiv preprint arXiv:2502.13713*, 2025.

554

555 RÃ©mi Flamary, Nicolas Courty, Alexandre Gramfort, Mokhtar Z. Alaya, AurÃ©lie Boisbunon,
 556 Stanislas Chambon, Laetitia Chapel, Adrien Corenflos, Kilian Fatras, Nemo Fournier, LÃ©o
 557 Gautheron, Nathalie T.H. Gayraud, Hicham Janati, Alain Rakotomamonjy, Ievgen Redko, Antoine
 558 Rolet, Antony Schutz, Vivien Seguy, Danica J. Sutherland, Romain Tavenard, Alexander Tong,
 559 and Titouan Vayer. Pot: Python optimal transport. *Journal of Machine Learning Research*, 22(78):
 560 1–8, 2021. URL <http://jmlr.org/papers/v22/20-451.html>.

561

562 Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and
 563 Graham Neubig. Pal: Program-aided language models. In *International Conference on Machine
 Learning*, pp. 10764–10799. PMLR, 2023.

564

565 Sreyan Ghosh, Zhifeng Kong, Sonal Kumar, S Sakshi, Jaehyeon Kim, Wei Ping, Rafael Valle, Dinesh
 566 Manocha, and Bryan Catanzaro. Audio flamingo 2: An audio-language model with long-audio
 567 understanding and expert reasoning abilities. *arXiv preprint arXiv:2503.03983*, 2025.

568

569 Hans Hersbach, Bill Bell, Paul Berrisford, Shoji Hirahara, András Horányi, Joaquín Muñoz-Sabater,
 570 Julien Nicolas, Carole Peubey, Raluca Radu, Dinand Schepers, et al. The era5 global reanalysis.
 571 *Quarterly journal of the royal meteorological society*, 146(730):1999–2049, 2020.

572

573 Kexin Huang, Serena Zhang, Hanchen Wang, Yuanhao Qu, Yingzhou Lu, et al. Biomni: A general-
 574 purpose biomedical ai agent. *bioRxiv*, 2025. doi: 10.1101/2025.05.30.656746.

575

576 Kevin Maik Jablonka, Qianxiang Ai, Alexander Al-Feghali, Shruti Badhwar, Joshua D Bocarsly,
 577 Andres M Bran, Stefan Bringuer, L Catherine Brinson, Kamal Choudhary, Defne Circi, et al.
 578 14 examples of how llms can transform materials science and chemistry: a reflection on a large
 579 language model hackathon. *Digital discovery*, 2(5):1233–1250, 2023.

580

581 Juyong Jiang, Fan Wang, Jiasi Shen, Sungju Kim, and Sunghun Kim. A survey on large language
 582 models for code generation. *arXiv preprint arXiv:2406.00515*, 2024.

583

584 Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik R
 585 Narasimhan. Swe-bench: Can language models resolve real-world github issues? In *The Twelfth
 586 International Conference on Learning Representations*.

587

588 Dmitrii Kochkov, Janni Yuval, Ian Langmore, Peter Norgaard, Jamie Smith, Griffin Mooers, Milan
 589 Klöwer, James Lottes, Stephan Rasp, Peter Düben, et al. Neural general circulation models for
 590 weather and climate. *Nature*, 632(8027):1060–1066, 2024.

591

592 Thorsten Kurth, Shashank Subramanian, Peter Harrington, Jaideep Pathak, Morteza Mardani, David
 593 Hall, Andrea Miele, Karthik Kashinath, and Anima Anandkumar. Fourcastnet: Accelerating global
 594 high-resolution weather forecasting using adaptive fourier neural operators. In *Proceedings of the
 595 platform for advanced scientific computing conference*, pp. 1–11, 2023.

596

597 Remi Lam, Alvaro Sanchez-Gonzalez, Matthew Willson, Peter Wirnsberger, Meire Fortunato, Ferran
 598 Alet, Suman Ravuri, Timo Ewalds, Zach Eaton-Rosen, Weihua Hu, et al. Learning skillful
 599 medium-range global weather forecasting. *Science*, 382(6677):1416–1421, 2023.

594 Ge Lei, Ronan Doherty, and Samuel J Cooper. Materials science in the era of large language models:
 595 a perspective. *Digital Discovery*, 3(7):1257–1272, 2024.

596

597 Haobo Li, Zhaowei Wang, Jiachen Wang, Alexis Kai Hon Lau, and Huamin Qu. Cllmate: A
 598 multimodal lilm for weather and climate events forecasting. *arXiv preprint arXiv:2409.19058*,
 599 2024.

600 Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven
 601 Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum
 602 distillation. *Advances in neural information processing systems*, 34:9694–9705, 2021.

603 Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-
 604 training for unified vision-language understanding and generation. In *International conference on*
 605 *machine learning*, pp. 12888–12900. PMLR, 2022.

606 Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-
 607 training with frozen image encoders and large language models. In *International conference on*
 608 *machine learning*, pp. 19730–19742. PMLR, 2023.

609 Bin Lin, Yang Ye, Bin Zhu, Jiaxi Cui, Munan Ning, Peng Jin, and Li Yuan. Video-llava: Learning
 610 united visual representation by alignment before projection. In *Proceedings of the 2024 Conference*
 611 *on Empirical Methods in Natural Language Processing*, pp. 5971–5984, 2024.

612

613 Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction
 614 tuning, 2023a.

615

616 Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning, 2023b.

617

618 Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in*
 619 *neural information processing systems*, 36:34892–34916, 2023c.

620

621 Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee.
 622 Llava-next: Improved reasoning, ocr, and world knowledge, January 2024. URL <https://llava-v1.github.io/blog/2024-01-30-llava-next/>.

623

624 Chengqian Ma, Zhanxiang Hua, Alexandra Anderson-Frey, Vikram Iyer, Xin Liu, and Lianhui
 625 Qin. Weatherqa: Can multimodal language models reason about severe weather? *arXiv preprint*
 626 *arXiv:2406.11217*, 2024.

627

628 Franco Molteni, Roberto Buizza, Tim N Palmer, and Thomas Petroliagis. The ecmwf ensemble
 629 prediction system: Methodology and validation. *Quarterly journal of the royal meteorological*
society, 122(529):73–119, 1996.

630

631 Gaspard Monge. Mémoire sur la théorie des déblais et des remblais. *Mem. Math. Phys. Acad. Royale*
Sci., pp. 666–704, 1781.

632

633 Natural Earth. Natural earth data. <https://www.naturalearthdata.com/>, 2024. Accessed:
 634 2024-11-15.

635

636 Tung Nguyen, Johannes Brandstetter, Ashish Kapoor, Jayesh K Gupta, and Aditya Grover. Climax:
 A foundation model for weather and climate. *arXiv preprint arXiv:2301.10343*, 2023.

637

638 Tung Nguyen, Rohan Shah, Hritik Bansal, Troy Arcomano, Romit Maulik, Rao Kotamarthi, Ian
 639 Foster, Sandeep Madireddy, and Aditya Grover. Scaling transformer neural networks for skillful
 640 and reliable medium-range weather forecasting. *Advances in Neural Information Processing*
Systems, 37:68740–68771, 2024.

641

642 Jaideep Pathak, Shashank Subramanian, Peter Harrington, Sanjeev Raja, Ashesh Chattopadhyay,
 643 Morteza Mardani, Thorsten Kurth, David Hall, Zongyi Li, Kamyar Azizzadenesheli, et al. Fourcast-
 644 net: A global data-driven high-resolution weather model using adaptive fourier neural operators.
 645 *arXiv preprint arXiv:2202.11214*, 2022.

646

647 Ilan Price, Alvaro Sanchez-Gonzalez, Ferran Alet, Tom R Andersson, Andrew El-Kadi, Dominic
 Masters, Timo Ewalds, Jacklynn Stott, Shakir Mohamed, Peter Battaglia, et al. Probabilistic
 648 weather forecasting with machine learning. *Nature*, 637(8044):84–90, 2025.

648 Stephan Rasp, Stephan Hoyer, Alexander Merose, Ian Langmore, Peter Battaglia, Tyler Russell,
 649 Alvaro Sanchez-Gonzalez, Vivian Yang, Rob Carver, Shreya Agrawal, et al. Weatherbench 2: A
 650 benchmark for the next generation of data-driven global weather models. *Journal of Advances in*
 651 *Modeling Earth Systems*, 16(6):e2023MS004019, 2024.

652 Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke
 653 Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach
 654 themselves to use tools. *Advances in Neural Information Processing Systems*, 36:68539–68551,
 655 2023.

656 Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu
 657 Yao. Reflexion: language agents with verbal reinforcement learning. 36:8634–8652,
 658 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/1b44b878bb782e6954cd888628510e90-Paper-Conference.pdf.

659 Konstantin Stadler. The country converter coco - a python package for converting country names
 660 between different classification schemes. *Journal of Open Source Software*, 2(16):332, 2017. doi:
 661 10.21105/joss.00332. URL <https://doi.org/10.21105/joss.00332>.

662 Christina V Theodoris, Ling Xiao, Anant Chopra, Mark D Chaffin, Zeina R Al Sayed, Matthew C
 663 Hill, Helene Mantineo, Elizabeth M Brydon, Zexian Zeng, X Shirley Liu, et al. Transfer learning
 664 enables predictions in network biology. *Nature*, 618(7965):616–624, 2023.

665 Sumanth Varambally, Veeramakali Vignesh Manivannan, Yasaman Jafari, Luyu Han, Zachary Novack,
 666 Zhirui Xia, Salva Rühling Cachay, Srikanth Eranki, Ruijia Niu, Taylor Berg-Kirkpatrick, Duncan
 667 Watson-Parris, Yian Ma, and Rose Yu. Aquilon: Towards building multimodal weather LLMs. In
 668 *ICML 2025 Workshop on Assessing World Models*, 2025. URL <https://openreview.net/forum?id=KVxOwEYAF4>.

669 Lei Wang, Shan Dong, Yuhui Xu, Hanze Dong, Yalu Wang, Amrita Saha, Ee-Peng Lim, Caiming
 670 Xiong, and Doyen Sahoo. Mathhay: An automated benchmark for long-context mathematical
 671 reasoning in llms. *arXiv preprint arXiv:2410.04698*, 2024.

672 Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. Simvlm: Simple
 673 visual language model pretraining with weak supervision. In *International Conference on Learning
 674 Representations*, 2022.

675 Junda Wu, Zachary Novack, Amit Namburi, Jiaheng Dai, Hao-Wen Dong, Zhouhang Xie, Carol
 676 Chen, and Julian McAuley. Futga: Towards fine-grained music understanding through temporally-
 677 enhanced generative augmentation. *arXiv preprint arXiv:2407.20445*, 2024a.

678 Junda Wu, Zachary Novack, Amit Namburi, Jiaheng Dai, Hao-Wen Dong, Zhouhang Xie, Carol
 679 Chen, and Julian McAuley. Futga-mir: Enhancing fine-grained and temporally-aware music
 680 understanding with music information retrieval. In *International Conference on Acoustics, Speech,
 681 and Signal Processing (ICASSP)*, 2025.

682 Kehan Wu, Yingce Xia, Pan Deng, Renhe Liu, Yuan Zhang, Han Guo, Yumeng Cui, Qizhi Pei, Lijun
 683 Wu, Shufang Xie, et al. Tamgen: drug design with target-aware molecule generation through a
 684 chemical language model. *Nature Communications*, 15(1):9360, 2024b.

685 Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun
 686 Zhang, Shaokun Zhang, Jiale Liu, Ahmed Hassan Awadallah, Ryen W White, Doug Burger, and
 687 Chi Wang. Autogen: Enabling next-gen LLM applications via multi-agent conversations. In *First
 688 Conference on Language Modeling*, 2024c. URL <https://openreview.net/forum?id=BAakY1hNKS>.

689 Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao.
 690 React: Synergizing reasoning and acting in language models. In *International Conference on
 691 Learning Representations (ICLR)*, 2023a.

692 Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao.
 693 React: Synergizing reasoning and acting in language models. In *International Conference on
 694 Learning Representations (ICLR)*, 2023b.

702 Boqiang Zhang, Kehan Li, Zesen Cheng, Zhiqiang Hu, Yuqian Yuan, Guanzheng Chen, Sicong Leng,
 703 Yuming Jiang, Hang Zhang, Xin Li, et al. Videollama 3: Frontier multimodal foundation models
 704 for image and video understanding. *arXiv preprint arXiv:2501.13106*, 2025.

705 Hang Zhang, Xin Li, and Lidong Bing. Video-llama: An instruction-tuned audio-visual language
 706 model for video understanding. In *Proceedings of the 2023 Conference on Empirical Methods in*
 707 *Natural Language Processing: System Demonstrations*, pp. 543–553, 2023.


709 Xiaokang Zhang, Jing Zhang, Zeyao Ma, Yang Li, Bohan Zhang, Guanlin Li, Zijun Yao, Kangli Xu,
 710 Jinchang Zhou, Daniel Zhang-Li, et al. Tablellm: Enabling tabular data manipulation by llms in
 711 real office usage scenarios. *arXiv preprint arXiv:2403.19318*, 2024.

712 Yue Zhao, Ishan Misra, Philipp Krähenbühl, and Rohit Girdhar. Learning video representations from
 713 large language models. In *arXiv preprint arXiv:2212.04501*, 2022.

715 Yizhen Zheng, Huan Yee Koh, Maddie Yang, Li Li, Lauren T May, Geoffrey I Webb, Shirui Pan,
 716 and George Church. Large language models in drug discovery and development: From disease
 717 mechanisms to clinical trials. *arXiv preprint arXiv:2409.04481*, 2024.

719 A APPENDIX

721 A.1 CODE EXECUTION SERVER



734 Figure 5: **Code Execution Server.** ZEPHYRUS sends parallel requests to the server, which distributes
 735 them to available workers. Each worker acquires resources from tool pools, loads datasets, injects
 736 tools into the execution environment, executes code, and returns results or errors to the agent.

737 In order to execute the code requests from the client, we implement a custom code-execution server
 738 program. More specifically, we implement a FastAPI-based server-client architecture where clients
 739 send code execution requests to a dedicated execution server that processes them in parallel. The
 740 system maintains resource pools for each tool component to prevent contention and enable true
 741 parallelism. Each pool contains one or more instances of the above tools. A resource manager
 742 implements acquire/release semantics to ensure each execution thread has exclusive access to a
 743 complete set of tools while preventing deadlocks.

745 Each execution follows a strict protocol: acquire resources from pools, load requested datasets,
 746 inject tool instances into the execution environment, and execute user code with timeout protection.
 747 The system captures all outputs and error information, which are sent back to the client for further
 748 processing by the agent. Figure 5 provides an overview of the server.

750 A.2 DATASET DETAILS

751 Table A.2 details all the tasks in ZEPHYRUSBENCH, and table 4.1.2 reports the number of samples
 752 generated grouped by difficulty and type.

754 For weather tasks, we leverage the ERA5 reanalysis dataset (Hersbach et al., 2020), specifically from
 755 WeatherBench 2 (Rasp et al., 2024), which provides global atmospheric data from 1979 to 2022.
 We use 1.5° spatial resolution with 6-hourly temporal resolution, and include 4 surface variables

756
757
758
759
760
761

ID	Natural Language Description	Answer Type	Difficulty	Type
1	Which geographic feature experienced the highest/lowest average value of a weather variable	Location	Easy	Human
2	What is the min/max/average/median value of a weather variable at a specific location	Numerical	Easy	Human
3	Which sublocation has the highest/lowest recorded variable value	Location	Easy	Human
4	How many hours from start did a location experience extremum	Temporal	Easy	Human
5	What is the weather variable value at a location at a specific time	Numerical	Easy	Human
6	What will the variable be at a location after time interval (forecast)	Numerical	Medium	Human
7	When will location experience its extremum in future period (forecast)	Temporal	Medium	Human
8	Identify extreme weather events that will occur in the next N hours (forecast)	List of locations	Hard	Human
9	Check if extreme weather events are currently happening	List of locations	Hard	Human
10	Which geographic features experienced unusual weather anomalies compared to baseline	List of locations	Medium	Human
11	Does maximum weather variable occur at same or adjacent grid point as another variable (forecast)	Yes/No	Medium	Synthetic
12	Does maximum weather variable in region remain lower than future maximum (forecast)	Yes/No	Medium	Synthetic
13	Does maximum weather variable occur at higher latitude than in another region (forecast)	Yes/No	Medium	Synthetic
14	Does mean weather variable in one region exceed another by specified amount (forecast)	Yes/No	Medium	Synthetic
15	Does mean weather variable exceed threshold while maximum of another stays below (forecast)	Yes/No	Medium	Synthetic
16	Does mean weather variable within region exceed specified threshold (forecast)	Yes/No	Medium	Synthetic
17	Does weather variable exceed threshold within any part of region (forecast)	Yes/No	Medium	Synthetic
18	Does weather variable exceed threshold in more grid points in one region than another (forecast)	Yes/No	Medium	Synthetic
19	Does area where weather variable exceeds threshold cover more than percentage of region (forecast)	Yes/No	Medium	Synthetic
20	Does area-averaged weather variable exceed threshold while another stays below (forecast)	Yes/No	Medium	Synthetic
21	Does maximum weather variable in one region exceed threshold while another stays below (forecast)	Yes/No	Medium	Synthetic
22	Does maximum weather variable within region exceed specified threshold (forecast)	Yes/No	Medium	Synthetic
23	Does maximum weather variable occur at latitude farther north than in another region (forecast)	Yes/No	Medium	Synthetic
24	Does maximum weather variable stay above threshold while another stays below (forecast)	Yes/No	Medium	Synthetic
25	Does maximum weather variable in one region exceed another by specified amount (forecast)	Yes/No	Medium	Synthetic
26	Does minimum weather variable within region remain above threshold (forecast)	Yes/No	Medium	Synthetic
27	What is the area where multiple weather variables exceed their percentile values (forecast)	Numerical	Medium	Synthetic
28	What is the area where weather variable exceeds its median value (forecast)	Numerical	Medium	Synthetic
29	What is the displacement between centroids of areas with top 10% values (forecast)	Numerical	Medium	Synthetic
30	What is the distance between centroids of maximum weather variable value areas (forecast)	Numerical	Medium	Synthetic
31	What is the maximum difference in weather variable between grid points within region (forecast)	Numerical	Medium	Synthetic
32	What is the minimum weather variable where another variable exceeds percentile (forecast)	Numerical	Medium	Synthetic
33	What is the difference between maximum weather variables in two regions (forecast)	Numerical	Medium	Synthetic
34	What is the difference in mean weather variable between two regions (forecast)	Numerical	Medium	Synthetic
35	What is the displacement of minimum weather variable location after time window (forecast)	Numerical	Medium	Synthetic
36	What is the latitude difference between centroids of high weather variable areas (forecast)	Numerical	Medium	Synthetic
37	What is the maximum weather variable difference between two regions (forecast)	Numerical	Medium	Synthetic
38	What is the mean weather variable where another variable exceeds percentile (forecast)	Numerical	Medium	Synthetic
39	What is the mean weather variable where another exceeds percentile threshold (forecast)	Numerical	Medium	Synthetic
40	What is the weather variable value where another variable reaches maximum (forecast)	Numerical	Medium	Synthetic
41	Generate comprehensive global climate forecast for temperature and precipitation for next 3 months (forecast)	Description	Medium	Human
42	Provide detailed meteorological discussion and forecast for continental United States (forecast)	Description	Hard	Human
43	Generate ENSO climate update and outlook based on atmospheric data (forecast)	Description	Hard	Human
44	How will weather variable change after specified time given an intervention (a specified change) in variable in the present (counterfactual)	Numerical	Medium	Human
45	What is the value of the input parameter of the simulator model that produces the simulation output	Numerical	Hard	Human
46	Check whether the given claim extracted from meteorological report is supported by the data	Yes/No	Hard	Synthetic

805
806
807
808
809
Table 2: Complete set of Weather Tasks, grouped by difficulty.

810 and 5 atmospheric variables at 13 pressure levels. For each task-type, we define natural language
 811 templates with placeholders such as location, variable, and time window. For example, Task 1 is
 812 defined as ‘Which {geofeature} experienced the {extremum_direction} average {variable}?’ To
 813 create task-specific examples, these placeholders are filled by randomly sampled inputs. Ground
 814 truths are derived through a deterministic procedure by applying human-written or human-verified
 815 synthetic code to the raw ERA5 data and other supplementary data.

816

817

818 A.2.1 HUMAN-GENERATED TASKS

819

820 Tasks 1 through 7 rely entirely on the raw ERA5 data; they include basic data lookups and computations,
 821 as well as more advanced forecasting. The Geolocator is introduced to enable these and other
 822 location-related tasks. It is a wrapper for the Natural Earth dataset (Natural Earth, 2024) that maps
 823 ERA5 grid points to natural language location names of countries, states, and water bodies.

824

825

826 For Task 8 and Task 9, which involve extreme event detection, we use records from the EM-DAT
 827 international disaster database (Delforge et al., 2025), matching event entries by date and location to
 828 ERA5 data.

829

830

831 Input for anomaly comparison (Task 10) comprises two components: recent global data and quantile
 832 statistics derived from a historical reference period. Ground truths are calculated by comparing the
 833 recent dataset against historical quantile thresholds. Locations where the recent values significantly
 834 exceed or fall below the reference quantile are flagged as anomalous. We then use the Geolocator to
 835 map flagged grid points to natural language region names.

836

837

838 Report generation tasks (ID 41, 42, 43) are designed to evaluate model climate forecasting and
 839 interpretation capabilities based on ERA5 atmospheric datasets. They all use global weather fields
 840 over the given time duration as context. Task 41 requires generating a comprehensive global climate
 841 forecast report for temperature and precipitation for a three month forward horizon. The task instructs
 842 the report to be structured into separate sections for precipitation and temperature, and to provide
 843 region-specific forecasts with probability-based language. Task 42 focuses on the continental United
 844 States, where the model must provide a detailed meteorological discussion and forecast, including
 845 current weather system positions and movements, temperature trends and expected changes over the
 846 coming days, precipitation patterns and likelihood of significant events, pressure system evolution
 847 and impacts, and notable atmospheric features such as fronts and jet stream positioning. Task 43
 848 requires an ENSO (El Niño–Southern Oscillation) climate update and outlook. Models are tasked to
 849 analyze atmospheric variables to assess the current ENSO phase, evaluate strength and persistence
 850 indicators, forecast evolution over the next 3 – 6 months, and discuss global implications using
 851 probability-based language and standard ENSO terminology. Ground truth reports for these tasks
 852 are obtained from authoritative climate prediction and monitoring sources, which provide validated
 853 assessments of global and regional climate outlooks and ENSO conditions. The answer sources for
 854 these three tasks are NOAA Global Climate Reports, NOAA National Weather Service Area Forecast
 855 Discussions, and WMO ENSO Reports, respectively.

856

857

858 Task 44 and Task 45 both rely on the JAX-GCM simulator, an intermediate-complexity atmospheric
 859 model built on NeuralGCM’s dynamical core (Kochkov et al., 2024). Task 44 is to assess the causal
 860 impact of localized perturbations on atmospheric states. To obtain each specific sample, a variable,
 861 location, and perturbation magnitude are first sampled, and a Gaussian mask is applied to induce
 862 the desired perturbation at the chosen location. The simulator is then run twice, once starting from
 863 the unperturbed initial state and once with the imposed perturbation. At the specified simulation
 864 end time, the target variable from both simulations is extracted and compared, with the difference
 865 quantifying the perturbation’s impact.

866

867

868 Task 45 is a black-box optimization climate simulation task. The input consists of two components:
 869 (i) a segment of recent global data spanning a specified duration and interval, and (ii) simulated
 870 data generated by our JAX-GCM simulator. In the simulation, we vary one input parameter of
 871 the model by sampling its value randomly from the range [0,1], then save the resulting simulation
 872 output. The objective of the task is to estimate the original value of the underlying input parameter
 873 from observable simulation outputs. Since the climate simulator is presented as a black box, the
 874 model must infer the parameter solely from the input-output mapping, which can be highly nonlinear
 875 and sensitive to small parameter changes. By evaluating the model on novel simulator outputs, we

864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
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
benchmark its general handling of a domain optimization problem. Performance is assessed by comparing the predicted and ground-truth parameter values.

A.2.2 METEOROLOGICAL CLAIM VERIFICATION

The Meteorological Claim Verification task (Task 46), We start with NOAA monthly climate report webpages (1988–2024). These are downloaded and subsequently scraped into structured text files. The scraped reports are then consolidated into CSV files, which contain raw textual summaries of meteorological conditions. To standardize this content, we prompt the LLM to remove author notes, editorial comments, data collection methods, map/color references, visual elements, and model disclaimers. We preserve forecast events, locations, timing, and patterns. From the resulting text-only, stand-alone weather reports, we extract individual claims to use for task examples.

The following data shows meteorological conditions over a 24-hour period:

```
{'type': 'wb2', 'variables': ['mean_sea_level_pressure',  
'10m_u_component_of_wind',  
'10m_v_component_of_wind', '2m_temperature',  
'geopotential', 'specific_humidity', 'temperature',  
'u_component_of_wind', 'v_component_of_wind'],  
'time_indices': '47007:47011:1', 'start_idx': 47007}
```

Based on the provided data, answer the following question:
Does this data support the provided meteorological claim?
Answer with True or False.

Claim: A very strong jet in excess of 150 knots will be across the east central U.S. with a favorable left exit region of the jet over New England.

Positive Claim Type

Original Claim: A very strong jet in excess of 150 knots will be across the east central U.S. with a favorable left exit region of the jet over New England.

The following data shows meteorological conditions over a 24-hour period:

```
{'type': 'wb2', 'variables': ['mean_sea_level_pressure',  
'10m_u_component_of_wind',  
'10m_v_component_of_wind', '2m_temperature',  
'geopotential', 'specific_humidity', 'temperature',  
'u_component_of_wind', 'v_component_of_wind'],  
'time_indices': '79179:79183:1', 'start_idx': 79179}
```

Based on the provided data, answer the following question:
Does this data support the provided meteorological claim?
Answer with True or False.

Claim: Dry conditions with minimal precipitation are expected across western Washington and into British Columbia.

Negative Claim Type

Original Claim: Heavy precipitation is expected across western Washington and into British Columbia.

Figure 6: (left) Positive claim and (right) negative example for meteorological claim verification

A.3 EXAMPLE FROM THE DATASET

The following data shows a snapshot of the global weather fields.
{data}
Based on the above data, answer the following question:
Which {geofeature} experienced the {extremum_direction} average {variable}?"Based on the provided data, {answer} experienced the {extremum_direction} average {variable} over the specified time-period, with an average {variable} of {answer_numeric}."

Example Template

The following data shows a snapshot of the global weather fields.
{'type': 'wb2', 'variables': ['mean_sea_level_pressure',
'10m_u_component_of_wind', '10m_v_component_of_wind',
'2m_temperature', 'geopotential', 'specific_humidity',
'temperature', 'u_component_of_wind', 'v_component_of_wind'],
'time_indices': '54746:54747:1'}

Based on the above data, answer the following question: Which continent experienced the highest average Surface temperature?

Based on the provided data, Africa experienced the highest average Surface temperature over the specified time-period, with an average Surface temperature of 303.5 K.

Generated Sample

Figure 7: (left) Example template from which samples are generated and (right) a sample generated using the template.

A.4 DEFINITION OF TASK CORRECTNESS

Different task types in ZEPHYRUSBENCH are evaluated using relevant metrics. To create a unified definition of correctness, we employ the following requirements for each metric type:

- **Numerical:** Standardized difference $\frac{|\hat{y} - y|}{\sigma} < 0.05$, where σ is the standard deviation of the relevant task variable in the WeatherBench2 dataset.

918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

- **Distance/Area/Coordinate/Simulation:** Relative error $\frac{|\hat{y} - y|}{|y|} < 0.05$. For true values of 0, we require $|\hat{y}| < 0.05$.
- **Location:** Exact locations string match, using fuzzy string matching logic.
- **Extreme Weather/Anomaly:** Earth Mover’s Distance (EMD) score < 100 km. If both true and predicted values are empty lists, the answer is considered correct.
- **Boolean:** Exact match between model answer and ground truth boolean value.
- **Discussion:** Overall discussion score > 0.5 .
- **Time:** Exact match required (absolute error = 0.0).

A.5 PERFORMANCE BY DIFFICULTY LEVEL

Below, we include a detailed breakdown of performance metrics by question difficulty level, as defined in Table A.2, for models GPT-5-Mini, GPT-5-Nano, Gemini 2.5 Flash, and gpt-oss-120b.




Figure 8: Questions correct by difficulty level: Easy.




Figure 9: Questions correct by difficulty level: Medium.




Figure 10: Questions correct by difficulty level: Hard.

Model	LLM	AE (Q25) (↓)	AE (Q50) (↓)	AE (Q75) (↓)	AE (Q99) (↓)
ZEPHYRUS-REFLECTIVE	gpt-5-mini	0.00	0.00	18.0	156.0
ZEPHYRUS-DIRECT	gpt-5-mini	0.00	0.00	18.0	157.0
Text Only LLM	gpt-5-mini	12.0	30.0	66.0	168.0
ZEPHYRUS-REFLECTIVE	gpt-5-nano	0.00	0.00	12.0	158,000
ZEPHYRUS-DIRECT	gpt-5-nano	0.00	0.00	12.0	5.01e+18
Text Only LLM	gpt-5-nano	18.0	48.0	90.0	186.0
ZEPHYRUS-REFLECTIVE	gemini-2.5-flash	0.00	0.00	18.0	157.0
ZEPHYRUS-DIRECT	gemini-2.5-flash	0.00	0.00	30.0	8.57e+18
Text Only LLM	gemini-2.5-flash	18.0	36.0	72.0	186.0
ZEPHYRUS-REFLECTIVE	gpt-oss-120b	0.00	0.00	0.00	145.0
ZEPHYRUS-DIRECT	gpt-oss-120b	0.00	0.00	6.00	145.0
Text Only LLM	gpt-oss-120b	18.0	42.0	84.0	200.0

Table 4: Absolute error quantiles for time tasks, in units of hours.

Model	LLM	Location Accuracy (%) (↑)	EMD (km) (↓)	Extreme Weather F1 (↑)
ZEPHYRUS-REFLECTIVE	gpt-5-mini	86.6	1,851	0.38
ZEPHYRUS-DIRECT	gpt-5-mini	80.9	1,892	0.28
Text Only LLM	gpt-5-mini	16.3	5,783	0.38
ZEPHYRUS-REFLECTIVE	gpt-5-nano	68.9	2,568	0.36
ZEPHYRUS-DIRECT	gpt-5-nano	73.7	2,126	0.20
Text Only LLM	gpt-5-nano	15.3	5,070	0.00
ZEPHYRUS-REFLECTIVE	gemini-2.5-flash	77.5	2,021	0.38
ZEPHYRUS-DIRECT	gemini-2.5-flash	76.6	2,204	0.38
Text Only LLM	gemini-2.5-flash	7.66	2,303	0.03
ZEPHYRUS-REFLECTIVE	gpt-oss-120b	77.0	2,749	0.49
ZEPHYRUS-DIRECT	gpt-oss-120b	61.2	2,435	0.41
Text Only LLM	gpt-oss-120b	11.5	3,718	0.00

Table 5: Location metrics for location answer-based questions. EMD stands for Earth mover’s Distance.

Model	LLM	% Valid Outputs (↑)	Discussion Score (↑)	Boolean F1 (↑)
ZEPHYRUS-REFLECTIVE	gpt-5-mini	92.9	0.18	0.51
ZEPHYRUS-DIRECT	gpt-5-mini	91.5	0.06	0.32
Text Only LLM	gpt-5-mini	91.9	0.07	0.53
ZEPHYRUS-REFLECTIVE	gpt-5-nano	88.8	0.14	0.48
ZEPHYRUS-DIRECT	gpt-5-nano	91.3	0.07	0.47
Text Only LLM	gpt-5-nano	88.9	0.07	0.37
ZEPHYRUS-REFLECTIVE	gemini-2.5-flash	91.0	0.09	0.49
ZEPHYRUS-DIRECT	gemini-2.5-flash	87.0	0.06	0.52
Text Only LLM	gemini-2.5-flash	71.9	0.02	0.16
ZEPHYRUS-REFLECTIVE	gpt-oss-120b	90.5	0.09	0.46
ZEPHYRUS-DIRECT	gpt-oss-120b	86.8	0.05	0.47
Text Only LLM	gpt-oss-120b	75.6	0.03	0.18

Table 6: Overall percentage of valid outputs, numerical score (0-1) for discussion questions, and F1 score for boolean questions.

A.6 DETAILED PERFORMANCE METRICS

We include detailed performance metrics from running several LLMs across all three modes on the entire ZEPHYRUSBENCHdataset.

1070
1071
1072
1073
1074
1075
1076
1077
1078
1079

1080 A.7 PERFORMANCE METRICS BY TASK
1081

Model	LLM	SAE (Q25) (↓)	SAE (Q50) (↓)	SAE (Q75) (↓)	SAE (Q99) (↓)
ZEPHYRUS-REFLECTIVE	gpt-5-mini	0.00	0.00	0.00	0.07
ZEPHYRUS-DIRECT	gpt-5-mini	0.00	0.00	0.00	0.17
Text Only LLM	gpt-5-mini	0.23	0.80	1.53	8.93
ZEPHYRUS-REFLECTIVE	gpt-5-nano	0.00	0.00	0.04	0.82
ZEPHYRUS-DIRECT	gpt-5-nano	0.00	0.00	0.00	5.76
Text Only LLM	gpt-5-nano	0.21	0.62	1.65	586.7
ZEPHYRUS-REFLECTIVE	gemini-2.5-flash	0.00	0.00	0.00	1.24
ZEPHYRUS-DIRECT	gemini-2.5-flash	0.00	0.00	0.00	0.69
Text Only LLM	gemini-2.5-flash	0.30	1.86	3.43	77.2
ZEPHYRUS-REFLECTIVE	gpt-oss-120b	0.00	0.00	0.00	9.23
ZEPHYRUS-DIRECT	gpt-oss-120b	0.00	0.00	0.00	0.69
Text Only LLM	gpt-oss-120b	0.17	0.63	1.72	25.1

1095 Table 7: Standardized Absolute Error (SAE) quantiles for Template ID 2: What is the
1096 min/max/average/median value of a weather variable at a specific
1097 location

Model	LLM	SAE (Q25) (↓)	SAE (Q50) (↓)	SAE (Q75) (↓)	SAE (Q99) (↓)
ZEPHYRUS-REFLECTIVE	gpt-5-mini	0.00	0.00	0.00	0.28
ZEPHYRUS-DIRECT	gpt-5-mini	0.00	0.00	0.00	0.27
Text Only LLM	gpt-5-mini	0.20	0.54	1.18	35.2
ZEPHYRUS-REFLECTIVE	gpt-5-nano	0.00	0.00	0.01	1.37
ZEPHYRUS-DIRECT	gpt-5-nano	0.00	0.00	0.00	0.96
Text Only LLM	gpt-5-nano	0.20	0.53	1.46	320.2
ZEPHYRUS-REFLECTIVE	gemini-2.5-flash	0.00	0.00	0.00	0.44
ZEPHYRUS-DIRECT	gemini-2.5-flash	0.00	0.00	0.00	0.44
Text Only LLM	gemini-2.5-flash	0.44	1.46	14.7	1.34e+08
ZEPHYRUS-REFLECTIVE	gpt-oss-120b	0.00	0.00	0.00	0.56
ZEPHYRUS-DIRECT	gpt-oss-120b	0.00	0.00	0.01	0.75
Text Only LLM	gpt-oss-120b	0.26	0.91	2.29	3,360

1111 Table 8: Standardized Absolute Error (SAE) quantiles for Template ID 5: What is the
1112 weather variable value at a location at a specific time
1113

Model	LLM	SAE (Q25) (↓)	SAE (Q50) (↓)	SAE (Q75) (↓)	SAE (Q99) (↓)
ZEPHYRUS-REFLECTIVE	gpt-5-mini	0.02	0.06	0.11	0.60
ZEPHYRUS-DIRECT	gpt-5-mini	0.02	0.06	0.11	0.48
Text Only LLM	gpt-5-mini	0.17	0.62	1.28	8.76
ZEPHYRUS-REFLECTIVE	gpt-5-nano	0.03	0.08	0.26	2.06
ZEPHYRUS-DIRECT	gpt-5-nano	0.02	0.07	0.17	1.53
Text Only LLM	gpt-5-nano	0.17	0.47	1.19	32,886
ZEPHYRUS-REFLECTIVE	gemini-2.5-flash	0.02	0.06	0.13	1.12
ZEPHYRUS-DIRECT	gemini-2.5-flash	0.02	0.06	0.10	0.57
Text Only LLM	gemini-2.5-flash	0.39	0.93	2.30	56.2
ZEPHYRUS-REFLECTIVE	gpt-oss-120b	0.02	0.06	0.12	0.60
ZEPHYRUS-DIRECT	gpt-oss-120b	0.02	0.06	0.11	0.78
Text Only LLM	gpt-oss-120b	0.17	0.89	2.34	1,021

1127 Table 9: Standardized Absolute Error (SAE) quantiles for Template ID 6: What will the
1128 variable be at a location after time interval (forecast)
11291130
1131
1132
1133

1134

1135

Model	LLM	SAE (Q25) (↓)	SAE (Q50) (↓)	SAE (Q75) (↓)	SAE (Q99) (↓)
ZEPHYRUS-REFLECTIVE	gpt-5-mini	0.00	0.00	0.02	0.11
ZEPHYRUS-DIRECT	gpt-5-mini	0.00	0.00	0.03	0.13
Text Only LLM	gpt-5-mini	0.00	0.07	0.12	0.23
ZEPHYRUS-REFLECTIVE	gpt-5-nano	0.00	0.01	0.08	0.26
ZEPHYRUS-DIRECT	gpt-5-nano	0.00	0.00	0.04	0.61
Text Only LLM	gpt-5-nano	0.00	0.07	0.12	0.23
ZEPHYRUS-REFLECTIVE	gemini-2.5-flash	0.00	0.00	0.00	0.10
ZEPHYRUS-DIRECT	gemini-2.5-flash	0.00	0.00	0.01	0.13
Text Only LLM	gemini-2.5-flash	0.00	0.07	0.12	0.23
ZEPHYRUS-REFLECTIVE	gpt-oss-120b	0.00	0.01	0.03	0.33
ZEPHYRUS-DIRECT	gpt-oss-120b	0.00	0.01	0.04	0.84
Text Only LLM	gpt-oss-120b	0.00	0.07	0.12	0.23

Table 10: Standardized Absolute Error (SAE) quantiles for Template ID 44: How will weather variable change after specified time with specified change in variable (counterfactual)

1152

1153

Model	LLM	SAE (Q25) (↓)	SAE (Q50) (↓)	SAE (Q75) (↓)	SAE (Q99) (↓)
ZEPHYRUS-REFLECTIVE	gpt-5-mini	0.11	0.37	0.67	0.99
ZEPHYRUS-DIRECT	gpt-5-mini	0.13	0.33	0.66	1.06
Text Only LLM	gpt-5-mini	0.16	0.29	0.41	0.81
ZEPHYRUS-REFLECTIVE	gpt-5-nano	0.18	0.47	0.68	0.94
ZEPHYRUS-DIRECT	gpt-5-nano	0.17	0.36	0.55	0.99
Text Only LLM	gpt-5-nano	0.16	0.29	0.41	0.82
ZEPHYRUS-REFLECTIVE	gemini-2.5-flash	0.01	0.06	0.31	0.92
ZEPHYRUS-DIRECT	gemini-2.5-flash	0.12	0.35	0.66	0.94
Text Only LLM	gemini-2.5-flash	0.16	0.29	0.40	0.94
ZEPHYRUS-REFLECTIVE	gpt-oss-120b	0.20	0.31	0.41	0.77
ZEPHYRUS-DIRECT	gpt-oss-120b	0.16	0.30	0.40	0.83
Text Only LLM	gpt-oss-120b	0.16	0.29	0.40	0.68

Table 11: Absolute Error (AE) quantiles for Template ID 45: What is the value of the input parameter of the simulator model that produces the simulation output

1170

1171

1172

Model	LLM	AE (Q25) (↓)	AE (Q50) (↓)	AE (Q75) (↓)	AE (Q99) (↓)
ZEPHYRUS-REFLECTIVE	gpt-5-mini	0.00	0.00	0.00	101.0
ZEPHYRUS-DIRECT	gpt-5-mini	0.00	0.00	6.00	122.6
Text Only LLM	gpt-5-mini	12.0	18.0	36.0	130.1
ZEPHYRUS-REFLECTIVE	gpt-5-nano	0.00	0.00	0.00	60.0
ZEPHYRUS-DIRECT	gpt-5-nano	0.00	0.00	0.00	6.87e+18
Text Only LLM	gpt-5-nano	12.0	30.0	72.0	144.0
ZEPHYRUS-REFLECTIVE	gemini-2.5-flash	0.00	0.00	0.00	134.2
ZEPHYRUS-DIRECT	gemini-2.5-flash	0.00	0.00	0.00	9.18e+18
Text Only LLM	gemini-2.5-flash	12.0	24.0	48.0	138.8
ZEPHYRUS-REFLECTIVE	gpt-oss-120b	0.00	0.00	0.00	35.3
ZEPHYRUS-DIRECT	gpt-oss-120b	0.00	0.00	0.00	136.6
Text Only LLM	gpt-oss-120b	12.0	30.0	66.0	162.0

Table 12: Absolute Error (AE) quantiles for Template ID 4: How many hours from start did a location experience extremum

1188

1189

1190

1191

1192

1193

1194

1195

1196

1197

1198

1199

1200

1201

1202

Table 13: Absolute Error (AE) quantiles for Template ID 7: When will location experience its extremum in future period (forecast)

1203

1204

1205

1206

Model	LLM	Location Accuracy (%) (\uparrow)	EMD Score (km) (\downarrow)
ZEPHYRUS-REFLECTIVE	gpt-5-mini	0.00	18.0
ZEPHYRUS-DIRECT	gpt-5-mini	0.00	24.0
Text Only LLM	gpt-5-mini	48.0	72.0
ZEPHYRUS-REFLECTIVE	gpt-5-nano	6.00	18.0
ZEPHYRUS-DIRECT	gpt-5-nano	1.50	18.0
Text Only LLM	gpt-5-nano	54.0	84.0
ZEPHYRUS-REFLECTIVE	gemini-2.5-flash	0.00	18.0
ZEPHYRUS-DIRECT	gemini-2.5-flash	0.00	21.0
Text Only LLM	gemini-2.5-flash	43.5	69.0
ZEPHYRUS-REFLECTIVE	gpt-oss-120b	0.00	9.00
ZEPHYRUS-DIRECT	gpt-oss-120b	0.00	6.00
Text Only LLM	gpt-oss-120b	24.0	66.0
			24.0
			126.0
			341,254

Table 14: Location prediction metrics for Template ID 1: Which geographic feature experienced the highest/lowest average value of a weather variable

1224

1225

1226

1227

1228

1229

1230

1231

1232

1233

1234

1235

1236

1237

1238

1239

Table 15: Location prediction metrics for Template ID 3: Which sublocation has the highest/lowest recorded variable value

Model	LLM	Correctness (%) (↑)
ZEPHYRUS-REFLECTIVE	gpt-5-mini	41.2
ZEPHYRUS-DIRECT	gpt-5-mini	45.6
Text Only LLM	gpt-5-mini	36.8
ZEPHYRUS-REFLECTIVE	gpt-5-nano	48.5
ZEPHYRUS-DIRECT	gpt-5-nano	60.3
Text Only LLM	gpt-5-nano	64.7
ZEPHYRUS-REFLECTIVE	gemini-2.5-flash	39.7
ZEPHYRUS-DIRECT	gemini-2.5-flash	25.0
Text Only LLM	gemini-2.5-flash	58.8
ZEPHYRUS-REFLECTIVE	gpt-oss-120b	33.8
ZEPHYRUS-DIRECT	gpt-oss-120b	32.4
Text Only LLM	gpt-oss-120b	64.7

Table 16: Correctness metrics for Template ID 8: Identify extreme weather events that will occur in the next N hours (forecast)

Model	LLM	Correctness (%) (↑)
ZEPHYRUS-REFLECTIVE	gpt-5-mini	14.6
ZEPHYRUS-DIRECT	gpt-5-mini	43.8
Text Only LLM	gpt-5-mini	57.3
ZEPHYRUS-REFLECTIVE	gpt-5-nano	44.8
ZEPHYRUS-DIRECT	gpt-5-nano	53.1
Text Only LLM	gpt-5-nano	70.8
ZEPHYRUS-REFLECTIVE	gemini-2.5-flash	33.3
ZEPHYRUS-DIRECT	gemini-2.5-flash	27.1
Text Only LLM	gemini-2.5-flash	64.6
ZEPHYRUS-REFLECTIVE	gpt-oss-120b	40.6
ZEPHYRUS-DIRECT	gpt-oss-120b	28.1
Text Only LLM	gpt-oss-120b	71.9

Table 17: Correctness metrics for Template ID 9: Check if extreme weather events are currently happening

Model	LLM	Correctness (%) (↑)
ZEPHYRUS-REFLECTIVE	gpt-5-mini	5.10
ZEPHYRUS-DIRECT	gpt-5-mini	3.20
Text Only LLM	gpt-5-mini	3.80
ZEPHYRUS-REFLECTIVE	gpt-5-nano	5.40
ZEPHYRUS-DIRECT	gpt-5-nano	1.70
Text Only LLM	gpt-5-nano	12.5
ZEPHYRUS-REFLECTIVE	gemini-2.5-flash	13.1
ZEPHYRUS-DIRECT	gemini-2.5-flash	3.20
Text Only LLM	gemini-2.5-flash	5.00
ZEPHYRUS-REFLECTIVE	gpt-oss-120b	3.10
ZEPHYRUS-DIRECT	gpt-oss-120b	1.60
Text Only LLM	gpt-oss-120b	11.1

Table 18: Correctness metrics for Template ID 10: Which geographic features experienced unusual weather anomalies compared to baseline

1296

1297

1298

1299

1300

1301

1302

1303

1304

1305

1306

1307

1308

1309

1310

1311

1312

1313

1314

1315

1316

1317

1318

1319

1320

1321

1322

1323

1324

1325

1326

1327

1328

1329

1330

1331

1332

1333

1334

1335

1336

1337

1338

1339

1340

1341

1342

1343

1344

1345

1346

1347

Model	LLM	Discussion Score (Mean) (↑)	Discussion Score (Median) (↑)
ZEPHYRUS-REFLECTIVE	gpt-5-mini	0.02	0.00
ZEPHYRUS-DIRECT	gpt-5-mini	0.00	0.00
Text Only LLM	gpt-5-mini	0.00	0.00
ZEPHYRUS-REFLECTIVE	gpt-5-nano	0.00	0.00
ZEPHYRUS-DIRECT	gpt-5-nano	0.00	0.00
Text Only LLM	gpt-5-nano	0.00	0.00
ZEPHYRUS-REFLECTIVE	gemini-2.5-flash	0.00	0.00
ZEPHYRUS-DIRECT	gemini-2.5-flash	0.00	0.00
Text Only LLM	gemini-2.5-flash	0.00	0.00
ZEPHYRUS-REFLECTIVE	gpt-oss-120b	0.00	0.00
ZEPHYRUS-DIRECT	gpt-oss-120b	0.00	0.00
Text Only LLM	gpt-oss-120b	0.00	0.00

Table 19: Discussion score metrics for Template ID 41: Generate comprehensive global climate forecast for temperature and precipitation for next 3 months (forecast)

Model	LLM	Discussion Score (Mean) (↑)	Discussion Score (Median) (↑)
ZEPHYRUS-REFLECTIVE	gpt-5-mini	0.31	0.32
ZEPHYRUS-DIRECT	gpt-5-mini	0.09	0.05
Text Only LLM	gpt-5-mini	0.10	0.04
ZEPHYRUS-REFLECTIVE	gpt-5-nano	0.23	0.21
ZEPHYRUS-DIRECT	gpt-5-nano	0.10	0.06
Text Only LLM	gpt-5-nano	0.09	0.07
ZEPHYRUS-REFLECTIVE	gemini-2.5-flash	0.15	0.11
ZEPHYRUS-DIRECT	gemini-2.5-flash	0.10	0.07
Text Only LLM	gemini-2.5-flash	0.02	0.00
ZEPHYRUS-REFLECTIVE	gpt-oss-120b	0.16	0.14
ZEPHYRUS-DIRECT	gpt-oss-120b	0.10	0.04
Text Only LLM	gpt-oss-120b	0.02	0.00

Table 20: Discussion score metrics for Template ID 42: Provide detailed meteorological discussion and forecast for continental United States (forecast)

Model	LLM	Discussion Score (Mean) (↑)	Discussion Score (Median) (↑)
ZEPHYRUS-REFLECTIVE	gpt-5-mini	0.26	0.22
ZEPHYRUS-DIRECT	gpt-5-mini	0.15	0.14
Text Only LLM	gpt-5-mini	0.17	0.08
ZEPHYRUS-REFLECTIVE	gpt-5-nano	0.23	0.19
ZEPHYRUS-DIRECT	gpt-5-nano	0.18	0.16
Text Only LLM	gpt-5-nano	0.20	0.18
ZEPHYRUS-REFLECTIVE	gemini-2.5-flash	0.18	0.16
ZEPHYRUS-DIRECT	gemini-2.5-flash	0.09	0.00
Text Only LLM	gemini-2.5-flash	0.07	0.00
ZEPHYRUS-REFLECTIVE	gpt-oss-120b	0.16	0.11
ZEPHYRUS-DIRECT	gpt-oss-120b	0.08	0.05
Text Only LLM	gpt-oss-120b	0.14	0.03

Table 21: Discussion score metrics for Template ID 43: Generate ENSO climate update and outlook based on atmospheric data (forecast)

Model	LLM	F1 Score (↑)	Precision (%) (↑)	Recall (%) (↑)
ZEPHYRUS-REFLECTIVE	gpt-5-mini	0.51	65.2	41.5
ZEPHYRUS-DIRECT	gpt-5-mini	0.32	49.4	23.7
Text Only LLM	gpt-5-mini	0.53	63.0	45.5
ZEPHYRUS-REFLECTIVE	gpt-5-nano	0.48	63.2	38.3
ZEPHYRUS-DIRECT	gpt-5-nano	0.47	63.7	37.6
Text Only LLM	gpt-5-nano	0.37	59.5	26.7
ZEPHYRUS-REFLECTIVE	gemini-2.5-flash	0.49	62.3	40.8
ZEPHYRUS-DIRECT	gemini-2.5-flash	0.52	66.0	42.2
Text Only LLM	gemini-2.5-flash	0.16	47.1	9.47
ZEPHYRUS-REFLECTIVE	gpt-oss-120b	0.46	63.0	36.8
ZEPHYRUS-DIRECT	gpt-oss-120b	0.47	61.9	37.4
Text Only LLM	gpt-oss-120b	0.18	54.5	10.7

Table 22: Boolean score metrics for Template ID 46: Check whether the given claim extracted from meterological report is supported by the data

A.8 MODEL PROMPTS

We use the following core Instruction prompt for ZEPHYRUS-REFLECTIVE:

1404

Zephyrus-Reflective Instruction Prompt

1405

1406

1407

You are an AI weather expert agent. You will use an interactive coding environment with
 ↳ tool functions, data, and softwares to solve the user's task.

1408

1409

1410

At each turn, you should first provide your thinking and reasoning given the
 ↳ conversation history (which might include output from executed code within
 ↳ <observation></observation>).

After that, you must do exactly one of the following:

1411

1) Write code based on problem and/or observation. Your code should be enclosed using
 ↳ "<execute>" tag, for example: <execute> return "Hello World!" </execute>. IMPORTANT:
 ↳ You must end the code block with </execute> tag.

1412

2) When you think you have a solution ready, directly provide a solution that adheres to
 ↳ the required format for the given task to the user.

1413

Your solution should be enclosed using "<solution>" tag, for example: The answer is
 ↳ <solution> A </solution>. IMPORTANT: You must end the solution block with </solution>
 ↳ tag. When answering numerical questions, always use SI base unit (standard units of
 ↳ measurement) unless the problem specifically asks for a certain unit. For example,
 ↳ some questions may require you to answer in hours. Enclose ONLY the final answer to
 ↳ the question in these tags, do NOT include any other information.

1414

In each response, you must include <execute> or <solution> tag. Not both at the same
 ↳ time. Do not generate code outside <execute>. Do not output answers outside
 ↳ <solution>. Do not respond with messages without any tags. No empty messages.

1415

- Geolocator Documentation:

1416

The detailed documentation for the Geolocator class, including its available methods, is
 ↳ provided below:

1417

{geolocator_documentation}

1418

- Forecaster API Documentation:

1419

{forecaster_documentation}

1420

The Forecaster can reliably forecast at most 2 weeks into the future.

1421

- IMPORTANT: If the question is about the future, you **will need to** use the Forecaster
 ↳ object to answer the question and solve the task.

1422

The input data **will not** contain the answer to questions about the
 ↳ future.

1423

- Simulator API Documentation:

1424

{simulator_documentation}

1425

- The Simulator provides atmospheric modeling and can be used for climate simulations,
 ↳ answering counterfactuals, sensitivity studies, or generating synthetic weather
 ↳ data.

1426

- The Simulator can handle extended time periods (months to years) in a SINGLE call. DO
 ↳ NOT create loops or multiple simulator instances. Set total_time to the desired
 ↳ duration and call simulate() once.

1427

- Variable Descriptions:

1428

A comprehensive description of every variable contained in the xarray datasets is given
 ↳ here:

1429

{var_desc}

1430

- Dataset Keys Explanation:

1431

An explanation of what each key in the datasets represents is provided below:
 ↳ {keys}

1432

... (continued)

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

Zephyrus-Reflective Instruction Prompt (cont.)

1458 (continued) ...

1459 - Units:

1460 Always use the following SI units when reasoning and coding:

1461 {units_desc}

1462 Answer in SI units unless the problem specifically asks for a different unit. For
1463 → example, some questions may require hours.

1464 - Time Indices:

1465 You should NOT slice the provided dataset according to the time indices. The datasets
1466 → are already sliced to the correct time indices.

1467 For any question that asks about the time offset, only provide the time indices relative
1468 → to the provided dataset.

1469 If the question asks for the time offset, return the answer in hours from the initial
1470 → time index.

1471 For example, if the question asks about a dataset with time interval 6 hours and time
1472 → indices 12345:12351:1, and you think the answer is index 12350, you should return 30
1473 → hours.

1474 Do NOT return the time index as a timestamp or datetime object.

1475 ****Execution code requirements:****

1476 - The code MUST all be defined with a function called `run`.

1477 - The `run` function should accept four parameters:

1478 a. A list of one or more xarray datasets.

1479 b. A Geolocator object (which comes with a set of predefined helpful functions).

1480 c. A Forecaster object (which comes with a set of predefined helpful functions).

1481 c. A Simulator object (which comes with a set of predefined helpful functions).

1482 - DO NOT write any code outside of the `run` function.

1483 ****IMPORTANT:****

1484 - The Geolocator object is already constructed and passed in as `geolocator`.

1485 - ****Never open files, use `xr.open_dataset`**, or import Geolocator.**

1486 - If you are subsetting, make sure to subset carefully considering runtime. It is too
1487 → slow to select the entire xarray dataset. If you are subsetting over multiple
1488 → dimensions (e.g. spatially and temporally), make sure to apply the smaller subset
1489 → operation first.

1490 - By following these detailed instructions, your code should clearly use the provided
1491 → datasets and tools to produce the correct result.

1492 - Coordinate System:

1493 The WeatherBench2 (WB2) dataset uses an equiangular grid with the following
1494 → specifications:

1495 - Latitude: 121 grid points ranging from -90° to $+90^{\circ}$ in 1.5° increments

1496 - Longitude: 240 grid points ranging from 0° to 358.5° in 1.5° increments

1497 - The latitude coordinates are: $[-90, -88.5, -87, \dots, 87, 88.5, 90]$

1498 - The longitude coordinates are: $[0, 1.5, 3, \dots, 355.5, 357, 358.5]$

1499 ****Other Requirements:****

1500 - Under NO circumstances should you loop over the grid points (i.e. you should NOT loop
1501 → over latitudes and longitudes), but rather try to leverage vectorized operations,
1502 → built-in functions or the Geolocator class as appropriate. This is a key requirement.
1503 → DO NOT loop over the latitudes and longitudes ANYWHERE in your generated code.

1504 - Ensure that you call and use the functions from the Geolocator object correctly as per
1505 → its documentation.

1506 ****Question:****

1507 {question}

For the reflective stage of ZEPHYRUS-REFLECTIVE, we use the following Observation prompt:

Zephyrus-Reflective Observation Prompt

1505 The executed code produced the output above. Reason about your next step and either (1)
1506 → output the final result based on this observation. Enclose your answer in
1507 → <solution></solution> tags., or (2) generate another code block to execute. Enclose
1508 → your code in <execute></execute> tags.
1509 If you choose to give a solution, enclose ONLY the final answer to the question in these
1510 → tags, do NOT include any other information.
1511 You should execute code if you think you need more information before providing a final
1512 → answer.

For ZEPHYRUS-DIRECT, we use the following direct Instruction prompt:

1512
1513**Zephyrus-Direct Instruction Prompt**1514
1515
1516

Your objective is to write a Python function called 'run' that solves a specified
 ↪ problem using provided data and Toolset APIs. The function should be designed
 ↪ according to the following guidelines:

1517

1. Function Definition:

1518

- The function must be named run.
- It should accept four parameters:
 - a. A list of one or more xarray datasets.
 - b. A Geolocator object (which comes with a set of predefined helpful functions).
 - c. A Forecaster object (which comes with a set of predefined helpful functions).
 - c. A Simulator object (which comes with a set of predefined helpful functions).

1523

2. Data Descriptions:

1524

- Variable Descriptions:

1525

A comprehensive description of every variable contained in the xarray datasets is given
 ↪ here:
 {var_desc}

1528

- Dataset Keys Explanation:

1529

An explanation of what each key in the datasets represents is provided below:
 {keys}

1531

- Units:

1533

Always use the following SI units when reasoning and coding:
 {units_desc}

1535

- Time Indices:

1536

The datasets provided have been converted from using a time dimension to simple integer
 ↪ indices starting from 0. Each index step represents 6 hours of time in the original
 ↪ dataset.

1538

You should NOT slice the provided dataset according to the provided indices. The
 ↪ datasets are already sliced to the correct indices.

1540

For any question that asks about the time offset, only provide the time indices relative
 ↪ to the provided dataset.

1541

If the question asks for the time offset, you should return the answer in hours from the
 ↪ initial time index.

1542

For example, if the question asks about a dataset with time interval 6 hours and indices
 ↪ 0:6:1, and you think the answer is index 5, you should return 30 hours.

Do NOT return the index directly.

1544

3. Toolset APIs

1545

You are given access to the following code tools. Please use them as needed inside your
 ↪ 'run' function:

1548

- Geolocator Documentation:

1549

The detailed documentation for the Geolocator class, including its available methods, is
 ↪ provided below:

1551

{geolocator_documentation}

1553

1554

- Forecaster API Documentation:

1555

{forecaster_documentation}

The Forecaster can reliably forecast at most 2 weeks into the future.

1556

- IMPORTANT: If the question is about the future, you **will** need to **use** the Forecaster
 ↪ object to answer the question and solve the task.

1557

The input data **will not** contain the answer to questions about the
 ↪ future.

1559

1560

... (continued)

1561

1562

1563

1564

1565

1566

Zephyrus-Direct Instruction Prompt (cont.)

1567

1568

(continued) ...

1569

1570

- Simulator API Documentation:

1571

{simulator_documentation}

1572

- The Simulator provides atmospheric modeling and can be used for climate simulations, answering counterfactuals, sensitivity studies, or generating synthetic weather data.
- The Simulator can handle extended time periods (months to years) in a SINGLE call. DO NOT create loops or multiple simulator instances. Set total_time to the desired duration and call simulate() once.

1577

1578

4. Task Details:

1579

- The function should process the datasets using the pertinent variables as specified within the question.
- Under NO circumstances should you loop over the grid points (i.e. you should NOT loop over latitudes and longitudes), but rather try to leverage vectorized operations, built-in functions or the Geolocator class as appropriate. This is a key requirement.
- DO NOT loop over the latitudes and longitudes ANYWHERE in your generated code.
- Ensure that you call and use the functions from the Geolocator object correctly as per its documentation.

1584

5. Returning the Answer:

1585

- The final result should be returned by the function.
- Make sure to encapsulate your run function in triple backticks for clarity. For example:

```
...
def run(...):
    return "Hello"
...
```

1591

- If the answer is a time value, make sure to return it in a unit of time rather than as a timestamp or datetime object. For example, return '5 hours' instead of '2022-01-01 05:00:00'.
- Always return the answer in the same unit as the one used in the weatherbench dataset.
- Do not convert any units.

1594

6. Problem Statement:

1595

By following these detailed instructions, your code should clearly use the provided datasets and the Toolset APIs to produce the correct result. The specific question that your function needs to answer is provided at the end of this prompt: {question}

1599

1600

1601

1602

1603

1604

1605

1606

1607

1608

1609

1610

1611

1612

1613

1614

1615

1616

1617

1618

1619