
Where Do Agents Differ? Interpretable Rule Discovery for Performance Differences Across Models and Data

Sascha Xu¹ Antoine Gautier² Jilles Vreeken¹

Abstract

Agentic systems are trained and evaluated with a variety of backend models over a heterogeneous collection of tasks. Typically, evaluation focuses on aggregate metrics over predefined categories and considers models in isolation. In this work, we aim to identify characteristics of the input space *where* the performance between two configurations, e.g. different backend models, differs substantially. To this end, we discover rules interpretable rules that describe task regimes with pronounced distributional differences. We demonstrate our approach in two agentic use cases: On `swe-bench`, we contrast a state-of-the-art coding agent across different backends, while on the `ChartQA` benchmark, we compare performance on synthetic and real data. Our results show that performance differences are highly structured and can reverse across task regimes, revealing insights that are not captured by standard evaluation.

1. Introduction

Evaluating agentic systems is challenging due to the diversity of tasks, data sources, and models involved. In practice, evaluation is typically performed using aggregate metrics, often computed over predefined subsets of data. Existing evaluation practices, such as model cards (Mitchell et al., 2019), focus on individual models and do not address comparative analysis across configurations. As a result, structured differences often remain hidden, such as when one system outperforms another on specific types of tasks.

Prior work has proposed error slicing as a method to identify regions of poor performance, supporting model debugging and validation (Chung et al., 2019; Sagadeeva & Boehm, 2021). These approaches however analyze a single model in isolation. In many practical settings however, evaluation

¹CISPA Helmholtz Center for Information Security ²QuantPi. Correspondence to: Sascha Xu <sascha.xu@cispa.de>.




Figure 1. Example rule discovered on `ChartQA`. For short questions with numeric answers over small tables, accuracy on generated questions is substantially higher than on human-written questions, revealing a systematic difference in task difficulty.

is inherently comparative. For example, when selecting between agentic backend models, for different kinds of tasks different backends are more effective than others. Similarly, when using synthetic data for training, we want to assess whether performance transfers to real-world data and for which tasks it doesn't.

In this work, we introduce a method that discovers interpretable rules describing *where two sources differ in performance*, by adopting a differential subgroup discovery approach (Xu & Vreeken, 2026). Rather than analyzing systems independently, we directly search for regions of the task space where performance distributions diverge and characterize these regions with concise, human-readable rules. We evaluate it in two complementary settings: on `swe-bench`, where we compare agent performance across backend models, and on `ChartQA`, where we analyze performance discrepancies between synthetic and real data.

We give an example of the type of insight our method provides in Figure 1 for the `ChartQA` benchmark for chart based question answering (Masry et al., 2022). Our method uncovers a subset of tasks defined by the rule "short questions with numeric answers over small tables" where the accuracy on synthetic generated questions is substantially higher than on human-written questions. An investigation into the specific questions in this subset, shown in Figure 2, reveals that these generated questions require only a direct lookup of a value from the chart, while the human-written questions require more complex reasoning. This uncovers a weakness of the data generating process and is not captured

by standard aggregate evaluation or error analysis tools. In general, our method reveals actionable insights, providing a practical tool for identifying when and why agent performance differs across tasks.

2. Method

Problem setup. We consider a collection of tasks, where each task is characterized by: (i) a feature vector $\mathbf{x} \in \mathbb{R}^d$ describing task properties, (ii) a performance outcome $y \in \mathcal{Y}$ (e.g., correctness or score), and (iii) a source indicator $a \in \{0, 1\}$ identifying the setting under which it was generated (e.g., model backend or data source). Given a collection of n tasks $\{(\mathbf{x}^{(i)}, y^{(i)}, a^{(i)})\}_{i=1}^n$, we seek to identify and characterize feature subspaces of X with a pronounced difference in performance y between source $a = 0$ and $a = 1$.

To model interpretable regions of the task space, we use rule-based descriptions over \mathbf{x} . In particular, we consider logical conjunctions $r : \mathbb{R}^d \rightarrow \{0, 1\}$ of interval-based conditions over individual features $x_j \in \mathbb{R}$ of the form

$$r(\mathbf{x}; S, \alpha, \beta) = \bigwedge_{j \in S} \mathbb{1}(\alpha_j < x_j < \beta_j)$$

where $S \subseteq [d]$ is a subset of feature on which conditions are imposed. To handle discrete features, we use one-hot encoding. A rule r thus describes task regimes that can be described using concise descriptions such as "Questions starting with 'which' about charts with 13-45 datapoints".

We base our approach on a method for differential subgroup discovery (Xu & Vreeken, 2026). It requires an objective function that quantifies the utility of a particular rule r , and which we define below for usage in the comparative evaluation setting.

Rule Quality. The primary measure for the quality of a rule is the observed empirical difference between the samples with $a^{(i)} = 0$ and $a^{(i)} = 1$ within the scope defined by r , i.e. $r(\mathbf{x}^{(i)}) = 1$. Hence, we contrast the rule-conditional performance distribution in population $a \in \{0, 1\}$


$$P_{a,r}(Y) := P(Y | A = a, r(X) = 1).$$

We estimate this distribution using all samples $\{y^{(i)} | r(\mathbf{x}^{(i)}) = 1, a^{(i)} = a\}$ covered by rule r . For example, for a binary performance measure $y \in \{0, 1\}$, the distribution $\hat{P}_{0,r}$ is estimated as the normalized counts of $y^{(i)} = 0$ and $y^{(i)} = 1$ under the condition of $a^{(i)} = 1$ and $r(\mathbf{x}^{(i)}) = 1$.

To quantify difference, we thus adopt a general distributional perspective and measure discrepancy using a divergence D . We use Jensen-Shannon divergence as it captures differences in full outcome distributions and remains well-defined even for non-binary performance metrics

$$D_{\text{JS}}(\hat{P}_{0,r}(Y), \hat{P}_{1,r}(Y)).$$


Generated questions



What was the highest expenditure on foreign military aid in 2009/10?

What was the unemployment rate in Sweden in 2018?

Human-written questions



What is the percentage of logistics division in 2013 by air?

What's the share of ethnic groups except Chinese and Indians?

Figure 2. Example question-chart pairs from ChartQA of the discovered rule in Figure 1.

The primary goal is to identify rules r where two sources exhibit substantially different performance distributions.

Support. To ensure that identified subgroups correspond to meaningful and statistically reliable task regimes, we require that they are sufficiently represented in both sources.

$$n_a(r) = \frac{|\{i | r(\mathbf{x}^{(i)}) = 1 \wedge a^{(i)} = a\}|}{|\{i | a^{(i)} = a\}|}$$

denotes the fraction of tasks from source a covered by rule r . To ensure balanced supports, we use the geometric mean from both sources $\sqrt{(n_0(r)n_1(r))}$ which penalizes low support in either source. We construct the combined objective as the product of Jensen-Shannon divergence and support. This ensures that we get rules both satisfy adequate support and divergence, as a rule that has only one property alone itself is not interesting, e.g. large coverage with no divergence. To find rules we optimize

$$\max_r \sqrt{(n_0(r)n_1(r))} \cdot D_{\text{JS}}(\hat{P}_{0,r}(Y), \hat{P}_{1,r}(Y)).$$

To optimize this objective, we use the DIFFSUB method (Xu & Vreeken, 2026), which uses gradient-based optimization to learn rules maximizing an objective based on the difference between two populations. We discover multiple rules through an iterative approach where we first learn a rule, then drop out all samples and re-run DIFFSUB for a total of three times. With the method established, we now turn to assessing its utility in agentic evaluation.




Figure 3. For tasks with long trajectories, we identify a rule where the OpenHands agent with the `claude-opus-4-5` backend outperforms the `gpt5` backend by 33.81%.

3. Applications

We now evaluate our method in two case studies of comparative analysis of agentic systems: (i) comparing performance across different backend models, and (ii) comparing performance across different data sources. In both settings, DIFFSUB is the primary analysis tool. For comparison, we also consult the Error Analysis Dashboard of the Microsoft Responsible AI (RAI) toolkit¹, which provides an interpretable exploration of model performance through a decision-tree based approach.

3.1. SWE-bench: Comparing Agentic Backends

First, we examine the performance difference of the same agent when using different backend language models. As benchmark, we use the SWE-bench benchmark (Jimenez et al., 2024), which consists of software engineering tasks derived from real-world GitHub issues. We evaluate the performance of the OpenHands agent (Wang et al., 2024), a SOTA system for this benchmark, across five different backend language models. We conduct pairwise comparisons between five backends with varying complexity and performance, using the results from the official SWE-bench repository. We consider the following backends: `claude-opus-4-5` (78% resolved rate), `gpt5` (72%), `claude-4-sonnet` (70%), `4x.scaled` (61%), and `devstral-small`(47%).

For each task, we combine benchmark metadata with runtime statistics. The task-side features include issue creation year, difficulty bucket, presence of hints, problem statement length, counts of tests to pass and pass-to-pass checks. The execution-side features include trajectory message count, number of terminal turns, and the size of the produced patch. This feature set lets us characterize behavior in terms of which tasks are hard and how the two backends behave while solving them.

¹<https://github.com/microsoft/responsible-ai-toolbox>




Figure 4. A second rule reveals a regime where `gpt5` achieves 88.89% while `claude-opus-4-5` drops to 36.36%.

Results. At the aggregate level, the five OpenHands backends span resolved rates from 47.66% to 77.91% (Figure 7). We focus on the comparison between `gpt5` and `claude-opus-4-5`, where the overall gap on the 497 shared tasks is 72.23% versus 77.87%. We show that this gap is far from uniform across the task space.

The rule with the largest positive gap covers 23.44% of the shared tasks and consists of runs with long trajectories (Figure 3). In this regime, `gpt5` as backend resolves only 55.88% of tasks, while `claude-opus-4-5` resolves 89.69%. This suggests the difference in performance is most pronounced in long, multi-turn tasks. Indeed, if the examples from this rule are removed, the aggregate comparison flips to 78.39% for `gpt5` versus 75.00% for `claude-opus-4-5`.

A second rule in Figure 4 reveals a regime where `gpt5` outperforms `claude-opus-4-5`: GitHub issues from [2016-2022] with hints, nominal difficulty 15 min - 1 hour and medium-sized patches, `gpt5` reaches 88.89% while `claude-opus-4-5` drops to 36.36%. On the other hand, the tree learned by the RAI Error Analysis Dashboard reveals task regimes where both models perform well respectively poorly. An analysis of each leaf in Appendix B.4 reveals that the maximum discrepancy between both backends only reaches up to 13.5% within the leaves. This demonstrates the added value of our comparative analysis

Pair	Overall	Max Pos+.Gap	Max Neg. Gap
opus vs gpt5	+5.63	+33.81 @ 23.4%	-52.53 @ 2.0%
opus vs sonnet	+7.24	+35.29 @ 12.6%	-2.50 @ 17.7%
opus vs 4x	+16.73	+94.12 @ 2.0%	-
opus vs devstral	+30.47	+64.47 @ 8.4%	-
gpt5 vs sonnet	+1.41	+14.28 @ 2.8%	-1.21 @ 28.0%
gpt5 vs 4x	+10.89	+45.80 @ 13.8%	-
gpt5 vs devstral	+24.74	+33.33 @ 2.5%	-2.86 @ 13.9%
sonnet vs 4x	+9.44	+12.71 @ 61.5%	-19.73 @ 6.5%
sonnet vs devstral	+23.22	+60.85 @ 15.7%	-49.26 @ 3.8%
4x vs devstral	+13.47	+20.39 @ 44.3%	-0.78 @ 28.2%

Table 1. Summary of all pairwise OpenHands backend comparisons on SWE-bench Verified. We report the overall gap in accuracy (left), the rule with the largest positive gap and its support (mid), and the subgroup with the largest negative gap (right).




Figure 5. Reverse-trend ChartQA rule. Questions starting with which asking for string entities are more accurately answered on human-written questions (87.26%) than generated ones (75.00%).

compared to monolithic error slicing.

In general, our rules reveal a rich landscape of performance differences across the task space. Table 1 summarizes the result for each pairwise comparison: the overall gap in accuracy, the rule with highest positive gap, and the strongest negative-sign rule. Across the ten backend pairs, seven contain a reverse-trend subgroup with at least 2% support, indicating that aggregate backend rankings often hide smaller but meaningful regions where the weaker backend overtakes the stronger one. We provide a comprehensive overview of the corresponding rules in Appendix B.3.

3.2. ChartQA: Synthetic vs. Real Data

Next, we study our method on the ChartQA benchmark (Masry et al., 2022), which consists of question-answer pairs grounded in chart images. The dataset contains both human-written and synthetic questions generated by a large language model. We treat these two subsets as distinct sources, with $a = 0$ corresponding to human-written questions and $a = 1$ to generated ones.

Each sample consists of a question and an associated chart. We extract a set of structured features x from the question and chart metadata for rule-based analysis. These features include properties of the question (e.g., length, first word, answer type) as well as characteristics of the data underlying the charts (e.g., number of columns, rows). We target the correct/incorrect response y of a Qwen 2.5-VL-7B-Instruct model (Qwen et al., 2025) on the test split, provided by the authors of the ChartQA benchmark².

Results. Overall, the Qwen model achieves a 80.72% accuracy on the human-written questions and a 94.96% accuracy on the generated questions. DIFFSUB shows that this difference is highly structured rather than uniform. We highlight two rules with contrasting trends.

In Figure 1, we displayed the most generated-favored rule covering 24.0% of the benchmark. It consists of short ques-

²<https://github.com/moured/qwen-vl2.5-chartqa/>

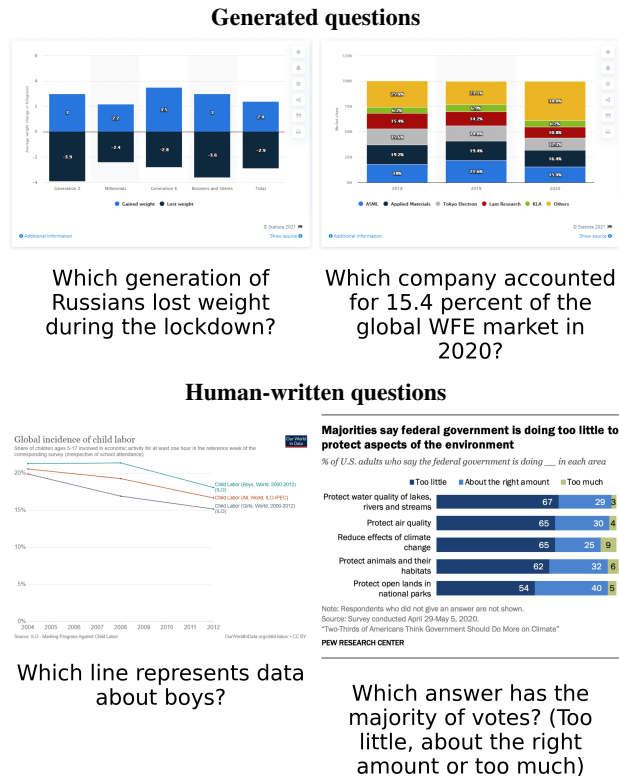


Figure 6. Example question-chart pairs of the discovered rule where accuracy on human-written questions is higher.

tions with numeric answers over charts with at most roughly 14 columns. In this region, accuracy is 95.04% on generated questions but only 62.44% on human questions. Figure 2 provides an insight into this rule by visualizing some generated and human-written questions. The generated questions can be answered by reading a numeric value from the chart, while the human-written questions require more complex reasoning. This exposes a weakness in the data generating process and provides an actionable suggestion for its improvement.

On the other hand, we also find a rule with the opposite trend, shown in Figure 5. This rule is defined by “which” questions with string answers over multi-column charts. Here, human accuracy is 87.26% while generated accuracy drops to 75.00%. This shows that synthetic questions do not simply make the task uniformly easier. Figure 6 shows a few examples. The generated questions require multiple reasoning steps from the raw data to the answer, while answering the human-written questions is less complex, and the model answers them more accurately.

We provide the full catalog of discovered rules and examples in the Appendix C. We also include the visual output of the RAI Error Analysis Dashboard for reference in Appendix C.3. It provides a detailed breakdown of performance strengths and weakness within the generated and

human question domain, but does not give any insight into performance differences *between* the two domains.

4. Conclusion

In this paper, we introduced a rule-based approach for comparative evaluation of AI agents and demonstrated its utility in two use cases. Compared to standard non-comparative error analysis tools, only our method reveals task regimes where overall performance differences stand out, providing a practical tool for understanding and improving agentic systems. Our method relies on predefined task features and is currently limited to pairwise comparisons. In future work, we plan to explore automatic task annotation using sparse autoencoders and LLMs (Gur-Arieh et al., 2025; Movva et al., 2025), and extend the approach to multiway settings.

References

- Chung, Y., Kraska, T., Polyzotis, N., Tae, K. H., and Whang, S. E. Slice finder: Automated data slicing for model validation. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*, pp. 1550–1553. IEEE, 2019.
- Gur-Arieh, Y., Mayan, R., Agassy, C., Geiger, A., and Geva, M. Enhancing automated interpretability with output-centric feature descriptions. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 5757–5778, 2025.
- Jimenez, C. E., Yang, J., Wettig, A., Yao, S., Pei, K., Press, O., and Narasimhan, K. R. SWE-bench: Can language models resolve real-world github issues? In *The Twelfth International Conference on Learning Representations*, 2024.
- Masry, A., Do, X. L., Tan, J. Q., Joty, S., and Hoque, E. Chartqa: A benchmark for question answering about charts with visual and logical reasoning. In *Findings of the association for computational linguistics: ACL 2022*, pp. 2263–2279, 2022.
- Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., and Gebru, T. Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*, pp. 220–229, 2019.
- Movva, R., Peng, K., Garg, N., Kleinberg, J., and Pierson, E. Sparse autoencoders for hypothesis generation. In *International Conference on Machine Learning*, pp. 44997–45023. PMLR, 2025.
- Qwen, :, Yang, A., Yang, B., Zhang, B., Hui, B., Zheng, B., Yu, B., Li, C., Liu, D., Huang, F., Wei, H., Lin, H., Yang, J., Tu, J., Zhang, J., Yang, J., Yang, J., Zhou, J., Lin, J., Dang, K., Lu, K., Bao, K., Yang, K., Yu, L., Li, M., Xue, M., Zhang, P., Zhu, Q., Men, R., Lin, R., Li, T., Tang, T., Xia, T., Ren, X., Ren, X., Fan, Y., Su, Y., Zhang, Y., Wan, Y., Liu, Y., Cui, Z., Zhang, Z., and Qiu, Z. Qwen2.5 technical report, 2025.
- Sagadeeva, S. and Boehm, M. Sliceline: Fast, linear-algebra-based slice finding for ml model debugging. In *Proceedings of the 2021 international conference on management of data*, pp. 2290–2299, 2021.
- Wang, X., Li, B., Song, Y., Xu, F. F., Tang, X., Zhuge, M., Pan, J., Song, Y., Li, B., Singh, J., et al. Openhands: An open platform for ai software developers as generalist agents. *arXiv preprint arXiv:2407.16741*, 2024.
- Xu, S. and Vreeken, J. Differential subgroup discovery: Characterizing where two populations differ, and why, 2026. URL <https://arxiv.org/abs/2604.27741>.




Figure 7. Resolved rate of the five OpenHands backends on SWE-bench Verified.

A. Rule Learning Method

For each comparative analysis, we apply DIFFSUB to a two-group dataset and use it as an iterative rule-learning procedure. In each run, DIFFSUB produces one explicit rule over the structured features. To obtain several non-redundant rules, we repeat the following steps:

1. Choose one directional setting: a `max` run to find rules to maximize the Jensen-Shannon divergence between the observed performance distributions, or a `min` run to find rules to minimize the divergence.
2. Run DIFFSUB with a fixed hyperparameter configuration for that dataset.
3. Extract the learned rule and record its support and group-wise performance statistics.
4. Remove all samples covered by that rule from both groups.
5. Re-run DIFFSUB on the remaining data.
6. Stop after three iterations.

This produces up to three rules for the `max` setting and up to three rules for the `min` setting. Below, we report the rules 1-3 for the `max` and `min` settings, respectively. The rules are reported in the order returned by this iterative procedure.

A.1. Experimental Setups

We use the same iterative procedure in both case studies, with dataset-specific feature sets and fixed hyperparameters.

SWE-bench. For each backend pair, we again run DIFFSUB once in the `max` setting and once in the `min` setting. Both runs use `max_reps=3`, `n_epochs=2000`, $\gamma = 0.3$, $\lambda = 0.75$, learning rate `lr_classifier=5 \times 10^{-3}`, and seed 42.

ChartQA. We run DIFFSUB once in the `max` setting and once in the `min` setting. Both runs use `max_reps=3` and seed 42. For the `max` run, we use `n_epochs=1000`, $\gamma = 0.3$, and $\lambda = 0.5$. For the `min` run, we use `n_epochs=1000`, $\gamma = 0.2$, and $\lambda = 0.5$.

B. Complete SWE-bench Rule Catalog

We conduct a pairwise comparison between five OpenHands backends on the SWE-bench Verified subset. The backends overall performance ranges from 47.66% (`devstral-small`) to 77.91% (`claude-opus-4-5`) in Figure 7.

We first provide the complete rule catalog for the `gpt5` versus `claude-opus-4-5` pairwise comparison, then provide an extended summary of the pairwise comparisons across all backends.

Features used. To characterize a task in the benchmark, the covariates are:

- `difficulty`: benchmark difficulty bucket; discrete; values `{<15 min fix, 15 min - 1 hour, 1-4 hours, >4 hours}`.
- `created_at_year`: GitHub issue creation year; continuous; range 2013–2023.
- `has_hints`: whether the task includes hints; discrete; values `{0, 1}`.

Interpretable Rule Discovery for Performance Differences Across Models and Data



- `problem_statement_words`: number of words in the problem statement; continuous; range 14–1892.
- `tests_to_pass_count`: number of tests that must pass; continuous; range 1–438.
- `pass_to_pass_count`: number of pass-to-pass checks; continuous; range 0–2476.
- `trajectory_messages`: number of messages in the execution trajectory; continuous; range 0–270.
- `trajectory_terminal_messages`: number of terminal messages in the trajectory; continuous; range 0–100.
- `trajectory_chars`: total characters in the execution trajectory; continuous; range 0–454573.
- `patch_files_changed`: number of files changed by the patch; continuous; range 1–66.
- `patch_changed_lines`: number of changed lines in the patch; continuous; range 1–16191.

B.1. Maximize Objective

Below, we list the three rules returned by the max run, which identifies subgroups with the largest performance gap between gpt5 and claude-opus-4-5.



Rule 1.

Rule
`trajectory_messages`
 $\in [58.12, 101.84]$
`trajectory_chars`
 $\in [31279.41, 335836.50]$





Rule 2.

Rule
`created_at_year`
 $\in [2015.05, 2022.31]$
`has_hints`
`patch_files_changed`
 $\in [1.69, 60.59]$
`difficulty =`
15 min -- 1 hour



Rule 3.

Rule
`created_at_year`
 $\in [2015.97, 2021.85]$
`trajectory_terminal_messages`
 $\in [2.60, 71.91]$
`difficulty =`
1--4 hours





B.2. Minimize Objective

Below, we list the three rules returned by the `min` run, which identifies subgroups with the smallest performance gap between `gpt5` and `claude-opus-4-5`.

Rule 1.

```



Rule
created_at_year
<= 2022.31
problem_statement_words
<= 1773.12
difficulty=<15 min fix
    
```



Rule 2.

```



Rule
created_at_year >= 2014.53
has_hints
problem_statement_words
>= 182.72
tests_to_pass_count <= 396.60
patch_files_changed <= 59.84
    
```



Rule 3.

```

Rule
created_at_year ∈ [2015.24, 2022.66]
problem_statement_words ∈ [85.19, 1729.90]
trajectory_messages ∈ [26.01, 248.18]
trajectory_terminal_messages ∈ [1.06, 82.22]
patch_files_changed ∈ [1.79, 59.84]
    
```



B.3. Pairwise Rule Summaries

For each backend pair in Table 1, we report the overall shared-task gap, the rule underlying the `Max Pos Gap` entry, and the rule underlying the `Max Neg Gap` entry when such a weaker-favored rule exists.

opus vs. gpt5.

Overall gap	497 shared tasks; opus 77.87% vs. gpt5 72.23% (+5.63 points).
Max rule	<code>trajectory_messages ∈ [58.12, 101.84]</code> and <code>trajectory_chars ∈ [31279.41, 335836.50]</code> Gap +33.81 points; support 23.44%; opus 89.69% vs. gpt5 55.88%.
Min rule	<code>created_at_year ∈ [2015.05, 2022.31]</code> , <code>has_hints</code> , <code>patch_files_changed ∈ [1.69, 60.59]</code> , and <code>difficulty = 15 min -- 1 hour</code> Gap -52.53 points; support 2.01%; opus 36.36% vs. gpt5 88.89%.

Interpretable Rule Discovery for Performance Differences Across Models and Data

opus vs. sonnet.

Overall gap	497 shared tasks; opus 78.07% vs. sonnet 70.82% (+7.24 points).
Max rule	<code>trajectory_messages ∈ [14.10, 299.79]</code> , <code>trajectory_terminal_messages ≤ 16.81</code> , <code>trajectory_chars ∈ [113131.41, 322326.12]</code> , and <code>difficulty = 15 min -- 1 hour</code> Gap +35.29 points; support 12.58%; opus 100.00% vs. sonnet 64.71%.
Min rule	<code>created_at_year ≥ 2013.62</code> , <code>problem_statement_words ≤ 1823.57</code> , <code>patch_files_changed ∈ [1.66, 22.43]</code> , <code>patch_changed_lines ≤ 1617.60</code> , and <code>difficulty = <15 min fix</code> Gap -2.50 points; support 17.71%; opus 81.25% vs. sonnet 83.75%.

opus vs. 4x.

Overall gap	496 shared tasks; opus 78.02% vs. 4x 61.29% (+16.73 points).
Max rule	<code>trajectory_messages ∈ [136.72, 237.33]</code> and <code>trajectory_chars ∈ [133724.69, 439362.91]</code> Gap +94.12 points; support 2.02%; opus 94.12% vs. 4x 0.00%.
Min rule	No weaker-favored rule appears among the saved rules for this pair. Gap -; support -.

opus vs. devstral.

Overall gap	489 shared tasks; opus 78.32% vs. devstral 47.85% (+30.47 points).
Max rule	<code>created_at_year ∈ [2013.22, 2020.45]</code> , <code>problem_statement_words ≤ 1497.41</code> , <code>trajectory_messages ≥ 117.43</code> , and <code>difficulty = 15 min -- 1 hour</code> Gap +64.47 points; support 8.38%; opus 75.00% vs. devstral 10.53%.
Min rule	No weaker-favored rule appears among the saved rules for this pair. Gap -; support -.

gpt5 vs. sonnet.

Overall gap	497 shared tasks; gpt5 72.03% vs. sonnet 70.62% (+1.41 points).
Max rule	<code>pass_to_pass_count ∈ [120.38, 2085.70]</code> , <code>trajectory_messages ≤ 284.69</code> , and <code>difficulty = 15 min -- 1 hour</code> Gap +14.28 points; support 2.82%; gpt5 100.00% vs. sonnet 85.72%.
Min rule	<code>problem_statement_words ≤ 1729.99</code> , <code>tests_to_pass_count ≤ 385.29</code> , <code>pass_to_pass_count ≤ 2304.96</code> , <code>trajectory_messages ≥ 40.00</code> , <code>trajectory_chars ≥ 50064.47</code> , <code>patch_files_changed ≤ 22.50</code> , and <code>patch_changed_lines ≤ 1616.85</code> Gap -1.21 points; support 27.97%; gpt5 78.12% vs. sonnet 79.33%.

gpt5 vs. 4x.

Overall gap	496 shared tasks; gpt5 71.98% vs. 4x 61.09% (+10.89 points).
Max rule	<code>created_at_year ≥ 2013.12</code> , <code>problem_statement_words ≤ 1564.78</code> , and <code>trajectory_messages ≤ 191.21</code> Gap +45.80 points; support 13.81%; gpt5 100.00% vs. 4x 54.20%.
Min rule	No weaker-favored rule appears among the saved rules for this pair. Gap -; support -.

gpt5 vs. devstral.

Overall gap	489 shared tasks; gpt5 72.39% vs. devstral 47.65% (+24.74 points).
Max rule	<code>has_hints</code> , <code>problem_statement_words ≤ 1517.41</code> , and <code>pass_to_pass_count ≥ 136.87</code> Gap +33.33 points; support 2.45%; gpt5 66.67% vs. devstral 33.33%.
Min rule	<code>problem_statement_words ≥ 214.42</code> , <code>trajectory_messages ∈ [52.07, 114.42]</code> , and <code>trajectory_chars ∈ [56865.86, 461782.00]</code> Gap -2.86 points; support 13.91%; gpt5 47.14% vs. devstral 50.00%.

sonnet vs. 4x.

Overall gap	498 shared tasks; sonnet 70.48% vs. 4x 61.04% (+9.44 points).
Max rule	<code>created_at_year ∈ [2018.52, 2022.47]</code> and <code>patch_changed_lines ≤ 84186.52</code> Gap +12.71 points; support 61.55%; sonnet 71.66% vs. 4x 58.95%.
Min rule	<code>tests_to_pass_count ∈ [3.84, 422.81]</code> , <code>pass_to_pass_count ≤ 1969.45</code> , <code>trajectory_messages ≤ 371.51</code> , and <code>trajectory_chars ≥ 28842.87</code> Gap -19.73 points; support 6.53%; sonnet 38.89% vs. 4x 58.62%.

Interpretable Rule Discovery for Performance Differences Across Models and Data

SWE-bench merged-model error-analysis baseline: surrogate tree

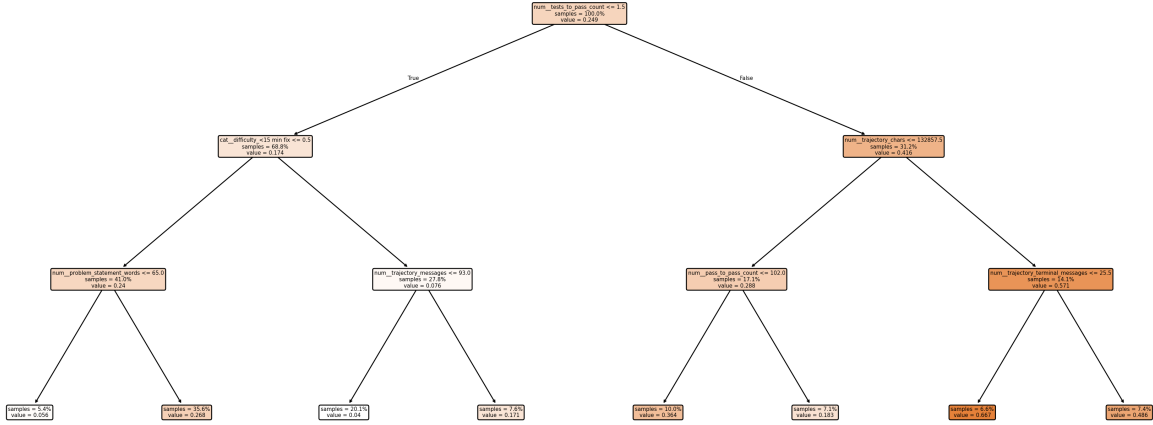


Figure 8. Surrogate decision tree as used in the RAI error analysis tool trained on the gpt5 vs.opus performance.

sonnet vs. devstral.

Overall gap	491 shared tasks; sonnet 70.88% vs. devstral 47.66% (+23.22 points).
Max rule	<code>created_at_year ∈ [2015.37, 2020.44]</code> and <code>trajectory_messages ≥ 129.31</code> Gap +60.85 points; support 15.68%; sonnet 78.85% vs. devstral 18.00%.
Min rule	<code>created_at_year ∈ [2015.63, 2019.39]</code> , <code>has_hints</code> , and <code>pass_to_pass_count ≤ 1834.71</code> Gap -49.26 points; support 3.77%; sonnet 10.00% vs. devstral 59.26%.

4x vs. devstral.

Overall gap	490 shared tasks; 4x 61.22% vs. devstral 47.76% (+13.47 points).
Max rule	<code>pass_to_pass_count ≤ 2373.31</code> , <code>trajectory_chars ≥ 31237.96</code> , <code>patch_files_changed ≤ 974.56</code> , and <code>difficulty = 15 min -- 1 hour</code> Gap +20.39 points; support 44.29%; 4x 59.22% vs. devstral 38.82%.
Min rule	<code>trajectory_messages ≤ 100.67</code> Gap -0.78 points; support 28.16%; 4x 68.52% vs. devstral 69.30%.

B.4. RAI Error Analysis SWE-bench

For comparison with the pairwise subgroup analysis, we also ran the RAI error-analysis baseline on the same 497 shared gpt5/opus tasks after stacking the two backends into one table with one row per (task, model) pair and adding model as a categorical feature. This yields 994 rows in total, with empirical accuracies 72.23% for gpt5 and 77.87% for opus. Figure 8 shows the corresponding surrogate tree.

Unlike the pairwise rules above, this baseline models absolute failure rather than relative advantage. In this merged setup, the depth-3 surrogate tree splits first on `tests_to_pass_count`, then refines by trajectory length and difficulty; notably, model does not appear among the top-level splits. This suggests that much of the variance captured by the RAI baseline is explained by broad task and execution properties before backend identity itself is used.

Leaf Rules.

- **Leaf 0.** `tests_to_pass_count > 1.5`, `trajectory_chars > 132857.5`, and `trajectory_terminal_messages ≤ 25.5`. Support 66; Accuracy 33.33%. Within this leaf, gpt5 reaches 33.90% and opus 28.57%.
- **Leaf 1.** `tests_to_pass_count > 1.5`, `trajectory_chars > 132857.5`, and

`trajectory_terminal_messages > 25.5`. Support 74; Accuracy 51.35%. All rows in this leaf come from opus.

- **Leaf 2.** `tests_to_pass_count > 1.5, trajectory_chars ≤ 132857.5, and pass_to_pass_count ≤ 102`. Support 99; Accuracy 63.64%. Within this leaf, gpt5 reaches 57.89% and opus 71.43%.
- **Leaf 3.** `tests_to_pass_count ≤ 1.5, not <15 min fix, and problem_statement_words > 65`. Support 354; Accuracy 73.16%. This is the largest mixed leaf, with gpt5 at 69.49% and opus at 76.84%.
- **Leaf 4.** `tests_to_pass_count > 1.5, trajectory_chars ≤ 132857.5, and pass_to_pass_count > 102`. Support 71; Accuracy 81.69%. Within this leaf, gpt5 reaches 76.92% and opus 87.50%.
- **Leaf 5.** `tests_to_pass_count ≤ 1.5, difficulty = <15 min fix, and trajectory_messages > 93`. Support 76; Accuracy 82.89%. Within this leaf, gpt5 reaches 76.92% and opus 84.13%.
- **Leaf 6.** `tests_to_pass_count ≤ 1.5, not <15 min fix, and problem_statement_words ≤ 65`. Support 54; Accuracy 94.44%. Within this leaf, gpt5 reaches 92.59% and opus 96.30%.
- **Leaf 7.** `tests_to_pass_count ≤ 1.5, difficulty = <15 min fix, and trajectory_messages ≤ 93`. Support 200; Accuracy 96.00%. Here both backends are strong, at 94.40% for gpt5 and 98.67% for opus.

Overall, the merged RAI baseline surfaces where *both* systems are easy or hard in absolute terms, but it is less direct than the pairwise subgroup analysis for isolating cohorts where one backend specifically outperforms the other.

C. Complete ChartQA Rule Catalog

Features used. For ChartQA, the rule discovery run uses the following structured covariates extracted from each question-chart pair:

- `query_length`: number of characters in the question; continuous; range 17–191.
- `label_kind`: answer label type; discrete; values {boolean, numeric, string}.
- `label_length`: number of characters in the answer label; continuous; range 1–10.
- `first_question_word`: first token of the question; discrete; values {how, in, is, other, what, when, which, who}.
- `n_rows_table`: number of rows in the chart table; continuous; range 1–49.
- `n_cols_table`: number of columns in the chart table; continuous; range 2–16.



Below we list all six ChartQA rules returned by DIFFSUB. For each rule, we show (i) the rule in textual form, (ii) the performance distribution for each source and (iii) a rule-condition plot, which visualizes each feature on which a condition is imposed and its marginal distribution in relation to the learned thresholds. After each rule, we display sampled examples from the human-written and generated subsets.

C.1. Maximize Objective



Below, we list the three rules returned by the `max` run, which identifies subgroups with the largest performance gap between generated and human-written questions.

Rule 1.

Rule
 $query_length \leq 177.72$
 $n_cols_table \leq 14.64$
 $label_kind = numeric$




Human-written examples




What is the percentage of logistics division in 2013 by air?

What's the share of ethnic groups except Chinese and Indians?

Generated examples





What was the turnover of the Greek manufacturing of rubber and plastic products in 2013?



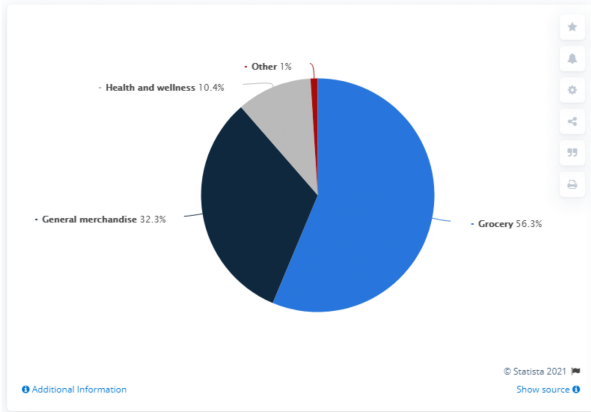
What was the highest expenditure on foreign military aid in 2009/10?

Rule 2.


Rule
 $query_length \leq 179.92$
 $n_cols_table \leq 14.83$



Human-written examples




What color does Blue indicate?




How many years does the poverty percentage rose above 11%?

Generated examples

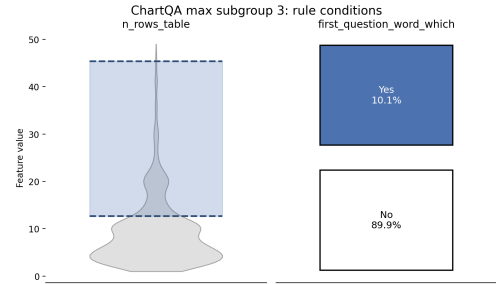



What was the average ticket price for Nashville Predators games in 2005/06?





How many Brazilian athletes participated in Sochi in 2014?

Rule
 $n_rows_table \in [12.70, 45.40]$
 $first_question_word = which$




Human-written examples




Which year saw the sharpest increase in the male percent?

Which year shows the largest decrease of employment rate?

Generated examples



Which region had the oldest grooms in Italy?



Which province had the highest relative incidence of the coronavirus?

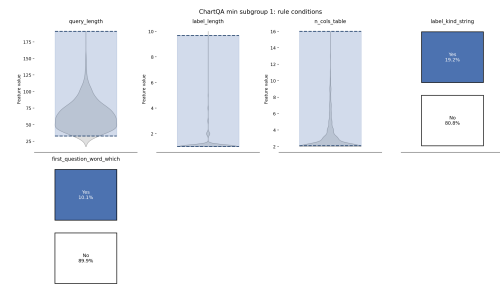

C.2. Minimize Objective

Below, we list the three rules returned by the `min` run, which identifies subgroups with the smallest performance gap between generated and human-written questions.

Rule 1.

```


Rule
query_length ≥ 32.82
label_length ≤ 9.67
n_cols_table ≥ 2.09
label_kind = string
first_question_word = which
    
```



Human-written examples

Global incidence of child labor


Share of children ages 5-17 involved in economic activity for at least one hour in the reference week of the corresponding survey (irrespective of school attendance)



Our World in Data

Majorities say federal government is doing too little to protect aspects of the environment

% of U.S. adults who say the federal government is doing ___ in each area



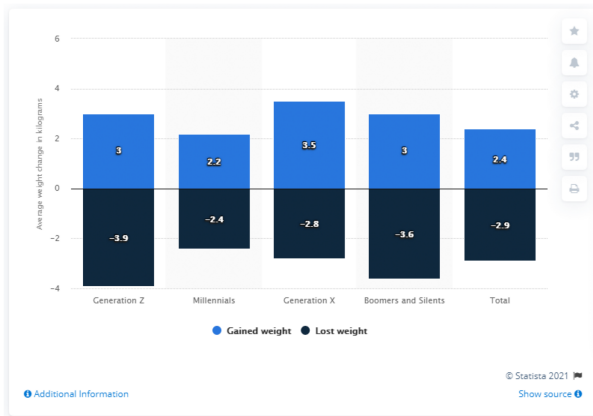
Note: Respondents who did not give an answer are not shown.
Source: Survey conducted April 29-May 5, 2020.
"Two-Thirds of Americans Think Government Should Do More on Climate"

PEW RESEARCH CENTER


Which line represents data about boys?

Which answer has the majority of votes? (Too little, about the right amount or too much)

Generated examples



Which generation of Russians lost weight during the lockdown?





Which company accounted for 15.4 percent of the global WFE market in 2020?

Rule

```


query_length ≥ 17.66
label_length ≤ 9.49
n_rows_table ∈ [3.14, 45.09]
n_cols_table ≥ 3.42
    
```



Human-written examples

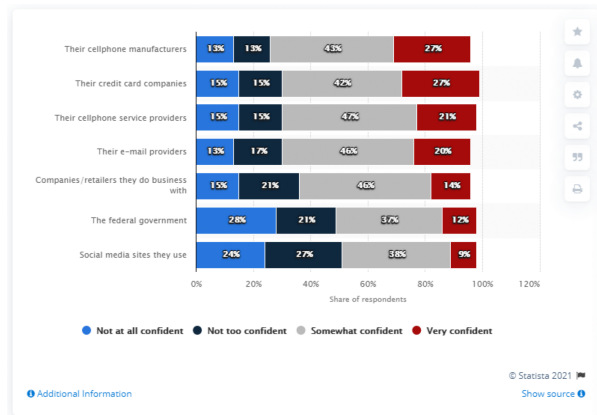
Freedom Caucus Members Have Less Seniority Than Other Republicans

Share of representatives among each group who have served ___ terms in the House (114th Congress)



Note: Partial terms are counted as full terms. Percentages may not sum to 100% because of rounding. Data are current as of Oct. 19, 2015.
 Source: Pew Research Center reporting (Freedom Caucus members) and analysis of the Congressional Biographical Directory



PEW RESEARCH CENTER



What's the value of the largest yellow bar?

What does Red color indicate?

Generated examples





What was Herman Miller's seating net sales in dollars in 2020?

What was Uruguay's literacy rate in 2018?

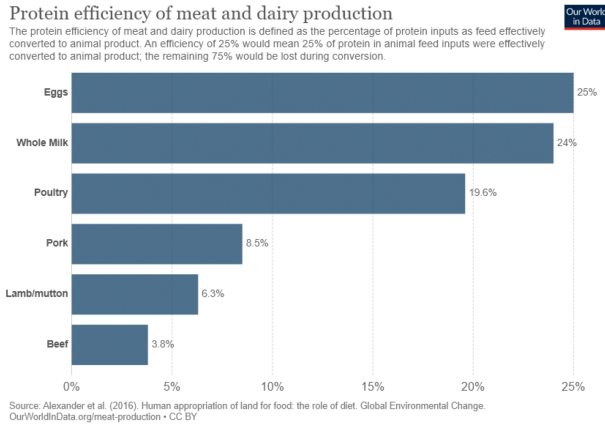
Rule 3.

```

Rule
query_length ≥ 20.38
label_length ≤ 9.48
n.rows.table ∈ [1.96, 41.55]
    
```




Human-written examples



Find the protein efficiency of Whole Milk?

Can the European Parliament's center hold?


Share of members of the European Parliament belonging to parties or political groups




PEW RESEARCH CENTER

Is the value of light green bar 35 in 1989?

Generated examples



What was Volkswagen's revenue in 2014?



How much money did George R.R. Martin earn between June 2014 and June 2015?

C.3. RAI Error Analysis Baseline

For reference, we also ran a decision-tree based error analysis baseline as used in the Microsoft Responsible AI Error Analysis Dashboard. Figure 9 shows the corresponding surrogate tree. This baseline largely partitions the data first by source (human versus generated), then refines leaves within each source. It is therefore very useful in identifying rules

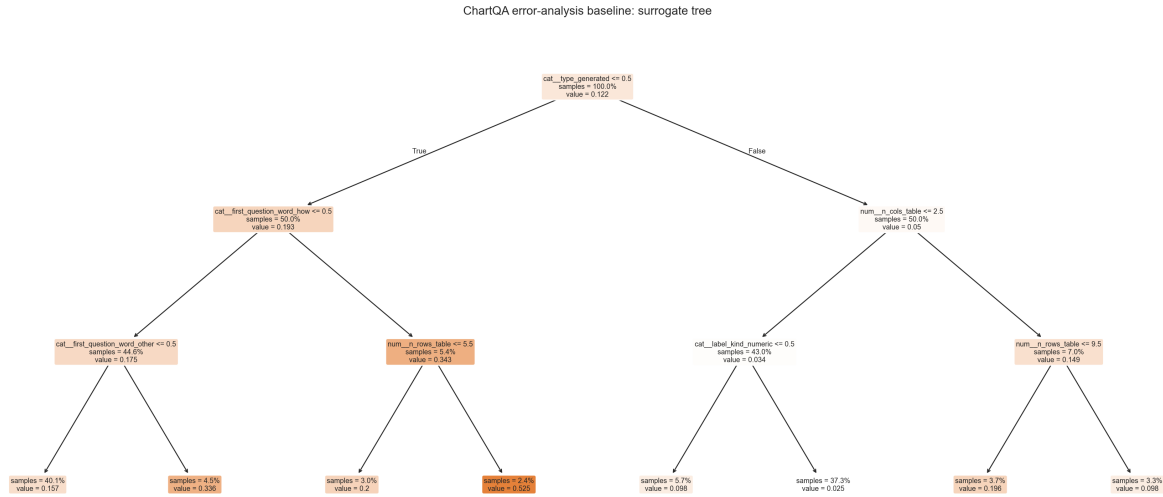


Figure 9. Surrogate decision tree used for the ChartQA error-analysis baseline.

for high/low error regions within each domain, but does not provide any information on the comparative differences of performance between generated and human-written questions.

Leaf Rules.

- **Leaf 0.** Human-written, `first_question_word = how`, and `n_rows_table > 5.5`. Support 59; Accuracy 47.46%.
- **Leaf 1.** Human-written, `first_question_word = other`. Support 113; Accuracy 66.37%.
- **Leaf 2.** Human-written, `first_question_word = how`, and `n_rows_table ≤ 5.5`. Support 75; Accuracy 80.00%.
- **Leaf 3.** Human-written, `first_question_word` neither `how` nor `other`. Support 1003; Accuracy 84.35%.
- **Leaf 4.** Generated, `n_cols_table > 2.5`, and `n_rows_table ≤ 9.5`. Support 92; Accuracy 80.43%.
- **Leaf 5.** Generated, `n_cols_table > 2.5`, and `n_rows_table > 9.5`. Support 82; Accuracy 90.24%.
- **Leaf 6.** Generated, `n_cols_table ≤ 2.5`, and answer type non-numeric. Support 143; Accuracy 90.21%.
- **Leaf 7.** Generated, `n_cols_table ≤ 2.5`, and answer type numeric. Support 933; Accuracy 97.53%.