

Appendix: Context-aware testing: a new paradigm for testing with large language models

Table of Contents

A	Extended related work	15
A.1	Enhanced overview of relevant literature	15
A.2	Components of the ML Testing pipeline	17
A.3	Comparison of the features of slice discovery methods	17
B	SMART Details	18
B.1	Contextual understanding	18
B.2	Operationalizing variables	18
B.3	Feasibility checks	20
B.4	Data adjustment queries	21
B.5	Requirements, automation, and outputs	21
B.6	Moving outside of IID data	22
B.7	SMART and multiple testing	22
C	Benchmarks & Experimental Details	23
C.1	Datasets	23
C.2	Baselines	23
C.3	LLM prompts	23
C.3.1	Generic templates.	23
C.3.2	Example prompts: OULAD Education.	24
C.4	Experimental details	29
C.4.1	Context-aware testing (Sec. 5.1.)	29
C.4.2	Requirements-constrained testing (Sec. 5.2.)	29
C.4.3	Targeting model failures (Sec. 5.3.)	30
C.4.4	Adaptive testing for a deployment environment (Sec. 5.4.)	30
C.4.5	Discovery of societally important groups and failure understanding (Sec. 5.5.)	30
D	Additional experiments	31
D.1	Requirements-constrained testing	31
D.2	Adaptive testing for a deployment environment	32
D.3	Effects of LLMs	33
D.3.1	Comparison over identified divergent slices	33
D.3.2	Performance across different models	33
D.3.3	Hypothesis generation	34
D.4	Effects of different tabular machine learning models	35
D.5	On the inductive biases of ML model testing	36
D.6	Context aware sensitivity	37
D.7	Cost of SMART	38
D.8	SMART with open-weight models	38
D.9	Understanding the importance of feature names	38
D.10	Example of a model report	40

A Extended related work

A.1 Enhanced overview of relevant literature

This work primarily engages with works on ML model testing. Consequently, we detail different paradigms of ML model testing next, which we summarize in Table 6 and Figure 5 showing that none of the existing paradigms satisfy all the properties for ML testing neither in terms of automation nor from the perspective of relevance as they are unable to incorporate context and/or requirements.

Table 6: Comparison of ML testing paradigms in terms of how tests are defined, the type of search space (i.e. relevance based on context and requirements), sensitivity to multiple testing and automation to enable scalability.

Paradigm	Objective	Test definition	Search space			Outputs	Examples
			Context-aware	Integrate Requirements	Automated testing		
Average testing	Overall performance	Data split	N.A.	N.A.	✓	Single score	
Behavioral	Test expected behavior	Pre-defined (Human)	N.A.	N.A.	✗	Multi-dimensional score	[16, 17, 5]
Data-only	Identify divergent groups	Search	✗	✗	✓	Multi-dimensional score	[14, 21, 15, 12, 13]
SMART (Ours)	Probe for model failure	Search	✓	✓	✓	Multi-dimensional score justifications, report	

Existing paradigms in ML model testing. The ML community has predominately approached ML model evaluation/testing via the use of held-out test datasets. On the basis of the test dataset a single performance metric (accuracy, AUC etc) is computed. This single average evaluation may mask nuances of the model’s performance along various dimensions. One approach to address this is to create better benchmark datasets when evaluating models on common benchmarking tasks. For example, manual corruptions like Imagenet-C [41] or by collecting additional real data such as the Wilds benchmark [8]. Benchmark datasets are labor-intensive to collect and their utility is limited to the specific benchmark tasks.

What if we want to evaluate models not confined to benchmarking tasks? To test in cases beyond benchmark tasks, the community has proposed trying to find slices or regions wherein the model fails (i.e. via stress tests). As mentioned in Sec. 1, this could be ► **Behavioral testing**: which requires human expertise and intuition to define the test scenarios (e.g. Checklist [16], HateCheck [17] or 3S-Testing [5]) — which is not automated and does not scale. Moreover, it runs a high risk of overlooking critical weaknesses due to human cognitive biases. ► **Data-only testing**: which does not account for context or requirements and searches exhaustively (e.g. Autostrat [14], SliceFinder [12] or DivExplorer [15]). This may slices focus on arbitrary, less important, or unrealistic/implausible scenarios that are unlikely to be seen in reality. Moreover, we run the risk of the multiple testing problem, where by virtue of the large number of tests evaluated, we might discover a divergent group by chance.

Hypothesis-driven ML model testing. In contrast, a hypothesis-driven approach to ML model testing brings about the falsifiability approach widely adopted in science. It begins with the formulation of specific, testable hypotheses based on theoretical understanding, context/domain knowledge, and the intended application of the model (i.e. requirements). The concept of hypothesis-driven ML model testing is deeply rooted in work by Popper [28] who proposed the principle of falsifiability as a driver of scientific progress. The progression of knowledge hinges on the formulation and rigorous testing of hypotheses which can either be falsified or supported by empirical evidence. In the context of hypotheses in ML model testing, testable statements about the model’s performance under various conditions, fairness, and robustness can improve our understanding of the model’s performance. By rigorously testing these hypotheses, we can uncover the strengths and limitations of ML models.

Contrast to data-only methods on unstructured data.

In this paper, we have discussed data-only methods for slice discovery or blindspot discovery applicable to structured data. Specifically, we focus on tabular data where metadata in the form of column names is explicitly encoded into the data. The data-only approaches covered in this work directly search over the raw feature space to identify slices with similar attributes wherein the model would exhibit underperforming predictions.

For completeness, we contrast to data-only approaches often applicable to data without explicit structure and metadata (images, text, audio). There have been numerous methods including [42–45]

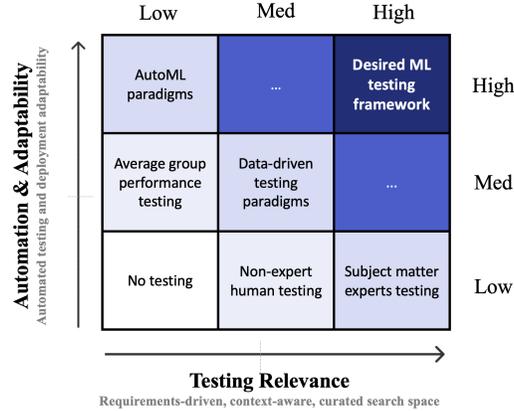


Figure 5: Contrasting paradigms of ML model testing along (i) Testing relevance which accounts for requirements and/or context and (ii) Degree of automation and adaptability when carrying out testing. We desire a new paradigm of ML testing to address both.

to identify slices of unstructured data with systematic failures. They all follow a similar pattern: (1) the data is embedded (often with a pre-trained) model into a representation space and (2) the underperforming slices are identified by clustering on the raw embedding space or after dimensionality reduction. We then need to post-hoc interpret the clusters in order to understand what they represent. This contrasts the tabular data setting where we find cohesive groups of features that provide an explicit interpretation. Moreover, in tabular regimes, we do not have access to pre-trained models in the same way as for domains such as images or text.

Specifically, to context-aware testing we also note that tabular data inherently includes context through interpretable/meaningful feature names and metadata. This context is not naturally present in images (which are tensors of pixel intensities). Hence, extending CAT to other domains like images or text would require incorporating metadata to provide necessary context — which is often unavailable. In addition, one would need to develop new ways to operationalize hypotheses on unstructured data.

Contrast to software testing.

Going beyond ML model testing, the idea of testing is also prevalent in software systems. For example, unit tests of functions in a codebase. We highlight that software testing of functional input-output correctness is also an area proposed in software testing, which could also theoretically be applied to ML systems. That said, we contrast between SMART and these paradigms in Table 7 below, demonstrating that we tackle a different testing problem.

Table 7: Comparison of SMART testing with software testing works

Criteria	SMART	Christakis et al. [18]	Sharma et al [19]	How is SMART different?
Test case generation	Automatic LLM-based test cases	Pre-Specified dimensions of functional correctness (k-safety properties)	Approximates the black-box model to test with a white-box decision tree model.	SMART uses the LLM to automatically hypothesize the test cases & execute them
What is the focus of testing	Identify model failures	Assesses I/O functional correctness	SMT solvers find violations to monotonicity properties . i.e. input increase!=output increase	SMART failures encompass a broader range of possibilities
Are tests context-aware?	Yes	No	No	SMART uniquely incorporates context awareness into testing, leading to more realistic assessments.

Additionally, a further orthogonal area to ML model testing is that of data validation [46, 47]. In contrast, to testing an ML model for failures, data validation aims to test input data pipelines for data quality problems or drift.

A.2 Components of the ML Testing pipeline

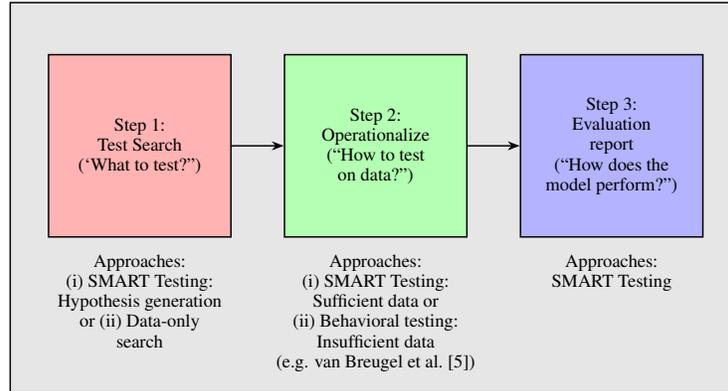


Figure 6: Components of ML Testing — (i) Test search, (ii) Operationalize, (iii) Evaluation report, with example approaches for each component.

- Test search: First, we need to decide “what to test”. This is a search problem to identify test dimensions. This could either be done via SMART which is targeted (via context and requirements) or data-only methods (where the space is larger). Alternatively, if human experts are available, humans could define the test dimensions.
- Operationalize: Second, we need to operationalize the test and address the challenge of “how to carry out this test on data”. If we have sufficient data — SMART could operationalize the tests on the data via an interpreter. Alternatively, if we don’t have sufficient data to run the test — then SMART could be augmented by behavioral testing approaches such as van Breugel et al. [5] — which once the test has been defined use synthetic data to augment small subgroups/slices.
- Evaluation report: Third, we carry out the test and evaluate the model to answer the question, “how does the model perform”. SMART can be used to produce a comprehensive report of failures and justifications in an automated manner.

A.3 Comparison of the features of slice discovery methods

Even as we zoom into the task of discovering slices where the model might underperform, we observe that SMART has features which are not supported by most of other discovery methods. We exemplify some of these features in Table 8

Table 8: Comparison of slice discovery methods

Criteria	SliceFinder	Pysubgroup	DivExplorer	Autostrat	SMART Testing
Integrates custom domain knowledge	✗	✗	✗	✗	✓
Constant slice discovery time	✗	✗	✗	✗	✓
Performance is resistant to irrelevant data	✗	✗	✗	✗	✓
Can capture rare slices	✗	✗	✗	✗	✓
Inherently supports logical ORs	✗	✗	✗	✓	✓
Resistant to overfitting the training set	✗	✗	✗	✗	✓

B SMART Details

We present a block diagram of the key components of SMART Testing in Figure 7. In addition, for each component we provide additional motivations and technical details not covered in the main paper. We further provide more technical information on certain implementation details.

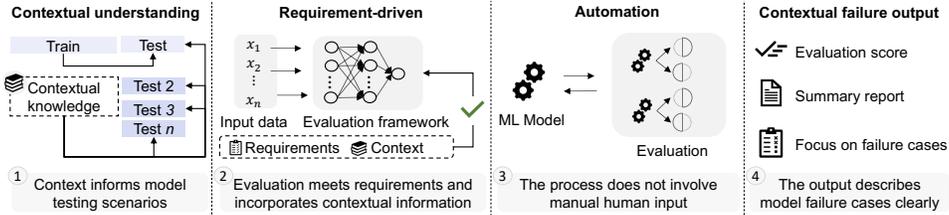


Figure 7: A strong machine learning testing framework should incorporate textual understanding during testing, should meet requirements, should be automated, and provide contextual reports with an emphasis on model failures.

B.1 Contextual understanding

A core component of SMART is leveraging LLMs to strategically identify *what to test for* in tabular ML models. This process is anchored on the premise that LLMs can effectively navigate the space of potential model failure slices, harnessing their contextual understanding to pinpoint slices where the model is most susceptible to failure. The operation of LLMs within the SMART framework involves three critical inputs:

1. **Context (\mathcal{C}):** A string describing the overarching scenario and the task at hand.
2. **Dataset Information (\mathcal{D}_c):** Extracted from the training data, this includes a description of the observations where the model did not fail ($\hat{y} = y$), or where it failed ($\hat{y} \neq y$). This could be characterized as a string with a description of the covariates of each group (e.g., mean, median, mode, and a textual description of the distribution).

Utilizing these inputs, the LLM generates a set of hypotheses (\mathcal{H}) and corresponding justifications (\mathcal{J}) regarding potential model failures. The framework also incorporates a *self-refine mechanism* to enhance hypothesis generation. This mechanism iteratively refines hypotheses based on the observed or provided context (\mathcal{C}) and data (\mathcal{D}_c), re-ranking them by their likelihood. The self-refine mechanism is introduced in order to generate hypotheses that are more likely to target specific model failures.

B.2 Operationalizing variables

Once the hypotheses have been proposed, it is important to operationalize them. This operationalization can be achieved through one of two methods:

1. **LLM-based Operationalization:** Hypothesizing possible operationalizations (e.g., $age > 70$ as a way to operationalize “elderly people”). In this case, the previous interactions and information is provided to the LLM, together with the relevant description of the data. It is then asked to provide possible ways to operationalize a specific hypothesis. This is done by passing an “operationalization prompt” which contains the aforementioned information.

The following is an example of an operationalization prompt in Python:

```
1     f"""
2         The following are hypotheses about which people within
3         a dataset the model might underperform on.
4         Propose specific ranges for each hypothesis. Hypotheses
5         : {hypotheses}.
6
7         Dataset information: {context}. {context_target}
8
9         The dataset contains {len(unique_values)} columns. The
10        columns are {', '.join(unique_values.keys())}. The values
11        are {str(unique_values.items())}
12
13        TASK: Propose specific variable ranges for each
14        hypothesis such that they are clearly operationalizable and
15        defined. Use this format: Hypothesis: <>; Operationalization
16        : <>.
17        """
```

Code Listing 2: Operationalization (LLM knowledge): General template

This is then converted into a specific operationalization for each hypothesis using an external compiler which maps the strings to a function.

2. **Data-driven Operationalization:** Utilizing training data to identify optimal splits for given covariates (e.g., $age \geq 82.32$ as a condition to split data based on the “age” covariate). This method receives the covariate (or set of covariates) as an input and returns the optimal split. The optimal split is defined as the split which can identify two groups which have the largest absolute difference in accuracies given some requirements (e.g. a minimum or maximum group size in each group). This splitting can be done by *any black box splitting function* which takes in a set of inputs and returns a splitting mechanism.

Black-box splitting function for data-driven operationalization. While any black-box splitting function could perform on the data, we provide details on the specific function used in SMART.

The function implemented is based on a decision tree model, which can be either a regressor or a classifier, depending on the nature of the outcome variable. In all our experiments, we perform classification-based tasks.

The function operates by fitting a decision tree to the data and recursively traversing the tree to find the split that yields the largest absolute difference in the outcome variable’s mean value between two slices. The split must also satisfy group size constraints, specified as minimum and maximum group sizes. The traversal process evaluates each potential split, calculating the mean outcome for each slice and the discrepancy between these means. The optimal split is the one that maximizes this discrepancy while adhering to the group size requirements.

The function returns a query string that represents this optimal split. This string can then be used to segment the dataset into the identified slices for further analysis or testing. The following is pseudocode for the splitting algorithm.

Algorithm 1 Optimal splitting mechanism within the SMART framework

Data: *dataframe, features, outcome, min_group_size, max_group_size*

Result: Optimal query string for data split

```
1 Function GetOptimalSplitQuery(dataframe, features, outcome, min_group_size,  
   max_group_size):  
2   Validate input features and outcome in dataframe Determine the type of decision tree  
   model based on outcome type Fit the decision tree model to dataframe using features  
   and outcome Initialize an empty list conditions for tracking split conditions return  
   TraverseTree(root, 0, conditions)  
3 Function TraverseTree(node, depth, conditions):  
4   if node is a leaf then  
5     Calculate the discrepancy in outcome between the two slices return the condition and  
     discrepancy if group size constraints are met  
6   end  
7   Determine left and right conditions based on the threshold at node return  
   TraverseTree(node.left, depth + 1, conditions  $\cup$  {left_condition}) // or  
   the right conditions, depending on which has the greater discrepancy
```

This splitting mechanism is used for continuous features. A separate black-box splitting mechanism is developed for categorical features based on iterating on different permutations of these features and evaluating them that helps with variable selection.

In practice, we employ the LLM-based operationalization for the SMART ablation and the data-driven operationalization within the original SMART framework. However, we highlight that this is a design choice that we have found works in practice; there is nothing stopping from using any operationalization framework within the main SMART framework.

B.3 Feasibility checks

SMART is built on many modules which can be toggled on or off. One of such modules is a module called “feasibility check” which evaluates whether there are any hypotheses which should be tested in the first place. The experiments presented in Sec. 5.1 highlight the importance of being able to identify when no relationship between covariates exist.

The feasibility check contains three steps. First, an LLM is queried to evaluate whether any relationships could exist between the covariates and an outcome variable. For instance, this could be whether a relationship could exist between loan default and the annual rainfall in a given region. Second, the answer is self-refined. This helps to evaluate feasibility because, in practice, initial responses tend to be over-optimistic (such as hypothesizing that annual rainfall is associated with loan default via a geographic proxy). This self-refinement helps to critically evaluate the previous answer. The number of steps in the self-refinement process is a hyperparameter. Lastly, the answer (whether or not a relationship could exist between the two variables and, hence, should be inspected) is converted to a boolean value via an external function.

The following is pseudocode which implements the feasibility module.

Algorithm 2 Feasibility for evaluating slices

Input: *unique_values*, *context*, *context_target*, *system_message*, *n_refine*

Output: *feasibility_boolean_response*

```
8 Function FeasibilityCheck(unique_values, context, context_target, system_message,  
  n_refine):  
  // Construct the feasibility task prompt  
9  task ← “Evaluate subgroups for model performance” task ← task ∪ “Context: ” ∪ context ∪  
    “Target: ” ∪ context_target task ← task ∪ “Columns: ” ∪ join(unique_values.keys())  
  // Get initial feasibility response  
10 feasibility_response ← GetLLMResponse(task, system_message)  
  // Refine the answer  
11 feasibility_response ← SelfRefine(unique_values, context, context_target,  
  feasibility_response, system_message, n_refine)  
  // Convert to boolean  
12 boolean_task ← “Based on analysis, provide yes/no answer” boolean_task ←  
  boolean_task ∪ “Analysis: ” ∪ feasibility_response  
13 feasibility_boolean_response ← GetLLMResponse(boolean_task)
```

B.4 Data adjustment queries

Given that SMART operates with an LLM, sometimes the framework outputs proposals which do not operationalize on the data well. For instance, even if a column “age” is a categorical variable, SMART might propose to operationalize the hypothesis “elderly people” as $age > 72$ which would cause an error.

To avoid this, we implement an additional data adjustment module which can handle such cases. It catches the error and re-prompts the LLM to find a group which could be operationalized given the data structure, and does so iteratively until such a group is found.

The following is a pseudocode function that explains how data adjustment is performed.

Algorithm 3 Pseudocode for adjusting subgroup/slice queries

Data: *X* (DataFrame), *n_subgroups* (integer)

Result: Adjusted subgroup queries in the dataset

```
14 Function AdjustSubgroupQuery(X, n_subgroups):  
15   for each subgroup condition do  
16     if CheckQueryExistence(X, condition) then  
17       UpdateSubgroupCondition(condition)  
18     else  
19       // Condition yields no rows, adjust it  
       adjusted_condition ← GetLLMResponse(condition) // Update the condition  
       in the subgroup  
20       UpdateSubgroupCondition(adjusted_condition)  
21     end  
22   end  
23 return
```

B.5 Requirements, automation, and outputs

Requirements. SMART can natively integrate user requirements into its framework. This is done by inputting requirements as a string and concatenating it together with the context. Some requirements are directly integrated into the framework itself (e.g. functionalities for determining the minimum or maximum sample size of a data split).

Automation. Fig. 1 showcases the pipeline of SMART and which components are automated. SMART is developed using an “sklearn” style fit, predict framework. The fit method automatically performs a feasibility check, generates hypotheses, justifications, operationalizes them, performs *self-falsification* using empirical data, re-ranks the hypotheses, and saves all intermediate results. The

method can then automatically be used on any piece of data to evaluate whether it underperforms on the groups that have been found to underperform.

Outputs. In addition to SMART outputting subgroups/slices or a scalar number, it can output a model report. An example report is provided in D.10. We note that this report is simply an example which employs both the findings of the fitting procedure and an additional LLM to summarize the outputs. More fine-grained outputs can be constructed.

B.6 Moving outside of IID data

In the paper, we propose that SMART can move outside of IID data and generalize better in the presence of covariate shift (refer to Sec. D.2). We highlight that this is done by performing an analysis on the original data and using an LLM propose possible hypotheses and splits for a different target domain where no data is present (hence, operationalizing only using the LLM with access to previous operationalizations). The following is an example prompt which is designed to do this.

```

1 f """
2     You have access to the following information.
3
4     Dataset information: {context}. {context_target}
5
6     The dataset contains {len(unique_values)} columns. The columns
7     are {'', '.join(unique_values.keys())}. The values are {str(
8     unique_values.items())}
9
10    However, you are no longer working with the same data as just
11    described. Rather, this is the context: {new_context}.
12
13    These are the hypotheses: {self._updated_hypotheses}.
14
15    TASK: Propose specific variable ranges for each hypothesis
16    such that they are clearly operationalizable and defined. Use this
17    format: Hypothesis: <>; Operationalization: <>.
18    """

```

Code Listing 3: Operationalization (LLM knowledge): General template

B.7 SMART and multiple testing

The reason why SMART performs testing via hypothesis generation is because testing for all slices is equivalent to generating and testing a hypothesis on the data. Here, we outline in greater detail how hypothesis generation is connected to multiple testing.

As outlined in the main manuscript, searching for f failures may bring up the challenge of multiple hypothesis testing. Specifically, when we evaluate the failure rate of model f across different slices $S_i \subseteq \mathcal{D}$, we are testing the null hypothesis $H_0^{(i)} : \mu_{S_i} = \mu_{\mathcal{D}}$ against the alternative $H_1^{(i)} : \mu_{S_i} \neq \mu_{\mathcal{D}}$. Then, the probability of making a Type I error increases with each test. This drastically inflates the family-wise error rate (FWER). For instance, assuming that each slice is independent, the probability of making one or more Type I errors across all tests is given by $1 - (1 - \alpha)^m$, where m is the total number of slices tested. While this can be addressed by adjusting for multiple testing, we run into the trade-off between the FWER control and statistical power. As we employ statistical correction methods to account for Type I errors, we increase the probability of Type II errors.

C Benchmarks & Experimental Details

We summarize all experimental details, datasets and benchmarks.

Code can be found at: https://github.com/pauliusrauba/SMART_Testing or https://github.com/vanderschaarlab/SMART_Testing

C.1 Datasets

We summarize the different datasets we use in this paper in Table 9. The datasets vary in number of samples, number of features and domain.

Table 9: Summary of the datasets used.

Name	n samples	n features	Domain	Task
Loan	5000	8	Finance	Loan default
Breast	116	8	Healthcare/Medicine	Breast cancer
Compas [33]	4743	8	Criminal justice	Criminal recidivism
OULAD [34]	5000	8	Education	Exam score
Diabetes	768	8	Healthcare/Medicine	Diabetes label
SEER [35]	2000	12	Healthcare/Medicine	Prostate cancer morality
CUTRACT [35]	20000	12	Healthcare/Medicine	Prostate cancer morality

Note that, depending on the experiment, we might vary the sample size or features used. This can be because of the added synthetic variables or certain slice discovery algorithms taking too long to run.

C.2 Baselines

We compare SMART testing with the following data-only baselines in Section 5.

- Autostrat [14]. We use the implementation from ⁷
- Pysubgroup [21]. We use the implementation from ⁸
- Divexplorer [15]. We use the implementation from ⁹
- Slicefinder [12]. We use the implementation from ¹⁰
- Sliceline [13]. We use the implementation from ¹¹

C.3 LLM prompts

C.3.1 Generic templates.

```
1 """Constructs the task description for the LLM."""
2 task = f"""Your task is to propose possible hypotheses as to which
3 subgroups within the dataset might have worse predictive
4 performance than on average because of societal bias in the
5 dataset, insufficient data, other relationships, or others. The
6 subgroups might be based on any of the provided characteristics,
7 as well as on any combination of such characteristics.

Dataset information: {context}. {context_target}

The dataset contains {len(unique_values)} columns. The columns are {'',
'.join(unique_values.keys())}.
```

⁷https://github.com/Adebayo-Oshingbesan/BPS_Autostrat_Code/

⁸<https://github.com/flemmerich/pysubgroup>

⁹<https://github.com/elianap/divexplorer>

¹⁰<https://github.com/yeounoh/slicefinder>

¹¹<https://github.com/DataDome/sliceline>

```

8 Task: Create {n} hypotheses as to which subgroups within the dataset
  the model will perform worse than on average because of societal
  biases or other reasons. Important: Your hypothesis can contain
  either one variable or two variables in the condition. Therefore,
  your goal is to find discrepancies in the model's performance, not
  the underlying data outcomes. Justify why you think that for each
  of the {n} hypotheses. Format of the output: Hypothesis: <>;
  Justification: <>.
9
10 """

```

Code Listing 4: Generic hypothesis prompt

```

1 """
2 The following are hypotheses about which people within a dataset the
  model might underperform on.
3 Propose specific ranges for each hypothesis. Hypotheses: {hypotheses}.
4
5 TASK: return a dictionary that contains an index number as the key and
  the column value as the value. If there are multiple columns in
  that hypothesis, return them in a list. There are the column names
  : {'', '.join(unique_values.keys())}.
6 """

```

Code Listing 5: Generic operationalization prompt

C.3.2 Example prompts: OULAD Education.

```

1 """
2 -----INPUT TEXT -----
3 Your task is to propose possible hypotheses as to which subgroups
  within the dataset might have worse predictive performance than on
  average because of societal bias in the dataset, insufficient
  data, other relationships, or others. The subgroups might be based
  on any of the provided characteristics, as well as on any
  combination of such characteristics.
4
5 Dataset information:
6 Open University Learning Analytics Dataset (OULAD) contains data about
  courses, students and their interactions with Virtual Learning
  Environment (VLE) for seven selected courses (called modules).
  Presentations of courses start in February and October - they are
  marked by B and J respectively. The dataset consists of tables
  connected using unique identifiers. All tables are stored in the
  csv format.
7
8 The task is to predict the students who are likely to fail or withdraw
  from a course.
9
10
11 The dataset contains 14 columns. The columns are gender, region,
  highest_education, imd_band, age_band, num_of_prev_attempts,
  studied_credits, disability, test, group_0, group_1, group_2,
  group_3, group_4. The values are dict_items([('gender', ['M', 'F
  ']), ('region', ['South Region', 'Scotland', 'Yorkshire Region', '
  West Midlands Region', 'East Anglian Region', 'East Midlands
  Region', 'South East Region', 'Wales', 'North Western Region', '
  London Region', 'South West Region', 'Ireland', 'North Region']),
  ('highest_education', ['Lower Than A Level', 'A Level or
  Equivalent', 'HE Qualification', 'No Formal quals', 'Post Graduate
  Qualification']), ('imd_band', ['70-80%', '40-50%', '0-10%',
  '10-20', '50-60%', '20-30%', '60-70%', '80-90%', '90-100%',
  '30-40%', nan]), ('age_band', ['0-35', '35-55', '55<=']), ('
  num_of_prev_attempts', [0, 1, 2, 3, 5, 4, 6]), ('studied_credits',
  {'min': 30, 'mean': 79.77122804326149, 'max': 655}), ('disability

```

```

', ['N', 'Y']), ('test', [False, True]), ('group_0', [False, True
]), ('group_1', [True, False]), ('group_2', [True, False]), ('
group_3', [False, True]), ('group_4', [False, True]))
12
13 Task: Create 5 hypotheses as to which subgroups within the dataset the
    model will perform worse than on average because of societal
    biases or other reasons. Therefore, your goal is to find
    discrepancies in the model's performance, not the underlying data
    outcomes. Justify why you think that. Format of the output:
    Hypothesis: <>; Justification: <>.
14 """
15
16 -----OUTPUT TEXT -----
17 Hypothesis 1: Students with disabilities will have worse predictive
    performance than average.
18 Justification: Students with disabilities may face additional
    challenges in their learning process, which could affect their
    performance in the course. They may require additional support or
    accommodations that are not adequately captured in the dataset,
    leading to lower predictive performance for this subgroup.
19
20 Hypothesis 2: Students from lower socioeconomic backgrounds (indicated
    by imd_band) will have worse predictive performance than average.
21 Justification: Students from lower socioeconomic backgrounds may have
    limited access to resources and support systems that are necessary
    for successful completion of the course. This lack of resources
    and support may not be adequately captured in the dataset, leading
    to lower predictive performance for this subgroup.
22
23 Hypothesis 3: Students with lower levels of education (indicated by
    highest_education) will have worse predictive performance than
    average.
24 Justification: Students with lower levels of education may have less
    prior knowledge and skills that are necessary for success in the
    course. This lack of foundational knowledge may not be adequately
    captured in the dataset, leading to lower predictive performance
    for this subgroup.
25
26 Hypothesis 4: Students who have attempted the course multiple times (
    indicated by num_of_prev_attempts) will have worse predictive
    performance than average.
27 Justification: Students who have attempted the course multiple times
    may have encountered difficulties or challenges in their previous
    attempts that hindered their success. These challenges may not be
    adequately captured in the dataset, leading to lower predictive
    performance for this subgroup.
28
29 Hypothesis 5: Students in certain regions (indicated by region) will
    have worse predictive performance than average.
30 Justification: Students in certain regions may have different
    educational systems, resources, and support networks that can
    impact their performance in the course. These regional differences
    may not be adequately captured in the dataset, leading to lower
    predictive performance for subgroups based on region.

```

Code Listing 6: Hypothesis generation: OULAD Dataset

```

1 """
2 -----INPUT TEXT -----
3
4 The following are hypotheses about which people within a dataset the
    model might underperform on.
5 Propose specific ranges for each hypothesis. Hypotheses: Hypothesis 1:
    Students with disabilities will have worse predictive performance
    than average.

```

6

7 Justification: Students with disabilities may face additional challenges in their learning process, which could affect their performance in the course. They may require additional support or accommodations that are not adequately captured in the dataset, leading to lower predictive performance for this subgroup.

8

9 Hypothesis 2: Students from lower socioeconomic backgrounds (indicated by imd_band) will have worse predictive performance than average.

10 Justification: Students from lower socioeconomic backgrounds may have limited access to resources and support systems that are necessary for successful completion of the course. This lack of resources and support may not be adequately captured in the dataset, leading to lower predictive performance for this subgroup.

11

12 Hypothesis 3: Students with lower levels of education (indicated by highest_education) will have worse predictive performance than average.

13 Justification: Students with lower levels of education may have less prior knowledge and skills that are necessary for success in the course. This lack of foundational knowledge may not be adequately captured in the dataset, leading to lower predictive performance for this subgroup.

14

15 Hypothesis 4: Students who have attempted the course multiple times (indicated by num_of_prev_attempts) will have worse predictive performance than average.

16 Justification: Students who have attempted the course multiple times may have encountered difficulties or challenges in their previous attempts that hindered their success. These challenges may not be adequately captured in the dataset, leading to lower predictive performance for this subgroup.

17

18 Hypothesis 5: Students in certain regions (indicated by region) will have worse predictive performance than average.

19 Justification: Students in certain regions may have different educational systems, resources, and support networks that can impact their performance in the course. These regional differences may not be adequately captured in the dataset, leading to lower predictive performance for subgroups based on region..

20

21 Dataset information:

22 Open University Learning Analytics Dataset (OULAD) contains data about courses, students and their interactions with Virtual Learning Environment (VLE) for seven selected courses (called modules). Presentations of courses start in February and October - they are marked by B and J respectively. The dataset consists of tables connected using unique identifiers. All tables are stored in the csv format.

23 .

24 The task is to predict the students who are likely to fail or withdraw from a course.

25

26

27 The dataset contains 14 columns. The columns are gender, region, highest_education, imd_band, age_band, num_of_prev_attempts, studied_credits, disability, test, group_0, group_1, group_2, group_3, group_4. The values are dict_items([('gender', ['M', 'F']), ('region', ['South Region', 'Scotland', 'Yorkshire Region', 'West Midlands Region', 'East Anglian Region', 'East Midlands Region', 'South East Region', 'Wales', 'North Western Region', 'London Region', 'South West Region', 'Ireland', 'North Region']), ('highest_education', ['Lower Than A Level', 'A Level or Equivalent', 'HE Qualification', 'No Formal quals', 'Post Graduate Qualification']), ('imd_band', ['70-80%', '40-50%', '0-10%',

```

28 '10-20', '50-60%', '20-30%', '60-70%', '80-90%', '90-100%',
29 '30-40%', nan]), ('age_band', ['0-35', '35-55', '55<=']), ('
num_of_prev_attempts', [0, 1, 2, 3, 5, 4, 6]), ('studied_credits',
{'min': 30, 'mean': 79.77122804326149, 'max': 655}), ('disability
', ['N', 'Y']), ('test', [False, True]), ('group_0', [False, True
]), ('group_1', [True, False]), ('group_2', [True, False]), ('
group_3', [False, True]), ('group_4', [False, True]))
30
31
32 -----OUTPUT TEXT-----
33 Hypothesis 1: Students with disabilities will have worse predictive
performance than average.
34 Operationalization: Students with disabilities can be identified by
the "disability" variable in the dataset. To test this hypothesis,
we can define the range for students with disabilities as those
who have a value of "Y" in the "disability" column.
35
36 Hypothesis 2: Students from lower socioeconomic backgrounds (indicated
by imd_band) will have worse predictive performance than average.
37 Operationalization: Socioeconomic background can be measured using the
"imd_band" variable in the dataset. To test this hypothesis, we
can define the range for students from lower socioeconomic
backgrounds as those who have an "imd_band" value in the range of
"0-30%".
38
39 Hypothesis 3: Students with lower levels of education (indicated by
highest_education) will have worse predictive performance than
average.
40 Operationalization: Education level can be determined using the "
highest_education" variable in the dataset. To test this
hypothesis, we can define the range for students with lower levels
of education as those who have a value of "Lower Than A Level" or
"No Formal quals" in the "highest_education" column.
41
42 Hypothesis 4: Students who have attempted the course multiple times (
indicated by num_of_prev_attempts) will have worse predictive
performance than average.
43 Operationalization: Previous attempts can be measured using the "
num_of_prev_attempts" variable in the dataset. To test this
hypothesis, we can define the range for students who have
attempted the course multiple times as those who have a value
greater than 1 in the "num_of_prev_attempts" column.
44
45 Hypothesis 5: Students in certain regions (indicated by region) will
have worse predictive performance than average.
46 Operationalization: Region can be determined using the "region"
variable in the dataset. To test this hypothesis, we can define
the range for students in certain regions as those who belong to
the regions of "North Region" or "Wales".

```

Code Listing 7: Operationalization (LLM knowledge): OULAD Dataset

```

1 """
2 -----INPUT TEXT -----
3
4 The following are groups that are defined based on the dataset.
Convert them into a Python dictionary format. Each group should be
represented as a key-value pair in the dictionary, where the key
is an index (0 to 4), and the value is a string representing the
group using Python syntax and logical operators. For multiple

```

```

5         conditions, use Python's logical 'and' ('&&') or 'or' ('||').
6         Ensure the format is a valid Python dictionary.
7
8 Examples:
9 - Single Condition: {0: 'X > 45'}
10 - Multiple Conditions: {1: '(X > 45) and (Y < 20)'}
11
12 Groups to summarize: Hypothesis 1: Students with disabilities will
13 have worse predictive performance than average.
14
15 Operationalization: Students with disabilities can be identified by
16 the "disability" variable in the dataset. To test this hypothesis,
17 we can define the range for students with disabilities as those
18 who have a value of "Y" in the "disability" column.
19
20 Hypothesis 2: Students from lower socioeconomic backgrounds (indicated
21 by imd_band) will have worse predictive performance than average.
22 Operationalization: Socioeconomic background can be measured using the
23 "imd_band" variable in the dataset. To test this hypothesis, we
24 can define the range for students from lower socioeconomic
25 backgrounds as those who have an "imd_band" value in the range of
26 "0-30%".
27
28 Hypothesis 3: Students with lower levels of education (indicated by
29 highest_education) will have worse predictive performance than
30 average.
31 Operationalization: Education level can be determined using the "
32 highest_education" variable in the dataset. To test this
33 hypothesis, we can define the range for students with lower levels
34 of education as those who have a value of "Lower Than A Level" or
35 "No Formal quals" in the "highest_education" column.
36
37 Hypothesis 4: Students who have attempted the course multiple times (
38 indicated by num_of_prev_attempts) will have worse predictive
39 performance than average.
40 Operationalization: Previous attempts can be measured using the "
41 num_of_prev_attempts" variable in the dataset. To test this
42 hypothesis, we can define the range for students who have
43 attempted the course multiple times as those who have a value
44 greater than 1 in the "num_of_prev_attempts" column.
45
46 Hypothesis 5: Students in certain regions (indicated by region) will
47 have worse predictive performance than average.
48 Operationalization: Region can be determined using the "region"
49 variable in the dataset. To test this hypothesis, we can define
50 the range for students in certain regions as those who belong to
51 the regions of "North Region" or "Wales".
52
53 Column names: gender, region, highest_education, imd_band, age_band,
54 num_of_prev_attempts, studied_credits, disability, test, group_0,
55 group_1, group_2, group_3, group_4
56 Column values: dict_items([('gender', ['M', 'F']), ('region', ['South
57 Region', 'Scotland', 'Yorkshire Region', 'West Midlands Region', '
58 East Anglian Region', 'East Midlands Region', 'South East Region',
59 'Wales', 'North Western Region', 'London Region', 'South West
60 Region', 'Ireland', 'North Region']), ('highest_education', ['
61 Lower Than A Level', 'A Level or Equivalent', 'HE Qualification',
62 'No Formal quals', 'Post Graduate Qualification']), ('imd_band',
63 ['70-80%', '40-50%', '0-10%', '10-20%', '50-60%', '20-30%',
64 '60-70%', '80-90%', '90-100%', '30-40%', nan]), ('age_band',
65 ['0-35', '35-55', '55<=']), ('num_of_prev_attempts', [0, 1, 2, 3,
66 5, 4, 6]), ('studied_credits', {'min': 30, 'mean':
67 79.77122804326149, 'max': 655}), ('disability', ['N', 'Y']), ('
68 test', [False, True]), ('group_0', [False, True]), ('group_1', [
69 True, False]), ('group_2', [True, False]), ('group_3', [False,
70 True]), ('group_4', [False, True]))

```

```

27 """
28
29 -----OUTPUT TEXT-----
30 {
31     0: 'disability == "Y"',
32     1: 'imd_band in ["0-30%"]',
33     2: 'highest_education in ["Lower Than A Level", "No Formal quals"]
34     ',
35     3: 'num_of_prev_attempts > 1',
36     4: 'region in ["North Region", "Wales"]'
37 }

```

Code Listing 8: Interpreter: OULAD Dataset

C.4 Experimental details

We note that all experiments were performed using two compute resources: a server with NVIDIA RTX A4000 GPU and 18-Core Intel Core i9-10980XE, as well as an Apple M1 Pro 32GB RAM. We exemplify SMART Testing using GPT-4 [29] as the LLM but run further experiments to test the sensitivity to the type of language model in Appendix D.3.

C.4.1 Context-aware testing (Sec. 5.1.)

Goal. We aim to underscore the role of context in ML model testing to prevent false positives, especially when dealing with tabular data where data may contain many irrelevant or uninformative features [30], persisting even post-feature selection [31, 32]. We contrast SMART which explicitly accounts for context, in contrast to data-only approaches which are context-unaware only operating on the data.

Setup. We fit a predictive model to the training dataset, varying the number of irrelevant, synthetically generated features contained in the dataset — where irrelevant features are drawn from different distributions. We then quantify the proportion of conditions in the identified slices that falsely include the irrelevant synthetic synthetically features.

Because different methods are sensitive to different types of irrelevant features, we developed a data generating processes that encompasses many types of variables. Over many runs, different data-only methods pick up on some of these variables, showcasing that all methods are susceptible to randomly sampled irrelevant features in the dataset.

Sampling mechanism. To evaluate the impact of irrelevant features, we enrich the dataset by adding synthetic categorical variables. The number of new variables is equal to the number of existing features in the dataset. For each new variable x_i , we determine its type by sampling from a Bernoulli distribution with probability 0.5. If the sampled value is 0, x_i is a Bernoulli variable with success probability 0.1; otherwise, x_i is a categorical variable with four categories, following a predefined probability distribution (e.g., {0.1, 0.3, 0.4, 0.2}):

$$\begin{aligned} \text{Type}(x_i) &\sim \text{Bernoulli}(0.5), \\ x_i &\sim \begin{cases} \text{Bernoulli}(0.1) & \text{if } \text{Type}(x_i) = 0, \\ \text{Categorical}(0.1, 0.3, 0.4, 0.2) & \text{if } \text{Type}(x_i) = 1. \end{cases} \end{aligned}$$

We note that there are synthetic data generating processes that completely break other data-only methods. As an example, creating a unique ID column for each sample breaks the Autostrat algorithm, as all the subgroups/slices identified are the unique IDs. The data generating process employed in our experiment reflects a broad variety of commonly encountered DGPs.

C.4.2 Requirements-constrained testing (Sec. 5.2.)

We test whether each of the methods can fulfil three requirements.

The first requirement involved the use of the variable “age” in each detected slice. This was passed as an input to SMART. The other methods do not accept context as input and, therefore, it was not

possible to fulfil these requirements. The numbers provided for other methods are simply how often they fulfilled the requirements by chance.

The second and third requirements involved obtaining a minimum and maximum sample size. This was passed as an input to SMART and the ablated SMART version. Based on this, SMART changed its hyperparameter within its function which asks to indicate a minimum and maximum sample size for the discovered slices. As with the previous experiments, this was not adjusted for the other groups because they do not take textual input.

C.4.3 Targeting model failures (Sec. 5.3.)

In order to evaluate the targeting of model failures, we try four different tabular models with pre-specified hyperparameters. We find discrepant slices on the training dataset and evaluate them on the testing dataset.

C.4.4 Adaptive testing for a deployment environment (Sec. 5.4.)

The goal of the adaptive testing experiment is to understand the extent to which SMART, as well as other data-only methods, can use data in a source domain to generalize to a new, target domain where a covariate shift has been detected. To this end, the datasets provided have a known covariate shift and can be evaluated.

Each method was trained on the UK dataset and the discovered slices were evaluated on the US dataset. No additional context was provided to data-only methods since they do not accept any text or context as inputs.

In contrast, we have provided SMART with the previously discovered slices and hypotheses, and have asked to re-evaluate these hypotheses in the context of the US dataset. Specifically, SMART re-hypothesized possible model failures for the US market but used the UK data to operationalize the variables. The ablated version, $SMART_{NSF}$, achieved the best overall performance on the US market. The ablated version (i) did not have access to the UK data and the failures of the models; and (ii) operationalized each covariate using the LLM alone (refer to Sec. B.2 for a discussion on operationalizations with different SMART versions). This provides evidence that, in the presence of covariate shift, using inductive knowledge or domain expertise might be more useful to finding meaningful model failures.

C.4.5 Discovery of societally important groups and failure understanding (Sec. 5.5.)

As discussed in the main paper, SMART provides both possible hypotheses and justifications for model failures which can be evaluated using a simple “fit” method. Furthermore, SMART prioritizes meaningful data slices which are of societal importance. Such slices can be inspected for any data input.

D Additional experiments

D.1 Requirements-constrained testing

Goal. Requirements are a crucial, yet neglected part of ML model testing, such as verifying performance on societally relevant dimensions or verifying specific aspects to meet compliance requirements. However, *no previous testing framework* has incorporated the notion of satisfying requirements when defining the test slices. This experiment illustrates how SMART integrates requirements provided in natural language, which then influences the hypotheses generated to satisfy testing requirements.

Setup. We cover three real-world requirements that end-users might have: one based on demographics and two based on sample size. *Requirement 1:* Each of the top 10 identified unique slices should involve the age of a person. *Requirement 2:* The sample size of the top 10 identified unique slices must have at least 150 observations. *Requirement 3:* The sample size of the top 10 identified unique slices has to be small, within 10 and 150 observations. We exemplify the experiment using a real-world prostate cancer dataset from the UK [35], as healthcare often mandates certain testing requirements (e.g. Collins et al. [48]).

Analysis. Table 10 shows that SMART which directly integrates requirements (via natural language), satisfies the requirements a greater number of times compared to data-only methods, which only satisfies requirements by chance (hence the low number of times). Beyond satisfying requirements, the SMART slices also represent model failures that almost always have statistically significant performance differences from average when evaluated on test data. Finally, while $SMART_{NSF}$ (ablation without self-falsification) can satisfy requirements, the number of statistically significant slices is lower than SMART, thus underscoring the value of our self-falsification mechanism. That said, $SMART_{NSF}$ still outperforms data-only baselines.

Table 10: Requirement satisfaction showing how many times the top 10 generated slices satisfied the requirements (Req) and how many of these slices had statistically significantly different performance from average (Sig) on a testing dataset. Maximum is 10. \uparrow is better.

	R_1 : Age		R_2 : Min sample size		R_3 : Max sample size	
	Req	Sig	Req	Sig	Req	Sig
Autostrat	0.00 \pm 0.00	0.00 \pm 0.00	1.00 \pm 0.00	1.00 \pm 0.00	0.00 \pm 0.00	0.00 \pm 0.00
PSG_B	2.00 \pm 0.00	2.00 \pm 0.00	2.00 \pm 0.00	2.00 \pm 0.00	0.00 \pm 0.00	0.00 \pm 0.00
PSG_A	2.00 \pm 0.00	2.00 \pm 0.00	2.00 \pm 0.00	2.00 \pm 0.00	0.00 \pm 0.00	0.00 \pm 0.00
Divexplorer	3.20 \pm 2.11	0.95 \pm 1.56	0.00 \pm 0.00	0.00 \pm 0.00	2.70 \pm 2.57	1.05 \pm 1.75
Slicefinder	4.50 \pm 1.72	1.75 \pm 0.99	4.55 \pm 1.28	3.05 \pm 0.80	3.20 \pm 1.33	1.10 \pm 0.94
Sliceline	1.20 \pm 0.40	1.20 \pm 0.40	1.35 \pm 0.48	1.35 \pm 0.48	0.15 \pm 0.48	0.15 \pm 0.48
SMART_NSF	9.90 \pm 0.30	5.65 \pm 0.65	6.00 \pm 0.00	6.00 \pm 0.00	4.15 \pm 0.65	2.60 \pm 0.86
SMART	9.70 \pm 0.56	9.50 \pm 0.59	9.85 \pm 0.36	9.80 \pm 0.51	8.15 \pm 1.11	6.10 \pm 1.70

Takeaway 2. SMART, unlike data-only methods, identifies slices that have significant performance differences, whilst also satisfying requirements — an important dimension not even considered by previous testing methods.

D.2 Adaptive testing for a deployment environment

Goal. Deploying an ML model often entails going beyond IID, such as a different deployment environment. We consider the case of deploying a model to a different country where there is a covariate shift¹² and evaluate the capabilities of testing frameworks to adapt across the different environment and identify model failures.

Setup. We use real-world prostate cancer datasets from different country’s cancer registries with known distribution shifts: SEER (US) [11] and CUTRACT (UK) [35]. We train predictor f on UK data, while our target deployment environment is the US.

Analysis. ▶ *Identifying model failures.* Table 11 shows that SMART better tests models *at deployment time* using the information provided. SMART identifies a much greater number of statistically significant model failures (*almost all possible*), both within the same environment (UK) and when shifting to a different one (US), even after adjusting for multiple comparisons using Bonferroni correction.

▶ *Sample size sensitivity.* We also assess *sensitivity to sample size*, see Fig. 8. Both variants of SMART are shown to consistently outperform data-only counterparts in identifying a much greater number of significant model failures. *Within domain (UK):* as expected, we find that for lower sample sizes, $SMART_{NSF}$ (without the self-falsification mechanism) is superior, however, given enough data we then find that SMART benefits from the self-falsification.

Table 11: Number of slices identified (out of a maximum of 10) that had significantly divergent performance from average (higher is better). S_α counts the number of significantly divergent groups at $\alpha = 0.05$; $S_{\alpha/n}$ applies the Bonferroni correction. \uparrow is better.

	\mathcal{D}_{train}^{UK}		\mathcal{D}_{test}^{UK}		\mathcal{D}_{test}^{US}	
	S_α	$S_{\alpha/n}$	S_α	$S_{\alpha/n}$	S_α	$S_{\alpha/n}$
Autostrat	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	0.95 ± 0.22	0.95 ± 0.22
PSG _B	2.00 ± 0.00	2.00 ± 0.00	2.00 ± 0.00	2.00 ± 0.00	1.75 ± 0.43	1.50 ± 0.50
PSG _A	2.00 ± 0.00	2.00 ± 0.00	2.00 ± 0.00	2.00 ± 0.00	1.75 ± 0.43	1.45 ± 0.50
Divexplorer	1.65 ± 1.88	0.45 ± 0.92	2.00 ± 2.28	0.85 ± 1.53	3.90 ± 2.57	2.75 ± 2.21
Slicefinder	3.65 ± 0.96	2.75 ± 0.62	3.85 ± 1.11	2.70 ± 0.46	6.80 ± 1.03	5.95 ± 1.02
Sliceline	1.00 ± 0.00	1.00 ± 0.00	1.35 ± 0.48	1.35 ± 0.48	1.35 ± 0.48	1.35 ± 0.48
SMART _{NSF}	8.30 ± 0.46	8.00 ± 0.00	8.45 ± 0.50	8.00 ± 0.00	9.20 ± 0.60	8.35 ± 0.48
SMART	9.60 ± 0.49	9.25 ± 0.54	9.45 ± 0.50	8.85 ± 0.57	8.85 ± 0.79	8.35 ± 0.79

Deployment environment (US): we find that self-falsification similarly requires sufficient samples; which we note is expected behavior. Interestingly, in the deployment setting (US), $SMART_{NSF}$ generally identifies the greatest number of significant failures (*almost all possible*) across different sample sizes. This suggests that under covariate shift, using inductive knowledge (via the LLM) or domain expertise might be more useful to find meaningful model failures.

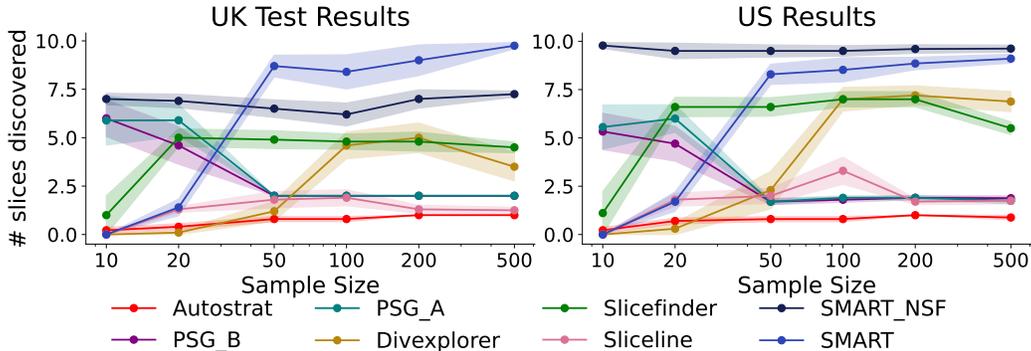


Figure 8: Number of significant groups discovered (out of a total of 10) based on the training dataset size. SMART can operate under any sample size; self-falsification mechanism requires a larger sample size to falsify hypotheses. \uparrow is better.

Overall, the results highlight the flexibility of SMART to handle different scenarios and sample sizes. From a practical perspective, while both SMART variants outperform data-only methods, the implication is that there is nuance in using different SMART variants for different scenarios.

Takeaway 4. SMART identifies more significant divergent failure slices in a deployment setting, outperforming data-only methods across environments and sample sizes.

¹² $p(X)$ changes, while $p(Y|X)$ remains the same

D.3 Effects of LLMs

In this section, we provide more experimental details which compare the effectiveness of two GPT models, GPT3.5 and GPT4. **We highlight that the goal is not to exhaustively test the framework with every LLM.** Rather, the goal is to showcase that SMART is feasible with at least the capabilities of GPT-4. We provide this section as a way to measure the sensitivity of the model’s performance with lower LLMs but highlight that we *do not* recommend using it with smaller LLMs, especially LLMs with fewer than 7B parameters.

D.3.1 Comparison over identified divergent slices

The following table reproduces the experiment from Sec. D.2 by directly comparing two models - GPT3.5 and GPT4. The setup is the same as in the original experiment.

Table 12: Number of slices identified (out of a maximum of 10) that had significantly divergent performance from average (higher is better). S_α counts the number of significantly divergent groups at $\alpha = 0.05$; $S_{\alpha/n}$ applies the Bonferroni correction.

	$\mathcal{D}_{\text{train}}^{\text{UK}}$		$\mathcal{D}_{\text{test}}^{\text{UK}}$		$\mathcal{D}_{\text{test}}^{\text{US}}$	
	S_α	$S_{\alpha/n}$	S_α	$S_{\alpha/n}$	S_α	$S_{\alpha/n}$
Autostrat	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	0.95 ± 0.22	0.95 ± 0.22
PSG _B	2.00 ± 0.00	2.00 ± 0.00	2.00 ± 0.00	2.00 ± 0.00	1.75 ± 0.43	1.50 ± 0.50
PSG _A	2.00 ± 0.00	2.00 ± 0.00	2.00 ± 0.00	2.00 ± 0.00	1.75 ± 0.43	1.45 ± 0.50
Divexplorer	1.65 ± 1.88	0.45 ± 0.92	2.00 ± 2.28	0.85 ± 1.53	3.90 ± 2.57	2.75 ± 2.21
Slicefinder	3.65 ± 0.96	2.75 ± 0.62	3.85 ± 1.11	2.70 ± 0.46	6.80 ± 1.03	5.95 ± 1.02
Sliceline	1.00 ± 0.00	1.00 ± 0.00	1.35 ± 0.48	1.35 ± 0.48	1.35 ± 0.48	1.35 ± 0.48
<i>SMART_NSF_GPT4</i>	8.30 ± 0.46	8.00 ± 0.00	8.45 ± 0.50	8.00 ± 0.00	9.20 ± 0.60	8.35 ± 0.48
<i>SMART_GPT4</i>	9.60 ± 0.49	9.25 ± 0.54	9.45 ± 0.50	8.85 ± 0.57	8.85 ± 0.79	8.35 ± 0.79
<i>SMART_NSF_GPT3.5</i>	8.20 ± 0.60	7.15 ± 0.65	8.20 ± 0.75	6.95 ± 0.67	9.25 ± 0.62	8.25 ± 0.43
<i>SMART_GPT3.5</i>	10.00 ± 0.00	9.85 ± 0.36	10.00 ± 0.00	9.75 ± 0.43	7.85 ± 0.48	7.15 ± 0.57

The table provides a measure of the model’s performance on the training dataset from the same environment ($\mathcal{D}_{\text{train}}^{\text{UK}}$), the testing dataset from the same environment ($\mathcal{D}_{\text{test}}^{\text{UK}}$), and a different deployment environment ($\mathcal{D}_{\text{test}}^{\text{US}}$).

Takeaway. Both GPT3.5 and GPT4 provide strong increases over benchmark methods with little variability between the two LLMs. One of the possible reasons why is that the hypothesis space of possible model failures is somewhat limited. This can be seen by the similar hypotheses that are generated by both GPT models.

D.3.2 Performance across different models

In this section, we vary different tabular machine learning model types and identify how well the ablated and original SMART, identified with GPT3.5 and GPT4, can identify slices with large performance discrepancies.

Table 13: The differences in accuracies between the top slice identified for each method on a testing dataset. The p-value computes the p-value associated with the difference in the accuracy. For the accuracy, higher values imply a greater ability to detect divergent slices (hence, higher is better). For the p-value, lower is better. Averages +- standard deviations are shown across 5 runs with random seeds and data splits

	Logistic Regression		SVM		XGBoost		Multi-layer Perceptron	
	$ \Delta Acc $	p-value	$ \Delta Acc $	p-value	$ \Delta Acc $	p-value	$ \Delta Acc $	p-value
SMART_NSF_GPT3.5	0.23 ± 0.03	0.00 ± 0.00	0.23 ± 0.03	0.00 ± 0.00	0.12 ± 0.06	0.14 ± 0.27	0.23 ± 0.03	0.00 ± 0.00
SMART_GPT3.5	0.34 ± 0.07	0.00 ± 0.00	0.34 ± 0.07	0.00 ± 0.00	0.28 ± 0.04	0.00 ± 0.00	0.34 ± 0.07	0.00 ± 0.00
SMART_NSF_GPT4	0.10 ± 0.01	0.00 ± 0.00	0.10 ± 0.01	0.00 ± 0.00	0.05 ± 0.04	0.25 ± 0.42	0.10 ± 0.01	0.00 ± 0.00
SMART_GPT4	0.40 ± 0.02	0.00 ± 0.00	0.40 ± 0.02	0.00 ± 0.00	0.29 ± 0.06	0.00 ± 0.00	0.40 ± 0.02	0.00 ± 0.00

SMART with deep learning models. SMART’s targeted sampling of hypotheses, is entirely independent of the downstream model used. i.e. SMART’s context-guided slice sampling mechanism is used to generate hypotheses independently of the downstream model.

We extend our analysis with Logistic Regression, SVM, XGBoost, and MLP to further include two tabular deep learning method: TabPFN and TabNet. As shown in Table 14, across all models SMART is the best at finding subgroups where the models are least reliable.

Table 14: Identifying slices with the highest performance discrepancies. We show differences in accuracies ($|\Delta Acc|$) between the top identified divergent slice and average performance across two state-of-the-art deep learning classifiers (over 5 runs) on the SEER dataset. \uparrow is better. 0.00 implies the evaluation method does not support the model.

Classifier	Evaluation Method						
	Autostrat	PSG_B	PSG_A	Divexplorer	Slicefinder	Sliceline	SMART
TabPFNClassifier	0.20 \pm 0.10	0.19 \pm 0.05	0.18 \pm 0.05	0.00 \pm 0.00	0.00 \pm 0.00	0.23 \pm 0.05	0.28 \pm 0.17
TabNet	0.10 \pm 0.09	0.10 \pm 0.04	0.10 \pm 0.04	0.02 \pm 0.04	0.00 \pm 0.00	0.17 \pm 0.08	0.17 \pm 0.12

Takeaway. GPT4 adds additional value when it comes to identifying slices with performance discrepancies. However, GPT3.5 is still able to find significant performance discrepancies across divergent slices, competitive and in many cases exceeding the performance of other methods (refer to Sec. D.4).

D.3.3 Hypothesis generation

A key component which determines the quality of the SMART method is the ability to navigate the search space by generating hypotheses. Therefore, we overview the kinds of hypotheses which are generated and how they differ based on the type of LLM used.

The table below provides an example of the top 10 hypotheses and justifications generated by GPT3.5 and GPT4 for the prostate cancer example in the UK, where the goal is to find specific slices where the model is likely to fail, and propose justifications for them.

Key takeaway. Many of the hypotheses for both models are similar. The underlying difference is in the ordering and justification of the hypotheses.

Table 15: Comparison of Hypotheses and Justifications generated by GPT4 and GPT3.5 for a breast cancer prediction task. The hypothesis target likely model failures.

GPT4		GPT3.5	
Hypothesis	Justification	Hypothesis	Justification
1 Patients with lower age	The model might perform worse for younger patients as prostate cancer is more common in older men, and there might be insufficient data for younger patients.	Subgroup of patients with older age (above 80) will have worse predictive performance due to higher likelihood of comorbidities and weaker immune systems, leading to higher mortality rates.	Older age is often associated with increased risk of comorbidities and weaker immune systems, which can negatively impact the prognosis of prostate cancer.
2 Patients with higher comorbidities	The model might perform worse for patients with multiple comorbidities as the complexity of their health condition might not be fully captured by the model.	Higher Gleason scores indicate more aggressive and advanced prostate cancer, which is generally associated with poorer prognosis and higher mortality rates.	Subgroup of patients with higher Gleason scores (4 or 5) will have worse predictive performance due to more aggressive and advanced cancer, leading to higher mortality rates.
3 Patients who underwent conservative management	The model might perform worse for these patients as the effectiveness of conservative management might vary greatly among individuals, and this variability might not be well captured by the model.	Subgroup of patients with stage 4 cancer will have worse predictive performance due to the advanced stage of the disease, indicating metastasis and limited treatment options.	Stage 4 prostate cancer is characterized by the spread of cancer to distant organs, indicating a more advanced and aggressive disease. The limited treatment options available for stage 4 cancer can result in worse predictive performance.
4 Patients with lower prostate-specific antigen levels	The model might perform worse for these patients as lower PSA levels might be associated with early-stage prostate cancer, which might be harder to predict due to lack of symptoms.	Patients with multiple comorbidities often have complex health conditions that can impact the effectiveness of prostate cancer treatment and management, leading to worse predictive performance.	Subgroup of patients with multiple comorbidities (e.g., 3 or 4) will have worse predictive performance due to the presence of other health conditions that may complicate the treatment and management of prostate cancer.
5 Patients with higher Gleason scores	The model might perform worse for these patients as higher Gleason scores indicate more aggressive cancer, which might progress unpredictably.	Subgroup of patients who received conservative management as their primary treatment option will have worse predictive performance due to the potential delay in receiving more aggressive treatments.	Conservative management may involve watchful waiting or active surveillance, which can result in delayed initiation of more aggressive treatments. This delay in treatment may lead to worse predictive performance.
6 Patients in stage 4	The model might perform worse for these patients as stage 4 prostate cancer is advanced and might progress unpredictably.	Primary hormone therapy is often used to control the growth of prostate cancer by reducing the levels of male hormones (androgens). Patients who do not receive this therapy may have uncontrolled cancer growth, resulting in worse predictive performance.	Subgroup of patients who did not receive primary hormone therapy will have worse predictive performance due to the potential lack of hormonal control over the cancer growth.
7 Patients who underwent radical prostatectomy	The model might perform worse for these patients as the outcomes of radical prostatectomy might vary greatly among individuals, and this variability might not be well captured by the model.	Subgroup of patients who did not undergo radical radiotherapy will have worse predictive performance due to the potential lack of localized cancer control.	Radical radiotherapy is a treatment option that aims to destroy cancer cells in the prostate gland. Patients who do not undergo this treatment may have a higher likelihood of localized cancer progression, leading to worse predictive performance.
8 Patients with lower Gleason scores	The model might perform worse for these patients as lower Gleason scores indicate less aggressive cancer, which might progress slower and be harder to predict.	Radical prostatectomy involves the surgical removal of the prostate gland. Patients who do not undergo this procedure may have residual cancer cells in the prostate, which can contribute to worse predictive performance.	Subgroup of patients who did not undergo radical prostatectomy will have worse predictive performance due to the potential presence of residual cancer cells in the prostate gland.
9 Patients in stage 1	The model might perform worse for these patients as stage 1 prostate cancer is early-stage and might be harder to predict due to lack of symptoms.	Subgroup of patients with higher prostate-specific antigen (PSA) levels (above 50) will have worse predictive performance due to the indication of more advanced cancer.	Higher PSA levels are often associated with more advanced prostate cancer. Patients with higher PSA levels may have a worse prognosis and higher mortality rates, leading to worse predictive performance.
10 Patients who underwent primary hormone therapy	The model might perform worse for these patients as the outcomes of primary hormone therapy might vary greatly among individuals, and this variability might not be well captured by the model.	The combination of older age, higher Gleason scores, and stage 4 cancer represents a subgroup with multiple negative prognostic factors. The cumulative effect of these factors is likely to result in the worst predictive performance.	Subgroup of patients with a combination of older age (above 70), higher Gleason scores (4 or 5), and stage 4 cancer will have the worst predictive performance due to the cumulative effect of advanced age, aggressive cancer, and metastasis.

D.4 Effects of different tabular machine learning models

The primary task of the SMART method is to evaluate a given, trained machine learning model. Thus far, we have been using a logistic regression model as the basis for evaluation in the main experiments. However, the results are *not* sensitive to the type of the tabular model. Therefore, in this section, we provide additional experiments where we vary the tabular model for the task. Specifically, we consider the following models initialized with their default hyperparameters for evaluation: Logistic Regression, Support Vector Machines, Boosting (implemented with XGBoost) and a multi-layer perceptron with 2 hidden layers and RELU activation functions in the hidden layers.

Goal. The primary goal is to understand whether the framework generalizes to other models which operate under different mapping mechanisms (e.g. a logistic regression, which is a linear model, compared to a tree-based model).

Setup. Given that not all the slice discovery or model evaluation algorithms output multiple slices, we constrain the evaluation to only focus on a single slice which might have discrepant performance. We use the UK prostate cancer dataset and evaluate the discrepancy of the top identified slice relative to the average across all identified model types. The discrepancy is calculated as the absolute differences of the average performance between the two groups, as well as the p-value associated with the difference. The results show the average performance \pm standard deviation of 5 random splits. A higher absolute difference indicates that the model fails on one of the slices more than average.

Discretizing the inputs. Many of the discovery methods, however, operate only on categorical data. The previously used dataset, however, has three continuous variables: age, prostate-specific antigen, and comorbidities. We therefore also assess the quality of these methods to discover slices on the testing dataset when these three variables are discretized into 10 bins each. The following

Table 16: The differences in accuracies between the top slice identified for each method on a testing dataset. The p-value computes the p-value associated with the difference in the accuracy. For the accuracy, higher values imply a greater ability to detect divergent slices (hence, higher is better). For the p-value, lower is better. Averages \pm standard deviations are shown across 5 runs with random seeds and data splits.

	Logistic Regression		SVM		XGBoost		Multi-layer Perceptron	
	$ \Delta Acc $	p-value	$ \Delta Acc $	p-value	$ \Delta Acc $	p-value	$ \Delta Acc $	p-value
Autostrat	0.24 ± 0.02	0.00 ± 0.00	0.24 ± 0.02	0.00 ± 0.00	0.09 ± 0.09	0.33 ± 0.46	0.24 ± 0.02	0.00 ± 0.00
<i>pysubgroup_beam</i>	0.23 ± 0.01	0.00 ± 0.00	0.23 ± 0.01	0.00 ± 0.00	0.11 ± 0.07	0.18 ± 0.40	0.23 ± 0.01	0.00 ± 0.00
<i>pysubgroup_apriori</i>	0.23 ± 0.01	0.00 ± 0.00	0.23 ± 0.01	0.00 ± 0.00	0.11 ± 0.07	0.19 ± 0.42	0.23 ± 0.01	0.00 ± 0.00
Divexplorer	0.05 ± 0.11	0.81 ± 0.43	0.09 ± 0.13	0.61 ± 0.53	0.14 ± 0.15	0.47 ± 0.49	0.02 ± 0.05	0.86 ± 0.32
Slicefinder	0.01 ± 0.00	0.00 ± 0.00	0.01 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.01 ± 0.01	0.00 ± 0.00
Sliceline	0.26 ± 0.06	0.00 ± 0.00	0.26 ± 0.06	0.00 ± 0.00	0.18 ± 0.09	0.00 ± 0.01	0.26 ± 0.06	0.00 ± 0.00
<i>SMART_NSF</i>	0.17 ± 0.01	0.00 ± 0.00	0.17 ± 0.01	0.00 ± 0.00	0.09 ± 0.05	0.07 ± 0.16	0.17 ± 0.01	0.00 ± 0.00
SMART	0.37 ± 0.03	0.00 ± 0.00	0.37 ± 0.03	0.00 ± 0.00	0.26 ± 0.06	0.00 ± 0.00	0.37 ± 0.03	0.00 ± 0.00

table reports the performance of all the methods (note: we did not discretize the dataset for SMART because it can work natively on continuous data).

Table 17: The differences in accuracies between the top slice identified for each method on a testing dataset when the datasets continuous features are discretized.

	Logistic Regression		SVM		XGBoost		Multi-layer Perceptron	
	$ \Delta Acc $	p-value	$ \Delta Acc $	p-value	$ \Delta Acc $	p-value	$ \Delta Acc $	p-value
Autostrat	0.02 ± 0.02	0.38 ± 0.27	0.02 ± 0.02	0.45 ± 0.37	0.02 ± 0.02	0.52 ± 0.32	0.02 ± 0.02	0.43 ± 0.33
<i>PSG_B</i>	0.01 ± 0.01	0.56 ± 0.29	0.01 ± 0.01	0.62 ± 0.38	0.01 ± 0.01	0.55 ± 0.25	0.01 ± 0.01	0.57 ± 0.23
<i>PSG_A</i>	0.01 ± 0.01	0.55 ± 0.29	0.01 ± 0.01	0.61 ± 0.38	0.01 ± 0.01	0.55 ± 0.26	0.01 ± 0.01	0.59 ± 0.26
Divexplorer	0.10 ± 0.08	0.47 ± 0.28	0.11 ± 0.10	0.41 ± 0.36	0.13 ± 0.11	0.38 ± 0.35	0.13 ± 0.10	0.38 ± 0.37
Slicefinder	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
Sliceline	0.09 ± 0.01	0.05 ± 0.06	0.08 ± 0.02	0.14 ± 0.09	0.11 ± 0.07	0.25 ± 0.24	0.09 ± 0.01	0.06 ± 0.08
<i>SMART_NSF</i>	0.17 ± 0.01	0.00 ± 0.00	0.17 ± 0.01	0.00 ± 0.00	0.09 ± 0.05	0.07 ± 0.16	0.17 ± 0.01	0.00 ± 0.00
SMART	0.37 ± 0.03	0.00 ± 0.00	0.37 ± 0.03	0.00 ± 0.00	0.26 ± 0.06	0.00 ± 0.00	0.37 ± 0.03	0.00 ± 0.00

Takeaway. SMART is able to consistently identify the greatest performing slices across a number of different tabular models.

D.5 On the inductive biases of ML model testing

Goal. We further observe that data-only testing methods implicitly assume the existence of slices with discrepancies in performance. While indeed ML models do fail — it is equally as problematic to highlight failures where there are none.

Setup. To evaluate this we propose a fully synthetic setup. Here, both the dependent variable (Y) and the independent variables (\mathbf{X}) are sampled from a predefined random distribution. Specifically, we predict loan default ($Y \in \{0, 1\}$) based on a set of independent variables $\mathbf{X} = \{N_{\text{runs}}, M_{\text{pref}}, A_{\text{rainfall}}, F_{\text{color}}, P_{\text{season}}\}$, which are conceptually and empirically independent of the outcome of interest. Ideally, if we account for context we should be able to identify these disparate features should not influence the account and hence without prior relationships we should not flag spurious slices.

We compare the data-only methods to SMART under three different data generating processes (scenarios) that capture diverse underlying dynamics denoted as $\mathcal{S}_{\text{uniform}}$, $\mathcal{S}_{\text{skewed}}$, and $\mathcal{S}_{\text{interactions}}$, where each DGP has a focus on uniform, skewed, and interactive effects, respectively.

The first scenario is given by the variables sampled from the following DGPs:

$$\begin{aligned}
N_{\text{runs}} &\sim \text{Uniform}\{1, 499\}, \\
M_{\text{pref}} &\sim \text{Uniform}\{1, 5\}, \\
A_{\text{rainfall}} &\sim \text{Uniform}\{20000, 99999\}, \\
F_{\text{color}} &\sim \text{Uniform}\{1, 6\}, \\
P_{\text{season}} &\sim \text{Uniform}\{0, 3\}, \\
Y &\sim \text{Uniform}\{0, 1\}.
\end{aligned}$$

The second scenario is given by the variables sampled from the following DGPs:

$$\begin{aligned}
N_{\text{runs}} &\sim \text{Uniform}\{1, 499\}, \\
M_{\text{pref}} &\sim \text{Binomial}(1, 0.5), \\
A_{\text{rainfall}} &\sim \text{Categorical}(0.1, 0.3, 0.4, 0.2), \\
F_{\text{color}} &\sim \text{Binomial}(1, 0.1), \\
P_{\text{season}} &\sim \text{Binomial}(1, 0.05), \\
Y &\sim \text{Uniform}\{0, 1\}.
\end{aligned}$$

The third scenario is given by the variables sampled from the following DGPs:

$$\begin{aligned}
N_{\text{runs}} &\sim \text{Uniform}\{1, 499\}, \\
M_{\text{pref}} &\sim \text{Binomial}(1, 0.5), \\
A_{\text{rainfall}} &\sim \text{Categorical}(0.1, 0.3, 0.4, 0.2), \\
A_{\text{music_hap}} &= M_{\text{pref}} \times A_{\text{rainfall}}, \\
A_{\text{run_hap}} &= N_{\text{runs}} \times M_{\text{pref}}, \\
Y &\sim \text{Uniform}\{0, 1\}.
\end{aligned}$$

Analysis. Table 18 shows the number of slices spuriously discovered, while Table 19 outlines the number of conditions within the slices. We can clearly see the pitfalls of data-only approaches which detect slices which in reality have no relation to one another — often surfacing few conditions per group which suggests they arise by chance. The rationale for this failure is simply because data-only approaches do not and cannot reason about the features and/or understand context and simply aim to find slices with discrepancies in performance — which of course could arise by chance. In contrast, we see that SMART by virtue of context-awareness can avoid surfacing groups — which in reality have no relationships.

Table 18: Number of discovered slices on a synthetic dataset with no prior relationships in three data generating process scenarios. slices capped at most 20. Average of 50 runs \pm standard deviations is shown.

Method	$\mathcal{S}_{\text{uniform}}$	$\mathcal{S}_{\text{skewed}}$	$\mathcal{S}_{\text{interactions}}$
Autostrat	1.00 \pm .0	1.00 \pm .0	1.00 \pm .0
<i>PSG_B</i>	20.00 \pm .0	20.00 \pm .0	20.00 \pm .0
<i>PSG_A</i>	20.00 \pm .0	20.00 \pm .0	20.00 \pm .0
divexplorer	20.00 \pm .0	20.00 \pm .0	20.00 \pm .0
slicefinder	20.00 \pm .0	20.00 \pm .0	20.00 \pm .0
SMART	0.00 \pm .0	0.00 \pm .0	0.00 \pm .0

D.6 Context aware sensitivity

We provide an additional experiment where we vary the sample size in the training dataset and observe how that affects the number of slices discovered for each method. We show that the SMART is not affected by irrelevant features regardless of the sample size of the training dataset. The result is shown in Figure 9.

Table 19: Number of conditions per discovered slice (false positives) in three data generating process scenarios. Average of 50 runs +- standard deviations is shown. Lower is better.

Method	$\mathcal{S}_{\text{uniform}}$	$\mathcal{S}_{\text{skewed}}$	$\mathcal{S}_{\text{interactions}}$
Autostrat	2.17 ± 0.46	1.73 ± 1.01	1.13 ± 0.35
PSG_B	1.03 ± 0.18	1.90 ± 0.71	1.27 ± 0.45
PSG_A	1.03 ± 0.18	1.90 ± 0.71	1.27 ± 0.45
divexplorer	2.10 ± 0.31	2.87 ± 0.68	2.40 ± 0.50
slicefinder	1.40 ± 0.50	1.50 ± 0.57	1.00 ± 0.00
SMART	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00

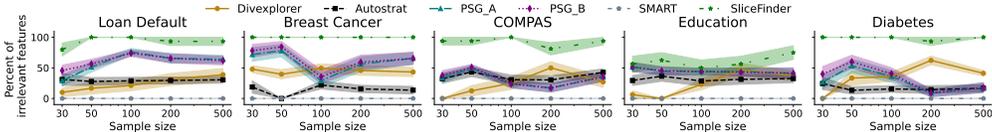


Figure 9: Proportion of irrelevant features (y) for each slice discovery method, based on the sample size. Lower is better.

D.7 Cost of SMART

We assess the cost of LLM hypothesis generation and scalability to larger datasets. Specifically, we demonstrate not only that SMART is cheap but also easily scalable to large datasets

- **Scalability:** SMART’s scalability depends on the number of hypotheses generated, not dataset size (unlike data-only methods). This allows SMART to easily scale to arbitrarily large datasets.
- **Cost Analysis:** In practical terms, cost then also scales primarily with the number of hypotheses generated, not dataset size. We provide a rough estimate based on token counts of input and outputs for 2 datasets (SEER and OULAD) in Table 20. This would be less than 0.1 USD for 5 hypotheses and less than 0.5 USD for 100 hypotheses for state-of-the-art

Table 20: Cost of SMART (USD) for different GPT LLMs and different numbers of hypotheses generated. The cost is estimated based on token counts. Note: GPT-4o models are post our paper and are even cheaper.

Model	Cost SEER		Cost OULAD	
	5 Hypothesis (USD)	100 Hypothesis (USD)	5 Hypothesis (USD)	100 Hypothesis (USD)
GPT-4	0.017	0.249	0.022	0.316
GPT-3.5	0.004	0.050	0.005	0.064
GPT-4o Mini	0.0003	0.005	0.0005	0.006
GPT-4o	0.008	0.125	0.011	0.158

D.8 SMART with open-weight models

SMART ideally should be used with the most capable LLM possible. That said, we assess the differences in hypotheses between open-weight models and GPT-4.

We assess Mistral-7b, Qwen-1.5-7b, Llama-3-8b, Llama-70b, where for the OULAD and SEER datasets we generate 5 hypotheses and assess overlap to the hypotheses generated by GPT-4. This is presented in Table 21.

To summarize, the overlap between open-source models and GPT-4 is between 60-80%. We find that open-source models propose similar hypotheses, but they are not replacements for more capable models. This highlights that less capable models might propose similar hypotheses, yet they still catch fewer model failures.

D.9 Understanding the importance of feature names

SMART uses the implicit context encoded in the interpretable feature names as a source of contextual information to guide hypothesis generation. For instance, in a medical dataset, features with names

Table 21: Comparison hypotheses by GPT-4 and overlap w/ open-weight models

Dataset	Factors (GPT Hypotheses)	Mistral-7b	Llama 3-8b	Qwen 1.5-7b	Llama 70b
Outlad Dataset	Disability	✓	✓	✓	✓
	IMD band	✓	✓	✓	✓
	Age		✓	✓	✓
	Number of previous attempts	✓	✓	✓	
	Test (boolean)				
	Oulad Overlap Percentage	60%	70%	60%	60%
SEER Dataset	Age		✓	✓	✓
	Prostate-specific antigen (PSA)				
	Comorbidities	✓	✓	✓	✓
	Treatment (conservative management)	✓	✓	✓	✓
	Cancer stage			✓	✓
	SEER Overlap Percentage	40%	60%	80%	80%

like age, sex, or patient covariate features provide context to guide LLM hypothesis generation. This contrasts with data-only approaches which only use the numerical data values alone and ignore the context surrounding the feature names.

We aim to assess the sensitivity to interpretable feature names to provide guidance on the use of SMART. First, we perform a qualitative study where we limit the data schema by hiding the feature names (such that they become uninformative) and inspect the hypotheses and justifications generated. We find that in the limited-schema case, SMART generates hypotheses based on inferences about the feature information (e.g. "the model might fail on feature_4 if feature_4 represents gender"). In contrast, informative names guide meaningful hypothesis generation. Such hypotheses and justifications are illustrated in Table 22.

Second, we evaluate whether limiting the data schema by hiding some feature names and leaving minimal external context affects detection rates of model failures. We compare two versions of SMART, original and with corrupted feature labels, in identifying data slices with high performance discrepancies from average (Fig. 10). We find that across two real-world private datasets, hiding the feature names hinders model evaluation. This highlights that feature names play an important role in finding model failures.

These results highlight while SMART does not rely on any additional feature descriptions, feature names play an important role in finding model failures, just as any human requires interpretable feature names to understand the data. That said, feature names (e.g. column labels such as sex, age, race etc) are present in almost all tabular datasets both in the research field and in the real world where data is stored in SQL tables with column names.

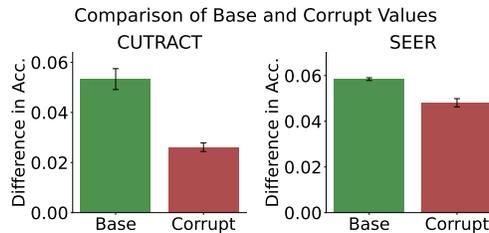


Figure 10: Identifying the importance of feature names as a source of information for context-aware testing. We compare two versions of SMART on CUTRACT and SEER datasets. The first version, “base”, is the normal SMART method. The second version, “corrupt”, has feature names changed to uninformative ones. The goal is to identify slices with the highest performance discrepancies between the average & subgroup prediction, calculated as the difference in their respective accuracies (y-axis). We see showing feature names enables to find subgroups that have a higher performance discrepancy from baseline performance than without feature names. ↑ is better.

Table 22: Example hypotheses and justifications when dataset column names are hidden. The hypotheses and justifications for features that do not encode meaningful information lose practical meaning. Context-aware testing benefits from rich feature names that can be used to guide meaningful hypothesis generation.

Model type	Hypothesis	Justification
Base	The model will perform worse on predicting the risk of disease for older individuals.	Age is a significant factor in many diseases, and older individuals often have more complex health situations with multiple comorbidities. This complexity could make it more difficult for the model to accurately predict disease risk.
Base	The model will perform worse on predicting the risk of disease for individuals with lower Gleason scores.	The Gleason score is a grading system used to determine the aggressiveness of prostate cancer. Lower scores indicate less aggressive cancer, which may be more difficult to predict due to its slower progression and less obvious symptoms.
Corrupted	{'1': ['feature_0']}	If feature_0 represents a characteristic such as race or ethnicity, the model might perform worse for certain subgroups due to societal biases. For instance, if the dataset is biased towards a particular race or ethnicity, the model's predictions for other races or ethnicities might be less accurate.
Corrupted	{'2': ['feature_4']}	If feature_4 represents a characteristic such as gender, the model might perform worse for certain subgroups due to societal biases. For example, if the dataset is biased towards one gender, the model's predictions for the other gender might be less accurate.

D.10 Example of a model report

We show an example of a model report which is automatically generated by SMART. A model report includes a summary of key hypotheses, justifications, sources, and summary findings as a written report. A model report fully satisfies the requirements of the user.

In this case, the model is generated to directly match latex output requirements. It is provided in the shaded box below as a part of the US prostate cancer (SEER) [11] evaluation.

A report on the performance of the model has been concluded. The following are the hypotheses, their justifications tested on the model with their conclusions on whether the hypothesis was supported.

Hypothesis	Justification	Operationalization	Hypothesis Supported
The model may perform worse for older patients	Older patients may have more comorbidities and complex health situations that are not fully captured by the dataset. Additionally, societal biases may lead to less aggressive treatment options being pursued for older patients, which could affect the model's predictions.	$age > 75$	Yes
The model may perform worse for patients with lower prostate-specific antigen levels	Lower levels of prostate-specific antigen may be associated with earlier stages of prostate cancer, which may be harder to predict due to less data and less obvious symptoms.	$prostate_specific_antigen < 10$	Yes
The model may perform worse for patients who have undergone conservative management	Conservative management is a less aggressive form of treatment, which may be chosen due to a variety of factors not captured in the dataset, such as patient preference or other health considerations. This could introduce additional complexity into the model's predictions.	$treatment_conservative_management == 1$	Yes
The model may perform worse for patients with a higher number of comorbidities	Patients with more comorbidities may have more complex health situations that are not fully captured by the dataset. Additionally, these patients may be more likely to die from causes other than prostate cancer, which could confuse the model's predictions.	$comorbidities > 2$	No
The model may perform worse for patients with a higher Gleason score	A higher Gleason score indicates more aggressive cancer, which may be harder to predict due to its rapid progression and the potential for other health factors to influence outcomes.	$gleason_score == 4$	No

A further analysis was run on the training and testing datasets to understand the effects of each identified slice. The following are two tables which summarize a part of this information:

	group_size	support	p_value_bootstrap	num_criteria	outcome_diff	accuracy_diff
H0	4915.00	0.31	0.00	1.00	0.23	0.05
H1	7749.00	0.48	0.00	1.00	0.26	0.01
H2	1456.00	0.09	0.00	1.00	0.06	0.09
H3	574.00	0.04	0.94	1.00	0.08	0.00
H4	94.00	0.01	0.14	1.00	0.36	0.09

	odds_ratio_outcome	odds_ratio_acc	lift_outcome	lift_acc	weighted_relative_y	weighted_relative_acc
H0	0.82	1.39	1.46	0.94	0.07	-0.02
H1	0.96	1.12	0.48	0.99	-0.13	-0.01
H2	0.99	1.41	1.11	0.89	0.01	-0.01
H3	0.98	1.01	0.85	1.00	-0.00	-0.00
H4	0.48	0.55	1.72	1.10	0.00	0.00

Recommendations:

1. The model appears to be less reliable for older patients ($age > 75$), patients with lower levels of prostate-specific antigen (< 10), and those who have undergone conservative management treatment. It also shows decreased performance for patients with more than two comorbidities and those with a Gleason score of 4. However, the model is more reliable when these conditions are not met.
2. Before deploying the model, the end user should be aware of the following:
 - The model's performance may be compromised for older patients and those with multiple comorbidities. Consider additional validation or alternative models for these groups.
 - Patients with lower prostate-specific antigen levels and those who have undergone conservative management treatment may also experience less accurate predictions. Additional clinical insights may be needed for these cases.
 - Although the model shows decreased performance for patients with a Gleason score of 4, this group is relatively small, so the impact on overall model performance may be limited. However, caution should be exercised when interpreting results for these patients.
 Remember, these recommendations are based on the training and test datasets from the UK. If deploying in a different geographical context, consider revalidating the model with local data.

Definitions of the metrics:

- **Group Size:**

$$\text{group_size} = |\text{slice}|$$

- **Support:**

$$\text{support} = \frac{|\text{slice}|}{|\text{dataset}|}$$

- **Number of Criteria:**

$$\text{num_criteria} = \text{Count}(\text{"and"}) + 1$$

- **Outcome Difference:**

$$\text{outcome_diff} = |\text{avg_outcome_dataset} - \text{avg_outcome_slice}|$$

- **Accuracy Difference:**

$$\text{accuracy_diff} = |\text{accuracy_dataset} - \text{accuracy_slice}|$$

- **Odds Ratio (Outcome):**

$$\text{odds_ratio_outcome} = \frac{p_1(1 - p_1)}{p_0(1 - p_0)}$$

- **Odds Ratio (Accuracy):**

$$\text{odds_ratio_acc} = \frac{p_1(1 - p_1)}{p_0(1 - p_0)}$$

(where p_1 and p_0 are accuracies in the slice and the rest of the dataset, respectively)

- **Lift (Outcome):**

$$\text{lift_outcome} = \frac{p_1}{p}$$

- **Lift (Accuracy):**

$$\text{lift_acc} = \frac{p_1}{p}$$

(where p_1 is accuracy in the slice and p is accuracy in the entire dataset)

- **Weighted Relative Outcome:**

$$\text{weighted_relative_outcome} = \text{support} \times \text{diff_outcomes}$$

- **Weighted Relative Accuracy:**

$$\text{weighted_relative_accuracy} = \text{support} \times \text{diff_accuracy}$$

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope?

Answer: [Yes]

Justification: The abstract accurately reflects the claims made in the paper. Our paper introduces the new paradigm of context-aware testing which we discuss in Sec. 3. We introduce our context-aware testing instantiation SMART in Sec. 4. We experimentally show how under a variety of different experimental conditions SMART outperforms data-only benchmarks in Sec. 5. We therefore believe the provided evidence fully supports the claims that context-aware testing is a new alternative to testing ML models relative to data-only methods.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We discuss the limitations of our work in Sec. 6. We also provide ways of addressing these limitations in Sec. 5.4 with an extended discussion on possible biases, mitigation strategies, and best practices, including how SMART provides unique opportunities to address such challenges with features such as model reports or transparent hypothesis testing and operationalization.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.

- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: The paper makes a claim that context can be used to guide the search for relevant and meaningful model failures via a context-guided slice sampling mechanism which uses context to prioritize likely slices. This is defined and discussed in Sec. 3.3.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can appear in either the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The paper provides the key experimental setups in the main experimental section (Sec. 5) together with additional experimental details in the appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.

- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We provide details about the algorithms and data in Appendix C, Sec. B and Sec. D. Code can be found at: https://github.com/pauliusrauba/SMART_Testing or https://github.com/vanderschaarlab/SMART_Testing

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: All the details on the experiments are either provided within Section 5, with full details provided in Appendix D.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Error bars (standard deviation) are included as relevant over multiple seeds for the experiments in Section 5 and additional experiments in Appendix D.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: All the compute details on the experiments are provided in Appendix C.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We have read the code of ethics do not violate any of the dimensions.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: Broader impacts of testing and implications are outlined in Sec 1 and 6. We discuss potential impacts and mitigations in Sec 5.4 and Sec 6.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Not applicable — our paper presents a new method for reliable and safe ML model testing.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: Appendix C provides details and citations for all assets (data and baselines) used in the paper.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not produce new assets such as datasets, but uses existing benchmarks/datasets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: We do not have crowdsourcing experiments or research with humans.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: We do not have crowdsourcing experiments or research with humans that would need an IRB.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.