

415 A Failure Modes

416 The design of the prompt plays a pivotal role in our entire process. Specifying instructions precisely
417 can create a significant difference, whether it comes to effectively describing tabular data, generating
418 reliable summaries or inferring accurate predictions (shown in Figures 1 and 2). While soft
419 prompting has been successful at instructing LLMs [23], it cannot be applied to our setting because
420 our classification algorithm learns a human-level summary as a prompt for classifying data, rather
421 than a soft prompt. Instead we chose to write prompts that ask the LLM to perform these tasks. In
422 this way, our prompting method is entirely gradient-free. Hand-written prompts also offer flexibility,
423 aligning well with our core methodology for creating weak learner. While carefully handcrafting
424 prompts this way might seem intensive, we will show that once we identify the ideal hyperparameter
425 settings they can be framed with little effort.

426 A.1 Data Conversion Challenges

427 Language models (LLMs) have demonstrated impressive performance on standard tasks with minimal
428 supervision [42, 2]. However, for converting tabular data to text, there were several considerations to
429 meet the requirements of our task. As highlighted in Section 3.1, we will refer to these texts as *data*
430 *descriptions*.

431 **Ensuring Uniform Length** Firstly, the data descriptions should not be too long or short, also be of
432 comparable length. Excessively long descriptions limit the number of examples that can be fit inside
433 the prompt and summarized. We observed in Figure 4 (right bottom) that the summary performance
434 also scales with more examples, so it makes sense to have descriptions of approximately uniform
435 length.

436 A straightforward way of achieving this uniformity would be by specifying a max word length as
437 part of the conversion prompt itself, as in “Describe in not more than 80 words”. However, we found
438 this approach can falter, sometimes leading to overly simplistic descriptions like “These are annual
439 spendings of a customer.” (in the `wholesale-customers` dataset).

440 Consequently, we adopt more nuanced strategy by first modifying the prompt with the terms “con-
441 cisely” and “accurately” to emphasize the brevity and preciseness of the generated descriptions
442 (shown in Figure 1). Then, we implement a resampling strategy, that generates descriptions until find-
443 ing the one with a desired length ranging between 20 to 80 words. This process achieves consistent
444 and uniformly long descriptions.

445 **Including Metadata** Prepending metadata to the prompt enhances the contextual awareness of the
446 task, resulting in higher-quality descriptions (shown in Figure 1).

447 **Separating Features from Labels** In our method, the data descriptions function dual role, both
448 as training examples and as query for inferring class labels. This suggests that, when converting
449 data to text, the features need to be described separately from the target label as illustrated in Figure
450 1. The resulting strings are then concatenated to form the data description. Instead, if the whole
451 tabular record were passed to the LLM, it often produces texts that assimilate the classification label
452 information in the meat of the description itself, rendering it difficult to extract a query for doing
453 inference.

454 While one might cleverly come up with prompts that can allow the LLM to describe the features
455 and target label in separate sentences, we found it to be more sensible to supply just the features for
456 describing and not reveal any information about the target task. Sometimes that can liberate the LLM
457 to hallucinate some facts about the task and form biased data to begin with.

458 **Natural-Sounding Descriptions** While the LLM generates a different-styled response every time,
459 to explicitly ensure that the generated descriptions are not template-like by chance, add a directive
460 at the end of the prompt: “Use your creativity”. This encourages the LLM to produce more natural
461 narratives of the record. Alternatively, setting a higher temperature during decoding achieves a similar
462 effect.

463 A.2 Summarization

464 There are several aspects worth considering that can contribute to high-quality summaries.

465 **Sampling candidate summaries** A well-crafted summary is a one that captures salient information
466 of the dataset, in a way that facilitates inferring predictions off it. However, the process of generating
467 summary using a LLM is inherently stochastic due to temperature sampling, as a result, the generated
468 summary can be noisy. From our experiments with tuning this temperature, we found 0.80 to be
469 ideal through Bayesian optimization. Even at this value, on average only 1 out of 3 summaries were
470 meaningful.

471 A noisy summary can be distinguished quite easily. For instance, on the `vehicle` dataset, the `tl;dr`
472 prompt elicits summaries as naive as “The given data describes a bus, van, or saab silhouette.” or
473 “The data in this table identifies a vehicle as a bus, saab, opel, or van. The compactness, circularity,
474 distance circularity, radius ratio, hollows ratio, and symmetry are all predictive of the vehicle’s type.”
475 which does not offer actionable insight.

476 This observation indicates that summaries need to be sampled quite a few times and the best one
477 can be determined based on the validation error. As a result, for the `Summary` learning procedure
478 in Section 3.2, we resample approximately 25 times to find a good summary. Also, given that our
479 datasets are small, it is not unusual for the summaries to have the same validation error. When tied,
480 we pick one having a higher training error rate, i.e. lower generalization gap.

481 Differently, in our `Summary` boosting procedure explained in Section 3.3, we resample only until
482 finding a summary whose training error is better than random guessing and return immediately.

483 **Ordering examples inside the *summarization* prompt** Unlike gradient descent, prompting is not
484 robust to the presentation of the examples to the learning algorithm. While we show via ablation
485 studies in Section 5 that there is no statistically significant difference in performance between either
486 shuffling the examples or listing them by class, we can generally expect that depending on the dataset
487 and the number of target classes, one might be preferred over the other.

488 For instance, in a multi-class setting, listing examples by class might be more helpful in reaching a
489 weak learner quickly. However, in a two-class setting, the summary might actually benefit from the
490 randomness in shuffled examples.

491 **Customizing the *summarization* prompt** The approach of asking the LLM to summarize examples
492 can also give rise to good/bad summaries. For instance, one can prompt the LLM with a simple
493 `tl;dr` or specify the task more elaborately. We will refer to the latter option as `explicit`. As we
494 demonstrated in Figure 4 (left), both are means to the goal and do not statistically differ in terms of
495 performance induced.

496 However, in our experiments on certain datasets, we would rather be incentivized choosing the
497 `explicit` over the `tl;dr` to attain a weak learner more quickly. This choice becomes important
498 purely for compute reasons as it will take relatively lesser resampling, while the `tl;dr` still works. For
499 instance, this scenario can happen when the LLM cannot decipher what the summary is supposed to
500 say, by just observing the examples. As examples, the `tl;dr` prompt suffices on datasets such as `iris`,
501 `diabetes`, and `wine` that are commonly encountered in prediction context, whereas the LLM might
502 not be very familiar with the goals of `vertebra-column` or `somerville-happiness-survey`
503 data, necessitating the use of the `explicit` prompt. For these other datasets, the fact that it is a
504 classification problem based on some features and target classes may not be very apparent from just
505 the examples and metadata. So, providing a directive such as “Summarize in detail how we can tell
506 apart people with normal and abnormal `vertebra-column`” reduces ambiguity in the task setup and
507 reduces probability of a noisy summary.

508 While manual intervention is necessary, framing this prompt can be done with little effort. We provide
509 a comprehensive list of these parameters for all datasets in Table 3.

510 **Including Metadata** Similar to data conversion, including meta-data information in the prompt
511 offers better penetration into the world of the dataset, as a result improves boosting performance.

Table 3: **Prompt design:** Prompt parameter settings for every dataset.

Dataset	Prompting hyperparameters
caesarian	<p><i>metadata:</i> This dataset contains information about caesarian section results of 80 pregnant women with the most important characteristics of delivery problems in the medical field. The goal is to predict whether a woman will undergo normal or caesarian delivery.</p> <p><i>classes:</i> [normal, caesarian]</p> <p><i>summary directive:</i> Tl;dr</p> <p><i>inference directive:</i> Hence this woman's delivery mode is likely to be (normal or caesarian):</p>
iris	<p><i>metadata:</i> This is the iris dataset, perhaps the best known database to be found in the pattern recognition literature. Fisher's paper is a classic in the field and is referenced frequently to this day. (See Duda & Hart, for example.) The data set contains 3 classes of 50 instances each, where each class refers to a type of iris plant. One class is linearly separable from the other 2; the latter are NOT linearly separable from each other. Predicted attribute- class of iris plant- setosa, versicolor, virginica</p> <p><i>classes:</i> [setosa, versicolor, virginica]</p> <p><i>summary directive:</i> Tl;dr</p> <p><i>inference directive:</i> Based on the above information, predict if this flower will be classified as setosa, versicolor, virginica</p>
tae	<p><i>metadata:</i> The data consist of evaluations of teaching performance over three regular semesters and two summer semesters of 151 teaching assistant (TA) assignments at the Statistics Department of the University of Wisconsin-Madison. The scores were divided into 3 roughly equal-sized categories ("low", "medium", and "high") to form the class variable.</p> <p><i>classes:</i> [low, medium, high]</p> <p><i>summary directive:</i> Tl;dr</p> <p><i>inference directive:</i> Predict whether this class will score low or medium or high:</p>
glass	<p><i>metadata:</i> This is the glass dataset from USA Forensic Science Service; 6 types of glass; defined in terms of their oxide content (i.e. Na, Fe, K, etc). The study of classification of types of glass was motivated by criminological investigation. At the scene of the crime, the glass left can be used as evidence...if it is correctly identified!</p> <p><i>classes:</i> [building_windows_float_processed, building_windows_non_float_processed, vehicle_windows_float_processed, containers, tableware, headlamps]</p> <p><i>summary directive:</i> Tl;dr</p> <p><i>inference directive:</i> There are 6 possible type of glass: building_windows_float_processed, building_windows_non_float_processed, vehicle_windows_float_processed, containers, tableware, headlamps. Predict which one will this sample be:</p>
breast-cancer	<p><i>metadata:</i> This is one of three domains provided by the Oncology Institute that has repeatedly appeared in the machine learning literature. This data set includes 201 instances of one class and 85 instances of another class. The instances are described by 9 attributes, some of which are linear and some are nominal. It contains information about women that had a recurrence or non-relapse of breast cancer after their first time.</p> <p><i>classes:</i> [recurrence, non-relapse]</p> <p><i>summary directive:</i> Based on the above examples, figure out under what conditions will a woman have recurrence or non-relapse of breast cancer?</p> <p><i>inference directive:</i> Predict whether this woman will have a recurrence or non-relapse:</p>
visualizing-environmental	<p><i>metadata:</i> This is the visualizing-environmental dataset, one of the 22 data sets from the book Visualizing Data published by Hobart Press (books@hobart.com). This data describes indicators for a positive/negative environment based on ozone, radiation and temperature.</p> <p><i>classes:</i> [positive, negative]</p> <p><i>summary directive:</i> Tl;dr</p> <p><i>inference directive:</i> There are clear signs of this environment being (positive or negative):</p>
analcadata-chlamydia	<p><i>metadata:</i> This chlamydia dataset is one of the data sets used in the book "Analyzing Categorical Data" by Jeffrey S. Simonoff, Springer-Verlag, New York, 2003. It contains results of individuals that tested for chlamydia.</p> <p><i>classes:</i> [positive, negative]</p> <p><i>summary directive:</i> Tl;dr</p> <p><i>inference directive:</i> Predict if this person will test positive or negative for chlamydia:</p>
wine	<p><i>metadata:</i> This is the Wine recognition data. Updated Sept 21, 1998 by C.Blake. It contains results of a chemical analysis of wines grown in the same region in Italy but derived from three different cultivars. The analysis determined the quantities of 13 constituents found in each of the three types of wines.</p> <p><i>classes:</i> [1, 2, 3]</p> <p><i>summary directive:</i> Using these examples and based on the contents of constituents, summarize what distinguishes wines of type 1 or 2 or 3?</p> <p><i>inference directive:</i> Hence this wine will be classified as ->type</p>
blood-transfusion-center	<p><i>metadata:</i> Data taken from the Blood Transfusion Service Center in Hsin-Chu City in Taiwan - this is a classification problem. The goal is to predict whether a given individual will consent or avoid donating blood.</p> <p><i>classes:</i> [consent, avoid]</p> <p><i>summary directive:</i> Tl;dr</p> <p><i>inference directive:</i> Therefore, this individual is likely to (avoid/consent):</p>
somerville-happiness-survey	<p><i>metadata:</i> This is the Somerville Happiness Survey Data Set. It has ratings collected from a survey of Somerville residents. From the responses of a resident, the goal is to predict whether they feel happy or unhappy about the place.</p> <p><i>classes:</i> [unhappy, happy]</p> <p><i>summary directive:</i> Based on the Somerville happiness survey, how can we predict whether a resident is happy or unhappy with their place?</p> <p><i>inference directive:</i> So this resident is (happy or unhappy):</p>
vehicle	<p><i>metadata:</i> This is the Statlog (Vehicle Silhouettes) Data Set. The purpose is to classify a given silhouette as one of four types of vehicle - bus, saab, opel or a van, using a set of features extracted from the silhouette. The vehicle may be viewed from one of many different angles.</p> <p><i>classes:</i> [bus, saab, opel, van]</p> <p><i>summary directive:</i> Using these examples, summarize how can we differentiate if a silhouette is that of a bus, saab, opel or a van.</p> <p><i>inference directive:</i> Out of saab, bus, van and opel, this vehicle is likely to be a</p>
statlog-heart	<p><i>metadata:</i> This dataset is a heart disease database similar to a database already present in the repository (Heart Disease databases) but in a slightly different form. It has data on individuals having and not having heart disease.</p> <p><i>classes:</i> [present, absent]</p> <p><i>summary directive:</i> Differentiate people with heart disease present from ones absent.</p> <p><i>inference directive:</i> In this case, heart disease is likely to be (present/absent):</p>
vertebra-column	<p><i>metadata:</i> This dataset contains values for six biomechanical features used to classify orthopaedic patients into 3 classes (normal, disk hernia or spondylololsthesis) or 2 classes (normal or abnormal). Biomedical data set built by Dr. Henrique da Mota during a medical residence period in the Group of Applied Research in Orthopaedics (GARO) of the Centre MÃ©dico-Chirurgical de RÃ©adaptation des Massues, Lyon, France. The task is to classify patients as belonging to one out of two categories: Normal (100 patients) or Abnormal (210 patients).</p> <p><i>classes:</i> [abnormal, normal]</p> <p><i>summary directive:</i> Based on the above examples, summarize how will you distinguish patients that have normal vs. abnormal vertebral column.</p> <p><i>inference directive:</i> Therefore, this individual's vertebral column is likely to be (abnormal or normal):</p>

ecoli	<i>metadata:</i> This data contains protein localization sites. Reference: "A Knowledge Base for Predicting Protein Localization Sites in Eukaryotic Cells", Kenta Nakai & Minoru Kanehisa, Genomics 14:897-911, 1992. <i>classes:</i> [1, 2] <i>summary directive:</i> Using these examples, how can we tell apart cells with protein localized in sites 1 and 2? <i>inference directive:</i> Hence protein localization will be at site ->
haberman-survival	<i>metadata:</i> The dataset contains cases from a study that was conducted between 1958 and 1970 at the University of Chicago's Billings Hospital on the survival of patients who had undergone surgery for breast cancer. <i>classes:</i> [survived, died] <i>summary directive:</i> Based on these examples, figure out what commonalities are predictive of patients surviving more than 5 years and less. <i>inference directive:</i> So, 5 years down the line, this person (survived/died):
diabetes	<i>metadata:</i> This dataset is originally from the National Institute of Diabetes and Digestive and Kidney Diseases. The objective is to predict based on diagnostic measurements whether a patient has high/low risk of developing diabetes. <i>classes:</i> [low, high] <i>summary directive:</i> Based on these examples, distinguish patients having low vs. high risk of diabetes. <i>inference directive:</i> Based on the reasoning, this patient is likely to have a (low/high):
visualizing-hamster	<i>metadata:</i> This is the visualizing-hamster dataset contains 22 data sets from the book Visualizing Data published by Hobart Press (books@hobart.com). It contains examples of hamsters that are ill and healthy. <i>classes:</i> [ill, healthy] <i>summary directive:</i> Using these examples, identify predictive indicators of ill and healthy hamsters. <i>inference directive:</i> Predict whether this hamster will be ill or healthy:
wholesale-customes	<i>metadata:</i> The data set refers to clients of a wholesale distributor. It includes the annual spending in monetary units (m.u.) on diverse product categories. This data gives information about spending patterns and region of operations of Retail and Horeca (Hotel/Restaurant/Café) customers of the wholesale distributor. <i>classes:</i> [retail, horeca] <i>summary directive:</i> Using these examples, summarize how can we differentiate Retail customers and Horeca customers. <i>inference directive:</i> Therefore, which one of Retail or Horeca this customer is likely to be:

512 A.3 Inference

513 **Mapping labels from LLM responses** Answer mapping refers to the process of assigning the
514 model's answer to a target output class. This step might be trivial when the answer is the class
515 itself, for example when the LLM responds with "non-relapse" or "recurrence" to a query on
516 the breast-cancer dataset. However, in other instances, it can become tricky when the LLM's
517 responses are "will not recur" or "has higher chance of non-relapse than recurrence", requiring a
518 more complex decoding logic to identify the target class.

519 Previous works have handled this problem by disguising the task as a Cloze prompt and learning
520 a verbalizer, i.e. MLP projection of hidden state of the [MASK] token, that maps the predicted
521 token to the target classes [4, 19]. By training a verbalizer, one can deterministically go from token
522 vocabulary to the label space. There also exist unsupervised statistical techniques for achieving label
523 mapping [41].

524 In our method however, we strictly interact with the LLM through prompts and do not access the
525 hidden state nor gradients. As a result, our inference process shown in Figure 2 focusses on inferring
526 the class label solely through prefix prompts, without relying on learning an explicit mapping.
527 Specifically, by conditioning on a suitable prefix, we constrain the LLM to return exactly the class
528 label string. For example, the prefix "Therefore this iris flower is likely to be (setosa, versicolor,
529 virginica):" works for the iris dataset. A key observation guiding the design of such a prefix prompt
530 is the fact that specifying the output classes entices the LLM to predict from among these classes.
531 With a rather plain prefix like "Predict what will be the type of this flower:", the LLM's answer space
532 is unconstrained and it might liberally go on to explain a chain of reasoning such as "The flower has
533 short petals and long sepals, hence it is versicolor, and not setosa." preventing a simple keyword
534 search for the class label.

535 For a full list of these inference prompts, refer Table 3.

536 **Two-step prompting, Davinci vs. Curie** It is worth mentioning that a two-step prompting trick,
537 by first calling "Lets think step by step" then concatenating the response with the prefix prompt
538 also results in accurate answers, as we have shown in Figure 4 (left). However, it could only be
539 implemented on the larger model Davinci but not Curie which is primarily used in our experiments.
540 Interestingly Davinci's chain of thought reasoning even outperforms its prefix prompt counterpart.
541 In all our experiments with Curie however, the prefix technique works reasonably well.

542 The Davinci API also offers a suffix argument which can be invoked for predicting in a more natural
543 way. For instance, for the breast-cancer dataset, the prompt can be posed with prefix "All in all,
544 this woman is more likely to have a " and a suffix "of breast cancer." directly expecting the LLM to
545 fill in with "recurrence" or "non-relapse."

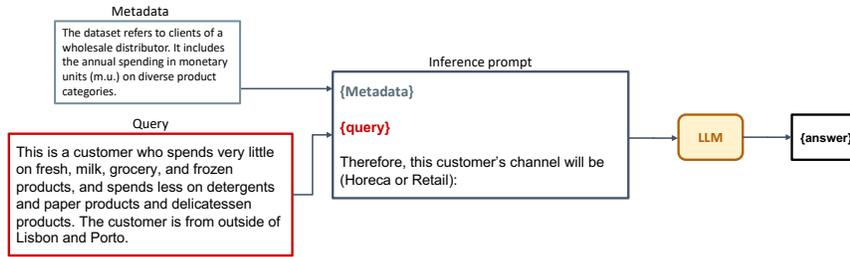


Figure 5: Steps of Zeroshot prompting

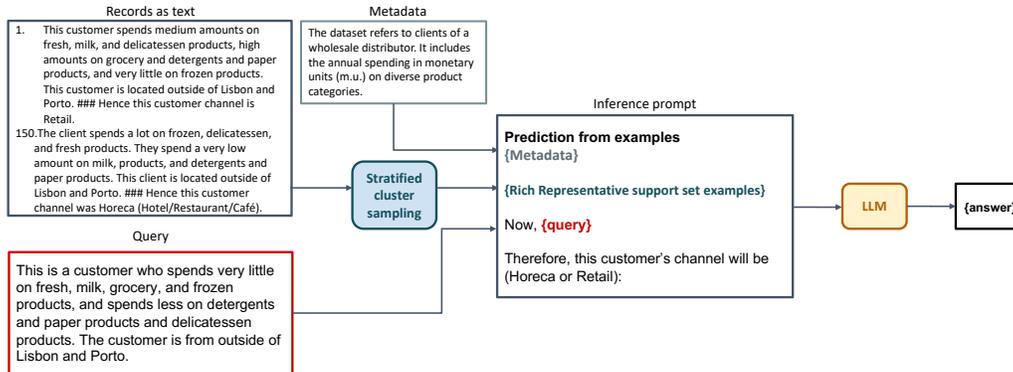


Figure 6: Workflow in fewshot prompting

546 A.4 Zeroshot Setting

547 We extend the analysis of prompting-based methods in Section 4.1 by further delving into the
 548 Zeroshot experiment. We illustrate this experimental setting in Figure 5. It only consists of the
 549 inference prompt, wherein the LLM is presented with the metadata and a query. To facilitate answer
 550 mapping, the output classes are also indicated in the prompt. Unlike Summary, the zeroshot process
 551 is not stochastic as there is no learning involved. For inference, the predicted tokens are sampled
 552 greedily at temperature = 0.

553 A.5 Few-shot Setting

554 Extending the analysis from Section 4, we explain the Fewshot setting. As illustrated in Figure 6,
 555 it is an instance of in-context learning, wherein the support set examples are enlisted in the prompt
 556 along with a query in the end. The support set was chosen from the training set through the stratified
 557 cluster sampling outlined in Algorithm 1. This results in a semantically diverse collection of examples
 558 evenly spread across the classes. Since we observe in Figure 4 (bottom right) that the Fewshot
 559 performance drops with more examples, we choose approximately 15 examples to fit in the prompt.

560 Similar to the Summary method, this inference prompt carries meta-information about the dataset,
 561 and also indicates the output classes. The prompt also injects support set examples that are stringed
 562 together as we show in Figure 6. The predictions are made greedily at a temperature of 0.

563 Again, the Fewshot method is a stochastic process whose outcome depends on the selection of the
 564 support set. So finding the ideal prompt, requires sampling different support sets from the training
 565 set quite a few times. We perform this resampling approximately 25 times and pick the best prompt
 566 based on the validation error.

567 A.6 Preprocessing continuous attributes

568 Extending from the ablation studies in Section 5, we demonstrate in Table 4, concrete examples of
 569 these encoding techniques applied to continuous features. Every numerical column in the dataset was
 570 subject to the transformation independently of the rest.

Table 4: Continuous variable transformations applied to an example from the wholesale-customers dataset. The raw tabular record is as follows: spending on fresh products: 6353.0, spending on milk products: 8808.0, spending on grocery products: 7684.0, spending on frozen products: 2405.0, spending on detergents and paper products: 3516.0, spending on delicatessen products: 7844.0 and customer’s region: Outside Lisbon and Porto.

Method	Data Representation	Example as text
4 bins + quantifiers {very low, low, high, very high}	<ul style="list-style-type: none"> - spending on fresh products : low - spending on milk products : very high - spending on grocery products : high - spending on frozen products : high - spending on detergents and paper products : high - spending on delicatessen products : very high - customer’s region : Outside Lisbon and Porto 	This customer spends low amounts on fresh products, very high amounts on milk products, high amounts on grocery products, frozen products, detergents and paper products, and very high amounts on delicatessen products. They are located outside of Lisbon and Porto.
5 bins + quantifiers {very low, low, medium, high, very high}	<ul style="list-style-type: none"> - spending on fresh products : medium - spending on milk products : very high - spending on grocery products : high - spending on frozen products : high - spending on detergents and paper products : high - spending on delicatessen products : very high - customer’s region : Outside Lisbon and Porto 	This customer from outside Lisbon and Porto spends medium on fresh products, very high on milk products, high on grocery products, high on frozen products, high on detergents and paper products, and very high on delicatessen products.
7 bins + quantifiers {extremely low, very low, low, medium, high, very high, extremely high}	<ul style="list-style-type: none"> - spending on fresh products : low - spending on milk products : very high - spending on grocery products : high - spending on frozen products : high - spending on detergents and paper products : very high - spending on delicatessen products : extremely high - customer’s region : Outside Lisbon and Porto 	This customer situated outside Lisbon and Porto spends low on fresh products, very high on milk products, high on grocery products, high on frozen products, very high on detergents and paper products, and extremely high on delicatessen products.
9 bins + quantifiers {lowest, extremely low, very low, low, medium, high, very high, extremely high, highest}	<ul style="list-style-type: none"> - spending on fresh products : low - spending on milk products : extremely high - spending on grocery products : high - spending on frozen products : high - spending on detergents and paper products : very high - spending on delicatessen products : highest - customer’s region : Outside Lisbon and Porto 	This customer spends low amounts on fresh products, extremely high amounts on milk products, high amounts on grocery products, frozen products, detergents and paper products, and highest amounts on delicatessen products. They are located outside Lisbon and Porto.
10 bins	<ul style="list-style-type: none"> - spending on fresh products : falls in the first out of ten bins of values - spending on milk products : falls in the second out of ten bins of values - spending on grocery products : falls in the first out of ten bins of values - spending on frozen products : falls in the first out of ten bins of values - spending on detergents and paper products : falls in the first out of ten bins of values - spending on delicatessen products : falls in the second out of ten bins of values - customer’s region : Outside Lisbon and Porto 	This customer spends relatively little on fresh, grocery, frozen and detergents/paper products, and more on milk and delicatessen products. They are based outside Lisbon and Porto.
Percentile	<ul style="list-style-type: none"> - spending on fresh products : falls in the forty-first percentile - spending on milk products : falls in the eighty-second percentile - spending on grocery products : falls in the sixty-fifth percentile - spending on frozen products : falls in the sixty-third percentile - spending on detergents and paper products : falls in the seventy-second percentile - spending on delicatessen products : falls in the ninety-eighth percentile - customer’s region : Outside Lisbon and Porto 	This customer has an annual spending of 41st percentile on fresh products, 82nd percentile on milk products, 65th percentile on grocery products, 63rd percentile on frozen products, 72nd percentile on detergents and paper products, and 98th percentile on delicatessen products, and is located outside of Lisbon and Porto.
Standard deviation	<ul style="list-style-type: none"> - spending on fresh products : is within one std-dev below the mean value - spending on milk products : is within one std-dev above the mean value - spending on grocery products : is within one std-dev below the mean value - spending on frozen products : is within one std-dev below the mean value - spending on detergents and paper products : is within one std-dev above the mean value - spending on delicatessen products : is two std-dev above the mean value - customer’s region : Outside Lisbon and Porto 	The customer has annual spending on fresh products, milk products, grocery products, frozen products, detergents and paper products, and delicatessen products within one standard deviation of the mean, except for delicatessen products which is two standard deviations above the mean. The customer is located outside Lisbon and Porto.
Quartiles	<ul style="list-style-type: none"> - spending on fresh products : is between the first quartile and median values - spending on milk products : is more than the third quartile value - spending on grocery products : is between median and third quartile values - spending on frozen products : is between median and third quartile values - spending on detergents and paper products : is between median and third quartile values - spending on delicatessen products : is more than the third quartile value - customer’s region : Outside Lisbon and Porto 	This customer spends more than the third quartile value on milk, delicatessen and detergents and paper products. The customer’s spending on fresh, grocery, and frozen products falls between the median and third quartile values, while the customer is located outside of Lisbon and Porto.

571 We applied several encoding techniques for continuous features, including binning, percentiles, and
572 standard deviations. Our approach involved using technical language terms to describe these ranges,
573 such as “falls in the *n*th bin/*n*th percentile or *n* deviations above/below mean”. We also characterize
574 them in a more naturalistic way by assigning quantifiers such as *low*, *medium*, and *high* to each level
575 in the binning technique.

576 To create effective textual descriptions, we examined three high-level approaches: 1. presenting
577 only numerical values, 2. using solely textual encodings, and 3. concatenating both. We observed
578 that utilizing textual encoding alone outperformed the other methods. As a result, we focused on
579 mainly comparing textual encoding methods as shown in Figure 4 (right bottom). Through Bayesian
580 optimization, we found that binning with “5” quantifiers was ideal for generating high-quality
581 summaries.

582 We describe each encoding technique as follows:

- 583 • **Binning:** It involves creating a histogram with the given number of bins. As outputs,
584 the values are directly described as “*falling in the n -th bin*” as illustrated in the 10 bins
585 experiment. However, in the presence of *degree quantifiers* which are categorical names
586 assigned to the bins, these tags are used instead. We found that as opposed to calling out the
587 bin number, describing in terms of these quantifiers further aids the LLM in comparing the
588 relative extent to which features match and improving the estimation of similarities. This
589 led us to tune the number of bins against these degree quantifiers, selecting values in the
590 range of 4, 5, 7, and 9 bins. The first four rows in Table 3 show how these tags get translated
591 into the record.
- 592 • **Percentile:** It is given by computing the percentile rank of a value relative to that series of
593 values. Then, the value is described as falling in that percentile rank in words. This is closer
594 to representation of the true numerical values per se, but helps the LLM draw comparisons
595 on a scale of 1-100.
- 596 • **Standard deviations:** In this procedure, the values are segmented into six ranges based on
597 distance from the mean, given by one/two/three standard deviations above/below the mean.
- 598 • **Quartiles:** Here, we consider the first and third quartiles, and the median as landmarks to
599 bucketize the values into four partitions.

600 Among these methods, the “5 bins with quantifiers” strikes a balance in granularity scale. It is not
601 excessively fine-grained as “percentile”, nor overly abstract, as the “4-bin” approach. This balance
602 ultimately leads to optimal performance.

603 A.7 Clustering Sampling components

604 We discuss more of the functions in Algorithm [1](#)

605 **GPT-Embedding** is OpenAI’s text similarity model `text-embedding-ada-002` that takes a max-
606 imum input size of 8191 tokens. It returns a 1536-dimensional embedding for text. OpenAI
607 recommends cosine distance for comparing *ada* embeddings in downstream tasks.

608 As a result, the **AgglomerativeClustering** algorithm applies hierarchical clustering over these features
609 using cosine distance, average linkage and a heuristically selected distance threshold of 0.05. It yields
610 a set of clusters C and each C_j contains a list of indices of data points that belong to that cluster j .

611 A.8 Adaboost Optimizations

612 We additionally apply several run-time optimizations to the boosting algorithm described in [3.3](#). Thus
613 we present its full version in Algorithm [3](#).

- 614 • **Raising the bar for a weak learner:** Our goal was to create high-quality summaries that
615 dramatically reduce the validation error rate and significantly accelerate the convergence
616 of the boosting procedure. Thus we raise the performance threshold to a notch slightly
617 higher than random guessing probability (see Step 15 in Algorithm [3](#)), provoking insightful
618 summaries.

619 We resample until finding a weak learner that satisfies this threshold.

620 The positive quantity μ is a hyperparameter that typically takes values 0.08 for 2-class
621 problem and 0.16 for 3-class problem, and so on.

622 Although this step increases compute, it yields better weak learners and improves conver-
623 gence overall.

- 624 • **Permuting the predicted class label assignments:**

625 We harness the potential of permuting the class assignments by exploring $K!$ different
626 mappings of predictions to classes using the **PermutedLabelMappings** function in steps
627 11-14. This process helps us identify the mapping that minimizes the training error to the
628 greatest extent.

629 By considering multiple permutations of predictions across the label space, as outlined in
630 Steps 11-14 of Algorithm [3](#), we obtain a hashmap p from the **PermutedLabelMappings**
631 function. This hashmap maps the predictions \hat{y} to the permuted label space. Selecting

Algorithm 3 Summary Boosting

```
1: Input:  $X$ , all training data;  $y$ , all training label;  $T$ : maximum number of rounds;  $s$ : size of the
   sampled subset.
2:  $h, P, \epsilon, \alpha \leftarrow$  empty array of size  $T$ .  $\triangleright h$  holds the hypotheses,  $P$  are the corresponding label
   mappings,  $\epsilon$  gathers the weighted train errors, and  $\alpha$  are coefficients of the hypotheses.
3:  $N \leftarrow \text{len}(X)$ 
4:  $c \leftarrow$  set of target classes
5:  $K \leftarrow \text{len}(c)$ 
6:  $w \leftarrow$  new array of size  $N$  filled with  $\frac{1}{N}$ .  $\triangleright w$  is the weighted data distribution
7: for  $r = 1$  to  $T$  do
8:    $(X_s, y_s) \leftarrow$  Cluster-sample  $s$  examples from training data using distribution  $w$ .
9:    $h[r] \leftarrow$  Summary  $(X_s, y_s)$   $\triangleright h[r]$  is the weak learner in the current round
10:   $\hat{y} \leftarrow h[r](X[i])$   $\triangleright \hat{y}$  refers to predictions on training set
11:   $\xi \leftarrow$  empty hashmap  $\triangleright \xi[p]$  will have error rate of the corresponding label mapping  $p$ 
12:  for  $p$  in PermutedLabelMappings $(c)$  do
13:     $\xi[p] \leftarrow \frac{\sum_{i=1}^N w[i] \times \mathbb{1}\{p[\hat{y}[i]] \neq y[i]\}}{\sum_{i=1}^N w[i]}$ 
14:  end for
15:   $p^* \leftarrow \arg \min_p \xi[p]$ 
16:  if  $\xi[p^*] > 1 - \frac{1}{K} - \mu$  OR AllSame $(\hat{y})$  then
17:    Resample  $(X_s, y_s)$  and Goto Step 8.
18:  else
19:     $P[r] \leftarrow p^*$ ;  $\epsilon[r] \leftarrow \xi[p^*]$ 
20:  end if
21:  if  $\epsilon[r] == 0$  then
22:    Break
23:  end if
24:   $\alpha[r] \leftarrow \log\left(\frac{1-\epsilon[r]}{\epsilon[r]}\right) + \log(K-1)$ 
25:  for  $i = 1$  to  $N$  do
26:     $w[i] = w[i] \times \exp(\alpha[r] \mathbb{1}\{P[r][h[r](X[i])] \neq y[i]\})$ 
27:  end for
28:   $w \leftarrow$  Normalize $(w)$ 
29: end for
30: Return  $h, \alpha$ 
```

632 the mapping that results in the lowest training error effectively diminishes the cumulative
633 training error during boosting iterations and proves to be an effective strategy for generating
634 strong weak learners. This technique is particularly advantageous in scenarios involving
635 more than two classes.

636 • **Sanity checks:** Finally, to ensure robustness of the weak learner when faced with skewed
637 datasets, we have implemented a policy that disallows a naive all-ones classifier. The
638 condition calling **AllSame** in Step 15 of Algorithm 3 performs this check.

639 A.9 Text Templates

640 In the ablation study which involves masking the attribute names, as illustrated in Figure 5 (second
641 from left), we transform the descriptive attributes into the textual format by applying a pre-defined
642 template. In Table 5 we provide examples of these templates for selected datasets.

643 A.10 Complexity Analysis

644 We provide the time complexity analysis comparing our boosting procedure to finetuning the LLM.

645 For finetuning, the complexity is $\mathcal{O}(TNf)$, where f is runtime of the LLM, T is number of epochs,
646 N is the number of data points.

647 For summary boosting, the complexity is $\mathcal{O}(TRf)$, where f is runtime of the LLM, T is number of
648 boosting rounds and R is the number of resampling per round.

Table 5: **Templatized descriptions:** Templates used to format examples for the ablation study between LLM-created data descriptions vs. template descriptions

Dataset	Descriptive attribute values	Template
caesarian	<i>age</i> : [very young, young, middle-aged, old, very old] <i>delivery_number</i> : [first, second, third, fourth, fifth] <i>delivery_time</i> : [timely, premature, latecomer] <i>blood_pressure</i> : [low, normal, high] <i>heart_problem</i> : [has, doesn't have] <i>delivery_mode</i> : [normal, caesarian]	This { <i>age</i> } woman is in her { <i>delivery_number</i> } delivery and it is { <i>delivery_time</i> }. She has a { <i>blood_pressure</i> } blood pressure and { <i>heart_problem</i> } heart problems. ### Based on these attributes, this woman is likely to deliver by { <i>delivery_mode</i> }
iris	<i>sepal_length, petal_length</i> : [very short, short, medium length, long, very long] <i>sepal_width, petal_width</i> : [very narrow, narrow, medium width, wide, very wide] <i>flower_type</i> : [setosa, versicolor, virginica]	This iris flower has { <i>sepal_length</i> } and { <i>sepal_width</i> } sepals. It also has { <i>petal_length</i> } and { <i>petal_width</i> } petals. ### Hence this flower is a { <i>flower_type</i> }
vertebral-column	<i>pelvic_incidence, pelvic_tilt, lumbar_lordosis_angle, sacral_slope, pelvic_radius, grade_of_spondylolisthesis</i> : [very low, low, medium, high, very high] <i>result</i> : [normal, abnormal]	This patient has a { <i>pelvic_incidence</i> } pelvic incidence, { <i>pelvic_tilt</i> } pelvic tilt, and { <i>lumbar_lordosis_angle</i> } lumbar lordosis angle, { <i>sacral_slope</i> } sacral slope, { <i>pelvic_radius</i> } pelvic radius and { <i>grade_of_spondylolisthesis</i> } grade of spondylolisthesis. ### As a result, the patient's vertebral-column is likely to be { <i>result</i> }
statlog-heart	<i>age</i> : [very young, young, middle-aged, old, very old] <i>sex</i> : [male, female] <i>chest_pain_type</i> : [asymptomatic, nonanginal pain, atypical angina, typical angina] <i>bp, cholesterol, st_depression, heart_rate, num_major_vessels</i> : [very low, low, medium, high, very high] <i>fasting_blood_sugar</i> : [high, low] <i>electrocardiographic_results</i> : [having left ventricular hypertrophy, normal, having ST-T wave abnormality] <i>slope_st_segment</i> : [flat, upsloping, downsloping] <i>exercise_induced_angina</i> : [has, do not have] <i>defect_type</i> : normal, reversible, fixed <i>presence_of_heart_disease</i> : [present, absent]	This individual is a/an { <i>age</i> } { <i>sex</i> } with { <i>chest_pain_type</i> } chest pain, { <i>bp</i> } resting blood pressure, and { <i>cholesterol</i> } serum cholesterol. Their fasting blood sugar { <i>fasting_blood_sugar</i> } >120 mg/dl, they are { <i>electrocardiographic_results</i> } and a { <i>heart_rate</i> } maximum heart rate. They { <i>exercise_induced_angina</i> } exercise-induced angina, and have a { <i>st_depression</i> } ST depression induced by exercise relative to rest. Their peak exercise ST segment has a { <i>slope_st_segment</i> } slope, and they have a { <i>num_major_vessels</i> } number of major vessels. The defect type is { <i>defect_type</i> }. ### Hence heart disease is likely to be { <i>presence_of_heart_disease</i> }.
haberman-survival	<i>age_at_time_of_op</i> : [very young, young, middle-aged, old, very old] <i>year_of_op</i> : [1964, 1962, 1965, 1959, 1958, 1960, 1966, 1961, 1967, 1963, 1969, 1968] <i>num_pos_axillary_nodes</i> : [very low, low, medium, high, very high] <i>survival_status</i> : [survived, died]	This patient was { <i>age_at_time_of_op</i> } at the time of operation in { <i>year_of_op</i> }. They had a { <i>num_pos_axillary_nodes</i> } number of positive axillary nodes detected. ### Therefore 5 years down the line, the patient { <i>survival_status</i> }

649 Concretely, for a dataset with 175 examples, finetuning takes $20 \text{ epochs} \times 175 \text{ examples} \times 2 = 7000$
 650 passes through the LLM. 2 stands for both forward and backward passes through the model.

651 For the same dataset boosting requires $50 \text{ rounds} \times 25 \text{ resampling on average} = 1250$ passes through
 652 the LLM.

653 Thus, we believe the complexity of our algorithm is at least comparable to, if not better than, that of
 654 finetuning (without considering the cost of the actual API calls).

655 A.11 Estimating the cost of API calls

656 While our method is applicable to any large language model (LLM), we primarily conducted experi-
 657 ments using GPT-3. Each API call to GPT-3 incurs a specific dollar cost.

658 After analyzing the running time complexity of summary boosting, which is $\mathcal{O}(TRf)$, we can provide
 659 a rough estimation of the cost associated with training a classifier on any given dataset.

660 To begin, when making a call to summarize examples, the prompt is filled up to the maximum context
 661 length, which is 2048 tokens for the query prompt and completion. We'll refer to these summary
 662 tokens as $S_t = 2048$.

663 Additionally, if N represents the size of the dataset and we allocate $(50 + 10)$

664 Now, to obtain a weak learner at boosting round r , we may need to resample up to R candidate
 665 summaries. Furthermore, we calculate the training error for each candidate summary to determine if
 666 it performs better than random guessing. Once the desired weak learner is found, we compute the
 667 validation error for that round only once. Therefore, each round requires querying $R \times (S_t + 0.5N \times$
 668 $P_t) + 0.1N \times P_t$ tokens.

669 Considering that the maximum number of rounds is denoted as T , the total number of tokens
 670 exchanged would be $T \times [R \times (S_t + 0.5N \times P_t) + 0.1N \times P_t]$.

671 For instance, let's consider a dataset with 175 examples. In this case, the cost would be $30 \text{ rounds} \times$
 672 $[20 \text{ resampling} \times (2048 \text{ summary tokens} + (0.5 \times 175 \text{ training examples}) \times 210 \text{ prediction tokens}) +$

Table 6: Comparing test error rate of Summary Boosting backed by Curie and ChatGPT on all datasets (\downarrow). Refer to caption of Table 1 for the notations.

Dataset	Data Type	Size	Curie	ChatGPT
caesarian [cae] (42901)	1c4d	80	0.300 \pm 0.04	0.406 \pm 0.03
iris (61)	4c0d	150	0.193 \pm 0.03	0.083 \pm 0.01
tae (48)	1c4d	151	0.454 \pm 0.03	0.443 \pm 0.04
glass (41)	9c0d	214	0.370 \pm 0.02	0.492 \pm 0.02
breast-cancer [bc] (13)	7c5d	277	0.288 \pm 0.02	0.360 \pm 0.01
visualizing-environmental [ve] (678)	3c0d	111	0.268 \pm 0.03	0.333 \pm 0.04
analcatdata-chlamydia [ac] (535)	2c2d	100	0.170 \pm 0.01	0.300 \pm 0.06
wine (43571)	13c0d	178	0.320 \pm 0.01	0.250 \pm 0.01
blood-transfusion-center [btc] (1464)	4c0d	748	0.240 \pm 0.04	0.433 \pm 0.01
somerville-happiness-survey [shs] [21]	0c7d	143	0.350 \pm 0.02	0.430 \pm 0.02
vehicle (54)	18c0d	846	0.410 \pm 0.04	0.350 \pm 0.16
statlog-heart [stath] [8]	6c7d	270	0.430 \pm 0.01	0.370 \pm 0.17
vertebra-column [vc] (1524)	6c0d	310	0.262 \pm 0.01	0.669 \pm 0.03
ecoli (1011)	7c0d	336	0.270 \pm 0.03	0.193 \pm 0.03
haberman-survival [hs] (43)	3c0d	306	0.250 \pm 0.01	0.415 \pm 0.03
diabetes [dia] (37)	8c0d	768	0.344 \pm 0.01	0.297 \pm 0.04
visualizing-hamster [hams] (708)	5c0d	73	0.207 \pm 0.00	0.400 \pm 0.08
wholesale-customers [wc] (1511)	6c1d	440	0.330 \pm 0.00	0.199 \pm 0.04

673 (0.1 \times 175 validation examples) \times 210 prediction tokens] = 12364050 tokens, which approximately
674 costs \$25 for Curie at a rate of \$0.002/1K tokens.

675 A.12 Can ChatGPT function as a weak learner?

676 One would expect that it is more advantageous to try newer LLMs such as ChatGPT that produce
677 increasingly more human-like text and are far more sample-efficient, i.e. can summarize more
678 examples since they come with a larger context length. To investigate this, we conduct experiments
679 by feeding ChatGPT with the same tabular data descriptions and using identical prompts to create
680 weak learners. The results are presented in Table 6.

681 Surprisingly ChatGPT outperforms Curie in classifying datasets with more numerical features, such
682 as wine, wholesale-customers, and iris. This observation suggests that LLMs are becoming
683 more adept at quantitative reasoning from finetuning with more data. However, the reinforcement
684 learning from human feedback (RLHF) [30] poses a limitation as it still ensures the generated text
685 does not deviate too much from its prior. The generated text distribution adheres closely to the
686 behavior programmed into the LLM induced by optimizing with such a reward model. Consequently
687 it becomes challenging to bias the LLM with adversarial examples that might occasionally emerge in
688 the training set.

689 For example, ChatGPT does not mostly generalize well on datasets with medical information such
690 as vertebra-column, breast-cancer, caesarian and blood-transfusion-center where
691 there can be examples contrary to common medical beliefs. In these cases, the RLHF is more
692 restrictive due to its conformity to human preferences and does not neutrally summarize examples at
693 hand from a classification standpoint. However, boosting imposes a significantly higher penalty on
694 examples that the model fails to classify correctly, causing ChatGPT to not decrease training error
695 after a few epochs. While these models exhibit promise in terms of higher-order problem-solving
696 skills, their capabilities can also be limited by their alignment with human preferences.