

Perspective

Statistical Hypothesis Testing versus Machine Learning Binary Classification: Distinctions and Guidelines

Jingyi Jessica Li^{1,*} and Xin Tong²¹Department of Statistics, University of California, Los Angeles, CA 90095-1554, USA²Department of Data Sciences and Operations, Marshall School of Business, University of Southern California, Los Angeles, CA 90089, USA*Correspondence: jjli@stat.ucla.edu<https://doi.org/10.1016/j.patter.2020.100115>

THE BIGGER PICTURE In data science education, two analysis strategies, hypothesis testing and binary classification, are mostly covered in different courses and textbooks. In real data application, it can be puzzling whether a binary decision problem should be formulated as hypothesis testing or binary classification. This article aims to disentangle the puzzle for data science students and researchers by offering practical guidelines for choosing between the two strategies.



Concept: Basic principles of a new data science output observed and reported

SUMMARY

Making binary decisions is a common data analytical task in scientific research and industrial applications. In data sciences, there are two related but distinct strategies: hypothesis testing and binary classification. In practice, how to choose between these two strategies can be unclear and rather confusing. Here, we summarize key distinctions between these two strategies in three aspects and list five practical guidelines for data analysts to choose the appropriate strategy for specific analysis needs. We demonstrate the use of those guidelines in a cancer driver gene prediction example.

INTRODUCTION

Making binary decisions is one of the most common human cognitive activities. Binary decisions are everywhere: from spam detection in IT technologies to biomarker identification in medical research. For example, facing the current COVID-19 pandemic, medical doctors need to make a critical binary decision: whether an infected patient needs hospitalization or not. Living in a big data era, how can we make rational binary decisions from massive data?

In data sciences, two powerful strategies have been developed to assist binary decisions: the statistical *hypothesis testing*¹ and the machine learning *binary classification*.² While both strategies are popular and have achieved profound successes in various applications, their distinctions are largely obscure to practitioners and even data scientists sometimes. An important reason is that the two strategies are usually introduced in different classes and covered by different textbooks, with few exceptions, such as “All of Statistics: A Concise Course in Statistical Inference” by Wasserman.³ Another source of confusion is the ambiguous use of the term “test” to represent both strategies in our daily life, such as in “statistical test” and “COVID-19 test,” where the latter is, in fact, bi-

nary classification and is referred to as “COVID-19 diagnosis” in this article.

There are online discussions about the relationship between binary classification and hypothesis testing; however, they focus on specific cases and are not unified into a coherent picture. For example, one discussion compares the Student’s t test, a specific statistical test, with the support vector machines, a specific binary classification algorithm.⁴ Another discussion compares the asymmetric nature of hypothesis testing with the general lack of asymmetry in binary classification.⁵ Besides online discussions, there are research works that borrow ideas from hypothesis testing to develop binary classification algorithms,^{6,7} but these works do not aim to link or compare the two strategies.

Here, we attempt to make the first efforts to summarize key distinctions between hypothesis testing in frequentist statistics (our discussion does not pertain to Bayesian hypothesis testing⁸) and binary classification in machine learning. We also provide five practical guidelines for data analysts to choose between the two strategies. In our discussion, we frequently use biomarker detection and disease diagnosis as examples of hypothesis testing and binary classification, respectively. In these two examples, *instances* refer to patients, and *features* refer to patients’ diagnostic measurements, such as blood pressure



and gene expression levels. Note that instances are often referred to as “individuals” in biomedical sciences, “objects” in engineering, “observations” in statistics, and “data points” in data sciences. Although many researchers outside of statistics refer to instances as “samples” (in the plural form), here we stick with the classic statistical definition: a “sample” is a collection of instances. Features are also referred to as “variables” and “covariates” in statistics and “attributes” in engineering.

DISTINCTIONS BETWEEN HYPOTHESIS TESTING AND BINARY CLASSIFICATION

Hypothesis testing and binary classification are rooted in two different cultures: *inference* and *prediction*, each of which has been extensively studied in statistics and machine learning, respectively, in the historical development of data sciences.⁹ In brief, an inferential task aims to infer an unknown truth from observed data, and hypothesis testing is a specific framework whose inferential target is a binary truth, i.e., an answer to a yes/no question. For example, deciding whether a gene is an effective COVID-19 biomarker in the blood is an inferential question, whose answer is unobservable. In contrast, a prediction task aims to predict an unobserved property of an instance, such as a patient or an object, based on the available features of this instance. Such prediction relies on building a trustworthy relationship, i.e., a *prediction rule*, from the input features to the target property, which must be based on human knowledge (throughout the human history) and/or established from data (after computing devices were developed). Binary classification is a special type of prediction whose target property is binary, and COVID-19 diagnosis is an example. In screening patients for COVID-19 exams, medical doctors make binary decisions based on patients’ symptoms (input features), and their decision rules are learned from previous patients’ diagnostic data and the medical literature.

Hypothesis testing is built upon the concept of *statistical significance*, which intuitively means that the data we observe present strong evidence against a presumed *null hypothesis*, the default. In the example of testing whether a gene is a COVID-19 biomarker in blood, the null hypothesis is that this gene does not exhibit differential expression in the blood of uninfected individuals and COVID-19 patients. This formulation reflects a conservative attitude: we do not want to call the gene a biomarker unless its expression difference is large enough between the healthy and diseased patients in our collected data. Statistical hypothesis testing provides a formal framework for deciding a threshold on the expression difference so that the gene can be identified as a biomarker with the desired confidence. A crucial fact about hypothesis testing is that the null and alternative hypotheses pertain to a property of an unseen population. As a result, we cannot know whether the null hypothesis holds or not. What we have access to is instances and their features, i.e., *data*, from the population, and hypothesis testing allows us to infer how unlikely the data are generated from the null hypothesis.

In machine learning, binary classification belongs to supervised learning, as it is supervised by quality *training data* that contain training instances from two classes, and each training instance is labeled as class 0 or 1 with a set of feature values.

A binary decision rule is first constructed from the training data and next applied to predict unobserved binary labels of new instances from their observed feature values. Binary classification embodies a large class of algorithms that automatically learn prediction rules from training data. In an ideal scenario, a prediction rule follows a scientific law, such as in Newton’s second law of motion, where the acceleration of an object is determined by the net force acting on the object and the mass of the object. However, most prediction tasks do not have scientific laws to follow, and the prediction rules learned from data could be useful but not necessarily revealing scientific truth.¹⁰ For example, we can effectively predict the coming of autumn from our observation of falling leaves, which, however, do not cause autumn to come. Nevertheless, the lack of scientific interpretation is often not a major concern in many industrial applications, such as spam detection and image recognition, where prediction algorithms have achieved tremendous successes, promoting machine learning to become a spotlight discipline with broad impacts on everyone’s life. Still, a necessary condition for binary classification to succeed is that training instances are good representatives of new instances to make predictions for. A notorious cautionary tale is Google Flu Trends, which mistakenly predicted a nonexistent flu epidemic because its training data did not well represent the long-term dynamics of flu outbreaks.^{11,8}

We summarize the key distinctions between hypothesis testing and binary classification in three aspects: data in relation to binary decisions, construction of decision rules, and evaluation criteria. Our discussion is centered around four concepts: *binary questions*, *binary answers*, *decision rules*, and *binary decisions*, which we define for each strategy in Table 1. We note that these four concepts belong to three stages in a typical data analysis: conceptual formulation (when binary questions and binary answers are formulated in a researcher’s mind), analysis (when a decision rule is constructed), and conclusion (when a binary decision is made).

Data in Relation to Binary Decisions

In this aspect, hypothesis testing and binary classification have two distinctions: (1) number of instances to make one decision given a decision rule and (2) availability of known binary answers in data. In hypothesis testing, each binary decision—rejecting a null hypothesis or not—is made from a collection of instances, called a *sample* in statistics. For example, to investigate whether a gene is a COVID-19 biomarker in blood, a researcher needs to collect blood from multiple uninfected and infected patients, whose number is called the *sample size*, and measure this gene’s expression within. Then the binary decision regarding whether to call this gene an informative biomarker is made jointly from the collected measurements. If multiple genes are tested simultaneously, we are in a situation called *multiple testing*,¹² which is commonly used in large-scale exploratory studies. No matter the number of tests being one or many, the number of instances used for each test should better exceed one in the practices of hypothesis testing. In fact, the greater the number of instances, the more we trust our decisions. We further discuss the impact of the number of instances on decision rules in the third aspect (see below: “Evaluation Criteria for Decision Rules”).

In contrast, binary classification makes a binary decision for every instance that needs a binary label. In COVID-19 diagnosis,

Table 1. Four Concepts under Hypothesis Testing and Binary Classification

| Concept | Hypothesis Testing | Binary Classification |
|-----------------|---|---|
| Binary question | Is the null hypothesis false? (unanswerable) | Does the instance have a label 1? |
| Binary answer | 0 (no) The null hypothesis is true (unobservable) | The instance has a label 0 |
| | 1 (yes) The null hypothesis is false (unobservable) | The instance has a label 1 |
| Decision rule | A statistical test that inputs data and outputs a p value, which is compared against a user-specified significance level α | A trained classifier that inputs an unlabeled instance's feature values and outputs a predicted label |
| Binary decision | 0 Do not reject the null hypothesis | Label the instance as 0 |
| | 1 Reject the null hypothesis | Label the instance as 1 |

a doctor needs to decide which patients should be hospitalized, and each patient will receive one decision. In other words, the number of instances in need of binary labels equals the number of decisions. Here training instances are not counted, because they already have binary labels. In practice, binary classification can be easily confused with multiple testing, as both strategies make multiple binary decisions (see the cancer driver gene prediction example in the section “Cancer Driver Gene Prediction: Hypothesis Testing or Binary Classification?” below). A way to distinguish the two strategies is to count the number of input instances used to make one decision given a decision rule, whose construction is discussed below in the section “Construction of Decision Rules.”

Another distinction is the availability of known answers to binary questions. Such answers are always lacking for hypothesis-testing questions but available in training data for binary classification. In hypothesis testing, a binary question is regarding the validity of a null hypothesis, and the answer to this question is an unobservable truth about an unseen population. For example, we do not know *a priori* whether a gene is a biomarker; otherwise, we would not need to do hypothesis testing. However, the unobservable binary answer is often mistaken as a binary decision—whether or not to reject the null hypothesis—an action taken based on a sample of instances and dependent upon a decision rule (Table 1). This mistake is commonly seen in scientific research papers that claim, “the null hypothesis is correct (or incorrect) because the p value is large (or small).” Here, we raise a strong caution against this misuse.

Unlike in hypothesis testing, a binary question in classification is regarding the binary label of an instance, and we already have known answers (labels) for training instances, which we utilize to build a decision rule to predict labels of unlabeled instances. For example, doctors diagnose new patients based on previous patients' data with diagnosis decisions. It is worth emphasizing that

a decision rule cannot be constructed if all training instances have the same label, say 0; hence, training data must contain both binary labels. In brief, hypothesis testing has no concept of training data, because data contain no answers to binary questions being asked; in contrast, training data serve as a critical component in binary classification.

Construction of Decision Rules

In hypothesis testing, the construction of a decision rule, also known as a *statistical test*, relies on three essential components: a test statistic that summarizes the data, the distribution of the test statistic under the null hypothesis, and a user-specified significance level α , which indicates the tolerable *type I error*, i.e., the conditional probability of mistakenly rejecting the null hypothesis given that it holds. (In many textbooks and practices, α is set to 0.05 by convention. However, we want to emphasize that this convention is just an arbitrary choice and should not be taken as the ritual. For example, there is a recent proposal to change the “default” value of α to 0.005,¹³ yet both 0.05 and 0.005 are arbitrary thresholds.) The first two components lead to a p value between 0 and 1, with a smaller value indicating stronger evidence against the null hypothesis. Then the null hypothesis is rejected if the p value does not exceed α . Numerous statistical tests have been developed since the advent of statistics, and a few of them, such as Student's t test and Wilcoxon's rank-sum test, have become standard practices in data analysis. Because of the wide popularity and meanwhile common misuses of hypothesis testing, there are recent in-depth and extensive discussions on the proper use and interpretation of p values in and outside of the statistics community.^{13,14}

In multiple testing, the choice of α value is determined by an overall objective on all tests together, and two widely used objectives are the family-wise error rate (FWER) (the probability of wrongly rejecting at least one null hypothesis) and the false discovery rate (FDR) (the expected proportion of falsely rejected hypotheses among all rejections).⁸ (In high-throughput data analyses common in genomics and proteomics, the FDR is the most popular objective, while the FWER is rarely used due to its over-conservativeness. However, the FWER is still frequently used in scientific research where a moderate [e.g., fewer than 20] number of hypothesis tests are performed together.) The Bonferroni correction is a conservative but guaranteed approach to control the FWER.¹⁵ The Benjamini-Hochberg procedure is a widely used approach to control the FDR,¹⁶ and there is a recent approach, *knockoffs*, to control the FDR when exact p values cannot be achieved.^{17,18} (Note that the FDR is a frequentist criterion. Under the Bayesian framework, empirical Bayes criteria, including the positive false discovery rate (pFDR),¹⁹ the local false discovery rate (FDR),²⁰ and the local false sign rate,²¹ have been developed to control the number of false positives in multiple testing.) It is worth noting that the construction of a decision rule in hypothesis testing does not necessarily require access to data. For example, in the classic Student's two-sample t test, under the assumption that the two samples (sets of instances) are generated from two normal distributions, the decision rule only depends on the two sample sizes and a user-specified α value (see Box 1). When researchers have collected a gene's expression data in many diseased and healthy patients and have verified that the two samples approximately follow normal distributions,

Box 1. Example: Decision Rule of the Student's Two-Sample t Test

Suppose that we have two samples of sizes 10 and 12 from two normal distributions, and we are interested in whether the two normal distributions have the same mean. Then the null hypothesis is that the two normal distributions have the same mean, and the alternative hypothesis is the opposite. The test statistic—the two-sample t statistic—follows the t distribution with 20 degrees of freedom (t_{20}) under the null hypothesis. Then given a significance level $\alpha \in (0, 1)$, we would reject the null hypothesis if the t statistic has an absolute value greater than or equal to $F_{t_{20}}^{-1}(1 - \alpha/2)$, i.e., the $(1 - \alpha/2)$ -th quantile of the t_{20} distribution. For example, if $\alpha = 0.05$, $F_{t_{20}}^{-1}(1 - \alpha/2) = 2.085963$. This *decision rule* is equivalent to that the p value is less than or equal to α . Note that this decision rule does not depend on any observed t statistic value calculated from a dataset. However, a *decision* requires an observed t statistic value. For example, if the t statistic has an observed value of 3, then we would reject the null hypothesis at $\alpha = 0.05$.

they can simply apply the two-sample t test, a readily available decision rule, to their data and decide whether this gene can be called a biomarker at their desired α value. If the normal distributional assumption does not seem to hold, researchers may use the Wilcoxon rank-sum test that does not have this assumption but only requires all the instances to be independent. Hence, in applications of hypothesis testing, the most critical step is to choose an appropriate statistical test, i.e., decision rule, by checking the test's underlying assumptions on data distribution. Meanwhile, the construction of valid new decision rules is mostly the job of academic statisticians.

In contrast to hypothesis testing, we do not usually have available decision rules to choose from in binary classification; instead, we need to construct a decision rule from training data in most applications. Image classification and speech recognition are probably two famous exceptions, where superb decision rules (classifiers) have been trained from industry-standard massive image and speech datasets that well represent almost all possible images and speeches in need of labeling (decisions) in daily applications. Yet in biomedical applications, such as COVID-19 diagnosis, a good decision rule is often lacking but needed to be constructed from in-house training data that represent future local patients in need of diagnosis. Despite its reliance on quality training data that contain a reasonable number of instances with accurate binary labels, binary classification is fortunate to have access to dozens of powerful algorithms that can be directly applied to training data to construct a decision rule. Famous algorithms include the logistic regression, support vector machines, random forests, gradient boosting, and the resurgent neural networks (and its buzzword version “deep learning”).^{2,22} As in hypothesis testing, the most critical step in applications of binary classification is the choice of an appropriate algorithm to build a decision rule from training data, while the development of new algorithms is a focus of data science researchers.

Evaluation Criteria for Decision Rules

Realizing the many possible ways of constructing decision rules in both hypothesis testing and binary classification, users face a challenging question in data analysis: How should I compare and evaluate decision rules? In hypothesis testing, statistical tests (decision rules) designed for the same null hypothesis are compared in terms of *power*: the conditional probability of correctly rejecting the null hypothesis given that it does not hold, e.g., correctly identifying an effective biomarker. Under

the same significance level α , the larger the power, the better the test. The *Neyman-Pearson lemma* provides the theoretical foundation for the most powerful test; however, in many practical scenarios, the Neyman-Pearson lemma does not apply and the most powerful test is not achievable, so statisticians have put continuous efforts into developing more powerful tests, such as in the flourishing field of statistical genetics.²³ For users, the power of a statistical test is not observable from data, which contain no information regarding the null hypothesis being true or not. Hence, the only evaluation criterion for users to choose among many statistical tests is whether their data seem to fit each test's underlying assumptions on data distribution, a check that can be quite tricky sometimes and may require consultations with statisticians. If many tests pass this check, most users would choose a popular test. An advanced user might opt for the test that gives the smallest p value, i.e., the strongest evidence against the null hypothesis. However, this option should be used with extreme caution, as it could easily become “p-hacking” or data dredging if without sufficient justification.²⁴

In binary classification, the evaluation criteria are more transparent and easier to understand, as they all rely on some sorts of prediction accuracy of a decision rule on *validation data*, which contain binary labels, represent future instances that need labeling, and most importantly are not part of the training data. Users may wonder: What if I only have one set of data with binary labels? A straightforward answer is to randomly split the data into training and validation cohorts, use the training cohort to construct a decision rule, and apply the rule to the validation cohort to evaluate a chosen measure of prediction accuracy. (However, as in the Google Flu Trends example, if the training data are not representative of future instances, this splitting idea would not work.) This answer is the core idea leading to *cross-validation*, the dominant approach for evaluating binary classification rules, and more generally, prediction rules.⁸ If users prefer not to split the data due to its limited number of instances, probabilistic approaches are available, and they allow users to use the whole dataset to train and subsequently evaluate a decision rule. Famous examples include the *Akaike's information criterion* and the *Bayesian information criterion*.⁸ However, there is no free lunch; most of these non-splitting approaches require assumptions on data distribution (if these probabilistic assumptions do not hold, there is no guarantee how the decision rule would perform on a future instance) and do not apply to binary classification algorithms that are not probability based, while cross-validation has no such restrictions. In terms of prediction accuracy, the most commonly used measure is

Box 2. Sample Sizes versus Evaluation Criteria

A general principle in data sciences is, if a sample is unbiasedly drawn from a population, the larger the sample size, the more information the sample contains about the population. This *large-sample principle* holds for both hypothesis testing and binary classification; for example, data from a larger number of representative patients would lead to better decision rules for both biomarker detection and disease diagnosis. However, between the two strategies there is an interesting but often neglected distinction: from a population with finite instances (e.g., the human population), the largest possible sample, which is equivalent to the whole population, would make a valid statistical test achieve a perfect power given any significance level α , while the largest possible training dataset might not lead to a classification rule with perfect prediction accuracy. While this distinction is fundamentally rooted in mathematics, an intuitive understanding can be obtained from our biomarker detection and disease diagnosis examples. Imagine that we have measured everyone in the world. If a gene is indeed a disease biomarker, we can for sure see a difference in this gene's expression between all the people carrying this disease and the rest of the population, achieving the perfect power. On the other hand, diseased patients and undiseased individuals may not be perfectly separated by diagnostic measurements. That is, two patients may have similar symptoms and lab test results, but one patient is diseased and the other is not. When this happens, even if we have training data from all but one person in the world, we still cannot be 100% sure whether the left-out individual has the disease just based on his or her diagnostic measurements.

the *overall accuracy*: the percentage of correctly labeled instances in the validation data, e.g., the percentage of correctly diagnosed patients in a cohort not used for training the decision rule. In many applications where the two classes corresponding to binary labels 0 and 1 have equal importance, this measure is reasonable. In disease diagnosis, however, the two directions of misdiagnosis: predicting a diseased patient as healthy versus predicting a healthy individual as diseased, are likely to have unequal importance, which would depend on the severity of the disease, the abundance of medical resources, and many other factors. For example, in countries with well-established health-care systems, diagnosis for high-mortality cancer patients should focus on reducing the false negative rate, i.e., the chance of missing a patient with a malignant tumor; hence, a more relevant prediction accuracy would be the true positive rate (one minus the false negative rate) in this context. Binary classification with unequal class importance is called *asymmetric classification*, to address which two frameworks have been developed: *cost-sensitive learning*²⁵ and *Neyman-Pearson classification*.^{26–28} (The Neyman-Pearson classification inherits its name from the Neyman-Pearson lemma due to a similar asymmetric nature: minimizing one type of error while controlling the other type of error.) Specifically, the cost-sensitive learning framework achieves a small false negative rate by placing on it a large weight relative to the weight on the false positive rate in the objective function; the Neyman-Pearson classification framework guarantees a high-probability control on the population-level false negative rate while minimizing the false positive rate. Another two commonly used accuracy measures for binary classification are the area under a receiver operating characteristic curve (AUROC) and the area under a precision-recall curve (AUPRC). (Important properties and distinctions between AUROC and AUPRC include but are not limited to: AUROC is invariant to the population sizes of the two classes, while AUPRC is not; AUROC can be overly optimistic if the two classes are extremely imbalanced in training data, while AUPRC does not have this issue. For detailed information, please refer to the two books^{29,30} and the survey article.³¹) However, these two measures are not evaluation criteria for one decision rule (classifier) but rather evaluate a trained classification algorithm (e.g., logistic regression with parameters estimated from training data) with varying

decision thresholds, each of which corresponds to a decision rule.

In summary, the evaluation of decision rules in hypothesis testing is less straightforward than in binary classification. To choose a statistical test for a specific dataset, users have to use subjective judgment to decide whether test assumptions are reasonably justified. On the other hand, classification algorithms can be compared on a more objective ground, the *Common Task Framework*,³² of which influential examples include the Kaggle competitions^{33–37} and the DREAM challenges.^{38–43} The Common Task Framework consists of three essential elements: training data, competing prediction algorithms, and validation data. A comparison is considered fair if all competing algorithms use the same training data to construct decision rules, which are subsequently evaluated on the same validation data using the same prediction accuracy measure. Furthermore, between hypothesis testing and binary classification, an interesting, technical distinction is how their evaluation criteria change with the sample size (see Box 2).

Table 2 summarizes the above distinctions between hypothesis testing and binary classification.

A CHECKLIST OF FIVE PRACTICAL GUIDELINES FOR CHOOSING BETWEEN HYPOTHESIS TESTING AND BINARY CLASSIFICATION

Based on the key distinctions between hypothesis testing and binary classification, we provide a checklist of five practical guidelines for data analysts to choose between the two strategies. Figure 1 provides an illustration.

Guideline 1: Decide on Instances and Features

Given a tabular dataset, the first and necessary step is to decide whether rows and columns should be considered as instances and features, respectively, or vice versa. The decision may seem trivial to experienced data scientists when columns represent variables in different units, e.g., gender, age, and body mass index, in which case columns should be considered as features for sure. However, the decision may become not-so-obvious in certain cases. For example, in Figure 1, a gene expression dataset has rows and columns corresponding to patients and genes,

Table 2. Side-by-Side Comparison of Hypothesis Testing and Binary Classification

| | Hypothesis Testing | Binary Classification |
|---|------------------------------------|---|
| Symmetry between binary answers | Asymmetric (default is 0) | Symmetric or asymmetric |
| No. of instances to make one decision given a decision rule | ≥ 1 (the larger the better) | = 1 |
| Available binary answers | No | Yes (training data) |
| Evaluation criteria | Power (given a significance level) | Prediction accuracy |
| With the largest possible no. of instances | Power = 1 | Prediction accuracy not necessarily perfect |

respectively, and all data values are in the same unit. The question is: Should we consider patients as instances or features? To answer this question, the key is to understand instances as either (1) repeated measurements in the data collection process or (2) a random sample from a population. Gene expression data are collected to understand gene expression patterns in healthy and diseased human sub-populations, so healthy and diseased patients are considered two random samples, each from one sub-population and satisfy (2). Hence, we conclude that patients are instances and genes are features. In general, the answer depends on the experimental design and the scientific question, both of which will determine what the underlying population is, as we see in the cancer driver gene prediction example (see below in the section “Cancer Driver Gene Prediction: Hypothesis Testing or Binary Classification?”).

Guideline 2: List the Binary Decisions to Be Made

The second step is to outline the binary decisions to be made from the data. Formulate analytical tasks, such as biomarker detection and disease diagnosis into binary questions, for which binary decisions will be made. Divide binary questions into those related to features and others concerning instances. For example, whether a gene is a biomarker is a feature-related question, and whether a patient has a disease is an instance-related question. Hypothesis testing can only answer feature-related questions, while binary classification can only address instance-related questions.

Guideline 3: Assess the Availability of Known Binary Answers in Data

After a list of binary questions is at hand, the next question is: Do the data contain any known answers? If we already have an answer to a binary question, we cannot formulate that question as a hypothesis-testing task. In the case where some instances contain known binary labels but we concern about the unknown labels of the rest of instances, we are facing a binary classification task, just as in disease diagnosis. Otherwise, if the data contain no binary labels, we do not have training data to

construct a classifier, which, if not given, would prohibit us from predicting unknown labels of instances.

Guideline 4: Count the Number of Instances for Making Each Binary Decision

Suppose that we are given a decision rule, i.e., a statistical test or a classifier in the form of a formula or a computer program that can take our data as input and output a binary answer. An easy check is to count the number of input instances needed to output each binary decision. If we are expecting one decision per input instance, it is likely a binary classification task. Otherwise, if each binary decision needs to be made from a group of instances together, the task cannot be binary classification but might be formulated as hypothesis testing.

Guideline 5: Evaluate the Nature of Binary Questions

The most fundamental check is to evaluate each binary question by its nature: Is the question regarding the unseen population of which our observed instances are a subset or regarding a particular instance? Asking whether a gene is a disease biomarker is a question of the former type, as it concerns whether this gene can distinguish the human sub-population carrying the disease from the rest of human population. In contrast, asking whether an individual has the disease is a question of the latter type. Hypothesis testing and binary classification are designed for answering questions of the former and latter type, respectively, as shown in Figure 1.

We recommend practitioners to check all the five guidelines before deciding whether hypothesis testing or binary classification is the correct strategy for a data analysis question.

CANCER DRIVER GENE PREDICTION: HYPOTHESIS TESTING OR BINARY CLASSIFICATION?

Finally, we present an important application example of cancer driver gene prediction to illustrate the distinction between hypothesis testing and binary classification. We will try to avoid using technical terms as much as possible for the ease of general readers. In the problem of cancer driver gene prediction, the goal is to utilize an individual gene’s mutational signatures such as the number of missense mutations to predict how likely the gene drives cancer. Note that the mutational signatures are aggregated from multiple patient databases, and individual patients’ data are unavailable. We have knowledge of a small set of cancer driver genes and neutral genes that are unlikely to drive cancer. The question is, how can we leverage this knowledge to predict whether a less-studied gene is a cancer driver gene? A famous algorithm, TUSON, addresses this question using a hypothesis-testing approach.^{44,45} In brief, it regards mutational signatures as features and uses hypothesis testing to assess how much an individual gene resembles known neutral genes based on each feature: the gene’s feature value is used as the test statistic, whose distribution under the null hypothesis (i.e., the gene is a neutral gene) is estimated from the feature values of known neutral genes; from the test statistic and the approximate null distribution, the gene receives a p value for that feature. Suppose that there are 10 features in total, then each gene receives 10 p values, which are subsequently combined into a single p value by Fisher’s method.⁴⁶




Figure 1. Illustration of a Gene Expression Dataset and Two Questions to Be Addressed by Hypothesis Testing and Binary Classification, Respectively

Hypothesis testing uses all the available instances to address a feature-related question: is a gene a biomarker with different expression levels in healthy and diseased patients? Binary classification answers an instance-related question: is a patient diseased?

From a statistical perspective, there are four apparent issues with this hypothesis-testing approach. First, each hypothesis test, one per gene per feature, only utilizes the known neutral genes (to estimate the null distribution) but does not fully capture the valuable information in known cancer driver genes. (To be exact, known cancer driver genes are used to select the predictive features before hypothesis testing is performed. However, in each test for one gene and one feature, the information of known cancer driver genes is only partially reflected in the direction of the gene’s p value: two-sided, larger-than, or smaller-than,^{44,45} which excludes the possibility that the known neutral genes may have feature values on the two sides of those of the known cancer driver genes.) Second, each hypothesis test is performed using a sample of size one (i.e., the test statistic is the feature value of one gene), which is known to be not powerful and thus undesirable (i.e., if the gene is a cancer driver gene, we may miss it with a high chance). This is the reason why we recommend using more than one instance for hypothesis testing (Table 2, and “Guideline 3” above). Third, combining multiple p values into a single p value is a difficult task, especially when p values are not independent of each other. The fact that mutational signatures are observed to be correlated features, their resulting p values are correlated for each gene, violating the independence assumption of Fisher’s method. Although there are methods for combining dependent p values,^{47–50} they cannot address the most fundamental

question—What is the population behind each hypothesis test?—leading to the last issue. Fourth, the population behind the null hypothesis is unspecified: for a given test, is the population about that particular gene or all the genes? Therefore, we conclude that the hypothesis-testing approach is inappropriate for this cancer driver gene problem, which, instead, should be formulated as a binary classification task by the reasoning below.

Here, we revisit this problem by following our checklist. Under Guideline 1, we consider genes as instances and mutational signatures as features, consistent with the existing studies. The reason is that we treat known cancer driver genes and neutral genes as a sample from the whole gene population of our interest, while we consider mutational signatures as given and we are not interested in the population they come from. Note that here genes are no longer treated as features as in biomarker detection where patients are instances. The contrast of the two examples suggests that a real quantity may be formulated as an instance or a feature depending on the data and the question of interest, and Guideline 1 provides a practical solution. Under Guideline 2, we conclude that the binary decisions to be made are instance-related because we would like to predict whether each gene is a cancer driver gene or not. Guideline 3 leads us to identify training data: known cancer driver genes and neutral genes. Next, if we already have a decision rule, we just need to input one gene to obtain its binary label: cancer driver gene or not.

Hence, we only need one instance for each binary decision, concluding Guideline 4. Finally, we evaluate the nature of each binary question, as suggested by Guideline 5, and we can see that each question is only concerning one instance (gene), not the gene population. After checking all the five guidelines, it becomes evident that this cancer driver gene prediction problem is better suited to be addressed by binary classification.

Why does TUSON adopt the hypothesis-testing approach? Its analysis results show that it aims to control the proportion of false discoveries among the predicted cancer driver genes, a criterion closely related to the FDR, which is widely used in multiple testing as we have discussed. (The difference is that the FDR is the expected proportion of false discoveries among discoveries, where the expectation is taken over possible input datasets, while in TUSON, the proportion is based on one dataset. However, this difference has been largely neglected in biomedical studies.) Our guess is that the TUSON authors formulate cancer driver gene prediction as a multiple testing problem because they want to apply the famous Benjamini-Hochberg procedure to control the FDR by setting a cutoff on p values, one per gene. However, this approach requires the validity of each p value, which must follow a uniform distribution between 0 and 1 under the null hypothesis. Due to the third issue we mentioned above (the violation of the assumption of Fisher's method), the combined p value of each gene has no guarantee to satisfy this requirement. Here, we would like to point out that the FDR concept is not restricted to multiple testing; in fact, it is a general evaluation criterion for multiple binary decisions, where each decision rule could be established by hypothesis testing or binary classification. Therefore, the goal of FDR control should not drive the choice between hypothesis testing and binary classification; instead, the choice should be based on the distinctions between the two strategies, as we have discussed in this article. Admittedly, the FDR has been rarely used as an evaluation criterion in binary classification; however, it is a closely related criterion: *precision*, the proportion of correct predictions among all positive predictions, is widely used, such as in AUPRC. (Note that precision is a criterion evaluated on a given set of validation data. Unlike the FDR, it is not an expected proportion. How to implement a theoretically guaranteed FDR control in binary classification is an open question for data science researchers.) For cancer driver gene prediction, if we adopt the binary classification approach, we may compare competing classification algorithms by evaluating their AUPRC values using cross-validation. After we choose the algorithm that achieves the largest AUPRC value, we can train it on known cancer driver genes and neutral genes using their mutational signatures, and we can set a threshold on the trained algorithm based on our desired precision level to obtain a classifier (decision rule). Then we can simply apply the classifier to predict whether a less-studied gene is a cancer driver gene from its mutational signatures. In fact, we have implemented this approach and shown that it leads to more accurate discoveries than previous studies do.⁵¹

DISCUSSION

In summary, hypothesis testing and binary classification are largely regarded as two separate topics that have rarely been compared with each other in data science education and research. However, their distinctions in applications are not as

apparent as in methodological research, where instances and features are well defined from the beginning. Instead, in applications how to formulate real quantities into instances or features is always a challenging task, a reason that obscures the distinctions between the two strategies. In this work, we attempt to summarize and compare the two strategies for the broad scientific community and the data science industry, and we provide five practical guidelines to help data analysts better distinguish between the two strategies in data analysis.

As an extension, when instances are categorized into more than two classes or groups, similar distinctions exist between hypothesis testing and classification. Multi-group comparison is a typical hypothesis-testing question. For example, if the question is whether a feature has the same expected value in all groups, it can be addressed by the analysis of variance. In contrast, multi-class classification seeks to assign every unlabeled instance one or more of the multiple class labels. In fact, compared with our previous discussion, here the distinction between hypothesis testing and classification is clearer: hypothesis testing still gives rise to a binary decision (reject or not the null hypothesis that a feature has the same expected value in all groups), while classification leads to a decision with more than two possible answers (every instance has more than two possible labels).

To conclude, we would like to emphasize again that hypothesis testing and binary classification are different in nature: the former concerns an unobservable population-level property of a feature, while the latter pertains to an observable label of an instance. Despite this inherent difference, it is possible to adopt ideas from one strategy to develop new decision rules for the other strategy, or use one strategy as a preceding step to enable the application of the other strategy. In one direction, powerful test statistics with theoretical foundations in hypothesis testing may inform the construction of classifiers in binary classification. For example, the likelihood ratio test statistic bears a similar mathematical form as the naive Bayes classifier. (The key distinction is that the likelihood ratio test statistic takes all the available instances as input to construct a decision rule for one test, while the naive Bayes classifier, if trained, takes one instance as input and outputs a predicted binary label.) There are other concrete examples that leverage test statistics to construct classifiers.^{6,7} In the other direction, successful classification algorithms may motivate new scientific questions that can be investigated by hypothesis testing. For example, convolutional and recurrent neural networks have demonstrated superb capacity to extract predictive features from unstructured data, such as images and texts.^{52–55} For those extracted features that are interpretable (such as a feature related to one brain region in fMRI images), researchers may want to investigate whether such a feature differs between two groups of images from different patient cohorts, and then they need to perform hypothesis testing on that particular feature. Furthermore, in some special and rare examples, an algorithm may serve the purposes of both strategies. The most famous example is logistic regression, which is both a classification algorithm and a testing approach for deciding whether associations exist between features and binary labels. In a binary classification task whose goal is to label instances, logistic regression is used to construct a classifier. Meanwhile, logistic regression and its accompanying Wald test can also be used to investigate how each feature influences binary labels of instances.⁵⁶ Ultimately,

effective data analysis requires appropriate usage of hypothesis testing and binary classification for suitable tasks, and this can only be realized when data analysts are well informed of the distinctions and connections between the two strategies.

ACKNOWLEDGMENTS

This work was supported by the following grants: National Science Foundation DBI-1846216, NIH/NIGMS R01GM120507, Johnson & Johnson WiSTEM2D Award, Sloan Research Fellowship, and UCLA David Geffen School of Medicine W.M. Keck Foundation Junior Faculty Award (to J.J.L.). We have received many insightful comments during the preparation of this manuscript. We thank Dr. Wei Li at University of California, Irvine, for bringing the cancer-driver gene prediction problem to our attention and commenting on our manuscript. We also thank the following people for their comments: Heather J. Zhou, Tianyi Sun, Kexin Li, Wenbin Guo, Yiling Chen, Dongyuan Song, and Ruochen Jiang in the Junction of Statistics and Biology (<http://jsb.ucla.edu>) at University of California, Los Angeles; Lijia Wang, Man Luo, Lan Gao, and Dr. Michael S. Waterman at University of Southern California; Ruhan Dong at Pandora Media; Dr. Han Chen at University of Texas Health; Dr. Mark Biggin at Lawrence Berkeley National Laboratory; and Dr. Thomas Burger at the Université Grenoble Alpes and the French National Center for Scientific Research (CNRS).

AUTHOR CONTRIBUTIONS

J.J.L. and X.T. conceived the idea. J.J.L. wrote the paper and prepared the visualization. X.T. thoroughly reviewed and edited the paper.

REFERENCES

- Lehmann, E.L., and Romano, J.P. (2006). *Testing Statistical Hypotheses* (Springer Science & Business Media).
- Murphy, K.P. (2012). *Machine Learning: A Probabilistic Perspective* (MIT press).
- Wasserman, L. (2013). *All of Statistics: A Concise Course in Statistical Inference* (Springer Science & Business Media).
- MaxG (2016). Distinguishing between two groups in statistics and machine learning: hypothesis test vs. classification vs. clustering. <https://stats.stackexchange.com/questions/262686/distinguishing-between-two-groups-in-statistics-and-machine-learning-hypothesis>.
- icurays1 (2016). Why is binary classification not a hypothesis test?. <https://stats.stackexchange.com/questions/240138/why-is-binary-classification-not-a-hypothesis-test>.
- Liao, S.-M., and Akritas, M. (2007). Test-based classification: a linkage between classification and statistical testing. *Stat. Probab. Lett.* **77**, 1269–1281.
- He, Z., Sheng, C., Liu, Y., and Zhou, Q. (2019). Instance-based classification through hypothesis testing. *arXiv* <https://arxiv.org/pdf/1901.00560.pdf>.
- Efron, B., and Hastie, T. (2016). *Computer Age Statistical Inference, volume 5* (Cambridge University Press).
- Breiman, L. (2001). Statistical modeling: the two cultures (with comments and a rejoinder by the author). *Stat. Sci.* **16**, 199–231.
- Riley, P. (2019). Three pitfalls to avoid in machine learning. *Nature* **572**, 27–29.
- Harford, T. (2014). Big data: a big mistake? *Significance* **11**, 14–19.
- Efron, B. (2012). *Large-scale Inference: Empirical Bayes Methods for Estimation, Testing, and Prediction, volume 1* (Cambridge University Press).
- Benjamin, D.J., Berger, J.O., Johannesson, M., Nosek, B.A., Wagenmakers, E.-J., Berk, R., Bollen, K.A., Brembs, B., Brown, L., et al. (2018). Redefine statistical significance. *Nat. Hum. Behav.* **2**, 6.
- Wasserstein, R.L., and Lazar, N.A. (2016). The ASA statement on p-values: context, process, and purpose. *Am. Stat.* **70**, 129–133.
- Bonferroni, C.E. (1936). Teoria statistica delle classi e calcolo delle probabilità. *Pubbl. R. Ist. Sup. Sci. Econ. Commer. Fir.* **8**, 1–62.
- Benjamini, Y., and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. R. Stat. Soc. Ser. B* **57**, 289–300.
- Barber, R.F., and Candès, E.J. (2015). Controlling the false discovery rate via knockoffs. *Ann. Stat.* **43**, 2055–2085.
- Candès, E., Fan, Y., Janson, L., and Lv, J. (2016). Panning for gold: model-X knockoffs for high-dimensional controlled variable selection. *arXiv:1610.02351*.
- Storey, J.D. (2003). The positive false discovery rate: a Bayesian interpretation and the q-value. *Ann. Stat.* **31**, 2013–2035.
- Efron, B. (2008). Microarrays, empirical Bayes and the two-groups model. *Stat. Sci.* **1**–22.
- Stephens, M. (2017). False discovery rates: a new deal. *Biostatistics* **18**, 275–294.
- Friedman, J., Hastie, T., and Tibshirani, R. (2001). *The Elements of Statistical Learning, volume 1* (Springer series in statistics).
- Sham, P.C., and Purcell, S.M. (2014). Statistical power and significance testing in large-scale genetic studies. *Nat. Rev. Genet.* **15**, 335–346.
- Head, M.L., Holman, L., Lanfear, R., Kahn, A.T., and Jennions, M.D. (2015). The extent and consequences of p-hacking in science. *PLoS Biol.* **13**, <https://doi.org/10.1371/journal.pbio.1002106>.
- Elkan, C. (2001). The foundations of cost-sensitive learning. *International Joint Conference on Artificial Intelligence, volume 17* (Lawrence Erlbaum Associates Ltd), pp. 973–978.
- Cannon, A., Howse, J., Hush, D., and Scovel, C. (2002). Learning with the Neyman-Pearson and Min-Max Criteria (Los Alamos National Laboratory), Tech. Rep. LA-UR 02-2951.
- Scott, C., and Nowak, R. (2005). A Neyman-Pearson approach to statistical learning. *IEEE Trans. Inf. Theor.* **51**, 3806–3819.
- Tong, X., Feng, Y., and Li, J.J. (2018). Neyman-Pearson classification algorithms and np receiver operating characteristics. *Sci. Adv.* **4**, eaao1659.
- He, H., and Ma, Y. (2013). *Imbalanced Learning: Foundations, Algorithms, and Applications* (John Wiley).
- Fernández, A., García, S., Galar, M., Prati, R.C., Krawczyk, B., and Herrera, F. (2018). *Learning from Imbalanced Data Sets* (Springer).
- Branco, P., Torgo, L., and Ribeiro, R.P. (2016). A survey of predictive modeling on imbalanced domains. *ACM Comput. Surv. (Csur)* **49**, 1–50.
- Donoho, D. (2017). 50 years of data science. *J. Comput. Graph. Stat.* **26**, 745–766.
- Ben Taieb, S., and Hyndman, R.J. (2014). A gradient boosting approach to the Kaggle load forecasting competition. *Int. J. Forecast.* **30**, 382–394.
- Graham, B. (2015). *Kaggle Diabetic Retinopathy Detection Competition Report* (University of Warwick).
- Iglovikov, V., Mushinskiy, S., and Osin, V. (2017). Satellite imagery feature detection using deep convolutional neural network: a Kaggle competition. *arXiv:1706.06169*.
- Zou, H., Xu, K., Li, J., and Zhu, J. (2017). The YouTube-8m Kaggle competition: challenges and methods. *arXiv:1706.09274*.
- Sutton, C., Ghiringhelli, L.M., Yamamoto, T., Lysogorskiy, Y., Blumenthal, L., Hammerschmidt, T., Golebiowski, J.R., Liu, X., Ziletti, A., and Scheffler, M. (2019). Crowd-sourcing materials-science challenges with the nomad 2018 Kaggle competition. *Npj Comput. Mater.* **5**, 1–11.
- Bansal, M., Yang, J., Karan, C., Menden, M.P., Costello, J.C., Tang, H., Xiao, G., Li, Y., Allen, J., Zhong, R., et al. (2014). A community computational challenge to predict the activity of pairs of compounds. *Nat. Biotechnol.* **32**, 1213–1222.

39. Eduati, F., Mangravite, L.M., Wang, T., Tang, H., Bare, J.C., Huang, R., Norman, T., Kellen, M., Menden, M.P., Yang, J., et al. (2015). Prediction of human population responses to toxic compounds by a collaborative competition. *Nat. Biotechnol.* *33*, 933–940.
40. Sieberts, S.K., Zhu, F., García-García, J., Stahl, E., Pratap, A., Pandey, G., Pappas, D., Aguilar, D., Anton, B., Bonet, J., et al. (2016). Crowdsourced assessment of common genetic contribution to predicting anti-TNF treatment response in rheumatoid arthritis. *Nat. Commun.* *7*, 1–10.
41. Guinney, J., Wang, T., Laajala, T.D., Winner, K.K., Bare, J.C., Neto, E.C., Khan, S.A., Peddinti, G., Airola, A., Pahikkala, T., et al. (2017). Prediction of overall survival for patients with metastatic castration-resistant prostate cancer: development of a prognostic model through a crowdsourced challenge with open clinical trial data. *Lancet Oncol.* *18*, 132–142.
42. Seyednasrollah, F., Koestler, D.C., Wang, T., Piccolo, S.R., Vega, R., Greiner, R., Fuchs, C., Gofer, E., Kumar, L., Wolfinger, R.D., et al. (2017). A dream challenge to build prediction models for short-term discontinuation of docetaxel in metastatic castration-resistant prostate cancer. *JCO Clin. Cancer Inform.* *7*, 1–15.
43. Salcedo, A., Tarabichi, M., Espiritu, S.M.G., Deshwar, A.G., David, M., Wilson, N.M., Dentre, S., Wintersinger, J.A., Liu, L.Y., Ko, M., et al. (2020). A community effort to create standards for evaluating tumor subclonal reconstruction. *Nat. Biotechnol.* *38*, 97–107.
44. Davoli, T., Xu, A.W., Mengwasser, K.E., Sack, L.M., Yoon, J.C., Park, P.J., and Elledge, S.J. (2013). Cumulative haploinsufficiency and triplosensitivity drive aneuploidy patterns and shape the cancer genome. *Cell* *155*, 948–962.
45. Tokheim, C.J., Papadopoulos, N., Kinzler, K.W., Vogelstein, B., and Karchin, R. (2016). Evaluating the evaluation of cancer driver genes. *Proc. Natl. Acad. Sci. U S A* *113*, 14330–14335.
46. Fisher, R.A. (1992). Statistical methods for research workers. In *Breakthroughs in Statistics* (Springer), pp. 66–70.
47. Brown, M.B. (1975). 400: a method for combining non-independent, one-sided tests of significance. *Biometrics* *31*, 987–992.
48. Kost, J.T., and McDermott, M.P. (2002). Combining dependent p-values. *Stat. Probab. Lett.* *60*, 183–190.
49. Wilson, D.J. (2019). The harmonic mean p-value for combining dependent tests. *Proc. Natl. Acad. Sci. U S A* *116*, 1195–1200.
50. Liu, Y., and Xie, J. (2020). Cauchy combination test: a powerful test with analytic p-value calculation under arbitrary dependency structures. *J. Am. Stat. Assoc.* *115*, 393–402.
51. Lyu, J., Li, J.J., Su, J., Peng, F., Chen, Y., Ge, X., and Li, Wei (2020). Dorge: discovery of oncogenes and tumor suppressor genes using genetic and epigenetic features. *bioRxiv*. <https://doi.org/10.1101/2020.07.21.213702>.
52. Krizhevsky, A., Sutskever, I., and Hinton, G.E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, pp. 1097–1105.
53. Hu, B., Lu, Z., Li, H., and Chen, Q. (2014). Convolutional neural network architectures for matching natural language sentences. *Adv. Neural Info. Process. Syst.* *2*, 2042–2050.
54. Mikolov, T., Kombrink, S., Burget, L., Černocký, J., and Khudanpur, S. (2011). Extensions of recurrent neural network language model. In *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (IEEE), pp. 5528–5531.
55. Sak, H., Senior, A.W., and Beaufays, F. (2014). Long short-term memory recurrent neural network architectures for large scale acoustic modeling. *arXiv*, 1402.1128.
56. McCullagh, P., and Nelder, J. (1989). *Generalized Linear Models* (Chapman & Hall/CRC).

About the Authors

Dr. Jingyi Jessica Li is an associate professor in the departments of statistics (primary), biostatistics, computational medicine, and human genetics (secondary) at University of California, Los Angeles (UCLA). Jessica is the director of the Center of Statistical Research in Computational Biology and a faculty member in the interdepartmental PhD program in bioinformatics. Leading a diverse research group—the Junction of Statistics and Biology—Jessica focuses on developing principled and practical statistical methods for tackling cutting-edge issues in biomedical data science. Prior to joining UCLA in 2013, Jessica obtained her PhD in biostatistics from University of California, Berkeley.

Dr. Xin Tong is an assistant professor in the Department of Data Sciences and Operations at the University of Southern California, where he teaches “Applied Statistical Learning Methods” in the master of science in business analytics program. His research focuses on asymmetric supervised and unsupervised learning, high-dimensional statistics, and network-related problems. He is an associate editor of *Journal of the American Statistical Association* and *Journal of Business and Economic Statistics*. Before joining the current position, he was an instructor of statistics in the Department of Mathematics at the Massachusetts Institute of Technology. He obtained his PhD in operations research from Princeton University.