

On the Existence of Simpler Machine Learning Models

Lesia Semenova
Duke University
Durham, North Carolina, USA
lesia@cs.duke.edu

Cynthia Rudin
Duke University
Durham, North Carolina, USA
cynthia@cs.duke.edu

Ronald Parr
Duke University
Durham, North Carolina, USA
parr@cs.duke.edu

ABSTRACT

It is almost always easier to find an accurate-but-complex model than an accurate-yet-simple model. Finding optimal, sparse, accurate models of various forms (linear models with integer coefficients, decision sets, rule lists, decision trees) is generally NP-hard. We often do not know whether the search for a simpler model will be worthwhile, and thus we do not go to the trouble of searching for one. In this work, we ask an important practical question: can accurate-yet-simple models be proven to exist, or shown likely to exist, before explicitly searching for them? We hypothesize that there is an important reason that simple-yet-accurate models often do exist. This hypothesis is that the *size of the Rashomon set is often large*, where the Rashomon set is the set of almost-equally-accurate models from a function class. If the Rashomon set is large, it contains numerous accurate models, and perhaps at least one of them is the simple model we desire. In this work, we formally present the *Rashomon ratio* as a new gauge of simplicity for a learning problem, depending on a function class and a data set. The Rashomon ratio is the ratio of the volume of the set of *accurate models* to the volume of the *hypothesis space*, and it is different from standard complexity measures from statistical learning theory. Insight from studying the Rashomon ratio provides an easy way to check whether a simpler model might exist for a problem before finding it, namely whether several different machine learning methods achieve similar performance on the data. In that sense, the Rashomon ratio is a powerful tool for understanding why and when an accurate-yet-simple model might exist. If, as we hypothesize in this work, many real-world data sets admit large Rashomon sets, the implications are vast: it means that simple or interpretable models may often be used for high-stakes decisions without losing accuracy.

CCS CONCEPTS

• **Computing methodologies** → *Supervised learning by classification; Classification and regression trees*; • **Human-centered computing** → *Empirical studies in visualization*.

KEYWORDS

Rashomon Set, Model Multiplicity, Simplicity, Generalization, Interpretable Machine Learning

ACM Reference Format:

Lesia Semenova, Cynthia Rudin, and Ronald Parr. 2022. On the Existence of Simpler Machine Learning Models. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*, June 21–24, 2022, Seoul, Republic of Korea. ACM, New York, NY, USA, 32 pages. <https://doi.org/10.1145/3531146.3533232>

1 INTRODUCTION

Following the principle of Occam’s Razor, one should use the simplest model that explains the data well. However, finding the simplest model, let alone any simple-yet-accurate model, is hard. As soon as simplicity constraints such as sparsity are introduced, the optimization problem for finding a simpler model typically becomes NP-hard. Thus, practitioners – who have no assurance of finding a simpler model that achieves the performance level of a black box – may not see a reason to attempt such potentially difficult optimization problems. Thus, sadly, what was once the holy grail of finding simpler models, has been, for the most part, abandoned in modern machine learning. In this work, we ask a question that is essential, and potentially game-changing, for this discussion: what if we knew, before attempting a computationally expensive search for a simpler-yet-accurate model, that one was likely to exist? Perhaps knowing this would allow us to justify the time and expense of searching for such a model. If it is true that many data sets have large enough Rashomon sets to admit simple models, then there are important implications for society – it means we may be able to use simpler or interpretable models for many high-stakes problems without losing accuracy.

Proving the existence of simpler models before aiming to find them differs from the current approach to machine learning in practice. We generally do not think about going from more complicated spaces to simpler ones; in fact, the reverse is true, where typical statistical learning theory and algorithms allowed us to maintain generalization when handling more complicated model classes (e.g., large margins for support vector machines with complex kernels or large margins for boosted trees) [11, 42]. We even build neural networks that are so complex that they can achieve zero training error, and try afterwards to determine why they generalize [4, 34]. However, because simple models are essential for many high-stakes decisions [40], perhaps we should return to the goal of aiming directly for simpler models. We will need new ideas in order to do this.

Decades of study about generalization in machine learning have provided many different mathematical theories. Many of them measure the complexity of classes of functions without considering the data [e.g., VC theory, 50], or measure properties of specific algorithms [e.g., algorithmic stability, see 5]. However, none of these theories seems to capture directly a phenomenon that occurs throughout practical machine learning. In particular, *there*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
FAccT '22, June 21–24, 2022, Seoul, Republic of Korea

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9352-2/22/06...\$15.00
<https://doi.org/10.1145/3531146.3533232>

are a vast number of data sets for which many standard machine learning algorithms perform similarly. In these cases, the machine learning models tend to generalize well. Furthermore, in these same cases, there is often a simpler model that performs similarly and also generalizes well.

We hypothesize that these three observations can all be explained by the same phenomenon: the “Rashomon effect,” which is the existence of many almost-equally-accurate models [6]. Firstly, following a key argument in our work, if there is a large *Rashomon set* of almost-equally-accurate models, a simple model may also be contained in it. Secondly, if the *Rashomon set* is large, many different machine learning algorithms may find different but approximately-equally-well-performing models inside it. An experimenter could then observe similar performance for different types of algorithms that produce very different functions. Thirdly, if the *Rashomon set* is large enough to contain simpler models, those models are guaranteed to generalize well. As we will show in Section 5.2, there are mathematical assumptions that allow us to *prove* existence of simpler models within the *Rashomon set*. If the assumptions are satisfied, a model from a simpler class is approximately as accurate as the most accurate model within the hypothesis space, which consequently leads to better generalization guarantees. The assumptions are based in approximation theory, which models how one class of functions can approximate another.

We quantify the magnitude of the *Rashomon effect* through the *Rashomon ratio*, which is the ratio of the *Rashomon set*’s volume to the volume of the hypothesis space. An illustration of the *Rashomon set* is shown in Figure 1; it does not need to be a connected or convex set. The *Rashomon ratio* can serve as a gauge of simplicity for a learning problem.¹ As a property of both a data set and a hypothesis space, it differs from the VC dimension [49] (because the *Rashomon ratio* is specific to a data set), it differs from algorithmic stability [see 22, 39] (as the *Rashomon ratio* does not rely on robustness of an algorithm with respect to changes in the data), it differs from local Rademacher complexity [2] (as the *Rashomon ratio* does not measure the ability of the hypothesis space to handle random changes in targets and actually benefits from multiple similar models), and it differs from geometric margins [50] (as the maximum margin classifier can have a small minimum margin yet the *Rashomon ratio* can be large, and margins are measured with respect to one model, whereas the *Rashomon ratio* considers the existence of many). We provide theorems that show simple cases when the *Rashomon ratio* disagrees with these complexity measures in Section 3 and Appendix B. The *Rashomon set* is not just functions within a flat minimum; it could consist of functions from many non-flat local minima as illustrated in Figure 4 in Appendix A, and it applies to discrete hypothesis spaces where gradients, and thus “sharpness” [16] do not exist. For linear regression, we derive a closed form solution for the volume of the *Rashomon set* in parameter space in Theorem 10 in Appendix B.1.

Our theory and empirical results have implications beyond cases where the size of the *Rashomon set* can be estimated in practice: they suggest computationally inexpensive ways to gauge whether

the *Rashomon set* is large without directly measuring it. *In particular, our results indicate that when many machine learning methods perform similarly on the same data set (without overfitting), it could be because the Rashomon set of the functions these algorithms consider is large. Thus, after running different machine learning methods and observing similar performance, our results indicate that it may be worthwhile to optimize directly for simpler models within the Rashomon set.*

We summarize the contributions of this work as follows: (i) We define the *Rashomon ratio* as an important characteristic of the *Rashomon set*. (ii) We provide generalization bounds for models from the *Rashomon set*, and show that the size of the *Rashomon set* serves as a barometer for the existence of accurate-yet-simpler models that generalize well. These are different from standard learning theory bounds that consider the distance between the true and empirical risks for the same function. (iii) We provide several approaches for estimating the size of the *Rashomon set*. (iv) We show empirically that when a large *Rashomon set* occurs, most machine learning methods tend to perform similarly, and also in these cases, simple or sparse (yet accurate) models exist. (v) We demonstrate that the *Rashomon ratio*, as a gauge of simplicity of a machine learning problem, is different from other known complexity measures such as VC-dimension, algorithmic stability, geometric margin, and Rademacher complexity. (vi) We show that larger *Rashomon sets* might occur in the presence of label or feature noise. An earlier version of this work appeared on arXiv under the title *A Study in Rashomon Curves and Volumes: A New Perspective on Generalization and Model Simplicity in Machine Learning* [43].

2 RELATED WORK

There are several bodies of relevant literature as discussed below.

Rashomon sets: *Rashomon sets* have been used for various purposes [6, 10, 19, 27, 32, 36, 44, 47]. For instance, Srebro et al. [44] consider a loss-restricted class of close-to-optimal models, and with an assumption of H -smoothness of a loss function, they obtain a tighter excess risk bound through local Rademacher complexity [2]. Our bounds do not work the same way and aim to prove a different type of result. Other works aim to search through the *Rashomon set* to find the most extreme models within it, rather than looking at the size of the *Rashomon set*, as we do in this work. Fisher et al. [19] leverages the *Rashomon set* in order to understand the spectrum of variable importance and other statistics across the set of good models. Our work considers the existence of models from simpler classes rather than exploring the *Rashomon set* to find a range of variable importance or other statistics. The work of Tulabandhula and Rudin [45, 46, 47] uses the *Rashomon set* to assist with decision making, by finding the range of downstream operational costs associated with the *Rashomon set*. *Rashomon sets* are related to p-hacking and robustness of estimation, because the *Rashomon set* is a set over which one might conduct a sensitivity analysis to choices made by an analyst [10]. Large *Rashomon sets* can occur when the machine learning pipeline is underspecified. D’Amour et al. [13] provides multiple examples of underspecification in computer vision, natural language processing, and healthcare domains; their work builds off of (an earlier version of) our work [43]. Madras et al. [29] proposed a post-hoc local-ensemble method that measures

¹Such measures are typically called “complexity” measures, but the *Rashomon ratio* measures simplicity, not complexity.

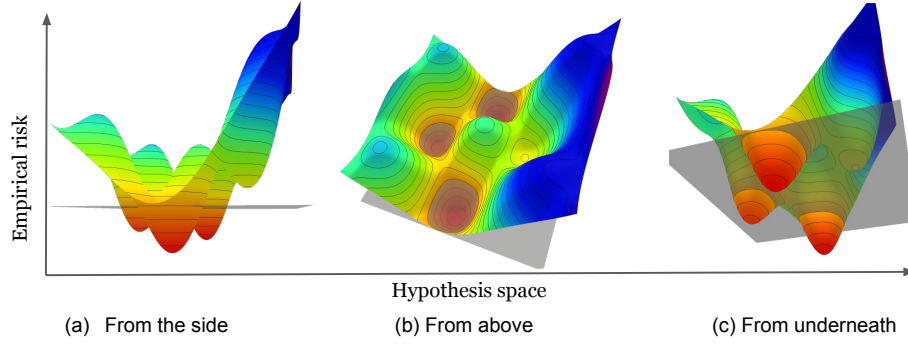


Figure 1: An illustration of a possible Rashomon set in two dimensional hypothesis space \mathcal{F} . Models below the gray plane belong to the Rashomon set $\hat{R}_{set}(\mathcal{F}, \theta)$, where the height of the gray plane is adjusted by the Rashomon parameter θ defined in Section 3.

underspecification for a given test datum. Marx et al. [31] studies conflicting predictions between models within the Rashomon set, while Coston et al. [12] investigates predictive disparities for algorithmic fairness.

Flat minima or wide valleys: The concept of flat minima (wide valleys) has been explored in the deep learning literature as a possible way to understand convergence properties of the complicated, non-convex loss functions that deep networks traverse during training [9, 16, 20, 23]. Based on a minimum-message-length argument [51], several works claim that flat loss functions lead to better generalization due to a robustness to noise around the minimum [9, 20, 23]. Following Hochreiter and Schmidhuber [20], Dinh et al. [16] define volume ϵ -flatness, which constitutes a special case of our Rashomon sets, as shown in Figure 4 in Appendix A. In particular, the Rashomon set is defined over the hypothesis (functional) space, while the volume ϵ -flatness is defined in a parameter space (though sometimes we use parameter space for ease of computation), and the Rashomon set is not necessarily a single connected component (although it might be in the case of a convex loss over a continuous domain), while the volume ϵ -flatness pertains only to a connected set. This means that the Rashomon set can contain models from different local minima, or can be defined on discrete spaces, while the volume ϵ -flatness is relevant only for continuous loss functions. Another way of quantifying flatness is σ -sharpness [16, 23], which measures the change of the loss function inside a σ -ball in a parameter space. In the case of a connected Rashomon set, this loss difference corresponds to the Rashomon parameter θ .

Statistical learning theory: Numerous works provide generalization bounds based on different complexity measures, and under different assumptions. Some discuss Rademacher [21, 44] and Gaussian complexities [21], PAC-Bayes theorems [25], covering numbers bounds [53], and margin bounds [24, 42, 49]. In contrast, under assumptions elaborated in Section 5, the Rashomon ratio provides a certificate of the existence of a simpler model that generalizes. The use of approximating sets, as used extensively in this paper, is used throughout the literature on learning theory [26, 28, 33, 42]. An example of this is the classical generalization bound for boosting and margins [42], which uses combinations of several random draws of

base classifiers to represent combinations of base classifiers. This is an instance of the so-called “Maurey’s lemma,” which provides an approximating set for linear model classes.

3 RASHOMON SET DEFINITIONS AND NOTATION

Consider a training set of n data points $S = \{z_1, z_2, \dots, z_n\}$, $z_i = (x_i, y_i)$ drawn i.i.d. from an unknown distribution \mathcal{D} on a bounded set $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$, where $\mathcal{X} \subset \mathbb{R}^p$ and $\mathcal{Y} \subset \mathbb{R}$ are an input and an output space respectively. Our hypothesis space is $\mathcal{F} = \{f : \mathcal{X} \rightarrow \mathcal{Y}\}$. We limit the hypothesis space \mathcal{F} to contain only models that vary within the bounded domain \mathcal{Z} where the data reside. We will assume that the hypothesis space is bounded and that there is a prior distribution ρ over functions in \mathcal{F} . To measure the quality of a prediction made by a hypothesis, we use a loss function $\phi : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}^+$. Specifically, for each given point $z = (x, y)$ and a hypothesis f , the loss function is $\phi(f(x), y)$. For a given f we will also overload notation by writing $l : \mathcal{F} \times \mathcal{Z} \rightarrow \mathbb{R}^+$ that takes f explicitly as an argument: $l(f, z) = \phi(f(x), y)$. We are interested in learning a model f that minimizes the *true risk* $L(f) = \mathbb{E}_{z \sim \mathcal{D}}[\phi(f(x), y)]$, which depends on unknown distribution \mathcal{D} and therefore is estimated with an *empirical risk*: $\hat{L}(f) = \frac{1}{n} \sum_{i=1}^n \phi(f(x_i), y_i)$.

The *empirical Rashomon set* (or simply *Rashomon set*) is a subset of models of the hypothesis space \mathcal{F} that have training performance close to the best model in the class, according to a loss function [6, 10, 19, 44, 47]. More precisely:

Definition 1 (Rashomon set). Given $\theta \geq 0$, a data set S , a hypothesis space \mathcal{F} , and a loss function ϕ , the *Rashomon set* $\hat{R}_{set}(\mathcal{F}, \theta)$ is the subspace of the hypothesis space defined as follows:

$$\hat{R}_{set}(\mathcal{F}, \theta) := \{f \in \mathcal{F} : \hat{L}(f) \leq \hat{L}(\hat{f}) + \theta\},$$

where \hat{f} is an empirical risk minimizer for the training data S with respect to loss function ϕ : $\hat{f} \in \arg\min_{f \in \mathcal{F}} \hat{L}(f)$.

If we want to specify the data set S that is used to compute the Rashomon set, we indicate the data set in the subscript, as $\hat{R}_{set_S}(\mathcal{F}, \theta)$. Fisher et al. [19]’s definition of Rashomon set is distinct from ours in that we typically use an empirical risk minimizer to

define the Rashomon set instead of a prespecified reference model which is independent of the sample.

The hypothesis space \mathcal{F} can be a well-defined hypothesis space, such as the space of decision trees of depth D or neural nets with D hidden layers, or it can be a more general space (a meta-hypothesis space) that contains models from different hypothesis spaces (e.g., linear functions, polynomials up to degree D , and piecewise constant functions).

We call θ the *Rashomon parameter*. Since hypothesis spaces can vary from one problem to another, we will often normalize the size of the Rashomon set via the *Rashomon ratio* $\hat{R}_{ratio}(\hat{R}_{set}(\mathcal{F}, \theta))$ which takes the Rashomon set as input and outputs a value between 0 and 1. Given a prior, ρ , on the hypothesis space, the Rashomon ratio measures the fraction of the hypothesis space contained in the Rashomon set. Unless explicitly specified, ρ is assumed to be uniform. For simplicity, we will denote the Rashomon ratio as $\hat{R}_{ratio}(\mathcal{F}, \theta)$. In general, the Rashomon ratio is $\hat{R}_{ratio}(\mathcal{F}, \theta) = \int_{\mathcal{F}} \mathbb{1}_{f \in \hat{R}_{set}(\mathcal{F}, \theta)} \rho(f) d\rho$. If the hypothesis space has a uniform prior, then the Rashomon ratio is the volume of the Rashomon set divided by the volume of the hypothesis space $\hat{R}_{ratio}(\mathcal{F}, \theta) = \frac{\mathcal{V}(\hat{R}_{set}(\mathcal{F}, \theta))}{\mathcal{V}(\mathcal{F})}$, where $\mathcal{V}(\cdot) : \mathcal{F} \rightarrow \mathbb{R}_+$ is the volume function. If the hypothesis space is discrete with a uniform prior, the Rashomon ratio can be computed as $\hat{R}_{ratio}(\mathcal{F}, \theta) = \frac{|\hat{R}_{set}(\mathcal{F}, \theta)|}{|\mathcal{F}|}$, where $|A| = \sum_{f \in \mathcal{F}} \mathbb{1}_{f \in A}$. The Rashomon ratio represents the fraction of models that are good (the fraction of models that fit the data about equally well). A larger Rashomon ratio implies that more models perform about equally well. The data set S is denoted in the subscript, as $\hat{R}_{ratio_S}(\mathcal{F}, \theta)$.

We consider *true Rashomon sets* that contain models with low true risk, relative to the optimal true risk value, with parameter $\gamma > 0$:

$$R_{set}(\mathcal{F}, \gamma) = \{f \in \mathcal{F} : L(f) \leq L(f^*) + \gamma\},$$

where $f^* \in \mathcal{F}$ minimizes the true risk. $R_{ratio}(\mathcal{F}, \theta)$ denotes the Rashomon ratio for the true Rashomon set.

A large true Rashomon set, as it turns out, can be a certificate of the existence of a simpler model. Though, since we can never actually explore the true Rashomon set, we would never know whether it will be (or has been) useful for a particular problem. We explain this in Section 5, and spend most of our effort considering *empirical* Rashomon sets, which are easier to work with in practice.

When the hypothesis space has a parameterized representation (denote \mathcal{F}_Ω), we assume that we can parameterize each model $f \in \mathcal{F}_\Omega$ with a unique parameter vector $\omega \in \Omega$ of finite length and denote $f(z) = f_\omega(z)$. In the next section, we discuss properties of the Rashomon ratio as a complexity measure.

4 RASHOMON RATIO AS A SIMPLICITY MEASURE

The Rashomon ratio, as a *property of a data set and a hypothesis space*, serves as gauge of simplicity of the learning problem. If the Rashomon set is large, many different reasonable optimization procedures could lead to a model from the Rashomon set. Therefore, for large Rashomon sets, accurate models tend to be easier to find (since optimization procedures can find them). *In other words, if the*

Table 1: Comparison of Rashomon ratio and other complexity measures. The Rashomon ratio considers the fact that there are multiple good models and is a property both of the hypothesis space and data.

Complexity measure	Property of	Depends on data	Considers set of good models
VC Dimension	hypothesis space	no	no
Algorithmic stability (Hypothesis stability [5])	algorithm, hypothesis space	no	no
Empirical algorithmic stability (Algorithmic hypothesis stability [5])	algorithm, hypothesis space	yes	no
Geometric margins	one function	yes	no
Empirical Local Rademacher Complexity [2]	hypothesis space	depends on features, not on labels	no
Rashomon ratio	hypothesis space	yes, but not always on labels (see Theorem 10)	yes

Rashomon ratio is large, the Rashomon set could contain many accurate and simple models, and the learning problem becomes simpler. On the other hand, smaller Rashomon ratios might imply a harder learning problem, especially in the case of few deep and narrow local minima.

The Rashomon ratio can give insight into the simplicity of a learning problem, though it was designed for a fundamentally different goal than well-known complexity measures from learning theory (see Table 1). While those complexity measures were designed to help us understand generalization, the Rashomon ratio (with additional assumptions) helps us understand whether simpler functions might exist with the same level of accuracy as complex functions. The Rashomon ratio depends on a loss function, the hypothesis space, and a data set, while the majority of other measures are either data-agnostic or focus on properties of a specific model in the space.

The Rashomon ratio is different from VC dimension. The VC dimension [49] shows the expressive power of a hypothesis space for *any* data set including *an extreme* arrangement of data points and labels. On the contrary, the Rashomon set depends on an empirical risk minimizer that we compute directly for a specific data set, which may not be extreme.

The Rashomon ratio is different from algorithmic stability. Algorithmic stability [5] (see Definition B.2) depends on a change to a data set, whereas Rashomon Ratio uses a fixed data set. As we will show in Theorem 10 in Appendix B.1, in the case of linear least squares regression, the Rashomon ratio depends on features (X) only, and does not depend on regression targets Y . In contrast, hypothesis stability depends heavily on Y . In fact, if we can control how we change the set of targets, hypothesis stability (a form of algorithmic stability) can be made to change by an arbitrarily large

amount. This is formalized in Theorem 2 with proof in Appendix B.2.

THEOREM 2 (RASHOMON RATIO IS NOT ALGORITHMIC STABILITY). *Consider a distribution P_X over a discrete domain $X = \{x_1, \dots, x_N\}$ and a learning algorithm A that minimizes the sum of squares loss: $\|X\omega - Y\|_2^2$, for a linear hypothesis space \mathcal{F}_Ω . For any $\lambda > 0$, there exist joint distributions P_{X, Y_1} and P_{X, Y_2} where for X drawn i.i.d. from P_X , Y_1 drawn from $P_{Y_1|X}$ over \mathcal{Y} , and Y_2 drawn from $P_{Y_2|X}$ over \mathcal{Y} , the expected Rashomon ratios are the same:*

$$\mathbb{E}_{P_{X, Y_1}}[\hat{R}_{ratio_{S_1}}(\mathcal{F}_\Omega, \theta)] = \mathbb{E}_{P_{X, Y_2}}[\hat{R}_{ratio_{S_2}}(\mathcal{F}_\Omega, \theta)],$$

yet hypothesis stability constants are different by our arbitrarily chosen value of λ : $\hat{\beta}_2 - \hat{\beta}_1 \geq \lambda$, where S_1 and S_2 denote data sets $S_1 = [X, Y_1]$ and $S_2 = [X, Y_2]$, $\hat{\beta}_1$ is the hypothesis stability coefficient of algorithm \mathcal{A} for distribution P_{X, Y_1} and $\hat{\beta}_2$ is the hypothesis stability coefficient for distribution P_{X, Y_2} .

The Rashomon ratio is different from geometric margins. The margin (i.e., the minimum margin of the maximum margin classifier) [7, 42] depends on points closest to the decision boundary (support vectors), while the Rashomon set does not necessarily rely on the support vectors and may depend on the full data set. Theorem 3 shows this with proof in Appendix B.3.

THEOREM 3 (RASHOMON RATIO IS NOT THE GEOMETRIC MARGIN). *For any fixed $0 < \lambda < 1$, there exists a fixed hypothesis space \mathcal{F}_Ω , a Rashomon parameter θ , and there exist two data sets S_1 and S_2 with the same empirical risk minimizer $\hat{f} \in \mathcal{F}_\Omega$ such that the geometric margin d is the same for both data sets, yet the Rashomon ratios are different:*

$$|\hat{R}_{ratio_{S_1}}(\mathcal{F}_\Omega, \theta) - \hat{R}_{ratio_{S_2}}(\mathcal{F}_\Omega, \theta)| > \lambda.$$

The Rashomon ratio is different from empirical local Rademacher complexity. Empirical Rademacher complexity [2] (see Definitions B.4) measures how well the hypothesis space can fit random assignments of the labels. The Rashomon ratio uses fixed labels. It measures the number of models that are close to optimal. In other words, the Rashomon set benefits from having multiple similar models, while Rademacher complexity treats them as equivalent. Please see Theorem 4 with definition and proof in Appendix B.4.

THEOREM 4 (RASHOMON RATIO IS NOT LOCAL RADEMACHER COMPLEXITY). *For $0 < \lambda < 1$, there exist two data sets S_1 and S_2 , a hypothesis space \mathcal{F}_Ω , and a Rashomon parameter θ such that the local Rademacher complexities defined on the Rashomon sets for S_1 and S_2 are the same: $\hat{R}_n^{S_1}(\hat{R}_{set}(\mathcal{F}_\Omega, \theta)) = \hat{R}_n^{S_2}(\hat{R}_{set}(\mathcal{F}_\Omega, \theta))$, yet the Rashomon ratios are different: $|\hat{R}_{ratio_{S_1}}(\mathcal{F}_\Omega, \theta) - \hat{R}_{ratio_{S_2}}(\mathcal{F}_\Omega, \theta)| > \lambda$.*

Now that we have established that the Rashomon ratio is not the same as other simplicity measures, we can now shift our focus to proving simplicity and generalization properties of models in the Rashomon set. This is critical to our thesis that simple-yet-accurate models exist.

5 RASHOMON SET MODELS: SIMPLICITY AND GENERALIZATION

Consider two hypothesis (functional) spaces with different levels of complexity, where the lower-complexity space serves as a good *approximating set* (i.e., a good *cover*) for the higher-complexity space. The hypothesis spaces are called \mathcal{F}_1 , for the simpler space, and \mathcal{F}_2 , for the more complex space, where $\mathcal{F}_1 \subset \mathcal{F}_2$. Here, to determine the complexity of a hypothesis space, we use traditional notions of complexity (conversely, simplicity) such as covering numbers or VC dimension. For a useful example of a simple and a more complex space, consider \mathcal{F}_2 to be the space of linear models with real-valued coefficients in a space of d dimensions, and consider \mathcal{F}_1 to be the space of scoring systems [48], which are sparse linear models, with at most d' nonzero integer coefficients, $d' \ll d$. Another example is if the more complex space \mathcal{F}_2 consists of boosted decision trees, and \mathcal{F}_1 consists of single trees. Generalization bounds would be tighter if we could use the lower complexity space \mathcal{F}_1 , but as we are considering functions from \mathcal{F}_2 , learning theory often has us include the complexity of \mathcal{F}_2 in the bound. Given this setup, we have several questions to answer:

- (1) What if the higher-complexity hypothesis space we chose were more complex than necessary for modeling the data? In that case, if we had instead used the simpler model class \mathcal{F}_1 , would we still get a model that is (almost) as good as we could have obtained using the more complex class \mathcal{F}_2 ? If so, perhaps we can leverage the complexity of the simpler model class \mathcal{F}_1 for generalization bounds on our model rather than the more complex class \mathcal{F}_2 . We answer this question in Section 5.1, where a property on the complex space that will help us is that **the true Rashomon set of \mathcal{F}_2 is large enough to admit a simpler model**. We do not need to know what this model is and we may never discover it (we would likely discover a different model using data).
- (2) Under what conditions on the complex and simpler model classes does the property we mentioned above (that the Rashomon set includes simpler models) hold? Does it hold often? As it turns out, under natural conditions on the function class and loss function, a large Rashomon set in the complex class does imply the existence of simple-yet-accurate models. We identify these conditions in Section 5.2, namely that the loss function is smooth, and that \mathcal{F}_1 serves as a cover for \mathcal{F}_2 . Thus, **under these natural conditions that occur in practice, a large Rashomon set for a complex class of functions implies the existence of a simple-yet-accurate model**.

The bounds we present in Section 5.1 do not serve the same purpose as standard statistical learning theoretic bounds, as they do not aim to bound generalization error for a single function (that is, the difference between training and test loss for a function). Rather, we are interested in bounding train loss of one function (a *simpler* function) with test loss of another (the optimal model in a *more complex* function class). Standard learning theory analysis handles the single function case nicely; we are concerned with other questions here.

5.1 The True Rashomon Set Can Be Very Helpful... But You Might Not Know When

As in classic Occam's razor bounds, we start with finite hypothesis spaces. Consider finite hypothesis spaces \mathcal{F}_1 and \mathcal{F}_2 , where $\mathcal{F}_1 \subset \mathcal{F}_2$. Consider the first question discussed above: Given \mathcal{F}_1 and \mathcal{F}_2 , can we have a guarantee that a model we produce using a simpler function class \mathcal{F}_1 on our data could be approximately as good as the test performance of the best model from \mathcal{F}_2 ? In the following theorem, we will make a key assumption that allows us to do this: we assume that the Rashomon set of \mathcal{F}_2 includes a member of the simpler class of functions, \mathcal{F}_1 , even if we do not know which function it is. Later, in Section 5.2, we show conditions under which simple models from \mathcal{F}_1 are proven to exist in the Rashomon set of \mathcal{F}_2 , which depends on the size of \mathcal{F}_2 's Rashomon set. Here, $|\mathcal{F}|$ denotes the cardinality of the finite space \mathcal{F} . These bounds can be generalized to infinite hypothesis spaces with a simple extension to covering numbers, but they are designed for intuition, which works nicely with finite hypothesis spaces. Again, this is different from a regular learning theory bound as it does not consider generalization of just one function.

THEOREM 5 (THE ADVANTAGE OF A TRUE RASHOMON SET). *Consider finite hypothesis spaces \mathcal{F}_1 and \mathcal{F}_2 , such that $\mathcal{F}_1 \subset \mathcal{F}_2$. Let the loss l be bounded by b , $l(f_2, z) \in [0, b] \quad \forall f_2 \in \mathcal{F}_2, \forall z \in \mathcal{Z}$. Define an optimal function $f_2^* \in \argmin_{f_2 \in \mathcal{F}_2} L(f_2)$. Assume that the true Rashomon set includes a function from \mathcal{F}_1 , so there exists a model $\tilde{f}_1 \in \mathcal{F}_1$ such that $\tilde{f}_1 \in R_{\text{set}}(\mathcal{F}_2, \gamma)$. (Note that we do not know \tilde{f}_1 .) In that case, for any $\epsilon > 0$ with probability at least $1 - \epsilon$ with respect to the random draw of data:*

$$L(f_2^*) - b\sqrt{\frac{\log |\mathcal{F}_1| + \log 2/\epsilon}{2n}} \leq \hat{L}(\tilde{f}_1) \leq L(f_2^*) + \gamma + b\sqrt{\frac{\log 1/\epsilon}{2n}}, \quad (1)$$

where $\hat{f}_1 \in \argmin_{f_1 \in \mathcal{F}_1} \hat{L}(f_1)$. (Unlike \tilde{f}_1 , we do know \hat{f}_1 because we can calculate it.)

That is, we can bound the best empirical model from \mathcal{F}_1 with the true risk of the best model within \mathcal{F}_2 . Thus, if the Rashomon set is large enough to include a single model from \mathcal{F}_1 , we can work with the simpler class \mathcal{F}_1 in practice and achieve strong performance guarantees.

The main assumption in Theorem 5 is about the population, and does not rely on the sample. It relies only on the existence of one special function in the true Rashomon set. There are no smoothness assumptions on the loss function. If the main assumption of this theorem holds, then we gain the benefit of guarantees on \mathcal{F}_2 from looking only at \mathcal{F}_1 empirically. We cannot check whether the assumption holds since it involves the true risk, but practitioners can reap the benefits of it anyway: The possibility of a large Rashomon set may embolden the practitioner to minimize over \mathcal{F}_1 , achieving test error close to the best of \mathcal{F}_2 if the conditions of Theorem 5 are indeed satisfied.

To make the connection of this result to Rashomon sets more explicit, we will choose a specific relationship between \mathcal{F}_1 and \mathcal{F}_2 , specifically, \mathcal{F}_1 will be a random sample of \mathcal{F}_2 that is chosen prior to, and separately from, learning. This is an artificial example in that \mathcal{F}_1 would never actually be chosen as a random sample from \mathcal{F}_2 in reality. However, the random sampling assumption permits

Table 2: Examples of the possible usage of Theorem 6.

If $ \mathcal{F}_1 = 100000$ then to get the bound (1) to hold with probability at least 99%
the Rashomon ratio should be $R_{\text{ratio}}(\mathcal{F}_2, \gamma) \geq 0.0053\%$.
If $ \mathcal{F}_1 = 10000$ then to get the bound (1) to hold with probability at least 99%
the Rashomon ratio should be $R_{\text{ratio}}(\mathcal{F}_2, \gamma) \geq 0.053\%$.
If $ \mathcal{F}_1 = 1000$ then to get the bound (1) to hold with probability at least 99%
the Rashomon ratio should be $R_{\text{ratio}}(\mathcal{F}_2, \gamma) \geq 0.53\%$.

\mathcal{F}_1 to be distributed fairly evenly within \mathcal{F}_2 , which, arguably, could approximate the way some simpler spaces are embedded in more complex spaces.

If \mathcal{F}_1 is a random sample of functions from \mathcal{F}_2 , and if \mathcal{F}_2 has a large true Rashomon set, then the true Rashomon set is likely to include at least one model from \mathcal{F}_1 . In that case, Theorem 5 applies. This is formalized below.

THEOREM 6 (EXAMPLE OF THE ADVANTAGE OF A LARGE TRUE RASHOMON SET). *Consider finite hypothesis spaces \mathcal{F}_1 and \mathcal{F}_2 , such that $\mathcal{F}_1 \subset \mathcal{F}_2$ and \mathcal{F}_1 is uniformly drawn from \mathcal{F}_2 without replacement. For loss l bounded by b , if the Rashomon ratio is at least*

$$R_{\text{ratio}}(\mathcal{F}_2, \gamma) \geq 1 - \epsilon^{\frac{1}{|\mathcal{F}_1|}}$$

then for any $\epsilon > 0$, with probability at least $(1 - \epsilon)^2$ with respect to the random draw of functions from \mathcal{F}_2 to form \mathcal{F}_1 and with respect to the random draw of data, the assumptions of Theorem 5 hold and thus the bound (1) holds.

Table 2 shows possible values of the lower bound on the Rashomon ratio, given $|\mathcal{F}_1|$ and ϵ . For example, the first line of the table states that if at least a tiny fraction (0.0053%) of the complex function space \mathcal{F}_2 consists of good models, and there exists at least 100,000 simple functions in \mathcal{F}_1 , then the chance that we will find an accurate-but-simple model on our data set is over 99%.

The intuition for Theorem 6 holds beyond the case when \mathcal{F}_1 is randomly sampled from \mathcal{F}_2 , it holds whenever \mathcal{F}_1 covers \mathcal{F}_2 sufficiently well. This intuition is that as the true Rashomon ratio increases, it is more likely that the empirical risk minimum of \mathcal{F}_1 will be close to the minimum of the true risk of \mathcal{F}_2 .

5.2 Proving the Existence of Simple-yet-Accurate Models with Good Generalization

Theorems 5 and 6 do not take advantage of the fact that we can investigate \mathcal{F}_2 empirically, and *more easily than we can investigate \mathcal{F}_1* ; these theorems instead only discuss exploration of \mathcal{F}_1 . Thus, the next analysis makes two improvements: (1) it studies empirical Rashomon sets instead of true Rashomon sets, (2) it substitutes the unrealistic random draw assumption for a realistic smoothness assumption. We now assume smoothness of the loss over the function space.

The field of Approximation Theory provides general conditions under which classes of functions can approximate each other. Given

a target function from one class, we want to know whether a sequence of functions from another class can converge to the target. Table 3 in Appendix C.4 shows classes of functions \mathcal{F}_2 that can be approximated by classes \mathcal{F}_1 . For instance, piecewise constant functions, such as decision trees, can approximate smooth functions.

For a hypothesis space \mathcal{F} and some $f' \in \mathcal{F}$, define the δ -ball of functions centered at f' as $B_\delta(f') = \{f \in \mathcal{F} : \|f' - f\|_p \leq \delta\}$. A loss $l : \mathcal{F} \times \mathcal{X} \rightarrow \mathcal{Y}$ is said to be K -Lipschitz, $K \geq 0$, if for all $f_1, f_2 \in \mathcal{F}$ and for all $z \in \mathcal{Z}$: $|l(f_1, z) - l(f_2, z)| \leq K\|f_1 - f_2\|_p$. The p -norm can be defined, for example, as $\|f\|_p = \left(\int_{\mathcal{X}} |f|^p d\mu\right)^{1/p}$, where μ is a measure on \mathcal{X} . Define a δ -packing as a finite set $\Xi = \{\xi_1, \dots, \xi_k | \xi_i \in \mathcal{F}\}$ such that $\|\xi_i - \xi_j\|_p > \delta$, meaning that $B_{\delta/2}(\xi_i) \cap B_{\delta/2}(\xi_j) = \emptyset$ for all $i \neq j$. The *packing number* $\mathcal{B}(\mathcal{F}, \delta)$ is the largest δ -packing.

Theorem 7 below uses the approximating set argument from the previous subsection, but now requires the Rashomon set to be large enough to include balls of functions rather than using the random draw assumption. As long as the set of simpler functions is distributed well among the full hypothesis space, each ball contains at least one function from the simpler class.

THEOREM 7 (EXISTENCE OF MULTIPLE SIMPLER MODELS). *For K -Lipschitz loss l bounded by b , consider hypothesis spaces \mathcal{F}_1 and \mathcal{F}_2 , $\mathcal{F}_1 \subset \mathcal{F}_2$. With probability greater than $1 - \epsilon$ w.r.t. the random draw of training data, if for every model $f_2 \in \hat{R}_{set}(\mathcal{F}_2, \theta)$ there exists $f_1 \in \mathcal{F}_1$ such that $\|f_2 - f_1\|_p \leq \delta$, then there exists at least $B = \mathcal{B}(\hat{R}_{set}(\mathcal{F}_2, \theta), 2\delta)$ functions $\bar{f}_1^1, \bar{f}_1^2, \dots, \bar{f}_1^B \in \hat{R}_{set}(\mathcal{F}_1, \theta)$ such that:*

- (1) *They are from the simpler space: $\bar{f}_1^1, \bar{f}_1^2, \dots, \bar{f}_1^B \in \mathcal{F}_1$.*
- (2) $\left|L(\bar{f}_1^i) - \hat{L}(\bar{f}_1^i)\right| \leq 2KR_n(\mathcal{F}_1) + b\sqrt{\frac{\log(2/\epsilon)}{2n}}$, for all $i \in [1, \dots, B]$, where $R_n(\mathcal{F})$ is the Rademacher complexity of a hypothesis space \mathcal{F} . (This is from standard learning theory.)

From Theorem 7, we see that since larger Rashomon sets have larger packing numbers, they contain more simpler models with good generalization guarantees. Note that in Theorem 7, other complexity measures from learning theory could be used. We chose Rademacher complexity as it provides the tightest bound among standard complexity measures.

Theorem 7 has practical implications. *If the Rashomon set is large, and the smoothness conditions are obeyed, Theorem 7 shows that many simple-yet-accurate models would exist, prior to actually finding them.* Knowledge that simple models exist implies it will be worthwhile to actually solve the difficult optimization problem to find a simple model.

Thus, if the Rashomon set is large, we have a guarantee. But how will we know when is the Rashomon set large? This is what we answer in the next section.

6 LARGER RASHOMON RATIOS CORRELATE WITH SIMILAR PERFORMANCE OF MACHINE LEARNING ALGORITHMS, AND GOOD GENERALIZATION

We expect that in many real-world applications of machine learning, properties similar to the assumptions behind our theorems hold,

i.e., that large enough Rashomon sets intersect simpler hypothesis spaces in ways that lead to or explain good performance. This conjecture is difficult to verify theoretically because it is not a mathematical conjecture about the structure of two specific function spaces, but a statement about many function spaces, and how they interact with commonly occurring data sets. Thus, we consider this question empirically.

Our experiments will demonstrate that, in the case where Rashomon sets are large, two conclusions follow that are consistent with our theoretical development. First, *training* performance in *simpler* hypothesis spaces is correlated with *test* performance in the more *complex* hypothesis spaces (Theorem 5), and second, that good *training* performance in a *simpler* space \mathcal{F}_1 correlates with good *generalization* performance of other models in the more *complex* space \mathcal{F}_2 . Most importantly, our experiments suggest an intriguing alternative to the often difficult computational problem of directly estimating the size of the Rashomon set, namely that *similar performance across a range of algorithms with different hypothesis spaces is strongly correlated with a large Rashomon set*.

Now we will describe our experimental setup for arriving at these conclusions.

6.1 Experimental Design

Data sets. We used 38 machine learning classification data sets from the UCI Machine Learning Repository [18], among which 16 have categorical features and 22 have real-valued features. The majority of the data sets are binary classification data sets and we adapted the rest to binary classification (as shown in Table 4 in Appendix D) to make importance sampling easier (as discussed in Appendix E). The number of features varies from 3 to 784, with the majority of the data sets being in the 15–25 feature range. Appendix D contains a description of the data sets we considered.

Definition of complex hypothesis space. For these experiments, we will consider \mathcal{F}_2 to be the union (\mathcal{F}_{union}) of the hypothesis spaces of five popular machine learning algorithms: logistic regression (LR), CART, random forests (RF), gradient boosted trees (GBT), and support vector machines with RBF kernels (SVM). CART, RF and GBT were regularized by varying the tree depth, the minimum number of samples required to split a node, the minimum number of samples required to create a leaf node, and the number of trees in the ensemble. SVMs were tuned by varying the regularization parameter and the kernel coefficient and LR by varying the regularization parameter. Appendix E discusses the effect of regularization on the model class. We chose algorithms that search hypothesis spaces of different complexity to ensure that these algorithms produce diverse models. The notion of \mathcal{F}_2 as a union of hypothesis spaces may seem surprising at first, but it is consistent with how many machine learning practitioners approach problems by running a collection of machine learning techniques in parallel and comparing the results, creating a *de facto* union space. Our experiment has three steps, as follows.

Step 1: Run all machine learning algorithms. We obtain training and generalization performance from all algorithms (logistic regression, CART, random forests, gradient boosted trees, and SVM with RBF kernels) on all data sets.

Step 2: Estimate the size of the Rashomon set. It is not possible to measure the Rashomon set of such a complex model space, so we will estimate its size by sampling from an approximating set, which is decision trees of bounded depth. Decision trees are easy to sample and can refine an input space arbitrarily finely as tree depth increases. With sufficient depth they can approximate many other types of hypothesis spaces, including those used by other machine learning methods. Thus, we will measure the size of the Rashomon set and Rashomon ratio in decision trees of depth seven as a surrogate for measuring these quantities in \mathcal{F}_2 . The suitability of these trees for this role is an empirical observation about the data sets we have used; they may not be a suitable surrogate for some other data sets, e.g., imagery data. We measure the size of the empirical Rashomon ratio as a surrogate for the true Rashomon ratio when referring to Theorem 5. To estimate the Rashomon ratio of depth seven decision trees, we used importance sampling. The proposal distribution assigns the correct labels to the leaves of the tree based on the training data. Since the data are populated on a bounded domain, to grow a tree up to a depth D fully, we make 2^{D-1} splits. For each data set and each depth, we average our results over ten folds for data sets with less than 200 points and over five folds for data sets with more than 200 points, and we sample 250,000 decision trees per fold. We choose the Rashomon parameter θ to be 5%, and, therefore, all the models in the Rashomon set have empirical risk not more than $\hat{L}(\hat{f}) + 0.05$, where $\hat{L}(\hat{f})$ is the lowest achievable empirical risk across all algorithms we considered. We further discuss experimental setup in Appendix E.

Step 3: See if a large Rashomon Set in Step 2 correlates with performance differences in Step 1. By construction, the hypothesis spaces of each of the machine learning algorithms we consider are embedded in \mathcal{F}_2 . RF and GBT both enjoy extremely rich hypothesis spaces that are likely close in size to \mathcal{F}_2 itself. LR and CART are less expressive than these others, so we will view LR and CART as simpler, \mathcal{F}_1 type, hypothesis spaces. Our question to answer is whether a large Rashomon set measured in Step 2 correlates with the functions from \mathcal{F}_1 (CART, LR) having performance as good as that of \mathcal{F}_2 (GBT, RF, SVM) as our theory predicts it will.

6.2 Experimental Results

Figure 2(a) shows the performance of the five machine learning algorithms on data sets for which the Rashomon ratio was largest, as measured in the space of decision trees of depth 7. Performance for all data sets is shown in Figures 14 and 15 in Appendix E. Across the 38 data sets considered, we observe **larger Rashomon ratios led to approximately similar training results across all algorithms** (within $\sim 5\%$ difference between algorithms). Here, large Rashomon ratios are on the order of $10^{-37}\%$ or $10^{-38}\%$, whereas small Rashomon ratios are $10^{-40}\%$ or less². Moreover, all of the models chosen by the algorithms, including simpler \mathcal{F}_1 type models, generalized well (the differences between training and test errors are within $\sim 5\%$). These results are consistent with our thesis that larger Rashomon sets lead to the existence of accurate-yet-simpler models (in agreement with the theory in Section 5.1), and that larger Rashomon sets lead to better generalization. *The results also*

imply that large Rashomon sets do occur in many data sets, with the Rashomon effect being large enough to include simpler models in practice (in agreement with Section 5.2).

Interestingly, the converse statement, that similar performance across different algorithms should lead to large Rashomon sets, does not always hold; sometimes, generalization occurs with small Rashomon ratios (see Figure 2(b)). This observation could be explained in several different ways. Mainly, the Rashomon ratio is not the only driver of good generalization performance. The amount of data is one obvious additional driver. Appendix F discusses this further. Quality of features is another driver, as discussed in Appendix G.

Our second main result is that in **all** cases where large Rashomon ratios were observed, **test performance was consistent with training performance across algorithms of varying complexity**. This correlation between the size of the Rashomon ratio and consistent generalization performance suggests an indirect means of assessing the size of the Rashomon ratio as an alternative to the computationally intensive approach of sampling. When consistent training and test performance across algorithms is observed, this may indicate a large Rashomon ratio.

One thing we notably did not observe were cases where algorithms did not generalize, performance differed across algorithms, and the Rashomon set was large. Across all 38 data sets, we did not observe cases where the Rashomon set was large and performance differed among algorithms.

Figure 2(c) shows small Rashomon sets, where we observe wildly different performance across algorithms, where sometimes the models generalize and sometimes they do not. We show one example of each of these cases in Figure 2(c). Our theory does not apply to the case of small Rashomon sets, and thus there is no guarantee for such data sets.

7 RASHOMON SETS IN THE PRESENCE OF NOISE

We have seen that we can empirically determine whether a data set is likely to have a large Rashomon set: as we showed, we simply run many algorithms, and if they all perform similarly and generalize, there could be a large Rashomon set. But what about before examining the data? Could we know, just from understanding what kind of data set it is, whether it is likely to have a large Rashomon set? We aim to answer this now.

A typical reason given for “underspecification” [13] (i.e., a large number of approximately-equally good models, a large Rashomon set) is the presence of substantial noise in the data. Intuitively, for data whose outcomes are essentially a random guess, it makes sense that no model would perform well, and many models would be equally poor. But what about more interesting cases? Does this intuition still hold? First, as we have shown above, large Rashomon sets exist for non-noisy data sets as well (see Figure 2 Voting and HTRU 2 data sets); *it is not just noise that determines the Rashomon set*. Second, it is not true that the Rashomon set always gets much larger under noise. In fact, if we add random classification noise [1, 35] (each label is flipped independently with some small probability), it is possible that the Rashomon set does not change at all. This is because the error rate of all models in the Rashomon set (assuming

²For other data sets and other metrics of measuring the Rashomon set, the results might be different.

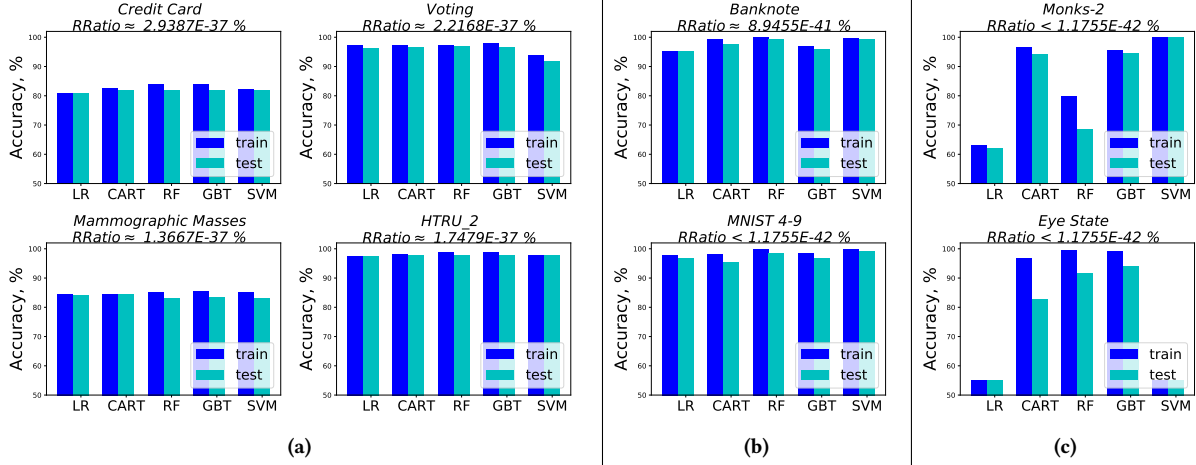


Figure 2: (a) Examples of experiments on four data sets showing that larger Rashomon ratios lead to similar performance of five machine learning algorithms with regularization. All the algorithms generalize well and have similar test accuracy. (b)-(c): Examples showing that smaller Rashomon ratios do not necessarily imply a performance difference between machine learning algorithms. Even with low Rashomon ratios, algorithms can be highly accurate and generalize well, as shown in Figure (b). On the other hand, when the Rashomon ratio is small, sometimes algorithms can perform differently or fail to generalize, as shown in Figure (c). In the figure, test accuracies, training accuracies and the Rashomon ratio are averaged over ten folds. We show all 38 data sets in the Appendix E.

they are all better than random guessing) increases by the same amount in expectation. At least, as we show in Theorem 8, the size of the true Rashomon set does not decrease if we add random classification noise.

THEOREM 8 (EXPECTED SIZE OF THE TRUE RASHOMON SET CANNOT DECREASE UNDER RANDOM CLASSIFICATION NOISE). Consider hypothesis space \mathcal{F} , data distribution $\mathcal{D} = \mathcal{X} \times \mathcal{Y}$, where, as before, $\mathcal{X} \in \mathbb{R}^p$, and $\mathcal{Y} \in \{-1, 1\}$. Let $\rho \in (0, \frac{1}{2})$ be a probability with which each label y_i is flipped independently, and \mathcal{D}_ρ denotes the noisy version of \mathcal{D} . If the loss function is $\phi(f(x), y) = \mathbb{1}_{[f(x) \neq y]}$, then in expectation, the true Rashomon set over \mathcal{D} is a subset of the true Rashomon set over \mathcal{D}_ρ , $R_{\text{set}_{\mathcal{D}}}(\mathcal{F}, \gamma) \subseteq R_{\text{set}_{\mathcal{D}_\rho}}(\mathcal{F}, \gamma)$.

In Theorem 8 we have shown that the size of the true Rashomon set does not decrease when adding random classification noise, but to prove that $R_{\text{set}_{\mathcal{D}}}(\mathcal{F}, \gamma) \subset R_{\text{set}_{\mathcal{D}_\rho}}(\mathcal{F}, \gamma)$, we would need at least one model \tilde{f} such that $\tilde{f} \notin R_{\text{set}_{\mathcal{D}}}(\mathcal{F}, \gamma)$, yet $\tilde{f} \in R_{\text{set}_{\mathcal{D}_\rho}}(\mathcal{F}, \gamma)$, and such a model may not actually exist; in fact, if all models have increased error rates when noise is added, it does not.

We still are left with finding a scenario where noise does impact the size of the Rashomon set in order to provide some proof to the intuition. Let us consider the setting of linear (Gaussian) discriminant analysis, where the data arise from two Gaussians, one with positive labels and one with negative labels. Instead of increasing label noise, we will increase feature noise by increasing the variances or changing the means of the Gaussians, making the two distributions overlap. In this case, will the Rashomon set increase in size? The answer to this question is conjectured in Conjecture 9.

CONJECTURE 9 (THE RASHOMON SET CAN INCREASE WITH FEATURE NOISE). Consider data distribution $\mathcal{D} = \mathcal{X} \times \mathcal{Y}$, where, $\mathcal{X} \in \mathbb{R}$, $\mathcal{Y} \in \{-1, 1\}$, and classes are balanced $P(Y = -1) = P(Y = 1)$

and generated by Gaussian distributions $P(X|Y = -1) = \mathcal{N}(\mu_1, \sigma^2)$, $P(X|Y = 1) = \mathcal{N}(\mu_2, \sigma^2)$, where $0 \leq \mu_1 < \mu_2$. For the hypothesis space $\mathcal{F} = \{f : f \in (\beta_1, \beta_2)\}$, where $(\mu_1, \mu_2) \subset (\beta_1, \beta_2)$, $\beta_1 \ll \mu_1$, and $\mu_2 \ll \beta_2$, and the Rashomon parameter $\gamma > 0$:

- (I) The volume of the Rashomon set is $\mathcal{V}(R_{\text{set}_{\sigma}}(\mathcal{F}, \gamma)) = |f_{\sigma}^{e_1} - f_{\sigma}^{e_2}|$, where $f_{\sigma}^{e_1}$ and $f_{\sigma}^{e_2}$ are the two solutions to Eqn. (2), where Φ is the CDF of the standard normal:

$$2\Phi\left(\frac{\mu_2 - \mu_1}{2\sigma}\right) - \Phi\left(\frac{\mu_2 - f}{\sigma}\right) - \Phi\left(\frac{f - \mu_1}{\sigma}\right) = \gamma. \quad (2)$$

- (II) We conjecture that³ for $\mathcal{F} = \{f : f \in (\mu_1, \mu_2)\}$, as we add feature noise to the data set by increasing the standard deviation σ , for all σ such that $\sigma > \bar{\sigma} = \frac{\mu_2 - \mu_1}{2\sqrt{2}}$, the volume of the Rashomon set increases as a function of σ .
- (III) Consider the setting where $\sigma = 1$ for both Gaussians, and we add or remove noise by moving the means μ_1 and μ_2 of the Gaussians towards or away from each other. For any $\gamma > 0$, the volume of the Rashomon set is minimized when $\mu_2 \approx \mu_1 + 2$. Moving the Gaussians either away from or towards each other increases the volume of the Rashomon set.

This conjecture is not a theorem because there is no analytical solution to the minimizer of the volume of the Rashomon set; the calculations are quite complex, involving differences of the CDF values of different Gaussians. However, all parts of the conjecture have been fully checked numerically. In part (II), we use an analytical derivation and exhaustive numerical computations to show that the derivatives of the left side of Eqn. (2) are either positive or negative

³The hypothesis space for Part II conservatively includes all reasonable candidates for the empirical risk minimizer. In other words, we assume that decision boundary can be anywhere between the means of the two distributions.

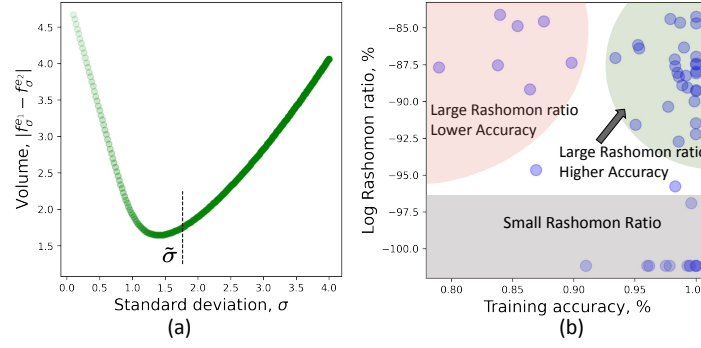


Figure 3: (a) Dependence of the Rashomon ratio on noise σ for the two Gaussians example in Conjecture 9. When $\sigma > \tilde{\sigma}$, as we add more noise, the size of the Rashomon set increases. (2). (b) Scatter plot of the log of Rashomon ratio versus the maximum accuracy across five different algorithms for 38 data sets in Section 6. We observe larger Rashomon ratios in both noisy and non-noisy data sets.

sign. For part (III), we transpose the left Gaussian to $N(0, 1)$ to form a canonical problem in which all possible solutions can be computed numerically. We exhaustively search over the range of γ , finding the optimal μ_2 and volume of the Rashomon set for each γ . We find that μ_2 is very close to 2 for all γ . We discuss this in Appendix I.

This conjecture suggests that **data that are approximately distributed according to two normal distributions, where the positive and negative normal distributions substantially overlap, will have a large Rashomon set.** Figure 3(a) shows the dependence of the Rashomon set on the noise level σ for $\mu_1 = 1, \mu_2 = 6$ and $\sigma \in [0.2, 4]$. Figure 3(b) plots maximum accuracy versus Rashomon ratio for 38 data sets considered in Section 6. These figures indicate that large Rashomon sets occur both in noisy and non-noisy data.

There can be many data sets with characteristics as in Conjecture 9. For example, let us consider criminal recidivism data, whose Rashomon sets have been studied [17, 19] and that admit simple-yet-accurate models [41, 52]. Each data point is generated based on a set of random events happening in the world; whether someone enters a job training program, whether someone associates with criminal associates after release, and whether someone commits a crime each day are all random variables whose random effects are cumulative over time, and thus could be modeled by Gaussians by the central limit theorem. By this logic, we would expect many criminal recidivism prediction problems to admit large Rashomon sets. Other high-stakes predictions such as loan defaults may have similar characteristics.

In a sense, this full analysis paints a much clearer picture as to why such problems admit simple yet similarly accurate models: their distributions are approximately Gaussian with significant overlap, such overlap leads to large Rashomon sets, and large Rashomon sets lead to the existence of simple yet similarly accurate models.

8 CONCLUSION AND IMPLICATIONS

We have proposed Rashomon sets and ratios as another perspective on the relationship between hypothesis spaces and data sets, and we have provided initial theoretical and experimental results

showing that this is a unique perspective that may help explain some phenomena observed in practice. More specifically, the main conclusions include: (1) Large Rashomon sets can embed models from simpler hypothesis spaces (Section 5); (2) Similar performance across different machine learning algorithms may correlate with large Rashomon sets (Section 6); (3) Large Rashomon sets correlate with existence of models that have good generalization performance (Section 6); (4) The Rashomon ratio is a measure of a learning problem’s complexity (Section 4), and that data that approximately arise from overlapping Gaussian distributions tend to have large Rashomon sets (Section 7).

Consider a researcher conducting a standard set of machine learning experiments in which the performance of several different algorithms are compared, and generalization is assessed. In the possible scenario where *all algorithms perform similarly*, and when their models tend to generalize well on validation data, the learning problem is likely to have a large Rashomon set. Based on the result in Section 5, simpler models are likely to exist in a large Rashomon set. If the researcher is interested in simpler models, they can search the simpler function class \mathcal{F}_1 , a subset of the larger class \mathcal{F}_2 , to locate simpler models within it. While optimizing for simplicity or interpretability constraints is usually much more computationally expensive than running standard machine learning algorithms, our thesis is that this search would be likely to succeed in the presence of a large Rashomon set. In the converse case, if the researcher’s algorithms perform *differently from each other*, the researcher might then select a more complex model class that achieves better performance yet does not overfit. Further, if the researcher knows that the data are likely to have arisen from overlapping Gaussian distributions, the researcher could assume that it is worthwhile to search for a simple model that performs well.

REFERENCES

- [1] Dana Angluin and Philip Laird. 1988. Learning from noisy examples. *Machine Learning* 2, 4 (1988), 343–370.
- [2] Peter L Bartlett, Olivier Bousquet, Shahar Mendelson, et al. 2005. Local Rademacher complexities. *The Annals of Statistics* 33, 4 (2005), 1497–1537.
- [3] Peter L Bartlett and Shahar Mendelson. 2002. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research* 3, Nov (2002), 463–482.

- [4] Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. 2019. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences* 116, 32 (2019), 15849–15854.
- [5] Olivier Bousquet and André Elisseeff. 2002. Stability and generalization. *Journal of Machine Learning Research* 2, Mar (2002), 499–526.
- [6] Leo Breiman et al. 2001. Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Statist. Sci.* 16, 3 (2001), 199–231.
- [7] Christopher JC Burges. 1998. A tutorial on support vector machines for pattern recognition. *Data Mining and Knowledge Discovery* 2, 2 (1998), 121–167.
- [8] Feilong Cao, Tingfan Xie, and Zongben Xu. 2008. The estimate for approximation error of neural networks: A constructive approach. *Neurocomputing* 71, 4-6 (2008), 626–630.
- [9] Pratik Chaudhari, Anna Choromanska, Stefano Soatto, Yann LeCun, Carlo Baldassi, Christian Borgs, Jennifer Chayes, Levent Sagun, and Riccardo Zecchina. 2019. Entropy-SGD: Biasing gradient descent into wide valleys. *Journal of Statistical Mechanics: Theory and Experiment* 2019, 12 (2019), 124018.
- [10] Beau Coker, Cynthia Rudin, and Gary King. 2021. A theory of statistical inference for ensuring the robustness of scientific results. *Management Science* 67 (2021), 5969–6627. Issue 10.
- [11] Corinna Cortes and Vladimir Vapnik. 1995. Support-vector networks. *Machine learning* 20, 3 (1995), 273–297.
- [12] Amanda Coston, Ashesh Rambachan, and Alexandra Chouldechova. 2021. Characterizing fairness over the set of good models under selective labels. In *Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research)*.
- [13] Alexander D'Amour, Katherine Heller, Dan Moldovan, Ben Adlam, Babak Alipanahi, Alex Beutel, Christina Chen, Jonathan Deaton, Jacob Eisenstein, Matthew D Hoffman, et al. 2020. Underspecification presents challenges for credibility in modern machine learning. *arXiv preprint arXiv:2011.03395* (2020).
- [14] Oleg Davydov. 2011. Algorithms and error bounds for multivariate piecewise constant approximation. In *Approximation Algorithms for Complex Systems*. Vol. 3. Springer, Berlin, Heidelberg, 27–45.
- [15] Ronald A DeVore. 1998. Nonlinear approximation. *Acta Numerica* 7 (1998), 51–150.
- [16] Laurent Dinh, Razvan Pascanu, Samy Bengio, and Yoshua Bengio. 2017. Sharp minima can generalize for deep nets. In *Proceedings of the 34th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 70)*. 1019–1028.
- [17] Jiayun Dong and Cynthia Rudin. 2020. Exploring the cloud of variable importance for the set of all good models. *Nature Machine Intelligence* 2, 12 (2020), 810–824.
- [18] Dheeru Dua and Casey Graff. 2019. UCI Machine Learning Repository. University of California, Irvine, School of Information and Computer Sciences.
- [19] Aaron Fisher, Cynthia Rudin, and Francesca Dominici. 2019. All models are wrong, but many are useful: Learning a variable's importance by studying an entire class of prediction models simultaneously. *Journal of Machine Learning Research* 20, 177 (2019), 1–81.
- [20] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Flat minima. *Neural Computation* 9, 1 (1997), 1–42.
- [21] Sham M. Kakade, Karthik Sridharan, and Ambuj Tewari. 2008. On the complexity of linear prediction: risk bounds, margin bounds, and regularization. In *Proceedings of the 21st International Conference on Neural Information Processing Systems* (Vancouver, British Columbia, Canada) (*NIPS'08*). Curran Associates Inc., Red Hook, NY, USA, 793–800.
- [22] Michael Kearns and Dana Ron. 1999. Algorithmic stability and sanity-check bounds for leave-one-out cross-validation. *Neural Computation* 11, 6 (1999), 1427–1453.
- [23] Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. 2016. On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836* (appeared at *ICLR 2017*) (2016).
- [24] Vladimir Koltchinskii and Dmitry Panchenko. 2002. Empirical margin distributions and bounding the generalization error of combined classifiers. *The Annals of Statistics* 30, 1 (2002), 1–50.
- [25] John Langford and John Shawe-Taylor. 2002. PAC-Bayes & margins. In *Proceedings of the 15th International Conference on Neural Information Processing Systems* (*NIPS'02*). MIT Press, Cambridge, MA, USA, 439–446.
- [26] Guillaume Lecué. 2011. *Interplay between concentration, complexity and geometry in learning theory with applications to high dimensional data analysis*. Ph.D. Dissertation. Université Paris-Est.
- [27] Benjamin Letham, Portia A. Letham, Cynthia Rudin, and Edward Browne. 2016. Prediction uncertainty and optimal experimental design for learning dynamical systems. *Chaos* 26, 6 (2016).
- [28] Gábor Lugosi and Marten Wegkamp. 2004. Complexity regularization via localized random penalties. *The Annals of Statistics* 32, 4 (2004), 1679–1697.
- [29] David Madras, James Atwood, and Alex D'Amour. 2019. Detecting underspecification with local ensembles. *arXiv preprint arXiv:1910.09573* (appeared at *ICLR 2020 under the title "Detecting Extrapolation with Local Ensembles"*) (2019).
- [30] Naresh Manwani and PS Sastry. 2013. Noise tolerance under risk minimization. *IEEE Transactions on Cybernetics* 43, 3 (2013), 1146–1151.
- [31] Charles Marx, Flavio Calmon, and Berk Ustun. 2020. Predictive Multiplicity in Classification. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119)*. PMLR, 6765–6774.
- [32] Nicolai Meinshausen and Peter Bühlmann. 2010. Stability selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 72, 4 (2010), 417–473.
- [33] Shahar Mendelson. 2003. A few notes on statistical learning theory. In *Advanced Lectures on Machine Learning*. Springer, 40 pages.
- [34] Preetum Nakkiran, Gal Kaplun, Yamini Bansal, Tristan Yang, Boaz Barak, and Ilya Sutskever. 2021. Deep double descent: Where bigger models and more data hurt. *Journal of Statistical Mechanics: Theory and Experiment* 2021, 12 (2021), 124003.
- [35] Nagarajan Natarajan, Inderjit S Dhillon, Pradeep K Ravikumar, and Ambuj Tewari. 2013. Learning with noisy labels. In *Advances in Neural Information Processing Systems*, C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger (Eds.), Vol. 26. Curran Associates, Inc.
- [36] Daniel Nevo and Ya'acov Ritov. 2017. Identifying a minimal class of models for high-dimensional data. *The Journal of Machine Learning Research* 18, 1 (2017), 797–825.
- [37] DJ Newman and TJ Rivlin. 1976. Approximation of monomials by lower degree polynomials. *Aequationes Mathematicae* 14, 3 (1976), 451–455.
- [38] Ramamohan Paturi. 1992. On the degree of polynomials that approximate symmetric boolean functions (preliminary version). In *Proceedings of the Twenty-Fourth Annual ACM Symposium on Theory of Computing* (Victoria, British Columbia, Canada) (*STOC '92*). Association for Computing Machinery, New York, NY, USA, 468–474.
- [39] William H Rogers and Terry J Wagner. 1978. A finite sample distribution-free performance bound for local discrimination rules. *The Annals of Statistics* 6, 3 (1978), 506–514.
- [40] Cynthia Rudin. 2019. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence* 1, 5 (2019), 206–215.
- [41] Cynthia Rudin, Caroline Wang, and Beau Coker. 2020. The Age of Secrecy and Unfairness in Recidivism Prediction. *Harvard Data Science Review* 2, 1 (31 1 2020).
- [42] Robert E Schapire, Yoav Freund, Peter Bartlett, Wee Sun Lee, et al. 1998. Boosting the margin: A new explanation for the effectiveness of voting methods. *The Annals of Statistics* 26, 5 (1998), 1651–1686.
- [43] Lesia Semenova, Cynthia Rudin, and Ronald Parr. 2019. A study in Rashomon curves and volumes: A new perspective on generalization and model simplicity in machine learning. *arXiv preprint arXiv:abs/1908.01755* (2019).
- [44] Nathan Srebro, Karthik Sridharan, and Ambuj Tewari. 2010. Smoothness, Low Noise and Fast Rates. In *Advances in Neural Information Processing Systems*, Vol. 23. Curran Associates, Inc., 2199–2207.
- [45] Theja Tulabandhula and Cynthia Rudin. 2013. Machine learning with operational costs. *The Journal of Machine Learning Research* 14, 1 (2013), 1989–2028.
- [46] Theja Tulabandhula and Cynthia Rudin. 2014. On combining machine learning with decision making. *Machine Learning (ECML-PKDD journal track)* 97, 1-2 (2014), 33–64.
- [47] Theja Tulabandhula and Cynthia Rudin. 2014. Robust optimization using machine learning for uncertainty sets. *arXiv preprint arXiv:1407.1097* (2014).
- [48] Berk Ustun and Cynthia Rudin. 2016. Supersparse linear integer models for optimized medical scoring systems. *Machine Learning* 102, 3 (2016), 349–391.
- [49] VN Vapnik and A Ya Chervonenkis. 1971. On the Uniform Convergence of Relative Frequencies of Events to Their Probabilities. *Theory of Probability and its Applications* 16, 2 (1971), 264.
- [50] Vladimir N Vapnik. 1995. *The Nature of Statistical Learning Theory*. Springer.
- [51] Christopher S Wallace and David M Boulton. 1968. An information measure for classification. *Comput. J.* 11, 2 (1968), 185–194.
- [52] Jiaming Zeng, Berk Ustun, and Cynthia Rudin. 2017. Interpretable classification models for recidivism prediction. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 180, 3 (2017), 689–722.
- [53] Ding-Xuan Zhou. 2002. The covering number in learning theory. *Journal of Complexity* 18, 3 (2002), 739–767.

A DIFFERENCE BETWEEN FLAT MINIMA AND RASHOMON SET

Figure 4 illustrates differences between the volume ϵ -flatness and the Rashomon set.

B CONNECTION TO SIMPLICITY MEASURES

We will use demonstrations to show the differences between the Rashomon ratio and other complexity measures mentioned in Section 4. However, first we discuss how to analytically compute the Rashomon ratio for ridge regression since the demonstration of the difference between algorithmic stability and Rashomon ratio relies on it.

B.1 Analytical Calculation of Rashomon Ratio for Ridge Regression

A special case of when the Rashomon ratio can be computed in closed form in a parameter space is ridge regression. For a space of linear models $\mathcal{F}_\Omega = \{\omega^T x, \omega \in \mathbb{R}^p\}$, ridge regression chooses a parameter vector by minimizing the penalized sum of squared errors for a training data set $S = [X, Y]$:

$$\min_{\omega} \hat{L}(\omega) = \min_{\omega} (X\omega - Y)^T(X\omega - Y) + C\omega^T\omega, \quad (3)$$

where the optimal solution of the ridge regression estimator is $\hat{\omega} = (X^T X + CI_p)^{-1} X^T Y$.

Geometrically, the optimal solution to ridge regression will be a parameter vector that corresponds to the intersection of ellipsoidal isosurfaces of the sum of squares term and a hypersphere centered at the origin, with the regularization parameter C determining the trade off between the loss and the radius of the sphere. More generally, isosurfaces of the ridge regression loss function are ellipsoids, and the volume of such an ellipsoid corresponds to the volume of the Rashomon set. For a hypothesis space with uniform prior and volume function \mathcal{V} , the Rashomon ratio is $\frac{\mathcal{V}(\hat{R}_{set}(\mathcal{F}_\Omega, \theta))}{\mathcal{V}(\mathcal{F}_\Omega)}$. Using the geometric intuition above, we compute the ratio of the Rashomon set in the parameter space by the following theorem:

THEOREM 10 (RASHOMON RATIO FOR RIDGE REGRESSION). *For a parametric hypothesis space of linear models $\mathcal{F}_\Omega = \{f_\omega(x) = \omega^T x, \omega \in \mathbb{R}^p\}$ with uniform prior and a data set $S = X \times Y$, the Rashomon set $\hat{R}_{set}(\mathcal{F}_\Omega, \theta)$ of ridge regression is an ellipsoid, containing vectors ω such that:*

$$(\omega - \hat{\omega})^T \frac{X^T X + CI_p}{\theta} (\omega - \hat{\omega}) \leq 1,$$

and the Rashomon ratio can be computed as:

$$\hat{R}_{ratio}(\mathcal{F}_\Omega, \theta) = \frac{J(\theta, p)}{\mathcal{V}(\mathcal{F})} \prod_{i=1}^p \frac{1}{\sqrt{\sigma_i^2 + C}}, \quad (4)$$

where σ_i are singular values of matrix X , $J(\theta, p) = \frac{\pi^{p/2} \theta^{p/2}}{\Gamma(p/2+1)}$ and $\Gamma(\cdot)$ is the gamma function.

PROOF. Consider all models $f_\omega \in \mathcal{F}_\Omega$ from the Rashomon set $\hat{R}_{set}(\mathcal{F}_\Omega, \theta)$. Then by Definition 1 we get:

$$\hat{L}(X, Y, \omega) \leq \hat{L}(X, Y, \hat{\omega}) + \theta. \quad (5)$$

Using $X^T Y = (X^T X + CI_p)\hat{\omega}$ from the optimal solution of the ridge regression estimator $\hat{\omega} = (X^T X + CI_p)^{-1} X^T Y$, and expanding the difference between empirical risks we have:

$$\begin{aligned} \theta &\geq \hat{L}(X, Y, \omega) - \hat{L}(X, Y, \hat{\omega}) \\ &= (X\omega - Y)^T(X\omega - Y) + C\omega^T\omega - (X\hat{\omega} - Y)^T(X\hat{\omega} - Y) - C\hat{\omega}^T\hat{\omega} \\ &= \omega^T X^T X \omega - 2\omega^T X^T Y + C\omega^T\omega - \hat{\omega}^T X^T X \hat{\omega} + 2\hat{\omega}^T X^T Y - C\hat{\omega}^T\hat{\omega} \\ &= \omega^T X^T X \omega - 2\omega^T (X^T X + CI_p)\hat{\omega} + C\omega^T\omega - \hat{\omega}^T X^T X \hat{\omega} \\ &\quad + 2\hat{\omega}^T (X^T X + CI_p)\hat{\omega} - C\hat{\omega}^T\hat{\omega} \\ &= \omega^T X^T X \omega + C\omega^T\omega - 2\omega^T (X^T X + CI_p)\hat{\omega} + \hat{\omega}^T X^T X \hat{\omega} + C\hat{\omega}^T\hat{\omega} \\ &= \omega^T (X^T X + CI_p)\omega - 2\omega^T (X^T X + CI_p)\hat{\omega} + \hat{\omega}^T (X^T X + CI_p)\hat{\omega} \\ &= (\omega - \hat{\omega})^T (X^T X + CI_p)(\omega - \hat{\omega}). \end{aligned}$$

Therefore the Rashomon set is an ellipsoid centered at $\hat{\omega}$:

$$(\omega - \hat{\omega})^T \frac{X^T X + CI_p}{\theta} (\omega - \hat{\omega}) \leq 1.$$

By the formula of the volume of a p -dimensional ellipsoid, the volume of the Rashomon set can be computed as:

$$\mathcal{V}(\hat{R}_{set}(\mathcal{F}_\Omega, \theta)) = \frac{\pi^{p/2} \theta^{p/2}}{\Gamma(p/2+1)} \prod_{i=1}^p \frac{1}{\sqrt{\sigma_i^2 + C}},$$

where σ_i are singular values of X . Therefore, the Rashomon ratio is $\hat{R}_{ratio}(\mathcal{F}_\Omega, \theta) = \frac{\mathcal{V}(\hat{R}_{set}(\mathcal{F}_\Omega, \theta))}{\mathcal{V}(\mathcal{F})} = \frac{J(\theta, p)}{\mathcal{V}(\mathcal{F})} \prod_{i=1}^p \frac{1}{\sqrt{\sigma_i^2 + C}}$, where

$$J(\theta, p) = \frac{\pi^{p/2} \theta^{p/2}}{\Gamma(p/2+1)}. \quad \square$$

Interestingly, from Theorem 10, it follows that for ridge regression, *the Rashomon ratio depends on the feature space only and does not depend on the regression targets Y* . Indeed, assume that every parameter vector ω such that $f_\omega \in \hat{R}_{set}(\mathcal{F}_\Omega, \theta)$ can be represented as $\omega = \hat{\omega} + \delta$. By a simple transformation, we have that $\hat{L}(f_\omega) - \hat{L}(f_{\hat{\omega}}) = \delta^T X^T X \delta$, meaning that if we take a step in parameter space, the empirical risk difference will depend only on the feature space and the step itself, and not on the targets of the problem. This observation can help us choose the parameter θ as $\theta = \delta^T X^T X \delta$ if we want to ensure some dependence between the optimal model $\hat{\omega}$ and a model of interest ω . Then, by choosing the direction as $\delta = \omega - \hat{\omega}$, we can compute the Rashomon parameter θ .

For other algorithms, the Rashomon ratio generally depends on the targets; in that sense, ridge regression is unusual.

B.2 Algorithmic Stability

The main motivation for algorithmic stability theory is to ensure robustness of a learning algorithm. Following Bousquet and Elisseeff [5], we define the hypothesis stability of a learning algorithm as follows.

Definition 11 (Hypothesis stability). A learning algorithm \mathcal{A} has β hypothesis stability with respect to the loss l if for all $i \in \{1, \dots, n\}$,

$$\mathbb{E}_{S, z} [|l(f_S, z) - l(f_{S^{(i)}}, z)|] \leq \beta,$$

where $\beta \in \mathbb{R}_+$, hypothesis f_S is learned by an algorithm \mathcal{A} on a data set S , loss $l(f_S, z) = \phi(f_S(x), y)$ for $z = (x, y)$, data set $S =$




Figure 4: Difference between the volume ϵ -flatness as defined in Dinh et al. [16] and the Rashomon set. The red line represents the volume ϵ -flatness in (a), and the Rashomon set in (b). The volume of the Rashomon set is the sum of lengths of red lines in (b). The height of the shaded area represents (a) the parameter ϵ or the 2σ -sharpness, and (b) the Rashomon parameter θ . The volume ϵ -flatness is defined by a connected component in a parameter space for a given local minimum, while the Rashomon set is defined with respect to an empirical risk minimizer over the full hypothesis space \mathcal{F} and may contain models from multiple local minima. Rashomon sets are also defined for discrete spaces.

$\{z_1, \dots, z_n\}$, and $S^{\setminus i}$ is modified from the training data by removing the i^{th} element of the data set: $S^{\setminus i} = \{z_1, \dots, z_{i-1}, z_{i+1}, \dots, z_n\}$.

The Rashomon ratio is fundamentally different from hypothesis stability, in case of linear least squares regression (which is discussed in Section B.1). This is formalized in Theorem 2.

THEOREM 2 (RASHOMON RATIO IS NOT ALGORITHMIC STABILITY). Consider a distribution P_X over a discrete domain $\mathcal{X} = \{x_1, \dots, x_N\}$ and a learning algorithm A that minimizes the sum of squares loss: $\|X\omega - Y\|_2^2$, for a linear hypothesis space \mathcal{F}_Ω . For any $\lambda > 0$, there exist joint distributions P_{X, Y_1} and P_{X, Y_2} where for X drawn i.i.d. from P_X , Y_1 drawn from $P_{Y_1|X}$ over $\mathcal{Y} | \mathcal{X}$, and Y_2 drawn from $P_{Y_2|X}$ over $\mathcal{Y} | \mathcal{X}$, the expected Rashomon ratios are the same:

$$\mathbb{E}_{P_{X, Y_1}}[\hat{R}_{\text{ratio}_{S_1}}(\mathcal{F}_\Omega, \theta)] = \mathbb{E}_{P_{X, Y_2}}[\hat{R}_{\text{ratio}_{S_2}}(\mathcal{F}_\Omega, \theta)],$$

yet hypothesis stability constants are different by our arbitrarily chosen value of λ : $\hat{\beta}_2 - \hat{\beta}_1 \geq \lambda$, where S_1 and S_2 denote data sets $S_1 = [X, Y_1]$ and $S_2 = [X, Y_2]$, $\hat{\beta}_1$ is the hypothesis stability coefficient of algorithm A for distribution P_{X, Y_1} and $\hat{\beta}_2$ is the hypothesis stability coefficient for distribution P_{X, Y_2} .

PROOF. Let us create our distribution. Consider the least squares regression $\min_{\omega} \sum_{i=1}^n l(\omega, z_i)^2$, where $\omega \in \mathbb{R}^p$, and loss $l(\omega, z) = \phi(\omega^T \mathbf{x}, y)$ for $z = (\mathbf{x}, y)$. For the marginal distribution P_X and $X = [x_1, \dots, x_n]$ drawn i.i.d. from P_X , we design distributions $P_{Y_1|X}$ and $P_{Y_2|X}$ as:

$$\begin{aligned} P_{Y_1|X}(y = 0 | \mathbf{x}) &= 1 \quad \forall \mathbf{x} \in \mathcal{X}, \\ P_{Y_2|X}(y = 0 | \mathbf{x} \neq \mathbf{x}_0) &= 1, \quad P_{Y_2|X}(y = 0 | \mathbf{x} = \mathbf{x}_0) = 0.5, \\ P_{Y_2|X}(y = H | \mathbf{x} = \mathbf{x}_0) &= 0.5, \end{aligned}$$

where $\mathbf{x}_0 \in \{x_1, \dots, x_N\}$ is some fixed point with a positive probability $P_X(\mathbf{x}_0)$ and we define $H \in \mathbb{R}$ later. That is, the two conditional distributions have $y = 0$ except when $\mathbf{x} = \mathbf{x}_0$ for Y_2 , when it is H with probability $1/2$.

As a first part of the proof, we show that the algorithmic stability constants are different. According to the definition of algorithmic stability, for P_{X, Y_1} we have:

$$\mathbb{E}_{S_1, z}[|l(f_{S_1}, z) - l(f_{S_1^{\setminus i}}, z)|] = 0 = \tilde{\beta}_1,$$

and for distribution P_{X, Y_2} :

$$\begin{aligned} \mathbb{E}_{S_2, z}[|l(f_{S_2}, z) - l(f_{S_2^{\setminus i}}, z)|] &= \sum_{S_2, z \sim P_{X, Y_2}} P_{X, Y_2}(S_2) P_{X, Y_2}(z) \\ &\times |l(f_{S_2}, z) - l(f_{S_2^{\setminus i}}, z)| \\ &\geq P_{X, Y_2}(S_2^s) P_{X, Y_2}(z^s) |l(f_{S_2^s}, z^s) - l(f_{S_2^{s, \setminus i}}, z^s)|, \end{aligned}$$

where S_2^s, z^s is a special draw such that $z^s = (\mathbf{x}_0, H)$, and where S_2^s includes one point at (\mathbf{x}_0, H) , one point at $(\mathbf{x}_0, 0)$, and the rest at other values $(\mathbf{x}, 0)$. Since the domain \mathcal{X} is discrete, the probabilities of a special draw are:

$$P_{X, Y_2}(z^s) = \frac{1}{2} \text{Bin}(1, n, P_X(\mathbf{x}_0)),$$

$$P_{X, Y_2}(S_2^s) = \frac{1}{4} \text{Bin}(1, n, P_X(\mathbf{x}_0))^2 \text{Bin}(n-2, n, 1 - P_X(\mathbf{x}_0)),$$

where $\text{Bin}(k, n, p_k) = \binom{n}{k} p_k^k (1 - p_k)^{n-k}$ is a binomial coefficient, namely the probability of getting exactly k successes from n trials, where each trial has a probability of success p_k . Denote $P_{(S_2^s, z^s)}$ as the probability of getting a special draw, then $P_{(S_2^s, z^s)} = P_{X, Y_2}(S_2^s) P_{X, Y_2}(z^s)$.

If S_2^s contains only two points $z_1 = (\mathbf{x}_0, H)$ and $z_2 = (\mathbf{x}_0, 0)$, the loss difference $|l(f_{S_2^s}, z^s) - l(f_{S_2^{s, \setminus i}}, z^s)|$ evaluated at z^s for all i will be at least $\frac{H^2}{4}$. To see this, note that the optimal function's value at \mathbf{x}_0 is: $f_{S_2^s}(\mathbf{x}_0) = \frac{H}{2}$, the optimal function's value at \mathbf{x}_0 after we remove the first point is $f_{S_2^{s, \setminus 1}}(\mathbf{x}_0) = 0$, and the optimal function's value at \mathbf{x}_0 after removing the second point is $f_{S_2^{s, \setminus 2}}(\mathbf{x}_0) = H$. Therefore, $l(f_{S_2^s}, z^s) = \frac{H^2}{4}$, $l(f_{S_2^{s, \setminus 1}}, z^s) = H^2$, $l(f_{S_2^{s, \setminus 2}}, z^s) = 0$. And we get that $|l(f_{S_2^s}, z^s) - l(f_{S_2^{s, \setminus 1}}, z^s)| = \frac{3H^2}{4}$, $|l(f_{S_2^s}, z^s) - l(f_{S_2^{s, \setminus 2}}, z^s)| = \frac{H^2}{4}$. As we add the rest of the points $(x_i, 0)$ to the data set S_2^s , the loss difference (from changing $f_{S_2^s}(z^s)$ to $f_{S_2^{s, \setminus i}}(z^s)$) in the special draw case will

only increase. Therefore for all i :

$$|l(f_{S_2^s, z^s}) - l(f_{S_2^s \setminus i, z^s})| \geq \frac{H^2}{4}.$$

If we choose H such that $H > 2\sqrt{\lambda} \left(P_{(S_2^s, z^s)}\right)^{-1/2}$, then from the definition of algorithmic stability we have:

$$\tilde{\beta}_2 \geq \mathbb{E}_{S_2, z} \left[\left| l(f_{S_2, z}) - l(f_{S_2 \setminus i, z}) \right| \right] \geq P_{(S_2^s, z^s)} \frac{H^2}{4} > \lambda.$$

Therefore for any given λ we get that $|\tilde{\beta}_1 - \tilde{\beta}_2| > \lambda$. This proves that the hypothesis stability constants are different and completes the first part of the proof.

We now need to prove that the expected Rashomon ratios are the same, which will constitute the second part of the proof. The Rashomon ratio for the hypothesis space \mathcal{F}_Ω of linear models does not depend on targets and can be calculated as in (4) for both S_1 and S_2 . Therefore the expected Rashomon ratios are the same:

$$\mathbb{E}_{P_{X, Y_1}} [\hat{R}_{ratio_{S_1}}(\mathcal{F}_\Omega, \theta)] = \mathbb{E}_{P_{X, Y_2}} [\hat{R}_{ratio_{S_2}}(\mathcal{F}_\Omega, \theta)].$$

Thus, both halves of our proof are complete. \square

B.3 Geometric Margin

For the parametric hypothesis space of linear models $\mathcal{F}_\Omega = \{f : f(x) = \omega^T x, \omega \in \mathbb{R}^p\}$ and binary classification, denote d_+ and d_- as the shortest distances from a decision boundary to the closest points with targets $y = 1$ and $y = -1$ respectively. Then the margin d is a sum of these distances $d = d_+ + d_-$ [7]. Moreover, for the model $f_{\hat{\omega}}$ that maximizes the margin, the margin width is $\frac{2}{\|\hat{\omega}\|_2}$.

Intuitively both the Rashomon ratio and the width of the geometric margin are data-dependent and show how expressive the hypothesis space is with respect to a given data set. However, the margin depends on support vectors while the Rashomon set depends on the full data set. Theorem 3 summarizes this idea.

THEOREM 3 (RASHOMON RATIO IS NOT THE GEOMETRIC MARGIN). *For any fixed $0 < \lambda < 1$, there exists a fixed hypothesis space \mathcal{F}_Ω , a Rashomon parameter θ , and there exist two data sets S_1 and S_2 with the same empirical risk minimizer $\hat{f} \in \mathcal{F}_\Omega$ such that the geometric margin d is the same for both data sets, yet the Rashomon ratios are different:*

$$|\hat{R}_{ratio_{S_1}}(\mathcal{F}_\Omega, \theta) - \hat{R}_{ratio_{S_2}}(\mathcal{F}_\Omega, \theta)| > \lambda.$$

PROOF. Consider two-dimensional separable data, $X \in [0, 1]^2$, and a parametrized hypothesis space of origin-centered linear models: $\mathcal{F} = \{\omega^T x, \omega = (k, -1), x \in \mathbb{R}^2, k \in \mathbb{R}\}$. Consider also 0-1 loss $\phi_\omega(x, y) = \mathbb{1}_{[y=sign(\omega^T x)]}$ and an empirical risk minimizer $\hat{f} = f_{\hat{\omega}}$ that maximizes the geometric margin. Since the data are populated in a $[0, 1]^2$ hypercube, as a hypothesis space we will consider all models that intersect the unit-hypercube.

For some positive constant $a \in (0, 1)$ that we choose later, consider the following regions of the feature space:

$$\begin{aligned} A &= \{x^1 \in [0, 1-a], x^2 > x^1 + (1-2a)\}, \\ B &= \{x^1 \in (a, 1], x^2 < x^1 - (1-2a)\}, \\ C &= \{x^1 \in [0, a], x^2 \in (1-a, 1]\}, \\ D &= \{x^1 \in (1-a, 1], x^2 \in [0, a)\}. \end{aligned}$$

Construct data set S_1 , such that $S_1 = \{(x_A, 1) \cup (x_B, -1) \cup (x_{S_1}^{s_1}, 1) \cup (x_{S_1}^{s_2}, -1)\}$, where $x_A \in A$ is any sample from the region A , $x_B \in B$ is any sample from the region B , $x_{S_1}^{s_1}$ and $x_{S_1}^{s_2}$ are special points for the data set S_1 such that $x_{S_1}^{s_1} = [1-2a, 1]$ and $x_{S_1}^{s_2} = [1, 1-2a]$. Please see Figure 5a for details.

Construct data set S_2 , such that $S_2 = \{(x_C, 1) \cup (x_D, -1) \cup (x_{S_2}^{s_1}, 1) \cup (x_{S_2}^{s_2}, -1)\}$, where $x_C \in C$ is any sample from the region C , $x_D \in D$ is any sample from the region D , $x_{S_2}^{s_1}$ and $x_{S_2}^{s_2}$ are special points for the data set S_2 such that $x_{S_2}^{s_1} = [a, 1-a]$ and $x_{S_2}^{s_2} = [1-a, a]$. Please see Figure 5b for details.

Note that the data sets we considered have the same width for the geometrical margin $d = \sqrt{2}(2a-1)$ (see Figures 5a, 5b). Now, we are left to show that the Rashomon ratios are different.

For the hypothesis space of origin-centered lines we have a unique parameterization and a one-to-one correspondence between an actual model and its parameterization. Therefore, if the Rashomon set is a single connected component, an angle α between the two most distant models in the Rashomon set gives us some information about the size of the Rashomon set. In particular, we can compute the Rashomon ratio as a ratio of the angle α that represents the Rashomon set and the angle β that corresponds to the hypothesis space as shown on Figure 5c. Since the hypothesis space is defined on the unit-hypercube, $\beta = \pi/2$ and for the Rashomon parameter $\theta = 0$ the Rashomon ratio is:

$$\hat{R}_{ratio}(\mathcal{F}, 0) = \frac{\alpha}{\beta} = \frac{2 \max_{f \in \hat{R}_{set}(\mathcal{F}_\Omega, 0)} |\arctan(f_{\hat{\omega}}) - \arctan(f_{\omega})|}{\pi/2}.$$

For data sets S_1 and S_2 Figures 5d and 5e show the Rashomon set and angles α_1 and α_2 that represent the volume of the Rashomon set. Given the special points in the data sets we can compute α_1 and α_2 exactly: $\alpha_1 = 2(\arctan(1) - \arctan(1-2a)) = \frac{\pi}{2} - 2\arctan(1-2a)$ and $\alpha_2 = 2(\arctan(1) - \arctan(\frac{a}{1-a})) = \frac{\pi}{2} - 2\arctan(\frac{a}{1-a})$. Then the Rashomon ratios difference is:

$$\begin{aligned} |\hat{R}_{ratio_{S_1}}(\mathcal{F}, 0) - \hat{R}_{ratio_{S_2}}(\mathcal{F}, 0)| &= \left| \frac{\alpha_1 - \alpha_2}{\pi/2} \right| \\ &= \left| \frac{4}{\pi} \left(\arctan(1-2a) - \arctan\left(\frac{a}{1-a}\right) \right) \right| \\ &= \left| \frac{4}{\pi} \arctan\left(1 - \frac{4a-2}{2a^2-1}\right) \right|. \end{aligned}$$

Now if we choose $a \in (0, 1)$ and such that $\left| \frac{4}{\pi} \arctan\left(1 - \frac{4a-2}{2a^2-1}\right) \right| > \lambda$, then the Rashomon ratio difference $|\hat{R}_{ratio_{S_1}}(\mathcal{F}, 0) - \hat{R}_{ratio_{S_2}}(\mathcal{F}, 0)|$ is at least λ . \square

B.4 Empirical Local Rademacher Complexity

The empirical Rademacher complexity is another complexity measure of the hypothesis space. Following Bartlett et al. [2], for binary classification we define it as follows.

Definition 12 (Empirical Rademacher complexity). Given a data set S , and a hypothesis space \mathcal{F} of real-valued functions, the empirical

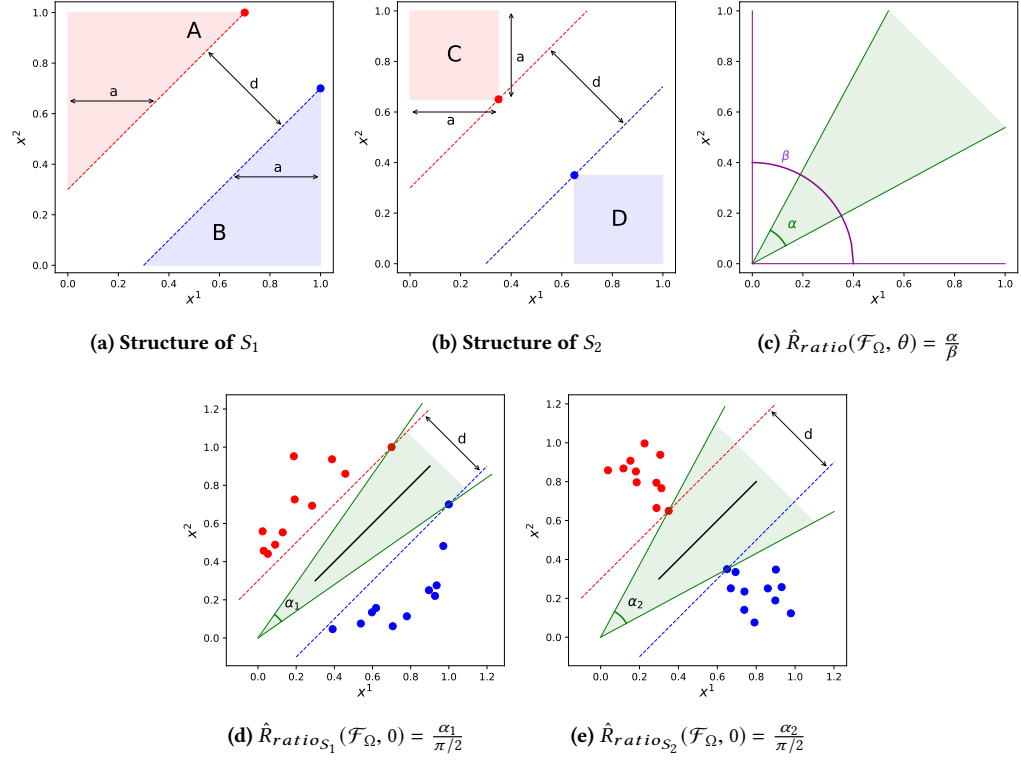


Figure 5: An illustration of different Rashomon ratios with identical geometric margins. (a) and (b) show the data sets S_1 and S_2 with identical margin d . The black line in (d) and (e) shows the optimal model, the shaded region in (c), (d), and (e) indicates the Rashomon set $\hat{R}_{set}(\mathcal{F}_\Omega, 0)$ with its boundaries represented by green lines. The hypothesis space consists of all origin-centered linear models that intersect the zero-one hypercube, where data reside. (c) shows that the Rashomon ratio can be computed as a ratio of angles α (represents the Rashomon set) and β (represents the hypothesis space). (d) and (e) illustrate that data sets S_1 and S_2 are represented by different angles α_1 and α_2 and therefore have different Rashomon ratios. Figure is best seen in color.

Rademacher complexity of \mathcal{F} is defined as:

$$\hat{R}_n^S(\mathcal{F}) = \frac{1}{n} \mathbb{E}_\sigma \left[\sup_{f \in \mathcal{F}} \sum_{i=1}^n \sigma_i f(z_i) \right],$$

where $\sigma_1, \sigma_2, \dots, \sigma_n$ are independent random variables drawn from the Rademacher distribution i.e. $P(\sigma_i = +1) = P(\sigma_i = -1) = 1/2$ for $i = 1, 2, \dots, n$.

Since we are interested only in models that are inside the Rashomon set, we will consider local empirical Rademacher complexity [2], which is defined using the Rashomon set $\hat{R}_{set}(\mathcal{F}, \theta)$. In the following theorem, we provide a simple example to show the discrepancy between the two measures.

THEOREM 4 (RASHOMON RATIO IS NOT THE LOCAL RADEMACHER COMPLEXITY). For $0 < \lambda < 1$, there exist two data sets S_1 and S_2 , a hypothesis space \mathcal{F}_Ω , and a Rashomon parameter θ such that the local Rademacher complexities defined on the Rashomon sets for S_1 and S_2 are the same: $\hat{R}_n^{S_1}(\hat{R}_{set}(\mathcal{F}_\Omega, \theta)) = \hat{R}_n^{S_2}(\hat{R}_{set}(\mathcal{F}_\Omega, \theta))$, yet the Rashomon ratios are different: $|\hat{R}_{ratio_{S_1}}(\mathcal{F}_\Omega, \theta) - \hat{R}_{ratio_{S_2}}(\mathcal{F}_\Omega, \theta)| > \lambda$.

PROOF. Consider two-dimensional separable symmetric data, $\mathcal{X} \in [0, 1]^2$, $\mathcal{Y} = \{0, 1\}$, 0-1 loss $\phi_f(x, y) = \mathbb{1}_{[y \neq \text{sign}f(x)]}$ with empirical risk minimizer \hat{f} , and a hypothesis space \mathcal{F}_Ω of decision stumps based on the first feature, where for $f \in \mathcal{F}_\Omega$: $f = 1$ if $x^1 > \omega$, $\omega \in \mathbb{R}$, $f = 0$ otherwise. We have a one-to-one correspondence between a function and its threshold parameter ω . Therefore, if the Rashomon set is a single connected component, we can compute the volume of the Rashomon set in parameter space by computing the difference between the largest and smallest threshold values of models within the Rashomon set, as illustrated in Figure 6a. For $\theta = 0$, the difference between the largest and the smallest threshold values will be equivalent to the minimal distance between points of opposite classes projected onto the first feature $d = \min_{x_i, x_j: y_i \neq y_j} |PR_1(x_i) - PR_1(x_j)|$, where PR_1 is the projection of point x onto first feature.

For the hypothesis space, we consider all decision stumps in the first dimension that are in the segment $[0, 1]$, where data are populated. The difference in thresholds for the hypothesis space is $\beta = 1$ and therefore $\mathcal{V}(\mathcal{F}_\Omega) = 1$. For $\theta = 0$, the volume of the Rashomon set will be equivalent to d —the projected minimal distance between points of opposite classes, and have that $\mathcal{V}(\hat{R}_{set}(\mathcal{F}_\Omega, 0)) = d$ and $\hat{R}_{ratio}(\hat{R}_{set}(\mathcal{F}_\Omega, 0)) = \frac{d}{1} = d$. Now consider any two separable

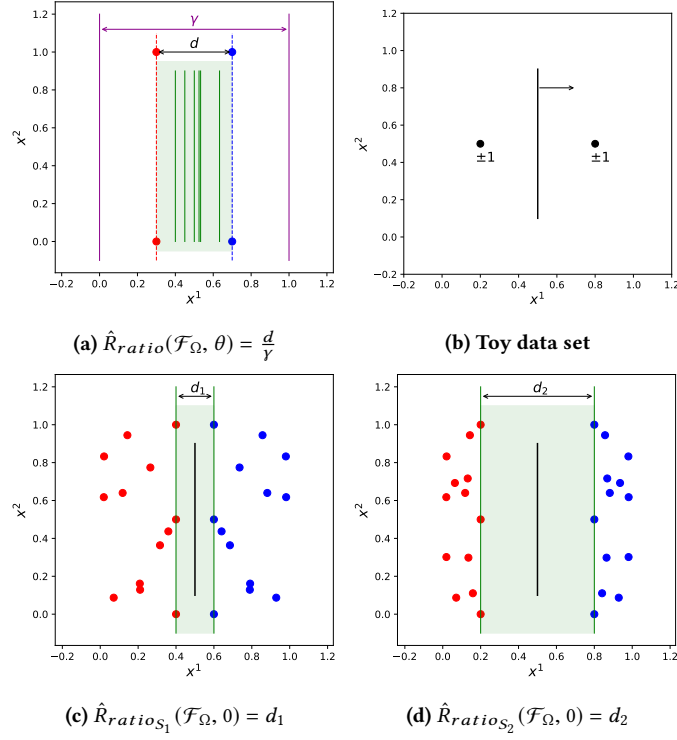


Figure 6: An illustration of different Rashomon ratios with equivalent empirical local Rademacher complexities. Black line shows the optimal model, shaded region indicates the Rashomon set $\hat{R}_{set}(\mathcal{F}_\Omega, 0)$ with its models represented by green lines, the magenta color indicates boundaries of the hypothesis space. (a) The projected minimal distance d is equivalent to the volume of the Rashomon set. (b) A toy data set that illustrates that the empirical local Rademacher complexity is zero for models in the Rashomon set. (c) data set S_1 , and (d) data set S_2 illustrate symmetric separable data sets with different Rashomon ratios. Best seen in color.

symmetric data sets S_1, S_2 with different projected minimal distances d_1 and d_2 , such that $|d_1 - d_2| > \lambda$. (Please see Figure 6c and 6d for details of the data sets S_1 and S_2 .) Consequently we get that:

$$\left| \hat{R}_{ratio_{S_1}}(\mathcal{F}_\Omega, 0) - \hat{R}_{ratio_{S_2}}(\mathcal{F}_\Omega, 0) \right| = |d_1 - d_2| > \lambda.$$

For a separable symmetric data S and 0-1 loss function, the Rashomon set $\hat{R}_{set}(\mathcal{F}_\Omega, 0)$ contains all models that separate data in the same way. Therefore the Rademacher complexity of the Rashomon set is $\hat{R}_n^S(\hat{R}_{set}(\mathcal{F}_\Omega))$ is:

$$\begin{aligned} \hat{R}_n^S(\hat{R}_{set}(\mathcal{F}_\Omega, 0)) &= \frac{1}{n} \mathbb{E}_\sigma \left[\sup_{f \in \hat{R}_{set}(\mathcal{F}_\Omega, 0)} \sum_{i=1}^n \sigma_i f(x_i) \right] \\ &= \frac{1}{n} \mathbb{E}_\sigma \left[\sum_{i=1}^n \sigma_i \hat{f}(x_i) \right] = 0, \end{aligned}$$

where in the penultimate equality we have used the fact that, in the case of separable data and $\theta = 0$, all models in the Rashomon set will perform identically on any permutation of the labels.

Equality of the empirical Rademacher complexity of the optimal model to zero follows from the symmetric data considered and symmetrical patterns of all possible target assignments. For example,

$$\text{for the toy data set in Figure 6b: } \hat{R}_n^S(\hat{R}_{set}(\mathcal{F}_\Omega, 0)) = \frac{1}{2} \frac{1}{4} \left[\left(\hat{f}(x_1) + \hat{f}(x_2) \right) + \left(\hat{f}(x_1) - \hat{f}(x_2) \right) + \left(-\hat{f}(x_1) + \hat{f}(x_2) \right) + \left(-\hat{f}(x_1) - \hat{f}(x_2) \right) \right] = 0.$$

Since both S_1 and S_2 are separable and symmetric we get that:

$$\hat{R}_n^{S_1}(\hat{R}_{set}(\mathcal{F}_\Omega, 0)) = 0 = \hat{R}_n^{S_2}(\hat{R}_{set}(\mathcal{F}_\Omega, 0)).$$

□

C PROOFS FOR GENERALIZATION RESULTS

C.1 Proof of Theorem 5

We recall and provide the proof of Theorem 5 after Proposition 13 that we use to prove Theorem 5.

Given a parameter $\eta \geq 0$, we call the Rashomon set with restricted empirical risk an *anchored Rashomon set*:

$$\hat{R}_{set}^{anc}(\mathcal{F}, \eta) := \{f \in \mathcal{F} : \hat{L}(f) \leq \eta\}.$$

We define also the *true anchored Rashomon set* based on the true risk as follows:

$$R_{set}^{anc}(\mathcal{F}, \eta) := \{f \in \mathcal{F} : L(f) \leq \eta\}.$$

PROPOSITION 13 (TRUE ANCHORED RASHOMON SET IS CLOSE TO EMPIRICAL). *For a loss l bounded by b and for any $\epsilon > 0$, and for a*

fixed f , if $f \in R_{set}^{anc}(\mathcal{F}, \eta)$ then with probability at least $1 - e^{-2n(\epsilon/b)^2}$ with respect to the random draw of training data,

$$f \in \hat{R}_{set}^{anc}(\mathcal{F}, \eta + \epsilon).$$

PROOF. For a fixed $f \in R_{set}^{anc}(\mathcal{F}, \eta)$ by Hoeffding's inequality:

$$\begin{aligned} P[\hat{L}(f) - L(f) > \epsilon] &= P\left[\frac{1}{n} \sum_{i=1}^n l(f, z_i) - \mathbb{E}[l(f, z)] > \epsilon\right] \\ &\leq e^{-2n(\epsilon/b)^2}. \end{aligned}$$

Therefore, with probability at least $1 - e^{-2n(\epsilon/b)^2}$ with respect to the random draw of data, $\hat{L}(f) - L(f) \leq \epsilon$.

Since $f \in R_{set}^{anc}(\mathcal{F}, \eta)$, then by definition of the Rashomon set, $L(f) \leq \eta$. Combining this with Hoeffding's inequality, we get that with probability at least $1 - e^{-2n(\epsilon/b)^2}$:

$$\hat{L}(f) \leq L(f) + \epsilon \leq \eta + \epsilon,$$

therefore $f \in \hat{R}_{set}^{anc}(\mathcal{F}, \eta + \epsilon)$. \square

Proposition 13 is based on the same intuition as Lemma 23 in the work of Fisher et al. [19], which is used to bound the probability with which a given model is not in the empirical Rashomon set; this is used in a proof of a bound for model class reliance. We use the proposition to indicate the probability with which the empirical anchored Rashomon set is as close as possible to the true anchored Rashomon set for a given model.

THEOREM 5 (THE ADVANTAGE OF A TRUE RASHOMON SET). Consider finite hypothesis spaces \mathcal{F}_1 and \mathcal{F}_2 , such that $\mathcal{F}_1 \subset \mathcal{F}_2$. Let the loss l be bounded by b , $l(f_2, z) \in [0, b] \quad \forall f_2 \in \mathcal{F}_2, \forall z \in \mathcal{Z}$. Define an optimal function $f_2^* \in \arg\min_{f_2 \in \mathcal{F}_2} L(f_2)$. Assume that the true Rashomon set includes a function from \mathcal{F}_1 , so there exists a model $\tilde{f}_1 \in \mathcal{F}_1$ such that $\tilde{f}_1 \in R_{set}(\mathcal{F}_2, \gamma)$. (Note that we do not know \tilde{f}_1 .) In that case, for any $\epsilon > 0$ with probability at least $1 - \epsilon$ with respect to the random draw of data:

$$L(f_2^*) - b\sqrt{\frac{\log |\mathcal{F}_1| + \log 2/\epsilon}{2n}} \leq \hat{L}(\tilde{f}_1) \leq L(f_2^*) + \gamma + b\sqrt{\frac{\log 1/\epsilon}{2n}},$$

where $\hat{f}_1 \in \arg\min_{f_1 \in \mathcal{F}_1} \hat{L}(f_1)$. (Unlike \tilde{f}_1 , we do know \hat{f}_1 because we can calculate it.)

PROOF. **Lower bound.** We apply the union bound and Hoeffding's inequality. The result is that with probability at least $1 - \epsilon$ for every $f_1 \in \mathcal{F}_1$ we have, for finite hypothesis space \mathcal{F}_1 :

$$L(f_1) \leq \hat{L}(f_1) + b\sqrt{\frac{\log |\mathcal{F}_1| + \log 2/\epsilon}{2n}}. \quad (6)$$

Combining this Occam's razor bound with the definition of $f_2^* \in \arg\min_{f \in \mathcal{F}_2} L(f)$ we get that, under the same conditions:

$$L(f_2^*) \leq \hat{L}(\hat{f}_1) \leq \hat{L}(\tilde{f}_1) + b\sqrt{\frac{\log |\mathcal{F}_1| + \log 2/\epsilon}{2n}}.$$

Upper bound. By the assumption of the theorem, we have that $L(\tilde{f}_1) \leq L(f_2^*) + \gamma$. Also, by the definition of an optimal model f_1^* , $L(f_1^*) \leq L(\tilde{f}_1)$. Combining these, we get that $L(f_1^*) \leq L(\tilde{f}_1) \leq L(f_2^*) + \gamma$. Thus f_1^* is in the true Rashomon set of \mathcal{F}_2 with parameter γ . Alternatively, f_1^* is in the true anchored Rashomon set of \mathcal{F}_2 with parameter $\eta = L(f_2^*) + \gamma$, $f_1^* \in R_{set}^{anc}(\mathcal{F}_2, \eta)$. Following

Proposition 13, we have that for any $\epsilon_1 > 0$ with probability at least $1 - e^{-2n(\epsilon_1/b)^2}$ with respect to the random draw of data, f_1^* is in the slightly larger anchored Rashomon set $\hat{R}_{set}^{anc}(\mathcal{F}_2, \eta + \epsilon_1)$, and therefore, with high probability, $\hat{L}(f_1^*) \leq \eta + \epsilon_1$. Or alternatively, by setting $\epsilon = e^{-2n(\epsilon_1/b)^2}$ we get that for any $\epsilon > 0$ with probability at least $1 - \epsilon$, we have $\hat{L}(f_1^*) \leq \eta + b\sqrt{\frac{\log 1/\epsilon}{2n}}$. Further, by definition of the empirical risk minimizer and given that $\eta = L(f_2^*) + \gamma$ we get:

$$\hat{L}(\hat{f}_1) \leq \hat{L}(f_1^*) \leq L(f_2^*) + \gamma + b\sqrt{\frac{\log 1/\epsilon}{2n}}.$$

Combining the previous two equations yields the statement of the theorem. \square

C.2 Proof of Theorem 6 via Lemma 14

Theorem 6 follows directly from Lemma 14 below and Theorem 5, which guarantees that with high probability, the sampled space \mathcal{F}_1 will contain at least one model from the true anchored Rashomon set.

LEMMA 14. For a finite hypothesis space \mathcal{F}_2 of size $|\mathcal{F}_2|$, we will draw $|\mathcal{F}_1|$ functions uniformly without replacement from \mathcal{F}_2 to form \mathcal{F}_1 . If the true Rashomon ratio of the hypothesis space \mathcal{F}_2 is at least

$$R_{ratio}(\mathcal{F}_2, \gamma) \geq 1 - \epsilon^{\frac{1}{|\mathcal{F}_1|}}$$

then with probability at least $1 - \epsilon$ with respect to the random draw of functions from \mathcal{F}_2 to form \mathcal{F}_1 , the Rashomon set contains at least one model \tilde{f}_1 from \mathcal{F}_1 .

PROOF. The probability of an individual sample from \mathcal{F}_2 missing the true Rashomon set is $1 - R_{ratio}(\mathcal{F}_2, \gamma)$. The probability if this happening $|\mathcal{F}_1|$ times independently is $(1 - R_{ratio}(\mathcal{F}_2, \gamma))^{| \mathcal{F}_1 |}$. Thus, for any $\epsilon > 0$, if the Rashomon ratio is at least $R_{ratio}(\mathcal{F}_2, \gamma) \geq 1 - \epsilon^{\frac{1}{|\mathcal{F}_1|}}$, the probability p_w of sampling, with replacement, at least one hypothesis from $R_{ratio}(\mathcal{F}_2, \gamma)$ is:

$$p_w = 1 - (1 - R_{ratio}(\mathcal{F}_2, \gamma))^{| \mathcal{F}_1 |} \geq 1 - \epsilon.$$

Let p_i be the probability, under sampling without replacement, that samples $1 \dots i$ have missed $R_{ratio}(\mathcal{F}_2, \gamma)$. $p_1 = 1 - R_{ratio}(\mathcal{F}_2, \gamma)$, and $p_i \leq (1 - R_{ratio}(\mathcal{F}_2, \gamma))^i$. The probability, under sampling without replacement, that at least one hypothesis from $R_{ratio}(\mathcal{F}_2, \gamma)$ in \mathcal{F}_1 is therefore $1 - p_{|\mathcal{F}_1|} \geq p_w$. Thus the statement of the lemma holds with probability at least $1 - \epsilon$. \square

Let us recall Theorem 6:

THEOREM 6 (EXAMPLE OF THE ADVANTAGE OF A LARGE TRUE RASHOMON SET). Consider finite hypothesis spaces \mathcal{F}_1 and \mathcal{F}_2 , such that $\mathcal{F}_1 \subset \mathcal{F}_2$ and \mathcal{F}_1 is uniformly drawn from \mathcal{F}_2 without replacement. For loss l bounded by b , if the Rashomon ratio is at least

$$R_{ratio}(\mathcal{F}_2, \gamma) \geq 1 - \epsilon^{\frac{1}{|\mathcal{F}_1|}}$$

then for any $\epsilon > 0$, with probability at least $(1 - \epsilon)^2$ with respect to the random draw of functions from \mathcal{F}_2 to form \mathcal{F}_1 and with respect to the random draw of data, the assumptions of Theorem 5 hold and thus the bound (1) holds.

PROOF. According to the Lemma 14, for any $\epsilon > 0$ with probability at least $1 - \epsilon$ with respect to the random draw of functions, if the Rashomon set is at least $R_{ratio}(\mathcal{F}_2, \gamma) \geq 1 - \epsilon^{\frac{1}{|\mathcal{F}_1|}}$, then the Rashomon set contains at least one model \tilde{f} from \mathcal{F}_1 . In that case, according to Theorem 5 with probability at least $1 - \epsilon$ with respect to the random draw of data, the bound (1) holds. Therefore with probability at least $(1 - \epsilon)^2$ we get the statement of the theorem. \square

C.3 Proof of Theorem 7

THEOREM 7 (EXISTENCE OF MULTIPLE SIMPLER MODELS). *For K -Lipschitz loss l bounded by b , consider hypothesis spaces \mathcal{F}_1 and \mathcal{F}_2 , $\mathcal{F}_1 \subset \mathcal{F}_2$. With probability greater than $1 - \epsilon$ w.r.t. the random draw of training data, if for every model $f_2 \in \hat{R}_{set}(\mathcal{F}_2, \theta)$ there exists $f_1 \in \mathcal{F}_1$ such that $\|f_2 - f_1\|_p \leq \delta$, then there exists at least $B = \mathcal{B}(\hat{R}_{set}(\mathcal{F}_2, \theta), 2\delta)$ functions $f_1^1, f_1^2, \dots, f_1^B \in \hat{R}_{set}(\mathcal{F}_1, \theta)$ such that:*

- (1) *They are from the simpler space: $f_1^1, f_1^2, \dots, f_1^B \in \mathcal{F}_1$.*
- (2) *$|L(f_1^i) - \hat{L}(f_1^i)| \leq 2KR_n(\mathcal{F}_1) + b\sqrt{\frac{\log(2/\epsilon)}{2n}}$, for all $i \in [1, \dots, B]$, where $R_n(\mathcal{F})$ is the Rademacher complexity of a hypothesis space \mathcal{F} . (This is from standard learning theory.)*

PROOF. Starting from the packing number of the Rashomon set $\mathcal{B}(\hat{R}_{set}(\mathcal{F}_2, \theta), 2\delta)$, there exists a 2δ -packing $\Xi = \{\xi_1, \dots, \xi_k | \xi_i \in \hat{R}_{set}(\mathcal{F}_2, \theta)\}$ such that $\|\xi_i - \xi_j\|_p > 2\delta$ for all $i \neq j$. On the other hand, for each $\xi_i \in \hat{R}_{set}(\mathcal{F}_2, \theta)$ there exists $\tilde{f}_1^i \in \mathcal{F}_1$ such that $\|\xi_i - \tilde{f}_1^i\|_p \leq \delta$ (this is the assumption that \mathcal{F}_1 serves as a good cover for \mathcal{F}_2). Therefore for each ball center ξ_i in the packing there is a distinct model \tilde{f}_1^i from the simpler hypothesis space \mathcal{F}_1 . Thus, the Rashomon set contains at least $B = \mathcal{B}(\hat{R}_{set}(\mathcal{F}_2, \theta), 2\delta)$ models from \mathcal{F}_1 .

The generalization bound follows Bartlett and Mendelson [3]. \square

C.4 Examples of function approximation in different hypothesis spaces

Table 3 shows examples of function classes where good approximating sets occur. More specifically, Table 3 describes classes of functions \mathcal{F}_2 that can be approximated with functions from classes \mathcal{F}_1 within δ using a specified norm.

D DATA SET DESCRIPTIONS

We provide a description of the data sets used in our experiments in Table 4. All of them were downloaded from the UCI Machine Learning Repository [18]. We show the number of features in each data set, sizes of the data set and any preprocessing steps that we used mainly to convert data to binary classification. For each data set, we performed cross-validation over ten folds for data sets with more than 200 points and over five folds for data sets with less than 200 points. We reserve one fold for testing, one for validation (e.g., hyper-parameter optimization) and the rest for training. All of the real-valued data sets were normalized to fit the unit-cube, and we did not standardize the data. During data processing, we omitted data records with missing values. We also omitted non-numerical

features (e.g., date or text) when there was not a natural way to convert them to categorical features.

Additionally, we performed experiments on twelve synthetic binary classification data sets. These data sets have two real features and represent different geometrical concepts for two-dimensional classification (e.g., large and small margins, concentric circles, half moons, etc.) as in Figure 7. Results and implications for synthetic data sets are consistent with those on the UCI data sets.

E PERFORMANCE OF DIFFERENT MACHINE LEARNING ALGORITHMS AND RASHOMON RATIO

Figure 14 and Figure 15 show a performance comparison of different machine learning algorithms with regularization for the categorical and real-valued data sets. Data sets shown in Figures 14 and 15 are shown in decreasing order of the Rashomon ratio, from the highest in Figure 14 to the Rashomon ratios that were so small that we were not able to measure them.

Regularization limits the hypothesis space and thus changes the nature of the Rashomon set's measurements. Each value of the regularization parameter corresponds to a soft constraint on the hypothesis space, which in turn can be realized as a hard constraint on this space. The Rashomon ratio in the regularized case will typically be larger or equal to the Rashomon ratio in the unregularized case. There are two reasons for this, explained below.

First, regularization reduces the hypothesis space. Hypotheses that were available when learning without regularization may be excluded when learning with regularization. As a result, the size of the hypothesis space decreases, which increases the Rashomon ratio.

Second, the empirical risk minimizer changes between the regularized and unregularized hypothesis sets, which means the criterion for falling into the Rashomon set changes as well. Recall that the Rashomon set is defined based on the best performing model on the training set. The regularized hypothesis space is less likely to contain overfitted models than the unregularized space. This means the regularized hypothesis space's empirical risk minimizer typically has higher empirical risk than that of the unregularized hypothesis space. Then, if the Rashomon parameter θ is fixed when comparing the two hypothesis spaces, there may be more models in the Rashomon set for the regularized case. Thus, in the regularized case, the size of the Rashomon set would be larger, and, therefore, the Rashomon ratio would be larger too.

Figure 17 and Figure 18 show a comparison of the performance of different machine learning algorithms without regularization for the categorical and real-valued data sets. Figure 16 and Figure 19 shows a performance comparison of different machine learning algorithms for the synthetic data sets with and without regularization.

As we mentioned before, we estimate the Rashomon ratio with importance sampling. For the proposal distribution, we generate a tree of depth D by randomly splitting on features. We assign labels to all 2^D leaves using the training data. If a leaf contains no training points, it acquires its label from the nearest ancestor that any training data pass through. The probability of sampling any tree from the proposal distribution is $p_p = p_f \times \prod_{i=1}^{2^D} 1$, where p_f is the

Table 3: Examples of function approximation in different hypothesis spaces: a function from space \mathcal{F}_1 approximates a function in space \mathcal{F}_2 with given guarantee δ .

\mathcal{F}_2	\mathcal{F}_1	δ (depends on parameters in bounds below)	Source
$f \in L_\infty(\Omega)$, $\ f\ _\infty \in [m, M]$	$s_N \in S(\Omega)$, s_N —piecewise constant, N —number of constants	$\ f - s_N\ _\infty \leq \frac{M-m}{2N}$	Davydov [14], DeVore [15]
$f \in W_p^1(\Omega)$, $1 \leq p \leq \infty$, where W_p^1 is a Sobolev space	$s_\Delta(f) \in S(\Omega)$, s_Δ —piecewise constant, Δ —fixed partition, $\Omega = (0, 1)^d$, N —number of constants	$\ f - s_\Delta(f)\ _p \leq CN^{-1/d} \ f\ _{W_p^1(\Omega)}$	Davydov [14]
$f \in \{x^k, k \in N\}$	$P(n)$ —polynomials of degree at most $n \in N$	$\ f - P(n)\ _\infty \leq \frac{1}{2^{k-1}} \sum_{j>(n+k)/2} \binom{k}{j}$	Newman and Rivlin [37]
$f \in C[0, 1]$ is a non-constant symmetric boolean function on x_1, \dots, x_n	$P(d)$ —algebraic polynomials of degree d	$\ f - P(d)\ _\infty \leq O(\sqrt{n(n - \Gamma(f))})$	Paturi [38]
$f \in Lip_M(\alpha)$, f is Lipschitz continuous with constant M	$N_n : [a, b] \rightarrow \mathbb{R}$ is a feedforward neural network with one layer and bounded, monotone and odd defined activation function, $n \in \mathbb{N}$	$\sup_{x \in [a, b]} f(x) - N_n(x) \leq \frac{5M}{2} \left(\frac{b-a}{n} \right)^\alpha$	Cao et al. [8]
$f \in L_p(I)$, where $I \subset \mathbb{R}^d$ is a cube in \mathbb{R}^d , $\ \cdot\ _{W^r(L_p(I))}$ —Sobolev semi norm	P_r —space of polynomials of order r in d , constant C depends on r	$\inf_{p \in P_r} \ f - p\ _{L_p(I)} \leq C I ^{r/d} \ f\ _{W^r(L_p(I))}$	DeVore [15]

probability of randomly sampling all of the features that comprise the splits of the tree. Our target distribution is a randomly sampled decision tree (both features and leaves) of depth D . Therefore, the probability of sampling a given tree from the target distribution is $p_t = p_f \times \prod_{i=1}^{2^D} \frac{1}{2}$, since we have two classification classes, where, as before, p_f is the probability of randomly sampling all features used within splits of the tree. Thus, for one tree of depth seven, its importance weight will be $\frac{p_t}{p_p} = \left(\frac{1}{2}\right)^{2^7} \approx 3 \times 10^{-39}$. The importance weight clearly dictates the order of magnitude of the Rashomon ratio in our experiments. The smallest possible non-zero Rashomon ratio ($\approx 1.175 \times 10^{-42}\%$) arises when we sample one model that is in the Rashomon set among 250,000 total models that were sampled. Therefore, we consider the Rashomon ratios of order $10^{-37}\%$ and $10^{-38}\%$ to be large, and Rashomon ratios of order $10^{-40}\%$, $10^{-41}\%$, etc., to be small. Note that if there are more than two classes in the data set, the importance weight will be even smaller, as the probability of sampling a random tree from a target distribution decreases. Thus, trees built on a data set with three classes will have lower probability than trees built on a data set with two classes. That is why we considered binary classification only and modified data as described in Table 4, as it is essential for us to compare ratios over different data sets.

If we choose another importance sampling method (for example with data assignment for only half of the leaves or with the guidance of both features and leaves) the Rashomon ratio may have different importance weights and therefore might have a different estimated size as well. This issue would be resolved if we sample a huge number of trees, which is hard to do in practice. Therefore, since our goal is to compare the Rashomon ratios across data sets and feature spaces, we use a consistent method of leaf-based importance sampling across all data sets and sample a manageable number of trees (250,000 in our case).

F LARGE RASHOMON RATIOS MAY APPEAR ARTIFICIALLY SMALL

Even when the Rashomon ratio is a good driver of generalization performance, it may appear artificially small because of a poor representation of data or poor choice of hypothesis space. For instance, if the features are highly correlated, this artificially deflates the size of the Rashomon ratio as discussed in Appendix G. Moreover, if the hypothesis space is poorly designed to include an overly large number of models, then the Rashomon ratio may appear artificially small. The issues with measuring the Rashomon ratio may be a possible explanation for some of the results in Figure 2(b), which includes some data sets with high-performing algorithms, yet (by the way we measured it) a small Rashomon ratio. In any case, small Rashomon ratios are not our main interest; here we are interested in what we would observe under large Rashomon ratios.

G QUALITY OF THE FEATURES

In our experiments, we observed a connection between the quality of the features and Rashomon ratios. The Rashomon ratio in its simplest form, under uniform prior on the hypothesis space, is the fraction of models that are inside the Rashomon set compared to the models in the hypothesis space. When a data set is augmented with additional features, the size of the hypothesis space grows. If the added features are completely irrelevant (consisting, for instance, of noise) then adding these features increases the size of the hypothesis space but does not increase the size of the Rashomon set. Thus, we might predict that the Rashomon ratio could decrease as irrelevant features are added to a data set.

Additionally, if we augment a data set with features that are highly correlated or identical to features that improve performance, then not only is the size of the hypothesis space increased, but also the size of the Rashomon set is likely to increase, as there exist more relevant models (even if the set becomes redundant with

Table 4: Classification data sets description and processing notes.

Data Set Name	Type of Features	Number of Features	Number of Data Points	Processing notes
Monks-1	Binary	15	556	
Monks-2	Binary	15	601	
Monks-3	Binary	15	554	
Voting	Binary	16	232	
SPECT	Binary	22	267	
Tic-tac-toe	Binary	27	958	
Hayes-Roth	Binary	12	160	Considered class 1 versus classes 2 and 3
Nursery-1	Binary	27	8586	Considered classes not_recom and priority
Nursery-2	Binary	27	8310	Considered classes priority and spec_prior
Mushroom	Binary	117	8124	
Breast Cancer	Binary	43	286	
Car Evaluation	Binary	21	1728	Converted to one vs all problem: class 1 versus all others
Primary Tumor	Binary	31	336	Converted to binary classification by considering classes 1, 2, 3, 4, 22, 10 versus all others
Mammographic Masses	Binary	25	830	
Phishing	Binary	23	1353	Considered classes 0 and 1 versus class -1
Balance	Binary	20	576	Considered classes L and R
Wine	Real	13	130	Considered classes 0 and 1
Iris	Real	4	100	Considered classes versicolour and virginica
Breast Cancer Wisconsin	Real	30	569	
Breast Cancer Coimbra	Real	9	116	
Digits 0-4	Real	64	363	Classes 0 and 4 considered only
Digits 6-8	Real	64	355	Classes 6 and 8 considered only
Student	Real	3	400	
Banknote	Real	4	1372	
Mapping	Real	28	10545	Converted to one vs all problem: class forest versus all others
Wifi Localization	Real	7	1000	Considered classes that represent rooms 2 and 3
Column 2C	Real	6	310	
Credit Card	Real	23	30000	
Planing Relax	Real	12	182	
Diabetic Retinopathy	Real	19	1151	
Survival	Real	3	306	
Skin Segmentation	Real	3	245057	
HTRU_2	Real	8	17898	
Magic	Real	10	19020	
Seeds	Real	7	140	Considered classes 1 and 2
Eye State	Real	14	14980	
MNIST 0-1	Real	784	13738	Considered classes 0 and 1
MNIST 4-9	Real	784	12752	Considered classes 4 and 9

models that predict equivalently). Thus, we might predict that the Rashomon ratio increases as we add copies of relevant features.

In general, these two examples of irrelevant and redundant features are corner cases, however, they do occur to a lesser degree in real world data sets, and we are interested in whether these cases have potentially influenced our experimental results in Section 6 in our observed Rashomon ratios. To investigate this, we augmented a data set with noise features, and separately, augmented the same data set with copies of useful features to see whether irrelevant or correlated features may have influenced our findings on the measurement of the Rashomon ratio. We used the Breast Cancer Wisconsin (Diagnostic) data set (shortly, BCW), which has approximately six important features. The results are shown in Figure 8. As before, our hypothesis space is decision trees of depth seven.

Irrelevant features. If the data set contains a lot of irrelevant or noisy features, we expect the Rashomon set to be relatively

small compared to the hypothesis space. Figure 8(a) shows how the Rashomon ratio changes as we iteratively decrease the number of features in the Wine data set, eliminating the least relevant features first, leaving the most significant ones (where relevance is determined according to a χ^2 test with the label). The Rashomon ratio grows as we first remove non-significant features, and after reaching a peak at around six features, it starts to decrease as we remove relevant features, and as models lose accuracy. Similarly, Figure 8(b) shows the influence of noisy features on the Rashomon ratio. Particularly, as we add more noisy irrelevant features, the Rashomon ratio starts to decrease. This is due to the same fact, that we artificially enlarge the hypothesis space while keeping the Rashomon set approximately the same. The noise features do not help improve the empirical risk, they only increase the size of the reasonable set.

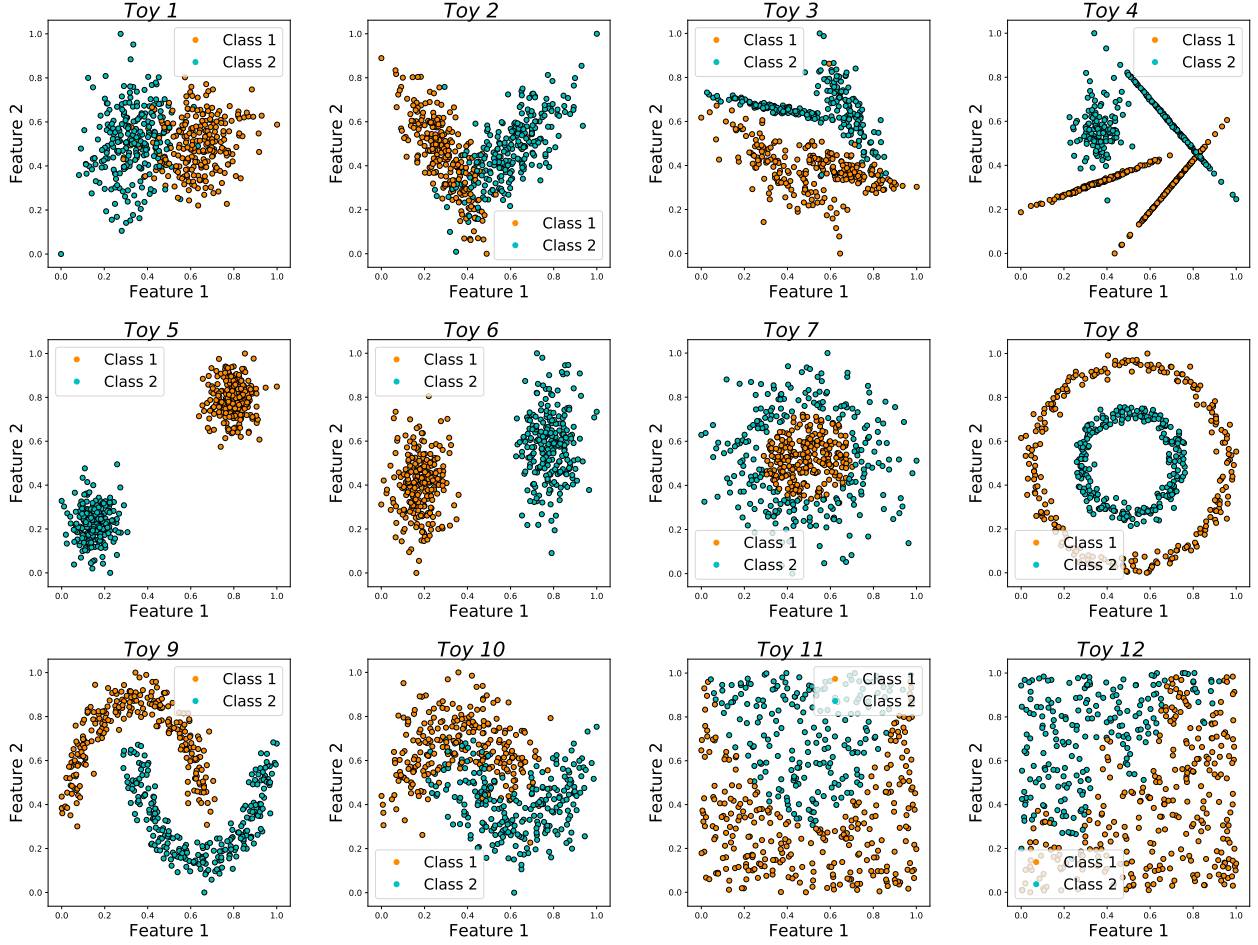


Figure 7: Synthetic two-dimensional data sets that we used in the experiments.

Redundant features. As a contrast to how we increased the hypothesis space in the previous experiment, we can increase the Rashomon set by adding more redundant, good features. Figure 8(c) shows how the Rashomon ratio changes for the BCW data set as we add more copies of the six most significant features. We observe that the Rashomon ratio increases. By adding copies of relevant features, we increased the number of trees at a given depth that could be good enough to be in the Rashomon set.

Our findings show a possible connection between the Rashomon ratio and feature analysis. In particular, in the case where different algorithms perform similarly, but the Rashomon ratio is observed to be small, it could be due to the reason that the data set contains noisy or irrelevant features. In that case, it may be possible to iteratively remove features to find those that produce the largest Rashomon ratio without changes to the empirical risk. The other extreme is less likely to be observed in practice, which is when the Rashomon ratio is extremely large due to redundant features. In that case, one could remove redundant (highly correlated) features before measuring the Rashomon ratio. The data sets with smaller numbers

of features induce easier learning/optimization problems in general. As we discussed earlier, the Rashomon ratio would generally not be measured in practice, and would be inferred in other ways. Thus, these results mainly pertain to an understanding of the experiments we did in Section 6 to provide a possible explanation for cases of small observed Rashomon ratios but where all methods perform the same and all functions generalize.

H PROOF OF THEOREM 8

THEOREM 8 (EXPECTED SIZE OF THE TRUE RASHOMON SET CANNOT DECREASE UNDER RANDOM CLASSIFICATION NOISE). Consider hypothesis space \mathcal{F} , data distribution $\mathcal{D} = \mathcal{X} \times \mathcal{Y}$, where, as before, $\mathcal{X} \in \mathbb{R}^p$, and $\mathcal{Y} \in \{-1, 1\}$. Let $\rho \in (0, \frac{1}{2})$ be a probability with which each label y_i is flipped independently, and \mathcal{D}_ρ denotes the noisy version of \mathcal{D} . If the loss function is $\phi(f(x), y) = \mathbb{1}_{[f(x) \neq y]}$, then in expectation, the true Rashomon set over \mathcal{D} is a subset of the true Rashomon set over \mathcal{D}_ρ , $R_{\text{set}_{\mathcal{D}}}(\mathcal{F}, \gamma) \subseteq R_{\text{set}_{\mathcal{D}_\rho}}(\mathcal{F}, \gamma)$.




Figure 8: An illustration of the influence of feature quality on the Rashomon ratio for the Breast Cancer Wisconsin data set (BCW). (a) shows the Rashomon ratio for the data set with different numbers of significant features according to a χ^2 test. Denote the BCW with the six most significant features as BCW6. (b) depicts the correspondence between the Rashomon ratio and different numbers of noisy features added to the BCW6 data set. The noise features are sampled from normal distribution $\mathcal{N}(0, 1)$ and then standardized to be in a hypercube of volume one. (c) shows the change in the Rashomon ratio as we add more redundant features to the BCW data set. We iteratively add one out of six features from the BCW6 data set at a time. Rashomon ratios in (a)–(c) are averaged over ten folds. The Rashomon parameter θ is set to 0.05. Rashomon ratios are computed with respect to the best sampled model across all variations of the data set.

PROOF. We are given that the probability with which each label is flipped is ρ , $P(\tilde{y}_i \neq y_i) = \rho$, where \tilde{y}_i is a flipped label of y_i . Recall that the true risk $L_{\mathcal{D}}(f) = \mathbb{E}_{(x,y) \sim \mathcal{D}}[\phi(f(x), y)] = \mathbb{E}_{(x,y) \sim \mathcal{D}}[\mathbb{1}_{f(x) \neq y}]$, and f^* is an optimal function, meaning that $f^* \in \operatorname{argmin}_{f \in \mathcal{F}} L_{\mathcal{D}}(f)$. Given the noisy distribution \mathcal{D}_{ρ} , denote $f_{\rho}^* \in \operatorname{argmin}_{f \in \mathcal{F}} L_{\mathcal{D}_{\rho}}(f)$. For any model from the true Rashomon set $f \in R_{\text{set}_{\mathcal{D}}}(\mathcal{F}, \gamma) = \{f : L_{\mathcal{D}}(f) \leq L_{\mathcal{D}}(f^*) + \gamma\}$, we have:

$$\begin{aligned}
 L_{\mathcal{D}_{\rho}}(f) &= \mathbb{E}_{(x, \tilde{y}) \sim \mathcal{D}} [\mathbb{1}_{f(x) \neq \tilde{y}}] \\
 &= \mathbb{E}_{(x, \tilde{y}) \sim \mathcal{D}} [\mathbb{1}_{f(x) \neq \tilde{y}} | \tilde{y} = y] P(\tilde{y} = y) \\
 &\quad + \mathbb{E}_{(x, \tilde{y}) \sim \mathcal{D}} [\mathbb{1}_{f(x) \neq \tilde{y}} | \tilde{y} \neq y] P(\tilde{y} \neq y) \\
 &= \mathbb{E}_{(x, y) \sim \mathcal{D}} [\mathbb{1}_{f(x) \neq y}] P(\tilde{y} = y) \\
 &\quad + \mathbb{E}_{(x, y) \sim \mathcal{D}} [\mathbb{1}_{f(x) = y}] P(\tilde{y} \neq y) \\
 &= \mathbb{E}_{(x, y) \sim \mathcal{D}} [\mathbb{1}_{f(x) \neq y}] (1 - P(\tilde{y} \neq y)) \\
 &\quad + (1 - \mathbb{E}_{(x, y) \sim \mathcal{D}} [\mathbb{1}_{f(x) \neq y}]) P(\tilde{y} \neq y) \\
 &= \mathbb{E}_{(x, y) \sim \mathcal{D}} [\mathbb{1}_{f(x) \neq y}] (1 - 2P(\tilde{y} \neq y)) + P(\tilde{y} \neq y) \\
 &= \mathbb{E}_{(x, y) \sim \mathcal{D}} [\mathbb{1}_{f(x) \neq y}] (1 - 2\rho) + \rho \\
 &= L_{\mathcal{D}}(f)(1 - 2\rho) + \rho.
 \end{aligned}$$

Since $L_{\mathcal{D}}(f^*) \leq L_{\mathcal{D}}(f_{\rho}^*)$ as f^* is an optimal model over \mathcal{D} , $0 < \rho < 0.5$ by assumption, and, again, f is in the true Rashomon set $R_{\text{set}_{\mathcal{D}}}(\mathcal{F}, \gamma)$, we have:

$$\begin{aligned}
 L_{\mathcal{D}_{\rho}}(f) - L_{\mathcal{D}_{\rho}}(f_{\rho}^*) &= L_{\mathcal{D}}(f)(1 - 2\rho) + \rho - L_{\mathcal{D}}(f_{\rho}^*)(1 - 2\rho) - \rho \\
 &= (L_{\mathcal{D}}(f) - L_{\mathcal{D}}(f_{\rho}^*))(1 - 2\rho) \\
 &\leq (L_{\mathcal{D}}(f) - L_{\mathcal{D}}(f^*))(1 - 2\rho) \\
 &\leq \gamma(1 - 2\rho) \leq \gamma.
 \end{aligned} \tag{7}$$

Therefore, f is in the true Rashomon set $R_{\text{set}_{\mathcal{D}_{\rho}}}(\mathcal{F}, \gamma)$. As this calculation holds for every model f from the true Rashomon set $R_{\text{set}_{\mathcal{D}}}(\mathcal{F}, \gamma)$, then in expectation $R_{\text{set}_{\mathcal{D}}}(\mathcal{F}, \gamma) \subseteq R_{\text{set}_{\mathcal{D}_{\rho}}}(\mathcal{F}, \gamma)$.

This proof uses similar technique as Theorem 1 in the work of [30], where the purpose is to show noise tolerance of the optimal model. \square

I EVIDENCE FOR CONJECTURE 9

While there is not yet a proof for the conjecture, there is substantial evidence, which we present in this section. Please see Figure 9 for illustration and details that will help with the intuition for these conjectures.

CONJECTURE 15 (THE RASHOMON SET CAN INCREASE WITH FEATURE NOISE). Consider data distribution $\mathcal{D} = \mathcal{X} \times \mathcal{Y}$, where, $\mathcal{X} \in \mathbb{R}$, $\mathcal{Y} \in \{-1, 1\}$, and classes are balanced $P(Y = -1) = P(Y = 1)$ and generated by Gaussian distributions $P(X|Y = -1) = \mathcal{N}(\mu_1, \sigma^2)$, $P(X|Y = 1) = \mathcal{N}(\mu_2, \sigma^2)$, where $0 \leq \mu_1 < \mu_2$. For the hypothesis space $\mathcal{F} = \{f : f \in (\beta_1, \beta_2)\}$, where $(\mu_1, \mu_2) \subset (\beta_1, \beta_2)$, $\beta_1 \ll \mu_1$, and $\mu_2 \ll \beta_2$, and the Rashomon parameter $\gamma > 0$:

- (I) The volume of the Rashomon set is $\mathcal{V}(R_{\text{set}_{\sigma}}(\mathcal{F}, \gamma)) = |f_{\sigma}^{e_1} - f_{\sigma}^{e_2}|$, where $f_{\sigma}^{e_1}$ and $f_{\sigma}^{e_2}$ are the two solutions to Eqn. (2), where Φ is the CDF of the standard normal:

$$2\Phi\left(\frac{\mu_2 - \mu_1}{2\sigma}\right) - \Phi\left(\frac{\mu_2 - f}{\sigma}\right) - \Phi\left(\frac{f - \mu_1}{\sigma}\right) = \gamma. \tag{2}$$

- (II) We conjecture that for $\mathcal{F} = \{f : f \in (\mu_1, \mu_2)\}$, as we add feature noise to the data set by increasing the standard deviation σ , for all σ such that $\sigma > \tilde{\sigma} = \frac{\mu_2 - \mu_1}{2\sqrt{2}}$, the volume of the Rashomon set increases as a function of σ .
- (III) Consider the setting where $\sigma = 1$ for both Gaussians, and we add or remove noise by moving the means μ_1 and μ_2 of the Gaussians towards or away from each other. For any $\gamma > 0$, the volume of the Rashomon set is minimized when $\mu_2 \approx \mu_1 + 2$. Moving the Gaussians either away from or towards each other increases the volume of the Rashomon set.




Figure 9: The setup for Conjecture 9. We show two Gaussians $\mathcal{N}(\mu_1, \sigma)$ and $\mathcal{N}(\mu_2, \sigma)$, with optimal model $f^* = \frac{\mu_1 + \mu_2}{2}$ (shown in red). Models $f_{\sigma}^{e_1}$ and $f_{\sigma}^{e_2}$ (shown in blue) correspond to the left and right edges of the Rashomon set (shown in purple). The loss of model $f_{\sigma}^{e_2}$ is computed as the sum of two Gaussian tails and is equal to the area of the green region. The objective $G(f_{\sigma}^{e_2}, \sigma) = L(f_{\sigma}^{e_2}) - L(f^*)$ corresponds to the area of the shaded region to the left of function $f_{\sigma}^{e_2}$.

PROOF. Without loss of generality, we will assume that $\mu_1 = 0$ and $\mu_2 = \mu > 0$. To get results for the original values μ_1 and μ_2 , we can simply add μ_1 to μ and f .

Let us show the first point of the conjecture and show how to compute the volume of the Rashomon set.

Evidence for Part (I). Denote ϕ_1 and ϕ_2 as probability density functions (PDF) and Φ_1 and Φ_2 as cumulative distribution functions (CDF) of classes $Y = -1$ and $Y = 1$ correspondingly. For a given model $f \in \mathcal{F}$, the loss can be computed as the sum of areas under the PDF corresponding to misclassification errors:

$$\begin{aligned} L(f) &= P(X > f | Y = -1) + P(X \leq f | Y = 1) \\ &= \int_f^{\infty} \phi_1(t) dt + \int_{-\infty}^f \phi_2(t) dt = 1 - \Phi_1(f) + \Phi_2(f) \\ &= 1 - \Phi\left(\frac{f-0}{\sigma}\right) + \Phi\left(\frac{f-\mu}{\sigma}\right) \\ &= 2 - \Phi\left(\frac{f}{\sigma}\right) - \Phi\left(\frac{\mu-f}{\sigma}\right), \end{aligned} \quad (8)$$

where Φ is the CDF of the normal distribution $\mathcal{N}(0, 1)$.

The optimal model, f^* , can be obtained when $P(X|Y = -1) = P(X|Y = 1)$, meaning that $\frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x-0)^2}{2\sigma^2}\right) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$, which leads to $f^* = \frac{\mu}{2}$. The loss of the optimal model is:

$$L(f^*) = 2 - 2\Phi\left(\frac{\mu}{2\sigma}\right).$$

Denote $G(f, \sigma)$ as the difference in loss between a model f from the Rashomon set and optimal model f^* ,

$$G(f, \sigma) := L(f) - L(f^*) = 2\Phi\left(\frac{\mu}{2\sigma}\right) - \Phi\left(\frac{\mu-f}{\sigma}\right) - \Phi\left(\frac{f}{\sigma}\right).$$

To find a model on the edge of the true Rashomon set, we set $G(f, \sigma) = \gamma$ and obtain Equation (2). As the problem is symmetric with respect to f^* and $\gamma > 0$, solutions $f_{\sigma}^{e_1}$ and $f_{\sigma}^{e_2}$ to Equation (2) correspond to the left and right edges of the true Rashomon set. Thus, for a given γ and σ , we can estimate the volume of the Rashomon set as $\mathcal{V}(\mathcal{R}_{set_{\sigma}}(\mathcal{F}, \gamma)) = |f_{\sigma}^{e_1} - f_{\sigma}^{e_2}|$.

Now we will show evidence for the second point of the conjecture that the volume of the Rashomon set increases when σ increases.

Evidence for Part (II). Let $\tilde{\sigma} = \frac{\mu}{2\sqrt{2}}$. Let's focus on the right edge of the true Rashomon set where $f \in (f^*, \mu)$. We will show that:

- (i) For any fixed σ , and for $f \in (f^*, \mu)$, $G(f, \sigma)$ is monotonically increasing in f .
- (ii) For any fixed $f \in (f^*, \mu)$ and $\sigma > \tilde{\sigma}$, $G(f, \sigma)$ is monotonically decreasing in σ .
- (iii) If (i) and (ii) hold, then for any σ_1 and σ_2 such that $\tilde{\sigma} < \sigma_1 < \sigma_2$, we have that $\mathcal{V}(\mathcal{R}_{set_{\sigma_1}}(\mathcal{F}, \gamma)) < \mathcal{V}(\mathcal{R}_{set_{\sigma_2}}(\mathcal{F}, \gamma))$.

We will prove (i) and (iii), and conjecture that (ii) is true, based on numerical experiments.

Let's focus on (iii) first. Consider σ_1 and σ_2 such that $\tilde{\sigma} < \sigma_1 < \sigma_2$. Let $f_{\sigma_1}^e, f_{\sigma_2}^e \in (f^*, \mu)$ be the right edge models of the corresponding true Rashomon sets $G(f_{\sigma_1}^e, \sigma_1) = \gamma$ and $G(f_{\sigma_2}^e, \sigma_2) = \gamma$. The volume of the Rashomon set is $\mathcal{V}(\mathcal{R}_{set_{\sigma}}(\mathcal{F}, \gamma)) = |f_{\sigma_1}^e - f_{\sigma_2}^e| = 2|f_{\sigma_2}^e - f^*| = 2|f_{\sigma_1}^e - f^*|$. Using monotonicity of $G(f, \sigma)$ with respect to σ given fixed f (the result of part (ii)), we get that:

$$G(f_{\sigma_1}^e, \sigma_2) < G(f_{\sigma_1}^e, \sigma_1) = \gamma = G(f_{\sigma_2}^e, \sigma_2),$$

which means that $f_{\sigma_1}^e < f_{\sigma_2}^e$ due to monotonicity of $G(f, \sigma)$ with respect to f given fixed σ (the result of part (i)). Therefore, as we increase the standard deviation from σ_1 to σ_2 , the Rashomon set increases as well:

$$\mathcal{V}(\mathcal{R}_{set_{\sigma_1}}(\mathcal{F}, \gamma)) = 2|f_{\sigma_1}^e - f^*| < 2|f_{\sigma_2}^e - f^*| = \mathcal{V}(\mathcal{R}_{set_{\sigma_2}}(\mathcal{F}, \gamma)).$$

Thus, we have proved part (iii).

Now, let us prove (i). Given fixed σ , the derivative of the objective $G(f, \sigma)$ with respect to f is:

$$\begin{aligned} G'_f(f, \sigma) &= \frac{1}{\sigma} \Phi'_f\left(\frac{\mu-f}{\sigma}\right) - \frac{1}{\sigma} \Phi'_f\left(\frac{f}{\sigma}\right) \\ &= \frac{1}{\sqrt{2\pi}\sigma} \left(e^{-\frac{(\mu-f)^2}{2\sigma^2}} - e^{-\frac{f^2}{2\sigma^2}} \right) > 0, \end{aligned}$$

since $\mu - f < f$ since $f > f^* = \frac{\mu}{2}$. Since the derivative $G'_f(f, \sigma) > 0$, the objective $G(f, \sigma)$ monotonically increases in f .

Finally, let's show (ii). Given fixed $f \in (f^*, \mu)$ and $\sigma > \tilde{\sigma}$, the derivative of the objective with respect to σ is:

$$\begin{aligned} G'_\sigma(f, \sigma) &= \frac{-2(\mu)}{2\sqrt{2\pi}\sigma^2} \Phi'_\sigma\left(\frac{\mu}{2\sigma}\right) + \frac{\mu-f}{\sqrt{2\pi}\sigma^2} \Phi'_\sigma\left(\frac{\mu-f}{\sigma}\right) + \frac{f}{\sqrt{2\pi}\sigma^2} \Phi'_\sigma\left(\frac{f}{\sigma}\right) \\ &= \frac{1}{\sqrt{2\pi}\sigma^2} \left[-\mu e^{-\frac{\mu^2}{8\sigma^2}} + (\mu-f)e^{-\frac{(\mu-f)^2}{2\sigma^2}} + f e^{-\frac{f^2}{2\sigma^2}} \right] \\ &= \frac{f}{\sqrt{2\pi}\sigma^2} \left[-\frac{\mu}{f} \left(e^{\frac{f^2}{2\sigma^2}} \right)^{-\frac{(\mu/f)^2}{4}} + \left(\frac{\mu}{f} - 1 \right) \left(e^{\frac{f^2}{2\sigma^2}} \right)^{-((\mu/f)-1)^2} \right. \\ &\quad \left. + \left(e^{\frac{f^2}{2\sigma^2}} \right)^{-1} \right]. \end{aligned}$$

Denote $u = \frac{\mu}{f}$. Since, $\frac{\mu}{2} < f < \mu$, then $u \in (1, 2)$. Denote $a = e^{\frac{f^2}{2\sigma^2}}$, then $\sigma^2 = \frac{f^2}{2\log(a)}$. Note that $a > 1$, and since $\sigma > \tilde{\sigma} = \frac{\mu}{2\sqrt{2}}$, $a < e^{\frac{4f^2}{\mu^2}} = e^{\frac{4}{u^2}} = s(u)$. Then $G'_\sigma(f, \sigma)$ can be expressed in terms of parameter u and variable a :

$$\begin{aligned} G'_\sigma(u, a) &= \frac{2\log(a)}{\sqrt{2\pi}f} \left[-ua^{-\frac{u^2}{4}} + (u-1)a^{-(u-1)^2} + a^{-1} \right] \\ &= \frac{2\log(a)}{\sqrt{2\pi}f} D(u, a). \end{aligned}$$

As $a > 1$ and $f > \frac{\mu}{2} > 0$, $\frac{2\log(a)}{\sqrt{2\pi}f} > 0$. To show that $D(u, a) < 0$ for any $u \in (1, 2)$ and any $a \in (1, s(u))$, we perform exhaustive numerical calculations spanning the possible values of $u \in (1, 2)$ and $a \in (1, s(u))$ in Figure 10. Indeed, $D(u, a) = 0$ when $u = 2$ or $a = 1$, and for all other values of u and a , $D(u, a) < 0$. Therefore, $G'_\sigma(f, \sigma) = G'_\sigma(u, a) < 0$. Since the derivative is negative, when $\sigma > \tilde{\sigma}$ the objective $G(f, \sigma)$ monotonically decreases in σ , which concludes our evidence for (ii), and thus part (II), of the conjecture.

Evidence for Part (III). Now we add or remove noise from the data set by moving the two means closer together or further apart (see Figure 11), for a fixed $\sigma = 1$. Recall that without a loss of generality, we take $\mu_1 = 0$ and denote $\mu = \mu_2$. The optimal model f^* is $f^* = \frac{\mu}{2}$. Given $\sigma = 1$ and $\gamma > 0$, we are interested in finding $\mu > 0$ and edge model $f \in (0, \mu/2)$ such that the volume of the Rashomon set is minimal. Therefore we have the following optimization problem:

$$\begin{aligned} \min_{\mu} \quad & \frac{\mu}{2} - f_\mu \text{ s.t. } f_\mu \text{ is defined by} \\ & 2\Phi\left(\frac{\mu}{2}\right) - \Phi(\mu - f_\mu) - \Phi(f_\mu) = \gamma. \end{aligned} \quad (9)$$

We cannot solve optimization problem (9) directly for the best μ for each γ , but we provide numerical solutions in Figure 12 for a range of values of γ . We observe that, regardless of the value of $\gamma > 0$, as we minimize the volume of the Rashomon set to find μ , the μ corresponding to the optimal solution is always approximately equal to 2. The edge model $f^e \in (0, 1)$ is then the one that satisfies the constraint: $2\Phi(1) - \Phi(2 - f^e) - \Phi(f^e) = \gamma$. (The edge model must vary with γ since the optimal μ does not.)

The value $\mu = 2$ might seem surprising as a solution for *any* $\gamma > 0$. However, when $\mu_1 = 0$ and $\mu_2 = 2$, $f^* = 1$, which is one standard deviation away from each mean and therefore corresponds

to a value where the inflection points of the two Gaussians coincide (normal distribution $\mathcal{N}(\mu, \sigma)$ has two inflection points $\mu \pm \sigma$). Given fixed γ , as we move either of the means of the Gaussians so that the inflection points no longer coincide, f^e moves outward to compensate so that the constraint in optimization problem (9) is satisfied. Therefore, the volume of the Rashomon set grows as we increase or decrease feature noise by moving μ away from 1 (in either direction) as shown in Figure 13. \square




Figure 10: Numerically showing that $D(u, a) < 0$ for any $u \in (1, 2)$, and $a \in (1, e^{\frac{4}{u^2}})$. Red plane corresponds to the value 0.




Figure 11: In part (III), we add or remove noise from the data by moving the mean of the right Gaussian. Both Gaussians have standard deviation 1.




Figure 12: Numerical solution to the optimization problem (9). For each fixed γ , we plot $f_{\sigma}^{e_1} = f^e \in (0, \mu/2)$, such that the volume of the Rashomon set is minimized. The color of the scatter plot points correspond to the value of the volume of the Rashomon set, where more intense color means higher value.




Figure 13: An example that shows that in part (III) as we move the right Gaussian away from a mean of 2 in either direction, the volume of the Rashomon set increases.

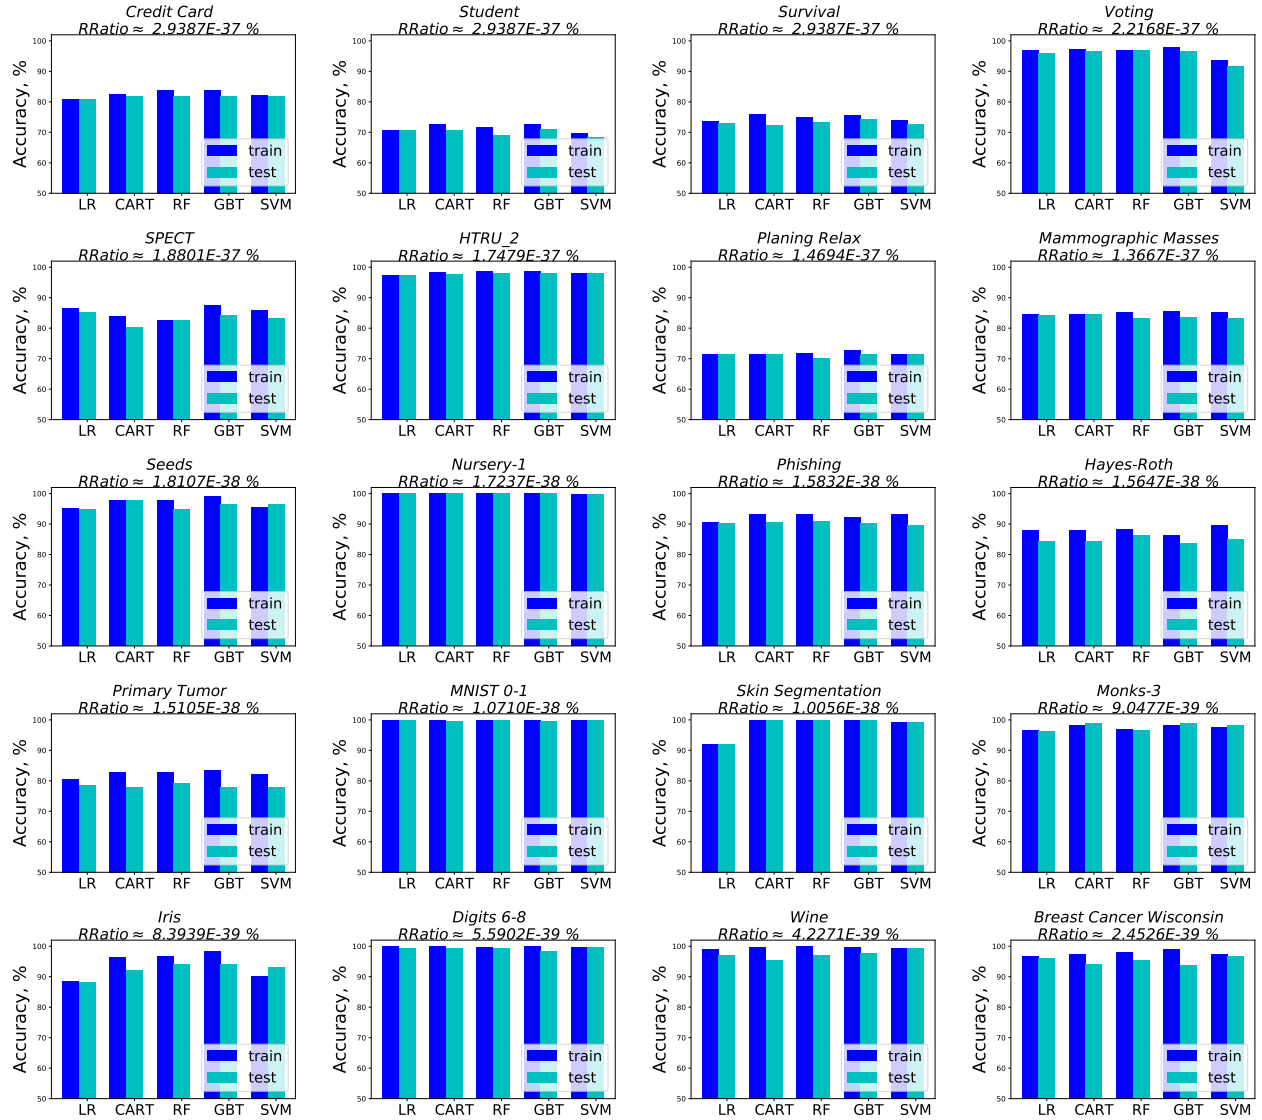


Figure 14: Performance of five machine learning algorithms with regularization for the UCI classification data sets. Data sets are listed in decreasing order of Rashomon ratio. Rashomon ratios, train and test accuracies are averaged over ten folds for data sets with more than 200 points and over five folds for data sets with less than 200 points. These plots continue in Figure 15. In these cases, test performance seems to be similar across algorithms. This will not be true in all cases as the Rashomon set becomes smaller, in Figure 15.

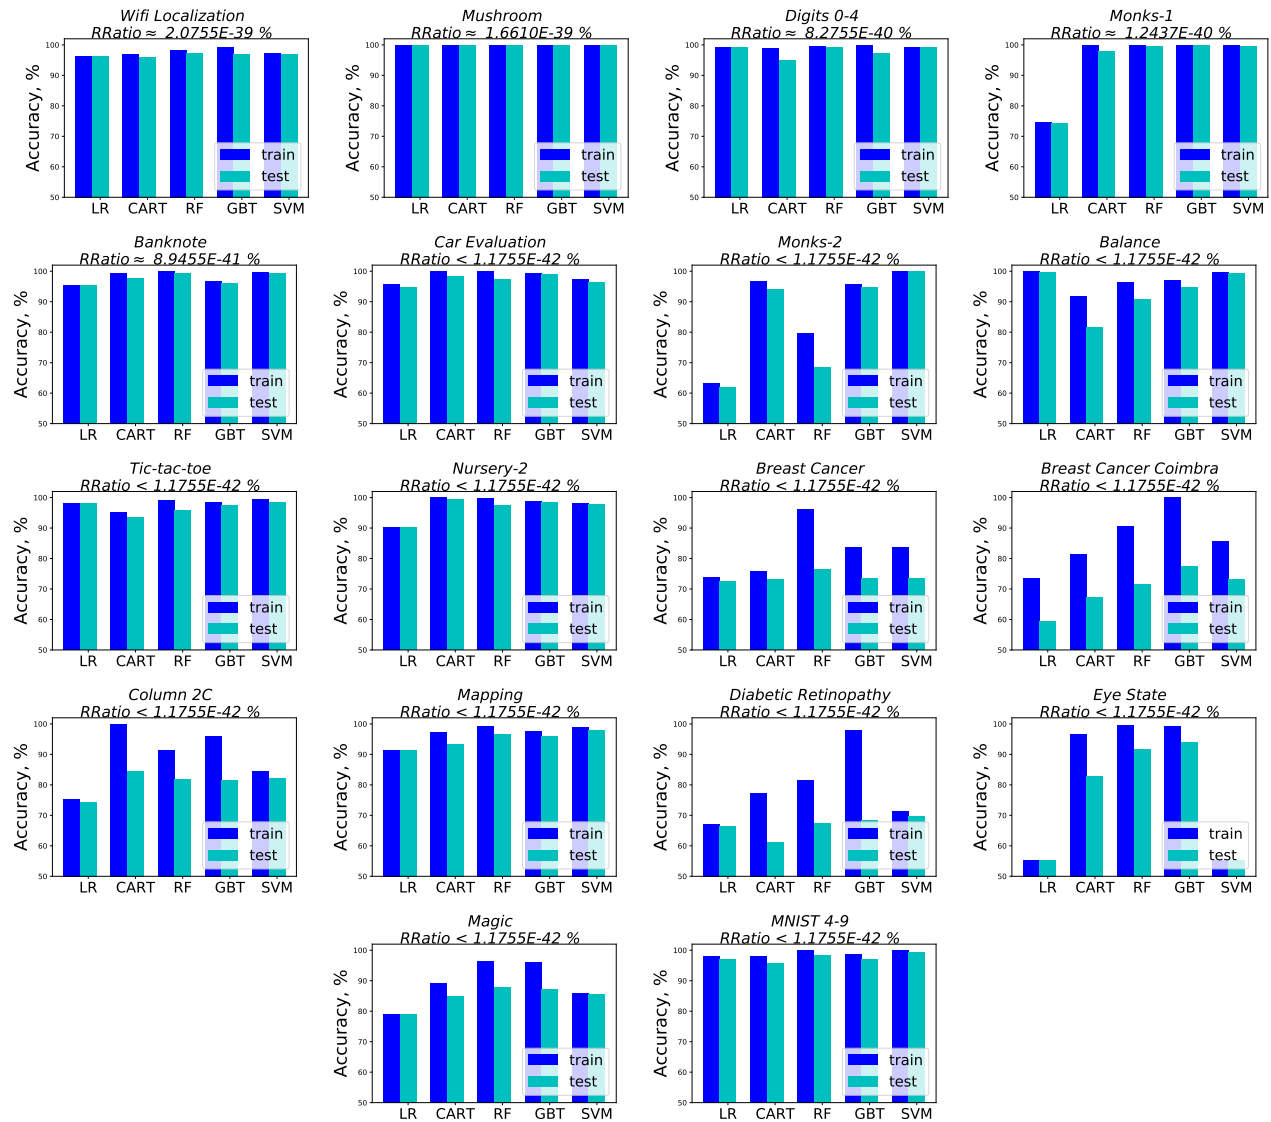


Figure 15: Performance of five machine learning algorithms with regularization for the UCI classification data sets. Data sets are listed in decreasing order of the Rashomon ratio, continued from Figure 14. Rashomon ratios, training accuracies, and test accuracies are averaged over ten folds for data sets with more than 200 points and over five folds for data sets with less than 200 points. Test performance sometimes varies across algorithms.

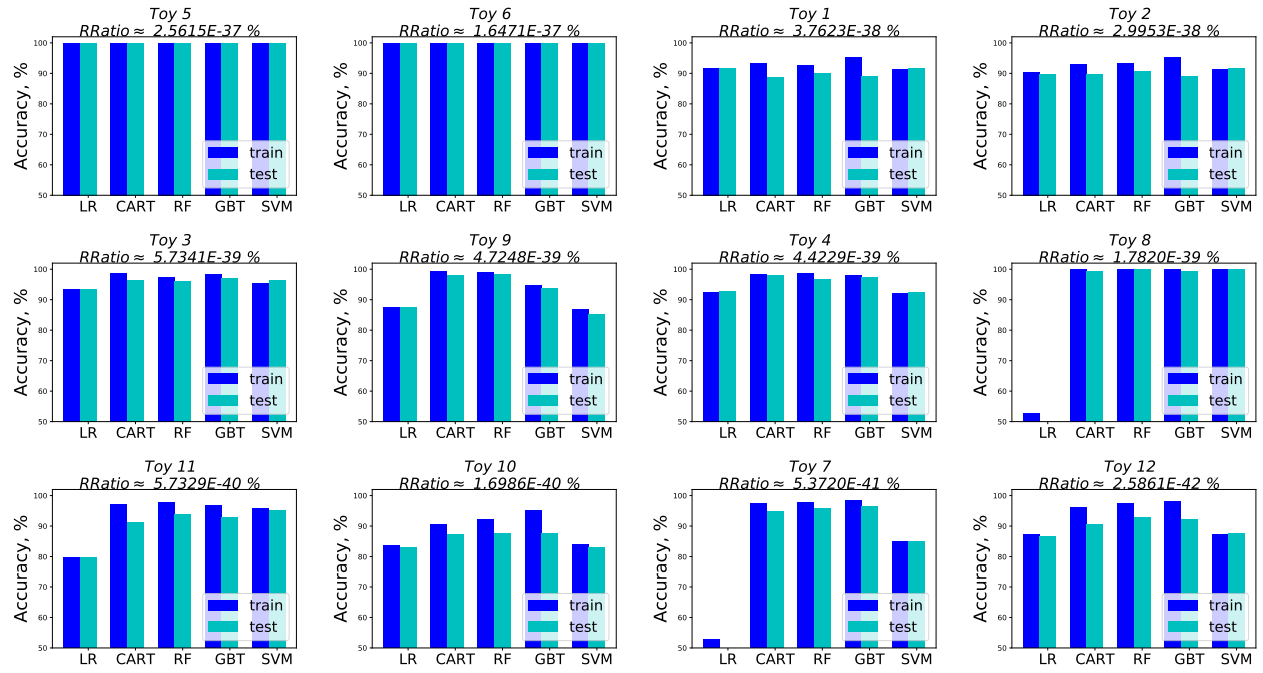


Figure 16: Performance of five machine learning algorithms with regularization for the synthetic data sets with real-valued features. Data sets are listed in decreasing order of the Rashomon ratio. Rashomon ratios, train and test accuracies are averaged over ten folds.




Figure 17: Performance of five machine learning algorithms without regularization for the UCI classification data sets. Data sets are listed in decreasing order of Rashomon ratio. Rashomon ratios, train and test accuracies are averaged over ten folds for data sets with more than 200 points and over five folds for data sets with less than 200 points. These plots continue in Figure 18. The data sets with larger Rashomon ratios correlate with similar performance of machine learning algorithms and good generalization.

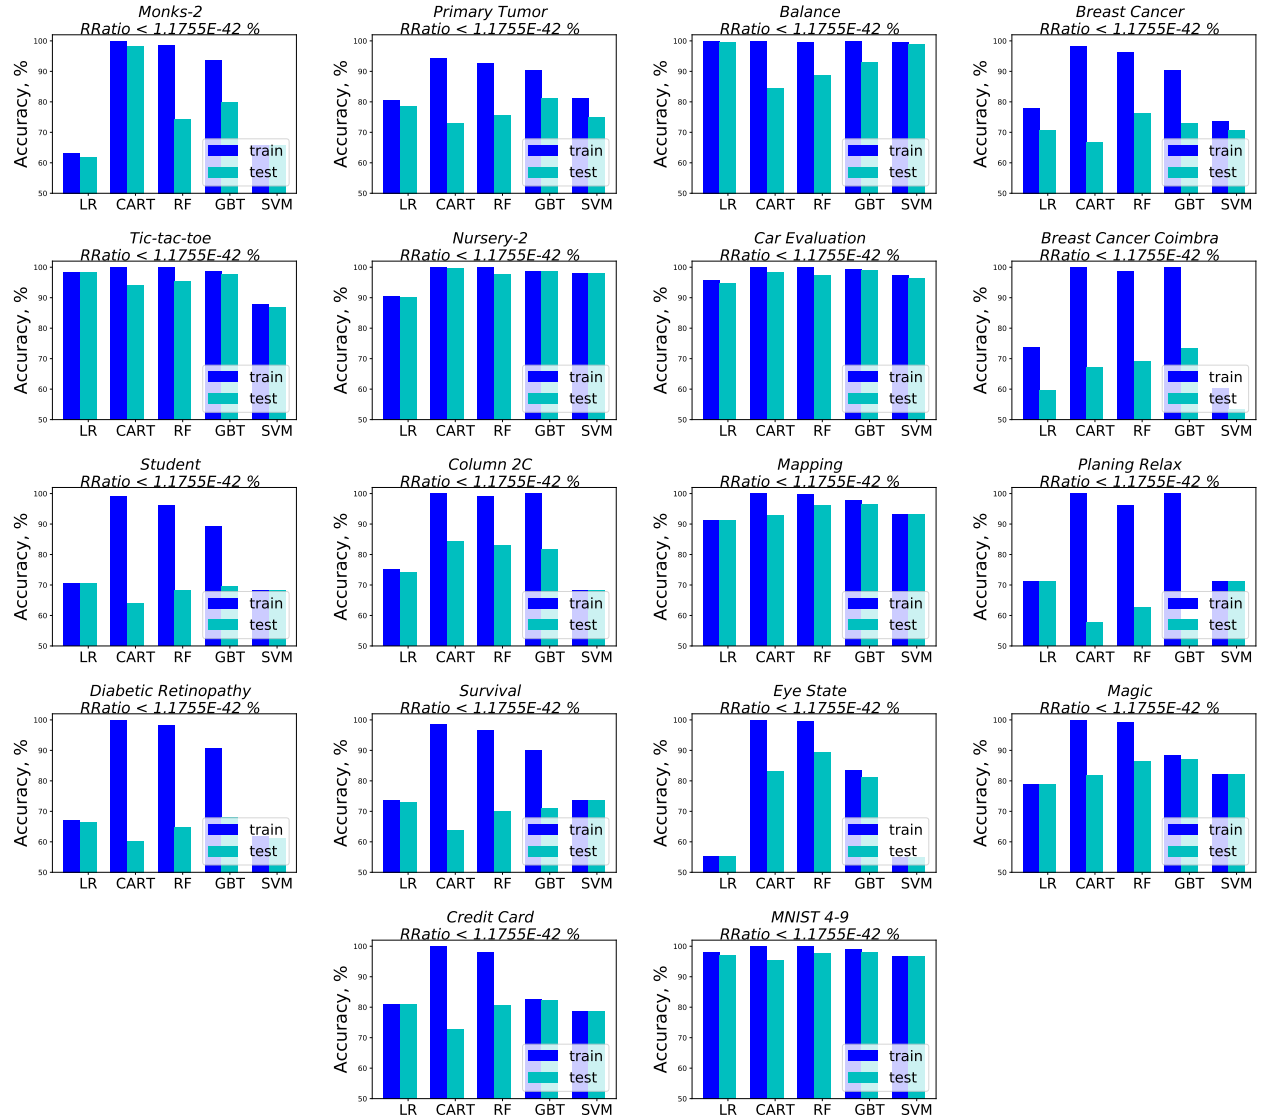


Figure 18: Performance of five machine learning algorithms without regularization for the UCI classification data sets. Data sets are listed in decreasing order of the Rashomon ratio continuing from Figure 17. Rashomon ratios, train and test accuracies are averaged over ten folds for data sets with more than 200 points and over five folds for data sets with less than 200 points

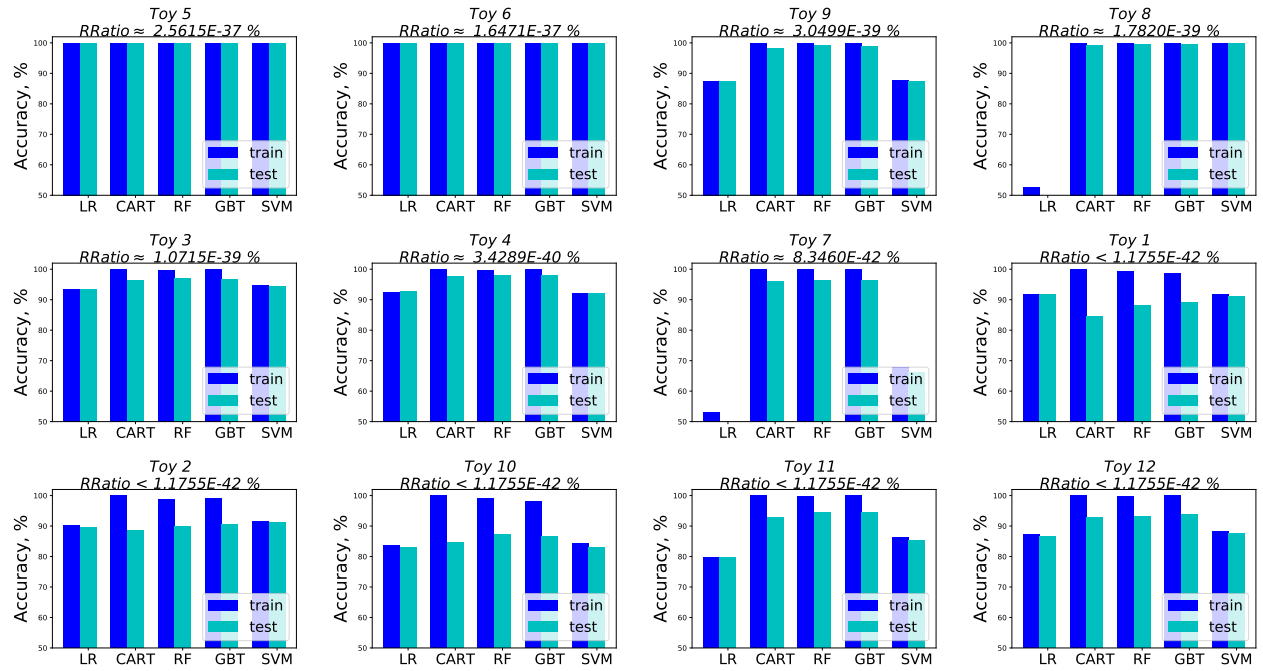


Figure 19: Performance of five machine learning algorithms without regularization for the synthetic data sets with real-valued features. Data sets are listed in decreasing order of the Rashomon ratio. Rashomon ratios, train and test accuracies are averaged over ten folds.