



INFORMS Journal on Data Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Conjecturing-Based Discovery of Patterns in Data

J. Paul Brooks, David J. Edwards, Craig E. Larson, Nico Van Cleemput

To cite this article:

J. Paul Brooks, David J. Edwards, Craig E. Larson, Nico Van Cleemput (2024) Conjecturing-Based Discovery of Patterns in Data. INFORMS Journal on Data Science

Published online in Articles in Advance 02 Feb 2024

<https://doi.org/10.1287/ijds.2021.0043>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2024, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.


For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Conjecturing-Based Discovery of Patterns in Data

J. Paul Brooks,^{a,*} David J. Edwards,^b Craig E. Larson,^c Nico Van Cleemput^d

^aDepartment of Information Systems, Virginia Commonwealth University, Richmond, Virginia 23284; ^bDepartment of Statistical Sciences and Operations Research, Virginia Commonwealth University, Richmond, Virginia 23284; ^cDepartment of Mathematics and Applied Mathematics, Virginia Commonwealth University, Richmond, Virginia 23284; ^dDepartment of Applied Mathematics, Computer Science and Statistics, Ghent University, 9000 Ghent, Belgium

*Corresponding author

Contact: jpbrooks@vcu.edu,  <https://orcid.org/0000-0003-0423-8422> (JPB); dedwards7@vcu.edu (DJE); clarson@vcu.edu (CEL); nico.vanCleemput@gmail.com (NVC)

Received: September 18, 2021
 Revised: July 28, 2022; March 19, 2023;
 August 13, 2023; December 2, 2023
 Accepted: December 15, 2023
 Published Online in *Articles in Advance*:
 February 2, 2024

<https://doi.org/10.1287/ijds.2021.0043>

Copyright: © 2024 INFORMS

Abstract. We propose the use of a conjecturing machine that suggests feature relationships in the form of bounds involving nonlinear terms for numerical features and Boolean expressions for categorical features. The proposed CONJECTURING framework recovers known nonlinear and Boolean relationships among features from data. In both settings, true underlying relationships are revealed. We then compare the method to a previously proposed framework for symbolic regression on the ability to recover equations that are satisfied among features in a data set. The framework is then applied to patient-level data regarding COVID-19 outcomes to suggest possible risk factors that are confirmed in the medical literature. Discovering patterns in data is a first step toward establishing causal relationships, which can be the basis for effective decision making.

Data Ethics & Reproducibility Note: Code and data to reproduce results are available here: <https://github.com/jpbrooks/conjecturing>. COVID-19 synthetic patient data were obtained as part of the Veterans Health Administration (VHA) Innovation Ecosystem and precisionFDA COVID-19 Risk Factor Modeling Challenge and are used here with permission from the Food and Drug Administration (FDA). The e-companion is available at <https://doi.org/10.1287/ijds.2021.0043>.

History: Olivia Sheng served as the senior editor for this article.

Keywords: [automated conjecturing](#) • [computational scientific discovery](#) • [interpretable artificial intelligence](#) • [nonlinear pattern discovery](#) • [Boolean pattern discovery](#)

1. Introduction

Modern machine learning methods allow one to leverage complex relationships present in data to generate accurate predictions but do not reveal them to the investigator. We propose an automated conjecturing framework for discovering nonlinear and Boolean relationships among the features in a given data set. Our primary goal is discovery—to provide the investigator with a manageable number of suggested relationships to inspire future investigation for validation.

The nonlinear relationships are produced in the form of bounds. Bounds are useful for scientific discovery from numeric data because they (1) suggest direct and indirect relationships among features, (2) suggest a functional form for the relationships, and (3) can subsequently be used as Boolean features (e.g., is this bound satisfied by an observation?) for discovering more complex Boolean relationships. Whereas previous related approaches seek to find equations for numeric data, our CONJECTURING method produces bounds for numeric data, Boolean expressions for discrete data, and bounds and Boolean expressions for mixed data.

Udrescu and Tegmark (2020) proposed a system called AI FEYNMAN that combines deep learning with methods for symbolic regression to recover nonlinear relationships in data. Impressively, they recover over 100 equations of varying complexity from data. In contrast to AI FEYNMAN, our CONJECTURING framework uses Fajtlowicz’s Dalmatian heuristic (Fajtlowicz 1995) to discover bounds rather than equations. Further, our framework can also be applied to categorical data to discover Boolean relationships among features and already-discovered bounds. This work represents the first application of the Dalmatian heuristic to learning both nonlinear and Boolean relationships from data. The bounds and conditions produce interpretable yet complex relationships.

2. Background and Previous Related Work

In this section, we provide background on our CONJECTURING framework, including examples of uses of bounds and sufficient conditions, a description of the core algorithm, and a survey of previous related work.

2.1. Conjectured Bounds and Sufficient Conditions

The algorithm we use to conjecture feature relationships is an adaptation of an algorithm that was originally designed to conjecture relationships for mathematical objects. To illustrate the potential value of bounds and sufficient conditions, we describe two problems and relevant results from graph theory. This paper extends these ideas regarding bounds and sufficient conditions to learning from data.

A graph is a collection of nodes, V , and edges, E , that are ordered pairs of nodes. Consider the problem of finding bounds for the independence number of a graph.¹ It is well-known that the linear programming (LP) relaxation of an appropriate integer program provides an upper bound on the independence number (Schrijver 2003). The Lovász ϑ number of a graph also provides an upper bound that is known to be no larger than the LP relaxation bound for any graph (Lovász 1979). Therefore, the Lovász ϑ bound dominates the LP relaxation bound, and such relationships are commonly pursued. However, relationships among bounds can be more nuanced. Consider a third bound on the independence number derived by Haemers (1979). For some graphs, it is a stronger bound than Lovász ϑ , whereas on other graphs, it is a weaker bound; for some graphs, Lovász ϑ is a sharp bound and Haemers's bound is not whereas for other graphs, Haemers's bound is a sharp bound, and Lovász ϑ is not. It remains an open question whether there are a "small" number of bounds where the largest value for any graph would provide a sharp bound on the independence number. In this paper, we describe a computational approach to discover bounds among numeric features in a data set. As with the independence number, collections of bounds can provide valuable insight into relationships for the system from which the data were collected.

Now consider the problem of finding sufficient conditions for a graph to be Hamiltonian.² Chvátal (1972) proved that for a graph G with certain conditions on the vertex degrees, G is Hamiltonian. Also, Chvátal and Erdős (1972) proved that if a graph satisfies a connectivity condition, then it is Hamiltonian. These are two conditions that are sufficient for a graph to be Hamiltonian, but neither implies the other. Some graphs satisfy both conditions, some graphs satisfy one condition, and some graphs satisfy neither condition. The existence and discovery of a (small) set of sufficient conditions that characterize all Hamiltonian graphs remain an open area of research. The pursuit of sufficient conditions of graph properties such as Hamiltonicity mirrors that of bounds (Larson and Van Cleemput 2017). In the context of learning from data, we show how categorical data, together with bounds discovered among numeric features, can be used as input to a computational approach for generating sufficient conditions for a property of interest.

2.2. The Dalmatian Heuristic


Our CONJECTURING framework is based on an implementation of Fajtlowicz's Dalmatian heuristic (Fajtlowicz 1995, Larson and Van Cleemput 2017). The heuristic was originally implemented in GRAFFITI (Fajtlowicz 1995), which was the first program to produce research conjectures that led to new mathematical theories. The program produces statements that are relations between mathematical *invariants*, which are numerical attributes of examples. Recent implementations of the Dalmatian heuristic have been applied to the discovery of relationships for graphs (Larson and Van Cleemput 2016) and game strategies (Bradford et al. 2020). The heuristic was adapted to work with *properties* that are Boolean attributes of examples by Larson and Van Cleemput (2017). We built our framework using a more recent implementation of the Dalmatian heuristic available.³

We now describe invariant conjecturing using Fajtlowicz's Dalmatian heuristic. The inputs include the following. Let E be a set of examples of a given type (e.g., graphs or data observations). Let $A = \{\alpha_1, \alpha_2, \dots, \alpha_m\}$ be real number invariants. In this work, the examples are n data observations, and the invariants are m numeric features. The real-numbered value of example i for invariant α_j is $\alpha_j(i) = x_{ij}$ for $i = 1, \dots, n$ and $j = 1, \dots, m$. Let O be a collection of *unary operators* and *binary operators*. Examples of unary operators include adding one, squaring, square-rooting, and division by two. Binary operators include addition, multiplication, and subtraction. Let $\alpha^* \in A$ be the invariant for which upper and lower bounds are of interest, and let $\alpha^*(i)$ be the value of the invariant of interest for example i .

The aim is to generate conjectured bounds that are true for any realization of input examples E . The Dalmatian heuristic provides criteria for generating conjectured bounds that are the best for E . Algorithm 1 provides a way to generate expressions of increasing complexity, apply the heuristic, and store conjectures. The *complexity* of an expression is the number of nodes in the corresponding expression tree (Figure 1) and is the sum of the number of invariants, number of unary operators, and number of binary operators. The algorithm proceeds by generating unlabeled trees and then labeling the nodes with operators and invariants. Expressions satisfying the Dalmatian heuristic conditions are retained as conjectures \mathcal{C} . For a conjecture $c \in \mathcal{C}$, let $c(i)$ be the conjectured bound for example i .

With examples E , invariants A , operators O , invariant of interest α^* , an upper limit on the proportion of missing values allowed for an invariant *skips*, and a direction indicating if the algorithm will produce upper or lower bounds (*UPPER* or *LOWER*), **procedure** CONJECTURING-INV is called (Algorithm 1, line 1). The number of unary nodes u and binary nodes b of an expression tree are initialized to zero, and the conjecture \mathcal{C} is initialized to the empty set (Algorithm 1, line 2). Line 3 of Algorithm 1

Figure 1. Expression Trees



Notes. (a) An upper bound on square footage $x_1/x_2 + x_3$, where x_1 is 300K, x_2 is pricePerSquareFoot, and x_3 is bathrooms. (b) Gravitational force km_1m_2/r^2 .

refers to the stopping criteria of the expression generator. For invariant conjecturing for upper bounds, if the minimum conjectured bound is tight for each example (i.e., $\min_{c \in \mathcal{C}} c(i) = \alpha^*(i)$ for $i \in E$), then the expression generator is stopped. Otherwise, expression generation continues until a time limit is reached. If exact bounds are not discovered for each example, more complex expressions are generated for larger time limits.

Line 4 calls a procedure to generate a tree, the branching nodes of which will be operators and the leaf nodes of which are invariants. Lines 5–11 enumerate every tree where each vertex connected to a leaf node has degree one or two. These branching nodes will correspond to unary or binary operators, respectively, when the tree is labeled. The leaf nodes will correspond to invariants. Unlabeled trees are grown recursively, and then the nodes are labeled with operators and invariants.

The **procedure** GENERATE TREE (Algorithm 1, line 4) creates a new tree with a single node and then calls **procedure** GENERATE TREE REC to add new nodes until there are u unary nodes and b binary nodes. The **procedure** GENERATE TREE REC (Algorithm 1, line 19) either calls GENERATE LABELED TREE to apply labels by assigning invariants to leaf nodes and operators to branching nodes to generate an

expression (Algorithm 1, line 21) or adds nodes to grow the tree (lines 23–32).

The **procedure** GENERATE LABELED TREE (Algorithm 1, line 35) takes as input a tree with u unary nodes and b binary nodes. Line 36 orders the nodes so that child nodes appear before their parent. Then line 37 creates a set of labeled trees. The leaf nodes are labeled with invariants, and the branching nodes are labeled with operators. Invariants with more than *skips* missing values among examples are not used for labeling. For the commutative binary operators, the left child is larger than the right if the left has more nodes. If the number of nodes is equal, we use the lexicographically largest string of labels. Because the suffix order guarantees that all subtrees are fully labeled before their parent is labeled, this is an unambiguous definition. Examples of labeled expression trees are given in Figure 1.

Lines 39 and 41 are the Dalmatian heuristic. A conjectured upper bound c is only retained in the database of conjectures \mathcal{C} if the bound passes the following two tests:

1. (*Truth test*). The candidate conjecture $\alpha^*(i) \leq c(i)$ is true for all examples $i \in E$, and
2. (*Nondominance test*.) There is an example i where $c(i) < \min\{c'(i) : c' \in \mathcal{C} \setminus \{c\}\}$. That is, the candidate conjecture would give a better bound for $\alpha^*(i)$ than any previously conjectured (upper) bound.

Line 41 ensures that the number of conjectures is no larger than the number of examples, that is, $|\mathcal{C}| \leq |E|$.

The procedure is the same for generating lower bounds, with the only difference being how the Dalmatian heuristic criteria are evaluated in lines 39 and 41.

The computational requirements of Algorithm 1 increase exponentially with the number of invariants and with the number of operators. The computation time per expression increases with the number of examples because of the check in step 39. To facilitate generation of more candidate expressions in less time, one can use fewer examples as input to the algorithm. To achieve additional efficiency, we implement the following design choice. In our implementation, when a tree is labeled, operators can be reused, but invariants cannot. We make this design choice so that more expressions can be generated in a smaller amount of time. In Section 5.4, we will demonstrate how this limitation can be overcome in situations where repeating invariants is warranted.

Algorithm 1 (Invariant Conjecturing)

Input: Examples E , Invariants A , operators O , invariant of interest α^* , invariant missing value limit *skips*, direction (*UPPER* or *LOWER*).

Output: Conjectured \mathcal{C} in the form of conjectured bounds on the invariant of interest α^* .

- 1: **procedure** CONJECTURING-INV
- 2: Set $u = 0, b = 0, \mathcal{C} = \emptyset$.
- 3: **while** not stopped **do**
- 4: GENERATE TREE(u, b).


```

5:   if  $b = 0$  then
6:     Set  $b = \lceil u/2 \rceil$ .
7:     Set  $u = 0$ .
8:   else
9:      $b - -$ .
10:     $u + = 2$ .
11:   end if
12: end while
13: return  $\mathcal{C}$ .
14: end procedure
15: procedure GENERATE TREE( $u, b$ )
16:   Set  $\text{tree} = \text{new tree}$  with single node.
17:   GENERATE TREE REC( $\text{tree}, u, b$ ).
18: end procedure
19: procedure GENERATE TREE REC( $\text{tree}, u, b$ )
20:   if number of unary nodes =  $u$  and number of
     binary nodes =  $b$  then
21:     GENERATE LABELED TREE( $\text{tree}$ ).
22:   else
23:     for all nodes  $v$  on the second-deepest level
       that have at most 1 child and have no nodes
       at the same level to their right with at least 1
       child do
24:       Add child to  $v$ .
25:       GENERATE TREE REC( $\text{tree}, u, b$ ).
26:       Remove that child from  $v$ .
27:     end for
28:     for all nodes  $v$  on the deepest level do
29:       Add child to  $v$ .
30:       GENERATE TREE REC( $\text{tree}, u, b$ ).
31:       Remove that child from  $v$ .
32:     end for
33:   end if
34: end procedure
35: procedure GENERATE LABELED TREE( $\text{tree}$ )
36:   Order the nodes in a suffix order.
37:   Recursively label each node in this ordered array
     with either an invariant, a unary operator, or a
     binary operator depending on its degree. For
     commutative binary operators we only label a
     vertex if its left child is larger than its right
     child.
38:   for each fully labeled tree do
39:     if the corresponding bound  $c$  is valid for all
       examples in  $E$  and is not dominated by exist-
       ing bounds in  $\mathcal{C}$ . then
40:       Set  $\mathcal{C} = \mathcal{C} \cup c$ .
41:       Remove dominated conjectures from  $\mathcal{C}$ .
42:     end if
43:   end for
44: end procedure

```

Figure 2 displays upper bounds in panel (a) and lower bounds in panel (b) derived for test instances for data generated based on a formula for gravity. The gray curves correspond to bounds, and each must be the best on at least one training example instance in order to be

retained. More details on this experiment are provided in Sections 3.1 and 4.1.

Algorithm 1 can be adapted for property conjecturing with few modifications. We now detail the differences. Let E be a set of examples, and let $\Pi = \{\pi_1, \pi_2, \dots, \pi_m\}$ be properties. The examples are n data observations, and the properties are m Boolean features. The truth value of example i for property π_j is $\pi_j(i)$. Let O be the following collection of operators: NOT (\neg), AND ($\&$), OR (\mid), XOR (exclusive or) (\oplus), and IMPLIES (\rightarrow). NOT is a unary operator, and the remaining operators are binary operators. Let $\pi^* \in \Pi$ be the property for which sufficient and/or necessary conditions are of interest, and let $\pi^*(i)$ be the truth value of the property of interest for example i .

The aim is to generate conjectured sufficient or necessary conditions for the property of interest that are valid for any realization of input examples E . The algorithm for property conjecturing **procedure** CONJECTURING-PROP generates unlabeled trees as in Algorithm 1 but then labels the nodes with operators and properties. Logical expressions satisfying the Dalmatian heuristic conditions are retained as conjectures \mathcal{C} . For a conjecture $c \in \mathcal{C}$, let $c(i)$ be the conjectured truth value for example i .

The inputs to property conjecturing are examples E , properties Π , operators O , a property of interest π^* , and a direction (*SUFFICIENT*, *NECESSARY*) indicating if the algorithm will produce sufficient or necessary conditions for the property of interest.

The stopping criterion for property conjecturing for the case that direction is *SUFFICIENT* is obtaining a set of conjectures where, for every example with $\pi^*(i) = \text{true}$, each example evaluates to true for at least one conjecture. Otherwise, expression generation continues until a time limit is reached.

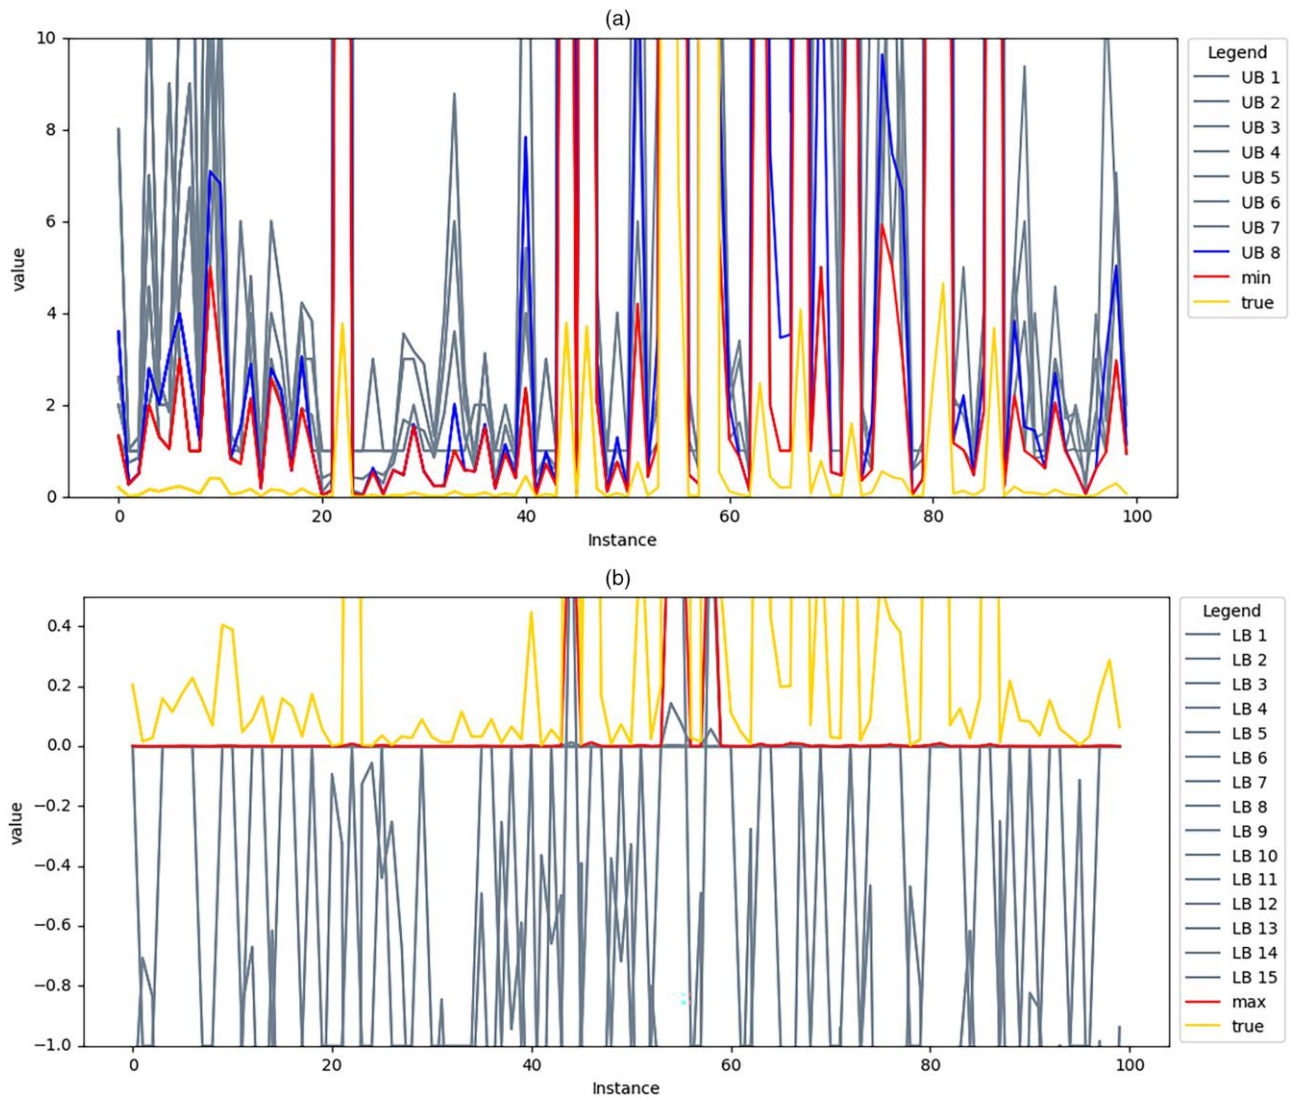
The Dalmatian heuristic for property conjectures is applied as follows. A conjectured sufficient condition c is only retained in the database of conjectures \mathcal{C} if the expression passes the following two tests:

1. (*Truth test*) For all examples $i \in E$ for which $c(i)$ is true, then $\pi^*(i)$ is also true, and
2. (*Nondominance test*.) The number of examples $i \in E$ for which $c(i)$ is true is not a subset of examples that evaluate to true for any previously conjectured sufficient condition.

To generate necessary conditions for π^* , one can generate sufficient conditions for $\neg\pi^*$ (NOT π^*).

Figure 3(a) depicts candidate conditions for examples with the property of interest (green) and those without (red). Conditions 1 and 2 are sufficient conditions for subsets of examples with the property of interest. Condition 3 evaluates to true for examples with and without the property of interest and would therefore not be retained. The goal of property conjecturing is to find a set of sufficient conditions that evaluate to true for all

Figure 2. (Color online) Invariant Conjecturing



Notes. (a) Upper bounds and (b) lower bounds generated for gravitational force using CONJECTURING-INV, the invariant version of the conjecturing algorithm. Instances from the training data are on the x axis.


examples with the property of interest and for none of the examples without the property of interest, as illustrated in Figure 3(b).

2.3. Other Related Work

In this section, we explain how our work is related to previous work in automated scientific discovery, machine learning interpretability, automated feature engineering, and empirical model building.

Symbolic regression has been used as a tool for automated scientific discovery. Symbolic regression is the use of genetic programming (GP) to approximate a target function on training data and generalize to produce predictions on new data (Nicolau and Agapitos 2021). Until the work of Schmidt and Lipson (2009), the focus was on improving prediction accuracy by approximating an

underlying function rather than a focus on discovering true functional relationships among features. Schmidt and Lipson (2009) extend previous work to develop a system for discovering laws for dynamical systems by considering relationships among derivatives. Their work led to the development of a software, EUREQA. More recently, Udrescu and Tegmark (2020) combined a variety of strategies, including dimensional analysis, symmetry identification, neural network training, and brute-force enumeration, into a framework called AI FEYNMAN to recover true physical functional forms from data. Petersen et al. (2021) propose a method for deep symbolic regression (DSR) that combines reinforcement learning with a recurrent neural network model. They compare their approach with methods based on priority queue training (PQT) proposed by Abolafia et al. (2018)

Figure 3. (Color online) Schematic of Property Conjecturing

Notes. In panel (a), conditions 1 and 2 evaluate to true for a subset of examples with the property of interest and for no samples that do not have the property of interest. Condition 3 evaluates to true for examples with and without the property of interest, and so, it is discarded. In panel (b), the union of sufficient conditions covers all examples with the property of interest.

and the traditional genetic programming approach. Our experiments include comparisons with each of these methods in the ability to recover equations from data.

Other frameworks have been proposed to computationally generate conjectures from data and discover scientific laws. Data smashing is introduced by Chattopadhyay and Lipson (2014) as a method for computing dissimilarities from streams of data (e.g., electroencephalogram data) to aid in revealing relationships among observations. Jantzen (2016) proposes an algorithm with the similar purpose of detecting types of dynamical systems called *dynamical kinds*. Subsequently, these kinds “are then targets for law-like generalization” (Jantzen 2016). Whereas Jantzen’s work

provides a method for discovering the kinds, it does not suggest how to recover the “laws.” It is these relationships that we aim to discover with the CONJECTURING framework.

Our work is distinguished from these previous works in that (1) we focus on generating bounds for invariants that serve as hypotheses for the investigator rather than recovering true functional forms or generating accurate predictions, (2) our invariant conjecturing algorithm is paired with a property conjecturing algorithm for discovering both nonlinear bounds and Boolean relationships, (3) our framework is designed for a given static observational data set rather than on discovering laws for dynamical systems, and (4) rather than a stochastic search

over the space of functional forms, our CONJECTURING system leverages sophisticated techniques for enumerating expressions of increasing complexity (described in Larson and Van Cleemput (2016) for “noiseless” data involving mathematical objects such as graphs). In our system, the human remains “in the loop” to evaluate the plausibility of suggested bounds and conditions.

Brunton et al. (2016) introduce SINDy for combining sparse regression and expert knowledge to develop models of dynamical systems. We adopt a similar approach to incorporating prior knowledge in Section 5.4 for the Nguyen benchmark suite (Nguyen et al. 2011) where we provide the CONJECTURING framework with candidate nonlinear functions as building blocks. Unlike our CONJECTURING framework, theirs is designed for recovering equations governing dynamical systems rather than bounds, and theirs is not capable of recovering Boolean relationships.

Langley (2019) provides a review of past efforts in computational scientific discovery. Several frameworks have their origins in analyzing mass spectroscopy and other electrochemical data. Bacon (Langley et al. 1987) is a general framework for scientific discovery based on suggesting and executing a series of designed experiments. Tallorin et al. (2018) proposed a method called POOL that uses Bayesian optimization and machine learning in an iterative fashion for experiments to discover peptide substrates for enzymes. Bacon and POOL both make recommendations regarding additional data to collect, whereas our system assumes that a fixed data set is provided that may or may not be the result of a designed experiment.

Precise definitions of “explainability” and “interpretability” are still being developed (Lu et al. 2019, Fürnkranz et al. 2020, Vilone and Longo 2020), as research in the area has rapidly accelerated. According to the convention of Rudin (2019), explainability is concerned with post hoc analyses of black box models to create simple explanations of model behavior. Motivated by observed accuracies of deep learning models, work in this area includes identifying important features for prediction, building simple local models, conducting sensitivity analyses, and deriving prototype examples (Samek and Müller 2019, Elton 2020). Tsang et al. (2018a, b, 2020) develop neural network frameworks for identifying sets of features for which there is an *interaction*—a nonadditive relationship among predictive features that influence a response value. These methods provide explainability in that they identify sets of features that interact, but the framework is not designed to reveal the functional form of the nonlinear interaction.

Rudin (2019) advocates the development of interpretable models where the mechanisms for predictions are simple relationships that are readily apparent to the investigator. Much of the recent work in this area is in

the development of decision rules (e.g., Hammer and Bonates 2006, Dash et al. 2018, Bellomarini et al. 2020) or decision lists and trees (e.g., Wang and Rudin 2015, Bertsimas and Dunn 2017, Wang et al. 2017, Rudin and Ertekin 2018, Verwer and Zhang 2019, Aghaei et al. 2021, Blanquero et al. 2021, Lemadjeng et al. 2023). Different from these works, our CONJECTURING framework automates the discovery of nonlinear features. In addition, as with work on decision rules in general, our framework can combine the discrete features in data with the discovered nonlinear features to discover a potentially richer set of Boolean relationships when compared with optimization-based trees and decision lists.

Khurana et al. (2018) propose a system that leverages reinforcement learning to search expression trees for predictive features. EXPLOREKIT (Katz et al. 2016) is a framework for automatic feature engineering that combines features using basic arithmetic operations and then uses machine learning to rank and select those with high predictive ability. THE DATA SCIENCE MACHINE (Kanter and Veeramachaneni 2015) automatically generates features for entities in relational databases with possible dependencies between tables followed by singular value decomposition. In none of these works is model transparency evaluated but, rather, only model performance. An important distinction of our work from these is that they focused on improving prediction accuracy, sometimes at the expense of understandable features, and not on scientific discovery.

Traditional statistical methods for empirical model building (e.g., regression analysis) tend to focus on first- and second-order polynomial models; interaction terms, up to a certain degree, are often included. Empirical models are intended to provide adequate prediction performance while also providing a simple assessment of feature importance via model coefficients. Techniques such as all-subsets, stepwise selection, and regularization methods (e.g., LASSO (Tibshirani 1996)) are commonly used to perform feature selection over model spaces of increasing complexity. However, domain knowledge is typically required for reciprocal or nonpolynomial relationships. Our CONJECTURING framework provides a search over a much broader class of nonlinear functions.

3. Two Motivating Examples

In this section, we describe two data sets where a “typical” knowledge discovery workflow fails to reveal important relationships among features.

Research on machine learning does, of course, lead to conjectured relationships between variables, which are, in turn, used to make predictions of one or more variables in terms of others. A trained neural net, for instance, can be viewed as a black box representing a function that produces an output for every input in its

domain. These functions are complex and of a different character than classical scientific laws: in particular, there is little hope of deriving these functions or relationships from simpler existing laws. Our CONJECTURING framework aims to help fill this gap in current capabilities.

3.1. Discovering Gravity

In this example, a numeric invariant of interest is determined by a more complex nonlinear relationship with three numeric predictors. Consider measurements that include the masses of two objects, m_1 and m_2 ; their distance, r ; and the gravitational force between them, F . The goal is to recover the dependence of F on m_1 , m_2 , and r , or

$$F = k \frac{m_1 m_2}{r^2},$$

where k is the gravitational constant. Following the demonstration by Langely et al. (1987), we create a fictional data set using a predefined value for k that is a random number between zero and one. For our illustrative example, we generated 1,000 training data points and 1,000 test data points with $k=0.057098$. The choice of k was randomly generated. Values for m_1 , m_2 , and r are samples from Uniform (1,100,000) distributions, and F is calculated for each sample with no noise.

A linear regression model will fail to capture the nonlinear interaction of the variables. Off-the-shelf machine learning methods such as random forests and neural networks can leverage the nonlinear relationship in the data but cannot present the relationship to the investigator. In the next section, we propose a framework for producing bounds on F that are functions of the other features.

3.2. Discovering an Interaction in Real Estate Valuation Data

The second example is a case where a Boolean variable of interest is almost completely determined by the product of two numeric features in the data set; that is, the second-order *interaction* term completely defines the relationship.

Consider a data set on residential real estate properties for sale obtained from <https://www.redfin.com>. The goal is to recover a relationship between price (above or below \$300,000), price per square foot, and total square footage.

This data set includes both the price per square foot and total square footage along with eight additional features such as the number of bathrooms and bedrooms. The property of interest (above versus below) can be determined (with some rounding error) by multiplying the price per square foot by square footage and setting a threshold. Thus, the interaction of price per square foot and square footage, hereafter called the *active interaction*, almost completely describes the relationship between the predictors and response. Data are partitioned into a

training data set with 1,000 houses and a test data set with 30,156 houses. In the next section, we leverage our framework for invariant bounds and then extend it to produce Boolean relationships to discover the active interaction term and how it is related to class membership.

Standard machine learning methods are able to achieve high rates of prediction accuracy, and some can identify the terms of the active interaction term as important with this data, but to our knowledge, none can help the investigator discover that the terms should be multiplied.

4. A Conjecturing Framework for Discovering Patterns in Data

We now describe a framework that leverages a conjecturing algorithm to discover nonlinear and Boolean feature relationships in data. All experiments were run on a computer with an Intel i7-2600 CPU @ 3.4GHz and 16 GB RAM.

4.1. Conjecturing for Nonlinear Relationships

The invariant version of the conjecturing method (**procedure** CONJECTURING-INV) can be used for discovering nonlinear relationships in data. Invariant conjectures are generated that provide upper and lower bounds on the invariant of interest. These conjectures are the nonlinear functions that can be used as new features and/or as a complete model for the system.

For the gravity case from Section 3.1, the invariants are $A = \{F, m_1, m_2, r\}$, and the invariant of interest is the force F . The examples, E , are the observations in the data.

The CONJECTURING framework is not designed to recover constants such as the gravitational constant k . In general, for a functional relationship with a constant k such that $0 < k < 1$, the expression without the constant provides an upper bound for the response. In cases where the constant is larger than one, the expression without the constant provides a lower bound.

For our example, CONJECTURING-INV returns 19 upper bounds and 24 lower bounds for F . Among the upper bounds is

$$F \leq m_1 m_2 / r^2,$$

which approximates the true gravity relationship used to generate the data. The bound does not include the constant k . Other bounds generated by CONJECTURING-INV include

$$F \leq 2m_2 / \sqrt{r}, \quad (1)$$

$$F \leq 2|m_1 - m_2|, \quad (2)$$

$$F \geq 8m_2 / r^2, \quad (3)$$

$$F \geq -1 / (r - 2m_2). \quad (4)$$

For ease in visualization, eight of the upper bounds and 15 of lower bounds for F were selected and are depicted

in Figure 2. The upper bound, $m_1 m_2 / r^2$, in Figure 2(a) is blue, whereas the true value, $km_1 m_2 / r^2$, is gold.

As the primary goal of our approach is discovery, the bounds produced are suggestions that require further validation. We consider it a success that the relationship $F \propto m_1 m_2 / r^2$ is included in one of the bounds. Among the other bounds produced, we see that true relationships $F \propto m_1$, $F \propto m_2$, and $F \propto 1/r^2$ are all suggested, along with false relationships $F \propto 1/\sqrt{r}$ and $F \propto 1/r$. Follow-up investigations can be used to inspect these relationships and potentially recover the gravitational constant. An approximation of the gravitational constant of 0.057 could be represented as $(+1 + 1 + 1 + 1 + 1) \times 10^{-1-1} + (+1 + 1 + 1 + 1 + 1 + 1 + 1) \times 10^{-1-1-1}$, which, by itself has complexity 22. The expression $m_1 m_2 / r^2$ has complexity six.

In this example, the CONJECTURING framework recovers the true nonlinear relationship up to a constant of proportionality along with 42 additional suggested bounds. Therefore, isolating a single true bound, in the case where the bound is unknown, can require additional analysis and/or experiments. The additional bounds can provide potential insight into feature interactions.

To investigate the potential dependence of results on the gravitational constant k , we conduct an additional experiment where for each of 10 replications, a different constant k is sampled from a Uniform(0,1) distribution. For each value of k , CONJECTURING-INV recovers $m_1 m_2 / r^2$ as a conjectured upper bound for F .

In another experiment, values for m_1 , m_2 , and r are sampled from Uniform(1,100) distributions, and the force F is calculated for each sample. Ten data sets with noise are created by sampling from a normal distribution with mean zero and standard deviation 10^{-t} multiplied by the root mean square of the F values in the training data for $t = 9, 8, \dots, 0$ so that the noise increases as t decreases. The root mean square of F is used so that the standard deviation of the noise distribution is adjusted to the scale of the target invariant as in Section 5.3. Noise is added to each calculated value for F . For $t \geq 4$, CONJECTURING-INV is able to recover the true nonlinear relationship up to a constant of proportionality, among other bounds.

4.2. Conjecturing for Nonlinear and Boolean Relationships with Mixed Data

Our CONJECTURING framework for mixed data leverages the invariant version (**procedure** CONJECTURING-INV) and the property version (**procedure** CONJECTURING-PROP) of the conjecturing algorithm. For mixed data, we propose a framework to produce conjectures of nonlinear and Boolean patterns. These conjectures can capture complex patterns while maintaining interpretability.

We assume that we are given a data set with numeric features N , Boolean features B , and a categorical feature of interest with levels \mathcal{Y} . Note that a categorical feature

with more than two levels can be converted to a series of Boolean features. Let π_y be the property that an observation has value y for $y \in \mathcal{Y}$.

For each level $y \in \mathcal{Y}$, the algorithm discovers bounds for the numeric features that are satisfied by each observation in the class (Algorithm 2, lines 4–14). These inequalities are converted to properties of the form “if the inequality is satisfied, then true; false, otherwise” (Algorithm 2, line 12). These new properties are combined with the original Boolean features in the data (Algorithm 2, line 13). The properties from across all classes are pooled together, the observations belonging to all classes are pooled together as examples, and then, for each level $y \in \mathcal{Y}$, the property version of conjecturing is applied to discover sufficient conditions for π_y (Algorithm 2, lines 15–20).

We now provide further details on Algorithm 2 using the real estate valuation case from Section 3.2 as an illustrative example. First, we convert the categorical feature *propertyType* into Boolean features *condo*, *mobileHome*, *singleFamily*, *townhouse*, *multiFamily2-4Unit*, *multifFamily5PlusUnit*, and *Other*. We also add a feature that is a constant value of 300,000 for each observation because it is the price cutoff, and we call it *300K*. The resulting 18 features are partitioned into numeric features $N = \{ \textit{bedrooms}, \textit{bathrooms}, \textit{squareFootage}, \textit{lotSize}, \textit{yearBuilt}, \textit{daysOnMarket}, \textit{pricePerSquareFoot}, \textit{hoaPerMonth}, \textit{latitude}, \textit{longitude}, \textit{and } 300K \}$ (Algorithm 2, line 2) and Boolean features $B = \{ \textit{condo}, \textit{mobileHome}, \textit{singleFamily}, \textit{townhouse}, \textit{multiFamily2-4Unit}, \textit{multifFamily5PlusUnit}, \textit{Other} \}$ (Algorithm 2, line 3).

In our training set, there are 1,000 observations that are used as examples. For each value of the property of interest, $\{ \textit{below}, \textit{above} \}$, the corresponding observations serve as the examples (Algorithm 2, line 6). For each numeric feature, upper and lower bounds on that feature are found that are functions of the other numeric features (Algorithm 2, lines 8 and 9). These are found by applying the invariant relations version of the conjecturing method (CONJECTURING-INV). For houses with property *below*, there are 1,280 bounds derived. Included are plausible relations concerning house features that are seemingly irrelevant to the classification task, such as

$$\textit{bathrooms} \leq 2 \times \textit{bedrooms} \tag{5}$$

$$\textit{bedrooms} \geq \textit{bathrooms} - 1 \tag{6}$$

$$\textit{lotSize} \geq (\textit{squareFootage} - \textit{yearBuilt}) \times \textit{bedrooms}. \tag{7}$$

Also included are less interpretable bounds, such as

$$\textit{yearBuilt} \geq \textit{hoaPerMonth} \times \log(10) / \log(2 \times \textit{daysOnMarket}) \tag{8}$$

$$\textit{daysOnMarket} \leq e^{\sqrt{2 \times \textit{lotSize}}} \tag{9}$$

$$\textit{hoaPerMonth} \leq 10^{2 \times \textit{bathrooms}} + \textit{squareFootage}. \tag{10}$$

There are also several bounds discovered that are close approximations of the relationship present in the active

interaction term, including

$$\text{squareFootage} \leq 300K/\text{pricePerSquareFoot} + \text{bathrooms} \quad (11)$$

$$\text{squareFootage} \leq 300K/\text{pricePerSquareFoot} + \text{bedrooms} \quad (12)$$

$$\text{squareFootage} \leq 300K/\text{pricePerSquareFoot} + \text{daysOnMarket} \quad (13)$$

$$\text{squareFootage} \leq 300K/(\text{pricePerSquareFoot} - 1) - 1 \quad (14)$$

$$\text{pricePerSquareFoot} \leq -300K/(\text{bedrooms} - \text{squareFootage}) \quad (15)$$

$$\text{pricePerSquareFoot} \leq \lceil 300K/\text{squareFootage} \rceil \quad (16)$$

$$300K \geq -(\text{bathrooms} - \text{squareFootage}) \times \text{pricePerSquareFoot}. \quad (17)$$

For houses with property *above*, there are 1,457 bounds derived, including a mix of simple relations and less intuitive relations. Also included are the following three relations that are nearly identical to the active interaction relation:

$$\text{squareFootage} \geq 300K/(\text{pricePerSquareFoot} + 1) \quad (18)$$

$$\text{pricePerSquareFoot} \geq 300K/\text{squareFootage} + 1 \quad (19)$$

$$300K \leq (\text{pricePerSquareFoot} + 1) \times \text{squareFootage}. \quad (20)$$

The resulting invariant relations are pooled together (Algorithm 2, line 10). The invariant relations are encoded as properties (Algorithm 2, line 12). The original binary features from the data are also encoded as properties for a total of $1,280 + 1,457 + 7 = 2,744$ properties. Examples of encoded properties from the invariant relations are

$$\text{bathrooms} \stackrel{?}{\leq} 2 \times \text{bedrooms} \quad (21)$$

$$(\text{yearBuilt} \stackrel{?}{\geq} \text{hoaPerMonth} \times \log(10)/\log(2 \times \text{daysOnMarket})) \quad (22)$$

$$(\text{squareFootage} \stackrel{?}{\leq} 300K/\text{pricePerSquareFoot} + \text{bathrooms}) \quad (23)$$

$$(\text{squareFootage} \stackrel{?}{\geq} 300K/(\text{pricePerSquareFoot} + 1)). \quad (24)$$

These properties can be used as Boolean features that indicate whether a nonlinear relationship among numeric features is satisfied for an observation.

The properties generated for each level $\{\textit{below}, \textit{above}\}$ are collected in a set Π along with π_y and the seven original Boolean features (Algorithm 2, line 13).

For each level $\{\textit{below}, \textit{above}\}$, apply the property version of conjecturing to the properties Π with the training data observations serving as the examples E and level as the property of interest (Algorithm 2, lines 15–20). The result is a set of properties that are sufficient conditions for the levels.

CONJECTURING-PROP returns only two properties. They both approximate the underlying active interaction.

$$\text{bathrooms} \geq -300K/\text{pricePerSquareFoot} + \text{squareFootage} \rightarrow \textit{below} \quad (25)$$

$$\text{squareFootage} \geq (300K + 1)/(\text{pricePerSquareFoot} - 1) \rightarrow \textit{above}. \quad (26)$$

An inspection of the data reveals that for some of the houses, there is some rounding error when comparing the price to the square footage multiplied by the price per square foot. The conjecturing algorithm compensates by using invariants as error terms. In the first property, the error term is $\text{bathrooms} \times \text{pricePerSquareFoot}$. In the second property, the error term is $\text{squareFootage} + 1$.

When these properties are applied as classification rules for predicting whether a house will be above or below \$300,000, they produce no error on the training data. The first property misclassifies 37 of 30,156 houses in the test data for an accuracy of 0.999. The second property misclassifies 26 houses. The misclassified houses are because of rounding error and miscoding of data. For example, one house in the test data is listed as having 31,248 bathrooms, and another is listed as having a price of \$459. Despite the noise and rounding error in the data, the CONJECTURING framework helped to discover the active interaction term.

Algorithm 2 (Conjecturing Framework for Nonlinear and Boolean Relationships with Mixed Data)

Input: Data observations $\{1, \dots, n\}$ with numeric features N , Boolean features B , and a categorical feature of interest with levels \mathcal{Y} ; a set of invariant operators O and a set of property operators P .

Output: A set of conjectured properties \mathcal{P} .

- 1: Set $\mathcal{P} = \emptyset$. /* Initialize properties set. */
- 2: Set $A = \{\alpha_j : j \in N\}$. /* Define the set of invariants to be the original numeric features in the data. */
- 3: Set $\Pi = \{\pi_j : j \in B\}$. /* Define the set of properties to be the original Boolean features in the data. */
- 4: **for** $y \in \mathcal{Y}$ **do** /* Loop on the levels of the categorical feature of interest */
 - 5: Set $\mathcal{R} = \emptyset$. /* Initialize invariant relations set. */
 - 6: Set $E = \{i : \pi_y\}$. /* Define the set of observations with level y as the examples. */
 - 7: **for** $j \in N$ **do** /* Loop on original numeric features. */
 - 8: Set $R_U = \text{CONJECTURING-INV}(E, A, O, \alpha_j, \text{UPPER})$ /* Submit examples, invariants, and the invariant of interest to the invariant version of CONJECTURING for upper bounds. */
 - 9: Set $R_L = \text{CONJECTURING-INV}(E, A, O, \alpha_j, \text{LOWER})$ /* Submit examples, invariants, and the invariant


```

of interest to the invariant version of CONJECTURING
for lower bounds. */
10: Set  $\mathcal{R} = \mathcal{R} \cup R_U \cup R_L$ .
11: end for
12: Convert the new invariant relations  $\mathcal{R}$  into
properties  $\Pi_{\mathcal{R}}$ .
13: Set  $\Pi = \Pi \cup \pi_y \cup \Pi_{\mathcal{R}}$ . /* Define the set of
properties to be the original Boolean features, the
level  $y$ , and the invariant relations properties. */
14: end for
15: for  $y \in \mathcal{Y}$  do /* Loop again on the levels of the
categorical feature of interest. */
16: Set  $E = \{1, \dots, n\}$ . /* Use all examples. */
17: Set  $\mathcal{P}_S = \text{CONJECTURING-PROP}(E, \Pi, P, \pi_y,$ 
SUFFICIENT). /* Submit examples, proper-
ties, and the level  $y$  as the property of interest to
the property version of CONJECTURING for suffi-
cient conditions. */
18: Set  $\mathcal{P}_N = \text{CONJECTURING-PROP}(E, \Pi, P, \pi_y,$ 
NECESSARY). /* Submit examples, proper-
ties, and the level  $y$  as the property of interest to
the property version of CONJECTURING for neces-
sary conditions. */
19: Set  $\mathcal{P} = \mathcal{P} \cup \mathcal{P}_S \cup \mathcal{P}_N$ .
20: end for
21: return  $\mathcal{P}$ .
    
```

Hereafter, we use “CONJECTURING framework” to imply:

1. In the case that all features are numeric, apply **procedure** CONJECTURING-INV.
2. In the case that all features are categorical, convert the features to a series of properties (Boolean features) and apply **procedure** CONJECTURING-PROP.
3. In the case of mixed data, apply Algorithm 2.

If there is no invariant of interest or no property of interest, each invariant and/or property can serve as the invariant/property of interest in turn, and conjectures can be generated for each.

5. Additional Computational Experiments

5.1. Sensitivity to the Number of Features

To investigate the impact of the number of features on the performance of the CONJECTURING framework, we conduct experiments adding noisy features to the gravity example. We use the same experimental setup as described in Section 4.1, including a time limit of five seconds. For $\ell = 0, \dots, 10$, we add ℓ noise invariants generated from a standard normal distribution and check

1. whether CONJECTURING-INV recovers $m_1 m_2 / r^2$,
2. the number of conjectures produced,
3. the number of expressions evaluated, and
4. the number of valid expressions produced. Valid expressions are bounds that are valid for all training examples.

For $\ell = 0, \dots, 6$, CONJECTURING-INV recovers $m_1 m_2 / r^2$ as a conjectured upper bound, and for $\ell = 7, \dots, 10$, the

bound is not recovered in the five-second time limit because of the additional noise invariants. The number of original invariants is five, including the force F , so in this experiment, we can add more than 100% additional invariants and still recover the true proportionality relationship $F \propto m_1 m_2 / r^2$.

Figure 4 contains plots of the number of conjectures produced, the number of expressions evaluated, and the number of valid expressions produced within the five-second time limit. The number of conjectures produced and expressions evaluated increases as the number of columns increases and tends to be larger for lower bounds than upper bounds. The number of conjectures produced ranges from 22 to 174. The number of valid expressions fluctuates between 75,000 and 135,000, and there is no discernible pattern effect of the number of noise invariants ℓ . As ℓ increases, there are more invariants available, and the number of low-complexity expressions increases exponentially, as does the number of low-complexity expressions comprised of the noise invariants. Low-complexity expressions can be generated and checked more quickly, which is why the number of expressions evaluated increases with ℓ . The number of expressions evaluated in five seconds ranges between about 300,000 and 1.4 million.

5.2. Sensitivity to Training Examples


To investigate the effect of different subsets of training examples on the ability of the CONJECTURING framework to recover true relationships, we apply the framework to the real estate experiment with 10 random samples of 1,000 training examples. We use the same experimental setup as described in Section 2, including a time limit of five seconds.

Table 1 contains the conjectures produced by the CONJECTURING framework for each of the 10 replications. As in the experiment in Section 4.2, the CONJECTURING framework makes use of invariants and operators to compensate for rounding error. The invariants employed as tolerances are *bathrooms* and *longitude*. The framework also employs operators $+1$, -1 , and $[\cdot]$ to account for deviations from the underlying active interaction. Each of the bounds can be rewritten in terms of *squareFootage* \times *pricePerSquareFoot* $-300K$ plus or minus a small error term containing, at most, one additional invariant. Therefore, we see that, in this instance, the method is not sensitive to the choice of training examples. Further experiments are needed to understand how well the CONJECTURING framework can recover underlying relationships in the presence of different kinds of noise. This is the subject of future work.

5.3. Comparison with AI FEYNMAN (Udrescu and Tegmark 2020)

In this section, we compare the ability of CONJECTURING to recover equations from data sets used by Udrescu and

Figure 4. The Number of Conjectures Produced, Expressions Tested, and Valid Expressions Produced as a Function of the Number of Noise Invariants Added to the Gravity Experiment



Tegmark (2020) with their algorithm AI FEYNMAN. We then apply the implementation of AI FEYNMAN to the gravity and real estate data sets described in Section 3. We note that the primary goal of CONJECTURING is for discovery of nonlinear and Boolean relationships, whereas the primary goal of AI FEYNMAN is recovery of equations.

5.3.1. Performance on Feynman Equations. We apply CONJECTURING to the first 10 equations listed in table 4 of Udrescu and Tegmark (2020) to draw comparisons based on solution time and noise tolerance. We used the data published by the authors.⁴ As in Udrescu and Tegmark (2020), for each instance, we apply CONJECTURING with three subsets of operators in turn: $\{+, -, \times, \div, +1, -1, ^2, \sqrt{\cdot}\}$; $\{+, -, \times, \div, +1, -1, ^2, \sqrt{\cdot}, \sin, \cos, \ln, ^{-1}, e\}$; and $\{+, -, \times, \div, +1, -1, ^2, \sqrt{\cdot}, \sin, \cos, \ln, ^{-1}, e, |\cdot|, \sin^{-1}, \tan^{-1}\}$. For instances where an equation includes the constant π , we include π as a constant invariant. For each instance, we use the first 10 samples in each data set and run CONJECTURING for 7,200 seconds, which is the same

time limit used for AI FEYNMAN and EUREQA reported in Udrescu and Tegmark (2020). We also run CONJECTURING for the noise tolerance of and time required by AI FEYNMAN to recover the equations as reported in table 4 of Udrescu and Tegmark (2020).

Table 2 and Table A.1 in the appendix contain the results of applying CONJECTURING to the data sets. CONJECTURING produces bounds that match the equation for five of the 10 instances. CONJECTURING finds a match for all equations with complexity 10 or less, matches one equation with complexity 11, and is unable to find a match for equations with higher complexity. Udrescu and Tegmark (2020) report that AI FEYNMAN resolves all of the equations, whereas EUREQA (Schmidt and Lipson 2009) resolves four of the 10 equations, two of which are different from those found by CONJECTURING. These results indicate that CONJECTURING is well suited for recovering equations of complexity 10 or less within 7,200 seconds. Higher-complexity formulas with more invariants require additional time.

Table 1. Sufficient Conditions Produced by CONJECTURING for Different Training Set Samples for the Real Estate Valuation Data

| Sample | Conjectures produced |
|--------|---|
| 1 | $bathrooms \geq -300K/pricePerSquareFoot + squareFootage \rightarrow below$ $squareFootage \geq (300K + 1)/(pricePerSquareFoot - 1) \rightarrow above$ |
| 2 | $squareFootage \leq \lceil 300K/pricePerSquareFoot \rceil \rightarrow below$ $squareFootage \geq (300K + 1)/(pricePerSquareFoot - 1) \rightarrow above$ |
| 3 | $bathrooms \geq -300K/pricePerSquareFoot + squareFootage \rightarrow below$ $squareFootage \geq (300K + 1)/(pricePerSquareFoot - 1) \rightarrow above$ |
| 4 | $bathrooms \geq -300K/pricePerSquareFoot + squareFootage \rightarrow below$ $squareFootage \geq (300K + 1)/(pricePerSquareFoot - 1) \rightarrow above$ |
| 5 | $bathrooms \geq -300K/pricePerSquareFoot + squareFootage \rightarrow below$ $squareFootage \geq (300K + 1)/(pricePerSquareFoot - 1) \rightarrow above$ |
| 6 | $squareFootage \leq \lceil 300K/pricePerSquareFoot \rceil \rightarrow below$ $squareFootage \geq (300K + 1)/(pricePerSquareFoot - 1) \rightarrow above$ |
| 7 | $squareFootage \leq (300K + longitude)/pricePerSquareFoot \rightarrow below$ $squareFootage \geq (300K + 1)/(pricePerSquareFoot - 1) \rightarrow above$ |
| 8 | $bathrooms \geq -300K/pricePerSquareFoot + squareFootage \rightarrow below$ $squareFootage \geq (300K + 1)/(pricePerSquareFoot - 1) \rightarrow above$ |
| 9 | $squareFootage \leq (300K - daysOnMarket)/pricePerSquareFoot \rightarrow below$ $squareFootage \geq (300K + 1)/(pricePerSquareFoot - 1) \rightarrow above$ |
| 10 | $bathrooms \geq -300K/pricePerSquareFoot + squareFootage \rightarrow below$ $squareFootage \geq (300K + 1)/(pricePerSquareFoot - 1) \rightarrow above$ |

Equations I.6.2 and I.6.2.b in Table 2 each have a repeated invariant σ . As noted in Section 2.2, CONJECTURING does not allow repeated invariants, and so, these equations will not be recoverable as bounds. In Section 5.4, we describe ways to address this deficiency and recover equations such as I.6.2 and I.6.2.b.

The normalized root mean square error (NRMSE) calculated for 100 test examples for the best-performing conjecture on the training data based on mean absolute error. NRMSE is calculated as $1/\sigma_f$ multiplied by the root mean square error on the test examples, where σ_f is the standard deviation of the invariant of interest for the test examples. CONJECTURING produced conjectures with NRMSE less than 1.0 for eight of the 10 equations. Udrescu and Tegmark (2020) report exact recovery of all of the equations by AI FEYNMAN, so the NRMSE is zero.

Despite the fact that CONJECTURING is not designed for recovery of equations, we see that it can be successful in doing so for lower-complexity nonlinear equations.

Table A.1 in the appendix contains the results of applying CONJECTURING to the data sets when noise is added to the invariant of interest. The noise added to the invariant of interest is sampled from a normal distribution with mean zero and standard deviation equal to the number in the table multiplied by the root mean square of the invariant of interest in the training data. The noise is the noise level tolerated by AI FEYNMAN as reported by Udrescu and Tegmark (2020). The time is the time reported by Udrescu and Tegmark (2020) for AI FEYNMAN to recover the equation.

CONJECTURING is unable to achieve exact recovery of the equations with the introduction of noise. The NRMSE is less than 1.0 for six of the 10 equations and does not exceed 1.625. Noise that results in target values above or below the ground truth will violate exact upper or lower bounds that could be produced by the Dalmatian heuristic. Even so, CONJECTURING is able to produce good approximations of the equations.

Table 2. Results for CONJECTURING on Data Sets from Udrescu and Tegmark (2020)

| Instance | Equation | Number of invariants | Complexity | NRMSE | Recovered by EUREQA? |
|----------|--|----------------------|------------|-------|----------------------|
| I.6.2.a | $f = e^{-\theta^2/2}/\sqrt{2\pi}$ | 2 | 9 | 0.000 | No |
| I.6.2 | $f = e^{-\theta^2/2\sigma^2}/\sqrt{2\pi\sigma^2}$ | 3 | 13 | 0.553 | No |
| I.6.2.b | $f = e^{-(\theta-\theta_1)^2/2\sigma^2}/\sqrt{2\pi\sigma^2}$ | 5 | 16 | 1.511 | No |
| I.8.14 | $d = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}$ | 4 | 10 | 0.000 | No |
| I.9.18 | $\frac{Gm_1m_2}{(x_2-x_1)^2+(y_2-y_1)^2+(z_2-z_1)^2}$ | 9 | 17 | 1.211 | No |
| I.10.7 | $m = \frac{m_0}{\sqrt{1-\frac{v^2}{c^2}}}$ | 3 | 9 | 0.000 | No |
| I.11.19 | $A = x_1y_1 + x_2y_2 + x_3y_3$ | 6 | 11 | 0.696 | Yes |
| I.12.1 | $F = \mu N_n$ | 2 | 3 | 0.000 | Yes |
| I.12.2 | $F = \frac{q_1q_2}{4\pi\epsilon r^2}$ | 5 | 12 | 0.671 | Yes |
| I.12.4 | $E_f = \frac{q_1}{4\pi\epsilon r^2}$ | 4 | 11 | 0.000 | Yes |

5.3.2. Performance of AI Feynman (Udrescu and Tegmark 2020) on Gravity and Real Estate Examples. We apply the implementation of AI FEYNMAN available⁵ to the gravity example described in Section 3.1 and the real estate example described in Section 3.2. A difference between our gravity example and the data sets used by Udrescu and Tegmark (2020) is that for the gravity example, the gravitational constant is the same for every data point, but for the data sets used by Udrescu and Tegmark (2020), it is treated as a variable and is different for each point.

For the real estate data, we apply AI FEYNMAN to the data to attempt to discover the relationship between the property of interest and the input features. The original numeric features are supplied along with Boolean features corresponding to the levels of the *property-Type* feature. Note that AI FEYNMAN is designed for recovering numeric functions and is therefore not suitable for Boolean relationships such as those in the real estate example.

For both instances, the AI FEYNMAN implementation aborts with an error regarding an eigenvalue calculation. We suspect that the source of the failure in both cases may be because of the difference in treatment of constants. In our gravity example, the gravitational constant is the same for all points, whereas in analogous examples, Udrescu and Tegmark (2020) treat constants as variables and generate a unique value for each observation. Our practice of treating the gravitational constant as the same for all observations may be contributing to an error in matrix calculations for AI FEYNMAN. In the real estate example, each observation has a feature with the same value (the \$300,000 cutoff). This constant column in the data matrix could also be contributing to an error in matrix calculations for AI FEYNMAN. In the electronic companion, we include the code and output for AI FEYNMAN applied to (1) their example 1, demonstrating that our installation is functional, (2) our gravity example, including the error message, and (3) our real estate example,

including the error message. These examples show that the CONJECTURING framework can provide useful insights on examples where AI FEYNMAN cannot.

5.4. Experiments with the Nguyen Benchmark Suite (Nguyen et al. 2011)

We apply our invariant conjecturing implementation to the Nguyen benchmark suite (Nguyen et al. 2011) so as to draw comparisons with symbolic regression methods described by Petersen et al. (2021). The Nguyen benchmark suite is a set of 12 equations. As mentioned before, our CONJECTURING framework is designed for discovering nonlinear relationships in the form of bounds and Boolean relationships, whereas symbolic regression methods are designed to recover equations. In these experiments, we investigate the ability of invariant conjecturing to recover equations (or approximations) among the discovered bounds.

The benchmark equations are in Table 3. We generate an instance for each using the protocols described by Petersen et al. (2021). For each equation, 20 training examples and 20 test examples are generated. For Equations (1) through (6), x is sampled from a Uniform(-1,1) distribution; for Equation (7), x is sampled from a Uniform(0,2) distribution; for Equation (8), x is sampled from a Uniform(0,4) distribution; and for Equations (9) through (12), x and y are sampled from a Uniform(0,1) distribution. For each equation, we allow a time limit of 10,000 seconds for generating upper and lower bounds. The operators include unary operators sine, cosine, natural log, and natural exponential and binary operators addition, subtraction, multiplication, and division.

As noted in Section 2.2, our CONJECTURING framework does not allow repeated invariants in conjectures. Therefore, most of the equations in Table 3 cannot be recovered by our framework. For each equation, we first evaluate the ability of the conjectured bounds to approximate the equation by reporting the normalized root mean squared error as defined and reported by Petersen

Table 3. Results for CONJECTURING in the Nguyen Benchmark Suite (Nguyen et al. 2011)

| Instance | Equation | Without additional invariants | | With additional invariants | |
|----------|---------------------------------------|-------------------------------|-------|----------------------------|-------|
| | | Recovered? | NRMSE | Recovered? | NRMSE |
| 1 | $f = x^3 + x^2 + x$ | No | 0.94 | Yes | 0.00 |
| 2 | $f = x^4 + x^3 + x^2 + x$ | No | 0.86 | Yes | 0.00 |
| 3 | $f = x^5 + x^4 + x^3 + x^2 + x$ | No | 1.01 | Yes | 0.00 |
| 4 | $f = x^6 + x^5 + x^4 + x^3 + x^2 + x$ | No | 0.63 | Yes | 0.00 |
| 5 | $f = \sin(x^2) \cos(x) - 1$ | No | 0.22 | Yes | 0.00 |
| 6 | $f = \sin(x) + \sin(x + x^2)$ | No | 0.44 | Yes | 0.00 |
| 7 | $f = \log(x + 1) + \log(x^2 + 1)$ | No | 0.58 | Yes | 0.00 |
| 8 | $f = \sqrt{x}$ | No | 1.08 | Yes | 0.00 |
| 9 | $f = \sin(x) + \sin(y^2)$ | No | 1.03 | Yes | 0.00 |
| 10 | $f = 2 \sin(x) \cos(y)$ | No | 0.60 | Yes | 0.00 |
| 11 | $f = x^y$ | Yes | 0.00 | Yes | 0.00 |
| 12 | $f = x^4 - x^3 + \frac{1}{2}y^2 - y$ | No | 0.83 | Yes | 0.00 |

et al. (2021) and described in Section 5.3. We report NRMSE for the conjecture with the lowest mean average error for the training examples.

The results in Table 3 indicate that our CONJECTURING framework is able to recover only Equation (11), $f = x^y$. It is able to do so despite the fact that the exponent operator is not included. CONJECTURING produces the expression $e^{y \log(x)}$ and simplified it to x^y . The NRMSE values for the best bounds for the other instances range from 0.22 to 1.08. These values are larger than those reported for symbolic regression methods as reported in table 10 of Petersen et al. (2021); the methods include deep symbolic regression (Petersen et al. 2021), priority queue training (Abolafia et al. 2018), vanilla policy gradient (VPG) (Petersen et al. 2021), genetic programming, and a method implemented in Mathematica based on Markov chain Monte Carlo and nonlinear regression.

For Expressions (1) through (8), the only invariants are the invariant of interest f and the input invariant x . Because of the fact that no invariants can be repeated, the CONJECTURING framework is limited to the application of only unary operators to x , and no expressions with binary operators are produced. As an example, for the first equation, the best-conjectured lower bound is

$$f(x) \geq \sin(e^{\cos(\sin(\sin(\sin(\log(\sin(\sin(\cos(\cos(\sin(e^x))))))))))}).$$

We now describe how our CONJECTURING framework can be adapted to allow for repeated invariants and report results for the adapted method. To address repeated invariants, we can add additional invariants using commonly occurring functional forms. For Equations (1) through (8), we add invariants x^2 , x^3 , x^4 , x^5 , x^6 , and $\sin(x)$, $\cos(x)$, \sqrt{x} along with two copies of constant 1 and constant 2. Two copies of constant 1 are included because it appears twice in Equation (7). For Equations (9) through (12), we also add y^2 as an invariant. Recall that in our implementation, whereas invariants cannot repeat in a conjecture, operators can. Therefore, an alternative approach to addressing the constants is to include the unary operator of addition by one.

As shown in Table 3, the CONJECTURING framework is able to exactly recover each equation as a bound when the additional invariants are included so that the NRMSE values are 0.000 for all equations.

The practice of adding the invariants that are nonlinear functions of the original input might appear to be impractical. However, as suggested by Brunton et al. (2016), specifying these invariants can reflect expert knowledge on the system being investigated. They note that identifying candidate functions for SINDy “must be a coordinated effort to incorporate expert knowledge, feature extraction, and other advanced methods.” CONJECTURING offers distinct capabilities for discovery, as nonlinear functions can be specified as invariants or may

still be discovered so long as they do not involve repeated input invariants.

6. Application to COVID-19 Data

In this section, we demonstrate the CONJECTURING framework on synthetic patient-level COVID-19 data that were provided as part of the Veterans Health Administration (VHA) Innovation Ecosystem and precisionFDA COVID-19 Risk Factor Modeling Challenge.⁶ The data include synthetic veteran patient health records, including medical encounters, conditions, medications, and procedures. All subjects are located in Massachusetts.

The goal of the challenge was to better understand risk and protective factors for COVID-19 outcomes. Participants were asked to predict alive/deceased status. In our experiments, we focused on investigating outcomes for subjects who had COVID-19. Because our goal is to discover potential risk and protective factors, we evaluate the performance of CONJECTURING by checking the performance of the feature relationships on holdout test data rather than on prediction accuracy. Establishing the risk and protective factors as causal would require additional controlled experiments.

The purpose of this demonstration is to show that the CONJECTURING framework can suggest patterns in data that are potentially valuable and useful. We demonstrate that the patterns discovered are validated by the medical literature and, further, that they suggest functional forms for the relationships among invariants.

Predictions were based on information obtained through December 31, 2019. In the training data, we drop all information pertaining to events on or after January 1, 2020, and drop subjects who died before January 1, 2020. The prediction horizon is January 1, 2020, through May 31, 2020, and there are 5,568 patients with deceased status.

Table B.1 in the appendix includes definitions of new features that we generated for each patient. For each numeric observation, we created invariants for the mean and most recent value. For each reported allergy, device, immunization, procedure, and discretely measured observation, we create a property corresponding to each level. In total, we use 309 invariants and 362 properties. We use a training set consisting of 100 subjects from each outcome class (deceased/alive). We compare the results of applying CONJECTURING with classification and regression trees (CART) (Breiman et al. 1984), which is another interpretable method.

6.1. Results for CONJECTURING

Upper and lower bounds are generated for each invariant and for each outcome. These bounds, along with the 362 properties in the data, are used as properties for CONJECTURING-PROP. Conjectures are generated for both outcomes. The parameter *skips* is set to 90%. We use the

remaining 73,497 subjects as a test data set, thereby allowing us to ascertain the effects of potential overfitting to the 200 subjects used for training.

Among those with COVID-19 in the test data, 5,468 (8.0%) have a status of deceased, and 68,029 (92.0%) are alive. There are 38 conjectures for sufficient conditions for alive status and 40 conjectures for sufficient conditions for deceased status produced by the framework. Tables C.1–C.4 in the appendix contain the conjectures and evaluations.

Tables C.2 and C.4 in the appendix contain quantitative evaluations of the performance of the conjectures. Each table contains the precision, support, and lift of each conjecture. Note that each conjecture is a sufficient condition expressed as a conditional statement. The precision is the percentage of test examples for which the conditional statement evaluates to true among those for which the antecedent is true. Precision may be thought of as the “hit rate” of the conjecture. The support is the number of test examples for which the antecedent evaluates to true. The lift is the ratio of the precision to the proportion of examples for which the consequent is true. If the lift is greater than one, then the conjecture is better at identifying people for which the consequent is true than a random selection from the population.

Of the 38 conjectures for alive status, 22 (57.9%) have lift of at least 1.00. The lift ranges from 0.81 to 1.07; note that the maximum possible lift for a conjecture for alive status is 1.08 ($1/(68,029/73,497)$). Of the 40 conjectures for deceased status, 34 (85%) have lift at least 1.00. The lift ranges from 0.17 to 4.15.

Consider the sufficient conditions for deceased status in Table C.3 in the appendix. The conjecture with the highest precision and lift is

$$\textit{longitude} > -\textit{age} \times \textit{medicationsLifetimePercCovered} \\ \rightarrow \textit{Deceased},$$

and it has a lift of 4.15, meaning that a subject for which $\textit{longitude} > -\textit{age} \times \textit{medicationsLifetimePercCovered}$ is 4.15 times as likely to die as a randomly selected subject. The conjecture indicates that subjects in the east who are older and have a larger percentage of medications covered by the payer are at higher risk of death. The presence of *longitude* in the conjecture could be an indication of higher risk in population centers in the east such as Boston, or it could just serve the purpose of a number so that the relationship of the other invariants is satisfied. The range for *longitude* is $(-73.49, -69.92)$. We conduct a follow up *t*-test for a difference in mean *longitude* by outcome, and the null hypothesis is not rejected ($p=0.42$), suggesting that *longitude* is serving as a tolerance factor. The percentage of medications covered by the VHA is higher for subjects with more preexisting conditions and for those with more expensive medications because there is a low copay annual cap (currently

\$700).⁷ Further, the conjecture produces a suggestion of a functional form for the relationship between these factors. The conjecture confirms the CDC guidance that older subjects and those with more preexisting conditions are a higher risk of death from COVID.⁸

The conjecture with the second-highest precision and lift is

$$\textit{medicationsActive} > [\textit{hemoglobinA1cHemoglobinTotalInBlood}] \\ \rightarrow \textit{Deceased},$$

with a lift of 3.29. The condition includes the number of active medications and the ratio of hemoglobin A1c to total hemoglobin (an HbA1c test). The conjecture indicates that those with more active medications than the HbA1c percentage are at higher risk of death, which again agrees with the CDC guidance concerning preexisting conditions. Values for HbA1c are typically between 5.7% and 10.0%.⁹ It could be that the HbA1c is serving as a constant so that when the conjecture is true for a patient, then that patient will have at least five active medications, indicating an increased number of additional conditions.

The conjecture with the third-highest precision and lift is

$$\textit{age} > \textit{carbonDioxide} \times [\textit{potassium}] \rightarrow \textit{Deceased},$$

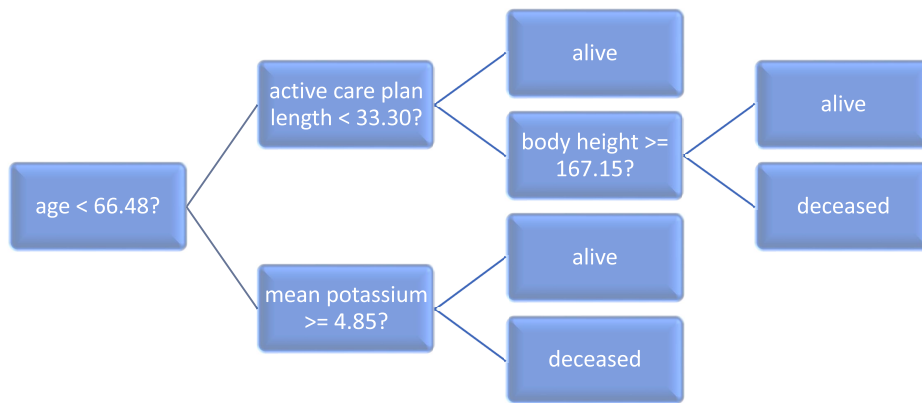
with a lift of 3.12. The conjecture suggests that older subjects with lower CO₂ levels and lower potassium are at higher risk of death. Lower CO₂ levels and abnormal potassium levels, particularly lower levels, have been independently studied and associated with COVID-19 morbidity and mortality (Hu et al. 2021, Noori et al. 2022). In addition to validating a role for these invariants, the conjecture suggests a potential nonlinear relationship among them.

For both outcomes, the CONJECTURING framework is able to generate new sufficient conditions that are true for the respective outcome at higher rates than would be expected for a patient selected at random. These results indicate that the conjecturing process is capturing relationships that hold across the population and are not merely reflective of the 200 training samples. In other words, overtraining appears to be mitigated. The discovered relationships, and the direct and indirect relationships that they indicate among features, are validated by CDC guidance and provide suggestions for deeper investigations into the functional form of the relationships and the extent of causality. The number of conjectures generated is not overwhelming for a human investigator to consider and further investigate.

6.2. Comparison with an Interpretable Model

We now consider the results of applying classification and regression trees (Breiman et al. 1984) to the COVID-19 data. A model is fit using the implementation in the R

Figure 5. (Color online) Tree Produced by CART for Predicting Alive/Deceased Status for COVID-19 Patients



Note. For each node, if the condition is satisfied, then the upper branch is taken.

library *rpart* (Therneau and Atkinson 2019). The tree produced by *rpart* is depicted in Figure 5.

Each leaf node corresponds to a sufficient condition for deceased status (True) or alive status (False). There are three sufficient conditions for alive status and two sufficient conditions for deceased status. Tables 4 and 5 contain the conditions and quantitative evaluations of the conditions produced by CART.

When comparing the results with those of CONJECTURING in Tables C.1–C.4 in the appendix, we note that CART produces conditions that are in conjunctive normal form where each clause consists of a single numeric bound, whereas CONJECTURING tends to leverage nonlinear relationships among invariants as the basis for conditions. In addition, we observe that CART (using the default control parameters):

1. produces many fewer conditions (3 versus 38 for alive status, 2 versus 40 for deceased status),
2. produces two conditions with much larger support in the test data than those produced by CONJECTURING (node 4 has 36,531, and node 7 has 12,443),
3. produces only two conditions that have lift greater than 1.0 (node 4 has lift 1.06, and node 7 has lift 2.85), and
4. does not produce conditions with better precision or lift than the best conditions produced by CONJECTURING.

We note that by using additional training data and adjusting the parameters for building the tree, we can obtain additional conditions that may produce conditions with higher precision and lift.

Both CART and CONJECTURING are able to leverage categorical variables for conditions, though CART does not do so for this training set. An example of such a condition is conjecture 27 in Table C.3 in the appendix.

Similar to many decision tree frameworks, CART leverages univariate bounds as component properties in its invariant clauses. CONJECTURING is unlikely to derive numeric bounds for individual features but instead produces more nonlinear relationships between invariants. Decision tree frameworks such as CART and CONJECTURING are complementary approaches for discovery of patterns among numeric and categorical features, but we see that CONJECTURING is capable of producing more complex yet interpretable relationships.

We did not compare CONJECTURING with association rule methods. Association rule methods for numerical data create bins using cutoff values in a manner similar to CART. Association rule methods might produce better-performing patterns but are unable to recover nonlinear relationships.

7. Conclusions

This work demonstrates that automated search for conjectured feature relations can support learning from data. The discovery of these kinds of feature relationships can also initiate new collaboration with domain scientists and lead to new scientific knowledge.

Our CONJECTURING framework is able to recover the functional form for gravity with only the measured force, masses, and distance. The framework is able to recover

Table 4. Conditions from CART for Alive/Deceased Status Among Those with COVID

| Node number | Sufficient condition |
|-------------|--|
| 4 | $age < 66.48$ and $activeCarePlanLength < 33.30 \rightarrow Alive$ |
| 6 | $age \geq 66.48$ and $meanPotassium \geq 4.85 \rightarrow Alive$ |
| 10 | $age < 66.48$ and $activeCarePlanLength \geq 33.30$ and $bodyHeight \geq 167.15 \rightarrow Alive$ |
| 7 | $age \geq 66.48$ and $meanPotassium < 4.85 \rightarrow Deceased$ |
| 11 | $age < 66.48$ and $activeCarePlanLength \geq 33.30$ and $bodyHeight < 167.15 \rightarrow Deceased$ |

Table 5. Evaluation of Conditions from CART Among Those with COVID

| Node number | Consequence | Precision (%) | Support | Lift |
|-------------|-----------------|---------------|---------|------|
| 4 | <i>Alive</i> | 98.28 | 36,531 | 1.06 |
| 6 | <i>Alive</i> | 81.69 | 2,070 | 0.88 |
| 10 | <i>Alive</i> | 91.57 | 5,371 | 0.99 |
| 7 | <i>Deceased</i> | 21.17 | 12,443 | 2.85 |
| 11 | <i>Deceased</i> | 5.37 | 3,746 | 0.72 |

the functional form for gravity when up to six additional noise features are added.

The framework also recovers an interaction between price per square foot, square footage, and price in real estate data. The interaction is recovered for each of 10 random samples of 1,000 training samples.

Using synthetic patient-level COVID-19 data, the framework produces conjectures that confirm CDC guidance and suggest new relationships regarding risk factors.

CONJECTURING provides a complementary functionality when compared with tree-based models. We demonstrate that CONJECTURING is able to produce conditions that involve nonlinear relationships among invariants, whereas CART produces conditions that are based on conjunctions of numeric bounds for individual invariants. Association rule methods that accommodate numerical data find rules based on ranges of individual invariants and would also fail to produce conditions that suggest nonlinear functional relationships among invariants.

In additional experiments, we observe that CONJECTURING can be adapted for other tasks such as symbolic regression. CONJECTURING is able to recover ground truth equations as bounds in some situations, but time limitations prevent recovery of more complex relationships. In our current implementation, repeated invariants cannot be produced. We demonstrate that a workaround is to add invariants for commonly occurring functional forms.

Even though CONJECTURING can be adapted for other tasks such as symbolic regression, it provides a fundamentally new functionality in terms of discovering patterns in data that symbolic regression methods, tree-based models, and association rule methods cannot provide. To the best of our knowledge, there is no other framework for discovering nonlinear relationships among features in the form of bounds and Boolean relationships involving these bounds and categorical features.

There is a trade-off in the proposed framework between the number of examples and the complexity of recovered expressions. With fewer examples, more candidate expressions can be checked in less time.

The current version of CONJECTURING requires that conjectures are true for every example. Future research will further investigate the ability of the framework to handle noisy data and investigate adjustments to the algorithm to

better handle noisy data such as generating conjectures that do not necessarily hold for all examples. If the handling of noise can be improved, then CONJECTURING may be able to be adapted to support predictive modeling efforts.

As indicated by the motivating examples from graph theory, it can be desirable to retain multiple bounds or conditions produced by conjecturing because different conjectures might be better descriptions of system behavior for certain examples and not others. In most of the experiments reported here, we evaluated conjectures as if they were each an entire model of the system and used standard ways of evaluating them. Additional metrics besides those considered in this work might also be helpful in evaluating conjectures. An interesting and important avenue of future research is how to better evaluate the ability of groups of conjectures to model system behavior.

If the CONJECTURING framework can provide functional relationships without constants of proportionality, the constant can be determined using regression with the original data. Suppose that the CONJECTURING framework indicates a relationship between the response y and predictors x of the form $y \leq b_1 f(x)$ for an unknown constant b_1 . A regression model can be fit of the form $\hat{y} = b_0 + b_1 f(x)$ using the data (x_i, y_i) , $i = 1, \dots, n$. The best strategy for determining constants of proportionality is another avenue for future research. Constants can provide unit consistency to conjectures. A method to search for unit consistent constants could facilitate the selection of meaningful conjectures.

Another area for potential research involves the so-called $p \gg n$ problem. That is, if the number of features is larger than the number of observations, then there are insufficient degrees of freedom to estimate a linear model with all p features or any more complex model. In such situations, feature and/or model selection tools are needed to search over potentially large model spaces. As desired model complexity increases (e.g., consideration of interaction terms), searching over such large model spaces can become computationally prohibitive. For instance, suppose an investigator seeks to identify a model by selecting the “best” subset from among 10 features and their associated 45 two-way interactions. In this example, simply considering models with only 10 variables requires searching over a model space larger than 2.9 billion. Future research will investigate the ability of the CONJECTURING framework to simplify model spaces and hence provide a mechanism for a more expeditious search of plausible models.

Acknowledgments

High-performance computing resources provided by the High Performance Research Computing (HPRC) Core Facility at Virginia Commonwealth University (<https://chipc.vcu.edu>) were used for conducting the research reported in this work.

Appendix A. Additional Results for CONJECTURING for Recovery of Equations

Table A.1. Additional Results for CONJECTURING on Data Sets from Udrescu and Tegmark (2020)

| Instance | Equation | Number of invariants | Complexity | Noise | Time (s) | NRMSE |
|----------|--|----------------------|------------|-----------|----------|-------|
| I.6.2.a | $f = e^{-\theta^2/2} / \sqrt{2\pi}$ | 2 | 9 | 10^{-2} | 16 | 0.154 |
| I.6.2 | $f = e^{-\theta^2/2\sigma^2} / \sqrt{2\pi\sigma^2}$ | 3 | 13 | 10^{-4} | 2,992 | 0.614 |
| I.6.2.b | $f = e^{-(\theta-\theta_1)^2/2\sigma^2} / \sqrt{2\pi\sigma^2}$ | 5 | 16 | 10^{-4} | 4,792 | 1.188 |
| I.8.14 | $d = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}$ | 4 | 10 | 10^{-4} | 544 | 1.323 |
| I.9.18 | $\frac{Gm_1m_2}{(x_2-x_1)^2+(y_2-y_1)^2+(z_2-z_1)^2}$ | 9 | 17 | 10^{-5} | 5,975 | 1.252 |
| I.10.7 | $m = \frac{m_0}{\sqrt{1-\frac{v^2}{c^2}}}$ | 3 | 9 | 10^{-4} | 14 | 0.034 |
| I.11.19 | $A = x_1y_1 + x_2y_2 + x_3y_3$ | 6 | 11 | 10^{-3} | 184 | 0.633 |
| I.12.1 | $F = \mu N_n$ | 2 | 3 | 10^{-3} | 12 | 0.002 |
| I.12.2 | $F = \frac{q_1q_2}{4\pi\epsilon r^2}$ | 5 | 12 | 10^{-2} | 17 | 0.885 |
| I.12.4 | $E_f = \frac{q_1}{4\pi\epsilon r^2}$ | 4 | 11 | 10^{-2} | 12 | 0.525 |

Appendix B. Features Generated for COVID-19 Data

Table B.1. Feature Definitions for COVID-19 Data

| Feature name | Definition |
|--|--|
| <i>healthcareExpenses</i> | The total lifetime cost of healthcare to the patient |
| <i>healthcareCoverage</i> | The total lifetime cost of healthcare services that were covered by payers |
| <i>latitude</i> | Latitude of patient’s home address |
| <i>longitude</i> | Longitude of patient’s home address |
| <i>age</i> | Current age of patient |
| <i>numAllergies</i> | Number of ongoing patient allergies |
| <i>activeCarePlans</i> | Number of current care plans |
| <i>lifetimeCarePlans</i> | Number of lifetime care plans |
| <i>activeCarePlanLength</i> | Length of time under current care plans |
| <i>lifetimeCarePlanLength</i> | Total lifetime length under care plans |
| <i>activeConditions</i> | Number of current health conditions |
| <i>lifetimeConditions</i> | Number of lifetime health conditions |
| <i>activeConditionLength</i> | Amount of time since current health condition(s) diagnosis |
| <i>lifetimeConditionLength</i> | Amount of time since first diagnosis of a health condition |
| <i>deviceLifetimeLength</i> | Total length of time using a medical device (e.g., pacemaker) |
| <i>encountersCount</i> | Total number of encounters with a healthcare professional |
| <i>encountersLifetimeTotalCost</i> | Total lifetime cost of healthcare encounters |
| <i>encountersLifetimeBaseCost</i> | Total lifetime cost of healthcare encounters, not including any line item costs related to medications, immunizations, procedures, or other services |
| <i>encountersLifetimePayerCoverage</i> | Total lifetime cost of healthcare encounters that were covered by payers |
| <i>encountersLifetimePercCovered</i> | Percentage of lifetime cost of healthcare encounters that were covered by payer |
| <i>imagingStudiesLifetime</i> | Number of lifetime imaging diagnostics (e.g., MRI) performed on patient |
| <i>immunizationsLifetime</i> | Number of lifetime immunizations received by patient |
| <i>immunizationsLifetimeCost</i> | Total lifetime cost of all immunizations received by patient |
| <i>medicationsLifetime</i> | Number of lifetime medications prescribed |
| <i>medicationsLifetimeCost</i> | Total lifetime cost of medications |
| <i>medicationsLifetimePercCovered</i> | Percentage of lifetime medication cost covered by payer |
| <i>medicationsLifetimeLength</i> | Total lifetime length on prescribed medications |
| <i>medicationsLifetimeDispenses</i> | Total lifetime number of prescription dispenses |
| <i>medicationsActive</i> | Number of current prescriptions |
| <i>proceduresLifetime</i> | Number of lifetime medical procedures (e.g., surgery) performed on patient |
| <i>proceduresLifetimeCost</i> | Total lifetime cost of all medical procedures performed on patient |

Appendix C. Results for CONJECTURING Applied to COVID Data

Table C.1. Conjectures for Alive Status Among Those with COVID

| | Conjecture |
|----|--|
| 1 | $activeConditionLength > age^2 / latitude \rightarrow Alive$ |
| 2 | $medicationsLifetime < -immunizationsLifetimeCost + 2 \times proceduresLifetime \rightarrow Alive$ |
| 3 | $medicationsActive < \min\{sodium, [QOLS]\} \rightarrow Alive$ |
| 4 | $activeCarePlans < e^{medicationsLifetimePercCovered} - immunizationsLifetime \rightarrow Alive$ |
| 5 | $activeCarePlanLength < 10^{encountersLifetimePercCovered} \times activeCarePlans \rightarrow Alive$ |
| 6 | $lifetimeConditionLength > \sqrt{QALY^{activeConditions}} \rightarrow Alive$ |
| 7 | $lifetimeCarePlans > encountersCount / 2 \rightarrow Alive$ |
| 8 | $lifetimeCarePlans > \min\{triglycerides, [lifetimeConditions]\} \rightarrow Alive$ |
| 9 | $encountersCount > activeConditions + medicationsLifetimeCost + 1 \rightarrow Alive$ |
| 10 | $diabetes \text{ and } latitude < \sqrt{totalCholesterol} + meanCarbonDioxide \rightarrow Alive$ |
| 11 | $activeCarePlans > medicationsLifetime^{medicationsLifetimeCost} \rightarrow Alive$ |
| 12 | $lifetimeCarePlans > \sqrt{encountersCount} + QOLS \rightarrow Alive$ |
| 13 | $age < lifetimeConditions \times \log(latitude) \rightarrow Alive$ |
| 14 | $anemiaDisorder \text{ and } activeConditionLength < \min\{ureaNitrogen, [activeCarePlanLength]\} \rightarrow Alive$ |
| 15 | $lifetimeCarePlans > \max\{DALY, [potassium]\} \rightarrow Alive$ |
| 16 | $healthcareCoverage > [DALY] \times lifetimeCarePlanLength \rightarrow Alive$ |
| 17 | $lifetimeConditionLength > (encountersLifetimeTotalCost - 1) / proceduresLifetime \rightarrow Alive$ |
| 18 | $healthcareCoverage > lifetimeConditionLength^2 / imagingStudiesLifetime \rightarrow Alive$ |
| 19 | $immunizationsLifetimeCost > (bodyHeight - 1)^{immunizationsLifetime} \rightarrow Alive$ |
| 20 | $numAllergies < activeCarePlanLength - age + 1 \rightarrow Alive$ |
| 21 | $activeCarePlans < [microalbuminCreatineRatio] - proceduresLifetimeCost \rightarrow Alive$ |
| 22 | $latitude < \sqrt{encountersLifetimeTotalCost} - medicationsLifetime \rightarrow Alive$ |
| 23 | $bodyMassIndex40 \rightarrow Alive$ |
| 24 | $activeCarePlanLength > age + proceduresLifetime - 1 \rightarrow Alive$ |
| 25 | $medicationsLifetimeLength < 2 \times activeCarePlan \times deviceLifetimeLength \rightarrow Alive$ |
| 26 | $carbonDioxide > respiratoryRate \times [hemoglobinA1cHemoglobinTotalInBlood] \rightarrow Alive$ |
| 27 | $osteoporosisDisorder \text{ and } lifetimeCarePlanLength < \min\{painSeverity, 2 \times activeCarePlanLength\} \rightarrow Alive$ |
| 28 | $healthcareCoverage < \sqrt{encountersLifetimePayerCoverage} \times medicationsLifetime \rightarrow Alive$ |
| 29 | $respiratoryRate < painSeverity + [leukocytesVolumeInBlood] \rightarrow Alive$ |
| 30 | $meanPainSeverity > \max\{proceduresLifetime, [hemoglobinA1cHemoglobinTotalInBlood]\} \rightarrow Alive$ |
| 31 | $prediabetes \text{ and } meanDiastolicBloodPressure > [carbonDioxide] + meanHeartRate \rightarrow Alive$ |
| 32 | $latitude < encountersCount \times [QOLS] \rightarrow Alive$ |
| 33 | $painSeverity < [meanPainSeverity] - 1 \rightarrow Alive$ |
| 34 | $lifetimeCarePlans > \max\{DALY, [potassium]\} \rightarrow Alive$ |
| 35 | $latitude < 2 \times DALY \times encountersLifetimePercCovered \rightarrow Alive$ |
| 36 | $activeCarePlanLength > \max\{sodium, [hematocritVolume]\} \rightarrow Alive$ |
| 37 | $medicationsLifetimePercCovered > latitude^2 / medicationsLifetimeDispense \rightarrow Alive$ |
| 38 | $healthcareCoverage > \frac{healthcareExpenses}{2 \times encountersLifetimePercCovered} \rightarrow Alive$ |

Table C.2. Evaluation of Conjectures for Alive Status Among Those with COVID

| Conjecture | Precision (%) | Support | Lift |
|------------|---------------|---------|------|
| 1 | 99.11 | 3,042 | 1.07 |
| 2 | 98.54 | 3,700 | 1.06 |
| 3 | 98.51 | 9,480 | 1.06 |
| 4 | 98.25 | 8,231 | 1.06 |
| 5 | 98.23 | 6,103 | 1.06 |

Table C.2. (Continued)

| Conjecture | Precision (%) | Support | Lift |
|------------|---------------|---------|------|
| 6 | 98.05 | 7,937 | 1.06 |
| 7 | 97.79 | 4,161 | 1.06 |
| 8 | 97.65 | 5,453 | 1.05 |
| 9 | 97.65 | 2,939 | 1.05 |
| 10 | 97.41 | 1,969 | 1.05 |
| 11 | 97.07 | 4,774 | 1.05 |
| 12 | 96.79 | 2,242 | 1.05 |
| 13 | 96.61 | 1,650 | 1.04 |
| 14 | 95.65 | 2,045 | 1.03 |
| 15 | 95.39 | 6,311 | 1.03 |
| 16 | 95.31 | 1,236 | 1.03 |
| 17 | 95.20 | 1,687 | 1.03 |
| 18 | 95.07 | 4,017 | 1.03 |
| 19 | 94.32 | 440 | 1.02 |
| 20 | 94.00 | 500 | 1.02 |
| 21 | 93.59 | 1,498 | 1.01 |
| 22 | 93.25 | 1,185 | 1.01 |
| 23 | 92.09 | 834 | 0.99 |
| 24 | 91.01 | 1,213 | 0.98 |
| 25 | 90.79 | 999 | 0.98 |
| 26 | 89.77 | 831 | 0.97 |
| 27 | 89.41 | 727 | 0.97 |
| 28 | 88.61 | 2,389 | 0.96 |
| 29 | 88.27 | 358 | 0.95 |
| 30 | 88.21 | 704 | 0.95 |
| 31 | 87.67 | 1,890 | 0.95 |
| 32 | 87.20 | 2,250 | 0.94 |
| 33 | 87.14 | 583 | 0.94 |
| 34 | 85.79 | 1,612 | 0.93 |
| 35 | 84.50 | 755 | 0.91 |
| 36 | 82.76 | 586 | 0.89 |
| 37 | 76.22 | 677 | 0.82 |
| 38 | 74.75 | 99 | 0.81 |

Table C.3. Conjectures for Deceased Status Among Those with COVID

| Conjecture | Sufficient condition |
|------------|---|
| 1 | $longitude > -age \times medicationsLifetimePercCovered \rightarrow Deceased$ |
| 2 | $medicationsActive$ (HTML translation failed) |
| 3 | $age > carbonDioxide \times [potassium] \rightarrow Deceased$ |
| 4 | $deviceLifetimeLength \leq 2 \times creatinine^{healthcareExpenses} \rightarrow Deceased$ |
| 5 | $implantableCardiacPacem \rightarrow Deceased$ |
| 6 | $latitude < \log(age)/\log(10)^{activeCarePlans} \rightarrow Deceased$ |
| 7 | $medicationsActive > [\log(alkalinePhosphataseEnzymaticActivity)/\log(10)] \rightarrow Deceased$ |
| 8 | $immunizationsLifetimeCost < age \times immunizationsLifetime^2 \rightarrow Deceased$ |
| 9 | $colonoscopy \text{ and } coronaryHeartDisease \rightarrow Deceased$ |
| 10 | $activeCarePlans < \min\{deviceLifetimeLength, medicationsActive\} \rightarrow Deceased$ |
| 11 | $glucose > [creatinine \times meanGlucose \rightarrow Deceased$ |
| 12 | $bodyWeight > [meanBodyWeight] + 1 \rightarrow Deceased$ |
| 13 | $healthcareExpenses < deviceLifetimeLength^2 \times lifetimeConditionLength \rightarrow Deceased$ |
| 14 | $lifetimeConditions > activeCarePlans + [ureaNitrogen] \rightarrow Deceased$ |
| 15 | $activeCarePlans > [\log(triglycerides)] \rightarrow Deceased$ |
| 16 | $healthcareExpenses < lifetimeConditionLength^2 + encountersLifetimeTotalCost \rightarrow Deceased$ |
| 17 | $overlappingMalignantNeo \rightarrow Deceased$ |
| 18 | $latitude > ureaNitrogen[albuminMassVolumeInSerumOrPlasma] \rightarrow Deceased$ |
| 19 | $activeConditionLength > erythrocytesVolumeInBlood \times [hemoglobinMassVolumeInBlood] \rightarrow Deceased$ |

Table C.3. (Continued)

| Conjecture | Sufficient condition |
|------------|--|
| 20 | $chronicObstructiveBronc \rightarrow Deceased$ |
| 21 | $longitude > \sqrt{healthcareCoverage} - encountersCount \rightarrow Deceased$ |
| 22 | $age > 10^{medicationsActive} - longitude \rightarrow Deceased$ |
| 23 | $chloride < [meanChloride] - lifetimeCarePlans \rightarrow Deceased$ |
| 24 | $medicationsActive > \max\{respiratoryRate, \log(latitude)\} \rightarrow Deceased$ |
| 25 | $localizedPrimaryOsteoa \rightarrow Deceased$ |
| 26 | $rheumatoidArthritis \rightarrow Deceased$ |
| 27 | $chronicPain \text{ and } smokesTobaccoDaily \rightarrow Deceased$ |
| 28 | $latitude < [erythrocyteDistributionWidth] - meanPainSeverity \rightarrow Deceased$ |
| 29 | $activeCarePlans > 10^{medicationsActive} / imagingStudiesLifetime \rightarrow Deceased$ |
| 30 | $tubalPregnancy \rightarrow Deceased$ |
| 31 | $activeConditions < medicationsActive^2 - medicationsLifetime \rightarrow Deceased$ |
| 32 | $alcoholism \text{ and } majorDepressionDisorder \rightarrow Deceased$ |
| 33 | $creatinine < [meanCreatinine] / lifetimeCarePlans \rightarrow Deceased$ |
| 34 | $healthcareCoverage < encoutnersLifetimePayerCoverage \times \log(latitude) / \log(10) \rightarrow Deceased$ |
| 35 | $encountersCount < \min\{DALY, 10^{immunizationsLifetime}\} \rightarrow Deceased$ |
| 36 | $lifetimeCarePlanLength > age + e^{medicationsLifetime} \rightarrow Deceased$ |
| 37 | $healthcareCoverage < activeCondionLength^2 - encountersLifetimeTotalCost \rightarrow Deceased$ |
| 38 | $age > 1/2 \times healthcareExpenses / immunizationsLifetimeCost \rightarrow Deceased$ |
| 39 | $activeCarePlanLength > activeConditionLength \times e^{DALY} \rightarrow Deceased$ |
| 40 | $medicationsLifetime < \sqrt{encountersLifetimePayerCoverage} - age \rightarrow Deceased$ |

Table C.4. Evaluation of Conjectures for Deceased Status Among Those with COVID

| Conjecture | Precision (%) | Support | Lift |
|------------|---------------|---------|------|
| 1 | 30.91 | 372 | 4.15 |
| 2 | 24.44 | 2,954 | 3.29 |
| 3 | 23.22 | 1,722 | 3.12 |
| 4 | 22.47 | 632 | 3.02 |
| 5 | 22.39 | 844 | 3.01 |
| 6 | 21.74 | 1,490 | 2.92 |
| 7 | 21.44 | 4,427 | 2.88 |
| 8 | 21.24 | 2,199 | 2.85 |
| 9 | 21.01 | 257 | 2.82 |
| 10 | 20.90 | 799 | 2.81 |
| 11 | 20.03 | 1,058 | 2.69 |
| 12 | 18.25 | 548 | 2.45 |
| 13 | 17.56 | 467 | 2.36 |
| 14 | 17.54 | 1,898 | 2.36 |
| 15 | 17.31 | 3,328 | 2.33 |
| 16 | 17.23 | 940 | 2.32 |
| 17 | 16.92 | 130 | 2.27 |
| 18 | 16.13 | 1,091 | 2.17 |
| 19 | 15.68 | 797 | 2.11 |
| 20 | 14.11 | 900 | 1.90 |
| 21 | 14.08 | 781 | 1.89 |
| 22 | 14.08 | 2,230 | 1.89 |
| 23 | 11.99 | 884 | 1.61 |
| 24 | 11.84 | 1,884 | 1.59 |
| 25 | 11.76 | 2,159 | 1.58 |
| 26 | 11.44 | 201 | 1.54 |
| 27 | 10.54 | 607 | 1.42 |

Table C.4. (Continued)

| Conjecture | Precision (%) | Support | Lift |
|------------|---------------|---------|------|
| 28 | 10.27 | 1,724 | 1.38 |
| 29 | 9.39 | 213 | 1.26 |
| 30 | 8.85 | 2,147 | 1.19 |
| 31 | 8.22 | 4,150 | 1.10 |
| 32 | 7.79 | 1,129 | 1.05 |
| 33 | 7.77 | 2,408 | 1.04 |
| 34 | 7.73 | 634 | 1.04 |
| 35 | 5.20 | 3,209 | 0.70 |
| 36 | 3.89 | 4,602 | 0.52 |
| 37 | 3.85 | 467 | 0.52 |
| 38 | 2.65 | 339 | 0.36 |
| 39 | 2.29 | 3,188 | 0.31 |
| 40 | 1.30 | 1,001 | 0.17 |

Endnotes

- ¹ The independence number is the largest number of nodes in a graph, no two of which are contained in an edge. The definition of independence number is not important for this example, only the fact that with every graph is associated a number called the “independence number.”
- ² A Hamiltonian graph is a graph with a spanning cycle (West 2001). The definition of Hamiltonian is not important for this example, only the fact that any graph either is or is not Hamiltonian.
- ³ See <http://nvcleemp.github.io/conjecturing/>.
- ⁴ See <https://space.mit.edu/home/tegmark/aifeynman.html>.
- ⁵ See <https://github.com/SJ001/AI-Feynman>.
- ⁶ See <https://precision.fda.gov/challenges/11/view>.
- ⁷ See <https://www.va.gov/health-care/copay-rates/>, accessed July 10, 2022.
- ⁸ See <https://www.cdc.gov/coronavirus/2019-ncov/need-extra-precautions/people-with-medical-conditions.html>, accessed July 10, 2022.
- ⁹ See <https://www.cdc.gov/diabetes/managing/managing-blood-sugar/a1c.html>, accessed July 10, 2022.

References

- Abolafia D, Norouzi M, Shen J, Zhao R, Le Q (2018) Neural program synthesis with priority queue training. Preprint, submitted January 10, <https://arxiv.org/abs/1801.03526>.
- Aghaei S, Gómez A, Vayanos P (2021) Strong optimal classification trees. Preprint, submitted March 29, <https://arxiv.org/abs/2103.15965>.
- Bellomari L, Benedetto D, Gottlob G, Sallinger E (2020) Vadalog: A modern architecture for automated reasoning with large knowledge graphs. *Inform. Systems* 105:101528.
- Bertsimas D, Dunn J (2017) Optimal classification trees. *Machine Learning* 106:1039–1082.
- Blanquero R, Carrizosa E, Molero-Río C, Romero Morales D (2021) Optimal randomized classification trees. *Comput. Oper. Res.* 132:105281.
- Bradford A, Day J, Hutchinson L, Larson CE, Mills M, Muncy D, Kaperick B, Van Cleemput N (2020) Automated conjecturing II: Chomp and intelligent game play. *J. Artificial Intelligence Res.* 68:447–461.
- Breiman L, Friedman JH, Olshen RA, Stone CJ (1984) *Classification and Regression Trees* (Routledge, New York).
- Brunton S, Proctor J, Kutz J (2016) Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proc. Natl. Acad. Sci. USA* 113:3932–3937.
- Chattopadhyay I, Lipson H (2014) Data smashing: Uncovering lurking order in data. *J. Royal Soc. Interface* 11:20140826.
- Chvátal V (1972) On Hamilton’s ideals. *J. Combin. Theory Ser. B* 12: 163–168.
- Chvátal V, Erdős P (1972) A note on Hamiltonian circuits. *Discrete Math.* 2(2):111–113.
- Dash S, Günlük O, Wei D (2018) Boolean decision rules via column generation. *32nd Conf. Neural Inform. Processing Systems (NeurIPS-18)* (Curran Associates, Red Hook, NY), 4660–4670.
- Elton D (2020) Self-explaining AI as an alternative to interpretable AI. Preprint, submitted February 12, <https://arxiv.org/abs/2002.05149>.
- Fajtlowicz S (1995) On conjectures of graffiti. *Graph Theory, Combinatorics, and Algorithms*, vol. 1 (Wiley, New York), 367–376.
- Fürnkranz J, Kliegr T, Paulheim H (2020) On cognitive preferences and the plausibility of rule-based models. *Machine Learning* 109:853–898.
- Haemers W (1979) On some problems of Lovász concerning the shannon capacity of a graph. *IEEE Trans. Inform. Theory* 25(2):231–232.
- Hammer P, Bonates T (2006) Logical analysis of data—An overview: From combinatorial optimization to medical applications. *Ann. Oper. Res.* 148: 203–225.
- Hu D, Li J, Gao R, Wang S, Li Q, Chen S, Huang J, et al. (2021) Decreased CO₂ levels as indicators of possible mechanical ventilation-induced hyperventilation in COVID-19 patients: A retrospective analysis. *Frontiers Public Health* 8:596168.
- Jantzen B (2016) Dynamical kinds and their discovery. Preprint, submitted December 15. <https://arxiv.org/abs/1612.04933>.
- Kanter JM, Veeramachaneni K (2015) Deep feature synthesis: Toward automating data science endeavors. *2015 IEEE Internat. Conf. Data Sci. Advanced Analytics* (Institute of Electrical and Electronics Engineers, Piscataway, NJ).
- Katz G, Shin ECR, Song D (2016) ExploreKit: Automatic feature generation and selection. *16th IEEE Internat. Conf. Data Mining* (Institute of Electrical and Electronics Engineers, Piscataway, NJ).
- Khurana U, Samulowitz H, Turaga D (2018) Feature engineering for predictive modeling using reinforcement learning. *32nd AAAI Conf. Artificial Intelligence (AAAI-18)* (Association for the Advancement of Artificial Intelligence, Palo Alto, CA).
- Langley P, Simon HA, Bradshaw GL, Zytkow JM (1987) *Scientific Discovery: Computational Explorations of the Creative Process* (MIT Press, Cambridge, MA).
- Langley P (2019) Scientific discovery, causal explanation, and process model induction. *Mind Soc.* 18:43–56.
- Larson CE, Van Cleemput N (2016) Automated conjecturing I: Fajtlowicz’s Dalmatian heuristic revisited. *Artificial Intelligence* 231: 17–38.
- Larson CE, Van Cleemput N (2017) Automated conjecturing III: Property-relations conjectures. *Ann. Math. Artificial Intelligence* 81(3):315–327.
- Lemadjeng AC, Rober T, Akyuz MH, Birbil SI (2023) Rule generation for classification: Scalability, interpretability, and fairness. Preprint, submitted August 30, <https://arxiv.org/abs/2104.10751v3>.
- Lovász L (1979) On the Shannon capacity of a graph. *IEEE Transactions Information Theory* 25(1):1–7.
- Lu J, Lee DK, Kim T, Danks D (2019) Good explanation for algorithmic transparency. Preprint, submitted November 11, <https://dx.doi.org/10.2139/ssrn.3503603>.
- Nguyen Q, Nguyen X, O’Neill M, McKay R, Galván-López E (2011) Semantically-based crossover in genetic programming: Application to real-valued symbolic regression. *Genetic Programming Evolvable Machines* 12:91–119.
- Nicolau M, Agapitos A (2021) Choosing function sets with better generalisation performance for symbolic regression models. *Genetic Programming Evolvable Machines* 22:73–100.
- Noori M, Nejadghaderi S, Sullman M, Carson-Chahhoud K, Kolahi AA, Safiri S (2022) Epidemiology, prognosis and management of potassium disorders in Covid-19. *Rev. Medical Virology* 32: e2262.
- Petersen B, Larma M, Mundhenk T, Santiago C, Kim S, Kim J (2021) Deep symbolic regression: Recovering mathematical expressions from data via risk-seeking policy gradients. *Proc. Internat. Conf. Learning Representation (ICLR)* (International Conference on Learning Representations, Appleton, WI).
- Rudin C (2019) Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence* 1:206–215.
- Rudin C, Ertekin S (2018) Learning customized and optimized lists of rules with mathematical programming. *Math. Programming Comput.* 10:659–702.
- Samek W, Müller KR (2019) Toward explainable artificial intelligence. Samek W, Montavon G, Vedaldi A, Hanson L, Müller KR, eds. *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning* (Springer Nature, Cham, Switzerland), 5–22.
- Schmidt M, Lipson H (2009) Distilling free-form natural laws from experimental data. *Science* 324(5923):81–85.

- Schrijver A (2003) *Combinatorial Optimization: Polyhedra and Efficiency*, vol. 24 (Springer-Verlag, Berlin, Heidelberg, Germany).
- Tallorin L, Wang JL, Kim WE, Sahu S, Kosa NM, Yang P, Thompson M, et al. (2018) Discovering de novo peptide substrates for enzymes using machine learning. *Nature Comm.* 9(1):1–10.
- Therneau T, Atkinson B (2019) rpart: Recursive partitioning and regression trees. R package version 4.1-15. Retrieved May 19, 2021, <https://CRAN.R-project.org/package=rpart>.
- Tibshirani R (1996) Regression shrinkage and selection via the LASSO. *J. Royal Statist. Soc. B* 58:267–288.
- Tsang M, Cheng D, Liu Y (2018a) Detecting statistical interactions from neural network weights. *Sixth Internat. Conf. Learn. Representations (ICLR-18)* (International Conference on Learning Representations, Appleton, WI).
- Tsang M, Rambhatla H, Liu Y (2020) How does this interaction affect me? Interpretable attribution for feature interactions. *34th Conf. Neural Inform. Processing Systems (NeurIPS-20)* (Curran Associates, Red Hook, NY), 6147–6159.
- Tsang M, Liu H, Purushotham S, Pavankumar M, Liu Y (2018b) Neural interaction transparency (NIT): Disentangling learned interactions for improved interpretability. *32nd Conf. Neural Inform. Processing Systems (NeurIPS-18)* (Curran Associates, Red Hook, NY), 5809–5818.
- Udrescu SM, Tegmark M (2020) AI Feynman: A physics-inspired method for symbolic regression. *Sci. Adv.* 6:eaay2631.
- Verwer S, Zhang Y (2019) Learning optimal classification trees using a binary linear program formulation. *33rd AAAI Conf. Artificial Intelligence (AAAI-19)* (Association for the Advancement of Artificial Intelligence, Palo Alto, CA).
- Vilone G, Longo L (2020) Explainable artificial intelligence: A systematic review. Preprint, submitted May 29, <https://arxiv.org/abs/2006.00093>.
- Wang F, Rudin C (2015) Falling rule lists. *18th Internat. Conf. Artificial Intelligence Statist. (AISTATS)* (Machine Learning Research Press, Ft. Lauderdale, FL).
- Wang T, Rudin C, Doshi-Velez F, Liu Y, Klampfl E, MacNeille P (2017) A Bayesian framework for learning rule sets for interpretable classification. *J. Machine Learning Res.* 18:1–37.
- West DB (2001) *Introduction to Graph Theory*. 2nd ed. (Prentice Hall, Hoboken, NJ).