



Motivated by the joint goals of prediction and explanation, we propose Structured Sum-of-Squares Decomposition (S3D) algorithm, a method for learning interpretable statistical models of behavioral data. The algorithm, a variant of regression trees [7], builds a single tree-like structure that is both highly *interpretable* and can be used for out-of-sample *prediction*. In addition, the learned models can be used to *visualize* data. S3D is a mathematically principled method that addresses the computational challenges associated with mining behavioral data. Such data is usually massive, containing records of many individuals, each with a large number of, potentially highly correlated, features. However, the data is also sparse (with only a few observations available per individual) and unbalanced (few examples of the behavior within each class). Yet another challenge is heterogeneity: data is composed of subgroups that vary widely in their behavior. For example, the vast bulk of social media users have very few followers and post a few messages, but a few users have millions of followers or are extremely prolific posters. Ignoring heterogeneity could lead analysis to wrong conclusions due to statistical paradoxes [8, 9].

S3D works as follows: given a set of features  $\{X_i\}_{i=1}^M$  and an outcome variable  $Y$ , S3D identifies a subset of  $m$  important features that are orthogonal in their relationship with the outcome and collectively explain the largest amount of the variation in the outcome. In addition to these *selected* features, S3D identifies correlations between all features, thus providing important insights into the effects of features that were not selected by the model. Similar to regression trees and other decision trees, the S3D algorithm recursively partitions the  $m$ -dimensional space defined by the selected features into smaller, more homogeneous subgroups or bins, where the outcome variable exhibits little variation within each bin but significant variation between bins; however, it does so in a structured way, by minimizing variation in the outcome *conditioned* on the existing partition. The decomposition effectively creates an approximation of the (potentially non-linear) functional relationship between  $Y$  and the features, while the structured nature of the decomposition gives the model interpretability and also reduces overfitting. The resulting model is parsimonious and, despite its low complexity, is a highly performant predictive tool.

To demonstrate the utility of the proposed method, we apply it to model a variety of datasets, from benchmarks to large-scale heterogeneous behavioral data that we collect from social platforms, including Twitter, Digg, Khan Academy, Duolingo, and Stack Exchange. Across datasets, performance of S3D is competitive to existing state-of-the-art methods on both classification and regression tasks, while it also offers several advantages. We highlight these advantages by showing how S3D reveals the important factors in explaining and predicting behaviors, such as experience, skill, and answer complexity when analyzing performance on the question answering site Stack Exchange or essential nodal attributes, such as activity and rate of receiving information on the social networks Digg and Twitter. Qualitatively, S3D visualizes the relationship between the outcome and features, and quantifies their importance via prediction task. Surprisingly, despite high heterogeneity of these relationships in many datasets, just a few important features identified by S3D can predict held-out data with remarkable accuracy.

Machine learning, data science, and social science communities have proposed different solutions to learning models of data. Popular among these are regression

methods, such as linear and logistic regression. Lasso [10] and elastic net [11] improve on standard regression, using regularization to reduce the dimensionality of the feature space and to prevent overfitting. Decision trees and their regression tree variants that inspired S3D, such as CART [7], BART [12], and MARS [13], partition data into non overlapping subsets to minimize the variance of response variable, or some other cost function, within these subsets. Ensemble tree-based methods, such as random forests [14], learn a model as a sum or ensemble of individual decision trees. However, while these approaches address one set of challenges, they often trip over the remaining ones. Regression models (e.g., ridge, lasso, elastic net) are limited by their specified functional form, and cannot capture relationships in data that do not adhere to this form. Decision trees, including random forests, on the other hand, are very effective at capturing non-linear and unbalanced data, but have limited interpretability. While they offer a measure of feature importance, the relationship between the outcome and features is not easily analyzable: explaining the predictions requires navigating the depths of many trees—with features potentially repeating at different levels—a herculean task at best. Our method, on the other hand, is fully transparent and allows for visualizing and explaining why predictions were made. S3D has lower complexity than random forests, which typically consist of dozens of trees; indeed S3D also has the advantage of having only two tunable parameters, as opposed to twelve for a typical random forest. Finally, S3D adds structure to the feature importance offered by random forests, showing which subset of features are sufficient for explaining variation in outcomes, as well constructing feature networks which show correlation between features and where the variance from the other features is absorbed.

S3D learns a low complexity model with only two hyperparameters. Its ability to capture non-linear and heterogeneous behavior, while remaining interpretable, makes it a valuable computational tool for the analysis of behavioral data.

## 2 Method

We present the Structured Sum-of-Squares Decomposition algorithm (S3D), a variant of regression trees [7, 12, 13], that takes as input a set of features  $\mathbf{X} = \{X_i\}_{i=1}^M$  and an outcome variable  $Y$  and selects a smaller set of  $m$  important features that collectively best explain the outcome. The method partitions the values of each important feature  $X_i^S$  to decompose the  $m$ -dimensional selected feature space into smaller non-overlapping blocks, such that  $Y$  exhibits significant variation between blocks but little variation within each block. These blocks allow us to approximate the functional relationship  $Y = f(\mathbf{X})$  as a multidimensional step function over all blocks in the  $m$ -dimensional selected feature space, thus capturing non-linear relationships between features and the outcome.

Our method chooses features recursively in a forward selection manner, so that features chosen at each step explain most of the variation in  $Y$  *conditioned* on the features chosen at the previous steps. Features that are correlated will explain much of the same variation in  $Y$ , and our approach of successively choosing features based on how much remaining variation in  $Y$  they explain results in a set  $\mathbf{X}^S = \{X_i^S\}_{i=1}^m$  of important features that are orthogonal in their relationships with  $Y$ . By decomposing data recursively, we create a parsimonious model that

quantifies relationships not only between the features  $X_i$  and the outcome variable  $Y$ , but also among the features themselves. We show that our model is able to achieve performance comparable to state-of-the-art machine learning algorithms on prediction tasks, while offering advantages over those methods: our algorithm uses only two tuning parameters, can represent non-linear relationships between variables, and creates an interpretable model that is amenable to analysis and produces insights into behavior that merely predictive models do not give.

## 2.1 Structured Feature Space Decomposition

A key concept used to describe variation in observations  $\{y_i\}_{i=1}^N$  of a random variable  $Y$  is the total sum of squares  $SST$ , which is defined as  $SST = \sum_{i=1}^N (y_i - \bar{y})^2$ , where  $\bar{y} = \sum_{i=1}^N y_i / N$  is the sample mean of the observations. The total sum of squares is intrinsically related to variation in  $Y$ ; indeed the sample variance  $\hat{\sigma}^2$  of  $Y$  can be directly obtained from this quantity as  $\hat{\sigma}^2 = SST / (N - 1)$ .

Given a feature  $X_j$ , one method of quantifying how much variation in  $Y$  can be explained by  $X_j$  is as follows: (1) partition  $X_j$  into a collection  $P_{X_j}$  of non-overlapping bins, (2) compute the number of data points  $N_p$  and the average value  $\bar{y}_p$  of  $Y$  in each bin  $p \in P_{X_j}$ , and (3) decompose the total sum of squares of  $Y$  as

$$\sum_{i=1}^N (y_i - \bar{y})^2 = \sum_{p \in P_{X_j}} N_p (\bar{y}_p - \bar{y})^2 + \sum_{p \in P_{X_j}} \sum_{i=1}^{N_p} (y_{p,i} - \bar{y}_p)^2, \quad (1)$$

where  $y_{p,i}$  here is the  $i$ 'th data point in bin  $p$ . The first sum on the right hand side of Eq. (1) is the explained sum of squares (or sum of squares between groups), a weighted average of squared differences between global ( $\bar{y}$ ) and local ( $\bar{y}_p$ ) averages that measures how much  $Y$  varies between different bins  $p$  of  $X_j$ . The second sum is the residual sum of squares (or sum of squares within groups), which measures how much variation in  $Y$  remains within each bin  $p$ . The  $R^2$  coefficient of determination is then the proportion of the explained sum of squares to the total sum of squares, given by

$$R^2 = \frac{\sum_{p \in P_{X_j}} N_p (\bar{y}_p - \bar{y})^2}{SST}. \quad (2)$$

The  $R^2$  measure takes values between zero and one, with large values of  $R^2$  indicating a larger proportion of the variation of  $Y$  explained by  $X_j$ . This method of approximating the functional relationship between  $Y$  and  $X_j$  as a step function with bins, or groups  $P_{X_j}$  and corresponding values  $\bar{y}_p$ , allows us to quantify the variation in  $Y$  explained by  $X_j$  through the  $R^2$  of the corresponding step function as given by Eq. (2).

### 2.1.1 Partitioning Values of a Feature

We now introduce a method to systematically *learn* the partition  $P_{X_j}$  of the feature  $X_j$  which will be central to our algorithm. Given the data, we can split the domain of the feature  $X_j$  at the value  $s$  into two bins:  $X_j \leq s$  and  $X_j > s$ . From Eq. (2),

we see that the proportion of variation in  $Y$  explained by such a split is

$$R^2(s; X_j) = \frac{N_{X_j \leq s}(\bar{y}_{X_j \leq s} - \bar{y})^2 + N_{X_j > s}(\bar{y}_{X_j > s} - \bar{y})^2}{SST}, \quad (3)$$

where  $N_{X_j \leq s}$  and  $\bar{y}_{X_j \leq s}$  are the number of data points and average value of  $Y$  in the bin  $X_j \leq s$ , and vice versa for  $N_{X_j > s}$  and  $\bar{y}_{X_j > s}$ .  $R^2(s; X_j)$  can be computed for each possible value of  $s$  in the domain of  $X_j$ , and we can choose the optimal split  $s_1$  as the split  $s$  that maximizes  $R^2(s; X_j)$  of Eq. (3). Choosing  $s_1$ , and partitioning the domain of  $X_j$  into  $P_{X_j} = \{[\min(X_j), s_1], (s_1, \max(X_j))\}$ , we can again find the next best split  $s_2$  to optimize the improvement in  $R^2$ . In general, having chosen  $n$  splits  $\{s_u\}_{u=1}^n$  with a resulting partition  $P_{X_j}$  of  $n+1$  bins, the next best split  $s_{n+1}$  can be chosen as the split  $s$  that maximizes the improvement in  $R^2$  as given by

$$\Delta R^2(s|P_{X_j}; X_j) = \frac{1}{SST} \left( N_{p(s)|X_j \leq s}(\bar{y}_{p(s)|X_j \leq s})^2 + N_{p(s)|X_j > s}(\bar{y}_{p(s)|X_j > s})^2 - N_{p(s)}(\bar{y}_{p(s)})^2 \right) \quad (4)$$

where  $p(s)$  in Eq. (4) is the bin in  $P_{X_j}$  that contains the point  $s$  and  $p(s)|X_j \leq s$  (resp.  $p(s)|X_j > s$ ) is the restriction of that bin to points  $X_j \leq s$  (resp.  $X_j > s$ ). In this manner, we recursively split the domain of  $X_j$  to create a partition of the feature.

However, splitting in this manner can continue indefinitely, resulting in a model that is too fine-grained and thus overfits the data. To prevent overfitting, we need a stopping criterion. To this end we introduce a loss function  $L(P_{X_j})$  that penalizes the size  $|P_{X_j}|$  of the partition, i.e., the number of bins:

$$L(P_{X_j}) = 1 - R^2(P_{X_j}) + \lambda |P_{X_j}|. \quad (5)$$

The parameter  $\lambda$  controls how fine-grained the bins are: smaller values of  $\lambda$  allow for more finer bins, and vice versa. The loss function of Eq. (5) reaches a minimum when the change in  $R^2(P_{X_j})$  from adding an extra split to  $P_{X_j}$  is less than  $\lambda$ —at this point we stop splitting and return the partition  $P_{X_j}$ .

Having formed the partition  $P_{X_j}$  of  $X_j$  with splits  $\{s_u\}_{u=1}^n$ , the total score  $R^2(X_j)$  can be calculated from Eq. (2), or by summing  $R^2(s_1; X_j)$  from Eq. (3) along with the  $\Delta R^2$  terms in Eq.(4) for each of the splits  $\{s_u\}_{u=2}^n$ . Completing this procedure for all features gives a measure of how much variation in  $Y$  each feature alone explains, and ranking these features allows us to choose the most important feature  $X_{C_1}$  that explains the largest amount of the variation in  $Y$ .

### 2.1.2 Selecting Additional Features

After choosing the most important feature, we search the rest of the features for one that explains most of the remaining variance in  $Y$ , then the third feature, and so on. Here, we describe the procedure for finding the next best feature having already chosen  $l$  features  $\mathbf{X}^S = \{X_1^S, \dots, X_l^S\}$  with a corresponding partition  $\mathcal{P}^S = P_1^S \times \dots \times P_l^S$ , where  $\times$  here is the cartesian product. In this case, a total  $R^2(\mathcal{P}^S) =$

$\sum_{p \in \mathcal{P}^S} N_p (\bar{y}_p - \bar{y})^2 / SST$  of the variation in  $Y$  has been explained, and we now look for the feature that best explains the remaining variation  $1 - R^2(\mathcal{P}^S)$ .

Given a remaining feature  $X_j$ , we partition the domain of  $X_j$  similarly to how we partitioned it when choosing the first feature. The first split  $s_1$  of  $X_j$  is chosen as the value  $s$  that maximizes the improvement in  $R^2$ , given by

$$\Delta R^2(s | \mathcal{P}^S; X_j) = \frac{1}{SST} \sum_{p \in \mathcal{P}^S} \left( N_{p|X_j \leq s} (\bar{y}_{p|X_j \leq s})^2 + N_{p|X_j > s} (\bar{y}_{X_j > s})^2 - N_p (\bar{y}_p)^2 \right), \quad (6)$$

where  $p|X_j \leq s$  (resp.  $p|X_j > s$ ) is the set of data points in  $p \in \mathcal{P}^S$  for which  $X_j \leq s$  (resp.  $X_j > s$ ). In general, given  $n$  splits and a corresponding partition  $P_{X_j}$  of  $X_j$ , the  $n + 1$ 'st split is chosen as the value  $s$  that maximizes

$$\Delta R^2(s | P_{X_j}, \mathcal{P}^S; X_j) = \frac{1}{SST} \sum_{p \in \mathcal{P}^S \times P_{X_j} | s \in p} \left( N_{p|X_j \leq s} (\bar{y}_{p|X_j \leq s})^2 + N_{p|X_j > s} (\bar{y}_{X_j > s})^2 - N_p (\bar{y}_p)^2 \right), \quad (7)$$

with the sum in Eq. (7) taken over all elements  $p$  of  $\mathcal{P}^S \times P_{X_j}$  that contain the point  $s$ . The loss function in the general setting is

$$L(P_{X_j} | \mathcal{P}^S) = \frac{1 - R^2(\mathcal{P}^S \times P_{X_j})}{1 - R^2(\mathcal{P}^S)} + \lambda |P_{X_j}|, \quad (8)$$

where the denominator of the fractional first term in Eq. (8) normalizes this term to be between zero and one, is the case in Eq. (5). This normalization is necessary because as we progress through the algorithm, subsequent features may explain less of the variance of  $Y$  (as features are chosen hierarchically), and so changes in  $1 - R^2(\mathcal{P}^S \times P_{X_j})$  by splitting  $X_j$  are smaller. The effect of this would be a coarser binning of the feature, and so the normalization ensures that this is not the case and that the feature is binned consistently at each stage of the algorithm. Again, the partition  $P_{X_j}$  of  $X_j$  is chosen that minimizes the loss function of Eq. (8), and the  $R^2$  improvement is calculated for this feature as

$$\Delta R^2(P_{X_j} | \mathcal{P}^S) = R^2(\mathcal{P}^S \times P_{X_j}) - R^2(\mathcal{P}^S). \quad (9)$$

This procedure is repeated for all remaining features  $X_j$  to select the feature  $X_{l+1}$  with the maximal  $R^2$  improvement.

This process of binning features, calculating their improvement  $\Delta R^2$  and choosing the one with the largest improvement continues until no further variation in  $Y$  can be explained or until an alternative stopping condition (such as a pre-specified maximum number of steps) is met. Our algorithm learns a hierarchy of important features that explain the variation in  $Y$  and a decomposition  $\mathcal{P}^S$  with corresponding values  $\{\bar{y}_p\}_{p \in \mathcal{P}^S}$  approximating the functional relationship between the outcome and the features. Note that, when binning a feature, once the vectors  $\{\{N_{p|X_j \leq s}\}_{p \in \mathcal{P}^S}\}_s$

and  $\{\{\bar{y}_p|X_j \leq s\}_{p \in \mathcal{P}^S}\}_s$  have been constructed (an operation that takes  $\mathcal{O}(N)$  time), the binning is independent of  $N$ , and instead depends on the number of unique values of the feature (as required to calculate the optimal splits  $s$  at each step). The fact that S3D scales linearly with the number of data points  $N$  allows us to apply the algorithm to large datasets, such as the Twitter and Digg (see Section 3).

### 2.1.3 Hyperparameters

The S3D model has two hyperparameters: (1)  $\lambda$  that controls granularity of feature binning; (2)  $k$  that specifies the number of features to use for prediction. Both are important to prevent overfitting — as left unrestricted, the algorithm can learn too fine-grained a model that fails to generalize to unseen data. We note that it is possible to stop early in the training phase by restricting the maximum number of features to select. In other words, the algorithm can stop when the number of selected features reaches a predefined value. Nonetheless, it is recommended *not* to lay any limit during the training phase but rather tune  $k$  in the validation step.

To tune the hyperparameters, we train S3D for various values of  $\lambda$ , in each case letting the algorithm continue until there is no further improvement in  $R^2$ . This results in a model with  $m_\lambda$  selected features and partition  $\mathcal{P}^{S_\lambda} = \{P_1^{S_\lambda}, \dots, P_{m_\lambda}^{S_\lambda}\}$ . Then, for  $k \in [1, m_\lambda]$ , we evaluate the predictive performance of the model using only the top  $k$  selected features and the sub-partition  $\{P_1^{S_\lambda}, \dots, P_k^{S_\lambda}\}$ . Performance is measured on held-out tuning data using a specified metric. The optimal hyperparameters  $(\lambda, k)$  are those that achieve the best performance on held-out data.

## 2.2 Applications of the Learned Model

Given a dataset, S3D learns an ordered set of important, orthogonal features  $\mathbf{X}^S$ , a partitioning  $\mathcal{P}^S$  of the selected feature space with corresponding  $\bar{y}_p$  and  $N_p$  values for each block  $p \in \mathcal{P}^S$ , and  $\Delta R^2$  for each remaining variable at each step of the algorithm. This decomposition serves as a parsimonious model of data and can be used for feature selection, feature correlation, prediction and analysis as described below.

### 2.2.1 Feature Selection and Correlations

The ordered set  $\mathbf{X}^S$  of important, orthogonal features allows us to quantify feature importance in heterogeneous behavioral data. The top-ranked features explain the largest amount of variation in the outcome variable, while each successive feature explains most of the remaining variation that is not explained by the features that were already selected.

Aside from the selected features  $\mathbf{X}^S$ , S3D provides insights into features that are *not* selected by the algorithm, quantifying variation that they explain in the outcome variable that is made redundant through the selected variables. This is calculated in the following manner. At a given step  $l$  of the algorithm, feature  $X_l^S$  is selected as the best feature with an  $R^2$  improvement of  $\Delta R^2(P_{X_l^S} | \mathcal{P}^{S(l-1)})$  (Eq. (9)), where  $\mathcal{P}^{S(l-1)}$  is the partition prior to step  $l$ . Meanwhile, a different remaining feature  $X_j$  has an  $R^2$  improvement of  $\Delta R^2(P_{X_j} | \mathcal{P}^{S(l-1)})$ . At the next stage of the algorithm, given  $X_l^S$  has been selected,  $X_j$  will have an  $R^2$  improvement

of  $\Delta R^2(P_{X_j}|\mathcal{P}^{S(l-1)} \times P_{X_i^S})$ , and thus the variation in  $X_j$  that is made redundant through the selection of  $X_i^S$  is the difference between these two  $\Delta R^2$ :

$$a_{X_j, X_i^S} = \Delta R^2(P_{X_j}|\mathcal{P}^{S(l-1)}) - \Delta R^2(P_{X_j}|\mathcal{P}^{S(l-1)} \times P_{X_i^S}). \quad (10)$$

The coefficients  $a_{X_j, X_i^S}$  facilitate our analysis of the effect of unselected features on the outcome variable, and we implement them as weights in a feature network that is weighted and directed. This network gives a tool for analysis—unselected features that otherwise can explain much of the variation in the outcome will have heavy links, and the selected features  $X_i^S$  to which these links point reveal correlations and through which selected features the unselected feature is made redundant.

### 2.2.2 Prediction

The learned model can be used as a predictive tool for both discrete and continuous-valued outcome variables. Given input data  $\mathbf{x}$ , the model predicts the expected value  $\hat{\mu}$  of  $Y$  as  $\hat{\mu}|\mathbf{x} = \bar{y}_{p(\mathbf{x})}$ , where  $p(\mathbf{x})$  is the block in decomposition  $\mathcal{P}^S$  to which  $\mathbf{x}$  belongs. For continuous-valued outcome variables, this predicted expected value  $\hat{\mu}$  will be the prediction of the outcome of  $Y$ , i.e.,  $\hat{y}|\mathbf{x} = \bar{y}_{p(\mathbf{x})}$ . For discrete-valued outcomes, the expected value has to be thresholded to predict an outcome class. For binary outcomes  $Y \in \{0, 1\}$ ,  $\hat{\mu}|\mathbf{x} = \bar{y}_{p(\mathbf{x})}$  is the maximum likelihood estimate of the probability of the outcome  $Y = 1$  in the block  $p(\mathbf{x})$ , and thus our model specifies that the outcome  $Y = 1|\mathbf{x}$  will occur with probability  $\bar{y}_{p(\mathbf{x})}$ . By choosing an appropriate *discrimination threshold*  $\theta$ , our model then makes the prediction  $\hat{y}$  as

$$\hat{y}|\mathbf{x} = \begin{cases} 1 & \bar{y}_{p(\mathbf{x})} \geq \theta \\ 0 & \bar{y}_{p(\mathbf{x})} < \theta. \end{cases} \quad (11)$$

*Unbalanced Data.* Two of the datasets that we study, Digg and Twitter, are highly unbalanced—their outcome variables (whether the user *adopts* a meme) are binary, and the proportions of positive outcomes in the data are 0.0025 and 0.0007 respectively. Using the standard discrimination threshold  $\theta = 0.5$  results in predicting an insufficient number of ones. To address this issue, we choose the discrimination threshold based on the training data, picking the largest value  $\theta = \theta_l$ , such that the number of predicted positive examples in the training data is greater than or equal to the actual number of positive examples in the training set. This threshold is then used for prediction on the held-out tuning data, as well as on the test data. Note that, for these two datasets, we also alter the discrimination threshold for the regression and random forest models in the same manner.

### 2.2.3 Analysis and Model Interpretation

One of the more novel contributions of S3D is its potential for model exploration. By selecting features sequentially, we create a model where typically lower amounts of variation are explained at successive levels, and so a visual analysis of the first few important dimensions of the model allows us to understand the effects of the important features on data and predictions. The expectations  $\bar{y}_p$  for each block  $p \in \mathcal{P}^S$



facilitate this exploration, approximating the relationship  $Y = f(\mathbf{X})$  between outcome and features. Furthermore, for binary data, the predicted outcomes obtained by thresholding the expectations show how predictions change as a function of the features, which allows for explaining predictions and visually exploring the data.

### 2.3 Comparison to State-of-the-art

We compare S3D to logistic regression (Lasso and elastic net [11]), random forests [14], and support vector machines with linear kernel (linear SVM) [15], using 5-fold cross validation (CV; Section 6.2). These algorithms are ideal benchmarks for S3D, with the ability to provide feature importance scores and therefore interpretability to the trained models. Further, we emphasize that S3D can produce sparse models that can do both feature selection and prediction. As shown in Section 3.2, S3D performs similarly to logistic regression in both classification and regression (Figures 2 and 3), but uses fewer features (Figure 4). Finally, both random forests [16] and linear SVM [17] are proven to be adequately powerful in prediction tasks while being relatively simple to train.

We partition data into five equal size folds,<sup>[1]</sup> each of which is rotated as a hold-out set for testing. For classification (except random forests), we standardize feature values by centering and scaling variance to one. For regression, both features and target values are standardized. Standardization of test set is based on training data. In each run, we train and tune the models on four folds, where three are used for training and one for validation. Finally we evaluate the performance of the tuned models on the remaining test fold. The final evaluation is therefore the average performance across each of the five folds. In other words, we tune the hyperparameters with 4-fold CV (hereafter *inner CV*) and evaluate the performance of the optimal models with 5-fold CV (hereafter *outer CV*.) For classification tasks, we evaluate performance using (1) accuracy (i.e., the percentage of correctly classified data points), (2) *F1* score (i.e., the harmonic mean of precision, the percentage of predicted ones that are correctly classified, and recall, the percentage of actual ones that are correctly classified), and (3) area under the curve (AUC.) For regression tasks, we employ (1) root mean squared error (RMSE), (2) mean absolute error (MAE-Mean), (3) median absolute error (MAE-Median). Note that for classification tasks, higher values of the metrics imply better performance. For regression tasks, lower values of the metrics imply better performance. During the inner CV phase (i.e., hyperparameter tuning), we optimize classification performance on AUC scores and regression on RMSE scores.

For Lasso regression we tune the strength of  $l_1$  regularization. For elastic net regression, we tune the strengths of both  $l_1$  and  $l_2$  regularizations. For random forests, we tune (1) the number of features to randomly select for each decision tree, (2) the minimum number of samples required to make a prediction, (3) whether or not to bootstrap when sampling data for each decision tree, and (4) criterion of the quality of a split (choices are Gini impurity or information gain for classification; MAE or MSE for regression.) For linear SVM, we tune the penalty parameter for regularization.

---

<sup>[1]</sup>For classification tasks, each fold is made by preserving the ratio of samples in each class (i.e., stratified.)

Finally, we compare performance of models of similar complexity. Since S3D makes predictions using a small number of features, we also retrain the random forest on a similarly small number of features (although, random forest is more complex, since it is an ensemble method). First, we use random forest to rank all the features and retrain the model on the top- $k$  (and top- $2k$ ) ranked features, where  $k$  is the number of important features selected by S3D. This enables us to compare the performance of models that use similar number of features. In addition, we investigate the effects of using S3D for *supervised feature selection*. Specifically, we train random forests using only features chosen by S3D (i.e.,  $k$  features).

### 3 Results

We apply S3D<sup>[2]</sup> to benchmark datasets from the UCI Machine Learning Repository [18] and Luís Torgo’s personal website.<sup>[3]</sup> In addition, we include five large-scale behavioral datasets, as described in the following paragraphs. ?? lists all 18 datasets used in both classification and regression tasks, along with their statistics. See Section 6 for more detailed description on data preprocessing.

Behavioral data came from various social platforms:

- *Stack Exchange*. The Q&A platform Stack Exchange enables users to ask and answer questions. Askers can also *accept* one of the answers as the best answer. This enables us to measure answerer performance by whether their answer was accepted as the best answer or not. The data we analyze includes a random sample of all answers posted on Stack Exchange from 08/2009 until 09/2014 that preserves the class distribution. Each record corresponds to an answer and contains a binary outcome variable  $Y \in \{0, 1\}$  (one indicates the answer was accepted, and zero otherwise), along with a variety of features. In total there are 14 features, including answer-based features, such as the number of *words*, *lines of code* and *hyperlinks* to Web content the answer contains, the number of other answers the question already has, its readability score; as well as such as answerer’s *reputation*, how long the answerer has been registered (*signup duration* in months) and as a percentile rank (*signup percentile*), the number of answers they have previously written, time since the previous answer, the number of answers written by the answerer in his or her current session, and answer’s position within the session.
- *Khan Academy*. The online educational platform Khan Academy enables users learn a subject then practice what they learned through a sequence of questions on the subject. We study performance during the practice stage by looking at whether users answered the questions correctly on their first attempt ( $Y = 1$ ) or not ( $Y = 0$ ). We study an anonymized sample of questions answered by adult Khan Academy users over a period from 07/2012 to 02/2014. For each question a user answers we have 19 features: as with Stack Exchange, these include answer-based, user-based, and other temporal features.
- *Duolingo*. The online language learning platform Duolingo is accessed through an app on a mobile device. Users are encouraged to use the app in short bursts

---

<sup>[2]</sup>S3D is available as C++ code and Python wrapper: see <https://github.com/peterfennell/S3D>.

<sup>[3]</sup><https://www.dcc.fc.up.pt/~ltorgo/Regression/DataSets.html>

during breaks and commutes. The data<sup>[4]</sup> was made available as part of a previous study [2]. The data contains a 2-week sample (02/28/2013 through 03/12/2013) of ongoing activity of users learning a language. All users in this data started lessons before the beginning of the data collection period. We focus on 45K users who completed at least five lessons. The median number of lessons was 7, although some had as many as 639 lessons. Performance on a lesson is defined as  $Y = 1$  if the user got *all* the words in the lesson correct; otherwise, it is  $Y = 0$ . Features describing the user include how many lessons and sessions the user completed, how many perfect lessons the user had, the month and day of the lesson, etc.

- *Digg*. The social news platform Digg allows users to post news stories, which their followers can like or “digg.” When a user diggs a story, that story is broadcast to his or her followers, a mechanism that allows for the diffusion of contents through the Digg social network. A further characteristic of Digg is its front page — stories that are popular are promoted to the front page and thus become visible to *every* Digg user. We study a dataset that tracks the diffusion of 3,500 popular Digg stories from their submissions by a single user to their eventual promotion to and residence on the Digg front page. We study information diffusion on Digg by examining whether or not ( $Y \in \{0, 1\}$ ) users “digg” (i.e., adopt) a story following exposure of the story from their friends, and thus share that information with their followers. The features associated with adoption include user-based features, such as indegree and out degree (number of followers and followees of the users), node activity (how often the user posts), information received, (the rate at which the user receives information for all followees), dynamics-related features such as the the number of times the user was exposed to the story, story-related features such as its global popularity in the previous hour, and diurnal-features, including the hour of the day and day of the week. Through this data, we can study the factors that are important in the spread of information in this social system.
- *Twitter*. On the online social network Twitter, users can post information, which is then broadcast to their followers, i.e., the other Twitter users that follow that user. This dataset tracks the spread of 65,000 unique URLs through the Twitter social network during one month in 2010. Similarly to *Digg*, we can study social influence and information diffusion by examining whether ( $Y = 1$ ) or not ( $Y = 0$ ) a user posts a URL after being exposed to it when one of his or her friends posts. The features associated with each exposure event are the same as those for *Digg*.

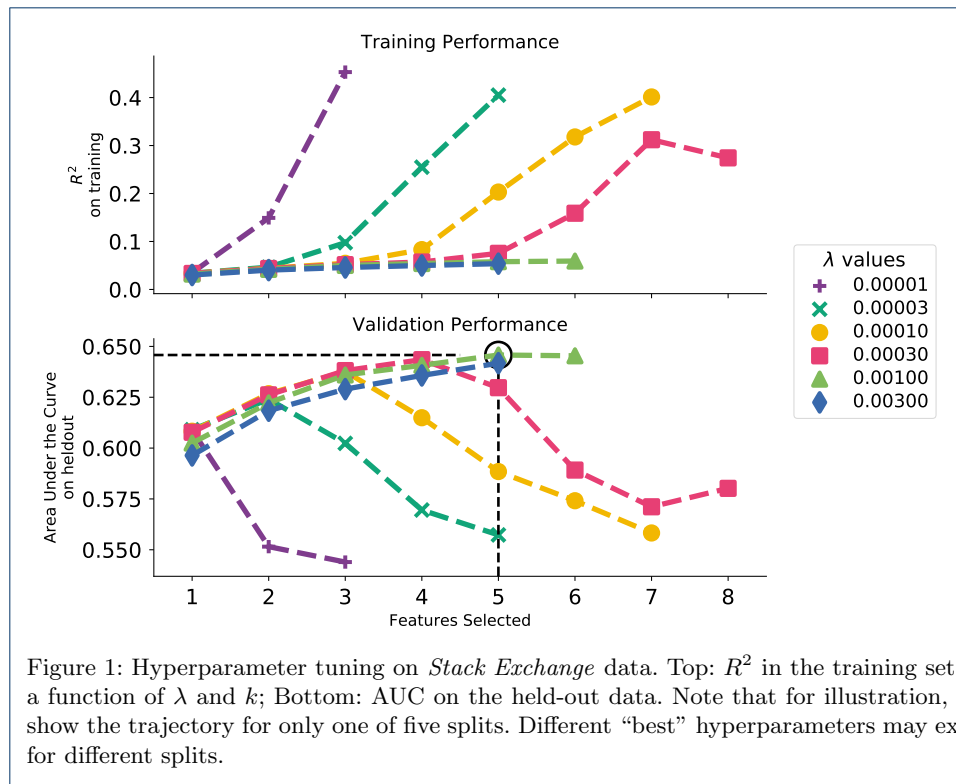
We compare the average predictive performance across 5 holdout sets of S3D to Lasso regression, elastic net regression, random forests, and linear SVM. We show that S3D can achieve competitive performance with the benchmark algorithms with a smaller set of features. Finally, we explore our tuned S3D models and demonstrate their utility to understanding human behaviors in Section 3.3. Relevant datasets and codes<sup>[5]</sup> can be used to replicate the following results.

### 3.1 Tuning hyperparameters

---

<sup>[4]</sup><https://github.com/duolingo/half-life-regression>

<sup>[5]</sup><https://github.com/peterfennell/S3D/tree/paper-replication>



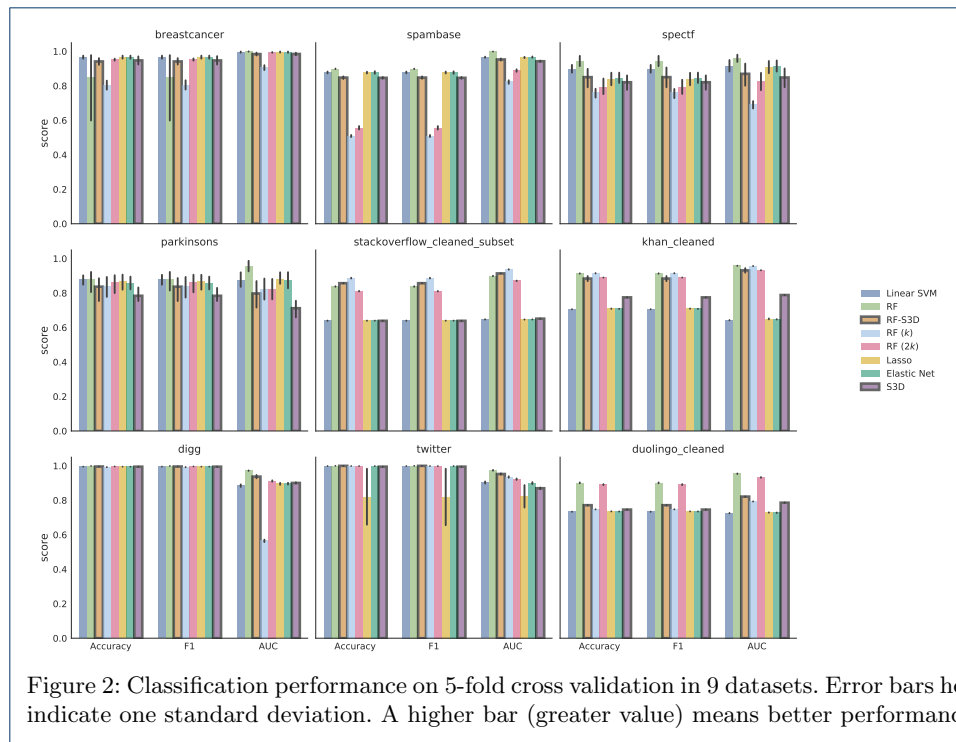
An essential part of training a statistical model is hyperparameter tuning — in the case of S3D, selecting the parameters  $\lambda$  and  $k$ . This procedure is illustrated for Stack Exchange data in Figure 1, where we show the total  $R^2$  at each step of the algorithm for various values of  $\lambda$ , as well as the AUCs at these steps computed on the held-out tuning data. Overly small values of  $\lambda$  perform quite poorly on the held-out data, as they produce very fine-grained partitions that overfit the data. Larger values of  $\lambda$  avoid being too fine-grained — the  $R^2$  on the training set increases initially but diverges again as extra features selected in additional steps overfit the data (as shown through the decreasing performance on the held-out data; Figure 1 bottom). Parameter  $k$  (x-axis in Figure 1) controls the number of steps of S3D, thus picking the optimal model between underfitting and overfitting. Supplementary file *s3d.hyperparameter\_df.csv* [6] reports the best hyperparameters for all datasets, across the 4-fold inner CV processes.

### 3.2 Prediction Performance

Figures 2 and 3 report performance on the outer CV for all datasets S3D, random forests, linear SVM, and logistic regressions (both Lasso and elastic net). Overall S3D achieves predictive performance comparable to other state-of-the-art machine learning methods.

In most cases, S3D’s performance is similar to that of logistic regression and linear SVM. Its performance relative to random forest is especially remarkable considering

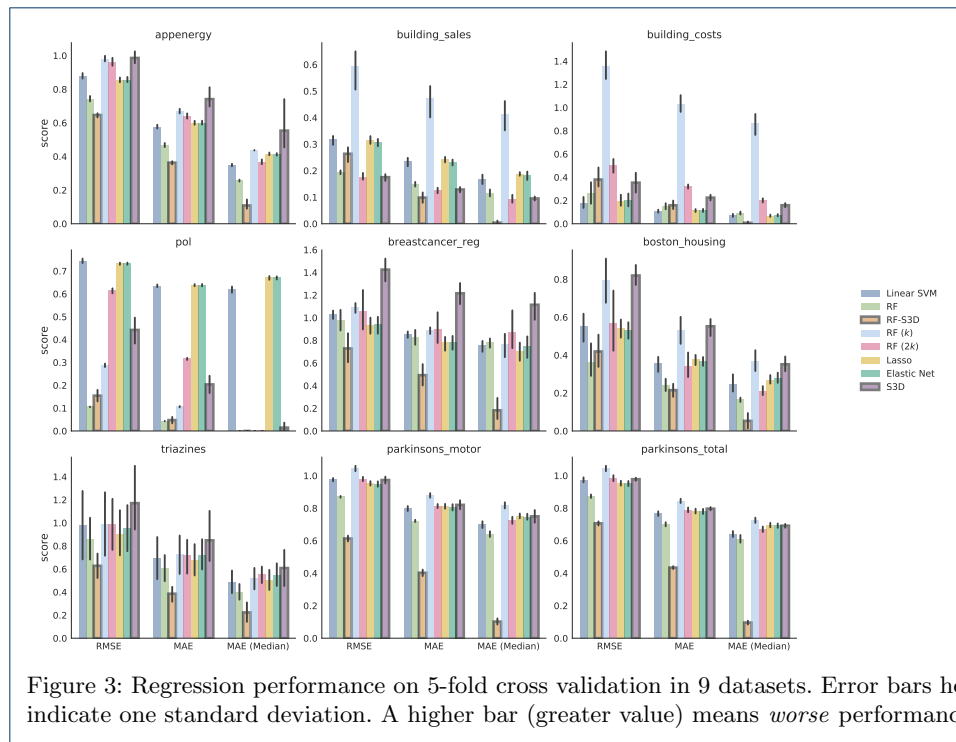
[6] [https://raw.githubusercontent.com/peterfennell/S3D/paper-replication/replicate/s3d\\_hyperparameter\\_df.csv](https://raw.githubusercontent.com/peterfennell/S3D/paper-replication/replicate/s3d_hyperparameter_df.csv)



the difference in complexity of the models. S3D uses a subset of features and a simple  $m$ -dimensional hypercube to make predictions, in stark contrast to random forests, which use all features and learn many decision trees. In contrast, S3D selects a small set of features, producing more parsimonious models as compared to Lasso and elastic net regression (Figure 4).

S3D is especially useful as a feature selection method. Using just the few features selected by S3D to train a random forest leads to highly competitive performance on the regression task for many datasets (RF-S3D bars in Figure 3). Remarkably, its performance is often better than that of the random forest trained on the full feature set. This is likely because features selected by S3D are uncorrelated with each other; while correlations among features used by the random forest reduce performance.

Finally, we show that the runtime of S3D is competitive to the other four algorithms (Figure 5). For each dataset, all models were trained using the best parameters found in inner CV and full training sets over each split (recall that there are five splits) repeatedly ten times. Therefore, each box in Figure 5 shows the distribution of training time over a total of 50 runs. Note that the Python package Scikit Learn [28] is used to implement logistic regression, random forest, and linear SVM, therefore producing superb runtime performance as it is highly optimized. We believe that the implementation of S3D can be further improved. For instance, the timing of S3D implementation includes reading the input file, whereas the other four have no such needs due to different setting. Furthermore, Figure 5 only reflects training time with one set of hyperparameters for each model. While random forests manifest outstanding efficiency, it is worth noting that the large amount of hyperparameters (in this study, we searched for four; there are at least four more) will



inevitably lead to undesirably long hours of grid search. On other hand, S3D only needs two (i.e.,  $\lambda$  and  $k$ ), which substantially reduce user effort in hyperparameter tuning.

### 3.3 Analyzing human behavior with S3D

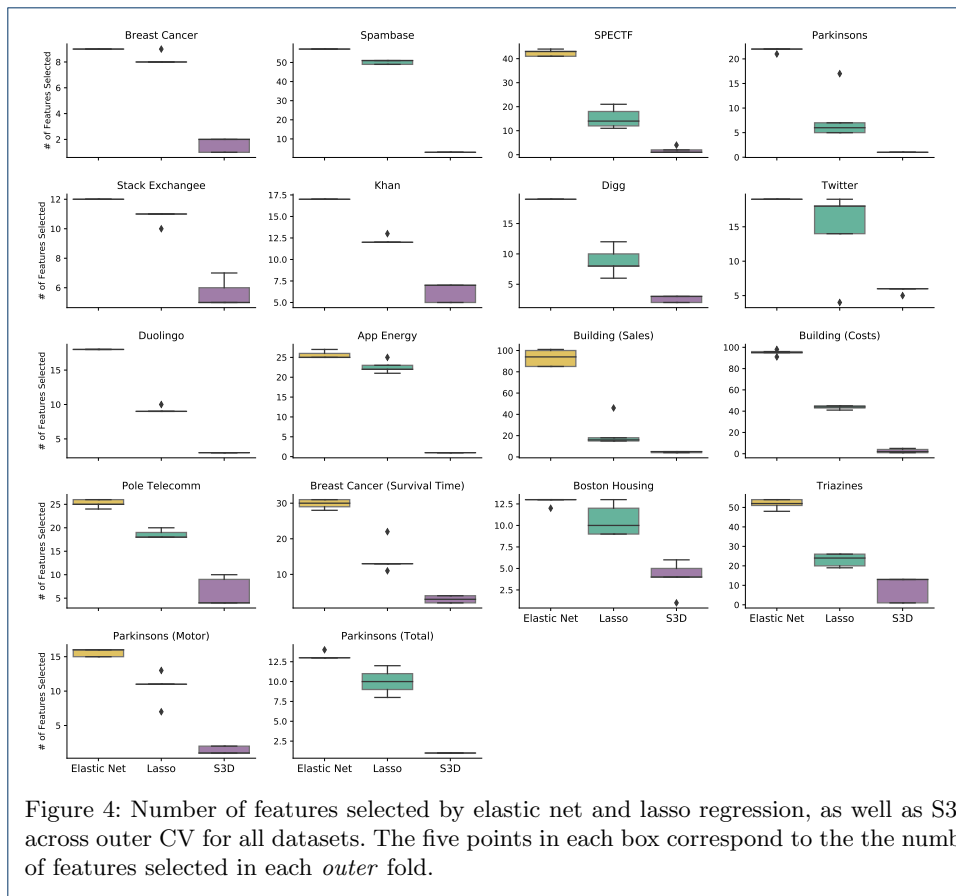
In this section, we present a detailed description of applying S3D to understand human behaviors using *Stack Exchange* data [7]. We additionally included *Digg* in Section 3.3.2 to demonstrate visualizations of the learned S3D model. Specifically, we used the best hyperparameters during cross validation:  $\lambda = 0.001$  and  $k = 5$  for *Stack Exchange*;  $\lambda = 0.001$  and  $k = 3$  for *Digg*. We note that, to explain the dataset, we applied S3D on *all the data available* using the best hyperparameters, which occurred most often during the cross validation processes.

#### 3.3.1 Feature Selection and Correlations

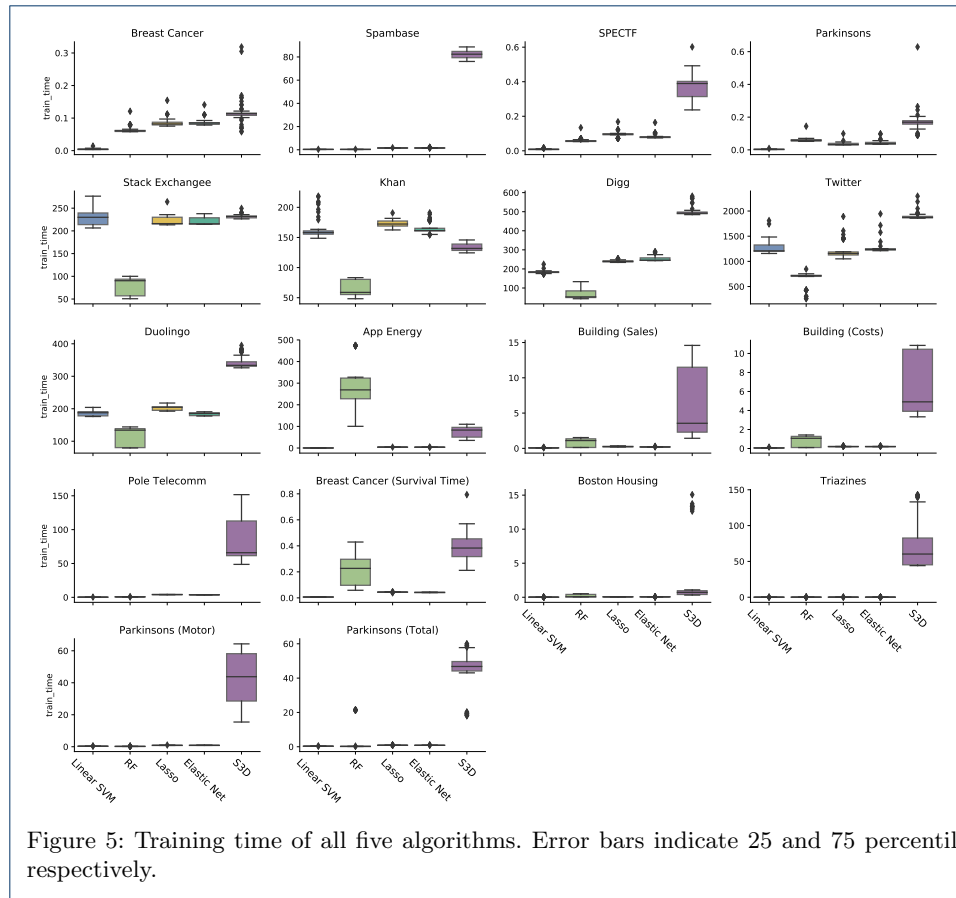
For the large scale behavioral data, S3D selects a subset of features that collectively explain the largest amount of the variation in the outcome variable. In the meantime, it quantifies correlations between selected features to unselected ones. In the following, we describe selected features and examine the effects of the unselected ones.

We give a detailed description of the features selected at each step (Figure 6) and the resulting feature network (Figure 7). Figure 6 visually ranks the features by showing the amount of variation explained by each feature at every step of the algorithm. The features selected at each step are outlined in black. The first and

[7] See <https://github.com/peterfennell/S3D/blob/paper-replication/replicate/3-visualize-models.ipynb> for results on the other four human behavior datasets.



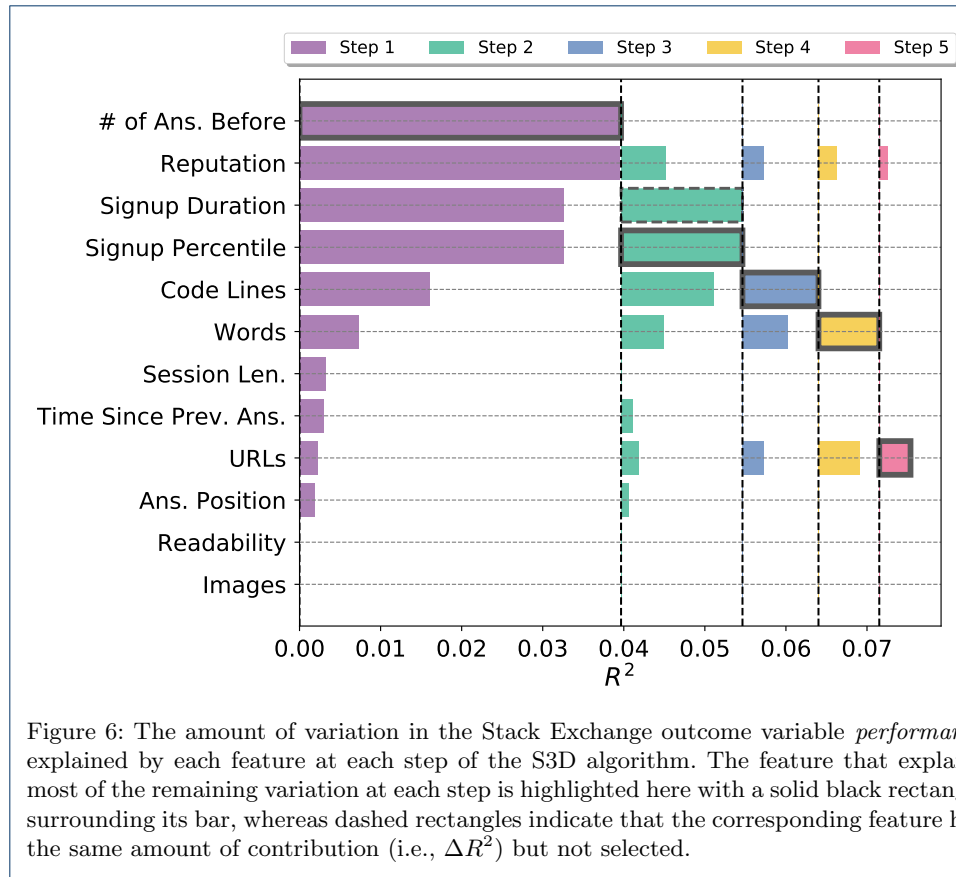
most important feature selected is *the number of answers provided before this question*. This feature, for one thing, indicates how active a user is in the community. For another, it implicitly reflects a user’s capability. Interestingly, there is obvious dependencies between the number of previous answers and (1) *reputation*, (2) *signup duration/percentile*, and (3) *code lines*. Given the amount of previous answers in the model, the contribution of these features decrease dramatically. The second feature S3D selects is *signup percentile*, which measures answerers’ “age” on Stack Exchange as a percentile rank. Intuitively, the longer a user stays in the system, the more likely they can accumulate their reputation and capability to produce a “good” answer. It is noteworthy that *signup duration* and *percentile* share the exact amount of explained variation, which echoes the fact that the Spearman correlation between them is **1**. Following user tenure, the number of *lines of codes* is selected as the third most important feature, followed by the *number of words* and *URLs*, which all, to some extent, manifest how informative an answer is. Note that the variation explained by the features *number of words* and *URLs* exceeds the variation explained by these features in the first step, leading to an interesting implication that there may exist an *interaction* effect. In particular, given answers with the same *number of code lines* and by answerers who signed up in similar time period and shared similar activeness, the *number of words* and *URLs* will contribute more to the final acceptance probability. The ability to identify moderation



effects among variables, in fact, is a fascinating characteristic of S3D when analyzing heterogeneous behavioral data.

With  $R^2 = 0.075$ , the five selected features collectively explain the largest amount of variation in whether an answer is accepted by the asker as the best answer to his or her question. The unselected features have been made redundant by the selected features. Such redundancies can be represented as a directed and weighted network through the coefficients of Equation (10), as shown in Figure 7. Specifically, links between selected features (purple) the unselected (pink) features show the variation in the outcome explained by the pink node can be explained by the purple node. The network visualizes the correlations and the significance of unselected features. While some of the correlations are obvious, such as those between the *number of answers*, *user reputation*, and *tenure length* (i.e., *signup duration/percentile*), others are less evident. For example, there are links from *reputation* to the number of *words*, and *code lines*, implying that reputable users may provide more detailed answers containing informative clues such as references to related webpages and sample codes. Although relatively weak, the link from answer position pointing towards the number of answers a user provided before and *signup percentile* alludes that senior users may tend to be more active and engaging in the community, therefore being early answerers to many questions. The feature network, in this manner, not only lets us analyze which unselected features are explanatory of an outcome



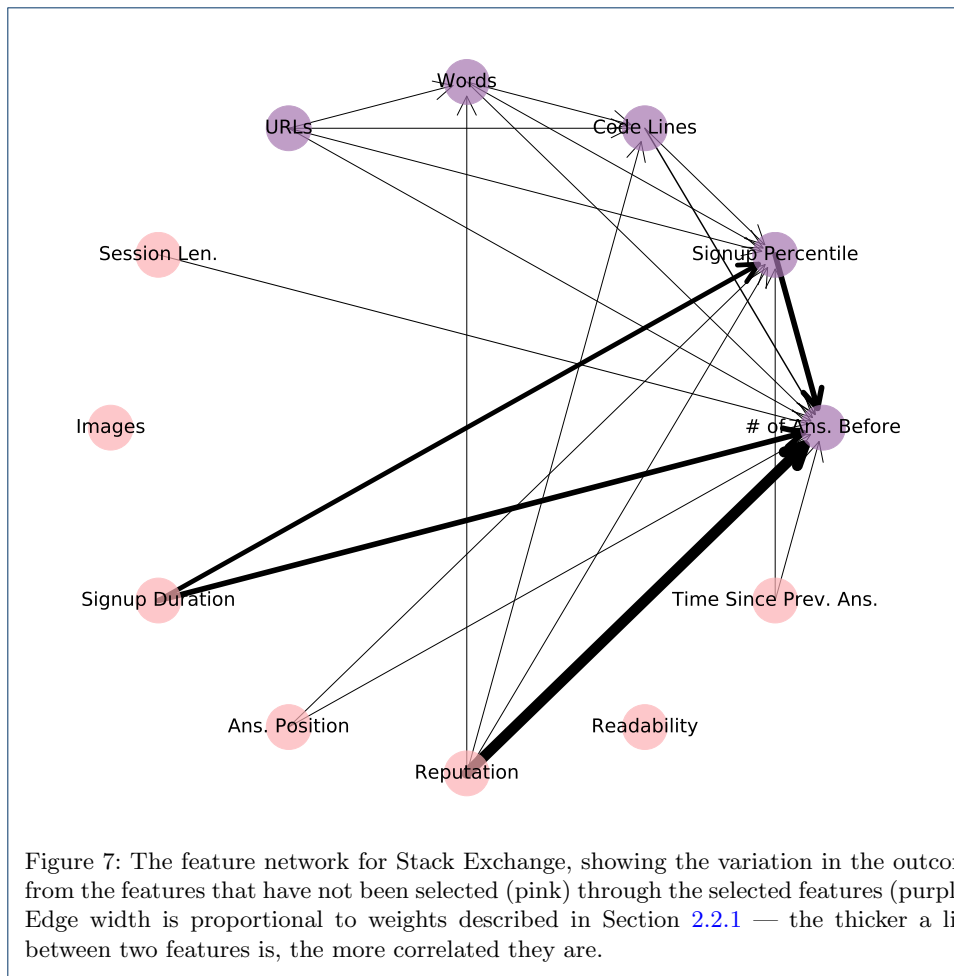


variable, but to which selected features they are correlated and are made redundant, providing a tool to suggest further exploration of correlations within the data.

In the Khan Academy dataset, S3D selects as important features: (1) the *time* it takes the user to solve the problem; (2) the *number of problems* that the user has solved on the first attempt without hints; (3) *time since previous problem*; (4) the number of *first five problems* solved correctly on the first attempt; (5) *index of the session* among all of that user’s sessions; (6) *index of the problem* within its session. It is noteworthy that the second and fourth features here are analogous to *signup duration* and *reputation* on Stack Exchange, as the number of problems that a user solves correctly on their first attempt is a combination of both skill and tenure.

For the Duolingo language learning platform, S3D picks similar skill-based features: (1) the number of *all lessons* completed perfectly; (2) the number of *prior lessons* completed; (3) the number of *distinct words* in the lesson. Similarly, the first feature here is equivalent to the second and fourth feature selected in Khan Academy, which quantifies both skill and tenure.

In the case of Digg social network, to explain whether a user will “digg” (or “like”) a story recommended by friends, S3D selects as important features: (1) user *activity* (how many stories this user recommended); (2) the amount of *information received* by this user from the people she follows; and (3) current popularity of this story. The first two features describe how a user processes and receives information, while the third one reflects how “viral” a story is.

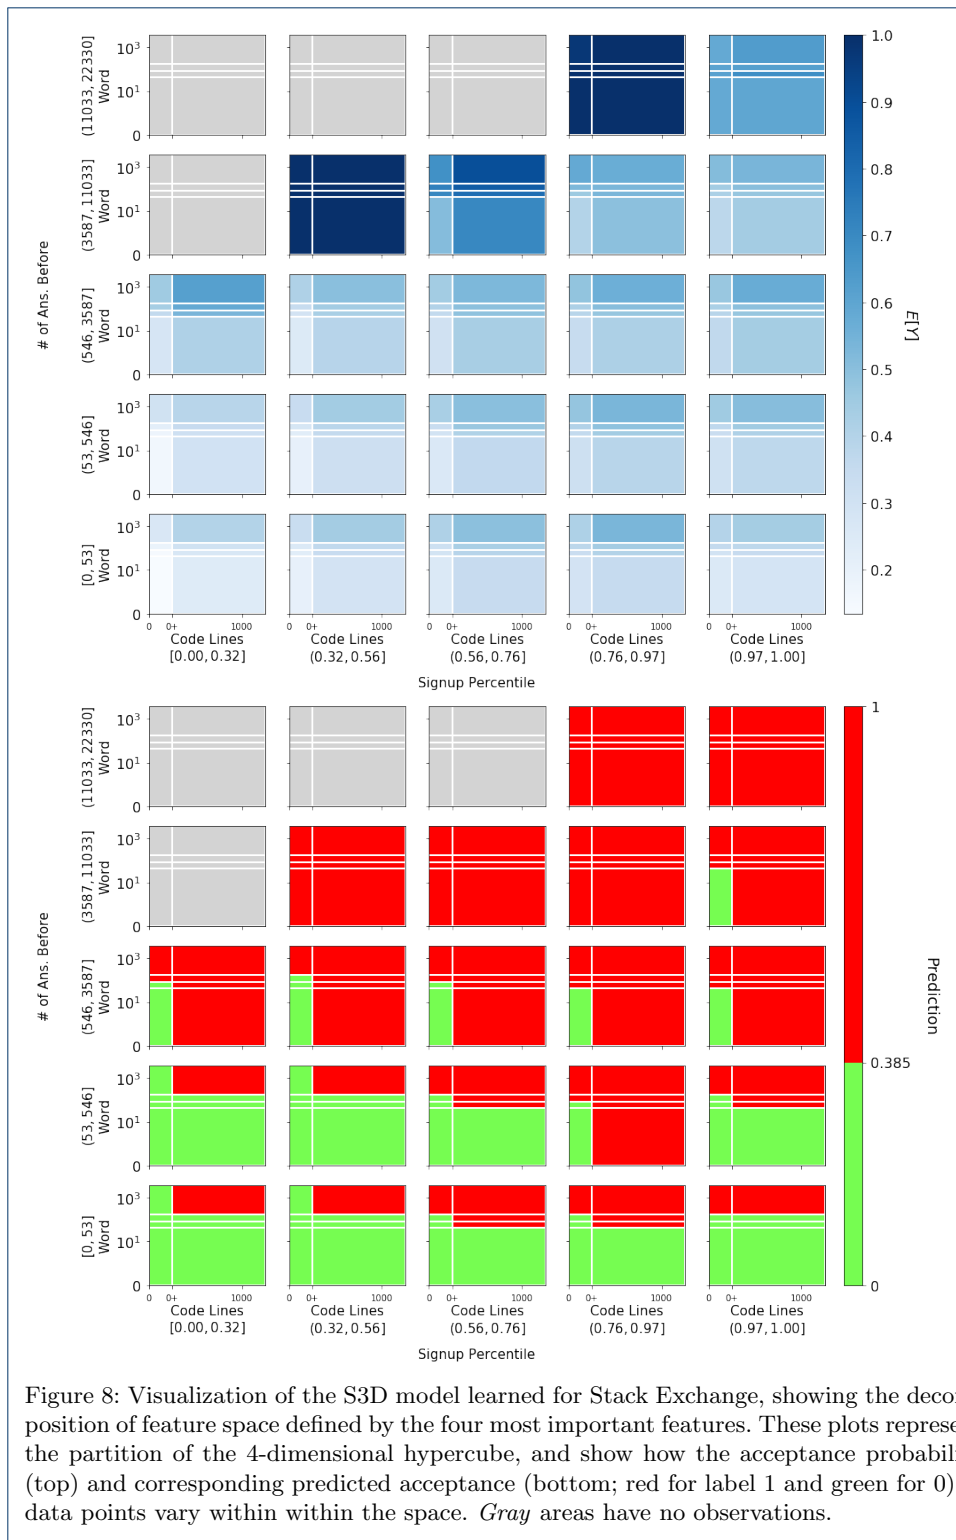


For Twitter, S3D selects (1) the amount of *information received* by this user; (2) *in-degree* (i.e., the number of followers and thus popularity) of friends; (3) the number of *times* this user has been exposed to this meme; (4) user *activity*; (5) the age of this tweet; (6) user’s *out-degree* (followees). S3D identifies the information received by the user as an important feature for both Digg and Twitter, which highlights important role that cognitive load plays in information spread online [29]. On the other hand, the differences such as the lesser importance of user activity and greater importance of a user’s friends in Twitter suggests interesting disparities in the manner of information diffusion on these two platforms.

### 3.3.2 Model Analysis

One of the more novel contributions of S3D is its potential for data exploration. By iteratively selecting features and measuring the amount of outcome variation they explain (Figure 6), we can visualize the important dimensions of the model to fully understand both effects of important features and the corresponding predictions. See Section 7 for a more detailed step-by-step illustration.

For Stack Exchange, S3D selects five important features. We visualized the model with the first four features in Figure 8, that unfolds the  $m = 4$  hypercube learned



by the model. It shows how the expectation (top plot) and the corresponding prediction (bottom plot) that the answer will be accepted as best answer, vary as a function of the four selected features. The prediction threshold selection is described

in Section 2.2.2. Each row of plots in Figure 8 corresponds to a single bin of the first selected feature *number of answers before*, while each column corresponds to bins of the second feature *signup percentile rank*. Individual plots vary according to the third and fourth features *code lines* and *words*. It is quite evident that variation in the outcome (i.e., Figure 8 top plot) is greater between plots than within plots, a result of the fact that features are picked *successively* to explain such variation. These plots show the collective effects of these four features: acceptance *increases* with the user’s experience (*number of answers before* feature) and tenure (*signup percentile rank*). Furthermore, longer answers with more *words* and *lines of code* are more likely to be accepted as best answers. Another interesting pattern emerges when the number of answers provided before is above 3,587: the acceptance rate rises when *signup percentile* goes down. In other words, given a high level of user engagement in the community, newer users tend to produce answers that have higher chances of being accepted. On the other hand, more senior users tend to have a higher probability of having their answers accepted, when the number of previous answers is lower.

We also examine the S3D model learned for Digg to illustrate its effectiveness on highly unbalanced and heterogeneous data. Here, S3D selects as important features *user activity* (i.e., how often a user posts per day), *information received* (the number of stories, or memes, a user’s friends recommend), and *current popularity of the story*, i.e., how many users have recommended it. The model, presented in Figure 9, shows extreme heterogeneity in data with values of features and adoption probabilities varying widely, and notable here is S3D’s ability to learn appropriate binnings of the features over many orders of magnitude. Specifically, the figure shows that the probability of a user to adopt a story *increases* when he/she is more active in the community (see also Figure S5), but *decreases* as users receive more information from friends (see also Figure S6). Specifically, given relatively low activity (e.g., adopting fewer than 505 stories), those users seeing less information from friends are more likely to adopt a new story (corresponding to higher color intensity on the left hand side in each plot). The highly active users, on the other hand, also tend to receive more information—around 1,000 stories per day—from friends. However, they too are less likely to adopt a new story as they receive more information from friends. These two features—*user activity* and *information received*—represent the interplay between information processing and cognitive load. Our analysis highlights the extent of to which information overload, which occurs when users receive more information than they can effectively process, inhibits adoption and spread of memes online [29–31].

The third feature *current popularity* shows the impact of story popularity (i.e., virality or stickiness) on adoption. Our model shows that more popular stories are more likely to be adopted by individuals, as would be expected of viral memes. Striking is the absence of features related to the number or timing of exposures, either as selected features or in the feature network. The exposure effects may be quite subtle or even non-existent. The latter suggests that information on Digg spreads as a simple contagion where the probability of adopting a meme is independent of the number of exposures [32, 33].

The large heterogeneity as a function of basic node features has important implications for the inference of social contagion, because heterogeneity and underlying

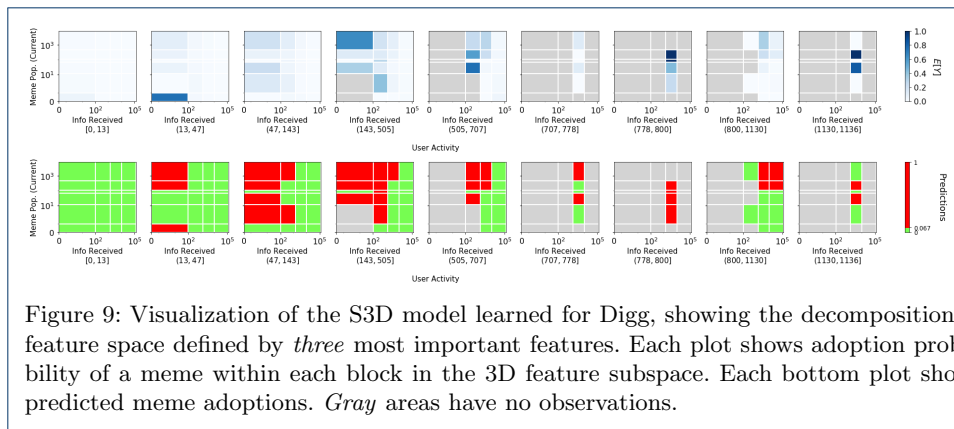


Figure 9: Visualization of the S3D model learned for Digg, showing the decomposition of feature space defined by *three* most important features. Each plot shows adoption probability of a meme within each block in the 3D feature subspace. Each bottom plot shows predicted meme adoptions. *Gray* areas have no observations.

confounders may distort analysis. A possible approach to such inference is to decompose the feature space, as in Figure 9, and statistically test the effect of multiple exposures in the resulting homogeneous blocks, an approach that would ensure that the most important factors that best explain the variation in the adoption of information have been conditioned on.

## 4 Conclusion

We introduced S3D, a statistical model with low complexity but strong predictive power and interpretability, that offers potential to greatly expand the scope of predictive models and machine learning. Our model gives a comprehensive description of the feature space, not only by selecting the most important features, but also by quantifying explained variation in the unselected features. This can be useful for practitioners, who are often concerned with the relationship between specific co-variables and an outcome variable. The decomposed feature space learned by the model presents a powerful tool for exploratory data analysis, and a medium for explaining model predictions. This is a positive step towards transparent algorithms that can be examined for bias, which presents a major stumbling block in the development and application of machine learning.

We have demonstrated the effectiveness of S3D on a variety of datasets, including benchmarks and real-world behavioral data, where it predicts outcomes in the held-out data at levels comparable to state-of-the-art machine learning algorithms. Its potential for interpreting complex behavioral data, however, goes beyond these alternate methods. S3D provides a ranking of features by their importance to explaining data, identifies correlated features, and can be readily used in visual data explorations. Our approach reveals the important factors in explaining human behavior, such as competition, skill, and answer complexity when analyzing performance on Stack Exchange or essential user attributes such as activity and information load in the social networks Digg and Twitter. Our method is self learning, involving minimal tuning of two hyperparameters, while also capable of systematically revealing the rich structure of complex data.

Aside from increasing our understanding of social systems, knowledge about what factors affect behavioral outcomes can also help us design of social platforms that improve human performance, including, for example, optimizing learning on educational platforms [2, 34] or fairer judicial decisions [35]. The insights gained from

the model can help design effective intervention strategies that change behaviors so as to improve individual and collective well-being.

Moving forward, there are two areas where S3D can make a considerable impact. The first is the mathematical modeling of human behavior, [36–38], where S3D can be used as a tool to extract model ingredients from data and help understand their functional form. The second is the use of S3D by practitioners to both explain predictions and analyze interventions based on these predictions. Transparency should be a key requirement for algorithms applied to sensitive areas such as predicting recidivism, and our work here shows that simple algorithms, such as S3D, can not only meet this requirement without sacrificing predictive accuracy. The development of machine learning tools should not be restricted to optimizing one single metric (predictive power), as other ingredients, such as interpretability, can improve how these methods are perceived by society.

### Abbreviations

Structured Sum-of-Squares Decomposition (S3D). Total sum of squares (SST). Support Vector Machine (SVM). Cross validation (CV). Area under the curve (AUC). Random Forest (RF). Question-answering (Q&A).

### Declarations

Availability of data and material

The code and data required for replicating reported results are available here:

<https://github.com/peterfennell/S3D/tree/paper-replication>

### Competing interests

The authors declare that they have no competing interests.

### Funding

This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) and the Army Research Office (ARO) under Contracts No. W911NF-17-C-0094 and No. W911NF-18-C-0011, and in part by the James S. McDonnell Foundation. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of DARPA and the ARO.

### Author's contributions

PF designed and implemented the algorithm; ZZ carried out the validation studies; KL designed the algorithm. All authors participated in writing the paper.

### Acknowledgements

Authors are grateful to Raha Moraffa for her help setting up the cross validation framework and useful suggestions for improving performance.

### Author details

<sup>1</sup>USC Information Sciences Institute, 4676 Admiralty Way, Marina del Rey, CA, USA. <sup>2</sup>University of Iowa, Iowa City, IA, USA.

### References

- Hofman, J.M., Sharma, A., Watts, D.J.: Prediction and explanation in social systems. *Science* **355**(6324), 486–488 (2017). doi:[10.1126/science.aal3856](https://doi.org/10.1126/science.aal3856)
- Settles, B., Meeder, B.: A Trainable Spaced Repetition Model for Language Learning. In: Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 1848–1858. Association for Computational Linguistics, Stroudsburg, PA, USA (2016). doi:[10.18653/v1/P16-1174](https://doi.org/10.18653/v1/P16-1174)
- Thaler, R.H., Sunstein, C.R.: *Nudge: Improving Decisions About Health, Wealth, and Happiness*, Rev and expanded edn. Penguin Books, New York (2009)
- Lazer, D., Pentland, A., Adamic, L., Aral, S., Barabasi, A.-L., Brewer, D., Christakis, N., Contractor, N., Fowler, J., Gutmann, M., Jebara, T., King, G., Macy, M., Roy, D., Van Alstyne, M.: Computational Social Science. *Science* **323**(5915), 721–723 (2009). doi:[10.1126/science.1167742](https://doi.org/10.1126/science.1167742)
- Lipton, Z.C.: The Mythos of Model Interpretability. *ACM Queue* **16**(3), 30 (2018). doi:[10.1145/3236386.3241340](https://doi.org/10.1145/3236386.3241340)

6. Hofman, J.M., Sharma, A., Watts, D.J.: Prediction and explanation in social systems. *Science* **488**(February), 486–488 (2017). doi:[10.1126/science.aal3856](https://doi.org/10.1126/science.aal3856)
7. Breiman, L., Friedman, J.H., Olshen, R.A., Stone, C.J.: *Classification and Regression Trees*, 1st edn. Wadsworth Publishing, New York (1984)
8. Blyth, C.R.: On Simpson's Paradox and the Sure-Thing Principle. *Journal of the American Statistical Association* **67**(338), 364–366 (1972). doi:[10.1080/01621459.1972.10482387](https://doi.org/10.1080/01621459.1972.10482387)
9. Alipourfard, N., Fennell, P.G., Lerman, K.: Can you Trust the Trend: Discovering Simpson's Paradoxes in Social Data. In: *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining - WSDM '18*, pp. 19–27. ACM Press, New York, New York, USA (2018). doi:[10.1145/3159652.3159684](https://doi.org/10.1145/3159652.3159684). [1801.04385](https://doi.org/10.1145/3159684.1801.04385)
10. Hastie, T., Tibshirani, R., Friedman, J.: *The Elements of Statistical Learning*, 2nd edn. Springer Series in Statistics. Springer, New York, NY (2009). doi:[10.1007/978-0-387-84858-7](https://doi.org/10.1007/978-0-387-84858-7). [1010.3003](https://doi.org/10.1007/978-0-387-84858-7)
11. Zou, H., Hastie, T.: Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **67**(2), 301–320 (2005). doi:[10.1111/j.1467-9868.2005.00503.x](https://doi.org/10.1111/j.1467-9868.2005.00503.x)
12. Chipman, H.A., George, E.I., McCulloch, R.E.: BART: Bayesian additive regression trees. *The Annals of Applied Statistics* **4**(1), 266–298 (2010). doi:[10.1214/09-AOAS285](https://doi.org/10.1214/09-AOAS285). [0806.3286](https://doi.org/10.1214/09-AOAS285)
13. Friedman, J.: Multivariate Adaptive Regression Splines. *The Annals of Statistics* **19**(1), 1–67 (1991). doi:[10.2307/2241837](https://doi.org/10.2307/2241837)
14. Breiman, L.: Random forests. *Machine Learning* **45**(1), 5–32 (2001). doi:[10.1023/A:1010933404324](https://doi.org/10.1023/A:1010933404324)
15. Boser, B.E., Guyon, I.M., Vapnik, V.N.: A training algorithm for optimal margin classifiers. In: *Proceedings of the Fifth Annual Workshop on Computational Learning Theory - COLT '92*, pp. 144–152. ACM Press, New York, New York, USA (1992). doi:[10.1145/130385.130401](https://doi.org/10.1145/130385.130401)
16. Fernández-Delgado, M., Cernadas, E., Barro, S., Amorim, D., Fernández-Delgado, A.: Do we Need Hundreds of Classifiers to Solve Real World Classification Problems? *Journal of Machine Learning Research* **15**, 3133–3181 (2014)
17. Chang, Y.-W., Hsieh, C.-J., Chang, K.-W., Ringgaard, M., Lin, C.-J.: Training and Testing Low-degree Polynomial Data Mappings via Linear SVM. *Journal of Machine Learning Research* **11**(Apr), 1471–1490 (2010)
18. Dheeru, D., Karra Taniskidou, E.: {UCI} Machine Learning Repository (2017). <http://archive.ics.uci.edu/ml>
19. Street, W.N., Wolberg, W.H., Mangasarian, O.L.: Nuclear feature extraction for breast tumor diagnosis. *ISandT/SPIE International Symposium on Electronic Imaging: Science and Technology* **1905**, 861–870 (1993). doi:[10.1117/12.148698](https://doi.org/10.1117/12.148698)
20. Kurgan, L.A., Cios, K.J., Tadeusiewicz, R., Ogiela, M., Goodenday, L.S.: Knowledge discovery approach to automated cardiac SPECT diagnosis. *Artificial Intelligence in Medicine* **23**(2), 149–169 (2001). doi:[10.1016/S0933-3657\(01\)00082-3](https://doi.org/10.1016/S0933-3657(01)00082-3)
21. Little, M.A., McSharry, P.E., Roberts, S.J., Costello, D.A.E., Moroz, I.M.: Exploiting nonlinear recurrence and fractal scaling properties for voice disorder detection. *BioMedical Engineering Online* **6**(1), 23 (2007). doi:[10.1186/1475-925X-6-23](https://doi.org/10.1186/1475-925X-6-23). [0707.0086](https://doi.org/10.1186/1475-925X-6-23)
22. Candanedo, L.M., Feldheim, V., Deramaix, D.: Data driven prediction models of energy use of appliances in a low-energy house. *Energy and Buildings* **140**, 81–97 (2017). doi:[10.1016/j.enbuild.2017.01.083](https://doi.org/10.1016/j.enbuild.2017.01.083)
23. Rafiei, M.H., Adeli, H.: A Novel Machine Learning Model for Estimation of Sale Prices of Real Estate Units. *Journal of Construction Engineering and Management* **142**(2), 04015066 (2016). doi:[10.1061/\(ASCE\)CO.1943-7862.0001047](https://doi.org/10.1061/(ASCE)CO.1943-7862.0001047)
24. Weiss, S.M., Indurkha, N.: Rule-based Machine Learning Methods for Functional Prediction. *Journal of Artificial Intelligence Research* **3**(1995), 383–403 (1995). doi:[10.1613/jair.199.9512107](https://doi.org/10.1613/jair.199.9512107)
25. Harrison, D., Rubinfeld, D.L.: Hedonic housing prices and the demand for clean air. *Journal of Environmental Economics and Management* **5**(1), 81–102 (1978). doi:[10.1016/0095-0696\(78\)90006-2](https://doi.org/10.1016/0095-0696(78)90006-2)
26. King, R.D., Hirst, J.D., Sternberg, M.J.E.: Comparison of artificial intelligence methods for modeling pharmaceutical QSARS. *Applied Artificial Intelligence* **9**(2), 213–233 (1995). doi:[10.1080/08839519508945474](https://doi.org/10.1080/08839519508945474)
27. Little, M.A., McSharry, P.E., Hunter, E.J., Spielman, J., Ramig, L.O.: Suitability of Dysphonia Measurements for Telemonitoring of Parkinson's Disease. *IEEE Transactions on Biomedical Engineering* **56**(4), 1015–1022 (2009). doi:[10.1109/TBME.2008.2005954](https://doi.org/10.1109/TBME.2008.2005954)
28. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, É.: Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* **12**(Oct), 2825–2830 (2011)
29. Hodas, N.O., Lerman, K.: How visibility and divided attention constrain social contagion. In: *Proceedings - 2012 ASE/IEEE International Conference on Privacy, Security, Risk and Trust and 2012 ASE/IEEE International Conference on Social Computing, SocialCom/PASSAT 2012*, pp. 249–257. IEEE, ??? (2012). doi:[10.1109/SocialCom-PASSAT.2012.129](https://doi.org/10.1109/SocialCom-PASSAT.2012.129)
30. Greg Ver Steeg, Ghosh, R., Lerman, K.: What stops social epidemics? In: *Proceedings of 5th International Conference on Weblogs and Social Media* (2011)
31. Gomez-Rodriguez, M., Gummadi, K.P., Schölkopf, B.: Quantifying Information Overload in Social Media and Its Impact on Social Contagions. In: *Proceedings of the 8th International Conference on Weblogs and Social Media, ICWSM 2014*, pp. 170–179 (2014)
32. Centola, D., Eguíluz, V.M., Macy, M.W.: Cascade dynamics of complex propagation. *Physica A: Statistical Mechanics and its Applications* **374**(1), 449–456 (2007). doi:[10.1016/j.physa.2006.06.018](https://doi.org/10.1016/j.physa.2006.06.018). [0504165](https://doi.org/10.1016/j.physa.2006.06.018)
33. Hodas, N.O., Lerman, K.: The simple rules of social contagion. *Scientific Reports* **4**(1), 4343 (2014). doi:[10.1038/srep04343](https://doi.org/10.1038/srep04343). [1308.5015](https://doi.org/10.1038/srep04343)
34. Rendle, S.: Factorization Machines with libFM. *ACM Transactions on Intelligent Systems and Technology* **3**(3), 1–22 (2012). doi:[10.1145/2168752.2168771](https://doi.org/10.1145/2168752.2168771)
35. Kleinberg, J., Lakkaraju, H., Leskovec, J., Ludwig, J., Mullainathan, S.: Human Decisions and Machine Predictions. *The Quarterly Journal of Economics* **133**(1), 237–293 (2017). doi:[10.1093/qje/qjx032](https://doi.org/10.1093/qje/qjx032)
36. Watts, D.J.: A simple model of global cascades on random networks. *Proceedings of the National Academy of*

Sciences **99**(9), 5766–5771 (2002). doi:[10.1073/pnas.082090499](https://doi.org/10.1073/pnas.082090499)

37. Fernández-Gracia, J., Suchecki, K., Ramasco, J.J., San Miguel, M., Eguíluz, V.M.: Is the Voter Model a Model for Voters? *Physical Review Letters* **112**(15), 158701 (2014). doi:[10.1103/PhysRevLett.112.158701](https://doi.org/10.1103/PhysRevLett.112.158701)

38. O’Sullivan, D.J.P., O’Keeffe, G.J., Fennell, P.G., Gleeson, J.P.: Mathematical modeling of complex contagion on clustered networks. *Frontiers in Physics* **3**, 71 (2015). doi:[10.3389/fphy.2015.00071](https://doi.org/10.3389/fphy.2015.00071)

**5 Appendix**

**6 Data**

We used 9 datasets for classification and regression tasks respectively (?? in the paper). Here we present a more detailed description of the data sources and preprocessing.

**6.1 Data Preprocessing**

We downloaded the original data files from UCI Machine Learning Repository [18] and Luís Torgo’s personal website[8]. In addition, we conducted further data cleaning before the experiments[9]:

- *Classification:* (1) For breast cancer (original), We dropped sample code number columns and removed samples with missing values (all in feature “Bare Nuclei”). (2) For parkinsons, we dropped ASCII subject name and recording number.
- *Regression:* (1) For appliances energy use, we dropped date columns; (2) For residential building, we dropped sales column when the target is costs and vice versa; (3) For parkinsons, we dropped columns subject number, age, sex, and test time, which are individual subject information not used for prediction in the original paper [27]. Finally, we note that log transformation of the target values were applied to appliances energy use and both residential building to reduce skewness. Specifically, we used  $\log_2(x + 1)$  to avoid the logarithm of zero values.

The five remaining human behavior datasets are thoroughly described in Section 3 of the paper.

**6.2 Cross Validation**

To evaluate prediction performance, we applied *nested cross validation (CV)* on all datasets (Figure 10). Specifically, we split each dataset into *five* equal-size folds (a.k.a., outer folds; 1 to 5 in Figure 10). Each of the five was picked as test set (i.e., held out from the training and validation process). Given each outer test set (fold 1 in the example), we conducted *inner CV* to find out the best hyperparameters for each predictive model based on the average performance on all four folds (fold 2, 3, 4, and 5.) During *inner CV*, we conducted 4-fold CV, where each of the four folds were held out as validation set (fold 2) and the rest (fold 3, 4, and 5) were inputs to the model. Different sets of *outer CV* may produce different “best” hyperparameters. The final prediction performance is the average value of those on each test set.

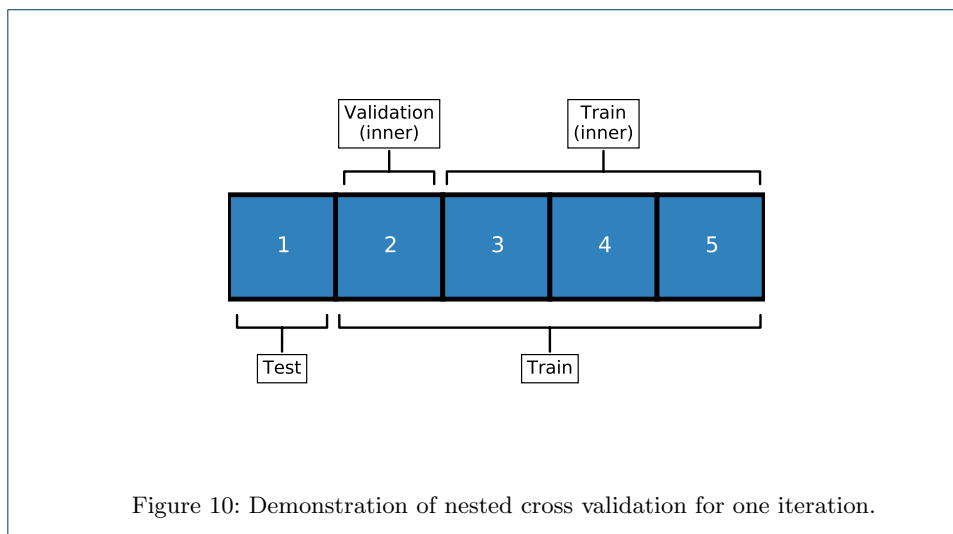


Figure 10: Demonstration of nested cross validation for one iteration.

**6.3 Data Standardization**

When applying logistic regressions [10, page 63] and linear support vector machine (SVM) [? ], we standardized all features (i.e., subtracting the mean and scaling to unit variance). For all regression data, regardless of regressors, we standardized all columns (i.e., features as well as targets). Note that in both cases, we transformed values in the test set based on the mean and variance of values in the training set.

[8] <https://www.dcc.fc.up.pt/~ltorgo/Regression/DataSets.html>

[9] <https://github.com/peterfennell/S3D/blob/paper-replication/data/download-datasets.ipynb>



## 7 Visualizing Data with S3D

One of the strengths of S3D is its ability to create visualizations of learned models for data exploration. In this section, we provide a detailed description of S3D visualizations of *Stack Exchange* and *Digg* data, starting from the simplest models that partition the data on the most important feature, to the increasingly more complex models that partition the feature space of several most important features.

### 7.1 Stack Exchange

The first important feature in Stack Exchange data, *the number of answers before*, reflects users' experience. Intuitively, the more answers a user has written in the past, the more likely that this user is to be active, knowledgeable, and experienced; therefore, the more likely the current answer is to be accepted as best answer by the asker. This is indeed reflected by the model learned by S3D, shown in Figure 11. Evidently, acceptance probability increases monotonically with user experience (Figure 11(a)). On the other hand, most users provide fewer than 1,000 answers. While exceptional activity leads to remarkable acceptance rate, these users are rare in the community (Figure 11(b)).

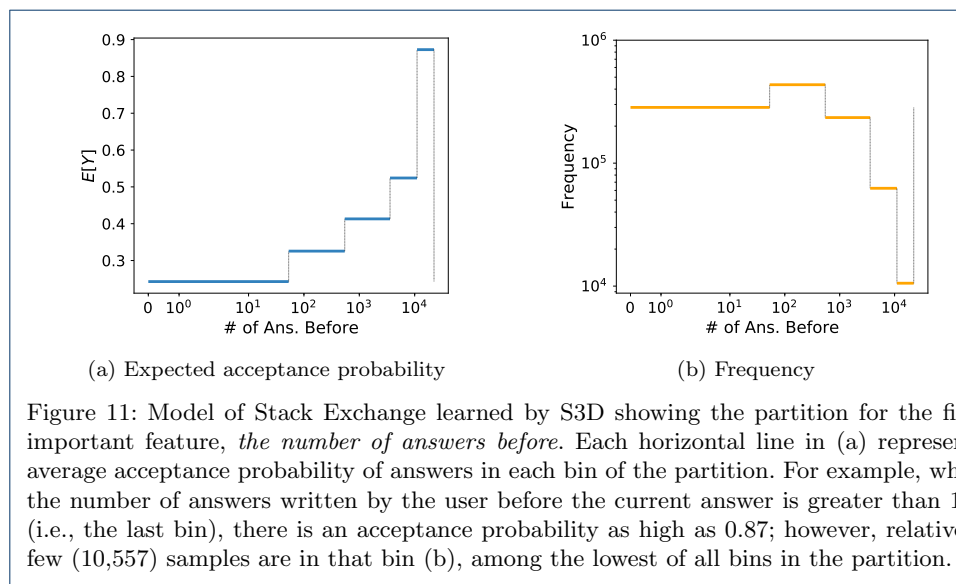


Figure 11: Model of Stack Exchange learned by S3D showing the partition for the first important feature, *the number of answers before*. Each horizontal line in (a) represents average acceptance probability of answers in each bin of the partition. For example, when the number of answers written by the user before the current answer is greater than  $10^4$  (i.e., the last bin), there is an acceptance probability as high as 0.87; however, relatively few (10,557) samples are in that bin (b), among the lowest of all bins in the partition.

In the second step, S3D included *signup percentile* as an additional feature to explain acceptance probability (Figure 12). This feature represents user's rank by length of tenure among all other answerers. Overall, we can see an increasing trend of acceptance probability from the bottom left corner to top right, implying a positive relationship where more senior and experienced users are more likely to get their answers accepted as best answers. Indeed, we can see from Figure 12b that senior users are usually those who have written more answers than newcomers. In the mean time, acceptance probability goes down when signup percentile is lower, given the number of answers beyond 3,587. This interesting pattern indicates that answers provided by highly productive users, newcomers, rather than seniors, are more likely to get their answers accepted as best answers.

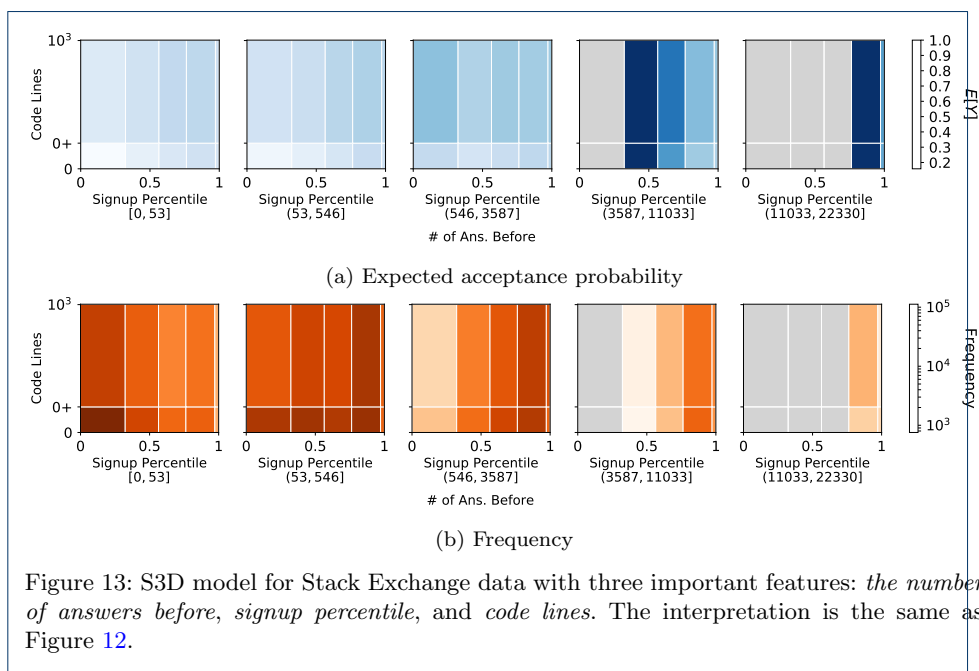
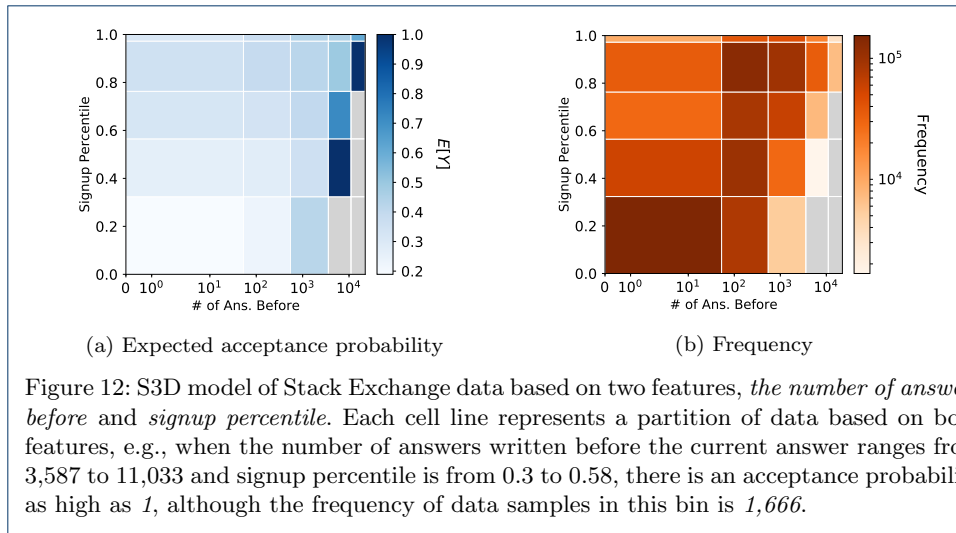
In Figure 13 we proceed to look at models that include the third important feature: *code lines*. Each plot in Figure 13(a) shows a bin with respect to the first important feature (number of answers written before this). The range of the bin is shown just under each plot. Note that in each plot there are two zeros on the y-axis. We manually expanded the range  $[0, 0]$  from a line to a band to make the binning more clear. In fact, S3D learns from the data that answers that include programming code have a better chance of being accepted (i.e., the color in Figure 13a is *less blue* when the number of code lines is above zero), although most answers do not contain code (Figure 13b). The model with four features can be found in the main paper in Figure 8.

Finally, Figure 14a shows the S3D model with the fourth important feature, *word*, which is the number of words in an answer. Generally, the longer—and presumably the more informative—the answer, the more likely it is to get accepted. Meanwhile, the distribution of words among answers mostly stayed in the bottom half (i.e., below 178 word counts) as shown in Figure 14b. See Section 3.3.2 for an in-depth discussion.

### 7.2 Digg

The first important feature identified by S3D in the Digg data is *user activity* (Figure 15). Recall that the target variable is whether or not a user will “digg” a story a story following an exposure by a friend. Digging a story is similar to retweeting a post on Twitter, hence, we generically refer to it as “adopting” a story or a meme. While, intuitively, more active users are more likely to “digg” stories, the increment as shown in Figure 15a is generally very small, implying the existence of a more complex behavioral mechanism of meme adoption.

In the next step, S3D finds that *information received by users* is an important feature (Figure 16), with which S3D discovers a more fine-grained partition of data. This feature describes the user's information load, that is the number of stories “dugg” or recommended by friends. Probability to “digg” increases when going from top left to bottom right. Indeed, active users who have lower information load are more likely to “digg” the stories they receive.



Nonetheless, most users are not active but receive a huge load of information from their neighbors (Figure 16b). See Figure 9 in the paper for the model with three features.

The last feature that S3D picked is *the current popularity of this meme*, which measures how “viral” a story is (Figure 17a). See Section 3.3.2 for a more detailed description on the learned model. Overall, users tend to “digg” a popular story, implying the *Matthew effect* on diffusion of Digg stories. In Figure 17b, it is interesting to see that most users are exposed to viral stories.

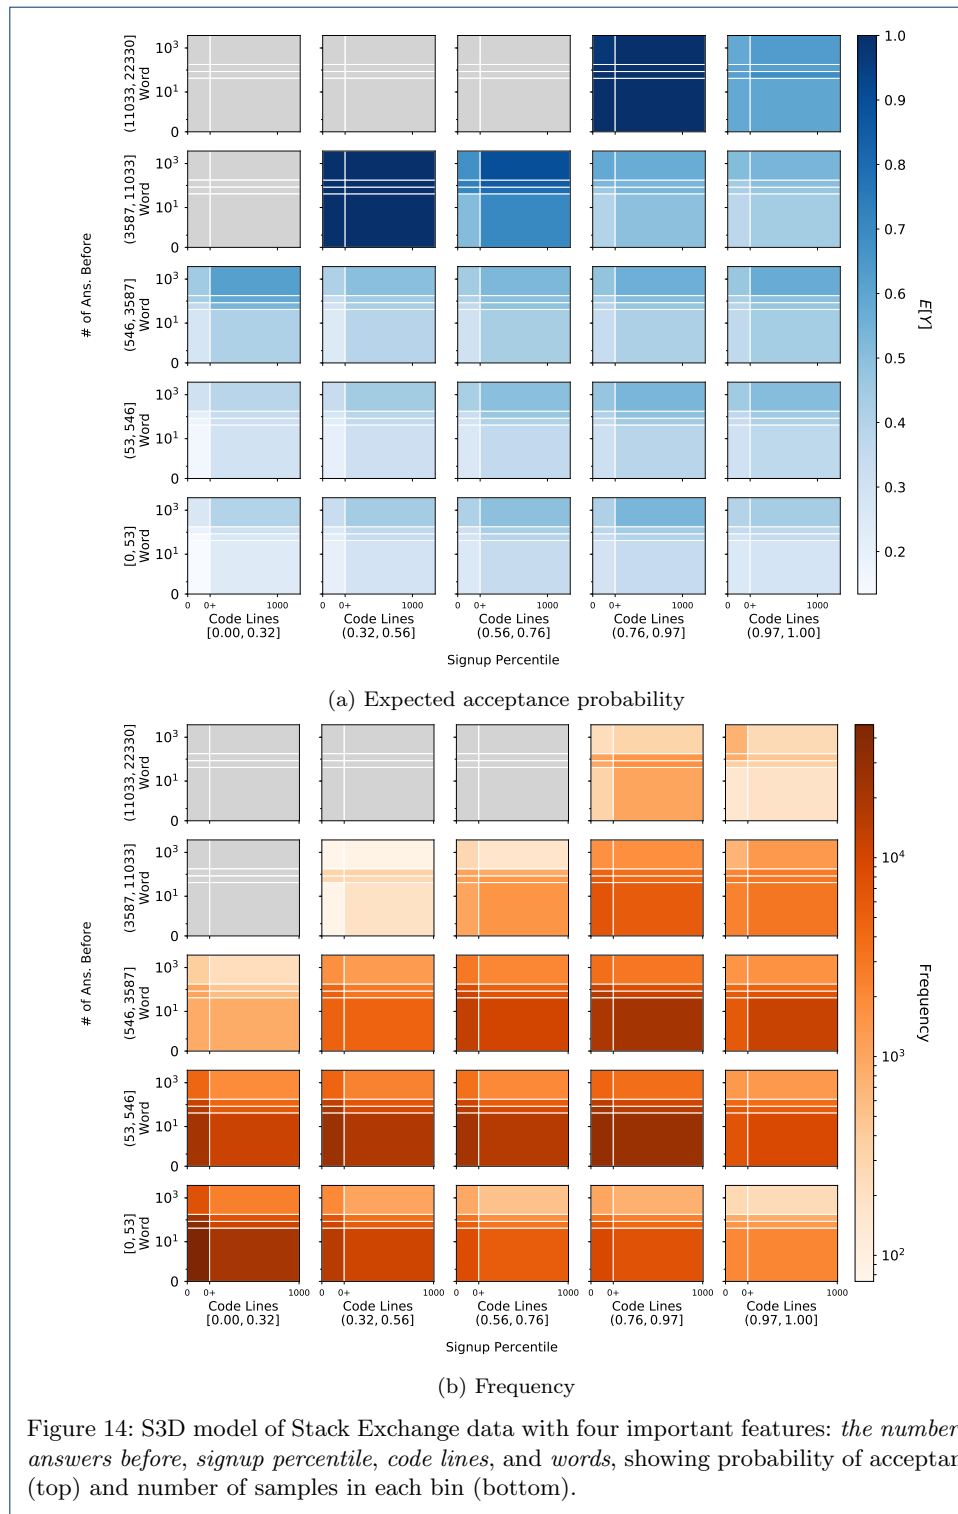


Figure 14: S3D model of Stack Exchange data with four important features: *the number of answers before*, *signup percentile*, *code lines*, and *words*, showing probability of acceptance (top) and number of samples in each bin (bottom).

