Topology and Geometry in Machine Learning for Logistic Regression

Colleen M. Farrelly, Independent Researcher (cfarrelly@med.miami.edu)

Abstract

Logistic regression plays an important role in medical research, and several machine learning extensions exist for this framework, including least angle regression (LARS) and least absolute shrinkage and selection operator (LASSO), which yield models with interpretable regression coefficients. Many machine learning algorithms have benefitted in the past few years from the inclusion of geometric and topological information, including manifold learning, shape-matching, and supervised learning extensions of generalized linear regression. This study demonstrates gains from the inclusion of differential geometric information in LARS models and of homotopy search in LASSO models above that of elastic net regression, a state-of-the-art penalized regression algorithm. Results hold across both simulated data and two real datasets, one predicting alcoholism risk and one predicting tumor malignancy. These algorithms also perform competitively with classification algorithms such as random forest and boosted regression, suggesting that machine learning methods which incorporate topological/geometric information about the underlying data may be useful on binary classification datasets within medical research. In addition, other hybrid techniques may outperform existing methods and provide more accurate models to understand disease. More work is needed to develop effective, efficient algorithms that explore the topology or geometry of data space and provide interpretable models.

Keywords: logistic regression, differential geometry, homotopy, machine learning, breast cancer

Introduction

Logistic regression is ubiquitous in medical and psychological research^{1,2,3}. However, logistic regression suffers with the problems of high dimensional data ($p \gg n$), sparsity, and collinearity--among others--which are common issues within medical datasets^{4,5}. Machine learning offers many extensions of logistic regression to these problems, either outputting a linear model with interpretable coefficients or outputting a model with ranked variable importance^{6,7,8}. These algorithms, such as penalized regression and its derivatives, random forest/tree-based methods, and superlearner ensembles that combine multiple models into a cohesive whole, have proven valuable in recent years and have been widely used in medical research from psychology to neuroscience to genomics to imaging^{9,10,11,12,13,14,15}.

In particular, machine learning methods based on penalizing a generalized linear model, such as logistic regression, have solved many problems inherent in medical datasets, including sparsity, high dimensionality ($n < p$), presence of outliers, and collinearity⁷. One advantage of these models is their basis in generalized linear models¹⁶, such that they return interpretable regression models from which information about the magnitude and directionality of effect can be estimated for each main effect/interaction term chosen by the model⁷. Some of the more popular medical research machine learning algorithms of this type include least angle regression (LARS)¹⁷, least absolute shrinkage and selection operator (LASSO)^{9,10}, and elastic net regression⁷. Many of these algorithms also have established convergence properties, which are not directly established or theoretically derived for other types of machine learning algorithms.

Prior attempts to incorporate geometric or topological information have improved machine learning capabilities across a wide variety of problems. Ricci curvature and its discrete analogue, Forman curvature, have added to the understanding of network properties and growth^{18,19}; topological metrics (such as Hausdorff distance) and clustering have also enhanced network analytic capabilities in recent years^{20,21,22}. The Hodge-Hemholtz decomposition has been applied to many ranking problems, particularly with sparse response data, and has theoretical connections to LASSO²³. Manifold learning has benefitted from Ricci-flow-based smoothing approaches, as well as geometrically-aware mapping methods^{24,25,26,27}. Lie group rotations²⁸, deconvolution techniques²⁹, and topological metrics/decompositions^{30,31} have improved algorithm performance. On matching problems and image analytics, particularly within protein/DNA

and neuroscience data. Recent topological and geometric contributions to supervised machine learning include piecewise regression based on Morse-Smale complexes³², homotopy-based parameter search³³ and outlier detection³⁴, neural network visualizations based on Morse theory³⁵, manifold-based random forest algorithms³⁶, and curvature-based reformulation of exponential families related to generalized linear modeling³⁷. Within unsupervised learning, techniques such as the mapper algorithm for topologically-based clustering^{38,39,40}, Morse-Smale decomposition³⁴, Reeb space algorithms⁴¹, nearest-neighbor-based storage for large k-mean problems⁴², and persistent homology matching/decomposition^{43,44} have found success in both extending statistical capabilities and finding meaningful results within medical data.

Despite these recent advances, very little work has been done to test such advances against other algorithms common in medical research to understand the nature and magnitude of gains achieved by incorporating geometry/topology into supervised learning. However, it seems possible that including topological or geometric information/methodology in penalized regression models may enhance model accuracy and capture more features in the data than is currently feasible. Further, methods that have the flexibility to explore the data space while creating a generalized linear model (such as boosted regression with linear base learners or multivariate adaptive regression splines, or MARS) may be able to leverage this information implicitly. Two recent extensions of penalized regression models include a geometric or topological component-homotopy-based LASSO⁴⁵ and a LARS algorithm leveraging the data's tangent space¹⁷--and these are compared with other penalized regression/machine learning extensions of logistic regression to assess the usefulness of incorporating topology/geometry into the model across several simulated datasets, a common medical data benchmark set, the UCI repository's Breast Cancer Wisconsin Dataset (BCWD), and the Collaborative Studies on Genetics of Alcoholism (COGA) demographic data for Caucasian participants. Results from BCWD and COGA data are also compared with machine learning methods that do not yield a generalized linear regression output, such as random forest and extreme gradient boosting, to assess relative abilities of the geometry/topology-based penalized regression models to nonparametric methods.

Methods

I) Overview of Algorithms

The main algorithms used in comparisons are presented both technically and through analogies, such that these algorithms are presented in sufficient detail for mathematically-inclined readers but accessible to others. This covers algorithm intuition, as well as statistical/computational nuances that impact algorithm usage and performance.

a) Logistic regression is an extension of multivariate regression to binary outcomes through the use of a link function, typically the logit, probit, or cloglog links¹⁶. For a logit link, the multivariate linear regression function

$$Y = X\beta + \varepsilon$$

(where β is a vector of coefficients that minimize the error of predicted outcome values and ε is a vector of normally-distributed errors) is extended to:

$$\ln\left(\frac{Y}{(1-Y)}\right) = X\beta + \varepsilon$$

where the logit function on the left-hand side refers to the log odds; thus, exponentiating the right-hand regression equation gives the log odds of the outcome occurring given the predictors observed for that individual. This allows one to derive odds ratios from each term's coefficient. As such, it falls under the exponential family of generalized linear models, and it inherits the limitations of this framework, including sample size requirements, assumptions of the independence and linearity of predictors, and assumptions of error term independence¹⁶. In many datasets, linearity/independence of predictors and sample size may pose issues to this method, necessitating the use of more flexible machine learning methods⁴.

b) Penalized linear models, such as elastic net and LASSO⁴⁵, extend logistic regression (as well as other generalized linear models) to samples in which predictors may outnumber observations ($p > n$). LASSO adds a penalty to the original least squares estimator to accomplish this:

$$\min_{\beta} ||\mathbf{y} - \mathbf{X}\beta|| \quad \text{s.t.} \quad ||\beta||_1 = \sum_{j=1}^p |\beta_j| \leq t$$

where p is the number of predictors and t is the regularization parameter, such that predictors with an estimated coefficient less than t are set to 0. Geometrically, this can be thought of as a cowboy standing at the origin casting his lasso rope of a given length to capture any predictors that fall within the radius of his cast, thus removing them from the set. Elastic net⁷ extends this ℓ_1 penalty to ℓ_1 and ℓ_2 penalties, such that the equation becomes:

$$\min_{\beta} ||\mathbf{y} - \mathbf{X}\beta|| \quad \text{s.t.} \quad J \leq t$$

where

$$J(\beta) = \alpha ||\beta||^2 + (1 - \alpha) ||\beta||_1$$

Thus, elastic net creates sparse models (ℓ_1 penalty) that also have a unique minimum as a result of convexity constraints (ℓ_2 penalty), combining the best of both types of penalties. Essentially, elastic net combines LASSO with a ridge regression to provide robust performance with imposed shrinkage, allowing the algorithm to handle $p > n$ situations, as well as situations where predictors are correlated with each other or exhibit deviations from linearity⁷. An analogy is clearing the area around the cowboy so that objects (posts, trees...) are removed from his field of vision, providing easier roping of variables near the origin.

c) One LASSO extension, homotopy LASSO⁴⁵, employs a topologically-based search for model coefficient estimation, which leverages properties of the Lagrange multiplier related to the ℓ_1 penalty, yielding an ordinary differential equation (ODE) solution. This solution is piecewise linear with a differentiable constraint⁴⁵. These have a known solution that involves using a known equation's solution and deforming it to a target equation solution that is unknown⁴⁵; methods employing this tactic are called homotopy methods.

Homotopy is an intrinsic topological property concerned with path equivalence between points, where path equivalence implies that they can be continuously deformed to one another (figure 1); the homotopy method has logged much success within ODE solutions recently⁴⁵. In figure 1, all paths on the sphere can be continuously deformed to each other; on the torus (the donut-shaped space), the hole prevents some paths from being deformed into each other. Thus, homotopy search avoids geometric pitfalls that might exist in the data spaces, which trap other algorithms, by searching for regression parameters through the use of easier predictor paths and deformation according to the data. Analogously, this is a bit like blindfolding someone and tasking them with finding a target guided by a rope connected to the target; he would likely be blocked by obstacles (temporarily or permanently).

d) Much like LASSO models, LARS models utilize the ℓ_1 norm to build a generalized linear regression model, specifically designed to introduce sparsity for datasets where $p \gg n$ ¹⁷. All coefficients are initially set to 0, and predictors are added stage-wise according to highest correlation with a given outcome. Coefficient estimates are increased in the direction of their correlation until another predictor has a greater correlation with the outcome. As coefficients are introduced, estimates are jointly increased through their joint least squares direction until the next coefficient is introduced. This process continues until all predictors are entered into the model; the optimal solution is then chosen according to a preset criterion, such as BIC or AIC of the models. As such, it is similar to forward regression selection schemes, only with an added penalty to drop terms from the model if their estimates are near zero.

A recent extension of LARS, DGLARS¹⁷, leverages the differential geometry of the data space (called the error tangent space) to scale the generalized linear model score function at each update by the square root of the conditional Fisher information. This gives the conditional Rao's score statistic:

$$r_j^u(\beta) = \frac{\partial_j \ell(\beta; \mathbf{y})}{\sqrt{I_j(\beta)}}$$

where $I_j(\boldsymbol{\beta})$ is the Fisher information ($I_j(\boldsymbol{\beta}) = E(\partial_j \ell(\boldsymbol{\beta}; \mathbf{y})^2)$) and $\partial_j \ell(\boldsymbol{\beta}; \mathbf{y})$ refers to the j^{th} directional derivative of the log-likelihood. If the derivative of a log-likelihood equation for a coefficient is 0, then this score statistic will equal 0, and it will be dropped from the solution equation, forming a set of non-selected model coefficients. Non-zero score statistics can be ranked by absolute value in a separate set of model coefficients.

Geometrically, this involves searching coefficient vectors with respect to the residual vector's tangent space, whereby predictors are added to the model according to the smallest angle relative to the residual vector tangent space (figure 2, as shown in⁷). Equiangularity ensures that collinear predictors are not added to an existing model, and the existing LARS ℓ_1 penalty ensures sparsity to deal with large numbers of predictors relative to observations. In essence, the geometry of the loglikelihood space partitions predictors into 3 sets: selected predictors (which correspond to good fits to a uniquely defined error tangent space), redundant predictors (which share a tangent space with already-selected predictors), and non-selected predictors (which have relatively large angles in the tangent space and, thus, do not represent good relationships to an outcome). This is analogous to a geologist sorting rock samples into different bags (igneous, metamorphic, and sedimentary) based on the visible properties of a given rock.

e) Multivariate adaptive regression splines (MARS) models⁵ build generalized linear models, such as logistic regression, by creating a weighted sum of basis functions ($B_j(x)$):

$$\mathbf{y} = \sum_{j=1}^p c_j B_j(\mathbf{x})$$

where each c_j indicates an optimized weight given to the basis function, $B_j(\mathbf{x})$. Basis functions allow for a wiggling of lines between anchor points, akin to a downhill skier's path between set flags. Thus, MARS builds a flexible model in which predictors can have nonlinear relations with an outcome and can adaptively handle samples in which predictors outnumber observations through the weighting term, which can give model sparsity by setting some coefficients to 0. This is similar to a slalom skier hitting each

flag on his way down but having the flexibility to choose his best route between flags, allowing him to minimize his final time.

f) Gradient boosting methods⁶ are an ensemble method based on an iteratively updated link function, $F_m(x)$, composed of a weighted sum of base learner functions, $h_i(x)$, and their weights, γ_i ,

$$F(x) = \sum_{i=1}^M \gamma_i h_i(x) + \text{constant}$$

with a greedy update that minimizes a certain loss function, such that data pieces contributing to error in a previous iteration are given preferential weighting in the subsequent iteration. This iteratively minimizes error and provides larger data coverage than other methods. Deducing the main parts of a puzzle from a growing set of pieces is an apt analogy for the boosting process.

Steepest descent methods (employing gradient calculation) are used to update the boosted model, as adaptive addition to the model is a difficult computational problem. This is similar to a climber descending a mountain according to quickest route (likely rappelling a cliff); it is represented mathematically as:

$$F_m(x) = F_{m-1}(x) - \gamma_m \sum_{i=1}^n \nabla(L(\gamma_i, F_{m-1}(x_i)))$$

where baselearner weight at a given iteration is:

$$\gamma_m = \arg \min_{\gamma} ||\mathbf{y} - \mathbf{X}\boldsymbol{\beta}|| \sum_{i=1}^m L\left(\gamma_i, F_{m-1}(x_i) - \gamma \frac{\partial L(\gamma_i, F_{m-1}(x_i))}{\partial f(x_i)}\right)$$

where $L(\gamma_i, F_{m-1}(x_i))$ is a loss function, such as ℓ_1 loss or Huber loss. Line search gives optimized weights:

$$\gamma_m = \arg \min_{\gamma} \sum_{i=1}^m L(\gamma_i, F_{m-1}(x_i) + \gamma h_m(x_i))$$

with gradient descent updates to the result. This creates a stable, robust final model after m iterations (chosen either by minimum change parameters or a maximum iteration number):

$$F_m(x) = F_{m-1}(x) + \gamma_m h_m(x)$$

Thus, boosting yields stable and flexible generalized linear model solutions that can handle nonlinear relationships between predictors and an outcome, as well as the $p > n$ situation.

g) Another ensemble method that extends generalized linear models is Bayesian model averaging (BMA)⁴⁶, whereby all possible models from a given set of predictors and outcome are combined into a weighted average according to the Bayes factor (B) of proposed model, M_1 , compared to the null model, M_0 , given the observed data, D :

$$B = \frac{pr(D|M_1)}{pr(D|M_0)}$$

With K possible models, the BMA posterior probability of a model, M_k , is given by:

$$pr(M_0|D) = \frac{\alpha_k B_{k0}}{\sum_{j=0}^K \alpha_j B_{j0}}$$

where α_k is the prior odds for M_k against M_0 . The advantages of this method include its ensemble formulation as an iterative weighted average (similar to boosting) and its ranking of solution models, such that the user can determine if one model is clearly best given a criterion (i.e. BIC, AIC...) or if several models seem to fit the data equally well⁴⁶. This is analogous to a painter blending a mix of palette colors to derive the right hue for a painting through an optimal combination of paints; some hues may be more complicated than others to create from a given palette.

h) Other machine learning algorithms for binary classification problems do not yield linear models with interpretable coefficients. Some nonparametric methods that have shown good performance on these types of problems in healthcare and medicine include random forest⁴⁷ (a bagged ensemble utilizing bootstrap methods to select observations and predictors upon which to build each individual model), conditional inference trees⁴⁸ (classification trees in which splits are selected through statistical testing), k-nearest neighbor regression⁴⁹ (a local regression method iteratively fitting the model at each point's topological neighborhood defined by the k-nearest points in the data space), single-layer feedforward neural network regression⁵⁰ (optimizing a series of parallel functional mappings according to inputs and outputs), and regularized boosted regression⁵¹ (a boosting framework incorporating the penalties of elastic net models with either tree or linear baselearners).

From these, three sets of comparisons were derived for the simulations and BCWD dataset. All simulations and models were conducted in R using the packages indicated in parentheses. First, linear main effects models were compared on simulations and BCWD data, including logistic regression (stats), elastic net (glmnet), homotopy LASSO (lasso2), DGLARS (dglars), BMA (BMA), MARS (earth), and boosted regression with linear baselearners (mboost); all hyperparameters were optimized to ensure best performance. Main effects plus two-way interaction terms models were compared on simulation and BCWD data for logistic regression, homotopy LASSO, MARS, and boosted regression with optimal tuning. Finally, a set of nonparametric models was tested against homotopy LASSO and DGLARS models on the BCWD dataset, including random forest (ranger), conditional inference trees (partkit ctee), k-nearest neighbor regression (caret with knnreg function), feedforward neural network (nnet), and regularized boosting with tree baselearners (xgboost).

II) Simulations and Datasets

Datasets were simulated with 13 predictors and a binary outcome, including four true predictors and nine noise variables. Relationships were characterized as linear (four main effects), nonlinear (two two-way interaction terms), and mixed (two main effects and one two-way interaction effect). Within each relationship condition, added noise level to the

outcome and overlap/outlier addition increased, such that three conditions existed within each relationship condition: low added Gaussian noise (0,0.25) and no outlier/overlap; medium added Gaussian noise (0,0.5) and low outlier/overlap (~5-10%); and high added Gaussian noise (0,0.75) and more outlier/overlap (~15-20%). Each of the nine conditions were replicated ten times, and model performance was averaged across replicates. All models were fit with a 70/30 train/test split to minimize bias and overfitting. Sample size within each replicate was set to 500, 1000, 2500, 5000, and 10000, respectively, such that convergence rates for each algorithm within each condition could be examined. This allowed for comparison of robustness to errors and misclassification/outliers within datasets across algorithms, as well as a way to test algorithms' ability to deal with small sample sizes that are common in medical research.

The BCWD dataset was downloaded from the UCI repository and included 569 individuals with 30 tumor attributes; the outcome of interest was malignant/non-malignant classification. All sets of algorithms were compared using the same 70/30 train/test split of this dataset, and a comparison of selected variables/coefficients was conducted across several algorithms yielding interpretable linear models.

The COGA dataset^{52, 53} focuses on genetic, demographic, and patient history information to predict alcoholism in a sample of Caucasian and African-American participants; previous results suggest family history and trauma history are important predictors of alcoholism onset and time-to-recovery⁵³. For this study, 23 demographic and patient history factors were considered for 163 Caucasian participants, including gender, ethnicity, parental alcoholism history (strict and relaxed criteria), total number of trauma types in childhood and adulthood, cigarette/marijuana/cocaine/opiate history (ever use and dependency), patient weight, educational attainment, income level, age at study, maximum number of drinks in one sitting, height, and professional attainment. All sets of algorithms were compared using a 70/30 train/test split, and selected variables/coefficients were compared across interpretable linear models.

Results

a) Main Effects Dataset Simulation

Several interesting trends emerge from the simulation trials (Figure 3; appendix Figure 7 shows homotopy LASSO and DGLARS relative to other algorithms without a color code). Most algorithms perform near optimally on datasets with only main effects and little noise/few outliers, particularly at sample sizes of at least 1000. This performance degrades as the noise level and overlap/outlier fraction increases, with boosted regression and homotopy LASSO retaining their strong performance for models including main effects and interaction terms.

b) Interaction Effects Dataset Simulation

Most algorithms performed similarly on the pure interaction term dataset with little noise/few outliers, though all performed worse on this task than on the main effects dataset. DGLARS and homotopy LASSO with main effects and interaction terms came out as the best performers on this task, with homotopy LASSO retaining its performance as noise and overlap/outliers were added. MARS and boosted regression with main effects and interaction terms caught up as more noise and overlap/outliers were added. DGLARS also showed noted improvement over other main effects models for the dataset on the condition with the most noise and overlap/outliers, suggesting that it is fairly robust.

c) Mixed Effects Dataset Simulation

For mixed effect datasets, particularly those with more noise and overlap/outlier presence, main effects DGLARS and both homotopy LASSO models show substantial gains over other algorithms, including both logistic regression and elastic net. Interestingly, MARS and boosted regression models including main effects and interaction terms do not show as good of performance as the main effects DGLARS model or either of the homotopy LASSO models at any of the noise conditions. Because real datasets tend to have main effects and interaction effects, as well as group overlap with imprecise measurements yielding noise in the outcome (particularly medical datasets

where diseases can have multiple subtypes/causes and outliers), the mixed effect dataset results suggest that algorithms incorporating topological or geometric information may yield better predictive performance and more accurate insights than other algorithms on real-world medical data.

The main effects DGLARS model results are particularly interesting, as this model performs relatively well even on data with only interaction effects or including both main effects and interaction effects. This performance was not reached by other main effects models in the mixed effect trials, and homotopy LASSO was the only other algorithm that came close to matching its performance for those mixed effect models. To investigate whether the noise or the group overlap/outlier addition (or the combination of both) underlay these relative improvements of DGLARS, additional simulations of the mixed high noise/overlap condition were run to isolate the cause of this phenomenon and better understand situations in which DGLARS is the clear favorite for deriving meaningful results and good prediction. Results suggest that it is the combination of group overlap and increasing noise that underlies this relative performance boost (appendix Figures 9 and 10), suggesting its efficacy on very messy problems where noise contaminates the data, linear or interaction-based relationships exist, and group overlap/outliers also exist.

For real world medical data, it is likely that DGLARS and homotopy LASSO will provide substantial gains over other linear model algorithms, particularly at small sample sizes common in GWAS/microarray studies or rare disease trials.

d) BCWD Dataset

Among main effects models (Figure 4), linear regression machine learning models yielded better prediction than logistic regression, with homotopy LASSO, MARS, and boosted regression models showing lower error and more balanced false positive/false negative rates. All linear regression machine learning fit statistics showed excellent fit, with DGLARS yielding a BIC of 113 and MARS yielding an R^2 value >0.9 . All models

gave an AUC over 0.90, with boosting and penalized regression models yielded AUCs >0.95.

Penalized main effects models generally selected similar predictors and reduced the set of predictors to less than half that of the original dataset in each individual model (Figures 5 and 6). Most extreme and mean concavity/concave points emerge as strong risk factors across models, suggesting that tumor geometry/shape is an important distinguisher between malignant and benign tumors. Computational issues existed for most algorithms, as odds ratios of >10000 were reported by many models. Homotopy LASSO was the only linear model algorithm that did not suffer from these singularities, suggesting that it is able to handle this type of data geometry effectively and return reasonable estimates of effect size.

The main effects plus interaction effects models demonstrate those computational issues, as issues occurred for logistic regression but not for boosted regression, MARS, or homotopy LASSO models. In fact, the logistic regression model outputted several errors about model fit and model estimate reliability. Given the relatively small sample size and large number of potential interaction terms, generalized linear models need computational help to identify and estimate effects.

The main effects homotopy LASSO algorithm out-performed nonparametric methods, and DGLARS performed comparably to these methods (Figure 4). In addition, both DGLARS and homotopy LASSO were able to minimize false negative rates relative to the nonparametric methods; in cancer detection tasks, false negatives miss sick patients who can benefit from early detection and treatment. Based on results, the best guidance would come from combining the homotopy LASSO model with a 2% false negative rate and the neural network with a 1% false positive rate, such that a negative on the homotopy LASSO would be classified as a negative and a positive on the neural network would be classified as a positive; thus, 97% of patients would be correctly classified. This strategy is similar to the superlearner framework, in which a diverse set of models is leveraged to find the best combination; in fact, combining these two models yields an AUC of 0.99, almost perfect prediction and better than the neural network model (0.91) or the main effects homotopy LASSO model (0.97).

e) COGA Dataset

Among main effects models (Figure 7), linear regression machine learning models perform better, posting accuracies >5% higher than logistic regression. All models show excellent fit statistics. DGLARS yielded a BIC of 25 (compared to a BIC of 46 for the logistic regression model), and MARS R^2 value was >0.9. Homotopy LASSO model showed the highest accuracy (AUC>0.95) and posted a much lower false negative rate than the other models; DGLARS complemented this model with a false positive rate of 0.

Main effects models generally selected similar predictors and reduced the set of predictors to less than half of the original set (Figure 8). Daily cigarette use and lifetime history of cocaine use were the strongest predictors of alcoholism. Total traumas experienced, maximum drinks, and lifetime history of marijuana use came out as important risk factors in several models, as well. Computational issues occurred for DGLARS (odds ratios >100) and the BMA model (no model could be fit). Homotopy LASSO did not suffer from any of these issues and generally showed overlap with many other models, particularly MARS.

The main effects plus interaction effects models demonstrated that sample size issues occurred for the logistic regression model but not boosted regression, MARS, or homotopy LASSO models. Homotopy LASSO emerged as the best model, with very low false positive rates and competitively low false negative rates. Given the extremely small sample size and large number of predictors, generalized linear models benefit from penalty methods and machine-learning-based selection strategies.

Compared to nonparametric methods, the main effects homotopy LASSO (and the full homotopy LASSO) model performed competitively, with all nonparametric methods other than neural networks posting AUCs >0.95 and 0 false negatives in models (Figure 7). Combining the main effects homotopy LASSO model with any of these other models provides a false negative rate of 0 and a false positive rate of 0.02, yielding almost perfect prediction. The neural network model struggled to separate the classes and seems to suffer from some of the same issues that the main effects plus interaction terms logistic regression model faced.

Discussion

These simulation, BCWD, and COGA results suggest that machine learning algorithms directly incorporating data topological or geometric information perform well relative to similar algorithms that do not consider this information, particularly on datasets with a mix of main effects and interaction terms or overlap between/outliers within classification groups. This mirrors recent successes of topologically- and geometrically-based algorithms for matching, dimensionality reduction, ranking, and partitioning problems^{25,27,30,34,39,44,54}.

In addition, results suggest that leveraging these types of algorithms may be useful in estimating effect sizes in datasets where singularities arise in model fitting. Homotopy LASSO was able to bound effect sizes and to avoid odds ratios that approach infinity; MARS, DGLARS, logistic regression, boosted regression, BMA, and elastic net models all output odds ratios approaching infinity for at least one selected predictor in BCWD or COGA datasets. This suggests the efficacy of topologically-based search methods in particular for estimating effect sizes within a generalized linear model framework; because topology focuses on global data characteristics, such as path equivalence/deformation, it can avoid the singular, saddle-point, or curved local geometry that can trap gradient search and hill-climbing solvers⁵³.

Several limitations exist in this study, including its non-exhaustive comparison of algorithms and focus on logistic regression, rather than other types of generalized linear modeling problems that arise in medical models. Future studies should test the efficacy of DGLARS and homotopy LASSO on Tweedie outcomes (such as behavioral count data with or without over-dispersion), normally-distributed outcomes (such as height or intelligence data), multinomial outcomes (such as cancer subtype classification), and survival outcomes (such as Cox regression for cancer survival data). Future simulations might focus on a broader range of algorithms than those yielding linear models with estimated coefficients. Another limitation is the limited testing on real-world medical datasets. Future studies may want to test these algorithms on microarray data or electronic health record data to confirm this study's findings and the potential use of these algorithms to analyze general healthcare data.

Despite these limitations, this study adds evidence that topologically- and geometrically-based machine learning algorithms have the potential to improve linear modeling prediction and estimation, particularly on medical datasets where noise is inherent and group overlap/outliers are common. More work should be done to investigate ways to incorporate this information to algorithms like LASSO, MARS, or boosted regression with linear baselearners.

Figure 1: Homotopic paths on a sphere (left) and homotopic vs. non-homotopic paths on a torus (right)

Figure 2: DGLARS geometry, angle between parameter likelihood vector and regression error tangent space

Figure 3: Color version of simulation results (simplified version in appendix that is color independent). Linear (left column), nonlinear (middle column), and mixed (right row) trials with low (top row), medium (middle row), and high (bottom row) noise/overlap; color code: Blue=interaction homotopy LASSO logistic regression, Dark Blue=main effects homotopy LASSO logistic regression, Red=interaction linear logistic boosted regression, Dark Red=main effects linear logistic boosted regression, Purple=main effects differential geometry LARS regression, Gold=main effects logistic regression, Yellow=interaction effects logistic regression, Green=interaction effects MARS logistic regression, Dark Green=main effects MARS logistic regression, Brown=main effects elastic net logistic regression, Pink=main effects Bayesian model averaged logistic regression; Black=maximum possible of aggregated linear models (superlearner estimate)

Figure 4: Performance of algorithm sets on BCWD

Figure 5: Blowup of odds ratios for models other than homotopy LASSO

Figure 6: Detailed odds ratios for overlapping terms in BCWD linear models

Figure 7: Performance of Algorithm Sets on COGA Data

Figure 8: Detailed odds ratios for overlapping terms in COGA Dataset linear models

References

Grant BF, Dawson DA. Age at onset of alcohol use and its association with DSM-IV alcohol abuse and dependence: results from the National Longitudinal Alcohol Epidemiologic Survey. *Journal of substance abuse*. 1997 Dec 31;9:103-10.

Hinkin CH, Hardy DJ, Mason KI, Castellon SA, Durvasula RS, Lam MN, Stefaniak M. Medication adherence in HIV-infected adults: effect of patient age, cognitive status, and substance abuse. *AIDS (London, England)*. 2004 Jan 1;18(Suppl 1):S19.

Andrews PJ, Sleeman DH, Statham PF, McQuatt A, Corruble V, Jones PA, Howells TP, Macmillan CS. Predicting recovery in patients suffering from traumatic brain injury by using admission variables and physiological data: a comparison between decision tree analysis and logistic regression. *Journal of neurosurgery*. 2002 Aug;97(2):326-36.

Heidema AG, Boer JM, Nagelkerke N, Mariman EC, Feskens EJ. The challenge for genetic epidemiologists: how to analyze large numbers of SNPs in relation to complex diseases. *BMC genetics*. 2006 Apr 21;7(1):23.

Fan J, Lv J. A selective overview of variable selection in high dimensional feature space. *Statistica Sinica*. 2010 Jan;20(1):101.

Friedman J, Hastie T, Tibshirani R. Additive logistic regression: a statistical view of boosting (with discussion and a rejoinder by the authors). *The annals of statistics*. 2000;28(2):337-407.

Zou H, Hastie T. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2005 Apr 1;67(2):301-20.

Friedman JH. Multivariate adaptive regression splines. *The annals of statistics*. 1991 Mar 1:1-67.

Wu TT, Chen YF, Hastie T, Sobel E, Lange K. Genome-wide association analysis by lasso penalized logistic regression. *Bioinformatics*. 2009 Mar 15;25(6):714-21.

Ayers KL, Cordell HJ. SNP Selection in genome-wide and candidate gene studies via penalized logistic regression. *Genetic epidemiology*. 2010 Dec 1;34(8):879-91.

Hillemacher T, Frieling H, Wilhelm J, Heberlein A, Karagülle D, Bleich S, Lenz B, Kornhuber J. Indicators for elevated risk factors for alcohol-withdrawal seizures: an analysis using a random forest algorithm. *Journal of neural transmission*. 2012 Nov 1;119(11):1449-53.

Sinisi SE, Polley EC, Petersen ML, Rhee SY, van der Laan MJ. Super learning: an application to the prediction of HIV-1 drug resistance. *Statistical applications in genetics and molecular biology*. 2007 Jan 1;6(1).

Pflueger MO, Franke I, Graf M, Hachtel H. Predicting general criminal recidivism in mentally disordered offenders using a random forest approach. *BMC psychiatry*. 2015 Mar 29;15(1):62.

Hua KL, Hsu CH, Hidayati SC, Cheng WH, Chen YJ. Computer-aided classification of lung nodules on computed tomography images via deep learning technique. *OncoTargets and therapy*. 2015;8.

Steyerberg EW, Mushkudiani N, Perel P, Butcher I, Lu J, McHugh GS, Murray GD, Marmarou A, Roberts I, Habbema JD, Maas AI. Predicting outcome after traumatic brain injury: development and international validation of prognostic scores based on admission characteristics. *PLoS Med*. 2008 Aug 5;5(8):e165.

McCullagh P. Generalized linear models. *European Journal of Operational Research*. 1984 Jun 1;16(3):285-92.

Augugliaro L, Mineo AM, Wit EC. A differential geometric approach to identify important variables in GLMs when $p \gg n$. *Statistics and the Life Sciences: High-dimensional inference and complex data*, University of Groningen, Groningen, September. 2009:9-11.

Weber M, Saucan E, Jost J. Characterizing Complex Networks with Forman-Ricci curvature and associated geometric flows. *arXiv preprint arXiv:1607.08654*. 2016 Jul 28.

Weber M, Jost J, Saucan E. Forman-Ricci flow for change detection in large dynamic data sets. *Axioms*. 2016 Nov 10;5(4):26.

Lee H, Ma Z, Wang Y, Chung MK. Topological Distances between Networks and Its Application to Brain Imaging. *arXiv preprint arXiv:1701.04171*. 2017 Jan 16.

Sardiu ME, Gilmore JM, Groppe B, Florens L, Washburn MP. Identification of topological network modules in perturbed protein interaction networks. *Scientific Reports*. 2017;7.

Gidea M. Topology data analysis of critical transitions in financial networks.

Xu Q, Jiang T, Yao Y, Huang Q, Yan B, Lin W. Random partial paired comparison for subjective video quality assessment via HodgeRank. In *Proceedings of the 19th ACM international conference on Multimedia* 2011 Nov 28 (pp. 393-402). ACM.

Huang Y, Kou G, Peng Y. Nonlinear manifold learning for early warnings in financial markets. *European Journal of Operational Research*. 2017 Apr 16;258(2):692-702.

Wang Y, Shi J, Yin X, Gu X, Chan TF, Yau ST, Toga AW, Thompson PM. Brain surface conformal parameterization with the Ricci flow. *IEEE transactions on medical imaging*. 2012 Feb;31(2):251-64.

Xu W, Hancock ER, Wilson RC. Rectifying non-euclidean similarity data using ricci flow embedding. In *Pattern Recognition (ICPR), 2010 20th International Conference on* 2010 Aug 23 (pp. 3324-3327). IEEE.

Li Y. Applying Ricci Flow to Manifold Learning. *arXiv preprint arXiv:1703.10675*. 2017 Mar 3.

Huang Z, Wan C, Probst T, Van Gool L. Deep Learning on Lie Groups for Skeleton-based Action Recognition. *arXiv preprint arXiv:1612.05877*. 2016 Dec 18.

Clarke J, Seo P, Clarke B. Statistical expression deconvolution from mixed tissue samples. *Bioinformatics*. 2010 Apr 15;26(8):1043-9.

Cazals F, Chazal F, Lewiner T. Molecular shape analysis based upon the Morse-Smale complex and the Connolly function. In *Proceedings of the nineteenth annual symposium on Computational geometry* 2003 Jun 8 (pp. 351-360). ACM.

Bakircioğlu M, Grenander U, Miller MI. Curve matching on brain surfaces using frenet distances. *Human Brain Mapping*. 1998 Jan 1;6(5-6):329-33.

Gerber S, Rübel O, Bremer PT, Pascucci V, Whitaker RT. Morse-smale regression. *Journal of Computational and Graphical Statistics*. 2013 Jan 1;22(1):193-214.

Suzumura S, Ogawa K, Sugiyama M, Takeuchi I. Outlier Path: A Homotopy Algorithm for Robust SVM. In *ICML 2014* (pp. 1098-1106).

Chen YC, Genovese CR, Wasserman L. Statistical inference using the Morse-Smale complex. *Electronic Journal of Statistics*. 2017;11(1):1390-433.

Pearson PT. Visualizing clusters in artificial neural networks using morse theory. *Advances in Artificial Neural Systems*. 2013 Jan 1;2013:6.

Tsagkraloulis D, Montana G. Random Forest regression for manifold-valued responses. *arXiv preprint arXiv:1701.08381*. 2017 Jan 29.

Amari SI. Differential geometry of curved exponential families-curvatures and information loss. *The Annals of Statistics*. 1982 Jun 1:357-85.

Lum PY, Singh G, Lehman A, Ishkanov T, Vejdemo-Johansson M, Alagappan M, Carlsson J, Carlsson G. Extracting insights from the shape of complex data using topology. *Scientific reports*. 2013 Feb 7;3:1236.

Dey TK, Memoli F, Wang Y. Topological Analysis of Nerves, Reeb Spaces, Mappers, and Multiscale Mappers. *arXiv preprint arXiv:1703.07387*. 2017 Mar 21.

Nielson JL, Cooper SR, Yue JK, Sorani MD, Inoue T, Yuh EL, Mukherjee P, Petrossian TC, Paquette J, Lum PY, Carlsson GE. Uncovering precision phenotype-biomarker associations in traumatic brain injury using topological data analysis. *PloS one*. 2017 Mar 3;12(3):e0169490.

Tierny J, Carr H. Jacobi Fiber Surfaces for Bivariate Reeb Space Computation. *IEEE Transactions on Visualization and Computer Graphics*. 2017 Jan;23(1):960-9.

Deng CH, Zhao WL. Fast k-means based on KNN Graph. *arXiv preprint arXiv:1705.01813*. 2017 May 4.

Moon C, Giansiracusa N, Lazar N. Persistence Terrace for Topological Inference of Point Cloud Data. *arXiv preprint arXiv:1705.02037*. 2017 May 4.

Bendich P, Gasparovic E, Tralie CJ, Harer J. Scaffoldings and Spines: Organizing High-Dimensional Data Using Cover Trees, Local Principal Component Analysis, and Persistent Homology. *arXiv preprint arXiv:1602.06245*. 2016 Feb 19.

Osborne MR, Presnell B, Turlach BA. A new approach to variable selection in least squares problems. *IMA journal of numerical analysis*. 2000 Jul 1;20(3):389-403.

Raftery AE, Madigan D, Hoeting JA. Bayesian model averaging for linear regression models. *Journal of the American Statistical Association*. 1997 Mar 1;92(437):179-91.

Breiman L. Random forests. *Machine learning*. 2001 Oct 1;45(1):5-32.

Hothorn T, Hornik K, Zeileis A. Unbiased recursive partitioning: A conditional inference framework. *Journal of Computational and Graphical statistics*. 2006 Sep 1;15(3):651-74.

Altman NS. An introduction to kernel and nearest-neighbor nonparametric regression. *The American Statistician*. 1992 Aug 1;46(3):175-85.

Bebis G, Georgiopoulos M. Feed-forward neural networks. *IEEE Potentials*. 1994 Oct;13(4):27-31.

Chen T, Guestrin C. Xgboost: A scalable tree boosting system. In *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* 2016 Aug 13 (pp. 785-794). ACM.

Van Der Maaten L, Postma E, Van den Herik J. Dimensionality reduction: a comparative. *J Mach Learn Res*. 2009 Oct 26;10:66-71.

Reich T. A genomic survey of alcohol dependence and related phenotypes: results from the Collaborative Study on the Genetics of Alcoholism (COGA). *Alcoholism: Clinical and Experimental Research*. 1996 Nov 1;20(s8).

Farrelly CM. The Role of Trauma in Alcoholism Risk and Age of Alcoholism Onset. *PsyArXiv preprint PsyArXiv:10.17605/OSF.IO/U3HG9*. 2017 September 26.

Du SS, Jin C, Lee JD, Jordan MI, Póczos B, Singh A. Gradient Descent Can Take Exponential Time to Escape Saddle Points. *arXiv preprint arXiv:1705.10412*. 2017 May 29.

Appendix

Figure 9: DGLARS (triangles) and main effects plus interaction effects homotopy LASSO (stars) comparison for simulations

Figure 10: DGLARS (triangles) and main effects plus interaction effects homotopy LASSO (stars) in the mixed simulations with high noise only (top), high overlap only (middle), high noise and high overlap (bottom)