# BRIEF COMMUNICATION OPEN (Check for updates) Prediction of the functional impact of missense variants in BRCA1 and BRCA2 with BRCA-ML

Steven N. Hart <sup>1</sup><sup>∞</sup>, Eric C. Polley<sup>1</sup>, Hermella Shimelis<sup>2</sup>, Siddhartha Yadav<sup>3</sup> and Fergus J. Couch<sup>1,2</sup>

In silico predictions of missense variants is an important consideration when interpreting variants of uncertain significance (VUS) in the *BRCA1* and *BRCA2* genes. We trained and evaluated hundreds of machine learning algorithms based on results from validated functional assays to better predict missense variants in these genes as damaging or neutral. This new optimal "BRCA-ML" model yielded a substantially more accurate method than current algorithms for interpreting the functional impact of variants in these genes, making BRCA-ML a valuable addition to data sources for VUS classification.

npj Breast Cancer (2020)6:13; https://doi.org/10.1038/s41523-020-0159-x

# INTRODUCTION

Failure to accurately predict the effects of missense variants in *BRCA1* and *BRCA2* confound interpretation of gene sequencing studies and clinical care. Until recently, few missense variants had been functionally evaluated using validated assays, so interpretations of pathogenicity have relied on in silico predictions of functional effect in combination with family-based data. Many in silico prediction models are derived from supervised learning methods using variants in many different genes across the genome. The objective of supervised learning is to identify and weight a set of input features to correctly predict whether a variant is damaging, neutral, or somewhere in between.

Machine learning (ML) is a suite of computational algorithms that are able to parse data, learn higher dimensional representations of that data, and ultimately make a prediction using that data. A subclass of ML, known as supervised learning, involves utilizing a training dataset with known outcomes and learning a function to be able to evaluate new unknown outcome observations and make predictions of the outcome. Examples of ML include logistic regression algorithms and more complex ones like random forests, gradient boosting machines, and neural networks. Choosing the algorithms most suited to a particular task is an active area of research, since no single algorithm outperforms all others on every task<sup>1</sup>. An efficient exploration of many different ML algorithms can be achieved through an automated machine learning (AutoML) approach. To what extent an optimal ensemble of methods could be determined by combining AutoML with these algorithms is the focus of this paper.

A key limitation to the application of existing in silico models to assessment of variants in a specific gene is the reliance on known damaging variants in other genes. Such variants are likely to cause a number of different effects (e.g., alternative splicing, disruption of protein–protein interactions, altered protein folding, etc.) that may or may not be relevant for a given gene of interest. Genespecific models will likely outperform any general model, but only a few genes have been characterized to a degree that would be informative for single gene models. Two such exceptions are *BRCA1* and *BRCA2*. The landscape of functionally characterized variants in *BRCA1* has dramatically increased because of three major analyses. Starita et al.<sup>2</sup> measured the impact of 1056 N-

terminal variants on the homologous recombination DNA repair activity of *BRCA1*. Findlay et al.<sup>3</sup> exploited the essentiality of BRCA1 for cell survival by used a saturating genome editing approach in HAP1 cells to evaluate nearly 4000 SNVs (n = 1837 distinct missense). Finally, Fernandes et al.<sup>4</sup> reported on analysis of 354 distinct missense variants (n = 79 in IARC classes 0 or 1 [benign] or 4,5 [pathogenic]) in the BRCT domain of BRCA1 using a validated transcriptional assay. Combined with results from a homology directed repair assay of 207 missense variants in the DNA-binding domain of BRCA2<sup>5</sup>, there are now sufficient numbers of variants to apply supervised learning methods to better predict damaging mutations in *BRCA1* and *BRCA2*.

# RESULTS

#### Optimal models for BRCA1 and BRCA2

An iterative process was used to build hundreds of predictive models with different algorithms (Linear models, Gradient Boosting Machines, XGBoost, Neural Networks, Random Forests, and Extremely Randomized Forests) and their associated hyperparameters. For BRCA1, 663 model/parameter combinations were tested with AutoML. The best performing model was a Gradient Boosting Machine with 48 trees of depth = 8 and between 16–33 leaves. The mean MCC was  $0.66 \pm 0.049$  s.d., corresponding to 89.5% sensitivity and 91.5% specificity. MutPred, AlignGVGD, and VEST3 contributed to 28% of the model overall, followed by CADD and REVEL at 7% each, and all others under 6.2%. Similarly, for BRCA2, 76 model/parameter combinations were tested. The best performing model being an XGBoosted Machine with 50 trees with a mean MCC of  $0.73 \pm 0.057$  s.d, sensitivity of 97.7%, and specificity of 85.1%. Variable importance for this model was driven by CADD, MutPred, and MCAP (all >10%), followed by EigenRaw and LRT, with the remaining predictors contributing below 6.4%. Throughout the remainder of the paper, we will simply refer to these gene-specific models as BRCA-ML.

For the individual predictors (excluding BRCA-ML), the best model as determined by Matthews correlation coefficient (MCC) for *BRCA1* was MutPredScore (MCC = 0.399, sensitivity = 94.5%, and specificity = 77.8%) and BayesDel (MCC = 0.673, sensitivity = 85.3%, and specificity = 83.0%) for *BRCA2*. More globally, we show





<sup>&</sup>lt;sup>1</sup>Department of Health Sciences Research, Mayo Clinic, Rochester, MN, USA. <sup>2</sup>Department of Laboratory Medicine and Pathology, Mayo Clinic, Rochester, MN, USA. <sup>3</sup>Department of Oncology, Mayo Clinic, Rochester, MN, USA. <sup>Elemail:</sup> hart.steven@mayo.edu



Fig. 1 Receiver operating curves (top) and precision-recall curves (bottom) for BRCA1 (right) and BRCA2 (left) for the hold out test set. The ideal location in the ROC curve is the top left corner, whereas the optimal position in the PR curve is the top right.

the receiver operating and precision-recall curves in Fig. 1, which demonstrates better performance of BRCA-ML compared with other prediction models. In particular, the high number of false negative calls in BRCA1 many of the models yielded low area under the precision-recall curves. BRCA-ML scores for every possible missense mutation caused by a single-nucleotide variation are also given in Supplementary Data Set 1.

Figure 2 shows the gene-level scores for every possible missense variant caused by a single-nucleotide variant in *BRCA1* and *BRCA2* using BRCA-ML and BayesDel<sup>9</sup>, a commonly used and highly accurate predictor. While BayesDel is correctly assigning higher scores to known functional domains, the higher scores are not much more than predicted benign variants across the gene. However, in BRCA-ML, the signal to noise ratio is considerably higher between damaging and neutral variants. This evidence suggests that, unlike BayesDel, changing the threshold for damaging mutations will not significantly affect the number of predicted damaging mutations in BRCA-ML.

## DISCUSSION

We have shown that AutoML methods are efficient means to derive optimal ML models for predicting damaging missense mutations in *BRCA1* and *BRCA2*. The final models derived for each gene, which we collectively term BRCA-ML, show marked improvements in MCC and other metrics with respect to individual missense prediction algorithms.

Even in the age of large-scale mutational scanning techniques like those from Findlay<sup>3</sup> and Starita<sup>2,3</sup>, in silico mutation analysis will likely continue to be relevant. While the number of variants functionally tested is impressive for both studies (1056 and 3893, respectively), there are over 12,520 and 22,772 possible single-nucleotide variants in *BRCA1* and *BRCA2*. Therefore, it could be several years before the technology exists to scale to all possible variants, hence a short term need for computational predictions.

It should be noted that there remain several limitations for these models. First, there are a limited number of known damaging mutations in BRCA1 and BRCA2 from which to build a model. The lack of damaging mutations limits the model ability to capture the complete variability of input data. Second, the training data are limited to characterized mutations in regions of the proteins known to be associated with impaired DNA damage repair. For example, the only missense variants in BRCA2 that are associated with disease are in the DNA-binding domain. It is not known if variants in other domains that we or others predict to cause damaging missense mutations are able to inhibit DNA repair. However, using the available functional data from BRCA1 in regions outside the BRCT and RING domains, BRCA-ML demonstrated a sensitivity of 74% and specificity of 98%, suggesting that extrapolation beyond the known domains may still perform well. Third, it is possible that there is some overfitting of the model due to the inherent biases in the 25 input features from dbNSFP. However, by keeping the test set isolated from the training data, this influence should be minimal. More known mutations in these genes will be necessary to guantify the amount of overfitting.

The data presented in this paper show that highly accurate prediction of missense variants in *BRCA1* and *BRCA2* are not only possible but simple to access (see Supplementary Data Set 1 for all possible SNVs in both genes). This improved performance in BRCA-ML should provide higher quality evidence to genetic counselors and researchers for interpreting deleteriousness of missense variants.

#### METHODS

# AutoML

We employed the AutoML approach with the *R* (version 3.4.2) package *h2o. ai* (version: 3.16.0.2)<sup>6</sup> to identify the optimal model for predicting the functional effect of missense variants in *BRCA1* and *BRCA2*. Variants were loaded in the following order: Hart<sup>5</sup>, Startia<sup>2</sup>, Fernandes<sup>4</sup>, and Findlay<sup>3</sup>; keeping only variants not observed in the previous studies. We also included new *BRCA2* functional data for 15 neutral (V2527A, G2544S,



Prediction · Damaging · Neutral

**Fig. 2** BayesDel (top) versus BRCA-ML (bottom) score distribution by gene. Distribution of missense prediction scores for *BRCA1* (left) and *BRCA2* (right) for each amino acid substitution. The blue line is a smoothed function to show where the density of scores is located.

I2627V, M2634T, Y2658H, A2671S, I2675V, V2728A, P2767S, A2770T, A2770D, S2806L, I2822F, S3123R, and Q2829R) and seven damaging mutations (F2562C, W2619G, K2657T, D2723N, L2753P, Y3006D, and L3101R) (Supplementary Data Set 2). Since not all variants could unequivocally be assigned to a given class, we selected variants for inclusion if they satisfied the following criteria: "FUNC" (neutral) or "LOF" (damaging)<sup>3</sup>, HDR score ≤0.33 (damaging) or ≥0.77<sup>2</sup>, or International Agency for Research on Cancer classes 0,1 (neutral) and 4,5 (damaging)<sup>4</sup>. Variants were excluded if they were not observed in known functional domains in *BRCA1* (BRCT: amino acids 1–109, RING: amino acids 1642–1855) or *BRCA2* (DNA Binding: amino acids 2479–3192). This left 1902 variants (n = 259 damaging) for *BRCA1* and 202 variants (n = 74 damaging) for *BRCA2*.

For training each gene, 80% of variants were selected and trained to maximize the per class accuracy, with robustness assessed using fivefold cross-validation. Input features were missense prediction models from dbNSFP (version 3.4)<sup>7</sup>, including SiftScore, Polyphen2HdivScore, Polyphen2HvarScore, LrtScore, MutationtasterScore, FathmmScore, Provean-Score, Vest3Score, MetasvmScore, MetalrScore, MCapScore, RevelScore, MutpredScore, CaddRaw, DannScore, FathmmMklCodingScore, GenocanyonScore, IntegratedFitconsScore, Gm12878FitconsScore, H1HescFitconsScore, HuvecFitconsScore, BayesDel, AlignGVGDPrior, EigenRaw, and EigenPcRaw. AlignGVGD<sup>8</sup> and BayesDel<sup>9</sup> were also added using the BioR framework<sup>10</sup>. Optimal cutpoints for each of the individual input features (n = 25) from dbNSFP, AlignGVGD, and BayesDel were determined using the same training data as used in AutoML so as to make a fair comparison.

### Evaluation

For the test set evaluation, statistical measures of sensitivity, specificity were computed with the *caret* package<sup>11</sup>. The MCC is used throughout as an optimal metric for gauging the performance of a binary classifier, as it represents a singular value that takes into consideration the proportion of each class. The values of MCC range from -1 to 1, where -1 represents the worst possible agreement and 1 representing perfect agreement. We also present traditional measures of performance for ML models such as receiver operating curves and precision-recall curves. All chromosomal locations and changes are relative to the hg19/GRCh37 human genome build.

#### DATA AVAILABILITY

The data generated and analyzed during this study are described in the following data record: https://doi.org/10.6084/m9.figshare.12018429<sup>12</sup>. The data and code are openly available via the author's GitHub, at https://github.com/Steven-N-Hart/BRCA-ML, and a static version of the data and code as it was at the time the related publication was published is available in the zipped folder attached to the above-mentioned figshare data record<sup>12</sup>.

#### CODE AVAILABILITY

Code is available at https://github.com/Steven-N-Hart/BRCA-ML

Received: 3 October 2019; Accepted: 6 April 2020; Published online: 29 April 2020

#### REFERENCES

- Olson, R. S., Cava, W. L., Mustahsan, Z., Varik, A. & Moore, J. H. Data-driven advice for applying machine learning to bioinformatics problems. In *Pacific Symposium* on *Biocomputing. Pacific Symposium on Biocomputing* 192–203 (World scientific, 2018).
- Starita, L. M. et al. A multiplex homology-directed DNA repair assay reveals the impact of more than 1,000 BRCA1 missense substitution variants on protein function. Am. J. Hum. Genet. 103, 498–508 (2018).
- Findlay, G. M. et al. Accurate classification of BRCA1 variants with saturation genome editing. *Nature* 562, 217–222 (2018).
- Fernandes, V. C. et al. Impact of amino acid substitutions at secondary structures in the BRCT domains of the tumor suppressor BRCA1: implications for clinical annotation. J. Biol. Chem. https://doi.org/10.1074/jbc.RA118.005274 (2019).
- Hart, S. N. et al. Comprehensive annotation of BRCA1 and BRCA2 missense variants by functionally validated sequence-based computational prediction models. *Genet. Med.* https://doi.org/10.1038/s41436-018-0018-4 (2018).
- CRAN—Package h2o. https://cran.r-project.org/web/packages/h2o/index.html. (Accessed 29 Jan 2018).
- Liu, X., Wu, C., Li, C. & Boerwinkle, E. dbNSFP v3.0: a one-stop database of functional predictions and annotations for human nonsynonymous and splicesite SNVs. *Hum. Mutat.* 37, 235–241 (2016).

- Tavtigian, S. V., Byrnes, G. B., Goldgar, D. E. & Thomas, A. Classification of rare missense substitutions, using risk surfaces, with genetic- and molecularepidemiology applications. *Hum. Mutat.* 29, 1342–1354 (2008).
- 9. Feng, B.-J. PERCH: a unified framework for disease gene prioritization. *Hum. Mutat.* **38**, 243–251 (2017).
- Kocher, J.-P. A. et al. The Biological Reference Repository (BioR): a rapid and flexible system for genomics annotation. *Bioinformatics* 30, 1920–1922 (2014).
- CRAN—Package caret. https://cran.r-project.org/web/packages/caret/index.html. (Accessed 11 Apr 2019).
- Hart, S. N., Polley, E. C., Shimelis, H., Yadav, S. & Couch, F. J. Prediction of the functional impact of missense variants in BRCA1 and BRCA2 with BRCA-ML. FigShare. https://doi.org/10.6084/m9.figshare.12018429 (2020).

# ACKNOWLEDGEMENTS

This work was supported by the Breast Cancer Research Foundation; National Institutes of Health grants [CA192393, CA176785, CA116167]; and a National Cancer Institute Specialized Program of Research Excellence (SPORE) in Breast Cancer to Mayo Clinic [P50 CA116201].

## AUTHOR CONTRIBUTIONS

H.S. generated the functional data. S.N.H. designed the analysis, analyzed the data, and drafted the paper. F.J.C., S.Y., and E.C.P. contributed a number of key suggestions for the present version. All authors edited and approved the paper.

# **COMPETING INTERESTS**

The authors declare no competing interests.

#### ADDITIONAL INFORMATION

Supplementary information is available for this paper at https://doi.org/10.1038/ s41523-020-0159-x.

Correspondence and requests for materials should be addressed to S.N.H.

Reprints and permission information is available at http://www.nature.com/ reprints

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit http://creativecommons. org/licenses/by/4.0/.

© The Author(s) 2020