

eXtasy: variant prioritization by genomic data fusion

Alejandro Sifrim^{1,2,4}, Dusan Popovic^{1,2,4}, Leon-Charles Tranchevent^{1,2}, Amin Ardeshirdavani^{1,2}, Ryo Sakai^{1,2}, Peter Konings^{1,2}, Joris R. Vermeesch³, Jan Aerts^{1,2}, Bart De Moor^{1,2}, Yves Moreau^{1,2}

¹ Department of Electrical Engineering, STADIUS Center for Dynamical Systems, Signal Processing and Data Analytics, KU Leuven, Leuven, Belgium

² iMinds Future Health Department, Leuven, Belgium

³ Laboratory of Molecular Cytogenetics and Genome Research, KU Leuven, Leuven, Belgium

⁴ These authors contributed equally to this work

Corresponding author: Yves Moreau, yves.moreau@esat.kuleuven.be

ABSTRACT

Massive parallel sequencing greatly facilitates the discovery of novel disease genes causing Mendelian and oligogenic disorders. However, many mutations are present in any individual genome, and identifying which ones are disease causing remains a largely open problem. We introduce a new approach to prioritize nonsynonymous single nucleotide variants (nSNVs) that substantially improves prediction of disease-causing variants in exome sequencing data by integrating variant impact prediction, haploinsufficiency prediction and phenotype-specific gene prioritization.

MAIN TEXT

Rare exomic variation identified by exome sequencing is particularly useful to discover the cause of many previously unresolved rare monogenic disorders. By filtering down the exome against nonsynonymous single nucleotide variants (nSNVs) and loss-of-function mutations that are not present in healthy populations or unaffected samples, we can discard a large proportion of the exomic variation as probably neutral. However, despite such aggressive filtering, several thousand candidate causal mutations remain and we need predictive methods to prioritize variants for further validation. Several computational methods have been proposed that take into account biochemical, evolutionary and structural properties of mutations to assess their potential deleteriousness¹⁻⁵. However, most of these methods suffer from high false positive rates when predicting the impact of rare nSNVs⁴. A plausible explanation for this poor performance is that many of the scrutinized variants are mildly deleterious and subject to weak

purifying selection^{6,7}, but not specific to the disease of interest. To assess this hypothesis and to further enhance variant prioritization, we propose a genomic data fusion methodology⁸ that integrates multiple strategies to detect deleteriousness of mutations and prioritizes them in a phenotype-specific manner. A key innovation that we incorporate into our strategy is a computational method for gene prioritization⁹, which scores mutated genes based on their similarity to known disease genes by fusing heterogeneous genomic information. We also integrate haploinsufficiency prediction scores¹⁰ that predict the probability that the function of a gene is affected if present in a functionally haploid state. To integrate or fuse these data sources, we use random forest learning¹¹ and train our model on the Human Gene Mutation Database (HGMD) of human disease-causing mutations¹² compared to three control sets: common polymorphisms and two independent sets of rare variation.

After generating and annotating the different data sets, we inspected the distribution of the different deleteriousness prediction scores across the positive and control sets (**Supplementary Fig. 1**). All deleteriousness prediction scores seemed to score the positive set high and thus showed a high sensitivity. When looking at the control sets, we observed that most methods correctly classified common polymorphisms as benign - yet classified a substantial proportion of rare variation as deleterious, leading to a low precision. Control variants seemed to occur more often in genes predicted to maintain functionality in a haploid state (being haplosufficient), whereas disease-causing variants showed no clear pattern. Disease-causing variants were primarily found in top-ranked genes after gene prioritization. By contrast, control variants showed a homogeneous distribution of gene prioritization ranks, which is to be expected under the assumption that they are prioritized for randomly selected phenotypes.

By integrating these different scores, we aimed to boost our ability to discriminate between putatively mildly deleterious rare variants and actual disease causing variation. We evaluated several commonly used classification approaches and chose a Random Forest (RF) classification algorithm because it outperformed all other classification algorithms on this task (**Supplementary Table 1**). This classifier is trained by comparing our positive set of disease-causing variants to the rare variant control sets. When compared to classical deleteriousness prediction scores, we observe a considerable improvement in all performance measures. This is the case when distinguishing between disease-causing and rare control variants (**Fig. 1a**, **Supplementary Fig. 2a,b** and **Supplementary Table 2**), as well as with distinguishing disease causing variants from common polymorphisms (**Fig. 1b**, **Supplementary Figure 2c,d** and **Supplementary Table 2**). The performance against common polymorphisms is in line with published results for deleteriousness prediction tools because these tools were trained using

common polymorphisms as controls. The performance of these tools against rare non-disease-causing variants is much lower than against common polymorphisms. Precision is the most improved performance measure, which is important when dealing with a prioritization task (**Supplementary Table 1, Supplementary Note**). However, performance measures obtained in retrospective benchmarks such as ours are usually optimistic estimates, because of the bias of prior information for gene prioritization predictions^{8,13}. This problem is difficult to address in an initial study and can only be truly resolved by long-term prospective benchmarks where predictions are made on the current state of knowledge and validated in future studies. To estimate to which degree such bias was present in our benchmark, we assessed classification performance based on the date of discovery of the causative variation using data for gene prioritization anterior to the year of discovery (**Supplementary Fig. 3**). If a bias is present, it would be less prominent in recent discoveries as these gene-disease associations would be less likely to be directly or indirectly incorporated in the gene prioritization data sources. Even though we see a slight decline in performance (in terms of Area Under the Curve) for more recent discoveries, the method still performs considerably better than classical deleteriousness prediction scores (**Fig. 1c, Supplementary Table 3**). The performance of deleteriousness prediction scores is not affected by the year of discovery of the association, which is expected because these scores do not integrate time-dependent information. Finally, we looked at feature importance by shuffling each feature across disease-causing and control variants evaluating the increase in classification error (**Supplementary Fig. 4**). This analysis showed that all included features were informative and improved classification to some degree.

In this study, we show a novel approach towards integrating biochemical, evolutionary and phenotype-specific information to prioritize mutations for follow-up validation studies. This approach helps to distinguish disease-causing mutations from neutral common polymorphisms as well as rare, potentially deleterious but phenotypically unrelated variation in the coding genome. We acknowledge that performance measures are likely overestimated because of the biased nature of retrospective benchmarks, yet we can clearly appreciate a marked improvement in prediction performance over frequently used deleteriousness prediction scores in recently published mutations. We envision that in the near future current initiatives such as the Critical Assessment of Genome Interpretation (CAGI), although currently focused on single phenotype benchmarks, will play a major role in providing unbiased prospective benchmarks to optimally assess the performance of methods such as the one described in this study.

Future research and the availability of larger public disease-causing variation data sets will likely widen the scope of the method to other types of mutations (mitochondrial, noncoding,

synonymous, nonsense, splice-site mutations, indels). Also, the addition of other data sources, such as locus-specific information (e.g., copy number variant prevalence, GWAS-associated loci), to our prioritization method and its integration into genetic association test across multiple samples^{14,15} will likely increase its power to discover the cause of genetic disease.

We freely provide public access to the described methodology through a web tool and also provide an offline standalone version (<http://homes.esat.kuleuven.be/~bioiuser/eXtasy/>, source code: <http://github.com/asifrim/eXtasy>).

ACKNOWLEDGEMENTS

The authors declare no conflicting points of interest. This research is supported by: Research Council KU Leuven: GOA/10/09 Manet, PFV/10/016 SymBioSys, IOF 3M120274 Immunosuppressive drugs; iMinds: SBO 2013; Hercules Stichting: Hercules III PacBio RS, Flemish Institute for Science and Technology: IWT-SB/093289, IWT-TBM Haplotyping; EU: Cost Action BM1006: NGS Data Analysis Network, FCT Neuroclinomics.

AUTHOR CONTRIBUTIONS

A.S., D.P., and Y.M. conceptually defined the project. A.S. and D.P. wrote the initial draft of the manuscript and performed the analyses. A.S. generated the data sets and developed the software tools. D.P. developed the benchmarks and trained the models. L.C.T. and A.S. computed the Endeavour gene prioritizations. A.A. and A.S. developed the webtool. R.S. and J.A. advised on data visualization and visual analytics. P.K. advised on statistical concerns. J.R.V. advised on genetical concerns. All authors revised and proofread the paper. B.D.M. co-supervised the project. Y.M. supervised the project.

REFERENCES

1. Adzhubei, I. A. *et al. Nature methods* **7**, 248–9 (2010).
2. Ng, P. & Henikoff, S. *Nucleic acids research* **31**, 3812 (2003).
3. Schwarz, J. M., Rödelberger, C., Schuelke, M. & Seelow, D. *Nature methods* **7**, 575–6 (2010).
4. Kumar, S., Sanderford, M., Gray, V. E., Ye, J. & Liu, L. *Nature methods* **9**, 855–6 (2012).
5. Chun, S. & Fay, J. C. *Genome research* **19**, 1553–61 (2009).

6. Asthana, S. *et al.* *Proceedings of the National Academy of Sciences of the United States of America* **104**, 12410–5 (2007).
7. Tennessen, J. A. *et al.* *Science* **337**, 64–9 (2012).
8. Moreau, Y. & Tranchevent, L.-C. *Nature reviews. Genetics* **13**, 523–536 (2012).
9. Aerts, S. *et al.* *Nature biotechnology* **24**, 537–44 (2006).
10. Huang, N., Lee, I., Marcotte, E. M. & Hurles, M. E. *PLoS genetics* **6**, e1001154 (2010).
11. Breiman, L. *Machine Learning* 1–33 (2001).
12. Stenson, P. D. *et al.* *Genome medicine* **1**, 13 (2009).
13. Myers, C. L., Barrett, D. R., Hibbs, M. A., Huttenhower, C. & Troyanskaya, O. G. *BMC genomics* **7**, 187 (2006).
14. Yandell, M. *et al.* *Genome research* **21**, 1529–42 (2011).
15. Ionita-Laza, I. *et al.* *American journal of human genetics* **89**, 701–12 (2011).

FIGURE LEGENDS

Figure 1: Receiver-Operator Curves (ROC) comparing eXtasy and classical deleteriousness prediction scores: ROC curves for **(a)** disease-causing vs. rare control variants **(b)** disease-causing vs. common polymorphisms. In both cases eXtasy outperforms other methods as can be seen by an increase in the area under the curve (AUC). **(c)** To test the effect of biases in our retrospective benchmark we compared obtained AUCs by stratifying disease-causing variants on the year of discovery. More recent disease variant associations

show a decrease in performance for eXtasy as biases decrease. The method however always outperforms classical deleteriousness prediction scores.

ONLINE METHODS

Data generation

To evaluate the effectiveness of our approach, we assembled and evaluated a data set consisting of disease-causing variants and different types of control variants. We used a positive data set consisting of 24,454 disease-causing nSNVs from the HGMD associated to 1,142 different Human Phenotype Ontology (HPO) terms¹⁶. Mapping between HGMD disease descriptors and HPO terms was performed using the Phenomizer tool¹⁷. Phenotypic terms for which fewer than three implicated genes were known were discarded (excluding the gene in which the variant was located) so as to allow meaningful subsequent gene prioritization. For the control data sets, we considered two classes of variants based on their minor allele frequency (MAF) in the 1000 Genomes Project data set: common polymorphisms (MAF > 1%) and rare variation (MAF ≤ 1%). Additionally we compiled a third control set of rare variants using in-house sequenced exomes of healthy individuals and retaining only high quality calls (coverage > 30X) not present in any publicly available variant database (NHLBI EVS, dbSNP, 1000G). After randomly assigning groups of 500 control variants from each control set to each HPO term represented in the disease-causing variants, we annotated all variant-phenotype combinations with functional information. For every variant, we extracted Polyphen2¹, SIFT², MutationTaster³ and LRT⁵ deleteriousness prediction scores from the dbNSFP database¹⁸. We calculated CAROL aggregate deleteriousness scores¹⁹ and added phyloP²⁰ and PhastCons²¹ evolutionary conservation scores across vertebrates, primates and placental mammals subsets from the UCSC genome browser database. We also included precalculated gene haploinsufficiency prediction scores¹⁰ where available. Finally, we computed gene prioritization scores with Endeavour⁹ using gene-phenotype associations obtained from HPO. (For additional information on the data generation and annotation process see **Supplementary Note**)

Classifier benchmarks

We set up a benchmark to determine the optimal classifier for genomic data fusion. To do this, we followed published guidelines²² applicable to these types of studies. Initially, we removed all records containing missing values and split the disease-causing variants into training (two thirds of the total number of genes) and testing sets (one third of the genes)(**Supplementary Fig. 5**). The data records were stratified at the gene level (the highest level of granularity) to assure that the algorithm is not overfitting the gene level information and thus does not overestimate

performance. Subsequently, we randomly subsampled the negative examples in the training set, so as to balance their number with the number of positive examples. We then trained 8 different classifiers using the same combinations of training and test subsets for all classifiers. In addition, we applied state-of-the-art deleteriousness prediction methods (Polyphen, SIFT, Mutation Taster and Carol scores) using their respective published or documented decision thresholds (0.85, 0.95, 0.5, 0.98). As there were several records per single variant in the data set (because of the different phenotypes per variant), we used the maximum score across phenotypes in testing to predict the outcome for each variant. We chose the maximum as this is robust to non-informative phenotypes (because of phenotypically variable diseases) which might be present in our benchmark data set of disease-causing variants. This did not apply for methods that produce a single score per variant because they are not phenotype specific (Polyphen, SIFT, Mutation Taster, and Carol).

The procedure was repeated 100 times on different random splits of data to obtain an estimate of the variance of the resulting performance measures. The results were computed in terms of accuracy, sensitivity, specificity, positive and negative predictive value, Matthews correlation coefficient, area under the Receiver/Operator Curve (ROC) and Precision/Recall (PR) Curves. We observed that all the classifiers built using all heterogeneous data sources perform better than state-of-the-art deleteriousness prediction methods, proving that these data sources contain additional information that facilitates better predictions.

Among all classification algorithms, the random forest (RF) outperformed all others in terms of all performance metrics, with the exception of sensitivity and negative predictive value (NPV) where linear discriminant analysis (LDA) performed marginally better (**Supplementary Table 1, Supplementary Fig. 6 and 7**). However, precision (positive predictive value or PPV) of the LDA is very low (0.35), indicating that the classifier is overly optimistic. The same observation holds for all of the state-of-the-art methods (precision between 0.20-0.23). Furthermore, the PR curve for the random forest shows that changing the default threshold (0.5) of the classifier results in a sharp increase in precision yet a small loss in sensitivity. This suggests that most of the true positives are highly ranked by the method. Also, the reported accuracy and other aggregate performance measures of the state-of-the-art tools depend greatly on where the decision threshold is set and the skewness of the class distribution. Even though the AUC for some of the ROC curves are higher on one tool compared to another, sometimes a point can be found where a certain measure is higher for the tool with the lower AUC.

For all classification algorithms, we use their respective Matlab 7.10 implementations – `classregtree` class for the decision trees, `TreeBagger` class for the random forests, function

knnclassify for k-nearest-neighbours, function classify (with argument “linear” or “quadratic”) for the LDA and QDA respectively, class NaiveBayes for Naive Bayes classification and functions svmtrain and svmclassify for the feed-forward neural networks. Most of the used functions and classes are part of the Statistical Toolbox, except for KNN and SVMs (Bioinformatics Toolbox). The details on parameter settings of particular classifiers are provided in **Supplementary Table 5**.

Control set benchmarks

We set up an additional benchmark to assess the behaviour of the final model under different classification schemes with regard to the different negative outcomes. We considered two different scenarios for training and testing. In a first scenario, the Random Forest model is trained using a subset of the rare non-disease causing variants as negatives, and tested against the rest of them (as in our standard setup). In a second scenario, the Random Forest model is trained using all of the rare non-disease causing variants available as negatives, and tested against the data set containing common polymorphisms.

In the case of the first scenario, the validation scheme is identical to the classifier benchmark: all data, including all positive and the negative examples, have been repeatedly (100 iterations) divided at random into training and testing subsets, grouped gene by gene. In the case of the second scenario, only the positive examples were randomly assigned, while the two distinct groups of negatives, rare and polymorphisms were used for respectively training and testing (**Supplementary Fig. 8**). We made sure that during a single iteration a split of positives stayed the same across the two scenarios.

Temporal stratification analysis

To analyze the sensitivity of the method with regard to a priori gene or disease association biases, we set up an additional benchmark to estimate the effect size of such biases. Under the assumption that recently published genes would be less likely to be biased by our gene prioritization step, we stratified our positive testing set of disease-causing variants by year of publication (2000-2012), while training the model only on data published prior to 2000. In this way we can measure performance before and after Endeavour’s data sources were last updated (last update occurred in 2008). This threshold applies to both variant and gene level of data granularity, so that variants that are discovered after 2000 but which are associated with genes that are part of the training (*i.e.*, for which the gene-phenotype association was discovered prior to 2000) are also removed from test sets. The negative examples (non-disease

causing variants) were randomly assigned to one given year between 2000 and 2012, in numbers matching the class distribution of the whole data (given the number of positives in a particular year), and with no overlap between training and test sets (see **Supplementary Fig. 9**). As before, the splits are performed gene-wise. After the training phase the classifier was used on cases from the subsequent years (2000-2011). The whole procedure of randomly assigning negatives was repeated 100 times to get stable estimates of performance metrics.

We observe a slight temporal decline in performance throughout all testing years. Nevertheless, eXtasy still significantly outperforms all classical deleteriousness prediction methods across all years. We attribute this decline in performance to the fact that over time some disease-causing genes are better described in the gene prioritization sources (*e.g.*, literature mining) and are therefore easier to classify. Although such effects are likely present and point to the fact that our main benchmark is an overoptimistic estimate of the real performance, this benchmark setting is a pessimistic estimate of the real performance due to various properties of the training/validation scheme. First, only a fraction of the positive training data can be used (data prior to 2000), leading to suboptimal learning of the features of disease-causing variants. Secondly, the training data itself contains well-described genes and is thus itself biased towards easier to classify examples. This leads to overly optimistic decision thresholds in the classifier, and thus degrades the performance when faced with more difficult examples of less well-described genes in the testing set. Finally, due to the gene-wise stratification of training and testing sets (to avoid overfitting specific genes), if a gene was published prior to 2000 but then later published in the light of a new phenotype, it was omitted from the test set. Well-described genes are often discovered to play important roles in new (and often related) physiological processes. This type of discoveries can greatly benefit from gene prioritization approaches but are excluded from this benchmark.

The real performance of the classifier depends greatly on the use case which is unknown to the researcher applying the method. In the case where the cause of the phenotype is a novel mutation in a previously described gene, the performance is likely to resemble, or even exceed, that of the cross-validation benchmark. When the phenotype is caused by a novel mutation in a gene not previously associated with the phenotype, or associated with a different phenotype, the real performance lies likely between the overoptimistic cross-validation benchmark and the pessimistic temporal analysis benchmark (the shaded light blue area in **Supplementary Fig. 3**).

Feature importance analysis

The Random Forest algorithm has the intrinsic ability to estimate the importance of features used for training¹¹. This is achieved by measuring the difference between the mean square error of the prediction on out-of-bag (OOB) examples when values of a given feature are shuffled compared to the error on undisturbed OOB data. This procedure is repeated for each and every tree in the ensemble with its corresponding OOB examples, providing a global measure of feature importance. Here, we analyzed how the different features in the data contributed to the overall classification. In particular, we ran 100 simulations (in line with previously described benchmarks) per feature during which ensembles are built using different random subsamples of the negative data. The result in terms of mean square error increase when shuffling the feature is displayed in **Supplementary Fig. 4**. From the plot, it appears that all features contribute to the classification to some degree. The increase of total MSE when one of them is randomized ranges from around 2% for PPI-HPRD to 12% for sequence similarity and functional annotations. Secondly, highly correlated features, such as various Endeavour scores or state-of-the-art methods (Carol with SIFT/Polyphen) usually form clusters of seemingly less important features with low yet non-zero increase in mean square error. This is expected as it has been shown that feature importance measures for random forests are strongly affected by the presence of correlation between features²³. In the absence of a particular feature, other correlated features partially “take over” the role of former, reducing the impact of the shuffling on the classification error. Hence, these variables are still individually very important - especially if data records contain missing values.

METHODS-ONLY REFERENCES

16. Robinson, P. N. *et al. American journal of human genetics* **83**, 610–5 (2008).
17. Köhler, S. *et al. American journal of human genetics* **85**, 457–64 (2009).
18. Liu, X., Jian, X. & Boerwinkle, E. *Human mutation* **32**, 894–9 (2011).
19. Lopes, M. C. *et al. Human Heredity* **73**, 47–51 (2012).
20. Pertea, M., Pertea, G. M. & Salzberg, S. L. *BMC bioinformatics* **12**, 274 (2011).
21. Siepel, A. *et al. Genome research* **15**, 1034–50 (2005).
22. Vihinen, M. *BMC genomics* **13 Suppl 4**, S2 (2012).
23. Strobl, C., Boulesteix, A.-L., Kneib, T., Augustin, T. & Zeileis, A. *BMC bioinformatics* **9**, 307 (2008).

COMPETING FINANCIAL INTERESTS

The authors declare no competing financial interests.