

Supplementary Appendix

This appendix has been provided by the authors to give readers additional information about their work.

Proactive variant effect mapping to accelerate genetic diagnosis for pediatric cardiac arrest

Section 1: Transformation of calmodulin fitness scores to log likelihood ratios of pathogenicity

The original scores in the calmodulin map (Weile et al. 2017) represent the fitness of yeast cells in a competitive growth assay, and reflect the ability of the expressed calmodulin variant to complement a temperature sensitive covering allele. The scores are scaled such that 0 represents a fitness analogous to nonsense variants, i.e. a complete loss of protein function, while 1 represents the fitness of the wild-type protein. However, the impact of a variant on protein function does not necessarily translate linearly to its propensity to cause disease. Indeed, given that the same calmodulin protein is encoded by three genes (*CALM1*, *CALM2*, and *CALM3*) and therefore by six alleles, it seems likely that calmodulin variants are dominant not due to gene dosage sensitivity, but rather due to variants with dominant negative ‘toxicity’. Consistent with this idea, despite the fact that 34 pathogenic or likely pathogenic missense variants have been collectively reported for *CALM1*, *CALM2*, and *CALM3* (Landrum et al. 2018), there have been no reports of pathogenic or likely pathogenic nonsense variants for these genes.

To relate protein function to pathogenicity, we estimated the densities of benchmark variant sets of putatively pathogenic and benign variants at given fitness scores. These were in turn used to calculate a log likelihood ratio (LLR) of pathogenicity, measuring how much more likely it is to observe a given fitness score for a variant that is pathogenic than one that is benign. Putatively pathogenic (‘Positive’) benchmark cases were curated from ClinVar using known pathogenic and likely pathogenic variants in *CALM1*, *CALM2*, and *CALM3* with at least 1-star quality, that is, the provision of assertion criteria. Putatively benign (‘Negative’) benchmark cases were curated from ClinVar using known benign or likely benign variants with from *CALM1*, *CALM2*, and *CALM3* with at least 1-star quality, and from gnomAD controls (Karczewski et al. 2020), using variants in *CALM1*, *CALM2*, and *CALM3* with an allele frequency in healthy control cases of greater than 10^{-6} .

To estimate densities of positive and negative variants at different fitness scores, we used an Epanechnikov kernel with automatic bandwidth determination via unbiased cross-validation, as implemented in the R-package “*kdensity*” v 1.0.1 (Moss and Tveten 2019). So that we might err on the side of conservatism, we mixed both positive and negative benchmark density functions with a uniform density, so that relative densities would tend towards an uninformative LLR of 0 where fitness scores fell far from benchmark variants. Because benchmark variants tended to have fitness scores between 0.3 to 1.3, LLRs approach 0 outside this range. Although for a recessive disorder where pathogenic variants are typically

loss-of-function, we might assume that LLR would be highest at lowest fitness scores, the likely dominant-negative mechanism of pathogenic calmodulin variants supports the more conservative LLR estimation scheme we applied.

The LLR function we estimated is positive (indicating a tendency towards pathogenicity) in a fitness range between 0.29 and 0.8, peaking at a fitness of 0.66, and is negative (indicating a tendency towards benignity) in the fitness range between 0.8 and 1.24, peaking at 1.15 (Figure S1).




Figure S1: Transformation of variant fitness scores to log likelihood ratios of pathogenicity. Bottom panel: Distribution of positive (red) and negative (green) benchmark cases and estimated density via Epanechnikov kernel. Top panel: Estimated log likelihood ratio (LLR) values as a function of fitness scores, in which fitness scores far from those of benchmark variants tend conservatively towards an uninformative LLR of 0.

Using the LLR function we transformed all fitness values in the original calmodulin map and calculated confidence intervals for each LLR using bootstrapping over the standard deviation of the fitness measurements. The full result can be found in Table S1.

An implementation of the method in the R programming language can be found on GitHub at <https://github.com/jweile/maveLLR>.

References

- Karczewski, Konrad J., Laurent C. Francioli, Grace Tiao, Beryl B. Cummings, Jessica Alföldi, Qingbo Wang, Ryan L. Collins, et al. 2020. "The Mutational Constraint Spectrum Quantified from Variation in 141,456 Humans." *Nature* 581 (7809): 434–43.
- Landrum, Melissa J., Jennifer M. Lee, Mark Benson, Garth R. Brown, Chen Chao, Shanmuga Chitipiralla, Baoshan Gu, et al. 2018. "ClinVar: Improving Access to Variant Interpretations and Supporting Evidence." *Nucleic Acids Research* 46 (D1): D1062–67.
- Moss, J., and M. Tveten. 2019. "Kdensity: An R Package for Kernel Density Estimation with Parametric Starts and Asymmetric Kernels." *Journal of Open Source Software*.
<https://joss.theoj.org/papers/10.21105/joss.01566.pdf>.
- Weile, Jochen, Song Sun, Atina G. Cote, Jennifer Knapp, Marta Verby, Joseph C. Mellor, Yingzhou Wu, et al. 2017. "A Framework for Exhaustively Mapping Functional Missense Variants." *Molecular Systems Biology* 13 (12): 957.