SuPreMo: a computational tool for streamlining *in silico* perturbation using sequence-based predictive models

Ketrin Gjoni[1,2] and Katherine S. Pollard[1,2,3]

[1]Gladstone Institute of Data Science and Biotechnology, San Francisco, CA 94158, USA

[2]Department of Epidemiology & Biostatistics, University of California, San Francisco, CA 94158, USA

[3]Chan Zuckerberg Biohub, San Francisco, CA 94158, USA

Summary

Computationally editing genome sequences is a common bioinformatics task, but current approaches have limitations, such as incompatibility with structural variants, challenges in identifying responsible sequence perturbations, and the need for vcf file inputs and phased data. To address these bottlenecks, we present Sequence Mutator for Predictive Models (SuPreMo), a scalable and comprehensive tool for performing *in silico* mutagenesis. We then demonstrate how pairs of reference and perturbed sequences can be used with machine learning models to prioritize pathogenic variants or discover new functional sequences.

Availability and Implementation

SuPreMo was written in Python, and can be run using only one line of code to generate both sequences and 3D genome disruption scores. The codebase, instructions for installation and use, and tutorials are on the Github page: https://github.com/ketringjoni/SuPreMo/tree/main.

Contact

katherine.pollard@gladstone.ucsf.edu

Supplementary information

Supplementary data are available at Bioinformatics online.

1 Introduction

Many machine learning (ML) models have been developed that predict cellular profiles from input DNA sequences (Supp Table 1). These sequence-to-profile models can predict biological features—including gene expression (Enformer (Avsec *et al.* 2021a), ExPecto (Zhou *et al.* 2018), Xpresso (Agarwal and Shendure 2020)), genome folding (Akita (Fudenberg, Kelley and Pollard 2020), C.origami (Tan *et al.* 2023), DeepC (Schwessinger *et al.* 2020), ORCA (Zhou 2022)), chromatin accessibility (Basenji (Kelley 2020), Basset (Kelley, Snoek and Rinn 2016)), and epigenetic marks (DeepFIGV (Hoffman *et al.* 2019), HyenaDNA (Nguyen *et al.* 2023), Sei (Chen *et al.* 2022a))—with incredible accuracy. These approaches are becoming increasingly popular for exploring biological questions at lower cost and higher throughput than experimental methods allow, and to address questions that are not possible to test experimentally. One exciting potential is to use sequence-to-profile models in tandem with *in silico* mutagenesis (ISM), in order to investigate how genomic alterations alter cellular profiles. This strategy generates testable, causal hypotheses about genotype-phenotype relationships (Chen *et al.* 2022b). ISM has been applied to the genomes of modern humans, archaic hominins (McArthur, Rinker and Capra 2021), and other species (Keough *et al.* 2023) to prioritize putative pathogenic variants for experimental studies (Benegas, Batra and Song 2023), decode the grammar of noncoding DNA sequences (Deng *et al.* 2023), discover new sequence motifs (Avsec *et al.* 2021b), design tissue-specific enhancers (Gosai *et al.* 2023), and uncover novel roles of sequence elements (Gunsalus, Keiser and Pollard 2023).

In theory, ISM is very high-throughput, making it feasible to quantify the effects of a large set of sequence perturbations, such as all variants in an individual's genome or a cohort of patients. However, the application of ISM at scale is currently limited by the process of generating sequences with and without perturbations. Existing tools incorporate variants into a reference genome in ways that are not compatible with ISM (bcftools consensus (Li 2011), GATK FastaAlternateReferenceMaker (Van der Auwera and O'Connor 2020), perEditor (Rivas-Astroza *et al.* 2011), etc.). One of the biggest limitations is that they incorporate all variants from an input variant call format (vcf) (Danecek *et al.* 2011) file into a single output fasta file, making it very difficult to isolate the effects of individual variants. Workarounds, such as generating an independent vcf file for each variant (or variant combination) and looping over these or post-processing the output fasta file to include one variant per locus, are extremely inefficient. Secondly, existing tools are made largely for single nucleotide polymorphisms (SNPs) or small insertions or deletions (indels), and cannot accommodate symbolic alleles—annotations in vcf files for structural variants (SVs). A possible workaround is to convert symbolic alleles into sequences by extracting them from a reference genome, but this becomes infeasible with large structural variants due to limitations with both variant complexity and memory allocation. One tool (perEditor(Rivas-Astroza *et al.* 2011)) is compatible with some complex variants but is not comprehensive and has stringent requirements. Finally, existing tools require the perturbations to be in a vcf format, which means that pseudo input files must be generated if one wishes to apply ISM to custom or simulated sequences (e.g., deleting all motifs for a given transcription factor or creating synthetic enhancers).

Due to these limitations, it is common practice for ISM practitioners to write their own code to generate input sequences for ISM studies. Indeed, the codebases for several ML models include code examples or frameworks for performing ISM (Enformer, Sei (Chen *et al.* 2022a), Basset), but these are restricted to simple variants (SNPs and indels) and do not generate sequence files for input into other models. SVs make good candidates for ISM since they span larger regions and are more likely to be damaging to the genes, regulatory regions or active sites they overlap or neighbor. For example, noncoding SVs have been shown to lead to cancer and developmental disorders by disrupting genomic contacts of key genes (Paik, Maule and Gallo 2021). SVs also alter more base pairs of the genome than any other type of genetic variation (1000 Genomes Project Consortium *et al.* 2015). One major challenge with SVs is that, to adhere to the fixed length input requirements of most ML models, input sequences must be padded, and consequently, model outputs require un-padding and masking. Another consideration is that—due to both biological effects and model artifacts related to making predictions for fixed width genomic windows—models can be highly sensitive to small changes in the input, such as masking, padding, and variant position in the window. Therefore, it is important to make predictions for augmented input sequences (shifted and/or reverse complement sequences) and evaluate them consistently across perturbations. Thus, incorporating perturbations into a reference genome becomes increasingly complicated and error-prone as variants get larger and more complex.

2 Tool description

To address these challenges, we developed SuPreMo, a framework for generating perturbed sequences for input into predictive models that is scalable, flexible, and comprehensive (Fig 1A). SuPreMo, which incorporates variants into the human reference genome one at a time and generates model-ready sequences (Supp Fig 1A), was extended to SuPreMo-Akita, which inputs those sequences into Akita (Fudenberg, Kelley and Pollard 2020), an ML model that predicts chromatin contact maps, and generates scores that measure variants' disruption to those maps (Supp Fig 1C). Both tools accept a variety of variant files—vcf (version 4.1 and 4.2), txt, bed-like, and tsv (generated from AnnotSV (Geoffroy *et al.* 2018))—making them flexible for use with real or synthetic perturbations. The following variant types (marked by their Manta (Chen *et al.* 2016) abbreviations) are supported: SNPs, indels, deletions (DEL), duplications (DUP), inversions (INV), and complex rearrangements with breakends (BNDs). Across a variety of datasets, including the 1K Genome Project, SuPreMo makes it possible to analyze over 50% of SVs that would not be accessible with existing tools (Fig 1B). In particular, symbolic alleles are now easily and uniformly processed (Fig 1B, orange). On the other hand, insertions, which make up <20% of SVs, remain inaccessible for sequence-based models since the precise inserted sequence is not provided by SV calling methods (Fig 1B, gray).

SuPreMo provides flexibility through the following parameters: prediction window shift, output sequence length, maximum variant length, and reverse complement. Variants

121 are centered in the output sequences unless they are too close to chromosome arm

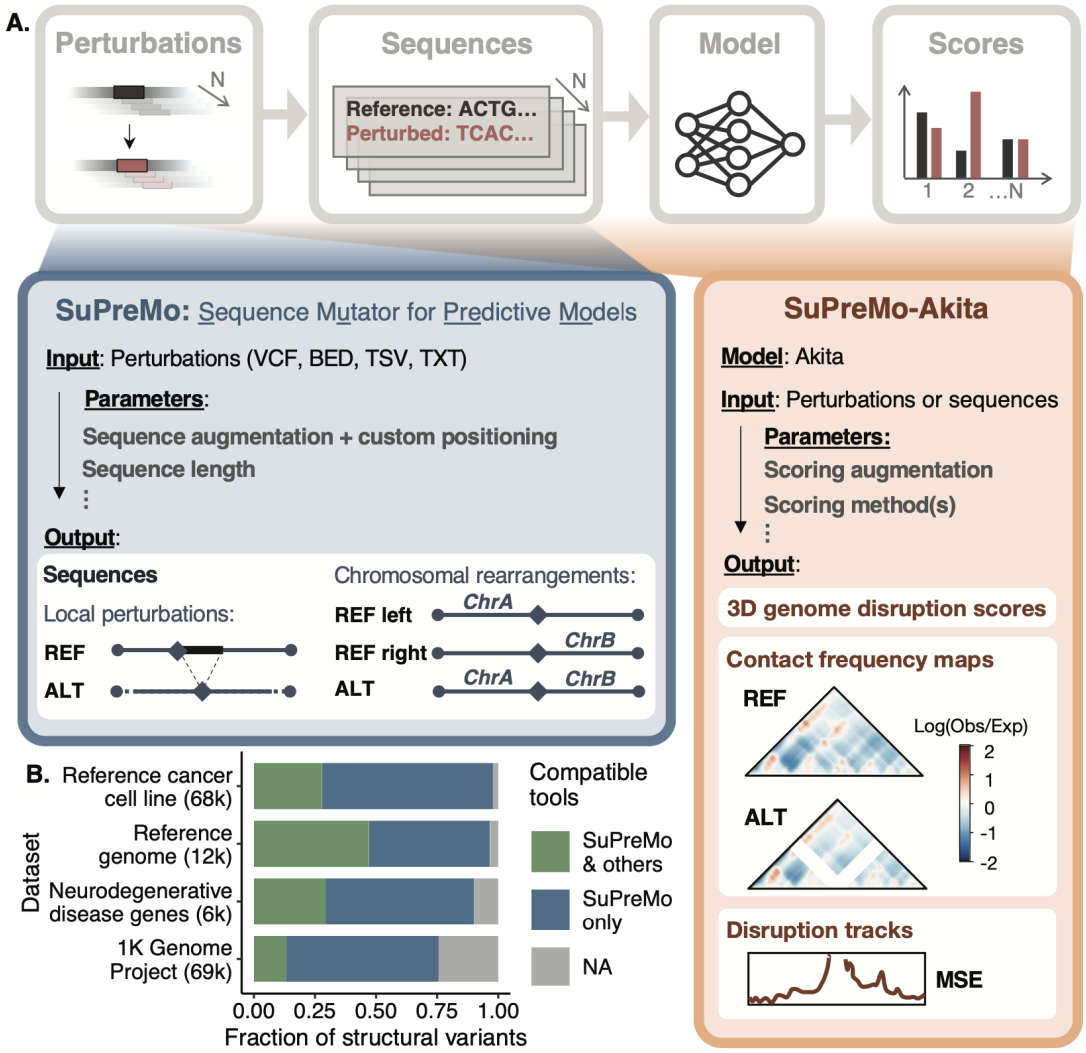

Figure 1: A: Schematic representation of SuPreMo. SuPreMo generates sequences by incorporating perturbations into the hg38 human reference genome. SuPreMo-Akita applies Akita to those sequences and generates 3D genome disruption scores (effect size of each perturbation) and, optionally, disruption tracks and predicted contact frequency maps. Parameters and outputs are specified. REF: derived from reference allele; ALT: derived from alternate allele; Log(Obs/Exp): log of observed over expected contacts; MSE: mean squared error. B: Categorization of SVs based on the ability of SuPreMo and other existing tools to incorporate them into a reference genome. SVs that other tools can already process include small indels (green); SVs that only SuPreMo can process include deletions, duplications, inversions, and chromosomal rearrangements (navy); SVs that no tool can process include insertions and copy number variants (CNVs) because the exact sequence is not provided by upstream variant calling pipelines (gray). Datasets are WGS/WES from healthy and disease individuals: a reference cancer cell line (Talsania *et al.* 2022), a reference genome (Zook *et al.* 2020), neurodegenerative disease gene sequences (Kaivola *et al.* 2023), and the 1K Genome Project (Mahmoud *et al.* 2019).

ends (Supp Fig 1B) or the shift parameter is used, which moves the prediction window around the variant and therefore changes its position in the sequence (Supp Fig 2). The length of the generated sequences can be customized to match the required input of the model of interest (Supp Table 1). For SV inputs, the maximum SV length considered is limited to two thirds of the input sequence, unless otherwise specified by the user. Lastly, the reverse complement of the generated sequence can be outputted. Generated sequences are accompanied by the relative position of the variant in each sequence. Thus, SuPreMo is a flexible tool for performing ISM that can be applied across sequence-based ML models.

SuPreMo-Akita generates an array of 3D genome disruption scores, predicted contact frequency maps for reference (wild type) and alternate (perturbed) sequences, and genomic tracks of disruption scores across the prediction window. Akita predicts contact frequency maps for a ~1 megabase (Mb) input sequence at a ~2 kilobase (kb) resolution. SuPreMo-Akita inputs variants as described above and optionally also takes in already generated sequences. Since methods for scoring contact maps are biased and sometimes only target certain features, we have made available 13 different predefined metrics (Gunsalus *et al.* 2023) to use with this tool, with the defaults being the most common measures: mean squared error (MSE) and Spearman's rank correlation coefficient (referred to here as just correlation). To assess the robustness of the generated disruption scores, the augmentation parameter optionally provides averages of scores from standard sequences, sequences with -1 bp and +1 bp shifts, and reverse complement sequences, or any other augmentations specified. Each

generated map will be accompanied by the start genomic position and the relative bin that the variant lies in.

Lastly, we considered computational efficiency. To enable customization to different hardware, the user can choose the number of rows to be processed at a time from the input file and what outputs to request, keeping in mind storage and memory limitations. We measured the run time, peak memory, and size of outputs on 3 GHz CPUs using a set of 100-1000 SVs of different types from the reference cancer cell line in Figure 1B. SuPreMo-Akita is fast and easily scaled up—with the augmentation parameter it takes approximately 19 seconds and accumulates ~ 0.15 megabytes (MB) of memory per variant (Supp Table 2).

We implemented SuPreMo using two models, although our framework is extendable to any model utilizing genome sequences as input. First, we used SuPreMo with DeepSEA (Zhou and Troyanskaya 2015) to rank a set of CTCF deletions based on their predicted effect on epigenetic marks. Second, we used SuPreMo-Akita on cancer SVs (Supp Fig 3A). SVs were scored using MSE and correlation, and the top 3 scoring variants for each SV type and scoring method were selected (Supp Fig 3B). We separately ranked variants by their type because their 3D genome disruption scores vary, and by the scoring method because each has unique biases. Using SuPreMo-Akita, contact frequency maps and disruption tracks were generated for these selected SVs and the most interesting variants, based on the structures they disrupt, were chosen (Supp Fig 3C). This method prioritized a deletion of an insulated site that is predicted to cause

increased contact frequency between neighboring regions (Supp Fig 3C, left panel). Step-by-step instructions for both implementations are available on Github.

3 Conclusion

SuPreMo is a software tool that facilitates ISM with predictive models and extends this principle with Akita to predict scores for 3D genome folding disruption. Potential use cases include scoring all variants in an individual or cohort for disruption to genome folding, generating predicted contact frequency maps to explore the effects of noncoding variants on regulatory interactions, performing ISM to evaluate or discover sequence motifs using Akita, and, more broadly, generating sequences for input into predictive models of interest to evaluate variant effects. SuPreMo is scalable to a large number of variants and only limited by the storage capacity the user has for the expected outputs. Overall, SuPreMo allows for easy, fast and broadly applicable analysis of simple variants, SVs, and chromosomal rearrangements in the context of sequence-based predictive models.

Acknowledgements

We thank Shu Zhang for helpful discussion and feedback and for reviewing the manuscript. We thank Maureen Pittman, Laura Gunsalus, and Shuzhen Kuang for helpful insight on using Akita.

Funding

This work was supported by the National Institutes of Health [grant number R03OD034499 and grant number U01HL157989] and Gladstone Institutes.

1000 Genomes Project Consortium, Auton A, Brooks LD *et al.* A global reference for human genetic variation. *Nature* 2015;**526**:68–74.

Agarwal V, Shendure J. Predicting mRNA Abundance Directly from Genomic Sequence Using Deep Convolutional Neural Networks. *Cell Rep* 2020;**31**:107663.

Avsec Ž, Agarwal V, Visentin D *et al.* Effective gene expression prediction from sequence by integrating long-range interactions. *Nat Methods* 2021a;**18**:1196–203.

Avsec Ž, Weilert M, Shrikumar A *et al.* Base-resolution models of transcription-factor binding reveal soft motif syntax. *Nat Genet* 2021b;**53**:354–66.

Benegas G, Batra SS, Song YS. DNA language models are powerful predictors of genome-wide variant effects. *bioRxiv* 2023:2022.08.22.504706.

Chen KM, Wong AK, Troyanskaya OG *et al.* A sequence-based global map of regulatory activity for deciphering human genetics. *Nat Genet* 2022a;**54**:940–9.

Chen V, Yang M, Cui W *et al.* Best Practices for Interpretable Machine Learning in Computational Biology. *bioRxiv* 2022b:2022.10.28.513978.

Chen X, Schulz-Trieglaff O, Shaw R *et al.* Manta: rapid detection of structural variants and indels for germline and cancer sequencing applications. *Bioinformatics* 2016;**32**:1220–2.

Danecek P, Auton A, Abecasis G *et al.* The variant call format and VCFtools. *Bioinformatics* 2011;**27**:2156–8.

Deng C, Whalen S, Steyert M *et al.* Massively parallel characterization of psychiatric

disorder-associated and cell-type-specific regulatory elements in the developing human cortex. *bioRxiv* 2023, DOI: 10.1101/2023.02.15.528663.

Fudenberg G, Kelley DR, Pollard KS. Predicting 3D genome folding from DNA sequence with Akita. *Nat Methods* 2020;**17**:1111–7.

Geoffroy V, Herenger Y, Kress A *et al.* AnnotSV: an integrated tool for structural variations annotation. *Bioinformatics* 2018;**34**:3572–4.

Gosai SJ, Castro RI, Fuentes N *et al.* Machine-guided design of synthetic cell type-specific cis -regulatory elements. *bioRxiv* 2023, DOI: 10.1101/2023.08.08.552077.

Gunsalus LM, Keiser MJ, Pollard KS. In silico discovery of repetitive elements as key sequence determinants of 3D genome folding. *Cell Genomics* 2023;**3**:100410.

Gunsalus LM, McArthur E, Gjoni K *et al.* Comparing chromatin contact maps at scale: methods and insights. *bioRxiv* 2023, DOI: 10.1101/2023.04.04.535480.

Hoffman GE, Bendl J, Girdhar K *et al.* Functional interpretation of genetic variants using deep learning predicts impact on chromatin accessibility and histone modification. *Nucleic Acids Res* 2019;**47**:10597–611.

Kaivola K, Chia R, Ding J *et al.* Genome-wide structural variant analysis identifies risk loci for non-Alzheimer's dementias. *Cell Genom* 2023;**3**:100316.

Kelley DR. Cross-species regulatory sequence activity prediction. *PLoS Comput Biol* 2020;**16**:e1008050.

Kelley DR, Snoek J, Rinn JL. Basset: learning the regulatory code of the accessible

genome with deep convolutional neural networks. *Genome Res* 2016;**26**:990–9.

Keough KC, Whalen S, Inoue F *et al.* Three-dimensional genome rewiring in loci with human accelerated regions. *Science* 2023;**380**:eabm1696.

Li H. A statistical framework for SNP calling, mutation discovery, association mapping and population genetical parameter estimation from sequencing data. *Bioinformatics* 2011;**27**:2987–93.

Mahmoud M, Gobet N, Cruz-Dávalos DI *et al.* Structural variant calling: the long and the short of it. *Genome Biol* 2019;**20**:246.

McArthur E, Rinker DC, Capra JA. Quantifying the contribution of Neanderthal introgression to the heritability of complex traits. *Nat Commun* 2021;**12**:4481.

Nguyen E, Poli M, Faizi M *et al.* HyenaDNA: Long-Range Genomic Sequence Modeling at Single Nucleotide Resolution. *ArXiv* 2023.

Paik S, Maule F, Gallo M. Dysregulation of chromatin organization in pediatric and adult brain tumors: oncoepigenomic contributions to tumorigenesis and cancer stem cell properties. *Genome* 2021;**64**:326–36.

Rivas-Astroza M, Xie D, Cao X *et al.* Mapping personal functional data to personal genomes. *Bioinformatics* 2011;**27**:3427–9.

Schwessinger R, Gosden M, Downes D *et al.* DeepC: predicting 3D genome folding using megabase-scale transfer learning. *Nat Methods* 2020;**17**:1118–24.

Talsania K, Shen T-W, Chen X *et al.* Structural variant analysis of a cancer reference

cell line sample using multiple sequencing technologies. *Genome Biol* 2022;**23**:255.

Tan J, Shenker-Tauris N, Rodriguez-Hernaez J *et al.* Cell-type-specific prediction of 3D chromatin organization enables high-throughput in silico genetic screening. *Nat Biotechnol* 2023;**41**:1140–50.

Van der Auwera GA, O'Connor BD. *Genomics in the Cloud: Using Docker, GATK, and WDL in Terra*. "O'Reilly Media, Inc.," 2020.

Zhou J. Sequence-based modeling of three-dimensional genome architecture from kilobase to chromosome scale. *Nat Genet* 2022;**54**:725–34.

Zhou J, Theesfeld CL, Yao K *et al.* Deep learning sequence-based ab initio prediction of variant effects on expression and disease risk. *Nat Genet* 2018;**50**:1171–9.

Zhou J, Troyanskaya OG. Predicting effects of noncoding variants with deep learning–based sequence model. *Nat Methods* 2015;**12**:931–4.

Zook JM, Hansen NF, Olson ND *et al.* A robust benchmark for detection of germline large deletions and insertions. *Nat Biotechnol* 2020;**38**:1347–55.