Mapping MAVE data for use in human genomics applications

Jeremy A. Arbesfeld[1,2], Estelle Y. Da[3], Kori Kuzma[2], Anika Paul[2], Tierra Farris[4], Kevin Riehle[4], Nuno Daniel Saraiva Agostinho[5], Jordan F. Safer[6], Aleksandar Milosavljevic[4], Julia Foreman[5], Helen V. Firth[5], Sarah E. Hunt[5], Sumaiya Iqbal[6], Melissa Cline[7], Alan F. Rubin[3,8], Alex H. Wagner[1,2]

[1]The Ohio State University, Department of Biomedical Informatics, Columbus, OH
[2]Nationwide Children's Hospital, The Steve and Cindy Rasmussen Institute for Genomic Medicine, Columbus, OH
[3]WEHI, Bioinformatics Division, Parkville, Australia
[4]Baylor College of Medicine, Department of Molecular and Human Genetics, Houston, TX
[5]European Molecular Biology Laboratory, European Bioinformatics Institute, Wellcome Genome Campus, Hinxton, Cambridge, CB10 1SD, United Kingdom
[6]The Broad Institute of MIT and Harvard, The Center for the Development of Therapeutics, Cambridge, MA
[7]University of California, Santa Cruz, BRCA Exchange, Santa Cruz, CA
[8]University of Melbourne, Department of Medical Biology, Parkville, Australia

Corresponding Authors:

Alan Rubin, PhD, WEHI, 1G Royal Pde, Parkville VIC 3052, Australia. Email: alan.rubin@wehi.edu.au

Alex Wagner, PhD, Nationwide Children's Hospital, Room WB3155, Research Building 3, 575 Children's Crossroad, Columbus, OH 43210. Phone: 614-355-1645, Email: Alex.Wagner@nationwidechildrens.org

Abstract

The Data Coordination and Dissemination Workstream of the Atlas of Variant Effects Alliance developed a novel method to map multiplexed assays of variant effect data to human reference sequences. This method was applied to 209 human score sets in MaveDB, across approximately 2.5 million protein and genomic variants, and successfully mapped 99.00% of examined variants. This effort enables the standardization and dissemination of MAVE data to precision medicine resources including the UCSC Genome Browser, ultimately assisting downstream users in the clinical assessment of variant pathogenicity.

Keywords

multiplexed assays of variant effect, variation representation specification, deep mutational scanning, massively parallel reporter assays, global alliance for genomics and health

Background

With the growing use of high-throughput sequencing technologies in the clinical setting, shortfalls in available genomic data to drive variant interpretation are increasingly observed. Specifically, clinical assessment of variants for pathogenicity are often inconclusive, with 49.80%(1) (1,091,123/2,191,109) of curated variants in ClinVar classified as "variants of uncertain significance" (VUS) due to the lack of clear evidence supporting or refuting pathogenicity(2). In silico prediction tools exist, but by themselves do not provide functional data that can help inform pathogenicity classification(3). In recent years, multiplexed assays of variant effect (MAVEs) have been introduced as a new line of functional evidence to support variant classification in a growing number of genes. MAVEs serve as useful tools for measuring the effects of variation on phenotype for thousands of variants in parallel(4). Commonly used MAVE designs include deep mutational scanning(5), in which the functional effects of protein variation in response to a selective pressure are described (6), and massively parallel reporter assays (MPRAs)(7), which validate the functions of different regulatory elements(8). As MAVEs can produce functional scores for many variants chosen systematically, they are able to identify variants of unknown significance (VUS) before they are detected in a clinical context, providing evidence that can ultimately assist in clinical variant interpretation (9)(10).

The increased use of MAVE experimental methods led to a need for central repositories for MAVE experiment metadata, assayed variants, and associated functional scores. In 2019, MaveDB(11) became the first such publicly-accessible resource. To improve data accessibility and discoverability, MaveDB organizes datasets in a hierarchical fashion using Experiment Sets, Experiments, and Score Sets. Experiment Sets are logical containers for linking multiple Experiment records, typically when multiple functional assays were performed on a single target as part of a study. Experiment records describe a single assay condition and its descriptive metadata, including the methods used and links to the raw sequence reads. Score Sets contain the score and optional count data for each variant measured in the assay and include details of the computational and statistical analysis performed as well as the target sequence. With nearly 300 submitted experimental datasets in MaveDB at the time of writing and more submitted every month, there is clear value in developing standards for the representation and exchange of these data, as well as standard methods for how these data may be calibrated and applied to support the clinical classification of genomic variants. The Atlas for Variant Effects Alliance (AVE; varianteffect.org) is a consortium working to realize these goals.

Under the auspices of the AVE Data Coordination and Dissemination (DCD) work stream, we have addressed the challenge of precisely mapping MAVE data to human reference sequences, and computationally represent computational homology mappings using the Global Alliance for Genomics and Health (GA4GH) Variation Representation Specification (VRS) (12) and the Sequence Ontology (SO)(13). For a given entry in MaveDB, a MAVE dataset describes sequence changes in MAVE-HGVS(14), a format similar to the Human Genome Variation Society (HGVS) variant nomenclature(15). MAVE-HGVS variants are then associated with their respective functional scores. These variants are represented with respect to a *target sequence* uploaded by the submitter; however, as this target sequence is not described with respect to human reference sequence sets (e.g. Ensembl/GENCODE(16), RefSeq(16,17), GRC genome assemblies), a challenge emerges concerning the standardized representation of variation. Specifically, while variants in HGVS are described in relation to accepted reference sequences (e.g. NP_003997.1:p.Trp24Cys), the vast majority of variants in MAVE score sets in MaveDB are described only in the context of the assayed sequence. For example, variant representations in MaveDB include "p.Ala40Ser" and "n.2G>A", describing changes with respect to a provided *target sequence* stored in the MAVE experiment data. Furthermore, the target sequences in MaveDB may not be reference-identical. While necessary for the precise description of observed variants in an experimental setting, this design presents challenges to interoperability between MAVE datasets and variants described on human reference sequences, including variants routinely reported by clinical sequencing pipelines.

To resolve this limitation, we developed a method for consistently mapping MAVE variants to human reference sequences while preserving the original sequence context, improving provenance and interoperability between MaveDB and existing applications based on human reference sets (**Figure 1**). Using open-source tools and databases(12,18–20), we generated a MAVE dataset mapping for the FAIR(21) and computable exchange of variation data with other datasets and tools. We describe the integration of these mapped data into several common tools used for human genomics research and clinical variant curation, including the UCSC Genome Browser(22), the Ensembl Variant Effect Predictor (VEP)(23), the Broad Institute Genomics 2 Proteins Portal(24), the ClinGen Data Platform(25), the DECIPHER resource(26), and Shariant(27). Our mapping approach closes an important gap for the application of MAVE data in genomic medicine and human health research.

Results

Composition of MaveDB Score Sets

Human score sets were selected for testing of the variant mapping functionality. At the time of writing, 209 of an available 299 score sets in the MaveDB had the *Homo sapiens* classification, together totaling 2,499,036 variants, providing a large and heterogeneous dataset upon which the mapping process could be developed and tested. Of the 209 selected score sets, 168 described protein coding elements while the remaining 41 covered regulatory and other noncoding elements, respectively. Among the 209 examined score sets, 176 contained DNA target sequences while the remaining 33 contained protein target sequences (**Figure 2a**). While the human score sets were designed to reveal insight into human biology, 62/159 unique experiments (38.99%) contained data from experiments that were conducted in genetic models not derived from human cells, including organisms such as yeast, mice, and bacteria (**Figure 2b** and **Supplemental Table 1**).

Most common among these experiment types were the 168 score sets describing protein coding variation, 95 of which reported associated UniProt(28) accessions. Variants from these score sets were described on either genomic (135 score sets) or protein (33 score sets) target sequences. The mean length of protein sequences (representing targeted regions of expressed human proteins) was approximately 573 amino acid residues, and the DNA target sequences had an average target sequence length of 1337 nucleotides. All 41 regulatory/other noncoding score sets reported DNA target sequences with an average length of 353 nucleotides. Gene symbols or aliases (e.g. CALM1, CBS) were provided for 125/168 protein coding score sets, and the remaining 43/168 reported specific domains and/or targets (e.g. hYAP65 WW domain, Src catalytic domain). All 41 regulatory/other noncoding category score sets described specific regulatory elements such as promoters (e.g PKLR promoter) or enhancers (e.g. IRF6 enhancer).

Mapping of MaveDB Score Set Variants to Human Reference Sequences

Variant mapping from human score sets was accomplished in three sequential steps (**Figure 3, Methods**). As homologous sequence annotations were not universally available across MAVE experiments, our first step was to use BLAT(20) to align MaveDB target sequences to the GRCh38 human genome assembly. Using this initial alignment data, we next computationally inferred compatible reference sequences (i.e. compatible RefSeq transcripts) associated with the score set, and the corresponding sequence offset for the MAVE target sequence. MAVE score set variants were then translated with respect to the associated RefSeq protein sequence or transcript-aligned genomic regions as appropriate (**Methods**) using the biocommons SeqRepo(18) python package and Universal Transcript Archive (UTA)(19) database and GenomicMedLab Common Operations on Lots of Sequences (cool-seq-tool)(29) python package. For score sets with regulatory/other noncoding elements, this second step was skipped, as the target could be directly mapped to a contiguous chromosomal region as aligned by BLAT.

Lastly, the MAVE variants and mapped homologous variants were each translated into VRS alleles or haplotypes as necessary (**Figure 3**) and output along with the associated score data to comprise the resultant

mapping set. Of the 209 human score sets that were potentially available for analysis, 207 were processed using the variant mapping algorithm. Of the two score sets that failed to map, the first (urn:mavedb:00000072-a-1) reported a UniProt accession that lacked any linkage to a corresponding RefSeq protein accession. The other score set (urn:mavedb:00000105-a-1) was unable to return a BLAT hit for the reference sequence. For the 207 human score sets in MaveDB, our method successfully mapped 99.00% (2,474,158/ 2,499,036) examined MAVE variants, where a successful mapping is defined as equivalence in the reference allele sequences of each pre-mapped and post-mapped variant pair. The remaining 1.00% (24,878/ 2,499,036) of discordant variants were caused by protein changes in the MAVE score set mapping across exon boundaries, non-homologous sequence content from MAVE target sequences preventing a reference match, and supplied variants occurring past the original target sequence length.

Integrating Data from Mapped MaveDB Variants into Genomic Databases and Tools

MaveDB API

MaveDB includes a FastAPI-driven API, specified at https://api.mavedb.org/docs. The API includes both pre- and post-mapped VRS objects for all variants mapped as part of this study, accessible as JavaScript Object Notation (JSON) using the MaveDB API `/mapped-variants` endpoint. These mappings are also available within the bulk downloads for mapped datasets. Future releases of MaveDB will further integrate native VRS support into the platform, by performing automatic mapping of submitted datasets using the protocols described here, and adding variant search capabilities using the contents of the mapped VRS objects. VRS objects will also power integration across different sources of genomic data as described here, and provide structured data for visualizations and other exploratory tools.

Genomics2Proteins Portal

Genomics 2 Proteins (G2P) portal is a discovery portal for linking genomic data to protein sequence and structure. The portal provides a user interface for exploring genetic variations, readouts from genetic perturbation assays, and protein features on protein sequence and structure at the amino acid residue level to help interpret the molecular effect of variations. The portal integrates data from large genomic (gnomAD, ClinVar and HGMD) and proteomic databases (UniProt, PDB, AlphaFold) as well as enables users to perform customized mapping of genetic variations to proteins (**Figure 4a**).

168 score sets, describing MAVEs for protein coding variations in 40 unique human genes, have been integrated into the G2P portal. In addition to single amino acid residue substitutions ("point mutation"), MAVEs were available for pairwise residue mutations ("pairwise mutation") for 7 out of 40 genes (Supplementary Figure 1a). While comparing to the length of the canonical protein isoform from Uniprot, MAVE data were mappable for >90% of the residues of 24 proteins (Supplementary Figure 1b).

All MAVEs for both point and pairwise mutations for each gene and score set are displayed in G2P as heatmaps (Supplementary Figure 1c-d). Additionally, for each gene and score set pair, mutations with top and bottom 99th percentile of MAVEs were displayed in the context of protein sequence annotations of structural (secondary structure, residues' solvent accessibility, etc.) and functional features (e.g., domain, active sites) (Supplementary Figure 1e). For example, MAVE readouts for *TP53* and score set urn:mavedb:00000068-b-1 range from -5.39 to 2.80. Mutations with scores >1.92 (top 99th percentile) and < -2.61 (bottom 99th percentile) were annotated in the "Protein sequence annotation" viewer of the portal. The filtering was performed for a clear visualization. These top and bottom 99th percentiles of MAVEs can also be downloaded as tabular format

from the portal. In the future release, MAVEs will be mapped on the protein structure positions of mutations. The integration of MAVEs with protein sequence and structural features will facilitate interpreting MAVE data.

UCSC Genome Browser

The UCSC Genome Browser(22) is a widely-used and highly-customizable web tool supporting genome research, and includes annotations from many datasets of clinical and research relevance. A genome browser Track Hub has been created supported by these mappings, which displays the protein variant in genomic coordinates, and the associated scores. The track hub renders these scores as a heat map, in which each column represents the mapped genomic location of the variants scored, each row represents an alternate allele and cells are colored on a blue/yellow color spectrum in proportion to the score. This supports quick visual analysis to relate the functional impact of variants to their genomic context (**Figure 4b**). The MAVE data track in the genome browser can be found at http://genome.ucsc.edu/s/mcline/mavedb.

Ensembl VEP

Ensembl VEP is an open-source tool for the annotation and prioritization of genomic variants. It aggregates annotations of currently available knowledge about variant loci via three interfaces, designed to suit different use cases: 1) a highly configurable command line tool, 2) a REST API, and 3) a simple web interface. The genomic mapping of MaveDB data from this work will be used to ingest MAVE score sets for use in Ensembl VEP. Extensions to report these data will initially be developed for the command line tool and we plan to update the Ensembl VEP REST and web interfaces to make open access MaveDB datasets available in our June 2023 release (**Figure 4c**). This integration will enable easy access to these data and convenient integration into large-scale annotation pipelines.

DECIPHER

DECIPHER (https://www.deciphergenomics.org/) is a global resource that shares phenotype linked variant data from rare disease patients to support research and diagnosis, and among other products provides variant interpretation interfaces (30)(26). The mapping of MAVE data to reference sequence genomic coordinates will enable DECIPHER to ingest MAVE data and display this data in user interfaces, enhancing accessibility of this information to the clinical community. MAVE data have already been incorporated into some ClinGen Expert Panel ACMG/AMP variant interpretation guidelines, e.g. for variants in TP53 associated with Li-Fraumeni syndrome(31). The MAVE data in DECIPHER will be displayed on the functional data tab (e.g. https://www.deciphergenomics.org/sequence-variant/2-47478337-G-A/annotation/functional) which is accessed from DECIPHER patient records in addition to variant pages and protein variant pages accessed through the site search tools. These data have a planned release in the DECIPHER resource in late 2023. Displaying these data in DECIPHER will increase the discoverability of the MAVE data for clinicians, clinical scientists, clinical researchers, research scientists and curators who use DECIPHER, empowering its use in variant interpretation and assisting DECIPHER's mission of mapping the clinically relevant elements of the human genome.

ClinGen Linked Data Hub

The ClinGen Linked Data Hub (LDH) (https://ldh.clinicalgenome.org) is a RESTful API service built on Linked Open Data principles(32) that aggregates excerpts of pertinent variant data from a variety of external sources to contribute supporting evidence required for use in the classification of variant pathogenicity in the ClinGen Data Platform(25). The LDH works in concordance with the ClinGen Allele Registry(33) which is a canonical

on-demand variant naming service. The mapped variants were submitted to the ClinGen Allele Registry and assigned Canonical Allele Identifiers (CA ID) and Protein Allele Identifiers (PA ID) for ingestion into the LDH.

Users can access the MaveDB data via the LDH API (note: UI also provides basic search functionality) by either using the LDH MaveDbMapping document's entity ID (score set accession + "#" + variant number; e.g. urn:mavedb:00000001-a-1#1) or by searching for the associated variant CA ID or PA ID. Accessing the MaveDbMapping documents using the variant CA ID or PA ID allows users to easily access MaveDB data for the variant of interest from multiple MaveDB experiments or score sets simultaneously alongside pertinent data from other sources. The LDH API can also be used to return all MaveDbMapping documents from a particular score set, enabling bulk usage. Leveraging both ClinGen CA IDs / PA IDs and GA4GH VRS IDs allows for straightforward data aggregation of variants by identifier from groups that leverage one or multiple data standards and provides the users with the level of specificity required for their application.

Shariant

The Shariant(27) platform is configured to consume Canonical Allele Identifiers from the ClinGen Allele Registry so the initial data exchange between MaveDB and Shariant will be accomplished using this system. This underscores the importance of integrating and supporting data standards, as the submission of VRS objects to the ClinGen Linked Data Hub is an essential step for generating the Canonical Allele Identifiers that Shariant requires. MaveDB data will be made available to Shariant users following CA ID and PA ID registration and incorporated into a planned platform update in late 2023.

Discussion

In this study, we address the challenge of mapping multiplexed assay of variant effect (MAVE) data to human sequence assemblies for use in human research and clinical applications. We introduced and evaluated a method for accomplishing this using the GA4GH Variation Representation Specification (VRS)(12) and associated open-source bioinformatics tools. Our approach is informed by FAIR data principles and enables semantically-precise representation of these homology maps for data provenance.

However, our approach fails to map a small percentage of variants (1.00%; 24,878/ 2,499,036) due to a lack of policy about how such variants should be resolved. For example, when a DNA target sequence representing a processed RNA aligns to disparate regions of the genome, it is difficult to map insertions or deletions that span intron-exon boundaries, as the corresponding reference sequence genomic coordinates can cover thousands of nucleotides. Should these be treated as multiple insertion/deletion events, or as a single, very large event that also covers the intronic space? In other cases, a segment of a target sequence may not align to known human reference sequences, limiting our ability to interpret variants that are reported at those unaligned positions of the target sequence. These are areas of ambiguity that would benefit from development of expert recommendations by the Atlas for Variant Effects Alliance (AVE).

Our study also revealed diversity in the way variant knowledge is consumed and used by various downstream tools. GA4GH VRS provides a precise mechanism for addressing the complexity of representing variants on assay-specific sequences, and has useful characteristics for addressing variant overprecision and globally unique variant identification(12). However, its relatively recent emergence as a variant representation standard required additional mechanisms for enabling downstream connectivity to other resources, representing human-mapped variants as HGVS. To address this, we have used open-source translation tools to annotate all mapped variants using HGVS(34). We have also developed methods for mapping protein variants to the genomic reference space. Altogether, our approach has enabled integration of MAVE data into the Ensembl VEP, UCSC Genome Browser, Broad Institute Genomics 2 Protein Portal, and ClinGen Linked Data Hub resources, with additional integrations forthcoming.

A remaining limitation for the applicability of these data to human clinical datasets is the development of expert guidelines for calibrating and scoring MAVE assays, work that is ongoing in the context of AVE. The responsible use of this data for clinical purposes relies on there being high confidence that the MAVE assay relates to the mechanism of disease. For each specific gene-disease relationship in question, ensuring that the MAVE assay is a valid predictor of pathogenicity will be essential. This is likely to be especially challenging for genes with multiple different disease associations and for proteins with multiple functional domains. We believe that the mapping of these data to human reference sequence assemblies provides a crucial foundation for the development of AVE guidelines for the use of MAVE data, by providing common sequence assemblies for the evaluation of MAVE score sets.

Conclusions

The impact of genome and exome sequencing on human research and clinical practice is hindered by challenges in variant interpretation. Multiplexed Assays of Variant Effect (MAVEs) provide a high-throughput functional assessment tool for variants in genes of relevance to human health and disease, and hundreds of MAVEs have been developed and submitted to the centralized MaveDB data repository. We developed a method for precise representation and mapping of MAVE data to human reference sequences and application of this method to score sets in MaveDB. We demonstrated how this process enables the mapping of 99.00% (2,474,158 / 2,499,036) of MAVE variants in MaveDB for ready integration into multiple variant annotation and evaluation platforms. We discussed current challenges in the use of these data in human genomics applications, and the need for expert communities like the Atlas of Variant Effects Alliance (AVE) to address these remaining gaps. We believe the mapped data from this study will help advance those efforts, and the data integrations at the UCSC Genome Browser, Genomics2Proteins Portal, Ensembl VEP, ClinGen Linked Data Hub, and others will be provide useful tools for advancing MAVE-informed genomic variant interpretation efforts.

Methods

Extraction of Metadata from the MaveDB API Score Sets Endpoint

For score sets whose listed target organism was *Homo sapiens*, seven variables were extracted from the MaveDB API score sets endpoint. These variables were: target sequence (string of nucleotides or amino acids), target sequence type (DNA or protein), target (e.g. CXCR4), assembly ID, UniProt ID, target type (e.g. protein coding), and URN (e.g. urn:mavedb:00000048-a-1). These data elements were extracted from the MaveDB API as they were the minimum information needed to determine the genomic coordinates targeted by an assay in a MaveDB score set.

Alignment of Target Sequences to Human Genome

Having extracted the necessary metadata from the MaveDB API, the initial step was to align target sequences to the human genome, allowing for the genomic coordinates of the examined sequences to be determined (**Figure 3**). To achieve this aim, all target sequences were run through BLAT against the GRCh38 human genome assembly. Depending on whether a target sequence was composed of nucleotides or amino acids, the BLAT query (q) argument was set to "dna" or "prot" to maximize the probability of returning high quality hits. In addition, the minimum score argument (minScore) was reduced to 20 from a default value of 30 to ensure that BLAT would return hits from short, specific target sequences.

After running BLAT, a series of steps were followed to ensure that suitable genomic coordinates were located. When running BLAT locally and outputting the results in a PSL file, different "hits" are reported at the chromosomal level (e.g. "chr3"). Within a hit are specific "HSP" (high-scoring pair) objects that describe regions of concordance between the queried sequence and the human genome. With the goal of ultimately selecting

the correct HSP object, the correct chromosomal hit needs to be chosen first. To maximize this probability, when a UniProt accession was available for a score set, the UniProt accession was supplied as a CURIE (e.g. uniprot:P12931) to the normalize method in the VICC Gene Normalization Service(35), returning an HGNC(36) consensus gene symbol and a chromosome number indicating where the gene occurs. Based on the chromosome number returned from the normalize method, the BLAT PSL file was filtered to only include the hit that contained the correct chromosome. If the filtered hit contained more than one HSP object, additional processing was required to select the correct genomic coordinates. First, the gene symbol corresponding to the provided UniProt accession was supplied to the search method within the Gene Normalization Service, returning location data, in genomic coordinates, from Ensembl(37) and NCBI(38). Specifically, these two sources had "start" and "end" attributes describing the location where the supplied gene of interest occurs on a chromosome. Using the position indicated in the start attribute, the HSP object with the minimum distance to the start position was selected.

Having performed a series of filtering and validation steps to ensure that the genomic coordinates with the highest potential accuracy were selected, alignment data was added to a dictionary that was keyed by score set URN. For each score set, the supplied data included: chromosome number, strand orientation, target name, target type, UniProt ID, percent coverage, percent identity, and a dataframe reporting the genomic coordinates supplied by the selected HSP object.

While the alignment procedure performs well for protein coding score sets with UniProt accessions, its efficacy is potentially limited for regulatory/other noncoding score sets. Specifically, these score sets lack UniProt accessions and often have more descriptive target names without gene names (e.g. hYAP65 WW domain); the ability to extract consensus gene symbols for these score sets is limited, thereby resulting in a potential inability to perform the hit/HSP filtering procedure. In these instances, the top scoring hit reported by BLAT was selected, but the additional validation steps described previously were not performed.

RefSeq Transcript Selection and Offset Computation

Having extracted genomic coordinates and other relevant information such as the chromosome number, there was sufficient data for selecting an appropriate human reference sequence for each score set. For the regulatory/other noncoding score sets that only reported genomic variation, the RefSeq chromosomal genomic reference sequence for GRCh38 was selected since the supplied genomic coordinates described locations on the sequence. However, for protein coding score sets that reported protein variants, a RefSeq protein reference sequence was needed, requiring use of the Biocommons UTA and SeqRepo databases, respectively, and associated cool-seq-tool(29) translation service.

First, using SeqRepo, a database that stores human reference sequences and links between different identifiers, chromosome numbers were converted to their GRCh38 RefSeq accession. Additionally, consensus gene symbols were derived by leveraging the Gene Normalization service that was utilized in the alignment procedure. Lastly, a query was run against UTA that extracted all the transcripts found within the genomic start and end coordinates, supplied by the HSP fragment, that were associated with the derived chromosome accession and gene symbol, and stored the identifiers in a list. This query was run for all rows in the HSP dataframe; for example, if an HSP object from a score set reported six fragments, six lists were generated. Once all lists were created and any non-coding transcript accessions were removed, the intersection of the lists was taken and supplied as input to the get_mane_from_transcripts() method in cool-seq-tool. If the method returned a nonempty list, the following prioritization ranking was followed for transcript selection: MANE Select, MANE Plus Clinical. If the list was empty, the lengths of all transcripts in the intersected list were taken, the longest transcript was chosen, and the protein RefSeq accession associated with the selected transcript was found.

To determine the exact location of the provided target sequence within the RefSeq protein sequence, the target sequence, if DNA, was converted to protein using the standard codon table and the first 10 amino acids were extracted. Then, by accessing the RefSeq sequence using SeqRepo and using the find() method, the substring's location was found; find() was run again a second time on the entire converted protein sequence, producing a boolean that indicated if the *entire* target sequence was an exact substring of the RefSeq sequence. When the intersection procedure returned a non-RefSeq sequence, web scraping was performed to extract the canonical UniProt sequence and find() was again run twice. Lastly, the protein reference sequence, offset, score set accession, transcript and MANE status, and boolean were stored as entries in a dictionary and saved to a pickle file.

After an initial offset was computed for each protein coding score set, an additional pass through the score sets was completed to check for possible discordance between the provided target sequences and corresponding variant matrices. Assume a target sequence is a string of amino acid residues and reports Leucine at position 100. If the score matrix reports a substitution such as "p.Arg100His," this would be an example of discordance as the expected reference amino acid would be Leucine instead of Arginine. To identify instances of such discordance, for each score set, the provided target sequence, if DNA, was first converted to a protein sequence using the standard codon table. Then, the variant matrix was parsed to assemble a dictionary reporting the expected reference amino acid at each position (ex. 1: M, 2: A, etc.). By comparing the expected amino acids from the dictionary to those in the corresponding positions in the target sequence, the correct start location in the target sequence was determined and the offset was modified to reflect this position.

The general offset computed by the described procedure was applicable for protein variants in protein coding score sets. For genomic variants in protein coding score sets, the correct mapped position was found by determining the alignment block that the variant was found in and computing the distance between the variant position and the start (for positive strand) or end (for negative strand) genomic coordinate of the selected block, depending on the orientation of the target sequence. The same logic was followed for regulatory/other noncoding score sets with genomic variants that had only one alignment block.

Mapping Variants using VRS

With the relevant human reference sequence data determined, variant locations reported in MaveDB score set matrices were updated and the variants themselves were expressed as VRS objects. First, for a given row in a score set, the variant was first converted to a VRS object as is, allowing for the VRS representation of the assayed variant to be generated. Specifically, the reported positions were left unchanged, a new sequence digest was computed using the sha512t24u digest algorithm, the allele was normalized using the SPDI Variant Overprecision Correction Algorithm (VOCA)(39), and the appropriate allele digest was determined using the VRS ga4gh_identify() method. In instances where multiple variants were indicated with the semicolon character, this process was run separately for each variant and a VRS haplotype was generated. When a variant type not supported by VRS appeared, although a computable identifier could not be assigned, VRS text variation was generated.

The process described above was then repeated using the information derived from the reference sequence selection procedure. Using the get_metadata() method in SeqRepo, the sequence digest was updated to describe the digest for the human reference sequence while the offset was added to start and end position values, respectively. A new allele digest was computed, allowing for the "mapped" variant to have a distinct identifier from the "assayed" variant. When multiple variants were reported, new VRS alleles were created for each variant and combined in a VRS haplotype. The assayed and mapped variant representations for each row were then converted to dictionaries, allowing for storage in a two-column dataframe.

It is possible that an individual uploading a score set to MaveDB could choose to represent variants using a human reference sequence. For example, score sets in experiment set urn:mavedb:00000097 report variants in the "hgvs_prot" column with a RefSeq protein accession (e.g. NP_009225.1:p.Pro1659Leu). When this occurred, the variant was directly converted to a VRS allele using the translate_from() method and the assayed and mapped representations of the variant were equivalent.

As the regulatory/other noncoding score sets only reported variants in the "hgvs_nt" column, one dataframe of VRS representations was produced for each of these score sets. However, as protein coding score sets were capable of having data in both the "hgvs_nt" column *and* "hgvs_prot" column, variants in both columns, if present, were mapped, resulting in two potential dataframes per score set. Once mappings were complete for all score sets, two dictionaries, keyed by score set URN and containing the mappings for protein coding and regulatory/other noncoding score sets were generated and saved as pickle files.

Annotating and Validating VRS Alleles/Haplotypes Reference Sequences

The pickle files containing the assayed and mapped variants were further annotated with the "vrs_ref_allele_seq" attribute that allows for the reference allele for a VRS object to be reported along with the modified allele. To ensure that all VRS objects in all processed score sets contained this attribute, a VRS object was instantiated for each pre-mapped and post-mapped variant, where the existing VA digest was supplied to the "id" attribute and the VRS allele was supplied to the "variation" attribute. SeqRepo was utilized to determine the reference allele sequence for the VariationDescriptor objects and the sequence was accordingly provided to the "vrs_ref_allele_seq" attribute. Additionally, for post-mapped variants, an HGVS string (e.g. "NC_000006.12:g.37808023C>A") describing the VRS allele was generated and supplied to the "expressions" attribute in the VariationDescriptor class.

The completion of this step validated the accuracy of the mapping procedure as it provided a check on the consistency between the reference alleles generated for the pre-mapped and post-mapped representations. For protein variants and genomic variants derived from positively oriented target sequences, we observed concordance between substrings provided by the "vrs_ref_allele_seq" attribute for the pre-mapped and post-mapped objects. However, as SeqRepo contains reference sequences derived from positive strands, the reference allele sequence was the reverse complement for genomic variants that came from negatively oriented target sequences (e.g. pre-mapped: "GAT", post-mapped: "ATC").

Mapping File Format

The newly created VariationDescriptor objects were added to a "mapped_scores" attribute for each score set. Additionally, a "computed_reference_sequence" attribute, storing the target sequence, sequence type, and SQ digest, and a "mapped_reference_sequence" attribute, storing the RefSeq accession, sequence type, and corresponding SQ digest, were added to each score set. Following the creation of the respective attributes, all processed score sets were saved as JSON files, gzipped, and uploaded to a publicly-accessible s3 bucket (mavedb-mapping).

Integrating MaveDB Data into the UCSC Genome Browser

Due to the large amount of protein-level variation data reported by MAVEs, additional tools needed to be employed to translate protein changes to codons. This was accomplished with the VICC protein to genome mapping method from cool seq tool(29) at https://normalize.cancervariants.org/variation/alignment_mapper/p_to_g to map the protein-level variants to the genomic coordinates associated with their corresponding codons.

Integrating the MaveDB Data into the ClinGen LDH

To integrate the MaveDB data into the ClinGen LDH, a MaveDbMapping document was created for each score set entry in the mapping files and added to the LDH as linked data for an LDH variant represented by a ClinGen canonical allele identifier. Because the ClinGen Allele Registry requires the use of standard human reference sequences (genome builds NCBI36, GRCh37, GRCh38 and transcripts from NCBI or ENSEMBL), each HGVS expression within the VariationDescriptor object from these score set entries was leveraged to either find the existing canonical allele identifier referenced in the score set entry or to register the variant with the ClinGen Allele Registry to obtain a new canonical allele identifier. MaveDbMapping documents were created by excerpting the MaveDB mapped scores object, score, MaveDB score set id (URN + entry number; eg. urn:mavedb:00000001-a-1#1*)*, captured provenance information (creation, modification and publish dates), and a link back to the referenced MaveDB score set page.

Declarations

Availability of data and materials

The code and datasets generated during this study are publicly available in the dcd_mapping repository at: **https://github.com/ave-dcd/dcd_mapping**.

Competing interests
The authors declare that they have no competing interests.

Funding
JAA and AHW were supported by award R35HG011949 from the National Human Genome Research Institute (NHGRI) of the National Institutes of Health (NIH). EYD was supported by NIH/NHGRI UM1HG011969. TF and KR received funding from NIH NHGRI Clinical Genome Resource (ClinGen) grant U24 HG009649. JFS was supported by the Merkin Institute of Transformative Technologies in Healthcare grant. AM receives funding from NIH NHGRI Clinical Genome Resource (ClinGen) grant U24 HG009649. HVF is supported by Wellcome Biomedical Resources Grant WT223718/Z/21/Z. SI was supported by the Merkin Institute of Transformative Technologies in Healthcare grant. MSC is supported by NCI grant U01CA242954. AFR was supported by NIH/NHGRI UM1HG011969 and RM1HG010461. This work was supported by the Australian government.

Authors' Contributions
All authors contributed to the design and execution of this work, and assisted in drafting and/or revising the manuscript.

Acknowledgements

We would like to thank William Ray and the Graphics Services team at Nationwide Children's Hospital for their assistance in developing the figures for this manuscript.

Figures

Fig 1. *Mapping MAVE variants to the Human Genome*

An overview of the MAVE variant mapping method. MAVE variants are described with respect to supplied target sequences, but the absence of linkages to human reference sequences limits the interoperability of MAVE data with human genomics applications (left). To overcome this limitation, we have developed a method to map MAVE variants to their corresponding human reference sequences (middle). Through the use of VRS, we are able to represent MAVE variants with respect to both assayed target sequences and versioned human reference sequences, creating robust homology maps (middle). The precise representation of MAVE variants using VRS ultimately facilitates the integration of MAVE data into downstream clinical and research applications (right).

Fig 2. *MaveDB Score Set Breakdown/Summary Statistics*

A summary of the MAVE data from MaveDB that was used for validation of the mapping method. **(A)** All score set entries in MaveDB are assigned an organism attribute (e.g. *Homo sapiens*, *Saccharomyces cerevisiae*). Score sets whose listed target organism was *Homo sapiens* (n = 209) were selected for testing of the mapping algorithm, and additional breakdowns describing the selected human score sets are presented. Made with SankeyMATIC. **(B)** MAVE experiments in MaveDB (n = 159) can be conducted in non-human cellular contexts, including yeast, bacteria, mice, and bacteriophage(40), (41). Experiments that do not report a cellular context are coded as "N/A" (n = 5).

Fig 3. *Variant Mapping Algorithm Workflow*

A depiction of the MAVE variant mapping workflow. For a given entry in MaveDB whose listed target organism is *Homo sapiens*, the provided target sequence is aligned to GRCh38 using BLAT(20), returning data including the chromosome number, gene symbol, and a set of genomic coordinates (**1**). If a score set describes a protein coding element, the outputted data can be supplied as a query to the Universal Transcript Archive (UTA) database, ultimately allowing for a RefSeq protein accession to be derived and for an offset to be computed (**2**). With a RefSeq sequence selected and offset calculated, the assayed variants in a MaveDB variant matrix are described with respect to their unique human reference sequence using the GA4GH Variation Representation Specification (VRS) (**3**). The resulting VRS objects are then annotated with descriptive metadata and integrated into specific score set JSON files. Lastly, the JSON files are gzipped and uploaded to the publicly-accessible mavedb-mapping s3 bucket for downstream integration. Created with BioRender.com.

Fig 4. *Downstream Integrations of MaveDB Data*

The mapping of MAVE data to the human genome permits downstream data integrations in various human genomics applications. **(A)** MAVE scores are visible as heatmaps for available genes in the Genomics 2 Proteins Portal. **(B)** MAVE data has been added as a track hub in the UCSC Genome Browser. MAVE protein variant positions are mapped to their corresponding genomic coordinates, and the score, chromosome band, genomic size, and strand are also reported for each variant. **(C)** MAVE scores and a link to the associated score set are reported, when available, for queried variants in the Ensembl Variant Effect Predictor.

Supplementary Figure 1. *Overview of MAVE data and visualization in the Genomics 2 Proteins portal.*

(a) Count of point mutations and pairwise mutations covered in MAVEs for 40 and 7 human genes, respectively. (b) Coverage of residues in MAVEs for each protein compared to the length of the sequence. (c) - (d) Heatmap visualization of MAVEs for gene and score set in the portal. (e) Annotation of protein sequence with mutations with top and bottom 99th percentile of MAVEs per score set in the context of protein physicochemical, structural and functional features.

Bibliography

1. Henrie A, Hemphill SE, Ruiz-Schultz N, Cushman B, DiStefano MT, Azzariti D, et al. ClinVar Miner: Demonstrating Utility of a Web-Based Tool for Viewing and Filtering ClinVar Data. Hum Mutat. 2018 Aug;39(8):1051.

2. Landrum MJ, Lee JM, Riley GR, Jang W, Rubinstein WS, Church DM, et al. ClinVar: public archive of relationships among sequence variation and human phenotype. Nucleic Acids Res. 2014 Jan 1;42(Database issue):D980.

3. Brnich SE, Abou Tayoun AN, Couch FJ, Cutting GR, Greenblatt MS, Heinen CD, et al. Recommendations for application of the functional evidence PS3/BS3 criterion using the ACMG/AMP sequence variant interpretation framework. Genome Med. 2019 Dec 31;12(1):1–12.

4. Tareen A, Kooshkbaghi M, Posfai A, Ireland WT, McCandlish DM, Kinney JB. MAVE-NN: learning genotype-phenotype maps from multiplex assays of variant effect. Genome Biol. 2022 Apr 15;23(1):98.

5. Fowler DM, Stephany JJ, Fields S. Measuring the activity of protein variants on a large scale using deep mutational scanning. Nat Protoc. 2014 Sep;9(9):2267–84.

6. Starita LM, Fields S. Deep Mutational Scanning: Library Construction, Functional Selection, and High-Throughput Sequencing. Cold Spring Harb Protoc. 2015 Aug 3;2015(8):777–80.

7. Klein JC, Agarwal V, Inoue F, Keith A, Martin B, Kircher M, et al. A systematic evaluation of the design and context dependencies of massively parallel reporter assays. Nat Methods. 2020 Nov;17(11):1083–91.

8. Inoue F, Ahituv N. Decoding enhancers using massively parallel reporter assays. Genomics. 2015 Sep;106(3):159–64.

9. Nielsen SV, Hartmann-Petersen R, Stein A, Lindorff-Larsen K. Multiplexed assays reveal effects of missense variants in MSH2 and cancer predisposition. PLoS Genet. 2021 Apr;17(4):e1009496.

10. Fayer S, Horton C, Dines JN, Rubin AF, Richardson ME, McGoldrick K, et al. Closing the gap: Systematic integration of multiplexed functional data resolves variants of uncertain significance in BRCA1, TP53, and PTEN [Internet]. Vol. 108, The American Journal of Human Genetics. 2021. p. 2248–58. Available from: http://dx.doi.org/10.1016/j.ajhg.2021.11.001

11. Esposito D, Weile J, Shendure J, Starita LM, Papenfuss AT, Roth FP, et al. MaveDB: an open-source platform to distribute and interpret data from multiplexed assays of variant effect. Genome Biol. 2019 Nov 4;20(1):223.

12. Wagner AH, Babb L, Alterovitz G, Baudis M, Brush M, Cameron DL, et al. The GA4GH Variation Representation Specification: A computational framework for variation representation and federated identification. Cell Genom [Internet]. 2021 Nov 10;1(2). Available from: http://dx.doi.org/10.1016/j.xgen.2021.100027

13. Eilbeck K, Lewis SE, Mungall CJ, Yandell M, Stein L, Durbin R, et al. The Sequence Ontology: a tool for the unification of genome annotations. Genome Biol. 2005 Apr 29;6(5):R44.

14. Rubin AF. mavehgvs [Internet]. Zenodo; 2023. Available from: https://zenodo.org/record/3405637

15. den Dunnen JT, Dalgleish R, Maglott DR, Hart RK, Greenblatt MS, McGowan-Jordan J, et al. HGVS Recommendations for the Description of Sequence Variants: 2016 Update. Hum Mutat. 2016 Jun;37(6):564–9.

16. Frankish A, Diekhans M, Ferreira AM, Johnson R, Jungreis I, Loveland J, et al. GENCODE reference annotation for the human and mouse genomes. Nucleic Acids Res. 2018 Oct 24;47(D1):D766–73.

17. O'Leary NA, Wright MW, Brister JR, Ciufo S, Haddad D, McVeigh R, et al. Reference sequence (RefSeq) database at NCBI: current status, taxonomic expansion, and functional annotation. Nucleic Acids Res. 2015 Nov 8;44(D1):D733–45.

18. Hart RK, Prlić A. SeqRepo: A system for managing local collections of biological sequences. PLoS One. 2020 Dec 3;15(12):e0239883.

19. Hart R. UTA: The Universal Transcript Archive [Internet]. Zenodo; 2013. Available from: https://zenodo.org/record/6975034

20. Kent WJ. BLAT--the BLAST-like alignment tool. Genome Res. 2002 Apr;12(4):656–64.

21. Wilkinson MD, Dumontier M, Aalbersberg IJJ, Appleton G, Axton M, Baak A, et al. The FAIR Guiding Principles for scientific data management and stewardship. Sci Data. 2016 Mar 15;3:160018.

22. Kent WJ. The Human Genome Browser at UCSC [Internet]. Vol. 12, Genome Research. 2002. p. 996–1006. Available from: http://dx.doi.org/10.1101/gr.229102.

23. McLaren W, Gil L, Hunt SE, Riat HS, Ritchie GRS, Thormann A, et al. The Ensembl Variant Effect Predictor. Genome Biol. 2016 Jun 6;17(1):1–14.

24. Iqbal S, Pérez-Palma E, Jespersen JB, May P, Hoksza D, Heyne HO, et al. Comprehensive characterization of amino acid positions in protein structures reveals molecular effect of missense variants. Proc Natl Acad Sci U S A. 2020 Nov 10;117(45):28201–11.

25. Dalton KP, Rehm HL, Wright MW, Mandell ME, Krysiak K, Babb L, et al. Accessing clinical-grade genomic classification data through the ClinGen Data Platform. Pac Symp Biocomput. 2023;28:531–5.

26. Foreman J, Brent S, Perrett D, Bevan AP, Hunt SE, Cunningham F, et al. DECIPHER: Supporting the interpretation and sharing of rare disease phenotype-linked variant data to advance diagnosis and research. Hum Mutat. 2022 Jun 1;43(6):682–97.

27. Tudini E, Andrews J, Lawrence DM, King-Smith SL, Baker N, Baxter L, et al. Shariant platform: Enabling evidence sharing across Australian clinical genetic-testing laboratories to support variant interpretation. Am J Hum Genet [Internet]. 2022 Nov 3 [cited 2023 May 6];109(11). Available from: https://pubmed.ncbi.nlm.nih.gov/36332611/

28. UniProt Consortium. UniProt: the universal protein knowledgebase in 2021. Nucleic Acids Res. 2021 Jan 8;49(D1):D480–9.

29. Kuzma K, Stevenson J, Wagner A. Cool Seq Tool [Internet]. 2023. Available from: https://zenodo.org/record/7884144

30. DECIPHER: Database of Chromosomal Imbalance and Phenotype in Humans Using Ensembl Resources. Am J Hum Genet. 2009 Apr 10;84(4):524–33.

31. Fortuno C, Lee K, Olivier M, Pesaran T, Mai PL, de Andrade KC, et al. Specifications of the ACMG/AMP variant interpretation guidelines for germline TP53 variants. Hum Mutat [Internet]. 2021 Mar [cited 2023 May 15];42(3). Available from: https://pubmed.ncbi.nlm.nih.gov/33300245/

32. Berners-Lee T. Linked Data [Internet]. The World Wide Web Consortium (W3C) - Design Issues. 2006 [cited 2023 May 12]. Available from: https://www.w3.org/DesignIssues/LinkedData.html

33. Pawliczek P, Patel RY, Ashmore LR, Jackson AR, Bizon C, Nelson T, et al. ClinGen Allele Registry links

information about genetic variants. Hum Mutat. 2018 Nov;39(11):1690–701.

34. Hart RK, Rico R, Hare E, Garcia J, Westbrook J, Fusaro VA. A Python package for parsing, validating, mapping and formatting sequence variants using HGVS nomenclature. Bioinformatics. 2015 Jan 15;31(2):268–70.

35. Kuzma K, Stevenson J, Wagner A. VICC Gene Normalization Service. 2022 May 11 [cited 2023 Jun 20]; Available from: https://zenodo.org/record/7508058

36. Yates B, Braschi B, Gray KA, Seal RL, Tweedie S, Bruford EA. Genenames.org: the HGNC and VGNC resources in 2017. Nucleic Acids Res. 2017 Jan 1;45(Database issue):D619.

37. Cunningham F, Allen JE, Allen J, Alvarez-Jarreta J, Amode MR, Armean IM, et al. Ensembl 2022. Nucleic Acids Res [Internet]. 2022 Jan 7 [cited 2022 Nov 21];50(D1). Available from: https://pubmed.ncbi.nlm.nih.gov/34791404/

38. Maglott D, Ostell J, Pruitt KD, Tatusova T. Entrez Gene: gene-centered information at NCBI. Nucleic Acids Res. 2011 Jan;39(Database issue):D52.

39. Bradley Holmes J, Moyer E, Phan L, Maglott D, Kattman B. SPDI: data model for variants and applications at NCBI. Bioinformatics. 2020 Mar 3;36(6):1902.

40. The R Project for Statistical Computing [Internet]. [cited 2023 Jun 5]. Available from: https://www.R-project.org/.

41. A Forward-Pipe Operator for R [R package magrittr version 2.0.3]. 2022 Mar 30 [cited 2023 Jun 5]; Available from: https://CRAN.R-project.org/package=magrittr

MAVE Library Generation

Assayed Variant Score Set

MAVE Target Sequence Representation

Mapped Variant Score Set

Human Reference Sequence Representation

Human Clinical and Research Applications

Model Organisms

A ⊂ B
A ⊆ B
A ≈ B

Sequence Ontology

A ≈ B

JSON

A   B

1

Align Target Sequence to Human Genome

Target Type is Protein Coding?

No

3

Yes

Yes

Target Organism is Human?

Standardize MaveDB Variants

2 Select Reference Sequence, Compute Offset

Tr. A

Tr. B

Tr. C

Block 1   Block 2   Block 3

A,B   B   B,C

intersection([A,B], [B], [B,C]) = **B**

Target Sequence   **Offset= 10**

NP Sequence

MaveDB Associated Metadata/Variant Score Matrix

MaveDB Data Available for Downstream Integration

C

Variant Effect Predictor results ⊞

Job details ⊞

Summary statistics ⊞

Results preview