Sample-specific haplotype-resolved protein isoform characterization via long-read RNA-seq-based proteogenomics

David Wissel[1,2,3], Gloria M. Sheynkman[4,5,6*†], Mark D. Robinson[1,3*†]

[1]Department of Molecular Life Sciences, University of Zurich, Zurich, Switzerland.
[2]Department of Computer Science, ETH Zurich, Zurich, Switzerland.
[3]SIB Swiss Institute of Bioinformatics, University of Zurich, Zurich, Switzerland.
[4]Department of Molecular Physiology and Biological Physics, University of Virginia, Charlottesville, VA, USA.
[5]Department of Biochemistry and Molecular Genetics, University of Virginia, Charlottesville, VA, USA.
[6]UVA Cancer Center, University of Virginia, Charlottesville, VA, USA.

*Corresponding author(s). E-mail(s): gs9yr@virginia.edu; mark.robinson@mls.uzh.ch;
†Equal supervision.

Abstract

Protein isoform inference from bottom-up mass spectrometry (MS) relies on database search strategies that assume the reference protein database accurately reflects the full repertoire of genetic and transcriptomic states present in the sample being analyzed. Long-read RNA sequencing (lrRNA-seq) now enables simultaneous recovery of complete transcript structures and the genetic variants present on each molecule, offering a direct route to allele-specific isoforms, yet this capability has not been fully leveraged to improve MS-based proteogenomics workflows. Here, we develop an end-to-end workflow for constructing and searching haplotype-resolved, sample-specific proteomes using matched lrRNA-seq and MS data. We benchmark multiple phasing algorithms on PacBio lrRNA-seq from Genome-in-a-Bottle samples and identify methods that achieve high phasing accuracy and completeness on transcriptomic reads. Our open-source, modular Snakemake pipeline performs variant calling, read-based phasing, isoform discovery, haplotype-resolved proteome construction, MS search, and downstream annotation. To demonstrate its utility, we apply the workflow to an induced pluripotent stem cell line (WTC11) and to an osteoblast differentiation time course, showing that haplotype-resolved databases enable detection of variant and splice peptides, allele-specific protein isoforms, and linked variants not detectable with reference-only proteomes. Together, our results demonstrate that lrRNA-seq-based phasing is feasible and effective for proteogenomics and provide a practical framework for allele-resolved proteome characterization in dynamic or disease-relevant settings.

Keywords: Long-read RNA-seq, Proteogenomics, Variant phasing

Background

The accurate identification and characterization of proteoforms has been a core focus of the proteomics and proteogenomics field [1]. Although top-down mass spectrometry (MS) methods have become increasingly popular for direct detection of proteoforms [2], the global characterization of protein isoforms primarily relies on bottom-up, or shotgun sequencing, regimes in which the existence of proteoforms can only be inferred from identified peptide fragments. In this schema, the short peptide fragments are mapped back to the annotated sequences in a reference proteome and the protein sequences with highest support, based on the inference algorithm, is reported. Such inference methods critically depend on the accuracy of the reference proteome

relative to the actual composition of the proteome in the sample subjected to analysis. However, reference proteomes are woefully incomplete; the number of proteoforms in the human proteome is estimated to be in the hundreds of thousands, if not millions, owing primarily to genetic variants, alternative splicing, and post-translational modifications [3]. Consequently, accurate protein inference requires protein annotations that faithfully capture genetic and transcriptomic complexity in a cell or tissue specimen [4].

In human cells, the protein inference challenge is compounded by the fact that primary protein sequences (N to C-terminal AA chain) contain different coding (e.g., missense) genetic variants across individuals. Variant calling attempts to identify single single nucleotide variants (SNVs) and insertions and deletions (indels) events that differ from a reference genome, within a particular sample.[1]

Genetic variation introduces not only single event amino-acid changes but also insertions and deletions (indels) of polypeptide segments as well as combinations of variants that co-occur and co-inherited on the same chromosome, generating paternal and maternal versions of a protein that may differ in their haplotype, or "phasing". Variant phasing methods aim to reconstruct these haplotypes, i.e., groups of variants that are jointly inherited.[2] However, these protein haplotypes are almost never represented in standard proteome databases, which typically assume a single reference sequence per isoform. Additionally, notwithstanding the genetic variation, transcriptomic complexity adds an additional layer of complexity as haplotype can produce multiple alternatively spliced transcripts. Genetic variants and splice choices jointly determine the final amino-acid sequence.

The true proteoform space for any given patient arises from the interplay of diploid genetic variation and sample-specific alternative splicing. The genome provides the primary sequence context from which all transcripts are transcribed, and the arrangement of variants along each chromosome determines how these changes are presented in downstream transcripts. Each allele can give rise to multiple splice isoforms, meaning that the same set of variants may be incorporated, excluded, or combined in distinct ways across known and novel transcripts. An accurate protein reference for MS-based inference therefore requires both (i) knowledge of the genetic variants present on each allele and (ii) the correct set of transcript structures expressed in that sample, so that the resulting protein models faithfully represent the allele-resolved sequences that cells actually produce. These considerations motivate the need for a systematic approach that identifies genetic variants, determines their phase (their co-occurrence on the same allele), and establishes the full length transcripts within which the genetic variants are expressed.

Phasing may be performed using several strategies, including population-based phasing (leveraging large reference panels), trio-based phasing (using parental genotypes), and read-based phasing (inferring haplotypes directly from sequencing data) Browning and Browning [5]. Of these, read-based phasing enables truly sample-specific reconstruction of allele structure because it uses the long or linked reads obtained from the sample itself to determine which variants co-occur.

In what follows, we give a short overview of recent proteogenomics approaches that enable the creation of sample-specific proteomes, in addition to variant phasing. Related work has developed along two largely independent axes: sample-specific proteogenomics workflows and protein haplotype prediction at population scales. We begin with proteogenomic frameworks that refine the MS search space by constructing customized proteomes for a given sample.

[1] We will use "small variant calling" and "variant calling" interchangeably hereafter, since we do not capture structural variants in this study.

[2] For simplicity, we restrict ourselves to the diploid setting, since we are primarily interested in developing this method as it applies to mammalian research in humans or more broadly mammals, and we disregard non-diploid disease settings, such as (commonly) cancer, for the moment.

Cesnik et al. [6] represents a consolidated pipeline that integrates short-read RNA-seq with variant calling to produce protein databases containing both variant and splice peptides. They used GATK [7] for small-variant discovery and StringTie2 [8] for novel transcripts reconstruction, and SnpEff [9] was used to project the variants onto protein sequences to generate proteins. Spritz, enables the identification of both variant peptides and, in some cases, proteoforms by utilizing its sample-specific database in bottom-up and top-down MS searches, respectively. Spritz processes short read RNA-seq data, so are limited to incomplete transcript reconstructions that do not capture the full length isoforms. In this regard, Miller et al. [4] introduced a pipeline for long-read RNA-seq (lrRNA-seq) data that constructs sample-specific, full-length proteomes. The authors used the Iso-Seq workflow [10] to assemble a sample-specific transcriptome, annotated open reading frame (ORF)s using CPAT [11], and construct the corresponding personalized proteome. Miller et al. [4] show such databases can enhance characterization of specific protein isoforms, with detailed isoform classification automatically performed based on an extension of SQANTI3, SQANTIProtein [12]. However, though capturing splicing complexity at full-length resolution, the pipeline omits genetic variation. More recently, Salz et al. [13] introduce a proteogenomic pipeline that indeed includes genetic variants in full-length isoform models. They report on the construction a personalized proteome using the HG001 Genome in a Bottle (GIAB) cell line [14]. Phased genetic variants were obtained from a previously published variant phasing dataset generated using an ancestry-based phasing approach [15] and novel transcripts structures obtained from an analysis of Oxford Nanopore data on the same cell line [16]. They then overlaid or "transferred" the phased variants onto the predicted proteins. However, their approach applies genomic variants uniformly across all transcript isoforms, without determining which isoforms are actually expressed or whether a given variant resides on the same allele as a particular transcript.

Beyond these proteogenomic efforts, a separate line of work has focused not on sample-specific database construction, but on defining the theoretical space of protein haplotypes arising from genetic variation across human populations.

Haplosaurus utilizes the popular tool Variant Effect Predictor (VEP) [17] to map all variants from the 1000 Genomes dataset [18] onto translated coding sequence (CDS)s (i.e., reference proteins) Spooner et al. [19]. This framework defines all possible protein sequences encoded by phased alleles and highlights how common haplotypes can differ substantially from canonical reference proteins across populations and lead to changes in amino acid that can affect protein functions as relevant to applications such as drug design. In a following study by another group Vašíček et al. [20], Haplosaurus was used to generate databases for MS searching. They generated haplotype protein FASTAs and performed MS database searches to assess the downstream detectability of peptide-level variants, identifying 176k peptides at 1% FDR, including 7.5K peptides containing one variant and 352 with multiple variants. This work is an early demonstration that multi-variant peptides from protein haplotypes can be observed in a systematic manner. The same authors built on this work by introducing a scalable bioinformatic pipeline, ProHap [21] and companion visualization tool [22], which constructs personalized proteomes by overlaying phased variants from population reference panels onto all Ensembl transcript annotations. The authors illustrate the usage of their method on the 1000 Genomes dataset, revealing cross-population proteome complexity. Collectively, all these studies establish the theoretical landscape of protein haplotypes and demonstrate their detectability in MS data. However, because they rely exclusively on reference transcript annotations, they cannot represent sample-specific or novel isoforms and cannot jointly model the allele-specific combinations of variants and splice choices that give rise to the proteoforms present in an actual sample.

Based on the advent of highly accurate long-read RNA sequencing, and, more recently, the increasing accessibility of highly accurate long-read RNA-seq datasets (e.g., PacBio Iso-Seq, ONT cDNA), we recognized an opportunity to build more sample-specific, personalized proteome. In long read RNA sequencing, each read, representing a transcript, simultaneously encodes information on the complete splicing structure of the expressed isoform and the set of genetic variants present along that same molecule. In principle, long-read RNA-seq should therefore allow direct recovery of both isoform identity and variant phase on the very same reads, which can be translated into protein sequences that contain allele-specific isoforms. Along with developments in long read sequencing, methods for phasing genetic variants have matured [23, 24] [25][26], and it is now possible to employ phasing on long read transcript data [27, 28], including methods that leverage allele specific expression (ASE) [29].

In this work, we address this gap by developing an end-to-end framework for constructing haplotype-resolved, sample-specific proteomes directly from long-read RNA sequencing. We briefly summarize the main contributions. First, in order to determine the ideal method for variant calling, we perform a benchmark of methods for variant phasing on PacBio lrRNA-seq data, focusing on raw phasing performance given the lack of diverse large benchmark datasets with matched MS and lrRNA-seq. Second, we implement a Snakemake workflow that enables the creation and search of haplotype-resolved, sample-specific proteomes from sample-matched lrRNA-seq and MS data [30]. Third, we describe the proteome complexity observed in an lrRNA-seq sample from the WTC11 cell line as well as in a different biological model of differentiation from iPSCs to osteoblasts, showing potential applications of our workflow in disease or other more complicated experimental designs. Finally, we highlight new aspects of how haplotype-resolved proteomes relate to protein isoform characterization. We report not only the prevalence of sample-specific haplotypic proteins, isoforms, and identified peptides, but also demonstrate how exact knowledge of all protein isoform products at the joint allele-isoform level has implications for accurate protein inference.

The main novel aspects of our work can be best highlighted in this way. Our workflow is designed to be end-to-end application. Previous work focused on the one-time usage of a particular sample [13] or on a subset of the scope of our pipeline, such as the creation of protein haplotypes only [19] without enabling, for example, sample-specific variant calling and phasing. Next, our workflow accounts for a broader swath of variant types, including phased SNVs, indels, and alternative splicing. To the best of our knowledge, previous work frequently disregarded the role of at least one of these possible determinants (indels, heterozygous variants, or alternative splicing). Third, we perform a benchmark of read-based phasing methods focused on PacBio LRRnA. As far as we know, previous work either focused on whole-genome or whole-exome data [23, 25, 26, 31], short-read RNA-seq [27, 29], or compared very few methods [28].

Results

Read-based phasing methods exhibit good performance on PacBio long-read RNA-seq data

We evaluated the performance of existing methods for variant phasing using PacBio lrRNA-seq on two GIAB datasets (HG002, HG005; see Methods). In particular, we wanted to know which phasing methods performed well in terms of phasing the 'ground-truth' GIAB variants, specifically using PacBio lrRNA-seq. For this purpose, we evaluated phasing accuracy via the switch error rate, as well as phasing completeness, using both the total number of phased variants and the proportion of fully phased CDSs (i.e., CDSs that consisted of a single contiguous phase block) (see Methods). We note that, when using lrRNA-seq data, the length of

phase blocks is limited by the length of a gene. Despite this, since our primary downstream application is MS search, methods that phase a large number of variants but do not fully phase a large proportion of CDSs are not very useful.

Overall, we found similar performance on both the HG002 and HG005 datasets. WhatsHap, HapCUT2, and Margin all performed at low switch error rates and performed broadly similar to each other, with WhatsHap achieving a slightly worse switch error rate compared to the other two, while phasing notably more variants overall, as well as phasing a larger percentage of CDSs fully (Fig. 1A). phASER did not perform well, underperforming the top three methods based on phasing accuracy, with similar performance to WhatsHap based on phasing completeness. HiPhase performed worse than all other methods regarding phasing accuracy on HG002 and HG005, but never fully phased less than 80% of all CDSs. There were no notable differences in results when including or excluding indels for HG002 and HG005.

Based on these results, for variant calling and phasing, in particular with the goal of performing MS search with PacBio, we recommend the usage of WhatsHap, Margin or HapCUT2 for PacBio data and use WhatsHap throughout our manuscript.

A modular workflow enables sample-specific haplotype-resolved proteogenomics

With the goal of easing maintainability and usability, we constructed a Snakemake pipeline that maximizes modularity. Importantly, we leverage various existing methods, augmenting them with custom scripts and tools only when necessary.

Users provide a reference genome and transcriptome, and sample-matched lrRNA-seq and MS data (Fig. 1B). Our pipeline begins by performing genome alignment on the lrRNA-seq data (see Methods). Afterward, our workflow performs variant calling (see Methods) and transcript discovery (see Methods) in parallel, accounting for genetic and alternative splicing variation, respectively. Variants are filtered for coverage and call quality, phased using WhatsHap, and post-processed to ensure that regions of interest (typically CDSs) consist of a contiguous phase block. Novel transcripts undergo an additional ORF calling step using ORFanage, followed by a post-processing step to ensure sufficient length of CDSs (see Methods).

Afterward, taking in both novel CDSs and phased variants, we construct protein haplotypes using haplosaurus, followed by a post-processing step and the generation of decoys for MS search (Fig. 1B). Specifically, first, we generate GENCODE reference protein haplotypes using gffread. Then, for our sample-specific proteome, we replace any protein isoform that has non-silent variants with the corresponding haplosaurus predicted protein haplotype. Lastly, we add the protein haplotypes corresponding to novel protein isoforms (whether they contain genetic variation or not).

Lastly, we perform a database search (see Methods), based on the sample-specific database and the sample-matched user-provided MS input data (Fig. 1B). Following the search, our pipeline performs post-search annotation of all protein isoforms, creates files to enable the creation of browser tracks, and performs protein inference.

Sample-specific haplotype-resolved proteogenomics enables the identification of a broad class of protein isoforms

Contrary to competing approaches, which either rely on population-based variant data and thus do not necessarily reflect sample-specific variation or do not jointly incorporate splice with phased genetic variation, our approach allows for the incorporation of a broad class of proteoform candidates. Specifically, our pipeline incorporates reference protein isoforms (which are identical to a reference database, such as GENCODE),

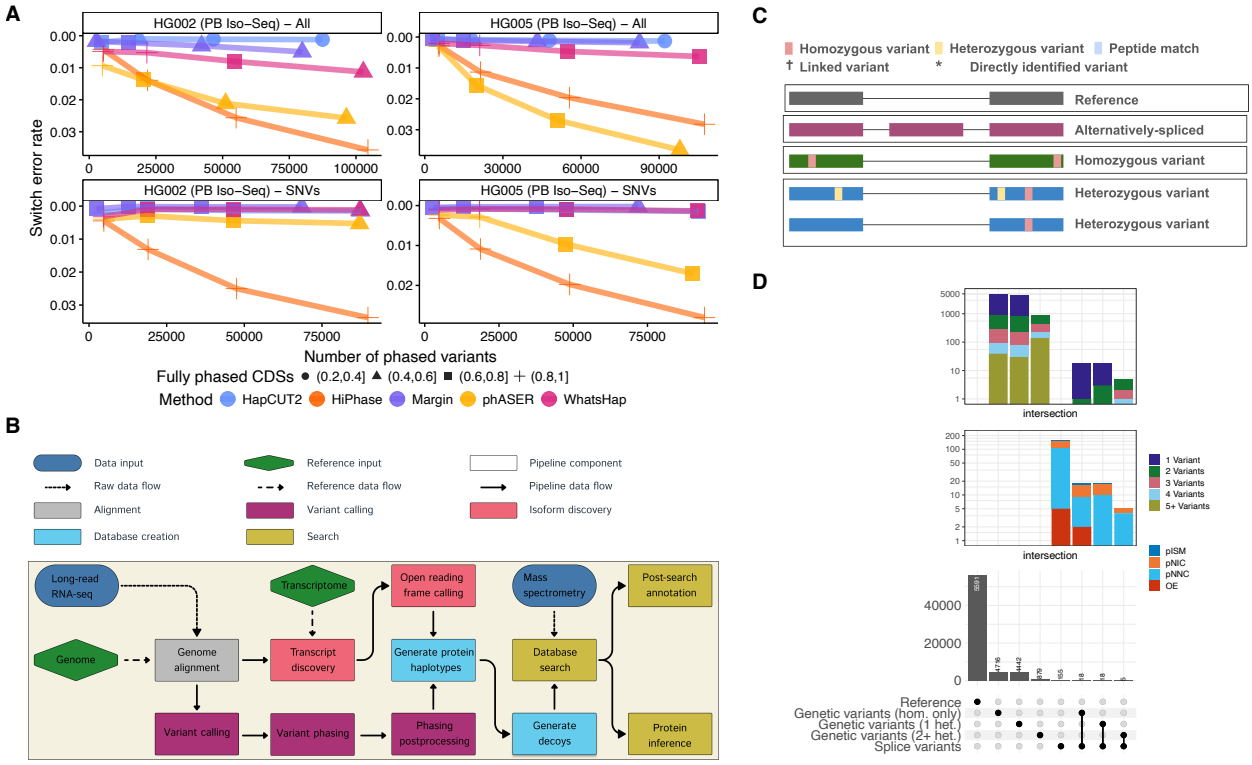

Fig. 1 Generating sample-specific proteomes using our workflow enables a comprehensive characterization of the theoretical proteome of a system, in our case, the WTC11 cell line. **A.** Phasing performance as measured by switch error rate of established phasing methods on lrRNA-seq data by the number of phased variants and fully phased CDSs, stratified by dataset and inclusion of indels. **B.** Schematic view of our pipeline to perform sample-specific haplotype-resolved proteogenomics. **C.** Overview of categories of protein isoforms identifiable with our workflow. Legend also used in Fig. 4. **D.** Number of protein isoforms belonging to each variant category and the reference, along with a breakdown of complexity within each variant category.

alternatively spliced protein isoforms (which are created using transcript discovery based on the inputted lrRNA-seq) (Fig. 1C). In addition, our pipeline enables sample-specific search of databases, including both protein isoforms that contain only homozygous variants (which rely only on variant calling and do not require phasing) and ones that contain heterozygous variants (which rely on variant phasing in order to prevent an enumeration of possible variant combinations) (Fig. 1C).

The proteome makeup of the WTC11 cell line is primarily driven by less complex events

We applied our pipeline to a WTC11 cell line sample (see Methods). Below and throughout the remainder of the text, we only consider non-silent genetic variation, due to the centering of our work on the proteome.

The majority of protein isoforms in our WTC11 database were identical to isoforms found in the GENCODE reference database from which our database was derived (55,911, 84.5%) (Fig. 1D). Further complexity in our sample was primarily driven by genetic variants, with only 196 (0.3 %) protein isoforms exhibiting alternative splicing (Fig. 1D), relative to the reference GENCODE proteome, while 10,078 (15.2%) protein isoforms presented one or more non-silent variations (Fig. 1D). In addition, in our sample, the majority of variant protein isoforms had either no heterozygous variants (4,734 or 47.0% of variant protein isoforms) or one heterozygous variant (4,460 or 44.3% of variant protein isoforms), with only a few presenting two or more heterozygous variants (884 or 8.8% of variant protein isoforms) (Fig. 1D).

There was only very moderate overlap between reference variation due to alternative splicing and genetic variation. In particular, a total of 41 protein isoforms (0.1%) exhibited both alternative splicing and any form of genetic variation (Fig. 1D).

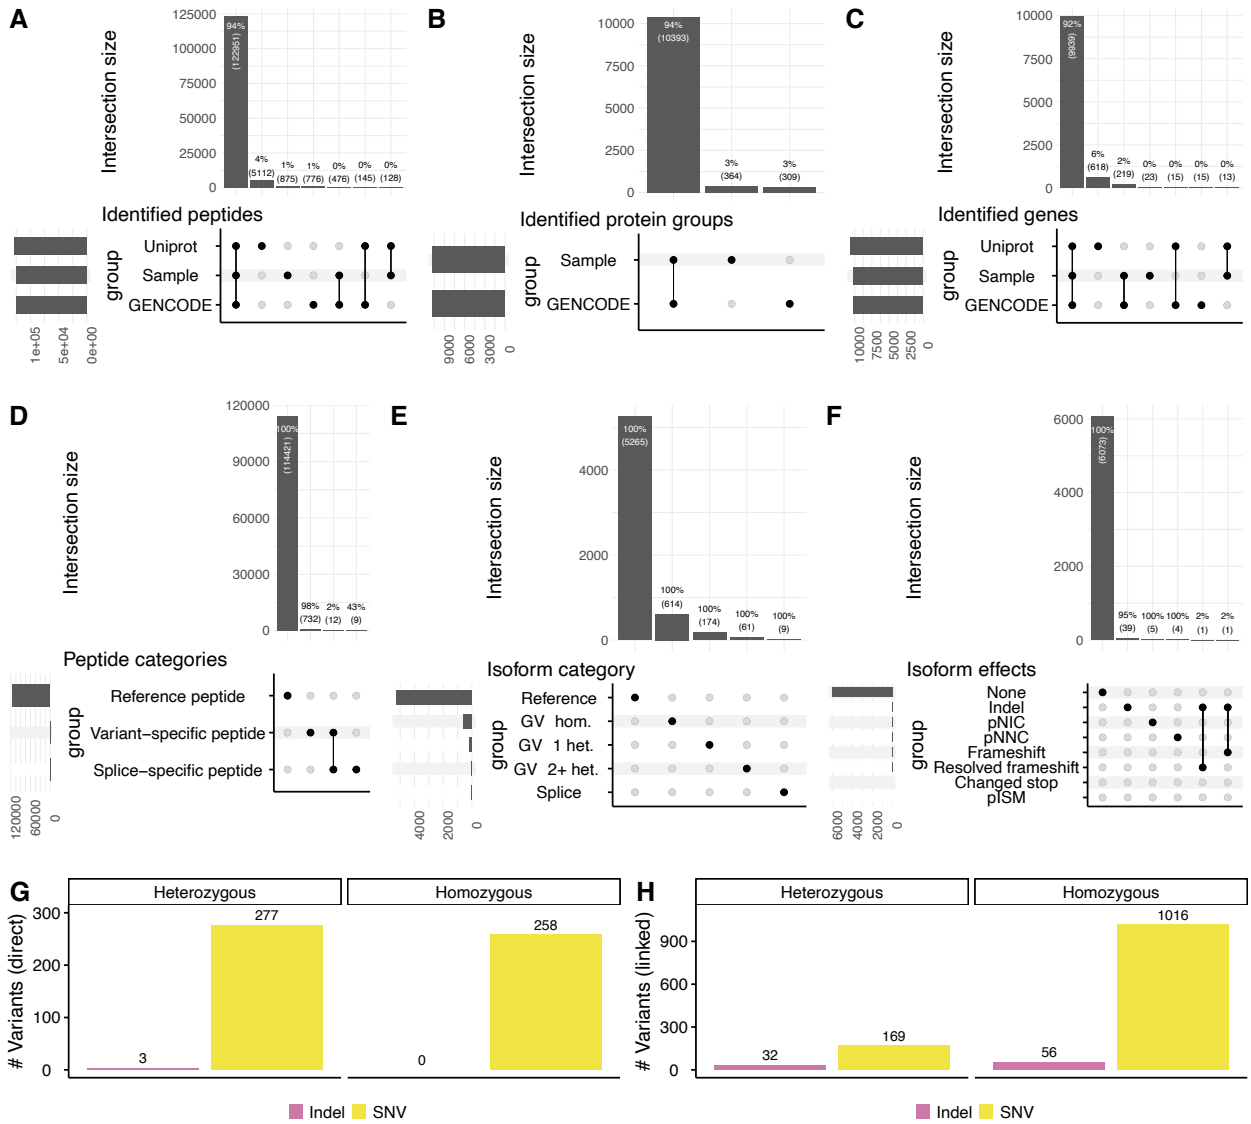

Fig. 2 Searching MS data against the sample-specific proteome of the WTC11 cell line reveals good concordance with reference proteomes, low linkage between protein isoform types and protein effects, and identification of both direct and linked variants. **A.** UpSet plot of peptide overlaps between sample-specific proteome, UniProt, and GENCODE reference. **B.** UpSet plot of protein group overlaps between the sample-specific proteome and GENCODE reference. **C.** UpSet plot of gene overlaps between the sample-specific proteome, UniProt, and GENCODE reference. **D.** UpSet plot of peptide categories identified in searching the sample-specific proteome (we note that this number does not account for PTMs and is thus slightly lower). **E.** UpSet plot of protein isoform types between alternative splicing and variant isoforms. **F.** UpSet plot of overlaps of protein effects caused by genetic variation and alternative splicing. **G.** Number of variants identified via direct peptide evidence, stratified by zygosity. **H.** Number of variants identified via linked peptide evidence, stratified by zygosity. We counted a peptide as identified if it had a q-value of less than 0.01 and was not a decoy or contaminant. Identified genes were not compared to UNIPROT due to non-trivial gene mapping for UNIPROT. **E** and **F** are based on uniquely identified protein isoforms.

Within protein isoforms that had genetic variants, most tended to have only one or two variants, when accounting for both homozygous and heterozygous variants (Fig. 1D). Splice variants produced a more nuanced picture, with the majority of splice variants being categorized as protein novel not in catalog (pNNC), followed by novel in catalog (pNIC), incomplete splice mismatches (pISM), and orphaned exons (OE) (Fig. 1D) when categorizing their protein isoform structure.

Long-read RNA-seq-based personalized proteomes enable improved protein isoform identification and characterization

In MS searches of our sample-specific haplotype-resolved database, as well as reference UniProt and GENCODE databases (see Methods), we found that all databases were largely concordant in terms of identified peptides (Fig. 2A). In particular, 122,951 identified peptides were common to all databases, representing,

respectively, 98.8% of identified sample database peptides, 95.8% of identified UniProt database peptides, and 99.4% of identified GENCODE database peptides. Overall, the UNIPROT database identified the most peptides, totaling 129,048, while the sample-specific database and GENCODE identified 3,923 and 4,700 fewer, respectively. Peptides that were detected (that is, had less than 1% FDR) in only one database were relatively frequently completely absent from other databases, especially UNIPROT, which had many peptides that were absent from both GENCODE and the sample-specific database (Fig. S1A-B). In contrast, those peptides detected in two databases were typically present in the remaining database but narrowly missed the FDR cutoff of 1%. Overall, in the search against the sample-specific database, we found peptide-level FDRs to be well calibrated in separating between decoys and non-decoy protein isoforms, both overall and stratified by peptide categories (reference peptides, alternative splicing-associated peptides, and variant peptides) (Fig. S1C).

Protein inference was compared only between our sample-specific database and GENCODE, due to the noise inherent when directly matching UniProt and GENCODE transcript identifiers. Protein inference identified 10,393 protein groups that were identical (that is, the protein groups contained exactly the same transcript identifiers) between the sample-specific search and the GENCODE search, with 364 protein groups being unique to the sample-specific search and 309 to GENCODE (Fig. 2B). Thus, the sample-specific search identified a total of 10,757 protein groups, 55 more than GENCODE. Upon further investigation of the reasons contributing to these protein group mismatches, we found that most cases were because the sample-specific database tended to identify protein groups in which at least one protein was absent from all protein groups in GENCODE (Fig. S1D). Cases where protein groups were not shared due to being a subset or superset of a protein group in the other database were rare (Fig. S1D).

We also performed gene-level analysis, determining that a gene was identified if a protein group was identified for which all proteins were products of this particular gene. Gene-level search results largely recapitulated our peptide-level search results. A total of 9,939 genes were identified in common between all three databases (sample-specific, GENCODE, and UniProt), with UniProt identifying the most genes with 10,585, and the sample-specific database and GENCODE trailing it slightly with 10,194 and 10,188 identified genes, respectively (Fig. 2C).

Most peptides observed in the sample-specific search results were reference peptides. Peptides that did not match the reference sequence were classified as variant peptides, splice peptides or both (see Methods). In particular, of the 114,421 annotatable total peptides identified with the sample-specific database, 114,421 (99%) were reference peptides under this definition, with 744 and 21 falling under the splice-specific and variant umbrella, under our definition (Fig. 2D). When investigating uniquely identified protein isoforms (that is, protein groups that only contained a single transcript id), we found no association between isoform categories (Fig. 2E) and only minor linkages between isoform effects (Fig. 2F), with one identified protein group presenting both an indel and a frameshift and one protein group presenting an indel and a resolved frameshift.

In addition to performing analyses on the protein isoform level, we also investigated the degree to which the search identified particular variants, i.e., unique peptide evidence directly confirmed a particular variant. Directly identified variants are those for which direct, unique peptide evidence was observed. Linked variants, conversely, are confirmed through their presence on the same haplotypes a directly identified variant. Thus, linked variants are not directly identified via peptide evidence but by a combination of peptide evidence and variant phasing. The search of our sample-specific database identified a total of 538 variants, of which 258 were homozygous and the rest heterozygous (Fig. 2G). Furthermore, the direct identification of indels was

Fig. 3 Sample-specific haplotype-resolved MS database searches enable the characterization and identification of broad categories of protein isoforms. **A.** Examples of protein models and corresponding inferred proteins of protein isoforms categories identifiable with our workflow. **B.** Example browser track for an alternatively spliced protein isoform on the TMEM63B gene. **C.** Example browser track for a homozygous variant protein isoform on the TYW3 gene.**D.** Example browser track for a heterozygous variant protein isoform on the USP36 gene.

extraordinarily rare, with only three heterozygous and no homozygous indel having been identified via direct peptide evidence in our search. Variants identified by linkage were more abundant overall, with 201 linked heterozygous variants being identified in our search, of which 32 were indels and 169 SNVs (Fig. 2H). Linked variants had significantly more homozygous variants than heterozygous, with a total of 1,072 heterozygous variants having been identified, of which 56 were indels and 1016 SNVs.

Haplotype-resolved proteogenomics enables the characterization of broad classes of protein isoforms

Overall, we show that haplotype-resolved proteogenomics approaches based on lrRNA-seq enable the characterization of broad classes of protein isoforms, including reference isoforms, alternatively-spliced isoforms, variant isoforms containing homozygous genetic variation, and variant isoforms containing one or more heterozygous variants (Fig. 3A).

Accounting for sample-specific alternative splicing events enables a more accurate protein model, which can enhance protein inference (by discarding reference protein isoform not present in the sample) and permit

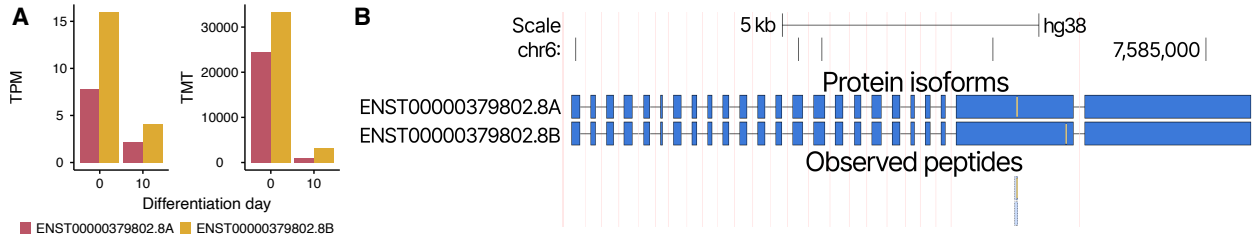

Fig. 4 Sample-specific haplotype-resolved MS database construction enables the investigation of proteomic complexity in arbitrary models. **A.** Transcript per Million (TPM) lrRNA-seq quantifications and TMT MS abundances for both haplotypes of the ENST00000379802.8 protein isoform on the DSP gene on days zero and then of the iPSC-osteoblast differentiation. **B.** Browser track highlighting the confirmed presence of both haplotypes of the ENST00000379802.8 transcript through unique peptides (shared peptides not shown).

more sensitivity (via the inclusion of novel protein isoforms not present in the reference) [4, 32, 33] (Fig. 3A, B). Genetic variation has also repeatedly been identified as a major contributor to proteomic complexity [34, 35] and enables more specific protein inference via additional variant peptides (Fig. 3A, C). Lastly, phased heterozygous genetic variation allows for not only potentially improved protein inference, but also the identification of protein haplotypes (Fig. 3A, D). The analysis of phased heterozygous genetic variation also enables the identification of linked variants (Fig. 3A), i.e., variants for which we have isoform-level peptide evidence but not necessarily variant-level evidence.

Sample-specific proteogenomics enables the analysis of condition-effect or differentiation models

Lastly, we present a vignette that highlights the applicability of our approach on arbitrary models, including ones that present a condition-specific effect or differentiation.

We obtained publicly available matched PacBio lrRNA-seq and TMT-labeled MS data from an iPSC into osteoblast differentiation [36]. We applied our pipeline to these data and highlight the ENST00000379802.8 transcript, belonging to the *DSP* gene. Consistent with the expected behavior of a gene important for cell cohesion [37], we find increased expression and TMT abundance in iPSCs (Day 0) relative to osteoblasts (Fig. 4A). In addition, lrRNA-seq and TMT data agree that both haplotypes of the ENST00000379802.8 are present, although the second (B) haplotype has lower abundance. The two haplotypes both exhibit a single (coding) heterozygous SNV, although their respective presences in the TMT are confirmed through peptides at an identical position, once with and once without the heterozygous variant in question (Fig. 4B).

Overall, our case study shows that our method can be used not only to describe and investigate the proteome of singular samples, but also to highlight differences between samples. Thus, our method can be useful for multiple biological applications, such as the quantification of protein levels contributed by A1AT M-haplotypes or the validation of neoantigens [38, 39].

Discussion

Our work presents, to the best of our knowledge, the first pipeline that creates and enables searching sample-specific haplotype-resolved proteomes from sample-matched lrRNA-seq and MS data. In particular, we have presented a Snakemake workflow that enables this task in a modular manner, relying where possible on existing open-source methods, easing future maintenance and extension to include new methods. In addition, we have demonstrated that variant phasing, similar to variant calling, can be performed well using lrRNA-seq, especially in areas with high coverage. We also illustrated the theoretical and observed proteome complexity observable in a WTC11 cell line sample. The majority of proteome complexity was still observed to be driven

by homozygous genetic variation and heterozygous genetic variation for which phasing was not required. Nevertheless, our results point toward the increased future importance of haplotype-resolved proteomes, especially as sequencing as well as MS depth and accuracy continue to increase. Lastly, we highlighted the applicability of investigating sample-specific proteomes of realistic biological systems in the ten-day differentiation from iPSCs to osteoblasts.

Concordant with previous work, genetic variation was a more dominant contributor to proteome complexity relative to alternative splicing, both for the observed and the theoretical proteome [13]. Interestingly, the number of protein isoforms inheriting two or more heterozygous variants was relatively small, although this was likely due to the low depth of our lrRNA-seq data. There was limited linkage between genetic variation and alternative splicing, with only a handful of protein isoforms presenting both types of variation in the theoretical proteome and none of them being identified during the search. These results suggest that the processes of genetic variation and alternative splicing may be largely independent, given that we do not find an increased proportion of alternatively spliced isoforms with variants, relative to all alternatively spliced isoforms. It should be noted that we did not attempt to identify haplotype-specific alternative splicing, as some methods have recently incorporated, which may impact this aspect [40]. Another point of note is the total number of identified variants, either via direct peptide evidence or via linkage. Interestingly, we found that the number of identified indels was much larger for the linked variants than the direct ones, both in absolute numbers and proportionally. To the best of our knowledge, this feature is caused by the fact that the direct identification of an indel requires coverage of the complete indel by a peptide, whereas linkage requires coverage of only a single SNVs on the same haplotype.

Our work has some limitations. First, we focused on PacBio Iso-Seq lrRNA-seq data from a WTC11 cell line in our work as well as PacBio Iso-Seq on the iPSC-osteoblast differentiation. In addition, our benchmarks were also performed on PacBio Iso-Seq data. As such, some of our conclusions may be specific to the Iso-Seq technology or one of the two systems we investigated. In addition, some of our statements regarding proteome complexity, while largely concordant with previous work [13], may also be a function of the lower sequencing depth of our data. The phasing benchmark we performed was based exclusively on GIAB data. While this is standard, it is possible that some of the methods we benchmarked for variant phasing were implicitly tuned to GIAB data during development, due to its frequent usage as a development dataset, potentially skewing some of our results. Our workflow currently does not account for changed start codons. However, the number of these cases was extremely small in our experiments, similar to changed stop codons ($< 0.1\%$). A final limitation in our current pipeline concerns discoverability. We leave an investigation of the discoverability of sample-specific variants in protein haplotypes to future work, noting that previous work has already investigated this point for population-based variation [20, 21].

There are several potential avenues for future work, following our introduction of a sample-specific haplotype-resolved proteogenomics pipeline and validation of phasing performance with lrRNA-seq data. First, variants phased with lrRNA-seq are now also increasingly used for other tasks, such as the quantification of ASE [28]. While our benchmark demonstrated that phasing variants with lrRNA-seq can be done relatively accurately, this is less clear for downstream tasks such as ASE and should thus be further investigated. Our work primarily highlights the usage of haplotype-resolved proteomes for improved isoform characterization. It would also be highly desirable to investigate the implications of haplotype-resolved transcriptomes and proteomes for other tasks, such as neoantigen prediction and validation. A last and perhaps particularly promising route for future investigation is further exploring protein inference algorithms in the

context of haplotype-resolved proteomes. Several current approaches discard shared peptides to achieve accurately calibrated FDRs. Given the fact that haplotype-resolved proteomes produce a larger number of shared peptides, these methods may want to account for this fact when running on haplotype-resolved proteomes.

Conclusion

We have introduced a sample-specific workflow that generates and performs MS search on haplotype-resolved proteomes, based on sample-matched MS and lrRNA-seq input data. Our results show that generating such proteomes works well and can produce enhanced protein isoform identification and characterization. Our work also reveals that for regions with relatively high coverage, lrRNA-seq data can not only be used to call variants accurately but also to phase them [41].

We envision our project as one possible go-to pipeline for lrRNA-seq proteogenomics, aided especially by our workflow choices that emphasize modularity. In this way, we believe our pipeline may be helpful both for basic research in proteogenomics as well as applications such as neoantigens prediction and validation.

Methods

Datasets

Miscellaneous

Haplotype-resolved quantification. For haplotype-resolved quantification (Fig. 4A), haplotype-resolved transcriptomes were generated using g2gtools 3.1.0, followed by quantification of the lrRNA-seq data using lr-kallisto 0.51.1 [42].

WTC11 cell line

PacBio lrRNA-seq data.

PacBio lrRNA-seq data (i.e., Iso-Seq) on the WTC11 cell line was downloaded from NCBI (SRR18130587) [41] (see Data Availability). Further experimental details are described in [41].

Mass spectrometry data.

Bottom-up, trypsin-digested mass-spectrometry data on the WTC11 cell line was acquired directly from Bollon et al. [43] as RAW files and subsequently converted to mzML for database searching using ProteoWizard `msconvert` using the `--zlib --32 --filter "peakPicking true 1-"` parameters [44] (see Data Availability). Further experimental details are available from Bollon et al. [43] upon request.

iPSC-osteoblast differentiation

PacBio lrRNA-seq data. PacBio lrRNA-seq data (i.e., Iso-Seq) data on the iPSC-osteoblast differentiation cell line were obtained directly from the authors [36].

Mass spectrometry data. TMT MS data on the iPSC-osteoblast differentiation cell line were obtained directly from the authors [36].

Genome in a Bottle

lrRNA-seq data.

uBAM files were downloaded from NCBI for the HG002 and HG005 Iso-Seq samples (see Data Availability).

References

Reference genome

We used the `GCA_000001405.15_GRCh38_no_alt_analysis_set.fna` hg38 genome as the reference genome throughout all of our analyses.

Reference transcriptome

The primary assembly of GENCODE V45 was used as our reference transcriptome throughout this work [45].

Contaminants

We used the common Repository of Adventitious Proteins, cRAP, as contaminants in all mass-spectrometry database searches across all of our experiments.

UniProt

The UP000005640 human reference proteome, including Swiss-Prot and TrEMBL proteins, as well as Swiss-Prot isoforms, was used as a reference proteome, in addition to the GENCODE V45-derived proteome, in our database searches [46].

Reproducibility

All products, including figures and tables, created and used within this paper are fully reproducible via a Snakemake pipeline available on the paper branch of our Github repository.

Experiments were run on AMD EPYC 7742 64-Core Processor CPUs under Ubuntu 22.04, using Snakemake 8.2.3 and SingularityCE 3.11.4. All other project-relevant software was run using mamba environments within Singularity containers via Snakemake. All software versions are pinned within their respective conda environments within the Snakemake pipeline and additionally logged on Zenodo (see Data Availability) [47].

Additional pipeline details

Alignment

Our pipeline begins with quality control and alignment. In particular, we first remove sequencing adapters and trim polyA tails from the lrRNA-seq data using `lima` and `isoseq refine`. Following this, we align the lrRNA-seq data to the genome using `minimap2` [48]. lima and isoseq refine were run using their default parameters. minimap2 was run using `minimap2 -ax splice:hq -uf`, as recommended by the manual for PB Iso-Seq data.

Variant calling

We call variants based on genome alignments using `Clair3-RNA` [49], followed by appropriate filtering of the called variants using R code. Afterward, we phase the variants using `whatshap` [23, 24]. Following this, we postprocess the phased variants by removing any heterozygous variants belonging to phase blocks that do not stretch across the full CDS region in which they are located. In addition, we arbitrarily assign phase blocks to variants overlapping only CDS regions that contain solely this variant, since phasing is arbitrary in these cases.

For variant calling with Clair3-RNA, we kept both indels and SNVs. In addition, we required a minimum quality of called variants of zero, a minimum allele frequency for called SNVs of 0.08, minimum allele frequency

for called indels of 0.15, a minimum coverage of three for all called variants, and a minimum mapping quality of five for an alignment to be considered.

Post-variant-calling, we filtered all variants not fulfilling a minimum GQ of five, in addition to removing multi-allelic variants.

Variant phasing using whatshap was performed using default parameters, except that we turned on `--distrust-genotypes`, as it has been shown to improve performance in some scenarios [25].

Isoform discovery

Using the genome-aligned lrRNA-seq reads, along with the reference transcriptome, we perform isoform discovery using `Bambu` [50]. On the Bambu-derived transcriptome, we call open-reading frames for all novel isoforms using `orfanage` [51]. We then post-process the resulting novel and reference isoforms by filtering for valid full CDS regions and a minimum number of nucleotides.

We ran Bambu with default parameters, letting it choose a novel discovery rate, which the method set to 0.067 for our WTC11 cell line sample.

For the reference GENCODE annotation, we standardized the input GTF using `gffread` [52], adjusting stop codons (`--adj-stop`) and filtering for complete coding sequences (`-J`). Following this, we filtered the resulting GENCODE coding transcripts for a minimum CDS length of 100 nucleotides.

For ORF calling, we inputted the union of all resulting filtered GENCODE coding transcripts and the novel Bambu isoforms, calling ORFs only for Bambu isoforms. orfanage was run with default parameters, using the filtered GENCODE coding transcripts as a reference and using the `BEST` mode. Following ORF calling, isoforms were filtered for coding sequences of a minimum length of 100 nucleotides.

Mass spectrometry database creation

To create a personalized proteome, based on our phased variants and partially novel transcriptome, we use `haplosaurus` [19]. In particular, we input a reference genome, the Bambu-derived transcriptome and the phased variants to generate protein haplotypes. We then post-process the haplosaurus output by first re-filtering the resulting protein isoforms for complete CDS region containment and a minimum codon length, since, for example, changed stop codons may have led to the truncation of some protein isoforms. Afterward, we generate unique IDs for each haplotype-resolved protein, appending A and B to the protein ID to signify the first and second haplotypes per protein isoforms. Then, we generate a sequence-unique set of protein isoforms, tracking which protein IDs result in the same protein sequence. Lastly, we concatenate contaminants and generate decoys for use in database searching using `DecoyPYrat` [53], to finalize the protein database.

Haplotype labels (A, B) are assigned per CDS, aiming to keep haplotypes together that inheret the same genotype of a particular heterozygous variant. There are some cases where the same heterozygous variant is assigned different haplotype labels, as haplotype labels are assigned per isoform, without looking back at all previously created labels. We've found these cases to be rarer than not and not to deter post-search visualization and annotation.

DecoyPYrat was run for trypsin, allowing a maximum of two missed cleaves, with a minimum peptide length of five and a maximum peptide length of 100, for a maximum of 100 iterations for decoy generation.

Mass spectrometry search

We then perform a database search using `sage` [54], inputting our generated MS database and the sample MS data. Based on the search results, we perform protein inference using a standard greedy set cover approach

[55] that we implement using `RcppGreedySetCover` [56] in R. In addition, we create BAM files that enable the creation of browser tracks for easy visualization of the search results by extracting genome-space sequences of both protein isoforms and observed peptides in R. Lastly, we perform post-search annotation in R, aggregating information from isoform discovery, variant calling, and protein haplotype generation to provide detailed information on if and how each protein isoform differs from the set of reference protein isoforms.

Search using sage was run using identical parameters as the decoy generation, where applicable. Full sage configuration files are available on Github and Zenodo for reproducibility (see Code Availability).

All post-search annotation was performed using an FDR of 0.01, using the q-value estimates produced by sage. Browser tracks for protein isoforms and matching peptides are created via custom matching in R to back-translate protein sequences to transcriptomic sequences, accounting only for non-silent variation. The resulting FASTA files were then assigned to a haplotype (homozygous, first haplotype, or second haplotype). This is trivial for protein isoforms. For peptides, we assigned them to a particular haplotype if they only matched one haplotype of one protein isoform and assigned them to be homozygous in all other cases. Protein isoform sequences were aligned to the genome using minimap2, using standard parameters for the spliced alignment of lrRNA-seq while peptide sequences were aligned using STAR [57], due to their short length.

The annotation file contains the following information per haplotype-resolved protein, constructed primarily out of haplosaurus and SQANTI protein output files, and aggregated via R code:

1. Protein splice category as determined by SQANTI protein [4].
2. Number of contained non-silent genetic variation, broken down by SNV/indel and homozygous/heterozygous.
3. List of all contained non-silent genetic variation.
4. Diff string between the observed protein isoform and the reference without any genetic variation.
5. Flags for several protein effects, including indel, changed stop codon, frameshift, and resolved frameshift.

Benchmark structure

We performed a phasing-only benchmark, in which we re-phased GIAB data for the HG002, and HG005 GIAB cell lines and estimated phasing accuracy- and completeness, relative to a ground truth of ancestry-derived phasing. Importantly, this benchmark evaluates only phasing, as we reuse the true unphased genetic variation as an input.

Preprocessing

To prepare the GIAB datasets for variant calling and phasing, we preprocessed them in the following way.

PB Iso-Seq.

We performed adapter removal and polyA tail trimming using lima and isoseq refine. Following this, data was aligned to the reference genome using standard settings for Iso-Seq data (`minimap2 -ax splice:hq -uf`).

Phasing-only benchmark.

In our phasing-only benchmark, we considered datasets from GIAB, HG002, and HG005, evaluating PB Iso-Seq for each.

As performance metrics to evaluate phasing, we considered both accuracy and phasing completeness. We measured accuracy via the switch error rate as computed by `whatshap compare` and phasing completeness

using both the total and relative number of phased variants as well as the total and relative number of fully phased regions of interest, which we consider to be CDS-, transcript- and gene-level.

In addition, similar to previous work on variant calling with lrRNA-seq, we stratify results by the required minimum coverage a considered variant has to have in the lrRNA-seq data to be included in the benchmark [41]. We considered minimum input coverages of 5, 10, 25, and 100. Furthermore, we stratified results by whether we phased only SNVs or SNVs and indels jointly, as this has been shown to impact performance for variant phasing [23].

Phasing methods. All methods were run with default parameters unless otherwise noted.

WhatsHap.

We turned on the `--distrust-genotypes` option, as it has been shown to increase performance in some cases [25].

HapCut2.

`extractHAIRS` was run with the `--ont` flag enabled, as this is recommended in the manual for long-read data.

Acknowledgements. The authors thank Madison M. Mehlferber and Elizabeth Tseng for helpful feedback and discussions throughout the project. The authors thank Saikat Bandyopadhyay for feedback on running a proof-of-concept iteration of the pipeline on Singularity. We would also like to acknowledge Mayank Murali for general feedback on an earlier iteration of the pipeline, Marina Panteli for exploring early possibilities for incorporating genetic variant information into long-read-derived proteomes, and Roger Volden for general feedback on the project.

Declarations

Ethics approval and consent to participate

This study used only previously sequenced, publically available datasets and was thus exempt from ethics approval.

Consent for publication

Not applicable.

Data availability

Raw data.

Raw lrRNA-seq data for the WTC11 cell line is available from NCBI (SRR18130587). Raw MS data for the WTC11 cell line are available from Bollon et al. [43] upon request.

Raw lrRNA-seq for the GIAB datasets is available as follows:

1. HG002 PB Iso-Seq: NCBI.
2. HG005 PB Iso-Seq: NCBI.

Reference data.

Reference data is downloaded as part of our Snakemake pipeline (see Code availability). In addition, the GENCODE V45 primary assembly is available from GENCODE, the hg38 reference genome from NCBI, the cRAP contaminants from GPM, and the UniProt human reference proteome from UniProt.

In addition, GIAB phased genetic variation files are available from NCBI for HG002 and HG005.

Produced result data.

Result data is cached and publically available on Zenodo [47].

Materials availability

Not applicable.

Code availability

All code necessary to reproduce the results in our paper is available on the `paper` branch of our Github repository. All code has, in addition, been cached on Zenodo Zenodo, where it is also freely accessible [47].

Author contribution

Study conception: GMS, MDR. Study design: DW, GMS, MDR. Benchmark design: DW. Software implementation: DW. Analysis and interpretation of results: DW, GMS, MDR. Draft manuscript preparation: DW, GMS, MDR. All authors reviewed the results and approved the final version of the manuscript.

References

[1] Smith, L.M., Kelleher, N.L.: Proteoforms as the next proteomics currency. Science **359**(6380), 1106–1107 (2018)

[2] Smith, L.M., Agar, J.N., Chamot-Rooke, J., Danis, P.O., Ge, Y., Loo, J.A., Paša-Tolić, L., Tsybin, Y.O., Kelleher, N.L., Top-Down Proteomics, C.: The human proteoform project: defining the human proteome. Science advances **7**(46), 0734 (2021)

[3] Aebersold, R., Agar, J.N., Amster, I.J., Baker, M.S., Bertozzi, C.R., Boja, E.S., Costello, C.E., Cravatt, B.F., Fenselau, C., Garcia, B.A., *et al.*: How many human proteoforms are there? Nature chemical biology **14**(3), 206–214 (2018)

[4] Miller, R.M., Jordan, B.T., Mehlferber, M.M., Jeffery, E.D., Chatzipantsiou, C., Kaur, S., Millikin, R.J., Dai, Y., Tiberi, S., Castaldi, P.J., *et al.*: Enhanced protein isoform characterization through long-read proteogenomics. Genome biology **23**(1), 69 (2022)

[5] Browning, S.R., Browning, B.L.: Haplotype phasing: existing methods and new developments. Nature Reviews Genetics **12**(10), 703–714 (2011)

[6] Cesnik, A.J., Miller, R.M., Ibrahim, K., Lu, L., Millikin, R.J., Shortreed, M.R., Frey, B.L., Smith, L.M.: Spritz: a proteogenomic database engine. Journal of proteome research **20**(4), 1826–1834 (2020)

[7] McKenna, A., Hanna, M., Banks, E., Sivachenko, A., Cibulskis, K., Kernytsky, A., Garimella, K., Altshuler, D., Gabriel, S., Daly, M., *et al.*: The genome analysis toolkit: a mapreduce framework for analyzing next-generation dna sequencing data. Genome research **20**(9), 1297–1303 (2010)

[8] Kovaka, S., Zimin, A.V., Pertea, G.M., Razaghi, R., Salzberg, S.L., Pertea, M.: Transcriptome assembly from long-read rna-seq alignments with stringtie2. Genome biology **20**, 1–13 (2019)

[9] Cingolani, P.: Variant annotation and functional prediction: Snpeff. In: Variant Calling: Methods and Protocols, pp. 289–314. Springer, New York, NY (2012)

[10] Gordon, S.P., Tseng, E., Salamov, A., Zhang, J., Meng, X., Zhao, Z., Kang, D., Underwood, J., Grigoriev, I.V., Figueroa, M., *et al.*: Widespread polycistronic transcripts in fungi revealed by single-molecule mrna sequencing. PloS one **10**(7), 0132628 (2015)

[11] Wang, L., Park, H.J., Dasari, S., Wang, S., Kocher, J.-P., Li, W.: Cpat: Coding-potential assessment tool using an alignment-free logistic regression model. Nucleic acids research **41**(6), 74–74 (2013)

[12] Pardo-Palacios, F.J., Arzalluz-Luque, A., Kondratova, L., Salguero, P., Mestre-Tomás, J., Amorín, R., Estevan-Morió, E., Liu, T., Nanni, A., McIntyre, L., *et al.*: Sqanti3: curation of long-read transcriptomes for accurate identification of known and novel isoforms. Nature Methods **21**(5), 793–797 (2024)

[13] Salz, R., Bouwmeester, R., Gabriels, R., Degroeve, S., Martens, L., Volders, P.-J., Hoen, P.A.: Personalized proteome: Comparing proteogenomics and open variant search approaches for single amino acid variant detection. Journal of Proteome Research **20**(6), 3353–3364 (2021)

[14] Zook, J.M., Catoe, D., McDaniel, J., Vang, L., Spies, N., Sidow, A., Weng, Z., Liu, Y., Mason, C.E.,

Alexander, N., *et al.*: Extensive sequencing of seven human genomes to characterize benchmark reference materials. Scientific data **3**(1), 1–26 (2016)

[15] Eberle, M.A., Fritzilas, E., Krusche, P., Källberg, M., Moore, B.L., Bekritsky, M.A., Iqbal, Z., Chuang, H.-Y., Humphray, S.J., Halpern, A.L., *et al.*: A reference data set of 5.4 million phased human variants validated by genetic inheritance from sequencing a three-generation 17-member pedigree. Genome research **27**(1), 157–164 (2017)

[16] Workman, R.E., Tang, A.D., Tang, P.S., Jain, M., Tyson, J.R., Razaghi, R., Zuzarte, P.C., Gilpatrick, T., Payne, A., Quick, J., *et al.*: Nanopore native rna sequencing of a human poly (a) transcriptome. Nature methods **16**(12), 1297–1305 (2019)

[17] McLaren, W., Gil, L., Hunt, S.E., Riat, H.S., Ritchie, G.R., Thormann, A., Flicek, P., Cunningham, F.: The ensembl variant effect predictor. Genome biology **17**, 1–14 (2016)

[18] Consortium, .G.P., *et al.*: A global reference for human genetic variation. Nature **526**(7571), 68 (2015)

[19] Spooner, W., McLaren, W., Slidel, T., Finch, D.K., Butler, R., Campbell, J., Eghobamien, L., Rider, D., Kiefer, C.M., Robinson, M.J., *et al.*: Haplosaurus computes protein haplotypes for use in precision drug design. Nature communications **9**(1), 4128 (2018)

[20] Vašíček, J., Skiadopoulou, D., Kuznetsova, K.G., Wen, B., Johansson, S., Njølstad, P.R., Bruckner, S., Käll, L., Vaudel, M.: Finding haplotypic signatures in proteins. GigaScience **12**, 093 (2023)

[21] Vašíček, J., Kuznetsova, K.G., Skiadopoulou, D., Njølstad, P.R., Johansson, S., Bruckner, S., Käll, L., Vaudel, M.: Prohap enables proteomic database generation accounting for population diversity. bioRxiv, 2023–12 (2023)

[22] Vašíček, J., Skiadopoulou, D., Kuznetsova, K.G., Käll, L., Vaudel, M., Bruckner, S.: Prohap explorer: Visualizing haplotypes in proteogenomic datasets. arXiv preprint arXiv:2504.06282 (2025)

[23] Patterson, M., Marschall, T., Pisanti, N., Van Iersel, L., Stougie, L., Klau, G.W., Schönhuth, A.: Whatshap: weighted haplotype assembly for future-generation sequencing reads. Journal of Computational Biology **22**(6), 498–509 (2015)

[24] Martin, M., Ebert, P., Marschall, T.: Read-based phasing and analysis of phased variants with whatshap. In: Haplotyping: Methods and Protocols, pp. 127–138. Springer, New York, NY (2022)

[25] Holt, J.M., Saunders, C.T., Rowell, W.J., Kronenberg, Z., Wenger, A.M., Eberle, M.: Hiphase: jointly phasing small, structural, and tandem repeat variants from hifi sequencing. Bioinformatics **40**(2), 042 (2024)

[26] Shafin, K., Pesout, T., Chang, P.-C., Nattestad, M., Kolesnikov, A., Goel, S., Baid, G., Kolmogorov, M., Eizenga, J.M., Miga, K.H., *et al.*: Haplotype-aware variant calling with pepper-margin-deepvariant enables high accuracy in nanopore long-reads. Nature methods **18**(11), 1322–1332 (2021)

[27] Castel, S.E., Mohammadi, P., Chung, W.K., Shen, Y., Lappalainen, T.: Rare variant phasing and haplotypic expression from rna sequencing with phaser. Nature communications **7**(1), 12817 (2016)

[28] Quinones-Valdez, G., Amoah, K., Xiao, X.: Long-read rna-seq demarcates cis-and trans-directed alternative rna splicing. bioRxiv, 2024–06 (2024)

[29] Berger, E., Yorukoglu, D., Zhang, L., Nyquist, S.K., Shalek, A.K., Kellis, M., Numanagić, I., Berger, B.: Improved haplotype inference by exploiting long-range linking and allelic imbalance in rna-seq datasets. Nature communications **11**(1), 4662 (2020)

[30] Mölder, F., Jablonski, K.P., Letcher, B., Hall, M.B., Tomkins-Tinch, C.H., Sochat, V., Forster, J., Lee, S., Twardziok, S.O., Kanitz, A., et al.: Sustainable data analysis with snakemake. F1000Research **10** (2021)

[31] Edge, P., Bafna, V., Bansal, V.: Hapcut2: robust and accurate haplotype assembly for diverse sequencing technologies. Genome research **27**(5), 801–812 (2017)

[32] Nilsen, T.W., Graveley, B.R.: Expansion of the eukaryotic proteome by alternative splicing. Nature **463**(7280), 457–463 (2010)

[33] Nesvizhskii, A.I.: Proteogenomics: concepts, applications and computational strategies. Nature methods **11**(11), 1114–1125 (2014)

[34] Dimitrakopoulos, L., Prassas, I., Diamandis, E.P., Nesvizhskii, A., Kislinger, T., Jaffe, J., Drabovich, A.: Proteogenomics: opportunities and caveats. Clinical chemistry **62**(4), 551–557 (2016)

[35] Sheynkman, G.M., Shortreed, M.R., Frey, B.L., Scalf, M., Smith, L.M.: Large-scale mass spectrometric detection of variant peptides resulting from nonsynonymous nucleotide differences. Journal of proteome research **13**(1), 228–240 (2014)

[36] Abood, A., Mesner, L.D., Jeffery, E.D., Murali, M., Lehe, M.D., Saquing, J., Farber, C.R., Sheynkman, G.M.: Long-read proteogenomics to connect disease-associated sqtls to the protein isoform effectors of disease. The American Journal of Human Genetics **111**(9), 1914–1931 (2024)

[37] Smith, E.A., Fuchs, E.: Defining the interactions between intermediate filaments and desmosomes. The Journal of cell biology **141**(5), 1229–1241 (1998)

[38] Vitiello, A., Zanetti, M.: Neoantigen prediction and the need for validation. Nature biotechnology **35**(9), 815–817 (2017)

[39] Jager, S., Cramer, D.A., Heck, A.J.: Normal alpha-1-antitrypsin variants display in serum allele-specific protein levels. Journal of Proteome Research **22**(4), 1331–1338 (2023)

[40] Tang, A.D., Felton, C., Hrabeta-Robinson, E., Volden, R., Vollmers, C., Brooks, A.N.: Detecting haplotype-specific transcript variation in long reads with flair2. Genome Biology **25**(1), 173 (2024)

[41] Souza, V.B., Jordan, B.T., Tseng, E., Nelson, E.A., Hirschi, K.K., Sheynkman, G., Robinson, M.D.: Transformation of alignment files improves performance of variant callers for long-read rna sequencing data. Genome Biology **24**(1), 91 (2023)

[42] Loving, R.K., Sullivan, D.K., Booeshagi, A.S., Reese, F., Rebboah, E., Sakr, J., Rezaie, N., Liang, H.Y., Filimban, G., Kawauchi, S., et al.: Long-read sequencing transcriptome quantification with lr-kallisto.

bioRxiv, 2024–07 (2025)

[43] Bollon, J., Shortreed, M.R., T JORDAN, B., Miller, R., Jeffery, E., Cavalli, A., Smith, L.M., Dewey, C., Sheynkman, G.M., Tiberi, S.: Isobayes: a bayesian approach for single-isoform proteomics inference. bioRxiv (2024)

[44] Holman, J.D., Tabb, D.L., Mallick, P.: Employing proteowizard to convert raw mass spectrometry data. Current protocols in bioinformatics **46**(1), 13–24 (2014)

[45] Frankish, A., Carbonell-Sala, S., Diekhans, M., Jungreis, I., Loveland, J.E., Mudge, J.M., Sisu, C., Wright, J.C., Arnan, C., Barnes, I., *et al.*: Gencode: reference annotation for the human and mouse genomes in 2023. Nucleic acids research **51**(D1), 942–949 (2023)

[46] Uniprot: the universal protein knowledgebase in 2023. Nucleic acids research **51**(D1), 523–531 (2023)

[47] Wissel, D., Sheynkman, G.M., Robinson, M.D.: Sample-specific haplotype-resolved isoform characterization via long-read RNA-seq-based proteogenomics. Zenodo (2025). https://doi.org/10.5281/zenodo.17674329 . https://doi.org/10.5281/zenodo.17674329

[48] Li, H.: Minimap2: pairwise alignment for nucleotide sequences. Bioinformatics **34**(18), 3094–3100 (2018)

[49] Zheng, Z., Yu, X., Chen, L., Lee, Y.-L., Xin, C., Wong, A.O.K., Jain, M., Kesharwani, R.K., Sedlazeck, F.J., Luo, R.: Clair3-rna: A deep learning-based small variant caller for long-read rna sequencing data. bioRxiv, 2024–11 (2024)

[50] Chen, Y., Sim, A., Wan, Y.K., Yeo, K., Lee, J.J.X., Ling, M.H., Love, M.I., Göke, J.: Context-aware transcript quantification from long-read rna-seq data with bambu. Nature methods **20**(8), 1187–1195 (2023)

[51] Varabyou, A., Erdogdu, B., Salzberg, S.L., Pertea, M.: Investigating open reading frames in known and novel transcripts using orfanage. Nature computational science **3**(8), 700–708 (2023)

[52] Pertea, G., Pertea, M.: Gff utilities: Gffread and gffcompare. F1000Research **9** (2020)

[53] Wright, J.C., Choudhary, J.S.: Decoypyrat: fast non-redundant hybrid decoy sequence generation for large scale proteomics. Journal of proteomics & bioinformatics **9**(6), 176 (2016)

[54] Lazear, M.R.: Sage: an open-source tool for fast proteomics searching and quantification at scale. Journal of Proteome Research **22**(11), 3652–3659 (2023)

[55] Zhang, B., Chambers, M.C., Tabb, D.L.: Proteomic parsimony through bipartite graph analysis improves accuracy and transparency. Journal of proteome research **6**(9), 3549–3557 (2007)

[56] Kaeding, M.: RcppGreedySetCover: Greedy Set Cover. (2018). R package version 0.1.0. https://CRAN.R-project.org/package=RcppGreedySetCover

[57] Dobin, A., Davis, C.A., Schlesinger, F., Drenkow, J., Zaleski, C., Jha, S., Batut, P., Chaisson, M., Gingeras, T.R.: Star: ultrafast universal rna-seq aligner. Bioinformatics **29**(1), 15–21 (2013)