

RESOURCE ARTICLE OPEN ACCESS

SwarmGenomics: A Unified Pipeline for Individual-Based Whole-Genome Analyses

Aure Kylmänen | Yu-Chi Chen | Sahar Javaheri Tehrani | Nikolas Vellnow | Justin J. S. Wilcox | Toni I. Gossmann 

Computational Systems Biology, Faculty of Biochemical and Chemical Engineering, TU Dortmund University, Dortmund, Germany

Correspondence: Toni I. Gossmann (toni.gossmann@tu-dortmund.de)**Received:** 19 September 2025 | **Revised:** 16 January 2026 | **Accepted:** 17 February 2026**Keywords:** bioinformatics pipeline | heterozygosity | mitochondrial genome | NUMT | PSMC | repeat annotation | runs of homozygosity | unmapped reads | whole-genome sequencing

ABSTRACT

Advances in sequencing technologies have made whole-genome data widely accessible, enabling research in population genetics, evolutionary biology, and conservation. However, analysing whole-genome sequencing (WGS) data remains challenging, often requiring multiple specialised tools and substantial bioinformatics expertise. We present SwarmGenomics, a modular, user-friendly command-line pipeline for reference-based genome assembly and individual-based genetic analyses. The pipeline integrates seven modules: heterozygosity estimation, runs of homozygosity detection, Pairwise Sequentially Markovian Coalescent (PSMC) analysis, unmapped reads classification, repeat analysis, mitochondrial genome assembly, and nuclear mitochondrial DNA segment (NUMT) identification. Each module can be run independently or as part of a complete workflow. We demonstrate the pipeline's utility with a case study on the giant panda (*Ailuropoda melanoleuca*), revealing insights into genetic diversity, inbreeding history, historical population size changes, transposable element activity, and microbial contamination. SwarmGenomics lowers the entry barrier for genomic analysis of diploid, non-model species, serving both as a research and teaching tool. The pipeline and documentation are available at <https://github.com/AureKylmanen/Swarmgenomics>.

1 | Motivation

Whole-genome sequencing (WGS) data is increasingly accessible due to significant advancements and declining costs of sequencing technologies. A primary catalyst for this development was the Human Genome Project, which spurred innovations in high-throughput sequencing and bioinformatics (Hood and Rowen 2013). This project also led to the establishment of publicly accessible genomic databases, such as GenBank, Ensembl and the UCSC Genome Browser (Birney et al. 2009). The rapid accumulation of genomic information in these databases has begun to outpace the availability of computational resources and analytical tools necessary for its effective use. Complementary initiatives, such as Bird10K (Prum et al. 2015) and Bat1K (Teeling et al. 2018), have further contributed to the expansion of genomic resources. Currently, the NCBI genome database contains genomic data for approximately 4264 mammalian,

2096 avian, and 2578 reptilian species (NCBI 2025a), representing a continually growing resource that supports research across more than 6400 mammalian species and approximately 11,000 bird and reptile species. This expanding genomic resource creates unprecedented opportunities for comparative and evolutionary studies across a wide range of species.

In parallel, large-scale genome consortia such as the Vertebrate Genomes Project (Rhie et al. 2021) and the Earth BioGenome Project (Lewin et al. 2022) are generating large volumes of high-quality reference genomes. As these comprehensive reference genomes become available, there is an increasing emphasis on moving beyond just genome assembly towards in-depth population genomics analyses. Such analyses are crucial for uncovering patterns of genetic variation within and between populations, identifying signatures of natural selection, and better understanding the evolutionary history of species. Moreover,

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2026 The Author(s). *Molecular Ecology Resources* published by John Wiley & Sons Ltd.

detailed population genomic data are essential for informing conservation strategies, particularly for endangered species facing threats from habitat loss, climate change, and human activity. As the number of reference genomes continues to grow along with improvements in sequencing technologies, the demand for accessible and reliable tools is becoming more pressing.

Despite significant progress, the analysis of whole-genome sequencing data remains technically and computationally demanding, often requiring extensive bioinformatics and programming expertise. This presents a substantial barrier for researchers from diverse scientific backgrounds. Laboratories are sometimes required to hire dedicated bioinformatics staff to overcome these challenges (Oakeson et al. 2017). The typical genomic analysis workflow involves multiple specialised software packages and several programming languages, leading to fragmented processes that can be especially daunting for beginners. To date, integrated pipelines that simplify and streamline genomic analyses into an accessible workflow remain limited.

Some tools such as the web-based platform Galaxy (Afgan et al. 2016) and pipelines like Genome After Party (Sanger 2025) have made significant progress lowering the entry barrier to bioinformatics by offering more user-friendly, integrated environments for genomic analyses. While Galaxy provides a graphical interface for many popular tools, Genome After Party focuses on streamlined, reproducible workflows for genome assembly and annotation. However, handling large whole-genome datasets or customising complex workflows may still require bioinformatics expertise. More recent reference-based pipelines, such as snpArcher (Mirchandani et al. 2024), provide highly efficient and reproducible workflows focused on variant calling, quality control, and downstream population genomic analyses in non-model organisms. While snpArcher is optimised for rapid and standardised variant detection, it primarily focuses on variant-centric analyses and does not encompass broader individual-based genomic analyses such as heterozygosity estimation, demographic history reconstruction, or repeat and mitochondrial genome characterisation.

To address these limitations, we developed SwarmGenomics, an integrated, user-friendly bioinformatics pipeline designed to facilitate reference-based genome assembly and genetic analyses. Specifically optimised for diploid, non-model organisms, SwarmGenomics broadens the accessibility of genomic analyses for species that remain understudied. The pipeline encompasses a straightforward, step-by-step assembly procedure followed by several analytical modules, including heterozygosity analysis, detection of runs of homozygosity, Pairwise Sequentially Markovian Coalescent (PSMC) analysis, characterisation of unmapped reads, repeat element annotation, mitochondrial genome assembly, and identification of nuclear mitochondrial DNA segments (NUMTs). SwarmGenomics is particularly well-suited for researchers with limited bioinformatics training and is intended to serve as both an educational resource and a self-learning tool.

To illustrate the practicality and educational value of SwarmGenomics, we demonstrate its application using real genomic data from the giant panda (*Ailuropoda melanoleuca*),

an iconic species of conservation interest. This case study highlights the pipeline's effectiveness as an accessible entry point into genomic research, providing insights into population genetics, evolutionary biology and conservation genetics.

1.1 | Individual-Based Genome Analysis

SwarmGenomics is explicitly designed for individual-based genome analysis, meaning that all downstream modules operate on whole-genome sequencing data from a single diploid individual rather than requiring population-level sampling. This design reflects a common practical situation in non-model organisms, where reference genomes are available but large population datasets are difficult or impossible to obtain.

Although many population genetic statistics are traditionally estimated from multiple individuals, a single diploid genome already contains substantial information about genetic diversity, inbreeding, and demographic history. Variation between the two homologous chromosomes of an individual captures the cumulative effects of mutation, recombination, and past population size changes. As a result, metrics such as heterozygosity, runs of homozygosity and coalescent-based demographic reconstructions can be inferred directly from individual genomes.

All analyses in SwarmGenomics therefore treat the sequenced diploid individual as the primary unit of inference. When multiple individuals are analysed independently using the same reference genome, their results can be compared in a consistent framework. In this way, SwarmGenomics supports both individual-level genomic inference and comparative analyses across individuals or species while remaining usable in settings where population-scale data are unavailable.

2 | Pipeline

SwarmGenomics is a versatile and user-friendly bioinformatics pipeline developed for whole-genome sequencing (WGS) analysis starting directly from raw paired-end reads by default, but can be adapted to single-end data as well. An overview of the pipeline's workflow is shown in Figure 1, with prospective future modules indicated in grey. Its modular design enables users to select and execute only the analyses relevant to their specific research needs (Table 1). The pipeline is implemented primarily in Bash, Python and R; however, no prior experience with these programming languages is necessary to use it effectively. We recommend running SwarmGenomics within a virtual machine (VM) to simplify software setup and enhance reproducibility in genomic analyses (Nocq et al. 2013). Alternatively, users without access to a VM can run the pipeline on a dedicated Linux workstation or a sufficiently powerful laptop.

2.1 | Customising Parameters

The SwarmGenomics pipeline is designed for reproducible genomic analyses and can be run either stepwise or through individual scripts that share a common params.txt configuration file. Users first edit params.txt to specify paths to input data

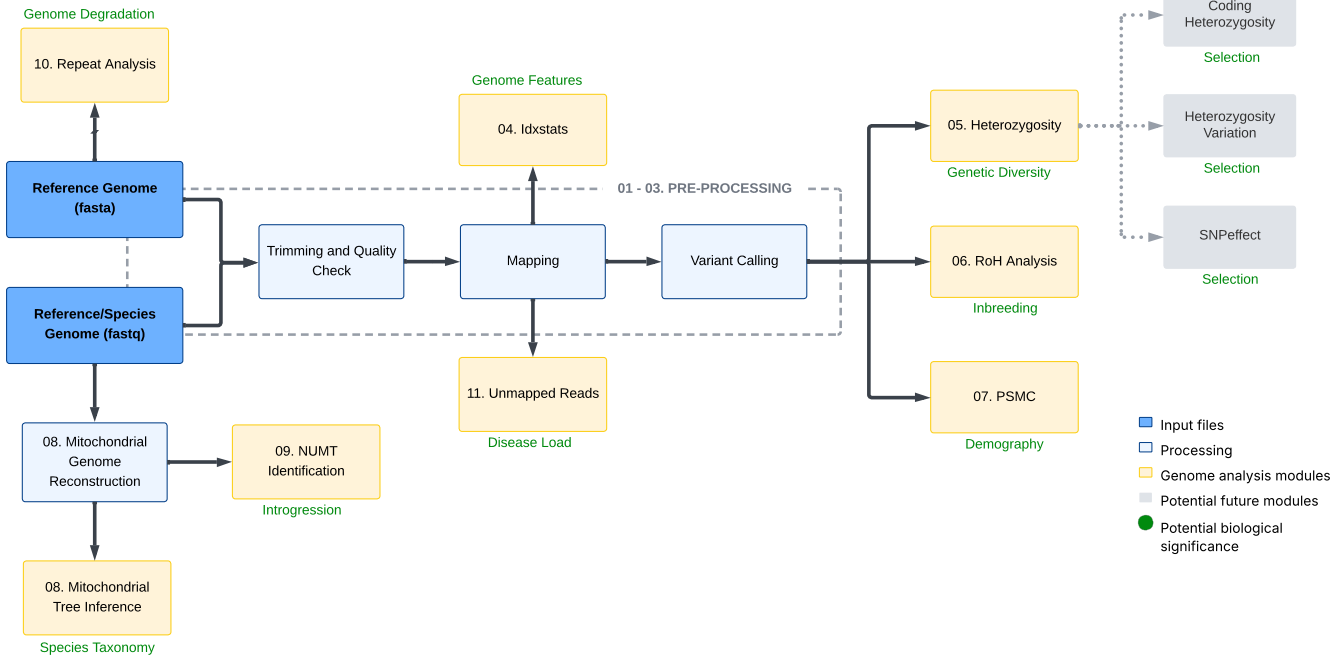


FIGURE 1 | Overview of the SwarmGenomics workflow. SwarmGenomics pipeline overview starting with reference and species genomes in FASTA and FASTQ format, followed by processing steps (in light blue) and modules for different genomic analyses (in yellow), with possible future modules in grey. Each module is numbered according to the corresponding page in the SwarmGenomics GitHub repository.

(e.g., FASTQ files, BAM files, or reference genomes), computational resources (threads, working directories), and tool-specific parameters. Each major analysis step such as preprocessing and the following genomic modules can be executed using its corresponding bash script. Scripts automatically read parameters from `params.txt`, enabling the same configuration file to be reused across all steps and ensuring consistency. This approach minimises user intervention, improves reproducibility and produces reliable outputs, allowing both novice and experienced users to efficiently run complex genomic analyses.

2.2 | Preprocessing

(Virtual) Machine	Data, software and time requirements
8–20 (V)CPUs	Input: reference genome (FASTA format) and reads from sequenced individual (FASTQ format)
+16 GB RAM	Installations: Miniconda3, FastQC, bwa-mem2, samtools, trimmomatic, SRA Tools
300–400 GB storage	Time: < 5 h

2.2.1 | Input Data

SwarmGenomics is designed to leverage the diploid genome of a single individual in combination with a reference genome, to

infer population genetic characteristics at the species or population level. The input genomic data (i.e., FASTQ files) may or may not originate from the same individual as the reference genome. It is important that datasets are clearly labelled to indicate whether they represent a reference individual or another sample from the same or a related species. For most modules, two genomic inputs are required: (1) a reference genome, ideally high-quality and species-specific, and (2) raw sequencing reads (in FASTQ or SRA format) from a single diploid individual. These can be derived from the same individual, although this is not strictly necessary. SwarmGenomics can be used to assess how well the input data match, and whether the individual aligns closely with the selected reference genome.

The reference genome is typically a well-annotated, curated sequence that represents the consensus DNA of a species, whereas the individual genome reflects the unique genetic composition of the sequenced sample. Both genome types are accessible through public databases such as NCBI Genomes and the Sequence Read Archive (SRA) (NCBI 2025a, 2025b), as well as other repositories like ENA, DDBJ or species-specific consortia (O’Cathail et al. 2025; Tanizawa et al. 2023; Zoonomia Consortium 2020; Smith et al. 2023). It is recommended to use a reference genome from RefSeq, as these assemblies are curated and annotated for accuracy (O’Leary et al. 2016). For the sequencing reads, whole-genome, paired-end Illumina short reads from a representative individual of the species (or a closely related species) are required. At present, SwarmGenomics is designed and validated exclusively for short-read Illumina data, as it relies on software tools optimised for this sequencing type. Third-generation long-read data from PacBio or Oxford Nanopore (Eid et al. 2009; Jain et al. 2016) are therefore not recommended as input for the current version of the pipeline.

TABLE 1 | Genome analysis modules in SwarmGenomics. Overview of each module's description, required software, input, output format and source code availability.

Module	Description	Software	Input	Output	GitHub
Genome features	Scaffold length, proportion of unmapped reads and coverage per chromosome	Samtools	BAM	A PNG file of bar plots	04. Genome features
Heterozygosity	Measuring genetic diversity within the population	Bcftools, Python, Pandas and Matplotlib	VCF	A text file and a barplot of heterozygosity in each chromosome	05. Heterozygosity
Runs of Homozygosity (RoH)	Detecting regions with identical haplotypes indicating inbreeding	Bcftools, Python, Pandas, Matplotlib, Numpy and Seaborn	VCF	Bar plots of the number and total length of RoH size classes	06. Runs of Homozygosity
Genome visualisation	Visualising genomic information on the 20 largest scaffolds	Circos	BAM, VCF and Ref	A PNG Circos plot	12. Genome Visualisation
PSMC	Estimating historical population sizes	PSMC and ms-PSMC	VCF	A PDF and PNG file of the PSMC graph	07. PSMC
Mitogenome assembly	Assembling the mitochondrial genome from WGS data	GetOrganelle	FASTQ	A FASTA file of the mitochondrial sequence	08. Mitochondrial Genome Reconstruction
NUMT identification	Detecting nuclear mitochondrial DNA segments (NUMTs)	R, Blast and Bedtools	Ref and mitogenome	A FASTA file with identified NUMTs	09. NUMT Identification
Repeat analysis	Identifying and classifying repetitive elements in the genome	HMMER, RM BLAST, TRF, RepeatMasker, RECON, Ninja, LTR Retriever, Mafft, CD-Hit, Genometools, RepeatModeler2 and RepeatScout.	Ref	An HTML with repeat landscape, genome fraction pie chart and output files compiling all data	10. Repeat Analysis
Unmapped reads	Identifying and analysing reads that do not map to the reference genome	Kraken2 and KrakenTools	BAM	An HTML for visualisation and text files with read classifications and counts per taxon	11. Unmapped Reads

2.2.2 | Data Preprocessing

Prior to downstream genetic analyses, raw sequencing data must be pre-processed and assembled. Genome assembly is a foundational step that reconstructs the genome from sequencing reads and enables the identification of variants and genomic features. While reference-based genome assembly can be complex, SwarmGenomics simplifies this process through a guided, step-by-step protocol and integrated preprocessing scripts.

Instructions for installing the required software are provided on the SwarmGenomics GitHub page (github.com/AureKylmanen/Swarmgenomics). For users starting with SRA files, these are first converted to FASTQ format, the standard format for storing raw sequencing reads and their associated quality scores (Cock et al. 2010). Quality control is performed using FastQC (Andrews 2010), which assesses read quality, base composition, and detects overrepresented sequences. Trimming is carried out with Trimmomatic (Bolger et al. 2014) to remove low-quality bases and adapter contamination. A second FastQC check is recommended to verify improvements in read quality.

2.2.3 | Read Mapping

Subsequently, reads are aligned to the reference genome using BWA (Li and Durbin 2009). The reference genome must first be indexed with `bwa index`. Paired-end reads are then aligned using `bwa mem`, producing a SAM file. This file is sorted and converted to BAM format using `samtools sort`, and then indexed with `samtools index` (Li et al. 2009). The reference genome is also indexed with `samtools faidx` to facilitate random access during variant calling.

2.2.4 | Variant Calling

(Virtual) Machine	Data, software and time requirements
16–28 (V)CPUs	Input: reference genome (FASTA format) and reads from sequenced individual (FASTQ format) or BAM file
16–64 GB RAM	Installations: Miniconda3, VCFtools, BEDtools, tabix, samtools, BCFtools
100–150 GB storage	Time: 3–10 h

Variant calling is a fundamental step in identifying genomic differences between an individual's genome and the reference. In SwarmGenomics, this is performed using `bcftools mpileup`, a lightweight and widely used tool that efficiently computes genotype likelihoods and generates a VCF file containing base-level information and variant calls across the genome (Danecek et al. 2011). BCFtools is used for its speed, simplicity and compatibility with standard short-read alignments, making it a practical default option for most users. The resulting VCF file serves as a key input for multiple downstream modules within the pipeline.

For users interested in more sophisticated or fine-tuned variant calling approaches, several alternative tools are available. FreeBayes is a haplotype-based caller that supports multi-sample joint calling and better resolves complex variants, although it typically requires more parameter tuning. GATK HaplotypeCaller is widely regarded for its robustness and accuracy, especially in well-annotated genomes, but comes with more stringent dependencies and a steeper learning curve. ANGSD (Analysis of Next Generation Sequencing Data) is particularly useful for population-level analyses and low-coverage data, as it focuses on estimating genotype likelihoods and allele frequencies rather than making discrete genotype calls (Korneliussen et al. 2014).

While SwarmGenomics uses BCFtools for its balance of speed and usability, users with more specific needs or higher data quality may benefit from exploring these alternatives. We encourage advanced users to consider other variant callers if their analyses demand increased sensitivity, support for population-scale datasets, or improved detection of structural variants.

SwarmGenomics supports workflows that begin either from raw FASTQ files or from pre-aligned BAM files. In cases where BAM files are used, it is important that the file represents only a single diploid individual. This is especially relevant when FASTQ data come from multiple sequencing runs, each labelled with a distinct read group (RG tag). Users must ensure that all read groups correspond to the same biological individual. Details about handling multiple read groups and ensuring proper sample-level variant calling are provided in the variant calling script section on the GitHub page.

Before proceeding to downstream analyses, repetitive genomic regions are filtered from the VCF file. This is done by generating a BED file of repeat elements based on the reference genome, which is then used to mask these regions. Filtering out repeats helps to avoid false variant calls and leads to more accurate estimates of heterozygosity and other metrics derived from the variant data.

```
perl -lne 'if(/^(>.*)/){ $head=$1 } else { $fa{$head} .= $_ } END{ foreach $s (sort(keys(%fa))){ print "$s\n$fa{$s}\n" }}' reference.fna |
perl -lne 'if(/^>(\S+)/){ $n=$1 } else { while(/[a-z]+)/g){ printf("%s\t%d\t%d\n", $n, pos($_) ) } printf("%s\t%d\t%d\n", $n, -length($1), pos($_) ) }' > repeats.bed
```

The `bedtools subtract -header` command is then used to remove the identified repeats from the VCF file. This file serves as the main input for many of the modules and analyses below. Lastly, only the largest chromosomes or scaffolds can be extracted to save time during analyses. In this example, scaffolds larger than 5 MB are extracted, though the size can be adjusted according to specific requirements.

```
$ cat reference.fna | awk '$0 ~ ">" {if (NR >1) {print c;} c=0;printf substr($0,2) "\t"; } $0 !~ ">" {c+=length($0);} END
```

```
{ print c; }' | awk '$(NF) >=5000000 {print $1}' > target_len.lst
```

This produces a file named `target_len.lst` that contains the names of sequences in the reference FASTA file that are at least 5,000,000-bp long. Afterwards, these specific chromosomes are selected using `tabix`, filtered based on depth and allele count using `vcftools`, and then compressed and indexed with `bgzip` and `bcftools`. Consensus sequences in FASTQ format for each chromosome are then generated using a reference genome and the filtered VCF files, and combined into a single diploid FASTQ file, which will be used for PSMC, and the separate VCF files are used for heterozygosity and runs of homozygosity analyses. Overall the preprocessing step takes around 8–20 h depending on the size of the genome and SRA files.

2.3 | List of Modules for Genome Analysis

The modules used for genome analysis are summarized in Table 1.

3 | Genome Analysis

3.1 | Genome Features

(Virtual) Machine	Data, software and time requirements
1 (V)CPU	Input: BAM file (or CRAM file + reference genome)
4 GB RAM	Installations: Samtools, R, ggplot2, dplyr
10 GB disk	Time: <1 min

When analysing genomic data, an initial overview of the mapping results is crucial for informing subsequent analyses. `Samtools idxstats` (Li et al. 2009) offers a fast and efficient method for summarising alignment statistics, providing chromosome- or scaffold-level details such as sequence length, the number of mapped reads, and the number of unmapped reads. These summary metrics can highlight patterns in sequencing coverage and mapping efficiency, and may also reveal anomalies, such as poorly assembled chromosomes or scaffolds with coverage patterns indicative of heterogametic chromosomes. In species with heteromorphic sex chromosomes (e.g., mammals), sequencing coverage can also be used to infer the biological sex of individuals by comparing read depth across sex chromosomes. Typically, males exhibit reduced or no coverage on the Y chromosome, while females lack it entirely.

These statistics can be generated using the following `samtools idxstats` command on a BAM file:

```
samtools idxstats file.bam > idxstats_output.txt
```

To better visualise and interpret this information, a bar plot of the 20 longest scaffolds or chromosomes (or fewer if fewer are available) is generated to highlight genomic scale. A normalised measure of unmapped reads per megabase is calculated to account for scaffold size, and an estimate of coverage per scaffold is derived based on the number of mapped reads and an assumed read length. Additionally, the total number of scaffolds and the number of scaffolds larger than 5 Mb can be obtained from the output using simple Bash commands.

The giant panda BAM file contains 22 scaffolds larger than 5 Mb. As shown in Figure 2A–C, the graphs display the size distribution of these scaffolds, the number of unmapped reads per Mbp, and the estimated coverage based on a read length of 150 bp, determined as the most frequent length among the first 1000 alignments. The overall genome coverage is approximately 50× and remains consistent across the scaffolds.

3.2 | Heterozygosity

(Virtual) Machine	Data, software and time requirements
16–28 (V)CPUs	Input: VCF file
32–64 GB RAM	Installations: Bcftools, R, ggplot2, dplyr
200–300 GB storage	Time: <2 h

Understanding genetic diversity within a population is crucial for addressing the complexities of evolution, adaptation, and conservation. Heterozygosity, one of the most significant indicators of genetic diversity, measures the presence of different alleles at a gene locus within an individual's genome (Nei 1973). High heterozygosity indicates healthy genetic variation, allowing populations to adapt to environmental changes and avoid inbreeding depression (Reed and Frankham 2003). Conversely, low heterozygosity may suggest inbreeding, genetic drift, or population bottlenecks, which can increase the risk of extinction (Charlesworth and Willis 2009). In the case of the giant panda, a species of conservation concern (IUCN 2023), analysing genetic diversity can provide insights into its potential to persist in a changing environment.

For estimating heterozygosity in the largest scaffolds, a bash script is used to count the number of heterozygous and homozygous genotypes, followed by calculating the proportion of heterozygous genotypes for each scaffold separately. The following Bash snippet estimates heterozygosity by counting heterozygous and homozygous genotypes from a VCF file:

```
heterozygosity=$(bcftools query -f ' [%GT\n]'
"$vcf_file" | \
awk 'BEGIN {het=0; hom=0} \
{if ($0 == "0/1" || $0 == "1/0") het+=1; \
else if ($0 == "0/0" || $0 == "1/1") hom+=1} \
END {if (het+hom > 0) print het/(het+hom); \
else print 0}')
```

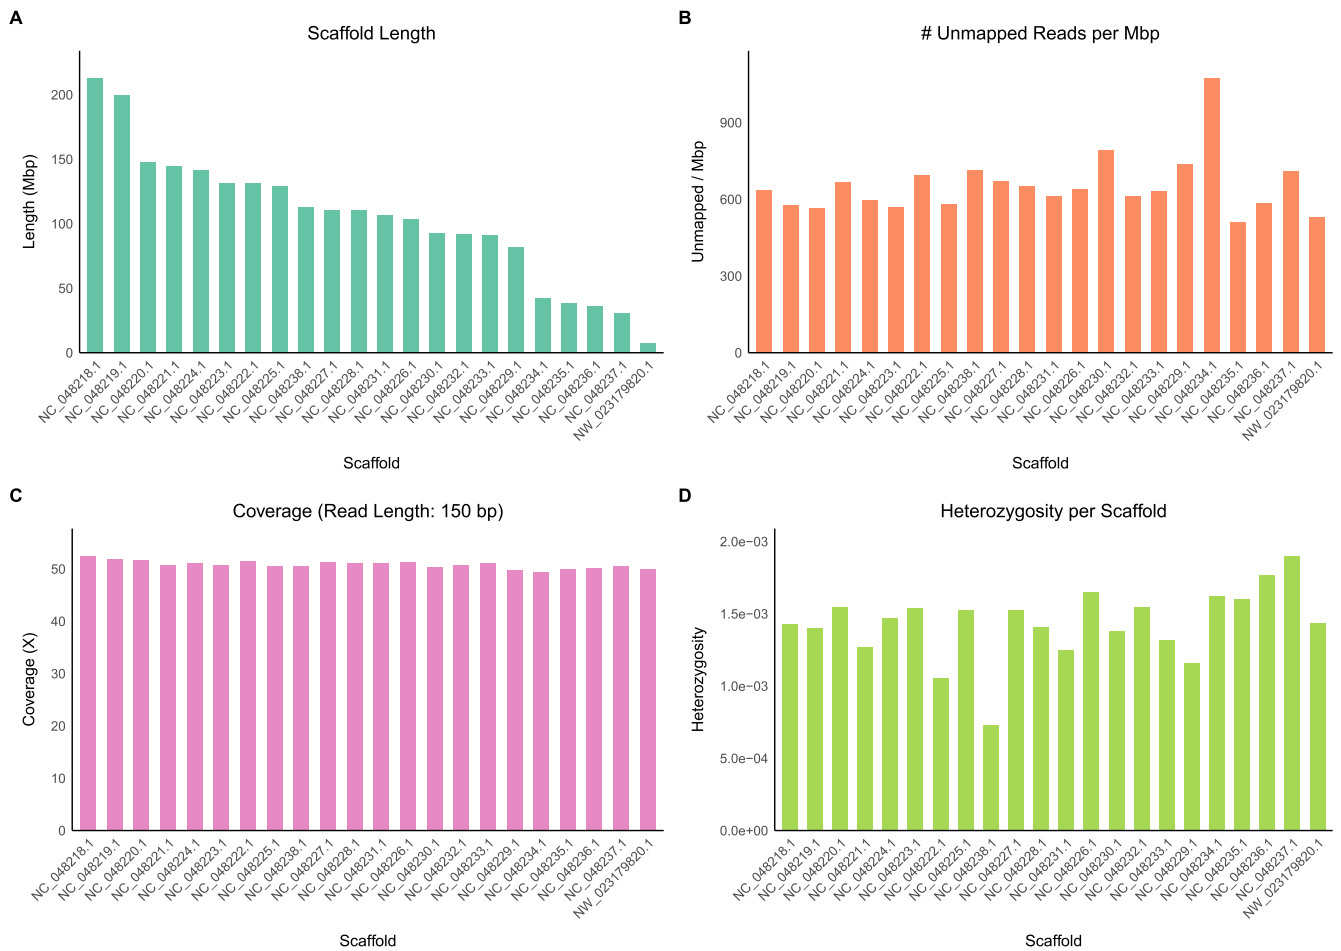


FIGURE 2 | Scaffold-level genome metrics of a single individual mapped to the giant panda genome. Shown are the 22 longest scaffolds/chromosomes. (A) Lengths in megabasepairs (Mbp) of each scaffold. (B) Number of unmapped reads (i.e., paired reads missing a mate) per megabase, indicating variation in mapping efficiency across scaffolds. (C) Estimated sequencing coverage (X) with 150bp read length, highlighting differences in sequencing depth among scaffolds (i.e., approximately half the coverage of the sex chromosome in an heterogametic individual). (D) Heterozygosity (heterozygous sites per Mbp) of the largest scaffolds of giant panda genome.

An R script can then be applied to plot the results as a bar plot, as shown in Figure 2D. Each bar represents a chromosome or a scaffold. In the case of the giant panda, heterozygosity per chromosome is moderate. However, scaffold NC_048238.1 exhibits much lower heterozygosity than the others.

3.2.1 | Individual-Level Interpretation of Heterozygosity

In SwarmGenomics, heterozygosity is estimated at the level of a single diploid individual by quantifying the proportion of heterozygous genotypes across the genome. This measure reflects the genetic divergence between the two homologous chromosomes carried by the individual and provides a direct summary of within-individual genetic diversity.

When calculated across large genomic regions or chromosomes, individual heterozygosity captures information about long-term effective population size and historical demographic processes affecting the lineage of the sampled individual. Because variant discovery is performed relative to a reference genome, heterozygosity estimates also reflect the relationship between the focal individual

and the chosen reference. While this does not yield population allele frequencies in the classical sense, it provides a biologically meaningful and comparable statistic that can be analysed across individuals processed using the same reference genome.

3.2.2 | Runs of Homozygosity

(Virtual) Machine	Data, software and time requirements
16–28 (V)CPUs	Input: VCF file
32–64 GB RAM	Installations: Bcftools, R; ggplot2, dplyr, gridExtra, tidyR
200–300 GB storage	Time: < 1 h

Similar to heterozygosity, runs of homozygosity (RoH) measure genetic diversity within populations, providing insights into genomic regions where individuals inherit identical haplotypes from both parents (Broman and Weber 1999). RoHs indicate

recent common ancestry or inbreeding events; therefore, the number and length of RoH can inform the demographic history of the species (Ceballos et al. 2018). High frequency and long length of RoH may suggest reduced genetic diversity and increased inbreeding depression, potentially impacting the species' ability to adapt and survive in changing habitats (Abascal et al. 2016; Ceballos et al. 2018). Conversely, shorter or fewer RoH segments reflect greater genetic heterozygosity and potential resilience to environmental stressors. For species like the giant panda, which faces conservation challenges due to fragmented populations and vulnerability to environmental changes, RoH analysis helps assess the genetic health of the population and its resilience to environmental change.

To analyse RoH in the genome, the BCFtools roh command is used. As with heterozygosity analysis, the input file should be a VCF file with FORMAT columns containing either genotype likelihoods (PL) or genotypes (GT). By default, the command uses genotype likelihoods unless otherwise specified with the `-G` option. The command also requires an estimate of the alternate allele frequency in the population for each site; here, the default alternate allele frequency value of 0.4 is used with `-AF-dflt 0.4`.

```
$ bcftools roh -G30 --AF-dflt 0.4 file.vcf
```

For visualising the results, an R script is used to create bar plots of the number of RoHs and the total length of RoHs of four size categories; 10kbp–0.1, 0.1–1, 1–3, and >3Mbp, as shown in Figure 3. Short ROHs (<0.1Mbp) result from ancient recombination and are common across populations. Short-to-medium ROHs (0.1–1 Mbp) indicate background relatedness within populations, while medium-to-long ROHs (1–3Mbp) suggest recent shared ancestry. Long ROHs (>3Mbp) arise from recent inbreeding, where parents share a close common ancestor.

To illustrate how RoH patterns can differ between species, the distribution in the giant panda genome, characterised by relatively few short RoHs, is shown in Figure 3A,B. In contrast, the RoH profile of the lion genome, which exhibits a greater number and total length of RoHs, is shown in Figure 3C,D. These visualisations highlight distinct genomic histories and demographic influences.

3.3 | Genome Visualisation

(Virtual) Machine	Data, software and time requirements
28 (V)CPUs	Input: reference genome, BAM file and VCF file
64 GB RAM	Installations: https://circos.ca/ bedtools, bcftools, mosdepth
50 GB Storage	Time: < 5 h

Circos plots are widely used in genomics to represent large-scale data in a circular format, allowing for the clear visualisation

of patterns and connections across entire genomes. This approach is particularly useful for displaying multiple layers of genomic information in a compact and interpretable way. For the SwarmGenomics pipeline, a Circos plot is generated to illustrate key genomic features, including coverage depth, GC content, repeat density, and heterozygosity. These metrics were chosen to provide insights into genome structure and variability. Coverage depth reflects sequencing uniformity, GC content describes sequence composition, repeat density highlights repetitive elements, and heterozygosity indicates genetic diversity.

All metrics are calculated in non-overlapping 1-Mb bins across the genome. Coverage is extracted from BAM files using mosdepth (Pedersen and Quinlan 2018), and repeat density is computed by intersecting repeat-masked regions with predefined genomic bins. GC content is calculated with bedtools nuc (Quinlan 2014), while heterozygosity is determined by using bcftools view `-g het` to filter variant calls from the VCF file and counting overlapping variants in each genomic bin.

These results are visualised in a Circos plot with the 20 largest scaffolds of the giant panda genome in Figure 4. GC content and genomic coverage are represented as line plots (purple and red, respectively), while heterozygosity and repeat density are shown as green and blue heatmaps.

3.4 | PSMC

(Virtual) Machine	Data, software and time requirements
	Input: VCF file
16–28 (V)CPUs	Installations: PSMC: https://github.com/lh3/psmc.git
32–64 GB RAM	
100–200 GB storage	Time: ~ 8 h

The Pairwise Sequentially Markovian Coalescent (PSMC) method provides insights into the historical population sizes of a species, extending back up to millions of years. By analysing patterns of genetic variation in a single genome, PSMC estimates historical population sizes based on coalescence events. This method is most accurate for periods ranging from 20,000 to 3 million years ago (Li and Durbin 2011; Palkopoulou et al. 2015). Estimating recent population sizes can be challenging due to a lack of mutations, while ancient periods are difficult to resolve because of the accumulation of mutations, which introduces noise into the genetic data. By applying PSMC to giant pandas, changes in their population size due to past climatic events, such as the last Ice Age, can be investigated.

PSMC was run following the official guidelines (<https://github.com/lh3/psmc>) on the diploid consensus sequence diploid.fq.gz generated previously. Unless stated otherwise, default settings were used. The `-N` option sets the number of EM iterations (typically 25). The `-t` option specifies the maximum coalescent time

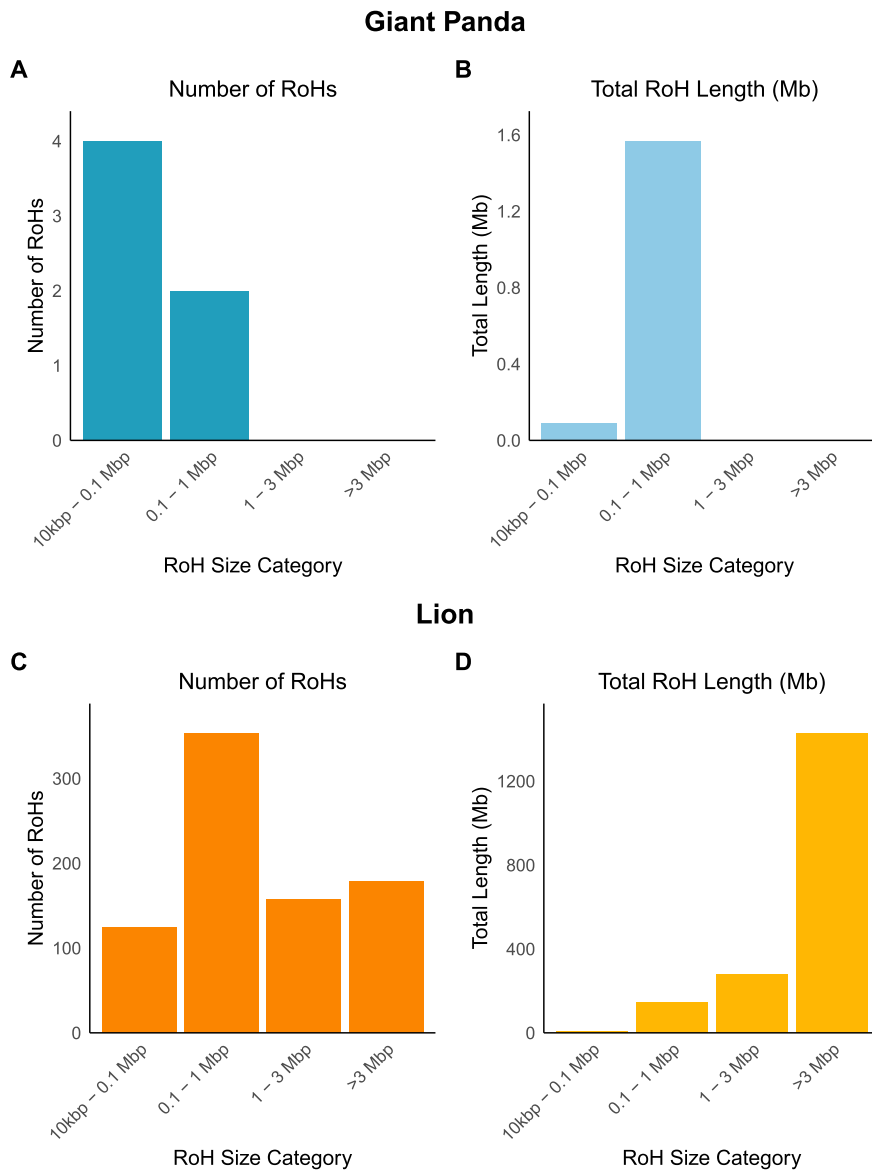


FIGURE 3 | Distribution of runs of homozygosity (RoHs) by size in two species. Bar plots of the number of RoHs and the total length of RoHs across four size categories: 10 kbp–0.1 Mbp, 0.1–1 Mbp, 1–3 Mbp, and > 3 Mbp. Graphs A and B show data for the giant panda, while graphs C and D show data for the lion. The left plots (A and C) show the count of RoHs in each category, while the right plots (B and D) display the corresponding total genomic length in megabases (Mb).

(T_{\max} , in units of $2N_0$), commonly set to 15. The `-r` option provides the initial estimate of the mutation-to-recombination rate ratio (θ/ρ) for the HMM/EM optimization; `-r 5` is a common starting value. The `-p` option defines the pattern of time intervals (free parameters) used to model the TMRCA distribution; a widely used pattern is `'4+25×2+4+6'`. However, this default may occasionally induce spurious recent peaks; splitting the first window (e.g., `'2+2+25×2+4+6'`) mitigates this artefact (Hilgers et al. 2025).

```
$ utils/fq2psmcfa -q20 diploid.fq.gz > dip-
loid.psmcfa
psmc -N25 -t15 -r5 -p "2+2+25*2+4+6" -o dip-
loid.psmc diploid.psmcfa
utils/psmc2history.pl diploid.psmc | utils/
history2ms.pl > ms-cmd.sh
```

```
utils/psmc_plot.pl diploid diploid.psmc
```

`psmc_plot.pl` from the PSMC program is used for plotting the results, it produces a pdf file with the graph. The script uses the `diploid.psmc` file, created during PSMC as the input file. To scale the graph according to the species, the user can define generation time (g) and mutation rate (μ). Table 2 gives few examples of mutation rates and generation times for different species, which can be used to guide the parameter selection.

```
$ utils/psmc_plot.pl diploid diploid.psmc
```

Scaled:

```
$ utils/psmc_plot.pl -u 1.29e-8 -g 12 diploid
diploid.psmc
```

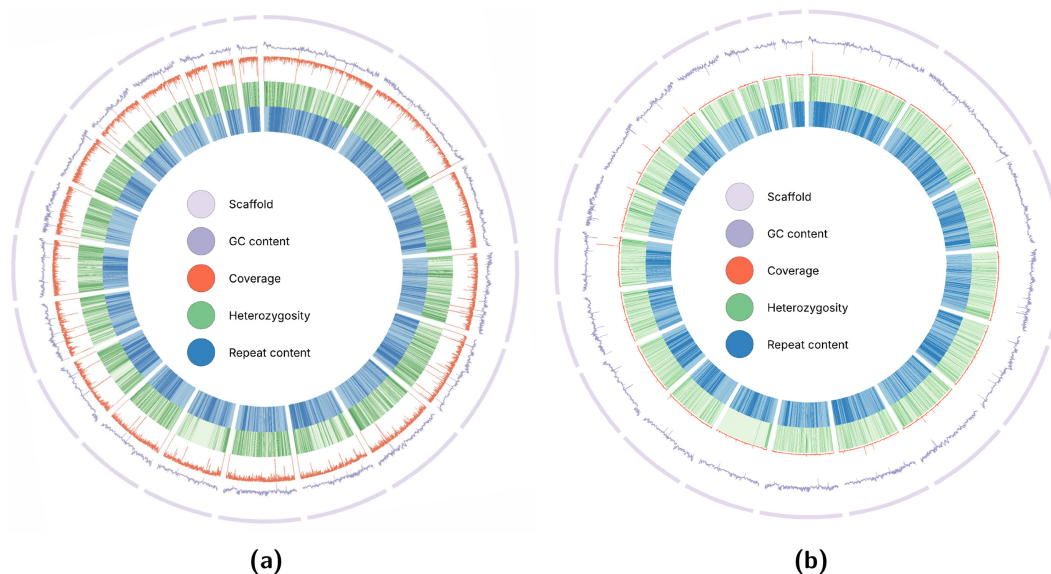


FIGURE 4 | Circos plot displaying genomic features in the 20 largest scaffolds of the giant panda genome: GC content as a purple line, genomic coverage as a red line, heterozygosity as a green heatmap, and repeat density as a blue heatmap. The heterozygosity heatmap ranges from [0–0.002]. (a) High coverage genome (50×). (b) Low coverage genome (4×).

Figure 5 shows the output of `psmc_plot.pl`. The *x*-axis represents time, with higher values further along the axis, while the *y*-axis is effective population size. The graph shows a decline in population size at 1000,000 to 500,000 years ago. This period occurs during the Pleistocene Epoch, which was characterised by repeated glaciations, or ice ages. This could have led to habitat loss or fragmentation reducing giant panda population sizes.

3.4.1 | Mitochondrial Genome Reconstruction

(Virtual) Machine	Data, Software and Time Requirements
16–28 (V)CPUs	Input: FASTQ files
32–64 GB RAM	Installations: GetOrganelle
100 GB storage	Time: < 10 h

The study of the mitochondrial genome (mitogenome) is of particular interest due to its unique properties and wide-ranging applications in genetic research, including genome-wide, epigenome-wide, and phenome-wide association studies (Ferreira and Rodriguez 2024). The mitochondrial genome's high mutation rate, heteroplasmy, copy number variation, maternal inheritance, and lack of recombination make the mitogenome especially valuable in studies of human disease, population structure, evolutionary biology and DNA barcoding (Lin et al. 2022; Wallace 2010; Song et al. 2022; Jin et al. 2020).


Advances in next-generation sequencing (NGS) technologies have driven the development of numerous specialised tools for organellar genome assembly, particularly for chloroplast and

mitochondrial genomes (Dierckxsens et al. 2017). Several assembly tools are described such as GetOrganelle (Jin et al. 2020), NOVOPlasty (Dierckxsens et al. 2017), MitoZ (Meng et al. 2019), MITObim (Hahn et al. 2013), SPAdes (Bankevich et al. 2012), MitoFinder (Allio et al. 2020), and Megahit (Li et al. 2016), but their assembly qualities vary widely. In a systematic comparison of human mitochondrial genome assembly tools, MitoFlex was the fastest, NOVOPlasty the most memory-efficient for simulated data, GetOrganelle and MitoFlex were best for SNP detection in human WGS data, and MToolBox and NOVOPlasty performed consistently across all sequencing depths (Mahar et al. 2023). Nevertheless, GetOrganelle's ability to accurately assemble genomes despite complex repeats and highly conserved regions—common challenges in genome assembly—makes it a preferred choice for researchers who prioritise accuracy and comprehensive analysis in mitochondrial genome assembly (Jin et al. 2020).

In this pipeline, GetOrganelle is used to assemble the mitogenome from previously generated FASTQ files using the script `'get_organelle_from_reads.py'` with a single line command. This script exploits Bowtie2, BLAST, and SPAdes, as well as the Python libraries NumPy, SciPy, and SymPy as dependencies to filter, align and assemble reads into contiguous sequences (contigs) representing the organelle genome. The process includes error correction, k-mer analysis, and graph-based assembly. Consequently, the output of GetOrganelle includes the assembled mitochondrial genome in FASTA format, along with detailed assembly statistics and quality metrics. These outputs enable users to perform further analyses, such as annotation of genes, comparative genomics, and evolutionary studies. These results are saved in `DESTINATION_PATH/mt/`, including the assembled mitogenome in FASTA format (`*_mitogenome.fasta`), the organelle-only assembly graph (`*.selected_graph.gfa`), log files, and various quality metrics. These outputs are suitable for downstream analyses such as gene annotation, phylogenetic

TABLE 2 | Mutation rates and generation times for various species used in PSMC analyses, as described in the literature.

Species	Mutation rate (μ)	Generation time (years)	References
Human (<i>Homo sapiens</i>)	1.25×10^{-8}	25	(Li and Durbin 2011)
Mouse (<i>Mus musculus</i>)	5.4×10^{-9}	0.5	(Fujiwara et al. 2022)
Dog (<i>Canis lupus familiaris</i>)	1×10^{-8}	3	(Freedman et al. 2014)
Great tit (<i>Parus major</i>)	2.0×10^{-9}	2	(Laine, Gossmann, Schachtschneider, et al. 2016)
Chicken (<i>Gallus domesticus</i>)	1.91×10^{-9}	1	(Wang et al. 2021)

**FIGURE 5** | PSMC plot of giant panda populations: Effective population size ($\times 10^4$) is shown on the y-axis, and time (in years) is shown on the x-axis. Parameters used are generation time $g = 12$ and mutation rate $\mu = 1.29 \times 10^{-8}$.

reconstruction, comparative genomics, and population studies (Jin et al. 2020).

```
$ get_organelle_config.py -a animal_mt
get_organelle_from_reads.py -l DESTINATION_PATH/*_1_paired.fastq.gz \
-2 DESTINATION_PATH/*_2_paired.fastq.gz -F animal_mt \
-o DESTINATION_PATH/mt/*_mitogenome -R 10 -t 26 -s
```

To place the reconstructed mitogenome in a phylogenetic context, a reference database of mitochondrial genome sequences was compiled and formatted using makeblastdb (Camacho et al. 2009). The query sequence was renamed for clarity and searched against this database using BLASTn (Altschul et al. 1990). The top 10 unique high-scoring matches were selected, combined with the query, and aligned using MAFFT (Katoh and Standley 2013). A maximum-likelihood phylogenetic tree was then inferred with FastTree based on the resulting alignment (Price et al. 2010). Full FASTA headers were used as tip labels, with the query sequence highlighted in red as shown in Figure 6.

Overall, GetOrganelle represents a significant advancement in the field of organelle genomics, providing researchers with a powerful tool to unlock the genetic secrets of mitochondria and other organelles. Its ability to deliver high-quality assemblies from short-read data makes it an invaluable resource for exploring the genetic and functional diversity of these critical cellular components (Jin et al. 2020).

3.4.2 | NUMT Identification

(Virtual) Machine	Data, Software and Time Requirements
16–28 (V)CPUs	Input: Mitogenome FASTA file
16–32 GB RAM	Installations: R, Blast, Bedtools
50–100 GB storage	Time: < 1 h

In this module, NUMT fragments harboured in nuclear genomes are identified. NUMT stands for nuclear mitochondrial DNA, which describes the mitochondrial DNA fragments detected in the nuclear genomes (Lopez et al. 1994). This continuous, ongoing and non-random (Wei et al. 2022; Tsuji et al. 2012) integration of mitochondrial DNA to the nuclear genome is identified across species and tissues (Hazkani-Covo 2022; Dayama et al. 2020; Liang et al. 2018; Calabrese et al. 2017). The NUMT insertion numbers and sizes vary not only between species but also within populations. In addition, these fragments can introgress from closely related species through hybridization or interbreeding (Bücking et al. 2019; Lammers et al. 2017; Pérez et al. 2017; Baldo et al. 2010), which further fuels the genetic variations in one population. Interestingly, after integration, the NUMT fragments can be preserved in an ancestral state and thus be used for evolutionary event identification and as phylogenetic markers (Uvizl et al. 2024; Lucas et al. 2022; Zhang et al. 2021; Miraldo et al. 2012; Soto-Calderón et al. 2014). This shows that NUMTs help unravel the evolutionary and population history. Furthermore, this mutagenic event could be under selection (Vendrami et al. 2022), hinting at the potential function of NUMT.

BLASTN is used to identify NUMTs in the nuclear genome with mitochondrial DNA sequences as queries. First, the user has to build the database for BLASTN with nuclear genomes.

```
$ makeblastdb -in nuclear_genome -parse_seqs -dbtype nucl
```

After building the database, the mitochondrial DNA sequences are blasted with the nuclear genome.

```
$ blastn -db nuclear_genome -query mito_seqs -outfmt 7 -word_size 20 -out blast_out
```

To generate the BLAST results in tabular output with comment lines, which facilitates the creation of the bed file for the next

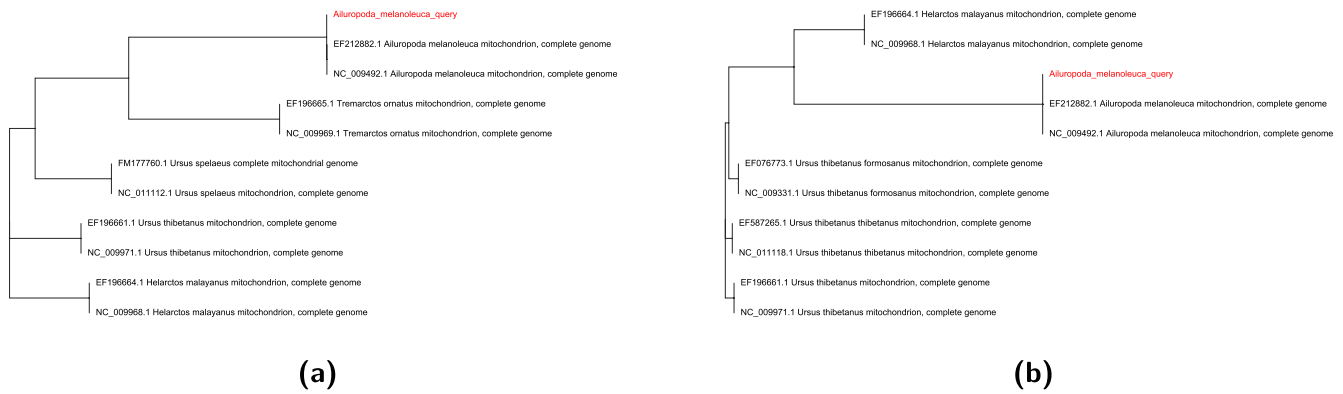


FIGURE 6 | Phylogenetic trees of mitochondrial genomes. (a) Reconstructed giant panda mitochondrial genome highlighted in red among related mitochondrial sequences based on sequence similarity. (b) Reference giant panda mitochondrial genome highlighted in red among related mitochondrial sequences based on sequence similarity.

step, the `-word_size` parameter is set to 20. This setting, used in previous research (Lammers et al. 2017; Vendrami et al. 2022), helps identify NUMTs. It should be noted that the `-word_size` parameter can influence the sensitivity and specificity of the BLAST search.

The information is then extracted from the BLAST result using the following command:

```
$ grep -v '#' blast_out | cut -d '\t'
-f2,3,4,9,10 > blast_clean_out
```

Next, BED files are generated using the provided script `get_bed2.R`. This script not only creates the bed file from the BLAST results but also filters out fragments smaller than 200 bp, to reduce noise in the results.

Finally, `bedtools` is used to generate FASTA files of NUMT sequences from the BED file and the genome file:

```
$ bedtools getfasta -fi nuclear_genome -bed
blast_clean_out.bed -fo numt.fasta
```

The output file is a FASTA file containing NUMT sequences. Each sequence header includes the scaffold or chromosome name followed by the NUMT position on that scaffold or chromosome. An example pipeline script, `runNUMTidentification2.3.cyc.py`, is available on the GitHub page. The number of NUMTs identified in this genome can be calculated and compared with other species or individuals (Figure 7). The results can be visualised in a circos plot showing the organisation of the mitochondrial genes (mtDNA) and NUMTs aligned in the reassembled mitogenome (Figure 7a) and in the reference mitogenome (Figure 7b). Fifty NUMTs were identified in the giant panda, which is consistent with findings in other mammalian species (Chen et al. 2025).

NUMT insertion is an ongoing process that provides valuable information about evolutionary events such as hybridization and interbreeding. After identifying NUMTs, the frequency and extent of this mutagenic event can be examined. The results can also be used for further analyses with other individuals or species, serving as markers for phylogenetic and population genetic studies.

3.4.3 | Repeat Analysis

(Virtual) Machine	Data, software and time requirements
12+ (V)CPUs	Input: Reference genome FASTA file
32–64 GB RAM	Installations: HMMER, RM BLAST, TRF, RepeatMasker, RECON, Ninja, LTR Retriever, Mafft, CD-Hit, Genometools, RepeatModeler2, RepeatScout
300–500 GB storage	Time: 24–48 h

Transposable elements (TEs) are ubiquitous genomic components and sources of evolutionary novelty across the eukaryotic tree of life. Indeed, TEs commonly account for the majority of genomic material in both plants and animals, making the study of these self-replicating ‘jumping genes’ an essential component of genomic analyses (Almojil et al. 2021; Slotkin 2018). As TEs evolve within lineages and can ‘jump’ across them, unique TEs often appear within specific clades and can be unexpectedly encountered in any clade. This makes de novo repeat identification necessary for detection of taxon-specific and horizontally transferred TEs (Rodriguez and Arkhipova 2022). Annotations of TEs may be of interest for studies on genomic architecture, host-parasite interactions within the genome, and sources of new or laterally transferred genomic material. TE annotations are also commonly considered in population genomic analyses in many species: as TEs commonly result in highly repetitive DNA motifs that interfere with alignment quality and estimates of genomic diversity (Cechova 2020; Rodriguez and Arkhipova 2022), TE masking may improve many analyses involving small variants (Patil and Vijay 2021).

This module identifies de novo TE families (repeats) within a single genome using RepeatModeler2 (Flynn et al. 2020), and annotates a single reference genome for these families, simple repeats, and satellite DNA using RepeatMasker (Tarailo-Graovac and Chen 2009). De novo repeat families are identified

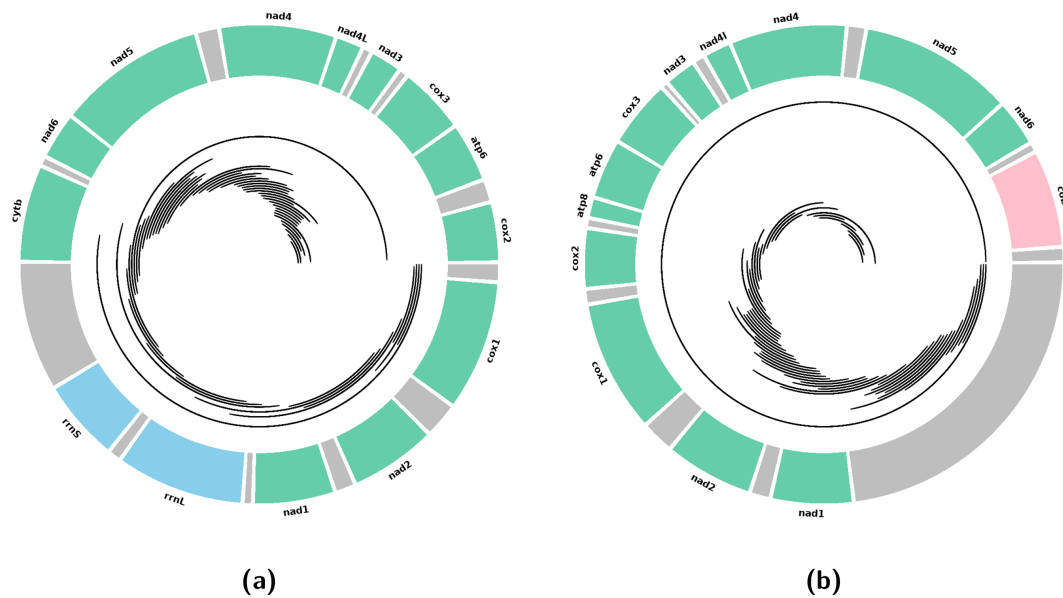


FIGURE 7 | Visualisation of NUMT alignments on the mitochondrial genome. The circos plots illustrate the organisation of mitochondrial DNA (mtDNA) genes and nuclear mitochondrial DNA segments (NUMTs) aligned to the mtDNA reference. The outermost track displays the annotated mitochondrial genome with genes colour-coded by category: protein-coding genes (green), ribosomal RNA (rRNA) genes (blue), and non-coding regions including introns (grey), other elements are in pink. The inner track shows NUMT alignments as black lines representing the start and end positions of mtDNA regions with homologous sequences in the nuclear genome. (a) shows the reassembled mitochondrial genome, while (b) depicts the reference mitochondrial genome.

through the occurrence of similar sequence motifs throughout the genome, eliminating the need for reference TE libraries. Reference sequences are, however, used to assign TE families to classes and subclasses whenever possible, but these are included as DFAM libraries with the RepeatMasker installation and do not require additional inputs from the user. The full script is available on the GitHub page.

Reference-based analysis can also be performed without de novo repeat annotation. This can be done very quickly and with minimal prerequisites. It is, however, most informative when studying narrow groups within well-annotated clades: Applying it to mammals ('mammalia') and birds ('aves') will provide fast and informative insights into the repeat landscapes of most species within these classes. Reference-based methods on poorly studied taxa may not yield meaningful results, and analyses with very broad categories of reference data (e.g., eukaryota) will likely exceed the run times required for de novo annotation.

Running this analysis requires only a single reference genome as input. The user may also wish to set the number of parallel threads to be used. This module will produce a FASTA file of de novo repeat families ('...-families.fasta') broken into subfamilies with greater than 80% identity to one another. It will also produce a file of repeat annotations on the reference genome ('...fna.out'), a summary table of repeats annotated within the reference genome ('...tbl'), an interactive summary graph of the TE landscape as the genomic proportion of repeats by divergence from consensus ('...html'), and a repeat masked reference genome for downstream analyses ('...fna.masked'). This will provide the user with the necessary files for a general analysis on TEs within an assembly, and allow for the exclusion of TEs from associated population genomic analyses as required.

The repeat landscape of the giant panda genome is visualised in Figure 8 with the fraction of different repeat families in the genome as a pie chart. Repeats make up 31% of the genome, with LINE families being most abundant. The y-axis of the repeat landscape shows the percentage of the genome that each repeat family makes, with the age of the repeats represented as the Kimura substitution level on the x-axis. In the giant panda genome, it shows a recent expansion of the LINE1 family, possibly indicating active repeats.

3.4.4 | Unmapped Reads Classification

(Virtual) Machine	Data, Software and Time Requirements
8–16 (V)CPUs	Input: FASTQ file
32–64 GB RAM	Installations: Kraken2, KrakenTools
~ 100 GB storage (depends on the database)	Time: < 1 h (but can be up to several days depending on the size of the database and the number of analysed unmapped reads)

WGS studies typically begin with assembling the genome of the organism of interest. The assembled genome then serves as a reference for subsequent whole-genome re-sequencing, in which sequenced reads from a newly generated sample are mapped to the reference genome. In such cases, a certain part of the produced reads cannot be mapped to the reference genome. These unmapped reads (UR) may result from methodological artefacts, sequencing errors, and they are typically ignored and excluded

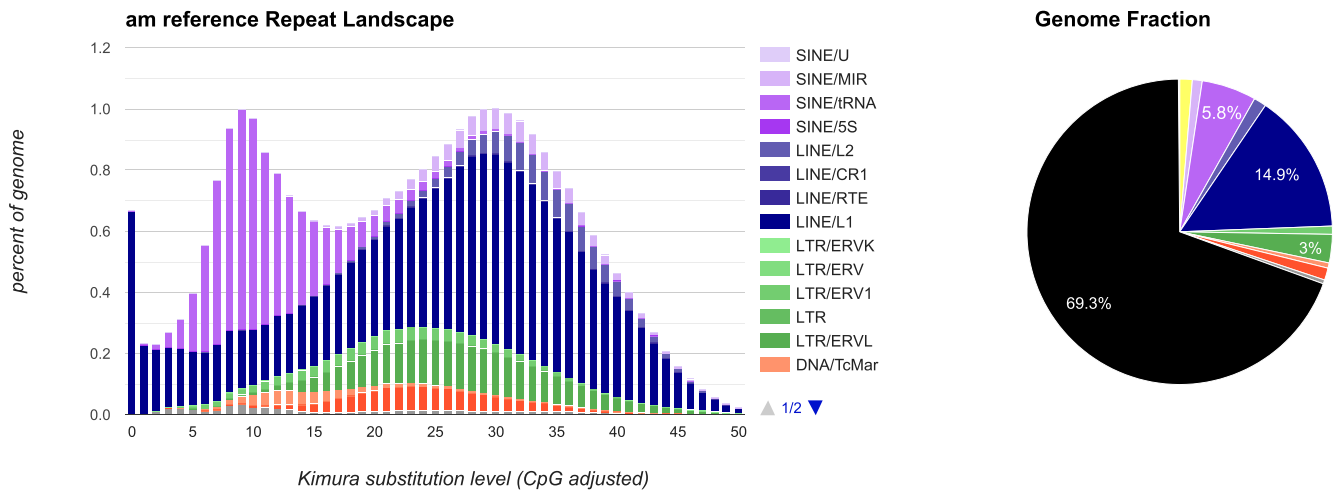


FIGURE 8 | Repeat landscape of the giant panda genome: The percentage of repeat families in the genome is shown on the y-axis, while the Kimura distance is plotted on the x-axis. A pie chart representing the percentage distribution of repeat families in the genome is also included.

from further analyses (Chrisman et al. 2022). However, many of them may still contain interesting biological information (Chrisman et al. 2022). For example, they were used to identify biotype-specific sequences in the aphid *Acyrtosiphon pisum* and its endosymbionts (Gouin et al. 2015), and to detect the blood parasites *Plasmodium* and *Trypanosoma* in RNA samples of the Great Tit *Parus major* (Laine et al. 2019).

An important source for UR may be microbes that are not contaminations but actually belong to the microbiome of the sampled organism. Bacteria are often important parts of the microbiome and may persist in body parts once considered sterile, such as eggs, blood, lung, and bladder (D'Aquila et al. 2021; O'Dwyer et al. 2016; Nyholm 2020; Trevelline et al. 2018). Viruses can also be studied by analysing UR (Moustafa et al. 2017). Such UR of viral origin might come from viruses that historically integrated into the host genome of a population, viruses that are integrated into the genome of only the sampled individual host, or viruses from an acute infection at the time point of sampling. Finally, genomic variation of the host can be a source of UR. For example, reads with complex structural variations, like large indels, inversions, and translocations may fail to map to the reference if reads are mapped to a reference genome that does not include such structural variations. This may happen when the reference individual and the sampled individual are from populations with divergent structural variants (Takayama et al. 2021; Wong et al. 2018, 2020). Altogether, there are multiple sources of variation in UR, which differ in their biological relevance, and their study can facilitate the understanding of genomic architecture and host-symbiont relationships.

In this module, Kraken 2 enables taxonomic classification of metagenomic reads (Wood et al. 2019). It compares read k -mers to a database of k -mers of known origin. For the analysis, Kraken 2 requires a FASTA file with the unmapped reads to be classified and the database to compare these reads to. For the following example, Kraken 2 and Krakentools have been pre-installed, for example, in a miniconda environment called 'kraken', which can be activated as follows.

```
$ conda activate kraken
```

First, the user must download and build a database containing the l -mers and the taxonomic information. One can (i) download a pre-built database from <https://benlangmead.github.io/aws-indexes/k2>, (ii) use the `kraken2-build -standard` command, or (iii) build one's own custom database with all the taxa that should be included. For simplicity, the use of a pre-built database will be described here. However, it should be noted that building a custom database may be necessary to achieve the best and most meaningful results.

Here, we used the PlusPFP-16 database, a database that adds sequences of protozoa, fungi, and plants to those in the standard database and is capped at 16GB. PlusPFP-16 is downloaded (by clicking on the .tar.gz file in the relevant row), and the tarball can then be extracted.

```
$ tar -xvzf k2_plusfpfp_16gb_20240112.tar.gz
```

For classifying the unmapped reads from a whole-genome sequencing project, the `kraken2` command is used. The placeholder `$DBPATH` should be replaced with the path where the database is saved, along with the name assigned to the database. The number of threads used by Kraken 2 can also be specified by replacing `$THREADNUM` with the desired number of threads. The flags `-classified-out` and `-unclassified-out` are optional.

```
$ kraken2 \
--db ${DBPATH} \
--output ${OUTPUTNAME} \
--use-names \
--report ${REPORTNAME} \
--classified-out ${CLASSIFIEDNAME} \
--unclassified-out ${UNCLASSIFIEDNAME} \
--confidence 0.1 \
--threads ${THREADNUM} \
${INPUTNAME}
```

Kraken 2 requires sufficient free memory to load the database into RAM that is, at least 16GB free RAM in this case. This command may take some time to complete and will generate a classification output file (.kraken) and a taxonomic summary

(.kreport). The output file (with a .kraken extension) contains a list of all reads and their classifications. The report (with a .kreport extension) is particularly useful, as it shows the number of reads classified under each taxon.

At this point, the script kreport2krona.py from KrakenTools can be used to create a visualisation of the results.

```
$ python3 kreport2krona.py -r REPORTNAME -o
SAMPLE.krona
ktImportText SAMPLE.krona -o SAMPLE.krona.html
```

The final results show that from the 1% classified unmapped reads, nearly all came from humans, which suggests that during the sample collection or laboratory work the giant panda sample was contaminated with human DNA. 0.09% of the classified unmapped reads had bacterial origin (Figure 9). In addition, a small fraction has come from viruses. Still, 99% of the unmapped reads were unclassified, suggesting that building a larger custom database might be beneficial.

3.4.5 | Comparative Genomics

Having access to existing genomic data greatly enhances the interpretability and context of newly analysed genomes. As users process their data through SwarmGenomics, they can directly compare key metrics, such as heterozygosity, repeat content, or inbreeding levels—to those from a growing database of diverse species.

To illustrate the utility of comparative genomic analyses, several summary plots are presented (Figure 10). Scaffold continuity is compared across species by showing the proportion of scaffolds above and below 5 Mb (Figure 10a). Heterozygosity levels are examined in relation to body mass, obtained from MammalBase (The MammalBase Community 2022) and Avibase (Lepage et al. 2014), providing insight into potential life-history patterns (Figure 10d). Runs of homozygosity are visualised to highlight extreme cases, such as the Indian lion, known for its inbred genome (Figure 10c) (De Manuel et al. 2020). The relationship between NUMTs and repeat content is explored across genomes (Figure 10b). Lastly, the most abundant bacterial phyla detected from unmapped reads are summarised, providing a snapshot of the potential microbiome or environmental signals in the raw sequencing data (Figure 10e).

4 | Discussion

The SwarmGenomics pipeline was developed to simplify whole-genome assembly and streamline downstream analyses, providing insights into species-level genomic diversity and evolutionary history. The pipeline comprises seven modular components: heterozygosity, runs of homozygosity, PSMC modelling, unmapped reads, repeat annotation, mitochondrial genome assembly and NUMT analysis. Together, these modules demonstrate how meaningful genomic characteristics can be extracted from a single whole-genome sequence, even for non-model species.




FIGURE 9 | Visualisation of the classified bacterial reads from the genomic sample of the giant panda: The graph shows the proportion of classified bacterial reads in the unmapped data.

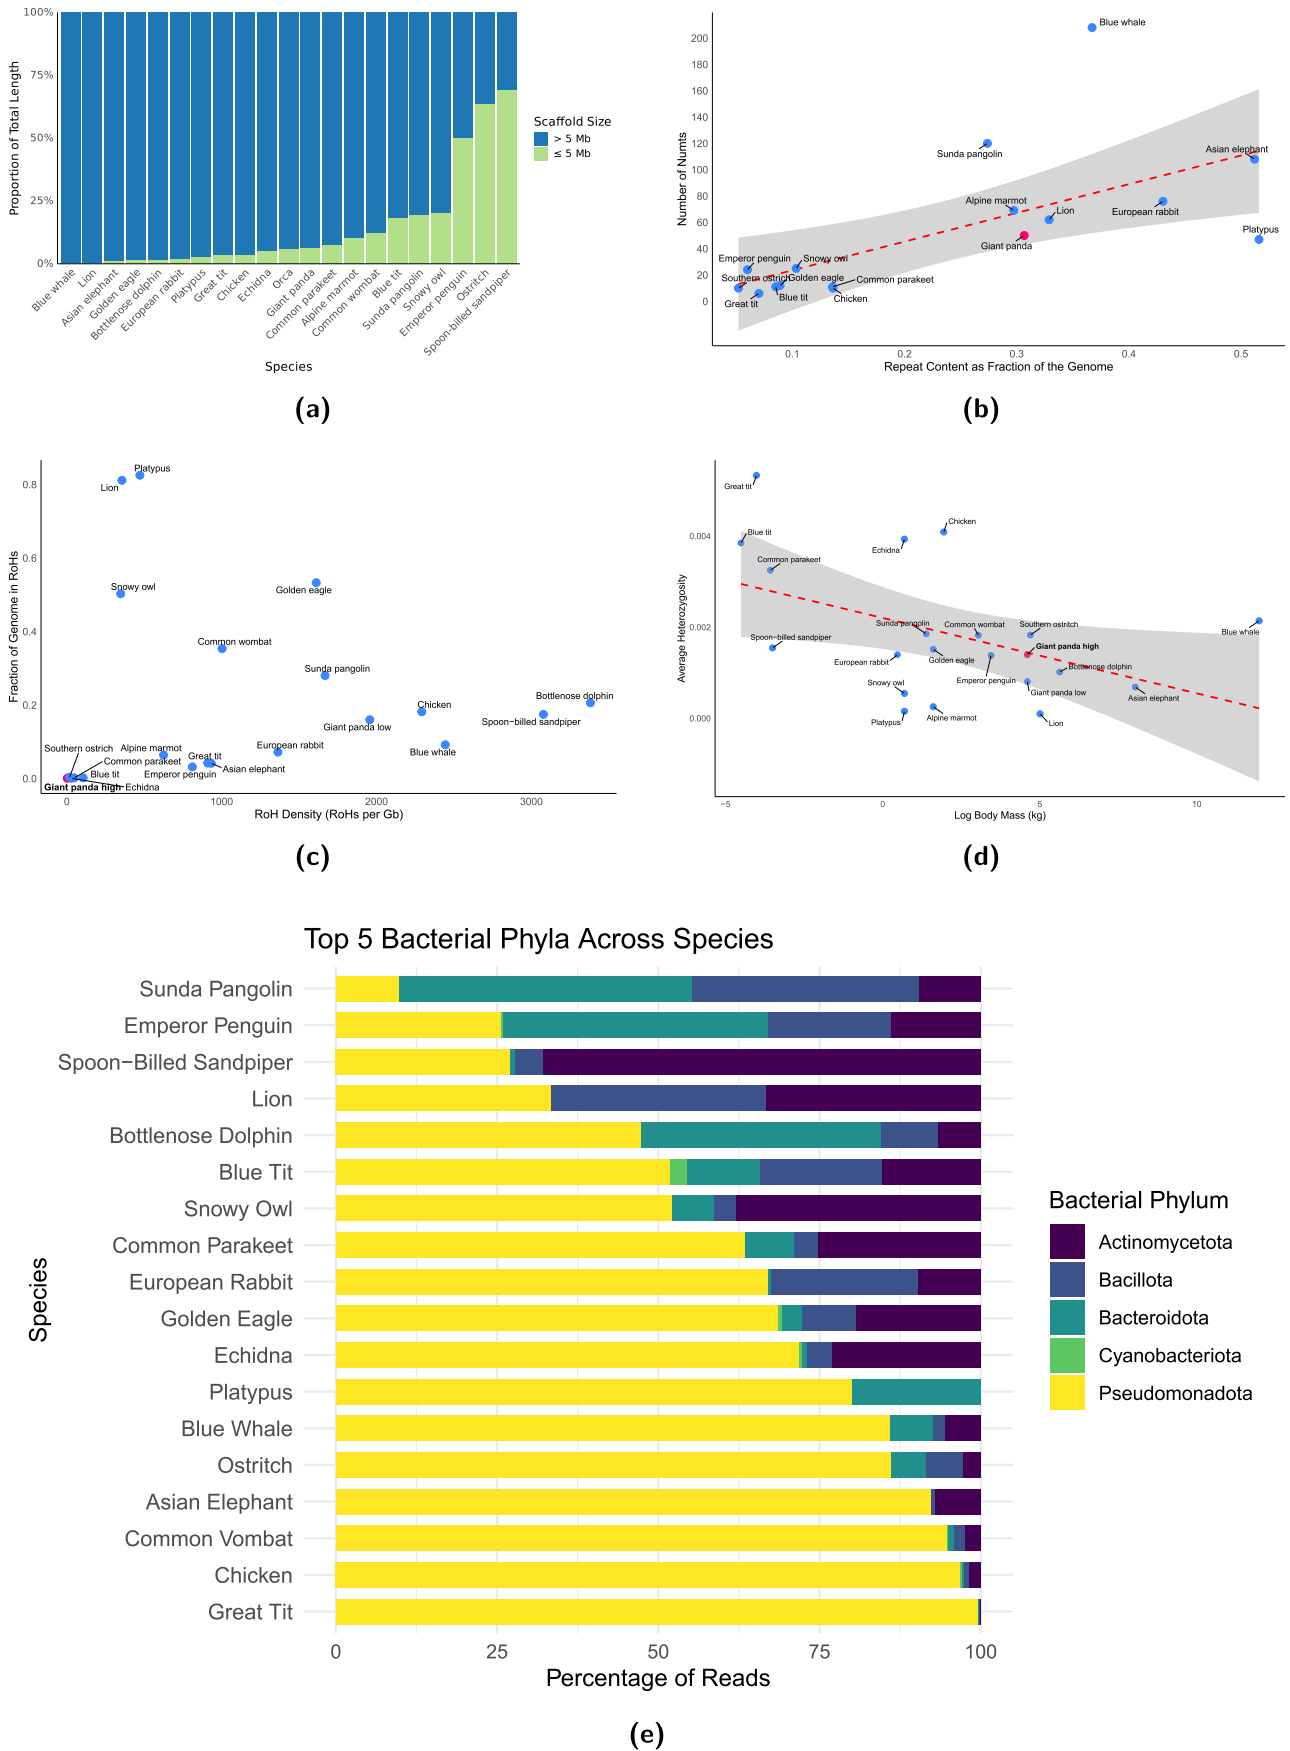


FIGURE 10 | Comparative genomics and associated analyses across studied species. (a) Proportion of scaffolds below and above 5 Mb across the genomes. (b) Correlation between NUMTs and repeat elements across species. (c) Comparison of Runs of Homozygosity (ROH) across species. (d) Correlation between heterozygosity and body mass. (e) Top 5 bacterial phyla across the species.

Using a single giant panda genome, we uncovered valuable insights into the genetic health and history of the species. The heterozygosity analysis across the largest scaffolds revealed moderate genetic diversity, although values varied by scaffold. This aligns with earlier studies using microsatellites, which have also shown variability in genetic diversity across panda populations (Lu et al. 2001; Shan et al. 2014; Li et al. 2023). The RoH analysis revealed few runs of homozygosity longer than 10 kb, suggesting an absence of recent inbreeding in the sampled individual. In contrast, application of the same analysis to a lion genome revealed numerous long RoHs, reflecting a different inbreeding history. The PSMC modelling indicated a major decline in effective population size between 500,000 and 1 million years ago, corresponding to two major Pleistocene glaciations in China, consistent with findings in (Zhao et al. 2013). Additionally, we identified 50 NUMTs, which could potentially be studied further to investigate historical hybridization or introgression events. The repeat annotation revealed a recent expansion of the LINE1 transposable elements, possibly indicating ongoing activity. TEs overall were found to constitute nearly one-third of the genome, consistent with previous estimates of approximately 37% (Li et al. 2010). The unmapped reads module identified signs of contamination by human DNA, as well as sequences of bacterial origin, with a large fraction attributed to the Pseudomonadota phylum. Given that only ~1% of unmapped reads could be classified, a larger, custom database would likely improve resolution. Altogether, the analysis illustrates how a single genome can reveal insights into genetic diversity, demographic history, symbiont presence, and evolutionary processes relevant to species conservation and biology.

4.1 | Pipeline Features and Future Development

A key strength of SwarmGenomics is its modular design, which enables users to select only the analyses most relevant to their study system and research questions. At present, the pipeline assumes access to a pre-assembled reference genome, but future releases will include a module for reference genome assembly, broadening applicability to a wider range of species.

Additional planned expansions include modules for population genetic inference, evolutionary genomics, and selection analysis, such as evaluating coding region heterozygosity and SNPeffect. We also aim to support multiple sequencing platforms, including PacBio and Oxford Nanopore, which will allow the incorporation of epigenomic data in future versions of the pipeline.

4.2 | Collaborative Development and Educational Use

The long-term vision of SwarmGenomics is to function as both a bioinformatics teaching platform and a collaborative research resource. The project's GitHub repository provides all materials needed to run analyses, with user-friendly documentation designed to support students and researchers new to genomics. It is freely available for use in university courses and self-guided learning.

A dedicated SwarmGenomics website is currently under development and will serve as a central database for user-generated

results. Users will be able to contribute intermediate and final output files, which will be made publicly accessible and downloadable, facilitating open data sharing.

'SwarmGenomics' refers to the project's emphasis on collaboration, where multiple researchers contribute genomic datasets and analyses to a common platform. This open-access model enhances the reuse of data and allows comparative studies across diverse taxa. By enabling shared contributions and reproducible workflows, SwarmGenomics aims to strengthen collaborative research in evolutionary and conservation genomics.

5 | Conclusion

In conclusion, SwarmGenomics offers a comprehensive and user-friendly pipeline that integrates whole-genome assembly with advanced genomic analyses. By providing insights into population genetics and evolutionary history, it serves as a valuable tool for both research and education. Planned developments, including modules for reference genome assembly and expanded analytical capabilities, will further enhance its utility across diverse taxa. We invite the community to use, adapt, and contribute to SwarmGenomics, promoting a shared platform for advancing genomic research and discovery.

Nomenclature

Diploid Organism An organism that has two sets of chromosomes, one from each parent.

Genome Assembly The process of piecing together short DNA sequences to reconstruct the full genome of an organism.

Heterozygosity A measure of genetic variation indicating the presence of different alleles at a gene locus within an individual or population.

Runs of Homozygosity (RoH) Contiguous stretches of DNA in the genome where the alleles at all loci are homozygous—indicating that identical copies of genetic material were inherited from both parents across this region.

PSMC (Pairwise Sequentially Markovian Coalescent) A computational method used to infer historical population sizes from a single diploid genome.

Unmapped Reads Sequences in a genomic dataset that cannot be aligned to a reference genome, potentially indicating novel or uncharacterized DNA.

Repeat Analysis The study of repetitive DNA sequences in the genome, which can have structural and regulatory functions.

Mitochondrial Genome Assembly Reconstruction of the complete mitochondrial DNA sequence from whole-genome sequencing data.

NUMTs (Nuclear Mitochondrial DNA Segments) Fragments of mitochondrial DNA that have been incorporated into the nuclear genome.

Effective Population Size The number of individuals that effectively participates in producing the next generation. It reflects the genetic diversity of a population, which can be different from the actual number of individuals in the population.

Reference Genome A high-quality, complete genome sequence used as a standard for mapping and comparing DNA sequences from other individuals or species.

PacBio/Nanopore Sequencing Advanced sequencing technologies that produce long-read sequences, useful for assembling complex genomes and studying epigenetics.

Variant Calling The process of identifying genetic variants, such as SNPs or indels, by comparing sequencing data to a reference genome.

Author Contributions

T.I.G. designed the study with input from all authors. A.K. implemented the pipeline and performed the analyses with help from all other authors. A.K. drafted the manuscript with input from all authors. T.I.G. provided supervision and critical feedback. A.K. led manuscript revisions with input from all authors. All authors read and approved the final manuscript.

Acknowledgements

This work was supported by funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme, grant agreement No. 947636. This work was supported by the de.NBI Cloud within the German Network for Bioinformatics Infrastructure (de.NBI) and ELIXIR-DE (Forschungszentrum Jülich and W-de.NBI-001, W-de.NBI-004, W-de.NBI-008, W-de.NBI-010, W-de.NBI-013, W-de.NBI-014, W-de.NBI-016, W-de.NBI-022). We would like to thank Ina Karabay for providing a baseline script for genome pre-processing. We also appreciate the students of the SwarmGenomics course at TU Dortmund University for their participation, which has allowed us to refine and enhance the course further. Moreover, we wish to express gratitude to other members of the Computational Systems Biology group for helpful feedback. We thank two anonymous reviewers for critical feedback that helped to improve this study. Open Access funding enabled and organized by Projekt DEAL.

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

SwarmGenomics is available on GitHub at <https://github.com/AureKylmanen/Swarmgenomics>.

References

Abascal, F., A. Corvelo, F. Cruz, et al. 2016. "Extreme Genomic Erosion After Recurrent Demographic Bottlenecks in the Highly Endangered Iberian Lynx." *Genome Biology* 17: 1–19.

Afgan, E., D. Baker, M. van den Beek, et al. 2016. "The Galaxy Platform for Accessible, Reproducible and Collaborative Biomedical Analyses: 2016 Update." *Nucleic Acids Research* 44, no. W1: W3–W10.

Allio, R., A. Schomaker-Bastos, J. Romiguier, F. Prosdocimi, B. Nabholz, and F. Delsuc. 2020. "MitoFinder: Efficient Automated Large-Scale Extraction of Mitogenomic Data in Target Enrichment Phylogenomics." *Molecular Ecology Resources* 20, no. 4: 892–905.

Almojil, D., Y. Bourgeois, M. Falis, I. Hariyani, J. Wilcox, and S. Boissinot. 2021. "The Structural, Functional and Evolutionary Impact of Transposable Elements in Eukaryotes." *Genes* 12, no. 6: 918.

Altschul, S. F., W. Gish, W. Miller, E. W. Myers, and D. J. Lipman. 1990. "Basic Local Alignment Search Tool." *Journal of Molecular Biology* 215, no. 3: 403–410.

Andrews, S. 2010. "FastQC: A Quality Control Tool for High Throughput Sequence Data."

Baldo, L., A. de Queiroz, M. Hedin, C. Y. Hayashi, and J. Gatesy. 2010. "Nuclear–Mitochondrial Sequences as Witnesses of Past Interbreeding and Population Diversity in the Jumping Bristletail *mesomachilis*." *Molecular Biology and Evolution* 28, no. 1: 195–210. <https://doi.org/10.1093/molbev/msq193>.

Bankevich, A., S. Nurk, D. Antipov, et al. 2012. "SPAdes: A New Genome Assembly Algorithm and Its Applications to Single-Cell Sequencing." *Journal of Computational Biology* 19, no. 5: 455–477.

Birney, E., T. J. Hudson, E. D. Green, et al. 2009. "Prepublication Data Sharing." *Nature* 461, no. 7261: 168–170.

Bolger, A. M., M. Lohse, and B. Usadel. 2014. "Trimmomatic: A Flexible Trimmer for Illumina Sequence Data." *Bioinformatics* 30, no. 15: 2114–2120.

Broman, K. W., and J. L. Weber. 1999. "Long Homozygous Chromosomal Segments in Reference Families From the Centre D'etude du Polymorphisme Humain." *American Journal of Human Genetics* 65, no. 6: 1493–1500.

Bücking, R., M. P. Cox, G. Hudjashov, et al. 2019. "Archaic Mitochondrial DNA Inserts in Modern Day Nuclear Genomes." *BMC Genomics* 20, no. 1: 1–14. <https://doi.org/10.1186/s12864-019-6392-8>.

Calabrese, F. M., D. L. Balacco, R. Preste, et al. 2017. "NumtS Colonization in Mammalian Genomes." *Scientific Reports* 7, no. 1: 16357. <https://doi.org/10.1038/s41598-017-16750-2>.

Camacho, C., G. Coulouris, V. Avagyan, et al. 2009. "BLAST+: Architecture and Applications." *BMC Bioinformatics* 10: 1–9.

Ceballos, F. C., P. K. Joshi, D. W. Clark, M. Ramsay, and J. F. Wilson. 2018. "Runs of Homozygosity: Windows Into Population History and Trait Architecture." *Nature Reviews Genetics* 19, no. 4: 220–234.

Cechova, M. 2020. "Probably Correct: Rescuing Repeats With Short and Long Reads." *Genes* 12, no. 1: 48.

Charlesworth, D., and J. H. Willis. 2009. "The Genetics of Inbreeding Depression." *Nature Reviews Genetics* 10, no. 11: 783–796.

Chen, Y.-C., D. L. J. Vendrami, M. L. Huber, et al. 2025. "Diverse Evolutionary Trajectories of Mitocoding DNA in Mammalian and Avian Nuclear Genomes." *Genome Research* 35, no. 6: 1313–1324.

Chrisman, B., C. He, J.-Y. Jung, et al. 2022. "The Human "Contaminome": Bacterial, Viral, and Computational Contamination in Whole Genome Sequences From 1000 Families." *Scientific Reports* 12: 9863.

Cock, P. J. A., C. J. Fields, N. Goto, M. L. Heuer, and P. M. Rice. 2010. "The Sanger FASTQ File Format for Sequences With Quality Scores, and the Solexa/Illumina FASTQ Variants." *Nucleic Acids Research* 38, no. 6: 1767–1771.

Danecek, P., A. Auton, G. Abecasis, et al. 2011. "The Variant Call Format and VCFtools." *Bioinformatics* 27, no. 15: 2156–2158.

D'Aquila, P., R. Giacconi, M. Malavolta, et al. 2021. "Microbiome in Blood Samples From the General Population Recruited in the MARK-AGE Project: A Pilot Study." *Frontiers in Microbiology* 12: 707515.

Dayama, G., W. Zhou, J. Prado-Martinez, T. Marques-Bonet, and R. E. Mills. 2020. "Characterization of Nuclear Mitochondrial Insertions in the Whole Genomes of Primates." *NAR Genomics and Bioinformatics* 2: 4. <https://doi.org/10.1093/nargab/lqaa089>.

De Manuel, M., R. Barnett, M. Sandoval-Velasco, et al. 2020. "The Evolutionary History of Extinct and Living Lions." *Proceedings of the National Academy of Sciences* 117, no. 20: 10927–10934.

Dierckxsens, N., P. Mardulyn, and G. Smits. 2017. "NOVOPlasty: De Novo Assembly of Organelle Genomes From Whole Genome Data." *Nucleic Acids Research* 45, no. 4: e18.

Eid, J., A. Fehr, J. Gray, et al. 2009. "Real-Time DNA Sequencing From Single Polymerase Molecules." *Science* 323, no. 5910: 133–138.

Ferreira, T., and S. Rodriguez. 2024. "Mitochondrial DNA: Inherent Complexities Relevant to Genetic Analyses." *Genes* 15, no. 5: 617.

Flynn, J. M., R. Hubley, C. Goubert, et al. 2020. "RepeatModeler2 for Automated Genomic Discovery of Transposable Element Families." *Proceedings of the National Academy of Sciences* 117, no. 17: 9451–9457.

- Freedman, A. H., I. Gronau, R. M. Schweizer, et al. 2014. "Genome Sequencing Highlights the Dynamic Early History of Dogs." *PLoS Genetics* 10, no. 1: e1004016. <https://doi.org/10.1371/JOURNAL.PGEN.1004016>.
- Fujiwara, K., Y. Kawai, T. Takada, et al. 2022. "Insights Into *Mus musculus* Population Structure Across Eurasia Revealed by Whole-Genome Analysis." *Genome Biology and Evolution* 14, no. 5: evac068.
- Gouin, A., F. Legeai, P. Nouhaud, et al. 2015. "Whole-Genome Re-Sequencing of Non-Model Organisms: Lessons From Unmapped Reads." *Heredity* 114, no. 5: 494–501.
- Hahn, C., L. Bachmann, and B. Chevreur. 2013. "Reconstructing Mitochondrial Genomes Directly From Genomic Next-Generation Sequencing Reads—A Baiting and Iterative Mapping Approach." *Nucleic Acids Research* 41, no. 13: e129.
- Hazkani-Covo, E. 2022. "A Burst of Numt Insertion in the Dasyuridae Family During Marsupial Evolution." *Frontiers in Ecology and Evolution* 10: 844443. <https://doi.org/10.3389/fevo.2022.844443>.
- Hilgers, L., S. Liu, A. Jensen, et al. 2025. "Avoidable False PSMC Population Size Peaks Occur Across Numerous Studies." *Current Biology* 35, no. 4: 927–930.e3. <https://doi.org/10.1016/j.cub.2024.09.028>.
- Hood, L., and L. Rowen. 2013. "The Human Genome Project: Big Science Transforms Biology and Medicine." *Genome Medicine* 5: 1–8.
- IUCN. 2023. "Giant Panda: *Ailuropoda melanoleuca*." <https://www.iucnredlist.org/species/712/121745669>.
- Jain, M., H. E. Olsen, B. Paten, and M. Akeson. 2016. "The Oxford Nanopore MinION: Delivery of Nanopore Sequencing to the Genomics Community." *Genome Biology* 17: 1–11.
- Jin, J.-J., W. B. Yu, J. B. Yang, et al. 2020. "GetOrganelle: A Fast and Versatile Toolkit for Accurate De Novo Assembly of Organelle Genomes." *Genome Biology* 21: 1–31.
- Katoh, K., and D. M. Standley. 2013. "MAFFT Multiple Sequence Alignment Software Version 7: Improvements in Performance and Usability." *Molecular Biology and Evolution* 30, no. 4: 772–780.
- Korneliusson, T. S., A. Albrechtsen, and R. Nielsen. 2014. "ANGSD: Analysis of Next Generation Sequencing Data." *BMC Bioinformatics* 15: 1–13.
- Laine, V. N., T. I. Gossmann, K. M. Schachtschneider, et al. 2016. "Evolutionary Signals of Selection on Cognition From the Great Tit Genome and Methylome." *Nature Communications* 7, no. 1: 10474.
- Laine, V. N., T. I. Gossmann, K. Van Oers, et al. 2019. "Exploring the Unmapped DNA and RNA Reads in a Songbird Genome." *BMC Genomics* 20: 1–12.
- Lammers, F., A. Janke, C. Rücklé, V. Zizka, and M. A. Nilsson. 2017. "Screening for the Ancient Polar Bear Mitochondrial Genome Reveals Low Integration of Mitochondrial Pseudogenes (Numts) in Bears." *Mitochondrial DNA Part B Resources* 2, no. 1: 251–254. <https://doi.org/10.1080/23802359.2017.1318673>.
- Lepage, D., G. Vaidya, and R. Guralnick. 2014. "Avibase—A Database System for Managing and Organizing Taxonomic Concepts." *ZooKeys* 420: 117–135.
- Lewin, H. A., S. Richards, E. Lieberman Aiden, et al. 2022. "The Earth BioGenome Project 2020: Starting the Clock."
- Li, D., R. Luo, C.-M. Liu, et al. 2016. "MEGAHIT v1.0: A Fast and Scalable Metagenome Assembler Driven by Advanced Methodologies and Community Practices." *Methods* 102: 3–11.
- Li, H., and R. Durbin. 2009. "Fast and Accurate Short Read Alignment With Burrows–Wheeler Transform." *Bioinformatics* 25, no. 14: 1754–1760.
- Li, H., and R. Durbin. 2011. "Inference of Human Population History From Individual Whole-Genome Sequences." *Nature* 475, no. 7357: 493–496.
- Li, H., B. Handsaker, A. Wysoker, et al. 2009. "The Sequence Alignment/Map Format and SAMtools." *Bioinformatics* 25, no. 16: 2078–2079.
- Li, R., W. Fan, G. Tian, et al. 2010. "The Sequence and De Novo Assembly of the Giant Panda Genome." *Nature* 463, no. 7279: 311–317.
- Li, W., C. Zhou, M. Cheng, et al. 2023. "Large-Scale Genetic Surveys for Main Extant Population of Wild Giant Panda (*Ailuropoda melanoleuca*) Reveals an Urgent Need of Human Management." *Evolutionary Applications* 16, no. 3: 738–749.
- Liang, B., N. Wang, N. Li, et al. 2018. "Comparative Genomics Reveals a Burst of Homoplasmy-Free Numt Insertions." In *Molecular Biology and Evolution*, edited by A. Yoder, vol. 35, 2060–2064. Oxford University Press. <https://doi.org/10.1093/molbev/msy112>.
- Lin, Y.-H., S.-N. Lim, C.-Y. Chen, H.-C. Chi, C.-T. Yeh, and W.-R. Lin. 2022. "Functional Role of Mitochondrial DNA in Cancer Progression." *International Journal of Molecular Sciences* 23, no. 3: 1659.
- Lopez, J., J. V. Lopez, N. Yuhki, R. Masuda, W. Modi, and S. J. O'Brien. 1994. "Numt, a Recent Transfer and Tandem Amplification of Mitochondrial DNA to the Nuclear Genome of the Domestic Cat." *Journal of Molecular Evolution* 39: 174–190. <https://doi.org/10.1007/BF00163806>.
- Lu, Z., W. E. Johnson, M. Menotti-Raymond, et al. 2001. "Patterns of Genetic Diversity in Remaining Giant Panda Populations." *Conservation Biology* 15, no. 6: 1596–1607.
- Lucas, T., V. Bretagnolle, and E. Pante. 2022. "Translocation of Mitochondrial DNA Into the Nuclear Genome Blurs Phylogeographic and Conservation Genetic Studies in Seabirds." *Royal Society Open Science* 9, no. 6: 211888. <https://doi.org/10.1098/rsos.211888>.
- Mahar, N. S., R. Satyam, D. Sundar, and I. Gupta. 2023. "A Systematic Comparison of Human Mitochondrial Genome Assembly Tools." *BMC Bioinformatics* 24, no. 1: 341.
- Meng, G., Y. Li, C. Yang, and S. Liu. 2019. "MitoZ: A Toolkit for Animal Mitochondrial Genome Assembly, Annotation and Visualization." *Nucleic Acids Research* 47, no. 11: e63.
- Miraldo, A., G. O. D. F. R. E. Y. M. Hewitt, P. A. U. L. H. Dear, O. C. T. A. V. I. O. S. Paulo, and B. C. Emerson. 2012. "Numts Help to Reconstruct the Demographic History of the Ocellated Lizard (*Lacerta lepida*) in a Secondary Contact Zone." *Molecular Ecology* 21, no. 4: 1005–1018. <https://doi.org/10.1111/j.1365-294X.2011.05422.x>.
- Mirchandani, C. D., A. J. Shultz, G. W. Thomas, et al. 2024. "A Fast, Reproducible, High-Throughput Variant Calling Workflow for Population Genomics." *Molecular Biology and Evolution* 41, no. 1: msad270. <https://doi.org/10.1093/MOLBEV/MSAD270>.
- Moustafa, A., C. Xie, E. Kirkness, et al. 2017. "The Blood DNA Virome in 8,000 Humans." *PLoS Pathogens* 13, no. 3: e1006292.
- NCBI. 2025a. "NCBI Genome Database." <https://www.ncbi.nlm.nih.gov/genome/>.
- NCBI. 2025b. "NCBI Sequence Read Archive (SRA)." <https://www.ncbi.nlm.nih.gov/sra>.
- Nei, M. 1973. "Analysis of Gene Diversity in Subdivided Populations." *Proceedings of the National Academy of Sciences* 70, no. 12: 3321–3323.
- Nocq, J., M. Celton, P. Gendron, S. Lemieux, and B. T. Wilhelm. 2013. "Harnessing Virtual Machines to Simplify Next-Generation DNA Sequencing Analysis." *Bioinformatics* 29, no. 17: 2075–2083.
- Nyholm, S. V. 2020. "In the Beginning: Egg–Microbe Interactions and Consequences for Animal Hosts." *Philosophical Transactions of the Royal Society, B: Biological Sciences* 375, no. 1808: 20190593.
- Oakeson, K. F., J. M. Wagner, M. Mendenhall, A. Rohrwasser, and R. Atkinson-Dunn. 2017. "Bioinformatic Analyses of Whole-Genome Sequence Data in a Public Health Laboratory." *Emerging Infectious Diseases* 23, no. 9: 1441–1445.

- O’Cathail, C., A. Ahamed, J. Burgin, et al. 2025. “The European Nucleotide Archive in 2024.” *Nucleic Acids Research* 53, no. D1: D49–D55.
- O’Dwyer, D. N., R. P. Dickson, and B. B. Moore. 2016. “The Lung Microbiome, Immunity, and the Pathogenesis of Chronic Lung Disease.” *Journal of Immunology* 196, no. 12: 4839–4847.
- O’Leary, N. A., M. W. Wright, J. R. Brister, et al. 2016. “Reference Sequence (RefSeq) Database at NCBI: Current Status, Taxonomic Expansion, and Functional Annotation.” *Nucleic Acids Research* 44, no. D1: D733–D745.
- Palkopoulou, E., S. Mallick, P. Skoglund, et al. 2015. “Complete Genomes Reveal Signatures of Demographic and Genetic Declines in the Woolly Mammoth.” *Current Biology* 25, no. 10: 1395–1400.
- Patil, A. B., and N. Vijay. 2021. “Repetitive Genomic Regions and the Inference of Demographic History.” *Heredity* 127, no. 2: 151–166.
- Pedersen, B. S., and A. R. Quinlan. 2018. “Mosdepth: Quick Coverage Calculation for Genomes and Exomes.” *Bioinformatics* 34, no. 5: 867–868.
- Pérez, T., F. Rodríguez, M. Fernández, J. Albornoz, and A. Domínguez. 2017. “Ancient Mitochondrial Pseudogenes Reveal Hybridization Between Distant Lineages in the Evolution of the *Rupicapra* Genus.” *Gene* 628: 63–71. <https://doi.org/10.1016/j.gene.2017.07.035>.
- Price, M. N., P. S. Dehal, and A. P. Arkin. 2010. “FastTree 2—Approximately Maximum-Likelihood Trees for Large Alignments.” *PLoS One* 5, no. 3: e9490.
- Prum, R. O., J. S. Berv, A. Dornburg, et al. 2015. “A Comprehensive Phylogeny of Birds (Aves) Using Targeted Next-Generation DNA Sequencing.” *Nature* 526, no. 7574: 569–573.
- Quinlan, A. R. 2014. “BEDTools: The Swiss-Army Tool for Genome Feature Analysis.” *Current Protocols in Bioinformatics* 47, no. 1: 11–12.
- Reed, D. H., and R. Frankham. 2003. “Correlation Between Fitness and Genetic Diversity.” *Conservation Biology* 17, no. 1: 230–237.
- Rhie, A., S. A. McCarthy, O. Fedrigo, et al. 2021. “Towards Complete and Error-Free Genome Assemblies of All Vertebrate Species.” *Nature* 592, no. 7856: 737–746.
- Rodríguez, F., and I. R. Arkhipova. 2022. “An Overview of Best Practices for Transposable Element Identification, Classification, and Annotation in Eukaryotic Genomes.” *Transposable Elements: Methods and Protocols* 2607: 1–23.
- Sanger. 2025. “Genome After Party.” https://pipelines.tol.sanger.ac.uk/genome_after_party.
- Shan, L., Y. Hu, L. Zhu, et al. 2014. “Large-Scale Genetic Survey Provides Insights Into the Captive Management and Reintroduction of Giant Pandas.” *Molecular Biology and Evolution* 31, no. 10: 2663–2671.
- Slotkin, R. K. 2018. “The Case for Not Masking Away Repetitive DNA.” *Mobile DNA* 9, no. 1: 15.
- Smith, T. P. L., D. M. Bickhart, D. Boichard, et al. 2023. “The Bovine Pangenome Consortium: Democratizing Production and Accessibility of Genome Assemblies for Global Cattle Breeds and Other Bovine Species.” *Genome Biology* 24, no. 1: 139.
- Song, N., H. Zhang, R.-e. Bai, and H.-g. Meng. 2022. “The Mitogenome of *Aleuroclava psidii* (Singh, 1931) (Hemiptera: Aleyrodidae) and Increased Number of Mitochondrial Gene Rearrangements in Whiteflies.” *Frontiers in Bioscience-Landmark* 27, no. 5: 154.
- Soto-Calderón, I. D., N. J. Clark, J. V. H. Wildschutte, K. DiMattio, M. I. Jensen-Seaman, and N. M. Anthony. 2014. “Identification of Species-Specific Nuclear Insertions of Mitochondrial DNA (Numts) in Gorillas and Their Potential as Population Genetic Markers.” *Molecular Phylogenetics and Evolution* 81: 61–70. <https://doi.org/10.1016/j.ympev.2014.08.018>.
- Takayama, J., S. Tadaka, K. Yano, et al. 2021. “Construction and Integration of Three De Novo Japanese Human Genome Assemblies Toward a Population-Specific Reference.” *Nature Communications* 12, no. 1: 226.
- Tanizawa, Y., T. Fujisawa, Y. Kodama, et al. 2023. “DNA Data Bank of Japan (DDBJ) Update Report 2022.” *Nucleic Acids Research* 51, no. D1: D101–D105.
- Tarailo-Graovac, M., and N. Chen. 2009. “Using RepeatMasker to Identify Repetitive Elements in Genomic Sequences.” *Current Protocols in Bioinformatics* 25, no. 1: 4–10.
- Teeling, E. C., S. C. Vernes, L. M. Dávalos, et al. 2018. “Bat Biology, Genomes, and the Bat1K Project: To Generate Chromosome-Level Genomes for All Living Bat Species.” *Annual Review of Animal Biosciences* 6, no. 1: 23–46.
- The MammalBase Community. 2022. “MammalBase—Database of Traits, Measurements and Diets of the Species in Class Mammalia. Licensed Under CC BY 4.0.” <https://www.mammalbase.net>.
- Trevelline, B. K., K. J. MacLeod, S. A. Knutie, T. Langkilde, and K. D. Kohl. 2018. “In Ovo Microbial Communities: A Potential Mechanism for the Initial Acquisition of Gut Microbiota Among Oviparous Birds and Lizards.” *Biology Letters* 14, no. 7: 20180225.
- Tsuji, J., M. C. Frith, K. Tomii, and P. Horton. 2012. “Mammalian NUMT Insertion Is Non-Random.” *Nucleic Acids Research* 40, no. 18: 9073–9088. <https://doi.org/10.1093/nar/gks424>.
- Uvizl, M., S. J. Puechmaille, S. Power, et al. 2024. “Comparative Genome Microsynteny Illuminates the Fast Evolution of Nuclear Mitochondrial Segments (NUMTs) in Mammals.” *Molecular Biology and Evolution* 41, no. 1: msad278. <https://doi.org/10.1093/molbev/msad278>.
- Vendrami, D. L. J., T. I. Gossmann, N. Chakarov, et al. 2022. “Signatures of Selection on Mitonuclear Integrated Genes Uncover Hidden Mitogenomic Variation in Fur Seals.” *Genome Biology and Evolution* 14, no. 7: 104. <https://doi.org/10.1093/gbe/evac104>.
- Wallace, D. C. 2010. “Mitochondrial DNA Mutations in Disease and Aging.” *Environmental and Molecular Mutagenesis* 51, no. 5: 440–450.
- Wang, M.-S., J. J. Zhang, X. Guo, et al. 2021. “Large-Scale Genomic Analysis Reveals the Genetic Cost of Chicken Domestication.” *BMC Biology* 19: 1–16.
- Wei, W., K. R. Schon, G. Elgar, et al. 2022. “Nuclear-Embedded Mitochondrial DNA Sequences in 66,083 Human Genomes.” *Nature* 611, no. 7934: 105–114. <https://doi.org/10.1038/s41586-022-05288-7>.
- Wong, K. H. Y., M. Levy-Sakin, and P.-Y. Kwok. 2018. “De Novo Human Genome Assemblies Reveal Spectrum of Alternative Haplotypes in Diverse Populations.” *Nature Communications* 9, no. 1: 3040.
- Wong, K. H. Y., W. Ma, C.-Y. Wei, et al. 2020. “Towards a Reference Genome That Captures Global Genetic Diversity.” *Nature Communications* 11, no. 1: 5482.
- Wood, D. E., J. Lu, and B. Langmead. 2019. “Improved Metagenomic Analysis With Kraken 2.” *Genome Biology* 20, no. 1: 257.
- Zhang, G., D. Geng, Q. Guo, et al. 2021. “Genomic Landscape of Mitochondrial DNA Insertions in 23 Bat Genomes: Characteristics, Loci, Phylogeny, and Polymorphism.” *Integrative Zoology* 4: 1–14. <https://doi.org/10.1111/1749-4877.12582>.
- Zhao, S., P. Zheng, S. Dong, et al. 2013. “Whole-Genome Sequencing of Giant Pandas Provides Insights Into Demographic History and Local Adaptation.” *Nature Genetics* 45, no. 1: 67–71.
- Zoonomia Consortium. 2020. “A Comparative Genomics Multitool for Scientific Discovery and Conservation.” *Nature* 587, no. 7833: 240–245.

Supporting Information

Additional supporting information can be found online in the Supporting Information section. **Data S1:** men70119-sup-0001-DataS1.docx.