Improving data archiving practices in ancient genomics

Anders Bergström
School of Biological Sciences, University of East Anglia, Norwich, UK
Correspondence: a.bergstrom@uea.ac.uk

Abstract

The sequencing of ancient DNA from preserved biological remains is producing a rich record of past genetic diversity in humans and other species. However, unless the primary data is made available in public archives in an appropriate fashion, its long-term value will not be fully realised. I surveyed publicly archived data from 28 recent ancient genomics studies. I found that half of the studies archived incomplete subsets of the generated genomic data, preventing accurate replication and representing a loss of data of potential use for future research. None of the studies met all archiving criteria that could be considered best practice. Based on these results, I make six recommendations for data producers: 1) archive all sequencing reads, not just those that can be aligned to a reference genome, 2) archive read alignments as well, but as secondary analysis files linked to the underlying raw read files, 3) provide correct experiment metadata on how samples, libraries and sequencing runs relate to each other, 4) provide informative sample metadata in the public archives, 5) publish and archive data from screening, low-coverage, poorly performing and negative experiments, and 6) review data archiving as part of peer review processes. Given the reliance on destructive sampling of finite material, I argue that ancient genomics studies have a particularly strong responsibility to ensure the longevity and reusability of generated data.

Introduction

The rich nature of genomic data means that new discoveries are often made using previously published data—including discoveries that the primary producers of the data might not even have envisioned. Ancient DNA has been described as a field with unusually good data sharing practices, as the vast majority of studies do share data in some form[1]. As the field has entered the era of high-throughput sequencing, however, enough attention has arguably not been given to how data is deposited and structured in public archives. Often, incomplete, filtered or mislabelled versions of the data are archived—for example, only the subset of sequencing reads that could be successfully aligned to some given reference genome. When incomplete versions of data are archived, published results can not be replicated from scratch by independent researchers, and the long-term value of the data to future science is reduced.

The way in which publicly archived data is organised, accessioned and associated with useful metadata, also greatly influences its utility to future users. The properties of findability, accessibility, interoperability, and reusability (FAIR) are often put forth as guiding principles for

data sharing[2]. Ideally, data and associated metadata should be machine-readable and processable by automated pipelines with minimal need for manual intervention. Many users of published ancient genomics datasets will have experienced that studies in this field generally fall short of the FAIR principles, as data reuse often requires quite a bit of manual detective work, sifting through the supplementary materials of papers, and sometimes contacting the primary data producers for clarifications. This mirrors the long-standing problem of insufficient or incorrect metadata that affects databases on genetic biodiversity more broadly[3]. Ancient genomic data will soon be available from tens of thousands of individuals, primarily from humans but also from other species as well as various environmental samples, and navigating this will become increasingly challenging if data archiving and documentation standards are not improved.

The primary public archives used for genome sequencing data are those provided by the International Nucleotide Sequence Database Collaboration (INSDC)[4]: the European Nucleotide Archive (ENA)[5], the Sequence Read Archive (SRA)[6], and the DNA Data Bank of Japan (DDBJ)[7], which mirror data between each other. The INSDC ecosystem also includes the BioProject and BioSamples databases[8]. This ecosystem can be challenging to navigate, but does provide a rigorous system for the organisation and documentation of biological samples and genomic sequencing data derived from them. The discussion that follows is primarily centred around the ENA system, but the principles apply more broadly.

I surveyed the genomic data deposited into public archives by 28 recent studies that sequenced ancient DNA from multicellular organisms through shotgun sequencing or large-scale capture[9–36]. Studies with primarily microbial or metagenomic targets were not included here, and practices and considerations might differ slightly for such studies[37,38]. I found widespread shortcomings in data archiving practices across the surveyed studies, and I present the results within the context of six recommendations for producers of ancient genomic data.

Results

1. Archive all sequencing reads

The most important thing to ensure is that all of the raw reads obtained in a sequencing experiment are archived, prior to any mapping to a reference genome, assembly, or other downstream analysis. This unfiltered version of the data represents the full information obtained from the experiment, enables any published findings to be replicated from scratch, and ensures that the long-term utility of the data to future research is maximised. The most commonly used format for raw reads is the FASTQ format[footnote 1]. When only reads that could successfully be

[1] The BAM format can also store unmapped reads, such that an unfiltered BAM file can contain the same complete information as a FASTQ file. If unfiltered BAM files that contain unmapped reads are deposited, as is often done in modern genomics studies, it would not be necessary to also deposit FASTQ files. Indeed, the ENA's primary recommendation is to deposit reads in the BAM or CRAM formats. However, I would recommend that ancient genomics studies deposit FASTQ files, as they are directly usable by most common software and eliminate ambiguity regarding whether all reads are in fact present.

aligned to a reference genome are archived, potentially valuable data is lost from the scientific record and unavailable for reuse by future researchers.

The fundamental argument for archiving raw reads is that even if the producers of that data do not envision much additional utility for the complete set of reads, future users of the data might. Some advantages to the availability of raw reads can be recognized already:

- The ability to remap reads from scratch, e.g. with different software or parameters. Most ancient genomics studies use the bwa aln software[39] to map reads, but if new methods that increase accuracy were to be developed, the benefit of that could only apply fully to previously published datasets if they include all reads.

- The ability to remap reads to different reference genomes. Higher-quality reference genomes can improve read mapping accuracy and enable analyses of previously inaccessible parts of the genome. Furthermore, alternative, multiple or modified reference genomes are sometimes used in ancient genomics to overcome technical challenges[40,41]. It is likely that the process of mapping short, damaged reads to a single reference genome is a major source of reference bias in downstream analyses, as reads that display differences from the reference are more likely to end up as unmapped[42]. Methods using variation graphs, incorporating genetic variation from multiple individuals, can reduce reference bias during read mapping[43]. But if the unmapped reads from previously published studies are not available, no current or future methodological advances can reverse the reference bias introduced during the original mapping process[42].

- The ability to analyse non-endogenous reads of interest, i.e. those from microorganisms and viruses. There are examples in the literature of microbial pathogen DNA being detected in previously published data[44,45]. As metagenomics detection methods improve, yet more discoveries might be made even in data that was published long ago.

My survey found: out of the 28 studies, 14 archived raw reads. The remaining 14 studies only archived read alignments to a reference genome, in no cases including unmapped reads, meaning that there has been a loss of data for those studies.

My survey found: out of the 21 studies making use of capture technology, only 7 archived raw reads from those experiments. An argument could be made that the targeted nature of capture technology reduces the relevance and future utility of reads not mapping to the targeted regions. However, concerns around remapping ability and reference genomes remain, and the fact that target enrichment processes are not 100% efficient means that off-target reads will be present and could prove useful. Non-targeted microbiota have been detected in the non-endogenous reads from modern DNA capture experiments[46]. The scarcity of raw data from ancient capture experiments means that it remains largely unknown to what extent microbial DNA is detectable in such data. I argue that raw reads should be archived for capture experiments too.

Various processing steps are often applied to reads in ancient genomics studies: trimming off technical adaptor sequences; merging overlapping read pairs resulting from paired-end sequencing of very short molecules and discarding those pairs that could not be merged; trimming off the last few bases of reads as these typically have higher rates of postmortem damage. No clear consensus exists on what, if any, processing to apply to reads prior to archiving. As a general principle, reads should be archived in a form that excludes technical artefacts arising from the sequencing experiment, while avoiding loss of information about the underlying biological sequence. Reads submitted for archiving should thus be demultiplexed into separate files per sample if needed, and any technical adapter sequences in the reads should be trimmed off—indeed, the ENA expects reads to not contain adapters.

However, further processing of the reads should arguably be avoided. This includes merging of overlapping pairs in paired-end sequencing experiments. An assumption is often made that if a given pair of reads do not overlap each other, it's because the underlying molecule is too long to be an endogenous ancient molecule, and the read pair should then be excluded. While this is a conservative assumption that likely benefits most downstream analyses, future users should be able to make their own choices about what reads to include or exclude. Depending on DNA preservation properties and read lengths, it should not be ruled out that some datasets might contain endogenous molecules that genuinely are longer than twice the read length and thus fail merging. Some current use cases for archived data also require the paired-end structure of reads to be intact, e.g. in metagenomics de-novo assembly. In a similar vein, while analysts might choose to trim off the last few bases of all reads to mitigate the effects of postmortem damage, this should not be applied to the files that are archived—it represents a loss of information about the underlying sequence, and future users should be able to make their own trimming decisions. In summary, my recommendation is thus to submit reads that have been demultiplexed and trimmed for adapters, but not merged, trimmed or filtered in any further way.

As a source of potential confusion, ENA automatically converts read alignments submitted in BAM format into unaligned reads in FASTQ format, and provides these as "*generated files*" alongside the "*submitted files*". A user downloading an auto-generated FASTQ file might not realise that it derives from a BAM file from which unmapped or other reads have been filtered out, and thus be under the incorrect impression that it contains all the reads obtained in the experiment. Users are thus recommended to check what files were actually submitted (e.g. using the "*submitted_format*" field).

It's worth pointing out that additional files can still be submitted to ENA studies after a paper has been published. In principle, it is thus never too late for data producers to go back and submit missing data to existing studies, if the data still exists somewhere locally (although extra care might be needed if doing so, so as to not create additional confusion about how different files relate to each other—see recommendation 3 below).

2. Archive read alignments as analysis files

The standard format for representing the alignments of reads to a reference genome is the BAM format[47], along with the more heavily compressed alternative CRAM[48]. Sharing alignments should not be considered a requirement for genomics studies in the same way that sharing of raw reads is, but doing so represents a useful service. Providing the exact alignments used for published analyses, including the processing and filtering deemed appropriate by the data producers, enable other researchers to quickly replicate and build upon those analyses. Archiving both raw reads and analysis-level read alignments can thus be considered best practice in ancient genomics data sharing.

It is not ideal to simply submit read alignment BAM files to public archives alongside, and as if they were equivalent to, raw FASTQ files, however. If this is done, the BAM files will be registered as separate experiments and be assigned independent sequencing run accessions, thereby duplicating data. Instead, ENA provides an “Analysis Files” functionality which allows for read alignments to be registered as the products of secondary analysis applied to the raw reads (though these ENA analysis files are not currently mirrored to the other INSDC databases). When submitted in this way, the read alignments are linked to the corresponding raw read files, are assigned “analysis accessions” instead of run accessions, and are listed in a dedicated section on the study page.

My survey found: out of the 28 studies, 22 archived read alignments. However, as discussed in the previous section, for 14 of the studies this was the only thing archived, as raw reads were not also deposited. Only 8 studies archived both raw reads and read alignments. Out of these, only 2 made use of the ENA “Analysis Files” functionality to appropriately categorise read alignment files as the products of secondary analysis rather than independent experiments.

The BAM format is designed to require little if any hard filtering—that is, the complete removal of reads from the file. Instead, information stored for each read on how it is aligned to the reference, including the mapping quality metric[49], allows for soft filtering to be performed on the fly. There is therefore little reason for the particular filtering choices preferred by the data producers to be applied to archived BAM files through hard filtering. For example, while a mapping quality threshold of 30 is often applied in ancient genomics analyses, some users might prefer to apply less stringent thresholds (for example 25[25], 20[26] or 10[9]). But if reads with mapping qualities below 30 have been removed from archived files, it is not possible to apply lower thresholds. Even reads with mapping qualities of 0—meaning they map equally well to more than one place in the reference genome—can be useful, for example for quantifying the total copy number of a genomic region that is represented more than once in the reference. The exclusion of reads that are not uniquely mappable has impeded systematic analyses of copy number variation in ancient human genomes[50]. Reads shorter than some threshold length, typically 35 nucleotides, are similarly often excluded from analyses, but shorter reads could potentially be of future use too[51].

My survey found: out of the 22 studies that archived read alignments in BAM format, all 22 had removed unmapped reads, as noted previously. 18 of the studies appear to have applied one or more additional hard filtering steps: 8 studies hard filtered on some mapping quality threshold of 25 or higher, and 15 studies hard filtered on a read length threshold of 30 or higher. For the 11 out of these 18 studies that did not also submit raw reads, this hard filtering results in the unavailability of potentially useful reads, representing an additional source of data loss on top of the loss resulting from the exclusion of unmapped reads[footnote 2].

3. Provide correct experiment metadata

Ancient genomics projects often involve the sequencing of multiple libraries for the same sample, multiple sequencing runs for the same library, capture of libraries that have previously been shotgun sequenced, etc. Keeping track of how samples, libraries and sequencing runs relate to each other is important to enable correct re-processing, and errors in experiment metadata documentation is a common source of confusion for users of published ancient genomic data. A key aspect is that in the metadata model used by the INSDC databases, samples exist independently of studies and datasets. Samples submitted through these databases are assigned sample accessions and are also automatically registered in the NCBI BioSample database[8]. Data from multiple different experiments, even from different studies, can then be linked to a given sample accession.

Keeping track of various technical details relating to how different libraries have been constructed can also be important in downstream analyses of ancient genomics data. Different methods—e.g. commonly employed 'double-stranded'[52] versus 'single-stranded'[53] library preparation protocols, or enzymatic removal of damaged nucleotides[54]—can impact data properties in different ways. Any relevant information like this can be included in the experimental metadata when submitting data files, e.g. in the optional free text "Library Construction Protocol" field.

My survey found: out of the 28 studies, for 9 studies the number of sample accessions with data files linked to the study in ENA did not match the number of samples cited in the corresponding paper. A user of any of these 9 datasets will thus need to engage in some detective work to figure out how archived files relate to samples described in the paper. For at least 5 studies, the discrepancy in numbers was due to experiments on the same sample, but different libraries, having been registered as separate samples. At least 3 studies reported new data on samples for which other studies had previously reported data, but none of them associated the new data files to the existing sample accessions, but instead registered new sample accessions. In some cases relationships to existing samples were seemingly indicated by providing sample names ending in e.g. "_new", however this does not constitute a formal link

[2] An argument for hard filtering of BAM files is that it reduces file sizes, but I would argue that such gains are typically marginal and do not outweigh the long-term data loss. The one hard filtering step that often does lead to substantial file size reductions is the removal of unmapped reads, owing to the large amounts of non-endogenous DNA typically recovered in ancient sequencing data. If raw reads are not also submitted separately, as discussed above, removing unmapped reads entails a loss of data and should thus be avoided. But if raw reads are also submitted separately as FASTQ files as recommended here, there is no loss of data, and providing smaller, analysis-ready BAMs without unmapped reads is in that case arguably the most practical for users.

within the metadata system and is a recipe for long-term confusion. No studies indicated technical library details, such as library preparation protocol or enzymatic damage removal, in the submitted experimental metadata.

The importance of experimental metadata also extends to read alignments in the BAM format. The BAM format is designed to keep track of relationships between samples, libraries and sequencing runs, through the use of so-called 'read groups'. Each read is assigned to a read group, which should represent a set of reads that share an experimental origin: the same sample, library and sequencing run. The identity of a read group can be described in the header of the BAM file through a number of tags ("SM" for sample, "LB" for library, "PL" for sequencing platform etc.). Most mainstream processing tools for BAM files are read group aware, and some tools will not run in the absence of certain read group tags. For example, most tools that identify PCR duplicates use the library ("LB") read group tag to only mark duplicates within a given library, even if sequenced across multiple runs, but not across different libraries from the same sample. Ideally, the sample and library IDs supplied with these read group tags within BAM files should exactly match the IDs provided in the corresponding ENA metadata, to avoid confusion and enable matching of IDs by automated processes rather than by human eyes.

My survey found: out of the 22 studies that archived BAM files, 19 made use of read groups. 18 made use of the "SM" read group tag indicating sample identity of reads. However, for 10 studies, the sample ID supplied in the "SM" tag did not match any of the sample IDs provided in the ENA metadata. For 3 studies, a single BAM file incorrectly contained more than one "SM" tag. 14 of the 22 studies made use of the "LB" tag indicating library identity of reads. However, for 10 studies, the library ID supplied in the "LB" tag did not match the library ID provided in the ENA metadata. For 7 of the studies, the archived BAM files produced one or more errors when tested using a BAM validation tool, which could relate to missing read groups but also other formatting issues. Overall, the archived BAM files from most of the studies would likely require manual verification or modification in one way or another before being safely usable in analyses.

4. Provide informative sample metadata

Users of publicly archived datasets also require metadata on the biological samples from which the data derives. In ancient genomics studies, key bits of metadata include the geographical location where biological remains were found, and when in the past the organism is thought to have lived (e.g. radiocarbon dates, if available, or other dating assessments). Other relevant metadata include the type of tissue that was sampled (e.g. what specific type of bone, tooth or other tissue), and, for remains from human-associated archaeological sites, the name and/or cultural context of a site. Another piece of information that helps long-term data provenance is reference to any relevant external museum, collection or excavation ID for the sampled specimen. Many studies that do not provide particularly informative sample metadata in ENA or similar archives often do so in the publication describing the study, in which case the information is still available in the scientific record for those that look hard enough for it. However, sample

metadata becomes more useful if it is also formally linked, in a standardised fashion, to the primary data in the public archives.

My survey found: out of the 28 studies, 17 provided a geographical label of one kind or another in the sample metadata submitted to public archives, but only 6 used a dedicated metadata field for this (typically "*geographic location (region and locality)*". Only 3 studies provided latitude and longitude coordinates. 4 studies provided some kind of dating information, but only 1 used a dedicated metadata field for this ("*date*"), and only 1 provided the radiocarbon lab codes that serve as unique identifiers for radiocarbon dates. 12 studies provided information on what tissue was sampled, though in 3 cases this was not higher resolution than just "Bone" versus "Tooth". Only 8 studies provided tissue information in a dedicated metadata field (typically "*tissue*"), while the others provided it in the general purpose "*Description*" field. 4 studies provided some kind of information on the identity or source of the sampled specimen (using the "*biological material*" or "*specimen voucher*" fields), though IDs were not always accompanied by information on what external collection or excavation those IDs refer to.

There is currently no clear consensus in ancient genomics on what metadata should be reported, or in what format various types of metadata should be provided. There is a strong need for standardisation in this area, perhaps including the adoption of controlled vocabularies (e.g. for tissues[55]), but that is a discussion that is outside the scope of this article. Ongoing efforts to e.g. establish a MIxS (*Minimum Information about any (x) Sequence*[56]) extension checklist for ancient samples would be very useful in this regard[57]. One reason for the poor state of sample metadata in the public archives is arguably that the standard metadata fields in the INDSC databases are not geared towards ancient samples, such that it's not necessarily clear to sample submitters in what fields different pieces of information should go. Once stable metadata standards have emerged in the ancient genomics community, these could likely also be adopted by the ENA through the creation of a dedicated sample checklist (the ENA has already adopted such checklists for various other types of samples with particular metadata requirements). In the meantime, a minimal recommendation to data submitters is: whatever metadata you include in the paper, make an effort to submit that same metadata to the public archive too, use metadata fields that are as appropriate as possible, and strive to make information as machine-readable as possible.

It's worth pointing out that, as samples exist independently of studies and data files, sample metadata in the INSDC databases can be updated even after data files have been submitted. It is thus never too late to go back to previously registered samples and supplement them with more informative metadata. Lastly, from mid-2023 the INSDC databases plan to make some minimal spatiotemporal metadata mandatory for newly submitted samples[58]. However, it should be noted that this minimal metadata concerns when and where (e.g. a museum) a sample was collected with the intention of DNA sequencing, and not the time and place where the organism lived[59]. This might become a source of confusion for ancient samples, and additional metadata fields will thus still be needed to provide this crucial information for ancient samples.

5. Archive data from low-coverage, screening and negative experiments

Many biological remains do not contain enough endogenous DNA to make genome sequencing worthwhile. Many ancient genomics projects first screen samples though very low-coverage sequencing, before deciding on what samples to perform deeper sequencing on. Data from this screening phase is typically not used for published analyses, and not deposited in public archives. Some papers mention how many samples were screened but not taken forward, but this is not always done. In a similar way, studies that do not necessarily employ a screening phase will sometimes nonetheless choose to exclude some samples from publications because the data obtained from them was poor, and not share the data from those experiments. While difficult to estimate, it is likely that data from the majority of ancient genome sequencing experiments performed to date remain unpublished and unavailable.

It could be argued that if one undertakes destructive sampling of any ancient biological remains, doing justice to those remains should involve a commitment to publish and make available the resulting data, whatever its quality turns out to be[footnote 3]. The practice of not archiving screening and other low-coverage or poor data can be seen as another instance of when data producers do not envision any utility of the data, but future users might. Already, several potential uses of very low-coverage sequencing data from ancient remains can be identified:

- Microbial and viral DNA can sometimes be detectable even in very low-coverage data. Detection in screening data often informs decisions on capture experiments targeting pathogenic microbes of interest. Beyond these, the non-pathogenic microbial content and what it says about post-mortem degradation processes remains largely unexplored. With continued advances in computational metagenomics methodology, and increased understanding of past pathogenic and non-pathogenic microbial diversity, samples that were initially deemed uninteresting might well be found to contain data of interest by future researchers.

- Data from samples with little or no endogenous DNA can contribute to the understanding of factors influencing DNA preservation. If data from large numbers of such ‘negative’ samples and their metadata was available, supervised models could be trained to predict DNA preservation as a function of various sample properties. Currently, however, the published literature primarily just contains results from relatively successful samples, thereby hindering systematic explorations in this area.

- Determination of biological sex through coverage analyses of the sex chromosomes does not require very high sequencing coverage, and is sometimes possible even for relatively poorly preserved samples[60,61]. While not necessarily very useful in most genomics analyses, the biological sex of individuals across regions, periods, sites etc. can inform on social, behavioural and cultural aspects[62–65].

[3] Another aspect of relevance is authorship fairness—archaeologists, museum curators etc. that provide samples that turn out negative for DNA will often not end up as authors on the publications that emerge from projects, even if their contributions were the same as that of providers of samples that turned out positive.

My survey found: out of the 28 studies, 21 mentioned in the paper the total number of studied samples, but did not necessarily archive data from all samples. 12 of the studies appear to have archived data from all studied samples, even if some were then left out of downstream analyses. For the remaining 16 studies, data from samples other than those included in the published analyses, whether such samples were explicitly mentioned or not discussed at all, did not seem to have been archived—meaning that data from some of the experiments conducted as part of these studies likely remain unpublished and unavailable[footnote 4].

Beyond questions of future utility, the practice of only publishing 'positive' and not 'negative' samples also has implications for the interpretation of scientific results. This is especially true for studies reporting the detection of pathogen DNA in ancient remains—if results from negative samples are not reported, the prevalence of a given pathogen across time and space can not be assessed. Some studies do mention the total number of screened samples, e.g. a study on hepatitis B described data from 137 samples produced after initially screening 7918 human samples of unspecified origin[66], and a study on smallpox described data from 13 samples produced after initially screening 1867 human samples from Eurasia and the Americas from the last ~32,000 years[67]. The latter study furthermore reported results on the number of positive and negative samples across broad time periods, thus providing valuable evidence on past prevalence. It could be argued that published results should not refer to data that is not publicly archived, as other researchers can not reproduce those results. I would instead argue that transparency about what was screened is never a bad thing, especially if it helps interpretation of the scientific findings. Even more useful, of course, would be if all the screening data is published and archived.

Some highly iconic and archaeologically important human remains have presumably been sampled for DNA without success, but those negative results have not been described in the scientific literature. It would arguably be valuable to make the apparent lack of endogenous DNA in some of these specimens part of the public, scientific record—and to make the screening data that was nonetheless produced available for anyone who wants to analyse its properties independently.

6. Peer review data archiving

Many of the common archiving shortcomings discussed here could be overcome relatively easily. A natural mechanism for the flagging of issues relating to the public archiving of data would be if editors and reviewers of papers scrutinised this aspect too, alongside their review of the scientific content of a paper, but this is rarely done. Papers submitted for review often

[4] There are practical challenges to sharing screening and negative data. For some samples there might not be a suitable paper where it would naturally fit, especially if all samples from a given project or archaeological site were negative. Even in such cases, data could be published in stand-alone publications, e.g. in archaeology journals as reports on DNA preservation at a set of sites, or in journals offering data-focussed articles that do not require interpretation of results (e.g. GigaScience, Scientific Data, F1000 Research). Stand-alone ENA study accessions could also be used to archive screening data, and these would not necessarily have to be linked to single publications—for example, a laboratory or project could create an ENA study titled "Ancient genome screening data from laboratory X", simply deposit all of their screening data there, and provide the study accession number in multiple publications.

withhold the archive study accession number until the paper is accepted, or authors set the study itself to not be publicly visible until the paper is published. Some authors might not be comfortable with making their data publicly available pre-publication, but doing so for the review process would contribute towards better standards in the field. It's also worth noting that this does not need to entail making any findings public pre-publication—the study description texts in ENA can be updated over time, and do not need to correspond to paper abstracts, or necessarily mention any scientific findings from the study at all. In fact, a study description is arguably more useful if it focuses on the technical details of the archived data, rather than simply duplicating the abstract of a published paper.

Conclusions

Shortcomings in how ancient genomics data is archived is likely largely the result of data submitters not being very familiar with how the public archives are structured, and the submission processes being quite complex. Some studies do put substantial efforts into providing comprehensive information on samples, libraries and sequencing experiments in the supplementary materials of papers, or in custom files hosted on institutional or other websites. However, a reliance on inconsistently structured supplementary materials, non-permanent websites, and direct requests to authors, is not a rigorous long-term solution for a data-intensive research field. Rather than treating the task of data archiving as a necessary inconvenience to meet a journal requirement, I would urge researchers in ancient genomics to treat this component of their research projects with the same level of attention and rigour as they treat other components. Building the steps necessary for appropriate archiving into standardised pipelines and workflows could be one way to make this process easier.

During the Human Genome Project, the so-called Bermuda principles prescribed the pre-publication release of genomic data as soon as it was produced. Along with the Fort Lauderdale Agreement that aimed to prevent data producers being scooped on principal findings, this has nurtured a long-standing culture of open and early data sharing in genomics[68]. Upholding these norms, and making appropriate use of existing archive infrastructure, will be increasingly important as ancient genomics grows. The dependence on destructive sampling of finite and sometimes very precious material arguably means that ancient genomics studies have a particularly strong responsibility to ensure that data is appropriately archived and reusable.

Methods

Studies to include in the data survey were identified as follows: I looked for papers published in 2021 or 2022 in the journals Nature, Science and Cell, by searching for "ancient DNA" or "ancient genome" using Google Scholar and the journals' own search tools. I included studies reporting shotgun or genome-wide capture data (but not e.g. only mitochondrial capture data) targeting any multicellular organism, but not studies focussed on microbial or metagenomic targets. This identified 28 studies, but it is possible that I missed some. I downloaded

experimental metadata on the archived data files for each study from the ENA (www.ebi.ac.uk/ena/, Feb 15th, 2023). I obtained associated metadata from the NCBI BioSample database, downloading the full metadata from all BioSample samples linked to the given ENA study / BioProject accession (https://www.ncbi.nlm.nih.gov/biosample/, Feb 15th, 2023). Three studies had archived data in the GSA database[69] rather than in the ENA, and I obtained corresponding information from there (https://ngdc.cncb.ac.cn/gsa/, Feb 15th, 2023). To survey properties of read and read alignment files, I downloaded the first FASTQ and/or the first BAM file listed in the ENA/GSA for each study. If there are differences in properties between files in the same study, this will thus not have been picked up here. One study had deposited raw, unaligned reads in BAM format rather than FASTQ format, and I treated this as equivalent to FASTQ files for reporting purposes. Information was extracted from BAM files using samtools[47]. BAM validation was performed using Picard's ValidateSamFile[70].

Data availability

The data analysed in this study was obtained from the following archive study accessions: PRJEB54831 (ENA)[9], PRJEB51180 (ENA)[10], PRJCA005576 (GSA)[11], PRJEB52849 (ENA)[12], PRJEB42656 (ENA)[13], PRJEB49291 (ENA)[14], PRJEB42781 (ENA)[15], PRJEB46734 (ENA)[16], PRJNA798381 (SRA)[17], PRJEB42269 (ENA)[18], PRJEB47891 (ENA)[19], PRJEB54899 (ENA)[20], PRJEB43715 (ENA)[21], PRJEB39134 (ENA)[22], PRJEB46162 (ENA)[23], PRJEB46875 (ENA)[24], PRJEB44430 (ENA)[25], PRJEB42199 (ENA)[26], PRJEB38555 (ENA)[27], PRJEB55327 (ENA)[28], PRJEB50368 (ENA)[29], PRJEB50857 (ENA)[30], PRJCA003870 (GSA)[31], PRJCA003699 (GSA)[32], PRJEB53475 (ENA)[33], PRJEB37782 (ENA)[34], PRJNA687817 (SRA)[35], PRJEB42372 (ENA)[36].

Acknowledgements

The author thanks James Fellows Yates, Josephine Burgin, Torsten Günther, Thiseas Lamnidis, Kendra Sirak and Pontus Skoglund for constructive comments on the manuscript.

References

1. Anagnostou, P. *et al.* When data sharing gets close to 100%: what human paleogenetics can teach the open science movement. *PLoS One* **10**, e0121409 (2015).
2. Wilkinson, M. D. *et al.* The FAIR Guiding Principles for scientific data management and stewardship. *Sci Data* **3**, 160018 (2016).
3. Toczydlowski, R. H. *et al.* Poor data stewardship will hinder global genetic diversity surveillance. *Proc. Natl. Acad. Sci. U. S. A.* **118**, (2021).
4. Cochrane, G., Karsch-Mizrachi, I., Takagi, T. & International Nucleotide Sequence

Database Collaboration. The International Nucleotide Sequence Database Collaboration. *Nucleic Acids Res.* **44**, D48–50 (2016).

5. Burgin, J. *et al.* The European Nucleotide Archive in 2022. *Nucleic Acids Res.* **51**, D121–D125 (2023).
6. Katz, K. *et al.* The Sequence Read Archive: a decade more of explosive growth. *Nucleic Acids Res.* **50**, D387–D390 (2022).
7. Ogasawara, O., Kodama, Y., Mashima, J., Kosuge, T. & Fujisawa, T. DDBJ Database updates and computational infrastructure enhancement. *Nucleic Acids Res.* **48**, D45–D50 (2020).
8. Barrett, T. *et al.* BioProject and BioSample databases at NCBI: facilitating capture and organization of metadata. *Nucleic Acids Res.* **40**, D57–63 (2012).
9. Lazaridis, I. *et al.* The genetic history of the Southern Arc: A bridge between West Asia and Europe. *Science* **377**, eabm4247 (2022).
10. Liu, Y.-C. *et al.* Ancient DNA reveals five streams of migration into Micronesia and matrilocality in early Pacific seafarers. *Science* **377**, 72–79 (2022).
11. Kumar, V. *et al.* Bronze and Iron Age population movements underlie Xinjiang population history. *Science* **376**, 62–69 (2022).
12. Todd, E. T. *et al.* The genomic history and global expansion of domestic donkeys. *Science* **377**, 1172–1180 (2022).
13. Vernot, B. *et al.* Unearthing Neanderthal population history using nuclear and mitochondrial DNA from cave sediments. *Science* **372**, (2021).
14. Lipson, M. *et al.* Ancient DNA and deep population structure in sub-Saharan African foragers. *Nature* **603**, 290–296 (2022).
15. Wang, C.-C. *et al.* Genomic insights into the formation of human populations in East Asia. *Nature* **591**, 413–419 (2021).
16. Spyrou, M. A. *et al.* The source of the Black Death in fourteenth-century central Eurasia.

Nature **606**, 718–724 (2022).

17. Klunk, J. *et al.* Evolution of immune genes is associated with the Black Death. *Nature* **611**, 312–319 (2022).
18. van der Valk, T. *et al.* Million-year-old DNA sheds light on the genomic history of mammoths. *Nature* **591**, 265–269 (2021).
19. Patterson, N. *et al.* Large-scale migration into Britain during the Middle to Late Bronze Age. *Nature* **601**, 588–594 (2022).
20. Gretzinger, J. *et al.* The Anglo-Saxon migration and the formation of the early English gene pool. *Nature* **610**, 112–119 (2022).
21. Carlhoff, S. *et al.* Genome of a middle Holocene hunter-gatherer from Wallacea. *Nature* **596**, 543–547 (2021).
22. Hajdinjak, M. *et al.* Initial Upper Palaeolithic humans in Europe had recent Neanderthal ancestry. *Nature* **592**, 253–257 (2021).
23. Robbeets, M. *et al.* Triangulation supports agricultural spread of the Transeurasian languages. *Nature* **599**, 616–621 (2021).
24. Zhang, F. *et al.* The genomic origins of the Bronze Age Tarim Basin mummies. *Nature* **599**, 256–261 (2021).
25. Librado, P. *et al.* The origins and spread of domestic horses from the Western Eurasian steppes. *Nature* **598**, 634–640 (2021).
26. Bergström, A. *et al.* Grey wolf genomic history reveals a dual ancestry of dogs. *Nature* **607**, 313–320 (2022).
27. Fernandes, D. M. *et al.* A genetic history of the pre-contact Caribbean. *Nature* **590**, 103–110 (2021).
28. Skov, L. *et al.* Genetic insights into the social organization of Neanderthals. *Nature* **610**, 519–525 (2022).
29. Gnecchi-Ruscone, G. A. *et al.* Ancient genomes reveal origin and rapid trans-Eurasian

migration of 7th century Avar elites. *Cell* **185**, 1402–1413.e21 (2022).

30. Marchi, N. *et al.* The genomic origins of the world's first farmers. *Cell* **185**, 1842–1859.e18 (2022).
31. Wang, T. *et al.* Human population history at the crossroads of East and Southeast Asia since 11,000 years ago. *Cell* **184**, 3829–3841.e21 (2021).
32. Mao, X. *et al.* The deep population history of northern East Asia from the Late Pleistocene to the Holocene. *Cell* **184**, 3256–3266.e13 (2021).
33. Waldman, S. *et al.* Genome-wide data from medieval German Jews show that the Ashkenazi founder event pre-dated the 14th century. *Cell* **185**, 4703–4716.e16 (2022).
34. Clemente, F. *et al.* The genomic history of the Aegean palatial civilizations. *Cell* **184**, 2565–2586.e21 (2021).
35. Liu, S. *et al.* Ancient and modern genomes unravel the evolutionary history of the rhinoceros family. *Cell* **184**, 4874–4885.e16 (2021).
36. Capodiferro, M. R. *et al.* Archaeogenomic distinctiveness of the Isthmo-Colombian area. *Cell* **184**, 1706–1723.e24 (2021).
37. Fellows Yates, J. A. *et al.* Community-curated and standardised metadata of published ancient metagenomic samples with AncientMetagenomeDir. *Sci Data* **8**, 31 (2021).
38. Der Sarkissian, C. *et al.* Ancient Metagenomic Studies: Considerations for the Wider Scientific Community. *mSystems* **6**, e0131521 (2021).
39. Li, H. & Durbin, R. Fast and accurate short read alignment with Burrows-Wheeler transform. *Bioinformatics* **25**, 1754–1760 (2009).
40. Peyrégne, S. *et al.* Nuclear DNA from two early Neandertals reveals 80,000 years of genetic continuity in Europe. *Sci Adv* **5**, eaaw5873 (2019).
41. Prüfer, K. *et al.* The complete genome sequence of a Neanderthal from the Altai Mountains. *Nature* **505**, 43–49 (2014).
42. Günther, T. & Nettelblad, C. The presence and impact of reference bias on population

genomic studies of prehistoric human populations. *PLoS Genet.* **15**, e1008302 (2019).

43. Martiniano, R., Garrison, E., Jones, E. R., Manica, A. & Durbin, R. Removing reference bias and improving indel calling in ancient DNA data analysis by mapping to a sequence variation graph. *Genome Biol.* **21**, 250 (2020).
44. Rasmussen, S. *et al.* Early divergent strains of Yersinia pestis in Eurasia 5,000 years ago. *Cell* **163**, 571–582 (2015).
45. Rascovan, N. *et al.* Emergence and Spread of Basal Lineages of Yersinia pestis during the Neolithic Decline. *Cell* **176**, 295–305.e10 (2019).
46. Kidd, J. M. *et al.* Exome capture from saliva produces high quality genomic and metagenomic data. *BMC Genomics* **15**, 262 (2014).
47. Li, H. *et al.* The Sequence Alignment/Map format and SAMtools. *Bioinformatics* vol. 25 2078–2079 Preprint at https://doi.org/10.1093/bioinformatics/btp352 (2009).
48. Bonfield, J. K. CRAM 3.1: advances in the CRAM file format. *Bioinformatics* **38**, 1497–1503 (2022).
49. Li, H., Ruan, J. & Durbin, R. Mapping short DNA sequencing reads and calling variants using mapping quality scores. *Genome Res.* **18**, 1851–1858 (2008).
50. Mathieson, S. & Mathieson, I. FADS1 and the Timing of Human Adaptation to Agriculture. *Mol. Biol. Evol.* **35**, 2957–2970 (2018).
51. Filippo, C. de, de Filippo, C., Meyer, M. & Prüfer, K. Quantifying and reducing spurious alignments for the analysis of ultra-short ancient DNA sequences. *BMC Biology* vol. 16 Preprint at https://doi.org/10.1186/s12915-018-0581-9 (2018).
52. Meyer, M. & Kircher, M. Illumina sequencing library preparation for highly multiplexed target capture and sequencing. *Cold Spring Harb. Protoc.* **2010**, db.prot5448 (2010).
53. Gansauge, M.-T. & Meyer, M. Single-stranded DNA library preparation for the sequencing of ancient or damaged DNA. *Nat. Protoc.* **8**, 737–748 (2013).
54. Briggs, A. W. *et al.* Removal of deaminated cytosines and detection of in vivo methylation in

ancient DNA. *Nucleic Acids Res.* **38**, e87 (2010).

55. Mungall, C. J., Torniai, C., Gkoutos, G. V., Lewis, S. E. & Haendel, M. A. Uberon, an integrative multi-species anatomy ontology. *Genome Biol.* **13**, R5 (2012).
56. Yilmaz, P. *et al.* Minimum information about a marker gene sequence (MIMARKS) and minimum information about any (x) sequence (MIxS) specifications. *Nat. Biotechnol.* **29**, 415–420 (2011).
57. MINAS: Minimum Information about any Ancient Sequence. https://mixs-minas.org/.
58. INSDC spatiotemporal metadata – minimum standards update (03-03-2023). https://www.insdc.org/news/insdc-spatiotemporal-metadata-minimum-standards-update-03-03-2023/ (2023).
59. European Nucleotide Archive (ENA). Spatiotemporal Metadata Standards. https://ena-docs.readthedocs.io/en/latest/faq/spatiotemporal-metadata.html.
60. Skoglund, P., Storå, J., Götherström, A. & Jakobsson, M. Accurate sex identification of ancient human remains using DNA shotgun sequencing. *J. Archaeol. Sci.* **40**, 4477–4482 (2013).
61. Mittnik, A., Wang, C.-C., Svoboda, J. & Krause, J. A Molecular Approach to the Sexing of the Triple Burial at the Upper Paleolithic Site of Dolní Věstonice. *PLoS One* **11**, e0163019 (2016).
62. Pečnerová, P. *et al.* Genome-Based Sexing Provides Clues about Behavior and Social Structure in the Woolly Mammoth. *Curr. Biol.* **27**, 3505–3510.e3 (2017).
63. Fages, A., Seguin-Orlando, A., Germonpré, M. & Orlando, L. Horse males became over-represented in archaeological assemblages during the Bronze Age. *Journal of Archaeological Science: Reports* **31**, 102364 (2020).
64. Elliott, E., Saupe, T., Thompson, J. E., Robb, J. E. & Scheib, C. L. Sex bias in Neolithic megalithic burials. *American Journal of Biological Anthropology* **180**, 196–206 (2023).
65. Anthony, D. W. & Brown, D. R. The dogs of war: A Bronze Age initiation ritual in the

Russian steppes. *Journal of Anthropological Archaeology* **48**, 134–148 (2017).

66. Kocher, A. *et al.* Ten millennia of hepatitis B virus evolution. *Science* **374**, 182–188 (2021).
67. Mühlemann, B. *et al.* Diverse variola virus (smallpox) strains were widespread in northern Europe in the Viking Age. *Science* **369**, (2020).
68. Kaye, J., Heeney, C., Hawkins, N., de Vries, J. & Boddington, P. Data sharing in genomics — re-shaping scientific practice. *Nat. Rev. Genet.* **10**, 331–335 (2009).
69. Chen, T. *et al.* The Genome Sequence Archive Family: Toward Explosive Data Growth and Diverse Data Types. *Genomics Proteomics Bioinformatics* **19**, 578–583 (2021).
70. Picard. http://broadinstitute.github.io/picard/.