

The Little Book of Smiley Plots

A collection of ancient DNA patterns and their causes

The SPAAM Community

2025-12-08

Table of contents

Preface	1
Contributors and Citation	3
Introduction	5
What are damage patterns	5
How are damage patterns analysed	6
Genomics	7
Metagenomics	7
I Valid Smiley Plots	9
1 Double stranded DNA libraries	13
2 Single stranded DNA Libraries	15
3 Partial UDG (double-stranded)	19
4 Partial UDG (single-stranded)	21
II Half Valid Smiley Plots	25
5 Proofreading enzymes	29
6 Internal Barcode Ligation Bias	33
7 Soft-clipping	35
8 Left over adapter	39
9 Excess of Short Reads	43

III Invalid Smiley Plots	45
10 Insufficient reads	49
11 Internal barcodes not removed	51
References	53

Preface

A key part of any ancient DNA project is to show that the DNA is exactly that - that the DNA is ancient, rather than from modern contamination.

A key authentication method is to show the presence of elevated C to T deamination patterns (and the complementary G to A) at the end of DNA molecules - known as damage patterns - originally reported by (Briggs et al. 2007).

These patterns can be plotted in what have been colloquially known as 'Smiley Plots. However, there can be a wide range of smiley plots, some which show valid ancient DNA, and others that do not - either due to not actually having true ancient DNA but also from laboratory and/or bioinformatic artifacts.

This book aims to act as a reference guide to interpreting ancient DNA damage plots, providing a wide range of example 'smiley plots', with descriptions of what they describe and what can cause them. As an added bit of fun, each type of 'smiley plot' comes with an artistic interpretation of the line shape contributed by members of the ancient DNA community.

Contributors and Citation

The following people contributed ¹ either data, caricatures, descriptions, or other suggestions to the little book.

- James A. Fellows Yates
- Tessa Zeibig
- Aida Andrades Valtueña
- Pete D. Heintzman
- Katerina Guschanski
- Jaelle C. Brealey
- Nihan D. Dagtas
- Zandra Fagernäs
- Patrícia Santos
- Alexander Hübner
- Anne Kathrine Wiborg Runge
- Petra Korlević
- Maja Vukovikj
- Nikos Psonis
- Christina Warinner
- Carl-Eric Wegner

To reference this book, please cite:

Fellows Yates, J. A., Zeibig, T., Andrades Valtueña, A., Heintzmann, P. D., Guschanski, K., Brealey, J. C., Dagtas, N. D., Fagernäs, Z., Santos, P., Hübner, A., Wiborg Runge, A. K., Korlević, P., Vukovikj, M., Psonis, N., & Warinner, C. (2025). *The Little Book of Smiley Plots: A collection of ancient DNA patterns and their causes*. Zenodo. <https://doi.org/10.5281/zenodo.17170743>

¹Flag represents institutional location at the time of contribution

Introduction

What are damage patterns

Damage patterns on ancient DNA molecules occur due to increased miscoding lesions at the end of molecules. When DNA molecules start to decompose (i.e., repair mechanisms are lost once an organism dies), the very long DNA molecules start to fragment due to ‘nicks’ occurring on the sugar-phosphate of one of the strands, weakening the structure and causing the molecule to cleave into two. However, this cleavage is not necessarily ‘clean’, i.e., occurs on both strands at the same position. Rather, when the two uneven ‘nicks’ cause the DNA molecule to cleave into two, this results in a ‘jagged’ break - with the molecules having ‘overhangs’ of one strand being longer than the other of each of the new two now-‘independent’ molecules.

The resulting single-stranded overhangs leave the nitrogenous-bases ‘exposed’ on the overhang to the surrounding environment. In such cases, of the four nucleotides, it was found that cytosines undergo deamination at a higher rate than the others via hydrolysis. The loss of a cytosines amine group results in a nucleotide structure normally found more often in RNA molecules - uracils. The reason why palaeogenomicists report ‘C to T’ damage patterns is because some polymerases will misread uracils as an adenine, and will incorporate a thymine on the opposite strand during DNA amplification. During each subsequent amplification cycle, the mis-incorporated T will propagate across the subsequent copies of the original DNA molecule.

To summarise, the unequal ends of fragmented DNA molecules results in the increased chance of damage to the nucleotides to the overhangs. This structural damage occurs more frequently in cytosines over the other bases, however these ‘damaged’ cytosines are misread by polymerases during DNA amplification to result on thymines on the opposite strand (rather than the expected complementary guanines).

It is important to note that the library construction method will influence damage, e.g., is the library constructed from double-stranded DNA or single-stranded DNA, is the polymerase in the initial library amplification proof-reading or not, and so on. Throughout the rest of this book, each damage

pattern will be described in the context of the library construction method of the data used to generate each damage pattern.

The increased *frequency* to C to T was detected at the end of molecules could only be detected with the invention of ‘Next Generation Sequencing’ or ‘NGS’ (Shendure and Ji 2008). NGS allowed palaeogenomicists to easily sequence thousands to millions of DNA molecules in one go in an untargeted manner, which subsequently meant that molecules from across *entire* genomes could be compared against a reference genome. Bioinformatically, the increase of C to T miscoding lesions were detected by measuring the frequency of mutations at each position across each read, where each read was derived from a different place on the reference genome. As many different places across the reference genome would have different base composition, one would expect to see an approximately random distribution of mutations across the genome. However it was observed in Neanderthals DNA libraries that the frequency of C to T mutation in the first ~10 base pairs of the 5p end of double-stranded library molecules had a higher frequency than the expected approximate equal distribution across each type of mutation, as seen in Figure 1.

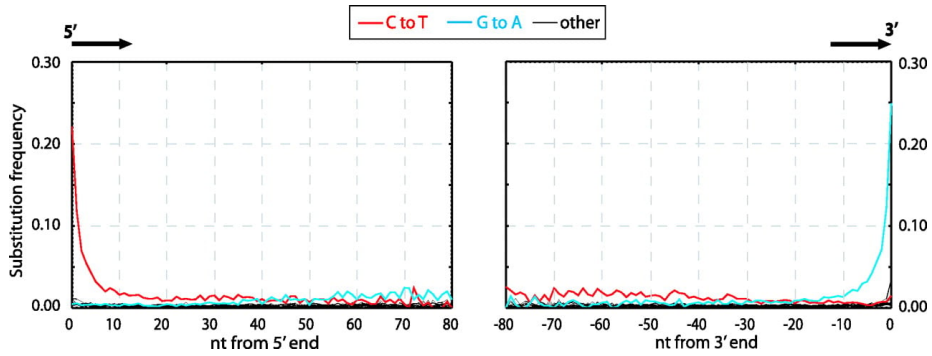


Figure 1: First reported misincorporation lesion ‘smiley plot’ from Neanderthal DNA (Briggs et al. 2007). Reproduced here under free access.

How are damage patterns analysed

There is a range of software that can generate damage pattern plots from ancient DNA NGS libraries. The vast majority of tools require to be of sequencing reads aligned to a reference genome or genomes. Here we make suggestions of some tools that you can use to generate such plots. The example damage patterns in this book will mostly be derived from genomics tools, as metagenomic damage plot generation may account for other factors than the ‘classical’ ancient DNA damage plot.

Genomics

These tools generally take BAM files as input (i.e., after mapping of FASTQ files to a reference genome using a short-read aligner):

- mapDamage
 - Source: <https://github.com/ginolhac/mapDamage>
 - Documentation: <https://ginolhac.github.io/mapDamage>
 - Citation: (Jónsson et al. 2013)
- PMDtools
 - Source: <https://github.com/pontussk/PMDtools>
 - Documentation: <https://github.com/pontussk/PMDtools>
 - Citation: (Skoglund et al. 2014)
- DamageProfiler
 - Source: <https://github.com/Integrative-Transcriptomics/DamageProfiler>
 - Documentation: <https://damageprofiler.readthedocs.io/en/latest/>
 - Citation: (Neukamm, Peltzer, and Nieselt 2021)

Metagenomics

These tools may take different approaches to generating their alignments (or even alignment free methods).

- MaltExtract
 - Source: <https://github.com/rhuebler/MaltExtract>
 - Documentation: <https://github.com/rhuebler/MaltExtract>
 - Citation: (Hübler et al. 2019)
- PyDamage
 - Source: <https://github.com/maxibor/pydamage>
 - Documentation: <https://pydamage.readthedocs.io/en/0.7/>
 - Citation: (Borry et al. 2021)
- MetaDMG
 - Source: <https://github.com/metaDMG-dev/metaDMG-core>
 - Documentation: <https://metadm-g-dev.github.io/metaDMG-core/>
 - Citation: (Michelsen et al. 2022)

Part I

Valid Smiley Plots

This section of the little book of smiley plots shows damage patterns as they *should* be from a molecular biology point of view.

These are the ones that will immediately make *you* smile as these give you good indications of valid ancient DNA in your sequencing library!

Chapter 1

Double stranded DNA libraries



Figure 1.1: Archibald by Tessa Zeibig




Figure 1.2: Example of a smiley plot of a double stranded DNA library. Data taken from library COD076E1bL1 (ERR1943600-ERR1943602) of (Star et al. 2017). Damage data generated using DamageProfiler and plotted using R and tidyverse packages (Wickham et al. 2019).

This is the ‘classical’ ancient DNA plot that you will see most often in palaeogenomics. You expect to see a smooth curve from the beginning of the read (position 1) to a flat line in the middle (e.g. positions 10-25 in mapDamage plots). At the 5’ end this will be indicated by the original C to T deamination, whereas the 3’ of the molecule will show the complementary G to A. You only see deamination the C to T (and complement G to A) at one end of the the molecule, as during typical double-stranded library construction protocols (Meyer and Kircher 2010) only *one* end of the single-ended overhangs of a DNA molecule is repaired by being ‘filled in’ (where the mis-reading of the deaminated C occurs). Overhangs at the other end of the molecule (which may also hold cytosine demination) are ‘blunt-ended’ by being trimmed off. Both fill-in and blunt-ending reactions are performed to allow ligation of next-generation-sequencing adapters and/or internal barcodes to both ends of the molecules. The highest frequency point of the curve can vary from 1% to 50% depending on the age and preservation of the sample.

If you get such a plot with smooth lines from ancient DNA double-stranded libraries, this is a good indication you have authentic ancient DNA!

Chapter 2

Single stranded DNA Libraries




Figure 2.1: Homage to the best moustache ever (Salvador Dalí) by Nihan D. Dagtas. Based upon 'Salvador Dalí' by Philippe Halsman (1953) from wikiart.org. Used under the 'Fair Use' as defined on Wikiart and is used for educational purposes




Figure 2.2: Example of a smiley plot of a double stranded DNA library. Data taken from library VEL003.B0101 of (Andrades Valtueña et al. 2022). Damage data generated using DamageProfiler and plotted using R and tidyverse packages (Wickham et al. 2019).

This is an increasingly common ancient DNA plot that is you will see more often as single-stranded library construction protocols become more popular. You expect to see a smooth curve from the beginning of the read (position 1) to a flat line in the middle (e.g. positions 10-25 in mapDamage plots). As with double-stranded ancient DNA libraries, the 5' end will have the expect original C to T deamination curve. In contrast to the double stranded protocol, the 3' of the molecule will *also* show the complementary C to T curve. You see the C to T deamination at both ends of the the molecule in this case, as during typical single-stranded library construction protocols (e.g. Gansauge et al. 2017) the entire DNA molecule is denatured, immobilised and then the complementary strand reconstructed - i.e., without any filling in or blunt ending. The highest frequency point of the curve can vary from 1% to ???% depending on the age and preservation of the sample.

If you get such a plot with smooth lines from ancient DNA single-stranded libraries, this is a good indication you have authentic ancient DNA!

Chapter 3

Partial UDG (double-stranded)

CARICATURE PLOT GOES HERE

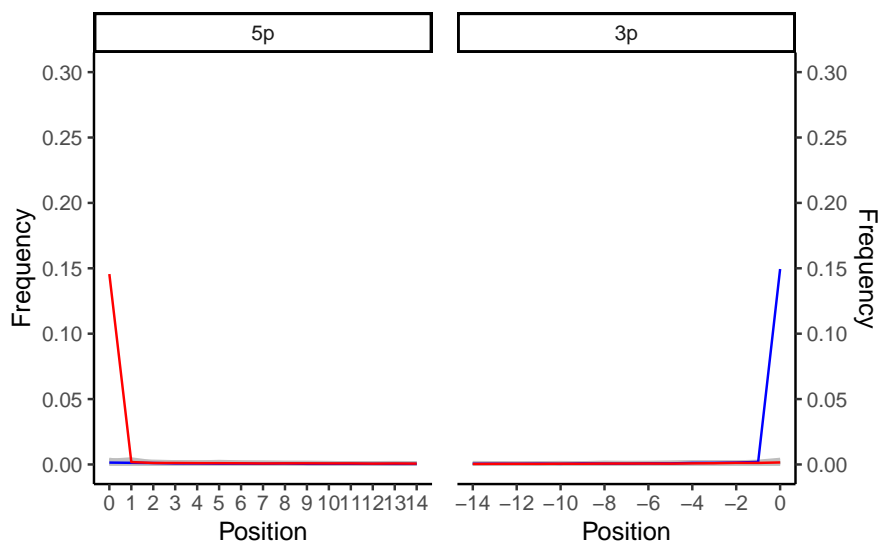


Figure 3.1: Example of a smiley plot of a double-stranded ‘partial UDG’ or ‘UDG half’ library. Data taken from library VLI092.A0101 (ERR8958796) of (Andrades Valtueña et al. 2022). Damage data generated using DamageProfiler and plotted using R and tidyverse packages (Wickham et al. 2019).

The smiley plot presented here has been generated from a double-stranded ‘partial UDG’ or ‘UDG-half’ library. UDG or USER Treatment consists of enzymatically cleaving off ends of molecules when a uracil is present. Whereas complete-UDG UDG full runs the protocol long enough for all DNA molecules to be enzymatically modified, partial UDG stops early, meaning that a small number of reads retain uracils on the last base. This allows for simultaneous authentication of damage, but makes it easier to then *in silico* remove damage by ‘trimming’ off one base from either end of each read, so you can ensure you do not incorporate damaged bases into downstream analyses.

Chapter 4

Partial UDG (single-stranded)



Figure 4.1: Prince Tjorven by Zandra Fagernäs

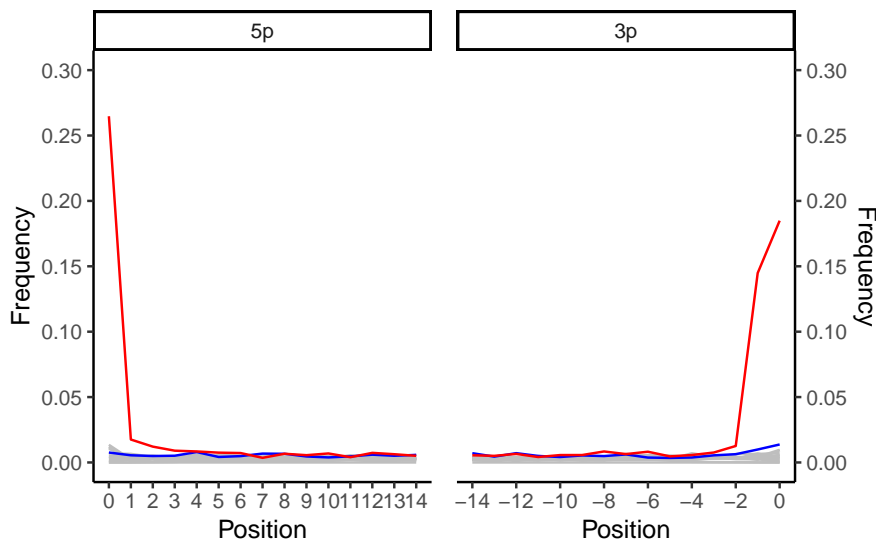


Figure 4.2: Example of a smiley plot of a partial UDG treated single-stranded library, from unpublished data. Damage data generated using DamageProfiler and plotted using R and tidyverse packages (Wickham et al. 2019). Damage-Profiler and plotted using R and tidyverse packages (Wickham et al. 2019).

The smiley plot presented here has been generated from a single-stranded ‘partial UDG’ or ‘UDG-half’ library. UDG or USER Treatment consists of enzymatically cleaving off ends of molecules when a uracil is present. As with single-stranded non-UDG treated libraries, the C to T pattern is retained at *both* ends due to the lack of a blunt-ending (i.e. clipping off of 3’ overhangs) during the library preparation protocol. Furthermore, as with partial-UDG treated libraries, the USER treatment reaction is not run for the whole time, meaning damage on terminal bases are left. This allows for simultaneous authentication of damage, but makes it easier to then *in silico* remove damage by ‘trimming’ off one base from either end of each read, so you can ensure you do not incorporate damaged bases into downstream analyses.

Part II

Half Valid Smiley Plots

This section of the little book of smiley plots shows all the slightly odd damage profiles that don't look as expected, but still retain enough of a pattern that you can take as a positive indicator of ancient DNA.

The descriptions hopefully will provide guidance on how to interpret and when to maybe consider re-doing libraries or mapping.

Chapter 5

Proofreading enzymes




Figure 5.1: Sebastian Sprudelwasser, on his way to a Berlin Rave. By Anne Kathrine Wiborg Runge




Figure 5.2: Example of a smiley plot of a double stranded library amplified after adapter ligation with a proofreading NEBNext Ultra II Q5 Master polymerase. Data taken from library HH06D of (Cai et al. 2022). Damage data generated using DamageProfiler and plotted using R and tidyverse packages (Wickham et al. 2019).

Sometimes the type of polymerase you use during library construction will influence the type of damage pattern you will receive.

In the example above, Cai et al. (2022) found a funny smiley plot, where while the 3' G to A patterns look like a classic non-UDG (i.e. non-USER treated, thus retaining damage) plot, the frequency of C to T on the 5' end was extremely reduced. This is because of the choice of polymerase used during library amplification was not conducive to the nature of ancient DNA.

In this paper, the Q5® High-Fidelity DNA Polymerase was as the polymerase for the initial amplification after adapter ligation. This polymerase has no tolerance for uracils.

During amplification, when it encounters a uracil, it cannot proceed with extending the newly forming complement strand. As the original strand with the uracil will not be copied. Thus the C-to-T deamination at the 5' termini of the reads will be absent in the plot.

In contrast, the original complementary DNA strand (that contains the adenine opposite uracil), which is filled-in by T4 DNA polymerase during the blunt-ending step of the double-stranded library preparation (Briggs et al. 2007), will still be amplified by the Q5 polymerase. This means the G-to-A misincorporation pattern is retained in 3' termini of the reads, even if the damage signal

were lost on the first strand.

While having a low error rate is great for modern genomics, this can be less optimal for preserving ancient DNA damage for profiling later on. Furthermore, you may have lost a fraction of true ancient DNA reads (those that have uracils in both strands).

In the case of this particular enzyme, it may not too much of a problem to prove authenticity visually as you retain the damage signal on the 3'. However this may be problematic for downstream aDNA validation tools that may have an expected 'model' of ancient DNA damage.

 Tip

The choice of enzyme only matters during the first amplification after adapter ligation. At subsequent amplifications, the (misincorporated) thymines have already been integrated into the template molecules, so it doesn't matter which enzyme you use.

Chapter 6

Internal Barcode Ligation Bias

CARICATURE PLOT GOES HERE




Figure 6.1: Example of a smiley plot of an internal barcoded double-stranded library with a ligation bias of certain barcodes. Data taken from sample Ua9 (ERS4545914) of (Brealey et al. 2020). Damage data generated using Damage-Profiler and plotted using R and tidyverse packages (Wickham et al. 2019).

In this smiley plot, you see slightly spiky ends of the damage curves; mainly that some of the last couple of bases are often lower than expected from the rest of the otherwise classical damage curve of a constant decrease. This was observed in Brealey et al. (2020), who associated this to the explanation of reduced ligation efficiency to DNA molecules of internal barcode with a terminal G or C (i.e., short synthetic oligos with known sequences added directly prior addition of library adapters and indices) as suggested by Rohland et al. (2015).

In other words, certain internal barcode sequences with a terminal G or C do not ligate as well to an ancient DNA molecule with a deaminated C on the terminus, and thus those reads will be lost during the second round of demultiplexing as it will not contain the barcode associated with that library.

Ultimately you should not be too worried about this plot if you get it - you probably still have aligned true aDNA reads, however you may have lost a small fraction of true ancient DNA reads in that particular library. If you wish to ensure you have retained as many aDNA reads as possible, you should rebuild the library from an extract but with different internal barcodes without a terminal G or C.

Chapter 7

Soft-clipping

CARICATURE PLOT GOES HERE




Figure 7.1: Example of a smiley plot of soft-clipped ancient DNA data as represented in PyDamage output. Data taken from an unpublished library. Plotted using R and tidyverse packages (Wickham et al. 2019). Note that the smiley plot above is not the ‘typical’ PyDamage output, however it is a simplified version of the ‘C to T transitions’ line in PyDamage plots represented here for illustrative purposes.

This smiley plot can often be seen when using certain short-read mapping settings. In particular researchers using the aligner `bowtie2` with one of the `local` modes will often see 0% damage on the first couple of positions from the end of the read, but then the subsequent frequencies along the remainder of the read will have a ‘typical’ damage pattern curve.

When in `local` mode, the aligner will allow ‘soft-clipping’. Soft-clipping was introduced to aligners when the length of DNA sequencing data increased and alignment issues occurred, e.g. transcriptome data could not be aligned due to the splicing of RNA. Therefore, the aligner gained the ability to keep alignments when only the inner-portion of the read maps optimally to the reference genome. In soft-clipping, the aligner will ‘ignore’ the ends of the reads and not use this information for evaluating the final alignment, however, it will retain those nucleotides in the alignment file. This is opposed to hard-clipping, during which these bases are entirely removed and are therefore ‘lost’ to downstream processes. This is commonly performed in modern DNA studies but can lead to issues in ancient DNA studies.

For example, the aligner may clip off read ends that have damage because it is alignment-wise better than having three consecutive bases that have damage.

In such cases, a researcher can try to use the `global` alignment mode in such such aligners (e.g. with `--sensitive` rather than `--sensitive-local` in `bowtie2`). Otherwise, if the pattern is sufficiently strong (and the alignments are trusted), a researcher can still use the plot and data as long as the pattern of the missing first few bases is described.

Chapter 8

Left over adapter




Figure 8.1: Menada by Maja Vukovikj (sketch converted to graphic with help from ChatGPT)




Figure 8.2: Example of a smiley plot with high frequency of a single last base, due to a remaining base from an insufficiently clipped off adapter (as it was missing from the sequenced given to the adapter clipping tool). Data taken from unpublished data . Damage data generated using DamageProfiler and plotted using R and tidyverse packages (Wickham et al. 2019).

This smiley plot was found from a paired-end sequenced, double-stranded, non-UDG treated library. It displays a very significant G to A frequency increase (exceeding 30%) on the first base of the 3' read termini, as well as a small increase for C to T misincorporations and other base changes. This is likely due to a remaining base from an adapter that was not fully clipped off during the adapter trimming step of read preprocessing.

In the example above, the library was preprocessed with AdapterRemoval (v2.1.7) for adapter trimming and paired-end merging based on sequence overlap. With paired end data where the DNA 'template' is shorter than the number of sequencing cycles (read length), AdapterRemoval first aligns the two reads to each other. It then uses the information from the assumed 'symmetrical overhang' from the alignment, to identify and clip the adapters. In other words, the sequence of the template should exactly match, and thus the rest of the read that does *not* overlap is assumed to be adapter sequence. A limitation of the tool is that asymmetric read pairs (e.g., reads with leading Ns or differing insert sizes that can 'shift' the template overlap) can result in incomplete adapter removal or loss of authentic fragment termini. In this example, a single 3' terminal adapter base remained after adapter trimming due to asymmetric overlap of the templates, leading to an artificial inflation of the damage frequency.

In other tools, this can be cause of the adapter sequence was not correctly specified (e.g., one base left off).

In case you get a similar plot and you still have your raw FASTQ files (and you probably do), you can easily reprocess them with a cleaner approach — e.g. another tool such as fastp (Chen et al. 2018), Cutadapt (Martin 2011) with Flash (Magoč and Salzberg 2011) (or another preferred tool for trimming/merging).

Chapter 9

Excess of Short Reads

CARICATURE PLOT GOES HERE

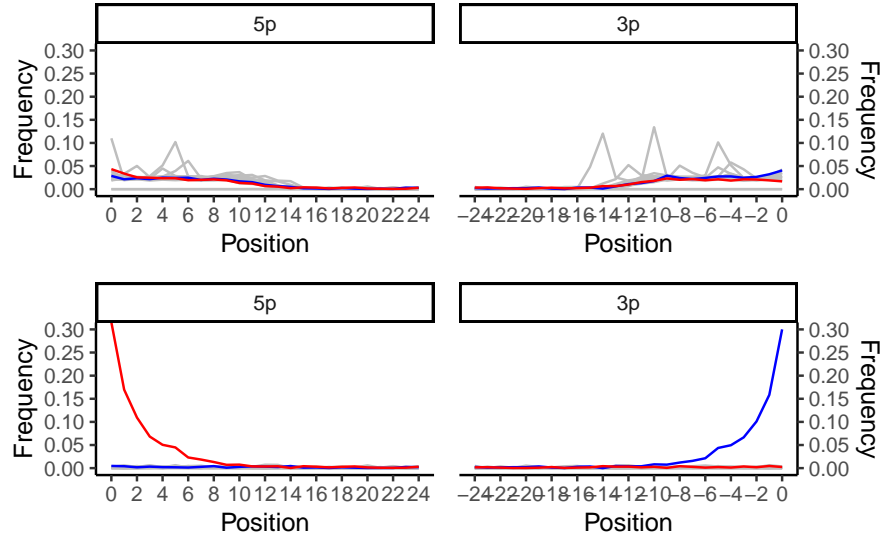


Figure 9.1: Example of a smiley plot with an excess of very short (<30bp) reads suppressing the damage signal (top), whereas removing reads less than 30bp reveals a perfect double-stranded non-UDG treated damage pattern (bottom). Data taken from mapping sequenced reads against the assembled *Cand. Rokubacteria* contigs of sample H22-8 of (Wegner et al. 2023). Damage data was re-generated from the original BAM files using *DamageProfiler* and plotted using R and *tidyverse* packages (Wickham et al. 2019).

In the top smiley plot you see an odd pattern of a set of very spikey background peaks despite smooth C to T/G and A lines, and rather than an exponential decay curve the C to T and G to A lines have a an elevated ‘hump’ between 0-14 bp and then a flat line after 14bp. In contrast, the bottom smiley plot of the same sample shows a perfect double-stranded non-UDG treated library damage pattern.

So what was the difference?!

In this case Wegner et al. (2023) identified the artefact in the top plot as a result of unfiltered ultra-short reads (<30bp) mapping to the assembled contigs of their assembled genomes. This can be seen in the the fact that much of the frequency of the ‘spikey’ misincorporations are concentrated in the first 14bp of the reads, but the remaining of the curve (i.e. in teh longer reads) is nice and smooth. Without sufficient length and mapping quality filtering, these ultra short reads resulted in unspecific mapping distributed across the contigs as well as in some cases in stacking in ‘conserved’ regions of the contigs. After re-filtering to remove all reads below 30bp, the artefact disappeared and a beautiful damage pattern emerged.

Part III

Invalid Smiley Plots

This section of the little book of smiley plots shows all the weird and wonderful *strange* smiley plots that represent problems or artefacts in the sequencing libraries which will make you bemused...

The descriptions hopefully will provide guidance on how to interpret and remedy such problems.

Chapter 10

Insufficient reads




Figure 10.1: Egbert by James A. Fellows Yates

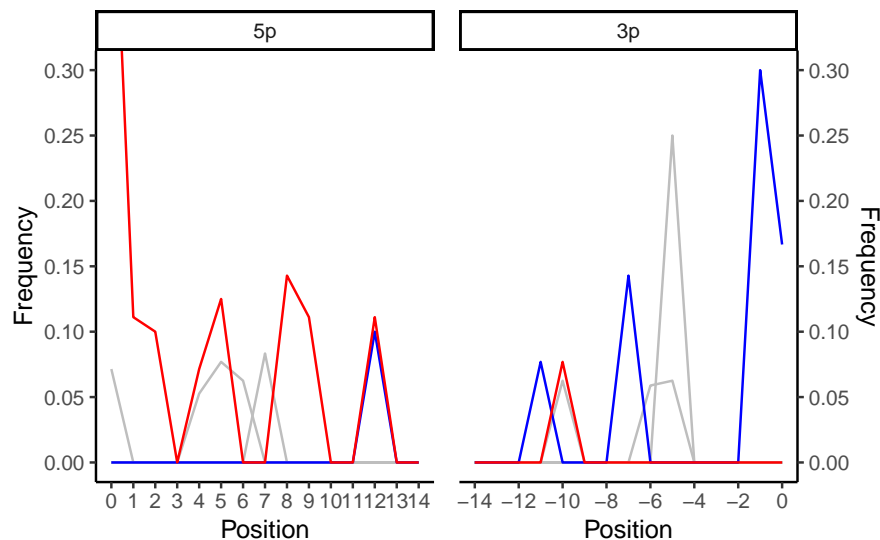


Figure 10.2: Example of a smiley plot of an alignment with insufficient reads to generate a confident smiley plot. Data taken from a non-UDG library of a captured Woolly Mammoth mitochondrial genome (JK2782) from (Fellows Yates et al. 2017), and sampled aligned reads down to 50 reads. Damage data generated using DamageProfiler and plotted using R and tidyverse packages (Wickham et al. 2019).

When you get random spikey lines in both 5p and 3p ends of the smiley plot, this more often than not indicates that insufficient reads are present to generate the damage profile. Given the plots are based on frequency, sufficient numbers of reads are needed to visualise the ‘fraction’ of C to T misincorporations versus the reference, if there are too few reads, this produces ‘noise’ in the line.

In this case of the example above, the aligned DNA *does* have a true damage signal (as indicated by the high frequency of the C-T misincorporations on the 0 and 1 positions of the 5p plot) so *may* give you a teeny-weeny hint of the presence of true ancient DNA. However the rest of line and also the 3p show random spikes making it very difficult to make any firm conclusion.

When you receive a plot like this, you normally need to increase the number of reads in your alignment against the reference genome (deeper sequencing, relaxing alignment parameters), or possibly you have the wrong reference genome (meaning it is not similar enough to align the reads in your library against it).

Chapter 11

Internal barcodes not removed




Figure 11.1: Sasha just can't get any work done while the office keeps blasting 10 different songs at once! By Petra Korlević

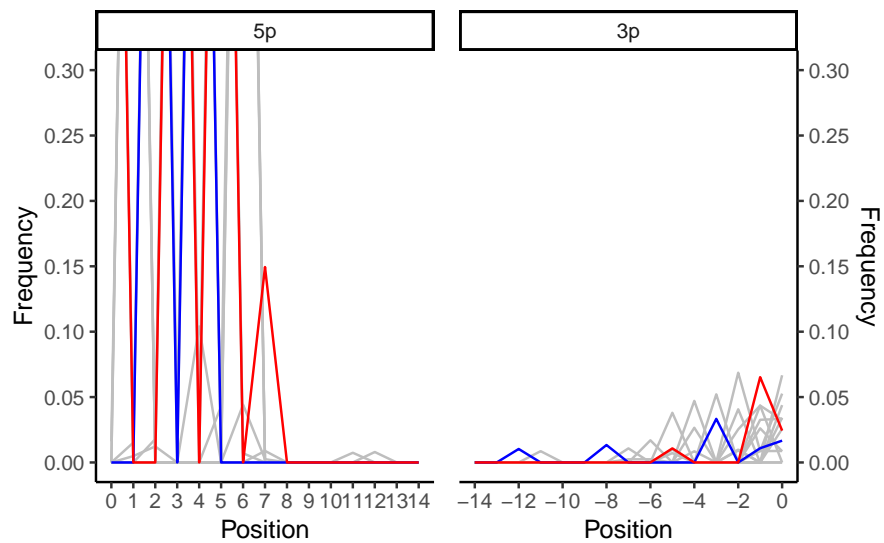


Figure 11.2: Example of a smiley plot of a double stranded library with barcodes that were not removed prior mapping. Data taken from library GEN72 (ERR2112579) of (Andrades Valtueña et al. 2017). Damage data generated using DamageProfiler and plotted using R and tidyverse packages (Wickham et al. 2019).

In this case, this smiley plot is not very smiley as the mapped reads have not had ‘internal barcode’ removed prior mapping. As a side effect, you see a very spiky initial ‘curve’, and then the rest of the read being flat.

In this case, this library only has a single barcode on the 5p end, but some labs may add internal barcodes to both ends of molecules.

References

- Andrades Valtueña, Aida, Alissa Mittnik, Felix M Key, Wolfgang Haak, Raili Allmãe, Andrej Belinskij, Mantas Daubaras, et al. 2017. “The Stone Age Plague and Its Persistence in Eurasia.” *Current Biology: CB* 27 (23): 3683–3691.e8. <https://doi.org/10.1016/j.cub.2017.10.025>.
- Andrades Valtueña, Aida, Gunnar U Neumann, Maria A Spyrou, Lyazzat Musralina, Franziska Aron, Arman Beisenov, Andrey B Belinskiy, et al. 2022. “Stone Age Yersinia Pestis Genomes Shed Light on the Early Evolution, Diversity, and Ecology of Plague.” *Proceedings of the National Academy of Sciences of the United States of America* 119 (17): e2116722119. <https://doi.org/10.1073/pnas.2116722119>.
- Borry, Maxime, Alexander Hübner, Adam B Rohrlach, and Christina Warinner. 2021. “PyDamage: Automated Ancient Damage Identification and Estimation for Contigs in Ancient DNA de Novo Assembly.” *PeerJ* 9 (July): e11845. <https://doi.org/10.7717/peerj.11845>.
- Brealey, Jaelle C, Henrique G Leitão, Tom van der Valk, Wenbo Xu, Katia Bougiouri, Love Dalén, and Katerina Guschanski. 2020. “Dental Calculus as a Tool to Study the Evolution of the Mammalian Oral Microbiome.” *Molecular Biology and Evolution* 37 (10): 3003–22. <https://doi.org/10.1093/molbev/msaa135>.
- Briggs, Adrian W, Udo Stenzel, Philip L F Johnson, Richard E Green, Janet Kelso, Kay Prüfer, Matthias Meyer, et al. 2007. “Patterns of Damage in Genomic DNA Sequences from a Neandertal.” *Proceedings of the National Academy of Sciences of the United States of America* 104 (37): 14616–21. <https://doi.org/10.1073/pnas.0704665104>.
- Cai, Dawei, Siqi Zhu, Mian Gong, Naifan Zhang, Jia Wen, Qiyao Liang, Weilu Sun, et al. 2022. “Radiocarbon and Genomic Evidence for the Survival of Equus Sussemionus Until the Late Holocene.” *eLife* 11 (May). <https://doi.org/10.7554/eLife.73346>.
- Chen, Shifu, Yanqing Zhou, Yaru Chen, and Jia Gu. 2018. “Fastp: An Ultra-Fast All-in-One FASTQ Preprocessor.” *Bioinformatics* 34 (17): i884–90. <https://doi.org/10.1093/bioinformatics/bty560>.
- Fellows Yates, James A, Dorothee G Drucker, Ella Reiter, Simon Heumos, Frido Welker, Susanne C Münzel, Piotr Wojtal, et al. 2017. “Central European Woolly Mammoth Population Dynamics: Insights from Late Pleis-

- tocene Mitochondrial Genomes.” *Scientific Reports* 7 (1): 17714. <https://doi.org/10.1038/s41598-017-17723-1>.
- Gansauge, Marie-Theres, Tobias Gerber, Isabelle Glocke, Petra Korlevic, Laurin Lippik, Sarah Nagel, Lara Maria Riehl, Anna Schmidt, and Matthias Meyer. 2017. “Single-Stranded DNA Library Preparation from Highly Degraded DNA Using T4 DNA Ligase.” *Nucleic Acids Research* 45 (10): e79. <https://doi.org/10.1093/nar/gkx033>.
- Hübner, Ron, Felix M Key, Christina Warinner, Kirsten I Bos, Johannes Krause, and Alexander Herbig. 2019. “HOPS: Automated Detection and Authentication of Pathogen DNA in Archaeological Remains.” *Genome Biology* 20 (1): 280. <https://doi.org/10.1186/s13059-019-1903-0>.
- Jónsson, Hákon, Aurélien Ginolhac, Mikkel Schubert, Philip L F Johnson, and Ludovic Orlando. 2013. “mapDamage2.0: Fast Approximate Bayesian Estimates of Ancient DNA Damage Parameters.” *Bioinformatics* 29 (13): 1682–84. <https://doi.org/10.1093/bioinformatics/btt193>.
- Magoč, Tanja, and Steven L Salzberg. 2011. “FLASH: Fast Length Adjustment of Short Reads to Improve Genome Assemblies.” *Bioinformatics (Oxford, England)* 27 (21): 2957–63. <https://doi.org/10.1093/bioinformatics/btr507>.
- Martin, Marcel. 2011. “Cutadapt Removes Adapter Sequences from High-Throughput Sequencing Reads.” *EMBNet.journal* 17 (1): 10. <https://doi.org/10.14806/ej.17.1.200>.
- Meyer, Matthias, and Martin Kircher. 2010. “Illumina Sequencing Library Preparation for Highly Multiplexed Target Capture and Sequencing.” *Cold Spring Harbor Protocols* 2010 (6): db.prot5448. <https://doi.org/10.1101/pdb.prot5448>.
- Michelsen, Christian, Mikkel Winther Pedersen, Antonio Fernandez-Guerra, Lei Zhao, Troels C Petersen, and Thorfinn Sand Korneliussen. 2022. “MetaDMG – a Fast and Accurate Ancient DNA Damage Toolkit for Metagenomic Data.” *bioRxiv*. <https://doi.org/10.1101/2022.12.06.519264>.
- Neukamm, Judith, Alexander Peltzer, and Kay Nieselt. 2021. “DamageProfiler: Fast Damage Pattern Calculation for Ancient DNA.” *Bioinformatics* 37 (20): 3652–53. <https://doi.org/10.1093/bioinformatics/btab190>.
- Rohland, Nadin, Eadaoin Harney, Swapan Mallick, Susanne Nordenfelt, and David Reich. 2015. “Partial Uracil-DNA-Glycosylase Treatment for Screening of Ancient DNA.” *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences* 370 (1660): 20130624. <https://doi.org/10.1098/rstb.2013.0624>.
- Shendure, Jay, and Hanlee Ji. 2008. “Next-Generation DNA Sequencing.” *Nature Biotechnology* 26 (10): 1135–45. <https://doi.org/10.1038/nbt1486>.
- Skoglund, Pontus, Bernd H Northoff, Michael V Shunkov, Anatoli P Derevianko, Svante Pääbo, Johannes Krause, and Mattias Jakobsson. 2014. “Separating Endogenous Ancient DNA from Modern Day Contamination in a Siberian Neandertal.” *Proceedings of the National Academy of Sciences of the United States of America* 111 (6): 2229–34. <https://doi.org/10.1073/pnas.1318934111>.
- Star, Bastiaan, Sanne Boessenkool, Agata T Gondek, Elena A Nikulina,

- Anne Karin Hufthammer, Christophe Pampoulie, Halvor Knutsen, et al. 2017. “Ancient DNA Reveals the Arctic Origin of Viking Age Cod from Haithabu, Germany.” *Proceedings of the National Academy of Sciences of the United States of America* 114 (34): 9152–57. <https://doi.org/10.1073/pnas.1710186114>.
- Wegner, Carl-Eric, Raphaela Stahl, Irina Velsko, Alex Hübner, Zandra Fagernäs, Christina Warinner, Robert Lehmann, Thomas Ritschel, Kai U Totsche, and Kirsten Küsel. 2023. “A Glimpse of the Paleome in Endolithic Microbial Communities.” *Microbiome* 11 (1): 210. <https://doi.org/10.1186/s40168-023-01647-2>.
- Wickham, Hadley, Mara Averick, Jennifer Bryan, Winston Chang, Lucy McGowan, Romain François, Garrett Golemund, et al. 2019. “Welcome to the Tidyverse.” *Journal of Open Source Software* 4 (43): 1686. <https://doi.org/10.21105/joss.01686>.

