

Long-read sequencing of diagnosis and post-therapy medulloblastoma reveals complex rearrangement patterns and epigenetic signatures

Tobias Rausch^{1,2,18}, Rene Snajder^{3,4,5,18}, Adrien Leger^{6,7}, Milena Simovic⁸, Mădălina Giurgiu^{15,16}, Laura Villacorta², Anton G. Henssen^{9,15,17}, Stefan Fröhling^{10,11,12}, Oliver Stegle^{1,3,13}, Ewan Birney⁶, Marc Jan Bonder^{3,18,19,*}, Aurelie Ernst^{8,18,*}, Jan O. Korbel^{1,6,14,18,*}

Summary

Initial submission:	Received : 2/20/22
	Scientific editor: Rosalind Mott, Laura Zahhn
First round of review:	Number of reviewers: 3 Revision invited : 6/14/22 Revision received : 9/23/22
Second round of review:	Number of reviewers: 3 Accepted : 2/22/23
Data freely available:	yes
Code freely available:	yes

This transparent peer review record is not systematically proofread, type-set, or edited. Special characters, formatting, and equations may fail to render properly. Standard procedural text within the editor's letters has been deleted for the sake of brevity, but all official correspondence specific to the manuscript has been preserved.

Referees' reports, first round of review

Reviewer #1: In Rausch et al. "Long-read sequencing of diagnosis and post-therapy medulloblastoma reveals

complex rearrangement patterns and epigenetic signatures", the authors utilized Oxford Nanopore (ONT) long reads, as well as Illumina short reads to study structural and epigenetic dysregulation in the cancer genome. Particularly, the authors describe a pattern of complex genomic structural variation which they term "templated insertion threads", that are detectable with both ONT and short-read data. To enable the resolution of chromothripsis and the identification of these templated insertion thread events, they present and rely primarily on two methods, Lorax and Rayas, respectively. Additionally the authors also apply the ability of ONT data to detect methylation status and identify differential methylation - including in a haplotype-aware manner. They find that the differential methylation is enriched for medulloblastoma (MB) driver genes in their MB sample. They also use Illumina short-read RNA-seq alongside the haplotype resolved methylation data to study allele-specific patterns of methylation and expression.

The work reflects the cutting edge use of sequencing technologies and sits at the forefront of research in complex structural variation. Particularly the haplotype-resolved analysis of structural variation and methylation signatures, individually and in concert, will enable important biological discoveries. There is an impressive amount of new material and methods presented, and the paper represents an important proof-of-concept for the field. However, there are simultaneously concerns that this reviewer has about several aspects of the work which warrant additional review and revision before publication.

Major issues:

1. In the section "Results: Haplotype-phased assembly of complex somatic rearrangements." The section concludes by speculating about why the HSR integration points could not be found. This reviewer agrees that ONT read length limitations may be one reason, however is it possible that heterogeneity of the HSR integration from cell to cell is a factor? If integration heterogeneity were abundant, it would perhaps also be very difficult to find consistent read-support anchoring HSRs into chromosomes. Moreover, in Figure 1e FISH image iii, the image caption states "colocalization of the centromere 17 probe with the RP11-651L9 probe", however the colocalization statement may apply to one of the red blots near the centromere, but there is a large diffuse pattern elsewhere which may suggest either extrachromosomal form or an extremely diverse integration form elsewhere. There is some concern by this reviewer with regards to exactly how much diversity there is in co-localization of the two guides from these two images alone. In what fraction of image was at least one 651L9 probe **not** co-localized with cen17?

2. The manuscript states the templated insertion thread presented in Figure 2a (primary tumor) was not detected again in the relapse. Were the majority of SV events present in primary found again in the relapse (suggestive of the relapse arising from a subclone of the primary)? If most all other SVs were conserved between primary and relapse, but this templated insertion thread was not, can the authors describe why it may have not reappeared? It is unclear whether we should have expected to see it in the relapse or not from how the data is presented in the manuscript.

3. As the authors describe in the introduction section - these templated insertion threads appear to be distinct from TST-jumps (Umbreit et al.) and similar mechanisms proposed by others (Li et al.), mostly by virtue of the increased copy number. If the samples with the other, similar, modes of insertion-like jumps are given to Rayas, how does it classify them? Specifically, if Rayas were run on the sequencing data generated by Umbreit et al. in the renal cell carcinoma sample described in the discussion section, would it discriminate between TST and templated insertion threads? Are the patterns described in Umbreit et al. perhaps an 'early' form of the more dramatic phenomenon described in this paper? Similarly, if appropriate thresholds are selected for running Rayas, how do the regions called as templated insertion threads by Rayas match with what was reported by Li et al. using PCAWG data?

This reviewer would like to know, definitively, if the authors are proposing that the templated insertion threads described in this paper are arising by the same mechanisms proposed in the prior works (albeit perhaps having far more mechanistic iterations to increase copy number)? Otherwise, can the authors clearly demonstrate/argue/propose the phenomenon presented in their paper arises by a different mechanism than TST?

4. To what extent does sequence homology exist along the junctions between segments and in self-to-self connections for templated insertion threads? Is there microhomology and is it greater/less than expected? This may help to reveal some of the biological signatures by which these rearrangements form.

5. The authors note that there was not a “significant relationship between ASM and proximity to somatic variants” (Results - “Allele specific methylation and expression”), and speculate that much of the ASM was due to germline variation. Does this imply it is necessary to also capture tissue of origin matched-normal to perform robust detection of ASM patterns in the context of somatic variants? Can the authors elaborate on what procedure might be done to help assist with the discovery of ASM that **are** linked to somatic variation?
6. This reviewer feels that additional validation of the Lorax method is warranted. Most genome-reconstruction tools are tested on multiple samples and/or simulated data, and then validated with an orthogonal technology (e.g. some form of cytogenetic validation, optical mapping, gel-based methods, etc.). However in this paper it appears it is only deployed on a single sample. Can the method be tested on other publicly available datasets? Does it support data from other long-read technologies, such as PacBio?
7. While Lorax simplifies the assembly of amplified chromothripsis regions by performing a targeted reconstruction, can the same conclusion/structure found by Lorax regarding the 1.55 Mbp event be found with other long-read based assembly methods (albeit perhaps less efficiently)? Such a demonstration may help convince the readers and this reviewer that the chromothripsis event reconstructed by Lorax in this case is both correct and complete.
8. Given that at least one of the telomere SV events was resolvable only using the T2T assembly, can they authors clarify in discussion to what extent they expect the more completely resolved reference genome to improve detection of templated insertion threads? It is hinted at in the discussion section - but can the authors clarify further? Does the finding presented in discussion about three of five SV telomere junctions being unresolvable with short reads even using T2T imply there exists a limited ability to improve SV calls on short read data using the T2T build, as those newly added/improved regions are too low complexity to map uniquely with short reads anyways? This reviewer feels that point deserves a more comprehensive discussion.
9. Reporting some runtime and resource requirement benchmarks on the analyzed cancer samples would be helpful for those considering using Lorax and Rayas.

Minor issues:

- a) The last sentence of the Summary seems vague or incomplete. “Our study shows the potential of long-read sequencing in cancer.” Can it be written more specifically?
- b) A clearer description of the haplotype phasing approach should be reported in the main text, to make it more clear which existing haplotyping tools are used to generate the phased blocks (Whatsmap) and how they are refined (ShapeIt).
- c) Please consider ordering the supplemental figures so they match their appearance in the main text for readability. E.g Figure S1 is not cited until well after many of the other supplemental figures. Same for Figure S3, and so on.
- d) Figures 1c, please more clearly indicate the chromosome number(s) associated with the coordinates on the x-axis. It took this reviewer a long time to determine that the chromosome cartoons from panel D matched the coordinates shown in panel C.
- e) In Figure 2A,B it is not clear which chromosomes the templated insertion coordinates match to in this plot. Are the genome regions shown in the top part of B matched directly to the genome regions shown in the bottom part of B? Placing the legend between them and not labeling chromosome names nearby coordinates on the x-axis makes it unclear. Matching the colors of the flanking regions from panel 2A to the colors of the highlighted genome regions shown in 2C may help clarify the relationship of panels A,B,C more naturally.
- f) Please correct grammatical errors throughout manuscript and figure captions (e.g, Fig 2c caption), and address tense issues, e.g. Results section, Haplotype-phased assembly of complex somatic rearrangements: “In the primary tumor, we find 697 somatic SVs,…” and also elsewhere “we identify”, “we separate” should be written in past tense so it is consistent with the rest of the tense in the section.

- g) In Rayas, what is the significance of filtering “connected subgraphs where all nodes of G_s are nearby in the genome”, and how was the default threshold of 10kbp decided upon? This reviewer is not asking for a rigorous justification of 10kbp, but rather a simple motivation and explanation for what kinds of events this filters out since it is not clear why this filter exists and how it affects discovery of templated insertion threads.
- h) Ensure all acronyms are defined. Haplotype “HP1” and “HP2” acronyms not explicitly defined. “DMR” not defined until Methods, after it appears many times.
- i) Figure S14, S19, consider including scale bars for the heatmaps.
- j) The use of the term “high-complexity regions” in the discussion section means something different than “high complexity” and “low complexity” with respect to sequence composition as described in the rest of the paper. Consider saying “regions of high structural rearrangement complexity” or similar to in the Discussion section to differentiate the two uses of the term.

The boost library appears to be required to compile Rayas and Lorax but this or any other compile requirements are not listed on the READMEs. A testing dataset would also be useful for assessing if a Rayas installation is working.

Reviewer #2: Rausch and colleagues present a genomic case report of a ONT sequenced medulloblastoma primary tumor and relapse along with its matched normal. They use the longer reads to assemble complex rearrangements including an interesting self-concatenation pattern that they dub a “templated insertion thread”. In addition to SV calling they use the ONT reads to highlight several allele-specific methylation patterns.

Overall, the study provides a somewhat limited data resource and anecdotal findings of unclear significance without a clear methodological advance. The templated insertion pattern while interesting is not pursued to sufficient technical or biological depth. Since the long read analysis is of a single case, the highlighted findings are inevitably cherry-picked and their biological or clinical significance is not apparent. Notably, the primary and relapse comparison particularly for methylation yields little insight into the genetics or epigenetics of medulloblastoma relapse. Finally, new algorithms are presented without benchmarking or comparison to previous tools.

Major critiques:

- The templated insertion thread pattern seems to be the most interesting finding in the study but is not pursued to sufficient technical or biological depth. The long read evidence for this (like for many of the other findings) comes from a single case, and actually a single sample since it is absent from the relapse. While intriguing, it is unclear how generalizable this pattern is. A possibility that needs to be considered is whether this is just a technical artifact.

* Limited correlation is done of long and short read data around the proposed templated insertion thread. However it seems like the rearrangements in the templated insertion thread should be detectable by short reads. Areas of focal increased sequencing depth like shown in Fig 2b are frequent in regions of the genome for example due to low mappability. How do the locations of the read depth increase correspond to the locations of the reconstructed thread? Do they match in quantity? Are they absent in the relapse and the normal? These basic analyses should be presented.

* The data shown in Fig 2 shows a single long read that is 50 Kb in size. Similarly the supplementary figures show single ONT reads each with distinct patterns. Do these distinct reads point to a single rearranged templated thread consensus sequence, e.g. through assembly or other integration approaches?

* Also the length of reads supporting the pattern appear to be 10 fold longer than the median read length in the library (5 kb). Is that the case for the other supporting reads and is that a problem?

* Are the templated thread adjacencies associated with the long reads found in short reads? It is unclear whether these patterns are happening at short read unmappable sequences. If they are mappable, they should be detectable in short reads - eg by aligning short reads to the long read-derived contigs.

* The pattern seems to very strongly resemble genomic shards, noted with BAC sequencing in 2007 (Bignell et al), however this paper is not referenced or discussed.

* The threading also seems to resemble the typhoons which were also noted in liposarcoma in 2020 (Hadi et al). These were also notable for a high level of tandem and fold back rearranged copies concentrated in an HSR, though maybe on a larger genomic scale (megabase instead of kilobase). How do these patterns relate, both biologically and in their distribution?

* The authors propose a new algorithm to nominate this pattern in short reads, but do not validate these predictions with long reads in any other samples. The algorithm they described in the supplement would seem to detect threads associated with many classes of complex SVs not just the highly self-similar patterns highlighted in Fig 2A. It would seem essential to get a sense of specificity, meaning how often does a pattern suggested to be a templated thread in short reads actually give you the long read pattern. Again insight into this is limited given the small long read dataset.

* The splicing analysis around the templated insertion thread is confusing since the primary tumor which harbors the SVs is the sample that lacks any aberrant splice isoforms. As a result it seems that the genomic change is either functionally inert, very subclonal, and/or not chromosomally integrated, as well as artifactual.

* Related to above, what evidence is there that these patterns are clonal and chromosomal? They are absent in the relapse, so perhaps they might be a transient and extrachromosomal amplification process like what has been noted by Jonathan Whetstone's group.

* Otherwise how do the authors explain the disappearance of the pattern in the relapse. Is there a deletion of this locus in the relapse, or did the relapse arise from a distant ancestor of the primary?

- Beyond the templated insertion finding, it is unclear what qualitative insight the long reads give here. In fact it seems that many of the findings would be evident using short reads.

* For example chromothripsis alleles like in 1C can be inferred from short reads through the use of allelic imbalance (eg RCK Aganezov Raphael Genome Research 2020) though potentially with some ambiguity ie multiple possible reconstructions.

* The comparison to short reads should also not just rely on a discordant pair based tool like DELLY, which only gives approximate breakpoints and limited resolution of complex regions. An assembly based algorithm, like GRIDSS or SVABA should work better as a comparison to short reads.

* While long reads can provide direct evidence to thread distant rearrangements, the relatively short long reads in these data (5kb) could only do this for nearby rearrangements ie those that are within 5kbp of each other. Given this limitation, the authors appear to be taking a more indirect / greedy approach to reconstruct the larger alleles eg those associated with chromothripsis in Fig 1, which though reasonable could be done just as well with short reads to the same effect. It appears therefore that this feature of the long read data and analysis are overstated.

* what is the broader relevance of these patterns to medulloblastoma? The PCAWG analysis does not seem to include a substantial number of medulloblastoma cases.

- As a data resource, the study is limited with one case and two tumor samples. Also the data quality is somewhat questionable, with read lengths that are somewhat short (5 kbp) even for ONT which is usually 10-50 Kbp even in the lower cost / high throughput mode. Larger and seemingly higher quality cancer long read datasets have been published over the past several years eg in genome research eg Nattstad 2018, aganezov..schatz 2020, sakamoto 2020.

- Given that there is one case profiled here, the authors are limited to cherry-picking anecdotal observations without statistical basis. As a result, there appears to be limited insight that can be drawn from these data into medulloblastoma or more generally cancer development and post-therapy relapse. It's unclear in particular what additional insight is drawn from the relapse sample. This includes the methylation findings, which are of uncertain significance and mostly validate previous findings about ASM that did not require long reads.

- The provided tools for amplicon assembly and templated rearrangements seem heuristic and do not show a clear advance over existing graph tools for SV analysis, including RCK (Aganezov), Linx (Shale et al 2022), JABBA (Hadi 2020), AmpliconArchitect (Deshpande 2019), and AmpliconReconstructor (Luebeck 2020). Also

it appears that the algorithm greedily looks for a single thread through the graph rather than trying to account for all the copies of segments in a region. This suggests that the authors are not considering alternate reconstructions, and thus possibly overstating their confidence in a given reconstruction.

Technical and minor comments:

- Comparison of primary and relapse in general does not seem to yield significant biology, but appears at least in part limited by the lower sequencing depth in the relapse. Unclear what is the significance of the observation that templated insertion threads are missing in the relapse? Were these alleles lost or just missed due to low read depth or perhaps because the samples were analyzed for SVs separately. One approach to integrate analyses of the two samples is to align the relapse reads to the primary tumor contigs, and vice versa.
- On page 12 "Since the CS11-17 rearrangement occurs in only one haplotype" .. is this a premise or a result of the analysis? ie are the parental haplotypes long enough to make this statement from the data or are the authors just applying a biological assumption. This should be clarified.
- Contig 2 is mysterious - it is somehow associated with the CS11-17 but does not seem to be rearranged (Fig 4)? Are there really no rearrangements around this segment? How is it known that it is associated with CS11-17 and not part of an independent event?
- Copy number of many segments in Fig 2B appears to be at 3.5. Is this a subclonal copy number change or due to a bad fit of purity/ploidy ?
- It is unclear whether the telomere fusions represent the connection of a telomere sequence to a chromosomal sequence, creating a new chromosome end, or the end to end fusion of two distinct telomeres. Usually "telomere fusions" refer to the latter pattern. It seems also that the former outcome should be also detectable with short reads?
- The observation of hypomethylation of contig 2 (Fig 4) is potentially interesting. Unclear whether this is the consequence of rearrangement or just a pre-existing aspect of this haplotype ie prior to cancer development. Since contig 2 does not appear to be connected to anything it is unclear
- ASM analysis (p11) does not seem to distinguish between promoter-associated CpG islands vs shore or open sea CpGs. How do the authors explain allele specific methylation outside of CGIs ? Also the comparison between primary and relapse seems to be fraught due to low coverage in the relapse (Fig 3A). The allele specific methylation changes do not seem to be associated with a consistent impact on expression.
- For Fig 4A the two regions are different
- How much more of the genome is mappable with 5 kb ONT reads relative to standard illumina 150 bp reads? The noisiness of ONT reads including the high rate of gaps makes them harder to map for a given length, so I wonder if there are many more rearrangements caught in repetitive regions for instance.
- Comparisons are made to bisulphite sequencing (eg p11) but no direct comparison to WGBS profiling appears to be done.
- p9 comparison to arrays is a bit of a straw man? "Overall, analysis of the ONT data provides a substantially more comprehensive picture of the tumor methylome, with 78% of the between sample DMRs inaccessible to the commonly used 450K array, and 65% inaccessible to the 850K array". It seems that a very similar statement could be made for WGBS.
- Fig 1C is an unnecessarily awkward and confusing visualization of a simple concept, which is the alignment of a multi-part contig to the reference genome. In particular, the placement of the purple / orange stacked bars on the x axis at 55.2 M is confusing. It appears that this should be placed outside of the x axis ie the x axis should be split into two axes. However it would seem better to replace this visualization with something more like Fig S9.
- The approach to coloring and patterning in figures seems suboptimal and in many cases collides across different figure elements which is confusing - examples include 4A, 4C, 4D. in particular the use of diagonal patterns to denote haplotypes eg in Fig 3 and 4 seems unnecessary, it would seem better to just use text labels.

Reviewer #3: Comments enter in this field will be shared with the author; your identity will remain anonymous. In their manuscript "Long-read sequencing of diagnosis and post-therapy medulloblastoma reveals complex rearrangement patterns and epigenetic signatures", Rausch et al. demonstrate the promise of long read sequencing for characterizing complex rearrangements patterns in tumor genomes. Specifically, the authors apply Oxford Nanopore Technologies nanopore sequencing to a primary medulloblastoma, its relapse and a matched normal blood sample from a patient with Li-Fraumeni syndrome. They leverage the long reads to obtain extensive phasing information and resolve the structure of various complex somatic rearrangements, most notably a novel type of pattern dubbed "templated insertion thread". By identifying its footprints in the short-read based PCAWG cohort, the authors demonstrate it is often associated with chromothripsis and highly frequent in liposarcoma. The authors also explore the DNA methylation patterns encoded by the single molecule nanopore reads to reveal differential and allele-specific methylation patterns, as well as integrate further datasets of short-read gene expression data.

The work is extremely comprehensive, and the paper reads fluently. The authors have gone through great lengths to ensure the manuscript and its associated methods represent a solid basis for both the long read and cancer genomics communities to build on. To this end, they have developed several new computational methods: (i) rayas which detects signatures of templated insertion threads in short-read sequencing data; (ii) wally, a package for visualization of read alignments from BAM/CRAM files, most notably at and around (structural) variants; (iii) lorax, a toolkit for analysis of long read data, specifically aimed at cancer genomics.

All of these tools, plus a set of further analysis scripts used in the manuscript, are versioned, well documented and made available on GitHub. For the users' convenience, precompiled binaries, as well as Dockerfiles and Singularity containers are provided, in what this reviewer considers a truly impactful and exemplary contribution to the scientific community.

I have no major concerns, but would like to see a few points extended or addressed in a revised version of the manuscript:

1) As indicated by the authors, nanopore sequencing coverage, especially for the normal and relapse samples is relatively low, and the median read length is somewhat limited compared to recent standards. Likewise, more recent base calling algorithms such as the super accuracy models incorporated in Guppy v5 and 6 could further improve analysis results (the authors use Guppy v4 and a high accuracy model). While I do not expect the authors to acquire additional sequencing data or rebasecall their raw data and redo the full analysis, these points could be included in the discussion to guide people in the field on what to aim for and what to expect when deciding to pursue long read sequencing of tumor genomes.

2) The study includes both a primary and matched relapse sample, with notable differences in copy number and structural variation (e.g. absence of the templated insertion threads from the relapse), potentially offering insights into the tumor's evolutionary history as well as the mechanisms and ordering of some of the rearrangements. These aspects are only briefly touched upon in the manuscript however.

- For instance, it would be informative to assess which somatic rearrangements are shared between the primary and relapse. Are there any such common variants within the chromothripsis regions on chromosomes 4,5,7,9,16,19 and X?

- Similarly, is the chr5p telomere fusion present or absent from the relapse? This could provide insights into the cause of this chromothripsis event.

- The authors also note potential ongoing evolution and rearrangements of the templated insertion threads themselves. To assess its stability, would it be possible to estimate how many different alleles are currently detected in this single bulk sample?

3) It may be helpful to provide a read length histogram and read length N50 for the nanopore sequenced samples (in addition to the stats in Table S1). This would further allow people considering long read experiments to readily compare data requirements and quality for the type of analysis presented.

4) Do the authors have a good estimate (either through analysis of the sequencing data or from pathology) of tumor purity? This would allow one to assess more accurately potential effects of normal cell admixture on the differential methylation and expression results.

5) From Figure 1A it appears several large CNVs, such as the losses on chr2 & 3, are not associated with a "large" SV. Were these not identified from the data or are they simply filtered in the display? In either case,

the authors may wish to elaborate.

6) Authors state: "The CS11-17 segment, present in both primary tumor and relapse, has a size of 1.55 Mbp; the 17p-arm region affected contains the TP53 locus, which has been lost on the chromothriptic haplotype." Authors may wish to clarify that it is the wild-type TP53 allele which was lost from the chromothriptic haplotype.

7) At the end of the section "Haplotype-phased assembly of complex somatic rearrangements", the authors note the following: "Yet, we failed to identify reads supporting reintegration of this structure into a chromosomal context, possibly due to limitations of ONT for resolving low-variant allele frequency SVs in conjunction with ITH, especially in complex regions that exhibit repetitive segments larger than the ONT read length". Is it not more likely that the limited coverage read lengths here are to blame for this?

8) In Figure 2B, it would be helpful if the top and bottom panels were (better) aligned.

9) Figure 1B shows non-integer copy number for the first segments on chr5. The authors may wish to elaborate on this (and other) subclonal copy number changes.

10) In the section "Pan-cancer landscape of templated insertion threads in 2,569 tumors.", the new rearrangement pattern is frequently referred to as "template insertion threads" rather than templated insertion threads.

11) When exploring telomeres, authors note that "Another telomere crisis event observed in the primary tumor likely fused chromosome 19 to the telomere of chromosome 16q". I don't think the presence of this telomere fusion necessarily implies a telomere crisis. It could simply derive from a subtelomeric dsDNA break and fusion.

12) Could the differential methylation calling be confounded by purity and/or copy number differences between the two tumor samples?

13) While generally carefully phrased, in parts of the manuscript where several specific genes are highlighted this can start to feel like cherry-picking. For instance, in the latter part of the differential methylation section, authors report finding >26,000 differentially methylated CpGs and note two specific genes, NRN1 and PTCH1. Further individual genes are highlighted in nearly all subsequent sections.

14) When resolving expression effects using ONT data, the authors leverage two distinct absolute log₂-fold change thresholds, >5 and >2, the latter set to compare to the promoter-linked differentially methylated regions. It is unclear why two thresholds are being used and what the total number of differentially expressed genes is under the weaker threshold (without exploring Table S7).

15) In line with the above, what fraction of the differentially expressed genes is likely due to being located on a segment with a copy number change between the two tumors?

16) Figure S19, while there is likely to be a negative correlation, the regression results in a poor fit and should be removed from the panel.

17) Are the methylation changes in the two promoter regions of TBX1 associated with any changes in transcript/isoform usage as well as total expression?

18) How were regions with loss of heterozygosity treated in the allele-specific methylation and expression analyses?

19) More general, what fraction of the allele-specific expression calls is due to allelic imbalance in copy number?

20) The following phrase is unclear to this reviewer: "... and ASM is increased in cancer, caused by disease associated regulatory SNPs"

21) Authors may wish to visualize the differential splicing pattern described in BASP1.

22) Authors highlight overexpression of CCND3 and BYSL in liposarcoma DO219967 in the PCAWG data set and suggest a nearby templated insertion thread may be driving this overexpression. Could the authors

exclude gene copy number as a driver for overexpression here?

23) MYPOP is mentioned to be subclonally amplified in the haplotype where the telomere associated SV is observed. Could the authors provide support for this subclonal event?

24) Could the tandem short template jump pattern recently described by Umbreit et al. represent a precursor event to the longer chained rearrangements here? That is, could genetic instability at the locus lead to an increase in copy number and further chaining as observed in the templated insertion threads?

25) For the DNA and RNA extractions, authors should further specify which "Qiagen kits" were used. Likewise, the "Illumina platform" used during RNA sequencing could be clarified.

26) In the Methods section "Gene fusion and validation using DNA long reads.", the hyphen can be dropped from "set-out".

Authors' response to the first round of review

General comments for all reviewers:

We would like to thank the three reviewers for their interest in our work. Reviewers one and three were generally very positive about our work. For example, Reviewer #1, writes: "The work reflects the cutting edge use of sequencing technologies and sits at the forefront of research in complex structural variation." & "There is an impressive amount of new material and methods presented, and the paper represents an important proof-of-concept for the field." Reviewer #3 writes: "The work is extremely comprehensive, and the paper reads fluently. The authors have gone through great lengths to ensure the manuscript and its associated methods represent a solid basis for both the long read and cancer genomics communities to build on." Reviewer #2 acknowledges that our long read based complex rearrangement assemblies include an "interesting/intriguing self-concatenation pattern" (the "templated insertion thread"). We summarize in the following some general revisions/additions to the manuscript included in response to the reviewers' comments, which have further strengthened our manuscript and the biological findings presented:

1. We generated additional long-read data and performed additional analyses of our longread dataset at paper revision stage, further substantiating our finding of the templated insertion thread (TI thread) characterized by complex self- and cross-concatenated DNA rearrangements, as requested by reviewer one and two. For example, we generated Oxford Nanopore Technology (ONT) long read sequence data on two liposarcoma samples from the NCT/DKTK Master cohort, following identification of the TI thread in previously generated short read (Illumina) datasets – and demonstrate validation of the presence of TI threads in the respective ONT datasets in both cases (new Supplementary Figure S25, S26). We also provide additional supporting evidence for the originally discovered TI thread in the medulloblastoma case using newly generated patient derived xenograft (PDX) data (new Figure S11). In addition, we updated our review of previous data from Umbreit et al. (2020) and Bignell et al. (2007), who report SV signatures that show clear differences from TI threads in terms of rearrangement complexity and segment copy number. In further support of our findings we also applied two assembly algorithms (Flye and Shasta) yielding contigs that recapitulated the TI thread computationally. Moreover, as requested by one of the reviewers, we applied short-read based amplicon reconstruction tools (Linx and RCK), which allow us to better highlight the unique benefits of long reads for resolving cancer genome structure.

2. In response to reviewer 2 and in an effort to validate as well as increase our understanding of TI threads and the assembled CS11-17 complex structure, we initiated a collaboration

with experts in extrachromosomal DNA (ecDNA). We generated Circle-seq data, with the aim to detect and characterize ecDNA structures – as well as to assess the potential overlap between the novel TI thread pattern and the other complex rearrangement structure (CS11-17) with circular structures. Importantly, we now provide evidence that the CS11-17 structure involved circular intermediates. The TI threads, by comparison, were not supported by Circle-seq reads, providing no evidence for their existence in the form of circular DNA structures. We performed additional FISH experiments showing the co-localization of probes specific for the CS11-17 on structures resembling marker chromosomes or ring chromosomes.

3. As requested by reviewer #2, we performed comprehensive benchmarking studies to compare complex SV reconstruction between ONT and short read (Illumina) sequence data – which underscore the advantages of ONT-based cancer genome analysis. We also performed a deeper comparison of the genomic variation between the primary tumor and the relapse sample. Additionally, we decided to re-base call and re-methylation call the ONT data, and given the rapid pace of development of base calling methods we obtained a marked increase of the read quality PHRED score (~ 2.64). As a result, de-multiplexing (11% increase in usable reads), variant calling, phasing and allele specific methylation analyses have improved, although none of the major conclusions of the work changed (e.g. illustrated by plotting read phase-ability rebuttal Fig. R1.1).

We provide a more detailed point-by-point response to all reviewer questions below.

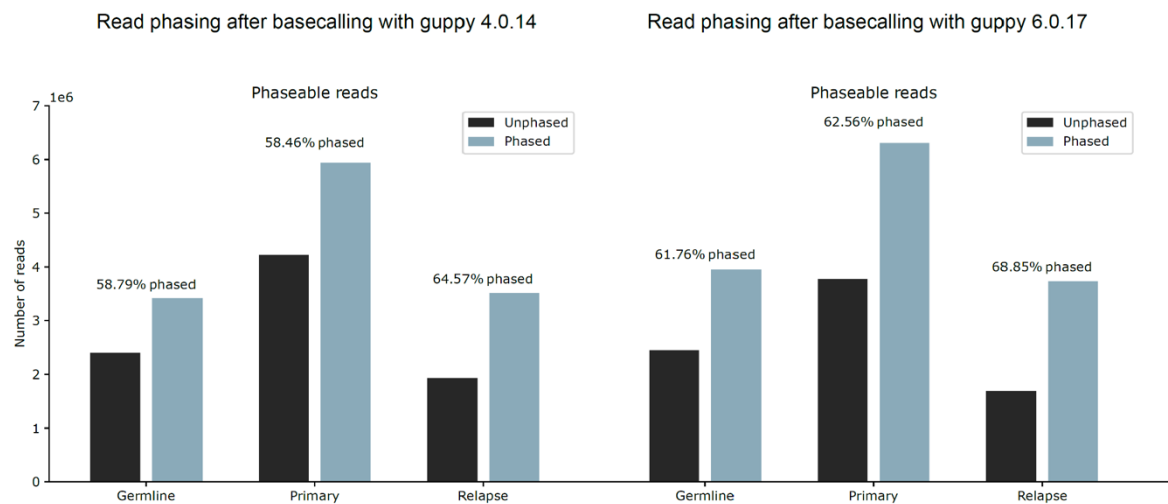


Figure R1.1: Read phasing before and after re-basecalling with guppy 6.0.17 shows a 2.3% increase in number of mapped reads after quality filtering and a 3.8% improvement to phase-ability of mapped reads

Point by point response:

Reviewer #1–general comments: In Rausch et al. “Long-read sequencing of diagnosis and post-therapy medulloblastoma reveals complex rearrangement patterns and epigenetic signatures”, the authors utilized Oxford Nanopore (ONT) long reads, as well as Illumina short reads to study structural and epigenetic dysregulation in the cancer genome. Particularly, the authors describe a pattern of complex genomic structural variation which they term “templated insertion threads”, that are detectable with both ONT and short-read data. To enable the resolution of chromothripsis and the identification of these templated insertion thread events, they present and rely primarily on two methods, Lorax and Rayas, respectively. Additionally the authors also apply the ability of ONT data to detect methylation status and identify differential

methylation - including in a haplotype-aware manner. They find that the differential methylation is enriched for medulloblastoma (MB) driver genes in their MB sample. They also use Illumina shortread RNA-seq alongside the haplotype resolved methylation data to study allele-specific patterns of methylation and expression.

The work reflects the cutting edge use of sequencing technologies and sits at the forefront of research in complex structural variation. Particularly the haplotype-resolved analysis of structural variation and methylation signatures, individually and in concert, will enable important biological discoveries. There is an impressive amount of new material and methods presented, and the paper represents an important proof-of-concept for the field. However, there are simultaneously concerns that this reviewer has about several aspects of the work which warrant additional review and revision before publication.

Response: We thank the reviewer for their positive and supportive review, and generally for their interest in our work. A point-by-point response to the comments raised is below.

Major issues:

Reviewer #1–1: In the section “Results: Haplotype-phased assembly of complex somatic rearrangements.” The section concludes by speculating about why the HSR integration points could not be found. This reviewer agrees that ONT read length limitations may be one reason, however is it possible that heterogeneity of the HSR integration from cell to cell is a factor? If integration heterogeneity were abundant, it would perhaps also be very difficult to find consistent read-support anchoring HSRs into chromosomes.

Moreover, in Figure 1e FISH image iii, the image caption states “colocalization of the centromere 17 probe with the RP11-651L9 probe”, however the colocalization statement may apply to one of the red blots near the centromere, but there is a large diffuse pattern elsewhere which may suggest either extrachromosomal form or an extremely diverse integration form elsewhere. There is some concern by this reviewer with regards to exactly how much diversity there is in co-localization of the two guides from these two images alone. In what fraction of image was at least one 651L9 probe not co-localized with cen17?

Response: As suggested by the Reviewer, we repeated the FISH experiment and specifically quantified the colocalization events between the RP11-651L9 FISH probe and the FISH probe for centromere 17. In total, 9% of the analyzed interphase nuclei show no colocalization and 91% show at least one colocalization. Additional pictures (Figure S5) as well as a supplementary table (Table S6) showing the quantification of the FISH signals are provided in the revised version of the manuscript.

In addition, we have now performed the FISH on metaphase spreads from the matched patient-derived xenograft model. This analysis shows the colocalization of the RP11-651L9 FISH probe and the FISH probe for centromere 17 on marker chromosomes (Figure 1E). These structures are chromosome fragments that do not look like typical double-minute chromosomes as they are larger, but could potentially be ring chromosomes.

We therefore agree that it is very well conceivable that heterogeneity of the HSR integration from cell to cell may contribute to the difficulty to find the integration site, and this heterogeneity is supported by the FISH data, as the number of signals and colocalizations vary between nuclei. We hence now discuss heterogeneity in HSR integration as a (likely) factor making it difficult to pinpoint integration points through ONT read analyses. The additional FISH analyses have been incorporated in the main text and in Figures 1E and S5.

Reviewer #1–2: The manuscript states the templated insertion thread presented in Figure 2a (primary tumor) was not detected again in the relapse. Were the majority of SV events present in primary found again in the relapse (suggestive of the relapse arising from a subclone of the

primary)? If most all other SVs were conserved between primary and relapse, but this templated insertion thread was not, can the authors describe why it may have not reappeared? It is unclear whether we should have expected to see it in the relapse or not from how the data is presented in the manuscript.

Response: We performed a more systematic comparison of somatic variants found in the primary tumor and relapse because multiple reviewers raised this point. These analyses show that the primary tumor and the relapse arose from a distant common ancestor – with both samples sharing only 34% of all somatic SNVs (see Figure R1.2, new Supplementary Figure S21). In particular, the relapse sample lacks 18% of the somatic SNVs present in the primary tumor, and it acquired a number of additional somatic SNVs (48%) absent in the primary tumor. Notably, the complex rearrangements involving chromosome 7 of the primary tumor are not present in the relapse sample, including the templated insertion threads (Figure R1.3, new Supplementary Figure S20) – but the CS11-17 chromothriptic structure has been preserved between both samples.

Therefore, for somatic SVs, we observe overall only 11% shared SVs between primary and relapse. The relapse sample gained 19% and lost 70% of all somatic SVs, which is mainly driven by the disappearance of many of the complex rearrangements involving chromosome 4,5,7,9,16,19 and X.

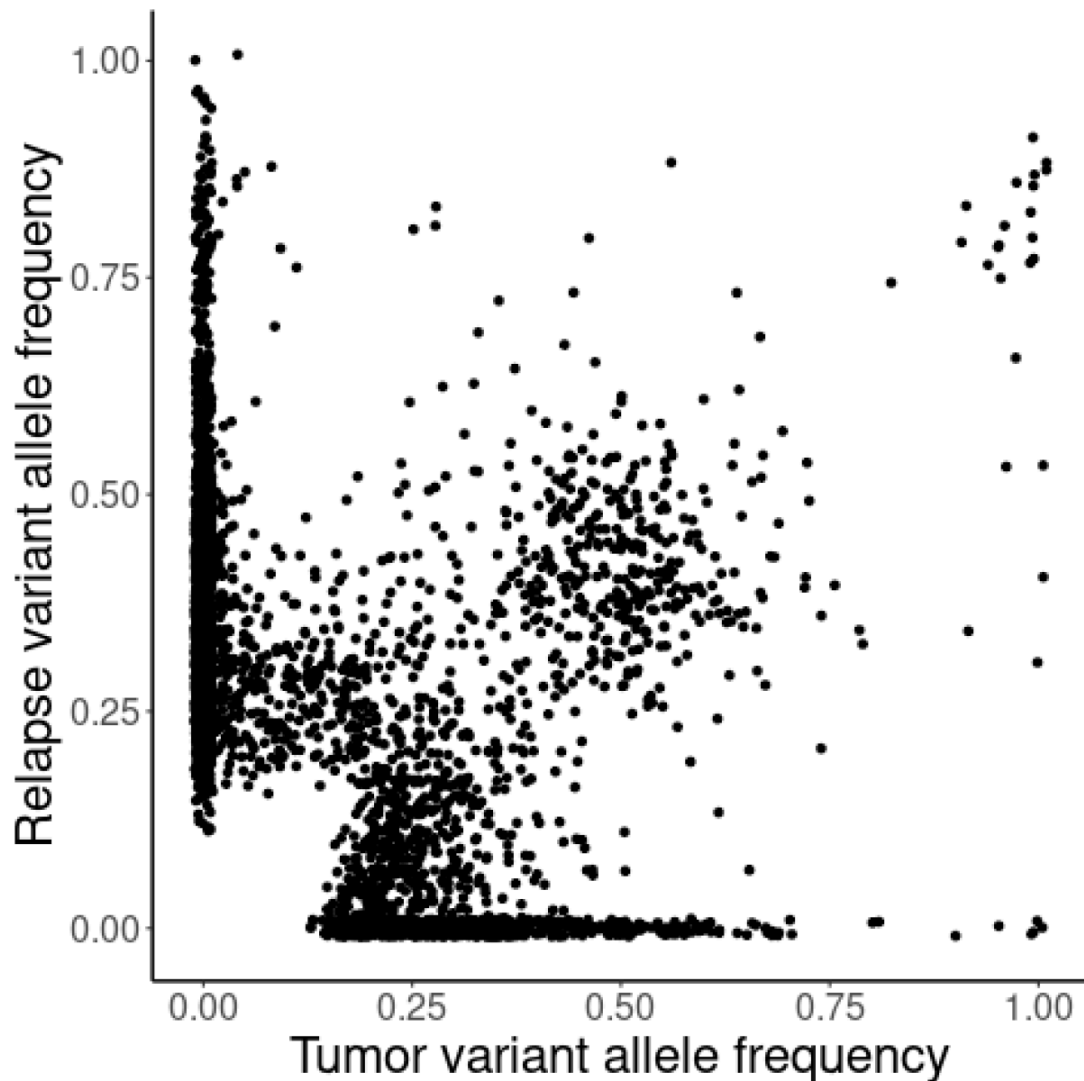


Figure R1.2. Comparison of the somatic variant allele frequency of all identified single-nucleotide mutations. We added randomly ± 0.01 in x- and y-direction to better highlight the variants that have been lost or acquired in relapse.




Figure R1.3. Read-depth ratio plots of the primary tumor and relapse computed by HaTCHet. Notably, the complex rearrangements on chr7 are not present in relapse, including the templated insertion thread involving segments of chromosome 7 and chromosome 4.

Reviewer #1–3: As the authors describe in the introduction section - these templated insertion threads appear to be distinct from TST-jumps (Umbreit et al.) and similar mechanisms proposed by others (Li et al.), mostly by virtue of the increased copy number. If the samples with the other, similar, modes of insertion-like jumps are given to Rayas, how does it classify them? Specifically, if Rayas were run on the sequencing data generated by Umbreit et al. in the renal cell carcinoma sample described in the discussion section, would it discriminate between TST and templated insertion threads? Are the patterns described in Umbreit et al. perhaps an ‘early’ form of the more dramatic phenomenon described in this paper? Similarly, if appropriate thresholds are selected for running Rayas, how do the regions called as templated insertion threads by Rayas match with what was reported by Li et al. using PCAWG data?

This reviewer would like to know, definitively, if the authors are proposing that the templated insertion threads described in this paper are arising by the same mechanisms proposed in the prior works (albeit perhaps having far more mechanistic iterations to increase copy number)? Otherwise, can the authors clearly demonstrate/argue/propose the phenomenon presented in their paper arises by a different mechanism than TST?

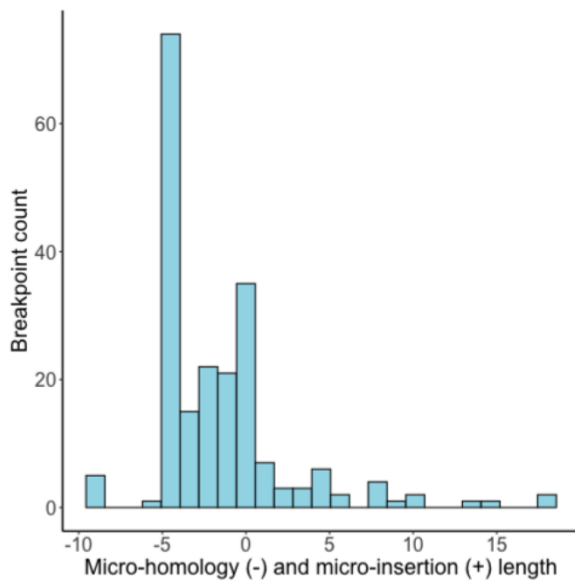
Response: We used bulk sequencing data of Primary Clone 1a displayed in Umbreit et al. Figure 5b (SRR10948021) to test whether rayas would detect TST jump patterns as templated insertion threads. Since rayas requires a control sample, and Umbreit et al. did not bulk-sequence the mother cell line, we obtained bulk-sequencing data of the mother cell line from Tourdot et al. (SRR1778442). Both samples were realigned to GRCh38 before running rayas. Rayas with default parameters failed to detect any templated insertion threads in these given data. Li et al. examined clusters of 2-10 SVs in PCAWG and then tried to identify a sequential path through the segments by following the breakpoint junctions. To our understanding, this procedure did not include self-concatenations because the authors considered individual segments in a

cluster to have similar copy-numbers. This is vastly different in templated insertion threads where some segments can reach a copy-number greater than 50 along other segments in the same thread that occur less than 5 times. Furthermore, Li et al. reports predominantly templated insertion events involving only two breakpoint junctions and the maximum observed is a cervical squamous cell cancer with seven templated insertions strung together, which is in stark contrast to templated insertion threads which can contain more than 100 breakpoint junctions. Because of this, we see solid overlap with Li et al. at the level of individual templated insertion source segments (78.7% called in PCAWG), but 0% overlap to templated insertion clusters in PCAWG with at least 3 events, presumably because highly complex self and cross-concatenations present in TI threads were not considered. This confirms that cycles, chains and bridges of templated insertions differ from templated insertion threads.

As alluded to in the Discussion section of our revised paper, our data in conjunction with prior work performed primarily using short reads (Umbreit et al. 2020; Li et al. 2020) suggests that a variety of templated insertion-associated classes of somatic SVs exist. These may or may not share common SV formation mechanisms yet to be discovered. Long read sequencing along with larger sample sizes will be ideal for determining the different types of templated insertions and characterizing their formation mechanisms in the future.

Reviewer #1–4: To what extent does sequence homology exist along the junctions between segments and in self-to-self connections for templated insertion threads? Is there microhomology and is it greater/less than expected? This may help to reveal some of the biological signatures by which these rearrangements form.

Response: Thanks for suggesting this breakpoint analysis. As shown in Figure R1.4, the breakpoint junctions tend to have 1-5bp of micro-homology, indicative of microhomologymediated end joining, or an exact breakpoint without micro-homology and micro-insertions. These junction patterns suggest the involvement of (alternative) nonhomologous end-joining, (alt-)NHEJ, or microhomology-mediated end joining (MMEJ). The peak at -4, a micro-homology of length 4, is due to a repeated junction that occurs 58 times in the TI thread exactly as shown at the bottom of Figure R1.4 (included as Supplementary Figure S16 in the revised manuscript).



chr4:168,398,369-168,398,412 (+) `GCTGTTTCCAGACATTGAGTGCCAGTGAGCATGGGAGGGAggccaagtgggggctaaga`
 TI-thread breakpoint sequence `GCTGTTTCCAGACATTGAGTGCCAGTGAGCATGGGAGGGAGGCCCCACCCTGGAGCCTCTCTGTGCCCAAGTTTTAA`
 chr7:7,805,260-7,805,299 (+) `tagaagaactattgatacctGGCCCCACCACCTGGAGCCTCTCTGTGCCCAAGTTTTAA`

Figure R1.4. Micro-homology and micro-insertion breakpoint junction analysis for the TI threads in the primary tumor. Microhomology length is plotted on the negative scale, microinsertions as positive lengths.

The bottom panel shows a characteristic micro-homology length of 4 (GGCC) for a junction that occurs 58 times in the TI thread.

Reviewer #1–5: The authors note that there was not a “significant relationship between ASM and proximity to somatic variants” (Results - “Allele specific methylation and expression”), and speculate that much of the ASM was due to germline variation. Does this imply it is necessary to also capture tissue of origin matched-normal to perform robust detection of ASM patterns in the context of somatic variants? Can the authors elaborate on what procedure might be done to help assist with the discovery of ASM that are linked to somatic variation?

Response: We thank the reviewer for raising this important question. Since methylation is highly tissue specific, we indeed believe that in order to confidently link somatic SNVs to methylation changes, it would require a matched-normal sample from the same tissue type. In our analysis we were hoping to detect subclonal somatic SNVs which split the methylation pattern of reads with higher confidence than germline SNVs (i.e. haplotype phasing). Furthermore, the methylation change must be close enough to the SNV to be able to map methylation and genotype on the same reads. While we were actually able to identify a small number of somatic SNVs with methylation changes that fulfilled this criteria (5 out of 1,793), it is difficult to say with certainty whether these are the result of poor haplotype phasing or other technical artifacts, without having a matched-normal sample. We therefore chose to be conservative in our reporting, stating that there is currently no significant relationship between ASM and proximity to somatic variants. We added a sentence on this limitation to the updated discussion.

Reviewer #1–6: This reviewer feels that additional validation of the Lorax method is warranted. Most genome-reconstruction tools are tested on multiple samples and/or simulated data, and then validated with an orthogonal technology (e.g. some form of cytogenetic validation, optical mapping, gel-based methods, etc.). However in this paper it appears it is only deployed on a single sample. Can the method be tested on other publicly available datasets? Does it support data from other long-read technologies, such as PacBio?

Response: Since multiple reviewers expressed their concern about using only a single sample we decided to screen another large cancer sample cohort, the NCT/DKTK Master cohort (Horak et al. 2017) for potential TI threads using rayas. We then generated additional (low coverage) ONT long-read data in two NCT/DKTK Master cohort samples to substantiate our finding of TI threads and at the same time validate the sensitivity of lorax.

In the first sample, P1 - obtained from a liposarcoma patient - lorax verified the predicted TI thread with self- and cross-concatenations as predicted by rayas on the matched short-read data (Figure R1.5 below, new Supplementary Figure S25). In the second sample P2, a skin metastasis sample of a primary liposarcoma, lorax likewise confirmed the characteristic copy-number increase and the self- and cross-concatenation of these templated insertions as a TI thread (Figure R1.6, new Figure S26), for this sample with multiple predicted genomic integration sites (new Figure S27). We emphasize that lorax validated the TI threads in both cases in spite of the relatively low ONT coverage.

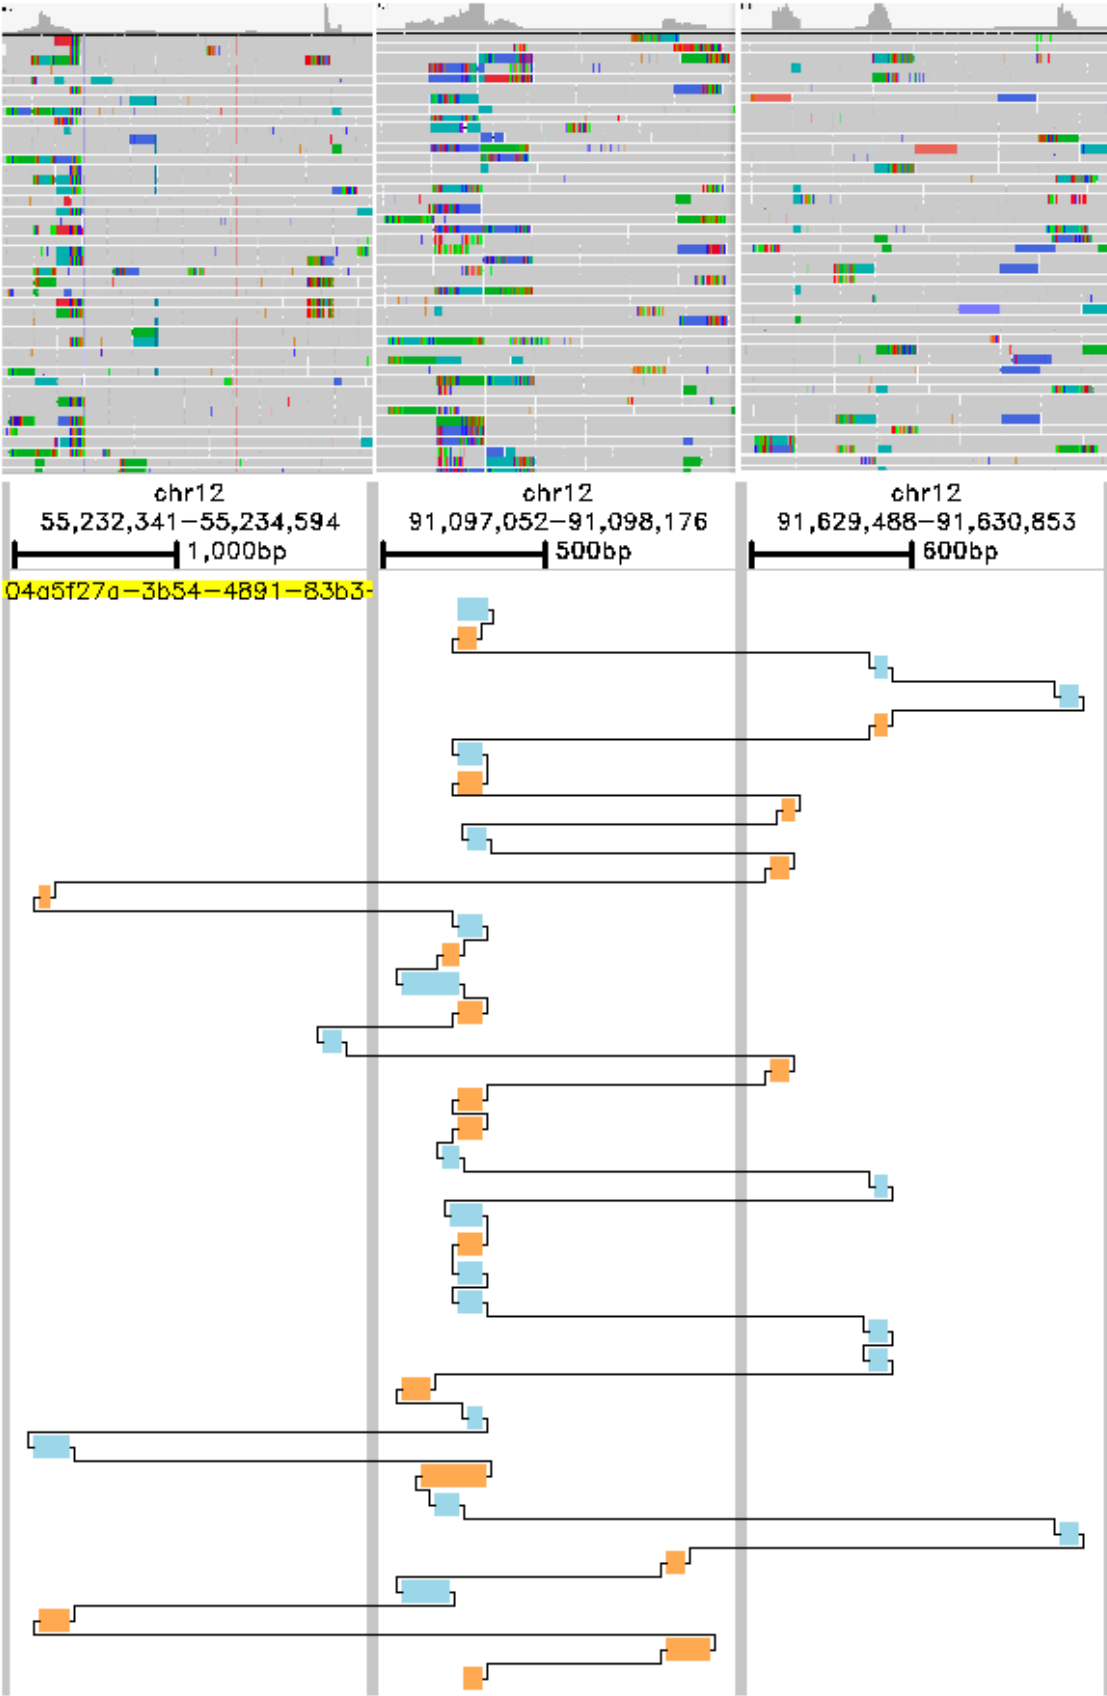


Figure R1.5. Raw ONT read supporting a TI thread (bottom panel) in liposarcoma sample P1. Alignment matches to GRCh38 are colored in blue (forward) or ochre (reverse). Aligned segments show strong coverage increases in the matched Illumina short-read data (top) as well as SV-supporting split-reads with soft-clips.

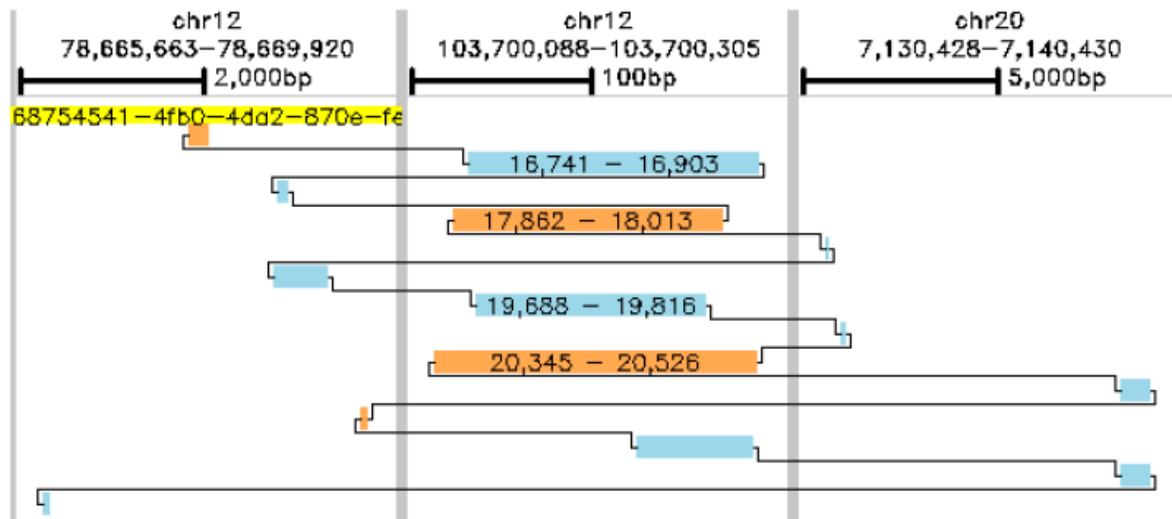



Figure R1.6. Raw ONT read supporting a TI thread (bottom panel) in P2. Alignment view is limited to source segments (separated by gray vertical lines) with at least 4 occurrences in the TI thread. Because of multiple integrations, aligned segments show strong coverage increases in the matched Illumina short-read data (top) as well as SV-supporting split-reads with soft-clips.

Additionally, we generated low coverage ONT data for a patient-matched xenograft model of the medulloblastoma primary tumor. Lorax identified parts of the TI thread present in the primary medulloblastoma with 6 out of 9 segments matching the primary tumor at 60% reciprocal overlap. Although none of the long reads captured the full 50kbp event, multiple reads supported parts of the TI thread structure observed in the primary tumor (see Figure R1.7, new Figure S11).




Figure R1.7. Pairwise dotplot of a PDX derived raw ONT read (x-axis) against the raw ONT read of the primary tumor (y-axis) presented in Figure 2 of the main manuscript. Both reads support the TI thread with matches to GRCh38 highlighted in the plot margins and colored by source chromosome. Lastly, we benchmarked lorax using simulated data (Figure R1.8, new Figure S40), which allowed us to assess its sensitivity and specificity in a controlled environment. Using the SV simulator Visor (Bolognini et al.), we implanted TI threads in chromosome 18 and then evaluated a range of calling parameters, technologies (ONT and PacBio) and coverage thresholds. In summary, lorax default parameters show 100% specificity at the expense of missing some (10%-30%) source segments. This miss rate is tolerable because it is at the level of individual source segments. As long as lorax identifies some source segments of a TI thread, the subsequent read extraction and local assembly is likely to rebuild the entire TI thread structure. All of the above analyses and additional samples have been included in the revised manuscript.




Figure R1.8. We simulated 10 TI threads on chr18 at 5x, 10x and 30x coverage, each with 5 randomly sampled source segments of a size smaller than 1kbp that were then concatenated and copied in random order to a final TI thread with 50 segments. We then applied lorax using a range of parameters, denoted as (A,B) in the legend. These parameters control the required increase in coverage (A) and the minimum number of required split-reads at the boundary of a segment (B). The default of lorax uses A=3 and B=3 (green plotting symbols). Simulated data with a typical error profile for ONT is on the left and simulated PacBio data on the right.

Reviewer #1–7: While Lorax simplifies the assembly of amplified chromothripsis regions by performing a targeted reconstruction, can the same conclusion/structure found by Lorax regarding the 1.55 Mbp event be found with other long-read based assembly methods (albeit perhaps less efficiently)? Such a demonstration may help convince the readers and this reviewer that the chromothripsis event reconstructed by Lorax in this case is both correct and complete.

Response: In response to this comment, we did run the Flye and Shasta assembler on the primary tumor sample. Both assemblers emitted contigs that fully spanned the TI threads of the primary tumor and an alignment of the contig against the raw read presented in Figure 2 confirms the correct assembly (Figure R1.9, new Figure S10). The CS11-17 chromothripsis event is much larger than the TI thread and both assemblers did not fully reconstruct this rearrangement, possibly because these assemblers compute a so-called squashed assembly of both haplotypes. Nevertheless, multiple contigs appear to confirm individual junctions of the CS11-17 structure (Figure R1.10, new Figure S41).

We have included Figure S41 as a new Supplementary Figure to illustrate the current need in cancer genomics for targeted assembly approaches that leverage the long-read phasing information in conjunction with allelic imbalances.




Figure R1.9. Alignment of the Shasta assembly contig against an ONT raw read supporting the TI thread. Colored bars at the bottom indicate unique segment alignments to GRCh38 with different colors representing different chromosomes from which the templated insertions were derived.

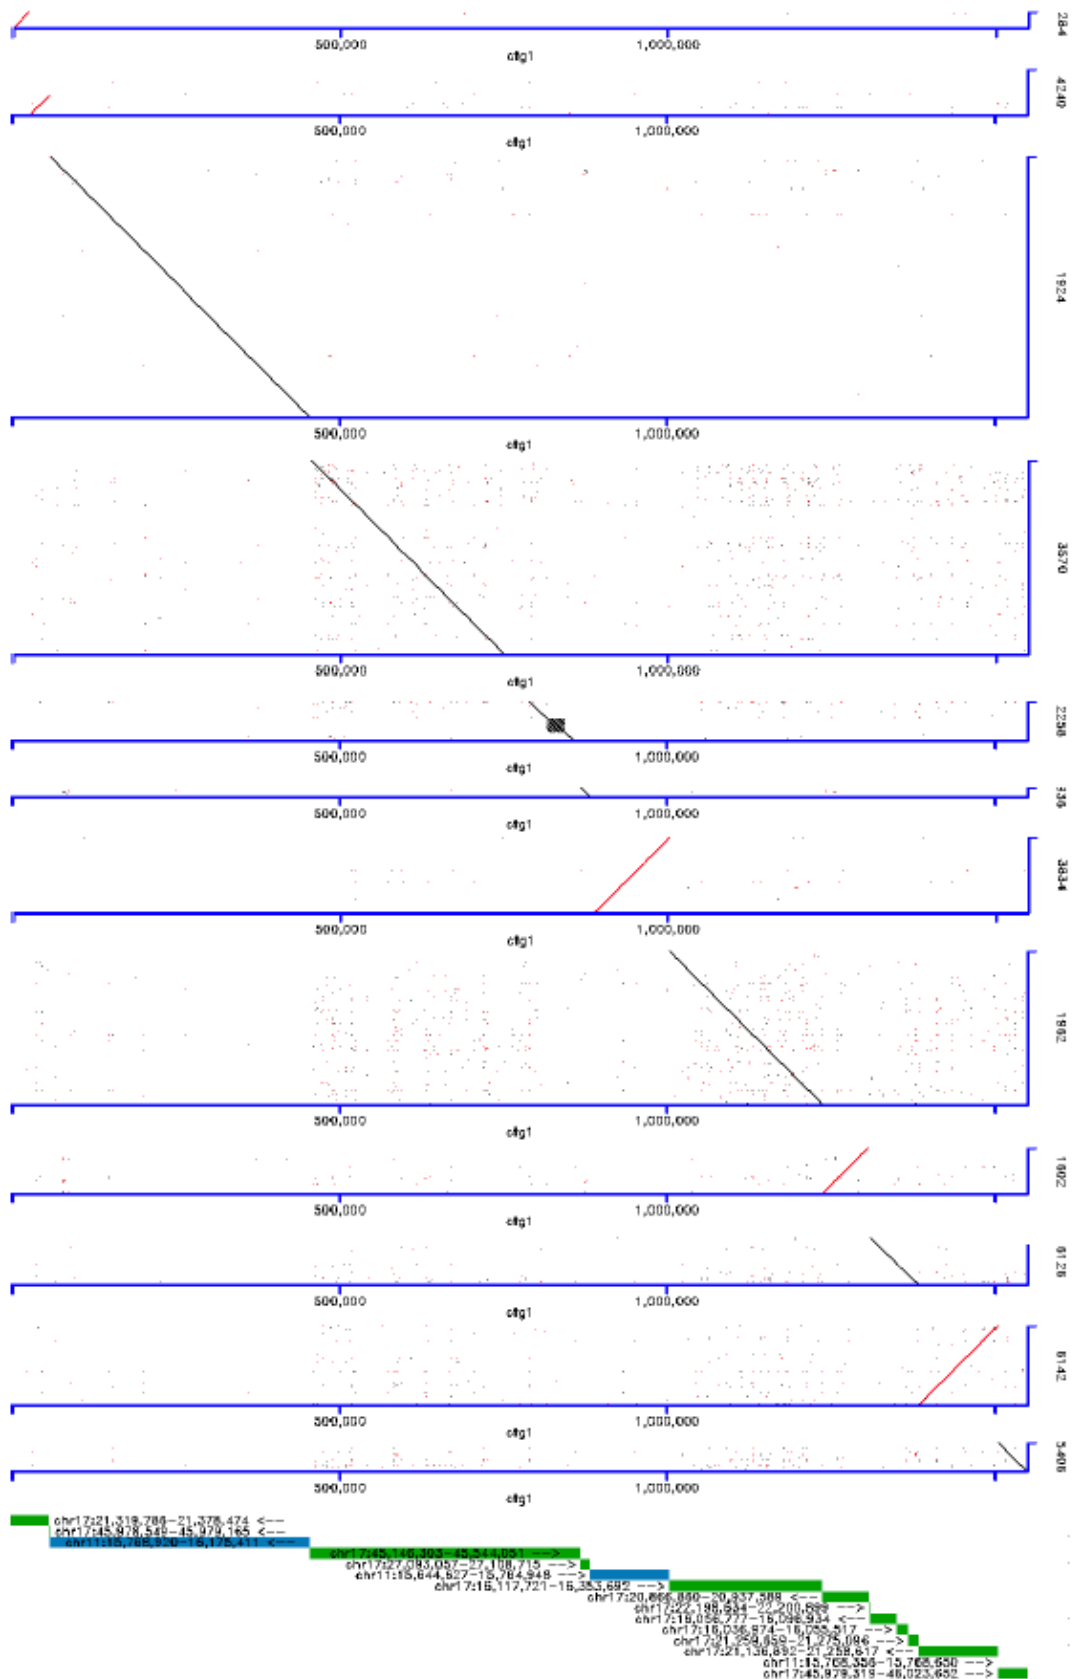


Figure R1.10. Contigs from a de novo Shasta assembly (y-axis) aligned to contig 1 (x-axis) of the targeted CS11-17 assembly. Multiple contigs fully overlap the CS11-17 assembly but the frequent breaks illustrate the need for targeted haplotype-resolved assembly approaches in cancer genomics. Colored bars at the bottom indicate unique segment alignments to GRCh38 with different colors representing different chromosomes from which the genomic segments were derived.

Reviewer #1–8: Given that at least one of the telomere SV events was resolvable only using the T2T assembly, can the authors clarify in discussion to what extent they expect the more completely resolved reference genome to improve detection of templated insertion threads? It is hinted at in the discussion section - but can the authors clarify further? Does the finding presented in discussion about three of five SV telomere junctions being unresolvable with short reads even using T2T imply there exists a limited ability to improve SV calls on short read data using the T2T build, as those newly added/improved regions are too low complexity to map uniquely with short reads anyway? This reviewer feels that point deserves a more comprehensive discussion.

Response: The choice of reference (GRCh38, T2T-CHM13 or pan-genome graphs) indeed affects the ability to map and interpret somatic variation. As reported by Aganezov et al. (2022), the T2T assembly improves genetic variant calling, revealing additional SVs in highly repetitive and previously unresolved genomic sequences – although even the T2T assembly retains regions too low in sequence complexity in which variant calling is still unsolved. Overall, we believe that a more comprehensive assessment of somatic variant calling accuracy using T2T versus GRCh38 would be a valuable contribution, but in the light of the design of this study with its relatively small sample size, and given the wide range of SV somatic mutation signatures across tumors, such a comparison is beyond the scope of our study. In the revised manuscript, we therefore restricted the respective section of our manuscript to the selective use of the T2T assembly (which we applied in cases where GRCh38 gave us alignment ambiguities).

Reviewer #1–9: Reporting some runtime and resource requirement benchmarks on the analyzed cancer samples would be helpful for those considering using Lorax and Rayas.

Response: For the short-read data of the primary tumor, rayas required 35 minutes at a peak memory usage of 3.26G RAM. For PCAWG, rayas' median runtime was 85.2 minutes for analyzing a paired tumor-normal cancer genome. We unfortunately did not track the peak memory usage in the PCAWG cohort but none of the cluster jobs was canceled at a maximum memory allocation per node of 16G.

For the TI thread discovery, lorax required 37 minutes and 3.3G RAM for the identification of connected components, 17 minutes to extract the reads (at <100MB RAM) and less than 5 minutes to assemble the small set of reads with Shasta. These runtimes and resource requirements are now mentioned in the Methods section.

Reviewer #1–Minor issues:

a) The last sentence of the Summary seems vague or incomplete. “Our study shows the potential of long-read sequencing in cancer.” Can it be written more specifically?

Response: Thanks for the suggestion to make it more specific. In our view, the main advantage of long reads is the complete genetic and epigenetic characterization of complex somatic rearrangements that can only be partially or indirectly ascertained using short reads, as exemplified in our study by TI threads. We reformulated the last sentence of the abstract accordingly, but had to keep it short due to the word limit.

b) A clearer description of the haplotype phasing approach should be reported in the main

text, to make it more clear which existing haplotyping tools are used to generate the phased blocks (WhatsHap) and how they are refined (ShapeIt).

Response: Thanks for pointing this out. As suggested, we included WhatsHap and ShapeIt in the main text.

c) Please consider ordering the supplemental figures so they match their appearance in the main text for readability. E.g Figure S1 is not cited until well after many of the other supplemental figures. Same for Figure S3, and so on.

Response: We thank the reviewer for pointing out the inconsistency in the ordering of the supplementary figures. We updated the ordering to reflect the order in the text.

d) Figures 1c, please more clearly indicate the chromosome number(s) associated with the coordinates on the x-axis. It took this reviewer a long time to determine that the chromosome cartoons from panel D matched the coordinates shown in panel C.

Response: Multiple reviewers suggested improvements for Figure 1C and 1D which led us to redesign these panels, including a clear indication of the involved chromosomes.

e) In Figure 2A,B it is not clear which chromosomes the templated insertion coordinates match to in this plot. Are the genome regions shown in the top part of B matched directly to the genome regions shown in the bottom part of B? Placing the legend between them and not labeling chromosome names nearby coordinates on the x-axis makes it unclear. Matching the colors of the flanking regions from panel 2A to the colors of the highlighted genome regions shown in 2C may help clarify the relationship of panels A,B,C more naturally.

Response: Figure 2B upper panel (short-reads) and lower panel (long reads) are indeed matched and we followed the reviewer's suggestion to clearly label the x-axis and to better link these panels. In response to this comment, we also added new plotting functions to wally to produce "enhanced" dot plots (Figure S27, for instance), that incorporate the minimap2 alignments to GRCh38 in the plot margins, and chained alignments of TI threads (Figure 2F, for instance). We hope that both of wally's visualization methods will improve understanding of TI threads and other future complex rearrangements in cancer.

f) Please correct grammatical errors throughout manuscript and figure captions (e.g, Fig 2c caption), and address tense issues, e.g. Results section, Haplotype-phased assembly of complex somatic rearrangements: "In the primary tumor, we find 697 somatic SVs,..." and also elsewhere "we identify", "we separate" should be written in past tense so it is consistent with the rest of the tense in the section.

Response: We thank the reviewer for pointing out these textual issues. We have addressed them in the revised version of the manuscript.

g) In Rayas, what is the significance of filtering "connected subgraphs where all nodes of G_s are nearby in the genome", and how was the default threshold of 10kbp decided upon? This reviewer is not asking for a rigorous justification of 10kbp, but rather a simple motivation and explanation for what kinds of events this filters out since it is not clear why this filter exists and how it affects discovery of templated insertion threads.

Response: Previous somatic SV studies in PCAWG (Li et al. 2020) and the Hartwig Medical Foundation cohort (Cameron et al. 2021) have shown the importance of short tandem duplication signatures in various tumor types. Short tandem duplications can have a similar size as the source segments in TI threads; likewise such tandem duplications show an increase in coverage as well as paired-end and split-read support. We therefore ensured that TI threads involve at least 2 source segments from distinct locations in the genome, to not confuse such duplication signatures with TI threads.

h) Ensure all acronyms are defined. Haplotype “HP1” and “HP2” acronyms not explicitly defined. “DMR” not defined until Methods, after it appears many times.

Response: We thank the reviewer for pointing out these textual issues and have addressed them in the revised version of the manuscript.

i) Figure S14, S19, consider including scale bars for the heatmaps.

Response: We have added scale bars to the presented heatmaps.

j) The use of the term “high-complexity regions” in the discussion section means something different than “high complexity” and “low complexity” with respect to sequence composition as described in the rest of the paper. Consider saying “regions of high structural rearrangement complexity” or similar to in the Discussion section to differentiate the two uses of the term.

Response: We have clarified the difference between high/low complexity regions as mentioned in the DNA-methylation analysis versus sample genomic complexity as mentioned in the discussion.

k) The boost library appears to be required to compile Rayas and Lorax but this or any other compile requirements are not listed on the READMEs. A testing dataset would also be useful for assessing if a Rayas installation is working.

Response: We thank the reviewer for this comment geared towards improvement of the usability of our tools. We have added these dependencies to the GitHub Readme. Since both tools (rayas and lorax) are available as Singularity or Docker containers and as a statically linked binary, we expect that most users will not need to compile the code from source.

Reviewer #2:

Reviewer #2—general comments: Rausch and colleagues present a genomic case report of a ONT sequenced medulloblastoma primary tumor and relapse along with its matched normal. They use the longer reads to assemble complex rearrangements including an interesting selfconcatenation pattern that they dub a “templated insertion thread”. In addition to SV calling they use the ONT reads to highlight several allele-specific methylation patterns.

Overall, the study provides a somewhat limited data resource and anecdotal findings of unclear significance without a clear methodological advance. The templated insertion pattern, while interesting is not pursued to sufficient technical or biological depth. Since the long read analysis is of a single case, the highlighted findings are inevitably cherry picked and their biological or clinical significance is not apparent. Notably, the primary and relapse comparison particularly for methylation yields little insight into the genetics or epigenetics of medulloblastoma relapse. Finally, new algorithms are presented without benchmarking or comparison to previous tools.

Response: We thank the reviewer for their review, and for their interest in the templated insertion thread (TI thread) complex SV pattern. We have addressed the specific comments

raised by the reviewer point-by-point below.

Major critiques:

Reviewer #2–1: The templated insertion thread pattern seems to be the most interesting finding in the study but is not pursued to sufficient technical or biological depth. The long read evidence for this (like for many of the other findings) comes from a single case, and actually a single sample since it is absent from the relapse. While intriguing, it is unclear how generalizable this pattern is. A possibility that needs to be considered is whether this is just a technical artifact.

Response: We thank the reviewer and do agree that it is important to further substantiate the observation of templated insertion threads (TI threads). In response to this reviewer comment, we generated additional data to further substantiate our finding of TI threads, both at the technical and biological level. We therefore generated ONT sequencing data on two liposarcoma samples from the NCT/DKTK Master project (Horak et al. 2017) as well as data on a patient-derived xenograft (PDX) model from the primary medulloblastoma. In addition to this, we teamed up with Anton Henssen's lab, experts in extrachromosomal circular DNA (ecDNA) (van Leen et al. 2022), to investigate TI threads for evidence of circular structures/intermediates using Circle-Seq (described in more detail in the methods section of the revised manuscript), coming up with no evidence of circular enrichment in TI threads. Lastly, we confirmed TI threads on the technical level using two assembly algorithms (Shasta and Flye) that generated contigs fully spanning the TI threads of the primary tumor.

More specifically, we first used rayas to predict the occurrence of TI threads in 17 previously generated short-read (Illumina WGS) liposarcoma patient datasets from the NCT/DKTK Master project (Horak et al. 2017) – a tumor type we chose given our prediction based on PCAWG that 74% of all liposarcomas exhibit TI threads. Rayas identified evidence for TI threads in 6 out of 7 (86%) dedifferentiated liposarcoma patients from the NCT/DKTK master project cohort, consistent with the results of the same method generated in the PCAWG data. In myxoid liposarcomas (N=10), rayas could not identify TI threads because myxoid liposarcomas are driven by a chimeric fusion gene (FUS-DDIT3) rather than genomic rearrangements affecting chromosome 12q. We next generated ONT data on two liposarcoma samples – a primary dedifferentiated liposarcoma (P1) and a skin metastasis of a different liposarcoma patient (P2), in order to achieve validation of the patterns identified using rayas (and pinpoint potential technical artifacts), as well as to enable further long read based analyses of this complex DNA rearrangement pattern. We sequenced P2 on a GridION instrument, achieving 4x coverage at a median read length of 6.8kbp (N50 read length: 11.9kbp) (Figure R2.1) and an estimated error rate of 4.4%. P1 was sequenced on a new ONT P2 Solo instrument, where the run resulted in 3x coverage (as a consequence of premature software termination) with a median read length of 4.9kbp (N50 read length: 8kbp). Lorax confirmed the TI thread pattern in both cases in the long read data, which verifies our ability to accurately call TI threads in short read datasets followed by subsequent verification using ONT sequencing.

We next performed further analyses of these newly generated complex SV data. The TI thread pattern of P1 closely resembled the medulloblastoma pattern with apparent self- and crossconcatenations observed even within a single read (Figure R2.2), although the coverage was too low to fully resolve this presumably long TI thread from start to end. Next, we analyzed the TI thread pattern in P2 (Figure R2.3). In contrast to the TI thread detected in the medulloblastoma sample, we found multiple occurrences of the TI thread structure in this liposarcoma genome, resulting in a high total copy number, with different integration sites contributing to the total copy number (see Figure R2.4). Again, short-read and long-read data are congruent (ruling out a technical artifact), with aligned long-read segments showing strong coverage increases in the short-read data as well as many SV-supporting reads with soft-clips (Figure R2.2, R2.3 and R2.4, new Supplementary Figures S25, S26 and S27).

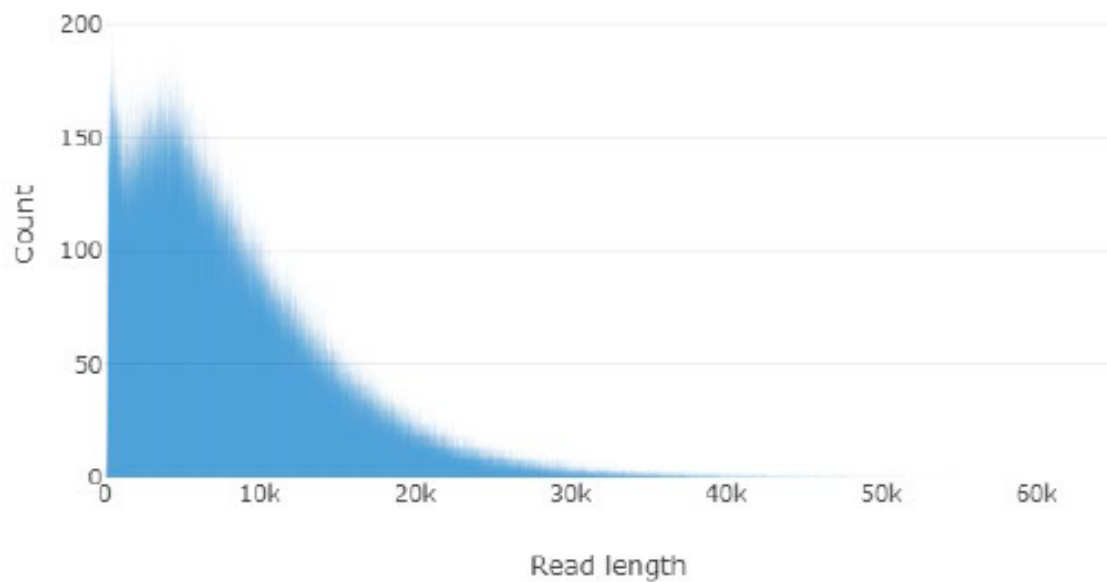


Figure R2.1. Read length distribution of the newly sequenced liposarcoma sample (P2).

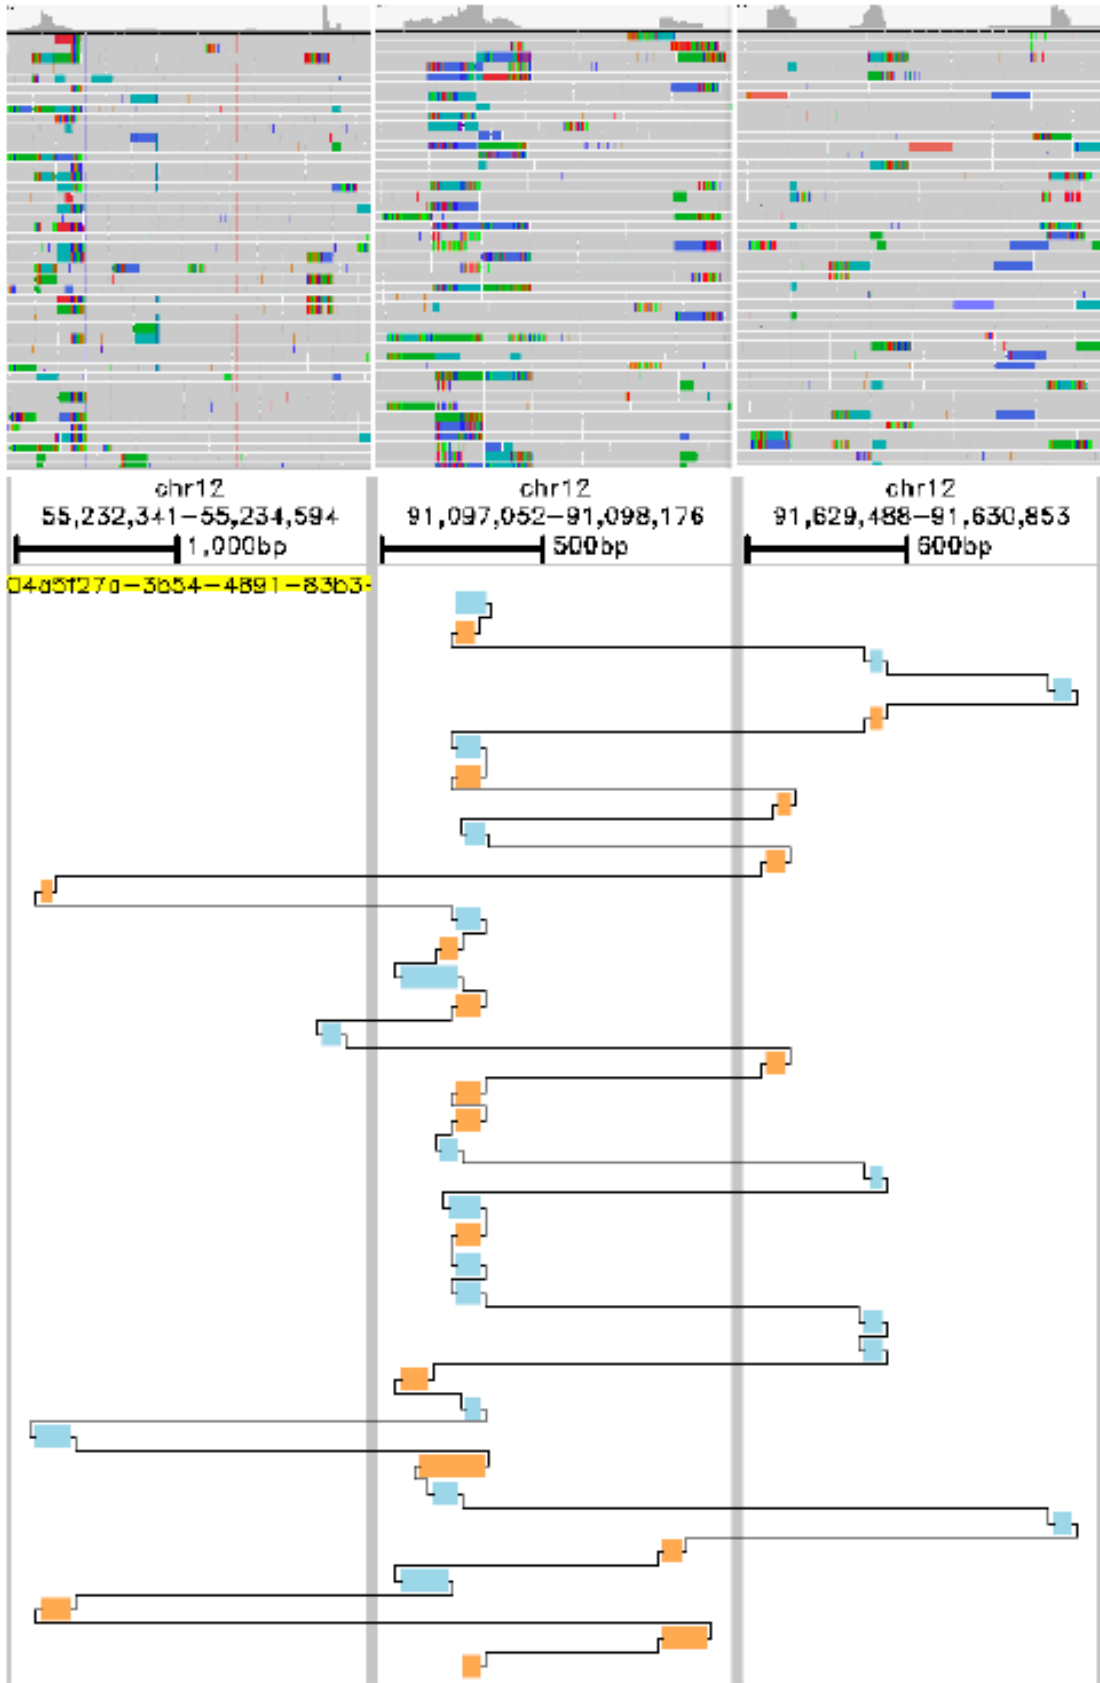


Figure R2.2. Raw ONT read supporting a TI thread (bottom panel) in liposarcoma sample P1. Alignment matches to GRCh38 are colored in blue (forward) or ochre (reverse). Aligned segments show strong coverage increases in the matched Illumina short-read data (top panel) as well as SV-supporting split-reads with soft-clips.

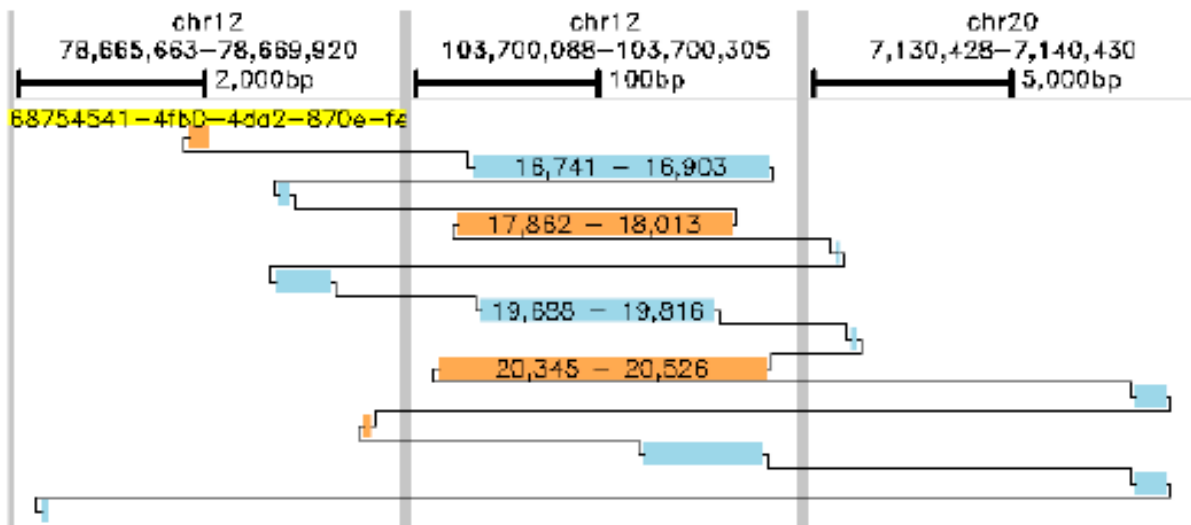




Figure R2.3. Raw ONT read supporting a TI thread (bottom panel) in P2. Alignment view is limited to source segments (separated by gray vertical lines) with at least 4 occurrences in the TI thread. Because of multiple integrations, aligned segments show strong coverage increases in the matched Illumina short-read data (top panel) as well as SV-supporting split-reads with soft-clips.

A)



R2.4. Lorax predicts multiple integration sites for a given TI thread. Pairwise dot plots of 3 raw ONT reads supporting the TI thread shown in Figure R2.3 with matches to GRCh38 highlighted in the plot margins and colored by source chromosome. Arrows indicate forward or reverse matches. The dot plot shows forward matches in black, reverse matches in red. TI thread structure is largely preserved but reads highlight different adjacent genomic sequences, indicative of multiple integrations.

We also generated low coverage sequencing data for a PDX model derived from the primary medulloblastoma, and in doing so obtained additional support for the detected TI thread (Figure R2.5, new Supplementary Figure S11). Although none of the reads captured the full 50kbp event, multiple reads supported parts of the TI structure and thus, these data further corroborate our findings made in the index patient.

Finally, at the technical level, we corroborated our TI thread findings by applying two assembly algorithms, Flye (Kolmogorov et al. 2019) and Shasta (Shafin et al. 2020), which both yielded contigs in agreement with the TI thread structure presented in Figure 2 (see assembly comment Reviewer #2-3 below, new Supplementary Figure S10). Collectively, these new data confirm our new observation of TI threads and rule out a technical artifact.

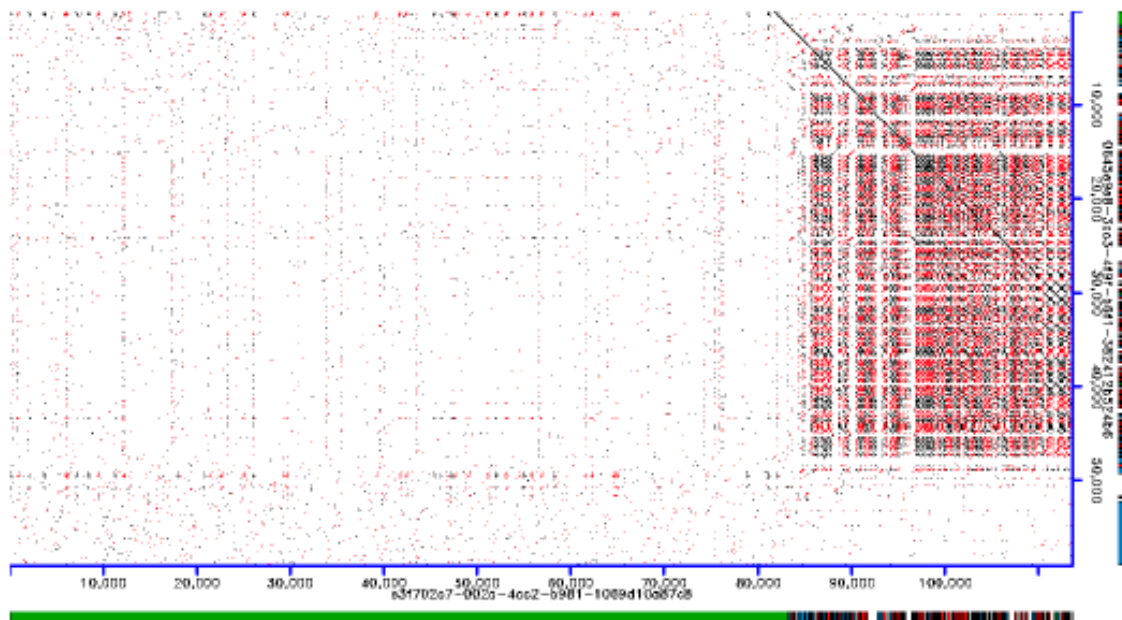


Figure R2.5. Alignment of a raw ONT read from the patient-derived xenograft (PDX) sample (x-axis) against the raw ONT read from the primary tumor (y-axis) that supports the TI-thread. The colored bars in the plot margin indicate individual aligned segments of the reads to GRCh38 with different colors representing alignments on different chromosomes.

Reviewer #2-2: Limited correlation is done of long and short read data around the proposed templated insertion thread. However it seems like the rearrangements in the templated insertion thread should be detectable by short reads. Areas of focal increased sequencing depth like shown in Fig 2b are frequent in regions of the genome for example due to low mappability. How do the locations of the read depth increase correspond to the locations of the reconstructed thread? Do they match in quantity? Are they absent in the relapse and the normal? These basic analyses should be presented.

Response: The correlation between short and long-read data is a legitimate point to address and

explore further. We made use of SvABA (Wala et al. 2018) to assemble TI thread junctions from the short-read data. Subsequently, we remapped the SvABA contigs to the TI thread structure, which verified several predicted self- and cross-concatenations (Figure R2.6, new Supplementary Figure S12). We stress, however, that none of the SvABA contigs could chain together more than 3 SV breakpoints (Figure R2.6), which highlights the importance of long reads to resolve the complexity of the TI thread structure.




Figure R2.6: Short-read SvABA contigs recapitulate individual SV breakpoints of the TI thread. Each of the 3 SvABA contigs (top panel, middle panel, bottom panel) support at most 3 SV breakpoints but the characteristic self- and cross-concatenation of TI threads is evident.

Besides short-read supported SV junctions, the overall Illumina read-based predicted copynumber is well-correlated with the observed occurrence count of segments in the TI thread, despite the short size of source segments (Figure R2.7, new Supplementary Figure S15). Visually the read-depth coverage flanks support the deviation in breakpoints and read-depth peaks within a larger segment correspond to nested alignments discovered using long-read sequencing (see, for instance, Figure R2.3).

We also performed additional analyses of the relapse sample, in which lorax did not identify TI threads. The absence of the original TI thread seen in the primary tumor is confirmed by readdepth analysis in the relapse (Figure R2.8) and by remapping the sequencing reads to the Shasta assembly (which did not yield a single confident mapping with $\text{MAPQ} > 1$ to the TI thread contig). As described in detail in our response to comment Reviewer #1–2, additional analyses we pursued at revision stage suggest that primary tumor and relapse originated from a distant ancestral clone that shared only a relatively small subset of somatic mutations (34%) – with the TI thread rearrangements not being amongst these shared somatic DNA alterations.




Figure R2.7. Each segment of the TI thread with a unique alignment in GRCh38 is annotated with its occurrence count in the TI thread and the read-depth based estimated copy-number in the matched short-read illumina data. Overall correlation is 0.99 (Pearson) and 0.85 (Spearman).




Figure R2.8. For the relapse sample, the read-depth based estimated copy-number supports the fact that the TI thread is not present in the relapse sample. The dashed line corresponds to copy-number 2 and all segments of the TI thread cluster close to this line.

Reviewer #2–3: The data shown in Fig. 2 shows a single long read that is 50 Kb in size. Similarly the supplementary figures show single ONT reads each with distinct patterns. Do these distinct reads point to a single rearranged templated thread consensus sequence, e.g. through assembly or other integration approaches?

Response: We applied the Flye (Kolmogorov et al. 2019) and Shasta (Shafin et al. 2020) assemblers to the long-read ONT data to respond to this comment (as alluded to above). One assembled Flye contig recaps the TI thread nearly perfectly, with the contig being only slightly truncated with respect to the original raw ONT read we present in Figure 2 (see Figure R2.9). This slight deviation suggests that the Flye assembly either collapsed the repetitive structure, or alternatively that there is tumor heterogeneity with respect to the underlying TI thread sequence (see comment Reviewer #3-2). For the Shasta assembly, we observed a single contig that fully aligns to the raw read (Figure R2.10, new Supplementary Figure S10) and 100% of its sequence is covered by aligning reads back to the assembly (Figure R2.11, new Supplementary Figure S19).

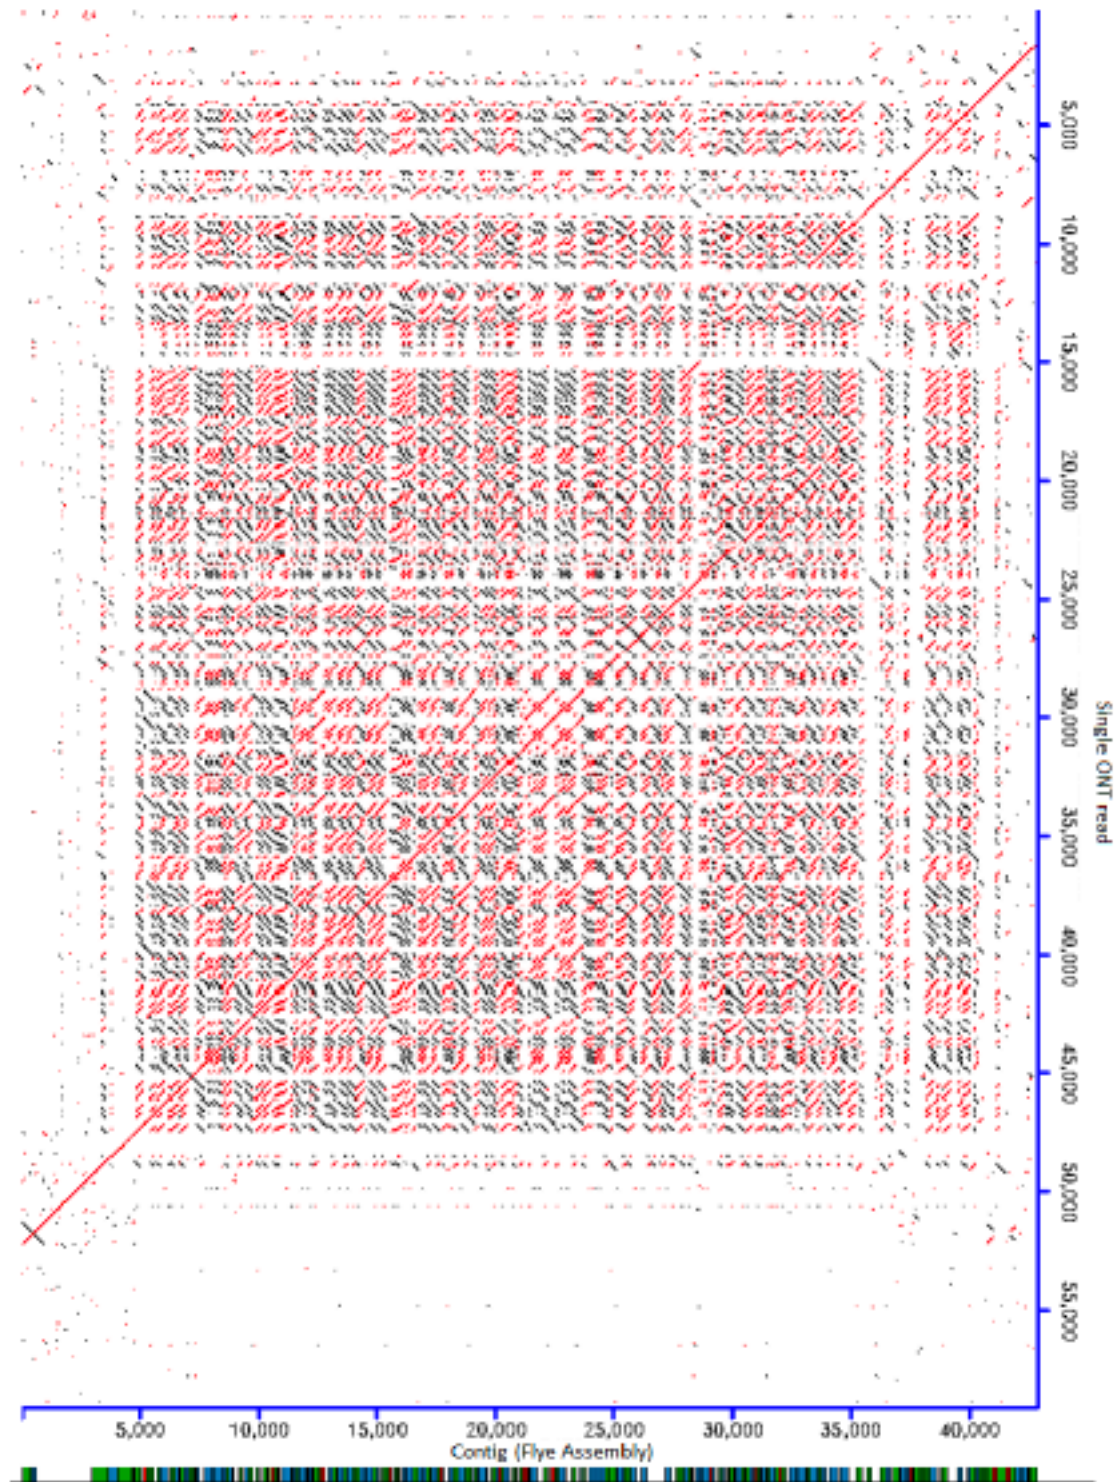


Figure R2.9. Alignment of the Flye assembly contig against a single ONT read supporting the TI thread. The assembly contains a deletion relative to the ONT read. Colored bars at the bottom indicate unique segment alignments to GRCh38 with different colors representing different chromosomes from which the

templated insertions were derived.




Figure R2.10. Alignment of the Shasta assembly contig against the single ONT read supporting the entire TI thread. Colored bars at the bottom indicate unique segment alignments to GRCh38 with different colors representing different chromosomes from which the templated insertions were derived.

Reviewer #2–4: Also the length of reads supporting the pattern appear to be 10 fold longer than the median read length in the library (5 kb). Is that the case for the other supporting reads and is that a problem?

Response: Compared to PacBio HiFi sequencing, ONT read length distributions typically have a long tail; in our case the ONT read data include a series of reads longer than 20kbp or even 50

kbp in length (e.g., see Figure R2.1). Although numerous reads below 10kbp support the TI thread, none of them captures the entire event; we therefore selected to always use the largest SV-spanning read for illustration purposes. In addition, aligning long reads back to the Shasta assembly shows that reads with a large variety in length confirm the (by and large uniformly covered) TI thread – in further support of our findings (Figure R2.11, new Supplementary Figure S19). Interestingly, a subset of reads actually supports the Flye assembly with a ~9kbp deletion relative to the original TI sequence found, suggesting the presence of heterogeneity within the tumor relative to this TI thread (Figure R2.11) – an analysis that has now been included in the revised manuscript in the TI thread section.

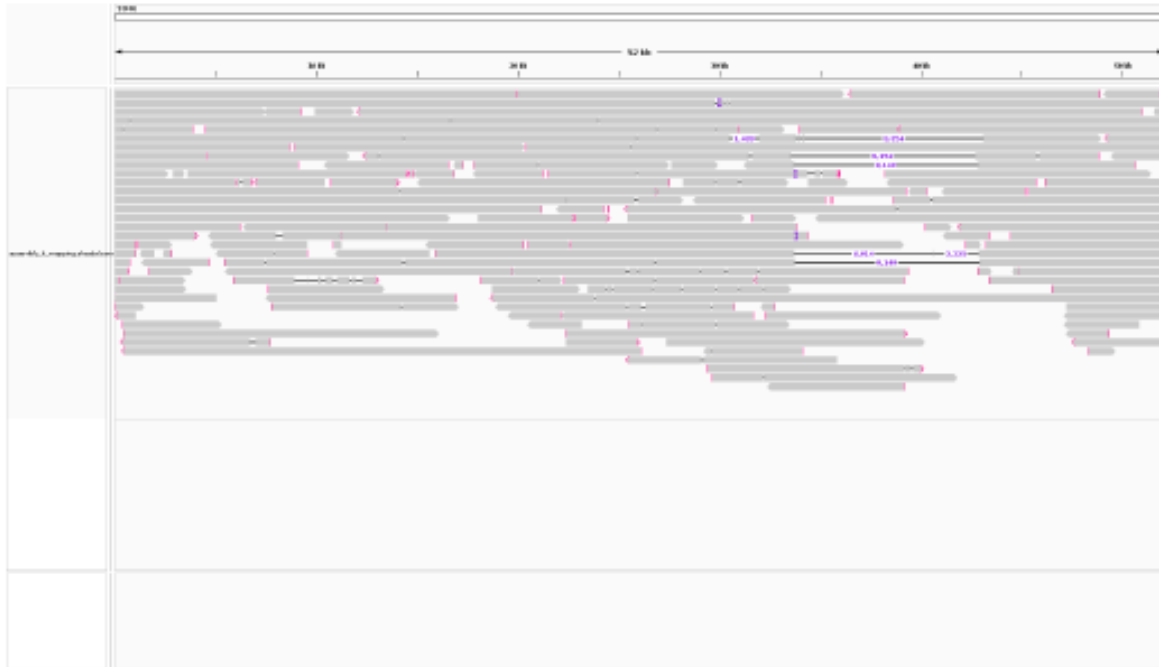


Figure R2.11. Alignment of the raw reads back to the assembled Shasta contig that supports the TI thread shown in Figure 2. The contig is evenly covered at the expected coverage and shows a subclonal ~9kbp deletion from 33kbp to 42kbp.

Reviewer #2–5: Are the templated thread adjacencies associated with the long reads found in short reads? It is unclear whether these patterns are happening at short read unmappable sequences. If they are mappable, they should be detectable in short reads - eg by aligning short reads to the long read-derived contigs.

Response: The source sequences of TI threads are mappable and thus, we do observe an increased copy-number in short read data as well as many SV-supporting split-reads – confirming the TI thread pattern (see, for instance, Figure R2.2 or R2.3, new Figure S25, S26). We specifically designed rayas to evaluate these patterns in short read data, which enabled us to study TI threads in short-read sequenced cancer genomics cohorts, including the entire PCAWG cohort as well as a subset of the NCT/DKTK master cohort (Horak et al., 2017).

Reviewer #2–6: The pattern seems to very strongly resemble genomic shards, noted with BAC sequencing in 2007 (Bignell et al), however this paper is not referenced or discussed.

Response: We are sorry for the oversight and now discuss and reference this relevant paper. As the reviewer noted, Bignell et al. describes small fragments of the genome, called genomic shards, strung end to end while investigating somatic rearrangements in bacterial artificial chromosomes (BACs) of HCC1954 and NCI-H2171. Source segments in this study were tightly clustered near chromosomal breaks (similar to the TST jumps described by Umbreit et al.), but no self-chaining with increased copy-number was reported – and therefore no pattern closely resembling TI threads has been reported by Bignell et al. 2007 (this might have, at least in part, been a resolution problem given the pre-WGS era of the Bignell et al. study).

Reviewer #2–7: The threading also seems to resemble the typhoons which were also noted in liposarcoma in 2020 (Hadi et al). These were also notable for a high level of tandem and fold back rearranged copies concentrated in an HSR, though maybe on a larger genomic scale (megabase instead of kilobase). How do these patterns relate, both biologically and in their distribution?

Response: We did not initially make a comparison with the typhons ('tyfonas') described in Hadi et al. because they are actually in a completely different length category/genomic scale compared with the TI threads described in our study - they even exceed the genomic mass of double minutes or breakage fusion bridge cycles. Most liposarcomas are characterized by giant marker chromosomes or supernumerary rings, which is likely the cause of tyfonas accumulation (Hadi et al. 2020) as well as seismic amplification accumulation (Rosswog et al. 2021) – the latter of which we already mentioned in our original manuscript. Based on the co-occurrence of typhoons with supernumerary ring chromosomes in liposarcoma, Hadi et al. (2020) proposed that tyfonas may represent the genomic footprint of these large marker chromosomes. In response to this reviewer comment, we now also included a citation to Hadi et al. 2020 to our revised manuscript.

Reviewer #2–8: The authors propose a new algorithm to nominate this pattern in short reads, but do not validate these predictions with long reads in any other samples. The algorithm *rayas* described in the supplement would seem to detect threads associated with many classes of complex SVs not just the highly self-similar patterns highlighted in Fig 2A. It would seem essential to get a sense of specificity, meaning how often does a pattern suggested to be a templated thread in short reads actually give you the long read pattern. Again insight into this is limited given the small long read dataset.

Response: We agree that validation of predicted TI threads based on short read datasets is important. In response to this reviewer comment, we thus set out to validate additional short-read based TI thread predictions made by *rayas* in two liposarcoma samples and the matched patient-derived xenograft (PDX) data of the primary tumor – by performing additional ONT long read sequencing (see our response to comment Reviewer #2-1, as well as Figure R2.2 and R2.3, for full details). In each case, we were able to successfully validate short read-based TI thread predictions in the newly generated ONT datasets. We also benchmarked *rayas* using simulated data (Figure R2.12, new Supplementary Figure S39), allowing us to assess its sensitivity and specificity in a controlled environment. This analysis has now been included in the revised manuscript in the Methods section.

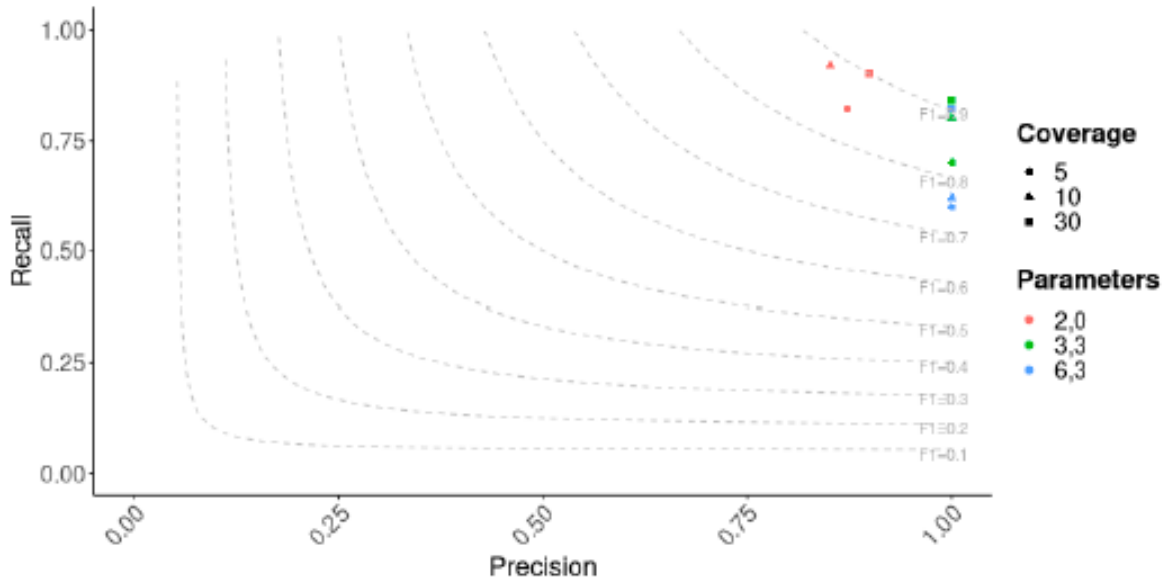


Figure R2.12. We simulated 10 TI threads on chromosome 18 at coverages 5x, 10x and 30x with 5 randomly sampled source segments of a size smaller than 1kbp that were concatenated and copied in random order to a TI thread with 50 segments. We then applied rayas using a range of parameters denoted as (A,B) in the legend that control the required increase in coverage (A) and the minimum number of required split-reads at the boundary of a segment (B). The default of rayas uses A=3 and B=3 (green plotting symbols).

We note that larger numbers of long-read cancer datasets will be required in the future to fully characterize the TI thread pattern that we describe in our study. We further caution that the currently known spectrum of bridges and cycles of templated insertions (Li et al.), tandem short template jumps (Umbreit et al.) and TI threads (our study) may represent only a subset of complex structural rearrangements occurring in cancer that involve templated DNA sequences. We therefore believe that long-read sequencing of cancer genomes will be paramount to unravel the rearrangement complexity of cancer genomic landscapes fully in the future.

Reviewer #2–9: The splicing analysis around the templated insertion thread is confusing since the primary tumor which harbors the SVs is the sample that lacks any aberrant splice isoforms. As a result it seems that the genomic change is either functionally inert, very subclonal, and/or not chromosomally integrated, as well as artifactual.

Response: We agree that the splicing analysis is not directly straightforward to interpret. To clarify, all of the splice junction start sites that are highlighted are novel sites (not present in Ensembl), as illustrated in the new Supplementary Figure S35. In this specific analysis we tried to show that there are functional effects near the TI thread insertion sites. To limit alternative options we tried to assess other potential confounders. (a) Not chromosomally integrated: We generated a Circle-Seq DNA library of the primary tumor and investigated whether we see any evidence of ecDNA signatures at the source segments of TI threads. As shown in Figure R2.13, this is not the case and we therefore assume TI threads are chromosomally integrated. (b) subclonal: Based on an analysis of the major and minor copy-number of all genomics segments, we indeed see strong evidence of large intra-tumor heterogeneity with many sub-clonal copynumber changes (Figure R2.14, new Supplementary Figure S3) but by remapping reads to the Shasta TI thread contig (Figure R2.11, new Supplementary Figure S19) we do observe a median coverage of 23x, exceeding the coverage one would expect for a subclonal event affecting only

one haplotype. We also generated additional matched patient-derived xenograft data. Analyzing the PDX data confirmed that the clone with the TI thread engrafted in the mouse and recapitulates the TI thread pattern (Figure R2.5, new Figure S11).




Figure R2.13: The source segments of TI threads do not show a signal in the new Circle-Seq data, indicating absence of TI threads on ecDNA.

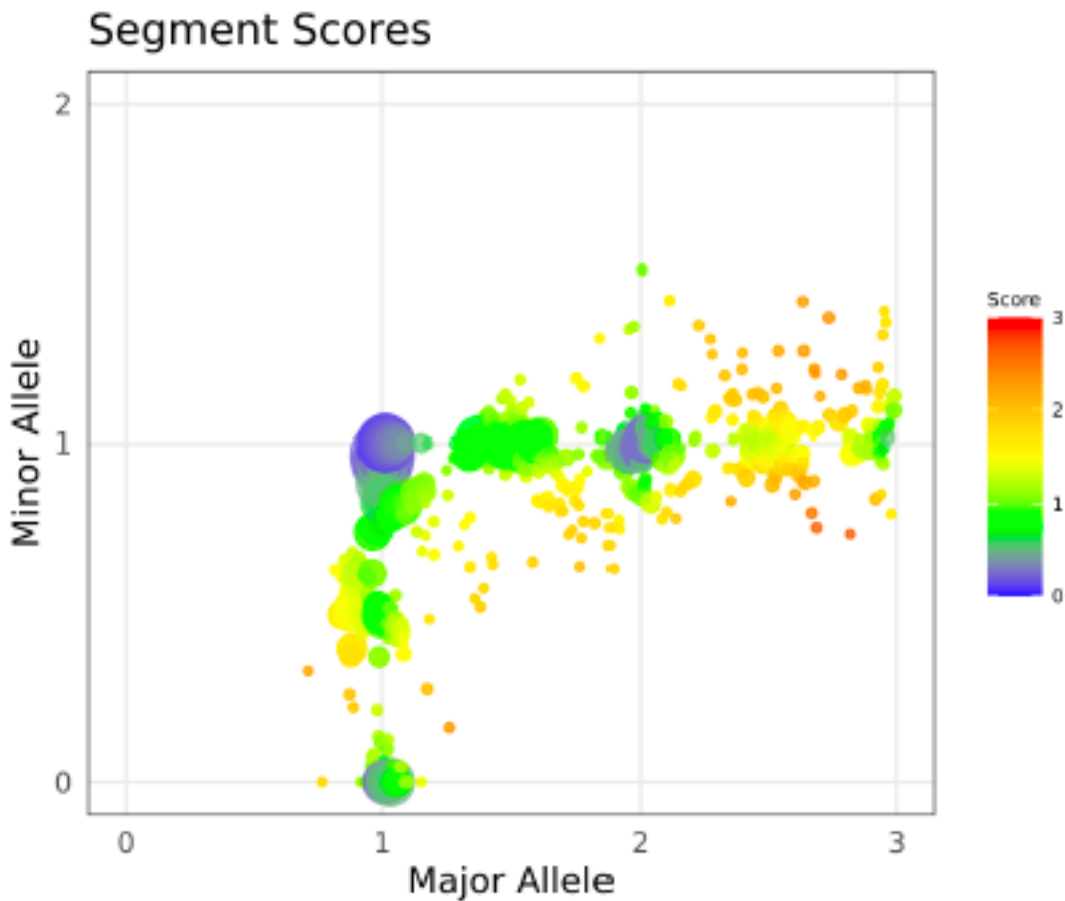


Figure R2.14: Major and minor allele of each copy-number segment as computed by Purple (Shale et al. 2022). Many segments do show a non-integer copy number as is the case for subclonal changes.

Reviewer #2–10: Related to above, what evidence is there that these patterns are clonal and chromosomal? They are absent in the relapse, so perhaps they might be a transient and amplification process like what has been noted by Jonathan whetstine's group.

Response: As mentioned above, the TI thread contig of the Shasta assembly has a median coverage of 23x, exceeding the expected coverage (<<15x) of a subclonal haploid rearrangement (see Figure R2.11, new Supplementary Figure S19). The confirmation of the TI thread in a matched patient-derived xenograft model suggests that the pattern is not a transient site-specific

DNA copy gain as described by Black et al. from Jonathan Whetstine's group. We further performed Circle-Seq of the primary tumor sample to assess the hypothesis that TI threads are part of extrachromosomal DNA, and found no overlap between Circle-Seq data and the source regions of the TI threads (experiments described in more detail in response to comment #2-9). The absence in the relapse is described in detail in our response to comment #2-11 below.

Reviewer #2-11: Otherwise how do the authors explain the disappearance of the pattern in the relapse. Is there a deletion of this locus in the relapse, or did the relapse arise from a distant ancestor of the primary?

Response: In response to this reviewer comment we performed a systematic comparison of somatic variants found in the primary tumor and relapse, which showed that the primary tumor and the relapse indeed arose from a distant common ancestor – with both samples sharing only 34% of all somatic SNVs (Figure R2.15, new Supplementary Figure S21). Additionally, both samples share the CS11-17 rearrangement. The relapse lacks 18% of the somatic SNVs present in the primary tumor, and acquired a number of additional somatic SNVs (48%) missing in the primary tumor, some of which are likely to be biologically relevant such as a frameshift variant in *SUFU* (Figure R2.16, new Supplementary Figure S22). Notably, the complex rearrangements involving chromosome 7 of the primary tumor are not present in relapse, including the TI threads (Figure R2.17, new Figure S20). These points are now clarified in our revised manuscript.




Figure R2.15. Comparison of the somatic variant allele frequency of all identified coding mutations. We added randomly ± 0.01 in x- and y-direction to better highlight the variants that have been lost or acquired in relapse.




Figure R2.16. A homozygous 2bp frameshift insertion in SUFU that is present in the relapse sample (top) but absent in blood (middle) and the primary tumor (bottom).




Figure R2.17. Read-depth ratio plots of the primary tumor and relapse computed by HaTCHet (Zaccaria et al. 2020). Notably, the complex rearrangements on chromosome 7 are not present in relapse, including the TI thread involving segments of chromosome 7 and chromosome 4.

Reviewer #2–12: Beyond the templated insertion finding, it is unclear what qualitative insight the long reads give here. In fact it seems that many of the findings would be evident using short reads.

Response: Several recent reports (Nattestad et al., Chaisson et al., Ebert et al.) have highlighted the fact that short-read sequencing is not well suited for the detection of somatic and germline

structural variants, and ultimately provides an incomplete view of complex SVs in the genome. TI threads exemplify the shortcoming of short reads particularly well, with prior studies based on Illumina sequence data not capturing this rearrangement pattern in spite of its high prevalence in certain types of cancer. We stress that owing to the limitations of Illumina sequencing, prior studies based on short reads have focused their in-depth analyses of SVs on structures exhibiting low to intermediate levels of complexity (see e.g. Li et al. 2020) or aggregate measures of complexity (e.g., junction copy number counts in Hadi et al. 2020), without providing full DNA reconstruction of such events. We are unaware of any short-read-based study that could contiguously assemble double-minute chromosomes or larger rearrangement structures, as we achieved for the CS11-17 assembly. The assembly of complex structural variants using ONT sequence data is likely to facilitate variant characterization efforts providing further insights into the genesis of somatic rearrangements in cancer and may potentially inform future clinical applications. Additionally, we made use of the long ONT reads to jointly analyze genetic and epigenetic information (which would require additional bisulfite sequencing in the context of Illumina read data), and have exemplified the utility of long reads for characterizing rearrangements involving telomeric sequences.

Reviewer #2–13: For example chromothripsis alleles like in 1C can be inferred from short reads through the use of allelic imbalance (eg RCK Aganezov Raphael Genome Research 2020) though potentially with some ambiguity ie multiple possible reconstructions.

Response: Short-read based approaches have indeed been very successful in inferring chromothripsis events (see e.g. Korbel & Campbell, Cell, 2013), which our revised Discussion now clarifies. Nevertheless, as alluded to above, short-read based approaches lack the contiguity to fully represent complex SV structures. Even local assembly approaches such as SvABA (Wala et al. 2018) provide only a limited view into such events, as exemplified in Figure R2.18 and R2.19 (below) where we remapped the SvABA short-read derived contigs to the Shasta assembly. Tools that employ allelic imbalance rely on the allelic depth of variants, a conceptual strategy which inevitably fails in the presence of short DNA segments that lack heterozygous SNP markers – which represents another limitation of short reads now discussed in our revised manuscript (see response to comment Reviewer #2–19).

Reviewer #2–14: The comparison to short reads should also not just rely on a discordant pair based tool like DELLY, which only gives approximate breakpoints and limited resolution of complex regions. An assembly based algorithm, like GRIDSS or SVABA should work better as a comparison to short reads.

Response: We thank the reviewer for this suggestion and we applied SvABA (Wala et al. 2018) to the primary tumor. We remapped all SvABA contigs to the Shasta assembly of the primary tumor to highlight the differences in contiguity and breakpoint junction reconstruction using short and long-reads (Figure R2.18). All SvABA contigs that mapped to the TI thread captured at most 3 SV breakpoints (Figure R2.19, new Supplementary Figure S12), which emphasizes the requirement for long reads to characterize the most complex DNA rearrangement structures (such as TI threads). This analysis has now been included in the revised manuscript in the TI thread section and under Methods.



Figure R2.18: Comparison of short-read derived SvABA contigs (top panel) aligned to the assembled TI thread. Aligned long-reads are shown in the bottom panel.

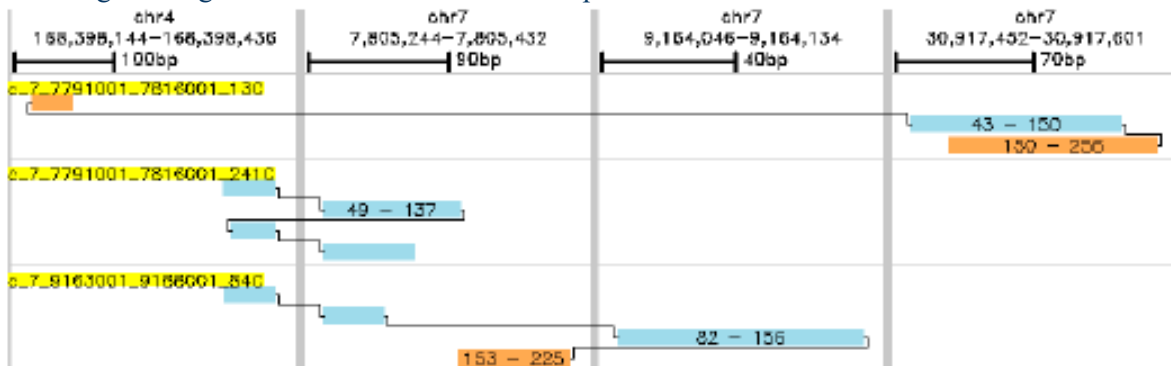


Figure R2.19: Short-read SvABA contigs recapitulate individual SV breakpoints of the TI thread. Each of the 3 contigs (top panel, middle panel, bottom panel) support at most 3 SV breakpoints but the characteristic self- and cross-concatenation of TI threads is evident.

Reviewer #2–15: While long reads can provide direct evidence to thread distant rearrangements, the relatively short long reads in these data (5kb) could only do this for nearby rearrangements ie those that are within 5kbp of each other. Given this limitation, the authors appear to be taking a more indirect / greedy approach to reconstruct the larger alleles eg those associated with chromothripsis in Fig 1, which though reasonable could be done just as well with short reads to the same effect. It appears therefore that this feature of the long read data and analysis are overstated.

Response: As alluded to above, 5kbp is the median read length and the typical read length distribution of ONT shows a long tail (see Figure R2.1). This enabled read-based phasing with an estimated N50 phased block length of 4.68Mbp. Contrary to statistical phasing using haplotype reference panels, which is often employed in short-read studies, long-read based phasing

provides a well-defined start and end position for each haplotype block which enables the targeted selection of reads for local assembly using lorax. In response to this reviewer comment, and several comments above (e.g. #2-12, #2-13, #2-14) we have amended the Discussion section to ensure proper discussion of the respective advantages and disadvantages of long- versus shortread sequencing.

Reviewer #2–16: What is the broader relevance of these patterns to medulloblastoma? The PCAWG analysis does not seem to include a substantial number of medulloblastoma cases.

Response: Based on our PCAWG analysis, TI threads are only occasionally seen in medulloblastoma patients. This is in line with the fact that TI threads are tightly linked with chromothripsis, which plays a major role in LFS-linked medulloblastoma but does not show a high prevalence across other (non-LFS) medulloblastoma subgroups (see Rausch et al. Cell 2012). TI threads do occur more commonly in several tumor types with notable enrichments in dedifferentiated liposarcoma samples – in which we undertook validation studies at revision stage – underscoring the broader relevance of this rearrangement pattern.

Reviewer #2–17: As a data resource, the study is limited with one case and two tumor samples. Also the data quality is somewhat questionable, with read lengths that are somewhat short (5 kbp) even for ONT which is usually 10-50 kbp even in the lower cost / high throughput mode. Larger and seemingly higher quality cancer long read datasets have been published over the past several years eg in genome research eg Nattestad 2018, aganezov..schatz 2020, sakamoto 2020.

Response: We stress that our study used primary tumor patient samples, in contrast to the cancer cell lines used by the previous studies the reviewer refers to – i.e. breast cancer cell lines (Nattestad 2018 and Aganezov 2020) and lung cancer cell lines (Sakamoto 2020). Cell lines allow the isolation of high molecular weight input DNA at significantly higher quantities than primary tumors, which at least in part is likely to explain differences in read length. Notably, our own ONT data from the matched patient-derived xenograft model (for which we were less limited in terms of amount of material as compared to the matched primary tumor) exhibits a median read length of 10,529bp (N50 read length: 41,030bp) – with some extremely long reads >100,000bp (see e.g. the TI thread confirming read shown in Figure R2.5). Nevertheless, in response to this reviewer comment we attempted to improve data quality by re-basecalling our entire data set with the latest guppy version, which improved the estimated read-error rate by more than 2% in all 3 samples (primary tumor, relapse and germline), improved the demultiplexing by 3.1%, increased the number of mapped reads by 2.3%, and improved phase-ability of mapped reads by 3.8%. None of our main conclusions changed as a result of these optimizations.

Reviewer #2–18: Given that there is one case profiled here, the authors are limited to Cherry-picking anecdotal observations without statistical basis. As a result, there appears to be limited insight that can be drawn from these data into medulloblastoma or more generally cancer development and post-therapy relapse. It's unclear in particular what additional insight is drawn from the relapse sample. This includes the methylation findings, which are of uncertain significance and mostly validate previous findings about ASM that did not require long reads.

Response: As alluded to above, there is a scarcity of primary tumor samples in the current longread literature. Given the limitations of short-read sequencing there is an acute need to develop methods to take full advantage of the increased read lengths provided by ONT and PacBio in the context of cancer patient primary samples. As exemplified by our study, short-reads have limitations for the reconstruction of complex SV events like TI threads – and our revised

manuscript now importantly shows that this pattern is not an anecdotal observation but can be successfully validated in an important sarcoma subtype.

In response to this reviewer comment, we have further amended the Discussion section to allude to the future need for a larger number of long-read primary cancer patient datasets to examine genetic (and epigenetic) patterns only partially accessible in short read datasets.

Reviewer #2–17: As a data resource, the study is limited with one case and two tumor samples. Also the data quality is somewhat questionable, with read lengths that are somewhat short (5 kbp) even for ONT which is usually 10–50 kbp even in the lower cost / high throughput mode. Larger and seemingly higher quality cancer long read datasets have been published over the past several years eg in genome research eg Nattestad 2018, aganezov..schatz 2020, sakamoto 2020.

Response: We stress that our study used primary tumor patient samples, in contrast to the cancer cell lines used by the previous studies the reviewer refers to – i.e. breast cancer cell lines (Nattestad 2018 and Aganezov 2020) and lung cancer cell lines (Sakamoto 2020). Cell lines allow the isolation of high molecular weight input DNA at significantly higher quantities than primary tumors, which at least in part is likely to explain differences in read length. Notably, our own ONT data from the matched patient-derived xenograft model (for which we were less limited in terms of amount of material as compared to the matched primary tumor) exhibits a median read length of 10,529bp (N50 read length: 41,030bp) – with some extremely long reads >100,000bp (see e.g. the TI thread confirming read shown in Figure R2.5). Nevertheless, in response to this reviewer comment we attempted to improve data quality by re-basecalling our entire data set with the latest guppy version, which improved the estimated read-error rate by more than 2% in all 3 samples (primary tumor, relapse and germline), improved the demultiplexing by 3.1%, increased the number of mapped reads by 2.3%, and improved phase-ability of mapped reads by 3.8%. None of our main conclusions changed as a result of these optimizations.

Reviewer #2–18: Given that there is one case profiled here, the authors are limited to Cherry-picking anecdotal observations without statistical basis. As a result, there appears to be limited insight that can be drawn from these data into medulloblastoma or more generally cancer development and post-therapy relapse. It's unclear in particular what additional insight is drawn from the relapse sample. This includes the methylation findings, which are of uncertain significance and mostly validate previous findings about ASM that did not require long reads.

Response: As alluded to above, there is a scarcity of primary tumor samples in the current longread literature. Given the limitations of short-read sequencing there is an acute need to develop methods to take full advantage of the increased read lengths provided by ONT and PacBio in the context of cancer patient primary samples. As exemplified by our study, short-reads have limitations for the reconstruction of complex SV events like TI threads – and our revised manuscript now importantly shows that this pattern is not an anecdotal observation but can be successfully validated in an important sarcoma subtype.

In response to this reviewer comment, we have further amended the Discussion section to allude to the future need for a larger number of long-read primary cancer patient datasets to examine genetic (and epigenetic) patterns only partially accessible in short read datasets.

Reviewer #2–19: The provided tools for amplicon assembly and templated rearrangements seem heuristic and do not show a clear advance over existing graph tools for SV analysis, including RCK (Aganezov), Linx (Shale et al 2022), JABBA (Hadi 2020), AmpliconArchitect (Deshpande 2019), and AmpliconReconstructor (Luebeck 2020). Also it appears that the algorithm greedily looks for a single thread through the graph rather than trying to account for all the copies of segments in a region. This suggests that the authors are not considering alternate

reconstructions, and thus possibly overstating their confidence in a given reconstruction.

Response: We thank the reviewer for suggesting further amplicon reconstruction tools, which we have explored when preparing our revised manuscript. We succeeded running both RCK and Linx, and we discuss the output of both methods below. We were unable to run AmpliconReconstructor because it requires BioNano optical mapping data which is not available for our Medulloblastoma sample. Jabba repeatedly failed after the installation and for AmpliconArchitect we ran into licensing issues for the Mosek optimizer.

RCK requires input in the form of segmented allele-specific copy numbers and novel genomic adjacencies. The somatic SVs of Manta were used as novel adjacencies and the allele-specific segment copy numbers were computed using HaTCHet – with both Manta and HaTCHet being the recommended input methods for RCK. However, HaTCHet, perhaps not unexpectedly because of the nature of TI threads, failed to identify templated insertion source segments at high-copy number because a significant number of these (short) segments do not even have a single (phased) germline SNP. Manta reported a multitude of somatic SVs connecting the templated insertion source segments, e.g. 20 somatic SVs for the region chr4:168,398,000-168,399,000 and 18 for the region chr7:7,805,000-7,806,000 in Figure 2B. We assume that the lack of a distinct copy-number segment matching these templated insertion source sequence regions and the high number of somatic SVs caused RCK to infer huge segments on chr4 and chr7 with non-plausible copy-number values (see Figure R2.20) – suggesting that high complex rearrangement patterns such as TI threads cause artifacts when using RCK. This indicates that RCK has an insufficient resolution for characterizing TI threads. Similar to RCK, Linx also requires genomic adjacencies and allele-specific copy-numbers as input. The recommended tools are GRIDSS and PURPLE which we could successfully install and run. For the CS11-17 assembly with larger segments, Purple confirmed our estimate of a median total copy-number of 3.88 with a major copy-number 2.7. The estimate for contig 2 which likely belongs to the CS11-17 assembly was a total copy-number of 3.77 and a major copynumber of 2.97. Linx also clustered the segments belonging to the CS11-17 structure into a single complex event, including segments overlapping contig 2. The chaining algorithm, however, failed to predict the entire derivative structure and outputted 14 independent chains for this cluster. Chain 1 was the longest that shared 11 out of 15 (73%) segments of contig 1 (see Figure R2.21, new Supplementary Figure S43). Like HaTCHet, Purple failed to identify the total and major copynumber for the segments involved in the TI thread. However, due to Linx' heuristic approach with several rules and clustering routines, the algorithm still managed to cluster all SVs related to the TI thread together in a giant complex event with 97 chains and 827 SVs together with many additional SVs from the massive chromothripsis event involving chr4, chr5 and chr7 (among others). Chain 2 best overlapped the TI thread presented in Figure 2 but greatly underestimated the true number of junctions with only 43 compared to 231 estimated from the Shasta assembled TI thread contig (see Figure R2.22, new Supplementary Figure S44) – suggesting clear limitations of Linx with respect to the characterization of such complex rearrangement structures.

In summary, our analyses show that TI threads clearly escape discovery with tools commonly used by the short read sequencing community, such as RCK and Linx. These allele-specific copynumber tools rely upon (phased) heterozygous germline SNPs, and thus biased against the structure of TI threads containing many small-scale amplicons/rearranged segments (which often do not contain any SNPs). Popular SV detection methods such as Delly, Manta and Gridds identify a multitude of somatic SVs at these loci by finding clusters of paired-ends and split-reads that do not support the reference allele. This clustering step can miss true adjacencies in close proximity and collapse adjacencies occurring more than once in a TI thread which ultimately leads to inconsistent input data of genomic adjacencies and allele-specific copy-number calls for genome reconstruction. For RCK, this fact posed a serious issue whereas Linx still managed to cluster the SVs albeit failed to resolve the TI thread structure.




Figure R2.20: A and B allele copy-number reconstruction output from RCK. Presumably the large number of somatic SVs incurred by templated insertions near chr4:168Mbp and chr4:123Mbp in conjunction with their expected absence in the output of allele-specific copy-number callers like HaTCHet caused an incorrect reconstruction of the A and B allele copy-number of chr4.




Figure R2.21: Linx visualization of the complex event overlapping the CS11-17 structure with chromosomal segments on the outer ring, chaining of variants as colored lines and copy-number gains (green) and losses (red) in the inner circle. For a detailed guide see Shale et al. 2022.




Figure R2.22: Linx visualization of the complex event overlapping the TI threads described in this manuscript. For a detailed visualization guide see Shale et al. 2022.

Importantly, lorax does not greedily collapse each connected component of the rearrangement graph into a single TI thread reconstruction but instead collects all supporting long reads for a targeted local assembly that can emit multiple contigs in case of TI thread heterogeneity. Our method thus leverages the long sequencing length to separate alleles or tear apart multiple distinct integrations as we could show for the validation sample P2. We are sorry for this misunderstanding and clarify in the revised methods description this important distinction to shortread reconstruction methods.

Technical and minor comments:

Reviewer #2–20: Comparison of primary and relapse in general does not seem to yield significant biology, but appears at least in part limited by the lower sequencing depth in the relapse. Unclear what is the significance of the observation that TI-threads are missing in the relapse? Were these alleles lost or just missed due to low read depth or perhaps because the samples were analyzed for SVs separately. One approach to integrate analyses of the two samples is to align the relapse reads to the primary tumor contigs, and vice versa.

Response: We followed the reviewer's suggestion and aligned the relapse reads back to the primary tumor contig from the Shasta assembly that confirmed the TI thread. This analysis did not identify a single confident mapping ($\text{MAPQ} > 1$) to the TI thread, which confirms its absence in the relapse tumor. We note that this observation is entirely consistent with our prediction that the primary tumor and the relapse arose from a distant common ancestor (see comment #2–11).

Reviewer #2–21: On page 12 "Since the CS11-17 rearrangement occurs in only one haplotype" .. is this a premise or a result of the analysis? ie are the parental haplotypes long enough to make this statement from the data or are the authors just applying a biological assumption. This should be clarified.

Response: The long-read phased blocks are generally long enough for haplotype separation with an estimated N50 phased block length of 2.29Mbp in the primary tumor using WhatsHap.

Reviewer #2–22: Contig 2 is mysterious - it is somehow associated with the CS11-17 but does not seem to be rearranged (Fig 4)? Are there really no rearrangements around this segment? How is it known that it is associated with CS11-17 and not part of an independent event?

Response: This is a good point. Our interpretation is that contig 2 is likely linked to CS11-17 as both share a similar copy-number and because reads aligning at the ends of contig 2 map to the peri-centromere of chr17 as is the case for the ends of contig 1. Despite the long reads, we could not resolve the peri-centromeric regions as none of the reads extended beyond the repeats (see Figure R2.23, new Supplementary Figure S6) – which underscores the need for further genomic technology developments that enable routine sequencing of cancer genomes with particularly long reads to reconstruct the most complex rearrangement classes (as we point out in detail in the discussion of our revised manuscript).

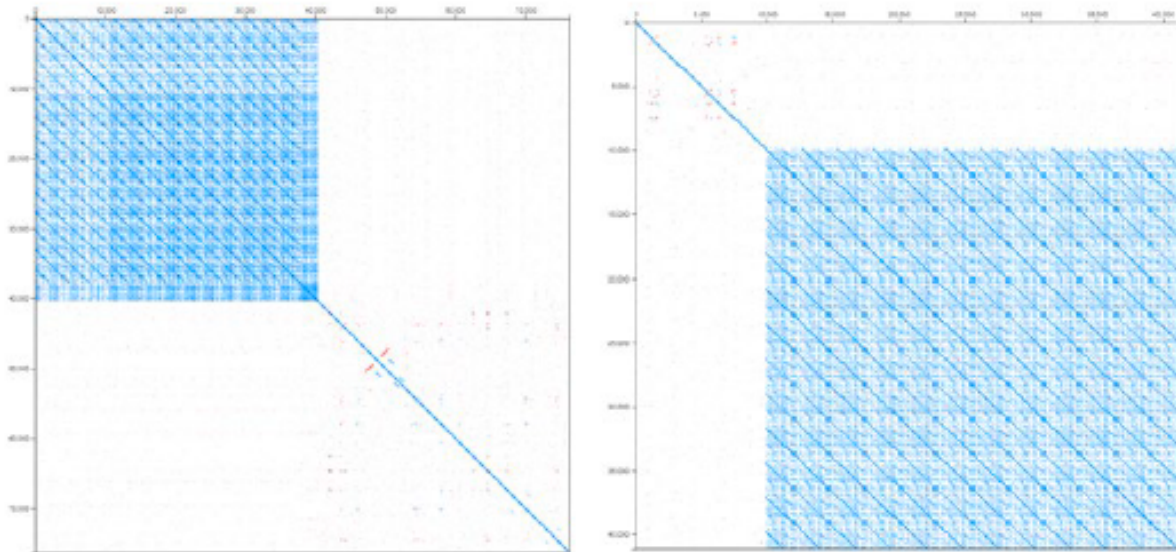


Figure R2.23: Dot plots of the longest read against itself at the left and right end of contig 2. For both reads, one can observe how reads start in a unique sequence on contig 2 but then extend into repetitive arrays.

Reviewer #2–22-B: Copy number of many segments in Fig 2B appears to be at 3.5. Is this a subclonal copy number change or due to a bad fit of purity/ploidy ?

Response: We applied the Gridss-Purple-Linx pipeline (Shale et al. 2022) of the Hartwig Medical Foundation which estimated the tumor purity to be close to 100% (range 93% - 100%). The tumor is mostly diploid (overall ploidy is 1.98) but shows extensive intra-tumor heterogeneity (see Figure R2.14) – suggesting that the estimated copy-number value of 3.5 results from the multi-clonality of the tumor.

Reviewer #2–23: It is unclear whether the telomere fusions represent the connection of a telomere sequence to a chromosomal sequence, creating a new chromosome end, or the end to end fusion of two distinct telomeres. Usually "telomere fusions" refer to the latter pattern. It seems also that the former outcome should be also detectable with short reads?

Response: We performed telomere FISH on metaphase spreads from the matched PDX model and did not detect any signal for the telomere probe within the chromosomes, only at the chromosome ends, which suggests a subtelomeric dsDNA break and fusion rather than End-To-End fusion.

Reviewer #2–24: The observation of hypomethylation of contig 2 (Fig 4) is potentially interesting. Unclear whether this is the consequence of rearrangement or just a pre-existing aspect of this haplotype ie prior to cancer development. Since contig 2 does not appear to be connected to anything it is unclear

Response: While causality (or order of causality) indeed cannot be inferred without a matched normal sample for the same tissue type (cerebellum), we have now made it a point in the manuscript to clarify that demethylation of this haplotype is only observed in the chromothriptic haplotype, in both cancer samples, and not in the blood sample (Figure R2.24).




Figure R2.24: Methylation state of CS11-17 ctg2 in germline, primary tumor and relapse sample, in both haplotypes. Demethylation of the chromothriptic haplotype can be observed in both primary tumor and relapse.

Reviewer #2–25: ASM analysis (p11) does not seem to distinguish between promoter-associated CpG islands vs shore or open sea CpGs. How do the authors explain allele-specific methylation outside of CGIs? Also, the comparison between primary and relapse seems to be fraught due to low coverage in the relapse (Fig 3A). The allele-specific methylation changes do not seem to be associated with a consistent impact on expression.

Response: When investigating the link between (allele-specific) methylation and expression, we focused on ASM effects intersecting with gene promoter regions in order to measure promoter methylation, regardless of whether these occur in annotated or putative CGIs. Aside from expression effects, however, we also set out to show the power to detect ASM independent of context, for example in order to study disease-associated demethylation (as shown in CS11-17). Regardless, in Figure R2.25 we present a breakdown of methylation effects into CGI, Shelf, Shore, and open sea, showing that we find predominantly open-sea methylation effects, albeit with a lesser effect size.

In order to address the issue with low coverage affecting discoverability in the relapse sample, we have revised the DMR and ASM analyses in the manuscript with a single consistent segmentation. Unlike in the original submission, where a separate segmentation has been obtained for each sample for ASM, now all analyses (DMR between samples and ASM) are performed on the same methylation segmentation, produced from both tumor samples combined, with sample and haplotype information (i.e. a 6-sample segmentation, primary HP1, primary HP2, primary unphased, relapse HP1, relapse HP2, relapse unphased). While low coverage is still a limiting factor for test power, this improved the discoverability of ASM in the relapse sample by 215%.

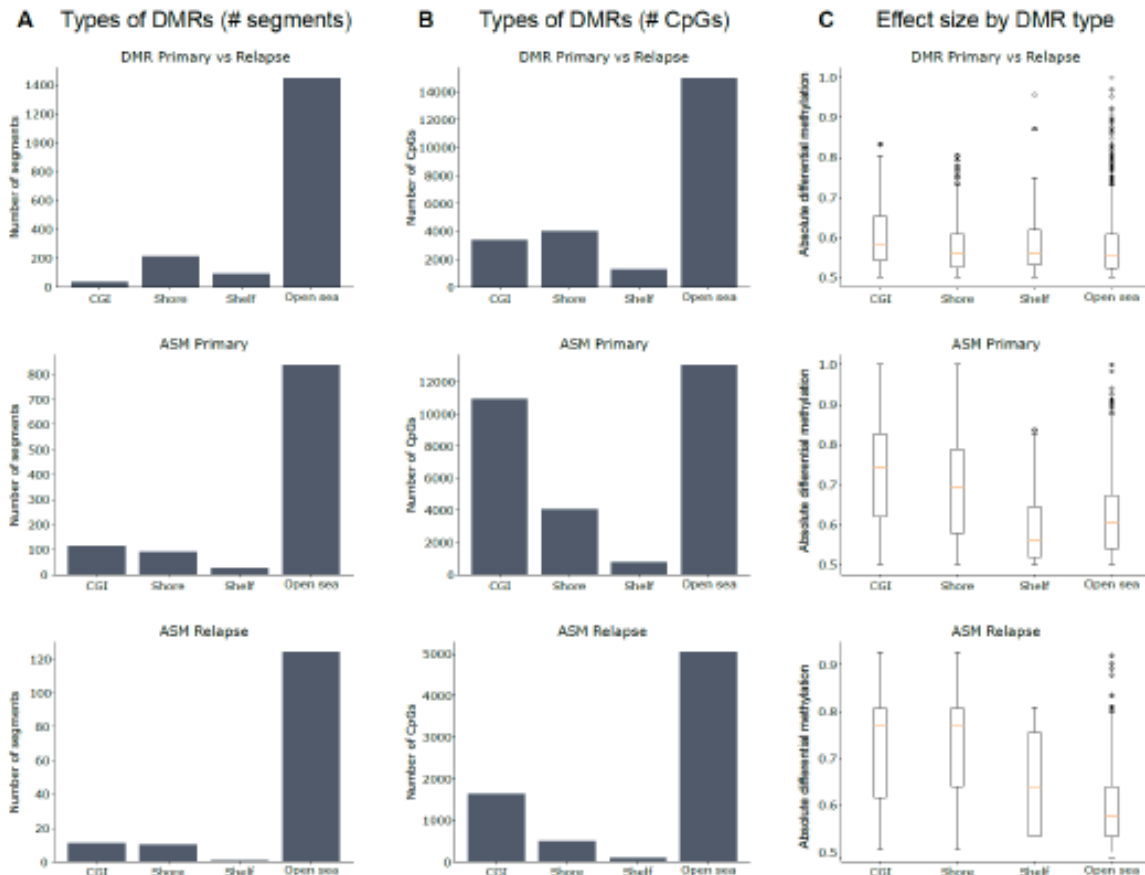


Figure R2.25: DMRs distinguished into CpG-Island, shelf, shore, and open-sea.

Reviewer #2-26. For Fig 4A the two regions are different

Response: We thank the reviewer for pointing out this in Figure 4A, which we have fixed.

Reviewer #2-27. How much more of the genome is mappable with 5 kb ONT reads relative to standard illumina 150 bp reads? The noisiness of ONT reads including the high rate of gaps makes them harder to map for a given length, so I wonder if there are many more rearrangements caught in repetitive regions for instance.

Response: For the primary tumor, 97.4% of the non N masked genome is mappable with short reads and 97.8% is mappable with long reads, considering a genomic position mappable if it is covered by at least 5 reads of mapping quality greater or equal to 10. It should be noted, however, that the strength of long-read sequencing is not only its increased assignment ability, but in particular its ability to reconstruct complex SV sequences and map insertions that cannot be captured with short reads. Consistent with the latter, our analyses show that the most pronounced difference in SV calling is for (germline) insertions where we successfully ascertained 6,561 insertions with ONT but less than 200 with short-reads.

Reviewer #2-28. Comparisons are made to bisulphite sequencing (eg p11) but no direct comparison to WGBS profiling appears to be done.

Response: We appreciate the question. While presenting our ASM analysis we attempted to

perform a synthetic comparison to what could have theoretically been achieved with WGBS. In this comparison, we assume a WGBS experiment with 150bps read-length such as on the HiSeq3000 platform. We make the generous assumption that a single heterozygous variant (which isn't a C>T variant) can lead to successfully phasing methylation calls from WGBS and detecting the ASM, as long as it's not repeat-masked. We have clarified this section of the results to make this synthetic comparison more clear.

Reviewer #2–29. p9 comparison to arrays is a bit of a straw man? "Overall, analysis of the ONT data provides a substantially more comprehensive picture of the tumor methylome, with 78% of the between sample DMRs inaccessible to the commonly used 450K array, and 65% inaccessible to the 850K array". It seems that a very similar statement could be made for WGBS.

Response: We agree with the reviewer that the phraseology in the original manuscript could be improved here. We now rephrased it to make it less of a note on comparing technologies, and more of a note about where methylation effects occur (largely out of scope of array technologies).

Reviewer #2–30. Fig 1C is an unnecessarily awkward and confusing visualization of a simple concept, which is the alignment of a multi-part contig to the reference genome. In particular, the placement of the purple / orange stacked bars on the x axis at 55.2 M is confusing. It appears that this should be placed outside of the x axis ie the x axis should be split into two axes. However it would seem better to replace this visualization with something more like Fig S9.

Response: We followed this reviewer's suggestion and now use a chained alignment in the revised Figure 1 (as in the old Fig. S9).

Reviewer #2–31: The approach to coloring and patterning in figures seems suboptimal and in many cases collides across different figure elements which is confusing - examples include 4A, 4C, 4D. in particular the use of diagonal patterns to denote haplotypes eg in Fig 3 and 4 seems unnecessary, it would seem better to just use text labels.

Response: We thank the reviewer for the feedback on our presentation. We agree, and have changed to using text-labels to distinguish between haplotypes throughout the manuscript, getting rid of the hatch pattern which can be hard to distinguish.

Reviewer #3:

Reviewer #3–general comments: "Long-read sequencing of diagnosis and post-therapy medulloblastoma reveals complex rearrangement patterns and epigenetic signatures", Rausch et al. demonstrate the promise of long read sequencing for characterizing complex rearrangements patterns in tumor genomes. Specifically, the authors apply Oxford Nanopore Technologies nanopore sequencing to a primary medulloblastoma, its relapse and a matched normal blood sample from a patient with Li-Fraumeni syndrome. They leverage the long reads to obtain extensive phasing information and resolve the structure of various complex somatic rearrangements, most notably a novel type of pattern dubbed "templated insertion thread". By identifying its footprints in the short-read based PCAWG cohort, the authors demonstrate it is often associated with chromothripsis and highly frequent in liposarcoma. The authors also explore the DNA methylation patterns encoded by the single molecule nanopore reads to reveal differential and allele-specific methylation patterns, as well as integrate further datasets of shortread gene expression data.

The work is extremely comprehensive, and the paper reads fluently. The authors have gone through great lengths to ensure the manuscript and its associated methods represent a solid

basis for both the long read and cancer genomics communities to build on. To this end, they have developed several new computational methods: (i) rayas which detects signatures of templated insertion threads in short-read sequencing data; (ii) wally, a package for visualization of read alignments from BAM/CRAM files, most notably at and around (structural) variants; (iii) lorax, a toolkit for analysis of long read data, specifically aimed at cancer genomics. All of these tools, plus a set of further analysis scripts used in the manuscript, are versioned, well documented and made available on GitHub. For the users' convenience, precompiled binaries, as well as Dockerfiles and Singularity containers are provided, in what this reviewer considers a truly impactful and exemplary contribution to the scientific community.

I have no major concerns, but would like to see a few points extended or addressed in a revised version of the manuscript:

Response: We thank the reviewer for their interest in our work and their very positive words. Please find a point-by-point response to the comments below.

Reviewer #3–1: As indicated by the authors, nanopore sequencing coverage, especially for the normal and relapse samples is relatively low, and the median read length is somewhat limited compared to recent standards. Likewise, more recent base calling algorithms such as the super accuracy models incorporated in Guppy v5 and 6 could further improve analysis results (the authors use Guppy v4 and a high accuracy model). While I do not expect the authors to acquire additional sequencing data or rebasecall their raw data and redo the full analysis, these points could be included in the discussion to guide people in the field on what to aim for and what to expect when deciding to pursue long read sequencing of tumor genomes.

Response: We agree with the reviewer that it is important to highlight all limitations of the study. We have clarified limitations pertaining to the read depth in the discussion (see below). Given the very rapid developments in the area of basecalling, we decided to re-basecall our data using Guppy V6.1.7 with the super high accuracy model with modbases – and we subsequently observed an increase of approximately 2.64 on read quality PHRED score. Through improvements in sequence quality and demultiplexing we now achieve an approximately 11% increase in the number of usable reads, we have added this information to the updated discussion.

Reviewer #3–2: The study includes both a primary and matched relapse sample, with notable differences in copy number and structural variation (e.g. absence of the templated insertion threads from the relapse), potentially offering insights into the tumor's evolutionary history as well as the mechanisms and ordering of some of the rearrangements. These aspects are only briefly touched upon in the manuscript however.

Response: As suggested by this reviewer, our revised manuscript now provides more details on the tumor's evolutionary history. Leveraging data from a previous study performing bulk short read sequencing and using allele frequencies to derive phylogenies, we identified the major clones in the primary tumor and the relapsed sample (Parra et al. 2021, Figure R3.1 below). Early clonal events include, for example, the loss of chromosome 17p and 3p. In contrast, chromothripsis on chromosome 7 is seen in the primary tumor but not in the patient-matched tumor relapse, perhaps due to the higher susceptibility of the chromothriptic clone to the cancer treatment applied. The templated insertion thread (TI thread), which we found to be tightly linked with the chromothriptic event on chromosome 7, was likewise not present in the relapse. Therefore, these somatic rearrangements, which may have played a driver role in the early stages of tumor development, do not appear to provide a selective advantage after treatment, or alternatively the cells bearing these rearrangements have been eradicated during treatment.




Figure R3.1: For convenience, we show here the Figure 5c panel of the bioRxiv preprint (<https://doi.org/10.1101/2021.06.25.449944>) that analyzed the clonal phylogenetic structure of the primary and relapsed tumor.

Reviewer #3–2–continued (B): For instance, it would be informative to assess which somatic rearrangements are shared between the primary and relapse. Are there any such common variants within the chromothripsis regions on chromosomes 4,5,7,9,16,19 and X?

We clarified the clonal relationship in the above comment #3-2 and discuss the copy-number differences between the primary and relapse in comment #2–11 (Figure R2.17, new Supplementary Figure S20). Briefly, both samples share the CS11-17 rearrangement but the complex rearrangements involving chromosome 7 of the primary tumor are not present in relapse, including the TI threads. Therefore, for somatic SVs, we observe overall only 11% shared SVs between primary and relapse. The relapse gained 19% and lost 70% of all somatic SVs, which is mainly driven by the disappearance of many of the complex rearrangements involving chromosome 4,5,7,9,16,19 and X. These points are now clarified in our revised manuscript.

Reviewer #3–2–continued (C): Similarly, is the chr5p telomere fusion present or absent from the relapse? This could provide insights into the cause of this chromothripsis event.

The primary and relapse originated from a distant common ancestor (see, e.g., comment reviewer #1-2) and many somatic SVs of the primary tumor are not present in relapse (70%), including the chr5p telomere fusion event. We verified the lack of such an SV call orthogonally by mapping the long-reads back to the Flye assembly that included a contig supporting the chr5p telomere fusion. This mapping showed only fragmented read alignments but no SV spanning reads, confirming the absence of the chr5p telomere fusion in relapse.

Reviewer #3–2–continued (D): The authors also note potential ongoing evolution and rearrangements of the templated insertion threads themselves. To assess its stability, would it be possible to estimate how many different alleles are currently detected in this single bulk sample?

Response: As described in our response to Reviewer Comment #2–4, an alignment of TI thread

supporting reads to the Shasta contig that captures the entire TI thread suggests at least 2 alleles that are detectable – at the current ONT read depth (Figure R3.2, new Supplementary Figure S19).



Figure R3.2. Alignment of the raw reads back to the assembled Shasta contig that supports the TI thread shown in Figure 2. The contig is evenly covered and shows a subclonal ~9kbp deletion from 33kbp to 42kbp suggesting at least 2 alleles for this TI thread.

Reviewer #3–3: It may be helpful to provide a read length histogram and read length N50 for the nanopore sequenced samples (in addition to the stats in Table S1). This would further allow people considering long read experiments to readily compare data requirements and quality for the type of analysis presented.

Response: We appreciate the suggestion and have supplemented Table S1 with the relevant statistics. We also included Supplementary Figure S45 (Figure R3.3 below) containing read lengths and PHRED scores for all ONT runs of the Medulloblastoma sample.



Figure R3.3: QC histograms for ONT runs. Left: PHRED quality scores as reported by guppy 6.1.7. Right: read length after QC filtering to reads with read average PHRED score of 7 or higher.

Reviewer #3–4: Do the authors have a good estimate (either through analysis of the sequencing data or from pathology) of tumor purity? This would allow one to assess more accurately potential effects of normal cell admixture on the differential methylation and expression results.

Response: The Gridss-Purple-Linx pipeline (Shale et al. 2022) of the Hartwig Medical Foundation estimated a tumor purity of close to 100% for the primary tumor (range 93%-100%) and 85% for the relapse sample (range 80%-87%). Both tumors are mostly diploid, with an overall ploidy of 1.98 for the primary tumor and 1.94 for the relapse sample. The primary tumor shows extensive intra-tumor heterogeneity (see also Figure R3.1 above), evident from subclonal copy-number segments (Figure R3.4, new Supplementary Figure S3). The relapse sample, due to the clonal selection upon treatment, shows less tumor sub-clonality. Methylation analysis was performed under strict effect size thresholding (only methylation changes of 0.5 or higher were included). This threshold should be sufficient to avoid tumor purity affecting the analysis. See also response #3–12 for a more detailed discussion.




Figure R3.4: Major and minor allele of each copy-number segment as computed by Purple (Shale et al. 2022). Many segments in the primary tumor do show a non-integer copy number as is the case for sub-clonal changes.

Reviewer #3–5: From Figure 1A it appears several large CNVs, such as the losses on chr2 & 3, are not associated with a "large" SV. Were these not identified from the data or are they simply filtered in the display? In either case, the authors may wish to elaborate.

Response: On chromosome 3 the p-arm is monosomic and the q-arm trisomic. We did not identify any large SV near the centromere of chr3. We also filtered SVs for the display and included only intra-chromosomal SVs >10Mbp, as stated in the Figure caption. The chromosome 2 deletion has a clear left-most breakpoint at chr2:57,101,210bp with a likely inter-chromosomal translocation to the chromosome 20 centromere but due to the repetitiveness of the underlying sequence neither delly nor sniffles identified this SV. The right-most breakpoint of the chromosome 2 deletion at chr2:89,753,994 is next to a region masked with Ns in GRCh38.

Reviewer #3–6: Authors state: "The CS11-17 segment, present in both primary tumor and relapse, has a size of 1.55 Mbp; the 17p-arm region affected contains the TP53 locus, which has been lost on the chromothriptic haplotype." Authors may wish to clarify that it is the wildtype TP53 allele which was lost from the chromothriptic haplotype.

Response: We have addressed this omission and rephrased the relevant passage to clarify that the wild-type TP53 allele had been located on the chromothriptic haplotype prior to deletion, this has been updated in the results section.

Reviewer #3–7: At the end of the section "Haplotype-phased assembly of complex somatic rearrangements", the authors note the following: "Yet, we failed to identify reads supporting reintegration of this structure into a chromosomal context, possibly due to limitations of ONT for resolving low-variant allele frequency SVs in conjunction with ITH, especially in complex regions that exhibit repetitive segments larger than the ONT read length". Is it not more likely that the limited coverage read lengths here are to blame for this?

Response: We have considered this point, and after examining the relevant region once again believe that the explanation is likely due to both the read-depth and read-length as well as limitations of ONT for resolving low-variant allele frequency SVs in conjunction with ITH. As alluded to in our response to Reviewer #2–17, we achieved a noticeably shorter read length in our primary tumor samples compared to the matched patient-derived xenograft sample as well as compared to prior studies that focused on cancer cell lines. Shorter read lengths affect both assembly contig lengths as well as the mapping of repetitive sequences. Contig 2 of the CS11-17 assembly is a prime example for this problem, where reads to the left and right of contig 2 map to the peri-centromere of chr17 but none of the reads extended beyond genomic repeat sequences (see Figure R3.5). We have updated the respective section in the manuscript to reflect these points, and have included Figure R3.5 (below) as the new Supplementary Figure S6.

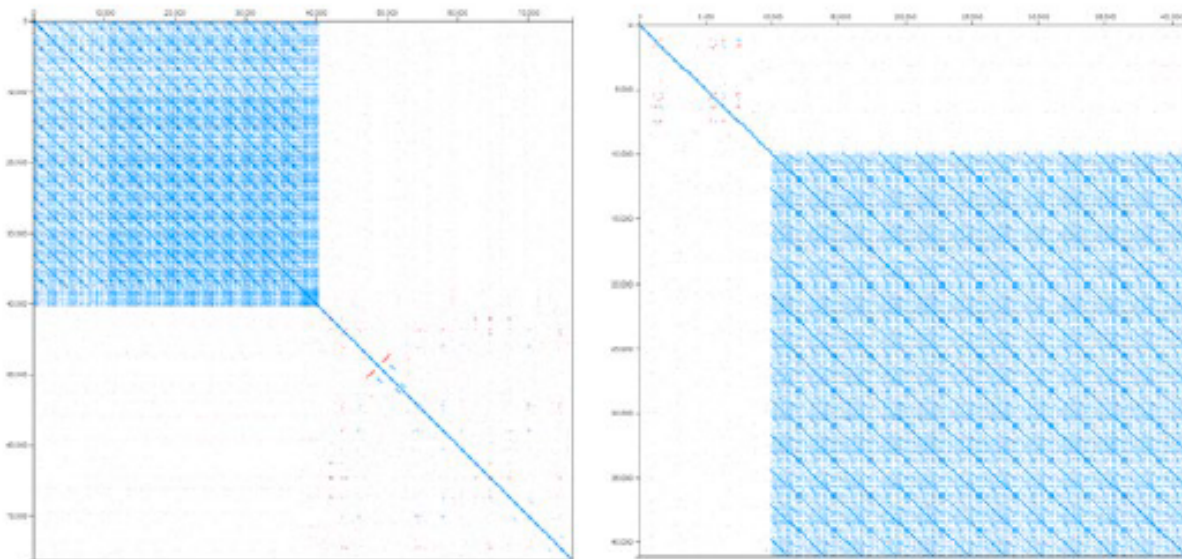


Figure R3.5: Dotplots of the longest read against itself at the left and right end of contig 2. For both reads, one can observe how reads start in unique sequence on contig 2 but then extend into repetitive sequence.

Reviewer #3–8: In Figure 2B, it would be helpful if the top and bottom panels were (better) aligned.

Response: We have fixed this in the revised manuscript, thanks for the suggestion.

Reviewer #3–9: Figure 1B shows non-integer copy number for the first segments on chr5. The authors may wish to elaborate on this (and other) subclonal copy number changes.

Response: As also alluded to above, we first assessed whether the non-integer copy numbers

could be due to a low tumor purity and applied the Gridss-Purple-Linx pipeline (Shale et al. 2022) of the Hartwig Medical Foundation which estimated the tumor purity to be close to 100% (range 93% - 100%). We then examined the major and minor copy-number of all segments, which indicates the presence of extensive intra-tumor heterogeneity (see Figure R3.4), and which can explain the non-integer copy-numbers.

Reviewer #3–10: In the section "Pan-cancer landscape of templated insertion threads in 2,569 tumors.", the new rearrangement pattern is frequently referred to as "template insertion threads" rather than templated insertion threads.

Response: We thank the reviewer for pointing out this textual issue. We moved to referring to templated insertion threads (TI threads) consistently and have updated the text throughout.

Reviewer #3–11: When exploring telomeres, authors note that "Another telomere crisis event observed in the primary tumor likely fused chromosome 19 to the telomere of chromosome 16q". I don't think the presence of this telomere fusion necessarily implies a telomere crisis. It could simply derive from a subtelomeric dsDNA break and fusion.

Response: We agree with the reviewer that the observed SV could be indeed due to a subtelomeric dsDNA break and fusion, and we have hence rephrased that sentence.

Reviewer #3–12: Could the differential methylation calling be confounded by purity and/or copy number differences between the two tumor samples?

Response: In order to make sure that purity/contamination would not affect DMR calling, we applied a rather stringent effect size threshold (> 0.5 methylation rate difference). This threshold limits the impact copy number variation could potentially have on DMR calling. Still, a region of allele specific methylation which is deleted in one sample in one allele (LOH) could potentially be mistaken as a between-sample DMR. To ensure that such effects do not drive our results, we compared DMRs between tumor samples with allele specific copy number changes between matching haplotypes. We found that among 1,638 DMRs between cancer samples, 449 regions have an allelic ratio difference between samples of greater than 0.25. Of those 449 regions, 317 regions show the same effect size direction in both haplotypes and are therefore not affected by allelic imbalance. For 90 regions one haplotype is deleted in one of the samples but the remaining haplotype sufficiently explains the discovered differential methylation. Only 31 regions cannot be sufficiently explained in this manner. Upon manual inspection of these 31 regions, we observed that 8 of them (0.4% of the total number of DMRs) appear to indeed be driven by copy number changes. Of these eight, five were near genic regions and only two were discovered as promoter-DMRs. We have flagged these two in the supplementary table S8 as false positives.

Reviewer #3–13: While generally carefully phrased, in parts of the manuscript where several specific genes are highlighted this can start to feel like cherry-picking. For instance, in the latter part of the differential methylation section, authors report finding $>26,000$ differentially methylated CpGs and note two specific genes, NRN1 and PTCH1. Further individual genes are highlighted in nearly all subsequent sections.

Response: We appreciate this comment, and we agree that a clearer distinction should be made between systematic analyses and examples of various types of effects discoverable using ONT seq. After careful consideration we thus removed BCAT1 from the mentioned examples containing differential methylation and expression, and we additionally removed the two differential expression examples (KCNA1 and DMBT1), which helped streamline our manuscript.

We further removed the detailed explanation of the gene fusion used as an example of SV-tracing for fusion validation from the main text and now only refer to the respective figure. We have also slightly altered the text to clarify that TBX1 represents an example of an effect type (alternate promoter methylation) not shown in the main figures. We have kept other specific gene examples, such as NCOR1, as well as genes on the chromothriptic CS11-17, the mentioning of which we believe is instructive to appreciate the likely biological effects of somatic alterations in this tumor sample.

Reviewer #3–14: When resolving expression effects using ONT data, the authors leverage two distinct absolute log₂-fold change thresholds, >5 and >2, the latter set to compare to the promoterlinked differentially methylated regions. It is unclear why two thresholds are being used and what the total number of differentially expressed genes is under the weaker threshold (without exploring Table S7).

Response: We have removed the stronger threshold in order to make this section clearer. We now refer only to the >2 log₂-fold change threshold.

Reviewer #3–15: In line with the above, what fraction of the differentially expressed genes is likely due to being located on a segment with a copy number change between the two tumors?

Response: We have now added two columns to supplementary table S7 depicting estimated copy numbers for the two tumor samples respectively. We checked how differential expression and copy number differences between the two samples relate and show that only 2 of the 41 genes with strong expression and promoter methylation differences also show large copy number changes; additionally, factoring out copy number differences even slightly improves correlation between differential promoter methylation and differential expression (from -0.30 Spearman-R to -0.33). We expanded the text accordingly to clarify these points in the manuscript.

Reviewer #3–16: Figure S19, while there is likely to be a negative correlation, the regression results in a poor fit and should be removed from the panel.

Response: We thank the reviewer, and have removed the regression line from the panel.

Reviewer #3–17: Are the methylation changes in the two promotor regions of TBX1 associated with any changes in transcript/isoform usage as well as total expression?

Response: We thank the reviewer for the question. We had a detailed look into the expression and splicing patterns of TBX1. We find that the expression of TBX1 in the relapse is right around 1 read per million, while the relapse has a read depth around 85 reads per million (log₂ fold change: 5.29). Given the low expression levels in the primary tumor, studying isoforms across the samples is difficult – our LeafCutter analysis does not show any split reads for the primary tumor. When examining the coverage plots (Figure R3.6) we do observe limited evidence for different splicing patterns versus the annotated isoforms in Ensembl, but given that it is hard to contrast these findings we feel this does not justify further highlighting.

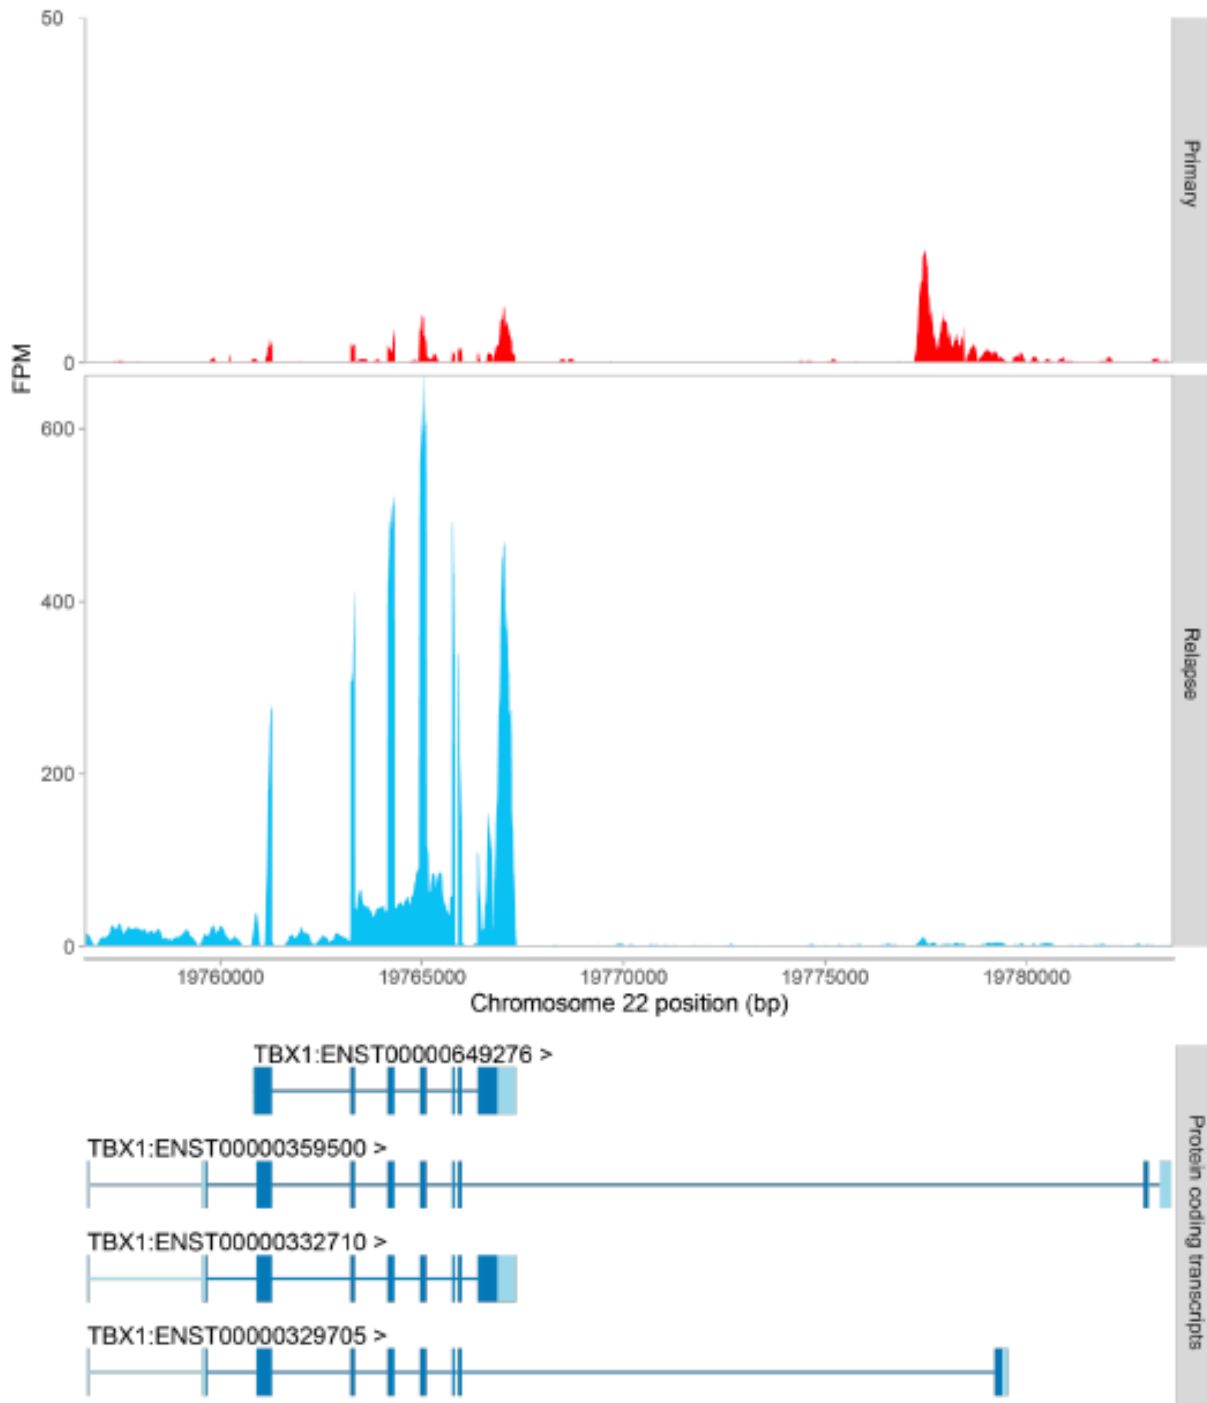


Figure R3.6: Expression of the TBX1 gene in the primary tumor (top) and relapse (middle) and protein-coding transcript derived from Ensembl. Shown are the expression levels, in features per million scale, relative to genomic location. The expression in both the primary and relapse samples indicate differential splicing from the annotated transcripts, but given the low expression in the primary tumor we can't compare the splicing ratios in a de novo fashion.

Reviewer #3–18: How were regions with loss of heterozygosity treated in the allele-specific methylation and expression analyses?

Response: Complete LOH of a gene (or gene promoter) would of course make one haplotype contain no methylation measurements and therefore fail to compute a methylation difference. As a result, no ASM would be detected. ASE, on the other hand, would be computed normally, assuming the locus is heterozygous in the Germline samples. When overlaying ASM and ASE effects we only include genes with significant ASM effects, so these would simply not show up in the results.

Reviewer #3–19: More general, what fraction of the allele-specific expression calls is due to allelic imbalance in copy number?

Response: We have rewritten this section in the manuscript to be more clear about the relationship between allelic copy number, ASE and ASM. While we see a significant relationship between ASE and allelic copy number, among 18 genes with strong ASM effects and significant ASE effects, only 4 (22%) appear to be driven by copy number imbalance. Even when broadening the definition of ASM from >0.5 abs methylation difference to >0.25 abs methylation difference, of the 51 genes with ASM and ASE, only 7 (13.7%) appear to be driven by copy number.

Reviewer #3–20: The following phrase is unclear to this reviewer: "... and ASM is increased in cancer, caused by disease associated regulatory SNPs"

Response: We appreciate the question. We refer to the phenomenon that disease associated (heterozygous) SNPs in cancer are associated with more frequent occurrence of allele specific methylation changes. We have rephrased the relevant section of the results.

Reviewer #3–21: Authors may wish to visualize the differential splicing pattern described in BASP1.

Response: We thank the reviewer for the suggestion and have added a supplementary figure, Figure S35, detailing the expression of BASP1 in the primary tumor and relapse. Figure S35 shows the three main differentially spliced junctions as well as the expression levels and reference annotation.

Reviewer #3–22: Authors highlight overexpression of CCND3 and BYSL in liposarcoma DO219967 in the PCAWG data set and suggest a nearby templated insertion thread may be driving this overexpression. Could the authors exclude gene copy number as a driver for overexpression here?

Response: The TI thread occurs in the vicinity of both genes but we did not intend to suggest it directly causes the overexpression. As outlined in the response to Reviewer #3–24, TI threads may be associated with genetic instability which could – in principle – include increases in the somatic copy-number of adjacent loci. We therefore examined the somatic gene copy-numbers of BYSL and CCND3, as determined by PCAWG (https://dcc.icgc.org/releases/PCAWG/consensus_cnv/gene_level_calls). In this case, there is indeed an increase in somatic copy-number with the TI thread being part of that amplified segment – as we now state in the revised manuscript.

Reviewer #3–23: MYPOP is mentioned to be subclonally amplified in the haplotype where the

telomere associated SV is observed. Could the authors provide support for this subclonal event?

Response: We thank the reviewer for pointing this out. While we cannot be certain that it is indeed a subclonal amplification, the distribution of reads (29 in major allele without SV, 9 in major allele with SV, and 18 in minor allele) cannot easily be explained by the tumor ploidy. We therefore assume subclonality. We have clarified the statement in the result section (“Functional annotation of the TI threads and telomere SVs”) of the manuscript.

Reviewer #3–24: Could the tandem short template jump pattern recently described by Umbreit et al. represent a precursor event to the longer chained rearrangements here? That is, could genetic instability at the locus lead to an increase in copy number and further chaining as observed in the templated insertion threads?

Response: This is an interesting hypothesis and with the additional liposarcoma samples we indeed see some evidence for multiple integrations, possibly caused by genetic instability. We first analyzed bulk sequencing data of the Primary Clone 1a displayed in Umbreit et al. Figure 5b (accession SRR10948021) to test whether rayas would detect TST jump patterns as templated insertion threads. Since rayas requires a control sample, and Umbreit et al. did not bulk-sequence the mother cell line, we obtained bulk sequencing data of the mother cell line from Tourdot et al. (accession SRR1778442). Both samples were realigned to GRCh38 before running rayas. As expected given the different appearance of both patterns, rayas with default parameters failed to detect any templated insertion threads in the given data. One of the validation samples in the Umbreit et al. study included a single long read from a primary renal cell carcinoma presented in Figure 5c, which contrary to the in-vitro data did show similarities to TI threads. We contacted the authors and they kindly shared the long read data with us. Interestingly, this long read indeed seems to provide direct evidence that a sequential set of templated insertions can occur almost unaltered twice within the same read, leading to an increased copy-number of all source segments (Figure R3.7).




Figure R3.7. Evidence of multiple copies for the templated insertions shown in Figure 5C of the Umbreit et al. study. A dot plot of the long read against itself reveals 2 similar copies of the same sequential order of templated insertions, one occurring at the beginning of the read and the other towards the end, suggesting that one instance is a reverse complement copy of the other (visible by the thin red line of matches in the bottom left and top right corner).

One of our own liposarcoma validation samples also showed clear evidence of TI threads (Figure R3.8, new Supplementary Figure S26) that are anchored in different parts of the genome. Although we lacked a single read spanning multiple instances, a pair-wise dot plot of the different integration events clearly shows how the sequential order of templated insertions is perfectly

conserved (Figure R3.9, new Supplementary Figure S27).

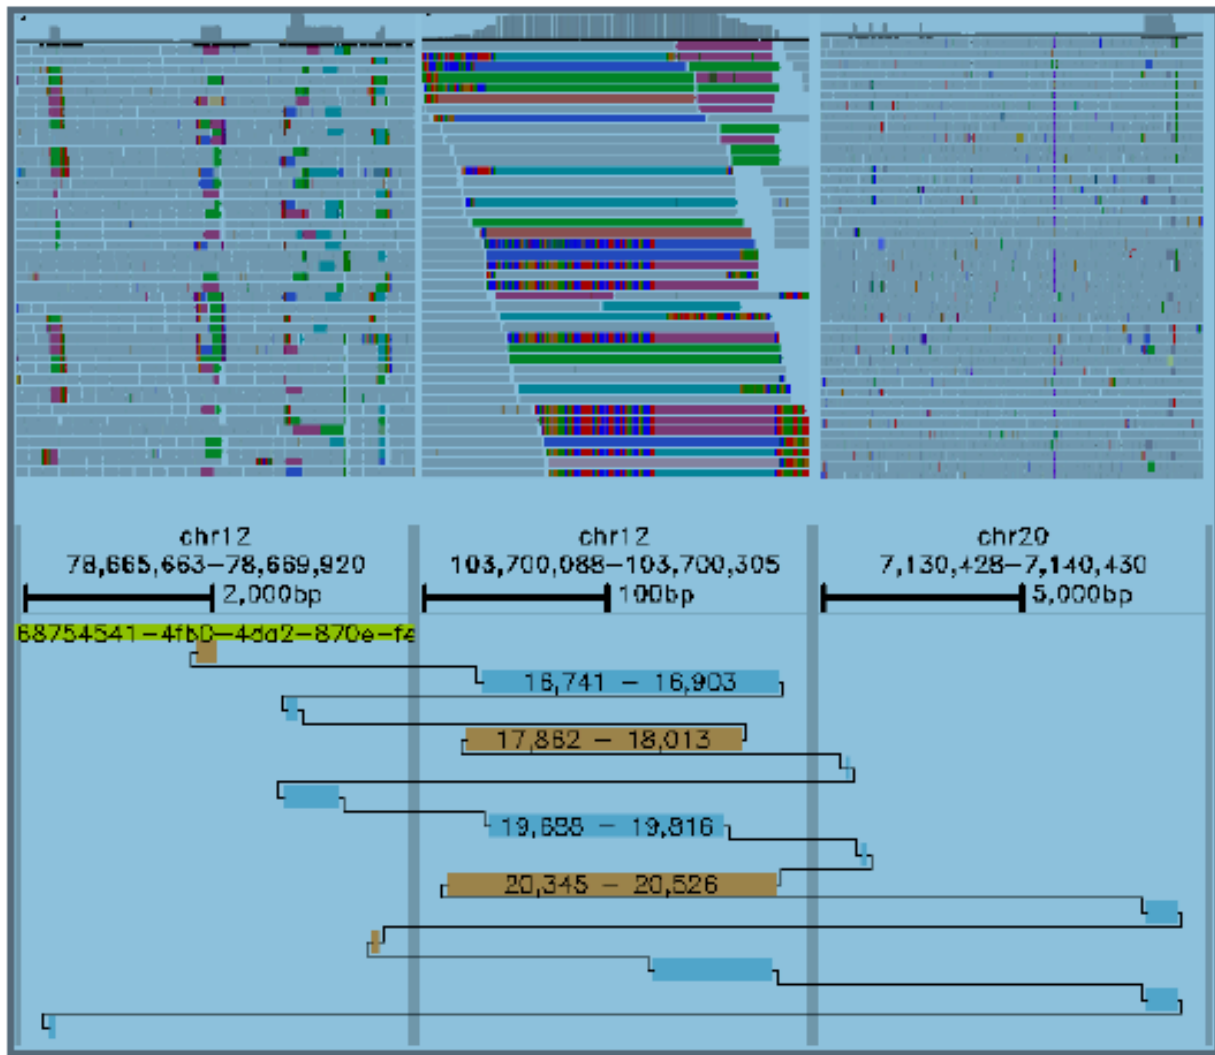


Figure R3.8. Raw ONT read supporting a TI thread (bottom panel) in liposarcoma sample P2. Alignment view is limited to source segments (separated by gray vertical lines) with at least 4 occurrences in the TI thread. Because of multiple integrations, aligned segments show strong coverage increases in the matched illumina short-read data (top panel) as well as SV-supporting split-reads with soft-clips.




Figure R3.9. Lorax predicts multiple integration sites for a given templated insertion thread. Pairwise dotplots of 2 raw ONT reads supporting the templated insertion thread shown in Figure R3.8 with matches to GRCh38 highlighted in the plot margins and colored by source chromosome. Arrows indicate forward or reverse matches. The dotplot shows forward matches in black, reverse matches in red.

In summary, it appears likely to us that templated insertions are associated with or potentially even result in genetic instability at the respective locus. TI threads may lead to further copynumber rearrangements, and potentially further chaining events that ultimately may contribute to the high copy-number of regions subject to TI threads. We added this hypothesis to the discussion.

Reviewer #3–25: For the DNA and RNA extractions, authors should further specify which "Qiagen kits" were used. Likewise, the "Illumina platform" used during RNA sequencing could be clarified.

Response: We have now added the information on the extraction kits and sequencer to the methods sections.

Reviewer #3–26: In the Methods section "Gene fusion and validation using DNA long reads.", the hyphen can be dropped from "set-out".

Response: We appreciate the correction and fixed it in the revised manuscript.

Referees' report, second round of review:

Reviewer #1:

The authors have done an exceptional job addressing my concerns and the amount of work they have added is impressive. In principle, this reviewer feels the manuscript is ready for publication, however there are a couple minor points the authors should address:

1. The fonts in the figures showing chained alignments (e.g. Figure 1D) are very poor quality. Is there anything that can be done to improve the readability of that text? While Figure 1 represents lots of important underlying work, the production quality of the figure does not match that - and it would stand to benefit from further unifying appearance of fonts and visual elements in the figure.
2. If it is possible, the authors may want to make their imaging data publicly available via FigShare or similar.
3. I agree with reviewers 2 & 3 that the selection of genes in the original manuscript may suffer from cherry-picking. That critique should be satisfactorily addressed.

Reviewer #2: The revised study of Rausch and colleagues includes new long read sequencing of a medulloblastoma PDX and two liposarcoma samples, and also has several additional analyses of existing and public data.

While TI threads indeed seem to be an interesting and novel pattern, much of the paper focuses on somewhat miscellaneous analyses (eg splicing, allele specific methylation), rather than developing this interesting finding. In addition, this finding still seems anecdotal and of unclear significance. Basic analyses to flesh out causes and consequences of the TI patterns are not done.

Rather the key thread linking the analyses in the paper seems to be to show the various uses of long reads in cancer. For many of these analyses, it still seems that similar conclusions could be reached with short reads. The technical / methodological advance continues to be unclear.

Overall the manuscript does not seem substantially changed and neither the revision or rebuttal substantially addresses my key points. In general, I think the study significantly overstates the utility of long reads. Data highlighted in the rebuttal seem to raise more concerns than they address.

Specific comments, including responses to rebuttal points:

- Many key results seem obtainable with short reads. Supporting this, the authors present a method rayas which they claim detects TI threads with specificity in short reads. While long reads are needed to fully reconstruct these rearranged TI thread contigs from end to end, it is unclear what biological insight that end-to-end reconstruction provides (see below)

- Fig. S1 is a nice example of overstating the utility of long reads. The haplotype blocks obtained from whatshap analysis of ONT reads are much smaller (panel A, 2-3 Mbp) than those that can be inferred through allelic imbalance and statistical phasing (panels B-C, chromosome scale). Both B-C can be easily obtained with short reads, since for a chromosome in allelic imbalance short read SNP VAF will give you fully

phased parental haplotypes. Applying this sort of logic to short read data would allow enable the sort of reconstructions shown in Fig 1. If a major purpose in the study is to show off the utility of long reads, a more sincere head-to-head comparison to short read approaches is warranted, eg through an analysis of samples with both short and long read sequencing.

- The authors gloss over a fundamental point - which is that long reads, particularly at the read lengths shown here, face many of the same challenges as short reads in the chromosome-scale reconstruction of complex SVs (Fig 1). While 10-100 kb long reads may be able to unambiguously reconstruct a TI thread, they are still too short to unambiguously reconstruct a multi megabase event like a chromothripsis, including the one shown in Fig 1.

It is not enough to use long read sequencing of the normal to infer megabase phase blocks and phase junctions to those blocks. Even with this parental phasing, there is still ambiguity in chromosome scale reconstruction. This is because cancer genomes often have multiple copies of each parental allele. Therefore it is not clear which of these distinct copies are connected to each junction and what is the order that these are arranged in a cancer chromosome, even with long reads phased to a parental haplotype.

The authors do not acknowledge this ambiguity (ie the classic karyotype reconstruction problem, cf Ron Shamir and others) and thus greatly overstate the utility of long reads.

- Figure 1 is a good example of this ambiguity, since all the segments in 1C are at copy 4 with three copies of the same parental allele; however, without directly observing an ONT read containing all of the involved junctions, it's impossible to know which parental allele copy is involved in each rearrangement. However the proposed reconstruction is over 1 megabase long and some of the segments (in window 1 and 5) are over 300 kilobases. There are no reads in their data that are this long, so how do they jump to the conclusion that this is the reconstruction? They must be applying some of the same heuristics that would be used in a short read analysis of the same genome. And indeed my guess is that a short read analysis of this genome would give the very same reconstruction of the CS11-17 locus in Figure 1.

- Related to comment R#2-3, it is still unclear how coherent the TI pattern is with respect to both contigs and reads. The dot plots (eg Fig R2.9) are not very helpful because they only show a pair of contigs / reads and are overwhelmed by the off diagonal signal which is mainly due to the self similarity in both sequences. To assess read support, it would be more helpful to just show the actual alignment of all reads (+/- contigs) to the best and longest contig i.e. in an IGV-style plot in contig coordinates. This includes the shorter ONT reads (related to comment R#2-4) or short reads (related to the comment below).

- Stepping back, it is fundamentally unclear what might be the functional relevance of the TI threads. Since they occur in very amplified regions, most of the expression changes of the genes around them can probably be explained by the increased copy number of the genes themselves. They don't appear to change splicing. Is there any interesting functional consequence that the PCAWG data point to? Do they land in enhancers? etc. I was hoping to see more of this sort of investigation in the revision.

- Related to above, how many TI threads occur outside of very high level amplification of the surrounding regions? Liposarcomas have very high copy amplifications, but the medulloblastoma shown here does not seem very amplified. Those TI threads that emerge in less amplified regions may be more interesting to probe for function.

- what sort of melanomas have TI threads? are these cutaneous / UV driven melanomas?

- Given the potential novelty of the TI thread pattern it is disappointing that there aren't more correlative analyses to better understand it. For example correlating it with specific drivers, where do these events land? are they associated with expression changes? Do they occur in tumors that have inactivate some DNA repair pathway? Again It would be more interesting and compelling I think to focus this paper on the TI threads, as I suggested in the initial review.

- The novelty or specificity of the "rayas" approach is still not clear. Fundamentally, since the TI thread junctions appear to be short read mappable (response to comment R#2-6) it is not clear why these junctions are missed by svaba or gridss.

If they are actually detected by these local assembly based callers, then is rayas just identifying clusters of such junctions? But then how does it differ from previous clustering approaches eg such as from Linx?

Otherwise is rayas actually finding new junctions that are missed by previous callers? If so, what explains that these rearrangements are missing from previous callers, if they are short read mappable?

- Related to above and Comment #2-19, the authors mention that LINX finds clusters at TI threads but does not reconstruct them correctly. It is unclear what the gold standard here is to say that they are incorrect? are these compared to the contig inferred from long reads? How far off are they from the truth? It is hard to get a sense of how much more sensitive or specific rayas is relative to these existing approaches. I suppose this could be addressed in simulations but ideally addressed in data.

- Related to Comment #2-19, the authors seem to conflate the TI thread detection, TI thread reconstruction in Fig 2, and the reconstruction of larger multi megabase events (Fig 1). These are separate analysis tasks. Related to the last task, it is unclear how much the phased long reads and the methods outlined in the study improve megabase scale reconstructions of cancer chromosomes, for examples in chromothripsis, related to the existing short read methods

- related to Comment #2-8, what I had in mind for an evaluation of rayas was a more comprehensive comparison of long vs short read data, particularly to understand specificity. The provided Visor simulation does not seem to fully evaluate specificity because it does not seem to simulate other sources of clustered rearrangements as well as aneuploidy.

- Related to comment R#2-5 if the junctional sequences are short read mappable, which they seem to be, then the junctions should be detectable with short reads. Which they seem to be (eg using their method rayas or even using svaba). In that case it is unclear whether rayas is detecting new junctions or just finding clusters of already mappable junctions?

- Related to above, if TI threads are actually detectable in short reads, then what is the real added value of long reads? In other words, what is the value of reconstructing these sequences from end to end. Perhaps there is some interesting insight to be drawn into the mutational process from the particular pattern of self concatenation that can be characterized only with long reads. However this is not pursued and there are perhaps too few examples shown here to pursue it.

- The presentation of Fig 2 still makes it unclear where this rather small (100 kb scale) structure fits into the larger many megabase chromothripsis event. Rebuttal Fig R2.4 makes it seem that there are many "integration sites", but presumably these fit into the larger chromothripsis event.

- related to comment #2-25: while the reviewer can appreciate the value of phasing methylation to parental alleles, most of the presented findings seem to be of unclear significance and ancillary to the SV portion of the paper or to the main observation which is TI threads. While the open sea hypomethylation in the chromothripsis allele is potentially interesting, the significance is unclear and an interpretation isn't provided to show the relevance.

- Related to comment R#2-9 it does not seem like the splicing data is relevant to the story. What is the reason to include it?

- Fig R1.7 what explains the multi-kb gaps in both PDX and primary ONT reads? ie kb regions in the read that are colored white meaning that there is no reference sequence (Similar issue with Fig R1.9).

- Fig R1.10 most of the de novo shasta contigs (y axis) do not bridge more than one reference contigs or in other words encompass no more than one rearrangement junction. It seems that the de novo shasta contigs are too short to support the long-range phasing implied by CS11-17.

- Related To #2-10, is there a positive control for the circle-seq analysis? The significance of a negative result is unclear without a positive control.

- The paper has "post-therapy" in its title but the conclusions still seem to offer little insight into the relapse. Related To #2-11 It is unusual and potentially interesting that the relapse is such a distant ancestor of the primary. Though without further evidence it seems that this is incidental to the presence of the TI thread, which at best seems dispensable for relapse.

- Incidentally, is it known what drove the relapse? Since the primary and relapse are so genetically distant, it seems the relapse might have very different drivers (in addition to TP53).

Reviewer #3: The authors have made extensive efforts to answer not only my own, but all reviewer comments. I have no remaining concerns.

Authors' response to the second round of review

Point by point response:

Reviewer #1:

Reviewer #1—general comments:

The authors have done an exceptional job addressing my concerns and the amount of work they have added is impressive. In principle, this reviewer feels the manuscript is ready for publication, however there are a couple minor points the authors should address.

Response: We thank the reviewer for their very positive assessment of our work.

Reviewer #1–1: The fonts in the figures showing chained alignments (e.g. Figure 1D) are very poor quality. Is there anything that can be done to improve the readability of that text? While Figure 1 represents lots of important underlying work, the production quality of the figure does not match that - and it would stand to benefit from further unifying appearance of fonts and visual elements in the figure.

Response: Thank you for the suggestion. We improved the fonts in the figure panels showing chained alignments. Please see updated Figures 1D, 2F and S26.

Reviewer #1–2: If it is possible, the authors may want to make their imaging data publicly available via FigShare or similar.

Response: Thanks for this suggestion, we uploaded the raw imaging data, related to figure 1E (metaphase spread), and figure S5 (interphase spread) to the BioImage Archive (<https://www.ebi.ac.uk/biostudies/bioimages/studies/S-BIAD611>) and Zenodo (<https://doi.org/10.5281/zenodo.7533768>). The repository links have been added to the manuscript section and data availability.

Reviewer #1–3: I agree with reviewers 2 & 3 that the selection of genes in the original manuscript may suffer from cherry-picking. That critique should be satisfactorily addressed.

Response: We are aware of the limitations of a study focusing mostly on one patient, and therefore a large part of our work was devoted to developing open source methods to advance the use of long reads in cancer genomics. Because of these development efforts, we were able to identify a novel TI thread pattern that could be confirmed in other specimens and tumour types, but for other reported rearrangement events and differentially methylated or expressed genes, we acknowledge that the biological significance is currently less clear. Following the Cell Genomics guidelines we now included a paragraph entitled "Limitations of the Study" to highlight these, and other, limitations of our study and made the context of the mentioned genes more clear.

Response: We are aware of the limitations of a study focusing mostly on one patient, and therefore a large part of our work was devoted to developing open source methods to advance the use of long reads in cancer genomics. Because of these development efforts, we were able to identify a novel TI thread pattern that could be confirmed in other specimens and tumour types, but for other reported rearrangement events and differentially methylated or expressed genes, we acknowledge that the biological significance is currently less clear. Following the Cell Genomics guidelines we now included a paragraph entitled "Limitations of the Study" to highlight these, and other, limitations of our study and made the context of the mentioned genes more clear.

Reviewer #2:

Reviewer #2—general comments

The revised study of Rausch and colleagues includes new long read sequencing of a medulloblastoma PDX and two liposarcoma samples, and also has several additional analyses of existing and public data.

While TI threads indeed seem to be an interesting and novel pattern, much of the paper focuses on somewhat miscellaneous analyses (eg splicing, allele specific methylation), rather than developing this interesting finding. In addition, this finding still seems anecdotal and of unclear significance. Basic analyses to flesh out causes and consequences of the TI patterns are not done.

Rather the key thread linking the analyses in the paper seems to be to show the various uses of long reads in cancer. For many of these analyses, it still seems that similar conclusions could be reached with short reads. The technical / methodological advance continues to be unclear.

Overall the manuscript does not seem substantially changed and neither the revision or rebuttal substantially addresses my key points. In general, I think the study significantly overstates the utility of long reads. Data highlighted in the rebuttal seem to raise more concerns than they address.

Response: To clarify, during manuscript revision we tripled the number of long-read cancer samples studied, rebasecalled all the original data, added new assay types (Circle-Seq) to clarify the cause of TI threads and performed a wealth of new computational analysis (de novo and targeted assemblies, xenograft analysis, and reconstruction of short-read amplicons, among many others). We investigated whether ecDNA is a possible cause of TI threads and also investigated co-occurrence with chromothripsis. We therefore respectfully disagree with the reviewer's opinion that TI threads represent merely an anecdotal observation; we have confirmed this pattern in independent specimens, tumor types and xenograft models, and have shown that we can infer their rearrangement signature indirectly from short reads followed by validation in long read data. We do acknowledge that there is work left needing to be done in the future outside of the scope of this study, which includes comprehensive characterization of common functional consequences of this interesting pattern of somatic DNA rearrangement.

Reviewer #2-1: Many key results seem obtainable with short reads. Supporting this, the authors present a method rayas which they claim detects TI threads with specificity in short reads. While long reads are needed to fully reconstruct these rearranged TI thread contigs from end to end, it is unclear what biological insight that end-to-end reconstruction provides (see below)

Response: Previous short read cancer genome sequencing efforts have failed to identify TI threads because of their structural complexity and the limitations of short-read amplicon reconstruction tools (as shown in our revised manuscript). We respectfully point out that in the past many other known cancer rearrangement processes such as chromothripsis were originally discovered in a new data type (short reads) but then studied and searched for in earlier technologies (microarrays). This was only possible due to a better understanding of a previously unknown genomic rearrangement pattern, and this principle is what we are using to screen short-read genomes on TI threads. Nevertheless, short reads cannot reconstruct TI thread patterns from end-to-end and lack resolution to distinguish between single and multiple integrations of potentially heterogeneous thread architectures which renders short-read analyses of TI threads incomplete. As already pointed out in the revised manuscript, the heterogeneous thread architectures shown in Figure S17 and S18 point towards a possible contribution of TI threads to genetic instability, a hallmark of cancer, exemplifying how end-to-end reconstruction of rearrangement events using long reads allows gaining biological insight into cancer.

We included a paragraph entitled "Limitations of the Study" that highlights the potential caveats of our TI thread analysis in the PCAWG cohort using short reads. We also suggest in that section additional long-read experiments to comprehensively characterize the templated insertion landscape in tumor genomes.

Reviewer #2–2: Fig. S1 is a nice example of overstating the utility of long reads. The haplotype blocks obtained from whatshap analysis of ONT reads are much smaller (panel A, 2-3 Mbp) than those that can be inferred through allelic imbalance and statistical phasing (panels B-C, chromosome scale). Both B-C can be easily obtained with short reads, since for a chromosome in allelic imbalance short read SNP VAF will give you fully phased parental haplotypes. Applying this sort of logic to short read data would allow enable the sort of reconstructions shown in Fig 1. If a major purpose in the study is to show off the utility of long reads, a more sincere head-to-head comparison to short read approaches is warranted, eg through an analysis of samples with both short and long read sequencing.

Response: For a short-read sequenced genome one can use allelic imbalance and statistical phasing, while for a long-read sequenced genome one can use allelic imbalance, statistical phasing and read-based phasing. Since read-based phasing does not require allelic imbalance and does not rely on a subset of SNPs being present in the haplotype reference panel one achieves (1) a greater density of phased markers (e.g. 19.5% of the phased markers of the Medulloblastoma are not present in the panel) and (2) a lower switch error rate compared to statistical phasing alone (Chaisson et al., 2019).

Reviewer #2–3: The authors gloss over a fundamental point - which is that long reads, particularly at the read lengths shown here, face many of the same challenges as short reads in the chromosome-scale reconstruction of complex SVs (Fig 1). While 10-100 kb long reads may be able to unambiguously reconstruct a TI thread, they are still too short to unambiguously reconstruct a multi megabase event like a chromothripsis, including the one shown in Fig 1. It is not enough to use long read sequencing of the normal to infer megabase phase blocks and phase junctions to those blocks. Even with this parental phasing, there is still ambiguity in chromosome scale reconstruction. This is because cancer genomes often have multiple copies of each parental allele. Therefore it is not clear which of these distinct copies are connected to each junction and what is the order that these are arranged in a cancer chromosome, even with long reads phased to a parental haplotype. The authors do not acknowledge this ambiguity (ie the classic karyotype reconstruction

problem, of Ron Shamir and others) and thus greatly overstate the utility of long reads.

Response: Contrary to statistical phasing, read-based phasing of long reads delivers boundaries of phased blocks and hence, long reads in the same phased block can be confidently split by haplotype, including in regions of the genome that do not show allelic imbalance (66.3% of the genome of the primary tumor). Deciphering a complete cancer genome with long reads will be a formidable challenge, but one that we believe can be overcome in the future, similar to recent achievements in telomere-to-telomere assembly of cell lines (Nurk et al., Science 2022). In this respect, we highlighted in the [“Limitation of the study”](#) section the need for developing production ready protocols for ultra-long read sequencing and further computational methods development for multi-allelic phasing and assembly – which we believe will help achieve the reconstruction of entire derivative chromosomes in cancer in the future.

Reviewer #2–4: Figure 1 is a good example of this ambiguity, since all the segments in 1C are at copy 4 with three copies of the same parental allele; however, without directly observing an ONT read containing all of the involved junctions, it's impossible to know which parental allele copy is involved in each rearrangement. However the proposed reconstruction is over 1 megabase long and some of the segments (in window 1 and 5) are over 300 kilobases. There are no reads in their data that are this long, so how do they jump to the conclusion that this is the reconstruction? They must be applying some of the same heuristics that would be used in a short read analysis of the same genome. And indeed my guess is that a short read analysis of this genome would give the very same reconstruction of the CS11-17 locus in Figure 1.

Response: We specifically looked at this question and our initially revised manuscript states that [“Linx also clustered the segments belonging to the CS11-17 structure into a single complex event, including segments overlapping contig 2. The chaining algorithm, however, failed to predict the entire CS11-17 structure and outputted 14 independent chains for this cluster.”](#) Therefore, short read analysis of this genome does not provide the same reconstruction of the CS11-17 locus.

Reviewer #2–5: Related to comment R#2-3, it is still unclear how coherent the TI pattern is with respect to both contigs and reads. The dot plots (eg Fig R2.9) are not very helpful because they only show a pair of contigs / reads and are overwhelmed by the off diagonal signal which is mainly due to the self similarity in both sequences. To assess read support, it would be more helpful to just show the actual alignment of all reads (+/- contigs) to the best and longest contig i.e. in an IGV-style plot in contig coordinates. This includes the shorter ONT reads (related to comment R#2-4) or short reads (related to the comment below).

Response: We thank the reviewer for this suggestion and point out that our revised manuscript contains such IGV-style plot, which is in Figure S19.

Reviewer #2–6: Stepping back, it is fundamentally unclear what might be the functional relevance of the TI threads. Since they occur in very amplified regions, most of the expression changes of the genes around them can probably be explained by the increased copy number of the genes themselves. They don't appear to change splicing. Is there any interesting functional consequence that the PCAWG data point to? Do they land in enhancers? etc. I was hoping to see more of this sort of investigation in the revision.

Response: The reviewer raises very interesting questions. It seems likely to us that TI threads

are associated with, or potentially even directly result in, genetic instability at the respective locus because of the observed heterogeneous architecture of TI threads and multiple integrations in the liposarcoma sample. We caution that resolving this possible association with genetic instability will require additional experiments in the future as described in the new **Limitations of the Study** section.

We also investigated the hypothesis that TI threads are formed via circular DNA intermediates using Circle-Seq data, which we found no evidence for. We found no significant enrichment of TI threads in enhancer regions, or at splice sites, but we note that other DNA rearrangement processes such as chromothripsis neither show such enrichment. We do agree that assessing the contribution of TI threads to the landscape of common cancer drivers will be a very important task for the future - which will be addressable once more genomes sequenced with long reads become available, allowing to systematically resolve the structure of TI threads at the nucleotide sequence level.

Reviewer #2-7: Related to above, how many TI threads occur outside of very high level amplification of the surrounding regions? Liposarcomas have very high copy amplifications, but the medulloblastoma shown here does not seem very amplified. Those TI threads that emerge in less amplified regions may be more interesting to probe for function.

Response: We apologize for not having been more explicit about this before. The liposarcoma sample P1 is indeed strongly amplified in the surrounding region with an estimated copy number (CN) of 41. The liposarcoma sample P2 shows multiple TI thread integration sites (see Figure S27) with widely varying copy-number (CN2 for chr4, CN28 for chr12 and CN31 for chr20). As in the medulloblastoma sample, the TI threads appear to be colocalized with rearrangement breakpoints where copy number changes occur (see new Figure S46, also shown below).

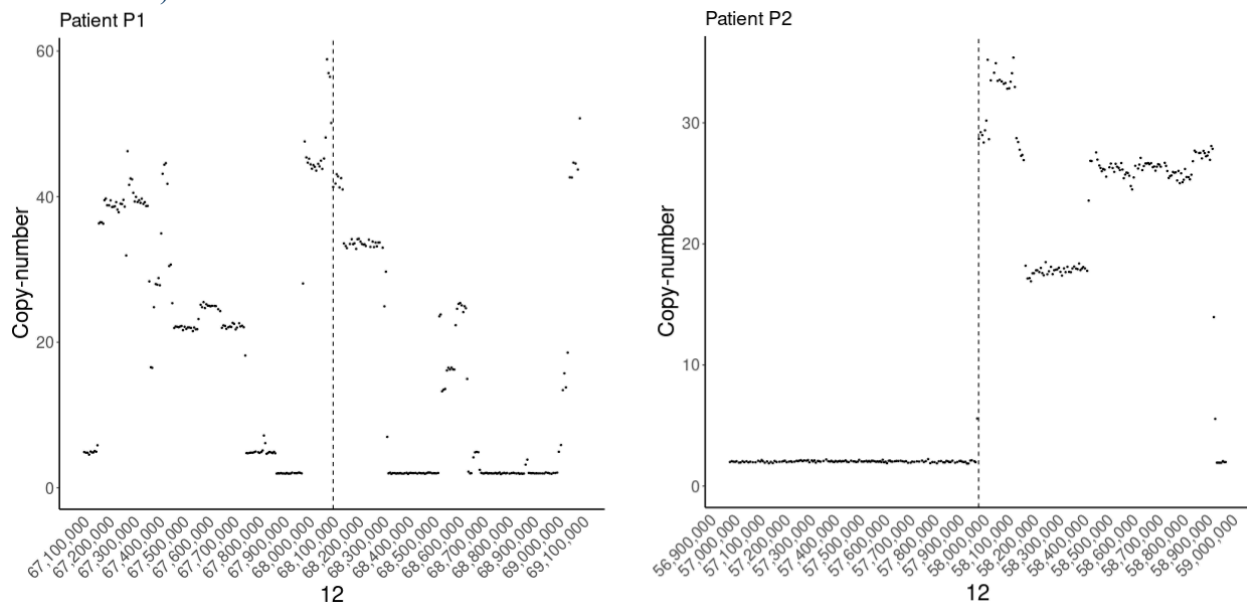
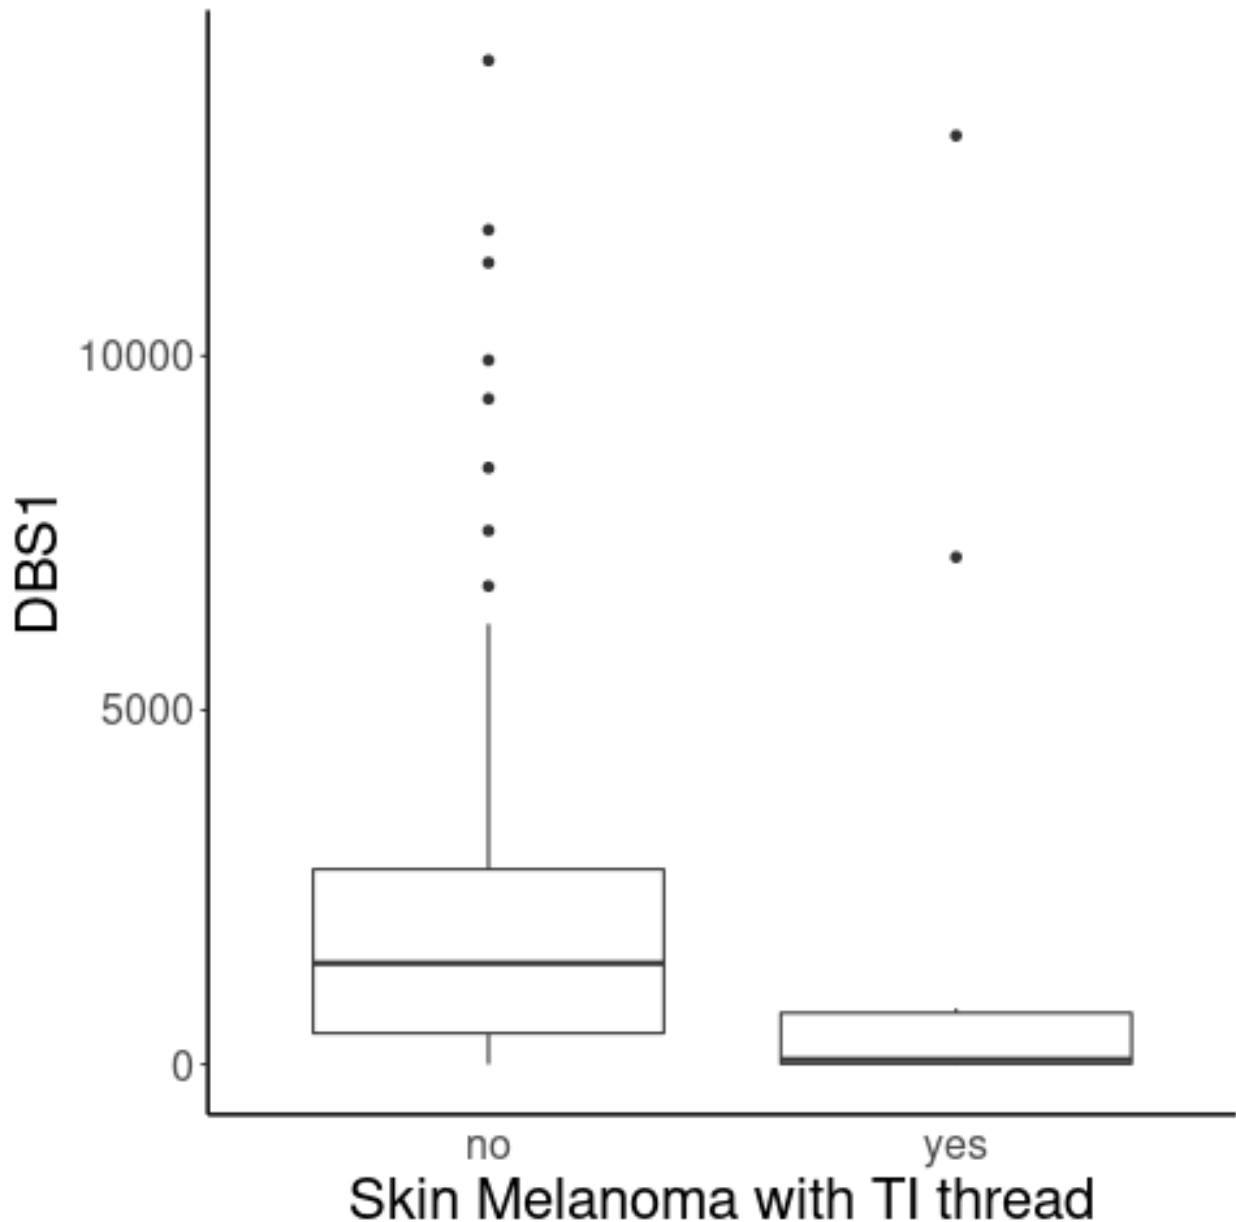


Figure S46. TI thread integration sites for patient P1 and P2. TI threads (dashed vertical line) colocalize with rearrangement breakpoints on chr12 where copy number changes occur. Coordinates are in GRCh37 and patient P1 is shown in the left panel and patient P2 in the right panel.

Reviewer #2-8: what sort of melanomas have TI threads? are these cutaneous / UV driven melanomas?

Response: Considering predicted exposure to doublet base substitution signature 1 (DSB1) from the PCAWG mutational signature analysis (Alexandrov et al., 2020) as a readout for UVdriven melanoma, the tumours predicted to harbour a TI thread do not appear to show strong UV mutational activity.



Reviewer #2-9: Given the potential novelty of the TI thread pattern it is disappointing that there aren't more correlative analyses to better understand it. For example correlating it with specific drivers, where do these events land? are they associated with expression changes? Do they occur in tumors that have inactivate some DNA repair pathway? Again It would be more interesting and compelling I think to focus this paper on the TI threads, as I suggested in the initial review.

Response: We thank the reviewer for this comment, and as already alluded to above we agree there is more work to be done on TI threads in the future. We note that the initial review report did not suggest to us that we should place more focus on relating inferred TI threads in

short read cancer genomes with drivers. Nonetheless, we emphasize as also stated in our response to #2-1, that short reads provide an incomplete view of the TI thread pattern, which we already alluded to in our prior point-by-point response (see previous comment #2-8): \approx We note that larger numbers of long-read cancer datasets will be required in the future to fully characterize the TI thread pattern that we describe in our study. \approx

Reviewer #2-10: The novelty or specificity of the "rayas" approach is still not clear. Fundamentally, since the TI thread junctions appear to be short read mappable (response to comment R#2-6) it is not clear why these junctions are missed by svaba or gridss.

Response: To clarify, individual junctions are not missed and we already included in the previous revised manuscript a Figure showing SvABA contigs aligned to these junctions (Figure S12).

Reviewer #2-11: If they are actually detected by these local assembly based callers, then is rayas just identifying clusters of such junctions? But then how does it differ from previous clustering approaches eg such as from Linx? Otherwise is rayas actually finding new junctions that are missed by previous callers? If so, what explains that these rearrangements are missing from previous callers, if they are short read mappable?

Response: We thank the reviewer for these questions. Similar to the point above, this question was addressed and answered in the revised manuscript, specifically: \approx However, due to Linx \approx heuristic approach with several rules and clustering routines, the algorithm still managed to cluster all SVs related to the TI thread together in a giant complex event with 97 chains and 827 SVs together with many additional SVs from the massive chromothripsis event involving chr4, chr5 and chr7 (among others). Chain 2 best overlapped the TI thread presented in Figure 2 but greatly underestimated the true number of junctions with only 43 compared to 231 estimated from the Shasta assembled TI thread contig (Figure S44) – suggesting clear limitations of Linx and short-reads in general with respect to the characterization of such complex rearrangement structures. \approx

Reviewer #2-12: Related to above and Comment #2-19, the authors mention that LINX finds clusters at TI threads but does not reconstruct them correctly. It is unclear what the gold standard here is to say that they are incorrect? are these compared to the contig inferred from long reads? How far off are they from the truth? It is hard to get a sense of how much more sensitive or specific rayas is relative to these existing approaches. I suppose this could be addressed in simulations but ideally addressed in data.

Response: In the initial manuscript, we presented a long DNA read spanning the entire TI thread; furthermore, following the reviewers \approx suggestion during the previous revision, (previous comment #2-3) we also included an assembled contig. The same TI thread was corroborated in an independent sample, the xenograft model established from the same tumor, and the long-read derived TI thread structure can therefore serve as a gold standard used for comparison with Linx.

Reviewer #2-13: Related to Comment #2-19, the authors seem to conflate the TI thread detection, TI thread reconstruction in Fig 2, and the reconstruction of larger multi megabase events (Fig 1). These are separate analysis tasks. Related to the last task, it is unclear how much the phased long reads and the methods outlined in the study improve megabase scale

reconstructions of cancer chromosomes, for examples in chromothripsis, related to the existing short read methods

Response: We emphasize that the CS11-17 reconstruction and the TI threads are separate sections in the manuscript, marked with distinct section titles and separate main Figures. We also evaluated separately the short-read amplicon reconstruction of these genomic rearrangement events to not conflate these analyses. With respect to the second question (on phasing and reconstruction), we refer to our response to Comment #2-2.

Reviewer #2-14: related to Comment #2-8, what I had in mind for an evaluation of rayas was a more comprehensive comparison of long vs short read data, particularly to understand specificity. The provided Visor simulation does not seem to fully evaluate specificity because it does not seem to simulate other sources of clustered rearrangements as well as aneuploidy.

Response: The original reviewer comment this refers to is: ³ It would seem essential to get a sense of specificity, meaning how often does a pattern suggested to be a templated thread in short reads actually give you the long read pattern.³ To address this request, we generated long read data for another 2 tumor samples and confirmed the TI thread pattern. In addition, we evaluated rayas using Visor simulations to evaluate sensitivity and specificity in a controlled setting.

Reviewer #2-15: Related to comment R#2-5 if the junctional sequences are short read mappable, which they seem to be, then the junctions should be detectable with short reads. Which they seem to be (eg using their method rayas or even using svaba). In that case it is unclear whether rayas is detecting new junctions or just finding clusters of already mappable junctions?

Response: To clarify, rayas finds clusters of short-read mappable junctions that self- and cross-connect highly amplified templated insertion source segments. As stated above however, the end-to-end reconstruction of the novel TI thread pattern requires long DNA reads.

Reviewer #2-16: Related to above, if TI threads are actually detectable in short reads, then what is the real added value of long reads? In other words, what is the value of reconstructing these sequences from end to end. Perhaps there is some interesting insight to be drawn into the mutational process from the particular pattern of self concatenation that can be characterised only with long reads. However this is not pursued and there are perhaps too few examples shown here to pursue it.

Response: As our revised manuscript clarifies, TI threads are incompletely characterized in short reads (see response to #2-1). The long reads were, as stated in the prior response, essential to discern TI threads as a novel complex rearrangement pattern. Also, only the end-to-end reconstruction through long reads revealed the heterogeneous thread architecture and multiple integration sites highlighted in the revised manuscript that may suggest a possible role of TI threads in fostering genetic instability.

Reviewer #2-17: The presentation of Fig 2 still makes it unclear where this rather small (100 kb scale) structure fits into the larger many megabase chromothripsis event. Rebuttal Fig R2.4 makes it seem that there are many "integration sites", but presumably these fit into the larger chromothripsis event.

Response: We thank the reviewer for this question. For clarification, Figure 2 refers to the primary medulloblastoma, whereas Fig. R2.2 - R2.4 refers to the liposarcoma validation samples (Fig. S25 - S27). Figure 2D therefore highlights the single integration site for the medulloblastoma sample.

Reviewer #2-18: related to comment #2-25: while the reviewer can appreciate the value of phasing methylation to parental alleles, most of the presented findings seem to be of unclear significance and ancillary to the SV portion of the paper or to the main observation which is TI threads. While the open sea hypomethylation in the chromothripsis allele is potentially interesting, the significance is unclear and an interpretation isn't provided to show the relevance.

Response: To clarify, the overarching goal of our manuscript was to provide an overview of different novel or improved analyses that can be used to interpret cancer genomes. The methylation analyses presented show the benefit of using ONT data for DNA methylation analyses and highlights important genes in the context of medulloblastoma. We also assessed the impact of complex genetic variation on DNA methylation but in our samples it is hard to solidify these links – larger studies operating on cancer genomes sequenced with long reads will be required to address this point in the future. Our manuscript reports on the possibilities offered by ONT to characterize the epigenetic landscape of cancers at an unprecedented scale, and provides future users the tools and means to do so. We added more context to the “Limitations of the study” section in response to this reviewer comment.

Reviewer #2-19: Related to comment R#2-9 it does not seem like the splicing data is relevant to the story. What is the reason to include it?

Response: Using the functional data we have on these samples, we tried to map the functional consequences of complex genetic effects. Specifically, the splicing effect that we describe in the *BASP1* gene is a potential functional consequence of the TI thread that is located in this exact region. Using these analyses we intended to address the request from this reviewer to provide a better understanding of what TI threads are and what kind of functional role they may have. As alluded to above, studies with large sample size will be required to address the question of functional relevance of TI threads as drivers, as described in our new “Limitations of the Study” section.

Reviewer #2-20: Fig R1.7 what explains the multi-kb gaps in both PDX and primary ONT reads? ie kb regions in the read that are colored white meaning that there is no reference sequence (Similar issue with Fig R1.9).

Response: Indeed, these sub-sequences are unmapped, suggesting non-templated DNA synthesis or very short concatenated templated sequences that cannot be mapped unambiguously in GRCh38 at the current error rate of ONT data. The revised manuscript included a statement about these unmapped sequences: “Alignment matches in the plot margins are interspersed with unmapped sequences, suggesting non-templated DNA synthesis or very short templated sequences that cannot be mapped unambiguously in GRCh38.”

Reviewer #2-21: Fig R1.10 most of the de novo shasta contigs (y axis) do not bridge more than one reference contigs or in other words encompass no more than one rearrangement junction. It seems that the de novo shasta contigs are too short to support the long-range

phasing implied by CS11-17.

Response: We fully agree, and we stress that exactly this observation motivated us to develop lorax which performs a targeted, haplotype-resolved assembly that takes advantage of the long-read phasing information to reconstruct the CS11-17 assembly.

Reviewer #2-22: Related To #2-10, is there a positive control for the circle-seq analysis? The significance of a negative result is unclear without a positive control.

Response: The FISH analyses suggest that the CS11-17 structure indeed involves circular DNA and this was confirmed by the Circle-Seq data (Figure S7). For TI threads, we do not have any additional positive control as these may not involve circular intermediates.

Reviewer #2-23: The paper has "post-therapy" in its title but the conclusions still seem to offer little insight into the relapse. Related To #2-11 It is unusual and potentially interesting that the relapse is such a distant ancestor of the primary. Though without further evidence it seems that this is incidental to the presence of the TI thread, which at best seems dispensible for relapse.

Response: As stated in our previous response to #2-11: " ... the complex rearrangements involving chromosome 7 of the primary tumor are not present in relapse, including the TI threads." Thus, the clone carrying complex rearrangements on chromosome 7 and TI threads seems to be more sensitive to treatment, possibly since it comprises faster growing (i.e. initially more aggressive) cancer cells. While we believe that our title is appropriate, we would be happy to discuss variations of the title currently in use with the editor.

Reviewer #2-24: Incidentally, is it known what drove the relapse? Since the primary and relapse are so genetically distant, it seems the relapse might have very different drivers (in addition to TP53).

Response: The TP53 mutation and loss of wild-type allele is preserved in relapse and thus, these early events appear to drive both the primary and secondary tumor. In addition, as stated in our revised manuscript, the relapse indeed acquired additional driver mutations, including a 2bp frameshift insertion in the tumor suppressor gene SUFU, which is a known medulloblastoma driver (Figure S22).

Reviewer #3

The authors have made extensive efforts to answer not only my own, but all reviewer comments. I have no remaining concerns.

Response: We thank the reviewer for the positive and highly supportive review