Scaffold assembly based on genome rearrangement analysis

Sergey Aganezov^{b,*}, Nadia Sitdykova^a, AGC Consortium, Max A. Alekseyev^b

^aAcademic University, St Petersburg, Russia

^bThe George Washington University, Washington, DC, USA

Abstract

Advances in DNA sequencing technology over the past decade have increased the volume of raw sequenced genomic data available for further assembly and analysis. While there exist many algorithms for assembly of sequenced genomic material, they often experience difficulties in constructing complete genomic sequences. Instead, they produce long genomic subsequences (*scaffolds*), which then become a subject to scaffold assembly aimed at reconstruction of their order along genome chromosomes. The balance between reliability and cost for scaffold assembly is not there just yet, which inspires one to seek for new approaches to address this problem.

We present a new method for scaffold assembly based on the analysis of gene orders and genome rearrangements in multiple related genomes (some or even all of which may be fragmented). Evaluation of the proposed method on artificially fragmented mammalian genomes demonstrates its high reliability. We also apply our method for incomplete anophelinae genomes, which expose high fragmentation, and further validate the assembly results with referenced-based scaffolding. While the two methods demonstrate consistent results, the proposed method is able to identify more assembly points than the reference-based scaffolding.

Keywords: scaffolding, genome rearrangements, breakpoint graph, MGRA, genome assembly

---

Background

Genome sequencing technology has evolved over time, increasing availability of sequenced genomic data. Modern sequencers are able to identify only short subsequences (*reads*) in the supplied genomic material, which then become an input to genome assembly algorithms aimed at reconstruction of the complete genome. Such reconstruction is possible (but not guaranteed) only if each genomic region is covered by sufficiently many reads. Lack of comprehensive coverage (particularly severe in single-cell sequencing [3, 13]) and presence of long similar subsequences (*repeats*) in genomes pose major obstacles for existing assembly algorithms. They therefore often are able to reliably reconstruct only long subsequences of the genome (inter-spaced with low-coverage regions and repeats), called *scaffolds*.

The challenge of reconstructing a complete genomic sequence from scaffolds is known as the *scaffolds assembly* problem. It is often addressed technologically by generating so-called long-jump libraries [15, 4] or by using a related complete genome as a reference. Unfortunately, the

---

*Corresponding author

Email address: aganezov@gwu.edu (Sergey Aganezov)

technological solution may be expensive and inaccurate [9], while the reference-based approach is obfuscated with structural variations across the genomes [5].

In the current study, we assume that the constructed scaffolds are accurate and long enough to allow identification of orthologous genes. The scaffolds then can be represented as ordered sequences of genes and we pose the scaffolds assembly problem as the reconstruction of the global gene order (along genome chromosomes) from the gene sub-orders defined by the scaffolds. We view such gene sub-orders as the result of both evolutionary events and technological fragmentation in the genome. Evolutionary events that change gene orders are *genome rearrangements*, most common of which are *reversals*, *fusions*, *fissions*, and *translocations*. Technological fragmentation can be modeled by artificial “fissions” that break genomic chromosomes into scaffolds. Scaffold assembly can therefore be reduced to the search for “fusions” that revert technological “fissions” and glue scaffolds back into chromosomes. This observation inspires us to employ the genome rearrangement analysis techniques for scaffolding purposes.

Rearrangement analysis of multiple genomes relies on the concept of the breakpoint graph. While traditionally the breakpoint graph is constructed for complete genomes, it can also be constructed for fragmented genomes, where we treat scaffolds as “chromosomes”. We will demonstrate that the breakpoint graph of multiple genomes possesses an important property that its connected components are robust with respect to genome fragmentation. In other words, connected components of the breakpoint graph mostly retain information about the complete genomes, even when the breakpoint graph is constructed on their scaffolds. We will show how to utilize connected components of the breakpoint graph for the scaffold assembly of fragmented genomes.

The paper is organized as follows. In the Methods section, we provide background information about breakpoint graphs and genome rearrangements, discuss connected components of breakpoint graphs with respect to genome fragmentation, and describe our scaffold assembly algorithm. In the Results and Discussion section, we evaluate our proposed algorithm on both simulated and real data. We summarize and discuss the paper results in the Conclusions section.

Methods

Genome and Breakpoint Graphs

We start with defining a graph representation for a single genome, which may consist of multiple chromosomes and/or scaffolds commonly referred to as *fragments*. We represent each fragment with n genes as an undirected graph on $2 \cdot n$ *regular* vertices representing gene extremities and several *irregular* vertices, labeled by ∞ , encoding fragment ends (telomeres, if a fragment is a chromosome). A gene a is represented by two regular vertices labeled as a^t and a^h denoting its *tail* and *head* extremities, respectively. Vertices corresponding to extremities of adjacent genes are connected by *regular* edges. Each vertex corresponding to a gene extremity at a fragment end is connected to an irregular vertex with an *irregular* edge. The *genome graph* of a genome is formed by the graph representing its fragments.

We remark that a single genome rearrangement affects the genome graph as follows: a pair of edges are removed and a new pair of edges on the same four vertices is created (Fig. 1). Genome rearrangements are therefore often modeled as DCJ [17] or 2-break [1] operations on graphs.

The *breakpoint graph* of k genomes composed of the same ℓ genes consists of $2 \cdot \ell$ regular vertices (representing gene extremities), a number of irregular vertices (representing fragment ends) and undirected edges of k colors (one color reserved for each of the genomes) encoding adjacencies between genes and/or fragment ends in the genomes. The breakpoint graph can be viewed as the

Figure 1: Fusion/fission operations between the genome graphs of two-chromosomal genome $P = [a, b] [-d, c]$ and unichromosomal genome $P' = [a, b, -c, d]$, where regular and irregular edges are represented as solid and dashed, respectively. Grey boxes enclose pairs of vertices representing genes.

superposition of k genome graphs of individual genomes (Fig. 2). All edges connecting a pair of vertices in the breakpoint graph form a *multiedge*, whose *multicolor* is the set of individual colors in the multiedge¹ (e.g., in Fig. 2 vertices a^h and b^t are connected by a multiedge of the red–black multicolor).

We remark that traditionally breakpoint graphs are constructed on synteny blocks whose end-points represent *breakpoints* (thus the name) in the genomes. In contrast, we construct the breakpoint graph directly on genes, whose extremities may or may not form breakpoints. Such graph therefore can contain *trivial multiedges* formed by parallel edges of all colors, which correspond to

¹Each regular vertex may be adjacent to at most one irregular vertex (e.g., in Fig. 2 the vertex a^t is connected to a single irregular vertex with red, black, and blue edges forming the red–black–blue multiedge (a^t, ∞)).

Figure 2: Breakpoint graph $\mathbb{B}G$ of genomes $A = [a, b, c, d]$ (black edges), $B = [a, b][c, d]$ (red edges), and $C = [a, -b, c][d]$ (blue edges). Regular edges in each genome are shown as solid, while irregular edges are shown as dashed.

gene adjacencies shared across all the genomes and would be hidden within synteny blocks. Clearly, each trivial multiedge in the breakpoint graph forms its own connected component, which we also call *trivial*.

Connected Components and Fragmentation

Under a *connected component* in the breakpoint graph we will understand any largest set of regular vertices such that any two of them are connected by a path consisting of regular edges of any colors. The connected components form a partition of the regular vertices. We will show that this partition is robust with respect to fragmentation of the genomes. Namely, we observe that the connected components of the breakpoint graph of multiple genomes are strongly connected and can be hardly broken by technological “fissions”. To support this observation, we applied a number of random fissions² to six complete mammalian genomes and analyzed how such fissions affected the connected components of the breakpoint graph.

Using Ensembl BioMart tool [10], we obtained the following six complete mammalian genomes and pairwise orthologous gene mappings between them: *Homo sapiens* (GRCh38), *Mus musculus* (GRCm38.p2), *Rattus norvegicus* (Rnor_5.0), *Canis familiaris* (CanFam3.1), *Macaca mulatta* (MMUL_1.0), and *Pan troglodytes* (CHIMP2.1.4). From the orthologous gene mappings, we constructed gene families and filtered some of them so that each genome was represented as sequences of the same 11816 genes, each appearing in a single copy.

²Unfortunately we do not have information about the actual mechanism of fragmentation of a genome into scaffolds.

Figure 3: Averaged number of connected components in the breakpoint graphs of multiple mammalian genomes fragmented by random “fissions”. Statistics for groups of 3, 4, 5, and 6 genomes is shown in green, brown, purple, and blue, respectively.

In order to determine how robustness of the connected components depends on the number of genomes, we analyzed different subsets of mammalian genomes of various sizes from 3 to all 6. For each subset of size ℓ , we considered all possible combinations of ℓ mammalian genomes and constructed their breakpoint graph. After the same number of random fissions³ was applied to every genome, we computed the averaged number of nontrivial connected components in the breakpoint graphs. The results in Fig. 3 demonstrate that as the number of genomes grows, random fissions are less likely to break connected components into smaller ones. In other words, most connected components in the breakpoint graph of fragmented genomes are likely not affected by random fissions and thus represent connected components also in the breakpoint graph of complete genomes. We will employ this robustness property of the connected components and use them to guide our scaffolding algorithm.

Scaffold Assembly Algorithm

Our scaffold assembly algorithm takes as an input a set of fragmented genomes, whose scaffolds are represented as sequences of genes. In the current study, we focus only on gene families that are present in each genome exactly once.

³A random fission operation uniformly selects a regular edge in the genome graph performs a fission on this edge (Fig. 1).

The first step of our algorithm is a construction of the breakpoint graph $\mathbb{BG}$ of the given genomes.

From the breakpoint graph perspective, scaffold assembly corresponds performing “fusions” in $\mathbb{BG}$, i.e., adding new regular edges (*assembly edges*) connecting vertices that represent scaffold ends. Since the connected components in the breakpoint graph are robust with respect to genome fragmentation, our algorithm adds assembly edges only within existing connected components and thus preserves them. So the second step of our algorithm identifies pairs of *matching vertices* that will then be connected by assembly edges. Namely, vertices $x, y : x \neq y$ form a matching pair in genome P if they satisfy each of the following conditions:

connectivity: x and y belong to the same connected component X of $\mathbb{BG}$;

extremity: there are multiedges (x, ∞) and (y, ∞) of multicolor $\{P\}$ in $\mathbb{BG}$;

uniqueness: there is no vertex $z \in X, z \neq x, z \neq y$ such that the multiedge (z, ∞) has multicolor $\{P\}$; in other words, there are exactly two irregular edges of multicolor $\{P\}$ in X : (x, ∞) and (y, ∞) .

reliability: there exists a multiedge (x, y) (of any multicolor) in $\mathbb{BG}$.

The **connectivity** condition preserves connected components in the breakpoint graph. The **extremity** condition ensures that each of x, y represents a scaffold end in the genome P but not in any other genome.⁴ The **uniqueness** condition ensures that there is a unique way to create an assembly edge for genome P inside the connected component X : if there is another multiedge (z, ∞) of multicolor $\{P\}$, it would be unclear which pair out of x, y, z to connect with an assembly edge. The **reliability** condition ensures that the new adjacency created by an assembly edge (x, y) in genome P is already present in some other genome(s).⁵

Once we obtained a list of matching vertex pairs for every genome P , we perform assembly of the corresponding fragment ends in P .

Integration with MGRA Framework

We remark that our algorithm can be integrated with the MGRA framework[2], which performs rearrangement analysis of multiple genomes, identifies reliable genome rearrangements and transforms their breakpoint graph into an *identity* breakpoint graph (of a single ancestral genome). The identity breakpoint graph consists of trivial multicycles, each forming its own connected component. In the process of this transformation MGRA can only break the connected components of the breakpoint graph into smaller ones, which can be viewed as a refinement of the original connected components. As a result, MGRA can make possible for two vertices to obtain the uniqueness property after a number of genome rearrangements (Fig. 5).

However, since the irregular edges in the breakpoint graph after a number of genome rearrangements may no longer correspond to fragment ends, the **extremity** condition does not anymore imply that x and y are fragment ends. Therefore, to integrate the scaffold assembly with MGRA, we modify the **extremity** condition to additionally test if x and y correspond to fragment ends in

⁴Under random fragmentation, it is more likely for two genomes to share a common chromosome end than a scaffold end, which is not a chromosome end. So, if x or y is a scaffold end in two or more genomes, it is more likely for this vertex to represent a chromosome end.

⁵This condition is optional. It shall be utilized when the given genomes are closely related; however if the genomes are rather diverse, this condition may result in only small number of conservative assembly edges.

```

input : breakpoint graph  $\mathbb{B}G$  on set  $G$  of  $k$  genomes
output: list of triples of  $(x, y, P)$  where  $x$  and  $y$  is a matching pair in genome  $P$ 

 $result \leftarrow$  empty list;
 $cgs \leftarrow$  connected components (sets of vertices) of  $\mathbb{B}G$  on regular edges of all colors;
foreach  $cc$  in  $cgs$  do
    foreach genome  $P$  in  $G$  do
         $ies \leftarrow$  all irregular multiedges of multicolor  $\{P\}$  connecting vertices from  $cc$ ;
        if  $|ies| = 2$  then
             $v1 \leftarrow$  regular vertex from first edge in  $ies$ ;
             $v2 \leftarrow$  regular vertex from second edge in  $ies$ ;
            if there is edge  $(v1, v2)$  in  $\mathbb{B}G$  then
                append  $(v1, v2, P)$  to  $result$ ;
            end
        end
    end
end
return  $result$ ;

```

Figure 4: Pseudocode for identification of matching pairs of vertices.

Figure 5: **Left panel:** A connected component of the breakpoint graph of blue, black, red and green genomes. In this component the green genome has four irregular edges that correspond to fragment ends: (b, ∞) , (c, ∞) , (d, ∞) , and (p, ∞) . Since its endpoints b, c, p do not satisfy the uniqueness property, while endpoint d does not satisfy the extremity property, our algorithm is not able to assemble any of corresponding fragments in the green genome.

Right panel: A possible refinement of the connected component with MGRA with a genome rearrangement in the black genome. It results in a split of the connected component into two, which makes vertices c and p obtain the uniqueness property with respect to the green genome.

the genome P . By similar reasons, the **reliability** condition for vertices x, y should be tested in the original breakpoint graph. So if $\mathbb{B}G$ denotes the original breakpoint graph, while $\mathbb{B}G'$ denotes this graph after a number of genome rearrangements performed by MGRA, then a matching pair (x, y) in $\mathbb{B}G'$ should satisfy the following conditions:

- connectivity'**: x and y belong to the same connected component X of $\mathbb{B}\mathbb{G}'$;
- extremity'**: there are multiedges (x, ∞) and (y, ∞) of multicolor $\{P\}$ in both $\mathbb{B}\mathbb{G}$ and $\mathbb{B}\mathbb{G}'$;
- uniqueness'**: for any vertex $z \in X$ such that $z \neq x$, $z \neq y$, and the multiedge (z, ∞) in $\mathbb{B}\mathbb{G}'$ has multicolor $\{P\}$, the multiedge (z, ∞) either is not present in $\mathbb{B}\mathbb{G}$ or has multicolor different from $\{P\}$;
- reliability'**: there exists a multiedge (x, y) (of any multicolor) in $\mathbb{B}\mathbb{G}$;

Integration with MGRA allows us to obtain more matching vertices (as compared to what we can recover from the original breakpoint graph). We also take into consideration all pairs of vertices that are endpoints of fusions reported by MGRA. If such vertices correspond to fragment ends, we interpret their fusion as assembly the corresponding fragments.

Results and Discussion

Artificially Fragmented Genomes

We start evaluation of the proposed scaffold assembly algorithm with running it on artificially fragmented mammalian genomes. We use the same set of six mammalian genomes and two different approaches for fragmenting them: random fragmentation and fragmentation based on repeats in the genomes. The random fragmentation allows us to overcome the lack of information about genome fragmentation mechanism. However, we may have better insight in the fragmentation model, if we assume that genome scaffolds were obtained from a conventional genome assembler having difficulties in reconstruction of the order of long DNA repeats. In this case, it becomes realistic to fragment the genomes based on locations of such repeats.

Random fragmentation.

To create instances of randomly fragmented genomes, we applied k random artificial “fissions” to each of the genomes. For each value of k , we created 10 different sets of fragmented genomes, executed our algorithm on each of the sets (both with and without MGRA integration), and reported the following normalized values (averaged over the 10 sets):

- *true positive*: $\frac{|A \cap B|}{|A \cup B|} \cdot 100\%$, the percentage of assembly edges that correspond to the breakpoints of artificial “fissions”.
- *false positive*: $\frac{|A \setminus B|}{|A \cup B|} \cdot 100\%$, the percentage of assembly edges that do not correspond to breakpoints of artificial “fissions” (such assembly edges may rather indicate the locations of actual fissions in mammalian evolution).

Here A is the set of pairs of fragment ends assembled by our algorithm, and B is the set of pairs of fragment ends resulted from the simulated fissions.

From the evaluation results in Fig. 6, we conclude that while our algorithm is not able to reconstruct complete genomes where fragmentation is high, it is able to reconstruct genomes almost completely, when the fragmentation is low. While the true positive rate of the assembly results decreases as fragmentation raises, the results still remain highly reliable with low false positive rate. This is not an unexpected property of our algorithm since it is based on the robustness of the connected components of the breakpoint graph. Integration with the MGRA framework further yields additional number of highly reliable fragment assemblies.

Figure 6: Accuracy of the proposed algorithm on artificially fragmented mammalian genomes both with (green bars) and without (blue bars) integration with MGRA. For each number of random “fissions” (with step 100) applied to every genome, blue and green bars give normalized true positives, while red bar gives normalized false positives.

Repeat-based fragmentation.

To create instances of repeat-based fragmented genomes, we removed all repeats longer than a fixed number of basepairs (from 1K to 4K bp with the step of 0.5K) and partitioned the genomes into fragments with no long repeats. We used the same set of six mammalian genomes, for which we obtained the repeats locations from RepeatMasker [14] database. We performed the following three experiments:

- (i) **de novo assembly of multiple genomes:** all six genomes are fragmented (Table 3).
- (ii) **assembly of multiple genomes with a single reference:** all genomes, but *dog* (the only representative of the carnivore clade) are fragmented (Table 4).
- (iii) **assembly of a single genome with multiple references:** only *dog* genome is fragmented (Table 5).

In each experiment we considered only fragments that contain genes that are present exactly once in each genome. We evaluated the proposed algorithm in the same way as in the random fragmentation experiment, both with and without MGRA integration (Fig. 7).

Figure 7: Accuracy of the proposed algorithm on artificially fragmented six mammalian genomes, that were broken at the positions of repeats of length at least L , with (green bars) and without (blue bars) integration with MGRA. For each value of L (with step 500 bp) blue and green bars give the *true positive* rate, while red bars give *false positive* rate for assembly results.

Experiments (i) and (ii) demonstrate that while in the presence of a reference genome our algorithm yields more true fragment adjacencies, it still performs relatively well in the case, when no reference is known. Experiment (iii) shows that our algorithm can be used as a highly reliable step for assembly of a single fragmented genome, when several complete reference genomes are known. Since DNA repeats are subject to genome rearrangements in the course of evolution, integration with MGRA yields additional true adjacencies.

Anophelinae Genomes

The second evaluation of our scaffold assembly algorithm was performed on highly fragmented genomes from anophelinae subfamily, followed by comparison of the results to a reference-based assembly approach. Namely, we considered six anophelinae genomes: *An. gambiae*, *An. arabiensis*,

Genomes	# scaffolds	
	before filtration	after filtration
<i>An. gambiae</i>	6	6
<i>An. arabiensis</i>	340	95
<i>An. quadriannulatus</i>	647	306
<i>An. merus</i>	1078	816
<i>An. dirus</i>	302	124
<i>A. albimanus</i>	57	39

Table 1: Statistics on the number of nonempty scaffolds in anophelinae genomes before and after gene family filtration.

Genomes	# assemblies	
	without MGRA	integrated with MGRA
<i>An. gambiae</i>	0	0
<i>An. arabiensis</i>	6	10
<i>An. quadriannulatus</i>	75	91
<i>An. merus</i>	466	550
<i>An. dirus</i>	30	45
<i>A. albimanus</i>	6	10

Table 2: Statistics on the number of reported scaffold assemblies, both with and without integration with MGRA.

An. quadriannulatus, *An. merus*, *An. dirus*, and *An. albimanus*, for which we constructed gene families using orthologous gene mapping from OrthoDB [16]. We then filtered out all gene families that are not present exactly once on each given genome, thus limiting ourselves to the case of uniform gene content across the genomes. After such filtration each genome was represented as sequences of the same 6,837 genes. We remark that filtration eliminated all genes from some scaffolds and thus we exclude such scaffolds from assembly. Table 1 gives the scaffold statistics for anophelinae genomes before and after filtration.

Order and orientation of gene families in each given genome was determined from the corresponding GFF3 annotation obtained from VectorBase [11], where a gene is represented by a sequence of coding exons of various length. We define the gene coordinate in a genomic fragment as the mean coordinate of all its coding exons start/end coordinates (i.e., $(\text{start} + \text{end})/2$ averaged over all exons). Table 2 reports the number of scaffold assemblies obtained by the proposed algorithm.

As mentioned above, we compared our assembly results to another anophelinae study (*comparison study*) led by Dr. Igor Sharakhov at Virginia Tech University. The comparison study performed analysis of *An. gambiae*, *An. arabiensis* genomes from the same source, where *An. gambiae* represents a complete genome, while *An. arabiensis* exposes rather high fragmentation. The genome data preparation was similar to ours. The relationships between these genes and their order on scaffold were visualized in genoPlotR [7] and further compared to the cytogenetic [8] and physical [6] maps identifying breakpoints of fixed reversals. The *An. gambiae* genome assembly was used as a reference for scaffolding in *An. arabiensis*.

Among 10 assemblies in *An. arabiensis* genome identified by our algorithm, the comparison study was able to identify and confirm 6. For example, our algorithm suggested assembly of scaffolds

Figure 8: genoPlotR visualization of gene order on scaffolds KB704374, KB704562, KB704518, and KB704685 for *A. arabiensis* genome. Courtesy of Dr. Igor Sharakhov.

KB704562 and KB704374 as well as of scaffolds KB704518 and KB704685 in the *A. arabiensis* genome, which was also identified by the comparison study with the gene reference-based (Fig. 8).

Conclusions

In current study, we proposed a scaffold assembly algorithm based on the genome rearrangement analysis, which can be used to assemble highly fragment genomes. The proposed algorithm relies on the properties of breakpoint graph of multiple genomes and can be further integrated with the MGRA framework. We evaluated the proposed algorithm by testing it on both real and simulated genomic data. In both cases, it significantly reduced fragmentation of the genomes and demonstrated high reliability.

While the proposed algorithm relies on unique gene content, we are currently expanding the algorithm with support of non-unique (inserted/deleted or duplicated) genes, which will potentially lead to even better quantity and quality of the scaffold assembly results. We implemented the proposed algorithm in a prototype software, which we plan to make user-friendly and publicly available in near future.

Acknowledgements

The paper was prepared in collaboration with the Anopheles Genomes Cluster Consortium [12]. We are particularly indebted to Dr. Robert Waterhouse at MIT for the orthologous gene mapping data for Anopheles genomes. We are further thankful to Dr. Igor Sharakhov and members of his laboratory Xiaofang Jiang and Brantley Hall at Virginia Tech University for cross-validation of our scaffolding results.

The first and fourth authors were supported by the National Science Foundation under Grant No. IIS-1462107.

Anopheles Genomes Cluster Consortium

Daniel E. Neafsey, George K. Christophides, Frank H. Collins, Scott J. Emrich, Michael C. Fontaine, William Gelbart, Matthew W. Hahn, Paul I. Howell, Fotis C. Kafatos, Daniel Lawson, Marc A. T. Muskavitch, Robert Waterhouse, and Nora J. Besansky.

References

- [1] Max A. Alekseyev and Pavel A. Pevzner. Multi-Break Rearrangements and Chromosomal Evolution. *Theoretical Computer Science*, 395(2-3):193–202, 2008.
- [2] Max A. Alekseyev and Pavel A. Pevzner. Breakpoint Graphs and Ancestral Genome Reconstructions. *Genome Research*, 19(5):943–957, 2009.
- [3] H. Chitsaz, J.L. Yee-Greenbaum, G. Tesler, M.J. Lombardo, C.L. Dupont, J.H. Badger, M. Novotny, D.B. Rusch, L.J. Fraser, N.A. Gormley, O. Schulz-Trieglaff, G.P. Smith, D.J. Evers, P.A. Pevzner, and R.S. Lasken. Efficient de novo assembly of single-cell bacterial genomes from short-read data sets. *Nat Biotechnol*, 29(10):915–921, 2011.
- [4] F S Collins and S M Weissman. Directional cloning of DNA fragments at a large distance from an initial probe: a circularization method. *Proceedings of the National Academy of Sciences*, 81(21):6812–6816, 1984.
- [5] Lars Feuk, Andrew R Carson, and Stephen W Scherer. Structural variation in the human genome. *Nature Reviews Genetics*, 7(2):85–97, 2006.
- [6] Phillip George, Maria V Sharakhova, and Igor V Sharakhov. High-resolution cytogenetic map for the african malaria vector anopheles gambiae. *Insect molecular biology*, 19(5):675–682, 2010.
- [7] Lionel Guy, Jens Roat Kultima, and Siv G. E. Andersson. genoPlotR: comparative gene and genome visualization in R. *Bioinformatics*, 26(18):2334–2335, 2010.
- [8] Robert A Holt, G Mani Subramanian, Aaron Halpern, Granger G Sutton, Rosane Charlab, Deborah R Nusskern, Patrick Wincker, Andrew G Clark, JoséM C Ribeiro, Ron Wides, et al. The genome sequence of the malaria mosquito anopheles gambiae. *Science*, 298(5591):129–149, 2002.
- [9] Martin Hunt, Chris Newbold, Matthew Berriman, and Thomas D Otto. A comprehensive evaluation of assembly scaffolding tools. *Genome biology*, 15(3):R42, 2014.
- [10] Arek Kasprzyk. BioMart: driving a paradigm change in biological data management. *Database*, 2011, 2011.
- [11] Karine Megy, Scott J Emrich, Daniel Lawson, David Campbell, Emmanuel Dialynas, Daniel ST Hughes, Gautier Koscielny, Christos Louis, Robert M MacCallum, Seth N Redmond, et al. VectorBase: improvements to a bioinformatics resource for invertebrate vector genomics. *Nucleic acids research*, 40(D1):D729–D734, 2012.
- [12] Daniel E Neafsey, George K Christophides, Frank H Collins, Scott J Emrich, Michael C Fontaine, William Gelbart, Matthew W Hahn, Paul I Howell, Fotis C Kafatos, Daniel Lawson, et al. The evolution of the anopheles 16 genomes project. *G3: Genes|Genomes|Genetics*, 3(7):1191–1194, 2013.
- [13] Sergey I. Nikolenko, Anton Korobeynikov, and Max A. Alekseyev. BayesHammer: Bayesian clustering for error correction in single-cell sequencing. *BMC Genomics*, 14:S7, 2013.

- [14] AFA Smit, R Hubley, and P Green. RepeatMasker Open-3.0. 1996-2010. <http://www.repeatmasker.org>.
- [15] Michael E Talkowski, Zehra Ordulu, Vamsee Pillalamarri, Carol B Benson, Ian Blumenthal, Susan Connolly, Carrie Hanscom, Naveed Hussain, Shahrin Pereira, Jonathan Picker, et al. Clinical diagnosis by whole-genome sequencing of a prenatal sample. *New England Journal of Medicine*, 367(23):2226–2232, 2012.
- [16] Robert M Waterhouse, Fredrik Tegenfeldt, Jia Li, Evgeny M Zdobnov, and Evgenia V Kriventseva. OrthoDB: a hierarchical catalog of animal, fungal and bacterial orthologs. *Nucleic acids research*, 41(D1):D358–D365, 2013.
- [17] Sophia Yancopoulos, Oliver Attie, and Richard Friedberg. Efficient sorting of genomic permutations by translocation, inversion and block interchange. *Bioinformatics*, 21(16):3340–3346, 2005.

<i>L</i>	Chimpanzee				Dog				Human				Macaca				Mouse				Rat			
	Orth	IDOrth	UOrth		Orth	IDOrth	UOrth		Orth	IDOrth	UOrth		Orth	IDOrth	UOrth		Orth	IDOrth	UOrth		Orth	IDOrth	UOrth	
1K	10631 (26.04)	6234 (19.6%)	1636 (7.7%)		9067 (33.2%)	4807 (24.2%)	1547 (9.9%)		20318 (36.8%)	5539 (16.4%)	1779 (6.9%)		12975 (33.1%)	5345 (22.8%)	1673 (9.7%)		12975 (38.3%)	4293 (21.1%)	1597 (10.1%)		10399 (34.9%)	5011 (27.3%)	1469 (10.1%)	
1.5K	8994 (39.01)	5577 (30.5%)	1939 (15.6%)		7150 (52.1%)	4066 (40.6%)	1663 (21.6%)		16835 (49.7%)	5256 (25.9%)	2147 (14.0%)		8586 (47.8%)	4844 (35.8%)	1943 (19.4%)		11088 (49.6%)	3994 (30.4%)	1868 (18.9%)		8743 (46.4%)	4246 (37.6%)	1672 (18.8%)	
2K	7496 (48.51)	4781 (38.8%)	2016 (23.8%)		5570 (65.9%)	3308 (53.8%)	1143 (34.5%)		13707 (60.6%)	4766 (34.9%)	2235 (21.9%)		6586 (60.8%)	4213 (48.5%)	1978 (30.5%)		9845 (56.4%)	3697 (36.1%)	2053 (23.2%)		7527 (54.1%)	3682 (44.2%)	1754 (26.4%)	
2.5K	6497 (55.9%)	4262 (45.5%)	2050 (30.5%)		4389 (75.8%)	2653 (64.2%)	1518 (46.9%)		11727 (67.5%)	4425 (41.2%)	2306 (28.2%)		3385 (71.9%)	3463 (59.2%)	1848 (41.7%)		9061 (60.3%)	3530 (39.6%)	2184 (29.9%)		6651 (59.7%)	3266 (49.1%)	1739 (32.7%)	
3K	5789 (64.9%)	3899 (54.6%)	2012 (38.9%)		3605 (81.9%)	2246 (70.9%)	1382 (55.2%)		10226 (72.3%)	4121 (45.9%)	2301 (33.1%)		4273 (79.3%)	2894 (67.5%)	1682 (50.6%)		8448 (63.5%)	3368 (42.2%)	2199 (33.1%)		5997 (63.9%)	2994 (52.4%)	1730 (37.3%)	
3.5K	5197 (69.3%)	3589 (58.6%)	1965 (43.4%)		3071 (85.6%)	1908 (74.7%)	1258 (61.8%)		9005 (75.9%)	3875 (50.3%)	2278 (37.4%)		3368 (84.9%)	2377 (74.4%)	1487 (58.6%)		7909 (66.2%)	3204 (44.3%)	2189 (35.8%)		5350 (67.9%)	2725 (55.7%)	1661 (41.5%)	
4K	4694 (73.6%)	3292 (63.1%)	1896 (48.1%)		2683 (88.4%)	1698 (78.1%)	1175 (66.8%)		8007 (79.2%)	3631 (54.1%)	2221 (41.2%)		2635 (89.3%)	1940 (80.1%)	1285 (65.8%)		7538 (68.3%)	3125 (45.9%)	2223 (38.1%)		4779 (71.7%)	2502 (58.9%)	1601 (45.6%)	

Table 3: Statistics on the number of fragments and coverage (in parenthesis) after removing all repeats of length at least L in the each of the six mammalian genomes. The column *Orth* accounts the fragments that contain at least one gene. Column *IDOrth* accounts for fragments that contain at least one non-duplicated gene. Similarly, column *UOrth* accounts for fragments that contain at least one gene present in every genome exactly once.

<i>L</i>	Chimpanzee				Dog				Human				Macaca				Mouse				Rat			
	Orth	IDOrth	UOrth		Orth	IDOrth	UOrth		Orth	IDOrth	UOrth		Orth	IDOrth	UOrth		Orth	IDOrth	UOrth		Orth	IDOrth	UOrth	
1K	10613 (26.1%)	6223 (19.6%)	1728 (8.1%)		40 (100%)	20381 (36.9%)	1871 (7.4%)		5528 (15.9%)	1055 (33.2%)	5334 (22.8%)		1766 (10.1%)	5334 (22.8%)	4283 (21.0%)		12975 (38.3%)	4283 (21.0%)	1676 (10.5%)		10399 (34.9%)	4988 (27.3%)	1538 (10.4%)	
1.5K	8994 (39.1%)	5572 (30.2%)	2015 (16.1%)		40 (100%)	16835 (49.8%)	2221 (14.5%)		5252 (25.4%)	8586 (47.8%)	4878 (35.8%)		2013 (20.1%)	4878 (35.8%)	3988 (30.4%)		11088 (49.6%)	3988 (30.4%)	1937 (18.7%)		8742 (46.4%)	4230 (37.5%)	1728 (21.5%)	
2K	7469 (48.5%)	4778 (38.8%)	2078 (24.3%)		40 (100%)	13797 (60.6%)	2312 (22.4%)		4767 (34.5%)	6865 (60.8%)	4209 (48.5%)		2034 (31.2%)	4209 (48.5%)	3691 (36.1%)		9845 (56.5%)	3691 (36.1%)	2113 (25.7%)		7527 (54.1%)	3671 (44.1%)	1805 (29.2%)	
2.5K	6497 (55.9%)	4262 (45.5%)	2092 (30.9%)		40 (100%)	11727 (67.6%)	2361 (28.6%)		4424 (41.2%)	5385 (71.9%)	3460 (59.2%)		1891 (42.2%)	3460 (59.2%)	3525 (39.5%)		9061 (60.5%)	3525 (39.5%)	2232 (30.4%)		6651 (59.7%)	3260 (49.1%)	1796 (35.4%)	
3K	5789 (69.9%)	3897 (54.6%)	2050 (39.3%)		40 (100%)	10226 (72.3%)	2340 (33.4%)		4120 (45.9%)	4723 (79.3%)	2893 (67.5%)		1712 (51.1%)	2893 (67.5%)	3363 (42.1%)		8448 (63.6%)	3363 (42.1%)	2236 (33.4%)		5997 (63.9%)	2990 (52.4%)	1762 (39.9%)	
3.5K	5197 (69.2%)	3588 (58.8%)	1996 (43.8%)		40 (100%)	9005 (75.9%)	2307 (37.8%)		3875 (50.3%)	3368 (84.9%)	2376 (75.4%)		1509 (59.1%)	2376 (75.4%)	3201 (44.3%)		7909 (66.2%)	3201 (44.3%)	2215 (36.1%)		5350 (67.9%)	2723 (55.7%)	1690 (44.2%)	
4K	4694 (73.6%)	3291 (63.1%)	1919 (48.3%)		40 (100%)	8007 (79.2%)	2249 (41.5%)		3630 (54%)	2635 (89.3%)	1940 (66.2%)		1298 (66.2%)	1940 (66.2%)	3122 (45.9%)		7538 (68.3%)	3122 (45.9%)	2247 (38.4%)		4779 (71.7%)	2501 (58.9%)	1625 (48.1%)	

Table 4: Statistics on the number of fragments and coverage (in parenthesis) after removing all repeats of length at least L in the each of the six mammalian genomes, except *dog*. The column *Orth* accounts the fragments that contain at least one gene. Column *IDOrth* accounts for fragments that contain at least one non-duplicated gene. Similarly, column *UOrth* accounts for fragments that contain at least one gene present in every genome exactly once.

L	Chimpanzee	Dog			Human	Macaca	Mouse	Rat
		Orth	IDOrth	UOrth				
1K	26 (100%)	9067 (33.3%)	4675 (23.7%)	3933 (21.4%)	24 (100%)	22 (100%)	21 (100%)	22 (100%)
1.5K	26 (100%)	7150 (52.1%)	3990 (40.1%)	3480 (37.4%)	24 (100%)	22 (100%)	21 (100%)	22 (100%)
2K	26 (100%)	5579 (65.9%)	3265 (53.4%)	2886 (50.4%)	24 (100%)	22 (100%)	21 (100%)	22 (100%)
2.5K	26 (100%)	4389 (75.8%)	2628 (63.8%)	2346 (60.9%)	24 (100%)	22 (100%)	21 (100%)	22 (100%)
3K	26 (100%)	3605 (81.9%)	2228 (70.7%)	2013 (68.1%)	24 (100%)	22 (100%)	21 (100%)	22 (100%)
3.5K	26 (100%)	3071 (85.6%)	1896 (74.5%)	1717 (72.3%)	24 (100%)	22 (100%)	21 (100%)	22 (100%)
4K	26 (100%)	2683 (88.4%)	1687 (77.9%)	1533 (75.8%)	24 (100%)	22 (100%)	21 (100%)	22 (100%)

Table 5: Statistics on the number of fragments and coverage (in parenthesis) after removing all repeats of length at least L in the *dog* genome. The column *Orth* accounts the fragments that contain at least one gene. Column *IDOrth* accounts for fragments that contain at least one non-duplicated gene. Similarly, column *UOrth* accounts for fragments that contain at least one gene present in every genome exactly once.