

Identifying Positional Homologs as Bidirectional Best Hits of Sequence and Gene Context Similarity

Melvin Zhang

Department of Computer Science
National University of Singapore
13 Computing Drive, Singapore 117417
Email: melvin@comp.nus.edu.sg

Hon Wai Leong

Department of Computer Science
National University of Singapore
13 Computing Drive, Singapore 117417
Email: leonghw@comp.nus.edu.sg

Abstract—Identifying corresponding genes (orthologs) in different species is an important step in genome-wide comparative analysis. In particular, one-to-one correspondences between genes in different species greatly simplify certain problems such as transfer of function annotation and genome rearrangement studies. Positional homologs are the direct descendants of a single ancestral gene in the most recent common ancestor and by definition form one-to-one correspondence. In this work, we present a simple yet effective method (BBH-LS) for the identification of positional homologs from the comparative analysis of two genomes. Our BBH-LS method integrates sequence similarity and gene context similarity in order to get more accurate ortholog assignments. Specifically, BBH-LS applies the bidirectional best hit heuristic to a combination of sequence similarity and gene context similarity scores. We applied our method to the human, mouse, and rat genomes and found that BBH-LS produced the best results when using both sequence and gene context information equally. Compared to the state-of-the-art algorithms, such as MSOAR2, BBH-LS is able to identify more positional homologs with fewer false positives.

I. INTRODUCTION

Genome-wide comparative analysis of different species is only possible if we can identify conserved elements across species boundaries [1]. For many studies, the elements under consideration is the set of protein coding genes. Therefore, the identification of corresponding genes (orthologs) between different species is an important step in any genome-wide comparative analysis. In particular, one-to-one correspondences between genes in different species simplify certain applications such as transfer of function annotation [2] and genome rearrangement studies [3].

Consider a set of extant genomes and their most recent common ancestor (MRCA). For each gene in the MRCA, there is at most one direct descendant of the gene in each of the extant genomes. The direct descendants of a gene in the MRCA form a set of *positional homologs* [4]. A single ancestral gene may have multiple descendants due to gene duplication, or no descendants because of gene loss. In the case of gene duplication, we distinguish between the gene that remains in the original location and the copy inserted into a new location. The gene that retains its ancestral location is the direct descendant.

Positional homologs represent a set of genes in one-to-one correspondence with each other where each member

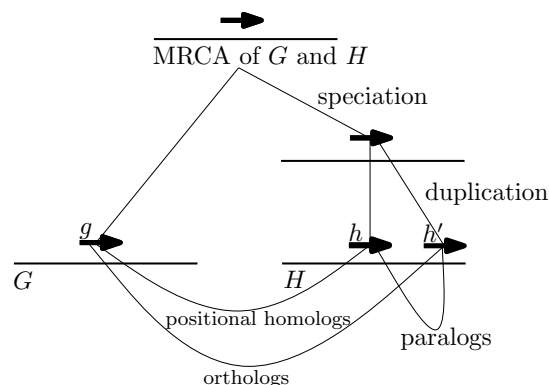


Fig. 1. The gene tree for three genes g , h , and h' that descended from a single ancestral gene in the most recent common ancestor (MRCA) of genome G and H . Gene g is orthologous to both h and h' , but only g and h are positional homologs because h is the original gene that was duplicated to get h' . Genes h and h' are paralogs as they are separated by a duplication event.

best reflect the original location of the ancestral gene in the MRCA. Similar concepts in the literature include exemplars [3], ancestral homologs [5], and main orthologs [6]. Orthologs are genes separated by a speciation event, while paralogs are genes separated by a duplication event. Orthologs and paralogs together make up the set of homologs [7]. Positional homologs are a subset of orthologs. Figure 1 shows the gene tree for three genes found in two genomes and it illustrates the concept of positional homologs, orthologs, and paralogs.

The problem of finding the set of positional homologs between two genomes is known as the ORTHOLOG ASSIGNMENT problem [6]. Current methods for the ORTHOLOG ASSIGNMENT problem fall into three categories: distance minimization, similarity maximization, and rule-based. *Distance minimization* methods relies on the parsimony principle. They assume that the removal of all the genes except for the positional homologs minimizes the genomic distance (usually some form of edit distance with genomic operations) between two genomes. Genomic distance measures such as the reversal distance [8] and breakpoint distance [9] have been considered using a branch-and-bound approach [3] as the corresponding computational problems are NP-hard [10]. MSOAR2 [11] uses a number of heuristic algorithms to assign positional homolog

pairs in several phases to minimize the number of reversals, translocations, fusions, fissions, and gene duplications between two genomes.

Closely related to distance minimization are the *similarity maximization* approaches. By identifying conserved structures between genomes, we can determine the similarity between them. We can model the ORTHOLOG ASSIGNMENT problem as finding the set of positional homologs that maximize the degree of similarity between two genomes. Bourque *et al.* [5] uses heuristics for the MAX-SAT problem to maximize the number of common or conserved intervals. The problem of maximizing the number conserved intervals is NP-hard [12]. Blin *et al.* [13] proposed a greedy method based on algorithms for global alignment that first finds a set of anchors and then recursively match genes found in large common intervals.

All of the preceding methods need a pre-processing step to compute gene families. This is typically accomplished using sequence similarity search followed by clustering of similar genes [14]. After that, sequence similarity is essentially reduced to a simple binary relation; two genes are the equivalent if they are in the same gene family and different otherwise. The main step uses heuristics to find a subset of genes that optimizes an NP-hard problem on gene orders. In short, the preceding methods use sequence similarity to build gene families and gene order information to further refine the gene families to get one-to-one gene matchings.

In contrast, *rule-based* methods do not need to build gene families. A widely used method for finding pairwise orthologs based on sequence similarity is the *bidirectional best hit (BBH)* heuristic. Two genes g and h form bidirectional best hits if the similarity between g and h is greater than that between g and any other gene (h is the best hit for g) and vice versa. In [4], a pair of BBHs are positional homologs if they are next to another pair of BBHs. Subsequently, [15] relaxed this condition and defined a local synteny test to determine whether a given pair of genes is a positional homolog pair. A gene pair passes the local synteny check if there are at least two pairs of genes (excluding the gene pair being tested) nearby with a sequence similarity above a certain threshold. Note that the local synteny test does not consider the sequence similarity between the gene pair being tested. Since positional homologs are a subset of all orthologs, other rule based methods designed for finding orthologs [16], [17] can also be used to identify positional homologs by restricting ourselves to one-to-one orthologous groups.

So far, existing methods have used sequence similarity information and gene order separately. In this paper, we propose to combine sequence similarity and gene context similarity into a single similarity score and identify positional homologs using the bidirectional best hit heuristic.

This has the advantage that the method is easy to implement and computationally efficient. Furthermore, we can easily vary the weightage of each type of similarity. We expected sequence similarity to play a larger role. Surprisingly, we get the best results using equal weightage for sequence and gene context similarity. Our method outperforms more complex methods,

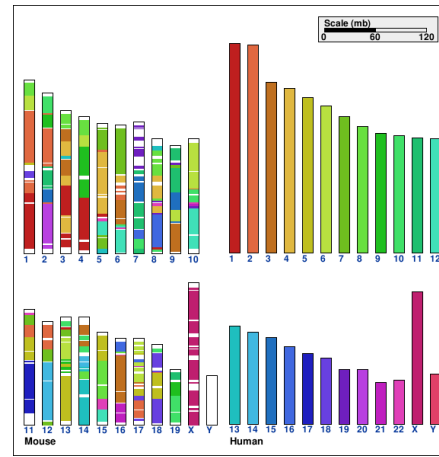


Fig. 2. Conserved synteny blocks between human and mouse genome generated by the Cinteny web server [19]

such as MSOAR2, in identifying positional homologs between human, mouse, and rat.

II. METHODS

Our approach is to approximate positional homologs as bidirectional best hits using a scoring scheme that integrates both sequence and gene context similarity scores.

Bidirectional (or reciprocal) best hits (BBH) is a widely used heuristic for finding orthologs between two species. Altenhoff *et al.* [18] compared a number of ortholog inference algorithms and found that BBH's overall performance is surprisingly good despite the simplicity of the method. In particular, they found that orthologs predicted by BBH show close functional relatedness. Another advantage of BBH is that it is easy to compute and commonly used in the literature.

However, using sequence similarity alone is *not enough* to identify the positional homolog among several orthologs [4]. In such cases, gene context can be used to disambiguate between the paralogs because positional homologs tend to have more similar gene context as evidenced by the presence of large synteny blocks (see Figure 2).

Furthermore, [20] showed that in 29–38 percent of the orthologs they investigated in bacteria, the gene pair with the lower sequence similarity have a higher gene context similarity. Hence, they advised combining gene context information with protein sequence information to predict functional orthologs. In this work, we integrate sequence similarity score with a gene context similarity score that reflects the shared gene neighborhood between two genes.

In the following subsections, we give the details for computing sequence and gene context similarity scores and explain how to combine them to compute bidirectional best hits.

A. Computing sequence similarity scores

We define the sequence similarity score between two genes as the Smith-Waterman alignment score between the respective peptide sequences. As a gene may have multiple transcripts,

we use the transcript with the longest peptide sequence to represent the gene. We use the SSEARCH program from the FASTA v36 package [21] to compute the Smith-Waterman alignment score between all pairs of peptide sequences using default parameters optimized for high sensitivity (BLOSUM50 substitution matrix and E-value cutoff of 10).

We use peptide sequences as the basis of sequence comparison as they have a number of advantages over using nucleotide sequences [22]. Peptide sequences are not affected by synonymous substitution and hence able to detect more distant homology. Furthermore, the alignments are faster to compute since the peptide sequence is only one third the length of the nucleotide sequence. Heuristic search algorithms, such as BLAST, are often used to find homologous sequences since they avoid computing the expensive dynamic programming alignment. However, a serious drawback is that the derived scores (bit score or E-value) are not symmetric and are difficult to easily integrate with other scores. On the other hand, the Smith-Waterman alignment score is symmetric and modern implementations are sufficiently fast for our purpose.

Since we want to integrate both sequence similarity and gene context similarity scores, we normalize the Smith-Waterman scores so that it ranges from 0 to 1 with a score of 1 indicating maximum sequence similarity. The Smith-Waterman alignment score is roughly linearly proportional to the length of the peptide sequences compared; longer peptide sequences tend to have higher alignment scores. Therefore, we remove this dependence on the length of the peptide sequence and normalize the score to range between 0 and 1 by dividing by the maximum Smith-Waterman score of the two self alignments. We formally defined the normalized Smith-Waterman score, sw_{norm} , as follows:

$$sw_{\text{norm}}(g, h) = \frac{sw(g, h)}{\max\{sw(g, g), sw(h, h)\}}$$

where $sw(g, h)$ is the Smith-Waterman alignment score between the peptide sequences of genes g and h .

B. Computing gene context similarity scores

Gene context similarity refers to the similarity in the genomic context of two genes. In contrast to sequence similarity, there is no widely accepted method to determine the level of gene context similarity between two genes. In this work, we make use of the concept of *local synteny* proposed in [15], [20].

Jun *et al.* [15] proposed a *local synteny test* that considers three genes upstream and downstream of two genes of interest to decide if they are orthologs. They modelled the sequence similarity between the two sets of six genes as a bipartite graph; there is an edge between two genes if their BLASTP E-value is less than $1e^{-5}$. They then compute a maximum matching of the graph. Two genes are *putative orthologs* if the size of the maximum matching is greater than one. In other words, they test if there is at least two other matching gene pairs in the vicinity of the gene pair of interest. They fixed the parameters (BLASTP threshold and size of gene neighborhood) in their




Fig. 3. Computing the local synteny score for g and h . We consider three genes upstream and downstream of the two genes of interest and add an edge between two genes if their BLASTP E-value is less than $1e^{-5}$. The thick edges show one of the possible maximum matching. The local synteny score of g and h is 4 since there are 4 edges in the maximum matching.

method by finding the values that maximizes the agreement with InParanoid [17] and Ensembl Compara [23] orthologs.) They showed that 93 percent of sampled inter-species pairs in five mammalian genomes (human, chimpanzee, mouse, rat, and dog) identified by their local synteny test are also found by InParanoid. By analyzing the remaining seven percent of the pairs, they conclude that the use of a local synteny test can resolve ambiguous many-to-many orthologous groups into one-to-one pairs.

While the binary test proposed in [15] detects the presence of other matching gene pairs in the vicinity of g and h , it does not capture the *strength* of this local synteny (or gene context similarity) nor does it make use of the sequence similarity of the gene pair being tested. Thus, it may cause errors in special cases: (a) *false positives* when the local context similarity is high, but the sequence similarity is low, or (b) *false negatives* when the local context similarity is low (only one other matching pair), while the sequence similarity is high.

In BBH-LS, we extend the binary test proposed by [15] by defining a *local synteny score*, $lss(g, h)$, that captures the *strength* of the local synteny between g and h . The *local synteny score* of two genes g and h is the size of the maximum matching between the six genes surrounding g and h (see Figure 3). This gives us a number between 0 and 6, which we normalize by dividing by 6. This is similar to the gene-neighborhood conservation score [20].

Formally, we define the normalized local synteny score, lss_{norm} , as follows:

$$lss_{\text{norm}}(g, h) = \frac{|\text{max matching of graph } G = (U \cup V, E)|}{6}$$

where U is the set of six genes around g , V is the set of six genes around h and there is an edge (u, v) in E if the BLASTP E-value of u and v is less than $1e^{-5}$.

C. Combining sequence and gene context similarity and computing bidirectional best hits

Given the normalized sequence similarity scores (sw_{norm}) and normalized gene context similarity scores (lss_{norm}), we combine them into a single similarity score (sim) with a parameter α to represent the weightage of gene context

similarity. Formally, we define the combined similarity score, sim , as follows:

$$\text{sim}(g, h) = (1 - \alpha) \times \text{sw}_{\text{norm}}(g, h) + \alpha \times \text{lss}_{\text{norm}}(g, h)$$

Using the combined score, we compute the set of bidirectional best hits by sorting all gene pairs in decreasing score and scanning this list once. A gene pair (g, h) is a bidirectional best hit if $\text{sim}(g, h)$ is strictly greater than $\text{sim}(g, h')$ for all other genes h' and $\text{sim}(g, h)$ is strictly greater than $\text{sim}(g', h)$ for all other genes g' . This guarantees that the set of bidirectional best hits is always one-to-one.

III. RESULTS AND DISCUSSION

We evaluate our BBH-LS method by applying it to the human, mouse, and rat genomes. For each pair of genome, we compared the performance of BBH-LS, BBH using only normalized Smith-Waterman score (BBH), MSOAR2 [11], InParanoid 4.0 [17], OMA [22], Ensembl Compara [23], and OrthoMCL [14].

Ideally, we should verify the inferred positional homologs by checking whether they perform the same function. Due to the lack of experimentally verified orthologs, we adopt the evaluation strategy used by MSOAR2. The idea is to make use of the official gene symbols for each gene as a proxy for its function. Gene symbols are manually curated based on gene function [24]. However, in the absence of experimentally verified function, genes may be manually assigned a gene symbol based on their sequence/structural similarity to other genes. This implies that there is some correlation between sequence similarity and gene symbols. Nevertheless, the additional manual curation put into the assignment of gene symbols makes them a reasonable measure.

Using this approach, we can classify the predicted positional homolog pairs into the following three categories:

- true positive: both genes share a common gene symbol
- false positive: gene symbols are completely different
- unknown: either one of the two genes have not been assigned a meaningful¹ gene symbol

The peptide sequences and locations of genes in each of three genomes were download from the Ensembl Release 60². There are 20801, 22842, and 22925 genes in the human, mouse and rat genome. We downloaded official gene symbols from the following species specific databases: HUGO Gene Nomenclature Committee³, Mouse Genome Informatics⁴, and Rat Genome Database⁵.

A. Parameter tuning for BBH-LS

Our scoring scheme uses the parameter α to controls the weightage of gene context similarity score. If α is 1, then we

¹We filter away symbols matching the regular expression “orf” in human genes, “Rik\$” or “GM[0-9]+\$” in mouse genes, and “LOC[0-9]+\$” or “RGD[0-9]+\$” in rat genes.

²retrieved from <ftp://ftp.ensembl.org/pub/release-60/fasta/> in Nov 2010

³retrieved from <http://www.genenames.org> in Dec 2010

⁴retrieved from <http://www.informatics.jax.org> in Dec 2010

⁵retrieved from <http://rgd.mcw.edu> in Dec 2010




Fig. 4. Performance of BBH-LS for different weightage of gene context similarity to sequence similarity on the *human-mouse* dataset. Left axis indicates the number of pairs of true positives and the right axis indicate the number of unknown pairs and false positives.




Fig. 5. Performance of BBH-LS for different weightage of gene context similarity to sequence similarity on the *mouse-rat* dataset. Left axis indicates the number of pairs of true positives and the right axis indicate the number of unknown pairs and false positives.

only use gene context similarity. If α is 0, then we only use sequence similarity.

We want to determine the optimal value of the parameter α on the human-mouse and mouse-rat dataset. To do this, we ran BBH-LS on the human-mouse and mouse-rat dataset over a range of values of α and tabulated the number of true positives, false positives and unknown pairs for each value. Figure 4 and Figure 5 shows how the number of true positives, false positives and unknown pairs varies as a function of α for each dataset.

For the human-mouse dataset (Figure 4) we observe that the number of true positives increases rapidly as α increases and then decreases at the same rate after reaching a maximum of 14133 when α is 0.53. However, the number of false positives

and unknown pairs also increased slightly as α increases. The same general trend for the true positives is observed for the mouse-rat dataset shown in Figure 5 (maximum of 12996 when α is 0.52).

We initially thought that the weightage of gene context similarity score should be much lower than that of the sequence similarity as many existing methods make use of sequence similarity but not gene context similarity. To our surprise, we found that setting α close to 0.50 maximizes the number of true positives for both datasets. In the following experiments, we set α as 0.50.

It is estimated that the last common ancestor of human and mouse existed 87 million years ago while the mouse-rat ancestor existed 16 million years ago [25], furthermore there are 297 large scale rearrangement events between human and mouse but only 106 rearrangement events between mouse and rat [26]. Despite the difference in the genomic distance in these two datasets, the best value of α is consistently around 0.50. Additional experiments using genomes of varying evolutionary distances will be necessary to determine whether this observation holds more generally.

B. Performance on Human-Mouse-Rat dataset

We obtained the output of the methods in our comparison by running the respective programs on the input data, except for OMA and Ensembl Compara as we did not have access to the programs. We downloaded the orthologs predicted by OMA⁶ and Ensembl Compara⁷ from their respective websites. InParanoid, OMA, and Ensembl Compara produces pairs of orthologous groups instead of positional homolog pairs. We get ortholog pairs by post-processing the output. InParanoid builds its groups from pairs of seed orthologs, we extract the seed orthologs from each group. For OMA, Ensembl Compara, and OrthoMCL, we use only the one-to-one groups.

Figure 6 shows the number of true positives (TP) and false positives (FP) for each method on three datasets. The results for OrthoMCL were not included in Figure 6 as is an outlier; for human-mouse dataset OrthoMCL has 8936 TP and 498 FP, for human-rat dataset there are 7409 TP and 530 FP, and for mouse-rat dataset there are 7812 TP and 819 FP.

For the human-mouse dataset, BBH-LS ($\alpha = 0.50$) identified the largest number of true positives (14125), followed by Ensembl Compara (13856), and MSOAR2 (13718). InParanoid which uses BLAST to compute sequence similarity does significantly worse than BBH using normalized Smith-Waterman alignment scores. In terms of the number of false positives, the methods we evaluated fall into three categories: low false positives (OrthoMCL, OMA, Ensembl Compara), medium false positives (InParanoid, BBH, BBH-LS), and high false positives (MSOAR2). OMA and Ensembl Compara performed surprisingly well given that we only consider the one-to-one groups that were generated.

The results for the human-rat dataset shown in Figure 6 is similar to that of the human-mouse data except that the

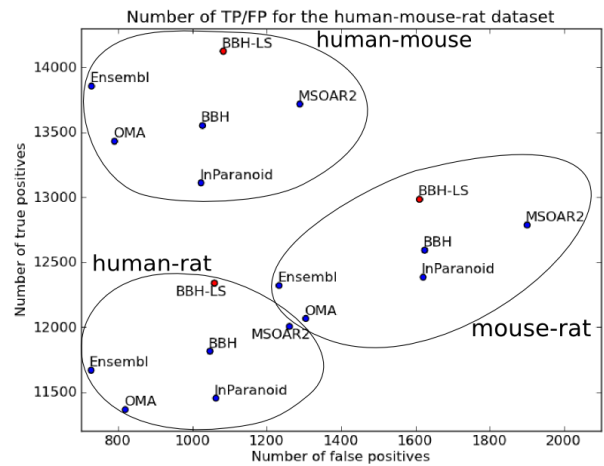


Fig. 6. Plot of number of true positives vs number of false positives in the output of BBH-LS, BBH, MSOAR2, InParanoid, OMA, and Ensembl Compara for the human-mouse, human-rat, and mouse-rat dataset




Fig. 7. Venn diagram showing the overlap between the true positives reported by BBH-LS, MSOAR2, and InParanoid for the human-mouse dataset.

number of true positives produced by Ensembl Compara and OMA has decrease relative to the other methods, but Ensembl Compara still has more true positives than InParanoid. For the mouse-rat dataset (Figure 6), OMA and Ensembl Compara is now worse than InParanoid. Another interesting characteristic of the mouse-rat dataset is the higher number of false positives, roughly doubled that of the human-mouse or human-rat dataset for all the methods.

Overall, in all three experiments, our BBH-LS method consistently produced the highest number of true positives as validated using gene symbols with a medium level of false positives.

Figure 7 shows a more detailed comparison of the true positives reported by BBH-LS, MSOAR2, and InParanoid. A total of 17081 true positive pairs are identified by at least one of the three methods and 57.5 percent (9322/17081) are identified by all three methods. Therefore, only slightly over half of the true positive pairs exhibit strong signals and are easy to detect. The rest require the combination of a number of different sources of information.

In the following, we illustrate a number of specific instances where gene context similarity made a significant differences. Figure 8 shows an instance where the high gene context

⁶retrieved from <http://omabrowser.org> in Dec 2010

⁷retrieved from <http://www.ensembl.org> in Dec 2010




Fig. 8. BBH erroneously paired RASGRF2 (human) to RASGRF1 (mouse) because of high Smith-Waterman score, this was corrected by BBH-LS with the help of local synteny score. Bold edges are the pairing from BBH-LS, thin edges are the pairing from BBH, sw = Smith-Waterman score, lc = local synteny score




Fig. 9. BBH-LS paired LILRA5 (human) with PIRA5 (mouse) and LAIR2 (human) with LIRA5 (mouse) due to the high local synteny produced by the five pairs of genes in between. The correct pairing should be LILRA5 (human) with LILRA5 (mouse) and this was picked up by BBH using just the normalized Smith-Waterman score.

similarity between related genes overcame the low sequence similarity between the true positives and converted what would be a false positive for BBH to a pair of true positives. There are a total of 12 of such cases. However, in four cases, the local synteny score caused a true positives identified by BBH to become a pair of false positives. Figure 9 illustrates one of these cases.

IV. CONCLUSION AND FUTURE WORK

The ORTHOLOG ASSIGNMENT problem is challenging in practice due to gene duplications and gene loss. Several sophisticated methods, which make use of complex heuristics (InParanoid) or require solving computationally hard problems (MSOAR2), have been proposed to tackle this problem. However, we show in this paper that the simple bidirectional best hit heuristic, coupled with a scoring scheme that combines both sequence and gene context similarity, is surprisingly good at identifying positional homologs. In all three pairwise comparison between human, mouse, and rat genomes, our BBH-LS method identified the most number of positional homolog (as validated using gene symbols) with a medium number of false positives.

We are currently working on reducing the number of false positives by considering the difference in the score of the best hit and second best hit.

ACKNOWLEDGEMENTS

We would like to thank the anonymous reviewers for many helpful suggestions and comments.

This work was supported in part by the National University of Singapore under ARF Grant R252-000-361-112.

REFERENCES

[1] E. Koonin, "Orthologs, Paralogs, and Evolutionary Genomics1," *Genetics*, vol. 39, no. 1, p. 309, 2005.
 [2] I. Friedberg, "Automated protein function prediction—the genomic challenge," *Briefings in bioinformatics*, vol. 7, no. 3, p. 225, 2006.
 [3] D. Sankoff, "Genome rearrangement with gene families," *Bioinformatics*, vol. 15, no. 11, pp. 909–17, Nov 1999.

[4] I. Burgetz, S. Shariff, A. Pang, and E. Tillier, "Positional homology in bacterial genomes," *Evolutionary Bioinformatics Online*, vol. 2, pp. 42–55, 2006.
 [5] G. Bourque, Y. Yacef, and N. El-Mabrouk, "Maximizing synteny blocks to identify ancestral homologs," in *Comparative Genomics*, ser. Lecture Notes in Computer Science, A. McLysaght and D. H. Huson, Eds., vol. 3678. Springer, 2005, pp. 21–34.
 [6] Z. Fu, X. Chen, V. Vacic, P. Nan, Y. Zhong, and T. Jiang, "Msoar: A high-throughput ortholog assignment system based on genome rearrangement," *Journal of Computational Biology*, vol. 14, no. 9, pp. 1160–1175, 2007.
 [7] W. Fitch, "Homology a personal view on some of the problems," *Trends in Genetics*, vol. 16, no. 5, pp. 227–231, 2000.
 [8] S. Hannenhalli and P. Pevzner, "Transforming cabbage into turnip (polynomial algorithm for sorting signed permutation by reversal)," *Journal of the ACM*, vol. 46, pp. 1–27, 1999.
 [9] G. A. Watterson, W. J. Ewens, T. E. Hall, and A. Morgan, "The chromosome inversion problem," *J. Theor. Biol.*, vol. 99, pp. 1–7, 1982.
 [10] D. Bryant, "The complexity of calculating exemplar distances," *Comparative Genomics: Empirical and Analytical Approaches to Gene Order Dynamics, Map Alignment, and the Evolution of Gene Families*, pp. 207–212, 2000.
 [11] G. Shi, L. Zhang, and T. Jiang, "MSOAR 2. 0: Incorporating tandem duplications into ortholog assignment based on genome rearrangement," *BMC bioinformatics*, vol. 11, no. 1, p. 10, 2010.
 [12] G. Blin and R. Rizzi, "Conserved interval distance computation between non-trivial genomes," in *COCOON*, ser. Lecture Notes in Computer Science, L. Wang, Ed., vol. 3595. Springer, 2005, pp. 22–31.
 [13] G. Blin, A. Chateau, C. Chauve, and Y. Gingras, "Inferring positional homologs with common intervals of sequences," in *Comparative Genomics*, ser. Lecture Notes in Computer Science, G. Bourque and N. El-Mabrouk, Eds., vol. 4205. Springer, 2006, pp. 24–38.
 [14] L. Li, C. Stoekert, and D. Roos, "OrthoMCL: identification of ortholog groups for eukaryotic genomes," *Genome research*, vol. 13, no. 9, p. 2178, 2003.
 [15] J. Jun, I. Mandoiu, and C. Nelson, "Identification of mammalian orthologs using local synteny," *BMC genomics*, vol. 10, no. 1, p. 630, 2009.
 [16] A. Schneider, C. Dessimoz, and G. Gonnet, "OMA Browserexploring orthologous relations across 352 complete genomes," *Bioinformatics*, vol. 23, no. 16, p. 2180, 2007.
 [17] G. Ostlund, T. Schmitt, K. Forslund, T. Kostler, D. Messina, S. Roopra, O. Frings, and E. Sonnhammer, "InParanoid 7: new algorithms and tools for eukaryotic orthology analysis," *Nucleic acids research*, vol. 38, no. Database issue, p. D196, 2010.
 [18] A. Altenhoff and C. Dessimoz, "Phylogenetic and functional assessment of orthologs inference projects and methods," *PLoS Comput Biol*, vol. 5, no. 1, p. e1000262, 2009.
 [19] A. Sinha and J. Meller, "Cinteny: flexible analysis and visualization of synteny and genome rearrangements in multiple organisms," *BMC bioinformatics*, vol. 8, no. 1, p. 82, 2007.
 [20] R. Notebaart, M. Huynen, B. Teusink, R. Siezen, and B. Snel, "Correlation between sequence conservation and the genomic context after gene duplication," *Nucleic acids research*, vol. 33, no. 19, p. 6164, 2005.
 [21] W. Pearson, "Searching protein sequence libraries: comparison of the sensitivity and selectivity of the Smith-Waterman and FASTA algorithms," *Genomics*, vol. 11, no. 3, pp. 635–650, 1991.
 [22] A. Roth, G. Gonnet, and C. Dessimoz, "Algorithm of OMA for large-scale orthology inference," *BMC bioinformatics*, vol. 9, no. 1, p. 518, 2008.
 [23] T. Hubbard, B. Aken, S. Ayling, B. Ballester, K. Beal, E. Bragin, S. Brent, Y. Chen, P. Clapham, L. Clarke *et al.*, "Ensembl 2009," *Nucleic acids research*, vol. 37, no. suppl 1, p. D690, 2009.
 [24] E. Bruford, M. Lush, M. Wright, T. Sneddon, S. Povey, and E. Birney, "The HGNC Database in 2008: a resource for the human genome," *Nucleic acids research*, vol. 36, no. suppl 1, p. D445, 2008.
 [25] M. Springer, W. Murphy, E. Eizirik, and S. O'Brien, "Placental mammal diversification and the cretaceous–tertiary boundary," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 100, no. 3, p. 1056, 2003.
 [26] G. Bourque, P. A. Pevzner, and G. Tesler, "Reconstructing the genomic architecture of ancestral mammals: Lessons from human, mouse, and rat genomes," *Genome Research*, vol. 14, no. 4, pp. 507–516, 2004.