

MSOAR: A High-Throughput Ortholog Assignment System Based on Genome Rearrangement

ZHENG FU,¹ XIN CHEN,² VLADIMIR VACIC,¹ PENG NAN,³ YANG ZHONG,³
and TAO JIANG^{1,3}

ABSTRACT

The assignment of orthologous genes between a pair of genomes is a fundamental and challenging problem in comparative genomics, since many computational methods for solving various biological problems critically rely on bona fide orthologs as input. While it is usually done using sequence similarity search, we recently proposed a new combinatorial approach that combines sequence similarity and genome rearrangement. This paper continues the development of the approach and unites genome rearrangement events and (post-speciation) duplication events in a single framework under the parsimony principle. In this framework, orthologous genes are assumed to correspond to each other in the most parsimonious evolutionary scenario involving both genome rearrangement and (post-speciation) gene duplication. Besides several original algorithmic contributions, the enhanced method allows for the detection of inparalogs. Following this approach, we have implemented a high-throughput system for ortholog assignment on a genome scale, called MSOAR, and applied it to human and mouse genomes. As the result will show, MSOAR is able to find 99 more true orthologs than the INPARANOID program did. In comparison to the iterated exemplar algorithm on simulated data, MSOAR performed favorably in terms of assignment accuracy. We also validated our predicted main ortholog pairs between human and mouse using public ortholog assignment datasets, synteny information, and gene function classification. These test results indicate that our approach is very promising for genome-wide ortholog assignment. Supplemental material and MSOAR program are available at <http://msoar.cs.ucr.edu>.

Key words: comparative genomics, gene duplication, genome rearrangement, inparalog, ortholog.

1. INTRODUCTION

ORTHOLOGS AND PARALOGS, originally defined by Fitch (1970), are two fundamentally different types of homologous genes. They differ in the way that they arose: orthologs are genes that evolved by speciation, while paralogs are genes that evolved by duplication. To better describe the evolutionary process and functional diversification of genes, paralogs are further divided into two subtypes: outparalogs, which

¹Department of Computer Science and Engineering, University of California, Riverside, California.

²School of Physical and Mathematical Science, Nanyang Technological University, Singapore.

³Shanghai Center for Bioinformatics Technology, Shanghai, China.

evolved via an ancient duplication preceding a given speciation event under consideration, and inparalogs, which evolved more recently, subsequent to the speciation event (Remm et al., 2001; Koonin, 2005). For a given set of inparalogs on a genome, there commonly exists a gene that is the direct descendant of the ancestral gene of the set, namely the one that best reflects the original position of the ancestral gene in the ancestral genome. Sankoff (1999) called such a gene the true exemplar of the inparalogous set. Given two genomes, two sets of inparalogous genes (one from each genome) are co-orthologous if they are descendants of the same ancestral gene at the time of speciation. These concepts are illustrated in Figure 1.

Clearly, orthologs are evolutionary and, typically, functional counterparts in different species. Therefore, many existing computational methods for solving various biological problems, e.g., the inference of functions of new genes and the analysis of phylogenetic relationship between different species, use orthologs in a critical way. A major complication with the use of orthologs in these methods, however, is that orthology is not necessarily a one-to-one relationship because a single gene in one phylogenetic lineage may correspond to a whole family of inparalogs in another lineage. More caution should be taken while such one-to-many and many-to-many relationships are applied to the transfer of functional assignments because inparalogs could have acquired new functions during the course of evolution. As a consequence, the identification of orthologs and inparalogs, especially those one-to-one orthology relationships, is critical for evolutionary and functional genomics, and thus a fundamental problem in computational biology.

It follows from the definition of orthologs and paralogs that the best way to identify orthologs is to measure the divergence time between homologous genes in two different genomes. As the divergence time could be estimated by comparing the DNA or protein sequences of genes, most of the existing algorithms for ortholog assignment, such as the well-known COG system (Tatusov et al., 1997, 2000) and INPARANOID (Remm et al., 2001), rely mainly on sequence similarity (usually measured via BLAST scores (Altschul et al., 1997)). An implicit, but often questionable, assumption behind these methods is that the evolutionary rates of all genes in a homologous family are equal. Incorrect ortholog assignments might be obtained if the real rates of evolution vary significantly between paralogs. On the other hand, we observe that molecular evolution proceeds in two different forms: local mutation and global rearrangement. Local mutations include base substitution, insertion and deletion, and global rearrangements include reversal (i.e., inversion), translocation, fusion, fission, and so on. Apparently, the sequence similarity-based methods for ortholog assignment make use of local mutations only and neglect genome rearrangement events that might contain a lot of valuable information.

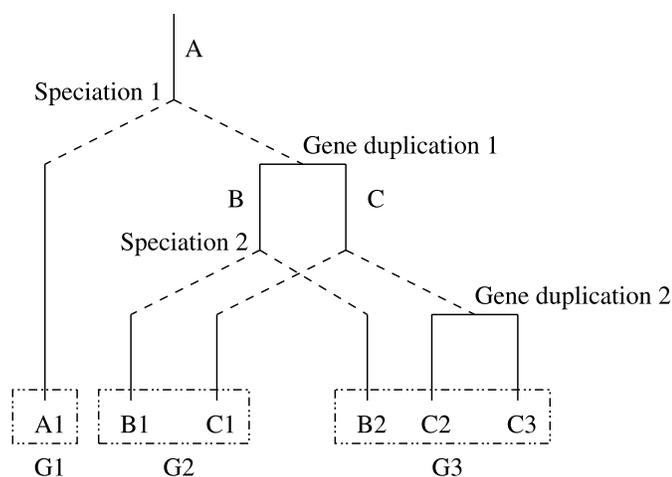


FIG. 1. An illustration of orthologous and paralogous relationships. After two speciation events and two gene duplications, three present genomes, $G_1 = (A_1)$, $G_2 = (B_1, C_1)$ and $G_3 = (B_2, C_2, C_3)$ are obtained. In this scenario, all genes in G_2 and G_3 are co-orthologous to gene A_1 . Genes B_1 and C_1 are outparalogs w.r.t. G_3 (i.e., the 2nd speciation), and are inparalogs w.r.t. G_1 (i.e., the 1st speciation). Gene C_2 is the direct descendant (i.e., true exemplar) of the ancestral gene C while C_3 is not, if C_3 is duplicated from C_2 . C_1 and C_2 are said to form a pair of main orthologs.

In our recent papers (Chen et al., 2005a, 2005b), we initiated the study of ortholog assignment via genome rearrangement and proposed an approach that takes advantage of evolutionary evidence from both local mutations and global rearrangements. It begins by identifying homologous gene families on each genome and the correspondence between families on both genomes using homology search. The homologs are then treated as copies of the same genes, and ortholog assignment is formulated as an optimization problem of rearranging one genome consisting of a sequence of (possibly duplicated) genes into the other with the minimum number of reversals, where the most parsimonious rearrangement process should suggest orthologous gene pairs in a straightforward way. A high-throughput system, called SOAR (System for Orthologous Assignment by Reversals), was implemented based on this approach. Though our preliminary experiments on simulated data and real data (the X chromosomes of human and mouse) have demonstrated that SOAR is very promising as an ortholog assignment method, it has the drawback of ignoring the issue of inparalogs. In fact, it assumed that there were no gene duplications after the speciation event considered. As a consequence, it only outputs one-to-one orthology relationships and every gene is forced to form an orthologous pair. Moreover, it is only able to deal with unichromosomal genomes. In this paper, we present several improvements that are crucial for more accurate ortholog assignment. In particular, the method will be extended to deal with inparalogs explicitly by incorporating a more realistic evolutionary model that allows duplication events after the speciation event. In summary, our main contributions in this study include

- We introduce a subtype of orthologs, called main orthologs, to delineate sets of co-orthologous genes. For two inparalogous sets of co-orthologous genes, the main ortholog pair is simply defined as the two true exemplar genes of each set (Fig. 1).¹ Since a true exemplar is a gene that best reflects the original position of the ancestral gene in the ancestral genome, main orthologs are therefore the positional counterpart of orthologs in different species. By definition, main orthologs form a one-to-one correspondence, thus allowing for the possibility of a direct transfer of functional assignments. We believe that, compared to other types of ortholog pairs, main orthologs are more likely to be functional counterparts in different species, since they are both evolutionary and positional counterparts.
- In our previous study, the evolutionary model assumes that there is no gene duplication subsequent to the given speciation event. Thus, no inparalogs are assumed to be present in the compared genomes, which is clearly inappropriate for nuclear genomes. In this paper, we propose a parsimony approach based on a more realistic evolutionary model that combines both rearrangement events (including reversal, translocation, gene fusion and gene fission) and gene duplication events. This will allow us to treat inparalogs explicitly. More specifically, in order to assign orthologs, we reconstruct an evolutionary scenario since the splitting of the two input genomes, by minimizing the (total) number of reversals, translocations, fusions, fissions and duplication events necessary to transform one genome into the other (i.e., by computing the rearrangement/duplication distance between two genomes). Such a most parsimonious evolutionary scenario should reveal main ortholog pairs and inparalogs in a straightforward way.
- Computing the rearrangement/duplication distance between two genomes is known to be very hard. We have developed an efficient heuristic algorithm that works well on large multichromosomal genomes like human and mouse. We strengthen and extend some of the algorithmic techniques developed in Chen et al. (2005a, 2005b), including (sub)optimal assignment rules, minimum common partition, and maximum graph decomposition, as well as a new post-processing step that removes “noise” gene pairs that are most likely to consist of inparalogs.
- Based on the heuristic algorithm mentioned above, we have implemented a high-throughput system for automatic assignment of (main) orthologs and the detection of inparalogs on a genome scale, called MSOAR. By testing it on simulated data and human and mouse genomes, the MSOAR system is shown to be quite effective for ortholog assignment. Namely, it is able to find 99 more true ortholog pairs between human and mouse genomes than INPARANOID (Remm et al., 2001).

¹Note that our definition of a main ortholog pair is different from the one in Remm et al. (2001), where it refers to a mutually best hit in an orthologous group.

Over the past decade, many computational methods for ortholog assignment have been proposed, most of which are based primarily on sequence similarity. These methods include the COG system (Tatusov et al., 1997, 2000), EGO (previously called TOGA) (Lee et al., 2002), INPARANOID (Remm et al., 2001), and OrthoMCL (Li et al., 2003), just to name a few. Some of these methods combine sequence similarity with the parsimony principle, such as the reconciled tree method (Yuan et al., 1998) and the bootstrap tree method (Storm and Sonnhammer, 2002), or make use of synteny information, such as OrthoParaMap (Cannon and Young, 2003) and the recent method proposed by Zheng et al. (2005). However, none of these papers use genome rearrangement. On the other hand, there have been a few papers in the literature that study rearrangements between genomes with duplicated genes, which is closely related to ortholog assignment. Sankoff (1999) proposed an approach to identify the true exemplar gene of each gene family, by minimizing the breakpoint/reversal distance between two reduced genomes that consist of only true exemplar genes. El-Mabrouk (2002) developed an approach to reconstruct an ancestor of a modern genome by minimizing the number of duplication transpositions and reversals. The work by Marron et al. (2004) and Swenson et al. (2005a) attempt to find a one-to-one gene correspondence between gene families based on conserved segments. Very recently, Swenson et al. (2005b) presented some algorithmic results on the cycle splitting problem in a combinatorial framework similar to the one introduced in Chen et al. (2005a, 2005b).

2. METHOD

The two genomes to be compared, denoted as Π and Γ , have typically undergone series of genome rearrangement and gene duplication events since they split from their last common ancestral genome. Clearly, we could easily identify main orthologs and inparalogs if given such an evolutionary scenario. Based on this observation, we propose an approach to reconstruct the evolutionary scenario on the basis of the parsimony principle (i.e., we postulate the minimal possible number of rearrangement events and duplication events in the evolution of two genomes since their splitting so as to assign orthologs). Equivalently, it can be formulated as a problem of finding a most parsimonious transformation from one genome into the other by genome rearrangements and gene duplications, without explicitly inferring their ancestral genome. Let $R(\Pi, \Gamma)$ and $D(\Pi, \Gamma)$ denote the number of rearrangement events and the number of gene duplications in a most parsimonious transformation, respectively, and $RD(\Pi, \Gamma)$ denotes the rearrangement/duplication (RD) distance between Π and Γ satisfying $RD(\Pi, \Gamma) = R(\Pi, \Gamma) + D(\Pi, \Gamma)$. Most genome rearrangement events will be considered in this study, including reversal, translocation, fusion and fission.

In practice, we will impose two constraints on this optimization problem, based on some biological considerations. First, we require that at least one member of each family that appears in the other genome be assigned orthology, because each family should carry out an essential function and the gene(s) retaining this function is more likely conserved during the evolution. Second, we observe that the assignment of orthologs that leads to the minimum rearrangement/duplication distance is not necessarily unique. Therefore, among all assignments with the minimum rearrangement/duplication distance, we attempt to find one that also minimizes $R(\Pi, \Gamma)$, in order to avoid introducing unnecessary false orthologous pairs.

Figure 2 presents a simple example to illustrate the basic idea behind our parsimony approach. Consider two genomes, $\Pi = -b - a_1 + c + a_2 + d + a_4 + e + f + g$ and $\Gamma = +a_1 + b + c + a_2 + d + e + a_5 + f + a_3 + g$, sharing a gene family a with multiple copies. As shown in Figure 2, both genomes evolved from the same ancestral genome $+a + b + c + d + e + f + g$, Π by one inversion and one gene duplication and Γ by two gene duplications, respectively. By computing the rearrangement/duplication distance $RD(\Pi, \Gamma) = 4$, the true evolutionary scenario can be reconstructed, which then suggests that the two genes a_1 form a pair of main orthologs, as well as the two genes a_2 . Meanwhile, a_3 , a_4 , and a_5 are inferred as inparalogs that were derived from duplications after the speciation event. It is interesting to see that here a_4 is not assigned orthology to a_3 or a_5 greedily. (Note that they are orthologs, but not main orthologs, by our definition.) This simple example illustrates that, by minimizing the reversal/duplication distance, our approach is able to pick correct main orthologs out of sets of inparalogs.

Note that, although gene loss may occur in the course of evolution, it actually has no impact on the capability of assigning ortholog by our method. If an inparalog is lost, the gene loss event can be simply ignored and this will not affect ortholog assignment. If some gene of a main ortholog pair is lost, our

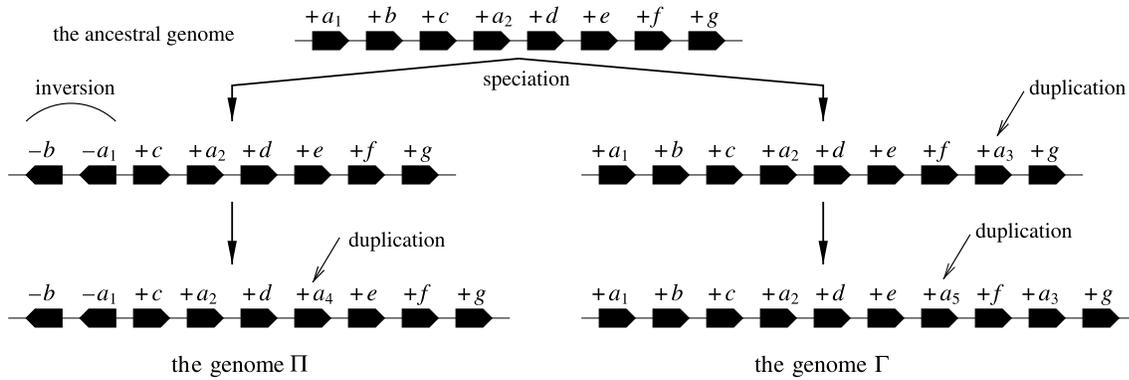


FIG. 2. An evolutionary history of two genomes Π and Γ . Π evolved from the ancestor by one inversion and one gene duplication, and Γ by two duplications.

approach attempts to identify the other gene as an inparalog rather than to assign it a wrong orthology, which also makes sense especially when considering the transfer of functional assignment.

Following the parsimony principle discussed above, we have implemented a high-throughput system for ortholog assignment, called MSOAR. It employs a heuristic to calculate the rearrangement/duplication distance between two genomes, which can be used to reconstruct a most parsimonious evolutionary scenario. In this section, we discuss in detail the heuristic algorithm.

We represent a gene by a symbol of some finite alphabet \mathcal{A} , and its orientation by the sign $+$ or $-$. A chromosome is a sequence of genes, while a genome is a set of chromosomes. Usually, a genome is represented as a set $\Pi = \{\pi(1), \dots, \pi(N)\}$, where $\pi(i) = \langle \pi(i)_1 \dots \pi(i)_{n_i} \rangle$ is a sequence of oriented genes in the i th chromosome. Recall the genome rearrangement problem between two genomes with distinct oriented genes. Hannenhalli and Pevzner (1995, 1999) developed algorithms for calculating genome rearrangement distance on both unichromosomal and multichromosomal genomes in polynomial time. The rearrangement distance between multichromosomal genomes is the minimum number of reversals, translocations, fissions, and fusions that would transform one genome into the other. Given two multichromosomal genomes Π and Γ , Hannenhalli and Pevzner (1995) gave a formula for calculating the genome rearrangement distance (called the HP formula in this paper). Tesler (2002) and Ozery-Flato and Shamir (2003) then suggested some corrections to the formula (called the revised HP formula):

$$d(\Pi, \Gamma) = b(\Pi, \Gamma) - c(\Pi, \Gamma) + p_{\Gamma\Gamma}(\Pi, \Gamma) + r(\Pi, \Gamma) + \left\lceil \frac{s'(\Pi, \Gamma) - gr'(\Pi, \Gamma) + fr'(\Pi, \Gamma)}{2} \right\rceil$$

where $b(\Pi, \Gamma)$ is the number of black edges in the breakpoint graph $G(\Pi, \Gamma)$, $c(\Pi, \Gamma)$ is the overall number of cycles and paths, $p_{\Gamma\Gamma}(\Pi, \Gamma)$ is the number of the $\Gamma\Gamma$ -paths, and r , s' , gr' , and fr' are some parameters in terms of real-knots (Ozery-Flato and Shamir, 2003). In practice, the dominant parts of the formula are the first three terms.

When the genomes Π and Γ contain duplicated genes, however, the rearrangement/duplication distance problem (i.e., $RD(\Pi, \Gamma)$) cannot be directly solved by the revised HP formula. Note that, once the main ortholog pairs are assigned and the inparalogs are identified, the rearrangement/duplication distance can be easily computed as follows. The number of duplications is determined by the number of inparalogs. After removing all the inparalogs, the rearrangement distance between the reduced genomes, which now have equal gene content, can be computed using the above formula since every gene can be regarded as unique. An outline of MSOAR is illustrated in Figure 3.

3. HOMOLOGY SEARCH AND GENE FAMILY CONSTRUCTION

MSOAR starts by calculating the pairwise similarity scores between all gene sequences of the two input genomes, via an all-*versus*-all gene sequence BLASTp comparison. As in INPARANOID (Remm et al., 2001), two cutoffs are applied to each pair of BLASTp hits. Two genes are considered homologous if

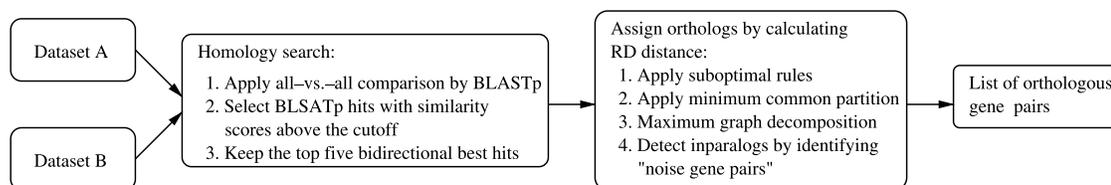


FIG. 3. An outline of MSOAR.

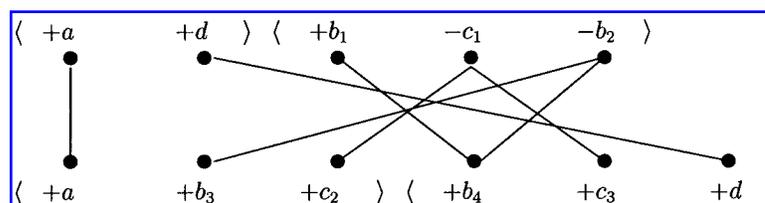


FIG. 4. A hit graph for genomes $\Pi = \{ \langle +a, +d \rangle, \langle +b_1, -c_1, -b_2 \rangle \}$ and $\Gamma = \{ \langle +a, +b_3, +c_2 \rangle, \langle +b_4, +c_3, +d \rangle \}$, each having two chromosomes.

(1) the bit score is no less than 50 and (2) the matching segment spans above 50% of each gene in length. In order to eliminate potential false main ortholog pairs, we take the top five bidirectional best hits of each gene as its potential main orthologs if their logarithmic E-value is less than the 80% of the best logarithmic E-value. By clustering homologous genes using the standard single linkage method, we obtain gene families. A gene family is said to be trivial if it has cardinality exactly 2 (i.e., with one occurrence in each genome). Otherwise it is said to be non-trivial. A gene belonging to a trivial (or non-trivial) family is said to be trivial (or non-trivial, respectively). We use a hit graph (denoted as \mathcal{H}) to describe the relationship between genes within each family. A hit graph is a bipartite graph illustrating the BLASTp hits between two genomes. Each vertex represents a gene and an edge connects two vertices from different genomes if they are potential main orthologs. Figure 4 gives an example of the hit graph. Adjacent genes in the hit graph are regarded as candidates for main ortholog pairs.

4. (SUB)OPTIMAL ASSIGNMENT RULES

We presented three assignment rules for identifying individual ortholog assignments that are (nearly) optimal in SOAR (Chen et al., 2005a, 2005b). In MSOAR, we will add two more assignment rules, which could make the system more efficient. The four rearrangement operations (reversal, translocation, fission and fusion) can be mimicked by reversals when we represent a multichromosomal genome by a *concatenate* (Hannenhalli and Pevzner, 1995; Tesler, 2002). This approach reduces the problem of computing rearrangement distance between two multichromosomal genomes to the problem of computing the reversal distance between two optimal concatenates. Since these two new rules are only concerned with segments of consecutive genes within a single chromosome, which also form gene segments in an optimal concatenate, the unichromosomal HP formula (Hannenhalli and Pevzner, 1999) can be used to prove their (sub)optimality. Let G and H be two chromosomes in genomes Π and Γ , respectively. A chromosome segment is defined as a substring of some chromosome (i.e., a consecutive sequence of genes). A chromosome segment $(g_{i_1} g_{i_2} \dots g_{i_n})$ in G matches a chromosome segment $(h_{j_1} h_{j_2} \dots h_{j_n})$ in H if g_{i_t} and h_{j_t} are connected by an edge in the hit graph and have the same orientations for all $1 \leq t \leq n$.

Rule 1. Suppose that a chromosome segment $(g_{i_1} g_{i_2} \dots g_{i_n})$ in G matches a chromosome segment $(h_{j_1} h_{j_2} \dots h_{j_n})$ in H or its reversal, where $g_{i_1}, g_{i_n}, h_{j_1}$ and h_{j_n} are trivial but the other genes are not. Define two new genomes Π' and Γ' by assigning orthology between g_{i_t} and h_{j_t} or g_{i_t} and $g_{j_{n+1-t}}$ (in the case of matching by a reversal), for all $1 \leq t \leq n$.

Rule 2. Suppose that for a chromosome segment $(g_{i_1}g_{i_2}\cdots g_{i_n})$ in G and a chromosome segment $(h_{j_1}h_{j_2}\cdots h_{j_n})$ in H , g_{i_1} matches h_{j_1} , g_{i_n} matches h_{j_n} , and $g_{i_2}\cdots g_{i_{n-1}}$ matches the reversal of $h_{j_2}\cdots h_{j_{n-1}}$, where g_{i_1} , g_{i_n} , h_{j_1} and h_{j_n} are trivial but the other genes are not. Define two new genomes Π' and Γ' by assigning orthology between g_{i_t} and $g_{j_{n+1-t}}$, for all $1 < t < n$.

We can prove that in both of the above cases, $RD(\Pi, \Gamma) \leq RD(\Pi', \Gamma') \leq RD(\Pi, \Gamma) + 2$ (Fu et al., 2006).

5. MINIMUM COMMON PARTITION

We extend the minimum common partition (MCP) problem, which was first introduced in our recent papers (Chen et al., 2005a, 2005b) to reduce the number of duplicates of each gene in ortholog assignment, to multichromosomal genomes. Use $\overline{\pi(i)}_j$ to represent a chromosome segment or its reversal in chromosome i of genome Π . A chromosome partition is a list $\{\overline{\pi(i)}_1, \overline{\pi(i)}_2, \dots, \overline{\pi(i)}_n\}$ of chromosome segments such that the concatenation of the segments (or their reversals) in some order results in the chromosome i . A genome partition is the union of some partitions of all the chromosomes. A list of chromosome segments is called a common partition of two genomes Π and Γ if it is a partition of both Π and Γ . Furthermore, a minimum common partition is a partition with the minimum cardinality (denoted as $L(\Pi, \Gamma)$) over all possible common partitions of Π and Γ . The MCP problem is the problem of finding the minimum common partition between two given genomes. Two genomes have a common partition if and only if they have equal gene content (i.e., they have the same number of duplications for each gene).

We can further extend MCP to an arbitrary pair of genomes that might have unequal gene contents. A gene matching \mathcal{M} between genomes Π and Γ is a matching between the genes of Π and Γ , which can be defined by a maximum matching in their hit graph \mathcal{H} . Given a gene matching \mathcal{M} , two reduced genomes (denoted as $\tilde{\Pi}_{\mathcal{M}}$ and $\tilde{\Gamma}_{\mathcal{M}}$) with equal gene content can be obtained by removing all the unmatched genes. The minimum common partition of Π and Γ is defined as the minimum $L(\tilde{\Pi}_{\mathcal{M}}, \tilde{\Gamma}_{\mathcal{M}})$ among all gene matchings \mathcal{M} .

Based on the fact that inserting two genes into the two genomes under consideration, one for each genome, will increase the rearrangement distance by at most three, the rearrangement/duplication distance can be approximated by the minimum common partition. An efficient heuristic algorithm for MCP on unichromosomal genomes was given in Chen et al. (2005a, 2005b). The algorithm constructs the so called “pair-match” graphs and then attempts to find a large independent set. We extend the method to multichromosomal genomes in a straightforward way.

6. MAXIMUM GRAPH DECOMPOSITION

After minimum common partition, the genomes Π and Γ may still contain duplicates, although the number of duplicates is expected to be small. In order to match all the genes, we define another problem, called maximum graph decomposition (MGD). The MGD problem is: among all pairs of reduced genomes of Π and Γ obtained by all possible gene matchings, find one with the maximum value of $c(\Pi, \Gamma) - p_{\Gamma\Gamma}(\Pi, \Gamma)$.

We design a greedy algorithm in MSOAR to solve MGD using a new graph, called the complete-breakpoint graph. The complete-breakpoint graph associated with Π and Γ is denoted as \mathcal{G} , which is adapted from the breakpoint graph of multichromosomal genomes of equal gene content consisting of only singletons (Hannenhalli and Pevzner, 1995). The prefix “complete” is added here to differentiate from the partial graphs in Chen et al. (2005a, 2005b). If Π and Γ have different numbers of chromosomes, add null chromosomes to the genome with fewer chromosomes to make them both have N_c chromosomes. As defined by Hannenhalli and Pevzner (1995), a cap is used as a marker that serves as a chromosomal end delimiter when we convert a multichromosomal genome into a unichromosomal genome. A capping of a chromosome $\pi(i)$ is $\pi(i) = \langle \pi(i)_0 \pi(i)_1 \cdots \pi(i)_{n_i} \pi(i)_{n_i+1} \rangle$, where $\pi(i)_0$ is the left cap of $\pi(i)$, called *lcap*, and $\pi(i)_{n_i+1}$ is the right cap, called *rcap*. Choose any capping and an arbitrary concatenation to transform Π and Γ into unichromosomal genomes $\hat{\pi}$ and $\hat{\gamma}$. In the complete-breakpoint graph \mathcal{G} , every

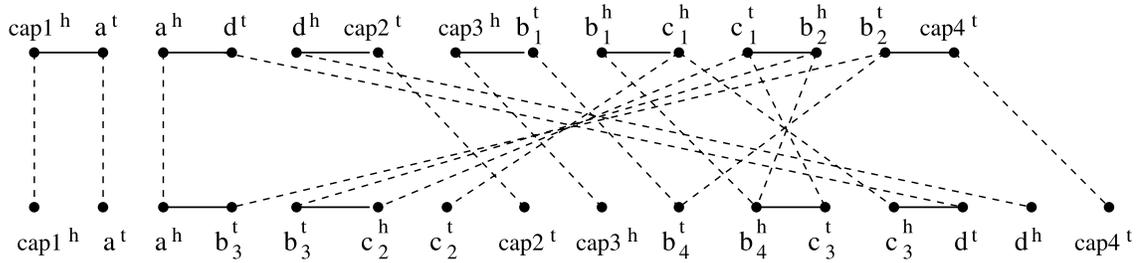


FIG. 5. The complete-breakpoint graph of two genomes with unequal gene contents $\Pi = \{(+a, +d), (+b_1, -c_1, -b_2)\}$ and $\Gamma = \{(+a, +b_3, +c_2), (+b_4, +c_3, +d)\}$. The hit graph of these two genomes is shown in Figure 4.

gene or cap g from each genome is represented by two ordered vertices $\alpha^t \alpha^h$ if α is positive or $\alpha^h \alpha^t$ if it is negative. Insert black edges between vertices that correspond to adjacent genes (or caps) in the same genome except the pairs of the form α_i^h and α_i^t from the same gene (or cap) α_i . Insert cross-genome grey edges $(\hat{\pi}_i^t, \hat{\gamma}_j^t)$ and $(\hat{\pi}_i^h, \hat{\gamma}_j^h)$ if gene $\hat{\pi}_i$ and gene $\hat{\gamma}_j$ are connected by an edge in the hit graph or they are the same caps. Next, we delete the left vertex of every lcap, the right vertex of every rcap and all the edges incident on them. The calculation of RD distance using the resulting graph no longer depends on the actual concatenations. Finally, we make the complete-breakpoint graph independent on the capping of Γ , by deleting the $2N_c$ black edges incident on the $\hat{\gamma}$ cap vertices. These cap vertices are called Π -caps. The vertex on the other end of a deleted black edge is called a Γ -tail unless the black edge arises from a null chromosome, in which case both of its ends are Π -caps. An example of the complete-breakpoint graph is shown in Figure 5. The complete-breakpoint graph contains both cycles and paths. Depending on whether the end points are both Π -caps, both Γ -tails or one of each, a path could be classified as a $\Pi\Pi$ -path, $\Gamma\Gamma$ -path or $\Pi\Gamma$ -path.

After the complete-breakpoint graph is constructed, we try to find small cycles and short $\Pi\Gamma$ -paths first, and then finish the decomposition by finding the rest of $\Pi\Pi$ -paths and $\Gamma\Gamma$ -paths. The decomposition has to satisfy the following three conditions: (1) every vertex belongs to at most one cycle or path (2) the two vertices representing each gene must be connected respectively to the two vertices of a single gene in the other genome by edges of the cycles or paths, otherwise both must be removed, i.e., the connections satisfy a pairing condition; and (3) the edges within a genome and across genomes alternate in a cycle or a path. Intuitively, small cycles may lead to large cycle decompositions, although it is not always the case. Moreover, the more $\Pi\Gamma$ -paths, the fewer $\Gamma\Gamma$ -paths, because the number of Γ -tail vertices is fixed and each vertex can only belong to at most one path. Note that during the cycle decomposition, some gene vertices might have all of their cross-genome edges removed since Π and Γ may have unequal gene contents and these genes are regarded as inparalogs. If two gene vertices α_i^t and α_i^h have no cross genome edges incident on them during the cycle decomposition, they need to be removed from the complete-breakpoint graph right away and a black edge need to be inserted between two endpoints of the deleted black edges arising from α_i^t and α_i^h .

Any feasible solution of the MCD problem gives a maximal matching between the genes of Π and the genes of Γ . The genes that have not been matched will be assigned as inparalogs of the matched ones in the same family. The matched genes suggest possible main ortholog pairs and a rearrangement scenario to transform Π to Γ by the operations reversal, translocation, fusion and fission.

7. “NOISE” GENE PAIR DETECTION

The maximum graph decomposition of a complete-breakpoint graph \mathcal{G} determines a one-to-one gene matching between two genomes. Unmatched genes are removed since either they are potential inparalogs or their orthology counterparts were lost during evolution. However, some individual paralogs might be forced to be assigned as main ortholog pairs because the maximum graph decomposition always gives a maximal matching between all the genes. Therefore, it is necessary to remove these “noise” gene pairs so that the output main ortholog pairs are more reliable.

After removing the unmatched genes, we obtain two reduced genomes with equal gene content. Remove all the gene pairs whose deletion would decrease the rearrangement distance of reduced genomes by at least two. Note that, in this case, the rearrangement/duplication distance will never increase since the deletion of a gene pair may only increase the number of duplications required in an optimal scenario by two. As mentioned before, we require that at least one main ortholog pair of each gene family be kept during this post-processing.

MSOAR combines the suboptimal ortholog assignment rules, heuristic MCP algorithm, heuristic MGD algorithm, and “noise” gene pair detection step to find all the potential main ortholog pairs and detect inparalogs.

8. RESULTS AND DISCUSSION

In order to test the performance of MSOAR as a tool of assigning orthologs, we have applied it to both simulated and real (human and mouse) genome sequence data, and compared its results with two well-known algorithms in the literature, namely, an iterated version of the exemplar algorithm and INPARANOID. We also validated our predicted ortholog pairs between human and mouse genomes using several public datasets on ortholog assignment, synteny information, and gene function classification.

8.1. MSOAR performance on simulated data

In order to assess the validity of our parsimony principle as a means of distinguishing main orthologs from inparalogs, we conduct two simple experiments to estimate the probability of inparalogs that may incorrectly be assigned orthology by transforming one genome into another with the minimum number of rearrangement and duplication events. The first experiment is done as follows. First, we simulate a genome G with 100 distinct genes, and then randomly perform k reversals on G to obtain another genome H . The boundaries of these reversals are uniformly distributed within the genome. Next, make a copy of some gene that is randomly selected from H and insert it back into H as a duplicate. Clearly, the inserted gene is an inparalog, by definition. In this case, there are only two possible ortholog assignments between G and H . Therefore, we can easily calculate the rearrangement/duplicate distance between G and H , and know the ortholog assignment that will be made by the parsimony principle. We repeat the above procedure on 100 random instances for each k , $0 \leq k < 60$, and count the number of instances for which the inparalogs are correctly identified. The distribution of these inparalogs against the number k of reversals is plotted in Figure 6 (the curve marked “one inparalog inserted”). The result shows that with a very high probability ($> 90\%$, for each $k < 55$) the main ortholog and inparalog can be correctly identified. This suggests that an orthology assignment between two genes that are the positional counterparts between two genomes tends to result in the smaller rearrangement/duplication distance, compared to the distance given by an assignment involving inparalogs.

The second experiment is conducted to estimate the probability that two inparalogs, one from each genome, would be identified as a main ortholog pair, instead of as two individual inparalogs. Its data is generated similarly as the first experiment, except that a same copy of the gene is also inserted into genome G , resulting in a non-trivial gene family of size four in both genomes. As before, we count the number of instances for which two inparalogs are correctly identified, and plot its distribution in Figure 6 (the curve marked “two inparalogs inserted”). The result shows that it is very unlikely that two inparalogs from different genomes are assigned as a (main) ortholog pair. This and the above findings provide some basic support for the validity of using the parsimony approach to identify main orthologs.

We further use simulated data to assess the performance of our heuristic algorithm for ortholog assignment. In order to make a comparison test, we implemented the exemplar algorithm (Sankoff, 1999) and extended it into a tool for ortholog assignment as described in Chen et al. (2005a, 2005b), called the iterated exemplar algorithm. The simulated data is generated as follows. Start from a genome G with n distinct symbols whose signs are generated randomly. Each symbol defines a single gene family. Then randomly combine two gene families into a new family until r singletons are left in the genome G . Perform k reversals on G to obtain a genome H as in the previous experiments. Finally, randomly insert c inparalogs (each is a copy of some gene randomly selected) into the two genomes. Note that some singletons may

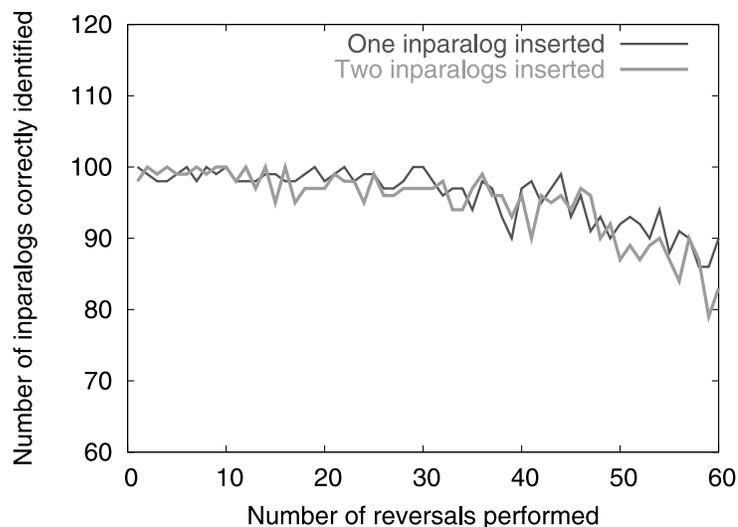


FIG. 6. Distribution of the number of inparalogs correctly identified by the parsimony principle.

be duplicated during this step so that more non-trivial gene families could be generated. The quadruple (n, r, k, c) specifies the parameters for generating two genomes as test data.

We run the iterated exemplar algorithm (Sankoff, 1999; Chen et al., 2005a, 2005b) and our heuristic algorithm on 20 random instances for each combination of parameters. The average performance of both algorithms is shown in Figure 7, in terms of the number of incorrectly assigned orthologs (i.e., genes in a genome that are not assigned orthology to their positional counterparts in the other genome) and inparalogs. As we can see, our heuristic algorithm is quite reliable in assigning orthologs and identifying inparalogs. On average, the number of incorrect assignments generally increases as the number of reversals k increases. While both algorithms perform equally well for inparalogous gene identification, our heuristic algorithm produces fewer incorrect ortholog assignments than the iterated exemplar algorithm, especially for the instances generated using parameters $n = 100$, $r = 80$, and $c = 5$ (Fig. 7).

8.2. MSOAR performance on a pairwise comparison of human and mouse proteomes

We consider two model genomes: human (*Homo sapiens*) and mouse (*Mus musculus*). Gene positions, transcripts and translations were downloaded from the UCSC Genome Browser (Karolchik et al., 2003) web site (<http://genome.ucsc.edu>). We use the canonical splice variants from the Build 36.1 human genome

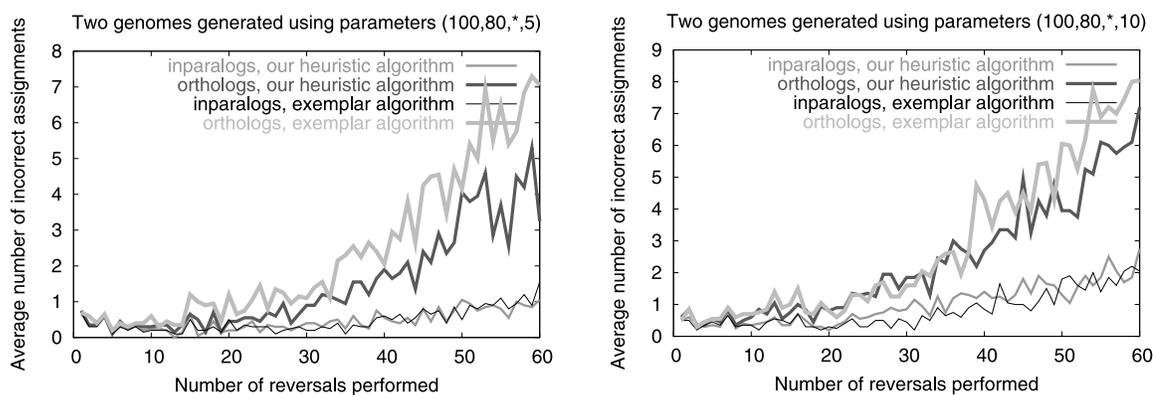


FIG. 7. Comparison of our heuristic and the exemplar algorithm on simulated data.

assembly (UCSC hg18, March 2006) and the Build 36 assembly of the mouse genome (UCSC mm8, February 2006). There are 20161 protein sequences in human genome assembly hg17 and 19199 sequences in mouse genome assembly mm6. Due to assembly errors and other reasons, 393 human and 121 mouse genes were mapped to more than one location in the respective genomes. For such a gene, we kept the first transcription start position which is closest to the 5' end as its start coordinate. A homology search was then performed, and a hit graph between human and mouse was built as described in Section 3.

We validate our assignments by using gene annotation, in particular, gene symbols. To obtain the most accurate list of gene symbols, we have cross-linked database tables from the UCSC Genome Browser with gene symbols extracted from UniProt (Bairoch et al., 2005) release 6.0 (September 2005). The official symbol of a gene is usually given to convey the character or function of the gene (Wain et al., 2002). Genes with identical symbols are most likely to be an orthologous pair, although we should keep in mind that many symbols were given mostly based on sequence similarity and erroneous/inconsistent symbols are known to exist in the annotation. Some genes have not yet been assigned official symbols or their functions have not been validated, so they might only have Riken identifiers or transcript identifiers available, or their gene symbols are constructed as "LOC" + the GeneID. If a pair of genes output by MSOAR have completely identical gene symbol, we count them as a true positive pair; if at least one of them has no valid gene symbol, it is an unknown pair; otherwise, it is counted as a false positive pair. We also calculate the total number of assignable pairs of orthologs (i.e., the total number of pairs of genes with identical symbols). For example, there are 11977 assignable orthologous gene pairs between human and mouse. Before removing "noise" gene pairs, MSOAR predicted 14362 ortholog pairs, among which 11064 are true positives, 1760 are unknown pairs and 1538 are false positives, resulting in a sensitivity of 92.38% and a specificity of 87.80%. After removing "noise" gene pairs, MSOAR predicted 14306 ortholog pairs, among which 11050 are true positives, 1748 are unknown pairs and 1508 are false positives, resulting in a sensitivity of 92.26% and a specificity of 87.99%. It is interesting to note that the last step of MSOAR identified 56 "noise" gene pairs, among which 75% (or 42) were false positives and unknown pairs. This result shows that the identification of "noise" gene pairs effectively detects false positives and could provide more reliable ortholog assignment.

The comparison result between MSOAR and INPARANOID (Remm et al., 2001) is shown in Table 1. MSOAR was able to identify 99 more true ortholog pairs than INPARANOID, although it also reported more false positives. Recall that INPARANOID is an ortholog assignment tool based solely on sequence similarity. It uses all-versus-all pairwise gene sequence comparison and bidirectional best hits (BBHs) to identify the main orthologs (and their associated inparalogs) between a given pair of genomes. Figure 8 illustrates two examples where INPARANOID failed to assign main ortholog pairs correctly because the two genes in each pair are not bidirectional best hits of each other, but MSOAR successfully resolved them.

In order to investigate the relationship between the syntenic blocks shared between human and mouse genomes and blocks defined by clustering consecutive MSOAR orthologous pairs, we applied MSOAR to UCSC hg16 and UCSC mm3, the most recent assemblies of human and mouse genomes for which synteny information is available at the UCSC Genome Browser. MSOAR blocks were obtained by clustering consecutive MOSAR-assigned orthologs with consistent pairs of chromosomal assignments (i.e., the genes involved in these orthologous pairs that are from the same genome must be on the same chromosome) and monotonically increasing or decreasing genomic coordinates on both genomes. In order to remove small noisy blocks, following the first iteration, all clusters with two or fewer genes were removed, and adjacent clusters were concatenated where it was possible. The genomic coordinates, orientations and chromosomal assignments of the MSOAR blocks are in agreement with syntenic blocks on all the chromosomes except the Y chromosome, where very few orthologs were identified and only (relatively) very small synteny

TABLE 1. COMPARISON OF ORTHOLOG ASSIGNMENTS BETWEEN MSOAR AND INPARANOID

	<i>Assignable</i>	<i>Assigned</i>	<i>True positive</i>	<i>Unknown</i>
MSOAR (before removing "noise" pairs)	11,977	14,362	11,064	1760
MSOAR (after removing "noise" pairs)	11,977	14,306	11,050	1748
INPARANOID	11,977	13,837	10,951	1615

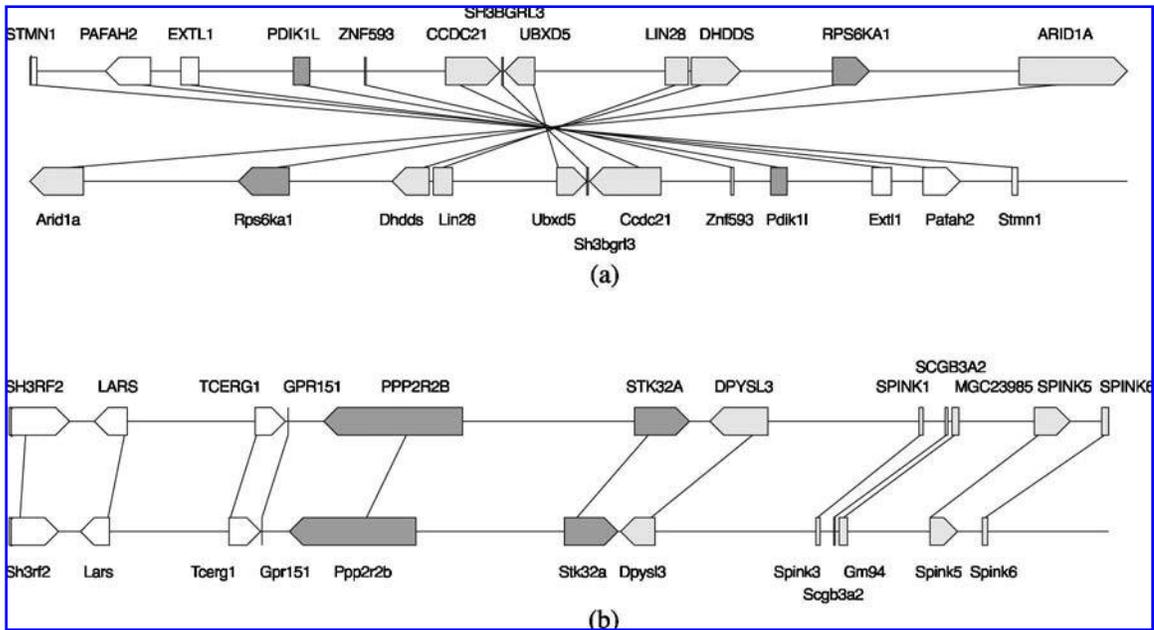


FIG. 8. (a) Assigned orthologs on the segment of human chromosome 1 (26,099,194–26,981,180) and the corresponding segment of mouse chromosome 4 (132,951,085–133,745,914), where MSOAR correctly identified mouse orthologs of genes PDIK1L and RPS6KA1 which were missed by INPARANOID, and (b) Segment of the human chromosome 5 (145,296,334–147,574,892) and mouse chromosome 18 (42,179,639–44,208,941), where MSOAR identified orthologs of PPP2R2B and STK32A which were missed by INPARANOID.

blocks exist. Thus, in essence, by using global genome rearrangements of genes MSOAR was able to capture the same information provided by the syntenic blocks, which were built by using unique sequence tags. However, because MSOAR blocks were built in a greedy fashion, it may happen that an MSOAR block spans several adjacent syntenic blocks, as is the case with human chromosome 4 and mouse chromosome 5. Figure 9 summarizes our findings. The interactive version of this figure, which displays genomic coordinates and allows zooming into the of the MSOAR blocks, can be found online at <http://msoar.cs.ucr.edu/>.

In order to test the effectiveness of orthologs assignment using MSOAR, we have performed extensive validation using HGNC Comparison of Orthology Predictions (HCOP) Search Tool (www.gene.ucl.ac.uk/cgi-bin/nomenclature/hcop.pl) (Wright et al., 2005; Eyre et al., 2006). HCOP contains the known orthology assertions for human and mouse genes as provided by HGNC (Wright and Bruford, 2006), MGI (Eppig et al., 2005), Homologene (Wheeler et al., 2006), Ensembl (Birney et al., 2006), PhIGs (<http://phigs.jgi-psf.org/>), and INPARANOID (Remm et al., 2001). Of the 14,306 main ortholog pairs for human and mouse identified by MSOAR, 14,083 pairs have valid Entrez Gene IDs (Maglott et al., 2005) which could be used to compare with the orthology assertions by the six methods on HCOP. Interestingly, 13,479 ortholog pairs (out of 14,083, or 95.7%) have at least one support from HCOP (i.e., are agreed on by at least one of the six other methods). Moreover, 7349 ortholog pairs (i.e., above 70%), are supported by at least five other methods. The distribution of the number of supports that the MSOAR assignment receives from HCOP is shown in Figure 10.

An important application of ortholog assignment is the functional inferences of proteins. In order to test the effectiveness of protein functional analysis using MSOAR, we have performed a validation using the PANTHER (Protein Analysis Through Evolutionary Relationships) classification system (www.pantherdb.org/) (Thomas et al., 2003). PANTHER is an online resource that classifies genes by their functions, based on a method that uses published scientific experimental evidence or evolutionary relationship to predict functions in the absence of direct experimental evidence. We examined the consistency of the main ortholog pairs output by MSOAR with respect to the PANTHER IDs assigned by the PANTHER classification system, with the assumption that proteins sharing the same PANTHER ID are classified into the same functional family and subfamily by expert biologists. The functional family and subfamily

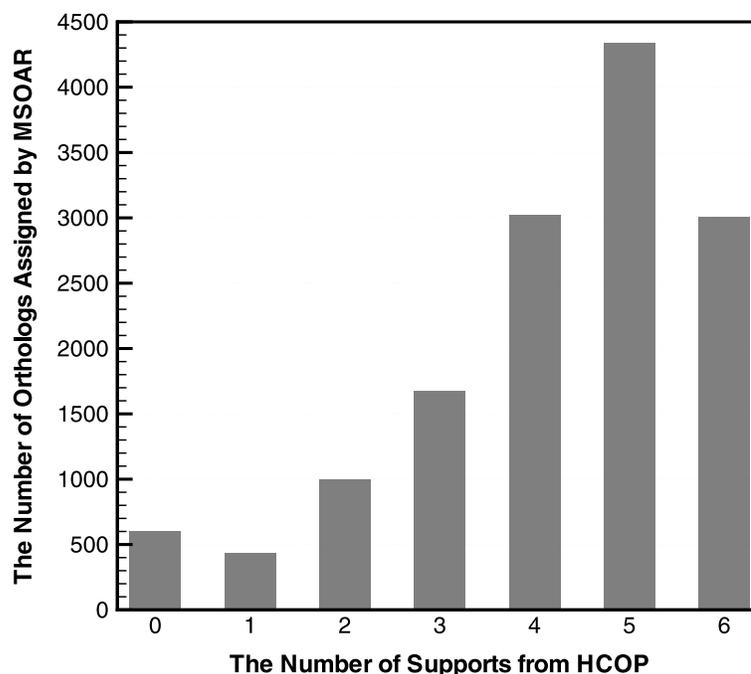


FIG. 10. Distribution of the number of supports that the MSOAR result gets from HCOP. The x -axis represents the number of the human-mouse main ortholog pairs assigned by MSOAR and the y -axis indicates the number of supports from HCOP.

information for all the human and mouse proteins was extracted (if available) from PANTHER. Out of the 14,083 orthologous gene pairs between human and mouse identified by MSOAR with valid Entrez Gene IDs, 11,887 (or 84.4%) pairs are classified into the same functional subfamily by PANTHER. This result demonstrates that although MSOAR relies mainly on genome rearrangements and gene duplications instead of sequence similarity, the main ortholog pairs obtained by MSOAR are very consistent with gene functional classification, indicating an accurate and meaningful ortholog assignment.

9. CONCLUSION

Although we anticipate that the system MSOAR will be a very useful tool for ortholog assignment, more systematic tests will be needed to reveal the true potential of this parsimony approach. Our immediate future work includes the incorporation of transpositions into the system and considering of weighting the evolutionary events. We plan to extend the method to multiple species comparisons.

ACKNOWLEDGMENTS

This project was supported in part by NSF (grant CCR-0309902), National Key Project for Basic Research (973; grant 2002CB512801), NSFC (grant 60528001), and a Changjiang Visiting Professorship at Tsinghua University. Also, the authors wish it to be known that the first two authors should be regarded as joint First Authors of this paper.

REFERENCES

Altschul, S.F., Madden, T.L., Schaffer, A.A., et al. 1997. Gapped blast and psi-blast: a new generation of protein database search programs. *Nucleic Acids Res.* 25, 3389–3402.

- Bairoch, A., Apweiler, R., Wu, C.H., et al. 2005. The universal protein resource (uniprot). *Nucleic Acids Res.* 33, D154–D159.
- Birney, E., Andrews, D., Caccamo, M., et al. 2006. Ensembl 2006. *Nucleic Acids Res.* 34 (Database Issue), D556–D561.
- Cannon, S.B., and Young, N.D. 2003. Orthoparamap: distinguishing orthologs from paralogs by integrating comparative genome data and gene phylogenies. *BMC Bioinform.* 4, 35.
- Chen, X., Zheng, J., Fu, Z., et al. 2005a. Computing the assignment of orthologous genes via genome rearrangement. *Proc. 3rd Asia Pac. Bioinform. Conf.* 363–378.
- Chen, X., Zheng, J., Fu, Z., et al. 2005b. The assignment of orthologous genes via genome rearrangement. *IEEE/ACM Trans. Comput. Biol. Bioinform.* 2, 302–315.
- El-Mabrouk, N. 2002. Reconstructing an ancestral genome using minimum segments duplications and reversals. *J. Comput. Syst. Sci.* 65, 442–464.
- Eppig, J.T., Bult, C.J., Kadin, J.A., et al. 2005. The mouse genome database (mgd): from genes to mice: a community resource for mouse biology. *Nucleic Acids Res.* 33, D471–D475.
- Eyre, T.A., Wright, M.W., Lush, M.J., et al. 2007. Hcop: a searchable database of human orthology predictions. *Brief Bioinform.* 8, 2–5.
- Fitch, W.M. 1970. Distinguishing homologous from analogous proteins. *Syst. Zool.* 19, 99–113.
- Fu, Z., Chen, X., Vacic, V., et al. 2006. A parsimony approach to genome-wide ortholog assignment. *RECOMB 2006* 578–594.
- Hannenhalli, S., and Pevzner, P.A. 1995. Transforming men into mice (polynomial algorithm for genomic distance problem). *Proc. 36th Annu. Symp. Found. Comput. Sci.* 581–592.
- Hannenhalli, S., and Pevzner, P.A. 1999. Transforming cabbage into turnip (polynomial algorithm for sorting signed permutations by reversals). *J. ACM* 46, 1–27.
- Karolchik, D., Baertsch, R., Diekhans, M., et al. 2003. The ucsc genome browser database. *Nucleic Acids Res.* 31, 51–54.
- Koonin, E.V. 2005. Orthologs, paralogs, and evolutionary genomics. *Annu. Rev. Genet.* 39, 309–338.
- Lee, Y., Sultana, R., Pertea, G., et al. 2002. Cross-referencing eukaryotic genomes: Tigr orthologous gene alignments (toga). *Genome Res.* 12, 493–502.
- Li, L., Stoeckert, C., and Roos, D. 2003. Orthomcl: identification of ortholog groups for eukaryotic genomes. *Genome Res.* 13, 2178–2189.
- Maglott, D., Ostell, J., Pruitt, K., et al. 2005. Entrez gene: gene-centered information at NCBI. *Nucleic Acids Res.* 33 (Database Issue), D54–D58.
- Marron, M., Swenson, K.M., and Moret, B.M.E. 2004. Genomic distances under deletions and insertions. *Theoret. Comput. Sci.* 325, 347–360.
- Ozery-Flato, M., and Shamir, R. 2003. Two notes on genome rearrangements. *J. Bioinform. Comput. Biol.* 1, 71–94.
- Remm, M., Storm, C.E., and Sonnhammer, E.L. 2001. Automatic clustering of orthologs and in-paralogs from pairwise species comparisons. *J. Mol. Biol.* 314, 1041–1052.
- Sankoff, D. 1999. Genome rearrangement with gene families. *Bioinformatics* 15, 909–917.
- Storm, C.E., and Sonnhammer, E.L. 2002. Automated ortholog inference from phylogenetic trees and calculation of orthology reliability. *Bioinformatics* 18, 92–99.
- Swenson, K.M., Marron, M., Earnest-DeYoung, J., et al. 2005a. Approximating the true evolutionary distance between two genomes. *Proc. 7th SIA Workshop Algorithm Eng. Exp.* 121–125.
- Swenson, K.M., Pattengale, N., and Moret, B.M.E. 2005b. A framework for orthology assignment from gene rearrangement data. *RECOMB 2005*, 153–166.
- Tatusov, R.L., Galperin, M.Y., Natale, D.A., et al. 2000. The cog database: a tool for genome-scale analysis of protein functions and evolution. *Nucleic Acids Res.* 28, 33–36.
- Tatusov, R.L., Koonin, E., and Lipman, D.J. 1997. A genomic perspective on protein families. *Science* 278, 631–637.
- Tesler, G. 2002. Efficient algorithms for multichromosomal genome rearrangements. *J. Comput. Syst. Sci.* 65, 587–609.
- Thomas, P.D., Campbell, M.J., Kejariwal, A., et al. 2003. Panther: a library of protein families and subfamilies indexed by function. *Genome Res.* 13, 2129–2141.
- Wain, H.M., Bruford, E.A., Lovering, R.C., et al. 2002. Guidelines for human gene nomenclature. *Genomics* 79, 464–470.
- Wheeler, D.L., Barrett, T., Benson, D.A., et al. 2006. Database resources of the national center for biotechnology information. *Nucleic Acids Res.* 34 (Database Issue), D173–D180.
- Wright, M.W., and Bruford, E.A. 2006. Human and orthologous gene nomenclature. *Gene* 369, 1–6.
- Wright, M.W., Eyre, T.A., Lush, M.J., et al. 2005. Hcop: the hgnc comparison of orthology predictions search tool. *Mamm. Genome* 16, 827–828.

- Yuan, Y.P., Eulenstein, O., Vingron, M., et al. 1998. Towards detection of orthologues in sequence databases. *Bioinformatics* 14, 285–289.
- Zheng, X.H., Lu, F., Wang, Z.Y., et al. 2005. Using shared genomic synteny and shared protein functions to enhance the identification of orthologous gene pairs. *Bioinformatics* 21, 703–710.

Address reprint requests to:

Dr. Zheng Fu

Department of Computer Science and Engineering

University of California

Riverside, CA 92521

E-mail: zfu@cs.ucr.edu