

# On Inferring Additive and Replacing Horizontal Gene Transfers Through Phylogenetic Reconciliation

Misagh Kordi  
 Department of Computer Science  
 and Engineering  
 University of Connecticut  
 Storrs, CT, USA  
 misagh.kordi@uconn.edu

Soumya Kundu\*  
 Department of Computer Science  
 and Engineering  
 University of Connecticut  
 Storrs, CT, USA  
 soumya.kundu@uconn.edu

Mukul S. Bansal†  
 Department of Computer Science  
 and Engineering  
 University of Connecticut  
 Storrs, CT, USA  
 mukul.bansal@uconn.edu

## ABSTRACT

Horizontal gene transfer is one of the most important mechanisms for microbial evolution and adaptation. It is well known that horizontal gene transfer can be either *additive* or *replacing* depending on whether the transferred gene adds itself as a new gene in the recipient genome or replaces an existing homologous gene. Yet, all existing phylogenetic techniques for the inference of horizontal gene transfer assume either that all transfers are additive or that all transfers are replacing. This limitation not only affects the applicability and accuracy of these methods but also makes it difficult to distinguish between additive and replacing transfers.

Here, we address this important problem by formalizing a phylogenetic reconciliation framework that simultaneously models both additive and replacing transfer events. Specifically, we (1) introduce the *DTRL* reconciliation framework that explicitly models both additive and replacing transfer events, along with gene duplications and losses, (2) prove that the underlying computational problem is NP-hard, (3) perform the first experimental study to assess the impact of replacing transfer events on the accuracy of the traditional DTL reconciliation model (which assumes that all transfers are additive) and demonstrate that traditional DTL reconciliation remains highly robust to the presence of replacing transfers, (4) propose a simple heuristic algorithm for DTRL reconciliation based on classifying transfer events inferred through DTL reconciliation as being replacing or additive, and (5) evaluate the classification accuracy of the heuristic under a range of evolutionary conditions. Thus, this work lays the methodological and algorithmic foundations for estimating DTRL reconciliations and distinguishing between additive and replacing transfers.

An implementation of our heuristic for DTRL reconciliation is freely available open-source as part of the RANGER-DTL software package from <https://compbio.engr.uconn.edu/software/ranger-dtl/>.

\*Currently with the Department of Computer Science at Stanford University.

†Also affiliated with the Institute for Systems Genomics at the University of Connecticut.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

ACM-BCB '19, September 7–10, 2019, Niagara Falls, NY, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.  
 ACM ISBN 978-1-4503-6666-3/19/09.

<https://doi.org/10.1145/3307339.3342168>

## KEYWORDS

Phylogenetic reconciliation; horizontal gene transfer; additive transfer; replacing transfer; xenologous gene displacement; DTRL reconciliation; microbial gene family evolution

### ACM Reference Format:

Misagh Kordi, Soumya Kundu, and Mukul S. Bansal. 2019. On Inferring Additive and Replacing Horizontal Gene Transfers Through Phylogenetic Reconciliation. In *10th ACM International Conference on Bioinformatics, Computational Biology and Health Informatics (ACM-BCB '19)*, September 7–10, 2019, Niagara Falls, NY, USA. ACM, New York, NY, USA, 10 pages.

## 1 INTRODUCTION

The transfer of genetic information between organisms that are not in a direct ancestor-descendant relationship, called *horizontal gene transfer* or simply *transfer* for short, is a crucial process in microbial evolution. The problem of detecting transfer events has been extensively studied and many different methods have been developed for the problem; see, e.g., [30] for a review. The two most widely used classes of methods are those based on atypical sequence composition and those based on phylogenetic discordance. Sequence composition methods look for atypical dinucleotide frequencies, codon usage biases, or other sequence features that might indicate instances of horizontally acquired genes, but are only effective at short evolutionary time scales and are unable to accurately identify the donors and recipients of transfer events [12, 30]. Phylogenetic methods rely on the fact that horizontal transfers leave tell-tale phylogenetic signatures in the topologies of the transferred genes. These methods construct gene trees for individual gene families and compare them to known species phylogenies to infer possible transfer events. It is well-understood that when a gene is horizontally transferred, it may either add itself as a new gene to the recipient genome, resulting in an *additive* transfer, or replace an existing homologous gene, resulting in a *replacing* transfer [9, 18, 19]. Yet, there do not currently exist any phylogenetic methods that simultaneously model both these types of transfers. This limitation not only affects the applicability and accuracy of these methods but also makes it difficult to distinguish between additive and replacing transfers.

Phylogenetic methods for inferring transfer events can be divided into two classes: (i) Those that implicitly assume that all transfers are replacing transfers and that all discordance between gene trees and species trees is due to these replacing transfer events, e.g., [1, 5, 6, 13, 15, 17, 23, 28], and (ii) those based on the Duplication-Transfer-Loss (DTL) reconciliation framework, which model gene

duplication and gene loss as additional sources of gene tree/species tree discordance, but implicitly assume that all transfers are additive transfers, e.g., [2, 8, 10, 11, 16, 20, 22, 24–27, 29]. Thus, no existing phylogenetic method models both additive and replacing transfers. And while methods based on DTL reconciliation represent a major advance in the ability to accurately detect transfer events, they are limited by their inability to properly handle replacing transfers.

**Our contribution.** In this work, we define and formalize a phylogenetic reconciliation framework that simultaneously models both additive and replacing transfer events. Our framework builds upon the standard parsimony-based DTL reconciliation model [2, 29], which assumes that the species tree is undated and seeks an optimal (and not necessarily time-consistent) reconciliation, by explicitly modeling replacing transfer events.<sup>1</sup> Specifically, we formally define the *Duplication-Additive-Transfer-Replacing-Transfer-Loss* (DTRL) reconciliation model that explicitly models both additive and replacing transfer events, along with gene duplications and losses. As with the underlying DTL reconciliation model, we formulate the DTRL reconciliation problem as one of finding a *most parsimonious* DTRL reconciliation, i.e., one with smallest total “reconciliation cost”. We prove that the problem of computing a most parsimonious DTRL reconciliation is NP-hard, using a reduction from the NP-hard *minimum rooted Subtree Prune and Regraft (rSPR) distance* problem, and perform the very first experiments to study the impact of replacing transfer events on the accuracy of DTL reconciliation itself. Surprisingly, we found that DTL reconciliation is highly robust to the presence of replacing transfer. Based on these results, we devise a simple heuristic to classify transfer events inferred through DTL reconciliation as being either additive or replacing, and study its classification accuracy using simulated datasets over a range of evolutionary conditions. Our experimental results show that, even though the problem of inferring optimal DTRL reconciliations is NP-hard, it should be possible to design effective heuristics for the problem based on the simpler, and efficiently solvable, DTL reconciliation model.

We note that the problem of integrating replacing transfers with DTL reconciliation has also been recently, and independently, studied by Hasic and Tannier in a recently published manuscript [14]. That manuscript proves that the problem of inferring replacing transfers through phylogenetic reconciliation is NP-hard when the species tree is dated. However, the results in that manuscript are largely complementary to the current work. Specifically, we provide a rigorous and precise formalization of the DTRL reconciliation framework, our proof of NP-hardness is not only completely different but applies to the undated version of the problem where the species tree is undated (arguably the more widely applicable version of the problem), we provide the first experimental results on the impact of replacing transfer on conventional DTL reconciliation, and we devise and evaluate the first heuristic algorithm for estimating optimal DTRL reconciliations.

<sup>1</sup>Note that the DTL reconciliation model [2, 29] on which our new model is based allows the inferred reconciliation to be time-inconsistent. This is simply because imposing time consistency makes the DTL reconciliation problem NP-hard [29], while allowing time-inconsistency makes the problem efficiently solvable with negligible impact on accuracy.

The remainder of the manuscript is organized as follows: Basic definitions, preliminaries, and a formal description of the DTRL reconciliation model appear in the next section. The NP-hardness proof appears in Section 3, and experimental results on the effect of replacing transfers on DTL reconciliation are described in Section 4, and our heuristic for classifying transfers is described and tested in Section 5. Concluding remarks appear in Section 6. For brevity, all proofs are deferred to the full version of this manuscript.


## 2 DEFINITIONS AND PRELIMINARIES

We follow basic definitions and notation from [2]. Given a rooted tree  $T$ , we denote its node, edge, and leaf sets by  $V(T)$ ,  $E(T)$ , and  $Le(T)$  respectively. The root node of  $T$  is denoted by  $rt(T)$ , the parent of a node  $v \in V(T)$  by  $pa_T(v)$ , its set of children by  $Ch_T(v)$ , and the (maximal) subtree of  $T$  rooted at  $v$  by  $T(v)$ . The set of *internal nodes* of  $T$ , denoted  $I(T)$ , is defined to be  $V(T) \setminus Le(T)$ . We define  $\leq_T$  to be the partial order on  $V(T)$  where  $x \leq_T y$  if  $y$  is a node on the path between  $rt(T)$  and  $x$ . The partial order  $\geq_T$  is defined analogously, i.e.,  $x \geq_T y$  if  $x$  is a node on the path between  $rt(T)$  and  $y$ . We say that  $y$  is an *ancestor* of  $x$ , or that  $x$  is a *descendant* of  $y$ , if  $x \leq_T y$  (note that every node is a descendant as well as ancestor of itself). We say that  $x$  and  $y$  are *incomparable* if neither  $x \leq_T y$  nor  $y \leq_T x$ . Given a non-empty subset  $L \subseteq Le(T)$ , we denote by  $lca_T(L)$  the last common ancestor (LCA) of all the leaves in  $L$  in tree  $T$ ; that is,  $lca_T(L)$  is the unique smallest upper bound of  $L$  under  $\leq_T$ . Given  $x, y \in V(T)$ ,  $x \rightarrow_T y$  denotes the unique path from  $x$  to  $y$  in  $T$ . We denote by  $dist_T(x, y)$  the number of edges on the path  $x \rightarrow_T y$ ; note that if  $x = y$  then  $dist_T(x, y) = 0$ . Given a set  $L \subseteq Le(T)$ , let  $T'$  be the minimal rooted subtree of  $T$  with leaf set  $L$ . We define the *leaf induced subtree* of  $T$  on leaf set  $L$ , denoted  $T[L]$ , to be the tree obtained from  $T'$  by successively removing each non-root node of degree two and adjoining its two neighbors. A tree is *binary* if all of its internal nodes have exactly two children. Throughout this work, the *term* tree refers to rooted binary trees.

A *species tree* is a tree that depicts the evolutionary relationships of a set of species. Given a gene family from a set of species, a *gene tree* is a tree that depicts the evolutionary relationships among the sequences encoding only that gene family in the given set of species. Thus, the nodes in a gene tree represent genes. Throughout this work, we denote the gene tree and species tree under consideration by  $G$  and  $S$ , respectively. We assume that each leaf of the gene tree is labeled with the species from which that gene (sequence) was obtained. This labeling defines a *leaf-mapping*  $\mathcal{L}_{G,S}: Le(G) \rightarrow Le(S)$  that maps a leaf node  $g \in Le(G)$  to that unique leaf node  $s \in Le(S)$  which has the same label as  $g$ . Note that the gene tree can have zero, one, or more than one gene from any species under consideration. The species tree contains at least all the species represented in the gene tree.

### 2.1 Additive and replacing transfers

When a gene is horizontally transferred, there are two possibilities for how it may incorporate itself into the recipient genome. The first possibility is that the transferred gene inserts itself to the recipient genome without overwriting any existing genes, thereby creating a new gene locus for itself. The second possibility is that



**Figure 1: Additive and replacing transfers.** This figure shows the evolution of two gene families inside the same species tree. Both gene families exist in the root of the species tree and evolve according to the topology of the species tree without any gene duplications or losses. Gene family 1 is affected by a replacing transfer event, as shown in the figure by the upper orange (dashed) arrow. Gene family 2 is affected by an additive transfer event, as shown by the lower orange (dashed) arrow. The topologies of the resulting gene trees for these two gene families are shown.

the transferred gene replaces an existing homologous copy of itself, preserving the total number of genes in the recipient genome; this type of horizontal transfer is sometimes also referred to as xenologous gene displacement [19].

**DEFINITION 2.1 (ADDITIVE TRANSFER).**

An additive transfer is a horizontal gene transfer that inserts itself into the recipient genome through the addition of a new gene locus.

**DEFINITION 2.2 (REPLACING TRANSFER).**

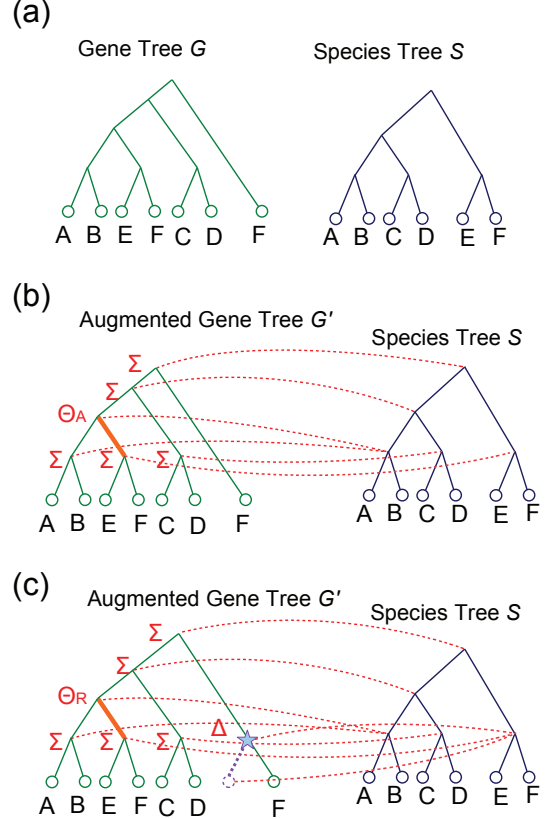
A replacing transfer is a horizontal gene transfer that inserts itself into the recipient genome by replacing a homologous gene at an existing gene locus.

Note that additive transfers result in an increase in the total number of genes in the recipient genome, while replacing transfers do not. We also point out that replacing transfers can only happen if the recipient genome already contains a homologous copy of the gene being transferred. Figure 1 illustrates how additive and replacing transfer events impact the resulting gene tree topology.

## 2.2 DTRL Reconciliation

The *D*uplication–*A*dditive–*T*ransfer–*R*eplacing–*T*ransfer–*L*oss (DTRL) Reconciliation model is based upon the well-studied parsimony-based DTL reconciliation framework [2, 29] (which implicitly assumes that all transfer events are additive). However, the introduction of replacing transfers into the model poses several challenges, as we describe below, and the DTL reconciliation framework must therefore be substantially extended to allow for replacing transfers. Specifically, to fully specify a DTRL reconciliation, we must (i) account for hidden duplication or transfer events that do not label any node of the gene tree, and (ii) include in the reconciliation

those gene lineages that have been lost (i.e., are no longer visible on the gene tree) but which played a role in the evolution of that gene family by participating in transfer events. We elaborate on these below.



**Figure 2: Hidden events and augmented gene trees.** Parts (b) and (c) of the figure show two alternative DTRL reconciliations for the gene tree  $G$  and species tree  $S$  shown in Part (a). Each reconciliation shows the augmented gene tree  $G'$ , the event type for each internal node in the augmented gene tree (where  $\Sigma$  represents speciation,  $\Delta$  represents duplication,  $\Theta_A$  represents additive transfer, and  $\Theta_R$  represents replacing transfer), and the red arcs show the mapping for each node of  $G'$  not in  $Le(G)$  (the mapping for each leaf node of  $G$  is implicitly defined by its leaf label). The bold orange edges represent transfer edges. The reconciliation in Part (b) invokes an additive transfer event and a loss event. For this reconciliation in Part (b),  $G'$  is the same as  $G$ . The reconciliation in Part (c) invokes a replacing transfer event, a hidden gene duplication event (marked by the blue star), and a loss event. The invisible lineage replaced by the replacing transfer event is shown by the purple dotted line in  $G'$ .

**Hidden events.** Unlike the DTL reconciliation model, where each speciation, duplication, or transfer event required by the reconciliation can be assigned to an individual gene tree node, a most-parsimonious DTRL reconciliation may postulate duplication and

transfer events (additive or replacing) that cannot be assigned to any node on the gene tree. Such *hidden events* may be required for most-parsimonious DTRL reconciliation but are invisible on the gene tree either because only descendants from one of the loci resulting from a duplication or additive transfer event survive in the gene family or because they appear on an invisible lineage. The reason hidden events can occur in optimal DTRL reconciliations is that one of the loci resulting from the hidden event is subsequently used (and overwritten) by one or more replacing transfers. This phenomenon is illustrated in Figure 2.

**Invisible gene lineages and augmented gene trees.** To properly recover replacing transfer events and correctly count the number of losses, it is necessary to postulate and account for those gene lineages that are no longer visible on the gene tree but which played a role in the evolution of that gene family by participating in replacing transfer events. Such invisible gene lineages can result from duplication, speciation, or transfer events, but become invisible because no descendants survive in the extant gene family. If these lineages do not participate in any transfer events that impacted the rest of the gene tree, then they can be safely ignored, but otherwise they must be accounted for if replacing transfers are to be recovered accurately and the number of losses counted correctly. We account for invisible lineages by augmenting the input gene tree with additional edges/subtrees, resulting in an *augmented gene tree*, and showing the DTRL reconciliation for this entire augmented gene tree. Figure 2 shows an example of an augmented gene tree and illustrates why it is important to consider invisible gene lineages.

The DTRL reconciliation model takes as input a rooted gene tree and a rooted species tree and defines a framework for reconciling the gene tree with the species tree by postulating duplication, additive transfer, replacing transfer, and gene loss events. The reconciliation creates an augmented gene tree, maps each augmented gene tree node to a unique species tree node, respecting the temporal constraints implied by the species tree topology, and designates each augmented gene tree node as representing either a speciation, duplication, additive transfer, or replacing transfer event. For any gene tree node, say  $g$ , that represents a transfer event, the reconciliation also specifies which of the two edges  $(g, g')$  or  $(g, g'')$ , where  $g', g''$  denote the children of  $g$ , represents the transfer edge and identifies the recipient species of the corresponding transfer. If  $g$  represents a replacing transfer event, the reconciliation also identifies the specific gene lineage that was lost as a result of that replacing transfer.

Next, we define what constitutes a valid DTRL reconciliation.

**DEFINITION 2.3 (DTRL-RECONCILIATION).**

A DTRL-reconciliation for  $G$  and  $S$  is a ten-tuple  $\langle \mathcal{L}, G', \mathcal{M}, \Sigma, \Delta, \Theta_A, \Theta_R, \Xi, \tau, \lambda \rangle$ , where  $\mathcal{L}: Le(G) \rightarrow Le(S)$  represents the leaf-mapping from  $G$  to  $S$ ,  $G'$  represents the augmented gene tree,  $\mathcal{M}: V(G') \rightarrow V(S)$  maps each node of  $G'$  to a node of  $S$ , the sets  $\Sigma, \Delta, \Theta_A$  and  $\Theta_R$  partition  $I(G')$  into speciation, duplication, additive transfer, and replacing transfer nodes, respectively,  $\Xi$  is a subset of  $E(G')$  that represents transfer edges (additive or replacing),  $\tau: \Theta_A \cup \Theta_R \rightarrow V(S)$  specifies the recipient species for each transfer

event, and  $\lambda: \Theta_R \rightarrow Le(G') \setminus Le(G)$  is an injective function that associates each replacing transfer event with a lost gene in the augmented gene tree, subject to the following constraints:

**Augmented gene tree constraint**

- (1)  $G = G'[Le(G)]$ .

**Mapping constraints**

- (2) If  $g \in Le(G)$ , then  $\mathcal{M}(g) = \mathcal{L}(g)$ .  
(3) If  $g \in I(G')$  and  $g'$  and  $g''$  denote the children of  $g$ , then,  
(a)  $\mathcal{M}(g) \not\prec_S \mathcal{M}(g')$  and  $\mathcal{M}(g) \not\prec_S \mathcal{M}(g'')$ ,  
(b) At least one of  $\mathcal{M}(g')$  and  $\mathcal{M}(g'')$  is a descendant of  $\mathcal{M}(g)$ .

**Event constraints**

- (4) Given any edge  $(g, g') \in E(G')$ ,  $(g, g') \in \Xi$  if and only if  $\mathcal{M}(g)$  and  $\mathcal{M}(g')$  are incomparable.  
(5) If  $g \in I(G')$  and  $g'$  and  $g''$  denote the children of  $g$ , then,  
(a)  $g \in \Sigma$  only if  $\mathcal{M}(g) = lca(\mathcal{M}(g'), \mathcal{M}(g''))$  and  $\mathcal{M}(g')$  and  $\mathcal{M}(g'')$  are incomparable,  
(b)  $g \in \Delta$  only if  $\mathcal{M}(g) \succeq_S lca(\mathcal{M}(g'), \mathcal{M}(g''))$ ,  
(c)  $g \in \Theta_A \cup \Theta_R$  if and only if either  $(g, g') \in \Xi$  or  $(g, g'') \in \Xi$ .  
(d) If  $g \in \Theta_A \cup \Theta_R$  and  $(g, g') \in \Xi$ , then  $\mathcal{M}(g)$  and  $\tau(g)$  must be incomparable, and  $\mathcal{M}(g')$  must be a descendant of  $\tau(g)$ , i.e.,  $\mathcal{M}(g') \leq_S \tau(g)$ .

**Replacing transfer constraint**

- (6) If  $g \in \Theta_A \cup \Theta_R$ , then  $g \in \Theta_R$  if and only if  $\mathcal{M}(\lambda(g)) = \tau(g)$ .

**Note:** This definition allows any invisible leaf node  $g$  (i.e.,  $g \in Le(G') \setminus Le(G)$ ) to map to a leaf node of  $S$ , say  $s \in Le(S)$ . However, gene  $g$  is not actually present in species  $s$  (otherwise it would not be invisible). Instead,  $\mathcal{M}(g) = s$  indicates that  $g$  existed in a predecessor species of  $s$  represented along the edge  $(pa(s), s) \in E(S)$ .

In the definition above, Constraint 1 specifies that the augmented gene tree,  $G'$ , must be consistent with the topology of the input gene tree  $G$ . Constraint 2 above ensures that the mapping  $\mathcal{M}$  is consistent with the leaf-mapping  $\mathcal{L}$ . Constraint 3a imposes on  $\mathcal{M}$  the temporal constraints implied by  $S$ , and Constraint 3b implies that any internal node in  $G'$  may represent at most one transfer event. Constraint 4 determines the edges of  $T$  that are transfer edges. Constraints 5a, 5b, and 5c state the conditions under which an internal node of  $G'$  may represent a speciation, duplication, and (additive or replacing) transfer respectively. Constraint 5d specifies which species may be designated as the recipient species for any given transfer event. Finally, constraint 6 specifies that a transfer event is labeled as a replacing transfer if and only if there exists a unique invisible leaf node in  $G'$  that represents the gene that is “replaced” by that replacing transfer. Note that constraints 2 through 5 are similar to those used in the DTL reconciliation model (e.g. [2]), except that they apply to the augmented gene tree  $G'$ , not to  $G$  as in DTL reconciliation, and take replacing transfers into account. Constraints 1 and 6 are unique to DTRL reconciliation.

While duplications, additive transfers, and replacing transfers are directly specified by any DTRL-reconciliation, losses are not. However, given a DTRL-reconciliation, the minimum number of losses implied by that reconciliation can be computed along the same lines as in the DTL reconciliation model [2], but with an adjustment to account for invisible lineages and replacing transfers. The adjustment is required to account for the implicit loss of a gene that occurs at each invisible leaf in the augmented gene tree

$G'$ . Some of these “losses” are due to replacing transfers, but those that are not must be explicitly counted as gene losses.

**DEFINITION 2.4 (LOSSES).** Given a DTRL-reconciliation  $\alpha = \langle \mathcal{L}, G', \mathcal{M}, \Sigma, \Delta, \Theta_A, \Theta_R, \Xi, \tau, \lambda \rangle$  for  $G$  and  $S$ , let  $g \in I(G')$  and  $\{g', g''\} = \text{Ch}(g)$ . The number of losses  $\text{Loss}_\alpha(g)$  at node  $g$ , is defined to be:

- $|\text{dist}_S(\mathcal{M}(g), \mathcal{M}(g')) - 1| + |\text{dist}_S(\mathcal{M}(g), \mathcal{M}(g'')) - 1|$ , if  $g \in \Sigma$ .
- $\text{dist}_S(\mathcal{M}(g), \mathcal{M}(g')) + \text{dist}_S(\mathcal{M}(g), \mathcal{M}(g''))$ , if  $g \in \Delta$ .
- $\text{dist}_S(\mathcal{M}(g), \mathcal{M}(g'')) + \text{dist}_S(\tau(g), \mathcal{M}(g'))$  if  $(g, g') \in \Xi$ .

The number of implicit losses at invisible leaves of  $G'$  (i.e., for the set  $\text{Le}(G') \setminus \text{Le}(G)$ ) is defined to be  $|\text{Le}(G') \setminus \text{Le}(G)| - |\Theta_R|$ .

The total number of losses in the DTRL-reconciliation  $\alpha$  is defined to be  $\text{Loss}_\alpha = |\text{Le}(G') \setminus \text{Le}(G)| - |\Theta_R| + \sum_{g \in I(G)} \text{Loss}_\alpha(g)$ .

In the DTRL reconciliation framework, each evolutionary event other than speciation is assigned a positive cost. Let  $P_\Delta, P_{\Theta_A}, P_{\Theta_R}$ , and  $P_{\text{Loss}}$  denote the gene duplication, additive transfer, replacing transfer, and gene loss costs, respectively. The reconciliation cost of a given DTRL-reconciliation is defined as follows.

**DEFINITION 2.5 (RECONCILIATION COST).** Given a DTRL-reconciliation  $\alpha = \langle \mathcal{L}, G', \mathcal{M}, \Sigma, \Delta, \Theta_A, \Theta_R, \Xi, \tau, \lambda \rangle$ , the reconciliation cost for  $\alpha$  is the total cost of all events invoked by  $\alpha$ . In other words, the reconciliation cost of  $\alpha$  is  $|\Delta| \times P_\Delta + |\Theta_A| \times P_{\Theta_A} + |\Theta_R| \times P_{\Theta_R} + \text{Loss}_\alpha \times P_{\text{Loss}}$ .

The goal is to find a DTRL-reconciliation that has minimum reconciliation cost. More formally:

**DEFINITION 2.6 (ODTRL PROBLEM).** Given  $G$  and  $S$ , along with  $P_\Delta, P_{\Theta_A}, P_{\Theta_R}$ , and  $P_{\text{Loss}}$ , the Optimal DTRL-Reconciliation Problem (ODTRL) problem is to find a DTRL-reconciliation for  $G$  and  $S$  with minimum reconciliation cost.

### 3 NP-HARDNESS OF ODTRL

We claim that the ODTRL problem is NP-hard and that the corresponding decision problem is NP-Complete. The decision version of the ODTRL problem is as follows:

**PROBLEM 1 (D-DTRL).**

**Instance:**  $G$  and  $S$ , along with event costs  $P_\Delta, P_{\Theta_A}, P_{\Theta_R}$ , and  $P_{\text{Loss}}$ , and a non-negative integer  $l$ .

**Question:** Does there exist a DTRL-reconciliation for  $G$  and  $S$  with reconciliation cost at most  $l$ ?

**THEOREM 3.1.** The D-DTRL problem is NP-Complete.

The D-DTRL problem is clearly in NP. In the remainder of this section we will show that the D-DTRL problem is NP-hard using a poly-time reduction from the decision version of the NP-hard *minimum rooted Subtree Prune and Regraft (rSPR) Distance* problem [7].

#### 3.1 Reduction from minimum rSPR distance

We begin by defining an rSPR operation and define the decision version of the minimum rSPR distance problem.

**DEFINITION 3.1 (rSPR OPERATION [7]).** Let  $T$  be a rooted binary tree and let  $e = \{u, v\}$  be an edge of  $T$  where  $u$  is the vertex that is in the path from the root of  $T$  to  $v$ . Let  $T'$  be the rooted binary tree obtained from  $T$  by deleting  $e$  and then adjoining a new edge

$f$  between  $v$  and the component  $C_u$  that contains  $u$  in one of the following two ways:

- Creating a new vertex  $u'$  which subdivides an edge in  $C_u$ , and adjoining  $f$  between  $u'$  and  $v$ . Then, either suppressing the degree-two vertex  $u$  or, if  $u$  is the root of  $T$ , deleting  $u$  and the edge incident with  $u$ , making the other end-vertex of this edge the new root.
- Creating a new root vertex  $u'$  and a new edge between  $u'$  and the original root. Then adjoining  $f$  between  $u'$  and  $v$  and suppressing the degree-two vertex  $u$ .

We say that  $T'$  has been obtained from  $T$  by a single rooted subtree prune and regraft (rSPR) operation.

**DEFINITION 3.2 (rSPR DISTANCE).** Given two trees  $T$  and  $T'$  with identical leaf sets, the rSPR distance between  $T$  and  $T'$ , denoted  $d_{\text{rSPR}}(T, T')$ , is defined to be the minimum number of rSPR operations required to transform  $T$  into  $T'$ .

The minimum rSPR distance problem is to find the rSPR distance between two trees. Its decision version can be stated as follows:

**PROBLEM 2 (D-rSPR PROBLEM).**

**Instance:** Two trees  $T$  and  $T'$  with identical leaf sets, and a non-negative integer  $k$ .

**Question:** Is  $d_{\text{rSPR}}(T, T') \leq k$ ?

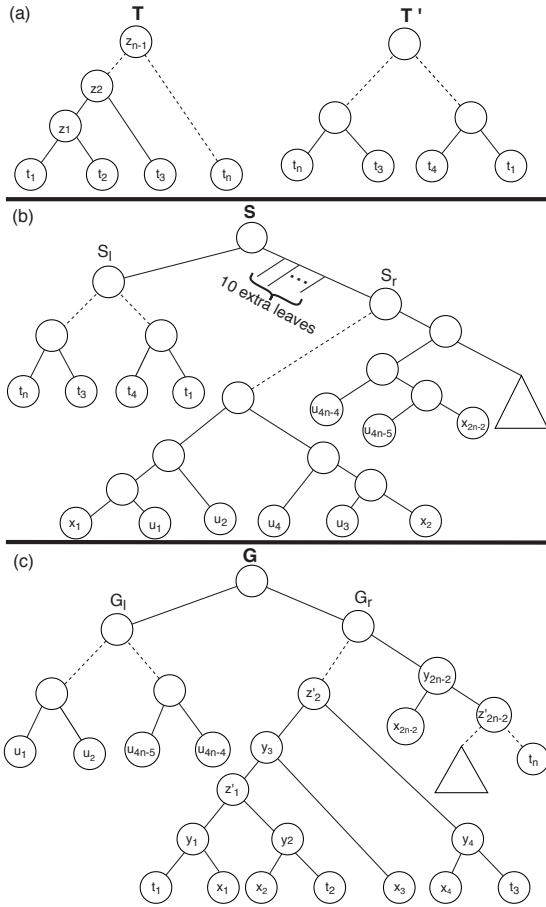
The D-rSPR problem is NP-Complete [7]. Consider any instance  $\rho$  of the D-rSPR problem with trees  $T$  and  $T'$  on the same leaf set of size  $n$  (i.e.,  $\text{Le}(T) = \text{Le}(T')$  and  $n = |\text{Le}(T)|$ ), and non-negative integer  $k$ . We will show how to transform  $\rho$  into an instance  $\delta$  of the D-DTRL problem by constructing  $G, S$ , and assigning the four event costs  $P_\Delta, P_{\Theta_R}, P_{\Theta_A}$ , and  $P_{\text{Loss}}$ , such that there exists a YES answer to the D-rSPR problem on  $\rho$  if and only if there exists a YES answer to the D-DTRL instance  $\delta$  with reconciliation cost at most  $l = 10n + 5k - 4$ .

#### 3.2 Gadget

We assume that the leaf set of  $T$  and  $T'$  is  $\{t_1, t_2, \dots, t_n\}$ . We also assume that the internal nodes of  $T$  are labeled  $\{z_1, z_2, \dots, z_{n-1}\}$ , as depicted in Figure 3(a). Next, we first show how to construct the species tree  $S$ , then the gene tree  $G$ , and then assign event costs. **Species tree.** The species tree  $S$ , is composed of two subtrees denoted  $S_l$  and  $S_r$  and ten extraneous leaf nodes (which are not represented in the gene tree). The root of subtree  $S_l$  is a child of  $rt(S)$ . The other subtree,  $S_r$ , is connected to  $rt(S)$  through a path to which the ten extraneous leaves are connected; these ten extraneous leaves ensure that no node of  $G$  maps to  $rt(S)$  in any optimal DTRL reconciliation. This is shown in Figure 3(b). The subtree  $S_l$  is identical to tree  $T'$ . Subtree  $S_r$  is a modified version of tree  $T$ , obtained as follows: We first perform a post-order traversal of tree  $T$  and number each node according to its position in the ordering, e.g. the left-most leaf node in  $T$  would be labeled with a 1, while  $rt(T)$  would be assigned the number  $2n - 1$ . Next, for each edge  $(pa(t), t) \in E(T)$ , if the number associated with  $t$  is  $i$ , we attach a subtree  $((x_i, u_{2i-1}), u_{2i})$  to edge  $(pa(t), t)$ . Thus,  $2n - 2$  subtrees are attached in all. Finally,

we delete all the original leaf nodes  $\{t_1, t_2, \dots, t_n\}$  from  $T$  and binarize the remaining tree by suppressing all non-root nodes of degree two. The resulting tree is  $S_r$ . This modification is depicted in Figure 3.

**Gene tree.** Gene tree  $G$  consists of two main subtrees, denoted  $G_l$  and  $G_r$ . Subtree  $G_l$  is obtained from species tree subtree  $S_r$  by removing all leaf nodes labeled with prefix  $x$  and then suppressing all non-root nodes of degree two. Subtree  $G_r$  is obtained by modifying  $T$  as follows: We consider again the post-order numbering of the nodes of  $T$  and, for each edge  $(pa(t), t) \in E(T)$ , if the number associated with  $t$  is  $i$ , we attach a leaf labeled  $x_i$  to edge  $(pa(t), t)$ . The new internal node created in attaching leaf  $x_i$  to the tree is denoted  $y_i$ . This construction is depicted in Figure 3(c)



**Figure 3: This figure illustrates the construction of species tree  $S$  (Part (b)) and gene tree  $G$  (Part (c)) for D-DTRL problem instance  $\delta$  based on trees  $T$  and  $T'$  (Part (a)) in the input instance  $\rho$  of the D-rSPR problem.**

Observe that each internal node of  $T$  has a corresponding node in  $G_r$ . We label these corresponding nodes of  $G_r$  as  $Z_G = \{z'_1, z'_2, \dots, z'_{n-1}\}$ , where node  $z'_i \in I(G_r)$  corresponds to node  $z_i \in I(T)$  for  $1 \leq i \leq n-1$ . We also define the following three subsets of  $V(G)$ :  $Y_G = \{y_1, y_2, \dots, y_{2n-2}\}$ ,  $X_G = \{x_1, x_2, \dots, x_{2n-2}\}$ , and  $T_G = \{t_1, \dots, t_n\}$ . Note that  $I(G_r) = Y_G \cup Z_G$ .

**Event costs.** Event costs are assigned as follows:  $P_\Delta = 4$ ,  $P_{\Theta_A} = 6$ ,  $P_{\Theta_R} = 5$ , and  $P_{Loss} = 3$ .

This completes our construction of instance  $\delta$  of the D-DTRL problem. Note that  $G$  and  $S$  can be both constructed in time polynomial in  $n = |Le(T)|$ .

**CLAIM 1.** *There exists a YES answer to the D-rSPR problem on  $\rho$  if and only if there exists a YES answer to the D-DTRL instance  $\delta$  with reconciliation cost  $l \leq 10n + 5k - 4$ .*

For brevity, a proof of this claim is deferred to the full version of this manuscript. The main idea behind this reduction can be explained briefly as follows. Each rSPR operation on instance  $\rho$  corresponds to exactly one replacing transfer event on gene tree  $G$  from instance  $\delta$ . Based on the structure of gene tree  $G$  and species tree  $S$ , we will be able to show that for each rSPR operation there is at least one way to get a valid corresponding replacing transfer.

The correctness of Theorem 3.1 follows immediately from Claim 1.

## 4 EXPERIMENTAL ANALYSIS

There do not currently exist any algorithms or heuristics to compute DTRL reconciliations, and it is not even known how algorithms for computing optimal DTL reconciliations perform when confronted with gene trees that have been affected by both additive and replacing transfers. Therefore, we first focused on answering two fundamental questions: (i) How is the accuracy of DTL reconciliation affected by the presence of replacing horizontal gene transfers? (ii) How well does DTL reconciliation perform at inferring replacing transfer events?

To answer these questions, we used the recently developed simulation framework SaGePhy [21] to stochastically evolve gene trees inside a given species tree under a model that allows for gene duplications, additive transfers, replacing transfers, and gene losses. Using this simulation framework we created a large number of gene trees with varying rates of evolutionary events, computed optimal DTL reconciliations for the gene/species tree pairs, and evaluated the accuracy of the inferred reconciliations by comparing them to the true evolutionary histories of those gene trees. To compute optimal DTL reconciliations we employed the widely-used RANGER-DTL [2, 3] software package.

**Simulated datasets.** We used SaGePhy [21] to generate 100 species trees, each containing 100 leaves and of height 1, under a birth-death process. Next, inside each of the species trees, we generated three different gene trees using low, medium, and high rates of duplication, additive transfer, replacing transfer, and loss events, resulting in three sets of 100 gene trees. To generate the low DTRL gene trees, we used duplication, additive transfer, replacing transfer, and loss rates of 0.133, 0.133, 0.133, and 0.266, respectively; for the medium DTRL gene trees we used rates of 0.3, 0.3, 0.3, and 0.6, respectively; and for the high DTRL gene trees we used rates of 0.6, 0.6, 0.6, and 1.2, respectively. Thus, the total transfer rate was twice the duplication rate, with an equal rate of additive and replacing transfers, and the loss rate was assigned to be equal to the sum of the duplication and additive transfer rates. These duplication, transfer, and loss rates are based on rates observed in real data and capture both datasets with lower rates of these events and datasets with a very high rate of these events [4].

For the low DTRL gene trees, the average gene tree leaf set size was 96.11, with an average of 2.37 additive transfers, 2.65 replacing transfers, and 2.19 duplication events per gene tree. For the medium DTRL gene trees, the average gene tree leaf set size was 94.75, with an average of 5.09 additive transfers, 5.01 replacing transfers, and 5.00 duplication events per gene tree. For the high DTRL gene trees, the average gene tree leaf set size was 110.22, with an average of 9.52 additive transfer events, 9.42 replacing transfer events, and 10.39 duplication events per gene tree.

#### 4.1 Impact of replacing transfers on DTL reconciliation

We evaluated the accuracy of DTL reconciliation in inferring the evolutionary event and species tree mapping for each internal node in the simulated gene trees. We computed a single optimal reconciliation for each gene tree using RANGER-DTL 2.0 [3] with default parameters (i.e., transfer cost of 3, duplication cost of 2, and loss cost of 1) and compared the computed reconciliation against the true evolutionary history of that gene tree. We observed very high accuracy for inferring the correct event type (speciation, duplication, or transfer) at each gene tree node. For instance, for the low DTRL gene trees, 99.67%, 96.35% and 96.22% of the gene tree nodes labeled as speciation, duplication, and transfer, respectively, in the computed reconciliations were inferred correctly. Even for the high DTRL gene trees, these percentages remained very high at 95.69%, 87.49%, and 95.25%, respectively. These results are shown in Figure 4(a).

Looking at the accuracy of mapping inference, we found that 99.09%, 97.11%, and 92.15% of all internal nodes were assigned the correct species node mapping for the low, medium, and high DTRL gene trees, respectively. Detailed results are shown in Figure 4(b).

We compared these results for event and mapping accuracy with results obtained on gene trees simulated with the same overall rates of duplication, transfer, and loss events but in which all simulated transfers were additive transfers (no replacing transfers). We found that the numbers were nearly identical, showing that the presence of replacing transfers does not negatively affect the accuracy of DTL reconciliation itself. For example, for the high DTRL gene trees, the percentage of speciation, duplication, and transfer nodes assigned the correct event type was 95%, 81%, and 95%, respectively, and 91% of all nodes were assigned the correct mapping. Note, however, that DTL reconciliation cannot distinguish between additive and replacing transfers, and both types of transfer events are simply inferred as “transfers”.

**Accuracy of inferring replacing transfers.** Next, we performed additional analysis to study if there was any discrepancy in the accuracies of inferring the correct event type (transfer) or mapping for additive transfers and those for replacing transfers. For the low DTRL gene trees, we found that additive transfers were assigned the correct event type 97.05% of the time and the correct mapping 89.45% of the time, while for replacing transfers these numbers were 95.47% and 85.28%, respectively. Likewise, for the medium DTRL gene trees, additive transfers were assigned the correct event type 95.87% of the time and the correct mapping 87.03% of the time, while for replacing transfers these numbers were 93.01% and 81.04%, respectively. For high DTRL gene trees,

these numbers were 95.38% and 75.53% for the additive transfers and 95.12% and 74.52% for the replacing transfers. Overall, this shows that replacing transfers are inferred and mapped with accuracy comparable to that of additive transfers. These results are shown in Parts (c) and (d) of Figure 4.

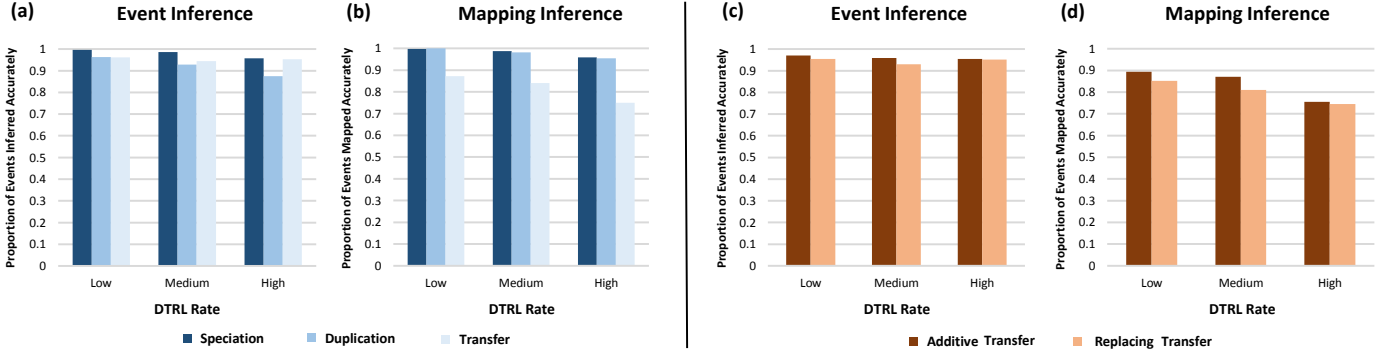
These results are highly significant and suggest that, to design an effective heuristic for DTRL reconciliation, it may suffice to first use DTL reconciliation to identify transfer events and then classify those transfer events as being either replacing or additive.

## 5 A HEURISTIC FOR CLASSIFYING TRANSFERS

To explore the feasibility of accurately classifying transfer events inferred through DTL reconciliation, we designed a simple heuristic for classifying inferred transfers and tested its accuracy on several simulated datasets. Given gene tree  $G$  and species tree  $S$ , our heuristic first computes an optimal DTL reconciliation, initially classifies all inferred transfer events as additive, and then greedily attempts to reclassify some of these transfer events as replacing. To determine if a transfer can be replacing, the heuristic checks if the resulting loss of a gene lineage in the recipient species will make it impossible to generate at least as many gene copies at each leaf descendant of the recipient species as are actually present. If it does then the transfer remains an additive transfer but otherwise is reclassified as a replacing transfer. This heuristic thus depends only on the actual counts (based on  $G$  and  $S$ ) and implied/inferred counts (based on the computed reconciliation) of genes at each leaf of the species tree.

More precisely, the heuristic works as follows:

- (1) Calculate the number of gene copies from  $Le(G)$  that are present in each extant species represented in the species tree. For a species  $s \in Le(S)$ , this count is represented by *actual-count*( $s$ ). Note that this is simply the number of leaf nodes of  $G$  that map to leaf  $s$ .
- (2) Compute an optimal DTL reconciliation for  $G$  and  $S$  (using RANGER-DTL 2.0 [3] with default parameters).
- (3) Classify each inferred transfer event as an additive transfer.
- (4) Based on the current reconciliation, compute the number of gene copies that would occur in each extant species if there were no gene losses. This can be counted easily as follows: Consider the path from the root of the species tree to the species (leaf) under consideration. Count the number of gene duplication nodes on the gene tree that map to a node on this path; let this number be denoted  $n_1$ . Count the number of additive transfer events on the gene tree whose recipient is a node on this path; let this number be denoted  $n_2$ . Determine if the root of the gene tree maps to a node on this path; if so, assign  $n_3 = 1$ , otherwise  $n_3 = 0$ . The final required count is simply  $n_1 + n_2 + n_3$ . For a species  $s \in Le(S)$ , this final count is represented by *inferred-count*( $s$ ).
- (5) For each node  $g$  in a pre-order traversal of  $G$ :
  - (a) If  $g$  is a transfer event:
    - (i) Let  $x \in V(S)$  denote the recipient species for that transfer event.



**Figure 4: Accuracy of DTL reconciliation in the presence of replacing transfers.** Part (a) shows the fraction of internal nodes across all low DTRL, medium, DTRL, and high DTRL gene trees, whose event types, speciation, duplication, or transfer, are inferred correctly through DTL reconciliation. Part (b) shows the corresponding fractions for correct mapping inference. Part (c) shows the fraction of additive transfer nodes and replacing transfer nodes across all low DTRL, medium, DTRL, and high DTRL gene trees, that are correctly inferred as transfer events by DTL reconciliation. Part (d) shows the corresponding fractions for correct mapping inference. For each DTRL rate, results are averaged across 100 datasets.

- (ii) Check if  $inferred-count(s) > actual-count(s)$  for each  $s \in Le(S(x))$ . If yes, reclassify  $g$  as a replacing transfer and reduce  $inferred-count(s)$  by 1 for each  $s \in Le(S(x))$ .
- (6) Output the resulting classification of inferred transfer events.

An implementation of the heuristic algorithm is freely available open-source as part of the RANGER-DTL software package: <https://compbio.engr.uconn.edu/software/ranger-dtl/>. Next, we illustrate this algorithm through an example.


**Illustration of the heuristic.** Consider the gene tree and species tree shown in Figure 5. Suppose the inferred reconciliation of the gene tree and species tree labels gene node  $g_3$  as a transfer event mapping to species node  $s_4$  and with recipient species  $s_6$ , gene node  $g_6$  as a gene duplication mapping to species  $D$ , gene node  $g_7$  as a transfer event mapping to species node  $G$  and with recipient species  $A$ , and all other nodes as speciations with the root of the species tree. The heuristic starts with this reconciliation and its task is to assign each of the two transfer nodes at  $g_3$  and  $g_7$  to be additive or replacing.

The *actual-counts* at the leaves  $A, B, C, D, E, F$ , and  $G$  of the species tree are easily computed to be, respectively: 2, 1, 0, 2, 1, 1, 1.

The heuristic starts by assuming that both transfer events are additive and that there are no losses, and computes the initial values of the *inferred-counts* at leaves  $A, B, C, D, E, F$ , and  $G$  of the species tree, respectively, as follows: 2, 1, 1, 2, 2, 2, 1.

Note that each of these *inferred-counts* is at least equal to the corresponding *actual-count*.

Next, the heuristic considers all transfer nodes on the gene tree one at a time in pre-order. Thus it first considers node  $g_3$ . It checks to see if the additive transfer at  $g_3$  could have been a replacing transfer instead. To make this determination, it checks if the *inferred-count* at each leaf descended from the recipient of this transfer event (i.e., all leaves descended from node  $s_6$ ) is strictly greater than its corresponding *actual-count*. In this example, we would



**Figure 5: Gene tree and species tree for illustrating heuristic.**

check to make sure that  $inferred-count(E) > actual-count(E)$  and  $inferred-count(F) > actual-count(F)$ . Both inequalities hold true in this case, and so the transfer at node  $g_3$  is labeled as a replacing transfer. The next step is to update the *inferred-counts* to account for this change. The new *inferred-counts* at the leaves  $A, B, C, D, E, F$ , and  $G$  of the species tree, respectively, are thus: 2, 1, 1, 2, 1, 1, 1.

Continuing the pre-order traversal the heuristic then considers the transfer at node  $g_7$ , and again checks to see if this additive transfer event could have been a replacing transfer. To make this determination, it checks if  $inferred-count(A) > actual-count(A)$ . This inequality does not hold since both counts are 2. Thus, the heuristic labels the transfer at  $g_7$  as an additive transfer.

## 5.1 Experimental Results

To evaluate the ability of this heuristic to classify transfer events accurately we applied it to several simulated datasets covering a wide range of evolutionary scenarios. We divide these datasets into three groups: Group 1 consists of the datasets described in Section 4 (consisting of the three sets of low, medium, and high DTRL trees) and we refer to these as *mixed datasets* since gene trees in these datasets contain both additive and replacing transfers. Group 2 consists of datasets in which all transfers are additive transfers.



**Table 1: Classification accuracy of the heuristic.** This table shows the results of applying the heuristic algorithm to the three groups of simulated datasets; Group 1 where gene trees contain both additive and replacing transfers, Group 2 where gene trees contain only additive transfers, and Group 3 where gene trees contain only replacing transfers. Each group is further divided into low, medium, and high DTRL datasets. For each dataset in each group, we measure (i) the percentage of true transfers that are inferred as transfers by the heuristic (“All transfers accuracy”), (ii) the percentage of true additive transfers that are inferred as additive transfers by the heuristic (“Additive transfers accuracy”), and (iii) the percentage of true replacing transfers that are inferred as replacing transfers by the heuristic (“Replacing transfers accuracy”). All numbers are averaged over the 100 gene tree/species tree pairs in each dataset. Note that the numbers reported for “All transfers accuracy” for Group 1 datasets are slightly different from those reported in Section 4. This is because there often exist multiple optimal reconciliations and each run of RANGER-DTL or the heuristic samples one of these optimal reconciliations at random.

	Mixed datasets (Group 1)			Additive datasets (Group 2)			Replacing datasets (Group 3)		
	Low DTRL	Med. DTRL	High DTRL	Low DTRL	Med. DTRL	High DTRL	Low DTRL	Med. DTRL	High DTRL
All transfers accuracy (%)	95.2	96.2	94.3	97.8	96.6	95.8	96.1	96.4	94.6
Additive transfers accuracy (%)	84.9	65.5	54.4	80.6	59.7	52.2	–	–	–
Replacing transfers accuracy (%)	86.1	80.7	70.1	–	–	–	93.9	88.1	75.0

As before, this group is composed of three sets of low, medium, and high DTRL trees. Finally, group 3 consists of datasets in which all transfers are replacing transfers, divided as before into three sets of low, medium, and high DTRL trees. For group 2 and group 3 datasets, the duplication, transfer, and loss rates used to generate the low, medium, and high DTRL trees are identical to those used for group 1 (described in detail in Section 4), except that in group 2 all transfers are additive and in group 3 all transfers are replacing. Thus, the total number of duplication and transfer events are roughly the same across the three groups.

To evaluate the classification accuracy of our heuristic, we measured the following for each dataset from each of the three groups: (1) What fraction of all transfer events in the true evolutionary history of a gene tree are correctly inferred as transfer events by the heuristic. (2) What fraction of additive transfer events in the true evolutionary history of a gene tree are correctly inferred as additive transfers by the heuristic. And (3) what fraction of all replacing transfer events in the true evolutionary history of a gene tree are correctly inferred as replacing transfers by the heuristic. Table 1 shows these results. As also seen in Section 4, transfer events can be identified with high accuracy across all three groups and all three DTRL rates. Results are more variable for classification of the inferred transfers as being either additive or replacing (which is the primary task of the heuristic). In general, over 80% of the additive transfers and 86% of the replacing transfers are classified correctly for the low DTRL datasets across the three groups, 60%–65% of additive and 80%–88% of replacing transfers are classified correctly for the medium DTRL datasets across the three groups, and approximately 52% of additive and 70% of replacing transfers classified correctly for the high DTRL datasets. These results show that the proposed heuristic is quite accurate at classification when DTRL rates are low but performance suffers as the rates increase. These results also show that, in general, the heuristic infers replacing transfers with greater accuracy than additive transfers. This is not entirely surprising given that the heuristic attempts to label as many of the transfers as replacing as possible; in particular, additive transfers may appear to be replacing due to the high loss rate in our datasets. As a result, the false negative rate for replacing

transfers is low but the false positive rate is high (since many additive transfers may be classified as replacing), while the opposite is true for additive transfers.

Overall, these experimental results demonstrate that there is often sufficient information in DTL reconciliations to be able to distinguish between additive and replacing transfers, and suggest that classification of transfer events inferred through DTL reconciliation is a promising approach for estimating optimal DTRL reconciliations. While our current heuristic is simple and has limited classification accuracy, our experimental results do also suggest that, in general, the ability to distinguish between replacing and additive transfers based purely on gene tree (and species tree) topology diminishes rapidly as the rate of evolutionary events increases. This is not surprising since evolutionary events that occur after an additive or replacing transfer can completely erase the phylogenetic (i.e., topological) signature of that additive or replacing transfer. Nonetheless, we expect more advanced heuristics to be more effective at distinguishing between additive and replacing transfers even for high rates of evolutionary events.

## 6 CONCLUSION

Accurate detection of both replacing and additive transfer events is crucial for understanding horizontal gene transfer in microbes and understanding microbial evolution in general. In this work, we address this problem by formalizing and experimentally studying the DTRL reconciliation framework that simultaneously models gene duplication, loss, and both additive and replacing transfer. Our framework builds upon the traditional DTL reconciliation model and extends it substantially to properly model replacing transfers. We prove that the underlying computational problem is NP-hard, and our proof establishes a close relationship between the rSPR distance problem and DTRL reconciliation. Our experimental results show that DTL reconciliation, which assumes all transfers are additive, is surprisingly robust to the presence of replacing transfer, and suggest that it should be possible to design effective heuristics for the DTRL reconciliation problem based on DTL reconciliation. To explore the feasibility of such an approach, we devised a simple heuristic to classify inferred transfer events as

being either additive or replacing and found that it achieves fairly good classification accuracy for low and medium rates of evolutionary events. This demonstrates the feasibility of estimating optimal DTRL reconciliations based on optimal DTL reconciliations followed by classification of inferred transfer events, and we expect improved heuristics to achieve greater classification accuracy. Our current heuristic has several limitations, of which the following two are particularly notable: First, it does not directly solve the DTRL reconciliation problem and its accuracy is therefore limited by the accuracy of the inferred DTL reconciliation (which does not model hidden events). Second, it ignores the presence of multiple optimal DTRL (or even DTL) reconciliations. Addressing these limitations may yield improved heuristics for DTRL reconciliation.

Our experimental results also suggest that, as expected, the ability to distinguish between replacing and additive transfers based purely on phylogenetic incongruence diminishes rapidly as the rate of evolutionary events increases, and therefore alternative approaches may be needed for such cases. One such alternative approach for estimating optimal DTRL reconciliations is to make use of available gene order information for the extant species in the analysis to classify each transfer event inferred through DTL reconciliation as being either additive or replacing based on genomic context. However, the applicability of such an approach is limited since it requires the use of complete genomic information and, due to genome rearrangements, can only be used for relatively closely related sets of species. A hybrid approach that uses both gene ordering information and phylogenetic incongruence may help overcome the limitations of the two separate approaches, and developing this hybrid approach is a promising research direction.

Finally, it would be useful to develop exact algorithms for the DTRL reconciliation problem. Even though we showed the problem to be NP-hard, it may be possible to design fixed parameter algorithms that can be efficiently applied to gene trees with small reconciliation cost (see. e.g., [14]), or to design effective branch and bound algorithms to rapidly compute optimal DTRL reconciliations for small gene trees.

**Acknowledgements:** The authors thank Abhijit Mondal for pointing out an error in an initial implementation of the heuristic.

**Funding:** This work was supported in part by NSF awards IIS 1553421, MCB 1616514, and EAR 1615573 to MSB.

## REFERENCES

- [1] Sophie S. Abby, Eric Tannier, Manolo Gouy, and Vincent Daubin. 2010. Detecting lateral gene transfers by statistical reconciliation of phylogenetic forests. *BMC Bioinformatics* 11, 1 (15 Jun 2010), 324.
- [2] Mukul S. Bansal, Eric J. Alm, and Manolis Kellis. 2012. Efficient algorithms for the reconciliation problem with gene duplication, horizontal transfer and loss. *Bioinformatics* 28, 12 (2012), 283–291.
- [3] Mukul S. Bansal, Manolis Kellis, Misagh Kordi, and Soumya Kundu. 2018. RANGER-DTL 2.0: rigorous reconstruction of gene-family evolution by duplication, transfer and loss. *Bioinformatics* 34, 18 (2018), 3214–3216. <https://doi.org/10.1093/bioinformatics/bty314>
- [4] Mukul S. Bansal, Yi-Chieh Wu, Eric J. Alm, and Manolis Kellis. 2015. Improved gene tree error correction in the presence of horizontal gene transfer. *Bioinformatics* 31, 8 (2015), 1211–1218. <https://doi.org/10.1093/bioinformatics/btu806>
- [5] Robert G. Beiko, Timothy J. Harlow, and Mark A. Ragan. 2005. Highways of gene sharing in prokaryotes. *Proceedings of the National Academy of Sciences of the United States of America* 102, 40 (2005), 14332–14337.
- [6] Alix Boc, Herve Philippe, and Vladimir Makarenkov. 2010. Inferring and Validating Horizontal Gene Transfer Events Using Bipartition Dissimilarity. *Syst. Biol.* 59, 2 (2010), 195–211. <https://doi.org/10.1093/sysbio/syp103>
- [7] M Bordewich and C Semple. 2005. On the computational complexity of the rooted subtree prune and regraft distance. *Annals of combinatorics* 8, 4 (2005), 409–423. <https://doi.org/10.1007/s0026-004-0229-z>
- [8] Zhi-Zhong Chen, Fei Deng, and Lusheng Wang. 2012. Simultaneous Identification of Duplications, Losses, and Lateral Gene Transfers. *IEEE/ACM Trans. Comput. Biology Bioinform.* 9, 5 (2012), 1515–1528.
- [9] Sang Chul Choi, Matthew D. Rasmussen, Melissa J. Hubisz, Ilan Gronau, and Michael J. Stanhope. 2012. Replacing and Additive Horizontal Gene Transfer in *Streptococcus*. *Molecular Biology and Evolution* 29 (2012), 3309–3320. <https://doi.org/10.1093/molbev/mss138>
- [10] Lawrence A. David and Eric J. Alm. 2011. Rapid evolutionary innovation during an Archaean genetic expansion. *Nature* 469 (2011), 93–96.
- [11] Jean-Philippe Doyon, Celine Scornavacca, K. Yu. Gorbunov, Gergely J. Szollosi, Vincent Ranwez, and Vincent Berry. 2010. An Efficient Algorithm for Gene/Species Trees Parsimonious Reconciliation with Losses, Duplications and Transfers. In *RECOMB-CG (Lecture Notes in Computer Science)*, Eric Tannier (Ed.), Vol. 6398. Springer, 93–108.
- [12] Robert Friedmann and Bert Ely. 2012. Codon Usage Methods for Horizontal Gene Transfer Detection Generate an Abundance of False Positive and False Negative Results. *Current Microbiology* 65, 5 (2012), 639–642.
- [13] Michael T. Hallett and Jens Lagergren. 2001. Efficient algorithms for lateral gene transfer problems. In *RECOMB*, 149–156.
- [14] Damir Hasić and Eric Tannier. 2019. Gene tree reconciliation including transfers with replacement is NP-hard and FPT. *Journal of Combinatorial Optimization* (22 Feb 2019). <https://doi.org/10.1007/s10878-019-00396-z>
- [15] Tobias Hill, Karl Nordstrom, Mikael Thollesson, Tommy Safstrom, Andreas Vernersson, Robert Fredriksson, and Helgi Skioth. 2010. SPRIT: Identifying horizontal gene transfer in rooted phylogenetic trees. *BMC Evolutionary Biology* 10, 1 (2010), 42. <https://doi.org/10.1186/1471-2148-10-42>
- [16] Edwin Jacox, Cedric Chauve, Gergely J. Szollosi, Yann Ponty, and Celine Scornavacca. 2016. ecceTERA: comprehensive gene tree-species tree reconciliation using parsimony. *Bioinformatics* 32, 13 (2016), 2056. <https://doi.org/10.1093/bioinformatics/btw105>
- [17] Guohua Jin, Luay Nakhleh, Sagi Snir, and Tamir Tuller. 2009. Parsimony Score of Phylogenetic Networks: Hardness Results and a Linear-Time Heuristic. *IEEE/ACM Trans. Comput. Biology Bioinform.* 6, 3 (2009), 495–505.
- [18] Slimane Khayi, Pauline Blin, Jacques Pedron, Teik-Min Chong, Kok-Gan Chan, Mohieddine Moumni, Valerie Helias, Frederique Van Gijsegem, and Denis Faure. 2015. Population genomics reveals additive and replacing horizontal gene transfers in the emerging pathogen *Dickeya solani*. *BMC Genomics* 16, 1 (2015), 788.
- [19] Eugene V. Koonin, Kira S. Makarova, and L. Aravind. 2001. Horizontal Gene Transfer in Prokaryotes: Quantification and Classification. *Annual Review of Microbiology* 55, 1 (2001), 709–742. PMID: 11544372.
- [20] Misagh Kordi and Mukul S. Bansal. In press. Exact Algorithms for Duplication-Transfer-Loss Reconciliation with Non-Binary Gene Trees. *IEEE/ACM Trans. Comput. Biology Bioinform.* (In press).
- [21] Soumya Kundu and Mukul S. Bansal. 2019. SaGePhy: An improved phylogenetic simulation framework for gene and subgene evolution. *Bioinformatics* (2019).
- [22] Daniel Merkle, Martin Middendorf, and Nicolas Wieseke. 2010. A parameter-adaptive dynamic programming approach for inferring cophylogenies. *BMC Bioinformatics* 11, Suppl 1 (2010), S60.
- [23] Luay Nakhleh, Derek A. Ruths, and Li-San Wang. 2005. RIATA-HGT: A Fast and Accurate Heuristic for Reconstructing Horizontal Gene Transfer. In *COCOON (Lecture Notes in Computer Science)*, Lusheng Wang (Ed.), Vol. 3595. Springer, 84–93.
- [24] Joel Sjostrand, Ali Tofigh, Vincent Daubin, Lars Arvestad, Bengt Sennblad, and Jens Lagergren. 2014. A Bayesian Method for Analyzing Lateral Gene Transfer. *Systematic Biology* 63, 3 (2014), 409–420. <https://doi.org/10.1093/sysbio/syu007>
- [25] Maureen Stolzer, Han Lai, Minli Xu, Deepa Sathaye, Benjamin Vernot, and Dannie Durand. 2012. Inferring duplications, losses, transfers and incomplete lineage sorting with nonbinary species trees. *Bioinformatics* 28, 18 (2012), 409–415.
- [26] Gergely J. Szollosi, Bastien Boussau, Sophie S. Abby, Eric Tannier, and Vincent Daubin. 2012. Phylogenetic modeling of lateral gene transfer reconstructs the pattern and relative timing of speciations. *Proceedings of the National Academy of Sciences* 109, 43 (2012), 17513–17518. <https://doi.org/10.1073/pnas.1202997109>
- [27] Gergely J. Szollosi, Eric Tannier, Nicolas Lartillot, and Vincent Daubin. 2013. Lateral Gene Transfer from the Dead. *Systematic Biology* 62, 3 (2013), 386–397. <https://doi.org/10.1093/sysbio/syt003>
- [28] Cuong Than, Derek A. Ruths, Hideki Inman, and Luay Nakhleh. 2007. Confounding Factors in HGT Detection: Statistical Error, Coalescent Effects, and Multiple Solutions. *J. Comput. Biol.* 14, 4 (2007), 517–535.
- [29] Ali Tofigh, Michael T. Hallett, and Jens Lagergren. 2011. Simultaneous Identification of Duplications and Lateral Gene Transfers. *IEEE/ACM Trans. Comput. Biology Bioinform.* 8, 2 (2011), 517–535.
- [30] Olga Zhaxybayeva. 2009. *Horizontal gene transfer: Genomes in flux*. Methods in Molecular Biology, Vol. 532. Humana Press, Chapter Detection and quantitative assessment of horizontal gene transfer, 195–213.