

The agreement distance of unrooted phylogenetic networks

Jonathan Klawitter

University of Würzburg, Germany

received 24th Aug. 2019, revised 17th Jan. 2020, 22nd May 2020, accepted 2nd June 2020.

A rearrangement operation makes a small graph-theoretical change to a phylogenetic network to transform it into another one. For unrooted phylogenetic trees and networks, popular rearrangement operations are tree bisection and reconnection (TBR) and prune and regraft (PR) (called subtree prune and regraft (SPR) on trees). Each of these operations induces a metric on the sets of phylogenetic trees and networks. The TBR-distance between two unrooted phylogenetic trees T and T' can be characterised by a maximum agreement forest, that is, a forest with a minimum number of components that covers both T and T' in a certain way. This characterisation has facilitated the development of fixed-parameter tractable algorithms and approximation algorithms. Here, we introduce maximum agreement graphs as a generalisations of maximum agreement forests for phylogenetic networks. While the agreement distance – the metric induced by maximum agreement graphs – does not characterise the TBR-distance of two networks, we show that it still provides constant-factor bounds on the TBR-distance. We find similar results for PR in terms of maximum endpoint agreement graphs.

Keywords: phylogenetic network, rearrangement operation, agreement distance, maximum agreement forest

1 Introduction

Phylogenetic trees and networks are graphs where the leaves are labelled bijectively by a set of taxa, for example a set of organisms, species, or languages (Semple and Steel, 2003; Dunn, 2014). They are used to model and visualise evolutionary relationships. While a phylogenetic tree is suited only for tree-like evolutionary histories, a phylogenetic network can also be used for taxa whose past includes reticulate events like hybridisation, horizontal gene transfer, recombination, or reassortment (Semple and Steel, 2003; Huson et al., 2010; Steel, 2016). Such reticulate events arise in all domains of life (Thomas and Nielsen, 2005; Rieseberg and Willis, 2007; Meier et al., 2017; Wagner et al., 2017). There is a distinction between rooted and unrooted phylogenetic networks. More precisely, in a rooted phylogenetic network the edges are directed from a designated root towards the leaves, thus modelling evolution along the passing of time. On the other hand, the edges of an unrooted phylogenetic network are undirected and the network thus represents the evolutionary relatedness of the taxa. In some cases, unrooted phylogenetic networks can be thought of as rooted phylogenetic networks in which the orientation of the edges has been disregarded. Janssen et al. (2018); Francis et al. (2018a); Huber et al. (2019) call such unrooted

phylogenetic networks proper. Here we focus on unrooted, binary, proper phylogenetic networks where binary means that all vertices except for the leaves have degree three.

A rearrangement operation makes a small graph-theoretical change to transform a phylogenetic network into another one. Since this induces neighbourhoods, rearrangement operations structure the set of phylogenetic networks on the same taxa into a space. Because of this property, they are used by several phylogenetic inference methods that traverse this space (Page, 1993; Bouckaert et al., 2014; Ronquist and Huelsenbeck, 2003; Guindon et al., 2010; Yu et al., 2013, 2014; Whidden and Matsen, 2015). Furthermore, the minimum number of rearrangement operations needed to transform one network into another induces a metric. This allows the comparison of results obtained for different data or from different inference methods, for instance, to evaluate their robustness or to find outliers or clusters.




Fig. 1: An SPR prunes the edge e in T_1 and regrafts it to the edge incident to leaf 1 to obtain T_2 . A TBR moves the edge e' to obtain T_3 from T_2 .

On unrooted phylogenetic trees, two popular rearrangement operations are subtree prune and regraft (SPR), which cuts (prunes) an edge at one side and then reattaches it, and tree bisection and reconnection (TBR), which removes an edge and then reconnects the two resulting smaller trees (Swofford et al., 1996). These are illustrated in Figure 1. Computing the SPR- and TBR-distance of two unrooted phylogenetic trees T and T' is NP-hard (Allen and Steel, 2001; Hickey et al., 2008). On the positive side, the TBR-distance of T and T' is characterised by a maximum agreement forest (MAF) of T and T' , which is a forest of smaller phylogenetic trees on which T and T' agree upon and that among all such forests has the minimum number of components (Allen and Steel, 2001). The idea is that a MAF captures all parts that remain unchanged by a shortest TBR-sequence that transforms T into T' . Figure 2 shows a maximum agreement forest F for T_1 and T_3 of Figure 1. Furthermore, a MAF F together with the edges that got moved by the TBR-sequence can be embedded into T and T' such that all edges are covered; see again Figure 2. Compared to a sequence of trees that describe a TBR-distance, MAFs provide a single structure and have therefore been utilised for NP-hardness proofs (Allen and Steel, 2001; Hickey et al., 2008), for fixed-parameter tractable and approximation algorithms (Allen and Steel, 2001; Hallett and McCartin, 2007; Rodrigues et al., 2007; Whidden and Zeh, 2009; Chen et al., 2015). So far, no characterisation of the SPR-distance in terms of such a structure has been found and Whidden and Matsen (2019) argue why such a characterisation is unlikely. In particular, they showed that an edge might be pruned twice and that common clusters (subtrees on a subset of the leaves) are not always maintained. However, Whidden and Matsen introduced maximum endpoint agreement forests (MEAF) (precisely defined in Section 4) as a variation of MAFs that bound the SPR-distance of two trees.

SPR and TBR on trees have recently been generalised to phylogenetic networks with the operations prune and regraft (PR) and TBR (Francis et al., 2018b; Janssen and Klawitter, 2019). In principle, these operations work the same on networks as on trees. A PR operation first prunes an edge at one side and




Fig. 2: The graph F is a maximum agreement forest for T_1 and T_3 of Figure 1. Together with the edges moved by the operations, F can be embedded into T_1 and T_3 such that all edges are covered.

then reattaches it at another edge; a TBR operation on a network may also first remove an edge and then add a new edge like a TBR operation on a tree. This is illustrated in Figure 3. However, both PR and TBR may also remove or add an edge to change the size of the network (see Figure 5). Janssen and Klavitter (2019) studied several properties of spaces of networks under PR and TBR and, among other results, showed that computing the TBR-distance of two networks is NP-hard.




Fig. 3: A PR prunes the edge $\{u, v\}$ at u in N_1 and regrafts it to the edge incident to leaf 3 to obtain N_2 . A TBR moves the edge $\{u, v\}$ to obtain N_3 from N_2 . The graph G is a maximum agreement graph for N_1 and N_3 shown with an embedding into N_3 .

Similar to the TBR-distance of unrooted phylogenetic trees, the SPR-distance of two rooted phylogenetic trees can be characterised by a rooted version of MAFs (Bordewich and Semple, 2005). This has again facilitated the development of fixed-parameter and approximation algorithms (Bordewich et al., 2008; Wu, 2009; Bonet and St. John, 2009; Whidden et al., 2013; Bordewich et al., 2017b). Prune and regraft (PR) and subnet prune and regraft (SNPR) are extensions of SPR for rooted phylogenetic networks (Bordewich et al., 2017a; Klavitter, 2019). Recently, we generalised MAFs to maximum agreement graphs (MAG) for networks. Similar to a MAF, the idea of a MAG is that its components model those parts of two phylogenetic networks on which they agree upon and on which they disagree upon (or the parts that stay unchanged and get changed under a sequence of rearrangement operations). Figure 3 illustrates this for the two unrooted phylogenetic networks N_1 and N_3 . While MAFs characterise the SPR-distance, we have shown that MAGs do not characterise the PR-distance (nor the SNPR-distance) of two rooted phylogenetic networks (Klavitter, 2019). This is due to similar reasons to why MAFs and MEAFs do not characterise the SPR-distance of two unrooted trees. However, we showed that MAGs induce a metric on phylogenetic networks, the agreement distance, which bounds the PR-distance of rooted networks (Klavitter, 2019).

In this paper, we look at how MAF and MEAF generalise for unrooted phylogenetic networks by introducing maximum agreement graphs and maximum endpoint agreement graphs and show that they

induce metrics (Section 3 and Section 4, respectively). We call these metrics the agreement distance (AD) and endpoint agreement distance (EAD), respectively. We then study the relations of AD, EAD, the TBR-distance, and the PR-distance in Section 5.

2 Preliminaries

This section provides the notation and terminology used in the remainder of the paper. In particular, we introduce notation in the context of phylogenetic networks as well as the PR and TBR operations.

Phylogenetic networks and trees. Let $X = \{1, 2, \dots, n\}$ be a finite set. An *unrooted binary phylogenetic network* N on X is a connected undirected multigraph such that the leaves are bijectively labelled with X and all non-leaf vertices have degree three. It is called *proper* if every cut-edge separates two labelled leaves (Francis et al., 2018a), and *improper* otherwise. Unless mentioned otherwise, we assume that a phylogenetic network is proper. Note that our definition permits the existence of parallel edges in N . An *unrooted binary phylogenetic tree* on X is an unrooted binary phylogenetic network that is a tree. See Figure 4 for examples. An edge of N is *external* if it is incident to a leaf, and *internal* otherwise.

To ease reading, we refer to a proper unrooted binary phylogenetic network (resp. unrooted binary phylogenetic tree) on X simply as a phylogenetic network or network (resp. phylogenetic tree or tree). Furthermore, let $u\mathcal{N}_n$ denote the set of all phylogenetic networks on X and let $u\mathcal{T}_n$ denote the set of all phylogenetic trees on X where $n = |X|$.




Fig. 4: An unrooted, binary phylogenetic tree $T \in u\mathcal{T}_6$ and an unrooted, binary proper phylogenetic network $N \in u\mathcal{N}_6$. The unrooted, binary phylogenetic network M is improper since the cut-edge e does not lie on a path that connects two leaves.

A network N has *reticulation number* r or, equivalently, is in *tier* r if r is the minimum number of edges that have to be deleted from N to obtain a spanning tree of N . Note that $r = |E| - (|V| - 1)$ where E and V are the edge and vertex set of N , respectively. This number is also known as the cyclomatic number of a graph (Diestel, 2017). Let $u\mathcal{N}_{n,r}$ denote tier r of $u\mathcal{N}_n$, that is, the set of networks in $u\mathcal{N}_n$ that are in tier r .

Suboperations and sprouts. Let G be an undirected graph. A degree-two vertex v of G with adjacent vertices u and w gets *suppressed* by deleting v and its incident edges, and adding the edge $\{u, w\}$. The reverse of this suppression is the *subdivision* of $\{u, w\}$ with a vertex v .

Let $\{u, v\}$ be an edge of G such that u either has degree one and is labelled (like a leaf of a network) or has degree three. A *pruning* of $\{u, v\}$ at u is the process of deleting $\{u, v\}$ and adding a new edge $\{\bar{u}, v\}$, where \bar{u} is a new (unlabelled) vertex. If u is now a degree two vertex, then we also suppress u . In reverse, the edge $\{\bar{u}, v\}$ gets *regrafted* to an edge $\{x, y\}$ by subdividing $\{x, y\}$ with a new vertex u and

then identifying \bar{u} and u . Alternatively, $\{\bar{u}, v\}$ may be *regrafted* to a labelled singleton u by identifying \bar{u} and u . The edge $\{u, v\}$ gets *removed* by deleting $\{u, v\}$ from N and suppressing any resulting degree-two vertices.

A *sprout* of G is an unlabelled degree one vertex of G . For example, applying a pruning to a phylogenetic network yields a graph with exactly one sprout.

Rearrangement operations. Let $N \in \mathcal{UN}_n$. The TBR operation is the rearrangement operation that transforms N into a phylogenetic network $N' \in \mathcal{UN}_n$ in one of the following four ways:

(TBR⁰) Remove an internal edge e of N , subdivide an edge of the resulting graph with a new vertex u , subdivide an edge of the resulting graph with a new vertex v , and add the edge $\{u, v\}$; or prune an external edge $e = \{u, v\}$ of N that is incident to leaf v at u , regraft the resulting sprout to an edge of the resulting graph.

(TBR⁺) Subdivide an edge of N with a new vertex u , subdivide an edge of the resulting graph with a new vertex v , and add the edge $e = \{u, v\}$.

(TBR⁻) Remove an edge e of N .

Note that a TBR⁰ can also be seen as the operation that prunes the edge $e = \{u, v\}$ at both u and v and then regrafts the two resulting sprouts. Hence, we say that a TBR⁰ *moves* the edge e . Furthermore, we say that a TBR⁺ *adds* the edge e and that a TBR⁻ *removes* the edge e . TBR is illustrated in Figure 5. Note that a TBR⁰ has an inverse TBR⁰ and that a TBR⁺ has an inverse TBR⁻, and that furthermore a TBR⁺ increases the reticulation number by one and a TBR⁻ decreases it by one. On trees, TBR⁰ equals the well known tree bisection and reconnection operation (Allen and Steel, 2001), which is also where the acronym comes from.




Fig. 5: Illustration of the TBR operation. The network N_2 can be obtained from N_1 by a TBR⁰ that moves the edge $\{u, v\}$ and the network N_3 can be obtained from N_2 by a TBR⁺ that adds the edge $\{u', v'\}$. Each operation has its corresponding TBR⁰ and TBR⁻ operation, respectively, that reverses the rearrangement.

Since a TBR operation has to yield a phylogenetic network, there are some restrictions on the edges that can be moved or removed. Firstly, if removing an edge by a TBR⁰ yields a disconnected graph, then in order to obtain a phylogenetic network an edge has to be added between the two connected components. For similar reasons, a TBR⁻ cannot remove a cut-edge. Secondly, the suppression of a vertex when removing an edge with a TBR⁻ may not yield a loop $\{u, u\}$. Thirdly, removing or moving an edge cannot create a cut-edge that does not separate two leaves. Otherwise the resulting network would be improper.

Let $N \in \mathcal{UN}_n$. A PR (*prune and regraft*) operation is the rearrangement operation that transforms N into a phylogenetic network $N' \in \mathcal{UN}_n$ with a PR⁺ = TBR⁺, a PR⁻ = TBR⁻, or a PR⁰ that prunes

and regrafts an edge e only at one endpoint, instead of at both like a TBR^0 (Janssen and Klawitter, 2019). Like for TBR, we say that the $\text{PR}^{0/+/-}$ moves/adds/removes the edge e in N . The PR operation is a generalisation of the well known SPR (*subtree prune and regraft*) operation on unrooted phylogenetic trees (Allen and Steel, 2001).

Distances. Let $N, N' \in \mathcal{UN}_n$. A TBR-*sequence* from N to N' is a sequence

$$\sigma = (N = N_0, N_1, N_2, \dots, N_k = N')$$

of phylogenetic networks such that N_i can be obtained from N_{i-1} by a single TBR for each $i \in \{1, 2, \dots, k\}$. The *length* of σ is k . The TBR-*distance* $d_{\text{TBR}}(N, N')$ between N and N' is the length of a shortest TBR-sequence from N to N' . The PR-distance is defined analogously. Janssen and Klawitter (2019, Corollary 4.4) have shown that the TBR- and PR-distance are well defined.

Embeddings and displaying. Let G be an undirected graph that is not necessarily simple; that is, G may contain parallel edges and loops. An edge $\{u, v\}$ of G is *subdivided* if $\{u, v\}$ is replaced by a path from u to v that contains at least one edge. A *subdivision* G^* of G is a graph that can be obtained from G by subdividing edges of G . If G has no degree two vertices, there exists a canonical mapping of vertices of G to vertices of G^* and of edges of G to paths of G^* .

Let N be an undirected graph, for example a network in \mathcal{UN}_n . Assume that G is connected. We say G has an *embedding* into N if there exists a subdivision G^* of G that is a subgraph of N . Now assume that G has components C_1, \dots, C_k . We say G has an *embedding* into N if the components C_i of G , for $i \in \{1, \dots, k\}$, have embeddings into N to pairwise edge-disjoint subgraphs of N . Note that these definitions imply that a labelled vertex of G^* is mapped to a labelled vertex of N with the same label.

We define a special type of embedding. Let n vertices of G be labelled bijectively with $X = \{1, 2, \dots, n\}$. We say G has an *agreement embedding* into N if there exists an embedding of G into N with the following properties.

- The pairwise edge-disjoint embeddings of components of G into N cover all edges.
- At most two vertices of G are mapped to the same vertex of N . In the case that exactly two vertices are mapped to the same vertex of N , one of these two vertices of G is a sprout and the other is a labelled, isolated vertex.
- For each labelled vertex v of N , there exists exactly one vertex \bar{v} with the same label in G and \bar{v} is mapped to v .

We make the observation that having an agreement embedding into a graph is a transitive property.

Observation 2.1. *Let G, H, N be undirected graphs such that G has an agreement embedding into H and H has an agreement embedding into N . Then G has an agreement embedding into N .*

Let $N, N' \in \mathcal{UN}_n$. We say N' *displays* N if N has an embedding into N' . For example, in Figure 4 the tree T is displayed by both networks N and M .

3 Agreement graph and distance

In this section we look at how agreement forests can be generalised for networks. Throughout this section, let $N, N' \in \mathcal{UN}_n$ be in tier r and r' , respectively. Without loss of generality, assume that $r' \geq r$ and let $l = r' - r$.

Suppose there is a TBR^0 that transforms N into N' by moving an edge e . This operation can be seen as removing e from N , obtaining a graph S , and then adding a new edge to S . We can interpret S as the part of N that remains unchanged or, in other words, N and N' agree on S . In general, we are interested in finding a graph that requires the minimal number of edge removals from N (or N') such that it has an embedding into N and N' . For two trees T and T' in \mathcal{UN}_n , this graph is precisely a *maximum agreement forest (MAF)* F . Allen and Steel (2001) showed that the number of components of F minus one is exactly the TBR-distance of T and T' , or, equivalently, the minimum number of edges that have to be removed from T (or T') to obtain F . If we consider again N and N' , then the removal of an edge must not necessarily increase the number of components. Therefore, instead of counting components, we are looking for a graph G consisting of components on which N and N' agree on and of additional edges that can be embedded into N and N' such that all edges are covered. In other words, we want that G has an agreement embedding into N and N' . Note that if N and N' are in different tiers, then we need additional edges for an agreement embedding into N' . We now make this precise.

Agreement graph. Let G be an undirected graph with connected components S_1, \dots, S_m and $E_1, \dots, E_{k-l}, E_{k-l+1}, \dots, E_k$ such that the S_i 's contain no sprouts and such that each E_j consist of a single edge on two unlabelled vertices. Then G is an *agreement graph* of N and N' if

- G without E_{k-l+1}, \dots, E_k has an agreement embedding into N , and
- G has an agreement embedding into N' .

For such an agreement graph, we refer to an S_i as *agreement subgraph* and to an E_j as a *disagreement edge*. A *maximum agreement graph (MAG)* G of N and N' is an agreement graph of N and N' with a minimal number of disagreement edges. See Figures 6 and 7 for two examples.




Fig. 6: A maximum agreement graph G for $N, N' \in \mathcal{UN}_n$. On the right, how G embeds into N' . Note that the disagreement edge E_2 is only needed for an agreement embedding into N' .

Note that if G contains m agreement subgraphs, then it also contains at least $m - 1$ disagreement edges since N and N' are connected graphs. Furthermore, unlike a MAF for two phylogenetic trees, G may contain agreement subgraphs without any labelled vertices and G may contain loops or parallel edges.

Let $T, T' \in \mathcal{UT}_n$. Let G be a maximum agreement graph of T and T' . Note that each agreement subgraph of G is a tree. Thus a (maximum) agreement graph of two trees is called a (maximum) agreement




Fig. 7: A maximum agreement graph G for $N, N' \in u\mathcal{N}_n$. On the right, how G embeds into N' . Note that G contains an agreement subgraph without labelled vertices.

forest. Note that the commonly used definition of agreement forests disregards disagreement edges (Allen and Steel, 2001).

Attached sprouts. Let V_N and E_N be the vertex and edge set of N , respectively. Let $G = (V_G, E_G)$ be an agreement graph of N and N' . Fix an agreement embedding of G into N . We say a sprout $\bar{u} \in V_G$ is *attached to* $\bar{e} \in E_G$ in N if \bar{u} is mapped to a vertex $u \in V_N$ that is an internal vertex of the path to which \bar{e} is mapped. Suppose G contains a labelled singleton \bar{x} . We say $\bar{u} \in V_G$ is *attached to* \bar{x} in N if \bar{u} and \bar{x} are mapped to the same leaf $u \in V_N$. This terminology can be extended from sprouts to disagreement edges. We say a disagreement edge E_i is *attached to an edge* $\bar{e} \in E_G$ in N if a sprout of E_i is attached to \bar{e} in N . Furthermore, we say E_i is *attached to an agreement subgraph* S_j in N if E_i is attached to an edge of S_j .

Note that a disagreement edge can be attached to itself. However, in general we would like to assume that an agreement embedding has nicer properties. This is what we look at next.

Ordered agreement embedding. Let G be an agreement graph of N and N' with agreement subgraphs S_1, \dots, S_m and disagreement edges E_1, \dots, E_k . Then an agreement embedding of G into N' is an *ordered agreement embedding* into N' if

- E_1 is attached to two distinct agreement subgraphs in N' ,
- E_i for $i \in \{2, \dots, m-1\}$ is attached to two distinct agreement subgraphs or an agreement subgraph and a disagreement edge E_j with $j < i$ in N' such that the subgraph of N' covered by S_1, \dots, S_m and E_1, \dots, E_i contains one connected component less than the subgraph of N' covered by S_1, \dots, S_m and E_1, \dots, E_{i-1} ,
- E_i for $i \in \{m, \dots, k\}$ is attached to agreement subgraphs or disagreement edges E_j with $j < i$ in N' .

An ordered agreement embedding of G into N is defined analogously but with the small difference that the third property only concerns the edges E_{m+1}, \dots, E_{k-l} .

Note that the first and second property of an ordered agreement embedding imply that the vertices and edges covered by the agreement subgraphs and the disagreement edges E_1, \dots, E_{m-1} form a connected subgraph of N . Moreover, in an ordered agreement embedding no disagreement edge is attached to itself. We now prove that an agreement graph always has an ordered agreement embedding.

Lemma 3.1. *Let $N, N' \in \mathcal{UN}_n$ be in tiers r and $r' \geq r$, respectively. Let $l = r' - r$. Let G be a maximum agreement graph for N and N' with m agreement subgraphs. Then G minus l disagreement edges has an ordered agreement embedding into N and G has an ordered agreement embedding into N' .*

Proof: The proof works the same for N and N' , so for simplicity we may assume that $l = 0$. Since G is a MAG of N and N' , there is an agreement embedding ϕ of G into N . Let $N_i, i \in \{1, \dots, m\}$, be the subgraphs of N to which the agreement subgraphs of G are mapped by ϕ . Colour all vertices and edges contained in these N_i 's black, and all other vertices and edges red. The red edges are thus the edges to which the disagreement edges of G are mapped. Note that the N_i 's are vertex-disjoint. Hence, since N is connected, it follows that the N_i 's are connected by red edges and paths. We use this fact to construct an ordered agreement embedding ϕ' of G into N .

For the ordered agreement embedding ϕ' map the agreement subgraphs of G into N like ϕ . Pick N_i and N_j such that there is path P from N_i to N_j with black end vertices and with red internal vertices and edges. Such a choice is possible by the observations above. Let ϕ' map E_1 to P . Colour the edges and vertices of P black, which makes N_i and N_j a single black subgraph N_i . Repeat this process for E_2, \dots, E_{m-1} . Note that this results in a single black component in N . Hence, for the remaining disagreement edges E_m, \dots, E_k we require from P only that it contains black end vertices and red internal vertices and edges, but not that P connects two distinct black components. As long as there remain red edges, we can find such P with a simple depth-first search in a red component that starts at a red edge incident to a black vertex and ends at a red edge incident to another black vertex. (Note that a red vertex always has degree three and a black, non-leaf vertex has at least degree two.) Therefore this process ends with all edges of N covered and coloured black. From a simple counting argument we get that we constructed exactly as many disagreement edges as G has. Hence, by construction the embedding ϕ' is an ordered agreement embedding of G into N . \square

We now define how to change an agreement embedding gradually. Let G be a MAG of N and N' . Let \bar{u} and \bar{v} be two sprouts of G with incident edges $\bar{e} = (\bar{u}, \bar{w})$ and $\bar{f} = (\bar{v}, \bar{z})$, respectively, such that \bar{u} is attached to \bar{f} in N . Let \bar{e} be mapped to the path $P = (y, \dots, w)$ in N and let \bar{f} be mapped to the path $P' = (x, \dots, y, \dots, z)$ in N . Then an *embedding change* of G into N with respect to \bar{u} and \bar{v} is the change of the embedding such that \bar{e} is mapped to the path (x, \dots, y, \dots, w) formed by a subpath of P' and the path P , and such that \bar{f} is mapped to the subpath (y, \dots, z) of P' ; see Figure 8.



Fig. 8: An embedding change with respect to \bar{u} and \bar{v} .

Agreement distance. Let $N, N' \in \mathcal{UN}_n$. Let G be a MAG of N and N' with k disagreement edges. We define the *agreement distance* d_{AD} of N and N' as

$$d_{AD}(N, N') = k.$$

Note that the agreement distance also equals half the number of sprouts of G .

Theorem 3.2. *The agreement distance d_{AD} on $u\mathcal{N}_n$ is a metric.*

Proof: Note that d_{AD} is symmetric, non-negative, and for all $M, M' \in u\mathcal{N}_n$ $d_{AD}(M, M') = 0$ if and only if $M = M'$. Therefore, to show that d_{AD} is a metric, it remains to show that d_{AD} satisfies the triangle inequality.

Let $N, N', N'' \in u\mathcal{N}_n$ and let $k' = d_{AD}(N, N')$ and $k'' = d_{AD}(N', N'')$. Let G' (resp. G'') be a MAG of N and N' (resp. N' and N'') with k' (resp. k'') disagreement edges. To show that the triangle inequality holds, we construct an agreement graph G for N and N'' with at most $k' + k''$ disagreement edges. The following construction is illustrated with an example in Figure 10.

For simplicity, assume for now that N, N' , and N'' are in the same tier. Fix ordered agreement embeddings of G' and G'' into N' , which is possible by Theorem 3.1. Based on the ordered agreement embedding we can construct a length- k'' sequence of graphs ($N' = M_0, M_1, \dots, M_{k''} = G''$) from N' to G'' where M_i is obtained from M_{i-1} for $i \in \{1, \dots, k''\}$ by removing an edge and adding a disagreement edge. Note that M_i has an agreement embedding into M_{i-1} and thus by the transitive property of agreement embeddings (recall Theorem 2.1) also an agreement embedding into N' . We use this sequence, to construct a sequence of graphs ($G' = G_0, G_1, \dots, G_{k''} = G$) such that G_i is obtained from G_{i-1} for $i \in \{1, \dots, k''\}$ either by setting $G_i = G_{i-1}$ or by the removal of an edge of an agreement subgraph and adding a disagreement edge. First, colour the disagreement edges of G_0 red. We will colour each newly added disagreement edge blue. Our construction will ensure the following properties:

- The only sprouts of G_i are in disagreement edges;
- G_i has an agreement embedding into M_i ,
- each blue disagreement edge of G_i is mapped to a disagreement edge of M_i .

Suppose from M_{i-1} to M_i the edge e gets removed and disagreement edge F added. We distinguish three cases, which are illustrated in Figure 9 (a) to (c). First, if there is an edge \bar{e} of G_{i-1} that is mapped to e by the agreement embedding of G_{i-1} into M_{i-1} and that is not incident to a sprout, then obtain G_i from G_{i-1} by removing \bar{e} and adding a blue disagreement edge E_j . This is shown in Figure 9 (a) and in the step from G_0 to G_1 in Figure 10. Note that G_i has an agreement embedding in M_i where E_j is mapped to F . Clearly G_i also has an agreement embedding into G_{i-1} and thus by Theorem 2.1 into N and N' .

Second, suppose that an edge \bar{e} of an agreement subgraph of G_{i-1} is mapped to a path P_e of M_{i-1} that contains $e = \{u, v\}$. If u (or v) lies within P_e , then a sprout of a disagreement edge of G_{i-1} is attached to it. Note that this sprout belongs to a red disagreement edge since blue disagreement edges are mapped to edges that got removed in an earlier step. Obtain G_i from G_{i-1} by removing \bar{e} and adding a blue disagreement edge E_j . This case also applies in the step from G_1 to G_2 in Figure 10. For the agreement embedding of G_i into M_i apply an embedding change (or embedding changes) as shown in Figure 9 (b). Then E_j is mapped onto F .

Third, suppose that a red disagreement edge \bar{e} of G_{i-1} is mapped to a path P_e that contains $e = \{u, v\}$. In this case set $G_i = G_{i-1}$. To obtain an agreement embedding of G_i into M_i apply again appropriate embedding changes as shown in Figure 9 (c) and in the step from G_2 to G_3 in Figure 10.

We claim that $G = G_{k''}$ is an agreement graph of N and N'' . By Theorem 2.1 G has an agreement embedding into G' and thus into N (and N'). Furthermore, G has an agreement embedding into $M_{k''} = G''$. Therefore, again by Theorem 2.1, we get that G has an agreement embedding into N'' . Concerning the




Fig. 9: How to obtain G_i (right) from G_{i-1} (left) with respect to the agreement embedding of G_{i-1} into M_{i-1} when an edge of an agreement subgraph is mapped precisely to e (a) or to a path containing e (b); or when a disagreement edge is mapped to a path containing e (c). It is also shown how embedding changes are applied to show that G_i has an agreement embedding into M_i .

components of G , note that G contains precisely k' red disagreement edges and at most k'' blue disagreement edges. By construction these disagreement edges contain all sprouts of G . Hence, G is an agreement graph of N and N'' that proves that $d_{\text{AD}}(N, N'') \leq k' + k'' = d_{\text{AD}}(N, N') + d_{\text{AD}}(N', N'')$. This concludes the proof for the case when N , N' , and N'' are in the same tier.

Cases where N , N' , and N'' are in different tiers work analogously. However, when for example N'' is in a higher tier than N' then in the construction of the sequence $(N' = M_0, M_1, \dots, M_{k''} = G'')$ from N' to G'' we stop removing edges at some point and only add disagreement edges. These extra disagreement edges are only needed for the agreement embedding into N'' but not N' . The same applies then to the construction of $(G' = G_0, G_1, \dots, G_{k''} = G)$. \square

Next, we show that if we restrict the agreement distance to the space of phylogenetic trees, then it equals the TBR-distance.

Proposition 3.3. *The agreement distance is equivalent to the TBR-distance on $u\mathcal{T}_n$.*

Proof: Let G be a maximum agreement forest of two trees $T, T' \in u\mathcal{T}_n$. Allen and Steel (2001) defined the function $m(T, T')$ as the number of agreement subgraphs of G minus one. If G contains k disagreement edges, then it contains $k + 1$ agreement subgraphs. Thus, $d_{\text{AD}}(T, T') = m(T, T')$. By Theorem 2.13 of Allen and Steel (2001), $m(T, T') = d_{\text{TBR}}(T, T')$. This concludes the proof. \square

Allen and Steel (2001) further showed that computing the TBR-distance of two phylogenetic trees is NP-hard. Janssen and Klawitter (2019, Theorem 6.1) showed that the TBR-distance of two trees in $u\mathcal{T}_n$ is the same as in $u\mathcal{N}_n$. These two results together with Theorem 3.3 give us the following corollary.

Corollary 3.4. *Computing the agreement distance on $u\mathcal{N}_n$ is NP-hard.*

4 Endpoint agreement graph and distance

While a TBR^0 prunes an edge at both ends, a PR^0 only prunes an edge at one side. Hence, agreement graphs are not suited to model PR^0 . In this section we introduce endpoint agreement graphs as a slight modification of agreement graphs which model PR^0 more closely. Let again $N, N' \in u\mathcal{N}_n$ be in tiers r and r' , respectively, and let $l = r' - r$.

Endpoint agreement graph. Let H be an undirected graph with connected components S_1, \dots, S_m and E_1, \dots, E_l such that each E_j consist of a single edge on two unlabelled vertices. Then H is an *endpoint agreement graph* (EAG) of N and N' if

- H without E_1, \dots, E_l has an agreement embedding into N , and

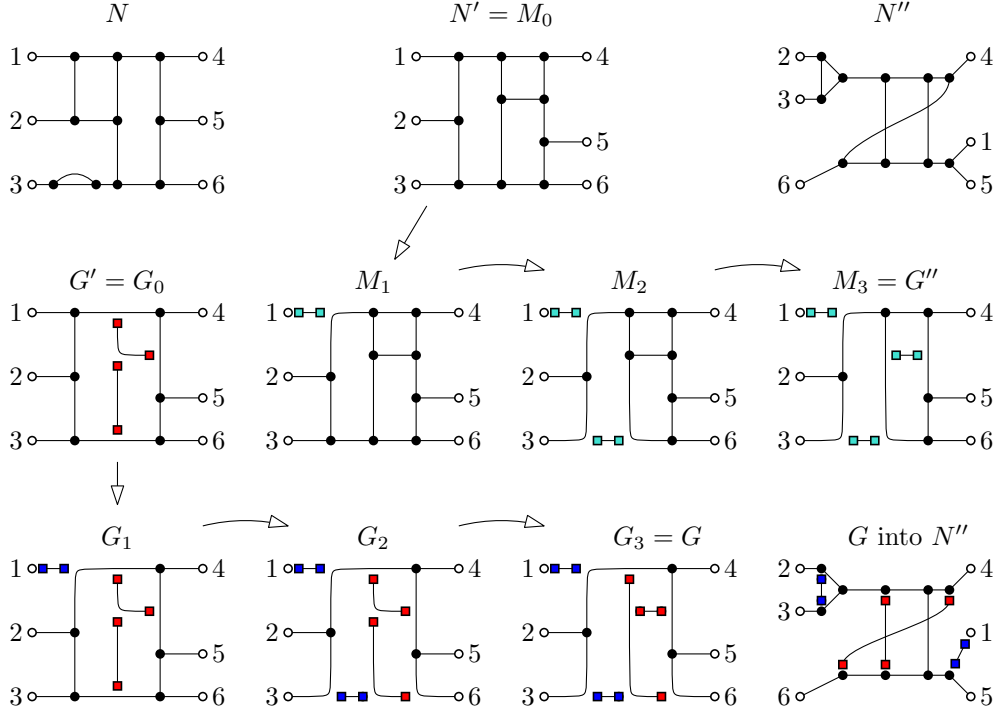


Fig. 10: An example for the constructions of the sequences $(N' = M_0, M_1, \dots, M_{k''} = G'')$ and $(G' = G_0, G_1, \dots, G_{k'} = G)$. Here, the networks N and N' have agreement distance 2 with maximum agreement graph G' ; the networks N' and N'' have agreement distance 3 with maximum agreement graph G'' . Lastly, an agreement embedding of the constructed G into N'' is shown.

- H has an agreement embedding into N' .

We refer to an S_i as (*endpoint*) *agreement subgraph* and to an E_j as *disagreement edge*. A *maximum endpoint agreement graph (MEAG)* H of N and N' is an endpoint agreement graph of N and N' with a minimal number of sprouts. See Figure 11 for an example. Note that, unlike to MAG, in a MEAG also endpoint agreement subgraphs can contain sprouts. We define an *ordered agreement embedding* of H into N' as an agreement embedding of H into N' such that

- no sprout of an agreement subgraph is attached to a disagreement edge and
- and the disagreement edges can be ordered (E_1, \dots, E_l) such that E_j may be attached to E_i only if $i \leq j$.

For an ordered agreement embedding of H into N only the first property has to hold.

A proof that ordered agreement embeddings exists works analogously to the proof of Theorem 3.1 and the proof of Lemma 3.2 (Klawitter, 2019), yet we outline the proof idea here. Starting with an agreement embedding of G into N' , apply embedding changes to any sprout of an endpoint agreement subgraph that




Fig. 11: A maximum endpoint agreement graph H for $N, N' \in u\mathcal{N}_n$. On the right, how H embeds into N' .

is attached to a disagreement edge. This way the first property can be enforced. For the second property, apply embedding changes if E_i that is attached to a disagreement edge E_j , $j > i$, for $i \in \{1, \dots, l\}$.

Endpoint agreement distance. Let H be a MEAG of N and N' . Let s be the number of sprouts of agreement subgraphs of H and let l be the number of disagreement edges of H . We define the *endpoint agreement distance* (EAD), denoted by d_{EAD} , of N and N' as

$$d_{\text{EAD}}(N, N') = s + l.$$

Following Whidden and Matsen (2019) we use their replug operation to show that the EAD is a metric.

Replug distance. We define a *replug network* M on X as an undirected multigraph such that the leaves and singletons are bijectively labelled with X and all non-leaf vertices have degree three. Unlike for a phylogenetic network, M may contain loops and be disconnected. Let $u\mathcal{M}_n$ be the set of all replug networks on X . Note that $u\mathcal{N}_n \subseteq u\mathcal{M}_n$.

Let $M \in u\mathcal{M}_n$. A *replug operation* is the rearrangement operation that transforms M into a replug network $M' \in u\mathcal{M}_n$ by pruning an edge at one vertex and then regrafting it again or by a vertical operation like a PR^+ or a PR^- . Unlike for PR , a replug operation does not have to ensure that the resulting network is connected or proper.

Let $N, N' \in u\mathcal{N}_n$. We define the *replug distance* d_R of N and N' as the distance of N and N' in $u\mathcal{M}_n$ under the replug operation. Note that since $u\mathcal{N}_n$ is connected under PR , it is also connected as subgraph of $u\mathcal{M}_n$ under the replug operation. Therefore, the replug distance is well defined and a metric. We now use the replug distance to prove that the endpoint agreement distance is a metric.

Proposition 4.1. *The endpoint agreement distance is equivalent to the replug-distance on $u\mathcal{N}_n$.*

Proof: Let $N, N' \in u\mathcal{N}_n$. We first prove that $d_R(N, N') \geq d_{\text{EAD}}(N, N')$. Let $d = d_R(N, N')$ and let $\sigma = (N = M_0, M_1, \dots, M_d = N')$ be a shortest replug sequence. Suppose that N and N' are in different tiers and that N' is above N . Note that we may assume that σ does not use any PR^- -like operation, as such an operation and the next PR^+ -like operation can be replaced with at most two replug operations that prune and regraft the same edge. We construct a sequence of graphs $(M_0 = H_0, H_1, \dots, H_d)$ such that H_i has an agreement embedding into M_i for $i \in \{0, 1, \dots, d\}$. The construction will also ensure that H_i has an agreement embedding into H_{i-1} (using the right number of disagreement edges) and thus by Theorem 2.1 also into $M_0 = N$.

Suppose M_i is obtained from M_{i-1} by a horizontal replug operation θ that prunes the edge $e = \{u, v\}$ at u . Consider the agreement embedding of H_{i-1} into M_{i-1} . Let $\bar{e} = \{\bar{u}, \bar{v}\}$ be the edge of H_{i-1} that is mapped to a trail $P = (w_1, w_2, \dots, w_k)$ containing e . Assume without loss of generality that \bar{u} is mapped

to w_1 and \bar{v} to w_k . (If $w_1 = w_k$, we further assume that \bar{e} imposed with the directed (\bar{u}, \bar{v}) is mapped from w_1 towards w_k .) Now, if $w_1 = u$, that is, \bar{u} is mapped to u , then we can prune \bar{e} at \bar{u} in H_{i-1} (unless \bar{u} already is a sprout) to obtain H_i . If \bar{u} already is a sprout, set $H_i = H_{i-1}$. The agreement embedding of H_i into M_i is derived from the agreement embedding of H_{i-1} into M_{i-1} and how θ regrafts e . If $w_1 \neq u$, a sprout \bar{x} of H_{i-1} is mapped to u . In this case, prune \bar{e} at \bar{u} (unless \bar{u} already is a sprout) to obtain H_i . If \bar{u} already is a sprout, set $H_i = H_{i-1}$. Then apply an embedding change with respect to \bar{x} and \bar{u} to obtain an agreement embedding into M_{i-1} . Derive an agreement embedding into M_i as in the previous case. Clearly H_i has an agreement embedding into H_{i-1} .

Next, suppose M_i is obtained from M_{i-1} by a vertical replug operation that adds the edge $e = \{u, v\}$ by subdividing the edges f and f' . Obtain H_i from H_{i-1} by adding a disagreement edge and obtain an agreement embedding of H_i into M_i by mapping the disagreement edge to e .

At the end of the sequence, H_d is an endpoint agreement graph of N and $M_d = N'$. Since we added at most d sprouts or disagreement edges, it follows that $d_R(N, N') \geq d_{\text{EAD}}(N, N')$.

We now prove that $d_{\text{EAD}}(N, N') \geq d_R(N, N')$. Let H be a maximum endpoint agreement graph of N and N' . Fix ordered agreement embeddings of H into N and N' , i.e., no sprout of an agreement subgraph is attached to a disagreement edge. Based on agreement embeddings of H into N and N' it is straightforward to use a replug operation for each sprout of H to prune an edge of N (or a resulting network) and regraft it according to the agreement embedding of H into N' . Lastly, if N' is in a tier above N , use a PR^+ -like replug operation for each disagreement edge of H to add an edge according to the agreement embedding of H into N' . \square

Corollary 4.2. *The endpoint agreement distance on $u\mathcal{N}_n$ is a metric.*

Whidden and Matsen (2019) showed that the endpoint agreement distance (or rather the replug distance) does not always equal the SPR-distance of two trees. Furthermore, they conjectured that computing the endpoint agreement distance is NP-hard for trees. This and whether it is NP-hard to compute the endpoint agreement distance of two networks remains open.

5 Relations of distances

In this section we look at the relations of the metrics induced by MAG, MEAG, TBR, and PR. We start by comparing the agreement distance with the TBR-distance. As we have seen in Theorem 3.3, they are equivalent on $u\mathcal{T}_n$. Furthermore, we can make the following observations.

Observation 5.1. *Let $N, N' \in u\mathcal{N}_n$. Then $d_{\text{AD}}(N, N') = 1$ if and only if $d_{\text{TBR}}(N, N') = 1$.*

Lemma 5.2. *Let $N, N' \in u\mathcal{N}_n$ be in tiers r and r' , respectively, such that N' displays N . Let $l = r' - r$. Then $d_{\text{AD}}(N, N') = d_{\text{TBR}}(N, N') = l$.*

Proof: The second equality follows from Corollary 5.6 by Janssen and Klawitter (2019). The equality also implies that there is a TBR^+ -sequence σ of length l from N to N' . Let G be the graph obtained from N by adding l disagreement edges. Then G without its disagreement edges has an agreement embedding into N and we can obtain an agreement embedding into N' from σ straightforwardly. Hence, G is a MAG of N and N' , which proves the first equality. \square

Lemma 5.3. *Let $T \in \mathcal{uT}_n$ and $N \in \mathcal{uN}_{n,r}$. Then $d_{AD}(T, N) = d_{TBR}(T, N)$.*

Proof: Janssen and Klavitter (2019, Theorem 4.13) showed that there is tree T' that is displayed by N such that $d_{TBR}(T, N) = d_{TBR}(T, T') + d_{TBR}(T', N)$. The tree T' is thus a tree that minimises the TBR-distance to T among all trees displayed by N . From Theorem 3.3 and Theorem 5.2 we thus get that

$$d_{AD}(T, N) \leq d_{AD}(T, T') + d_{AD}(T', N) = d_{TBR}(T, T') + d_{TBR}(T', N) = d_{TBR}(T, N).$$

For the converse direction, consider a maximum agreement graph G of T and N with k disagreement edges. From an ordered agreement embedding of G into N , we get that G with $k - r$ disagreement edges embeds onto a tree T' displayed by N . Hence,

$$d_{AD}(T, N) \geq d_{AD}(T, T') + d_{AD}(T', N) = d_{TBR}(T, N).$$

□

After these three cases, where the agreement distance and the TBR-distance are equivalent, we show with the following example that this is in general not the case.




Fig. 12: Two networks $N, N' \in \mathcal{uN}_n$ with $d_{AD}(N, N') = 2$, but with $d_{TBR}(N, N') = 3$ as proven in Theorem 5.4 (for example with a TBR-sequence via M). The graph G illustrated with an agreement embedding into N is a MAG of N and N' .

Lemma 5.4. *The networks N and N' in Figure 12 have $d_{AD}(N, N') = 2$, $d_{EAD}(N, N') = 2$, and $d_{TBR}(N, N') = 3$.*

Proof: Concerning the agreement distance, observe that $d_{AD}(N, N') > 1$. Next, note that the graph G in Figure 12 has agreement embeddings into N and N' . This also yields an agreement embedding of G into N' by swapping the singletons labelled 1 and 2. Hence, G with two disagreement edges is a MAG of N and N' , which proves that $d_{AD}(N, N') = 2$.

Concerning the endpoint agreement distance, we see that the leaves 1 and 2 can be swapped with two replug operations.

Concerning the TBR-distance, observe that there is no length two TBR^0 -sequence from N to N' . This can be seen as with only two TBR^0 the leaves 1 and 2 cannot be swapped nor can the two biconnected components be transformed into each other within \mathcal{uN}_n . To see that $d_{TBR}(N, N') = 3$, note that $d_{TBR}(N, M) = 1$ and that the leaves 1 and 2 can be swapped with a single TBR^0 in M resulting in a network M' with $d_{TBR}(M', N') = 1$. □

Next, we show that the agreement distance provides a lower and an upper bound on the TBR-distance of any two networks N and N' .

Lemma 5.5. *Let $N, N' \in \mathcal{uN}_n$. Then*

$$d_{AD}(N, N') \leq d_{TBR}(N, N').$$

Proof: Let $d = d_{TBR}(N, N')$ and let $\sigma = (N = M_0, M_1, \dots, M_d = N')$ be a TBR-sequence from N to N' . To prove the lemma, we show how to obtain an agreement graph G of N and N' with at most d disagreement edges from σ .

We construct a sequence of graphs $(N = G_0, G_1, \dots, G_d = G)$ such that G_i is an agreement graph of M_0 and M_i for $i \in \{0, 1, \dots, d\}$. This holds trivially for $i = 0$. In the following, when we consider agreement embeddings of G_i into M_0 and M_i where M_0 and M_i are in different tiers then we ignore, for simplicity, that one of the embeddings needs less disagreement edges.

Suppose M_i is obtained from M_{i-1} by a TBR^0 that moves the edge $e = \{u, v\}$. Let \bar{e} be the edge of G_{i-1} that is mapped to a path P that contains e by the agreement embedding of G_{i-1} into M_{i-1} . We distinguish four cases, namely whether \bar{e} is part of an agreement subgraph and whether P contains only e . (They are comparable to the cases in the proof of Theorem 3.2; see also Figure 9 again.)

1. Assume that \bar{e} is part of an agreement subgraph and mapped precisely to e . Then obtain G_i by removing \bar{e} and adding a disagreement edge. Clearly G_i has an agreement embedding into N and M_i .
2. Assume that \bar{e} is part of an agreement subgraph and P has length at least two. Further assume without loss of generality that neither u nor v is an end vertex of P . Then there are sprouts \bar{u} and \bar{v} that are attached to \bar{e} in M_{i-1} and that are mapped to u and v , respectively. Again obtain G_i by removing \bar{e} and adding a disagreement edge $\{\bar{x}, \bar{y}\}$. However, for an agreement embedding of G_i into M_{i-1} map $\{\bar{x}, \bar{y}\}$ to P and then apply embedding changes with respect to \bar{x} and \bar{u} and with respect to \bar{y} and \bar{v} . We can derive from this an agreement embedding of G_i into M_i .
3. Assume that \bar{e} is a disagreement edge and mapped precisely to e . Then set $G_i = G_{i-1}$ and it is straightforward to obtain agreement embeddings.
4. Last, assume that \bar{e} is a disagreement edge and P has length at least two. There are then again without loss of generality two sprouts attached to \bar{e} . Set $G_i = G_{i-1}$ and obtain an agreement embedding of G_i into M_{i-1} (and M_i) by applying embedding changes as in the second case. Hence, in either case, we obtain an agreement graph G_i of M_0 and M_i .

Next, suppose M_i is obtained from M_{i-1} by a TBR^- that removes the edge $e = \{u, v\}$. Like for a TBR^0 , if an edge \bar{e} of an agreement graph is mapped to e , we obtain G_i from G_{i-1} by removing \bar{e} . Otherwise we set $G_i = G_{i-1}$. Furthermore, if M_{i-1} is in a higher tier than M_0 we also remove a disagreement edge. Using again embedding changes if sprouts were attached to \bar{e} , it is straightforward to construct an agreement embedding of G_i into M_i . Thus G is an agreement graph of M_0 and M_i .

Lastly, suppose that M_i is obtained from M_{i-1} by a TBR^+ . If M_i is in a higher tier than M_0 , then obtain G_i from G_{i-1} by adding a disagreement edge. Otherwise, set $G_i = G_{i-1}$. In either case, it is clear that G_i is an agreement graph of M_0 and M_i .

Note that for each TBR of σ we added at most one disagreement edge. Hence, G_d , which is an agreement graph of M_0 and M_d , is an agreement graph of N and N' with at most d disagreement edges. This concludes the proof. \square

Lemma 5.6. *Let $N, N' \in \mathcal{UN}_n$. Then $d_{\text{TBR}}(N, N') \leq 2 d_{\text{AD}}(N, N')$.*

Proof: Suppose N and N' are in tier r and r' , respectively, and that $r' \geq r$. Let $l = r' - r$. Let $d = d_{\text{AD}}(N, N')$ and $k = d - l$. Let G be a MAG of N and N' with agreement subgraphs S_1, \dots, S_m and disagreement edges $E_1, \dots, E_k, E_{k+1}, \dots, E_d$. To prove the theorem, we construct a TBR-sequence

$$\sigma = (N = M_0, M_1, \dots, M_d, \dots, M_{d+k} = N')$$

such that M_i is obtained from M_{i-1} by a TBR^+ for $i \in \{1, \dots, d\}$ and by a TBR^- for $i \in \{d+1, \dots, d+k\}$. Along σ we maintain a series of graphs G_0, \dots, G_{d+k} such that G_i has an agreement embedding into M_i .

Fix ordered agreement embeddings of G into N and N' , which is possible by Theorem 3.1. Let N_1, \dots, N_m be the subgraphs of N to which S_1, \dots, S_m of G are mapped, respectively. We define N'_1, \dots, N'_m analogously for N' . Note that the disagreement edges of G are mapped to paths in N (resp. N') that (as a whole) pairwise connect the N_i 's (resp. N'_i 's). The idea is now as follows. From M_0 to M_d we add d edges to reconstruct the paths that connect the N'_i 's as in N' while maintaining the paths that connect the N_i 's as in N . From M_d to M_{d+k} we then remove edges guided by how the paths connect the N_i 's in N . This is illustrated in Figure 13. We now define the graphs G_i formally and explain how to construct σ .

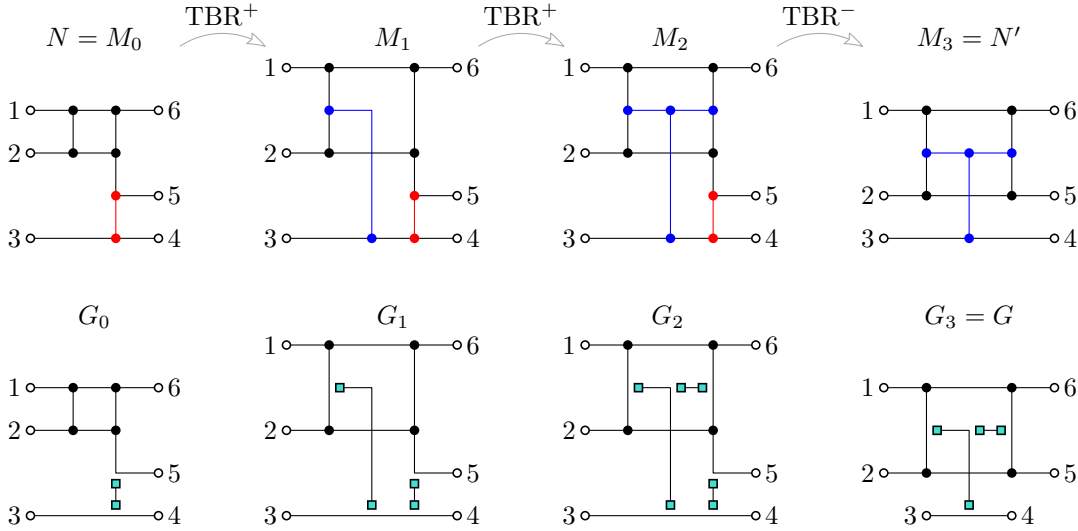


Fig. 13: Construction of a TBR-sequence from N to N' based on a MAG G of N and N' .

Let G_0 be G without l disagreement edges. Therefore, G_0 has k disagreement edges and an agreement embedding into $N_0 = N$ without spare disagreement edges. For $i \in \{1, \dots, d\}$ let G_i be G_{i-1} plus one disagreement edge. Next, for $i \in \{d+1, \dots, d+k\}$ let G_i be G_{i-1} minus one disagreement edge. Let E_j^i for $i \in \{0, 1, \dots, d+k\}$ and $j \in \{1, \dots\}$ denote the disagreement edges of G_i .

In M_0 colour the subgraph to which agreement subgraphs of G_0 are mapped black. Colour all other vertices and edges red. Obtain M_1 from M_0 as follows. First, assume that E_1^{d+k} is attached to edges

\bar{e} and \bar{e}' of agreement subgraphs in M_{d+k} . Consider the paths $P_{\bar{e}}$ and $P_{\bar{e}'}$ in N_0 to which \bar{e} and \bar{e}' are mapped. Ignoring vertices on these paths that are incident to red edges, we can perceive $P_{\bar{e}}$ and $P_{\bar{e}'}$ as edges e and e' . Then add with a TBR^+ an edge f from e to e' . Next assume that E_1^{d+k} is attached to a labelled singleton \bar{u} of G in M_{d+k} . Then let u be the leaf of M_0 to which \bar{u} is mapped. Note that u is incident to a red edge, say e , in M_0 . Obtain e' as in the previous case. If E_1^{d+k} is attached to two labelled singletons, then obtain a second red edge e' analogous to how we obtained e . In either case, apply the TBR^+ that adds an edge f from e to e' . Let M_1 be the resulting network. Colour the new edge f blue. Obtain an agreement embedding of G_1 into M_1 by extending the agreement embedding of G_0 into N_0 by mapping E_{k+1}^1 to f . Note that M_1 is a proper phylogenetic network since adding an edge (with a TBR^+) to a proper network yields a proper network. In particular, edges obtained from subdividing e and e' still lie on paths between leaves and so does thus f . Repeat this process to obtain M_i from M_{i-1} for $i \in \{2, \dots, d\}$ based on how E_i^{d+k} embeds into M_{d+k} .

Observe that M_d and G_d with its agreement embedding into M_d can also be obtained by applying the construction we used to obtain M_d from M_0 by starting from M_{d+k} and considering the agreement embedding of G_0 into M_0 (instead of the agreement embedding of G_{d+k} into M_{d+k}). The only difference is that in the two resulting agreement embeddings of G_d into M_d blue disagreement edges might be attached to red edges or vice versa, wherever there is a labelled singleton (leaf). Nevertheless, this shows that we can construct the full TBR-sequence σ . To conclude the proof, note that $d + k \leq 2d$. \square

From Theorem 5.5 and Theorem 5.6 we get the following theorem.

Theorem 5.7. *Let $N, N' \in \mathcal{uN}_n$. Then*

$$d_{AD}(N, N') \leq d_{\text{TBR}}(N, N') \leq 2 d_{AD}(N, N').$$

Janssen and Klawitter (2019, Corollary 3.3) showed that the PR-distance is bound from below by the TBR-distance and from above by at most twice the TBR-distance. Hence, we get the following corollary.

Corollary 5.8. *Let $N, N' \in \mathcal{uN}_n$. Then*

$$d_{AD}(N, N') \leq d_{\text{PR}}(N, N') \leq 4 d_{AD}(N, N').$$

We now turn to the endpoint agreement distance and look at its relation to the agreement distance and the PR-distance.

Proposition 5.9. *Let $N, N' \in \mathcal{uN}_n$. Then*

$$d_{AD}(N, N') \leq d_{\text{EAD}}(N, N') \leq 2 d_{AD}(N, N').$$

Proof: We start with the first inequality. Let H be a maximum endpoint agreement graph of N and N' . Suppose H has s sprouts in agreement subgraphs and l disagreement edges. We prove that there is an agreement graph G of N and N' with at most $s + l$ disagreement edges. For this, we construct a sequence of graphs $(G_s, G_{s-1}, \dots, G_0)$ such that each G_i is an endpoint agreement graph of N and N' with at most i sprouts in agreement subgraphs and at most $l + (s - i)$ disagreement edges. (We slightly abuse the definition of disagreement edges here and consider any edge incident with two sprouts of G_i as a disagreement edge.) Therefore, setting $G = G_0$ will prove the proposition.

Suppose \bar{u} is a sprout of an agreement subgraph of G_{i+1} . Let $\bar{e} = \{\bar{u}, \bar{v}\}$ be the edge incident to \bar{u} . If \bar{v} is also a sprout, set $G_i = G_{i+1}$ and classify \bar{e} as a disagreement edge. Furthermore, we can also

set $G_{i-1} = G_i$ since we eliminated two sprouts of agreement subgraphs at once. Otherwise, obtain G_i from G_{i+1} by pruning \bar{e} from \bar{v} . Since \bar{v} is either a degree vertex or a labelled leaf, we can directly derive agreement embeddings of G_i into N and N' from the agreement embeddings of G_{i+1} . Since every step reduces the number of sprouts in agreement subgraphs by at least one and adds at most one disagreement edge, G_0 is as desired.

For the second inequality, note that a MAG of N and N' with k disagreement edges is also an EAG of N and N' with $2k$ sprouts. \square

Theorem 5.10. *Let $N, N' \in \mathcal{UN}_n$. Then*

$$d_{EAD}(N, N') \leq d_{PR}(N, N') \leq 3 d_{EAD}(N, N').$$

Proof: Without loss of generality, assume that N is not in a higher tier than N' . For the lower bound, consider a shortest PR-sequence σ from N to N' . Note that σ is also a replugin sequence. There is thus a replugin sequence from N to N' whose length is at most the length of σ . The lower bound now follows from Theorem 4.1.

Next, we prove the upper bound. Let H be a MEAG for N and N' with s sprouts in agreement subgraphs and l disagreement edges. Fix ordered endpoint agreement embeddings of H into N and N' . Let $d = d_{EAD}(N, N') = s + l$. We construct a PR-sequence $\sigma = (N = M_0, M_1, \dots, M_{d'} = N')$ with $d' \leq 3d$. Along σ , we maintain a sequence of graphs $(H = H_0, H_1, \dots, H_{d'})$ that consist of H plus possibly extra disagreement edges such that H_i has an agreement embedding into M_i . We call these extra disagreement edges *ghost disagreement edges*.

Let E_1, \dots, E_l be the disagreement edges of H . For $i \in \{1, \dots, l\}$ obtain M_i from M_{i-1} by adding an edge e with a PR^+ according to where the disagreement edge E_i is attached to in the agreement embedding of H into N' . If E_i is attached to a labelled singleton \bar{v} , then both \bar{v} and a sprout \bar{u} are mapped to a leaf v of M_{i-1} . In this case, attach e to the edge incident to v . For an agreement embedding of H into M_i , map the disagreement edge E_i to the newly added edge. If E_i should be attached to v , apply the appropriate embedding change with \bar{u} . Set $H_i = H$.

We now use PR^0 to move edges according from where sprouts are attached to in M_i to where they are attached to in N' . If we have done this for a sprout, we call it *handled* and *unhandled* otherwise. Let \bar{u} be an unhandled sprout of H_i . Let \bar{u} and its incident edge be mapped to u and $e = \{u, v\}$ of M_{i-1} , respectively. If e can be pruned at u and attached to the edge according to where \bar{u} is mapped to in N' such that the result is a proper phylogenetic network, then apply this PR^0 to obtain M_i . Set $H_i = H_{i-1}$. Note that in the case that \bar{u} is mapped to a leaf w in N' , then e is attached the edge incident to w , and we apply the appropriate embedding change for H_i . (Apply this to each unhandled sprout where possible). Otherwise, use a PR^+ to add a (ghost) edge f from e to the edge incident to leaf l to obtain M_i . Obtain H_i from H_{i-1} by adding a ghost disagreement edge F . Map F to f and apply an embedding change with respect to F and \bar{u} . Note that now the first case applies for \bar{u} and e . Thus we also obtain M_{i+1} and $H_{i+1} = H_i$. When all sprouts are handled, the agreement embedding of H_i without ghost disagreement edges is mapped to the subgraph of M_i that is precisely a subdivision of N' . We thus need at most s further PR^- to remove all ghost edges.

In total, this process requires at most $l + 3s = d' \leq 3d$ PR. This proves the upper bound. \square

6 Concluding remarks

In this paper, we defined maximum agreement graphs (MAG) for two unrooted, proper, binary phylogenetic networks. Like maximum agreement forests for trees, a MAG models how two networks agree on subgraphs that stay untouched when moving edges with TBR operations. If the two networks are in different tiers, then a MAG also models how the networks disagree on that. Based on MAGs, we defined the agreement distance of phylogenetic networks. By showing that this new metric is equivalent to the TBR-distance for two trees, we obtained that it is NP-hard to compute the agreement distance.

We have seen that the agreement distance and the TBR-distance are equivalent for trees and for networks with distances of at most one. Furthermore, we know that the agreement distance of a tree and a network equals their TBR-distance. On the other hand, there are networks N and N' , as in Figure 12, with agreement distance two but higher TBR-distance. However, note that N and N' are in tier seven. It is therefore of interest to further study when exactly the agreement distance is equivalent to the TBR-distance and when not. In general, we showed that the agreement distance of two networks provides a natural lower bound and an upper bound with factor two on their TBR-distance. If we drop the requirement that networks have to be proper, it is also open whether the agreement distance and the TBR-distance are equivalent or not.

Like SPR on trees has been generalised to PR on networks, we have generalised maximum endpoint agreement forests of Whidden and Matsen (2019) to maximum endpoint agreement graphs (MAEGs) for networks. We showed that MAEGs induce a metric, called endpoint agreement distance, which bounds the PR-distance naturally from below and with a factor of three from above. Furthermore, we showed that the agreement distance provides bounds on the PR-distance either via the TBR-distance or via its relation to the endpoint agreement distance.

MAFs and MAEFs have been used to develop algorithms that compute the TBR-distance and PR-distance of two trees, respectively. It is thus of interest to see whether MAGs can be utilised to develop approximation algorithms for the agreement distance. Note that such an algorithm would also be an approximation algorithm of the TBR- and the PR-distance.

Acknowledgements

The author would like to thank the New Zealand Marsden Fund for their financial support and the anonymous reviewers for their very helpful comments.

References

- B. L. Allen and M. Steel. Subtree Transfer Operations and Their Induced Metrics on Evolutionary Trees. *Annals of Combinatorics*, 5(1):1–15, 2001. doi:10.1007/s00026-001-8006-8.
- M. L. Bonet and K. St. John. Efficiently Calculating Evolutionary Tree Measures Using SAT. In O. Kullmann, editor, *Theory and Applications of Satisfiability Testing - SAT 2009*, pages 4–17, 2009. ISBN 978-3-642-02777-2. doi:10.1007/978-3-642-02777-2_3.
- M. Bordewich and C. Semple. On the Computational Complexity of the Rooted Subtree Prune and Regraft Distance. *Annals of Combinatorics*, 8(4):409–423, 2005. doi:10.1007/s00026-004-0229-z.

- M. Bordewich and C. Semple. A universal tree-based network with the minimum number of reticulations. *Discrete Applied Mathematics*, 250:357–362, 2018. doi:10.1016/j.dam.2018.05.010.
- M. Bordewich, C. McCartin, and C. Semple. A 3-approximation algorithm for the subtree distance between phylogenies. *Journal of Discrete Algorithms*, 6(3):458–471, 2008. doi:10.1016/j.jda.2007.10.002.
- M. Bordewich, S. Linz, and C. Semple. Lost in space? Generalising subtree prune and re-graft to spaces of phylogenetic networks. *Journal of Theoretical Biology*, 423:1–12, 2017a. doi:10.1016/j.jtbi.2017.03.032.
- M. Bordewich, C. Scornavacca, N. Tokac, and M. Weller. On the fixed parameter tractability of agreement-based phylogenetic distances. *Journal of Mathematical Biology*, 74(1):239–257, 2017b. doi:10.1007/s00285-016-1023-3.
- R. Bouckaert, J. Heled, D. Kühnert, T. Vaughan, C.-H. Wu, D. Xie, M. A. Suchard, A. Rambaut, and A. J. Drummond. BEAST 2: A Software Platform for Bayesian Evolutionary Analysis. *PLOS Computational Biology*, 10(4):1–6, 2014. doi:10.1371/journal.pcbi.1003537.
- J. Chen, J.-H. Fan, and S.-H. Sze. Parameterized and approximation algorithms for maximum agreement forest in multifurcating trees. *Theoretical Computer Science*, 562:496–512, 2015. doi:10.1016/j.tcs.2014.10.031.
- R. Diestel. *Graph Theory*. Springer-Verlag Berlin Heidelberg, 5 edition, 2017. ISBN 978-3-662-53622-3. doi:10.1007/978-3-662-53622-3.
- M. Dunn. Language phylogenies. In C. Bower and B. Evans, editors, *The Routledge Handbook of Historical Linguistics*, chapter 7. Routledge, 2014. doi:10.4324/9781315794013.ch7.
- J. Felsenstein. *Inferring phylogenies*, volume 2. Sinauer Associates, 2004.
- A. Francis, K. T. Huber, and V. Moulton. Tree-Based Unrooted Phylogenetic Networks. *Bulletin of Mathematical Biology*, 80(2):404–416, 2018a. doi:10.1007/s11538-017-0381-3.
- A. Francis, K. T. Huber, V. Moulton, and T. Wu. Bounds for phylogenetic network space metrics. *Journal of Mathematical Biology*, 76(5):1229–1248, 2018b. doi:10.1007/s00285-017-1171-0.
- P. Gambette, L. van Iersel, M. Jones, M. Lafond, F. Pardi, and C. Scornavacca. Rearrangement moves on rooted phylogenetic networks. *PLOS Computational Biology*, 13(8):1–21, 2017. doi:10.1371/journal.pcbi.1005611.
- S. Guindon, J.-F. Dufayard, V. Lefort, M. Anisimova, W. Hordijk, and O. Gascuel. New Algorithms and Methods to Estimate Maximum-Likelihood Phylogenies: Assessing the Performance of PhyML 3.0. *Systematic Biology*, 59(3):307–321, 2010. doi:10.1093/sysbio/syq010.
- D. Gusfield. *ReCombinatorics: the algorithmics of ancestral recombination graphs and explicit phylogenetic networks*. MIT Press, 2014.

- M. Hallett and C. McCartin. A Faster FPT Algorithm for the Maximum Agreement Forest Problem. *Theory of Computing Systems*, 41(3):539–550, 2007. doi:10.1007/s00224-007-1329-z.
- G. Hickey, F. Dehne, A. Rau-Chaplin, and C. Blouin. SPR Distance Computation for Unrooted Trees. *Evolutionary Bioinformatics*, 4:EBO.S419, 2008. doi:10.4137/EBO.S419.
- K. T. Huber, L. van Iersel, R. Janssen, M. Jones, V. Moulton, Y. Murakami, and C. Semple. Rooting for phylogenetic networks. *arXiv preprint arXiv:1906.07430*, 2019.
- D. H. Huson, R. Rupp, and C. Scornavacca. *Phylogenetic networks: concepts, algorithms and applications*. Cambridge University Press, 2010. doi:10.1093/sysbio/syr055.
- R. Janssen and J. Klawitter. Rearrangement operations on unrooted phylogenetic networks. *Theory and Applications of Graphs*, 6(2), 2019. doi:10.20429/tag.2019.060206.
- R. Janssen, M. Jones, P. L. Erdős, L. van Iersel, and C. Scornavacca. Exploring the tiers of rooted phylogenetic network space using tail moves. *Bulletin of Mathematical Biology*, 80(8):2177–2208, 2018. doi:10.1007/s11538-018-0452-0.
- J. Klawitter. The agreement distance of rooted phylogenetic networks. *Discrete Mathematics & Theoretical Computer Science*, 21(3), 2019. doi:10.23638/DMTCS-21-3-19.
- J. Klawitter and S. Linz. On the Subnet Prune and Regraft Distance. *Electronic Journal of Combinatorics*, 26(2):329–355, 2019. doi:10.37236/7860.
- J. I. Meier, D. A. Marques, S. Mwaiko, C. E. Wagner, L. Excoffier, and O. Seehausen. Ancient hybridization fuels rapid cichlid fish adaptive radiations. *Nature Communications*, 8:14363, 2017. doi:10.1038/ncomms14363.
- R. D. M. Page. On Islands of Trees and the Efficacy of Different Methods of Branch Swapping in Finding Most-Parsimonious Trees. *Systematic Biology*, 42(2):200–210, 1993. doi:10.2307/2992542.
- L. H. Rieseberg and J. H. Willis. Plant Speciation. *Science*, 317(5840):910–914, 2007. doi:10.1126/science.1137729.
- E. M. Rodrigues, M.-F. Sagot, and Y. Wakabayashi. The maximum agreement forest problem: Approximation algorithms and computational experiments. *Theoretical Computer Science*, 374(1):91–110, 2007. doi:10.1016/j.tcs.2006.12.011.
- F. Ronquist and J. P. Huelsenbeck. MrBayes 3: Bayesian phylogenetic inference under mixed models. *Bioinformatics*, 19(12):1572–1574, 2003. doi:10.1093/bioinformatics/btg180.
- C. Semple and M. A. Steel. *Phylogenetics*, volume 24. Oxford University Press on Demand, 2003.
- K. St. John. Review Paper: The Shape of Phylogenetic Treespace. *Systematic Biology*, 66(1):e83–e94, 2017. doi:10.1093/sysbio/syw025.
- M. Steel. *Phylogeny: discrete and random processes in evolution*. Society for Industrial and Applied Mathematics, 2016. ISBN 978-1-611974-47-8.

- D. L. Swofford, G. J. Olsen, and P. J. Waddell. Phylogenetic Inference. In D. M. Hillis, C. Moritz, and B. K. Mable, editors, *Molecular Systematics*, chapter 11, pages 407–514. Sinauer Associates, 1996.
- C. M. Thomas and K. M. Nielsen. Mechanisms of, and Barriers to, Horizontal Gene Transfer between Bacteria. *Nature Reviews Microbiology*, 3(9):711–721, 2005. doi:10.1038/nrmicro1234.
- A. Wagner, R. J. Whitaker, D. J. Krause, J.-H. Heilers, M. van Wolferen, C. van der Does, and S.-V. Albers. Mechanisms of gene flow in archaea. *Nature Reviews Microbiology*, 15(8):492–502, 2017. doi:10.1038/nrmicro.2017.41.
- C. Whidden and F. A. Matsen. Quantifying MCMC Exploration of Phylogenetic Tree Space. *Systematic Biology*, 64(3):472–491, 2015. doi:10.1093/sysbio/syv006.
- C. Whidden and F. A. Matsen. Calculating the Unrooted Subtree Prune-and-Regraft Distance. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 16(3):898–911, 2019. doi:10.1109/TCBB.2018.2802911.
- C. Whidden and N. Zeh. A Unifying View on Approximation and FPT of Agreement Forests. In S. L. Salzberg and T. Warnow, editors, *Algorithms in Bioinformatics*, pages 390–402, 2009. doi:10.1007/978-3-642-04241-6_32.
- C. Whidden, R. G. Beiko, and N. Zeh. Fixed-Parameter Algorithms for Maximum Agreement Forests. *SIAM Journal on Computing*, 42(4):1431–1466, 2013. doi:10.1137/110845045.
- Y. Wu. A practical method for exact computation of subtree prune and regraft distance. *Bioinformatics*, 25(2):190–196, 2009. doi:10.1093/bioinformatics/btn606.
- Y. Yu, R. M. Barnett, and L. Nakhleh. Parsimonious Inference of Hybridization in the Presence of Incomplete Lineage Sorting. *Systematic Biology*, 62(5):738–751, 2013. doi:10.1093/sysbio/syt037.
- Y. Yu, J. Dong, K. J. Liu, and L. Nakhleh. Maximum likelihood inference of reticulate evolutionary histories. *Proceedings of the National Academy of Sciences*, 111(46):16448–16453, 2014. doi:10.1073/pnas.1407950111.