

Chow-Liu++: Optimal Prediction-Centric Learning of Tree Ising Models

Enric Boix-Adserà*

Guy Bresler†

Frederic Koehler‡

February 17, 2023

Abstract

We consider the problem of learning a tree-structured Ising model from data, such that subsequent predictions computed using the model are accurate. Concretely, we aim to learn a model such that posteriors $P(X_i|X_S)$ for small sets of variables S are accurate. Since its introduction more than 50 years ago, the Chow-Liu algorithm, which efficiently computes the maximum likelihood tree, has been the benchmark algorithm for learning tree-structured graphical models. A bound on the sample complexity of the Chow-Liu algorithm with respect to the prediction-centric local total variation loss was shown in [9]. While those results demonstrated that it is possible to learn a useful model even when recovering the true underlying graph is impossible, their bound depends on the maximum strength of interactions and thus does not achieve the information-theoretic optimum. In this paper, we introduce a new algorithm that carefully combines elements of the Chow-Liu algorithm with tree metric reconstruction methods to efficiently and optimally learn tree Ising models under a prediction-centric loss. Our algorithm is robust to model misspecification and adversarial corruptions. In contrast, we show that the celebrated Chow-Liu algorithm can be arbitrarily suboptimal.

1 Introduction

Undirected graphical models (also known as Markov random fields) are a flexible and powerful paradigm for modeling high-dimensional data, where variables are associated to nodes in a graph and edges indicate interactions between the variables. One of the primary strengths of graphical models is that the graph structure often facilitates efficient inference computations [47], including computation of marginals, maximum a posteriori (MAP) assignments, and posteriors. Popular algorithms for these tasks include belief propagation and other message-passing algorithms, variational methods, and Markov chain Monte Carlo. Once a graphical model representing the data of interest has been determined, the prototypical use scenario entails observing values for a subset S of the variables, computing posteriors of the form $P(X_i|X_S)$ for i in some set I , and then forming predictions for the variables in I .

In modern applications an approximate graphical model must typically be learned from data, and it therefore makes sense to evaluate the learned model with respect to accuracy of posteriors

*Massachusetts Institute of Technology. Department of EECS. Email: eboix@mit.edu. Supported by an NSF Graduate Fellowship, a Siebel Scholarship, and NSF TRIPODS award 1740751.

†Massachusetts Institute of Technology. Department of EECS. Email: guy@mit.edu. Supported by MIT-IBM Watson AI Lab and NSF CAREER award CCF-1940205.

‡Massachusetts Institute of Technology. Department of Mathematics. Email: fkoeher@mit.edu. This work was supported in part by NSF CAREER Award CCF-1453261, NSF Large CCF-1565235 and Ankur Moitra’s ONR Young Investigator Award.

$P(X_i|X_S)$ [22, 46, 9]. Aside from facilitating computation of posteriors and other inference tasks, the underlying structure of a graphical model is also sometimes of interest in its own right, for instance in certain biological or economic settings. In this context it makes sense to try to reconstruct the graph underlying the model assumed to be generating the data, and most papers on the topic have used the 0-1 loss (equal to zero if the learned graph \hat{G} is equal to the true graph G and one otherwise). In this paper we will be primarily concerned with learning a model such that posteriors $P(X_i|X_S)$ are accurate, even if the structure is not, but we will also deduce results for reconstructing the graph.

The posteriors $P(X_i|X_S)$ of interest often have small conditioning sets S , for instance in a model capturing a user's preferences (like or dislike) for items in a recommendation system, the set S would consist of items for which a user has already provided feedback. It turns out that a *local* version of total variation, as studied in [39, 9], captures accuracy of posteriors $P(X_i|X_S)$ averaged over the values X_S . Specifically, let

$$\text{locTV}^{(k)}(P, Q) = \max_{|A|=k} \text{TV}(P_A, Q_A) \quad (1)$$

be the maximum over sets of variables A of size k of the total variation between the marginals P_A and Q_A . As noted in Section 2.3, this quantity controls accuracy of posteriors $P(X_i|X_S)$ for $|S| \leq k - 1$.

Instead of learning a model close in local total variation, one might use (global) total variation as done in the recent papers [6, 13] which determined the optimal sample complexity for learning tree-structured models to within small total variation. Total variation distance also translates to a guarantee on the accuracy of posteriors $P(X_i|X_S)$, but does not exploit the small size of S . Indeed, as we show in Section D, learning to within small total variation has a high sample complexity of at least $\Omega(n \log n)$ for trees on n nodes, and does not yield meaningful implications in the high-dimensional setting of modern interest (where number of samples is much smaller than number of variables). Furthermore, both [6] and [13] are based on the Chow-Liu algorithm, and as we will show in this paper, the Chow-Liu algorithm can be arbitrarily suboptimal with respect to local total variation. Thus, accuracy with respect to local versus global total variation results in a completely different learning problem with distinct algorithmic considerations.

There is a large literature on structure learning (see [38, 45, 8, 26, 49, 10, 25, 20, 41] and references therein), but prediction-centric learning remains poorly understood. In this paper, we focus on tree Ising models. In this case, the minimax optimal algorithm for structure learning is the well-known Chow-Liu algorithm introduced more than 50 years ago ([11, 9, 44]), which efficiently computes the maximum likelihood tree. Learning tree Ising models with respect to the prediction-centric locTV was studied in [9], where the main take-away was that it is possible to learn a good model even when there is insufficient data to recover the underlying graph. They proved upper and lower bounds on the sample complexity, but left a sizable gap. Their upper bounds were based on analyzing the Chow-Liu algorithm, and left open several core questions: Can the optimal sample complexity for learning trees under the locTV loss be achieved by an efficient algorithm? In particular, the guarantee in [9] depends on the maximum edge strength of the model generating the data – is this necessary? Finally, [9] assumed that the observed data was from a tree Ising model. Can one give sharp learning guarantees under model misspecification, where the learned model is a tree Ising model but the model generating the samples need not be? More broadly, can one robustly learn from data subject to a variety of noise processes, including from adversarially corrupted data? An initial foray towards learning with noise was made in [33], but their approach is based on modifying the techniques of [9], which as discussed below cannot yield sharp results.

1.1 Our Contribution

In this paper we give an algorithm, Chow-Liu++, that is *simultaneously optimal* for both the prediction-centric locTV loss defined in (1) and the structure learning 0-1 loss, and that is also robust to model misspecification and adversarial corruptions. Our algorithm takes as input the n^2 pairwise correlations between variables and has $O(n^2)$ runtime, linear in the input.

Theorem 1.1 (Informal statement of Theorem 3.1). *If P is the distribution of a tree Ising model (T, θ) on n nodes with correlations $\mu_{ij} = \mathbb{E}_P[X_i X_j]$, then when given as input approximate correlations $\tilde{\mu}_{ij}$ satisfying $|\tilde{\mu}_{ij} - \mu_{ij}| \leq \epsilon$ the Chow-Liu++ algorithm outputs a tree Ising model $(\hat{\mathsf{T}}, \hat{\theta})$ with distribution \hat{P} such that $\text{locTV}^{(2)}(\hat{P}, P) \leq C\epsilon$.*

Corollary 1.2. *Given data from a distribution P , Chow-Liu++ learns a tree Ising model with distribution \hat{P} such that $\text{locTV}^{(k)}(\hat{P}, P)$ is within a constant factor of optimal. Moreover*

- (i) *Chow-Liu++ is robust to model misspecification, or equivalently, adversarial corruptions of a fraction of the data.*
- (ii) *In the case that P is a tree Ising model, then Chow-Liu++ returns the correct tree with optimal number of samples up to a constant factor.*

These results, as well as others, are stated more formally in Section 3.

1.2 Techniques and Main Ideas

Our Chow-Liu++ algorithm combines the Chow-Liu algorithm with tree metric reconstruction methods, together with new analysis for both. Our main insight is that the Chow-Liu algorithm and tree metric algorithms work in complementary ranges of interaction strengths – by carefully combining them we obtain an algorithm that works regardless of the minimum and maximum edge strengths. Somewhat counterintuitively, we use the Chow-Liu algorithm indirectly to help us to partition the problem instance and only use actual edges from the algorithm that may well be incorrect. Below, we highlight some of the main challenges and the ideas we introduce to overcome them. We focus on the case of $k = 2$ in the loss $\text{locTV}^{(k)}$, since the result for general $k > 2$ can be obtained by means of an inequality for tree-structured distributions proved in [9] that bounds the $\text{locTV}^{(k)}$ distance in terms of the $\text{locTV}^{(2)}$ distance.

Challenge 1: Failure of Chow-Liu. The main difficulty in proving our result is that the classical Chow-Liu algorithm, which is well-known to be optimal for structure learning and was recently shown to be optimal for total variation [6, 13], does not work for learning to within locTV . Concretely, we show that the Chow-Liu algorithm is not robust to model misspecification – i.e., if the model is not a tree-structured model, but is instead only *close* in locTV to a tree model – then Chow-Liu can fail. We prove this as Theorem 3.3.

Theorem 1.3 (Informal statement of Theorem 3.3). *Chow-Liu fails with model misspecification.*

To prove this theorem we construct a distribution that is close to a tree Ising model in *local* total variation, such that even when given population (infinite sample) values for the correlations the Chow-Liu algorithm makes local errors that accumulate at global scales. This issue does not arise in the analysis of [6, 13], since they assume a model is ϵ -close in total variation to a tree model, and thus the input distribution is extremely close to a tree model in a global sense. Furthermore, in converting this to a finite sample result they assume access to at least $n \log n$ samples, which results in extremely small local estimation errors that do not appreciably accumulate.

Main Idea 1: Structural lemma for Chow-Liu. When there is insufficient data for the Chow-Liu algorithm (or any other algorithm) to reconstruct the correct tree, it turns out that there is still a precise sense in which the output of the Chow-Liu algorithm *resembles* the correct tree. We will momentarily describe the notion of approximation upon which our structural lemma is based, but we first explain specifically which issue it addresses.

Suppose that all edge correlations of the model generating the samples lie in the interval $[\alpha, 1-\beta]$. As shown in [9], learning the structure of tree Ising models requires more samples as α or β get smaller, and no algorithm can reconstruct the exact structure of the tree given $O((\alpha\beta)^{-2} \log n)$ samples. Turning to the locTV loss, our counterexample in Theorem 3.3 shows that strong edges remain problematic for Chow-Liu. However, surprisingly, it turns out that Chow-Liu is useful in a precise sense, even in the presence of weak edges that preclude structure learning, as shown by our structural lemma. See Section 5.2 for more details.

Lemma 1.4 (Structural Lemma for Chow-Liu). *Let T^{CL} be the tree returned by the Chow-Liu algorithm and W be the subset of weak edges with empirical correlation $\leq 1/10$. Let V_1, \dots, V_m denote the connected components of vertices in T^{CL} upon removal of edges in W . Then*

- (i) *The sets of vertices V_1, \dots, V_m are each connected within the true tree T , and*
- (ii) *If we contract the vertex sets V_1, \dots, V_m in T^{CL} , then we obtain the same tree as when we contract these vertex sets in T , as long as all edges have correlation $\Omega(\epsilon)$.*

In other words, the weak edges learned by Chow-Liu correctly partition the vertices into connected components. The subgraphs on these connected components can be badly erroneous, and the weak edges themselves may be between incorrect vertices, but they are between the correct *sets* of vertices. The constant $1/10$ in the definition of weak edges is an arbitrary, albeit sufficiently small, constant. Note also that item (ii) of the structural lemma comes with the qualifier that all edges have correlation at least of order ϵ , where ϵ as in Theorem 3.1 above denotes the accuracy of approximate correlations. The other case turns out to be easy to deal with via a direct argument, because all correlations along paths that include the edge are very weak.

The next two challenges address separately the problem of global error accumulation due to incorrectly placed weak edges and accurate reconstruction within components.

Challenge 2: Accumulation of errors on long paths. The structural lemma that we prove for Chow-Liu has seemingly nothing to do with the objective of reconstructing the tree in $\text{locTV}^{(k)}$ distance – it is merely a combinatorial result. It is necessary to leverage this structural result to resolve the core issue of error accumulation.

Main Idea 2: Error composition theorem for Chow-Liu. We postpone the challenge of accurately reconstructing a model within each component $V_i, i \in [m]$, to be dealt with in a moment. Our composition lemma states that if for each $i \in [m]$ we have a good $\text{locTV}^{(2)}$ reconstruction (T'_i, θ'_i) of the model on the subtree $T[V_i]$ induced by V_i , then we can stitch these together along with the weak edges W of the Chow-Liu tree defined in Lemma 1.4, in order to obtain a good $\text{locTV}^{(2)}$ reconstruction of the overall tree.

Lemma 1.5 (Informal statement of composition lemma for Chow-Liu (Proposition 5.2)). *Suppose that for each $i \in [m]$, (T'_i, θ'_i) is ϵ -close in $\text{locTV}^{(2)}$ distance to the tree model (T, θ) on vertex set V_i . Then the tree model $(\hat{T}, \hat{\theta})$ with tree $\hat{T} = W \cup (\cup_{i \in [m]} T'_i)$ and edge weights $\hat{\theta} = \tilde{\theta}_W \cup (\cup_{i \in [m]} \theta'_i)$ is $O(\epsilon)$ -close in $\text{locTV}^{(2)}$ distance to (T, θ) .*

The proof of this composition result requires careful control over accumulation of errors along paths. Consider vertices u, v and the paths $\hat{\pi}(u, v)$ and $\pi(u, v)$ between these vertices in \widehat{T} and T , respectively. We can decompose each of these paths into subpaths that visit different subtrees on the sets V_1, \dots, V_m . Crucially, the structural lemma guarantees that these decompositions are consistent. For the portion of the path $\hat{\pi}(u, v)$ that lies within a set V_i of the partition, we can bound its error relative to $\pi(u, v)$ by ϵ by using the $\text{locTV}^{(2)}$ guarantee between T'_i and $T[V_i]$ on this subtree. However, a path $\pi(u, v)$ between vertices u and v in T may visit many distinct sets V_i of the partition, but in this case the path must contain weak edges between the sets and this suppresses the errors and counteracts the accumulation that would otherwise occur.

Challenge 3: Dealing with strong edges. The error composition lemma for Chow-Liu reduces the learning problem to reconstructing an accurate tree model on each subset V_i of vertices. Note that the edges W whose removal defined the partition have correlation $\leq 1/10$, so assuming the input correlations have accuracy $\epsilon \leq 1/20$ implies that the true model on V_i has edge strengths lower-bounded by $1/20$. Thus, the reconstruction problem on V_i is potentially simpler. However, our counterexample construction (which has lower-bounded edge strengths) implies that the Chow-Liu algorithm will not succeed even on such models, so we need a different algorithm.

Main Idea 3: Tree metric reconstruction. We show how to adapt and repurpose classical tree metric reconstruction algorithms to solve the learning problem in the case that all edge strengths are lower bounded. There are two central difficulties in applying tree metric reconstruction algorithms. The first is that the tree metric reconstruction problem formulation entails outputting a tree that has *Steiner nodes*, which in graphical models terminology are latent unobserved variables, and our problem formulation requires outputting a tree model without any Steiner nodes. We overcome this difficulty via a novel desteinerization procedure that removes Steiner nodes and is guaranteed to preserve the quality of approximation. The second difficulty is that the sort of approximation guarantee produced by tree metric reconstruction algorithms, i.e., small additive error in evolutionary distance (defined in Section 2.5), is impossible to obtain in our setting. Instead, we devise a variant of the algorithm that obtains a multiplicative approximation to the distance between far away vertices, which turns out to be enough to ensure reconstruction in locTV .

Combining ideas: Chow-Liu++. The Chow-Liu++ algorithm combines Chow-Liu with an approach based on tree metric reconstruction to produce a tree accurate in $\text{locTV}^{(2)}$ distance. Our algorithm is simple and robust to noise. We remark that neither the Chow-Liu algorithm nor the tree metric reconstruction algorithm is sufficient on its own. Recall that the Chow-Liu algorithm fails when there is no upper bound on the strengths of the edges in the tree model (Section 4). We show in Appendix C that the tree reconstruction algorithm fails when there is no lower bound on the strengths of the edges in the tree model, since then the obtained approximation of the evolutionary distance is too weak. It is therefore necessary to combine the two approaches in order to obtain the best of both worlds.

1.3 Related work

Learning in global TV or KL. Several papers have studied the problem of learning graphical models to within small total variation or KL-divergence. As we show in Appendix D, learning tree Ising models to within a given total variation requires $\Omega(n \log n)$ samples. Thus, guarantees for learning to within total variation (or KL-divergence, which by Pinsker's inequality is even more

stringent) are not relevant to the high-dimensional regime (with far fewer samples than number of variables) of interest in modern applications. We summarize the results in this line of work below.

Abbeel et al. [1] considered the problem of learning general bounded-degree factor graphical models to within small KL-divergence and showed polynomial bounds on time and sample complexity. Devroye et al. [14] bounded the minimax (information theoretic) rates for learning various classes of graphical models in total variation. For the case of tree Ising models on n nodes, they showed that given m samples the total variation distance of the learned model from the one generating the samples decays at rate $O\left(\sqrt{\frac{n \log(n)}{m}}\right)$. In Appendix D of the current paper we show this to be tight. The statistical rates analyzed by [14] are achieved by exhaustive search algorithms and the task of finding efficient algorithms remained. Recently Bhattacharyya et al. [6] and Daskalakis and Pan [13] studied exactly this problem, showing that the classical Chow-Liu algorithm achieves the optimal sample complexity $O(n \log n)$. They also showed robustness to model misspecification. As we show in Section 4, the Chow-Liu algorithm can be arbitrary suboptimal under the local total variation loss, and moreover, the technical challenges we outlined in Section 1.2 are unique to local total variation.

Learning trees with noise. A number of papers have recently considered the problem of learning Ising models (both trees and general graphs) with either stochastic or adversarial noise.

In the stochastic noise setting, Nikolakakis et al. [32] give guarantees for structure learning using the Chow-Liu algorithm, when the variables are observed after being passed through a binary symmetric channel. Goel et al. [20] show that an appropriate modification of the Interaction Screening objective from [45] yields an algorithm that can learn the structure of Ising models on general graphs of low width (which generalizes node degree) when each entry of each of the observed samples is missing with some fixed probability. Katiyar et al. [24] observe that learning the structure of tree Ising models with stochastic noise of unknown magnitude is in general non-identifiable, but they characterize a small equivalence class of models and give an algorithm for learning a member of the equivalence class.

Lindren et al. [28] give non-matching lower and upper bounds for learning to within small ℓ_2 error in parameters, where the samples are subjected to adversarial corruptions. As discussed in [26], this also gives the graph structure under natural assumptions on the magnitudes of weights. Prasad et al. [34] consider learning Ising models under Huber's contamination model (in which the adversary is weaker than in the ϵ -corruption model). They restrict attention to the high-temperature regime (where Dobrushin's condition holds), and prove bounds on the ℓ_2 error. Diakonikolas et al. [15] study a noise model with a stronger adversary than Huber's contamination model, but weaker than ϵ -corruption, and again give bounds on the ℓ_2 error of the parameters when the original model satisfies Dobrushin's condition. (Under Dobrushin's condition accuracy of parameters in ℓ_2 translates also to accuracy in total variation.)

Learning for predictions. Wainwright [46] observed that when using an approximate inference algorithm such as loopy belief propagation, learning an incorrect graphical model may yield more accurate predictions than in the true model. Heinemann and Globerson [22] considered the task of learning a model from data such that subsequent inference computations are accurate, which is our goal in this paper. They considered the class of models on graphs with high girth, lower-bounded interaction strength, and also sufficiently weak interactions that the model satisfies correlation decay, in which case belief propagation is approximately correct. They proposed an algorithm which always returns a model from this class and showed finite sample guarantees for its success.

Bresler and Karzand [9] aimed to remove the conditions in [22] on the interaction strengths and

formulated the problem of learning a tree Ising model to within local total variation. They proved guarantees on the performance of the Chow-Liu algorithm which were independent of the minimum edge interaction strength, but were suboptimal in their dependence on maximum interaction strength. Nikolakakis et al. [33] generalized the results of [9] to the setting where each variable is flipped with some known probability q . Their algorithm preprocesses the data to estimate the correlations which are then input to the Chow-Liu algorithm.

Phylogenetic reconstruction. A related area with a different focus is *phylogenetic tree reconstruction*. In this setting, we assume that the data is generated by an unknown tree graphical model and want to estimate the tree topology given access only to the leaves (corresponding to e.g. extant species). The Steel evolutionary metric (see Section 2.5) and the closely related *quartet test* have been used extensively in this literature: see for example [42, 12, 29, 43, 17, 2]. In this setting the main goal is to recover the structure (because it encodes e.g. the history of evolution); furthermore, since most of the nodes in the tree model are latent, it is information-theoretically impossible to achieve meaningful reconstruction in this setting unless edge interactions are sufficiently strong, see e.g. [12]. This is not a concern in our setting because there are no latent nodes. Applications to phylogeny motivated work on the closely related problem of estimating a tree metric (with latent/Steiner nodes) from approximate distances [2, 3, 19]. The most relevant work in this line is [2], which showed how to reconstruct a tree metric given distances within an ℓ_∞ ball; as we explain later, this does not hold for far away nodes when we estimate the evolutionary metric from $O(\log n)$ samples, so we cannot apply their algorithm. A key insight of theirs which we do build on is a connection between tree and ultrametric reconstruction: see Section 7 for further discussion.

1.4 Outline

The rest of the paper is organized as follows. In Section 2 we give necessary background on tree-structured Ising models, formally define the learning problem we consider, and review the Chow-Liu algorithm and tree metric reconstruction. In Section 3 we state our results. We show that the Chow-Liu algorithm is not robust to model misspecification in Section 4. We describe our algorithm, and give the high-level ideas of the proof, in Section 5. Most of the proofs appear in Sections 6 and 7.

Acknowledgment

GB thanks Mina Karzand for many discussions on topics related to those of this paper. We also thank the anonymous reviewers for their helpful and constructive feedback. This work was done in part while the authors were participating in the program *Probability, Geometry, and Computation in High Dimensions* at the Simons Institute for the Theory of Computing.

2 Preliminaries

2.1 Notation

We write $[t] = \{1, 2, \dots, t\}$. Let $\text{Rad}(q)$ denote the distribution such that $X \sim \text{Rad}(q)$ has $\mathbb{P}(X = 1) = q$ and $\mathbb{P}(X = -1) = 1 - q$. The graphs in this paper are all undirected. For an undirected graph $G = (V, E)$ and two vertices $u, v \in V$, we denote by (u, v) the undirected edge between these vertices. For a subset of vertices $U \subseteq V$, $G[U] = (U, E')$ denotes the subgraph of G induced by U ,

where $E' = \{(u, v) \in E : u, v \in U\}$. Also, let $V(G)$ and $E(G)$ denote the vertex set and edge set of G , respectively.

Throughout, ϵ will always be a positive real number less than or equal to 10^{-5} . We use hat, as in \widehat{T} , $\widehat{\theta}$, or $\widehat{\mu}$, to denote quantities associated to an estimator. We use tilde, as in \tilde{d} or $\tilde{\mu}$, to denote approximate quantities typically given as inputs to algorithms.

2.2 Tree-structured Ising models

In this paper, we consider tree-structured Ising models with no external field. It will be convenient to use a nonstandard parameterization in terms of edge correlations. The models are defined by a pair (T, θ) , where T is a tree $T = (\mathcal{V}, E)$ on n nodes and each edge $e \in E$ has corresponding parameter $\theta_e \in [-1, 1]$. Each node $i \in \mathcal{V}$ is associated to a binary random variable $X_i \in \{-1, +1\}$, and each configuration of the variables $x \in \{-1, +1\}^n$ is assigned probability

$$P(x) = \frac{1}{Z} \exp \left(\sum_{(i,j) \in E} \tanh^{-1}(\theta_{ij}) x_i x_j \right) \quad (2)$$

where Z is the normalizing constant. For each edge $(i, j) \in E$ the correlation between the variables at its endpoints is given by $\mathbb{E}_P[X_i X_j] = \theta_{ij}$. One obtains the standard parameterization $P(x) = Z^{-1} \exp(\sum_{(i,j) \in E} J_{ij} x_i x_j)$ upon substituting $\theta_{ij} = \tanh J_{ij}$.

In general, we could have an external field term $\sum_{i \in \mathcal{V}} \theta_i x_i$ in the exponent of (2). The assumption of no external field (i.e., $\theta_i = 0$) implies that the nodewise marginals satisfy $P(X_i = +1) = P(X_i = -1) = 1/2$ for all i . As in [9], this assumption helps to make the analysis tractable and at the same time captures the central features of the problem.

We write μ_{uv} for the correlation $\mathbb{E}_P[X_u X_v]$, and also write $\mu_e = \mu_{uv}$ for an edge $e = (u, v) \in E$. It turns out that in tree models each correlation μ_{uv} is given by the product of correlations along the path connecting u and v . This fact is an easy consequence of the Markov property implied by the factorization (2), and will be used often.

Fact 2.1 (Correlations in tree models). *For any tree Ising model (T, θ) with distribution P as in (2), the correlation μ_{uv} between the variables at nodes u and v is given by*

$$\mu_{uv} = \mathbb{E}_P[X_u X_v] = \prod_{e \in \pi(u,v)} \mu_e = \prod_{e \in \pi(u,v)} \theta_e.$$

Here $\pi(u, v)$ is the path between u and v in T .

Since the correlation between a pair of vertices has at most unit magnitude, Fact 2.1 implies a triangle inequality for the correlations:

Fact 2.2 (Triangle inequality). *For any tree Ising model (T, θ) with distribution P as in (2), and any vertices u, v, w , it holds that $|\mu_{uv}| \leq |\mu_{uw}| |\mu_{vw}|$.*

2.3 Prediction-centric learning

2.3.1 The basic learning problem

We now formally describe the learning problem in the simpler well-specified case. We are given m samples $X^{(1)}, \dots, X^{(m)}$ generated i.i.d. from a tree-structured Ising model (T, θ) on n nodes with distribution P . The goal is to use these samples to construct a tree-structured Ising model $(\tilde{T}, \tilde{\theta})$ with distribution \tilde{P} such that the posterior $\tilde{P}(x_i | x_S) \approx P(x_i | x_S)$ for any $i \in \mathcal{V}$ and assignment

x_S to a small subset of variables S . As discussed in [9], the accuracy of posteriors $\tilde{P}(x_i|x_S)$ for $|S| = k - 1$ is captured by the *local total variation* distance of order k , defined as

$$\text{locTV}^{(k)}(P, Q) := \max_{A:|A|=k} \text{TV}(P_A, Q_A).$$

(Here P_A is the marginal of P on the set of variables in A , and analogously for Q .) From the definition of conditional probability that the absolute deviation of a given posterior averaged over assignments is controlled as

$$\mathbb{E}|P(x_i|X_S) - Q(x_i|X_S)| \leq 2 \cdot \text{locTV}^{(|S|+1)}(P, Q) \quad \text{for any } i \text{ and } S.$$

Therefore, the goal is to learn a tree-structured Ising model \tilde{P} that is close to P in $\text{locTV}^{(k)}$.

In order to achieve this, we will use the following basic facts. One may readily check that $\text{locTV}^{(k)}$ for any k satisfies the triangle inequality. Additionally, it was shown in [9] that for tree models P and Q ,

$$\text{locTV}^{(k)}(P, Q) \leq k2^k \cdot \text{locTV}^{(2)}(P, Q). \quad (3)$$

Finally, the following is useful in order to bound $\text{locTV}^{(2)}$.

Fact 2.3. *The joint distribution of a pair of unbiased ± 1 binary variables x_u, x_v is a function of their correlation given by $P(x_u, x_v) = (1 + x_u x_v \mathbb{E}_P[X_u X_v])/4$. It follows that if the singleton marginals of distributions P and Q over $\{-1, +1\}^n$ are unbiased, then*

$$\text{locTV}^{(2)}(P, Q) = \frac{1}{2} \max_{u, v} |\mathbb{E}_P[X_u X_v] - \mathbb{E}_Q[X_u X_v]|.$$

2.3.2 Learning from pairwise statistics: the ϵ -tree-learning problem

We now consider a general estimation scenario that we call the *ϵ -tree-learning problem*. We emphasize that this problem is a *fully deterministic* algorithmic problem. Our results in both the well-specified and misspecified settings will follow from our result for this more general problem.

Definition 2.4 (ϵ -tree-learning problem). An algorithm is given $0 < \epsilon < 1$ as well as estimates $\tilde{\mu}_{ij}$ for $1 \leq i, j \leq n$ satisfying $|\tilde{\mu}_{ij} - \mu_{ij}| \leq \epsilon$, for some unknown tree Ising model (T, θ) with distribution P having correlations $\mu_{ij} = \mathbb{E}_P[X_i X_j]$. The algorithm is said to *solve the ϵ -tree-learning problem* (with some approximation constant C) if it is guaranteed to output a tree Ising model $(\hat{\mathsf{T}}, \hat{\theta})$ with distribution \hat{P} satisfying $\text{locTV}^{(2)}(\hat{P}, P) \leq C\epsilon$.

It is worth emphasizing that these estimates can be chosen adversarially subject to the accuracy constraint, potentially even in a way that is not realizable by any probability distribution. Note that the problem is well-defined: while in general there can be many tree models P with correlations within ϵ of the given estimates, all are within $\text{locTV}^{(2)}$ of at most 2ϵ of one another by the triangle inequality. Hence it suffices to have the output be close in $\text{locTV}^{(2)}$ to *any* tree Ising model.

What this means is that exhaustively searching over tree models with an appropriate grid over parameters and outputting one with pairwise correlations close to the given $\tilde{\mu}$ estimates constitutes a valid solution of the ϵ -tree-learning problem. Obviously such a procedure is not computationally feasible, and the primary technical contribution of this paper is to give an efficient algorithm for ϵ -tree-learning.

An algorithm for ϵ -tree-learning immediately implies that one can efficiently learn from $O((\log n)/\epsilon^2)$ samples in the PAC model, even under possible model misspecification.

Claim 2.5 (ϵ -tree-learning implies learning under model misspecification). *Let P be an arbitrary distribution over $\{+1, -1\}^{\mathcal{V}}$, and let $\Delta = \inf_{Q \in \mathcal{T}} \text{locTV}^{(k)}(Q, P)$ be the amount of misspecification in the problem, where here the infimum is over tree Ising models. Given an approximation parameter $\epsilon > 0$, the misspecification $\Delta > 0$, and $m = O(\log(|\mathcal{V}|/\delta)/\epsilon^2)$ samples from P , there is a procedure that takes $O(|\mathcal{V}|^2m)$ time and one call to an algorithm solving ϵ -tree-learning with approximation constant C , and with probability at least $1 - \delta$ yields a tree Ising model $(\hat{\mathcal{T}}, \hat{\theta})$ with distribution \hat{P} satisfying*

$$\text{locTV}^{(k)}(\hat{P}, P) \leq Ck2^k(\epsilon + \Delta).$$

Proof. First, from the samples x_1, \dots, x_m compute the empirical correlation $\tilde{\mu}_{u,v} = \frac{1}{m} \sum_{i=1}^m x_{i,u} x_{i,v}$ for each $u, v \in \mathcal{V}$. Condition on the event that $|\tilde{\mu}_{u,v} - \mathbb{E}_P[X_u X_v]| \leq \epsilon$ for all $u, v \in \mathcal{V}$. This occurs with probability $1 - \delta$ by a Chernoff bound and a union bound. Let Q be a tree-structured distribution achieving $\text{locTV}^{(k)}(Q, P) = \Delta$ which exists by compactness of \mathcal{T} . Let $\mu_{u,v}$ be the pairwise correlations of Q and note that for all $u, v \in \mathcal{V}$

$$|\tilde{\mu}_{u,v} - \mu_{u,v}| \leq \epsilon + \Delta/2$$

by Fact 2.3 and the triangle inequality. Run the ϵ -tree-learning algorithm with approximation parameter $\epsilon + \Delta/2$ and inputs $\tilde{\mu}$, to obtain a tree model $(\hat{\mathcal{T}}, \hat{\theta})$ with distribution \hat{P} that satisfies the guarantee $\text{locTV}^{(2)}(\hat{P}, Q) \leq C(\epsilon + \Delta/2)$. By (3), this means $\text{locTV}^{(k)}(\hat{P}, Q) \leq Ck2^k(\epsilon + \Delta/2)$, so by triangle inequality $\text{locTV}^{(k)}(\hat{P}, P) \leq Ck2^k(\epsilon + \Delta/2) + \Delta \leq Ck2^k(\epsilon + \Delta)$, since $Ck2^k \geq 2$. \square

We remark that the amount of misspecification Δ does not need to be known, since we can essentially run an exponential search with an extra runtime factor overhead of $O(\log(1/\Delta))$, and lose only a constant factor in the quality of the reconstruction.

Remark 2.6 (Distributional Robustness Interpretation). The proof of Claim 2.5 has a natural interpretation in terms of *distributionally robust optimization* (DRO), see e.g. [48, 27, 18]. We started with the empirical measure P_m given by m samples from P and observed that P lies within the *ambiguity set*

$$\mathcal{P} = \{R : \text{locTV}^{(2)}(P_m, R) \leq \epsilon, \inf_{Q \in \mathcal{T}} \text{locTV}^{(k)}(Q, R) \leq \Delta\}.$$

Since the true measure P is unknown, the DRO approach would consider the minimax optimization

$$\inf_{\hat{P} \in \mathcal{T}} \sup_{R \in \mathcal{P}} \text{locTV}^{(k)}(R, \hat{P}). \quad (4)$$

The assumptions guarantee, by triangle inequality, that the minimizer \hat{P} of (4) has objective value at most $O(k2^k\epsilon + \Delta)$. However, because the set of tree measures \mathcal{T} is nonconvex (e.g. as a subset of the space of all probability measures), there is no obvious computationally efficient way to solve this minimization. Instead of solving it directly, we observed that the combination of an ϵ -learning algorithm and the triangle inequality enable us to compute an approximate minimizer \hat{P} satisfying

$$\sup_{R \in \mathcal{P}} \text{locTV}^{(k)}(R, \hat{P}) \leq Ck2^k(\epsilon + \Delta)$$

which in particular implies the result when we plug in P for R . Important to our application, the uncertainty parameter ϵ can be taken to depend only logarithmically in the dimension $|\mathcal{V}|$ while containing the true measure P ; this contrasts with the usual “curse of dimensionality” which requires ϵ to have an exponential dependence on the dimension (e.g. for Wasserstein uncertainty, see the discussion in [7]).

2.4 The Chow-Liu algorithm

The Chow-Liu algorithm is a classical algorithm that finds the maximum-likelihood tree model given observations. In the case of Ising models with no external field that we consider, it can be motivated based on the following observation.

Claim 2.7. *Given a tree model (T, θ) with correlations μ , T is a maximum weight spanning tree for weights $|\mu|$.*

Proof. T is a maximum spanning tree if for any vertices u, v in T , all of the edges e on the path $\pi(u, v)$ from u to v have correlation $|\mu_e| \geq |\mu_{u,v}|$. This is the case, since $|\mu_{u,v}| = \prod_{e \in \pi(u,v)} |\mu_e| \leq \min_{e \in \pi(u,v)} |\mu_e|$. \square

Based on this observation, given estimates $\tilde{\mu}$ for the correlations, it makes sense to try to learn a tree based on the max weight spanning tree with weights $|\tilde{\mu}|$, and indeed as shown in [9] the Chow-Liu algorithm (which computes the maximum likelihood tree Ising model with no external field) does exactly this. We take this as the definition.

Definition 2.8. Given approximate correlations $\tilde{\mu}$, the Chow-Liu tree T^{CL} is a maximum spanning tree for weights $|\tilde{\mu}|$, breaking ties arbitrarily.

As we will show, the Chow-Liu algorithm does not solve the ϵ -tree-learning problem, as it can err when there are edges with strong correlation in the true model. However, it will be an important ingredient of our final algorithm, **Chow-Liu++**.

2.5 Evolutionary distance and tree metric reconstruction

In the Chow-Liu++ algorithm, we will also use a connection between tree models and tree metrics via the insight of Steel [42], that on all tree models there is a natural measure of distance between the vertices called the *evolutionary distance* which, furthermore, can be estimated from samples without knowing the tree topology.

Definition 2.9 (Evolutionary distance). In the special case of interest (pairwise interactions with unbiased binary spins), the evolutionary distance of a tree model (T, Θ) with distribution P is simply

$$d(u, v) := -\log \mathbb{E}_P[X_u X_v] = -\log \mu_{uv}.$$

It is well-known that the evolutionary distance is a tree metric, defined formally as follows.

Definition 2.10 (Tree Metric). Let T be a tree with vertex set \mathcal{V} . We say that $d : \mathcal{V} \times \mathcal{V} \rightarrow \mathbb{R}_{\geq 0}$ is a tree metric over T if for any $u, v \in \mathcal{V}$ and $u = u_0, \dots, u_k = v$ the unique path from u to v in T that $d(u, v) = \sum_{i=0}^{k-1} d(u_i, u_{i+1})$. Equivalently, d is the shortest path metric when edges $(u, v) \in E(T)$ are weighted by $d(u, v)$.

We will also consider tree metrics that are allowed to have additional Steiner nodes (i.e. latent nodes in the graphical models terminology). The metric associated to a tree with Steiner nodes is called an additive metric.

Definition 2.11 (Additive metric). Let \mathcal{V} be a finite set. We call $d : \mathcal{V} \times \mathcal{V} \rightarrow \mathbb{R}_{\geq 0}$ an *additive metric* if it is the restriction to \mathcal{V} of some tree metric $d' : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ over a tree with vertex set $\mathcal{X} \supseteq \mathcal{V}$. The nodes in $\mathcal{X} \setminus \mathcal{V}$ are called *Steiner nodes*.

3 Results

Our main result is the following theorem, the main ideas of which are presented in Section 5.

Theorem 3.1 (Guarantee for Chow-Liu++). *Let P be the distribution of a tree Ising model (T, θ) on n binary ± 1 variables X_1, \dots, X_n with correlations $\mu_{ij} = \mathbb{E}_P[X_i X_j]$. Given $0 < \epsilon < 10^{-5}$ and estimates $\tilde{\mu}_{ij}$ for $1 \leq i, j \leq n$ satisfying $|\tilde{\mu}_{ij} - \mu_{ij}| \leq \epsilon$, Chow-Liu++ runs in $O(n^2)$ time and solves the ϵ -tree-learning problem, i.e., there is a universal constant C such that it returns a tree Ising model $(\hat{\mathsf{T}}, \hat{\theta})$ with distribution \hat{P} such that*

$$\text{locTV}^{(2)}(\hat{P}, P) \leq C\epsilon.$$

By Claim 2.5, this theorem implies the following oracle inequality result (in terms of $\text{locTV}^{(k)}$) for estimation from i.i.d. samples. We emphasize, however, that our main theorem applies even in the case that the samples are not generated i.i.d. according to *any* distribution.

Corollary 3.2 (Oracle inequality for $\text{locTV}^{(2)}$). *Let P be an arbitrary distribution on n binary ± 1 variables. Let $\Delta = \inf_{Q \in \mathcal{T}} \text{locTV}^{(2)}(Q, P)$, where the infimum is over tree Ising models. There is a numerical constant $C > 0$ such that given $m \geq C \log(n/\delta)/\epsilon^2$ i.i.d. samples from P and the misspecification error Δ , with probability at least $1 - \delta$, there is an algorithm that runs in $O(n^2m)$ time and returns a tree Ising model $(\hat{\mathsf{T}}, \hat{\theta})$ with distribution \hat{P} satisfying*

$$\text{locTV}^{(k)}(\hat{P}, P) \leq Ck2^k(\Delta + \epsilon).$$

The sample-complexity dependence on n and ϵ is optimal up to constant factors, by Corollary 3.6 of [9], and the runtime complexity matches that of the Chow-Liu algorithm. Furthermore, we remark that the result can be adapted so that the amount of model misspecification Δ does not need to be known, as we can run exponential search over Δ , and only invoke an extra runtime factor of $O(\log(1/\Delta))$.

Theorem B.1 in Appendix B shows that sufficient accuracy in $\text{locTV}^{(3)}$ implies correct structure recovery; thus the learning for predictions problem *generalizes* the structure learning problem. This is used to show Corollary B.3 which states that Chow-Liu++ is optimal also for structure learning. Thus, given data from a tree Ising model, Chow-Liu++ learns the correct structure if that is possible, and regardless it learns a model that yields good predictions.

We also prove in the next section the surprising result that the Chow-Liu algorithm fails under model misspecification, even when given population values for the correlations.

Theorem 3.3. *For any $\delta > 0$, there is a distribution P such that*

- (i) *there is a tree-structured distribution P' with $\text{locTV}^{(2)}(P', P) \leq \delta$, and*
- (ii) *the Chow-Liu algorithm, given population values for the correlations, outputs a tree T^{CL} such that all T^{CL} -structured distributions Q have $\text{locTV}^{(2)}(Q, P) \geq 0.05$.*

4 The Chow-Liu Algorithm Fails with Model Misspecification

In this section, we establish Theorem 3.3, which shows a surprising brittleness of the Chow-Liu algorithm to model misspecification. In particular, suppose that the distribution P is not a tree-structured model, but is δ -close in $\text{locTV}^{(2)}$ distance to a tree-structured model P' for some small $\delta > 0$. We show that the Chow-Liu algorithm does not achieve nontrivial reconstruction guarantees even when run with infinitely many samples from P (i.e., with the population values for the correlations).

Informally, the reason for this failure is that the greedy Chow-Liu algorithm can be tricked into taking too many shortcuts when constructing its tree, resulting in significant distortion of the correlation between distant nodes. Note that we prove a rather strong failure of the Chow-Liu algorithm: not only is the specific distribution prescribed by the Chow-Liu algorithm (given by a tree and edge weights) a poor approximation to the one generating the data, but so too is *any* tree-structured distribution that has the Chow-Liu tree structure.

The proof of Theorem 3.3 is as follows. The construction of P is given in Section 4.1. Its closeness to a tree-structured distribution is shown in Section 4.2. The failure of the Chow-Liu algorithm is proved in Section 4.3.

4.1 Construction of distribution P

Fix a parameter $0 < \delta < 0.1$ and let $n = \lceil 1/\delta \rceil$. Let $P = P_\delta$ denote the distribution over the variables $(X_1, \dots, X_n, Y_1, \dots, Y_n)$ that is sampled as follows. Let A_1, \dots, A_n, B, C be independent $\{+1, -1\}$ -valued random variables such that for all $i \in [n]$ we have $\mathbb{E}[A_i] = e^{-\delta}$, $\mathbb{E}[B] = e^{-2\delta}$ and $\mathbb{E}[C] = 0$. Now let

$$X_i = C \prod_{t=1}^i A_t \quad \text{and} \quad Y_i = BC \prod_{t=1}^i A_t \quad \text{for all } i \in [n].$$

By a straightforward calculation for each $i, j \in [n]$ the first two moments are:

$$\begin{aligned} \mathbb{E}_P[X_i] &= \mathbb{E}_P[Y_i] = 0, \\ \mathbb{E}_P[X_i X_j] &= \mathbb{E}_P[Y_i Y_j] = \exp(-\delta|i - j|), \\ \mathbb{E}_P[X_i Y_j] &= \exp(-\delta|i - j| - 2\delta). \end{aligned}$$

4.2 Closeness to tree-structured distribution P'

We construct a tree-structured distribution P' over the variables $(X_1, \dots, X_n, Y_1, \dots, Y_n)$ as follows. Let W, Z_1, \dots, Z_n be independent $\{+1, -1\}$ -valued random variables such that for all $i \in [n]$ we have $\mathbb{E}[Z_i] = e^{-\delta}$ and $\mathbb{E}[W] = 0$. Let

$$X_i = Y_i = W \prod_{t=1}^i Z_t.$$

This can be represented by a tree-structured distribution, where the tree structure is a path in order $X_1, Y_1, X_2, Y_2, \dots, X_n, Y_n$ (see panel (a) in Figure 1). For all $i, j \in [n]$, the first two moments are:

$$\begin{aligned} \mathbb{E}_{P'}[X_i] &= \mathbb{E}_{P'}[Y_i] = 0, \\ \mathbb{E}_{P'}[X_i X_j] &= \mathbb{E}_{P'}[X_i Y_j] = \mathbb{E}_{P'}[Y_i Y_j] = \exp(-\delta|i - j|). \end{aligned}$$

Therefore, the correlations of P and P' match up to error 2δ . For any $i, j \in [n]$:

$$\mathbb{E}_P[X_i X_j] = \mathbb{E}_{P'}[X_i X_j], \quad \mathbb{E}_P[Y_i Y_j] = \mathbb{E}_{P'}[Y_i Y_j]$$

$$|\mathbb{E}_P[X_i Y_j] - \mathbb{E}_{P'}[X_i Y_j]| = |e^{-\delta|i - j| - 2\delta} - e^{-\delta|i - j|}| \leq 1 - e^{-2\delta} \leq 2\delta$$

Since the correlations under P and P' are 2δ -close and the variables are unbiased, Fact 2.3 implies

$$\text{locTV}^{(2)}(P, P') \leq 2\delta. \tag{5}$$

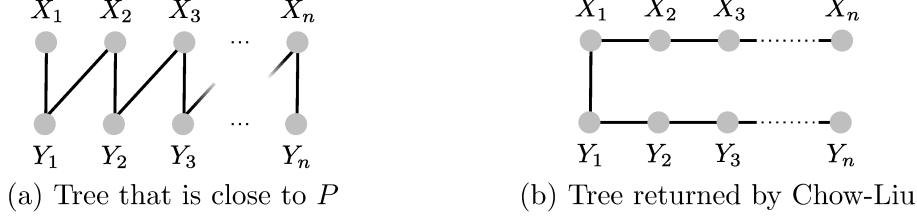


Figure 1: Example capturing the failure of Chow-Liu under model misspecification. The true distribution P is within $\text{locTV}^{(2)}$ distance δ of a tree model P' structured as in (a), however Chow-Liu returns a tree model (b) for which all parameter choices result in $\Omega(1)$ distance in $\text{locTV}^{(2)}$ to P . We remark that the picture (b) is only one of the possible trees returned by Chow-Liu, since there are many ties in our construction, but we show that all of them fail.

4.3 Failure of the Chow-Liu algorithm

We now show that the Chow-Liu algorithm, when given as input the population correlations under P , outputs a tree T^{CL} such that all T^{CL} -structured distributions are 0.1-far from P in $\text{locTV}^{(2)}$ distance. This shows the failure of the Chow-Liu algorithm. First, define the set $S = \{(X_i, X_{i+1})\}_{i \in [n-1]} \cup \{(Y_i, Y_{i+1})\}_{i \in [n-1]}$. Note that all pairs of variables in S have correlation $e^{-\delta}$ under P , and all other pairs have correlation at most $e^{-2\delta}$. Since the edges in S yield no cycles, the Chow-Liu (max-weight spanning) tree T^{CL} contains the edges S . Thus, the Chow-Liu tree T^{CL} consists of the paths X_1, \dots, X_n and Y_1, \dots, Y_n as well as one edge (X_i, Y_j) between these paths.

Fix any T^{CL} -structured distribution Q . We wish to show that $\text{locTV}^{(2)}(Q, P) \geq 0.05$. Suppose that $i \leq n/2$ (a symmetric argument using $\mathbb{E}_Q[X_1 Y_1]$ works if $i > n/2$):

$$\begin{aligned}
 |\mathbb{E}_Q[X_n Y_n]| &= |\mathbb{E}_Q[X_n X_i] \mathbb{E}_Q[X_i Y_n]| && \text{using Fact 2.1} \\
 &\leq |\mathbb{E}_Q[X_n X_i]| \\
 &\leq |\mathbb{E}_P[X_n X_i]| + 2\text{locTV}^{(2)}(Q, P) && \text{using Fact 2.3} \\
 &\leq e^{-\delta n/2} + 2\text{locTV}^{(2)}(Q, P), && \text{using } i \leq n/2 \text{ and Fact 2.3}
 \end{aligned}$$

In the other direction, again by Fact 2.3

$$|\mathbb{E}_Q[X_n Y_n]| \geq |\mathbb{E}_P[X_n Y_n]| - 2\text{locTV}^{(2)}(Q, P) = e^{-2\delta} - 2\text{locTV}^{(2)}(Q, P),$$

so putting these inequalities together

$$\text{locTV}^{(2)}(Q, P) \geq \frac{e^{-2\delta} - e^{-\delta n/2}}{4} > 0.05,$$

by our assumption that $n \geq 1/\delta$ and $\delta < 0.1$. □

5 Algorithm

In this section, we present our robust algorithm, Chow-Liu++, for prediction-centric learning. The formal proof of correctness is given in this section, but most of the details are deferred to Sections 6, 7, and 8.

5.1 Overview

The two most popular classical approaches for reconstructing tree models under noisy data observations are (i) the Chow-Liu algorithm, and (ii) tree metric reconstruction based on the evolutionary distance. Surprisingly, neither of these algorithms solves the ϵ -tree-learning problem by itself. As we have shown in Section 4, the Chow-Liu algorithm fails when the edges are allowed to be arbitrarily strong. And, on the other hand, as we show in Section C, the tree metric reconstruction approach fails when the edge strengths are allowed to be arbitrarily weak.

However, by a stroke of luck, the algorithms happen to be effective in contrasting edge strength regimes. The Chow-Liu algorithm is known to work when all edge strengths are *upper-bounded* by a constant [9]¹. And we prove that a tree metric reconstruction approach works when all edge strengths are *lower-bounded* by a constant. By stitching these algorithms together, we obtain the best of both worlds: a single algorithm, **Chow-Liu++**, that succeeds for models with arbitrary edge weights.

For simplicity, we present our algorithm in the restricted case that the model is ferromagnetic (the edge correlations are nonnegative). The following proposition, proved in Appendix A shows that restricting to ferromagnetic models comes at no loss of generality:

Proposition 5.1. *The ϵ -tree-learning problem for general models on n vertices can be solved in $O(n^2)$ time and one call to an oracle LEARNFERROMODEL for the ϵ -tree-learning problem restricted to ferromagnetic models.*

In Section 5.2, we present our algorithm LEARNFERROMODEL, which solves the reconstruction problem for ferromagnetic models, and is based off of the Chow-Liu algorithm. This algorithm has a key subroutine, LEARNLWRBDDMODEL, which we present in Section 5.3. The subroutine LEARNLWRBDDMODEL is inspired by the tree metric reconstruction literature, and learns a model assuming that the model's edge strengths are lower-bounded by a constant. The results in Sections 5.2 and 5.3 are summarized by:

Proposition 5.2. *LEARNFERROMODEL($\tilde{\mu}, \epsilon$) solves the ferromagnetic ϵ -tree-learning problem in $O(n^2)$ time and $m = O(n)$ calls to LEARNLWRBDDMODEL. Furthermore, letting t_1, \dots, t_m denote the number of input vertices in the calls to LEARNLWRBDDMODEL, we have $\sum_{i=1}^m t_i \leq n$.*

Proposition 5.3. *The algorithm LEARNLWRBDDMODEL solves the ϵ -tree-learning problem on tree models with n variables with edge correlations $\geq 1/20$ in time $O(n^2)$.*

Combining these propositions gives the **Chow-Liu++** algorithm, and allows us to formally prove Theorem 3.1.

Proof of Theorem 3.1. By Proposition 5.3, the LEARNLWRBDDMODEL algorithm is guaranteed to solve the ϵ -tree-learning problem in $O(n^2)$ time when the edge weights are lower-bounded by $1/20$. By Proposition 5.2, we can use it as a subroutine of the LEARNFERROMODEL algorithm to solve the ϵ -tree-learning problem in $O(n^2)$ time for ferromagnetic models. And by Proposition 5.1 we can use LEARNFERROMODEL as a subroutine of LEARNMODEL (i.e., the **Chow-Liu++** algorithm) given in Appendix A to solve the ϵ -tree-learning problem in $O(n^2)$ time for general models. \square

5.2 Chow-Liu strikes back: handling weak edges

We present our algorithm LEARNFERROMODEL, which learns any ferromagnetic model (i.e., any model with arbitrary nonnegative edge strengths). The overall scheme of this algorithm is to (i) use

¹Our proof of this fact is different from the proof given by [9] and we can handle model misspecification.

the Chow-Liu algorithm to break the learning problem into small subproblems, each corresponding to a subtree with lower-bounded edge strengths, and (ii) solve each subproblem with the subroutine `LEARNLWRBDDMODEL`. In this scheme, the Chow-Liu algorithm handles the weak edges of the model, and the `LEARNLWRBDDMODEL` subroutine handles the strong edges. A surprising takeaway is that Chow-Liu algorithm redeems itself – the failure of Chow-Liu proved in Section 4 is due only to strong edges in the tree model, as Chow-Liu can handle the weak edges well. Pseudocode for `LEARNFERROMODEL` is given in Algorithm 1.

Algorithm 1: `LEARNFERROMODEL`($\tilde{\mu}, \epsilon$)

Input: $0 < \epsilon < 10^{-5}$
Output: ϵ -approximate correlations $\tilde{\mu}$ for ferromagnetic model (\mathbf{T}, θ)
Model: $(\widehat{\mathbf{T}}, \widehat{\theta})$ that $O(\epsilon)$ -approximates (\mathbf{T}, θ) in $\text{locTV}^{(2)}$ distance.

```

// Run Chow-Liu
1  $\mathbf{T}^{\text{CL}} \leftarrow \text{MaximumSpanningTree}(\tilde{\mu})$ 
    // Use the result to partition the vertices  $\mathcal{V}$  into  $\sqcup_i V_i$ 
2  $W \leftarrow \{e \in E(\mathbf{T}^{\text{CL}}) : \tilde{\mu}_e \leq 1/10\}$ 
3  $V_1, \dots, V_m \leftarrow \text{ConnectedComponents}(\mathbf{T}^{\text{CL}} \setminus W)$ .
    // Learn a tree model on each set  $V_i$ 
4  $(\widehat{\mathbf{T}}^{(i)}, \theta^{(i)}) \leftarrow \text{LEARNLWRBDDMODEL}(\{\tilde{\mu}_{uv} : u, v \in V_i\}, \epsilon)$  for each  $i \in [m]$ 
    // Stitch together the learned models with the weak edges  $W$  from Chow-Liu
5  $(\widehat{\mathbf{T}}, \widehat{\theta}) \leftarrow ((\sqcup_i \widehat{\mathbf{T}}^{(i)} \cup W), (\sqcup_i \theta^{(i)} \cup \tilde{\mu}_W))$ 
6 Return  $(\widehat{\mathbf{T}}, \widehat{\theta})$ 

```

In words, `LEARNFERROMODEL` proceeds as follows. We first run Chow-Liu, obtaining the tree \mathbf{T}^{CL} . Denoting by W the weak edges in \mathbf{T}^{CL} with empirical correlations $\leq 1/10$, we partition the vertices \mathcal{V} into sets V_1, \dots, V_m according to the connected components in the forest with edges $E(\mathbf{T}^{\text{CL}}) \setminus W$. Now, perhaps counterintuitively, we keep only the weak edges W in the Chow-Liu tree, and discard the stronger edges whose only purpose was to define the groups of relatively strongly dependent variables. Crucially, we can prove that each of these vertex sets V_i induces a subtree $\mathbf{T}[V_i]$ of the original model (or of any model close in $\text{locTV}^{(2)}$ in the misspecified case), with edge weights lower-bounded by $1/10 - \epsilon \geq 1/20$. We then reconstruct the model separately on each subtree $\mathbf{T}[V_i]$ using the `LEARNLWRBDDMODEL` procedure, taking advantage of the lower bound on the edge strength within the subtree. Finally, we glue together the reconstructed trees on vertex sets V_1, \dots, V_m using the weak edges W from the Chow-Liu tree. Figure 2 illustrates the algorithm.

The algorithm's formal correctness guarantee is given in Proposition 5.2. We sketch the proof, deferring details to Section 6. We wish to prove that for any vertices u and v , that the correlation in the reconstructed tree model approximates the true correlation:

$$|\hat{\mu}_{u,v} - \mu_{u,v}| \leq O(\epsilon).$$

Our proof separates the problem into two different cases, depending on the edge strengths of the path $\pi(u, v)$ in \mathbf{T} between u and v . In the first case, all of the edges e of $\pi(u, v)$ have strength at least $\theta_e \geq C\epsilon$ for some large enough constant C ; in the second case there is some edge e with correlation strength $\theta_e < C\epsilon$.

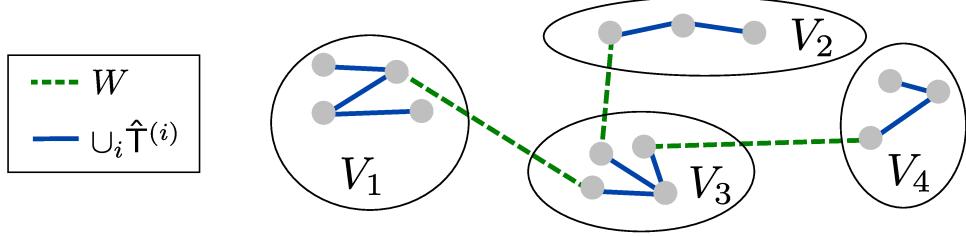


Figure 2: Diagram illustrating the tree \hat{T} returned by LEARNFERROMODEL (Algorithm 1). The green dashed edges are the weak edges W of the Chow-Liu tree. They connect between sets of vertices V_1, \dots, V_m . The solid blue edges within each set V_i are the tree $\hat{T}^{(i)}$ constructed by the LEARNLWRBDDMODEL subroutine run on V_i .

Case 1 Suppose that the edges e in $\pi(u, v)$ all have true correlation $\theta_e \geq C\epsilon$. Then we can prove the following critical structural lemma: the reconstructed path $\hat{\pi}(u, v)$ in \hat{T} between u and v visits sets V_{l_1}, \dots, V_{l_t} in the same order as the path $\pi(u, v)$ in T . Specifically, we can decompose the paths $\pi(u, v)$ and $\hat{\pi}(u, v)$ as a concatenation of paths and edges,

$$\hat{\pi}(u, v) = \hat{\pi}(\hat{a}_1, \hat{b}_1) \circ (\hat{b}_1, \hat{a}_2) \circ \hat{\pi}(\hat{a}_2, \hat{b}_2) \circ \dots \circ (\hat{b}_{t-1}, \hat{a}_t) \circ \hat{\pi}(\hat{a}_t, \hat{b}_t),$$

$$\pi(u, v) = \pi(a_1, b_1) \circ (b_1, a_2) \circ \pi(a_2, b_2) \circ \dots \circ (b_{t-1}, a_t) \circ \pi(a_t, b_t),$$

where for any $i \in [t]$, the paths $\pi(a_i, b_i)$ and $\hat{\pi}(\hat{a}_i, \hat{b}_i)$ stay within the same set V_{l_i} . Here a_i and b_i denote the first and last vertices visited by $\pi(u, v)$ in V_{l_i} . And \hat{a}_i and \hat{b}_i are similarly defined for $\hat{\pi}(u, v)$. (In particular $a_1 = \hat{a}_1 = u$ and $b_t = \hat{b}_t = v$.) See Figure 3 for an illustration.

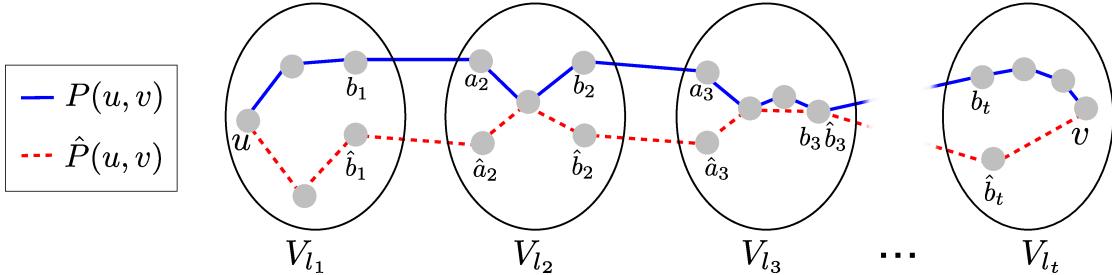


Figure 3: If all edges e in the true path $\pi(u, v)$ have strength at least $C\epsilon$, then the paths $\pi(u, v)$ and $\hat{\pi}(u, v)$ visit the sets of the partition $\mathcal{V} = \sqcup V_i$ in the same order.

Using this decomposition of $\hat{\pi}(u, v)$ and $\pi(u, v)$ into a concatenation of paths and edges, we can write $\hat{\mu}_{u,v}$ and $\mu_{u,v}$ as a product of correlations of the edges and paths in the decomposition:

$$\hat{\mu}_{u,v} = \hat{\mu}_{\hat{a}_1, \hat{b}_1} \hat{\mu}_{\hat{b}_1, \hat{a}_2} \hat{\mu}_{\hat{a}_2, \hat{b}_2} \dots \hat{\mu}_{\hat{b}_{t-1}, \hat{a}_t} \hat{\mu}_{\hat{a}_t, \hat{b}_t}. \quad (6)$$

$$\mu_{u,v} = \mu_{a_1, b_1} \mu_{b_1, a_2} \mu_{a_2, b_2} \dots \mu_{b_{t-1}, a_t} \mu_{a_t, b_t}, \quad (7)$$

For our analysis, we also define the quantity $\Upsilon_{u,v}$ by replacing each $\hat{\mu}$ with μ in equation (6):

$$\Upsilon_{u,v} = \mu_{\hat{a}_1, \hat{b}_1} \mu_{\hat{b}_1, \hat{a}_2} \mu_{\hat{a}_2, \hat{b}_2} \dots \mu_{\hat{b}_{t-1}, \hat{a}_t} \mu_{\hat{a}_t, \hat{b}_t}. \quad (8)$$

We are now ready to prove our goal, which is to show the error bound $|\hat{\mu}_{u,v} - \mu_{u,v}| \leq O(\epsilon)$. We bound the error by the triangle inequality, writing

$$|\hat{\mu}_{u,v} - \mu_{u,v}| \leq |\hat{\mu}_{u,v} - \Upsilon_{u,v}| + |\Upsilon_{u,v} - \mu_{u,v}|.$$

This bounds the error by two terms where each term has a nice intuitive interpretation: (i) the first term $|\hat{\mu}_{u,v} - \Upsilon_{u,v}|$ represents the error caused because the edge weights $\hat{\theta}$ for \hat{T} are based on approximate correlations rather than on the true population correlations, and (ii) the second term $|\Upsilon_{u,v} - \mu_{u,v}|$ represents the error caused because the paths $\hat{\pi}(u, v)$ and $\pi(u, v)$ are different, since the tree \hat{T} might not equal the true tree T . We now bound both terms by $O(\epsilon)$, which will prove the desired result.

First term: We show $|\hat{\mu}_{u,v} - \Upsilon_{u,v}| \leq O(\epsilon)$. For each $i \in [t-1]$, the edge $(\hat{b}_i, \hat{a}_{i+1})$ lies between distinct sets $V_{l_i}, V_{l_{i+1}}$, so it belongs to the set W of weak edges of the Chow-Liu tree. So by construction it has weight equal to the empirical correlation: $\hat{\mu}_{\hat{b}_i, \hat{a}_{i+1}} = \tilde{\mu}_{\hat{b}_i, \hat{a}_{i+1}}$. Since $\tilde{\mu}$ and μ are ϵ -close by assumption, this means

$$|\hat{\mu}_{\hat{b}_i, \hat{a}_{i+1}} - \mu_{\hat{b}_i, \hat{a}_{i+1}}| \leq O(\epsilon).$$

Furthermore, for each $i \in [t]$, recall that the path $\pi(a_i, b_i)$ lies fully in V_{l_i} , so it lies in a subtree of \hat{T} outputted by the reconstruction subroutine LEARNLWRBDDMODEL. By the correctness guarantee of the subroutine (Proposition 5.3), we also have

$$|\hat{\mu}_{\hat{a}_i \hat{b}_i} - \mu_{\hat{a}_i \hat{b}_i}| \leq O(\epsilon).$$

Therefore, comparing formula (6) for $\hat{\mu}_{u,v}$ and formula (8) for $\Upsilon_{u,v}$, we notice that both quantities are a product of $2t-1$ factors, such that each factor in the formula (6) is $O(\epsilon)$ -close to the corresponding factor in the formula (8). Since all factors have magnitude bounded by 1, this immediately implies that $|\hat{\mu}_{u,v} - \Upsilon_{u,v}| \leq O(\epsilon t)$. Unfortunately, this crude bound is insufficient, since we do not have any control on t , the number of vertex sets V_{l_1}, \dots, V_{l_t} visited by the path $\hat{\pi}(u, v)$. The bound fails because it simply adds up the $O(\epsilon)$ approximation error incurred each time we transition from one set V_{l_i} to the next set $V_{l_{i+1}}$.

Fortunately, a more refined argument can show that the errors do not accumulate linearly in t . Notice that for any $i \in [t-1]$ the edge $(\hat{b}_i, \hat{a}_{i+1})$ is one of the weak edges W of the Chow-Liu tree, since \hat{b}_i and \hat{a}_{i+1} are in distinct sets V_{l_i} and $V_{l_{i+1}}$. This means that $\hat{\mu}_{\hat{b}_i, \hat{a}_{i+1}} \leq 1/10$. Hence, $\hat{\mu}_{u,v}$ and $\Upsilon_{u,v}$ decay geometrically in the number of sets t visited by the path:

$$\hat{\mu}_{u,v} \leq (1/10)^t, \quad \Upsilon_{u,v} \leq (1/10 + O(\epsilon))^t.$$

Thus, if $t \geq \log(1/\epsilon)$, then $|\hat{\mu}_{u,v} - \Upsilon_{u,v}| \leq |\hat{\mu}_{u,v}| + |\Upsilon_{u,v}| \leq O(\epsilon)$. Combined with the crude error bound of $O(\epsilon t)$ for $t \leq \log(1/\epsilon)$, this shows that $|\hat{\mu}_{u,v} - \Upsilon_{u,v}| \leq O(\epsilon \log(1/\epsilon))$ independently of t . In the formal proof, we shave off the factor of $\log(1/\epsilon)$ with a more careful analysis and show $|\hat{\mu}_{u,v} - \Upsilon_{u,v}| \leq O(\epsilon)$.

Second term: To show $|\Upsilon_{u,v} - \mu_{u,v}| \leq O(\epsilon)$, we perform a similar analysis. We consider the formula (7) for $\mu_{u,v}$ and the formula (8) for $\Upsilon_{u,v}$ and again compare these formulas factor by factor. However, this analysis is substantially more delicate than the analysis for the first term. A main obstacle is that some factors in (7) may not $O(\epsilon)$ -approximate the corresponding factors of (8). Concretely, we may have the bad case that $|\mu_{a_i, b_i} - \mu_{\hat{a}_i, \hat{b}_i}| \gg \epsilon$ for some $i \in [t]$. To overcome this obstacle, we prove that if this bad case happens, then $\mu_{u,v}$ and $\Upsilon_{u,v}$ decay sufficiently quickly to zero that they are $O(\epsilon)$ -close anyway. We defer further details of the bound of this second term to the formal proof in Section 6.

Case 2 Now suppose that the path $\pi(u, v)$ contains a very weak edge e with strength $\theta_e < C\epsilon$. In this case, we no longer have the key structural lemma that $\hat{\pi}(u, v)$ and $\pi(u, v)$ visit vertex sets V_{l_1}, \dots, V_{l_t} in the same order. For example, if all vertices are perfectly uncorrelated (all edge strengths are 0), then any true tree T is possible, so such a structural result on the reconstruction \hat{T} is impossible to guarantee.

Although we cannot prove the same structural lemma, we can prove that the algorithm still works. We prove that $\hat{\pi}(u, v)$ must also have an edge \hat{e} of strength $\hat{\theta}_{\hat{e}} \leq O(\epsilon)$. Since $\mu_{u,v} \leq \theta_e < C\epsilon$ and $\hat{\mu}_{u,v} \leq \hat{\theta}_{\hat{e}} \leq O(\epsilon)$, then $|\mu_{u,v} - \hat{\mu}_{u,v}| \leq |\mu_{u,v}| + |\hat{\mu}_{u,v}| \leq O(\epsilon)$, as desired.

5.3 Tree metric reconstruction: learning lower-bounded models

It remains to present the subroutine `LEARNLWRBDDMODEL`, which solves the learning problem in the case that the model's edge weights are lower bounded by a constant. Pseudocode is given in Algorithm 2.

Algorithm 2: `LEARNLWRBDDMODEL`($\tilde{\mu}, \epsilon$)

Input: $0 < \epsilon < 10^{-5}$
Input: ϵ -approximate correlations $\tilde{\mu}$ for model (T, θ) with edge strengths $\theta \geq 1/20$
Output: Model $(\hat{T}, \hat{\theta})$ that $O(\epsilon)$ -approximates (T, θ) in $\text{locTV}^{(2)}$ distance.

```

// Construct approximation to evolutionary distance,  $\tilde{d}_{\text{pre}} : \mathcal{V} \times \mathcal{V} \rightarrow [0, \infty]$ 
1  $\tilde{d}_{\text{pre}}(u, v) \leftarrow -\log(\max(0, \tilde{\mu}_{u,v} - \epsilon))$  for all  $u, v \in \mathcal{V}$ 
// Run tree metric reconstruction on evolutionary distance
2  $(\hat{T}, \hat{d}) \leftarrow \text{TREEMETRICRECONSTRUCTION}(\tilde{d}_{\text{pre}}, \log(1/1/20), 200000\epsilon)$ 
// Convert back to Ising model
3  $\hat{\theta}_e \leftarrow \exp(-\hat{d}(e))$  for all edges  $e \in E(\hat{T})$ 
4 Return  $(\hat{T}, \hat{\theta})$ 

```

The algorithm is based on our solution to a particular tree metric reconstruction problem. The connection to tree metrics is via the beautiful insight of Steel [42], that on all tree models there is a natural measure of distance between the vertices called the *evolutionary distance* which can be estimated from samples without knowing the tree topology. Recall from Definition 2.9 that the evolutionary distance between $u, v \in \mathcal{V}$ is given by:

$$d(u, v) = -\log \mu_{u,v}.$$

It is known that d is a tree metric on the tree T (see Definition 2.10), so this suggests applying a tree metric reconstruction algorithm to learn T given an approximation \tilde{d}_{pre} of d . Unfortunately, as we discuss below in Section 5.3.2, classical tree metric reconstruction algorithms are insufficient for this task. Instead we provide our own algorithm, which we call `TREEMETRICRECONSTRUCTION`, to reconstruct the evolutionary distance with the following guarantee.

Theorem 5.4 (Tree Metric Reconstruction). *There exists an absolute constant $C > 0$ such that the following result is true for arbitrary $\epsilon, L > 0$. Let (T, d) be an unknown tree metric on vertex set \mathcal{V} . Suppose that neighboring vertices u, v satisfy $d(u, v) \leq L$, and suppose we have access to a local distance estimate \tilde{d}_{pre} such that*

$$\begin{aligned} d(u, v) &\leq \tilde{d}_{\text{pre}}(u, v), & \text{for all } u, v \in \mathcal{V} \\ \tilde{d}_{\text{pre}}(u, v) &\leq d(u, v) + \epsilon & \text{for all } u, v \in \mathcal{V} \text{ such that } d(u, v) \leq 3L_{\text{upp}} \end{aligned}$$

Given this information, there exists a $O(n^2)$ -time algorithm TREEMETRICRECONSTRUCTION which constructs a tree \widehat{T} equipped with tree metric \widehat{d} such that

$$|\widehat{d}(u, v) - d(u, v)| \leq C(d(u, v)/L + 1)\epsilon \quad (9)$$

for all vertices u, v .

In words, for vertices u, v that are at a low distance $d(u, v) = O(L)$, the reconstructed tree metric \widehat{d} offers a $O(\epsilon)$ -additive approximation to $d(u, v)$. For vertices u, v that are far apart, at distance $d(u, v) = \Omega(L)$, the metric \widehat{d} offers a $(1 + O(\epsilon/L))$ -multiplicative approximation to $d(u, v)$. In Section 5.3.1 we prove that Theorem 5.4 is sufficient to guarantee correctness of the LEARNLWRBDDMODEL algorithm. In Section 5.3.2, we prove Theorem 5.4, relying on results from Sections 7 and 8, and explain how it differs from and extends known tree metric reconstruction algorithms.

5.3.1 Proof of Proposition 5.3, correctness of LearnLwrBddModel

We prove the correctness of LEARNLWRBDDMODEL assuming Theorem 5.4. The runtime is $O(n^2)$ since the TREEMETRICRECONSTRUCTION subroutine runs in $O(n^2)$ time by Theorem 5.4. It remains to prove that the reconstruction $(\widehat{T}, \widehat{\theta})$ returned by the algorithm is accurate in local total variation. In what follows, let

$$L_{\text{upp}} = -\log(1/20).$$

We prove that the function $\tilde{d}_{\text{pre}} : \mathcal{V} \times \mathcal{V} \rightarrow [0, \infty]$ constructed by LEARNLWRBDDMODEL satisfies the following approximation guarantee with respect to d .

Lemma 5.5 (Approximation guarantee of \tilde{d}_{pre}). *Define $\tilde{d}_{\text{pre}} : \mathcal{V} \times \mathcal{V} \rightarrow [0, \infty]$ as in Step 1 of LEARNLWRBDDMODEL. Then*

$$\begin{aligned} d(u, v) &\leq \tilde{d}_{\text{pre}}(u, v) && \text{for all } u, v \\ \tilde{d}_{\text{pre}}(u, v) &\leq d(u, v) + 200000\epsilon && \text{for all } u, v \text{ such that } d(u, v) \leq 3L_{\text{upp}} \end{aligned}$$

Proof. Recall that $|\tilde{\mu}_{u,v} - \mu_{u,v}| \leq \epsilon$ for all $u, v \in \mathcal{V}$. Thus, the first inequality follows from

$$d(u, v) = -\log(\mu_{u,v}) \leq -\log(\max(0, \tilde{\mu}_{u,v} - \epsilon)) \leq \tilde{d}_{\text{pre}}(u, v).$$

For the second inequality, consider u, v such that $d(u, v) \leq 3L_{\text{upp}}$. Then $\mu_{u,v} = \exp(-d(u, v)) \geq \exp(-3L_{\text{upp}}) = (1/20)^3 > 2 \cdot 10^{-5} > 2\epsilon$, so

$$\tilde{d}_{\text{pre}}(u, v) = -\log(\tilde{\mu}_{u,v} - \epsilon) \leq -\log(\mu_{u,v} - 2\epsilon) \leq -\log(\mu_{u,v}) + \frac{2\epsilon}{\mu_{u,v} - 2\epsilon} \leq d(u, v) + 200000\epsilon.$$

The second-to-last inequality comes from $1/(\mu_{u,v} - 2\epsilon)$ -Lipschitzness of $\log(t)$ for $t \geq 1/(\mu_{u,v} - 2\epsilon)$. The last inequality comes from $\mu_{u,v} > (1/20)^3 > 2 \times 10^{-5} > 2\epsilon$. □

We observe that L_{upp} is an upper bound on the edge length in the evolutionary tree metric (T, d) , since each edge has correlation at least $1/20$. Thus, we can plug the approximation guarantee of Lemma 5.5 into Theorem 5.4, and ensure that the tree metric approximation $(\widehat{T}, \widehat{d})$ computed in LEARNLWRBDDMODEL satisfies, for some constant C independent of ϵ :

$$|\widehat{d}(u, v) - d(u, v)| \leq C(d(u, v)/L_{\text{upp}} + 1)\epsilon \leq C(d(u, v) + 1)\epsilon, \quad (10)$$

where we have used that $L_{\text{upp}} \geq 1$ is a uniformly lower-bounded constant. This guarantee turns out to be sufficient to solve the problem of learning in total variation, since the model $(\widehat{T}, \widehat{\theta})$ outputted by LEARNLWRBDDMODEL, which has pairwise correlations $\widehat{\mu}_{u,v} = \exp(-\widehat{d}(u, v))$, is guaranteed to be close to the true model by the following lemma:

Lemma 5.6. *Let $\mu_{uv} = e^{-d(u, v)}$ and $\widehat{\mu}_{uv} = e^{-\widehat{d}(u, v)}$. If $\widehat{d}(u, v)$ satisfies the approximation guarantee (10) with respect to $d(u, v)$, then $|\mu_{uv} - \widehat{\mu}_{uv}| \leq 2C\epsilon$.*

Proof. Suppose that $2C\epsilon < 1$ without loss of generality, since $\mu_{uv}, \widehat{\mu}_{uv} \in [0, 1]$. From $|\mu_{uv} - \widehat{\mu}_{uv}| = e^{-d(u, v)} |e^{d(u, v) - \widehat{d}(u, v)} - 1|$ and the guarantee in (9) we see that

$$e^{-d(u, v)} (e^{-(d(u, v) + 1) \cdot C\epsilon} - 1) \leq \mu_{uv} - \widehat{\mu}_{uv} \leq e^{-d(u, v)} (e^{(d(u, v) + 1) \cdot C\epsilon} - 1).$$

Now we differentiate the RHS to get

$$\frac{d}{dt} (e^{-t} (e^{(t+1)C\epsilon} - 1)) = e^{-t} (1 - (1 - C\epsilon)e^{(t+1)C\epsilon}).$$

This is positive for $t \in [0, \frac{1}{C\epsilon} \log \frac{1}{1-C\epsilon} - 1]$ and is negative for larger t , so it suffices to bound the RHS for $d(u, v) = \frac{1}{C\epsilon} \log \frac{1}{1-C\epsilon} - 1$, and it is easily verified that the RHS is at most $2C\epsilon$.

For the LHS, we differentiate to get

$$\frac{d}{dt} (e^{-t} (e^{-(t+1)C\epsilon} - 1)) = e^{-t} (1 - (1 + C\epsilon)e^{-(t+1)C\epsilon}),$$

and we see that the LHS is decreasing precisely on $[0, \frac{1}{C\epsilon} \log(1 + C\epsilon) - 1]$. Plugging into the LHS we see that the minimum value it can take is at least $(1 + C\epsilon)^{-1} - 1 \geq -C\epsilon$. \square

Since all pairwise correlations are $2C\epsilon$ -close, the model $(\widehat{T}, \widehat{\theta})$ is $2C\epsilon$ -close to (T, θ) in $\text{locTV}^{(2)}$, and Proposition 5.3 follows. \square

5.3.2 Proof of Theorem 5.4, tree metric reconstruction algorithm

Here we present the TREEMETRICRECONSTRUCTION algorithm and prove the correctness guarantee, Theorem 5.4. The pseudocode for TREEMETRICRECONSTRUCTION is given as Algorithm 3. In the first step, ADDITIVEMETRICRECONSTRUCTION (given in Section 7) reconstructs an additive metric approximation. In the second step, DESTEINERIZE (given in Section 8) converts the reconstructed additive metric to a tree metric without Steiner nodes.

Algorithm 3: TREEMETRICRECONSTRUCTION($\tilde{d}_{\text{pre}}, L, \epsilon$)

Input: $\tilde{d}_{\text{pre}}, L, \epsilon$ satisfying the conditions of Theorem 5.4

Output: Tree metric $(\widehat{T}, \widehat{d})$ satisfying approximation guarantee (9)

- 1 $(\widehat{T}', \widehat{d}') \leftarrow \text{ADDITIVEMETRICRECONSTRUCTION}(\tilde{d}_{\text{pre}}, L, \epsilon)$
 - 2 $(\widehat{T}, \widehat{d}) \leftarrow \text{DESTEINERIZE}(\widehat{T}', \widehat{d}')$
 - 3 Return $(\widehat{T}, \widehat{d})$
-

Although there is a rich literature that studies tree metric reconstruction problems, known algorithms do not suffice to obtain the guarantee of Theorem 5.4. The most relevant result from the literature is the landmark algorithm of Agarwala et al. [2] for additive metric approximation

(recall from Definition 2.11 that additive metrics are tree metrics with Steiner nodes). Given a $O(\epsilon)$ -approximation \tilde{d} to an additive metric d , the algorithm outputs an additive metric \hat{d} that $O(\epsilon)$ -approximates d . Two major obstacles arise when applying this result. The first obstacle is that the algorithm of Agarwala et al. outputs an additive metric as opposed to a tree metric. If we convert this additive metric into a tree model, then the Steiner nodes in the additive metric correspond to latent variables in the model. This is undesirable, as we wish to learn a model without latent variables. The second obstacle is that it is impossible to compute a $O(\epsilon)$ -additive approximation to the evolutionary distance between all pairs of nodes. The culprits are far-away nodes: as the correlations decrease to 0, the evolutionary distance grows to infinity, and so a $O(\epsilon)$ -additive approximation in the correlations does not yield a $O(\epsilon)$ -additive approximation in the evolutionary distance.

Nevertheless, although the additive metric reconstruction theorem of Agarwala et al. [2] does not directly apply, our proof of Theorem 5.4 leverages the important connection that Agarwala et al. found between the tree metric reconstruction problem and reconstruction of ultrametrics. Through this connection to ultrametric reconstruction, in Section 7 we derive the algorithm ADDITIVEMETRICRECONSTRUCTION, which obtains an additive metric reconstruction of the evolutionary distance satisfying the approximation guarantee (9), where distances between far-away nodes are multiplicatively-approximated. Then, in Section 8 we show how to convert this additive metric to a tree metric with an algorithm DESTEINERIZE that clusters the Steiner nodes to their closest non-Steiner node. The guarantees that we prove for ADDITIVEMETRICRECONSTRUCTION and DESTEINERIZE are stated below:

Theorem 5.7 (Additive Metric Reconstruction). *There exists a constant $C > 0$ such that the following result is true for arbitrary $\epsilon > 0$ and any $L \geq 100\epsilon$. Let (T, d) be an unknown tree metric on vertex set \mathcal{V} . Suppose that neighboring vertices $u, v \in \mathcal{V}$ satisfy $d(u, v) \leq L$, and suppose we have access to a distance estimate $\tilde{d}_{\text{pre}} : \mathcal{V} \times \mathcal{V} \rightarrow \mathbb{R} \cup \{\infty\}$ such that*

$$\begin{aligned} d(u, v) &\leq \tilde{d}_{\text{pre}}(u, v), & \text{for all } u, v \in \mathcal{V} \\ \tilde{d}_{\text{pre}}(u, v) &\leq d(u, v) + \epsilon & \text{for all } u, v \in \mathcal{V} \text{ such that } d(u, v) \leq 3L_{\text{upp}} \end{aligned}$$

Given this information, there exists a $O(|\mathcal{V}|^2)$ -time algorithm ADDITIVEMETRICRECONSTRUCTION which constructs an additive metric (\hat{T}, \hat{d}) with $|V(\hat{T})| \leq O(|\mathcal{V}|)$ Steiner nodes such that

$$|\hat{d}(u, v) - d(u, v)| \leq C(d(u, v)/L + 1)\epsilon$$

for all $u, v \in \mathcal{V}$.

Theorem 5.8 (Removing Steiner nodes). *There exists a constant $C > 0$ such that the following is true. Let $L \geq 2\epsilon \geq 0$ and suppose (T, d) is an unknown tree metric on vertex set \mathcal{V} with maximum edge length L . Suppose \hat{d} is an additive metric on \mathcal{V} with Steiner tree representation \hat{T} such that*

$$|d(x, y) - \hat{d}(x, y)| \leq (d(x, y)/L + 1)\epsilon$$

for all $x, y \in \mathcal{V}$. Then DESTEINERIZE with (\hat{T}, \hat{d}) as input runs in $O(|\mathcal{V}|^2 + |V(\hat{T})|)$ time and outputs a tree metric (T', d') on \mathcal{V} such that

$$|d(x, y) - d'(x, y)| \leq (d(x, y)/L + 1)\epsilon + C\epsilon$$

for all $x, y \in \mathcal{V}$.

Theorems 5.7 and 5.8 are proved in Sections 7 and 8, respectively. Combined, these two theorems immediately imply the correctness and runtime of TREEMETRICRECONSTRUCTION, Theorem 5.4:

Proof of Theorem 5.4. By Theorems 5.7 and 5.8, TREEMETRICRECONSTRUCTION runs in $O(n^2)$ time. Furthermore, there are constants C', C'' such that if $\epsilon > 0$, $L \geq 100\epsilon$, $L \geq 2C'\epsilon$, then the outputted tree metric satisfies the approximation guarantee

$$|\mathbf{d}(x, y) - \hat{\mathbf{d}}(x, y)| \leq C'(\mathbf{d}(x, y)/L + 1)\epsilon + C'C''\epsilon.$$

Taking $C = \max(100, 2C', C'(C'' + 1))$, we can therefore ensure (9) for any $\epsilon, L > 0$. (Since if $L < 100\epsilon$ or $L < 2C'\epsilon$, outputting $\hat{\mathbf{d}}(x, y) = 0$ for all x, y works.) \square

Remark 5.9. Theorem 5.4 is much easier to prove if we do not require the output $\hat{\mathbf{d}}$ to be a tree metric, since essentially looking at the shortest path metric of $\hat{\mathbf{d}}$ gives a global metric satisfying a guarantee like (9); see Lemma 7.9 for a proof. However, straightforward approaches to extract a tree from this metric — e.g. taking a shortest path tree containing all shortest paths from some arbitrary root node ρ , or a minimum cost spanning tree — are easily seen to fail on a path graph with edge distances $\Theta(\epsilon)$. If we use a minimum cost spanning tree then, in the case where \mathbf{d} is the evolutionary metric, the resulting tree is the Chow-Liu Tree and its failure is explained in Section 4. If we use a shortest path tree started from the first vertex of the path, then the tree can fail in the same way as Chow-Liu: by modifying all distances by $\Theta(\epsilon)$ we can make the shortest path to the final vertex in the path go through every other vertex in the graph, except possibly at the end, making the tree match the Chow-Liu tree in Figure 1 and making distances between some neighbors off by the entire diameter of the graph.

6 Chow-Liu: handling weak edges

This section is devoted to the proof of Proposition 5.2, which shows correctness of the algorithm LEARNFERROMODEL, based on the Chow-Liu algorithm. Section 6.1 formally proves the proposition, using claims whose proofs are deferred to Section 6.2.

6.1 Proof of Proposition 5.2

First, the output $\hat{\mathbf{T}}$ of LEARNFERROMODEL is a tree, because each T'_i is a tree on the set V_i , and by construction the edge set W forms a tree when the sets of the partition V_1, \dots, V_m are contracted. The runtime guarantee of the algorithm is also straightforward, since the maximum spanning tree and connected components can be found in $O(|\mathcal{V}|^2)$ time. It remains to show the $\text{locTV}^{(2)}$ approximation guarantee. By Fact 2.3, this reduces to proving that there is an absolute constant $C > 0$ such that

$$|\hat{\mu}_{uv} - \mu_{uv}| \leq C\epsilon \text{ for all } u, v \in \mathcal{V}.$$

To show this, let $\pi(a, b) \subset \mathbf{T}$ and $\hat{\pi}(a, b) \subset \hat{\mathbf{T}}$ denote the paths from vertex a to vertex b in \mathbf{T} and $\hat{\mathbf{T}}$, respectively. Then consider any pair of vertices $u, v \in \mathcal{V}$. If there is a very weak edge $e \in \pi(u, v)$ such that $\theta_e \leq 4\epsilon$, then Claim 6.6 below shows that there is an edge $e' \in \hat{\pi}(u, v)$ such that $\hat{\theta}_{e'} \leq 5\epsilon$, and hence

$$|\hat{\mu}_{uv} - \mu_{uv}| \leq |\hat{\mu}_{uv}| + |\mu_{uv}| \leq |\hat{\theta}_{e'}| + |\theta_e| \leq 9\epsilon \leq C\epsilon.$$

So without loss of generality $\theta_e > 4\epsilon$ for all edges $e \in \pi(u, v)$. Under this edge strength lower bound condition, Claim 6.5 below shows that $\pi(u, v)$ and $\hat{\pi}(u, v)$ are equal if we contract the vertex sets V_1, \dots, V_m . In other words, we can decompose the paths as

$$\pi(u, v) = \pi(a_1, b_1) \circ (b_1, a_2) \circ \pi(a_2, b_2) \circ \dots \circ (b_{t-1}, a_t) \circ \pi(a_t, b_t),$$

$$\hat{\pi}(u, v) = \hat{\pi}(a'_1, b'_1) \circ (b'_1, a'_2) \circ \hat{\pi}(a'_2, b'_2) \circ \dots \circ (b'_{t-1}, a'_t) \circ \hat{\pi}(a'_t, b'_t),$$

where $a_i, b_i, a'_i, b'_i \in V_{\sigma(i)}$ for all $i \in [t]$ for some permutation $\sigma : [m] \rightarrow [m]$.

Hence,

$$\begin{aligned} \mu_{uv} &= \prod_{i=1}^t \mu_{a_i b_i} \prod_{i=1}^{t-1} \mu_{b_i a_{i+1}}, \\ \hat{\mu}_{uv} &= \prod_{i=1}^t \hat{\mu}_{a'_i b'_i} \prod_{i=1}^{t-1} \hat{\mu}_{b'_i a'_{i+1}}. \end{aligned}$$

In order to bound the difference between these two quantities, we also define

$$\Upsilon_{uv} = \prod_{i=1}^t \mu_{a'_i b'_i} \cdot \prod_{i=1}^{t-1} \mu_{b'_i a'_{i+1}},$$

which is the product along the estimated path in $\hat{\mathbf{T}}$ of the correlations as computed in \mathbf{T} . Now applying the triangle inequality yields

$$|\hat{\mu}_{uv} - \mu_{uv}| \leq |\hat{\mu}_{uv} - \Upsilon_{uv}| + |\Upsilon_{uv} - \mu_{uv}|.$$

Claim 6.7 below shows that $|\hat{\mu}_{uv} - \Upsilon_{uv}| \leq C''\epsilon$ for some absolute constant C'' . The rough idea is that, for each $i \in [t]$, $\hat{\mu}_{a'_i b'_i}$ is an estimated correlation within the partition $V_{\sigma(i)}$ and $\mathbf{T}[V_{\sigma(i)}]$ has all edge weights of magnitude at least $1/20$ by Claim 6.2 below. Hence $\hat{\mu}_{a'_i b'_i}$ is a $C'\epsilon$ -approximation of $\mu_{a'_i b'_i}$ for a universal constant C' by the correctness of LEARNLWRBDDMODEL. By the fact that $\tilde{\mu}$ is an ϵ -approximation of μ , we also know that for each $i \in [t-1]$, $\hat{\mu}_{b'_i a'_{i+1}}$ is a 2ϵ -approximation of $\mu_{b'_i a'_{i+1}}$. Furthermore, as discussed in the overview 5.2, the errors from each term in the product do not accumulate as t increases, but rather add with geometric attenuation, because $\mu_{b'_i a'_{i+1}} \leq 1/10 + \epsilon \leq 1/5$ for all $i \in [t-1]$.

Claim 6.9 shows that $|\Upsilon_{uv} - \mu_{uv}| \leq 4\epsilon$. This claim is proved in two steps. First, it is easy to show that $\Upsilon_{uv} \leq \mu_{uv}$ by noting that $\mu_{ab}\mu_{bc} \leq \mu_{ac}$ for any vertices a, b, c in a tree-structured model. Second, to show that $\mu_{uv} \leq \Upsilon_{uv} + 4\epsilon$, we prove that a_i is close to a'_i and b_i is close to b'_i for all $i \in [t]$. This latter step requires more care – in particular, we use that we may assume that $\theta_e > 4\epsilon$ for all edges $e \in \pi(u, v)$ without loss of generality.

So overall $|\hat{\mu}_{uv} - \mu_{uv}| \leq (C'' + 4)\epsilon \leq C\epsilon$ for an absolute constant C . \square

6.2 Deferred details

The claims in this section fill out the details of the proof of Proposition 5.2.

6.2.1 Correctness within each set V_i

First, we show that within each vertex set V_i the correctness of LEARNLWRBDDMODEL ensures that $\hat{\mu}$ $O(\epsilon)$ -approximates μ .

Claim 6.1 (Correlations are weak between sets). *If $u \in V_i, v \in V_j$ for some $i \neq j$, then $\tilde{\mu}_{uv} \leq 1/10$.*

Proof. Since u and v are not in the same set of the partition, there is an edge e (in the set of weak edges W) on the path from u to v in T^{CL} with $\tilde{\mu}_e \leq 1/10$. If $\tilde{\mu}_{uv} > 1/10 \geq \tilde{\mu}_e$, then removing e and adding (u, v) to T^{CL} gives a spanning tree of strictly higher weight, contradicting maximality of T^{CL} . \square

Recall (T, θ) is a tree model with correlations μ within ϵ of the given approximate values $\tilde{\mu}$.

Claim 6.2 (Each set induces a subtree with lower-bounded edge strengths). *$\mathsf{T}[V_i]$ is a subtree of T for each $i \in [m]$. Furthermore, $\theta_e \geq 1/10 - \epsilon$ for all $e \in E(\mathsf{T}[V_i])$.*

Proof. Assume for the sake of contradiction that $\mathsf{T}[V_i]$ is not connected. Then there is a vertex $w \notin V_i$ and partition $V_i = V_{i,1} \sqcup V_{i,2}$ into nonempty sets such that all paths in T from $V_{i,1}$ to $V_{i,2}$ go through w . So for all $u \in V_{i,1}$ and $v \in V_{i,2}$,

$$\mu_{uv} = \mu_{uw}\mu_{wv} \leq (\tilde{\mu}_{uw} + \epsilon)(\tilde{\mu}_{wv} + \epsilon) \leq (1/10 + \epsilon)(1/10 + \epsilon) \leq 1/20.$$

We have used $\tilde{\mu}_{uw}, \tilde{\mu}_{wv} \leq 1/10$ by Claim 6.1, since $w \notin V_i$, $u, v \in V_i$. Thus $\tilde{\mu}_{uv} \leq \mu_{uv} + \epsilon \leq 1/10$, for all $u \in V_{i,1}$ and $v \in V_{i,2}$. On the other hand, there must be $u \in V_{i,1}, v \in V_{i,2}$ such that $\tilde{\mu}_{uv} > 1/10$ (otherwise the only possible edges between $V_{i,1}$ and $V_{i,2}$ in T^{CL} would be in W and hence V_i would not have ended up as one part), which is a contradiction.

Furthermore, if there is $e \in E(\mathsf{T}[V_i])$ such that $\theta_e < 1/10 - \epsilon$, then let $V_i = V'_{i,1} \sqcup V'_{i,2}$ be the partition of V_i formed by removing e from $\mathsf{T}[V_i]$. For any $u \in V'_{i,1}, v \in V'_{i,2}$, we have $\tilde{\mu}_{uv} \leq \mu_{uv} + \epsilon \leq \theta_e + \epsilon < 1/10$, which by the argument in the previous paragraph contradicts the construction of V_i . \square

Claim 6.3 (Correlations are correct within each set). *There is an absolute constant C' such that if $u, v \in V_i$ for some i , then $|\hat{\mu}_{uv} - \mu_{uv}| \leq C'\epsilon$.*

Proof. Claim 6.2 shows that $\mathsf{T}[V_i]$ is a tree with $\theta_e \geq (1/10 - \epsilon) \geq 1/20$ for each edge e , so Proposition 5.3 applies, showing that LEARNLWRBDDMODEL solves the ϵ -reconstruction problem. In other words, there is an absolute constant C' such that LEARNLWRBDDMODEL returns a tree-structured model (T'_i, θ'_i) such that all correlations satisfy $|\mu'_{i,uv} - \mu_{uv}| \leq C'\epsilon$ for $u, v \in V_i = V(T'_i)$, where $\mu'_{i,uv}$ is the correlation computed in the model (T'_i, θ'_i) . The proof follows by noting that correlations $\mu'_{i,uv}$ computed within each model (T'_i, θ'_i) are exactly the same as the correlations $\hat{\mu}_{uv}$ in the final estimated model $(\hat{\mathsf{T}}, \hat{\theta})$, since $\hat{\mathsf{T}}[V_i]$ is connected and thus correlations within V_i entail only edges within $\hat{\mathsf{T}}[V_i] = T'_i$. \square

This latter claim implies the desired local reconstruction error guarantee when $u, v \in V_i$ for some $i \in [m]$.

6.2.2 Topology roughly correct

In order to translate the error guarantee within each set V_i (proved by Claim 6.3) into an error guarantee between any pair of distinct sets V_i and V_j , we first prove that the topologies of T and $\hat{\mathsf{T}}$ are equal (except possibly very weak edges) when the sets V_i are contracted. In other words, except for very weak edges, the edges in $\hat{\mathsf{T}}$ between sets V_1, \dots, V_m go between the correct sets.

Claim 6.4. *Let $(u, v) \in E(\mathsf{T})$ such that $u \in V_i, v \in V_j$ for distinct i, j . If $\theta_{uv} = \mu_{uv} \geq 4\epsilon$ then there is $(u', v') \in E(\hat{\mathsf{T}})$ such that $u' \in V_i, v' \in V_j$.*

Proof. Consider the path $\hat{\pi}(u, v)$ in $\widehat{\mathbf{T}}$. Let $e = (a, b)$ be the first edge on this path such that a and b are in separate components of $\mathbf{T} \setminus (u, v)$. We show that $a \in V_i$ and $b \in V_j$, proving the theorem. Otherwise, without loss of generality $a \in V_k$ for some $k \neq i$. We show that $\tilde{\mu}_{ab} < \tilde{\mu}_{uv}$:

$$\begin{aligned}
\tilde{\mu}_{ab} &\leq \mu_{ab} + \epsilon \\
&\leq \mu_{au}\mu_{uv}\mu_{vb} + \epsilon && \text{by Fact 2.1, since } (u, v) \in \pi(a, b) \\
&\leq \mu_{au}\mu_{uv} + \epsilon && \text{using } \mu_{vb} \leq 1 \\
&\leq (\tilde{\mu}_{au} + \epsilon)(\tilde{\mu}_{uv} + \epsilon) + \epsilon \\
&\leq (1/10 + \epsilon)(\tilde{\mu}_{uv} + \epsilon) + \epsilon && \text{by Claim 6.1, since } u \in V_i, a \in V_k \text{ for } i \neq k \\
&\leq (1/10 + \epsilon)\tilde{\mu}_{uv} + (11/10 + \epsilon)\epsilon \\
&= \tilde{\mu}_{uv} - (9/10 - \epsilon)\tilde{\mu}_{uv} + (11/10 + \epsilon)\epsilon \\
&< \tilde{\mu}_{uv}, && \text{since } \tilde{\mu}_{uv} \geq \mu_{uv} - \epsilon \geq 3\epsilon.
\end{aligned}$$

We now show that $\mathbf{T}^{\text{CL}'} = \mathbf{T}^{\text{CL}} \cup (u, v) \setminus (a, b)$ is a tree. Since $\mathbf{T}^{\text{CL}'}$ has strictly greater $\tilde{\mu}$ weight than \mathbf{T}^{CL} , this would contradict maximality of the Chow-Liu tree and prove the claim. (i) First, $(a, b) \in W \subset E(\mathbf{T}^{\text{CL}})$, since $(a, b) \in E(\widehat{\mathbf{T}})$ and $\tilde{\mu}_{ab} < \tilde{\mu}_{uv} \leq 1/10$ by Claim 6.1. (ii) Second, u and v are in distinct components of $\mathbf{T}^{\text{CL}} \setminus (a, b)$, since u is connected to a and v is connected to b in $\mathbf{T}^{\text{CL}} \setminus (a, b)$. Indeed, for any edge $f = (c, d) \neq (a, b)$ in $\hat{\pi}(u, v)$ we prove that c and d are connected in $\mathbf{T}^{\text{CL}} \setminus (a, b)$: either $c \in V_l, d \in V_{l'}$ for distinct l, l' , in which case $(c, d) \in W \subset E(\mathbf{T}^{\text{CL}})$; or, if $c, d \in V_l$ for some l , then c and d are connected in $\mathbf{T}^{\text{CL}} \setminus (a, b)$ because V_l is a connected component of $\mathbf{T}^{\text{CL}} \setminus W$ and $(a, b) \in W$. Combined, (i) and (ii) prove that $\mathbf{T}^{\text{CL}'}$ is a tree. \square

This immediately implies the following:

Claim 6.5 (Reconstructed topology correct – not-very-weak edges case). *Let $u, v \in \mathcal{V}$. If all edges e of $\pi(u, v)$ have weight $\theta_e \geq 4\epsilon$, then we can decompose the paths $\pi(u, v) \subseteq \mathbf{T}$, $\hat{\pi}(u, v) \subseteq \widehat{\mathbf{T}}$ as*

$$\pi(u, v) = \pi(a_1, b_1) \circ (b_1, a_2) \circ \pi(a_2, b_2) \circ \cdots \circ (b_{t-1}, a_t) \circ \pi(a_t, b_t),$$

$$\hat{\pi}(u, v) = \hat{\pi}(a'_1, b'_1) \circ (b'_1, a'_2) \circ \hat{\pi}(a'_2, b'_2) \circ \cdots \circ (b'_{t-1}, a'_t) \circ \hat{\pi}(a'_t, b'_t),$$

where $a_i, b_i, a'_i, b'_i \in V_{\sigma(i)}$ for all $i \in [t]$ for some permutation $\sigma : [m] \rightarrow [m]$.

Proof. Since each set V_i induces a subtree $\mathbf{T}[V_i]$ in \mathbf{T} (by Claim 6.2), we can uniquely decompose any path into a union of paths over the subtrees and edges between the subtrees:

$$\pi(u, v) = \pi(a_1, b_1) \circ (b_1, a_2) \circ \pi(a_2, b_2) \circ \cdots \circ (b_{t-1}, a_t) \circ \pi(a_t, b_t),$$

such that there exists a permutation $\sigma : [m] \rightarrow [m]$ with $V(\pi(a_i, b_i)) \subseteq V_{\sigma(i)}$ for all $i \in [t]$.

Let $a'_1 = a_1 = u$, $b'_t = b_t = v$. By Claim 6.4, for each $i \in [t-1]$ there is an edge (b'_i, a'_{i+1}) such that $b'_i \in V_{\sigma(i)}$, $a'_{i+1} \in V_{\sigma(i+1)}$, i.e. there is a primed edge connecting each pair of groups $V_{\sigma(i)}, V_{\sigma(i+1)}$ that is connected by a non-primed edge. Because for each $i \in [t]$, $\widehat{\mathbf{T}}[V_i]$ is a subtree of $\widehat{\mathbf{T}}$ and hence is connected, we see that $\widehat{\mathbf{T}}[V_{\sigma(1)} \cup \cdots \cup V_{\sigma(t)}]$ together with edges $(b'_1, a'_2) \cup \cdots \cup (b'_{t-1}, a'_t)$ forms a connected graph. It follows that $\hat{\pi}(u, v)$ cannot have any additional edges between V_i groups, i.e., the path $\hat{\pi}(a'_i, b'_i)$ stays in $V_{\sigma(i)}$ for all $i \in [t]$. Thus

$$\hat{\pi}(a'_1, b'_1) \circ (b'_1, a'_2) \circ \hat{\pi}(a'_2, b'_2) \circ \cdots \circ (b'_{t-1}, a'_t) \circ \hat{\pi}(a'_t, b'_t)$$

is a path from u to v in $\widehat{\mathbf{T}}$, and it satisfies the required properties. \square

The next claim shows that if the path has a weak edge, then the reconstruction will still be reasonable:

Claim 6.6 (Reconstructed topology correct – very weak edge case). *Let $u, v \in \mathcal{V}$ such that $\pi(u, v) \subseteq \mathsf{T}$ has an edge e with $\theta_e \leq 4\epsilon$. Then $\hat{\pi}(u, v) \subseteq \hat{\mathsf{T}}$ has an edge e' with $\hat{\theta}_{e'} \leq 5\epsilon$.*

Proof. Consider the tree T^{CL} constructed in LEARNFERROMODEL and let $\pi^{\text{CL}}(u, v)$ be the path from u to v in T^{CL} . Note that by the construction of $\hat{\mathsf{T}}$, the path $\hat{\pi}(u, v)$ includes all edges $e \in \pi^{\text{CL}}(u, v)$ with $\tilde{\mu}_e \leq 1/10$ as these are in W . It follows that if there is some edge \tilde{e} in $\pi^{\text{CL}}(u, v)$ with $\mu_{\tilde{e}} \leq 4\epsilon$, then $\tilde{\mu}_{\tilde{e}} \leq \mu_{\tilde{e}} + \epsilon \leq 5\epsilon \leq 1/10$, and hence $\tilde{e} \in W$, so $\hat{\pi}(u, v)$ would contain this edge \tilde{e} and we would be finished. Henceforth assume that $\mu_{\tilde{e}} > 4\epsilon$ for all \tilde{e} in $\pi^{\text{CL}}(u, v)$. In particular, $\theta_{\tilde{e}} > 4\epsilon$ for all \tilde{e} in $\pi^{\text{CL}}(u, v)$.

Now, the hypothesis of the claim is that $\pi(u, v) \subseteq \mathsf{T}$ has an edge e of weight $\theta_e \leq 4\epsilon$. This contradicts the fact that T is a maximum spanning tree on the weights μ (Claim 2.7), because $\pi^{\text{CL}}(u, v) \cup \pi(u, v)$ must form a cycle containing e , and removing e from T and adding in an edge of weight greater than 4ϵ from the cycle would strictly increase the weight of the tree. \square

To summarize, Claims 6.5 and 6.6 show that between the sets V_1, \dots, V_m of the partition, $\hat{\mathsf{T}}$ roughly recovers the topology of T .

6.2.3 Bounding $|\hat{\mu}_{uv} - \Upsilon_{uv}|$

We will now prove the claimed upper bound on $|\hat{\mu}_{uv} - \Upsilon_{uv}|$, where we recall that $\hat{\mu}_{uv} = \prod_{i=1}^t \hat{\mu}_{a'_i b'_i} \prod_{i=1}^{t-1} \hat{\mu}_{b'_i a'_{i+1}}$ and $\Upsilon_{uv} = \prod_{i=1}^t \mu_{a'_i b'_i} \prod_{i=1}^{t-1} \mu_{b'_i a'_{i+1}}$ for a'_i, b'_i decomposing the path $\hat{\pi}(u, v)$ as in the proof of Proposition 5.2. The main result is:

Claim 6.7. *If all edges e in the path $\pi(u, v)$ have weight $\theta_e \geq 4\epsilon$, then $|\hat{\mu}_{uv} - \Upsilon_{uv}| \leq (C' + 1)\epsilon$, where C' is the constant from Claim 6.3.*

The proof of the claim will rely on the following technical bound:

Lemma 6.8 (Telescoping bound). *Let $0 \leq \rho \leq 1$. For any $\alpha, \beta \geq 0$, $\gamma_1, \dots, \gamma_t \in [0, \rho]$, and $\gamma'_1, \dots, \gamma'_t \in [0, \rho]$ such that $|\gamma_i - \gamma'_i| \leq \delta$ for all $i \in [t]$, we have*

$$\left| \alpha \cdot \prod_{i=1}^t \gamma'_i - \beta \cdot \prod_{i=1}^t \gamma_i \right| \leq |\alpha - \beta| \rho^t + \alpha t \delta \rho^{t-1}$$

Proof. By the triangle inequality and the assumption that $\gamma_i, \gamma'_i \leq \rho$,

$$\left| \alpha \cdot \prod_{i=1}^t \gamma'_i - \beta \cdot \prod_{i=1}^t \gamma_i \right| \leq |\alpha - \beta| \prod_{i=1}^t \gamma_i + \alpha \left| \prod_{i=1}^t \gamma'_i - \prod_{i=1}^t \gamma_i \right| \leq |\alpha - \beta| \rho^t + \alpha \left| \prod_{i=1}^t \gamma'_i - \prod_{i=1}^t \gamma_i \right|.$$

Finally, the last term in the absolute value can be bounded by the triangle inequality by considering intermediate terms $S_j := \left(\prod_{i=1}^j \gamma'_i \right) \left(\prod_{i=j+1}^m \gamma_i \right)$ and using that $|S_j - S_{j+1}| \leq \delta \rho^{m-1}$. \square

We are now ready to prove Claim 6.7.

Proof of Claim 6.7. Applying Lemma 6.8 with $\alpha = \prod_{i=1}^t \mu_{a'_i b'_i}$, $\beta = \prod_{i=1}^t \hat{\mu}_{a'_i b'_i}$, $\gamma_i = \mu_{b'_i a'_{i+1}}$, $\gamma'_i = \hat{\mu}_{b'_i a'_{i+1}}$, $\rho = (1/10 + \epsilon)$,

$$|\hat{\mu}_{uv} - \Upsilon_{uv}| \leq \left| \prod_{i=1}^t \mu_{a'_i b'_i} - \prod_{i=1}^t \hat{\mu}_{a'_i b'_i} \right| \cdot (1/10 + \epsilon)^{t-1} + \epsilon \cdot (t-1)(1/10 + \epsilon)^{t-2}, \quad (11)$$

where we have used $\alpha, \beta \leq 1$, $\gamma'_i = \hat{\mu}_{b'_i a'_{i+1}} = \tilde{\mu}_{b'_i a'_{i+1}} \leq 1/10$ (the equality is due to (b'_i, a'_{i+1}) being a edge of $\widehat{\mathsf{T}}$ and the inequality by Claim 6.1 since b'_i and a'_{i+1} are in distinct sets $V_{\sigma(i)}, V_{\sigma(i+1)}$ of the partition), and $|\gamma_i - \gamma'_i| \leq \epsilon$ since $\tilde{\mu}$ ϵ -approximates μ .

Again applying Lemma 6.8, this time with $\alpha = \beta = 1$, $\gamma_i = \mu_{a'_i b'_i}$, $\gamma'_i = \hat{\mu}_{a'_i b'_i}$, $\rho = 1$,

$$\left| \prod_{i=1}^t \mu_{a'_i b'_i} - \prod_{i=1}^t \hat{\mu}_{a'_i b'_i} \right| \leq t C' \epsilon, \quad (12)$$

where we have used $|\gamma_i - \gamma'_i| \leq C' \epsilon$ by Claim 6.3 (approximation guarantee for paths within the set $V_{\sigma(i)}$). Combining (11) and (12) yields

$$|\hat{\mu}_{uv} - \Upsilon_{uv}| \leq \epsilon \cdot (C' \cdot t(1/10 + \epsilon)^{t-1} + (t-1)(1/10 + \epsilon)^{t-2}) \leq (C' + 1)\epsilon.$$

The last inequality can be obtained by a simple calculus argument applied to tx^{t-1} showing that for $x \leq 1/e$ the maximum over integers t is obtained at $t = 1$, and then separately bounding each of the two terms. \square

6.2.4 Bounding $|\Upsilon_{uv} - \mu_{uv}|$

In this section, we will prove the claimed upper bound on $|\Upsilon_{uv} - \mu_{uv}|$, where recall that $\Upsilon_{uv} = \prod_{i=1}^t \mu_{a'_i b'_i} \prod_{i=1}^{t-1} \mu_{b'_i a'_{i+1}}$ and $\mu_{uv} = \prod_{i=1}^t \mu_{a_i b_i} \prod_{i=1}^{t-1} \mu_{b_i a_{i+1}}$ for a'_i, b'_i decomposing $\hat{\pi}(u, v)$ and a_i, b_i decomposing $\pi(u, v)$ as in the proof of Proposition 5.2. The case that there is a very weak edge e' with $\theta_{e'} \leq 4\epsilon$ in $\pi(u, v)$ is addressed at the start of the proof of Proposition 5.2 in Section 6.1, so as noted there, we assume that all edges in $\pi(u, v)$ have weight at least 4ϵ .

Claim 6.9. *If all edges e in the path $\pi(u, v)$ have weight $\theta_e \geq 4\epsilon$, then $|\Upsilon_{uv} - \mu_{uv}| \leq 4\epsilon$.*

Proof. It is clear that $\Upsilon_{uv} \leq \mu_{uv}$ via a sequence of applications of the triangle inequality (Fact 2.2), since $a'_1 = u, b'_t = v$. It remains to show that $\Upsilon_{uv} \geq \mu_{uv} - 4\epsilon$. By Claim 6.5, $a_i, b_i, a'_i, b'_i \in V_{\sigma(i)}$ for all $i \in [t]$, for some permutation $\sigma : [m] \rightarrow [m]$. Applying the triangle inequality (Fact 2.2) for each $i \in [t]$ we have

$$\mu_{a'_i b'_i} \geq \mu_{a'_i a_i} \mu_{a_i b_i} \mu_{b_i b'_i}. \quad (13)$$

Using Fact 2.1, we can also rewrite $\mu_{b'_i a'_{i+1}}$ for any $i \in [t-1]$ as

$$\mu_{b'_i a'_{i+1}} = \mu_{b'_i b_i} \mu_{b_i a_{i+1}} \mu_{a_{i+1} a'_{i+1}}, \quad (14)$$

since (b_i, a_{i+1}) is an edge in $\pi(b'_i, a'_{i+1})$. To see this, note that removing (b_i, a_{i+1}) from T splits it into two components, $U \supseteq V_{\sigma(i)} \ni b'_i$ and $U' \supseteq V_{\sigma(i+1)} \ni a'_{i+1}$. Combining these last two equations gives

$$\begin{aligned} \Upsilon_{uv} &\geq \prod_{i=1}^t \mu_{a_i b_i} \prod_{i=1}^{t-1} \mu_{b'_i a'_{i+1}} \mu_{b'_i b_i} \mu_{a'_{i+1} a_{i+1}} && \text{by (13), since } a'_1 = a_1, b'_t = b_t \\ &= \prod_{i=1}^t \mu_{a_i b_i} \prod_{i=1}^{t-1} (\mu_{b'_i a'_{i+1}})^2 / \mu_{b_i a_{i+1}} && \text{by (14) and } \mu_{b_i a_{i+1}} \geq 4\epsilon > 0 \\ &\geq \prod_{i=1}^t \mu_{a_i b_i} \prod_{i=1}^{t-1} (\mu_{b_i a_{i+1}} - 4\epsilon) && \text{by deferred Claim 6.10 below.} \end{aligned}$$

Because $\Upsilon_{uv} \leq \mu_{uv}$, it follows that

$$\begin{aligned} |\Upsilon_{uv} - \mu_{uv}| &\leq \left| \prod_{i=1}^t \mu_{a_i b_i} \prod_{i=1}^{t-1} (\mu_{b_i a_{i+1}} - 4\epsilon) - \mu_{uv} \right| \\ &= \left| \prod_{i=1}^t \mu_{a_i b_i} \prod_{i=1}^{t-1} (\mu_{b_i a_{i+1}} - 4\epsilon) - \prod_{i=1}^t \mu_{a_i b_i} \prod_{i=1}^{t-1} \mu_{b_i a_{i+1}} \right|. \end{aligned}$$

We now apply the bound from Lemma 6.8 with $\alpha = \beta = \prod_{i=1}^t \mu_{a_i b_i}$, $\gamma_i = \mu_{b_i a_{i+1}}$, $\gamma'_i = \mu_{b_i a_{i+1}} - 4\epsilon$, and $\rho = 1/10 + \epsilon$, to obtain

$$|\Upsilon_{uv} - \mu_{uv}| \leq (4\epsilon)(t-1)(1/10 + \epsilon)^{t-2} \leq 4\epsilon,$$

where we have used $\alpha \leq 1$ and $\gamma_i, \gamma'_i \in [0, 1/10 + \epsilon]$ by Claim 6.1. \square

Claim 6.10. *In the setup of Claim 6.9, $(\mu_{b'_i a'_{i+1}})^2 / \mu_{b_i a_{i+1}} \geq \mu_{b_i a_{i+1}} - 4\epsilon$ for all $i \in [t-1]$.*

Proof. First, by the maximum spanning tree construction in LEARNFERROMODEL,

$$\tilde{\mu}_{b'_i a'_{i+1}} = \max_{w \in V_{\sigma(i)}, w' \in V_{\sigma(i+1)}} \tilde{\mu}_{ww'},$$

and, by the fact that T is a maximum spanning tree on μ (Claim 2.7),

$$\mu_{b_i, a_{i+1}} = \max_{w \in V_{\sigma(i)}, w' \in V_{\sigma(i+1)}} \mu_{ww'}.$$

Because $\tilde{\mu}$ ϵ -approximates μ and the maximums of two functions that are point-wise within ϵ are also within ϵ , it follows that

$$|\tilde{\mu}_{b'_i a'_{i+1}} - \mu_{b_i a_{i+1}}| \leq \epsilon.$$

As a consequence,

$$|\mu_{b'_i a'_{i+1}} - \mu_{b_i a_{i+1}}| \leq |\tilde{\mu}_{b'_i a'_{i+1}} - \mu_{b_i a_{i+1}}| + |\tilde{\mu}_{b'_i a'_{i+1}} - \mu_{b'_i a'_{i+1}}| \leq 2\epsilon.$$

(Here for the second term we directly used that $\tilde{\mu}$ and μ are within ϵ .) Using the last equation together with the assumption $\mu_{b_i a_{i+1}} = \theta_{b_i a_{i+1}} \geq 4\epsilon > 2\epsilon$ shows that $(\mu_{b'_i a'_{i+1}})^2 \geq (\mu_{b_i a_{i+1}} - 2\epsilon)^2$, so

$$(\mu_{b'_i a'_{i+1}})^2 / \mu_{b_i a_{i+1}} \geq (\mu_{b_i a_{i+1}} - 2\epsilon)^2 / \mu_{b_i a_{i+1}} = \mu_{b_i a_{i+1}} - 4\epsilon + 4\epsilon^2 / \mu_{b_i a_{i+1}} \geq \mu_{b_i a_{i+1}} - 4\epsilon,$$

which proves the claim. \square

7 Additive metric reconstruction: handling strong edges, step 1

In this section, we present our ADDITIVEMETRICRECONSTRUCTION algorithm, which is the first step in the tree metric reconstruction procedure underlying LEARNLWRBDDMODEL which handles the strong edges of the model (see Section 5.3). We prove the following guarantee, recalling from Definition 2.11 that an additive metric is a tree metric with Steiner nodes:

Theorem 7.1 (Restatement of Theorem 5.7). *There exists a constant $C > 0$ such that the following result is true for arbitrary $\epsilon > 0$ and any $L \geq 100\epsilon$. Let (T, d) be an unknown tree metric on vertex set \mathcal{V} . Suppose that neighboring vertices $u, v \in \mathcal{V}$ satisfy $d(u, v) \leq L$, and suppose we have access to a distance estimate $\tilde{d}_{\text{pre}} : \mathcal{V} \times \mathcal{V} \rightarrow \mathbb{R} \cup \{\infty\}$ such that*

$$d(u, v) \leq \tilde{d}_{\text{pre}}(u, v), \quad \text{for all } u, v \in \mathcal{V} \quad (15)$$

$$\tilde{d}_{\text{pre}}(u, v) \leq d(u, v) + \epsilon \quad \text{for all } u, v \in \mathcal{V} \text{ such that } d(u, v) \leq 3L_{\text{upp}} \quad (16)$$

Given this information, there exists a $O(|\mathcal{V}|^2)$ -time algorithm ADDITIVEMETRICRECONSTRUCTION which constructs an additive metric (\hat{T}, \hat{d}) with $|V(\hat{T})| \leq O(|\mathcal{V}|)$ Steiner nodes such that

$$|\hat{d}(u, v) - d(u, v)| \leq C(d(u, v)/L + 1)\epsilon \quad (17)$$

for all $u, v \in \mathcal{V}$.

Before proving Theorem 5.7, let us review related results in the literature and state how our theorem differs. Reconstruction of tree metrics is a well-studied problem related to phylogenetic reconstruction (see e.g. [12] and references within). In the setting where all pairwise distances are known to be estimated accurately, the additive metric (Steiner node) version of the problem has a known solution by Agarwala et al. [2]. This follows from an important connection between the tree metric reconstruction problem and reconstruction of ultrametrics (defined later). This is useful because optimal ultrametric tree reconstruction in ℓ_∞ can (remarkably) be solved in polynomial time [19]; this follows from the existence of *subdominant ultrametrics* which we will discuss later. This reduction leads to the following 3-approximation algorithm for reconstructing additive metrics (i.e. tree metrics with Steiner nodes):

Theorem 7.2 (Agarwala et al. [2]). *Suppose that d is an additive metric on \mathcal{V} and \tilde{d} satisfies $|\tilde{d}(x, y) - d(x, y)| \leq \epsilon$ for all $x, y \in \mathcal{V}$. Then an additive metric (\hat{T}, \hat{d}) such that*

$$\max_{x, y} |\hat{d}(x, y) - d(x, y)| \leq 3\epsilon,$$

can be computed from \tilde{d} in time $O(|\mathcal{V}|^2)$.

Unfortunately, the additive metric reconstruction algorithm of Agarwala et al. [2] cannot be used in our setting. In contrast to Theorem 7.2, our Theorem 5.7 is not guaranteed that $\tilde{d}(u, v)$ is ϵ -close to $d(u, v)$ for every pair of vertices u, v – the approximation only holds for close vertices. Thus, in our setting for two vertices that are far away the distance cannot be estimated up to $O(\epsilon)$ additive error, and our algorithm settles for reconstructing the distance up to a $(1 + O(\epsilon))$ multiplicative factor. (In the statistical setting we consider, reconstructing the distances between far-away vertices up to $O(\epsilon)$ error would require $\Omega(n)$ samples, because the relevant correlations can be as small as $O(1/n)$, whereas we aim to use $O(\log n)$ samples).

We discuss the main idea behind Theorem 5.7. While the additive metric reconstruction algorithm of Agarwala et al. [2] cannot be applied in our setting, the underlying mathematical tools from metric geometry remain useful and we discuss these in the first subsection. At a high level, the key idea we take from [2] is that we can reconstruct an additive metric close to d provided we: (1) compute a close-to-tight upper bound \tilde{d} on the true metric d and (2) show the existence of a “witness tree” d' which is close to d and matches root-to-leaf distances in \tilde{d} from some arbitrary choice of root ρ . Here “close” needs to be in the same sense as (17), i.e. close nodes are additively close and far-away nodes are multiplicatively close. For step (1), it is unacceptable in our setting

to use the approximate distances \tilde{d}_{pre} directly, as those distances may be very inaccurate between far-away nodes. Instead, we show for (1) that taking \tilde{d} as the *shortest-path metric* over *upper (confidence) bounds* on the estimated distances suffices to construct a suitable \tilde{d} — see Lemma 7.9. For (2), we construct a witness tree based on d, \tilde{d} using a new two-step procedure, which carefully balances the opposing goals of matching root-to-leaf distances in \tilde{d} while preserving the overall metric approximation of d — see Lemma 7.11. Combining key inputs (1) and (2) lets us use a similar combinatorial algorithm to [2], based on reduction to computation of a *subdominant ultrametric*, to compute an additive metric approximation.

7.1 Preliminaries from (ultra)metric geometry

Pseudometrics. A tuple (\mathcal{V}, d) , where $d : \mathcal{V} \times \mathcal{V} \rightarrow \mathbb{R}_{\geq 0}$, is a *pseudometric space* if d satisfies the following axioms:

1. $d(u, u) = 0$ for all $u \in \mathcal{V}$,
2. $d(u, v) = d(v, u)$ for all $u, v \in \mathcal{V}$, and
3. $d(u, v) \leq d(u, w) + d(w, v)$ for all $u, v, w \in \mathcal{V}$.

Normally a pseudometric is called a metric when it satisfies the additional axiom that $d(u, v) = 0$ iff $u = v$. In what follows this distinction is counterproductive; both terms will refer to a pseudometric unless explicitly noted otherwise.

Ultrametrics. Ultrametrics are a special kind of metric where a very strong form of the triangle inequality holds (18); they play a fundamental role in metric geometry and its applications to probability theory, physics, computer science, etc. — see the survey [31].

Definition 7.3 (Ultrametric). Suppose (\mathcal{V}, d) is a metric space. We say that d is an *ultrametric* if

$$d(u, v) \leq \max\{d(u, w), d(w, v)\} \quad (18)$$

for all $u, v, w \in \mathcal{V}$.

Definition 7.4 (Subdominant ultrametric). Let $a : \mathcal{V} \times \mathcal{V} \rightarrow \mathbb{R}_{\geq 0}$ be an arbitrary function. We say that an ultrametric $d : \mathcal{V} \times \mathcal{V} \rightarrow \mathbb{R}_{\geq 0}$ is the *subdominant ultrametric* of a if it satisfies the following properties:

1. For all $u, v \in \mathcal{V}$, $d(u, v) \leq a(u, v)$.
2. For any other ultrametric e such that $e(u, v) \leq a(u, v)$ for all $u, v \in \mathcal{V}$, it holds true that $e(u, v) \leq d(u, v)$ for all $u, v \in \mathcal{V}$.

Remarkably, for any $a : \mathcal{V} \times \mathcal{V} \rightarrow \mathbb{R}_{\geq 0}$ on a finite² set \mathcal{V} , a subdominant ultrametric always exists and can be efficiently computed. This result has a long history and shows up in the statistical physics literature [36, 37], the literature on numerical taxonomy [23], and studies on hierarchical clustering [30] among other places. A simple algorithm for computing the subdominant ultrametric based off of minimum spanning trees, which gives the result stated below, appears at least in [19] and [36].

Theorem 7.5 ([19, 36]). *Let $a : \mathcal{V} \times \mathcal{V} \rightarrow \mathbb{R}_{\geq 0}$ be an arbitrary function and suppose \mathcal{V} is finite. Then the subdominant ultrametric of a exists and it can be computed in time $O(|\mathcal{V}|^2)$.*

²When \mathcal{V} is infinite, subdominant ultrametrics may not exist. See [5] for a more general existence result.

Ultrametric = Additive metric + Centroid metric. An ultrametric on a finite set \mathcal{V} can be visualized as a tree where the leaves are the elements of \mathcal{V} , the leaves are all equidistant from the root of the tree, and the least common ancestor of $u, v \in \mathcal{V}$ is at height $\mathbf{d}(u, v)/2$ above them. This description makes it clear that ultrametrics are a special case of additive metrics. It turns out there is a simple way to convert an additive metric to an ultrametric by adding an appropriate *centroid metric* [4, 2].

Definition 7.6 (Centroid Metric). Suppose that for all $u \in \mathcal{V}$ we are given $\ell_u \geq 0$. Then we say that $\mathbf{c} : \mathcal{V} \times \mathcal{V} \rightarrow \mathbb{R}_{\geq 0}$ is the corresponding *centroid metric* if $\mathbf{c}(u, v) = \ell_u + \ell_v$ for $u \neq v$.

Note that the centroid metric is also an additive metric; it can be visualized as the graph metric on a star graph where the leaves are $u \in \mathcal{V}$ and the edges have length ℓ_u .

The following lemma (established in the proof of Claim 3.5B of [2]) shows that we can convert any additive metric into an ultrametric by adding an appropriate centroid metric. Equivalently, this ultrametric can be defined by the equation $\mathbf{e}(u, v) = 2D_{\max} - 2\mathbf{d}(\rho, \text{LCA}_{\rho}(u, v))$ below.

Lemma 7.7 ([2]). Suppose \mathbf{d} is an additive metric on \mathcal{V} , $\rho \in \mathcal{V}$, and $\ell_u = D_{\max} - \mathbf{d}(\rho, u)$ for all u where $D_{\max} \geq \max_u \mathbf{d}(\rho, u)$. Let \mathbf{c} be the centroid metric formed from $\{\ell_u\}_{u \in \mathcal{V}}$. Then $\mathbf{e} = \mathbf{d} + \mathbf{c}$ is an ultrametric such that

$$2D_{\max} \geq \mathbf{e}(u, v) = 2D_{\max} - 2\mathbf{d}(\rho, \text{LCA}_{\rho}(u, v)) \geq 2 \max\{\ell_u, \ell_v\}.$$

In the reverse direction, if we have an ultrametric and a centroid metric which satisfy the appropriate conditions, then subtracting the centroid metric from the ultrametric gives an additive metric. This is established in the proof of Claim 3.5A in [2].

Lemma 7.8 ([2]). Suppose that \mathbf{e} is an ultrametric on \mathcal{V} , $\mathbf{c}(u, v) = \ell_u + \ell_v$ is a centroid metric on \mathcal{V} , and $\rho \in \mathcal{V}$. Let $D_{\max} = \ell_{\rho}$. Suppose that

$$\mathbf{e}(\rho, u) = 2D_{\max}$$

for all $u \in \mathcal{V}$, and for all $u, v \in \mathcal{V}$

$$2D_{\max} \geq \mathbf{e}(u, v) \geq 2 \max\{\ell_u, \ell_v\}.$$

Then $\mathbf{d} = \mathbf{e} - \mathbf{c}$ is an additive metric on \mathcal{V} and $\ell_u = D_{\max} - \mathbf{d}(\rho, u)$ for all u .

7.2 Reconstructing an accurate additive metric

The pseudocode for ADDITIVEMETRICRECONSTRUCTION is given as Algorithm 4.

We prove correctness, Theorem 5.7. In the analysis, we let $\tilde{\mathbf{d}}$ be the shortest path metric with respect to $\tilde{\mathbf{d}}$:

$$\tilde{\mathbf{d}}(u, v) = \min_{\pi \in \text{paths}(u, v)} \sum_{e=(a, b) \in \pi} \tilde{\mathbf{d}}_{\text{pre}}(a, b). \quad (19)$$

This is consistent with the algorithm, although in the algorithm only shortest paths from ρ are actually used or computed; the shortest paths between other pairs of nodes are useful in the analysis. We start with a lemma relating \mathbf{d} and $\tilde{\mathbf{d}}$.

Lemma 7.9. Under the assumptions of Theorem 5.7, for all $u, v \in \mathcal{V}$

$$\mathbf{d}(u, v) \leq \tilde{\mathbf{d}}(u, v) \leq \mathbf{d}(u, v) + (\mathbf{d}(u, v)/L + 1)\epsilon.$$

Algorithm 4: ADDITIVEMETRICRECONSTRUCTION($\tilde{d}_{\text{pre}}, L, \epsilon$)

Input: $\tilde{d}_{\text{pre}}, L, \epsilon$ as in Theorem 5.7

Output: Additive metric (\hat{T}, \hat{d}) satisfying the approximation guarantee (17)

- 1 Fix an arbitrary root $\rho \in \mathcal{V}$.
- 2 Using Dijkstra's algorithm, compute shortest path distances $\tilde{d}(\rho, u)$ for all u from the weighted graph corresponding to \tilde{d}_{pre} .
- 3 Define $D_{\text{max}} := \max_u \tilde{d}(\rho, u)$, $\ell_u := D_{\text{max}} - \tilde{d}(\rho, u)$ for all $u \in \mathcal{V}$, and let $c(u, v) = \ell_u + \ell_v$ be the corresponding centroid metric.
- 4 Define $a(u, \rho) := a(\rho, u) = 2D_{\text{max}}$. For all $u, v \in \mathcal{V} \setminus \{\rho\}$ such that $\tilde{d}_{\text{pre}}(u, v) < 3L$, define

$$a(u, v) := \min\{2D_{\text{max}}, \tilde{d}_{\text{pre}}(u, v) + c(u, v) + 44\epsilon\}.$$

- 5 Let e be the subdominant ultrametric of a , computed using the algorithm from Theorem 7.5.

- 6 Return $\hat{d} := e - c$, and the corresponding tree with Steiner nodes \hat{T} .
-

Proof. By (15) and the triangle inequality, we know $d(u, v) \leq \tilde{d}(u, v)$ which proves the lower bound.

The upper bound follows by considering the path from u to v in T : since the maximum edge length in the ground truth tree T is L , we can choose vertices $u = w_1, \dots, w_k$ along the path from u to v such that $d(w_i, w_{i+1}) \in [L, 2L]$ for all $i < k$ and $d(w_k, v) \leq L$. (Note that $k-1 \leq d(u, v)/L$.) Now because \tilde{d} is the shortest path metric,

$$\tilde{d}(u, v) \leq \sum_{i < k} \tilde{d}_{\text{pre}}(w_i, w_{i+1}) + \tilde{d}_{\text{pre}}(w_k, v) \stackrel{(16)}{\leq} \sum_{i < k} d(w_i, w_{i+1}) + d(w_k, v) + k\epsilon.$$

Using $d(u, v) = \sum_{i < k} d(w_i, w_{i+1}) + d(w_k, v)$ due to d being a tree metric as well as the fact that $k-1 \leq d(u, v)/L$ shows the upper bound. \square

The following lemma shows that estimated root-to-leaf distances are consistent in the following sense: for any node u with nearby descendant v , $\tilde{d}(\rho, v) - \tilde{d}(\rho, u)$ is a $O(\epsilon)$ additive approximation to $d(u, v)$. Note that if $\tilde{d} = d$ then this would be an exact equality. This is used in Lemma 7.11 below to argue that a slightly modified version of d can match root-to-leaf distances with \tilde{d} exactly.

Lemma 7.10. *Under the assumptions of Theorem 5.7 the following is true. Choose an arbitrary root ρ for T . For any u, v such that $d(u, v) \leq 2L$ and u is an ancestor of v in T ,*

$$\tilde{d}(\rho, u) + d(u, v) - 4\epsilon \leq \tilde{d}(\rho, v) \leq \tilde{d}(\rho, u) + d(u, v) + \epsilon.$$

Proof. To prove the upper bound, we observe that $\tilde{d}(\rho, v) \leq \tilde{d}(\rho, u) + \tilde{d}(u, v)$ by the triangle inequality for d and then by (16) $\tilde{d}(u, v) \leq \tilde{d}_{\text{pre}}(u, v) \leq d(u, v) + \epsilon$.

Now we prove the lower bound on $\tilde{d}(\rho, v)$. Consider the shortest path $\rho = a_1, a_2, \dots, a_k = v$ from ρ to v in the weighted graph \tilde{d}_{pre} . Let a_i be the last node on this path that is not a descendant of u in T . Note that since $d(u, v) \leq 2L$ by assumption, (16) implies that $\tilde{d}_{\text{pre}}(u, v) \leq 2L + \epsilon \leq 3L$. By (15) this gives $d(a_i, a_{i+1}) \leq 3L$. Since a_i is not a descendant of u , but a_{i+1} is a descendant of u , the path in T between a_i and a_{i+1} must contain u , so $d(u, a_i) \leq 3L$.

By definition of \tilde{d} and a_i , we have $\tilde{d}(\rho, v) = \tilde{d}(\rho, a_i) + \tilde{d}(a_i, v)$, hence

$$\begin{aligned} \tilde{d}(\rho, v) &\geq \tilde{d}(\rho, a_i) + d(a_i, v) = \tilde{d}(\rho, a_i) + d(a_i, u) + d(u, v) \\ &\geq \tilde{d}(\rho, a_i) + \tilde{d}(a_i, u) + d(u, v) - (3L/L + 1)\epsilon \geq \tilde{d}(\rho, u) + d(u, v) - 4\epsilon \end{aligned}$$

where in the first step we used Lemma 7.9, in the second step we know $d(a_i, v) = d(a_i, u) + d(u, v)$ because v is an ancestor of u and a_i is not, in the third step we used Lemma 7.9 and $d(u, a_i) \leq 3L$, and in the last step we used the triangle inequality. \square

The following Lemma 7.11 is crucial. It proves that there exists an additive metric d' which matches the root-to-leaf distances in \tilde{d} and is also a good approximation to the ground truth d . This is crucial for applying the tools for ultrametric reconstruction from Section 7.1: more specifically, the existence of d' ensures that when we convert from \tilde{d} to a and extract the subdominant ultrametric e , that the distances in e are not shrunk by too much (see the proof of Theorem 5.7).

We prove the Lemma with an explicit construction of d' . This construction assumes knowledge of the ground truth tree T , and therefore cannot be run by the algorithm ADDITIVEMETRICRECONSTRUCTION, which only has knowledge of \tilde{d}_{pre} . Nevertheless, the constructive proof suffices to ensure the existence of d' , which we will later show is sufficient for ADDITIVEMETRICRECONSTRUCTION to succeed.

Lemma 7.11. *Under the assumptions of Theorem 5.7 the following is true. There exists an additive metric d' such that*

$$|d(u, v) - d'(u, v)| \leq (36 + 2d(u, v)/L)\epsilon \quad (20)$$

for all $u, v \in \mathcal{V}$ and

$$d'(\rho, u) = \tilde{d}(\rho, u) \quad (21)$$

for all $u \in \mathcal{V}$.

Proof. We construct d' in two phases, which we now overview. In the first phase, we subdivide T into a union of subtrees of diameter $\Theta(L)$. The roots of these subtrees are a subset of the nodes $S \subset \mathcal{V}$. By shortening or lengthening edges of T appropriately, we construct a T -structured metric d'' such that the condition $d''(\rho, v) = \tilde{d}(\rho, v)$ holds for all $v \in S$ and such that d'' approximates d in the sense of (20). In the second phase, we create the additive metric d' by moving each node in d'' by $O(\epsilon)$ distance.

First phase. Consider the ground truth tree T with edge lengths corresponding to the tree metric d . Run a breadth-first search (BFS) from the root node ρ , inductively constructing a set of nodes $S \subset \mathcal{V}$, as follows. First let $\rho \in S$. Then, when the BFS reaches node v , add v to S if and only if the closest ancestor to v that has been added to S is at distance at least L from v .

For any node $w \in \mathcal{V} \setminus \{\rho\}$, let $\text{pa}(w)$ denote the parent of w . Also, let $\text{pa}_S(s)$ denote the closest (strict) ancestor of w that lies in S . Note that $\text{pa}_S(w)$ is guaranteed to exist, since the root ρ is in the set S . By construction, $d(\text{pa}_S(s), s) \in [L, 2L]$ for any node $s \in S \setminus \{\rho\}$, since the maximum edge length is at most L in the tree. Similarly, we have $d(\text{pa}_S(u), u) \leq 2L$ for any $u \in \mathcal{V} \setminus \{\rho\}$.

We define d'' by modifying the edge lengths in d . First, shorten certain edges of d : for any edge $(\text{pa}(w), w)$ in T , let

$$d''(\text{pa}(w), w) \leftarrow \begin{cases} d(\text{pa}(w), w), & d(\text{pa}_S(w), \text{pa}(w)) \geq 2\epsilon \\ \max(0, d(\text{pa}(w), w) - 2\epsilon + d(\text{pa}_S(w), \text{pa}(w))), & d(\text{pa}_S(w), \text{pa}(w)) < 2\epsilon \end{cases}$$

The reason for this shortening step is that for any $s \in S \setminus \{\rho\}$, we have $d''(\text{pa}_S(s), s) = \max(0, d(\text{pa}_S(s), s) - 2\epsilon)$. After applying the edge shortening step, also update the length of the parent edge of each $s \in S \setminus \{\rho\}$:

$$d''(\text{pa}(s), s) \leftarrow d''(\text{pa}(s), s) + (\tilde{d}(\rho, s) - \tilde{d}(\rho, \text{pa}_S(s))) - (d(\text{pa}_S(s), s) - 2\epsilon).$$

This latter update is well-defined since the edge is lengthened (indeed, Lemma 7.10 implies $\tilde{d}(\rho, s) - (\tilde{d}(\rho, \text{pa}_S(s)) + d(\text{pa}_S(s), s) - 2\epsilon) \geq 0$) and thus the length remains nonnegative. Together, these edge shortenings and lengthenings ensure that for any $v \in \mathcal{V} \setminus S$,

$$\begin{aligned} d''(\text{pa}_S(v), v) &= \sum_{(\text{pa}(w), w) \in \pi(\text{pa}_S(v), v)} d''(\text{pa}(w), w) \\ &= \max(0, d(\text{pa}_S(v), v) - 2\epsilon), \end{aligned} \quad (22)$$

and similarly for any $s \in S \setminus \{\rho\}$,

$$\begin{aligned} d''(\text{pa}_S(s), s) &= \max(0, d(\text{pa}_S(s), s) - 2\epsilon) + (\tilde{d}(\rho, s) - \tilde{d}(\rho, \text{pa}_S(s))) - (d(\text{pa}_S(s), s) - 2\epsilon) \\ &= \tilde{d}(\rho, s) - \tilde{d}(\rho, \text{pa}_S(s)), \end{aligned} \quad (23)$$

where to obtain the last line we use $d(\text{pa}_S(s), s) \geq L \geq 2\epsilon$. Since $d''(\rho, s) = d''(\rho, \text{pa}_S(s)) + d''(\text{pa}_S(s), s)$, applying the above equation inductively on the depth of s , we obtain

$$d''(\rho, s) = \tilde{d}(\rho, s) \text{ for all } s \in S. \quad (24)$$

Furthermore, for any $u \in \mathcal{V} \setminus \{\rho\}$, we have

$$|d''(\text{pa}_S(u), u) - d(\text{pa}_S(u), u)| \leq 2\epsilon. \quad (25)$$

If $u \in \mathcal{V} \setminus S$, this follows from (22). If $u \in S \setminus \{\rho\}$, this follows from (23), Lemma 7.10, and $d(\text{pa}_S(u), u) \leq 2L$.

We now prove that for any $u, v \in \mathcal{V}$,

$$|d''(u, v) - d(u, v)| \leq (2d(u, v)/L + 12)\epsilon. \quad (26)$$

Since d and d'' are both tree metrics on d , we must have $d(u, v) = d(u, w) + d(v, w)$ and $d''(u, v) = d''(u, w) + d''(v, w)$ for $w = \text{LCA}_\rho^\top(u, v)$. Hence, it suffices to prove that, for any $u, w \in \mathcal{V}$ where w is an ancestor of u ,

$$|d''(u, w) - d(u, w)| \leq (2d(u, w)/L + 6)\epsilon. \quad (27)$$

Let s_1, \dots, s_k be the vertices in S along the path from w to u (i.e., $S \cap \pi(w, u)$). Note that

$$\begin{aligned} |d''(s_1, s_k) - d(s_1, s_k)| &= \left| \sum_{i=2}^k d''(\text{pa}_S(s_i), s_i) - d(\text{pa}_S(s_i), s_i) \right| \leq \sum_{i=2}^k |d''(\text{pa}_S(s_i), s_i) - d(\text{pa}_S(s_i), s_i)| \\ &\leq 2(k-1)\epsilon \leq 2d(s_1, s_k)\epsilon/L \end{aligned}$$

by (25) and the fact that $d(s_1, s_k) \geq (k-1)L$ since each pair s_i, s_{i-1} is at least at distance L . Finally, if $w = s_1$, then $d''(w, s_1) = d(w, s_1) = 0$ and if $w \neq s_1$, then $\text{pa}_S(w) = \text{pa}_S(s_1)$, so $|d''(w, s_1) - d(w, s_1)| = |d''(\text{pa}_S(s_1), s_1) - d''(\text{pa}_S(w), w) - d(\text{pa}_S(s_1), s_1) + d(\text{pa}_S(w), w)| \leq 4\epsilon$ by (25). Similarly, by (25) we have $|d''(s_k, u) - d(s_k, u)| \leq 2\epsilon$. So overall

$$\begin{aligned} |d''(w, u) - d(w, u)| &\leq |d''(w, s_1) - d(w, s_1)| + |d''(s_1, s_k) - d(s_1, s_k)| + |d''(s_k, u) - d(s_k, u)| \\ &\leq (2d(s_1, s_k)/L + 6)\epsilon \leq (2d(w, u)/L + 6)\epsilon, \end{aligned}$$

proving (27), and hence proving (26).

Second phase. Construct T' as follows: first, begin with a copy of T with edge lengths induced by d'' . However, let all the nodes in this copy of T now be Steiner nodes: i.e., node $u \in \mathcal{V}$ in T

corresponds to a Steiner node u^{stein} in the copy T' . Second, for each $u \in \mathcal{V}$, place node $u \in \mathcal{V}$ at distance $\tilde{d}(\rho, u)$ from ρ as follows: if $\tilde{d}(\rho, u) \leq d''(\rho, u)$, then place u along the path from ρ to u^{stein} . Otherwise, if $\tilde{d}(\rho, u) \geq d''(\rho, u)$, attach u to u^{stein} with an edge of length $\tilde{d}(\rho, u) - d''(\rho, u)$. This induces a T' -structured tree metric d' , which is an additive metric when restricted to the set \mathcal{V} . Note that $d'(u, u^{\text{stein}}) \leq 12\epsilon$ for all $u \in \mathcal{V}$, because if $u \in S$ then $d'(u, u^{\text{stein}}) = 0$ by (24). And if $u \in \mathcal{V} \setminus S$, then

$$\begin{aligned} d'(u, u^{\text{stein}}) &= |d''(\rho, u) - \tilde{d}(\rho, u)| = |d''(\rho, \text{pa}_S(u)) + d''(\text{pa}_S(u), u) - \tilde{d}(\rho, u)| \\ &= |\tilde{d}(\rho, \text{pa}_S(u)) + d''(\text{pa}_S(u), u) - \tilde{d}(\rho, u)| \\ &\leq |\tilde{d}(\rho, \text{pa}_S(u)) + d(\text{pa}_S(u), u) - \tilde{d}(\rho, u)| + 10\epsilon \leq 12\epsilon, \end{aligned}$$

where in the second line we have used (24) for the first equality, (27) and $d(\text{pa}_S(u), u) \leq 2L$ for the first inequality, and Lemma 7.10 for the last inequality. Hence, between any pair of vertices $u, v \in \mathcal{V}$, we have $|d'(u, v) - d''(u, v)| \leq d'(u, u^{\text{stein}}) + d'(v, v^{\text{stein}}) + |d'(u^{\text{stein}}, v^{\text{stein}}) - d''(u, v)| = d'(u, u^{\text{stein}}) + d'(v, v^{\text{stein}}) \leq 24\epsilon$, proving the lemma when combined with (26). \square

We now apply this existence lemma to prove the correctness of ADDITIVEMETRICRECONSTRUCTION.

Proof of Theorem 5.7. Define $\epsilon' = 44\epsilon$. We claim that $d'(u, v) \leq \tilde{d}_{\text{pre}}(u, v) + \epsilon'$ for any $u, v \in \mathcal{V}$. Indeed, if $\tilde{d}_{\text{pre}}(u, v) > 3L$, then $\tilde{d}_{\text{pre}}(u, v) = \infty$ so the claim holds. Otherwise, if $\tilde{d}_{\text{pre}}(u, v) \leq 3L$, then $d(u, v) \leq 3L + \epsilon$, so $d'(u, v) \leq d(u, v) + (36 + 2(3L + \epsilon)/L)\epsilon \leq \tilde{d}_{\text{pre}}(u, v) + 44\epsilon$, by Lemma 7.11 and $\epsilon \leq L$.

This guarantees $d'(u, v) + c(u, v) \leq a(u, v)$ for all u, v . Moreover, $d' + c$ is an ultrametric by Lemma 7.7, since $\ell_u = \tilde{d}(\rho, u) = d'(\rho, u)$ for all u . So the subdominant ultrametric e satisfies $d'(u, v) + c(u, v) \leq e(u, v)$, i.e. $d'(u, v) \leq \hat{d}(u, v)$, which by the guarantee of Lemma 7.11 ensures that

$$\hat{d}(u, v) \geq d(u, v) - (2d(u, v)/L + 36)\epsilon.$$

Also, $2 \max\{\ell_u, \ell_v\} \leq d'(u, v) + c(u, v) \leq e(u, v)$ guarantees that we can apply Lemma 7.8 to show \hat{d} is an additive metric.

On the other hand, by construction $\hat{d}(u, v)$ is a metric such that $\hat{d}(u, v) \leq \tilde{d}_{\text{pre}}(u, v) + \epsilon'$. Therefore applying the upper bound argument from Lemma 7.9 shows that

$$\hat{d}(u, v) \leq \tilde{d}(u, v) + (\hat{d}(u, v)/L + 1)\epsilon',$$

so by Lemma 7.9

$$\hat{d}(u, v) \leq (\tilde{d}(u, v) + \epsilon')/(1 - \epsilon'/L) \leq ((d(u, v)/L + 1)\epsilon + \epsilon')/(1 - \epsilon'/L) \leq (2d(u, v)/L + 90)\epsilon,$$

where we used used $\epsilon'/L \leq 1/2$ and $\epsilon' = 44\epsilon$. This completes the proof. \square

8 Desteinerization: handling strong edges, step 2

In this section we present our DESTEINERIZE algorithm, which is the second ingredient in the tree metric reconstruction procedure underlying LEARNLWRBDDMODEL which handles the strong edges of the model (see Section 5.3). The essence of the DESTEINERIZE result is that if we reconstruct an additive metric \hat{d} (with Steiner nodes) which is close to a ground truth tree metric d *without* Steiner nodes, then it is possible to remove the Steiner nodes from \hat{d} without significantly changing the distances between nodes. In the process of proving this, we establish several useful basic lemmas (especially Lemma 8.4) which show that an accurate additive metric must “essentially” match the topology of the ground truth tree. In the following theorem, DESTEINERIZE refers to Algorithm 5.

Theorem 8.1 (Restatement of Theorem 5.8). *There exists a constant $C > 0$ such that the following is true. Let $L \geq 2\epsilon \geq 0$ and suppose (T, d) is an unknown tree metric on vertex set \mathcal{V} with maximum edge length L . Suppose \hat{d} is an additive metric on \mathcal{V} with Steiner tree representation \hat{T} such that*

$$|d(x, y) - \hat{d}(x, y)| \leq (d(x, y)/L + 1)\epsilon \quad (28)$$

for all $x, y \in \mathcal{V}$. Then DESTEINERIZE with (\hat{T}, \hat{d}) as input runs in $O(|\mathcal{V}|^2 + |V(\hat{T})|)$ time and outputs a tree metric (T', d') on \mathcal{V} such that

$$|d(x, y) - d'(x, y)| \leq (d(x, y)/L + 1)\epsilon + C\epsilon$$

for all $x, y \in \mathcal{V}$.

Remark 8.2. We contrast this theorem with the famous result of [21], which showed that any additive metric can be approximated by a tree metric (i.e. an additive metric without steiner nodes) while distorting distances by only a multiplicative factor of 8. Thus, [21] showed that if we are satisfied with a constant factor in the distance approximation, then we never need to use Steiner nodes in the tree. In contrast, Theorem 5.8 makes a stronger assumption, that the input additive metric is close to some tree metric, and makes an essentially stronger conclusion: the output $O(\epsilon/L)$ -additively approximates short distances and $(1 + O(\epsilon/L))$ -multiplicatively approximates long distances. This stronger approximation guarantee is needed downstream to reconstruct the tree model in local total variation distance.

Algorithm 5: DESTEINERIZE(\hat{T}, \hat{d})

Input: Additive metric \hat{d} on \mathcal{V} with Steiner tree representation \hat{T}

Output: Tree metric d' on \mathcal{V} with tree representation T'

- 1 While there exists a Steiner node of degree 1 or 2 in \hat{T} , delete this node; this can be done without changing \hat{d} on elements of \mathcal{V} by either deleting or merging edges.
 - 2 // Define clustering of nodes
 - 2 For every node v in \hat{T} define $f(v)$ to be the closest non-Steiner node according to \hat{d} , breaking ties in favor of the node with lowest graph distance and then among those ties in favor of the lowest index node. (In Claim 8.5 below we show that for every $w \in \mathcal{V}$, $f^{-1}(w)$ induces a subtree in \hat{T} .)
 - 3 // Construct new tree on non-Steiner nodes
 - 3 Let T' have vertex set \mathcal{V} . For each $u, v \in \mathcal{V}$, include edge (u, v) in $E(T')$ if there is a $u' \in f^{-1}(u)$ and $v' \in f^{-1}(v)$ such that $(u', v') \in E(\hat{T})$.
 - 4 Fix an arbitrary root $\rho \in \mathcal{V}$ and for all neighbors u of ρ let $d'(\rho, u) = \hat{d}(\rho, u)$.
 - 5 Assign lengths to the edges of T' in breadth-first order as follows. When assigning length to a given edge (u, v) , note that all edges on the path between ρ and u have already been assigned a length, so $d'(\rho, u)$ is well-defined. Let $d'(u, v) = \max(0, \hat{d}(\rho, v) - d'(\rho, u))$. (Observe that $d'(\rho, v) \geq \hat{d}(\rho, v)$.)
 - 6 Return (T', d')
-

We now prove correctness of DESTEINERIZE. In what follows, let \mathcal{V} denote the non-Steiner nodes of \hat{T} , and let $\mathcal{X} = V(\hat{T})$ denote the Steiner nodes. Let C_{radius} denote the maximum distance in \hat{d} between a Steiner node to the closest non-Steiner node:

$$C_{\text{radius}} = \max_{x \in \mathcal{X} \setminus \mathcal{V}} \min_{v \in \mathcal{V}} \hat{d}(x, v).$$

Equivalently, C_{radius} is the smallest radius such that for every Steiner node $x \in \mathcal{X} \setminus \mathcal{V}$, there exists a non-Steiner node $v \in \mathcal{V}$ with $\hat{d}(x, v) \leq C_{\text{radius}}$. Theorem 5.8 immediately follows from the next two lemmas. Essentially, Lemma 8.3 implies that the desteinerization procedure distorts distances by at most $O(C_{\text{radius}})$. Lemma 8.4 guarantees that, since \hat{d} is a good approximation to a tree metric d without Steiner nodes, every Steiner node has a non-Steiner node at distance $O(\epsilon)$:

Lemma 8.3. *DESTEINERIZE with input \hat{d} outputs a tree metric d' such that $|d'(x, y) - d(x, y)| \leq 12C_{\text{radius}}$ for all $x, y \in \mathcal{V}$.*

Lemma 8.4. *Let \hat{d} be an additive metric that approximates a tree metric d (without Steiner nodes) as in the setting of Theorem 5.8. Then, after Step 1 of DESTEINERIZE, $C_{\text{radius}} \leq 30\epsilon$.*

8.1 Proof of Lemma 8.3: distances changed by $O(C_{\text{radius}})$

Lemma 8.3 states that DESTEINERIZE outputs a tree metric that distorts distances by at most $O(C_{\text{radius}})$. To prove it, we first show that the graph T' constructed by DESTEINERIZE is indeed a tree.

Claim 8.5. *In Step 2 of DESTEINERIZE, for any input additive metric \hat{d} and any $w \in \mathcal{V}$, the subgraph of \hat{T} induced by $f^{-1}(w)$, $\hat{T}[f^{-1}(w)]$, is a subtree.*

Proof. $\hat{T}[f^{-1}(w)]$ is nonempty because $w \in f^{-1}(w)$. So to show that $\hat{T}[f^{-1}(w)]$ is a subtree it suffices to show that this subgraph is connected. To this end, for each node $u \in f^{-1}(w)$, it suffices to show that all nodes on the path $\hat{\pi}(u, w) \subseteq \hat{T}$ between u and w are in $f^{-1}(w)$. Denote by $a_1 = u, a_2, \dots, a_t = w$ the path $\hat{\pi}(u, w)$ and let $i = \min\{j : f(a_j) \neq w\}$ be the first node that is *not* mapped to w . Write $w' = f(a_i)$. Then by the triangle inequality,

$$\hat{d}(a_{i-1}, w') \leq \hat{d}(a_i, w') + \hat{d}(a_{i-1}, a_i) \stackrel{(*)}{\leq} \hat{d}(a_i, w) + \hat{d}(a_{i-1}, a_i) = \hat{d}(a_{i-1}, w),$$

where the inequality $(*)$ follows by the definition of f , and the last equality step uses that $a_{i-1}, a_i, \dots, a_t = w$ form a path in \hat{T} . If $(*)$ is strict, then $\hat{d}(a_{i-1}, w') < \hat{d}(a_{i-1}, w)$ and this contradicts the definition of i as being the first in the path that is not mapped to w . Otherwise, if $\hat{d}(a_{i-1}, w') = \hat{d}(a_{i-1}, w)$, then $(*)$ is an equality. However, by the definition of f , this means that w' is at least as close as w to a_i in graph distance. If w' is strictly closer than w to a_i in graph distance, then w' is also strictly closer than w to a_i in graph distance, meaning that $f(a_{i-1}) = w'$, a contradiction. Otherwise, if w' is as close as w to a_i in graph distance, then $w' < w$ because $f(a_i) = w'$, and moreover w' is at least as close as w to a_{i-1} in graph distance, so $f(a_{i-1}) = w'$, a contradiction. \square

Claim 8.6. *The graph T' constructed in Step 3 is a tree.*

Proof. As shown in the previous claim, the sets $f^{-1}(w)$ for $w \in \mathcal{V}$ partition the vertex set of \hat{T} into subtrees $\hat{T}[f^{-1}(w)]$. The graph T' is obtained from \hat{T} by contracting the sets $\hat{T}[f^{-1}(w)]$. In general, contraction of subsets of vertices maintains connectivity. Moreover, considering the contraction of sets performed sequentially, a first such contraction of $\hat{T}[f^{-1}(w)]$ leading to a cycle would necessarily imply existence of a cycle before contraction due to connectivity of $\hat{T}[f^{-1}(w)]$, contradicting that \hat{T} is a tree. \square

Now that we know by that T' is a tree and so d' is a tree metric, it remains to bound the error. We next show that all distances to ρ are accurately approximated.

Claim 8.7. *After completing Step 5, for all $u \in \mathcal{V}$ we have $\mathbf{d}'(\rho, u) - 2C_{\text{radius}} \leq \hat{\mathbf{d}}(\rho, u) \leq \mathbf{d}'(\rho, u)$.*

Proof. The bound $\hat{\mathbf{d}}(\rho, u) \leq \mathbf{d}'(\rho, u)$ is a consequence of the edge length assignment in step 5, chosen specifically to ensure this inequality, so we proceed with proving the other inequality.

Consider an arbitrary node $u \in \mathcal{V}$ and denote by $a_0 = \rho, \dots, a_t = u$ the path in \mathbf{T}' between ρ and u . If $\hat{\mathbf{d}}(\rho, u) \geq \mathbf{d}'(\rho, u)$ then there is nothing to prove, so assume that $\hat{\mathbf{d}}(\rho, u) < \mathbf{d}'(\rho, u)$. Let

$$i = \min\{i \in [t] : \hat{\mathbf{d}}(\rho, a_j) < \mathbf{d}'(\rho, a_j) \text{ for all } i \leq j \leq t\}$$

be the earliest index such that a_i, a_{i+1}, \dots, a_t are all too far from ρ in metric \mathbf{d}' as compared to $\hat{\mathbf{d}}$. Note that $i \geq 2$ since $\hat{\mathbf{d}}(\rho, a_1) = \mathbf{d}'(\rho, a_1)$. Moreover, step 5 of DESTEINERIZE assigns $\mathbf{d}'(a_{j-1}, a_j) = 0$ for j in the range $i \leq j \leq t$, so $\mathbf{d}'(\rho, a_{i-1}) = \mathbf{d}'(\rho, a_i) = \dots = \mathbf{d}'(\rho, a_t)$.

By step 3 of DESTEINERIZE, for each j , $i \leq j \leq t$, there is an edge (a_{j-1}, a_j) only if the edge $(b_{j-1}^{\text{out}}, b_j^{\text{in}})$ is in $\hat{\mathbf{T}}$ for some $b_{j-1}^{\text{out}} \in f^{-1}(a_{j-1}), b_j^{\text{in}} \in f^{-1}(a_j)$, which implies that the path in $\hat{\mathbf{T}}$ from ρ to b_t^{in} passes through the sequence $b_{i-1}^{\text{out}}, b_i^{\text{in}}, b_i^{\text{out}}, \dots, b_{t-1}^{\text{out}}, b_t^{\text{in}}$, whence $\hat{\mathbf{d}}(\rho, b_{i-1}^{\text{out}}) \leq \dots \leq \hat{\mathbf{d}}(\rho, b_t^{\text{in}})$.

Finally, the definition of C_{radius} implies that $\hat{\mathbf{d}}(v, w) \leq C_{\text{radius}}$ for any node $v \in f^{-1}(w)$, it follows that

$$\begin{aligned} \hat{\mathbf{d}}(\rho, a_t) &\geq \hat{\mathbf{d}}(\rho, b_t^{\text{in}}) - C_{\text{radius}} \geq \hat{\mathbf{d}}(\rho, b_{i-1}^{\text{out}}) - C_{\text{radius}} \geq \mathbf{d}'(\rho, a_{i-1}) - 2C_{\text{radius}} \\ &= \mathbf{d}'(\rho, a_t) - 2C_{\text{radius}}, \end{aligned}$$

which proves the claim. \square

To prove that the distance between any pair of nodes $u, v \in \mathcal{V}$ is well-approximated, we need to introduce some terminology.

Definition 8.8 (Lowest common ancestor (LCA)). Let \mathbf{T} be a tree, S be a subset of the vertices, and ρ a (root) vertex. We define $\mathbf{LCA}_{\rho}^{\mathbf{T}}(S)$ to be the *lowest common ancestor* of S with respect to ρ . More explicitly, it is the furthest vertex from ρ in the intersection of all paths from u to ρ for $u \in S$. If $S = \{y, z\}$ is a set of size 2, we write $\mathbf{LCA}_{\rho}^{\mathbf{T}}(\{y, z\}) = \mathbf{LCA}_{\rho}^{\mathbf{T}}(y, z)$ as shorthand.

The following crucial claim shows that if we know the distances between three nodes u, v , and ρ , then the distance of all of these points to their “midpoint” $\mathbf{LCA}_{\rho}^{\mathbf{T}}(u, v) = \mathbf{LCA}_{\rho}^{\mathbf{T}}(\rho, v) = \mathbf{LCA}_{\rho}^{\mathbf{T}}(\rho, u)$ is also determined. This allows us to “triangulate” the location of nodes in the tree.

Claim 8.9. *Let \mathbf{d} be a metric on tree \mathbf{T} rooted at node ρ . Then for any $u, v \in V(\mathbf{T})$,*

$$\mathbf{d}(\mathbf{LCA}_{\rho}^{\mathbf{T}}(u, v), \rho) = \frac{\mathbf{d}(u, \rho) + \mathbf{d}(v, \rho) - \mathbf{d}(u, v)}{2}.$$

Proof. This follows from $\mathbf{d}(u, \rho) = \mathbf{d}(u, \mathbf{LCA}_{\rho}^{\mathbf{T}}(u, v)) + \mathbf{d}(\mathbf{LCA}_{\rho}^{\mathbf{T}}(u, v), \rho)$, the matching expression for $\mathbf{d}(v, \rho)$, and $\mathbf{d}(u, v) = \mathbf{d}(u, \mathbf{LCA}_{\rho}^{\mathbf{T}}(u, v)) + \mathbf{d}(\mathbf{LCA}_{\rho}^{\mathbf{T}}(u, v), v)$. \square

We now complete the proof of Lemma 8.3 by applying Claim 8.9 to both $\hat{\mathbf{d}}(u, v)$ and $\mathbf{d}'(u, v)$. As shown in Claim 8.7, for every pair of nodes $u, v \in \mathcal{V}$, the constructed tree metric \mathbf{d}' approximates distances $\hat{\mathbf{d}}(u, \rho)$ and $\hat{\mathbf{d}}(v, \rho)$ up to at most $2C_{\text{radius}}$. We also claim that

$$|\hat{\mathbf{d}}(\mathbf{LCA}_{\rho}^{\hat{\mathbf{T}}}(u, v), \rho) - \mathbf{d}'(\mathbf{LCA}_{\rho}^{\mathbf{T}'}(u, v), \rho)| \leq 4C_{\text{radius}}, \quad (29)$$

which implies the $|\hat{d}(u, v) - d'(u, v)| \leq 12C_{\text{radius}}$ by Claim 8.9. To show (29), first

$$\begin{aligned} |\hat{d}(\text{LCA}_{\rho}^{\hat{T}}(u, v), \rho) - d'(\text{LCA}_{\rho}^{T'}(u, v), \rho)| &\leq |\hat{d}(\text{LCA}_{\rho}^{\hat{T}}(u, v), \rho) - \hat{d}(\text{LCA}_{\rho}^{T'}(u, v), \rho)| \\ &\quad + |\hat{d}(\text{LCA}_{\rho}^{T'}(u, v), \rho) - d'(\text{LCA}_{\rho}^{T'}(u, v), \rho)|. \end{aligned}$$

The second quantity is at most $2C_{\text{radius}}$ by Claim 8.7 (these are distances from ρ under \hat{d} and d' of the same node $\text{LCA}_{\rho}^{T'}(u, v) \in \mathcal{V}$). For the first quantity, note that all paths in T' are the same as in \hat{T} if we contract the sets of vertices $f^{-1}(w)$ for all $w \in \mathcal{V}$. Thus, both $\text{LCA}_{\rho}^{\hat{T}}(u, v)$ and $\text{LCA}_{\rho}^{T'}(u, v)$ lie in the same subtree $\hat{T}[f^{-1}(w)]$ for some w , and the bound follows from the fact that these subtrees have diameter at most $2C_{\text{radius}}$, since by construction all nodes in $\hat{T}[f^{-1}(w)]$ are at distance at most C_{radius} from w . \square

8.2 Proof of Lemma 8.4: C_{radius} is small

Lemma 8.4 states that if \hat{d} is a good estimate of a tree metric d in the sense of Theorem 5.8, then after Step 1 of DESTEINERIZE every Steiner node of \hat{d} has a $O(\epsilon)$ -close non-Steiner node. First, we prove the following technical claim.

Claim 8.10. *Suppose that d is a tree metric on tree T and that y is on the unique path from x to z . Suppose that y' satisfies $|d(y', x) - d(y, x)| \leq \epsilon$ and $|d(y', z) - d(y, z)| \leq \epsilon$. Then $d(y, y') \leq \epsilon$.*

Proof. Consider the paths from x to y' and from z to y' . One of these paths must pass through y , since their concatenation forms a (non-simple, i.e., with repeated nodes) path from x to z . Without loss of generality assume this is the path from x to y' . Then

$$d(x, y') = d(x, y) + d(y, y'),$$

which means that $d(y, y') = d(x, y') - d(x, y) \leq \epsilon$ using the assumption of the lemma. \square

Now we are ready to prove Lemma 8.4. Let T be the tree associated to tree metric d in the statement of Lemma 8.4 and let \hat{T} be the tree associated with \hat{d} after step 1 of DESTEINERIZE.

Proof of Lemma 8.4. Fix an arbitrary root $\rho \in \mathcal{V}$. Pick an arbitrary Steiner node $v \in \mathcal{X}$ of \hat{d} . The node v must have at least two children w_1, w_2 with descendants in \mathcal{V} , since otherwise Step 1 of DESTEINERIZE would have removed v . We claim that in \hat{T} there must be at least one descendant $x_1 \in \mathcal{V}$ of w_1 such that $\hat{d}(v, x_1) \leq L + 2\epsilon$. To prove this, take an arbitrary node $x'_1 \in \mathcal{V}$ in the \hat{T} -subtree under w_1 , then by assumption there exists a path from x'_1 to ρ in T using edges of length at most L , and by assumption these nodes are at most $(L + 2\epsilon)$ -far apart in \hat{T} ; therefore going up this path, the last node which is a \hat{T} -descendant of w_1 must satisfy the condition to be x_1 . By the symmetrical argument, we can find a node x_2 descending from w_2 satisfying $\hat{d}(v, x_2) \leq L + 2\epsilon$. By a similar argument, we may find an ancestor $w \in \mathcal{V}$ of v in the tree \hat{T} such that $\hat{d}(v, w) \leq L + 2\epsilon$.

Note that by construction $v = \text{LCA}_{\rho}^{\hat{T}}(x_1, x_2)$, and furthermore that $v = \text{LCA}_w^{\hat{T}}(x_1, x_2)$, since w lies on the path from ρ to v in \hat{T} . Define $v' = \text{LCA}_w^T(x_1, x_2)$. By Claim 8.9, we have

$$\hat{d}(v, w) = \frac{\hat{d}(x_1, w) + \hat{d}(x_2, w) - \hat{d}(x_1, x_2)}{2} \text{ and } d(v', w) = \frac{d(x_1, w) + d(x_2, w) - d(x_1, x_2)}{2}. \quad (30)$$

Now use the fact that for any $a, b \in \{x_1, x_2, w\}$, we have $\hat{d}(a, b) \leq 2L + 4\epsilon$ by the triangle inequality. By applying (28), since $\epsilon/L \leq 1/2$ we obtain

$$d(a, b) \leq (2L + 5\epsilon)/(1 - \epsilon/L) \leq 4L + 10\epsilon,$$

and consequently applying (28) again

$$|\mathbf{d}(a, b) - \hat{\mathbf{d}}(a, b)| \leq ((4L + 10\epsilon)/L + 1)\epsilon \leq 10\epsilon.$$

Plugging this approximation guarantee into (30) yields

$$|\hat{\mathbf{d}}(v, w) - \mathbf{d}(v', w)| \leq 15\epsilon$$

So $\mathbf{d}(v', w) \leq 15\epsilon + L + 2\epsilon \leq 2L$, and thus by (28), $|\mathbf{d}(v', w) - \hat{\mathbf{d}}(u, v)| \leq 4\epsilon$, which means that

$$|\hat{\mathbf{d}}(v, w) - \hat{\mathbf{d}}(v', w)| \leq 20\epsilon. \quad (31)$$

Note that v' is on the \mathbf{T} path from x_1 to w , so $\mathbf{d}(x_1, w) = \mathbf{d}(x_1, v') + \mathbf{d}(v', w)$, and also v is on the $\hat{\mathbf{T}}$ path from x_1 to w , so $\hat{\mathbf{d}}(x_1, v) = \hat{\mathbf{d}}(x_1, v) + \hat{\mathbf{d}}(v, w)$. Combined with the fact that $|\hat{\mathbf{d}}(x_1, w) - \mathbf{d}(x_1, w)| \leq 10\epsilon$ and (31), this means

$$|\hat{\mathbf{d}}(v, x_1) - \hat{\mathbf{d}}(v', x_1)| \leq 30\epsilon. \quad (32)$$

Therefore, applying Claim 8.10 with (31) and (32) proves $\hat{\mathbf{d}}(v', v) \leq 30\epsilon$. Since v was arbitrary, $C_{\text{radius}} \leq 30\epsilon$. \square

9 Discussion

While this paper settles the problem of efficiently learning tree Ising models to within local total variation, including with model misspecification or adversarial modification of the data, a variety of questions remain. For one, we think of k in the order of $\text{locTV}^{(k)}$ as constant, and we leave open the task of efficiently learning a tree achieving optimal dependence on k in the $\text{locTV}^{(k)}$ error.

We assumed an Ising model without external field and it would be interesting to generalize the results to allow for arbitrary external field or to arbitrary higher cardinality alphabets and continuous models such as Gaussian models. Beyond the binary alphabet case, the component of Chow-Liu++ based on tree metric reconstruction algorithms may need some new ingredients.

Perhaps the most compelling open direction is developing efficient algorithms for estimating non-tree graphical models within local total variation.

A Proof of Proposition 5.1

We prove Proposition 5.1, which states that learning general models reduces to learning ferromagnetic models. The reduction is given in Algorithm 6.

Algorithm 6: LEARNMODEL($\tilde{\mu}, \epsilon$), i.e., the Chow-Liu++ algorithm

Input: ϵ -approximate correlations $\tilde{\mu}$ for an unknown model (\mathbf{T}, θ) , and $\epsilon > 0$

Output: Model $(\hat{\mathbf{T}}, \hat{\theta})$ that $O(\epsilon)$ -approximates (\mathbf{T}, θ) in $\text{locTV}^{(2)}$ distance.

- 1 $(\hat{\mathbf{T}}, \theta') \leftarrow \text{LEARNFERROMODEL}(|\hat{\mu}|, \epsilon)$
 - 2 $\hat{\theta}_e \leftarrow \text{sgn}(\tilde{\mu}_e) \cdot \theta'$ for all $e \in E(\hat{\mathbf{T}})$
 - 3 **return** $(\hat{\mathbf{T}}, \hat{\theta})$
-

Proposition A.1 (Restatement of Proposition 5.1). *The ϵ -tree-learning problem for general models on n vertices can be solved in $O(n^2)$ time and one call to an oracle LEARNFERROMODEL for the ϵ -tree-learning problem restricted to ferromagnetic models.*

Proof of Proposition 5.1. The runtime guarantee is straightforward, as there are n^2 correlations $\tilde{\mu}_{uv}$. For correctness, we begin by noting that $\|\tilde{\mu}_{uv} - \mu_{uv}\| \leq |\tilde{\mu}_{uv} - \mu_{uv}| \leq \epsilon$ for each u, v , and that the ferromagnetic model $(\mathsf{T}, |\theta|)$ has correlations $|\mu_{uv}|$. In words, the entries of $|\tilde{\mu}|$ are within ϵ of the correlations for $(\mathsf{T}, |\theta|)$ and therefore by the guarantee of LEARNFERROMODEL the returned model $(\widehat{\mathsf{T}}, \theta')$ has pairwise correlations μ' such that $\max_{u,v} \|\mu'_{uv} - \mu_{uv}\| \leq C\epsilon$ for some absolute constant C .

Let $\hat{\mu}$ denote the correlations in the final model $(\widehat{\mathsf{T}}, \hat{\theta})$. Since $|\hat{\theta}| = |\theta'|$ and hence $|\hat{\mu}| = |\mu'|$, we have

$$\max_{u,v} \|\hat{\mu}_{uv} - \mu_{uv}\| \leq C\epsilon. \quad (33)$$

It remains to show that $\hat{\mu}_{uv}$ has roughly the same sign as μ_{uv} for any u, v . We consider two cases.

Case 1. If $|\hat{\mu}_{uv}| \leq (C+1)\epsilon$, then we do not need the correct sign, since by the triangle inequality and (33)

$$|\hat{\mu}_{uv} - \mu_{uv}| \leq |\hat{\mu}_{uv}| + |\mu_{uv}| \leq 2|\hat{\mu}_{uv}| + C\epsilon \leq (3C+2)\epsilon.$$

Case 2. If $|\hat{\mu}_{uv}| > (C+1)\epsilon$, then it suffices to prove that $\text{sgn}(\mu_{uv}) = \text{sgn}(\hat{\mu}_{uv})$, because then

$$|\hat{\mu}_{uv} - \mu_{uv}| = \|\hat{\mu}_{uv} - \mu_{uv}\| \leq C\epsilon.$$

We consider the vertices w_1, \dots, w_t on the path from u to v in $\widehat{\mathsf{T}}$, where $u = w_1$ and $v = w_t$. We first show that the true correlations on this path are not too small: for any $i \in [t-1]$,

$$|\mu_{w_i w_{i+1}}| \geq |\hat{\mu}_{w_i w_{i+1}}| - C\epsilon \geq |\hat{\mu}_{uv}| - C\epsilon > \epsilon. \quad (34)$$

The first inequality follows from (33), and the second uses $|\hat{\mu}_{uv}| = \prod_{j=1}^{t-1} |\hat{\mu}_{w_j w_{j+1}}| \leq |\hat{\mu}_{w_i w_{i+1}}|$. Inequality (34) implies that the approximate correlations $\tilde{\mu}$ have the correct sign, that is,

$$\text{sgn}(\mu_{w_i w_{i+1}}) = \text{sgn}(\tilde{\mu}_{w_i w_{i+1}}). \quad (35)$$

This yields

$$\begin{aligned} \text{sgn}(\mu_{uv}) &= \prod_{i=1}^{t-1} \text{sgn}(\mu_{w_i w_{i+1}}) && \text{by Claim A.2 since model is tree-structured} \\ &= \prod_{i=1}^{t-1} \text{sgn}(\tilde{\mu}_{w_i w_{i+1}}) && \text{by (35)} \\ &= \prod_{i=1}^{t-1} \text{sgn}(\hat{\theta}_{w_i w_{i+1}}) = \text{sgn}(\hat{\mu}_{uv}) && \text{by construction of } (\widehat{\mathsf{T}}, \hat{\theta}). \end{aligned}$$

We see that in both cases $|\mu'_{uv} - \mu_{uv}| \leq (3C+2)\epsilon$. \square

Claim A.2. *Let μ the pairwise correlations for tree-structured model (T, θ) with vertex set \mathcal{V} . Let $w_1, \dots, w_t \in \mathcal{V}$ be some sequence of vertices. If $|\mu_{w_i w_{i+1}}| > 0$ for all $i \in [t-1]$, then*

$$\text{sgn}(\mu_{w_1 w_t}) = \prod_{i=1}^{t-1} \text{sgn}(\mu_{w_i w_{i+1}}).$$

We emphasize that it is not the case that $\mu_{w_1 w_t} = \prod \mu_{w_i w_{i+1}}$ (or there would be nothing to prove).

Proof. For each $i \in [t-1]$, define $\pi_i \subset \mathsf{T}$ to be the path between w_i and w_{i+1} . Let $\pi \subset \mathsf{T}$ be the path between w_1 and w_t . The key observation is that set $E(\pi)$ of edges in π is the set of edges that occurs in an odd number of the paths π_1, \dots, π_{t-1} . From this we obtain the desired statement,

$$\operatorname{sgn} \left(\prod_{i=1}^{t-1} \mu_{w_i w_{i+1}} \right) = \operatorname{sgn} \left(\prod_{i=1}^{t-1} \prod_{e \in E(\pi_i)} \theta_e \right) = \operatorname{sgn} \left(\prod_{e \in E(\pi)} \theta_e \right) = \operatorname{sgn}(\mu_{w_1 w_t}).$$

□

B Prediction-Centric Learning Generalizes Structure Learning

As discussed in the introduction, it is not possible to guarantee structure recovery when the strength of interactions is unbounded: if we allow arbitrarily strong interactions, then any spanning tree can perfectly represent the uniform distribution over the set of n all-equal bits (when $n = 3$, the model $X \sim \text{Uni}\{(1, 1, 1), (-1, -1, -1)\}$); if we allow arbitrarily weak interactions, then very weak edges cannot be distinguished from non-edges with any bounded number of samples.

In this section only, we consider the case where interactions are bounded, so structure recovery is possible. We show in this setting that prediction-centric learning automatically implies we get the correct structure. This justifies the claim that prediction-centric learning is the natural generalization of structure learning to models with arbitrarily weak or strong interactions, including settings with hard constraints or nearly independent variables. Afterwards, we show combining the reduction with our main result gives an algorithm with optimal sample complexity for structure learning.

Formally, we show the following:

Theorem B.1. *Suppose that (T, θ) is an Ising model P on the tree T with edge weights θ , and that for all edges $i \sim j$ that $|\theta_{ij}| \in [\alpha, 1 - \beta]$ with $\alpha, \beta \in (0, 1)$. Let $\epsilon = \alpha\beta/8$ and suppose (T', θ') is any tree Ising model P' with $\text{locTV}^{(3)}(P, P') < \epsilon$, then $\mathsf{T} = \mathsf{T}'$.*

Proof. We use the following notational convention: random vector X denotes a sample from the Ising model (T, θ) and random vector X' denotes a sample from the Ising model (T', θ') .

Suppose that $i \sim j$ in T but they are not neighbors in T' , so there exists at least one other node k which lies on the path from i to j . Also suppose without loss of generality that $E[X_i X_j] \geq 0$ (otherwise, flip the sign of X_j). Using the local TV assumption, by Fact 2.3, we have

$$E[X_i X_j] - 2\epsilon \leq E[X'_i X'_j] = E[X'_i X'_k] E[X'_k X'_j] \leq (E[X_i X_k] + 2\epsilon)(E[X_k X_j] + 2\epsilon). \quad (36)$$

In T , either the path from k to j passes through i or the path from k to i passes through j , depending on which node k is closer to. Assume without loss of generality (by symmetry) that the former case holds: then $E[X_j X_k] = E[X_i X_j] E[X_i X_k]$. Expanding the RHS of (36) then gives

$$(E[X_i X_k] + 2\epsilon)(E[X_i X_j] E[X_i X_k] + 2\epsilon) = E[X_i X_k]^2 E[X_i X_j] + 2\epsilon(E[X_i X_k] + E[X_i X_j]) + 4\epsilon^2$$

and combining with (36) we see that

$$E[X_i X_j] \leq E[X_i X_k]^2 E[X_i X_j] + 8\epsilon$$

and so

$$\alpha \leq \mathbb{E}[X_i X_j] \leq \frac{8\epsilon}{1 - \mathbb{E}[X_i X_k]^2} \leq \frac{8\epsilon}{1 - |\mathbb{E}[X_i X_k]|} \leq \frac{8\epsilon}{\beta}$$

which gives a contradiction using $\epsilon < \alpha\beta/8$. \square

Recall from (3) that $\text{locTV}^{(3)}$ and $\text{locTV}^{(2)}$ are equivalent up to constants for tree models. Using this, we establish optimality of the result.

Remark B.2 (Optimality). The following example shows the optimality, up to constants, of Theorem B.1. Suppose in T that i is the parent of j and k , $\theta_{ik} = 1 - \beta$ and $\theta_{ij} = \alpha$. Suppose that in T' , k is the parent of i and j with $\theta'_{ik} = 1 - \beta$ and $\theta'_{kj} = \alpha$. Then the two models match in all degree 2 moments except two: $\mathbb{E}[X_i X_j] = \alpha$ whereas $\mathbb{E}[X'_i X'_j] = (1 - \beta)\alpha$, and $\mathbb{E}[X_j X_k] = (1 - \beta)\alpha$ whereas $\mathbb{E}[X'_j X'_k] = \alpha$. Since the two cases are symmetrical, by Fact 2.3 we see

$$\text{locTV}^{(2)}(P, P') = (\alpha - (1 - \beta)\alpha)/2 = \alpha\beta/2.$$

Theorem B.1 above combined with our Chow-Liu++ guarantee in Corollary 3.2 implies that Chow-Liu++ successfully recovers the structure of a tree Ising model at the same optimal sample complexity as the Chow-Liu algorithm, up to constant factors.

Corollary B.3. *Let (T, θ) be a tree Ising model with edge correlations $|\theta_{ij}| \in [\alpha, 1 - \beta]$ with $\alpha, \beta \in (0, 1)$. There is a constant C such that given $m \geq C \log(n/\delta)/\alpha^2\beta^2$ samples, Chow-Liu++ returns the correct tree T with probability at least $1 - \delta$.*

Proof. Let $\epsilon = \alpha\beta/8$ and suppose the constant C describing m in the statement is sufficiently large. By Corollary 3.2, the learned tree model \hat{P} satisfies $\text{locTV}^{(3)}(\hat{P}, P) < ck2^k\epsilon < \epsilon$, where c can be chosen to be an arbitrarily small constant by growing C and P is the distribution of the original model (T, θ) . Theorem B.1 now implies that the tree underlying \hat{P} is correct. \square

Theorem 3.1 of [9], when converted to our parameterization, shows that $(\log n)/\alpha^2\beta^2$ samples are necessary to reconstruct the correct tree with constant probability. Thus the guarantee in the corollary just above proves the sample complexity is optimal up to constant factors.

C Failure of pure tree metric reconstruction approach

As explained in Section 5.2 above, our final learning algorithm combines the tree metric reconstruction based algorithm with the Chow-Liu algorithm, where the Chow-Liu algorithm is used as a kind of “clustering algorithm” which handles edges with weak correlation (large evolutionary distance). We already saw in Section 4 that purely using Chow-Liu fails (in a quite strong sense) to prove a guarantee similar to Theorem 3.1. In this section, we give a complementary result, showing that the general tree metric reconstruction approach outlined in Section 5.3 cannot by itself obtain the guarantee of Theorem 3.1 — in other words, the Chow-Liu preprocessing step is truly necessary to handle weak/long edges.

There are obvious technical issues trying to adapt the analysis of ADDITIVEMETRICRECONSTRUCTION to deal with misspecification, since δ -misspecification (or equivalent finite sample limitations) mean that evolutionary distances of size $\Omega(\log(1/\delta))$ are not at all accurate compared to the target tree model. Beyond those issues in the analysis, we next explain two different ways in which the pure tree metric reconstruction algorithm will fail, even with an infinite number of samples, both related to difficulties handling long edges; these failures show that if we wanted to remove the Chow-Liu step, then we would at the very least have to make the tree metric reconstruction procedure more complex.

First cause of failure: infinite distances. Since the tree metric reconstruction algorithm relies on reconstructing the evolutionary distance $d(u, v) = \log(1/\mu_{uv})$, the most obvious point of failure is that the chosen root node ρ may have correlation $\mu_{\rho u} = 0$ with all other nodes u , in which case the algorithm fails when run with population values for the correlations. Indeed, the evolutionary distance satisfies $d(\rho, u) = \infty$ and the centroid metric constructed by the algorithm satisfies $c(u, v) = \infty$ for all $u, v \neq \rho$. The algorithm fails since it must subtract an infinite value from an infinite value in order to output the reconstruction.

Second cause of failure: desteinerization and latent variables. The failure of the tree metric reconstruction algorithm is not simply due to the pathological case when there are perfectly-uncorrelated vertices. We demonstrate the issue with a counterexample of a different nature, where pairwise correlations are all nonzero.

We start by describing a distribution P over $\{\pm 1\}$ -valued random variables X, Y_1, Y_2 with $E[X] = E[Y_1] = E[Y_2] = 0$ and $E[Y_1 Y_2] = 1/4$, $E[XY_1] = E[XY_2] = \delta$. This distribution is not a tree Ising model, but it can be realized as the marginal law of a tree Ising model with an additional latent variable (i.e. Steiner node in the metric interpretation) Z :

1. X is sampled from a Rademacher distribution, i.e. uniform on $\{\pm 1\}$.
2. Z equals X with probability 2δ , otherwise Z is sampled independently from a Rademacher distribution.
3. With probability $1/2$, $Y_1 = Z$ and otherwise Y_1 is sampled from a Rademacher distribution. Y_2 is generated in the same way.

It's clear that P is within $\text{locTV}^{(2)}$ distance δ of a tree Ising model where X is independent of Y_1, Y_2 . We now explain why P is a problematic example for the pure tree metric reconstruction algorithm.

Outside of the present work, reconstruction algorithms based on the evolutionary metric are largely used to reconstruct tree models with latent variables, see e.g. [12], in which case it is considered a virtue that the output metric includes the Steiner nodes. However, in the present context it is a problem: when our algorithm runs the initial tree metric reconstruction step (ADDITIVEMETRICRECONSTRUCTION), the output will contain the Steiner node Z which is at distance $\log(1/2)$ from Y_1, Y_2 and at distance $\log(1/\delta)$ from X . Essentially, the problem is that the Steiner node is not close to any of the observable nodes, and this is an issue as the final output cannot have latent variables. More precisely, there is no way to desteinerize this tree in the sense of either (1) deleting the Steiner node if it unnecessary (i.e. of degree 2) or (2) replacing the steiner node/latent variable with an observed variable while keeping the other distances in the tree unchanged (as in DESTEINERIZE): replacing Z by any of Y_1, Y_2 , or X will result in a tree Ising model with $\text{locTV}^{(2)}$ distance $\Omega(1)$ from P . For example, replacing Z by Y_1 results in a model with $E[Y_1 Y_2] = 1/2$ instead of $1/4$.

D Optimal Lower bound for TV Learning

In recent work, Devroye et al [14] studied minimax rates for learning various kinds of graphical models in Total Variation (TV) distance. They proved that for a tree Ising model on n nodes, the minimax rate for reconstructing the tree model in TV from m samples is upper bounded by $O\left(\sqrt{\frac{n \log(n)}{m}}\right)$ and posed the tightness of this result as an open question. In this Appendix, we

resolve this open question by proving a matching information-theoretic lower bound, showing that the minimax rate is $\Theta\left(\sqrt{\frac{n \log n}{m}}\right)$ up to a universal constant.

Construction. Recall from Stirling’s formula that $\log(n!) \sim n \log(n)$. This motivates the following simple construction for the lower bound:

1. \mathcal{S} is a family of permutations on $[n]$ to be specified later.
2. Pick a permutation π from family of permutations \mathcal{S} .
3. Build an Ising model on the matching graph with covariance α/\sqrt{n} between vertices i and $n + \pi(i)$ for every $i \in [n]$. (This corresponds to edge weight essentially α/\sqrt{n} since the correlation for edge weight β is $\tanh(\beta) \approx \beta$)

This yields a family of distributions $\{P_\pi\}_{\pi \in \mathcal{S}}$ on $2n$ nodes.

For the set of permutations, we choose a set which satisfies $\log |\mathcal{S}| = \Omega(n \log(n))$ and such that every set of permutations has Hamming distance at least $n/4$, where the Hamming distance is $|\{x : \pi(x) \neq \pi'(x)\}|$. In the survey of Quistorff [35], such a result is given as (6) and attributed to Deza.

Analysis. The proof reduces to the following claims:

1. Every two elements have large total variation distance. More specifically, to distinguish models from π_1, π_2 look at statistic $\sum_{i=1}^n X_i X_{\pi_1(i)}$. Under P_{π_1} the expectation is $\alpha\sqrt{n}$ and variance is $\Theta(n)$, whereas under P_{π_2} the expectation is less than $(3/4)\alpha\sqrt{n}$ and variance is $\Theta(n)$. Applying the Central Limit Theorem, this shows the total variation distance is $\Omega(\alpha)$.
2. Every two elements have reasonably small KL. Recall by the Gibbs variational formula [16] that

$$KL(P_{\pi_1}, P_{\pi_2}) = \left(\frac{1}{2} \mathbb{E}_{\pi_2}[X^T J_{\pi_2} X] + H_{\pi_2}(X)\right) - \left(\frac{1}{2} \mathbb{E}_{\pi_1}[X^T J_{\pi_2} X] + H_{\pi_2}(X)\right).$$

Since the entropies are the same, this is just a difference of expectations. By similar reasoning to above, it is of order $\Theta(\alpha^2)$ (there are $n/4$ missing edges, and they each contribute $(\alpha/\sqrt{n}) \cdot (c\alpha/\sqrt{n})$).

3. Any algorithm given $m \leq C_2 n \log(n)/\alpha^2$ samples fails to reconstruct with probability at least $1/2$. Since between any two models the KL for m samples is of order $m\alpha^2$ by tensorization and (2), and $\log |\mathcal{S}| \sim n \log(n)$, this follows directly from Fano’s inequality (see e.g. [40]).

Combining these claims shows that $\sqrt{n \log(n)/m}$ is the tight rate for learning tree Ising models in TV distance.

References

- [1] Pieter Abbeel, Daphne Koller, and Andrew Y Ng. “Learning factor graphs in polynomial time and sample complexity”. In: *The Journal of Machine Learning Research* 7 (2006), pp. 1743–1788.

- [2] Richa Agarwala, Vineet Bafna, Martin Farach, Mike Paterson, and Mikkel Thorup. “On the approximability of numerical taxonomy (fitting distances by tree metrics)”. In: *SIAM Journal on Computing* 28.3 (1998), pp. 1073–1085.
- [3] Nir Ailon and Moses Charikar. “Fitting tree metrics: Hierarchical clustering and phylogeny”. In: *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS’05)*. IEEE. 2005, pp. 73–82.
- [4] Jean-Pierre Barthélémy and Alain Guénoche. *Trees and proximity representations*. John Wiley & Sons, 1991.
- [5] José M Bayod and J Martinez-Maurica. “Subdominant ultrametrics”. In: *Proceedings of the American Mathematical Society* 109.3 (1990), pp. 829–834.
- [6] Arnab Bhattacharyya, Sutanu Gayen, Eric Price, and NV Vinodchandran. “Near-Optimal Learning of Tree-Structured Distributions by Chow-Liu”. In: *arXiv preprint arXiv:2011.04144* (2020).
- [7] Jose Blanchet, Yang Kang, and Karthyek Murthy. “Robust Wasserstein profile inference and applications to machine learning”. In: *Journal of Applied Probability* 56.3 (2019), pp. 830–857.
- [8] Guy Bresler. “Efficiently learning Ising models on arbitrary graphs”. In: *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*. 2015, pp. 771–782.
- [9] Guy Bresler and Mina Karzand. “Learning a tree-structured Ising model in order to make predictions”. In: *Annals of Statistics* 48.2 (2020), pp. 713–737.
- [10] Guy Bresler, Frederic Koehler, and Ankur Moitra. “Learning restricted Boltzmann machines via influence maximization”. In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. 2019, pp. 828–839.
- [11] C Chow and Cong Liu. “Approximating discrete probability distributions with dependence trees”. In: *IEEE transactions on Information Theory* 14.3 (1968), pp. 462–467.
- [12] Constantinos Daskalakis, Elchanan Mossel, and Sébastien Roch. “Evolutionary trees and the Ising model on the Bethe lattice: a proof of Steel’s conjecture”. In: *Probability Theory and Related Fields* 149.1-2 (2011), pp. 149–189.
- [13] Constantinos Daskalakis and Qinxuan Pan. “Tree-structured Ising models can be learned efficiently”. In: *arXiv preprint arXiv:2010.14864* (2020).
- [14] Luc Devroye, Abbas Mehrabian, Tommy Reddad, et al. “The minimax learning rates of normal and Ising undirected graphical models”. In: *Electronic Journal of Statistics* 14.1 (2020), pp. 2338–2361.
- [15] Ilias Diakonikolas, Daniel M Kane, Alistair Stewart, and Yuxin Sun. “Outlier-Robust Learning of Ising Models Under Dobrushin’s Condition”. In: *arXiv preprint arXiv:2102.02171* (2021).
- [16] Richard S Ellis. *Entropy, large deviations, and statistical mechanics*. Springer, 2007.
- [17] Péter L Erdős, Michael A Steel, László A Székely, and Tandy J Warnow. “A few logs suffice to build (almost) all trees (I)”. In: *Random Structures & Algorithms* 14.2 (1999), pp. 153–184.
- [18] Peyman Mohajerin Esfahani and Daniel Kuhn. “Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations”. In: *Mathematical Programming* 171.1 (2018), pp. 115–166.
- [19] Martin Farach, Sampath Kannan, and Tandy Warnow. “A robust model for finding optimal evolutionary trees”. In: *Algorithmica* 13.1-2 (1995), pp. 155–179.

- [20] Surbhi Goel, Daniel M Kane, and Adam R Klivans. “Learning Ising models with independent failures”. In: *Conference on Learning Theory*. PMLR. 2019, pp. 1449–1469.
- [21] Anupam Gupta. “Steiner points in tree metrics don’t (really) help”. In: *Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms*. 2001, pp. 220–227.
- [22] Uri Heinemann and Amir Globerson. “Inferning with high girth graphical models”. In: *International Conference on Machine Learning*. PMLR. 2014, pp. 1260–1268.
- [23] Nicholas Jardine and Robin Sibson. *Mathematical taxonomy*. Tech. rep. 1971.
- [24] Ashish Katiyar, Vatsal Shah, and Constantine Caramanis. “Robust Estimation of Tree Structured Ising Models”. In: *arXiv preprint arXiv:2006.05601* (2020).
- [25] Jonathan Kelner, Frederic Koehler, Raghu Meka, and Ankur Moitra. “Learning some popular gaussian graphical models without condition number bounds”. In: *Advances in Neural Information Processing Systems*. 2021.
- [26] Adam Klivans and Raghu Meka. “Learning graphical models using multiplicative weights”. In: *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2017, pp. 343–354.
- [27] Jaeho Lee and Maxim Raginsky. “Minimax statistical learning with wasserstein distances”. In: *arXiv preprint arXiv:1705.07815* (2017).
- [28] Erik M Lindgren, Vatsal Shah, Yanyao Shen, Alexandros G Dimakis, and Adam Klivans. “On robust learning of ising models”. In: *NeurIPS Workshop on Relational Representation Learning*. 2019.
- [29] Elchanan Mossel and Sébastien Roch. “Learning nonsingular phylogenies and hidden Markov models”. In: *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*. 2005, pp. 366–375.
- [30] Fionn Murtagh. “A survey of recent advances in hierarchical clustering algorithms”. In: *The computer journal* 26.4 (1983), pp. 354–359.
- [31] Assaf Naor. “An introduction to the Ribe program”. In: *Japanese Journal of Mathematics* 7.2 (2012), pp. 167–233.
- [32] K. Nikolakakis, D. Kalogerias, and A. D. Sarwate. “Learning Tree Structures from Noisy Data”. In: *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics (AISTATS)*. Vol. 89. 16–18 April 2019, pp. 1771–1782.
- [33] Konstantinos E Nikolakakis, Dionysios S Kalogerias, and Anand D Sarwate. “Predictive Learning on Hidden Tree-Structured Ising Models”. In: *Journal of Machine Learning Research* 22.59 (2021), pp. 1–82.
- [34] Adarsh Prasad, Vishwak Srinivasan, Sivaraman Balakrishnan, and Pradeep Ravikumar. “On Learning Ising Models under Huber’s Contamination Model”. In: *Advances in Neural Information Processing Systems* 33 (2020).
- [35] Jörn Quistorff. “A survey on packing and covering problems in the Hamming permutation space”. In: *the electronic journal of combinatorics* 13.1 (2006), A1.
- [36] R Rammal, JC Angles d’Auriac, and Benoît Douçot. “On the degree of ultrametricity”. In: *Journal de Physique Lettres* 46.20 (1985), pp. 945–952.
- [37] Rammal Rammal, Gérard Toulouse, and Miguel Angel Virasoro. “Ultrametricity for physicists”. In: *Reviews of Modern Physics* 58.3 (1986), p. 765.

- [38] Pradeep Ravikumar, Martin J Wainwright, and John D Lafferty. “High-dimensional Ising model selection using l1-regularized logistic regression”. In: *The Annals of Statistics* 38.3 (2010), pp. 1287–1319.
- [39] Patrick Rebeschini, Ramon Van Handel, et al. “Can local particle filters beat the curse of dimensionality?” In: *Annals of Applied Probability* 25.5 (2015), pp. 2809–2866.
- [40] Phillippe Rigollet and Jan-Christian Hütter. “High dimensional statistics”. In: *Lecture notes for course 18S997* (2015).
- [41] Abhin Shah, Devavrat Shah, and Gregory Wornell. “On Learning Continuous Pairwise Markov Random Fields”. In: *International Conference on Artificial Intelligence and Statistics*. PMLR. 2021, pp. 1153–1161.
- [42] Michael Steel. “Recovering a tree from the leaf colourations it generates under a Markov model”. In: *Applied Mathematics Letters* 7.2 (1994), pp. 19–23.
- [43] Mike Steel. *Phylogeny: discrete and random processes in evolution*. SIAM, 2016.
- [44] Vincent YF Tan, Animashree Anandkumar, Lang Tong, and Alan S Willsky. “A large-deviation analysis of the maximum-likelihood learning of Markov tree structures”. In: *IEEE Transactions on Information Theory* 57.3 (2011), pp. 1714–1735.
- [45] Marc Vuffray, Sidhant Misra, Andrey Y Lokhov, and Michael Chertkov. “Interaction screening: efficient and sample-optimal learning of ising models”. In: *Proceedings of the 30th International Conference on Neural Information Processing Systems*. 2016, pp. 2603–2611.
- [46] Martin J Wainwright. “Estimating the” Wrong” Graphical Model: Benefits in the Computation-Limited Setting.” In: *Journal of Machine Learning Research* 7.9 (2006).
- [47] Martin J Wainwright and Michael I Jordan. *Graphical models, exponential families, and variational inference*. Now Publishers Inc, 2008.
- [48] Wolfram Wiesemann, Daniel Kuhn, and Melvyn Sim. “Distributionally robust convex optimization”. In: *Operations Research* 62.6 (2014), pp. 1358–1376.
- [49] Shanshan Wu, Sujay Sanghavi, and Alexandros G Dimakis. “Sparse logistic regression learns all discrete pairwise graphical models”. In: *Advances in Neural Information Processing Systems* 32 (2019), pp. 8071–8081.