

---

# Interactive Bayesian Hierarchical Clustering

---

Sharad Vikram  
Sanjoy Dasgupta

SVIKRAM@CS.UCSD.EDU  
DASGUPTA@CS.UCSD.EDU

Computer Science and Engineering, UCSD, 9500 Gilman Drive, La Jolla, CA 92093

## Abstract

Clustering is a powerful tool in data analysis, but it is often difficult to find a grouping that aligns with a user’s needs. To address this, several methods incorporate constraints obtained from users into clustering algorithms, but unfortunately do not apply to hierarchical clustering. We design an *interactive Bayesian* algorithm that incorporates user interaction into hierarchical clustering while still utilizing the geometry of the data by sampling a constrained posterior distribution over hierarchies. We also suggest several ways to intelligently query a user. The algorithm, along with the querying schemes, shows promising results on real data.

## 1. Introduction

*Clustering* is a basic tool of exploratory data analysis. There are a variety of efficient algorithms—including  $k$ -means, EM for Gaussian mixtures, and hierarchical agglomerative schemes—that are widely used for discovering “natural” groups in data. Unfortunately, they don’t always find a grouping that suits the user’s needs.

This is inevitable. In any moderately complex data set, there are many different plausible grouping criteria. Should a collection of rocks be grouped according to value, or shininess, or geological properties? Should animal pictures be grouped according to the Linnaean taxonomy, or cuteness? Different users have different priorities, and an unsupervised algorithm has no way to magically guess these.

As a result, a rich body of work on *constrained* clustering has emerged. In this setting, a user supplies guidance, typically in the form of “must-link” or “cannot-link” constraints, pairs of points that must be placed together or apart. Introduced by Wagstaff & Cardie (2000), these constraints have since been incorporated into many differ-

ent *flat* clustering procedures (Wagstaff et al., 2001; Bansal et al., 2004; Basu et al., 2004; Kulis et al., 2005; Biswas & Jacobs, 2014).

In this paper, we introduce constraints to *hierarchical clustering*, the recursive partitioning of a data set into successively smaller clusters to form a tree. A hierarchy has several advantages over a flat clustering. First, there is no need to specify the number of clusters in advance. Second, the tree captures cluster structure at multiple levels of granularity, simultaneously. As such, trees are particularly well-suited for exploratory data analysis and the discovery of natural groups.

There are several well-established methods for hierarchical clustering, the most prominent among which are the bottom-up agglomerative methods such as average linkage (see, for instance, Chapter 14 of Hastie et al. (2009)). But they suffer from the same problem of under-specification that is the scourge of unsupervised learning in general. And, despite the rich literature on incorporating additional guidance into flat clustering, there has been relatively little work on the hierarchical case.

What form might the user’s guidance take? The usual must-link and cannot-link constraints make little sense when data has hierarchical structure. Among living creatures, for instance, should `elephant` and `tiger` be linked? At some level, yes, but at a finer level, no. A more straightforward assertion is that `elephant` and `tiger` should be linked in a cluster that does not include `snake`. We can write this as a *triplet* (`{elephant, tiger}, snake`). We could also assert (`{tiger, leopard}, elephant`). Formally,  $(\{a, b\}, c)$  stipulates that the hierarchy contains a subtree (that is, a cluster) containing  $a$  and  $b$  but not  $c$ .

A wealth of research addresses learning taxonomies from triplets *alone*, mostly in the field of phylogenetics: see Felsenstein (2004) for an overview, and Aho et al. (1981) for a central algorithmic result. Let’s say there are  $n$  data items to be clustered, and that the user seeks a particular hierarchy  $T^*$  on these items. This  $T^*$  embodies at most  $\binom{n}{3}$  triplet constraints, possibly less if it is not binary. It was pointed out in Tamuz et al. (2011) that roughly  $n \log n$

carefully-chosen triplets are enough to fully specify  $T^*$  if it is balanced. This is also a lower bound: there are  $n^{\Omega(n)}$  different labeled rooted trees, so each tree requires  $\Omega(n \log n)$  bits, on average, to write down—and each triple provides  $O(1)$  bits of information, since there are just three possible outcomes for each set of points  $a, b, c$ . Although  $n \log n$  is a big improvement over  $n^3$ , it is impractical for a user to provide this much guidance when the number of points is large. In such cases, a hierarchical clustering cannot be obtained on the basis of constraints alone; the geometry of the data must play a role. We consider an interactive process during which a user incrementally adds constraints.

- Starting with a pool of data  $X \subseteq \mathbb{R}^d$ , the machine builds a candidate hierarchy  $T$ .
- The set of constraints  $C$  is initially empty.
- Repeat:
  - The machine presents the user with a small portion of  $T$ : specifically, its restriction to  $O(1)$  leaves  $S \subset X$ . We denote this  $T|_S$ .
  - The user either accepts  $T|_S$ , or provides a triplet constraint  $(\{a, b\}, c)$  that is violated by it.
  - If a triplet is provided, the machine adds it to  $C$  and modifies the tree  $T$  accordingly.

In realizing this scheme, a suitable clustering algorithm and querying strategy must be designed. Similar issues have been confronted in flat clustering—with must-link and cannot-link constraints—but the solutions are unsuitable for hierarchies, and thus a fresh treatment is warranted.

#### THE CLUSTERING ALGORITHM

What is a method of hierarchical clustering that takes into account the geometry of the data points as well as user-imposed constraints?

We adopt an *interactive Bayesian* approach. The learning procedure is uncertain about the intended tree and this uncertainty is captured in the form of a distribution over all possible trees. Initially, this distribution is informed solely by the geometry of the data but once interaction begins, it is also shaped by the growing set of constraints.

The nonparametric Bayes literature contains a variety of different distributional models for hierarchical clustering. We describe a general methodology for extending these to incorporate user-specified constraints. For concreteness, we focus on the Dirichlet diffusion tree (Neal, 2003), which has enjoyed empirical success. We show that triplet constraints are quite easily accommodated: when using a Metropolis-Hastings sampler, they can efficiently be enforced, and the state space remains strongly connected, assuring convergence to the unique stationary distribution.

#### THE QUERYING STRATEGY

What is a good way to select the subsets  $S$ ? A simple option is to pick them at random from  $X$ . We show that this strategy leads to convergence to the target tree  $T^*$ . Along the way, we define a suitable distance function for measuring how close  $T$  is to  $T^*$ .

We might hope, however, that a more careful choice of  $S$  would lead to faster convergence, in much the same way that intelligent querying is often superior to random querying in active learning. In order to do this, we show how the Bayesian framework allows us to quantify which portions of the tree are the most uncertain, and thereby to pick  $S$  that focuses on these regions.

Querying based on uncertainty sounds promising, but is dangerous because it is heavily influenced by the choice of prior, which is ultimately quite arbitrary. Indeed, if only such queries were used, the interactive learning process could easily converge to the wrong tree. We show how to avoid this situation by interleaving the two types of queries.

Finally, we present a series of experiments that illustrate how a little interaction leads to significantly better hierarchical clusterings.

#### 1.1. Other related work

A related problem that has been studied in more detail (Zoller & Buhmann, 2000; Eriksson et al., 2011; Krishnamurthy et al., 2012) is that of building a hierarchical clustering where the only information available is pairwise similarities between points, but these are initially hidden and must be individually queried.

In another variant of interactive flat clustering (Balcan & Blum, 2008; Awasthi & Zadeh, 2010; Awasthi et al., 2014), the user is allowed to specify that individual clusters be merged or split. A succession of such operations can always lead to a target clustering, and a question of interest is how quickly this convergence can be achieved.

Finally, it is worth mentioning the use of triplet constraints in learning other structures, such as Euclidean embeddings (Borg & Groenen, 2005).

## 2. Bayesian hierarchical clustering

The most basic form of hierarchical clustering is a rooted binary tree with the data points at its leaves. This is sometimes called a *cladogram*. Very often, however, the tree is adorned with additional information, for instance:

1. An ordering of the internal nodes, where the root is assigned the lowest number and each node has a higher number than its parent.

This ordering uniquely specifies the induced  $k$ -clustering (for any  $k$ ): just remove the  $k - 1$  lowest-numbered nodes and take the clusters to be the leaf-sets of the  $k$  resulting subtrees.

## 2. Lengths on the edges.

Intuitively, these lengths correspond to the amount of change (for instance, time elapsed) along the corresponding edges. They induce a *tree metric* on the nodes, and often, the leaves are required to be at the same distance from the root.

## 3. Parameters at internal nodes.

These parameters are sometimes from the same space as the data, representing intermediate values on the way from the root to the leaves.

Many generative processes for trees end up producing these more sophisticated structures, with the understanding that undesired additional information can simply be discarded at the very end. We now review some well-known distributions over trees and over hierarchical clusterings.

Let's start with cladograms on  $n$  leaves. The simplest distribution over these is the uniform. Another well-studied option is the *Yule model*, which can be described using either a top-down or bottom-up generative process. The top-down view corresponds to a continuous-time pure birth process: start with one lineage; each lineage persists for a random exponential(1) amount of time and then splits into two lineages; this goes on until there are  $n$  lineages. The bottom-up view is a coalescing process: start with  $n$  points; pick a random pair of them to merge; then repeat. [Aldous \(1995\)](#) has defined a one-parameter family of distributions over cladograms, called the *beta-splitting model*, that includes the uniform and the Yule model as special cases. It is a top-down generative process in which, roughly, each split is made by sampling from a Beta distribution to decide how many points go on each side. To move to arbitrary (not necessarily binary) splits, a suitable generalization is the Gibbs fragmentation tree ([McCullagh et al., 2008](#)).

In this paper, we will work with joint distributions over both tree structure and data. These are typically inspired by, or based directly upon, the simpler tree-only distributions described above. Our primary focus is the Dirichlet diffusion tree ([Neal, 2003](#)), which is specified by a birth process that we will shortly describe. However, our methodology applies quite generally. Other notable Bayesian approaches to hierarchical clustering include: [Williams \(2000\)](#), in which each node of the tree is annotated with a vector that is sampled from a Gaussian centered at its parent's vector; [Heller & Ghahramani \(2005\)](#), that defines a distribution over flat clusterings and then specifies an agglomerative scheme for finding a good partition with respect to this

distribution; [Adams et al. \(2008\)](#), in which data points are allowed to reside at internal nodes of the tree; [Teh et al. \(2008\)](#); [Boyles & Welling \(2012\)](#), in which the distribution over trees is specified by a bottom-up coalescing process; and [Knowles & Ghahramani \(2015\)](#), which generalizes the Dirichlet diffusion trees to allow non-binary splits.

## 2.1. The Dirichlet diffusion tree

The Dirichlet diffusion tree (DDT) is a generative model for  $d$ -dimensional vectors  $x_1, x_2, \dots, x_N$ . Data are generated sequentially via a continuous-time process, lasting from time  $t = 0$  to  $t = 1$ , when they reach their final value.

The first point,  $x_1$ , is generated via a Brownian motion, beginning at the origin, i.e.  $X_1(t + dt) = X_1(t) + \mathcal{N}(0, \sigma^2 I_d dt)$  where  $X_1(t)$  represents the value of  $x_1$  at time  $t$ . The next point,  $x_2$ , follows the path created by  $x_1$  until it eventually *diverges* at a random time, according to a specified acquisition function  $a(t)$ . When  $x_2$  diverges, it creates an internal node in the tree structure which contains both the time and value of  $x_2$  when it diverged. After divergence, it continues until  $t = 1$  with an independent Brownian motion. In general, the  $i$ -th point follows the path created by the previous  $i - 1$  points. When it reaches a node, it will first sample one of two branches to enter, then either 1) diverge on the branch, whereupon a divergence time is sampled according to the acquisition function  $a(t)$ , or 2) recursively continue to the next node. Each of these choices has a probability associated with it, according to various properties of the tree structure and choice of acquisition function (details can be found in [Neal \(2003\)](#)). Eventually, all points will diverge and continue independently, creating an internal node storing the time and intermediate value for each point at divergence. The DDT thus defines a binary tree over the data (see [Figure 1](#) for an example). Furthermore, given a DDT with  $N$  points, it is possible to sample the possible divergence locations of a  $(N + 1)$ -th point, using the generative process.

Sampling the posterior DDT given data can be done with the Metropolis-Hastings (MH) algorithm, an MCMC method. The MH algorithm obtains samples from target distribution  $p(x)$  indirectly by instead sampling from a conditional "proposal" distribution  $q(x|x')$ , creating a Markov chain whose stationary distribution is  $p(x)$ , assuming  $q$  satisfies some conditions. Our choice of proposal distribution modifies the DDT's tree structure via a *subtree-prune and regraft* (SPR) move, which has the added benefit of extending to other distributions over hierarchies.

An SPR move consists of first a *prune* then a *regraft*. Suppose  $T$  is a binary tree with  $n$  leaves. Let  $s$  be a non-root node in  $T$  selected uniformly at random and  $S$  be its corresponding subtree. To prune  $S$  from  $T$ , we remove  $s$ 's parent  $p$  from  $T$ , and replace  $p$  with  $s$ 's sibling.




Figure 1. An example DDT with 1-dimensional data. Blue lines represent paths taken by each data point, and black dots represent nodes of the tree. The rightmost dots are leaves and the others are nodes created when points diverged. When drawing the hierarchy, typically the top stem is omitted.

Regrafting selects a branch at random and attaches  $S$  to it as follows. Let  $(u, v)$  be the chosen branch ( $u$  is the parent of  $v$ ).  $S$  is attached to the branch by creating a node  $p$  with children  $s$  and  $v$  and parent  $u$  (see Figure 2 for an example).

The MH proposal distribution for the DDT is an augmented SPR move, where the time and intermediate value at each node are sampled in addition to tree structure. The exchangeability of the DDT enables efficient sampling of regraft branches by simulating the generation process for a new point and returning the branch and time where it diverges. The intermediate values for the entire tree are sampled via an interleaved Gibbs sampling move, as all conditional distributions are Gaussian.

### 3. Adding interaction

Impressive as the Dirichlet diffusion tree is, there is no reason to suppose that it will magically find a tree that suits the user’s needs. But a little interaction can be helpful in improving the outcome.

Let  $T^*$  denote the target hierarchical clustering. It is not necessarily the case that the user would be able to write this down explicitly, but this is the tree that captures the distinctions he/she is able to make, or wants to make. Figure 3 (left) shows an example, for a small data set of 5 points. In this case, the user does not wish to distinguish between points 1, 2, 3, but does wish to place them in a cluster that excludes point 4.

We could posit our goal as exactly recovering  $T^*$ . But in many cases, it is good enough to find a tree that captures all the distinctions within  $T^*$  but also possibly has some extraneous distinctions, as in the right-hand side of Figure 3.

Formally, given data set  $X$ , we say  $S \subseteq X$  is a *cluster* of tree  $T$  if there is some node of  $T$  whose descendant leaves are exactly  $S$ . We say  $T$  is a *refinement* of  $T^*$  if they have the same set of leaves, and moreover every cluster of  $T^*$  is also a cluster of  $T$ . This, then, is our goal: to find a refinement of the target clustering  $T^*$ .

#### 3.1. Triplets

The user provides feedback in the form of triplets. The constraint  $(\{a, b\}, c)$  means that the tree should have a cluster containing  $a$  and  $b$  but not  $c$ . Put differently, the lowest common ancestor of  $a, b$  should be a strict descendant of the lowest common ancestor of  $a, b, c$ .

Let  $\Delta(T)$  denote the set of all proper triplet constraints embodied in tree  $T$ . If  $T$  has  $n$  nodes, then  $|\Delta(T)| \leq \binom{n}{3}$ . For non-binary trees, it will be smaller than this number. Figure 3 (left), for instance, has no triplet involving 1, 2, 3.

Refinement can be characterized in terms of triplets.

**Lemma 3.1.** *Tree  $T$  is a refinement of tree  $T'$  if and only if  $\Delta(T') \subseteq \Delta(T)$ .*

*Proof.* See supplement.  $\square$

In particular, *any triplet-querying scheme that converges to the full set of triplets of the target tree  $T^*$  is also guaranteed to produce trees that converge to a refinement of  $T^*$ .*

With this lemma in mind, it is natural to measure how close a tree  $T$  is to the target  $T^*$  with the following (asymmetric) distance function, which we call *triplet distance* (TD):

$$\text{TD}(T^*, T) = \frac{\sum_{c \in \Delta(T^*)} \mathbb{I}(c \notin \Delta(T))}{|\Delta(T^*)|} \quad (1)$$

where  $\mathbb{I}$  is the indicator function. This distance is zero exactly when  $T$  is a refinement of  $T^*$ , in which case we have reached our goal.

A simple strategy for obtaining triplets would be to present the user with three randomly chosen data points and have the user pick the odd one out. This strategy has several drawbacks. First, some sets of three points have no triplet constraint (for instance, points 1, 2, 3 in Figure 3). Second, the chosen set of points might correspond to a triplet that has already been specified, or is *implied* by specified triplets. For example, knowledge of  $(\{a, b\}, c)$  and  $(\{b, c\}, d)$  implies  $(\{a, c\}, d)$ . Enumerating the set of implied triplets is non-trivial for  $n > 3$  triplets (Bryant & Steel, 1995), making it difficult to avoid these implied triplets in the first place.

We thus consider another strategy—rather than the user arranging three data points into a triplet, the user observes




Figure 2. In the subtree-prune and regraft (SPR) move, a subtree  $S$  is selected uniformly at random and is then pruned from the tree. Next, a regraft location is selected from the valid regraft locations, and  $S$  is re-attached at that location.




Figure 3. Target tree  $T^*$  (left) and a refinement of it.

the hierarchy induced over some  $O(1)$ -sized subset  $S$  of the data and corrects an error in the tree by supplying a triplet. We call this a *subtree query*. Finally, we note that in this work we only consider the *realizable* case where the triplets obtained from a user do not contain contradictory information and that there is a tree that satisfies all of them.

### 3.2. Finding a tree consistent with constraints

We start with a randomly initialized hierarchy  $T$  over our data and show an induced subtree  $T|_S$  to the user, obtaining the first triplet. The next step is constructing a new tree that satisfies the triplet. This begins the feedback cycle; a user provides a triplet given a subtree and the triplet is incorporated into a clustering algorithm, producing a new candidate tree. A starting point is an algorithm that returns a tree consistent with a set of triplets.

The simplest algorithm to solve this problem is the BUILD algorithm, introduced in Aho et al. (1981). Given a set of triplets  $C$ , BUILD will either return a tree that satisfies  $C$ , or error if no such tree exists. In BUILD, we first construct the Aho graph  $G_C$ , which has a vertex for each data point and an undirected edge  $\{a, b\}$  for each triplet constraint  $(\{a, b\}, c)$ . If  $G_C$  is connected, there is no tree that satisfies all triplets. Otherwise, the top split of the tree is a partition of the connected components of  $G_C$ : any split is

fine as long as points in the same component stay together. Satisfied triplets are discarded, and BUILD then continues recursively on the left and right subtrees.

BUILD satisfies triplet constraints but ignores the geometry of the data, whereas we wish to take both into account. By incorporating triplets into the posterior DDT sampler, we obtain high likelihood trees that still satisfy  $C$ .

### 3.3. Incorporating triplets into the sampler

In this section, we present an algorithm to sample candidate trees from the posterior DDT, constrained by a triplet set  $C$ . It is based on the subtree prune and regraft move.

The SPR move is of particular interest because we can efficiently enforce triplets to form a *constrained-SPR move*, resulting in a sampler that only produces trees that satisfy a set of triplets. A *constrained-SPR move* is defined as an SPR move that assigns zero probability to any resulting trees that would violate a set of triplets. Restricting the neighborhood of an SPR move runs the risk of partitioning the state space, losing the convergence guarantees of the Metropolis-Hastings algorithm. Fortunately, a constrained-SPR move does not compromise strong connectivity. For any realizable triplet set  $C$ , we prove the constrained-SPR move Markov chain’s aperiodicity and irreducibility.

Consider the Markov chain on the state space of rooted binary trees that is induced by the constrained sampler.

**Lemma 3.2.** *The constrained-SPR Markov chain is aperiodic.*

*Proof.* A sufficient condition for aperiodicity is the existence of a “self-loop” in the transition matrix: a non-zero probability of a state transitioning to itself. Supposed we have pruned a subtree already. When regrafting, the ordinary SPR move has a non-zero probability of choosing any branch, and a constrained-SPR move cannot regraft to branches that would violate triplets. Since the current tree




Figure 4. Visualized is a constrained-SPR move. Pruning is identical but a regraft location is selected from the valid regraft locations limited by triplets. In this image, we are constrained by the sole triplet  $\{(2, 3), 4\}$ .

in the Markov chain satisfies triplet set  $C$ , there is a non-zero probability of regrafting to the same location. We thus have an aperiodic Markov chain.  $\square$

**Lemma 3.3.** *A constrained-SPR Markov chain is irreducible.*

*Proof.* (sketch) To show irreducibility, we show that a tree  $T$  has a non-zero probability of reaching an arbitrary tree  $T'$  via constrained-SPR moves where both  $T$  and  $T'$  satisfy a set of triplets  $C$ . Our proof strategy is to construct a canonical tree  $T_C$ , and show that there exists a non-zero probability path from  $T$  to  $T_C$ , and therefore from  $T'$  to  $T_C$ . We then show that for a given constrained-SPR move, the reverse move has a non-zero probability. Thus, there exists a path from  $T$  to  $T_C$  to  $T'$ , satisfying irreducibility.

Recall that the split at a node in a binary tree that satisfies triplet set  $C$  corresponds to a binary partition of the Aho graph at the node (see Section 3.2).  $T_C$  is a tree such that every node in  $T_C$  is in *canonical form*. A node is in canonical form if it is a leaf node, or, the partition of the Aho graph at that node can be written as  $(l, r)$ .  $l$  is the single connected component containing the point with the minimum data index, and  $r$  is the rest of the components.

To convert a particular node  $s$  into canonical form, we first perform “grouping”, which puts  $l$  into a single descendant of  $s$  via constrained-SPR moves. We then make two constrained-SPR moves to convert the partition at  $s$  into the form  $(l, r)$  (see Figure 5). We convert all nodes into canonical form recursively, turning an arbitrary tree  $T$  into  $T_C$ .

Finally, the reverse constrained-SPR move has a non-zero probability. Suppose we perform a constrained-SPR move on tree  $T_1$ , converting it into  $T_2$  by detaching subtree  $s$  and attaching it to branch  $(u, v)$ . A constrained-SPR move on  $T_2$  can select  $s$  for pruning and can regraft it to form  $T_1$  with a non-zero probability since  $T_1$  satisfies the same constraint set as  $C$ . For a full proof, please refer to the supplement.  $\square$

The simplest possible scheme for a constrained-SPR move would be rejection sampling. The Metropolis-Hastings algorithm for the DDT would be the same as in the unconstrained case, but any trees violating  $C$  would have accept probability 0. Although this procedure is correct, it is impractical. As the number of triplets grows larger, more trees will be rejected and the sampler will slow down over time.

To efficiently sample a tree that satisfies a set of triplets  $C$ , we modify the regraft in the ordinary SPR move. The constrained-SPR move must assign zero probability to any regraft branches that would result in a tree that violates  $C$ . This is accomplished by generating the path from the root in the same manner as sampling a branch, but avoiding paths that would result in violated triplets.

#### DESCRIPTION OF CONSTRAINED-SPR SAMPLER

Recall that in the DDT’s sampling procedure for regraft branches, a particle at a node picks a branch, and either diverges from that branch or recursively samples the node’s child. Let  $s$  be the root of the subtree we are currently grafting back onto tree  $T$ , let  $C$  be the triplet set, and let  $\text{leaves}(u)$  denote the descendant-leaves of node  $u$ . Suppose we are currently at node  $u$ , deciding whether to diverge at the branch  $(u, v)$  or to recursively sample  $v$ . Consider any triplet  $(\{a, b\}, c) \in C$ . If all—or none—of  $a, b, c$  are in  $\text{leaves}(s)$ , then the triplet is unaffected by the graft, and can be ignored. Otherwise, checks are needed:

1.  $c \in \text{leaves}(s)$

Then we know  $a, b \notin \text{leaves}(s)$ . If  $a$  and  $b$  are split across  $v$ ’s children, we are banned from sampling  $v$ .

2.  $a \in \text{leaves}(s)$  but  $b, c \notin \text{leaves}(s)$

If both  $b$  and  $c$  are in  $\text{leaves}(v)$ , we are required to sample  $v$ . If just  $c$  is in  $\text{leaves}(v)$ , we are banned from both diverging at  $(u, v)$  and sampling  $v$ . Otherwise we can either diverge at  $(u, v)$  or sample  $v$ .




Figure 5. The process of converting  $s$  into canonical form. We first group nodes from  $l$  into their own isolated subtree, then perform two constrained-SPR moves to put  $s$  into canonical form.

(The case where  $b \in \text{leaves}(s)$  is symmetric to case 2.) If we choose to sample  $v$ , we remove constraints from our current set  $C$  that are now satisfied, and continue recursively. This defines a procedure by which we can sample a divergence branch that does not violate constraints.

While the constrained-SPR sampler can produce a set of trees given a set of static constraints, the BUILD algorithm is useful in adding new triplets into the sampler. Suppose we have been sampling trees with constrained-SPR moves satisfying triplet set  $C$  and we obtain a new triplet  $u = (\{a, b\}, c)$  from a user query. We take the current tree  $T$  and find the least common ancestor (call it  $z$ ) of  $a$  and  $b$ . We then call BUILD( $C + \{u\}$ ) on just the nodes in  $\text{leaves}(z)$ , and we substitute the resulting subtree at position  $z$  in tree  $T$ .

### 3.4. Intelligent subset queries

We now have a method to sample a constrained distribution over candidate trees. Given a particular candidate tree  $T$ , our first strategy for subtree querying is to pick a random subset  $S$  of the leaves of constant size, and show the user the induced subtree over the subset,  $T|_S$ . We call this *random subtree querying*. But can we use a set of trees produced by the sampler to make better subtree queries? If tree structure is ambiguous in a particular region of data, i.e. there are several hierarchies that could explain a particular configuration of data, the MH algorithm will sample over these different configurations. A query over points in these ambiguous regions may help our algorithm converge to a better tree faster. By looking for these regions in our samples, we can choose query subsets  $S$  for which the tree structure is highly variable, and hopefully the resulting triplet from the user will reduce the ambiguity.

More precisely, we desire a notion of tree variance. Given a set of trees  $\mathcal{T}$ , what is the variance over a given subset of the data  $S$ ? We propose using the notion of tree distance

as a starting point. For a given tree  $T$ , the tree distance between two nodes  $a$  and  $b$ , denoted  $\text{treedist}_T(a, b)$ , is the number of edges of  $T$  needed to get from  $a$  to  $b$ . Consider two leaves  $u$  and  $v$ . If the tree structure around them is static, we expect the tree distance between  $u$  and  $v$  to change very little, as the surrounding tree will not change. However, if there is ambiguity in the surrounding structure, the tree distance will be more variable. Given a subset of data  $S$  and a set of trees  $\mathcal{T}$ , the tree distance variance (TDV) of the trees over the subset is defined as:

$$\text{TDV}(\mathcal{T}, S) = \max_{u, v \in S} \text{Var}_{T \in \mathcal{T}}[\text{treedist}_{T|_S}(u, v)] \quad (2)$$

This measure of variance is the *max* of the variance of tree distance between any two points in the subset. Computing this requires  $O(|\mathcal{T}||S|^2 \log |S|)$  time, and since  $|S|$  is constant, it is not prohibitively expensive.

Given a set of trees from the sampler  $\mathcal{T}$ , we now select a high-variance subtree by instantiating  $L$  random subsets of constant size,  $S_1, \dots, S_L$  and picking  $\text{argmax}_l \text{TDV}(\mathcal{T}, S_l)$ . We call this *active subtree querying*. Although using tree variance will help reconcile ambiguity in the tree structure, if a set of samples from a tree all violate the same triplet, it is unlikely that active querying will recover that triplet. Thus, interleaving random querying and active querying will hopefully help the algorithm converge quickly, while avoiding local optima.

## 4. Experiments

We evaluated the convergence properties of five different querying schemes. In a “simple query”, a user is presented with three random data and picks an odd one out. In a “smart query”, a user is unrealistically shown the entire candidate tree and reports a violated triplet. In a “random query”, the user is shown the induced candidate tree over a random subset of the data. In an “active query”, the user

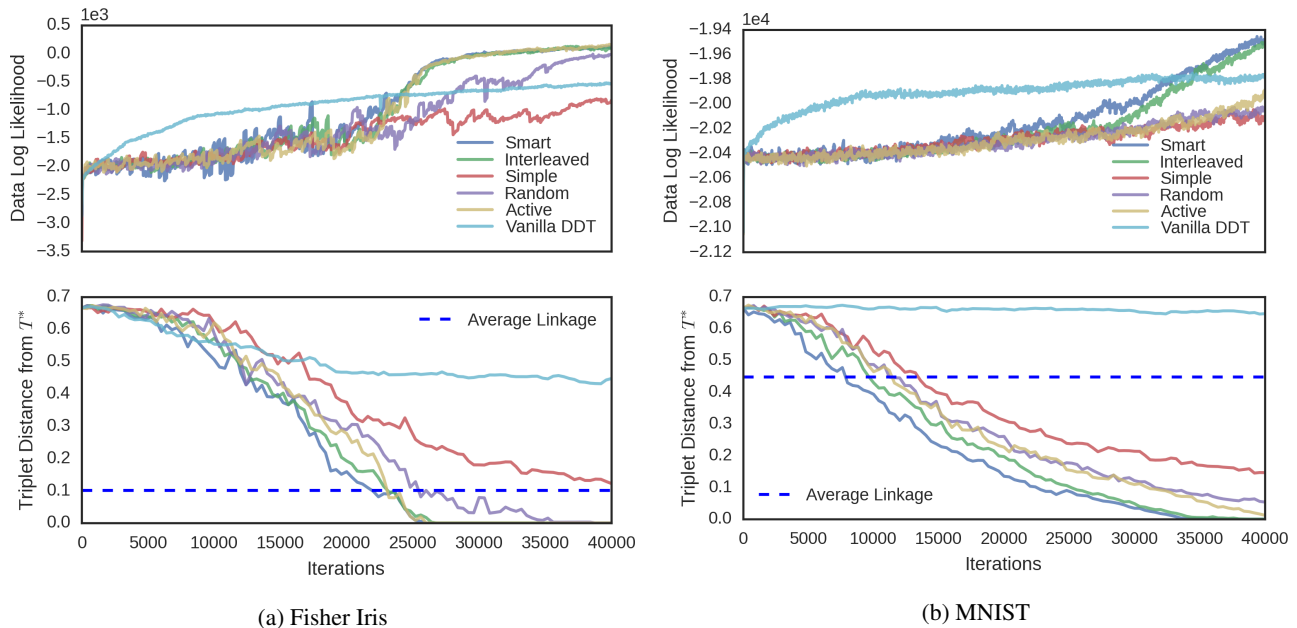


Figure 6. The average of four runs of constrained-SPR samplers for the Fisher Iris dataset and the MNIST dataset, using 5 different querying schemes. A query was made every 100 iterations.

is shown a high variance subtree using tree-distance variance. Finally, in an “interleaved query”, the user is alternatively shown a random subtree and a high variance subtree. In each experiment,  $T^*$  was known, so user queries were simulated by picking a triplet violated by the root split of the queried tree, and if no such triplet existed, recursing on a child. Each scheme was evaluated on four different datasets. The first dataset, MNIST (Lecun et al., 1998), is an 10-way image classification dataset where the data are  $28 \times 28$  images of digits. The target tree  $T^*$  is simply the  $K$ -way classification tree over the data. The second dataset is Fisher Iris, a 3-way flower classification problem, where each of 150 flowers has five features. The third dataset, Zoo (Lichman, 2013), is a set of 93 animals and 15 binary morphological features for each of animals, the target tree being the induced binary tree from the Open Tree of Life (Hinchliff et al., 2015). The fourth dataset is 20 Newsgroups (Joachims, 1997), a corpus of text articles on 20 different subjects. We use the first 10 principal components as features in this classification problem. All datasets were modeled with DDT’s with acquisition function  $a(t) = 1/(1-t)$  and Brownian motion parameter  $\sigma^2$  estimated from data. To better visualize the different convergence rates of the querying schemes, MNIST and 20 Newsgroups were subsampled to 150 random points.

For each dataset and querying scheme, we instantiated a SPR sampler with no constraints. Every one hundred iterations of the sampler, we performed a query. In subtree queries, we used subsets of size  $|S| = 10$  and in ac-

tive querying, the highest-variance subset was chosen from  $L = 20$  different random subsets. As baselines, we measured the triplet distance of the vanilla DDT and the average linkage tree. Finally, results were averaged over four runs of each sampler. The triplet distances for Fisher Iris and MNIST can be seen in Figure 6. Results for the other datasets can be found in the supplement. Although unrealistic due to the size of the tree shown to the user, the smart query performed the best, achieving minimum error with the least amount of queries. Interleaved followed next, followed by active, random, and simple. In general, the vanilla DDT performed the worst, and the average linkage score varied on each dataset, but in all cases, the subtree querying schemes performed better than both the vanilla DDT and average linkage.

In three datasets (MNIST, Fisher Iris and Zoo), interactive methods achieve higher data likelihood than the vanilla DDT. Initially, the sampler is often restructuring the tree with new triplets and data likelihood is unlikely to rise. However, over time as less triplets are reported, the data likelihood increases rapidly. We thus conjecture that triplet constraints may help the MH algorithm find better optima.

## 5. Future Work

We are interested in studying the non-realizable case, i.e. when there does not exist a tree that satisfies triplet set  $C$ . We would also like to better understand the effect of constraints on searching for optima using MCMC methods.



## Acknowledgements

Sharad Vikram and Sanjoy Dasgupta are supported by the National Science Foundation under grant CNS-1446912. The authors are grateful for feedback from the anonymous reviewers and for help and advice from Suqi Liu, Stefanos Poulis, and Christopher Tosh.

## References

- Adams, R.P., Ghahramani, Z., and Jordan, M.I. Tree-structured stick breaking for hierarchical data. In *Advances in Neural Information Processing Systems*, 2008.
- Aho, A.V., Sagiv, Y., Szymanski, T.G., and Ullman, J.D. Inferring a tree from lowest common ancestors with an application to the optimization of relational expressions. *SIAM Journal on Computing*, 10:405–421, 1981.
- Aldous, D. Probability distributions on cladograms. In Aldous, D. and Pemantle, R. (eds.), *Random Discrete Structures (IMA Volumes in Mathematics and its Applications 76)*, pp. 1–18, 1995.
- Awasthi, P. and Zadeh, R.B. Supervised clustering. In *Advances in Neural Information Processing Systems*, 2010.
- Awasthi, P., Balcan, M.-F., and Voevodski, K. Local algorithms for interactive clustering. In *Proceedings of the 31st International Conference on Machine Learning*, 2014.
- Balcan, M.-F. and Blum, A. Clustering with interactive feedback. In *Algorithmic Learning Theory (volume 5254 of the series Lecture Notes in Computer Science)*, pp. 316–328, 2008.
- Bansal, N., Blum, A., and Chawla, S. Correlation clustering. *Machine Learning*, 56(1–3):89–113, 2004.
- Basu, S., Banerjee, A., and Mooney, R. Active semi-supervision for pairwise constrained clustering. In *SIAM International Conference on Data Mining*, 2004.
- Biswas, A. and Jacobs, D. Active image clustering with pairwise constraints from humans. *International Journal of Computer Vision*, 108(1):133–147, 2014.
- Borg, I. and Groenen, P.J.F. *Modern Multidimensional Scaling: Theory and Applications*. Springer Verlag, 2005.
- Boyles, L. and Welling, M. The time-marginalized coalescent prior for hierarchical clustering. In *Advances in Neural Information Processing Systems*, 2012.
- Bryant, D. and Steel, M. Extension operations on sets of leaf-labeled trees. *Advances in Applied Mathematics*, 16(4):425–453, 1995.
- Eriksson, B., Dasarathy, G., Singh, A., and Nowak, R. Active clustering: robust and efficient hierarchical clustering using adaptively selected similarities. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics*, 2011.
- Felsenstein, J. *Inferring Phylogenies*. Sinauer, 2004.
- Hastie, T., Tibshirani, R., and Friedman, J. *The Elements of Statistical Learning*. Springer, 2nd edition, 2009.
- Heller, K. and Ghahramani, Z. Bayesian hierarchical clustering. In *Proceedings of the 22nd International Conference on Machine Learning*, 2005.
- Hinchliff, Cody E., Smith, Stephen A., Allman, James F., Burleigh, J. Gordon, Chaudhary, Ruchi, Coghill, Lyndon M., Crandall, Keith A., Deng, Jiabin, Drew, Bryan T., Gazis, Romina, Gude, Karl, Hibbett, David S., Katz, Laura A., Laughinghouse, H. Dail, McTavish, Emily Jane, Midford, Peter E., Owen, Christopher L., Ree, Richard H., Rees, Jonathan A., Soltis, Douglas E., Williams, Tiffani, and Cranston, Karen A. Synthesis of phylogeny and taxonomy into a comprehensive tree of life. In *Proceedings of the National Academy of Sciences*, 2015.
- Joachims, Thorsten. A probabilistic analysis of the rocchio algorithm with tfidf for text categorization. In *Proceedings of the Fourteenth International Conference on Machine Learning, ICML '97*, pp. 143–151, 1997. ISBN 1-55860-486-3.
- Knowles, D.A. and Ghahramani, Z. Pitman-Yor diffusion trees for Bayesian hierarchical clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(2):271–289, 2015.
- Krishnamurthy, A., Balakrishnan, S., Xu, M., and Singh, A. Efficient active algorithms for hierarchical clustering. In *Proceedings of the 29th International Conference on Machine Learning*, 2012.
- Kulis, B., Basu, S., Dhillon, I., and Mooney, R. Semi-supervised graph clustering: a kernel approach. In *Proceedings of the 22nd International Conference on Machine Learning*, 2005.
- Lecun, Yann, Bottou, Lon, Bengio, Yoshua, and Haffner, Patrick. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*, pp. 2278–2324, 1998.
- Lichman, M. UCI machine learning repository, 2013. URL <http://archive.ics.uci.edu/ml>.
- McCullagh, P., Pitman, J., and Winkel, M. Gibbs fragmentation trees. *Bernoulli*, 14(4):988–1002, 2008.

- Neal, R.M. Density modeling and clustering using Dirichlet diffusion trees. In Bernardo, J.M. et al. (eds.), *Bayesian Statistics 7*, pp. 619–629. Oxford University Press, 2003.
- Tamuz, O., Liu, C., Belongie, S., Shamir, O., and Kalai, A.T. Adaptively learning the crowd kernel. In *Proceedings of the 28th International Conference on Machine Learning*, 2011.
- Teh, Y.W., III, H. Daume, and Roy, D.M. Bayesian agglomerative clustering with coalescents. In *Advances in Neural Information Processing Systems*, 2008.
- Wagstaff, K. and Cardie, C. Clustering with instance-level constraints. In *Proceedings of the 17th International Conference on Machine Learning*, 2000.
- Wagstaff, K., Cardie, C., Rogers, S., and Schroedl, S. Constrained k-means clustering with background knowledge. In *Proceedings of the 18th International Conference on Machine Learning*, 2001.
- Williams, C.K.I. A MCMC approach to hierarchical mixture modeling. In *Advances in Neural Information Processing Systems*, 2000.
- Zoller, T. and Buhmann, J.M. Active learning for hierarchical pairwise data clustering. In *Proceedings of the 15th International Conference on Pattern Recognition*, 2000.