

# Phylogeographic Ancestral Inference Using the Coalescent Model on Haplotype Trees

Ioanna Manolopoulou and Brent C. Emerson

October 7, 2011

## **Abstract**

Phylogeographic ancestral inference is a question frequently arising in population ecology, aiming to understand the geographical roots and structure of species. Here we specifically address relatively small scale mtDNA datasets (typically less than 500 sequences with less than 1000 nucleotides), focusing on ancestral location inference. Our approach uses a coalescent modelling framework projected onto haplotype trees in order to reduce computational complexity, at the same time adhering to complex evolutionary processes. Statistical innovations of the last few years have allowed for computationally feasible yet accurate inferences in phylogenetic frameworks. We implement our methods on a set of synthetic datasets and show how, despite high uncertainty in terms of identifying the root haplotype, estimation of the ancestral location naturally encompasses lower uncertainty, allowing us to pinpoint the Maximum A Posteriori estimates for ancestral locations. We exemplify our methods on a set of synthetic datasets and then combine our inference methods with the phylogeographic clustering approach presented in Manolopoulou et al. (2011) on a real dataset from weevils in the Iberian peninsula in order to infer ancestral locations as well as population substructure.

Phylogeographic ancestral inference is a question frequently arising in population ecology, aiming to understand the geographical roots and structure of species. While many of the cutting-edge approaches (Bloomquist et al., 2010; Lemey et al., 2010, 2009) can provide powerful inferences using highly sophisticated evolutionary models, the amount of data (especially when sampling resources are limited) naturally leads to high levels of uncertainty. At the same time, although parsimonious approaches (Swofford and Berlocher, 1987), typically represented by gene trees, are computationally efficient, it is well-known that they can misrepresent evolution, leading to biased inferences (Felsenstein, 1978). Here we specifically address relatively small scale mtDNA datasets (typically less than 500 sequences with less than 1000 nucleotides), focusing on ancestral location inference.

Our approach uses a coalescent modelling framework projected onto haplotype trees in order to reduce computational complexity, but adheres to complex evolutionary processes. Statistical innovations of the last few years have allowed for computationally feasible yet accurate inferences in phylogenetic frameworks. Although our approach does not employ a rigorous geographical migration model, analysis of synthetic datasets shows that our methods can provide valid and computationally efficient results. Despite high uncertainty in terms of identifying the root node, estimation of the ancestral location naturally encompasses lower uncertainty, allowing us to pinpoint the Maximum A Posteriori estimates for ancestral locations. We combine our inferences with the phylogeographic clustering methods presented in Manolopoulou et al. (2011) on a real dataset from weevils in the Iberian peninsula in order to infer ancestral locations as well as population substructure.

The paper is organized in the following sections: Section 1 presents the core of the evolutionary model used and its projection onto haplotype trees, and Section 2 proceeds to incorporate ancestral location inference. Section 3 presents analysis of a set of synthetic datasets, followed by the analysis of a real dataset in Section 4.

# 1 The haplotype tree model

Coalescent theory (Kingman, 1982), in its basic form, retrospectively relates a set of sequences back to their Most Recent Common Ancestor (MRCA) through a series of coalescence events, based on the assumptions of constant population and random mating. In a sample of  $N$  sequences viewed backwards into the past, the length of time to the latest coalescence event is proportional to  $\binom{N}{2}$ . In order to use haplotype trees as a representation of evolution faithful to the coalescent model (Wakeley, 2008; Kingman, 1982), we describe a theoretical framework whereby haplotype trees are translated into sets of coalescent trees.

Assuming that mutations occur independently as a Poisson process at rate  $\theta/2$ , they can be thought of as being poured down the coalescent tree (Tavaré, 1986, 2003) resulting in a combined Markov process involving both coalescence and mutation events. A rooted haplotype tree only partially determines a set of mutation and coalescence events which occurred in history, and places restrictions on their ordering; see Appendix A. Calculating probabilities over rooted haplotype trees therefore requires summing over all possibilities and orderings of past events. We denote the temporal order of mutation and split events with  $\mathcal{H}$ , with  $\mathcal{H}^t$  being the  $t^{\text{th}}$  event and  $H$  the total number of events. We use a simple evolutionary model with equal exponential mutation process rates across all nucleotide sites and between all possible mutations, but the flexibility of our methods easily extends to more complex evolutionary processes. Conditional on the root  $r$  and the total number of events  $H$  we have

$$\mathbb{P}(\mathcal{H} | H, r) = \prod_t \mathbb{P}(\mathcal{H}^t | r), \quad (1)$$

where the events, forwards in time (Ethier and Griffiths, 1987; Stephens and Donnelly, 2000; Tavaré, 2003), are independent and have probabilities given by the ratios of the underlying

exponential rates  $\binom{N_t}{2}$  and  $\frac{\theta N_t}{2}$  ( $N_t$  being the number of sequences present at time  $t$ ), so that

$$\begin{aligned}\mathbb{P}(\text{a given sequence splits}) &= \frac{2}{N_t} \frac{N_t - 1}{(N_t - 1 + \theta)}, \\ \mathbb{P}(\text{a given sequence mutates at any of its sites}) &= \frac{1}{N_t} \frac{\theta}{(N_t - 1 + \theta)}.\end{aligned}\tag{2}$$

This implies that the probability of a haplotype tree, denoted by  $T$ , given the root  $r$ , can be calculated by summing over the probabilities (1) of all temporal orderings  $\mathcal{H}_j$  which are consistent with the tree. In other words,

$$\mathbb{P}(T | r) = \sum_j \mathbb{P}(\mathcal{H}_j | H_j, r),\tag{3}$$

where  $\mathcal{H}_j$  is consistent with  $T$ . This model implicitly incorporates information about the number of copies of each haplotype within the sample.

We begin by assuming that, in the absence of any information about the mutation process, any haplotype tree  $T$  is equiprobable, and place an Inverse-Gamma prior on the mutation rate  $\theta \sim IG(a_\theta, b_\theta)$ . Similarly to the haplotype tree, we assume that all tree topologies are equally likely a priori given the root, so that  $p(T|r) \propto 1$ , and any sequence (as opposed to haplotype) is equally likely to be the root, so that  $p(r) \propto 1$ . The distribution of the sequence data  $\mathcal{S}$  becomes

$$P(\mathcal{S} | T, r) = \begin{cases} 1 & \text{if } T \text{ consistent with } \mathcal{S} \\ 0 & \text{otherwise} \end{cases}$$

Although only one sequence set  $\mathcal{S}$  is consistent with the fully specified haplotype tree  $T$ , there are several haplotype trees represented by the set  $\Omega$  consistent with the sequences (see Example in the Supplementary Material). The posterior distribution of the haplotype tree then becomes

$$\mathbb{P}(T | \mathcal{S}, r) \propto \begin{cases} \mathbb{P}(T | r) & \text{if } T \in \Omega \\ 0 & \text{otherwise,} \end{cases}$$

where the normalization constant may be calculated as  $\sum_{T_i \in \Omega} \mathbb{P}(T_i | r)$ . Similarly, we can calculate the posterior distribution for the root

$$P(r | \mathcal{S}, T) \propto P(T | r).$$

Calculation and exploration of the infinite state space  $\Omega$  which is consistent with  $\mathcal{S}$  is computationally challenging. We contend that, under an argument of relaxed parsimony, it is possible to reduce the state space to a finite (but vast) set  $\Omega := \Omega(\mathcal{S})$  of realistic haplotype trees (see Manolopoulou et al., 2011).

Finally, the probability of a tree can only be calculated conditional on the total number of mutation and split events, here denoted by  $H$ . However, the set  $\Omega$  may contain trees involving a different number of events. In order to calculate the probability of any one of those trees, we require

$$\mathbb{P}(T | r) \propto \mathbb{P}(T | H, r) \times \mathbb{P}(H | r).$$

Assuming a uniform prior on  $H$  such that

$$\mathbb{P}(H | r) \propto 1,$$

we see that  $\mathbb{P}(T | r)$  can be calculated using (3) and simply multiplying over all the events for any size of tree  $H$ .

This model provides several contributions to inference on the rooted haplotype tree. Based on explicit distributions about haplotype trees, it supplies a rigorous mathematical framework for estimation. It is consistent with many of the theoretical properties of ancestral inference developed by Griffiths and Tavaré (1994) and empirical predictions on haplotype trees raised by Crandall and Templeton (1993) and Posada and Crandall (2001). For example, older alleles have a greater probability of becoming interior (as opposed to leaf) haplotypes: this may be directly derived from (3), since interior haplotypes naturally allow a much larger number of orderings in which events may have occurred. Furthermore,

haplotypes of greater frequency are more likely to have a higher degree (i.e., more mutational connections in the tree): the probability of a mutation increases according to frequency of the haplotype. Perhaps the most important advantage of this model is that the posterior probability of a haplotype tree (1) can be explicitly expressed, which allows for backward rather than forward inference. Finally, although the estimates for the root are inherently unreliable because of the variation in the model, the synthetic data analysis in Section 3 shows that when the sequence data  $\mathcal{S}$  are combined with geographical data for each individual, ancestral locations may be estimated with a much higher probability of success.

In order to draw inferences about the haplotype tree under this model, calculation of  $\mathbb{P}(T | \mathcal{S}, r)$  is required. We employ an auxiliary variable approach following Beaumont (2003), whereby a single temporal ordering of haplotype tree events is considered an auxiliary variable in our inferences, allowing for computational feasibility. An improved approach would allow a set of several temporal orderings (Beaumont, 2003) to be considered as latent in order to reduce Monte Carlo error; previous analyses (Manolopoulou, 2009) showed that, in our case, this was not necessary. Our framework then fully specifies the distribution

$$p(r, \mathcal{H}, T | \mathcal{S}).$$

In order to explore the space of temporal orderings of events, we describe Algorithm 1.1 similar to (Ethier and Griffiths, 1987), which generates temporal orderings consistent with a fixed haplotype tree  $T$ .

**Algorithm 1.1.**

This algorithm generates temporal orderings by mimicking the ancestral history of the sample, starting with the root and ending with the observed sequences.

1. Start at the root. Initially only one copy of the root haplotype is present. Split it into two copies and repeat the next step until all mutation or split events determined by the haplotype tree have occurred.

2. For all sequences present, consider all mutations and splits that are consistent with the haplotype tree. Draw one of those events proportional to the probabilities given by (2), and repeat this step until the tree is complete.

For each temporal ordering  $\mathcal{H}$  generated from this algorithm, the proposal probability  $q(\mathcal{H} | T, r)$  can be calculated by multiplying the normalized terms in (2) over all events. By construction, any  $\mathcal{H}$  consistent with  $T$  may be generated under Algorithm 1.1 because at all the steps, the consistent events have non-zero probability.

In order to simulate samples from the posterior distribution of  $(r, T)$ , we run a Markov chain Monte Carlo sampler; details may be found in the Supplementary Material.

## 2 Ancestral locations in phylogeographic analysis

One of the objectives of phylogeographic ancestral analysis is to identify the location(s) from where a population originated. Although the analysis presented here does not assume a geographical model for the spread of populations in time, it allows for inferring ancestral locations by tracing the locations of the root haplotype, or, if extinct, the locations of its following descendants. This may be better understood through Figure 1.

Figure

In cases where the root haplotype is missing, as in the right-hand panel of Figure 1, we know that its original location will, on average, contain haplotypes from all possible descendant branches. As a result, if a location contains all three pink haplotypes, it is more likely to be ancestral than a location which only contains haplotypes from one of the branches. This approach is consistent with many descriptive characteristics of an ancestral area, for example as presented by Emerson and Hewitt (2005).

Our heuristic approach calculates the contribution of each of the oldest haplotypes along each descendant branch of the root (if the root haplotype is observed in our sample, then we simply have the root only), and then we add the contribution of each of those haplotypes for each location. For example, referring back to the right-hand panel of Figure 1, and assuming that the three pink haplotypes appeared in 1, 2 and 4 locations accordingly, then the contribution of each haplotype to each of the locations it appears in would be  $1/7$ . Although we do not take into account distance from the root, geographical location, or number of times each haplotype appears in each location, implicitly assuming standing variation in the population, our synthetic trials have shown that our approach provides valuable results in inferring ancestral locations.

### 3 Synthetic data analysis

We generate a set of 100 replicate synthetic datasets and assess the performance of our algorithm. Each dataset is initiated by a sequence of length  $l=500$ , at an initial geographical location  $y_{11} = (0, 0)$  and with mutation rate  $\theta \sim IG(1, 10)$ . Each new sequence  $j$  of haplotype  $i$  then is assumed either to stay in its current location, or move to a new location: with probability 0.9 it stays in the geographical location of its ancestor  $a_{ij}$  such that  $y_{ij} = y_{a_{ij}}$ ; otherwise, it moves to a new location  $y_{ij} = N(y_{a_{ij}}, 0.1)$ . The new sequence is forced to start a new location if the location of its ancestor contains 15 or more sequences. These tuning generative parameters were chosen in order for the synthetic datasets to match the real dataset at hand as much as possible. The iterative algorithm stops when it reaches 100 observed sequences (not including ones which are extinct in the process), corresponding to a variable number of haplotypes, locations and geographical clusters. Locations and haplotypes are ordered from oldest to most recent.

We assume both a known and an unknown tree, showing the results below. As expected, the lack of sufficient data results in weak performance at inferring root haplotypes, as indicated by Figure 2 (a). However, the methods are successful in inferring ancestral locations. In the case of a known haplotype tree, the top three ancestral locations cover over 90% of the cases, and the corresponding success rate in the case of an unknown haplotype tree is roughly 85%. Note that assessment of the ancestral haplotype inference in the case of an unknown tree is not possible, as haplotypes cannot always be uniquely identified.

Figure 2

## 4 Real dataset implementation

We apply our algorithms to a mitochondrial DNA dataset of weevils in the Iberian peninsula. *Rhinusa vestita* is a seed parasite weevil feeding and reproducing on snapdragons. It is believed to have been present in Portugal, Spain, France and Italy. The complete nucleotide sequence for the mitochondrial COII gene (722 bp) was obtained for 275 *Rhinusa vestita* individuals. Previous studies investigating the association of weevils with three host plant species, combined with knowledge about the glaciation history of the Iberian peninsula (Hewitt, 2000), led to the biological prediction that the species originated from the Rhône valley to the east and west.

We combine our methods with the analysis presented in Manolopoulou et al. (2011), in order to infer both ancestral locations but also population substructure. The results confirm the biological hypothesis of the location of origin; the top four locations, collecting 75% of posterior mass, are shown in Table 1.

Table 1
---------

Figure
--------

Figure
--------

## 5 Discussion

We have presented a statistical framework whereby the coalescent model is used in order to draw inferences about haplotype trees through Markov chain Monte Carlo. In addition, we have described methods for inferring ancestral locations in phylogeographic settings. Our results were validated by simulated synthetic datasets, and were successful in confirming the biological hypothesis in the real dataset.

Although more sophisticated evolutionary models may be used to account for a variable population size (see Slatkin, 2001), selection (see Neuhauser and Krone, 1997) and recombination (see Hudson and Kaplan, 1988), prior implementations we ran showed that in small-scale datasets such as the one at hand, the data are very weakly informative about many of the additional evolutionary parameters. Perhaps the most valuable extension would allow for the coalescence rate to vary across population clusters, in order to represent local proliferations. Similarly, rigorous theoretical calculations relating ancestral haplotypes with geographical locations (Bloomquist et al., 2010; Lemey et al., 2010, 2009), perhaps through the use of an explicit migration model, can provide a solid basis for an improved estimator of ancestral locations.

Finally, our methods are freely available through an R package “Bayesian Phylogeographic Clustering”, available at <http://www.stat.duke.edu/~im30/software.html>.

## References

- M. Beaumont. Estimation of population growth or decline in genetically monitored populations. **Genetics**, 164:1139–1160, 2003.
- E.W. Bloomquist, P. Lemey, and M.A. Suchard. Three roads diverged? Routes to phylogeographic inference. **Trends in Ecology & Evolution**, 2010.
- K. Crandall and A. Templeton. Empirical tests of some predictions from coalescent theory with applications to intraspecific phylogeny reconstruction. **Genetics**, 134:959–969, 1993.
- B. Emerson and G. Hewitt. Phylogeography. **Current Biology**, 15:R367–371, 2005.
- S. N. Ethier and R. C. Griffiths. The infinitely-many-sites model as a measure-valued diffusion. **The Annals of Probability**, 15:515–545, 1987.
- J. Felsenstein. Cases in which parsimony or compatibility methods will be positively misleading. **Systematic Zoology**, 27:401–410, 1978.
- R. Griffiths and S. Tavaré. Ancestral inference in population genetics. **Statistical Science**, 9:307–319, 1994.
- G. Hewitt. The genetic legacy of the Quaternary ice ages. **Nature**, 405:907–913, 2000.
- R. Hudson and N. Kaplan. The coalescent process in models with selection and recombination. **Genetics**, 120, 1988.
- J. Kingman. The coalescent. **Stochastic Processes and their Application**, 1982.
- P. Lemey, A. Rambaut, A.J. Drummond, and M.A. Suchard. Bayesian phylogeography finds its roots. **PLoS Computational Biology**, 5(9):e1000520, 2009.

- P. Lemey, A. Rambaut, J.J. Welch, and M.A. Suchard. Phylogeography takes a relaxed random walk in continuous space and time. **Molecular Biology and Evolution**, 27(8): 1877, 2010.
- I. Manolopoulou. **A Bayesian Approach to Nested Clade Analysis**. PhD thesis, University of Cambridge, 2009.
- I. Manolopoulou, L. Legarreta, B. C. Emerson, S. P. Brooks, and S. Tavaré. A Bayesian approach to phylogeographic clustering. **Journal of the Royal Society Interface Focus**, 2011.
- C. Neuhauser and S. Krone. The genealogy of samples in models with selection. **Genetics**, 145, 1997.
- D. Posada and K. Crandall. Intraspecific gene genealogies: trees grafting into networks. **Trends in Ecology and Evolution**, 16:37–45, 2001.
- M. Slatkin. Simulating genealogies of selected alleles in a population of variable size. **Genetic Research**, 78:49–57, 2001.
- M. Stephens and P. Donnelly. Inference in molecular population genetics. **Journal of the Royal Statistical Society. Series B (Methodological)**, 62:605–655, 2000.
- D.L. Swofford and S.H. Berlocher. Inferring evolutionary trees from gene frequency data under the principle of maximum parsimony. **Systematic Biology**, 36(3):293, 1987.
- S. Tavaré. Some probabilistic and statistical problems in the analysis of DNA sequences. **Lectures on Mathematics in the Life Sciences**, 17:57–86, 1986.
- S. Tavaré. **Nature Encyclopedia of the Human Genome**. Nature Publishing Group, 2003.

J. Wakeley. **Coalescent Theory: An Introduction**. Roberts and Company Publishers, 2008.

## A Haplotype Tree Example

Suppose the haplotype tree is given by the top tree of Figure 5. For ease of exposition, the numbers on the nodes here represent the sample sizes of each haplotype rather than the label of each haplotype, and we represent each event by updating the numbers on each haplotype according to the number of times it is observed at each time-point in the sample.

Figure 5

Simulating a temporal ordering implies that, starting with the ancestral sequence, we specify a series of split and mutation events which occurred by mimicking evolution, eventually resulting in the fixed haplotype tree. For example, the bottom panel of Figure 5 is a possible temporal ordering of the observed tree given in the top panel.

Observe now that, for example, the root node could not have split any further: this would result in three copies of the ancestral haplotype, which is inconsistent with the haplotype tree which specifies precisely two. In addition, it would not have been possible for the intermediate haplotype to mutate after Step 3 above, since then it would disappear from the ancestral sequences, and another mutation would not have been possible. In other words, consistent events are defined as follows.

- A split event is consistent with the haplotype tree, if it does not imply that the sample size of that haplotype will exceed the number of times it appears in the complete haplotype tree, plus the number of mutations that haplotype will be forced to undergo in following steps (so, in the example, the intermediate haplotype after Step 5 will be forced to undergo exactly one more mutation).
- Similarly, a mutation is possible if (a) is true, and (b) OR (c) are true:
  - (a) it is represented by an edge on the haplotype tree, where the ancestral sequence of the edge has already appeared in the ancestral sample;

- (b) the ancestral sequence of the edge corresponding to that mutation does not go extinct;
- (c) the ancestral sequence of the edge goes extinct, and there are not more events involving that sequence which have not yet occurred but are forced by the haplotype tree.

## B Markov chain Monte Carlo sampler

The complete model contains the tree topology  $T$ , the root  $r$  and mutation rate  $\theta$ , and also includes the temporal ordering  $\mathcal{H}$  as a latent variable. In order to draw samples from the posterior distribution of the parameters of interest  $p(r, T | \mathcal{S})$ , we construct a Markov chain Monte Carlo sampler. The chain is initialized by drawing a mutation rate  $\theta^{(0)}$ , generating a tree  $T^{(0)}$ , and picking root  $r^{(0)}$  uniformly from  $T^{(0)}$ .

1. Propose a new root by using the prior distribution as a proposal kernel over all available sequences  $q(r \rightarrow r') = p(r')$ , and sample a latent temporal ordering  $\mathcal{H}' = \{\mathcal{H}'_1, \dots, \mathcal{H}'_J\}$  according to Algorithm 1.1 with probability  $q(\mathcal{H}' | T, r')$ . Accept or reject  $(r', \mathcal{H}')$  according to the corresponding Metropolis-Hastings ratio  $\min(1, A_r)$ , where

$$A_r = \frac{P(\mathcal{H}' | \mathcal{S}, H, r', \theta)}{P(\mathcal{H} | \mathcal{S}, H, r, \theta)} \times \frac{q(\mathcal{H} | T, H, r, \theta)}{q(\mathcal{H}' | T, H, r', \theta)}$$

2. Propose a new tree topology  $T'$  at random (implying a number of events  $H'$ ), and sample  $\mathcal{H}' = \{\mathcal{H}'_1, \dots, \mathcal{H}'_J\}$  according to Algorithm 1.1 with probability  $q(\mathcal{H}' | T', r)$ . Accept or reject the new tree topology and latent ordering according to the corresponding Metropolis-Hastings ratio  $\min(1, A_T)$ , where

$$A_T = \frac{P(\mathcal{H}' | \mathcal{S}, H', r, \theta)}{P(\mathcal{H} | \mathcal{S}, H, r, \theta)} \times \frac{q(\mathcal{H} | T, H, r, \theta)}{q(\mathcal{H}' | T', H', r, \theta)}$$

3. Propose new mutation rate from the prior  $\theta \sim IG(a_\theta, b_\theta)$ , and accept according to the corresponding Metropolis-Hastings ratio  $\min(1, A_\theta)$ , where

$$A_\theta = \frac{P(\mathcal{H} | \mathcal{S}, H, r, \theta')}{P(\mathcal{H} | \mathcal{S}, H, r, \theta)}$$



## List of Figures

1	Two possible genealogy scenarios, where coloured points represent observed haplotypes (with the colour representing the location) and small black circles are unsampled. In the figure on the left, the oldest haplotype is the pink one at the top. In the figure on the right, the oldest haplotype is missing, and the next possible descendants are the three pink haplotypes. . . . .	21
2	Panel (a) shows a histogram of posterior MAP estimates of the top 20 ancestral haplotypes in the case of a known tree. Panels (b) and (c) show the posterior MAP estimates of ancestral locations for a known and unknown haplotype tree, respectively. . . . .	23
3	One of the non-unique MAP estimates of the haplotype tree using our approach, where colour corresponds to cluster and size to the number of individuals sampled with each sequence. . . . .	25
4	Corresponding bivariate normal contour plots evaluated at the posterior means for the weevil dataset. The black dots indicate sampling locations, and colours correspond to the clusters shown in Figure 3. The larger dot corresponds to the MAP ancestral location. . . . .	27

5 Top panel: In this tree the MRCA of the sample (the top haplotype) is observed twice in the sample. Note that one of the intermediate haplotypes is not observed in the sample (and hence has zero sample size). Bottom panel: a possible scenario for how the present sample came about. Nodes without a number represent haplotypes that have not arisen yet. At first one sequence is present, the ancestral sequence, which split into two (remember that the first event is always a split). Then one of those two identical sequences split again to give us a total of three. One of those three then mutates to give rise to the intermediate haplotype, which in turn splits and then mutates (and goes extinct) to give us the right-hand leaf. Finally, the intermediate haplotype mutates again to give us the left-hand leaf, which subsequently splits to give another copy of itself. . . . .





Ioanna Manolopoulou and Brent C. Emerson

Figure 1 (of 5)

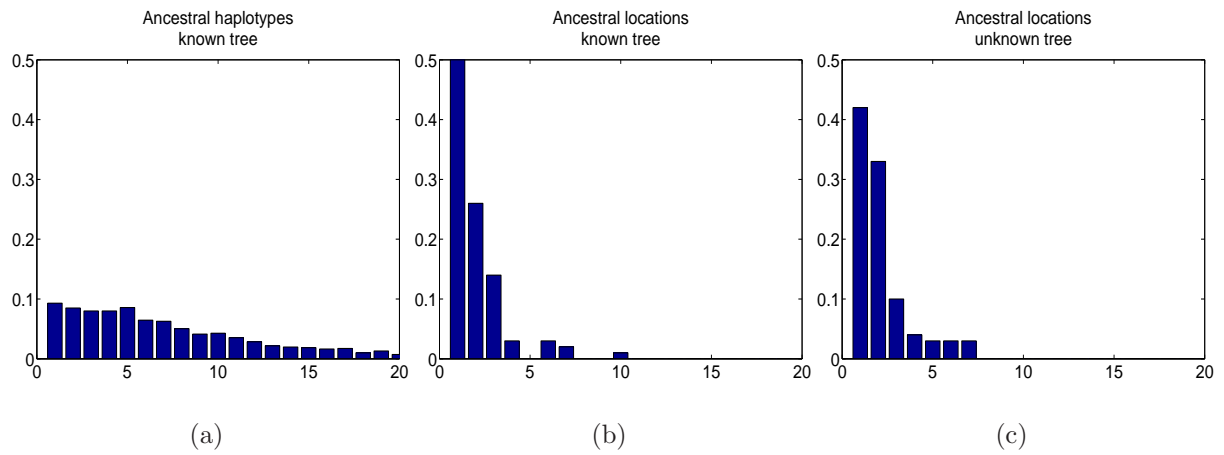


Figure 2: Panel (a) shows a histogram of posterior MAP estimates of the top 20 ancestral haplotypes in the case of a known tree. Panels (b) and (c) show the posterior MAP estimates of ancestral locations for a known and unknown haplotype tree, respectively.



Ioanna Manolopoulou and Brent C. Emerson

Figure 2 (of 5)




Figure 3: One of the non-unique MAP estimates of the haplotype tree using our approach, where colour corresponds to cluster and size to the number of individuals sampled with each sequence.



Ioanna Manolopoulou and Brent C. Emerson

Figure 3 (of 5)




Figure 4: Corresponding bivariate normal contour plots evaluated at the posterior means for the weevil dataset. The black dots indicate sampling locations, and colours correspond to the clusters shown in Figure 3. The larger dot corresponds to the MAP ancestral location.



Ioanna Manolopoulou and Brent C. Emerson

Figure 4 (of 5)




Figure 5: Top panel: In this tree the MRCA of the sample (the top haplotype) is observed twice in the sample. Note that one of the intermediate haplotypes is not observed in the sample (and hence has zero sample size). Bottom panel: a possible scenario for how the present sample came about. Nodes without a number represent haplotypes that have not arisen yet. At first one sequence is present, the ancestral sequence, which split into two (remember that the first event is always a split). Then one of those two identical sequences split again to give us a total of three. One of those three then mutates to give rise to the intermediate haplotype, which in turn splits and then mutates (and goes extinct) to give us the right-hand leaf. Finally, the intermediate haplotype mutates again to give us the left-hand leaf, which subsequently splits to give another copy of itself.



Ioanna Manolopoulou and Brent C. Emerson

Figure 5 (of 5)

## List of Tables

1	Posterior ancestral probabilities of the top four sampling locations of the <i>R. vestita</i> data. . . . .	33
---	---	----



location	posterior mass
Brissac	0.26
Petit Luberon	0.21
La Clape	0.14
Grotte Petit	0.14

Table 1: Posterior ancestral probabilities of the top four sampling locations of the *R. vestita* data.



Ioanna Manolopoulou and Brent C. Emerson

Table 1 (of 1)