

Discrete trait analysis

Kenneth B. Hoehn

2022-04-13

Contents

Discrete trait statistics	1
Caveats and interpreting results	1
Setting up IgPhyML	2
Set up data structures and trees	2
Discrete trait analysis with fixed trees	3
Accounting for uncertainty in tree topology	4
Within and between lineage permutations	4
Controlling false positive rate through downsampling	5

These features are currently only in the development version of Dowser, not the version on CRAN. Please see install page for details on the latest version.

Dowser implements recently developed phylogenetic tests to characterize B cell migration, differentiation, and isotype switching. These are “discrete” traits because their values are not continuous.

Discrete trait statistics

There are two main statistics implemented by Dowser to characterize the distribution of trait values at the tips of trees. These all derive from using a maximum parsimony algorithm to reconstruct trait values at the internal nodes of each tree. Then “switches” between discrete characters are quantified along the tree topologies.

1. PS: Parsimony score, the total number of switches among all trait values.
2. SP: Switch proportion, the proportion of switches from one trait to another.

The significance of each of these statistics is estimated using a permutation test, which randomizes trait values at the tree’s tips.

Full published details on these methods is available [here](#)

Caveats and interpreting results

Before using these tests, it’s important to understand how the test works and what potential biases can affect the results. These tests are detailed in full [here](#). These tests summarize the distribution of trait values along trees. They don’t directly test for cellular migration and differentiation. As such, it’s important to understand potential caveats and explore alternative explanations.

What does a significant SP test mean? A significant SP test from tissue A to tissue B indicates that there is a greater proportion of switches from A to B along the trees than expected by random

tip/trait relationships. This can be due to B cells preferentially migrating from tissue A to B, but there are other possibilities.

As explored in the original paper, a significant SP test from tissue A to B could be due to:

1. Biased migration from A to B.
2. Biased sampling, especially under-sampling tissue A.
3. Significantly higher mutation rate over time in the cells in B than in A, placing it further down the tree.
4. Extremely low rate of switching events (if no downsampling used, see below).

Usually option #1 is most biologically interesting, but it is important to consider the other options as well.

Setting up IgPhyML

All of the switch count statistics on this page require IgPhyML to be installed and passed to the `findSwitches` function. To set up IgPhyML, please visit the docs page. This is true even if IgPhyML is not used to build the trees. Note, if you're not using a Linux computer, your best option may be to use the Immcantation Docker container.

Set up data structures and trees

This step proceeds as in tree building, but it is important to specify the column of the discrete trait you want to analyze in the `formatClones` step. In this example we are using simulated data from nose and lung biopsies. However, this could be any discrete trait value such as cell types. Filtering out clones that contain only a single trait value type is not strictly necessary but can dramatically improve computing time.

```
library(dowser)

# load example AIRR tsv data
data(ExampleAirr)

trait="biopsy"

# Process example data using default settings
clones = formatClones(ExampleAirr,
  traits=trait,num_fields="duplicate_count", minseq=3)

# Column shows which biopsy the B cell was obtained from
print(table(ExampleAirr[[trait]]))
#Lung Nose
# 145  244

# Calculate number of tissues sampled in tree
tissue_types = unlist(lapply(clones$data, function(x)
  length(unique(x@data[[trait]]))))

# Filter to multi-type trees
```

```
clones = clones[tissue_types > 1,]

# Build trees using maximum likelihood (can use alternative builds if desired)
trees = getTrees(clones, build="pml")
```

Discrete trait analysis with fixed trees

Once we've set up the tree objects, we can calculate the switches along these trees using a maximum parsimony algorithm implemented in IgPhyML. We only need to perform this computationally intensive step once. All discrete trait tests use the resulting object. Note the IgPhyML location must be properly configured for your setup. Note also that your results may differ slightly from the ones shown below, due to the stochasticity of this test.

```
# the location of the igphyml executable
# this is location in Docker image, will likely be different if you've set it up yourself
igphyml_location = "/usr/local/share/igphyml/src/igphyml"
```

```
# calculate switches along trees compared to 100 random permutations
# this may take a while, and can be parallelized using nproc
switches = findSwitches(trees, permutations=100, trait=trait,
  igphyml=igphyml_location, fixtrees=TRUE)
```

```
# Perform PS test on switches
ps = testPS(switches$switches)
print(ps$means)
# A tibble: 1 × 6
#   RECON PERMUTE   PLT DELTA STAT  REPS
#   <dbl>   <dbl> <dbl> <dbl> <chr> <int>
#1     8     8.6  0.4  -0.6 PS     100
```

```
sp = testSP(switches$switches, alternative="greater")
print(sp$means)
# A tibble: 2 × 8
# Groups:   FROM [2]
#   FROM TO   RECON PERMUTE  PGT DELTA STAT  REPS
#   <chr> <chr> <dbl>   <dbl> <dbl> <dbl> <chr> <int>
#1 Lung  Nose  0.131   0.323   1 -0.192 SP     100
#2 Nose  Lung  0.869   0.677   0  0.192 SP     100
```

PS test The object `ps` contains the results of the PS test. The `means` dataframe contains the test summary statistics, while the `reps` dataframe contains information from each repetition. `RECON` shows the number of switches across all trees, `PERMUTE` shows the mean number of switches across all permutations. `PLT` is the P value that there are at least as many switches in the observed trees as in the permuted data. If `PLT < 0.05`, this indicates that there are significantly fewer switches among biopsies than expected by random tree/trait association. This indicates that sequences within the sample biopsies are more clustered together within the trees than expected. In this case, `PLT` is 0.4, so we can't draw that conclusion.

SP test The object `sp` contains the results of the SP test. As with the PS test, the `means` dataframe contains summary statistics. The columns are largely the same as in the PS test, but the SP test is performed in a particular direction. In this case, we have an SP test from Lung to Nose and then from Nose to Lung. We can see that the SP value from Nose to Lung in our trees (`RECON`) is much higher than the mean SP statistic in our permuted trees (`PERMUTE`). The p value in the `PGT` column (0) shows that the SP statistic from Nose to Lung is significant. This indicates that there are a greater proportion of switches from the Nose to the Lungs within these trees than expected from random tree/trait association. This is consistent with biased movement from Nose to Lungs in this dataset (but see Caveats and interpreting results section).

Accounting for uncertainty in tree topology

The previous analysis used fixed tree topologies, which will speed up calculations but does not account for uncertainty in tree topology. To account for this, be sure `fixtrees=FALSE` (the default option). For this option, the trees will be re-built for each permutation in the manner specified (same parameters as `getTrees`).

```
# calculate switches along bootstrap distribution of trees
# build using the 'pml' maximum likelihood option
# in a real analysis it's important to use at least 100 permutations
switches = findSwitches(trees, permutations=10, trait=trait,
  igphyml=igphyml_location, fixtrees=FALSE, build="pml")

sp = testSP(switches$switches, alternative="greater")
print(sp$means)
# A tibble: 2 × 8
# Groups:   FROM [2]
#   FROM TO     RECON PERMUTE   PGT DELTA STAT  REPS
#   <chr> <chr> <dbl>   <dbl> <dbl> <dbl> <chr> <int>
#1 Lung  Nose  0.168   0.358   1  -0.190 SP      10
#2 Nose  Lung  0.832   0.642  0.1  0.190 SP      10
```

Within and between lineage permutations

In some cases it may be preferable to permute trait values among trees rather than within them. This will detect association between traits within a tree as well as directional relationships. In general it is harder to interpret. To perform this test, set `permuteAll=TRUE`.

```
sp = testSP(switches$switches, alternative="greater", permuteAll=TRUE)
print(sp$means)
# A tibble: 2 × 8
# Groups:   FROM [2]
#   FROM TO     RECON PERMUTE   PGT DELTA STAT  REPS
#   <chr> <chr> <dbl>   <dbl> <dbl> <dbl> <chr> <int>
#1 Lung  Nose  0.168   0.241   0.6 -0.0736 SP      10
#2 Nose  Lung  0.832   0.759   0.4  0.0736 SP      10
```

Controlling false positive rate through downsampling

The SP test has been shown to have a high false positive rate if switching events are rare along very large trees (see paper). To reduce this effect, by default a downsampling algorithm in `findSwitches` will downsample all trees to have a maximum tip-to-switch ratio of 20. This ratio can be toggled by altering the `tip_switch` parameter. This feature can also be turned off by setting `downsample=FALSE`, but this is not recommended.

```
# Downsample each tree to a tip-to-switch ratio of 10 instead of 20
# this will reduce the false positive rate but also (likely) power
switches = findSwitches(trees, permutations=100, trait=trait,
  igphyml=igphyml_location, fixtrees=TRUE, tip_switch=10)

# didn't have much effect for this dataset
sp = testSP(switches$switches, alternative="greater")
print(sp$means)
# A tibble: 2 × 8
# Groups:   FROM [2]
# FROM TO RECON PERMUTE PGT DELTA STAT REPS
# <chr> <chr> <dbl> <dbl> <dbl> <dbl> <chr> <int>
#1 Lung Nose 0.168 0.358 1 -0.190 SP 10
#2 Nose Lung 0.832 0.642 0.1 0.190 SP 10
```