

# Shazam: Quantification of selection pressure

*Namita Gupta & Jason Anthony Vander Heiden*

2016-08-05

## Contents

Example data . . . . .	1
Calculate selection PDFs for individual sequences . . . . .	1
Group and convolve individual sequence PDFs . . . . .	4
Plot and compare selection scores for groups . . . . .	5

BASELINE quantifies selection pressure by calculating the posterior probability density function (PDF) based on observed mutations compared to expected mutation rates derived from an underlying SHM targeting model. Selection is quantified via the following steps:

1. Calculate the selection scores for individual sequences.
2. Group by relevant fields for comparison and convolve individual selection PDFs.
3. Plot and compare selection scores of different groups of sequences.

## Example data

A small example Change-O database is included in the `alakazam` package. The example dataset consists of a subset of Ig sequencing data from an influenza vaccination study (Laserson and Vigneault et al., PNAS, 2014). The data include sequences from multiple time-points before and after the subject received an influenza vaccination. Quantifying selection requires the following fields (columns) to be present in the Change-O database:

- SEQUENCE\_ID
- SEQUENCE\_IMGT
- GERMLINE\_IMGT\_D\_MASK

```
# Load example data  
library(shazam)  
data(ExampleDb, package="alakazam")
```

## Calculate selection PDFs for individual sequences

Selection scores are calculated with the `calcBaseline` function. This can be performed with a single call to `calcBaseline`, which performs all required steps. Alternatively, one can perform each step separately for greater control over the analysis parameters.

## Calculating selection in multiple steps

Individual sequences within clonal groups are not, strictly speaking, independent events and it is generally appropriate to only analyze selection pressures on an effective sequence for each clonal group. The `collapseClones` function provides one strategy for generating an effective sequences for each clone. It reduces the input database to one row per clone and appends a `CLONAL_SEQUENCE` column which contains the consensus sequence for each clone.

```
# Collapse clonal groups into single sequences
clones <- collapseClones(ExampleDb, regionDefinition=IMGT_V_NO_CDR3, nproc=1)
```

Following construction of an effective sequence for each clone, the observed and expected mutation counts are calculated for each sequence of the `CLONAL_SEQUENCE` column adding in the previous step. `observedMutations` is used to calculate the number of observed mutations and `expectedMutations` calculates the expected frequency of mutations. The underlying targeting model for calculating expectations can be specified using the `targetingModel` parameter. In the example below, the default `HS5FModel` is used. Column names for sequence and germline sequence may also be passed in as parameters if they differ from the Change-O defaults.

Mutations are counted by these functions separately for complementarity determining (CDR) and framework (FWR) regions. The `regionDefinition` argument defines whether these regions are handled separately, and where the boundaries lie. There are four built-in region definitions in the `shazam` package, which are all dependent upon the V segment being IMGT-gapped:

- `IMGT_V`: All regions in the V segment grouped as either CDR or FWR.
- `IMGT_V_BY_REGIONS`: The CDR1, CDR2, CDR3, FWR1, FWR and FWR3 regions in the V segment treated as individual regions.
- `IMGT_V_NO_CDR3`: All regions in the V segment, excluding CDR3, grouped as either CDR or FWR (default).
- `IMGT_V_BY_REGIONS_NO_CDR3`: The CDR1, CDR2, FWR1, FWR and FWR3 regions in the V segment treated as individual regions.

Users may define other region sets and boundaries by creating a custom `RegionDefinition` object.

```
# Count observed mutations and append OBSERVED columns to the output
observed <- observedMutations(clones,
                             sequenceColumn="CLONAL_SEQUENCE",
                             regionDefinition=IMGT_V_NO_CDR3, nproc=1)
# Count observed mutations and append EXPECTED columns to the output
expected <- expectedMutations(observed,
                              sequenceColumn="CLONAL_SEQUENCE",
                              targetingModel=HS5FModel,
                              regionDefinition=IMGT_V_NO_CDR3, nproc=1)
```

The counts of observed and expected mutations can be combined to test for selection using `calcBaseline`. The statistical framework used to test for selection based on mutation counts can be specified using the `testStatistic` parameter.

```
# Calculate selection scores using the output from expectedMutations
baseline <- calcBaseline(expected, testStatistic="focused",
                        regionDefinition=IMGT_V_NO_CDR3, nproc=1)
```

## Calculating selection in one step

It is not required for `collapseClones`, `observedMutation` and `expectedMutations` to be run prior to `calcBaseline`. If the output of the previous three steps does not appear in the input `data.frame`, then `calcBaseline` will call the appropriate functions prior to calculating selection scores.

```
# Calculate selection scores from scratch on subset
baseline <- calcBaseline(ExampleDb, testStatistic="focused",
                        regionDefinition=IMGT_V_NO_CDR3, nproc=1)

# Subset the original data to switched isotypes
db_sub <- subset(ExampleDb, ISOTYPE %in% c("IgA", "IgG"))
# Calculate selection scores from scratch on subset
baseline_sub <- calcBaseline(db_sub, testStatistic="focused",
                            regionDefinition=IMGT_V_NO_CDR3, nproc=1)
```

## Using alternative mutation definitions and models

The default behavior of `observedMutations` and `expectedMutations`, and by extension `calcBaseline`, is to define a replacement mutation in the usual way - any change in the amino acid of a codon is considered a replacement mutation. However, these functions have a `mutationDefinition` argument which allows these definitions to be changed by providing a `MutationDefinition` object that contains alternative replacement and silent criteria. `shazam` provides the following built-in `MutationDefinitions`:

- `CHARGE_MUTATIONS`: Amino acid mutations are defined by changes in side chain charge class.
- `HYDROPATHY_MUTATIONS`: Amino acid mutations are defined by changes in side chain hydrophobicity class.
- `POLARITY_MUTATIONS`: Amino acid mutations are defined by changes in side chain polarity class.
- `VOLUME_MUTATIONS`: Amino acid mutations are defined by changes in side chain volume class.

The default behavior of `expectedMutations` is to use the human 5-mer mutation model, `HS5FModel`. Alternative SHM targeting models can be provided using the `targetingModel` argument.

```
# Calculate selection on charge class with the mouse 5-mer model
baseline <- calcBaseline(ExampleDb, testStatistic="focused",
                        regionDefinition=IMGT_V_NO_CDR3,
                        targetingModel=MRS5NFModel,
                        targetingModel=CHARGE_MUTATIONS,
                        nproc=1)
```

## Group and convolve individual sequence PDFs

To compare the selection scores of groups of sequences, the sequences must be convolved into a single PDF representing each group. In the example dataset, the `SAMPLE` field corresponds to samples taken at different time points before and after an influenza vaccination and the `ISOTYPE` field specifies the isotype of the sequence. The `groupBaseline` function convolves the `BASELINE` PDFs of individual sequences/clones to get a combined PDF. The field(s) by which to group the sequences are specified with the `groupBy` parameter. The `groupBaseline` function automatically calls `summarizeBaseline` to generate summary statistics based on the requested groupings, and populates the `stats` slot of the input `Baseline` object with the number of sequences with observed mutations for each region, selection scores, 95% confidence intervals, and p-values that each group differs from zero.

```
# Combine selection scores by time-point
grouped_1 <- groupBaseline(baseline, groupBy=c("SAMPLE"))

# Combine selection scores by time-point and isotype
grouped_2 <- groupBaseline(baseline_sub, groupBy=c("SAMPLE", "ISOTYPE"))
```

## Convolving variables at multiple level

To make selection comparisons using two levels of variables, you would need two iterations of groupings, where the first iteration of `groupBaseline` groups on both variables, and the second iteration groups on the “outer” variable. For example, if a data set has both case and control subjects, annotated in `STATUS` and `SUBJECT` columns, then generating convolved PDFs for each status would be performed as:

```
# First group by subject and status
subject_grouped <- groupBaseline(baseline, groupBy=c("STATUS", "SUBJECT"))

# Then group the output by status
status_grouped <- groupBaseline(subject_grouped, groupBy="STATUS")
```

## Testing the difference in selection PDFs between groups

The `testBaseline` function will perform significance testing between two grouped `BASELINE` PDFs, by region, and return a data.frame with the following information:

- `REGION`: The sequence region, such as “CDR” and “FWR”.
- `TEST`: The name of the two groups compared.
- `PVALUE`: Two-sided p-value for the comparison.
- `FDR`: FDR corrected p-value.

```
testBaseline(grouped_1, groupBy="SAMPLE")

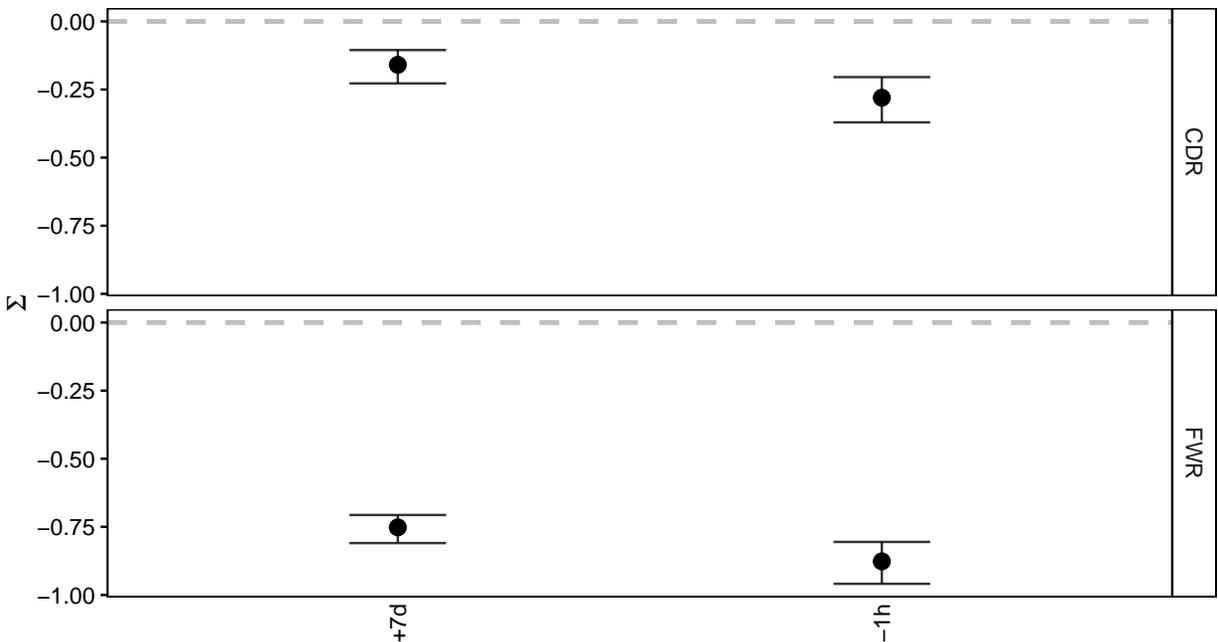
##   REGION      TEST      PVALUE      FDR
## 1   CDR -1h != +7d 0.011937345 0.01193735
## 2   FWR -1h != +7d 0.005143058 0.01028612
```

## Plot and compare selection scores for groups

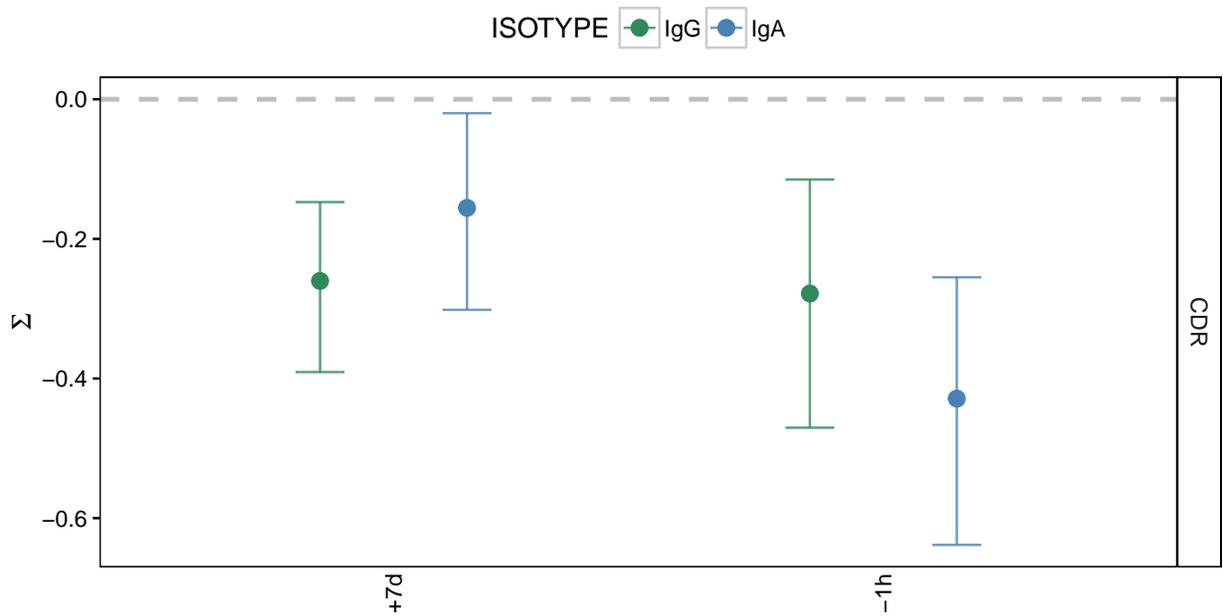
`plotBaselineSummary` plots the mean and confidence interval of selection scores for the given groups. The `idColumn` argument specifies the field that contains identifiers of the groups of sequences. If there is a secondary field by which the sequences are grouped, this can be specified using the `groupColumn`. This secondary grouping can have a user-defined color palette passed into `groupColors` or can be separated into facets by setting the `facetBy="group"`. The `subsetRegions` argument can be used to visualize selection of specific regions. Several examples utilizing these different parameters are provided below.

```
# Set sample and isotype colors
sample_colors <- c("-1h"="seagreen", "+7d"="steelblue")
isotype_colors <- c("IgM"="darkorchid", "IgD"="firebrick",
                   "IgG"="seagreen", "IgA"="steelblue")

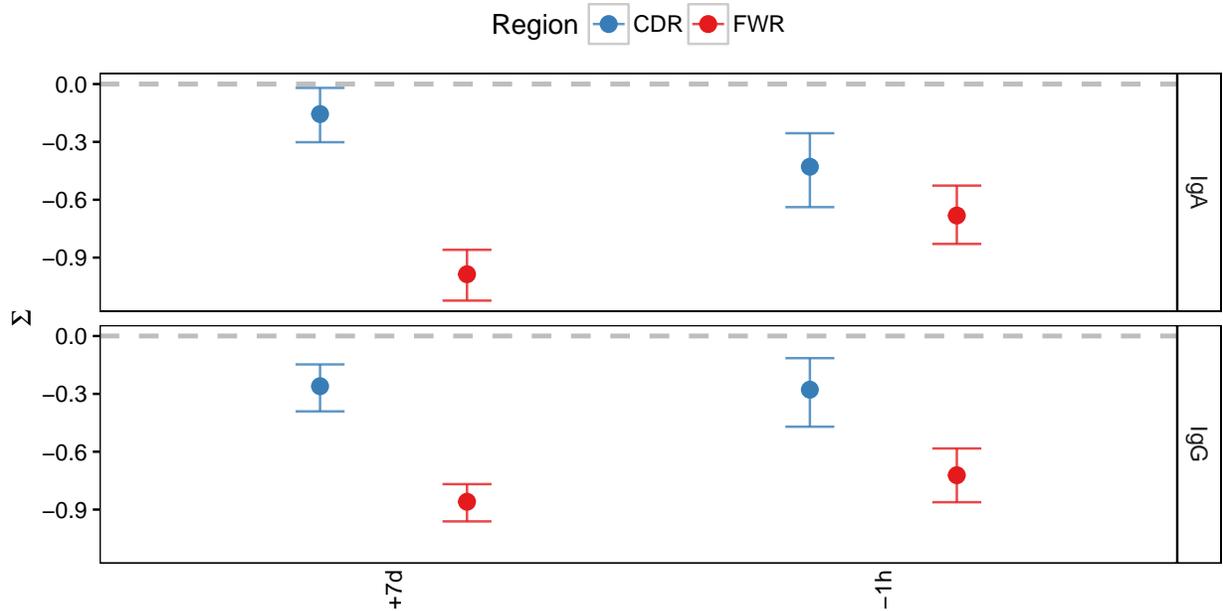
# Plot mean and confidence interval by time-point
plotBaselineSummary(grouped_1, "SAMPLE")
```



```
# Plot selection scores by time-point and isotype for only CDR
plotBaselineSummary(grouped_2, "SAMPLE", "ISOTYPE", groupColors=isotype_colors,
                   subsetRegions="CDR")
```



```
# Group by CDR/FWR and facet by isotype
plotBaselineSummary(grouped_2, "SAMPLE", "ISOTYPE", facetBy="group")
```



plotBaselineDensity plots the full Baseline PDF of selection scores for the given groups. The parameters are the same as those for plotBaselineSummary. However, rather than plotting the mean and confidence interval, the full density function is shown.

```
# Plot selection PDFs for a subset of the data
plotBaselineDensity(grouped_2, "ISOTYPE", groupColumn="SAMPLE", colorElement="group",
```

```
colorValues=sample_colors, sigmaLimits=c(-1, 1))
```

