

# Ensemble NMA of *E.coli* DHFR structures

Lars Skjaerven, Xin-Qiu Yao, Guido Scarabelli & Barry J. Grant

September 28, 2016

## Background

Bio3D<sup>1</sup> is an R package that provides interactive tools for structural bioinformatics. The primary focus of Bio3D is the analysis of bimolecular structure, sequence and simulation data (Grant et al. 2006).

Normal mode analysis (NMA) is one of the major simulation techniques used to probe large-scale motions in biomolecules. Typical application is for the prediction of functional motions in proteins. Version 2.0 of the Bio3D package now includes extensive NMA facilities (Skjaerven et al. 2015). These include a unique collection of multiple elastic network model force-fields, automated ensemble analysis methods, and variance weighted NMA (see also the [NMA Vignette](#)). Here we provide an in-depth demonstration of ensemble NMA with working code that comprise complete executable examples<sup>2</sup>.

## Requirements

Detailed instructions for obtaining and installing the Bio3D package on various platforms can be found in the [Installing Bio3D Vignette](#) available both on-line and from within the Bio3D package. In addition to Bio3D the *MUSCLE* multiple sequence alignment program (available from the [muscle home page](#)) must be installed on your system and in the search path for executables. Please see the installation vignette for further details.

## About this document

This vignette was generated using **Bio3D version 2.3.0**.

## 1 Part I: Ensemble NMA of *E.coli* DHFR structures

In this vignette we perform *ensemble NMA* on the complete collection of *E.coli* Dihydrofolate reductase (DHFR) structures in the protein data-bank (PDB). Starting from only one PDB identifier (PDB ID 1rx2) we show how to search the PDB for related structures using BLAST, fetch and align the structures, and finally calculate the normal modes of each individual structure in order to probe for potential differences in structural flexibility.

---

<sup>1</sup>The latest version of the package, full documentation and further vignettes (including detailed installation instructions) can be obtained from the main Bio3D website: <http://thegrantlab.org/bio3d/>

<sup>2</sup>This vignette contains executable examples, see `help(vignette)` for further details.

## 1.1 Search and retrieve DHFR structures

Below we perform a blast search of the PDB database to identify related structures to our query *E.coli* DHFR sequence. In this particular example we use function `get.seq()` to fetch the query sequence for chain A of the PDB ID 1RX2 and use this as input to `blast.pdb()`. Note that `get.seq()` would also allow the corresponding UniProt identifier.

```
library(bio3d)
```

```
aa <- get.seq("1rx2_A")
blast <- blast.pdb(aa)
```

```
## Searching ... please wait (updates every 5 seconds) RID = YR3W564X014
## .
## Reporting 576 hits
```

Function `plot.blast()` facilitates the visualization and filtering of the Blast results. It will attempt to set a seed position to the point of largest drop-off in normalized scores (i.e. the biggest jump in E-values). In this particular case we specify a cutoff (after initial plotting) of 225 to include only the relevant *E.coli* structures:

```
hits <- plot(blast, cutoff=225)
```

```
## * Possible cutoff values:    253 25
##           Yielding Nhits:    125 576
##
## * Chosen cutoff value of:    225
##           Yielding Nhits:    125
```

The Blast search and subsequent filtering identified a total of 101 related PDB structures to our query sequence. The PDB identifiers of this collection are accessible through the `pdb.id` attribute to the `hits` object (`hits$ pdb.id`). Note that adjusting the cutoff argument (to `plot.blast()`) will result in a decrease or increase of hits.

We can now use function `get.pdb()` and `pdbslit()` to fetch and parse the identified structures. Finally, we use `pdbaln()` to align the PDB structures.

```
# fetch PDBs
```

```
raw.files <- get.pdb(hits$ pdb.id, path = "raw_pdb", gzip=TRUE)
```

```
# split by chain ID
```

```
files <- pdbslit(raw.files, ids = hits$ pdb.id, path = "raw_pdb/split_chain", ncore=4)
```

```
# align structures
```

```
pdb.all <- pdbaln(files, fit=TRUE)
```

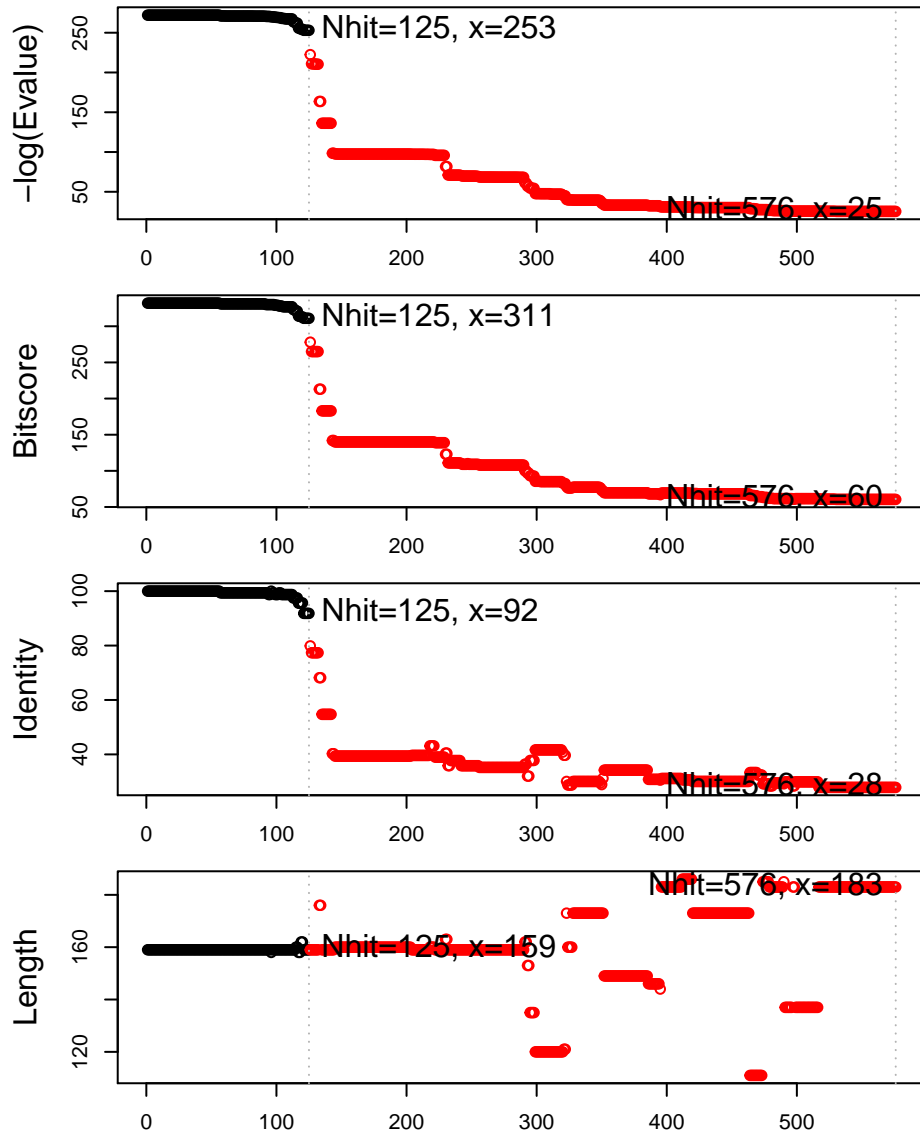


Figure 1: Blast results. Visualize and filter blast results through function `plot.blast()`. Here we proceed with only the top scoring hits (black).

The *pdb.all* object now contains *aligned* C-alpha atom data, including Cartesian coordinates, residue numbers, residue types, and B-factors. The sequence alignment is also stored by default to the FASTA format file 'aln.fa' (to view this you can use an alignment viewer such as SEAVIEW, see *Requirements* section above). In cases where manual edits of the alignment are necessary you can re-read the sequence alignment and coordinate data with:

```
aln <- read.fasta('aln.fa')
pdbs.all <- read.fasta.pdb(aln)
```

At this point the *pdb.all* object contains all identified 101 structures. This might include structures with missing residues, and/or multiple structurally redundant conformers. For our subsequent NMA missing in-structure residues might bias the calculation, and redundant structures can be useful to omit to reduce the computational load. Below we inspect the connectivity of the PDB structures with a function call to **inspect.connectivity()**, and **trim.pdbs()** to filter out those structures from our *pdb.all* object. Similarly, we omit structures that are conformationally redundant to reduce the computational load with function **filter.rmsd()**:

```
# remove structures with missing residues
conn <- inspect.connectivity(pdbs.all, cut=4.05)
pdbs <- trim(pdbs.all, row.inds=which(conn))

# which structures are omitted
which(!conn)

# remove conformational redundant structures
rd <- filter.rmsd(pdbs$xyz, cutoff=0.1, fit=TRUE)
pdbs <- trim(pdbs, row.inds=rd$ind)

# remove "humanized" e-coli dhfr
excl <- unlist(lapply(c("3QL0", "4GH8"), grep, pdbs$id))
pdbs <- trim(pdbs, row.inds=which(!(1:length(pdbs$id) %in% excl)))

# a list of PDB codes of our final selection
ids <- unlist(strsplit(basename(pdbs$id), split=".pdb"))
```

Use **print()** to see a short summary of the *pdbs* object:

```
print(pdbs, alignment=FALSE)

##
## Call:
## trim.pdbs(pdbs = pdbs, row.inds = which(!(1:length(pdbs$id) %in%
## excl)))
##
## Class:
## pdbs, fasta
##
```

```
## Alignment dimensions:
## 108 sequence rows; 166 position columns (159 non-gap, 7 gap)
##
## + attr: id, xyz, resno, b, chain, ali, resid, call
```

## 1.2 Annotate collected PDB structures

Function `pdb.annotate()` provides a convenient way of annotating the PDB files we have collected. Below we use the function to annotate each structure to its source species. This will come in handy when annotating plots later on:

```
anno <- pdb.annotate(ids)
print(unique(anno$source))
```

```
## [1] "Escherichia coli"      "Klebsiella pneumoniae"
```

## 1.3 Principal component analysis

A principal component analysis (PCA) can be performed on the structural ensemble (stored in the `pdb`s object) with function `pca.xyz()`. To obtain meaningful results we first superimpose all structures on the *invariant core* (function `core.find()`).

```
# find invariant core
core <- core.find(pdb)s)

# superimpose all structures to core
pdb$sxyz = pdbfit(pdb)s, core$c0.5A.xyz)

# identify gaps, and perform PCA
gaps.pos <- gap.inspect(pdb)sxyz)
gaps.res <- gap.inspect(pdb)sali)
pc.xray <- pca.xyz(pdb)sxyz[,gaps.pos$f.inds])
```

```
# plot PCA
plot(pc.xray)
```

Note that a call to the wrapper function `pca()` would provide identical results as the above code:

```
pc.xray <- pca(pdb)s, core.find=TRUE)
```

Function `rmsd()` will calculate all pairwise RMSD values of the structural ensemble. This facilitates clustering analysis based on the pairwise structural deviation:

```
rd <- rmsd(pdb)
hc.rd <- hclust(dist(rd))
grps.rd <- cutree(hc.rd, k=4)
```

```
plot(pc.xray$z[,1:2], col="grey50", pch=16, cex=1.3,
     ylab="Prinipcal Component 2", xlab="Principal Component 1")
points(pc.xray$z[,1:2], col=grps.rd, pch=16, cex=0.9)
```




Figure 2: Projection of the X-ray conformers shows that the E.coli DHFR structures can be divided into three major groups along their two first eigenvectors: closed, open, and occluded conformations. Each dot is colored according to their cluster membership: occluded conformations (green), open conformations (black), and closed conformations (red).

To visualize the major structural variations in the ensemble function `mktrj()` can be used to generate a trajectory PDB file by interpolating along the eigenvector:

```
mktrj(pc.xray, pc=1,
      resno=pdb$resno[1, gaps.res$f.inds],
      resid=pdb$resid[1, gaps.res$f.inds])
```

Function `pdbfit()` can be used to write the PDB files to separate directories according to their cluster membership:

```
pdbfit(trim(pdb, row.ind=which(grps.rd==1)), outpath="grps1") ## closed
pdbfit(trim(pdb, row.ind=which(grps.rd==2)), outpath="grps2") ## open
pdbfit(trim(pdb, row.ind=which(grps.rd==3)), outpath="grps3") ## occluded
```

## 1.4 Normal modes analysis

Function `nma.pdb()` will calculate the normal modes of each protein structure stored in the `pdb` object. The normal modes are calculated on the full structures as provided by object `pdb`. With the default argument `rm.gaps=TRUE` unaligned atoms are omitted from output:

```
modes <- nma(pdb, rm.gaps=TRUE, ncore=4)
```

The `modes` object of class `enma` contains aligned normal mode data including fluctuations, RMSIP data, and aligned eigenvectors. A short summary of the `modes` object can be obtained by calling the function `print()`, and the aligned fluctuations can be plotted with function `plot.enma()`.

```
print(modes)
```

```
##
## Call:
##   nma.pdb(pdb = pdb, rm.gaps = TRUE, ncore = 4)
##
## Class:
##   enma
##
## Number of structures:
##   108
##
## Attributes stored:
##   - Root mean square inner product (RMSIP)
##   - Aligned atomic fluctuations
##   - Aligned eigenvectors (gaps removed)
##   - Dimensions of x$U.subspace: 477x471x108
##
## Coordinates were aligned prior to NMA calculations
##
## + attr: fluctuations, rmsip, U.subspace, L, full.nma, xyz,
##         call
```

```
# plot modes fluctuations
plot(modes, pdb=pdb, col=grps.rd, label=NULL)
```

```
## Re-reading PDB (3DRC_A) to extract SSE
```

```
## Warning in dssp.pdb(pdb.ref, ...): Non-protein residues detected in input
## PDB: CL, MTX, HOH
```

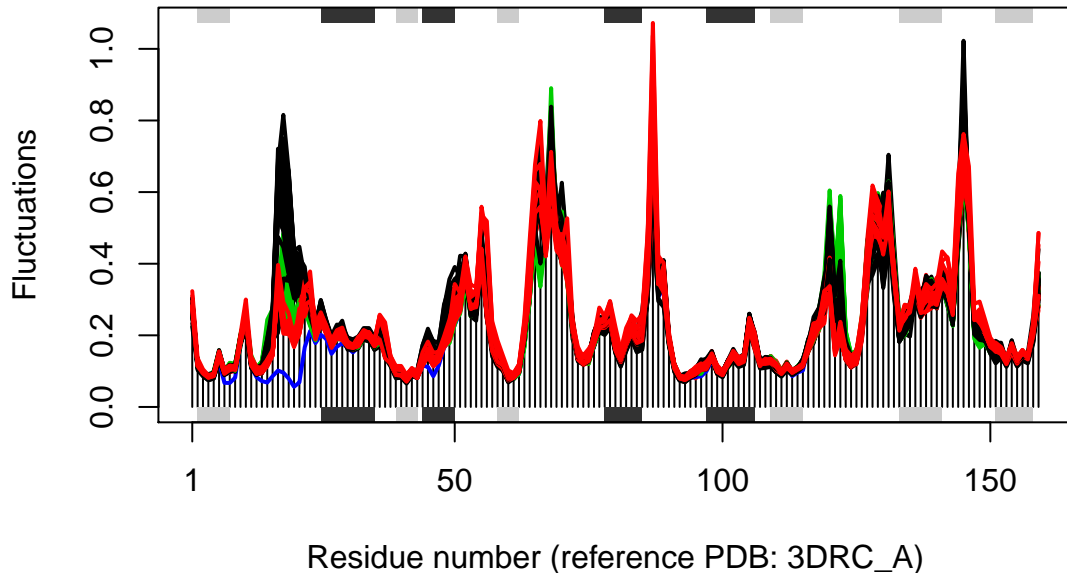


Figure 3: Normal mode fluctuations of the E.coli DHFR ensemble. Each line represent the mode fluctuations for the individual structures. Color coding according to structural clustering: green (occluded), black (open), and red (closed).

```
hc.nma <- hclust(as.dist(1-modes$rmsip))
grps.nma <- cutree(hc.nma, k=4)
```

```
heatmap(1-modes$rmsip, distfun = as.dist, labRow = ids, labCol = ids,
        ColSideColors=as.character(grps.nma), RowSideColors=as.character(grps.nma))
```

## 1.5 Fluctuation analysis

Comparing the mode fluctuations of two groups of structures can reveal specific regions of distinct flexibility patterns. Below we focus on the differences between the open (black), closed (red) and occluded (green) conformations of the *E.coli* structures:

```
cols <- grps.rd
cols[which(cols!=1 & cols!=2)]=NA
plot(modes, pdbs=pdbs, col=cols, signif=TRUE)
```



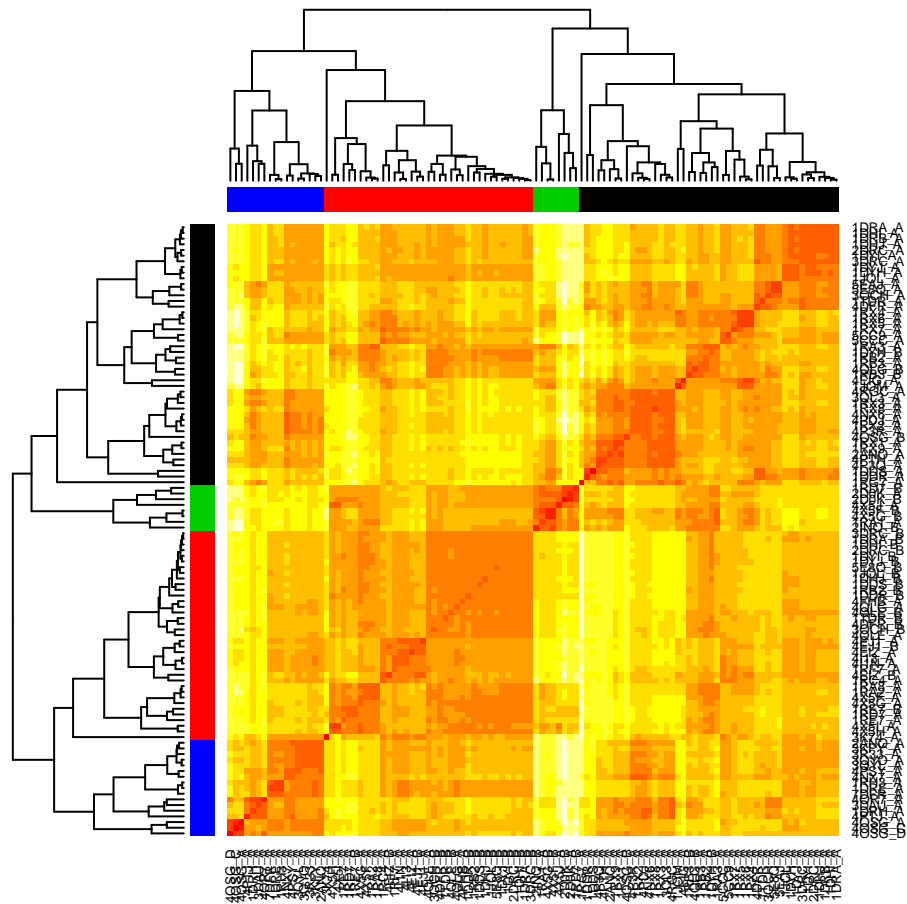


Figure 4: Clustering of structures based on their normal modes similarity (in terms of pair-wise RMSIP values).

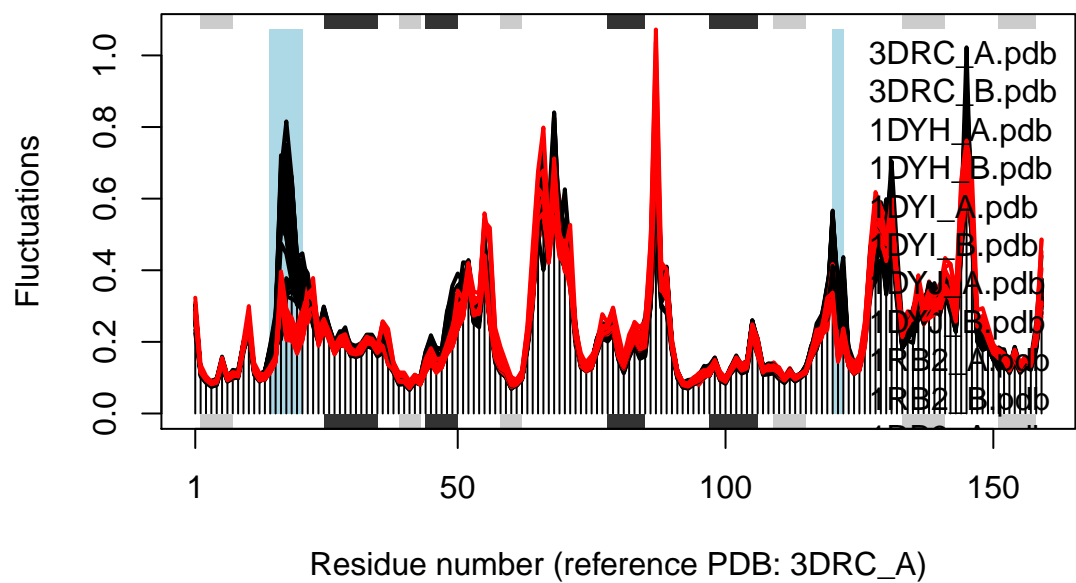


Figure 5: Comparison of mode fluctuations between open (black) and closed (red) conformers. Significant differences among the mode fluctuations between the two groups are marked with shaded blue regions.

```
## Re-reading PDB (3DRC_A) to extract SSE
```

```
cols <- grps.rd  
cols[which(cols!=1 & cols!=3)]=NA  
plot(modes, pdbs=pdbs, col=cols, signif=TRUE)
```

```
## Re-reading PDB (3DRC_A) to extract SSE
```




Figure 6: Comparison of mode fluctuations between open (black) and occluded (green) conformers. Significant differences among the mode fluctuations between the two groups are marked with shaded blue regions.

```
cols <- grps.rd  
cols[which(grps.rd!=2 & grps.rd!=3)]=NA  
plot(modes, pdbs=pdbs, col=cols, signif=TRUE)
```

```
## Re-reading PDB (2ANO_A) to extract SSE
```

```
## PDB has ALT records, taking A only, rm.alt=TRUE
```

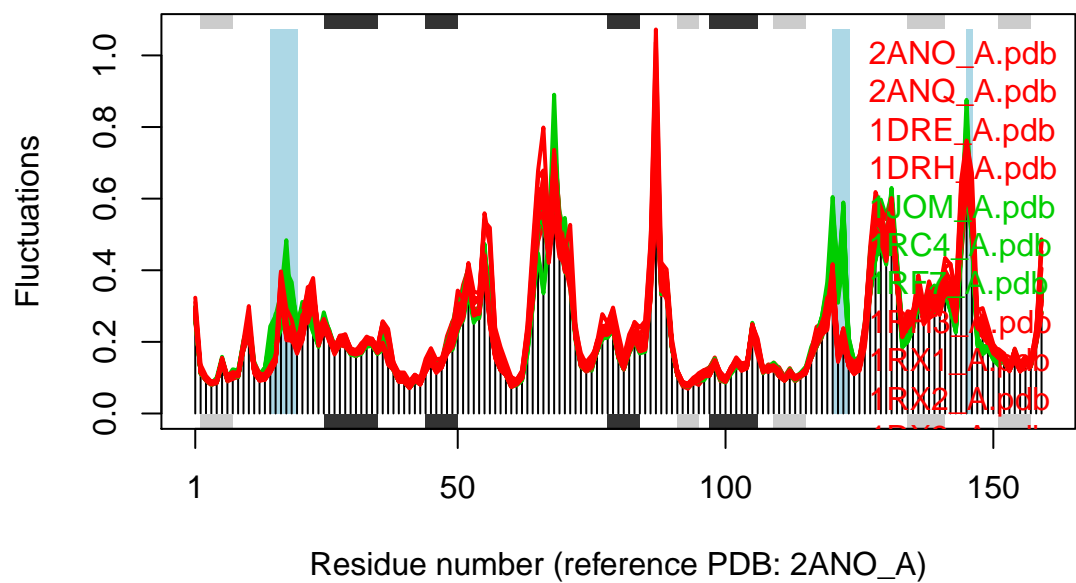


Figure 7: Comparison of mode fluctuations between closed (red) and occluded (green) conformers. Significant differences among the mode fluctuations between the two groups are marked with shaded blue regions.

## 1.6 Compare with MD simulation

The above analysis can also nicely be integrated with molecular dynamics (MD) simulations. Below we read in a 5 ns long MD trajectory (of PDB ID 1RX2). We visualize the conformational sampling by projecting the MD conformers onto the principal components (PCs) of the X-ray ensemble, and finally compare the PCs of the MD simulation to the normal modes:

```
pdb <- read.pdb("md-traj/1rx2-CA.pdb")
trj <- read.ncdf("md-traj/1rx2_5ns.nc")

inds <- pdb2aln.ind(pdb, gaps.res$f.inds)
inds.core <- pdb2aln.ind(pdb, core$c0.5A.atom)

# inds$a is for the pdbs object
# inds$b is for the MD structure

trj.fit <- fit.xyz(pdb$xyz[1,], trj, core$c0.5A.xyz, inds$core$b$xyz)
proj <- project.pca(trj.fit[, inds$b$xyz], pc.xray)
cols <- densCols(proj[,1:2])
```

```
plot(proj[,1:2], col=cols, pch=16,
      ylab="Prinipcal Component 2", xlab="Principal Component 1",
      xlim=range(pc.xray$z[,1]), ylim=range(pc.xray$z[,2]))
points(pc.xray$z[,1:2], col=1, pch=1, cex=1.1)
points(pc.xray$z[,1:2], col=grps.rd, pch=16)
```

```
# PCA of the MD trajectory
pc.md <- pca.xyz(trj.fit[, inds$b$xyz])

# compare MD-PCA and NMA
r <- rmsip(pc.md$U, modes$U.subspace[, ,1])

print(r)
```

```
## $overlap
##      b1  b2  b3  b4  b5  b6  b7  b8  b9  b10
## a1  0.529 0.008 0.003 0.027 0.000 0.000 0.001 0.001 0.000 0.019
## a2  0.017 0.070 0.019 0.005 0.026 0.027 0.006 0.000 0.014 0.057
## a3  0.035 0.207 0.006 0.004 0.097 0.002 0.006 0.010 0.001 0.033
## a4  0.003 0.165 0.139 0.105 0.045 0.018 0.051 0.052 0.009 0.000
## a5  0.003 0.013 0.011 0.008 0.065 0.003 0.009 0.049 0.008 0.001
## a6  0.002 0.046 0.000 0.031 0.041 0.082 0.018 0.082 0.059 0.000
## a7  0.003 0.086 0.106 0.071 0.004 0.001 0.047 0.009 0.051 0.003
## a8  0.018 0.023 0.220 0.331 0.027 0.102 0.006 0.003 0.000 0.005
## a9  0.082 0.083 0.007 0.000 0.009 0.001 0.064 0.069 0.011 0.001
## a10 0.001 0.003 0.121 0.011 0.193 0.006 0.015 0.002 0.002 0.005
```




Figure 8: Projection of MD conformers onto the X-ray PC space provides a two dimensional representation of the conformational sampling along the MD simulation (blue dots).

```
##
## $rmsip
## [1] 0.6443585
##
## attr(,"class")
## [1] "rmsip"
```

```
plot(r, xlab="MD PCA", ylab="NMA")
```




Figure 9: Overlap map between normal modes and principal components of a 5 ns long MD simulation. The two subsets yields an RMSIP value of 0.64, where a value of 1 would indicate identical directionality.

```
# compare MD-PCA and X-rayPCA
r <- rmsip(pc.md, pc.xray)
```

## 1.7 Domain analysis with GeoStaS

Identification of regions in the protein that move as rigid bodies is facilitated with the implementation of the GeoStaS algorithm (Romanowska, Nowinski, and Trylska 2012). Below we demonstrate the

use of function `geostas()` on data obtained from ensemble NMA, an ensemble of PDB structures, and a 5 ns long MD simulation. See `help(geostas)` for more details and further examples.

**GeoStaS on a PDB ensemble:** Below we input the `pdbs` object to function `geostas()` to identify dynamic domains. Here, we attempt to divide the structure into 2 sub-domains using argument `k=2`. Function `geostas()` will return a `grps` attribute which corresponds to the domain assignment for each C-alpha atom in the structure. Note that we use argument `fit=FALSE` to avoid re-fitting the coordinates since. Recall that the coordinates of the `pdbs` object has already been superimposed to the identified invariant core (see above). To visualize the domain assignment we write a PDB trajectory of the first principal component (of the Cartesian coordinates of the `pdbs` object), and add argument `chain=gs.xray$grps` to replace the chain identifiers with the domain assignment:

```
# Identify dynamic domains
gs.xray <- geostas(pdbs, k=2, fit=FALSE)

# Visualize PCs with colored domains (chain ID)
mktrj(pc.xray, pc=1, chain=gs.xray$grps)
```

**GoeStaS on ensemble NMA:** We can also identify dynamic domains from the normal modes of the ensemble of 82 protein structures stored in the `modes` object. By using function `mktrj.enma()` we generate a trajectory from the first five modes for all 82 structures. We then input this trajectory to function `geostas()`:

```
# Build conformational ensemble
trj.nma <- mktrj.enma(modes, pdbs, m.ind=1:5, s.ind=NULL, mag=10, step=2, rock=FALSE)

trj.nma
```

```
##
## Total Frames#: 5940
## Total XYZs#: 477, (Atoms#: 159)
##
## [1] 25.125 59.135 5.302 <...> 33.862 47.305 6.474 [2833380]
##
## + attr: Matrix DIM = 5940 x 477
```

```
# Fit to invariant core
trj.nma <- fit.xyz(trj.nma[1,], trj.nma,
                  fixed.ind=core$c0.5A.xyz,
                  mobile.ind=core$c0.5A.xyz)
```

```
# Run geostas to find domains
gs.nma <- geostas(trj.nma, k=2, fit=FALSE)
```

```
## .. 'xyz' coordinate data with 5940 frames
## .. coordinates are not superimposed prior to geostas calculation
## .. calculating atomic movement similarity matrix ('amsm.xyz()')
```



```
## .. dimensions of AMSM are 159x159
## .. clustering AMSM using 'kmeans'

# Write NMA generated trajectory with domain assignment
write.pdb(xyz=trj.nma, chain=gs.nma$grps)
```

Note that function `geostas.enma()` wraps these three main functions calls to facilitate this calculation. See `help(geostas)` for more information.

**GeoStaS on a MD trajectory:** The domain analysis can also be performed on the trajectory data obtained from the MD simulation (see above). Recall that the MD trajectory has already been superimposed to the invariant core. We therefore use argument `fit=FALSE` below. We then perform a new PCA of the MD trajectory, and visualize the domain assignments with function `mktrj()`:

```
gs.md <- geostas(trj, k=2, fit=FALSE)
pc.md <- pca(trj, fit=FALSE)
mktrj(pc.md, pc=1, chain=gs.md$grps)
```




Figure 10: Visualization of domain assignment using function `geostas()` on the first five normal modes of the entire ensemble of 82 DHFR structures.

## 1.8 Measures for modes comparison

Bio3D now includes multiple measures for the assessment of similarity between two normal mode objects. This enables clustering of related proteins based on the predicted modes of motion. Below

we demonstrate the use of root mean squared inner product (RMSIP), squared inner product (SIP), covariance overlap, bhattacharyya coefficient, and PCA of the corresponding covariance matrices.

```
# Similarity of atomic fluctuations
```

```
sip <- sip(modes)
hc.sip <- hclust(as.dist(1-sip), method="ward.D2")
grps.sip <- cutree(hc.sip, k=3)
```

```
hclustplot(hc.sip, k=3, colors=grps.rd, labels=ids, cex=0.7, main="SIP", fillbox=FALSE)
```

```
par(fig=c(.65, 1, .45, 1), new = TRUE)
plot(pc.xray$z[,1:2], col="grey50", pch=16, cex=1.1,
     ylab="", xlab="", axes=FALSE, bg="red")
points(pc.xray$z[,1:2], col=grps.sip, pch=16, cex=0.7)
box()
```

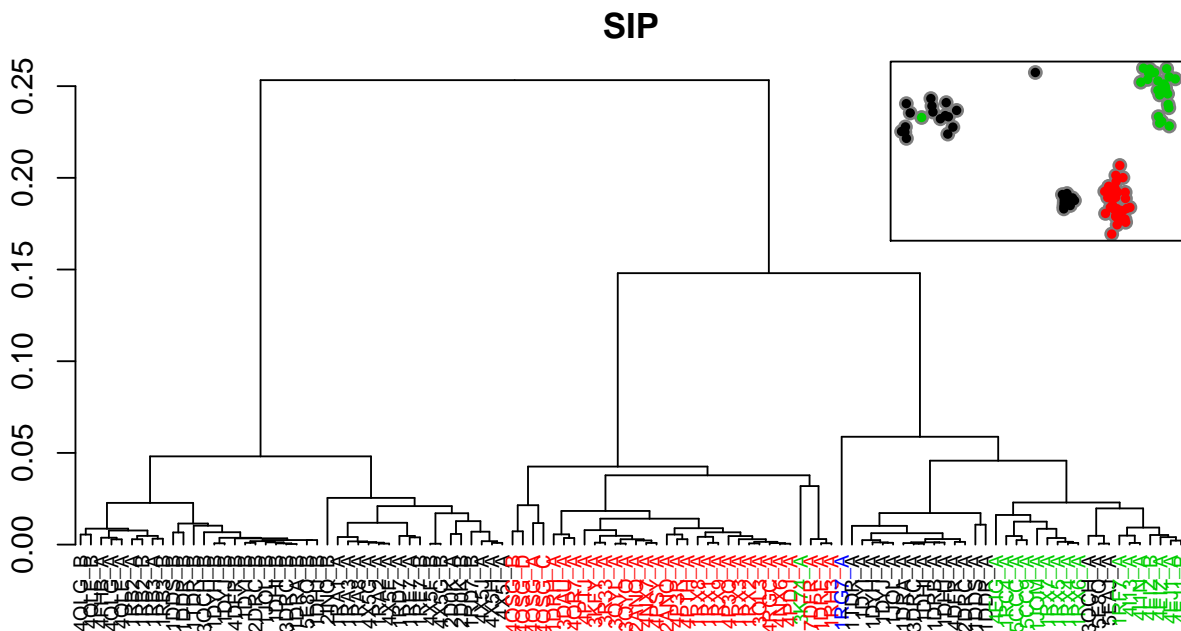


Figure 11: Dendrogram shows the results of hierarchical clustering of structures based on the similarity of atomic fluctuations calculated from NMA. Colors of the labels depict associated conformational state: green (occluded), black (open), and red (closed). The inset shows the conformerplot (see Figure 2), with colors according to clustering based on pairwise SIP values.

```
# RMSIP
```

```
rmsip <- rmsip(modes)
hc.rmsip <- hclust(dist(1-rmsip), method="ward.D2")
grps.rmsip <- cutree(hc.rmsip, k=3)
```

```

hclustplot(hc.rmsip, k=3, colors=grps.rd, labels=ids, cex=0.7, main="RMSIP", fillbox=FALSE)

par(fig=c(.65, 1, .45, 1), new = TRUE)
plot(pc.xray$z[,1:2], col="grey50", pch=16, cex=1.1,
     ylab="", xlab="", axes=FALSE)
points(pc.xray$z[,1:2], col=grps.rmsip, pch=16, cex=0.7)
box()

```




Figure 12: Dendrogram shows the results of hierarchical clustering of structures based on their pairwise RMSIP values (calculated from NMA). Colors of the labels depict associated conformational state: green (occluded), black (open), and red (closed). The inset shows the conformerplot (see Figure 2), with colors according to clustering based on the pairwise RMSIP values.

```

# Covariance overlap
co <- covsoverlap(modes, subset=200)
hc.co <- hclust(as.dist(1-co), method="ward.D2")
grps.co <- cutree(hc.co, k=3)

```

```

hclustplot(hc.co, k=3, colors=grps.rd, labels=ids, cex=0.7, main="Covariance overlap", fillbox=
par(fig=c(.65, 1, .45, 1), new = TRUE)
plot(pc.xray$z[,1:2], col="grey50", pch=16, cex=1.1,
     ylab="", xlab="", axes=FALSE)
points(pc.xray$z[,1:2], col=grps.co, pch=16, cex=0.7)
box()

```

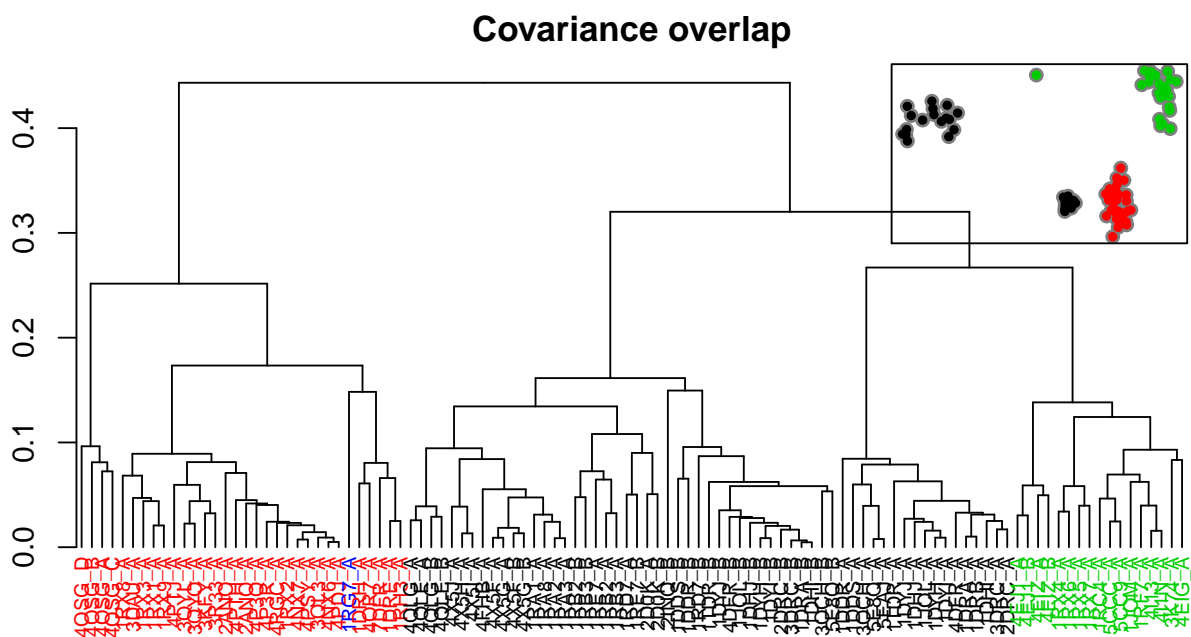


Figure 13: Dendrogram shows the results of hierarchical clustering of structures based on their pairwise covariance overlap (calculated from NMA). Colors of the labels depict associated conformational state: green (occluded), black (open), and red (closed). The inset shows the conformerplot (see Figure 2), with colors according to clustering of the Covariance overlap measure.

```

# Bhattacharyya coefficient
covs <- cov.enma(modes)
bc <- bhattacharyya(modes, covs=covs)
hc.bc <- hclust(dist(1-bc), method="ward.D2")
grps.bc <- cutree(hc.bc, k=3)

hclustplot(hc.bc, k=3, colors=grps.rd, labels=ids, cex=0.7, main="Bhattacharyya coefficient",
par(fig=c(.65, 1, .45, 1), new = TRUE)
plot(pc.xray$z[,1:2], col="grey50", pch=16, cex=1.1,
      ylab="", xlab="", axes=FALSE)
points(pc.xray$z[,1:2], col=grps.bc, pch=16, cex=0.7)
box()

```




Figure 14: Dendrogram shows the results of hierarchical clustering of structures based on their pairwise Bhattacharyya coefficient (calculated from NMA). Colors of the labels depict associated conformational state: green (occluded), black (open), and red (closed). The inset shows the conformerplot (see Figure 2), with colors according to clustering of the pairwise Bhattacharyya coefficients.

```

# PCA of covariance matrices
pc.covs <- pca.array(covs)
hc.covs <- hclust(dist(pc.covs$z[,1:2]), method="ward.D2")
grps.covs <- cutree(hc.covs, k=3)

```

```

hclustplot(hc.covs, k=3, colors=grps.rd, labels=ids, cex=0.7, main="PCA of covariance matrices")

par(fig=c(.65, 1, .45, 1), new = TRUE)
plot(pc.xray$z[,1:2], col="grey50", pch=16, cex=1.1,
     ylab="", xlab="", axes=FALSE)
points(pc.xray$z[,1:2], col=grps.covs, pch=16, cex=0.7)
box()

```

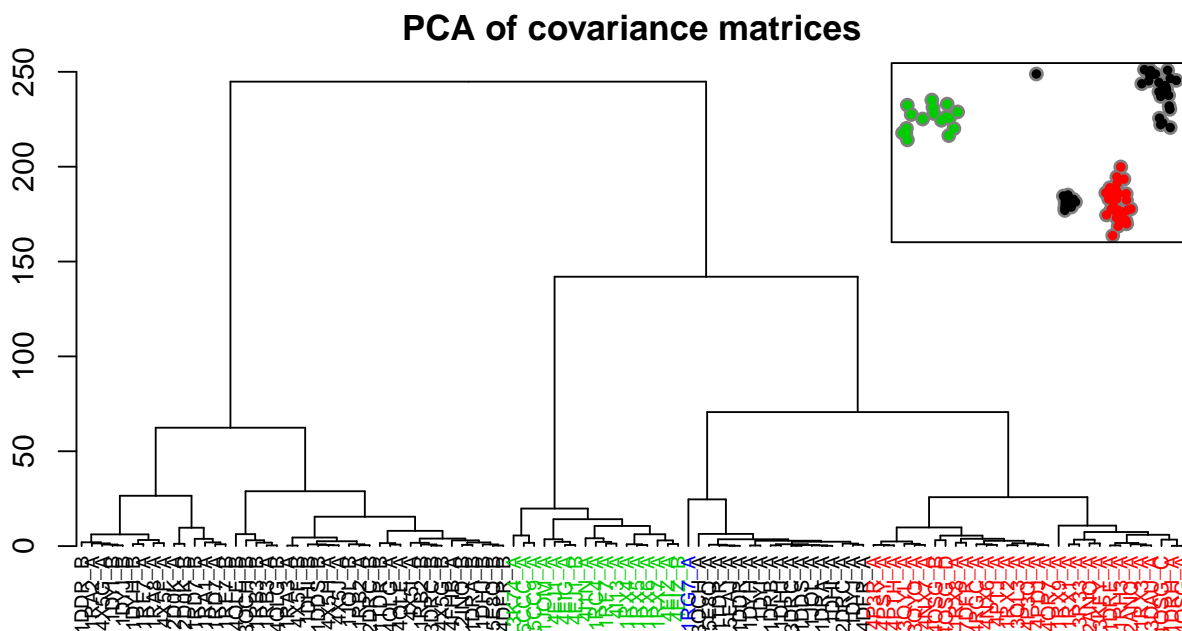


Figure 15: Dendrogram shows the results of hierarchical clustering of structures based on the PCA of covariance matrices (calculated from NMA). Colors of the labels depict associated conformational state: green (occluded), black (open), and red (closed). The inset shows the conformerplot (see Figure 2), with colors according to clustering based on PCA of covariance matrices.

## Document Details

This document is shipped with the Bio3D package in both R and PDF formats. All code can be extracted and automatically executed to generate Figures and/or the PDF with the following commands:

```

library(rmarkdown)
render("Bio3D_nma-dhfr-partI.Rmd", "all")

```

## Information About the Current Bio3D Session

```
print(sessionInfo(), FALSE)
```

```
## R version 3.3.1 (2016-06-21)
## Platform: x86_64-redhat-linux-gnu (64-bit)
## Running under: Fedora 24 (Twenty Four)
##
## attached base packages:
## [1] stats      graphics  grDevices  utils      datasets  methods   base
##
## other attached packages:
## [1] rmarkdown_1.0 bio3d_2.3-0
##
## loaded via a namespace (and not attached):
## [1] Rcpp_0.12.7      codetools_0.2-14  digest_0.6.10
## [4] grid_3.3.1       formatR_1.4       magrittr_1.5
## [7] evaluate_0.9     ncdf4_1.15        KernSmooth_2.23-15
## [10] stringi_1.1.1    tools_3.3.1       stringr_1.0.0
## [13] yaml_2.1.13      parallel_3.3.1    htmltools_0.3.5
## [16] knitr_1.14
```

## References

- Grant, B.J., A.P.D.C Rodrigues, K.M. Elsayy, A.J. Mccammon, and L.S.D. Caves. 2006. “Bio3d: An R Package for the Comparative Analysis of Protein Structures.” *Bioinformatics* 22: 2695–6. doi:[10.1093/bioinformatics/btl461](https://doi.org/10.1093/bioinformatics/btl461).
- Romanowska, Julia, Krzysztof S. Nowinski, and Joanna Trylska. 2012. “Determining geometrically stable domains in molecular conformation sets.” *Journal of Chemical Theory and Computation* 8 (8): 2588–99. doi:[10.1021/ct300206j](https://doi.org/10.1021/ct300206j).
- Skjaerven, L., X.Q. Yao, G. Scarabelli, and B.J. Grant. 2015. “Integrating Protein Structural Dynamics and Evolutionary Analysis with Bio3D.” *BMC Bioinformatics* 15: 399. doi:[10.1186/s12859-014-0399-6](https://doi.org/10.1186/s12859-014-0399-6).