Discovering and analyzing DNA sequence motifs
The rGADEM package.

Arnaud Droit *and Raphael Gottardo †

October 17, 2016

A step-by-step guide in the analysis of DNA sequence motifs using the rGADEM package in R

Contents

---

*arnaud.droit@crchuq.ulaval.ca
†rgottard@fhcrc.org>

Part I

Licensing

rGADEM is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version. We ask you to cite the following paper if you use this software for publication.

> L. Leiping. GADEM: A Genetic Algorithm Guided Formation of Spaced Dyads Coupled with an EM Algorithm for Motif Discovery. J Comput Biology

Part II

Introduction

In our guide, we include examples of code that we hope will help you when using the rGADEM package. The examples are kept at the basic level for ease of understanding. Some of the options in the functions have been set by default. To learn more about the exact parameters and usage of each function, you may type `help(FUNCTION_NAME)` of the function of interest in R after the rGADEM package is loaded.

Genome-wide analyses of protein binding sites generate large amounts of data; a ChIP data-set might contain 10,000 sites. Unbiased motif discovery in such datasets is not generally feasible using current methods that employ probabilistic models. We propose an efficient method, rGADEM, which combines spaced dyads and an expectation-maximization (EM) algorithm. Candidate words (four to six nucleotides) for constructing spaced dyads are prioritized by their degree of overrepresentation in the input sequence data. Spaced dyads are converted into starting position weight matrices (PWMs). rGADEM then employs a genetic algorithm (GA), with an embedded EM algorithm to improve starting PWMs, to guide the evolution of a population of spaced dyads toward one whose entropy scores are more statistically significant. Spaced dyads whose entropy scores reach a pre-specified significance threshold are declared motifs.

To use rGADEM, the user provide a set of coordinate include in the BED file or set of sequences in the FASTA format. The BED file contains location on chromosome, start and end position on the chromosome. For each line on

the BED file, we convert coordinate on a FASTA sequence. rGADEM has been developed in order to facilitate the discover and analysis of transcriptor factors and it is designed to works with the identification of motif package : MotIV. Thus, you can use the object returns by the rGADEM package to use MotIV (see MotIV package in Bioconductor for more detail).

The C code in rGADEM makes use of Grand Central Dispatch on Mac OS X 10.6 (Apple Inc) and openMP (openMP.org), with no user configuration, which greatly facilitates parallel processing and improve computing time on multicore machines (i.e. most modern computers). This implementation can actually reduce the computing time by a factor of 10, from several hours to a few minutes (depending on the number of input sequences).

Part III
Step-by-step Guide

1 rGADEM package and Packages

To load the rGADEM package, you should use this commande :

```
> library(rGADEM)
```

In the case of your data provide from Homo Sapiens :

```
> library(BSgenome.Hsapiens.UCSC.hg19)
```

2 Loading in the data

The next step is to load data from BED files or FASTA format in the R environment. The sequences are stored in some of the basic containers defined in the `Biostrings` package (see Biostring's document for more information). So ,the data can be manipulated in a consistent and easy way.
The maximum number of sequences allowed is 44 000. But it is possible to change the default parameters by editing the defines.h file and recompiling it.
The data used in this example are available in : extdata folder. The path for the data are : /rGADEM/inst/extdata.

2.1 From a BED file

Each line on the BED file contain the location, start and end position on the chromosome.

```
> pwd<-"" #INPUT FILES- BedFiles, FASTA, etc.
> path<- system.file("extdata/Test_100.bed",package="rGADEM")
> BedFile<-paste(pwd,path,sep="")
> BED<-read.table(BedFile,header=FALSE,sep="\t")
> BED<-data.frame(chr=as.factor(BED[,1]),start=as.numeric(BED[,2]),end=as.numeric(BED
```

Once the data have been read, we can create the list of 'RangedData' object (from IRanges package) where each element of the list corresponds to a different chromosome.

```
> rgBED<-IRanges(start=BED[,2],end=BED[,3])
> Sequences<-RangedData(rgBED,space=BED[,1])
```

2.2 From a FASTA file

In the case you have a FASTA file, you can load the data as follow :

```
> pwd<-"" #INPUT FILES- BedFiles, FASTA, etc.
> path<- system.file("extdata/Test_100.fasta",package="rGADEM")
> FastaFile<-paste(pwd,path,sep="")
> Sequences <- read.DNAStringSet(FastaFile, "fasta")
```

3 rGADEM analysis

At this time, we are now ready to start rGADEM analysis. If you want more details about rGADEM parameters, you may type help(gadem) in R environment. In this example, we have defined two parameters for rGADEM :

- verbose =1 : Print immediate results on screen.

- genome = Hsapiens : specify the genome.

    We also describe two important parameters :

- P-Value cutoff: The P-Value cutoff controls the number of binding site in a motif. By default, the P-value cutoff is : 0.0002

- E-Value cutoff: The E-Value cutoff controls the number of motifs to be identified. By default, the E-value cutoff is : 0.0

```
> gadem<-GADEM(Sequences,verbose=1,genome=Hsapiens)

Retrieving sequences... Done.
*** Start C Programm ***
=======================================================================================
input sequence file:
number of sequences and average length:                                  50 202.0
Use pgf method to approximate llr null distribution
parameters estimated from sequences in:

number of GA generations & population size:                              5 100

PWM score p-value cutoff for binding site declaration:                    2.000000e-04
ln(E-value) cutoff for motif declaration:                    0.000000

number of EM steps:                                          40
minimal no. sites considered for a motif:                    2

[a,c,g,t] frequencies in input data:                         0.289937 0.210063
=======================================================================================
*** Running an unseeded analysis ***
GADEM cycle  1: enumerate and count k-mers... top 3  4, 5-mers: 18 22 32
Done.
Initializing GA... Done.
GADEM cycle[  1] generation[  1] number of unique motif: 3
   spacedDyad: cccttnnnnnnnnnnccacc motifConsensus: CyCyTwTyTCwGmvyCCnnA  0.80 fitnes
   spacedDyad: aagannnnggga         motifConsensus: AAGrrGywGrGr           1.00 fitnes
   spacedDyad: ccannnnnnnctc        motifConsensus: kCACyCCAryCTC          0.50 fitnes

GADEM cycle[  1] generation[  2] number of unique motif: 3
   spacedDyad: cccttnnnnnnnnnnccacc motifConsensus: CyCyTwTyTCwGmvyCCnnA  0.80 fitnes
   spacedDyad: aagannnnggga         motifConsensus: AAGrrGywGrGr           1.00 fitnes
   spacedDyad: ccannnnnnnctc        motifConsensus: kCACyCCAryCTC          0.50 fitnes

GADEM cycle[  1] generation[  3] number of unique motif: 5
   spacedDyad: cccttnnnnnnnnnnccacc motifConsensus: CyCyTwTyTCwGmvyCCnnA  0.80 fitnes
   spacedDyad: aagannnnggga         motifConsensus: AAGrrGyTGrGG           0.90 fitnes
```

6

```
    spacedDyad: ctcccnnnnnnnnncct      motifConsensus: yyCTsCwsCCTvArCCy     0.40 fitnes
    spacedDyad: ccannnnnnctc           motifConsensus: kCACyCCAryCTC         0.50 fitnes
    spacedDyad: cagnnnnnnnntccc         motifConsensus: CTkmTCyCTnsrCCy       0.90 fitnes

GADEM cycle[  1] generation[  4] number of unique motif: 4
    spacedDyad: cccttnnnnnnnnnccacc motifConsensus: CyCyTwTyTCwGmvyCCnnA  0.80 fitnes
    spacedDyad: aagannnnggga           motifConsensus: AAGrrGyTGrGG          0.90 fitnes
    spacedDyad: tcttnnnnnnnntgg         motifConsensus: TGTTGsCCAGGCTGG       0.20 fitnes
    spacedDyad: tctnnnnnnnnnncct        motifConsensus: yCTymTsyCTnmrCCy      0.40 fitnes

GADEM cycle[  1] generation[  5] number of unique motif: 5
    spacedDyad: cccttnnnnnnnnncca       motifConsensus: CyCyTryyTCwGmmyCyn    0.80 fitnes
    spacedDyad: aagannnnggga           motifConsensus: AAGrrGyTGrGG          0.90 fitnes
    spacedDyad: tctnnnnnnnnnncctc       motifConsensus: hCTCCwGCCTsrGCmTC     0.40 fitnes
    spacedDyad: tctnnnnnnnnnncct        motifConsensus: yCTymTsyCTnmrCCy      0.40 fitnes
    spacedDyad: tctgnnnnncctcc          motifConsensus: wCTGyAyyCCmAsC        0.50 fitnes

*** Running an unseeded analysis ***
GADEM cycle  2: enumerate and count k-mers... top 3  4, 5-mers: 12 20 22
Done.
Initializing GA... Done.
GADEM cycle[  2] generation[  1] number of unique motif: 1
    spacedDyad: ccctnnctg              motifConsensus: GCCAGGCTG             0.40 fitnes

GADEM cycle[  2] generation[  2] number of unique motif: 1
    spacedDyad: cagccnnnnnntct          motifConsensus: CAGmCTGGsywwCA        0.80 fitnes

GADEM cycle[  2] generation[  3] number of unique motif: 1
    spacedDyad: tgggnnnnnngctgg         motifConsensus: TGTnGbCCAGGCTrs       0.50 fitnes

GADEM cycle[  2] generation[  4] number of unique motif: 1
    spacedDyad: tgggnnnnnngctgg         motifConsensus: TGTnGbCCAGGCTrs       0.50 fitnes

GADEM cycle[  2] generation[  5] number of unique motif: 1
    spacedDyad: tgggnnnnnngctgg         motifConsensus: TGTnGbCCAGGCTrs       0.50 fitnes

*** Running an unseeded analysis ***
GADEM cycle  3: enumerate and count k-mers... top 3  4, 5-mers: 12 20 14
Done.
Initializing GA... Done.
```

```
GADEM cycle[  3] generation[  1] number of unique motif: 1
    spacedDyad: tttctnnnnnnnngggag   motifConsensus: TTTnwTyTTTTkAAArAr     1.00 fitnes

GADEM cycle[  3] generation[  2] number of unique motif: 1
    spacedDyad: tttctnnnnnnnngggag   motifConsensus: TTTnwTyTTTTkAAArAr     1.00 fitnes

GADEM cycle[  3] generation[  3] number of unique motif: 1
    spacedDyad: tttctnnnnnnnngggag   motifConsensus: TTTnwTyTTTTkAAArAr     1.00 fitnes

GADEM cycle[  3] generation[  4] number of unique motif: 1
    spacedDyad: tttctnnnnnnnngggag   motifConsensus: TTTnwTyTTTTkAAArAr     1.00 fitnes

GADEM cycle[  3] generation[  5] number of unique motif: 1
    spacedDyad: tttctnnnnnnnngggag   motifConsensus: TTTnwTyTTTTkAAArAr     1.00 fitnes
```

4 Seeded analysis

In a seeded analysis rGADEM does not generate the starting PWMs through spaced dyads and optimize them through a Genetic Algorithm. This makes seeded runs much faster than unseeded. The efficiency of seeded runs makes it practical, even for sequence sets consisting of thousands of ChIP-Seq peak cores, to assess several alternative seed PWMs when prior knowledge suggests that this may be advisable (for example, when several database motifs are plausible candidate seeds).

The main advantage of a seeded analysis over an unseeded analysis is its computational efficiency.We recommend a seeded analysis whenever a reasonable starting PWM is available. However, for de novo motif discovery , an unseeded analysis is necessary.

First step is to prepare a text file with your PWM. It could be a general database (JASPAR, Transfac©,...). Only STAT1 have been selected in our example but it is possible to select a list of PWMs.

```
> path<- system.file("extdata/jaspar2009.txt",package="rGADEM")
> seededPwm<-readPWMfile(path)
> grep("STAT1",names(seededPwm))
> STAT1.PWM=seededPwm[103]
```

At this step, we have two choice :

Only seeded analysis :

```
> gadem<-GADEM(Sequences,verbose=1,genome=Hsapiens,Spwm=STAT1.PWM, fixSeeded=TRUE)
```

or seeded analysis following by unseeded analysis :

```
> gadem<-GADEM(Sequences,verbose=1,genome=Hsapiens,Spwm=STAT1.PWM)
```

5   rGADEM output

At the end of analysis, gadem object have been created in your R current session. This object contain all of your data information about your analysis (sequence consensus, pwm, chromosome, pvalue...). In fact, gadem object is a list of object.

- align : This object contains the individual motifs identified but and the location (seqID and position) of the sites in the original sequence data.
- motif : This object contains contains PWM, motif consensus, motif length and all aligned sequences for a specific motif.
- parameters : This object contains contains parameters of rGADEM analysis.

For more details, please see the RD files for each object.

5.1   Acces to pwm

To view all PWM

```
> nOccurrences(gadem)
```

```
[1] 27 16
```

To view pwm for the motif 1 :

```
> nOccurrences(gadem)[1]
```

```
[1] 27
```

5.2 Acces to sequence consensus

To view all sequences consensus :

```
> consensus(gadem)
```

```
[1] "nCyyyTryyTCwGymTCyy" "mTGTnGsCCAGGCTrsw"
```

5.3 Acces to sequence consensus

To access to the first sequence :

```
> consensus(gadem)[1]
```

```
[1] "nCyyyTryyTCwGymTCyy"
```

5.4 Acces to chromosome position

To view start position on chromosome :

```
> startPos(gadem)
```

```
$CCTGTAATCTCCTAGTCTC
[1] 43552181
```

5.5 Acces to chromosome position

To view end position on chromosome :

```
> endPos(gadem)
```

```
$CCTGTAATCTCCTAGTCTC
[1] 43552382
```

5.6 Acces to parameters

And finaly, if you want to show parameters for this analysis :

```
> gadem@parameters
```