

Sequence manipulation and scanning

Benjamin Jean-Marie Tremblay*

14 March 2020

Abstract

Sequences stored as XStringSet objects (from the Biostrings package) can be used by several functions in the universalmotif package. These functions are demonstrated here and fall into two categories: sequence manipulation and motif scanning. Sequences can be generated, shuffled, and background frequencies of any order calculated. Scanning can be done simply to find locations of motif hits above a certain threshold, or to find instances of enriched motifs.

Contents

1	Introduction	1
2	Creating random sequences	2
3	Calculating sequence background	3
4	Shuffling sequences	4
5	Miscellaneous string utilities	5
6	Scanning sequences for motifs	6
7	Enrichment analyses	8
8	Testing for motif positional preferences in sequences	9
9	Motif discovery with MEME	10
	Session info	12
	References	14

1 Introduction

This vignette goes through generating your own sequences from a specified background model, shuffling sequences whilst maintaining a certain k -let size, and the scanning of sequences and scoring of motifs. For an introduction to sequence motifs, see the introductory vignette. For a basic overview of available motif-related functions, see the motif manipulation vignette. For a discussion on motif comparisons and P-values, see the motif comparisons and P-values vignette.

*b2tremblay@uwaterloo.ca

2 Creating random sequences

The **Biostrings** package offers an excellent suite of functions for dealing with biological sequences. The **universalmotif** package hopes to help extend these by providing the **create_sequences()** and **shuffle_sequences()** functions. The first of these, **create_sequences()**, in it's simplest form generates a set of letters in random order, then passes these strings to the **Biostrings** package. The number and length of sequences can be specified. The probabilities of individual letters can also be set.

The **freqs** option of **create_sequences()** also takes higher order backgrounds. In these cases the sequences are constructed in a Markov-style manner, where the probability of each letter is based on which letters precede it.

```
library(universalmotif)
library(Biostrings)

## Create some DNA sequences for use with an external program (default
## is DNA):

sequences.dna <- create_sequences(seqnum = 500,
                                freqs = c(A=0.3, C=0.2, G=0.2, T=0.3))

## writeXStringSet(sequences.dna, "dna.fasta")
sequences.dna
#> DNAStringSet object of length 500:
#>      width seq
#> [1] 100 CTTATGATACGTAAGTAGTAGGGCCGCCGCCGC...TTTAGCATGGGATTTATCCATTCCAATCCATC
#> [2] 100 TCCAGAAATATCACTACTAAAAATTATGAGTAT...GTTACTAGCCACGTTGTTGAGGTTAAGTTCTG
#> [3] 100 AAACCGCTAATGGAATAAGATGTCAGCAGACGC...AGTAAAAGATGGGTGTTCAAGCCCCAAAAATGT
#> [4] 100 GTTCAACTATCGGATTATCCCACTCATCACTAT...ACTATGACTGATCTGAGTCCCATTAAAGTTGA
#> [5] 100 TTTGTGGTAAACTCCTACCCTGACTTAGGGAGA...TCTACTAACAGAATCAGCTCATCCAATAAACG
#> ... ...
#> [496] 100 CACCGAACTTGTAATCGGTGATGGAAGCAGCT...TTCTTTGTTAGAAATTTCTATCAGACTATGTA
#> [497] 100 TTAGCTCCAAAGAATTCAGTAGGAAGCATGATC...GTCTGAGGGTAGTATATTACGAGTACGTCTTC
#> [498] 100 AGTGAACCACTCCCGCTCATAAGCATCTGCT...CGCTAGGGCCGTGAGCGCCCCTACAGCATAGT
#> [499] 100 GCGTTTGCAACCGCAGCGAACGTATGCGTATTA...AGCATTGCAAAATGAGGGTTGATGTAGAGAAGA
#> [500] 100 TACTGATCACTATCTGCAAATACGAATAGGCCG...TCCAGCGATTGACACTCAAGACACAGTTGGCC

## Amino acid:

create_sequences(alphabet = "AA")
#> AAStringSet object of length 100:
#>      width seq
#> [1] 100 QQPFGDNKVHHMYHPKIGFWTEVTTFYFKWHEAL...SVPFTLAGPALDRDHCKCKWNCHKWQYETRIN
#> [2] 100 HIFNLGDVSPPCYQGVSMKVPIRPPWL VVQIAW...MRGKQYCNGCYFLFPD VDVDDPVVIYHPKRD
#> [3] 100 YATTRKQGPWYNYVVDSPTKNNLK VQHTANAI...GNVPLLCVWCLHEHYEHEHTQEWGTCYMLDCQ
#> [4] 100 QRKDYNGRMFGDYHMSMCTRFSGKFVWSSISCT...AKMTHWDENDWLWKGFSFSSFFFRRSYQGTKG
#> [5] 100 HKAIERTEINPQYPEFVHAQAYHCCTEFRRYCG...SGDAEKDLFDKNPMPGFHFQSHNDQLWVCMST
#> ... ...
#> [96] 100 VIVMYHIGIVDAIMPMFKTVGKNNWRERSYLG...TFFPRPLSETKEPAVLNVTCQWWWPRYRLPAE
#> [97] 100 NCLTFLYRGEKMHVGYQPATCLKRQKERHQHQ...PCTTNDMCTTVHHEMAYGAFYGYIMKYVGHIS
#> [98] 100 FSCCLPMDDMSCHFVKYWFVTHFMQPPQPVHD...IWLAINNILVIKCFPLARYSANS LDYCDARH
#> [99] 100 VKQIRSCNATCNHMMVIEKRAECPVCPYCHN...DSCFFCNQDVMSHTQWCDWICVFKSWFWRDV
#> [100] 100 MDGPWWQAVDKDHVEHRLPPMFAVENDMNHGHA...VPRKANPYSWHQMKDRIDNTWEEPHLWKRLK

## Any set of characters can be used
```

```

create_sequences(alphabet = paste0(letters, collapse = ""))
#> BStringSet object of length 100:
#>      width seq
#> [1] 100 xawpugikhoetkjsknawxvclwlfmrkzuna...dndyqplbxxamgjdqfxumaibpkpmchh
#> [2] 100 vbsetnrvocimutlvbatvqfwtwlyjvzoaa...hbgxhfxcvyaymsephhlvrzbrcfueyfop
#> [3] 100 scotqtagwqmfecdgpbqsmiipiribfyina...lpjvyvjdswaktcgwlyrtmcaeuftkiw
#> [4] 100 qdkjnajrdeqzpmurcbmqhlmtxtqydaa...pcmulveqaxzlieopxqjzcfjkpiwlde
#> [5] 100 ndgykhscksuszwocqbjodoejfdklayxna...tqptfbhgowaajfukmsgdofmdshzaxiokl
#> ... ..
#> [96] 100 xuijunveciedbokjpkvjcsxqdywzwrval...emqxyivlggbnlxodsqqlputqbilnuklzz
#> [97] 100 uveysuepjwiwmdudkshxvjnoejrhqqnm...iztwoygmdbargqlvhrmqgrjkaxjwogg
#> [98] 100 swanpanarkmpuhvfqloetyujakvjrqqam...mnwufosobebmyqstzyxkmtsslqiyiqoo
#> [99] 100 pwxmdmhwllyyqjgroqellcobggqlhacqenm...qbtzweepzdbyezuadpdiigtbnfnsnutvu
#> [100] 100 natsjnewfnucragasmiakerdxusmpzam...uocrnuqqwdblkjwigfjgettkovccgwcd

```

3 Calculating sequence background

Sequence backgrounds can be retrieved for DNA and RNA sequences with `oligonucleotideFrequency()` from "Biostrings". Unfortunately, no such `Biostrings` function exists for other sequence alphabets. The `universalmotif` package proves `get_bkg()` to remedy this. Similarly, the `get_bkg()` function can calculate higher order backgrounds for any alphabet as well. It is recommended to use the original `Biostrings` for very long DNA and RNA sequences whenever possible though, as it is much faster than `get_bkg()`.

```

library(universalmotif)

## Background of DNA sequences:
dna <- create_sequences()
get_bkg(dna, k = 1:2, list.out = FALSE)
#>      A      C      G      T      AA      AC      AG
#> 0.24580000 0.25300000 0.24880000 0.25240000 0.06191919 0.06272727 0.06121212
#>      AT      CA      CC      CG      CT      GA      GC
#> 0.05979798 0.06373737 0.06171717 0.06303030 0.06464646 0.05979798 0.06636364
#>      GG      GT      TA      TC      TG      TT
#> 0.06242424 0.06010101 0.06020202 0.06212121 0.06222222 0.06797980

## Background of non DNA/RNA sequences:
qwerty <- create_sequences("QWERTY")
get_bkg(qwerty, k = 1:2, list.out = FALSE)
#>      E      Q      R      T      W      Y      EE
#> 0.16530000 0.16790000 0.16210000 0.16740000 0.16700000 0.17030000 0.02363636
#>      EQ      ER      ET      EW      EY      QE      QQ
#> 0.02616162 0.02686869 0.02929293 0.02666667 0.03262626 0.02727273 0.03040404
#>      QR      QT      QW      QY      RE      RQ      RR
#> 0.02898990 0.02727273 0.02676768 0.02717172 0.02909091 0.02686869 0.02656566
#>      RT      RW      RY      TE      TQ      TR      TT
#> 0.02787879 0.02565657 0.02595960 0.02676768 0.02888889 0.02868687 0.02525253
#>      TW      TY      WE      WQ      WR      WT      WW
#> 0.02979798 0.02808081 0.02838384 0.02878788 0.02646465 0.02767677 0.02848485
#>      WY      YE      YQ      YR      YT      YW      YY
#> 0.02717172 0.03030303 0.02707071 0.02464646 0.03030303 0.02919192 0.02888889

```

4 Shuffling sequences

When performing *de novo* motif searches or motif enrichment analyses, it is common to do so against a set of background sequences. In order to properly identify consistent patterns or motifs in the target sequences, it is important that there be maintained a certain level of sequence composition between the target and background sequences. This reduces results which are derived purely from differential letter frequency biases.

In order to avoid these results, typically it is desirable to use a set of background sequences which preserve a certain *k*-let size (such as dinucleotide or trinucleotide frequencies in the case of DNA sequences). Though for some cases a set of similar sequences may already be available for use as background sequences, usually background sequences are obtained by shuffling the target sequences, while preserving a desired *k*-let size. For this purpose, a commonly used tool is `uShuffle` (Jiang et al. 2008). The `universalmotif` package aims to provide its own *k*-let shuffling capabilities for use within R via `shuffle_sequences()`.

The `universalmotif` package offers three different methods for sequence shuffling: `euler`, `markov` and `linear`. The first method, `euler`, can shuffle sequences while preserving any desired *k*-let size. Furthermore 1-letter counts will always be maintained. However in order for this to be possible, the first and last letters will remain unshuffled. This method is based on the initial random Eulerian walk algorithm proposed by Altschul and Erickson (1985) and the subsequent cycle-popping algorithm detailed by Propp and Wilson (1998) for quickly and efficiently finding Eulerian walks.

The second method, `markov` can only guarantee that the approximate *k*-let frequency will be maintained, but not that the original letter counts will be preserved. The `markov` method involves determining the original *k*-let frequencies, then creating a new set of sequences which will have approximately similar *k*-let frequency. As a result the counts for the individual letters will likely be different. Essentially, it involves a combination of determining *k*-let frequencies followed by `create_sequences()`. This type of shuffling is discussed by Fitch (1983).

The third method `linear` preserves the original 1-letter counts exactly, but uses a more crude shuffling technique. In this case the sequence is split into sub-sequences every *k*-let (of any size), which are then re-assembled randomly. This means that while shuffling the same sequence multiple times with `method = "linear"` will result in different sequences, they will all have started from the same set of *k*-length sub-sequences (just re-assembled differently).

```
library(universalmotif)
library(Biostrings)
data(ArabidopsisPromoters)

## Potentially starting off with some external sequences:
# ArabidopsisPromoters <- readDNASTringSet("ArabidopsisPromoters.fasta")

euler <- shuffle_sequences(ArabidopsisPromoters, k = 2, method = "euler")
markov <- shuffle_sequences(ArabidopsisPromoters, k = 2, method = "markov")
linear <- shuffle_sequences(ArabidopsisPromoters, k = 2, method = "linear")
k1 <- shuffle_sequences(ArabidopsisPromoters, k = 1)
```

Let us compare how the methods perform:

```
o.letter <- get_bkg(ArabidopsisPromoters, 1, as.prob = FALSE, list.out = FALSE)
e.letter <- get_bkg(euler, 1, as.prob = FALSE, list.out = FALSE)
m.letter <- get_bkg(markov, 1, as.prob = FALSE, list.out = FALSE)
l.letter <- get_bkg(linear, 1, as.prob = FALSE, list.out = FALSE)

data.frame(original=o.letter, euler=e.letter, markov=m.letter, linear=l.letter)
#>   original euler markov linear
#> A   17384 17384  17451  17384
#> C    8081  8081   8121   8081
```

```

#> G      7583 7583 7616 7583
#> T      16952 16952 16862 16952

o.counts <- get_bkg(ArabidopsisPromoters, 2, as.prob = FALSE, list.out = FALSE)
e.counts <- get_bkg(euler, 2, as.prob = FALSE, list.out = FALSE)
m.counts <- get_bkg(markov, 2, as.prob = FALSE, list.out = FALSE)
l.counts <- get_bkg(linear, 2, as.prob = FALSE, list.out = FALSE)

data.frame(original=o.counts, euler=e.counts, markov=m.counts, linear=l.counts)
#>      original euler markov linear
#> AA      6893 6893 6151 6488
#> AC      2614 2614 2781 2693
#> AG      2592 2592 2627 2647
#> AT      5276 5276 5874 5536
#> CA      3014 3014 2791 2932
#> CC      1376 1376 1364 1309
#> CG      1051 1051 1211 1131
#> CT      2621 2621 2747 2701
#> GA      2734 2734 2705 2657
#> GC      1104 1104 1210 1171
#> GG      1176 1176 1186 1150
#> GT      2561 2561 2508 2598
#> TA      4725 4725 5787 5292
#> TC      2977 2977 2757 2899
#> TG      2759 2759 2585 2650
#> TT      6477 6477 5716 6096

```

5 Miscellaneous string utilities

Since biological sequences are usually contained in `XStringSet` class objects, `get_bkg()` and `shuffle_sequences()` are designed to work with such objects. For cases when strings are not `XStringSet` objects, the following functions are available:

- `count_klets()`: alternative to `get_bkg()`
- `shuffle_string()`: alternative to `shuffle_sequences()`

```

library(universalmotif)

string <- "DASDSDSASDSSA"

count_klets(string, 2)
#>      klets counts
#> 1      AA      0
#> 2      AD      0
#> 3      AS      2
#> 4      DA      1
#> 5      DD      1
#> 6      DS      3
#> 7      SA      2
#> 8      SD      3
#> 9      SS      1

shuffle_string(string, 2)
#> [1] "DSDSSDASDDASA"

```

Finally, the `get_klets()` function can be used to get a list of all possible k-lets for any sequence alphabet:

```
library(universalmotif)

get_klets(c("A", "S", "D"), 2)
#> [1] "AA" "AS" "AD" "SA" "SS" "SD" "DA" "DS" "DD"
```

6 Scanning sequences for motifs

There are many motif-programs available with sequence scanning capabilities, such as HOMER and tools from the MEME suite. The `universalmotif` package does not aim to supplant these, but rather provide convenience functions for quickly scanning a few sequences without needing to leave the R environment. Furthermore, these functions allow for taking advantage of the higher-order (`multifreq`) motif format described here.

Two scanning-related functions are provided: `scan_sequences()` and `enrich_motifs()`. The latter simply runs `scan_sequences()` twice on a set of target and background sequences. Given a motif of length `n`, `scan_sequences()` considers every possible `n`-length subset in a sequence and scores it using the PWM format. If the match surpasses the minimum threshold, it is reported. This is case regardless of whether one is scanning with a regular motif, or using the higher-order (`multifreq`) motif format (the `multifreq` matrix is converted to a PWM).

Before scanning a set of sequences, one must first decide the minimum logodds threshold for retrieving matches. This decision is not always the same between scanning programs out in the wild, nor is it usually told to the user what the cutoff is or how it is decided. As a result, `universalmotif` aims to be as transparent as possible in this regard by allowing for complete control of the threshold. For more details on PWMs, see the introductory vignette.

One way is to set a cutoff between 0 and 1, then multiplying the highest possible PWM score to get a threshold. The `matchPWM()` function from the `Biostrings` package for example uses a default of 0.8 (shown as "80%"). This is quite arbitrary of course, and every motif will end up with a different threshold. For high information content motifs, there is really no right or wrong threshold; as they tend to have fewer non-specific positions. This means that incorrect letters in a match will be more punishing. To illustrate this, contrast the following PWMs:

```
library(universalmotif)
m1 <- create_motif("TATATATATA", nsites = 50, type = "PWM", pseudocount = 1)
m2 <- matrix(c(0.10,0.27,0.23,0.19,0.29,0.28,0.51,0.12,0.34,0.26,
               0.36,0.29,0.51,0.38,0.23,0.16,0.17,0.21,0.23,0.36,
               0.45,0.05,0.02,0.13,0.27,0.38,0.26,0.38,0.12,0.31,
               0.09,0.40,0.24,0.30,0.21,0.19,0.05,0.30,0.31,0.08),
             byrow = TRUE, nrow = 4)
m2 <- create_motif(m2, alphabet = "DNA", type = "PWM")
m1["motif"]
#>           T           A           T           A           T           A           T
#> A -5.672425  1.978626 -5.672425  1.978626 -5.672425  1.978626 -5.672425
#> C -5.672425 -5.672425 -5.672425 -5.672425 -5.672425 -5.672425 -5.672425
#> G -5.672425 -5.672425 -5.672425 -5.672425 -5.672425 -5.672425 -5.672425
#> T  1.978626 -5.672425  1.978626 -5.672425  1.978626 -5.672425  1.978626
#>           A           T           A
#> A  1.978626 -5.672425  1.978626
#> C -5.672425 -5.672425 -5.672425
#> G -5.672425 -5.672425 -5.672425
#> T -5.672425  1.978626 -5.672425
m2["motif"]
#>           S           H           C           N           N           N
```

```
#> A -1.3219281 0.09667602 -0.12029423 -0.3959287 0.2141248 0.1491434
#> C 0.5260688 0.19976951 1.02856915 0.6040713 -0.1202942 -0.6582115
#> G 0.8479969 -2.33628339 -3.64385619 -0.9434165 0.1110313 0.5897160
#> T -1.4739312 0.66371661 -0.05889369 0.2630344 -0.2515388 -0.4102840
#>
#> R N N V
#> A 1.0430687 -1.0732490 0.4436067 0.04222824
#> C -0.5418938 -0.2658941 -0.1202942 0.51171352
#> G 0.0710831 0.5897160 -1.0588937 0.29598483
#> T -2.3074285 0.2486791 0.3103401 -1.65821148
```

In the first example, sequences which do not have a matching base in every position are punished heavily. The maximum logodds score in this case is approximately 20, and for each incorrect position the score is reduced approximately by 5.7. This means that a threshold of zero would allow for at most three mismatches. At this point, it is up to you how many mismatches you would deem appropriate.

This thinking becomes impossible for the second example. In this case, mismatches are much less punishing, to the point that one could ask: what even constitutes a mismatch? The answer to this question is much more difficult in cases such as these. An alternative to manually deciding upon a threshold is to instead start with maximum P-value one would consider appropriate for a match. If, say, we want matches with a P-value of at most 0.001, then we can use `motif_pvalue()` to calculate the appropriate threshold (see the comparisons and P-values vignette for details on motif P-values).

```
motif_pvalue(m2, pvalue = 0.001)
#> [1] 4.8779
```

Furthermore, the `scan_sequences()` function offers the ability to scan using the `multifreq` slot, if available. This allows to take into account inter-positional dependencies, and get matches which more faithfully represent the original sequences from which the motif originated.

```
library(universalmotif)
library(Biostrings)
data(ArabidopsisPromoters)

## A 2-letter example:

motif.k2 <- create_motif("CWWWCC", nsites = 6)
sequences.k2 <- DNAStringSet(rep(c("CAAAACC", "CTTTTCC"), 3))
motif.k2 <- add_multifreq(motif.k2, sequences.k2)
```

Regular scanning:

```
head(scan_sequences(motif.k2, ArabidopsisPromoters, RC = TRUE,
                    threshold = 0.9, threshold.type = "logodds"))
#> Note: detected non-finite positional weights, normalising
#> Note: since this motif has a pseudocount of 0, 1 will be used
#> DataFrame with 6 rows and 12 columns
#>   motif motif.i sequence start stop score match
#>   <character> <integer> <character> <integer> <integer> <numeric> <character>
#> 1 motif 1 AT4G28150 621 627 9.08 CTAAACC
#> 2 motif 1 AT1G19380 139 145 9.08 CTTATCC
#> 3 motif 1 AT1G19380 204 210 9.08 CTAAACC
#> 4 motif 1 AT1G03850 203 209 9.08 CTAATCC
#> 5 motif 1 AT5G01810 821 827 9.08 CATATCC
#> 6 motif 1 AT5G01810 840 846 9.08 CAAATCC
#>   thresh.score min.score max.score score.pct strand
#>   <numeric> <numeric> <numeric> <numeric> <character>
```

```
#> 1      8.172 -19.649  9.08  100      +
#> 2      8.172 -19.649  9.08  100      +
#> 3      8.172 -19.649  9.08  100      +
#> 4      8.172 -19.649  9.08  100      +
#> 5      8.172 -19.649  9.08  100      +
#> 6      8.172 -19.649  9.08  100      +
```

Using 2-letter information to scan:

```
head(scan_sequences(motif.k2, ArabidopsisPromoters, use.freq = 2, RC = TRUE,
                    threshold = 0.9, threshold.type = "logodds"))
#> Note: detected non-finite positional weights, normalising
#> Note: since this motif has a pseudocount of 0, 1 will be used
#> DataFrame with 6 rows and 12 columns
#>      motif motif.i sequence start stop score match
#>   <character> <integer> <character> <integer> <integer> <numeric> <character>
#> 1      motif         1 AT4G12690     938    943  17.827 CAAAAC
#> 2      motif         1 AT2G37950     751    756  17.827 CAAAAC
#> 3      motif         1 AT1G49840     959    964  17.827 CTTTTC
#> 4      motif         1 AT1G77210     184    189  17.827 CAAAAC
#> 5      motif         1 AT1G77210     954    959  17.827 CAAAAC
#> 6      motif         1 AT3G57640     917    922  17.827 CTTTTC
#>   thresh.score min.score max.score score.pct strand
#>   <numeric> <numeric> <numeric> <numeric> <character>
#> 1      16.0443  -16.842   17.827     100      +
#> 2      16.0443  -16.842   17.827     100      +
#> 3      16.0443  -16.842   17.827     100      +
#> 4      16.0443  -16.842   17.827     100      +
#> 5      16.0443  -16.842   17.827     100      +
#> 6      16.0443  -16.842   17.827     100      +
```

As an aside: the previous example involved calling `create_motif()` and `add_multifreq()` separately. In this case however this could have been simplified to just calling `create_motif()` and using the `add.multifreq` option:

```
library(universalmotif)
library(Biostrings)

sequences <- DNASTringSet(rep(c("CAAAACC", "CTTTTCC"), 3))
motif <- create_motif(sequences, add.multifreq = 2:3)
```

7 Enrichment analyses

The `universalmotif` package offers the ability to search for enriched motif sites in a set of sequences via `enrich_motifs()`. There is little complexity to this, as it simply runs `scan_sequences()` twice: once on a set of target sequences, and once on a set of background sequences. After which the results between the two sequences are collated and run through enrichment tests. The background sequences can be given explicitly, or else `enrich_motifs()` will create background sequences on its own by using `shuffle_sequences()` on the target sequences.

Let us consider the following basic example:

```
library(universalmotif)
data(ArabidopsisMotif)
data(ArabidopsisPromoters)
```



```

enrich_motifs(ArabidopsisMotif, ArabidopsisPromoters, shuffle.k = 3,
              threshold = 0.001, RC = TRUE)
#> DataFrame with 1 row and 11 columns
#>      motif motif.i target.hits target.seq.hits target.seq.count
#>      <character> <integer> <integer> <integer> <integer>
#> 1 YTTYTTTTTTYTTY 1      681      50      50
#>      bkg.hits bkg.seq.hits bkg.seq.count Pval Qual Eval
#>      <integer> <integer> <integer> <numeric> <numeric> <numeric>
#> 1      290      48      50 2.17634e-37 2.17634e-37 4.35268e-37

```

Here we can see that the motif is significantly enriched in the target sequences. The Pval was calculated by calling `stats::fisher.test()`.

One final point: always keep in mind the `threshold` parameter, as this will ultimately decide the number of hits found. (A bad threshold can lead to a false negative.)

8 Testing for motif positional preferences in sequences

The `universalmotif` package provides the `motif_peaks()` function, which can test for positionally preferential motif sites in a set of sequences. This can be useful, for example, when trying to determine whether a certain transcription factor binding site is more often than not located at a certain distance from the transcription start site (TSS). The `motif_peaks()` function finds density peaks in the input data, then creates a null distribution from randomly generated peaks to calculate peak P-values.

```

library(universalmotif)
data(ArabidopsisMotif)
data(ArabidopsisPromoters)

hits <- scan_sequences(ArabidopsisMotif, ArabidopsisPromoters, RC = FALSE)

res <- motif_peaks(hits$start,
                   seq.length = unique(width(ArabidopsisPromoters)),
                   seq.count = length(ArabidopsisPromoters))


## Significant peaks:
res$Peaks
#> DataFrame with 1 row and 2 columns
#>      Peak      Pval
#>      <numeric> <numeric>
#> 1      892 4.26875e-14

```

Using the datasets provided in this package, a significant motif peak was found about 100 bases away from the TSS. If you'd like to simply know the locations of any peaks, this can be done by setting `max.p = 1`.

The function can also output a plot:

```
res$Plot
```



In this plot, red dots are used to indicate density peaks and the blue line shows the P-value cutoff.

9 Motif discovery with MEME

The `universalmotif` package provides a simple wrapper to the powerful motif discovery tool **MEME** (Bailey and Elkan 1994). To run an analysis with **MEME**, all that is required is a set of `XStringSet` class sequences (defined in the `Biostrings` package), and `run_meme()` will take care of running the program and reading the output for use within R.

The first step is to check that R can find the **MEME** binary in your `$PATH` by running `run_meme()` without any parameters. If successful, you should see the default **MEME** help message in your console. If not, then you'll need to provide the complete path to the **MEME** binary. There are two options:

```
library(universalmotif)

## 1. Once per session: via `options()``
options(meme.bin = "/path/to/meme/bin/meme")

run_meme(...)

## 2. Once per run: via `run_meme()``
run_meme(..., bin = "/path/to/meme/bin/meme")
```

Now we need to get some sequences to use with `run_meme()`. At this point we can read sequences from disk or extract them from one of the Bioconductor `BSgenome` packages.

```

library(universalmotif)
data(ArabidopsisPromoters)

## 1. Read sequences from disk (in fasta format):

library(Biostrings)

# The following `read*()` functions are available in Biostrings:
# DNA: readDNAStringSet
# DNA with quality scores: readQualityScaledDNAStringSet
# RNA: readRNAStringSet
# Amino acid: readAAStringSet
# Any: readBStringSet

sequences <- readDNAStringSet("/path/to/sequences.fasta")

run_meme(sequences, ...)

## 2. Extract from a `BSgenome` object:

library(GenomicFeatures)
library(TxDb.Athaliana.BioMart.plantsmart28)
library(BSgenome.Athaliana.TAIR.TAIR9)

# Let us retrieve the same promoter sequences from ArabidopsisPromoters:
gene.names <- names(ArabidopsisPromoters)

# First get the transcript coordinates from the relevant `TxDb` object:
transcripts <- transcriptsBy(TxDb.Athaliana.BioMart.plantsmart28,
                             by = "gene")[gene.names]

# There are multiple transcripts per gene, we only care for the first one
# in each:

transcripts <- lapply(transcripts, function(x) x[1])
transcripts <- unlist(GRangesList(transcripts))

# Then the actual sequences:

# Unfortunately this is a case where the chromosome names do not match
# between the two databases

seqlevels(TxDb.Athaliana.BioMart.plantsmart28)
#> [1] "1" "2" "3" "4" "5" "Mt" "Pt"
seqlevels(BSgenome.Athaliana.TAIR.TAIR9)
#> [1] "Chr1" "Chr2" "Chr3" "Chr4" "Chr5" "ChrM" "ChrC"

# So we must first rename the chromosomes in `transcripts`:
seqlevels(transcripts) <- seqlevels(BSgenome.Athaliana.TAIR.TAIR9)

# Finally we can extract the sequences
promoters <- getPromoterSeq(transcripts,
                             BSgenome.Athaliana.TAIR.TAIR9,

```

```
upstream = 1000, downstream = 0)
```

```
run_meme(promoters, ...)
```

Once the sequences are ready, there are few important options to keep in mind. One is whether to conserve the output from MEME. The default is not to, but this can be changed by setting the relevant option:

```
run_meme(sequences, output = "/path/to/desired/output/folder")
```

The second important option is the search function (`objfun`). Some search functions such as the default `classic` do not require a set of background sequences, whilst some do (such as `de`). If you choose one of the latter, then you can either let MEME create them for you (it will shuffle the target sequences) or you can provide them via the `control.sequences` parameter.

Finally, choose how you'd like the data imported into R. Once the MEME program exits, `run_meme()` will import the results into R with `read_meme()`. At this point you can decide if you want just the motifs themselves (`readsites = FALSE`) or if you'd like the original sequence sites as well (`readsites = TRUE`, the default). Doing the latter gives you the option of generating higher order representations for the imported MEME motifs as shown here:

```
motifs <- run_meme(sequences)
motifs.k23 <- mapply(add_multifreq, motifs$motifs, motifs$sites)
```

There are a wealth of other MEME options available, such as the number of desired motifs (`nmotifs`), the width of desired motifs (`minw`, `maxw`), the search mode (`mod`), assigning sequence weights (`weights`), using a custom alphabet (`alph`), and many others. See the output from `run_meme()` for a brief description of the options, or visit the online manual for more details.

Session info

```
#> R version 4.0.0 beta (2020-04-14 r78227)
#> Platform: x86_64-w64-mingw32/x64 (64-bit)
#> Running under: Windows Server 2012 R2 x64 (build 9600)
#>
#> Matrix products: default
#>
#> locale:
#> [1] LC_COLLATE=C
#> [2] LC_CTYPE=English_United States.1252
#> [3] LC_MONETARY=English_United States.1252
#> [4] LC_NUMERIC=C
#> [5] LC_TIME=English_United States.1252
#>
#> attached base packages:
#> [1] stats4      parallel  stats      graphics  grDevices  utils      datasets
#> [8] methods    base
#>
#> other attached packages:
#> [1] TFBSTools_1.25.2      Logolas_1.11.0        dplyr_0.8.5
#> [4] ggtree_2.1.6          ggplot2_3.3.0         MotifDb_1.29.11
#> [7] GenomicRanges_1.39.3  GenomeInfoDb_1.23.17  Biostrings_2.55.7
#> [10] XVector_0.27.2        IRanges_2.21.8        S4Vectors_0.25.15
#> [13] BiocGenerics_0.33.3   universalmotif_1.5.36
#>
#> loaded via a namespace (and not attached):
```

```

#> [1] VGAM_1.1-2 colorspace_1.4-1
#> [3] grImport2_0.2-0 ellipsis_0.3.0
#> [5] base64enc_0.1-3 applot_0.0.4
#> [7] rGADEM_2.35.1 farver_2.0.3
#> [9] bit64_0.9-7 AnnotationDbi_1.49.1
#> [11] fansi_0.4.1 motifStack_1.31.2
#> [13] splines_4.0.0 R.methodsS3_1.8.0
#> [15] knitr_1.28 ade4_1.7-15
#> [17] jsonlite_1.6.1 splitstackshape_1.4.8
#> [19] Rsamtools_2.3.7 seqLogo_1.53.0
#> [21] gridBase_0.4-7 annotate_1.65.1
#> [23] GO.db_3.10.0 png_0.1-7
#> [25] R.oo_1.23.0 httr_1.4.1
#> [27] BiocManager_1.30.10 readr_1.3.1
#> [29] compiler_4.0.0 rvcheck_0.1.8
#> [31] assertthat_0.2.1 Matrix_1.2-18
#> [33] lazyeval_0.2.2 cli_2.0.2
#> [35] htmltools_0.4.0 tools_4.0.0
#> [37] gtable_0.3.0 glue_1.4.0
#> [39] TFMpvalue_0.0.8 GenomeInfoDbData_1.2.3
#> [41] reshape2_1.4.4 tinytex_0.21
#> [43] Rcpp_1.0.4.6 Biobase_2.47.3
#> [45] vctrs_0.2.4 ape_5.3
#> [47] nlme_3.1-147 rtracklayer_1.47.0
#> [49] ggseqlogo_0.1 gbrd_0.4-11
#> [51] xfun_0.13 CNEr_1.23.0
#> [53] stringr_1.4.0 ps_1.3.2
#> [55] lifecycle_0.2.0 powerLaw_0.70.4
#> [57] gtools_3.8.2 XML_3.99-0.3
#> [59] zlibbioc_1.33.1 MASS_7.3-51.5
#> [61] scales_1.1.0 BSgenome_1.55.4
#> [63] hms_0.5.3 SummarizedExperiment_1.17.5
#> [65] RColorBrewer_1.1-2 yaml_2.2.1
#> [67] memoise_1.1.0 MotIV_1.43.0
#> [69] stringi_1.4.6 RSQLite_2.2.0
#> [71] SQUAREM_2020.2 highr_0.8
#> [73] tidytree_0.3.3 caTools_1.18.0
#> [75] BiocParallel_1.21.2 bibtex_0.4.2.2
#> [77] Rdpack_0.11-1 rlang_0.4.5
#> [79] pkgconfig_2.0.3 matrixStats_0.56.0
#> [81] bitops_1.0-6 evaluate_0.14
#> [83] lattice_0.20-41 purrr_0.3.3
#> [85] htmlwidgets_1.5.1 GenomicAlignments_1.23.2
#> [87] treeio_1.11.3 patchwork_1.0.0
#> [89] labeling_0.3 bit_1.1-15.2
#> [91] processx_3.4.2 tidyselect_1.0.0
#> [93] plyr_1.8.6 magrittr_1.5
#> [95] bookdown_0.18 R6_2.4.1
#> [97] DelayedArray_0.13.12 DBI_1.1.0
#> [99] pillar_1.4.3 withr_2.1.2
#> [101] KEGGREST_1.27.1 RCurl_1.98-1.1
#> [103] tibble_3.0.0 crayon_1.3.4
#> [105] rmarkdown_2.1 jpeg_0.1-8.1
#> [107] grid_4.0.0 data.table_1.12.8

```

```
#> [109] blob_1.2.1          digest_0.6.25
#> [111] xtable_1.8-4        tidyr_1.0.2
#> [113] R.utils_2.9.2       munsell_0.5.0
#> [115] DirichletMultinomial_1.29.0
```

References

- Altschul, Stephen F., and Bruce W. Erickson. 1985. “Significance of Nucleotide Sequence Alignments: A Method for Random Sequence Permutation That Preserves Dinucleotide and Codon Usage.” *Molecular Biology and Evolution* 2 (6): 526–38.
- Bailey, T. L., and C. Elkan. 1994. “Fitting a Mixture Model by Expectation Maximization to Discover Motifs in Biopolymers.” *Proceedings of the Second International Conference on Intelligent Systems for Molecular Biology* 2: 28–36.
- Fitch, Walter M. 1983. “Random Sequences.” *Journal of Molecular Biology* 163 (2): 171–76.
- Jiang, M., J. Anderson, J. Gillespie, and M. Mayne. 2008. “uShuffle: A Useful Tool for Shuffling Biological Sequences While Preserving K-Let Counts.” *BMC Bioinformatics* 9 (192).
- Propp, J. G., and D. W. Wilson. 1998. “How to Get a Perfectly Random Sample from a Generic Markov Chain and Generate a Random Spanning Tree of a Directed Graph.” *Journal of Algorithms* 27: 170–217.