

Detecting Obscure Tandem Repeats in Sequences

Erik S. Wright

April 15, 2025

Contents

1	Introduction	1
2	Getting Started	3
2.1	Startup	3
3	Taking a first look at Obscurin	3
4	Detecting repeats automatically with DetectRepeats	6
5	Inferring the History of Repeats in Obscurin	8
6	Session Information	10

1 Introduction

Jacques Monod famously penned an essay in 1977 likening the evolutionary process to tinkering and boldly claiming that “evolution does not produce novelties from scratch.” He cites molecular evidence:

This is exemplified by the finding that large segments of genetic information, that is, of DNA, turn out to be homologous, not only in the same organism, but also among different organisms, even among those that are phylogenetically distant. Similarly, as more is known about amino acid sequences in proteins, it appears not only that proteins fulfilling similar functions in different organisms have frequently similar sequences, but also that proteins with different functions also exhibit rather large segments in common.

We now know that proteins can, and regularly do, evolve *de novo*. Yet, repetition of existing sequence, followed by divergence (i.e., tinkering), is still a major source of evolutionary innovation. DNA slippage during replication is one mechanism that can result in repeated sequences, in particular adjacent repeats in the DNA. These tandem repeats, especially in the form of short microsatellites, underly several human diseases including fragile X syndrome and Huntington’s disease.

Microsatellites are easy to spot by eye, as they contain repetitive stretches of a few nucleotides repeated many times in a row (e.g., CAG CAG CAG CAG CAG [...] in the protein Huntingtin). Harder to detect are tandem repeats that have evolved considerably since their origin, such that the repeated copies no longer look anything like their shared ancestral sequence. It is estimated that over half of the human proteome contains tandem repeats, many of which share an ancient origin and are undetectable by eye.



Figure 1: Partial predicted structure of the human protein Obscurin showing the repeated Ig domains.

Obscurin is one human protein containing several repeats that are difficult to spot by eye. It contains 68 Ig (Immunoglobulin) domains, which each contain multiple antiparallel β -sheets (Fig. 1). Many of these Ig domains are found adjacent to each other and likely arose from a common ancestor. This vignette describes how to detect these obscure repeats using the `DetectRepeats` function in the R package DECIPHER.

2 Getting Started

2.1 Startup

To get started we need to load the DECIPHER package, which automatically loads a few other required packages.

```
> library(DECIPHER)
```

Detecting repeats is performed with the `DetectRepeats` function. Its help page can be accessed through:

```
> ? DetectRepeats
```

Once DECIPHER is installed, the code in this tutorial can be obtained via:

```
> browseVignettes("DECIPHER")
```

3 Taking a first look at Obscurin

Here, we will take a deep dive into the human protein Obscurin, which is expressed in muscle tissue. Obscurin is a large protein primarily composed of Ig domains, an example of which looks like:

```
> Ig <- AAStringSet("VRALPARFTEGLRNEEAMEGATATLQCELSKAAPVEWRKGLEALRDGDKYSLRQDGAVCELQIHGLA")
> Ig
AAStringSet object of length 1:
      width seq
[1]      88 VRALPARFTEGLRNEEAMEGATATLQCELSKAAP...GAVCELQIHGLAMADNGVYSCVCGQERTSATLT
```

Since we know the form of the repeat, we can find others like it in the sequence by aligning the repeat to a sliding window along the entire protein sequence. First, we load our sequences with:

```
> # specify the path to your file of sequences:
> fas <- "<<path to training FASTA file>>"
> # OR use the example DNA sequences:
> fas <- system.file("extdata",
                     "LongHumanProteins.fas.gz",
                     package="DECIPHER")
> # read the sequences into memory
> DNA <- readDNAStringSet(fas)
> DNA
DNAStringSet object of length 11:
      width seq                                     names
[1] 107976 ATGACAACCTCAAGCACCGACGT...TATACATATTCGATCCATTTAA titin isoform IC
[2]  26778 ATGGATCAGCCACAGTTCAGCG...CAACCTGGCCCAGGTGCGCTGA obscurin isoform c
[3] 16908  ATGATTTTCCTGGGAAGTTGTCC...TGTGTCCGCCTATCCATACTAA hemicentin-1 prec...
[4]  5847  ATGGCGCCACCTGGGGCCCTG...CTTTGACCACTATGCAACCTAA receptor-type tyr...
[5]  5745  ATGGGGGATGTGAAGCTGTTG...AGGGGAAGAGGAAGAAGAGTGA myosin light chai...
...    ...
[7]  3825  ATGCCTGAGCCGGGGAAGAAGC...GGAGGTGCGAGTGCCTCAGTGA myosin-binding pr...
[8]  1497  ATGTTTAACTACGAACGTCCAA...CTATGAAAGTGAAGAAGTTTAA myotilin isoform a
[9]  1263  ATGCTCAAAGCGGTGATCCTGA...CACCAACCAGATCATCCTCTGA mannose-1-phospha...
[10] 1101  ATGGTGCTGCACCTACTGCTCT...GCAGCACTGTGCTTGTATCTAA inhibin alpha cha...
[11]   717  ATGTCTAAAGCAGCTCCTGCCA...CGCTCGCGTCGACGTGGCCTGA SPEG neighbor pro...
```

Although `DetectRepeats` works on nucleotide and protein sequences, it is generally easier to find tandem repeats in amino acid sequences than it is to find them in DNA sequences. We can easily translate our proteins:

```
> AA <- translate(DNA)
> AA
AAStringSet object of length 11:
      width seq
[1] 35992 MTTQAPTFTQPLQSVVVLEGSTA...LSLGNFEGSDSATVNIHIRSI* titin isoform IC
[2]  8926 MDQPQFSGAPRFLTRPKAFVVS...VRNREKRRALLYKRHNLAQVR* obscurin isoform c
[3]  5636 MISWEVVHTVFLFALLYSSLAQD...NGTIEYQTTFIVYIAVSAYPY* hemicentin-1 prec...
[4]  1949 MAPTWGPGMVSVVGPMLLVLL...EYQFCYQAALEYLGSFDHYAT* receptor-type tyr...
[5]  1915 MGDVKLVASSHISKTSLSDPSR...TAELIVETMEEGEGEGE***** myosin light chai...
...
[7]  1275 MPEPGKKPVSAFSKKPRSVVEAA...RATNLQGEARCECRLEVRVPQ* myosin-binding pr...
[8]   499 MFNYERPKEHFIQSQNPCGSRLQP...NPEGEFQRLAAQSGLYESEEL* myotilin isoform a
[9]   421 MLKAVILIGGPQKGTRFRPLSFE...LNSIVLPHEKLSRSFTNQIIL* mannose-1-phospha...
[10]  367 MVLHLLLFLLLTPQGGHSCQGLE...GGYSFKYETVPNLLTQHCACI* inhibin alpha cha...
[11]  239 MSKAAPAKKPVAVAPAGCTLDI...YEVYVENS LGMDQSFARVDVA* SPEG neighbor pro...
> names(AA)
[1] "titin isoform IC"
[2] "obscurin isoform c"
[3] "hemicentin-1 precursor"
[4] "receptor-type tyrosine-protein phosphatase S isoform 1 precursor"
[5] "myosin light chain kinase, smooth muscle isoform 1"
[6] "roundabout homolog 1 isoform a precursor"
[7] "myosin-binding protein C, cardiac-type"
[8] "myotilin isoform a"
[9] "mannose-1-phosphate guanyltrtransferase alpha"
[10] "inhibin alpha chain preproprotein"
[11] "SPEG neighbor protein"
> index <- 2
> AA <- AA[index]
```

We see that *Obscurin* is the second protein in the *AAStringSet*. The others are related human proteins containing Ig domains. It is possible to analyze many proteins at the same time, but this vignette focuses on only a single protein so the examples run quickly. To detect repeats in a different protein, or multiple proteins, simply change the value of `index` above.

We can find the repeated sequence using pairwise alignment by aligning each of the windows along the sequence to the Ig domain and plot their scores. The alignment score quantifies how well each window along the protein matches the Ig domain. Windows matching Ig domains will show up as peaks in the score.

We can make a couple of observations from the plot. First, many of the Ig domains are located adjacent to each other, likely as a result of tandem repeats, although some are isolated to other parts of the protein. Second, the tandem repeats around the main peak have higher scores, suggesting these orthologous repeats share a more recent evolutionary ancestor. Many of the other repeats have negative scores even though they are also Ig domains.

```

> windows <- extractAt(AA[[1]],
  IRanges(seq_len(width(AA[1]) - width(Ig)),
    width=width(Ig)))
> p <- AlignPairs(windows,
  Ig,
  pairs=data.frame(Pattern=seq_along(windows), Subject=1),
  verbose=FALSE)
> plot(p$Score, ylab="Pairwise alignment score", xlab="Amino acid position", type="l")

```



Figure 2: Peaks marking the start of Ig domains in the human protein Obscurin.

4 Detecting repeats automatically with DetectRepeats

With Obscurin we have the advantage that we already know to look for repeated Ig domains. But how can we find tandem repeats when we don't know what to expect? This is when it comes in handy to use the `DetectRepeats` function to find tandem repeats. It takes sequences as input and then outputs the predicted tandem repeats and their associated scores. Let's take a look at how it does with Obscurin:

```
> reps <- DetectRepeats(AA, allScores=TRUE)
=====

Time difference of 39.44 secs
> reps[which.max(reps[, "Score"]),]
   Index Begin  End      Left      Right  Score
14      1      5 5657 5, 137, .... 136, 268.... 503.8159
```

The result is a *data.frame* giving the beginning and ending positions of the tandem repeat, as well as every repeat within it. The top scoring repeat spans the Ig domains we found earlier (i.e., from `Begin` to `End`). The `Left` and `Right` boundaries of each repeat unit are listed, and well as the `Index` of the sequence containing the repeat.

Note that `DetectRepeats` scores are calibrated such that repeats are expected to be found by chance every e^{Score} positions. That is, approximately every 3,000 positions would be within a repeat meeting the default *minScore* of 8 by chance. Since we asked for *allScores*, the output contains all repeats with scores above *minScore*, even if they are overlapping. This allows us to easily plot the different repeats that were identified in Obscurin:

Here, the repeats are shown as separate colored line segments within each identified repeat. All of the adjacent Ig domains were identified multiple times, often with different beginning and ending positions since the start and end boundaries of each unit are difficult to determine. If we set *allScores* to `FALSE` (the default) then we would only obtain the top scoring repeat in each region.

```

> plot(NA,
      xlim=c(0, max(width(AA))),
      ylim=range(0, reps[, "Score"]),
      xlab="Position in protein",
      ylab="Tandem repeat score")
> for (i in seq_len(nrow(reps)))
  segments(reps[[i, "Left"]],
          reps[i, "Score"],
          reps[[i, "Right"]],
          reps[i, "Score"],
          col=seq_along(reps[[i, "Left"]]))

```

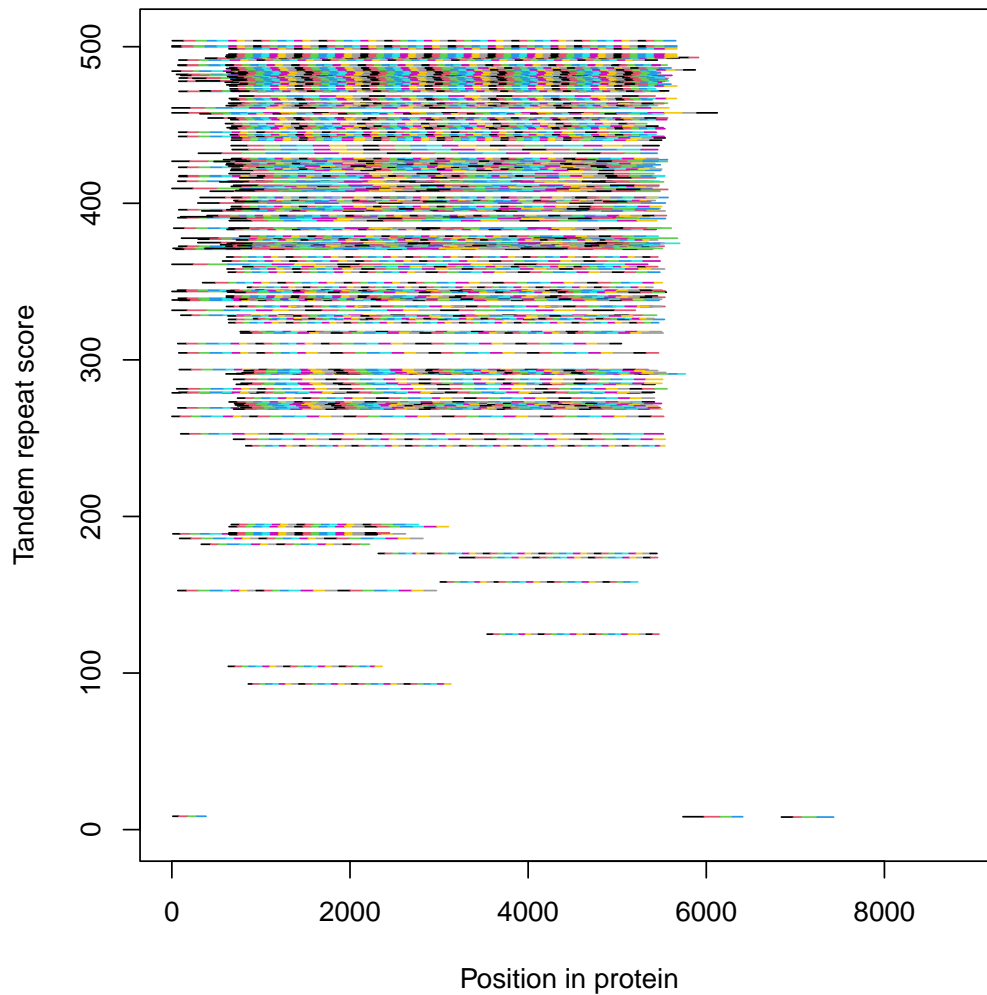


Figure 3: The different repeats detected in the human protein Obscurin.

5 Inferring the History of Repeats in Obscurin

It is feasible to infer the order of repeats from a phylogeny built from the repeat sequences. Repeats that duplicated more recently will presumably be more closely related on a phylogenetic tree. To construct a tree, first we must extract the repeats.

```
> i <- which.max(reps[, "Score"])
> seqs <- extractAt(AA[[reps[i, "Index"]]],
  IRanges(reps[[i, "Left"]], reps[[i, "Right"]]))
> seqs <- AlignSeqs(seqs, verbose=FALSE) # align the repeats
> names(seqs) <- seq_along(seqs) # number from left to right
> seqs
AAStringSet object of length 60:
      width seq                                     names
[1]    279 -----...----- 1
[2]    279 -----...-----E 2
[3]    279 -----...GIEEEGTERRLTVRNVSADDDA 3
[4]    279 VYIC-ETPEGSRTVAELAVQGNL...----- 4
[5]    279 -----...----- 5
...    ...    ...
[56]    279 -----...----- 56
[57]    279 -----...----- 57
[58]    279 -----...----- 58
[59]    279 -----...----- 59
[60]    279 WRLCHELVPGPECVVDGLAPGET...----- 60
```

Then we can infer a phylogenetic tree using the `Treeline` function in `DECIPHER`. Note that the repeats are numbered from leftmost to rightmost in the tandem repeat. There is a signal of adjacent repeats branching later from each other in the phylogeny. From this it is possible to hypothesize the set of duplication events that resulted in the observed set and locations of repeats in *Obscurin*.


```

> dend <- Treeline(seqs, verbose=FALSE)
> dend <- reorder(dend, seq_along(seqs))
> layout(matrix(1:2))
> plot(dend)
> plot(unlist(dend), xlab="Position on tree", ylab="Repeat number")

```

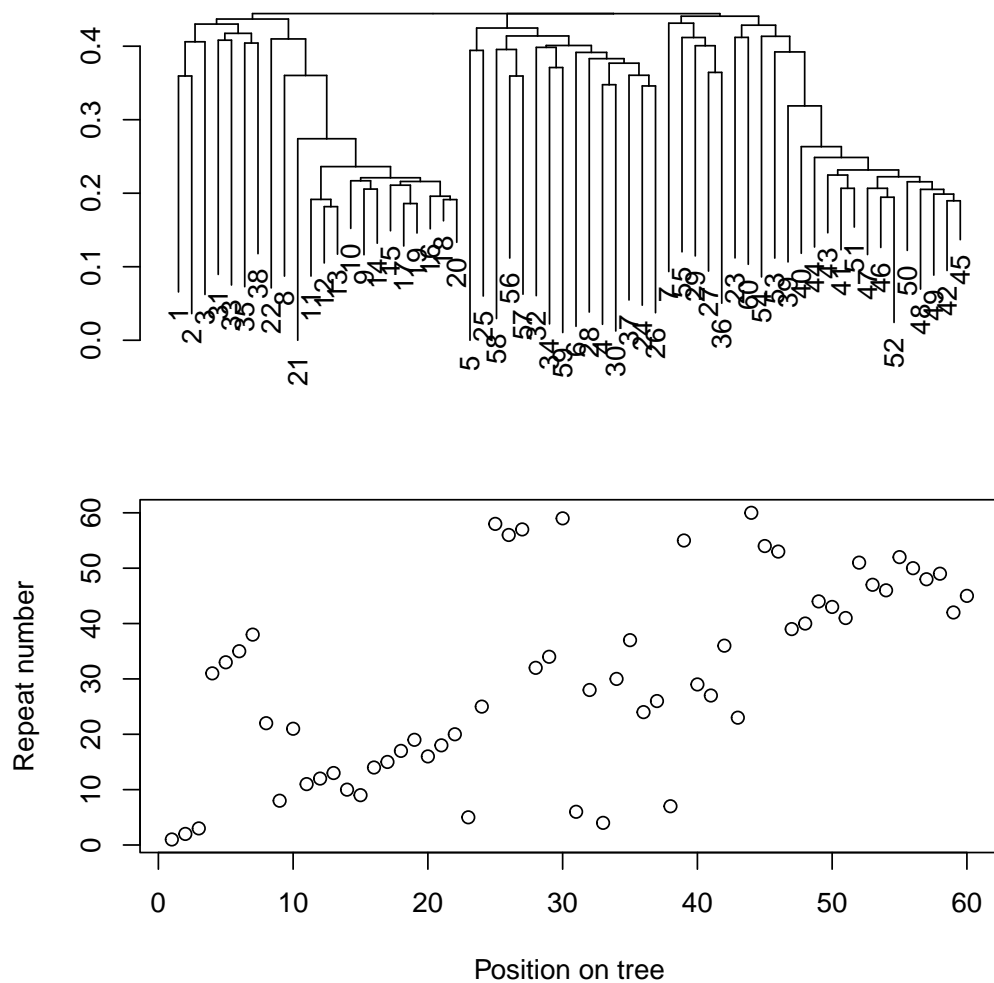


Figure 4: Inferred phylogeny of Ig domain repeats detected in Obscurin.

6 Session Information

All of the output in this vignette was produced under the following conditions:

- R version 4.5.0 RC (2025-04-04 r88126), aarch64-apple-darwin20
- Running under: macOS Ventura 13.7.1
- Matrix products: default
- BLAS:
/Library/Frameworks/R.framework/Versions/4.5-arm64/Resources/lib/libRblas.0.dylib
- LAPACK:
/Library/Frameworks/R.framework/Versions/4.5-arm64/Resources/lib/libRlapack.dylib
; LAPACK version 3.12.1
- Base packages: base, datasets, grDevices, graphics, methods, stats, stats4, utils
- Other packages: BiocGenerics 0.54.0, Biostrings 2.76.0, DECIPHER 3.4.0, GenomeInfoDb 1.44.0, IRanges 2.42.0, S4Vectors 0.46.0, XVector 0.48.0, generics 0.1.3
- Loaded via a namespace (and not attached): DBI 1.2.3, GenomeInfoDbData 1.2.14, KernSmooth 2.23-26, R6 2.6.1, UCSC.utils 1.4.0, compiler 4.5.0, crayon 1.5.3, httr 1.4.7, jsonlite 2.0.0, tools 4.5.0