

# Sequence Alignment of Short Read Data using Biostrings

Patrick Aboyoun  
Fred Hutchinson Cancer Research Center  
Seattle, WA 98008

27 July 2009

## Contents

1	Introduction	1
2	Setup	2
3	Finding Possible Contaminants in the Short Reads	3
4	Aligning Bacteriophage Reads	16
5	Session Information	18

## 1 Introduction

While most researchers use sequence alignment software like ELAND, MAQ, and Bowtie to perform the bulk of short read mappings to a target genome, BioConductor contains a number of string matching/pairwise alignment tools in the Biostrings package that can be invaluable in answering complex scientific questions. These tools are naturally divided into four groups (`matchPDict`, `vmatchPattern`, `pairwiseAlignment`, and `OTHER`) that contain the following functions:

`matchPDict` : `matchPDict`, `countPDict`, `whichPDict`, `vmatchPDict`, `vcountPDict`, `vwhichPDict`

`vmatchPattern` : `matchPattern`, `countPattern`, `vmatchPattern`, `vcountPattern`, `neditStartingAt`, `neditEndingAt`, `isMatchingStartingAt`, `isMatchingEndingAt`

`pairwiseAlignment` : `pairwiseAlignment`, `stringDist`

**OTHER** : `matchLRPatterns` (finds singleton paired-end matches), `trimLRPatterns` (trims left and/or right flanking patterns), `matchProbePair` (finds theoretical amplicons), `matchPWM` (matches using a position weight matrix)

For detailed information on any of these functions, use `help(<< function name >>)` from within R.

Of the functions listed above, the `pairwiseAlignment` function stands out for its production of the most complex output object. When producing more than just the alignment score, this output (either a *PairwiseAlignedXStringSet* or a *PairwiseAlignedFixedSubject*) can be processed by a number of helper functions including those listed in Tables 1 & 2 below.

Table 3 shows the relative strengths and weaknesses of the three main functional families and hints at how they can be used sequentially to find answers to multi-faceted questions.

Function	Description
[	Extracts the specified elements of the alignment object
alphabet	Extracts the allowable characters in the original strings
compareStrings	Creates character string mashups of the alignments
deletion	Extracts the locations of the gaps inserted into the pattern for the alignments
length	Extracts the number of patterns aligned
mismatchTable	Creates a table for the mismatching positions
nchar	Computes the length of "gapped" substrings
nedit	Computes the Levenshtein edit distance of the alignments
indel	Extracts the locations of the insertion & deletion gaps in the alignments
insertion	Extracts the locations of the gaps inserted into the subject for the alignments
nindel	Computes the number of insertions & deletions in the alignments
nmatch	Computes the number of matching characters in the alignments
nmismatch	Computes the number of mismatching characters in the alignments
pattern, subject	Extracts the aligned pattern/subject
pid	Computes the percent sequence identity
rep	Replicates the elements of the alignment object
score	Extracts the pairwise sequence alignment scores
type	Extracts the type of pairwise sequence alignment

Table 1: Functions for *PairwiseAlignedXStringSet* and *PairwiseAlignmentFixedSubject* objects.

## 2 Setup

This lab is designed as series of hands-on exercises where the students follow along with the instructor. The first exercise is to load the required packages:

### Exercise 1

Start an R session and use the `library` function to load the *ShortRead* software package and *BSgenome.Mmusculus.UCSC.mm9* genome package along with its dependencies using the following commands:

```
> suppressMessages(library("ShortRead"))
> library("BSgenome.Mmusculus.UCSC.mm9")
```

### Exercise 2

Use the `packageDescription` function to confirm that the loaded version of the *Biostrings* package is  $\geq 2.13.28$  and the *IRanges* package is  $\geq 1.3.44$ .

```
> packageDescription("Biostrings")$Version
```

```
[1] "2.13.29"
```

```
> packageDescription("IRanges")$Version
```

```
[1] "1.3.44"
```

Seek assistance from one of the course assistants if you need help updating any of your *BioConductor* packages.

This lab also requires you have access to sample data.

### Exercise 3

Copy the data from the distribution media to your local hard drive. Change the working directory in R to point to the data location.

```
> setwd(file.path("path", "to", "data"))
```

Function	Description
<code>aligned</code>	Creates an <i>XStringSet</i> containing either “filled-with-gaps” or degapped aligned strings
<code>as.character</code>	Creates a character vector version of <code>aligned</code>
<code>as.matrix</code>	Creates an “exploded” character matrix version of <code>aligned</code>
<code>consensusMatrix</code>	Computes a consensus matrix for the alignments
<code>consensusString</code>	Creates the string based on a 50% + 1 vote from the consensus matrix
<code>coverage</code>	Computes the alignment coverage along the subject
<code>mismatchSummary</code>	Summarizes the information of the <code>mismatchTable</code>
<code>summary</code>	Summarizes a pairwise sequence alignment
<code>toString</code>	Creates a concatenated string version of <code>aligned</code>
<code>Views</code>	Creates an <i>XStringViews</i> representing the aligned region along the subject

Table 2: Additional functions for *PairwiseAlignedFixedSubject* objects.

<code>matchPDict</code>	<code>vmatchPattern</code>	<code>pairwiseAlignment</code>
Utilizes a fast string matching algorithm for multiple patterns.	Uses a fast string matching algorithm for multiple subjects.	Not practical for long strings.
Finds all occurrences with up to the specified # of mismatches.	Finds all occurrences with up to the specified # of mismatches / edit distance.	Returns only one of the best scoring alignment.
Supports removal of repeat masked regions.	Supports removal of repeat masked regions.	Cannot handle masked genomes.
Produces limited output: # of times a pattern matches and where they occur.	Produces limited output: # of times a pattern matches and where they occur.	Allows various summaries of alignments.
Does not support insertions or deletions.	Supports insertions and deletions.	Supports insertions and deletions.
Uses a mismatch penalty scheme.	Uses a mismatch penalty or edit distance penalty scheme.	Provides a flexible alignment framework, including quality-based scoring.

Table 3: Comparisons of string matching/alignment methods.

### 3 Finding Possible Contaminants in the Short Reads

The raw base-called sequences that are produced by high-throughput sequencing technologies like Solexa (Illumina), 454 (Roche), SOLiD (Applied Biosystems), and Helicos tend to contain experiment-related contaminants like adapters and PCR primers as well as “phantom” sequences like poly As. Functions like `countPDict`, `vcountPattern`, and `pairwiseAlignment` from the `Biostrings` package allow for the discovery of these troublesome sequences.

These raw base-called sequences can be read with functions like the `readXStringColumns` function and processed with functions like `tables`, which find the most common sequences, from the `ShortRead` package. While this course will be using pre-processed data for this exercise, the code to find the top short reads looks something like:

```
> library(ShortRead)
> sp <- list(experiment1 = SolexaPath(file.path("path", "to", "experiment1")),
+           experiment2 = SolexaPath(file.path("path", "to", "experiment2")))
```

```

> patSeq <- paste("s_", 1:8, ".*_seq.txt", sep = "")
> names(patSeq) <- paste("lane", 1:8, sep = "")
> topReads <- lapply(structure(seq_len(length(sp)), names = names(sp)),
+   function(i) {
+     print(experimentPath(sp[[i]]))
+     do.call(SplitDataFrameList, lapply(structure(seq_len(length(patSeq)),
+       names = names(patSeq)), function(j, n = 1000) {
+         cat("Reading", patSeq[[j]], "...")
+         x <- tables(readXStringColumns(baseCallPath(sp[[i]]),
+           pattern = patSeq[[j]], colClasses = c(rep(list(NULL),
+             4), list("DNAStrng")))[[1]], n = n)[["top"]]
+         names(x) <- chartr("-", "N", names(x))
+         cat("done.\n")
+         DataFrame(read = DNAStrngSet(names(x)), count = unname(x))
+       })
+   })
+ })

```

#### Exercise 4

Use the `load` function to load the pre-processed top short reads object from the data directory into your R session.

```

> load(file.path("data", "topReads.rda"))

```

#### Exercise 5

Use the `class` function to find the class of the `topReads` object and then print out the object.

```

> class(topReads)
[1] "list"
> topReads
$experiment1
SimpleSplitDataFrameList: 8 elements
names(8): lane1 lane2 lane3 lane4 lane5 lane6 lane7 lane8
$experiment2
SimpleSplitDataFrameList: 8 elements
names(8): lane1 lane2 lane3 lane4 lane5 lane6 lane7 lane8

```

#### Exercise 6

The `topReads` object is a list of `SimpleSplitDataFrameList` objects. Extract the data for experiment 1, lane 1 to find out its content.

```

> topReads[["experiment1"]][["lane1"]]
DataFrame: 1000 rows and 2 columns
colnames(2): read count
> head(topReads[["experiment1"]][["lane1"]][["read"]])
A DNAStrngSet instance of length 6
width seq
[1] 36 AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

```

```
[2] 36 GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTGAAA
[3] 36 GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTAGAT
[4] 36 GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTTGAT
[5] 36 GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTGAT
[6] 36 GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTATAT
```

```
> head(topReads[["experiment1"]][["lane1"]][["count"]])
```

```
[1] 81237 62784 57519 16286 11849 10927
```

### Exercise 7

Extract the most common read in each of the 8 lanes for both experiments by nesting an `lapply` function call in an `sapply` function call.

```
> sapply(topReads, lapply, function(x) as.character(x[["read"]][[1]]))
```

```
      experiment1
lane1 "AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA"
lane2 "AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA"
lane3 "GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTGAAA"
lane4 "GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTGAAA"
lane5 "AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA"
lane6 "GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTGAAA"
lane7 "GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTGAAA"
lane8 "AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA"
      experiment2
lane1 "GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTGAAA"
lane2 "GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTGAAA"
lane3 "GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTGAAA"
lane4 "GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTGAAA"
lane5 "AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA"
lane6 "GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTGAAA"
lane7 "GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTGAAA"
lane8 "GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTGAAA"
```

The pre-processed data, `topReads`, loaded in the previous exercise, are in a *list* of *SimpleSplitDataFrameList* objects that represent the read and its corresponding number of occurrences. At a high level, the *list* elements represent two Solexa experiments and the *SimpleSplitDataFrameList* elements representing the 8 lanes of a Solexa run. In both of these experiments, lanes {1-4, 6-8} contain mouse-related experimental data and lane 5 contains data from bacteriophage  $\phi$ X174.

The `sapply` function call in the above example, which extracts the most prevalent sequence in each of the lanes, shows that the top read is either `GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTGAAA` or all As. Given that the former sequence is the 33 base pairs of Solexa's genomic DNA/ChIP-seq adapter plus 3 As and the latter sequence of 36 As, it would appear that As are called when there is little information about a particular base.

### Finding Poly N Sequences

When data are acquired through the `ShortRead` package, poly N sequences can be removed using the `polyn-Filter` function. Since we are operating on pre-processed data, we will have to remove poly N sequences using more rudimentary tools.

### Exercise 8

Use the following steps to find the top sequences with with at least 34 nucleotides of a single type (A, C, T, G):

1. Extract the named vector corresponding to the top sequence counts for experiment 1, lane 1.
2. Use the `alphabetFrequency` function to find the alphabet frequencies of the reads.
3. Use the parallel max, `pmax`, function to find the maximum number of occurrences for each of the four bases.
4. Create a `DNAStrngSet` whose elements contain at least 34 bases of a single type.

```
> lane1.1TopReads <- topReads[["experiment1"]][["lane1"]]
> alphabetCounts <- alphabetFrequency(lane1.1TopReads[["read"]],
+   baseOnly = TRUE)
> lane1.1MaxLetter <- pmax(alphabetCounts[, "A"], alphabetCounts[,
+   "C"], alphabetCounts[, "G"], alphabetCounts[, "T"])
> lane1.1PolySingles <- lane1.1TopReads[["read"]][lane1.1MaxLetter >=
+   34]
> length(lane1.1PolySingles)
```

```
[1] 115
```

```
> head(lane1.1PolySingles)
```

```
  A DNAStrngSet instance of length 6
  width seq
[1] 36 AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
[2] 36 CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
[3] 36 AAAAAAAAAAAAAAAAAAAAAACAAAAAAAAAAAAAAAAA
[4] 36 AAAAAAAAAAAAAAAAAAAAAACAAAAAAAAAAAAAAAAA
[5] 36 AAAAAAAAAAAAAAAAAAAAAAGAAAAAAAAAAAAAAAAA
[6] 36 AAAAAAAAAAAAAAAAAAAAAATAAAAAAAAAAAAAAAAAA
```

### Finding Adapter-Like Sequences

While the Solexa's adapter is known not to map to the mouse genome,

### Exercise 9

Show that Solexa's DNA/ChIP-seq adapter doesn't map to the mouse genome.

```
> adapter <- DNAStrng("GATCGGAAGAGCTCGTATGCCGTCTTCTGCTTG")
```

Search the *Mmusculus* genome by first setting up a `BSPParams` parameter object that utilizes the `countPattern` function and then using the `bsapply` function to loop over the chromosomes. For more information, type `help("BSPParams")` and `help("bsapply")`.

```
> bsParams <- new("BSPParams", X = Mmusculus, FUN = countPattern,
+   simplify = TRUE)
> bsapply(bsParams, pattern = adapter)
```

<code>chr1</code>	<code>chr2</code>	<code>chr3</code>	<code>chr4</code>	<code>chr5</code>	<code>chr6</code>
0	0	0	0	0	0
<code>chr7</code>	<code>chr8</code>	<code>chr9</code>	<code>chr10</code>	<code>chr11</code>	<code>chr12</code>

	0	0	0	0	0	0
chr13	chr14	chr15	chr16	chr17	chr18	
	0	0	0	0	0	0
chr19	chrX	chrY	chrM	chr1_random	chr3_random	
	0	0	0	0	0	0
chr4_random	chr5_random	chr7_random	chr8_random	chr9_random	chr13_random	
	0	0	0	0	0	0
chr16_random	chr17_random	chrX_random	chrY_random	chrUn_random		
	0	0	0	0	0	0

repeated sequencing of the adapter is a great inefficiency within an experiment. These adapter-like sequences can distort quality assurance of the Solexa data and removing them upstream can help prevent distortions in downstream QA conclusions.

### Exercise 10

Use the following steps to find the adapter-like sequences within the top reads:

1. Create a `DNAStrngSet` object containing the distinct reads by first extracting the top read sequences through nested `lapply` operations, then removing the names of the experiments using the `unname` function, then using the `unique` function to find the distinct set of reads, and then using the `sort` function to sort the sequences in alphabetical order.
2. Use the `isMatchingAt` function to find the adapter-like sequences.
3. Obtain the subset of adapter-like sequences.

```
> distinctReads <- DNAStrngSet(sort(unique(unname(unlist(lapply(topReads,
+   lapply, function(x) as.character(x[["read"]]))))))))
> whichAdapters <- isMatchingAt(adapter, distinctReads, max.mismatch = 4,
+   with.indels = TRUE)
> adapterReads <- distinctReads[whichAdapters]
> length(adapterReads)
```

```
[1] 819
```

```
> head(adapterReads)
```

```
A DNAStrngSet instance of length 6
width seq
[1] 36 AATCGGAAGAGCTCGTATGCCGTCTTCTGCTTAGAA
[2] 36 AATCGGAAGAGCTCGTATGCCGTCTTCTGCTTAGAT
[3] 36 AATCGGAAGAGCTCGTATGCCGTCTTCTGCTTATAT
[4] 36 AATCGGAAGAGCTCGTATGCCGTCTTCTGCTTGAAA
[5] 36 AATCGGAAGAGCTCGTATGCCGTCTTCTGCTTGGAT
[6] 36 AATCGGAAGAGCTCGTATGCCGTCTTCTGCTTGTA
```

As the results above show, Solexa's 33-mer adapter is closely related to 819 distinct short reads from the top reads lists.

### Exercise 11

Use the following steps to find the number of distinct adapter-like reads and the total number of these reads in each of the 8 lanes for the two experiments:

1. Use nested `lapply` function calls to extract the adapter-like sequences from each of the Solexa lanes.

2. Use nested `sapply` function calls to get the number of distinct adapter-like sequences.

3. Use nested `sapply` function calls to get the total number of adapter-like sequences.

```
> topAdapterReads <- lapply(topReads, lapply, function(x) x[x[["read"]] %in%
+   adapterReads, ])
> sapply(topAdapterReads, sapply, nrow)
```

	experiment1	experiment2
lane1	500	226
lane2	303	235
lane3	462	323
lane4	547	305
lane5	0	0
lane6	464	275
lane7	516	284
lane8	343	206

```
> sapply(topAdapterReads, sapply, function(x) sum(x[["count"]]))
```

	experiment1	experiment2
lane1	265463	158678
lane2	225519	178534
lane3	308251	303996
lane4	456932	290159
lane5	0	0
lane6	343988	255142
lane7	360014	252049
lane8	233244	177058

These adapter-like sequences are not wholly without value because they can provide some insight in where base call errors are most likely to occur for a particular sequence.

## Exercise 12

Find the distinct sequences from lane 1 of experiment 1 and their associated counts.

```
> lane1.1AdapterCounts <- topAdapterReads[["experiment1"]][["lane1"]][["count"]]
> lane1.1AdapterReads <- topAdapterReads[["experiment1"]][["lane1"]][["read"]]
> length(lane1.1AdapterReads)
```

```
[1] 500
```

```
> head(lane1.1AdapterReads)
```

```
A DNASTringSet instance of length 6
width seq
[1] 36 GATCGGAAGAGCTCGTATGCCGCTTCTGCTTGAAA
[2] 36 GATCGGAAGAGCTCGTATGCCGCTTCTGCTTAGAT
[3] 36 GATCGGAAGAGCTCGTATGCCGCTTCTGCTTTGAT
[4] 36 GATCGGAAGAGCTCGTATGCCGCTTCTGCTTGGAT
[5] 36 GATCGGAAGAGCTCGTATGCCGCTTCTGCTTATAT
[6] 36 GATCGGAAGAGCTCGTATGCCGCTTCTGCTTAGAA
```



### Exercise 13

Use the `pairwiseAlignment` function to fit the pairwise alignments of the adapter-like sequences against the adapter then summarize the results using the `summary` function.

```
> lane1.1AdapterAligns <- pairwiseAlignment(lane1.1AdapterReads,
+     adapter, type = "local-global")
> summary(lane1.1AdapterAligns, weight = lane1.1AdapterCounts)
```

Local-Global Fixed Subject Pairwise Alignment  
Number of Alignments: 265463

Scores:

Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
27.75	57.52	57.52	59.09	65.40	65.40

Number of matches:

Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
30.00	32.00	32.00	32.27	33.00	33.00

Top 10 Mismatch Counts:

	SubjectPosition	Subject	Pattern	Count	Probability
1	33	G	A	106988	0.403024150
2	33	G	T	41812	0.157505942
3	20	C	A	12558	0.047306028
4	33	G	C	7298	0.027491590
5	29	G	T	5686	0.021419181
6	20	C	N	2038	0.007677153
7	20	C	T	1996	0.007518939
8	20	C	G	1595	0.006008370
9	14	C	A	1487	0.005601534
10	14	C	T	902	0.003397837

### Finding Over-Represented Sequences

Another potential source of data contamination is over-represented sequences. These sequences can be found by clustering the short reads.

### Exercise 14

First find the distinct sequences from lane 1 of experiment 2 and their associated counts.

```
> lane2.1TopCounts <- topReads[["experiment2"]][["lane1"]][["count"]]
> lane2.1TopReads <- topReads[["experiment2"]][["lane1"]][["read"]]
> length(lane2.1TopReads)
```

```
[1] 1000
```

```
> head(lane2.1TopReads)
```

```
A DNASTringSet instance of length 6
width seq
[1] 36 GATCGGAAGAGCTCGTATGCCGCTTCTGCTTGAAA
[2] 36 GATCGGAAGAGCTCGTATGCCGCTTCTGCTTAGAT
[3] 36 AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
```



```

[1] 36 AAATGAGAAATACACACTTTAGGACGTGAAATATGG
[2] 36 AATGAGAAATACACACTTTAGGACGTGAAATATGGC
[3] 36 TGAAAATCACGAAAAATGAGAAATACACACTTTAGG
[4] 36 AGAAATACACACTTTAGGACGTGAAATATGGCGAGG
[5] 36 AATATGGCAAGAAAACTGAAAAATCATGAAAAATGAG
[6] 36 AAAATCACGAAAAATGAGAAATACACACTTTAGGAC
[7] 36 AGAAAATGAAAAATCACGAAAAATGAGAAATACACA
[8] 36 AGGACGTGGAATATGGCAAGAAAACTGAAAAATCATG
[9] 36 AAAATGAGAAATACACACTTTAGGACGTGAAATATG
... ..
[192] 36 CTGAAAAAGGTGGAAAAATTTAGAAATGTCCACTGTA
[193] 36 AATGGAAAAATGAGAAACATCCACTTGACGACTTGAA
[194] 36 GAGAGAAAACTGAAAAATCACGAAAAATGAGAAATAC
[195] 36 AAAATAATGGAAAAATGAGAAACATCCACTTGACGAC
[196] 36 AGTGAATATGGCGAGGAAAACTGAAAAAGGTGGAA
[197] 36 AAAACTGAAAAATCATGAAAAATGAGAAACATCCACT
[198] 36 GGCGAGGAAAACTGAAAAAGGTGGAAAAATTTAGAAA
[199] 36 AGAGAAACATCCACTTGACGACTTGAAAAATGACGA
[200] 36 AGGAAAAATGAGAAATACACACTTTAGGACGTGAAAT

```

```
> lane2.1TopReads[lane2.1Groups == 8]
```

```
A DNASTringSet instance of length 197
width seq
```

```

[1] 36 ACTGAAAAATCACGAAAAATGAGAAATACACACTTTA
[2] 36 AAACATCCACTTGACGACTTGAAAAATGACGAAATC
[3] 36 TAGGACGTGGAATATGGCAAGAAAACTGAAAAATCAT
[4] 36 GGAATATGGCAAGAAAACTGAAAAATCATGAAAAATG
[5] 36 GTAGGACGTGGAATATGGCAAGAAAACTGAAAAATCA
[6] 36 TGAAAAATCACGAAAAATGAGAAATACACACTTTAGG
[7] 36 CGTGAATATGGCGAGGAAAACTGAAAAAGGTGGAA
[8] 36 GAATATGGCAAGAAAACTGAAAAATCATGAAAAATGA
[9] 36 GAAAAATCACGAAAAATGAGAAATACACACTTTAGGA
... ..
[189] 36 ATCATGAAAAATGAGAAACATCCACTTGACGACTTG
[190] 36 ATTTAGAAATGTCCACTGTAGGACGTGGAATATGGC
[191] 36 TAGAAATGTCCACTGTAGGACGTGGAATATGGCAAG
[192] 36 TCCACTTGACGACTTGAAAAATGACGAAATCACTAA
[193] 36 TTAGAAATGTCCACTGTAGGACGTGGAATATGGCAA
[194] 36 AATTTAGAAATGTCCACTGTAGGACGTGGAATATGG
[195] 36 CGAAATCACTAAAAAACGTGAAAAATGAGAAATGCA
[196] 36 GAAATATGGCGAGGAAAACTGAAAAAGGTGGAAAAAT
[197] 36 TGTCCACTGTAGGACGTGGAATATGGCAAGAAAACT

```

```

> unknownSeqs <- intersect(reverseComplement(lane2.1TopReads[lane2.1Groups ==
+ 9]), lane2.1TopReads[lane2.1Groups == 8])
> length(unknownSeqs)

```

```
[1] 155
```

```
> head(unknownSeqs)
```

```

A DNASTringSet instance of length 6
width seq
[1] 36 AAATGAGAAATACACACTTTAGGACGTGAAATATGG
[2] 36 AATGAGAAATACACACTTTAGGACGTGAAATATGGC
[3] 36 TGAAAATCACGGAAAATGAGAAATACACACTTTAGG
[4] 36 AGAAATACACACTTTAGGACGTGAAATATGGCGAGG
[5] 36 AATATGGCAAGAAAATGAAAATCATGGAAAATGAG
[6] 36 AAAATCACGGAAAATGAGAAATACACACTTTAGGAC

```

### Exercise 17

Create a set of interesting sequences and associated counts based upon the intersection created above.

```

> unknownCounts <- lane2.1TopCounts[match(unknownSeqs, lane2.1TopReads)] +
+   lane2.1TopCounts[match(reverseComplement(unknownSeqs), lane2.1TopReads)]
> unknownSeqs <- unknownSeqs[order(unknownCounts, decreasing = TRUE)]
> unknownCounts <- unknownCounts[order(unknownCounts, decreasing = TRUE)]
> length(unknownCounts)

[1] 155

> head(unknownCounts)

[1] 387 375 358 357 354 345

```

These sequences of unknown origin may be related and could potential assemble into a more informative larger sequence. This assembly can be performed using functions from the Biostrings package by first finding a starter, or seeding, sequences that can be grown using pairwise alignments of the starter sequences and the remaining sequences.

### Exercise 18

Use the following step to find a starter or seed sequence to use in an assembly process by finding the distinct sequence that closest related to the set of unknown sequences:

1. Use the `stringDist` function to find the number of matches amongst the reads using an overlap alignment with a scoring scheme of `{match = 1, mismatch = -Inf, gapExtension = -Inf}` then convert the results into a *matrix* and loop over the rows to count how many times each distinct read overlap with other distinct reads at least 24 bases in the 36 bases reads.
2. Choose the distinct sequence with the most similar distinct sequences using the metric developed in the previous step.

```

> submat <- nucleotideSubstitutionMatrix(match = 1, mismatch = -Inf)
> whichStarter <- which.max(apply(as.matrix(stringDist(unknownSeqs,
+   method = "substitutionMatrix", substitutionMatrix = submat,
+   gapExtension = -Inf, type = "overlap")), 1, function(x) sum(x >=
+   24)))
> starterSeq <- unknownSeqs[[whichStarter]]
> starterSeq

```

```

36-letter "DNASTring" instance
seq: TGAAAATCACGGAAAATGAGAAATACACACTTTAGG

```

### Exercise 19

Use the `pairwiseAlignment` function to generate the pairwise alignments of all sequences against the starter sequence.

```
> starterAlign <- pairwiseAlignment(unknownSeqs, starterSeq, substitutionMatrix = submat,
+   gapExtension = -Inf, type = "overlap")
```

### Exercise 20

Assemble a sequence by using the starter sequence created above and the set of interesting sequences you found. The first step in this assembly is to create a function that generates a sequence through unanimous vote in a consensus matrix.

```
> unanimousChars <- function(x) {
+   letters <- c("A", "C", "G", "T")
+   mat <- consensusMatrix(x)[letters, , drop = FALSE]
+   paste(apply(mat, 2, function(y) {
+     z <- which(y != 0)
+     ifelse(length(z) == 1, letters[z], "?")
+   }), collapse = "")
+ }
```

### Exercise 21

The next step is to find which alignments are in the “prefix” of the starter sequence. These are the sequences that overlap to the left of the start sequence.

```
> whichInPrefix <- (score(starterAlign) >= 10 & start(subject(starterAlign)) ==
+   1 & start(pattern(starterAlign)) != 1)
> prefix <- narrow(unknownSeqs[whichInPrefix], 1, start(pattern(starterAlign[whichInPrefix])) -
+   1)
> prefix <- DNASTringSet(paste(sapply(max(nchar(prefix)) - nchar(prefix),
+   polyn, nucleotides = "-"), as.character(prefix), sep = ""))
> consensusMatrix(prefix, baseOnly = TRUE)
```

	[,1]	[,2]	[,3]	[,4]	[,5]	[,6]	[,7]	[,8]	[,9]	[,10]	[,11]	[,12]	[,13]
A	1	2	0	0	5	0	0	0	0	0	11	12	0
C	0	0	0	0	0	6	7	0	0	0	0	0	0
G	0	0	3	4	0	0	0	0	9	10	0	0	0
T	0	0	0	0	0	0	0	8	0	0	0	0	13
other	25	24	23	22	21	20	19	18	17	16	15	14	13
	[,14]	[,15]	[,16]	[,17]	[,18]	[,19]	[,20]	[,21]	[,22]	[,23]	[,24]	[,25]	
A	14	0	0	0	0	0	20	0	22	23	24	25	
C	0	0	0	0	18	0	0	0	0	0	0	0	
G	0	0	16	17	0	19	0	21	0	0	0	0	
T	0	15	0	0	0	0	0	0	0	0	0	0	
other	12	11	10	9	8	7	6	5	4	3	2	1	
	[,26]												
A	0												
C	26												
G	0												
T	0												
other	0												

```
> unanimousChars(prefix)
[1] "AAGGACCTGGAATATGGCGAGAAAAC"
```

### Exercise 22

The corresponding step is to find which alignments are in the “suffix” of the starter sequence. These are the sequences that overlap to the right of the start sequence.

```

> whichInSuffix <- (score(starterAlign) >= 10 & end(subject(starterAlign)) ==
+ 36 & end(pattern(starterAlign)) != 36)
> suffix <- narrow(unknownSeqs[whichInSuffix], end(pattern(starterAlign[whichInSuffix])) +
+ 1, 36)
> suffix <- DNASTringSet(paste(as.character(suffix), sapply(max(nchar(suffix)) -
+ nchar(suffix), polyn, nucleotides = "-"), sep = ""))
> consensusMatrix(suffix, baseOnly = TRUE)

```

	[,1]	[,2]	[,3]	[,4]	[,5]	[,6]	[,7]	[,8]	[,9]	[,10]	[,11]	[,12]	[,13]
A	26	0	0	0	0	21	20	19	0	17	0	0	0
C	0	25	0	0	0	0	0	0	0	0	0	0	0
G	0	0	24	0	22	0	0	0	0	0	0	15	14
T	0	0	0	23	0	0	0	0	18	0	16	0	0
other	0	1	2	3	4	5	6	7	8	9	10	11	12

	[,14]	[,15]	[,16]	[,17]	[,18]	[,19]	[,20]	[,21]	[,22]	[,23]	[,24]	[,25]
A	0	0	11	0	0	8	7	6	5	0	0	0
C	13	0	0	0	0	0	0	0	0	4	0	0
G	0	12	0	10	9	0	0	0	0	0	0	2
T	0	0	0	0	0	0	0	0	0	0	3	0
other	13	14	15	16	17	18	19	20	21	22	23	24

	[,26]
A	1
C	0
G	0
T	0
other	25

```

> unanimousChars(suffix)

```

```

[1] "ACGTGAAATATGGCGAGGAAAACCTGA"

```

### Exercise 23

Combine the prefix and suffix with the starter sequence.

```

> extendedUnknown <- DNASTring(paste(unanimousChars(prefix), as.character(starterSeq),
+ unanimousChars(suffix), sep = ""))
> extendedUnknown

```

```

      88-letter "DNASTring" instance
seq: AAGGACCTGGAATATGGCGAGAAAACCTGAAAATCAC...ACACTTTAGGACGTGAAATATGGCGAGGAAAACCTGA

```

### Exercise 24

Align the set of unknown sequences against the extended sequence.

```

> unknownAlign <- pairwiseAlignment(unknownSeqs, extendedUnknown,
+ substitutionMatrix = submat, gapExtension = -Inf, type = "overlap")
> table(score(unknownAlign))

```

```

 0  1  2  3  4  5  6  7  8  9 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36
12 26 26  2  1  1  1  1  1  1  2  2  2  2  2  2  2  2  2  2  2  2  2  2  2  2  53

```

### Exercise 25

Use the countPDict function within nested sapply/lapply function calls to show the number of reads that map to the unknown sequence in the 8 lanes from the 2 experiments.

```

> sapply(topReads, lapply, function(x) {
+   whichNoNs <- (alphabetFrequency(x[["read"]])[, "N"] == 0)
+   x <- x[whichNoNs, ]
+   pdict <- PDict(x[["read"]])
+   whichMapped <- (countPDict(pdict, extendedUnknown) + countPDict(pdict,
+     reverseComplement(extendedUnknown))) > 0
+   sum(x[whichMapped, "count"])
+ })

```

	experiment1	experiment2
lane1	1577	10855
lane2	4627	10482
lane3	1284	10633
lane4	2219	8400
lane5	0	0
lane6	1659	13095
lane7	1823	11099
lane8	4657	14916

### Exercise 26

Use the `countPattern` function within a `bsapply` loop to find to which chromosome the extended unknown sequence maps.

```

> params <- new("BSPParams", X = Mmusculus, FUN = countPattern,
+   simplify = TRUE)
> unknownCountPattern <- bsapply(params, pattern = extendedUnknown)
> unknownCountPattern

```

chr1	chr2	chr3	chr4	chr5	chr6
0	1	0	0	0	0
chr7	chr8	chr9	chr10	chr11	chr12
0	0	0	0	0	0
chr13	chr14	chr15	chr16	chr17	chr18
0	0	0	0	0	0
chr19	chrX	chrY	chrM	chr1_random	chr3_random
0	0	0	0	0	0
chr4_random	chr5_random	chr7_random	chr8_random	chr9_random	chr13_random
0	0	0	0	0	0
chr16_random	chr17_random	chrX_random	chrY_random	chrUn_random	
0	0	0	0	0	

### Exercise 27

Finally use the `matchPattern` function to find the exact location on chromosome that it maps to.

```

> mm9Chr2 <- Mmusculus[["chr2"]]
> mm9Ch2View <- matchPattern(extendedUnknown, mm9Chr2)
> mm9Ch2View

```

```

Views on a 181748087-letter DNASTring subject
subject: NNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNN...GGTTAGGTCTAGGGTTTGC GCCTGGATTACGGGT
views:
      start      end width
[1] 98507289 98507376    88 [AAGGACCTGGAATATGGCGAGAAA...TGAAATATGGCGAGAAAAC TGA]

```





```
5420-letter "DNAStrng" instance
seq: GAGTTTTATCGCTTCCATGACGCAGAAGTTAACT...CAGAGTTTTATCGCTTCCATGACGCAGAAGTTAACA
```

### Exercise 32

Show an aligned/unaligned breakdown of the read counts in the “Hoover” Solexa QA plot. This can be accomplished through the following steps:

1. Use the `PDict` function to create pattern dictionaries for the cleaned reads and their reversed complement.
2. Use the `countPDict` function to find which reads map at least once to the phage genome.
3. Create an indicator variable that states whether or not a distinct sequence maps to the phage genome.

```
> posPDict <- PDict(DNAStrngSet(names(cleanReadTable)), max.mismatch = 2)
> negPDict <- PDict(reverseComplement(DNAStrngSet(names(cleanReadTable))),
+   max.mismatch = 2)
> whichAlign <- rep(FALSE, length(phageReadTable))
> whichAlign[-whichNotClean] <- (countPDict(posPDict, nebPhage,
+   max.mismatch = 2) + countPDict(negPDict, nebPhage, max.mismatch = 2) >
+   0)
```

### Exercise 33

Count the number of distinct reads that map to the genome as well as the overall percentage of reads that map to the genome.

```
> table(whichAlign)

whichAlign
FALSE TRUE
312787 196626

> round(sapply(split(phageReadTable, whichAlign), sum)/sum(phageReadTable),
+   2)

FALSE TRUE
0.19 0.81
```

### Exercise 34

Create a histogram, conditioned on alignment status, that shows the “Hoover” plot mentioned in the Short-Read vignette.

```
> print(histogram(~log10(phageReadTable[phageReadTable > 1]) |
+   whichAlign[phageReadTable > 1], xlab = "log10(Read Counts)",
+   main = "Read Counts by IS(Aligned to Phage)"))
```




Figure 2: Hoover Plot Deconstructed

```
> toLatex(sessionInfo())
```

- R version 2.10.0 Under development (unstable) (2009-07-25 r48998), i386-apple-darwin9.7.0
- Locale: en\_US.UTF-8/en\_US.UTF-8/C/C/en\_US.UTF-8/en\_US.UTF-8
- Base packages: base, datasets, graphics, grDevices, methods, stats, utils
- Other packages: Biostrings 2.13.29, BSgenome 1.13.10, BSgenome.Mmusculus.UCSC.mm9 1.3.11, IRanges 1.3.44, lattice 0.17-25, ShortRead 1.3.22
- Loaded via a namespace (and not attached): Biobase 2.5.5, grid 2.10.0, hwriter 1.1, tools 2.10.0

Table 4: The output of `sessionInfo` while creating this vignette.

## 5 Session Information