

Lab exercises. Rare variant concepts and tools with Bioconductor.

VJ Carey

November 18, 2009

1 Acknowledgment

The material covered in this lab depends heavily on software designs and creations of Martin Morgan and Sean Davis; the IRanges, BSgenome and AnnotationDbi infrastructures are also critical for allowing these problems to be stated and solved concisely, so great thanks are due to Patrick Aboyoun, Marc Carlson, Mike Lawrence and Hervé Pages.

2 Resources

Using the 0.1.6 version of samtools, we created a pileup of 1000 genomes reads from NA19240's solexa image:

```
samtools pileup -cvf \  
  human_b36_female.fa NA19240.chrom17.SLX.maq.SRP000032.2009_07.bam > \  
  n240_17.pup
```

This text file is available to you for parsing as follows:

```
> library(ind1KG)  
> library(Rsamtools)  
> pup17 = gzfile(system.file("pileups/n240_17.pup.gz", package = "ind1KG"))  
> c17p.i = readPileup(pup17, variant = "indel")  
> c17p.i
```

RangedData with 22640 rows and 11 value columns across 1 space

	space	ranges	referenceBase	consensusBase	consensusQuality
	<character>	<IRanges>	<factor>	<factor>	<integer>
1	17	[55518, 55518]	A	A	98
2	17	[55994, 55994]	T	T	70

3	17	[56014, 56014]		A	R	37
4	17	[57801, 57801]		G	G	62
5	17	[59631, 59631]		C	C	72
6	17	[62489, 62489]		G	G	61
7	17	[62491, 62491]		C	C	123
8	17	[62495, 62495]		C	C	107
9	17	[62498, 62498]		A	A	106
10	17	[64559, 64559]		T	T	68

	snpQuality	maxMappingQuality	coverage	alleleOne	alleleOneSupport
	<integer>	<integer>	<integer>	<character>	<integer>
1	0	47	31	-T	2
2	0	52	20	-A	6
3	37	46	18	+G	7
4	0	57	27	+A	17
5	0	60	21	-AT	11
6	0	56	36	+AC	7
7	0	56	32	+AA	1
8	0	56	33	+AG	1
9	0	56	38	+CG	1
10	0	51	33	+C	1

	alleleTwo	alleleTwoSupport	additionalIndels
	<character>	<integer>	<integer>
1	*	29	0
2	*	14	0
3	*	11	0
4	*	10	0
5	*	10	0
6	*	29	0
7	*	31	0
8	*	32	0
9	*	37	0
10	*	31	1

...

<22630 more rows>

Information on dbSNP SNP locations is available in

```
> library(SNPlocs.Hsapiens.dbSNP.20080617)
> c6 = getSNPlocs("chr6")
> c6[1:5, ]
```

RefSNP_id	alleles_as_ambig	loc
1	6922869	Y 92596

```

2 6905277          S 92646
3 6923601          Y 92941
4 12195140         Y 93039
5 12202968         R 93040

```

Resource-oriented exercise: After reviewing the ‘Setup’ material in the next section, estimate the frequencies of dbSNP-catalogued SNP per base pair in intronic vs. exonic DNA on chromosome 17. Estimate frequencies stratified by GC content (i.e., tabulate by 0, 1, 2 bases G or C in SNP).

Special design exercise (attempt only after all other exercises below have been solved correctly): Consider alternative representations of the SNP location/value data. The allele data could be represented as a single DNAStrng, and the location information as an IRanges instance. Assess the resource consumption and query resolution performance of these representations in comparison to the existing data.frame. Consider also a representation rooted in an RDBMS such as SQLite.

3 Exercises

3.1 Check for coding indels and SNP

- Setup: According to data distributed by Shendure, NA19240 has two copies of a triple insert in the coding region of CDRT4. Verify.

```

> library(org.Hs.eg.db)
> egid = get("CDRT4", revmap(org.Hs.egSYMBOL))
> kgid = get(egid, org.Hs.egUCSCKG)
> library(GenomicFeatures)
> library(GenomicFeatures.Hsapiens.UCSC.hg18)
> library(IRanges)
> genes = geneHuman()
> cdrt4info = genes[genes$name == kgid, ]
> trloc = transcripts(cdrt4info)
> trr = ranges(trloc)[["chr17"]]
> c17p.i[which(ranges(c17p.i)[["17"]] %in% trr), ]

```

RangedData with 12 rows and 11 value columns across 1 space

	space	ranges	referenceBase	consensusBase
	<character>	<IRanges>	<factor>	<factor>
1	17	[15283313, 15283313]	G	G
2	17	[15283315, 15283315]	T	T
3	17	[15284249, 15284249]	C	C
4	17	[15286452, 15286452]	G	G
5	17	[15286466, 15286466]	A	A

```

6           17 [15288786, 15288786] |           G           G
7           17 [15290129, 15290129] |           T           T
8           17 [15295510, 15295510] |           C           C
9           17 [15304438, 15304438] |           C           C
10          17 [15304444, 15304444] |           A           A
      consensusQuality snpQuality maxMappingQuality  coverage  alleleOne
      <integer>      <integer>      <integer> <integer> <character>
1           93           0           59           38           -T
2          135           0           59           36           -A
3           42           0           55           19           +CTT
4            3           0           47           24           -A
5           90           0           52           31           *
6          115           0           54           40           *
7           88           0           58           40           -C
8           31           0           57           21           +CCATGGCT
9           86           0           58           31           +TAA
10          80           0           58           29           +TAC
      alleleOneSupport  alleleTwo alleleTwoSupport  additionalIndels
      <integer> <character>      <integer>      <integer>
1           1           *           37           0
2          17           *           19           0
3           6           *           13           0
4            1           *           23           0
5          30           -G           1           0
6          39           -TC          1           0
7          32           *           8           0
8           3           *           18           0
9          10           *           21           0
10          1           *           28           0
...
<2 more rows>

```

On the basis of current annotation, none of these indels are in exons:

```

> exloc = exons(cdrt4info)
> exr = ranges(exloc)[["chr17"]]
> c17p.i[which(ranges(c17p.i)[["17"]] %in% exr), ]

```

RangedData with 0 rows and 11 value columns across 1 space

- Exercise: Write a function with parameters identifying a RangedData instance generated from a pileup, a gene symbol, a variant type, and a specification of feature scope, that reports on the variants present in the gene. Discuss infelicities

of data structure in the code segment above that should be ameliorated to simplify solution of this exercise.

- Exercise: Whether or not you solve the previous exercise, characterize the variants in gene MYH3 for NA19240 in some concise way. It is advisable to focus on SNP; show that there are coding SNP present for this individual that are not identified in dbSNP.
- Exercise: Introduce and justify a mechanism for filtering variant reporting using quality information.

3.2 Compare Solexa calls with Sanger sequencing

- Setup: The 4 million phase II HapMap genotype calls for NA19240 are available to you in package hmyriB36. A selection confined to chromosome 6 is available in the ind1KG package.

```
> library(ind1KG)
> library(snpMatrix)
> data(yri240_6)
> yri240_6$hm
```

```
A snp.matrix with 1 rows and 265955 columns
Row name: NA19240
Col names: rs4097465 ... rs4599694
```

```
> yri240_6$supp[1:10, ]
```

	dbSNPalleles	Assignment	Chromosome	Position	Strand
rs4097465	G/T	G/T	chr6	37012	-
rs7754266	A/G	A/G	chr6	94609	+
rs9393087	C/T	C/T	chr6	94901	+
rs12192290	A/T	A/T	chr6	95272	+
rs11962658	A/C	A/C	chr6	96774	+
rs7742004	C/G	C/G	chr6	97749	+
rs2107722	G/T	G/T	chr6	98500	-
rs1929630	A/C	A/C	chr6	99536	+
rs12524398	C/G	C/G	chr6	99694	+
rs10484790	C/T	C/T	chr6	99750	-

- Exercise: Assess how many of the MAQ-based SNP calls using the chromosome 6 pileup data are found at dbSNP locations. Is the distribution of quality scores for variants identified at dbSNP locations similar to that of putatively de novo variants?

3.3 *de novo* SNPs in probes: effects on expression microarrays

Exercise: Acquire the probe sequences for the illumina Human v1 expression array, perhaps by inverting the nuids found in the lumiHumanIDMapping metadata package.

```
> library(lumiHumanIDMapping)
```

```
This is mgcv 1.5-6 . For overview type `help("mgcv-package")'.
```

```
> con = lumiHumanIDMapping_dbconn()
```

```
> dbListTables(con)
```

```
[1] "HumanHT12_V3_0_R1_11283641_A" "HumanHT12_V3_0_R2_11283641_A"
[3] "HumanRef8_V2_0_R1_11223162_A" "HumanRef8_V2_0_R2_11223162_A"
[5] "HumanRef8_V3_0_R0_11282963_A" "HumanRef8_V3_0_R1_11282963_A"
[7] "HumanRef8_V3_0_R2_11282963_A" "HumanWG6_V1"
[9] "HumanWG6_V2_0_R1_11223189_A" "HumanWG6_V2_0_R2_11223189_A"
[11] "HumanWG6_V2_11223189_B"      "HumanWG6_V3_0_R0_11282955_A"
[13] "metadata"                    "nuID_MappingInfo"
```

```
> dbGetQuery(con, "select * from HumanWG6_V1 limit 5")
```

	Search_key	Target	ProbeId	Accession	Symbol	nuID
1	PLAC3	GI_23097300-A	0002360044	NM_021936.1	PLAC3	cn0dn1Sqdb0UHE4nEY
2	COG4	GI_21070955-A	0003940446	NM_015386.1	COG4	ik1S1J.eTo60t35XQE
3	GI_4505876-A	GI_4505876-A	0006420736	NM_000445.1	PLEC1	NBHBefupq1_azWVUMA
4	PTPRD	GI_18860893-A	0002630279	NM_130393.1	PTPRD	KcS1fQzU6Ld941MSpE
5	HS6ST2	GI_27597081-A	0003120162	NM_147174.2	HS6ST2	ZeMrPvoCSjgl41LoAk

Determine the genomic positions of all probes interrogating genes on chromosome 17 using Biostrings matchPDict against the consensus genomic sequence for chromosome 17. Find all probes (on chr17) corresponding to sequence for which NA19240 is found by MAQ to harbor a variant (use the pileup noted previously). We will call these probes “associated with sequence variants”. Compute expression Z-scores for expression levels obtained for NA19240 using mean and standard deviation based on log expression for the 89 individuals in hmyriB36 excluding NA19240. Can the distribution of expression Z-scores for probes associated with sequence variants be distinguished from the distribution of expression Z-scores for probes not associated with sequence variants.

Extra credit extension: Some probes define sequence associated with splice junctions. These 50mers will not align to consensus genomic sequence, but will align once introns are removed. Can you identify probes associated with splice junctions that are also associated with sequence variants? Does the expression Z-score for splice-junction-associated probes differ in distribution from the general distribution of expression Z-scores?

Additional exercises. Retrieve the SOLiD or 454-based short read archives for NA19240 and check the consistency of conclusions obtained in prior Solexa-based exercises with results based on these platforms.

4 Session information

```
> sessionInfo()
```

```
R version 2.10.0 (2009-10-26)  
x86_64-unknown-linux-gnu
```

```
locale:
```

```
[1] LC_CTYPE=en_US.UTF-8      LC_NUMERIC=C  
[3] LC_TIME=en_US.UTF-8      LC_COLLATE=en_US.UTF-8  
[5] LC_MONETARY=C            LC_MESSAGES=en_US.UTF-8  
[7] LC_PAPER=en_US.UTF-8     LC_NAME=C  
[9] LC_ADDRESS=C             LC_TELEPHONE=C  
[11] LC_MEASUREMENT=en_US.UTF-8 LC_IDENTIFICATION=C
```

```
attached base packages:
```

```
[1] splines  tools      stats      graphics  grDevices  utils      datasets  
[8] methods  base
```

```
other attached packages:
```

```
[1] lumiHumanIDMapping_1.4.0  
[2] lumi_1.12.1  
[3] MASS_7.3-3  
[4] preprocessCore_1.8.0  
[5] mgcv_1.5-6  
[6] affy_1.24.2  
[7] annotate_1.24.0  
[8] SNPlocs.Hsapiens.dbSNP.20080617_0.99.1  
[9] org.Hs.eg.db_2.3.6  
[10] RSQLite_0.7-3  
[11] DBI_0.2-4  
[12] AnnotationDbi_1.8.1  
[13] Biobase_2.6.0  
[14] snpMatrix_1.10.2  
[15] survival_2.35-7  
[16] GenomicFeatures.Hsapiens.UCSC.hg18_0.1.0  
[17] GenomicFeatures_0.2.0  
[18] rtracklayer_1.6.0  
[19] RCurl_1.3-0  
[20] bitops_1.0-4.1  
[21] ind1KG_0.1.2  
[22] Rsamtools_0.1.19  
[23] BSgenome_1.14.1
```

[24] Biostrings_2.14.6

[25] IRanges_1.4.7

loaded via a namespace (and not attached):

[1] affyio_1.14.0 grid_2.10.0 lattice_0.17-26 nlme_3.1-96

[5] XML_2.6-0 xtable_1.5-6