Application of VAM to 10x PBMC 3k scRNA-seq data using Seurat log normalization for the MSigDB Hallmark collection.

H. Robert Frost

1 Load the VAM package

Loading VAM will also load the required packages MASS and Matrix. Seurat is referenced via suggests so must be directly loaded to enable access to Seurat functions.

```
> library(VAM)
> library(Seurat)
```

2 Load and process the 10x PBMC scRNA-seq data

This example uses the same 10x PBMC scRNA-seq data set that is used in the Seurat Guided Clustering vignette
(https://satijalab.org/seurat/v3.1/pbmc3k_tutorial.html). The Cell Ranger files for this data set can be downloaded from
https://s3-us-west-2.amazonaws.com/10x.files/samples/cell/pbmc3k/pbmc3k_filtered_gene_bc_matrices.tar.gz. This data is loaded and processed using the same Seurat logic found in the Guided Clustering vignette. In particular, the Seurat log normalization method implemented by NormalizeData() is used with variable genes determined by FindVariableFeatures(). This method for variable feature determination decomposes the measured variance for each gene into biological and technical components and provides the values of technical variance input to the VAM algorithm.

```
> # update the data.dir argument to reflect the local location of the PBMC data
> pbmc.data = Read10X(data.dir = "./filtered_gene_bc_matrices/hg19/")
> pbmc = CreateSeuratObject(counts = pbmc.data, project = "pbmc3k", min.cells = 3,
+         min.features = 200)
> pbmc[["percent.mt"]] = PercentageFeatureSet(pbmc, pattern = "^MT-")
> pbmc = subset(pbmc, subset = nFeature_RNA > 200 & nFeature_RNA < 2500 & percent.mt < 5)
> pbmc = NormalizeData(pbmc)
> pbmc = FindVariableFeatures(pbmc, selection.method = "vst", nfeatures = 2000)
> pbmc = ScaleData(pbmc, features = rownames(pbmc))
> pbmc = RunPCA(pbmc, features = VariableFeatures(object = pbmc))
> pbmc = RunUMAP(pbmc, dims = 1:10)
> pbmc = FindNeighbors(pbmc, dims = 1:10)
> pbmc = FindClusters(pbmc, resolution = 0.5)

Modularity Optimizer version 1.3.0 by Ludo Waltman and Nees Jan van Eck

Number of nodes: 2638
Number of edges: 95927

Running Louvain algorithm...
```

```
Maximum modularity in 10 random starts: 0.8728
Number of communities: 9
Elapsed time: 0 seconds
```

```
> # Assign the known cell type labels to the clusters (this follows the Seurat vingette)
> cell.types = c("Memory CD4 T",  "CD14+ Mono", "Naive CD4 T", "B", "CD8 T", "FCGR3A+ Mono",
+                         "NK", "DC", "Platelet")
> names(cell.types) = levels(pbmc)
> pbmc = RenameIdents(pbmc, cell.types)
> pbmc
```

```
An object of class Seurat
13714 features across 2638 samples within 1 assay
Active assay: RNA (13714 features, 2000 variable features)
 3 layers present: counts, data, scale.data
 2 dimensional reductions calculated: pca, umap
```

3 Load Ensembl IDs

The Ensembl IDs and gene names must be read in from the genes.tsv file and filtered to match genes left after the quality control steps performed in the prior section.

```
> feature.data = read.delim("./filtered_gene_bc_matrices/hg19/genes.tsv",
+         header = FALSE, stringsAsFactors = FALSE)
> ensembl.ids = feature.data[,1]
> gene.names = feature.data[,2]
> genes.after.QC = rownames(pbmc)
> indices.to.keep = unlist(sapply(genes.after.QC, function(x) {which(gene.names == x)[1]}))
> ensembl.ids = ensembl.ids[indices.to.keep]
> gene.names = gene.names[indices.to.keep]
```

4 Load the MSigDB Hallmark collection

The following logic loads the MSigDB Hallmark collection using the msigdbr R package. Because the MSigDB gene sets are defined in terms of Entrez gene IDs, these must be mapped to Ensembl IDs using the org.HS.eg.db R package (from Bioconductor). The data frame returned by msigdbr is then converted into a list of gene ID vectors (each list element corresponds to a gene set and is a vector of Ensembl IDs). This list is transformed into the structure required by vamForSeurat() using the createGeneSetCollection() helper function. This helper function filters out genes not also contained in the PBMC scRNA-seq data and generates a list whose elements are vectors of gene indices in the scRNA-seq data.

```
> library(msigdbr)
> library(org.Hs.eg.db)
> # Load the MSigDB Hallmark collection using the msigdbr package
> H.collection = msigdbr(category="H")
> # Get the entrez gene IDs that are mapped to an Ensembl ID
> entrez2ensembl = mappedkeys(org.Hs.egENSEMBL)
> # Convert to a list
> entrez2ensembl = as.list(org.Hs.egENSEMBL[entrez2ensembl])
> # Convert Entrez IDs to Ensembl IDs using the org.Hs.eg.db package
```

```
> # Doing this a rather simplistic (and inefficient) way for clarity
> msigdb.entrez.ids = H.collection$entrez_gene
> num.ids = length(msigdb.entrez.ids)
> msigdb.ensembl.ids = rep(NA, num.ids)
> for (i in 1:num.ids) {
+         entrez.id = msigdb.entrez.ids[i]
+         id.index = which(names(entrez2ensembl) == entrez.id)
+         if (length(id.index > 0)) {
+           # only use the first mapped ensembl id
+           msigdb.ensembl.ids[i] = entrez2ensembl[[id.index]][1]
+         }
+ }
> # Save the ensembl IDs in the data frame
> H.collection$ensembl_gene = msigdb.ensembl.ids
> # Create a gene.set.collection list of Ensembl IDs
> gene.set.names = unique(H.collection$gs_name)
> num.sets = length(gene.set.names)
> gene.set.collection = list()
> for (i in 1:num.sets) {
+         gene.set.name = gene.set.names[i]
+         gene.set.rows = which(H.collection$gs_name == gene.set.name)
+         gene.set.ensembl.ids = H.collection$ensembl_gene[gene.set.rows]
+         gene.set.collection[[i]] = gene.set.ensembl.ids
+ }
> names(gene.set.collection) = gene.set.names
> # Create the collection list required by vamForSeurat()
> gene.set.collection = createGeneSetCollection(gene.ids=ensembl.ids,
+         gene.set.collection=gene.set.collection)
> length(gene.set.collection)

[1] 50
```

5 Execute VAM method

Since the scRNA-seq data has been processed using Seurat, we execute VAM using the vamForSeurat() function. We have set return.dist=T so that the squared adjusted Mahalanobis distances will be returned in a "VAMdist" Assay.

```
> pbmc = vamForSeurat(seurat.data=pbmc,
+     gene.set.collection=gene.set.collection,
+     center=F, gamma=T, sample.cov=F, return.dist=T)
```

Look at the first few entries in the "VAMdist" and "VAMcdf" Assays.

```
> pbmc@assays$VAMdist@data[1:5,1:5]

5 x 5 sparse Matrix of class "dgCMatrix"
                            AAACATACAACCAC-1 AAACATTGAGCTAC-1 AAACATTGATCAGC-1
HALLMARK-ADIPOGENESIS              98.166655       620.737174       466.226316
HALLMARK-ALLOGRAFT-REJECTION     1068.181897      1478.494977      1022.962722
HALLMARK-ANDROGEN-RESPONSE        436.991228       309.464019       384.337446
```

```
HALLMARK-ANGIOGENESIS                5.985118         8.191464         7.099042
HALLMARK-APICAL-JUNCTION           189.136024       203.550925       260.118588
                            AAACCGTGCTTCCG-1 AAACCGTGTATGCG-1
HALLMARK-ADIPOGENESIS              281.82782        109.42689
HALLMARK-ALLOGRAFT-REJECTION      1447.04239       1434.74848
HALLMARK-ANDROGEN-RESPONSE         302.40801        589.97580
HALLMARK-ANGIOGENESIS               28.40593         52.57082
HALLMARK-APICAL-JUNCTION           233.90846        251.14443
```

```
> pbmc@assays$VAMcdf@data[1:5,1:5]
```

```
5 x 5 sparse Matrix of class "dgCMatrix"
                            AAACATACAACCAC-1 AAACATTGAGCTAC-1 AAACATTGATCAGC-1
HALLMARK-ADIPOGENESIS             0.08326084        0.9499964        0.8516699
HALLMARK-ALLOGRAFT-REJECTION      0.22272295        0.7519254        0.1726399
HALLMARK-ANDROGEN-RESPONSE        0.62849455        0.2892127        0.4949607
HALLMARK-ANGIOGENESIS             0.24851191        0.2985509        0.2747268
HALLMARK-APICAL-JUNCTION          0.38892756        0.4371007        0.6091210
                            AAACCGTGCTTCCG-1 AAACCGTGTATGCG-1
HALLMARK-ADIPOGENESIS              0.5492316        0.1058668
HALLMARK-ALLOGRAFT-REJECTION       0.7190928        0.7055768
HALLMARK-ANDROGEN-RESPONSE         0.2705019        0.8818932
HALLMARK-ANGIOGENESIS              0.5790371        0.7489537
HALLMARK-APICAL-JUNCTION           0.5335335        0.5841469
```

6 Compute DE pathways

Use the Seurat FindAllMarkers() function to identify Hallmark pathways whose VAM scores are enriched within each cell type cluster according to a Wilcoxon test.

```
> library(dplyr)
> pbmc.markers = FindAllMarkers(pbmc, assay="VAMcdf", only.pos = TRUE, logfc.threshold = 0.01)
> #pbmc.markers %>% group_by(cluster) %>% top_n(n = 3, wt = avg_log2FC)
```

7 Visualize VAM scores

Visualize VAM scores using Seurat DoHeatmap(). The default Assay must first be changed to "VAMcdf" and the slot parameter must be set to "data" in the call to DoHeatmap().

```
> library(ggplot2)
> DefaultAssay(object = pbmc) = "VAMcdf"
> top.pathways <- pbmc.markers %>% group_by(cluster) %>% top_n(n = 3, wt = avg_log2FC)
> DoHeatmap(pbmc, slot="data", features = top.pathways$gene,  size=3, label=T) + NoLegend()
```

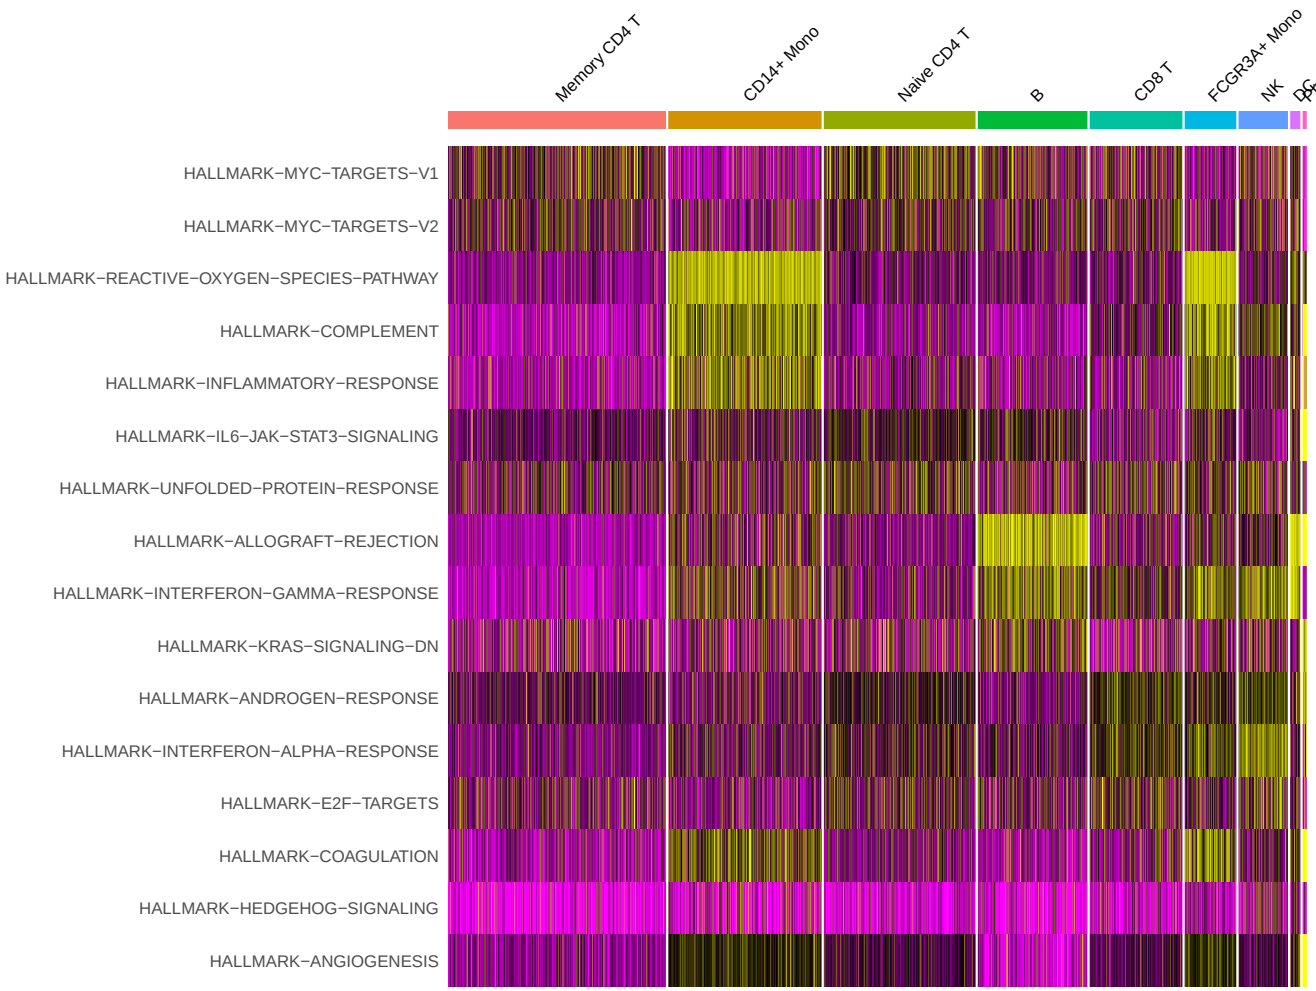
Memory CD4 T
CD14+ Mono
Naive CD4 T
B
CD8 T
FCGR3A+ Mono
NK
HALLMARK−MYC−TARGETS−V1
HALLMARK−MYC−TARGETS−V2
HALLMARK−REACTIVE−OXYGEN−SPECIES−PATHWAY
HALLMARK−COMPLEMENT
HALLMARK−INFLAMMATORY−RESPONSE
HALLMARK−IL6−JAK−STAT3−SIGNALING
HALLMARK−UNFOLDED−PROTEIN−RESPONSE
HALLMARK−ALLOGRAFT−REJECTION
HALLMARK−INTERFERON−GAMMA−RESPONSE
HALLMARK−KRAS−SIGNALING−DN
HALLMARK−ANDROGEN−RESPONSE
HALLMARK−INTERFERON−ALPHA−RESPONSE
HALLMARK−E2F−TARGETS
HALLMARK−COAGULATION
HALLMARK−HEDGEHOG−SIGNALING
HALLMARK−ANGIOGENESIS