

# Supplemental methods for: A limited set of transcriptional programs define major cell types

Alessandra Breschi<sup>\*1,2,3</sup>, Manuel Muñoz-Aguirre<sup>\*1,4</sup>, Valentin Wucher<sup>\*1</sup>, Carrie A. Davis<sup>5</sup>, Diego Garrido-Martín<sup>1,2</sup>, Sarah Djebali<sup>1,2,6</sup>, Jesse Gillis<sup>3</sup>, Dmitri D. Pervouchine<sup>1,7</sup>, Anna Vlasova<sup>8</sup>, Alexander Dobin<sup>5</sup>, Chris Zaleski<sup>5</sup>, Jorg Drenkow<sup>5</sup>, Cassidy Danyko<sup>5</sup>, Alexandra Scavelli<sup>5</sup>, Ferran Reverter<sup>1,2</sup>, Michael P. Snyder<sup>3</sup>, Thomas R. Gingeras<sup>†5</sup> and Roderic Guigó<sup>†1,2</sup>

<sup>1</sup>Centre for Genomic Regulation (CRG), The Barcelona Institute of Science and Technology, Dr. Aiguader 88, 08003 Barcelona, Catalonia, E-08003

<sup>2</sup>Universitat Pompeu Fabra (UPF), Barcelona, Catalonia, E-08003

<sup>3</sup>Stanford University, Department of Genetics, Stanford, 94305, USA

<sup>4</sup>Universitat Politècnica de Catalunya. Departament d'Estadística i Investigació Operativa. 08034 Barcelona, Catalonia, E-08003

<sup>5</sup>Cold Spring Harbor Laboratory, Cold Spring Harbor, NY 11742

<sup>6</sup>IRSD, Université de Toulouse, INSERM, INRA, ENVT, UPS, Toulouse, France.

<sup>7</sup>Skolkovo Institute for Science and Technology, 3 Nobel st., Moscow, Russia 143025

<sup>8</sup>Research Institute of Molecular Pathology (IMP), Vienna Biocenter (VBC), 1030 Vienna, Austria

---

\*A. Breschi, M. Muñoz-Aguirre and V. Wucher contributed equally to this work.

†Correspondence should be addressed to E-mail: roderic.guigo@crg.cat (Roderic Guigó) and gingeras@cshl.edu (Thomas R. Gingeras)

# Contents

|   |           |
|---|-----------|
| <b>1 Gene expression analyses</b>   | <b>1</b>  |
| 1.1 Hierarchical clustering and t-SNE . . . . .   | 1         |
| 1.2 Clustering statistics . . . . .   | 1         |
| 1.2.1 Silhouette score . . . . .  | 1         |
| 1.2.2 Elbow method . . . . .  | 1         |
| 1.3 Network modularity . . . . .  | 2         |
| 1.4 Estimation of proportions of explained variance . . . . .                               | 2         |
| 1.5 Identification of cell type specific genes . . . . .                                    | 3         |
| 1.6 Neural cell clustering . . . . .  | 3         |
| 1.6.1 Identification of neural specific genes . . . . .                                     | 3         |
| 1.6.2 Sample selection and processing . . . . .   | 3         |
| 1.6.3 Identification of neural subtype specific genes . . . . .                             | 5         |
| 1.7 Identification of cell type specific TFs . . . . .                                      | 5         |
| 1.8 Identification of driver genes through PCA . . . . .                                    | 6         |
| 1.9 Relative contribution of gene expression to variability in isoform abundances . . . . . | 6         |
| 1.10 Conservation of cell type specific genes . . . . .                                     | 7         |
| 1.10.1 Sequence conservation . . . . .  | 7         |
| 1.10.2 Expression conservation . . . . .  | 7         |
| <b>2 Analysis of Cap Analysis of Gene Expression data (CAGE)</b>                            | <b>7</b>  |
| <b>3 Analysis of DNase-seq data</b>   | <b>8</b>  |
| 3.1 DHS filtering . . . . .   | 8         |
| 3.2 t-SNE and PCA based on DHSs . . . . .   | 8         |
| 3.3 TF binding prediction in DHSs . . . . .   | 9         |
| <b>4 Analysis of Tabula Muris data</b>  | <b>9</b>  |
| <b>5 Analysis of chromatin marks</b>  | <b>11</b> |
| <b>6 Splicing</b>   | <b>11</b> |
| 6.1 Computation of exon inclusion levels . . . . .  | 11        |
| 6.2 Gene Expression vs Exon inclusion . . . . .   | 12        |

|          |   |           |
|----------|---|-----------|
| 6.3      | Cell type specific alternative splicing . . . . .                                 | 12        |
| <b>7</b> | <b>Analysis of alternative TSS usage with RAMPAGE data</b>                        | <b>13</b> |
| <b>8</b> | <b>Analysis of GTEx data</b>  | <b>14</b> |
| 8.1      | Selection of tissue-specific genes . . . . .                                      | 14        |
| 8.2      | Estimation of cellular enrichments . . . . .                                      | 14        |
| 8.3      | Classification of stomach histological samples . . . . .                          | 15        |
| 8.4      | Identification of variable genes among stomach samples . . . . .                  | 16        |
| 8.5      | Automated classification of colon samples . . . . .                               | 16        |
| 8.6      | Analysis of single cell RNA-seq data . . . . .                                    | 17        |
| 8.6.1    | Analysis of muscle single cell RNA-seq data . . . . .                             | 17        |
| 8.7      | Relationship between estimated cellular enrichments and histological phenotypes . | 17        |
| 8.7.1    | Parsing histopathological annotation . . . . .                                    | 17        |
| 8.7.2    | Testing the relationship . . . . .  | 18        |
| 8.8      | Estimation of fat proportions in muscle based on pathology annotations . . . . .  | 18        |
| 8.9      | Estimation of fat proportion in Adipose - Subcutaneous histology images . . . . . | 18        |
| <b>9</b> | <b>Analysis of cancer expression data</b>   | <b>19</b> |

## Supplemental methods

### 1 Gene expression analyses

#### 1.1 Hierarchical clustering and t-SNE

Hierarchical clustering based on gene expression (Fig. 1B) was performed on log10-transformed RPKM (with a pseudocount of 0.01) after filtering and batch correction. Complete clustering algorithm is applied to the vectors of Pearson's correlation coefficients between each pair of samples. The distance between two vectors is computed as  $\text{abs}(1-\text{cc})$ , where  $\text{cc}$  is again the Pearson's correlation coefficient between the vectors.

t-distributed stochastic neighbor embedding (t-SNE, R package `Rtsne` v0.11) is computed over the distance matrix between each pair of samples with perplexity 5, where the distance is calculated as  $\text{abs}(1-\text{cc})$  on the same log10-transformed RPKM as before.

#### 1.2 Clustering statistics

##### 1.2.1 Silhouette score

For a set of samples partitioned into  $k$  clusters, the silhouette score (Rousseeuw 1987)  $s_i$  of a sample  $i$  is calculated as:  $s_i = \frac{(b_i - a_i)}{\max(a_i, b_i)}$ , where  $a_i$  is the average distance between  $i$  and all samples of the cluster to which  $i$  belongs to, and  $b_i$  being the minimal average distance between  $i$  and all samples from the rest of the clusters to which  $i$  does not belong. The silhouette score ranges from -1 to 1, where a score of 1 means that the data is clustered accurately. The silhouette score was calculated with the `silhouette()` function from the `cluster` R package.

##### 1.2.2 Elbow method

The elbow plot (Thorndike 1953) for the number of clusters in ENCODE data (Fig. S2B) was generated as follows: first, we perform hierarchical clustering using  $1 - \text{cor}(X)$  as the distance criterion, where  $X$  is the matrix of Pearson's correlation between ENCODE primary cells (Fig. 1B). Then, for each  $k \in (2, \dots, 10)$ , we cut the hierarchical clustering tree into  $k$  clusters. For each cluster, we compute the sum of the squared differences between the samples and the corresponding centroid. We take the sum of the  $k$  scalar values to obtain the total within group sum of squares.

Both the Silhouette and the Elbow methods weakly support four over five clusters, which will result from splitting the epithelial cluster into two. In any case, these two epithelial subclusters can be grouped together into a single cluster, clearly differentiated from the rest of the clusters/types.

### 1.3 Network modularity

Modularity was computed similarly to what was described in (Breschi et al. 2016). Briefly, we built a graph where vertices (or nodes) are samples and where two vertices, samples, are connected if the Pearson's correlation coefficient between the corresponding samples, computed on the gene expression values or the xCell enrichment, is higher than a certain threshold (excluding connections of a sample with itself). The correlation thresholds have been computed independently for the RPKM and the xCell enrichment values (endothelial, epithelial, mesenchymal, neural and blood cell types), and each threshold corresponds to a specific density value, where the density is calculated as the ratio between the number of edges in the graph over the number of possible edges. The density was sampled from 20% to 45% with a step of 2.5%. Like in hierarchical clustering, gene expression values are log10-transformed RPKM after adding a pseudocount of 0.01. When computing the modularity, the membership of each node/sample is the GTEx organ. To compute the modularity we used the modularity() function from the igraph v0.7.1 R package. The modularity is plotted as a function of the density and a line was fitted using the default parameters of the geom\_smooth() function from the ggplot2 R package (Fig. S28).

### 1.4 Estimation of proportions of explained variance

To estimate the proportion of variance explained by the factors body location, germ layer and major cell type, we built a separate linear model for each gene and factor. The proportion of explained variance is usually defined as the ratio of the variance across levels of a factor over the total variance. However, to account for different number of levels in each factor, we compute  $\omega^2$ , which takes into account also the degrees of freedom of that factor:  $\omega^2 = \frac{SSQ - (k-1)MSE}{SST + MSE}$

Where  $SSQ$  is the sum of squares for a given factor,  $k - 1$  is the number of degrees of freedom,  $MSE$  is the mean squared error and  $SST$  is the total sum of squares.

## 1.5 Identification of cell type specific genes

Cell cluster specificity was surveyed with the edgeR package (Robinson et al. 2010). Since edgeR relies on a negative binomial model which requires discrete read counts, for this analysis we used the read counts of the filtered 16,265 genes, and kept the counts for individual samples without averaging the replicates. To find genes specific of each major cell type, endothelial, epithelial, melanocytes, and mesenchymal, we performed pairwise differential expression between samples of a given cluster and all the others. Genes with false discovery rate (FDR)  $<0.01$  and at least 4-fold change were considered cell-cluster-specific. The number of cell cluster-specific genes changes depending on the cluster: 627 in the endothelial cluster, 966 epithelial, 438 melanocytes and 840 mesenchymal (2,871 in total, Table S7). Gene Ontology (GO) enrichment was performed with the R package GOstat (Falcon and Gentleman 2007).

## 1.6 Neural cell clustering

### 1.6.1 Identification of neural specific genes

To identify neural specific genes we compared the expression inferred from CAGE data for 11,000 genes between neural cells, including neurons, astrocytes, and neuroepithelial cells, and all other cell types. We retained only genes with more than 50 reads in at least 3 neural samples. For consistency with the identification of the other cell type specific genes, we used the edge R package (Robinson et al. 2010) to perform pairwise differential expression between neural cells and the rest of the samples. We identified 333 neural specific genes with false discovery rate (FDR)  $<0.01$  and at least 4-fold change. Gene Ontology (GO) enrichment was performed with the R package GOstat (Falcon and Gentleman 2007).

### 1.6.2 Sample selection and processing

To perform neural cell clustering we collected publicly available data from GEO (<http://www.ncbi.nlm.nih.gov/gds>) and the ENCODE portal (<https://www.encodeproject.org/>).

Read counts for ENCODE RNA-seq data were obtained for a subset of the initial primary cells with the addition of neural cells from experiments ENCSR968WKR (bipolar neurons originated from GM23338), ENCSR244ISQ (neural progenitor cells originated from H9) and ENCSR233IJT (astrocytes). Read counts for ARPE-19 retinal pigmented epithelial cells and HFF fibroblasts were

downloaded from GEO series GSE120891 (Nogalski et al. 2019). Gene ids were mapped between Gencode version 24 to Gencode version 19 to match our annotation. The read counts were concatenated to ENCODE RNA-seq reads and normalized to RPKM as a single matrix following TMM normalization. They are collectively referred to as RNA-seq in Figure S6.

Read counts for single cell RNA-seq from adult brain were collected from GEO series GSE67835 (Darmanis et al. 2015). We relied on the original paper classification for neural cell subtype definition. Read counts for all available samples were normalized to RPKM following TMM normalization. Only postnatal non-hybrid cells were included in the analysis, and cell GSM1657976 was removed, since it was an outlier for number of expressed genes. The average RPKM within each neural subtype was used to summarize the expression values across multiple single cells, when needed. They are collectively referred to as scRNA-seq in Figure S6.

Raw fastq reads from RNA-seq of purified human brain cell populations were collected from GEO series GSE73721 (Zhang et al. 2016). Since the available processed data included only FPKM calculated from a very different pipeline, we processed the raw reads directly and quantified gene expression with the ENCODE pipeline and parameters. Similarly to the previous datasets, read counts were normalized to RPKM after TMM normalization. These samples are collectively referred to as brainRNA-seq in Figure S6.

Microarray data was obtained for a subset of samples for which DNase-seq was also available from the ENCODE portal. Raw .CEL files were downloaded from the GEO series GSE19090 (Sheffield et al. 2013). CEL files were converted from CEL format version 3 to CEL format version 4, and genes were quantified with the R packages *oligo* 1.46.0 (Carvalho and Irizarry 2010) and *pd.huex.1.0.st.v2* (Carvalho 2017) with *rma* normalization on core targets. Genes probeset ids were mapped to Gencode gene ids with Affymetrix annotation file ([http://www.affymetrix.com/Auth/analysis/downloads/na36/wtexon/HuEx-1\\_0-st-v2.na36.hg19.transcript.csv.zip](http://www.affymetrix.com/Auth/analysis/downloads/na36/wtexon/HuEx-1_0-st-v2.na36.hg19.transcript.csv.zip)). These samples are collectively referred to as Array in Figure S6.

Because the dynamic range of gene expression measured with microarrays is smaller than with RNA-seq, we applied a modified quantile normalization to enable comparison between the two platforms to generate Figure S6. RNA-seq RPKM and microarray expression values were separately binned into 1000 bins for cell type specific and neural specific genes across all cell types when available in both datasets. Then microarray values were mapped to the RPKM values of the corresponding bin in the ENCODE data. We could not apply other scaling procedures for this specific case, because other RNA-seq data, including single cell and bulk RNA-seq, did

not include any cell types other than neural cells, while the rest of RNA-seq and microarray data included endothelial, epithelial and mesenchymal cells as well. Thus, genes specific of the latter cell types, would not be assigned scaled values comparable across the different experimental settings.

Overall, we found that neural specific genes are highly expressed in all neural samples, but not in samples from other major cell types (Fig. S5A,B, Fig. S6), supporting the fact that neural cells constitute a major distinct cell type. Moreover, we found that brain endothelial cells also exhibit a neural specific profile, despite expressing endothelial specific genes as well (Fig. S6). Microglia cells, on the other hand, although still clustering with other neural cells, seem to have a less evident neural specific signature, which could be related to their different embryological origin, from the mesoderm layer instead of neuroectoderm, like the other neural cells (Fig. S6).

### 1.6.3 Identification of neural subtype specific genes

We used single cell RNA-seq data (Darmanis et al. 2015) to identify genes specific of each neural cell subtype, namely astrocytes, neurons, oligodendrocytes, oligodendrocyte precursors (OPC), brain endothelial cells, and microglia. Pairwise differential expression was performed with edgeR between cells each cell type and the rest of the cells. We identified 3,136 differentially expressed genes with false discovery rate (FDR)  $<0.01$  and at least 4-fold change, consistently expressed in multiple cells of the same type ( $> 1$  cpm in at least 10 cells). Specifically, we identified 722 astrocyte, 264 brain endothelial, 285 microglia, 1,287 neuron, 174 OPC and 404 oligodendrocyte specific genes, after removing cell type and neural specific genes (Fig. S7). These results show that despite sharing a common neural transcriptional program, different neural cells, express different sets of genes, which further characterize each subtype.

## 1.7 Identification of cell type specific TFs

We intersected our list of 2,871 cell type specific genes with a list of 1,558 TFs from the AnimalTFDB database (Zhang et al. 2011) to get a set of 167 cell type specific TFs. To identify the most correlated TFs (Fig. 2C) we computed the Pearson's correlation coefficient between log10-transformed RPKM values (pseudocount 0.01) of pairs of TFs across all samples. A final set of 56 TFs had a correlation of at least 0.85 with at least another TF. Of these, 38 have at least one known or novel sequence motif in (Kheradpour and Kellis 2014).

## 1.8 Identification of driver genes through PCA

In addition to differential gene expression, PCA was used to identify driver genes through their loadings. As a result of PCA, each gene is scored for its correlation with each principal component (loadings). Major cell type clusters are better separated in a space defined by more than one principal component: for example, endothelial cells segregate from the other cell types both along PC1 and PC2. Thus, to find driver genes, we computed the cosine similarity between each gene vector of loadings and the centroid of each major cell type in principal component coordinates. We selected 233 genes with cosine  $> 0.9$ , of which 228 (98%) are cell type specific according to differential expression analysis. The 5 remaining genes, while significantly differentially expressed, did not pass the chosen logFC threshold. To facilitate interpretation, we manually curated a list of 43 driver genes with reported cell type specific functions in the literature (Table S8).

## 1.9 Relative contribution of gene expression to variability in isoform abundances

Gene expression contribution in the transcript abundance variation across all samples was computed following the methodology presented in (Gonzalez-Porta et al. 2012) and further improved in (Melé et al. 2015). In a nutshell, for each gene, samples are represented in a multidimensional space using their transcript abundances as coordinates. The contribution of gene expression in the transcript abundance variation is computed by the variation of transcript abundance after projecting the samples into a model of constant splicing (a line in the multidimensional space) divided by the total variation of transcript abundance without projection. If this ratio is close to 1, projecting into the "no splicing" model didn't reduce the transcript abundance variation, pointing at mainly gene expression contribution. Inversely, if close to 0, alternative splicing (or post-transcriptional regulation, in general) is mostly responsible for the major part of the transcript abundance variation.

A generalization of this approach allows to estimate the effect of a given factor, in this case the cell type, to the contribution of gene expression in transcript abundance variation. Precisely, we asked how much of the transcript variation attributed to cell type is due to changes in gene expression. In practice we compare the proportion of variation explained by the cell type classification after and before projecting the samples into the "no splicing" model. The proportion of variance explained by cell type classification is derived from the classical ANOVA decomposition. The "no splicing" model is represented by a line in the multidimensional space formed by the dif-

ferent transcripts abundances. Like in the analysis across all samples, a value around 1 means that the projection didn't affect the estimate of variance explained, supporting a full contribution of gene expression. A ratio around 0 means that the variance explained was greatly reduced after projection, supporting a major contribution of alternative splicing.

## 1.10 Conservation of cell type specific genes

### 1.10.1 Sequence conservation

As a measure of sequence conservation of cell type specific genes, we computed the fraction of genes in each set which have a one-to-one orthology relationship with human and other ten vertebrates, as well as worm and fly. In Figure 3C, species are sorted by increasing evolutionary distance from human (Kumar et al. 2017). The list of orthologous genes was retrieved from Ensembl Compara (Herrero et al. 2016) v75, which is compatible to Gencode v19 (<http://feb2014.archive.ensembl.org/biomart/martservice>). As a control, the fraction of orthologs was computed also for all protein coding genes (20,731 genes) and for the set of already defined orthologous genes in (Barbosa-Morais et al. 2012) and filtered in (Breschi et al. 2016). As this latter set of orthologous genes was from a previous version of Ensembl, we used a subset of them which are retained through the version 75 (6,268 out of 6,283).

### 1.10.2 Expression conservation

Using a similar approach to what is described in (Breschi et al. 2016), expression conservation was measured as the Pearson's correlation coefficient between expression in human and any other vertebrate species in each organ for the different gene sets. Expression data were obtained from (Breschi et al. 2016; Barbosa-Morais et al. 2012). Only cell type specific genes with orthologous genes in that dataset were used: 207 of 635 endothelial genes, 277 of 950 epithelial, 130 of 353 melanocyte, 295 of 935 mixed.

## 2 Analysis of Cap Analysis of Gene Expression data (CAGE)

Gene expression data from CAGE in human primary cells was obtained using FANTOM5 CAGE data (The FANTOM Consortium and the RIKEN PMI and CLST (DGT) 2014) and through a private collaboration. Read counts for each gene were obtained by summing up the read counts for all the

promoters of that gene and then normalized to cpm. Only samples that passed FANTOM5 quality filters were retained and two biological replicates were selected at random, when more than two were available for a given cell line, resulting in 142 cell lines. We applied npIDR filtering (Djebali et al. 2012) on read counts between the two replicates, set to 0 the cpm values when the read counts were not reproducible ( $npIDR > 0.1$ ), and averaged the cpm between replicates. Genes with  $cpm < 5$  in all samples were discarded, resulting in 21,269 genes.

t-SNE was computed on the distance matrix between each pair of samples (Fig. 1D) using log10-transformed cpm after adding a pseudocount of 1. The distance between two samples is computed as  $abs(1-cc)$ , where  $cc$  is again the Pearson's correlation coefficient between the vectors of peak quantifications.

PCA was computed on the same cpm values as t-SNE, after mean-centering and scaling for each gene across all samples (Fig. S1C).

### 3 Analysis of DNase-seq data

#### 3.1 DHS filtering

DHSs for 127 human primary cells were downloaded from the ENCODE portal as narrowPeak files (Table S6). The peaks were filtered within each sample to remove the ones with a read count lower than the minimum number of reads for which there is the maximum p-value. The filtered peaks are merged across experiment if they are separated by at most one nucleotide, resulting in 742,099 merged peaks. Read counts are normalized to counts per million with TMM normalization (Robinson and Oshlack 2010). The sum of read counts within each sample after filtering is used as library size. Finally, we retained 555,693 peaks after removing the ones that do not have 1 cpm in at least two samples.

#### 3.2 t-SNE and PCA based on DHSs

As with CAGE peaks, t-SNE was computed on the distance matrix between each pair of samples (Fig. 1E) using log10-transformed cpm after adding a pseudocount of 1. The distance between two samples is computed as  $abs(1-cc)$ , where  $cc$  is again the Pearson's correlation coefficient between the vectors of peak quantifications.

PCA was computed on the same cpm values as t-SNE, after mean-centering and scaling for

each gene across all samples (Fig. S1B).

### 3.3 TF binding prediction in DHSs

We used the 742,099 merged peaks to look for cell type specific binding sites. First, for each DNase-seq sample we computed the proportion of cell type specific genes with a peak within their TSS region, extended 10kb upstream and 5kb downstream. TSSs for all transcripts of a given gene are considered to take into account for alternative TSS usage. Next, we scanned the peaks for known and novel motifs (Kheradpour and Kellis 2014) of our defined cell type specific TF, with FIMO (Grant et al. 2011) (default parameters). We used CENTIPEDE (Pique-Regi et al. 2011) to filter for motif hits with high likelihood of predicted binding (simple model,  $\lambda=0.1$ , posterior probability  $> 0.95$ ). The proportions shown in Figure 2D for each sample are calculated as the number of cell type specific genes with predicted TF binding in their TSS over the total number of cell type specific genes with DNase peaks in their TSS.

## 4 Analysis of Tabula Muris data

The Tabula Muris project generated a compendium of single-cell transcriptomic data from more than 100,000 cells isolated from 20 organs from three female and four male three-month-old mice (The Tabula Muris Consortium 2018). For each mouse, several organs were analyzed, controlling in this way for age, environment and epigenetic effects, enabling direct comparison of cell type composition between organs. Two distinct technological approaches were employed: microfluidic droplet-based 3'-end counting using the 10x Genomics platform, and FACS-based full length transcript analysis using SMART-Seq2 libraries (hereafter "10x" and "FACS", respectively). These platforms generated gene expression information of 55,656 and 44,949 cells, assigned to 55 and 81 known cell types, respectively.

To preprocess the data, we used the original "data ingest" scripts provided by the Tabula Muris Consortium (<https://github.com/czbiohub/tabula-muris>), but excluding cells that were labeled generically: bladder cells (1,203 in 10x and 695 in FACS), kidney cells (45 in 10x), and cells with no assigned ontology class (791 in 10x and 170 in FACS). This results in 53,617 and 44,084 cells, grouped into 53 and 80 cell types, respectively. The generated *Seurat* objects were upgraded to version 3.0.1 using the *UpgradeSeuratObject* function. For downstream analyses,

only the sets of variable genes identified by *Seurat* across the cells were used (2,010 and 4,694 genes, respectively).

We assigned to the cells a major cell type (blood, epithelial, mesenchymal and endothelial, plus neural in the case of the FACS data) based on the known histology of the cells. Cells with specialized functions were annotated as tissue-specific (TS) cells, and were not assigned a major cell type category. In 10x, these were bladder urothelial cells, mast cells and neuroendocrine cells; for FACS, this was the case for 11 cell types: Kupffer cells, Slamf1-negative multipotent progenitor cells, Slamf1-positive multipotent progenitor cells, pancreatic A cells, pancreatic D cells, pancreatic PP cells, pancreatic acinar cells, pancreatic ductal cells, pancreatic stellate cells, professional antigen presenting cells, and type B pancreatic cells.

PCA embeddings of the cells based on the expression of variable genes were computed using *Seurat* (Fig. S8A and S9A). Most cells clearly grouped into the blood, mesenchymal, endothelial, and epithelial major cell types, plus neural in the case of FACS. The exception was hepatocytes (1,764 and 391 cells in 10x and FACS, respectively) which formed a clearly separated transcriptional group. For this reason, we excluded them from further analyses, recomputed variable genes (1,894 and 4,671 for 10x and FACS, respectively) and regenerated the PCA for the cells in both platforms (Fig. S8B and S9B).

We computed Spearman's correlation between the mean expression profiles of the cells assigned to each cell type - tissue pair, and then performed hierarchical clustering using Ward2's clustering criterion with euclidean distance (Fig. S8C and S9C).

In 10x cells, nearly all cell types clustered within the corresponding major cell types, with the exception being cardiac muscle cells which we labelled as mesenchymal, but cluster within the epithelial types. In FACS cells, most cells also properly clustered within the major cell types. Indeed, blood, mesenchymal, endothelial and epithelial cells clustered mostly within their major cell types. Neural cells clustered together with different types of pancreatic endocrine cells, consistent with the strong morphological, and physiological similarities between these two types of cells (Arntfield and van der Kooy 2011), and with a few epithelial cells. This resembles the neuroepithelial cluster that we found in human. The exception is oligodendrocytes, that we labelled as neural, but that unexpectedly cluster within the epithelial type. Finally, a very few cell types that we labelled as mesenchymal and epithelial clustered together separated from the rest, but connected to the epithelial cluster.

## 5 Analysis of chromatin marks

ChIP-seq histone modification data was obtained from the ENCODE portal (The ENCODE Project Consortium 2012) for 15 cell lines for which RNA-seq data was also available (Table S9). The data was generated with antibodies for the following histone marks: H3K27ac, H3K27me3, H3K36me3, H3K4me1, H3K4me2, H3K4me3, H3K79me2, H3K9ac, H3K9me1, H3K9me3. For each gene in each experiment, we averaged the ChIP-seq signal every 10 nucleotides in the region around the TSS (+/- 2kb), with the software bwtool (Pohl and Beato 2014).

To assess whether cell type specific genes have stronger activating marks (H3K4me3, H3K27ac) in the corresponding cell type compared to the other cell types (Fig. S16), we normalized the signal over the sum of the signal at the TSS of all genes.

To compare chromatin modifications for the different sets of cell type specific genes (Fig. S16B), we selected a subset of genes, that are not cell type specific, but have comparable expression levels to each set of cell type specific genes separately. This was achieved by binning the reference set of cell type specific genes into 200 bins of expression and by randomly selecting an equivalent number of non cell type specific genes belonging to those expression bins. In addition, because we aimed to compare chromatin signal within the same cell, the signal was normalized between 0 and 1, so within-sample differences are on similar scales across samples.

## 6 Splicing

### 6.1 Computation of exon inclusion levels

Exon inclusion levels, measured as percent spliced-in (PSI) values (Pervouchine et al. 2013), were computed for 439,778 internal exons with IPSA Splicing Analysis Pipeline (<https://github.com/pervouchine/ipsa>). PSI values were computed only for splicing events supported by at least 10 reads. PSI values for exons with an absolute difference in PSI values larger than 0.1 between two replicates were set to NA in both samples. PSI values were averaged between replicates, when available.

We also used LeafCutter (Li et al. 2018) with the default parameters and the protocol described by the authors. First, we converted the BAM files described in the "RNA-seq processing pipeline" section into junction files. Then, we perform intron clustering using the default parameters of 30 minimum reads in a cluster, 100,000 bp as maximum intron length and 0.001 for minimal fraction of reads in a cluster that support a junction. In order to be consistent with the analysis on IPSA,

the average intron clustering value is computed for each pair of replicates.

## 6.2 Gene Expression vs Exon inclusion

We compared the differences of gene expression and of exon inclusion within and between cell types. The difference of gene expression within and between cell types was computed as the number of genes differentially expressed (edgeR (Robinson et al. 2010),  $FDR < 0.01$ ,  $\log_2\text{fold-change} > 2$ ) between each pair of samples belonging to the same or different cell types, respectively. The difference of exon inclusion within and between cell types, instead, was computed as the number of genes containing at least one exon differentially included between pairs of the same or different cell types, respectively, when using IPSA. We defined an exon as differentially included between two samples if the absolute difference between its PSIs in the two samples was larger than 0.2. We also used other cutoffs, 0.1, 0.3, 0.4, 0.5, obtaining similar results (Fig. S19B). Finally, we computed the ratio between the number of differentially expressed genes (DE genes) and the number of genes with differentially included exons (DS genes) for each pair of samples. As there was evident bias between PSI values from the first two batches compared to the third, we restricted the analysis to the first and the second batches, which include enough representative samples for the different cell types (Fig. 3B).

For LeafCutter, the differential splicing (DS) analysis has then been performed using the LeafCutter script between each pair of primary cell lines, using a threshold on the FDR of 0.01 to define the differential intron cluster. The DS analysis has been run using the following parameters: `-min_samples_per_group=1` and `-min_samples_per_intron=1` to be able to run DS with only two samples. Like for IPSA, a gene is considered as DS if it contains at least one differentially spliced intron cluster with a  $\Delta\text{PSI} > 0.2$ , and only the first and second batches were used.

## 6.3 Cell type specific alternative splicing

We compute event-specific PSI values for well defined alternative splicing events using an in-house script, from exclusion and inclusion read counts obtained by the IPSA pipeline. Alternative splicing events were defined based on the annotation between pairs of transcripts with the AStalavista tool (Sammeth et al. 2008). We ensure reproducibility between replicates by requiring that the event is supported by at least 10 reads in both replicates and that the difference in PSI between the replicates is less 0.15. PSI values are then averaged between replicates. Events that do not pass

these filters are assigned NA in that pair of sample. Because we are interested in cell type specific events, we filter for events with no more than 75% missing values of the samples within each cell type. Events are further filtered by total variance  $> 0.01$  across all samples. Finally, missing values are imputed for the remaining 1,113 events with `impute.knn()` function (Troyanskaya et al. 2001) from the R package `impute` which employs the k-nearest neighbors algorithm. The 1,113 events are tested for cell type specificity with the Kruskal-Wallis test, a non-parametric alternative to one-way ANOVA, and p-values are corrected for multiple testing with the Benjamini-Hochberg procedure. With 1% FDR we identified: 3 alternative acceptor events, 3 alternative donor events, 46 multiple exon skipping events, 14 mutually exclusive exon events, 7 retained introns and 157 single cassette exons (Table S10). Sashimi plots were created with `ggsashimi` (Garrido-Martín et al. 2018). In Figure S20C, the exon has been previously reported to be more often included in muscle cells (Lenz et al. 1989).

## 7 Analysis of alternative TSS usage with RAMPAGE data

We analyzed RAMPAGE data from 28 experiments, corresponding to 14 primary cells with 2 replicates. The TSS peaks were downloaded from the ENCODE portal (Table S3) in the gff format. The TSS peaks are associated to genomic regions, but not directly to genes. In order to associate a TSS peak to a gene, we assign the TSS to the closest gene with the largest overlap with the associated genomic region. If the closest gene is farther than 10kb, the TSS is annotated as novel and is not assigned to any gene. TSSs that are overlapping at least by one nucleotide are merged across samples, yielding 34,802 TSSs. Read counts are normalized to counts per million of mapped reads. We further retained 9,068 genes (16,918 TSSs), which have at least one TSS with  $\text{cpm} > 0.5$  in at least 2 samples. Of these, 3,856 genes which had more than one TSS were tested for differential TSS usage with an approach similar to the one described in the previous section.

For each gene, we estimated the variance in TSS quantifications between cell types that can be attributed to changes in gene expression (`var.bwGp.ge`) or to changes in TSS usage (`var.bwGp`). To select genes with differential TSS usage between cell types, we binned them based on their variation in gene expression (`var.bwGp.ge`) in 10 quantiles and computed the Z-score of the variation in TSS usage with respect to each quantile (Fig. S21). We defined 17 genes whose variation in TSS usage is larger than the variation in expression and with a Z-score  $> 1.5$ .

## 8 Analysis of GTEx data

### 8.1 Selection of tissue-specific genes

Tissue-specific genes were selected according to a tissue-specificity (ts) score described in (Yang et al. 2018), which accounts for the fact that genes can be expressed in a selective way in multiple tissues rather than just in a single one. In order to select only the genes with restricted expression to few tissues, we used a more stringent threshold than the one used in the original work (ts score $>4$  instead of ts score $>3$ ). In addition, genes are considered specific of a given tissue if they have a median RPKM $>1$  across the samples of that tissue. These filters led to a set of 10,365 tissue-specific genes for all the GTEx tissues and sub-tissues, most of which are testis and brain specific (4,693 and 1,980 genes, respectively).

### 8.2 Estimation of cellular enrichments

In this analysis we aimed to estimate the cellular enrichments of the five major cell types described so far in whole organs, using gene expression data from the GTEx Consortium (Lonsdale et al. 2013).

To this end, we used the xCell 1.1.0 R package, which is signature-based cellular enrichment method that depends on lists of marker genes that are representative of a given cell type. The methodology details for the enrichment score estimation can be found in the original xCell publication (Aran et. al 2017). We replace the xCell default signature gene sets by a signature composed of six gene sets: four sets correspond to the genes specific to three major cell types (endothelial, epithelial and mesenchymal), and melanocytes, obtained as previously described (see section "Identification of cell type specific genes"). The fifth set corresponds to the Neural cell type, which is the union of the genes specific to the 13 GTEx brain subregions (see section "Selection of tissue-specific genes"). The last set corresponds to the Blood major cell type, and it is the set of genes specific to the Whole Blood tissue (see section "Selection of tissue-specific genes").

xCell authors state that "xCell performs best with heterogeneous datasets. Thus it is recommended to use all data combined in one run, and not break down to pieces (especially not cases and control in different runs)" (see <http://xcell.ucsf.edu/>). For this reason, we have combined the gene expression matrices of ENCODE primary cells, GTEx tissues, and PCAWG samples into a single matrix in order to perform the enrichment estimation, keeping only the genes

that are present in at least one of the six gene sets.

Results are presented in Figure 4A. As a control, we also include the enrichments in the ENCODE primary cells monitored here. As expected, these samples have the highest enrichment for major cell types (endothelial, epithelial and mesenchymal) to which they correspond. Since xCell works on ranks, it follows that the enrichment of the GTEx and PCAWG samples should be between 0 and the enrichment of the ENCODE samples, as they should be the "maximally enriched" ones, providing a fair way to compare tissue samples within the same major cell type. However, it is important to realize that these enrichment scores are not comparable across cell types (for example, an enrichment score of 1000 in Endothelial cell type might not be comparable with an enrichment score of 1000 in Epithelial cell type). The same applies for brain and whole blood GTEx samples and their Neural and Blood major cell type enrichment estimations.

We projected the tissue samples on a three-dimensional space where the axes correspond to the enrichment for the major cell types (endothelial, epithelial and mesenchymal). The projection shows that each tissue has its own characteristic cellular composition (Fig. 5A). This is observed as well when performing dimensionality reduction (t-SNE and UMAP) over the enrichment scores of the five major cell types (Fig. S27).

The matrix of RPKM values for tissues was downloaded from the GTEx portal (<http://www.gtexportal.org/home/datasets/>, file: `GTEx_Analysis_v6_RNA-seq_RNA-SeQCv1.1.8_gene_rpkm.gct.gz`) for 8,527 samples (after filtering) and 56,319 genes of the v6 release. Samples extracted from tissues with less than 20 individuals were removed from the analysis. These included bladder, cervix and fallopian tubes.

### 8.3 Classification of stomach histological samples

Histological images for the corresponding samples with RNA-seq data can be publicly found on the Biospecimen Research Database website (<https://brd.nci.nih.gov/brd/image-search/searchhome>). Images of histological slides of stomach sections were classified based on the presence of the mucosa and muscularis layers (Fig. 4B). To each slide we manually assign two binary vectors, one for the mucosa and one for the muscularis layer, where 1 and 0 indicate presence or absence of the layer, respectively. The length of the vectors depends on the number of tissue sections on the slide, and the order reflect the order of the tissue sections from left to right and from top to bottom. Then, we computed the proportion of sections with a given layer and rounded the proportion to be binary. Thus, each slide will have one of the following possible com-

positions: mc1ms1, mc1ms0, mc0ms1, mc0ms0. Finally, we focused on the samples with either one or the other layer, i.e. mc1ms0 (mostly mucosa) and mc0ms1 (mostly muscularis).

#### **8.4 Identification of variable genes among stomach samples**

To find the genes at the base of the transcriptional differences we observed amongst stomach samples, we selected in an unsupervised fashion the most variable genes, by using the projection score (Fontes et al. 2011). We identified 500 most variable genes which maximize the projection score for the first three principal components. Of these, 96 are cell type specific (84 epithelial-specific, 9 mesenchymal, 2 endothelial, 1 melanocyte), according to our definition. The expression of these 96 genes clearly discriminates stomach samples with only mucosa or muscularis layer (Fig. 4D).

#### **8.5 Automated classification of colon samples**

We used a machine learning approach to automatically classify the histological slides based on the presence of the mucosa and/or the muscularis layer.

First, we implemented a script in Python 2.7 using the OpenCV2 library in order to preprocess the images, find the tissue piece contours, and create binary masks that allowed to save each piece into a separate image. Because the two layers have distinctive color patterns in the histological slides, we used all the channels from the RGB color model and the first two channels from the HSV color model as features. For each image, we computed the normalized histogram for each color channel. The concatenation of these histograms will be the feature set for each observation.

We used the manually annotated labels (muscularis/mucosa/both) for stomach samples to train a support vector machine (SVM) model with radial basis function (RBF) kernel. The SVM was fitted with 10-fold cross validation using the caret package in R. 75% of the samples (550 images) were used to train and validate the model and the other 25% (182 images) for testing, obtaining an accuracy of 87% over the test set (Fig. S26B). The 95% CI for the prediction is (81.64%, 91.82%), the No Information Rate (NIR, which is the proportion of the largest class) is 48.35%, and the p-value [Acc > NIR] is  $< 2.2 \times 10^{-16}$ .

Since the color composition of stomach and colon tissues is very similar, we applied the classifier to colon pieces as well in order to produce the class labels for each image. The descriptive features for colon were generated in the same way as the ones for stomach.

## 8.6 Analysis of single cell RNA-seq data

### 8.6.1 Analysis of muscle single cell RNA-seq data

Muscle single cell RNA-seq data was downloaded from the R package *HSMMSingleCell* (Trapnell et al. 2014). The data consists of gene expression measurements from an induced differentiation of primary human myoblasts and was processed following the authors' guidelines (<http://monocle-bio.sourceforge.net/monocle-vignette.pdf>). Briefly, the gene expression matrix was filtered to remove lowly expressed genes and cells with few detected genes. Differentially expressed genes to use as reference for the pseudotemporal sorting were inferred from a predefined list of 21 marker genes known in myoblast differentiation. Pseudotemporal sorting was performed with the R package 'monocle' v0.99.0 (Trapnell et al. 2014). For cell type specific genes and muscle specific genes from GTEx, the Pearson's correlation coefficient is computed between their expression in each single cell and the average expression from the muscle gene signature matrix, i.e. the average expression within each major cell type and median expression in skeletal muscle samples from GTEx (Fig. S25C).

t-SNE was performed on the joined gene expression matrix for single cells and for the ENCODE primary cells, after normalization. Specifically, for ENCODE data, the batch-normalized gene expression matrix was used, generated as described above. The expression values for each gene of matrices were then mean-centered and scaled separately for each matrix. Principal component analysis was used to reduce the dimensionality of the joined dataset. Finally, t-SNE was applied on the resulting principal components with a perplexity of 10 using the R package *Rtsne* v0.13 (Krijthe 2015).

## 8.7 Relationship between estimated cellular enrichments and histological phenotypes

### 8.7.1 Parsing histopathological annotation

Each GTEx histological image has an associated free-form text comment detailing the findings of the pathology review. We seek to automatically recover morphological and/or medical phenotype information from these comments and structure it in a friendly way to perform statistical analyses, however, the free-text poses a challenge since the format is inconsistent and difficult to standardize. To solve these issues, we developed a text processing procedure in Python 2.7.

First, we search which annotations contain a specific phenotype within a certain tolerance threshold, using Levenshtein distance to account for typographical errors (insertions, deletions, substitutions). Next, we perform part-of-speech tagging to find dependencies between the words, using the TextBlob NLP library. With these part-of-speech tags, we can build rules to keep words on which we are interested. Using this methodology we are able to reshape automatically many of the free-form text annotations into a friendly format to perform statistical analyses (Table S13), however, a small fraction of the annotations still had to be manually curated due to some annotations not having a correct part-of-speech tag.

### **8.7.2 Testing the relationship**

We tested the significance of alterations in the xCell enrichments for 61 phenotypes. We performed a Mann-Whitney *U* test, and corrected for multiple testing (Benjamini-Hochberg correction (Benjamini and Hochberg 1995)).

## **8.8 Estimation of fat proportions in muscle based on pathology annotations**

Each histological image is accompanied by a pathology review comment. For the case of muscle tissue, many of these comments contain a numerical estimation of the proportion of fat present in the tissue. We were able to extract these estimations using regular expressions and spaCy 2.0.7, which is a Python API for text processing. For samples that contain fat estimations that are expressed as a range (for example, 5%-10%) we take the midpoint (7.5%) as an approximation of the value for the sample. Once we have the estimates at the sample level, we compute the summary statistics stratifying by atrophic and non-atrophic samples.

## **8.9 Estimation of fat proportion in Adipose - Subcutaneous histology images**

Each histology image is encoded as an SVS file which contains several layers that contain the image at different resolutions. We extract a downscaled version of the histology images from 305 SVS files corresponding to Adipose - Subcutaneous samples in order to perform the estimation. After which, we preprocess the images to delete small components and reduce noise. Then, we generate a mask over the regions of interest, which are the tissue cuts. We infer the two most dominant colors inside the masked regions, and then we assign to each pixel in the masked region one of these two most dominant colors with respect to their color proximity. Once the image has

been transformed, we can calculate the proportions of the colors to obtain an estimation of the proportion of fat. With this procedure, we obtain a mean fat proportion of 84%. All the image manipulations were performed using Mathematica 11.0.

## 9 Analysis of cancer expression data

We used gene expression data for 1,963 cancer and normal samples, measured as fragments per million mapped reads (FPKM), from the Cancer Genome Atlas Pan-Cancer analysis project (PCAWG) (The Cancer Genome Atlas Research Network et al. 2013) encompassing 30 cancer sequencing projects and 20 organs and tissues. To have a comparable expression measure to the other data sets in this study, we converted FPKM to RPKM by multiplying each value by a factor of 2. The cellular enrichments of cancer and normal samples from the PCAWG project were estimated at the same time that the GTEx enrichments, as previously described.

## References

Aran D, Hu Z, Butte AJ. 2017. xCell: digitally portraying the tissue cellular heterogeneity landscape. *Genome Biology* **18**. <https://doi.org/10.1186/s13059-017-1349-1>

Arntfield ME, van der Kooy D. 2011.  $\beta$ -Cell evolution: How the pancreas borrowed from the brain. *BioEssays* **33**: 582–587.

Barbosa-Morais NL, Irimia M, Pan Q, Xiong HY, Gueroussov S, Lee LJ, Slobodeniuc V, Kutter C, Watt S, Colak R, et al. 2012. The Evolutionary Landscape of Alternative Splicing in Vertebrate Species. *Science* **338**: 1587–1593.

Benjamini Y, Hochberg Y. 1995. Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society: Series B (Methodological)* **57**: 289–300.

Breschi A, Djebali S, Gillis J, Pervouchine DD, Dobin A, Davis CA, Gingeras TR, Guigó R. 2016. Gene-specific patterns of expression variation across organs and species. *Genome Biology* **17**. <https://doi.org/10.1186/s13059-016-1008-y>

The Cancer Genome Atlas Research Network, Weinstein JN, Collisson EA, Mills GB, Shaw KRM, Ozenberger BA, Ellrott K, Shmulevich I, Sander C, Stuart JM. 2013. The Cancer Genome Atlas Pan-Cancer analysis project. *Nature Genetics* **45**: 1113–1120.

Carvalho BS, 2017. pd.huex.1.0.st.v2. <https://doi.org/doi:10.18129/B9.bioc.pd.huex.1.0.st.v2>

Carvalho BS, Irizarry RA. 2010. A framework for oligonucleotide microarray preprocessing. *Bioinformatics* **26**: 2363–2367.

Darmanis S, Sloan SA, Zhang Y, Enge M, Caneda C, Shuer LM, Hayden Gephart MG, Barres BA, Quake SR. 2015. A survey of human brain transcriptome diversity at the single cell level. *Proceedings of the National Academy of Sciences* **112**: 7285–7290.

Djebali S, Davis CA, Merkel A, Dobin A, Lassmann T, Mortazavi A, Tanzer A, Lagarde J, Lin W, Schlesinger F, et al. 2012. Landscape of transcription in human cells. *Nature* **489**: 101–108.

The ENCODE Project Consortium. 2012. An integrated encyclopedia of DNA elements in the human genome. *Nature* **489**: 57–74.

Falcon S, Gentleman R. 2007. Using GOstats to test gene lists for GO term association. *Bioinformatics* **23**: 257–258.

The FANTOM Consortium and the RIKEN PMI and CLST (DGT). 2014. A promoter-level mammalian expression atlas. *Nature* **507**: 462–470.

Fontes M, Soneson C. 2011. The projection score - an evaluation criterion for variable subset selection in PCA visualization. *BMC Bioinformatics* **12**. <https://doi.org/10.1186/1471-2105-12-307>

Garrido-Martín D, Palumbo E, Guigó R, Breschi A. 2018. ggsashimi: Sashimi plot revised for browser- and annotation-independent splicing visualization. *PLoS Computational Biology* **14**. <https://doi.org/10.1371/journal.pcbi.1006360>

Gonzalez-Porta M, Calvo M, Sammeth M, Guigo R. 2012. Estimation of alternative splicing variability in human populations. *Genome Research* **22**: 528–538.

Grant CE, Bailey TL, Noble WS. 2011. FIMO: scanning for occurrences of a given motif. *Bioinformatics* **27**: 1017–1018.

Harrow J, Frankish A, Gonzalez JM, Tapanari E, Diekhans M, Kokocinski F, Aken BL, Barrell D, Zadissa A, Searle S, et al. 2012. GENCODE: The reference human genome annotation for The ENCODE Project. *Genome Research* **22**: 1760–1774.

Herrero J, Muffato M, Beal K, Fitzgerald S, Gordon L, Pignatelli M, Vilella AJ, Searle SMJ, Amode R, Brent S, et al. 2016. Ensembl comparative genomics resources. *Database* **2016**. <https://doi.org/10.1093/database/bav096>

Kheradpour P, Kellis M. 2014. Systematic discovery and characterization of regulatory motifs in ENCODE TF binding experiments. *Nucleic Acids Research* **42**: 2976–2987.

Krijthe JH. 2015. Rtsne: T-Distributed Stochastic Neighbor Embedding using Barnes-Hut Implementation. R package version 0.13. URL: <https://github.com/jkrijthe/Rtsne>

Kumar S, Stecher G, Suleski M, Hedges SB. 2017. TimeTree: A Resource for Timelines, Timetrees, and Divergence Times. *Molecular Biology and Evolution* **34**: 1812–1819.

Lenz S, Lohse P, Seidel U, Arnold HH. 1989. The alkali light chains of human smooth and nonmuscle myosins are encoded by a single gene. *Tissue-specific expression by alternative splicing pathways. Journal of Biological Chemistry* **264** (15): 9000-9015.

Li YI, Knowles DA, Humphrey J, Barbeira AN, Dickinson SP, Im HK, Pritchard JK. 2018. Annotation-free quantification of RNA splicing using LeafCutter. *Nature Genetics* **50**: 151–158.

Lonsdale J, Thomas J, Salvatore M, Phillips R, Lo E, Shad S, Hasz R, Walters G, Garcia F, Young N, et al. 2013. The Genotype-Tissue Expression (GTEx) project. *Nature Genetics* **45**: 580–585.

Melé M, Ferreira PG, Reverter F, DeLuca DS, Monlong J, Sammeth M, Young TR, Goldmann JM, Pervouchine DD, Sullivan TJ, et al. 2015. The human transcriptome across tissues and individuals. *Science* **348**: 660–665.

Nogalski MT, Solovyov A, Kulkarni AS, Desai N, Oberstein A, Levine AJ, Ting DT, Shenk T, Greenbaum BD. 2019. A tumor-specific endogenous repetitive element is induced by herpesviruses. *Nature Communications* **10**. <https://doi.org/10.1038/s41467-018-07944-x>

Robinson MD, McCarthy DJ, Smyth GK. 2010. edgeR: a Bioconductor package for differential expression analysis of digital gene expression data. *Bioinformatics* **26**: 139–140.

Robinson MD, Oshlack A. 2010. A scaling normalization method for differential expression analysis of RNA-seq data. *Genome Biology* **11**. <https://doi.org/10.1186/gb-2010-11-3-r25>

Rousseeuw PJ. 1987. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics* **20**: 53–65.

Sammeth M, Foissac S, Guigó R. 2008. A General Definition and Nomenclature for Alternative Splicing Events. *PLoS Computational Biology* **4**. <https://doi.org/10.1371/journal.pcbi.1000147>

Sheffield NC, Thurman RE, Song L, Safi A, Stamatoyannopoulos JA, Lenhard B, Crawford GE, Furey TS. 2013. Patterns of regulatory activity across diverse human cell types predict tissue identity, transcription factor binding, and long-range interactions. *Genome Research* **23**: 777–788.

The Tabula Muris Consortium. 2018. Single-cell transcriptomics of 20 mouse organs creates a Tabula Muris. *Nature* **562**: 367–372.

Thorndike RL. 1953. Who belongs in the family? *Psychometrika* **18**: 267–276.

Trapnell C, Cacchiarelli D, Grimsby J, Pokharel P, Li S, Morse M, Lennon NJ, Livak KJ, Mikkelsen TS, Rinn JL. 2014. The dynamics and regulators of cell fate decisions are revealed by pseudotemporal ordering of single cells. *Nature Biotechnology* **32**: 381–386.

Pervouchine DD, Knowles DG, Guigo R. 2013. Intron-centric estimation of alternative splicing from RNA-seq data. *Bioinformatics* **29**: 273–274.

Pique-Regi R, Degner JF, Pai AA, Gaffney DJ, Gilad Y, Pritchard JK. 2011. Accurate inference of transcription factor binding from DNA sequence and chromatin accessibility data. *Genome Research* **21**: 447–455.

Pohl A, Beato M. 2014. bwtool: a tool for bigWig files. *Bioinformatics* **30**: 1618–1619.

Troyanskaya O, Cantor M, Sherlock G, Brown P, Hastie T, Tibshirani R, Botstein D, Altman RB. 2001. Missing value estimation methods for DNA microarrays. *Bioinformatics* **17**: 520–525.

Yang RY, Quan J, Sodaei R, Aguet F, Segrè AV, Allen JA, Lanz TA, Reinhart V, Crawford M, Hasson S, et al. 2018. A systematic survey of human tissue-specific gene expression and splicing reveals new opportunities for therapeutic target identification and evaluation. bioRxiv doi: 10.1101/311563

Zhang HM, Chen H, Liu W, Liu H, Gong J, Wang H, Guo AY. 2011. AnimalTFDB: a comprehensive animal transcription factor database. *Nucleic Acids Research* **40**: D144–D149.

Zhang Y, Sloan SA, Clarke LE, Caneda C, Plaza CA, Blumenthal PD, Vogel H, Steinberg GK, Edwards MSB, Li G, et al. 2016. Purification and Characterization of Progenitor and Mature Human Astrocytes Reveals Transcriptional and Functional Differences with Mouse. *Neuron* **89**: 37–53.