A cell-type based model explaining co-expression patterns of genes in the brain

Pascal Grange[1], Jason Bohland[2], Hemant Bokil[1], Sacha Nelson[3], Benjamin Okaty[3], Ken Sugino[3], Lydia Ng[5], Michael Hawrylycz[5], Partha P. Mitra[1]

[1] Cold Spring Harbor Laboratory,
One Bungtown Road, Cold Spring Harbor, New York 11724, United States
[2] Health Sciences Department, Boston University, Boston, Massachussets, United States
[3] Department of Biology, Brandeis University, Waltham, Massachusetts, United States
[4] Department of Genetics, Harvard Medical School, Boston, Massachussets, United States
[5] Allen Institute for Brain Science, Seattle, Washington 98103, United States

∗ E-mail: pascal.grange@polytechnique.org

Abstract

Much of the genome is expressed in the vertebrate brain with individual genes exhibiting different spatially-varying patterns of expression. These variations are not independent, with pairs of genes exhibiting complex patterns of co-expression, such that two genes may be similarly expressed in one region, but differentially expressed in other regions. These correlations have been previously studied quantitatively, particularly for the gene expression atlas of the mouse brain, but the biological meaning of the co-expression patterns remains obscure. We propose a simple model of the co-expression patterns in terms of spatial distributions of underlying cell types. We establish the plausibility of the model in terms of a test set of cell types for which both the gene expression profiles and the spatial distributions are known.

Contents

1 The Allen Gene Expression Atlas

1.1 Gene-expression energies

The gene expression energies we analyzed were obtained from ISH images of thousands of genes in the Allen Gene Expression Atlas [1, 2]. For each of the genes, an eight-week old C57Bl/6J male mouse brain was prepared as unfixed, fresh-frozen tissue. The following steps were taken in an automatized pipeline[1]:

- **Colorimetric** *in situ* **hybridization**;

- **Automatic processing of the resulting images.** Find tissue area eliminating artifacts, look for cell-shaped objects of size $\simeq 10 - 30$ microns to minimize artefacts;

---

[1]For more details on the processing of the ISH image series, see the NeuroBlast User Guide, http://mouse.brain-map.org/documentation/index.html

- **Aggregation of the raw pixel data into a grid.** The mouse brain is partitioned into $V = 49,742$ cubic voxels of side 200 microns. For every voxel v, the *expression energy* of the gene g is defined as a weighted sum of the greyscale-value intensities of pixels p intersecting the voxel:

$$E(v, g) := \frac{\sum_{p \in v} M(p) I(p)}{\sum_{p \in v} 1},$$

  where $M(p)$ is a Boolean mask worked out by step 2 with value 1 if the pixel is expressing and 0 if it is non-expressing.

The present analysis is focused on 4104 genes for which sagittal and coronal data are available. We computed the correlation coefficients between sagittal and coronal data and selected the genes in the top-three quartiles of correlation (3041 genes) for further analysis.

1.2 Classical neuroanatomy: systems of annotation

Partitions of the brain (or of the left hemisphere) into regions, at the same resolution (200 microns) as the gene-expression data, are available from the Allen Reference Atlas [6]. The Allen reference Atlas was obtained using the same mouse strain and methodology as for the the gene-expression data. Each voxel in the mouse brain therefore comes with a label containing the name of the brain region to which it belongs.

As separable parts of the brain (connected components of the gene-expression data in gene space) are of special interest to us, we will make use of two non-hierarchical systems of annotation, available for the left hemisphere. We will refer to them as:
- 'Big 12',consisting of the 12 regions of the left hemisphere (together with a more patchy group of voxels termed 'basic cell groups of regions') whose names, sizes and shapes are shown in Table 1;
- 'Fine', a refinement of 'Big 12' into 94 regions.

2 The Allen Gene Expression Atlas and cell types

2.1 Gene expression energies and cell-types

A major challenge for neurobiology in this era of complete genomes is to map gene expression patterns with cellular resolution in the brain. The Allen Gene Expression Atlas takes a gene-by-gene approach to this problem, obtaining a brainwide expression profile for each gene. Co-expression patterns of genes can then be studied by comparing those maps. A complementary approach uses microarray experiments to study co-expression patterns in a small set of neurons of the same type. We studied two datasets using these two approaches, each of which boils down to a matrix of positive numbers estimating numbers of mRNAs:

- **Allen Gene Expression Atlas (co-registered ISH data in the whole brain).** As explained in the first section, the mouse brain is partitioned into $V = 49,742$ cubic voxels
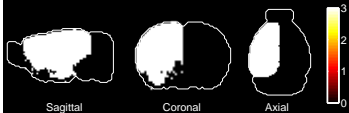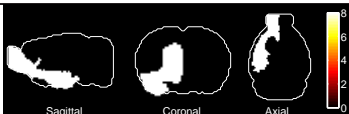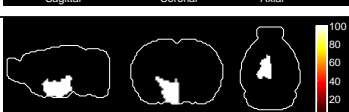
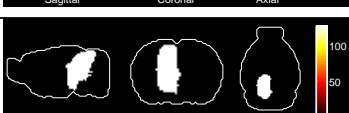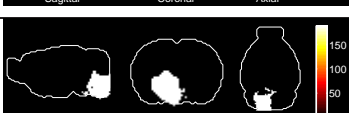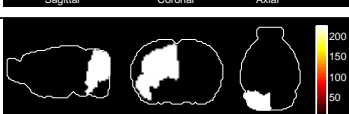| Name of region	Percentage of hemisphere	Profile of region (maximal-intensity projection)
Basic cell groups and regions	4.6	
Cerebral cortex	29.5	
Olfactory areas	9.2	
Hippocampal region	4.3	
Retrohippocampal region	4	
Striatum	8.6	
Pallidum	1.9	
Thalamus	4.3	
Hypothalamus	3.5	
Midbrain	7.8	
Pons	4.6	
Medulla	6.2	
Cerebellum	11.5	

Table 1: Brain regions in the coarsest annotation of the left hemisphere in the Allen Reference Atlas, with the relative volume of the hemisphere they occupy, and maximal-intensity projections

of side 200 microns. For a given gene g and a given voxel v, the AGEA provides the expression energy of g at voxel v, or $E(v, g)$, defined as the average intensity of pixels in the set of ISH images that intersect v. It is an estimator of the number of mRNAs for gene g in voxel v. We focus on about $G_A = 3041$ genes that have the best correlation of signals between sagittal and coronal sections.

- **Cell-type specific microarray data.** We analyzed microarray data for $T = 64$ different cell types [8–12]. Each of these cell types is characterized by expression of $G_T = 14580$ genes. These data can be arranged in a T by G_T matrix C, where $C(t, g)$ is the expression level of gene g in type t.

Studying these two datasets together can help us gain some insight into brainwide co-expression patterns in terms of neuronal types.

We worked out the intersection of the two sets of genes for corresponding to the two datasets. We picked the columns of both matrices corresponding to genes that are in both datasets (let $G = 2131$ denote the number of such genes), and rearranged them in the same order. From now on, E is assumed to be a V times G matrix, and C to be a T times G matrix, with columns of both matrices corresponding to the same set of genes, ordered in the same way. The quantitative techniques used in this section do not depend on the precise value of G, V and T, and the computations can be repeated with richer datasets involving more genes (larger G), higher resolution (larger V), and/or larger numbers of cell types (larger T).

2.2 Correlation between gene-expression data for cell-types and the AGEA

We first studied the correlations between the AGEA and cell-type data. For each cell-type t, we have a vector in gene-space with components $C_t(g) = C(t, g)$, $1 \le g \le G$ from microarray data. We also have a vector in gene space at every voxel v from the AGEA, with components $E(g, v)$, $1 \le g \le G$. We can compute the correlation coefficient between cell-type t and voxel v as the cosine of the angle between the deviation from the average cell-type data and from the average voxel, in gene space:

$$\text{Corr}(t, v) = \frac{\sum_{g=1}^{G}(C(t, g) - \bar{C}(g))(E(v, g) - \bar{E}(g))}{\sqrt{\sum_{g=1}^{G}(C(t, g) - \bar{C}(g))^2}\sqrt{\sum_{g=1}^{G}(E(v, g) - \bar{E}(g))^2}}, \tag{1}$$

$$\bar{C}(g) = \frac{1}{T}\sum_{t=1}^{T} C(t, g), \tag{2}$$

$$\bar{E}(g) = \frac{1}{V}\sum_{v=1}^{V} E(v, g). \tag{3}$$

For a fixed type t, the correlations with all the voxels v can be rearranged into the volume of the brain, and plotted (see the second column of the tables in section 5).
It is quite encouraging to observe that correlation patterns for some cell types exhibit striking

resemblance with neuroanatomical patterns: for instance correlations between the AGEA and the microarray data for medium spiny neurons are highest in the striatum (see Figure 6).

2.3   Fitting a linear combination of cell types to the AGEA

The maps from genes to the whole brain given by the expression energies obtained gene by gene in the AGEA must come from the gene-expression activity of the brain cells. Their co-expression properties must come from the co-expression of genes in given cell types, and from the spatial variation of the abundance of cell types.

We would like to decompose the signal in the AGEA into its cell-type specific components. Let us introduce the quantity $\rho_t(v)$ denoting the contribution of cell type t at voxel v, and propose the following linear model:

$$E(v,g) = \sum_{t=1}^{T} \rho_t(v)C_t(g) + \text{Err}(v,g). \tag{4}$$

Both sides are estimators of the number of mRNAs for gene g at voxel v. The quantity denoted by Err is an error term (reflecting noise in the measurements, reproducibility issues, the non-linearity of the relations between numbers of mRNAs, expression energies and microarray data, and the fact that $T = 64$ types are not enough to sample the whole diversity of cell types in the mouse brain).

In order to find the best fit of the model to the AGEA, we have to solve the following minimization problem under positivity constraints, that corresponds to the least squares of the error function:

$$(\rho_t(v))_{1\leq t\leq T, 1\leq v\leq V} = \text{argmin}_{\phi\in\mathbf{R}_+^T\times\mathbf{R}_+^V}\mathcal{E}_{E,C}(\phi),$$

where T is the total number of types available from the microarray data, V is the total number of voxels in the brain at a resolution of 200 microns, and the function $\mathcal{E}_{E,C}$ is taken to be the sum of squares of the duifferences between the expresison energies and the superposition of cell types weighted by the positive coefficients of the matrix ϕ:

$$\mathcal{E}_{E,C}(\phi) = \sum_{v=1}^{V}\left(\sum_{g=1}^{G}\left(E(v,g) - \sum_{t=1}^{T}\phi(t,v)C_t(g)\right)^2\right).$$

As the terms of the sum over voxels that contain coefficients of a fixed line voxel v of the matrix E, involve only terms of the test function that with the same voxel index v, the problem can be solved voxel by voxel. At each voxel one has to minimize a quadratic function of a vector with T positive components. As the function is quadratic and the domain of the unknown is convex, the problem at each voxel is a convex optimization problem, which admits a global minimum:

$$\forall v \in [1..V], (\rho_t(v))_{1\leq t\leq T} = \text{argmin}_{\nu\in\mathbf{R}_+^T}\sum_{g=1}^{G}\left(E(v,g) - \sum_{t=1}^{T}\nu(t)C_t(g)\right)^2.$$

For each cell-type index t, a heat map of the maximal-intensity projections of $(\rho_t(v))_{1 \leq v \leq V}$ is shown in the rightmost column of the tables of section 5 (N/A indicates that the optimization returned a zero at all voxels for the corresponding type).

2.4   Results

The maximal-intensity projections of correlations and fittings are presented in the tables of section 5 for all cell types in our study. Some striking neuroanatomical patterns appear in the results, both in correlations and in fittings. We only made use of the microarray data and of the AGEA in the above analysis. We compared the results of the computations to the Allen Reference Atlas, and found out that some major cell-types have strongly inhomogeneous correlations and fittings across the brain. So it is interesting to compare the observed neuroanatomical patterns to the metadata available from [8–12], specifying from which regions of the brain the cell samples were taken.

2.4.1   Consistent neuroanatonical patterns

The inspection of the tables of results in section 5 yields several classes of patterns, both in correlations and fittings, that exhibit similarity with neuroanatomical structures, such as cerebral cortex (and subregions thereof), cerebellum, hippocampus, striatum thalamus, and white matter (including *arbor vitae*). We studied a few examples in more details, using sections in addition to maximal-intensity projections.

- **Cortical patterns.** See Figure 1 for a layer-specific class of pyramidal neurons in the cortex [8], and Figure 2 for a class of pyramidal neurons taken from the amygdala and whose fitting coefficients indeed exhibit a pattern similar to the amygdala. The patterns of fitting coefficients are much sparser than those of correlations.

- **Cerebellar patterns.** See Figure 3 for a class of Purkinje cells [8] and Figure 4 for a class of mature oligodendrocytes [11], both extracted from the cerebellum. Their correlation and fitting coefficients are indeed mostly localized in cerebellum.

- **Hippocampal pattern.** See Figure 5 for a class of pyramidal neurons [8] extracted from the hippocampus. Its correlation pattern is highest in the hippocampal region and fitting coefficients are mostly localized in the hippocampal region.

- **Striatal pattern.** See Figure 6 for a class of medium spiny neurons [11] extracted from the striatum. Its correlation and fitting coefficients are indeed mostly localized in the striatum.

- **White-matter pattern.** See Figure 7 for a class of astrocytes [10] extracted from the cortex. Its correlation and fitting coefficients look exhibit a singular pattern that looks like white matter, with the most caudal component corresponding to the *arbor vitae*.

These examples of agreement between the support of the fitting coefficients and the metadata establish the plausibility of the linear model we proposed in Equation 4. Some of the fitting
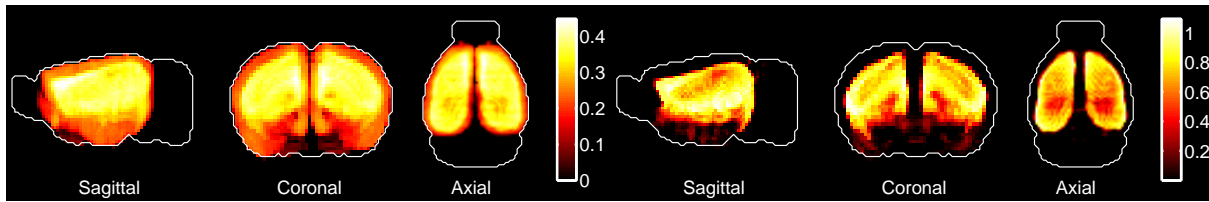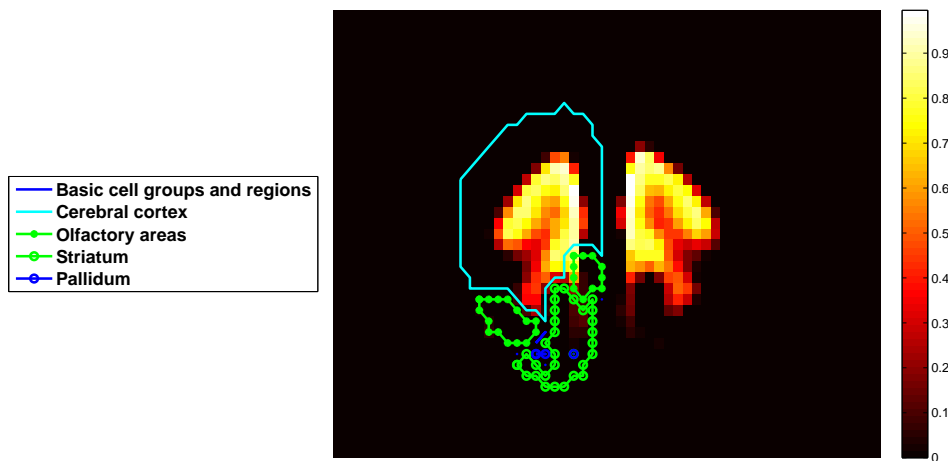
(a) Maximal-intensity projection of correlations.

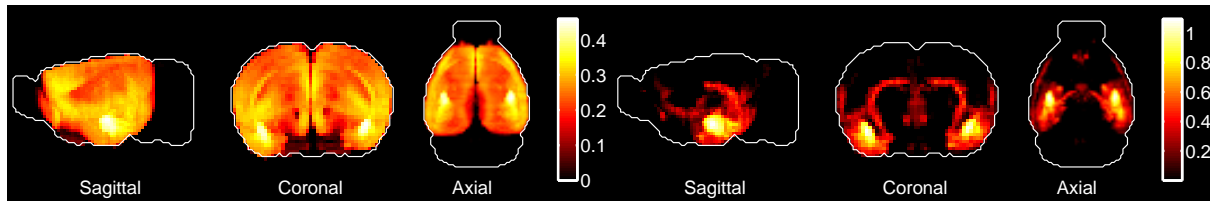(b) Maximal-intensity projection of the linear fitting.
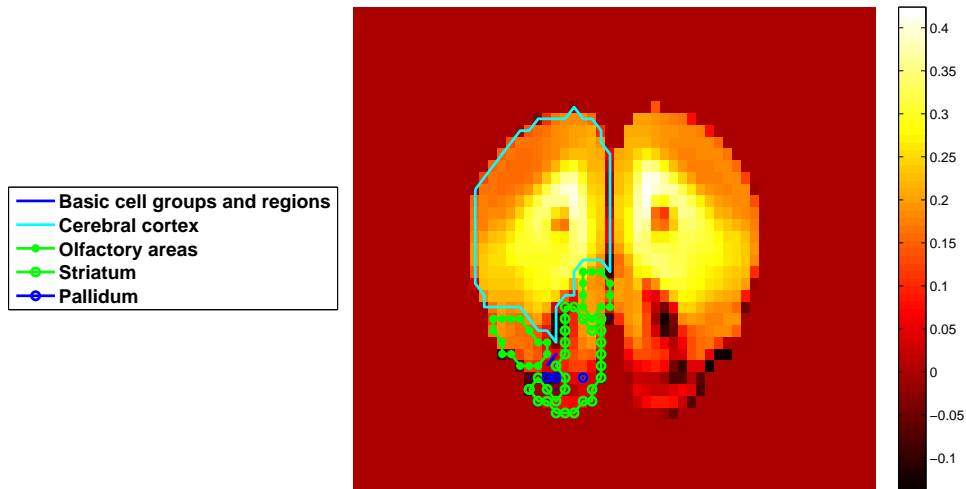
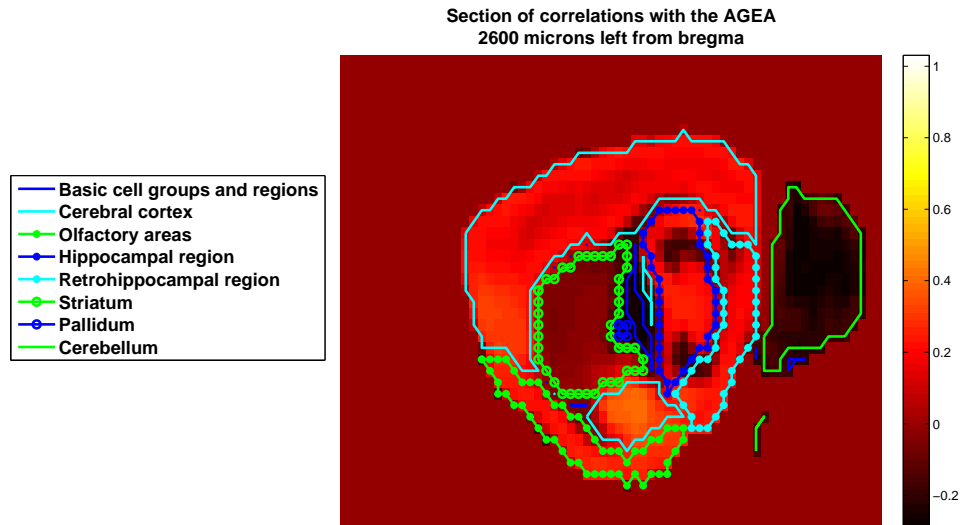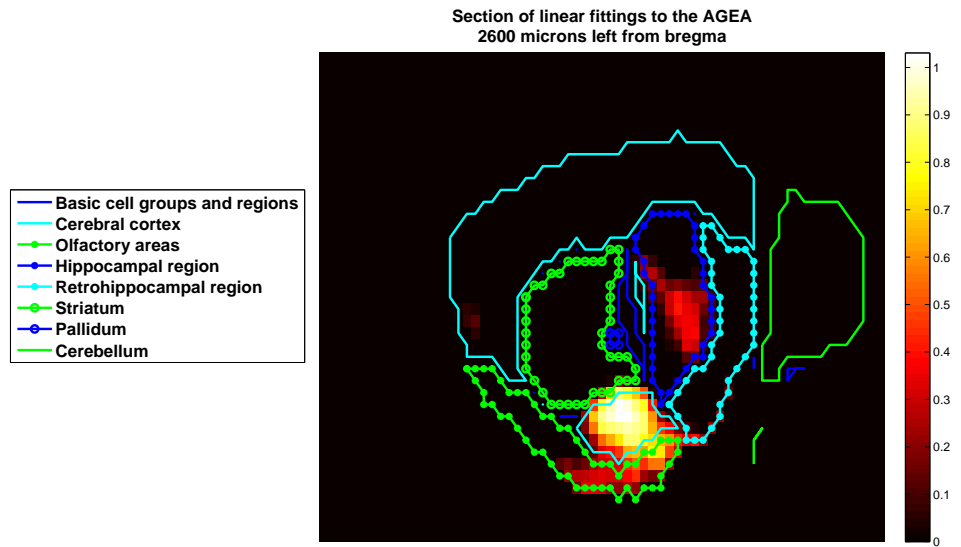(c) A section of correlations, with atlas boundaries

(d) A section of fitting coefficients, with atlas boundaries

Figure 1: Correlations and fitting coefficients between the AGEA and pyramidal neurons taken from primary somatosensory area, layer 5 [8].

(a) Maximal-intensity projection of correlations.
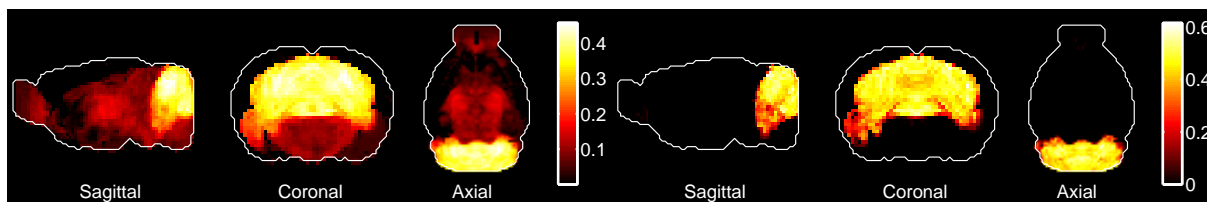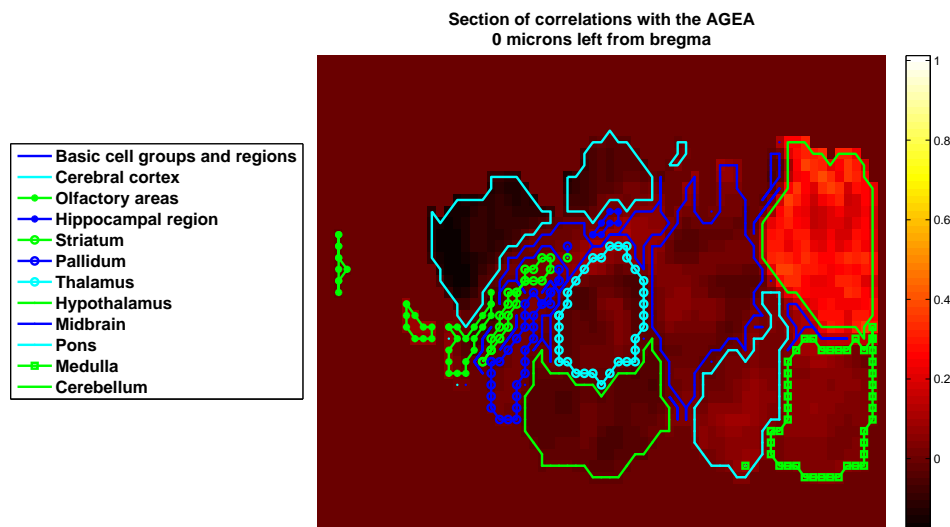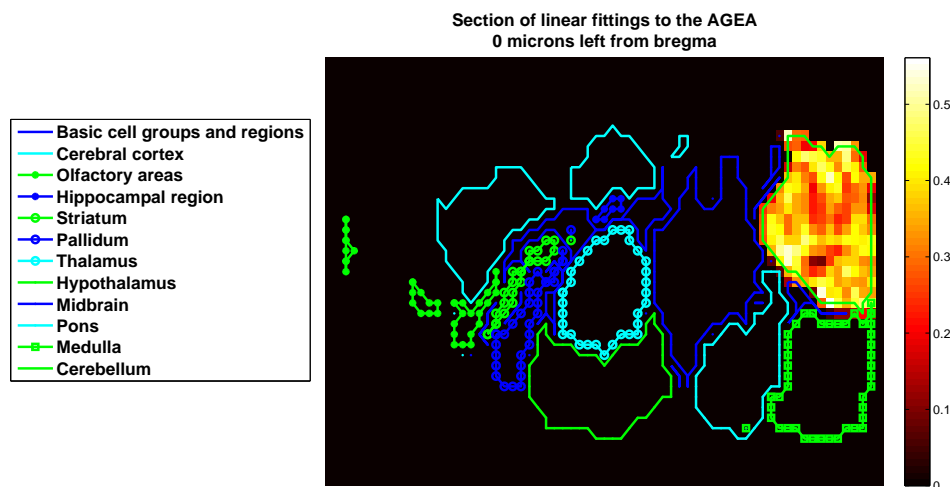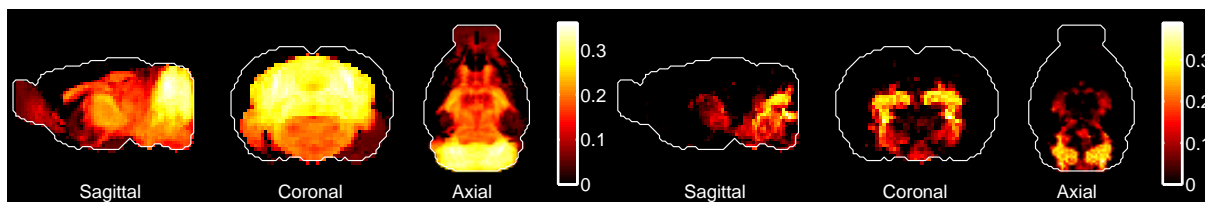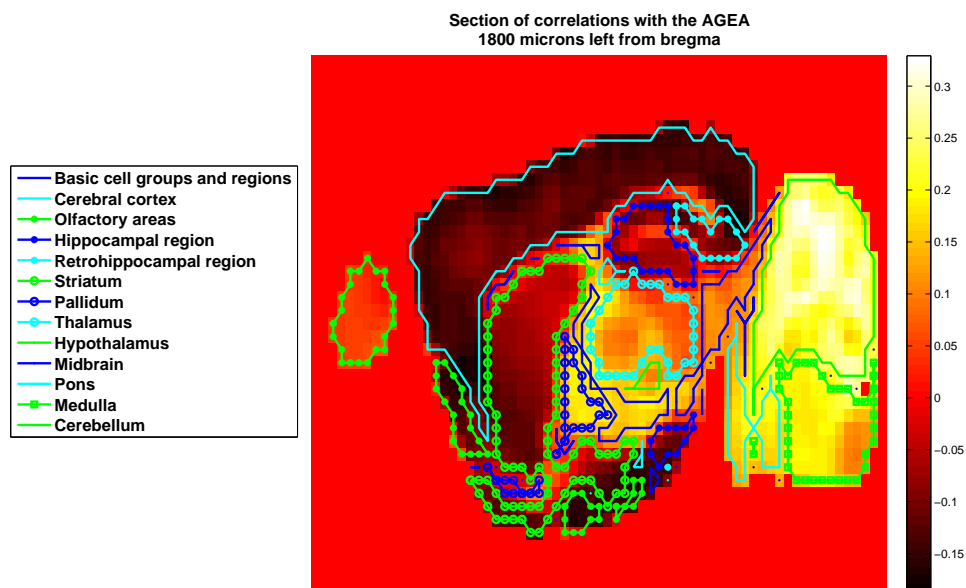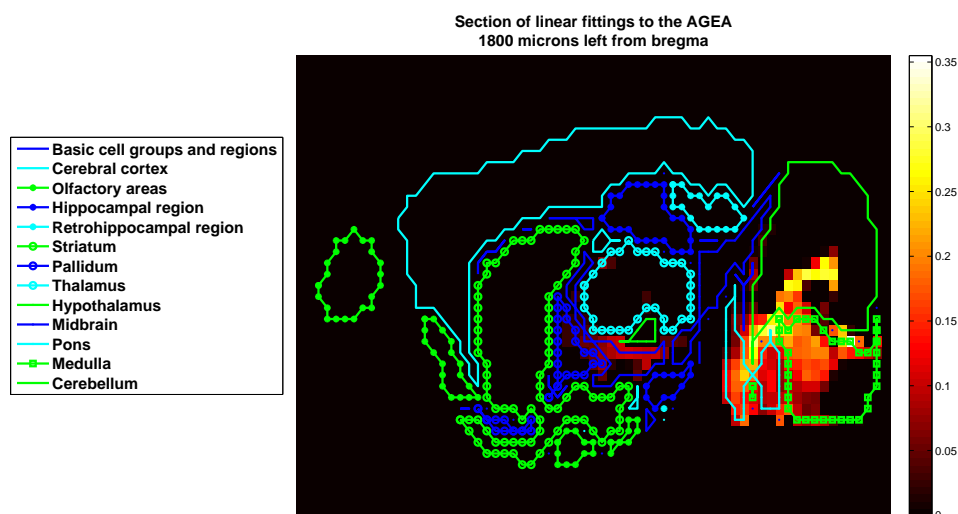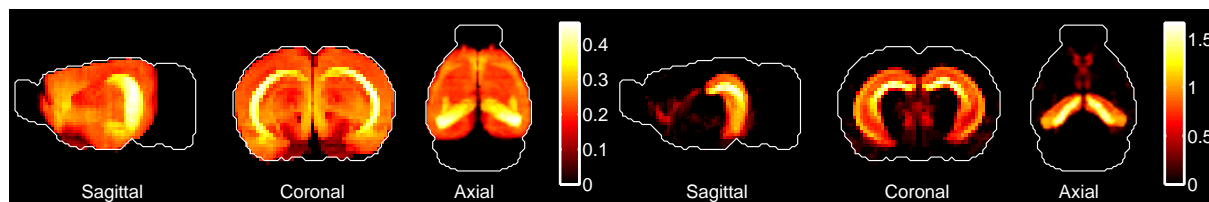
(b) Maximal-intensity projection of the linear fitting.
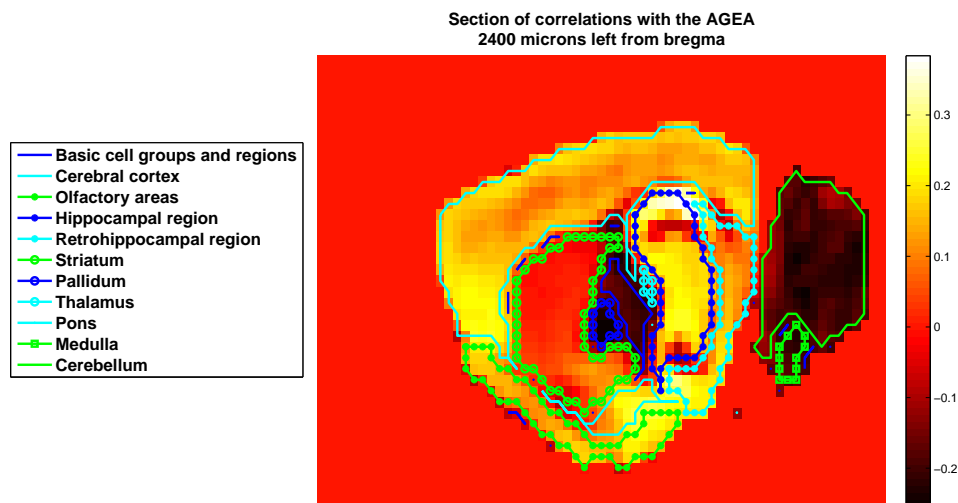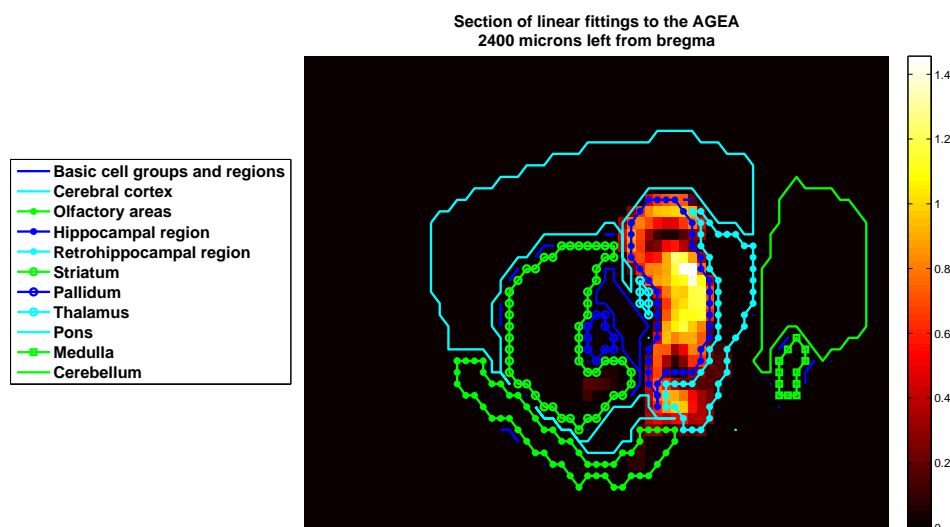

(c) A section of correlations, with atlas boundaries

(d) A section of fitting coefficients, with atlas boundaries

Figure 2: Correlations and fitting coefficients between the AGEA and pyramidal neurons taken from the amygdala (which falls into the cerebral cortex, layer 6B in the Allen Reference Atlas) [8].

(a) Maximal-intensity projection of correlations.

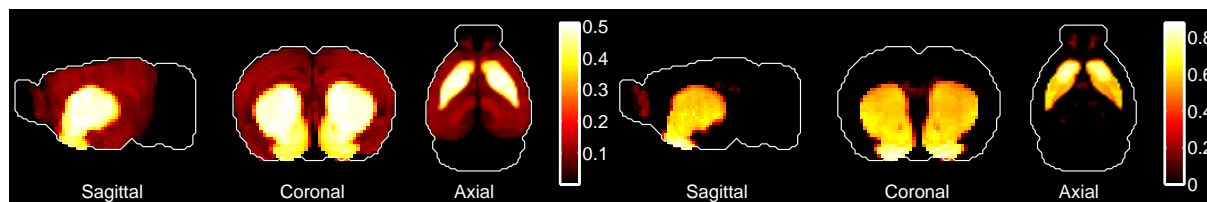(b) Maximal-intensity projection of the linear fitting.

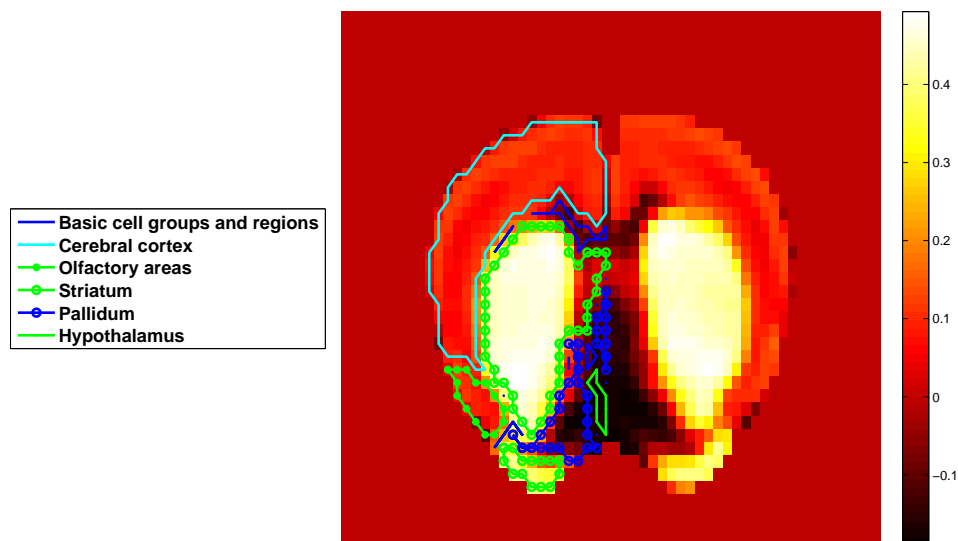(c) A section of correlations, with atlas boundaries

(d) A section of fitting coefficients, with atlas boundaries

Figure 3: Correlations and fitting coefficients between the AGEA and Purkinje cells [9] (index 1 in the table of cell-type results).
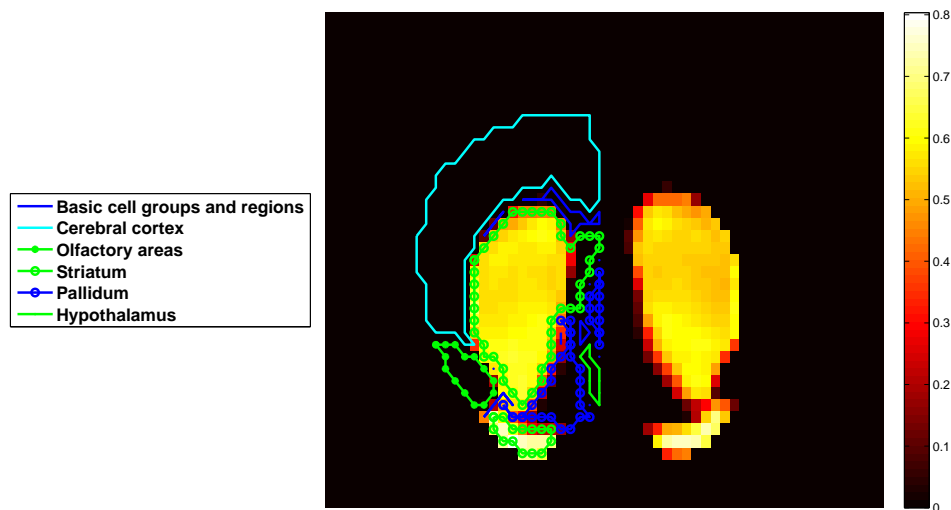
(a) Maximal-intensity projection of correlations.

(b) Maximal-intensity projection of the linear fitting.
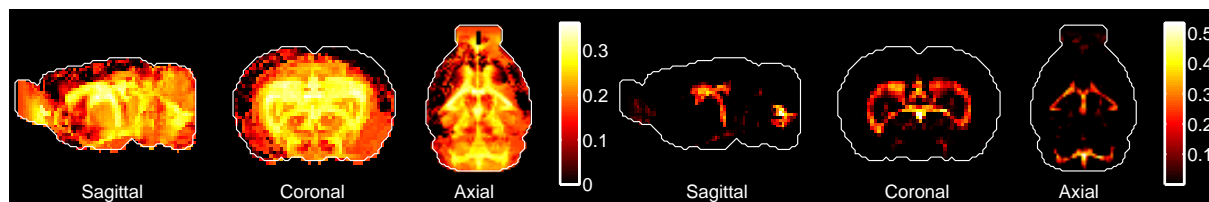

(c) A section of correlations, with atlas boundaries

(d) A section of fitting coefficients, with atlas boundaries

Figure 4: Correlations and fitting coefficients between the AGEA and mature oligodendrocytes extracted from the cerebellum [11] (index 21 in the table of cell-type results).

(a) Maximal-intensity projection of correlations.

(b) Maximal-intensity projection of the linear fitting.

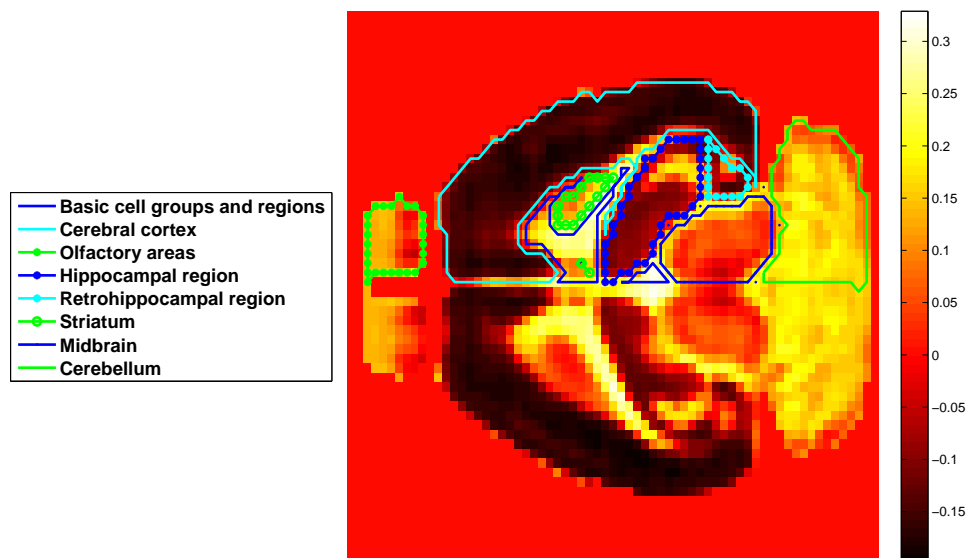(c) A section of correlations, with atlas boundaries

(d) A section of fitting coefficients, with atlas boundaries

Figure 5: Correlations and fitting coefficients between the AGEA and pyramidal neurons taken from the hippocampus (Ammon's horn) [8] (index 49 in the table of cell-type results).

(a) Maximal-intensity projection of correlations.　(b) Maximal-intensity projection of the linear fitting.

(c) A section of correlations, with atlas boundaries

(d) A section of fitting coefficients, with atlas boundaries

Figure 6: Correlations and fitting coefficients between the AGEA and *Drd2* medium spiny neurons taken from the striatum [11], index 16 in the table of cell types.

(a) Maximal-intensity projection of correlations.   (b) Maximal-intensity projection of the linear fitting.

(c) A section of correlations, with atlas boundaries
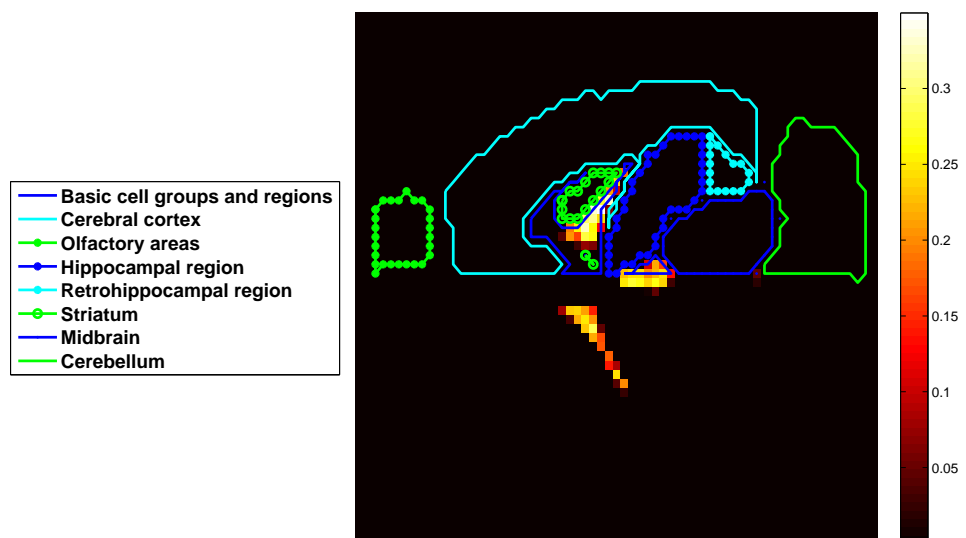
(d) A section of fitting coefficients, with atlas boundaries

Figure 7: Correlations and fitting coefficients between the AGEA and astrocytes taken from the cerebral cortex (Ammon's horn) [10] (index 31 in the table of cell-types).

profiles indeed have a considerably cleaner profile than the correlation, and are able to reveal fine neuroanatomical details The microarray dataset was rich enough to distinguish each of the regions described in the patterns above from the rest of the brain. It is interesting to note that four of the non-cortical patterns described above emerge from bi-clustering of the data at seven biclusters (see section 3 on clustering analysis), if the list of genes is restricted to the most localized genes in the sense of the Kullback–Leibler distance. So the diversity of cell-types found in the microarray dataset reveals partitions of the brains by sets of genes that can be obtained from (a special subset of) the AGEA without using extra data. The major patterns listed above are therefore not a mere accident of the microarray dataset, but reflect a reasonable first list of data-driven neuroanatomical structures.

2.4.2   Unexpected fitting patterns

The results of the model contain some surprises that can be linked to the relative paucity of cell types in the study, compared to the whole diversity of cells in the mouse brain, and to the fact that the cell samples are not distributed uniformly across the brain. Two of the anatomical patterns (thalamus and olfactory areas) that emerge from biclustering of localized genes (see section 3 on clustering analysis) are not as convincingly illustrated by the results of our analysis. A class of Gabaergic Interneurons (index 55, [8]) has fitting coefficients well-localized in the olfactory areas, but the cells were extracted from the somatosensory cortex (primary somatosensory area in the Allen reference Atlas). As the olfactory areas are not represented in our microarray dataset, this gives an indication of the cell types in the dataset that are closest in terms of gene expression to cell-types represented in the olfactory areas. As for the thalamus, it is represented by only one cell type in our microarray dataset (GABAergic interneurons, PV+, index 60, [8]). The cells were extracted from the lateral geniculate complex, and the fitting coefficients are positive in the olfactory areas and in midbrain as well as in the thalamus. Again clustering suggests that some cell types are specific to thalamus, and a richer microarray dataset with better sampling of the thalamus will move some of the fitting coefficients to these cell types.

A few cell-types have surprising fitting coefficients, some of which may be traced to developmental properties.

- **A class of Purkinje cells does not fit to the cerebellum.** See Figure 8 for a class of Purkinje cells (index 52, unpublished) that correlate best with the AGEA both in the thalamus and in the cerebellum, but that fits only in the thalamus. This indicates that thalamus must contain cell types whose gene expression profile is closest to Purkinje cells in the present dataset, but that thalamus is not sampled in enough detail by our microarray dataset for these cell types to be distinguished from this class of Purkinje cells.

- **Some pyramidal neurons, including corticospinal, fit poorly to the AGEA.** Some pyramidal neurons, extracted from the cortex, have very sparse fitting coefficients, even though their correlation profiles with the AGEA have a maximal-intensity projection that looks like the cerebral cortex (see indices 38, 39, 41, 42, 43). It could be that the different non-linearities between numbers of mRNAs, ISH data and microarray data, prevent these
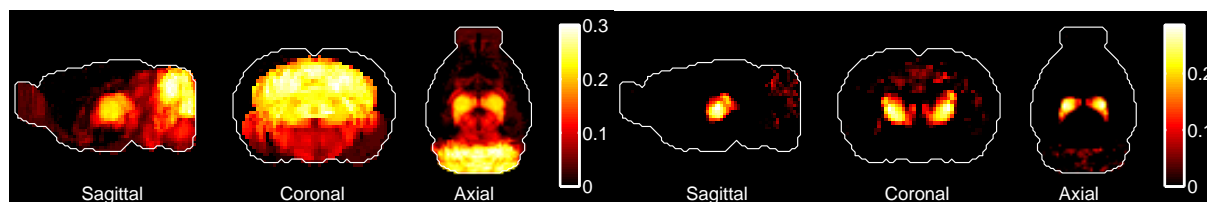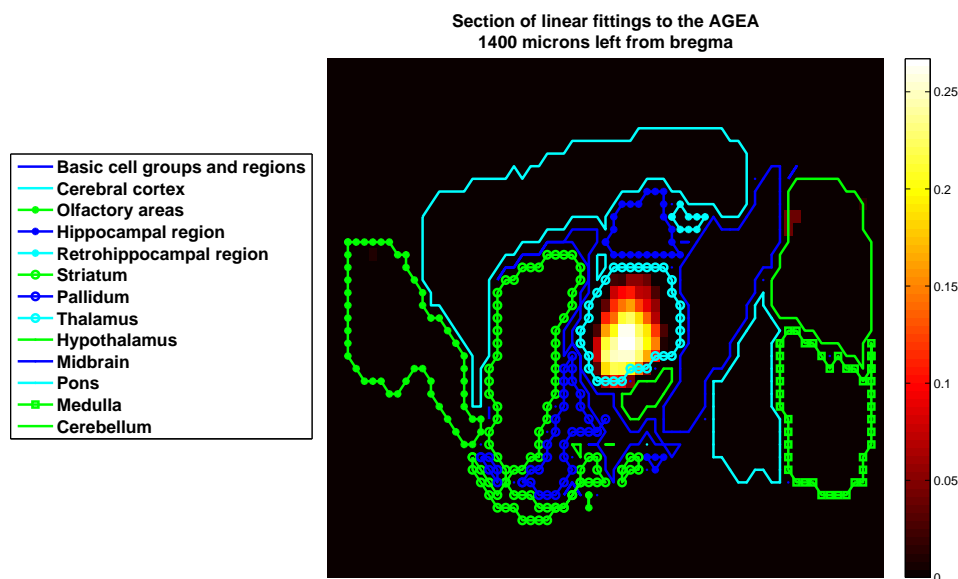
(a) Maximal-intensity projection of correlations.

(b) Maximal-intensity projection of the linear fitting.
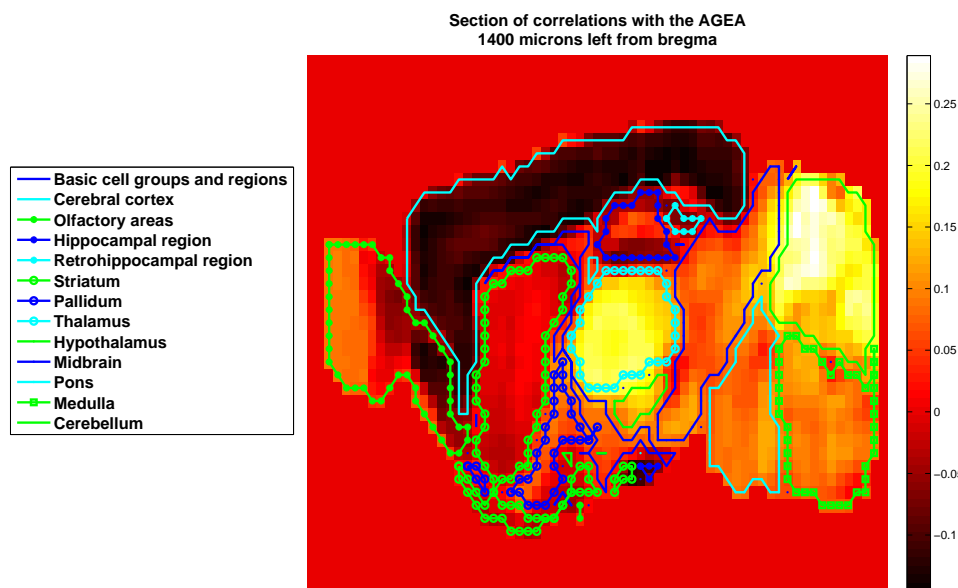
(c) A section of correlations, with atlas boundaries

(d) A section of fitting coefficients, with atlas boundaries

Figure 8: Correlations and fitting coefficients between the AGEA and a class of Purkinje cells extracted from the cerebellum (index 52 in the table of cell-type results). The correlation profile is high both in the cerebellum and in the thalamus, but the fitting coefficients are localized in

cell types to be correctly detected by our linear model. However, one can notice that the ages of the mice are P3, P6, P14 for these classes of pyramidal neurons. The poor fitting to the AGEA could therefore come from a developmental effect, with corticospinal pyramidal neurons and callosally-projecting pyramidal neurons maturing late, while the AGEA corresponds to the adult mouse brain.

3   Clustering analysis

3.1   Biclustering of genes and voxels

Some of the cell types for which we have microarray data are predicted by the linear model 4 to be very localized in the brain. However, the validity model is limited by the number of cell types in the study. The T different cell types only span a T-dimensional subspace of gene space, while the AGEA contains data for thousands of genes. However, in the absence of microarray data for more cell types, we can still ask if the Allen Gene Expression Atlas has some sets of genes whose expression is strongly localized in some sets of voxels. The sets of voxels would be the brain regions in which some cell types are expected to be localized, and the sets of genes would be the genes whose expression is highest in those cell types.

Mathematically, this problem is a biclustering problem: we need to partition the set of voxels *and* the set of genes in the AGEA in an optimal way given the gene expression energies. Gene expression energies can be used to turn voxels and genes into a bipartite graph, and this graph can be partitioned using a biclustering algorithm.

3.1.1   From the AGEA to a bipartite graph

The AGEA can be mapped to a *weighted bipartite graph* in the following way:

- the first set of vertices consists of voxels, numbered from 1 to V,

- the second set of vertices consists of genes, numbered from 1 to G,

- each of the edges connect one voxel to one gene, and has a weight given by the expression energy of the gene at the voxel. If the expression energy $E(v, g)$ is zero, there is no edge between voxel v and gene g.

We looked for partitions of this weighted bipartite graph into subgraphs such that the weights of the internal edges of the subgraphs are strong compared to the weights of the edges between the subgraphs. This is the isoperimetric problem addressed by the algorithm of [23] (the graph need not be bipartite to apply this algorithm).

Given a weighted graph, the algorithm cuts some of the links, thus partitioning the graph into a subset S and its complementary $\bar{S}$, such that the sum of weights (in this case expression energies) in the set of cut edges is minimized relative to the total weight of internal edges in

S. The sum of weights in the set of cut edges is analogous to a boundary term, while the total weight of internal edges is analogous to a volume term. In that sense the problem is an isoperimetric optimization problem, and the optimal set S minimizes the *isoperimetric ratio* ρ over all the possible subgraphs:

$$S = \mathrm{argmin}_{\mathrm{Vol}(s) \leq \mathrm{Vol}(\bar{s})}\, \rho(s), \tag{5}$$

$$\rho(s) := \frac{|\partial s|}{\mathrm{Vol}(s)}, \tag{6}$$

$$|\partial s| = \sum_{i \in s, j \in \bar{s}} W_{ij}, \tag{7}$$

$$\mathrm{Vol}(s) = \sum_{i \in s, j \in s} W_{ij}, \tag{8}$$

where the quantity W_{ij} is the weight of the link between vertex i and vertex j. Once S has been worked out, the algorithm can be applied to S and its complementary $\bar{S}$. This recursive application goes on until the isoperimetric ratio reaches a stopping ratio, representing the highest allowed isoperimetric ratio. This value is a parameter of the algorithm. Rising it results in a higher number of clusters, as it rises the number of acceptable cuts.

3.1.2 Restriction of the Allen Atlas to the most localized genes

We studied the localization properties of the genes in the Allen Gene Expression Atlas by comparing their expression energies to a uniform function over the brain. We computed the Kullback-Leibler divergence of each gene from a uniform distribution:

$$KL(g) = -\sum_{v=1}^{V} \frac{1}{V} \log\left(V E_{\mathrm{prob}}(v, g)\right), \tag{9}$$

$$E_{\mathrm{prob}}(v, g) = \frac{E(v, g)}{\sum_{v=1}^{V} E(v, g).} \tag{10}$$

where V is the total number of voxels in the brain and E_{prob} is a normalized version of E such that the expression energy of each gene is a probability density over the volume of the brain.

We took the highest 5 percentiles of genes by KL-divergence, and thresholded the expression energies, thus reducing the number of voxels. These are the most localized genes by that measure, so solving the ispoerimetric problem on the graph constructed with those genes only will be more likely to produce the localized biclusters we are looking for. Thresholding the expression value gets rid of many voxels with low expression that would be (uninteresting) optimal subgraphs of their own because they are hardly connected to any genes.
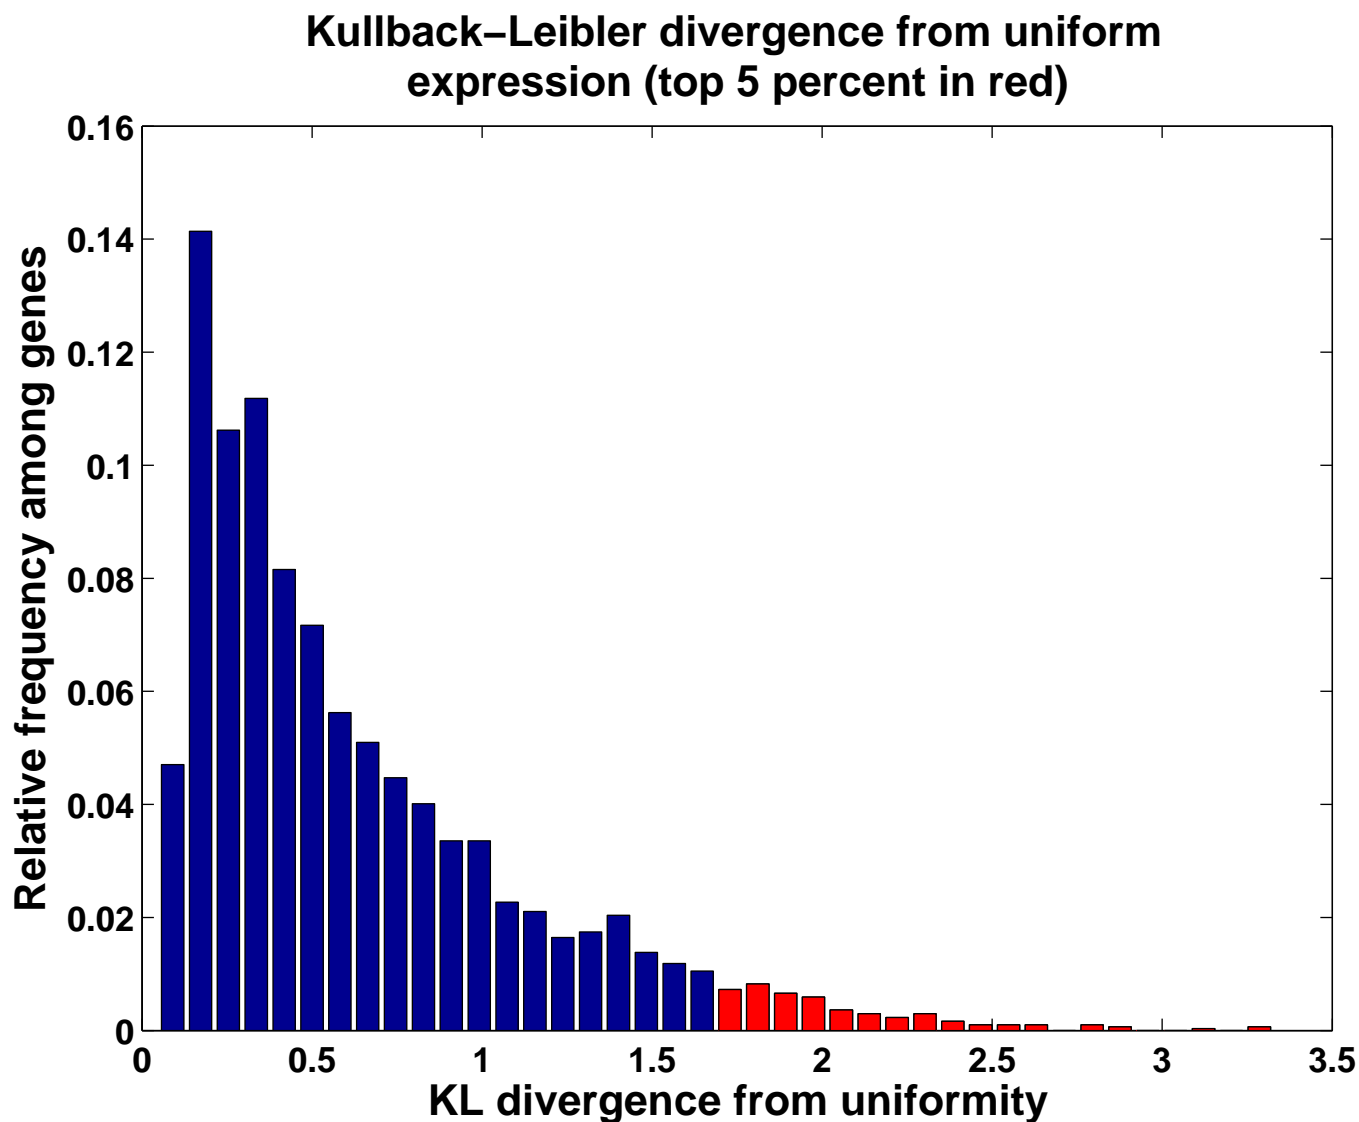
Figure 9: A histogram of the occurence of Kullback–Leibler divergences from a uniform expression across 3041 genes. The top 5 percentiles of KL divergences are highlighted in red. They consist of the genes selected for biclustering analysis. The names of those genes appear in the tables describing each of the biclusters.

3.2 Results

We illustrate biclusters corresponding to the plateau at $B = 7$ biclusters in the number of bi-cluster as a function of the stopping ratio.

For each of the biclusters, the algorithm returns a set of voxels and a set of genes. Let is call the set of voxels $\mathcal{V}_b$ and the set of genes $\mathcal{G}_b$. For each of the biclusters, we plotted:

- the voxel profile of the bicluster: the maximal-intensity projections of the set of voxels $\mathcal{V}_b$ in the bicluster;

- the expression profile of the bicluster: the maximal-intensity projection of the the sum of the expression energies of the set of genes $\mathcal{G}_b$ in the bicluster.

The expression profiles follow the shapes of the sets of voxels quite nicely. For each gene, we computed the fraction of the total expression of the bicluster that it represents.
Moreover, we noticed many of the biclusters look similar to brain regions defined in the 'Big 12' partition of the left hemisphere in the Allen Reference Atlas [6], illustrated on Table 1. In order to assess this resemblance quantitatively, we took the Allen Reference Atlas of the left hemisphere and for each region R, we computed the fraction of the voxels in $\mathcal{V}_b$, and the fraction of the expression energy that is contained in the region. These quantities are expressed as:

$$\text{Fraction of voxels}(R, b) = \frac{|\mathcal{V}_b \cap R|}{|R|}, \text{ for brain region } R \text{ and bicluster } b, \tag{11}$$

$$\text{Fraction of expression}(R, b) = \frac{\sum_{v \in \mathcal{V}_b \cap R} \sum_{g \in \mathcal{G}_b} E(v, g)}{\sum_{v \in \mathcal{V}_b} \sum_{g \in \mathcal{G}_b} E(v, g)}, \text{ for brain region } R \text{ and bicluster } b. \tag{12}$$

Only regions that have non-zero overlap with the biclusters are shown in tables. These regions are ordered by decreasing value of the fraction of voxels. The biclusters are described by the following figures:

- **Bicluster 1: 'cerebellum-like'**, see Figure 10.

- **Bicluster 2: 'white-matter like'**, see Figure 11.

- **Bicluster 3: 'hippocampus-like'**, see Figure 12.

- **Bicluster 4: 'olfactory-like'**, see Figure 13.

- **Bicluster 5:** a disconnected bunch with medulla and other regions, see Figure 14.

- **Bicluster 6: 'thalamus-like'**, see Figure 15.

- **Bicluster 7: 'striatum-like'**, see Figure 16.
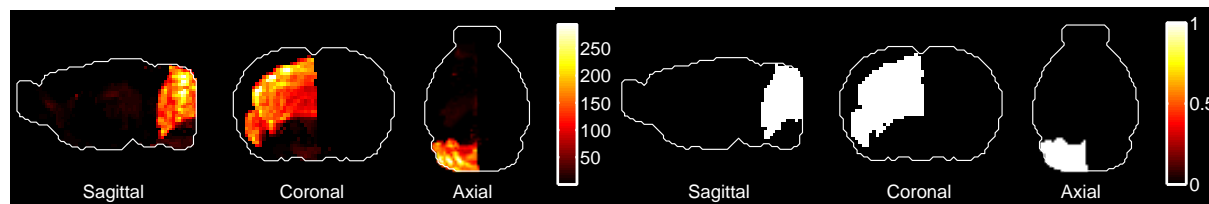
(a) Sum of gene-expression energies in the cerebellum-like bicluster

(b) The set of voxels in the cerebellum-like bicluster

GeneNames	Fraction of cluster energy (pct)
Il16	19.04
Neurod1	17.74
Cnksr3	14.48
Slc6a5	9.35
Gng13	6.55
3110001A13Rik	6.44
Gabra6	6.19
Col18a1	3
Plxdc1	2.45
En2	2.14
AW049765	2.12
B3gnt5	2.07
2010106G01Rik	1.96
Lhx5	1.64
Htr5b	1.44
Ptprm	1.16
Slc22a3	0.72
Zfp423	0.63
Gbx2	0.33
Ghrh	0.27
Nr5a1	0.26

(c) The list of genes in the cerebellum-like bicluster, ordered by decreasing expression

Brain region	Percentage of cluster
Cerebellum	96.58
Medulla	2.87
Basic cell groups and regions	0.56

(d) The set of voxels in the cerebellum-like bicluster

Figure 10: Bicluster 1: 'cerebellum-like'.
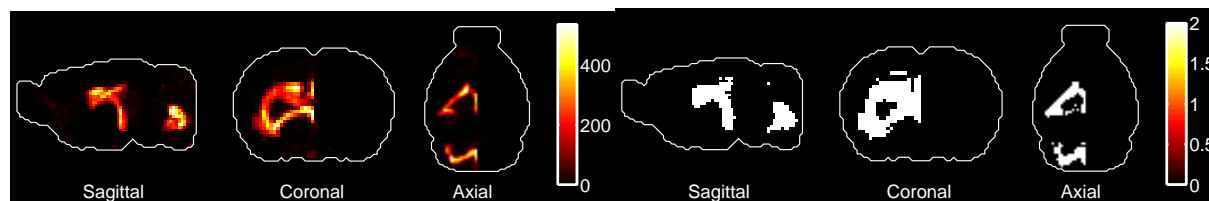
(a) Sum of gene-expression energies in bicluster 2

(b) The set of voxels in bicluster 2

GeneNames	Fraction of cluster energy (pct)
Kl	11.22
Clic6	9.24
Ace	8.33
Cab39l	6.34
F5	5.99
Tgfb2	5.88
Rdh5	5.4
Prlr	4.69
AI987712	4.58
Acaa2	4.37
E030013G06Rik	4.31
Tcn2	3.96
Gm967	2.44
Frmpd2	2.36
Trpv4	2.34
Cd59a	2.2
LOC436099	2.14
Vil2	2.12
Sntb1	2.07
Fzd4	2.06
6820408C15Rik	1.92
D9Ertd280e	1.71
Thbs4	1.68
B3gat2	1.45
Itpkb	1.2

(c) The list of genes in bicluster 2 ordered by decreasing expression

Brain region	Percentage of cluster
Basic cell groups and regions	36.85
Medulla	15.61
Hippocampal region	12.86
Cerebral cortex	12.43
Striatum	9.39
Cerebellum	6.79
Thalamus	5.49
Midbrain	0.29
Pons	0.29

(d) The set of voxels in bicluster 2

Figure 11: Bicluster 2: 'white-matter like'. The pattern is disconnected and looks like white matter, with the most caudal component coinciding with the *arbor vitae*. Some of the white matter is annotated by the Allen reference Atlas as 'Basic cell groups and regions'.
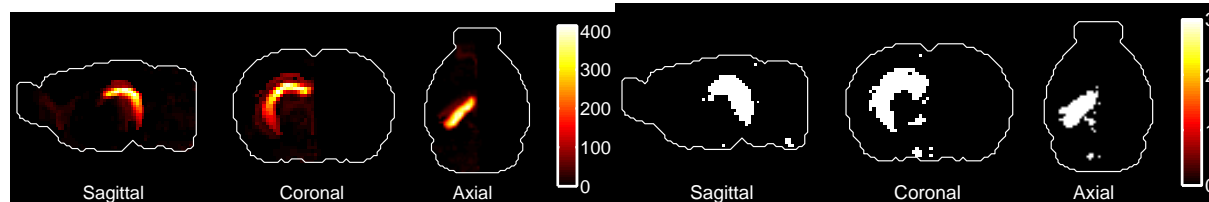
(a) Sum of gene-expression energies in bicluster 3

(b) The set of voxels in bicluster 3

GeneNames	Fraction of cluster energy (pct)
C1ql2	11.64
TC1412430	7.35
Zbtb20	6.48
Sipa1l2	6.27
C630041L24Rik	5.79
Dsp	5.48
Crlf1	5.46
A330019N05Rik	5.25
Slc39a6	5.12
C78409	5.07
Tnfrsf25	4.71
Cyp7b1	4.36
Lct	3.71
Klk8	3.31
Gpc4	2.94
Prox1	2.66
Nr3c2	2.47
Pkp2	2.47
Arl15	2.16
Itga7	1.82
Fat4	1.49
Sema5a	1.47
Csf2rb2	1.37
Rnf19	1.14

Brain region	Percentage of cluster
Hippocampal region	80.39
Thalamus	9.48
Midbrain	3.59
Basic cell groups and regions	2.94
Medulla	1.8
Cerebral cortex	1.47
Retrohippocampal region	0.16
Hypothalamus	0.16

(d) The set of voxels in bicluster 3

(c) The list of genes in bicluster 3 ordered by decreasing expression

Figure 12: Bicluster 3: 'hippocampus-like'.
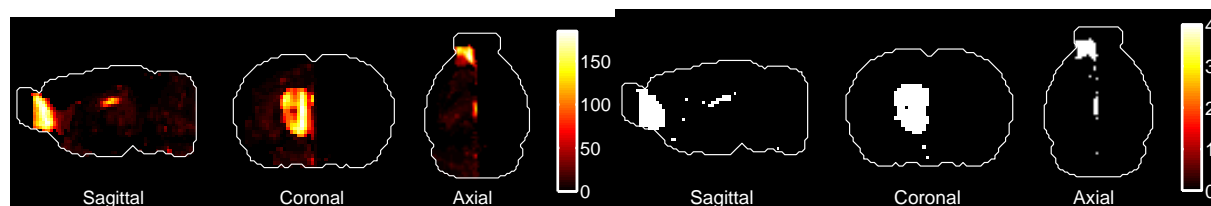
(a) Sum of gene-expression energies in bicluster 4

(b) The set of voxels in bicluster 4

GeneNames	Fraction of cluster energy (pct)
Lrrc55	10.29
Stard8	8.21
9830123M21Rik	6.79
Nmb	6.6
C230040D10Rik	5.89
Kcnb2	5.84
4930589M24Rik	5.45
Adamts19	5.26
Raver2	4.46
Syt6	4.35
Chrna3	3.45
Cdh23	3.01
Sp8	3
B930011P16Rik	2.75
Nmbr	2.62
A230065H16Rik	2.31
Vipr2	2.17
Jam2	2.12
Eomes	1.95
B430201A12Rik	1.89
mCG1049722.1	1.66
Atp6v1c2	1.61
Rgnef	1.42
Atp10a	1.4
Dnahc11	1.19
Scube2	1.17
AW456874	1.08
Tmem16a	1.05
Lhx8	1.02

(c) The list of genes in bicluster 4 ordered by decreasing expression

Brain region	Percentage of cluster
Olfactory areas	86.68
Cerebral cortex	7.38
Thalamus	3.48
Basic cell groups and regions	1.64
Striatum	0.61
Medulla	0.2

(d) The set of voxels in bicluster 4

Figure 13: Bicluster 4: 'olfactory-like'
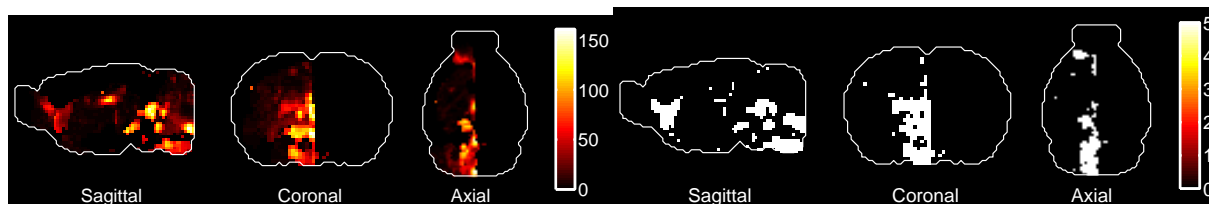
(a) Sum of gene-expression energies in bicluster 5

(b) The set of voxels in bicluster 5

GeneNames	Fraction of cluster energy (pct)
Glra1	22.11
Slc18a3	5.75
Slc18a2	4.86
Slc5a7	4.78
Ddc	4.39
Slc6a3	3.79
1810023C24Rik*	3.48
Chrna6	3.24
Pou4f1	3.09
LOC244958	2.36
Chrnb4	2.26
Gpr151	2.24
Tph2	2.24
Avp	2.23
Calca	2.19
Anxa2	2.07
Frmd6	2.05
C130034I18Rik	1.98
Slc6a4	1.85
Epha1	1.79
Postn	1.69
Stk24	1.66
Il13ra1	1.64
C130021I20Rik	1.63
Calcb	1.61
A330102H22Rik	1.51
Dbh	1.27
Ppap2a	1.26
Gpr3	1.22
Hspb1	1.15
Tal1	0.98
Layn	0.97
Cdh6	0.95
Tspan12	0.93
Wif1	0.86
Scml2	0.72
Pscdbp	0.48
Gucy2c	0.41
Ntsr1	0.3

(c) The list of genes in bicluster 5 ordered by decreasing expression

Brain region	Percentage of cluster
Medulla	47.11
Midbrain	16.33
Pons	14.46
Cerebral cortex	8.5
Olfactory areas	6.8
Basic cell groups and regions	4.08
Thalamus	1.53
Striatum	0.51
Cerebellum	0.51
Hypothalamus	0.17

(d) The set of voxels in bicluster 5

Figure 14: Bicluster 5: a disconnected bunch with medulla and other regions.
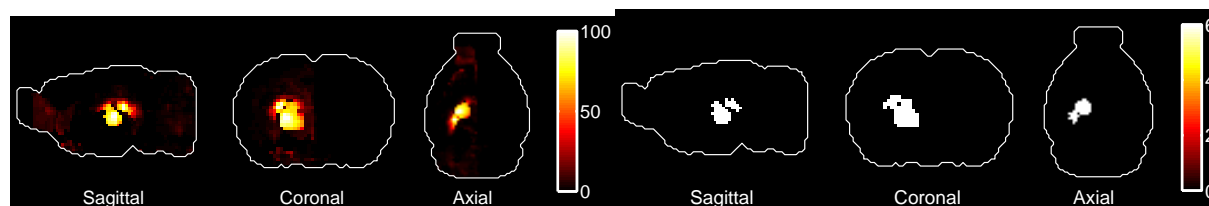
(a) Sum of gene-expression energies in bicluster 6   (b) The set of voxels in bicluster 6

GeneNames	Fraction of cluster energy (pct)
Plekhg1	26.95
1110069I04Rik*	22.71
Tnnt1	16.72
Lef1	14.68
Rab37	14.08
Slitrk6	4.86

Brain region	Percentage of cluster
Thalamus	98.4
Basic cell groups and regions	0.53
Hippocampal region	0.53
Hypothalamus	0.53

(d) The set of voxels in bicluster 6

(c) The list of genes in bicluster 6 ordered by decreasing expression

Figure 15: Bicluster 6: 'thalamus-like'.

4   Dimensionality estimates of brain regions in gene space

4.1   Scaling argument and algorithm

Due to the large number of genes that are analyzed in the AGEA atlases, our data live in very high-dimensional space. On the other hand, many genes have non-zero expression in the gene-expression profiles of cell-types. So the multiplicity of genes expressed may not reflect the complexity of the underlying biological objects.

A fixed cell-type is just a one-dimensional subspace of gene space. From the inspection of the projection of the SVD of gene-expression energies onto the subspace spanned by the first three singular vector in gene space, there seems to be clouds of voxels that are localized on low-dimensional subspaces. How can we detect such localization properties? The following scaling argument has been used to discover the intrinsic dimension of datasets that are only known by their pairwise dissimilarities [25].
If the points are drawn from an underlying submanifold of the ambient space, of dimension d, then a cone with tip at the center of the cloud of points and an opening of 90 degrees should intercept a fraction of the points that decreases exponentially with the dimension d. This is illustrated on Figure (17) for a disc embedded into three dimensions.

For a given region R_r, the dimension of the underlying subspace of gene space is estimated by the following procedure:
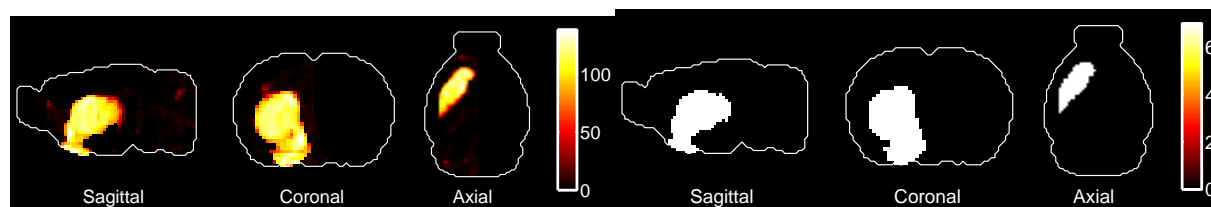
(a) Sum of gene-expression energies in bicluster 7

(b) The set of voxels in bicluster 7

GeneNames	Fraction of cluster energy (pct)
Rgs9	25.68
Pde1b	20.98
Adora2a	15.7
Rarb	12.92
Gprin3	8.32
Gpr6	8
Serpina9	5.72
Cd4	2.69

(c) The list of genes in bicluster 7 ordered by decreasing expression

Brain region	Percentage of cluster
Striatum	95.23
Basic cell groups and regions	1.87
Cerebral cortex	1.73
Pallidum	1.18

(d) The set of voxels in bicluster 7

Figure 16: Bicluster 7: 'striatum-like'.
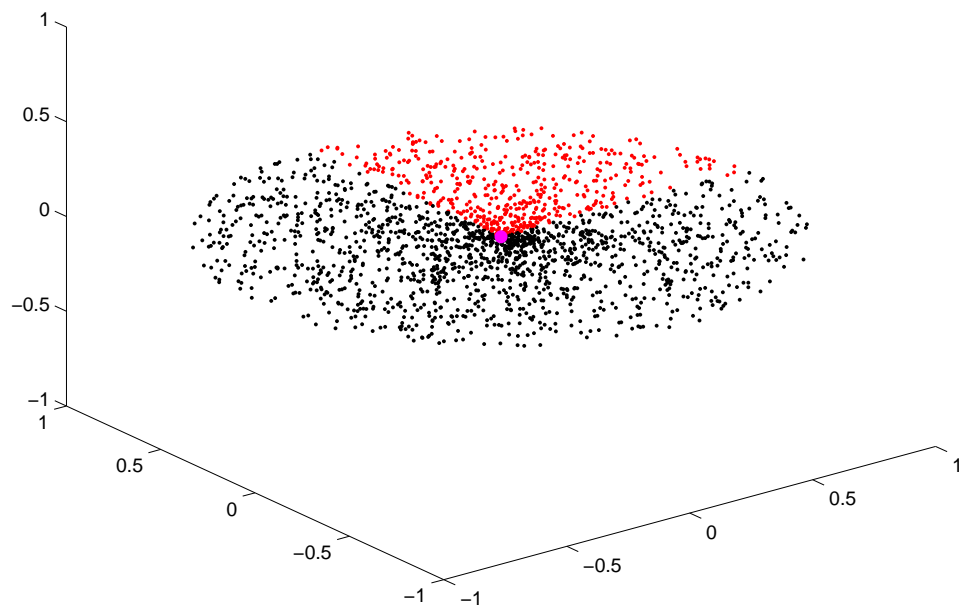
Figure 17: **A disc sampled by 2000 random points with uniform probability, embedded into a three-dimensional space.** A cone with tip at the center of the disc (shown in magenta), with an opening of 90 degrees, intercepts roughly a quarter of the points. This fraction would be the same if the disc was embedded into a higher-dimensional space.

1. Compute the pairwise distances in the region. Compute the pairwise distances between the voxels $(v_i)_{1,\leq r \leq V_r}$ belonging to region R_r (these are the relevant lines of the data matrix D):

$$\forall i \in [1..V_r], \ v_i := (D_{ig})_{1 \leq g \leq G} \in \mathbf{R}^G,$$

$$\forall i, j \in [1..V_r], \ d_{ij} := ||v_i - v_j||_{L^2}.$$

2. Find the center of the region. Find out the center C of this cloud defined as the point *in the set* which minimizes the dispersion of points around itself. This is as close to the average as we get, within the set (this is to pathologies with very hollow, non-convex regions, where the ordinary mean may be far from all the voxels). Let us call c the index of the center, meaning $C = v_c$ with:

$$c = \operatorname{argmin} \sum_j d_{cj}^2.$$

This is a Fréchet mean with uniform weight in the image of the region R_r under D.

3. Define a typical cone in the region. Pick a point P (with index p, or $P = v_p$) in the region, which is at one standard deviation away from the center (or as close to it as you can get while staying within the region):

$$p = \operatorname{argmin} \left(d_{cp}^2 - \sigma_d^2 \right),$$

$$\sigma_d^2 = \sum_j d_{cj}^2.$$

Consider the cone whose tip is at C, whose axis points towards P, and whose opening angle is $\pi/2$.

3. Compute the number of voxels in the region that fall within the cone. For a point M with index m, the property of being inside the cone is defined in terms of its distances to C and P, as the critical distance from P is attained [2] when the angle $\widehat{PCM}$ equals $\pi/4$:

$$d_{mp} < \sqrt{d_{pc}^2 + d_{mc}^2 - \sqrt{2} d_{pc} d_{mc}}.$$

We can compute the fraction of points in the cloud that is inside the cone:

$$f_{C,P} = \frac{1}{V} \left| \left\{ m \in [1..V], d_{mp} < \sqrt{d_{pc}^2 + d_{mc}^2 - \sqrt{2} d_{pc} d_{mc}} \right\} \right|.$$

4. The fraction of points intercepted by the cone scales exponentially with the dimension of the underlying subspace. If the points are taken from a δ-dimensional manifold which is not too anisotropic (see below for a quantitative study of the dispersion of the

---

[2]This is obtained by Pythagoras' theorem in the triangle CMP: the angle $\widehat{PCM}$ is smaller than a chosen angle β (here $\beta = \pi/4$), if its cosine is larger than $\cos \beta$. Pythagoras' theorem applied to the triangle PCM expresses this condition as $\frac{d_{pc}^2 + d_{mc}^2 - d_{mp}^2}{2 d_{pc} d_{mc}} < \cos \beta$, i.e. $d_{mp} < \sqrt{d_{pc}^2 + d_{mc}^2 - 2 \cos \beta d_{pc} d_{mc}}$

values that are returned by this algorithm for different choices of cones), this fraction should scale as

$$f_{C,P} \simeq 2^{-\delta},$$

which gives a global estimate of the dimension of the (putative) manifold underlying the cloud of points:

$$\delta_P = -\frac{\log f_{C,P}}{\log 2}.$$

The results are quite small (typically below ten) for all the structures in the 'Fine' annotation of the Allen Reference Atlas [6], but should *not* be compared directly to G, as the output of the calculation described above is constrained by the sample size, and many regions have quite small size.

4. Probabilistic interpretation of the dimensions. Repeat the above computation for all points P distinct from the center C. This induces a distribution of estimates of the dimension for the cloud of points, with a value of δ_P associated to each point P. The more peaked this distribution is, the better the cloud of point is localized on a manifold. The value of the peak estimates the dimension of this manifold.

$$\mathrm{Prob}(\delta = \delta_0) = \mathrm{Prob}(f_{C,P} = 2^{-\delta_0}V) = \frac{1}{V-1}\left|\left\{P, f_{C,P} = 2^{-\delta_0}V\right\}\right|. \tag{13}$$

The values of the dimension are binned into integer values. If the brain structure considered is 1) well-localized and 2) reasonably homogeneous, the plot of the probability distribution should be 1) well-separated from a Dirac mass placed at the critical dimension of the structure, and 2) be peaked. Results for caudoputamen are shown on Figure 18.

4.2 Results and limitations

What is the maximal value of the dimension that can be found using the conical trick described above? This value, which we will call *critical dimension* is attained in situations where the cloud of points is so dilute that the cone is essentially empty: it intercepts no point in the region apart from its tip C and the point P that defines the orientation of the cone. This critical dimension, or δ_{crit}, is therefore defined by

$$2^{-\delta_{\mathrm{crit}}} = \frac{2}{V},$$

hence

$$\delta_{\mathrm{crit}} = \frac{\log V}{\log 2} - 1.$$

When a cloud of V points has critical dimension, it must therefore be considered to be a genuine cloud of points, living in high dimension, rather than a sample taken from a lower-dimensional space.

Of course, the critical dimension is a decreasing function of the sample size, and at the resolution of 200 microns we are working with, the critical dimensions of the brain structures are all orders of magnitude below the total number of genes $G = 3041$, which is a big threat to
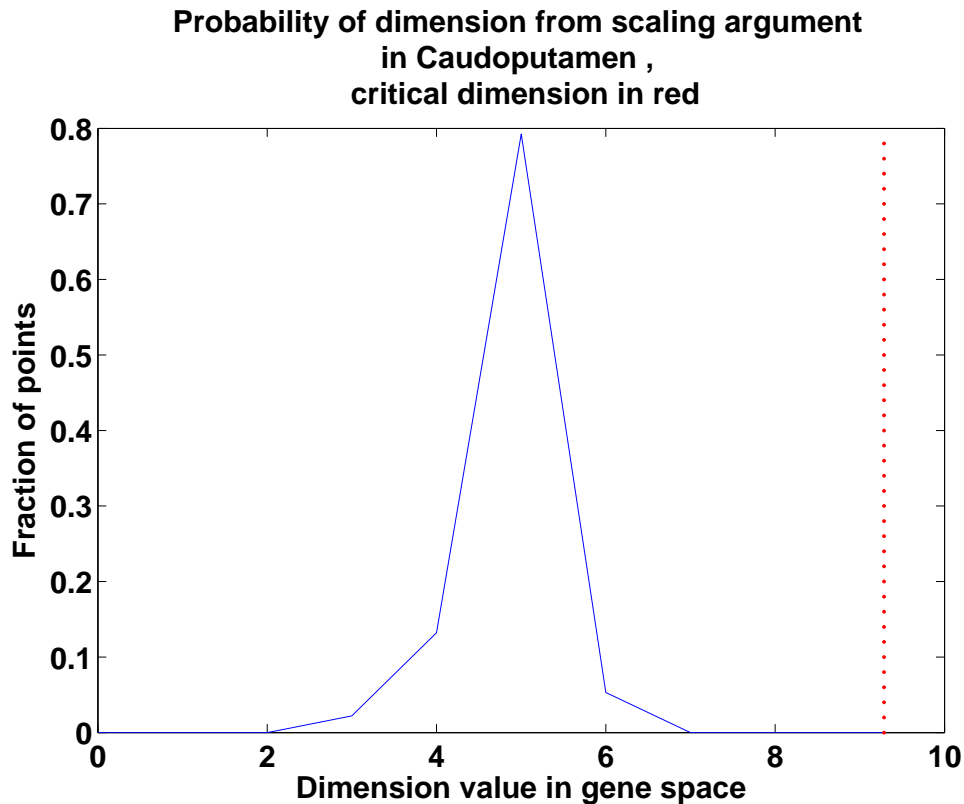
Figure 18: Distribution of the dimension estimates for the voxels belonging to caudoputamen. At a resolution of 200 microns, caudoputamen consists of 1248 voxels, which induces a value of 9.28 for the critical dimension, the maximal dimension that can be returned by the algorithm. 80 percent of the voxels correspond to an underlying dimension of 5, and the probablity of the critical dimension is zero.

our method: as the value of δ fluctuates with the orientation of the cone, it may well be that cones that are close in terms of Euclidean distances induce fluctuation in the dimension so big that they saturate the critical dimension. In what follows, the distribution of dimensions across all possible orientations of the cone is studied.

Can we trust this estimate? If the manifold from which the voxels are taken had dimension larger than the $\log_2$ of the number of voxels for structure, we would not be able to detect it. So for the many small structures in the fine annotation, the value of the dimension is much limited by the resolution scale. But this is an inherent characteristic of dimension for physical object defined through observations of clouds of points: dimension is not a well-defined quantity and depends on scale.

4.3 What distribution of dimensions would be expected by chance?

The center and shape of this probability distribution have to be compared to the ones that would be obtained by drawing a set of random points in G dimensions, with the same sample size. Consider a fixed center C and a second point P defining a reference quadrant. The G-dimensional ambient space is partitioned into 2^G quadrants by completing the axis $[CP)$ into an orthogonal basis. When drawing $V - 2$ points randomly, we draw each of them from one of those quadrant, with equal probability

$$p = \frac{1}{2^G}.$$

For a given quadrant, the number of drawn points that falls into it follows a binomial law of parameters $V - 2$ (as two points are used to define the quadrant) and p. The dimension $\delta_{V,G}$ of the set of drawn points is therefore distributed as the (opposite of) the $\log_2$ of a binomial variable:

$$d_{V,G} = \frac{1}{\log 2} \left(\log V - \log(N_{V,G} + 2) \right), \quad N_{V,G} \sim B(V - 2, 2^{-G}),$$

$$P(d_{V,G} = D) = P(N_{V,G} = 2^{-D}(V - 2)),$$

$$\text{i.e. } P(d_{V,G} = D) = \frac{V'!}{(2^{-D}V')!(V' - 2^{-D}V')!} p^{2^{-D}V'}(1 - p)^{V' - 2^{-D}V'}, \quad V' = V - 2.$$

Any sample of the size of those we are dealing with (V up to a few thousands of points, with fixed $G = 3041$) is insufficient to parse a given dimension: the probability distribution of $N_{V,G}$ is extremely peaked at zero: $d_{V,G}$ is therefore extremely peaked at the critical dimension:

$$P(N_{V,G} = 0) = (1 - p)^{V-2} = 1 - (V - 2) \times p + o(p).$$

So, in order to have one percent of the probability away from the critical dimension put the following lower bound $V_{\min}^{pc}$ on the sample size:

$$P(N_{V,G} = 0) < 0.99 \Rightarrow V > V_{\min}^{pc} = \frac{0.01}{p} + 2 \simeq 2^{3039},$$
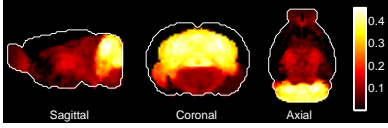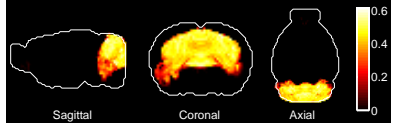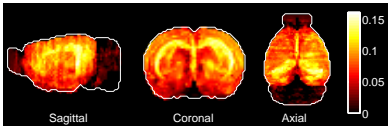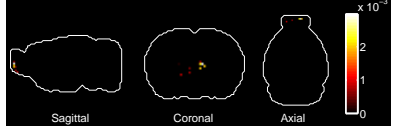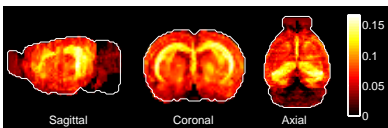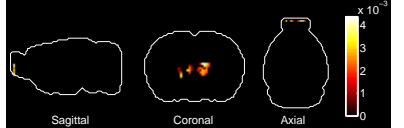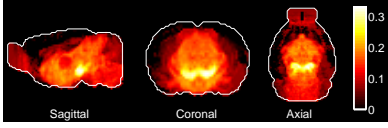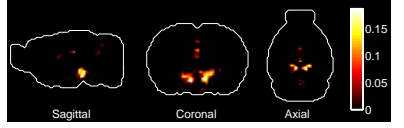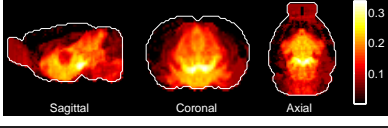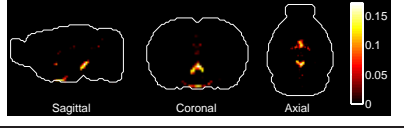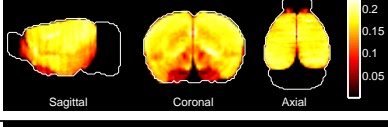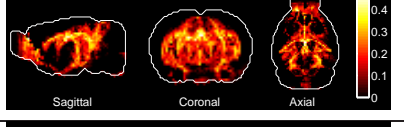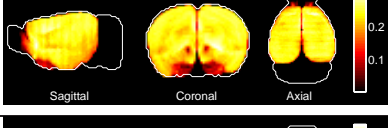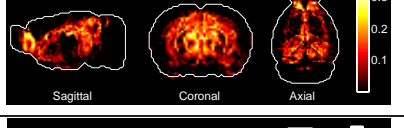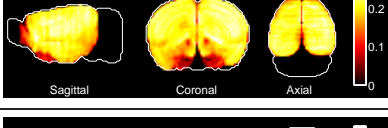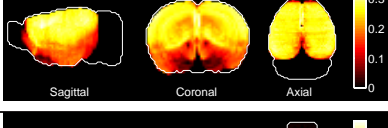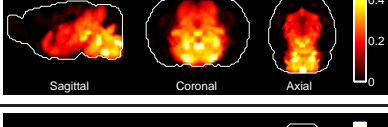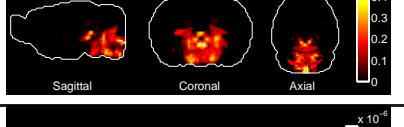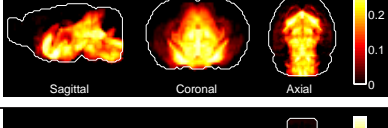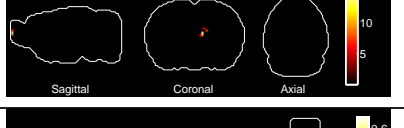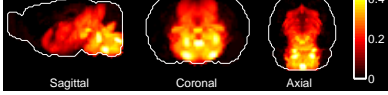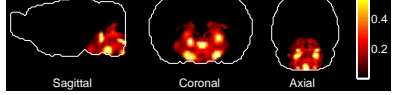
which is hopeless, as the whole brain has $V = 49742 \simeq 2^{15.6}$ voxels, so the voxel size would have to be $2^{(3039-15.6)/3} \simeq 2^{1008}$ times smaller, way out of bounds of the size of known physical objects. Anther way of putting one percent of the probability away from the critical dimension is to lower the number of genes:
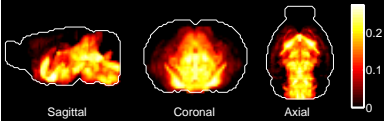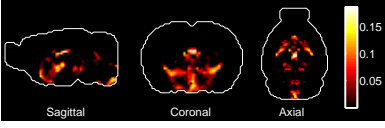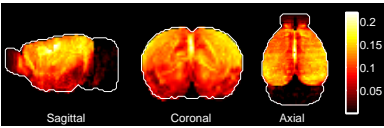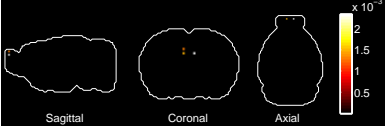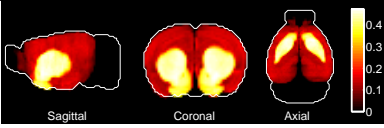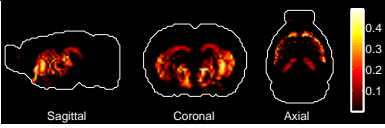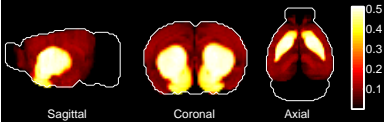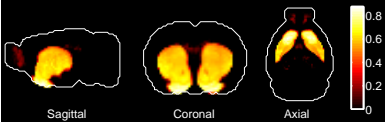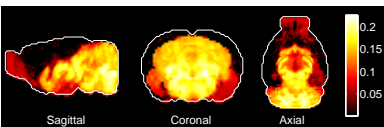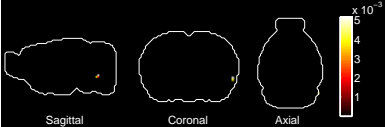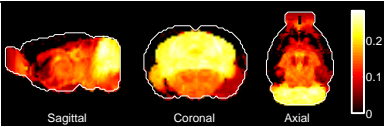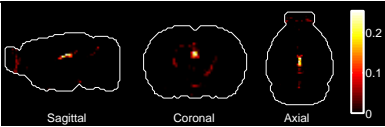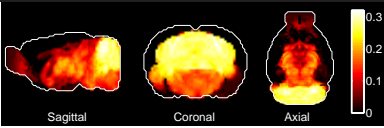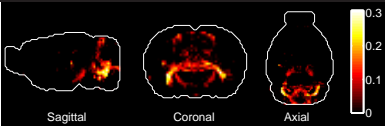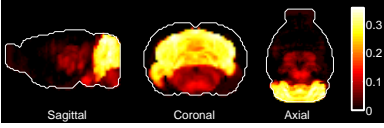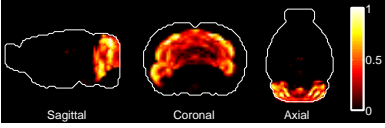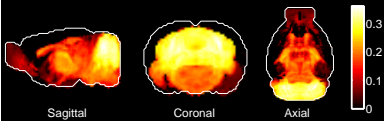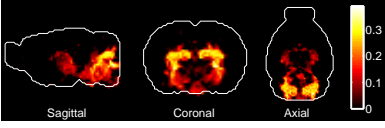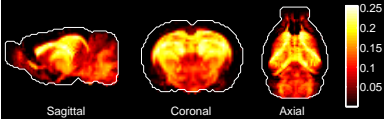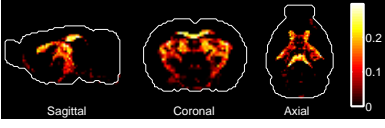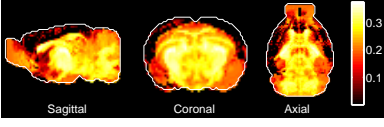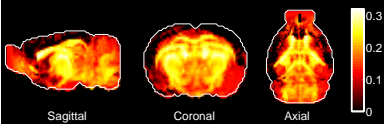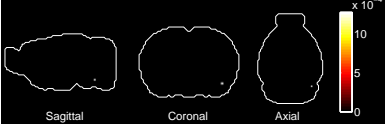
$$G < G_{\max} \frac{1}{\log(2)} \times (\log(V - 2) - \log(0.01)) \simeq 16,$$

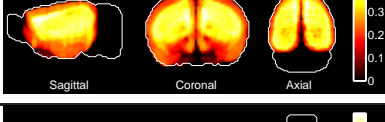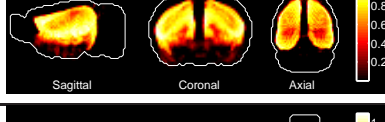which is a very low bound on the dimension of gene space, compared to the size of our data.
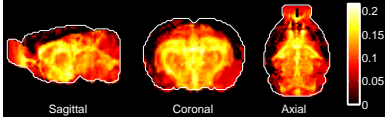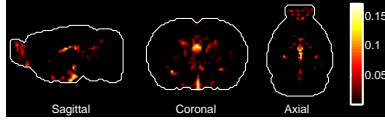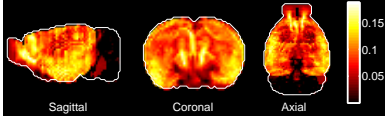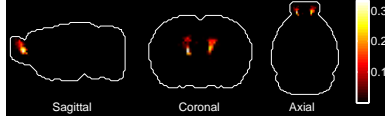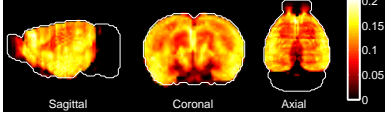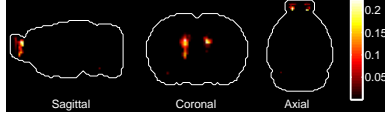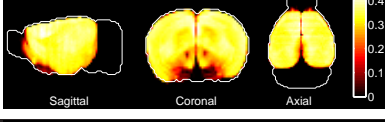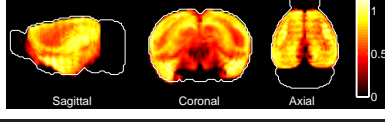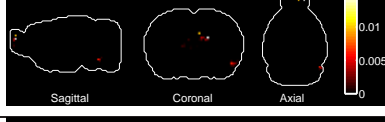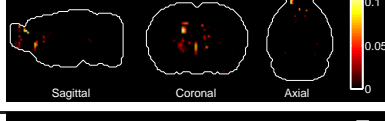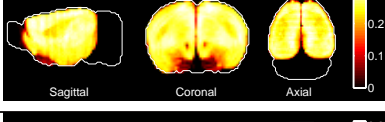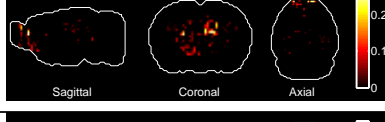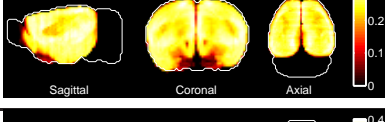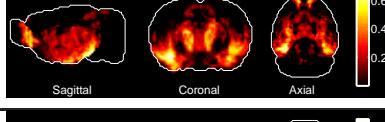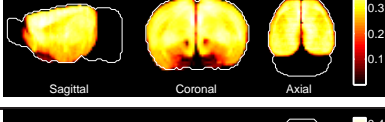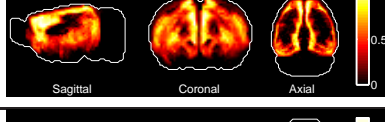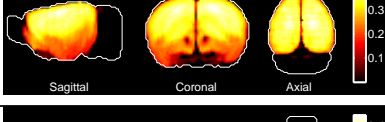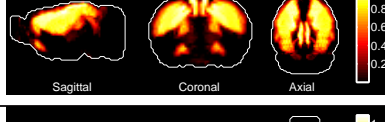
So our null hypothesis of points drawn binomially into a flat gene space is not anywhere near the actual distribution of gene expression energies, even if we processed the data at a resolution of one micron. So even if the distribution of dimensions for brain structures in gene space is not very sharply peaked because of noise or anisotropy or underlying disconnected components, the very fact that the probability is concentrated away from the critical dimension makes the gene expression data extremely packed in gene space, with a number of independent degrees of freedom much lower than the total number of genes in our data.
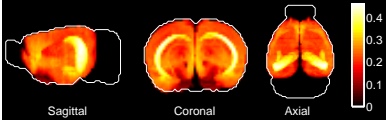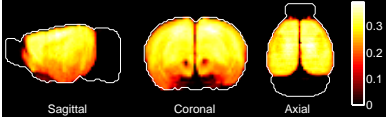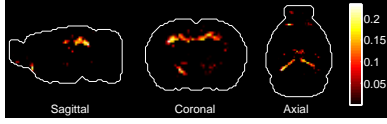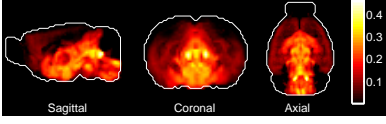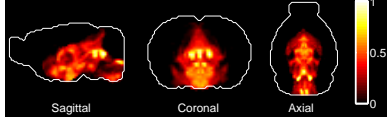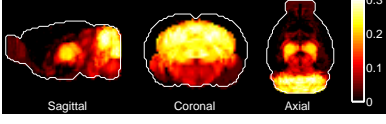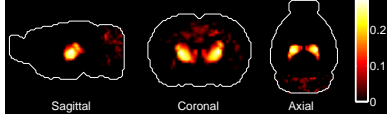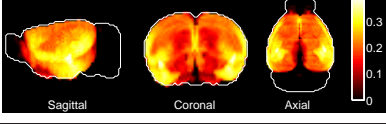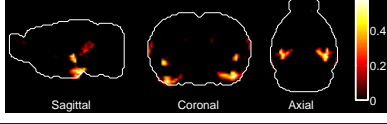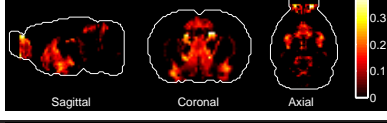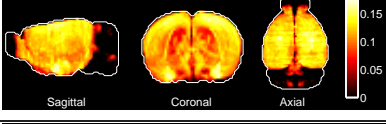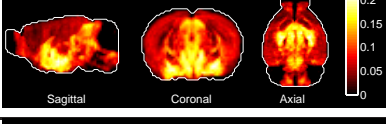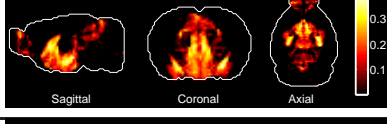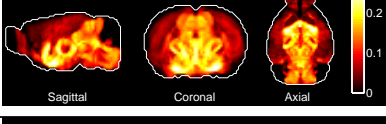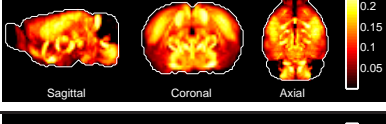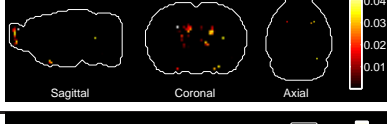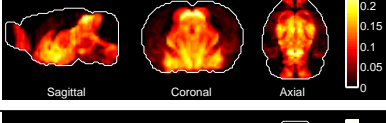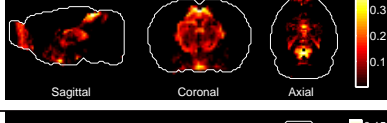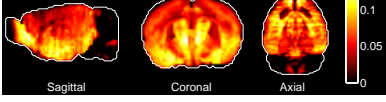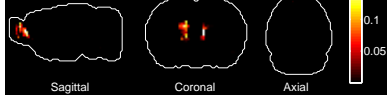
5 Tables of correlations and fittings between cell-types and gene-expression energies
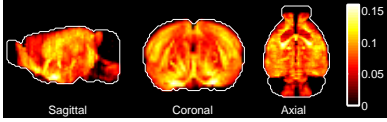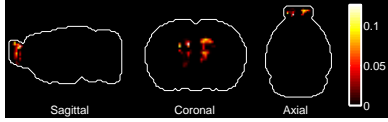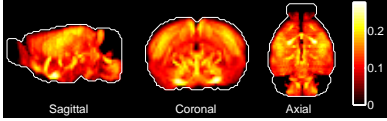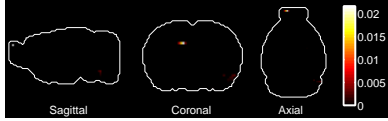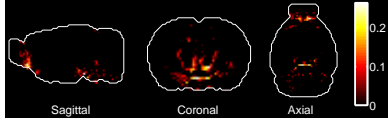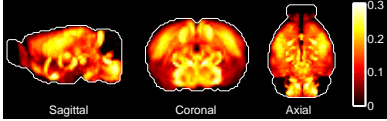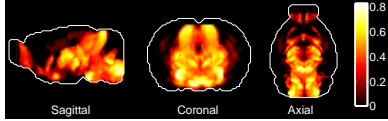
index	Cell type	Heat map of correlations	Heat map of weight
1	Purkinje Cells		
2	Pyramidal Neurons		
3	Pyramidal Neurons		
4	A9 Dopaminergic Neurons		
5	A10 Dopaminergic Neurons		
6	Pyramidal Neurons		
7	Pyramidal Neurons		
8	Pyramidal Neurons		
9	Mixed Neurons		
10	Motor Neurons, Midbrain Cholinergic Neurons		
11	Cholinergic Projection Neurons		
12	Motor Neurons, Cholinergic Interneurons		

index	Cell type	Heat map of correlations	Heat map of weight
13	Cholinergic Neurons		
14	Interneurons		
15	Drd1+ Medium Spiny Neurons		
16	Drd2+ Medium Spiny Neurons		
17	Golgi Cells		
18	Unipolar Brush cells (some Bergman Glia)		
19	Stellate Basket Cells		
20	Granule Cells		
21	Mature Oligodendrocytes		
22	Mature Oligodendrocytes		
23	Mixed Oligodendrocytes		N/A
24	Mixed Oligodendrocytes		

index	Cell type	Heat map of correlations	Heat map of weight
25	Purkinje Cells		
26	Neurons		
27	Bergman Glia		
28	Astroglia		
29	Astroglia		
30	Astrocytes		
31	Astrocytes		
32	Astrocytes		
33	Mixed Neurons		
34	Mixed Neurons		
35	Mature Oligodendrocytes		
36	Oligodendrocytes		
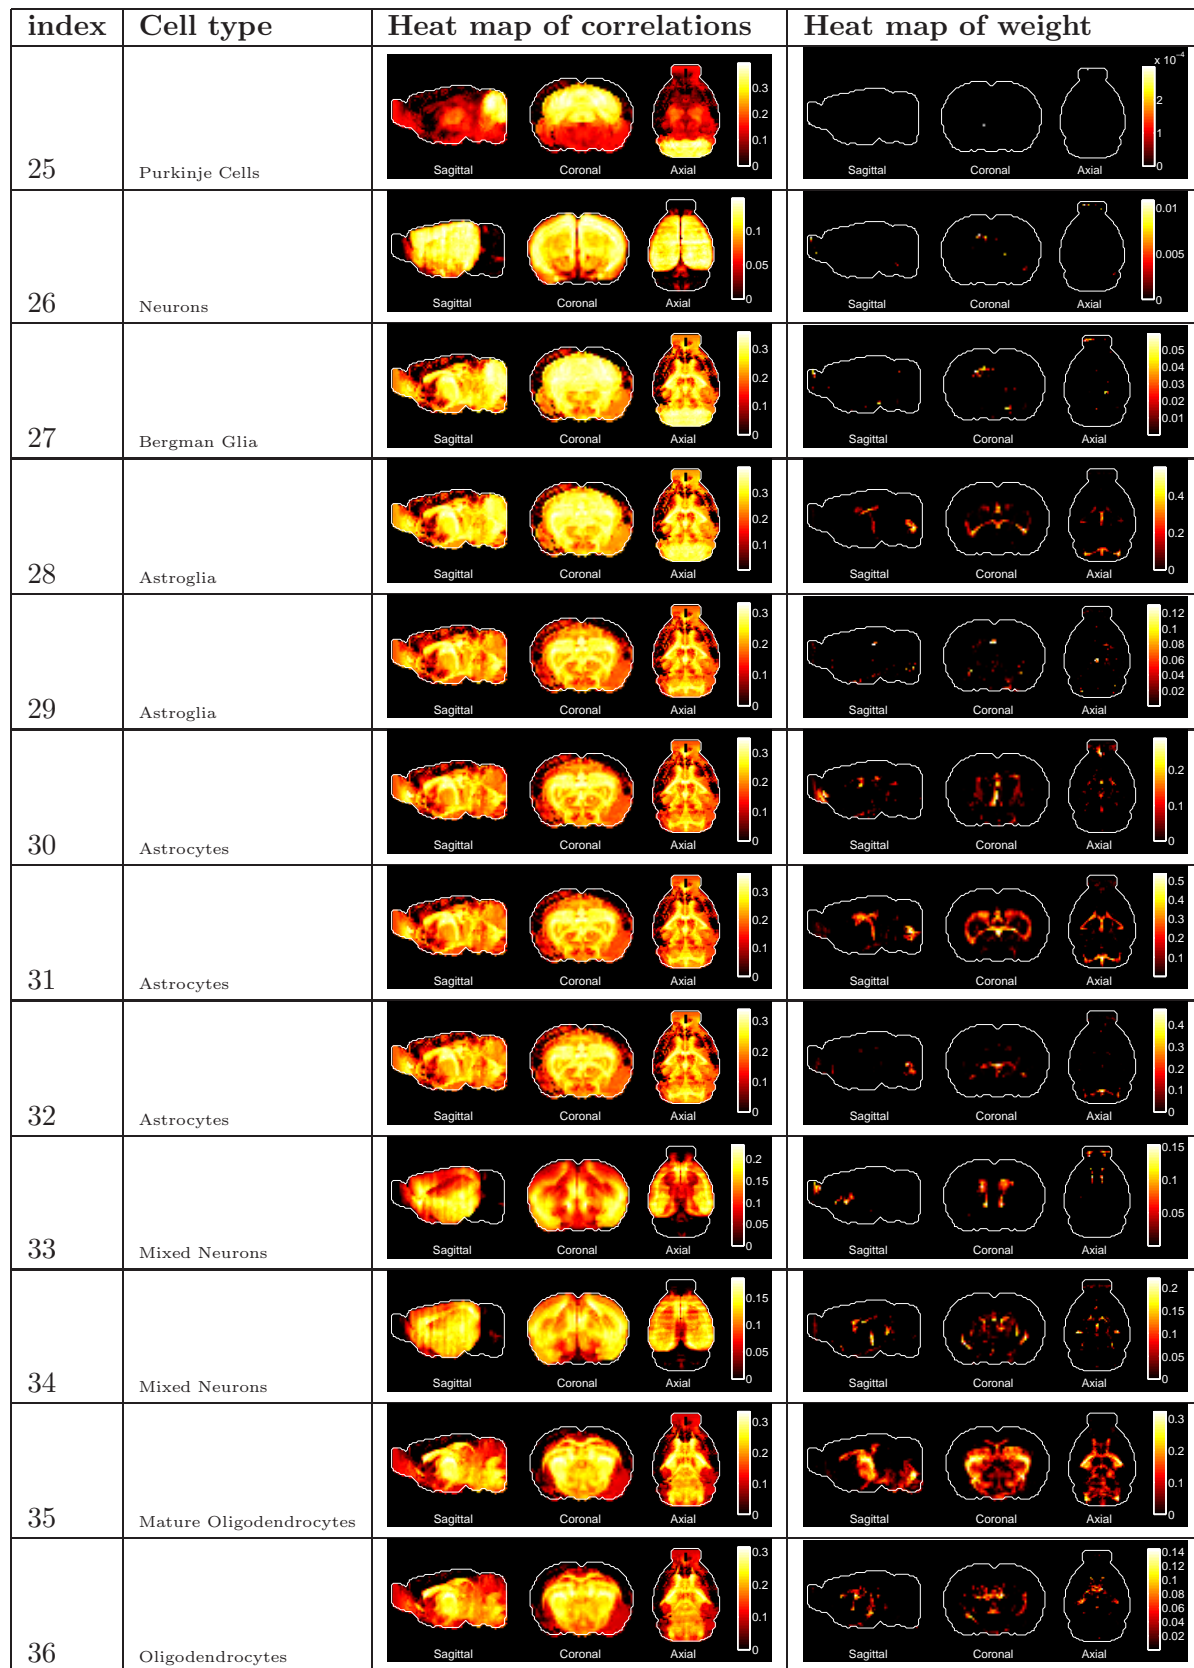
index	Cell type	Heat map of correlations	Heat map of weight
37	Oligodendrocyte Precursors		
38	Pyramidal Neurons, Callosally projecting, P3		
39	Pyramidal Neurons, Callosally projecting, P6		
40	Pyramidal Neurons, Callosally projecting, P14		
41	Pyramidal Neurons, Corticospinal, P3		
42	Pyramidal Neurons, Corticospinal, P6		
43	Pyramidal Neurons, Corticospinal, P14		
44	Pyramidal Neurons, Corticotectal, P14		
45	Pyramidal Neurons		
46	Pyramidal Neurons		
47	Pyramidal Neurons		
48	Pyramidal Neurons		

index	Cell type	Heat map of correlations	Heat map of weight
49	Pyramidal Neurons		
50	Pyramidal Neurons		
51	Tyrosine Hydroxylase Expressing		
52	Purkinje Cells		
53	Glutamatergic Neuron (not well defined)		
54	GABAergic Interneurons, VIP+		
55	GABAergic Interneurons, VIP+		
56	GABAergic Interneurons, SST+		
57	GABAergic Interneurons, SST+		
58	GABAergic Interneurons, PV+		
59	GABAergic Interneurons, PV+		
60	GABAergic Interneurons, PV+, P7		

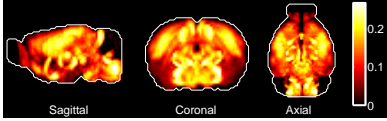| index	Cell type	Heat map of correlations	Heat map of weight
61	GABAergic Interneurons, PV+, P10		
62	GABAergic Interneurons, PV+, P13-P15		
63	GABAergic Interneurons, PV+, P25		
64	GABAergic Interneurons, PV+		

References

1. E.S. Lein *et al.*, *Genome-wide atlas of gene expression in the adult mouse brain*, Nature **445** (2006) 168-176.

2. L. Ng, A. Bernard, C. Lau, C.C. Overly, H.-W. Dong, C. Kuan, S. Pathak, S.M. Sunkin, C. Dang, J.W. Bohland, H. Bokil, P.P. Mitra, L. Puelles, J. Hohmann, D.J. Anderson, E.S. Lein, A.R. Jones and M. Hawrylycz, *An anatomic gene expression atlas of the adult mouse brain*, Nature Neuroscience **12**, 356 - 362 (2009).

3. M. Hawrylycz *et al.*, *Digital Atlasing and Standardization in the Mouse Brain*, PLoS Computational (2011).

4. M. Bota, H.-W. Dong and L.W. Swanson, *From gene networks to brain networks*, Nature neuroscience (2003).

5. L. Ng, M. Hawrylycz and D. Haynor, *Automated high-throughput registration for localizing 3D mouse brain gene expression using ITK*, Insight-Journal (2005).

6. H.-W. Dong, *The Allen reference atlas: a digital brain atlas of the C57BL/6J male mouse*, Wiley, 2007.

7. A. Gupta, Y. Wang, and H. Markram, *Organizing Principles for a Diversity of GABAergic Interneurons and Synapses in the Neocortex*, Science **14** (2000), Vol. 287 no. 5451 pp. 273-278 .

8. K. Sugino, C.M. Hempel, M.N. Miller, A.M. Hattox, P. Shapiro, C. Wu, Z.J. Huang and S.B. Nelson, *Molecular taxonomy of major neuronal classes in the adult mouse forebrain*, Nature Neuroscience **9**, 99-107 (2005).

9. Rossner MJ, Hirrlinger J, Wichert SP, Boehm C, Newrzella D, Hiemisch H, Eisenhardt G, Stuenkel C, von Ahsen O, Nave KA, *Global transcriptome analysis of genetically identified neurons in the adult cortex* J. Neurosci. 2006 **26(39)** 9956-66.

10. Cahoy JD, Emery B, Kaushal A, Foo LC, Zamanian JL, Christopherson KS, Xing Y, Lubischer JL, Krieg PA, Krupenko SA, Thompson WJ, Barres BA, *A transcriptome database for astrocytes, neurons, and oligodendrocytes: a new resource for understanding brain development and function*, J. Neurosci. 2008 **28(1)** 264-78.

11. Doyle JP, Dougherty JD, Heiman M, Schmidt EF, Stevens TR, Ma G, Bupp S, Shrestha P, Shah RD, Doughty ML, Gong S, Greengard P, Heintz N, *Application of a translational profiling approach for the comparative analysis of CNS cell types* Cell (2008) **135(4)** 749-62.

12. Okaty BW, Miller MN, Sugino K, Hempel CM, Nelson SB, *Transcriptional and electrophysiological maturation of neocortical fast-spiking GABAergic interneurons*, J. Neurosci. (2009) **29(21)** 7040-52.

13. D.D. O'Leary and Y. Nakagawa, *Patterning centers, regulatory genes and extrinsic mechanisms controlling arealization of the neocortex*, Curr. Opin. Neurobiol. **12**, (2002), 14-25.

14. R.W. Guillery, *On counting and counting errors*, Journal of Comparative Neurology, **447**:1-7 (2002).

15. P. Grange and P.P. Mitra, *Marker genes for Anatomical Regions in the Brain: insights from the Allen Gene Expression Atlas in the mouse brain*, `arXiv:1105.1217`.

16. C.L. Thompson *et al.*, *Genomic anatomy of the hippocampus*, Neuron (2008).

17. J.W. Bohland, H. Bokil, C.B. Allen and P.P. Mitra, *The Brain Atlas Concordance Problem: Quantitative Comparison of Anatomical Parcellations*, PLoS ONE (2009).

18. J.W. Bohland, H. Bokil, C.-K. Lee, L. Ng, C. Lau, C. Kuan, M. Hawrylycz and P.P. Mitra, *Clustering of spatial gene expression patterns in the mouse brain and comparison with classical neuroanatomy*, Methods, Volume **50**, Issue 2, February 2010, Pages 105-112.

19. T. Hastie, R. Tibshirani and J. Friedman, *The elements of statistical learning: data mining, inference, and prediction*, Springer Series in Statistics (2009).

20. A. Rahimi and B. Recht, *Clustering with Normalized Cuts is Clustering with a Hyperplane*, in Statistical Learning in Computer Vision (2004).

21. L.W. Swanson, *Brain Maps: Structure of the Rat Brain*, Academic Press.

22. K.B.J. Franklin and G. Paxinos, *The Mouse Brain in Stereotaxic Coordinates*, Elsevier Science and Technology.

23. L. Grady and E. Schwartz, *Isoperimetric partitioning: A new algorithm for graph partitioning*, SIAM Journal of Scientific Computing **27** (2006), no. 6, 1844-1866.

24. J.W. Bohland, H. Bokil, C.B. Allen and P.P. Mitra, *The Brain Atlas Concordance Problem: Quantitative Comparison of Anatomical Parcellations*, PLoS ONE (2009).

25. M. Brand, *Charting a manifold*, Advances in neural information processing systems (2003).

26. I. Vastrik, P. D'Eustachio, E. Schmidt, G. Gopinath, D. Croft, B. de Bono, M. Gillespie, B. Jassal, S. Lewis, L. Matthews, G. Wu, E. Birney, and L. Stein, *Reactome: a knowledge base of biologic pathways and processes*, Genome Biology **8** (2007), no. 3, R39.