Published in final edited form as:

J Neurosci Methods. 2019 January 15; 312: 162–168. doi:10.1016/j.jneumeth.2018.12.003.

A deep learning based method for large-scale classification, registration, and clustering of in-situ hybridization experiments in the mouse olfactory bulb

Alexander Andonian^{a,1}, Daniel Paseltiner^{b,1}, Travis J. Gould^c, and Jason B. Castro^{b,*}

^aComputer Science and Artificial Intelligence Laboratory (CSAIL), Massachusetts Institute of Technology, Cambridge, MA 02139, United States

^bBates College, Program in Neuroscience, Lewiston, ME 04240, United States

^cBates College, Physics Department, Lewiston, ME 04240, United States

Abstract

Background: The Allen Mouse Brain Atlas allows study of the brain's molecular anatomy at cellular scale, for thousands genes. To fully leverage this resource, one must register histological images of brain tissue—a task made challenging by the brain's structural complexity and heterogeneity, as well as inter-experiment variability.

New method: We have developed a deep-learning based methodology for classification and registration of thousands of sections of brain tissue, using the mouse olfactory bulb (OB) as a case study.

Results: We trained a convolutional neural network (CNN) to derive an image similarity measure for in-situ hybridization experiments, and embedded these in a low-dimensional feature space to guide the design of registration templates. We then compiled a high quality, registered atlas of gene expression for the OB (the first such atlas for the OB, to our knowledge). As proof-of-principle, the atlas was clustered using non-negative matrix factorization to reveal canonical expression motifs, and to identify novel, lamina-specific marker genes.

Comparison with existing methods: Our method leverages virtues of CNNs for a set of important problems in molecular neuroanatomy, with performance comparable to existing methods.

Conclusion: The atlas we have compiled allows for intra- and inter-laminar comparisons of gene expression patterns in the OB across thousands of genes, as well identification of canonical expression profiles through clustering. We anticipate that this will be a useful resource for investigators studying the bulb's development and functional topography. Our methods are publicly available for those interested in extending them to other brain areas.

*Corresponding author. jcastro@bates.edu (J.B. Castro).

¹These authors contributed equally.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at <https://doi.org/10.1016/j.jneumeth.2018.12.003>.

Keywords

Convolutional neural network; Olfaction; Mitral cell; Image processing; Neuroinformatics

1. Introduction

Histologists have long appreciated the brain's molecular heterogeneity, and have sought to understand its functional organization by cataloging region-, layer-, and cell-specific patterns of gene enrichment. While this has historically been a piecemeal effort involving study of one or a handful of genes at a time, recently developed resources allow investigation the brain's molecular organization, for thousands of genes. Perhaps the most notable of these resources is the Allen Brain Atlas (ABA)-a compendium of $\sim 10^5$ highly standardized in-situ hybridization (ISH) experiments (2D images) mapping the expression patterns of all $\sim 20,000$ genes in the mouse brain, at cellular resolution (Lein et al., 2007; Ng et al., 2009).

To fully capitalize on resources such as the ABA, one must compare expression levels from common anatomical locations pointwise across thousands of histological images, collected from thousands of experimental subjects. For tissue as complex and heterogeneous as the brain, this poses a challenging registration problem, as neural structures can vary markedly in their composition, shapes, and relative orientations even in closely spaced tissue sections. While these difficulties are greatly mitigated through experimental standardization, establishing strict correspondences across subjects ultimately requires either gridding sections to a common atlas, or nonlinear (deformable) registration of images (reviewed in (Toga and Thompson, 1999) and (Maintz and Viergever, 1998)). The former method has been used to investigate molecular anatomy at the mesoscopic scale of brain systems and structures (Bohland et al., 2010; Dong et al., 2009; Ko et al., 2013; Ng et al., 2009; Thompson et al., 2008). The latter method was used recently by (Ramsden et al., 2015) to register ISH sections of entorhinal cortex from the ABA, as well as by (Ng et al., 2005) and (Jagalur et al., 2007) in more general frameworks for registering ISH brain images.

To extend the set of available tools for atlas expression data, we have developed a deep-learning framework employing a convolutional neural network (CNN) to classify and register ISH imaging data from the Allen Brain Atlas. CNNs have advanced the state of the art in a wide range of image processing tasks, and have achieved or surpassed expert-level performance in many medical imaging domains (reviewed in (Shen et al., 2017)). Here, we illustrate their virtues in the context of atlas large numbers of histological sections with laminar precision—an application for which they have been used comparatively little (for notable exceptions, see (Cheng et al., 2018; Miao et al., 2016; Shen et al., 2017; Simonovsky et al., 2016; Yang et al., 2016)). We use our methods to assemble the first (to our knowledge) comprehensive and high-quality atlas of gene expression for the mouse olfactory bulb (OB) – a histologically complex and heterogeneous brain region that is the first CNS structure of the ascending olfactory pathway. The topography of gene expression in the OB is of great interest to anatomists and physiologists studying this structure's functional subdivisions and early development (for general reviews, see (Mori and Sakano, 2011; Murthy, 2011)).

Consequently, we anticipate that having tools to map expression in the bulb with laminar/intralaminar specificity, for thousands of genes will be an important resource for the field.

While several studies have used ISH data to explore the large-scale (thousands of genes) molecular organization of the hippocampus, neocortex, and hypothalamus, the OB has received much less attention. This is regrettable, as the OB's comparatively peripheral location (it receives direct sensory input from the periphery), and its strongly topographic inputs make it an excellent model system for relating molecular anatomy to sensory function. As a case study for developing classification/registration methods, it presents both virtues and challenges. Its well-delineated, high contrast layers are more salient than the comparatively diffuse layers of the neocortex. At the same time, the bulbar perimeter is densely packed with high spatial-frequency structures called glomeruli, which are a challenge for registration. Additionally, because the bulb is a peripheral structure, any tissue tears or wrinkles will affect section quality (unlike the hippocampus, which is an internal structure). Indeed, we observed that OB tissue quality was highly variable (Fig. 1)

Our workflow uses a CNN to derive a similarity measure for image pairs from the ABA, which we use to embed a large set of images (~30,000) in a low-dimensional space. This embedding is then used to designate sets of candidate registration partners (those from 'similar enough' planes of section that their registration is biologically meaningful), and to design a set of registration templates that tile the image space. All of our code for data fetching, registration, and clustering are available at <https://github.com/CastroLab/ImageRegistrationPipeline> to facilitate the application of these methods to other systems.

2. Methods

2.1. Workflow and data fetching

In brief, our workflow consists of data fetching, pre-processing, classification and registration template design, registration, and clustering. We first downloaded all (~ 32,000) in-situ hybridization (ISH) experiments (coronal sections) containing olfactory bulb from the Allen Brain Atlas (ABA) using the Allen Institute's API and our own custom-written Python scripts (code available on github; see above). In parallel, we also downloaded the expression image masks for each ISH experiment. Briefly, these are thresholded and segmented image masks derived from the raw data that map expression differences to a common 8-bit scale. Images were JPEGs ~ 2MB each and were maintained on local servers after downloading. Details on the ABA tissue-processing pipeline have been described in detail in previous publications (Lein et al., 2007), and are also documented on publicly available white-papers from the Allen Brain Institute: (<http://help.brain-map.org/display/mousebrain/Documentation>).

Each ABA coronal tissue section is annotated with a unique section-ID indicating rostro-caudal position relative to a reference brain sectioned at 100 μ m intervals. We grabbed all tissue sections with section IDs > 400 (larger values are more rostral, and the bulb is the most rostral structure of the mouse brain). This range was determined empirically and was chosen to conservatively ensure we missed no sections with olfactory bulb (OB). Because of this conservative criterion, we initially retained sections that did not contain OB, and which

included portions of frontal cortex, piriform cortex, and the anterior olfactory nucleus (AON).

2.2. Olfactory bulb histology

The OB is histologically complex and heterogeneous, and its cross-sectional anatomy can change considerably in successive sections. Coronal sections taken from the most rostral pole are roughly elliptical in shape, and comprise 5 sharply delineated histological layers. A large, inner nuclear-region containing granule cells (the granule cell layer, GCL) forms the central core of the OB, and 4 closed ‘ensheathing’ layers wrap around the GCL. These layers are, from deep to superficial: the inner plexiform layer (non-nuclear, containing mostly dendrites), the mitral cell layer (MCL) (containing cell bodies of the OB’s principal neurons), the outer plexiform layer (non-nuclear, containing mostly MCL dendrites), and the glomerular layer (Fig. 1). The glomerular layer forms the periphery of the OB, and contains numerous spherical structures called glomeruli-sites where olfactory sensory neurons synapse onto mitral cell dendrites. As sections are taken more caudally, the dorsal portion of the main olfactory bulb is invaded by the accessory olfactory bulb (AOB)-a discrete olfactory subsystem with inputs and projections that are non-overlapping with those of the main olfactory bulb. In coronal sections from the most caudal aspect of the bulb, the contiguity of the mitral cell layer is interrupted along the OB’s medial aspect by the anterior olfactory nucleus (AON), and frontal cortex begins to appear in the dorsal aspect of tissue sections.

2.3. Preprocessing and characterization of tissue section heterogeneity

We first classified images so that only high quality sections from anatomically similar locations would be registered to one another. In the ideal case, this would be made straightforward by simply registering images with similar or identical section ID numbers. In practice, however, we observed that classification on the basis of given section numbers was highly unreliable. Even two tissue sections (from two different animals) with identical section IDs often came from very different anatomical locations. Conversely, two tissue sections from the same anatomical location could also have very different section numbers. The basic problem is illustrated in Fig. 1, which shows data from downloaded tissue sections with the same section ID number (in this case, section ID#440).

These sections clearly derive from different anatomical locations, and a classifier relying heavily on the section ID# would therefore be unreliable. Additionally, we observed a large number of unusable tissue sections containing tears, holes, and bubbles. Fig. 1A shows a random and representative set of 25 images from our downloaded image set, highlighting the heterogeneity and variable quality of the data; these low-quality sections were reliably identified by our classifier and were excluded from the later registration pipeline.

Our basic pre-processing routine employed custom-written Matlab and Python scripts, and consisted of segmentation, thresholding, cropping, and affine transformations to correct for differences in image scale as well as slight differences in image orientation. Images were converted to grayscale, centered, and downsampled to 220×220 . After thresholding, the binary segmentation mask underwent a morphological closing operation

(`skimage.morphology.binary_closing`) with a disk structuring element of radius 2. This mask was used to set all background pixels to pure white, which was critical for registration.

2.4. Classification and derivation of the image similarity metric

We trained a convolutional neural network (CNN) to classify OB tissue sections into one of 6 groups (5 anatomical section types & 1 ‘other’ category comprising damaged, low-quality, or otherwise un-classifiable sections). The network was trained on a random set of 2,227 images that were expert-labeled (by JBC) using the scoring rubric shown in Fig. 2. Our classifier was implemented in MatConvNet (vlfeat.org) (Vedaldi and Lenc, 2015)-a deep learning framework for Matlab developed by the Oxford visual geometry group. We used a transfer learning approach using a pre-trained architecture (the popular Alexnet architecture (Krizhevsky et al., 2012), trained on ImageNet (www.image-net.org)), and fine-tuned only the last fully connected layer and softmax classifier. Briefly, the architecture comprises 5 Convolutional/Rectified Linear Unit (ReLU) layers, three of which are followed by max pooling, and three fully connected (“dense”) layers, for a total of 7 hidden layers and an output layer.

The network was trained on a custom-built machine with 2 x SLI Dual GeForce GTX Titan x 12GB GPUs, and training took ~2–3h. All downloaded ABA ISH images were passed through the trained CNN, and the weights of the 7th (and final) fully-connected layer (fc7) were used as feature vectors to describe the images. Pairwise Euclidian distances between these fc7 vectors were used as a proxy for image similarity. Classification was highly accurate, as indicated by the confusion matrix shown in Fig. 2C (derived from a 10-fold cross validation). Inspection of misclassified sections revealed that many of these were infact edge-cases that were ambiguous even for expert human classifiers (not shown).

To identify ‘neighborhoods’ of candidate registration partners in image-space, we used t-distributed stochastic neighbor embedding (t-SNE) (van der Maaten and Hinton, 2008) to create a two-dimensional representation of the set of classified ISH sections. Unlike projection-based methods of clustering (such as PCA), t-SNE uses local relationships between points (defined probabilistically) to create a low-dimensional mapping, allowing it to capture potential non-linear structure in the data. Fig. 3 shows a t-SNE embedding (van der Maaten and Hinton, 2008) of fc7 features derived from the ISH images. Neighboring images closely resembled one another, and in some cases images segregated into discrete clusters-examples of which are shown in Fig. 3.

Moreover, the distance metric we derived was far superior to other commonly used metrics, including the sum of squared differences, and mutual information between image pairs. We also directly compared our metric to one derived from the Scale Invariant Feature Transform (SIFT)-a popular feature-based metric. Briefly, we computed SIFT features at 4 spatial scales for all downloaded images (using VLFeat, as in Ramsden et al., 2015), to represent images as SIFT feature vectors. We fetched 100 seed images at random, and for each such image grabbed the 200 closest neighbors in both CNN-space and SIFT-space. To quantify image similarity of these neighbors, we computed the mean pairwise mutual information between the image seed and all 200 neighbors. In 68% of cases, this mutual information score was larger (indicating more similar images) for nearest-neighbors in the CNN space.

An example of one of these SIFT vs. CNN comparisons is shown in Fig. 4, and several others are shown in Supplementary Fig. 1.

2.5. Nonlinear registration

From inspection of the t-SNE map of images, we grabbed all images containing OB sections in which the mitral cell layer was closed and contiguous (i.e. not ‘invaded’ by the AOB dorsally, or by the AON medially). The goal here was simply to designate a set of sections that were registerable—that is, sections for which registration would be biologically meaningful. The methods described are generic and can be readily applied to other seed locations and section-criteria within the bulb. For our subsequent registration and analysis, we worked with a set of 4518 images (out of the original 32,768) images that had fully-closed mitral cell layers. To design a set of registration templates, we chose 12 uniformly distributed seed-points across the image space of closed MCLs to define a grid (Fig. 5A). We then hand-picked 10 individual ISH experiments within each of these 12 neighborhoods, choosing images for symmetry, tissue quality, and contrast.

These sets of ten images were registered groupwise using the Medical Image Registration Toolbox (MIRT) (Myronenko and Song, 2010)—a Matlab-based package for performing parametric (B-spline) registrations. In groupwise registration, individual images are iteratively registered to a recalculated mean image. Images in each neighborhood were then each registered pairwise to their closest template. In general, groupwise registration tends to be more accurate than pairwise registration, and in principle could be applied to the set of all images. In practice, however, processing thousands of images in parallel like this was computationally prohibitive and impractical, even with GPU-accelerated matrix operations. Registrations (both groupwise and pairwise) were performed using cubic B-splines for image deformation, and similarity between the fixed and target images was assessed by calculating material similarity and mutual information (Maes et al., 1997), as these metrics showed the least sensitivity to differences in contrast. Registration was performed with mostly default parameters, with the exception of regularization weight, which we increased from the default 0.01 to 0.1 to prevent ‘catastrophic’ and biologically unmeaningful image deformations. The image transformations were discovered for ISH images and also applied to the corresponding expression images for subsequent quantification and analysis of gene expression.

3. Results and discussion

A summary of registration results is shown in Fig. 6, which shows median projections for all 908 images that successfully registered to one of the templates. Whereas the median projection of pre-processed images shows relatively little differentiation between cell layers (Fig. 6A), the median projection of the nonlinearly registered images shows a sharp delineation between the bulb’s major laminae (Fig. 6B).

A single registered ISH image (Fig. 6C) is shown for comparison with the median projection. The full-width at half maximum (FWHM) for the mitral cell layer is nearly identical for the single shown image and the median projection, indicating that the resolution is sufficient to resolve laminar details. Registration quality was ground-truthed by hand, with

images scored as successful or not using a custom Python GUI that superimposed template and moving images on different color channels. Of the 908 registrations performed, 72% were manually scored as successful. There are unfortunately no directly comparable data sets or analyses to compare against, but this performance is similar to other reports in different systems (Ramsden et al., 2015). To quantify potential bias in registration successes/failures owing to differences in expression levels across images, we plotted the percentage of registration successes for images in each expression quartile (calculated from median pixel intensity over the entire expression image). We observed no trend in this relationship (Fig. S2), indicating that our atlas of registered images is not biased toward high or low-expressing genes.

As a proof of principle application, we performed a component-based analysis on the registered maps of gene expression for the fifth registration template. Each expression mask formed a column of a 3200 (pixels) by 908 (genes) matrix, which was factorized into a matrix of basis vectors (W) and a matrix of weights (H) using non-negative matrix factorization (NMF) (Lee and Seung, 1999)-a dimensionality reduction technique frequently used for genomic data (Cheng and Church, 2000; Thompson et al., 2008). Briefly, NMF seeks to represent an $m \times n$ data matrix, A , as the product of two lower-rank matrices, W ($m \times s$), and H ($s \times n$). The choice of s is application-specific. For a factorization with $s = 2$ (i.e. a 2 dimensional representation of the expression data), the OB was strongly clustered into dorsal and ventral halves (not shown), consistent with our own previous work (Noto et al., 2017) using the voxellated (200 micron³) expression data from the Allen Brain Institute, and also consistent with known topography of molecularly dissociable inputs to the dorsal vs. ventral bulb (Bozza et al., 2009; Kobayakawa et al., 2007; Yoshihara et al., 1997). For a factorization with $s = 3$ (Fig. 7), the basis vectors strongly demarcated the OB's major nuclear cell layers (the glomerular, mitral, and granule cell layers). Notably, the factorization does not have an explicit spatial component, meaning that the laminar organization and contiguity observed in the basis vectors are driven strictly by correlated expression patterns across genes. This facility of NMF in extracting laminar components from ISH data was also observed previously in the hippo-campus (Thompson et al., 2008).

Inspecting the matrix of weights (H) (Fig. 7B), we were able to easily cluster genes by their resemblance to one of the three basis vectors (Fig. 7A), and screen for genes showing distinct laminar expression profiles. In Fig. 7A, we also show sample genes exhibiting strong clustering to one of the three basis vectors. We identify three novel candidate marker genes for each of the bulb's major cell layers. The ubiquitin c-terminal hydrolase, *uchl1*, shows strong and selective labeling in the mitral cell layer (with some labeling of glomeruli as well). The extracellular matrix protein spondin 1 (*spon1*) shows strong expression in the granule cell layer, and neurexin 3 (*nrxn3*) showed selective labeling of the glomerular layer. These methods allow for large-scale and straightforward screening of genes showing laminar enrichment in the olfactory bulb (OB). As such, they may facilitate the search for genes important for laminar and areal patterning within the OB, as well as for genes that specify intrinsic and biophysical specializations of different cell types.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

We thank the Allen Institute for making these data available. JBC is supported by an NSF CAREER award (1553279). TJG is supported by an Institutional Development Award (IDeA) from the National Institute of General Medicine at the National Institutes of Health under grant number P20GM103423.

References

- Bohland JW, Bokil H, Pathak SD, Lee C-K, Ng L, Lau C, Kuan C, Hawrylycz M, Mitra PP, 2010 Clustering of spatial gene expression patterns in the mouse brain and comparison with classical neuroanatomy. *Methods* 50, 105–112. [PubMed: 19733241]
- Bozza T, Vassalli A, Fuss S, Zhang J-J, Weiland B, Pacifico R, Feinstein P, Mombaerts P, 2009 Mapping of class I and class II odorant receptors to glomerular domains by two distinct types of olfactory sensory neurons in the mouse. *Neuron* 61, 220–233. [PubMed: 19186165]
- Cheng Y, Church GM, 2000 Biclustering of expression data. *Proc. Int. Conf. Intell. Syst. Mol. Biol. ISMB Int. Conf. Intell. Syst. Mol. Biol.* 8, 93–103.
- Cheng X, Zhang L, Zheng Y, 2018 Deep similarity learning for multimodal medical images. *Comput. Methods Biomech. Biomed. Eng. Imaging Vis.* 6, 248–252.
- Dong H-W, Swanson LW, Chen L, Fanselow MS, Toga AW, 2009 Genomic-anatomic evidence for distinct functional domains in hippocampal field CA1. *Proc. Natl. Acad. Sci. U. S. A.* 106, 11794–11799. [PubMed: 19561297]
- Jagalur M, Pal C, Learned-Miller E, Zoeller RT, Kulp D, 2007 Analyzing in situ gene expression in the mouse brain with image registration, feature extraction and block clustering. *BMC Bioinformatics* 8, S5.
- Ko Y, Ament SA, Eddy JA, Caballero J, Earls JC, Hood L, Price ND, 2013 Cell type-specific genes show striking and distinct patterns of spatial expression in the mouse brain. *Proc. Natl. Acad. Sci. U. S. A.* 110, 3095–3100. [PubMed: 23386717]
- Kobayakawa K, Kobayakawa R, Matsumoto H, Oka Y, Imai T, Ikawa M, Okabe M, Ikeda T, Itoharu S, Kikusui T, et al., 2007 Innate versus learned odour processing in the mouse olfactory bulb. *Nature* 450, 503–508. [PubMed: 17989651]
- Krizhevsky A, Sutskever I, Hinton GE, 2012 ImageNet classification with deep convolutional neural networks In: Pereira F, Burges CJC, Bottou L, Weinberger KQ (Eds.), *Advances in Neural Information Processing Systems 25*. Curran Associates, Inc, pp. 1097–1105.
- Lee DD, Seung HS, 1999 Learning the parts of objects by non-negative matrix factorization. *Nature* 401, 788–791. [PubMed: 10548103]
- Lein ES, Hawrylycz MJ, Ao N, Ayres M, Bensinger A, Bernard A, Boe AF, Boguski MS, Brockway KS, Byrnes EJ, et al., 2007 Genome-wide atlas of gene expression in the adult mouse brain. *Nature* 445, 168–176. [PubMed: 17151600]
- Maes F, Collignon A, Vandermeulen D, Marchal G, Suetens P, 1997 Multimodality image registration by maximization of mutual information. *IEEE Trans. Med. Imaging* 16, 187–198. [PubMed: 9101328]
- Maintz JBA, Viergever MA, 1998 A survey of medical image registration. *Med. Image Anal.* 2, 1–36. [PubMed: 10638851]
- Miao S, Wang ZJ, Liao R, 2016 A CNN regression approach for real-time 2D/3D registration. *IEEE Trans. Med. Imaging* 35, 1352–1363. [PubMed: 26829785]
- Mori K, Sakano H, 2011 How is the olfactory map formed and interpreted in the mammalian brain? *Annu. Rev. Neurosci.* 34, 467–499. [PubMed: 21469960]
- Murthy VN, 2011 Olfactory maps in the brain. *Annu. Rev. Neurosci.* 34, 233–258. [PubMed: 21692659]

- Myronenko A, Song X, 2010 Intensity-based image registration by minimizing residual complexity. *IEEE Trans. Med. Imaging* 29, 1882–1891. [PubMed: 20562036]
- Ng L, Hawrylycz M, Haynor D, 2005 Automated high-throughput registration for localizing 3D mouse brain gene expression using ITK. *Insight J.* 14.
- Ng L, Bernard A, Lau C, Overly CC, Dong H-W, Kuan C, Pathak S, Sunkin SM, Dang C, Bohland JW, et al., 2009 An anatomic gene expression atlas of the adult mouse brain. *Nat. Neurosci.* 12, 356–362. [PubMed: 19219037]
- Noto T, Barnagian D, Castro JB, 2017 Genome-scale investigation of olfactory system spatial heterogeneity. *PLoS One* 12, e0178087.
- Ramsden HL, Surmeli G, McDonagh SG, Nolan MF, 2015 Laminar and dorsoventral molecular organization of the medial entorhinal cortex revealed by large-scale anatomical analysis of gene expression. *PLoS Comput. Biol.* 11, e1004032.
- Shen D, Wu G, Suk H-I, 2017 Deep learning in medical image analysis. *Annu. Rev. Biomed. Eng.* 19, 221–248. [PubMed: 28301734]
- Simonovsky M, Gutierrez-Becker B, Mateus D, Navab N, Komodakis N, 2016 A Deep Metric for Multimodal Registration. *ArXiv160905396 Cs.* .
- Thompson CL, Pathak SD, Jeromin A, Ng LL, MacPherson CR, Mortrud MT, Cusick A, Riley ZL, Sunkin SM, Bernard A, et al., 2008 Genomic anatomy of the hippo-campus. *Neuron* 60, 1010–1021. [PubMed: 19109908]
- Toga AW, Thompson P, 1999 Chapter 1—an introduction to brain warping In: Toga W (Ed.), *Brain Warping*. Academic Press, San Diego, pp. 1–26.
- van der Maaten L, Hinton G, 2008 Visualizing data using t-SNE. *J. Mach. Learn. Res.* 9, 2579–2605.
- Vedaldi A, Lenc K, 2015 MatConvNet: convolutional neural networks for MATLAB. *Proceedings of the 23rd ACM International Conference on Multimedia* 689–692.
- Yang X, Kwitt R, Niethammer M, 2016 Fast Predictive Image Registration. *ArXiv160702504 Cs.* .
- Yoshihara Y, Kawasaki M, Tamada A, Fujita H, Hayashi H, Kagamiyama H, Mori K, 1997 OCAM: a new member of the neural cell adhesion molecule family related to zone-to-zone projection of olfactory and vomeronasal axons. *J. Neurosci.* 17, 5830–5842. [PubMed: 9221781]

Fig. 1. Heterogeneity of olfactory bulb sections from the Allen Brain Atlas (ABA).

A) Montage of 25 randomly chosen in-situ hybridization experiments from OB-containing portions of the ABA (which we defined as having section ID > 400; see text). Note that sections are of highly variable size, scale and quality; many tissue sections contain folds, bubbles, and tears. This selection is representative of tissue quality and variability for the bulb at large. **B1)** Illustration of low reliability of ABA section ID number as an indicator of true anatomical location. All 9 images shown have the same section number (440), but can clearly be separated into at least three groups on the basis of histological features. Cartoons on the left show landmark features for each row of images. Abbreviations: MCL: mitral cell layer; AOB(gr): granule cell layer of the accessory olfactory bulb (AOB); AOB(mit & glom): mitral and glomerular layers of the AOB; Sez: sub-ependymal zone; Aco: anterior commissure. **B2)** Quantification of tissue variability. Percent of tissue sections (of 452 images) from each of the three anatomical locations indicated by the cartoons (sections categorized by an expert annotator (JBC)). As in B1, all data are from sections with a section ID of 440.

Fig. 2. Workflow and CNN pipeline for olfactory bulb image classification.

A) Overall workflow. (1) ISH images and thresholded expression masks were downloaded from the Allen Brain Institute; all coronal sections containing olfactory bulb (section ID > 400) were fetched. (2) Images were filtered and preprocessed, classified (3) using a CNN, and embedded in 2D (4) for inspection using tSNE. (5) Neighborhoods of interest in the tSNE map were identified for registration and design of registration templates. (6) Images in a neighborhood were registered pairwise to their respective template. **B)** Scoring rubric used for tissue classification to generate the labeled data set (see text). Key differentiating criteria

included: i) presence/absence of a ‘closed’ mitral cell layer; ii) shape of the anterior commissure; ii) presence of AOB cell layers (mitral and granule cells). C) Schematic showing the basic CNN architecture used for image classification (see text). Node values for the final fully connected layer (fc7; penultimate layer shown in the schematic) were used as feature vectors for computing image similarity. C) Confusion matrix summarizing performance of the network.

Fig. 3. Image space of olfactory bulb histological sections.

A) t-SNE embedding of the 7th fully connected layer (fc7) for all ~ 32,000 coronal section images from the Allen Brain Atlas with section ID > 400. These indices span the rostral portion of the brain, and include all olfactory bulb sections. **B)** Centroid (right) and median-projection (left) images for representative sub-regions shown in (A); only select sub-regions are shown for visual clarity. Color labels on the numbered sections in (B) correspond to clusters identified in the t-SNE map shown in (A). **C)** Depiction of all fc7 features (same data that are embedded in (A), with each image shown as a column vector. Images have been

sorted into non-overlapping blocks of columns (demarcated by vertical black lines), with each block corresponding to a given tSNE cluster. Each fc7 vector has been reduced to 10D (using non-negative matrix factorization) to improve visual clarity; note that clusters can be distinguished by their strong combinatorial expression of a small number of features.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Fig. 4. Comparison of CNN-based classification with SIFT-based classification.

A) Montages showing the 48 nearest-neighbors to a seed image (solid box in the upper left), with image similarity computed using the CNN fc7 metric (top montage) vs. a SIFT metric (bottom montage). Note that the seed image is the same for both montages, allowing a direct comparison. Only 48 neighbors are shown here for purposes of visualization, but quantification was done on the 200 closest neighbors (see text). Dashed boxes indicate images that are considerably different from the seed image (that is, their registration would not be biologically meaningful). **B)** Histograms of pairwise mutual information scores between the seed image shown in (A) and SIFT-neighbors (dark, filled histogram), and CNN-neighbors (red). See text for details.

Fig. 5. Image space of olfactory bulb histological sections with ‘closed’ mitral cell layers.
A) t-SNE embedding of all images for which the mitral cell layer was closed and contiguous (roughly corresponding to cluster # 3 from Fig. 3). The superimposed grid identifies 12 neighborhoods for which reference images (templates) were created. **B)** Registration templates derived by groupwise registration of 10 selected histological sections from each of the grid sections shown in (A).

Fig. 6. Registration results.

A) Median projection of 908 preprocessed (i.e. scaled and cropped) but not registered coronal sections. **B)** Median projection of the same images following preprocessing and registration to template #5. **C)** Example of a single in-situ-hybridization image registered to template #5. **D)** Line profiles confirming that registration has resolution sufficient to resolve laminar details. Line profiles were calculated along the colored line segments indicated in panels A,B and C. pre-only= pre-processing only (panel A); reg = registered (panel B); example = single example in-situ image (panel C). Abbreviations: OPL = outer plexiform layer; MCL = mitral cell layer; IPL = inner plexiform layer; GCL = granule cell layer.

Fig. 7. Proof-of-principle investigation of olfactory bulb molecular anatomy using dimensionality reduction on registered images.

A) top row: three olfactory bulb basis vectors (w1–3) derived from performing non-negative matrix factorization on registered images. middle row: in-situ hybridization images for sample genes clustering to each of the 3 basis vectors. uch1: ubi-quitin c-terminal hydrolase; spon1: spondin 1; nrnx3: neurexin 3. Bottom row: expression masks (thresholded & segmented masks of the ISH data, provided by the Allen Institute) of same genes. **B)** Plot of NMF weights (h) for all 908 genes. Many genes are located close to the origin, indicating

low expression and/or a lack of strict clustering to one of the three lamina-defining basis vectors in (A). Strongly clustered genes are those ‘snapping’ to one of the three axes; three examples of such genes are shown in red (same genes as in A).

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript