Supporting Information for:

Automated Anatomical Interpretation of Ion Distributions in Tissue: Linking Imaging Mass Spectrometry to Curated Atlases

Nico Verbeeck[1,2], Junhai Yang[3], Bart De Moor[1,2], Richard M. Caprioli[3], Etienne Waelkens[4,5], and Raf Van de Plas[,3]*

1. Department of Electrical Engineering (ESAT), STADIUS Center for Dynamical Systems, Signal Processing, and Data Analytics, KU Leuven, Kasteelpark Arenberg 10, box 2446, 3001 Leuven, Belgium.

2. iMinds Medical IT, Kasteelpark Arenberg 10, box 2446, 3001 Leuven, Belgium.

3. Mass Spectrometry Research Center and Departments of Biochemistry, Chemistry, Pharmacology, and Medicine, Vanderbilt University, 465 21st Avenue South, MRB III Suite 9160, Nashville, TN 37232, USA.

4. Sybioma, KU Leuven, Campus Gasthuisberg O&N 2, Herestraat 49, box 802, 3000 Leuven, Belgium.

5. Department of Cellular and Molecular Medicine, KU Leuven, Campus Gasthuisberg O&N 1, Herestraat 49, box 901, 3000 Leuven, Belgium.

Contents:

1. IMS Data Acquisition and Microscopy Staining Details

Materials

Ethanol, acetonitrile (ACN), and acetic acid were purchased from Fisher Scientific (Suwanee, GA), xylene was purchased from Acros (Morris Plains, NJ), chloroform, 1,5-diaminonaphthalene (DAN), and methanol were purchased from Sigma-Aldrich (Milwaukee, WI). Sinapinic acid (SA) was purchased from Oakwood Products, Inc (SC) and recrystallized twice with 70% ACN. Conductive indium tin oxide (ITO) coated microscope glass slides were purchased from Delta Technologies (Stillwater, MN). Cresyl Violet (for Nissl stain) was purchased from Electron Microscope Sciences (Hatfield, PA). Milli-Q water was from a Milli-Q Advantage A10 Ultrapure Water Purification System (Millipore, Billerica, MA).
Carnoy's fluid was prepared from 60 mL of ethanol, 30 mL of chloroform, and 10 mL of acetic acid. Fresh frozen mouse brain was purchased from Pel-Freez Biologicals (Rogers, AZ), and was sectioned using a Leica CM3050 cryostat (Leica Microsystems GmbH, Wetzlar, Germany). Frozen tissue sections were thaw mounted on cold ITO coated microscope slides. Serial sections were prepared for MS imaging and optical imaging.

Sample preparation for imaging mass spectrometry of proteins

Briefly, mouse brain sections of 12 µm thickness were rinsed with 70% ethanol 30 s, 100% ethanol 30 s, Carnoy's fluid 2 min, 100% ethanol 30 s, H_2O 30 s, and 100% ethanol 30 s. The sections were then dried for 10 min under ambient conditions and sublimated with sinapinic acid to obtain a coating of 0.2 mg/cm^2. The SA coated slides were then treated with H_2O: acetic acid (1 mL: 50 µL) vapor to recrystallize[1] the SA coating for 3.5 minutes under 85 °C.

Sample preparation for imaging mass spectrometry of lipids

Tissue sections of 12 µm thickness were sublimated with DAN to obtain a matrix coating of 0.2 mg/cm^2 and were then treated with vapor of water:methanol:chloroform (40:30:10) at 85°C for 2 min. The procedure is similar as in the literature mentioned above, but with 2 min and a total volume of 30 µL of liquid.

Nissl staining of tissue sections for microscopy

Literature reference:
Histological and Histochemical Methods: Theory and Practice, 4th edition, edited by J.A. Kiernan, Publisher: Cold Spring Harbor Laboratory Press; 4th edition (March 1, 2008), ISBN-10: 1904842429

[1] Yang, J.; Caprioli, R. M. Anal. Chem. 2011, 83, 5728–5734.

Procedure:

1. Dehydrate:
75% ethanol for 30 sec
95% ethanol for 30 sec
100% ethanol for 30 sec
100% ethanol for 30 sec
Xylene for 30 sec
Xylene for 30 sec
Xylene for 30 sec

2. Rehydrate and stain:
Xylene for 30 sec
Xylene for 30 sec
100% ethanol for 30 sec
100% ethanol for 30 sec
95% ethanol for 30 sec
75% ethanol for 30 sec
Milli-Q H_2O for 30 sec
Cresyl violet for 10 min

3. Destain:
Milli-Q H_2O
75% ethanol 30 sec
95% ethanol 30 sec
100% ethanol for 30 sec
100% ethanol for 30 sec
Xylene for 30 sec
Xylene for 30 sec
Xylene for 45 min

4. Coverslip

The stained sections were scanned using a Mirax slide scanner from Zeiss with 0.33 µm spatial resolution.

Mass spectrometry and data analysis
MALDI MS analyses were performed on a Bruker Autoflex Speed MALDI TOF mass spectrometer in positive linear mode (proteins) and a Bruker Autoflex Speed TOF/TOF mass spectrometer in negative reflector mode (lipids) using FlexControl 3.3 software. Approximately 100 shots/spot were acquired with a 1 kHz repetition rate Smartbeam II Nd:YAG laser. Image acquisition was carried out using FlexImaging 2.1 and spectral analysis was performed with FlexAnalysis 3.3 and MATLAB.

2. Registration Process Details

This section gives a more detailed description of the registration process.

1. Registration of IMS data to experimental histology

The first registration step maps IMS locations (at 80 and 100 µm resolution for lipids and proteins respectively) to the experimentally acquired microscopy image of the neighboring Nissl-stained tissue section (0.3 µm res.), further referred to as the experiment histology.

While it is common practice in IMS experiments to overlay ion images on a neighboring tissue section for visualization purposes, extra care is taken since here this alignment will become part of a larger registration pipeline.
The alignment is based on the use of an affine transformation matrix and involves several processing steps:

- Selecting an m/z that clearly co-localizes with a well defined anatomical structure; Here, ion m/z 18411 was selected, which clearly delineated the *corpus callosum* (Figure S-1), a good reference structure for registration.
- Loading both the ion image (m/z 18411) and the experiment histology into MATLAB and selecting landmark points using the cpselect function.
- Using these landmarks to create the affine transformation matrix that determines the registration by employing the MATLAB function cp2tform.

Figure S-1. Affine registration of m/z 18411 to the experimental histology.

2. Registration of atlas data to reference histology

The second step registers the anatomical atlas to the microscopy of the reference Nissl stain, which is included with the atlas and hereafter referred to as the reference histology. The registration from atlas to reference histology is taken care of by the Allen Mouse Brain Atlas, which provides these modalities already in a registered state, making step 2 for our purposes a straightforward identity transformation. Out of the reference stains provided with the atlas, the closest structural match to the experiment histology is manually selected, identical to how a medical professional would select a reference tissue in the atlas for

anatomical interpretation. Alternatively, the closest matching tissue can be selected automatically using methods such as those proposed by Abdelmoula *et al*.[2] Automatic selection can be helpful in certain situation (e.g. when a large number of tissues needs to be registered at once), but is not always desirable (e.g. in the case of partially matching tissues due to a skewed cut, where a winner-takes-all mechanism may give incorrect results). Since the selection needs to happen only once for each IMS experiment, taking care of the spatial mapping for potentially thousands of ion images in a single action, the selection mechanism is typically not a bottleneck to the interpretation efforts and its implementation can be tailored to the particular study at hand. The reference stains are all in gray-scale and have been down-sampled to 25 µm resolution.

3. Registration of experimental histology to reference histology

The registration between the experiment histology and the reference histology is the most difficult step in the registration process, as it has to account for tissue deformation during extraction and freezing of the brain, and cutting artifacts. As can be seen in Figs. S-2 and S-4, the experiment histology has a relatively high deformation compared to the reference histology from the Allen Brain Atlas. While this may seem problematic for a correct registration of both images, this deformation can be accounted for by using non-rigid image registration techniques.

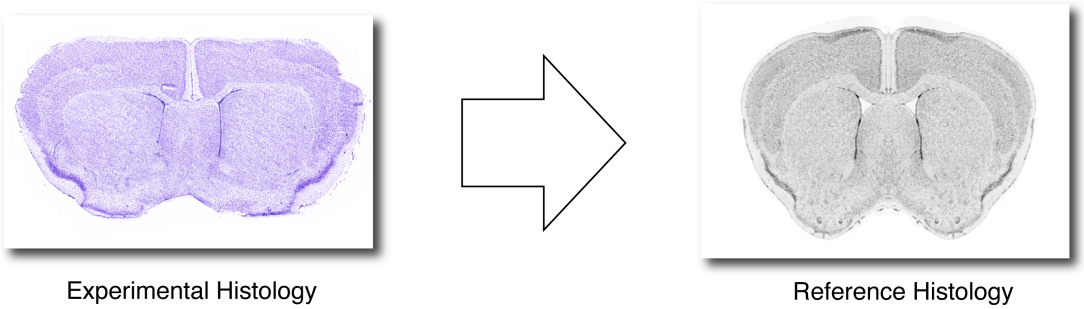

Figure S-2. The experiment histology has a relatively high deformation compared to the reference histology in the Allen Brain Atlas. The need for non-rigid registration is apparent.

Preliminary processing

- Difference in resolution

As the reference histology is provided at 25 µm resolution, and the experiment histology has a resolution of 0.33 µm, an intermediate resolution of 5 µm is used for the registration. Thus, the experiment microscopy picture is down-sampled with a factor of 15 and the reference microscopy is up-sampled with a factor of 5,

[2] Abdelmoula, W. M.; Carreira, R. J.; Shyti, R.; Balluff, B.; van Zeijl, R. J. M.; Tolner, E. A.; Lelieveldt, B. F. P.; van den Maagdenberg, A. M. J. M.; McDonnell, L. A.; Dijkstra, J. *Anal. Chem.* **2014**, 86, 3947–54.

using the imresize function provided in MATLAB.

- Difference in color scheme

As the reference microscopy is provided in grayscale, the experiment microscopy is converted to grayscale as well.

Rigid registration

The experiment histology is first centered and scaled to the reference histology using a rigid registration, in order to get a good starting point for the non-rigid registration. This is done in MATLAB by selecting landmark points using the *cpselect* function. These landmarks are then used to create the affine transformation matrix that determines the registration by employing the MATLAB function *cp2tform*.

Non-rigid registration

The non-rigid registration is performed using the algorithms from the Medical Image Registration Toolbox (MIRT) by Myronenko[3]. A non-rigid registration contains a similarity measure and a transformation model. The similarity measure defines how pixels are compared between images, while the transformation model determines the types of non-rigid transformations allowed.

- Similarity measure

The squared correlation coefficient is used as a similarity measure to compare the intensities of pixels across the images. The squared correlation coefficient is maximized when the images are linearly related, and is defined as

$$E_{CC}(T) = \frac{\left(\sum_{n=1}^{N}\left(I_n - \bar{I}\right)\left(J_n - \bar{J}^{\mathrm{T}}\right)\right)^2}{\sum_{n=1}^{N}\left(I_n - \bar{I}\right)^2 \sum_{n=1}^{N}\left(J_n - \bar{J}^{\mathrm{T}}\right)^2}$$

where $\bar{I}$ is the intensity mean of image I and $\bar{J}$ is the intensity mean of image J. The medical image processing literature shows that the best choice of similarity measure is dependent on the particularities of the data at hand. For these case studies, the squared correlation coefficient empirically showed sufficiently good results, so no further benchmarking of similarity measures needed to be pursued. Further improvements of the registration phase are welcome but not essential to the correct functioning of the anatomical interpretation phase, since its optimization formulation allows the interpretation output to degrade gracefully as the number of registration errors rises. The entire registration procedure was run on a desktop computer.

[3] Myronenko, A. Non-rigid Image Registration, PhD. Dissertation, Oregon Health & Science University, 2010.

- Transformation model

As a transformation model we use the Free Form Deformation (FFD), which is a popular non-rigid registration model that uses a mesh of control points (example in Figure S-3). In this mesh, each control point is connected to its neighbor, and each location in between these control points is interpolated via B-spline basis functions. The number of admissible transformations is limited by the spacing of the control points and by the smoothness of the B-spline basis functions. The degree of non-rigid deformation that can be modeled depends on the resolution of the mesh of control points, which in turn determines the computational complexity. In order to have an optimal trade off between computational complexity and admissible deformation, the MIRT uses a hierarchical multi-resolution approach, as suggested by Rueckert et al.[4], in which the resolution of the control mesh is increased, along with the image resolution, in a coarse to fine fashion. The initial mesh size we use consists of 8x8 control points, with 4 consecutive hierarchical levels.

Using the results of the affine registration as a starting point, the experimental histology is registered to the reference histology using the above settings. The result of the registration can be seen in Figure S-4, which shows very good matching despite the relatively large initial deformation.

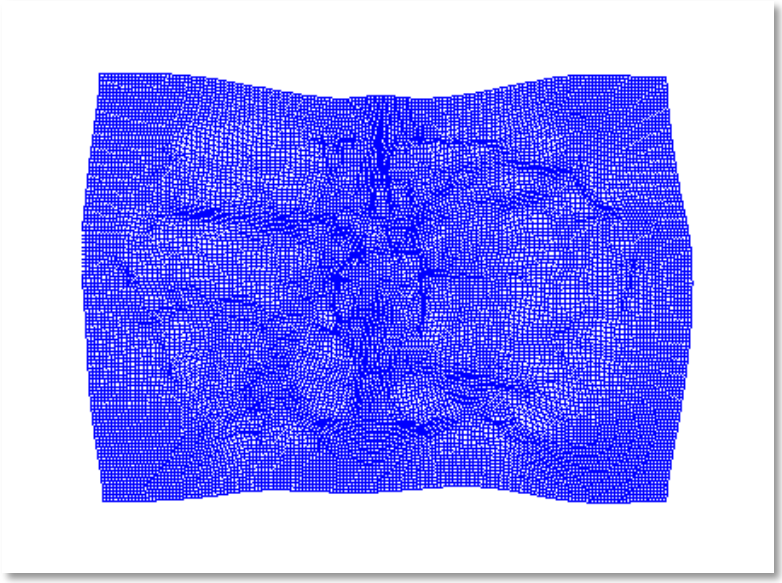

Figure S-3. The Free Form Deformation registration method uses a mesh of control points to register the experimental histology to the reference histology. The number of admissible transformations is limited by the spacing of the control points and by the smoothness of the B-spline basis functions.

[4] Rueckert, D.; Sonoda, L. I.; Hayes, C.; Hill, D. L. G.; Leach, M. O.; Hawkes, D. J. *IEEE Trans. Med. Imaging* **1999**, *18*, 712–721.

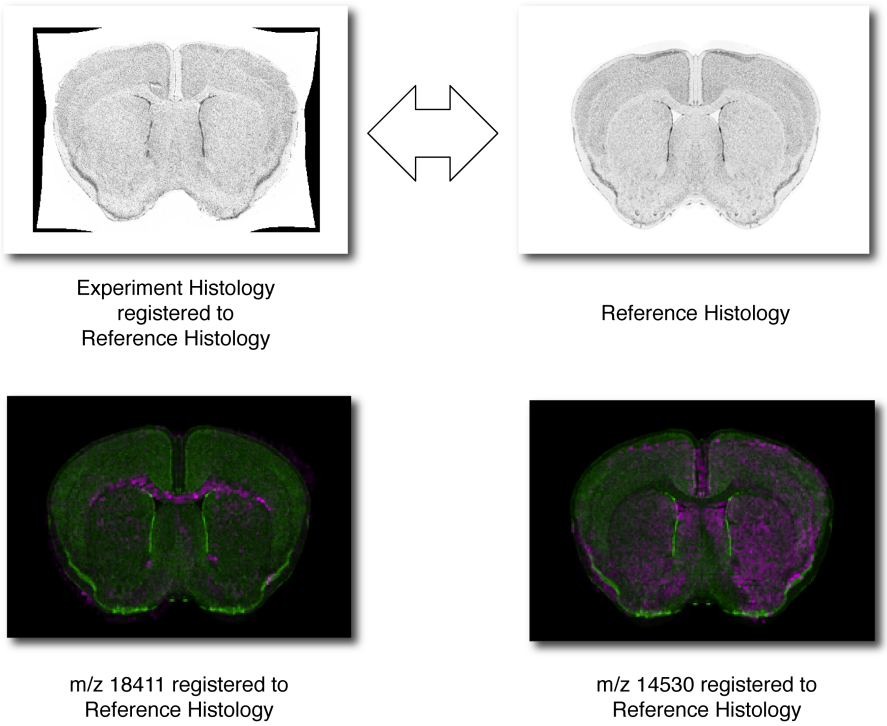

Figure S-4. Results of the non-rigid registration process. Despite the high initial deformation of the experiment histology (see Fig. S-2), the registered version looks very similar to the reference histology. The lower images show examples of using this non-rigid transformation to register IMS data to the reference histology.

3. Automated Anatomical Interpretation Details

This section gives a more detailed description of the automated anatomical interpretation process.

The goal of the anatomical interpretation method is to take a mapped ion image, and to decompose that ion distribution pattern into a combination of atlas-provided anatomical structures, without the need for human intervention. At its core, it is a problem of approximating the spatial pattern of an ion with a combination of patterns selected from a provided vocabulary of anatomical patterns. The algorithm we developed therefore considers the ion image to be a sum of contributions from a finite set of anatomical patterns, and uses a least squares argument to approach this challenge as a multivariate optimization problem. Consider the set of anatomical structure images previously defined in the section on correlation-based querying, which contain high values in pixels that are a member of the structure and low values in pixels that are not. We represent each of those anatomical images ϕ^m as a vector of length $K \in \mathbb{N}_0$ where K denotes the total number of pixels in the image, or all of them combined as

$$\{\phi^m \in \mathbb{R}_+^K\}_{m=1}^M$$

where $M \in \mathbb{N}$ is the total number of anatomical images. Note that the anatomical images are positive by construction since they represent membership to an anatomical structure. For our examples this formulation can even be simplified to $\{\phi^m \in [0,1]^K\}_{m=1}^M$. Similarly, let the ion image we want to decompose be described by a positive vector $q = (q_1, \dots, q_K)^{\mathrm{T}} \in \mathbb{R}_+^K$ of length K. Due to the IMS-atlas link, the anatomical images ϕ^1 to ϕ^M contain the same number of pixels as the ion image q. The ion image q is also positive by construction since its pixels encode ion counts.

The algorithm now seeks the optimal (and smallest) combination of anatomical images that, when multiplied by their contribution coefficients, add up to the target ion image. It is this optimal profile of intensity contribution coefficients p that answers our anatomical decomposition question. Anatomical images, and thus structures, corresponding to high absolute coefficients are important for approximating the ion distribution of interest, and are therefore considered part of the anatomical interpretation of that ion image. Anatomical structures with low absolute coefficients are considered unrelated to the ion image of interest. The following linear model is adopted for any $k = 1, \dots, K$

$$q_k = \sum_{m=1}^{M} \phi_k^m p_m + \varepsilon_k$$

where the coefficients $p = (p_1, \dots, p_M) \in \mathbb{R}^M$ encode the assumption that the ion

image is essentially a weighted sum of the anatomical images/structures, up to the residual $\varepsilon = (\varepsilon_1, \dots, \varepsilon_K)^{\mathrm{T}} \in \mathbb{R}^K$. Our implementation uses a classical approach to approximate the linear coefficients we are looking for, minimizing the squared norm of the residuals, or

$$p^* = \arg\min_p \sqrt{\sum_{k=1}^{K} \left(\sum_{m=1}^{M} \phi_k^m p_m - q_k \right)^2}$$

This optimization problem considers the optimal contribution coefficients to be the ones that minimize the difference between the measured ion image and its approximation using anatomical patterns. However, since the anatomical structures are stored in the Allen Mouse Brain Atlas in the form of a hierarchical tree, many anatomical patterns are very similar to each other, exhibiting large overlap and often differing by only a small sub-area. In the current formulation, this could cause multiple similar structures to be selected, dividing that area's contribution among them. This could result in unnecessarily complex anatomical interpretations consisting of a large numbers of structures. Instead, we want the algorithm to select only the best fitting anatomical structures for its interpretation, disregarding similar but less optimal ones, and keeping the total number of anatomical structures in the interpretation to a minimum. To accomplish this, we introduce an L1 regularization term for an appropriate choice of $\gamma \geq 0$, turning the anatomical interpretation into

$$p_\gamma^* = \arg\min_p \left[\sqrt{\sum_{k=1}^{K} \left(\sum_{m=1}^{M} \phi_k^m p_m - q_k \right)^2} + \gamma \sum_{m=1}^{M} |p_m| \right]$$

Similar to the LASSO[5] algorithm, the 1-norm makes the optimization problem pursue sparsity in the solution of coefficients. This can be interpreted as follows: if the optimization problem does not know which solution p^* to prefer up to a numerical quantity, then choose the solution with the smallest 1-norm. Parameter γ regulates what is meant by such a numerical quantity: if γ is large, a sparse solution containing only a few significant anatomical contribution coefficients will be preferred over an exact least squares fit of the ion image, while the reverse holds true for when γ is small. In our case studies, we set $\gamma = 1$. Given the number of pixels versus the number of anatomical structures involved in these data sets, this setting ensures that approximation of the ion image is more important than sparsity in the anatomical interpretation, while the sparsity term does keep the anatomical explanation to a minimum if multiple valid interpretations exist.

Note that the coefficients are not constrained to positive values in this

[5] Tibshirani, R. J. R. *Stat. Soc. Ser.* B **1996**, 58, 267–288.

formulation. This allows the anatomical interpretation to say things like “the ion seems to be present in zone A plus zone B minus zone C.” Although it will usually result in a more intricate interpretation using lots of little structures, if desired, it is possible to constrain the algorithm to only consider additive interpretation by setting a positivity constraint on the coefficients in the optimization problem.

Our implementation uses CVX, a package for specifying and solving convex programs[6,7] to solve the optimization problem for each ion image we want interpreted. The method delivers an approximation of each ion image using anatomical patterns, and a set of coefficients that report which anatomical structures are involved. A nice feature of the method is that an anatomical contribution coefficient tends to be proportional to the ion intensity in that anatomical structure, inherently assigning a notion of importance to each anatomical zone involved.

---

[6] Grant, M.; Boyd, S. In *Recent Advances in Learning and Control*; Blondel, V.; Boyd, S.; Kimura, H., Eds.; Springer-Verlag Limited, **2008**; pp. 95–110.

[7] Grant, M.; Boyd, S. *CVX: Matlab Software for Disciplined Convex Programming, version 2.0 beta;* **2012**.

4. Additional Correlation-based Querying Examples

Protein case study examples

Figure S-5 shows the correlation table that results from the spatial correlation analysis between the anatomical structure images and the ion images of the peak picked IMS protein data. The y-axis shows all the anatomical structures at this depth of the brain (119 in total). The x-axis shows the peak picked m/z values. Anatomical structure/ion abundance combinations that have a high positive correlation (range ρ= [0.54 0.76]) are shown in dark red. Combinations that have a high negative correlation are indicated in dark blue. The other combinations exhibit low correlation (|ρ|<0.15).

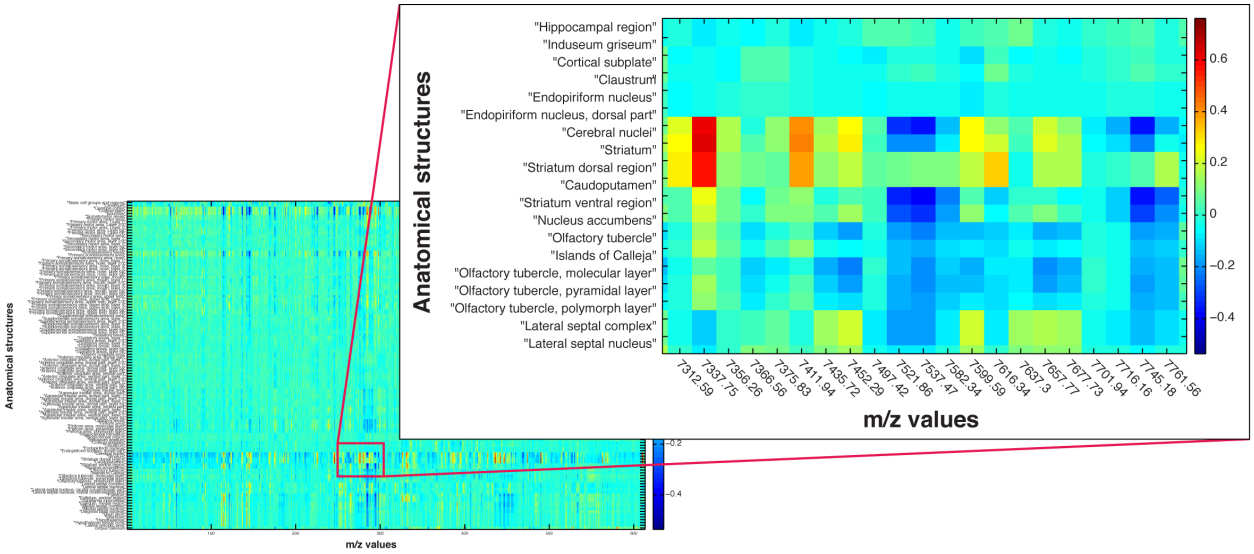

Figure S-5. The correlation table displaying spatial correlation between the anatomical structure images and ion images of the peak picked IMS protein data. The y-axis shows all the anatomical structures that are present at the depth in the brain that this IMS experiment was acquired at. The x-axis shows the peak picked m/z values. Red = high positive correlation, blue = high negative correlation, green = low correlation.

Figure S-6 shows an extended overview of the correlation-based queries presented in the main paper, including examples of negative and near-zero correlation. Ions m/z 6279 and m/z 11247 are two randomly selected ion image examples that exhibit low spatial correlation (|ρ|<0.15) with the caudoputamen, while ions m/z 11842 and m/z 6961 describe ion images that have high negative correlation with the target structure, exhibiting a clear absence of these ions therein. The anti-correlating examples demonstrate an interesting feature of the algorithm, which is that the anatomical query can be used not only to find ions specifically present in the anatomical structure of interest, but that it can also be used to search for ions specifically absent from the region of interest.

Figure S-7 shows several examples of positively correlated anatomical structure/ion image pairs in the protein data set. Example S-7A shows the correlation of the *pallidum* (ventral region) and m/z 4538, which is almost exclusively located in this area. The congruence of these relatively small areas demonstrates the accuracy of the registration. Figure S-7C shows m/z 18483, which is well localized in the *corpus callosum*. Unfortunately, the *corpus callosum* is currently not available as a separate anatomical area in the Allen Brain Atlas, and is a substructure of the fiber tracts and ventricles, which shows high correlation with this ion in the current tissue slice. Figure S-7G shows the correlation between the lateral septal complex and m/z 9751. From this image we can see that m/z 9751 is not exclusively located in the septal complex, but rather in multiple anatomical structures. This example emphasizes the need for a multi-membership approach.

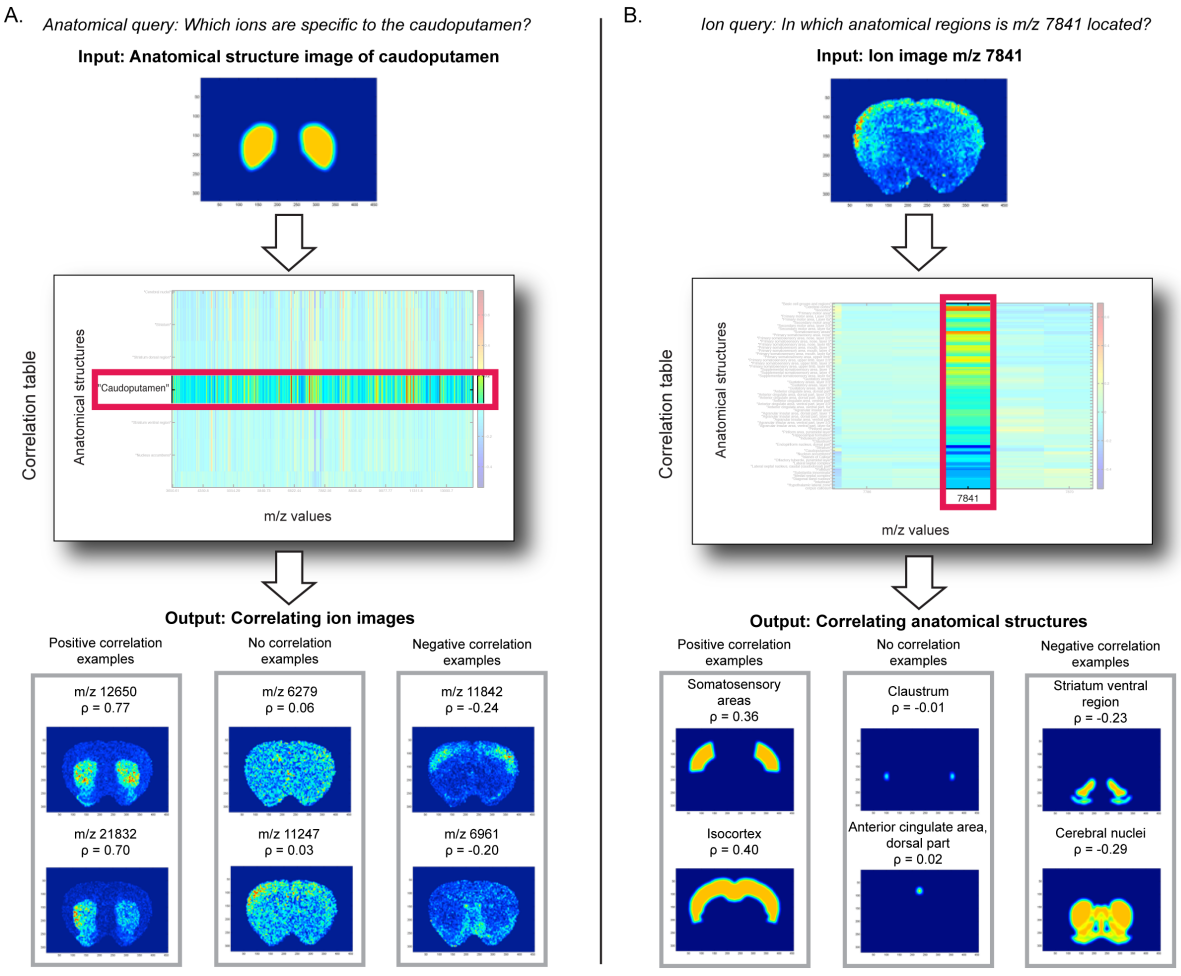

Figure S-6. (Panel A) Example of an anatomical query, finding ions specific to the *caudoputamen*. The anatomical structure image of the *caudoputamen* is given as input and the correlation table returns the spatial correlation to this structure for each ion image. Two examples of ion images that positively correlate with the target anatomical structure are displayed on the left. Examples of no correlation are shown in the middle,

with examples of negative correlation on the right. (Panel B) Example of an ion query, finding anatomical regions in which m/z 7841 is highly expressed. The ion image of m/z 7841 is given as an input and the correlation table returns the spatial correlation to this ion image for all the anatomical regions. Two examples of anatomical structure images that positively correlate with the target ion image are displayed on the left. Examples of no correlation are shown in the middle, with examples of negative correlation on the right.

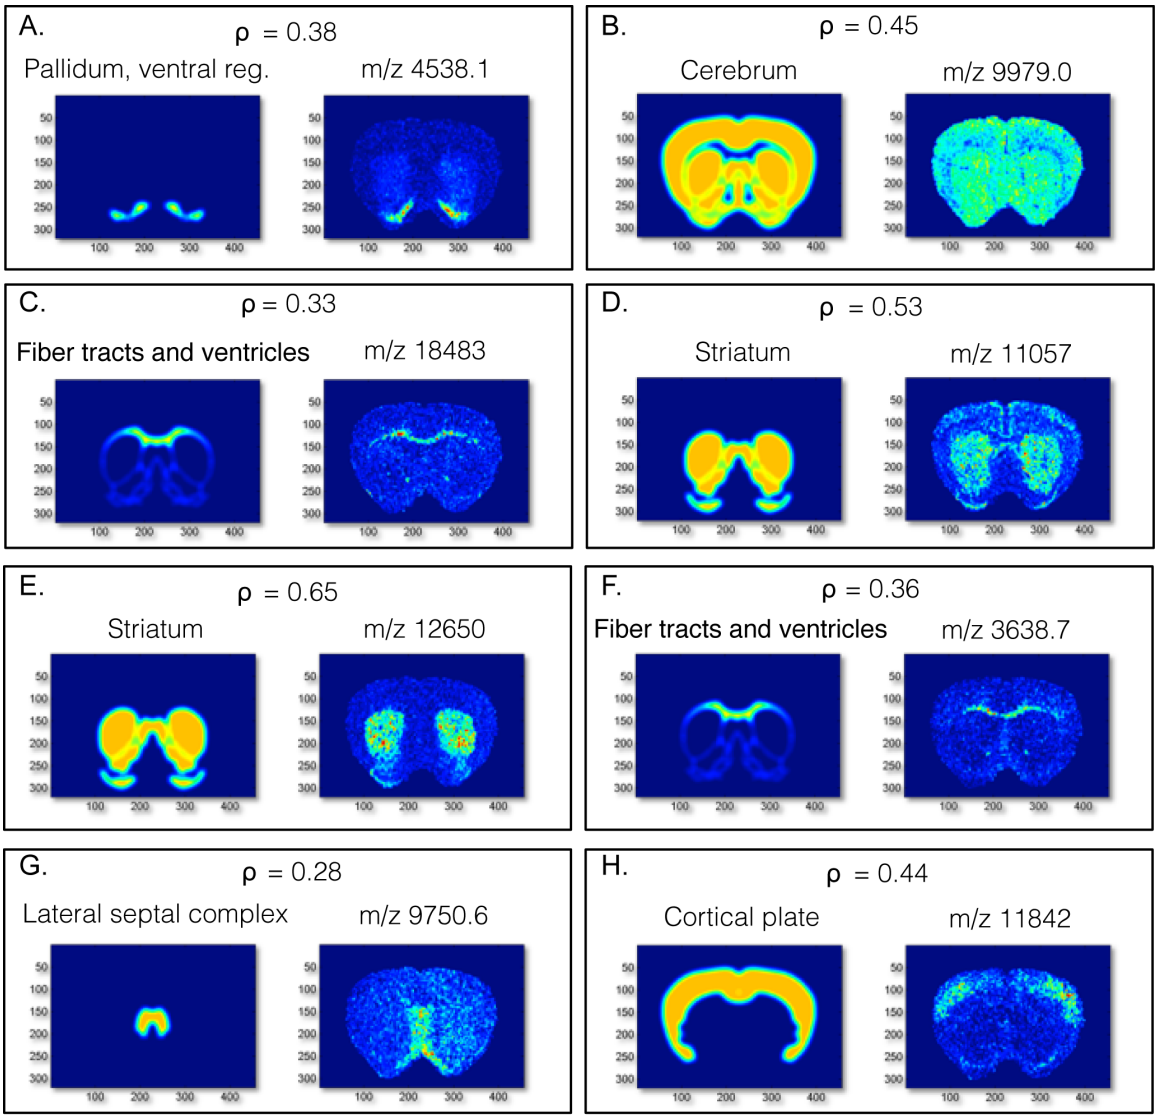

Figure S-7. Examples of positively correlating anatomical structure/ion image combinations. In each image the anatomical structure image is shown on the left and the correlating ion image is shown on the right.

Lipid case study examples

Figure S-8 shows the correlation table for the lipid data (negative mode) with the atlas data. It shows areas of high positive, high negative, and low correlation. Figure S-9 displays several examples of anatomical structure/ion abundance combinations with high positive correlation.

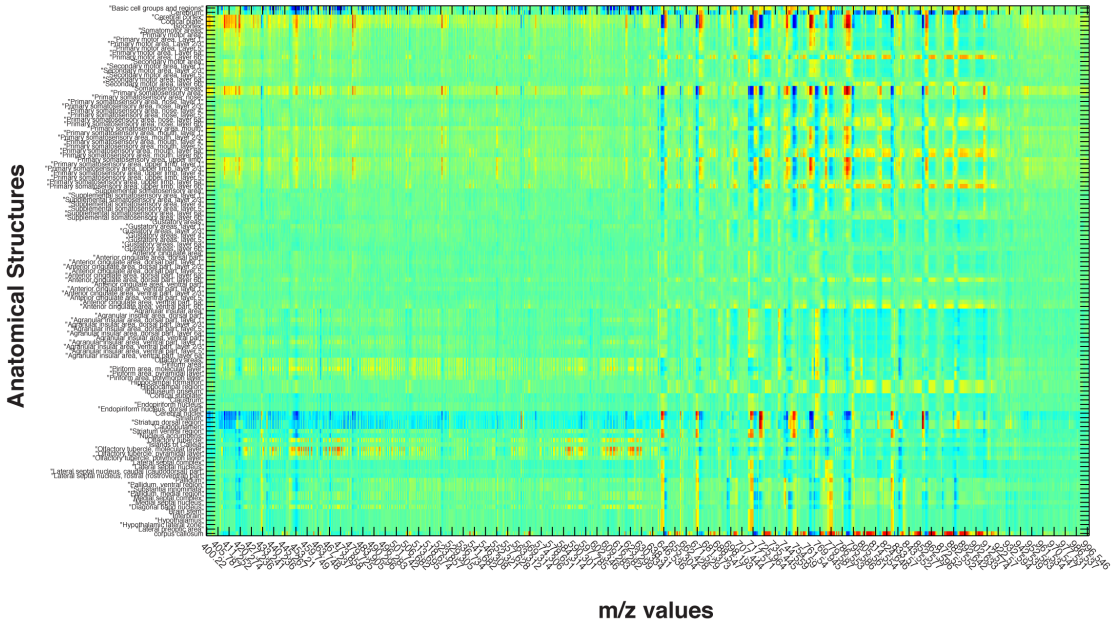

Figure S-8. The correlation table resulting from the spatial correlation analysis between the anatomical structure images and the ion images of the peak-picked IMS lipid data. The y-axis shows all the anatomical structures that are present in the tissue slice. The x-axis shows the peak-picked m/z values. Red = high positive correlation, blue = high negative correlation, and green = low correlation.

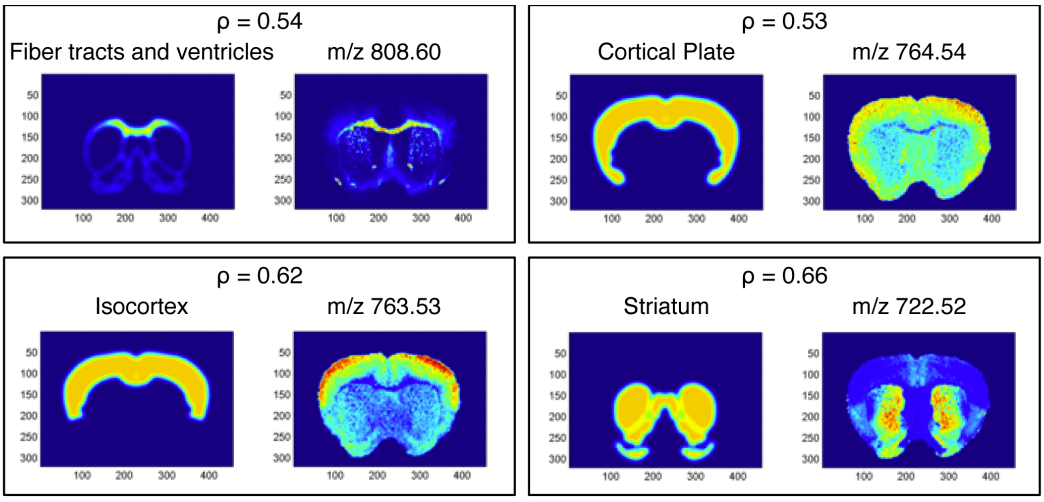

Figure S-9. Examples of several positively correlating lipid ion images. In each image the anatomical structure image is shown on the left and the correlating ion image is shown on the right.

5. Additional Automated Anatomical Interpretation Examples

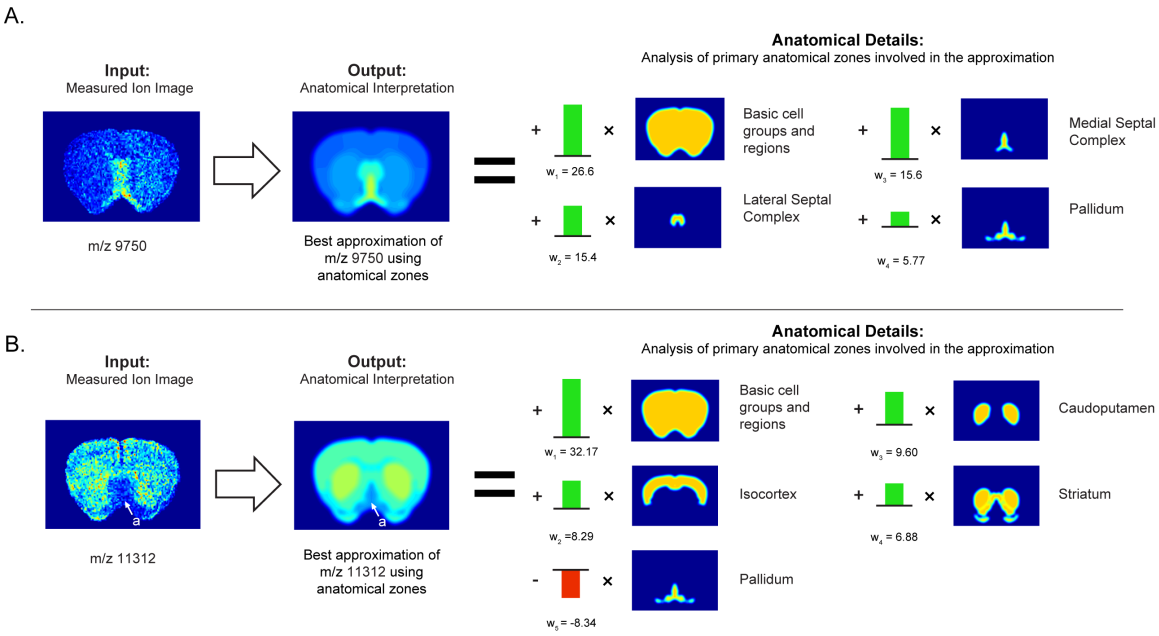

Figure S-10. Examples of automated anatomical interpretation applied to the protein-focused IMS case study. When an ion image is given as input (left), the interpretation method provides an optimal anatomical explanation for the observed ion pattern (right), using the library of provided anatomical structures. Specifically, the ion intensity pattern is decomposed without user intervention into an optimal combination of contributing anatomical structures from the atlas. The interpretation method provides (i) the closest approximation of the measured ion image using atlas structures (middle), and (ii) a linear combination of the contributing anatomical structures, specifying name, reference location, and contributing intensity or weight in the interpretation (right). This visualization shows for each ion image the most contributing anatomical zones, preceded by their weight, which signifies their relative importance in the overall anatomical interpretation. Negative weights indicate a relative decrease of the ion in those areas. (panel A) Ion m/z 9750 is highly expressed in the medial septal complex and the lateral septal complex. (panel B) Ion m/z 11312 is specifically absent in the *pallidum*. This is indicated by its negative contribution coefficient or weight, and is also clearly visible in the ion image. The empirical ion distributions show good congruence with the boundaries of the anatomical structures defined in the atlas, indicating good spatial mapping between the data sources, and strong biological signals in IMS measurements also in the protein-focused case study.

6. Sparsity in Automated Anatomical Interpretation

From the equation of the convex optimization formulation (p. S-10), it is clear that γ is the only parameter in the anatomical interpretation model. Parameter γ controls how much emphasis is placed on the approximation of the measured ion pattern versus the simplicity (sparsity) of the anatomical explanation. This parameter functions essentially as a dial that allows the user to choose whether a broad overview anatomical interpretation is provided, capturing the major anatomical trends of the ion pattern, or whether a more detailed interpretation is returned, that breaks those larger anatomical zones into more precise sub-zones. The γ dial of the algorithm has a range from 0% sparsity (closest approximation, most detailed anatomical interpretation) to 100% sparsity (approximation of only the broader trends, less complex anatomical interpretation).

Which level of interpretation complexity is preferable is dependent on the particular purpose that the anatomical interpretation will serve. If it is meant for human consumption, it is often not preferable to have a very intricate interpretation using twenty or more anatomical labels. Instead, a medical professional typically wants a broad breakdown into the five or six major anatomical zones involved, unless they have a particular case-study reason to request otherwise. If the interpretation is to be used as annotations for a computational follow-up (e.g. classification or clustering that includes anatomical labels as extra information), there is little reason not to go for highly detailed interpretations since the computer will easily handle these. Regardless of the particular use case, the same model is capable of delivering these interpretations.
Given these examples, it is clear that there is not one particular 'best' value for γ, since the best value for one application is not necessarily the best for another. We therefore prefer to provide the user access to this dial, so that he/she may determine directly which sort of interpretation complexity is desired for a particular use case. However, note that the ability to request different levels of complexity in an anatomical interpretation does not jeopardize the analytical accuracy of the interpretation. This is because a less complex anatomical interpretation (higher γ) is directly related to a more complex interpretation (lower γ) in the sense that large anatomical zones are simply broken down into subdivisions that allow for a more granular description. As such, changing γ only changes the complexity of the list of anatomical labels involved, it does not change which anatomical space is tied to the ion pattern.

For the proof-of-concepts and introduction of the multivariate model and algorithm in this paper, we have set γ to one since it allows us to demonstrate the capabilities of the algorithm sufficiently to make our point. Due to the dimensions of our data sets (and in fact most MALDI IMS data sets) a γ of one ensures that the sparsity is secondary to the approximation. It basically ensures that only if two equally close anatomical interpretations are available, the sparser one wins.

Another note is that due to the speed and parallelism of the algorithm, it is perfectly feasible to explore multiple γs and determine the 'optimal' γ for a particular data set and application. In fact, this is the common approach to setting such parameters in the optimization literature. An example of this is the use of 'optimal trade-off curves' in Boyd and Vandenberghe[8].

With regards to comparing anatomical interpretations between different ion images, it is also good to note that the γ parameter is equal across all the hundreds of ion images within an IMS experiment. This makes these interpretations directly comparable regardless of the particular value of γ, an advantage not available from human interpretation (where the level of interpretation complexity is a free-floating variable, subjectively controlled at best).

[8] S. P. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, Cambridge, UK, 2004.

7. Automated Anatomical Interpretation and the Structure of Annotation Labels

With section S-6 discussing the role of sparsity in the anatomical interpretation, it is valuable to also have a closer look at the structure of the AMBA and specifically how the structure of the annotations in the AMBA pertains to the interpretation. As mentioned in the manuscript on page 6-7, the AMBA is organized in a hierarchical fashion. This annotation structure sometimes does not accurately reflect the region subdivisions encountered in nature, and since the anatomical interpretation method can only work with the anatomical annotations provided by the atlas, an optimal interpretation sometimes returns more anatomical zones than would be expected at first sight.

Below, we illustrate this hierarchical subdivision of the atlas with several screenshots (Fig. S-11 to S-13), taken from the interactive atlas viewer available on the AMBA website[9]. As an example, these figures show that the *isocortex* is hierarchically first sub-divided into "radial" sections, and only at a deeper level do these radial structures get split further into concentric sub-divisions. The hierarchical annotation structure prevents the AMBA from providing a single named annotation that represents all concentric regions of the *isocortex* as one, regardless of the radial locations involved. As such, the automated anatomical interpretation method does not have a concentric radial-independent sub-division pattern available in its vocabulary of anatomical structures, and can only represent such a structure as a combination of concentric sub-divisions of radial sub-divisions. Studies such as Ko et al.[10] on the other hand show that in biology a concentric radial-independent division exists in the mouse brain, and such a division can also be seen in ion image m/z 7841 of Figure 4.B (the intensity of the ion is higher in the outer-most concentric layer of the brain).

[9] http://atlas.brain-map.org/

[10] Ko, Y.; Ament, S. A.; Eddy, J. A.; Caballero, J.; Earls, J. C.; Hood, L.; Price, N. D. Proc. Natl. Acad. Sci. U. S. A. 2013, 110, 3095–100.

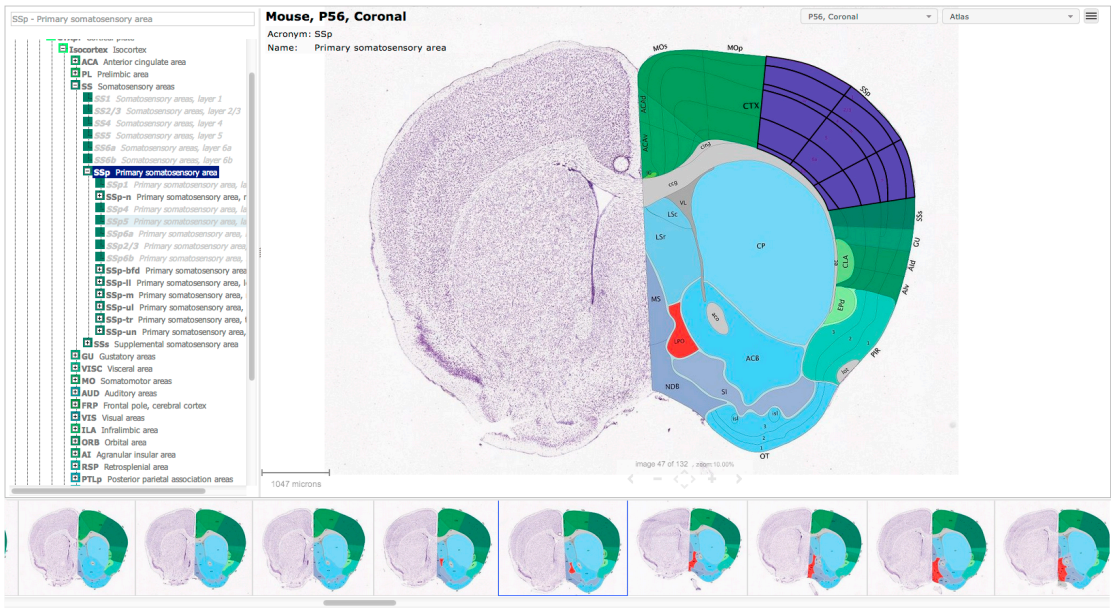

Figure S-11: Primary somatosensory area

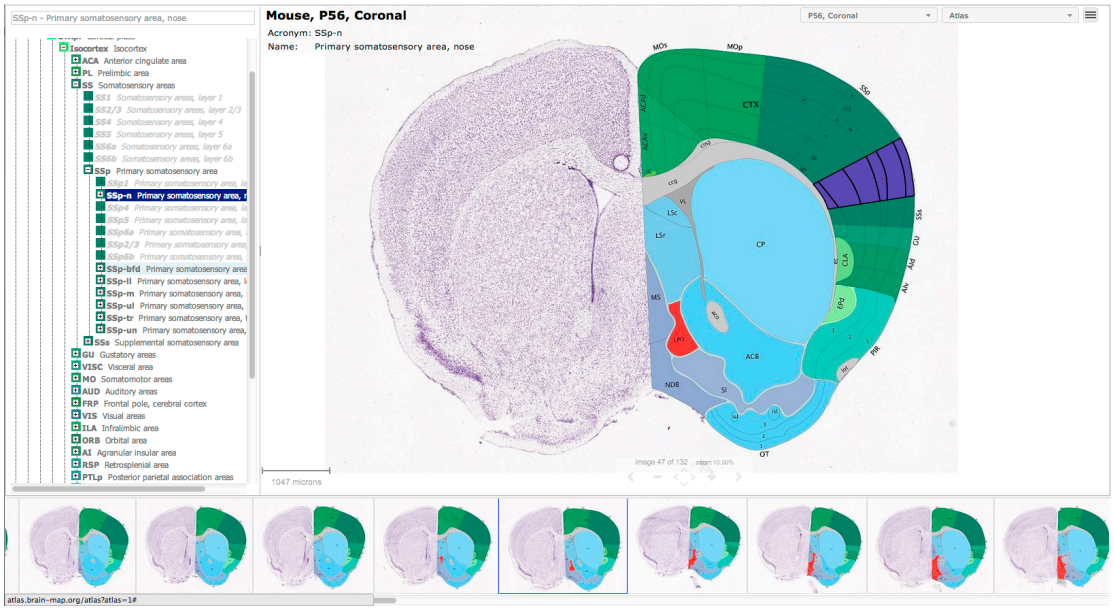

Figure S-12: Primary somatosensory area, nose (a radial sub-area of the "Primary somatosensory area" parent shown in the previous figure)

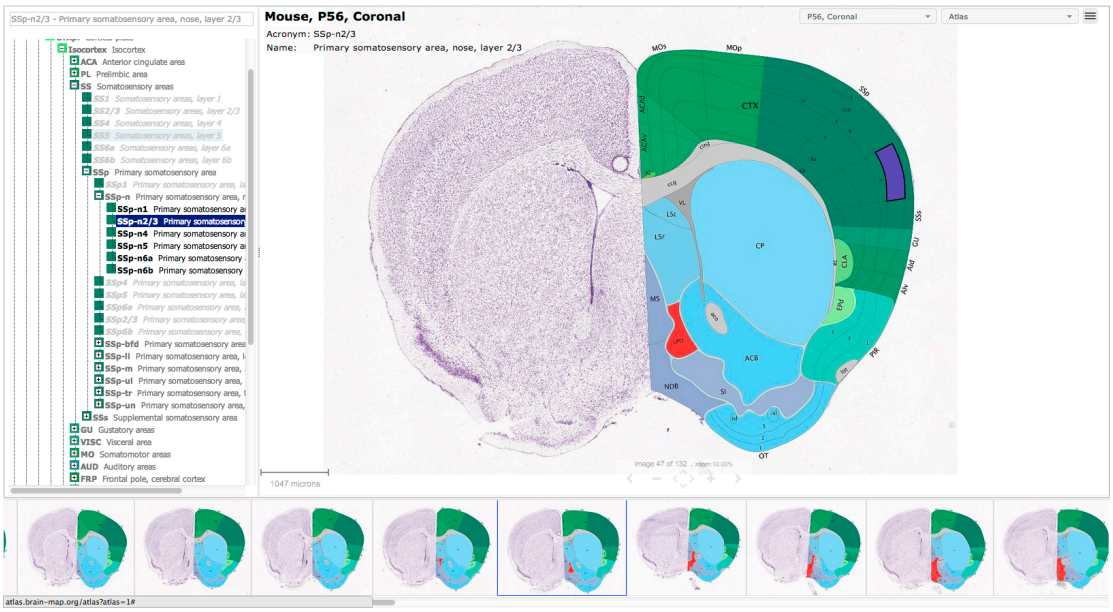

Figure S-13: Primary somatosensory area, nose, layer 2/3 (a concentric sub-area within a radial sub-area of the "Primary somatosensory area")

This example shows that ions specific to a concentric layer of the brain, but not specific to a radial sub-division of that brain, are difficult to approximate using the AMBA. The AMBA essentially forces a non-hierarchical anatomy structure into a hierarchical representation, and as a result loses some of the parent-daughter region relationships. This is an example of a biological region that can be easily recognized, but is not explicitly represented in the atlas and therefor hard to label with a single name. This sort of biological region cannot be captured adequately using a univariate approach, and this example further illustrates the need for multivariate interpretation algorithms capable of handling multi-membership properly.

With a multi-membership aware interpretation method, such as the one developed in this paper, there are currently two ways of handling these types of regions:

1. Complement the AMBA with additional patterns that group sub-patterns in different combinations (could be done exhaustively). For our proof-of-concept demonstrations we wanted to use the AMBA as-is, so this option is not explicitly discussed here.
2. Use the regularization parameter γ of our method as a dial that can be set to "make the interpretation as specific as possible" (which would grab the cortical layers specifically) or "make the interpretation high-level" (which would grab the overall encompassing region), as discussed in section S-6.

It is valuable to note that due to the speed and parallelism of the interpretation method, it is perfectly feasible to collect interpretations using multiple regularization parameter values (from global all the way to very specific). It would just require storing more anatomical membership tables, but would be a means

of ‘auto-discovering’ the sub-categorizations that have been lost due to the hierarchical structure of the AMBA.

In short, although the hierarchical structure of the AMBA is a feature external to our method, our interpretation method provides the parameter to either ask for an interpretation of a particular kind (from global to more specific), or the speed and performance to interrogate the same IMS measurement from a bunch of different viewpoints and collect the data for assessment in a later phase.

8. Automated Anatomical Interpretation of Non-Conforming Ion Images

This section describes how the automated anatomical interpretation method handles ion images that do not conform to anatomy, or that do not exhibit uniform intensity distribution within known anatomical structures.
The interpretation method we developed finds for each empirically acquired ion image in an IMS experiment, the best possible approximation using only the vocabulary of reference anatomy patterns in the atlas. This means that for each ion image in the IMS experiment, there is a measure of 'distance' between the measured ion image and its approximation using anatomy from the atlas. This distance is expressed in terms of ion counts, and is implicitly provided by our algorithm for each individual ion species in the study. These distances can be directly compared across all ion images in an IMS experiment.

Two possible scenarios unfold:

1. Small distance between ion image and anatomical approximation:
   The spatial distribution of this ion can be well approximated using the atlas
   ⇒ strong automated anatomical interpretation of this ion species is possible using the standard anatomy provided by the atlas.

2. Large distance between ion image and anatomical approximation:
   The spatial distribution of this ion cannot be well approximated using the atlas
   ⇒ the anatomy atlas does not provide the tools to adequately describe the spatial distribution of this ion species.

In other words, the automated anatomical interpretation algorithm can be used as an automated filter that separates ion species that conform to known anatomy from those ion species that deviate from known anatomy.

This means that the algorithm we describe here does not only provide a means for automated anatomical interpretation of ion images, but implicitly also provides a means of discovering which ion species deviate substantially from known histology. As such, our method can be used to automatically detect which ion species are distributed into subdivisions, subareas, or gradients within anatomical regions that according to the atlas are supposed to be homogeneous. Essentially, our method gives an automated anatomical interpretation for those ions that can be explained using the atlas, and additionally highlights those ions in which the IMS measurements show patterns that cannot be explained by known anatomy (e.g. disease-specific chemical patterns, previously unknown chemical subdivisions within tissue areas, etc.).