

# On Gabor Wavelet-based Image Processing for Nissl-Stained Rat Brain Slices

Kazunori Okada<sup>†</sup> and Michael J. Lyons<sup>‡</sup>

<sup>†</sup> Laboratory of Computational and Biological Vision  
University of Southern California  
HNB228 Los Angeles, CA 90089-2520 U.S.A.

kazunori@organic.usc.edu

<sup>‡</sup> ATR Media Information Science Lab  
2-2-2 Hikaridai Seika-cho, Soraku-gun  
Kyoto, Japan, 619-0288

mlyons@atr.co.jp

## Abstract

This article presents our investigation on applying the Gabor wavelet-based image encoding and processing technique to the nissl-stained rat brain slice images. In past, the Gabor wavelet-based method has been successfully applied to the task of face recognition. Our main question is, however, how universal the Gabor wavelet-based method is; can we apply the method to other classes of object without object-specific optimizations? We tested a modification of the original face recognition system in two tasks with the brain slice images: slice classification and automatic slice registration. We empirically showed that the modified classification system with the voting algorithm performed very well with 100% success rate and our preliminary system for the automatic registration performs satisfactorily. Results of our similarity analyses indicated that these successful performance appeared to be attributed to information in the high-frequency range (subcortical level) rather than in the low-frequency range (cortical contour). These results not only supports universality of the Gabor wavelet-based method but

also encourage further development of image processing systems for other biomedical data and tasks using the same technique.

## 1 Introduction

In the field of brain science, there has recently been increasing demand for useful image analysis tools to aid their investigations [3]. There are many previous studies that have addressed this issue especially in the area of biomedical image processing. These studies often target specialized images, such as MRI or PET images, that require high operational and computational costs. In the area of anatomical studies of brains, on the other hand, various staining techniques have been developed to visualize subcortical areas within sliced brain tissues in the ordinary 2D digital image format. These techniques are widely used due to its low cost, however, there is a general lack of computational studies which aim to develop tools for analyzing these images. Automation of otherwise labor-intensive manual analysis of them will help to standardize the technique and facilitate the further advancement of the field.

A decade ago, Lades et al. [2] have proposed a general 2D view-based object recognition method which exploited the dynamic link architecture. The groups lead by Dr. von der Malsburg at University of Southern California and Ruhr-Universität Bochum have successfully demonstrated applicability of this method for human face recognition problem [6, 4]. In this method, as its core, a filtering operation with Gabor kernels is used to extract relevant information of objects in the images. The resulting Gabor jet-based object representation allows a pixel in the 2D image to be associated with a feature vector called *jet* containing spatial frequency information of surrounding local region of the pixel. This representation scheme enriches general representation power and makes it more flexible than other conventional encoding methods.

The purpose of this study is to investigate feasibility of this method for analyzing 2D brain images created by a specific staining technique: nissl-stained rat brain slice images (for example, see Fig. 3). These nissl-stained images are, however, greatly different from the human face images. Fig. 1 illustrates qualitative difference of these two object classes. Can an image




Figure 1: Comparison of the nissl-stained images and facial images.

processing system optimized for recognizing the face images deal with the nissl-stained slice images? A series of careful numerical experiments are conducted to answer this question.

In the next section, general background of the method using the Gabor kernel filtering is described. The third section introduces our prototype system for identifying the depth position ( $z$ -coordinates within a 3D stack of 2D atlas images) of test images and presents its empirical evaluations. The fourth section proposes a novel method for automatically registering, or aligning, atlas images into a physical 3D coordinate system. We discuss our results and future work in the last section.

## 2 Gabor Wavelet-based Image Representation

In this study, we utilize a linear filter operation, a convolution of a 2D image with a bank of multi-scale and orientation kernels, for extracting pictorial and structural feature information of brain regions,

$$(WI)(\vec{k}, \vec{x}_0) := (\Psi_{\vec{k}} * I)(\vec{x}_0) := \iint \Psi_{\vec{k}}(\vec{x}_0 - \vec{x}) I(\vec{x}) d^2 x. \quad (1)$$

A family of kernels  $\Psi_{\vec{k}}$  used in this study is Gabor-based wavelets with a DC free component. They take the form of a 2D complex plane wave

restricted by a Gaussian envelope function,

$$\Psi_{\vec{k}}(\vec{x}) := \frac{\vec{k}^2}{\sigma^2} \exp\left(-\frac{\vec{k}^2 \vec{x}^2}{2\sigma^2}\right) (\exp(i\vec{k}\vec{x}) - \exp(-\frac{\sigma^2}{2})). \quad (2)$$

Image representation scheme using this approach has been shown to have an advantage in general pattern recognition tasks. Moreover, the sensitivity profile of neurons found in visual cortex of higher vertebrates can be well approximated by Gabor wavelets [1]. After transforming a 2D image by the filtering operation, each pixel in the image will be associated with a feature vector called *jet*,  $\mathcal{J}_{(i,j)}$ . A jet consists of complex projection coefficients with a bank of kernels whose size and 2D orientation vary. As a result, a jet contains multi-scale and orientation feature information of a local region around the sampling pixel. For example, a jet sampled on a right eye of a facial image may include information from left eye or nose. This overlap, or redundancy, helps us to compare a pair of images when there is a partial occlusion or noisy correspondence of sampling points. This makes the scheme a very robust representation method.

After this encoding process, a grid graph is placed on the input image, and only jets at nodes of the grid are sampled and stored in order to conserve memory space. We call this set of discretely sampled jets an *entry*. In order to compare a pair of entries, local similarity scores between pairs of corresponding jets are computed by the cosine (normalized dot-product) of the jet pairs,

$$S_v(\mathcal{J}^I, \mathcal{J}^M) := \frac{\mathcal{J}^I \cdot \mathcal{J}^M}{\|\mathcal{J}^I\| \|\mathcal{J}^M\|}. \quad (3)$$

The similarity score ranges between 0 to 1, where 1 indicates that the pair of jets are identical. Note that the jet is a vector of complex numbers therefore each jet coefficient has magnitude and phase components. For this study, we subjected only magnitudes of the jet to this similarity function, ignoring the phases.




Figure 2: Database of the nissl-stained rat brain slice images. The first row: plate 15 to 24; the second row: plate 26 to 35.

### 3 Identification of Nissl-Stained Brain Slice Images

#### 3.1 Problems

In this section, we construct and evaluate a classification (identification) system, which finds the slice image in the brain atlas most similar to an arbitrary input slice. This is analogous to the face recognition task where each probe (input slice) is searched against album entries (atlas slice image stack) of known people. One of the major difference between images of human face and rat brain slice is the treatment about the coordinate system. Given a stack of face images of different people, there does not exist an intrinsic geometric coordinate system that relates these images. On the other hand, brain slice images have intrinsic geometrical relationship to a 3D Euclidean coordinate system since they are parts of a single object. In fact, the main purpose of such a brain atlas is to aid researchers to locate specific brain tissues referring to the coordinate system. In this study, *slice depth*, the plate number of each slice, is used as the third axis on top of the 2D image plane, forming a 3D coordinate space. This section provides performance evaluation of our classification system as well as analysis of characteristics of the brain slice images.




Figure 3: 2D image of nissl-stained rat brain slice: the plate 30 of the Swanson's atlas.

### 3.2 Methods

The brain slice images used in this study are electronically scanned from original positive copies of the Swanson's Rat Brain Atlas [5]. 21 images between the plate 15 to 35 are used in this study. Fig. 2 summarizes the database of these digitized images (only 20 are shown). Note that, in comparison with the facial images in Fig. 1, some of consecutive nissl-stained images are so similar that it is hard to distinguish one from others. Fig. 3 shows the plate 30 in a larger resolution. In our experiments, we used geometrically fixed grid graph across the slices since all the images are said to be registered properly to a fixed 3D coordinate system. Before our classification experiments, each of these 21 images is processed in the following manner.

1. normalize the size of image into one of the three power of two squares: 128 by 128, 256 by 256, 512 by 512.
2. Fourier transform the image.

3. perform the convolution of the image in the Fourier domain (Morlet Transformation) with 40 (5 scales and 8 orientations) Gabor kernels. We set  $\sigma$  by  $2\pi$ , maximum norm of vector  $\vec{k}$  by  $\frac{\pi}{2}$ , and step factor of scale by  $\frac{1}{\sqrt{2}}$ .
4. place a fixed rectangular grid graph with 25 (5x5) nodes.
5. create an entry by sampling 25 jets from the graph nodes.

The classification test is performed, for all the 21 entries, by taking one entry from the album and matching it against other 20 entries. We call the album entry most similar to an input the *best match*. When the best match is one of the two adjacent slices of the probe, we consider it was a correct classification. The best match is searched by using two algorithms: Averaging and Voting,

**Averaging:** averaging locally classified slice depths.

1. given a probe entry, compute jet similarities on each node over album entries.
2. at each node, sort 20 similarities across the album entries.
3. at each node, take slice depth of the local best match.
4. average the local slice depths over the nodes.

**Voting:** voting among locally classified slice depths.

1. perform the above 1-3.
2. perform the majority voting among the local slice depths and take the majority slice depth as the depth estimate of the probe.

The second algorithm is more relaxed than the first one in the sense that it can exclude the grossly wrong local slice depth estimate (suppressing outliers).

### 3.3 Results

#### 3.3.1 The Averaging Algorithm

The averaging algorithm is first evaluated. Fig. 4 shows error deviations of slice depth between the average estimate and each probe's ground-truth (correct) depth. Results show that 1) the beginning and the end of the slice stack




Figure 4: Deviation of slice depth between estimate and ground truth using the averaging algorithm.

tend to be identified towards the middle slices and 2) the error deviations seem to be linearly correlated to the slice depths. The first observation can be explained by that the distribution of album entries poses a probabilistic bias towards the middle of the sequence so that errors of the average depth pull the estimates towards the middle. The linearity indicates that structural information in the nissl-stained images is detectable by our method, however this algorithm is not appropriate for the classification task due to the significance of the errors.

### 3.3.2 The Voting Algorithm

Next, we evaluate the voting algorithm. Fig. 5 shows percentages of correctly classified nodes over the total of 25 nodes for each probe entry. Results show that 1) the correct percentages deviate largely across the probe entries and 2) smaller image resolution seems to give higher percentages. Average of the percentages across the probes were 57.0%, 41.1%, and 32.4% for the 128 by 128, 256 by 256, and 512 by 512 image resolutions, respectively. When the image resolution was the smallest (128 x 128), the percentages for some






Figure 5: Percentage of correctly estimated nodes for each probe entry.

probes were significantly increased, marking 84% as highest. Fig. 6 shows the error range of the local voting results: histograms of locally estimated slice depths as a function of signed error (slice depth distance between the local estimate and ground truth). Regardless the image resolution, the figure showed two high and sharp peaks next to zero that indicates no error.

Since the above experiments strongly indicated a favor to the voting algorithm, we constructed a classification system based on the algorithm and evaluated its performance. Tab. 3.3.2 provides the correct classification/identification rates for different image resolutions. Results were promising; the system gave perfect classification when the 256 by 256 image resolution was used. When the smaller resolution was used, the correct classification rate was lowered but remained relatively high. On the other hand, when the larger resolution was used, the rate dropped largely. The fact that the system performed better in the 256 by 256 resolution than in the 128 by 128 resolution was surprising because our results shown in Fig. 5 predicted otherwise. In the 128 by 128 resolution, probe 20 and 30 were misclassified as 18 and 32, respectively. The error is insignificant as far as the performance is concerned.




Figure 6: Histogram of locally estimated slice depths as a function of corresponding signed errors for the 3 different image resolutions.

Table 1: Correct classification rates by the voting system.

Input Image Size	128 x 128	256 x 256	512 x 512
Recog. Rate (%)	90.5	100	76.2

In order to investigate the question posed by the previous experiments, we next studied distribution of the vote counts over each album entry. Fig. 7 displays 2D contour plots of the vote count distribution (histogram of locally estimated slice depths) over 21 probes and 21 albums for the 3 different image resolutions. In the perfect classification situation, for each column, the immediately adjacent region of the diagonal should show the highest peak. The plot for the 128 by 128 resolution showed almost no peaks in the non-diagonal area, clearly agreeing with the results shown in Fig. 5. The plot for the 256 by 256 resolution had more peaks in the non-diagonal area than one for 128 by 128, however the diagonal area of 256 by 256 seems to be more distinctive, or stronger, than that of 128 by 128, facilitating the higher

classification performance. The plot for the 512 by 512 resolution showed more peaks in the non-diagonal area and less strength of the diagonal area, causing the low classification performance.

### 3.3.3 Implications from the Relationship between System’s Performance and Image Resolution

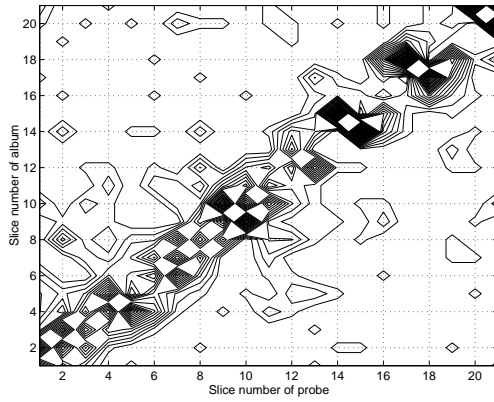
The observed relationship between the system’s performance and image resolution provides interesting insight for the characteristics of the nissl-stained images.

Note that parameters of the Gabor kernels described in Sec 3.2 were given without an object-specific tuning; they were tuned for maximizing face recognition performance in the 128 by 128 image resolution. Therefore, the frequency-domain range covered by the kernel set used in the experiments are tuned to describe the specific object in the specific size. The fact that increasing image resolution helped to improve the classification performance indicates that higher frequency-domain information contributes to characterize intrinsic nature of the target object more in the brain slice images than facial images. This agrees with the fact that subcortical regions in the nissl-stained images are visualized by the random-dot-like high-frequency patterns (for example, see Fig. 3 and 2) while face images are more obviously structural, providing mainly smoother pixel-value gradients in lower-frequency range.


This observation leads to two favorable implications: 1) our Gabor jet-based image representation can extract meaningful structural information from seemingly random high-frequency patterns and 2) it was the structural information in sub-region level of the brain images rather than global contour of the brain which contributed more to the successful classification performance. The former suggests universality of the representation scheme, while the latter indicates that our method may be applicable for analyzing subcortical regions which is more challenging task.

The fact that the highest performance was reached when the median resolution was used is favorable by a practical reason. A possible shortcoming of the representation scheme is its high computational complexity. The cost increases exponentially with the image resolution due to Fourier transformation. Although the current state-of-the-art computers can perform the process by a fraction of second, this leads to a problem when computing a

128 x 128



256 x 256



512 x 512




Figure 7: 2D contour plots of node-wise vote counts (histogram of locally estimated slice depth) over probe and album entries for 3 different image resolutions.




Figure 8: Histogram of Gabor jet similarity value distributions over a cross-match for the 3 different image resolutions.

large number of high resolution images. The ceiling of the effective image resolution helps to mitigate this scaling problem.

### 3.3.4 Similarity Values of Gabor Jets from Different Entries

Results of the above experiments successfully demonstrated the effectiveness of the Gabor jet-based representation for the task of our focus. In order to further investigate properties of this representation scheme, we now analyze Gabor jet similarity values computed during the classification task.

Fig. 8 shows histograms of all (local) jet similarity values computed for a cross-matching: comparing each probe entry with the rest of 20 entries. Such a plot illustrates a range of similarity values that all pairs of jets in this dataset cover. Regardless the different image resolutions, the average similarity value was about 0.76. There was no significant difference in the distribution property especially between the 256 by 256 and 512 by 512 image resolutions. The wide range of similarity values suggests a good descriptive power of the metric and representation scheme of our system.

Fig. 9 displays average similarity values as a function of slice depth dis-




Figure 9: Average Gabor jet similarity values over slice depth distances for 3 different image resolutions.

tances, illustrating the relationship between the jet similarity values and z-distance in the brain coordinate space. All the jet similarity values in Fig. 8 are reorganized and averaged for each slice depth distance. Results show clear tendency that nearby slices are more similar than farther-distanced slices. Before the distance reaches 10 slices, the similarity value smoothly decreases as the distance increases. After passing the threshold, the value seems to reach the minimum. The plots for the 256 by 256 and 512 by 512 resolutions are very similar despite the significant difference in performance of the classification system based on the voting algorithm. This indicates that absolute values of the raw jet similarities may not be the main contributor to the successful classification performance.

Note that our classification system was based on the nearest neighbor scheme (winner-take-all). In the averaging algorithm, such a scheme can be effective if similarity value of the best match is distinctively higher than the second best match in general. Although this argument does not directly apply to the voting algorithm, such situation does not harm the voting method.

Fig. 10 illustrates such discriminability of the best match from the second

best match for the 3 image resolutions. Each plot compares two histograms of similarity values computed from pairs of entries with distance of 1 slice and of 2 slices. Regardless of the resolution, those two similarity histograms were enormously overlapped. For the 128 by 128 and 256 by 256 resolutions, the two distributions appeared to be slightly shifted. This result agrees with the fact that the averaging algorithm did not perform well, however we cannot draw much insight towards the voting algorithm. In order to improve the averaging algorithm, we at least require more sophisticated classifier than our nearest neighbor one.

### 3.3.5 Similarity Values of Gabor Jets within the Same Image

Next, we directly address characteristics of the nissl-stained images by studying Gabor jet similarity values of arbitrary pairs of jets within a single image. Each Gabor kernel has its support: the range of image region around the sampling pixel in which information is taken into account or “seen”. Because of this, two jets sampled at nearby locations have certain amount of information overlap whose extent depends on the support size and the sampling distance. This makes the similarity distribution curve as a function of pixel distance smoother than by using methods with correlation-based metric over raw pixel values, providing a basis for more accurate tracking of landmark points. With this knowledge, one can investigate what Gabor jets are seeing by studying the similarity distribution over different sampling distances, given a fixed kernel support. We performed the same analysis to three different image types: 1) the nissl-stained image, 2) facial image, and 3) random-dot image. Comparison of the analysis results can clarify characteristics of these images.

Fig. 11 shows average Gabor jet similarities as a function of pixel distance in a brain slice image. We used the 256 by 256 image resolution in this experiment. As an example, the plate 21 from the brain atlas is used to compute Gabor jets at all pixel locations. Then similarity values in all permutation of jet pairs are computed without using the grid graph and are reorganized as a function of pixel distance. Results show that 1) the similarity value rapidly decreases as the distance within the image increases, 2) after reaching 50 pixel distance, the similarity value continues to decrease in a somewhat lower rate, 3) after reaching 250 pixel distance, the similarity increases. The last observation is due to an artifact of the wrap-around boundary condition used in our preprocessing. The first and second observa-

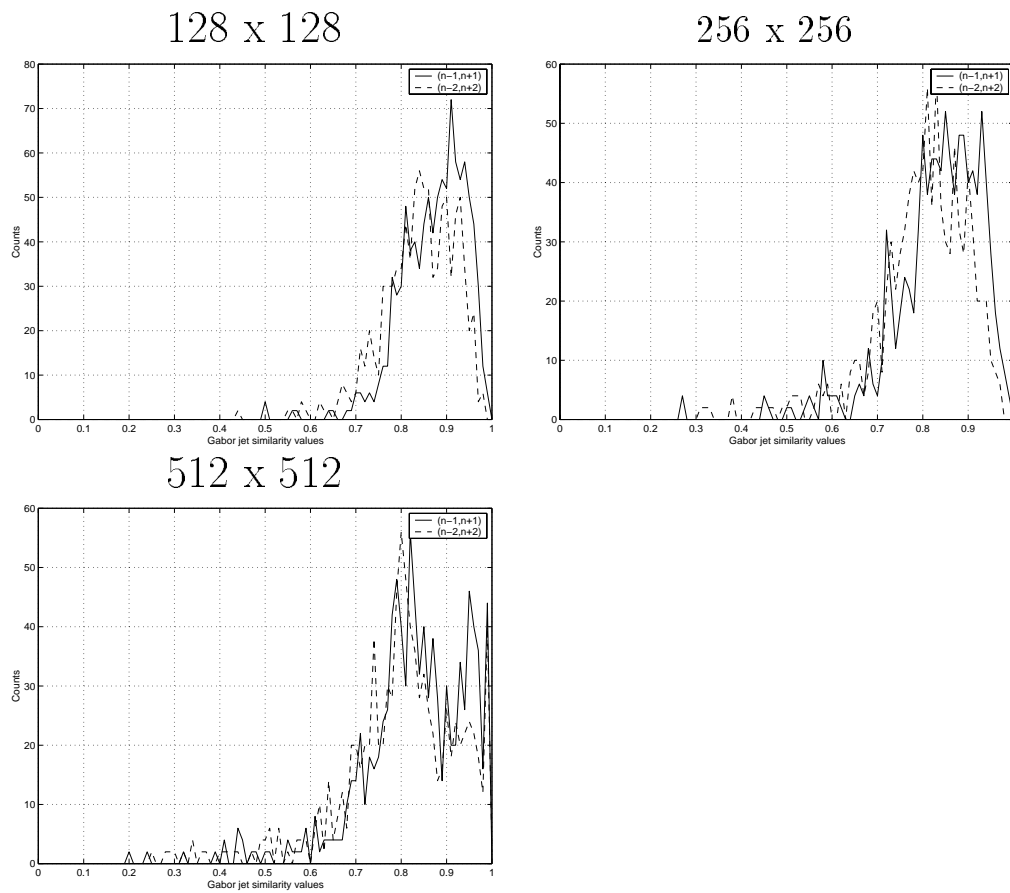


Figure 10: Comparison of Similarity Histograms computed from entry pairs with distance of 1 slice and 2 slices for the 3 image resolutions.



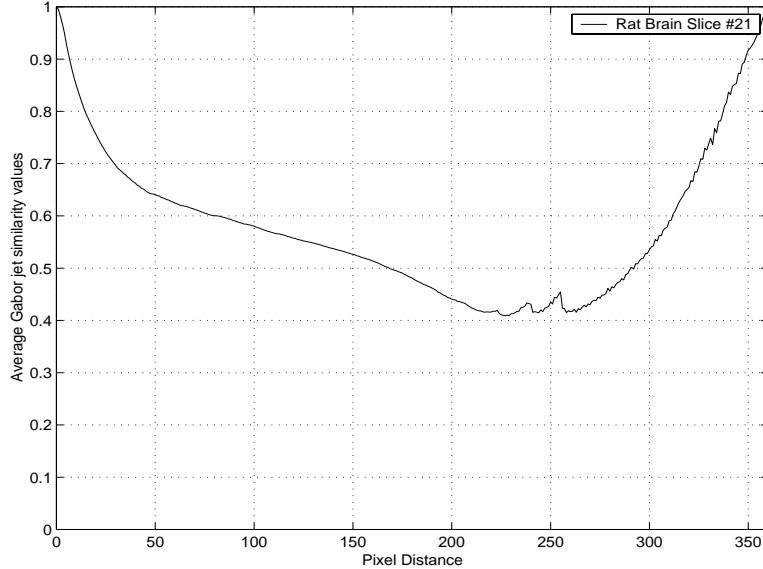


Figure 11: Average Gabor jet similarities as a function of pixel distance in the Swanson’s plate #21.

tions indicate that Gabor jets seem to well capture structural information of different regions in the image; jets at different location see innately different patterns therefore lowering similarity values.

Fig. 12 compares the average similarity distributions for the 3 different image types: the brain slice image, human face image (from FERET database), and random-dot 2D image. Within a 15 pixel-distance range, property of the brain image was in-between the human face and the random-dot images. The information overlap of nearby jets explains the relatively high similarity values in a small distance range. The case of the random-dots exemplifies this point most clearly as there are no pixel-level correlations. Similarity values were higher than 0.9 within 9, 7, and 5 pixel distances for human face, rat brain slice, and random-dots, respectively. Beyond 15 pixel distance, however, the similarities of the brain image and the human face continued to decrease similarly, while that of the random-dots became distinctively different (a constant similarity at 0.8 regardless of pixel distance). These results indicate that structural information of the brain image is well captured by Gabor jets especially in lower-frequency range (>15-25 pixels). The similar




Figure 12: Average Gabor jet similarities of the rat brain slice, human face, and random-dot images.

results for the brain and face images suggest that our method’s success on the face recognition task may translate well to the nissl-stained images.

## 4 Automatic Registration of Nissl-Stained Brain Slice Images

### 4.1 Problems

Accurate registration of brain slice images into a fixed 3D coordinate space is a very critical issue since the main purpose of the brain atlas is to precisely point a tissue in the space in order to locate the same region in different individuals. Currently, this registration requires meticulous manual work and careful use of heuristics for defining such a 3D coordinate system. This is a very difficult task because each brain has individual differences and it is a non-rigid object. Our purpose here is to propose an automatic method for this brain slice registration task which provides an alternative to the heuristically defined 3D coordinate system.

## 4.2 Methods

We assume that 1) we are provided with a stack of brain slice images in a correct order but 2) each slice was subjected to scaling and translation during the slicing and imaging processes. The registration problem in this case is reduced to solving the correspondence problem of the successive pairs of the slice images. In order to demonstrate a feasibility of this approach, we constructed a simple prototype of rigid graph matching algorithm similar to the one used in [2], but with less degree of freedom: only considering the translation factor. The prototype utilized the 256x256 image resolution and the same relevant parameters as the previous section, including the 5x5 square grid. The simple hill-climbing algorithm was used to solve the 2D search problem,

1. compute jets on all pixels in a starting image (create a jet-image).
2. for the first image, place the grid manually, resulting in an entry.
3. create a jet-image from the next image of the sequence.
4. search the second jet-image with the previous entry for finding a grid position that gives the maximum similarity.
5. continue the processes 3 and 4 until processing all the images.

After placing a grid graph on the first image, the registration process proceeds automatically. Search space can be limited into a fixed window due to the proximity of consecutive slices. After we placed the grid on every images, we can register them by translating each image (morphing if we also consider scaling and other non-rigid distortion in future) such that the grid graph appears at the fixed position across all images. We consider the resulting stack as a definition of the 3D coordinate system.

## 4.3 Results

Fig. 13 displays several examples of the 2D search space as contour plots of 2D similarity surfaces. The plots are shown as a function of the horizontal (i-step) and vertical (j-step) displacement of the graph for 4 different pairs of consecutive slices. The search was restricted within 60 and 20 pixels along

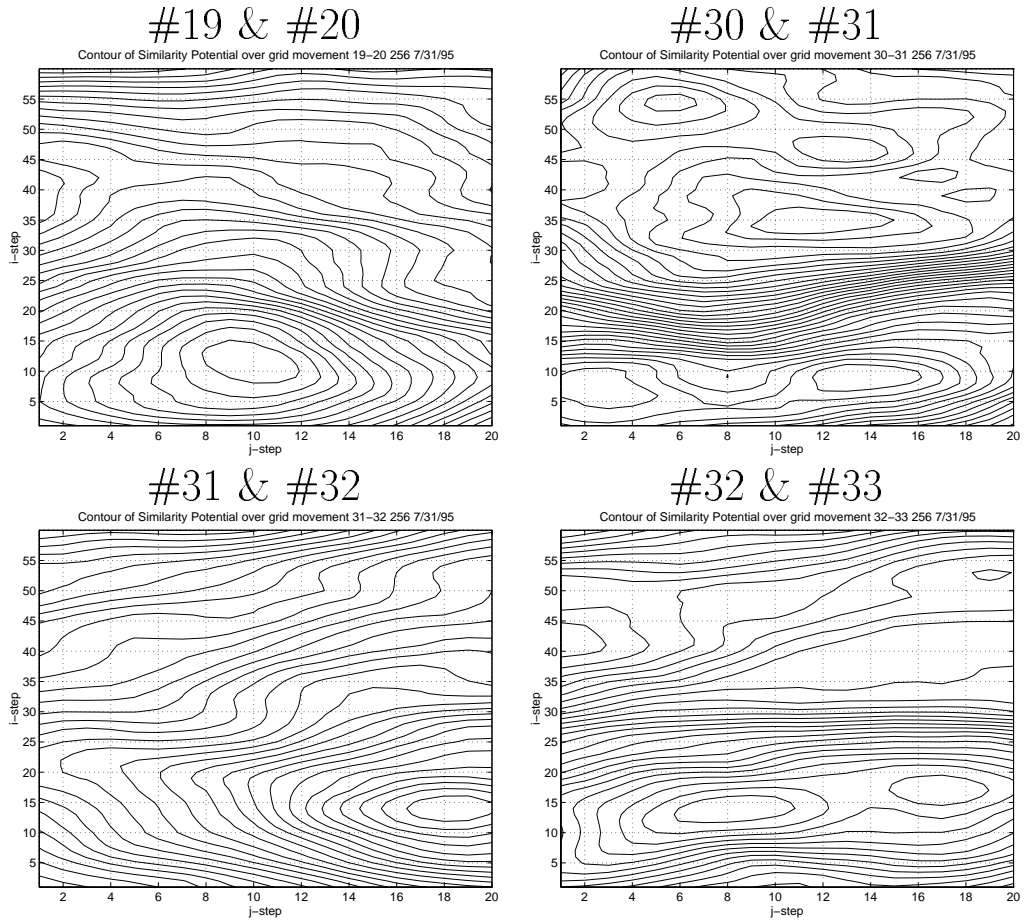


Figure 13: Examples of 2D search space visualization for the automatic registration.

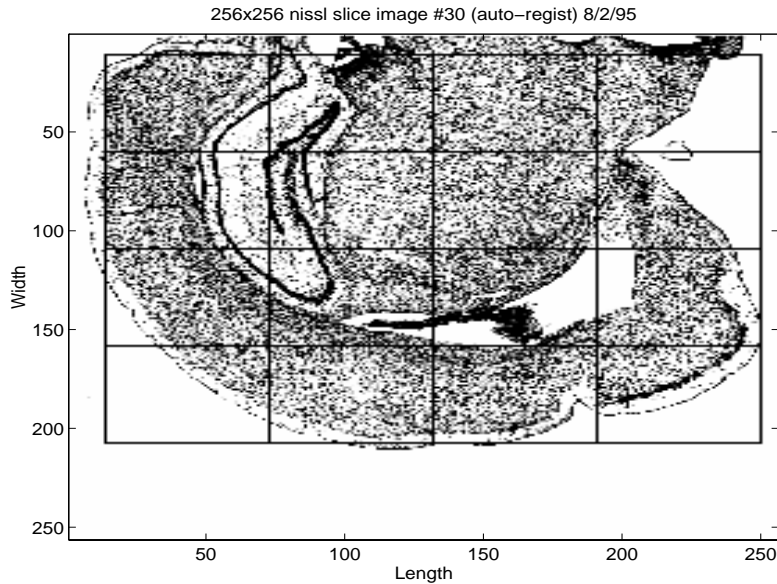


Figure 14: An example of automatically located grid on the brain slice image (#30).

the  $i$  and  $j$  direction, respectively. In general, the similarity surfaces were very smooth. This helps our very simple search algorithm to find the global maximum readily without using high cost search algorithms. The examples in the left column show cases where there were no local maxima for the search to be trapped. On the other hand, those in the right column show the worst cases in our experiment. In our experiment with the 21 image sequence, we observed that our method can find the global maxima in the efficient cost. In general, ill-placed initial conditions may be a potential problem, however we did not experience this problem because the slice covers large portion of an image. Fig. 14 displays an example of the automatic registration result, showing the found graph location in the image.

## 5 Discussion

We have investigated for applying the Gabor wavelet-based image representation and processing method, whose advantage has been shown for the face recognition task, to process images of a completely different object:

nissl-stained rat brain slices. We have demonstrated two applications of our method: classification among known slices from the atlas (resulting in information of its rough position in the brain) and automatic registration of slice sequence.

Our experimental results for the classification task were very promising. We have achieved 100% correct classification and qualitatively demonstrated its advantages by various analyses. Note that this was achieved by using the system without any object-specific optimizations for its encoding part! Our system used the voting algorithm for classification while the averaging algorithm was typically used in the Gabor wavelet-based face recognition systems [2, 4]. However, parameter settings of our system's encoding part (computation of Gabor jets) were not modified from the original that was optimized for maximizing performance of the face recognition task. This suggests universality of the image encoding method with Gabor wavelets. The actual performance statistics must be treated with a caution because we used only a small number of images. We must further evaluate the proposed system with much larger data set, however we are optimistic that some problem-specific optimizations can improve the system's performance to the level where location information with a physical unit can be provided with a certain error bound.

The current state of the automatic registration system is still preliminary. Although we have successfully demonstrated an automation of the task, there is so far no guarantee that the resulting 3D coordinate system is faithful to the actual brain coordinate system which is required for the result to be useful. We could not perform error analysis, beside our visual inspection of the results, because of the lack of data to compare our result with. The rigid-body assumption was also a very strict one. For our system to be practical, we must consider more sophisticated matching method in order to account for the complicated image distortions. These issues remain as open problems.

Topics for our future research include 1) if our classification method can find the slope of given slices by looking at 2D distribution of the locally estimated slice depths, 2) if the method can accommodate variations in different individuals and animals, and 3) if the method can be extended to analyze subcortical regions explicitly. Our analyses of the classification system indicated that our successful classification performance appeared to be attributed to information in the high-frequency range, which is favorable towards the third topic. However, all these issues must be resolved in order to make our

method applicable to more practical scenarios.

## References

- [1] J. G. Daugman. Complete discrete 2-D Gabor transforms by neural networks for image analysis and compression. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 36:1169–1179, 1988.
- [2] M. Lades, J. C. Vorbrüggen, J. Buhmann, J. Lange, C. von der Malsburg, R. Würtz, and W. Konen. Distortion invariant object recognition in the dynamic link architecture. *IEEE Transactions on Computers*, 42:300–311, 1993.
- [3] M. I. Miller, G. E. Christensen, Y. Amit, and U. Grenander. Mathematical textbook of deformable neuroanatomies. In *Proceedings of National Academy of Science U.S.A.*, volume 90, pages 11944–11948, 1993.
- [4] K. Okada, J. Steffens, T. Maurer, H. Hong, E. Elagin, H. Neven, and C. von der Malsburg. The Bochum/USC face recognition system: And how it fared in the FERET phase III test. In *Face Recognition: From Theory to Applications*, pages 186–205. Springer-Verlag, 1998.
- [5] L. W. Swanson. *Brain maps: Structure of the Rat Brain*. Elsevier Science Publishers, Netherlands, 1992.
- [6] L. Wiskott, J.-M. Fellous, N. Krüger, and C. von der Malsburg. Face recognition by elastic bunch graph matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19:775–779, 1997.