A strategy based on maximum spanning trees to stitch together microscope images

Philippe Thévenaz[a], David Lambiel[b], and Michael Unser[a]

[a]Biomedical Imaging Group, École polytechnique fédérale de Lausanne (EPFL), Switzerland
[b]Chaire de recherche en géomatique cognitive, Université Laval, Canada

ABSTRACT

Assembling partial views is an attractive means to extend the field of view of microscope images. In this paper, we propose a semi-automated solution to achieve this goal. Its intended audience is the microscopist who desires to scan a large area while acquiring a series of partial views, but who does not wish to—or cannot—planify the path of the scan. In a first stage, this freedom is dealt with by interactive manipulation of the resulting partial views, or tiles. In a second stage, the position of the tiles is refined by a fully automatic pairwise registration process. The contribution of this paper is a strategy that determines which pairs of tiles to register, among all possible pairs.

The central tenet of our proposed strategy is that two tiles that happen to possess a large common area will register with higher accuracy than two tiles with a smaller overlap. Our strategy is then to minimize the number of pairwise registrations while maximizing the global amount of overlap, and while ensuring that the local registration efforts are sufficient to link all tiles together to yield a global mosaic. By stating this requirement in a graph-theoretic context, we are able to derive the optimal solution thanks to Kruskal's algorithm.

Keywords: Registration, Mosaic, Graph.

1. DESCRIPTION OF PURPOSE

Microscope images come in a large variety of modalities, but all suffer from the same, universal problem: the field of view is often extremely narrow when compared to the size of the sample to investigate. There exists, however, a trivial real-time solution that every microscopist quickly masters, which consists in scanning the sample and integrating the outcome. Unfortunately, the resulting mental representation cannot be communicated. To overcome this limitation, it is necessary to acquire and store a series of overlapping partial views, or tiles; then only, an image with an extended field of view can be assembled from them.

The assembly task is non-trivial, especially if high geometric accuracy is desired. Even when the scanning procedure involves motorized devices or piezoelectric actuators, it is not reasonable to expect that the overlapping tiles can be superimposed without some sub-pixel tuning. In a microscopy context, this sub-pixel tuning is particularly relevant because the scanning procedure is often organized in such a way that any given tile is surrounded by, and overlaps, several other tiles, which creates multiple cross-relations. To obtain a globally consistent image, it would then be desirable to perform the assembly by considering all the tiles simultaneously; unfortunately, this approach would outlast the user's patience, either because the number of all possible pairwise registration operations to perform would grow too large, or because the number of dimensions of the search space associated with the simultaneous registration of all the tiles would lead the optimizer to stall since the geometric parameters are strongly coupled. Therefore, for efficiency purpose, we argue that the tiles must be assembled pairwise to yield a local relative transformation, and that the number of pairs to consider has to be minimized; the global position of each tile is then given by stringing local transformations together. The downside is that the residual errors add up and propagate.

This situation differs from the creation of panoramic images (typically, landscapes) for which it is often true that only left and right neighboring tiles must be considered. In a microscopy context, the geometry requires more

Corresponding author: P. Thévenaz, EPFL/STI/IOA/LIB, Building BM-4.137, Station 17, CH–1015 Lausanne VD, Switzerland; phone: +41(21)693.51.61 (CET); fax: +41(21)693.37.01; e-mail: philippe.thevenaz@epfl.ch; home page: http://bigwww.epfl.ch/.

Medical Imaging 2006: Image Processing, edited by Joseph M. Reinhardt, Josien P. W. Pluim,
Proc. of SPIE Vol. 6144, 61442A, (2006) · 0277-786X/06/$15 · doi: 10.1117/12.651320

accuracy than for creating a panoramic landscape because of more exacting consistence requirements; conversely, panoramas face the difficult issue of blending and color grading, which are essentially irrelevant for microscope images because illumination is much better controlled. Another difference between the creation of panoramic images and the creation of microscope mosaics is that, in the latter case, every care is usually taken to ensure that the images are distortion-free, while this is much less the case for panoramas; moreover, the bulk of the geometric issues is a simple translation (and perhaps a limited amount of rotation) for microscopy images, while panoramas often require a richer transformation called a homography. In summary, the creation of microscope mosaics is made easier than that of panoramas by the absence of the need for color and distortion corrections (only translation and rotation are cared for), while it is made more difficult by the necessity of a high geometric accuracy to limit the propagation of registration errors. The sensitivity to these errors is much higher in mosaics than in panoramas because it is likely that the microscopist will choose a scanning path that folds back on itself.

The purpose of this paper is to present a semi-automated solution to facilitate the assembly of overlapping tiles in a microscopy context. The scenario that we have in mind is as follows: a microscopist explores a sample and acquires overlapping images during this free, unplanned exploration. At a later stage, the images are subjected to an interactive rough positioning before our solution kicks in and refines it. A final mosaic ensues.

2. METHODS

Suppose the three mutually overlapping tiles (A, B, C) must be assembled. Let U and V refer to any two tiles out of (A, B, C); their true pairwise geometric relation $\hat{T}_{U\to V}$ can be estimated by a registration procedure that yields a perhaps noisy $T_{U\to V}$. Moreover, let A be the tile that anchors the assembly. Certainly, we can assign the position of B with respect to A thanks to $T_{B\to A}$; same can be said of C and $T_{C\to A}$. But if we do so, we have placed B and C without benefitting from the contribution of the knowledge of $T_{B\to C}$ or $T_{C\to B}$; we have also ignored $T_{A\to B}$ and $T_{A\to C}$. (Because of measurement uncertainties, it is not necessarily the case that $(T_{V\to U} \circ T_{U\to V})$ results in the identity transformation.) In conclusion, by using just $T_{B\to A}$ and $T_{C\to A}$, we have created a mosaic but we have no insurance that it is consistent with the set of all potentially available pairwise estimated geometric relations, or that it is the best possible use of only two of them. For instance, we could have kept $T_{B\to A}$, while replacing $T_{C\to A}$ by $(T_{B\to A} \circ T_{C\to B})$ to obtain a perhaps slightly different result.

Suppose now that we can determine the *a posteriori* consistence of the solution, perhaps by measuring by how much a circular chain of transformations differs from the identity. Then, one possible way to do better in general would be to build a graph where the nodes correspond to the tiles, and where the edges correspond to the availability of pairwise geometric relations. We would then examine every possible spanning tree, and we would measure the degree to which the solution given by a spanning tree would be consistent with the remaining edges—those that do not belong to the spanning tree under examination. Unfortunately, maximizing this consistence over all possible solutions is not a realistic approach as the number of spanning trees is usually too large in practice, and it is required that every possible registration attempt be performed beforehand.

In this paper, we propose a solution that is inspired by the spanning-tree approach. We replace the direct determination of the *a posteriori* consistence of the solution by an *a priori* estimate of the accuracy of the measured relation $T_{U\to V}$. More precisely, we make the hypothesis that the degree of accuracy is proportional to the area of the overlap between U and V, which is assigned as a weight to the graph edge that links U to V. Then, Kruskal's algorithm[1] can be used to obtain the maximum spanning tree (see Appendix A). A major advantage of this *a priori* approach is that a rough guess $T^{(0)}_{U\to V}$ of the relative position of U and V is sufficient to estimate their amount of overlap, and no registration effort must be committed at this stage. (The fine-tuning of $T^{(0)}_{U\to V}$ to get $T_{U\to V}$ might result in long-range consequences, particularly if rotation is involved, while it doesn't have much of an impact on the local amount of overlap of $(U^{(0)}, V)$ compared to (U, V).) We defer the computational cost of high-quality registration until after the maximum spanning tree has been determined; but then, its knowledge gives us the opportunity to perform the minimum number of pairwise registrations, which is one less than the number of tiles.

Since a local misregistration between two tiles might have global implications, it is essential to strive for the highest-possible registration accuracy. To register U to V, we have therefore chosen an algorithm that relies on a high-quality interpolation model to ensure that any residual geometric error remains largely sub-pixel.[2] To gain

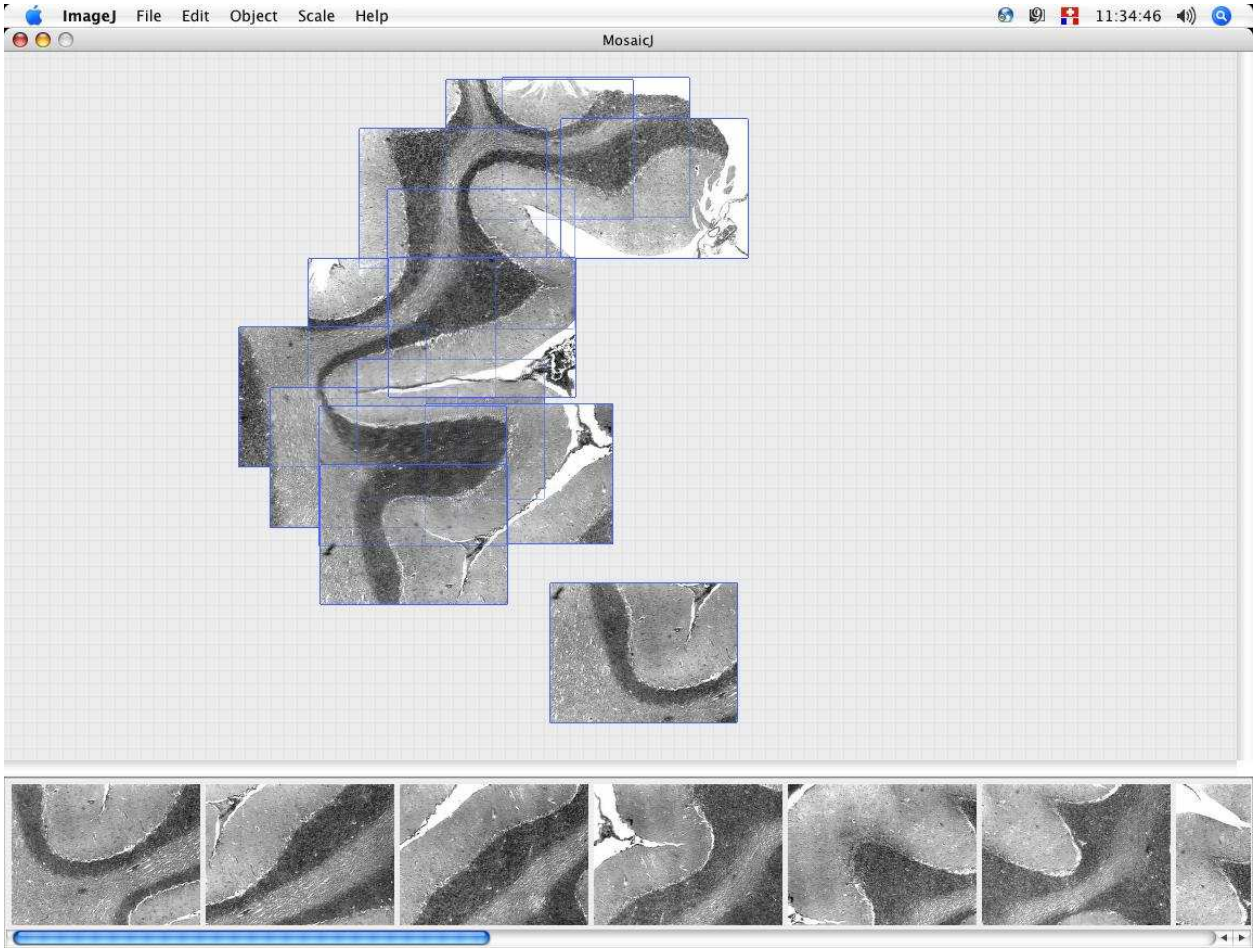

Figure 1. Interface of MosaicJ. The building of a rough mosaic of thirty tiles is under progress; about a dozen tiles have been placed already, and the next tile is been manually dragged to fit them. Each visible tile is the thumbnail of a (1,392 × 1,040) microscope image; the resolution of the thumbnails can be adjusted. Here, the tiles have been outlined for clarity.

in efficiency, only the area of overlap—along with some additional arbitrary margin—takes part in the registration effort. Robustness is achieved by a multiresolution strategy. The registration criterion is least-squares.

3. RESULTS

The method we propose is called "MosaicJ". We have implemented it as a plugin for the public-domain, platform-independent software ImageJ.[3] It relies on our registration engine, in the form of another plugin called "TurboReg". The two plugins are made freely available at http://bigwww.epfl.ch/, along with source code.

The interface has been streamlined to improve its accessibility to the non-specialist. Once a series of tiles has been made available, the user can organize them simply by dragging the mouse, thus interactively providing the rough initial solution that is used to determine the minimum spanning tree. This initial solution is translational only. Once satisfied with it, the user can request that a mosaic be produced. MosaicJ then orchestrates the registration operations performed by TurboReg, which involve translational and rotational fine-tuning of the mutual geometric relations between tiles. An optional blending mechanism ensures their smooth visual transition.

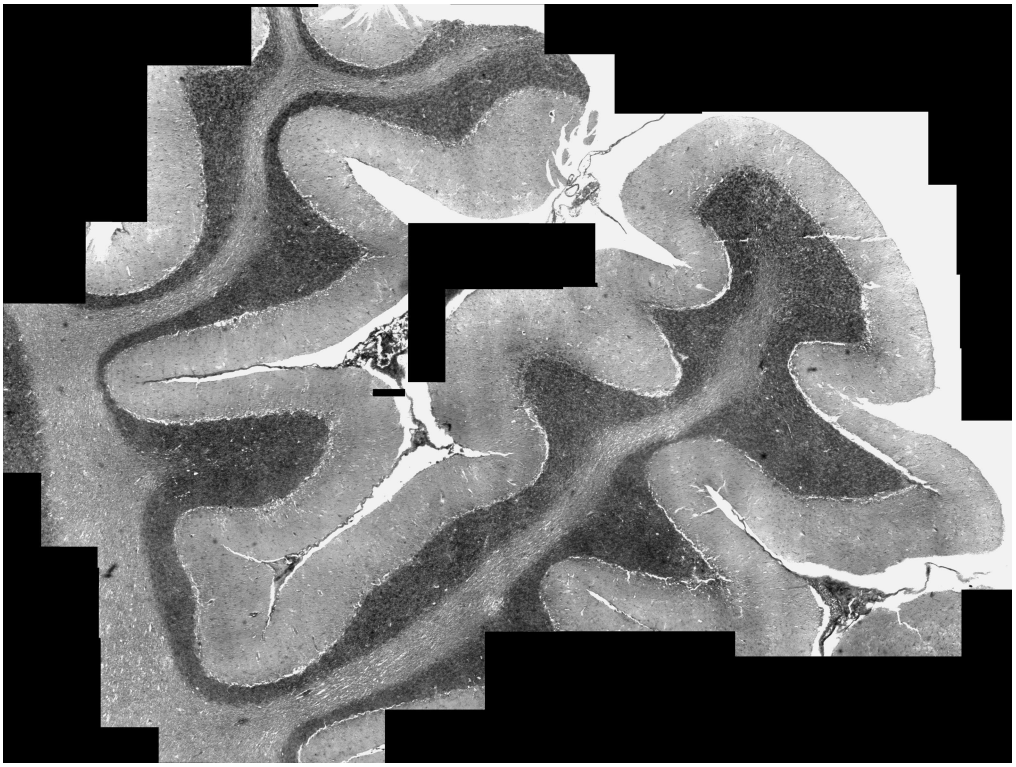

Figure 2. Final mosaic of a histological cryo-section of a human cerebellum (human t.s. routine stained). Thirty ($1{,}392 \times 1{,}040$) tiles were acquired while delineating part of a gyrus; their assembly resulted in a ($6{,}208 \times 4{,}669$) image. Black areas indicate locations where no tiles have been acquired.

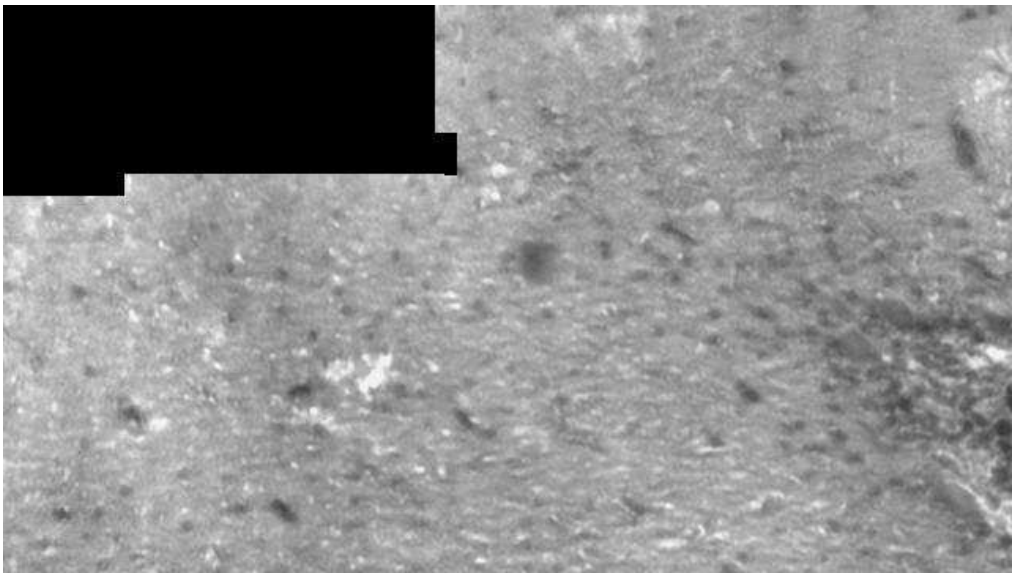

Figure 3. Full-resolution mosaic cut-out (640×360). The location of several tile borders can be inferred by the way they delineate the black area, where no data is available. This hints at the likely location of possible artifacts that would have resulted from misaligned tiles; despite this *a priori* knowledge, such artifacts are not visually discernable, which gives a qualitative testimony to the quality of our proposed method.

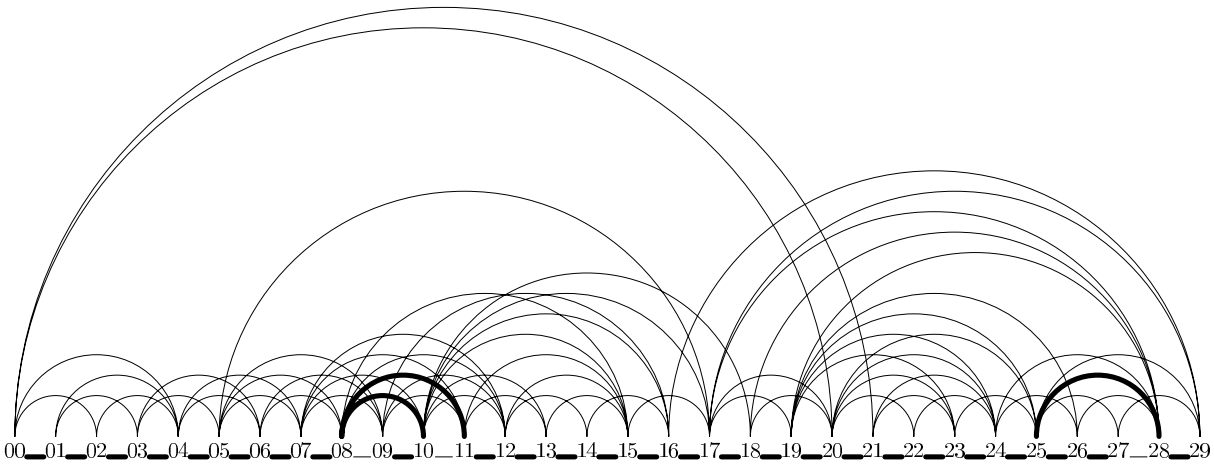

Figure 4. Graph of the overlaps. The existence of a thin line between nodes i and j indicates that the tile indexed by i overlaps the jth tile. A thick line is used to represent the maximum spanning tree, as obtained by Kruskal's algorithm on the data corresponding to Figure 2. We see that this algorithm found best to occasionally depart from the natural ordering of the tiles while building the mosaic. This is more likely to happen when the scanning path folds back on itself.

We illustrate in Figure 1 the interface seen by the user. A well at the bottom of the screen contains the tiles that have not yet been positioned. The main area is used to drag the tiles; a transparency mechanism helps the user to visualize the rough version of the mosaic. We give in Figure 2 the resulting high-quality mosaic after automatic fine-tuning. A full-scale cut-out of the central part, where several tiles meet, is presented in Figure 3; in this figure, no artifact can be seen in the prolongation of the edges of the areas not covered by any tile, which indicates that the assembly process achieved excellent visual quality. We report in Figure 4 the maximum spanning tree chosen by Kruskal's algorithm. As the scanning path follows the windy convolutions of a gyrus, a tile may overlap not only its immediate predecessor and successor, but also other tiles much farther down the path.

4. CONCLUSIONS

We have proposed a strategy to minimize the number of pairwise registration steps required for building a mosaic out of overlapping tiles in a microscopy context. In designing this strategy, we have made the reasonable assumption that the registration accuracy increases with the amount of overlap, out of which we have derived a graph-theoretic optimal solution. We have built MosaicJ, a platform-independent software that implements this strategy; we distribute this software freely. It benefits from TurboReg, a high-quality registration engine that has also been released in the public domain. We hope that microscopists will be able to put MosaicJ and TurboReg to good use.

ACKNOWLEDGMENTS

This work has been supported in part by the CIBM (Centre d'imagerie biomédicale, http://www.cibm.ch/).

APPENDIX A. KRUSKAL'S ALGORITHM

Let a weighted graph $G = \{V(G), E(G)\}$ consist of a set of vertices $V(G)$ and of a set of weighted edges $E(G)$, where each edge is assigned a weight $w_G(uv) \in \mathbb{R}$, with $u \in V(G)$, $v \in V(G)$, and $uv \in E(G)$. The number of vertices of G is $|G|$. The number of edges of G is $\varepsilon(G)$. A uv-path on G is a sequence of K distinct vertices $u_k \in V(G)$ such that $u_0 = u$, $u_{K-1} = v$, $u_k u_{k+1} \in E(G)$. The graph is said to be connected if at least one uv-path on G exists for all distinct u and v in $V(G)$. A cycle is a uv-path such that $vu \in E(G)$. A tree is a connected graph that has no cycles. A forest is a disconnected graph that has no cycles.

Given a connected and weighted graph G with $|G| > 1$, the purpose of Kruskal's algorithm is to build a tree F that satisfies $V(F) = V(G)$ and that maximizes $\sum_{uv \in E(F)} w_G(uv)$. The procedure is to start with a forest F of $|G|$ trees that initially contain a distinct vertex only, and to gradually merge and link these trees in a greedy fashion until the forest F becomes a tree with $\varepsilon(F)+1 = |F|$. This procedure can be proved to yield an optimal solution. Letting T and T_u be auxiliary trees, and letting H be a copy of G, the algorithm can be formally described by

$$
\begin{aligned}
H &\leftarrow G \\
F &\leftarrow \{V(H), \emptyset\} \\
\forall u \in V(H) : T_u &\leftarrow \{u, \emptyset\}
\end{aligned}
$$

$$
\text{while } \varepsilon(F) + 1 < |F| : \begin{cases}
xy \leftarrow \arg\max_{uv \in E(H)} w_G(uv) \\
H \leftarrow \{V(H), E(H) \setminus xy\} \\
T_x \neq T_y \;\Rightarrow \begin{cases}
T \leftarrow T_x \cup T_y \cup \{\emptyset, xy\} \\
T_x \leftarrow T \\
T_y \leftarrow T \\
F \leftarrow F \cup T.
\end{cases}
\end{cases}
$$

REFERENCES

1. J. Kruskal Jr., "On the shortest spanning subtree of a graph and the traveling salesman problem," *Proceedings of the American Mathematical Society* **7**, pp. 48–50, February 1956.
2. P. Thévenaz, U. Ruttimann, and M. Unser, "A pyramid approach to subpixel registration based on intensity," *IEEE Transactions on Image Processing* **7**, pp. 27–41, January 1998.
3. W. Rasband, "ImageJ." U.S. National Institutes of Health, Bethesda MD, USA, 1997-2005. http://rsb.info.nih.gov/ij/.