Visualization in Medicine and Life Sciences (2013)
L. Linsen, H. -C. Hege, and B. Hamann (Editors)

Short Papers

Hierarchical Poisson-Disk Sampling for Fiber Stipples

Mario Hlawitschka[†1], Mathias Goldau[†1,2], Alexander Wiebel[‡3,4], Christian Heine[§5], and Gerik Scheuermann[†1]

[1]Leipzig University, Germany
[2]Max-Planck-Institute Neurological Research Cologne, Germany
[3]Zuse Institute Berlin (ZIB), Germany
[4]Coburg University of Applied Sciences, Germany
[5]ETH Zurich, Switzerland

Abstract

To understand neural tracts of the brain, neuroscientists use visualizations of diffusion data. Fiber stippling – a technique inspired by illustrations – accommodates probabilistic tracts, main diffusion direction, and anatomical context in the same slice image. It uses stratified sampling to place stipples, but this can result in overlaps and undersampled areas that distort the perception of tract probability. Moreover, when changing slices in an interactive setting, resampling leads to visual noise that distracts from real changes in the data. In this paper, we propose to use Poisson-disk samplings to ensure adequate pattern perception inside slices and a hierarchy of samplings to ensure coherence among slices. We also port the algorithm to the GPU to achieve interactive frame rates. Our modifications are appreciated by neuroscientists, who can now investigate white-matter structures more confidently.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Line and curve generation

1 Introduction and Background

Neuroscientists study the neural “wiring” scheme of the human brain to better understand its functions. To probe for structural connectivity [Spo07], diffusion-weighted magnetic resonance imaging (dMRI) – measuring the diffusion strength of water inside the brain – can be used. It allows to reconstruct white-matter fiber bundles (e.g. [JJB11, FDG*11]) by following diffusion directions. One technique, probabilistic tractography, estimates structural connectivity by computing a connectivity score (CS), which measures the likelihood for each voxel to be connected to a given seed region in the brain. This score can be modeled by a scalar field, but its depiction in anatomical context is difficult.

Goldau et al. [GWG*11] propose to use a technique called *fiber stippling* for rendering probabilisitc tracts, motivated by illustrations in the book *Fiber pathways of the brain* [SP06]. Goldau et al.'s method generates a set of small glyphs, reminiscent of line stipples known from computer graphics, such that the perceived stipple density and color signifies connectivity scores, and their orientation the main local diffusion direction. Anatomical context is provided via outlining the gray- and white-matter boundaries as well as underlays.

According to our collaborators, neuroscience experts of the Max-Planck-Institute for Neurological Research in Cologne, Germany, as well as neuroscience experts from the Max-Planck-Institute for Human Cognitive and Brain Sciences in Leipzig, Germany, fiber stippling is very helpful due to the following reasons: Medical staff and neuroscientists are familiar with and prefer medical data on slices, as they alleviate occlusions. Anatomical context, even plain gray-matter and white-matter boundaries, is crucial when evaluating tractography data. Fiber stipples' visual sparseness enables their use as overlays to rich anatomical context, e.g. T1 images, and supports comparing multiple probabilistic tracts in the same image. Stipples are presented at sub-

† {math,hlawit,scheuer}@informatik.uni-leipzig.de
‡ alexander.wiebel@hs-coburg.de
§ cheine@inf.ethz.ch

DOI: 10.2312/PE.VMLS.VMLS2013.019-023

voxel precision, making it possible to track structures near the resolution limit.

The same experts also see some drawbacks of Goldau et al. [GWG*11]'s technique. First, stipple placement using stratified sampling introduces random perturbations in the perception of connectivity scores (CS). Stipple overlaps interfere with the opacity of stipples and visual "holes" are created in regions with low CS, see Figure 2c. This complicates following white matter tracts at the very limit of image resolution and in regions of low CS. Second, when moving slices through the volume, *frame-incoherence* can be observed: the statified sampling places stipples such that they appear to jump rather than change smoothly. Finally, stipple reorganization and drawing take between one and two seconds. The latter two points complicate following tracts when moving the slice interactively. Hence we propose the following improvements:

- replace stratified sampling by Poisson-disk sampling for better perception of connectivity scores,
- use a hierarchy of Poisson-disk samplings to ensure frame coherence, and
- port the new algorithm to the GPU to achieve frame rates suitable for interaction.

2 Related Work

For the visualization of probabilistic tracts often only the CS is visualized. Therefore, three-dimensional techniques [DDA07, SEI10, vKRC*10] as well as two dimensional techniques [MN07, SSA*08, ASH*09, SHH*10] are available. Generally, all three-dimensional techniques suffer from imprecise anatomical context – if it is present at all. For two-dimensional techniques, anatomical context is possible as underlays but often occluded by the tract representation. Besides the visualization of the CS, there are also techniques just for the diffusion pattern [CTJC10, BH04, Kin04, HES08].

Fiber stippling [GWG*11] combines CS with diffusion directions while still supporting anatomical context. Therefore, a scalar field (CS), a vector field (major diffusion orientations), and a scalar field (anatomical context) are needed. Given an arbitrary slice of the brain, the stipples are displayed only in regions of significant CS. The stipple orientation is obtained from the vector field where as the opacity and stipple density reflects the CS. For stipple placement, stratified sampling is used. This means that the slice is partitioned into fixed buckets, each associated with a CS and hosting a number of stipples. The bucket's CS decides whether any stipples are rendered or not. The stipple placement inside each bucket is random and the number of stipples depends on the desired density, where each stipple queries its own orientation and CS.

The placement of fiber stipples is very similar to the placement of glyphs [KW06, ROP11].The perfectly dense placement of fiber stipples at each slice would take the stipple's orientation into account, which conflicts with the frame coherence or would have to consider animations for the relocation of stipples (cf., e.g., the interactive glyph placement strategies presented in [HSH07]). More recent approaches use the metric tensor to display similar dense patterns [ASKH12]. We focus on improving the sparse and locally changing placement of stipples.

3 Methods

Before describing our improvements for the fiber stipples, we should state first, that we do not enforce any additional restrictions of the original technique and support all previously known features which include:

- Visual encoding of diffusion orientation as stipples
- Visual encoding of CS in density and opacity
- Parameters for scaling of stipple size and density
- Anatomical context as boundary lines and underlays
- Arbitrary slice resolution and position

3.1 Poisson-Disks for stipple placement

In order to reduce misinterpretation of fiber stipples, the most important improvement is the stipple placement with regard to frame coherence and performance improvements. Originally, fiber stipples were placed with stratified sampling. That means that they partition the slice into fixed buckets, each associated with a connectivity score, and then draw and place a number of samples randomly depending on the desired density, where each sample queries its own orientation and CS.

To overcome unwanted overlaps of stipples we are allocating a certain fixed radius for each stipple. This reduces the uniformity of the spatial coverage but allows for a simplified calculation of a nearly optimal sampling. Poisson-disk sampling guarantees a uniform stipple distribution with a minimum distance between all pairs of stipple centers. For a straightforward Poisson-disk sampling, the additional need for non-uniform sampling and the change of resolution is difficult. Multiresolution Poisson-disk samplings [GM09] can be used, but they have shown to be a time-consuming task and lack the possibility to introduce frame coherence.

To further be able to choose the density of the visualization, we chose a multilevel approach in which we prepare N sets of points where the union of the first n sets has Poisson-disk characteristics.

These sets are computed using the following criteria: First, the set of all points is a set of Poisson-disk distributed points in the plane with a minimal distance

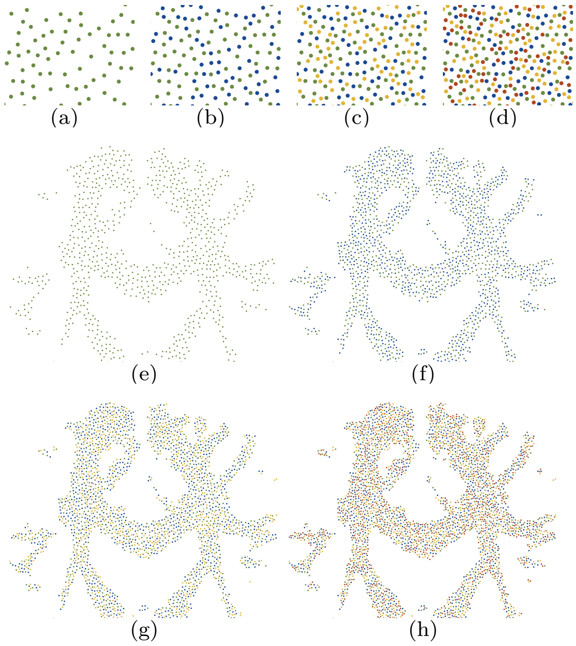

Figure 1: *Internal patterns for the point seeding strategy. (a)-(d): the combination of the first 1 to 4 point sets to obtain point sets with gradually increasing density and gradually decreasing point distance. (e)-(h): The point set as it is filtered in the vertex shader to obtain a representation of a fiber tract. In all images, the seed points are represented by dots obtaining their color from the point set they belong.*

r that guarantees a minimal overlap of glyphs. Second, the points in each of the N sets are chosen from the set of initial points randomly so that all points of the first n sets are spaced with a minimal radius r_N as shown in Figure 1.

To adapt the density of stipples to the connectivity score (CS), the first k point sets are used in the visualization with k being proportional to the density and the points within the slice are chosen using a von-Neumann rejection sampling [Sob06] based on the CS of the probabilistic tract.

Even though the new stipple placement reduces overlapping, it is still possible to “violate” the non-overlapping constraint by changing the parameters to obtain different styles for the visualization (see Figure 3b).

3.2 Establishing frame coherence

To establish frame coherence, we have to ensure that the set of seed points in neighboring voxels are alike. This can be done by, first, using the same point sets and, second, reducing the stochastic component of the von-Neumann approach by a fixed “random value” computed from the stipple index. For a given CS c, we use all points of the first $n = \lfloor cN \rfloor$ point sets and the first $k = \lceil (cN - n)|S_{n+1}| \rceil$ points of the $|S_{n+1}|$ points of the next set.

In this stage, the stochastic selection of seed points became a purely deterministic process that guarantees frame coherence in the visualization. The acceptance or rejection can be determined for each point individually without the generation of random numbers.

When moving the visualization to neighboring slices, the same initial point set is used. New glyphs will appear or glyphs will vanish based on the change in the CS, but the overall layout does not change.

3.3 Porting fiber stippling to the GPU

To improve the overall performance, we implemented fiber stipples as well as the anatomical context (isolines) generation completely on the GPU.

For the fiber stipples, we store all multilevel point sets for the seed points on the GPU. In case the sample will be discarded, a degenerate quad is produced in the vertex shader of the GPU and, therefore, will be ignored in the following steps. Otherwise, a quadrilateral is produced that contains the circumference of the stipple. Afterwards, the two endpoints (or focal points) of the stipple's rectangle are passed to the fragment shader which renders the stipple with a rectangle and two disks in the given color.

The implementation of the anatomical context, instead, has proven to be more complex. Our initial idea to implement the isoline algorithm based 2D textures and the display based on the fragment's values did not work for this data due to the limited numerical precision when low gradients were present in the texture. Eventually, we used a marching-quad [GWG*11] algorithm on a grid of arbitrary resolution. For each grid cell, represented as quadrilateral, we computed positions of the isovalue on the edges. The very same way as the vertex shader filters quadrilaterals based on the CS when rendering the stipples, the vertex shader also filters here quadrilaterals when no edge of a grid cell contains the isovalue. Implementation detail to note here: we used overlapping grid cells to circumvent some line-connection artifacts when employing thicker isolines.

4 Results

Figure 2 shows the advantages of our new seeding strategy compared to the previous approach. With the new placement strategy, stipples are more regularly distributed and visual artifacts such as overlapping

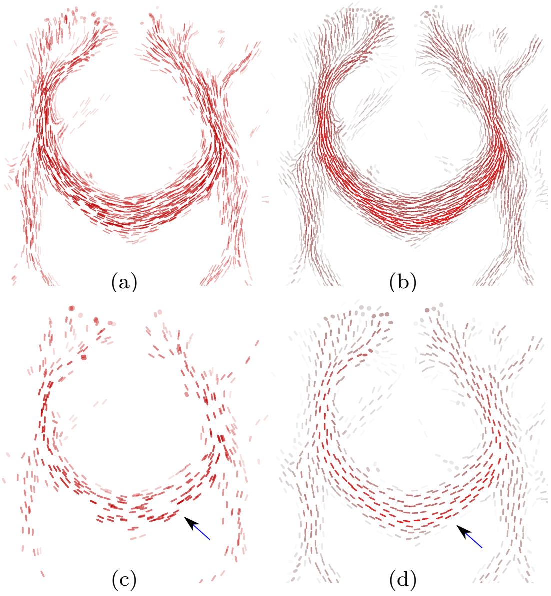

Figure 2: *Comparison of the new seeding strategy with stratified sampling: (a),(c) original stratified sampling and (b),(d) new multilevel Poisson-disk based sampling with approximately the same number of samples. (a),(b) show higher density while (c),(d) show lower density.*

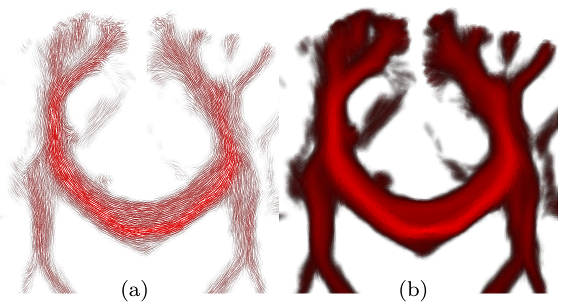

Figure 3: *Variation of different parameters. (a) long but thin stipples, and (b) bold stipples, with explicit strong overlap.*

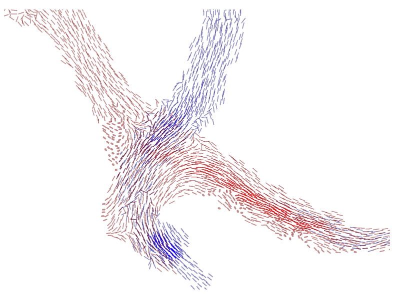

Figure 4: *Two crossing probabilistic tracts, Cortico Spinal Tract (blue) and Corpus Callosum (red), displayed at the same time.*

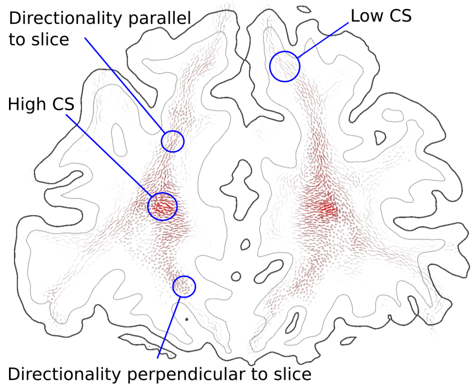

Figure 5: *Improved fiber stipples with boundary lines rendering a probabilistic tract in coronal view. Similar images of the previous technique can be found in Goldau et al. [GWG*11], Figs 11(c,d).*

stipples or holes are reduced. See, for example, the regions marked by arrows in Figure 2c and Figure 2d.

A depiction including anatomical context can be found in Figure 5. The GPU implementation on current hardware (CPU: Intel E5620, GPU: Nvidia GTX 560 - 1GB) renders this image at 300fps@2560x1600, which is interactive performance. In contrast, the original fiber stipples need around 1fps instead.

5 Discussion and Conclusion

By the techniques introduced in this paper, we are able to significantly reduce visual artifacts in the rendering which carried the potential for misinterpretations in the original fiber stippling method. The interactive experience of fiber stippling has been improved by ensuring frame coherent display, and the new implementation allows for a wider range of interactive parameter changes. The latter gives the user the opportunity to strengthen the visibility of the shown features (Figure 3). In comparison to the original work, we achieve better interactive performance while still supporting all previous features. The improved sampling also helps perception of crossing tracts, because by using the same sampling, stipple centers align (Figure 4). Preliminary evaluation by neuroscientists indicate that they can now investigate white-matter structures with much higher confidence.

Acknowledgements

The authors would like to thank the OpenWalnut project [EGRW13] for providing a very flexible, powerful and modern software framework supporting neuroscience visualization.

References

[ASH*09] Anwander A., Schurade R., Hlawitschka M., Scheuermann G., Anderson T., Knösche T. R.: White matter imaging with virtual Klinger dissection. *NeuroImage 47*, Supplement 1 (2009), S105 – S105. Organization for Human Brain Mapping 2009 Annual Meeting. 2

[ASKH12] Auer C., Stripf C., Kratz A., Hotz I.: Glyph- and texture-based visualization of segmented tensor fields. In *International Conference on Information Visualization Theory and Applications (IVAPP'12)* (2012). 2

[BH04] Benger W., Hege H.-C.: Tensor splats. In *Electronic Imaging 2004* (2004), International Society for Optics and Photonics, pp. 151–162. 2

[CTJC10] Calamante F., Tournier J.-D., Jackson G. D., Connelly A.: Track-density imaging (TDI): Super-resolution white matter imaging using whole-brain track-density mapping. *NeuroImage 53*, 4 (2010), 1233 – 1243. 2

[DDA07] Descoteaux M., Deriche R., Anwander A.: Deterministic and probabilistic Q-ball tractography: From diffusion to sharp fiber distribution. 2

[EGRW13] Eichelbaum S., Goldau M., Reichenbach A., Wiebel A.: OpenWalnut. http://www.openwalnut.org, 2009–2013. 4

[FDG*11] Fillard P., Descoteaux M., Goh A., Gouttard S., Jeurissen B., Malcolm J., Ramirez-Manzanares A., Reisert M., Sakaie K., Tensaouti F., Yo T., Mangin J.-F., Poupon C.: Quantitative evaluation of 10 tractography algorithms on a realistic diffusion MR phantom. *NeuroImage 56*, 1 (2011), 220 – 234. 1

[GM09] Gamito M. N., Maddock S. C.: Accurate multidimensional Poisson-disk sampling. *ACM Trans. Graph. 29*, 1 (2009), 8:1–8:19. 2

[GWG*11] Goldau M., Wiebel A., Gorbach N., Melzer C., Hlawitschka M., Scheuermann G., Tittgemeyer M.: Fiber stippling: An illustrative rendering for probabilistic diffusion tractography. In *Biological Data Visualization (BioVis), 2011 IEEE Symposium on* (2011), pp. 23 –30. 1, 2, 3, 4

[HES08] Hlawitschka M., Eichelbaum S., Scheuermann G.: Fast and memory efficient GPU-based rendering of tensor data. In *Proceedings of the IADIS International Conference on Computer Graphics and Visualization 2008* (2008), pp. 36–42. 2

[HSH07] Hlawitschka M., Scheuermann G., Hamann B.: Interactive glyph placement for tensor fields. In *Advances in Visual Computing: Third International Symposium, ISVC, Lake Tahoe, Nevada/California (Lecture Notes in Computer Science)* (2007), vol. LNCS 4841 and LNCS 4842, pp. 331–340. 2

[JJB11] Jbabdi S., Johansen-Berg H.: Tractography: Where do we go from here? *Brain Connectivity 1*, 3 (2011), 169–183. 1

[Kin04] Kindlmann G.: Superquadric tensor glyphs. In *Proceedings of IEEE TVCG/EG Symposium on Visualization 2004* (04 2004), pp. 147–154. 2

[KW06] Kindlmann G., Westin C.-F.: Diffusion tensor visualization with glyph packing. *Visualization and Computer Graphics, IEEE Transactions on 12*, 5 (2006), 1329–1336. 2

[MN07] McGraw T., Nadar M.: Stochastic dt-mri connectivity mapping on the gpu. *IEEE TRANSACTIONS ON VISUALIZATION AND COMPUTER GRAPHICS* (2007), 1504–1511. 2

[ROP11] Ropinski T., Oeltze S., Preim B.: Survey of glyph-based visualization techniques for spatial multivariate medical data. *Computers & Graphics 35*, 2 (2011), 392–401. 2

[SEI10] Svetachov P., Everts M. H., Isenberg T.: DTI in context: Illustrating brain fiber tracts in situ. *Computer Graphics Forum 29*, 3 (2010), 1024–1032. 2

[SHH*10] Schurade R., Hlawitschka M., Hamann B., Scheuermann G., Knösche T. R., Anwander A.: Visualizing white matter fiber tracts with optimally fitted curved dissection surfaces. In *VCBM 2010 Eurographics Workshop on Visual Computing for Biology and Medicine* (2010), Bartz D., Botha C. P., Hornegger J., Machiraju R., Wiebel A., Preim B., (Eds.), Eurographics Association, pp. 41–48. 2

[Sob06] Sobol I. M.: *Die Monte-Carlo Methode.* Deutscher Verlag der Wissenschaften, 2006. 3

[SP06] Schmahmann J. D., Pandya D. N.: *Fiber Pathways of the Brain.* Oxford University Press, 2006. 1

[Spo07] Sporns O.: Brain connectivity. *Scholarpedia 2*, 10 (2007), 4695. 1

[SSA*08] Schultz T., Sauber N., Anwander A., Theisel H., Seidel H.-P.: Virtual klingler dissection: Putting fibers into context. In *Computer Graphics Forum* (2008), vol. 27, pp. 1063–1070. 2

[vKRC*10] von Kapri A., Rick T., Caspers S., Eickhoff S. B., Zilles K., Kuhlen T.: Evaluating a visualization of uncertainty in probabilistic tractography. In *Proc. SPIE* (2010), Wong K. H., Miga M. I., (Eds.), vol. 7625, SPIE, p. 762534. 2