



Mathematical morphology

Originally developed for binary images (like masks), morphological filters work on shapes. For example, they can remove small features, bridge gaps, or even remove all elongated shapes. Typically the “shapes” are taken to be connected components of level sets, but other schemes are definitely possible.

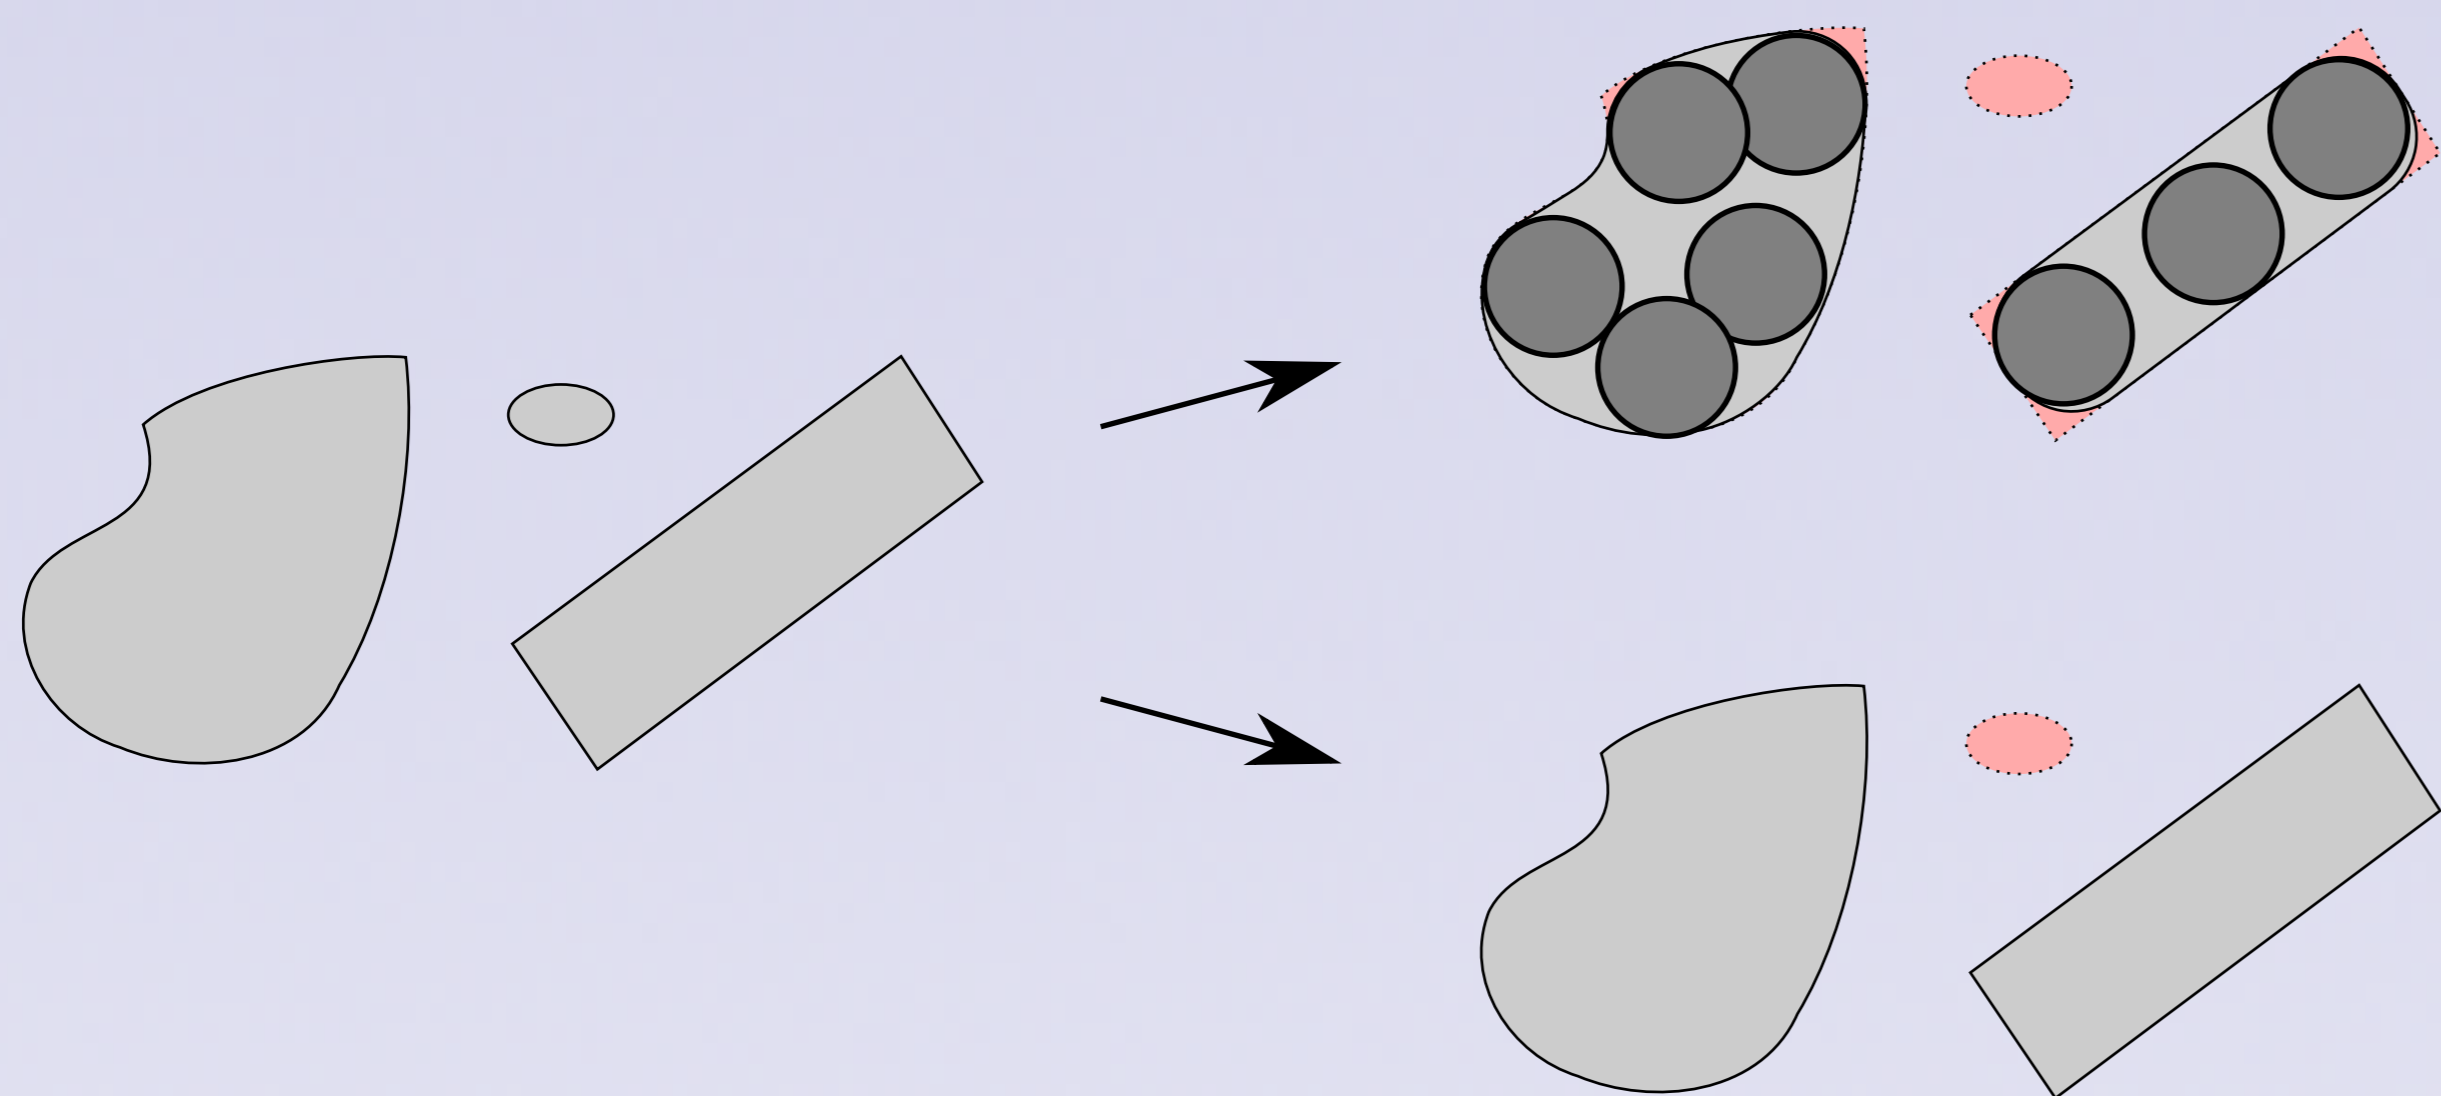


Figure 2: The shapes on the left are filtered using two morphological filters. Top: a morphological “opening” by a disc only keeps those parts of the shapes where the disc fits. Bottom: a “connected” filter does not alter the shapes, it either keeps them or removes them in their entirety.

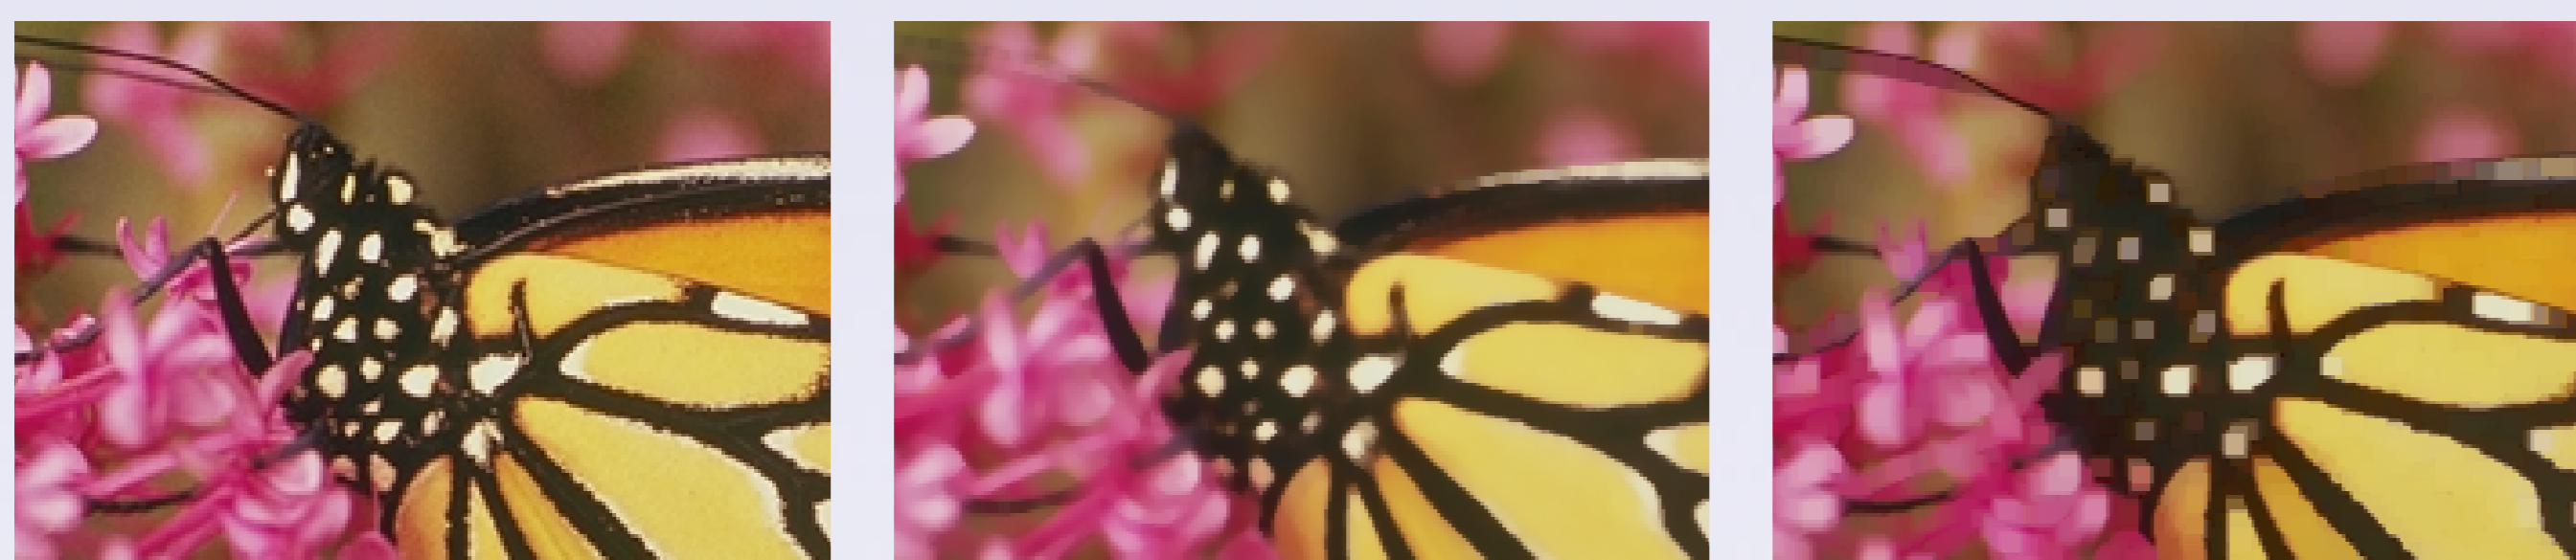


Figure 3: From left to right: the original image, a median filtered version, and the result of a morphological opening. Both filters are considered morphological operators.

Tensors

Vectors and matrices are the most common examples of tensors. Important examples where these show up are colour images and so-called “Diffusion Tensor Imaging”. The latter technique uses MRI scans to determine how water diffuses at each position in the brain (or other tissue). This information is encoded in a matrix (a tensor) that essentially gives an estimate of the ease with which water diffuses in each direction. Since water cannot (easily) diffuse perpendicular to nerve fibres, this data can be used to explore the structure of the brain in-vivo.

Diffusion tensor tensor images are often noisy and/or unclear. Post-processing is therefore a must, and morphological filters seem an appealing option. Also, this would enable many of the image analysis tools that rely on morphological filters to carry over to the tensorial case.

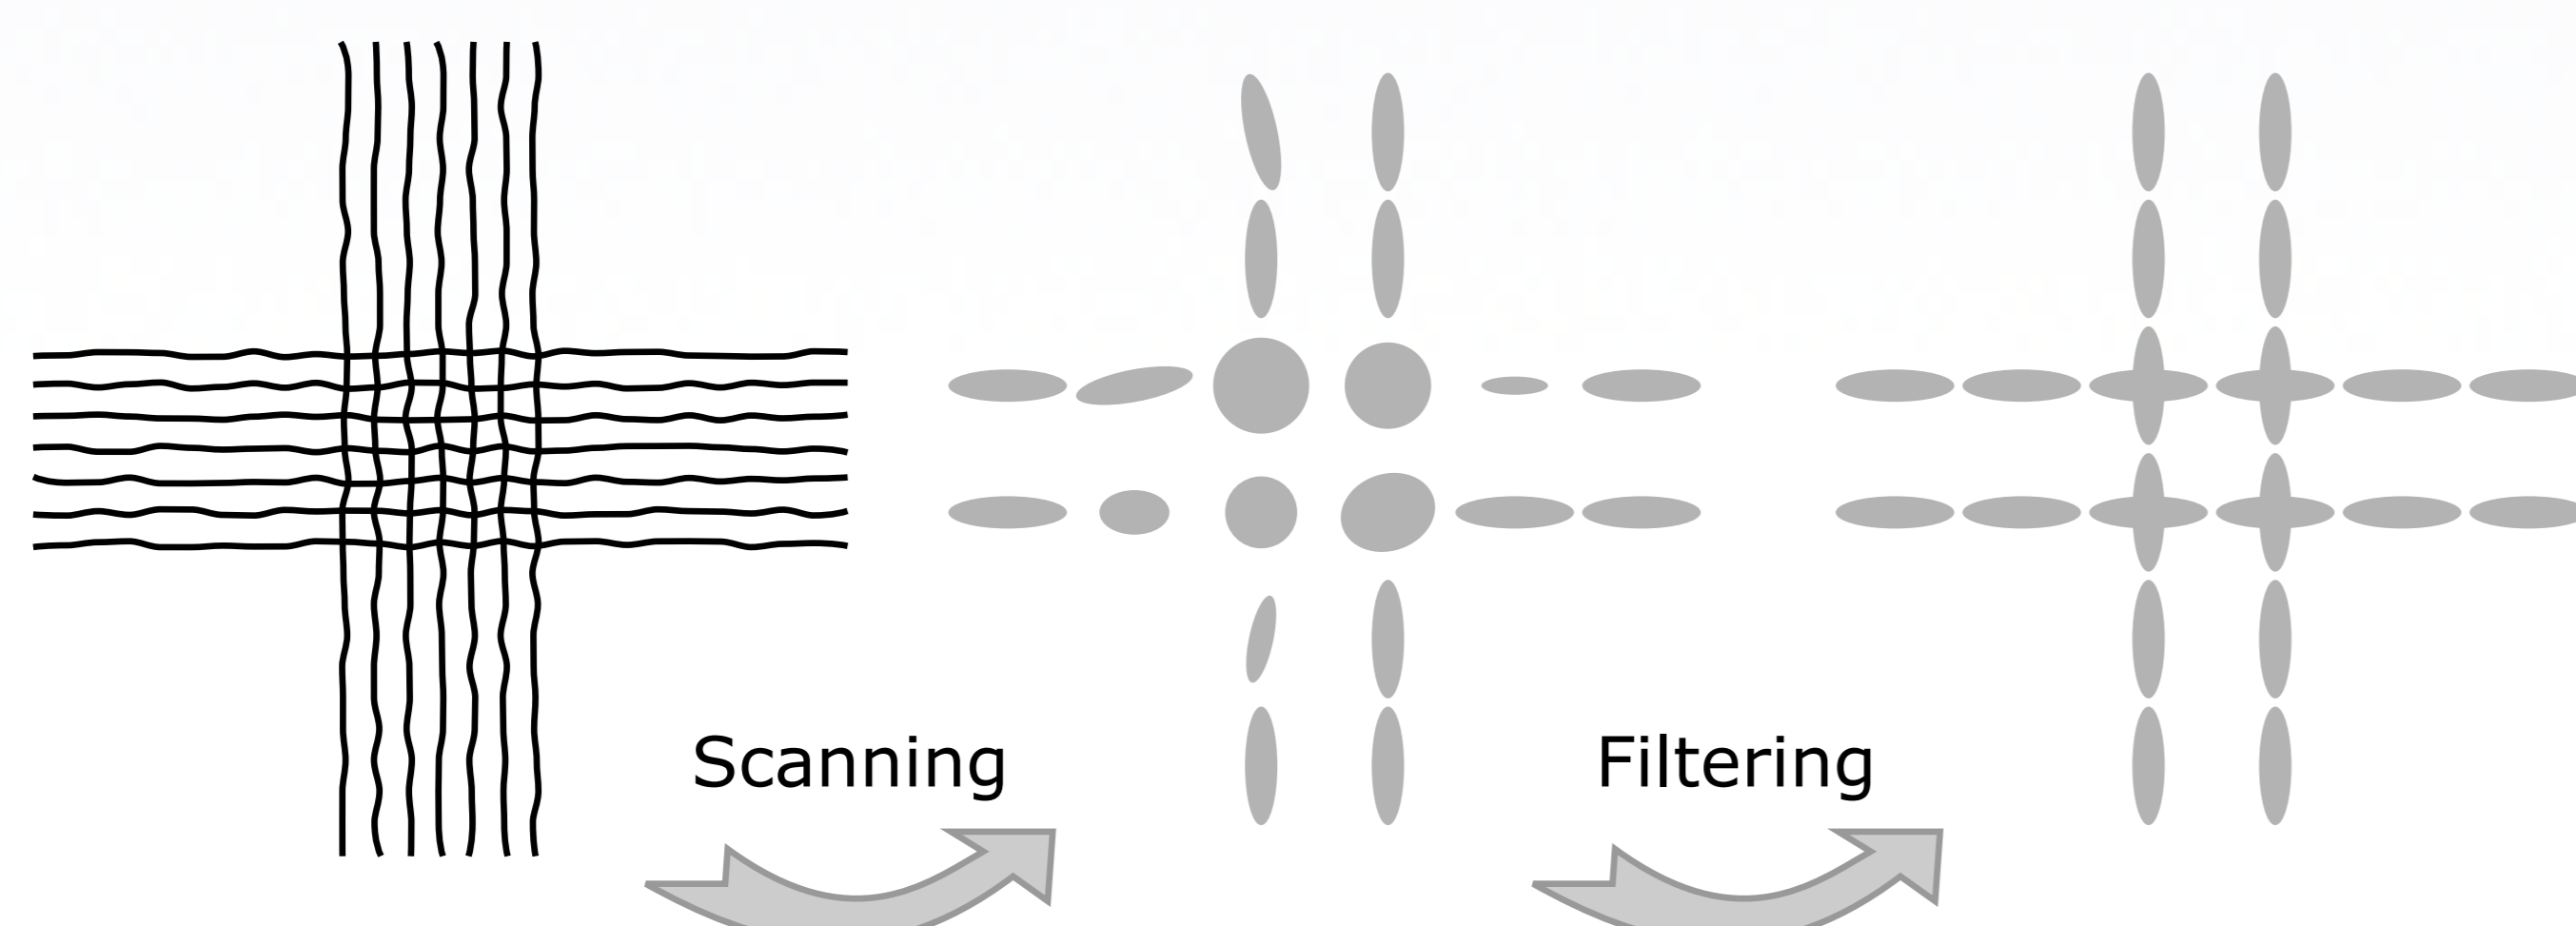


Figure 4: Raw scans of nerve fibre bundles are often noisy, and contain areas where the fibre orientations are not directly apparent. Morphological filters should be able to help clean up these images, so further processing becomes much easier.

Invariance

Nobody wants a filter to work on a particular image, only to fail on a shifted copy. Similarly, a change in hue in the input should often just result in a change in hue in the output. And when scanning fibre bundles, filters should work regardless of the orientation of the fibres.

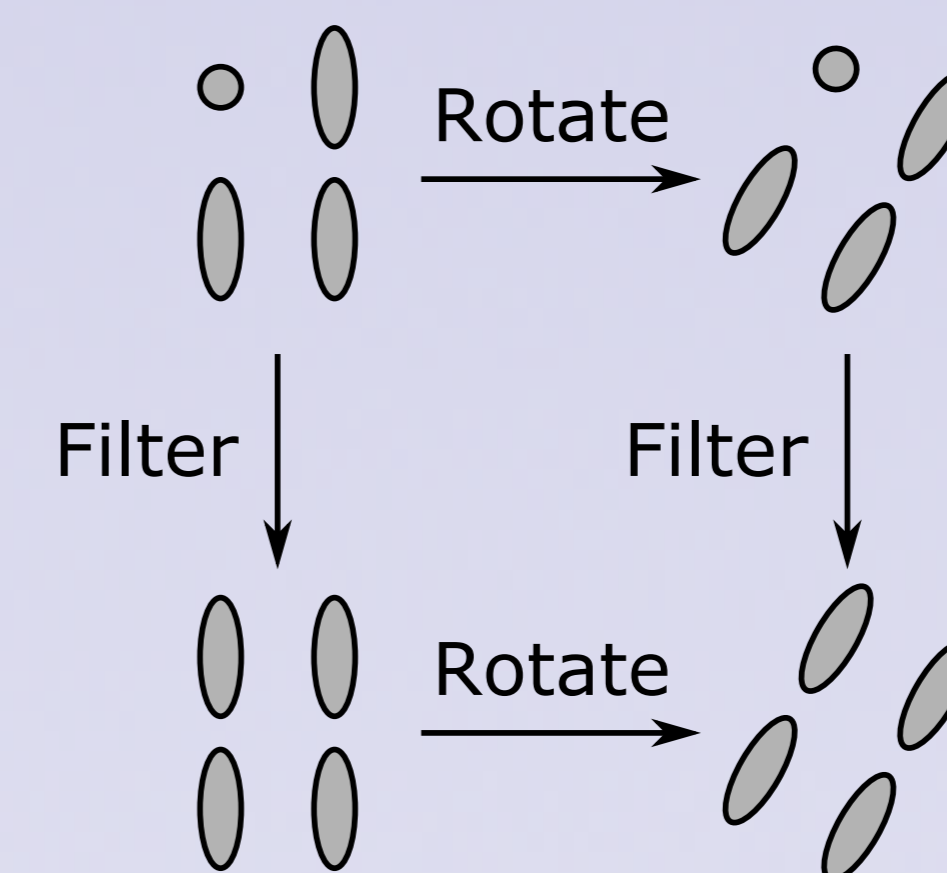


Figure 5: First rotating and then filtering a diffusion tensor scan should have the same result as first filtering and then rotating.

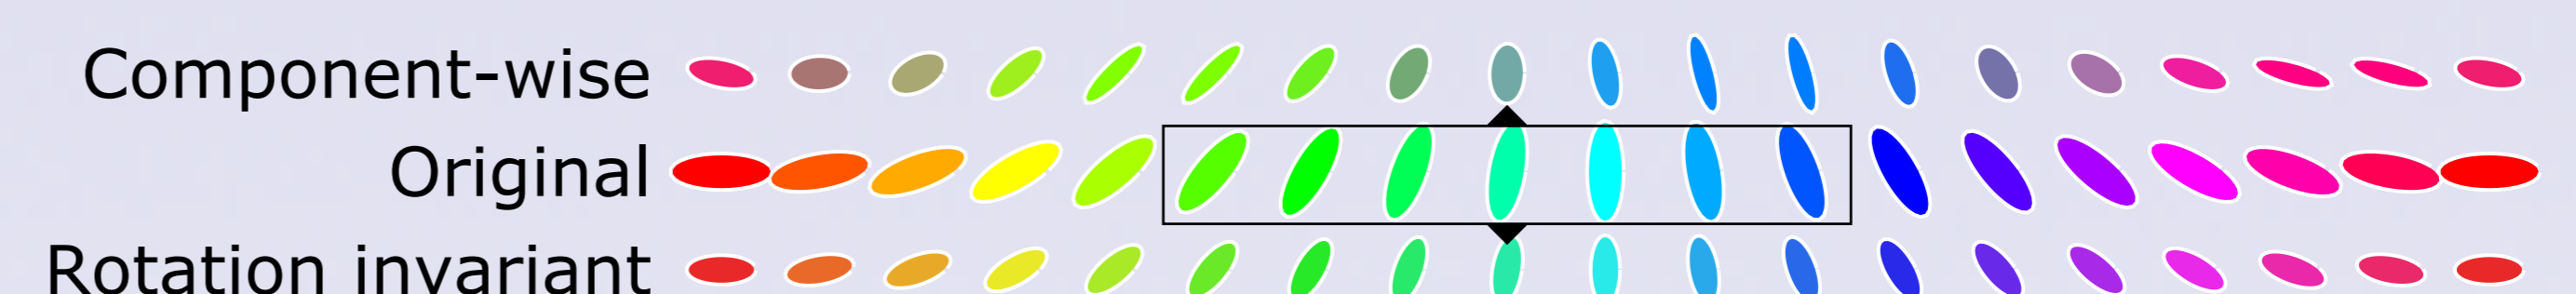


Figure 6: Straightforward component-wise and rotation invariant “minima” over different, equally wide, ranges of tensors. Hue: direction of principal eigenvector. Saturation: ratio between eigenvalues.

Future work

The goal is to develop efficient and useful tools based on morphological filters for tensor images. This will involve local filters (for noise reduction, detection of crossings, and so on), as well as filters that handle entire shapes at a time. Also, morphological techniques could be used to analyse tensor images.

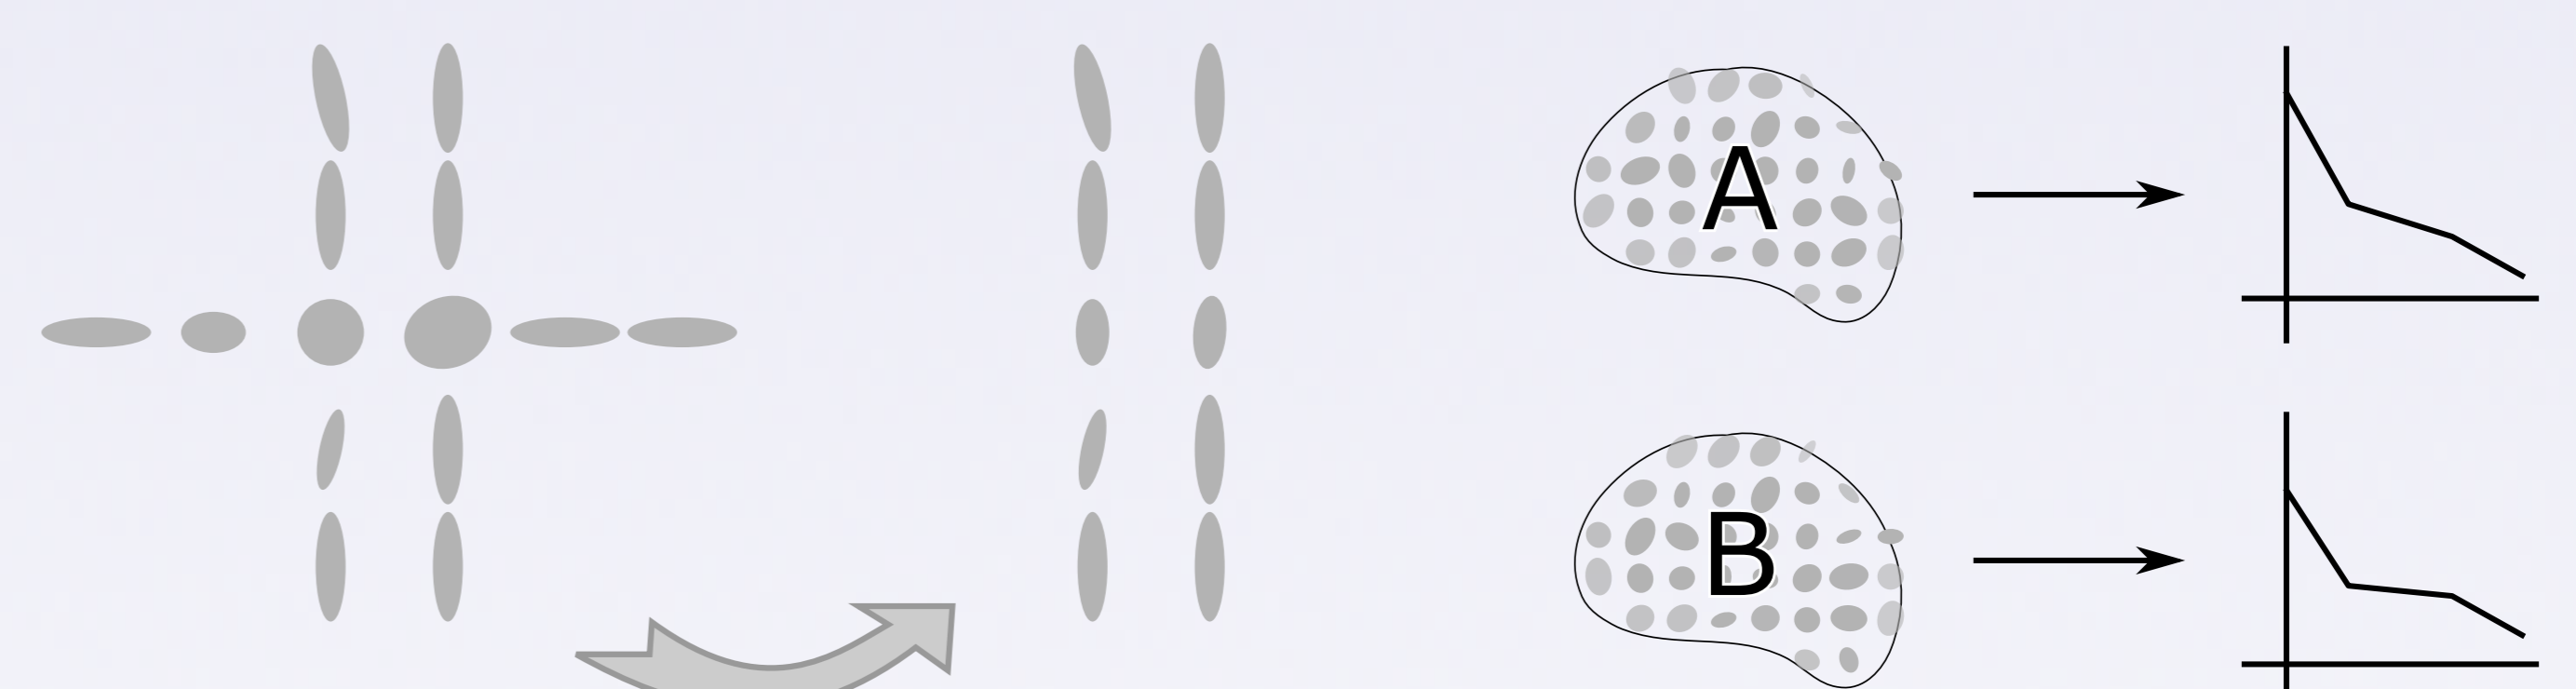


Figure 7: (Left) So-called “connected” filters could be used to essentially select only those fiber bundles that satisfy criteria like thickness, length, etc. (Right) Morphological analysis of tensor images could provide a way to quantitatively compare different images, enabling classification for example.

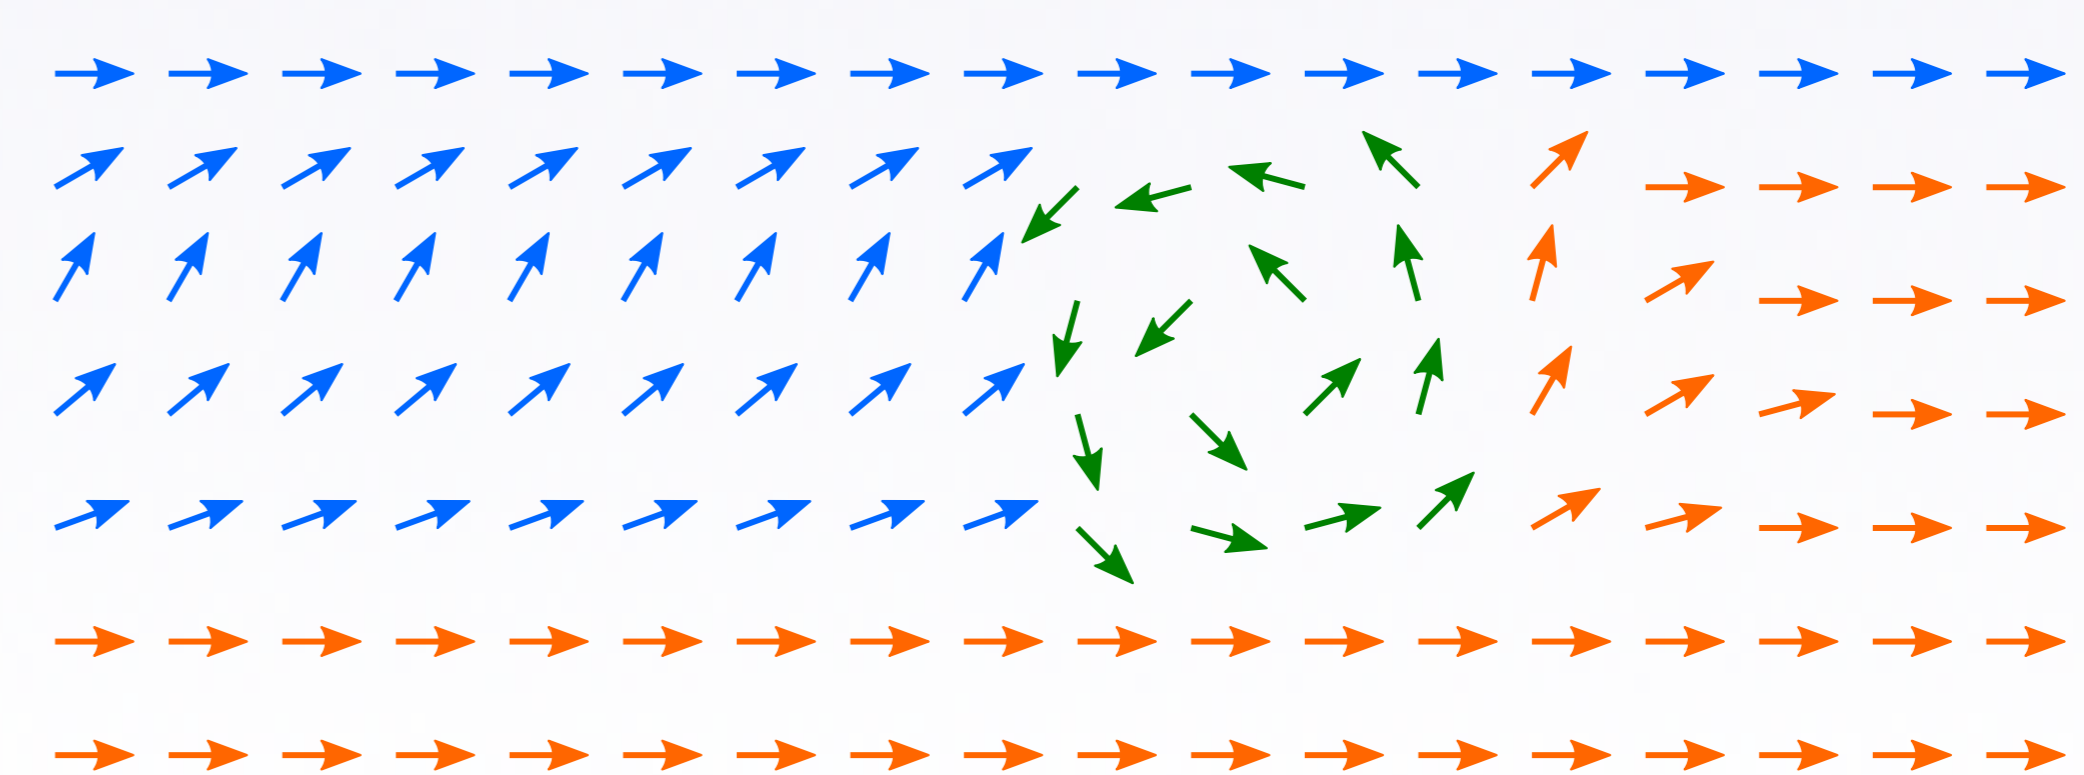


Figure 8: Flow fields are another potential application domain. It might be interesting to provide (hierarchical) segmentations of flow fields, for example. Those could then also be combined with pattern matching to find certain features in a flow.

Contact

Jasper van de Gronde

Johann Bernoulli Institute for Mathematics and Computer Science
University of Groningen
Nijenborgh 9
9747 AG Groningen

e-mail: j.j.van.de.gronde@rug.nl url: <http://www.cs.rug.nl/svcg/>