Shuffle Distance and Symmetry
Some Ramblings

R. S. Schestowitz*

Research Student

Imaging Science and Biomedical Engineering

Stopford Building

University of Manchester

United Kingdom

April 28th, 2005

- We currently evaluate shuffle distance from shuffle distance images that work in a single direction.
- The direction is dependent on the implementation and is arbitrary.
- We might wish to use information from a shuffle distance image that goes in both direction.
- Averaging of the resulting two images seems like the most basic idea.
- The aggregated shuffle distance image seems to be affected slightly.
- In the case of the brain:
  - Skull misalignment is highlighted equally well for both original images.
  - More brighter shades appear, but their intensity is lower because of the averaging.

*Contact: `r@schestowitz.com`

- Evaluation and comparison between symmetrical and asymmetrical measures is non-trivial.
- One can evaluate the performance from the distance (average pixel intensity) or from the raw images, which are too large to handle *en masse*.
- By taking an image set and slicing it into two groups, a large number of pairings becomes available.
- Evaluation can also be made by using the existing task of model evaluation. However, correct solutions are scarcely known.
- If ground truth was available, then performance of a symmetric measure and an asymmetric one would be comparable.
- Can an artificial and quick test be conducted? Perhaps with synthetic data?
- One problem with measuring performance is that the nature of shuffle is quite 'organic'. It does not compute anything that is an inherent characteristic the data.
- One possibility is to plot results of some kind for shuffle in either direction and in both. By looking at the 3 curves, conclusions can be drawn. But the question then becomes: "What results should be plotted?".
- Mean intensity of the images is expected to have the following relationship: $(mean(image1)+mean(image2))/2 = mean(shuffle(image1+image2))$
- Since the above holds, the question then becomes: "How do the different intensity values spread within the shuffle distance images?". Also, it would interesting to decide on where high intensities are most helpful. For example, does one want high intensities around the skull or in the centre of the brain? Many parameters control the behaviour for a given set of data and, quite clearly, there is an element of art in parameter selection.
- A more correct way of handling distances and using the shuffle transform is by computing it in both direction. It is potentially twice as expensive in terms of resources, but in practice, possible speed-ups exist.
- Averaging of the two shuffle distance images is *weak*. A more cunning approach will use the correlation between the two images to produce meaningful and valuable information.

- For instance, the inability to ’fit’ pixels (a discrepancy) in both directions, implies that the penalty should perhaps be raised.
- Salt-and-pepper noise might in some cases work in favour of one side, but not in favour of the other. In a sense, shuffle distance can detect inconsistency in the data as it detects unexpected pixels without a local region and assumes a pair of images should be similar.

Ways to Proceed

When results that are somehow bound to ground-truth become available, different shuffle distance approaches can be compared. The aim is to fit a shuffle-based method to a somewhat ’correct’ solution so that it emulates more reliable methods.