Video Textures Using the Auto-Regressive Process

Neill Campbell *, Colin Dalton, David Gibson and Barry Thomas
University of Bristol, UK

1 Introduction

Recently, there have been attempts at creating 'video textures', that is, synthesising new video clips based on existing ones. Schodl et al. showed new video clips by carefully choosing sub-loops of an original video sequence that could be replayed.

We present a different approach to recreating (potentially infinitely long) new sequences. We transform each frame of the video into an eigenspace using Principal Components Analysis (PCA) so that the original sequence can be viewed simply as a signature through this low-dimensional space (see Gibson et al. for another example of using PCA to assist animators). A new sequence can be generated by moving through this space and creating 'similar' signatures. This similarity is derived using the auto-regressive process (ARP) as discussed by Blake and Isard. A 2nd-order process is used and provides a statistical framework for assessing the quality of the new sequence. The new signature has the 2nd-order properties of the original and is much more than simply a random-walk through the space. New sequences created with our approach can contain images never present in the original sequence and are very convincing.

2 Results

Our first example is that of a campfire sequence. A plot of the original signature is shown below in red, with the synthesised signature superimposed in blue. The sequence is reconstructed using eight eigenvectors (using fewer may lead to blurring) and well characterises the flickering of the fire.

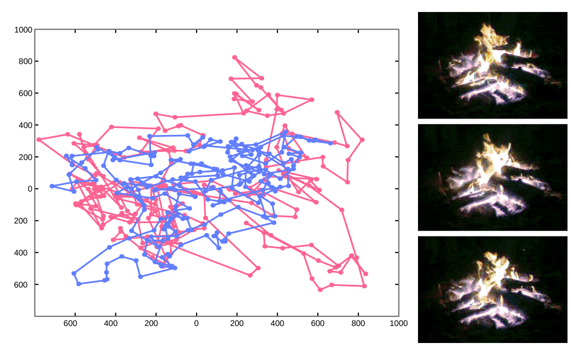

Figure 1: (Left) The original campfire sequence (red) vs the newly synthesised sequence (blue). For clarity, only the 2 dominant eigenvectors are shown. (Right) three of the resulting frames from the new sequence.

The ARP assumes that the signature is a Gaussian process, but sometimes image sequences have highly non-gaussian characteristics. An example of this are candle flames, where, due to their bending, correlations between images are not well explained using a linear analysis technique such as PCA. To overcome this we use the Combined Appearance Model of Cootes and Taylor to decompose the appearance into a shape description (a polymesh) and a texture map (warped into a neutral pose). The combined PCA space is treated by us in the same way as the image eigenspace used for the campfire sequence, but exhibits the Gaussian nature required by the ARP. Shown below is a synthesised sequence of a man laughing created using such an appearance model. Once again the sequence captures the feeling of the original sequence well with the man smiling, laughing, shaking his head and rocking back and forth. The minor distortions present are due to the image warping technique used and we will overcome these. A second non-gaussian example fits a spline through the data space a synthesises a horse running on a treadmill using the distance along the spline, and an offset from the spline.

Figure 2: Four frames of laughing man. All of the images shown here are new and were not present in the original sequence.

3 Conclusion

We have presented a novel approach to creating video textures of arbitrary length. Such technology has a wide applicability, including the automatic generation of computer backgrounds and crowd scenes. This image-based rendering technique has been demonstrated for video sequences and also combined appearance models. The signatures on which the resulting images are based have the same properties as the original signatures and the sequences capture the look-and-feel of the original ones.

*e-mail: Neill.Campbell@bris.ac.uk