Fast and Accurate Active Appearance Models

Iain Matthews

The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213, USA
Email: iainm@cs.cmu.edu

Abstract

Active Appearance Models (AAMs) are generative, non-rigid, parametric models of a particular visual phenomena. They are most often applied to faces, which are perhaps the most popular class of deformable objects in computer vision.

There are two elements to using AAMs: first, how to create the model, and second: given a model, how to fit it to a given image? This keynote will give an overview of the standard methods of model construction, then go into detail on the fast, analytical fitting solutions we have developed.

Keywords: Active Appearance Models, non-rigid face tracking, gradient descent image alignment, 2D+3D AAM.

1 Brief Overview of the Keynote

The computer vision community and industry have made great progress on face detection and recognition. The current popular approaches are now mostly based on single image analysis using sliding detection windows and boosted classifiers. They work impressively well for frontal face images. However, there are many more adjuvant application areas for facial analysis if we are able to accurately locate and describe faces in real-time through video sequences. This allows analysis of what a face is "doing", rather than just where it is and who it belongs to. This is a more appealing and difficult problem and is the motivation for our research.

The approach we use is based on Active Appearance Models (Cootes, Edwards & Taylor 2001) (AAMs) but uses an efficient gradient descent fitting algorithm (Matthews & Baker 2004). This approach has led to several extensions to basic algorithm. For example, the 2D+3D algorithm allows us to recover 3D head pose and 3D face shape but can still be fit in real-time (Xiao, Baker, Matthews, & Kanade 2004). We can extend this further to make good use of 3D shape constraints across multiple, simultaneous images (Koterba, Baker, Matthews, Hu, Xiao, Cohn, & Kanade 2005).

The same mathematical framework has also proven useful for automated model construction (Baker, Matthews & Schneider 2004), and robust fitting under occlusion (Gross, Matthews & Baker 2006).

The keynote will give an overview of this recent work as well as ongoing extensions and applications, and the current limitations we encounter.

2 Acknowledgements

Many people have contributed to the work that is presented: Simon Baker, Ralph Gross, Jing Xiao, Seth Koterba, Krishnan Ramnath, Changbo Hu, Simon Lucey, Barry Theobald, Takahiro Ishikawa, Junya Inada, Fernando de la Torre, Alvaro Collet, Jeffrey Cohn, Takeo Kanade.

The research was supported in part by DENSO Corporation, the U.S. Department of Defense contract N41756-03-C4024, and the National Institute of Mental Health grant R01 MH51435.

References

Baker, S., Matthews, I., & Schneider, J. (Oct. 2004), 'Automatic construction of active appearance models as an image coding problem', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 26, No. 10, pp. 1380–1384.

Cootes, T.F., Edwards, G.J., & Taylor, C.J. (June 2001), 'Active Appearance Models', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 23, No. 6, pp. 681–685.

Gross, R., Matthews, I., & Baker, S. (June 2006), 'Active appearance models with occlusion', *Image and Vision Computing*, Vol. 24, No. 6, pp. 593–604.

Koterba, S., Baker, S., Matthews, I., Hu, C., Xiao, J., Cohn, J., & Kanade, T. (2005), 'Multi-view AAM fitting and camera calibration', *in Proc. 10th IEEE International Conference on Computer Vision (ICCV2005)*, Beijing, China, Vol. 1, pp. 511–518.

Matthews, I., & Baker, S. (Nov. 2004), 'Active appearance models revisited', *International Journal of Computer Vision*, Vol. 60, No. 2, pp. 135–164.

Xiao, J., Baker, S., Matthews, I., & Kanade, T. (June 2004), 'Real-time combined 2D+3D active appearance models', *in Proc. 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR2004)*, Washington DC, USA, Vol. 2, pp. 535–542.