

The AAM-API: An Open Source Active Appearance Model Implementation

Mikkel B. Stegmann

Informatics and Mathematical Modelling, Technical University of Denmark,
Richard Petersens Plads, Building 321, DK-2800 Kgs. Lyngby, Denmark
mbs@imm.dtu.dk, <http://www.imm.dtu.dk/~mbs/>

Abstract. This paper presents a public domain implementation of the Active Appearance Model framework and gives examples using it for segmentation and analysis of medical images. The software is open source, designed with efficiency in mind, and has been thoroughly tested and evaluated in several medical and non-medical applications.

1 Introduction

In recent years, analysis through synthesis has proven itself as a very robust and widely applicable approach to computerised interpretation of images. In particular, the Active Appearance Models (AAMs) [3, 1] have enjoyed much success in both medical and non-medical applications. Examples of the former include segmentation and analysis of cardiac cine MRI, cardiac perfusion MRI, echocardiogram time series, diaphragm dome CT, brain MRI, hand radiographs and face images. Refer to the review in [4] for the references to these cases.

Active Appearance Models are generative models capable of synthesising images of a given object class. By estimating a compact and specific basis from a training set, model parameters can be adjusted to fit unseen images and hence perform image interpretation. The modelled object properties are shape and pixel intensities. Training objects are defined by marking up each example image with points of correspondence. Using prior knowledge of the optimisation space, AAMs can be rapidly fitted to unseen images, given a reasonable initialisation. For further details on AAMs refer to [1, 2].

This paper presents an open source implementation of the Active Appearance Models and gives examples of its use in several medical applications.

2 The AAM-API

The AAM-API is a C++ implementation of the Active Appearance Model framework. Emphasis is put on performance to yield responsiveness for interactive applications, while yet maintaining a clear structure. The API is focused on single-image 2D applications, but the modularity enables extensions, such as coupled 2D time series and fusion of 2D images from different modalities, e.g. PET, CT and MRI. Such extensions have previously been described in the literature.




Fig. 1: Examples of the medical applications of the AAM-API. From top, left: Cardiac cine MRI, brain MRI, metacarpal radiograph, face image and cardiac perfusion MRI.

To facilitate transparency in the model building process a vast range of documentation can be exported. This makes the package well suited for educational purposes. Further, a method for model initialisation is included, thus rendering image segmentation and analysis a fully automated process. Example segmentations obtained using the API are shown in Figure 1.

Alongside the API are scripts for annotating images etc. and the binary AAMLab; a graphical user interface for manipulating annotations and running model searches. The latter is very useful for building new models bootstrapped from a small set of annotated examples. On OpenGL-compliant machines the AAMLab also features an exploratory real-time visualisation of the AAM deformation modes. A more thorough introduction to the package and accompanying reference data sets is given in [4]. Finally, several performance benchmarks can be generated in e.g. a leave-one-out study as a part of a clinical validation.

The software runs on the Microsoft Windows platform and requires two 3rd part libraries, namely LAPACK and VisionSDK, both available in source code form. Microsoft Visual Studio is required to modify or recompile the software. The complete package is described in more detail and is available for download from <http://www.imm.dtu.dk/~aam/>. Source code documentation and papers describing the applications shown in Figure 1 are also available from this address.

References

1. T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. In *Proc. European Conf. on Computer Vision*, volume 2, pages 484–498. Springer, 1998.
2. T. F. Cootes and C. J. Taylor. *Statistical Models of Appearance for Computer Vision*. Tech. report, University of Manchester, 2001. <http://www.isbe.man.ac.uk/~bim/>.
3. G. J. Edwards, C. J. Taylor, and T. F. Cootes. Interpreting face images using active appearance models. In *Proc. 3rd IEEE Int. Conf. on Automatic Face and Gesture Recognition*, pages 300–5. IEEE Comput. Soc, 1998.
4. M. B. Stegmann, B. K. Ersbøll, and R. Larsen. FAME – a flexible appearance modelling environment. *IEEE Trans. on Medical Imaging*, 2003 (to appear).