

Digital Clones and Virtual Celebrities

Facial Tracking, Gesture Recognition and Animation for the Movie Industry

Barnabás Takács¹, Thomas Fromherz, Steve Tice²
Virtual Celebrity Productions
3679 Motor Ave., Suite 200
Los Angeles, CA, 90034, U.S.A.

Dimitris Metaxas
University of Pennsylvania
200 South 33rd St.
Philadelphia, PA, 19104-6389, U.S.A.

Abstract

We describe a high performance facial tracking and animation solution, called the Digital Cloning SystemTM, specifically designed to meet the needs of the entertainment industry. Special and digital effects are now being used in most feature films, thus representing an important new application domain for facial and body tracking technologies. Virtual Celebrity's system consists of (i) a real-time module responsible for pre-visualization of the tracked data on the face of an animated celebrity, and (ii) an off-line component that suitably meets the needs of high quality special effects and film production. We briefly review the state-of-the-art and present a comparative analysis of needs, methods, and existing solutions in film production vs. Internet and video telephony applications. Experimental results and demonstrations prove the effectiveness of our method.

1. Introduction

Automated analysis and synthesis of facial expressions have been in the focus of face recognition research for the last decade. Experimental systems have been developed and tested for model-based coding techniques that achieve visual communication at very low data rates [1,2] and high compression of 3D animated data. More recently, the advent of Internet-driven avatar technology has created a strong need among users to assume a particular identity in virtual spaces and to freely interact with each other. For both applications, the principal front end is a real-time facial motion tracker that follows the motion of a small set of characteristic points on the face (typically less than 20), and uses gesture recognition to analyze, transmit, and re-animate events related to expressions and emotions [3-6].

We describe a new domain for facial feature tracking, expression analysis and animation which, although similar in formulation, poses a set of completely different technical problems. The digital effects industry (in film production) frequently faces the challenge of having to seamlessly integrate animated 3D models into 2D live-action footage. Such was the case in recent Block Buster films like *Titanic*, *Batman and Robin* and *Face-Off*⁷. In this paper, we present a brief overview of our achievements in advancing the state-of-the-art in facial feature tracking, gesture recognition, and photo-realistic facial animation for this particular application domain.

Virtual Celebrity Productions has designed and built a proprietary *Digital Cloning SystemTM* (DCS) which utilizes a dual tracker architecture and advanced facial animation tools to create photo-realistic/'seamless' human characters⁸. The two major components of our tracker module are: (i) a real-time tracking & animation system called *Geppetto*, and (ii) a very high quality off-line tracking/rendering pipeline. The DCS module accepts a sequence of film resolution images (4K by 4K pixels, 16 bit / color) and outputs tracking data and high quality 3D animation ready for film production. A detailed description of the DCS can be found in [8]. The specific system incorporates:

- A real-time facial tracker and gesture recognizer coupled with an animation software used to pre-visualize the effects of current tracking data.
- A set of specialized facial feature trackers that extract all the necessary information relevant for creating a photo-real *Digital CloneTM*.
- A camera tracker that derives the 3D motion of the camera as needed for special effects.
- A highly accurate off-line facial feature tracker that is capable of tracking up to 500 points without markers on an actor's face.

¹ Corresponding author - TakacsB@virtualceleb.com

² Currently with QuantumWorks Corporation.

- A deformable model-based 3D face tracker based on the integration of optical flow and edges extracted from the given image sequence [9].
- Interfaces to the major computer animation packages allowing for direct control of any muscle or surface deformation-based facial animation model.

2. State of the Art in Facial Tracking

The various facial motion estimation approaches found in the literature can be categorized either as model-based or non model-based. Despite the number of non model-based techniques [11], in this paper we focus on model-based approaches, thus taking full advantage of the context of the tracking problem to achieve the highest possible quality. In the model-based domain a wide variety of face models have been used. Among them are 2D face models based on splines or deformable templates and 3D models such as the one by Terzopoulos and Waters [12] where face motion is measured in terms of muscle activations. Other simple 3D models have been used to address the problem of image coding. Commonly all these techniques are based on either edge information or optical flow and perform well under limited circumstances. On the other hand, none of the above approaches can accommodate large head rotations due to (i) the simplicity of the underlying models used, (ii) the use of a single visual cue, and (iii) the lack of handling self-occlusion. Contrary to the above techniques the methodology originally introduced by DeCarlo and Metaxas [13] allows for simultaneous estimation of 3D shape and motion based on the integration of information from optical flow, edges and other shape constraints.

In general, feature tracking is a two stage process. The tracking routine must first find the exact location of the features in each frame, and subsequently solve the correspondence problem as defined by these features. Global constraints on the motion itself, such as partial rigidity and prediction schemes [14-16] (e.g. Kalman filtering) can also be used to improve accuracy. Still a critical bottleneck in this process is the feature selection algorithm [61,72].

Some of the most frequently used solutions group the pixels surrounding the tracked point, into a single feature signature. Examples are found in the vast literature and extensive studies on image correlation [4,60,89], sum of square difference methods (SSD) [22], adaptive window size [23], on using optimization criterion with respect to translation and linear image deformations [24,25], and on the quality of the features being tracked [26].

Recently Gabor filters were successfully used to locate and track features in face images [27-29]. The basic visual feature is a local image descriptor, called a “jet” [30], of which each component is the filter response of a Gabor wavelet of specific frequency and orientation. The tracking system would typically use wavelets of 4-6 different frequencies and 8-10 orientation, thus defining a total of 32-60 complex components characterizing the local patch around the tracked point.

Optical flow in an image is produced by motion, and therefore it calls for low-level processing of high density pixel data. The key advantage of this technique lies in its holistic nature, i.e. it does not require preliminary feature detection, and can be used as a direct representation of facial action after proper post-processing of the motion field. However, the technique is quite sensitive to noise and is not well suited for capturing rapid motion. Despite these difficulties, optical flow was successfully used to track a dense set of points of facial features conveying subtle changes in expressions and emotions. As an example, in [9] the authors incorporate the optical flow equations into a deformable model formulation combined with Kalman filtering and achieve highly accurate tracking results. Other researchers have also coupled multi-grid flow calculations with a physical model of a face [31,32].

3. High Accuracy Face Tracker and Facial Animation System

The overall architecture of the *VCP Face Tracker* is presented in Figure 1. The basic API consists of multiple layers starting from a set of basic tracking algorithms to 2D/3D and camera motion trackers. On the top of this general purpose API, a set of application specific routines were developed that extract highly detailed 3D data of head motion (yaw, pitch, roll, 3D position) and facial features such as eyes (pupils, eyelids), mouth corners, lips, nose, etc. The user can also specify to track a number of occasionals including wrinkles and miscellaneous points. Once all subtle facial movement is extracted the tracker proceeds to generate animation data, first by rendering a real-time version using VCP’s *Geppetto* system, and subsequently creating direct animation data for high end tools such as SoftImage, Houdini, and/or Maya.

We now describe the basic building blocks of the system in detail. The real-time module of the tracking system was developed for producing low-cost animation for Internet and TV. For feature films, the high quality tracking is an off-line process. However, the user can still utilize the power of the real-time mode, as a pre-

visualization engine to show what the final animation

characters. The motion capture module was built from




Figure 1: Overview of Virtual Celebrity Production's face tracker architecture.

will look like. In the off-line mode, any frame in a sequence can be processed at any time and as many times as needed. The processing time is not of major importance; accuracy, however, is. A typical input image would have a 4096x4096 pixel resolution and 16 bits per color channel. When formulating the tracking algorithm, we needed to take into consideration that the presence of highlights, shadows, overlaps, and multiple (often colored) light sources inherently limit the performance of any point-tracker algorithm. Due to these limitations of the imaging process and the total lack of control over the environment, one is required to design an integrated tracking solution comprising of multiple basic algorithms. Table 1. summarizes the differences and similarities of facial tracking and gesture recognition in the entertainment industry vs. the classical domain, i.e. video telephony, avatars, and human-computer interfaces. As is easily seen, the technical requirements of these two seemingly similar problems often require fundamentally different solutions.

Real-time Facial Tracking and Gesture Recognition:

As an important component of the *Digital Cloning System*TM, a real-time gesture recognizer and animator module serves to pre-visualize the quality and integrity of current tracking data. It is based on a performance animation system we have developed, called the *Geppetto*TM Animation System, which is able to produce *Interactive Digital Performers*. Unlike other cartoon-like character systems, it captures the subtleties of minute actor eye movements, and while adding subtle animation behaviors, produces a realistic human quality in our

off-the-shelf components [10].

Although *Geppetto* has its own built-in real-time tracker currently tracking 12 points on a face, it can also be fed with data from the off-line module of the *VCP Face Tracker*. The points are connected to an underlying muscle-expression system. From this tracking data a mixture of 19 basic muscles and multitudes of expressions are recognized and rendered. Gesture recognition is a two-stage process consisting of (i) an off-line calibration and (ii) real-time tracking & animation. During calibration, a sequence of example facial gestures like 'right side smile' or 'jaw down' are processed, building a relationship between tracked data positions and person specific gesture values. A learning algorithm is responsible for adapting the particular point configuration on the actor's face to the generic muscle model of the celebrity face model. In the animation stage, the pre-visualization module receives tracking data from the *VCP Face Tracker* and translates it into gesture values. This translation is supported by a rule table built up in the preceding learning step, thus making the translation robust and less ambiguous. The resulting gesture values consist of morph targets, which are applied to a corresponding 3D model of the animated character.

Basic Tracking Algorithms: We implemented a set of basic tracking algorithms that deliver subpixel accuracy. Specific algorithms include region-based tracking algorithms, such as normalized correlation, texture-based trackers (Gabor jets), and optical flow. We found that each of these solutions exhibit serious limitations in a

real-world production environment. Figure 2 shows a typical frame from a TV commercial in which the actor

acts in front of a green screen later to be replaced by a

	Photo-realistic Film Production	Video Telephony & Virtual Telepresence
GOAL:	Convey acting, emotions, personality	Communication, anonymity & fun
IMAGING:		
# cameras	Single or multiple	Single
Camera Type	Unknown / High performance professional	Unknown / Low-cost commercial
Camera properties	Uncalibrated (mostly unknown)	Uncalibrated (don't care)
Camera motion	YES	NO (stationary)
Image size	Very high resolution	Low to mid-resolution
	4Kx4K / 16bits * 4 color	128x128 or 256x256
Color	YES (principle movie camera)	YES
	NO (high speed & resolution tracking cameras)	
Capture speed	24-30 fps for film	10-30 fps
	30-120 fps for additional tracking cameras	
ENVIRONMENT:		
Conditions	Uncontrolled	Uncontrolled
Illumination	Changing / multiple light sources, different color	Fairly constant
TRACKING:		
Face Finding:	Semi-automatic / automatic	Automatic
Camera motion	YES - high accuracy 3D tracking	NO
Light source estimation	YES	NO
Markers	Markerless or with some make-up	NO
# tracked points	100-500	5-20
Face Tracking:	Automatic / subpixel accuracy	Automatic / coarse
Global head motion	0.1 degrees accuracy	1-5 degs accuracy for user interaction
Facial features	Highly detailed with subpixel accuracy	Coarse to medium accuracy
	Needs to track unstructured facial regions	Mostly tracks well defined fiducial points
	Pupils, eyelids, gaze	Eyes, blink detection
	Mouth region, inner-outer lips	Mouth corners, lip tracking
	Facial surface / exact 3D data estimation	NO - expression detection only
	Wrinkles & occasional	NO
Speed:	Real-time and/or off-line	Real-time (15-30 fps)
Processing	Multiple passes, bi-directional	Single-pass, unidirectional
ANIMATION:		
Model construction	Manual / pre-production	Automatic, off-line
Model adaptation	Semi-automatic	Automatic
	Animated individual is a different person	Animated individual is the same person
Gesture recognition	YES (previsualization only)	YES
Advanced Animation	YES (Houdini, Maya, SoftImage, etc.)	NO
Photo-quality	YES	Currently NO, but required in the future
Compositing	YES / semi-automatic	YES / automatic
Postproduction	YES	NO
MISC.:		
Voice data:	Not always available	YES
Multi-modal	YES (visual & voice)	YES (visual & voice)

Table 1: Comparison of facial tracking specifications for the movie vs. the communication / Internet industry.

background image using chroma-key methods. The over 350 points (marked in the figure) were automatically tracked by our system. Note the dark chin line and the low-contrast of the lip contour. Although the face itself is well lit, there are several image regions where highlights and dark shadows might cause the above algorithms to lose points. To overcome the problems associated with each of the above techniques, we designed a hybrid tracking algorithm that combines the advantages of these methods while minimizing the sensitivity to noise and imaging condition. The *VCP Point Tracker* adaptively selects the best method for each tracking point in a given frame. The image is then processed multiple times and the results are integrated, minimizing a cost function defined over all tracking points. Since this is the first layer in the API, we did not put much emphasis on achieving 100% accuracy. Instead, we designed higher level 2D/3D trackers which impose a geometric model on the tracking process and used this model to correct for

inaccuracies from lower API layers. This geometric correction layer operates on top of region-based, Gabor jet, and optical flow tracking algorithms.

The second layer of our architecture implements a set of *structural trackers* that, instead of following individual points in the image plane, take advantage of an underlying model geometry to which these (individual) points belong. Thus, one can track point groups, 2D structures, and/or splines as well as true 3D objects, imported in the form of a geometry description file. The 3D tracker can also be set to maintain rigid shape or accommodate non-rigid deformations.

Real-time Animation: The *Geppetto* system mentioned above is used as a real-time visualization tool to display how the tracking data modulates the expressions of the animated celebrity. As described earlier, this sub-system uses gesture recognition to estimate muscle parameters from a small number of tracked points in a form similar

to Facial Action Coding (FAC) units, and utilizing morph-targets translates these values to animation data.

updating, and improving our 3D models. Furthermore, by means of this off-line process, we create a very detailed




Figure 2: 350+ tracked facial points, fitted splines and global head model used by the VCP tracker.

These morph-targets are built automatically from the high resolution models which our animators create for each film project.

Film-quality Animation: For the film-quality animation, the *VCP Face Tracker* was interfaced with high performance animation packages, such as SoftImage, Houdini, and Maya. The interface controls an iterative off-line process that converts the tracking data into animation data ready for further production stages. The algorithm iteratively deforms the high resolution facial surface using a set of effectors; muscles or channels. First, it estimates the intrinsic camera parameters used to shoot the footage, then using these settings proceeds to find the best muscle configurations responsible for a particular expression on the actor's face.

In the present version of the software, deformations of the facial surface are controlled through a total of 23 channels, 19 of which correspond to the mouth and cheek areas and 4 control the eye movements. While some channels directly relate to anatomical muscles found on the human face, others are linked into virtual muscle groups to help create all possible facial expressions. Each muscle value is estimated to an accuracy of 0.5% on a full range of 0-100%. In addition, we also allow for over-stretching (up to 200%) if the valid range of a particular muscle does not accommodate the tracking data. This feature was found to be very useful in creating,

and accurate data set supporting the learning module of the gesture recognizer above³.

In summary, we have developed a dual animation architecture bringing the high speed and robustness of gesture recognition together with precise parameter estimation in the model space. The gesture recognition engine (*Geppetto*) offers real-time pre-visualization of the tracked data on the animated person's face using only a limited set of 12 points and low resolution (subsampling) images. It also generates seed expressions for the subsequent film-quality process. The high performance interface (SoftImage, Houdini, Maya), on the other hand, works in an off-line manner, typically processing a given film sequence overnight. The module accepts high resolution images (max. 4Kx4K), and presents its output in an environment suitable for film production professionals. Thus, the conversion of 2D markerless film footage to 3D animation happens in a fully automated way, thus saving time and effort of the visual artists working on the project.

4. Experiments & Results

³ Values of 23 muscles creating facial expressions within 0.5% accuracy are stored together with 2D tracked coordinates and 3D geometry.

We have tested the performance and accuracy of our facial tracking and animation solution in the framework of VCP's *Digital Cloning SystemTM*. We have conducted a series of tests using a variety of cameras and performance scenarios. The real-time module tracks 12 markers and calculates 23 basic muscles and a multitude of expressions at 30 frames per second. The tracking speed of the off-line production system is 1-30 seconds per 4Kx4K frame depending on the number of tracked points (currently we allow a maximum of 500 points), and, most importantly, depending on the required accuracy of the calculated data. Specifically, these processing times correspond to subpixel accuracy tracking in the image domain, 0.1 degree accuracy of global head rotation (yaw, pitch, roll), and 10 mm accuracy in global head and camera motion along the Z axis. Once tracking is complete, the high performance animation module applies the extracted motion and fills the animation channels at a speed of 1-10 minutes per frame.

Figure 3 (left) shows frames from a typical input sequence used in a TV commercial. The actor is placed in front of a green screen to help later compositing. Key facial points are tracked and applied to the high resolution model of another person. The final results corresponding to the input are shown on the right. Figure 4 shows two examples of VCP's digital celebrities. W.C. Fields is a famous American comedian and entertainer. Two of his original photographs (above) and the corresponding digital model (below) can be seen on the left. On the right, Marlene Dietrich, a German-born actress who became a global icon in the 30's. Two of her original photographs (above) and the corresponding digital model is shown (below). We are currently working on a new project potentially featuring Marlene in a commercial. For more information, please visit our websites at <http://www.virtualceleb.com> and <http://www.globalicons.com>.

5. Conclusions

In this paper, we described a high performance facial tracking and animation solution, called the *Digital Cloning SystemTM*, specifically designed to meet the needs of the entertainment industry. The system consists of the real-time component *Geppetto*, responsible for pre-visualization of the tracked data on the face of an animated celebrity, and an *off-line* module that suitably meets the needs of high quality special effects and film production. We developed this dual animation architecture after reviewing in detail the technical advantages and disadvantages offered to the tracking problem by various solutions. Our experimental results and demonstrations prove the usability of our technique

in the targeted application environment. We believe this technology is the first step towards a new age in making films in Hollywood.

6. References

- [1] Pearson, D.E. and J.A. Robinson (1983), Visual Communication at Very Low Data Rates, *Proc. IEEE*, **93**(4), 795-812.
- [2] Choi, C.S, K. Aizawa, H. Harashima, and T. Takebe (1994), Analysis and Synthesis of Facial Image Sequences in Model-Based Image Coding, *IEEE Trans. On Circuits and Systems for Video Technology* **4**(3), 257-275.
- [3] D. Terzopoulos, F. Parke, D. Sweetland, K. Waters, M. M. Cohen (1997), SIGGRAPH 97 Panel on Facial Animation: Past, Present and Future, <http://mambo.ucsc.edu/psl/sig97/siggraph97-panel.html>, Los Angeles.
- [4] K. Waters, T. Levergood (1995), DECface: A System for Synthetic Face Applications, *Multimedia Tools and Applications*, **1**, 349-366.
- [5] K. Waters, J. Rehg, M. Loughlin, S. B. Kang, D. Terzopoulos (1996), Visual Sensing of Humans for Active Public Interfaces, *Digital Cambridge Research Lab TR 95/6*.
- [6] Y. Yacoob, L. Davis (1994), *Computer Vision and pattern Recognition Conference*, chapter Computing Spatio-Temporal Representations of Human Faces, *IEEE Computer Society*, 70-75.
- [7] Robertson, B (1998), Right on Track, *Computer Graphics World*, **5**, 54-66.
- [8] Takács, B. and T. Fromherz (1999), VCP Face Tracker and Animation Software, *Tech. Rep. TR-VCP99-01*, Virtual Celebrity Productions, Los Angeles, California.
- [9] DeCarlo, D. and D. Metaxas (1996), Deformable Model-Based Face Shape and Motion Estimation, *ICAFGR*, 146-150.
- [10] Tice, S.E. and M.J. Fusco (1993), Digital Content Creation Systems, *Character Motion Systems Course #1, SIGGRAPH '93*.
- [11] Y. Yacoob, L. Davis, *Computer Vision and pattern Recognition Conference*, chapter Computing Spatio-Temporal Representations of Human Faces, pp.70-75, *IEEE Computer Society*, 1994.
- [12] Terzopoulos, D. and K. Waters (1993), Analysis and Synthesis of Facial Image Sequences Using Physical and Anatomical Models, *IEEE Trans. PAMI*, **15**(6), 569-579.
- [13] DeCarlo, D. and D. Metaxas (1996), Deformable Model-Based Face Shape and Motion Estimation, *ICAFGR*, 146-150.
- [14] Jebara, T.S. and A. Pentland (1996), Parametrized Structure from Motion for 3D Adaptive Feedback Tracking of Faces, *Tech. Rep. MIT Media Laboratory #410*, Massachusetts Institute of Technology.
- [15] Basu, S., N. Oliver, and A. Pentland (1998), 3D Modeling and Tracking of Human Lip Motions, *ICCV*.

- [16] Metaxas, D. and D. Terzopoulos (1993), Shape and Nonrigid Motion Estimation through Physics-Based Synthesis, *IEEE Trans. PAMI*, **15**(6), 580-591.
- [17] Shi, J. and C. Tomasi (1994), Good Features to Track, *CVPR*.
- [18] Tian, Q. and M.N. Huhna (1986), Algorithms for Subpixel Registration, *CVGIP*, **35**, 220-233.
- [19] Burt, P.J., C. Yen, and X. Xu (1982), Local Correlation Measures for Motion Analysis, *IJCV*, **2**(3), 283-310.
- [20] Ryan, T.W., T.R. Gray, and B.R. Hunt (1980), Prediction of Correlation Errors in Stereo Pair Images, *Optical Engineering*, **19**(3), 312-322.
- [21] Wood, G.A. (1983), Realities of Automatic Correlation Problem, *Photogram. Eng. And Rem. Sens.* **49**, 537-538.
- [22] Burt, P.J., C. Yen, and X. Xu (1982), Local Correlation Measures for Motion Analysis, *IJCV*, **2**(3), 283-310.
- [23] Okutomi, M. and T. Kanade (1992), A Locally Adaptive Window for Signal Matching, *IJCV*, **7**(2), 143-162.
- [24] Lucas, B.D. and T. Kanade (1981), An Iterative Image Registration Technique with an Application to Image Vision, *IJCAI*.
- [25] Manmath, R. and J. Olinesis (1993), Extracting Affine Deformations from Image Patches – I: Finding Scale and Rotation, *CVPR*, 754-755.
- [26] Tomasi, C. and T. Kanade (1991), Detection and Tracking of Point Features, *Tech. Rep. CMU-CS-91-132*, Carnegie Mellon University.
- [27] Maurer, T. and C. von der Malsburg (1996), Tracking and Learning Graphs and Pose on Image Sequences of Faces, *ICAFGR*, 176-181.
- [28] Elagin, E., J. Steffens and H. Neven (1998), Automatic Pose Estimation System for Human Faces based on Bunch Graph Matching Technology, *ICAFGR*, 136-141.
- [29] Eyematic Interfaces Inc. (1998), Visual Interface Technology API Overview, personal communication.
- [30] Lades, M. *et.al.* (1993), Distortion Invariant Object Recognition in the Dynamic Link Architecture, *IEEE Trans. Comp.*, **42**(3), 300-311.
- [31] Cohn, J.F. and J.J. Lien (1998), Beyond Prototypic Expressions: Discriminating Subtle Changes in the Face, *IEEE Robot and Human Communication Workshop*, Takamatsu, Japan.
- [32] I. A. Essa, A. Pentland, A Vision System for Observing and Extracting Facial Action Parameters, *Proc. of Computer Vision and Pattern Recognition (CVPR 94)*, pp.76-83, 1994.



Figure 3: Typical input (left) and rendered output (right) sequence of a talking head used in a TV commercial.



Figure 4: Original photographs of W.C. Fields (above left), a famous american comedian, and Marlene Dietrich (above right), with their respective Digital ClonesTM (below).