This paper appears as it did in the proceedings. Since that publication, a bug in the implementation has been corrected. For corrected results, please see the web-based version of this paper pointed to by http://www.interval.com/~covell/pubs.html.

Autocorrespondence:
Feature–based Match Estimation and Image Metamorphosis

Michele Covell
Interval Research Corporation; 1801 Page Mill Road, Bldg. C; Palo Alto, CA 94304; USA
covell@interval.com

ABSTRACT

This paper addresses the problem of matching distinct but related objects across images, for applications such as view-based model capture, virtual presence, and morphing special effects. A new approach to match estimation is presented for estimating correspondances between distinct objects. This approach minimizes the effects of lighting variations and non-rigid deformations. A sparse set of features are used for alignment, followed by image-plane warping, or morphing, based on these constraints.

1. INTRODUCTION

Inter-image interpolation is a common element of many diverse applications: in film-to-television rate conversions, in interpolated display of slow-motion sequences and time-lapse photography, in MPEG video decompression, in camera-viewpoint interpolation[1] and in morphing special effects.[2]

In its simplest form, inter-image interpolation is completed by cross-fading two images. This approach results in obvious misalignment artifacts in many situations. Motion artifacts are one well-known example. A similar artifact can occur when the original images present two distinct but related objects. For example, if the two originals are distinct portraits, motion artifacts will be seen in the cross fade, if they are not correctly aligned.

To minimize misalignment artifacts, inter-image interpolation is completed after some type of matching operation. Both stereo mapping and motion estimation have received vast amounts of attention from the vision and signal processing communities in the last ten years and now constitute fairly mature technologies. Most of these techniques have been developed to match views of a scene, from closely-spaced points in time *and* viewpoint (motion estimation on video images or stereo matching on simultaneous images). In contrast, this paper concentrates on the problem of matching views of distinct objects, such as portrait images of different people.

When distinct but related images are used as the source and target images, match estimation must deal with potentially large, locally discontinuous offsets, with areas of occlusion and disocclusion and with differences in lighting and surface reflectance. Scale-space methods[1,2], non-linear local transform methods[3], feature-based methods[4,5], and non-linear distance measures[5] have all been used in motion/stereo to address similar difficulties. Unfortunately, their underlying assumptions (about the stability of shape and the rigidity of the motion, or about the lack of rotation/expansion/compression) will be violated within the more general correspondence problem, since there is little coherence between the images. A new approach to match estimation is presented which reduces the severity of these assumptions.

[1] Camera-viewpoint interpolation is needed for interactive walk-throughs and other virtual presence applications and for some types of view-based model capture.

[2] Morphing special effects provide a natural looking transition between distinct images. Examples can be seen in the movie *Terminator II* and in Jackson's video *Black or White*.

This new matching procedure is applied to the general correspondence problem. The applications which require matching between distinct objects all allow some amount of manual intervention. These applications include those listed under camera-viewpoint interpolation and morphing special effects. In these applications, the automatically estimated correspondences can be used to reduce the amount of manual labelling (Figure 1). Since some manual intervention is expected, it is both the number of matching errors and the ease with which they can be corrected that is used as quality criteria for the correspondence estimates.

Matching errors can be easily corrected if the number of independent matching constraints is minimized and the image alignment is determined from a sparse set of constraints. Feature-based image metamorphosis[6], or morphing, fulfills these requirements. This paper and our previous results[7] are the first to combine match estimation with feature-based image morphing. By combining match estimation and image morphing, a natural transformation between distinct images can be created, without the labor intensive process of manually placing all constraint-segment pairs.

The next section quickly outlines feature-based image morphing. A new approach to match estimation is described in Section 3. Section 4 shows the results of morphing based on the correspondence constraints estimated by this approach. These results are compared to those reported previously. Finally, conclusions and future directions are discussed in Section 5.

2. FEATURE-BASED IMAGE METAMORPHOSIS

Feature-based image metamorphosis, or morphing, is a combination of image warping and crossfading. While a variety of algorithms exist for image warping based on rectangular-gridpoint correspondences[8] and triangle correspondences[9], the algorithm used here is based on line-segment correspondences[6]. First, consider a single line-segment correspondence constraint, with endpoints A and B in the source image and endpoints C and D in the target image. Conceptually, the constraint results in a translation of the source image plane so that A is superimposed on C; followed by a rotation, so the line (A,B) lines up with the line (C,D); followed

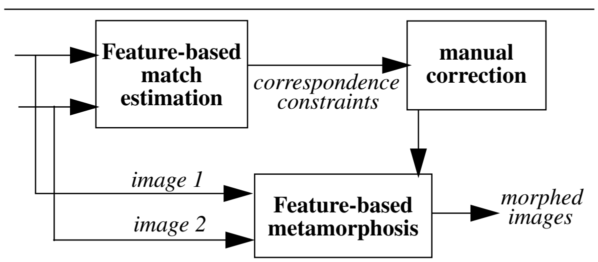

Figure 1: Overview of feature-based match estimation and image metamorphosis. Feature-based matching yields a set of correspondence constraints, which map line segments between images. Feature-based image metamorphosis uses these correspondence constraints to control the geometry of its image warping. For some applications, manual intervention is possible between matching and morphing.

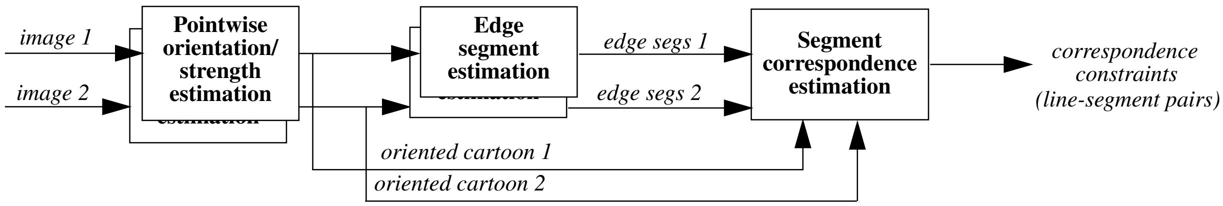

Figure 2: Overview of feature-based match estimation. An oriented cartoon of each image is created by estimating the local orientation and the strength of that orientation at each point. The strongly oriented points are linked together into edge segments. Each edge segment image is matched against the other cartoon, with the resulting correspondence constraints being pooled together to constrain future matches of other segments. Finally, these correspondence constraints are used to define the geometry for image morphing.

by a compression or expansion of the image plane in the direction of (A,B), so that B is superimposed on D. The direction orthogonal to the rotated (A,B) is neither expanded or compressed, based on this one constraint. Multiple constraints are combined by taking a weighted average of their image-plane transformations, where the weighting is proportional to the distance from the line segment.

This approach to image morphing allows us to use simple line-segment correspondence estimates as our input. The next section describes a method for estimating these line-segment correspondence constraints.

3. FEATURE-BASED MATCH ESTIMATION

Figure 2 shows a flow graph of the proposed approach to feature-based matching. The first step in the process is to estimate an orientation for each point and the strength of that orientation. I will refer to the images of these features as oriented cartoons. The points of strongest orientation are then combined and summarized in the form of edge segments. Each edge segment covers neighboring points which are approximately aligned with the segment, both in location and in orientation. These edge segments, along with the original oriented cartoons, are used to estimate the image correspondences. The remainder of this section discusses each of these steps in detail.

3.1. Oriented cartoons: estimating salience and orientation

Oriented cartoons are dense, pointwise estimates of orientation and the strength (or salience) of that orientation. The oriented cartoon will determine both the location of potential correspondence constraints and the degree of match or mismatch between the images. Therefore, the oriented cartoons need to be stable under lighting changes. The features should also be conservative in how they measure salience: lower importance should be attached to strong unstable features than to fainter stable features. Finally, the oriented cartoons should emphasize compact regions around these features, since spatially compact regions will result in sharper match/mismatch metrics. This subsection describes a way of computing oriented cartoons to fulfill these requirements. Figure 3 provides an example of the result.

To emphasize compact regions around features, I use the orientation and the amplitude of the maximum curvature at each point in the image plane. To improve the lighting invariance, I normalize the curvature amplitude by an estimate of the local average illumination. Since measures of curvature and local illumination both implicitly require a chosen scale, I estimate each at multiple scales and combine the normalized curvature and its orientation across scale. This avoids the need to choose a single analysis scale. It also provides conservative estimates of salience. Areas with oriented energy at multiple scales are emphasized over ones with energy at a single scale, thus increasing the likelihood that the area corresponds to a stable edge feature.

I use steerable filters[10] to estimate the orientation and amplitude of the maximum curvature at each scale. The three basis functions for all oriented second-derivative of a circular Gaussian kernel and the four basis functions for its quadrature pair are given in [10]. The orientation of maximum curvature is then given by:

$$\cos 2\theta = (G_{xx} - G_{yy})\,\mathrm{sign}(G_{xx} + G_{yy})$$

$$\sin 2\theta = -2G_{xy}\,\mathrm{sign}(G_{xx} + G_{yy})$$

where G_{xx}, G_{xy}, and G_{yy} are the image responses of the three basis function defined in Table III of [10].[3]

The amplitude of the maximum curvature is then:

$$G_{\theta\theta} = G_{xx}\cos^2\theta - 2G_{xy}\cos\theta\sin\theta + G_{yy}\sin^2\theta$$

The maximum curvature is a different measure of edge strength from that given in [10]. Using the maximum curvature results in a narrower line-impulse response and narrower but spatially offset line-step responses.

Additional information about the local topology can be derived from the sign of the response of the quadrature filter. Using the definition:

$$\theta_{max} = \begin{cases} \theta & H_{\theta\theta} \geq 0 \\ \theta + \pi & H_{\theta\theta} < 0 \end{cases}$$

then θ_{max} will point perpendicular to the edge, in the direction of increasing brightness.

The salience is then defined as the sum across scales of $\frac{1}{\sigma}G_\theta / G_{avg}$. Weighting by inverse scale compensates for the 1/f energy distribution which tends to occur in natural imagery. It should be noted that the salience will take on both positive values (marking dark lines and the dark side of step edges) and negative values (marking bright lines and the light side of step edges). A single estimate for orientation is similarly derived. This estimate is computed by vector addition on $(\frac{1}{\sigma}G_\theta, \theta_{max})$, starting with the coarsest scale and successively adding in each finer scale.

3.2. Edge segments: clustering and summarizing an oriented cartoon

Edge segments are both groupings and summaries of the strongly salient points in the oriented cartoons. Each edge segment is a grouping of salient points in the segment-based matching. Points are clustered together by their assignment to a shared segment but it is their true locations, orientations and strengths that are matched (as a group) against the oriented cartoon of the other image. The edge segments which are successfully matched are also

3. The definitions in [10] are for $\sigma^2 = 1/2$: the indices and amplitude of the basis functions must be scaled appropriately for other values of σ^2.

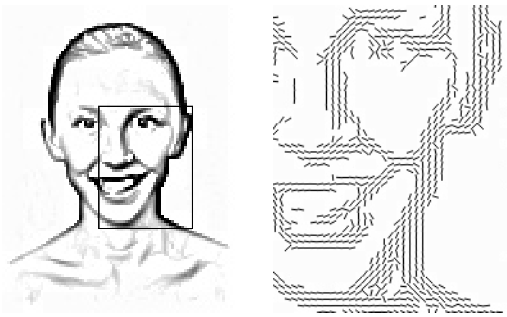

Figure 3: Simple rendering of an oriented cartoon. The left image shows the positive values of local salience. The right image shows the local orientation at points with positive salience values. The lines parallel to the edge orientation.

Figure 4: Examples of chained curves and their approximating line segments. The left image shows the chains of positive-valued points that are found from the oriented cartoon. The right image shows their summarizing line segments.

used in the morphing process. On these segments, the endpoints of the segments and of its (virtual-segment) pair summarize the locations and the destinations of the underlying points. These two roles for edge segments give rise to a number of requirements.

Since each cluster of points is treated as a coherent unit in matching, the grouping of points into segments should only include points which suffer a common fate: that is, points which as a group undergo a single rigid translation, rotation and scale change. Since the endpoints of the edge segments and their matching pairs are used in morphing to summarize a single constraint on all the intervening points, the edge segments should trace the locations of the underlying points with reasonable fidelity. Edge segments which intersect one another anywhere other than their endpoints should also be avoided. This is because, if both segments result in correspondence constraints, these constraints are unlikely to be consistent at their intersection. Assuming the earlier requirements are met, the number of edge segments should be minimized while still describing the salient areas. The advantages of this are seen both in matching and in morphing. In matching, the larger the cluster (and therefore the fewer clusters), the less ambiguity will be seen in the choice of the best match.[4] In morphing, fewer constraint lines require less computation and, since local matching errors will be spread over fewer constraints, less effort is needed to correct mismatches.

This list of requirements can not be easily satisfied using standard line fitting techniques[11]. Instead, I have used a two-step process for finding edge segments. The first step chains together points in the oriented cartoons to form edge curves. The second step assigns the links of the curves to edge segments.

I first chain together similar points within the oriented cartoon to form curves. The left image in Figure 4 provides an example of edge curves. To insure simple chains of points which can be easily subdivided into edge segments, only two links are allowed at any single location: one “to” the location and one “from” it. The roles of source and destination of a link are symmetric—they are used only as a way of distinguishing a point orientation from its 180º rotation (for which the direction of increasing brightness is flipped).

Since the edge curves determine many of the characteristics of the final edge segments, the requirements listed for edge segments need to be considered in chaining together edge curves. To help insure that a edge segment will move rigidly between images, the links within an edge curve should join points that:

– are both above some minimum salience;
– are near to one another;
– share similar orientation estimates; and
– share a baseline approximately aligned with the right-hand perpendicular to their maximum curvature.[5]

As mentioned above, edge curves are only strung between points in the oriented cartoon which have salience values above some minimum. This insures that the points underlying the edge curves all have a well-defined orientation: reliable orientation estimates will be needed for matching. This threshold forces the edge curves into a compact region near edge features. Non-maximal suppression could be used to further reduce the support of edge curves near a single feature. However, this did not seem necessary in view of the data reductions which are provided by subsequent processing.

The other three requirements mentioned above increase the likelihood of a common fate within the curve and decrease the likelihood of intersections between curves. The chances of a common fate within an edge curve are increased by linking nearby points that align with each other and with their baseline. Points with these shared characteristics are likely to all mark one side of a single, underlying edge feature. Staying to one side of an edge feature avoids linking together features on separate occluding surfaces (such as the outline of a head and the edges of the background). It also avoids linking together points which are on independent but abutting surfaces (such as pursed lips or closed eye lids). This improves the chances that the linked points are neighbors on a single object surface and will therefore move in a coherent manner.

Linking nearby points that align with each other and with their baseline also decreases the chances of two curves crossing one another. Often, two edge curves will parallel each other for most of their length, with little variation in their separation. When two edge curves approach each other without aligning, one or the other curve will typically end near an implied point of convergence. These collisions are detected in segmenting edge curves and become the endpoint locations for edge segments on the intersecting curves.

Finally, the salience of the weaker side of a link is boosted, encouraging another connection to be made to or from that point. This encourages longer chains over many short chains. Longer chains allow for longer edge segments and, ultimately, for a smaller, more concise set of constraint lines describing the same

4. This is true only if the larger cluster also shares a common fate.

5. The baseline of a potential link is the vector between its source and destination. By using a handed perpendicular and by introducing the roles of source and destination to links, edge curves which switch-back on themselves are avoided.

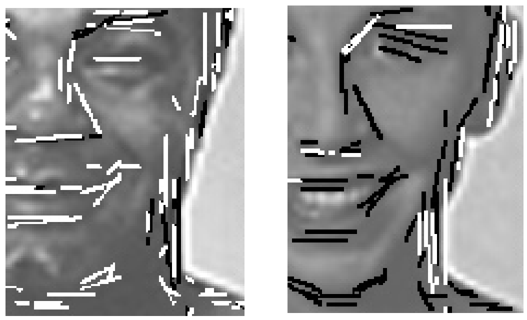

Figure 5: Examples of correspondence constraints. These two images show some of the correspondence constraints found using the algorithm described in Section 3.3. Each constraint is made up of one edge segment (shown in white) and one virtual edge segment (shown in black).

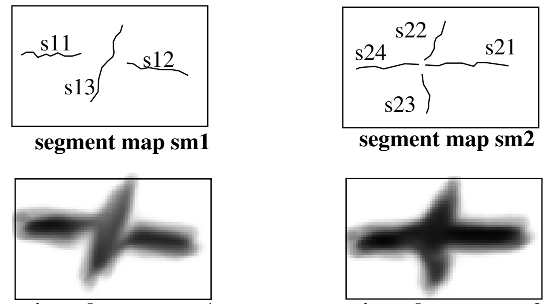

oriented cartoon oc1 **oriented cartoon oc2**
(only the amplitudes are sketched here)

Figure 6: A simple illustrative example for segment-based matching. The example shown here is used as a simple illustration of the process used in segment-based matching.

underlying warping.

Once edge curves have been linked together using this process, they are subdivided into edge segments. Figure 4 provides an example of edge curves divided into edge segments.

The edge segment locations are already highly constrained by the edge curves that were chained together in the previous step. The only remaining decisions are where along each edge curve to place the endpoints of the segments, in order to break the curve into nearly linear, non-intersecting sections. This segmentation uses the following criteria:

– Curve ambiguity: Endpoints are fixed at any intersection points between edge curves. This avoids inconsistent constraints at the intersection points. It also removes the ambiguity represented by these intersections. At the intersection of two curves the best way to connect the curves across the intersection is often ambiguous. By putting segment endpoints at the intersection, the need to trace which curve goes which way is avoided. For the same reason, segment endpoints are placed at implied convergence points. A convergence point is implied where two edge curves approach each other without aligning, and with one curve ending just short of the other.

– Location fidelity: Edge segments are forced to provide some minimum fidelity in their description of the curve location. This insures some minimum level of accuracy in how the correspondence constraints derived from these segments trace out the underlying edge feature. Above this minimum, the location fidelity of the edge segments also helps select the placement of their endpoints.

– Minimal description: Long edge segments are preferred over many short ones, in order to ultimately provide a concise set of correspondence constraints. An absolute minimum length is also used to remove excessively short segments.

The way that edge segments are used in correspondence estimation is described in the next section.

3.3 Correspondence constraints: matching edge segments to oriented cartoons

Correspondance constraints are defined by an edge segment in one image and a transformation into the other image. The result of this transformation is a new, or virtual, edge segment in the second image. Figure 5 provides an example of some edge segments and their corresponding virtual segments which were found using the approach described here.

The correspondance estimation procedure needs to find the best possible set of transformations so that edges in one image are transformed into the similar features in the other. Each segment can be transformed along four degrees of freedom. The segment can translate in two dimensions, rotate, and scale.

Correspondance constraints are found by searching through all possible transformations of all edge segments in each image. Each edge segment will have one transformation that will provide the best fit between its underlying feature points in its originating image and the features in the opposite image. Then, the edge segment with the strongest best-fit transformation is frozen, anchoring that edge segment and its pair (the virtual segment) in each image. The best-fit transformations of the remaining edge segments are reevaluated and the process is repeated on the unanchored (or floating) segments. At each iteration one more edge segment and its virtual-segment pair are anchored in the images. The process continues until all edge segments with valid best-fit transformations are anchored.

To simplify the explanation, consider the example shown in Figure 6. The cartoon-amplitude images show the strength of the oriented cartoon features in the two images. The segment maps show the locations of the points which are grouped into edge segments. Each edge-segment location refers to the same point in the underlying oriented cartoon for its orientation and strength estimates.

Transformations are evaluated based on the strength of the resulting match at the edge segment and based on a bias term from neighboring segments. The isolated–segment match will be discussed first and the bias term will be consided in the next paragraph. The strength of the isolated–segment match is the strength of the match between feature points underlying the edge segment in the originating image and the features in the opposite image. For example, in Figure 6, the best transformation for line segment *s13* is chosen by evaluating, for each of the allowed transformations, the correlation between the transformed salience/orientation values underlying *s13* and the salience/orientation values at the transformed locations in oriented cartoon *oc2*.[6] The transformed salience/orientation values underlying *s13* are provided by the oriented cartoon *oc1*. For each point location underlying *s13*, the salience value is taken directly from the oriented cartoon *oc1* but its

6. The correlation between the oriented cartoon points can not use a simple vector dot product since the points in the oriented cartoon with negative salience (on the bright side of a step) are fundamentally different from points in the oriented cartoon with positive salience (on the dark side of a step) and the opposite orientation vector. Instead, a “signed-amplitude vector product” is defined as:

$$(G_1, \theta_1) \bullet (G_2, \theta_2) = \min((G_1 G_2), (G_1 G_2 \cos(\theta_1 - \theta_2)))$$

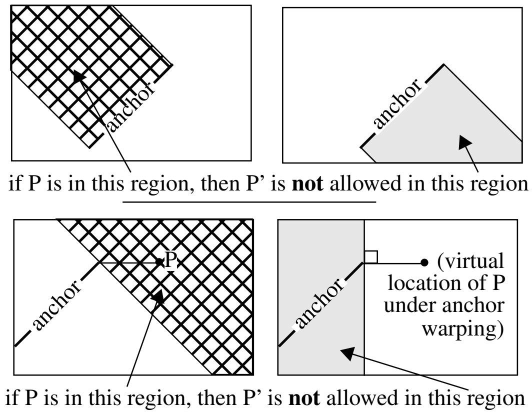

Figure 7: Examples of disallowed geometric warpings. Two classes of warpings which are not allowed are shown here. Two images have been warped as described in Section 2 to align the anchored constraint segments. When point P is in the cross-hashed area, then its transformation pair P' is **not** allowed in the dotted area.

orientation is changed to reflect the rotation included in the current transformation. The best transformation for line segment *s13* is the one which results in the highest correlation between these transformed salience/orientation values and the salience/orientation values from oriented cartoon *oc2*.

Image warpings are smooth and well behaved, with discontinuities across a limited number of edges. The bias term in the match strength encourages coherent image warpings at neighboring segments. When evaluating the match strength assigned to a line–segment transformation, the neighboring edge segments are transformed with the same transformation to check their fit. Anchored edge segments add a mismatch term equal to an elastic force between their frozen location and their location under the trial transformation. Floating edge segments add a term equal to their own isolated–segment match strength using the trial transformation. In summary, the isolated–segment match from a segment/transformation pair is biased by the contributions from the anchored and floating segments. The weighting of these bias term changes the stiffness of connections between the edge segments.

Combinations of constraints which "tear" the image geometry should also be avoided. For example, in Figure 6, if segment *s13* is anchored at the vertical feature in oriented cartoon *oc2*, then the virtual segment for *s11* in *oc2* should not be allowed to move to the right of the virtual segment for *s13*. Similarly, since segment *s24* is to the left of the virtual segment for *s13*, the virtual segment for *s24* in *oc1* should not be allowed to move to the right of segment *s13*. The types of transformations which an anchored segment can prevent are shown in Figure 7. Both endpoints for the current segment and both endpoints for the segment under the trail transformation are considered in the role of point P. This insures that the current segment is not assigned a transformation which is geometrically inconsistent with the previously anchored segments.

4. RESULTS

To demonstrate the capabilities and the types of errors of the feature-based match estimation described in Section 3, this section considers the matching of portrait images.[7] This type of image matching is particularly useful for extending databases for face recognition[12]. Figure 8 shows an example of matching faces across race, gender and facial expression, but under similar orientation. In this figure, I show:

– the two original images with their salience images;

– the correspondence constraints, as found by match estimation (without manual intervention);

– the halfway morph using the automatic estimates for the correspondance constraints (no manual intervention); and

– the halfway morph after correcting the indicated number of correspondence constraints (some manual intervention).

The halfway point in the morph sequences was chosen since it shows the errors in the correspondence constraints most clearly.

The image morphed in Figure 8 using uncorrected constraints can be compared directly to earlier results, as shown in Figure 9. The errors near the eyes and the mouth are different but equally severe in Figure 8 (using the uncorrected constraints) and Figure 9. There are additional errors in Figure 9 distributed across many of the constraints. This spatially diffuse error may be due to the multiscale matching procedure used in [7]: an error in matching at a coarse scale will often result in widely distributed errors in the final match estimate. Because of this diffuse error, the interpolation results using the approach described here are preferable to those reported in [7], even using the unmodified constraints.

The value of the method reported here is even higher with manual correction. Only a few constraint lines needed to be modified to correct the errors in Figure 8. The last morphed image in Figure 8 was generated by correcting ten constraint lines: four and three around the left and right eyes, respectively, and three near the mouth. This contrasts sharply with what would be required to correct the errors in Figure 9. More than twice that number near the eyes and mouth would need to be changed in Figure 9 to correct the errors in those regions alone. To correct the diffuse error in Figure 9 would require modification of an order of magnitude more (nearly every constraint line).

5. CONCLUSIONS

This paper addresses the problem of matching distinct but related objects across images, for applications such as view-based model capture, virtual presence, and morphing special effects. A new approach to match estimation was presented for estimating correspondances between distinct objects. The correspondances are estimated at a sparse set of features. This feature based approach allowed us to align the images at stable, distinctive points on the objects. Image morphing using these constraints completes the alignment process. This new approach succeeds at aligning faces without regard to lighting, race or gender. Difference in facial expressions have proven more difficult.

Sparsely estimated alignment constraints allow us to use feature-based matching. In addition, they have the advantage of simplifying and speeding manual correction of the constraints. For our target applications, some manual intervention between match estimation and image morphing is possible. Segment-based matching is particularly quick and simple to correct manually, since the number of separate constraints which must be manipulated is minimized.

My future research in this area will center on using active contours ("snakes") in place of individual segments; on modifying the morphing algorithm to accommodate regions of occlusion and disocclusion; and on model-based matching for specific types of objects, such as faces.

---

7. Due to length constraints, I only include one example image pair. More examples will be included in the parallel technical report.

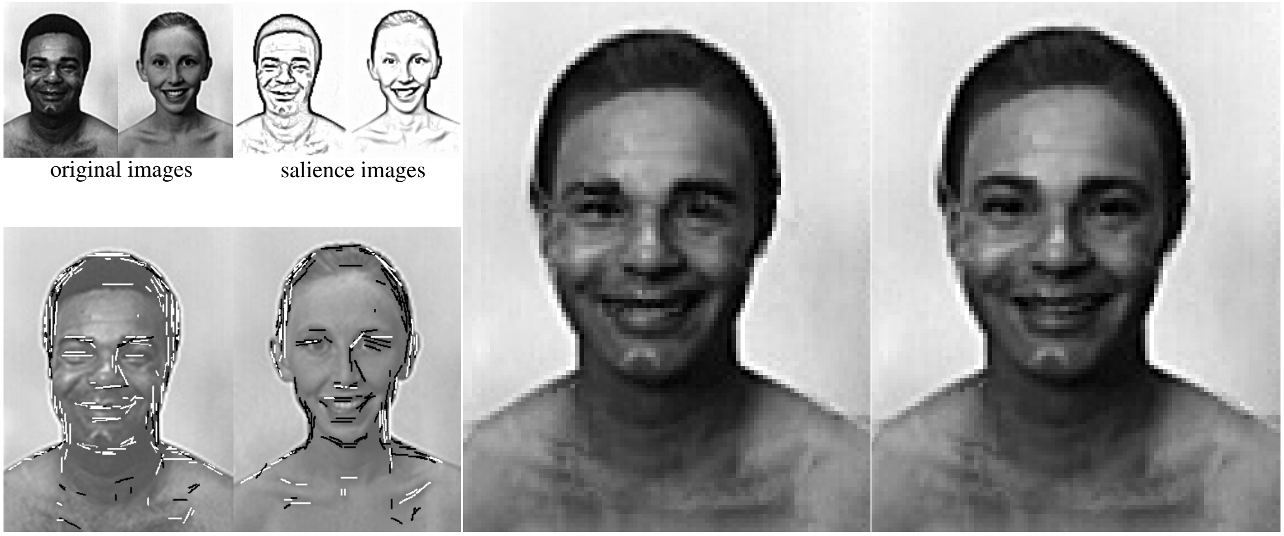

original images salience images

locations of uncorrected correspondence constraint segments (shown in white)

morph using uncorrected constraints

morph image using corrected constraints (number corrected = 10)

Figure 8: Match estimation and morphing across race, gender and facial expression under similar pose. The corrections which were made to the correspondence constraints were: four and three constraint lines at the left and right eyes, respectively and three constraint lines at the mouth. These are locations where the differences in the facial expressions are the most pronounced.

REFERENCES

[1] M. Luettgen, W. Karl, A. Willsky, "Efficient Multiscale Regularization with Applications to the Computation of Optical Flow," *IEEE Trans Image Processing*, vol. 3, pp. 41–64, 1994.

[2] P. Anandan, "A Computational Framework and an Algorithm for the Measurement of Visual Motion," *International Journal of Computer Vision*, vol. 2, pp. 283–310, 1989.

[3] R. Zabih, J. Woodfill, "Non-parametric local transforms for computing visual correspondence," Proc. European Conference on Computer Vision, pp. 151–158, 1994.

[4] B. Bhanu, O. Faugeras, "Shape Matching of Two-Dimensional Objects," IEEE Trans. PAMI, vol 6, pp. 137–156, 1984.

[5] D. Huttenlocher, J. Noh, W. Rucklidge, *Tracking Non-Rigid Objects in Complex Scenes*, Cornell University TR 92-1320, 1992.

[6] T. Beier, S. Neely. "Feature-Based Image Metamorphosis," *Proc. SIGGRAPH'92*, pp. 35-42, 1992.

[7] M. Covell, M. Withgott. "Spanning the gap between motion estimation and morphing," *Proc. IEEE ICASSP'94*, vol. 5, pp. 213–216, 1994.

[8] J. Nieweglowski, T. Campbell, P. Haavisto, "A novel video coding scheme based on temporal prediction using digital image warping," IEEE Trans. Consumer Electronics, vol. 39, pp. 141–150, 1993.

[9] M. Hall, public domain software *morphine, available* for anonymous ftp (run *archie* for the substring *morphine*), 1992.

[10] W. Freeman, E. Adelson, "Design and Use of Steerable Filters," IEEE Trans. PAMI, vol. 13, pp. 891–906, 1991.

[11] R. Duda, P. Hart, *Pattern Classification and Scene Analysis*, Wiley & Sons, 1973.

[12] D. Beymer, "Face Recognition from One Model View," to appear in *Proc. ICCV'95*, 1995.

Figure 9: Results of morphing using match estimation approach reported in [7]. The errors in constraint line placement using the approach reported in [7] tend to be more widely distributed across the constraints. In this example, the errors are not only concentrated at the eyes and mouth, as they were using the uncorrected constraints in Figure 8. They are also spread across the nose, the ears, and the silhouette of the head and shoulders. This diffuse error is difficult to correct manually.

ACKNOWLEDGEMENTS

Meg Withgott started me along this line of inquiry as well as helping with the initial research. Malcolm Slaney has spent many hours discussing both this paper and the underlying research. His comments have greatly clarified and improved the presentation. Many thanks!