

Local Geometric Indexing of High Resolution Data for Facial Reconstruction from Sparse Markers

Matthew Cong
Industrial Light & Magic
mcong@ilm.com

Lana Lan
Industrial Light & Magic
llan@ilm.com

Ronald Fedkiw
Stanford University
Industrial Light & Magic
rfedkiw@stanford.edu

Abstract

When considering sparse motion capture marker data, one typically struggles to balance its overfitting via a high dimensional blendshape system versus underfitting caused by smoothness constraints. With the current trend towards using more and more data, our aim is not to fit the motion capture markers with a parameterized (blendshape) model or to smoothly interpolate a surface through the marker positions, but rather to find an instance in the high resolution dataset that contains local geometry to fit each marker. Just as is true for typical machine learning applications, this approach benefits from a plethora of data, and thus we also consider augmenting the dataset via specially designed physical simulations that target the high resolution dataset such that the simulation output lies on the same so-called manifold as the data targeted.

1. Introduction

Realistic facial animation has a wide variety of applications in both computer vision and the entertainment industry [33]. It is typically achieved through a combination of keyframe animation, where an animator hand-adjusts controls corresponding to the motion of different parts of the face, and facial performance capture, which uses computer vision to track the motion of an actor’s face recorded from one or more cameras. Despite the many techniques developed over the years, facial performance capture remains a difficult task, and the high degree of accuracy required to generate realistic facial animation severely suppresses its widespread impact.

One class of techniques which has a proven track record uses markers painted on an actor’s face in conjunction with a stereo pair of head mounted cameras [6, 33]. These markers are tracked in each camera and triangulated to obtain a sparse set of animated 3D bundle positions representing the motion of the actor’s face. In order to reconstruct a

full 3D facial pose for each frame of 3D bundles, one often uses a parameterized (blendshape) model [21, 26]. However, these parameterized models often have large infeasible spaces. While a skilled animator can aim to avoid these infeasible spaces, an optimization algorithm would need them explicitly specified which is typically not practical. Another commonly used approach interpolates bundle displacements across the face [6]. However, this results in reconstructed geometry that is overly smooth since the sparse bundle positions cannot represent high-resolution details between the bundles, especially details that appear during expressions, e.g. folds, furrows, and wrinkles. To address these shortcomings, we follow the current trend in the deep learning community of adding more and more data by using a large dataset of facial shapes to inform the reconstruction of the face surface geometry from the tracked bundle positions.

Our approach to this problem can be thought of as local geometric indexing wherein each bundle needs to identify relevant associated geometry from the dataset. To accomplish this, we envision the dataset as a separate point cloud for each bundle; this point cloud is obtained by evaluating the 3D position of the relevant bundle on each face shape in the dataset. These point clouds are then used to index the dataset in order to figure out the most relevant shapes given a bundle position. A bundle position that lies outside of its associated point cloud indicates a lack of data and can be projected back towards the point cloud. On the other hand, it is also possible for many candidate points to exist in the point cloud in which case neighboring bundles and their associated point clouds can be used to disambiguate. Finally, the shapes chosen for each bundle are combined to obtain a high-resolution dense reconstruction of the facial geometry.

We begin the exposition by describing the creation of our facial shape dataset, which is initially bootstrapped via a combination of dense performance capture and hand sculpting for a small set of expressions and is further augmented using physical simulation. Then, we detail our local geometric indexing scheme and show how it can be used to

find the shapes that are most relevant to a bundle given its position. This is followed by a discussion of the various smoothness considerations that are used to inform our approach for spatially blending the relevant shapes across the face to recover a high-resolution dense reconstruction of the full face. Finally, we apply our algorithm to a series of feature film production examples and compare the results to other popular approaches.

2. Prior Work

Capture: High-resolution facial geometry can be captured using dense performance capture techniques such as [1, 3, 16, 17]. However, these methods typically require environments with controlled lighting and dedicated camera hardware. These restrictions, along with the limitations on the actor’s motion, often make these techniques unsuitable for on-set capture where an actor often needs to interact with the set and/or other actors. On-set capture typically involves painting a marker pattern on an actor’s face and recording the actor’s performance with a set of helmet mounted cameras. The markers can be tracked in the resulting footage and triangulated to recover a sparse set of bundle positions that follow the actor’s facial performance.

Reconstruction: In order to animate the neutral mesh of an actor, one could compute a smooth deformation of the face mesh by interpolating the motion of the bundles (see e.g. the corrective shape computed in [6]). However, this usually results in a deformed mesh that contains too much of the high-frequency detail of the neutral mesh and too little of the high-frequency detail associated with a particular expression. Several approaches have been proposed for augmenting the resulting deformed mesh with additional high frequency details including [5], [7], [8], [10], and [18] which use deformation gradients, a nonlinear shell energy, feature graph edge strains, polynomial displacement maps, and neural networks respectively. Masquerade [23] combines some of these approaches for facial performances solved from helmet mounted cameras. However, such approaches do not remove high-frequency details in the neutral mesh that are not present in the expression. Furthermore, if the smooth deformation interpolates the bundles, the addition of fine scale details in this manner can potentially move the surface farther away from the bundles.

Blendshapes: Instead of interpolating the motion of the bundles directly, one could use the bundles to drive a blendshape facial rig [21] which specifies the deformation of the face as a linear combination of facial shapes. These facial shapes are acquired using dense performance capture (see e.g. [1, 3, 16, 17]) and/or sculpted by an experienced modeler [13, 20]. Then, one can optimize for the shape weights that minimize the difference between the bundle positions and their associated surface positions on the resulting mesh

[6, 22]. However, such an approach often results in unnatural combinations of shapes with weights that are difficult to interpret [9]. These infeasible combinations can be avoided by experienced animators but are extremely problematic for optimization algorithms. In order for an optimization algorithm to avoid these combinations, one would need to specify all such invalid combinations in the high-dimensional Cartesian space of facial shapes, which is intractable.

Patch-Based Approaches: The patch-based model of [34] is particularly notable because it uses a smaller number of facial shapes compared to a traditional blendshape rig. Despite the small number of facial shapes, the resulting per-patch shape in this model still lies in the Cartesian product of the input shapes. Thus, as the size of the dataset increases, one would still expect the model to overfit on a per-patch basis. The FaceIK editing technique of [35] also uses a localized blendshape deformation model by adaptively segmenting the face mesh based on user specified control points, solving for blendshape weights for each control point based on its position, and spatially blending the resulting weights across the mesh using a radial basis function. In order to improve sparsity of the blendshape weights and reduce overfitting, blendshapes that are farther away from the control points are penalized. Unlike [35], which uses an interpolatory approach, our approach uses a non-manifold mesh and other considerations to boost the domain from \mathbb{R}^3 into higher dimensions. Other localized models have also been proposed such as [32] which uses PCA-based patches.

3. Dataset

Given the high-resolution mesh of an actor in the neutral or rest pose, we construct a dataset of high-quality facial shapes that sufficiently samples the actor’s range of motion and expression. We bootstrap this process by acquiring high-resolution facial geometry for a selection of the actor’s (extreme) facial poses taken from a range of motion exercise using the Medusa performance capture system [1, 2, 3]. For each facial pose, Medusa both deforms the neutral mesh to the pose based on images from multiple cameras and estimates the cranium rigid frame associated with the deformed mesh. The cranium rigid frame is manually refined (if necessary), validated against the images from each of the cameras, and then used to stabilize the associated deformed mesh. Each stabilized deformed mesh is then stored as a per-vertex displacement from the neutral mesh.

These stabilized facial shapes are further improved using physical simulation. Starting from the high-resolution neutral mesh, we build a simulatable anatomical face model by morphing an anatomically and biomechanically accurate template model following the approach of [11]. Then, we use the art-directed muscle simulation framework of [12] to

target each captured facial shape to obtain a corresponding simulated facial shape with improved volume conservation, more realistic stretching, and a more plausible response to contact and collision. The captured and simulated facial shape are then selectively blended together by a modeler to obtain a combined facial shape that incorporates both the high degree of detail obtained from capture as well as the physical accuracy obtained from simulation. Finally, this combined facial shape is further refined by a modeler based on the images in order to resolve any remaining artifacts before being added to the dataset. See [13, 20].

At this point, the dataset consists of facial shapes corresponding to various extreme poses. We augment the dataset with in-betweens to better represent subtle motions and combinations of expressions. To do this, one could construct a blendshape system using the facial shapes already in the dataset and evaluate this blendshape system at fixed intervals in the high-dimensional Cartesian space; however, the resulting in-betweens would suffer from well-known blendshape artifacts such as volume loss due to linear interpolation. Instead, one could use the aforementioned process targeting the high-dimensional Cartesian space blendshapes with the art-directed muscle simulation framework of [12], or alternatively one could use the approach of [12] alone to move between various extreme facial poses creating in-betweens. We utilize a combination of these options to add anatomically motivated nonlinear in-betweens to the dataset.

4. Local Geometric Indexing

Our local geometric indexing scheme begins by constructing a separate point cloud for each bundle, accomplished by evaluating the surface position of the bundle on each facial shape in the dataset. The brute force version of our algorithm would tetrahedralize each point cloud with all possible combinations of four points resulting in a non-manifold tetrahedralized volume (See Sec. 4.1). Then, given a bundle position, we find all the tetrahedra in the associated tetrahedralized volume that contain it. Since the tetrahedralized volumes are only dependent on the dataset, this process can be accelerated by precomputing a uniform grid spatial acceleration structure [15, 27]. For each of these tetrahedra, we compute the convex barycentric weights (w_i, w_j, w_k, w_l) of the bundle position and use these to blend together the four facial shapes $\vec{b}_i, \vec{b}_j, \vec{b}_k,$ and \vec{b}_l corresponding to the vertices of the tetrahedron. The resulting candidate shape is given by

$$\vec{x} = \vec{x}_0 + \sum_{n \in \{i,j,k,l\}} w_n \vec{b}_n \quad (1)$$

where \vec{x}_0 represents the neutral mesh positions. By construction, the candidate surface geometry is guaranteed to

intersect the bundle position and lie within the convex hull of the facial shapes.

If there are no tetrahedra that contain the bundle position, we project the bundle position to the convex hull of the associated point cloud by using the barycentric coordinates for the closest point on the associated tetrahedralized volume. The lack of tetrahedra containing a bundle position indicates a need for additional facial shapes in the dataset; however, this projection approach gives reasonable results in such scenarios.

Local geometric indexing can be viewed as a piecewise linear blendshape system, although the pieces are difficult to describe due to overlapping non-manifold tetrahedra and various nonlinear, nonlocal, and higher dimensional strategies for choosing between multiple overlapping tetrahedra. Still, by augmenting the dataset with more in-betweens, we can insert so-called Steiner points [4] allowing for increased efficacy – stressing the importance of collecting more and more data.

4.1. Tetrahedralization

As the size of the point cloud increases, the construction of all possible tetrahedra quickly becomes unwieldy. Thus, we aggressively prune redundancies from the point cloud, e.g. removing points corresponding to expressions that do not involve them. For example, we do not add bundle evaluations to a forehead bundle’s point cloud from expressions that only involve the lower half of the face. Besides reducing the number of points, we may also eliminate tetrahedra especially those that are poorly shaped: too thin, too much spatial extent, etc. Moreover, tetrahedra which are known to be problematic admitting shapes that are locally off-model can also be deleted. Similar statements hold for unused or rarely used tetrahedra, etc. Importantly, through continued use and statistical analysis, our tetrahedral database can evolve for increased efficiency and quality.

Instead of considering all possible combinations of four points, one could tetrahedralize each point cloud using a space-filling tetrahedralization algorithm such as constrained Delaunay tetrahedralization [29]. However, this would restrict a bundle position to lie uniquely within a single tetrahedron and create a bijection between a bundle position and local surface geometry. This is problematic because different expressions can map to the same bundle position with different local curvature. For example, a bundle along the midline of the face on the red lip margin can have the same position during both a smile and a frown. Thus, it is better to construct an overlapping non-manifold tetrahedralization in order to allow for multiple candidate local surface geometries for a bundle position, later disambiguating using additional criteria. Moreover, as discussed later, one may create more than one point cloud for an associated bundle with each point cloud corresponding to different cri-

teria. For example, the shapes one uses for an open jaw could differ significantly when comparing a yawn and an angry yell; different point clouds for sleepy, angry, happy, etc. would help to differentiate in such scenarios.

Again, we stress that a space-filling manifold tetrahedralized volume allows a bundle only three degrees of freedom as it moves through the manifold tetrahedralized volume in \mathbb{R}^3 , whereas overlapping non-manifold tetrahedra remove uniqueness in \mathbb{R}^3 boosting the domain to a higher dimensional space; then, other considerations may be used to ascertain information about other dimensions and select the appropriate tetrahedron.

5. Smoothness Considerations

Our local geometric indexing scheme generates local surface geometry for each bundle independently, and we subsequently sew the local surface geometry together to create a unified reconstruction of the full face. Because only local geometry is required, we only need to store small surface patches (and not the full face geometry) for each point in the point cloud making the method more scalable. To sew the local patches together, we first construct a Voronoi diagram on the neutral mesh using the geodesic distances to the surface position of each bundle in the rest pose. See Figure 1 (Top Left). These geodesic distances are computed using the fast marching method [19]. The local surface geometry for each bundle could then be applied to its associated Voronoi cell on the mesh, although the resulting face shape would typically have discontinuities across Voronoi cell boundaries as shown in Figure 1 (Top Right).

We have experimented with a number of scattered interpolation methods aimed at smoothing the local patches across Voronoi cell faces including, for example, radial basis functions as in [34]. We experimentally achieved the best results leveraging our Voronoi diagram using natural neighbor interpolation [25, 30]. For a given vertex on the neutral mesh, natural neighbor weights are computed by inserting the vertex into the precomputed Voronoi diagram, computing the areas stolen by the new vertex’s Voronoi cell from each of the pre-existing neighboring Voronoi cells, and normalizing by the total stolen area. For each vertex, the natural neighbor weights are used to linearly blend the shapes used for each surrounding bundle. Note that a vertex placed at a bundle position would not change the Voronoi regions of surrounding bundles and would merely adopt the Voronoi region from the bundle it is coincident with; this guarantees that the resulting blended surface still exactly interpolates the bundle positions. In this way, we obtain a C^0 continuous reconstructed surface [28] that passes through all of the bundle positions. See Figure 1 (Bottom Right). We found that constructing the Voronoi diagram and calculating the natural neighbor weights in UV/texture space and subsequently mapping them back onto the 3D mesh




Figure 1. Top Left: Voronoi diagram on the neutral mesh. Top Right: Applying the locally indexed surface geometry to each Voronoi cell results in discontinuities across cell boundaries. Bottom Left: Natural neighbor weights for a single bundle. The weight is 1 at the bundle surface position and 0 at the surface positions corresponding to neighboring bundles. Bottom Right: Using natural neighbor weights, we obtain a smoother reconstruction that interpolates the bundle positions.

yielded smoother natural neighbor weights than performing the equivalent operations on the 3D mesh directly.

5.1. Choosing Tetrahedra

In order to minimize kinks in the C^0 continuous reconstructed surface, we use an additional smoothness criterion when choosing between overlapping tetrahedra. If there are multiple tetrahedra which contain the bundle position, we choose the tetrahedron that results in local surface geometry that minimizes the distances from neighboring bundle positions to their respective surface positions. This indicates that the local surface geometry is representative of the bundle as well as the neighborhood between the bundle and its neighboring bundles.

In the case where no tetrahedra contain the bundle position, one can apply a similar criterion to project the bundle back to the dataset in a smooth manner. When deciding which tetrahedron to project to, one could consider not only the distance from the bundle under consideration to the re-

sulting surface, but also the distances that neighboring bundles would be from the resulting surface.

In the case of an animated bundle with time-varying position, we apply additional criteria to prevent disjoint sets of shapes from being chosen in neighboring frames, ameliorating undesirable oscillations in the animated reconstructed surface. To do this, we assign higher priority to tetrahedra which share more points and therefore facial shapes with the tetrahedron used on the previous frame, biasing towards a continuous so-called winding number on the non-manifold representation.

6. Jaw Articulation

So far, we have considered facial shapes and bundle positions relative to the neutral mesh. However, these shapes and bundle positions may include displacements due to rotational and prismatic jaw motion [31, 36]. This can result in significant linearized rotation artifacts in the reconstruction which reduces the generalizability of our approach. In order to address this, we hybridize our approach by using linear blend skinning to account for the jaw pose.

To do this, we modify Eq. 1 with a block diagonal matrix of spatially varying invertible transformations $T(\theta)$ calculated using linear blend skinning from the jaw parameters θ and a set of unskinned facial shapes \vec{b}_n^* to obtain

$$\vec{x} = T(\theta) \left(\vec{x}_0 + \sum_{n \in \{i,j,k,l\}} w_n \vec{b}_n^* \right). \quad (2)$$

For a shape with known jaw parameters θ_n , setting Eq. 1 equal to Eq. 2 and rearranging terms gives an expression for the unskinned facial shape

$$\vec{b}_n^* = T(\theta_n)^{-1} \left(\vec{x}_0 + \vec{b}_n \right) - \vec{x}_0$$

as a function of the facial shape \vec{b}_n . See [21, 24]. In order to utilize this approach, every shape in the database needs the jaw parameters θ_n estimated so that we may store \vec{b}_n^* instead of \vec{b}_n . Similarly for each frame, θ must be estimated using one of the usual methods for head and jaw tracking so that the bundle positions can be unskinned before indexing into the point cloud.

As mentioned in Sec. 4.1, having a large number of points can result in an unwieldy number of tetrahedra. Thus, one could bin points into different point clouds based on a partition computed using the jaw parameters θ ; each point cloud would only contain a range of jaw parameters and would therefore be smaller. Moreover, it makes more sense to interpolate between shapes with similar jaw parameters as opposed to significantly different jaw parameters. One should likely still unskin all of the shapes in the point cloud to have the same jaw parameter value for better efficacy;

however, choosing a non-neutral reference shape for the unskinning (e.g. in the middle of the relevant jaw parameter range) could be wise.

7. Experiments

Verification: In order to verify our algorithm, we calculated a set of 3D bundles for each facial shape in our dataset by evaluating the surface position of each bundle on the facial shape. Then, we inputted each set of bundle positions into our local geometric indexing algorithm, and verified that the resulting reconstruction is nearly identical to the original facial shape. See Figure 2.

High-Resolution Capture Comparison: Next, we evaluate our algorithm on a high-resolution performance outputted from the Medusa performance capture system [1, 2, 3]. The jaw is tracked using the lower teeth during the portions of the performance where they are visible and interpolated to the rest of the performance using the chin bundles as a guide. Like the previous experiment, we calculate a set of 3D bundles for each frame of the performance and use this animated set of 3D bundles as input into our local geometric indexing algorithm. The resulting high-resolution reconstruction of the performance using our dataset is very similar to the original performance. See Figure 3. The differences in the mouth corners and lips are due to artifacts in the Medusa performance capture. By indexing the most relevant cleaned up shapes in our dataset, we obtain a cleaner reconstruction while also adding detail sculpted by a modeler such as lip wrinkles. Other differences, such as those on the forehead and side of the face, occur because there are no bundles in those locations and thus our algorithm extrapolates from the nearest bundle.

Comparison to Other Approaches: In Figure 4, we compare our approach to other popular approaches on a




Figure 2. In order to verify our approach, we input 3D bundle positions from each shape in our library into our local geometric indexing algorithm; the results obtained are nearly identical to the original shapes. Top: Scan. Bottom: Local geometric indexing. A video showing the results on a larger dataset is available in the supplementary material.




Figure 3. Top: A high-resolution facial performance processed using the Medusa performance capture system [1, 2, 3]. Bottom: Reconstruction obtained using local geometric indexing driven by the bundle positions on the captured geometry. None of the high-resolution facial shapes in the performance were included in the dataset used by our algorithm. A number of the differences, such as those in the mouth corners and eyebrows, are actually due to artifacts in the Medusa performance capture geometry that are cleaned up by our reconstruction indicating that our approach provides some degree of regularization. The remaining differences are outside of the region spanned by the bundles where we would expect less accuracy due to limited data. A video showing the results on the entire performance is available in the supplementary material.

performance captured using two vertically stacked helmet mounted fisheye cameras. Footage from the top camera placed at nose level is shown in Figure 4 (Far Left). The images from both cameras are undistorted and the cameras are calibrated using the markers on the helmet. The calibrated cameras are used to triangulate bundle positions which are then rigidly aligned to the neutral mesh using a combination of the bundles on the nose bridge, forehead, and the cheeks with varying weights based on the amount of non-rigid motion in those regions. The jaw is tracked in the same manner as the previous experiment. As shown in Figure 4 (Middle Left), interpolating the bundle displacements across the mesh using [6] reconstructs a yawn instead of the angry face in the corresponding helmet mounted camera footage because it does not contain any additional high-resolution detail beyond that of the neutral mesh. Since the neutral mesh represents one’s face while expressionless, similar to that when asleep, using the displacements of the neutral mesh and its features often leads to expressions that appear tired. In order to obtain Figure 4 (Middle), we first constructed a blendshape rig using the facial shapes in our dataset. Then, we solved for the blendshape weights that minimize the Euclidean distances from the bundles to their relevant surface points subject to a soft constraint that penalizes the weights to lie between 0 and 1. The result incorporates more high-resolution details than Figure 4 (Middle) but suffers from overfitting resulting in severe artifacts around the mouth and eyes. Even though the resulting weights lie between 0 and 1, they are neither convex nor sparse which leads to unnatural combinations. Of course, increased regularization would smooth the artifacts shown in the figure creating a result that looks more like Figure 4 (Middle Left). In comparison, the

reconstruction obtained using our local geometric indexing algorithm shown in Figure 4 (Far Right) captures many of the high-resolution details that are not present in the neutral mesh including the deepened nasolabial folds, jowl wrinkles, and lip stretching without the overfitting artifacts of Figure 4 (Middle).

RBF Interpolation: Alternatively, instead of using natural neighbor interpolation, one could use radial basis functions to smooth with our local geometric indexing algorithm. As long as the radial basis function is applied on the facial shape weights as opposed to the vertex positions themselves, this still yields high-resolution features from the dataset in the reconstruction; however, the reconstructed surface will typically not pass through the bundles. This can be corrected by smoothly interpolating the remaining displacements needed to properly interpolate the bundles across the mesh (with e.g. [6]). As shown in Figure 4 (Middle Right), the reconstruction obtained using a combination of radial basis function interpolation and a smoothly interpolated deformation has a higher degree of detail than smoothly interpolating the deformation from neutral mesh. In a similar manner, additional rotoscoped constraints such as lip occlusion contours [6, 14], markers visible in only a single camera, etc. can be incorporated as a postprocessing step on top of our approach; in fact, we utilized [6] to incorporate lip occlusion contours in Figure 4.

Temporal Smoothing: Figure 5 demonstrates the ability of our method to capture subtle expressions while also maintaining temporal coherency in the presence of bundle positions with random and systematic errors (e.g. errors in depth due to the limited parallax between the two cameras).




Figure 4. Far Left: Helmet mounted camera footage. Middle Left: The reconstruction obtained by interpolating the bundle displacements across the mesh using [6] conveys a yawn as opposed to the anger/tension because it does not utilize any additional high-resolution detail beyond that of the neutral mesh. Middle: The typical overfitting symptomatic of blendshape rigs; with enough regularization, one would expect the detail to fade similar to the result using [6]. Middle Right: Using Gaussian RBF interpolation instead of natural neighbor interpolation in our approach results in additional high-resolution detail but does not interpolate the bundle positions. Far Right: Our approach passes through the bundles, conveys the expression, and captures high-resolution details that are not present in the neutral mesh.

If necessary, one can obtain a smoother performance by either temporally smoothing the input bundle positions or smoothing the barycentric weights on each bundle. In this performance, we apply temporal smoothing by taking a central moving average of the barycentric weights associated with each bundle relative to the jaw skinned neutral mesh in order to avoid smoothing the jaw animation. Because transitions between different sets of shapes typically occur when the same bundle position is achievable using multiple tetrahedra, we found this straightforward temporal smoothing scheme to have negligible impact on the ability for the reconstruction to interpolate the bundles.

Data Augmentation via Simulation: Figure 5 also illustrates the efficacy of augmenting the dataset using the art-directed muscle simulation framework of [12]. Figure 5 (Middle Left) was the result obtained without augmenting

and Figure 5 (Middle Right) was the improved result obtained by adding a number of new facial shapes via [12] as outlined in Section 3.

Generating/Correcting Rigs: Our local geometric indexing algorithm can also be used to generate actor-specific facial rigs. Given a generic template blendshape rig applied to the actor neutral mesh, we evaluate bundle positions for individual blendshapes and use these bundle positions as input into our local geometric indexing algorithm to reconstruct corresponding actor-specific blendshapes. We apply the same approach to combinations of blendshapes in order to obtain corresponding actor-specific corrective shapes [6] that do not exhibit the artifacts commonly found in combinations of blendshapes. See Figure 7. These actor-specific blendshapes and corrective shapes can be incorporated into an actor-specific nonlinear blendshape facial rig for use in




Figure 5. Far Left: Helmet mounted camera footage. Middle Left: The reconstruction obtained using our approach captures a subtle expression in the helmet mounted camera footage. This performance also shows the effectiveness of our temporal smoothness constraints. See video in supplementary material. Middle Right: Adding simulated in-betweens allows us to improve the smoothness of the reconstruction in the philtrum and the right jowl while also improving the lift in the upper right cheek. Far Right: Heatmap highlighting the differences between (Middle Left) and (Middle Right).




Figure 6. Top: Helmet mounted camera footage. Bottom: Reconstructions obtained using our method.

keyframe animation and other facial capture applications.

Usability: Our local geometric indexing calculations can be performed independently for each bundle and as expected, our parallel CPU implementation using Intel Threading Building Blocks scales linearly. Given this degree of parallelism in the local geometric indexing scheme as well as the GPU implementation of natural neighbor interpolation demonstrated in [25], our algorithm has the potential to run at interactive rates on the GPU. Already, our




Figure 7. Left: The combination of sneer, snarl, and upper lip raiser blendshapes leads to severe pinching artifacts in the cheeks and excessive deformation in the nose. These blendshapes are often used in conjunction to animate an angry face. Right: Reconstruction obtained using our local geometric indexing algorithm with the bundles calculated from (Left) as input. The reconstruction fixes the aforementioned artifacts in the cheek and nose, improves the shape of the upper lip, and preserves the emotional intent associated with each of the individual blendshapes.

approach is general and efficient enough to have been incorporated for use in the production of a major feature film. It has been tested on a wide range of production examples by a number of different users with significant creative and technical evaluation on the results. We show a small selection of our test results in Figure 6.

8. Conclusion

We have presented a data-driven approach for high-resolution facial reconstruction from sparse marker data. Instead of fitting a parameterized (blendshape) model to the input data or smoothly interpolating a surface displacement to the marker positions, we use a local geometric indexing scheme to identify the most relevant shapes from our dataset for each bundle using a variety of different criteria. This yields local surface geometry for each bundle that is then combined to obtain a high-resolution facial reconstruction.

We have applied our method to real-world production helmet mounted camera footage to obtain high-quality reconstructions. Rotoscoped features, including lip occlusion contours, can be readily incorporated as a postprocess. Finally, our approach has already been deployed for use in a film production pipeline for a major feature film where it has been leveraged by many users to obtain production quality results.

References

- [1] T. Beeler, B. Bickel, P. Beardsley, B. Sumner, and M. Gross. High-quality single-shot capture of facial geometry. In *ACM SIGGRAPH 2010 Papers*, SIGGRAPH '10, pages 40:1–40:9, New York, NY, USA, 2010. ACM. 2, 5, 6

- [2] T. Beeler and D. Bradley. Rigid stabilization of facial expressions. *ACM Trans. Graph.*, 33(4):44:1–44:9, July 2014. [2](#), [5](#), [6](#)
- [3] T. Beeler, F. Hahn, D. Bradley, B. Bickel, P. Beardsley, C. Gotsman, R. W. Sumner, and M. Gross. High-quality passive facial performance capture using anchor frames. In *ACM SIGGRAPH 2011 Papers*, SIGGRAPH '11, pages 75:1–75:10, New York, NY, USA, 2011. ACM. [2](#), [5](#), [6](#)
- [4] M. d. Berg, O. Cheong, M. v. Kreveld, and M. Overmars. *Computational Geometry: Algorithms and Applications*. Springer-Verlag TELOS, Santa Clara, CA, USA, third edition, 2008. [3](#)
- [5] A. H. Bermanto, D. Bradley, T. Beeler, F. Zund, D. Nowrouzezahrai, I. Baran, O. Sorkine-Hornung, H. Pfister, R. W. Sumner, B. Bickel, and M. Gross. Facial performance enhancement using dynamic shape space analysis. *ACM Trans. Graph.*, 33(2):13:1–13:12, Apr. 2014. [2](#)
- [6] K. S. Bhat, R. Goldenthal, Y. Ye, R. Mallet, and M. Koperwas. High fidelity facial animation capture and retargeting with contours. In *Proceedings of the 12th ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, SCA '13, pages 7–14, New York, NY, USA, 2013. ACM. [1](#), [2](#), [6](#), [7](#)
- [7] B. Bickel, M. Botsch, R. Angst, W. Matusik, M. Otaduy, H. Pfister, and M. Gross. Multi-scale capture of facial geometry and motion. In *ACM SIGGRAPH 2007 Papers*, SIGGRAPH '07, New York, NY, USA, 2007. ACM. [2](#)
- [8] B. Bickel, M. Lang, M. Botsch, M. A. Otaduy, and M. Gross. Pose-space animation and transfer of facial details. In *Proceedings of the 2008 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, SCA '08, pages 57–66, Aire-la-Ville, Switzerland, Switzerland, 2008. Eurographics Association. [2](#)
- [9] E. Chuang and C. Bregler. Performance driven facial animation using blendshape interpolation. Technical report, Stanford University, January 2002. [2](#)
- [10] W. chun Ma, A. Jones, J. yuan Chiang, T. Hawkins, S. Frederiksen, P. Peers, M. Vukovic, M. Ouhyoung, and P. Debevec. Facial performance synthesis using deformation-driven polynomial displacement maps. *ACM Trans. Graphics*, pages 121–1121, 2008. [2](#)
- [11] M. Cong, M. Bao, J. L. E, K. S. Bhat, and R. Fedkiw. Fully automatic generation of anatomical face simulation models. In *Proceedings of the 14th ACM SIGGRAPH / Eurographics Symposium on Computer Animation*, SCA '15, pages 175–183, New York, NY, USA, 2015. ACM. [2](#)
- [12] M. Cong, K. S. Bhat, and R. Fedkiw. Art-directed muscle simulation for high-end facial animation. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, SCA '16, pages 119–127, Goslar, Germany, 2016. Eurographics Association. [2](#), [3](#), [7](#)
- [13] M. Cong, L. Lan, and R. Fedkiw. Muscle simulation for facial animation in Kong: Skull Island. In *ACM SIGGRAPH 2017 Talks*, SIGGRAPH '17, pages 21:1–21:2, New York, NY, USA, 2017. ACM. [2](#), [3](#)
- [14] D. Dinev, T. Beeler, D. Bradley, M. Bächer, H. Xu, and L. Kavan. User-guided lip correction for facial performance capture. *Comput. Graph. Forum*, 37(8):93–101, 2018. [6](#)
- [15] A. Fujimoto, T. Tanaka, and K. Iwata. ARTS: Accelerated ray-tracing system. *IEEE Computer Graphics and Applications*, 6(4):16–26, April 1986. [3](#)
- [16] G. Fyffe, A. Jones, O. Alexander, R. Ichikari, and P. Debevec. Driving high-resolution facial scans with video performance capture. *ACM Trans. Graph.*, 34(1):8:1–8:14, Dec. 2014. [2](#)
- [17] A. Ghosh, G. Fyffe, B. Tunwattanapong, J. Busch, X. Yu, and P. Debevec. Multiview face capture using polarized spherical gradient illumination. In *Proceedings of the 2011 SIGGRAPH Asia Conference*, SA '11, pages 129:1–129:10, New York, NY, USA, 2011. ACM. [2](#)
- [18] L. Huynh, W. Chen, S. Saito, J. Xing, K. Nagano, A. Jones, P. Debevec, and H. L. Mesoscopic facial geometry inference using deep neural networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [2](#)
- [19] R. Kimmel and J. A. Sethian. Computing geodesic paths on manifolds. *Proceedings of the National Academy of Sciences*, 95(15):8431–8435, 1998. [4](#)
- [20] L. Lan, M. Cong, and R. Fedkiw. Lessons from the evolution of an anatomical facial muscle model. In *Proceedings of the ACM SIGGRAPH Digital Production Symposium*, DigiPro '17, pages 11:1–11:3, New York, NY, USA, 2017. ACM. [2](#), [3](#)
- [21] J. P. Lewis, K. Anjyo, T. Rhee, M. Zhang, F. Pighin, and Z. Deng. Practice and theory of blendshape facial models. In S. Lefebvre and M. Spagnuolo, editors, *Eurographics 2014 - State of the Art Reports*. The Eurographics Association, 2014. [1](#), [2](#), [5](#)
- [22] H. Li, T. Weise, and M. Pauly. Example-based facial rigging. In *ACM SIGGRAPH 2010 Papers*, SIGGRAPH '10, pages 32:1–32:6, New York, NY, USA, 2010. ACM. [2](#)
- [23] L. Moser, D. Hendler, and D. Roble. Masquerade: Fine-scale details for head-mounted camera motion capture data. In *ACM SIGGRAPH 2017 Talks*, SIGGRAPH '17, pages 18:1–18:2, New York, NY, USA, 2017. ACM. [2](#)
- [24] V. Orvalho, P. Bastos, F. Parke, B. Oliveira, and X. Alvarez. A Facial Rigging Survey. In M.-P. Cani and F. Ganovelli, editors, *Eurographics 2012 - State of the Art Reports*. The Eurographics Association, 2012. [5](#)
- [25] S. W. Park, L. Linsen, O. Kreylos, J. D. Owens, and B. Hamann. Discrete Sibson interpolation. *IEEE Transactions on Visualization and Computer Graphics*, 12(2):243–253, Mar. 2006. [4](#), [8](#)
- [26] F. I. Parke. *A Parametric Model for Human Faces*. PhD thesis, The University of Utah, 1974. [1](#)
- [27] M. Pharr and G. Humphreys. Primitives and intersection acceleration. In M. Pharr and G. Humphreys, editors, *Physically Based Rendering*, chapter 4, pages 182 – 258. Morgan Kaufmann, Boston, second edition, 2010. [3](#)
- [28] B. Piper. Properties of local coordinates based on Dirichlet tessellations. In G. Farin, H. Hagen, H. Noltemeier, and W. Knödel, editors, *Geometric Modelling*, pages 227–239. Springer-Verlag, London, UK, UK, 1993. [4](#)
- [29] J. R. Shewchuk. Constrained delaunay tetrahedralizations and provably good boundary recovery. In *Eleventh International Meshing Roundtable*, pages 193–204, 2002. [3](#)

- [30] R. Sibson. A vector identity for the Dirichlet tessellation. *Mathematical Proceedings of the Cambridge Philosophical Society*, 87(1):151155, 1980. [4](#)
- [31] E. Sifakis, I. Neverov, and R. Fedkiw. Automatic determination of facial muscle activations from sparse motion capture marker data. In *ACM SIGGRAPH 2005 Papers*, SIGGRAPH '05, pages 417–425, New York, NY, USA, 2005. ACM. [5](#)
- [32] J. R. Tena, F. De la Torre, and I. Matthews. Interactive region-based linear 3d face models. In *ACM SIGGRAPH 2011 Papers*, SIGGRAPH '11, pages 76:1–76:10, New York, NY, USA, 2011. ACM. [2](#)
- [33] L. Williams. Performance-driven facial animation. In *ACM SIGGRAPH 2006 Courses*, SIGGRAPH '06, New York, NY, USA, 2006. ACM. [1](#)
- [34] C. Wu, D. Bradley, M. Gross, and T. Beeler. An anatomically-constrained local deformation model for monocular face capture. *ACM Trans. Graph.*, 35(4):115:1–115:12, July 2016. [2](#), [4](#)
- [35] L. Zhang, N. Snavely, B. Curless, and S. M. Seitz. Spacetime faces: High resolution capture for modeling and animation. In *ACM SIGGRAPH 2004 Papers*, SIGGRAPH '04, pages 548–558, New York, NY, USA, 2004. ACM. [2](#)
- [36] G. Zoss, D. Bradley, P. Bérard, and T. Beeler. An empirical rig for jaw animation. *ACM Trans. Graph.*, 37(4):59:1–59:12, July 2018. [5](#)