Supplementary Material for
Computational Design of Custom-Fit PAP Masks

This supplementary material is composed of four parts. The first part provides details about computing the mask-face alignment for cushion evaluation. The second part provides details about formulating the optimization-based design. The third part provides implementation details of the optimization solver. The fourth part provides details and procedure of fabricating our designed custom-fit PAP mask.

1 Mask-Face Alignment

Since our custom-fit mask does not alter the mask frame, the aim of mask-face alignment is to position the mask reasonably as per its design. Hence, we align the scanned mask with the human face, mimicking how a person would wear it during sleep. In the cushion design, while the cushion shape can vary, the positions of the connector and the human face remain fixed.

One method to achieve proper alignment involves additionally 3D scanning the individual while they are wearing the mask, and aligning the separately 3D scanned face and mask to this position. However, as this can be time-consuming and labour-intensive, we propose a computational approach to approximate the alignment.

Mask global position. We fix the mask in the global coordinate and align it with any given human face. In the fixed position, the center of the mask interface bounding box is at the origin, with the mask cushion oriented toward the negative z-axis and the connector facing the positive z-axis. The mask's upward orientation corresponds to the y-axis. Therefore, when aligning the human face with the mask, it faces along the z-axis. The human head's upward orientation corresponds to the y-axis. The normal vector $\mathbf{n}$ to the human coronal plane is defined in the z-axis direction.

Alignment. For mask-face alignment, we first identify the human face symmetry plane, position it at the YZ plane, and find a 2D rigid transformation within this plane by minimizing the curve-face distance. We require two feature points on the human face, selected by the user: one at the nose tip $\mathbf{S}_1$, and the other on the nose bridge between the eyes $\mathbf{S}_2$.

We determine the parameters of the symmetry plane through an optimization process. The plane is parameterized by its normal direction and a point on the plane. For a given plane, we assess the optimization objective function as the Hausdorff distance between the human face surface and its reflection relative to the plane. We uniformly sample the parameters for optimization. Using the feature points, we search the normal direction around the vectors perpendicular to $\mathbf{S}_1 - \mathbf{S}_2$, and the plane's position around $\mathbf{S}_1$. After determining the human face's symmetry plane, we position it within the YZ plane, and perform rotation in the plane so that the upward orientation of the human face corresponds to the y-axis, and the forward orientation corresponds to the z-axis. The upward and forward orientation of the human face can be determined by the direction of the vector $\mathbf{S}_2 - \mathbf{S}_1$, i.e., examining its y-, z-values.

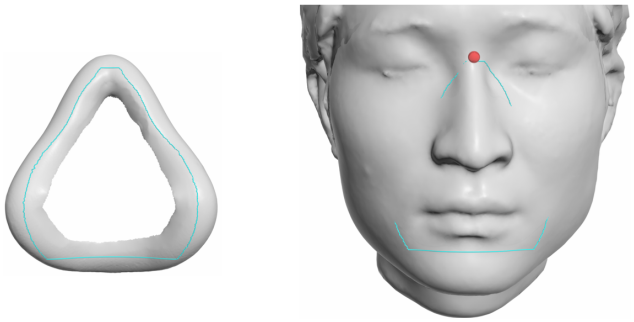

Figure 1: *(Left) Extract a centerline curve from the reconstructed BMC F2 mask cushion. (Right) Align the centerline curve with the face model, where the symmetry planes of the human face and the mask are aligned, and the highest point of the centerline is anchored at the midpoint (in red) of the subject's two eyes.*

The alignment problem now reduces to finding a 2D rigid transformation within the symmetry plane. We extract a centerline curve $\mathbf{q}_0$ from the reconstructed mask model for alignment; see Figure 1. We project the feature point $\mathbf{S}_2$ onto the symmetry plane along the plane normal, and anchor the projected point at the highest intersection point of the centerline curve and the symmetry plane. We then rotate the human face within the symmetry plane to minimize the curve-face distance, achieved by uniformly sampling the rotation angle. The curve-face distance is defined as $\max_{\mathbf{x}\in Q} |\phi(\mathbf{x})|$ where Q is the set of all the points in curve $\mathbf{q}_0$, and ϕ is the signed distance field of the human face. We compute the distance by using a sampled point set $Q_1 \subset Q$ of curve $\mathbf{q}_0$, i.e., $\max_{\mathbf{x}\in Q_1} |\phi(\mathbf{x})|$.

Under this alignment, due to the minimum distance of the curve and face, the cushion just touches the human face. Subjects should tighten the mask straps to create contact pressure across the cushion surface. For virtual mask-face fitting, we achieve this by translating the human face in its coronal plane normal direction and increasing interpenetration. The translation length is typically set to $\frac{1}{4}$ of the human face distance to the extracted curve $\mathbf{c}$ on the connector. This is sufficient to allow the cushion surface to penetrate into the human face across the trajectory curve.

2 Optimization Formulation

In this section, we provide more details of formulating the regularization term in the objective function in our optimization-based design (Section 5 of the paper).

Regularization. $E_{\text{regu}} = \delta_1 R_{\text{width}} + \delta_2 R_{\text{size}} + \delta_3 R_{\text{length}} + R_{\text{similarity}}$ regularizes the cushion surface shape by constraining the key profile curves' shape and length, as well as distance between adjacent key profile curves. Among these terms, R_{length} and $R_{\text{similarity}}$ are particularly important for maintaining a cushion-like shape for the swept surface, while the other two terms, R_{width} and R_{size}, are usually given much smaller weights. We provide detailed explainations on each component below.

1. *Maintaining cushion width.* Cushion width refers to the size of the cushion along the normal direction of the human face coronal plane $\mathbf{n}$. Since the local frame y-axis is exactly $\mathbf{n}$, we define

$$R_{\text{width}} = \Big(\sum_{i=1}^{K} \mathbf{b}_0^y(t_i)\Big)^2. \tag{1}$$

It can keep the cushion width within an appropriate range. The weight δ_1 typically set as 1×10^{-3}.

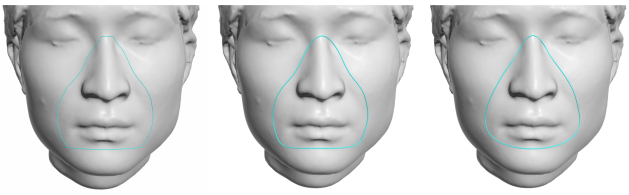

Figure 2: *Initializing the trajectory curve. (Left) Project the centerline curve of the BMC mask cushion onto the face model. (Middle) Approximate the projected centerline curve using a parametric curve. (Right) Adjust the trajectory curve to reduce curvature, enhance symmetry and maintain its distance to the mask connector.*

2. *Maintaining cushion size*. We control the overall size of th e cushion by restricting the change of the $\mathbf{b}_0^x(t_i)$ by term:

$$R_{\text{size}} = \sum_{i=1}^{K} (\mathbf{b}_0^x(t_i))^2. \tag{2}$$

The weight δ_2 typically set as 1×10^{-3}.

3. *Even distribution of profile curve lengths*. To ensure an even distribution of profile curve lengths, we define

$$R_{\text{length}} = \frac{1}{K} \sum_{i=1}^{K} (\mathbf{b}_2^{2\,x}(t_i) - \frac{1}{K} \sum_{j=1}^{K} \mathbf{b}_2^{2\,x}(t_j))^2 \tag{3}$$

to limit the variance of the x-values for the curve endpoint $\mathbf{b}_2^2(t_i)$. We use the endpoint x-value instead of its norm because this allows the profile curve to move more freely along $\mathbf{n}$. The weight δ_3 typically set as 10.

4. *Similar adjacent key profile curves*. To prevent sudden changes on the cushion surface, adjacent key profile curves should be similar. This is achieved by penalizing the parameter vector distance between adjacent key profile curves:

$$R_{\text{similarity}} = \frac{1}{K} \sum_{i=1}^{K} \|\mathbf{p}_i - \mathbf{p}_{i+1}\|^2, \tag{4}$$

where $\mathbf{p}_i$ is the parameters of ith key profile curve, and $\mathbf{p}_{K+1} = \mathbf{p}_1$.

3 Optimization Solver

To solve the optimization-based design problem, we initialize the cushion surface by modifying the BMC mask cushion to fit the static face model (Section 3.1). Then, we use gradient-based optimization with backtracking line search for optimizing the cushion surface. To this end, we address potential feasibility issues from the initialization, and formulate the barrier and penalty functions for handling constraints (Section 3.2).

3.1 Initialization

It is important to find a good initial cushion surface for optimization. The key idea of our initialization is to make the trajectory curve adapt to the human face, and create the profile curves based on the reconstructed BMC mask cushion surface. The initialization involves two steps: identifying the trajectory curve and determining the key profile curves.

Trajectory curve initialization. For the trajectory curve initialization, we project the extracted centerline curve of the BMC mask onto the human face along the coronal plane normal $\mathbf{n}$, and fit the projected polyline to obtain the initial trajectory curve; see Figure 2.

We make further adjustment on the trajectory curve to reduce curvature, enhance symmetry and maintain its distance to the mask connector, while keeping its shape close to the human face.To this end, we define and minimize the energy $E^q = \eta_1 E^q_{\text{comf}} + \eta_2 E^q_{\text{regu}}$, where E^q_{comf} approximates the comfort of the cushion, E^q_{regu} regularizes the trajectory curve, and η_1 and η_2 are the weights(typically set as 1000 and 0.01 in our experiments). The comfort energy is defined as $E^q_{\text{comf}} = \frac{1}{n_s} \sum_{i=1}^{n_s} (\phi(\mathbf{q}(x_i)) - \frac{1}{n_s} \sum_{i=1}^{n_s} \phi(\mathbf{q}(x_i)))^2$, where ϕ is the signed distance field of the human face, and $\{x_i\}_{i=1}^{n_s} \subset [0, 1)$ are sampled parameters(n_s is set as 500 in our experiments). We derive this metric to minimize curve-face distance while accommodating a wider range of potentially suitable trajectory curves. The regularization term $E^q_{\text{regu}} = \delta_1^q R^q_{\text{curvature}} + \delta_2^q R^q_{\text{width}} + \delta_3^q R^q_{\text{angle}} + \delta_4^q R^q_{\text{align}} + R^q_{\text{symmetry}}$, We provide details on each component below:

1. *Small curvature*. To reduce collisions among profile curves, we approximate the integral of squared curvature with:

$$R^q_{\text{curvature}} = \int_0^1 \|\mathbf{q}''(t)\|^2 \, dt. \tag{5}$$

Given that $\mathbf{q}''(t)$ is a linear spline, we can compute the integral piecewise using Gaussian quadrature. The weight δ_1^q is usually set as 5×10^{-4} in our experiments.

2. *Maintaining cushion width*. Based on the trajectory curve, we can measure the cushion width as the distance from the trajectory curve to the connector: $\sum_{i=1}^{n_s} \psi(\mathbf{q}(x_i))$, where ψ is denoted as the signed distance field of the mask connector. We measure the difference between the cushion width and the BMC mask cushion width as follows:

$$R^q_{\text{width}} = (\sum_{i=1}^{n_s} \psi(\mathbf{q}(x_i)) - \sum_{i=1}^{n_s} \psi(\mathbf{q}_0(x_i)))^2, \tag{6}$$

where $\mathbf{q}_0$ is the extracted centerline curve on the BMC mask cushion; see again Figure 1. The weight δ_2^q is usually set as 2×10^{-2} in our experiments.

3. *Large control polygon angle*. The upper section of the trajectory curve is prone to having a large curvature. So we define R^q_{angle} as the angle of the point at the very top of the control polygon of the initial trajectory curve:

$$R^q_{\text{angle}} = -\arccos\left(\frac{(\mathbf{d}_{k+1} - \mathbf{d}_k)^T}{\|\mathbf{d}_{k+1} - \mathbf{d}_k\|} \frac{\mathbf{d}_{k-1} - \mathbf{d}_k}{\|\mathbf{d}_{k-1} - \mathbf{d}_k\|}\right), \tag{7}$$

where $\mathbf{d}_k$ is the highest control point(i.e. having the largest y-value). The index k is determined by the initial trajectory curve for the optimization. The weight δ_3^q is usually set as 1.2×10^4 in our experiments.

4. *Aligned with human face*. In the mask-face alignment, the top of the curve is anchored at a feature point located on the nose bridge between the eyes. It is preferable for the optimized trajectory curve to remain close to this feature point. This also helps control the overall size of the cushion. In our experiment, without this regularization, the highest point of the trajectory curve may drop, resulting in a smaller overall size of the cushion. Therefore, we control the y-value of the highest point on the trajectory curve through regularization:

$$R^q_{\text{align}} = (\max_t\{\mathbf{q}^y(t)\} - \max_t\{\mathbf{q}_0^y(t)\})^2. \tag{8}$$

The weight δ_4^q is usually set as 1×10^{10} in our experiments.

5. *Symmetry*. We consider trajectory curve symmetry because the human face is nearly symmetric. To enable the curve to

better adapt to the human face, perfect symmetry of the curve is not required; instead, the symmetry of its projection in the human coronal plane is considered. Given our alignment, the symmetry plane should be YZ plane. So the symmetry can be measured by the distance between the curve $\mathbf{P}_n(\mathbf{q})$ and its reflection $\mathbf{T}(\mathbf{P}_n(\mathbf{q}))$ with respect to the YZ plane:

$$R^q_{\text{symmetry}} = \text{dist}(\mathbf{P}_n(\mathbf{q}), \mathbf{T}(\mathbf{P}_n(\mathbf{q}))). \tag{9}$$

There can be different measures of distance. In our experiment, we calculate the distance as: $\sum_{i=1}^{n_y}(I_i^{1\ x} + I_i^{2\ x})$, where I_i^1 and I_i^2 are two intersection points of the curve $\mathbf{P}_n(\mathbf{q})$ and a plane parallel to ZX plane: $y = y_i$. And $\{y_i\}_{i=1}^{n_y}$ are sampled in interval $[\min_t\{\mathbf{P}_n(\mathbf{q})^y(t)\},\ \max_t\{\mathbf{P}_n(\mathbf{q})^y(t)\}]$

We minimize E^q using the Nelder-Mead algorithm [Nelder and Mead 1965].

Profile curve initialization. For the key profile curve initialization, we first fit the reconstructed BMC mask cushion surface with our swept surface to obtain its key profile curves, set those curves as initial key profile curves, and move the control points $\{\mathbf{b}_2^1(t_i)\}$ onto the extracted curve on the connector in the local frame for the fixed cushion boundary constraint. To fit the BMC cushion surface, we fit its extracted centerline curve as the trajectory curve, obtain the intersection polylines of the BMC cushion surface and the local frame XY planes at $\{t_i\}$, and fit these polylines as the key profile curves.

3.2 Optimization

Addressing feasibility issues. The fixed cushion boundary constraint is already satisfied in the above initialization process. We must ensure that the initialized cushion surface satisfies both the height-field and collision-free constraints, as we use barrier functions for the two constraints.

We find our initial cushion surface generally satisfies the height-field constraint, since the inner part of the BMC mask cushion surface is a height field. As for the collision-free constraint, the initialized cushion surface may not meet the requirement. The reason is that after we offset the profile curve along its normal, self-intersection is introduced in the cushion shell. In this case, self-intersection can happen only if the offset length is larger than the minimum radius of curvature of the profile curve. Therefore, we can solve this problem by gradually reducing the curvature of the profile curve, until the maximum curvature is smaller than $\frac{1}{\tau}$ (τ is the offset length of the profile curves). Notice that our design parameter is the key profile curves. If a problematic profile curve is not among the key profile curves, we reduce the curvature of the two adjacent key profile curves. Since the profile curve is interpolated from the key profile curves, this action also reduces its curvature.

To reduce the key profile curve curvature, we adjust the control points to straighten the curve, and use different strategies on the Bézier curves $\mathbf{p}^1$ and $\mathbf{p}^2$. To straighten curve $\mathbf{p}^1$, we gradually move point $\mathbf{b}_1^1$ towards the midpoint of $\mathbf{b}_0$ and $\mathbf{b}_2^1$. Concurrently, curve $\mathbf{p}^2$ rotates to maintain the G^1 smoothness of the entire profile curve $\mathbf{p}$. To straighten curve $\mathbf{p}^2$, we rotate point $\mathbf{b}_2^2$ around the point $\mathbf{b}_1^2$.

Handling constraints. We convert the constraints to penalty terms and barrier terms and added to the objective function with weights for optimization. So the optimization is turned to an unconstrained problem, and the solution approaches to the solution of the constrained problem when the weights of the barrier terms approach 0 and the weight of the penalty term approaches $+\infty$. In our experiments, we fix the weights for simplicity.

1. *Convex profile curves*. We use quadratic penalty function for handling convexity constraints:

$$\sum_{i=1}^{K}(p_i^1 + p_i^2), \tag{10}$$

where $p_i^1 = (\min\{\det\left(\mathbf{b}_0(t_i) - \mathbf{b}_1^1(t_i), \mathbf{b}_2^1(t_i) - \mathbf{b}_1^1(t_i)\right), 0\})^2$, and $p_i^2 = (\min\{\det\left(\mathbf{b}_2^2(t_i) - \mathbf{b}_1^2(t_i), \mathbf{b}_0(t_i) - \mathbf{b}_1^2(t_i)\right), 0\})^2$. The weight for this term is set as 0.01 in our experiments.

2. *Collision-free tetrahedral mesh of the cushion shell*. Denote $\mathcal{T}$ as the set of all the tetrahedra in the tetrahedral mesh of the cushion shell, and $\{\mathbf{v}_i^e\}_{i=1}^4$ is the four vertices of a tetrahedron $e \in \mathcal{T}$. So the volume of e is $\mathrm{V}^e = \frac{1}{6}\det(\mathbf{v}_2^e - \mathbf{v}_1^e,\ \mathbf{v}_3^e - \mathbf{v}_1^e,\ \mathbf{v}_4^e - \mathbf{v}_1^e)$. Then this term is expressed by a barrier function that prevents the tetrahedral volume become negative:

$$\sum_{e\in\mathcal{T}} B(6\ \mathrm{V}^e, b^e), \tag{11}$$

where $B(x, \hat{x})$ is a barrier function. Here we use the barrier function proposed by [Li et al. 2020], as this barrier function vanishes when x exceeds a parameter $\hat{x}$:

$$B(x,\hat{x}) = \begin{cases} -(x-\hat{x})^2\ \log(\frac{x}{\hat{x}}), & if\ 0 < x < \hat{x}, \\ 0, & if\ x >= \hat{x}. \end{cases} \tag{12}$$

The parameter $b^e = 6\ \mathrm{V}^e$ is evaluated on the initial tetrahedral mesh of the cushion shell and determined at the beginning of the optimization. The weight for this term is set as 5×10^{-4} in our experiments.

3. *Height-filed cushion surface*. Given the mask-face alignment, this constraint can be achieved by making sure that the curve $\mathbf{p}_t^2$ increases monotonically along x-axis of the local frame $\mathbf{r}_1(t)$:

$$\mathbf{b}_2^{2\ x}(t) > \mathbf{b}_1^{2\ x}(t) > \mathbf{b}_0^x(t),\ t \in [0,1). \tag{13}$$

Notice that this is a constraint for every profile curves, we can deal with it by considering a number of sampled profile curves. In our experiment, since the cushion surface is interpolated from the key profile curves, it is enough to consider the constraint only on the key profile curves. We use the same barrier function in Equation 12 for defining this constraint:

$$\sum_{i=1}^{K}\Big\{B(\mathbf{b}_2^{2\ x}(t_i) - \mathbf{b}_1^{2\ x}(t_i), b_i) + B(\frac{\pi}{2} - \theta(t_i), \frac{\pi}{4}) + B(\theta(t_i) + \frac{\pi}{2}, \frac{\pi}{4})\Big\}, \tag{14}$$

where $b_i = \mathbf{b}_2^{2\ x}(t_i) - \mathbf{b}_1^{2\ x}(t_i)$ is evaluated on the initial cushion surface of the optimization. The weight for this term is set as 1 in our experiments.

4 Custom-fit PAP Mask Fabrication

Mold design for fabricating the mask interface. The fabrication material of the mask interface must be biomedically safe since it directly contacts human skin. We use a skin-safe silicone (Dragonskin 30, Smooth-On) for fabricating our designed mask interface. This material is commonly used in medical prosthetics and cushioning applications.

We fabricate the mask interface by casting silicone into a mold with rigid pieces. We model the mold with 3 pieces. Pieces 1, 2, and 3 are shown in red, green, and blue, respectively, in Figure 3. In broad terms, piece 1 contains the connector, and piece 2 and 3 are the inside and outside negative space of the cushion part respectively.

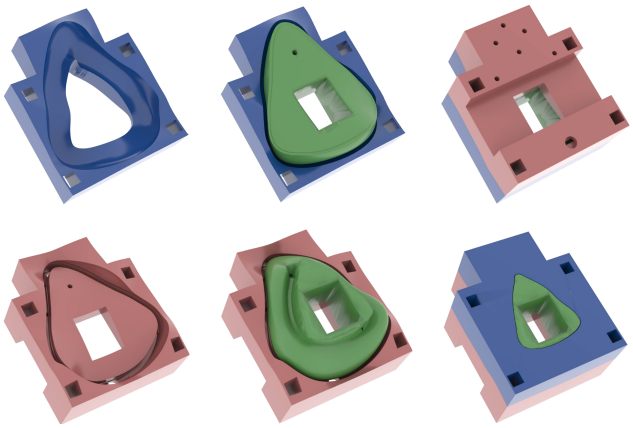

Figure 3: *We model a 3-piece mold for fabricating our designed mask interface. The square holes are used for aligning the mold pieces. The large circular hole on the red piece is for casting silicone materials. The small circular holes on the red and green pieces are for getting rid of air in the mold.*

The parting surface between pieces 1 and 2 approximates the minimal surface with its boundary curve on the connector. The parting surface between piece 1 and 3 is chosen so that the cushion surface in contact with piece 3 forms a height field along the z-axis direction. The parting surface of piece 2 and 3 is a cylinder surface as well as half of the cushion surface. The upper boundary of the cylinder surface is at the inner edge of the cushion.

All the air vents for casting and for air escape are standard cylinders. Most air vents are connected to the mask interface's connector. Only one air vent connects to the cushion inner edge, which is the hole on the green mold piece in Figure 3.

We also experimented with a 2-piece mold, which was obtained by merging pieces 2 and 3. The issue is that the profile curve bends inward, creating an undercut in the merged piece. Although the soft cushion can deform and escape from the undercut, it still poses challenges during pre-processing, such as when removing the internal support materials. That is why we divide this merged piece into two parts.

Fabrication procedure of the custom-fit PAP mask. We first use our computational tool to design a custom-fit PAP mask interface for a specific subject. Next, we model the mold geometry based on the mask interface, and 3D print the mold pieces with Ultimaker S5 printer and tough PLA material. Then, we perform pre-processing for casting, including: removing the water-soluble support material from the mold, cleaning and applying mold release to the mold, assembling the mold, and mixing the silicone materials. Finally, we inject the silicone into the mold through a syringe, applying air pressure with the Ultimus V dispenser from Nordson EFD. Due to the high viscosity of the silicone, it has difficulty flowing in the mold. That's why we have to apply air pressure to the syringe. After the silicone material is cured, we de-mold and clean the mask interface.

The custom-fit mask is obtained by assembling the BMC F2 mask frame with the fabricated mask interface. The BMC F2 mask frame can be reused with different subjects. Generally, it takes 2 days to fabricate one custom-fit mask. The most time-consuming parts are 3D printing the mold pieces (around 30 hours) and curing the silicone (around 16 hours).

References

LI, M., FERGUSON, Z., SCHNEIDER, T., LANGLOIS, T., ZORIN, D., PANOZZO, D., JIANG, C., AND KAUFMAN, D. M. 2020. Incremental potential contact: Intersection- and inversion-free, large-deformation dynamics. *ACM Trans. on Graph. (SIGGRAPH) 39*, 4, 49:1–49:20.

NELDER, J. A., AND MEAD, R. 1965. A simplex method for function minimization. *The Computer Journal 7*, 4, 308–313.