Fabricating 3D Figurines with Personalised Faces

J.Rafael Tena Moshe Mahler Thabo Beeler Max Grosse Hengchin Yeh
Iain Matthews

Disney Research

Abstract

We present a semi-automated system for fabricating figurines with faces that are personalised to the individual likeness of the customer. The efficacy of the system has been demonstrated by commercial deployments at Walt Disney World Resort and Star Wars Celebration VI in Orlando Florida. Although the system is semi automated, human intervention is limited to a few simple tasks to maintain the high throughput and consistent quality required for commercial application. In contrast to existing systems that fabricate custom heads that are assembled to pre-fabricated plastic bodies, our system seamlessly integrates 3D facial data with a predefined figurine body into a unique and continuous object that is fabricated as a single piece. The combination of state-of-the-art 3D capture, modelling, and printing that are the core of our system provide the flexibility to fabricate figurines whose complexity is only limited by the creativity of the designer.

1 Introduction

A long-standing goal in business is to provide consumers with goods and services that are personally tailored to them. Many companies are currently experiencing demands from their customers for the delivery of personalised products [1]. While consumer tailored services have become ubiquitous, the elevated costs of producing *one-of-a-kind* prototypes have comparatively hindered the delivery of customised fabricated goods. Technological advances in 3D printing have the potential to enable the fabrication of highly customised goods for a fraction of the traditional cost, opening new avenues for research and commercial enterprise. In the realm of customised merchandise, one opportunity is to use the face of the consumer to personalise artifacts. In particular, multiple companies are pursing the production of customised figurines, such as *bobble heads* and action figures, with faces that intend to capture the likeness of the costumer.

We present a semi-auotmated system for fabricating figurines with faces that are personalised to the individual likeness of the costumer. Our system is robust and was commercially deployed for Disney's *D-Tech Me* experiences, which were hosted at Walt Disney World Resort, a vacation destination in Orlando, and Star Wars Celebrations VI, an annual gathering of fans of the Star Wars franchise. During deployment, young girls could have their facial likeness captured in a figurine of a Disney Princess of their choice, while general

Figure 1: Examples of figurines fabricated by our system.

customers could have their facial likeness ”frozen” in *carbonite* as represented in one of the classical scenes of the Star Wars movie saga.

Although the system is semi-automated, human intervention is limited to a few simple tasks to maintain the high throughput and consistent quality required for a commercial application. In contrast to existing systems that fabricate custom heads that are then attached to pre-fabricated plastic bodies, our system seamlessly integrates 3D facial data with a predefined figurine body into a unique and continuous object that is fabricated as a single piece by 3D printing. The core components of our system, *i.e.* State-of-the-art 3D capture, modelling and printing, provide the flexibility to fabricate figurines whose complexity is only limited by the creativity of the designer. Our system is a flexible platform for creating customised products, and it is also an experience that tangibly engages the general public with state-of-the-art computer vision and computer graphics techniques. Several example figurines created by our system are shown in Figure 1.

The system comprises five distinct steps: *i)* acquiring the 3D geometry and colour of the customer's face; *ii)* registering the customer's 3D data to a face template; *iii)* attaching the facial 3D data to a figurine body; *iv)* colour matching facial 3D data to the figurine body; and *v)* 3D printing of the completed figurine. Figure 2 shows a block diagram of the system. Our main contribution is the demonstration of a complete system that applies computer vision, face modeling, and analysis to fabricate high quality products at a commercial volume and with little user input. It is the only system to date capable of mass producing fully integrated, seamless, single-piece figurines. As an additional contribution, our system engages the general public with state-of-the-art computer vision, graphics, and 3D printing.

2 Figurine Modelling

In order to fabricate figurines, first we require a human-like computer generated (CG) 3D figurine whose face will be used capture the likeness of the customer. A professional 3D modeler creates offline a CG

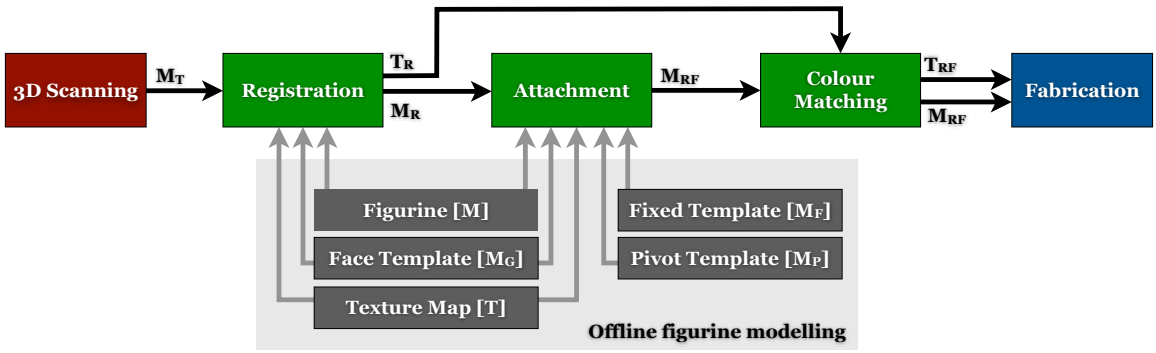

Figure 2: A system for fabricating figurines with personalised faces. Coloured boxes are processes while grey boxes are assets generated offline. Arrows represent flows of assets (see text and Table 1) through the pipeline.

model, for example, a figurine of *Belle* from Disney's *Beauty and the Beast* is shown in Figure 3. The face of the figurine has generic human features and its body has been sculpted to match the anatomy of a young girl. The CG model is expressed as a triangulated mesh, i.e. its surface is approximated by a set of triangle primitives each of which is described by the cartesian coordinates of its three vertices. The modeler also creates a texture map, an image that describes the colour of each of the triangle primitives of the figurine. This is accomplished by projecting each of the triangles to the image plane and filling the corresponding pixels with the desired colour. The texture map of our *Belle* figuring is shown in Figure 3. Once the CG figurine is completed the modeller extracts from it three additional templates: the face template, and what we call the pivot and fixed templates. Each template is a subset of the figurine. The three templates of our *Belle* figurine are also shown in Figure 3. The templates are used to guide different processes of our system and provide an additional layer of quality control as we will show in the following sections. For brevity we associate the CG figurine, its templates and other 3D CG assets with unique mathematical variables as described in Table 1. This variables are also used in Figure 2 to illustrate how the different templates move through the system.

3 3D Facial Capture

The first process of our system is to capture the geometry and colour of the face of the customer. We use the single-shot 3D scanner described by Beeler and colleagues [2]. Their scanner consists of a multi-camera, passive stereo setup that can capture the 3D geometry and colour of the customer's face in a single shot under standard light sources, and produces sub-millimeter accuracy. The scanner outputs the acquired facial geometry as a triangulated mesh with per-vertex colours specified as RGB triplets. A typical facial scan contains over half a million vertices. In comparison to laser-based 3D scanners, computer-vision based scanners as the one used by our system do not require the subject to remain still for prolonged periods of time. This characteristic makes them suitable for capturing the faces of customers of any age and allows

Asset	Variable
Customer's facial 3D scan	M_T
CG figurine	M
CG figurine's texture map	T
CG figurine's face template	M_G
CG figurine's pivot template	M_P
CG figurine's fixed template	M_F
CG figurine's body template	M_B
Registered face template	M_R
Registered texture map	T_R
Customised figurine	M_{RF}
Customised figurine's texture map	T_{RF}

Table 1: Variables used to designate the assets that flow through the pipeline of our system. This variables are used throughout the text and figures.

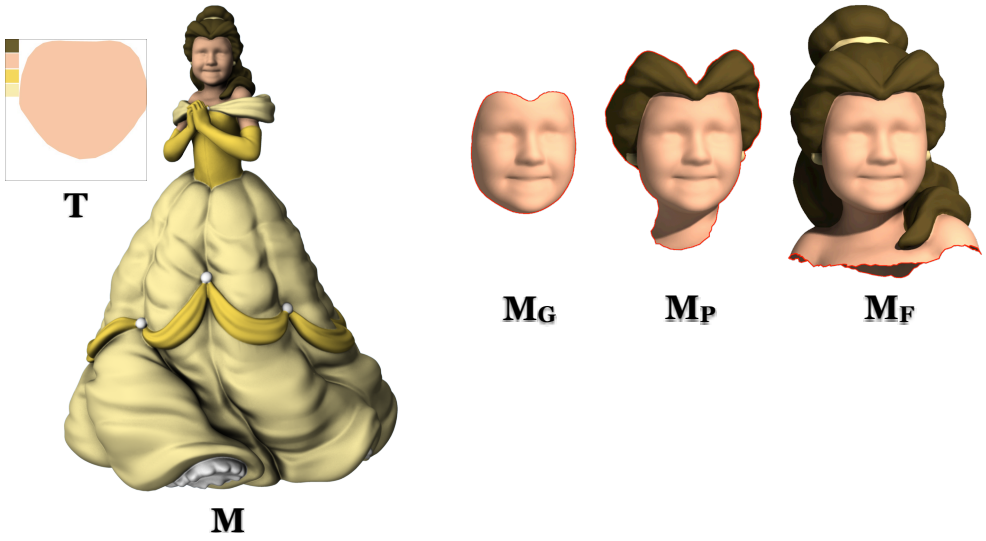

Figure 3: A figurine (M) of *Belle* from Disney's *Beauty and the Beast*, and its corresponding texture map (T). The figurine has generic facial features and the anatomy of a young girl. The modeler also extracts the face (M_G), pivot (M_P) and fixed(M_F) templates. The system uses the boundary vertices highlighted in red to compute the thin plate spline interpolant that constructs the figurine.

Figure 4: Example of a customer's 3D facial scan (M_T). The 3D scanner presented by Beeler and colleagues [2] was used to acquire the data. The scanner utilizes the images shown in the bottom row to reconstruct the 3D geometry and colour.

for fast, high throughput, acquisition. Figure 4 shows an example of 3D data captured by our system. The scanner simultaneously captures eight photographs which it uses to reconstruct the 3D geometry of the customer's face. Our system is not married to a particular 3D acquisition device, any method or hardware that can accurately capture 3D geometry and colour is acceptable.

4 Registration

Once a customer has been 3D scanned, the next step is to register the acquired geometry to the figurine's face template. Registration is a process by which point-to-point correspondences are established between two 3D objects, traditionally two instances of the same class. There are multiple approaches to 3D registration. We use and adaptation of the non-rigid template-based 3D facial registration method described by Tena and colleagues [3]. Non-rigid template-based methods typically optimise an energy function that deforms the template into the shape of a target object. In our case we want to deform the face template M_G, extracted during figurine modelling, into the shape of the customer's 3D scan M_T, acquired during the scanning process. To guarantee high registration accuracy even in the presence of unpredictable facial expressions, the method relies on a human operator to annotate several key feature points, such as the eye and mouth corners, on the template and the target.

The registration method of Tena *et al.* [3] achieves sub-millimeter accuracy and proceeds in three stages: *i) global mapping*, *ii) local matching*, and *iii) energy minimisation*. During *global mapping*, thin plate spline interpolation (TPS) as introduced by Bookstein [4] is used to warp the face template so that the set

of k corresponding feature points annotated by the human operator on both the template and the 3D scan are brought into alignment. For example, let $\mathbf{p}_i = (x_{pi}, y_{pi}, z_{pi})$ be the i^{th} feature point on M_G; and $\mathbf{q}_i = (x_{qi}, y_{qi}, z_{qi})$ the corresponding feature point on M_T. If the distances between the feature points on M_G are defined as $d_{ij} = |\mathbf{p}_i - \mathbf{p}_j|$, the interpolating function that warps $\mathbf{p}_i$ to $\mathbf{q}_i$ is:

$$F(x, y, z) = [f_x(x, y, z), f_y(x, y, z), f_z(x, y, z)] \tag{1}$$

$$f_{k=x,y,z}(x, y, z) = a_1^k + a_x^k + a_y^k y + a_z^k z + \sum_{i=1}^{n} w_i^k B(|\mathbf{p}_i - (x, y, z)|), \tag{2}$$

where $B(d)$ is a *biharmonic function* ($B(d) = d$ in our implementation), and $(\mathbf{w}^k, a_1^k, a_x^k, a_y^k, a_z^k z)^T = f(B(d_{ij}), \mathbf{p}, \mathbf{q})$ are the parameters of the thin plate spline interpolant. A detailed description for computing the interpolant parameters can be found in Bookstein's published work [4].

The *local matching* step consists of finding for each vertex $\mathbf{v_i}$ in the face template M_G a corresponding point on the surface of the 3D scan M_T. The registration algorithm of Tena and colleagues [3], defines the corresponding point to vertex $\mathbf{v_i}$ as the closest vertex $\tilde{\mathbf{v}}_\mathbf{i}$ in M_T. Instead, we define the corresponding point $\tilde{\mathbf{s}}_\mathbf{i}$ as the intersection of a ray, projected from $\mathbf{v_i}$, and the surface of M_T. Because M_T is a triangulated mesh, any point on its surface can be described using *barycentric coordinates* [5]. If the euclidean distance between $\mathbf{v_i}$ and $\tilde{\mathbf{s}}_\mathbf{i}$ is greater than $\frac{1}{10}$ of the length of the 3D facial scan, or their surface normals differ by more than $90°$, the correspondence of $\mathbf{v_i}$ to $\tilde{\mathbf{s}}_\mathbf{i}$ is flagged as unreliable.

The *energy minimisation* step consists of optimising an objective function with a data term E_{ext} that penalises the euclidean distance between the correspondences determined during *local matching*, and a smoothness term E_{int} that penalises deformations of the face template:

$$E_t = E_{ext} + \lambda E_{int} \tag{3}$$

$$E_{ext} = \sum_{i=1}^{n} w_i |\tilde{\mathbf{s}}_i - \mathbf{v}_i|^2 \tag{4}$$

$$E_{int} = \sum_{i=1}^{n} \sum_{j=1}^{m} |(\mathbf{v}_i - \mathbf{v}_j) - (\bar{\mathbf{v}}_i - \bar{\mathbf{v}}_j)|^2, \tag{5}$$

where λ is a constant that weights the contribution of each term, n is the number of vertices in the face template M_G, w_i is a weight set to 0 for correspondences flagged as unreliable and 1 otherwise, m is the number of vertices $\mathbf{v}_j$ connected by an edge to $\mathbf{v}_i$, and $\bar{\mathbf{v}}_i$ and $\bar{\mathbf{v}}_j$ are their original positions in M_G. In oder words, the data term pushes the vertices of the face template towards their corresponding points on the surface of the 3D face scan while the smoothness term restrains neighbouring vertices in the face template to maintain their relative positions to each other. Accordingly the parameter λ controls the rigidity if the face template.

The objective function is minimised using the conjugate gradient method. As described by Tena *et al.* [3], we follow an iterative approach to increase the accuracy of the registration process. Specifically we optimise Equation 3, apply the non-shrinking filter proposed by Taubin [6] to further regularise the face template, set

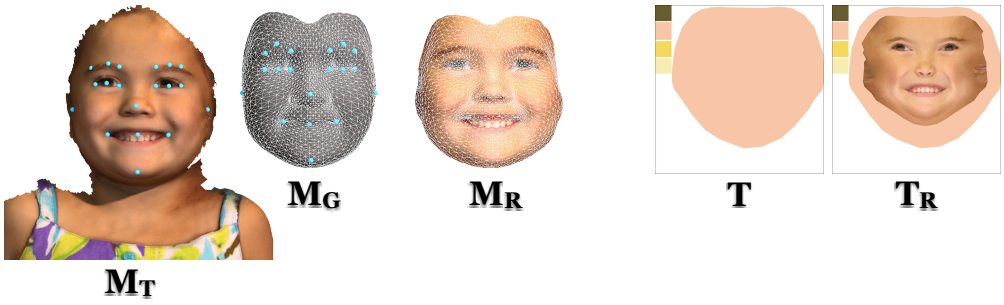

Figure 5: Registration example: The system registers the customer's face scan (M_T) to the face template (M_G) and creates the registered template (M_R). The system also extracts the colour information of the customer's 3D scan and adds it to the figurines texture map (T) to generate the registered texture map (T_R). The registration algorithm is initialised by the feature points highlighted in blue.

the values of $\bar{\mathbf{v}}$ to the current positions of $\mathbf{v}$, and repeat *local matching*, *energy minimisation*, and filtering. Applying the non-shrinking filter to the face template allows to use lower values of λ *i.e.* a less rigid face template, while maintaining its smoothness. After registration, we refer to the face template as the *registered template* M_R. The registered template has the same mesh topology as the face template, but the 3D shape of the 3D face scan. The registration algorithm is shown in Algorithm 1.

Algorithm 1 Registration algorithm

```
Annotate feature points on M_T
Global mapping
for i = 1 to 2 do
    Local matching
    Energy minimisation
    Filter M_G
    Reset v̄ to the current v values
end for
Transfer texture from M_T to M_R
```

After registration, it is necessary to transfer the colour of the 3D face scan M_T to the registered template M_R. The transfer is accomplished by texture mapping. Remember that a modeller has already created a texture map to which all triangle primitives of the figurine, including those of the face template, were mapped. More specifically, each vertex of the face template is parameterised by a (u, v) pair that maps its 3D coordinates to a point in the 2D image. Because the face template and the registered template share the same topology, they both map to the same (u, v) pairs in the texture map. Accordingly, for each triangle of the registered template in uv space we find all the pixels $\mathbf{t_i}$ of the texture map that fall within the triangle. Using barycentric coordinates we find for each $\mathbf{t_i}$ its corresponding 3D coordinate $\mathbf{s_i}$ on the surface of the registered template. We then project a ray from $\mathbf{s_i}$ onto the surface of the 3D face scan to find the corresponding point $\tilde{\mathbf{s}}_\mathbf{i}$ and its barycentric coordinates, remember that the face scan is also a triangulated mesh. The vertex colours (RGB triplets) of the triangle on which point $\tilde{\mathbf{s}}_\mathbf{i}$ lies, along with the point's barycentric coordinates are used to find

the RGB value that is assigned to pixel $\mathbf{t_i}$ on the texture map. Figure 5 shows an example of registration and texture extraction. The 3D face scan M_T is shown on the left, and the face template M_G on the middle. The set of feature points used for global mapping are shown in blue. The registered template M_R is shown on the right.

4.0.1 System Implementation details

The registration process of our system for fabricating personalised figurines takes as inputs a 3D, vertex-coloured, high-resolution (over half a million vertices), triangulated face scan M_T; and a CG textured figurine M with texture map T and corresponding face template M_G (a few thousand vertices). The system automatically removes the triangle primitives corresponding to the face template from the figurine M to create the body template M_B. Mathematically $M_G \cup M_B = M$ and $M_G \cap M_B \neq \emptyset$, because M_G and M_B share vertices at their boundaries (see Figure 6). A human operator manually annotates a set of 19 feature points on the 3D face scan, which have also been pre-annotated on the face template offline (see Figure 5). Annotating these feature points is the only human intervention required for the registration process. The system applies *ordinary procrustes analysis* [7] on the set of feature points to find the transformation that scales and aligns the face scan to the face template. After alignment, the registration algorithm deforms the face template into the shape of the face scan to create the registered template M_R and its corresponding texture map T_R, as shown in Figure 5. The constant λ that weights the contribution of the smoothness term is set to 0.25 to emphasise the data term and obtain greater deformation. Finally, a set of 1000 vertices is randomly selected on the face and registered templates and the system uses procrustes analysis to align the latter to the former. This final step ensures the scale and orientation of the registered template are correctly matched to the body template M_B.

5 Attachment

At this stage of our system's pipeline, we have a registered template that carries the 3D facial geometry of the customer and a body template, which is our faceless figurine. The next step is to attach the registered template to the body template. Both templates are complementary to each other and have corresponding vertices on their boundaries. However, they cannot be directly rejoined because their shapes do not match as we show in Figure 6. There are several methods in the literature to perform mesh editing and deformation operations [8]. Some of the more favoured approaches fall in the family of gradient domain methods. For instance the work by Sorkine *et al.* [9] and Yu *et al.* [10] describes the use of gradient domain methods to merge two different meshes. Because we have corresponding vertices at the boundaries of the registered and body templates, there is a one-to-one mapping to create a seam, which allows for a simpler and less computationally expensive solution. To close the gap between the registered template M_R and the body template M_B, we used the thin plate spline interpolation technique [4], that was also used for registration, to warp M_B onto M_R. The choice of warping the body M_B onto the face M_R ensures that the registered template is not deformed, which preserves the facial identity of the customer. To further restrict the interpolant, we resort to the pivot template M_P and the fixed M_F template that were created offline during the figurine

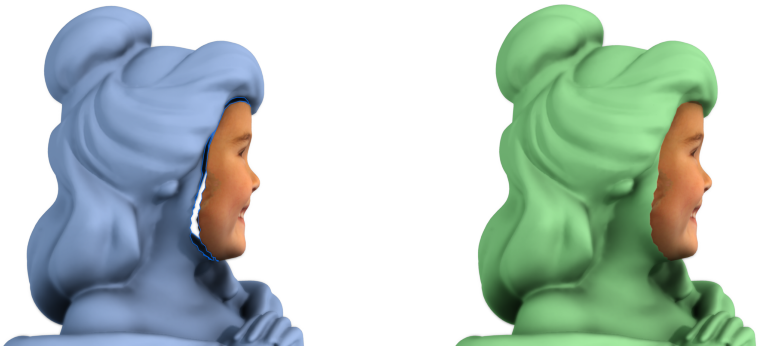

Figure 6: The system warps the body template (M_B) using thin plate spline interpolation to fit it to the registered template (M_R). The body template is shown in light blue before the warp. The fully assembled figurine (M_{RF}) is shown in green. The boundary vertices highlighted in blue, and the boundary vertices of the fixed and pivot templates (highlighted in red in Figure 3) are used as control points for the warping interpolant.

modeling phase described earlier in this work. The boundary vertices of the fixed and pivot templates, which are highlighted in red in Figure 3, and the boundary vertices of the registered and body templates, which are highlighted in blue in Figure 6 are all used as control points for the thin plate spline interpolant. Let $\mathbf{p_R}$, $\mathbf{p_B}$, $\mathbf{p_P}$, and $\mathbf{p_F}$ be the sets of vertices at the boundaries of M_R, M_B, M_P and M_F respectively. Then we find the interpolant F(x,y,z) that maps the set $(\mathbf{p_B}, \mathbf{p_P}, \mathbf{p_F})^T$ to the set $(\mathbf{p_R}, \mathbf{p_P}, \mathbf{p_F})^T$. This interpolant warps the body template while minimising the deformation around the vertices $\mathbf{p_P}$ and $\mathbf{p_F}$. In practical terms, the pivot and fixed templates allow the modeler to control which section of the figurine will be warped in order to accommodate the geometry of the customer's face. For our *Belle* figurine, the fixed and pivot templates shown in Figure 3 ensure that all the vertices below the figurine's shoulder line are not warped and that the overall shape of *Belle's* characteristic hairstyle is preserved. Figure 6 shows the body template before and after warping. Once the warping is completed, the boundary vertices of the registered template are fused to those of the body template to create a single manifold mesh, which is the completed customised figurine M_{RF}. Finally, the system smooths the seam by filtering the fused vertices using the Laplacian filter proposed by Vollmer and colleagues [11]. The filtering proceeds in an iterative fashion. On a first pass, the filter is applied to the seam vertices and their 2^{nd} order neighbours *i.e.* vertices that are separated from the seam by up to two edges. Second and third passes are applied up to the 1^{st} order neighbourhood, and the seam only respectively. The iterative approach achieves smoothing while preventing denting.

5.0.2 System Implementation Details

The attachment process of our system is fully automatic. The inputs are the registered templateM_R, the body template M_B, the pivot template M_p, the fixed template M_f, and the original texture map T created by the modeler. We usually define the pivot template to include half of the head and neck in the coronal plane,

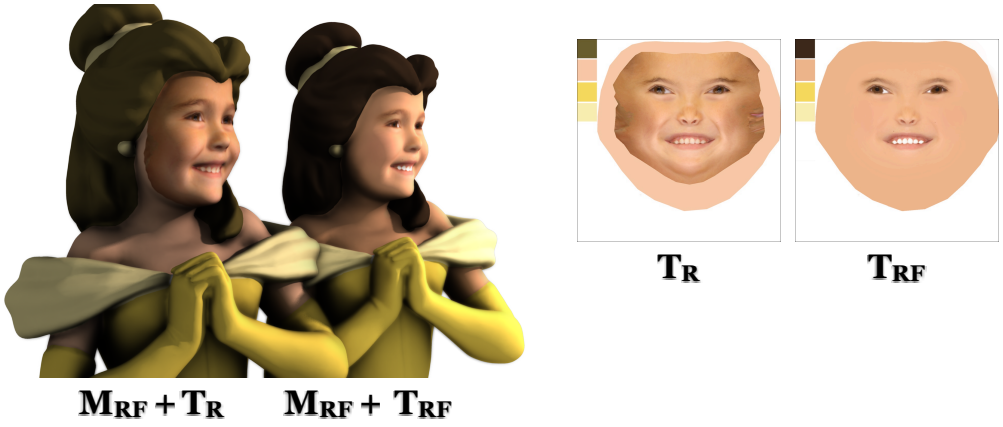

Figure 7: Colour matching for *D-Tech Me Disney Princess*. An artist edits the registered texture map (T_R) to create the customised figurine's texture map (T_{RF}) . The customised figurine (M_{RF}) is shown with the two different texture maps. The artist edits the skin tone and hair color to match those of the customer.

which ensures the warping function maintains the shape of the head. We define the fixed template to include everything above the shoulders of the figurine to ensure the overall shape of the body shape is preserved. To avoid smoothing sharp geometric details intended by the modeller, prior to filtering the seam, the system extracts, from texture map T which contains the original figurine colours prior to registration, the colours of the vertices selected for filtering and their 1^{st} order neighbours A vertex that has a neighbour with a colour different than its own is located at a detail boundary and is therefore deselected for filtering.

6 Colour Matching

With the digital customised figurine fully assembled, the next process employs a human artist to smooth and edit the figurine's texture map to either match the skin and hair color of the scanned customer or adapt it to match the figurine's style. Keeping an artist in the loop gives the system flexibility to produce figurines with different aesthetic requirements while guaranteeing consistent quality. For our *Belle* figurine, the colour matching process includes: i) whitening the eyeballs and teeth to increase contrast, ii) selecting a single skin tone that best approximates the customer's, iii) blending the selected skin tone into the texture map, and iv) selecting a single hair tone that best approximates customer's. We account for the limitations in the colour gamut of 3D printing by allowing the artist to select only skin and hair tones that are known to produce aesthetically pleasing results. For other figurines, the colour matching process might be completely different. For instance, the process for our *D-Tech Me Carbon Freeze Me* figurines, shown in Figure 8, consists only of adjusting the texture map RGB balance and blending in a grey tone. By keeping this module under the control of an artist, the system can handle complex editing processes without the need of adding new specialised algorithms. Figure 7 shows a figurine before and after the colour matching process.

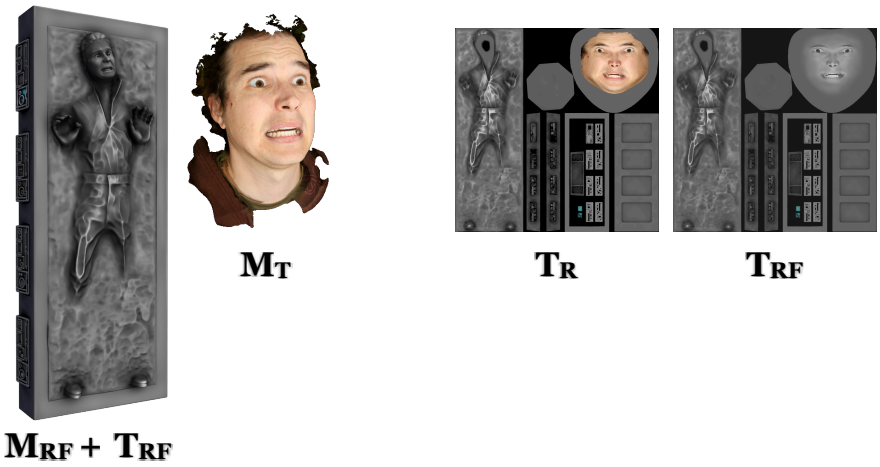

Figure 8: Colour matching for *D-Tech Me Carbon Freeze Me*. An artist edits adjusts the RGB colour balance of the registered texture map (T_R) and blends in a *carbonite* grey tone to create the customised figurine's texture map (T_{RF}) . The customised figurine (M_{RF}) is shown with its customised texture map.

7 Fabrication

Fabrication is the last process of our pipeline. A 3D printer brings the completed digital customised figurine into the real world. Any 3D printing device capable of handling manifold triangulated meshes is applicable. Our system uses the *ZPrinter® 650*, a state-of-the-art colour printer manufactured by *3D Systems* [12] that uses composite powder as printing material. This printer has a resolution of 600×540 dpi, a printing volume of $254 \times 381 \times 203$ mm, and can resolve features as small as 0.1 mm. After printing, it also provides a range of finishing options to increase durability and provide different aesthetics. An example of our *Belle* figurine is shown in Figure 9.

Our system prints figurines that range in height from 76.2 to 177.8 mm. We found smaller sizes to be fragile at the limbs and with not enough facial detail to be easily recognisable. The printing throughput is highly variable and dependent of the figurine size and model, from an approximate maximum of 6 figurines per hour for the 76.2 mm figurines (48 figurines/batch) to a minimum of 0.75 figurines per hr for wide 177.8 mm figurines (6 figurines/batch). After printing, a matte resin finishing is applied to increase durability.

8 Deployment

We have successfully deployed our system for testing and commercialisation. During this process we have fabricated thousands of figurines encompassing 20 different models and 5 different colour matching treatments. In addition, we also experimented with creating different body types to reflect the age of the customer. There are four different body types available for female figurines to better approximate female bodies of different ages. The figurines in Figure 10 are examples of figurine models with different aesthetics and

Figure 9: The system fabricates the completed customised figurine (M_{RF}) with a customised texture map (T_{RF}) using a colour 3D printer. We use a *ZPrinter® 650*.

age adjusted body types. There have been 3 commercial deployments in Orlando, Florida for Disney's *D-Tech Me* experiences. World of Disney, a merchandise store located within Walt Disney World Resort, hosted *D-Tech Me Disney Princess*, while *D-Tech Me Carbon Freeze Me* was hosted at Hollywood Studios, also in Walt Disney World Resort, and at Star Wars Celebrations IV. During the experience, customer's get 3D scanned in a setup themed according to the *D-Tech Me* modality. With the help of a couple of human operators, customers are scanned a few times and then choose their most favourable portrait based on the images acquired by the scanner. The system reconstructs the selection and an operator shows costumers a low resolution version of their 3D facial scan, an experience that often generates excitement. Customers participating on *D-Tech Me Disney Princess* can chose out of seven available princess figurines and a desired hair tone. The selections of skin tone and body type are left to the artist performing colour matching and the 3D scanner operator respectively. *D-Tech Me Carbon Freeze Me* offered a single figurine model representative of Star Wars' Captian Solo frozen in *carbonite*. The rest of the fabrication process is transparent to the customers who receive their customised figurine a couple of weeks later by post. Once the order is completed, the system performs a full 3D reconstruction of the image data collected by the scanner to create the customer's facial 3D scan. Reconstruction times can take up to 20 minutes on a standard computer. The system sends the face scan to a station manned by a human operator who annotates the feature points required for registration. A trained operator can do the job in less than a minute. When annotation is completed the system executes the registration and attachments processes. The corresponding computation time is approximately 50s on a standard computer. After registration and attachment, the system sends the assembled figurine to the artist's work station for colour matching. An artist takes approximately two minutes to edit the texture map. The actual time is not only dependent on the skill but also on the figurine's aesthetics requirements. It takes longer to process a princess figurine than a carbonite one. Finally the system sends the completed figurine to the 3D printer for fabrication, followed by finishing, packaging and shipping.

Figure 10: Digital examples of figurines created by our system for *D-Tech Me Carbon Freeze Me* and *D-Tech Me Disney Princess*. The artist applied different aesthetic treatments during colour matching. The male figurines have bodies of different heights to reflect different ages, while the female figurines have different chest to waist ratios.

9 Conclusion

We present a complete system for fabricating custom figurines with personalised faces. It is the only system to date capable of mass producing fully integrated, seamless, single-piece figurines. The commercial deployment of our system and the production of thousands of figurines attests to its robustness and efficacy. As an additional bonus, the system has engaged the general public with state-of-the-art computer vision, graphics, and 3D fabrication technology in an intuitive and tangible manner. We believe that customised production systems, such as the one here described, will become more common and have a broad impact, beyond the merchandise domain explored in this work, as 3D printing technology becomes more prevalent.

Although our system is automated enough to allow for high throughput, a fully automated system is our future goal. Current 3D facial registration techniques still require some kind of human interaction for initialisation when high precision is required and no constraints are placed on the facial expression of the target face. An avenue for future work is advancing the state-of-the-art to achieve fully automatic facial registration. Additionally, we would like to further automate the colour matching process by performing feature extraction, segmentation, and expression analysis on the texture map. This would allow to perform aesthetic editing operations without human intervention. The downside of automation is that algorithms tailored for the aesthetics of each figurine model would be necessary, adding to the system's complexity. There are still many questions to be answer with respect to the self-recognition process that motivates a customer to identify with the fabricated figurine. Our current system ignores the customer's full body shape and hairstyle when producing the figurine, however these elements are very important for the self-perception of many individuals. In the future we will explore techniques that would allow us to incorporate the customer's body and hair into the process of figurine fabrication. It is not only a question of developing the right technology, but also of preference and perception to find the right balance between complexity and customer satisfaction.

10 Acknowledgements

The commercial deployment of our system required a lot of work from multiple parties across The Walt Disney Company. We would like to acknowledge Tom Ngo, Cyd Tetro, and the rest of the staff at Disney Research that have worked in the project. Parks & Resorts Merchandise, and the New Technolgy Group and Industrial Engineering at Walt Disney World Resort in Orlando. We also acknowledge io-development for developing the software for deployment and 3D Systems for providing the 3D printing hardware and expertise during development and deployment.

References

[1] L. Hvam, N. H. Mortensen, and J. Riis, *Product Customization*. Springer-Verlag, 2008.

[2] T. Beeler, B. Bickel, P. Beardsley, B. Sumner, and M. Gross, "High-quality single-shot capture of facial geometry," *ACM Transactions on Graphics*, vol. 29, pp. 40:1–40:9, July 2010.

[3] J. R. Tena, M. Hamouz, A. Hilton, and J. Illingworth, "A validated method for dense non-rigid 3D face registration," in *Proceedings of the IEEE International Conference on Video and Signal Based Surveillance*, 2006.

[4] F. Bookstein, "Principal warps: thin-plate splines and the decomposition of deformations," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 11, pp. 567 –585, jun 1989.

[5] E. Weisstein, "Barycentric coordinates." From Mathworld A Wolfram Web Resource http://mathworld.wolfram.com/BarycentricCoordinates.html, August 2007.

[6] G. Taubin, "Curve and surface smoothing without shrinkage," in *Computer Vision, 1995. Proceedings., Fifth International Conference on*, pp. 852 –857, jun 1995.

[7] I. L. Dryden and K. V. Mardia, *Statistical Shape Analysis*. John Wiley & Sons, 2002.

[8] M. Botsch and O. Sorkine, "On linear variational surface deformation methods," *IEEE Transactions on Visualization and Computer Graphics*, vol. 14, pp. 213–230, Jan./Feb. 2008.

[9] O. Sorkine, D. Cohen-Or, Y. Lipman, M. Alexa, C. Rössl, and H.-P. Seidel, "Laplacian surface editing," in *Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, pp. 175–184, 2004.

[10] Y. Yu, K. Zhou, D. Xu, X. Shi, H. Bao, B. Guo, and H.-Y. Shum, "Mesh editing with poisson-based gradient field manipulation," *ACM Trans. Graph.*, vol. 23, pp. 644–651, Aug. 2004.

[11] J. Vollmer, R. Mencl, and H. Muller, "Improved laplacian smoothing of noisy surface meshes," *Computer Graphics Forum*, vol. 18, pp. 131–138, Sept. 1999.

[12] 3D Systems, "3Dsystems." http://www.3dsystems.com, September 2012.