Demonstration of Style2Fab: Functionality-Aware Segmentation for Fabricating Personalized 3D Models with Generative AI

Citation: Faruqi, Faraz, Katary, Ahmed, Hasic, Tarik, Abdel-Rahman, Amira, Rahman, Nayeemur et al. 2023. "Demonstration of Style2Fab: Functionality-Aware Segmentation for Fabricating Personalized 3D Models with Generative AI."

As Published: https://doi.org/10.1145/3586182.3615769

Publisher: ACM|The 36th Annual ACM Symposium on User Interface Software and Technology

Persistent URL: https://hdl.handle.net/1721.1/152894

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Demonstration of Style2Fab: Functionality-Aware Segmentation for Fabricating Personalized 3D Models with Generative AI

Faraz Faruqi
ffaruqi@mit.edu
MIT CSAIL
Cambridge, MA, USA

Ahmed Katary
akatary@mit.edu
MIT CSAIL
Cambridge, MA, USA

Tarik Hasic
thasic@mit.edu
MIT CSAIL
Cambridge, MA, USA

Amira Abdel-Rahman
amira.abdel-rahman@cba.mit.edu
Center for Bits and Atoms, MIT
Cambridge, MA, USA

Nayeemur Rahman
nayeem31@mit.edu
MIT CSAIL
Cambridge, MA, USA

Leandra Tejedor
leandra0@mit.edu
MIT CSAIL
Cambridge, MA, USA

Mackenzie Leake
leake@mit.edu
MIT CSAIL
Cambridge, MA, USA

Megan Hofmann*
Khoury College of Computer Sciences, Northeastern University
Boston, MA, USA

Stefanie Mueller*
stefanie.mueller@mit.edu
MIT CSAIL
Cambridge, MA, USA

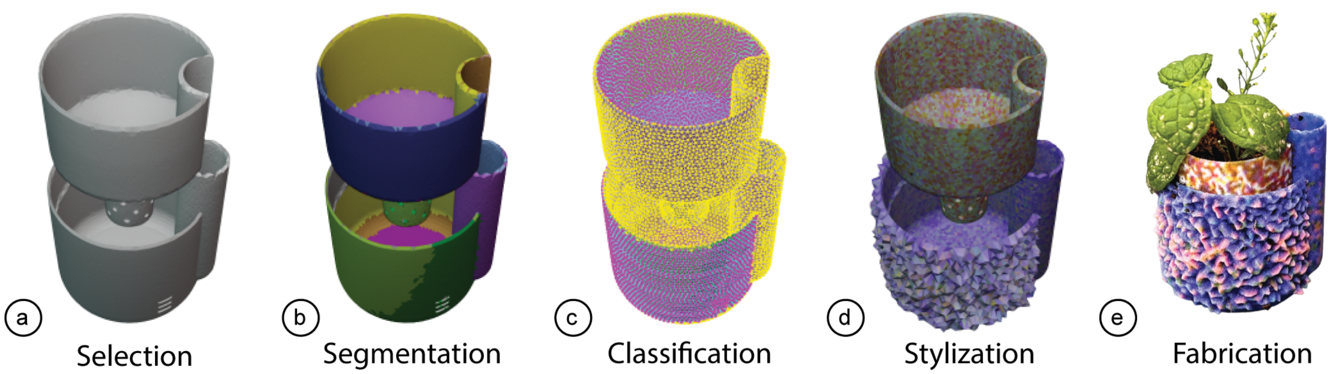

Figure 1: To stylize a 3D model without affecting its functionality a user: a) selects a model to stylize, b) segments the model, c) automatically classifies the segments into aesthetic or functional, d) selectively styles only the aesthetic segments, and e) fabricates their stylized model.

ABSTRACT

With recent advances in Generative AI, it is becoming easier to automatically manipulate 3D models. However, current methods tend to apply edits to models globally, which risks compromising the intended functionality of the 3D model when fabricated in the physical world. For example, modifying functional segments in 3D models, such as the base of a vase, could break the original functionality of the model, thus causing the vase to fall over. We introduce Style2Fab, a system for automatically segmenting 3D models into functional and aesthetic elements, and selectively modifying the aesthetic segments, without affecting the functional segments. Style2Fab uses a semi-automatic classification method to decompose 3D models into functional and aesthetic elements, and differentiable rendering to selectively stylize the functional segments. We demonstrate the functionality of this tool with six application examples across domains of Home Interior Design, Medical Applications, and Personal Accessories.

*Equal contribution

UIST '23 Adjunct, October 29–November 01, 2023, San Francisco, CA, USA

ACM ISBN 979-8-4007-0096-5/23/10.
https://doi.org/10.1145/3586182.3615769

CCS CONCEPTS

• **Human-centered computing → Human computer interaction (HCI).**

KEYWORDS

personal fabrication; digital fabrication; 3d printing; generative AI.

ACM Reference Format:
Faraz Faruqi, Ahmed Katary, Tarik Hasic, Amira Abdel-Rahman, Nayeemur Rahman, Leandra Tejedor, Mackenzie Leake, Megan Hofmann, and Stefanie Mueller. 2023. Demonstration of Style2Fab: Functionality-Aware Segmentation for Fabricating Personalized 3D Models with Generative AI. In *The 36th Annual ACM Symposium on User Interface Software and Technology (UIST '23 Adjunct), October 29–November 01, 2023, San Francisco, CA, USA.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3586182.3615769

1 INTRODUCTION

A key challenge for many makers is modifying or "*stylizing*" [1, 11] open source designs shared in online repositories [2, 8] (e.g., Thingiverse [3]). While these platforms provide numerous ready-to-print 3D models, customization is limited to changing predefined parameters [1]. While recent advances in deep-learning methods enable aesthetic modifications in 3D models with *styles* [10, 14], remixing existing designs with these styles presents new challenges. Beyond aesthetics, 3D printed models often have designed functionality that is directly related to geometry. Manipulating an entire 3D model, which can change the whole geometry, may break this functionality. Styles can be selectively applied, but this requires the novice maker to identify which pieces of a 3D model affect the functionality and which are purely aesthetic — a daunting task for users remixing unfamiliar designs. In some cases, users can label functionality in CAD tools [5, 15], however, most of the models shared in online repositories are 3D models that have lost this key meta-data. Alternatively, datasets for 3D printing [4] allow users to easily explore 3D printing settings for their 3D models, but they lack support for estimating print results for 3D models after manipulation.

To help makers make use of emerging 3D manipulation tools, we present "Style2Fab", a system that allows users to decompose 3D models designed for 3D printing into components based on their functional and aesthetic parts. This method allows makers to selectively stylize 3D models while maintaining the desired original functionality. Style2Fab decomposes 3D models into three classes of segments: 1) *aesthetic*, contributing only to model aesthetics; 2) *internally*-functional, related to assembly of component-based models; or 3) *externally*-functional, related to an interaction with the environment. The system uses a topology-based method that automatically segments 3D meshes, and classifies the functionality of those segments into the aforementioned three categories. Finally, Style2Fab, allows users to stylize the aesthetic segments based on text descriptions, using differentiable rendering [7] as proposed in Text2Mesh [10]. Our work demonstrates how we extend these methods to enable complex remixing of open-source 3D meshes for 3D printing without modifying their original functionality.

In the remaining sections, we present the Style2Fab system which uses a functionality-aware segmentation method to help makers stylize 3D printable designs. Next, we present several application examples from domains ranging from Home Interior Design to Medical Applications.

2 STYLE2FAB SYSTEM DESIGN

We developed Style2Fab as a plugin to the open-sourced 3D design software tool Blender. We designed the system to selectively stylize 3D models without changing their functionality beyond aesthetics, demonstrating the utility of our functionality-aware segmentation method to automatically identify the segments of a model that will have the greatest effect on its functionality. Style2Fab uses Text2Mesh [10], a differentiable-rendering-based stylization method, allowing users to describe their desired styles with text prompts.

2.1 Stylization of Segments

Figure 2: A demonstration of the differences between functionality-aware and global styling. A "flexi cat" model (Thing: 3576952) is shown (a) without styles, (b) with functionality-aware styles, and (c) functionally broken by global styles.

To stylize individual segments based on text prompts, we use Text2Mesh [10]. Text2Mesh uses a neural network architecture that leverages CLIP [12]. The system considers a 3D model as a collection of vertices, where each vertex has a color channel (RGB) and a 3D position that can move along its vertex normal. Text2Mesh reduces the loss between the 3D model rendering (CLIP representation) from different angles and the CLIP representation of the textual prompt using gradient descent. Text2Mesh makes small manipulations in both the color channel and vertex displacement along the vertex normal for each of the vertices in order to make it look more similar to the text prompt. This allows the system to generate a stylized 3D model that reflects the user's desired style.

By default, this method stylizes the entire mesh and changes the geometry of the functional segments. In Figure 2a-c, we show that global stylization can render a functional object, in this case, an articulated cat, inoperable. Thus, we augment this system by adding an additional step of masking the gradients and setting functional vertices to zero. This allows manipulation of the color and displacement channels while preserving the functional segments of the model. As specified in Text2Mesh [10], we run this optimization for 1500 iterations.

2.2 User Interface and Workflow

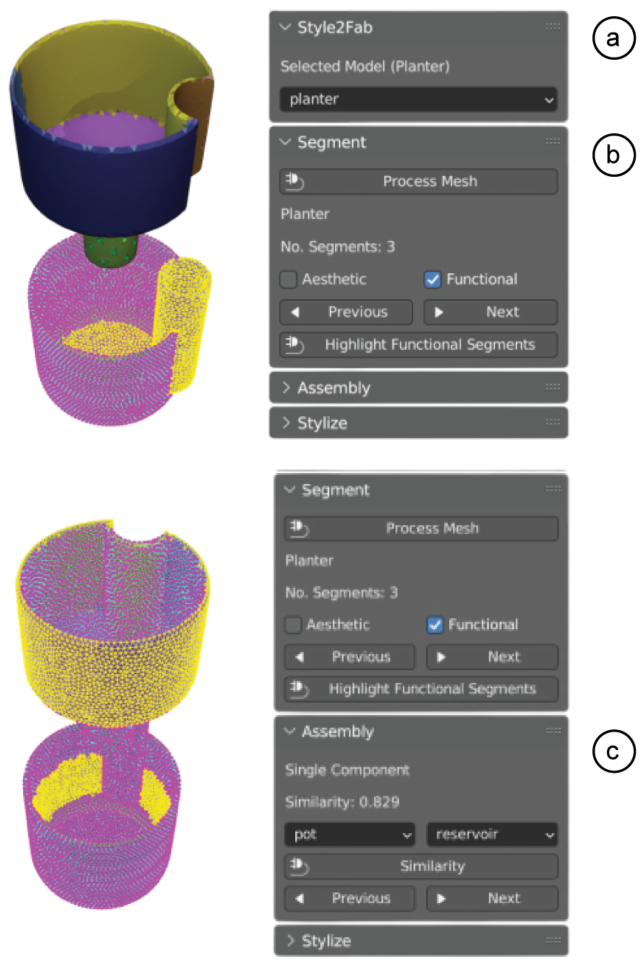

Figure 3: In the Style2Fab UI, the user (a) loads their model and (b) browses the functionality labels after segmentation. (c) In the case of multi-component models, the user can examine and adjust linked segments between components.

To stylize a model with Style2Fab, the user must: (1) pre-process their model for segmentation and stylization, (2) segment and classify the functionality of each segment, (3) selectively apply a style to segments based on functionality, and (4) review their stylized model. We break these tasks up into four menus in the user interface ((Figure 3)).

2.2.1 *Pre-Processing and Segmentation.* Once the user has loaded an OBJ file of their 3D mesh into the plugin, they "Process" the model to standardize its resolution and automatically detect the number of segments (i.e., k) needed to classify functionality across the model (Figure 3b). By default, the resolution is set to 25k faces, based on our evaluation, but the user is free to adjust this value. After segmentation, the system assigns each segment a unique color to help with visual identification. If the user wants more or fewer segments they can modify the value of k in the interface and re-segment the model. For multi-component models and mechanisms, the user can load multiple meshes representing each component and segment them in parallel.

2.2.2 *Functionality-Verification of Segments.* After segmentation, the plugin opens panels displaying the functionality classification of each segment. Users can review each segment and determine if they agree with the classification. The user's goal is to identify the set of segments that should not be stylized to preserve the desired functionality of the design. The user may have a different intended use of the model than the original designer (e.g., hanging a vase from the ceiling instead of placing it on a table) and thus change the functional classification of specific segments (e.g., labeling the base as aesthetic). To simplify this process, the user can select "highlight all functional segments" (Figure 3b) to have all segments classified as functional highlighted in the user interface. If they agree with this segmentation, they can move on to stylization. Otherwise, they can individually review all segments. Note that when walking through the different segments, the segment is highlighted on the model.

When working with multiple components, the user can review the segments that were classified as having internal context based on linkages between components of the model. In the "Assembly" panel, the user can walk through connected faces on distinct models (Figure 3c). If the user disagrees with this classification, they can adjust this similarity parameter between 0 and 1.

2.2.3 *Selective Stylization of Aesthetic Elements.* Users can stylize aesthetic segments of a 3D model by entering a natural language description of their desired style and clicking "Stylize Mesh". The completed model is then rendered alongside the original for review. Users can iterate on this process and apply new styles using new text prompts, or re-segment the model as needed.

3 DEMONSTRATIONS

In this section, we showcase Style2Fab's functionality-aware stylization through six application scenarios across four categories: Home Decor, Personalized Health Applications, Tactile Educational Artifacts, and Personal Accessories. These examples highlight the versatility of tools that use functionality-aware segmentation to personalize models.

3.1 Home Interior Design

Interior design is a popular domain for personalized fabrication. Here, we demonstrate customizing a Self-Watering Planter[1], a two-component model with internal and external functionality. The

[1] https://www.thingiverse.com/thing:903411/

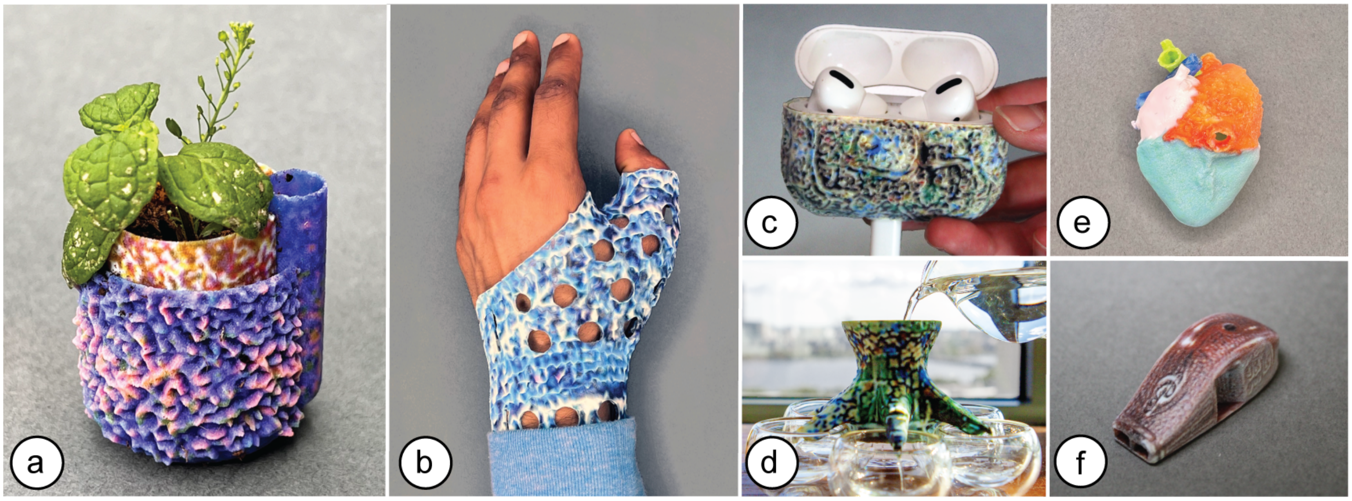

Figure 4: Application scenarios for Style2Fab (all models sourced from Thingiverse): (a) A multi-component self-watering planter styled as "*A rough multi-color Chinoiserie Planter*". (b) A personalized Thumb Splint styled like "*a blue knitted sweater*". (c) A personalized AirPods cover "*in the style of Moroccan Art*". (d) A Drinks Dispenser model styled as "*made of vintage mosaic glass tiles*". (e) A color-coded, textured educational model of the human heart. (f) A functional whistle styled as "*A beautiful whistle made of mahogany wood*".

internal functionality relates to the assembly of the pot and reservoir, while external functional segments include the base and watering cavity. Using Style2Fab, we segmented the model, verified the functional aspects, and applied the *"rough multi-color Chinoiserie Planter"* style. The fabricated model showcases the desired aesthetics without compromising its self-watering capabilities or stability. Another home decor example is the Indispensable Dispenser[2], a drink dispenser that distributes liquid into six containers via interior cavities and spouts. We used Style2Fab to preserve the functional segments (base and interior cavities) and applied a *"colorful water dispenser made of vintage mosaic glass tiles"* style to the aesthetic segments. The resulting dispenser, shown in Figure 4a, retains its functionality while exhibiting the desired visual appearance.

3.2 Medical/Assistive Applications

"Medical Making" [9] and "DIY Assistive Technology" [2] are emerging and critical domains for personalized fabrication by non-technical experts. Social accessibility research [13] shows that considering both the aesthetic and functional features of medical/assistive devices increases their adoption. However, individuals with disabilities and their clinicians may not have the time or expertise to personalize devices [6]. We first demonstrate stylizing a geometrically-complex thumb splint source from Thingiverse[3], to appear as "*A beautiful thumb splint styled like a blue knitted sweater]*". Like Hofmann et al's [6] participant, we wanted the model to blend into the sleeve of a sweater (Figure 4b). Our functionality-aware segmentation method preserved the smooth internal surface that contacts the skin and holes that increase breathability. The exterior is stylized and attractive. Our second example stylizes a tactile graphic of a human heart to apply unique textures to each region of the heart. This would make identifying each region easier for a blind person who accesses the model through touch. Our automatic segmentation and classification method identifies different segments of the heart that can be stylized with different textures (Figure 4e).

3.3 Personalizing Accessories

In this application scenario, we demonstrate how functionality-aware segmentation can be used to personalize accessories, such as an Airpods Cover and a Whistle. We demonstrate Style2Fab's personalization capabilities using a Thingiverse Airpods case[4]. The interface segmented and preserved the functional aspects, including internal geometry and charging cable hole even though our classification method has no specific information about the external objects the cover interacts with. We stylized this model with the prompt: "*A beautiful antique AirPods cover in the style of Moroccan Art* (Figure 4e). The resulting model fits the AirPods Pro case and allowed charging while featuring Moroccan Art-inspired patterns. Next, we applied styles to the popular V29 Whistle from Thingiverse[5] without compromising its acoustic functionality. The system preserved the whistle's resonant chamber and mouthpiece while styling the exterior with a prompt: *"A beautiful whistle made of mahogany wood"* (Figure 4f). The functionality-aware styled whistle sounds like the original while the globally-styled whistle lost its functionality due to internal geometry manipulation.

[2] https://www.thingiverse.com/thing:832751/
[3] https://www.thingiverse.com/thing:5259956
[4] https://www.thingiverse.com/thing:4105467/files
[5] https://www.thingiverse.com/thing:1179160

4 CONCLUSION

In this demonstration, we present a novel system for functionality-aware segmentation and classification of 3D models. This system allows users to modify and stylize 3D models while preserving their functionality. We presented the user interface implemented as a plugin to the popular 3D modeling tool, Blender, along with the technical details for the classification and segmentation methods. Finally, we demonstrate the utility of this tool across several application domains. This work speaks more broadly to the goal of working with generative models to produce functional physical objects and empowering users to explore digital design and fabrication.

REFERENCES

[1] Celena Alcock, Nathaniel Hudson, and Parmit K. Chilana. 2016. Barriers to Using, Customizing, and Printing 3D Designs on Thingiverse. In *Proceedings of the 2016 ACM International Conference on Supporting Group Work* (Sanibel Island, Florida, USA) *(GROUP '16)*. Association for Computing Machinery, New York, NY, USA, 195–199. https://doi.org/10.1145/2957276.2957301

[2] Erin Buehler, Stacy Branham, Abdullah Ali, Jeremy J. Chang, Megan Kelly Hofmann, Amy Hurst, and Shaun K. Kane. 2015. Sharing is Caring: Assistive Technology Designs on Thingiverse. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) *(CHI '15)*. Association for Computing Machinery, New York, NY, USA, 525–534. https://doi.org/10.1145/2702123.2702525

[3] Thingiverse. com. 2023. Thingiverse - Digital Designs for Physical Objects. https://www.thingiverse.com [Online; accessed 4. Apr. 2023].

[4] Faraz Faruqi, Kenneth Friedman, Leon Cheng, Michael Wessely, Sriram Subramanian, and Stefanie Mueller. 2021. SliceHub: Augmenting Shared 3D Model Repositories with Slicing Results for 3D Printing. *arXiv preprint arXiv:2109.14722* (2021).

[5] Megan Hofmann, Gabriella Hann, Scott E. Hudson, and Jennifer Mankoff. 2018. Greater than the Sum of Its PARTs: Expressing and Reusing Design Intent in 3D Models. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3173574.3173875

[6] Megan Hofmann, Kristin Williams, Toni Kaplan, Stephanie Valencia, Gabriella Hann, Scott E. Hudson, Jennifer Mankoff, and Patrick Carrington. 2019. "Occupational Therapy is Making": Clinical Rapid Prototyping and Digital Fabrication. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3290605.3300544

[7] Hiroharu Kato, Deniz Beker, Mihai Morariu, Takahiro Ando, Toru Matsuoka, Wadim Kehl, and Adrien Gaidon. 2020. Differentiable Rendering: A Survey. *CoRR* abs/2006.12057 (2020). arXiv:2006.12057 https://arxiv.org/abs/2006.12057

[8] Stacey Kuznetsov and Eric Paulos. 2010. Rise of the Expert Amateur: DIY Projects, Communities, and Cultures. In *Proceedings of the 6th Nordic Conference on Human-Computer Interaction: Extending Boundaries (NordiCHI '10)*. ACM, 295–304. https://doi.org/10.1145/1868914.1868950

[9] Udaya Lakshmi, Megan Hofmann, Stephanie Valencia, Lauren Wilcox, Jennifer Mankoff, and Rosa I. Arriaga. 2019. "Point-of-Care Manufacturing": Maker Perspectives on Digital Fabrication in Medical Practice. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 91 (nov 2019), 23 pages. https://doi.org/10.1145/3359193

[10] Oscar Michel, Roi Bar-On, Richard Liu, Sagie Benaim, and Rana Hanocka. 2022. Text2mesh: Text-driven neural stylization for meshes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 13492–13502.

[11] Lora Oehlberg, Wesley Willett, and Wendy E. Mackay. 2015. Patterns of Physical Design Remixing in Online Maker Communities. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) *(CHI '15)*. Association for Computing Machinery, New York, NY, USA, 639–648. https://doi.org/10.1145/2702123.2702175

[12] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 8748–8763. https://proceedings.mlr.press/v139/radford21a.html

[13] Kristen Shinohara, Cynthia L. Bennett, Wanda Pratt, and Jacob O. Wobbrock. 2018. Tenets for Social Accessibility: Towards Humanizing Disabled People in Design. *ACM Trans. Access. Comput.* 11, 1, Article 6 (mar 2018), 31 pages. https://doi.org/10.1145/3178855

[14] Yawar Siddiqui, Justus Thies, Fangchang Ma, Qi Shan, Matthias Nießner, and Angela Dai. 2022. Texturify: Generating Textures on 3D Shape Surfaces. *SpringerLink* (Nov. 2022), 72–88. https://doi.org/10.1007/978-3-031-20062-5_5

[15] Tom Veuskens, Florian Heller, and Raf Ramakers. 2021. CODA: A Design Assistant to Facilitate Specifying Constraints and Parametric Behavior in CAD Models. In *Graphics Interface 2021*. https://openreview.net/forum?id=1dLDPJeafRZ