INTRODUCTION & PREVIOUS SUCCESSES

Ground-breaking *hardware* advances from the past decade or two have blurred the dividing line between the real and virtual worlds beyond recognition. Nowadays, one can scan a real-world object using a simple smartphone app or Kinect and render it in a virtual scene, and even recover full 3D virtual scenes from two-dimensional real-world pictures using Machine Learning [1]. The inverse direction of this challenge has been even more dramatically revolutionized by the democratization of 3D printing technology, accessible today to many consumers.

However, *software* has been lagging largely behind this revolutionary change in the Computer Graphics pipeline. Most algorithms in the Geometry Processing toolbox are designed to work on certain, fixed input geometries. This fails to incorporate the uncertainties and errors associated with real-world geometric data capture, as well as the possibility of including further capturing steps as intermediate steps in the pipeline. Similarly, virtual 3D modelling tools rarely account for whether the designed object is actually fabricable or not, leading to disappointing outcomes when attempting to bring the virtual asset into the real world.

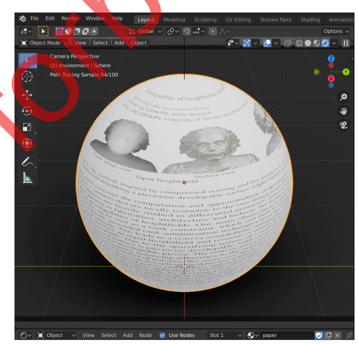

As an example, take a sheet of paper and attempt to fold it into a sphere. You will find that this is impossible; in other words, a paper sphere is not a *fabricable* shape. However, existing 3D design software like Maya or Blender (see inset), unaware of this fact, allows one to build a virtual model of this sphere, and then render it with a texture that replicates that of a piece of paper to produce a realistic image. The same applies to wood that is bent for ship construction or steel sheets for architecture. In [2], published at ACM SIGGRAPH, we show that one characterize these fabricable shapes mathematically and use that insight to either constrain the user to them during the design process or to project a freely designed existing surface into a fabricable one.

Modelling these non-stretchable surfaces is an example of using properties of real-world materials to guide fabrication-aimed design. Another possibility is that the latter is not constrained by the object's material, but by the method by which it will be fabricated. For instance, we may design an object with the intention to then build it with the aid of a CNC machine with a round end mill. In [3], published in ACM SIGGRAPH Asia, we study the set of 3D surfaces that can be constructed in this way, and propose a way of projecting any given object into a CNC-millable one without the aliasing and lossy nature of previous methods.

My vision is that of an accessible and seamless integration between the real and virtual worlds: virtual design must be fabrication-aware and the Geometry Processing pipeline must be revolutionized to fully exploit data captured from the real world. My previous research tackles some of the specific challenges one encounters when making this vision a reality, and sets the stage for my proposed future work.

MY PLAN FOR THE FUTURE

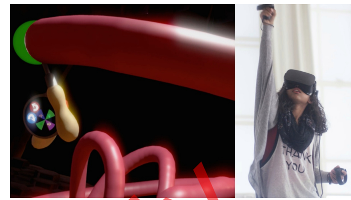

I have a long-term vision with short and mid term deliverables.

Efficient and robust Swept Volumes

A common way of constructing virtual objects is by sweeping an input shape across 3D space, and outputting the solid region of space covered by the shape as it moves [4][5]. An exciting application of this design framework is to Virtual Reality sculpting (see image, from Adobe Oculus Medium). Currently existing methods for computing these swept volumes lack efficiency and/or robustness (see [6][7][8][9][10][11]): if they are forced to run near to real time, the outputs present stroboscopic aliasing. By drawing insights from physics, I believe the problem can be rephrased with a change in the frame of reference. Instead of a surface moving rigidly across space, the surface can be thought of as fixed and every particle in space moving around it. Therefore, computing the swept volume can be rephrased in terms of a particle-surface collision detection problem, which I expect to solve efficiently and robustly by drawing from the fields of non-linear ray tracing [12][13][14] and importance sampling [15][16][17].

Uncertain Captured Geometry

Input (real scan) Our output

Many challenges remain in the question of transferring real world geometry into the virtual world. Often, this is done via 3D scanning, and many works in the literature [18][19][20] have been proposed for smoothing or post-processing the outputs of these scanners (among them, our [2], a figure of which is shown as inset). Dealing with this scan data, which is often noisy and filled with artifacts, often means making deterministic, final decisions to questions like "Is this depth data point an outlier, or does it really represent the real-world geometry?" or "Is this a hard crease present in the scanned object, or is it noise provoked by the user while scanning?"

Today, the answers to these questions are final and can dramatically impact what we conclude about the object further down the Geometry Processing pipeline. If we wrongly identified some extreme depth data points as outliers and discarded them, we may wrongly conclude the object to be more smooth than it really is, leading to very different material properties. Similarly, if we identified one-directional noise to be a crease in the real geometry instead of the result of noise, our elasticity or resistance calculations about the object may be completely off. In the future, I will experiment with adding a measure of uncertainty or stochasticity to the capture of real-world geometry. If we can measure the *possibility* of a certain data point being an outlier or a product of noise instead of making a deterministic choice, we should be able to make *probabilistic* conclusions about the properties of the object. Ideally, this would even inform the user to re-scan or put more emphasis on certain parts of the shape, leading to more precise results.

In summary, my plan for the future is to tackle specific challenges associated with the interaction between real-world and virtual-world geometry, while keeping an open mind and also taking on more broad ideas with the potential of revolutionizing the way we think of transporting geometry from one to the other. My previous work and publications in the field show a promising start in this direction, and I find myself thrilled to fulfill that promise in the following years.

REFERENCES - SILVIA GONZÁLEZ SELLÁN

[1] Yao, Y., Schertler, N., Rosales, E., Rhodin, H., Sigal, L., & Sheffer, A. (2020). Front2Back: Single View 3D Shape Reconstruction via Front to Back Prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 531-540).

[2] Sellán, Silvia, Noam Aigerman, and Alec Jacobson. "Developability of heightfields via rank minimization." *ACM Transactions on Graphics (SIGGRAPH)* 39.4 (2020): 109-1.

[3] Sellán, Silvia, Jacob Kesten, Ang Yan Sheng, and Alec Jacobson. "Opening and Closing Surfaces." *ACM Transactions on Graphics (SIGGRAPH Asia)* 396 (2020): 1-1.

[4] Peternell, M., Pottmann, H., Steiner, T., & Zhao, H. (2005). Swept volumes. *Computer-Aided Design and Applications*, *2*(5), 599-608.

[5] Abdel-Malek, K., Yang, J., Blackmore, D., & Joy, K. (2006). Swept volumes: fundation, perspectives, and applications. *International Journal of Shape Modeling*, *12*(01), 87-127.

[6] Rossignac, J., Kim, J. J., Song, S. C., Suh, K. C., & Joung, C. B. (2007). Boundary of the volume swept by a free-form solid in screw motion. *Computer-Aided Design*, *39*(9), 745-755.

[7] Menon, J., Marisa, R. J., & Zagajac, J. (1994). More powerful solid modeling through ray representations. *IEEE Computer Graphics and Applications*, *14*(3), 22-35.

[8] Campen, M., & Kobbelt, L. (2010, May). Exact and robust (self-) intersections for polygonal meshes. In *Computer Graphics Forum* (Vol. 29, No. 2, pp. 397-406). Oxford, UK: Blackwell Publishing Ltd.

[9] Von Dziegielewski, A., Hemmer, M., & Schömer, E. (2013). High precision conservative surface mesh generation for swept volumes. *IEEE Transactions on Automation Science and Engineering*, *12*(1), 183-191.

[10] Kim, Y. J., Varadhan, G., Lin, M. C., & Manocha, D. (2004). Fast swept volume approximation of complex polyhedral models. *Computer-Aided Design*, *36*(11), 1013-1027.

[11] Zhang, X., Kim, Y. J., & Manocha, D. (2009, October). Reliable sweeps. In *2009 SIAM/ACM joint conference on geometric and physical modeling* (pp. 373-378).

[12] Mo, Q., Yeh, H., & Manocha, D. (2015). Tracing analytic ray curves for light and sound propagation in non-linear media. *IEEE transactions on visualization and computer graphics*, *22*(11), 2493-2506.

[13] Weiskopf, D., Schafhitzel, T., & Ertl, T. (2004, September). GPU-based nonlinear ray tracing. In *Computer graphics forum* (Vol. 23, No. 3, pp. 625-633). Oxford, UK and Boston, USA: Blackwell Publishing, Inc.

[14] Stam, J., & Languénou, E. (1996, June). Ray tracing in non-constant media. In *Eurographics Workshop on Rendering Techniques* (pp. 225-234). Springer, Vienna.

[15] Cornuet, J. M., Marin, J. M., Mira, A., & Robert, C. P. (2012). Adaptive multiple importance sampling. *Scandinavian Journal of Statistics*, *39*(4), 798-812.

[16] Agarwal, S., Ramamoorthi, R., Belongie, S., & Jensen, H. W. (2003). Structured importance sampling of environment maps. In *ACM SIGGRAPH 2003 Papers* (pp. 605-612).

[17] Colbert, M., Premoze, S., & François, G. (2010, July). Importance sampling for production rendering. In *ACM SIGGRAPH*.

[18] Taubin, G. (1995, September). A signal processing approach to fair surface design. In Proceedings of the 22nd annual conference on Computer graphics and interactive techniques (pp. 351-358).

[19] Jones, T. R. (2003). *Feature preserving smoothing of 3D surface scans* (Doctoral dissertation, Massachusetts Institute of Technology).

[20] Weyrich, T., Pauly, M., Keiser, R., Heinzle, S., Scandella, S., & Gross, M. H. (2004). Post-processing of Scanned 3D Surface Data. *SPBG*, *4*, 85-94.