INTERACTING WITH PERSONAL FABRICATION MACHINES

By Stefanie Mueller, X-Consortium Career Development Assistant Professor of Electrical Engineering and Computer Science; Member, Computer Science and Artificial Intelligence Laboratory (CSAIL)

Personal fabrication tools, such as 3D printers, are on the way to enabling a future in which non-technical users will be able to create custom objects. With the recent drop in price for 3D printing hardware, these tools are about to enter the mass market: While the average consumer 3D printer was priced at $14,000 in 2007, today's hardware costs, on average, only $500[10]. Given the decreasing price, it is not surprising that the number of consumer 3D printers sold has doubled every year[10].

While the hardware is now affordable and the number of people who own a 3D printer is increasing, only a few users actually create new 3D models. Most download models from a platform, such as Thingiverse, and fabricate them on their 3D printers. At most, users adjust a few parameters of the model, such as changing its color or browsing between predetermined shape options.

I believe that personal fabrication has the potential for more: I envision a future in which, rather than just consuming existing content, 3D printers will allow non-technical users to create objects that only trained experts can create today. While there are many open challenges, I will use this article to discuss how we can improve the interaction model underlying current fabrication devices.

1.1 Interaction Model with 3D Printers Today

In the current interaction model, users sit at a computer and use a digital 3D editor to create a digital 3D model. Only at the end of the design process do users send the file to the 3D printer, which creates the object in one go. Because 3D printing is slow, this process tends to take hours of printing time for small objects and may even require overnight printing.

1.2 Drawing a Parallel to Personal Computing

Looking back in history, this interaction model with the delayed feedback was also common with early computers[1]. In the early '60s, computers were so slow that the average program had to be executed overnight. Feedback was delayed until the next morning and if a program failed, users had to repeat the entire process, potentially waiting another night for results.

Similar to 3D printers today, early computers were limited to expert users because when programs were executed in one go overnight, users had to know what they were doing to succeed.

1.3 Towards Turn-Taking and Direct Manipulation

However, today we are at a point at which even non-technical users can use personal computers. Beside many technical developments, two advances in the interaction model enabled this: (1) the move from executing in one go to *turn-taking*, and (2) the move from turn-taking to *direct manipulation*[3].

1) Turn-taking: By decreasing the interaction unit to single requests, turn-taking systems, such as the command line, provided users with feedback *after every input*. This enabled the trial-and-error process that non-technical users tend to employ: quickly iterating through potential solutions and building each step onto the results of previous ones[2]. However, while the turn-taking interaction model provided a great step forward to making the technology available for non-technical users, the feedback cycle was still limited in that it consisted of two discrete steps: users first had to create an input and only *afterwards* received feedback.

2) Direct manipulation: With the invention of direct manipulation[9] that further decreased the interaction unit to a single feature, users finally received real-time feedback: Input by the user and output by the system are so tightly coupled that no visible lag exists. This tightened feedback cycle has many benefits, among others that "novices can learn basic functionality quickly" and "retain operational concepts"[8]. (See Figure 1.)

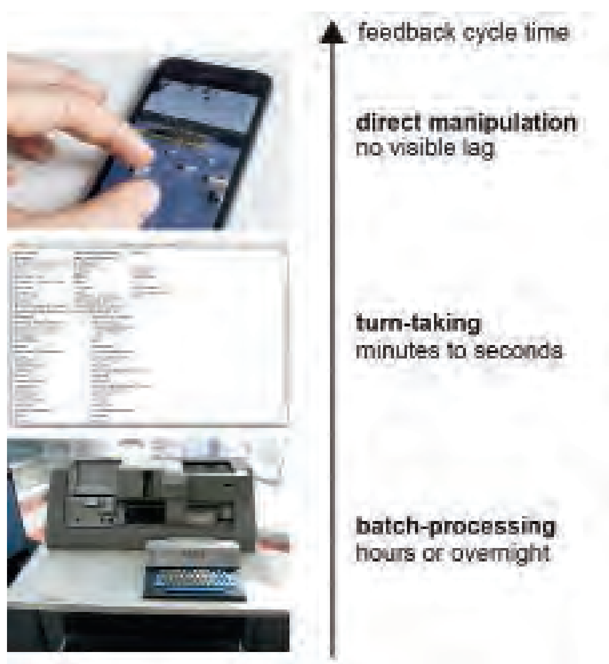

Figure 1

Figure 1. Server Feedback

As described above, the current interaction model of 3D printers requires objects to be fabricated in one go. Thus, from a human-computer interaction standpoint, we are today at the point at which we were with personal computers in the 1960s: Only few users are able to use the technology, and even for experts, it is a cumbersome process due to the delayed feedback.

1.4 Bringing Direct Manipulation to Fabrication

I argue that by repeating the evolution of the interaction model from personal computing, we will see the same benefits for personal fabrication: Direct manipulation will allow non-technical users to create physical objects as easily as they manipulate digital data with today's personal computers.

A direct manipulation system for personal fabrication needs to have four main characteristics: **1)** the physical environment is the workspace, not a digital editor; **2)** users work hands-on on the physical workpiece through physical tools as known from traditional crafting; **3)** each physical action results in immediate physical change, which can also be reversed; and **4)** in contrast to traditional crafting, users receive support from a computer system that helps to achieve precision.

In the following section, we show examples of two systems that implement the requirements listed above and iteratively decrease the interaction unit from entire objects, to single elements, to features to achieve real-time physical feedback.

1) Turn-taking: Interactive Laser-Cutting. To illustrate what a turn-taking system for personal fabrication might look like, we decrease the interaction unit from entire objects to single elements. In our system constructable[6], users draw with a laser pointer onto the workpiece inside a laser cutter. The drawing is captured with a camera. As soon as the user finishes drawing an element, such as a line, the constructable system beautifies the path and cuts it, resulting in physical output after every editing step. Different tools allow users to accomplish different tasks, such as copy-pasting physical shapes or creating matching finger joints between two edges. In addition, constructable ensures that all physical output is aligned (see Figure 2).

Figure 2

Figure 2: constructable

While constructable allows for fast physical feedback, the interaction is still best described as turn-taking because it consists of two discrete steps: users first perform a command and *then* the system responds with physical feedback.

2) Direct Manipulation: Continuous Forming. By decreasing the interaction unit even further to a single *feature*, we explore how to make the workpiece change while the user is manipulating it, resulting in real-time physical feedback: Input by the user and output by the fabrication device are so tightly coupled that no visible lag exists. Our system FormFab[7] provides such continuous physical feedback (see Figure 3). To accomplish this, FormFab neither adds nor subtracts material, but instead reshapes it (formative fabrication). A heat gun attached to a robotic arm warms up a thermoplastic sheet until it becomes compliant; users then control a pneumatic system that applies either pressure or vacuum, thereby pushing the material outwards or pulling it inwards. As users interact, they see the workpiece change continuously.

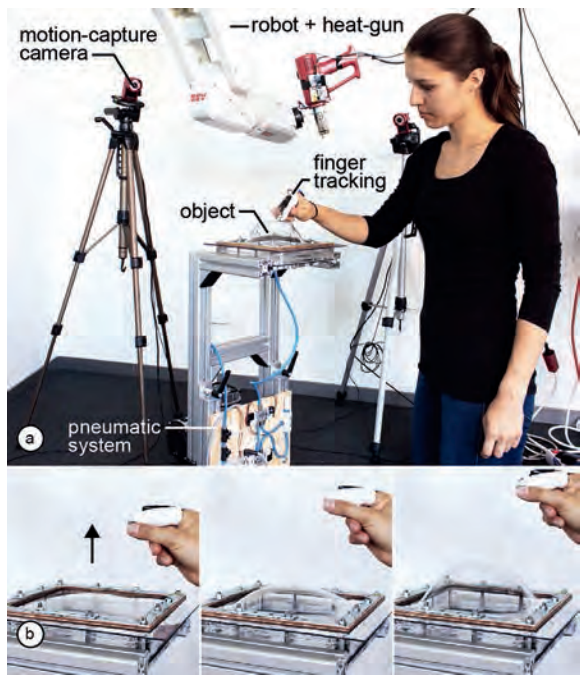

Figure 3

"I envision a future in which, rather than just consuming existing content, 3D printers will allow non-technical users to create objects that only trained experts can create today."

1.5 Discussion

Direct manipulation systems for personal fabrication extend the range of problems novice users can tackle, but they are subject to the same limitations as those for personal computing: While they are useful for some design problems, they are less so for others. As Norman et al.[4] point out, direct manipulation interfaces are limited to operations that can be done on "visible objects" and have "difficulties handling variables" and "distinguishing an individual element from a representation of a set or class of elements." Thus design problems that require more abstract thinking for which users must first sit down with a piece of paper and make a detailed plan are better handled with traditional digital 3D editing. In addition, the systems presented in this article inherently scale 1:1 and do not offer a way of inspecting a detail in magnification, which limits users to projects that fit a particular scale. The same way that a saw and a wood chisel cannot replace a detailed design process, our systems cannot replace a complex 3D editing tool for trained engineers.

1.6 Outlook

We focused on using technology available today to explore interaction paradigms that will become possible in the future when fabrication actually gets faster. In recent years, we have begun rapidly advancing towards such a future. The recently introduced 3D printer Carbon3D, for instance, speeds up fabrication by up to 200 times.

While there is little data today that could prove a Moore's law for 3D printers, an executive from 3D Systems, a leading manufacturer, noted in 2014 that 3D printing speed had, on average, doubled every 24 months over the previous 10 years[5]. If such a trend should materialize, it is not far-fetched to assume that fabrication technology will be able to provide feedback even for large high-resolution objects within seconds or even in real time, thereby enabling a future in which digital displays will be replaced with physical displays that transform virtual reality into actual physical reality.

"By repeating the evolution of the interaction model from personal computing, we will see the same benefits for personal fabrication: direct manipulation will allow non-technical users to create physical objects as easily as they manipulate digital data with today's personal computers."

References

[1] P. Ceruzzi, A History of Modern Computing, 2nd edition. The MIT Press (2003).

[2] M. Csikszentmihalyi, Flow: The Psychology of Optimal Experience. Harper Perennial Modern Classics (1990).

[3] J. Grudin, A Moving Target: The Evolution of Human-Computer Interaction. Human-Computer Interaction Handbook, 3rd edition, Taylor and Francis (2012).

[4] E. Hutchins, J. Hollan, and D. Norman. Direct Manipulation Interfaces. Human-Computer-Interaction (1985), vol. 1, pp. 311-338.

[5] Moore's law for 3D printing: https://3dprint.com/7543/3d-printing-moores-law

[6] S. Mueller, P. Lopes, P. Baudisch. Interactive Construction: Interactive Fabrication of Functional Mechanical Devices. Proceedings of the ACM UIST 2012, pp. 599-606.

[7] S. Mueller, A. Seufert, H. Peng, R. Kovacs, K. Reuss, T. Wollowski, F. Guimbetiere, P. Baudisch. Continuous Interactive Fabrication. Under review for ACM UIST 2017.

[8] B. Shneiderman. Direct Manipulation: A Step Beyond Programming Languages. Computer (1983), vol. 16, issue 8, pp. 57-69.

[9] B. Shneiderman. The Future of Interactive Systems and the Emergence of Direct Manipulation. Proceedings of the NYU Symposium on User Interfaces on Human Factors and Interactive Computer Systems (1984), pp. 1–28.

[10] Wohlers Report (2016). https://wohlersassociates.com/2016report.htm

SOCIAL DATA PROCESSING

By Devavrat Shah, Professor of Electrical Engineering and Computer Science; Member, Laboratory for Information and Decision Sciences (LIDS), Institute for Data, System and Society (IDSS); Director, Statistics and Data Science Center (SDSC)

Background. What went wrong with the election polls in the 2016 U.S. presidential election? How can the online activity of the population help curate better life experiences for all? Can we utilize online personas for reaching out to individuals in a targeted manner? What about predicting the demand for espadrilles this summer? Or ranking the performance of your favorite sports team? And what happened to the promise of using collective wisdom for stopping the spread of ''fake news'' on Facebook?

Answers to all those questions depend on our ability to process "social" data to extract meaningful information. For the past few decades, and even more so recently, everything online is being recorded. If aliens came to earth and inspected the social data — generated by us as a society (not by machines) — they would learn, for instance, that Patriots' Day is when the Boston Marathon is held.

Put another way: Social data presents us with an enormous opportunity for making data-driven decisions for better living, more efficient operations, more effective policy making, and overall uplifting of societies. Here, data is the enabler. Access to social data has been democratized; the key to success lies in the ability to process it so that we can extract meaningful information from it. This makes it feasible for someone like me as an "ivory-tower" academic to carefully think through a solution, test it out, and then have a chance of making an impact in the real world — and, in the process, advance the foundations of statistics and machine learning. In that sense, social-data processing presents an unusual, potentially once-in-a-generational, opportunity that can lead to a remarkable convergence of academia and industry.

Challenge. The standard approach for data-driven decisions following statistical decision theory is to use an appropriate model that connects data to decision variables, helps make desired predictions, and eventually facilitates optimization over decision choices. Here, data is generated by humans, so modeling social behavioral aspects is essential. Any social scientist can attest that modeling human behavior is an extremely intricate task and the resulting models can be highly context-dependent. That makes it especially challenging to come up with effective, meaningful models. The only hope is for gaining access to *lots* of social data to decipher the right model for a particular interest from a large class of models. In other words, we need a model that is flexible enough to capture a wide array of social scenarios. But the model must be sufficiently tractable so that with *enough* data it can capture the ground truth faithfully, and it is important that such a system can computationally scale along with the data.

In short, the key intellectual challenge is in finding a sufficiently flexible model for social data that is both *statistically* and *computationally* tractable. This is a major challenge, and its successful resolution can have substantial impact on all the previously mentioned scenarios — and many others.

Turning Weakness to Strength. To progress toward such a grand challenge, it is essential to identify the properties of social data that are ubiquitous across a variety of scenarios and that can be captured to develop meaningful models.

We have identified one such property: social data is (or should be) *anonymous*. That is, from the data-processing perspective, it should not matter who has generated the data. To put it another way, the overall conclusion should remain invariant if we re-name the individuals who have generated the data. For example, the results of a democratic election should not change even if the voters' names change, as long as the total number of votes for each candidate remains the same. In the same way, the popularity of a specific style of espadrilles does not depend on which specific individuals bought them, only on how many pairs are being purchased.

Anonymity seems like a constraint or a weakness from any angle you look at it.

After all, anonymity and privacy protections restrict the type of information that we can mine from data. But we derive strength from this apparent weakness. It will help us address the challenge of developing tractable and flexible models for social data.

Mathematically, anonymity can be viewed as the underlying "probabilistic model" having a certain "exchangeability" property. A remarkable development in mathematical statistics, starting with the work of de Finetti in the 1930s with further developments in the 1970s and 1980s, provides a crisp non-parametric characterization for such models: the Latent Variable Model. We utilize the Latent Variable Model for social-data processing for a variety of scenarios, including some of those discussed in the questions that opened this article.

Taking the First Steps. We start by examining the question of designing personalization or recommendation systems such as those used by Netflix, YouTube, Amazon, and Spotify. Here, the goal is using the history of an entire population's preferences to predict which movies, music, books, or other products that individual consumers may like and that they have not already experienced. On one hand, the question is: What is the best algorithm to design for that end goal using the non-parametric model emerging from exchangeability? On the other hand, that question has been with us since the dawn of the e-commerce era.

There is a popular algorithm, called Collaborative Filtering[0], that has been with us from the start and that continues to be used due to its simplicity and empirical success. In a nutshell,