Foundations and Trends® in
Theoretical Computer Science

Personal Fabrication

Patrick Baudisch and Stefanie Mueller

now
the essence of knowledge

Foundations and Trends® in Human-Computer
Interaction
Vol. 10, No. 3–4 (2016) 165–293

DOI: 10.1561/1100000055

Personal Fabrication

Patrick Baudisch
Hasso Plattner Institute, Potsdam, Germany
baudisch@hpi.de

Stefanie Mueller
Hasso Plattner Institute, Potsdam, Germany
MIT CSAIL, Cambridge, MA, USA
stefanie.mueller@mit.edu

Contents

Abstract

While fabrication technologies have been in use in industry for several decades, expiring patents have recently allowed the technology to spill over to technology-enthusiastic “makers”. The big question now is whether the technology will further progress towards consumers, which would allow the technology to scale from hundreds of thousands of users to hundreds of millions of users.

Such a transition would enable consumers to use computing not just to process data, but for physical matter. This holds the promise of democratizing a whole range of fields preoccupied with physical objects, from product design to interior design, to carpentry, and to some areas of mechanical and structural engineering. It would bring massive, disruptive change to these industries and their users.

We analyze similar trends in the history of computing that made the transition from industry to consumers, such as desktop publishing and home video editing, and come to the conclusion that such a transition is likely.

Our analysis, however, also reveals that any transition to consumers first requires a hardware + software system that embodies the skills and expert knowledge that consumers lack: (1) *hardware and materials* that allow fabricating the intended objects, (2) software that embodies *domain knowledge,* (3) software that embodies the know-how required to operate the *machinery,* and (4) software that provides immediate *feedback* and supports interactive exploration. At the same time, sustained success will only be possible if we also consider future implications, in particular (5) *sustainability* and (6) *intellectual property.* We argue that researchers in HCI and computer graphics are well equipped for tackling these six challenges. We survey the already existing work and derive an actionable research agenda.

P. Baudisch and S. Mueller. *Personal Fabrication.* Foundations and Trends® in Human-Computer Interaction, vol. 10, no. 3–4, pp. 165–293, 2016.
DOI: 10.1561/1100000055.

1

Introduction

In HCI and computer graphics, research on fabrication technology tends to be perceived as a recent trend. The truth, however, is that the technology itself has been in use for decades.

The reason that we as researchers may have missed the beginning of the field is that the field initially took place behind closed doors — as a small, high-margin market in industry that was protected by patents. Starting in the 1960s with computer-controlled laser cutters and milling machines and later on in the 1980s with 3D printing, the relevant technologies were initially conceived as a fast way for creating prototypes for product development. At the time, it was called "rapid prototyping technology."

The first industrial 3D printer, the *SLA-1* from *3D Systems*, was introduced in 1987 (Figure 1.1). Many other industrial systems followed with the invention of additional 3D printing techniques. With all patents being filed in the 1980s and 1990s by the future CEOs of large companies, such as *3D Systems* and *Stratasys*, the market was locked down for several decades.

In 2009, however, the first major patent expired, thereby initiating the transition of the technology from industry to the world

Figure 1.1: The first 3D printer: The *SLA-1 from 3D Systems*.

outside. Technology enthusiasts who grew out of hacker spaces and the crafting-oriented *DIY* culture had already created their own fabrication hardware (e.g., see the *RepRap* project, 2005) and now started commercializing their low-cost devices with products such as the MakerBot Cupcake CNC [2009]. These companies entered the market with the declared goal of targeting a market segment that industrial 3D printing companies had overlooked: low-cost 3D printers.

With more and more patents expiring, we currently see an increasing number of the 1980s and 1990s fabrication technologies becoming available outside of industry. While the last decade was marked by low-cost 3D printers that extruded plastic filament, we now see a diverse spectrum, including low-cost printers based on curing resins [e.g., the Form1. *Formlabs*, 2012] and sintering powder [e.g., Sintratec, 2014]. As a result, newly founded companies picked up the technologies and are now competing in the market, resulting in fast progress and price drops by several orders of magnitude.

Makers are playing a key role in this transition, as they make their own fabrication machines. This has resulted in hundreds of freely available 3D printer designs, as of today [Price Comparison 3D Printers].

These new fabrication machines are no longer closed-source industrial 3D printers that companies encapsulated to protect their IP, but instead open-source 3D printers that can easily be "hacked", which has given even further momentum to the evolution of these devices.

In the wake of this evolution, the maker movement continues to pick up additional fabrication technologies, including laser cutters [e.g., Glowforge, 2016], milling machines [e.g., Shapeoko, 2013], and water jet cutters [e.g., Wazer, 2016].

1.1 The promise of fabrication in the hands of consumers

The fact that fabrication technologies are already looking back at a 30+ year history seems to suggest that personal fabrication cannot be novel. This is *not* the case. What is novel about "personal fabrication" is not the "fabrication" thought, but the "personal".

There is no universally agreed upon definition for personal fabrication yet. In 2005, Neil Gershenfeld described personal fabrication as "the ability to design and produce your own products, in your own home, with a machine that combines consumer electronics with industrial tools." However, as of today, these are the homes of a selected few — the homes of technology enthusiasts.

The big question today is whether this evolution will continue, i.e., will fabrication transition not only from industry to technology enthusiasts, but will it continue to consumers[1]? The latter would promise to empower hundreds of millions of new users and could give the field of personal fabrication enormous impact.

So what would that impact be — what would consumers do with personal fabrication technology?

Our immediate reaction might be to look at today's makers, seeing the somewhat ad-hoc projects they create and to discard the potential

---

[1] There is no agreed upon name for this group of people. We use the term *consumers* here because all we know about them is that their intent is to "consume" the outcome of what they make, unlike makers who are interested in the technical process [Hudson et al., 2016]. Hudson et al. refer to consumers as "casual makers" but we argue this is not the best term as these people have little in common with makers. Also, the fact that they care about the *outcome* arguably makes them *less* casual than makers

of personal fabrication as a whole. This would be a mistake, because early adopters historically have never been good indicators for the following consumer market (a gap that has been referred to as the *chasm* [Moore, 2006]). This gap tends to be even larger for early adopters that are driven by technology enthusiasm, because their projects tend to revolve around exploring the technological possibilities rather than the applications. Makers today might reason "I have a 3D printer... let me find out what I can do with it . . . ", then look at a database, such as Thingiverse or Instructables, and download a project. Consequently, the threshold for the expected utility of the outcome can be arbitrarily low, as this group of users tends to perceive the technical challenge per se as rewarding.

This process stands in stark contrast with consumers who are motivated exclusively by the utility of the expected outcome [Hudson et al., 2016]. Consumers, who are in it for the result, thus share fewer values with the makers as they might appear to at first glance. So when we see makers today download and replicate interesting "proof-of-concept" objects, such as an interlocking gear mechanism, it gives us little indication of the types of problems consumers may tackle using the technology.

So what problems can we expect consumers to tackle? We argue that candidate problems come from several professional fields, in particular those fields that are primarily concerned with physical output, such as product design [Kim and Bae, 2016] as well as some areas of mechanical and structural engineering. If larger fabrication machines should become mass available as well, applications will also come from interior design, furniture construction [Lau et al., 2011], and related fields.

Any of these fields account for multi-billion dollar markets. If personal fabrication should enter these markets, personal fabrication could be expected to grow to the size of these markets.

In addition to the fields listed above, *new* fields may form around personal fabrication. This is an open-ended question and we may continue to see new applications over time. In 1968, Doug Engelbart asked what value could be derived if intellectual workers had access to an instantly responsive computer system 24 hours a day [Engelbart, 1962]. With personal fabrication we are facing the same type of question: what will intellectual workers do with a personal computer system *if that system also allowed creating immediate physical output*?

1.2 Personal fabrication and its underlying AD/DA pattern

In order to understand personal fabrication, we may compare personal fabrication with similar technologies from the history of interactive computing. In order to determine which technologies to consider, we will first try to understand what it is that characterizes personal fabrication.

We use the simple example of a copy machine for physical keys. Figure 1.2 shows the traditional workflow *before* personal fabrication. A key maker places the original key into the tracer unit of a mechanical key copy machine, and a blank key into the machine's milling unit. Both the tracer and the mill are tightly coupled. As the key maker traces the cuts of the original key, the milling part follows the same path, engraving the same pattern into the blank key.

The key copy machine is a highly specialized machine in that it replicates nothing but keys. It also is an analog machine, as we can tell from the fact that copies of copies eventually will not open the door anymore, as inaccuracies accumulate from generation to generation leading to larger and larger errors.

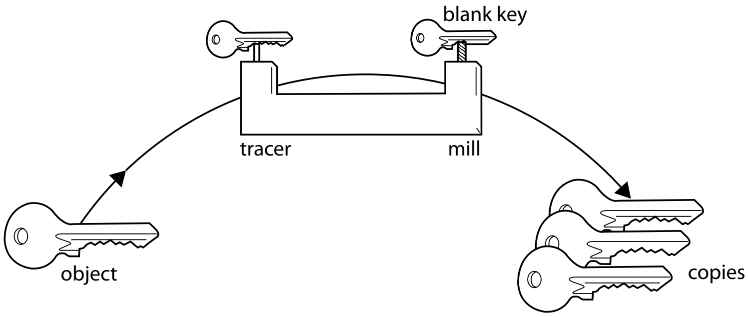

Figure 1.2: The traditional analog way of replicating keys.

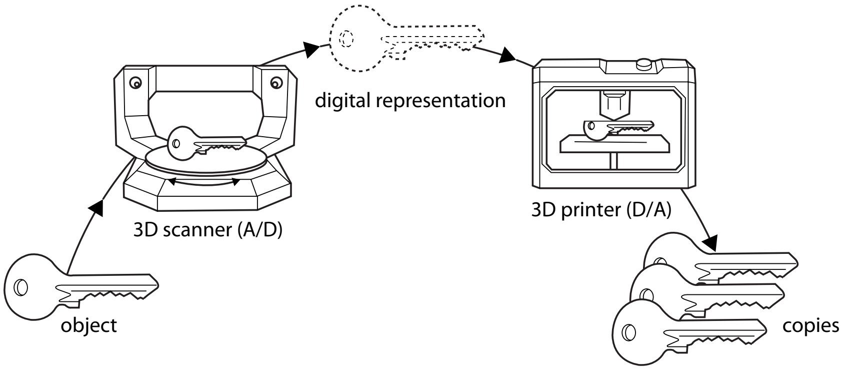

Figure 1.3: The digital solution consisting of scanner and printer that forms the basis for personal fabrication.

As shown in Figure 1.3, the personal fabrication workflow is essentially the same, except that it replaces the specialized key copy machine with a combination of a general-purpose 3D scanner with a general-purpose 3D printer.

This is what we think of as the schema underlying personal fabrication: (1) The scanner is a hardware unit that turns physical objects into digital objects, an “analog-to-digital converter” (AD). (2) The printer is a hardware unit that turns virtual objects into physical objects, a “digital-to-analog converter” (DA). In the shown “AD/DA” setup, these two units create a copy machine for physical objects, as first demonstrated in 1991 [Reyes, 1991] and commercially available today [ZEUS].

While the scanner/printer configuration is more complex than the specialized analog solution it replaces, the extra complexity pays off quickly as the setup is more flexible. For example, it applies to a wide variety of objects, rather than just keys.

More important, however, the two-machine solution and its intermediate digital representation allow creating additional workflows by merely adding software. For example, by inserting a software filter capable of re-inserting missing geometry, we can create a machine that repairs physical objects [Teibrich et al., 2015].

This illustrates the general pattern this setup is capable of: convert whatever problem needs to be solved to the digital domain, solve it in software, and convert the result back to a physical world. This is beneficial because developing and deploying new software tends to be faster and cheaper than creating and deploying new hardware.

The simple workflows that scan and produce in one go may not be the most interesting ones after all — the truly impactful workflows tend to involve digital storing and digital sharing. The new workflow, for example, allows using the same setup to make backups of physical objects, share designs in online repositories (such as the aforementioned *Thingiverse*), or distribute designs using a file sharing network. Any of these add tremendous impact to the original idea of a "copy machine" that goes way beyond what its analog counterpart was capable of.

1.3 Personal fabrication, like other AD/DA technologies before it, will result in disruptive change

If we assume that the transition of personal fabrication to consumers will actually happen, our next question naturally is to ask "how will it be?" Will personal fabrication lead to a big disruptive change or will it just add a small new commodity to people's lives? Where will personal fabrication ultimately lead?

In order to predict the future of personal fabrication, we now look at past innovations that structurally resemble personal fabrication in that they follow the same AD/DA pattern and see how these turned out.

Picking relevant past technologies is easy, because we have seen the AD/DA pattern before. Examples include desktop publishing, digital video editing, and digital music editing.

Desktop publishing: In 1969, the invention of the laser printer by Gary Starkweather at *Xerox* allowed for high-quality print output, which added the DA component to the already available AD image scanners. Before the introduction of this AD/DA pattern, users had to compose print layouts by photographing image and text elements literally *laid out* on a table. Layout based on personal computers (e.g., *Type Processor One*, 1983) allowed all this process to take place in software, which enabled fast iteration. Physical snippets and the camera disappeared from the process and the only memento of its existence is that publishers to date still require a "camera-ready version" of papers accepted for publication. The transition to software allowed a wider audience to gain access to desktop publishing or simplified word processing. As of today, *Microsoft Word* and *Google Docs* have brought the concept to over a billion consumers.

Digital Video Editing: Analog video editing in the early 1950s required users to locate the edit points by shuttling the physical tape to the desired location, carefully slicing the tape with a razor blade, and reconnecting it to the other desired tape parts with splicing tape. This process was time-consuming and limited in that it did not allow enhancing the video. Early computerized systems in the 1960s allowed synchronizing tape from different scenes by marking the scenes on the physical tape. In 1972, *SuperPaint* [Hiltzik, 2000] was the first graphics program that used [Frame grabbing] to convert analog video into digital images. This allowed rearranging segments and enhancing frames with digital data (e.g., changing hue, saturation, and value, or using different paintbrushes and pencils to draw on the frames), thereby laying the foundation for an entire new industry on digital editing and post-processing. As of today, hundreds of millions of mobile devices provide consumers not only with a built-in camera, but also with preinstalled digital video software (e.g., *iMovie* on iOS).

Digital Music Editing: Similarly, analog audio editing required users to cut tape and to manually reconnect it to the other desired parts. This made multi-track assemblies difficult, as it was hard to move one track in time relative to another. With the invention of the digital sound recording (*Pulse-code modulation* (PCM)) and new software for digital audio editing, the entire audio industry was transformed. As of today, hundreds of millions of mobile devices ship with the ability to record and play back audio, as well as consumer-friendly audio editing programs (such as *GarageBand* on *iOS*).

If one really wanted to trace back the AD/DA pattern to its beginning, one might even consider text. In the early 1960s, text was replicated by first encoding the data into an analog punch card, which was then replicated using an analog teleprinter (e.g., *Teletype Model 33 ASR*, 1963). In the mid-1960s, keyboards (AD) were introduced as a more flexible means to edit text on a computer, as they made changes a matter of retyping a small part of the input instead of ripping up and retyping an entire card. Raster screens (Michael Noll at *Bell Labs* in 1968 [Ragnet, 2008]) allowed for real-time output (DA), transforming how people exchanged information using computers.

In summary, in all of these examples from interactive computing, the AD/DA pattern led to massive, disruptive change to both the field it affected and to the new user base it empowered. And in all these cases, there was a transition from industry to technology enthusiasts to consumers, which allowed the respective fields to assume the massive scale they have today.

If these previous developments should be any indication, they would suggest that personal fabrication will be going down the same route, leading to disruptive change as it reaches new users and ultimately consumers, at which point it could be expected to grow by several orders of magnitude.

1.4 How past AD/DA media transitioned to consumers

If we look at these examples of past AD/DA patterns, we see that the transition to consumers could only take place once conditions had been created that allowed the respective tasks to be performed by *consumers* — tasks previously performed only by professionals in industry or at least by technology enthusiasts. Overall, we argue it always took at least the following four elements to get the technology ready for consumers — and we already briefly mentioned them above.

1. *Hardware and materials.* The transition from specialized analog machines to AD/DA machines helped commoditize the hardware. In particular, the transition allowed individual technologies to "piggyback" onto personal computing. First, the personal computer inherently offered a wide spectrum of technology that one might not necessarily have built into the new machines otherwise, such as access to a backup

system and network access. These added benefits added momentum to the evolution of the new technology.

Second, the connection to the personal computer reduced the required upfront hardware investment. As more and more users owned personal computers in the first place, users only needed to buy a peripheral device in order to get access to the new technology. These peripherals could be simple and cheap, as they could use the resources of the personal computer. Early *PostScript* printers, for example, went as far as to leverage the personal computer for rasterizing the print image in the personal computer's RAM — which is exactly what we are seeing today with 3D printers that convert their document to a machine representation ("slicing") on the personal computer.

2. *Domain knowledge.* Industry professionals have expertise in the target domain, i.e., they know how to edit video, how to layout print, and so on. Consumers, in contrast, lack this expertise. So, in order to enable consumers to perform these tasks, software systems need to embody the lacking domain knowledge. For example, when movie editing transitioned to computers, the early systems were 1:1 replications of the editing environments common with physical videotape ([Quantel Harry] in 1985, and *Avid Technology's Avid/1 Media Composer* [3D Hubs] in 1987). Twenty years later, automatic video editing software (e.g., *Muvee's autoProducer* [Muvee]) automatically creates entire movies from users' raw footage based on default settings alone; more ambitious users can tweak this preliminary result, but they do not have to. In another example, Adobe Photoshop Elements retouches red eyes in photographs at the push of a button. Microsoft PowerPoint and Apple Pages allow users to create presentations and documents simply by filling in their contents into pre-designed templates. More recently, users have gained access to even more domain knowledge by downloading solutions from shared repositories [Lau et al., 2011].

3. *Feedback through interactivity.* Systems that embody domain knowledge can only go so far — there are always factors left that are not covered by the system, such as the user's assessment of the esthetics of a layout. Even with systems that embody various kinds of domain knowledge this continues to require exploration — trial and error. To reduce the number of iterations, software systems build on

the *what-you-see-is-what-you-get* principle (e.g., *Bravo*, 1974 [Hiltzik, 2000]) provide users with a sense of their final output along the way. During exploration, users receive immediate feedback, and are also able to *undo* steps.

4. *Machine knowledge.* The DA machines in AD/DA systems generally make the workflow easier. In particular, they eliminate the need for physical skill. Manually cutting film is challenging; so is manually creating a carefully aligned layout with scissors and glue. Digital video editing software and desktop publishing software eliminate these physical tasks, allowing everyone to produce a correct cut or a perfectly aligned layout. However, the new machines also bring their own challenges, as they require users to express their ideas in appropriate digital representations that they may not be familiar with. This is historically where an additional software layer comes in that embodies the required "machine knowledge."

Along the same lines, such software may also help users obtain the best results by providing additional expert know-how about the device. For example, while everyone may be able to print images, obtaining best results may require knowledge of the color spectrum (gamut) and resolution the printer is able to reproduce. Historically, additional software layers, such as *PostScript* would abstract these issues away by allowing users to produce machine-independent descriptions of print documents. Documents would be shared in this abstract format, knowing that the *PostScript* interpreter in the target printer would translate the abstract description into the best possible representation for the respective printer.

Combined, we argue that it is these four elements that allowed the previous AD/DA media to get ready for consumers.[2]

---

[2] Arguably, the same four elements were also necessary to allow personal computing as a whole to transition to consumers. Computing also started in industry and transitioned to technology enthusiasts (in the 1970s). If we look at personal computing in the hands of consumers today, we see the same four elements: (1) Consumer-friendly hardware, more and more in the form of self-contained "appliances", (2) Application programs that embody domain knowledge, including the programs we just discussed, (3) Feedback through interactivity, here in the form of the graphical user interface and its use of direct manipulation. (4) Operating systems that abstract away the necessity to know about the hardware. The resulting

1.5 Transitioning personal fabrication: the six challenges

Given the structural similarities to previous AD/DA media, we argue that it will take *exactly the same* four elements to transition from fabrication in industry to consumers (Figure 1.4): (1) *Hardware and material* developments will have to ensure that users will be able to fabricate the objects they want to create. (2) Systems will have to embody the *domain knowledge* (e.g., physics simulations) users need in order to obtain functional results. (3) Alternatively, objects designed with subjective (e.g., esthetic) considerations in mind are better assessed by human judgment. Accordingly, systems have to provide users with *feedback* along the design process. (4) Finally, systems will encapsulate the *machine-specific knowledge* required to fabricate the object on a specific machine.

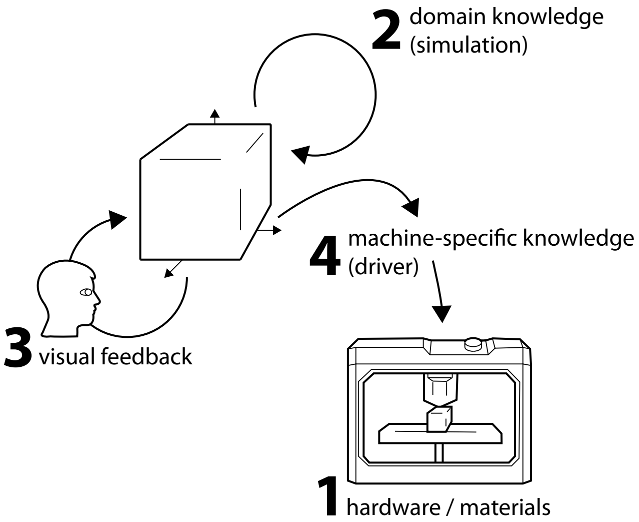

Figure 1.4: The four main challenges: (1) hardware/materials, (2) domain knowledge, (3) visual feedback, (4) machine-specific knowledge.

---

transition to consumers was, by the sheer numbers, clearly the biggest transition in the history of computing.

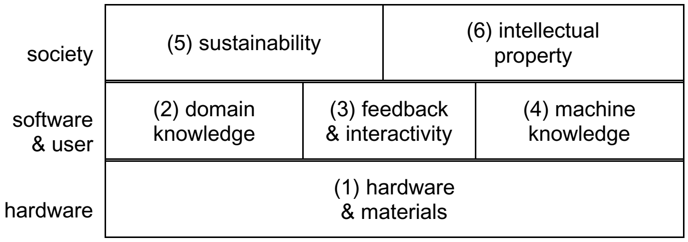

Figure 1.5: The six challenges of personal fabrication.

This means that if we as researchers and engineers want fabrication to make the transition to consumers and thereby empower hundreds of millions of new users, these are the conditions *we* need to create.

In addition to the four challenges discussed above, we see two additional challenges: (5) *sustainability*, including factors such as trash, material, and energy consumption and (6) *intellectual property*, including approaches that tackle the difficulties resulting from the sharing of protected designs.

While these two challenges may not be necessary for AD/DA fields to reach consumers in the first place, they tend to emerge as the field grows in size. It thus seems safe to expect that fabrication will face these issues as well eventually. We therefore argue that we should consider these challenges now — before they have a chance to grow out of proportion.

In Figure 1.5, we summarize all six challenges grouped into a *hardware* layer at the bottom, a *software and user* layer in the middle, and a *society* layer on top.

Naturally, the main challenges for researchers in human computer interaction can be found in the *user* level in the middle of our chart, which is all about establishing a successful connection between users and the system and more specifically about abstracting away any challenges that could prevent consumers from performing the work traditionally performed by experts. Given that the transition of personal computing to consumers (“discretionary use”) has been one of the core concerns of the HCI community for decades, HCI researchers are well equipped to tackle these challenges.

Our survey of the related work, however, shows that HCI researchers are making contributions to all three levels. The *hardware and materials* level offers plenty of opportunity not only for mechanical engineers and material scientists, but also for HCI researchers with a hardware angle (as found, for example at the *User Interface Software and Technology* (UIST) conference [Hudson, 2014]). Questions involving the societal impact of personal fabrication provide a great challenge for researchers on the empirical and ethnographic side of HCI.

In addition, we see researchers in computer graphics making major contributions around various challenges, but especially around the challenge of embodying *domain knowledge* and *machine-specific knowledge* into software. Projects in this space not only involve the simulation of forces, but also build heavily on processing 3D geometries, which makes computer graphics researchers particularly well equipped to tackle this class of problems. However, similar to researchers in HCI, researchers in computer graphics have tackled challenges in several of the other categories as well.

In the following chapters, we try to obtain a deeper understanding of the state of the art with respect to the six challenges by surveying the related work on personal fabrication. If we look at some of the main conferences on human–computer interaction, we see that research on personal fabrication is just starting out, but is growing quickly (e.g., CHI 2013 first five papers on fabrication, CHI 2016 seventeen papers, UIST 2012: first three papers on fabrication, UIST 2016: a quarter of the program was on fabrication).

We present the work grouped by the challenge it addresses. For each challenge, we relate it to previous instances of the AD/DA pattern and use this analogy to extrapolate the current trends towards the questions and opportunities researchers in personal fabrication are about to encounter. While we focus on human–computer interaction and computer graphics, we also include selected works from adjacent fields such as mechanical engineering, material science, and robotics.

2

Hardware and Materials

In order for a field to transition to a digital workflow, the involved AD and DA converters have to be able to translate all relevant aspects of the involved artifacts to data and back. As outlined in the introduction, for physical objects the AD component is a 3D scanner, the DA component of a fabrication device. Achieving a "perfect" conversion that would make a scanned and refabricated object indistinguishable from its original is still subject to research at this time. In this chapter, we survey the current state of the art and point out the resulting challenges.

We discuss recent developments in the order suggested by Figure 2.1. We begin with techniques that create a specific appearance of 3D objects, such as achieving a desired *shape, color*, and *reflectance*. We then move on to discuss techniques that attempt to reproduce the tactile qualities of objects, either by printing *tactile textures* or by using *soft materials*. Finally, we survey different techniques to make an object perform a desired function. This typically includes novel printing materials, such as *conductive* materials that allow for printed electronics [Ahn et al., 2011] and *optical clear* materials that allow printing light pipes [Willis et al., 2012]. Functional properties, however, can also be

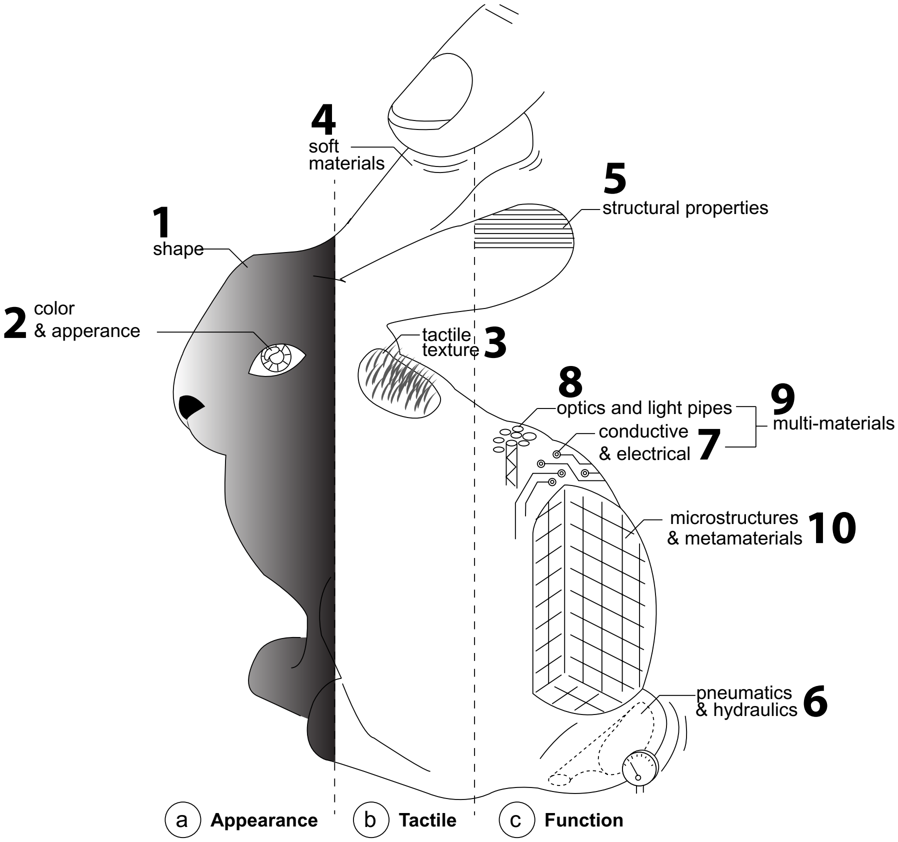

Figure 2.1: Fabrication machines allow users to design three aspects of an object: (a) appearance, (b) tactile qualities, and (c) function.

achieved by designing the *internal structure* of an object. For instance, redistributing an object's infill allows shifting its center of mass, which can be used to make it stand [Prévost et al., 2013]. Finally, by creating objects from repeating cell structures researchers have shown how to emulate a range of material properties from a single material (so-called *metamaterials* [Lee et al., 2012]).

Before looking at each technique in detail, we provide a short overview of the most relevant fabrication technologies to cover the necessary background knowledge.

2.1 Personal fabrication technologies

Personal fabrication technologies can be grouped into three main categories as illustrated by Figure 2.2: (a) additive, (b) subtractive, and (c) formative fabrication.

Subtractive fabrication technologies, such as milling and laser cutting, cut objects from a block or sheet of material. One of the key benefits of this approach is that it preserves most of the properties of the materials — the structure of wood, for example, persists, when a part of wood is cut into pieces. However, since each block or sheet is made from a single material, the process is generally limited to one material per part. In addition, subtractive technologies have only limited abilities to create structures inside 3D objects, as, for instance, the milling head has trouble reaching inside the volume.

Similar to subtractive fabrication, *formative* fabrication, such as vacuum forming and blow molding, uses a single sheet or block of material. Rather than cutting it, however formative tools *reshape* the material into a new form generally by stretching it. The main benefit of this approach is that it allows modifying objects very quickly. However, formative fabrication is subject to the same limitations as subtractive fabrication and it is even more limited in terms of its abilities to create internal structures.

Unlike subtractive and formative fabrication, *additive* fabrication techniques, such as 3D printing, generally start with an empty build platform. On this platform additive fabrication creates objects by adding material typically voxel-by-voxel and layer-by-layer as in the

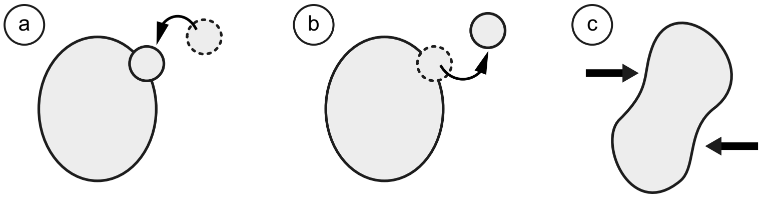

Figure 2.2: The three main fabrication processes (a) additive, (b) sub-tractive, and (c) formative.

case of the most common 3D printers. This process enables many degrees of freedom with respect to designing both the overall shape and the internal structure of objects. However, this particular versatility also comes at a cost: 3D printing requires the involved materials to be broken down into the form of filament, powder, or resin, which causes the material to lose most of its structure (e.g., consider an object printed from wood filament that contains grinded wood particles compared to the same object made from a real piece of wood with natural structures).

In this chapter, we focus mainly on additive fabrication techniques, and in particular 3D printing due to the flexibility it offers. There are several different 3D printing technologies, such as those that extrude filament through a hot nozzle (*Fused Deposition Modeling*), sinter powder using heat (*Selective Laser Sintering*), bind powder using a liquid (*Inkjet 3D Printing*), solidify liquids using light (*Stereolithography, Polyjet Printing*), or cut layers into shape before laminating them onto each other (*Layered Object Manufacturing*). See Thompson [2007] for a more detailed explanation.

With this overview in mind, let us look at what these technologies can accomplish today and the improvements that came out of recent research.

2.2 Shape

The most obvious design dimension fabrication machines offer is to allow users to design the *shape* of objects. Current commercial high-resolution 3D printers already exceed the resolution of the human eye (e.g., *Objet Connex*, 16 μm layers = 1,600 dpi). Researchers have even been able to 3D print at resolutions of up to 1 μm (Figure 2.3) allowing users to make entire objects smaller than a human hair. This is enabled through a process called Two photon lithography, and already commercially available through companies such as Nanoscribe. These new high-resolution printers eliminate all visible artifacts that were common with old low-resolution 3D printers, such as being able to see the different layers stacked onto each other. Consequently, current 3D

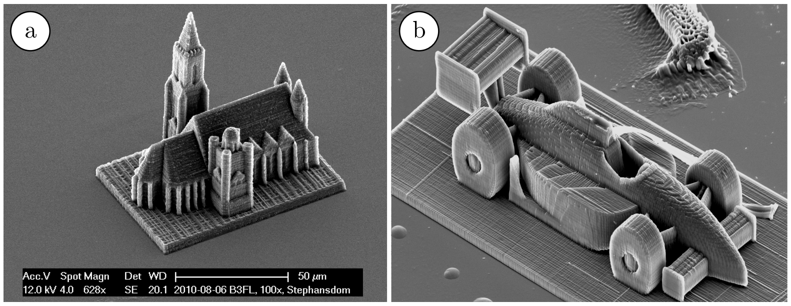

Figure 2.3: Current high-resolution printers already exceed the resolution of what the human eye can see. (a, b) This 3D printed building and racecar are around the size of a human hair [Two photon lithography].

printers offer everything required for fabricating objects whose shape looks right.

2.3 Color and appearance

Coloring a 3D object in its entirety involves coloring the surface and its volume. There are only few technologies today that can achieve a full color spectrum: One such technology is 3D printers that use regular office paper as their printing material [e.g., MCor Technologies]. These devices use a 2D printer to create the color on a paper layer, then cut and laminate the layer onto the existing stack (based on *layered-object manufacturing*).

Similarly, good results can be achieved for powder-based 3D printers with an inkjet head that release a binder to locally harden the powder (Figure 2.4a). Differently colored binders can be used like the CMYK cartridges in 2D printing to achieve any desired color.

The color abilities of other 3D printing technologies are much more restricted. However, various research projects developed techniques to increase their color abilities. For instance, FDM 3D printers generally extrude only a single strand of filament per printing nozzle and are thus limited to a one color per nozzle. Reiner et al. [2014] show how to achieve a color gradient by interleaving filaments (Figure 2.4b).

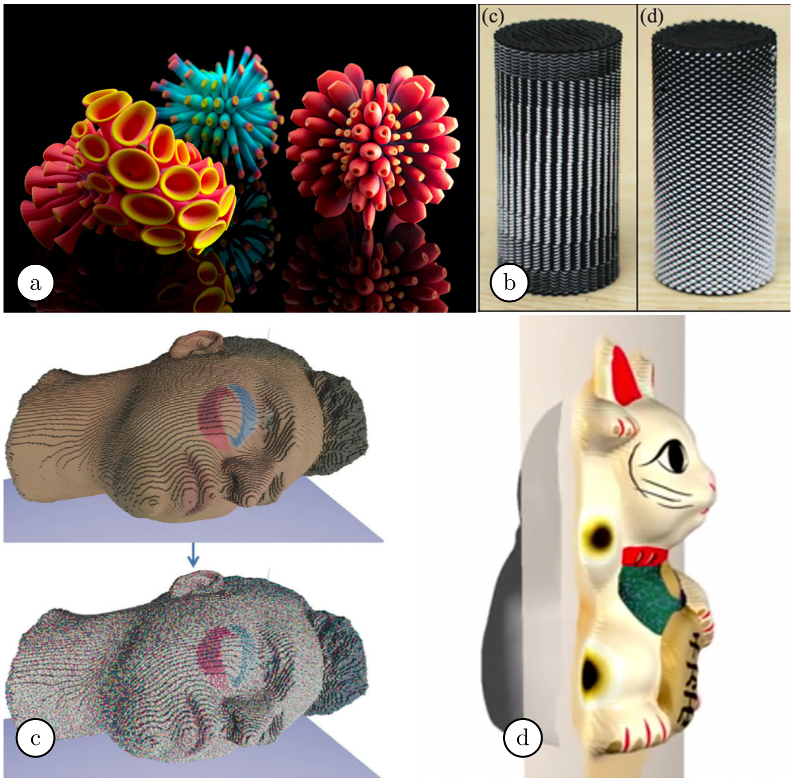

Figure 2.4: Color: (a) full color powder-based 3D printing with inkjet head (3D Systems ProJet × 60 series), (b) interleaving two filaments with FDM 3D printing [Reiner et al., 2014], (c) half-toning with poly-jet 3D printing [Brunton et al., 2015], (d) post-processing color with hydrographic printing [Panozzo et al., 2015].

Similarly, in polyjet printing the range of colors is limited by the available polymers. Brunton et al. [2015] increase the spectrum by half-toning the available colors (Figure 2.4c).

Another way to color objects is to first fabricate the shape, and then color the object in a post-process as shown in Figure 2.4d. This approach generally requires the desired color pattern to be printed on a 2D sheet, which is then applied to the surface, for instance, by dipping the 3D object into a liquid to pick up a transcription film that floats on the surface (*hydrographic printing* [Panozzo et al., 2015]). Besides

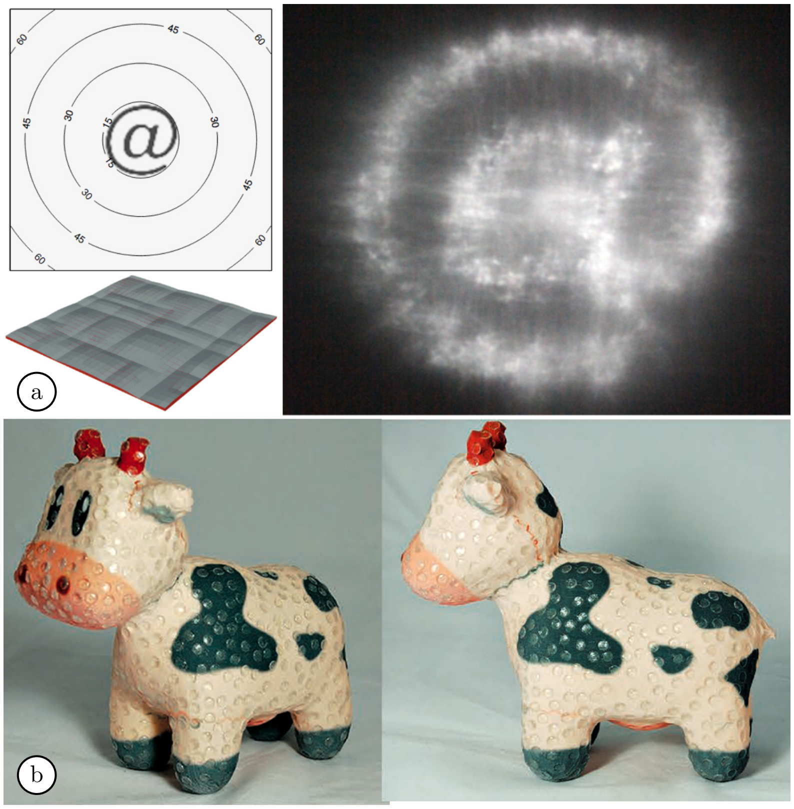

Figure 2.5: Reflectance: (a) milling small microstructures into a surface [Weyrich et al., 2009], (b) using 3D printing [Rouiller et al., 2013].

post-processing color, researchers also showed how to pre-process color: In *Computational Thermoforming* [Schüller et al., 2016] an image is printed onto a thermoplastic sheet, which in a subsequent thermoforming process is shaped into a 3D object that carries the texture. These approaches, however, only color the surface.

In addition to color, researchers have explored how to fabricate objects that reflect light in a particular way (Figure 2.5).

Weyrich et al. [2009] show how to achieve a desired homogeneous reflectance by milling small microstructures into the surface of an object (Figure 2.5a). Rouiller et al. [2013] extend this concept to 3D printing

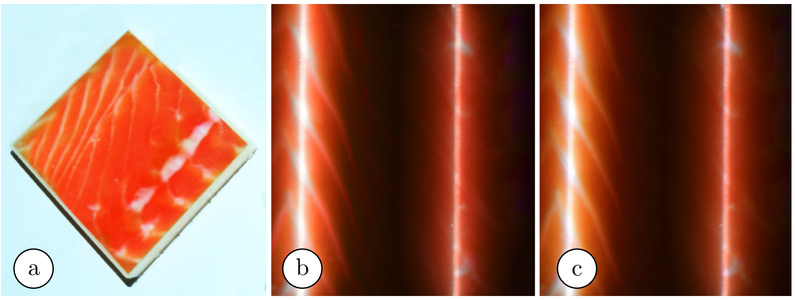

Figure 2.6: Refractance: (a) 3D printed piece of salmon. It's inner structure scatters light (b) in the same way as (c) real salmon when photographed with diffused lighting [Dong et al., 2010].

and expand it to spatially varying reflectance: The black dots on the toy cow shown in Figure 2.5b, for example, are highly reflective while the nose is matt diffuse.

Finally, researchers have explored how to control the scattering of light, i.e., how the light is refracted *inside* the material (Figure 2.6). By stacking layers of materials with different scattering properties, Hašan et al. [2010] and Dong et al. [2010] demonstrate how to create realistic appearances. While their approaches create a homogeneous scattering behavior, Peers et al. [2006] show how to vary scattering across an object's volume.

For more information on appearance fabrication please see Hullin et al. [2013] for an overview.

2.4 Tactile textures

In addition to defining an object's visual appearance, users may want to design how an object feels when touched (Figure 2.7).

The resolution of today's 3D printers is sufficient to print even the smallest features humans can perceive. To help designers enhance 3D models with surface textures, researchers in HCI have provided specialized design tools. For instance, *Haptic Print* [Torres et al., 2015] allows users to select from different surface textures, which are then automatically applied to the object's surface.

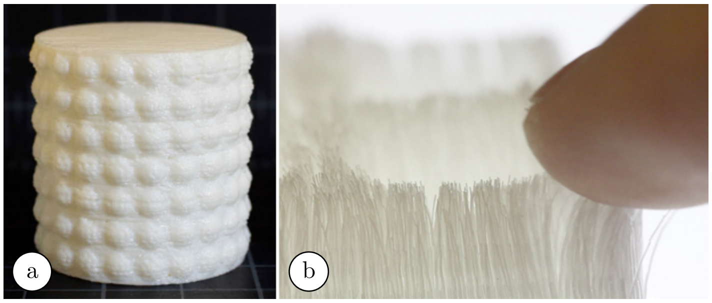

Figure 2.7: Surface texture: (a) *Haptic Print* [Torres et al., 2015] enables designers to apply surface textures to 3D model, (b) *Cillia* [Ou et al., 2016] prints high-resolution textures without the need of support material.

Another track of research explored how to texture 3D objects with bristles and hair. On FDM 3D printers, the key idea is to extrude small amounts of material while pulling away quickly, thereby stretching the extruded material (*3D Printed Hair* [Laput et al., 2015]). In contrast to regular layer-wise printing, this approach works without support material. On stereolithography 3D printers researchers achieved bristle-like features by printing a series of layers each of which shows dots of decreasing size [Ou et al., 2016].

2.5 Soft materials

The materials used with 3D printing generally range from stiff (*ABS*, *PLA*, as well as metals and glass) to elastic (silicones, e.g., Ninjaflex). While silicones are certainly soft (i.e., compress when pressed), what users perceive as soft is not just a matter of the material, but also of the 3D structure of the material, in particular around its surface.

Since these microscopic structures are still hard to replicate with 3D printers, researchers proposed fabricating with materials that have such "soft" surface textures (e.g., felt). Traditional fabrication devices, however, cannot process these materials and thus processing these materials requires developing new fabrication machinery.

In *Printing Teddy Bears*, for instance, Hudson [2014] shows how to 3D print using soft thread. A felting needle on a custom 3D printer

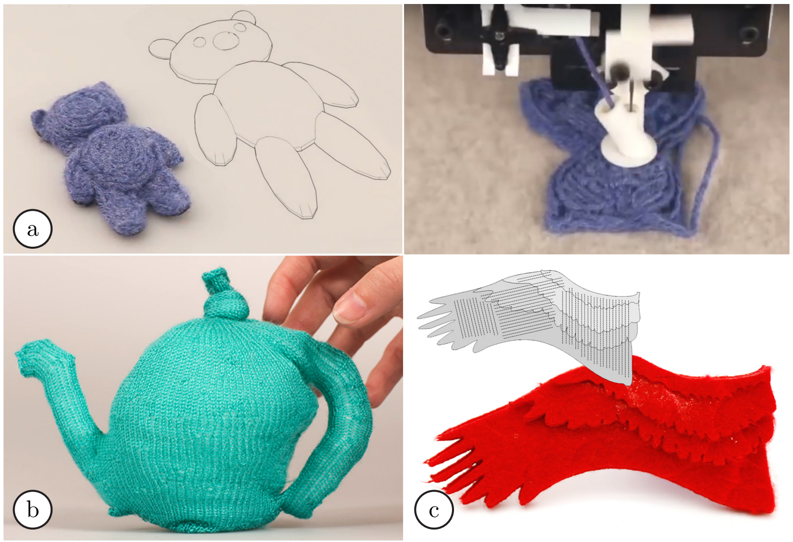

Figure 2.8: Soft 3D Printing: (a) printing with a felting needle [Hudson, 2014], (b) a 3D object created with a knitting machine [McCann et al., 2016], (c) cutting fabric from a roll and laminating it [Peng et al., 2015a].

entangles thread, resulting in a felt-like structure (Figure 2.8a). Similarly, Sosanya creates soft objects by expanding a loom to allow it to weave in three dimensions *(3D Weaver* [Sosanya]). In contrast, McCann et al. [2016] use an existing knitting machine, but provide a design software that makes it easy to translate a custom design into machine instructions for 3D knitting (Figure 2.8b).

Instead of using thread, Peng et al. start with sheets of felt in the first place (Figure 2.8c). Their custom device first laser cuts individual felt shapes from a roll of fabric; and after cutting a layer glues it onto the layer stack on the build plate, thereby assembling the 3D shape. They refer to this special form of *layered-object-manufacturing* as *Layered Fabric Printer* [Peng et al., 2015a].

In their quest to produce soft surfaces, researchers have also explored how to use living organisms as part of fabrication machinery. In *CNSilk* [Tsai et al., 2012] researchers use silkworms that spin silk around a template structure to create soft objects.

In contrast to the above techniques, Fabrican creates soft surface textures using a liquid material that develops a soft structure by cross-linking between fibers when being sprayed.

2.6 Structural properties

While the technologies and techniques discussed above are generally relevant to 3D objects for which appearance and tactile qualities are important, 3D printing also opens up exciting opportunities in the realm of *functional* objects, such as objects that implement *mechanical* functions.

One of the key challenges when creating mechanical devices, such as a walking automata [Jansen], is that they have to withstand forces. These tend to be substantially higher than with objects for which only appearance and feel is important.

While studies find that 3D printed objects can generally perform similar to their mass-manufactured counterparts [Tymraka et al., 2014], one issue in layer-based 3D printing is that layers tend to delaminate more easily when tension is applied perpendicular to the layers. A more recent 3D printing process called *continuous liquid interface production* (commercialized with Carbon3D [2015]) addresses this issue. The technology builds on the stereolithography 3D printing process, but cures material continuously. This results in objects of equal tensile strengths along all three dimensions (Figure 2.9).

2.7 Pneumatics and hydraulics

While the vast majority of machinery created using 3D printers builds on mechanical mechanisms, such as gears and levers, 3D printed *pneumatics* can be used to create soft machines, such as soft robots. This is generally accomplished using rubberlike materials (e.g., *Tango* printing materials from Stratasys). In addition to fabricating air chambers with 3D printing, air chambers can also be fabricated using a wide range of other processes, such as silicone casting (*PneUI* [Yao et al., 2013]).

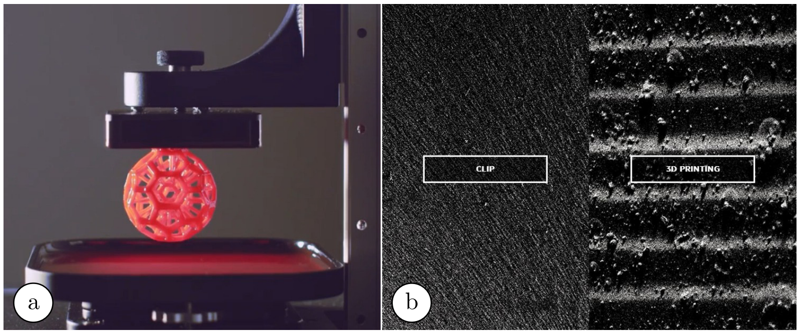

Figure 2.9: (a) To prevent delamination of layers along the printing direction, *CLIP* [Carbon3D, 2015] cures layers continuously rather than one-by-one. (b) The left side shows CLIP's continuously cured object; the right side shows traditional layer-wise 3D printing.

Pneumatic devices generally use one or more compressors to build up air pressure. Computer-controlled valves then determine where the air is directed, i.e., which 3D printed air chambers are inflated using the compressed air. By embedding air pressure sensors these pneumatic devices can also serve as input components, for instance, to detect if a user squeezes a soft robot's arm [Slyper and Hodgins, 2012]. The shape of the air chamber determines the airflow and thus allows pneumatic devices to differentiate between manipulation types. Vázquez et al. [2015] extend this principle to also sensing how much force the user applies.

Researchers have also showed how to print hydraulic devices that allow handling bigger forces. Roumen et al. [2016] for instance, print hydraulics by filling the chambers with water droplets *during* 3D printing (Figure 2.10).

2.8 Conductive and electrical

3D-printable conductive materials allow integrating the functionality traditionally offered by printed circuit boards.

The development of 3D FDM-printable conductive materials required solving several challenges. Initially, high-conductivity silver

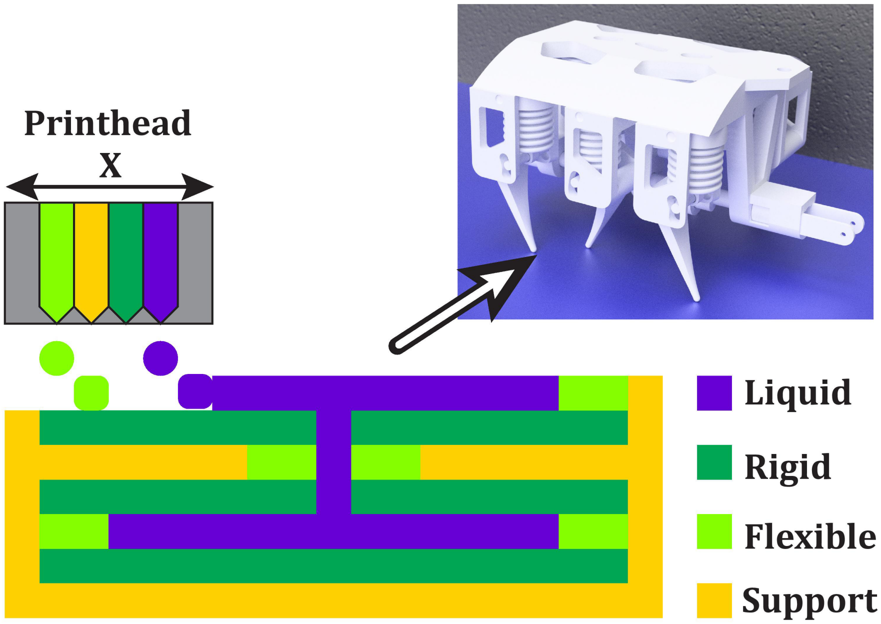

Figure 2.10: Printing liquids and solids in a single process to create this hydraulic robot [Roumen et al., 2016].

inks had to be extruded at high temperatures, which tended to melt the plastic materials typically used for the remainder of the object. Composite materials, in particular plastic mixed with *Carbon black* filler (*CarboMorp*h [Leigh et al., 2012]) helped overcome the issue, as they can be extruded at much lower temperatures. This material, however, offered only low conductivity limiting its use to driving low-power components, such as LEDs (Figure 2.11a). More recently, Ahn et al. [2011] solved this problem by developing a highly conductive silver ink that can be extruded at low-temperatures. The research is now commercialized with the 3D printer Voxel8, the first consumer printer for printing electronics (Figure 2.11b). Lopes et al. [2012] made conductive printing available to stereolithography 3D printing by adding a head that locally sprays conductive paint onto the object while it is being printed. With inkjet 3D printing, the next step might be to 3D

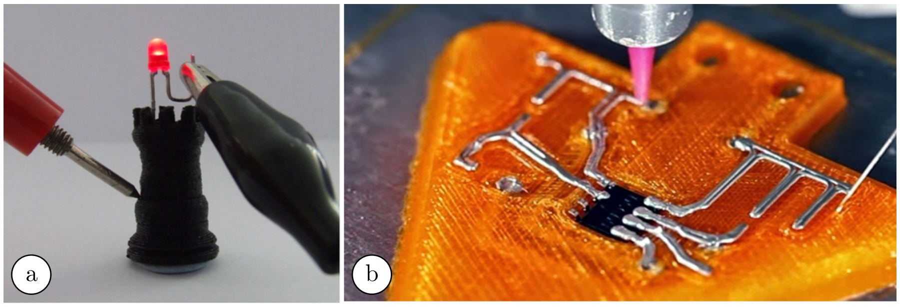

Figure 2.11: Electronics: (a) low-extrusion temperature, low-conductivity LEDs (*CarboMorph* [Leigh et al., 2012]), (b) low-extrusion temperature, high-conductivity [Voxel8].

print stretchable electronics, such as those outlined in the 2D fabrication process *Stretchi* [Wessely et al., 2016].

While the technology described above allows fabricating PCB-like functionality, the electronic components, such as resistors, capacitors, and diodes, are still mass-produced and (typically manually) inserted after or during the 3D printing process.

Ultimately, we see 3D printing not only subsume printed circuit boards, but also electronic components to achieve 3D objects with integrated circuits all in a single integrated process. This vision stands at the very beginning, but researchers have started to build prototypes to illustrate the potential of integrated printing of electronics (see examples in, e.g., *SteelSense* [Vasilevitsky and Zoran, 2016]). In addition, recently, researchers started to 3D print the first few types of electronic components: [Peng et al., 2016a], for instance, demonstrated how to create simple coils and motors by combining FDM 3D printing and a custom print head that lays out either copper or soft iron wire. In addition, Lewis and Ahn [2015] have shown early designs for 3D printing LEDs.

2.9 Optics and light pipes

While clear materials have decorative applications (e.g., predefined light patterns called *caustics* [Schwartzburg et al., 2014]), they also allow printing light pipes. *Printed Optics* [Willis et al., 2012] shows

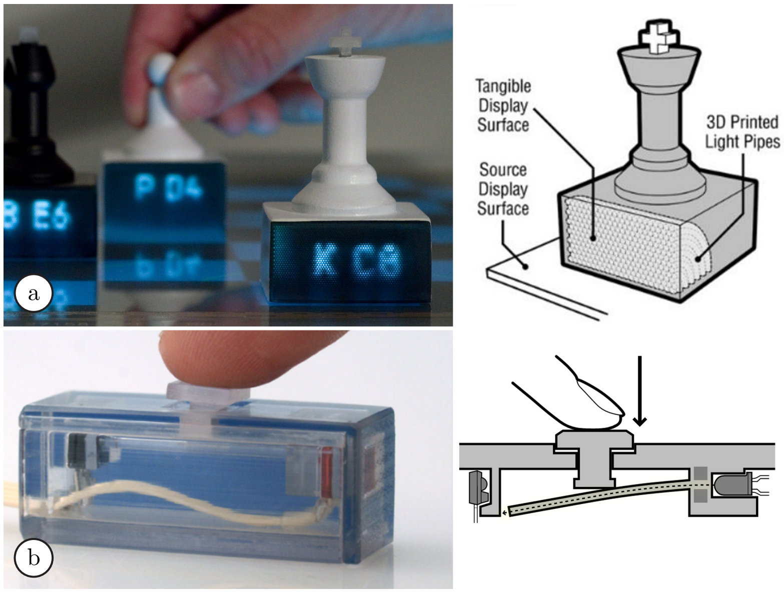

Figure 2.12: Printed Optics [Willis et al., 2012]: (a) novel display element, (b) button based on optics, i.e., when the button is pushed down, the light from the LED no longer hits the sensor.

how to use this property to integrate display and sensor elements into 3D objects (Figure 2.12).

Computational Light Routing [Pereira et al., 2014] provides an algorithm that optimizes the light transmission (i.e., minimizes fiber curvature and maximizes fiber separation, while taking into account fiber arrival angle).

2.10 Multi-material printing

The objects that surround us every day are diverse and so is the range of materials they are made of. This presents a challenge as the number of materials available for 3D printing and especially for a given 3D printer tends to be small.

A basic approach to the problem is to increase the number of materials available on a given 3D printer (e.g., 10 different materials with *MultiFab* Sitthi-Amorn et al. [2015]). However, this process does not scale as the number of extruders quickly becomes large and having all the materials on stock is not feasible.

Researchers showed how to tackle both problems by creating additional materials on the fly. The key idea is to create new materials by *mixing* a small set of base materials (so-called Digital Materials). This process allows, for instance, mixing soft and hard base materials in order to produce a material of intermediate hardness. The mixture can vary for every voxel in the print, resulting in so-called *functionally graded materials* (Figure 2.13).

Similar problems exist for other fabrication techniques, such as laser cutting (only a single material sheet at a time). With *Foldem* [Perumal and Wigdor, 2016] show how to use a sheet consisting of different materials stacked on top of each other and the laser's capabilities to selectively cut into these layers to create a range of different physical properties. However, the capabilities are still lacking behind compared to what is possible with multi-material 3D printing today.

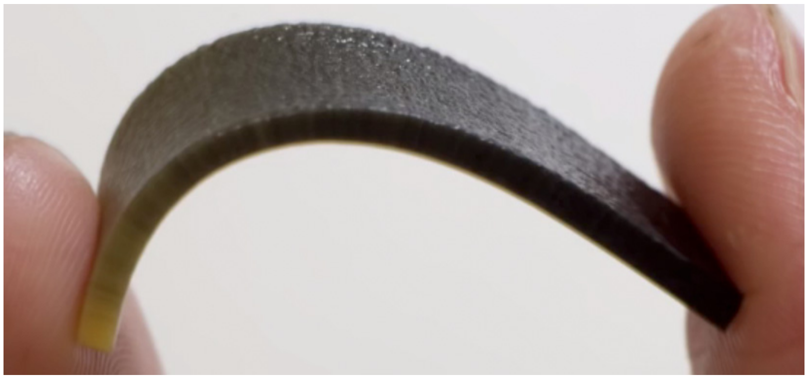

Figure 2.13: A digital material printed with MultiFab [Sitthi-Amorn et al., 2015], it has a gradient from hard to soft here illustrated by also mixing two col-ors at different proportions.

2.11 Infill, microstructures and metamaterials

Different materials are just one approach to varying the physical properties of a 3D printed object. The other approach is to vary the infill, an approach that is surprisingly versatile.

Early work on infill optimization aimed at optimizing objects' strength-to-weight ratio. In its simplest form, this means to fill the inside of objects with a sparse infill pattern, such as a honeycomb grid, in order to achieve reasonably high stiffness while saving material and printing time. This approach can be optimized further by replacing the honeycomb grid with a custom grid that considers an object's shape and the mechanical load applied to different regions of the object [Lu et al., 2014]. As shown in Figure 2.14, material can be further reduced by leaving out infill altogether and instead printing skin-frame structures underneath an object's surface Wang et al. [2013].

More recently, researchers started to use the infill to control object properties that can be perceived from the outside (Figure 2.15). For instance, to prevent an object from falling over, researchers used the infill to shift the object's center of mass (*Make it stand* [Prévost et al.,

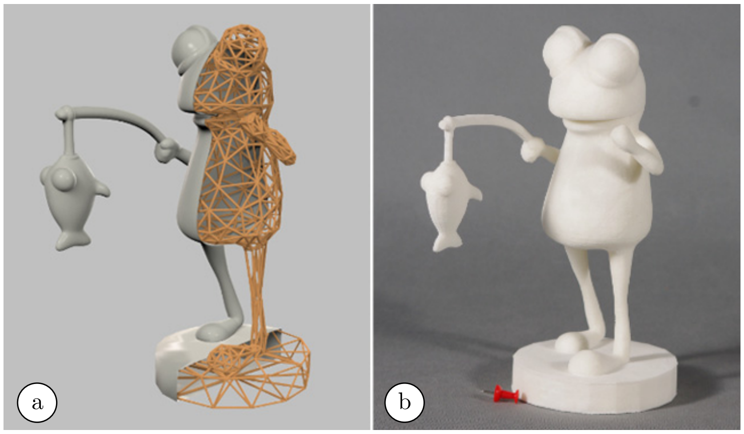

Figure 2.14: Infill: Structures printed underneath the surface are sturdy but do not require as much printing material and thus also less printing time [Wang et al., 2013].

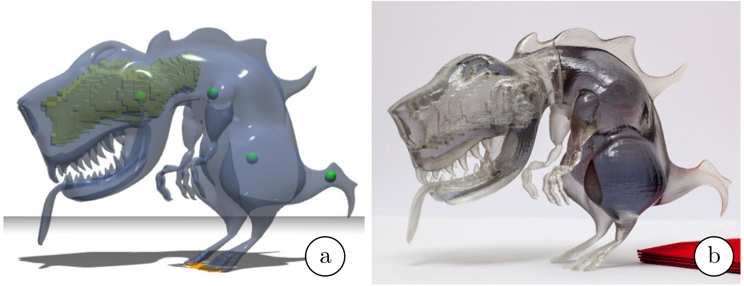

Figure 2.15: Modifying infill in order to balance objects (*Make it Stand* [Prévost et al., 2013]).

2013]). The same approach can be applied to adjust an object's principal axis of rotation to allow it to spin (*Spin-it* [Bächer et al., 2014]), or its buoyant equilibrium to make it float upright (*Buoyancy Optimization* [Wang and Whiting, 2016]).

While the examples in the previous paragraph only control a few degrees of freedom of the object, controlling infill can be pushed much further. The general idea, initially developed in mechanical engineering [Lee et al., 2012], is to subdivide the inside of an object into 3D cells. Since each cell can now be designed individually, the resulting objects literally offer thousands of degrees of freedom.

Bickel et al. [2010] show how to use such structures to reproduce the stiffness distribution of an existing shoe with the limited set of materials available on a 3D printer: They stack 2D layers of different 'hole' patterns to vary how the material compresses in z-direction — layers with big holes compress easier than those with small holes. The concept of varying stiffness was pushed further by Schumacher et al. [2015]: They used cells on a 3D grid, each cell having a different softness/hardness (Figure 2.16).

Panetta et al. [2015] use a similar approach and explore the space of different elastic textures. Their softest pattern is over a thousand times softer than their stiffest pattern, and the patterns' ability to expand (a.k.a. *Poisson ratio*) ranges from below zero to nearly 0.5. Instead of using regular cells, Martínez et al. [2016] explore the space of Voronoi patterns that lead to foam-like structures.

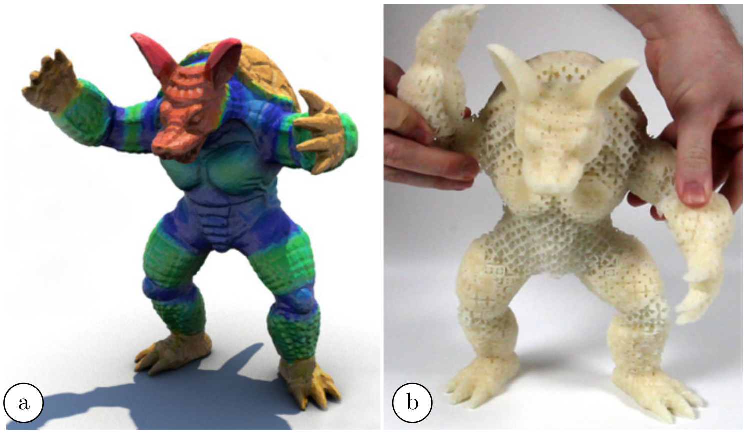

Figure 2.16: (a) Schumacher et al. [2015] achieve varying degrees of softness by (b) changing the microstructure of the material.

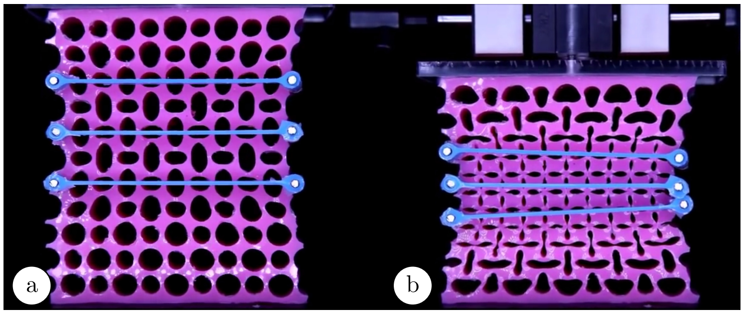

Figure 2.17: Metamaterials enable properties not found in nature. The metamaterial shown here collapses when compressed [Florijn et al., 2014].

A particularly exciting line of work in this area pursues the objective of creating objects which properties cannot be found in nature. These have been referred to as *metamaterials* (Figure 2.17).

Examples include materials that expand in two dimensions upon one-dimensional stretch (a.k.a. *auxetic materials* [Mir et al., 2014]) and materials that "pull" in the direction of compression rather than resisting it (a.k.a. *negative stiffness* [Rafsanjani et al., 2015]).

While most metamaterials have been explored only on a 2D grid, Shim et al. [2012] provide first examples for 3D geometries. For a comprehensive overview of different *mechanical* metamaterials and their properties, see Lee et al. [2012] and Elipe and Lantada [2012]. Metamaterials also exist for other properties, such as *thermal, electromagnetic* and *acoustic*, we refer to Kadic et al. [2013] for an overview.

Ion et al. [2016] take a different perspective on microstructures by thinking of the resulting objects as machines rather than materials. Their main idea is to group cells into simple mechanisms that transform movement and forces. Such basic mechanisms can then be further assembled into simple, yet self-contained machines, such as a door latch or a pair of pliers (Figure 2.18). Unlike traditional machines, such *metamaterial mechanisms* can be fabricated as one single part, thus require no assembly.

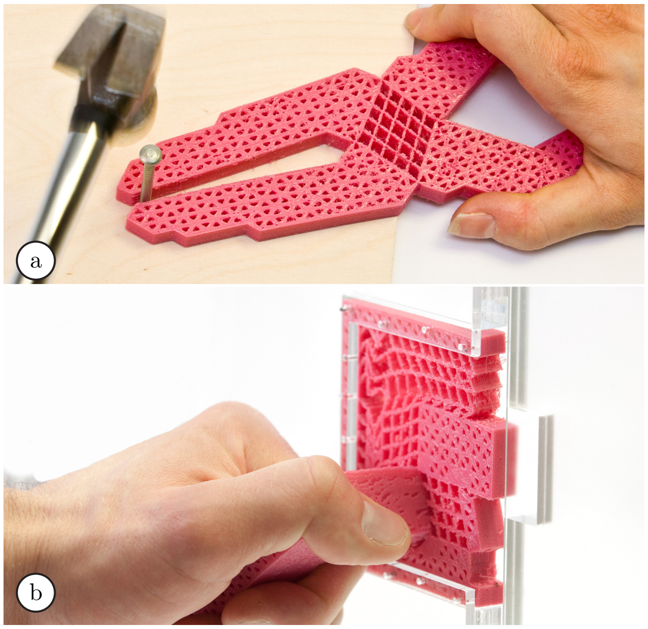

Figure 2.18: Metamaterial Mechanisms [Ion et al., 2016] transfer input movement and forces into output movement and forces using mechanisms based on cells: (a) a pair of pliers, (b) a door latch.

One of the strengths of creating object properties by means of microstructures is that this can be achieved using a single material. As a result, such objects can typically be fabricated even on simple machines, such as FDM printers with only a single extruder.

2.12 Conclusion and open research questions

As discussed in the introduction to this chapter, in order for a field to transition to a digital workflow, the involved AD and DA converters have to be able to translate all relevant aspects of the involved objects to data and back.

Since the field is large, we had to limit the survey in this chapter to the research efforts most closely linked to personal fabrication technologies, and in particular 3D printing. However, attempts to arrange physical matter in high detail have been made by various other research fields, such as *nanotechnology*.

Zooming out, we can say that research has covered most qualities required to fabricate objects for which *appearance* and *feel* is important (one could call these "decorative" objects). The reason is that the qualities required for decorative objects only have to appeal to the human senses and thus only have to match human perceptual abilities. However, the fabrication machinery required to fabricate these qualities is typically only found in industrial-strength machines. Arguably, the biggest challenge for fabricating decorative objects thus lies in bringing high-end qualities, such as high resolution and digital materials, to consumer 3D printers.

In contrast, with respect to functional objects, much more research is required. We see research opportunities particularly in 3D printing electronics. While the current evolution is fast-paced, we are today at a point at which we can barely fabricate conductive paths. Next, there will be plenty of research opportunities in developing methods for 3D printing mechanical and electronic components, such as resistors, capacitors, and transistors.

For the long-term future, we may ask ourselves what it will take to replicate various types of consumer objects. Ultimately, this would require manipulating physical matter on a molecular level Service [2015], but there will certainly be many steps in between.

In order to get a sense for what is next, we can use the simple and somewhat naïve approach of mapping what we see in 3D printing today to the 20th century technologies 3D printing emulates. We might point to injection molding (invented in 1872), printed circuit boards (1903), and the first LEDs (1927). And then consequently, we might choose to follow down the electronics path and put transistors (1947) and integrated circuits (1949) on our research agenda. Or we try to continue the work on light pipes (1920s), as well as pneumatics and hydraulics both of which started in the 17th century.

3

Domain Knowledge

In order to enable the transition from industrial use and maker use to consumer use, AD/DA software systems need to eliminate the need for expertise by embodying all necessary *domain knowledge.*

In the traditional analog workflow, users had to know where to place a cut in the film footage to create an exciting movie or how to layout text and images for a given type of content to obtain an appealing page. Today, this knowledge is still required when trying to produce high-end contents; however getting started is easy as digital systems embody this domain knowledge. For a digital personal fabrication system, the software may need to know how to design an object that is structurally sound or how to solve a given mechanical problem.

There are multiple ways to offer such domain knowledge to the user as illustrated by existing software from other domains, such as desktop publishing or video editing. In its simplest form, software offers known solutions as a starting point, e.g., by providing libraries of parts. Slightly more advanced systems offer templates that users can customize. Other approaches allow users to simulate the final outcome at various stages along the process. Finally, such software may even autonomously analyze what users are creating in order to support them with suggestions, as is the case with spell/grammar checkers in desktop publishing software, color gamut warnings in image processors, and clipping warnings in audio processing packages.

3.1 Objectives: Domain knowledge in personal fabrication

In personal fabrication, the tools and techniques that embody domain knowledge are an active field of research. While computer graphics has examined many techniques that allow users to 3D model the shape and movement of *virtual* characters, physical modeling is different. Unlike virtual objects that tend to live in a world governed by simplified and idealized rules, fabricated objects are subject to the laws of actual physical reality. Designing for fabrication thus means designing objects that will perform in a certain expected way in the physical world.

Besides helping users to integrate parts and mechanisms required for *motion*, modeling for physical output requires dealing with *forces.* Thus, when we talk about *domain* knowledge, the *domains* we are referring to include those disciplines that deal with forces, such as *structural engineering* and *mechanical engineering.*

For inexperienced users, understanding, estimating, and designing with forces can be prohibitively difficult. To illustrate this point, let us consider the chair you are using in your office: How much torque can the chair's backrest withstand before it breaks? If made twice as thick, how much more force will it withstand now? And if subjected to too high a load, which part will break first and why?

One of the reasons why users may find it hard to grasp forces is they are invisible, so that users experience them only indirectly. This makes forces different from shape and motion, which users can perceive directly and in their daily lives.

The difficulties of designing with forces can be directly observed by looking at shared online repositories: The majority of 3D models found in the online database *Thingiverse* for example, are *decorative* objects, i.e., objects that are desirable because of their shape and appearance, but that exhibit no *functional* behavior.

Existing structural and mechanical engineering software tools, such as Autodesk Inventor and SolidWorks, give engineers full control over

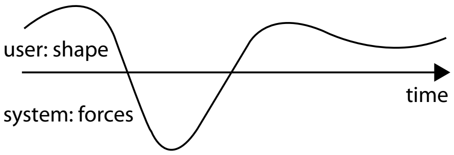

Figure 3.1: One successful approach to building systems that provide users with domain knowledge is to let users specify shape and motion, then handle all resulting forces under the hood and respond back to the user in terms of updating the shape.

how to solve a particular engineering challenge including the involved forces. On the flipside, these tools were designed for industrial users that have an in-depth understanding of the domain. Since consumers lack this understanding, these tools in their current form are unlikely to be picked up by them.

The challenge is thus to develop tools that allow consumers to design objects involving forces, yet to abstract away the domain knowledge that mechanical engineers tend to rely on. Achieving this would enable consumers to design complex mechanical objects and machines.

Figure 3.1 illustrates how systems embody domain knowledge involving forces. Such systems let users specify the shape and motion of a desired object. Then the system simulates the physical behavior including forces under the hood. The system then responds by critiquing the user's design by suggesting where shape needs to be changed and how to achieve the desired physical behavior (e.g., *Design by Example* [Schulz et al., 2014]). Alternatively, the system can automatically adjust the user's design to make it comply with the objectives and with respect to forces (e.g., *Pteronyms* [Umetani et al., 2014]). The benefit of conducting all communication between system and user in terms of shape and motion instead of forces is that it allows users to stay in the realm they understand.

We now go over the individual systems grouped by complexity of mechanical knowledge embodied, as shown in Table 3.1.

Table 3.1: A rough classification of the types of domain knowledge involved in designing mechanical machines.

	Stationary	In motion
Force	3. Statics	4. Dynamics
Shape	1. Shape	2. Kinematics

3.2 Shape

As discussed above, there is generally no need to support users in dealing with shape. However, for functional shapes, such as enclosures that have to fit a particular object, many design steps can be automated for convenience. *Enclosed* [Weichel et al., 2013], for example, computes enclosures for electronic prototypes (Figure 3.2): Users first build their circuitry inside an integrated software package, then specify what elements should go inside of an enclosure and which ones should be facing outwards, such as displays. At this point, the system has all the necessary information and can generate a matching enclosure automatically. *Enclosed* thereby abstracts away the skill of layouting parts in a 3D space, while making sure the result fits the enclosure and can be assembled.

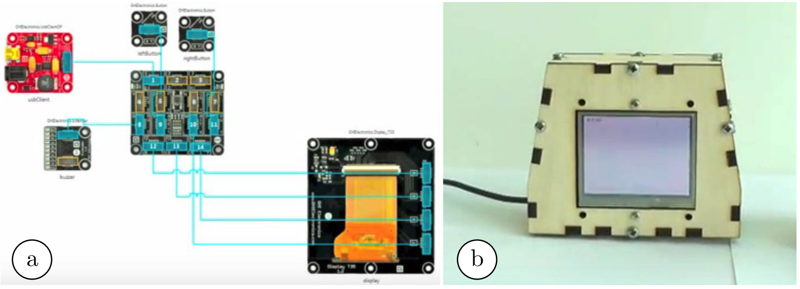

Figure 3.2: Shapes fitting with respect to each other: automatically generating enclosures for electronic prototypes [Weichel et al., 2013].

Automatically determining fit gets more complicated if the involved parts are compliant. For instance, when designing a cover to store a camera, the opening has to be large enough to remove the camera. Igarashi et al. [2009] provide a tool that automates the workflow: users simply indicate where the opening should be, the system then constructs the matching geometry.

The process of generating enclosures is further complicated if the enclosed object contains moving parts, such as the doors of a shelve (Figure 3.3). Here software needs to ensure that the parts do not collide and fulfill their functional relationships, i.e., that doors cover the shelf Koo et al. [2014].

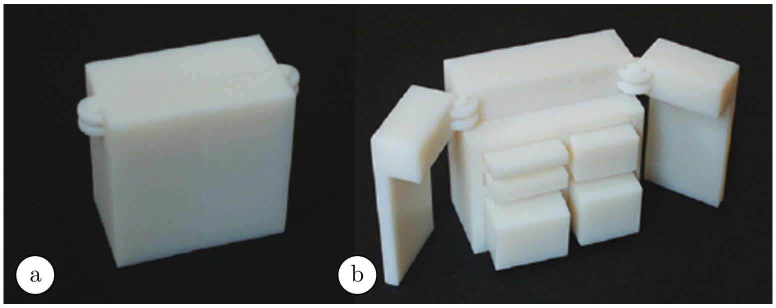

Figure 3.3: Maintaining functional relationships: The doors need to cover the shelve [Koo et al., 2014].

Similar needs to avoid collision arise when designing pose-able figurines, such as the one shown in Figure 3.4. Researchers proposed several methods for finding the optimal placement of joints given weak geometries so as to provide each joint with the space required to move and sufficient thickness to withstand the involved forces [Bächer et al., 2012].

3.3 Kinematics

Several researchers have presented systems that help users create *kinematic* systems. Kinematics describes the motion of objects without considering the forces that have caused the motion.

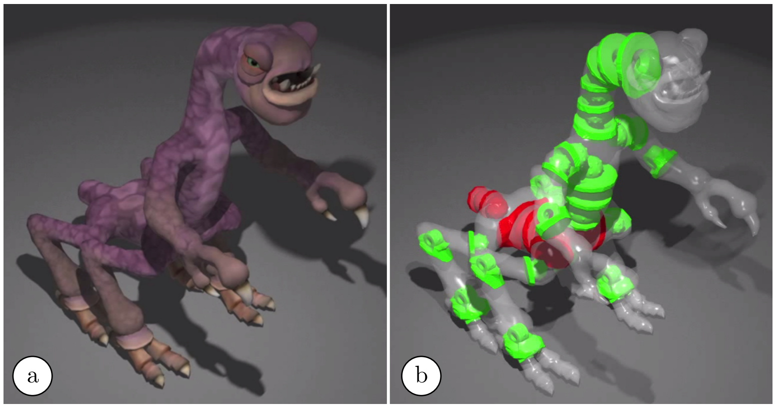

Figure 3.4: To convert this 3D model into a mechanical character whose limbs can be moved, the joints need to be placed on geometries that can withstand the forces that are applied to them [Bächer et al., 2012].

As mentioned earlier, kinematics in personal fabrication is inspired by computer animation. However, in digital animation, the mechanism that results in motion can be as simple as a set of key frames on a digital timeline. *Physical* motion, in contrast, requires a chain of volumetric interconnected mechanical elements, such as cranks, gears, and pulleys to convert an input motion (and force) into a specific output motion (and force). Such elements are often referred to as *mechanisms*.

Deciding which mechanisms to use, determining their parameters, and how they should be arranged is a non-trivial task. It becomes even more complex as additional constraints are added, such as minimizing assembly time, minimizing the space required to house the mechanisms, and minimizing the number of motors driving the mechanisms.

Kinematic design tools simplify the task: they allow users to simply specify the desired motion of their 3D model; the system then fills in the mechanism that implements this motion (Figure 3.5). To specify the motion, users can either key-frame different poses of the model [Zhu et al., 2012], sketch the motion path [Coros et al., 2013], or use data from a motion capture system [Ceylan et al., 2013]. These systems then insert various mechanisms, such as cams and followers and crank-sliders [Zhu et al., 2012], gears and linkages [Coros et al., 2013], or

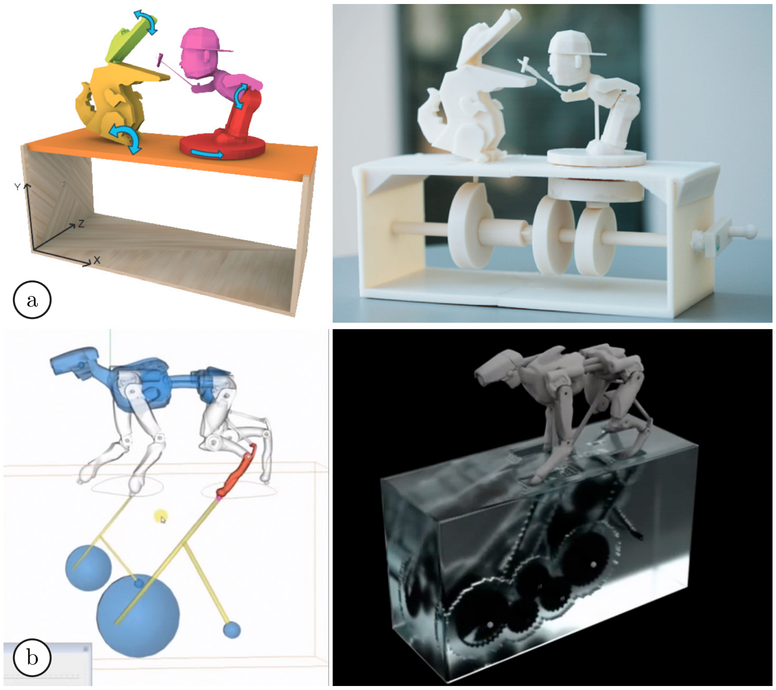

Figure 3.5: Examples of automatically generated kinematic objects: Users specify the desired output motion, and then the system creates the mechanism to produce the desired motion. (a) cams and followers and crank-sliders [Zhu et al., 2012], (b) gears and linkages [Coros et al., 2013].

specialized mechanical oscillators that are placed at every joint of the character, oscillating at different phases and frequencies [Ceylan et al., 2013].

While the systems illustrated above place the driving mechanism in a separate box below the animated character (Figure 3.5), Coros et al. [2013] point out that it is more desirable to integrate the driving mechanism with the character. For characters with a large surface area, the gears can be covered with the character's shape. However, Since gears are often too large to be included in a character's shape, Thomaszewski et al. [2014] and Megaro et al. [2014] propose using linkages instead. Linkages can be given any shape, as long as the connection points

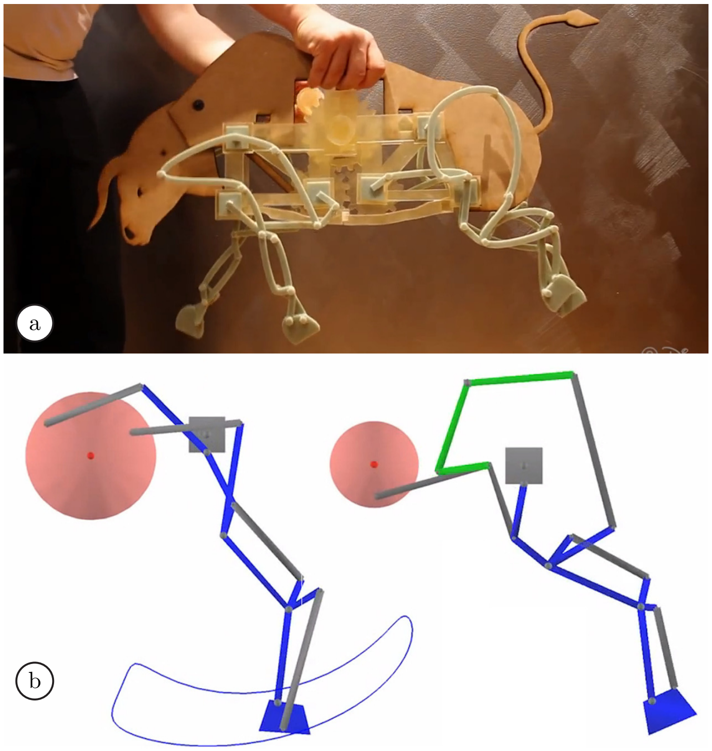

Figure 3.6: (a) Using linkages instead of gears allows integrating the functionality with the shape of the object [Thomaszewski et al., 2014]. (b) The required linkages to implement the motion are blue, the green linkages have no function but are added to better approximate the original shape.

remain the same. Thus, they can be used to define the outline of the character, such as the bull's leg in Figure 3.6. *LinkEdit* [Bächer et al., 2015] allows users to optimize such linkages.

3.4 Statics

We now go back to stationary objects — however, instead of only looking at shape, we look at the forces that act on the shape. The branch of mechanics that considers force, but in which objects are at rest and do not experience acceleration, is called *statics*.

A simple aspect of statics is balance. Unlike objects in the digital world, physical objects need to be balanced in order to not tip over. Objects are in balance whenever their center of mass is located above the convex hull of points touching the ground. Researchers proposed software that helps balance an object by changing its weight distribution. This can be achieved either through selectively hollowing the inside, using materials of different weights, or slightly changing the object's geometry (*Make it Stand* [Prévost et al., 2013]). This approach has also be extended to simple dynamic phenomena, such as "spinnable" objects that need to be balanced around a rotational axis (*Spin-It* [Bächer et al., 2014]) and floating objects that have to be balanced while swimming (*Buoyancy Optimization* [Wang and Whiting, 2016]).

SketchChair [Saul et al., 2011] considers the problem of balance for objects that are subject to additional weight (and thus additional force), here chairs subjected to a person sitting on them. The system allows users to place a virtual person onto the chair designs, which will cause poorly designed chairs to fall over (Figure 3.7).

A more elaborate aspect of statics is stiffness, i.e., if objects will break. For instance, when creating decorative hole patterns across a

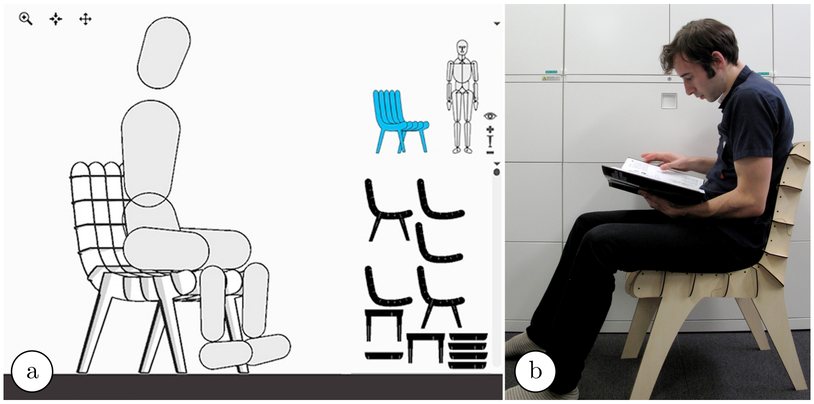

Figure 3.7: Balance: (a) *SketchChair* [Saul et al., 2011] allows quickly testing if a chair can carry the user.

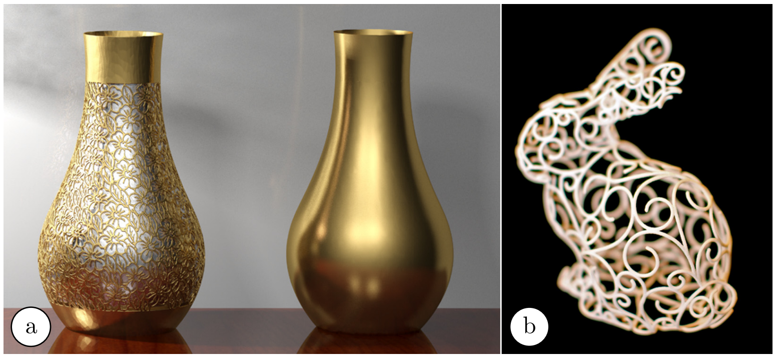

Figure 3.8: Adding decorative patterns without generating weak geometries (a) synthesis of filigrees [Chen et al., 2016] and (b) *ornamental curve networks* [Zehnder et al., 2016].

surface, such as those shown in Figure 3.8, the patterns need to be created in a way that avoids generating weak geometries (see *stenciling* [Schumacher et al., 2016], *ornamental curve networks* [Zehnder et al., 2016], and *synthesis of filigrees* [Chen et al., 2016] for examples of algorithms).

If an object already contains weak geometries, the *stress-relief* system [Stava et al., 2012] locates the weak parts in a 3D model and then either automatically thickens the respective parts or inserts stabilizing struts (Figure 3.9).

While the stress-relief system only covers gravity loads and picking objects with two fingers, Zhou et al. propose a more general algorithm [2013]. Recently, Langlois et al. [2016] have developed a more elaborate approach that calculates failure probabilities based on stochastic structural analysis.

Approaches that optimize infill to minimize weight while avoiding weak geometries include the use of wireframe structures [Wang et al., 2013] and optimized interior tessellations [Lu et al., 2014].

3.5 Dynamics

We now move on from considering forces for stationary objects to considering forces for moving objects. While the previously shown

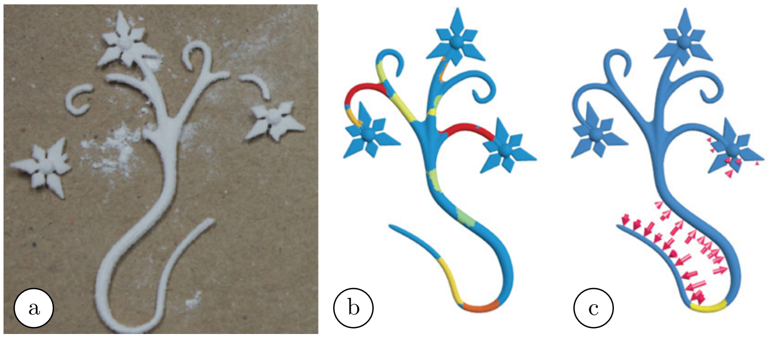

Figure 3.9: [Zhou et al., 2013] find weak parts in the 3D model and strengthen the parts accordingly.

kinematic systems consider motion, they ignore forces, i.e., they provide their models with essentially unlimited power: The model of the bull shown in Figure 3.6, for example, is actuated using a 300 W power drill. These designs also assume that the resulting heat can simply dissipate into the environment.

The systems shown in this section, in contrast, take a different approach: they consider forces and their effect on the motion, resulting in more efficient and elegant designs. In mechanical engineering, this field has been referred to as *dynamics*.

Several dynamic systems tackle the challenge of *walking*. Bharaj et al. [2015a] help users design 3D printed characters that actually walk. Building on a library of template mechanisms, their algorithm optimizes the assembly for walkability by simulating how they will perform in the physical world (Figure 3.10). Megaro et al. [2015] allow simulating characters that do not only walk straight but also take turns. They accomplish this by simplifying computational expensive dynamic walking algorithms to an approach that can be executed in real-time.

Note how these systems are different from the kinematic systems presented earlier. While the kinematic systems conveyed the *impression* of walking, they had to be held up during operation — if they were set down on the ground while performing their movement sequences, they would most likely fall over. The reason is that any movement of

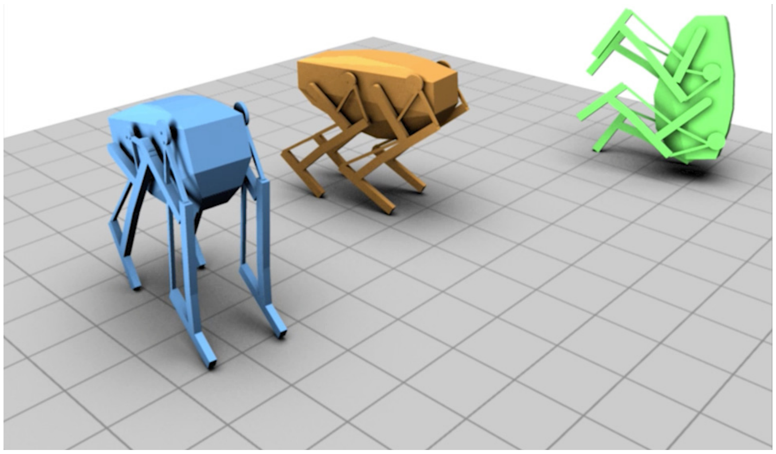

Figure 3.10: Making characters walk requires optimizing dynamic forces to prevent them from falling over [Bharaj et al., 2015a].

a limb not only shifts the character's center of mass, but also subjects the character to counter forces resulting from the acceleration of the limb. Dynamic systems, in contrast, consider these forces.

Similar challenges exist when designing flying objects (aero *dynamic* systems). *Pteromys* [Umetani et al., 2014], for instance, helps users design custom gliders: When flying, gliders are subject to drag (forces that make the glider resist the airflow) and lift (forces that move the glider upwards). Since these forces depend on the shape, velocity, and orientation of the glider and since they change constantly as the glider moves through the air, the resulting parameter space is too complex for users to tackle manually. *Pteromys* abstracts away this complex domain knowledge by automatically tweaking the user's design for optimal flight performance between editing steps (Figure 3.11).

While *Pteromys* is limited to flat symmetrical two-dimensional gliders, *OmniAD* [Martin et al., 2015] allows creating freeform 3D kites. Their user interface visualizes stabilizing forces as blue arrows and destabilizing forces as red arrows, helping users to build up an understanding of how forces interact with their design. Recently, Du et al. [2016] extended the concept to multi-copters.

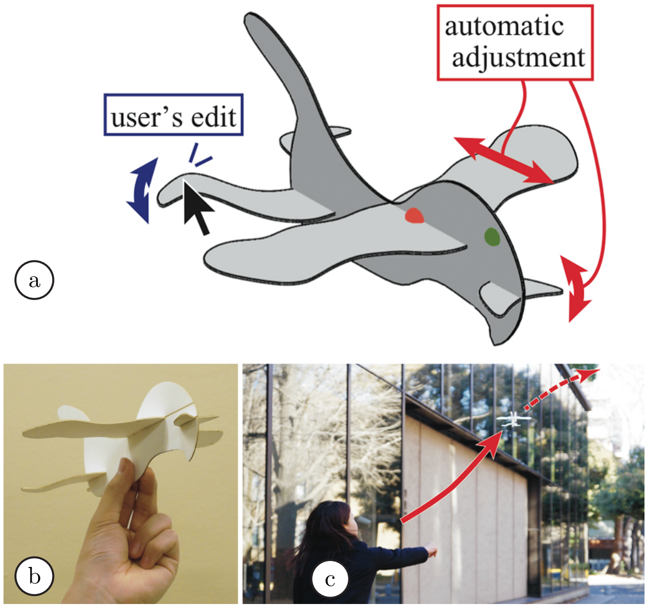

Figure 3.11: Aerodynamics: (a) In *Pteromys* [Umetani et al., 2014] users define the shape of the paper airplane, the system then optimizes for lift and drag forces to (b, c) create the best flight path.

Similar concepts can be applied to the dynamic phenomenon *oscillation*. Umetani et al. [2010] provide a tool that helps users design custom metallophones. While users design the shape of a plate, the tool simulates the acoustic frequency spectrum that would be created if the part was struck (Figure 3.12a). Bharaj et al. [2015b] take the inverse approach: they let users specify a desired input shape and sound; their system then tweaks the shape of the metal parts so as to produce the desired sound (Figure 3.12b).

Recently, several projects have extended the concept to 3D. [Li et al., 2016] present a system for designing acoustic filters that can not only produce a desired sound pitch but can also attenuate undesired noise. Since the influence of the shape on the filtered frequency bands is non-intuitive, they provide a tool that automates the process. *Printone* [Umetani et al., 2016] is a system that allows users to create functional

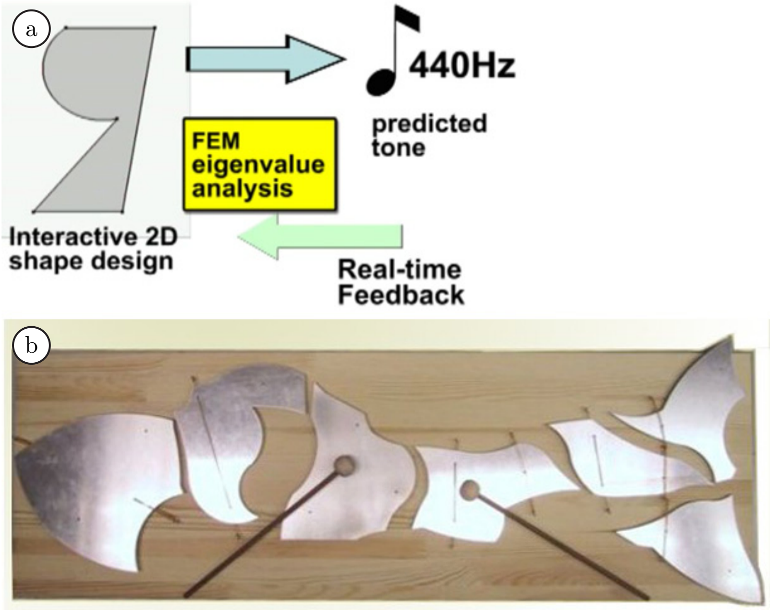

Figure 3.12: Custom instruments with desired acoustic properties: (a) the user adjusts the shape and gets feedback on how this influences the sound [Umetani et al., 2010]. (b) The inverse approach: given a shape and sound, the system slightly changes the shape to match the sound [Bharaj et al., 2015b].

three-dimensional wind instruments — users only load a 3D mesh and place the finger holes and the fipple, the system then automatically resizes the holes and creates the inner cavities.

3.6 Conclusions and open research questions

The systems mentioned in this chapter offer substantial help to users in that they allow them to create physical objects that would otherwise be very difficult to achieve. They thereby have the potential to enable inexperienced users, such as consumers, to design in a design space they would otherwise not be capable of designing in. Thus, we expect to see many more of such systems in the years to come.

However, the biggest opportunity for future research may lie in supporting users more broadly: Unlike the domain-specific and reasonably narrow solutions we see today, consumers may be able to find value in systems that help them design in a much broader design space.

We see some early systems that point into the direction of such broader "expert systems". *Design by Example* [Schulz et al., 2014], for instance, is a generic approach for automatically generating fabricatable assemblies from a database of arbitrary parts and a collection of examples that allow the system to conclude how those parts can be put together (Figure 3.13).

Similarly, *Fab Forms* [Shugrina et al., 2015] is a generic approach that allows users to browse a parametric design space (such as the heel height of a shoe) while seeing only the physically valid options (Figure 3.14).

MetaMorphe [Torres and Paulos, 2015] takes a first step at making a big design space manageable by decomposing designs along multiple dimensions: it uses *html* to define the shape of the desired object, *CSS* for its appearance, and *Javascript* for defining its function (Figure 3.15).

Achieving both, i.e., breadth and depth obviously is a tall order and it would be hard for any individual research or engineering team

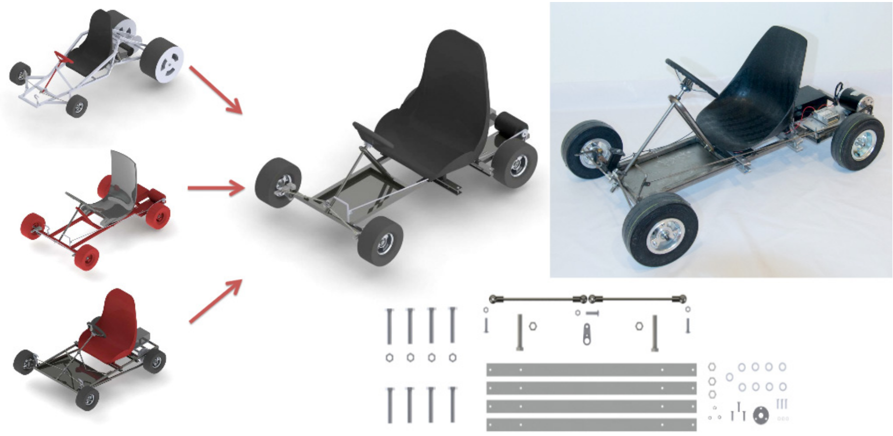

Figure 3.13: Towards generic methods for generating fabricatable parts: *Design by Example* [Schulz et al., 2014] generates designs based on a database of parts and a collection of examples using the parts.

Estimated cost: $384.41
Estimated weight: 256.27g

Figure 3.14: Parametrized *Fab Forms* [Shugrina et al., 2015] allow exploring only the physically valid options.

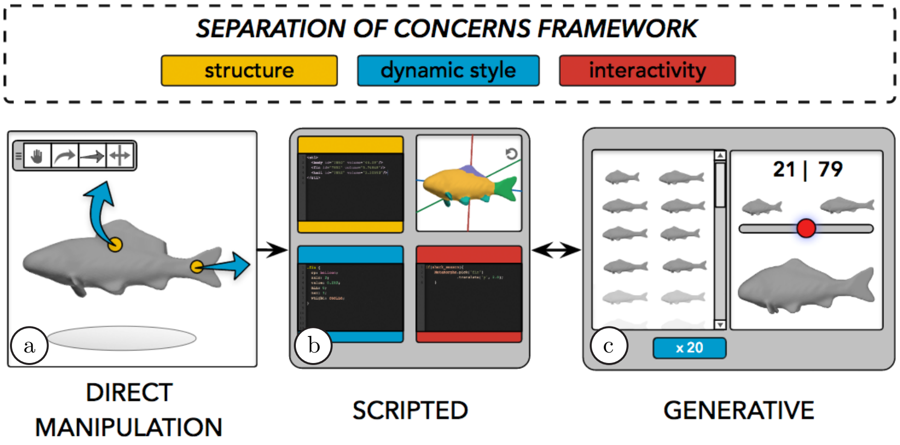

Figure 3.15: MetaMorphe [Torres and Paulos, 2015] splits object description into shape, appearance, and function.

to create such a system. However, this challenge may be solved in a distributed fashion. If we look at previous AD/DA media, the approach to tackling this appears to be plug-in architectures, such as the hundreds of plug-ins available for *Adobe Photoshop*, each of which serves a different, specialized purpose. The authors of the plug-ins benefit from

Photoshop in that it provides the overarching systems, allowing the plug-in authors to focus on their effort on their respective objective. In order to allow for such synergies in the context of personal fabrication, the next step one would expect to see is plug-in architectures as well as the standardization in general that could then allow combining systems in a plug-in architecture. The result could be a much broader notion of "assets" that can be assembled and re-configured in various ways (as demonstrated today in the context of computer game, by engines such as *Unity* [Unity Games Engine]).

4

Visual Feedback and Interactivity

Systems that embody domain knowledge can only go so far — there are always factors left that cannot be covered with an automated workflow, such as the assessment of the esthetic qualities of an object. Even with systems that embody various kinds of domain knowledge, this requires the user's judgment and often an interactive exploration of trial and error.

Consequently, AD/DA systems have always offered interactive functionality that allowed users to either preview their result (e.g., thumbnail overviews in video editing [Uchihashi et al., 1999]) or to even work directly with the preview (e.g., 'what-you-see-is-what-you-get' desktop publishing [Hiltzik, 2000]).

4.1 Editing objects with the help of visual previews

Commercially available 3D modeling systems, such as [SketchUp] and *TinkerCAD* [Thomaszewski et al., 2014], have traditionally relied on 2D input, such as mouse and keyboard, and 2D output on a computer screen. This 2D input and output, however, results in a mismatch with the objects that are being manipulated, because they are typically three-dimensional.

In order to make the input more intuitive, researchers created systems that accept user input in 3D space, e.g., allowing users to perform gestures as if they were actually manipulating an (invisible) physical workpiece. In *Spatial Modeling* [Willis et al., 2010], for example, users define the shape of a lamp by describing its surface using their hands. *Virtual Pottery* [Cho et al., 2012] allows users to shape a virtual piece of clay by moving their hands as if they were physically shaping clay (Figure 4.1).

Dress-up [Wibowo et al., 2012] enhances gestures with physical props: users sketch the shape of dresses using handheld tools on the body of a physical mannequin. Each tool has a different functionality, such as creating or removing surfaces (Figure 4.2). Similarly, *ToolDevice* [Arisandi et al., 2012] provides users with a knife prop to cut objects and a hammer to join them.

Extending this approach, researchers made output 3D as well, essentially resulting in augmented reality systems. In *Situated Modeling* [Lau et al., 2012], for instance, users wear an augmented reality headset that allows them to visually evaluate new furniture designs as they sketch them in place (Figure 4.3). *MixFab* [Weichel et al., 2014] follows the same approach, but does not require a head mounted display. Instead,

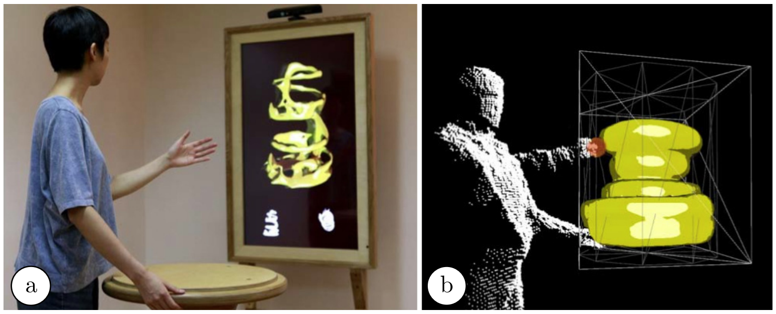

Figure 4.1: Modeling with gestures: In *Virtual Pottery* [Cho et al., 2012] users shape a virtual piece of clay with their hands.

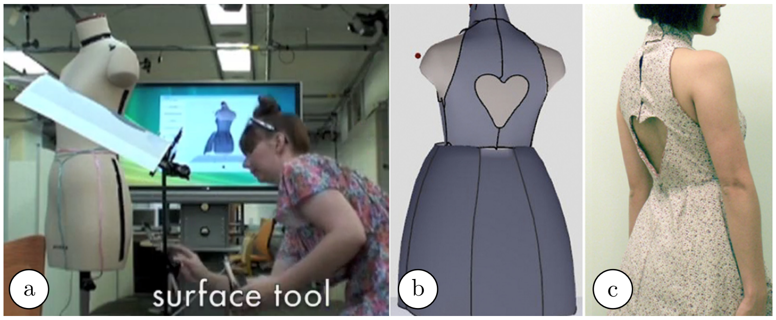

Figure 4.2: Modeling with physical props: DressUp Wibowo et al. [2012] allows users to sketch dress designs directly around a physical mannequin.

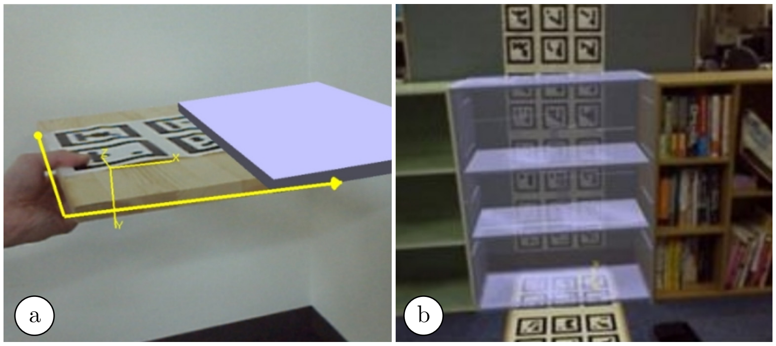

Figure 4.3: Modeling with in-context augmented reality feedback: *Situated Modeling* [Lau et al., 2012].

the *MixFab* table uses a beam splitter and a display mounted at 45° to overlay physical and virtual content.

4.2 When the object itself is required as feedback

Visual feedback is beneficial for virtually all types of physical design tasks. However, visual feedback *alone* is not always sufficient. For instance, when designing handles and tools, the exact size and ergonomics matter [Mueller et al., 2014a]. The same holds for furniture design, where users want to experience the object with their own bodies as reference [Adobe Photoshop Elements]. When designing shoes, the tactile qualities of the object may be crucial [Bickel et al., 2010] and when creating fine art sculptures users may want to know what the actual object will look and feel like [Zoran et al., 2013]. In these cases, users may prefer the actual *physical* object over a preview display, as the object itself allows them to judge any of its qualities.

To have a physical version for evaluation, users may need to fabricate at least one prototype per design iteration. This can be problematic as many of today's fabrication machines require hours to days to produce 3D objects. The duration of a design process that requires test prints thus easily gets dominated by fabrication time as every few minutes of design work may imply another overnight print, resulting in a drawn out, inefficient design process.

One way to save time is to optimize the fabrication machinery. Strategies range from optimized slicing [Wang et al., 2015] to parallel fabrication using additional print heads (e.g., Hansen et al. [2013] for FDM, *HP Jet Fusion* for inkjet printing [HP Jet Fusion]) and new printing techniques that print sequences of layers in a single continuous motion [Carbon3D, 2015].

Another approach to optimization is to render intermediate versions as simplified "previews"; only the final version is then fabricated in full detail. This multi-stage process has been referred to as *low-fidelity fabrication* [Mueller et al., 2015a] or *low-fab* for short (Figure 4.4).

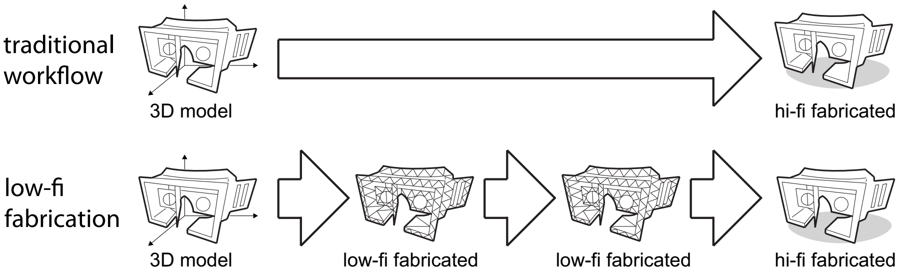

Figure 4.4: Low-fidelity fabrication prints intermediate versions as fast, low-fidelity previews [Mueller et al., 2015a].

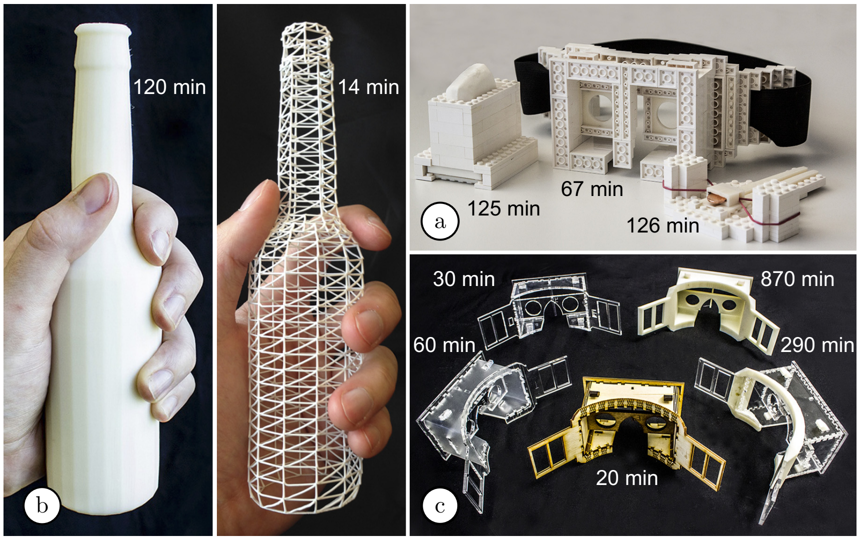

Figure 4.5: Low-fidelity fabrication techniques:: (a) *faBrickator* [Mueller et al., 2014b], (b) *WirePrint* [Mueller et al., 2014a, Wu et al., 2016, Huang et al., 2016], (c) Platener [Baechler et al., 2013].

The key idea behind low-fab is to produce previews that focus on those aspects of the object that are required by the current design iteration. *FaBrickator* [Mueller et al., 2014b], for example, is a low-fab technique that allows users to focus on selected parts of the model by replacing the rest of the geometry with LEGO bricks (Figure 4.5a). *WirePrint* [Mueller et al., 2014a, Wu et al., 2016, Huang et al., 2016], in contrast, produces objects as wireframe previews, which preserve an object's shape and thus allow users to validate the object's ergonomic qualities (Figure 4.5b). *Platener* [Baechler et al., 2013], in contrast, allows fabricating models designed for 3D printing using fast laser cutters; by preserving rectilinear elements it generally performs well on objects designed to perform a mechanical function (Figure 4.5c). All three techniques save about 90% of the fabrication time, depending on object type.

If we consider low-fab in a wider sense, any fabrication device that produces a fast, yet coarse rendition of an object can be part of a low-fab process, such as systems that fabricate using a coarse glue gun

[Lee, 2011], coarse clay extrusion [Peng et al., 2015b], or foam [Willis et al., 2011]. One would then complement these systems with a separate high-quality fabrication machine in order to produce the final version.

4.3 Interactive fabrication

While low-fidelity fabrication allows users to create, inspect, and redo objects quickly, redoing an object in its entirety may not necessarily be the most effective approach in the first place. Arguably, feedback is most beneficial when making key design decisions *along the way.* This concept of letting users experience the workpiece along the way has been referred to as *interactive fabrication* [Willis et al., 2011]. Figure 4.6 illustrates the main concept, which was inspired by direct manipulation [Shneiderman, 1983] principles.

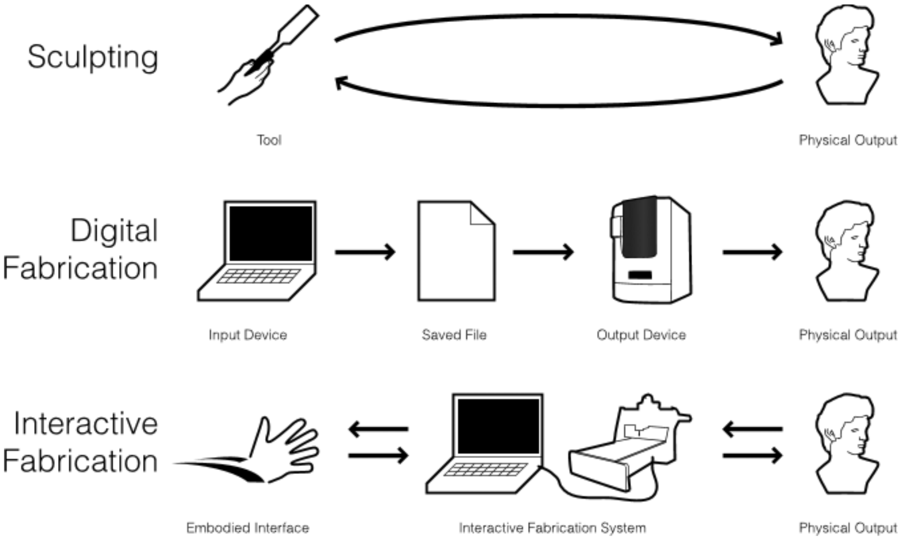

Figure 4.6: How interactive fabrication differs from crafting and the traditional 3D editing workflow according to Willis et al. [2011].

Interactive fabrication can be summarized as follows:

(1) *Continuous representation of the object of interest*: "avoid a 'representation' of the object of interest, but instead allow the user

to look directly at the fabricated form". This also means that the system responds with an updated version of the physical workpiece, which distinguishes interactive fabrication from systems that annotate the workpiece, rather than actually manipulate it.

(2) *Physical actions:* "the user interacts through an embodied interface where their physical actions are sensed and interpreted in real-time. Physical action again determines the embodied output in the form of digital fabrication.". More recent systems extend this notion with interactions that take place on the physical workpiece, for instance, by using direct touch interactions.

(3) *Rapid, incremental, reversible operations:* "Optimizing the speed and response time [. . .] The ability to reverse a physical action."

(4) The diagram further clarifies the distinction from traditional wood working and crafting tools, and their digital fabrication equivalents. The key here is the *computer system* that helps users achieve their design objective. Fabrication tools for freehand sketching, such as the 3Doodler and *Protopiper* [Agrawal et al., 2015] shown in Figure 4.7, therefore do *not* classify as interactive fabrication.

Precursors to interactive fabrication go back as far as to 2009. Figure 4.8 shows *ModelCraft* [Song et al., 2009]. This system allows users to manipulate a 3D model folded from paper by drawing change requests directly onto the paper model using an Anoto Pen. The pen allows the computing system in *ModelCraft* to track the change requests, update the 3D model, and 2D print a new paper model with folding instructions, which users assemble by hand. While *ModelCraft* does not (1) provide a continuous representation of the object but instead refabricates each new version from scratch, it was one of the first systems that allowed users to (2) work directly on the physical workpiece.

Another early prototype, *Shaper* [Willis et al., 2011] explores how to provide a (1) continuous representation through fast physical feedback, albeit at the expense of not offering (2) on-object input. By touching a touchscreen, users instruct the system where to extrude

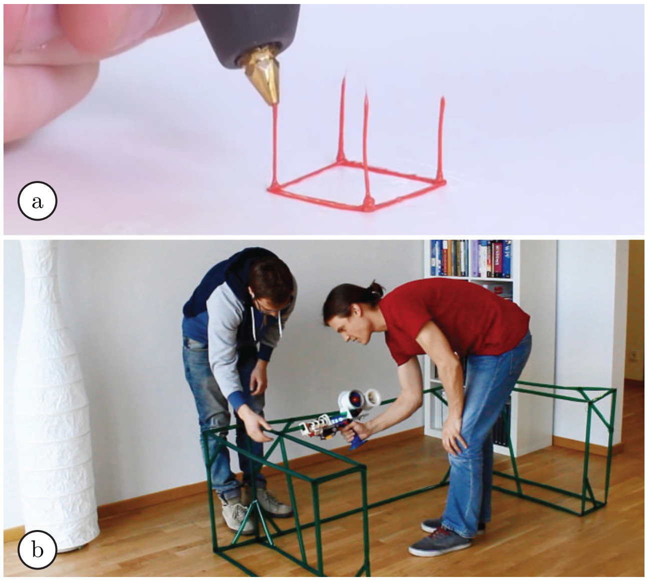

Figure 4.7: Handheld extruders for freeform sketching are *not* interactive fabrication because they lack built-in computer support: (a) the 3Doodler pen allows users to sketch in space, (b) *Protopiper* [Agrawal et al., 2015] scales this up by using light-weight tape as material.

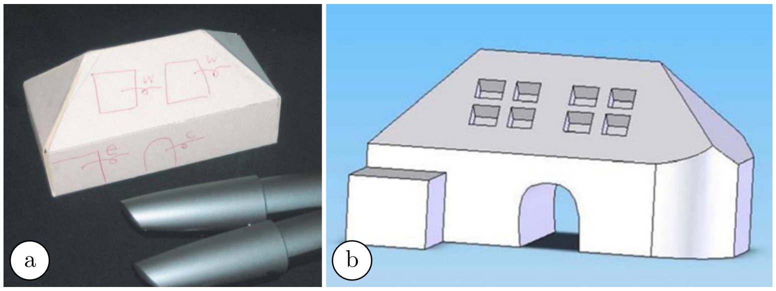

Figure 4.8: A precursor to interactive fabrication: In *ModelCraft* [Song et al., 2009] users annotate directly on a paper model, the changes are tracked, and a new paper sheet printed, which users manually fold into the changed 3D shape.

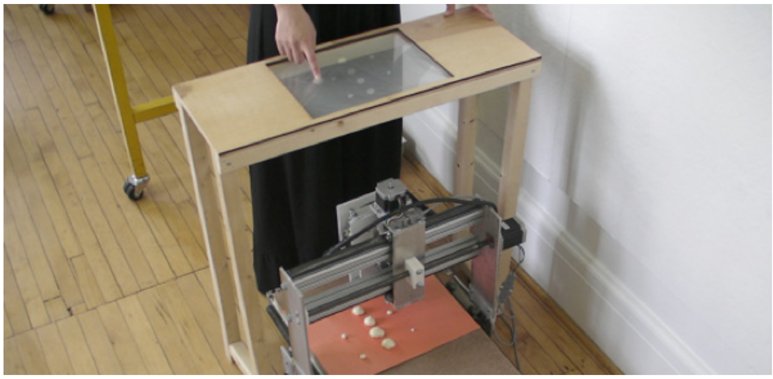

Figure 4.9: A touch on *Shaper's* [Willis et al., 2011] touchscreen causes the extruder to create a foam drop at the corresponding location.

a drop of foam onto the 2 1/2 D foam model three feet below the screen (Figure 4.9).

CopyCAD [Follmer et al., 2010] is one of the first systems that combines (1) a continuous representation through automated fabrication (here a computer-controlled milling machine) and (2) on-object interaction. As illustrated by Figure 4.10, *CopyCAD* users sketch onto the physical workpiece with a pen, a camera captures these annotations and then operates the mill accordingly. The system also allows capturing the shape of other objects using a camera, which helps users remix designs.

Constructable [Mueller et al., 2012] is the first system to introduce precision into interactive fabrication (Figure 4.11). *Constructable* allows users to point at the work piece through the safety glass enclosure using laser pointers, which enables users to work directly on the workpiece despite the enclosure. *Constructable* captures the bright dot of the laser pointer using an overhead camera, beautifies the observed path according to pre-programmed constraints, and then immediately cuts using the laser cutter. *Constructable* achieves precise interaction by means of

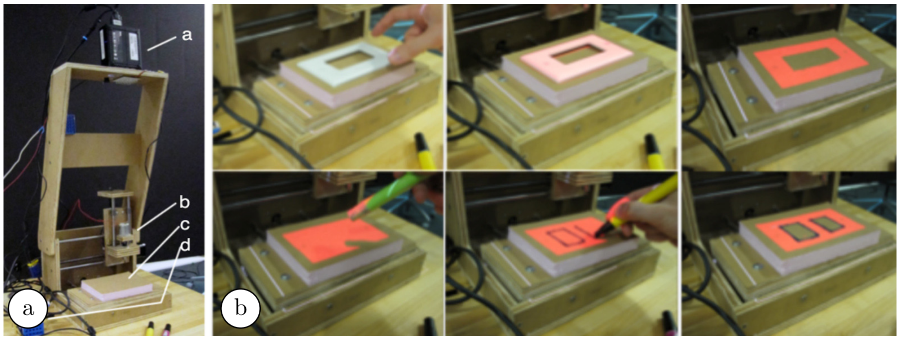

Figure 4.10: In *CopyCAD* [Follmer et al., 2010], users draw with a pen on the workpiece, the milling machine then mills the path accordingly.

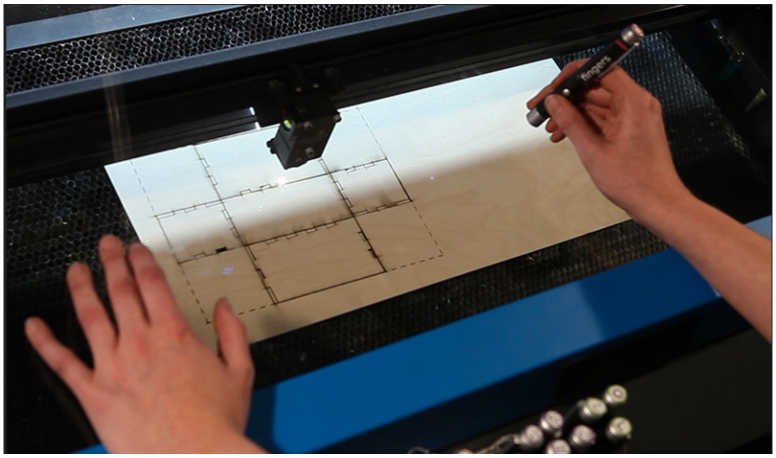

Figure 4.11: In *constructable* [Mueller et al., 2012], users draw with a laser pointer onto the workpiece, the laser then cuts the path. *Constructable* offers functionality to achieve precise mechanical constructions, such as the gearbox being constructed in this image.

tool-specific constraints; the polyline tool, for example, always produces straight lines.

Constructable is limited in that it can only produce 2D parts — the 3D model does not materialize until assembled *after* the design process is complete. To allow users to create 3D models *during* the interactive fabrication process, *LaserOrigami* [Mueller et al., 2013] extends the *constructable* system by bending 2D parts into 3D shapes while still inside the lasercutter. *LaserOrigami* achieves this by defocusing the laser to the point where it does not cut anymore, but instead heats up the workpiece so as to make it compliant and bend under gravity. Similarly, *LaserStacker* [Umapathi et al., 2015], allows users to create 3D objects in the laser cutter; however, it achieves this by welding stacks of acrylic sheets.

The concept of interactive fabrication has also been extended to on-body design. *Tactum* [Gannon et al., 2015], for instance, uses a projector to visualize, e.g., a bracelet the user is currently designing directly on the user's arm. Users can modify the design by manipulating the projection using their fingers. *ExoSkin* [Gannon et al., 2016] extends this concept by adding fabricated output: using a skin-safe cold-extrusion polymer clay, users can extrude a physical design directly on and around a body part following the displayed projection.

By combining additive and subtractive fabrication, several research prototypes introduced the notion of reversible operation into interactive fabrication. *DCoil* [Peng et al., 2015b], for instance, is a hand-held wax extruder that allows users to coil up 3D models while specifying constraints in the coiling process. It also supports a cutting knife to remove geometry from the prototype (Figure 4.12).

Similarly, *Reform* [Weichel et al., 2015] consists of a clay extruder that can add geometry and a milling machine that can remove geometry, allowing the system to reverse the last step. Finally, *On-the-Fly Print* [Peng et al., 2016b] consists of an extruder that prints *WirePrint low-fab prototypes* [Mueller et al., 2014a] and a cutting knife, which can remove outdated geometries.

4.4 Continuous interactive fabrication

The systems described above all implement an interaction style that can be described as *turn-taking:* the user produces some input, *then* the system updates the physical object, *then* the user provides another round of input, and so on. The next logical step to interactive fabrication is

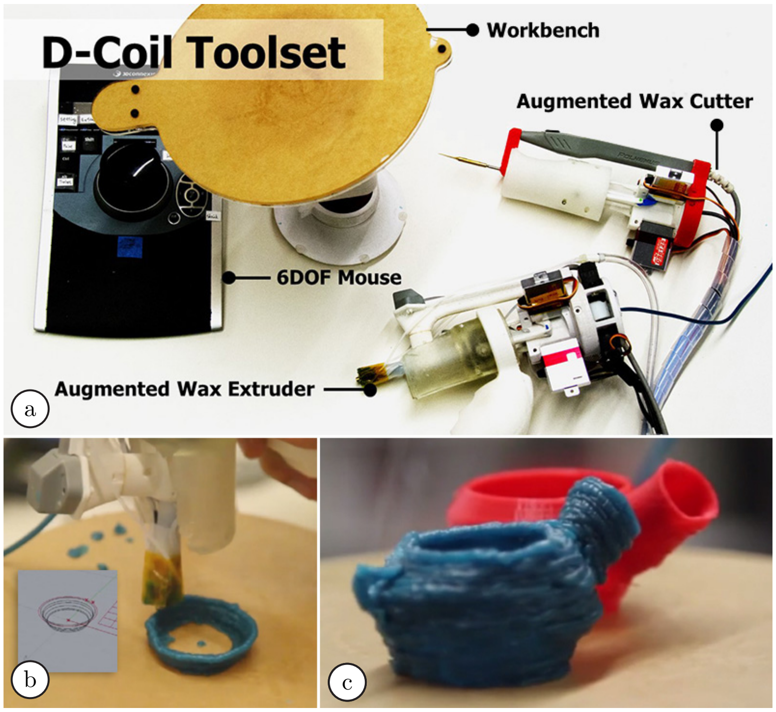

Figure 4.12: Towards forgiveness: (a) DCoil [Peng et al., 2015b] consists of a wax extruder and wax cutter to add and remove geometry. (b) The coiling process. (c) Since every interaction is tracked, the resulting model can be beautified and 3D printed at the end of the process.

to create systems that fabricate continuously *while* the user interacts with the workpiece (as described in (1) *continuous representation* from Willis et al.'s [2011] interactive fabrication definition). We would like to call this *continuous* interactive fabrication.

One group of precursors to continuous interactive fabrication are systems that allow for such continuous interaction albeit at the expense of offering only a rudimentary form of physical user action. The first of these systems is *Haptic Intelligentsia* [Lee, 2011], which is a force-feedback device with an attached hot glue gun that only extrudes material when the user is following a predetermined path. Since the path is

Figure 4.13: Continuous interactive fabrication by replicating an existing model: (a) *Haptic Intelligentsia* [Lee, 2011].

determined in a separate 3D editor prior to interacting, *Haptic Intelligentsia* is limited to users determining where to start along the path, in which direction to traverse it, and how to perform small deviations that result in texture (Figure 4.13).

Similarly, *Position-Correcting Router* [Rivers et al., 2012] allows for continuous interaction, albeit again users are limited to following a path. Similarly, *Enchanted Scissors* [Yamashita et al., 2013] only cut when the user follows a predefined cutting path; and *Augmented Airbrush* [Shilkrot et al., 2015] only sprays when the user holds it into the correct location. See *The Wise Chisel* [Zoran et al., 2014a] for a more comprehensive overview of tools of this type.

More recent projects in this line of work still use the concept of a predetermined 3D model; however, they allow users to modify this model during interaction (Figure 4.14): *FreeD* [Zoran and Paradiso, 2013], for instance, is a hand-held milling tool that stops carving when the user is close to hurting the predefined model. However, users can push a button to override the constraint and thereby keep modifying the physical object — the underlying digital model however is not adjusted.

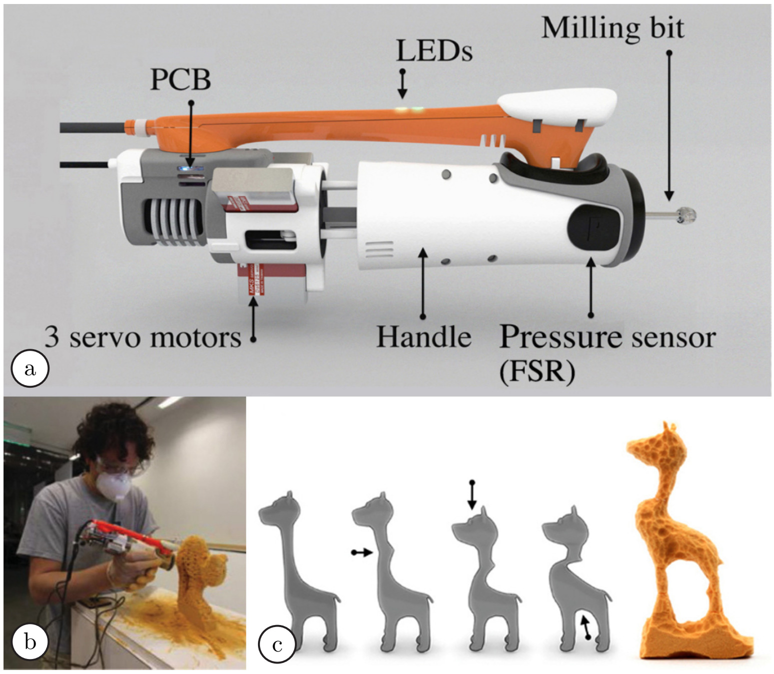

Figure 4.14: Continuous fabrication with variation to the digital model. As the user carves, the model adjusts automatically to the remaining available material (*Hybrid Carving* [Zoran et al., 2013]).

Hybrid Carving [Zoran et al., 2013] is an extension of this work that adjusts the underlying digital model on the fly: As shown in Figure 4.14c, when the user removes too much material from the workpiece, the system responds by adjusting the shape of the neck of the giraffe. This allows the system to give the user guidance for all subsequent steps based on the adjusted model. While users are guided in the process, there is still some artistic freedom, as a recent case study shows (*Hybrid Artisans* [Zoran et al., 2014b]).

FormFab [Mueller et al., 2017] is the first continuous interactive fabrication system that offers reversible operations. As shown in Figure 4.15, *FormFab* allows users to sculpt by interactively pulling bubbles out of a thermoplastic sheet.

Figure 4.15: Continuous interactive fabrication: *FormFab* [Mueller et al., 2017] allows users to push and pull the material while receiving continuous feedback about the current state.

FormFab works by first heating up the area to deform. Once the material is compliant, users interactively control a pneumatic system, thereby regulating the air pressure below the workpiece so as to make its deformation follow the hand. Such *formative* fabrication not only allows for interactive rates during deformation, but also allows users to undo their changes by pushing back in what was pulled out and by pulling back out what was pushed in, thereby overcoming a key limitation of subtractive techniques.

4.5 Conclusions and open research questions

Visual feedback and the overall concept of direct manipulation are of paramount importance when editing 3D objects for fabrication. The promise of interactive fabrication is to one step further, i.e., to deliver feedback that is more encompassing than the merely visual feedback provided by screens, thereby allowing users to judge the whole gamut of the object's qualities, including tactile qualities, ergonomic fit, and esthetics.

The design and engineering of interactive fabrication systems remains a research challenge. The main hurdle as of today still revolves around speeding up fabrication machines to interactive rates. Thus, interactive fabrication will likely trail the evolution of fast fabrication. But the vision is already here and in the meantime, placeholder technologies, such as today's fabrication devices, will allow researchers to explore what these future systems will be like.

So where will interactive fabrication go ultimately? Once more we can use personal computing as an analogy, as the evolution of interactive fabrication systems so far appears to closely mirror the history of interactive computing [Cardinal, 2011]. Up to the 1950s, executing programs in one go was the dominant way of interacting with computer systems (Figure 4.16a). In the early 1960s, command line/teletype systems allowed for turn-taking request and response interaction, leading to interactive computing (Figure 4.16b). Still in the 1960s, visionaries such as Sutherland [1963] and Engelbart [1962] introduced the concept

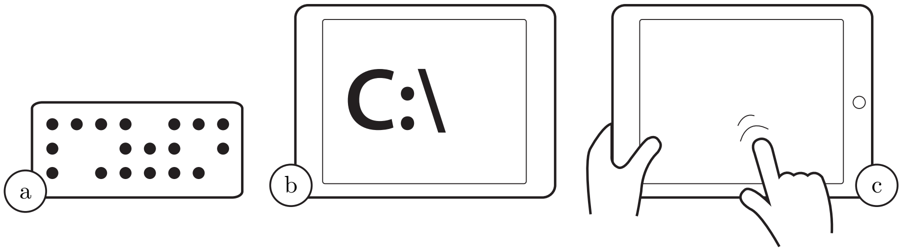

Figure 4.16: Evolution of interactive computing: (a) executing programs in one go using punch cards, (b) command lines allowed for turn-taking, and (c) direct manipulation allowed for continuous feedback. The evolution of fabrication systems follows the same pattern and ultimately arrives at continuous interactive fabrication.

of direct manipulation, which is the dominant interaction concept seen today (Figure 4.16c).

When we map this analogy back to personal fabrication, one might argue that fabrication lives roughly in the late 1960s. The interaction with 3D printers still resembles the punch card era, as 3D printers are still primarily operated by sending the object in one go and take hours to fabricate with no opportunity to modify the process along the way. At the same time, we already see a good number of research prototypes based on turn-taking interaction and we start to see systems based on direct manipulation, suggesting that the future of modeling systems for fabrication could go this way.

Open challenges in interactive fabrication revolve around the four main qualities of interactive fabrication mentioned earlier: (1) continuous representation through updating the workpiece, (2) physical actions and hands-on editing on the workpiece, (3) rapid, incremental, and reversible operations, and (4) the computing system that supports the user. As of today, various subsets of these qualities have been demonstrated, however, the ultimate goal is to achieve all of them in one single system. Such a system with immediate physical response and also capable of reversing actions will allow user and system to morph the workpiece into any shape — instantaneously and continuously, thereby providing user and system with full control over physical matter.

At this point, the technology behind personal fabrication will fuse with several related fields, all of which aspire to or build on the same vision of gaining control over physical matter (Figure 4.17), including modular robotics [Fukuda and Kawauchi, 1990], smart matter [Goldstein and Mowry, 2004], shape-changing interfaces [Coelho and Zigelbaum, 2011], tangible computing [Ishii and Ullmer, 1997], and shape displays [Follmer et al., 2013] — an exciting perspective.

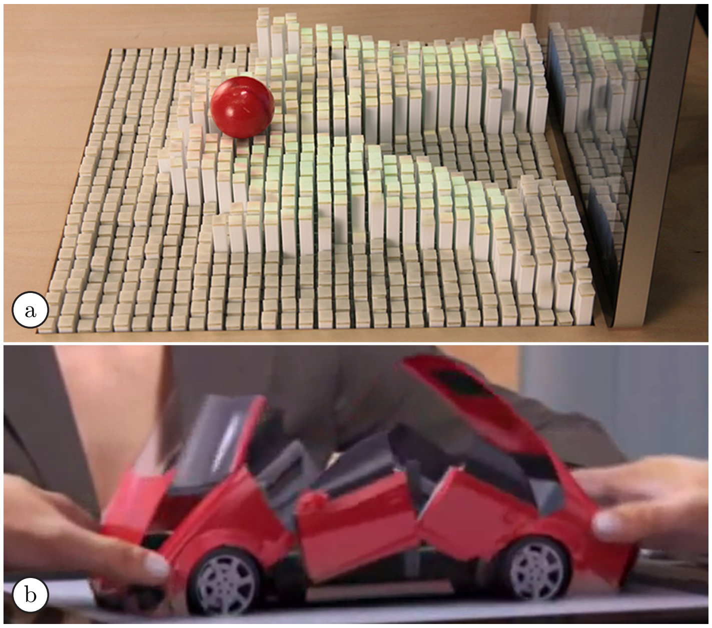

Figure 4.17: Two examples of other fields attempting real-time physical change: (a) shape displays (e.g., *InFORM* [Follmer et al., 2013]), (b) smart matter (e.g., *Claytronics* envisionment [Goldstein and Mowry, 2004]).

5

Machine-Specific Knowledge

Once a physical object has been designed, be it with the help of a system embodying *domain knowledge* or *feedback and interactivity* or both, users will want to fabricate it.

As with the DA component in any of the previous AD/DA technologies, personal fabrication machines make users' lives easier as they eliminate the need for users to operate mechanical tools, such as saws and wood chisels, and thus eliminate the need for mechanical skill.

In exchange, however, the new tools introduce their own challenges. In particular, these machines require objects to be converted to a representation specific to the machine and often also to the used material. Typically, there are multiple ways for how a given model can be converted and a naïve approach often leads to sub-par results, such as higher material consumption, higher production of scrap and support material, the need for a larger build volume, or the resulting objects may break more easily as 3D printed layers tend to delaminate.

Consequently, researchers have developed tools that help users overcome these hurdles. Unlike the techniques we discussed in the *Domain Knowledge* chapter, i.e., techniques that embody knowledge about mechanical and structural engineering, the techniques discussed in this

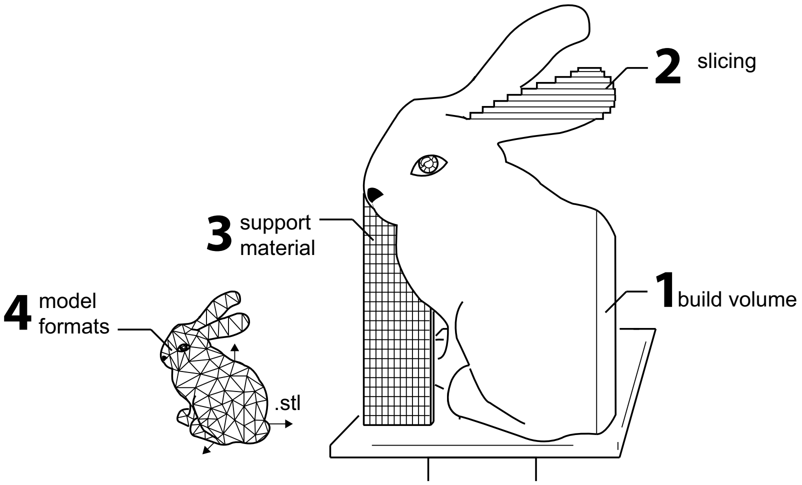

Figure 5.1: Some examples of machine-specific knowledge: machine specific build volume, slicing heights, support material generation, and model formats.

chapter embody knowledge about and are specific to the fabrication machine at hand (Figure 5.1).

5.1 Fitting a 3D model into a machine-specific build volume

Chopper [Luo et al., 2012] splits up models in a way that makes them fit into the limited-size build chamber of their 3D printer while ensuring the resulting parts are easy to assemble and the seams are unobtrusive (Figure 5.2a). *Dapper* [Chen et al., 2015] also splits models into several parts, but tries to minimize the number of printed layers (Figure 5.2b).

5.2 Optimizing slicing for stability

Three-dimensional printed objects are sturdier against forces in some directions than in others. The reason is that most types of 3D printers 'glue' one layer of material onto the next, so that objects are more likely to break along the interface between two layers. To minimize this issue, Umetani and Schmidt [2013] re-orient objects before 3D printing

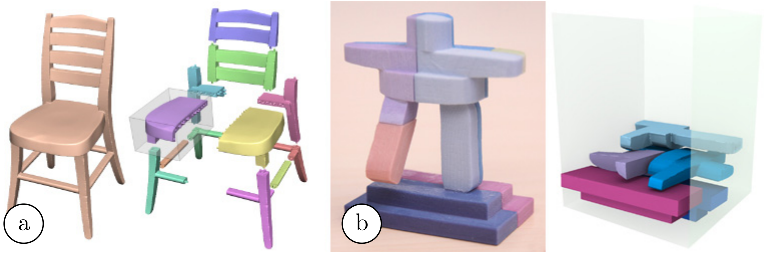

Figure 5.2: (a) *Chopper* [Luo et al., 2012] splits a model to fit the build volume of a given 3D printer. (b) *Dapper* [Chen et al., 2015] optimizes for number of printed layers.

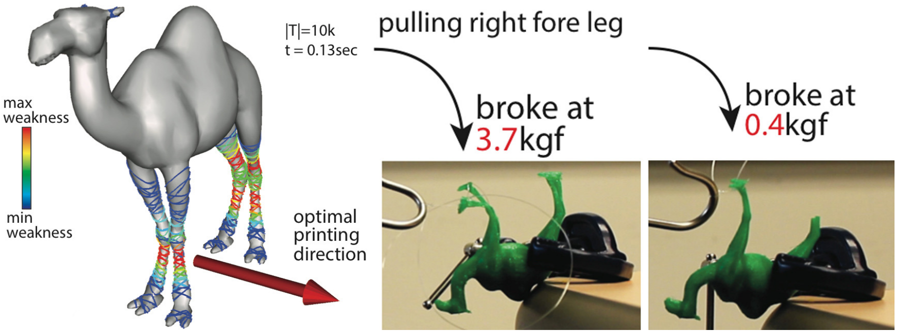

Figure 5.3: Optimizing the fabrication process by changing the printing direction to optimize for structural strength [Umetani and Schmidt, 2013].

according to the direction of the expected forces (Figure 5.3). *Orthogonal Slicing* by [Hildebrand et al., 2013] takes this a step further. Instead of slicing along one axis, they propose splitting the model into parts, which allows them to optimize the slicing direction for each part individually. The parts are then assembled to form the object.

5.3 Optimizing FDM printing speed

Wang et al. [2015] save 30–40% printing time by implementing low-detail regions as thicker layers, i.e., adjusting the slicing height accordingly, which results in less extruded material (Figure 5.4).

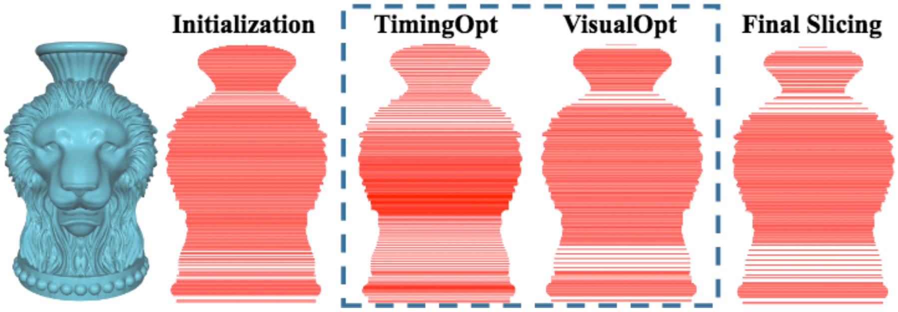

Figure 5.4: Optimizing the fabrication process to save printing time by printing with different layer heights [Wang et al., 2015].

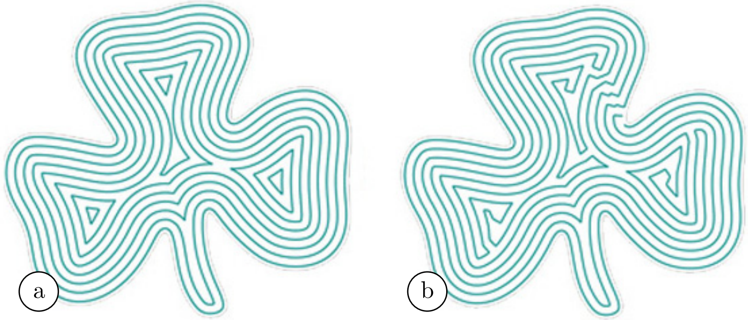

Figure 5.5: By reconsidering the print path for each layer, researchers showed how to save printing time and improve overall print quality: (a) Many closed curves. (b) One continuous curve [Zhao et al., 2016].

A different approach to speeding up fabrication time is to reconsider how each layer is produced by the 3D printer. *Connected Fermat Spirals* [Zhao et al., 2016] is a new approach to laying out the filament on each layer in FDM 3D printing: since the new path is one globally continuous curve, the layer is not only printed faster but is also of higher quality (Figure 5.5).

5.4 Calibrating joint geometries

Creating joints requires the involved parts to not only have enough space to be assembled and move, but they also have to fit to produce

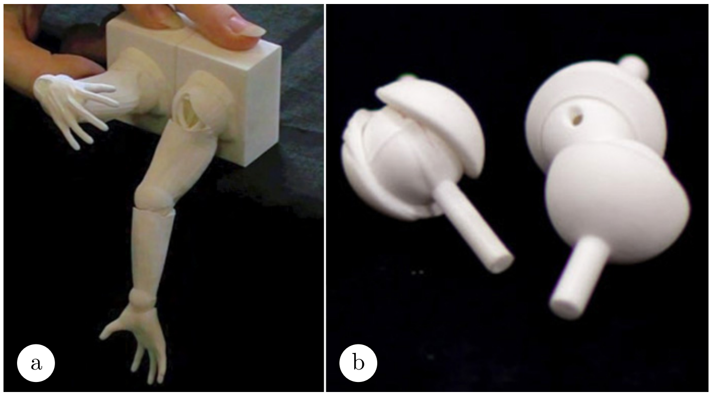

Figure 5.6: (a) Friction joints allow users to position a part in a desired pose [Calì et al., 2012]. (b) Easy to print joint-geometries, also from [Calì et al., 2012].

sufficient friction forces that hold the parts together. For instance, when assembling a box from 2D laser-cut parts that have finger joints, the joints on each part have to be carefully calibrated to each other to hold together. In laser-cutting, this requires considering the thickness of the cut line and how much it is slanted, which is determined by the focal length of the laser, the distance to the workpiece, and the thickness of the workpiece.

Similarly, when creating ball joints that should hold a part in a desired position, all parts of the joint need to be carefully calibrated to each other to hold the part in place (see the holding vs. the hanging arm in Figure 5.6a). Calì et al. [2012] describe different joint geometries that are particularly easy to calibrate and suitable for 3D printing as they require only minimal support (Figure 5.6b).

5.5 Minimizing support material

FDM and stereolithography 3D printing require support material structures to allow them to print overhangs. These are 3D printed along with the actual object geometry and generally removed after printing completes. Support material can account for a substantial portion of printing time and also requires time to be removed (Figure 5.7a right side).

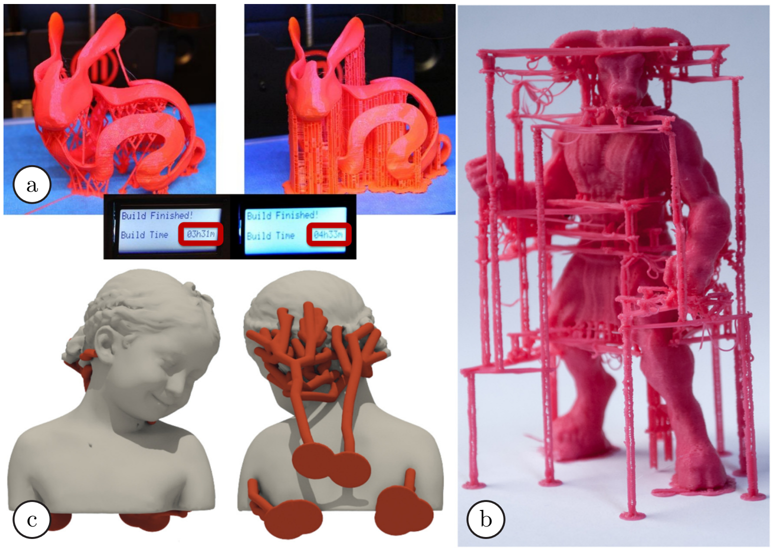

Figure 5.7: Saving support material: (a) right: traditional support, left: branching support [Schmidt and Umetani, 2014], (b) bridging support [Dumas et al., 2014], (c) optimizing for visual saliency, i.e., fewer artifacts after support removal [Zhang et al., 2015].

In order to reduce the need for support material, researchers show how *branching support structures* [Schmidt and Umetani, 2014] require only a fraction of the material since they connect to the surface of the object at just a few points (Figure 5.7a). Vanek et al. [2014] extend the concept of branching support by reorienting the object to minimize the overall area that requires support (support structures are only required for steep overhangs, typically larger than 45°).

Dumas et al. propose creating support structures by *bridging* [2014], i.e., printing horizontally by pulling the print head sideways (Figure 5.7b). They argue that this printing process is more reliable than branching as reconnecting to thin branching structures on the next layer tends to be error-prone.

Since removing support leaves visual artifacts, Zhang et al. [2015] show how to place support material in a way that least interferes with the dominant visual features on the model (Figure 5.7c).

Another approach to avoiding support material is to reorient the object so as to convert any overhang that requires support into one that can be printed without. The *RevoMaker* [Gao et al., 2015] solves this challenge using a rotating build platform (Figure 5.8). The build platform consists of a cube that during the printing process becomes the interior of the object.

CofiFab [Song et al., 2016] avoids support by splitting the model in several parts that have no overhangs. After printing, the user assembles the parts by attaching them to a laser cut infill also generated by the system (Figure 5.9).

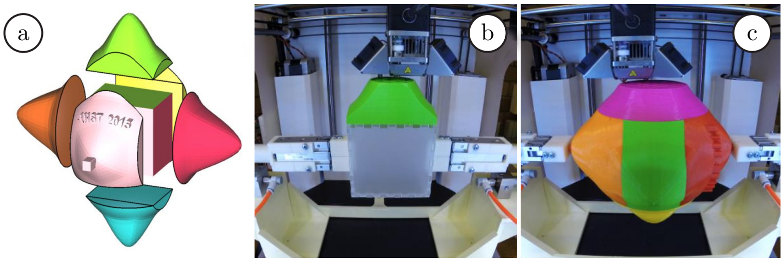

Figure 5.8: Avoiding support material: the *RevoMaker* [Gao et al., 2015] uses a rotating build platform.

Figure 5.9: Avoiding support material: *CofiFab* [Song et al., 2016] splits the model into parts that do not require support, but the user has to assemble them at the end.

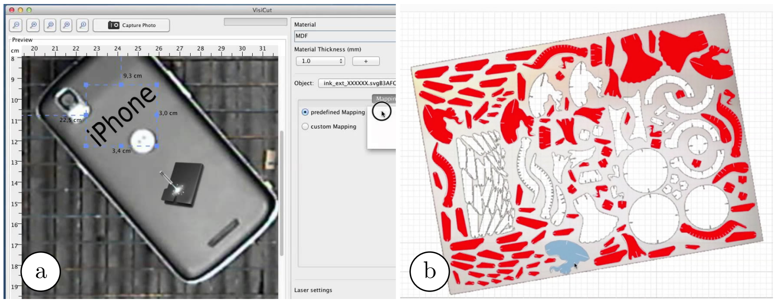

Figure 5.10: Optimizing use of material sheets using a camera to display the inside of the laser-cutter (*VisiCut* [Oster and Bohne, 2011]). (b) Using efficient interaction techniques for packing (*PacCam* [Saakes et al., 2013]).

5.6 Minimizing scrap material

Along the same lines, several researchers proposed methods for minimizing scrap material in *subtractive* manufacturing, especially laser cutting. *VisiCut* [Oster and Bohne, 2011], for instance, helps users pack their laser cut parts more tightly, thereby saving scrap. *VisiCut* takes a snapshot of the material sheet inside the laser cutter, displays it to the user, and lets the user re-arrange parts to make optimal use of the available material (Figure 5.10a). *PacCam* [Saakes et al., 2013] extends *VisiCut* [Bharaj et al., 2015a] by providing interaction techniques that make rearranging parts more efficient (Figure 5.10b).

5.7 Conversion of 3D models to 2D fabrication machines

Currently, models for digital fabrication are fairly process-specific. This makes, for instance, fabricating 3D objects on machines that are limited to producing planar parts, such as laser cutters, challenging. The traditional approach is to leave the conversion to users, i.e., users create the 2D cutting plans. This, however, requires users to imagine how to break down the 3D object they have in mind into 2D workpieces. It also makes it difficult for the user to make modifications to a model, as most modifications will affect multiple parts as well as their interplay. The same holds for all other 2D fabrication machines that follow the

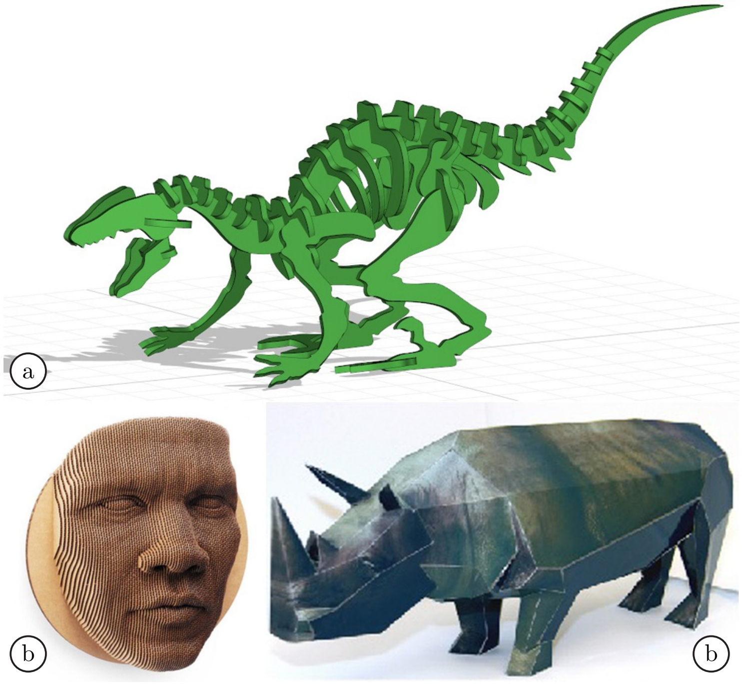

Figure 5.11: Three-dimensional modeling for laser cutting: (a) *FlatFitFab* [McCrae et al., 2014] constrains the modeling space to laser-cut-able 2D plates. Other tools convert the 3D model to (b) stacks [Authentise], or (c) foldable strips [Chen et al., 2013b].

same working principle, such as plasma cutters, water jet cutters, and 3-axis milling machines.

FlatFitFab [McCrae et al., 2014] offers a different interaction model in that it allows users to edit a laser cut model in 3D. On export, *FlatFitFab* automatically converts the 3D model to 2D parts with matching joints (Figure 5.11a). To guarantee that the conversion will work out reliably, *FlatFitFab* constrains 3D modeling to 2D plates arranged in 3D space, thus to those 3D models that have a clear laser-cut-able representation.

Other systems let users design objects in an arbitrary 3D editor, resulting in arbitrary 3D shapes. Only afterwards these systems perform a (lossy) conversion of the 3D model into a 2D representation, resulting either in intersecting planes (*Fabrication-aware Design* [Schwartzburg and Pauly, 2013]), stacks of slices (*Autodesk 123 Make* [Authentise]), foldable strips (*Multiplanar models* [Chen et al., 2013b], *Principal Strips* [Takezawa et al., 2016]), or plates (*Platener* [Baechler et al., 2013]). This line of work builds on earlier systems that produced similar conversions for (manual) crafts processes, such as paper folding (*Making PaperCraft Toys* [Mitani and Suzuki, 2004]) and sewing (*Plushie* [Mori and Igarashi, 2007]).

5.8 Machine-independent object specifications

Traditionally, models have been described as *geometry*: For 3D printing as large collections of triangles (*.stl* format), for laser cutting as lines in the 2D plane (*.svg* format). Several researchers have proposed to instead represent 3D objects by describing their *functional* properties [Vidimče et al., 2013], such as how the object responds to forces. The benefit of such a format is that the conversion to a specific fabrication machine can be done while preserving the desired functional properties.

Bickel et al. [2010], for instance, show how to replicate the deformation behavior of existing objects using the limited set of materials on a 3D printer. As an example, they use a shoe sole. In order to obtain a function-driven specification, they measure how much the material deforms in response to being pushed by a prong with a specific force. Their algorithm then automatically generates a stack of layers made from different materials that will produce the same response (Figure 5.12). In Bickel et al. [2012] show how to automatically compute the material composition for 3D printable synthetic skin that deforms in the same way as the corresponding real skin.

Conceptually similar, users of *Deformable Characters* [Skouras et al., 2013] specify the geometry of an object and a set of target poses this object is supposed to assume. The system then automatically computes the internal material composition and the location for actuation points that will allow the object to be posed as specified (Figure 5.13). Their algorithm proceeds by exploring the search space of all possible material mappings, simulates the outcome for each one, and picks the one that best reproduces the desired properties.

Figure 5.12: High-level functional specifications for materials: Deformation behavior of a shoe sole [Bickel et al., 2010].

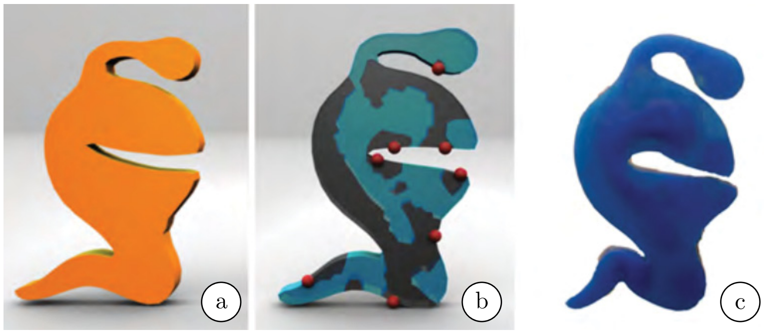

Figure 5.13: High-level functional specifications for materials: Deformation behavior to achieve a set of target poses for animation [Skouras et al., 2013].

Mesh2Fab [Yang et al., 2015] takes a different approach to achieve a desired object function: Instead of adjusting the material composition, it adjusts the object's geometry to work with the specific material by either thickening parts or correcting the contact angles between them (Figure 5.14).

Spec2Fab [Chen et al., 2013a] provides a generalized framework for a function-to-material translation pipeline, thus works for a wide variety of high-level functions, such as mechanical and optical properties (Figure 5.15).

Figure 5.14: *Mesh2Fab* [Yang et al., 2015] preserves an object's function by adapting the model geometry based on the available material.

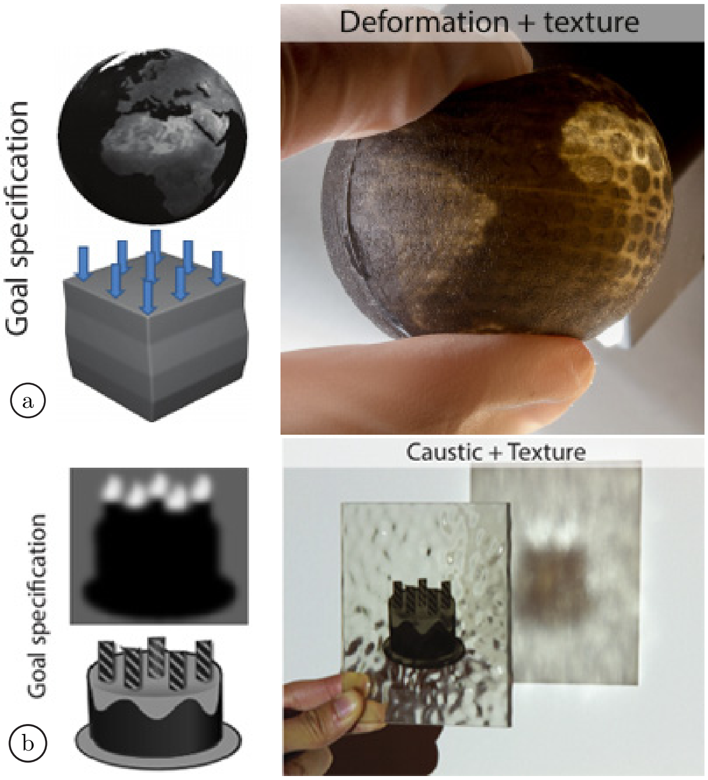

Figure 5.15: *Spec2Fab* [Chen et al., 2013a] provides a generalized optimization framework for high-level functional material assignments.

With *OpenFab* Vidimče et al. [2013], propose to integrate the printing material into shaders (so-called *fablets*) that can be assigned to different object geometries. Finally, *Foundry* [Vidimče et al., 2016] provide a user interface that allows users to define materials' properties by composing a set of operators into an operator graph.

5.9 Machine-independent specification of physical controls

The concept of machine-independent specifications goes a long way. Savage et al. apply this concept to physical controls, so that users describe what physical controls they want and where to place them; the system then generates their geometry and adds the tracking abilities required to capture their position. Users of *Lamello* [Savage et al., 2015], for instance, drag high-level functional components, such as sliders and dials from a library into their 3D model (Figure 5.16). Upon export to a 3D printer, *Lamello* enhances these components with tines that create different sounds when being struck. During use, a clip-on-microphone allows *Lamello* to sense which control is operated.

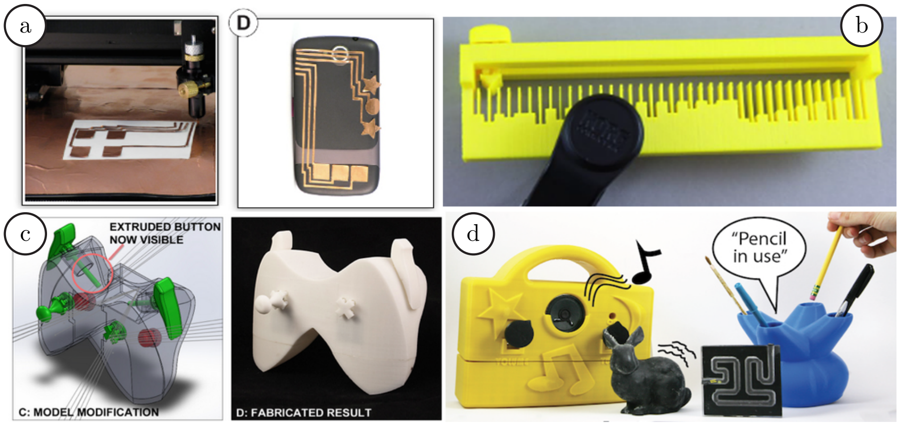

Figure 5.16: Fabrication-independent physical controls: (a) using a vinyl cutter and conductive tape (*Midas* [Savage et al., 2012]), (b) a microphone and tines (*Lamello* [Savage et al., 2015]), (c) a camera and markers (*Sauron* [Savage et al., 2013]), (d) tubes and infrared signals (A Series of Tubes [Savage et al., 2014]).

Sauron [Savage et al., 2013] approaches the same challenge with optical tracking. Users insert controls into the 3D model of, for example, a game controller, and attach a virtual camera. If the camera is not able to see all the controls, *Sauron* automatically extends the geometry of the components so as to be within the view cone. During use, the camera observes the controls to track the user interaction.

A Series of Tubes [Savage et al., 2014] extends this concept further by including IR sensing and pneumatics. The system automatically carves pipes into the object geometry to make space for the IR signal and the air stream to pass through.

Finally, *Midas* [Savage et al., 2012] generates layouts of touch sensors that are then implemented using conductive tape and a vinyl cutter. Users only specify the touch points; the system then automatically generates the routing to connect the touch points with a sensor.

5.10 Conclusion and open research questions

Machine knowledge covers many of the basic issues that need to be taken care of in order to allow consumers to use the machinery.

Personal fabrication is still at an early stage and as a consequence, we see a lot of these projects still as separate software systems. In the future, we expect anything machine-specific to be integrated in the operating systems of the fabrication machines, such as the drivers, the same way that *Postscript* interpretation today takes place inside of the printers, rather than inside the personal computer.

One of the key challenges we see for the future is to make the descriptions of physical 3D objects machine-independent. On the one hand, users might want to share models with people who own a different fabrication machine. But at least as important, users will update their own fabrication machines and the new machines will differ from the old ones. While this will have only a moderate effect on decorative objects, it can affect functional objects and mechanisms to the extent that they stop working.

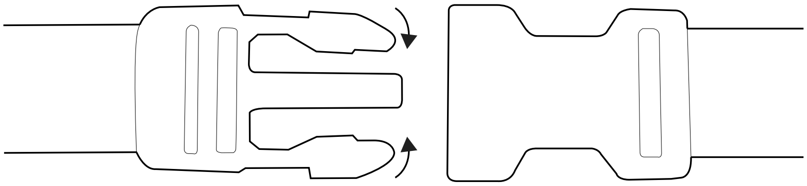

Figure 5.17: The problem with geometry-only file formats: This snap fit might not work when printed on a different machine or with a different material.

Figure 5.17 illustrates the problem at the example of a simple mechanism: a snap fit closure commonly found, for example, as part of backpacks. These devices are typically fabricated from elastic plastic materials. When printed on a different machine and in particular with a different material, the device may stop working. The two springy tips that bend inwards to snap into the cover become hard to operate when printed from a stiffer material; they may break if the material is too brittle; or it may jam if printed from a material with a tacky surface. This problem arises, because the currently used file exchange formats, such as *.stl* and *.svg* describe geometry only, which makes them material- and machine-specific.

Along the lines of machine-independent specifications, which we already discussed to some extent above, we would much prefer a format that describes an object's function rather than its shape. This would allow each machine to replace the original object with an object of a different shape but of similar function, such as two cantilever springs of specified stiffness in the case of the snap fit closure.

Once again there is plenty of historical precedent in the world of computing. Early programming took place in machine-specific machine language, but quickly transitioned to portable high-level languages, such as *C* (1972), and even to byte-code languages, such as *Java* (1991). Similarly, early descriptions of print were machine-specific, but eventually transitioned to high-level descriptions, such as *PostScript* (1976). This transition, which took place in computing from the 1970s to the 1990s, is what we now need to see in fabrication.

Currently, new description languages are emerging, such as the *.3mf* [3MF File Format, 2015] and *.amf* [AMF File Format] file formats, and their declared goal is to replace the overused low-level format *.stl*. The declared intention is to make 3D models more portable by, e.g., allowing to specify multiple model parts, their relations, and possible transformations in a single file. This is certainly a step in the right direction, but a truly function-driven specification should go even further and is thus a great research opportunity.

And for those who (rightfully) expect hand-optimized solutions to perform better, maybe we can draw inspiration from desktop publishing one more time and borrow concepts, such as "font hinting" (hand-created specification to adjust the display of a font so that it lines up with the raster of a raster display).

If we zoom out even further, we would expect to see a diversification of description languages ranging from hand-optimized high-performance low-level descriptions (that integrate domain knowledge and machine knowledge, allowing users to express domain knowledge in the context of a given machine) to less efficient, but highly portable descriptions.

6

Sustainability

As discussed in the introduction, the main premise of this monograph is to investigate the transition of personal fabrication to consumers. If the numbers should follow the lead of the previous AD/DA media, this could mean a transition from hundreds of thousands of users to hundreds of millions of users. While we share the excitement about this future evolution, we are worried about potential implications on sustainability.

Even though the previous instances of the AD/DA pattern, such as text, images, music, and video, were essentially just about data — they still resulted in substantial physical waste. For instance, desktop publishing resulted in several trillion pages being printed and disposed of every year [Ragnet, 2008].

Personal fabrication may come at an even larger environmental impact. The reason is that the traditional AD/DA media tend to consume fewer physical resources and thus create less waste as they go more and more digital — the creation of physical objects, in contrast, *inherently* consumes material and creates waste.

We thus feel this is the right time to ask whether and how personal fabrication can be made sustainable. Can we reduce the consumption of resources per fabricated object? Can we recycle? In what cases can we avoid fabricating altogether? And — is the material consumption really the main challenge with respect to sustainable fabrication?

6.1 Reducing material consumption

One approach to more sustainable fabrication is to reduce the amount of material consumed for a given 3D print.

One solution is to optimize the distribution of infill material since this leaves the appearance of the resulting objects unchanged. Any of the infill techniques from section 'infill' in the chapter *material* work, such as those that use skin-frame structures [Wang et al., 2013] or reduce the honeycomb infill pattern [Lu et al., 2014].

Material savings can also be realized by rendering surfaces as mere wireframes (*WirePrint* [Mueller et al., 2014a]). Such a reduction will typically be acceptable when fabricating prototypes for evaluating shape, such as in ergonomic testing (as part of a *low-fidelity fabrication* process [Mueller et al., 2015a]).

When designing purely functional objects, even object geometry tends to be flexible, allowing for further optimization. *SolidThinking* [SolidThinking], for example, uses this observation to reshape all structural elements of an object into trusses that require less material. Similarly, Galjaard et al. [Fukuda and Kawauchi, 1990] show how to optimize the geometry of joints so as to require minimal printing material while still being structurally sound.

Another approach to reducing the amount of material consumed is to embed reusable objects. *FaBrickator* [Mueller et al., 2014b], for instance, complements 3D prints with reusable building blocks.

6.2 Reducing support and scrap material

Personal fabrication machines typically consume not only the material that accounts for the final object, but may also process materials that end up directly as waste.

While 3D printing by means of laser sintering has the desirable property that unused material stays in powder form, allowing it to be reused right away, many other forms of 3D printing, such as FDM, require support material in order to fabricate overhanging structures, as discussed earlier. Therefore, any of the techniques from section "Minimizing support material" in the previous chapter *machine knowledge* help to improve sustainability.

Support material can cause additional environmental issues, in that its removal tends to require the use of a solvent, some of which are toxic. The recent invention of water-soluble support materials is a step in a more sustainable direction [Water soluble support]. However, the water still contains the plastic particles and additional solutions are necessary to filter them out.

Along the same lines, scrap material causes environmental issues. Any of the techniques from section "Minimizing scrap material" in chapter *machine knowledge* help to improve sustainability.

6.3 Recycling

Once a fabricated object becomes obsolete, the object itself becomes waste. This is an opportunity for material recycling. Filament extruders, such as the *Recyclebot* [Baechler et al., 2013] (now commercialized under the name *FilaBot* [Faludi et al., 2015]), allow users to shred their obsolete objects, melt the resulting pieces, and extrude them to form new filament (Figure 6.1).

A similar process can be applied to objects created by thermoforming: The *Dishmaker* [Bonanni et al., 2005], for instance, creates bowls from flat sheets of plastic by first warming them up and then forming them with a heat stamp. When not used anymore the *Dishmaker* returns them to their original flat shape by applying heat and pressure (Figure 6.2).

Recycling is much more difficult for multi-material prints, as the materials have to be separated from each other in the process. Hiller and Lipson [2007] propose addressing this issue by assembling objects from prefabricated magnetic/non-magnetic "voxels" using a water-soluble binder. As can be seen in Figure 6.3, after dissolving objects in water

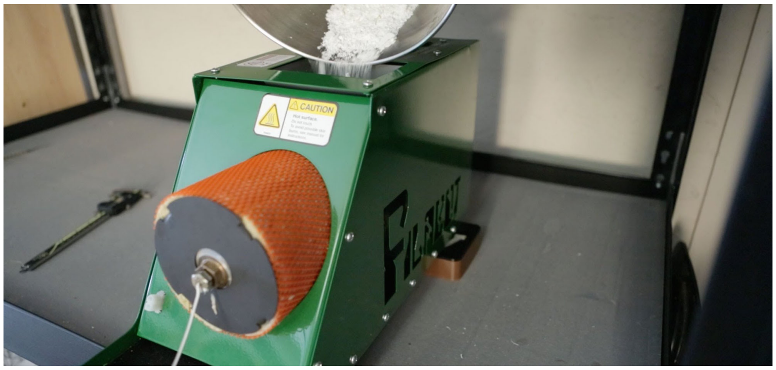

Figure 6.1: Filament extruders re-melt objects into new filament, here: *Filabot* [Faludi et al., 2015].

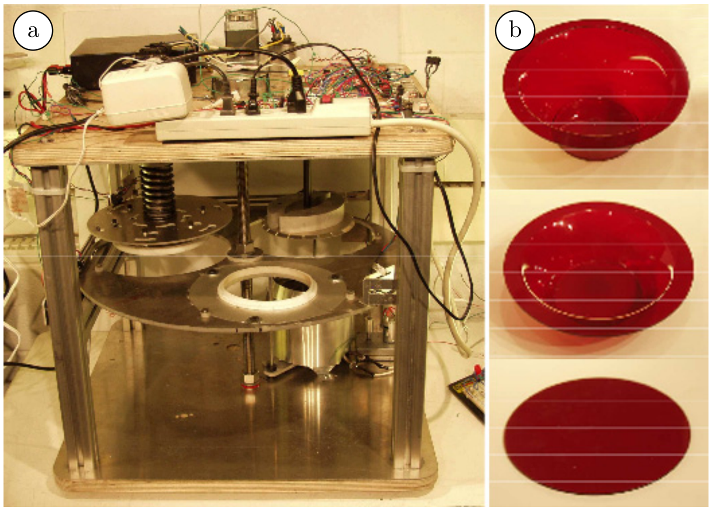

Figure 6.2: The *Dishmaker* [Bonanni et al., 2005] recycles dishes using thermoforming.

this allows separating voxels using a magnet in order to recycle them [Hiller and Lipson, 2009].

A more immediate route is to avoid multi-material 3D printing in the first place and instead emulate multi-material properties using

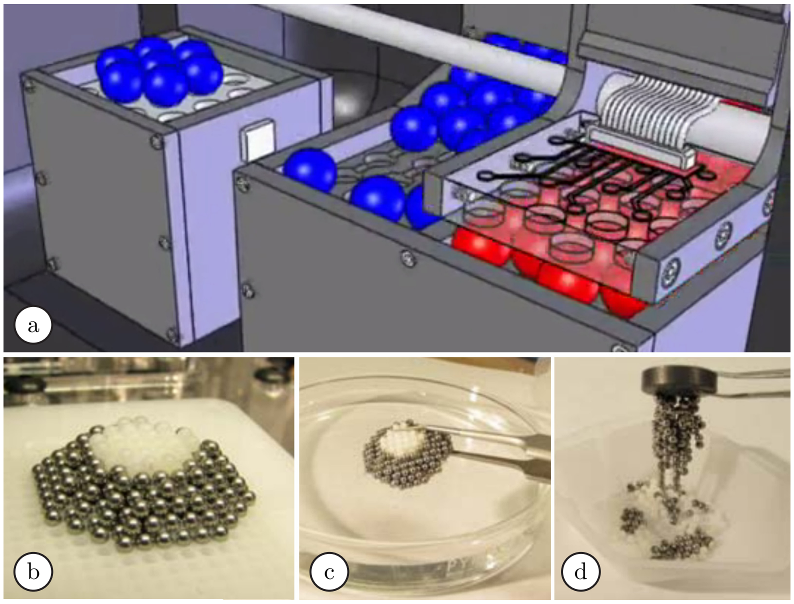

Figure 6.3: Hiller and Lipson [2007] propose multi-material printing magnetic/not magnetic elements for easy recycling using a magnet.

metamaterials [Bickel et al., 2010] (see section "Metamaterials" in chapter *material*). Objects made from metamaterials are easy to recycle as they tend to consist of a single material.

But even in situations where recycling is possible, it is worth checking whether it will actually result in a positive net balance. Filament recycling, for example, consumes energy and requires additional machinery. Furthermore, the plastics used in FDM get brittle after a few rounds of recycling, which leads to decreased material performance, and may thus not allow for useful recycling anymore.

Recycling, as applied to mass produced goods today, works in large parts because of the scale and uniformity of the involved objects — the mass customization and distributed nature of personal fabrication, in contrast, seem to be at odds with this process — at least with how we think of recycling today.

Biodegradable materials, such as *PLA*, may sometimes serve as an alternative approach here. While the distributed nature of personal fabrication can result in material quantities too small to recycle, letting material degrade works even for small quantities.

In contrast, we might choose to avoid materials, for which there is no known recycling process, such as stereolithography; this process is based on curing photosensitive polymers, an essentially irreversible process.

6.4 Avoiding fabrication

Another approach to reducing the impact of personal fabrication on sustainability is to reduce the number of objects we produce. For instance, we may try to iterate less during the design phase. We can achieve this by using software that embodies domain knowledge or by using software that provides previews, as described in the previous chapters.

In addition, today, we choose to discard and refabricate object at various occasions: when the print mechanically fails, when the resulting object does not meet the requirements during testing, when our requirements change over time, or when an object breaks [Teibrich et al., 2015]. For any of these reasons, we tend to reprint the object in its entirely. Teibrich et al. [2015] propose *patching* objects instead, which can be accomplished with less material. Figure 6.4 shows their system: a milling head removes flawed or otherwise undesired geometry; the FDM 3D printing head then adds the new desired geometry.

6.5 Conclusion and open research questions

In this chapter, we discussed how to make the 3D printing process more sustainable by reducing waste and material use. There are still a lot of open research questions, including the recycling of objects produced by means of stereolithography, laser cutting, etc. Another factor worth discussing is the higher energy use compared to mass manufacturing, which is discussed in mechanical engineering and economic analysis (see Faludi et al. [2015]).

However, if these issues should be solved well, personal fabrication may even hold a sustainability promise. In environmental science,

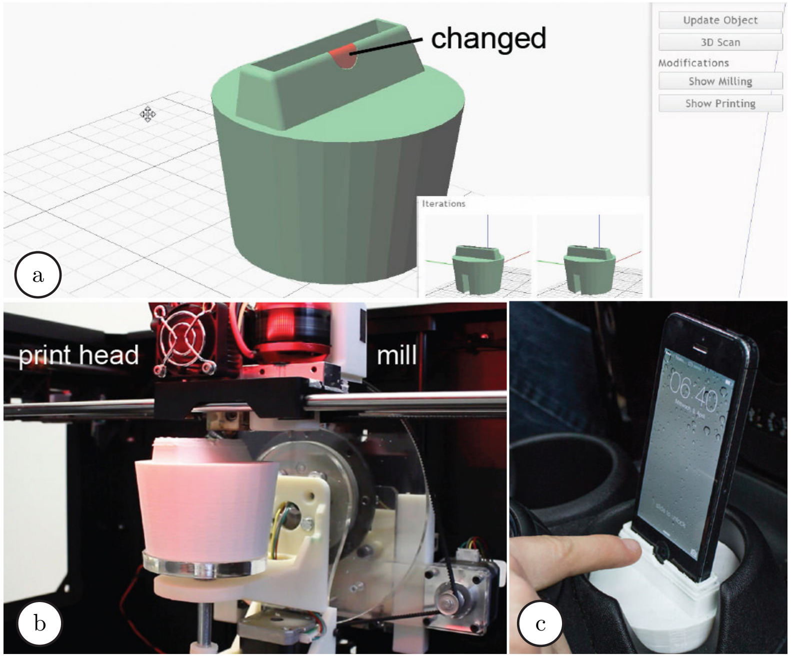

Figure 6.4: *Patching Physical Objects* [Teibrich et al., 2015] removes and reprints only the unsatisfactory parts instead of printing the object from scratch.

researchers are currently discussing the question whether 3D printing might have the potential to be more sustainable than mass manufacturing [Gebler et al., 2014]. Factors include reduced shipping volumes, as raw materials rend to be more compact, reduced storage, as objects are typically made on demand, and extended lifespan, as users tend to be more attached to objects custom made for them [Diegel et al., 2010].

However, the true sustainability problem with personal fabrication may be elsewhere. While the research projects discussed above identify the sustainability problem on a per-object level, we argue that the biggest sustainability problem may well be posed by the fabrication hardware itself. In particular, 3D printers appear to be continuously undergoing substantial technological advances, which could motivate users to replace these machines frequently, as the next device offers better printing resolution, more printable materials, speed and interactivity.

Our concern is once again inspired by the field of personal computing, where fast advances in hardware continue to cause millions of users to replace their computing devices on a regular basis [Yu et al., 2010]. This has resulted in substantial amounts of *e-waste* [Widmer et al., 2005], a trend that perpetuates itself today with the rapid replacement of mobile phones.

While there is little data today that could prove a Moore's law for 3D printers, 3D systems executive Merrill Lynch stated in his 2014 keynote that 3D printing speed for their products on average had doubled every 24 months over the last 10 years [Krassenstein, 2014]. If such a trend should indeed materialize, an increased distribution of personal fabrication hardware among consumers could ultimately lead to a correspondingly large resource and trash problem as in the computing industry.

In the future, we might tackle this problem by sharing fabrication machinery [e.g., 3D Hubs] and by upgrading hardware instead of replacing it. If we look at the history of personal computing, upgradable hardware did not always produce the desired effect, as mass-produced machines tended to be cheaper than selected upgrades. Personal fabrication might actually come out ahead here, if we manage to create fabrication machines capable of upgrading themselves and each other [Vinge, 1993].

7

Intellectual Property

Historically, the AD/DA pattern has always resulted in intellectual property rights disputes. As of today, intellectual property lawyers are sending out cease and desist orders for images, music, and movies. In the future, will we sue each other over the replication of physical objects?

7.1 A historical perspective

To use an example, recording music from the radio to compact cassette was illegal in Germany during large parts of the 1970s. At the same time, it was common practice to do so, as audio cassette recorders were available to consumers since 1963 and had received substantial market success with more than 2.4 million sold players within the first five years [Braun, 2002]. As a result, many owners of cassette recorders found themselves performing copyright violations on a regular basis.

Just a decade earlier, the situation had been much simpler, when the possession of music and the ownership of music were one and the same: People were able to listen to music if and only if they owned the album. One could argue that the album physically incorporated not only the audio medium, but also the respective license to play the music. If one sold the album, the buyer would obtain the medium *and the license*.

This fact changed when cassette recorders appeared on the mass market, as they made it possible to possess the music without owning a license. We might say that medium and license “dissociated”.

This unsettling situation was eventually resolved in 1965 in Germany when the *German society for musical performing and mechanical reproduction rights* (GEMA) managed to charge a few cents for each blank audio cassette sold, which served as an in-advance compensation for the music that was expected to be recorded on the cassette [Private Copying Levy]. At least, users were now allowed to record music from the radio.

This settled the topic for two decades — until in 1982 people started copying compact discs and then moved on to sharing digital music online, reviving the 1960s situation on a much bigger scale and with perfect copies through lossless, digital reproduction. Sharing platforms, such as [Napster], became the primary means for the distribution of illegal copies at scale.

Today, the conflict has turned into tens of thousands of cease and desist orders automatically going out each month.

With cease and desist orders also being sent out for sharing movies and for sharing copyrighted images [Getty Images], this pattern obviously applies not only to music, but also to the other occurrences of the AD/DA pattern discussed earlier.

7.2 Copyright issues involved in sharing 3D models

It is thus no surprise that the advent of 3D scanners and 3D printers has brought copyright infringement and the resulting cease and desist response to physical objects. While the idea may still feel somewhat unfamiliar, the pattern is exactly the same.

When we buy a physical object today, we may not realize that we are actually buying both the physical instance and the license to use it. In contrast, 3D printing allows us to fabricate an object without owning a license. Once a design is available in digital format, it can, for example, be shared and fabricated millions of times for the price of the material cost, infringing on the designer's copyright. For example, we may obtain a 3D model of a tablet computer stand that the designer/manufacturer normally sells to consumers for $10. If users can 3D print this model for $2 worth of material, there is a clear incentive for users to do it as they save substantial cost. At the same time, however, the manufacturer loses the margin and the designer who created the object in the first place will not get paid.

As a result, 3D model platforms, such as *Thingiverse*, are starting to experience the same backlash that file-sharing services, such as *Napster*, experienced. For instance, in 2011, designer Ulrich Schwanitz successfully sued *Thingiverse* because one of his designs had been reverse-engineered and a similar file posted on the site [Weinberg, 2013]. A year later in 2012, *The Pirate Bay* [Rideout, 2011], a popular page for illegal sharing of files, announced that it is adding a 'Physibles' section for sharing 3D printable content.

Generally not subject to copyright issues is the "*remixing*" of 3D models. In the same way that *Youtube* allows users to download videos, remix them, and re-upload, 3D models can nowadays be remixed and shared easily. *Thingiverse*' 'remix' section for a 3D model, for instance, provides insight into how the model evolved from existing models in the collection.

Recently, a survey found that 80% of top 3D designers do not share their design for fear of theft [Authentise Interview]. This is understandable in the light of a study by Gartner that predicts that by 2018, 3D printing will result in the loss of at least $100 billion per year due to illegally shared content [Gao et al., 2015]. Thus, arguably, digital rights questions around personal fabrication are currently delaying adoption.

Several research and engineering projects are therefore attempting to address the intellectual property rights question and we discuss them in the following. We do not tackle questions that traditionally fall into the realm of patent law, but instead refer to the white papers on intellectual property rights by Weinberg [2010, 2013].

7.3 Digital rights management

The Authentise platform makes it difficult for users to print a 3D object purchased online more than once. *Authentise* sends out 3D models in the form of *gCode*, i.e., low-level instructions that instruct a *specific* machine where to extrude material, how fast and how much, in order to produce the respective object. While users *can* share *gCode* with each other, the machine-specificity makes the received code useless, unless the recipient of the code has access to exactly the same type of machine.

In addition, *Authentise* does not offer the *gCode* for download, but instead streams it in increments, again making it harder for customers to illegally capture the code.

Finally, *Authentise* encrypts *gCode* before streaming it and decrypts it on the target 3D printer, which prevents men-in-the-middle from intercepting and capturing the *gCode* stream.

Authentise' approach is not dissimilar from copy protections used in other AD/DA domains, such as DRM for streaming music or streaming movies. It also comes with the same side effects, for example, it requires a permanent connection to the server. For more details see Kerikmäe and Rull's [2016] article on *3D Printing Using Secured Streaming*.

While these approaches provide a great step forward towards protecting intellectual property rights, they cannot offer perfect security. While DRM mechanisms can make it hard to obtain object files when downloading, there is at the very least always the possibility of analog reproduction, i.e., users may scan (at least the outside) of objects they have purchased, and then re-fabricate them.

7.4 Certifying object authenticity

Intellectual property does not only protect the designer, but also the user of the design: When downloading a specific object, users might want to know that the object is indeed from the declared source, even more so if the object is supposed to perform a critical function.

One approach is to physically print hidden identifiers, such as *watermarks*, into each object, which is also known as *steganography*. The

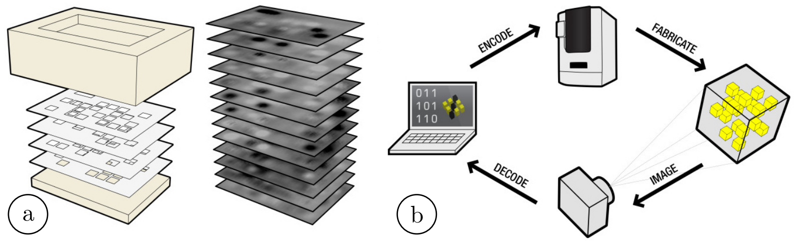

Figure 7.1: Integrating markers to identify physical objects (*Infrastructs* [Willis and Wilson, 2013]): (a) the marker, (b) the encode, fabricate, image, decode process.

[Microtrace Microtaggant technology] proposes *chemical* watermarking: During the 3D printing process, the filament is mixed, e.g., with special inks or particles that do not influence the function or appearance of the object, but can easily be used to identify it. Willis et al. [2013] propose watermarking by embedding geometric markers into the infill of objects (Figure 7.1). Invisible from the outside, such markers can be read using *TeraHertz* scanners (*InfraStructs* [Willis and Wilson, 2013]).

While embedding the ID of the manufacturer can be used to guarantee the authenticity of a part, embedding the ID of the customer who licensed the part can be used to trace back illegal copies for the purpose of prosecution.

7.5 Transferring a license

Selling 3D objects has to work differently when the seller is a consumer rather than the copyright owner discussed above. This situation occurs for example when a consumer is reselling a used object on a trading site, such as *eBay*. With web-based search and instant electronic payment, the entire process of finding and paying for used objects can be done within minutes; only, the buyer has to wait days or weeks for the object to arrive via physical mail. Fabrication-based machines, such as the 3D fax machine invented in 1991 [Reyes, 1991], could offer quasi-instant delivery: sellers would simply place the sold object into their 3D fax machine and the receiving 3D fax unit starts to fabricate a copy. With

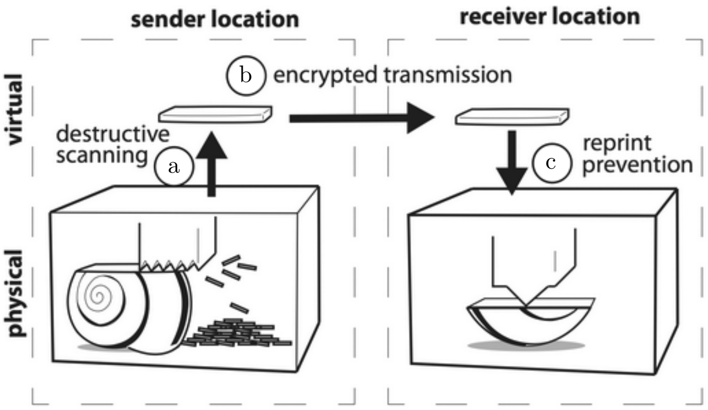

Figure 7.2: (a) "Transfer" of a 3D object from sender to receiver using two *Scotty* [Mueller et al., 2015b] units. (b) The devices ensure that the object remains unique during scanning, transmission, and refabrication.

up-to-date 3D printing technology [e.g., Carbon3D, 2015], this could be as fast as a few minutes.

What holds this scenario back is the fact that it creates an illegitimate copy. While the copyright owners discussed earlier have the right to make copies and to create additional licenses, consumers generally own only the one license they have previously acquired with the physical instance. Thus, when consumers sell, they naturally have to sell their own license, i.e., they stop owning their license. The 3D fax machine violates this, as it produces a situation with seller and buyer both in possession of the object.

Scotty [Mueller et al., 2015b] is a transmission device that addresses this situation by destroying the seller's object in the process. As illustrated by Figure 7.2, the destruction proceeds in three steps.

First, during scanning, *Scotty* destroys the seller's object as part of the scanning process, i.e., by shaving off one layer at a time with the built-in milling machine. Each layer is captured with the built-in camera. Second, during transmission, *Scotty* deletes the sender's local copy. *Scotty* also encrypts the object using the receiver's public key, which prevents men-in-the-middle intercepting the transmission from fabricating a copy of the object. Third, once the object has been refabricated, the receiving *Scotty* unit deletes its digital representation of the object. It also maintains an eternal log of objects already fabricated, preventing a re-submission from producing additional copies. As a result, *Scotty* "transfers" the object rather than copying it, resulting in a single object at the buyer's site — this object that now bears the legitimate license, as the destruction of the original freed the license there.

7.6 Conclusion and open research questions

Three-dimensional printing is subject to similar challenges with respect to intellectual property rights as other types of media characterized by the AD/DA pattern.

We see two possible directions for the future: The most obvious direction would be to ask what we can do to get the digital rights management of personal fabrication to the same level as the DRM of images, videos, and audio. We may (continue to) analyze the technologies that have evolved in the contexts of these types of media, such as how to decouple the rights for a digital content from download in music and movies (*Apple* [Zweig and Woodyatt, 2015]) and port them to the specific nature of fabrication technologies (see this DRM patent on 3D printing [Jung et al., 2012] as an example). This would lead to the development of techniques for streaming encrypted content to laser cutters and milling machines and watermarking the resulting objects. In analogy to media rental platforms such as the *iTunes* store and media subscription services, such as *Spotify*, we could then design rental and subscription services for 3D objects.

As an alternative direction, we may ask ourselves whether there are better ways to handle intellectual property rights than we have seen in the past with images, audio, and video. Arguably, the digital rights management of those media is all but solved in the first place, with cease and desist orders going out in the ten thousands every month. This is not a desirable situation and the same way that it hampers progress in the realm of other types of media, an unresolved intellectual property situation could slow down the evolution and adoption of personal fabrication technology. Addressing intellectual property questions is therefore, arguably, of paramount importance.

8

Conclusions and Outlook

In this monograph, we investigated the transition of fabrication technology to consumers. We analyzed similar transition in the past — in personal computing and in what we call AD/DA media, i.e., images, audio, and video. We observed that for all types of media, the transition to consumers was preceded by the emergence of hardware + software systems that embodied the skills and expert knowledge that consumers lacked. We concluded that research in personal fabrication will have to address the same challenges if personal fabrication is to ever reach consumers.

We broke the challenge of reaching consumers down into four + two individual challenges: (1) *hardware and materials* that allow fabricating the intended object, (2) software that embodies *domain knowledge*, (3) software that provides immediate feedback and supports *interactive exploration*, (4) software that embodies the know-how required to operate the new *machinery;* as well as the two challenges required for long-term success: (5) *sustainability*, and (6) *intellectual property.*

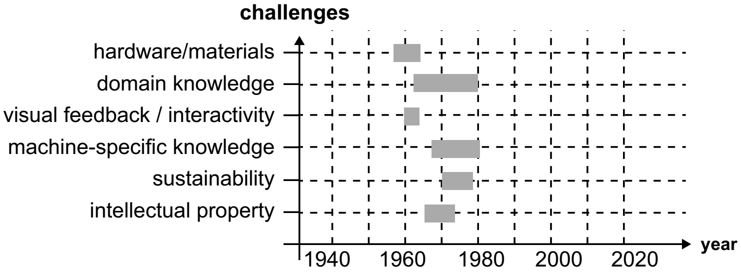

Figure 8.1: Where personal fabrication stands today with respect to personal computing and in particular previous AD/DA media.

We used these six challenges to lay out the related work in personal fabrication. The fact that the related work fits into this structure suggests that this is an appropriate categorization, and also implies that we are already on the way to solving these challenges.

We argued that the analogy to personal computing that we made for each challenge will continue to hold in the future. In the six chapters above, we already exploited this assumption in that we extrapolated how various aspects of personal fabrication might evolve and what specific research questions could present themselves as a result. In order to do so, we mapped the personal fabrication timeline to the personal computing timeline, as summarized by Figure 8.1.

With this reasoning in mind, everything located right of the markers in Figure 8.1 has the potential to be future work for personal fabrication. With other words, if we are trying to answer "where might personal fabrication go next and what research questions should we tackle" one strategy might be to find the corresponding point in time in the history of personal computing and ask "where did personal computing go after it had reached this point in its evolution and what research questions had to be answered in order to get there?"

While we applied this approach so far only to the six challenges in isolation, the approach also applies to personal fabrication as a whole: If personal fabrication today maps roughly to the 1970s in personal computing — what might happen next, i.e., what happened next in computing?

The answer is: a lot. Here are a few selected trends that took place in computing and that might be worth exploring in the context of fabrication.

1. Mobile fabrication: In 1981, the first mass-produced *portable* microcomputer was released, the [Osborne 1 Portable Microcomputer, 1981]. While it offered only 64 Kb memory while weighing 24 pounds, it laid the foundation for an entire new industry reaching from today's laptop computers and tablets, to mobile phones. Once computing was available anywhere and not just at the stationary desktop, new fields such as context-aware computing emerged that processed data depending on the user's surroundings. When applying this to personal fabrication, we may ask what will happen once (high-speed) fabrication devices are miniaturized to the extent that users will carry them at all times, allowing them to fabricate anything anywhere anytime. We refer to *Mobile Fabrication* [Roumen et al., 2016] for a first exploration in this direction.

2. Collaborative fabrication and Social fabrication: In 1984, Irene Greif and Paul Cashman defined the term *Computer-Supported Cooperative Work*. It addresses how collaborative activities and their coordination can be supported by means of computer systems. In the context of fabrication, it might be interesting to explore what will happen when users are connected by personal fabrication systems that produce synchronized physical output in real-time.

3. Ubiquitous fabrication, pervasive fabrication: In 1988, Mark Weiser [1999] coined the term *Ubiquitous computing*. In the context of personal fabrication, we might investigate what happens when consumers do not own a *single* personal fabrication device anymore, but if many devices distributed across their household offer the ability to fabricate.

4. Shared repositories of physical objects. In 1986, the first shared repository for writing software was developed by Dick Grune. In the context of fabrication, we start to see users share 3D models intended for fabrication in online repositories, such as *Thingiverse* and *Instructables*. However, there is tremendous potential for advancement here, including version control, editing by the crowds (in analogy to

Wikipedia, 2001), and rigorous reputation management (as in *Stack Overflow*, 2008). In the future, such carefully designed sharing sites might well become points of exchange between technology enthusiastic makers, who contribute 3D models, and consumers, who retrieve such models.

5. Physical object synthesis. In 2006, photo-synthesis software systems, such as *Photo Tourism* [Snavely et al., 2006], introduced the ability to aggregate thousands of images into coherent panoramas and videos of locations. As fabrication transitions to consumers, the sheer numbers of 3D models will create new opportunities for data mining, machine learning, and related aggregation technologies. Unlike makers, consumers may not necessarily work together; however, the data they produce could potentially be harvested and aggregated by computer systems, including parametric models aggregated from hundreds of models of the same object.

6. Open source/crowd fabrication. In 2009, *crowd computing* introduced the idea of using the ability of the world's population to collaborate on large and often global projects. In programming, large open source projects have been successful at solving hard challenges, such as the creation and maintenance of operating systems (e.g., *Linux)*. In the context of fabrication, crowds might be leveraged to solve hard mechanical engineering problems, in particular those that catch the interest of large audiences (crowd-funded hardware projects on platforms such as *Kickstarter* could potentially provide insights into popular topics). Candidate questions might reach from sustainability, alternative energy, and adaptation of technology to local environment to moonshot projects, such as devices intended to be used in space exploration missions, in which people participate out of curiosity. A first project on crowd-sourced fabrication [Lafreniere et al., 2016] targeting the collaborative assembly of large structures provides initial insights into the potential of this research area.

Obviously, these are just a few selected examples — any analysis of the history of personal computing allows us to extract many more.

8.1 When will we get there?

So how long will it take to get there — when will personal fabrication reach consumers?

To consider personal computing as a reference one last time, let us see how long it took personal computing to reach consumers. If we pick Doug Engelbart's 1986 online system (*Mother of all demos* [Engelbart, 1962]) as the beginning of personal computing, we could argue that it took only a bit over a decade until the main interaction paradigms were worked out (here the graphical user interfaces) and the first commercial system that implemented the concepts shipped (*Xerox Star* in 1980, *Apple Macintosh* in 1984).

These devices were not marketed to consumers, though; they were designed to address office workers; people with less expertise than engineers in industry, yet certainly still professionals.

If we ask when personal computing truly reached consumers, one might argue that it took another *four* decades to reach that point, a period during which personal computing adopted several additional key concepts, including (1) narrower, more focused functionality, a concept first seen with "information appliances" (Jeff Raskin, 1979 [Norman, 1998]), (2) easy maintenance as the result of standardized hardware, a concept first seen with network computers (Oracle, 1996), and (3) a new post-GUI interface, in which content was in the center of attention, rather than functions (also referred to as "natural user interfaces" [Wigdor and Wixon, 2011]). The first time all these concepts came together, arguably was in the form of the mobile touch devices of the late 2000s. These not only offered a particularly easy to use post-GUI touch interface, but were also backed by an app store, and finally could be operated and maintained stand-alone as the owner's sole device (iOS 5's *PC Free* feature, 2012). *54 years* after Doug Engelbart's original demo, consumers were finally able not only to operate, but also to maintain their personal computing devices.

If we now consider our journey towards personal fabrication in the hands of consumers, we need to acknowledge that our journey has only just begun. If we pick 2009 as our starting point, i.e., when the first FDM patent ran out and the *MakerBot Cupcake CNC* appeared on the market, we can see that we are clearly still at the very beginning. And if personal fabrication today feels like a niche technology for hobbyists, it is most likely because we still have decades to go. We should look at the in-between progress with patience: The success of personal computing, as we see it today, i.e., not in the form of desktop computers, but in the form of mobile devices, could certainly not have been predicted until decades after its inception. And if personal fabrication should turn out anything like it, we have an amazing journey ahead of us.

References

3D Hubs. https://www.3dhubs.com.

3Doodler. http://the3doodler.com.

3MF File Format. http://3mf.io/wp-content/uploads/2015/04/3MFcoreSpec_1.0.1.pdf, 2015.

Adobe Photoshop Elements. http://www.adobe.com/products/photoshop-elements.html.

H. Agrawal, U. Umapathi, R. Kovacs, J. Frohnhofen, H.-T. Chen, S. Mueller, and P. Baudisch. Protopiper: Physically sketching room-sized objects at actual scale. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology* (UIST '15), pages 427–436, 2015.

B. Y. Ahn, S. B. Walker, S. C. Slimmer, A. Russo, A. Gupta, E. B. Duoss, T. F. Malkowski, and J. A. Lewis. Planar and three-dimensional printing of conductive inks. *Journal of Visualized Experiments*, 58(3189):1–8, 2011.

AMF File Format. http://amf.wikispaces.com.

Anoto Pen. http://www.anoto.com.

Apple Pages. http://www.apple.com/mac/pages.

R. Arisandi, Y. Takami, M. Otsuki, A. Kimura, F. Shibata, and H. Tamura. Enjoying virtual handcrafting with tool device. In *Adjunct Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology* (UIST Adjunct Proceedings '12), pages 17–18, 2012.

Authentise. Secure. http://authentise.com/tech/secure.pdf.

Authentise Interview. http://3dprinting.com/news/talking-3d-printing-authentise.

Autodesk Inventor. http://www.autodesk.com/products/inventor/overview.

M. Bächer, B. Bickel, D. James, and H. Pfister. Fabricating articulated characters from skinned meshes. *ACM Transactions on Graphics* Article 47, 31 (4), 2012.

M. Bächer, E. Whiting, B. Bickel, and O. Sorkine-Hornung. Spin-it: Optimizing moment of inertia for spinnable objects. *ACM Transactions on Graphics* Article 96, 33(4), 2014.

M. Bächer, S. Coros, and B. Thomaszewski. Linkedit: Interactive linkage editing using symbolic kinematics. *ACM Transactions on Graphics* Article 99, 34(4), 2015.

C. Baechler, M. DeVuono, and J. M. Pearce. Distributed recycling of waste polymer into reprap feedstock. *Rapid Prototyping Journal*, 19(2):118–125, 2013.

G. Bharaj, S. Coros, B. Thomaszewski, J. Tompkin, B. Bickel, and H. Pfister. Computational design of walking automata. In *Proceedings of the 14th ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (SCA '15), pages 93–100, 2015a.

G. Bharaj, D. I. W. Levin, J. Tompkin, Y. Fei, H. Pfister, W. Matusik, and C. Zheng. Computational design of metallophone contact sounds. *ACM Transactions on Graphics* Article 223, 34(6), 2015b.

B. Bickel, M. Bacher, M. A. Otaduy, H. R. Lee, H. Pfister, M. Gross, and W. Matusik. Design and fabrication of materials with desired deformation behavior. *ACM Transactions on Graphics* Article 63, 29(4), 2010.

B. Bickel, P. Kaufmann, M. Skouras, B. Thomaszewski, D. Bradley, T. Beeler, P. Jackson, S. W. Marschner, and M. Gross. Physical face cloning. *ACM Transactions on Graphics* Article 118, 31(4), 2012.

L. Bonanni, S. Sarcia, S. Paudel, and T. Selker. Dishmaker: Personal fabrication interface. In *Proceedings of the 23rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems* (CHI EA '05), 2005.

H.-J. Braun. Music and technology in the twentieth century. *JHU Press*, 161, 2002.

A. Brunton, C. A. Arikan, and P. Urban. Pushing the limits of 3D color printing: Error diffusion with translucent materials. *ACM Transactions on Graphics* Article 4, 35(1), 2015.

J. Calì, D. A. Calian, C. Amati, R. Kleinberger, A. Steed, J. Kautz, and T. Weyrich. 3d-printing of non-assembly, articulated models. *ACM Transactions on Graphics* Article 130, 31(6), 2012.

Carbon3D. `http://carbon3d.com`, 2015.

D. Cardinal. From punched cards to siri. the history (and future) of input devices. `http://www.extremetech.com/computing/98287-from-punchcards-to-ipads-the-history-of-input-devices`, 2011.

D. Ceylan, W. Li, N. J. Mitra, Agrawala M., and M. Pauly. Designing and fabricating mechanical automata from mocap sequences. *ACM Transactions on Graphics* Article 186, 32(6), 2013.

D. Chen, D. I. W. Levin, P. Didyk, P. Sitthi-Amorn, and W. Matusik. Spec2Fab: A reducer-tuner model for translating specifications to 3D prints. *ACM Transactions on Graphics* Article 135, 32(4), 2013a.

D. Chen, P. Sitthi-Amorn, J. T. Lan, and W. Matusik. Computing and fabricating multiplanar models. *Computer Graphics Forum*, 32(2):205–315, 2013b.

W. Chen, X. Zhang, S. Xin, Y. Zia, S. Lefebvre, and W. Wang. Synthesis of filigrees for digital fabrication. *ACM Transactions on Graphics* Article 98, 35(4), 2016.

X. Chen, H. Zhang, J. Lin, R. Hu, L. Lu, Q. Huang, B. Benes, D. Cohen-Or, and B. Chen. Dapper: Decompose-and-pack for 3D printing. *ACM Transactions on Graphics* Article 213, 34(6), 2015.

S. Cho, Y. Heo, and H. Bang. Turn: A virtual pottery by real spinning wheel. In *ACM SIGGRAPH 2012 Emerging Technologies* (SIGGRAPH '12), Article 25, 2012.

M. Coelho and J. Zigelbaum. Shape-changing interfaces. *Personal Ubiquitous Computing*, 15(2):161–173, 2011.

S. Coros, B. Thomaszewski, G. Noris, S. Sueda, M. Forberg, B. Sumner, W. Matusik, and B. Bickel. Computational design of mechanical characters. *ACM Transactions on Graphics* Article 83, 32(4), 2013.

O. Diegel, S. Singamneni, S. Reay, and A. Withell. Sustainable product design through additive manufacturing. *KMUTNB: International Journal of Applied Science and Technology*, 4(1):39–45, 2010.

Digital Materials. `http://www.stratasys.com/materials/polyjet/digital-materials`.

Y. Dong, J. Wang, F. Pellacini, X. Tong, and B. Guo. Fabricating spatially-varying subsurface scattering. *ACM Transactions on Graphics* Article 62, 29(4), 2010.

T. Du, A. Schulz, B. Zhu, B. Bickel, and W. Matusik. Computational multicopter design. *ACM Transaction on Graphics* Article 227, 35(6), 2016.

J. Dumas, J. Hergel, and S. Lefebvre. Bridging the gap: Automated steady scaffoldings for 3D printing. *ACM Transactions on Graphics* Article 98, 33 (4), 2014.

J. C. Á. Elipe and A. D. Lantada. Comparative study of auxetic geometries by means of computer-aided design and engineering. *Smart Materials and Structures*, 21(10):105004, 2012.

D. C. Engelbart. Augmenting human intellect: A conceptual framework. *SRI Summary Report* AFOSR-3223, 1962.

Fabrican. http://www.fabricanltd.com.

J. Faludi, Z. Hu, S. Alrashed, C. Braunholz, S. Kaul, and L. Kassaye. Does material choice drive sustainability of 3D printing? *International Journal of Mechanical, Aerospace, Industrial and Mechatronics Engineering*, 9(2), 2015.

B. Florijn, C. Coulais, and M. van Hecke. Programmable mechanical metamaterials. *Physical Review Letters*, 113(17):175503, 2014.

S. Follmer, D. Carr, E. Lovell, and I. Hiroshi. CopyCAD: Remixing physical objects with copy and paste from the real world. In *Adjunct Proceedings of the 23nd Annual ACM Symposium on User Interface Software and Technology* (UIST '10), pages 381–382, 2010.

S. Follmer, D. Leithinger, A. Olwal, A. Hogge, and H. Ishii. inFORM: Dynamic physical affordances and constraints through shape and object actuation. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology* (UIST '13), pages 417–426, 2013.

Form1. *Formlabs*. https://www.kickstarter.com/projects/formlabs/form-1-an-affordable-professional-3d-printer, 2012.

Frame grabbing. https://en.wikipedia.org/wiki/Frame_grabber.

T. Fukuda and Y. Kawauchi. Cellular robotic system (CEBOT) as one of the realization of self-organizing intelligent universal manipulator. In *Proceedings of the IEEE International Conference on Robotics and Automation* (ICRA '90), pages 662–667, 1990.

M. Gannon, T. Grossman, and G. Fitzmaurice. Tactum: A skin-centric approach to digital design and fabrication. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI '15), pages 1779–1788, 2015.

M. Gannon, T. Grossman, and G. Fitzmaurice. ExoSkin: On-body fabrication. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (CHI '16), pages 5996–6007, 2016.

W. Gao, Y. Zhang, D. C. Nazzetta, K. Ramani, and R. J. Cipra. RevoMaker: Enabling multi-directional and functionally-embedded 3D printing using a rotational cuboidal platform. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology* (UIST '15), pages 437–446, 2015.

M. Gebler, A. J. S. Uiterkamp, and C. Visser. A global sustainability perspective on 3D printing technologies. *Energy Policy*, 74:158–167, 2014.

N. A. Gershenfeld. *Fab: The Coming Revolution on Your Desktop — from Personal Computers to Personal Fabrication*. Basic Books, 2005.

Getty Images. http://www.gettyimages.de.

Glowforge. Low-cost laser cutter. https://glowforge.com/, 2016.

S. C. Goldstein and T. C. Mowry. Claytronics: An instance of programmable matter. In *Wild and Crazy Ideas Session of ASPLOS*, 2004.

C. J. Hansen, R. Saksena, D. B. Kolesky, J. J. Vericella, S. J. Kranz, G. P. Muldowney, K. T. Christensen, and J. A. Lewis. High-throughput printing via microvascular multinozzle arrays. *Advanced Materials*, 25:1, 2013.

M. Hašan, M. Fuchs, W. Matusik, H. Pfister, and S. Rusinkiewicz. Physical reproduction of materials with specified subsurface scattering. *ACM Transactions on Graphics* Article 61, 29(4), 2010.

K. Hildebrand, B. Bickel, and M. Alexa. Orthogonal slicing for additive manufacturing. *Computers and Graphics*, 37(6):669–675, 2013.

J. Hiller and H. Lipson. Methods of parallel voxel manipulation for 3D digital printing. In *Proceedings of the 18th Solid Freeform Fabrication Symposium*, pages 200–211, 2007.

J. Hiller and H. Lipson. Fully recyclable multi-material printing. In *Proceedings of the 20th Solid Freeform Fabrication Symposium*, 2009.

M. A. Hiltzik. *Dealers of Lightning: Xerox PARC and the Dawn of the Computer Age*. HarperBusiness, 2000.

HP Jet Fusion. http://www8.hp.com/us/en/pdf/printers/3d-printers/4AA6-4892ENA-P.pdf.

Y. Huang, J. Zhang, X. Hu, G. Song, Z. Liu, L. Yu, and L. Liu. FrameFab: Robotic fabrication of frame shapes. *ACM Transactions on Graphics* Article 224, 35(6), 2016.

N. Hudson, C. Alcock, and P. K. Chilana. Understanding newcomers to 3D printing: Motivations, workflows, and barriers of casual makers. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (CHI '16), USA, pages 384–396, 2016.

S. E. Hudson. Printing teddy bears: A technique for 3D printing of soft interactive objects. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '14), pages 459–468, 2014.

M. B. Hullin, I. Ihrke, W. Heidrich, T. Weyrich, G. Damberg, and M. Fuchs. State of the art in computational fabrication and display of material appearance. *Eurographics 2013 State-of-the-Art Report* (STAR), 2013.

Y. Igarashi, T. Igarashi, and H. Suzuki. Interactive cover design considering physical constraints. *Computer Graphics Forum*, 28(7):1965–1973, 2009.

Instructables. www.instructables.com.

A. Ion, J. Frohnhofen, L. W. Wall, R. Kovacs, M. Alistar, J. I. C. Lindsay, P. Lopes, H.-T. Chen, and P. Baudisch. Metamaterial mechanisms. In *Proceedings of the 29th Annual ACM Symposium on User Interface Software and Technology* (UIST '16), pages 529–539, 2016.

H. Ishii and B. Ullmer. Tangible bits: Towards seamless interfaces between people, bits and atoms. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems* (CHI '97), pages 234–241, 1997.

T. Jansen. Propeller for animaris geneticus bellus. http://www.shapeways.com/product/RXW7P5ZLU/propeller-for-animaris-geneticus-bellus?optionId=55918691.

E. K. Y. Jung, R. W. Lord, M. A. Malamud, J. D. Rinaldo, Jr., C. T. Tegreene, and L. L. Wood, Jr. Manufacturing control system. *U.S. Patent 8286236*, filed January 31, 2008, issued September 10, 2012.

M. Kadic, T. Bückmann, R. Schittny, and M. Wegener. Metamaterials beyond electromagnetism. *Reports on Progress in Physics*, 76(12):126501, 2013.

T. Kerikmäe and Addi Rull. The future of law and etechnologies. In *Intellectual Property Protection of 3D Printing Using Secured Streaming.* Springer Book, 2016.

Y. Kim and S.-H. Bae. SketchingWithHands: 3D sketching handheld products with first-person hand posture. In *Proceedings of the 29th Annual ACM Symposium on User Interface Software and Technology* (UIST '16), pages 797–808, 2016.

B. Koo, W. Li, J. Yao, M. Agrawala, and N. J. Mitra. Creating works-like prototypes of mechanical objects. *ACM Transactions on Graphics* Article 217, 33(6), 2014.

B. Krassenstein. 3D Printing and Moore's law. `https://3dprint.com/7543/3d-printing-moores-law`, 2014.

B. Lafreniere, T. Grossman, F. Anderson, J. Matejka, H. Kerrick, D. Nagy, L. Vasey, Evan Atherton, N. Beirne, M. H. Coelho, N. Cote, S. Li, A. Nogueira, L. Nguyen, T. Schwinn, J. Stoddart, D. Thomasson, R. Wang, T. White, D. Benjamin, M. Conti, A. Menges, and G. Fitzmaurice. Crowdsourced fabrication. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology* (UIST '16), pages 15–28, 2016.

T. Langlois, A. Shamir, D. Dror, W. Matusik, and D. I. W. Levin. Stochastic structural analysis for context-aware design and fabrication. *ACM Transactions on Graphics* Article 226, 35(6), 2016.

G. Laput, X. A. Chen, and C. Harrison. 3D printed hair: Fused deposition modeling of soft strands, fibers, and bristles. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology* (UIST '15), pages 593–597, 2015.

M. Lau, A. Ohgawara, J. Mitani, and T. Igarashi. Converting 3D furniture models to fabricatable parts and connectors. *ACM Transaction on Graphics* Article 85, 30(4), 2011.

M. Lau, M. Hirose, A. Ohgawara, J. Mitani, and T. Igarashi. Situated modeling: A shape-stamping interface with tangible primitives. In *Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction* (TEI '12), pages 275–282, 2012.

J.-H. Lee. Haptic intelligentsia. `http://studio-homunculus.com/portfolio/haptic-intelligentsia-human-prototyping-machine`, 2011.

J.-H. Lee, J. P. Singer, and E. L. Thomas. Micro-/nanostructured mechanical metamaterials. *Advanced Materials*, 24(36):4782–4810, 2012.

S. J. Leigh, R. J. Bradley, C. P. Purssell, D. R. Billson, and D. A. Hutchins. A simple, low-cost conductive composite material for 3D printing of electronic sensors. *PLoS ONE*, 7(11), 2012.

J. A. Lewis and B. Y. Ahn. Device fabrication: Three-dimensional printed electronics. *Nature*, 518:42–43, 2015.

D. Li, D. I. W. Levin, W. Matusik, and C. Zheng. Acoustic voxels: Computational optimization of modular acoustic filters. *ACM Transactions on Graphics* Article 88, 35(4), 2016.

A. J. Lopes, MacDonald E., and R. B. Wicker. Integrating stereolithography and direct print technologies for 3D structural electronics fabrication. *Rapid Prototyping Journal*, 18(2):129–143, 2012.

L. Lu, A. Sharf, H. Zhao, Y. Wei, Q. Fan, X. Chen, Y. Savoye, C. Tu, D. Cohen-Or, and B. Chen. Build-to-last: Strength to weight 3D printed objects. *ACM Transactions on Graphics* Article 97, 33(4), 2014.

L. Luo, I. Baran, S. Rusinkiewicz, and W. Matusik. Chopper: Partitioning models into 3D-printable parts. *ACM Transactions on Graphics* Article 129, 31(6), 2012.

MakerBot Cupcake CNC. `http://www.makerbot.com/blog/tag/cupcake-cnc`, 2009.

T. Martin, N. Umetani, and B. Bickel. OmniAD: Data-driven omni-directional aerodynamics. *ACM Transaction on Graphics* Article 113, 34(4), 2015.

J. Martínez, J. Dumas, and S. Lefebvre. Procedural voronoi foams for additive manufacturing. *ACM Transactions on Graphics* Article 44, 35(4), 2016.

J. McCann, L. Albaugh, V. Narayanan, A. Grow, W. Matusik, J. Mankoff, and J. Hodgins. A compiler for 3D machine knitting. *ACM Transactions on Graphics* Article 49, 35(4), 2016.

J. McCrae, N. Umetani, and K. Singh. FlatFitFab: Interactive modeling with planar sections. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology* (UIST '14), pages 13–22, 2014.

MCor Technologies. `http://mcortechnologies.com`.

V. Megaro, B. Thomaszewski, D. Gauge, S. Coros, E. Grinspun, and M. Gross. ChaCra: An interactive design system for rapid character crafting. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (SCA '14), pages 123–130, 2014.

V. Megaro, B. Thomaszewski, M. Nitti, O. Hilliges, M. Gross, and S. Coros. Interactive design of 3D-printable robotic creatures. *ACM Transactions on Graphics* Article 216, 34(6), 2015.

Microsoft PowerPoint. `https://products.office.com/en-us/powerpoint`.

Microtrace Microtaggant technology. `http://www.microtracesolutions.com/taggant-technologies/`.

M. Mir, M. N. Ali, J. Sami, and U. Ansari. Review of mechanics and applications of auxetic structures. *Advances in Materials Science and Engineering 2014,* Article 753496, pages 1–17, 2014.

J. Mitani and H. Suzuki. Making papercraft toys from meshes using strip-based approximate unfolding. *ACM Transactions on Graphics*, 23(3):259–263, 2004.

G. A. Moore. *Crossing the Chasm: Marketing and Selling High-Tech Products to Mainstream Customers*. HarperBusiness, 2006.

Y. Mori and T. Igarashi. Plushie: An interactive design system for plush toys. *ACM Transactions on Graphics* Article 45, 26(3), 2007.

S. Mueller, P. Lopes, and P. Baudisch. Interactive construction: Interactive fabrication of functioning mechanical devices. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology* (UIST '12), pages 599–606, 2012.

S. Mueller, B. Kruck, and P. Baudisch. LaserOrigami: Laser-cutting 3D objects. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI'13), pages 2585–2592, 2013.

S. Mueller, S. Im, S. Gurevich, A. Teibrich, L. Pfisterer, F. V. Guimbretière, and P. Baudisch. WirePrint: 3D printed previews for fast prototyping. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology* (UIST '14), pages 273–280, 2014a.

S. Mueller, T. Mohr, K. Guenther, J. Frohnhofen, and P. Baudisch. FaBrickation: Fast 3D printing of functioning objects by integrating construction kit building blocks. In *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems* (CHI'14), pages 3827–3834, 2014b.

S. Mueller, D. T. S. Beyer, A. Teibrich, L. Pfisterer, K. Guenther, J. Frohnhofen, H.-T. Chen, P. Baudisch, S. Im, and F. V. Guimbretière. Low-fidelity fabrication: Speeding up design iteration of 3D objects. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems* (CHI EA '15), pages 327–330, 2015a.

S. Mueller, M. Fritzsche, J. Kossmann, M. Schneider, J. Striebel, and M. Schneider. Scotty: Relocating physical objects across distances using destructive scanning, encryption, and 3D printing. In *Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction* (TEI '15), pages 233–240, 2015b.

S. Mueller, A. Seufert, and Baudisch. P. FormFab. In preparation, 2017.

Muvee. `http://www.muvee.com`.

Nanoscribe. `http://www.nanoscribe.de/en/`.

Napster. `http://napster.com`.

Ninjaflex. `http://www.3dprima.com/en/ninjaflex.html`.

D. A. Norman. *The Invisible Computer: Why Good Products Can Fail, the Personal Computer is so Complex, and Information Appliances are the Solution*. MIT Press, 1998.

Osborne 1 Portable Microcomputer. `http://www.computinghistory.org.uk/det/504/osborne-1/`, 1981.

T. Oster and R. Bohne. Visicut — a userfriendly tool for laser-cutting. In *Bachelor Thesis RWTH Aachen*, 2011. `https://hci.rwth-aachen.de/visicut`.

J. Ou, G. Dublon, C.-Y. Cheng, F. Heibeck, K. D. D. Willis, and H. Ishii. Cilllia: 3D printed micro-pillar structures for surface texture, actuation and sensing. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (CHI '16), pages 5753–5764, 2016.

J. Panetta, Q. Zhou, L. Malomo, N. Pietroni, P. Cignoni, and D. Zorin. Elastic textures for additive fabrication. *ACM Transactions on Graphics* Article 135, 34(4), 2015.

P. Panozzo, O. Diamanti, S. Paris, M. Tarini, E. Sorkine, and O. Sorkine-Hornung. Texture mapping real-world objects with hydrographics. In *Proceedings of the Eurographics Symposium on Geometry Processing* (SGP '15), pages 65–75, 2015.

P. Peers, K. V. Berge, W. Matusik, R. Ramamoorthi, J. S. Lawrence, and P. Dutré. A compact factored representation of heterogeneous subsurface scattering. *ACM Transactions on Graphics*, 25(3):746–753, 2006.

H. Peng, J. Mankoff, S. E. Hudson, and J. McCann. A layered fabric 3D printer for soft interactive objects. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI '15), pages 1789–1798, 2015a.

H. Peng, A. Zoran, and F. V. Guimbretière. D-coil: A hands-on approach to digital 3D models design. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI '15), pages 1807–1815, 2015b.

H. Peng, F. Guimbretière, J. McCann, and S. E. Hudson. A 3D printer for interactive electromagnetic devices. In *Proceedings of the 29th Annual ACM Symposium on User Interface Software and Technology* (UIST '16), pages 553–562, 2016a.

H. Peng, R. Wu, S. Marschner, and F. V. Guimbretière. On-the-fly print: Incremental printing while modelling. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (CHI '16), pages 887–896, 2016b.

T. Pereira, S. Rusinkiewicz, and W. Matusik. Computational light routing: 3D printed optical fibers for sensing and display. *ACM Transactions on Graphics* Article 24, 33(3), 2014.

V. Perumal and D. Wigdor. Foldem: Heterogeneous object fabrication via selective ablation of multi-material sheets. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (CHI '16), pages 5765–5775, 2016.

R. Prévost, E. Whiting, S. Lefebvre, and O. Sorkine-Hornung. Make it stand: Balancing shapes for 3D fabrication. *ACM Transactions on Graphics* Article 81, 32(4), 2013.

Price Comparison 3D Printers. `http://www.3ders.org/pricecompare/3dprinters/`.

Private Copying Levy. `https://en.wikipedia.org/wiki/Private_copying_levy`.

Quantel Harry. Suite demo: Compositing 80's style. `https://vimeo.com/146640097`.

A. Rafsanjani, A. Akbarzadeh, and D. Pasini. Snapping mechanical metamaterials under tension. *Advanced Materials*, 27:5931–5935, 2015.

F. Ragnet. *The 'Less Paper' Office: How to Reduce Costs, Enhance Security and be a Better Global Citizen.* White Paper Xerox Global Services, 2008.

T. Reiner, N. Carr, R. Měch, O. Štava, C. Dachsbacher, and G. Miller. Dual-color mixing for fused deposition modeling printers. *Computer Graphics Forum*, 33(2):479–486, 2014.

RepRap. `http://reprap.org/`, 2005.

R. Reyes. World's first 3D fax machine. `https://webspace.utexas.edu/reyesr/self/3D_fax.htm`, 1991.

B. Rideout. Printing the impossible triangle: The copyright implications of three-dimensional printing. *Journal of Business Entrepreneurship & Law* Article 6, 5(1), 2011.

A. Rivers, I. E. Moyer, and F. Durand. Position-correcting tools for 2D digital fabrication. *ACM Transactions on Graphics* Article 88, 31(4), 2012.

O. Rouiller, B. Bickel, J. Kautz, W. Matusik, and M. Alexa. 3D printing spatially varying BRDFs. *IEEE Computer Graphics and Applications*, 33 (6):48–57, 2013.

T. J. Roumen, B. Kruck, T. Dürschmid, T. Nack, and P. Baudisch. Mobile fabrication. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology* (UIST '16), pages 3–14, 2016.

D. Saakes, T. Cambazard, J. Mitani, and T. Igarashi. PacCAM: Material capture and interactive 2D packing for efficient material usage on CNC cutting machines. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology* (UIST '13), pages 441–446, 2013.

G. Saul, M. Lau, J. Mitani, and T. Igarashi. Sketchchair: An all-in-one chair design system for end users. In *Proceedings of the Fifth International Conference on Tangible, Embedded, and Embodied Interaction* (TEI '11), pages 73–80, 2011.

V. Savage, X. Zhang, and B. Hartmann. Midas: Fabricating custom capacitive touch sensors to prototype interactive objects. In *Proceedings of the 25th annual ACM symposium on User interface software and technology* (UIST '12), pages 579–588, 2012.

V. Savage, C. Chang, and B. Hartmann. Sauron: Embedded single-camera sensing of printed physical user interfaces. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology* (UIST '13), pages 447–456, 2013.

V. Savage, R. Schmidt, T. Grossman, G. Fitzmaurice, and B. Hartmann. A series of tubes: Adding interactivity to 3D prints using internal pipes. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology* (UIST '14), pages 3–12, 2014.

V. Savage, A. Head, B. Hartmann, D. B. Goldman, G. Mysore, and W. Li. Lamello: Passive acoustic sensing for tangible input components. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI '15), pages 1277–1280, 2015.

R. Schmidt and N. Umetani. Branching support structures for 3D printing. In *ACM SIGGRAPH 2014 Studio* (SIGGRAPH '14), Article 9, 2014.

C. Schüller, D. Panozzo, A. Grundhöfer, H. Zimmer, E. Sorkine, and O. Sorkine-Hornung. Computational thermoforming. *ACM Transactions on Graphics* Article 43, 35(4), 2016.

A. Schulz, A. Shamir, D. I. W. Levin, P. Sitthi-amorn, and W. Matusik. Design and fabrication by example. *ACM Transactions on Graphics* Article 62, 33 (4), 2014.

C. Schumacher, B. Bickel, J. Rys, S. Marschner, C. Daraio, and M. Gross. Microstructures to control elasticity in 3D printing. *ACM Transactions on Graphics* Article 136, 34(4), 2015.

C. Schumacher, B. Thomaszewski, and M. Gross. Stenciling: Designing structurally-sound surfaces with decorative patterns. In *Symposium on Geometry Processing (SGP)*, 2016.

Y. Schwartzburg and M. Pauly. Fabrication-aware design with intersecting planar pieces. *Computer Graphics Forum*, 32(2):317–326, 2013.

Y. Schwartzburg, R. Testuz, A. Tagliasacchi, and M. Pauly. High-contrast computational caustic design. *ACM Transactions on Graphics* Article 74, 33(4), 2014.

Service, Robert F. The synthesis machine. *Science*, 347(6227):1190–1193, 2015.

Shapeoko. Low-cost milling machine. `http://carbide3d.com/shapeoko`, 2013.

R. Shilkrot, P. Maes, J. A. Paradiso, and A. Zoran. Augmented airbrush for computer aided painting (CAP). *ACM Transactions on Graphics* Article 19, 34(2), 2015.

J. Shim, C. Perdigou, E. R. Chen, K. Bertoldi, and P. M. Reis. Buckling-induced encapsulation of structured elastic shells under pressure. *Proceedings of the National Academy of Sciences*, 109(16):5978–5983, 2012.

B. Shneiderman. Direct manipulation: A step beyond programming languages. *Computer*, 16(8):57–69, 1983.

M. Shugrina, A. Shamir, and W. Matusik. Fab forms: Customizable objects for fabrication with validity and geometry caching. *ACM Transactions on Graphics* Article 100, 34(4), 2015.

Sintratec, 2014. `https://www.indiegogo.com/projects/sintratec-world-s-first-desktop-laser-sinterer{#}/`.

P. Sitthi-Amorn, J. E. Ramos, Y. Wangy, J. Kwan, J. Lan, W. Wang, and W. Matusik. Multifab: A machine vision assisted platform for multi-material 3D printing. *ACM Transactions on Graphics* Article 129, 34(4), 2015.

SketchUp. `http://www.sketchup.com`.

M. Skouras, B. Thomaszewski, S. Coros, B. Bickel, and M. Gross. Computational design of actuated deformable characters. *ACM Transactions on Graphics* Article 82, 32(4), 2013.

R. Slyper and J. Hodgins. Prototyping robot appearance, movement, and interactions using flexible 3D printing and air pressure sensors. In *Proceedings of the 21st IEEE International Symposium on Robot and Human Interactive Communication* (Ro-Man'12), pages 6–11, 2012.

N. Snavely, S. M. Seitz, and R. Szeliski. Photo tourism: Exploring photo collections in 3D. *ACM Transactions on Graphics*, 25(3):835–846, 2006.

SolidThinking. `http://www.solidthinking.com`.

SolidWorks. http://www.solidworks.com.

H. Song, F. Guimbretière, and H. Lipson. The modelcraft framework: Capturing freehand annotations and edits to facilitate the 3D model design process using a digital pen. *ACM Transactions on Computer-Human Interaction* (ToCHI) Article 14, 16(3), 2009.

P. Song, B. Deng, Z. Wang, Z. Dong, W. Li, C.-W. Fu, and L. Liu. CofiFab: Coarse-to-fine fabrication of large 3D objects. *ACM Transactions on Graphics* Article 45, 35(4), 2016.

O. Sosanya. 3D weaver. http://www.sosafresh.com/3d-weaver.

O. Stava, J. Vanek, B. Benes, N. Carr, and R. Měch. Stress relief: Improving structural strength of 3D printable objects. *ACM Transaction on Graphics* Article 48, 31(4), 2012.

I. Sutherland. *Sketchpad: A Man-machine Graphical Communication System*. PhD Thesis MIT, 1963.

M. Takezawa, T. Imai, K. Shida, and T. Maekawa. Fabrication of freeform objects by principal strips. *ACM Transactions on Graphics* Article 225, 35 (6), 2016.

A. Teibrich, S. Mueller, F. Guimbretière, R. Kovacs, S. Neubert, and P. Baudisch. Patching physical objects. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology* (UIST '15), pages 83–91, 2015.

Thingiverse. https://www.thingiverse.com/.

B. Thomaszewski, S. Coros, D. Gauge, V. Megaro, E. Grinspun, and M. Gross. Computational design of linkage-based characters. *ACM Transaction on Graphics* Article 64, 33(4), 2014.

R. Thompson. *Manufacturing Processes for Design Professionals*. Thames & Hudson, 2007.

C. Torres and E. Paulos. Metamorphe: Designing expressive 3D models for digital fabrication. In *Proceedings of the 2015 ACM SIGCHI Conference on Creativity and Cognition* (C&C '15), pages 73–82, 2015.

C. Torres, T. Campbell, N. Kumar, and E. Paulos. HapticPrint: Designing feel aesthetics for digital fabrication. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology* (UIST '15), pages 583–591, 2015.

E. Tsai, M. Firstenberg, J. Laucks, Y. Sterman, B. Lehnert, and N. Oxman. CNSilk: Spider-silk inspired robotic fabrication. In *Proceedings of Robotic Fabrication in Architecture* (Rob-Arch '12), pages 160–165, 2012.

Two photon lithography. `http://www.tuwien.ac.at/en/news/news_detail/article/7444`.

B. M. Tymraka, M. Kreiger, and J. M. Pearcea. Mechanical properties of components fabricated with open-source 3-D printers under realistic environmental conditions. *Materials & Design*, 58:242–246, 2014.

S. Uchihashi, J. Foote, A. Girgensohn, and J. Boreczky. Video manga: Generating semantically meaningful video summaries. In *Proceedings of the Seventh ACM International Conference on Multimedia* (MULTIMEDIA '99), pages 383–392, 1999.

U. Umapathi, H.-T. Chen, S. Mueller, L. Wall, A. Seufert, and P. Baudisch. Laserstacker: Fabricating 3D objects by laser cutting and welding. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology* (UIST '15), pages 575–582, 2015.

N. Umetani and R. Schmidt. Cross-sectional structural analysis for 3D printing optimization. In *SIGGRAPH Asia 2013 Technical Briefs* (SA '13) Article 5, 2013.

N. Umetani, J. Mitani, T. Igarashi, and K. Takayama. Designing custommade metallophone with concurrent eigenanalysis. In *New Interfaces for Musical Expression* (NIME'10), 2010.

N. Umetani, Y. Koyama, R. Schmidt, and T. Igarashi. Pteromys: Interactive design and optimization of free-formed free-flight model airplanes. *ACM Transaction on Graphics* Article 65, 33(4), 2014.

N. Umetani, A. Panotopoulou, R. Schmidt, and E. Whiting. Printone: Interactive resonance simulation for free-form print-wind instrument design. *ACM Transactions on Graphics* Article 184, 35(6), 2016.

Unity Games Engine. `https://unity3d.com/`.

J. Vanek, J. A. G. Galicia, and B. Benes. Clever support: Efficient support structure generation for digital fabrication. *Computer Graphics Forum*, 33 (5):117–125, 2014.

T. Vasilevitsky and A. Zoran. Steel-sense: Integrating machine elements with sensors by additive manufacturing. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (CHI '16), pages 5731–5742, 2016.

M. Vázquez, E. Brockmeyer, R. Desai, C. Harrison, and S. E. Hudson. 3D printing pneumatic device controls with variable activation force capabilities. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI '15), pages 1295–1304, 2015.

K. Vidimče, S.-P. Wang, J. Ragan-Kelley, and W. Matusik. OpenFab: A programmable pipeline for multi-material fabrication. *ACM Transaction on Graphics* Article 136, 32(4), 2013.

K. Vidimče, A. Kaspar, Y. Wang, and W. Matusik. Foundry: Hierarchical material design for multi-material fabrication. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology* (UIST '16), pages 563–574, 2016.

V. Vinge. The upcoming singularity or how to survive in a post human era. In *VISION-21 Symposium*, 1993. `http://www.frc.ri.cmu.edu/\simhpm/book98/com.ch1/vinge.singularity.html`.

Voxel8. `http://www.voxel8.com`.

L. Wang and E. Whiting. Buoyancy optimization for computational fabrication. *Computer Graphics Forum*, 35(2), 2016.

W. Wang, T. Y. Wang, Z. Yang, L. Liu, X. Tong, W. Tong, J. Deng, F. Chen, and X. Lia. Cost-effective printing of 3D objects with skin-frame structures. *ACM Transaction on Graphics* Article 177, 32(6), 2013.

W. Wang, H. Chao, J. Tong, Z. Yang, X. Tong, H. Li, X. Liu, and L. Liu. Saliency-preserving slicing optimization for effective 3D printing. *Computer Graphics Forum*, 34(6):148–160, 2015.

Water soluble support. `http://www.3dsystems.com/shop/cartridges/infinity`.

Wazer. Low-cost waterjet cutter. `https://www.kickstarter.com/projects/1294137530/the-first-desktop-waterjet-cutter`, 2016.

C. Weichel, M. Lau, and H. Gellersen. Enclosed: A component-centric interface for designing prototype enclosures. In *Proceedings of the 7th International Conference on Tangible, Embedded and Embodied Interaction* (TEI '13), pages 215–218, 2013.

C. Weichel, J. Hardy, J. Alexander, and H. Gellersen. ReForm: Integrating physical and digital design through bidirectional fabrication. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology* (UIST '15), pages 93–102, 2015.

M. Weichel, M. Lau, D. Kim, N. Villar, and H. W. Gellersen. MixFab: a mixed-reality environment for personal fabrication. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '14), pages 3855–3864, 2014.

M. Weinberg. *It Will Be Awesome If They Don't Screw It Up: 3D Printing, Intellectual Property, and the Fight Over the Next Great Disruptive Technology.* White Paper, Public Knowledge, 2010.

M. Weinberg. *What's the Deal with Copyright and 3D Printing?* White Paper, Public Knowledge, 2013.

M. Weiser. The computer for the 21st century. In *Mobile Computing and Communications Review — Special issue dedicated to Mark Weiser* (SIGMOBILE) 3, volume 3, pages 3–11, 1999.

M. Wessely, T. Tsandilas, and W. E. Mackay. Stretchis: Fabricating highly stretchable user interfaces. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology* (UIST '16), pages 697–704, 2016.

T. Weyrich, P. Peers, W. Matusik, and S. Rusinkiewicz. Fabricating microgeometry for custom surface reflectance. *ACM Transaction on Graphics* Article 32, 28(3), 2009.

A. Wibowo, D. Sakamoto, J. Mitani, and T. Igarashi. DressUp: A 3D interface for clothing design with a physical mannequin. In *Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction* (TEI '12), pages 99–102, 2012.

R. Widmer, H. Oswald-Krapf, D. Sinha-Khetriwal, M. Schnellmann, and H. Böni. Global perspectives on e-waste. *Environmental Impact Assessment Review*, 25(5):436–458, 2005.

D. Wigdor and D. Wixon. *Brave NUI World: Designing Natural User Interfaces for Touch and Gesture.* Morgan Kaufmann Publishers Inc., first edition, 2011.

K. D. D. Willis and A. D. Wilson. InfraStructs: Fabricating information inside physical objects for imaging in the terahertz region. *ACM Transaction on Graphics* Article 138, 32(4), 2013.

K. D. D. Willis, J. Lin, J. Mitani, and T. Igarashi. Spatial sketch: Bridging between movement & fabrication. In *Proceedings of the fourth International Conference on Tangible, Embedded, and Embodied Interaction* (TEI '10), pages 5–12, 2010.

K. D. D. Willis, C. Xu, J. K. Wu, G. Levin, and M. D. Gross. Interactive fabrication: New interfaces for digital fabrication. In *Proceedings of the Fifth International Conference on Tangible, Embedded, and Embodied Interaction* (TEI '11), pages 69–72, 2011.

K. D. D. Willis, E. Brockmeyer, S. E. Hudson, and I. Poupyrev. Printed optics: 3D printing of embedded optical elements for interactive devices. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology* (UIST '12), pages 589–598, 2012.

R. Wu, H. Peng, F. Guimbretière, and S. Marschner. Printing arbitrary meshes with a 5DOF wireframe printer. *ACM Transaction on Graphics* Article 101, 35(4), 2016.

M. M. Yamashita, J. Yamaoka, and Y. Kakehi. Enchanted scissors: A scissor interface for support in cutting and interactive fabrication. In *ACM SIGGRAPH 2013 Posters* (SIGGRAPH '13), Article 33, 2013.

Y. L. Yang, J. Wang, and N. J. Mitra. Reforming shapes for material-aware fabrication. In *Proceedings of the Eurographics Symposium on Geometry Processing* (SGP '15), pages 53–64, 2015.

L. Yao, R. Niiyama, J. Ou, S. Follmer, C. Della Silva, and H. Ishii. PneUI: Pneumatically actuated soft composite materials for shape changing interfaces. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology* (UIST '13), pages 13–22, 2013.

J. Yu, E. Williams, M. Ju, and Y. Yang. Forecasting global generation of obsolete personal computers. *Environmental Science & Technology*, 44(9): 3232–3237, 2010.

J. Zehnder, S. Coros, and B. Thomaszewski. Designing structurally-sound ornamental curve networks. *ACM Transactions on Graphics*, 35(4):99, 2016.

ZEUS. The world's first all-in-one 3D printer/copy machine. `http://www.zeus.aiorobotics.com/`.

X. Zhang, X. Le, A. Panotopoulou, E. Whiting, and C. C. L. Wang. Perceptual models of preference in 3D printing direction. *ACM Transaction on Graphics* Article 215, 34(6), 2015.

H. Zhao, F. Gu, Q.-X. Huang, J. Garcia, Y. Chen, C. Tu, B. Benes, H. Zhang, D. Cohen-Or, and B. Chen. Connected fermat spirals for layered fabrication. *ACM Transactions on Graphics*, 35(4):100, 2016.

Q. Zhou, J. Panetta, and D. Zorin. Worst-case structural analysis. *ACM Transaction on Graphics* Article 137, 32(4), 2013.

L. Zhu, W. Xu, J. Snyder, Y. Liu, G. Wang, and B. Guo. Motion-guided mechanical toy modeling. *ACM Transaction on Graphics* Article 127, 31 (6), 2012.

A. Zoran and J. A. Paradiso. FreeD: A freehand digital sculpting tool. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '13), pages 2613–2616, 2013.

A. Zoran, R. Shilkrot, and J. A. Paradiso. Human-computer interaction for hybrid carving. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology* (UIST '13), pages 433–440, 2013.

A. Zoran, R. Shilkrot, P. Goyal, P. Maes, and J. A. Paradiso. The wise chisel: The rise of the smart handheld tool. *Pervasive Computing, IEEE*, 13(3): 48–57, 2014a.

A. Zoran, R. Shilkrot, S. Nanayakkara, and J. A. Paradiso. The hybrid artisans: A case study in smart tools. *ACM Transactions on Computer-Human Interaction* (ToCHI) Article 15, 21(3), 2014b.

J. M. Zweig and J. H. Woodyatt. Decoupling rights in a digital content unit from download. In *Patent US8934624 B2,* filed December 27, 2011, issued Jan 13, 2015.