

Catalyzing the Agility, Accessibility, and Predictability of the Manufacturing-Entrepreneurship Ecosystem through Design Environments and Markets for Virtual Things

Alexander Brodsky^a, Yotam Gingold^b, Thomas D. LaToza^c, Lap-Fai Yu^d and Xu Han^e
Department of Computer Science, George Mason University, 4400 University Drive, Fairfax, U.S.A.

Keywords: Product and Service Networks, Manufacturing, Design Environments, Markets, Decision Optimization.

Abstract: Proposed is a fundamentally new approach to manufacturing as a service based on a market of virtual things: parameterized products and services that can be searched, composed and optimized, while hiding the underlying complexity of product designs and manufacturing service networks. The approach includes (1) a mathematical framework, composition and decision guidance for virtual things; (2) bootstrapping the market with novel computational techniques and tools to reuse the distributed wealth of existing product and process designs by generalizing them into models of virtual things; and, (3) intelligent computational design tools for entrepreneurs. The goal is to catalyze the agility, accessibility and predictability of the manufacturing-entrepreneurship ecosystem, transforming the Future of Manufacturing.

1 INTRODUCTION

There is a critical disconnect between entrepreneurs who envision new products and manufacturers who might build them. To bridge the disconnect, in this position paper we propose a fundamentally new approach to manufacturing as a service based on a market of virtual things: parameterized products and services that can be searched, composed and optimized, while hiding the underlying complexity of product designs and manufacturing service networks. Our approach bootstraps the market with novel computational techniques and tools to reuse the distributed wealth of existing product and process designs by generalizing them into models of virtual things. This will catalyze the agility, accessibility and predictability of the manufacturing-entrepreneurship ecosystem, transforming the Future of Manufacturing.

Entrepreneurs use their domain knowledge and market insights to conceptualize innovative products, but may fail to realize their ideas due to insufficient

design and manufacturing knowledge. They lack agility (getting a product to market fast), access (to manufacturing and supply chain resources), and predictability. Manufacturers' specialized knowledge in their vertical domains amounts to a distributed volume of existing expert-crafted product and process designs, which assure predictable outcomes. However, they lack agility and access to markets and revenue opportunities provided by entrepreneurial ideas outside of existing rigid supply-chain pyramids. As a result, both entrepreneurs and manufacturers, especially small and medium enterprises (SMEs), miss opportunities to create value.

There has been significant research in manufacturing product and process design (Gingold, Igarashi, and Zorin, 2009; Yu, Yeung, Tang, Terzopoulos, Chan, and Osher, 2011; LaToza, Shabani, and Van Der Hoek, 2013; Shin, Kim, Shao, Brodsky, and Lechevalier, 2017), analysis and optimization (Egge, Brodsky, and Griva, 2013; Shao, Brodsky, and Miller, 2018). Recently, a number of startups have taken important complementary steps to bridge this gap. Companies such as Xometry offer

^a <https://orcid.org/0000-0002-0312-2105>

^b <https://orcid.org/0000-0002-5381-2104>

^c <https://orcid.org/0000-0002-9564-3337>

^d <https://orcid.org/0000-0002-2656-5654>

^e <https://orcid.org/0000-0002-3347-3627>

easy access to manufacturing as a virtual service, where entrepreneurs may enter a CAD file and receive a price and commitment in real time. Behind the scenes, this is enabled through an accurate predictive pricing model and a network of manufacturers with various capabilities, such as CNC machining, injection molding, and 3D printing. However, combining these unit processes into a composite manufacturing process to come up with a finished consumer product is out of their scope. Companies like Kerfed improve agility in manufacturing response to customer demand by accepting a CAD design of an assembly out of standard components, and performing analysis to discover the characteristics of its components and their interconnection so that they can be (semi-) automatically sourced from suppliers, and so that the assembled product could be priced for a customer order. Companies like Physna boost manufacturers' agility in responding to customer demand by searching for similar CAD designs in a large design database using not only meta-data of existing designs, but also their geometric and functional properties, significantly simplifying the creation of a new CAD design via re-use. CAD/CAM software, like OnShape, has been widely used for product and process design in increasingly more diverse vertical domains, enabling designers to specify the blueprints of their idea with high precision.

However, major challenges remain. First, entrepreneurs do not typically have CAD modelling skills; even when starting from a similar design, they may not understand the design complexity and intent, and still need to rely on professional CAD designers. Second, when using a fixed CAD product design for sourcing manufacturers, the design is typically not optimized to consider manufacturability and supply chain and manufacturing costs. Yet it is often possible, via small modifications to a product's CAD design, to make manufacturing significantly simpler and less expensive with little or no effect on desirable customer-facing product characteristics. Third, and perhaps most important, US manufacturers, especially SMEs, are still primarily selling low-margin manufacturing capacity, because they face stiff competition due to lower labor costs and increasing quality of foreign, especially East and South-East Asian, manufacturing. Manufacturing SMEs typically do not offer new innovative products with high profit margins because they lack access to these innovative product ideas and the agility to respond to the market opportunities they present.

This paper is organized as follows. In Section 2 we overview our approach; in Section 3 we illustrate

the approach by a real-world example. We discuss a range of research questions to be addressed to realize the new approach in Section 4, and conclude in Section 5.

2 OUR APPROACH

We propose a fundamentally new approach to, and a novel productivity framework for, the manufacturing-entrepreneurship ecosystem based on bootstrapped markets of virtual products and services (see Figure 1, middle funnel layer), which we collectively call V-things. A virtual product is represented by a parameterized CAD design, e.g., to characterize a customizable consumer product, part or raw material. A virtual service represents a parameterized transformation of virtual products into other virtual products, e.g., to characterize a customizable manufacturing process, supply, transportation, logistics or a composed service network. Each V-thing—product or service—is associated with an analytic model that describes the product and/or service's feasibility and customer-facing characteristics (e.g., weight, durability, strength, volume for a product; and cost, delivery time and default risk for a service) as a function of the product and/or service's decision and fixed parameters (e.g., dimensions, position of fixtures, type and properties of materials for a product; and settings for manufacturing processes, selection of and ordered quantities from suppliers and manufacturers).

The purpose of the Decision Guidance System over a repository of V-things (Figure 1) is to enable manufacturers and entrepreneurs to (1) search for relevant V-things (products and services) in the market, (2) compose them into more complex V-things (e.g., assembled products or service networks) and, most importantly, (3) guide decisions, activity that involves model training, predictions, optimization and trade-off analysis, i.e., recommending users Pareto-optimal choices on V-thing parameter instantiation (corresponding to specific products and services), while eliciting and acting on preferences among possibly competing objectives, such as cost, reliability and time to market.

To manufacturers, V-thing markets offer an order-of-magnitude more agility in response to customer demand and access to entrepreneurs with ideas. More speculatively, V-thing markets may allow manufacturers to expand their business model, from selling low-margin manufacturing capacity to agile supply of high-margin on-demand products in their vertical markets, boosting their global

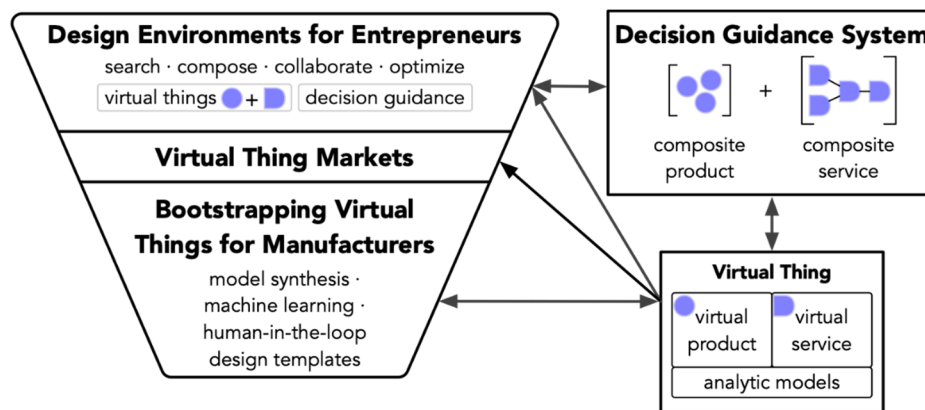


Figure 1: The funnel of the manufacturing/entrepreneurship ecosystem.

competitiveness. To scale up the creation of V-things—products and services—beyond the traditional limits of generative design, we envision research and development of novel computational techniques and V-thing design tools for manufacturers (see bottom part of the funnel in Figure 1). These techniques and tools will support search, reuse, and generalization of manufacturers’ existing product and process designs into models of V-things, leveraging their domain expertise to manufacture similar things. We envision an extensive use of repositories of examples created in a CAD system as well as physically scanned examples. The creation of new V-things will also need to leverage available V-things in the market. We envision associating virtual things with multi-aspect descriptions to aid their discovery by entrepreneurs.

To entrepreneurs, V-thing markets offer the agility to realize their ideas for a new product or service through flexible search, composition, optimization and Pareto trade-off analysis using available V-things, while hiding the complexity of underlying product designs and manufacturing service networks—both process and supply chain. We expect to design and develop Intelligent Design Tools for Entrepreneurs (see the top layer of the funnel in Figure 1), as possible extensions to existing CAD/CAM tools, using paradigms such as design-by-sketch and by example, and leveraging V-thing markets. This agility will drive manufacturing demand.

3 MOTIVATING EXAMPLE

Dentists re-opening their practices after closures due to COVID-19 need to overcome a major exposure risk. Many dental procedures—those that require the

use of a high-speed handpiece or an ultrasonic scaler—generate a pressurized spread of aerosol, which may carry microorganisms, including the novel coronavirus. The main mitigating solution offered by dental suppliers is an extra-oral suction, based on repurposed dust vacuums. This is too bulky, noisy, and expensive for a dental operator. A dentist entrepreneur comes up with a much smarter idea: she wants to repurpose an existing HVE (high-volume evacuation) line already available in the dental unit and normally used for dental suction—but not for the collection of aerosol in the air. What is missing is a specially designed funnel (Figures 2 & 3) that can be attached to an existing HVE line and be held in close proximity to the patient's mouth during a dental procedure. This funnel must satisfy a number of properties: (1) it must be of geometry and size that maximize the suction of aerosol (too small will not be effective for aerosol cloud; too big will not generate sufficient suction pressure); (2) it must be light, yet strong and autoclavable, i.e., withstand sterilization temperatures of 175°C; and (3) it must be attachable to both a cheek retractor and an external adjustable arm. In addition, the adjustable arm (Figure 2) must be designed to hold the funnel attached to the HVE line in the required position to enable hands-free operation, as well as an optional transparent shield. The entrepreneur dentist envisions that, if introduced to the market quickly, this new aerosol collection funnel can easily be sold for \$70-80 per part, which is a small fraction compared to \$2000-3000 per one bulky and noisy extra-oral suction device currently on the market. She and her dentist colleagues would certainly find this offering extremely useful and relatively inexpensive.

This motivating example was found in the wild, suggested by a dentist in a Facebook group for dentists. However, that dentist’s idea would never get anywhere beyond a Facebook post. Xuction Dental—

a start-up company in Virginia—invested significant engineering, material science and manufacturing expertise to implement it.

As lead users, dentists are very familiar with their needs and current technologies, but are generally not technical in the sense of manufacturers or engineers (Hippel, 1988). The road to idea realization has high barriers to entry. Expertise is distributed and siloed: The entrepreneur must find product designers and manufacturers with whom to partner. Communication is difficult and tools are inaccessible: Everyone must communicate their capabilities and their needs to each other, using very different languages and perspectives. Entrepreneurs cannot participate in the digital design of the product. Product designers may not have access to manufacturing decisions. Siloed decision-making results in sub-optimal designs. Work is wasted: The work of designers and manufacturers is delivered bespoke for a particular product. Without opportunities for discovery and re-use of designs and processes, productivity is suppressed and capacity is underused.

In our vision, the dentist entrepreneur uses an accessible Intelligent Design Environment for Entrepreneurs (Figure 4) in collaboration with other innovators (such as product designers) who may suggest improvements. The entrepreneur creates a rough sketch of their product vision (e.g., the aerosol suction funnel) and provides some free text describing it. The design tool constructs a 3D model approximating the funnel sketch and uses it to search for relevant V-products (e.g., for vacuum polymer funnels) and associated V-services (e.g., manufacturer who produces them) in the V-thing market. The dentist explores one V-product that looks relevant, and the design tool displays a 3D-model of the V-product fitted to the dentist's sketch. The 3D depiction is annotated with customer-facing characteristics, which can be used to express known constraints and objectives/criteria to be considered. For example, the dentist may provide funnel product constraints, such as the diameter of connecting hose, the maximum allowed weight, the minimal temperature of 175°C to withstand, and service constraints such as the number of units to be produced and the maximum delivery time window. She also chooses objectives to be considered, such as vacuuming efficiency, weight, cost-per-unit and delivery time. The design tool leverages the V-service and V-product analytic models and uses the Decision Guidance System to recommend and display a few Pareto-optimal alternatives in terms of the specified objectives while soliciting comparison responses. After a number of iterations, the dentist converges to

a specific instance of a vacuum funnel and specific service terms. The dentist initially orders a couple of samples, tries them out, makes adjustments, and then places a production order of 10,000 units to the V-service provider to be sold to dental practices.



Figure 2: Dental aerosol funnel connected to HVE suction line. © Xuction Dental.



Figure 3: Dental aerosol collection funnel. © Xuction Dental.




Figure 4: The Entrepreneur Design Environment.

The creator of the vacuum polymer funnel V-product and associated manufacturing V-service may be a small injection molding manufacturer, who happen to produce similar products, and who decided to extend its business model from selling manufacturing (injection molding) capacity to wholesale of some V-products, such as on-demand vacuum polymer funnels. To do that, the manufacturer uses in-house and/or hired expertise to specify V-product and V-service designs, leveraging many specific expert-crafted CAD/CAM product and process designs of similar things produced in the past. Design Tool for V-things helps manufacturers to search for relevant specific designs, and generalize

them with analytic model, that expresses feasibility and customer facing characteristics such as vacuuming efficiency, weight, cost-per-unit and delivery time as a function of internal product and process parameters: geometry, dimensions, type and density of polymer material, as well as process settings. While this task requires expertise, even with the help of the V-thing Design Tool, the outcome is highly reusable and allows the manufacturer significant agility and access to otherwise unavailable markets, such as for the dental aerosol collection funnel, which can be sold at much higher profit margins. In turn, the manufacturer may use some other existing V-things in the market, e.g., polymer material V-product and associated V-service. The manufacturer of polymer material, in turn, leverages its expertise in designing and producing special polymers with unique properties, such as low density and the ability to withstand high temperature. Of course, behind V-things in the market may also be engineering and technology firms that want to expand their business model from selling consulting to becoming virtual manufacturers, while generating demand for manufacturing capacity in the external service network.

4 TECHNICAL PROBLEMS

To realize this new paradigm, we need to overcome a number of mathematically and computationally challenging research problems.

4.1 V-things Math Framework, Composition, Search and Decision Guidance

The framework will include mathematical formalization of V-things—products and services—including their design specs, customer facing specs, customer requirements specs, and the notions of feasible and optimal parameter instantiation based on analytic models associated with V-things. To support the creation of V-things by manufacturers one needs to design recursive compositional models—e.g., for product assembly and service networks—in such a way that compositions would be easy (e.g., graphically) to specify by (non-mathematical) domain users, yet can be interpreted as formal analytic models by the system.

We envision a virtual product to be represented by a parameterized CAD design, e.g., to characterize a customizable consumer product, part or raw material.

A virtual service represents a parameterized transformation of virtual products into other virtual products, e.g., to characterize a customizable manufacturing process, supply, transportation, logistics or a composed service network. Each V-thing—product or service—is associated with an analytic model that describes the product and/or service’s feasibility and customer-facing metrics/characteristics as a function of the product and/or service’s (fixed and decision) parameters. For V-products, examples of customer-facing metrics include external dimensions, weight, durability and vacuum efficiency; while examples of internal parameters include internal dimensions, position of fixtures, and type and properties of materials. For V-services, examples of customer-facing metrics include cost-per-unit, total ordered quantities per item, delivery time, carbon emissions per unit, and default risk; while examples of internal parameters include settings for unit manufacturing processes (e.g., CNC machining, injection molding or 3D printing) and selection of and ordered quantities from suppliers and manufacturers. Intuitively, V-things’ customer-facing metrics are all that customers care about when selecting products and services; whereas, customers do not care about, or even understand, V-thing parameters outside the set of customer-facing metrics.

Consider an example of a manufacturing service network (Figure 5) for a heat sink product (Brodsky, A., Krishnamoorthy, M., Nachawati, M. O., Bernstein, W. Z., and Menascé, 2017; Brodsky, Nachawati, Krishnamoorthy, Bernstein, and Menascé, 2019), produced by Birmingham Aluminum Ltd. This product is an assembly of aluminum and the covering plastic frame using accessories. Both the product and the service are composite. The service to produce the finished heat sink product (HS) involves a hierarchical service network, which includes supply, manufacturing and demand services; in turn, manufacturing is also a service network, composed of aluminum plate contract manufacturer, smelting, HS base production line, HS base contract manufacturer and HS production line. In turn, HS production line is a service network composed of HS shearing, anodizing, CNC machining, quality inspection, and final assembly, etc. The challenge here is to avoid hard-wired and time-consuming development of analytic models for every composition of V-products (like the assembled heat sink) and V-services (like the heat sink service network). To address this challenge, one needs to design (re-usable) recursive compositional models—across both product assembly and service

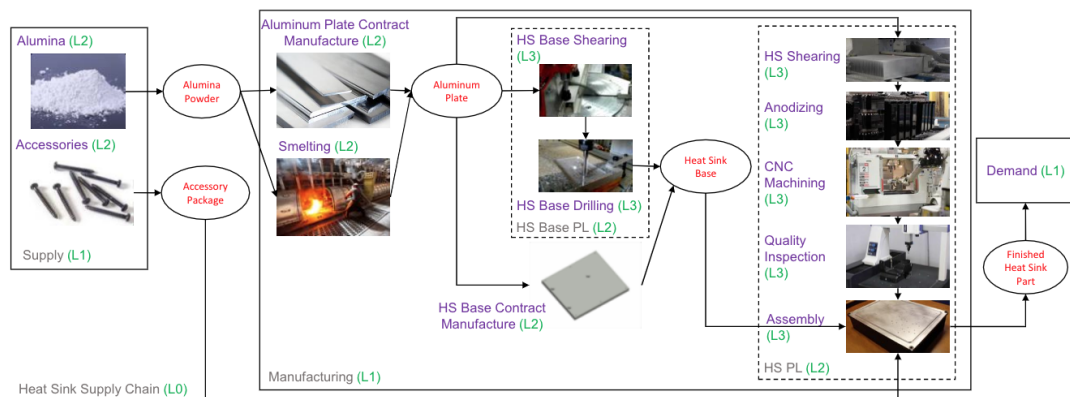


Figure 5: Example of manufacturing service network for heat sink product.

network compositional hierarchies—in such a way that compositions would be easy to specify (e.g., graphically) by (non-technical) domain users, yet can be interpreted as formal analytic models by the system. To achieve this goal, we envision leveraging techniques from the Factory Optima system, which was designed and developed for NIST (Brodsky, Krishnamoorthy, Bernstein, and Nachawati, 2016; Brodsky et al, 2017; Brodsky et al, 2019; Brodsky, Shao, Krishnamoorthy, Narayanan, Menascé, and Ak, 2016), but which has not considered parameterized or composed product designs.

While searching for V-things in the market repository is conceptually similar to searching for regular products and services, it is fundamentally different and more challenging computationally. Just a match between a user requirement spec and a particular V-thing customer facing spec is a constraint satisfaction problem, which, like the corresponding optimization problem, may be both non-linear and combinatorial in high-dimensional space.

To scale up online optimization for practical size problems within manageable computational time, one idea is to design pre-processing algorithms that generate differentiable surrogates for (combinatorial components of) analytic models used in optimization problems. To scale-up search for V-things, we will need to design offline pre-processing algorithms to generate bounding polyhedral set approximations that are amenable to efficient (multi-dimensional) indexing techniques for search. Another major challenge we need to overcome has to do with the fact that composable and modular analytic models—against which optimization is applied—are expressed using object-oriented code (e.g., in Python); yet the best mathematical programming algorithms require, as input, a closed-form-arithmetic (“white-box”) optimization model (as opposed to simulation-like “black-box” model). This can be done by leveraging

and further developing symbolic computation techniques to machine generate closed-form-arithmetic optimization models from software code in order to use the best existing, as well as develop extensions to, mathematical programming algorithms (Brodsky and Wang, 2008; Brodsky and Luo, 2015).

4.2 Design Tools for Virtual Things for Manufacturers

The goal is to design computational techniques to generalize manufacturers’ existing designs (products and services) as V-things. Bootstrapping the v-things repository involves identifying its decision parameters and analytic models that express feasibility and customer-facing characteristics as a function of these parameters. The challenge is that black-box data-driven approaches may fail to find straightforward and reliable shape designs or governing equations. To solve this problem, we envision the need to leverage and extend the techniques of program synthesis (Solar-Lezama, 2008) to enable the creation of analytic models by non-programmers, resulting in “grey-box” models that are partly physics-based and partly data-driven. Since there are many model and non-decision parameter alternatives, we envision the use of machine learning algorithms to train, validate, and select the best model alternatives.

We propose example-based techniques for generating parametric CAD models. Rather than requiring manufacturers to re-train with a new design tool, we envision the need to analyze a set of existing shapes and semi-automatically find parameters to define a family of shapes comprising a V-product. For example, an engineer with CAD experience could create multiple instances of a design with their favorite CAD tool. Alternatively, a machine operator can create multiple variations of a physical object. We

envision the need to analyze and filter these shapes and propose variables. For example, parameters could be continuous like repeated lengths, which may appear in whole multiples, or discrete like symmetry relationships or choice of materials.

We envision developing approaches for enabling users to author models expressing performance characteristics of v-things, such as strength, stability, manufacturing expense and feasibility, and material waste. The metrics associated with each v-thing can be used for e.g., Pareto front discovery and optimization. The metrics can measure mass, strain under load, manufacturing material waste, torque, etc. Users can design finite element simulations involving the part. Importantly, it is desirable for metrics to be differentiable when possible, allowing their use in gradient-based optimization applications (e.g., decision guidance, Pareto front discovery, and deep learning).

Rather than requiring knowledge of programming, which manufacturers may not possess, we envision creating novel end user programming techniques which enable performance models to be created through examples. To solve this problem, one idea is to use “grey-box” models that are partly analytical or physics-based and partly data-driven. To do that one can leverage and extend the techniques of program synthesis (Solar-Lezama, 2008) to enable the creation of analytic or physics-based models with meaningful parameters by non-programmers. The resulting programs will have an overly large set of parameters. We propose to put the user in charge of suggesting and filtering possible parameters to be user-facing. To do this, one can explore Programming by Demonstration approaches (Cypher and Halbert, 1993; Lieberman, 2001). Users can describe examples of performance for specific inputs or mark measurements (solo or repeated) and components with a symmetry relationship. Since there are many model and non-decision parameter alternatives, we will use machine learning regression and classification algorithms to train, validate, and select the best model alternatives. It is also important to explore ways in which shapes and performance characteristics can be visualized, helping communicate the effects of choices on the model.

Manufacturer’s knowledge and experience in manufacturing also make them uniquely well-suited to create v-services for a v-product’s creation. The v-service for manufacturing a v-product entails sourcing raw materials and arranging the manufacturing process. We propose to leverage our prior work using flow diagrams to specify v-services. This will be integrated into the v-thing designer,

allowing manufacturers to design a shape’s parameters simultaneously with its manufacturing process. To enable manufacturers, who are not expected to be programmers, to design the analytical models for the v-services, we envision the use and exploration of Programming by Demonstration approaches (Cypher and Halbert, 1993) (see End-User Authoring of Performance Analytic Models).

4.3 Intelligent Computational Design Tools for Entrepreneurs

Intelligent computational design tools for entrepreneurs and their collaborators must enable them to turn ideas into virtual things and then into prototypes without having expertise in CAD or engineering. We envision intuitive search approaches based on sketching, similar-product search, and assembly-based modelling to enable entrepreneurs to find and compose virtual things within the marketplace intuitively. Such approaches will also encourage the reuse and adaptation of existing virtual things to unleash their potential. The computational design tools driven by decision guidance will also perform optimization and Pareto trade-off analysis to automatically suggest design alternatives. Entrepreneurs will be able to select between alternatives, providing preferences which the system uses to iteratively elicit the utility function and use it to generate new alternatives, and collaborate with the tools in the ideation process.

Sketching is a natural, straightforward way of expression for illustrating creative ideas. Compared to using traditional, sophisticated CAD software (e.g., 3ds Max) for creating 3D model designs, which requires a steep learning curve to master, it is much easier for people to sketch their ideas on a tablet. A sketch-based design interface allows people to focus on envisioning the design of their products rather than operating the sophisticated interface of CAD software. There are many challenges in creating a convenient and effective sketch-based design interface. One challenge is due to the irregularity of sketches: most people are not artists and they can only sketch their ideas roughly. One approach to tackle this problem is to devise machine learning approaches for inferring a clean and valid design from a user’s sketches. Recent work using generative adversarial networks (GAN) for inferring 3D models from sketches provides a promising solution (Guérin, Digne, Galin, Peytavie, Wolf, Benes, and Martinez, 2017; Portenier, Hu, Szabó, Bigdeli, Favaro, and Zwicker, 2018). Sketch-based interfaces have also been proposed for creating furniture designs (Xie, Xu,

Mitra, Cohen-Or, Gong, Su, and Chen, 2013) (Figure 6).




Figure 6: A sketch-based interface for furniture design.

For most people, it is much simpler to design a virtual service or product guided by suggestions. For example, when renovating homes, people often refer to a magazine showing many examples of home renovation projects to get inspiration, rather than designing from scratch. Akin to this, we envision suggestive design interfaces to help entrepreneurs with design. For instance, consider the design of a chair. A suggestive user interface may work like this:

- 1) The user first specifies the high-level goals of the chair design, such as the style (e.g., classic or modern?), context of use (e.g., a dining chair or a desk chair?), physical properties (e.g., dimensions, weights), and functionalities (e.g., adjustable?). The user may also sketch his or her rough idea or provide an existing similar design.

- 2) According to the user's specification from step (1), the system samples a number of feasible design solutions that match with the user's preference;

- 3) The user chooses one of the suggestions;

- 4) The user may modify the suggested design to better match with the user envisions;

- 5) The system generates new suggestions based on the specified modifications.

- 6) Repeat steps (3) to (5) until the user obtains a desired final design.

Here, the research challenge lies in inferring what the user wants from the high-level description or rough sketch in step 1). A promising strategy to overcome such a challenge involves applying a data-driven approach to learn statistical patterns of design from a large database of existing designs. For example, given a database of 3D chair designs, one can train machine learning classifiers to determine perceptual shape style similarity (Lun, Kalogerakis, and Sheffer, 2015). Given a rough sketch or a partially finished chair design, a suggestive interface can infer and recommend possible full designs according to styles and assembly schemes learned from existing chair designs (Xie et al, 2013).

Another promising strategy to help entrepreneurs create designs is assembly-based 3D modeling. The

idea is to provide users with simple primitive shapes that they can assemble into a complex object.

Another promising strategy to help entrepreneurs create designs is assembly-based 3D modeling. The idea is to provide users with simple primitive shapes that they can assemble into a complex object. Akin to the furniture design of IKEA, the algorithm automatically decomposes a furniture product (e.g., a chair) into a number of manufacturable, modular components which customers can easily assemble into the full products. Compared to the traditional approach of creating 3D objects from scratch using low-level mesh or primitive manipulation tools in CAD software, assembly-based modeling is much simpler to learn and perform.

A major research challenge of realizing assembly-based modeling lies in designing a set of compatible primitive shapes that the user can conveniently assemble into many objects. A trivial solution is to design a set of very general-purpose primitive shapes, like LEGO bricks of different dimensions, which give a high-degree of freedom and hence high flexibility with respect to the objects they can assemble. However, it typically takes many very general-purpose primitive shapes to assemble a desired object, and hence the physical assembly process could be time-consuming and complex.

To tackle such challenges, we will employ a recently devised approach called "hands-on assembly-based modeling" (Duncan, Yu, and Yeung, 2016). The key idea is to create an algorithm to automatically extract and generate a set of compatible, interchangeable, and semantically meaningful primitive shapes given a set of existing objects. Such primitive shapes can be used for assembling many variations of the original objects. Given a small set of chair 3D models, which can be easily found on the Internet, the algorithm automatically decomposes the chairs into a set of compatible, interchangeable, and 3D-printable primitive components—such as legs, bases, and backs—that a lay user can easily assemble into different new chairs.

It is important to impose physical and functional constraints on the generated primitive shapes, as well as the final product assembled using these primitive shapes. Such constraints have practical implications. For example, realized as virtual products that are traded on our platform, the primitive shapes should be compact and regular to facilitate manufacturing, packaging, and transportation; while the final object assembled should possess desirable physical properties (e.g., the assembled chair must be sturdy).

5 CONCLUSIONS

We envision new design environments and markets for virtual things as the bridge over the gap between unmatched entrepreneurial initiatives and manufacturing capabilities of the value creation ecosystem today. The goal is to catalyze the agility, accessibility, and predictability of the ecosystem. As discussed in Section 3, significant research problems need to be overcome, including (1) V-things Math Framework, Composition, Search and Decision Guidance; (2) Design Tools for V-things for Manufacturers; (3) Intelligent Computational Design Tools for Entrepreneurs.

ACKNOWLEDGEMENTS

Lap-Fai Yu is supported by an NSF CAREER grant (award number: 1942531) in this work.

Yotam Gingold was supported by the United States National Science Foundation (IIS-1453018) and a gift from Adobe Systems Inc.

REFERENCES

- Gingold, Y., Igarashi, T., and Zorin, D., 2009. Structured annotations for 2D-to-3D modeling. In *ACM SIGGRAPH Asia '09*.
- Yu, L-F., Yeung, S. K., Tang, C-K., Terzopoulos, D., Chan, T. F., and Osher, S., 2011. Make it home: automatic optimization of furniture arrangement. In *ACM transactions on graphics* 30, 4: 86.
- LaToza, T. D., Shabani, E., and Van Der Hoek, A., 2013. A study of architectural decision practices. In *6th International Workshop on Cooperative and Human Aspects of Software Engineering '13 (CHASE)*, 77–80.
- Shin, S., Kim, D. B., Shao, G., Brodsky, A., and Lechevalier, D., 2017. Developing a decision support system for improving sustainability performance of manufacturing processes. In *Journal of Intelligent Manufacturing* 28, 1421–1440.
- Shao, G., Brodsky, A., and Miller, R., 2018. Modeling and Optimization of Manufacturing Process Performance using Modelica Graphical Representation and Process Analytics Formalism. In *Journal of intelligent manufacturing* 29, 6: 1287–1301.
- EGGE, N., Brodsky, A., and Griva, I., 2013. An Efficient Preprocessing Algorithm to Speed-Up Multistage Production Decision Optimization Problems. In *46th Hawaii International Conference on System Sciences '13*.
- Hippel, E., 1988. *The Sources of Innovation*, Oxford University Press. USA.
- Brodsky, A., Krishnamoorthy, M., Nachawati, M. O., Bernstein, W. Z., and Menascé, D. A., 2017. Manufacturing and contract service networks: Composition, optimization and tradeoff analysis based on a reusable repository of performance models. In *2017 IEEE International Conference on Big Data*.
- Brodsky, A., Nachawati, M. O., Krishnamoorthy, M., Bernstein, W. Z., and Menascé, D. A., 2019. Factory optima: a web-based system for composition and analysis of manufacturing service networks based on a reusable model repository. In *International Journal of Computer Integrated Manufacturing* 32, 206–224.
- Brodsky, A., Krishnamoorthy, M., Bernstein, W. Z., and Nachawati, M. O., 2016. A system and architecture for reusable abstractions of manufacturing processes. In *2016 IEEE International Conference on Big Data*.
- Brodsky, A., Shao, G., Krishnamoorthy, M., Narayanan, A., Menascé, D. A., and Ak, R., 2016. Analysis and Optimization based on Reusable Knowledge Base of Process Performance Models. In *International Journal of Advanced Manufacturing Technology: 1–21*.
- Solar-Lezama, A., 2008. *Program Synthesis by Sketching*. University of California, Berkeley.
- Cypher, A., Halbert, D. C., 1993. *Watch what I Do: Programming by Demonstration*. MIT Press.
- Lieberman, H., 2001. *Your Wish is My Command: Programming By Example*. Elsevier.
- Guérin, É., Digne, J., Galin, É., Peytavie, A., Wolf, C., Benes, B., and Martinez, B., 2017. Interactive example-based terrain authoring with conditional generative adversarial networks. In *ACM Transactions on Graphics* 36, 1–13.
- Portenier, T., Hu, Q., Szabó, A., Bigdeli, S. A., Favaro, P., and Zwicker, M., 2018. Faceshop: deep sketch-based face image editing. In *ACM Transactions on Graphics* 37, 1–13.
- Xie, X., Xu, K., Mitra, N. J., Cohen-Or, D., Gong, W., Su, Q., and Chen, B., 2013. Sketch-to-design: Context-based part assembly. In *Computer Graphics Forum*, 233–245.
- Lun, Z., Kalogerakis, E., and Sheffer, A., 2015. Elements of style: learning perceptual shape style similarity. In *ACM Transactions on Graphics '15*.
- Duncan, N., Yu, L-F., and Yeung, S-K., 2016. Interchangeable Components for Hands-On Assembly Based Modelling. In *ACM transactions on graphics* 35, 6.
- Brodsky, A., Wang, X. S., 2008. Decision-Guidance Management Systems (DGMS): Seamless Integration of Data Acquisition, Learning, Prediction and Optimization. In *Proceedings of the 41st Annual Hawaii International Conference on System Sciences (HICSS 2008)*, 71.
- Brodsky, A., Luo, J., 2015. Decision Guidance Analytics Language (DGAL): Toward Reusable Knowledge Base Centric Modeling. In *17th International Conference on Enterprise Information Systems (ICEIS 2015)*.