Intelligent Tools for Creative Graphics
(SIGGRAPH 2020 Course)

Ariel Shamir, Niloy J. Mitra, Nobuyuki Umetani, and Yuki Koyama

Course Notes

Abstract

In recent years, much research has been dedicated to the development of "intelligent tools" that can assist both professionals as well as novices in the process of creation. Using the computational power of the machine, and involving advanced techniques, the tools handle complex and tedious tasks that were difficult or even impossible for humans, thereby freeing the human creator of many constraints and allowing her to concentrate on the creative process, while ensuring high-quality and valid design. This course is aimed at presenting some of the key technologies used to assist interactive creative processes. The course allows researchers and practitioners to understand these techniques more deeply, and possibly inspire them to research this subject and create intelligent tools themselves. More specifically, the course will concentrate on four main enabling technologies: geometric reasoning, physical constraints, data-driven techniques and machine learning, and crowdsourcing. In each of these areas the course will survey several recent papers and works and provide examples of using these in the creation of a variety of outputs: 3D models, animations, images, videos and more.

Citation: Ariel Shamir, Niloy J. Mitra, Nobuyuki Umetani, and Yuki Koyama. 2020. Intelligent Tools for Creative Graphics. In *ACM SIGGRAPH 2020 Courses*. Article 15. DOI: `https://doi.org/10.1145/3388769.3407498`

Contents

Course Schedule

This 3-hour course is divided into subparts based on the methods that make the tools intelligent. The full course schedule is in Table 1.

Table 1: Course schedule.

Time	Topic	Instructor
10 min	Introduction	Ariel Shamir
40 min	Geometric Reasoning	Niloy J. Mitra
40 min	Physical Constraints	Nobuyuki Umetani
10 min	(Short Break)	
40 min	Data-Driven Techniques and Machine Learning	Ariel Shamir
25 min	Crowdsourcing	Yuki Koyama
15 min	Summary and Discussion	

About the Lecturers

- Prof. **Ariel Shamir** is the Dean of the Efi Arazi school of Computer Science at the Interdisciplinary Center in Israel. He received his Ph.D. from the Hebrew University in Jerusalem, and spent two years as PostDoc in the computational visualisation centre at UT Austin. Ariel has numerous publications and a number of patents. He was listed on the Thomson Reuters highly cited researchers in 2015. Ariel has a broad commercial experience consulting various companies including Disney research, Mitsubishi Electric, PrimeSense (now Apple), Verisk and more. He specializes in geometric modeling, computer graphics, image processing and machine learning.
- **Niloy J. Mitra** is a Professor of Geometry Processing in the Department of Computer Science, University College London (UCL). His research interests include shape analysis, data-driven geometry processing, and computational design and fabrication. Niloy received the ACM Siggraph Significant New Researcher Award in 2013, the BCS Roger Needham award in 2015, and the Eurographics Outstanding Technical Contributions Award in 2019.
- **Nobuyuki Umetani** is a project lecturer in the creative informatics department at the University of Tokyo, in Japan. His interests are computational fabrication, physics simulation, and interactive techniques.
- **Yuki Koyama** is a Researcher at National Institute of Advanced Industrial Science and Technology (AIST), Japan. He received his Ph.D. from The University of Tokyo in 2017. His research fields are mainly computer graphics and human-computer interaction. In particular, he is interested in enhancing various design activities by employing computational techniques such as mathematical optimization.

1 Introduction

Since its beginning, Computer Graphics has been concerned with the creation of things. Be it 3D virtual models for rendering, digital images—both photorealistic and non-photorealistic, videos, films, animations and recently also physical objects that can be fabricated. The challenges in creating such media or objects are numerous. Not only does the creator need high technical abilities and skills, she also needs ideas and inspiration. Many digital tools and applications were implemented over the years to help the creative process, and interactive techniques were at the core of computer graphics. Still, many of these interactive tools today require a high-level of domain-specific expertise to use.

In recent years, much research has been dedicated to the development of *intelligent tools* that can assist both professionals as well as novices in the process of creation. Using the computational power of the machine, and involving advanced techniques, the tools handle complex and tedious tasks that were difficult or even impossible for humans, thereby freeing the human creator of many constraints and allowing her to concentrate on the creative process, while ensuring high-quality and design validity.

This course is aimed at presenting some of the key technologies used to assist interactive creative processes. The course allows researchers and practitioners to understand these techniques more deeply, and possibly inspire them to research this subject and create intelligent tools themselves.

More specifically, the course will concentrate on four main enabling technologies:

- Geometric reasoning
- Physical constraints
- Data-driven techniques and machine learning
- Crowdsourcing

In each of these areas the course will survey several recent papers and works and provide examples of using these in the creation of a variety of outputs: 3D models, animations, images, videos and more.

Geometric and Physical constraints assist mostly in creating valid 3D objects both for virtual worlds and for actual realization using computational fabrication, however we will also show examples where these techniques can help image editing. Data-driven techniques utilize recent advances in machine learning and deep-learning to assist in the creation, editing and modification of images, videos but also 3D objects. Crowdsourcing allows the exploration and use of high dimensional design parameter spaces, and design alternatives.

Using these techniques, an interactive tool becomes more "intelligent" in a sense that it can make independent decisions and checks while the object or media are edited or created. Such independent behaviour is based on knowledge from any one of the four abovementioned technologies. Creating with such an "intelligent agent" creates a more natural interaction process and opens up opportunities that were previously limited only to experts.

2 Geometric Reasoning

An intelligent agent should be able to generate novel variations that are valid and diverse—geometrically, semantically, and topologically. Examples include generating furniture layout, floor plan generation, garment authoring, 3D models, and more generally virtual 3D environments. In this part of the course, we propose to focus on: generative 3D models with fine-scale details and generative layouts incorporating both topological and geometric variations.

2.1 Learning from physical simulators

Configurations of objects in real world obey laws of physics. Examples include (stable) stacks of objects, dynamics of garments on human avatars, or even interaction of light and shadow with objects. While much progress has been made in developing realistic physical models and real-time systems, matching simulators to real data remains a challenge. In this course, we would discuss scenarios [37; 38; 12] where we can train machines to learn directly from physical simulations and subsequently result in intelligent tools that add realistic details at design-time without incurring unnecessary cost and computational overhead.

2.2 Learning to manipulate images

A common goal in Computer Graphics is to realistically manipulate images. The fundamental challenge in this setting is the lack of full scene description in the form of 3D objects, illumination setting, and material parameters. Geometric reasoning, in the form of primitive fitting or learning directly from training data, has been shown to produce geometric proxies as abstractions of such (unknown) scenes, and subsequently enable realistic and perspective-correct manipulations. We will show examples with cuboids [40] and generalized cylinders [4].

2.3 Learning to generate topological variations

The ability to generate novel, diverse, and realistic 3D shapes along with associated part semantics and structure is central to many applications requiring high-quality 3D assets or large volumes of realistic training data. A key challenge towards this goal is how to accommodate diverse shape, including both continuous deformations of parts as well as structural or discrete alterations which add to, remove from, or modify the shape constituents and compositional structure. Such object structure can typically be organized into a hierarchy of constituent object parts and relationships, represented as a hierarchy of n-ary graphs. We will discuss StructureNet [21], a hierarchical graph network which (i) can directly encode shapes represented as such n-ary graphs; (ii) can be robustly trained on large and complex shape families; and (iii) be used to generate a great diversity of realistic structured shape geometries. We will also discuss extensions [22] to learning domains for manipulation where both geometric and topological changes can be unified.

3 Physical Constraints

Computer controlled fabrication tools such as 3D printers, laser cutters and CNC milling machines have become wildly available to the consumers. However, it is difficult to design original customized shapes while considering its physical functionality. In this talk, we present the studies to incorporate real-time physics simulation into help the design of creative functional shapes. In particular, we focus on the (i) modeling of physics (ii) data-driven approach to facilitate the interactive modeling. Related publications: [34; 32; 33; 31].

4 Data-Driven Techniques and Machine Learning

Recent years have seen a revolution in the amount and ways data is used in all realms of science in general and computer science specifically. Complex problems are solved using machine learning techniques that rely on data. For example, in computer vision neural networks are now used for almost every task successfully. The question is how to make use of such techniques in the context of an intelligent tool for the creation of graphical content. To this end, there are several machine learning techniques that enable the tool to make independent decisions without explicit input from the human user. These include:

- Classifiers or clustering algorithms that assist filtering large amounts of data to bring only relevant data to the user.
- Feature extraction from media (images, 3D objects, videos) that allow faster search and retrieval, and faster comparison between objects.
- Cross-modal analysis that allows connecting between different media types.
- Generative models that enable the creation of novel content either for inspiration or for further refinement.

In this part of the course we will present the basics of these technologies and show examples how they are used inside intelligent tools for the creation of various media such as images, video, and 3D objects.

4.1 Images

In recent years we have seen a large growth in image fabrication techniques - especially using Generative Adversarial Neural Networks (GANs) [9]. Although these methods produce very realistic results, one of the key challenges that remain to be solved is their editability. Usually GANs are best at creating random images and influencing or modifying the resulting output image is very difficult. This is in contrast to earlier methods that presented tools where images were fabricated e.g., based on sketching their outline [3]. The use of direct search and retrieval based on the sketch allows greater flexibility and usefulness, which is more difficult to achieve using GANs. Some new trends emerge today that allow to fabricate real images using GANs based either on sketch inputs or rough segmentation maps [24; 2]. Still, the challenge remains of allowing editing operations on images to build tools using neural networks.

4.2 Video

If synthesizing an image involves photo-realism or feasible quality, synthesizing and editing movies is much more of a challenge. This is not only because of the added temporal dimension - many frames need to be fabricated, but also because semantic understanding and narrative meaning should be addressed. The challenge becomes not only synthesizing content but also cutting, framing, and assembling it. Editing a video is an task that needs expertise. Hence, intelligent tools that allow the editing or creating videos try to alleviate the problem by defining operations in a different domain. For example, text editing is a task that almost anyone can handle. Yet again, searching, classifying and mapping text and imagery to the same latent domain based on learning, can assist the definition of tools for creating meaningful montages [36], or for altering the words of a talking-head video [6].

4.3 3D Objects

When synthesizing 3D objects another large challenge comes up if one wants to actually fabricate the object. Physical and geometric constraints come into play and were dealt with in a previous section. Data-driven approaches can still assist by building semantic databases and using search and replace mechanisms to assemble objects, add connectors to existing models, or fabricate connectors between two objects from scratch [26; 15]. Tools built on such technologies relieve the user from tedious tasks such as exact positioning, the need for attention to small details and more.

5 Crowdsourcing

5.1 Course Overview

Tools can be more intelligent if they can adequately model human perception (*e.g.*, visual preference and perceived semantics). Crowdsourcing is an effective way to gather perceptual feedback on visual designs on demand, and researchers in computer graphics have proposed digital content creation tools with crowd intelligence in this decade. In this part of the course, especially focusing on parametric design scenarios, we describe

- how to model human preference via crowdsourcing (*e.g.*, ways to design crowdsourcing queries and learn a latent function from gathered data) and
- how to use the preference model to make tools intelligent.

For the latter, we especially discuss two approaches: embedding crowds into a user interface to facilitate users' exploration and decision making [14] and embedding crowds into an optimization engine to achieve automatic perception-aware parameter adjustment [16], with case studies on photo color enhancement and material BRDF design.

5.2 Background

The term, *crowdsourcing*, was first introduced by Howe [10] in 2006. He defined it as

> *"the act of taking a job traditionally performed by a designated agent (usually an employee) and outsourcing it to an undefined, generally large group of people in the form of an open call"* [11].

Amazon Mechanical Turk (MTurk)[1] is one of the most popular crowdsourcing platforms, where we can ask many crowdworkers to perform *microtasks* on demand to gain human intelligence. One of the most popular usages is to develop training datasets for machine learning, where researchers or practitioners ask crowdworkers to gather data (*e.g.*, gathering relevant texts from the Web) or annotate data (*e.g.*, classify images into one of some categories).

Human computation [25] is a concept that is closely related to crowdsourcing. This term was described by von Ahn, one of the pioneers of this paradigm, as

> *"a paradigm for utilizing human processing power to solve problems that computers cannot yet solve"* [35].

This paradigm does not necessarily assume that human processing power is supplied by crowdworkers. However, by using crowdsourcing platforms and formulating the process that needs human intelligence as a set of microtasks, we can reasonably compose algorithms that systematically incorporate human intelligence and can run on demand. We call this paradigm *crowdsourced human computation* [13] and algorithms (or systems) in this paradigm *crowd-powered* algorithms (or systems) [14]. If an algorithm is iterative and it involves crowdsourcing in each iteration, then it is also called a *crowds-in-the-loop* algorithm [16; 13] as a special case of *human-in-the-loop* settings. To the best of our knowledge, the work by Gingold *et al.* [8] was the first one that investigated this direction for creative graphics.

5.3 Types of Intelligence

Many researchers have proposed intelligent tools for creative graphics enabled by crowdsourcing. These works have been presented mainly in the computer graphics and human-computer interaction communities. About the types of intelligence, they can be roughly classified into the following categories:

- **General preference**: Consider the task of assessing, given a visual artwork (*e.g.*, a photograph), how subjectively good it is. This task is important in developing tools for creative graphics since such subjective goodness can be the primary objective in many scenarios. However, it is difficult

[1] https://www.mturk.com/

for computers to perform this task. For this, crowdsourcing enables us to systematically gather preferential feedbacks by many crowdworkers and then train an estimator to estimate general human preference (*i.e.*, preference learning) [41; 27; 14; 16]. The learned preference can be used for assisting exploration [14] or finding better options efficiently [16]. This course will discuss this type of intelligence in detail.

- **Semantic attributes**: Semantic attributes are human-understandable concepts that are often represented as adjectives (*e.g.*, "modern", "strong", and "romantic"), and they can help designers intuitively explore different visual options in digital content creation. However, determining the semantic attributes of a given visual option needs human intelligence. Chaudhuri *et al.* [1] used crowdsourcing to estimate semantic attributes of 3D models (*e.g.*, how "scary" does this animal head model looks?), and they implemented an attribute-based intelligent tool for part-based 3D modeling. Similarly, researchers have investigated how semantic attributes could facilitate digital content creation in various domains, including font selection [23], physically based animation [29], and more [30; 39; 5; 28].
- **Perceptual similarity**: Recommending alternative visual options to users would be a desirable function for content creation tools. One of the ways to achieve this function is to calculate similarities among visual options and then provide options that are similar to the current one. A challenge here is that the similarity measure should be perceptually reasonable; otherwise, the recommendation becomes meaningless. We can use crowdsourcing to obtain a perceptually meaningful similarity measure; we can do so by gathering perceptual feedbacks about visual similarity from crowds and then performing metric learning [23; 7; 20].
- **Perceptual compatibility**: Compatibility is defined among multiple objects, and we say that they are compatible if they go well together without conflict. In digital content creation, compatibility often matters when multiple visual elements exist in the same scene. Crowdsourcing can be used to evaluate or estimate perceptual compatibility between visual elements [19; 18].

5.4 Further Reading

We suggest those who want to learn more to refer to a book chapter [13], which discusses in detail how to formulate a typical visual design process as an optimization problem, how to use crowdsourcing for the problem, and how to design effective queries to crowdworkers. It is also notable that some algorithms developed for crowds-in-the-loop settings have the potential to be adapted to single-user (*i.e.*, user-in-the-loop) settings as well [17].

References

[1] Siddhartha Chaudhuri, Evangelos Kalogerakis, Stephen Giguere, and Thomas Funkhouser. Attribit: Content creation with semantic attributes. In *Proc. UIST '13*, pages 193–202, 2013. doi: https://doi.org/10.1145/2501988.2502008.

[2] Shu-Yu Chen, Wanchao Su, Lin Gao, Shihong Xia, and Hongbo Fu. DeepFaceDrawing: Deep generation of face images from sketches. *ACM Trans. Graph.*, 39(4):72:1–72:16, July 2020.

[3] Tao Chen, Ming-Ming Cheng, Ping Tan, Ariel Shamir, and Shi-Min Hu. Sketch2photo: Internet image montage. *ACM Trans. Graph.*, 28(5):1–10, December 2009. doi: https://doi.org/10.1145/1618452.1618470.

[4] Tao Chen, Zhe Zhu, Shi-Min Hu, Daniel Cohen-Or, and Ariel Shamir. Extracting 3d objects from photographs using 3-sweep. *Commun. ACM*, 59(12):121–129, December 2016. doi: https://doi.org/10.1145/3007175.

[5] Ruta Desai, Fraser Anderson, Justin Matejka, Stelian Coros, James McCann, George Fitzmaurice, and Tovi Grossman. Geppetto: Enabling semantic design of expressive robot behaviors. In *Proc. CHI '19*, pages 369:1–369:14, 2019. doi: https://doi.org/10.1145/3290605.3300599.

[6] Ohad Fried, Ayush Tewari, Michael Zollhöfer, Adam Finkelstein, Eli Shechtman, Dan B Goldman, Kyle Genova, Zeyu Jin, Christian Theobalt, and Maneesh Agrawala. Text-based editing of talking-head video. *ACM Trans. Graph.*, 38(4):68:1–68:14, July 2019. doi: https://doi.org/10.1145/3306346.3323028.

[7] Elena Garces, Aseem Agarwala, Diego Gutierrez, and Aaron Hertzmann. A similarity measure for illustration style. *ACM Trans. Graph.*, 33(4):93:1–93:9, July 2014. doi: https://doi.org/10.1145/2601097.2601131.

[8] Yotam Gingold, Ariel Shamir, and Daniel Cohen-Or. Micro perceptual human computation for visual tasks. *ACM Trans. Graph.*, 31(5):119:1–119:12, September 2012. doi: https://doi.org/10.1145/2231816.2231817.

[9] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Proc. NIPS '14*, pages 2672–2680, 2014. URL: http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf.

[10] Jeff Howe. The rise of crowdsourcing. https://www.wired.com/2006/06/crowds/, 2006. Last checked: October 23, 2016.

[11] Jeff Howe. Crowdsourcing: A definition. http://crowdsourcing.typepad.com/cs/2006/06/crowdsourcing_a.html, 2006. Last checked: October 23, 2016.

[12] Carlo Innamorati, Bryan Russell, Danny Kaufman, and Niloy J. Mitra. Neural re-simulation for generating bounces in single images. In *Proc. ICCV '19*, pages 8718–8727, 2019. doi: https://doi.org/10.1109/ICCV.2019.00881.

[13] Yuki Koyama and Takeo Igarashi. Computational design with crowds. In Antti Oulasvirta, Per Ola Kristensson, Xiaojun Bi, and Andrew Howes (Eds.), *Computational Interaction*, chapter 6, pages 153–184. Oxford University Press, 2018. doi: https://doi.org/10.1093/oso/9780198799603.003.0007.

[14] Yuki Koyama, Daisuke Sakamoto, and Takeo Igarashi. Crowd-powered parameter analysis for visual design exploration. In *Proc. UIST '14*, pages 65–74, 2014. doi: https://doi.org/10.1145/2642918.2647386.

[15] Yuki Koyama, Shinjiro Sueda, Emma Steinhardt, Takeo Igarashi, Ariel Shamir, and Wojciech Matusik. AutoConnect: Computational design of 3d-printable connectors. *ACM Trans. Graph.*, 34(6):231:1–231:11, October 2015. doi: https://doi.org/10.1145/2816795.2818060.

[16] Yuki Koyama, Issei Sato, Daisuke Sakamoto, and Takeo Igarashi. Sequential line search for efficient visual design optimization by crowds. *ACM Trans. Graph.*, 36(4):48:1–48:11, July 2017. doi: https://doi.org/10.1145/3072959.3073598.

[17] Yuki Koyama, Issei Sato, and Takeo Igarashi. Sequential gallery for interactive visual design optimization. *ACM Trans. Graph.*, 38(4):88:1–88:12, July 2020. doi: https://doi.org/10.1145/3386569.3392444.

[18] Lasse Farnung Laursen, Yuki Koyama, Hsiang-Ting Chen, Elena Garces, Diego Gutierrez, Richard Harper, and Takeo Igarashi. Icon set selection via human computation. In *Pacific Graphics Short Papers*, 2016. doi: https://doi.org/10.2312/pg.20161326.

[19] Tianqiang Liu, Aaron Hertzmann, Wilmot Li, and Thomas Funkhouser. Style compatibility for 3d furniture models. *ACM Trans. Graph.*, 34(4):85:1–85:9, July 2015. doi: https://doi.org/10.1145/2766898.

[20] Zhaoliang Lun, Evangelos Kalogerakis, and Alla Sheffer. Elements of style: Learning perceptual shape style similarity. *ACM Trans. Graph.*, 34(4):84:1–84:14, July 2015. doi: https://doi.org/10.1145/2766929.

[21] Kaichun Mo, Paul Guerrero, Li Yi, Hao Su, Peter Wonka, Niloy J. Mitra, and Leonidas J. Guibas. StructureNet: Hierarchical graph networks for 3D shape generation. *ACM Trans. Graph.*, 38(6):242:1–242:19, November 2019. doi: https://doi.org/10.1145/3355089.3356527.

[22] Kaichun Mo, Paul Guerrero, Li Yi, Hao Su, Peter Wonka, Niloy Mitra, and Leonidas Guibas. Structedit: Learning structural shape variations. In *Proc. CVPR '20*, pages 8859–8868, 2020.

[23] Peter O'Donovan, Jānis Lībeks, Aseem Agarwala, and Aaron Hertzmann. Exploratory font selection using crowdsourced attributes. *ACM Trans. Graph.*, 33(4):92:1–92:9, July 2014. doi: https://doi.org/10.1145/2601097.2601110.

[24] T. Park, M. Liu, T. Wang, and J. Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proc. CVPR '19*, pages 2332–2341, 2019. doi: https://doi.org/10.1109/CVPR.2019.00244.

[25] Alexander J. Quinn and Benjamin B. Bederson. Human computation: A survey and taxonomy of a growing field. In *Proc. CHI '11*, pages 1403–1412, 2011. doi: https://doi.org/10.1145/1978942.1979148.

[26] Adriana Schulz, Ariel Shamir, David I. W. Levin, Pitchaya Sitthi-amorn, and Wojciech Matusik. Design and fabrication by example. *ACM Trans. Graph.*, 33(4):62:1–62:11, July 2014. doi: https://doi.org/10.1145/2601097.2601127.

[27] Adrian Secord, Jingwan Lu, Adam Finkelstein, Manish Singh, and Andrew Nealen. Perceptual models of viewpoint preference. *ACM Trans. Graph.*, 30(5):109:1–109:12, October 2011. doi: https://doi.org/10.1145/2019627.2019628.

[28] Ana Serrano, Diego Gutierrez, Karol Myszkowski, Hans-Peter Seidel, and Belen Masia. An intuitive control space for material appearance. *ACM Trans. Graph.*, 35(6):186:1–186–12, November 2016. doi: https://doi.org/10.1145/2980179.2980242.

[29] Leonid Sigal, Moshe Mahler, Spencer Diaz, Kyna McIntosh, Elizabeth Carter, Timothy Richards, and Jessica Hodgins. A perceptual control space for garment simulation. *ACM Trans. Graph.*, 34(4):117:1–117:10, July 2015. doi: https://doi.org/10.1145/2766971.

[30] Stephan Streuber, M. Alejandra Quiros-Ramirez, Matthew Q. Hill, Carina A. Hahn, Silvia Zuffi, Alice O'Toole, and Michael J. Black. Body Talk: Crowdshaping realistic 3d avatars with words. *ACM Trans. Graph.*, 35(4):54:1–54:14, July 2016. doi: https://doi.org/10.1145/2897824.2925981.

[31] Nobuyuki Umetani, Danny M. Kaufman, Takeo Igarashi, and Eitan Grinspun. Sensitive couture for interactive garment modeling and editing. *ACM Trans. Graph.*, 30(4):90:1–90:12, July 2011. doi: https://doi.org/10.1145/2010324.1964985.

[32] Nobuyuki Umetani, Takeo Igarashi, and Niloy J. Mitra. Guided exploration of physically valid shapes for furniture design. *ACM Trans. Graph.*, 31(4):86:1–86:11, July 2012. doi: https://doi.org/10.1145/2185520.2185582.

[33] Nobuyuki Umetani, Yuki Koyama, Ryan Schmidt, and Takeo Igarashi. Pteromys: Interactive design and optimization of free-formed free-flight model airplanes. *ACM Trans. Graph.*, 33(4):65:1–65:10, July 2014. doi: https://doi.org/10.1145/2601097.2601129.

[34] Nobuyuki Umetani, Athina Panotopoulou, Ryan Schmidt, and Emily Whiting. Printone: Interactive resonance simulation for free-form print-wind instrument design. *ACM Trans. Graph.*, 35(6):184:1–184:14, November 2016. doi: https://doi.org/10.1145/2980179.2980250.

[35] Luis von Ahn. *Human Computation*. PhD thesis, Carnegie Mellon University, Pittsburgh, PA, USA, 2005. URL: http://reports-archive.adm.cs.cmu.edu/anon/2005/CMU-CS-05-193.pdf. CMU-CS-05-193.

[36] Miao Wang, Guo-Wei Yang, Shi-Min Hu, Shing-Tung Yau, and Ariel Shamir. Write-a-video: Computational video montage from themed text. *ACM Trans. Graph.*, 38(6):177:1–177:13, November 2019. doi: https://doi.org/10.1145/3355089.3356520.

[37] Tuanfeng Y. Wang, Duygu Ceylan, Jovan Popović, and Niloy J. Mitra. Learning a shared shape space for multimodal garment design. *ACM Trans. Graph.*, 37(6):203:1–203:13, December 2018. doi: https://doi.org/10.1145/3272127.3275074.

[38] Tuanfeng Y. Wang, Tianjia Shao, Kai Fu, and Niloy J. Mitra. Learning an intrinsic garment space for interactive authoring of garment animation. *ACM Trans. Graph.*, 38(6):220:1–220:12, 2019. doi: https://doi.org/10.1145/3355089.3356512.

[39] Mehmet Ersin Yumer, Siddhartha Chaudhuri, Jessica K. Hodgins, and Levent Burak Kara. Semantic shape editing using deformation handles. *ACM Trans. Graph.*, 34(4):86:1–86:12, July 2015. doi: https://doi.org/10.1145/2766908.

[40] Youyi Zheng, Xiang Chen, Ming-Ming Cheng, Kun Zhou, Shi-Min Hu, and Niloy J. Mitra. Interactive images: Cuboid proxies for smart image manipulation. *ACM Trans. Graph.*, 31(4):99:1–99:11, 2012. doi: https://doi.org/10.1145/2185520.2185595.

[41] Jun-Yan Zhu, Aseem Agarwala, Alexei A. Efros, Eli Shechtman, and Jue Wang. Mirror mirror: Crowdsourcing better portraits. *ACM Trans. Graph.*, 33(6):234:1–234:12, November 2014. doi: https://doi.org/10.1145/2661229.2661287.