A conversation with CHCCS 2018 achievement award winner Alla Sheffer

Alla Sheffer
University of British Columbia

ABSTRACT

A 2018 CHCCS Achievement Award from the Canadian Human-Computer Communications Society is presented to Dr. Alla Sheffer for her numerous highly impactful contributions to the field of computer graphics research. Her diverse research addresses geometric modeling and processing problems both in traditional computer graphics settings and in multiple other application domains, including product design, mechanical and civil engineering, and fashion design. CHCCS invites a publication by the award winner to be included in the proceedings, and this year we continue the tradition of an interview format rather than a formal paper. This permits a casual discussion of the research areas, insights, and contributions of the award winner. What follows is an edited transcript of a conversation between Alla Sheffer and Paul Kry that took place on 13 March, 2018, via Skype.

THE INTERVIEW

CHCCS: Hello Alla, and congratulations on your CHCCS achievement award!

Alla: Thank you Paul! I am truly honoured to receive this award, and thank you again for organizing and conducting this interview.

CHCCS: I believe you have been at UBC now for, is it 15 years? Looking back, would you say that a lot of things have changed, at UBC and in general with Canadian research in computer graphics?

Alla: It is 2018, and I started in 2003, so... yes, 15 years! In many ways UBC is the same place I joined 15 years ago - very collegial, very friendly, good at balancing research and teaching. What has somewhat changed is my perspective on things. I am in a very different spot now - instead of being the most junior professor, I'm suddenly one of the most senior! It's an interesting spot to be in, as both others and myself expect more from me - broader impact, mentoring, leadership... and I am trying to do more to fulfill these expectations. I also see a bunch of positive changes in our department, UBC and in general across Canada when it comes to research support and research ambition. More team building, more opportunities for team grant applications, more focus on excellence and impact.

In addition, our department is now going through a major hiring spree, a trend that I see across other Canadian schools as well. We had the big hiring bump, back when I was hired, and then things slowed down, and now we are again at a point where we have the opportunity to bring in a new generation of graphics faculty.

CHCCS: With many computer science faculties growing these days, would VR/AR and machine learning be strategic areas for new graphics faculty?

Alla: Yes, there are a few areas which are suddenly very hot. VR/AR is definitively one, where we are slowly starting to move from using VR/AR as a cool toy to exploring its less self-evident applications. For me one of the most interesting questions in this space is how to use the extra degrees of freedom provided by the 3D VR/AR interfaces for creating content. I believe there are lots of interesting opportunities in this space.

Deep learning is another interesting space, as we have yet to figure out what it can or cannot do. and where its inherent limitations are.

CHCCS: What are your thoughts on machine learning, and how would you say it intersects with things in your work?

Alla: Well, I think that machine learning in general and in particular deep learning are important optimization tools. These are wonderful tools for what I would call - and learning experts may not like it - principled parameter optimization or tuning. As I see it one can think of a lot of the learning machinery as a principled process for computing the coefficients of some functional defined over a large family of complex basis functions. In many spaces, it is very much the right thing to do, because you do very often have a sense of the right variables, or features, to "feed" into the learning framework, but you need to understand how to correctly combine them. And that's where I think machine learning in general and deep learning in particular are very natural. I have used these tools very successfully in such contexts.

My first foray into machine learning was in 2008 back when it was just starting to gain popularity. In this paper we looked at detecting the correct, or natural, upright orientation of man-made objects [6]. This is something that humans are very good at, even when they don't know the function of the object; somehow it easy for people to figure out how to place something upright. When we wanted to algorithmically detect this orientation we were able to identify a bunch of factors that were likely to affect whether you place the object this way, or that way, or another way. However, it was far from obvious how to combine these sometimes contradictory cues. Machine learning provided us with an effective framework to combine those cues in a principled way.

More recently, we used a similar approach when we looked at style analysis and transfer for man made shapes [14, 15]. We wanted to both identify style - for instance, if you look at a building, we want to know its architectural style - is this is a Byzantine church, or a Japanese temple? How does one detect a given style or assess stylistic similarity? Or, more challenging, how do I model a chair that matches the style of a given table? This is again a space where learning is a very natural thing to do: you have lots of cues as to what humans look at when assessing style or functionality. But one needs to figure out how to combine them... Learning tools are ideal for such tasks. Most recently we used ML tools to assist in parsing and vectorization of raster region boundaries on low resolution clipart and icon images [9].

CHCCS: On the topic of previous work, what would you consider to be your most significant contribution over these years, or perhaps your favorite result?

Alla: Well, probably the first one was the ABF, the angle based flattening [21, 22], mostly because I kind of stumbled on the problem of parameterization almost by accident. I came up with a solution, which is really simple and literally uses high-school level trigonometry; you can explain the ABF algorithm to anyone who

learned basic formulas such as the sine rule. Despite its simplicity it is a surprisingly effective, powerful and useful algorithm.

I also really like the True2Form [23], and CrossShade [20] frameworks. They both target interpretation of designer sketches, a new and very intriguing space. One of the reasons I like this line of work, is that it was very spontaneous. CrossShade started from an almost accidental discussion between me and Adrian Boisseau. I visited the graphics group at Inria Sophia-Antipolis, and Adrien and I had this long friendly chat where we brought up the problem of understanding shape from cross-hairs and I found myself thinking, "wait a minute, this is cool and interesting" and so we started looking at cross-section drawings in more depth, and came up first with CrossShade and then True2Form.

I like both those two projects and the entire series of works on sketch analysis that follow up from them [3, 7, 2, 18, 4, 11]. The common thread through all these methods is an attempt to understand the choices designers make. Our key insight was that the choices designers make in creating content are very strongly connected to the perceptual cues viewers employ when observing and parsing this content. Thus when developing algorithms for parsing designer sketches or other inputs, we heavily rely on perception studies, as well as on observations about specific choices designers make and specific techniques they employ. It is a very interesting space, since there is this connection between psychology, design, and math. While from an algorithmic perspective it all boils down to math, the challenge is in figuring out which math, and why, since we somehow need to mathematically formulate the different cues and then and then combine them in a principled way.

CHCCS: That's a very interesting comment about stumbling upon both interesting problems and graceful solutions. Is this part of the key to success, or are there strategies to problem solving or problem selection that you would pass on to others?

Alla: I don't know if it is the key to success, but for me this is key to keep being excited about and engaged in research. I know lots of colleagues who are much more systematic than I and are very successful. For me, this discovery of new problems is really what makes research interesting. You find yourself thinking, "Hey! This a problem I've never thought about before. It is outside of my narrow comfort zone, but it is cool and interesting, and it looks like, maybe, what I know can in fact be used to solve it."

CHCCS: So it sounds like the suggestion is follow your heart and work on the fun problems that are just outside of your comfort zone then? Maximize your research happiness?

Alla: Right! Try to look for these problems, and look for collaborations. Most of the stuff you stumble onto... well, a lot of it comes from talking to people who you hadn't necessarily collaborated with before, and then going "Yay - this is interesting!"

CHCCS: Do you think you have a most under-appreciated paper? Is there a hidden gem out there that got published at a less visible venue, but is something that people should be aware of?

Alla: I have a couple of papers that, looking back, I should have done a better job of promoting, and pushing to better venues. One such paper is is my first foray into modeling garments using a sketching interface [19]. The paper is getting reasonable citation counts these days, so I think it is sufficiently visible, but I think it is somewhat under appreciated. In many ways this research was a precursor to much of the work on perception-driven modeling by myself and others. It asked (and answered) the question of why, when people look at a drawing of a garment in front view, they have a very clear idea in their head of how that garment will look in profile or any other view. The 3D garment shape humans perceive is related to the physics of garments, but also to how people interpret drawings based on what they know about both cloth physics and garment designs. For me, looking back, I learned a lot when I did that work. It is very interesting work because it sort of asks a lot of exactly those perceptual questions very early on, before me and others started working more in this space.

CHCCS: What parts of your previous work get the most use in industry? What do you think about tech transfer and commercialization?

Alla: Not as much as I would like. I know that several of my early methods such as ABF [22] are used in commercial softwares, based on what I heard from folks at companies such as Dassault or Pixar. In the past tech transfer was relatively low on my priority list. However, these days we are looking at trying to license our IP more; we have several methods which has been patented, such as our work on hex mesh generation [13, 8, 12], and garment transfer [5]. I'd like to see our algorithms more heavily used. One of my goals in the next few years is to have a bigger industry impact. I think that is an interesting and important challenge. Too often it takes decades for ideas to percolate from academia to industry.

CHCCS: You mentioned that there were papers that came very easily, for instance the ABF, but it is likely more common that research is the result of more difficult process. Do you have any "behind the research" stories about results that were surprisingly more difficult than what you would have otherwise expected?

Alla: No, not really. When you look back, most papers are about trial and error. You start with a problem, you come up with a potential a solution, you try the solution, and you typically find that it doesn't work or doesn't work well enough. Thats how scientific research works. You try something, you look at the results, you go "Oh, it is wrong!". Then you see why it is wrong and you try the next idea, and the next idea. From my experience, most research fits into this iterated model. I can't think of a single recent instance where the first tested solution worked.

CHCCS: So it *never* works on the first try?

Alla: Very rarely. Interesting research problems are rarely easy to solve. There is the famous statement by Albert Einstein: If we knew what we were doing, it wouldn't be called research.

CHCCS: There are many examples in your work that look at the connections between perception and math. Obviously there are challenges in connecting such disparate topics and I'm wondering if you could comment on this.

Alla: We kind of touched on this already, but yes, I find this space very interesting. There are lots of insights to be had. When you look at perception, math, and design, and how those things interact... for instance, look at sketch interfaces.

Humans can effectively communicate shape via sketches, but if you want to try to use a sketch as an input for a computer algorithm then you need to understand what makes it effective. It is an interesting domain because the reason that artist sketches are effective is not because one day somebody woke up and said "this is how I should do sketching". Instead we have many centuries during which artists developed techniques which they found effective, and these techniques were taught in art and design schools and passed along from one generation to another. Their choices are clearly motivated by how humans perceive sketched shapes. Perception research provides many cues in this regard, but too often these are tested on simple stimuli and can't be applied directly to interpret actual complex sketches. Finally, you have us computer scientists who ask "OK, how do I take all this know-how from design, and perception, and try to combine it into tools?"

Work in this domain is very much about translation between different languages, because neither artist nor perception research talk in terms of mathematical formulas. Often when we ask designers why they make the choices they make, they can't necessarily explain them. So it is kind of this interesting space where it really is about translating from the language of design to the language of mathematics, and also trying to understand where the translation breaks down and where you need to fill in the missing pieces.

A lot of computer graphics has well-prescribed problem statements. If you think about parameterization, the problem is very well defined. You know you want to take A, produce B, and in the process you have a bunch of metrics you want to optimize. So the challenge is just to come up with the algorithm. In this new space, the challenge is to come up with the mathematical definition of the problem, because you need to formulate the problem before you can try solving it. But just formulating it, is a big challenge in itself.

CHCCS: I imagine the future holds a number of interesting challenges in the coming years in this area.

Alla: Yeah, I think there are lots of questions. The human ability to understand and to reason about things that are much more, let's say, amorphous, than mathematical is much higher than that of existing algorithms. I don't think computers will ever reach the same level as humans, but we can try to use reasoning about perception and artist choices to provide tools for people which will make it easier for them to create or manipulate shapes.

CHCCS: Are there any other big challenges in coming years that you might think would be a good thing for new grad students to get involved with now?

Alla: Specific challenges? Well... one of the interesting ones would be develop tools that use VR interfaces for content creation. A lot of our tools for interaction and content creation in recent work have been about moving from menu driven interfaces to more intuitive, sketch-based ones. These interfaces operate in a two-dimensional space, and one of the core challenges in using them is to imbue the 2D user input with the third, depth dimension. VR interfaces are three-dimensional, thus one should be able to employ them to create and manipulate content directly in 3D... The research question here is: does this extra degree of freedom make things easier? Maybe, but maybe not. I think there is lots of interesting stuff there.

Another space I find intriguing is developing content creation tools geared toward non-experts. Until recently, tablets were something that only very few people could afford, but now all of a sudden you have these touch and pen interfaces everywhere. So trying to use those to allow a much broader user base to create content is an intriguing research direction. A lot of our content creation tools were originally targeted towards expert modelers, who can spend years learning how to use ZBrush, or Maya, or 3D Studio Max, but if you want to provide tools for somebody at home to design their own dress, or design a piece of jewelery that they will then send to a 3D printer, you need very different interaction frameworks.

So I think these are very interesting research directions. Both a different target user set, and different possibilities for interfaces. Another intriguing space is design tools targeting digital manufacturing and 3D printing. Designing manufacturable objects requires accounting for physics in addition to geometry.

CHCCS: Are there are elements in your background that you found helped you in your career? Something unexpected?

Alla: Well, obviously there is math and geometry and computer science, but one of the things that these days that I'm finding surprisingly useful is familiarity with a range of application domains, such as fashion or art history.

One of the reasons I had been successful in coming up with effective tools for garment design [5, 1, 10] is that I do have experience with garment making. It's something I learned back when I was a teenager, and I never thought it would ever be useful, beyond a hobby. This experience helps me recognize the problems and have intuition about what kinds of solutions designers would want.

Similarly, I was always interested in art, reading piles of books, going to museums, and so on. This background knowledge was extremely useful when we worked on tasks such as computational calligraphy [16, 24], or when we worked on problems such as shape abstraction [17], or style analysis for artifacts such as furniture or building. For instance, in the latter case, my ability to reason about style, and to know or identify styles, and therefore try to reason back from this knowledge was extremely handy. So bits of general knowledge, which become very useful for computer graphics research.

CHCCS: Do you have time to do much tailoring these days?

Alla: I did sew a doll dress for our SIGGRAPH 2016 paper on physics-drive pattern adjustment [1]. For the fold design paper that was just accepted to SIGGRAPH 2018 [10] we had a professional make a few fabric pieces. I sadly no longer have the time for such tasks.

CHCCS: What sources do you look at when addressing such problems as sketch based modeling, or garment modeling?

Alla: The most important aspect for me is learning from domain experts. Reading design literature or fashion literature, and trying to understand what domain experts do and why. It is an interesting task because it is not something that a computer science background really prepares you for. So it is interesting and useful at the same time, but it is also a bit of a challenge. You need to try to understand what domain experts who are speaking a different language are trying to say.

CHCCS: What might students like to know in preparing for successful careers in geometry processing or computer graphics? What words of wisdom or advice would you give to new graduate students?

Alla: There are clearly a lot of obvious skills which are good to have, such as programming or math, but on top of those I would suggest they learn how to express their ideas in writing. Learn to explain why what they do is important, and why it is new. Learn to identify, after the fact, what is cool in what you did. You need to look back at your project and ask "OK, what ended up being cool and novel here?" Not what took me the most time, and what I thought was going to be the highlight at the beginning of the project, but try to reverse engineer what actually ended up being the key thing. It is not easy but it's something you really need to learn.

CHCCS: Would you say there be any interesting emerging directions that might influence future computer graphics research? Where do you see the changing trends or exciting problems?

Alla: There is a lot of research which is very linearly moving forward. We solve problem A, now let's solve problem A better. While this work is clearly very valuable, the works I find most exciting are those that find not only new solutions but new problems. I am most fascinated by papers that solve simple problem that nobody has ever looked at before. The problems they solve are obvious in hindsight, but still had never been addressed. For me, those are the best papers.

CHCCS: Instead of looking ahead 50 years, I suppose there is a big 50 year retrospective birthday party coming up at UBC for the computer science department?

Alla: Yeah, UBC CS is going to be 50! So that is exciting, because we are slowly becoming a mature department and a mature field. I think it is an interesting point for all of computer science, because we need to figure out what is next. We have made a lot of progress in many spaces, but still have lots of open problems to explore. There is a lot of momentum and excitement in graphics right now: virtual and augmented reality, digital fabrication, applications of machine learning tools... I expect us to see lots of cool ideas and technologies coming up. UBC is very well positioned to be a major player in this field - we already have one of the strongest graphics groups in the world and are actively hiring new faculty.

CHCCS: It will be exciting to see what comes next. Thank you for the conversation, and congratulations again on your CHCCS achievement award. I will look forward to seeing you at at Graphics Interface!

Alla: Same here! Will be great to see you again.

REFERENCES

[1] A. Bartle, A. Sheffer, V. G. Kim, D. M. Kaufman, N. Vining, and F. Berthouzoz. Physics-driven pattern adjustment for direct 3d garment editing. *ACM Trans. Graph.*, 35(4):50:1–50:11, July 2016.

[2] M. Bessmeltsev, W. Chang, N. Vining, A. Sheffer, and K. Singh. Modeling character canvases from cartoon drawings. *Transactions on Graphics (2015)*, 34(5), 2015.

[3] M. Bessmeltsev, N. Vining, and A. Sheffer. Gesture3d: Posing 3d characters via gesture drawings. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH Asia 2016)*, 35(6), 2016.

[4] M. Bessmeltsev, C. Wang, A. Sheffer, and K. Singh. Design-driven quadrangulation of closed 3d curves. *Transactions on Graphics (Proc. SIGGRAPH ASIA 2012)*, 31(5), 2012.

[5] R. Brouet, A. Sheffer, L. Boissieux, and M.-P. Cani. Design preserving garment transfer. *ACM Trans. Graph.*, 31(4):36:1–36:11, July 2012.

[6] H. Fu, D. Cohen-Or, G. Dror, and A. Sheffer. Upright orientation of man-made objects. In *ACM transactions on graphics (TOG)*, volume 27, page 42. ACM, 2008.

[7] G. Gori, A. Sheffer, N. Vining, E. Rosales, N. Carr, and T. Ju. Flowrep: Descriptive curve networks for free-form design shapes. *ACM Transaction on Graphics*, 36(4), 2017.

[8] J. Gregson, A. Sheffer, and E. Zhang. All-hex mesh generation via volumetric polycube deformation. *Computer Graphics Forum*, 30(5):1407–1416.

[9] S. Hoshyari, E. Dominici, A. Sheffer, N. Carr, D. Ceylan, Z. Wang, and I.-C. Shen. Perception-driven semi-structured boundary vectorization. *ACM Transactions on Graphics (to appear)*, 2018.

[10] M. Li, A. Sheffer, E. Grinspun, and N. Vining. Foldsketch: Enriching garments with physically reproducible folds. *ACM Transactions on Graphics (to appear)*, 2018.

[11] C. Liu, E. Rosales, and A. Sheffer. Strokeaggregator: Consolidating raw sketches into artist-intended curve drawings. *ACM Transactions on Graphics (to appear)*, 2018.

[12] M. Livesu, A. Sheffer, N. Vining, and M. Tarini. Practical hex-mesh optimization via edge-cone rectification. *Transactions on Graphics (Proc. SIGGRAPH 2015)*, 34(4), 2015.

[13] M. Livesu, N. Vining, A. Sheffer, J. Gregson, and R. Scateni. Polycut: Monotone graph-cuts for polycube base-complex construction. *Transactions on Graphics (Proc. SIGGRAPH ASIA 2013)*, 32(6), 2013.

[14] Z. Lun, E. Kalogerakis, and A. Sheffer. Elements of style: Learning perceptual shape style similarity. *ACM Trans. Graph.*, 34(4):84:1–84:14, July 2015.

[15] Z. Lun, E. Kalogerakis, R. Wang, and A. Sheffer. Functionality preserving shape style transfer. *ACM Trans. Graph.*, 35(6):209:1–209:14, Nov. 2016.

[16] R. Maharik, M. Bessmeltsev, A. Sheffer, A. Shamir, and N. Carr. Digital micrography. *Transactions on Graphics (Proc. SIGGRAPH 2011)*, pages 100:1–100:12, 2011.

[17] R. Mehra, Q. Zhou, J. Long, A. Sheffer, A. Gooch, and N. J. Mitra. Abstraction of man-made shapes. *ACM Transactions on Graphics*, 28(5), 2009.

[18] H. Pan, Y. Liu, A. Sheffer, N. Vining, C. Li, and W. Wang. Flow aligned surfacing of curve networks. *ACM Trans. Graph. (SIGGRAPH)*, 34(4), 2015.

[19] C. Robson, R. Maharik, A. Sheffer, and N. Carr. SMI 2011: Full paper: Context-aware garment modeling from sketches. *Comput. Graph.*, 35(3):604–613, June 2011.

[20] C. Shao, A. Bousseau, A. Sheffer, and K. Singh. CrossShade: shading concept sketches using cross-section curves. *ACM Transactions on Graphics*, 31(4), 2012.

[21] A. Sheffer and E. de Sturler. Parameterization of faceted surfaces for meshing using angle-based flattening. *Engineering with computers*, 17(3):326–337, 2001.

[22] A. Sheffer, B. Lévy, M. Mogilnitsky, and A. Bogomyakov. Abf++: Fast and robust angle based flattening. *ACM Trans. Graph.*, 24(2):311–330, Apr. 2005.

[23] B. Xu, W. Chang, A. Sheffer, A. Bousseau, J. McCrae, and K. Singh. True2Form: 3d curve networks from 2d sketches via selective regularization. *ACM Transactions on Graphics*, 33(4), 2014.

[24] C. Zou, J. Cao, W. Ranaweera, I. Alhashim, P. Tan, A. Sheffer, and H. Zhang. Legible compact calligrams. *ACM Trans. Graph.*, 35(4):122:1–122:12, July 2016.