REVISTA DE CIENCIAS DE LA COMPUTACIÓN

Dr. Alex Bronstein

New dimensions of media

En este interesante artículo, el Dr. Bronstein, explica la evolución, importancia, y las aplicaciones del análisis de modelos 3D. Así, como sus principales aportes al desarrollo de esta fascinante área. Bronstein es profesor de Ingeniería Eléctrica de la Universidad de Tel Avid, autor de más de 100 publicaciones en Journals y Conferencias, con una docena de patentes y aplicaciones patentadas y autor del libro *Numerical geometry on non-rigid shapes*. Continua en la 2

El camino de la computación

Dr. Omar Flores

En este artículo el Dr. Flores, científico investigador del grupo personalizado de la computación en *Intel Labs*, describe como dos jóvenes con espíritu y ganas de superación, deciden viajar a otro país, donde la ciencia y la tecnología está más desarrollada, con el objetivo de convertirse en mejores personas y mejores profesionales gracias al poder de una educación de alto nivel. Por otro lado, insta a estudiantes relacionados a la computación a sumergirse en el mundo de la ciencia y la investigación, además del fascinante mundo de la programación. Continua en la página 7

Cursos: Análisis computacional de modelos deformables

IPRODAM GROUP

La Universidad La Salle de Arequipa - Perú, mediante su grupo de investigación IPRODAM, organizaron, por primera vez en nuestro país, el curso de verano en análisis computacional de modelos deformables realizado del 2 al 10 de febrero. El curso fué diregido por el Dr. Alex Bronstein y tuvo una duración de 4 horas diarias. Estuvo dirigido a alumnos de últimos años de pregrado y, especialmente, para alumnos de postgrado con sólidos conocimientos en matemática y programación. El curso culminó con exposiciones por parte de los miembros del grupo de investigación en procesamiento de imágenes y minería de datos (IPRODAM) de la Universidad la Salle.

ÍNDICE

Universidades de La Salle en el Mundo

Universidad La Salle México | Universidad La Salle Brasil UNILASALLE | Universidad de la Salle Bogotá | La Salle Universities | Escuelas Cristianas La Salle

EDITOR
Dr. Cristian López Del Alamo
clopez@ulasalle.edu.pe

DISEÑADOR GRÁFICO
Jorge Luis Contreras Cano
http://ulasalle.edu.pe

New dimensions of media

Dr. Alex Bronstein is Associate Professor of Electrical Engineering at Tel Aviv University. He received his Ph.D. in Computer Science, Technion in 2007. His main research interests are theoretical and computational methods in metric geometry and their application. He has authored over 100 publications in leading journals and conferences, over two dozens of patents and patent applications, and the book Numerical geometry on non-rigid shapes

Por Dr. Alex Bronstein

Among all products of the modern era, electronic media has left a remarkably profound imprint on our society, probably only comparable to that of the internal combustion engine or antibiotics. In its turbulent history, dating back only some century and a half ago (roughly starting from Morse's early experiments in the 1830's), pivotal technology breakthroughs are related to an increase in the number of dimensions. A bold example is the transition from telephone and radio to television – first black and white and then color. Today, we are past another such transition and are witnessing the rapidly emerging three-dimensional (3D) media ecosystem. This ecosystem is powered by several technological breakthroughs. The first one is, doubtlessly, the growing availability of depth or "RGBD" (red, green, blue and depth) image sensors – cameras, in which each pixel is capable of sensing distance to the object in the scene, on top of its usual ability to measure the amount of light in the three visible spectral bands.

When I started my graduate studies at the Technion in Israel in the early 2000's, depth sensors were largely the prerogative of high-end precision manufacturing costing frightening four- or five-digit figures. Starving for large quantities of 3D data for my Ph.D. thesis, I had no choice but to built my own depth camera based on structured light (i.e., triangulation from code projected onto the scene and captured by a regular camera). The first prototype was bulky and had rather long acquisition times. Since I was interested in 3D face recognition, I asked my subjects to stand still and not to breathe, and when they refused to collaborate, I had to put them into a menacingly-looking wooden machine suspiciously resembling a guillotine. Even with these primitive tools, me and my twin brother Michael were able to build a 3D-based face recognition system capable of distinguishing between ourselves, as well as between many other identical tweens, including an amazing triplet of joyful young ladies similar as three drops of water. The fact that a machine can distinguish between its identical twin creators appealed much to the reporters, and for about six months that followed our lab was frequently visited by people with cameras and microphones, publishing a few dozens of articles and reportages in a dozen of languages, culminating in the ridiculously-titled CNN Technology report Twins crack the face recognition puzzle.

A few years later, after making our 3D camera smaller, faster, and more accurate, we took it together with our face recognition system onto a roadshow to the East and the West coasts in an attempt to raise money for a commercial project. We toured all the top-tier venture capital funds on Sandhill road, trying to convince them to invest in 3D technology. And, like the Little Prince with his drawings, we heard the same answer over and over – that 3D had no market, and that we were trying to boil the ocean. (We eventually raised money for a different enterprise, but this is an entirely different story.)

QWERTY keyboard of a Remington typewriter.

Everything changed in 2010 when an Israeli startup company PrimeSense has partnered with Microsoft to develop what was initially called Project Natal. Xbox makers were looking for ways to reinvent themselves, but I doubt that anybody could ever foresee the success of what has become the Kinect that literally disrupted human-machine interaction paradigms in the post-iPhone world, and made seamless gesture-based gaming a reality. Controlling a game using your bare hands and your body – no keyboard, no remote, just you and your digital alter ego – is, indeed, quite natural. It has always puzzled me how much time it took such a natural idea to mature. For the past 140 years, since its introduction in 1873 by Remington, we have been using the QWERTY keyboard as the principal means of interaction with our machines. The seemingly random layout of this keyboard has been studied to avoid mechanical jams of typewriter keys. Typewriters have long become museum exhibits and electronic keyboards do not jam anymore, but mental inertia covered up by the fancy term "backward compatibility" is hard to die.

Playing videogames with Microsoft Kinect. Your body is the ultimate game controller.

About the same time, I joined three other nerds in founding the startup Invision that was attempting to miniaturize our structured light depth sensor. I attribute it to an incredible strike of luck that our effort has come to the radars of the semiconductor supergiant Intel who acquired the company after eighteen months on the road with only slides and barely working prototypes, and entrusted us with creating world's first integrated depth sensor. In 2014, after two years of fighting against the laws of physics, our technology came to the world under the RealSense brand.

Intel's Senior Vice President Mooly Eden holding the RealSense camera module at CES 2014. Technology made in Israel. Image credit: Intel.

Intel's RealSense camera has set a new technological landmark: due to its ultra-compact form factor and low price it can be placed into the lid of a laptop or the bezel of a tablet, promising to make depth acquisition truly affordable and ubiquitous. RealSense is already being integrated in ultrabooks such as Lenovo Yoga and many other devices. With Apple (after the purchase of PrimeSense), Google (with its Tango project), and Samsung (with their signature Asian secrecy) following suit, RGBD sensors have reached a pivotal point that I find similar to the transition from the bulky webcam peripherals to tiny cameras integrated into laptops and hand-held devices that occurred about 15 years ago. The consequence of the latter transition that turned RGB cameras into our intimate companions was literally an explosion of all kinds of previously unimaginable applications that profoundly affected our society and changed our everyday life's habits (Skype and Facebook selfies are just two trivial examples). The availability of integrated depth sensors is likely to have a similar impact, most of which we are yet unable to foresee and fully appreciate.

Intel RealSense sensor integrated into a laptop allowing gesture-based interaction. Image credit: Intel.

Combined with the exponentially growing computing power on all platforms (the famed Moore's law), the availability of integrated RGB cameras has driven many image processing and computer vision tasks out of the domain of science to the realm of engineering. Applications involving automatic image enhancement, face detection and recognition, optical character recognition, or image-based search are nowadays ubiquitous. Depth sensors promise to solve numerous other, currently unsolved, image processing and computer vision tasks. For example, segmenting foreground from background or refocusing an image at a given depth in real time is still largely impractical with RGB cameras, but is almost trivial with the addition of depth information. Other tasks like gesture and facial expression recognition and gaze tracking become an order of magnitude more precise when using an RGBD sensor.

Depth sensing is a formidable tool allowing our machines to better understand and interact with the physical world around them, which is the fundamental challenge of computer vision. However, it brings forth the need to process and analyze three-dimensional information that is remarkably more complicated than processing and interpreting 2D images. One of the reasons is that 3D data are inherently non-Euclidean, obliterating many natural notions such as regular sampling and global coordinate systems. Another reason is that 3D geometries surrounding us are highly deformable – we live in a non-rigid world from macro to nano. We applaud the art of a ballerina who by a mere posture of her body can convey an entire world of emotion, or are amazed by the strength and technique of an acrobat performing seemingly impossible exercises. Yet, analyzing the enormous number of degrees of freedom such bending, twisting, folding, and articulation allows is literally a nightmare.

My interest in the analysis of deformable objects dates back to the last year of my undergraduate studies at the Technion in early 2000's. At that time, the field was nearly virgin and reputed a tough nut in machine vision. The lack of simple parametrization of the deformations rendered many problems apparently intractable. In my Ph.D. research that I did under the supervision of Prof. Ron Kimmel, who since then has become a life-long mentor, colleague, business partner, and friend, I was probably among the first to introduce a framework for modeling deformable shapes as metric spaces and deal with the theoretical and computational aspects of deformable shape similarity, symmetry, and correspondence. We were able to show that under a fairly rich family of approximately distance-preserving deformations, these important problems could be efficiently addressed through a generalization of multidimensional scaling, a popular tool from statistics. These tools became a reference in the field, and laid the foundations of its first systematic treatise that I co-authored under the name of Numerical geometry of non-rigid shapes (published by Springer in 2008). Today, I am pleased to see that in the decade that has since passed the field is flourishing and has attracted many talented scientists and engineers. These developments will certainly have a large impact on the entire 3D ecosystem, and they again remind me the avalanche effect that digital signal processing underwent after the invention of the compact disc, and image processing underwent after the advent of the digital camera.

Another important driver of the 3D ecosystem are 3D printing technologies, which in a certain sense close the loop, allowing a digitally processed scan of a real-world object to become tangible again. In the past few years, thermal extrusion technology made a sub-$1000 3D printer a reality, in part thanks to the incredible success of MakerBot (founded in 2009 and acquired by Stratasys in 2013 for $400 million). While the choice of materials and printing quality, speed, and cost are still unsatisfactory for many needs, 3D printing technologies are rapidly improving.

MakerBot founders (left to right) Adam Mayer, Zach Smith and Bre Pettis with the f nal MakerBot Cupcake prototypes. Image credit: Wikipedia.

Here, another analogy with the 2D world seems unavoidable: In the past two decades, the rapid development of digital imaging has revolutionized the printing market. In the 90's, a typical consumer like myself would use a 9-dot matrix printer to print simple text documents or low-quality images and would resort to the services of a photographic studio or a professional printer for high-quality photographs and documents. In the 2000's, everyone can print high-quality photos and documents on an affordable desktop printer, thanks to the inkjet and laser printing technologies. While today's thermal extrusion desktop 3D printers still strikingly resemble the old dot matrix printer, companies like 3D Systems and Stratasys are building industrial-grade printers capable of combining tiny drops of different substances at a microscopic level, producing completely new, previously unthinkable, materials with continuously varying mechanical, optical, and electromagentic properties. According to many analysts, the entrance of the printing giant HP into the 3D printing business promises the transition of these technologies from multimillion dollar machines to desktop printers. The impact of these technologies on our society is hard to grasp; I personally believe that sooner or later it will spawn an industrial 'anti-revolution,' in which many manufacturing and production processes will become decentralized again and performed on the consumer's desk.

I have no doubt that we are living through a revolution that will profoundly impact our society and permanently change our everyday habits. The new dimension of information brings formidable scientif c, technological, engineering and business challenges, and creates amazing new opportunities. The "killer application" of 3D is yet terra nondum cognita, and I invite my students and my fellow professors, engineers, and entrepreneurs to invest their brain energy in this incredible f eld. ■

Curso de *Shape Analysis* con el Dr. Alex Bronstein en la Universidad la Salle de Arequipa