Ren Ng – Research Statement

It is only within the last decade that digital photography has finally eclipsed film in sales, unit volume and number of pictures taken. But when the tipping point arrived, the societal effects of the paradigm change were swift, profound and ubiquitous. In recent years, something has begun to emerge that many consider the likely successor to digital. It goes by the name "computational photography," and I believe it has the potential to be just as revolutionary a change in paradigm. Early successes include high dynamic range imaging, video stabilization and panoramic stitching, which are now found in mobile phone cameras; and computational browsing over collections of images, such as gigapixel imaging and photo tourism. In spite of these early successes, computational photography today is limited by working from 2D representations – that is, each of these examples works by combining conventional digital photographs.

Light field photography, introduced in my doctoral thesis [1], yields a fundamentally richer representation, capturing the 4D light field (set of all light rays) in every shot, not just a 2D projection (photograph). The richer representation is necessary to more deeply realize the potential of integrating advanced computation into the photography pipeline. In fact, it has enabled a solution to some of the oldest problems in photography. For example, since the beginning in the 1830s, we have had to focus the camera and choose the depth of field by setting the aperture *before* we take the shot, and lens designs have been severely constrained by the need to control aberrations in hardware. With light field photography, we can refocus and choose a different depth of field in software *after* we take the shot, and lens designs have been freed by the new ability to compensate for general aberrations in light field processing.

Nevertheless, realizing the full potential of light field photography is hard, and there is a long way to go. One of the main challenges is that today's systems [8] offer lower resolution and require longer processing times than traditional digital cameras. I believe the overall benefits of light field photography will begin to eclipse digital photography at approximately 100 million rays (today's systems capture 11 million rays), and will grow exponentially to at least 1 billion rays. This will be analogous to digital photography eclipsing film as it grew from several million pixels to tens of millions.

Achieving the stated growth for light field photography requires a rethinking of camera design, addressing new problems and research opportunities over the entire imaging pipeline, including sensor design, optical design, theory and algorithms for representing and processing light fields, new user interaction models for devices and photographs, and imaging applications across the arts and sciences. My research has attacked, and will continue to attack, this problem on all these fronts because it is necessary to realize the full potential of the new paradigm in photography. Multidisciplinary collaboration is core to my approach, as discussed below.

Overview of Experience - Preparation for Faculty Research

My dissertation introduced the light field photography system, comprising a camera with a fundamentally different kind of sensor that records the light field in every exposure, and a computational engine to render photographs flexibly from the recorded light field. My research made contributions in system design, by explaining the optical principles and trade-offs in designing light field cameras; in mathematical analysis, by deriving theoretical tools for calculating performance and designing algorithms; and in experimental validation, by building prototype cameras and using them to test theory and explore the new art of light field photography. The first light field art projects attracted media interest, and led to initial public awareness of the research.

Following my graduate studies, I started a company, today called Lytro, to bring my dissertation research to market. I learned a tremendous amount technically and entrepreneurially as I ran the company as CEO through 70 employees and shipping the first consumer light field camera for $399 in 2012. This milestone accomplished my goal of putting the technology in the hands of everyday people. The company's success has helped to elevate the profile of the field. Lytro has been featured in hundreds of news articles [9,10,11] and won numerous awards [12,13,14], which has attracted attention from millions of people to computational photography. As one measure, my Ph.D. thesis was downloaded at a peak rate of almost 8,000 times a day after we launched the product.

Although I did not have the opportunity to publish papers while running the company, I have remained deeply engaged in technical leadership, architecture and invention. I have developed a deep understanding of the global supply chain for consumer imaging, and have architected major parts of the company's strategy for developing critical component supply. This has given me an intimate working view of the real needs of the field. A list of patents and patent applications can be found in my CV. Lytro has developed an exciting technology and product roadmap to execute in the coming years, and I am proud to say that the company truly represents the state-of-the-art in computational photography today.

I believe that my experiences as a theoretician, engineer, entrepreneur and executive in the global camera market have given me a unique perspective and preparation to become a professor in computational photography. For my technical work, I won the ACM Dissertation Award, Arthur Samuel Award at Stanford, the Selwyn Award, the TR35, the PMDA Technical Achievement Award, and I was inducted into RIT's Imaging Hall of Fame. For my entrepreneurial work, I was recognized with MIT Tech Review's Entrepreneur of the Year, Silicon Valley Journal's 40 Under 40, and Fast Company's 100 Most Creative People in Business. I have a deep commitment to the pursuit of excellence in my work, and I look forward to building a world-class research program that drives the state of the art in computational photography, and maximizes the transfer of benefits to society at large. I anticipate my research agenda and multidisciplinary collaborations will grow over time into all the following areas, which are critical for realizing the full potential of computational photography.

Computational Camera Systems Engineering

To fully eclipse digital, light field cameras will need to grow to hundreds of millions of light rays in each exposure - I believe the opportunity is clear to scale exponentially to 1 billion rays and beyond. The high potential for growth comes from the image sensor industry, which has shrunk the mass production pixel design to 1.1 micron pitch, but only ships that pixel on small dies. In contrast, a full-frame (24 x 36 mm) light field sensor with 1.1 micron pixels would theoretically contain over 700 million rays (and cross the 1 billion mark with the next reduction in pixel size). Collaboration with industry sensor designers has shown me that it makes sense to plan for exponential growth in the near term. On the research side, I believe that this new approach presents a plethora of challenges and opportunities in the design of sensors, and I will seek out opportunities to build deep collaborations with faculty experts in this area.

With exponential growth will come a fast growing challenge transferring, encoding, processing and storing light field data. Representations and algorithms are discussed below. Computing architectures today are increasingly organized around the input of high-bandwidth data streams into high-compute kernels. With a long-term vision of 1 billion light rays per frame at video rates, light field capture and processing represents one of the highest performance use cases that I am aware of for systems research. In my research program, I envision long-term collaborations with systems and architecture faculty to leverage computational photography applications in their research of high throughput computing platforms.

One of the applications that would be enabled by such a platform is correction of aberrations for lens design. Light fields present the opportunity to design entirely new lenses that go beyond conventional limits. Since the very first "computed" lens design by Josef Petzval in 1839, the foremost geometrical constraint in lens design has been control of optical aberrations. The usual approach is physically-based: optimization of the physical lens design before any photograph is taken. In contrast, my thesis work introduced a computationally-based approach to compensate for aberrations *after* the photograph is taken [2,1]. Simulation on existing lens designs showed effective resolution could in some cases be improved by over an order of magnitude (and the improvement grows with light field resolution).

In this way, light field photography fundamentally enlarges the space of lens designs that meet a design spec, and therefore opens the door to entirely new lenses that are lower cost and/or higher performance (larger zoom range, more powerful aperture, lower weight and complexity). The company has taken a critical first step in this direction, and has shown me that we are poised for a revolution in the field of lens design. Fundamental research is needed, however, to develop the theoretical framework, build optimization tools, and design some breakthrough lenses. I have, and will continue, collaborations on these fronts, and will continue to evangelize the new opportunities this offers to the optics community.

In summary, the computational photography paradigm demands a rethinking of the entire camera machine, from lens to sensor to computing architecture.

Theoretical Foundations and Signal Processing

To fully exploit light fields as resolutions rise exponentially, there will also be a pressing need for compact representations, compression, and fast algorithms, as well as new theoretical tools to inform tradeoffs and predict performance. Based on my background in applied mathematics, I believe that multi-resolution (wavelets) and Fourier-domain techniques are promising. Before my research on the light field camera, I worked on fast relighting of virtual scenes using nonlinear wavelet approximations. I published two SIGGRAPH papers [6,7] on these topics, and together these papers have been cited over 500 times according to Google Scholar. My research program will seek to apply these kinds of techniques in 4D to encode light fields compactly, and ideally to process them efficiently directly in the compressed domain.

I have also found Fourier-domain approaches powerful and insightful. One of the theoretical contributions of my thesis research was a mathematical analysis of the light field flowing inside the camera body. A major result was the Fourier Slice Photography Theorem, which revealed the elegant mathematical relationship that, in the Fourier domain, 2D photographs lie along 2D hyper-planes of the 4D light field's spectrum. Applying this theorem yielded an asymptotically fast Fourier algorithm for refocusing from light fields. The theorem also enabled closed-form analysis of expected sharpness of refocused images, a prediction of exiting performance that motivated us to build the original prototype camera. The prototype camera work was published as a tech report [5], and the theoretical work was published as a SIGGRAPH 2005 paper [4]. Together, these two papers have been cited over 650 times. I intend to deepen and broaden these kinds of applications of Fourier methods to the theory of light fields, and computational imaging at large.

New theoretical foundations, signal representations and processing algorithms are necessary to inform and power the transformation of digital to computational photography.

Applications

Photography is a ubiquitous tool in the arts and the sciences, and computational photography has the opportunity to revolutionize many fields. The following is a brief survey of areas where I would like to collaborate with domain experts to apply computational photography to their problems.

Light field video has the potential to revolutionize cinematography. Today, focusing in video is impossible without a script and a person to pull the focus. By providing the flexibility to change lens settings in software after the shoot, light field cameras bring new creative freedom and efficiency on set. There are also obvious implications for security, machine vision and other video applications.

Holography represents the ultimate 3D viewing technology, but capturing holograms has historically been difficult. Every light field can be processed into a holographic stereogram with full two-dimensional parallax. Therefore, light field cameras provide the data source for, and motivate the development of, low-cost holographic printers and real-time holographic displays.

Computer graphics and vision researchers will benefit from a rich new set of opportunities with plentiful light field cameras. As one example, it will make sense to pursue inverse rendering algorithms that attempt to derive shape, materials and lights from light fields. Success would provide a framework for broadly applying graphics and vision techniques to our everyday photographs.

Scientific imaging is an area where light field photography has already begun to make an impact. Marc Levoy showed that a light field microscope [3] enables 3D tomographic reconstruction of the specimen from a single photographic exposure, because it records rays that are attenuated through the semi-transparent specimen at a range of viewing angles (similar to a CT scan). Telescopes present another intriguing opportunity - would it be possible to discard the adaptive optics used in ground-based telescopes, instead recording light field video and correcting for atmospheric aberrations in software? In general, every imaging device that uses optics in front of a sensor may benefit from recording and processing light field data according to the principles laid out in my dissertation.

In closing, computational photography systems research is highly multidisciplinary and requires integrated thinking across the entire imaging stack. My work has bridged, and will continue to bridge disciplines, seeking to bring together experts from different fields to collaborate on the challenges and opportunities of this important research frontier. While I have deliberately focused this research statement on my areas of expertise, my education has given me broad interests across the humanities and sciences. Over time I would love to find a way to work on, among other things, neuroscience (particularly bio-electrical integration with, and direct cognitive control of, digital memory and computation), the science of musical instruments (I played the viola), and the philosophical question of what technology progress means for the human experience (I'm human, and a humanist).

1. Digital Light Field Photography, by Ren Ng. Ph.D. Dissertation in Computer Science, Stanford Unviersity, 2006.
2. Digital Correction of Lens Aberrations in Light Field Photography, by Ren Ng and Pat Hanrahan. SPIE Vol 6342 (Proceedings of the International Optical Design Conference), 2006.
3. Light Field Microscopy, by Marc Levoy, Ren Ng, Andrew Adams, Matthew Footer and Mark Horowitz. ACM Transactions on Graphics (Proceedings of SIGGRAPH) 25(3):924-934, 2006.
4. Fourier Slice Photography, by Ren Ng. ACM Transactions on Graphics (Proceedings of SIGGRAPH 2005) 24(3): 735-744, 2005.
5. Light Field Photography with a Hand-Held Plenoptic Camera, by Ren Ng, Marc Levoy, Mathieu Brédif, Gene Duval, Mark Horowitz, Pat Hanrahan. Stanford University Computer Science Tech Report CSTR 2005-02.
6. Triple Product Wavelet Integrals for All-Frequency Relighting, by Ren Ng, Ravi Ramamoorthi and Pat Hanrahan. ACM Transactions on Graphics (Proceedings of SIGGRAPH) 23(3): 475-485, 2004.
7. All-Frequency Shadows Using Non-linear Wavelet Lighting Approximation, by Ren Ng, Ravi Ramamoorthi and Pat Hanrahan. ACM Transactions on Graphics (Proceedings of SIGGRAPH) 22(3): 376-381, 2003.
8. Lytro Light Field Camera, www.lytro.com.
9. Steve Lohr, A Start-Up's Camera Lets You Take Shots First and Focus Later, New York Times, June 21, 2011, p. B1, http://www.nytimes.com/2011/06/22/technology/22camera.html.
10. Walt Mossberg, Radical Camera Lets You Pick What's Blurry And What's Not, Wall Street Journal, March 1, 2012, p. D1, http://allthingsd.com/20120229/radical-camera-lets-you-pick-whats-blurry-and-whats-not/
11. Bloomberg TV, Lytro CEO Ng Describes Camera's Image Technology, http://www.bloomberg.com/video/78288272-lytro-ceo-ng-describes-camera-s-image-technology.html
12. Popular Science Innovation of the Year, Popular Science Magazine, December 2011.
13. Time Magazine 50 Best Inventions, Time Magazine, Nov 28, 2011.
14. 10 Emerging Technologies 2012, Technology Review, June 2012.