I study imaging from the perspective of computer science and optics. I believe this joint approach will become necessary, as the interaction between billions of light rays, millions of pixels, and thousands of materials becomes increasingly high-dimensional and data-intensive. In my research, I develop theories on how to extract information from light, instantiate my designs as prototype cameras, and strengthen latent ties between computer vision, optics, computer graphics, and data science. *Computational imaging* and *computational photography* are widely used terms to describe my research area.

It is an exciting time to work in imaging. The field has rapidly evolved, particularly in the last 5 years. Self-driving cars, VR/AR, and medical imaging are among the industries that will continue to benefit from imaging advances. As one measure of growth, new computational imaging journals, conferences, and workshops have arisen in the last decade.

A: My Approach

In contrast to previous approaches that separately study optics and computing, I study these topics as a cohesive system. There is a logic to my hybrid strategy. Longstanding problems in optics that were previously considered intractable can be addressed with state-of-the-art computation. I develop a theory of *light transport* to model the complex interaction between light and materials in a tractable space. In doing so, I am able to formulate forward and inverse algorithms that can be physically instantiated in *programmable imaging hardware*.

My PhD thesis adopted recent advances in signal processing and data-driven approaches to re-examine optical multipath interference. By intent, I wanted my doctoral work to be focused in computational imaging, and not in applications. With this accrued experience, I will broaden into two application areas:

- Application Area 1: **Next-generation imaging sensors for autonomous driving and systems.**
- Application Area 2: **Next-generation imaging sensors to sense bio-signals.**

I believe that my skillset in optics and algorithms can transform the vector of of autonomous systems and health technologies.

B: Contributions and Impact

My previous contributions appear in venues of **computational photography** (ICCP), **computer graphics** (SIGGRAPH, SIGGRAPH Asia), **computer vision** (CVPR, ICCV, IJCV) and **applied optics** (Optics Letters, COSI). I organize my past preparation into three areas of research:

1. **Multipath Interference**: Interference is a fundamental phenomena in the imaging sciences. When light propagates from a light emitter to a light sensor there are multiple potential light paths a photon may take. Analogous to interference in wireless networks, my research aims to use programmable pixels and illumination sources to combat multipath (when it is undesirable) and exploit multipath (to make the invisible, visible).

2. **Next-generation LIDAR and 3D imaging**: 3D sensors can be found in self-driving cars, VR headsets, and manufacturing facilities. However, traditional 3D systems are fundamentally limited in their range precision. My doctoral research demonstrated some of the highest precision forms of 3D imaging ever introduced in machine vision, by exploiting the rotation of light waves (polarization).

3. **New connections between computational imaging and healthcare**: During my PhD, I started a long-term effort with Harvard Medical School to understand how we can capture and process additional degrees of freedom in visible light to see inside the body without X-rays.

I believe the above choice of past topics has been fruitful. Over 15 U.S. patents have been filed during my time as a PhD student. A few months ago, my research in ICCV appeared in the best papers special issue in IJCV (May 2017). With university faculty and industry collaborators from Microsoft Research, I have co-organized courses and tutorials at SIGGRAPH '14, SIGGRAPH '15, and ICCV '15. While my research has also led to numerous personal awards, I defer this recognition to the prominence and growth of my research area.

Contribution 1: Multipath Interference

Light misbehaves in real-scenes, scattering in a chaotic fashion before reaching the camera sensor. In certain cases, this form of multipath causes deleterious effects. For example, a time of flight camera is not geared to measure scattered light paths from fog. In these cases, we seek to *combat* the effects of multipath interference. In other cases, the scattered light contains scene information. There we seek to *embrace* multipath.

Combating Multipath As a PhD student, my first set of projects studied the separation of light paths using time of flight cameras. These cameras operate by using an opto-electronic emitter (like a laser or LED) that blinks rapidly in a sinusoidal pattern (the Google Tango and Microsoft Kinect are prominent examples). In order to generalize these cameras to settings like fog, transparent objects, or scattering objects it is necessary to consider multipath effects. In SIGGRAPH Asia 2013, we built a customized time of flight camera that reprograms the light to blink in a carefully chosen pattern [7]. In post-processing we can then deconvolve with knowledge of the temporal blur kernel to overcome some of the deleterious effects of multipath. There is, of course, a design exercise to identify the optimal coded sequences in context of multipath (an analysis inspired by multipath wireless channels). When we are successful, we are able to separate a dense assortment of light paths at each pixel and re-arrange them based on their time of arrival. This enables "light-in-flight" imaging, which we show at this URL: (`http://web.media.mit.edu/~achoo/nanoguide/`).

Once we proposed the initial paper on coded time of flight imaging, we were able to explore a number of extensions. In ICCP 2014, we added multiple, programmable lasers which are sensed by a single image sensor [1]. Since interference is more prone to occur with multiple lasers, we extended our coding scheme to trace light paths back from the laser that they originated at. In CVPR 2015 we studied how to combat multipath interference by using both space and time, to capture smooth multipath (as would occur in for subsurface scatterers like soap or wax) [9]. The next year, in CVPR 2016, we connected our suite of multipath interference solutions with analogous work in coherent microscopic imaging. We derived a theoretical proof on how far apart reflections could be in their time delay, and still be separated. Equipped with this holistic understanding, we were able to demonstrate new results, like removing reflections from windows and 3D imaging of concave objects [4].

Embracing Multipath In many cases, multipath light can be an invaluable source of information and should be embraced. I highlight one such example: seeing around corners. Light from an occluded object, by definition, bounces multiple times before reaching the camera. In order to "see" around the corner, it is necessary to extract information from multipath. Within my community, there has been an explosion of interest in this topic. We now know that it is possible to form blurry images of occluded objects around the corner. However, what was not addressed prior to my work was provable bounds on the resolution of occluded targets ("how blurry is the image around the corner?"). We presented a method at SIGGRAPH 2016, that theoretically characterized the blur kernel when looking around the corner [8]. We made some interesting findings. The blur depends both on the precision in the time of flight measurement as well as the bidirectional reflectance distribution function (BRDF) of the line-of-sight reflector. Since the BRDF is often non-Lambertian, we were able to create a tailored algorithm that exploited non-Lambertian BRDFs to see around the corner in real-time, with centimeter localization. Please see the video linked in the caption of Figure 1 for a real-time demo.

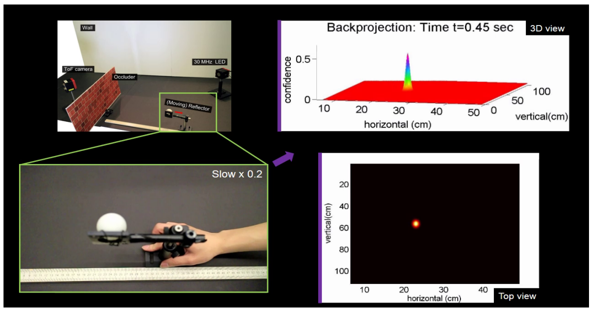

Fig. 1: Real-time object tracking around corners. `www.media.mit.edu/~achoo/occluded/`

Contribution 2: Next-generation LIDAR and 3D imaging.

Multipath is only one of the roadblocks toward realizing robust 3D vision. In follow-up papers, I targeted additional limitations of 3D sensors like range precision, low-power operation, and long-range detection.

High Resolution 3D scanning: In precision manufacturing, bio-prosthetics, and cultural heritage, it is often necessary to voxelize a 3D object at micron resolution. Unfortunately, 3D cameras today do not support this level of precision. Allow me to illustrate with a visual example. A coffee cup was scanned in Figure 3 using three methods. The first two methods are existing technologies: the Microsoft Kinect, and an industrial quality NextEngine 3D multistripe laser scanner. Neither technique recovers the fine geometry in the Styrofoam cup. To solve this problem, we proposed "Polarized 3D" at ICCV in 2015 [5] and IJCV in 2017 [6]. We leverage the change in the polarization of reflected light caused by object slope. This physical principle is not new, but before our work it was an underconstrained problem. We use a coarse depth skeleton to constrain the problem and obtain a high-quality 3D range map, shown in Figure 3c.

High Frequency Imaging addresses Low Power and Long Range Operation: Most types of active 3D sensors are limited to short ranges due to power restrictions. To enable low-power and long-range operation, some of my recent work has repurposed principles from fiberoptic telecommunications to the realm of optical LIDAR. We use a heterodyne scheme (existing time of flight imagers are homodyne) to obtain a 100X improvement in strobing speed and a corresponding increase in precision at low power levels.

Contribution 3: New connections between computational imaging and healthcare

With programmable pixels, illumination, and optics we can attack bio-sensing problems. Since MIT does not have a medical school, I pursue this path primarily with **Professor Rajiv Gupta** and other faculty at Harvard Medical School. I highlight two past findings in context of imaging through skin.

Imaging through skin without X-rays: Imagine the possibility of seeing through human skin without using X-rays. We can study this problem in a restricted setting. Today, doctors detect blood clots using X-ray angiograms. However, near-infrared light (NIR) is a promising alternative since it penetrates deeper through tissue and is absorbed by hemoglobin in the blood. Unfortunately, it has been hard to effectively use NIR for clot detection in the past due to scattering. To descatter images, we projected high-frequency checkerboard illumination patterns on a subject's skin, repeating the procedure at different wavelengths (sampling 5 dimensions: 2D camera pixels, 2D projector pixels, and wavelength) [3]. By fusing the captured images, we were able to obtain a coarse map of the most important blood vessels. We see this technique as a screening and triage tool, which has been tested with medical collaborators in the developing world (e.g. India).

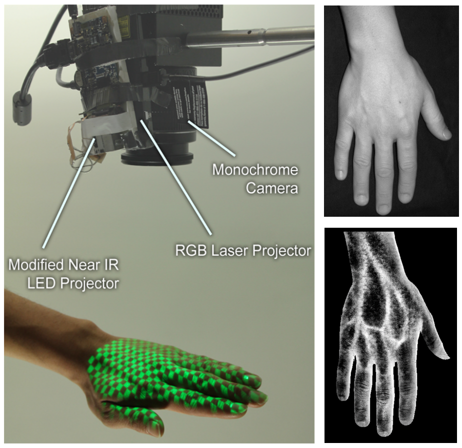

Fig. 2: Seeing through skin sans X-rays.

Low-cost X-ray imaging for Global Health: Although the project in Figure 2 tried to supplant X-rays in a specific scenario, X-rays are irreplaceable for many other procedures (e.g. cancers, tuberculosis, etc). It is a pressing global health challenge that today nearly 70 percent of the world does not have access to X-ray machines. To democratize X-ray imaging, we need to move beyond the specialized and elaborate nature of the X-ray image detector. Along with Professor Rajiv Gupta (Harvard Medical School) and Sr. Research Scientist Richard Lanza (MIT Nuclear Engineering), I explored whether we could repurpose ordinary document scanners for X-ray imaging. We modified the optics of a Canon LIDE 220 scanner to enable X-ray sensitivity and used post-processing computation to denoise the images [2]. The early feedback we are receiving from physicians is that the images we captured are of "diagnostic quality". Based on these pilot results, we have secured funding for further development of scanner-based X-ray imaging.

C: My Future Themes

We have only started probing the theoretical and practical possibilities that programmable pixels, lights, and optics will enable. I will generally orient my future contributions to the two application areas of future interest (Section A). However, my future research program will draw from three primary themes.

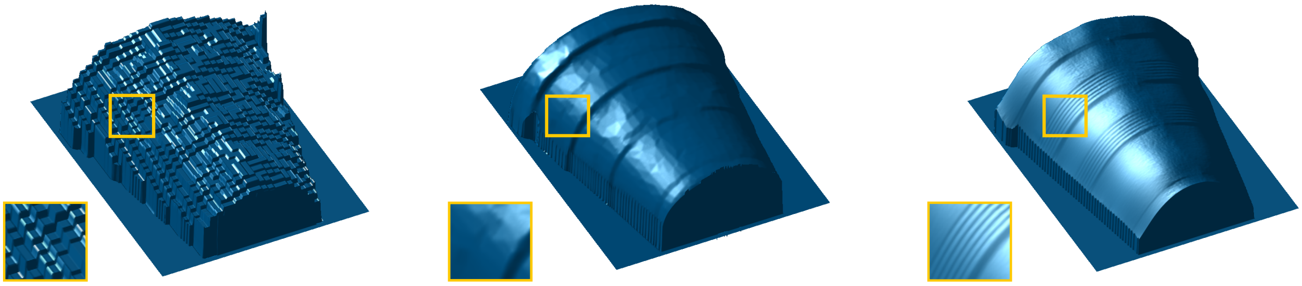

(a) Depth from Microsoft Kinect **(b)** Depth from Laser Scanner **(c)** Polarized 3D (Proposed)

Fig. 3: Instantiating better 3D LIDAR by studying the additional dimension of light polarization.

Strengthening the ties between optics and data science: In my past work, I looked at light transport primarily along the dimensions of time and polarization (4D ray modulation). However, the physics of light is far more complex. With advances in optical capture systems and data storage, one can analyze light transport in an even broader realm that includes wave interference, diffraction, hyperspectral image stacks, holography, light fields, and much more. Efficient numerical algorithms, data structures, and machine learning methods will play as big of a role as fundamental physics and optics in probing these light signals.

Connecting light transport to wireless and telecommunications: In many ways, wireless sensing and computational imaging are two sides of a coin. Both fields take a full-stack approach to jointly capture and process light. It is the six orders of magnitude difference in the frequency of the carrier signal (GHz for wireless, PHz for computational imaging) that divides the two fields. My own research in time of flight cameras incorporates ideas from wireless sensing, like multipath interference and orthogonal frequency-division multiplexing. In such fashion, light transport theory will increasingly draw from aspects from wireless sensing, and vice versa.

Tighter theoretical bounds on imaging: Ordinarily, imaging tends toward an applied science. While it is not uncommon to observe theoretical bounds in imaging literature, it is somewhat uncommon to observe progressive tightening and re-examination of theoretical bounds. Nevertheless, I believe that theoretical bounds will be an active topic in computational imaging in coming years. With the introduction of fundamentally new imaging systems (like seeing around corners), it is more critical than ever to prove when methods will work.

A. Kadambi, A. Bhandari, R. Whyte, A. Dorrington, and R. Raskar. Demultiplexing illumination via low cost sensing and nanosecond coding. In *IEEE ICCP*, 2014.

A. Kadambi, A. Cramer, R. Lanza, R. Raskar, and R. Gupta. Using document scanners for x-ray imaging. In *IEEE ISBI**, 2018 (*Under Review).

A. Kadambi, H. Ikoma, X. Lin, G. Wetzstein, and R. Raskar. Subsurface enhancement through sparse representations of multispectral direct/global decomposition. In *OSA COSI*, 2013.

A. Kadambi, J. Schiel, and R. Raskar. Macroscopic interferometry: Rethinking depth estimation with frequency-domain time-of-flight. In *IEEE CVPR*, 2016 (3 percent acceptance rate).

A. Kadambi, V. Taamazyan, B. Shi, and R. Raskar. Polarized 3d: High-quality depth sensing with polarization cues. In *IEEE ICCV*, 2015 (3 percent acceptance rate).

A. Kadambi, V. Taamazyan, B. Shi, and R. Raskar. Depth sensing using geometrically constrained polarization normals. *IJCV*, 2017 (Best Papers of ICCV Special Issue).

A. Kadambi, R. Whyte, A. Bhandari, L. Streeter, C. Barsi, A. Dorrington, and R. Raskar. Coded time of flight cameras: sparse deconvolution to address multipath interference and recover time profiles. *ACM Transactions on Graphics (TOG)*, 2013.

A. Kadambi, H. Zhao, B. Shi, and R. Raskar. Occluded imaging with time-of-flight sensors. *ACM Transactions on Graphics (TOG)*, 2016.

N. Naik, A. Kadambi, C. Rhemann, S. Izadi, R. Raskar, and S. B. Kang. A light transport model for mitigating multipath interference in tof sensors. In *IEEE CVPR*, 2015.