

SIDDHARTH KHULLAR

Research Statement

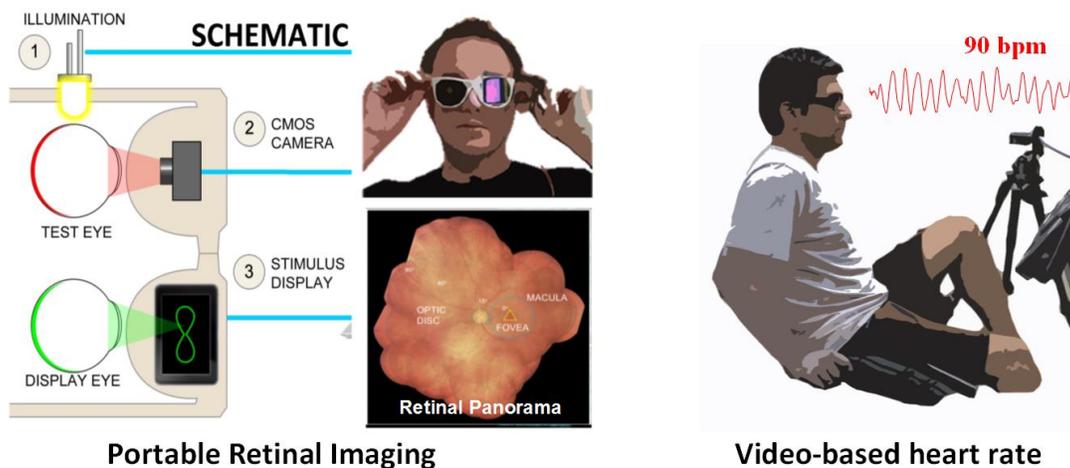
Post-doctoral Researcher
Microsoft Research, Redmond, WA, USA

Ph.D., Imaging Science
Rochester Institute of Technology, Rochester, NY, USA

My research focuses on **developing readily-available sensing systems and algorithms for personalized wellness**. The projects described here draw on my core skill set that includes computer vision, sensor fusion, optics, signal processing, and rapid prototyping. First, I describe a **portable retinal imaging** device developed at the MIT Media Lab [1], followed by my post-doctoral work at Microsoft Research on **unobtrusive heart rate sensing** [2]. Next, I introduce my doctoral work at RIT and the Mind Research Network on **preprocessing and enhancement algorithms for functional brain imaging** [3, 4].

IMAGING SYSTEMS FOR WELLNESS MONITORING

Many essential health screening devices – such as ECG machines, ophthalmoscopes, and ultrasound scanners – are used routinely in clinical settings. However, because these tools are expensive and require significant expertise, they are typically used only in the context of clinical visits where patients are already presenting with symptoms. In other words, screening devices are not being employed as *preventive* tools due to lack of access outside of the clinic. In my research, I have challenged current norms and designs (e.g. expensive and bulky medical equipment) to **explore and create wellness monitoring devices that are easy to use, providing new opportunities for preventive screening**.



Retinal Imaging (MIT Media Lab)

Retinal images are a powerful, non-invasive method for early detection, diagnosis, and treatment of diseases such as diabetes and hypertension. Unfortunately, current retinal diagnostics require expensive equipment, and specific expertise, making access, portability, and regular examinations challenging. To solve this, I worked with a team of researchers at the MIT Media Lab to develop a portable retinal imaging device called *PRISM* – Portable Retinal Imaging for Self-Monitoring.

PRISM illuminates the retina from the temporal side of the eye in order to exploit the transparent nature of retinal tissues as well as translucent properties of the skin. A white LED is placed near the eye, but out of the user’s line of sight to avoid any visual strain from the light. This enables light to pass through the periphery of the eye and illuminate the back of the retina (retro-illumination) without dilating the pupil. In other words, the back of the eye now behaves as a projector, which alleviates the need for complex focusing optics. Imaging is realized through a simple micro-camera (CMOS CCD) setup placed in front of the user’s eye that captures focused images of the retro-illuminated retina. One challenge we had to overcome was maintaining a focused view of the retina as the combined focal length of the camera and the ocular lens changed (the ocular lens accommodated behind the cornea). This was tackled as a bounded-limits problem; we estimated

the upper and lower possible focal lengths along with the depth of field of the camera, which allowed us to derive an optimal range of focal lengths. Together, a low-power illumination technique and simple imaging optics enable a fully functional, low-cost, easy-to-use retinal imaging system. In the future, proliferation of similar light-weight solutions may help people gain wider access to preventive screening technologies.

This project demonstrated that there are open opportunities to apply my knowledge and skills to rapidly advance sensing systems for non-intrusively monitoring human wellness.

Heart Rate from Videos (Microsoft Research)

Most heart diseases (e.g. hypertension) show no visible symptoms, and patients suffering from such diseases are often unaware of their condition. Heart rate is an important indicator of cardiovascular health, making it essential to track along other indicators (e.g. BP, weight, cholesterol etc.) on a regular basis. Optical photoplethysmography (PPG) and electrocardiography (ECG) have been the two primary approaches to obtain pulse waveforms in a clinical setting. Unfortunately, the expertise required to make these measurements, along with the hassle of instrumenting our bodies (e.g. hooking electrodes to chest or fingers), have made such techniques somewhat impractical outside of a clinic.

Today, virtually all computing devices are either already equipped (e.g. phones, tablets etc.) or can be easily accessorized (e.g. desktop/console cameras) with at least one camera pointed toward the user's face. Recent techniques have demonstrated that by analyzing intensity changes due to blood flow in the face, reasonable estimates of pulse rate can be obtained [2]. Unfortunately, these approaches tend to be sensitive to (1) user motion due to exercise or postural changes, and (2) dynamic or poor lighting conditions. Such limitations render this modality much less useful outside of controlled laboratory settings.

I have developed a video-based heart rate monitoring algorithm that can overcome these limitations and enable continuous HR monitoring, even during daily activities such as desk work or exercise. This algorithm aimed to robustly separate noise (lighting changes and user motion) from signal (small changes in the face, orders of magnitude smaller than noise, due to pulsatile blood flow). To achieve this, my algorithm employs a combination of data-driven techniques such as motion filtering, blind signal separation (independent component analysis – ICA), feature extraction, and classification. Motion filtering and ICA initiate pulse extraction from a person's face, captured using a RGB video camera. Machine learning then enables automatic identification of the correct waveform amongst a large group of "pulse-like" candidate waveforms. Features extracted from each candidate signal, specifically autocorrelation, kurtosis, and periodicity, are strong low-level indicators of shape of the waveforms. A classifier trained on similar data is applied to choose the most suitable candidate and report the heart rate. This approach is shown to robustly handle pragmatic limitations posed by user motion (exercise etc.), varying lighting, and varying skin tones [2]. This enables the utilization of heart rate in non-clinical scenarios as an indicator of human engagement (gaming, watching digital content etc.), behavior (affect and emotion), and long-term wellness (heart rate statistics over days or longer).

These two projects exemplify ways to combine general-purpose cameras with the appropriate algorithms to conceive new forms of wellness tracking devices. I believe that ***creating such unconventional, easy-to-use technologies may facilitate preventive health tracking.***

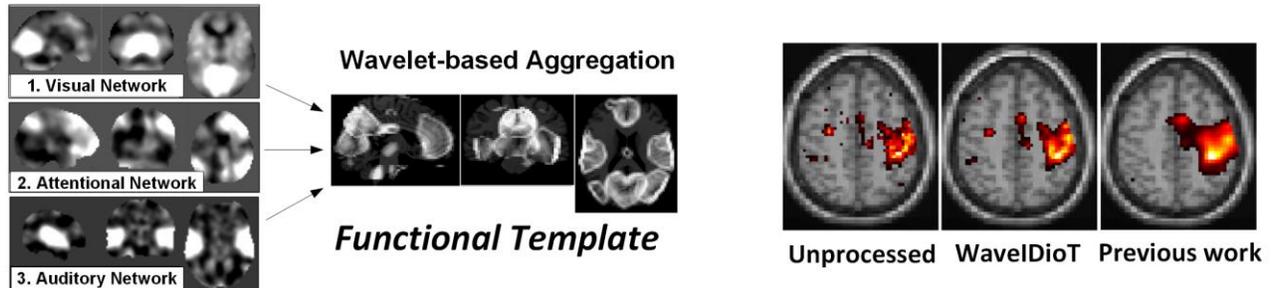
DATA-DRIVEN SIGNAL PROCESSING FOR fMRI

Functional MRI (fMRI) images represent blood flow in the brain: a powerful tool for exploring brain function. In most fMRI experiments, researchers want to address two important questions associated with brain function:

- 1) How does functional activity map spatially onto the corresponding anatomical brain regions?***
- 2) How consistent (or different) is this spatial mapping across a sample population?***

In my doctoral work, I developed signal processing methods for functional brain imaging by applying two concepts – Wavelets and Independent Component Analysis (ICA). By extending these, I formulated new methods for denoising volumetric images, normalizing multi-subject fMRI (functional magnetic resonance imaging), and synthesizing spatial templates of brain networks. Preprocessing steps in fMRI analysis play a pivotal role in influencing almost every decision in

designing, performing, and analyzing results from an fMRI experiment. Related concerns such as inter-subject variability, function-to-structure mapping, template correspondence, and poorly modeled imaging noise have been significant open challenges since the advent of fMRI. The most widely adopted preprocessing solutions largely fail to address these issues, and are chosen by neuroimaging researchers based on their availability, familiarity, and ease of use (e.g. SPM or FreeSurfer [3,4]). As part of my doctoral work, ***I developed a novel image pre-processing framework that focuses on data-driven approaches and reduces model dependence for fMRI analysis.***



Specifically, I created **ICA-based Functional Normalization (ICA-fNORM)** [4], a new framework for spatial normalization of fMRI data from a group of subjects that utilizes their resting state functional maps (*default* state of the brain) – rather than the structure – as the reference. I achieve this by weaving together two distinct algorithms that benefit each other: ICA *unmixes* the complex fMRI data into multiple components or networks, and wavelets adaptively combine these networks into a single aggregate *functional template*. This “intelligent” signal processing approach combines information from multiple brain states and forms a stronger foundation for decoding interaction patterns in patients affected by neurological disorders such as Schizophrenia and Alzheimer’s. As a near-future possibility, research labs will be able to generate scanner-specific, study-specific, and population-specific templates of brain function. This could help overcome sources of noise from a single hardware system used repeatedly and reveal new information about population-based brain function.

I also developed a noise-removal technique using Wavelet transforms and Bayesian signal estimation theory that optimizes fMRI activation specificity and shape accuracy. This technique is available as a software toolbox – *WaveIDioT* [5] – and been applied to a variety of fMRI datasets (> 500 downloads). The algorithm is spatially adaptive and utilizes inter-scale and intra-scale dependencies in fMRI images without sacrificing resolution or distorting intensity values. This denoising scheme preserves spatial details (edges, corners, etc.) and 3-D morphology of intensity images while maintaining the homogeneity of the original signal values across an fMRI image sequence. This method has been used as favorable alternative to existing denoising methods and improve specificity to accurately decipher bio-markers for brain disorders such as schizophrenia [3].

Through my dissertation work, I acquired a deep understanding of signal processing and machine learning theory, in addition to an ability to organize, analyze, and visualize large data sets. Through this work, I demonstrated successful applications of established signal processing concepts (such as Wavelets and ICA) and vastly encourage seeking algorithmic innovation in pursuit to better understand functional neuroimaging datasets. .

FUTURE AGENDA

I have collaborated with researchers across a variety of fields – perceptual image quality, optics, computer vision, eye tracking, neuroimaging, human computer interaction and wearable medical sensing. These collaborations, combined with the range of projects I have completed so far, have motivated me to explore and innovate in broad domains for the future.

More specifically, I would like to further explore the different forms of imaging (visible, near IR, and short-wave IR) and wearable sensors for characterizing bio-signal quantities (e.g. HR, BP, oxygenation, glucose etc.). Possibilities in this space are exciting and will continue to become prevalent with large leaps in fabrication and computing technology. For example, hyper-spectral imaging (new optics and smaller sensors) may enable rapid evaluation of skin pigmentation (e.g. melanomas), oxygenation (e.g. respiratory disorders), and hemoglobin density (e.g. anemia) by simply looking at the skin, or even produce rapid representations of plaque and tartar within the mouth (e.g. high-contrast under blue/UV light).

Proliferation of inertial sensors embedded in mobile and wearable devices will increasingly call upon innovation in the signal processing and machine learning communities. As an intern at Microsoft, I experimented with inertial data for complex types of activity monitoring and realized there is a great potential in utilizing these sensors to enable interesting new ways of interaction by developing algorithms tailored to specific user scenarios. Finally, sensor fusion and machine learning for analyzing data from multiple sources (cameras, accelerometers, gyroscopes, and magnetometers) can be largely helpful in advancing wellness tracking applications and also enabling new forms of ubiquitous computing. I will to explore and conduct more of this research in the future.

The emergence of additive manufacturing, non-expert 3D printing, and sensor miniaturization may impact the future landscape of personal digital devices. I seek opportunities to explore the fast-growing, cloud-connected digital ecosystem known as the *Internet of Things*. There are rich opportunities for innovation in sensing systems which may have broader and sustained impact (e.g.: occupancy sensing, smart appliances etc.).

I aim to simplify personal sensing systems to a point where the user can be placed within the loop. Identifying entry points in our daily routines is rapidly becoming a design principle. A system that improves itself by leveraging human intelligence and habits has the potential to become a leading example for improving the overall user experience. This likely involves innovating upon increasingly ubiquitous mobile devices as sensing and computation platforms, but perhaps a bigger challenge comes from making data acquisition seamless and unobtrusive. Overall, integrative approaches where sensor functionalities are tightly coupled with algorithmic innovation will increase usefulness of existing sensors and establish new paradigms in human machine interaction, and wellness computing.

SELECTED PUBLICATIONS

1. M.E. Lawson, **S. Khullar**, et al., “[Computational Retinal Imaging via Binocular Coupling and Indirect Illumination](#)”, *Late breaking technical talk/poster, ACM SIGGRAPH 2012*.
2. **S. Khullar**, N. Joshi, T. S. Saponas, D. Morris, D.S. Tan, “[Making Pulse Measurement from Video More Practical](#)”, *Submitted to ACM CHI 2014*.
3. **S. Khullar**, A.M. Michael, N. Correa, T. Adali, S. Baum, V.D. Calhoun, “[Wavelet-based fMRI analysis: 3-D denoising, signal separation and validation metrics](#)”, *Neuroimage, vol. 54, no. 4, Oct. 2010*.
4. **S. Khullar**, A.M. Michael, N.D. Cahill, K.A. Kiehl, G.D. Pearlson, S. Baum, V.D. Calhoun, “[ICA-fNORM: Spatial Normalization of fMRI data using Intrinsic Group ICA-networks](#)”, *Frontiers in Sys. Neurosci. vol. 5, Nov. 2011*.
5. **S. Khullar**, V.D. Calhoun, “[WaveIDiot: Wavelet-based Image Denoising Toolbox for fMRI data.](#)” 2011 URL: <http://mialab.mrn.org/software/waveidiot/index.html>.
6. Center for Disease Control and Prevention (2012). URL: <http://www.cdc.gov/features/vitalsigns/hypertension/>.



CONTACT

Siddharth Khullar

sidk@microsoft.com

For more information, please visit: <http://www.siddharthkhullar.net/>