Vikram Ravindra

Research Statement

Summary

My areas of interest are broadly in machine learning, and data science techniques, particularly in the context of biomedical imaging. In the course of my research work, I have explored the areas of network science, randomized numerical linear algebra, and optimization. My recent work focuses on the analysis of biomedical images and connectomics – a field dedicated to the understanding of the human brain [1, 2, 3, 4, 5, 6, 7]. My research goal is to develop and apply novel techniques in computer science and statistics to find solutions to practical problems, such as characterization and diagnosis of neurodegenerative diseases and the development of innovative Human-Computer Interfaces that exploit direct access to neurological signals. Beyond these areas, I have also contributed to diverse applications, including transcriptomics [8], healthcare [9], methods [10, 11] and robotics [12, 13].

Current Research

My current research aims to develop novel methods and statistical models for human brain connectomes derived from functional MRIs (fMRI). Connectomes characterize regions of the brain that are co-active during specific tasks, as reflected in the time-series fMRI data. Analyses of connectomes has a number of important applications. I provide some examples from my recent work:

Modeling response to audio-visual stimulus: Functional MRI studies are typically focused on understanding of neuronal responses to specific stimulus. The characterization of such responses gives us a population-wide response to a given sensory input (or stimulus). An accurate representation of responses to individual elements that constitute a generic (and complicated) stimulus gives us the means to estimate the expected response for a new stimulus. I explore these questions using an fMRI dataset from a cohort of subjects watching BBC's Sherlock. Our aim is to find a cohort-wide stable representation of neuronal response that characterizes high-level aspects of the input. Such a representation would capture various features of the inputs, such as background information in the video stimulus, or presence/ absence of music. It would also capture contextual features such as the level of excitement or boredom, and anticipation or fear, which require background information (such as the story line).

I formulate this problem in the framework of Archetypal Analysis. I express the response at

each point in time as a convex combination of "archetypal" responses to the pure (elemental) inputs. I show that the archetypes themselves provide accurate representations of neuronal responses across a population [4]. Moreover, I show that the information encoded in the archetypes is strongly correlated with the information encoded by expert annotation of the video (Pearson Correlation > 0.9). Finally, I show that I can predict the response from new subjects with an accuracy in excess of 93%. This work was the first study that relates video features (expert annotations) with neuronal responses in a rigorous algorithmic and statistical framework based on deconvolution of observed mixed imaging signals.

Brain signatures: Population-wide functional brain studies often focus on signals that are shared amongst a cohort, and those that discriminate across different cohorts. This approach, while being useful in certain contexts, results in decimating the differences between individual brains in a cohort. This loss of information may have substantial effect on downstream applications, and recommendations from such studies. An important component of functional brain networks is individual specific signals, or *signatures*. The differences between individuals enables personalized healthcare by providing targeted treatment on the basis of signatures. In our work [1, 3], we are interested in locating regions in the brain that best code for the signatures. I use a *matrix sampling* [1] approach to identify a small set of features (edges in the functional brain graph) that best describe the identity of a subject. I show that an intelligent selection of ~ 100 features from over 64k features can predict the identity of a subject (at rest) with an accuracy of 94%. The (sub)-regions that I identify can be used to accurately sample cells for technologies such as single-cell transcriptomics to find genes that are differentially expressed in these regions. This guides development of **personalized medicine** to cure, or manage various neurodegenerative disorders. I also use these features to recover associated regions that cover less than 5% of the cortical surface. In [2], I use a *Network Alignment* approach to align pairs of unregistered functional brain networks. My formulation of network alignment takes advantage of the fact that the physical distance between nodes in a graph is roughly fixed (i.e., the graph is rigid). I show that an alternating approach that minimizes the error for network alignment and structural alignment efficiently distinguishes between pairs of networks belonging to the same subject and pairs of networks belonging to different subjects. The advantage of my approach is that distinct anatomical features of each subject are highlighted, and not suppressed. This affords downstream fMRI analyses to derive insights about differences within a cohort, which could lead to the *discovery of new (possibly rare) sub-types of neurodegenerative diseases*.

Functional brain signatures, while important to the understanding of behavior and disease can also be a privacy violation for the human subjects who participate in projects that release neuroimaging datasets for public use. In [2], I demonstrate a **de-anonymization attack** to deanonymize the identities (name, DOB, SSN, contact details) of such individuals. Specifically, I show that signatures of patients in private repositories can be used to identify subjects in public datasets (where subjects are initially de-identified). This is a real, rapidly evolving problem compounded by increasing resolution in images (both temporal and spatial) and the number of publicly available repositories. My work is currently cited by European agencies, as they revamp data release standards for neuroimages.

Research Goals

My career goal is to make fundamental contributions to different application domains through novel data science/ ML techniques.

Short Term Research Goals In the short term, I plan to work on a number of challenging problems that impact the fields of data science, functional connectomics and neuroinformatics. (i) **Integrating meso-scale connectomes and transcriptomes.** The rapid advances in technologies that map brains in the meso-scale and those of single-cell transcriptomics have reduced the gap in the resolutions at which the two fields usually operate. As I have worked with both types of data, I am ideally placed to integrate insights of the two modes of data. An example where such an analysis would yield success is in the study of brain lesions. The connectome reveals structural details of the cell, which serves as phenotype for related (dis-)regulation in gene expression of specific regions that are adversely affected by disease. (ii) **Explainable brain activity.** Complex, continuous inputs have not been characterized in terms of the responses they elicit from different neurons and different regions in the brain. Most current efforts focus on simple stimulus (e.g., horizontal bars excite a particular kind of neurons, which are not excited when exposed to vertical bars). My work in [4] is the first step in expressing response in terms of explainable factors. I believe that further studies would require a combined effort in algorithms and in hardware. Novel deconvolution methods, coupled with different imaging modalities would help explain observed responses. As the inputs are realistic, the findings of this study help us predict behaviour of healthy and diseased persons when exposed to new inputs. (iii) **Privacy-Preserving functional brain studies.** Building on my work in [2], which shows privacy related vulnerabilities in neuroimaging datasets, I plan on building systems that manage biomedical data in a secure manner. Differential privacy based methods are promising and will promote work in data security and data management along with image processing and computer vision. I look forward to collaborating with interested colleagues within the department as well across departments on these topics.

The success of these projects require development of methods in core data science and machine learning, and suitable application to specific contexts. These problems have sufficient visibility in high impact journals such as Nature, Science and Cell, as well as machine learning venues such as NeurIPS and ICML. To achieve sufficient freedom in the types of projects I work on, I plan on collaborating across departments to build infrastructure required to collect data (imaging hardware), as well as a lab that develops novel prototypes to record activity.

Long Term Research Goals The big questions that I aim to tackle with my research focus on Human-Computer Interfaces that are controlled by, and provide feedback to the brain. This is motivated by the observation that information conveyed via sensory inputs (visual, auditory, olfactory, etc) are of substantially higher bandwidth than the output conveyed to the world by an individual (speech, movement of motor muscle etc). The most important factor that exacerbates this asymmetry of information is the limitation posed by our natural end effectors (i.e., mouth, limbs). In cases where complex physical operations need to be accomplished in real-time, interfacing the brain directly to hardware is more efficient. Examples of such situations

include control of fighter jets, robot arms by surgeons, etc. With rapid improvements in imaging and sensory hardware, this goal is feasible over the next decade. My experience with different modalities in imaging and biosensors, along with an appreciation for a wide range of techniques in data science gives me unique perspective that would let me make considerable advances in this field.

Publications

V. Ravindra, P. Drineas, and A. Grama, "Constructing compact brain connectomes for individual fingerprinting," *Frontiers in Neuroscience*, 2021.

V. Ravindra and A. Grama, "De-anonymization attacks on neuroimaging datasets," in *ACM Conference for Management of Data (SIGMOD '21)*, 2021.

V. Ravindra, H. Nassar, D. F. Gleich, and A. Grama, "Rigid graph alignment," in *International Conference on Complex Networks and Their Applications*, pp. 621–632, Springer, 2019.

V. Ravindra and A. Grama, "Characterizing similarity of visual stimulus from associated neuronal response," in *IJCAI 2020*.

V. Ravindra, G. Sanders, and A. Grama, "Identifying coherent subgraphs in dynamic brain networks," in *ICIP 2021*.

V. Ravindra, H. Nassar, D. F. Gleich, and A. Grama, "Aligning spatially constrained graphs," *In Review, IEEE Transactions on Knowledge and Data Engineering*.

V. Ravindra, C.-H. Fang, and A. Grama, "May i see what you see? predicting visual stimulus from functional brain images," in *Review*.

S. Mohammadi, V. Ravindra, D. F. Gleich, and A. Grama, "A geometric approach to characterize the functional identity of single cells," *Nature communications*, vol. 9, no. 1, pp. 1–10, 2018.

C.-H. Fang, V. Ravindra, S. Akhter, M. Adibuzzaman, P. Griffin, S. Subramaniam, and A. Grama, "Identifying and analyzing sepsis states: A retrospective study on patients with sepsis in icus," *In Review*, 2021.

T. A. Runkler and V. Ravindra, "Fuzzy graph clustering based on non-euclidean relational fuzzy c-means," in *2015 Conference of the International Fuzzy Systems Association and the European Society for Fuzzy Logic and Technology (IFSA-EUSFLAT-15)*, pp. 91–97, Atlantis Press, 2015.

C.-H. Fang, V. Ravindra, A. Grama, J. Davila-Velderrain, and S. Mohammadi, "Supervised archetypal analysis," in *Review*.

V. Ravindra and C. Castellini, "A comparative analysis of three non-invasive human-machine interfaces for the disabled," *Frontiers in neurorobotics*, vol. 8, p. 24, 2014.

C. Castellini and V. Ravindra, "A wearable low-cost device based upon force-sensing resistors to detect single-finger forces," in *5th IEEE RAS/EMBS International Conference on Biomedical Robotics and Biomechatronics*, pp. 199–203, IEEE, 2014.