Academic Research Statement

"In theory there is no difference between theory and practice, but in practice there is" - Jan L.A. Van de Snepscheut

Experience and Achievements

My research broadly spans machine learning (ML), deep learning (DL) and signal processing for applications in computer vision, healthcare, graph analysis and scientific machine learning. I have focused on both theoretical and algorithmic aspects of machine learning methods for challenging real-world problems. With a strong publication record (Google Scholar) and multiple patents, I have successfully contributed to a number of application areas. Furthermore, I have secured multiple research grants from different agencies (DOE, Advanced Scientific Computing Research program, NNSA, ARPA-E), and my research findings have transitioned into federal business opportunities. As a recognized ML expert, I have been a panelist in several DOE workshops and visioning meetings, an invited participant in prestigious forums around the globe (STS Young Leaders program, Dagstuhl, Banff International Research Station) and a distinguished speaker at multiple reserach organizations. As a result of my contributions to the field, I received the LLNL Early Career award in 2020 (only awardee in AI/ML among all scientific disciplines).

Data-centric learning has been a core research focus right from my Ph.D., where I worked on sparse representations and dictionary learning. By analyzing the intriguing connection between dictionary learning and unsupervised clustering problems, I developed a *stable* learning algorithm (with respect to the perturbations in training data) for inferring dictionaries. Interestingly, our approach adopted a hierarchical learning strategy for its implementation, similar to modern deep networks. This work has led to exciting applications in speech source separation, inverse imaging problems, synthetic aperture radar imaging and medical image analysis. With the emergence of deep neural networks and novel representation learning paradigms, my research has broadened to these topics and I have made several important contributions:

• **Core ML**: An important direction of my research has been to understand and improve generalization of ML models "in the wild", *i.e.*, make predictions under covariate shifts, concept shifts, unknown corruptions, label distribution shifts and semantic novelty. This encompasses both predictive models trained from scratch and those that are built upon pre-trained representations, e.g., self-supervised learning. In this context, I have explored a variety of formulations, including unsupervised domain adaptation, zero-shot domain generalization, transfer learning, adversarial robustness, and fully test-time adaptation. Furthermore, motivated by science and engineering applications which use experiment designs to bootstrap ML pipelines, our team developed a novel statistical mechanics framework for studying ML generalization under the lens of sample design, and made foundational advances in space-filling sampling strategies with guaranteed generalization. Finally, through my works in learning with structured data, I have designed algorithmic solutions for multi-layer graph modeling, representation learning and improving robustness to adversarial corruptions. My efforts in Core ML have been supported by multiple research grants from DOE and DNN, and has led to several publications at premier venues including Neurips, ICML, CVPR, AAAI, ICCV, JMLR, Nature, PNAS and TNNLS. (highlight);

• **Uncertainty quantification in deep learning**: The intricate interactions between data sampling, model selection and the inherent randomness in complex systems strongly emphasize the need for a rigorous characterization of ML algorithms. Hence, I am interested in characterizing prediction uncertainties in deep models to enable safe ML practice – identifying out-of-distribution (OOD) samples, delegating high-risk predictions to experts, defending against adversarial attacks and incorporating real-world priors. I carried out this work as part of a research grant from the Advanced Scientific Computing Research program (2 projects selected among $100+$ submissions). Key milestones from this research include scalable techniques for characterizing different uncertainty sources, advanced black-box optimization with neural network surrogates and state-of-the-art semantic novelty/covariate shift detectors. In addition to multiple publications, our work was featured on editor's highlights from the prestigious Nature Communications journal (highlight);

• **Inverse problems**: Solving ill-posed inverse problems is a central challenge in science and engineering, and it hinges on effective prior specification. In my research, I have extensively explored the design of deep image priors for different data modalities and the use of generative models as data priors (autoencoders, GANs, diffusion models). In addition to developing novel optimization strategies, for the first time, we studied the utility of generative priors for solving inversion problems on OOD data (examples: MimicGAN (IJCV 2020) and SPHInX (ICML 2022)). With applications ranging from inverse imaging, CT reconstruction, history matching and audio restoration, my solutions have had significant practical impact (highlight);

• **Explainability**: This has emerged as a crucial aspect of building trust in machine learning systems and enabling practitioners to interact with complex neural networks that are otherwise opaque. Given that explainability is a key design objective in several high-impact applications, I have developed a number of key frameworks to - (i) introspect how a model works (TreeView, MARGIN); (ii) support counterfactual reasoning (DISC, TraCE); (iii) explain distribution shifts; and (iv) reliably explain predictions for distribution-shifted data (ProFILE). Furthermore, I have collaborated with visualization researchers to design novel interactive tools for analyzing high-dimensional data and corresponding ML models (highlight);

• **AI4Sciences**: Across scientific disciplines, researchers commonly design and evaluate experiments by comparing empirical observations with simulated predictions from numerical models. Simulations can provide insights into the underlying phenomena and are often instrumental to effective experiment design. In practice, these simulations can be evaluated at multiple fidelities, contain measurements from diverse modalities and are often accompanied with a variety of physical constraints and known relationships. Along with the team at LLNL, I pioneered the area of cognitive simulations to effectively leverage AI in enabling simulation-driven scientific discovery. Our novel advances to scientific representation learning, surrogate modeling, multi-fidelity optimization and transfer learning have led to key milestones in high-energy physics (highlight) and traumatic brain injury characterization (highlight);

• **AI4Health**: Finally, a research area that I have been personally passionate about is in applying

AI technologies for clinical diagnosis and patient data modeling. Through fruitful collaborations with clinical experts, I have made several fundamental contributions to handle crucial challenges posed by clinical problems – weak supervision, data biases, disparate data modalities and need for reliable evaluation. My contributions have been widely recognized through high-impact publications, best paper awards (KDD Data Science for Healthcare 2019 workshop, SPIE Medical Imaging 2021), patents, news articles (highlight) and external media (podcast);

Research Vision

While the current advances in data-driven learning and reasoning have enabled unprecedented capabilities, a critical path forward is to develop the next generation of prescriptive AI systems – that can integrate decision support, effectively leverage external knowledge and enable human-in-the-loop inferencing. With the overarching goal of building high-impact AI solutions for challenging real-world problems, I want to work on ML theory and algorithms pertinent to the following thrust areas:

• **Towards Reliable, Robust and Safe Models**: Understanding the behavior of learning algorithms under different families of distribution shifts (often specified using structural causal models) has been an ongoing topic of research. While a number of recent empirical findings are shedding light into the generalization in real-world settings, there is a need for new theoretical frameworks to both express different uncertainty sources into the learning process (Bayesian perspective) as well as determine performance bounds on OOD generalization and ML safety metrics (calibration, consistency, outlier rejection, adversarial robustness) simultaneously. Furthermore, coupling the properties of pre-trained representations with downstream learning tasks to automatically handle biases in the target dataset is an important problem;

• **Knowledge-Driven AI**: Inspired by the success in computer vision and NLP, there is increased interest in creating foundation models to create baseline representations for arbitrary downstream tasks. With the emergence of multimodal models (e.g., Imagen, Dall-E or Parti) and general purpose encoders (e.g., CLIP or StyleGAN mappers), there is a strong need for a new generation of knowledge-integrated AI solutions. In many real-world applications, a wide-range of knowledge encodings are becoming available (e.g., knowledge graphs, non-differentiable simulators, family of data generation processes, specification of physical models based on PDEs) and there is a need for unified learning frameworks to effectively leverage those knowledge sources. This will require advances in optimization techniques (e.g., zeroth order optimization) and heterogeneous knowledge distillation. In addition to producing physically grounded models (improved generalization), this new class of solutions can help integrate decision support into ML pipelines. Furthermore, using this framework, we can revisit the fundamental problems of few-shot learning and test-time adaptation;

• **Designing Closed Loop Systems**: Over the last decade, we have witnessed significant progress in deep learning-powered sequential optimization and model-based reinforcement learning. Going forward, I am excited about new opportunities in realizing powerful, human-in-loop systems, that can effectively guide humans to make locally optimal decisions using only partial observations, but also leverage feedback from the "world" to refine the decision support systems (e.g., continually adapt or even create knowledge encodings). This naturally requires interfaces for models to communicate uncertainties, multitude of solutions with different short- and long-term rewards and trust-promoting explanations. Furthermore, I also want to explore new approaches for conditional exploration of data manifolds and inverse optimization, which can be integral components of such a closed-loop system.

As a practitioner, I strongly believe that democratizing technology and scientific advances is imperative for positively impacting the society at large. Consequently, I have been a strong proponent of "open" research, and will continue to actively pursue collaborations and community-wide engagement to produce the most effective AI solutions. Finally, and most importantly, I want to leverage my experience to contribute to the broader vision of the AI fraternity within and beyond my organization.

Outreach and Service

As part of my outreach, I routinely collaborate with AI researchers from academia and industry, physical/life science researchers, clinical experts, postdoctoral researchers/PhD students that I have mentored (20+ students around the globe), and practitioners advocating for fair adoption of technologies in the society. Furthermore, in order to facilitate community-wide conversations about the field and its growth, I have worked with other researchers to organize workshops at prestigious forums such as Dagstuhl, and also been part of several peer-review committees in AI/ML conferences/journals and funding agencies. I have also been part of several student mentoring programs (AAAI doctoral consortium, NSF MSGI program, DOE CSGF program).

Furthermore, I strongly believe in the need to improve diversity and inclusion in AI research and work environments. Hence, I have engaged in mentoring women researchers and doctoral/postdoctoral scholars from minority groups, as well as supporting their participation in top AI/ML conferences for better exposure (e.g., co-authoring papers, proof-reading their grant applications, providing them with career counseling, helping with proposal building for early career researchers). In order to be part of a positive societal impact, I have contributed to open source initiatives (MLCommons, DOE Applied Math Visioning Committee, DOE AI4EarthSciences committee) and to AI proposals for social projects with non-profit organizations (Rand Corp). During the COVID-19 pandemic, I was part of the LLNL team that worked closely with CDC and MITRE to leverage AI for epidemiology modeling and policy optimization. In recognition of this effort, I received a Gold award from the Weapons and Complex Integration program at LLNL. Overall, I have passionately engaged with researchers and practitioners to foster an active research community.