Towards Reliable Learning in the Wild: Generalization and Adaptation

Huaxiu Yao

Department of Computer Science & School of Data Science and Society
UNC-Chapel Hill

The past decade has witnessed a remarkable success of machine learning, particularly in settings where engineers frame a task to a specific function, then collect a large amount of human-labeled, independent and identically distributed (i.i.d.) examples to fit this function, and finally deploy the fitted function to make predictions on new examples within the same distribution. While this standardized pipeline has proven to be immensely beneficial for fields where data labeling is relatively inexpensive and the training and deployment steps are under the same environment (e.g., classical image classification), it represents just a fraction of the possible real-world scenarios.

To deploy ML systems in broader realistic scenarios (e.g., medical imaging, drug discovery), I assert that such systems should relax the i.i.d. assumption and not focus on only one particular environment or task. Instead, we should design machine learning systems that are capable of acquiring commonly shared knowledge from versatile source environments and tasks, while being generalizable and adaptable to new environments and new tasks. Such capabilities will allow ML systems to be more seamlessly deployed to a wide range of settings. For example, to apply ML in healthcare, an entire engineering team is dedicated to deploying and monitoring the ML systems: they need to regularly refresh the models, mitigate racial bias between subpopulations, and adapt the models to a new environment (e.g., a new hospital) with limited patient records. My research vision focuses on automating this process by developing machine learning systems that are reliable and widely generalizable to changing environments and tasks with limited supervision. To develop such systems, I made the following contributions.

- **Mitigating Spurious Correlation and Enhancing Out-of-Distribution Robustness.** For ML systems to perform reliably under distribution shifts, it's crucial to address the spurious correlations present in the training data. To this end, I developed a series of methods to (1) directly capture environment-agnostic information through selective augmentation, and (2) posthoc rectify inaccurately predicted results. In addition to these methodological contributions, we introduced the Wild-Time benchmark, designed to promote the advancement of models that can adapt to natural temporal distribution shifts.
- **Effective and Compositional Adaptation with Limited Supervision.** In addition to enhancing out-of-distribution robustness, gathering a Small amount of target data for adaptation offers an alternative approach to tackle changing environments. To ensure both effective and efficient adaptation, I introduced several methods that (1) facilitate compositional generalization and adaptation by extracting relationships from historical observations or by integrating relationships derived from meta-data; (2) bolster reliable adaptation even in the face of imperfect observations; and (3) enable precise adaptation through selective fine-tuning.
- **Interdisciplinary Applications.** Multiple of the aforementioned contributions have sound theoretical guarantees and provide guidance for building ML systems that are more reliable to changing environments. To further make such ML models more practical in reality, I collaborated with interdisciplinary stakeholders in computational biology, medicine, natural language processing, computer science education, transportation, and hydrogeochemistry. We leveraged domain knowledge to reduce the dependency on a substantial amounts of labeled data and to make our models more robust to changing distributions, leading to data science applications in healthcare and medicine, smart cities, and e-commerce.

Future Plans. The objective of my future research is to design the next generation of machine learning systems capable of continuously and reliably adapting to complex environments and tasks. Building upon our previous results, there are three key areas and next steps essential for successfully achieving this goal: strengthening the power of pre-training and fine-tuning, achieving unbiased and fair learning, and performing continual learning under changing environments. By making machine learning systems well-generalizable and adaptable, I believe that these systems can be deployed everywhere with strong reliability. Such a wide spectrum AI deployments further require expertise and collaboration in many interdisciplinary fields.