Overview:

- Distribution Shift in the Real World:
- In Sequential Decision-Making
  - Non-stationary Environment
  - Close-Loop Interaction

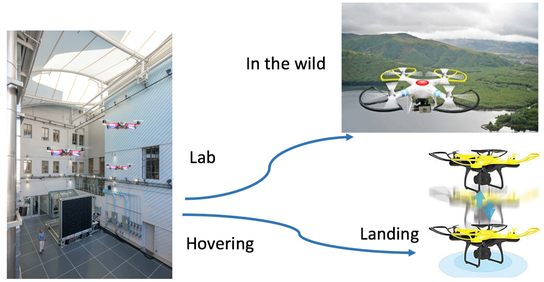

- Exploration as Covariate Shift
  - P(y|x) is the same
  - P(x) is different

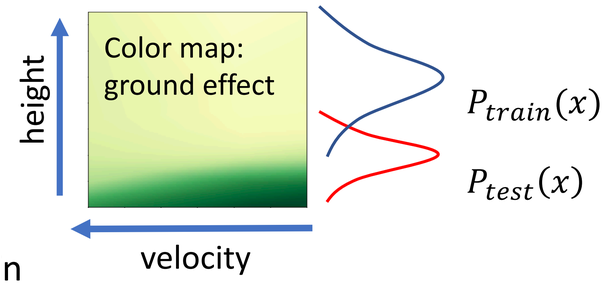

- Goal: Safe learning and control to achieve automatic data collection in real world systems.

Distributionally Robust Extrapolation:

- Learning residual dynamics model: $s_{t+1} = \underbrace{f(s_t, a_t)}_{\text{Nominal physics model}} + \underbrace{\hat{f}_a(s_t, a_t)}_{\text{Regression Model}}$
- Distributionally robust learning framework with relative loss under covariate shift

$$\operatorname{argmin}_{\hat{P}(y|\boldsymbol{x})} \max_{\check{P}(y|\boldsymbol{x}) \in \tilde{\Phi}_{train}} E_{P_{test}(x)\check{P}(y|\boldsymbol{x})}\left[-\log \frac{\hat{P}(Y|X)}{P_0(Y|X)}\right]$$

- Feature constraints enforced by training data, for example:

$$\left| E_{P_{train}(\boldsymbol{x})\check{P}(y|\boldsymbol{x})}[\phi(X,Y)] - \frac{1}{n}\sum \phi(\boldsymbol{x}_i, y_i) \right| < \epsilon$$

- Exponential family as solution: $\hat{P}(y|x) \propto P_0(y|x) \exp[\frac{P_{train}(x)}{P_{test}(x)} \theta \cdot \phi(x,y)]$
- In deep learning case, use the top layer as $\phi(x,y)$

Integration to Control Systems:

- Learning error translates to control tracking error:
- $\lim_{t\to\infty} ||\tilde{x}(t)|| = \gamma \cdot \epsilon$
- Controller: Contraction Mapping Controller
- Safety:
  - Hard requirements
  - Chance constraints
- Pool of trajectories can be replaced by principled planning algorithms.

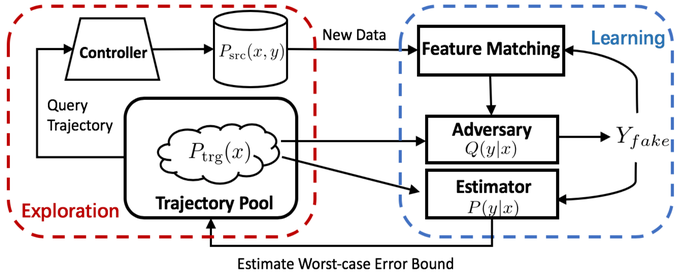

Safe exploration: only collect data that is safe under the worst-case model!

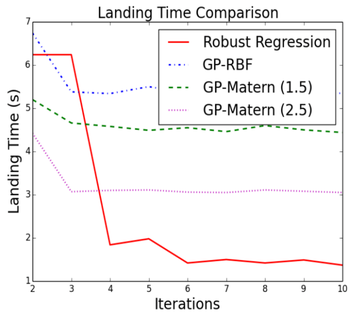

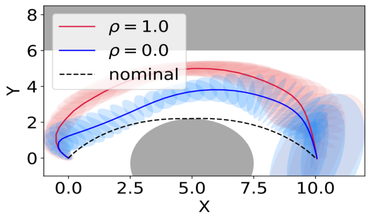

ρ controls how "exploratory" the trajectory is.

My Future Research Agenda:

Develop principled machine learning algorithms to facilitate their integration to the real-world decision-making systems.

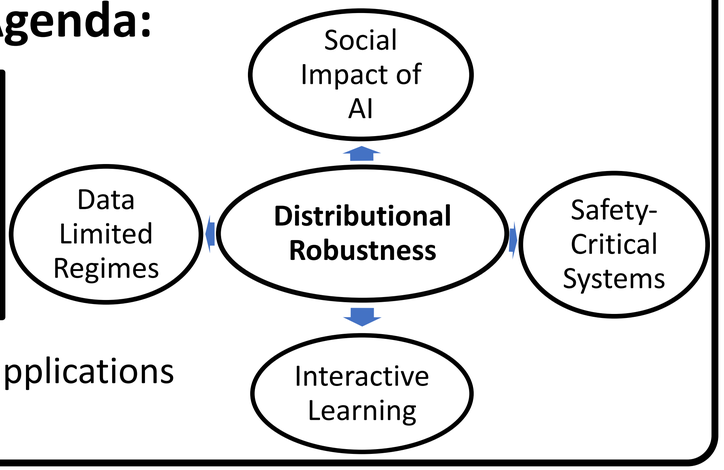

- Bridge the theory and the applications
- Work with domain experts