Integrating Models and Data for Robust Manipulation with and around people

Siddhartha Srinivasa
Boeing Endowed Professor
Personal Robotics Lab
Computer Science and Engineering
University of Washington


Carnegie Mellon University
HERB
Carnegie Mellon University

Manipulation

Personal Robotics Lab

Carnegie Mellon University

Physics-based Manipulation

DARPA ARM-S

CMU 2010-13

Harness the Mechanics of Manipulation to Funnel Uncertainty

Quasi-Static Pushing

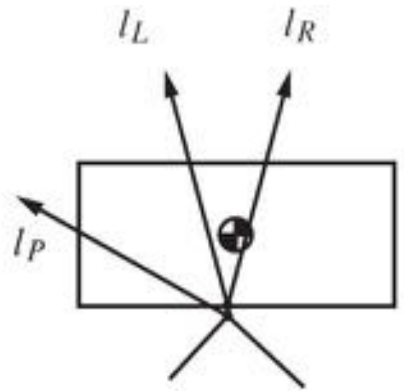

The Voting Theorem
[Mason'81]

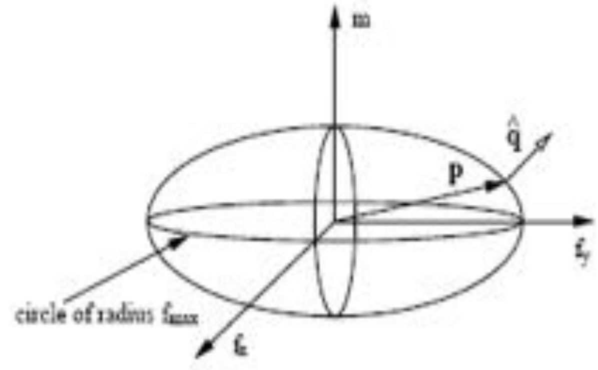

The Limit Surface
[Goyal et al.'91, Howe and Cutkosky'96]

How much should the robot know?

- Object mass? No.
- Object-surface friction? No.
- Object pressure distribution? Pick conservatively.
- Finger-object friction? Pick conservatively.

Analytical Capture Regions

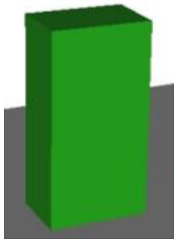

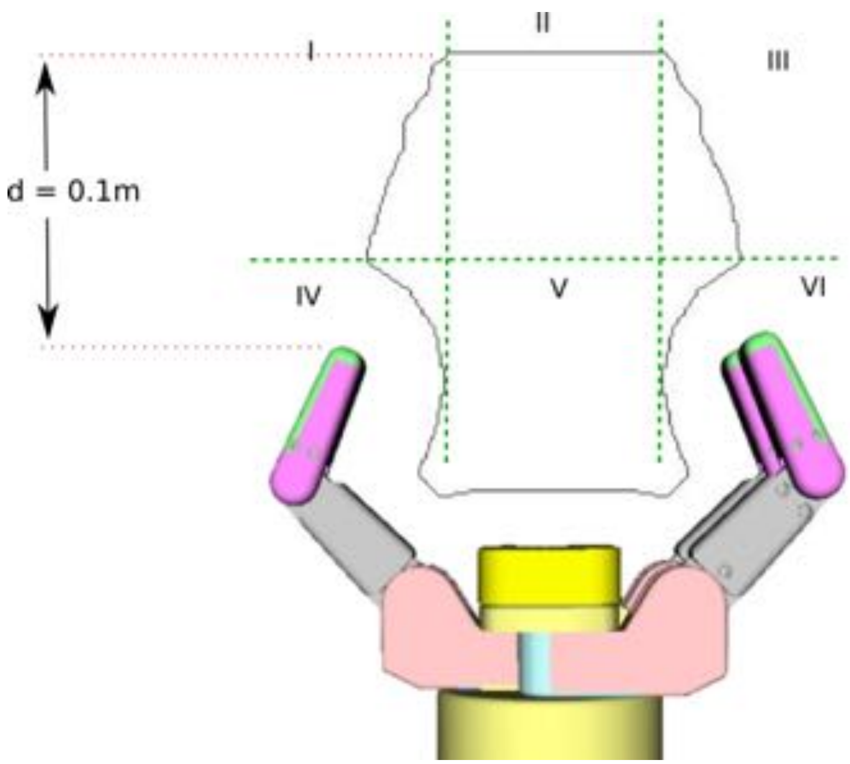

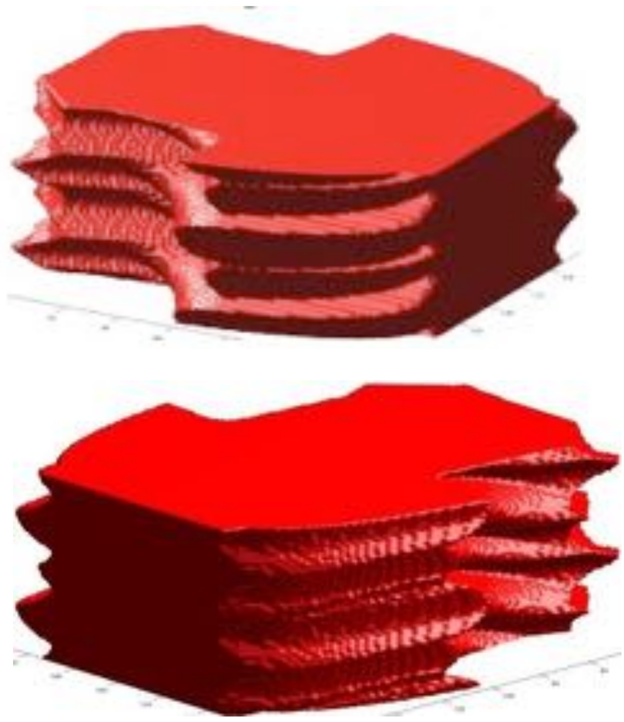

Addressing Object Pose Uncertainty

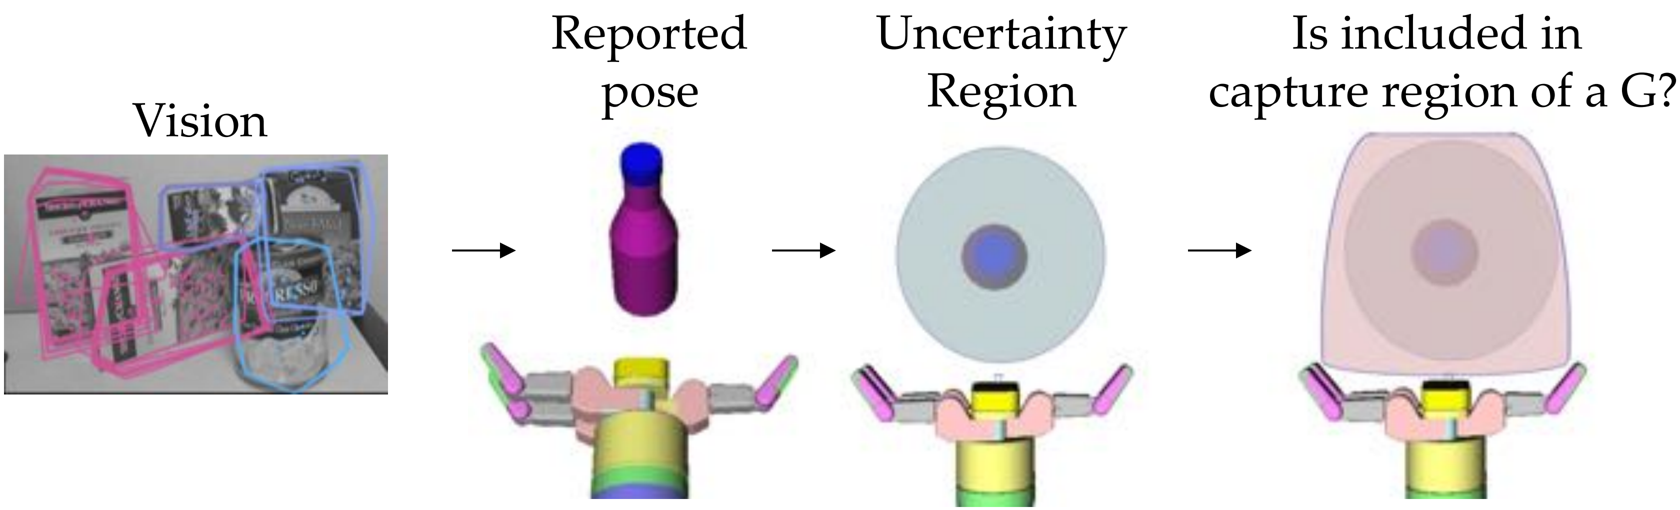

Dogar. “Physics-based Manipulation Planning in Cluttered Human Environments”
Doctoral Dissertation, 2013.

Physics-based Manipulation

Exploiting physics to manipulate objects

Dogar, Srinivasa “A Planning Framework for Non-Prehensile Manipulation under Clutter and Uncertainty”, AuRo 2012.

Autonomous control of complex dynamical systems

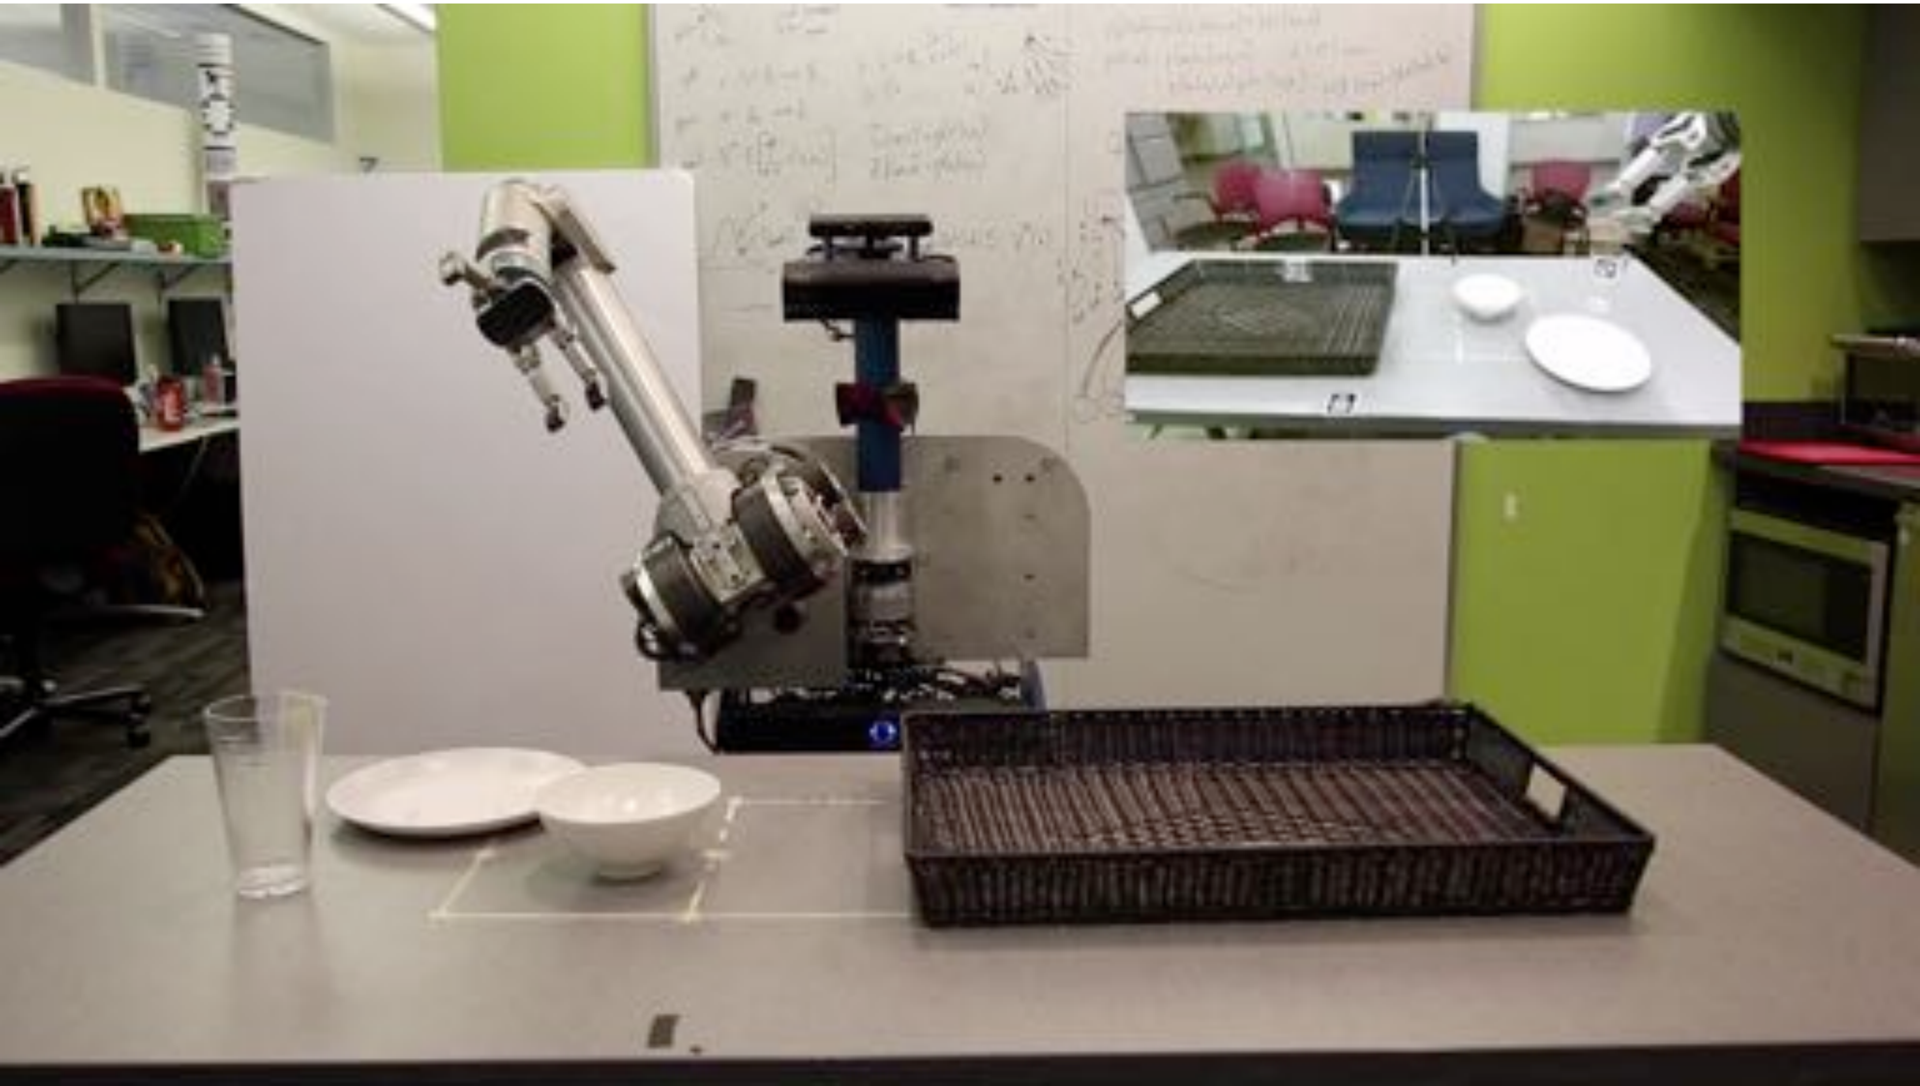

Global models are often
only partially correct

Optimal Control ⟷ Model-based RL

	Global, analytical models		
Model	Inaccurate		
Policy	No uncertainty		
Data	No data collection		
Training	Fast convergence		

Optimal Control ⟷ Model-based RL

	Global, analytical models		Locally learned models
Model	Inaccurate		Locally more accurate, captures uncertainty
Policy	No uncertainty		Uncertainty-aware
Data	No data collection		Requires data collection
Training	Fast convergence		Slow convergence

Optimal Control ⟷ Model-based RL

	Global, analytical models	Hybrid model	Locally learned models
Model	Inaccurate	Globally available, Locally accurate, captures uncertainty	Locally more accurate, captures uncertainty
Policy	No uncertainty	Uncertainty-aware	Uncertainty-aware
Data	No data collection	Use data to correct analytical models	Requires data collection
Training	Fast convergence	Moderate	Slow convergence

Learn the residual between simulation and reality

Gaussian Process Regression

Gilwoo Lee

Optimal Control ⟷ Model-based RL

	Global, analytical models	Hybrid model	Locally learned models
Model	Inaccurate	Globally available, Locally accurate, **captures uncertainty**	Locally more accurate, **captures uncertainty**
Policy	No uncertainty	Uncertainty-aware	Uncertainty-aware
Data	No data collection	Use data to correct analytical models	Requires data collection
Training	Fast convergence	Moderate	Slow convergence

Policy Search:
Iterative Linear Quadratic Regulator

Linear dynamics,

Quadratic cost,

Iterative local improvements

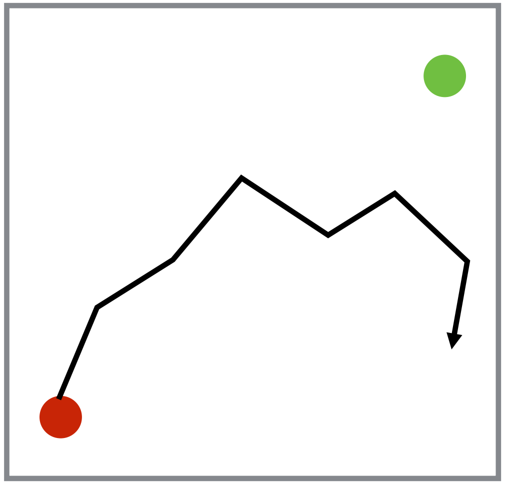

Li, Weiwei, and Emanuel Todorov. "Iterative linear quadratic regulator design for nonlinear biological movement systems." *ICINCO (1)*. 2004.

Policy Search:
Iterative Linear Quadratic Regulator

Linear dynamics,

Quadratic cost,

Iterative local improvements

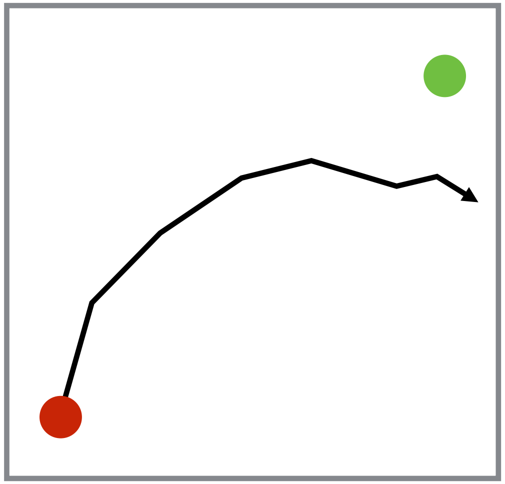

Li, Weiwei, and Emanuel Todorov. "Iterative linear quadratic regulator design for nonlinear biological movement systems." *ICINCO (1)*. 2004.

Policy Search:
Iterative Linear Quadratic Regulator

Linear dynamics,

Quadratic cost,

Iterative local improvements

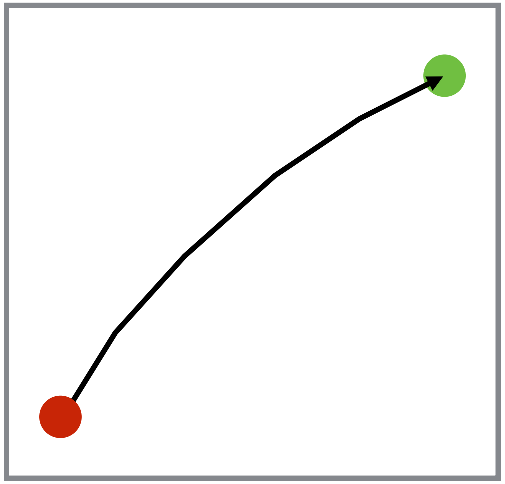

Li, Weiwei, and Emanuel Todorov. "Iterative linear quadratic regulator design for nonlinear biological movement systems." *ICINCO (1)*. 2004.

Policy Search:
Iterative Linear Quadratic Gaussian Control

Linear stochastic dynamics,

Quadratic cost,

Iterative local improvements

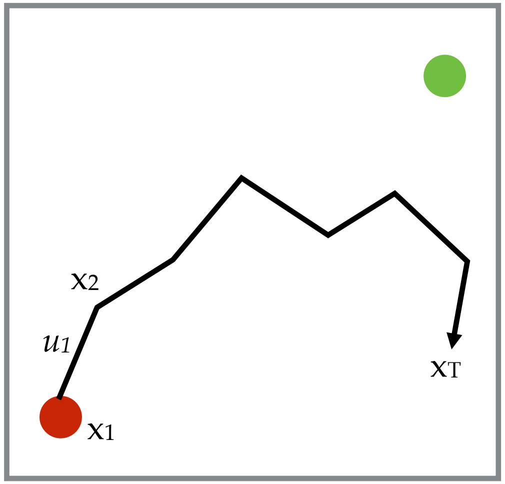

Todorov, Emanuel, and Weiwei Li. "A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems."

American Control Conference, 2005. Proceedings of the 2005. IEEE, 2005.

Incorporate model uncertainty into the Iterative Linear Quadratic Regulator

Robust Iterative Linear Quadratic Regulator

Policy Search:
Iterative Linear Quadratic Gaussian Control

$$x_{t+1} = f(x_t, u_t) + \zeta(x_t, u_t), \quad \zeta \sim N(0, \Gamma)$$

Bellman update:

$$V(x_t) = \min_{u_t} \underbrace{\left[l(x_t, u_t) + E[V'(f(x_t, u_t))] \right]}_{Q}$$

$$Q(\delta x, \delta u) = Q_x \delta x + Q_u \delta u + + 1/2(\delta x^T Q_{xx} \delta x + \delta u^T Q_{uu} \delta u + 2\, \delta x^T Q_{xu} \delta u)$$

$$Q_x = l_x + E[\, f_x^T V'_x \,] = l_x + E[(f + \zeta)_x{}^T V'_x] = l_x + f_x{}^T V'_x$$

$$Q_u = l_u + E[\, f_u^T V'_x \,]$$

$$Q_{xx} = l_{xx} + E[\, f_x^T V'_{xx} f_x \,] = l_{xx} + E[(f+\zeta)_x^T V'_{xx}(f+\zeta)_x \,] = l_{xx} + f_x^T V'_{xx} f_x + E\,[\zeta_x^T V'_{xx} \zeta_x \,]$$

$$Q_{uu} = l_{uu} + E[\, f_u^T V'_{xx} f_u \,]$$

$$Q_{ux} = l_{ux} + E[\, f_u^T V'_{xx} f_x \,]$$

Todorov, Emanuel, and Weiwei Li. "A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems." *American Control Conference, 2005. Proceedings of the 2005*. IEEE, 2005.

Policy Search: Robust ILQG

$$x_{t+1} = f(x_t, u_t) + \zeta(x_t, u_t)$$
$$= f_{global}(x_t, u_t) + \mu(x_t, u_t) + \xi(x_t, u_t) + \zeta(x_t, u_t)$$

$\zeta \sim N(0, \Gamma)$, $\xi \sim N(0, \Sigma)$

Policy Search:
Robust ILQG

$$x_{t+1} = f(x_t, u_t) + \zeta(x_t, u_t)$$
$$= f_{global}(x_t, u_t) + \mu(x_t, u_t) + \xi(x_t, u_t) + \zeta(x_t, u_t) \qquad \zeta \sim N(0, \Gamma),\ \xi \sim N(0, \Sigma)$$

Bellman update:

$$V(x_t) = \min_{u_t} \underbrace{\boxed{l(x_t, u_t) + E[V'(f(x_t, u_t))]}}_{Q}$$

$$Q(\delta x, \delta u) = Q_x \delta x + Q_u \delta u + + 1/2(\delta x^T Q_{xx} \delta x + \delta u^T Q_{uu} \delta u + 2\, \delta x^T Q_{xu} \delta u)$$

$$Q_x = l_x + E[\, f_x^T V'_x\,] = l_x + E[(f_{global} + \mu + \xi + \zeta)_x{}^T V'_x]$$
$$Q_u = l_u + E[\, f_u^T V'_x\,]$$
$$Q_{xx} = l_{xx} + E[\, f_x^T V'_{xx} f_x\,] = l_{xx} + E[(f_{global} + \mu + \xi + \zeta)_x^T V'_{xx}(f_{global} + \mu + \xi + \zeta)_x\,]$$
$$Q_{uu} = l_{uu} + E[\, f_u^T V'_{xx} f_u\,]$$
$$Q_{ux} = l_{ux} + E[\, f_u^T V'_{xx} f_x\,]$$

GP-ILQG

GP-ILQG: Data-driven Robust Optimal Control for Uncertain Nonlinear Dynamical Systems. Lee, Srinivasa, and Mason, 2017

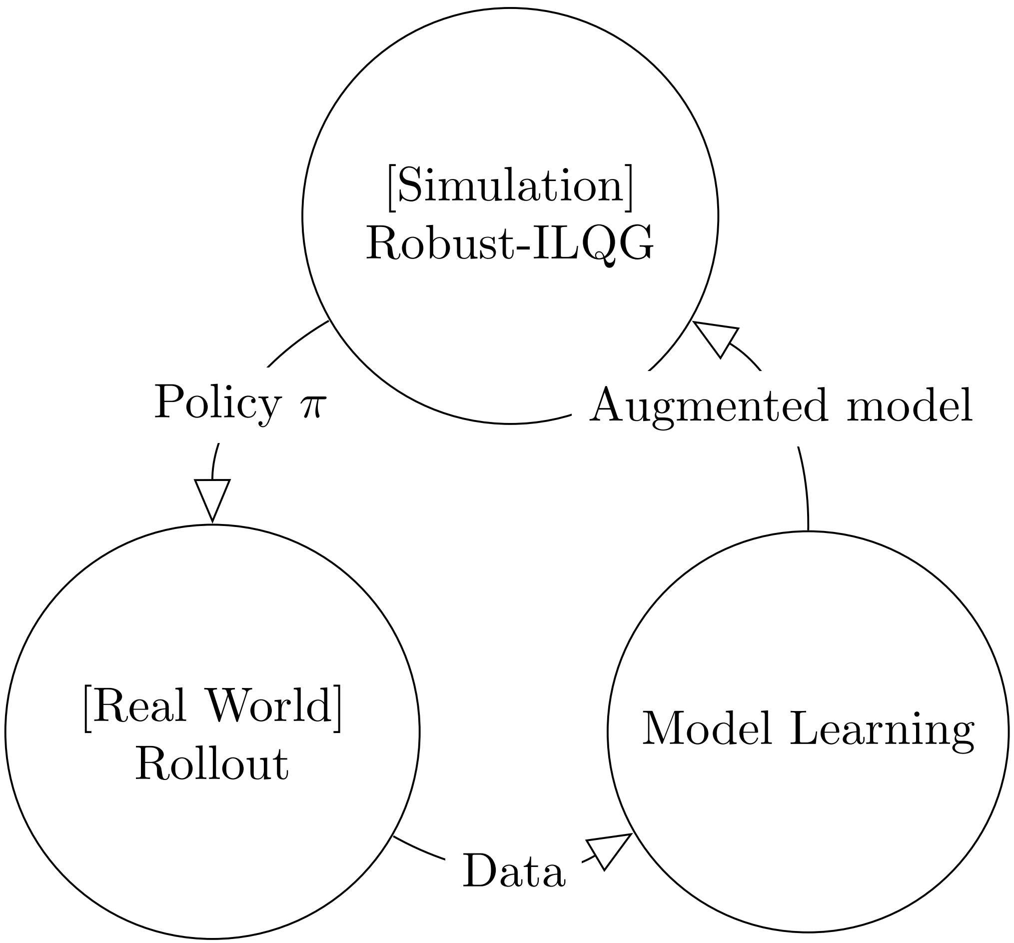

Comparison

Optimal Control ⟷ Model-based RL

	Global, analytical models	Hybrid model	Locally learned models
Algorithm	ILQG	GP-ILQG	Probabilistic-DDP
	Assumes perfect model		**Needs data**
Initializatio n	**Initialized with random policy**	**Initialized with ILQG policy & random trajectories**	**Random policy & Demonstrated trajectories**

Todorov, Emanuel, and Weiwei Li. "A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems." *American Control Conference, 2005. Proceedings of the 2005*. IEEE, 2005.

Pan, Yunpeng, and Evangelos Theodorou. "Probabilistic differential dynamic programming." *Advances in Neural Information Processing Systems*. 2014.

Optimal Control ⟷ Model-based RL

	Global, analytical models	GP-ILQG	Locally learned models
Model	Inaccurate	Globally available, Locally accurate, captures uncertainty	Locally more accurate, captures uncertainty
Policy	No uncertainty	Uncertainty-aware	Uncertainty-aware
Data	No data collection	Use data to correct analytical models	Requires data collection
Training	Fast convergence	Moderate	Slow convergence

Limitations

- Scaling of GPs
  Use any learner that reports uncertainty

- Robustness to model error discourages exploration
  Encourage exploration, e.g. posterior sampling

Takeaways

- Models are nice, but not perfect
  Bootstrap models with data

- Strive for minimalism
  Task-driven model learning

Integrating Models and Data for Robust Manipulation with and around people

Siddhartha Srinivasa
Boeing Endowed Professor
Personal Robotics Lab
Computer Science and Engineering
University of Washington