Optimal Kinodynamic Planning for Compliant Mobile Manipulators

Katie Byl

UC Santa Barbara

On ^ Optimal Kinodynamic Planning for Compliant Mobile Manipulators

Katie Byl

UC Santa Barbara

Overview

- **What** should be **optimized**
  - Robots have stepped outside factory environments
  - Challenges and goals in robotics have changed
- Scalar metrics and their estimation
  - How to judge overall task performance
  - How to **estimate low-level performance**
- Tools for policy improvement
  - On-going challenges in machine learning

Overview

- **What** should be **optimized** : **Reliability**
  - Robots have stepped outside factory environments
  - Challenges and goals in robotics have changed
- Scalar metrics and their estimation
  - How to judge overall task performance
  - How to **estimate low-level performance**
- Tools for policy improvement
  - On-going challenges in machine learning

Kinodynamic Planning

- **Kinematics: geometric** constraints
  - feasible, and without collisions
- **Dynamics: differential** constraints
  - velocity and torque limits
- Traditional **optimization** in robotics?
  - **minimum time** solutions

Canny, Donald, Reif and Xavier (1988-1993)

Kinodynamic Planning

- For **high-DOF robots** (e.g., humanoids):
  - Finding any feasible solutions is challenging
  - ***Typically, such solutions are not optimal***

Example:

- Rapidly-exploring Randomized Trees (RRTs)
  - Sequential process:
    - First, find feasible kinematic trajectories
    - Then, ensure differential constraints (adjust speed)

LaValle and Kuffner (1998-2001)

What are our planning goals?

- (Agree on) simple definitions:
- What is success?
- What are appropriate metrics?
- How to accomplish these optimizations

What are our planning goals?

- (Agree on) simple definitions:

  Optimal = maximizing probability of success

- What is success?

- What are appropriate metrics?

- How to accomplish these optimizations

What are our planning goals?

- (Agree on) simple definitions:

  Optimal = maximizing probability of success

- What is success?

  Varies, to encompasses task-specific concepts

- What are appropriate metrics?

- How to accomplish these optimizations

What are our planning goals?

- (Agree on) simple definitions:

  Optimal = maximizing probability of success

- What is success?

  Varies, to encompasses task-specific concepts

- What are appropriate metrics?

  Again, probability of success

- How to accomplish these optimizations

What are our planning goals?

- (Agree on) simple definitions:

  Optimal = maximizing probability of success

- What is success?

  Varies, to encompasses task-specific concepts

- What are appropriate metrics?

  Again, probability of success

- How to accomplish these optimizations

  Model low-level (Markov) transition probabilities

Why use such a broad goal…

- *Success* definitions vary (across applications):
  - Optimizing probability *can* encompass anything

- Other metrics are indirect / imprecise / vague
  - Example: a geometric “safety margin” is a proxy
  - In stochastic models, exact safety margins emerge

Why use such a broad goal…

- *Success* definitions vary (across applications) :
  - Optimizing probability *can* encompass anything
- Other metrics are indirect / imprecise / vague
  - Example: a geometric “safety margin” is a proxy
  - In stochastic models, exact safety margins emerge
- **Stochasticity and underactuation dominate**

Locomotion ↔ manipulation

Locomotion ↔ manipulation

Shared challenges

- **Mechanical impedance:** $$Z(s) = \frac{F(s)}{V(s)}$$
  - stiff / compliant
  - precise / adaptable

  $\Big\}$ Tradeoffs
- Significant uncertainty and noise
  - "mixing" effects cause perturbations to propagate
- Sensing the environment
  - Visual and tactile; geometry and impedance
- Contact / interactions difficult to model
  - Rolling, sliding, deformations

Motivation

- **Real humans sometimes fail:**

  Trip, fall, drop objects, bump into furniture…

  …but usually recover.

- **Same performance expectations for robots**

  A balance between **risk** and **reward**.

Toy Example

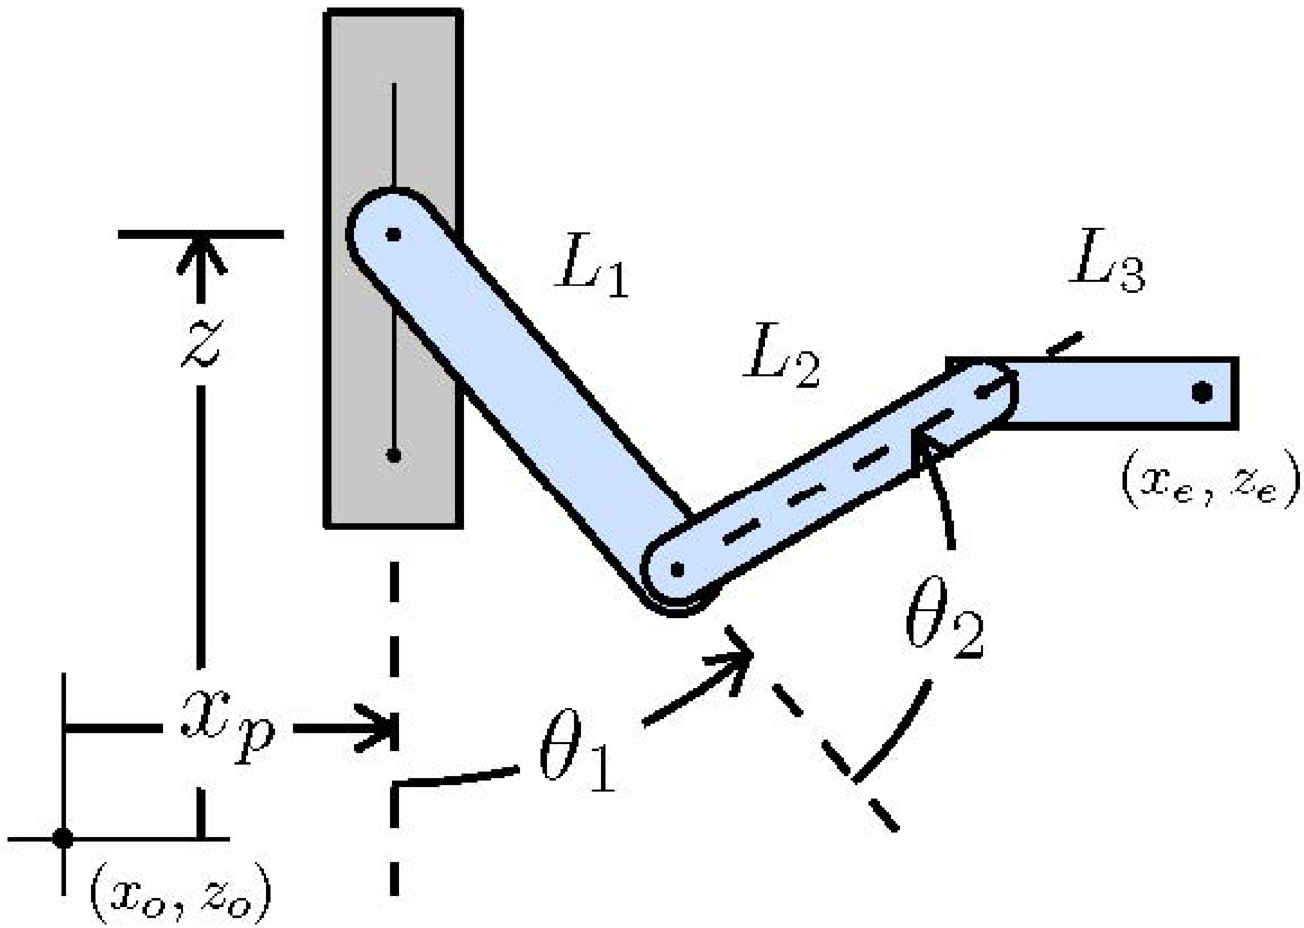

Toy Example

Two tasks:

1. Move and replace vase
2. Lower shelf

Assumed Known (for now):

- Geometry
- Dynamics of shelf: spring-mass-damper (shelf impedance)

Unknown:

- Mass of vase (vase impedance)

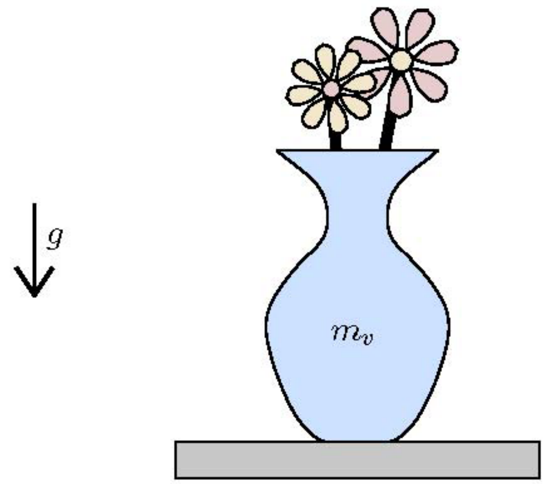

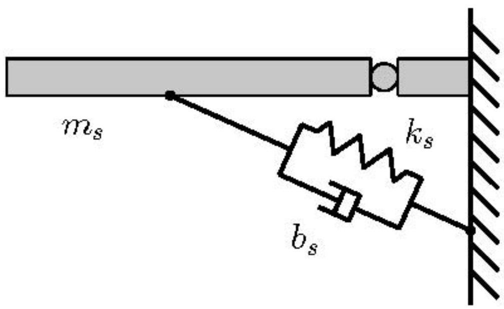

Toy Example

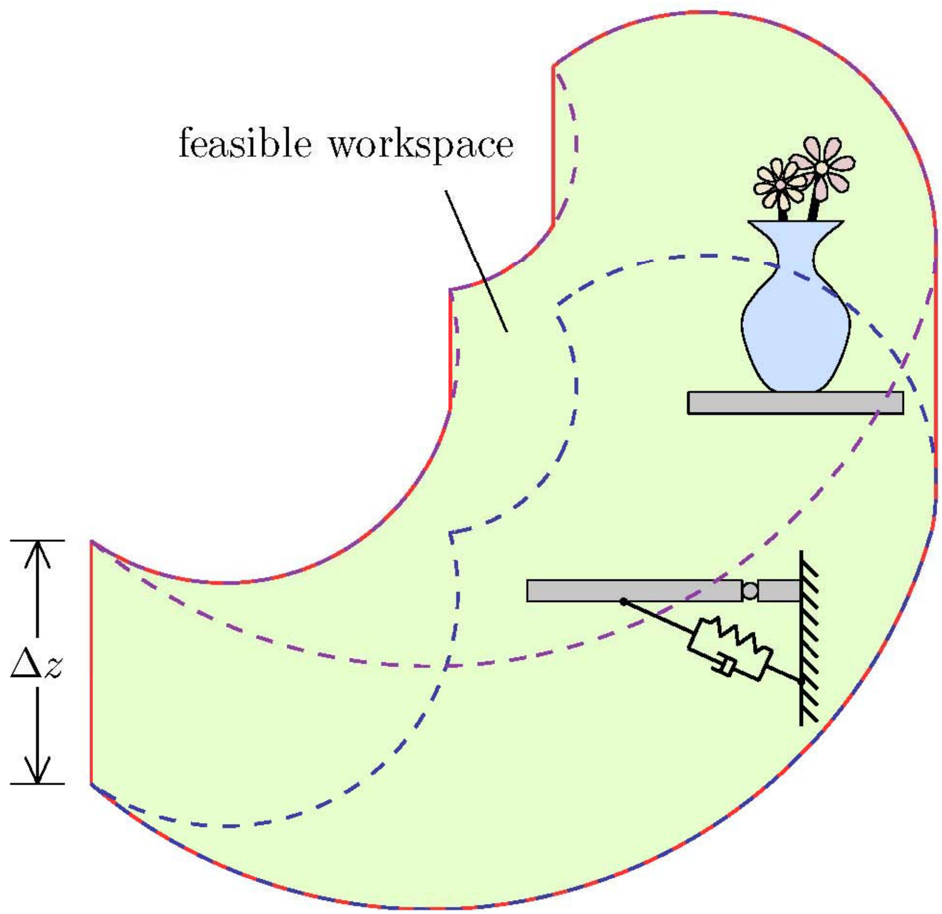

Toy Example

- Planner selects:
  - Initial body pose
  - Joint trajectories
  - Variable torso height

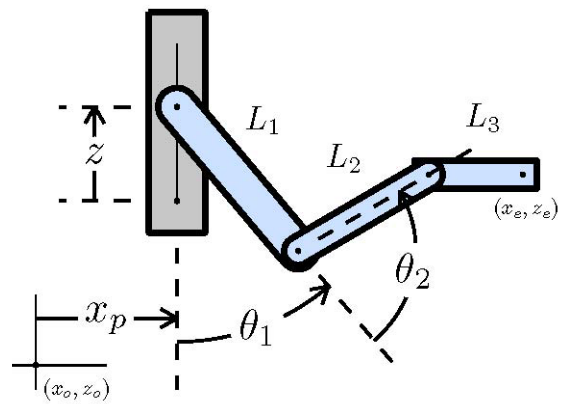

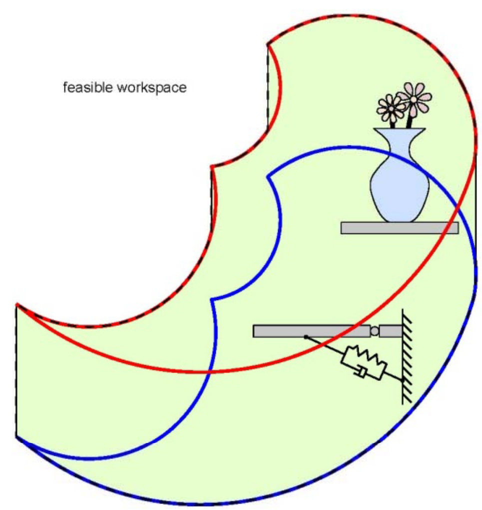

Toy Example

- Planner selects:
  - Initial body pose
  - Joint trajectories
  - Variable torso height
- **No speed requirements**
- **Failure modes:**
  - Insufficient forces
  - Infeasible kinematics

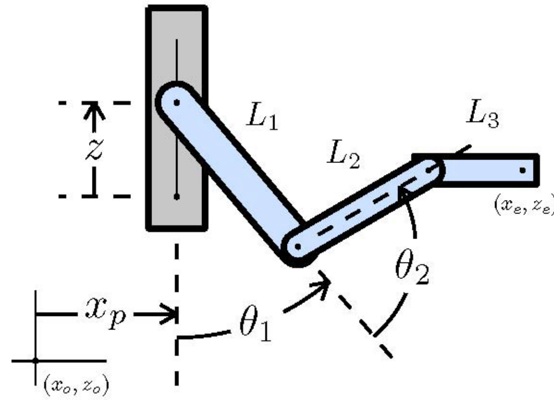

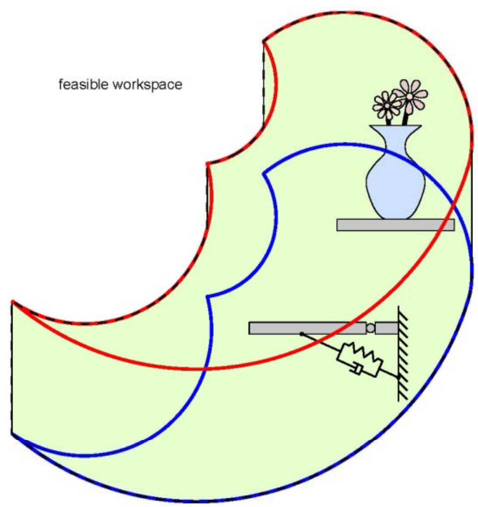

Toy Example

- Model variability in mass of vase
  - Probability density function (PDF)

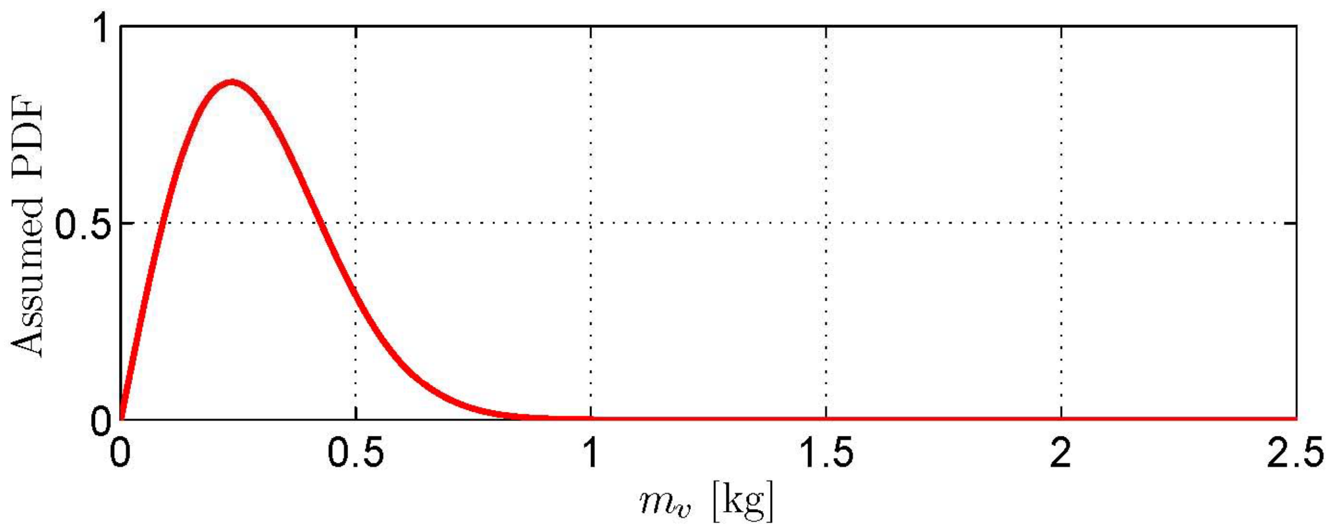

Toy Example

Feasible 2D poses for invariant kinodynamic constraints

- Considering ***only known properties*** : **DETERMINISTIC** feasibility

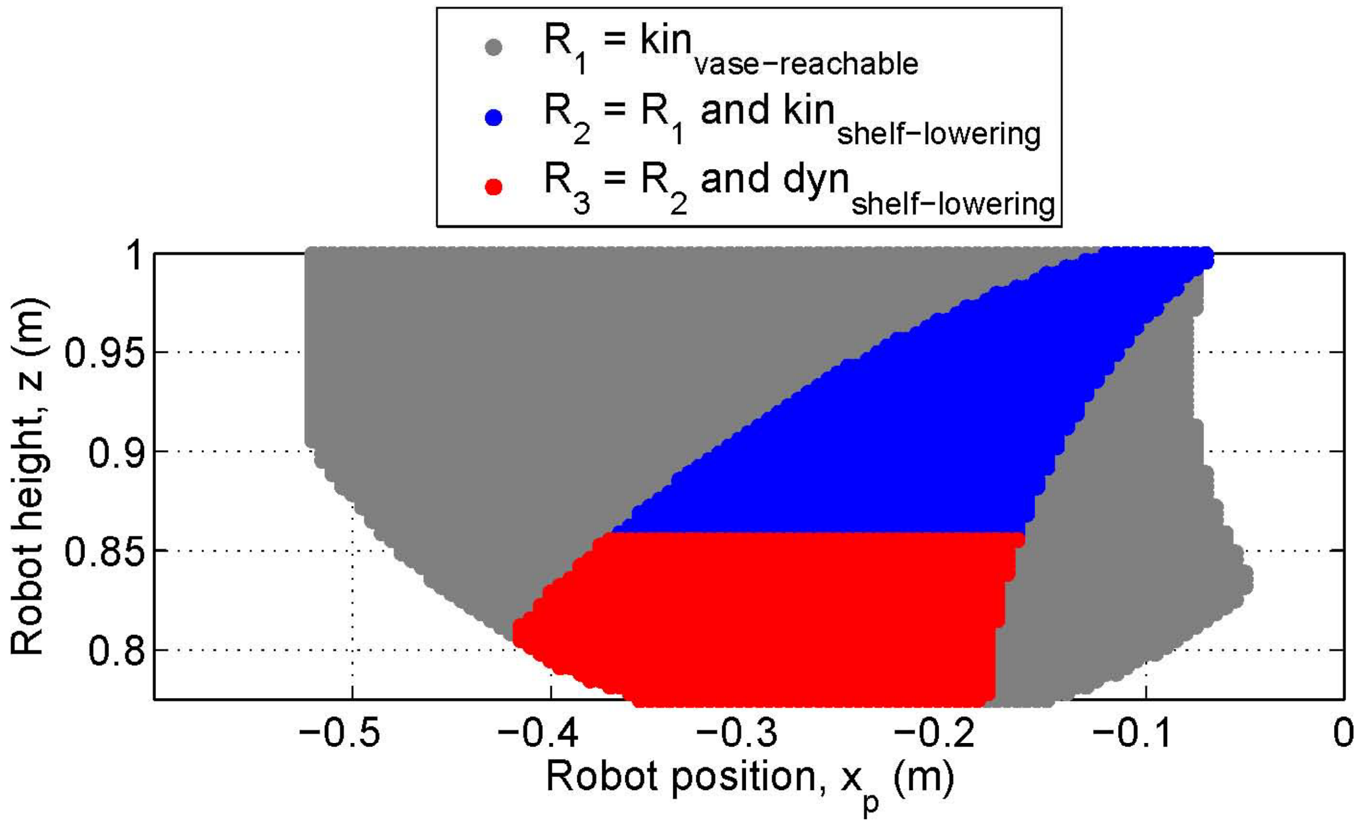

Toy Example

Kinematically feasible, with **uncertainty in mass of vase**

- Plot show probability of success : **STOCHASTIC** feasibility

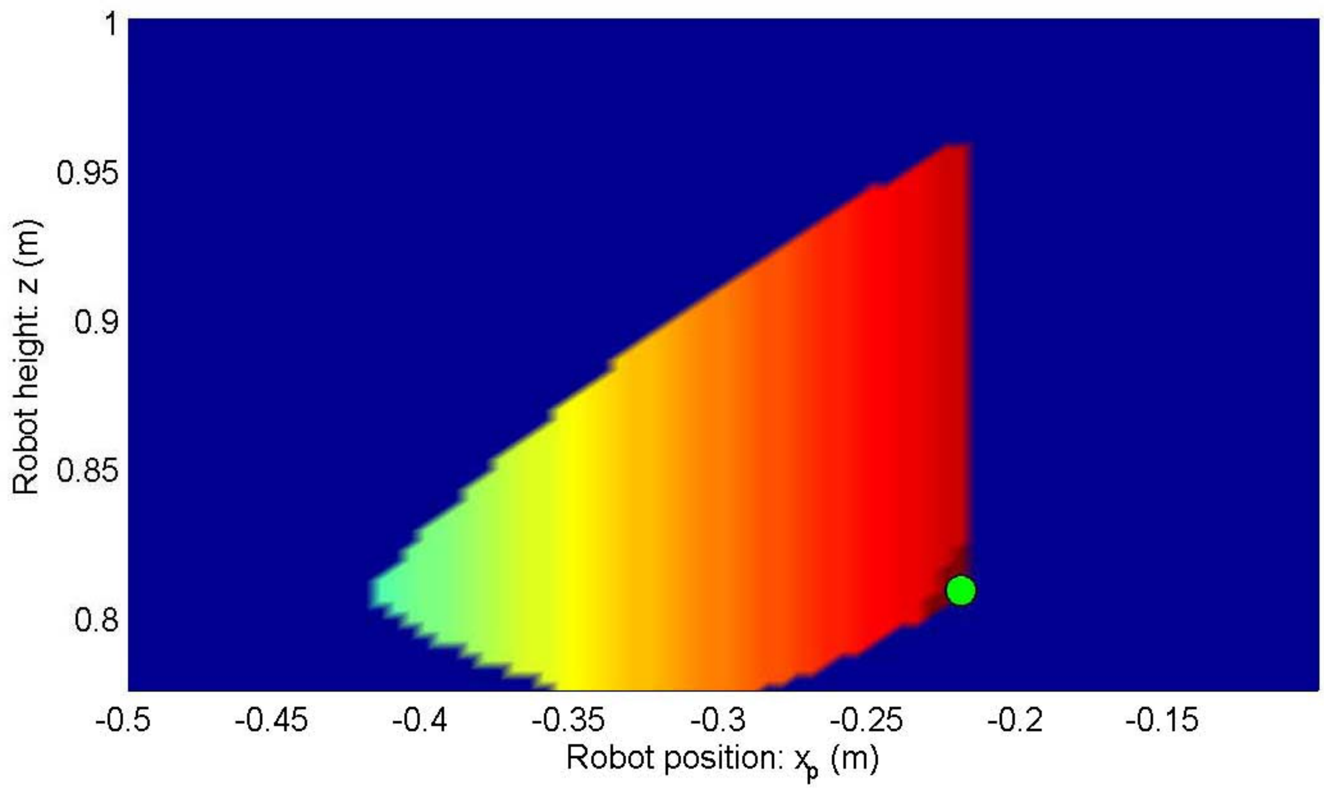

Toy Example

Kinematically feasible, with **uncertainty in mass of vase**

- Plot show probability of success : **STOCHASTIC** feasibility

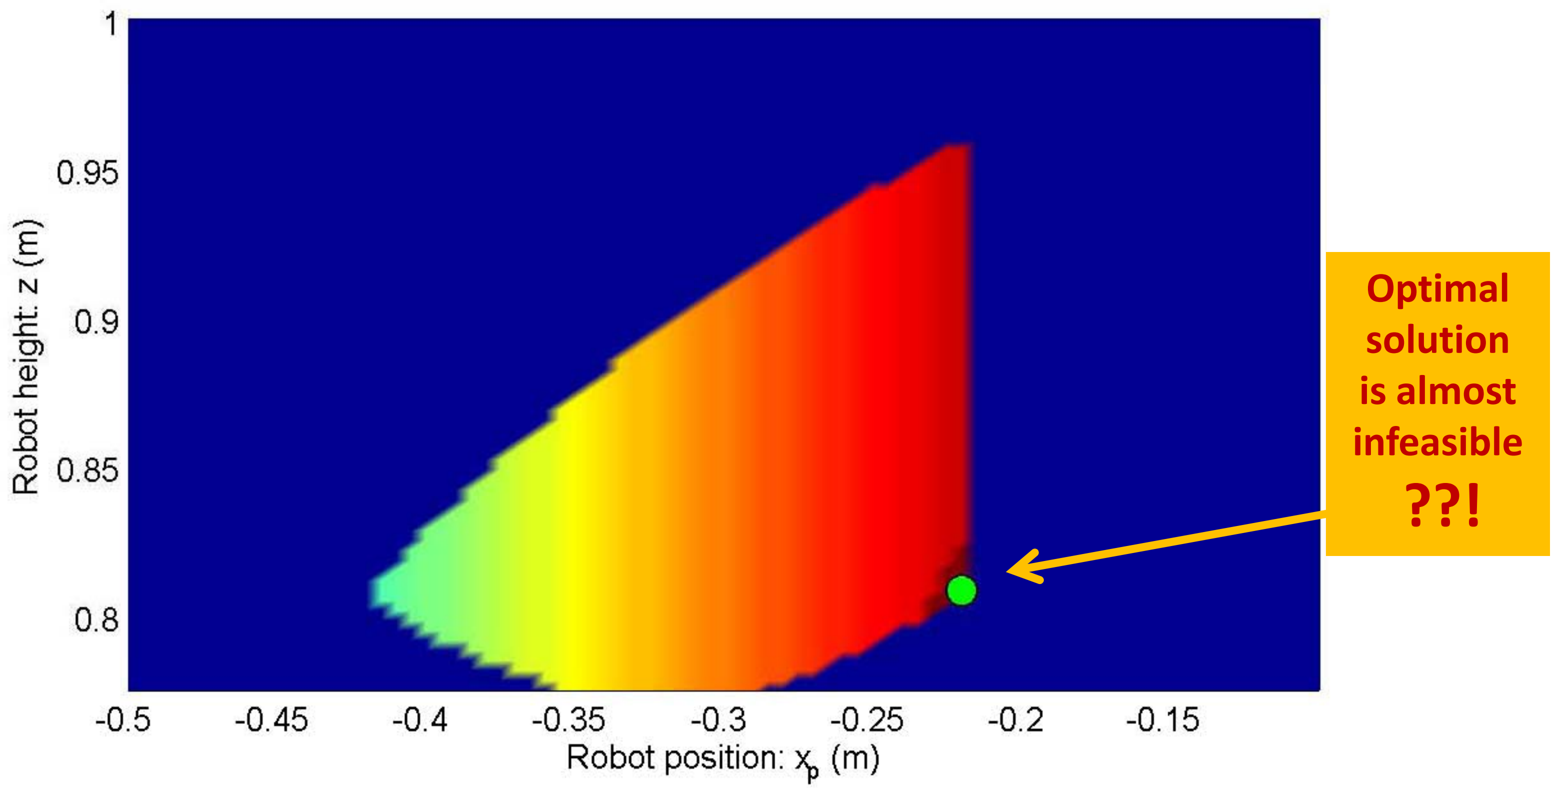

Toy Example

- Add some **uncertainty in geometry**, too
  - Here, uncertainty is in **vase position**

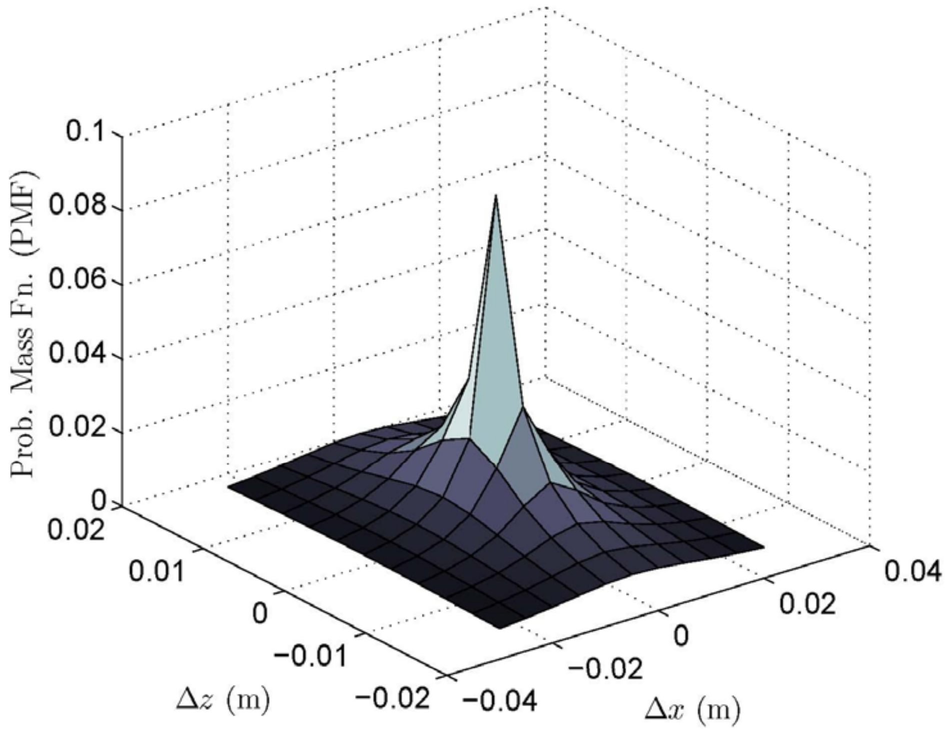

Toy Example

- Safety margins now emerge (naturally)

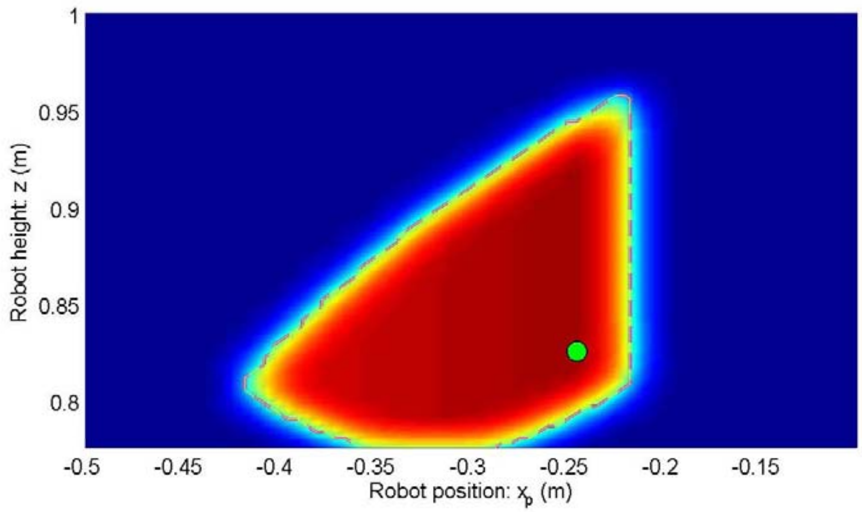

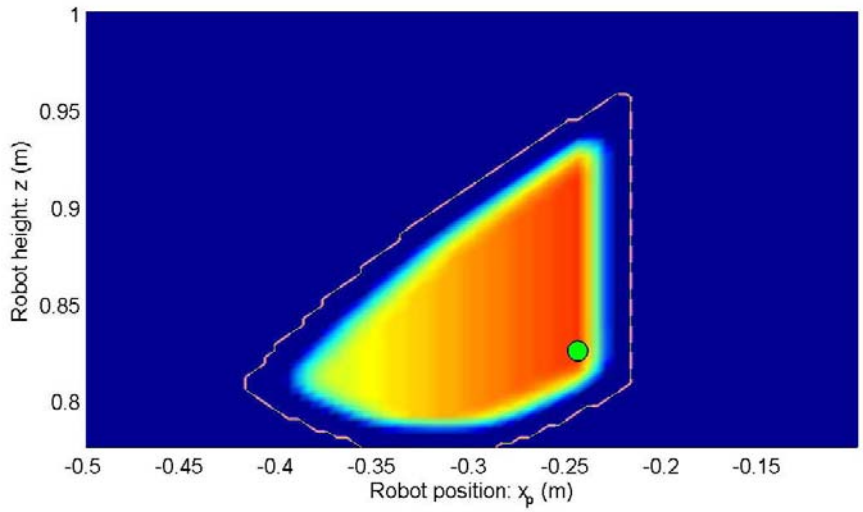

At left: Color scale goes from 0 (blue) to near 100% (red).

At right: All poses with probability of success less that 80% are dark blue.

Generalized Approach

- **Markov model:**
  - Robot kinodynamics
  - Environment
- **Transition matrix:**
  - Depends on policy
  - Time-to-failure metric (2^{nd} eigenvalue)
- **Policy improvement:**
  - Improve low-level models
  - Machine learning: iterate

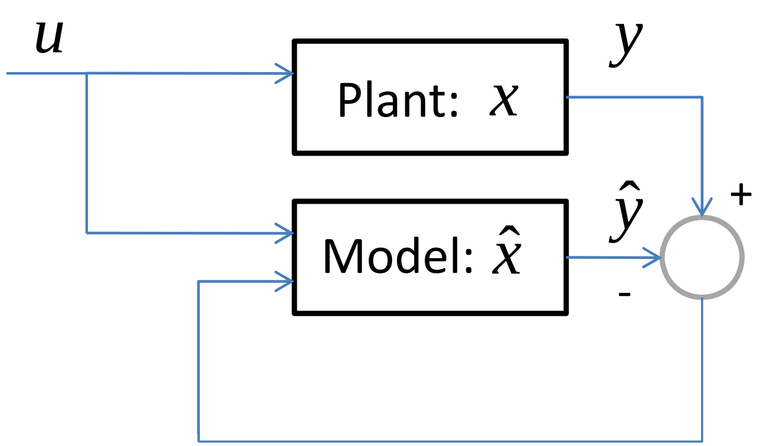

As in much of control, differences between measured states, y, and expected states, $\hat{y}$, can be used (fed back) to update a model of the dynamics.

Potential Issues

- How sensitive are solutions to noise models?
  - For a particular task, overall probability varies
  - But resulting policy seems much less sensitive (?)
  - Relative magnitudes of uncertainty matter

- How do we estimate failure rates?
  - It may take too long to observe actual failures
  - Models must capture dominant failure modes

Brief aside: **Service** vs. **Disservice**

- The **ultimate** planning issues:
  - How will and should home robotics be used?
    - **Planner capabilities must match applications**
  - Toward enabling safe human-robot interaction

- What **future applications** are we **enabling**?
  Robots working ***with*** humans, not ***for*** them
  - **Rehabilitation**
  - **Retuning sensory-motor capabilities**
    (games and play – with a purpose)

Conclusions

- **Goal: Maximize probability of success**
- **Methods: Markov modeling**
  - Include uncertainties in our modeled dynamics
  - Estimate and adjust transition probabilities online

Discussion

- A good approach? (What do humans optimize?)
- Why not? (Too challenging…)
- **What are better alternatives???**