# Guest Lecture

METR 4202: **Robotics** & Automation


Dr Surya Singh -- Lecture # 12 (Week 13)


**October 21, 2019**


**metr4202-staff@itee.uq.edu.au**

http://robotics.itee.uq.edu.au/~metr4202/

[http://**metr4202.com**]

# Probabilistic Robotics: Control
(LQR, Value Functions, Q-Learning,
Deep Reinforcement Learning, etc.)
# & The Future of Robotics/Automation
(Open Challenges)

METR 4202: **Robotics** & Automation

Dr Surya Singh -- Lecture # 12 (Week 13)

**October 21, 2019**

**metr4202-staff@itee.uq.edu.au**

http://robotics.itee.uq.edu.au/~metr4202/

[http://**metr4202.com**]

# Reference Material

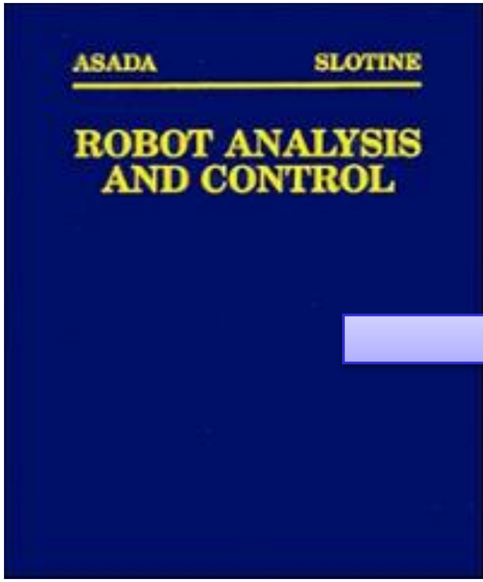


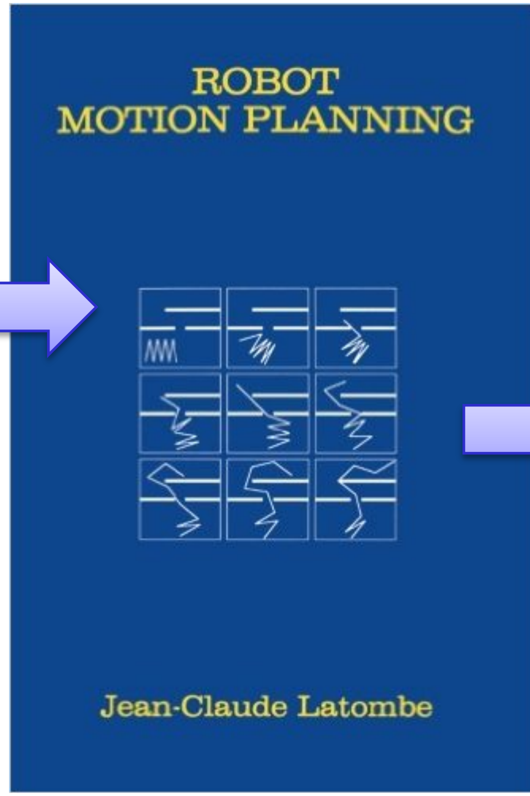


UQ Library
(TJ211.4 .L38 1991)

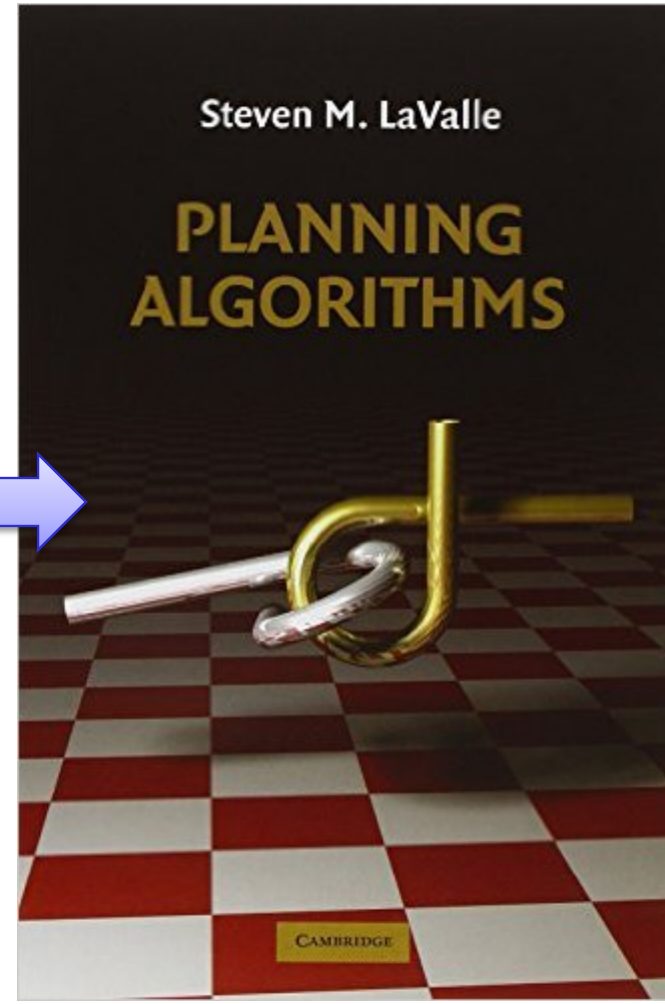


UQ Library / Online (PDF)

# Probabilistic Robotics

What about variation?

# Inherent Sources of Uncertainty

There are some common sources of uncertainty:

- Inherent stochasticity in the system being modeled
- Incomplete modeling
- Incomplete observability
- Incompetent control (or action)

# Two Views of Probability

## Frequentist probability

- Relates directly to the rates at which events occur
- Basically, things that are directly "repeatable"
- e.g. Drawing a hand of cards

## Belief (Bayesian) probability

- Degree of a Belief
- Qualitative levels of uncertainty
- For more details about why a small set of common sense
- Ramsey (1926) –
  The same set of axioms control both kinds of probability

# II.1 Basic Probability Theory

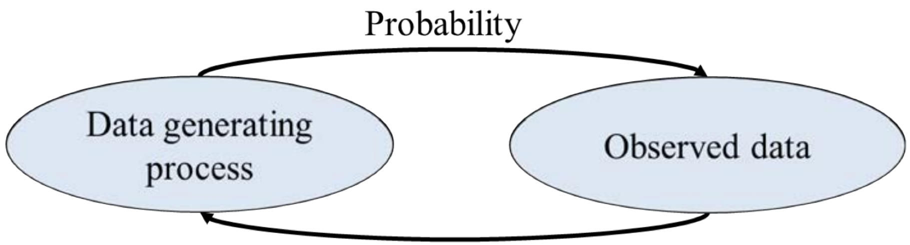


- Probability Theory
  - Given a data generating process, what are the properties of the outcome?
- Statistical Inference
  - Given the outcome, what can we say about the process that generated the data?
  - How can we generalize these observations and make predictions about future outcomes?

# In Particular: Bayes' Rule

$$P(\mathrm{x} \mid \mathrm{y}) = \frac{P(\mathrm{x})P(\mathrm{y} \mid \mathrm{x})}{P(\mathrm{y})}$$

- What does Bayes' Formula helps to find?
  - Helps us to find:

$$P(B \mid A)$$

  - By having already known:

$$P(A \mid B)$$

# Example: Inference / Generative Models

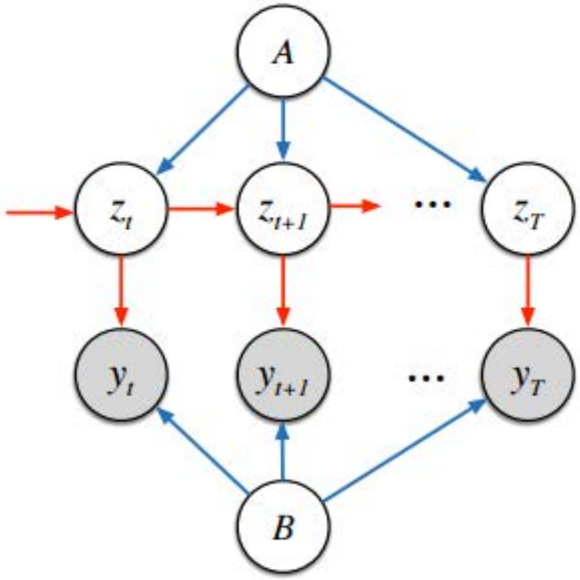


**Gaussian Linear State Space Model**
**Kalman Filter**

$$z_t \sim \mathcal{N}(z_t | A z_{t-1}, \sigma_z^2 I)$$

$$y_t \sim \mathcal{N}(y_t | B z_t, \sigma_y^2 I)$$

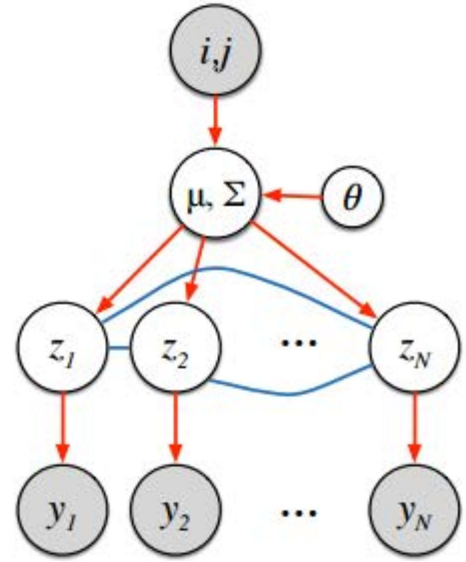


**Latent Gaussian Cox Point Process**

$$x \sim \mathcal{N}(x | \mu(i,j), \Sigma(i,j))$$

$$y_{ij} \sim \mathcal{P}(c \exp(x_{ij}))$$

# Robot Controls

## Controls Example I: Cart & Pole
## (& Obstacles)

# Classic State-Space Problem: Cart & Pole

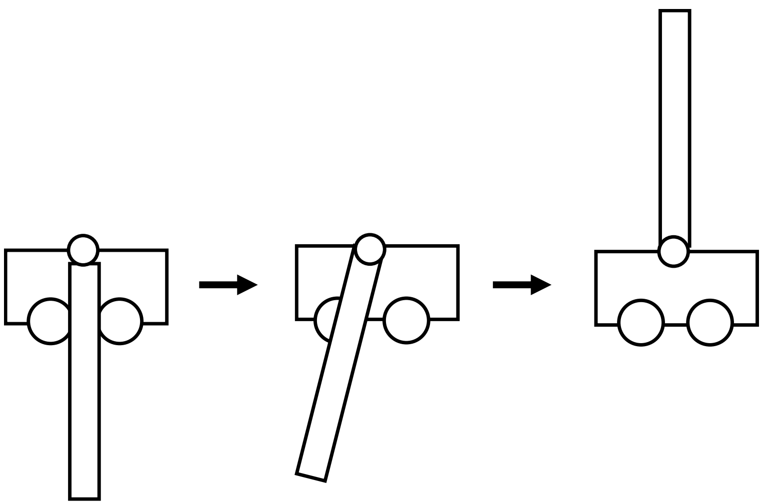

- Equations of State:

$$(M+m)\,\ddot{x} + mL\ddot{\theta}\cos\theta - mL\dot{\theta}^2\sin\theta = f$$
$$mL\ddot{x}\cos\theta + mL^2\ddot{\theta} - mgL\sin\theta = 0$$

- Solution:

$$\ddot{x} = \frac{f + m\sin\theta\left(L\dot{\theta}^2 - g\cos\theta\right)}{M + m\sin^2\theta}$$
$$\ddot{\theta} = \frac{-f\cos\theta - mL\dot{\theta}^2\sin\theta\cos\theta + (M+m)g\sin\theta}{L(M + m\sin^2\theta)}$$

$$\dot{X} = \begin{bmatrix}\dot{x}\\ \ddot{x}\\ \dot{\theta}\\ \ddot{\theta}\end{bmatrix} = \begin{bmatrix}\dot{x}\\ \dfrac{f + m\sin\theta(L\dot{\theta}^2 - g\cos\theta)}{M + m\sin^2\theta}\\ \dot{\theta}\\ \dfrac{-f\cos\theta - mL\dot{\theta}^2\sin\theta\cos\theta + (M+m)g\sin\theta}{L(M+m\sin^2\theta)}\end{bmatrix} = \begin{bmatrix}\dot{x}\\ \dfrac{f - m\theta g}{M}\\ \dot{\theta}\\ \dfrac{-f + (M+m)g\theta}{LM}\end{bmatrix} = \begin{bmatrix}0 & 1 & 0 & 0\\ 0 & 0 & -\dfrac{mg}{M} & 0\\ 0 & 0 & 0 & 1\\ 0 & 0 & \dfrac{(M+m)g}{LM} & 0\end{bmatrix}\begin{bmatrix}x\\ \dot{x}\\ \theta\\ \dot{\theta}\end{bmatrix} + \begin{bmatrix}0\\ \dfrac{1}{M}\\ 0\\ -\dfrac{1}{LM}\end{bmatrix}f$$

See Also: Tutorial 12: Cart-Pole Inverted Pendulum (State-Space)

# Motion Planning & Control: a Unified Design Tool?

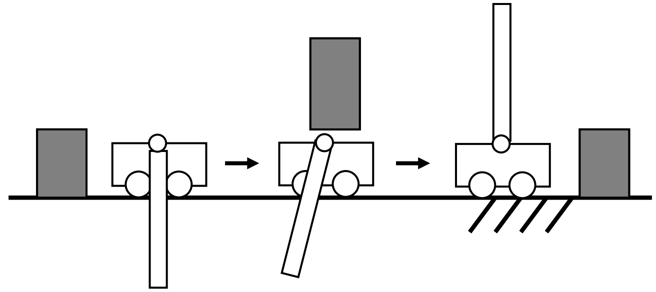

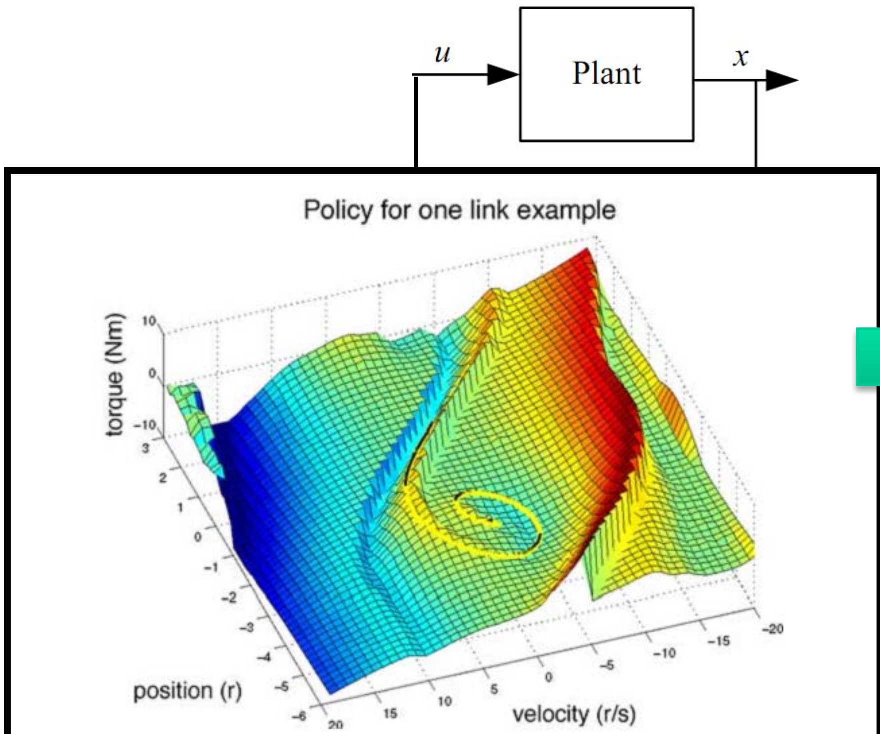


Atkeson, C. & Stephens, B.
*"Random sampling of states in dynamic programming"* IEEE T. Syst. Man Cybern, 2008

Maeda, Singh, Durrant-Whyte, *ICRA* 2011

POMDP Framework:

$$\pi^*(b) = \underset{a \in \mathcal{A}}{\operatorname{argmax}} \left( \int_{s \in \mathcal{S}} R(s,a)b(s)ds + \mathbb{E}\left[ \sum_{t=1}^{\infty} \gamma^t R(s_t, a_t) | \tau(b,a,O),\ \pi^* \right] \right)$$

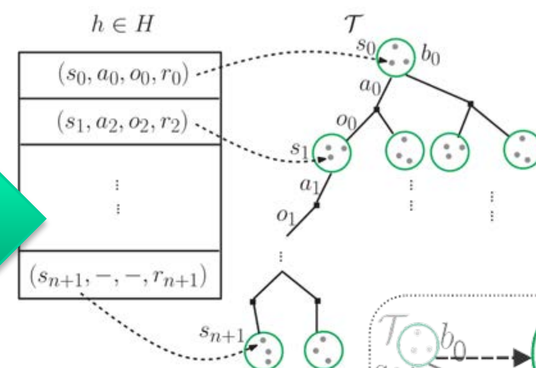


Belief tree

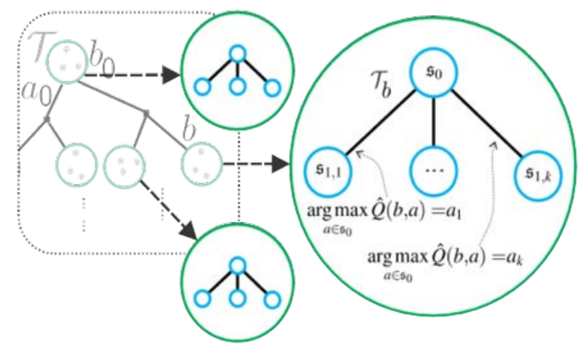


GPS tree

Seiler, *et al*. ICRA, May 2015, 2290-2297

# Why? Bellman's Principle of Optimality

- Suppose we are given a DT ODE

$$\Delta x = F(x, u)$$
$$x \in \mathbb{R}^n, u \in \mathbb{R}^m$$

- Optimal Control:

Find $u: [0, t] \to \mathbb{R}^m$ that minimizes the cost (or said more optimistically: *maximizes the reward*):

$$c(x, u) = \underbrace{\boxed{l(t, x(t))}}_{\text{Final cost}} + \underbrace{\boxed{\sum_{s=0}^{t-1} \mathcal{L}(s, x(s), u(s))}}_{\text{"Running" cost}}$$

- Bellman's insight (1957):

* the optimal control at time s depends only on $\boldsymbol{x}(\boldsymbol{s})$

(i.e. not any prior states!)

∴ It's an MDP!

# Motion Planning & Control:
## Trajectory Generation with Constraints

$$\dot{y}\cos(\theta) - \dot{x}\sin(\theta) = 0$$

**Motion <u>Planning</u> Methods:**

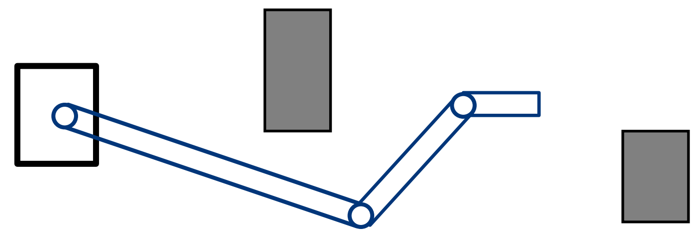

**Steering/<u>Feedback Control</u> Methods:**

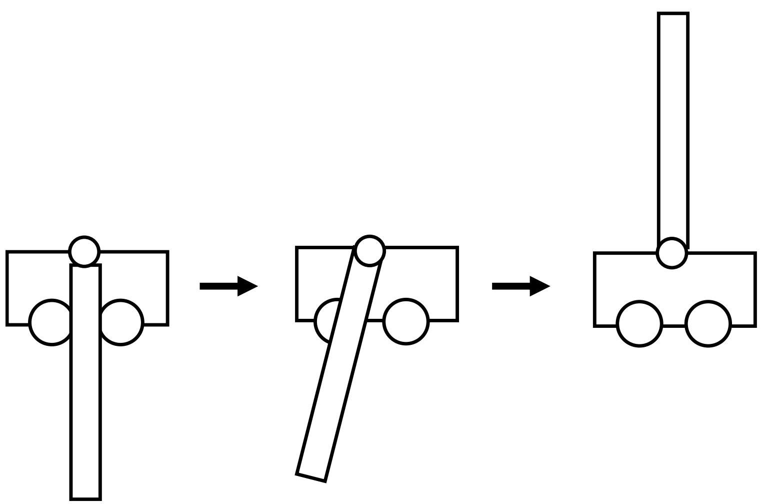

# One Approach: Gain-Scheduled RRT

Core Idea:

- Basic approach: decoupling → tractable
- Integrated approach: use feedback to shortcut the planning phase

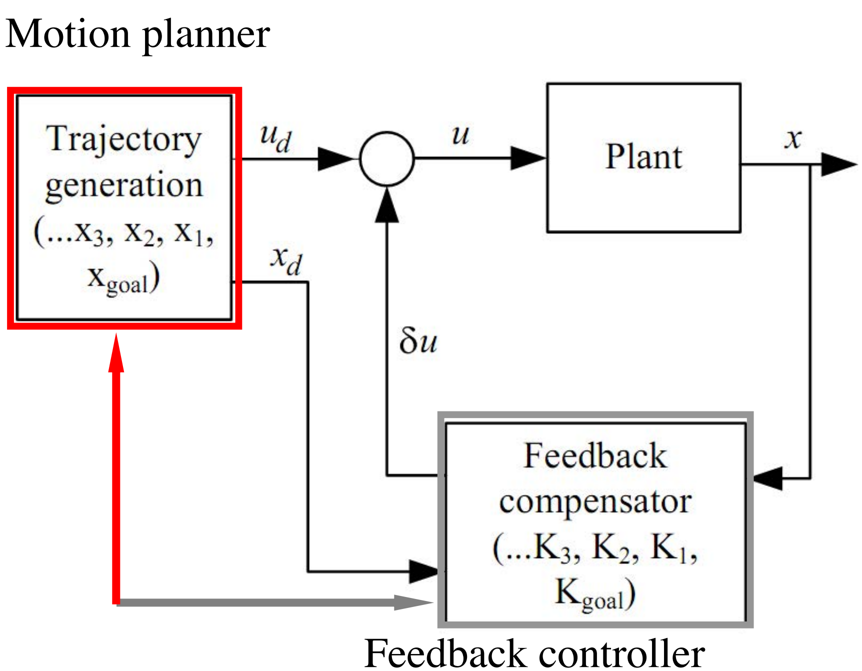


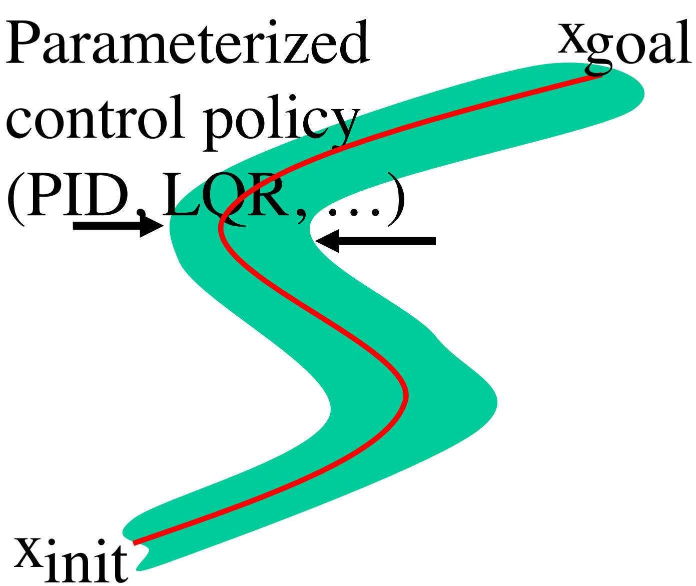


* Maeda, G.J; Singh, S.P.N; Durrant-Whyte, H. *"Feedback Motion Planning Approach for Nonlinear Control using Gain Scheduled RRTs"*, IROS 2010

# Gain-Scheduled RRT: Algorithm

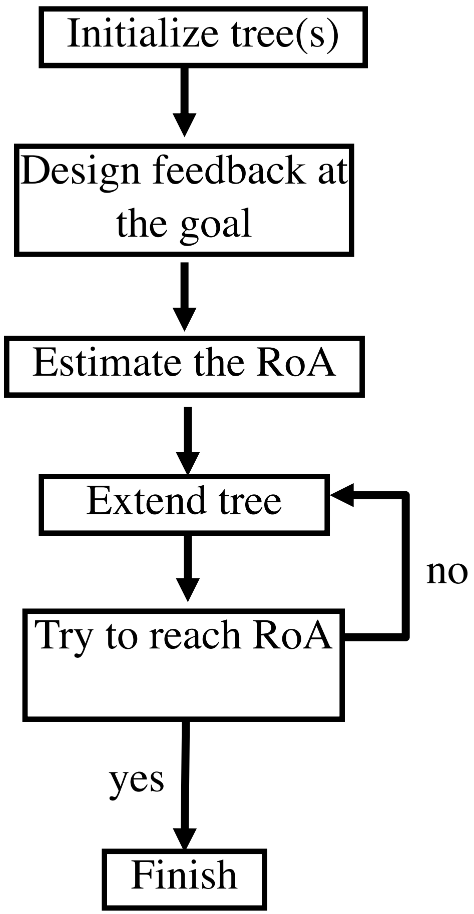


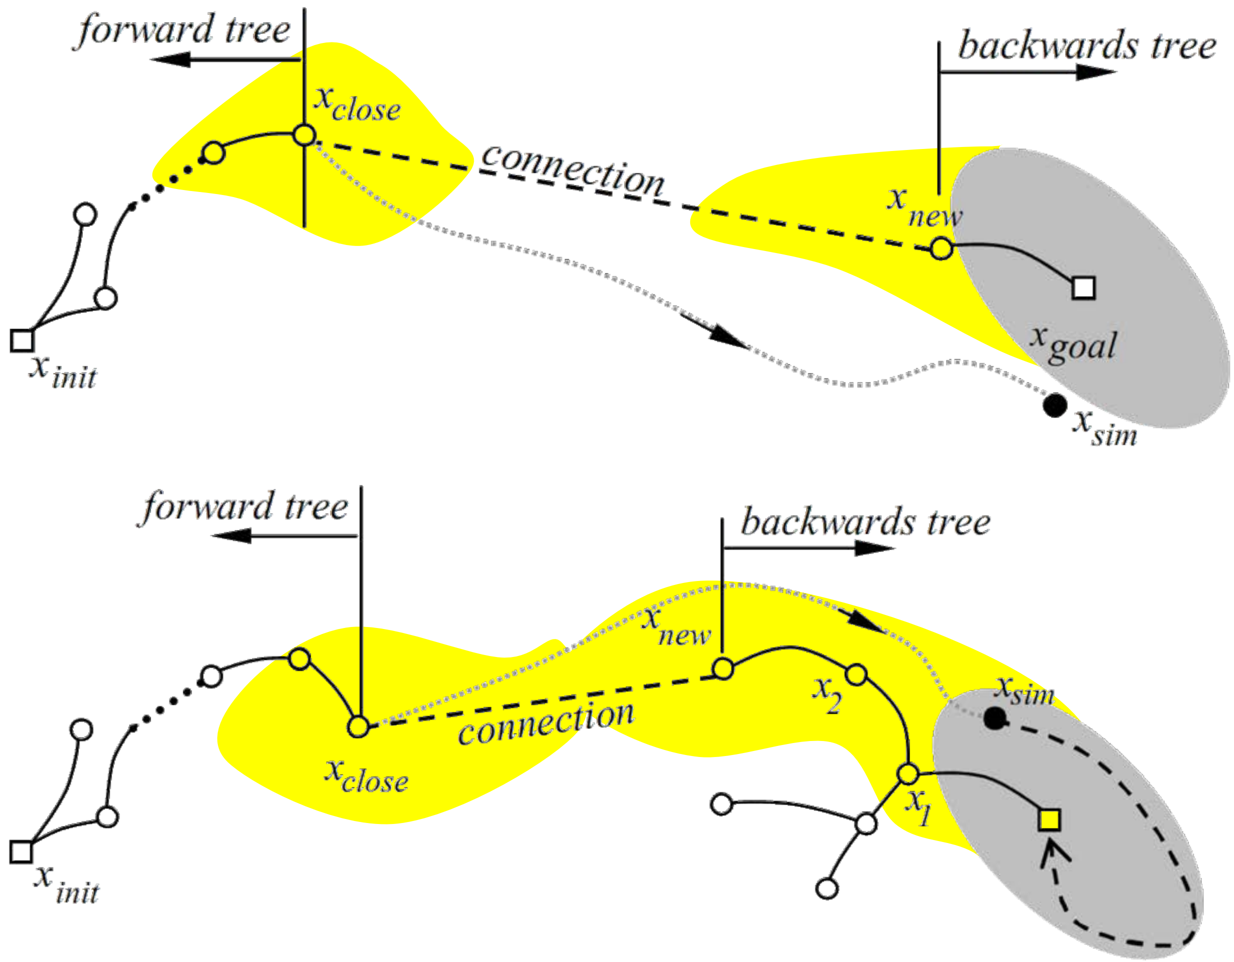

# Gain-Scheduled RRT

Core Idea:

- Basic approach: decoupling → tractable
- Integrated approach: use feedback to shortcut the planning phase

Holonomic case

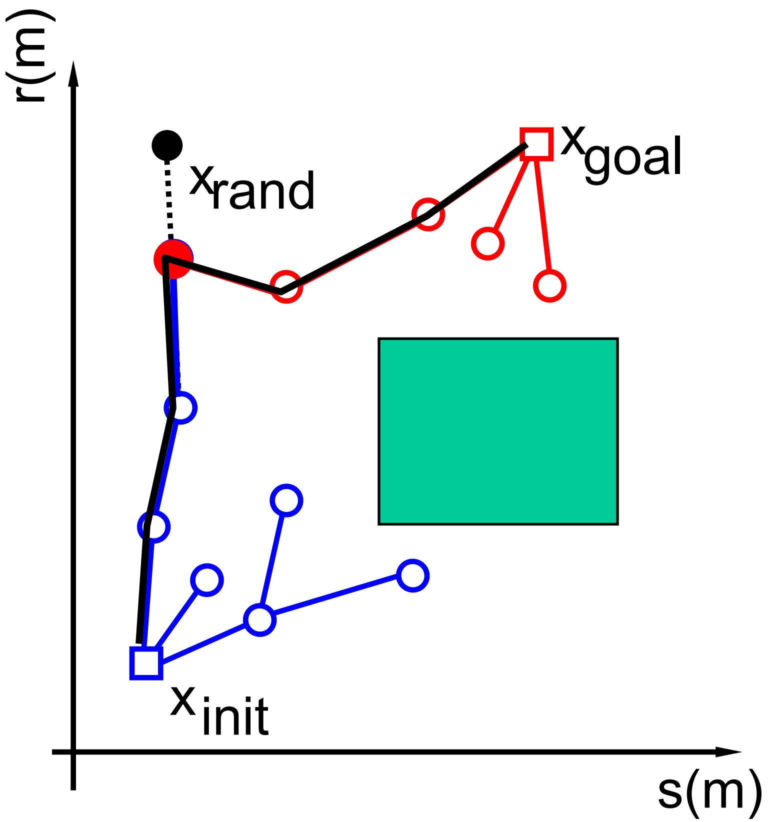


* Maeda, G.J; Singh, S.P.N; Durrant-Whyte, H. *"Feedback Motion Planning Approach for Nonlinear Control using Gain Scheduled RRTs"*, IROS 2010

# Gain-Scheduled RRT: RRT Connection Gap

A RRT solution rarely reaches the goal (or connect the two trees) with zero error

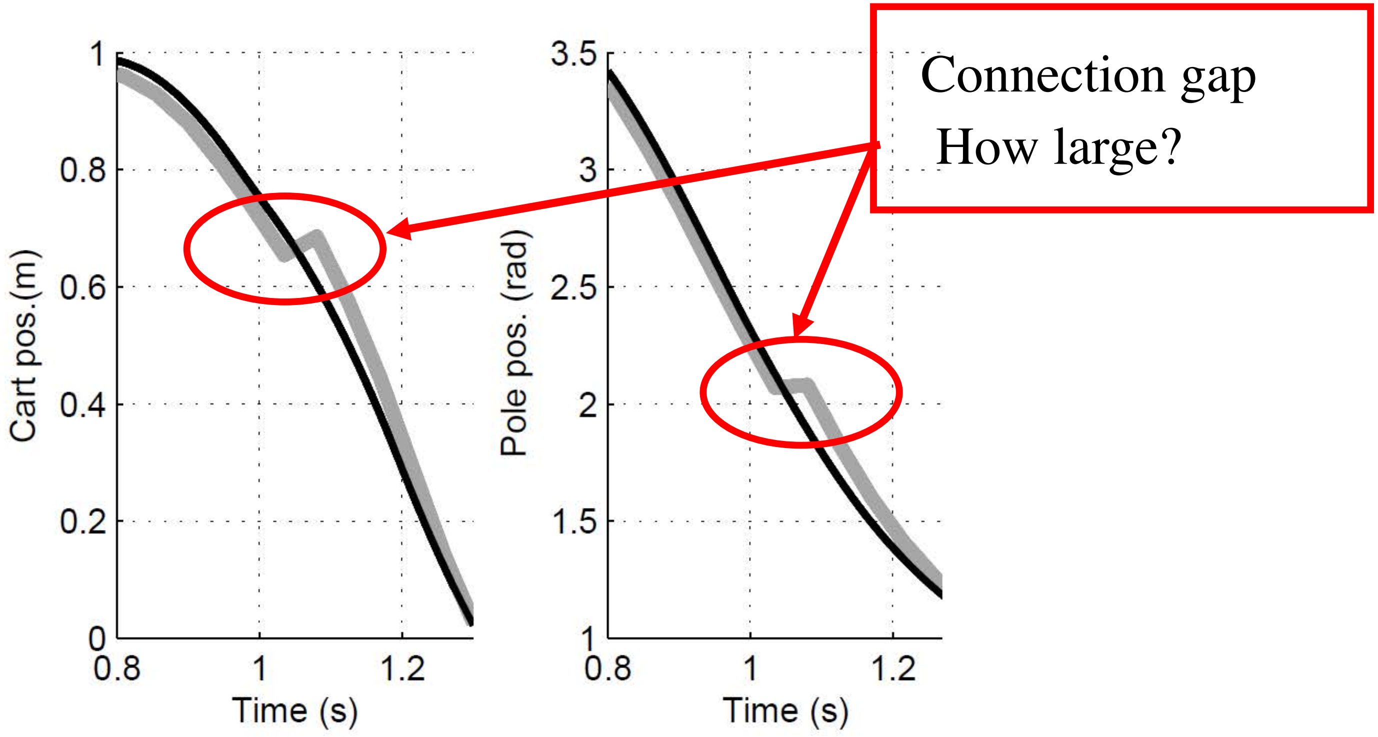

# Gain-Scheduled RRT: Search

Single state search:
Extensive exploration

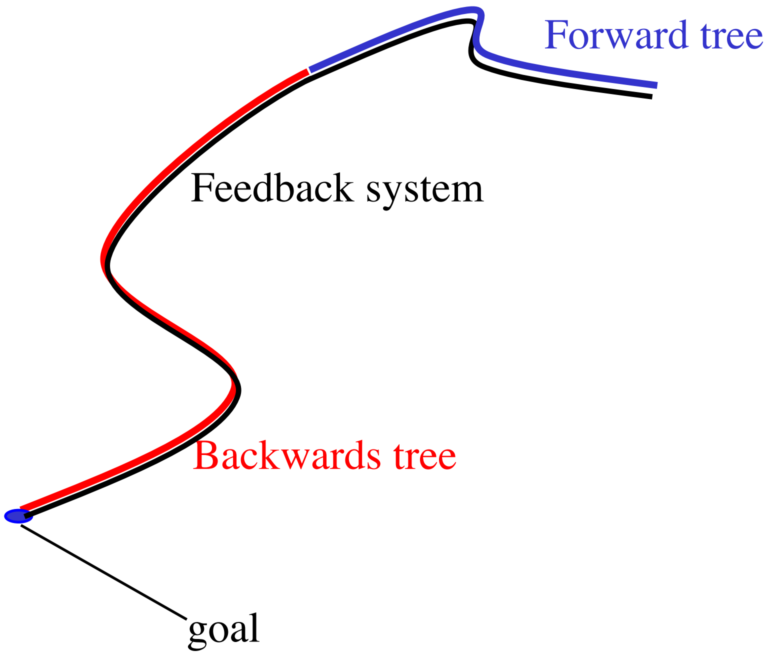


Region search
RRT connection is relaxed

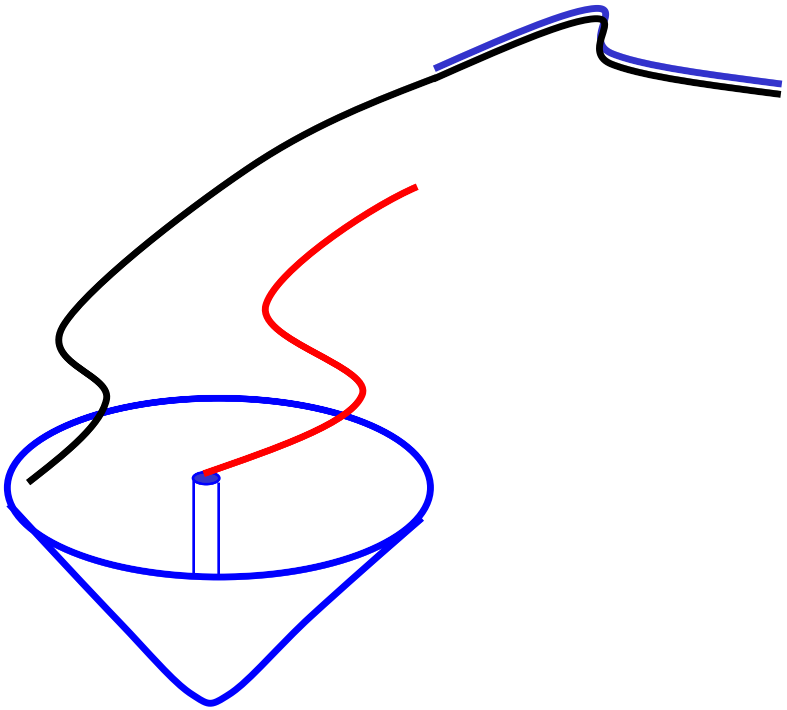

# GS-RRT: RoA & Verification

- Find a candidate

In the LQR case: J: optimal cost-to-go S: Algebraic Ricatti Eq.

$$V(x) = J^*(x_{goal} - x) = (x_{goal} - x)^T S (x_{goal} - x)$$

- Maximize

$$\dot{J}^*(\overline{x}) + h(\overline{x})(\rho - J^*(\overline{x})) < 0$$

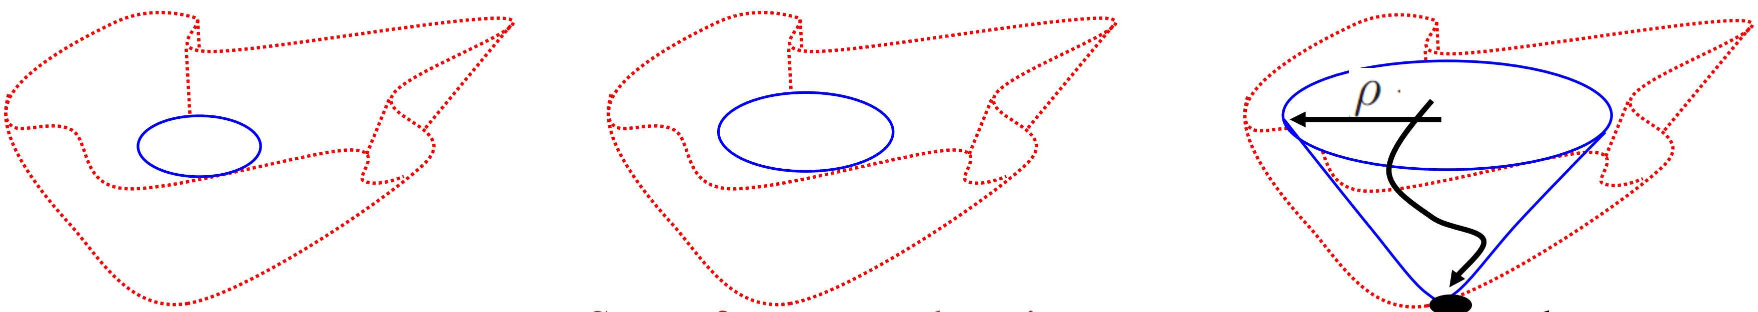


Sum of squares relaxation

$$V(x) \text{ locally } \boxed{\text{positive definite}} \text{ in } B_\rho$$

$$\dot{V}(x) \text{ locally } \boxed{\text{negative definite}} \text{ in } B_\rho$$

- Verify candidate

**R. Tedrake, *"LQR-Trees: Feedback Motion Planning on Sparse Randomized Trees"*, RSS 2009

# Gain-Scheduled RRT:

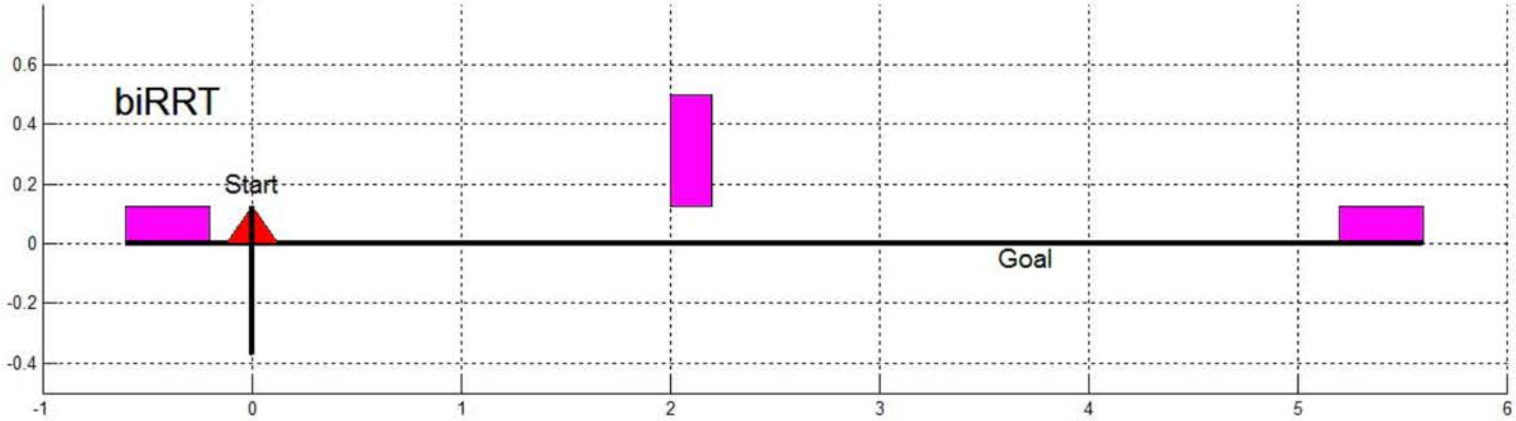


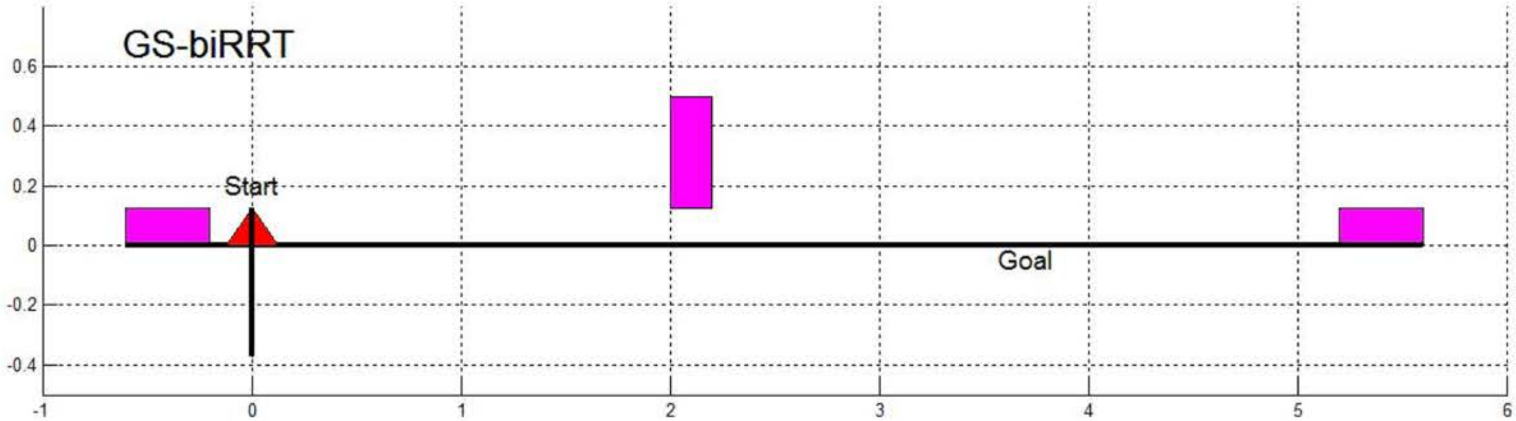

# Automatic Robot Controls

## Controls Example II:
## Underactuated Robotics

# The Jitterbug Problem

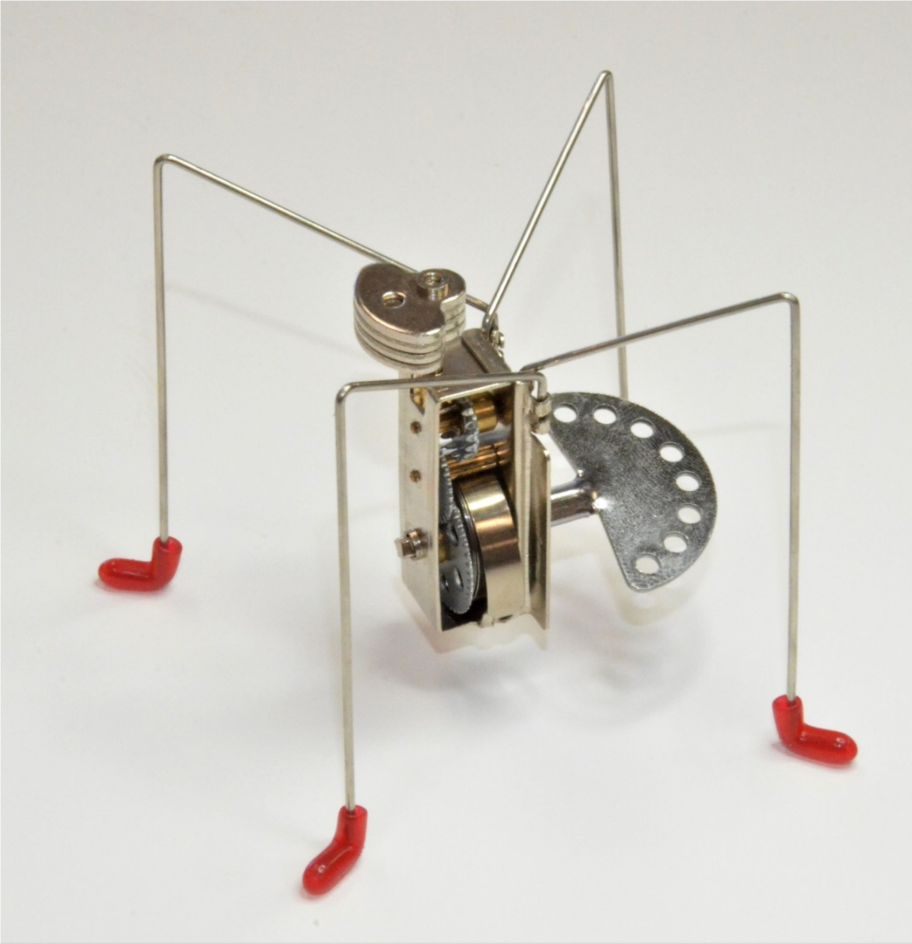

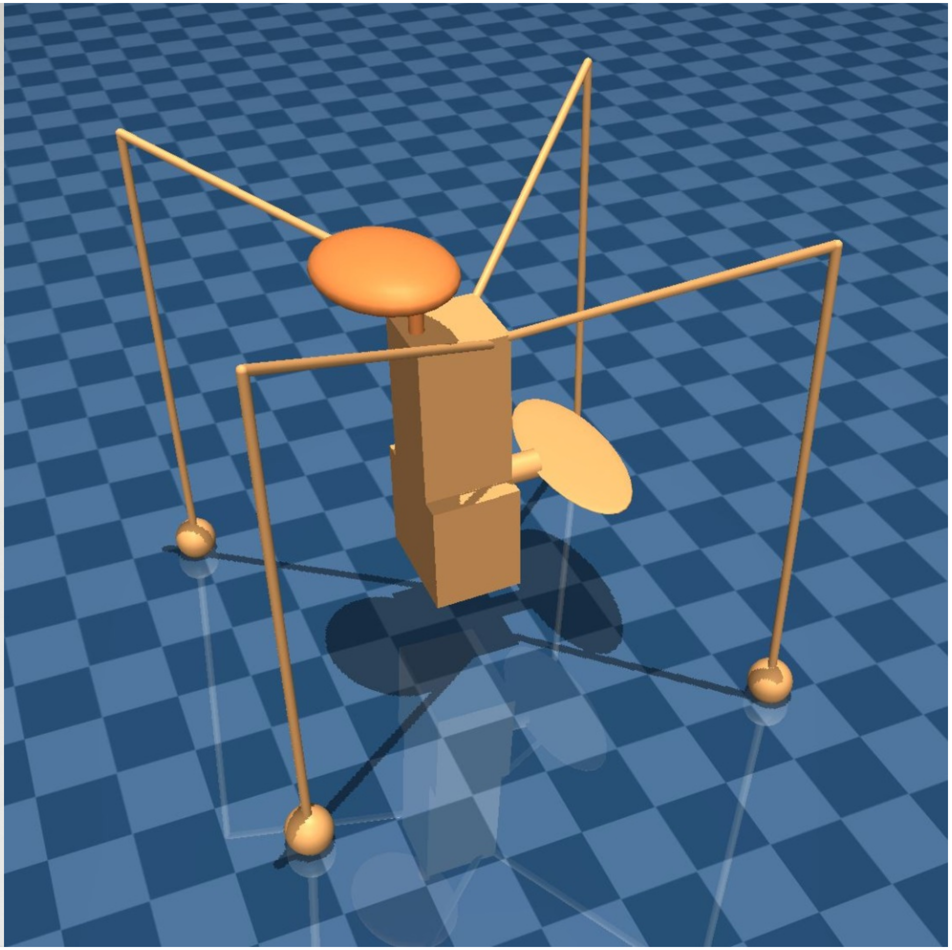

# Solving the Jitterbug Problem:
# Continuous Action Deep Reinforcement Learning

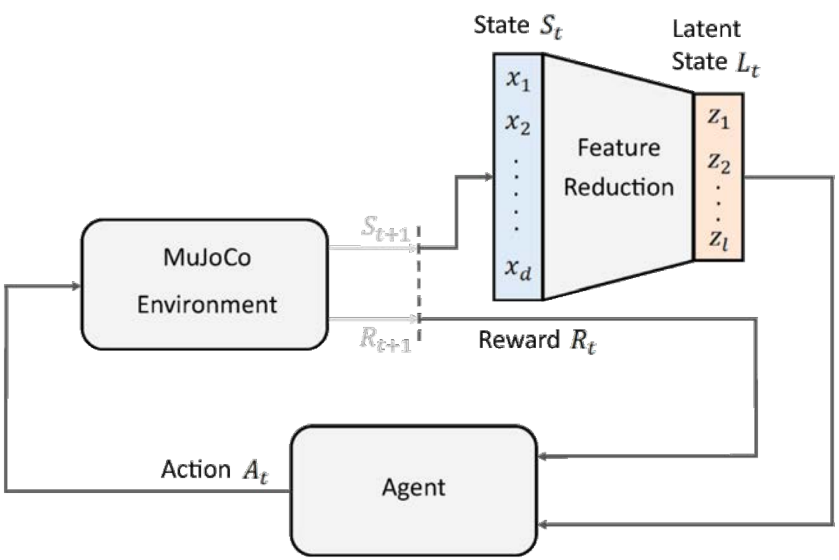


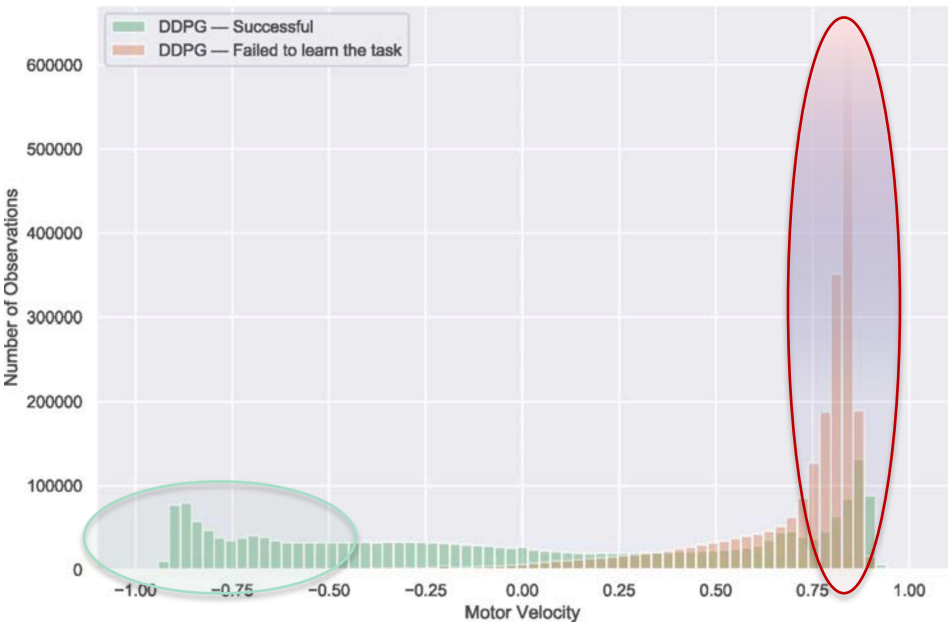

# Solving the Jitterbug Problem:

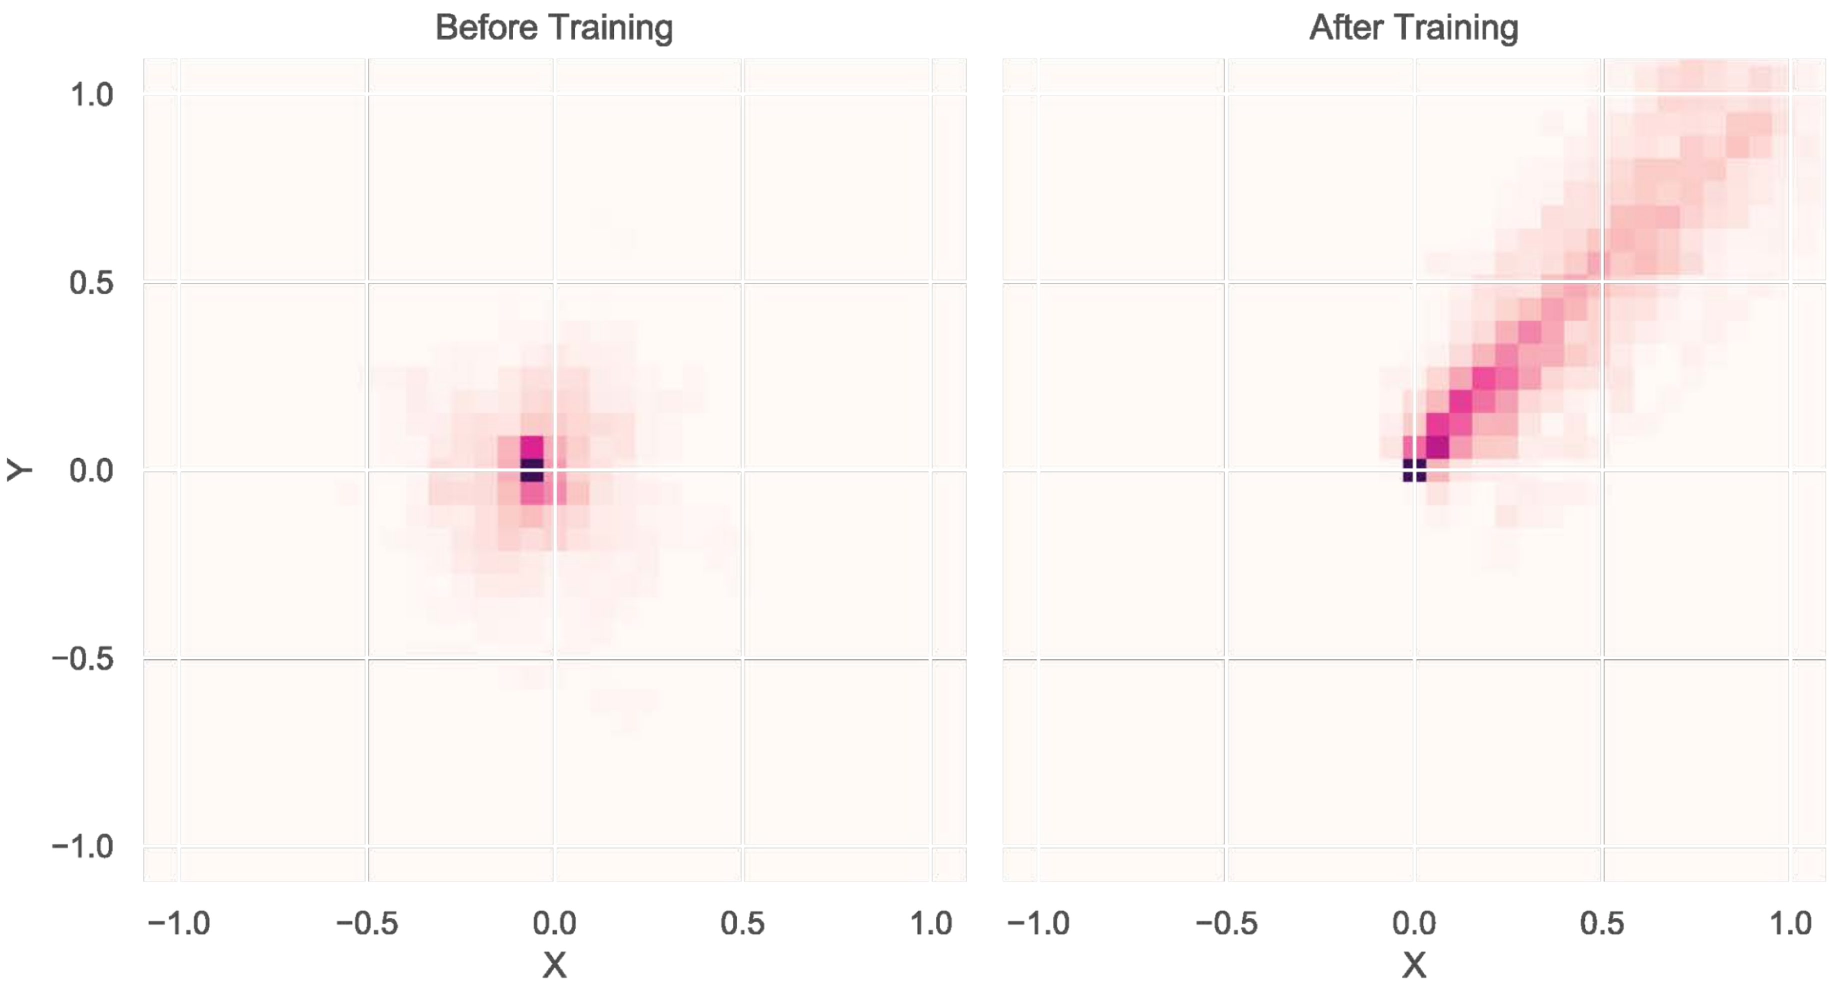

# Future of Robotics

Some Research Examples ☺

# Computer Aided Surgery: R/C Toolholders?

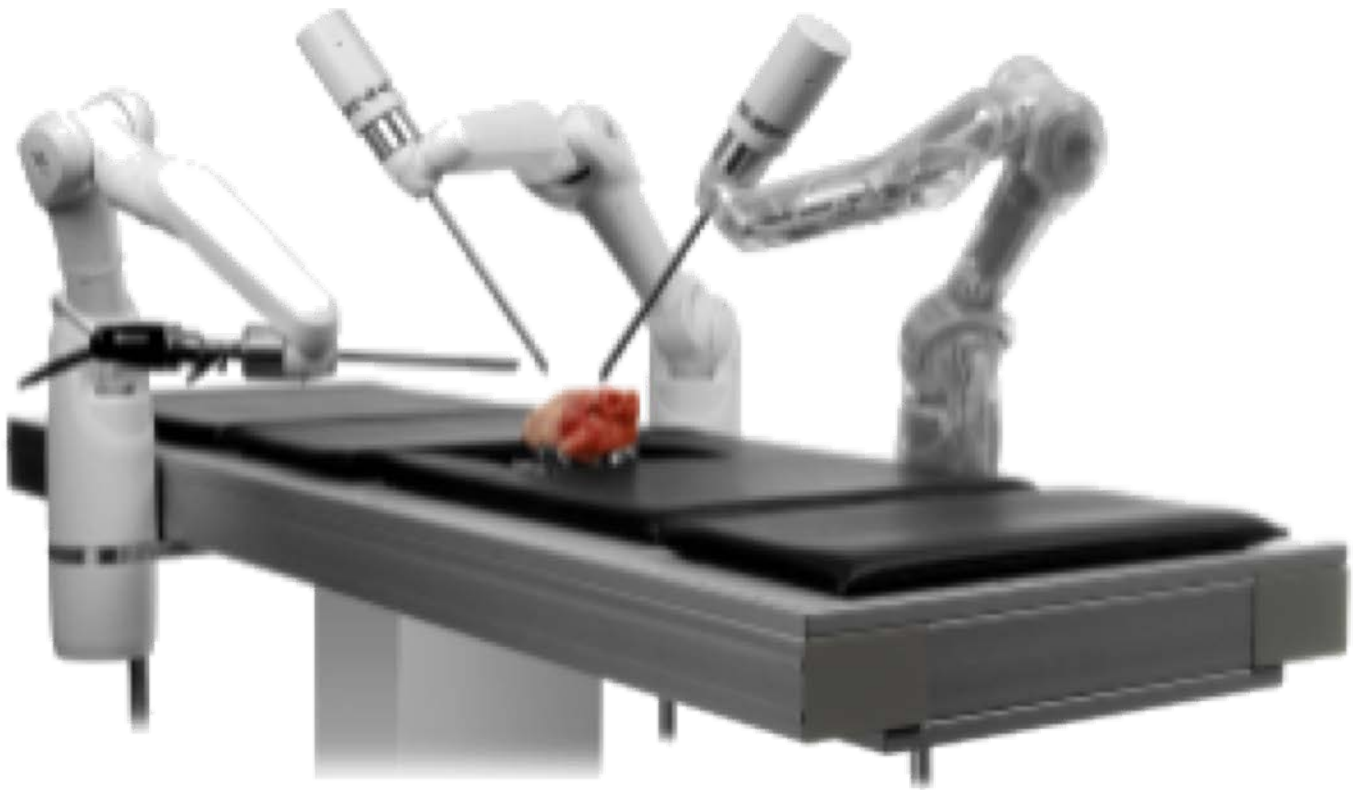

➔ Move in tandem with heart: Cardiac procedures without stopping it

- Unstructured environment (patient) makes this harder

# Modern (Tele)Surgical Robotics:

- **Biomechanics approach**: Predict expected tissue trajectories
- (Stochastic) **Robot Motion Planning / Control Methods**!

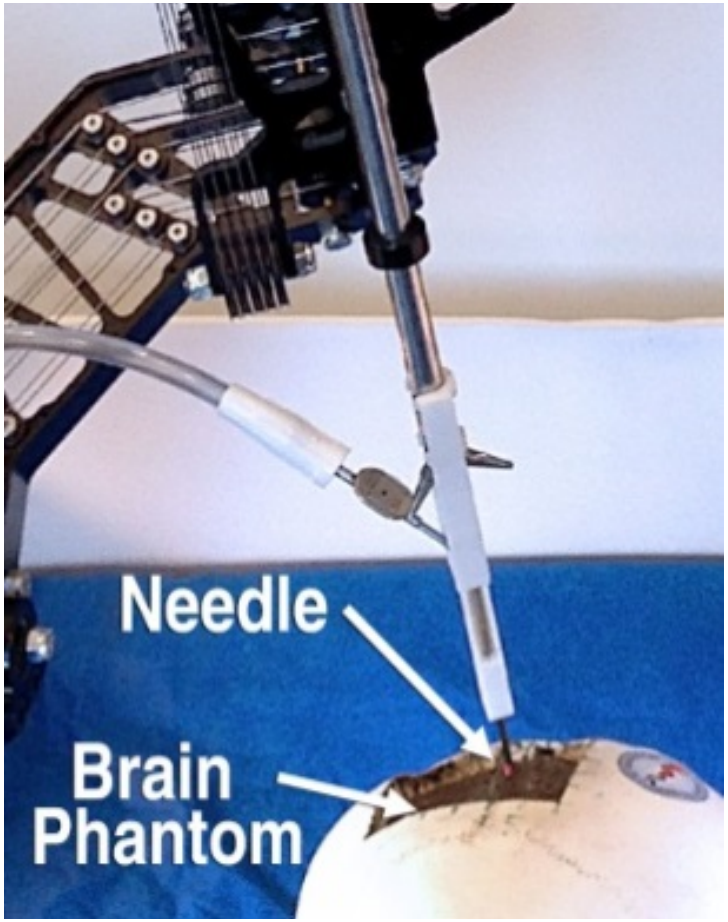


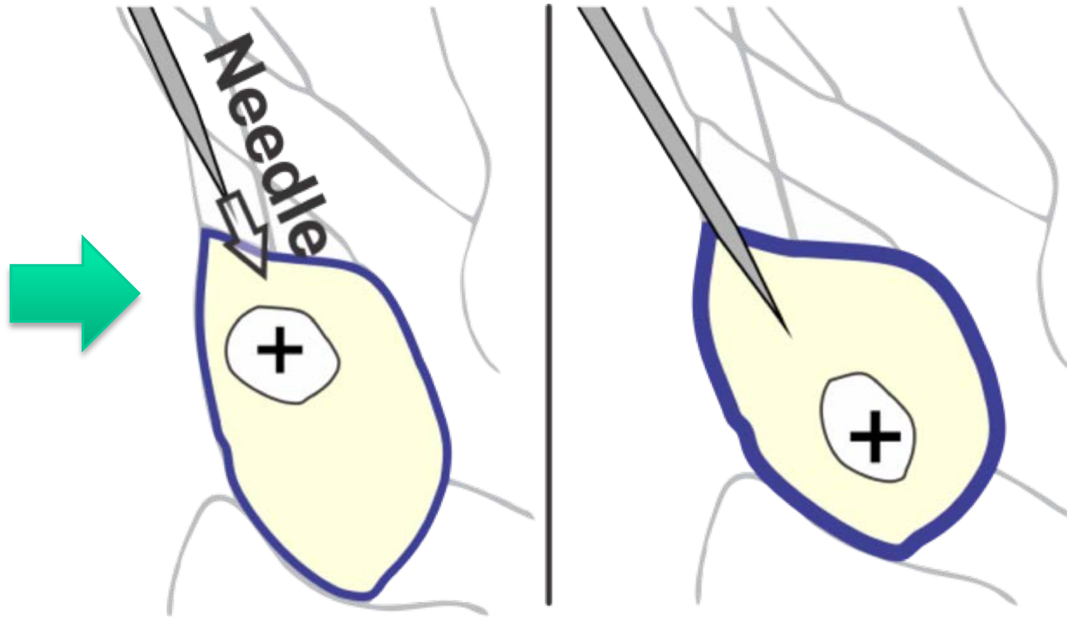


ARC DP160100714

# Computer Aided Surgery: “**Soft**” is “**Hard**”!

# **Research:** Incorporating Stiffness (Haptics):
## Visual Deformable Object Analysis

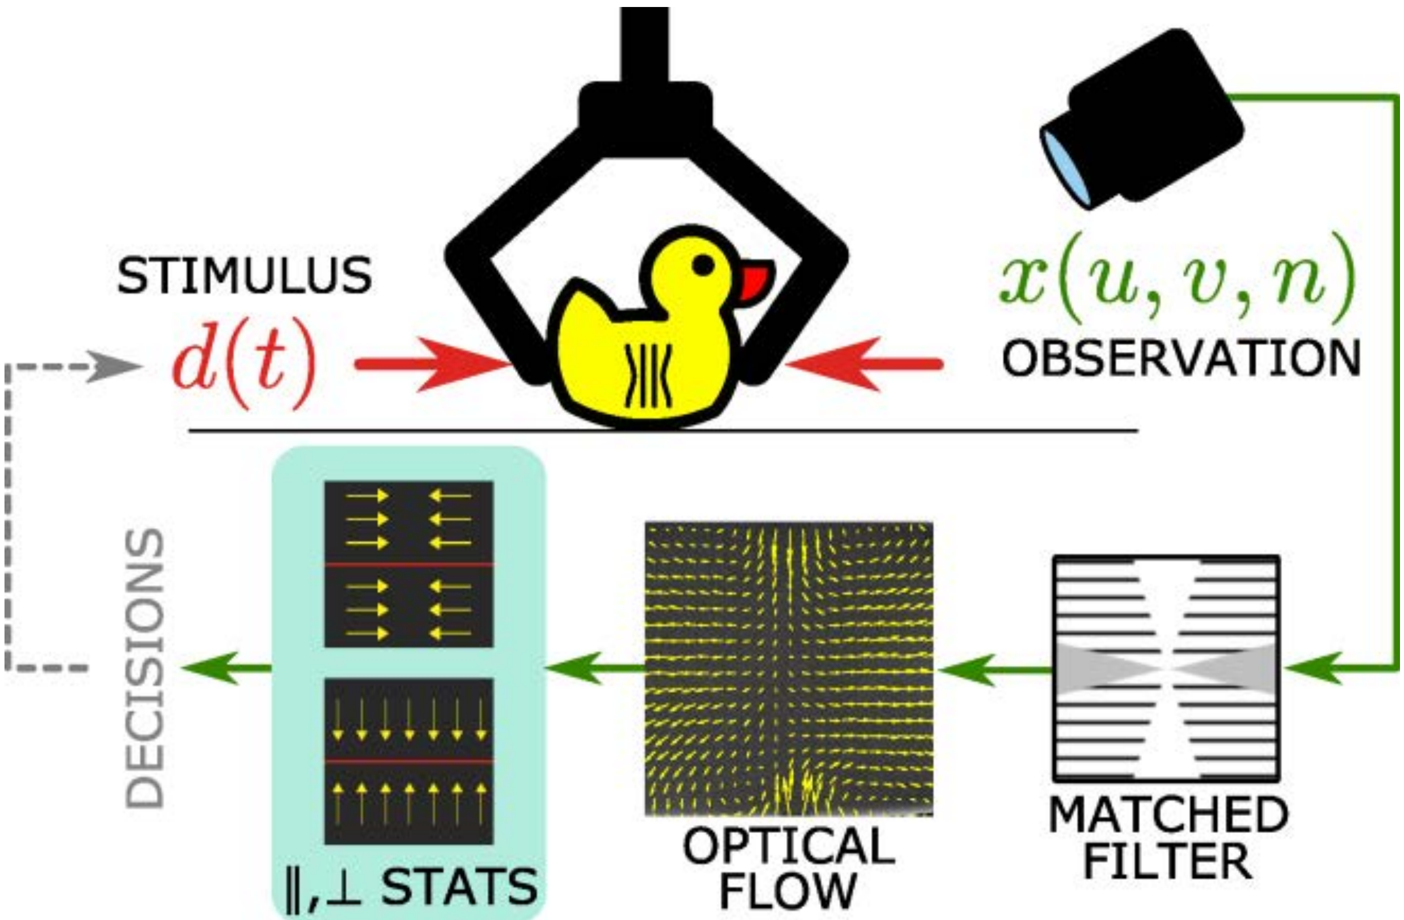


Dansereau, Singh, Leitner, *ICRA 2016*

# Iceberg to Titanic: Take Advantage of Information

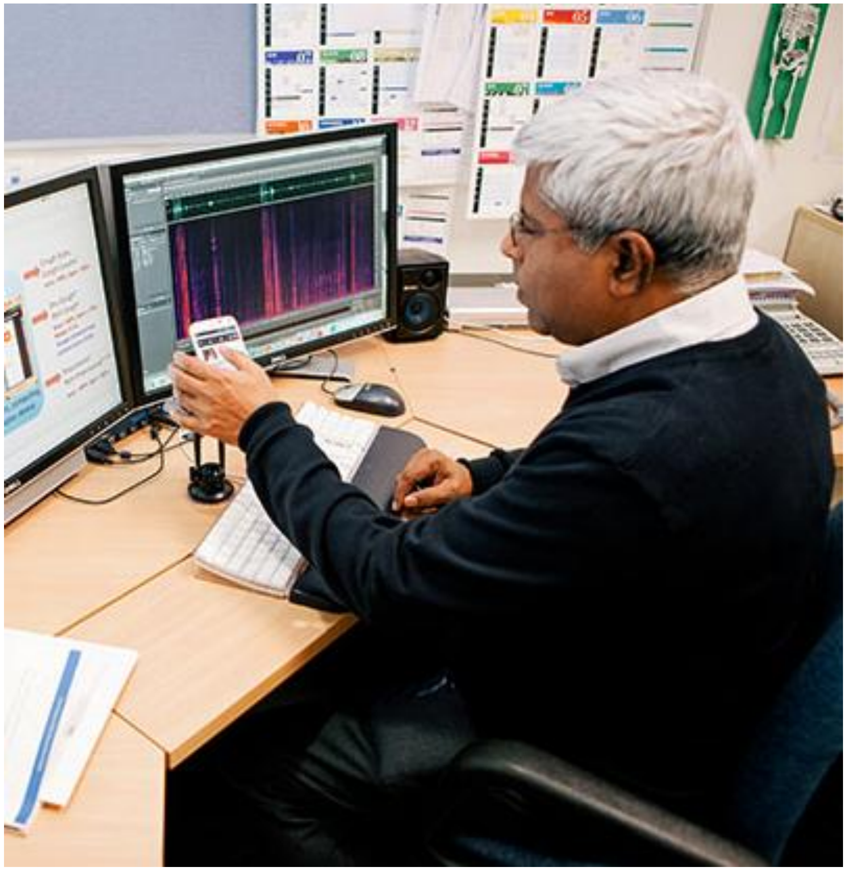

- 30 Min/Day Talking on Phone
  - 5.5 days/year of audio samples
  - Track this (notably the pauses) over time to detect onset of dimentia

- 150 Photos/Month

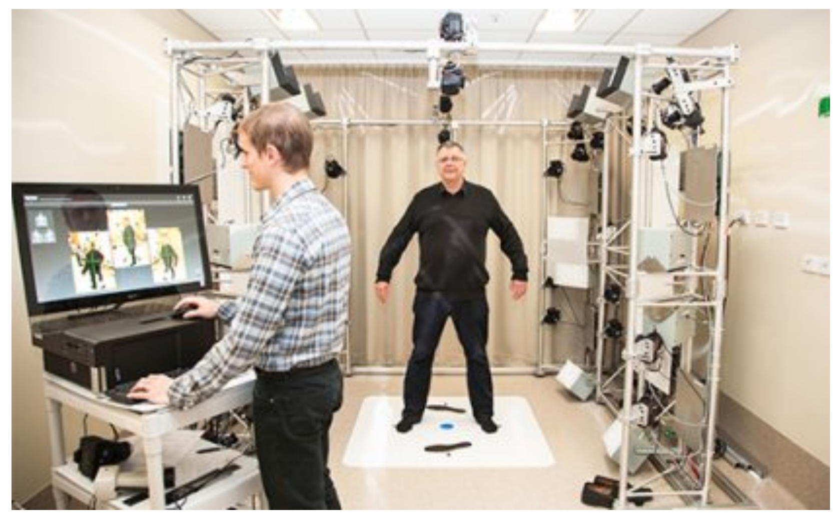

# How?

- More Signals

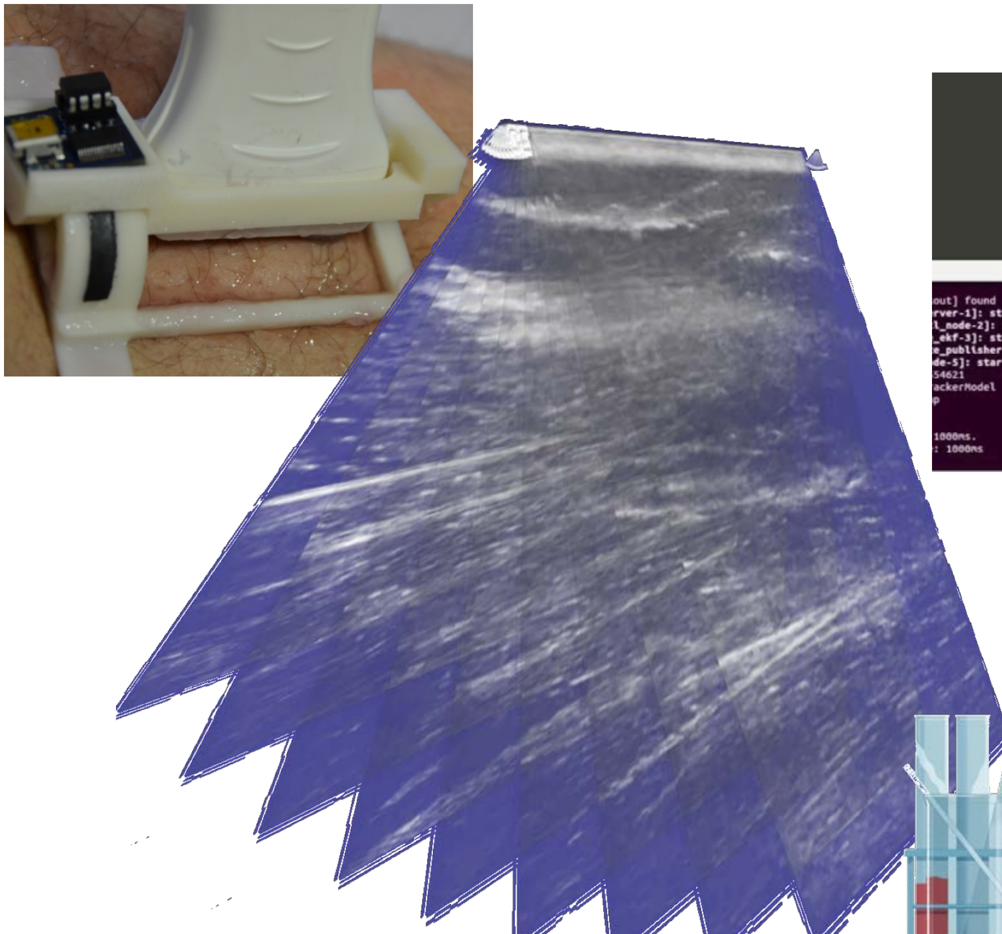

- Stochastic Processing (Think **TAPIR!**)

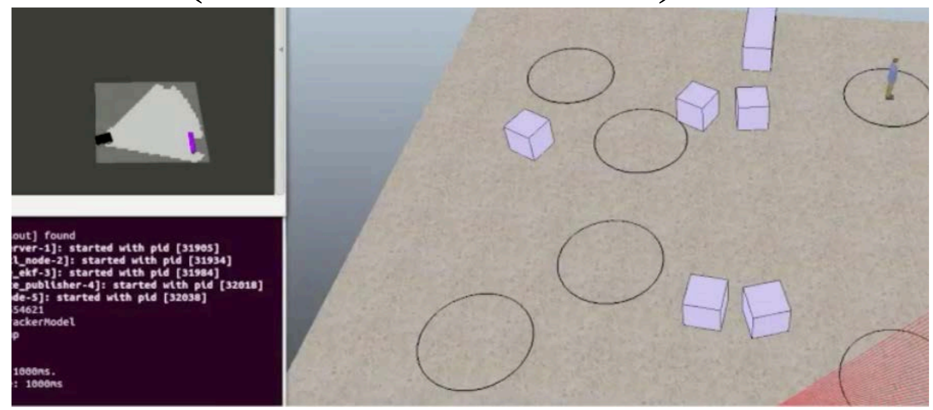


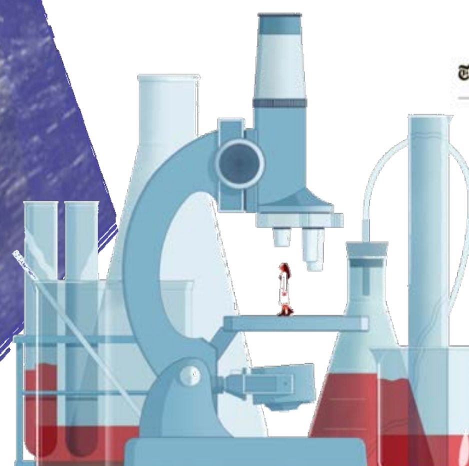

The New York Times http://nyti.ms/22UymHG

SundayReview | OPINION

## The Tampon of the Future

By PAGAN KENNEDY APRIL 1, 2016

IS it possible to extract blood from people without causing pain? For decades, this problem has stumped the medical industry. In an effort to replace the old-fashioned needle, companies are trying to deploy laser beams and tiny vacuums to draw blood.

In 2014, an engineer at Harvard named Ridhi Tariyal hit on a far simpler workaround. "I was trying to develop a way for women to monitor their own fertility at home," she told me, and "those kinds of diagnostic tests require a lot of blood. So I was thinking about women and blood. When you put those words together, it becomes obvious. We have an opportunity every single month to collect blood from women, without needles."

Together with her business partner, Stephen Gire, she has patented a method

# Robotics & Health: A Friendly Touch!

# ∴ Planning to Inform "Novel Robot" Design Strategy

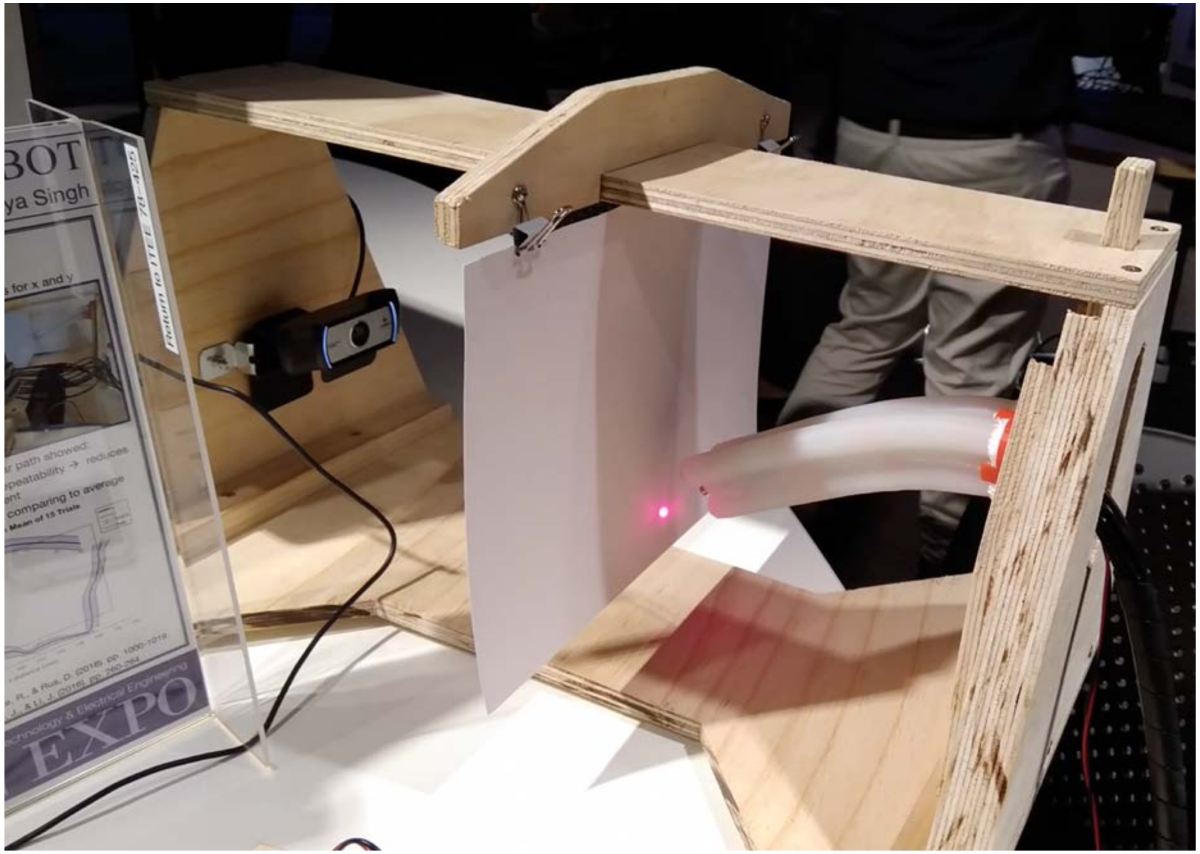

Morley-Drabble & Singh, *AIM* 2018

Ham, Singh & Lucey, *WACV* 2017

# THE
# Future of Robotics…

You!

☺

# It’s all up to you!

*If you want to build a ship, don't drum up the men to gather wood, divide the work and give orders.* ***Instead, teach them to yearn for the vast and endless sea.***

Antoine de Saint-Exupery, "The Wisdom of the Sands"

© National Geographic. Suruga Bay, Japan
“The Magic Starts Here: Kenji’s Workshop of Camera Wizardry”, December 4, 2014

# UQ Robotics: Dynamic Systems in Motion

Diverse international research group

***Pauline Pounds*** *(ANU/Yale)*
***Surya Singh*** *(Stanford/Syd)*

Aerial Systems

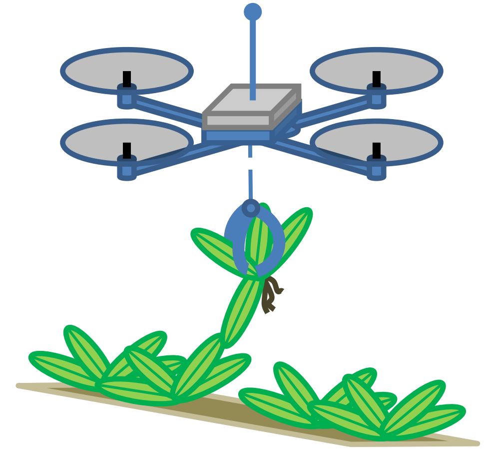

Systems

Decision-Making/
Planning

Mechanics
of motion