

Workshop Proceedings: Opportunities and Challenges of Joint Inference and Control in Mobile Robotics

Daniel Koditschek, Vijay Kumar, Stefano Soatto, Avik De

Full-day workshop on May 31, 2014 (Sat) at ICRA 2014

Contents

1 Introduction	1
1.1 Objectives of this Workshop	1
1.2 Intended Audience	2
2 Invited Talks	2
2.1 Dynamic Belief States and Information-Theoretic Decision Making in Adversarial Environments // Daniel Lee (University of Pennsylvania)	2
2.2 Bayesian Path Planning for Learning Spatial-temporal Processes // Fabio Ramos (University of Sydney)	7
2.3 Autonomous Exploration of Large Scale Natural Environments // Stefan Williams (ACFR, Sydney)	12
2.4 Active Inference of Representations: Control's Role in Visual Perception and Vice-versa // Stefano Soatto (UCLA)	12
2.5 Designing Efficient Low-latency Sensorimotor Control // Andrea Censi (MIT)	12
2.6 Information-based, Multi-target Localization Using Small Teams of Mobile Sensors // Philip Dames (University of Pennsylvania)	35
2.7 Toward Dynamical Sensor Management: Reactive Wall-following on RHex // Avik De (University of Pennsylvania)	40
2.8 The Impact of Perception Capabilities on Agile Robot Motion: A Statistical Mechanics Perspective // Sertac Karaman (MIT)	44

1 Introduction

1.1 Objectives of this Workshop

Progress in robotics has yielded physical platforms that are slowly becoming more competent in the unstructured physical world. These growing capabilities raise the prospects for realizing their huge potential value as autonomous sensors: for survivors in search and rescue settings; for anomalies or threats in surveillance applications; for novelty or scientifically motivated collection in extraterrestrial exploration. At the same time, sensor technology has taken off, and the materials, communications and computational technology underlying its advance now raise the prospect of data torrents so vast that they cannot even be reasonably stored, much less processed and interpreted without some active, real-time interpretive control. Whether in electromagnetic ranging, electro-optical (laser, cameras), or, chemotaxis modalities, the multiplicity of tuning affordances and resulting highly variable focus of attention invites and even demands that algorithms for autonomous sensor management move into the realm of real time feedback control. Addressing such control problems raises novel questions of how to formulate tasks whose goals have as much to do with the agents state of information as with its material situation; how to couple internal variables such as belief state with physical degrees of freedom; and how to develop new representations that facilitate that integration and promote the expression of information-sensitive mechanical goals. This day-long workshop will sample the range of new opportunities, questions, and issues that arise as sensors become robots, and robots become sensors. New sensing modalities such as chemo-sensitive nanoscale devices

raise the prospect of unparalleled access to perceptual domains long the unique province of animals: do we know how to use them? Traditionally “high-end” modalities such as radar have been transformed both regarding cost (in footprint and dollars) as well as realtime tunability by the advance of electronics and computation: can the sophisticated offline designs that emerged over nearly a century of waveform and receiver engineering be adapted for closed loop operation on mobile robots? Decades following the initial push for active vision in robotics, what is today’s state of the art, what theoretical insights have emerged, with what implications for practice, and how close to realtime implementation? Even assuming a nicely adaptive and computationally tractable sensorium, how should information-sensitive tasks be formulated to express the appropriate tradeoff between exploration and exploitation? How should strategic operation shift this tradeoff in the face of adversarial environments? How does a “distributed body” enhance or complicate the opportunities for joint inference and control over the sensorium? How does an imperfectly actuated body subject to a highly irregular, unpredictable environment support and benefit from the tunable sensorium?

1.2 Intended Audience

We target robotics researchers working in the traditional area of active sensing as well as experts in technology and policy seeking to understand emerging opportunities for multidisciplinary advances bearing upon robotics. There has been a great deal of interest in this topic arising from various research communities and so we have a very full day of speakers planned. The format for the workshop would be roughly one dozen 20 min individual talks (e.g., two 1.5 hr sessions in the morning and afternoon respectively) followed by a panel discussion with audience participation at the end of the day.

2 Invited Talks

2.1 Dynamic Belief States and Information-Theoretic Decision Making in Adversarial Environments

Speaker Daniel Lee (University of Pennsylvania)

Abstract The need to properly account for uncertainty in sensing has resulted in the recent interest in robotic applications of probabilistic inference techniques. In these algorithms, the role of maintaining a dynamic belief state which describes the distribution over potential states as they evolve in time is critical. These belief states can then be used as inputs to policies that attempt to choose optimal actions. I will discuss recent computational approaches to handling the unbounded dimensionality of these belief states. I will also show how recent information-theoretic approaches to bounded rationality can be interpreted as optimal stochastic policies in an adversarial environment.


Dynamic Belief States and Bounded Rationality

Daniel D. Lee
Pedro Ortega



Belief state

Distribution over possible poses:




Measurement update (Bayes rule): $p(s'_t) \propto p(s_t)p(o_t | s_t)$

Motion update (convolution): $p(s_{t+1}) = \int p(s_t)p(s_{t+1} | s_t, u_t) ds_t$

- ◆ Probabilistic filter of state over time

Sensory data




- ◆ Noisy high-dimensional signals from a variety of multimodal sources

Kalman filter

Multidimensional Gaussian: $p(s_t) = N(\mu_t, C_t)$


Measurement update: $p(s'_t) \propto p(s_t)p(o_t | s_t)$

Gaussians are closed under multiplication.




Motion update: $p(s_{t+1}) = \int p(s_t)p(s_{t+1} | s_t, u_t) ds_t$

Gaussians are closed under convolution.



- ◆ With Gaussians, just need to keep track of changes to mean and covariance (or inverse covariance)

States




Pose state:
 $s_t = (x_t, y_t, \theta_t)$

Map: $\bar{m} = (m_1, m_2, \dots)$


Dynamics model:

$$s_{t+1} = f(s_t, u_t) + \eta_t$$



Measurement model:

$$o_t = g(s_t, \bar{m}) + \eta'_t$$



Particle filter

$$p(s) \approx \sum_i w_i \delta(s - s_i)$$

Effective number of particles: $\bar{N} = \frac{1}{\sum_i w_i^2}$

Resampling:

$$\sum_i w_i \delta(s - s_i) \rightarrow \sum_i \frac{1}{N} \delta(s - s'_i)$$

- ◆ Sample-based approach to approximate arbitrary distribution function

Rao-Blackwellized pose filter


Factor belief state: $p(x, y, \theta) = p(\theta)p(x, y | \theta)$

Discrete orientation samples:

$$p(\theta) = \frac{1}{N} \sum_i \delta(\theta - \theta_i)$$

Conditional Gaussian likelihood:

$$p(x, y | \theta_i) = N(x, y | \theta_i)$$




- ◆ Hybrid filter combines flexibility of particle filter with efficiency of Kalman filter

Optimal policy

Maximize expected utility:

$$p^*(x) = \arg \max_{p(x)} \sum_x p(x) U(x)$$



Deterministic policy:



$$p_i^* = \begin{cases} 1 & \text{if } U_i \geq U_j \\ 0 & \text{otherwise} \end{cases}$$

- ◆ With full control of outcome, deterministic policy is always optimal (also for MDPs)

Utility functions


Outcomes: $X = \{x_i\}$

Utility function: $U : X \rightarrow R$



- ◆ Utility of consumption of an item in economics

Relative entropy



$$F = \langle U \rangle + \frac{1}{\beta} S$$

$$S = - \sum_x p(x) \log \frac{p(x)}{q(x)}$$


Optimal policy:

$$p^*(x) = \arg \max_{p(x)} F \propto q(x) e^{\beta U(x)}$$

- ◆ Optimal stochastic policy at finite temperature

Expected utility

Probabilistic outcomes: $p(x_i)$



Expected utility:

$$\langle U \rangle = \sum_i p(x_i) U_i$$

- ◆ Can incorporate risk aversion in expected utility


Adversarial interpretation

$$\begin{aligned} & \max_p \sum_x p(x) U(x) - \frac{1}{\beta} p(x) \log \frac{p(x)}{q(x)} \\ & = \max_p \min_C \sum_x p(x) [U(x) - C(x)] + q(x) e^{\beta C(x)} \\ & \quad + const \end{aligned}$$

Adversarial environmental costs: $C(x)$

- ◆ Interpretation of free energy as min-max game

Legendre-Fenchel transform




Convex conjugate: $f^*(s) = \max_x [s \cdot x - f(x)]$

Adversarial costs

$$F = \max_p \min_C \sum_x p(x) [U(x) - C(x)] + q(x) e^{\beta C(x)}$$

Costs: $C(x)$

Actual Utility: $U(x) - C(x)$



- ◆ Environmental adversary can change utility by subtracting costs with diminishing returns

KL divergence conjugate

$$-\frac{1}{\beta} \sum_x p(x) \log \frac{p(x)}{q(x)}$$

$$= \min_{C(x)} \sum_x -p(x)C(x) + q(x)e^{\beta C(x)} - \frac{1}{\beta}(\log \beta + 1)$$

Optimal condition: $\frac{\partial}{\partial C(x)} : p(x) = \beta q(x) e^{\beta C(x)}$


- ◆ Convex conjugate of KL-divergence regularization

Minimax iterations

Policy:



Utility-Cost:



Cost constraint: $\sum_x q(x) e^{\beta C(x)} = 1$

- ◆ Policy and cost iterations in min-max game

Cost constraint

$$\min_C \sum_x -p(x)C(x) + q(x)e^{\beta C(x)}$$

Or as cost constraint: $\sum_x q(x) e^{\beta C(x)} = 1$

Optimal cost allocation: $C^*(x) = \frac{1}{\beta} \log \frac{p(x)}{q(x)}$

- ◆ Adversary cannot make costs arbitrary large

Dual equilibrium



$U(x) - C(x) = \text{const}$

$U^* = U(x) - C(x)$

Cost constraint: $\sum_x q(x) e^{\beta(U(x) - U^*)} = 1$

$\exp(\beta U^*) = \sum_x q(x) e^{\beta U(x)} = Z$

$C^*(x) = U(x) - \frac{1}{\beta} \log \sum_x q(x) e^{\beta U(x)}$

Discussion

- ◆ Sensory data and belief states
- ◆ Belief state representations
- ◆ Optimal policies with beliefs
- ◆ Adversarial interpretation of free energy
- ◆ Exponential cost penalty
- ◆ Dual interpretation

2.2 Bayesian Path Planning for Learning Spatial-temporal Processes

Speaker Fabio Ramos (University of Sydney)

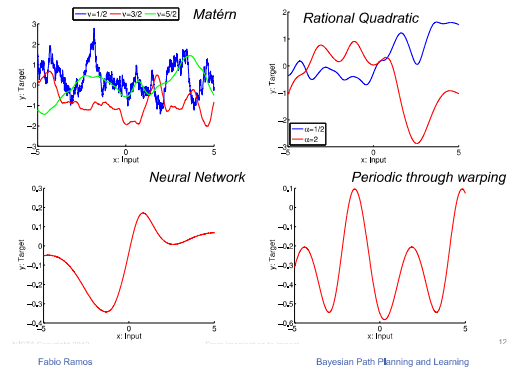
Abstract In this talk I will present a novel technique for learning spatial-temporal environment processes such as air pollution or wind speed with a mobile robot. The method is based on Bayesian optimisation and is able to select paths that maximise the prediction performance for processes where tracking peaks is crucial (such as air pollution), trading exploration-exploitation in a principled statistical manner. I will show applications in air pollution monitoring, vibration modelling while navigating on uneven terrains, and lightening changes to illustrate the benefits of the approach. Finally I will show a web-based app built for the Environmental Protection Agency in Australia for real-time air pollution forecast in the Hunter Valley region, where coal mines, urban centres and vineyards need to coexist.

Bayesian Path Planning for Environment Modelling of Spatial-Temporal Processes

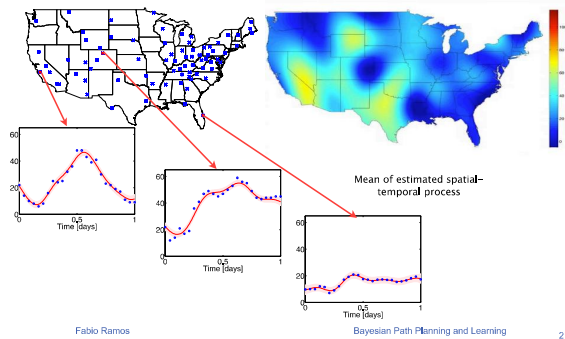
Fabio Ramos



Nonparametric models



Pollution monitoring



Gaussian process regression

Given a set of samples

$$X = [x_0, \dots, x_{N-1}]^T$$

$$y = [y_0, \dots, y_{N-1}]^T$$

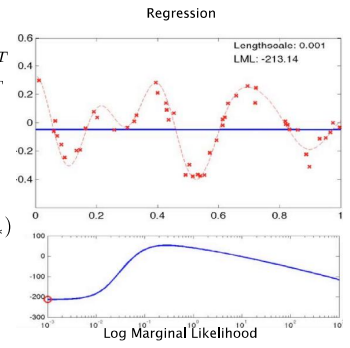
Choose a covariance function

$$k(x_i, x_j | \theta)$$

Predict the value of $f(x_*)$


Mean $\mu(f(x_*))$

Variance $\sigma(f(x_*))$



Prior over functions

$$\text{posterior} \propto \text{prior} \times \text{likelihood}$$




- A prior over the weights induces a prior over functions
 - e.g. smooth functions
 - Closeness in input space \rightarrow closeness in output space

Planning: Bayesian Optimisation

Find the maximum of an unknown, noisy and costly to evaluate function f . **Idea:** Choose next sampling location x by maximising an acquisition function S over the domain of the GP model of the function.

Algorithm 1 Bayesian Optimisation

- x_i : chosen sampling point at iteration i .
- s : acquisition function.
- f : unknown function.
- for $i = 1, 2, 3, \dots$ do
 - Find $x_i = \arg \max_x s(x)$
 - Acquire a sample from f at location x_i .
 - Update the GP model of f with the new sample.
- end for



Acquisition functions

Probability of Improvement

$$PI(x) = \Phi\left(\frac{\mu(x) - f(x^+) - \xi}{\sigma(x)}\right)$$

$$DPI(x|x^-) = \Phi\left(\frac{\sigma(x) - f(x^+) - \xi}{\sigma(x)}\right) + \gamma \cdot d(x, x^-)$$

Expected Improvement

$$EI(x) = \sigma(x) [Z\Phi(Z) + \phi(Z)]$$

$$DEI(x|x^-) \triangleq \sigma(x) [Z\Phi(Z) + \phi(Z)] + \gamma \cdot d(x, x^-)$$


$$Z = \frac{\mu(x) - f(x^+) - \xi}{\sigma(x)}$$

Upper Confidence Bound

$$UCB(x) \triangleq \alpha \cdot \mu(x) + \kappa \cdot \sigma(x)$$

$$DUCB(x|x^-) \triangleq \alpha \cdot \mu(x) + \kappa \cdot \sigma(x) + \gamma \cdot d(x, x^-)$$

Introduce **penalty on distance** to last sampled location x^-
With $\gamma < 0$




[Brochu, Cora, Freitas 2010]

Avoiding high vibration


[Souza, et al, ICRA 2014]



(a) Map matching



(b) Mean



(c) Variance

The mean of the vibration estimate shows a clear distinction between two explored terrains, grass and asphalt.

Ozone concentration

Mean of Estimation at $t=0,62$


IG (Information Gain)

UCB

DUCB




Trajectories



RMSE 13.78
Distance 717.18

13.04
718.25

12.33
639.06



RED = Goal areas
BLUE = Avoid sampling here

Indoor ST-Luminosity model

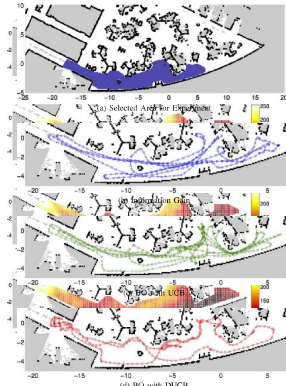


TABLE II
RESULTS FOR REAL EXPERIMENT

Indicator	Method	Mean	Distance [m]
RMSE	IG	16.44	157.01
RMSE	UCB	17.16	217.45
RMSE	DUCB	16.91	98.08
WRMSE	IG	5.41	157.01
WRMSE	UCB	4.60	217.45
WRMSE	DUCB	4.70	98.086



[Marchant, Ramos, IROS 2012]
Bayesian Path Planning and Learning

BO for continuous paths

We use parametrised curves


$$C_1(u) = a_x u^3 + b_x u^2 + c_x u + d_x$$

$$C_2(u) = a_y u^3 + b_y u^2 + c_y u + d_y$$

$$C_1(u=0) = p_{x_i} = d_x,$$

$$C_2(u=0) = p_{y_i} = d_y,$$

$$\left. \frac{\partial C_2 / \partial u}{\partial C_1 / \partial u} \right|_{u=0} = \frac{c_y}{c_x} = \tan \alpha_i.$$



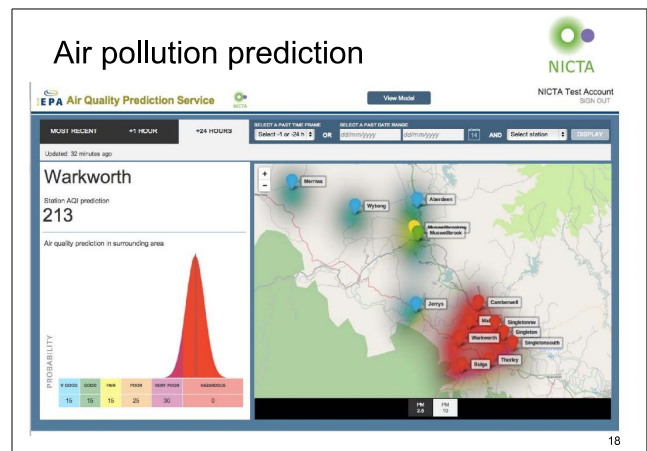
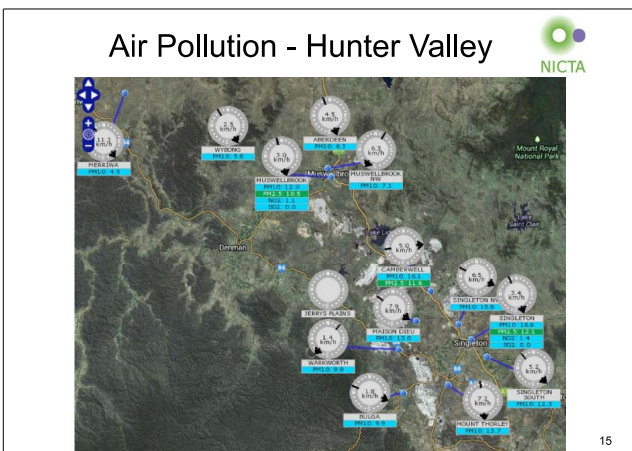
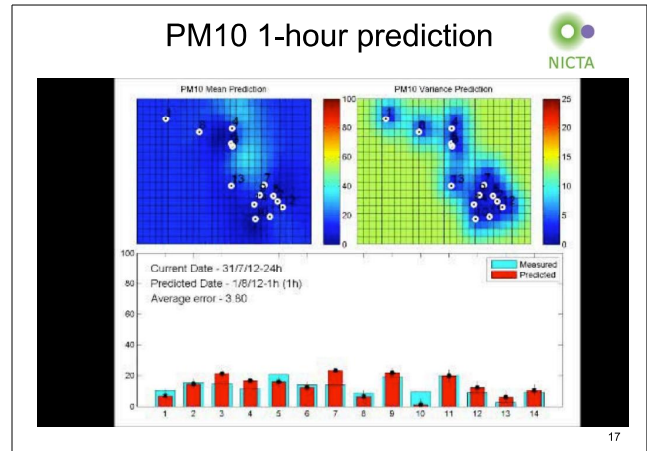
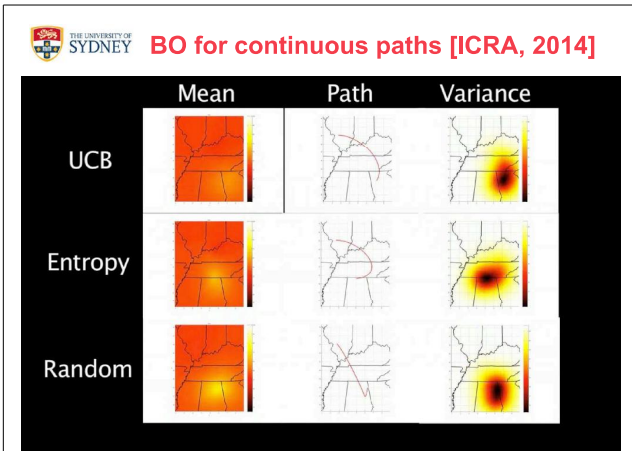
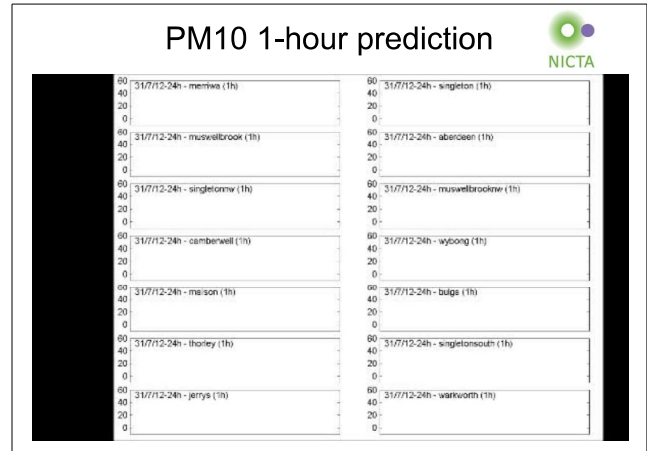
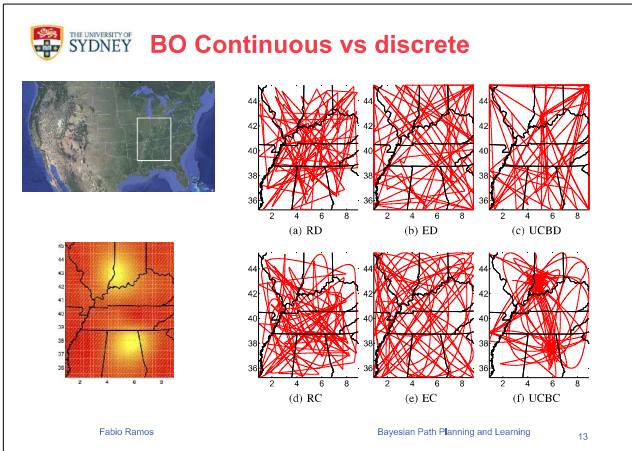
Modified cost for sampling over paths

[Marchant, Ramos, ICRA 2014]

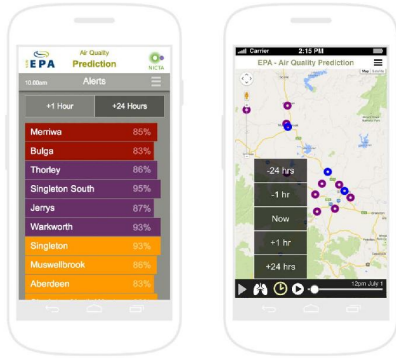
$$r(C(u, \beta)|h) = \int_{C(u, \beta)} h(v) dv$$

$$= \int_0^1 UCB(C(u, \beta)) \|C'(u, \beta)\| du$$

$$= \int_0^1 [\mu(C(u, \beta)) + \kappa \sigma(C(u, \beta))] \|C'(u, \beta)\| du.$$



Air pollution App



19



Summary

- We model air pollution as a spatial temporal model that is incrementally built.
- We use a BO approach for choosing sampling locations, automatically trading off exploration and exploitation.
- We propose a continuous generalisation for BO for general path planning problems.
- There are many applications, and a lot of industry and government interest.

2.3 Autonomous Exploration of Large Scale Natural Environments

Speaker Stefan Williams (ACFR, Sydney)

2.4 Active Inference of Representations: Control’s Role in Visual Perception and Vice-versa

Speaker Stefano Soatto (UCLA)

Abstract The “state” of a system or agent, understood as a function of measured data that is “useful” towards a control or decision task, should ideally “separate” sensing and control: Sensing would infer the function of all past data that is “sufficient” and hand it off to a control or decision module – agnostic of how the state or “representation” is inferred – to accomplish the task. While this is indeed possible for linear systems in Gaussian noise, complex sources of uncertainty make the separation imperfect if not impossible. Specifically, when uncertainty is due to sensing mechanisms that involve occlusion and scaling – such as visual sensing, whether in the visible or other spectra – control is actually necessary to infer a state that is sufficient to accomplish even elementary decision tasks. In addition, there may be uncertainty on the task itself. In this talk, we will explore ways of formalizing the properties that an “ideal representation” should have to support a variety of decision, control and interaction tasks with physical space, where sensing is provided by visual as well as other modalities. We will then see how some drastic simplifications yields to methods that are currently in use today, and point to ways to improve them. We will show applications in visual recognition (finding a known object in an unknown environment) as well as reconstruction (building a model of the environment to support navigation tasks) exploiting visual and inertial sensors.

Speaker Bio Stefano Soatto is the founder and director of the UCLA Vision Lab (vision.ucla.edu). He received his Ph.D. in Control and Dynamical Systems from the California Institute of Technology in 1996; he joined UCLA in 2000 after being Assistant and then Associate Professor of Electrical and Biomedical Engineering at Washington University, Research Associate in Applied Sciences at Harvard University, and Assistant Professor in Mathematics and Computer Science at the University of Udine, Italy. He received his D.Ing. degree (highest honors) from the University of Padova- Italy in 1992. Dr. Soatto is the recipient of the David Marr Prize (with Y. Ma, J. Kosecka and S. Sastry) for work on Euclidean reconstruction and reprojection up to subgroups. He also received the Siemens Prize with the Outstanding Paper Award from the IEEE Computer Society for his work on optimal structure from motion (with R. Brockett). He received the National Science Foundation Career Award and the Okawa Foundation Grant. He is a Member of the Editorial Board of the International Journal of Computer Vision (IJCV), the International Journal of Mathematical Imaging and Vision (JMIV) and Foundations and Trends in Computer Graphics and Vision.


2.5 Designing Efficient Low-latency Sensorimotor Control

Speaker Andrea Censi (MIT)

Designing efficient low-latency sensorimotor control

Andrea Censi

Laboratory for Information and Decision Systems




Monday, 12:30 — MoB09

Joint inference and control: opportunities and challenges

Designing efficient low-latency sensorimotor control

Andrea Censi


Laboratory for Information and Decision Systems



Monday, 12:30 — MoB09

Perception is solved!

Designing efficient low-latency sensorimotor control



Monday, 12:30 — MoB09

Perception is solved!



- ▶ As a robotics researcher, you shouldn't compete with people doing "passive" perception.

perception **in** robotics

- ▶ As a robotics researcher, you shouldn't compete with people doing "passive" perception.



Zhou, Koltun. *Color Map Optimization for 3D Reconstruction with Consumer Depth Cameras*, SIGGRAPH 2014.

perception **in** robotics



perception **for** robotics

- ▶ As a robotics researcher, you shouldn't compete with people doing "passive" perception.



Zhou, Koltun. *Color Map Optimization for 3D Reconstruction with Consumer Depth Cameras*, SIGGRAPH 2014.

What is Robotics?

What is Robotics?

1. The business of adapting cool techniques in other fields to obtain a cute demo with a robot.

What's embodied intelligence about?


- ▶ It's (also) about doing well in the world using **limited resources**.

What is Robotics?

1. The business of adapting cool techniques in other fields to obtain a cute demo with a robot.
2. The scientific quest of understanding and replicating embodied intelligence.

What's embodied intelligence about?


- ▶ It's (also) about doing well in the world using **limited resources**.

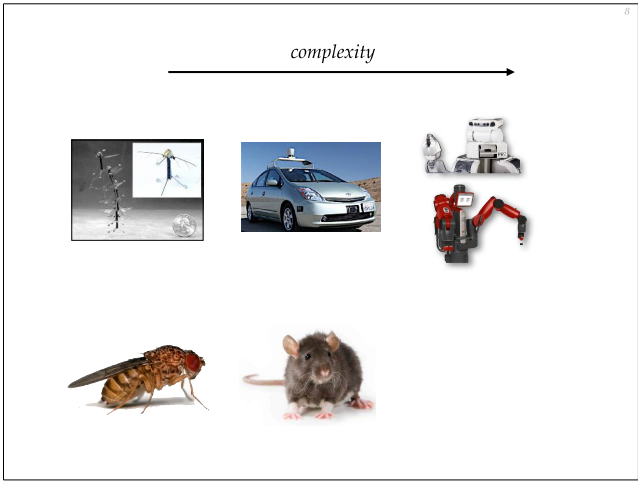
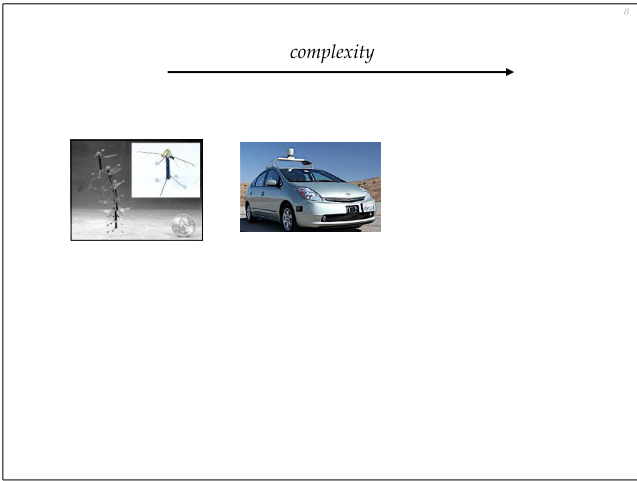
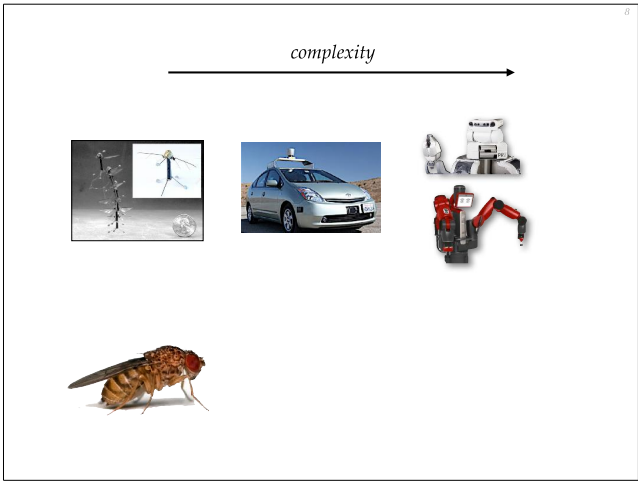
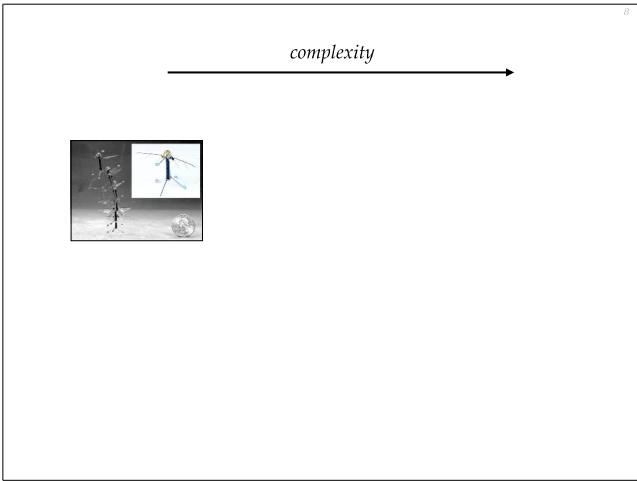
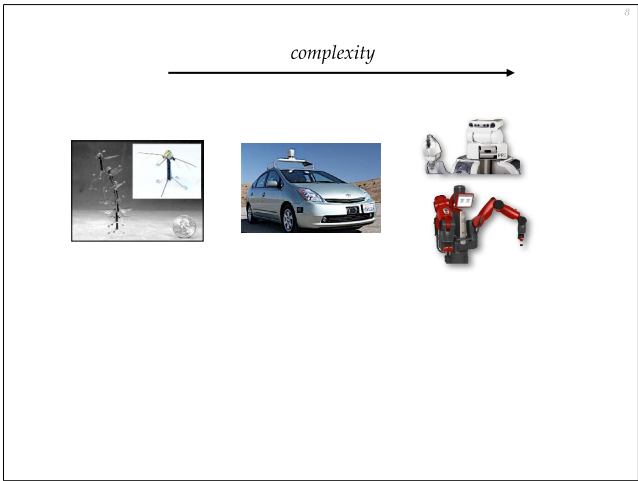



What's embodied intelligence about?

What's embodied intelligence about?

- ▶ It's (also) about doing well in the world using **limited resources**.





complexity →

Doing well with limited resources

► Here's a task T; X watts of power; and Z bytes of memory.
Design something that gives a reasonable answer in Y seconds.

$X = 500 \text{ watts}$
 $Y = 100 \text{ milliseconds}$

Doing well with limited resources

Doing well with limited resources

► Here's a task T; X watts of power; and Z bytes of memory.
Design something that gives a reasonable answer in Y seconds.

$X = 50 \text{ milliwatts}$
 $Y = 1 \text{ millisecond}$

$X = 500 \text{ watts}$
 $Y = 100 \text{ milliseconds}$

Doing well with limited resources


► Here's a task T; X watts of power; and Z bytes of memory.
Design something that gives a reasonable answer in Y seconds.

Joint inference and control: opportunities and challenges


Joint inference and control: opportunities and challenges

solving the joint problem
is more resource-efficient


- ▶ The world/plant is a causal black box from u to y .




- ▶ We need to design an agent/controller as a causal black box from y to u .




- ▶ Markov assumption: S are the states




- ▶ The world/plant is a causal black box from u to y .




- ▶ Markov assumption: S are the states




- ▶ Markov assumption:
 S are the states




- ▶ A **(deterministic) agent** is a tuple (Γ, f, g) where Γ is any set representing the agent memory; $f: \Gamma \times \mathcal{Y} \rightarrow \Gamma$ defines the memory dynamics;




- ▶ A **(deterministic) agent** is a tuple (Γ, f, g)




- ▶ A **(deterministic) agent** is a tuple (Γ, f, g) where Γ is any set representing the agent memory; $f: \Gamma \times \mathcal{Y} \rightarrow \Gamma$ defines the memory dynamics; $g: \Gamma \times \mathcal{Y} \rightarrow \mathcal{U}$ is the memory-to-command map.



- ▶ A **(deterministic) agent** is a tuple (Γ, f, g) where Γ is any set representing the agent memory;




- ▶ The "canonical" probabilistic agent:



► The “canonical” probabilistic agent:

Γ = beliefs (probability distributions on world’s state)




► The “canonical” probabilistic agent:

Γ = beliefs (probability distributions on world’s state)

γ_k = belief about world’s state

f = Bayesian filter


g = solver of a POMDP



► The “canonical” probabilistic agent:

Γ = beliefs (probability distributions on world’s state)

γ_k = belief about world’s state




► The ^{realistic} “~~canonical~~” probabilistic agent:

Γ = beliefs (probability distributions on world’s state)

γ_k = belief about world’s state

f = Bayesian filter

g = solver of a POMDP




► The “canonical” probabilistic agent:

Γ = beliefs (probability distributions on world’s state)

γ_k = belief about world’s state

f = Bayesian filter



► The ^{realistic} “~~canonical~~” probabilistic agent:


Γ = beliefs (probability distributions on world’s state)

γ_k = belief about world’s state

f = Bayesian filter

^{approximated}

g = solver of a POMDP



15

realistic
 ▶ The **"canonical"** probabilistic agent:

Γ = beliefs (probability distributions on world's state)
 γ_k = belief about world's state

f = Bayesian filter
 approximated

g = solver of a POMDP
 approximated

16

▶ The "canonical" probabilistic agent:

Γ = beliefs (probability distributions on world's state)
 γ_k = belief about world's state

alternatives?

15

realistic
 ▶ The **"canonical"** probabilistic agent:

Γ = beliefs (probability distributions on world's state)
 γ_k = belief about world's state

f = Bayesian filter
 approximated

g = solver of a ~~POMDP~~ stochastic optimal control problem
 approximated

17

**Joint inference and control:
 opportunities and challenges**

↙

solving the joint problem
 is more resource-efficient

15

realistic
 ▶ The **"canonical"** probabilistic agent:

Γ = beliefs (probability distributions on world's state)
 γ_k = belief about world's state

f = Bayesian filter
 approximated

g = solver of a ~~POMDP~~ stochastic optimal control problem
 approximated based on certainty-equivalence

17

**Joint inference and control:
 opportunities and challenges**

↙

solving the joint problem
 is more resource-efficient

↓

There are many formalizations
 (only partially compatible)

17

Joint inference and control: opportunities and challenges

↓

solving the joint problem
is more resource-efficient

↓

There are many formalizations
(only partially compatible)

1. ...
2. ...
3. ...
4. ...
5. ...

18

Doing well with limited resources

1. Find an **optimal agent** that uses the **fewest resources**.
2. Find a **suboptimal agent** with given **resources bounds**.

18

Doing well with limited resources

19

Offline design vs online execution

18

Doing well with limited resources

1. Find an **optimal agent** that uses the **fewest resources**.

19

Offline design vs online execution

problem spec


↓

synthesis algorithm

↓


agent description

Offline design vs online execution




1. Minimality of sensing / control

- ▶ What can you do with minimal sensing / control?




Offline design vs online execution




1. Minimality of sensing / control

- ▶ What can you do with minimal sensing / control?

O'Kane, LaValle. *On comparing the power of robots*. IJRR 2008
Localization with limited sensing. TRO 2007



1. Minimality of sensing / control




1. Minimality of sensing / control

- ▶ What can you do with minimal sensing / control?


O'Kane, LaValle. *On comparing the power of robots*. IJRR 2008
Localization with limited sensing. TRO 2007

- ▶ Sensing data is very redundant for place recognition




1. Minimality of sensing / control

- ▶ What can you do with minimal sensing / control?
O'Kane, LaValle. *On comparing the power of robots*. IJRR 2008
Localization with limited sensing. TRO 2007
- ▶ Sensing data is very redundant for place recognition




Milford. *Vision-based place recognition: how low can you go?* IJRR 2013



3. Penalizing the control information

2. Penalizing the cost of computation

Ortega, Braun. *Thermodynamics as a theory of decision-making with information-processing costs*, 2013
Braun, Ortega, Theodorou, Schaal. *Path Integral Control and Bounded Rationality*, 2011




3. Penalizing the control information




2. Penalizing the cost of computation

Ortega, Braun. *Thermodynamics as a theory of decision-making with information-processing costs*, 2013
Braun, Ortega, Theodorou, Schaal. *Path Integral Control and Bounded Rationality*, 2011




updating information state takes work
(physical work)

↓
variational problem




3. Penalizing the control information



3. Penalizing the control information

23

Rubin, Shamir, Tishby. *Trading value and information in MDPs*. 2010




3. Penalizing the control information


23

$$= \text{deviation from random policy} = \mathbb{E} \left\{ \log \frac{\pi_s(u)}{\bar{\rho}(u)} \right\}$$

\nearrow policy
 \searrow a blind random policy



Rubin, Shamir, Tishby. *Trading value and information in MDPs*. 2010




3. Penalizing the control information

23

= deviation from random policy =

Rubin, Shamir, Tishby. *Trading value and information in MDPs*. 2010



4. Minimizing the agent-world bandwidth

24


3. Penalizing the control information

23

$$= \text{deviation from random policy} = \mathbb{E} \left\{ \log \frac{\pi_s(u)}{\bar{\rho}(u)} \right\}$$

\nearrow policy
 \searrow a blind random policy

Rubin, Shamir, Tishby. *Trading value and information in MDPs*. 2010



4. Minimizing the agent-world bandwidth

24




4. Minimizing the agent-world bandwidth




24

4. Minimizing the agent-world bandwidth




24

4. Minimizing the agent-world bandwidth




24

4. Minimizing the agent-world bandwidth



25


4. Minimizing the agent-world bandwidth



24

4. Minimizing the agent-world bandwidth

Tishby, Polani. *Information Theory of Decisions and Actions*. 2011




25

4. Minimizing the agent-world bandwidth

Tishby, Polani. *Information Theory of Decisions and Actions*. 2011

"information to go" =




4. Minimizing the agent-world bandwidth

Tishby, Polani. *Information Theory of Decisions and Actions*. 2011

"information to go" = $D_{\text{KL}}(p_\pi \parallel \hat{p})$


\ distribution of states, actions under random policy
distribution of states, actions under policy



4. Minimizing the agent-world bandwidth

Tishby, Polani. *Information Theory of Decisions and Actions*. 2011

"information to go" = $D_{\text{KL}}(p_\pi \parallel \hat{p})$




4. Minimizing the agent-world bandwidth

Tishby, Polani. *Information Theory of Decisions and Actions*. 2011

"information to go" = $D_{\text{KL}}(p_\pi \parallel \hat{p})$

\ distribution of states, actions under random policy
distribution of states, actions under policy




4. Minimizing the agent-world bandwidth

Tishby, Polani. *Information Theory of Decisions and Actions*. 2011

"information to go" = $D_{\text{KL}}(p_\pi \parallel \hat{p})$

\ distribution of states, actions under random policy



4. Minimizing the agent-world bandwidth

Tishby, Polani. *Information Theory of Decisions and Actions*. 2011

"information to go" = $D_{\text{KL}}(p_\pi \parallel \hat{p})$


\ distribution of states, actions under random policy
distribution of states, actions under policy

4. Minimizing the agent-world bandwidth

Tishby, Polani. *Information Theory of Decisions and Actions*. 2011

$$\text{"information to go"} = D_{\text{KL}}(p_{\pi} \parallel \hat{p})$$

↙ distribution of states, actions under random policy
↘ distribution of states, actions under policy



5. Minimality of representation (size of agent state)



1. Find an **optimal agent** that uses the **fewest resources**.
2. Find a **suboptimal agent** with given **resources bounds**.

4. Minimizing the agent-world bandwidth

Tishby, Polani. *Information Theory of Decisions and Actions*. 2011

$$\text{"information to go"} = D_{\text{KL}}(p_{\pi} \parallel \hat{p})$$


↙ distribution of states, actions under random policy
↘ distribution of states, actions under policy



5. Minimality of representation (size of agent state)

1. Find an **optimal agent** that uses the **fewest resources**.
2. Find a **suboptimal agent** with given **resources bounds**.

Roy, Gordon, Thrun. *Finding Approximate POMDP Solutions Through Belief Compression*. JAIR 2005




(a) A common belief

(b) An unlikely belief

5. Minimality of representation (size of agent state)

- ▶ Most of the computation cost is in updating the representation.
- ▶ Penalize size of representation:

$$\min |\Gamma|$$



5. Minimality of representation (size of agent state)


1. Find an **optimal agent** that uses the **fewest resources**.
2. Find a **suboptimal agent** with given **resources bounds**.

5. Minimality of representation


31

Tovar, Guilamo, LaValle *Gap Navigation Trees: Minimal Representation for Visibility-based Tasks*. WAFR 2004

- ▶ A range-finder can be abstracted as a "gap sensor"



- ▶ Map can be represented as graphs




Q: Can we automatically synthesize minimal representations?

5. Minimality of representation

32

▶ Task: find-object

- A robot must find a static object in a known environment.




5. Minimality of representation

32

▶ Task: find-object

- A robot must find a static object in a known environment.



5. Minimality of representation


32

▶ Task: find-object

- A robot must find a static object in a known environment.

▶ Sensors:

- Camera that detects object on sight.
- Observable robot position (to be relaxed)




5. Minimality of representation

32

▶ Task: find-object

- A robot must find a static object in a known environment.



5. Minimality of representation

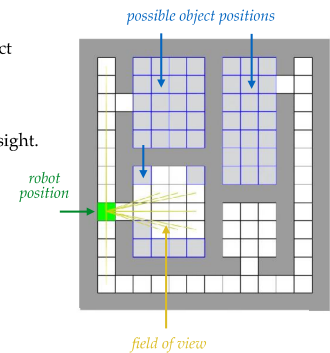
32

▶ Task: find-object

- A robot must find a static object in a known environment.

▶ Sensors:

- Camera that detects object on sight.
- Observable robot position (to be relaxed)



5. Minimality of representation

32

► **Task: find-object**


- A robot must find a static object in a known environment.

► **Sensors:**

- Camera that detects object on sight.
- Observable robot position (to be relaxed)

► **Actions:**

- move (up, down, left, right)
- declare where the intruder is



5. Minimality of representation


34

- Optimal agent only needs to represent optimally reachable beliefs.

5. Minimality of representation

33

- Formalized as POMDP.
 ► Solution obtained from the MDP in belief space.



5. Minimality of representation

34

- Optimal agent only needs to represent optimally reachable beliefs.





Fig. 1. Belief space \mathcal{B} , reachable space $\mathcal{R}(b_0)$, and optimally reachable space $\mathcal{R}^*(b_0)$. Note that $\mathcal{R}^*(b_0) \subseteq \mathcal{R}(b_0) \subseteq \mathcal{B}$.

Kurniawati, Hsu, Lee. *SARSOP: Efficient Point-Based POMDP Planning by Approximating Optimally Reachable Belief Spaces*. RSS 2008

5. Minimality of representation

33

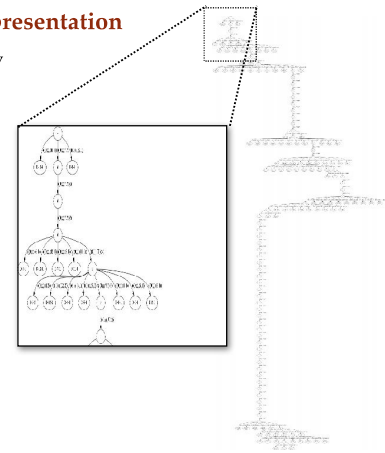
- Formalized as POMDP.
 ► Solution obtained from the MDP in belief space.



5. Minimality of representation

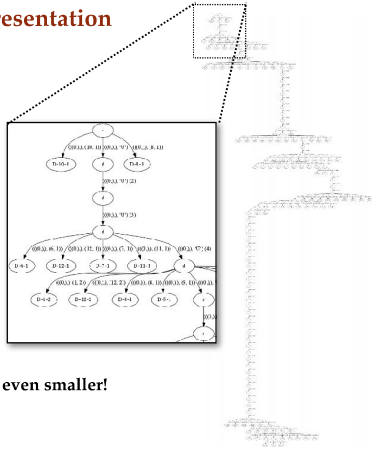
35

- "policy graph": optimally reachable beliefs and corresponding optimal commands



5. Minimality of representation

- ▶ “policy graph”: optimally reachable beliefs and corresponding optimal commands





- ▶ Minimal representation is even smaller!

$$\min |\Gamma|$$

35



- ▶ Q: What is the size of the minimal representation?
A: $|\Gamma| = 3$ states

Here's a minimal representation that we obtain automatically.



36

- ▶ Q: What is the size of the minimal representation?
A: $|\Gamma| = 3$ states





36

- ▶ The size of the agent's representation depends on the sensorium power.

more powerful ← less powerful

37

- ▶ Q: What is the size of the minimal representation?
A: $|\Gamma| = 3$ states



36

- ▶ The size of the agent's representation depends on the sensorium power.

more powerful ← less powerful

Observable robot position

37

37

► The size of the agent's representation depends on the **sensorium power**.

more powerful ← ————— → less powerful

Observable robot position Use a range-finder for localization

37

► The size of the agent's representation depends on the **sensorium power**.

more powerful ← ————— → less powerful

Observable robot position Use a range-finder for localization

horizon = 3 horizon = 2 horizon = 1

37

► The size of the agent's representation depends on the **sensorium power**.

more powerful ← ————— → less powerful

Observable robot position Use a range-finder for localization

horizon = 3

38

► The size of the agent's representation depends on the **sensorium power**.

more powerful ← ————— → less powerful

Observable robot position Use a range-finder for localization

horizon = 3 horizon = 2 horizon = 1

3 states 3 states 5 states 8 states

0 0 0 0 0 0 0 0 0 0 0 0 0

37

► The size of the agent's representation depends on the **sensorium power**.

more powerful ← ————— → less powerful

Observable robot position Use a range-finder for localization

horizon = 3 horizon = 2

38

► The size of the agent's representation depends on the **sensorium power**.

more powerful ← ————— → less powerful

Observable robot position Use a range-finder for localization

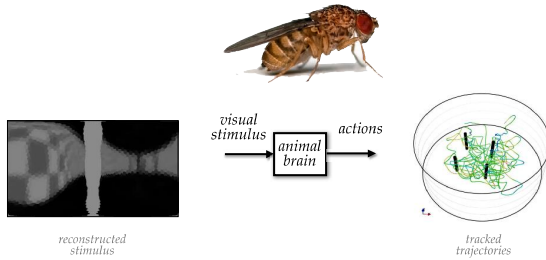
horizon = 3 horizon = 2 horizon = 1

3 states 3 states 5 states 8 states

0 0 0 0 0 0 0 0 0 0 0 0 0

5. Minimality of representation

- ▶ “What is the simplest neural process that realizes the observed behavior?”



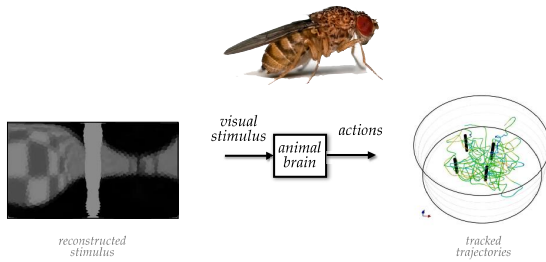
39

Joint inference and control: opportunities

solving the joint problem
is more resource-efficient

5. Minimality of representation

- ▶ “What is the simplest neural process that realizes the observed behavior?”



39

Joint inference and control: opportunities

solving the joint problem
is more resource-efficient

There are many formalizations
(only partially compatible)

Joint inference and control:

40

Joint inference and control: opportunities

solving the joint problem
is more resource-efficient

There are many formalizations
(only partially compatible)

1. Minimality of sensing / control
2. Penalizing computation
3. Penalizing control information
4. Penalizing agent-world bandwidth
5. Minimality of representation

40

40

**Joint inference and control:
opportunities and challenges**

↓

solving the joint problem
is more resource-efficient

↓

There are many formalizations
(only partially compatible)

1. Minimality of sensing / control
2. Penalizing computation
3. Penalizing control information
4. Penalizing agent-world bandwidth
5. Minimality of representation

41

**Joint inference and control:
opportunities and challenges**

↓

solving the joint problem
is more resource-efficient

↓

There are many formalizations
(only partially compatible)

1. Minimality of sensing / control
2. Penalizing computation
3. Penalizing control information
4. Penalizing agent-world bandwidth
5. Minimality of representation

Death by generality

Q: What is robotics?

Q: What's special about embodied intelligence?

41

**Joint inference and control:
opportunities and challenges**

↓

solving the joint problem
is more resource-efficient

↓

There are many formalizations
(only partially compatible)

1. Minimality of sensing / control
2. Penalizing computation
3. Penalizing control information
4. Penalizing agent-world bandwidth
5. Minimality of representation

Death by generality

Q: What is robotics?

Q: What's special about embodied intelligence?

Death by abstraction

Q: What can we integrate within realistic architectures?

2.6 Information-based, Multi-target Localization Using Small Teams of Mobile Sensors

Speaker Philip Dames (University of Pennsylvania)

Abstract There are many situations in which teams of robots can be used for active information acquisition, such as security and surveillance, infrastructure inspection, target tracking, and search and rescue. All of these scenarios share a common problem: while the types of objects of interest are known (e.g., a cell phone signal from a trapped individual) the number of such objects in the environment will not be a priori. The estimation problem is further complicated by the sensors returning false positive measurements, missing detections, and returning noisy estimates of true objects. We utilize a mathematical tool called the probability hypothesis density (PHD) filter that allows us to simultaneously estimate the number of objects in the environment and their positions while dealing with imperfect sensors. Then using the resulting estimate of the target set, the robot team follows a receding horizon, information-based control law which maximizes the mutual information between the target set and the binary event of getting no target detections, effectively hedging against non-informative actions in a computationally tractable manner.

In some of these scenarios, such as surveillance, the robot team operates in an environment with existing communication infrastructure. In such instances, the robots may leverage that infrastructure to quickly disseminate information across the robot team without requiring direct peer-to-peer links to other robots in the team. In this case we model the information trade off between directly taking measurements of objects in the environment and receiving measurements through communication channels with base stations.

Suggested reading [[EWBS09](#), [VSD03](#), [HT10](#), [Pul05](#), [Gro06](#), [DK13a](#)]

Information-based, Multi-target Localization Using Small Teams of Mobile Sensors

Philip Dames and Vijay Kumar
31 May 2014

Challenges


- Unknown number of objects a priori
- Imperfect sensors
 - False positive and false negative detections
 - Limited field of view (FoV)
 - Noisy measurements
- Potentially limited communication range
- Unknown data association

31 May 2014

4

Problem Statement

- Goal: determine the number of objects of interest and their locations within the environment using a small team of mobile robots



31 May 2014

Assumptions

- Obstacle map is known
- Robots are able to self-localize
- Targets are stationary
- There may be a base station with one or more access points in the environment

31 May 2014

5

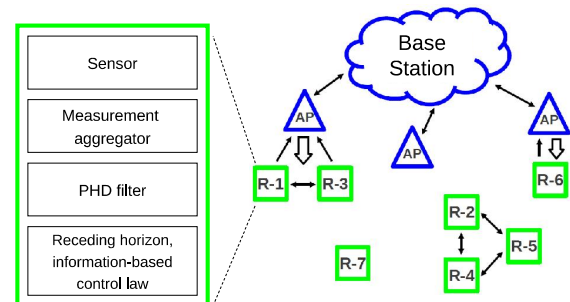
Example Scenarios

- Search and rescue
 - Localize trapped individuals via cell phone signals
- Infrastructure inspection
 - Localize sensors in a smart building
- Environmental monitoring
 - Search for a species of interest

31 May 2014

3

System Overview



31 May 2014

6

Related Work

- DDF
 - Cooperative search [Grocholsky 2002], [Cole 2009]
- Receding horizon control
 - Survey [Mayne and Michalska 1990], [Mayne, et al 2000]
 - Target tracking [Ryan 2009]
- Information-based control
 - Target tracking [Grocholsky 2002], [Bourgault, et al 2002]
 - Tractable information approximations for target localization [Hoffmann, Tomlin 2010], [Charrow, et al 2013]
- Control using finite set statistics
 - Renyi divergence [Ristic, Vo 2010], [Ristic, Vo, Clark 2011]
 - MI gradient [Schwager, et al 2011], [Dames, et al 2012]

31 May 2014

7

PHD Filter Assumptions

- For computational tractability, [Mahler 2003] assumes
 - New/existing targets independent
 - Targets i.i.d.
 - Targets move independently
 - Targets generate measurements independently
 - Clutter i.i.d. and independent of measurements
 - Number of targets, clutter is Poisson

31 May 2014

10

Sensing – Range-only

- Robot pose q
- Detection model


$$p_d(x; q) = \begin{cases} f(\|x - q\|_2) & \|x - q\|_2 \leq d_{\max} \\ 0 & \text{else} \end{cases}$$

- Measurement model

$$g(z | x; q) = \mathcal{N}(z; \|x - q\|_2, \sigma_g^2)$$

- Clutter model

$$\kappa(z) = \begin{cases} \frac{1}{d_{\max}} & 0 \leq z \leq d_{\max} \\ 0 & \text{else} \end{cases}$$



31 May 2014

8

PHD Filter

- Update equation

$$v^t(x) = (1 - p_d(x))v^{t-1}(x) + \sum_{z \in Z^t} \frac{p_d(x)g(z | x)}{\kappa(z) + \int p_d(\xi)g(z | \xi)v^{t-1}(\xi) d\xi} v^{t-1}(x)$$

- Low complexity updates

- Linear in number of particles and measurements
- Avoids data association

31 May 2014

11

Probability Hypothesis Density

- Probability hypothesis density (PHD) $v(x)$
 - Target density field over environment
 - Integral over any region is expected number of targets in that region
- Typically implemented as:
 - Weighted particle set [Vo, Singh, Doucet 2003]
 - GMM [Vo, Ma 2006]

31 May 2014

9

Joint Estimation

- Robots and base station maintain their own filters
- Robots exchange measurement sets
- Measurement aggregator ensures robot doesn't double count information

31 May 2014

12

Control Law

- Maximize mutual information between target set and binary detection events

$$q^* = \arg \max_{q \in \mathcal{Q}} I[X, Y; q]$$

- Using full measurement sets is computationally intractable

- Use binary approximation $Y = \begin{cases} 0 & Z = \emptyset \\ 1 & \text{else} \end{cases}$

31 May 2014

13

Information Value of Communication

- Robots can receive measurements by:
 - Directly taking them
 - Communicating with a base station

- Updated control law

$$q^* = \max_{q \in \mathcal{Q}} (\arg \max_{q \in \mathcal{Q}} I[X, Y; q], I[X, Y_{\text{comm}}])$$

- If second term is higher, robot visits base station

31 May 2014

16

Joint Control

- Split team into coalitions
 - Members of a coalition have overlapping sensor FoV

- Computational complexity

$$\mathcal{O}((PRT + 2^{RT})2^{RT} A^R)$$

- A = # of actions per robot
- P = # of particles in PHD
- R = # of robots in coalition
- T = length of planning horizon

31 May 2014

14

Base Station Data

- Data available at the base station, Y_{comm} , is unknown until robot reaches the access point


- Number of new measurements
 - Assume geometric rate of return to base station
- Locations measurements taken
 - Average over possible locations
- Correlations between measurements
 - Measurements independent

31 May 2014

17

Trajectory Generation


- For each robot:
 - Select all points a specified distance away
 - Remove points sufficiently close together
 - Maintain action diversity while reducing complexity
 - Interpolate paths at a fixed distance to get waypoints




31 May 2014

Example Maps


- 3 robots, 3 step horizon, infinite communication range



(a) Environment 1



(b) Environment 2

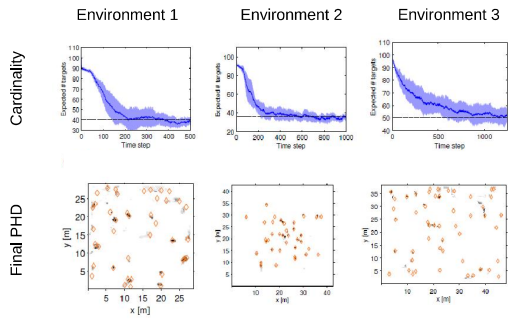


(c) Environment 3

31 May 2014

18

Results



31 May 2014

19


Conclusion

- Multi-target estimation
 - Avoids data association
- Communication strategy
 - Allows for a consistent estimate across the team
 - Prevents double counting information
 - Allows for decentralized estimation and control
- Information-based, receding horizon control law
 - Computationally tractable
 - Trades off information benefit of sensing versus communication
- Simulation results show performance with 1's to 100's of targets in indoor environments

31 May 2014

22

Results – High Target Density



31 May 2014

20


Questions?

31 May 2014

23

Results – Low Target Density

- Limited communication range
- 4 robots, 5 access points



31 May 2014

21

2.7 Toward Dynamical Sensor Management: Reactive Wall-following on RHex

Speaker Avik De (University of Pennsylvania)

Abstract We propose a new paradigm for reactive wall-following by a planar robot taking the form of an actively steered sensor model that augments the robot’s motion dynamics. We postulate a foveated sensor capable of delivering third-order infinitesimal (range, tangent, and curvature) data at a point along a wall (modeled as an unknown smooth plane curve) specified by the angle of the ray from the robot’s body that first intersects it. We develop feedback policies for the coupled (point or unicycle) sensorimotor system that drive the sensor’s foveal angle as a function of the instantaneous infinitesimal data, in accord with the trade-off between a desired standoff and progress-rate as the wall’s curvature varies unpredictably in the manner of an unmodeled noise signal. We prove that in any neighborhood within which the third-order infinitesimal data accurately predicts the local “shape” of the wall, neither robot will ever hit it. We empirically demonstrate with comparative physical studies that the new active sensor management strategy yields superior average tracking performance and avoids catastrophic collisions or wall losses relative to the passive sensor variant.

This work was presented in poster form at ICRA 2013.

Suggested reading [[Cow06](#), [JK05](#), [DK13b](#), [DBK14](#), [KR90](#)]

Toward Dynamical Sensor Management: Reactive Wall-following and Beyond

Avik De, Daniel Koditschek
 Joint Inference and Control Workshop at ICRA 2014
 This work was supported by AFOSR MURI FA9550-10-1-0567.

1

End Result: Wall-following with RHex




4

Inspiration for Wall-following

Why is wall-following useful?


- Useful behavior with sparse sensing
- Online navigation problem with obstacle

Some Exemplar Prior Work



[Cowan 06] Antenna-enabled rapid wall following in American cockroach

[Lamperski 05] Passive antenna robotic wall-following



[Matveev 11] Boundary patrol using "disk" sensors



Could an actively-controlled foveating sensor help the control

2

Plant Model




Locomotion Model

- RHex: kinematic unicycle
- Non-holonomically constrained

Sensor Model

- Hardware: fixed-pitch Hokuyo laser scanner
- Idealized sensor has an *infinitesimal sensorium*: range, tangent, curvature (e.g. antenna)
- Implemented with numerical pre-processing
- Sensor is *actively steerable*



5




Related: Biological & Bio-inspired Active Tactile Sensing

In Biology,

- Rats "whisk" [Hartmann 01]
- Cockroaches?


In bio-inspired mechanics,

- Active feedforward whisking => better SNR on depth [Kim & Moller 06]
- Active antennae sense depth and texture [Kaneko 98]



3


Environment / Task Model



- Wall = unknown smooth simple plane curve of bounded curvature
 - (violation of latter doesn't affect performance in practice)
- Task is to
 - maintain a desired offset, q_2 , from the wall (with q_1 small),
 - and a desired rate-of-progress, ρ .

6

Illustrative Result for Particle Robot (1)



Proposition 1 (Point robot convergence). *With active sensing, we can assure (a) $\rho = 1$ (desired rate of progress), (b) $\tilde{q}_2 \rightarrow 1$, and (c) $q_1 \rightarrow 0$, whereas with passive sensing we can only guarantee (a) and (b).*

- Start with a point robot (no unicycle constraint)
- Simple kinematics:
 - The **active sensor** system is fully actuated!
 - In contrast, the **passive sensor** system for $\rho \neq 1$ is

$$\dot{q} = \begin{bmatrix} 0 \\ \tilde{u}_\perp \end{bmatrix} - \bar{\kappa} J q,$$
 Easy to see that q is *uncontrollable!*
- Rate-of-progress related by trigonometry:

$$\rho = \tilde{u}_\parallel + \frac{\|q\|_2}{q_2} v.$$
- Local frame kinematics even simpler:

$$\dot{q} = \tilde{u} + \rho n,$$
- Only difficulties are unmodeled "noise" vector,

$$n := -e_1 - \bar{\kappa} J q,$$
- ...and ρ .

7

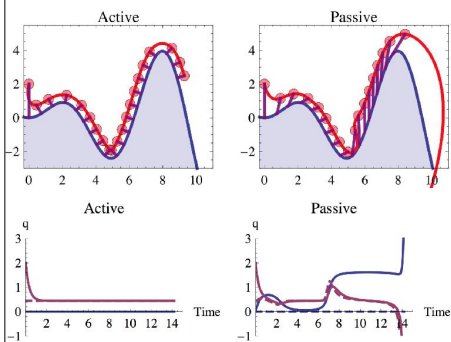
Extension to Kinematic Unicycle



- Cannot "cancel" the noise vector as before, but can control its *magnitude* by ρ (rate-of-progress)
- Needed for proof, but experiments still succeed for rapid rates

10

Illustrative Result for Particle Robot (2)




- Intuitively,
- Active sensing: three dof (q_x, q_y, ρ) and three inputs (u_x, u_y, v) vs.
 - Passive sensing: three dof and two inputs (u_x, u_y only).

8

Experiment: Unmodeled Hallway



- Tracking error is highly correlated with curvature spikes (**corners**)
- Traditionally, sharp corners and clutter have been challenging for smooth wall-following controllers




11

Result: Local Wall-avoidance Guarantee

Proposition: The feedback controller guarantees safety in a local neighborhood of radius ϵ , around the sensed point $b(\sigma)$.

Proof: Wall is repelling in a neighborhood where a quadratic approx is valid



9

Experiment: Active vs. Passive Sensing (1)

Passive sensor with small look-ahead (right-looking) fails at concave corner

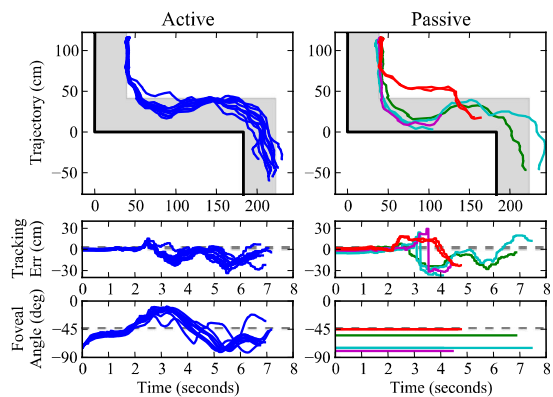
Passive sensor with large look-ahead (forward-looking) fails at convex corner (recall, system is memoryless)

Active sensor automatically adjusts look-ahead according to rate-of-progress, curvature, state.



12

Experiment: Active vs. Passive Sensing (2)




13

Looking Ahead: The Problem of Local Navigation from Infinitesimal Data

Navigation Functions (NF's)


- Dynamical navigation (goal-seeking while avoiding obstacles) is a solved problem (right [Rimon 89])—when the world is known *a priori*
- Real-world problems:
 - underactuation,
 - exteroception is local



Local Navigation

- Possit ability to measure local **gradients**

$$G := [\nabla\gamma, \nabla\beta]$$
- (as well as its rate-of-change for either dynamic systems or dynamic tasks)
- Key ideas:
 - Work in local frame
 - *Compose* the requirements as a weighting $G^T W$
- Problem when G loses column rank
 - WF: columns always orthogonal!
- Eg: Hill climbing (simulation)



14

Active (Foveated / Pointed) Sensors

- How do we estimate $\nabla\beta$? Assuming the world is static, obstacle surface must be in the direction of motion.
- Dithering / curvature estimation

Conclusion



We

- presented a novel outlook to real-time control of a **coupled sensorimotor system**,
- provided analytical proofs of stability, convergence and guarantees against failure as long as the robot stays near the sensed point, and
- demonstrated qualitatively favorable real-world performance

15

2.8 The Impact of Perception Capabilities on Agile Robot Motion: A Statistical Mechanics Perspective

Speaker Sertac Karaman (MIT)

The Impact of Perception Capabilities on Agile Robot Motion: A Statistical Mechanics Perspective

Sertac Karaman

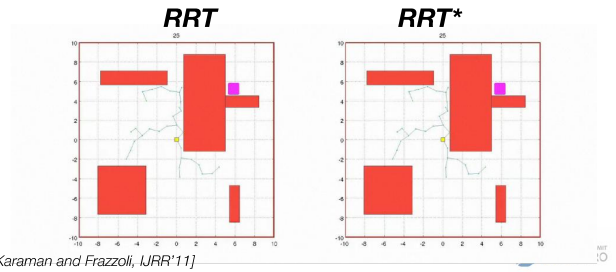
Department of Aeronautics and Astronautics
Massachusetts Institute of Technology

karaman@mit.edu
<http://karaman.mit.edu>



Sampling-based Algorithms for Optimal Planning

- **RRT** fails to converge to optimal solutions
- **RRT*** guarantees **asymptotic optimality**: almost-sure convergence to optimal solutions
- **RRT** and **RRT*** have the same asymptotic computational complexity



Practical Algorithms for Motion Planning

Robot motion planning:


Find a path from point A to point B

- Fundamental problem in robotics.
- Computationally challenging; Practical algorithms exist


Rapidly-exploring Random Tree (RRT)

[LaValle, Kuffner, '01]

- Randomly sample states
- Connect samples into a tree



Sampling-based Algorithms for Motion Planning



[Jeon, Karaman, and Frazzoli, '12]

RRT in the Real World: The DARPA Urban Challenge

- Drive 60 miles in urban traffic in less than 6 hours
- Abide by the rules of the road, e.g., lanes, intersections, passing, u-turns, etc.



[Leonard et al., '08]




Anytime Sampling-based Algorithms

Anytime flavor:

1. Quickly find an approximate solution
2. Improve the approximation provably towards an exact solution

- Differential games
- Provably-correct trajectory synthesis from high-level LTL specifications
- Continuous-time stochastic optimal control (including sensing uncertainty).

Anytime computation of the value function



[Huynh, Karaman, Frazzoli, '12]

[Chaudhari, Karaman, Frazzoli, '13]



Opportunities and Challenges in Sampling-based Algorithms for Joint Sensing and Control

- High-dimensional state spaces.
- Amenable to anytime computation.
- Formal guarantees, e.g., computational complexity, probabilistic completeness, and asymptotic optimality.



- Algorithms can benefit from better inference algorithms, better optimization methods, etc.
- Meaningful special cases: Minimal predictive models of potential environments.
- Parallel computation, e.g., on GP-GPUs.



I. High-speed Navigation in Cluttered Environments


On Perception Capabilities of Agile Robotic Vehicles

BBC Documentary: Goshawk Flight in Woodland



On Bio-inspired Agile Robotics

Agile



Inspired by the Nature, the community aspires to build small, agile, efficient robots

Small

- Tremendous progress in designing/building hardware.
- Can we characterize the fundamental limits of agile motion, e.g., in terms of the perception capabilities of the robot?



Engineering High-speed Robotic Vehicles?




- What is the maximum speed that this robot can achieve maintain for a long time?
- **How does this performance depend on perception, actuation, and computation capabilities of the robot?**



Phase Transitions: The General Case

Forest process: A **marked point process** generates *locations* and *sizes* of trees.



An instantiation of the Poisson model with uniformly random tree radius





FIG. 2. 134 spruce of age 60 years in a 56 m x 38 m stand of Tharandter Wald (Germany). The trees are shown as circles, where the diameters are proportional to the height. Figure taken from [Stoyan and Penttinen '00]



The Single-integrator Bird in a Poisson Forest

Bird dynamics: $\dot{x}(t) = \begin{pmatrix} \nu \\ u(t) \end{pmatrix}; |u(t)| \leq 1$
 $y(t) = x(t)$



- Forest process:**
- **Locations** of the trees are generated by a **Poisson process** with intensity ρ (tree density).
 - **Radii** of the trees are the same, say r (tree radius).

Super-critical Flight

Theorem
 $\mathbb{P}(\text{"An infinite collision-free trajectory exists"}) = 0$
 when $\frac{\rho r^2}{\sin(\alpha)} > 0.219450$

Sub-critical Flight

Theorem
 $\mathbb{P}(\text{"An infinite collision-free trajectory exists"}) = 1$
 when $\frac{\rho r^2}{\sin(\alpha)} < 0.071921$




Phase Transitions: The General Case

Bird Dynamics:

$$\dot{x}(t) = f_\nu(x(t), u(t)); \quad x(t) \in \mathbb{R}^n, y(t) \in \mathbb{R}^2, \\ y(t) = g_\nu(x(t)) \quad \nu \in \mathbb{R}_{>0}$$

Speed Model: Non-decreasing path sets with decreasing ν .
Assumption: Dynamics of the bird is translation invariant.




- We look for infinite collision-free trajectories



Phase Transitions in a Poisson Forest

Empirical probability of tracing a long forest segment constructed using computational simulations.



Phase Transitions: The General Case

Forest process: A **spatial marked point process** generates *locations* and *sizes* of trees.

Bird dynamics: $\dot{x}(t) = f_\nu(x(t), u(t)); \quad x(t) \in \mathbb{R}^n, y(t) \in \mathbb{R}^2, \\ y(t) = g_\nu(x(t)) \quad \nu \in \mathbb{R}_{>0}$

Speed Model: Non-decreasing path sets with decreasing ν .
Assumption: Dynamics of the bird is translation invariant.

Theorem: Phase transitions in ergodic forests

Suppose the forest generating process is **ergodic**.

Then, there exists a critical speed, ν_{crit} , such that

- For all $\nu > \nu_{crit}$,
 $\mathbb{P}(\text{"An infinite collision-free trajectory exists"}) = 0$.
- For all $\nu < \nu_{crit}$,
 $\mathbb{P}(\text{"An infinite collision-free trajectory exists"}) = 1$.




Proof Techniques from Percolation Theory

- **Percolation theory** studies global properties of random graphs (in particular, random vertices in discrete/continuous domains)

Global properties:

- Connectivity
- Average degree
- Giant components




- Main results in terms of phase transitions and phase diagrams.

- Immediate applications in statistical mechanics

- More recent applications in communication/social/financial/political networks.

- We heavily utilize percolation theoretic arguments.

- Often, the resulting percolation models are unique.



How are Perception Capabilities and Performance Related?


- The performance increases rapidly with increasing perception range:
 - Assuming the forest distribution is "light-tailed," an optimal planning algorithm achieves the following performance:

Theorem

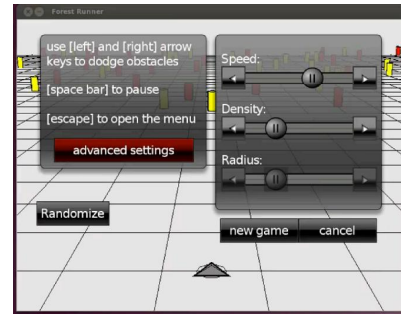
Let R denote the perception range and $L(R)$ denote the maximum distance that the robot can travel without collisions. Then, there are constants c and γ such that

$$\mathbb{P}(L(R) \geq cR^\gamma) \geq p$$

for all $R > 0$.




Performance and Agility: Experiments




How about Agility?

- Performance drops sharply with decreasing agility





Performance and Agility: Experiments



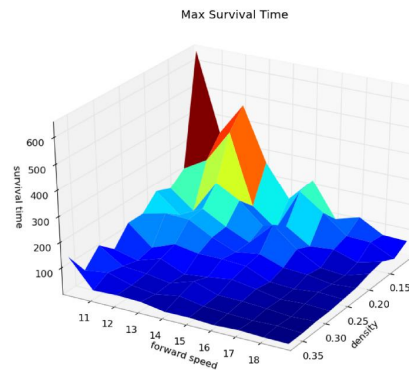
- Significant interest; Presentations encourage people to play.

Experiments with Pigeons at Harvard University

by Andrew Biewener at the Harvard University Department of Evolutionary Biology.



Performance and Agility: Experiments



II. On High-performance Information Collection

A Different Interpretation: On the Power of (even a little) Local Perception/Planning

- Logarithmic sensing range (planning horizon) is enough, under mild assumptions:

Corollary

With range $O(\log(L))$, the robot can travel a distance of L almost as good as if it had infinite perception range.

- Example: Consider i.i.d. rewards distributed $\text{Exp}(1)$.


	No Range	$O(\log(L))$ range	L range
Time-average Reward	1	-2	2

- Rather than rewards, we can consider data collection and inference (e.g., adaptive sampling for hypothesis testing, sparse recovery, ...)

[Ma and Karaman, WAFR'14 (to appear)]



Online Planning and Sensing for Information Gathering



Model

- Target locations are not known a priori; they are discovered on the fly.
- Spatial statistics for the target locations are available, e.g., from past experience.

Fundamental Questions

- How much information can the smartest algorithm possibly collect?
- How does performance scale with perception, actuation, computation capabilities?




Novel Connections between Planning/Sensing and Nonequilibrium Statistical Mechanics (NESM)

The Methodology:

- Establish novel connection between this problem and fundamental problems of nonequilibrium statistical mechanics.
- Robot treated like a particle moving in a field.


NESM in biological transport



The Results regarding perception capabilities:

- The problem is in the Kardar-Parisi-Zhang universality class (not Gaussian).
- The optimal reward converges to the Tracy-Widom distribution (not Gaussian).
- Fluctuations from the expected value is of order $1/3$ (not $1/2$ as in Gaussian).

Tracy-Widom Distribution



- Results also for actuation and computation capabilities.



[Ma and Karaman, WAFR'14 (to appear)]



On Performance and Perception Range

- In other words, we can go very far optimally with little perception range.

Thm: With range r the vehicle can travel $c_1 \exp(c_2 \delta^{2/3} r)$ distance and collect $c^* - \delta$ reward per unit time with high probability.



[Ma and Karaman, WAFR'14 (to appear)]



Conclusions and Remarks

- The design of agile robotic vehicles for
 - high-speed navigation through cluttered environments;
 - rapid information collection.
- Tradeoff between performance and robot's capabilities, e.g., in terms of perception (as well as actuation and computation).
- Novel connections with statistical mechanics.



Conclusions and Remarks

- Deeper connections with statistical mechanics.
- Planning in the sensor/information space, in particular realistic sensor models should include *sensing uncertainty*, *occlusions*.
- Performance with respect to actuation and computation capabilities as well as perception.



References

- [Cow06] N. J. Cowan. Task-level control of rapid wall following in the american cockroach. *Journal of Experimental Biology*, 209(9):1617–1629, May 2006.
- A description of the wall-following behavior in cockroaches.
- [DBK14] Avik De, Karl S. Bayer, and Daniel E. Koditschek. Active sensing for dynamic, non-holonomic, robust visual servoing. In *IEEE International Conference on Robotics and Automation*, May 2014.
- A description of the benefits of an actively panning camera to the stabilizability of a unicycle robot (written by the speaker).
- [DK13a] P. Dames and V. Kumar. Cooperative multi-target localization with noisy sensors. In *Proc. IEEE Int. Conf. on Robotics and Automation*, Karlsruhe, Germany, May 2013.
- An article by the speaker.
- [DK13b] Avik De and Daniel E. Koditschek. Toward dynamical sensor management for reactive wall-following. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, page 24002406. IEEE, 2013.
- A description of the controllability benefits of an active antenna in robotic wall-following (written by the speaker).
- [EWBS09] Ozgur Erdinc, Peter Willett, and Yaakov Bar-Shalom. The bin-occupancy filter and its connection to the phd filters. *Signal Processing, IEEE Transactions on*, 57(11):4232–4246, 2009.
- An intuitive derivation of the PHD filter equations.
- [Gro06] Ben Grocholsky. Information-theoretic control of multiple sensor platforms. 2006.
- This PhD thesis is a nice resource on information-theoretic control of multi-robot teams.
- [HT10] Gabriel M Hoffmann and Claire J Tomlin. Mobile sensor network control using mutual information methods and particle filters. *Automatic Control, IEEE Transactions on*, 55(1):32–47, 2010.
- An overview of information based control for teams of robots when using a particle set to approximate the distribution over the target state.
- [JK05] E. W. Justh and P. S. Krishnaprasad. Natural frames and interacting particles in three dimensions. *arXiv:math/0503390*, March 2005.
- An instance of the “moving frame” technique for a local analysis.
- [KR90] Daniel E Koditschek and Elon Rimon. Robot navigation functions on manifolds with boundary. *Advances in Applied Mathematics*, 11(4):412–442, December 1990.
- Construction of provably correct navigation functions when global information is available.
- [Pul05] GW Pulford. Taxonomy of multiple target tracking methods. In *Radar, Sonar and Navigation, IEE Proceedings-*, volume 152, pages 291–304. IET, 2005.
- This provides a fairly comprehensive, though slightly outdated, overview of multi-target tracking methods.
- [VSD03] Ba-Ngu Vo, Sumeetpal Singh, and Arnaud Doucet. Sequential monte carlo implementation of the phd filter for multi-target tracking. In *Proc. Intl Conf. on Information Fusion*, pages 792–799, 2003.
- An overview of how to implement PHD filtering.