Neural Robot Learning

MetaNook 2018

April, 4th 2018 by Prof. Dr. Elmar Rueckert

latest updated Nov. 2nd 2018

Studying Robotics and Autonomous Systems (RAS)

Humanoid robots are among the most complex machines on earth.

And you will learn here how to build, teach and program them.

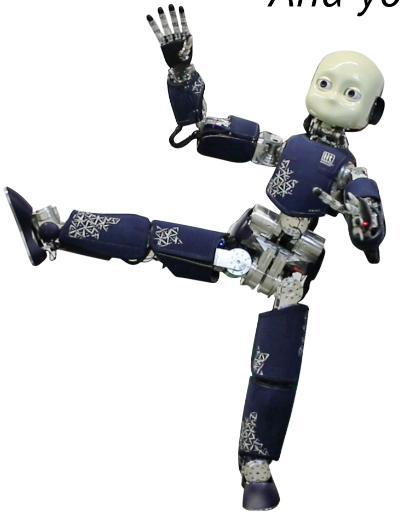

What is a hu

"Humanoid robots a umans", page 58
in the tec.org panel

The pepper robot by Softbank Robotics.

Th
m

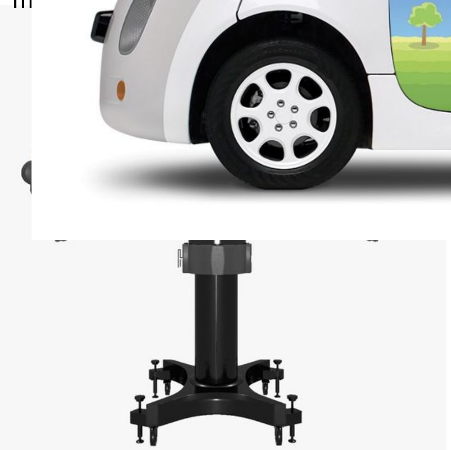

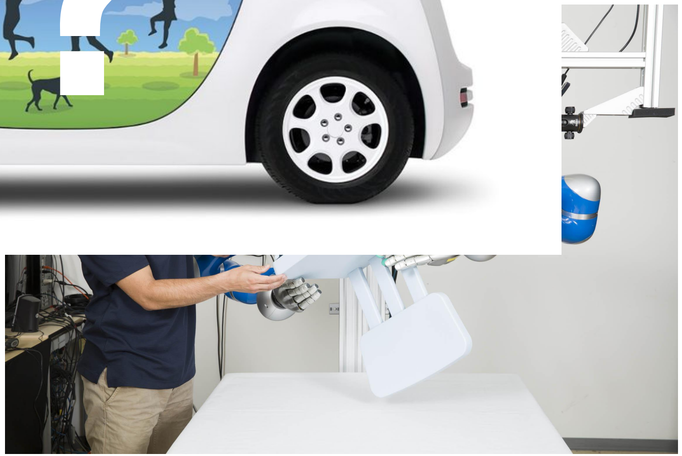

Prof. Jan Peters from the Intelligent Autonomous Systems Lab at Technische Universität Darmstadt in 2017.

A great job

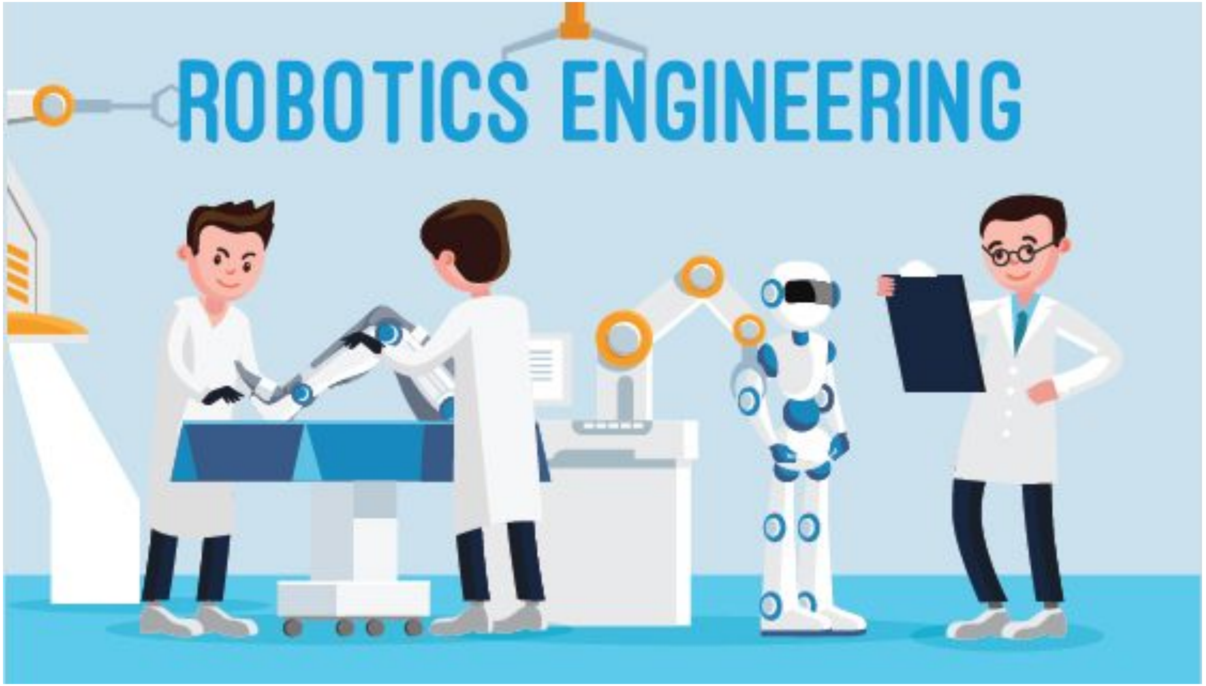

With amazing job prospects

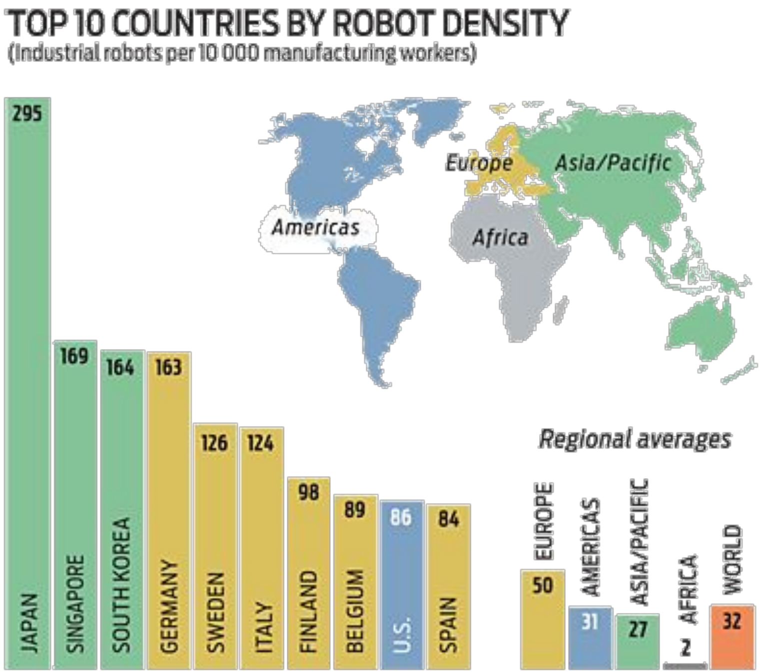

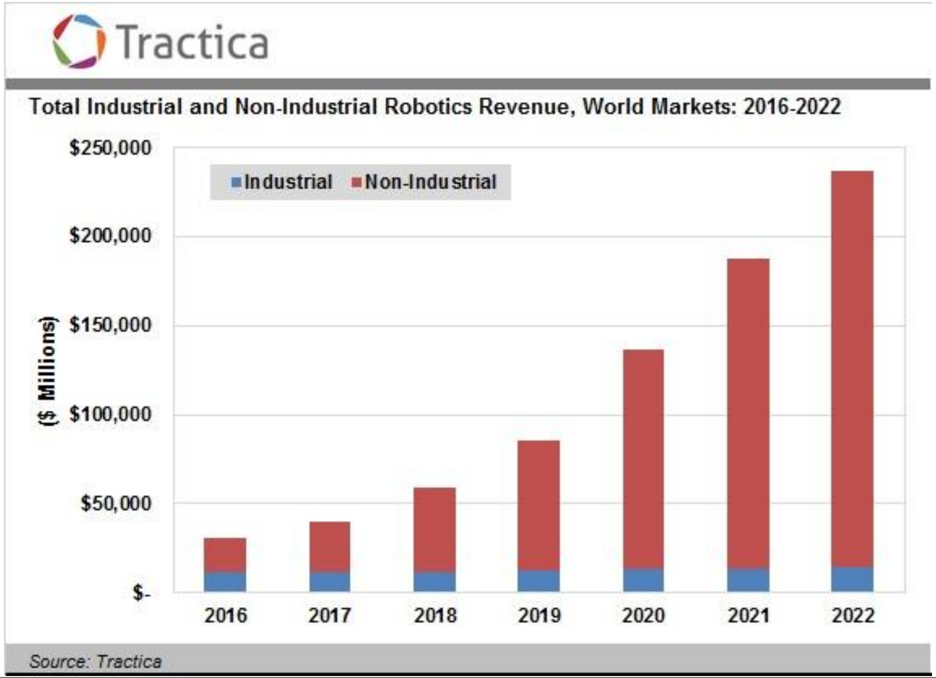

More specifically, humanoids ...

- can interact with humans in environments made for humans.
- can autonomously acquire new skill or self-improve existing skills.
- can reason about the world and create abstractions.

⇨ Humanoid robotics is a subfield of Artificial Intelligence!

More than robotics …

The challenges in understanding humans and in building intelligent humanoids are converging!

~ 700 muscles
~ 100 joints
~ 100 x 10^6 photo receptors
~ 10^2 FA-I receptors per fingertip

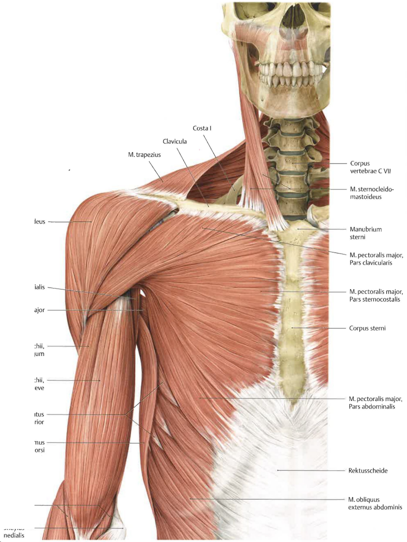

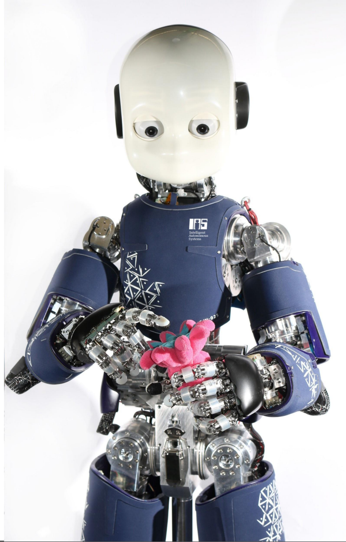

- 53 degrees of freedom
- 4 force/torque sensors
- 1.8 x 10^6 photo receptors
~ 2000 tactile sensors

More than robotics …

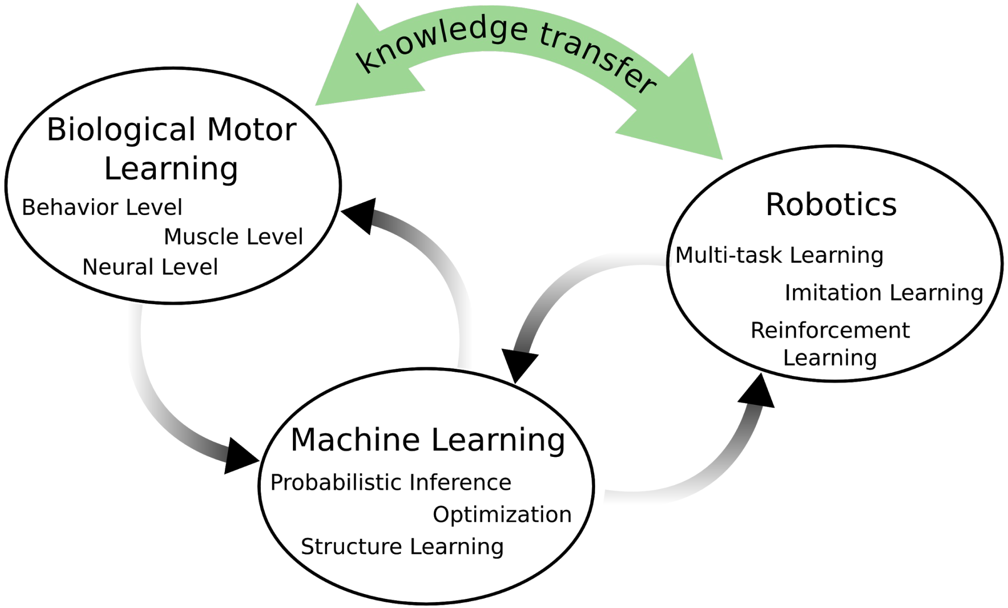

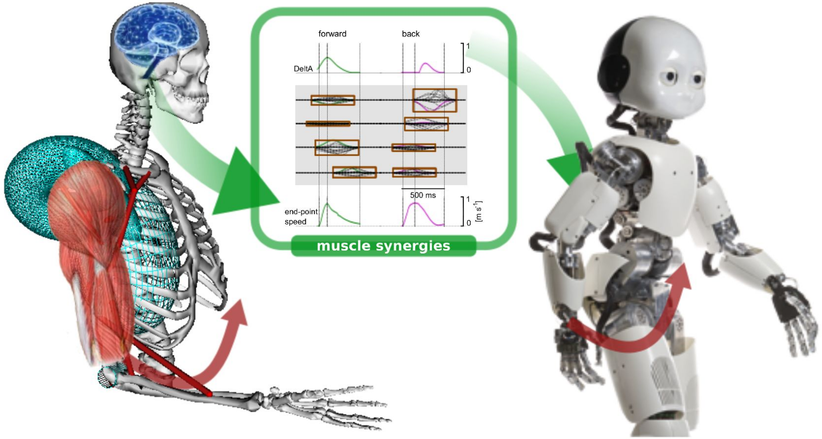

A brief historical review

Link to a more detailed history review

1920 **Karek Capek**: “robot” in his play “R.U.R.” (Rossum’s Universal Robots).

1941 **Isaac Asimov**: Three laws of “robotics”:

1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.
2. A robot must obey orders given it by human beings except where such orders would conflict with the First Law.
3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.

A brief historical review

1968 "**Shakey**" of the "Stanford Research Institute" defines a landmark in robotics:

- basic planning and navigation skills.
- object detection and manipulation capabilities.

A brief historical review

1973 **Ichiro Kato** develops the first “full-scale” antrophomorphic humanoid, WABOT I.

A brief historical review

1996 **Honda** presents its P2

they started with E0 in 1986

the history of Hunda’s humanoids

A brief historical review

2004 The Italian Institute of Technologie presents the **ICub** (intelligent man-cub).

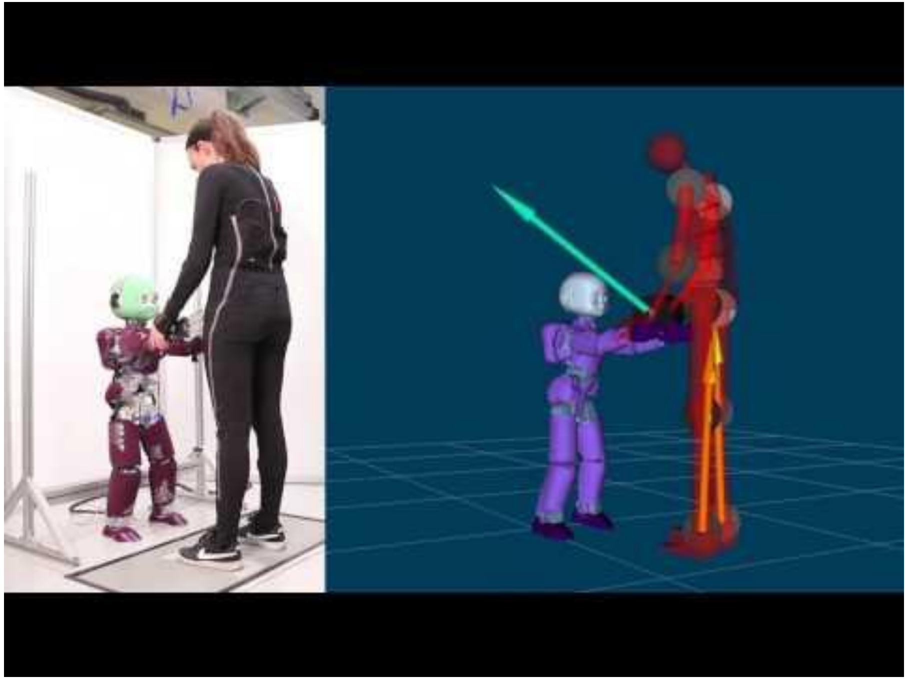

A brief historical review

2017 Boston dynamics’ **Atlas** impresses the robotics community.

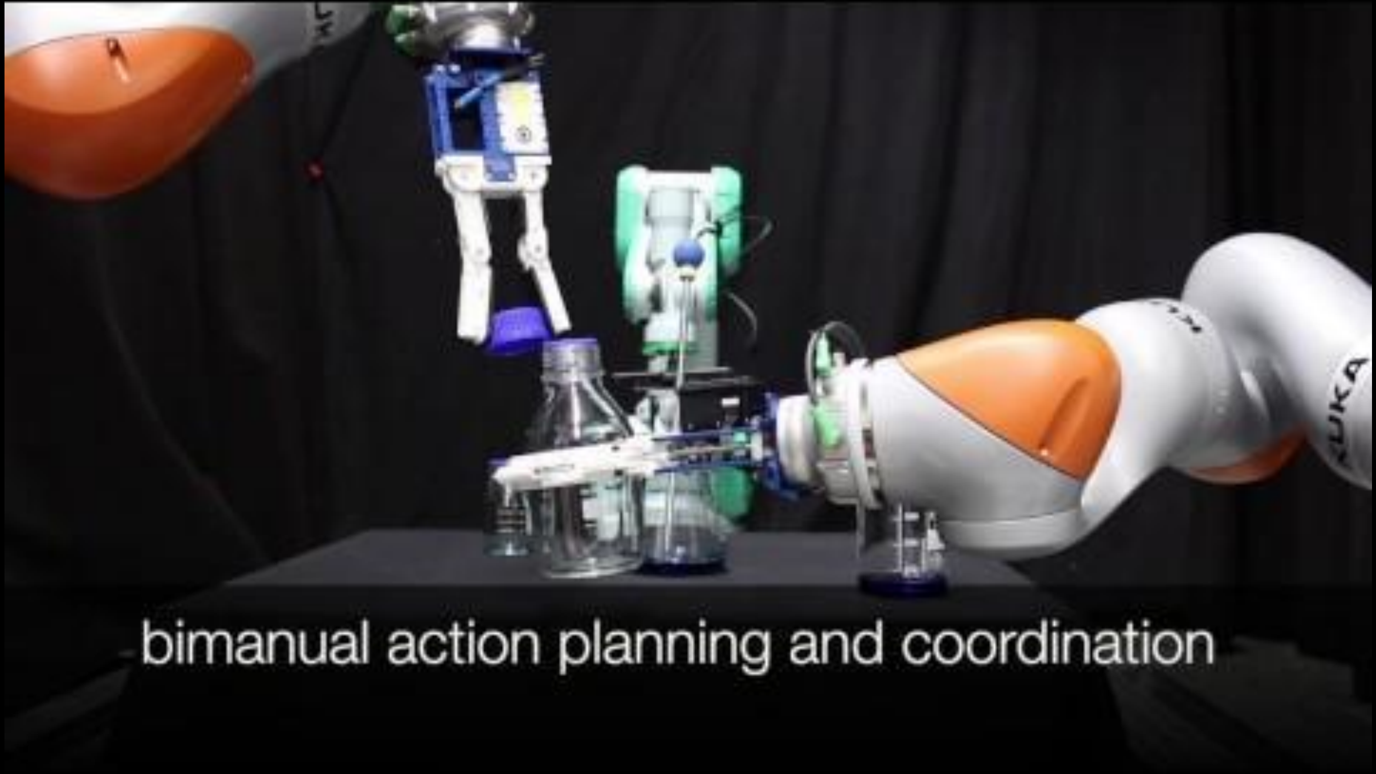
KUKA
bimanual action planning and coordination

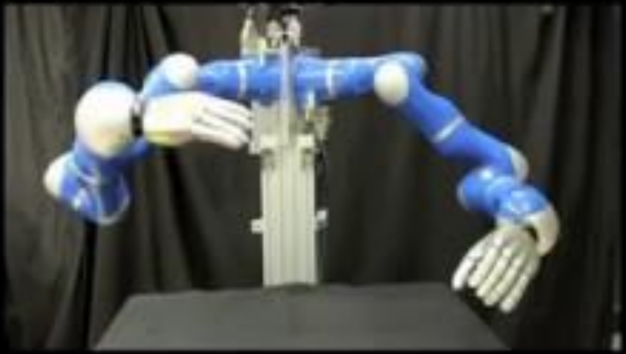



9:33:02 05/06/2015 UTC

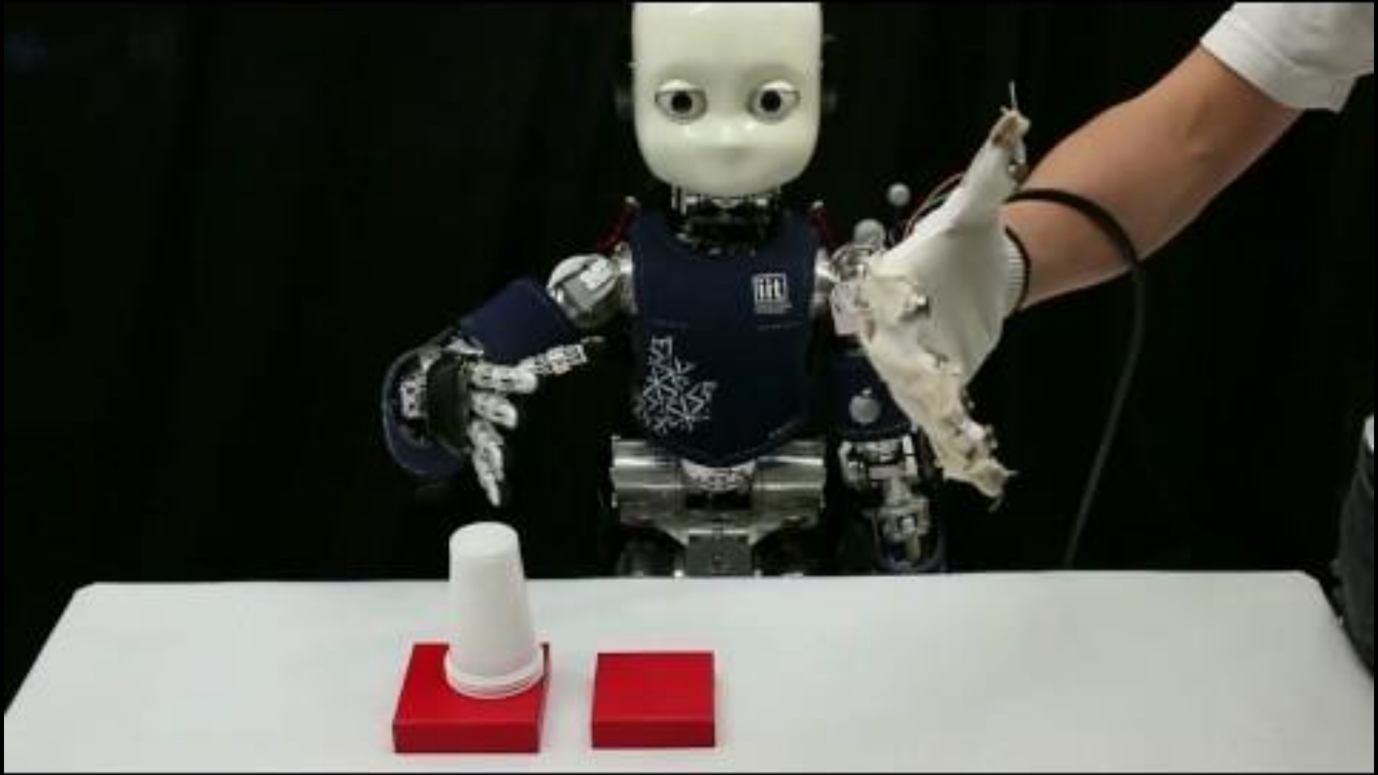

Research questions

1. How can humans learn new motor skills within a few trials?
   a. “control only when necessary” - motor variability
   b. exploiting kinematic and task redundancy
   c. transfer of related skills

2. How do humans solve cognitive reasoning tasks in huge spaces?
   a. planning in stochastic environments
   b. inferring multiple solutions in milliseconds
   c. online model adaptation from intrinsic motivation signals.

II. Representations of Skills & Imitation Learning

V. Cognitive Reasoning & Planning

I. Kinematics, Dynamics & Model Learning

III. Feedback, Priorities & Torque Control

IV. Reinforcement Learning & Policy Search

VI. Sensor Integration & Fusion

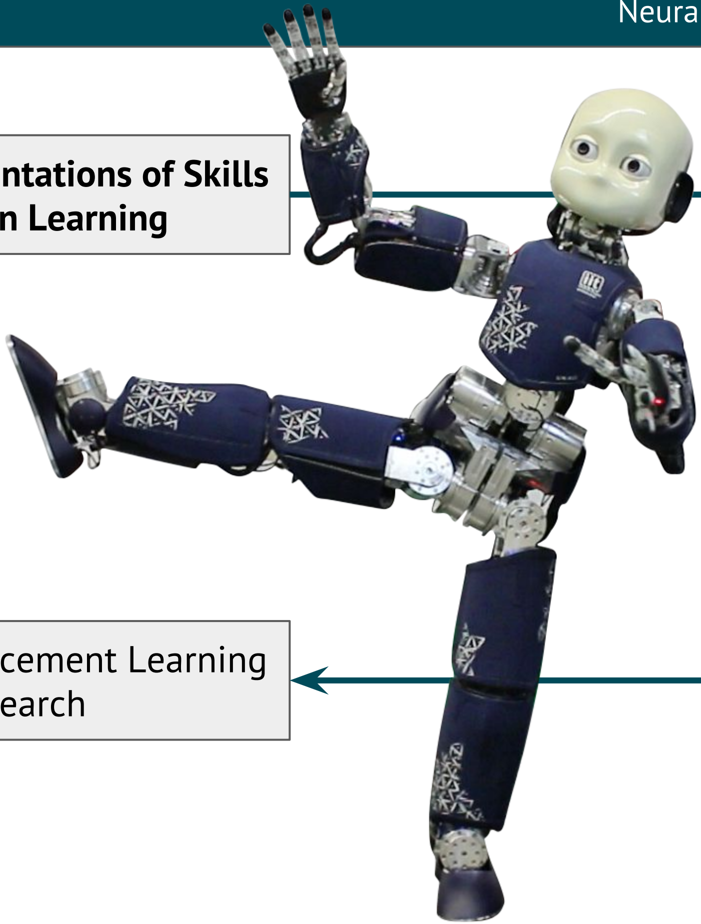

Where do we need representations of skills?

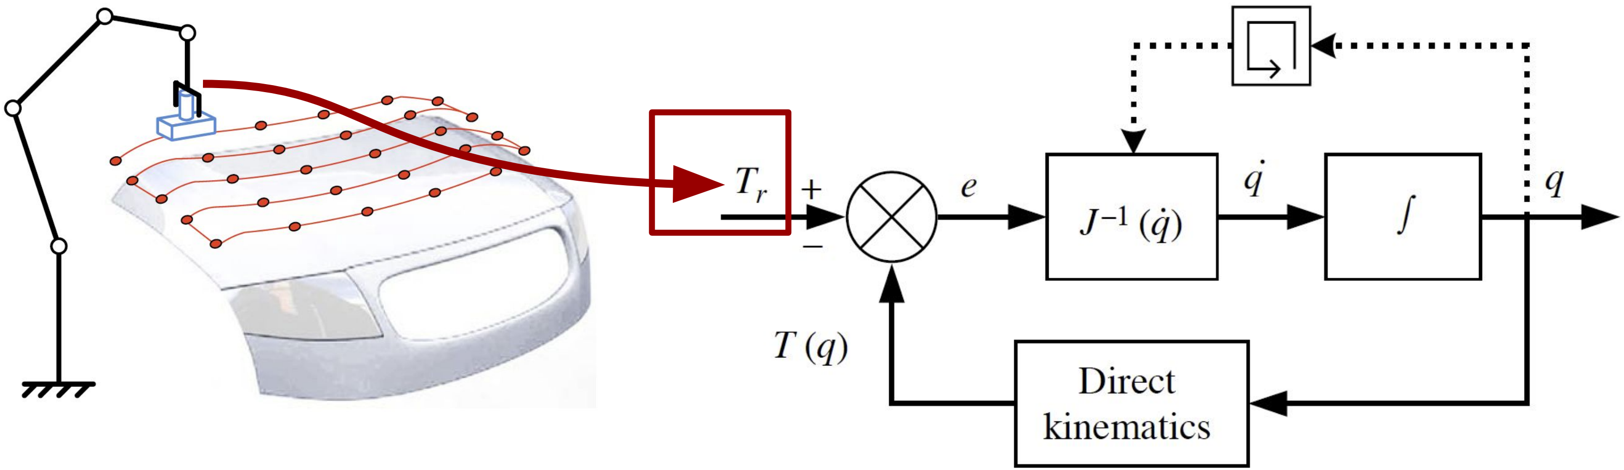

Fan Zeng, Beshah Ayalew and Mohammed Omar: Roboticc automotive paint curing using thermal signature feedback, 2009

Naive vector/matrix representation

scales in **O(d x T x K),** where d … number of joints, force plates or markers,
T … number of time steps per trial k = 1…K

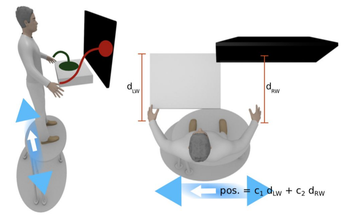

50 Mio. data pts stored with 64 bits per double **> 3 GByte for movements of 1 second!**

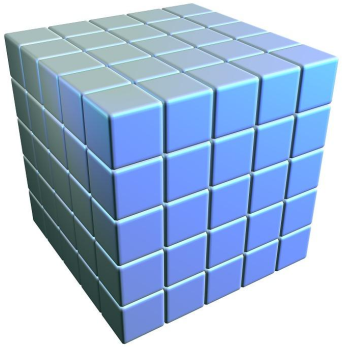

Even when we average over all 9600 trials we would need to store **5500 data points per second!**

Desired features of skill representations

- **Compact** (few parameters to learn).
- **Smooth** (need to compute derivatives for velocities and controls).
- **Flexible** generalizables to different tasks (goal locations, orientations, etc.).
- Can be learnt from the data through **imitation learning (IM)**.
- Self-improvement through **reinforcement learning (RL)**.
- **Composable** through sequencing and **co-activation.**
- **Stochastic**, can model the variance of the data.
- **Coupled**, can model the coupling of joints.

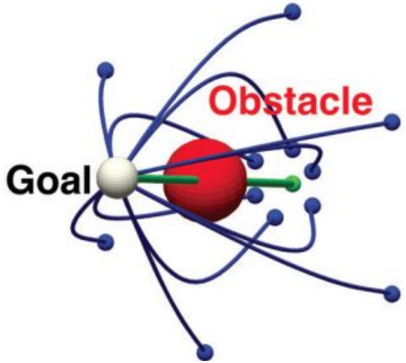

Ex. flexibility to start at different poses.

My approach: learning probabilistic models

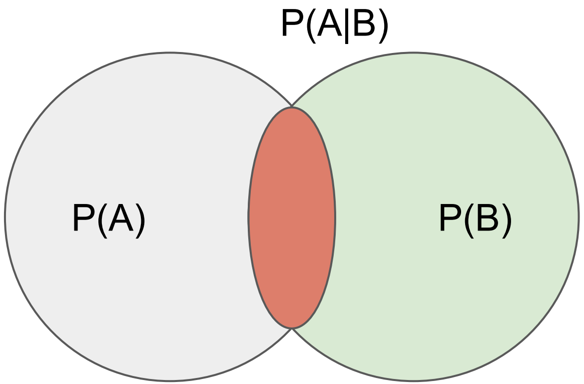

<u>Learning problem:</u>

$P(A|B) = P(A, B)/P(B)$

given data samples from $P(A, B)$
assuming priors $P(A)$, $P(B)$

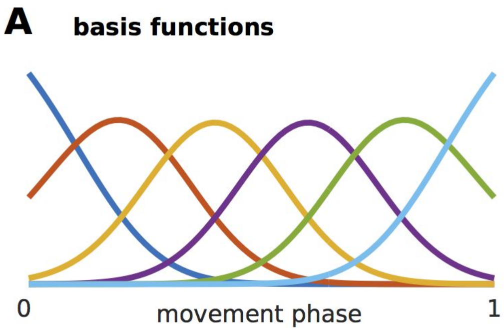
A
basis functions
0
movement phase
1

[1] Generative Model: $\boldsymbol{y}_t = \boldsymbol{\Phi}_t \, \boldsymbol{w}$

[2] Gaussian Features: $\phi_{t,i} = \frac{1}{\mathscr{Z}} \exp\left(-\frac{1}{2h}(z(t) - c_i)^2\right)$,

[3] Learning the Prior: $\boldsymbol{w}^{[i]} = \left(\boldsymbol{\Phi}_{1:T}^T \boldsymbol{\Phi}_{1:T} + \lambda \boldsymbol{I}\right)^{-1} \boldsymbol{\Phi}_{1:T}^T \boldsymbol{\tau}^{[i]}$.

[4] Model:

$$
\begin{aligned}
p(\boldsymbol{\tau}) &= \int p(\boldsymbol{\tau} | \boldsymbol{w}) p(\boldsymbol{w}) d\boldsymbol{w} \\
&= \int \mathscr{N}(\boldsymbol{y}_{1:T} | \boldsymbol{\Phi}_{1:T} \boldsymbol{w}, \boldsymbol{\Sigma}_y) \underbrace{\boxed{\mathscr{N}(\boldsymbol{w} | \boldsymbol{\mu}_w, \boldsymbol{\Sigma}_w)}}_{\text{learned prior}} d\boldsymbol{w} \\
&= \mathscr{N}\left(\boldsymbol{y}_{1:T} | \boldsymbol{\Phi}_{1:T} \boldsymbol{w},\ \boldsymbol{\Phi}_{1:T} \boldsymbol{\Sigma}_w \boldsymbol{\Phi}_{1:T}^T + \boldsymbol{\Sigma}_y\right) .
\end{aligned}
$$

Learning from demonstrations

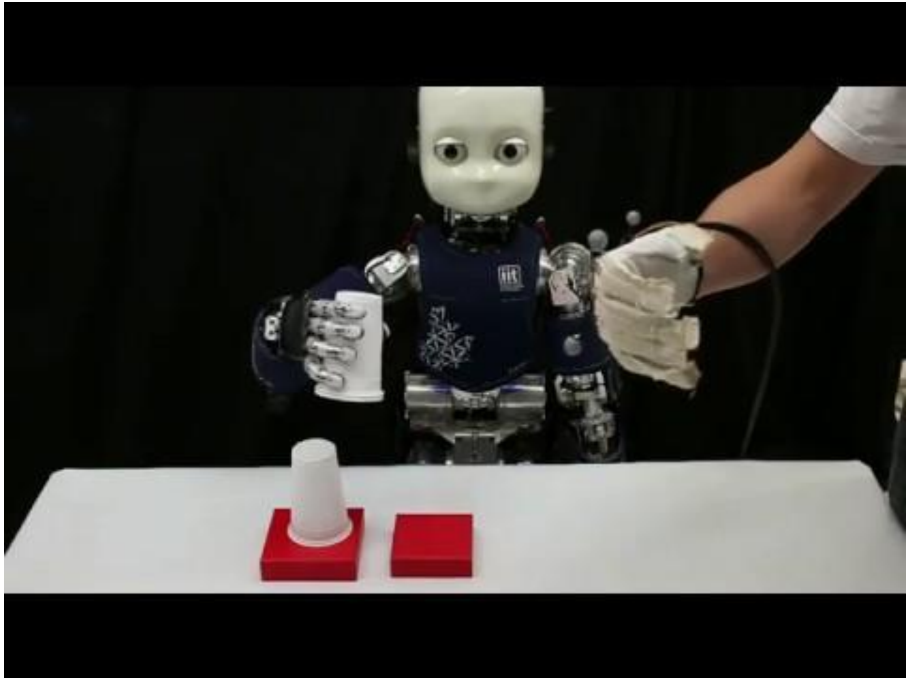

Classical Motion Analysis

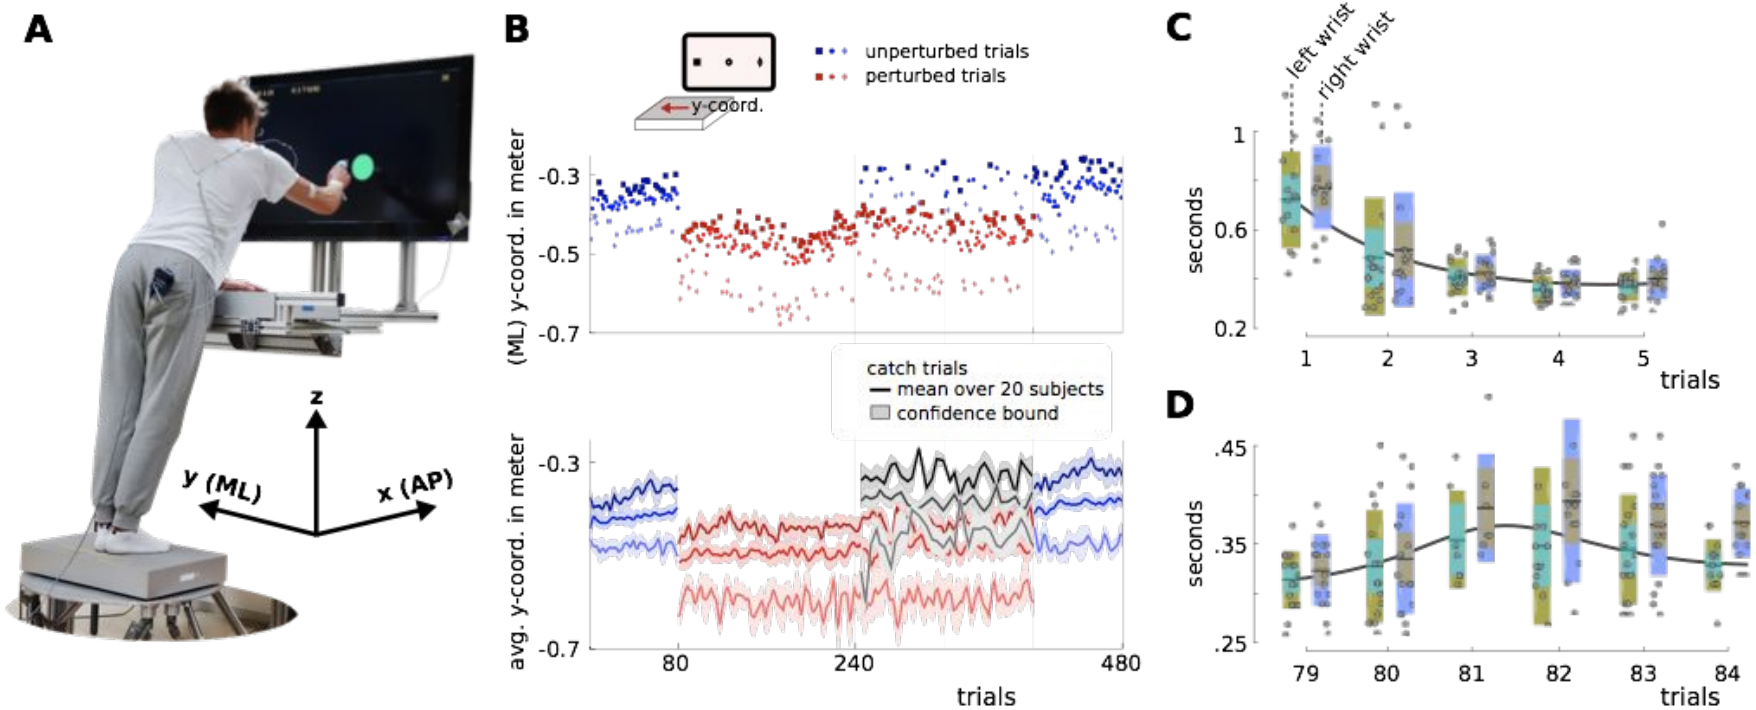

Probabilistic Motion Analysis

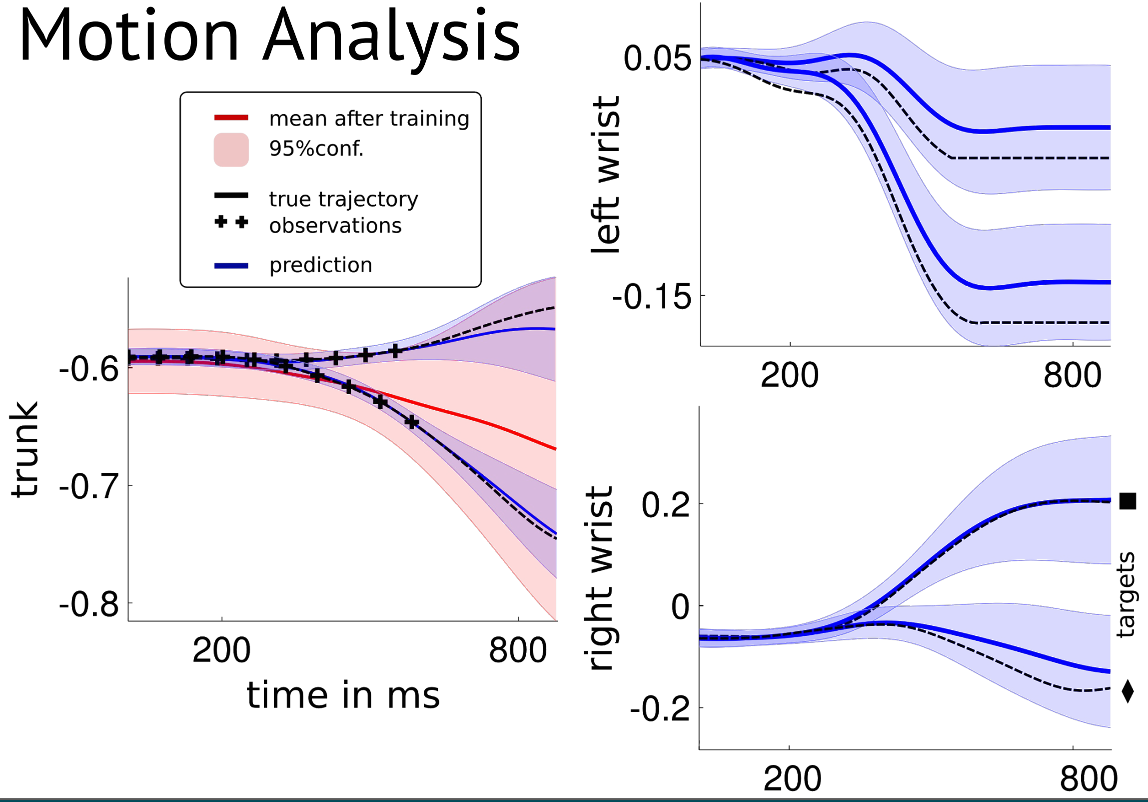

P(A|B)

Predictive models of EMGs

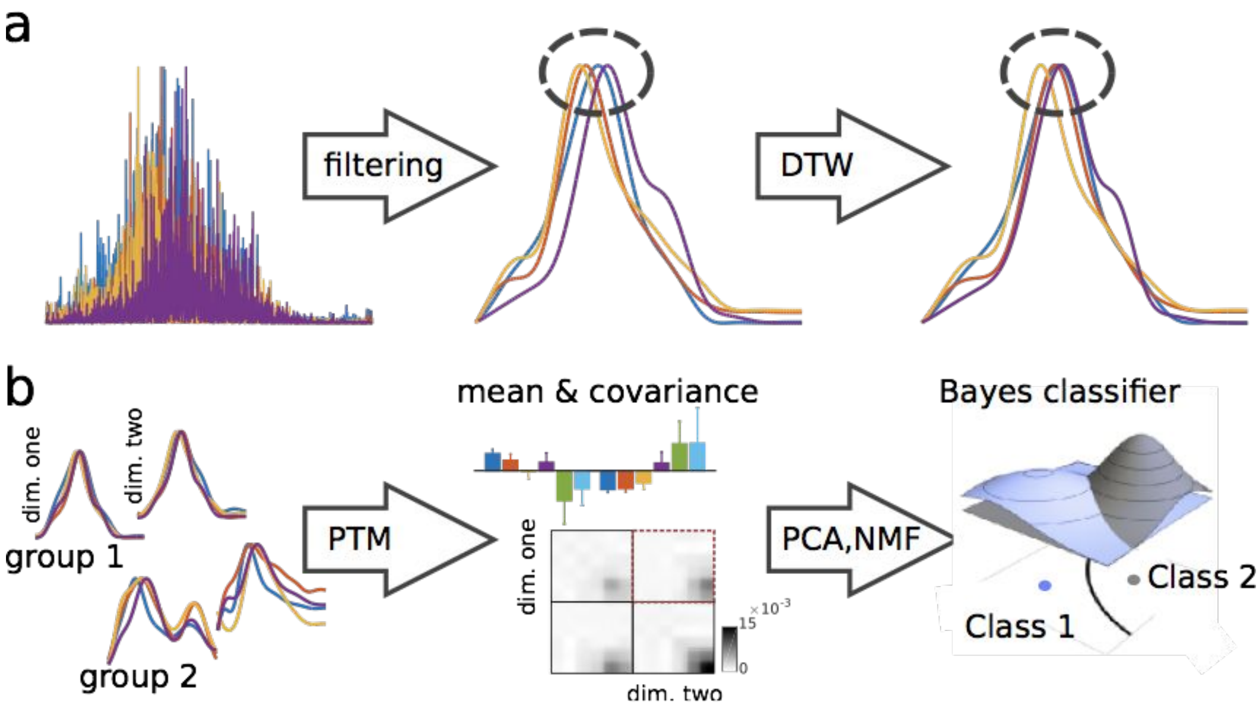

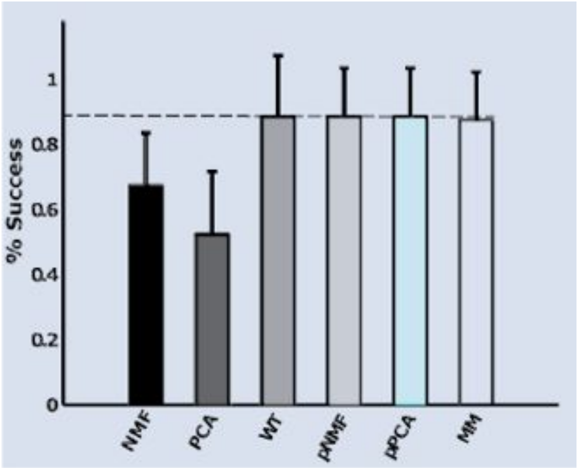

more at: https://rob.ai-lab.science/publications/

Rueckert, Elmar; Camernik, Jernej; Peters, Jan; Babic, Jan
Probabilistic Movement Models Show that Postural Control Precedes and Predicts Volitional Motor Control Journal Article
Nature Publishing Group: Scientific Reports, 6 (28455), 2016.

Rueckert, Elmar; Mundo, Jan; Paraschos, Alexandros; Peters, Jan; Neumann, Gerhard
Extracting Low-Dimensional Control Variables for Movement Primitives Inproceedings
Proceedings of the International Conference on Robotics and Automation (ICRA), 2015.

Rueckert, Elmar; Lioutikov, Rudolf; Calandra, Roberto; Schmidt, Marius; Beckerle, Philipp; Peters, Jan
Low-cost Sensor Glove with Force Feedback for Learning from Demonstrations using Probabilistic Trajectory Representations Inproceedings
ICRA 2015 Workshop on Tactile and force sensing for autonomous compliant intelligent robots, 2015.

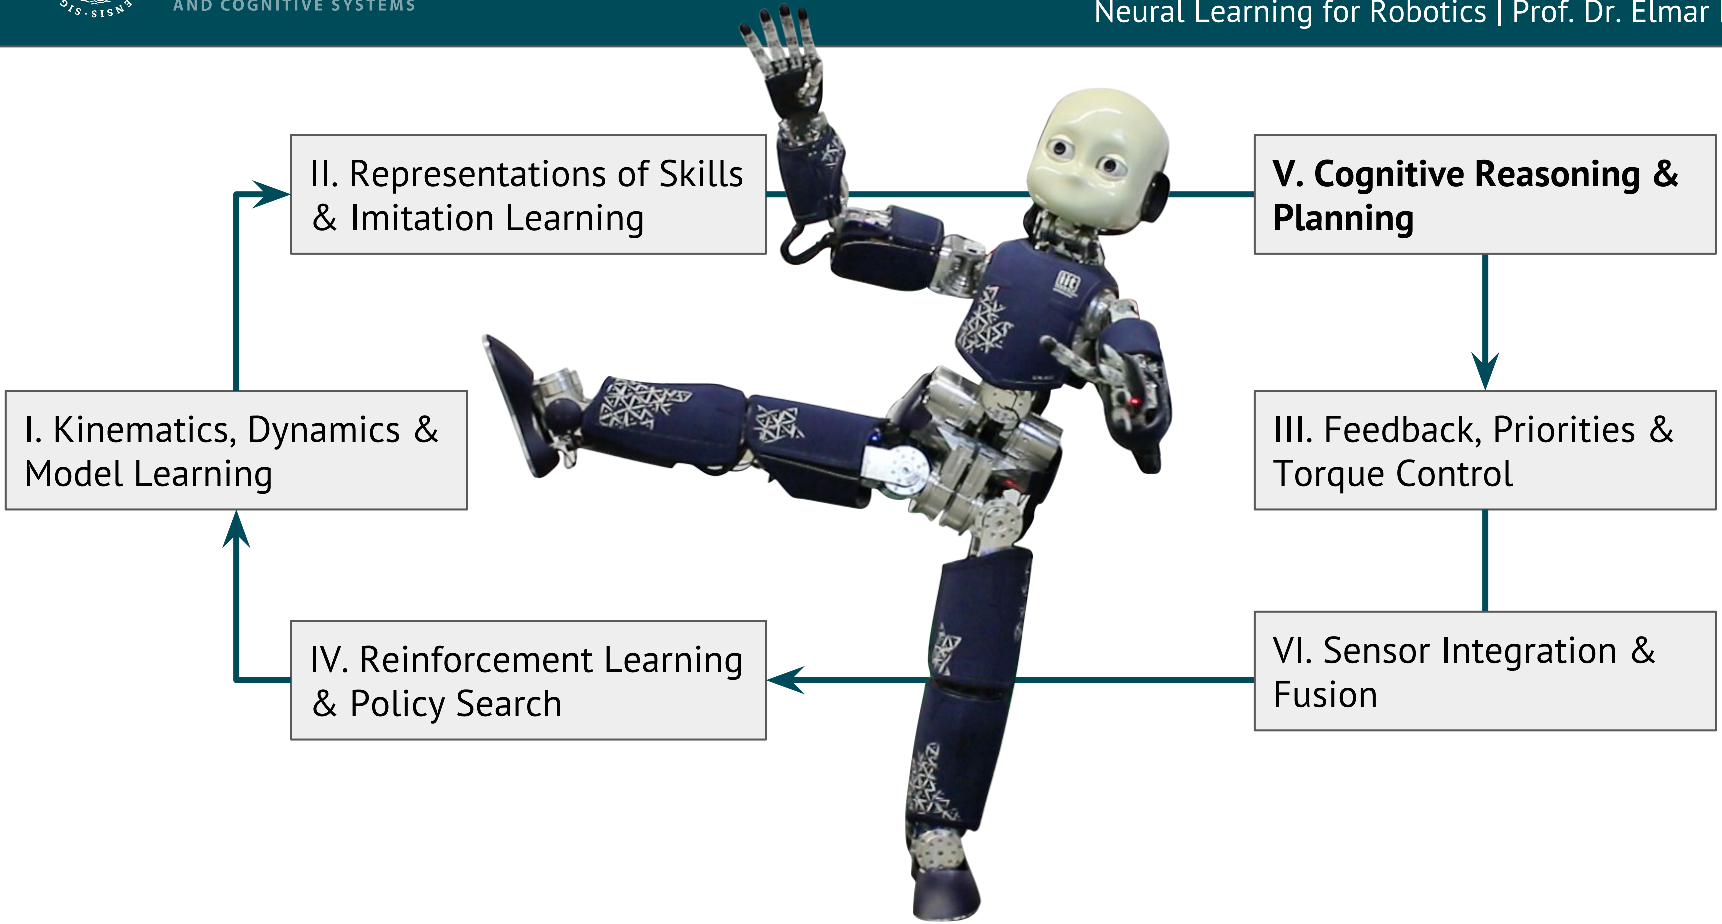
II. Representations of Skills & Imitation Learning
V. Cognitive Reasoning & Planning
I. Kinematics, Dynamics & Model Learning
III. Feedback, Priorities & Torque Control
IV. Reinforcement Learning & Policy Search
VI. Sensor Integration & Fusion

Predictive models of rats' navigation skills

Predictive models of rats' navigation skills

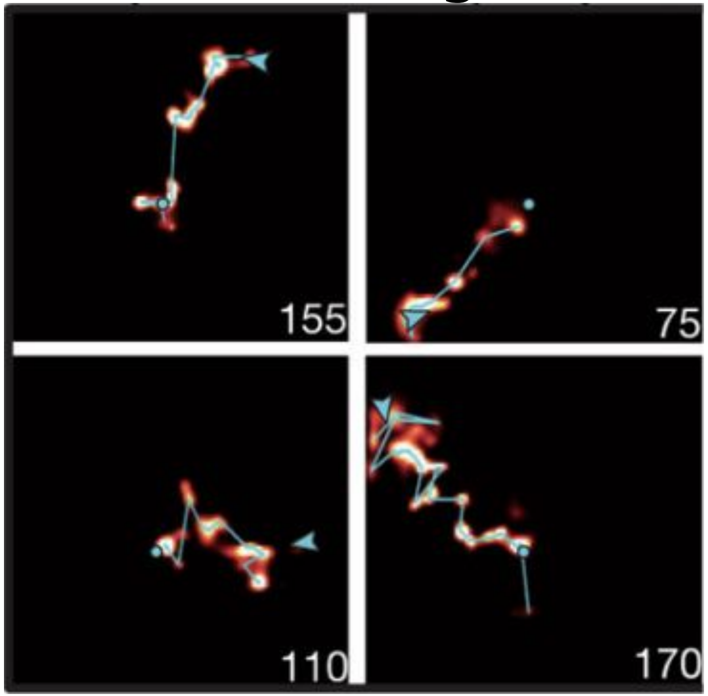

Neural Planning

$$q(\underline{\boldsymbol{\nu}}; \boldsymbol{\theta}) = p(\boldsymbol{\nu}_0) \prod_{t=1}^{T} \prod_{k=1}^{K} {\rho_{t,k}}^{\nu_{t,k}} \left(1 - \rho_{t,k}\right)^{1-\nu_{t,k}}$$

$$= p(\mathbf{v}_0) \prod_{t=1}^{T} \mathcal{T}(\mathbf{v}_t \mid \mathbf{v}_{t-1}) \phi_t(\mathbf{v}_t; \boldsymbol{\theta})$$

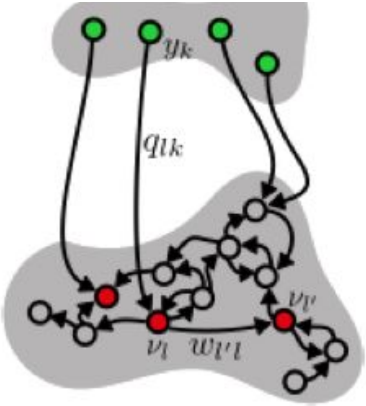

$$\mathcal{T}(\mathbf{v}_t \mid \mathbf{v}_{t-1}) = \exp\left(\sum_{i=1}^{K} w_{ki} \, v_{t-1,i} \, v_{t,k}\right)$$

$$\phi_t(\mathbf{v}_t; \boldsymbol{\theta}) = \frac{\exp\left(\sum_{j=1}^{N} \theta_{kj} \, y_{t-1,j} \, v_{t,k}\right)}{\sum_{l=1}^{K} \exp(u_{t,l})}$$

For real robot control without smoothing

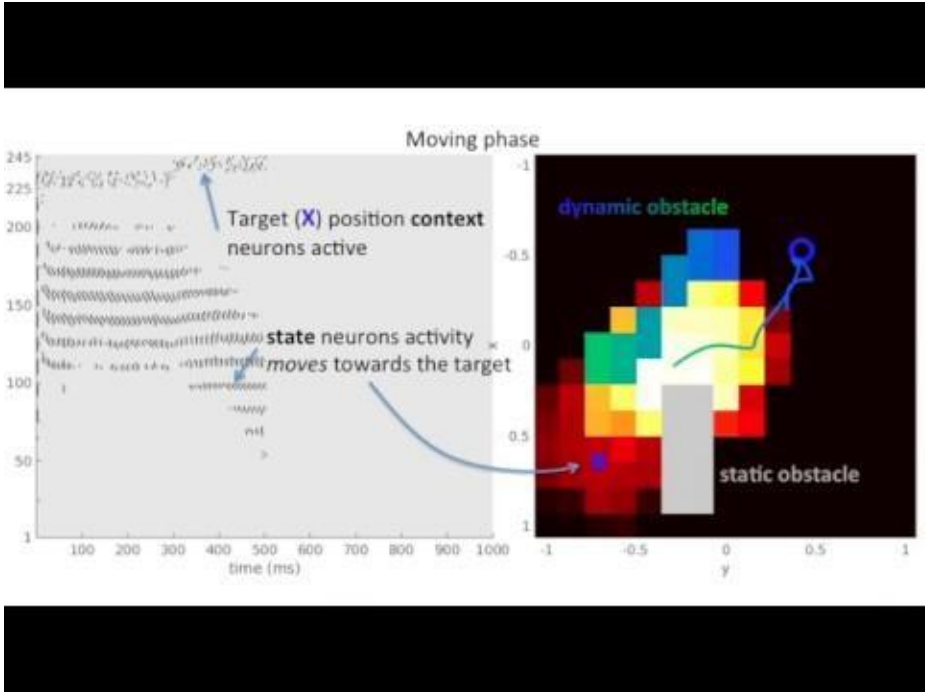

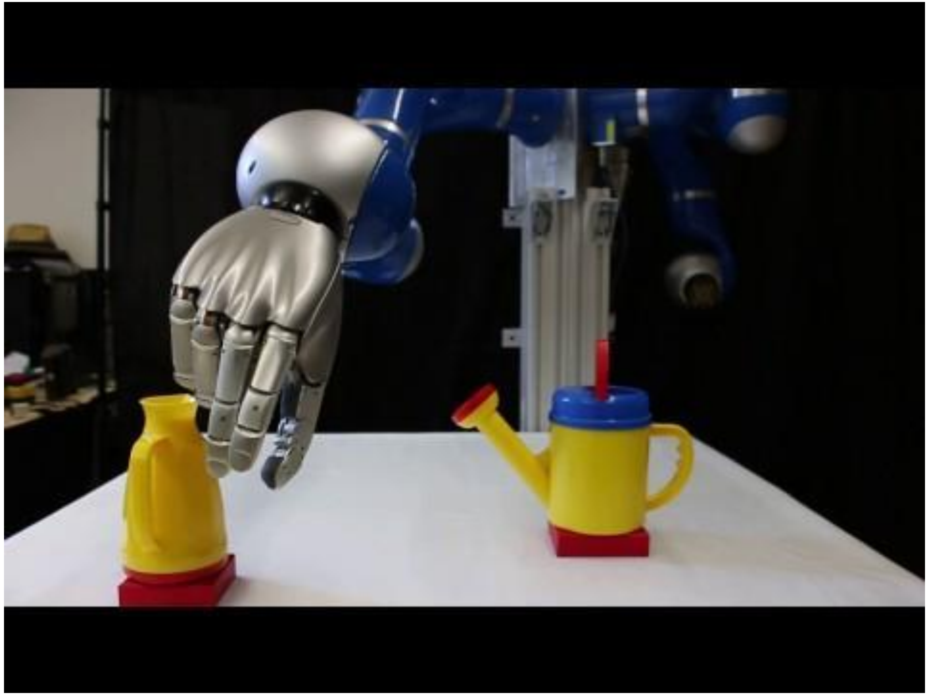

Model Learning in 15 Minutes

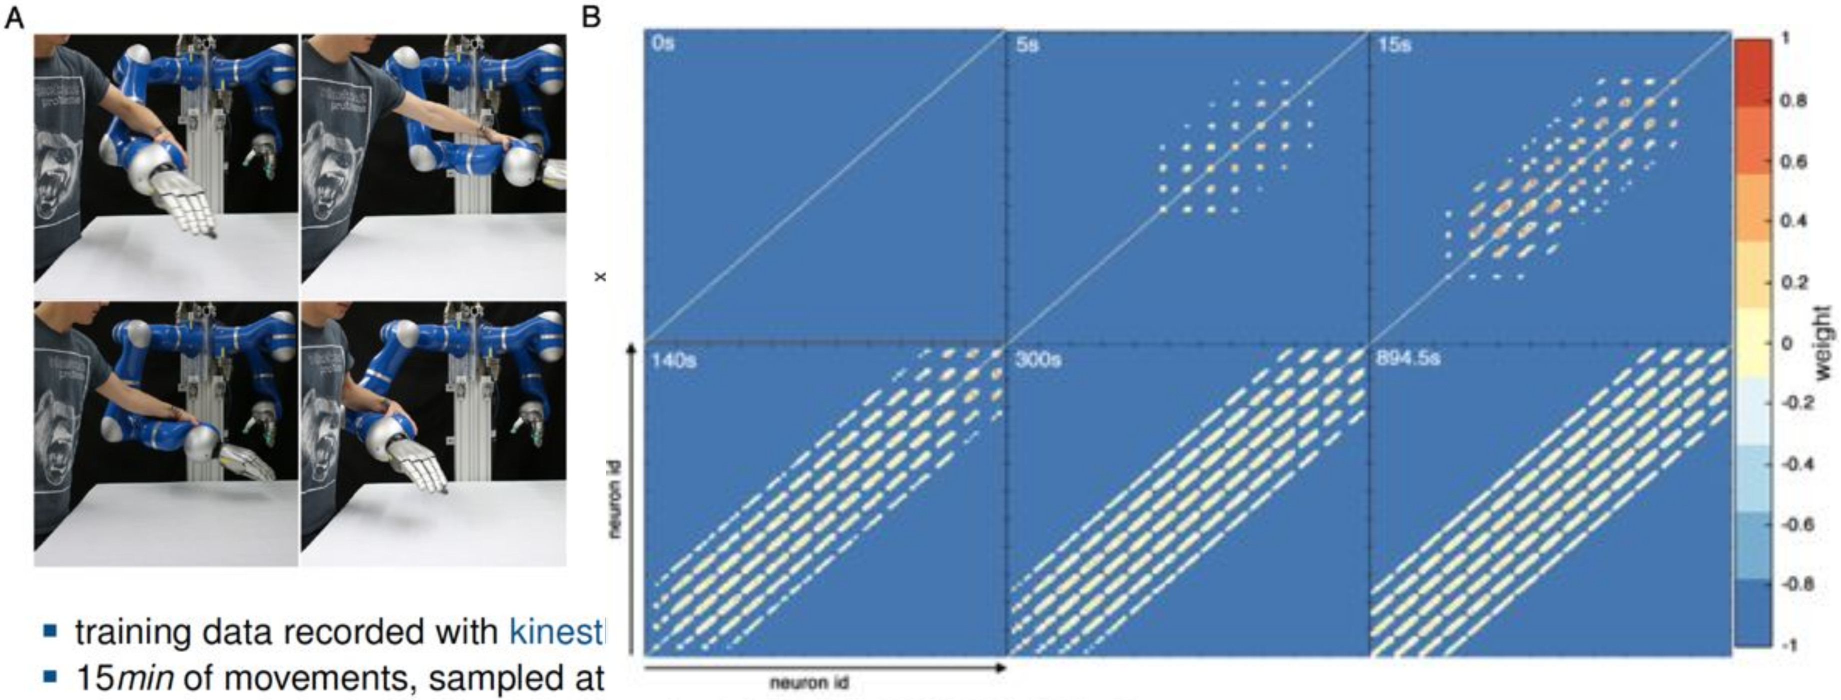

- training data recorded with kinest
- 15*min* of movements, sampled at

Real Time Adaptation and Control

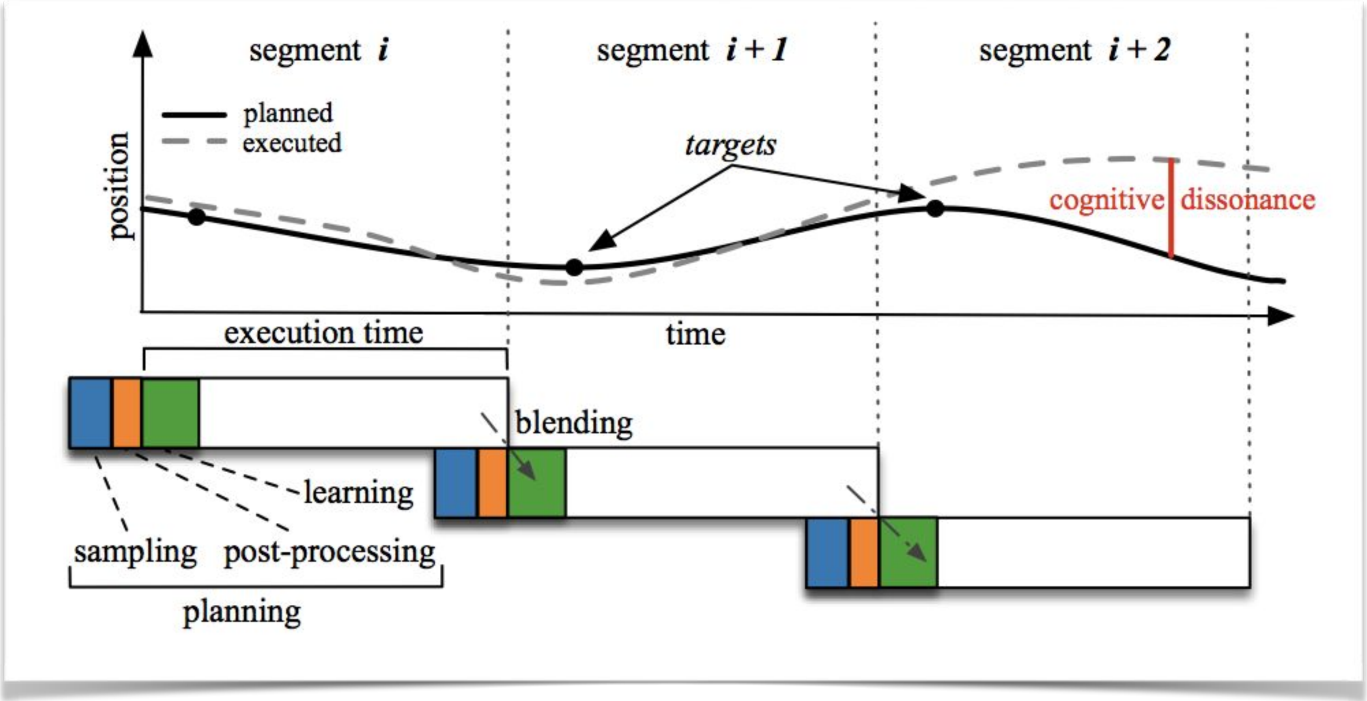

more at: https://rob.ai-lab.science/publications/

Paraschos, Alexandros; Rueckert, Elmar; Peters, Jan; Neumann, Gerhard
Probabilistic Movement Primitives under Unknown System Dynamics Journal Article
Advanced Robotics (ARJ), 2018.

Tanneberg, Daniel; Peters, Jan; Rueckert, Elmar
Online Learning with Stochastic Recurrent Neural Networks using Intrinsic Motivation Signals Inproceedings
Proceedings of the Conference on Robot Learning (CoRL), 2017.

Rueckert, Elmar; Kappel, David; Tanneberg, Daniel; Pecevski, Dejan; Peters, Jan
Recurrent Spiking Networks Solve Planning Tasks Journal Article
Nature Publishing Group: Scientific Reports, 6 (21142), 2016.

Rueckert, Elmar; Neumann, Gerhard; Toussaint, Marc; Maass, Wolfgang
Learned graphical models for probabilistic planning provide a new class of movement primitives Journal Article
Frontiers in Computational Neuroscience, 6 (97), 2013.

Summary

1. How can humans learn new motor skills within few trials?

> Learning probabilistic generative models that capture the correlations of multiple joints/signals.

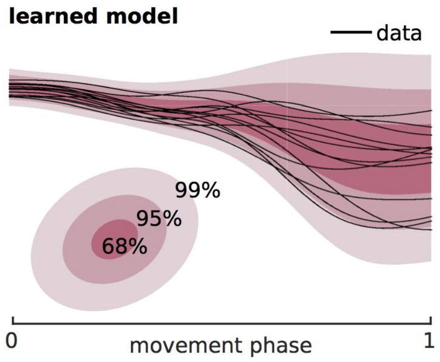

- For **noisy** and **high** dimensional **human** and **robot** data.
- Can exploit **correlations** for **predictions.**
- Low dimensional **feature** representation for **learning.**
- Generative model of **stroke-based** and **rhythmic** movements with **feedback**.

Summary

1. How do humans solve cognitive reasoning tasks in huge spaces?

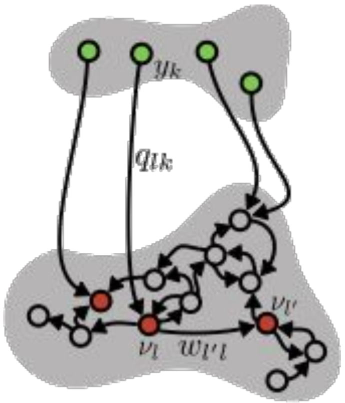

Learning stochastic neural networks grounded in the probabilistic inference framework.

- Simultaneously learning **forward, inverse kinematics** and **state transition models** through kinesthetic teaching.
- Implements **optimal planning** through reinforcement learning.
- **Online adaptation** in few seconds from **intrinsic motivation** signals.
- Model **predictive control** implementation on **real robots**.

If you are interested in the latest developments or you are looking for a job ...

- Join the linked youtube channels.
- Join mailing list like robotics-worldwide@usc.edu, researchers@pascal-network.org, ml-news@googlegroups.com .
- Write a short paper during your **bachelor or master thesis** and visit one of the major international robotics conferences.
- Visit us in our offices.

Channel of the ROB group at University of Lübeck

Thank you for your attention!

Contact:

Universität zu Lübeck
Institute for Robotics and Cognitive Systems
Ratzeburger Allee 160
Building 64, Room 94
23538 Lübeck, Deutschland

Telefon: +49 (0) 451 3101 5201
E-Mail: rieckhoff@rob.uni-luebeck.de

Disclaimer: