UNIVERSITÄT ZU LÜBECK

Artificial Intelligence for Autonomous Systems

Prof. Dr. Elmar Rueckert

February 3rd, 2020. KI-Kolloquium, Universität zu Lübeck

Now for two years Juniorprofessor for Robotics at the University of Luebeck

- Winner of the *German AI-Young Researcher Award 2019* with 15,000€.
- Chair of the experts committee on *fundamentals of intelligent learning systems*.

Introduction & Motivation

Humanoid robots are among the most complex machines on earth.

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

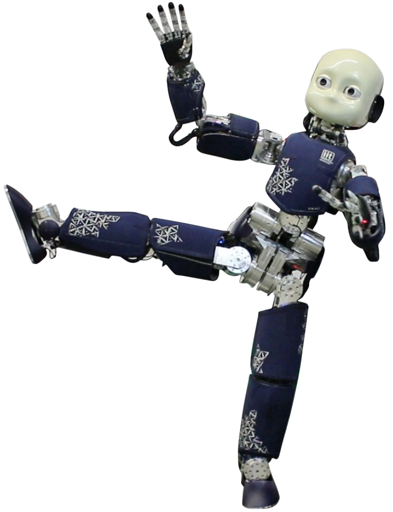

The challenges in understanding humans and in building intelligent humanoids are converging, **but …**

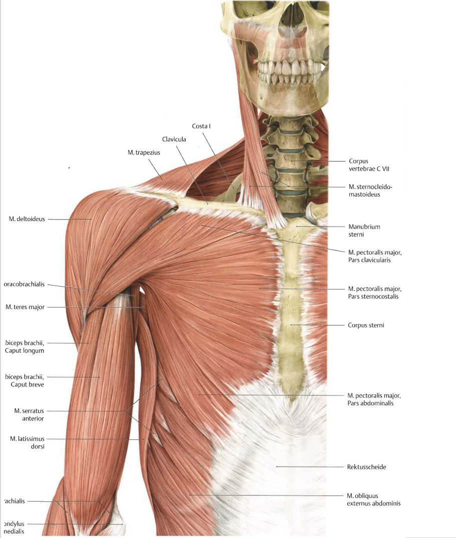

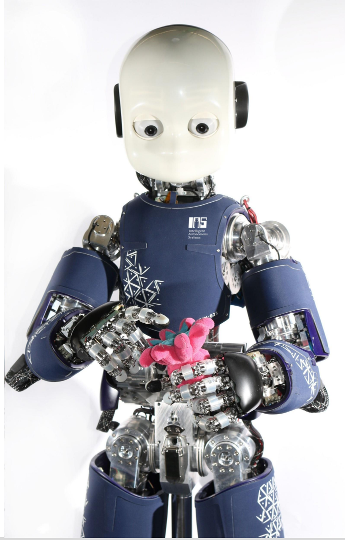

~ 100 / 53	joints
~ 100 / 1.8 10^6	photo receptors
~ 100 (finger tips) / 2000	tactile receptors

robot vision is richer & more precise.
robot motion is faster & more accurate.

However, their motor skills are **inferior to humans, why?**

Our understanding of the human motor control system is limited.

My goal is to build *Intelligent Learning Systems* from the **interaction** of

models of human motor control,
robotics implementations
and machine learning methods.

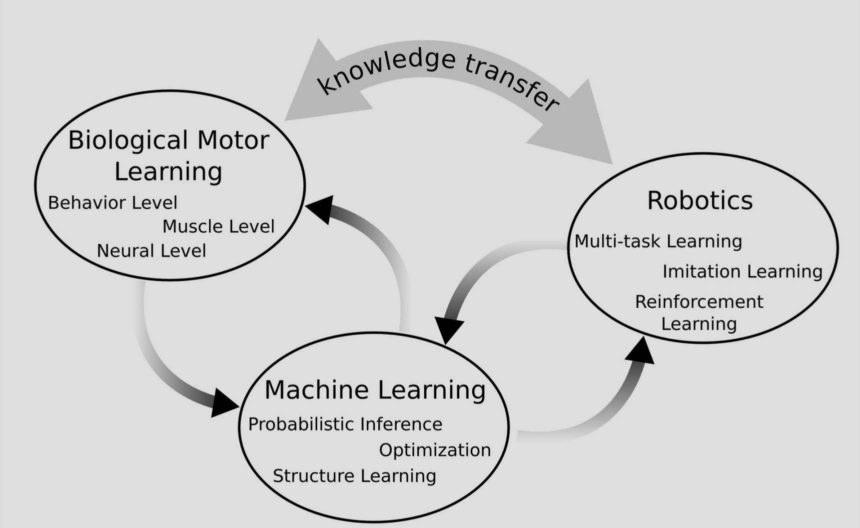

Interdisciplinary research:

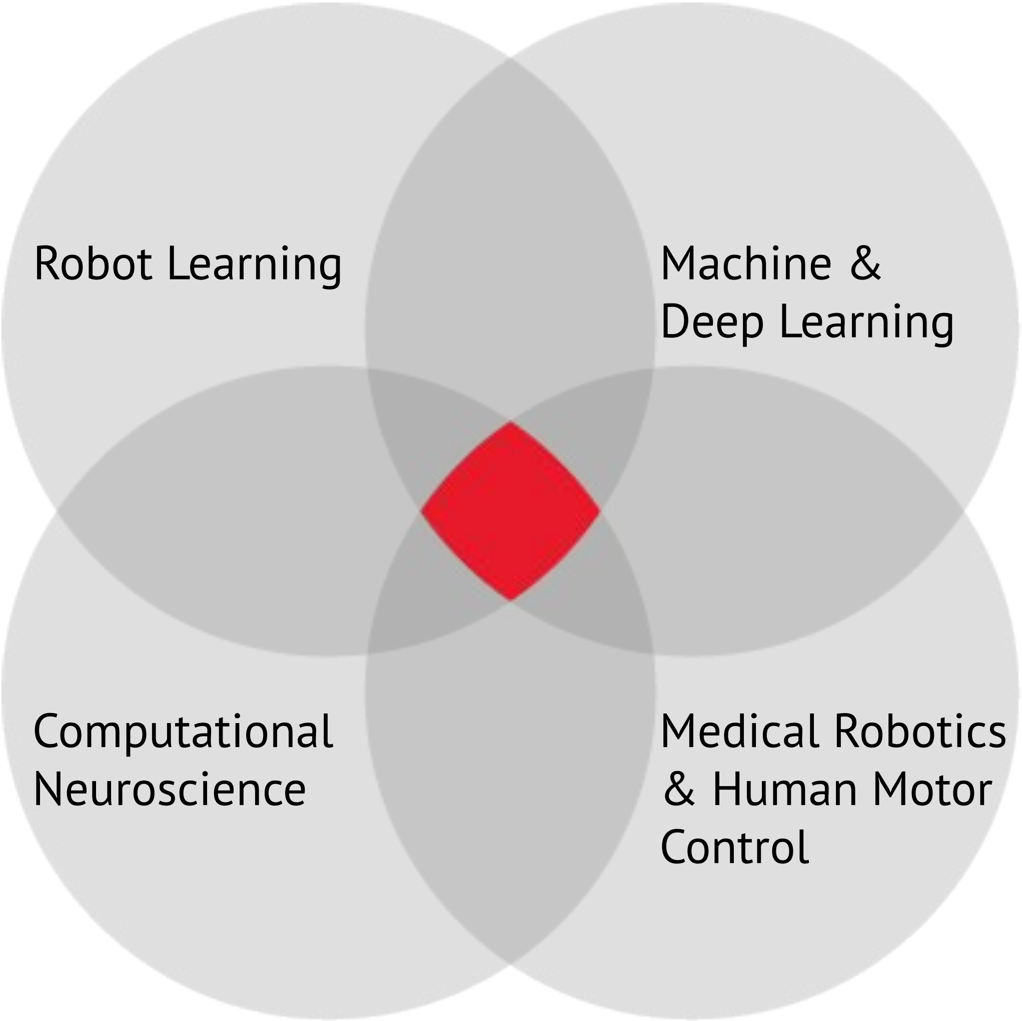

Challenges in Motor Skill Learning

Challenges in Robot Control

- High dimensional systems - Bernstein’s redundancy problem.
- Noisy environment - stochastic dynamic process.
- Limited training data - robot hardware is fragile.
- Physical interactions cannot be simulated accurately.
- Complex systems - vision, control & cognition.
- Need for efficient methods - real time processing and control.
- Scalability issues - learning millions of skills, negative transfer, etc.

We need to combine AI & Machine Learning with the development of smart Sensors!

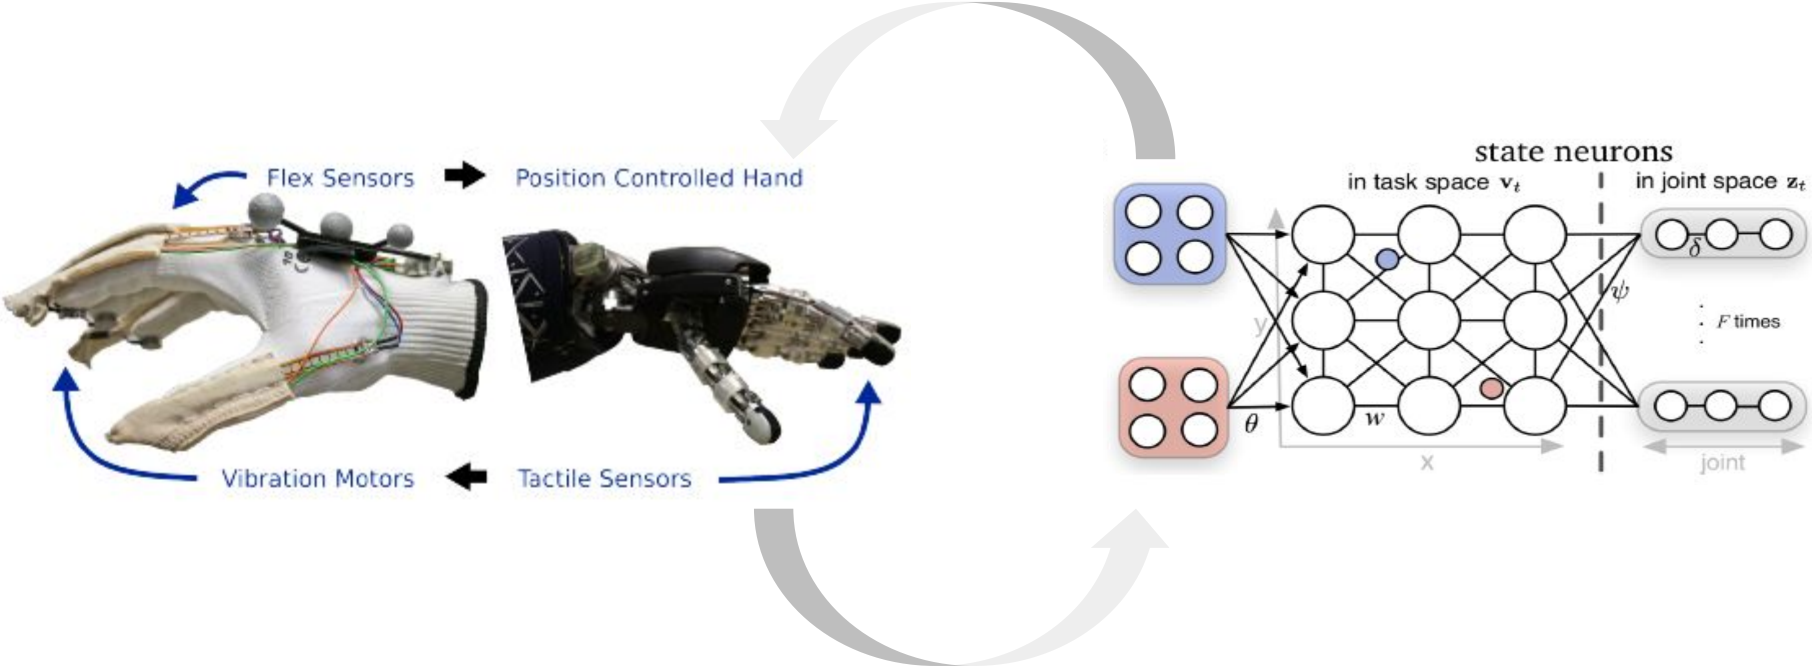

Overview:

- **High-level cognitive functions** like symbolic planning, reasoning and inference.

- **Transfer Learning of skill repertoires** like movement primitive libraries, language grammars, expert systems, multi-task transfer.

- **Task & Motion representations** and skill learning like movement primitives, single task learning with deep nets.

- **Dynamic feedback control**: closed-loop control, adaptive control, torque control.

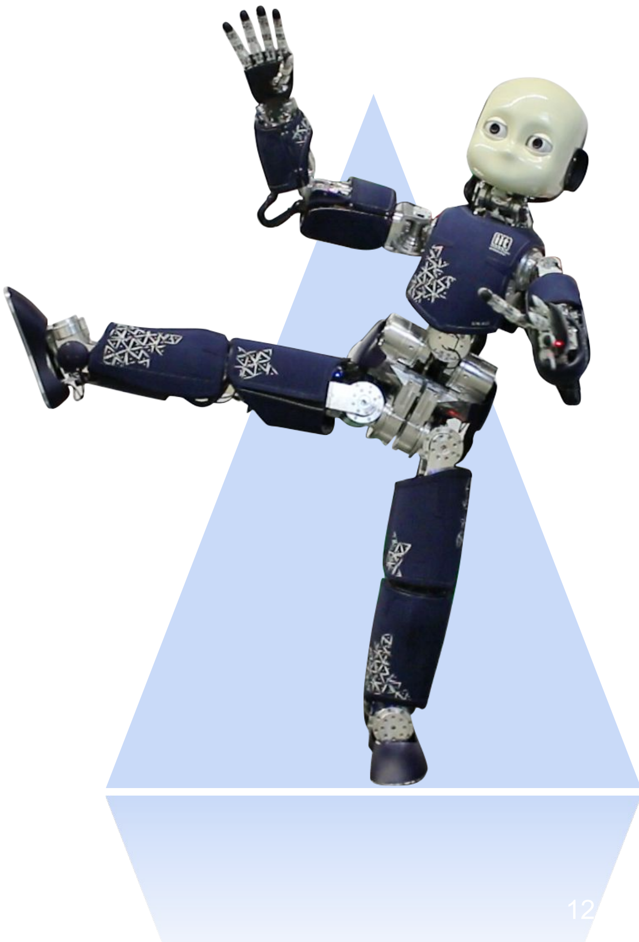

Learning (inv.) Dynamics Models in O(n) time

$$\boldsymbol{\tau} = M(\mathbf{q})\ddot{\mathbf{q}} + c(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} + g(\mathbf{q})$$

- Motor commands, e.g., torques.
- Mass matrix.
- Coriolis and centrifugal forces.
- Gravity term.

$$\boldsymbol{\tau} = f_{ID}(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}})$$

Learning (inv.) Dynamics Models in O(n) time

with Recurrent Neural Networks (LSTMs)

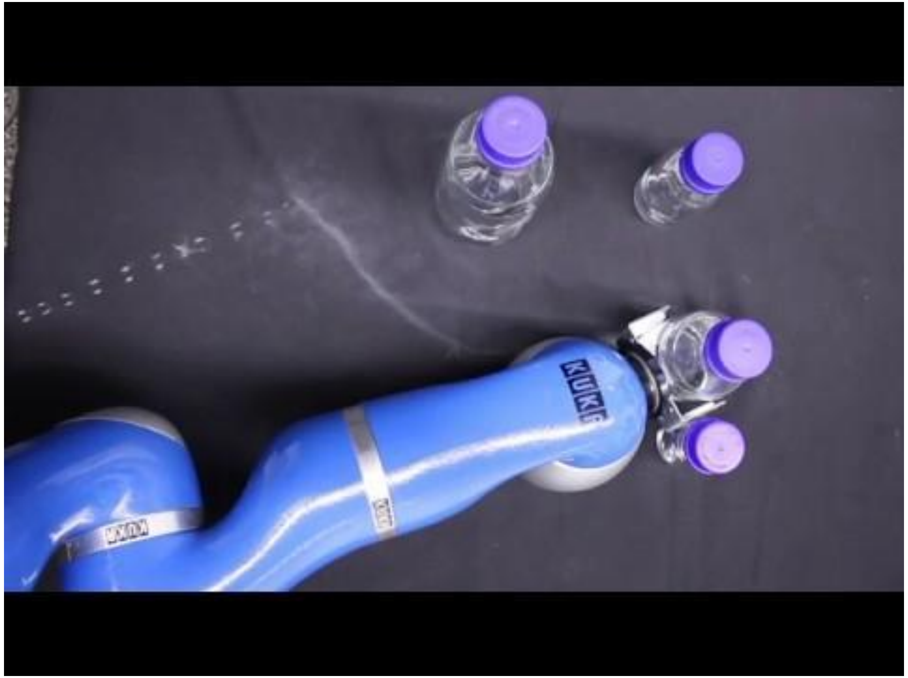

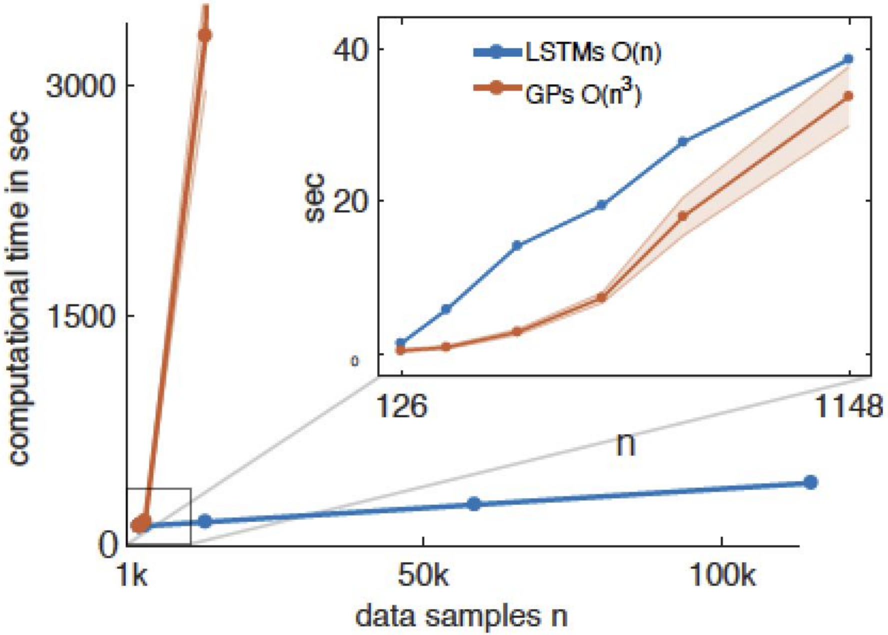

Learning (inv.) Dynamics Models with Neural Networks outperforms Gaussian Processes

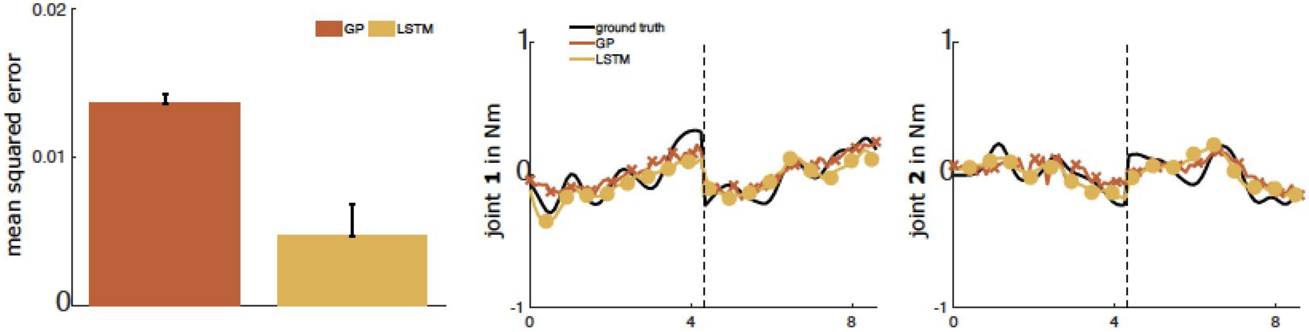

Example: Quadruped Balancing with learned dynamics models

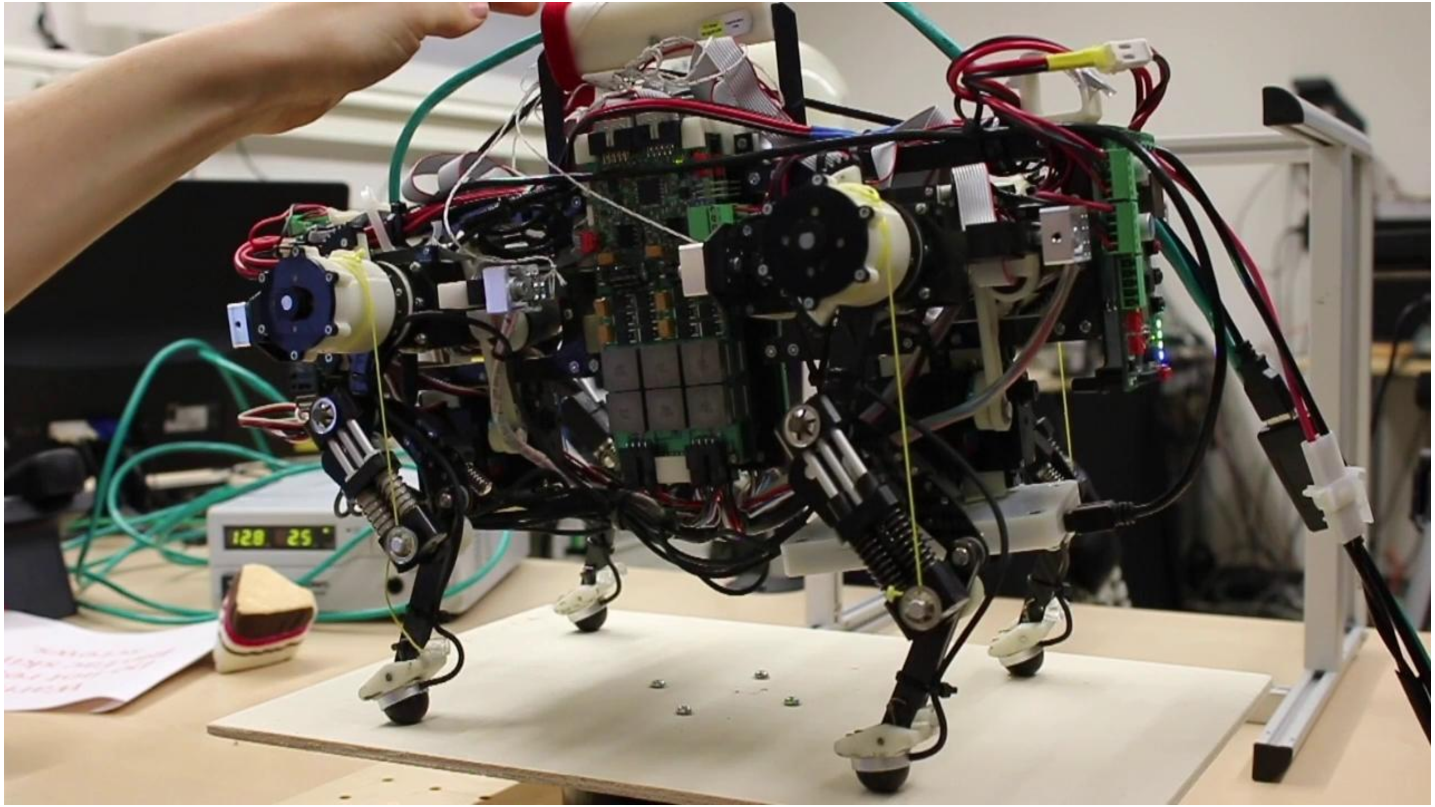

Research Overview: *3 Core Ideas illustrated in a Humanoid Robot*

- **High-level cognitive functions** like symbolic planning, reasoning and inference.
- **Transfer Learning of skill repertoires** like movement primitive libraries, language grammars, expert systems, multi-task transfer.
- **Task & Motion representations** and skill learning like movement primitives, single task learning with deep nets.
- **Dynamic feedback control**: closed-loop control, adaptive control, torque control.

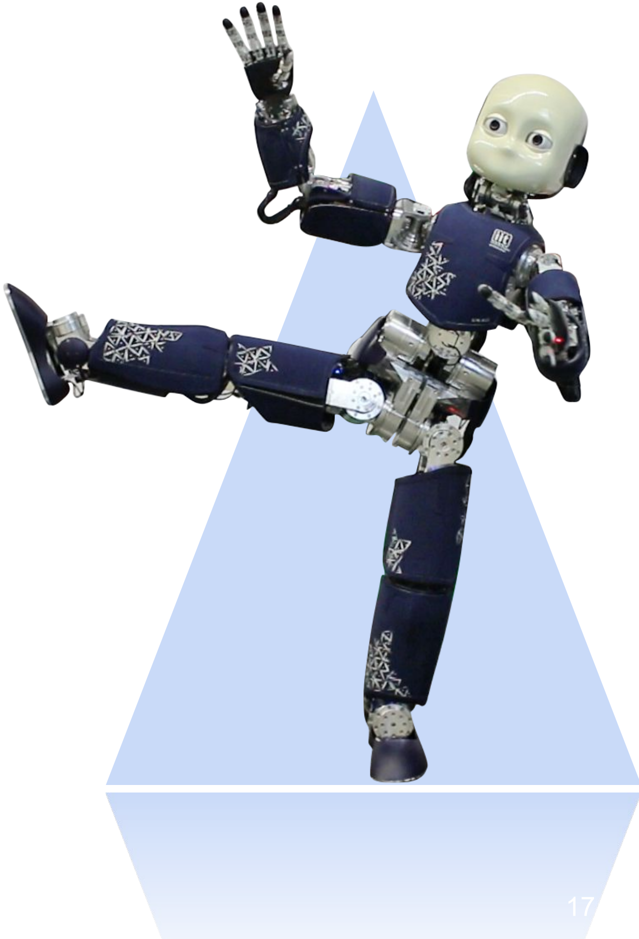

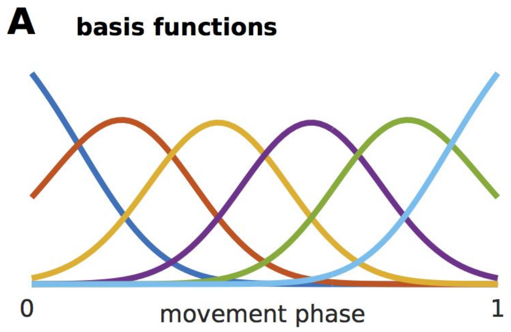
A
basis functions
0
movement phase
1

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

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

Learning through Least Squares Regression

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

learned prior

[4] Model:

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

Paraschos, Alexandros; **Rueckert, Elmar**; Peters, Jan; Neumann, Gerhard. Probabilistic Movement Primitives under Unknown System Dynamics. *Advanced Robotics (ARJ), 32 (6), pp. 297-310*, 2018. **Advanced Robotics Best Paper Award 2019.**

Sensor Glove with Vibro-Tactile Feedback

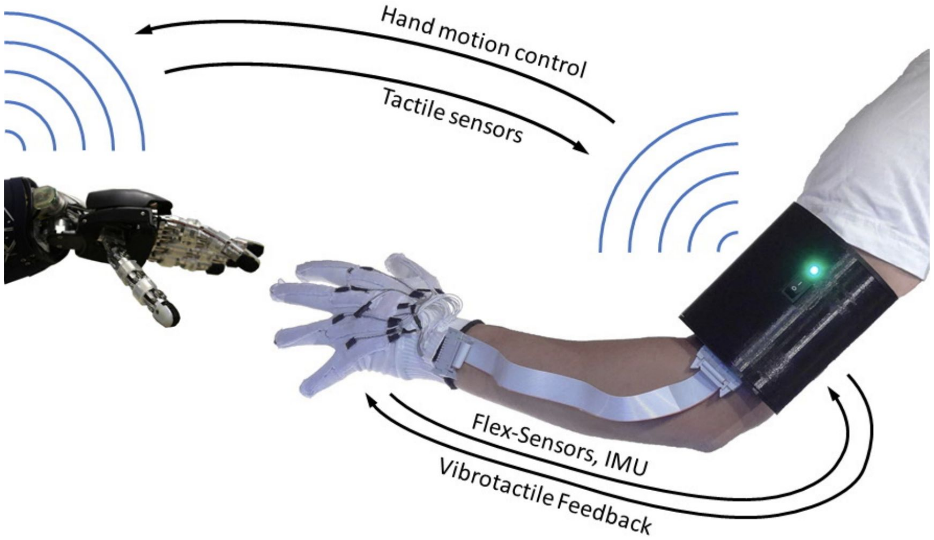

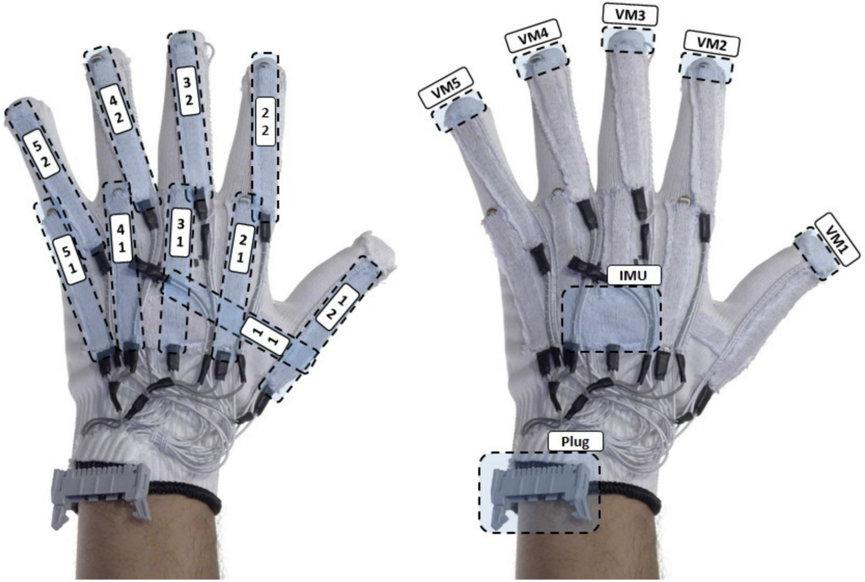

Learning Humanoid Skills from demonstrations

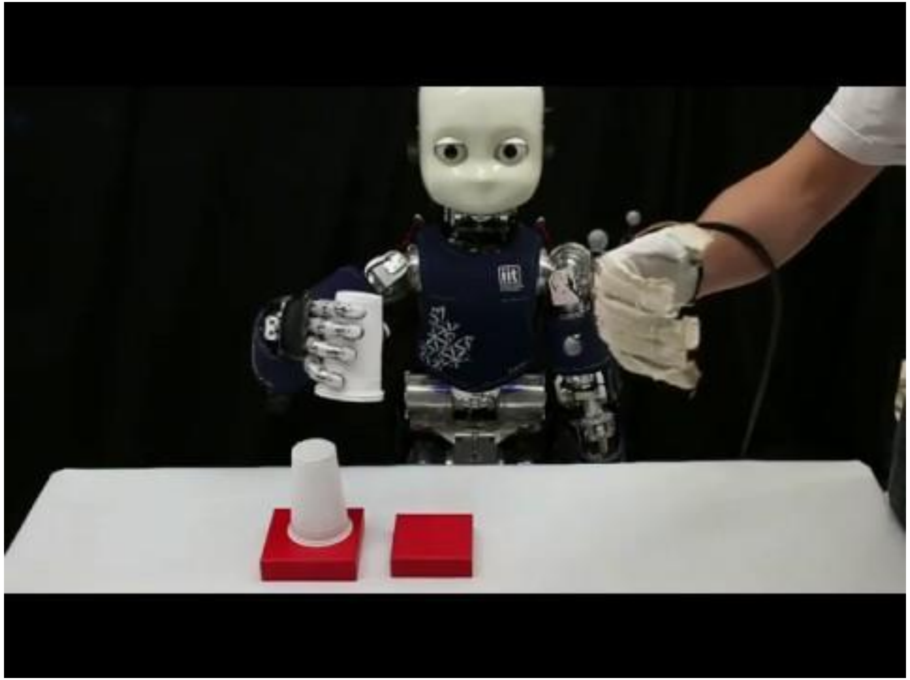

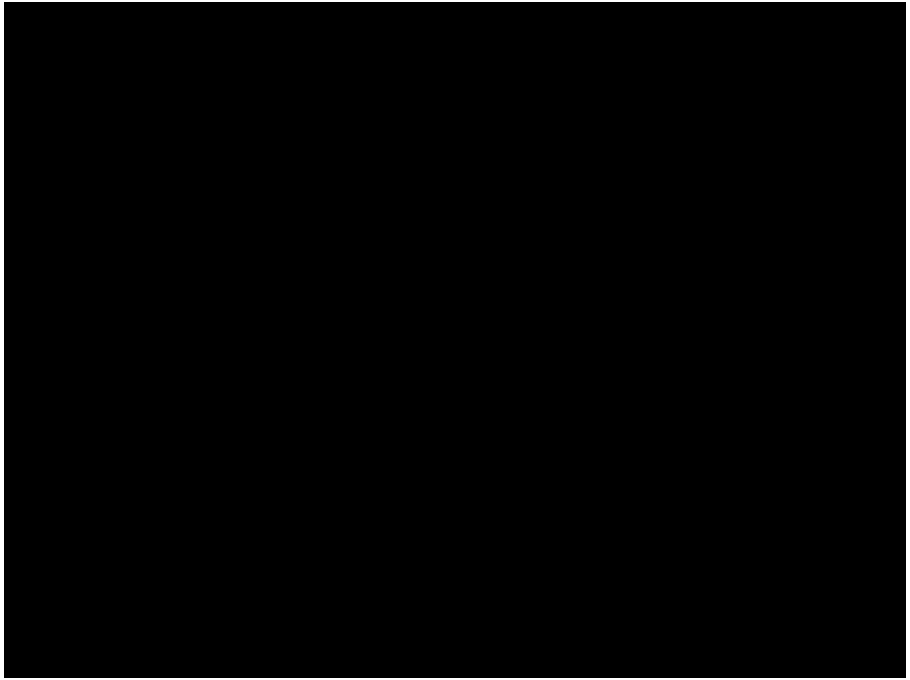

Rueckert, Elmar; Lioutikov, R.; Calandra, R.; Schmidt, M.; Beckerle, P.; Peters, J.. Low-cost Sensor Glove with Force Feedback for Learning from Demonstrations using Probabilistic Trajectory Representations. *International Conference on Robotics and Automation, Workshop Paper (ICRA)*, 2015.

Learning Human Skills from observations

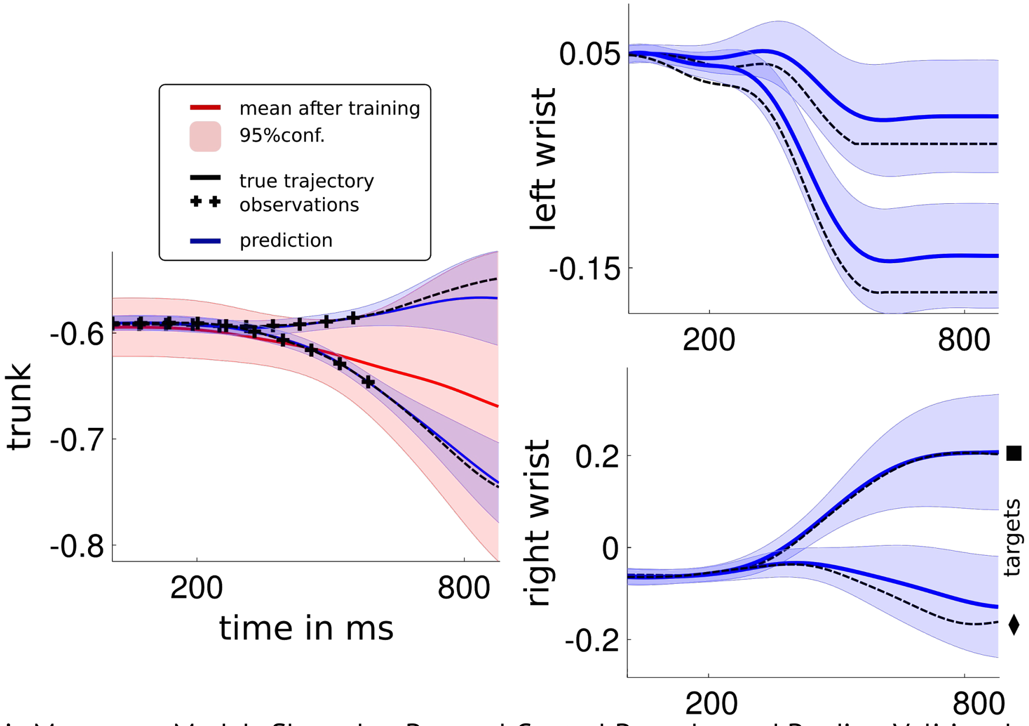

Rueckert, Elmar; Camernik, Jernej; Peters, Jan; Babic, Jan. Probabilistic Movement Models Show that Postural Control Precedes and Predicts Volitional Motor Control. ***Nature Publishing Group: Scientific Reports***, *6 (28455)*, 2016.

Learning Salient Features of Motor Skills

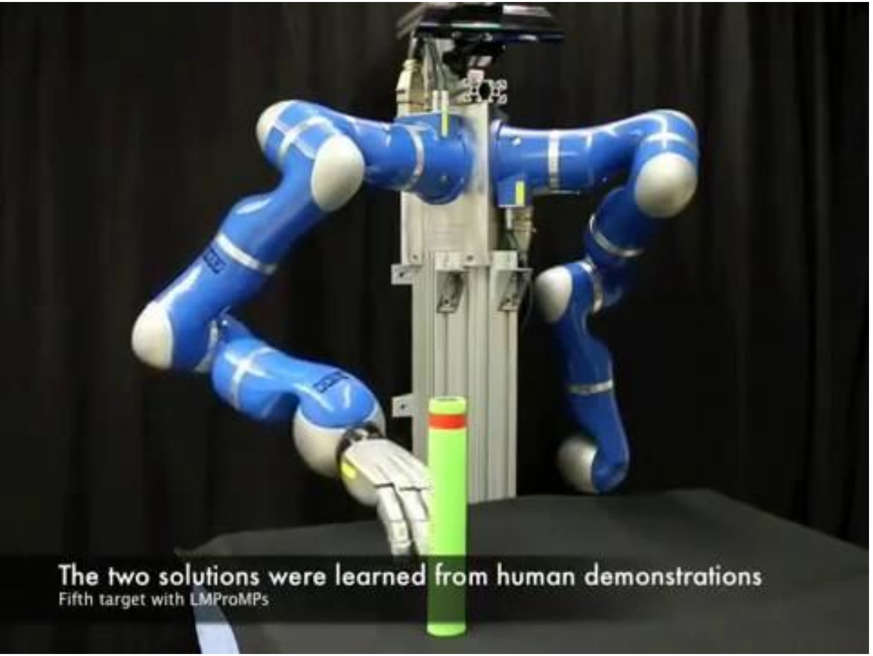

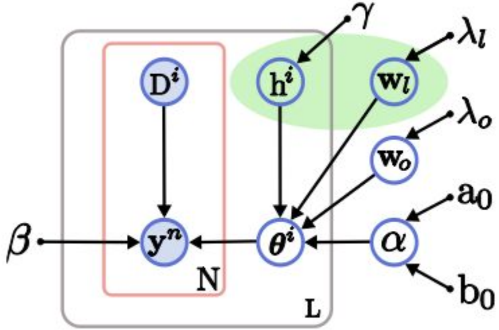

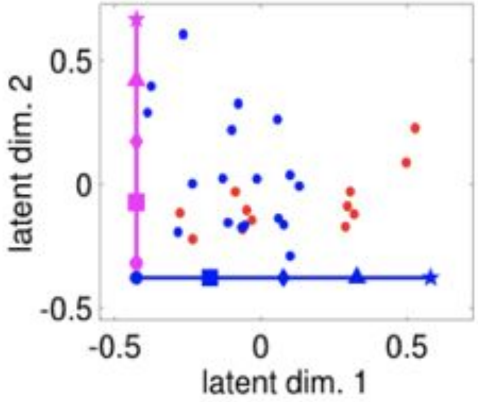

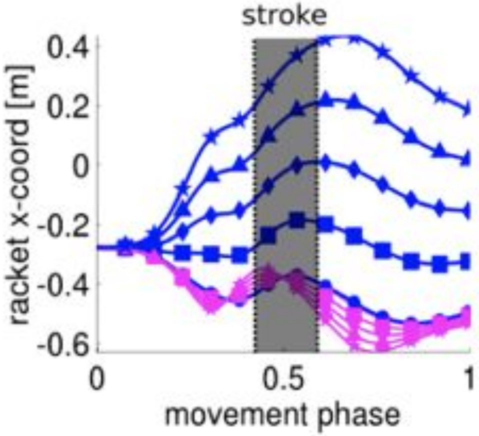

Rueckert, Elmar; Mundo, Jan; Paraschos, Alexandros; Peters, Jan; Neumann, Gerhard. Extracting Low-Dimensional Control Variables for Movement Primitives. *In Proceedings of the International Conference on Robotics and Automation (ICRA)*, 2015.

Research Overview: *3 Core Ideas illustrated in a Humanoid Robot*

- **High-level cognitive functions** like symbolic planning, reasoning and inference.
- **Transfer Learning of skill repertoires** like movement primitive libraries, language grammars, expert systems, multi-task transfer.
- **Task & Motion representations** and skill learning like movement primitives, single task learning with deep nets.
- **Dynamic feedback control**: closed-loop control, adaptive control, torque control.

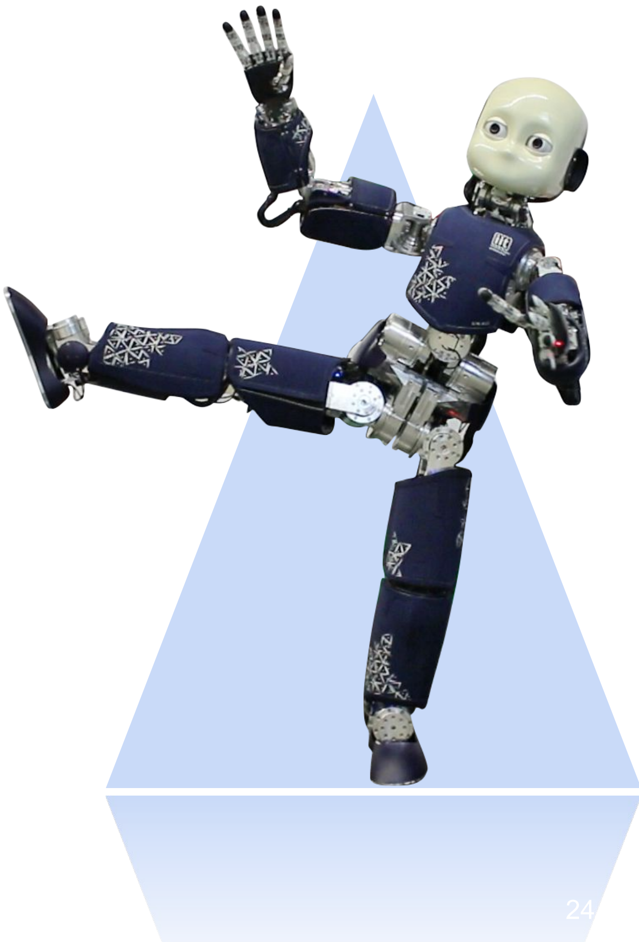

Incremental Imitation learning a primitive library

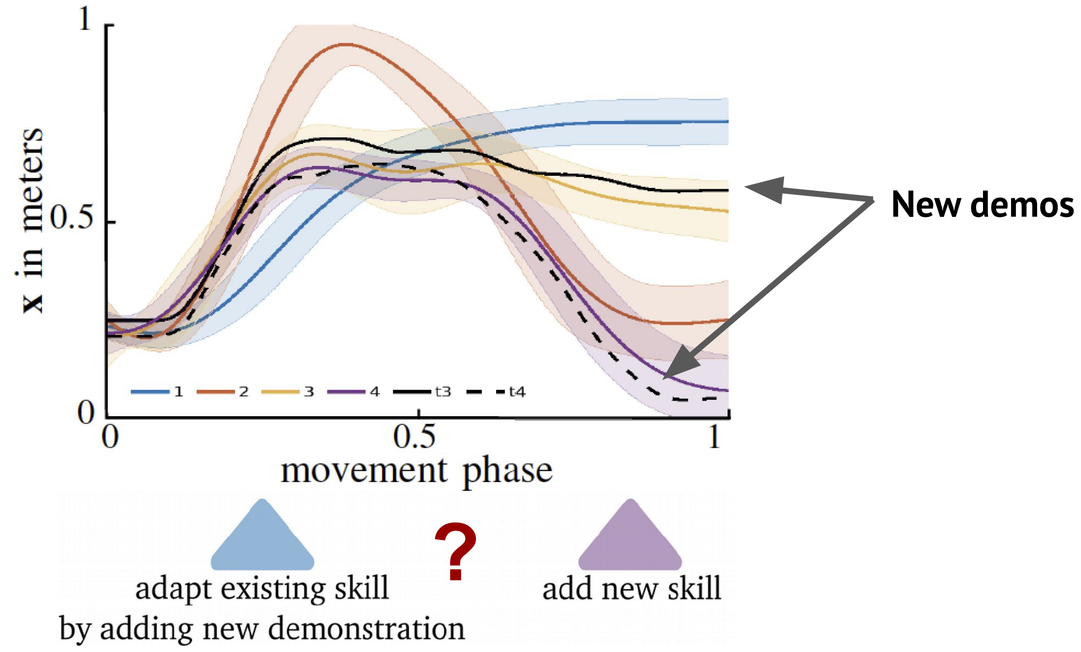

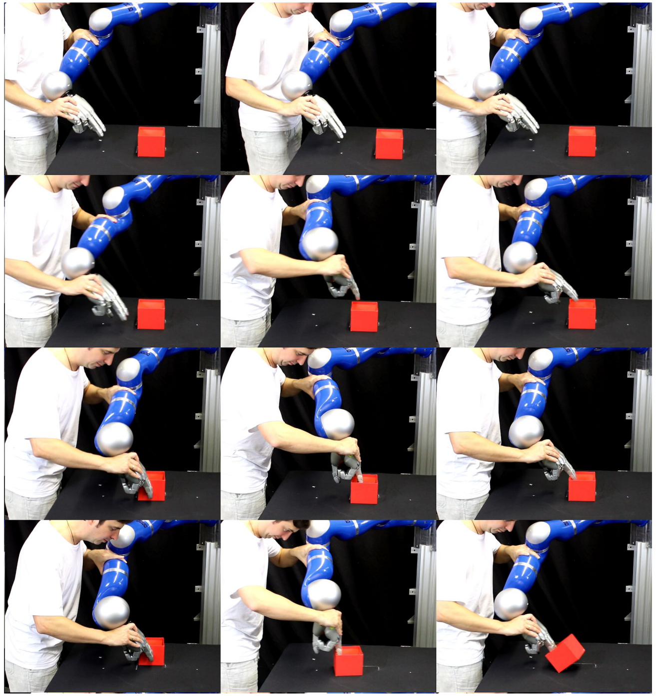

Stark, Svenja; Peters, Jan; **Rueckert, Elmar**. A Comparison of Distance Measures for Learning Nonparametric Motor Skill Libraries. *Proceedings of the International Conference on Humanoid Robots (HUMANOIDS)*, 2017.

Transfer Learning with Movement Primitives

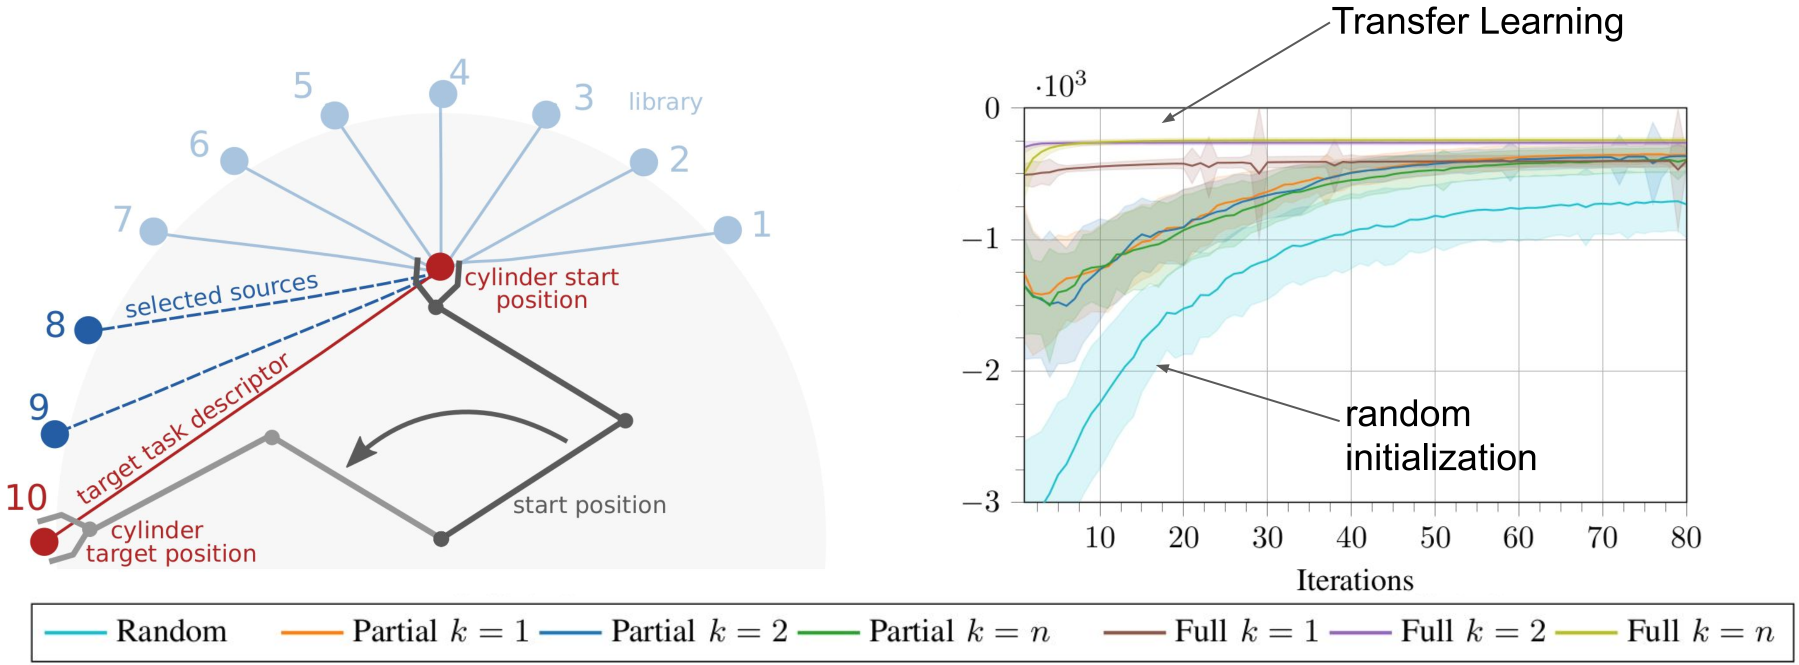

Svenja Stark, Jan Peters and **Elmar Rueckert**. Experience Reuse with Probabilistic Movement Primitives. *Proceedings of the 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019.

Research Overview: *3 Core Ideas illustrated in a Humanoid Robot*

- **High-level cognitive functions** like symbolic planning, reasoning and inference.
- **Transfer Learning of skill repertoires** like movement primitive libraries, language grammars, expert systems, multi-task transfer.
- **Task & Motion representations** and skill learning like movement primitives, single task learning with deep nets.
- **Dynamic feedback control**: closed-loop control, adaptive control, torque control.

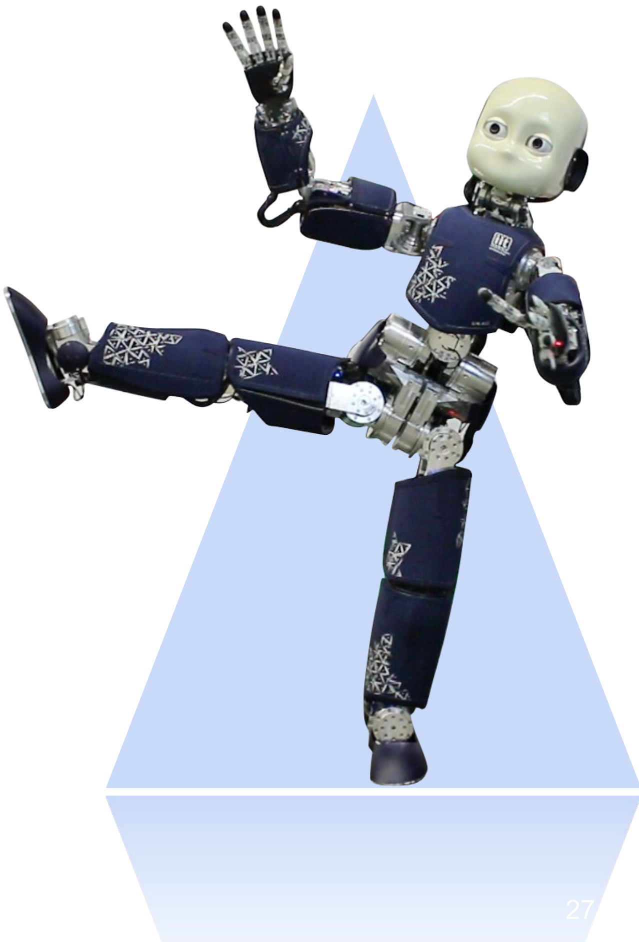

A neuroinspired plannig approach

Pfeiffer, B. & Foster, D. Hippocampal place-cell sequences depict future paths to remembered goals. *Nature 497, 74–79 (2013)*.

A neuroinspired plannig approach

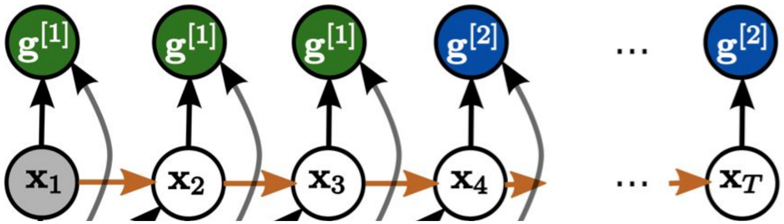

Planning as Inference

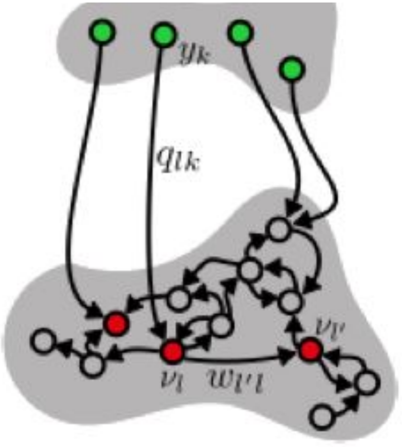

Planning in Spiking
Neural Networks

$$p\left(\underline{\boldsymbol{x}}|r=1\right)\frac{1}{\mathscr{Z}}p\left(r|\underline{\boldsymbol{x}}\right)p\left(\boldsymbol{x}_0\right)\prod_{t=1}^{T}\mathscr{T}\left(\boldsymbol{x}_t|\boldsymbol{x}_{t-1}\right)$$

Proof for optimal planning as inference in recurrent neural networks!

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

A neuroinspired plannig approach

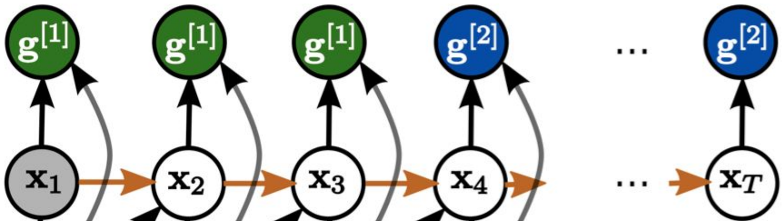

Planning as Inference

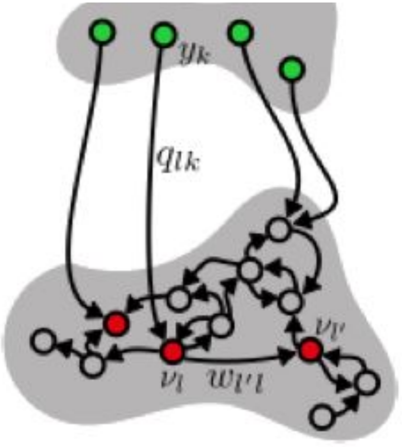

Planning in Spiking Neural Networks

$$p(\underline{x}|r=1)\,\frac{1}{\mathcal{Z}}\,p(r|\underline{x})\,p(x_0)\prod_{t=1}^{T}\mathcal{T}(x_t|x_{t-1})$$

Reward modulated Hebbian Learning

Supervised Learning

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

A neuroinspired plannig approach

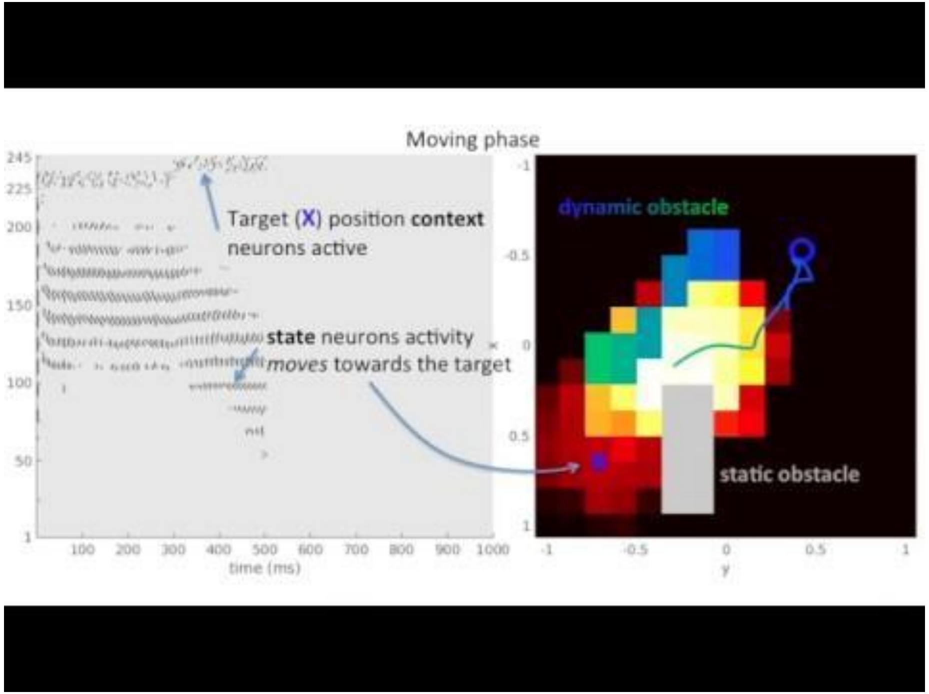

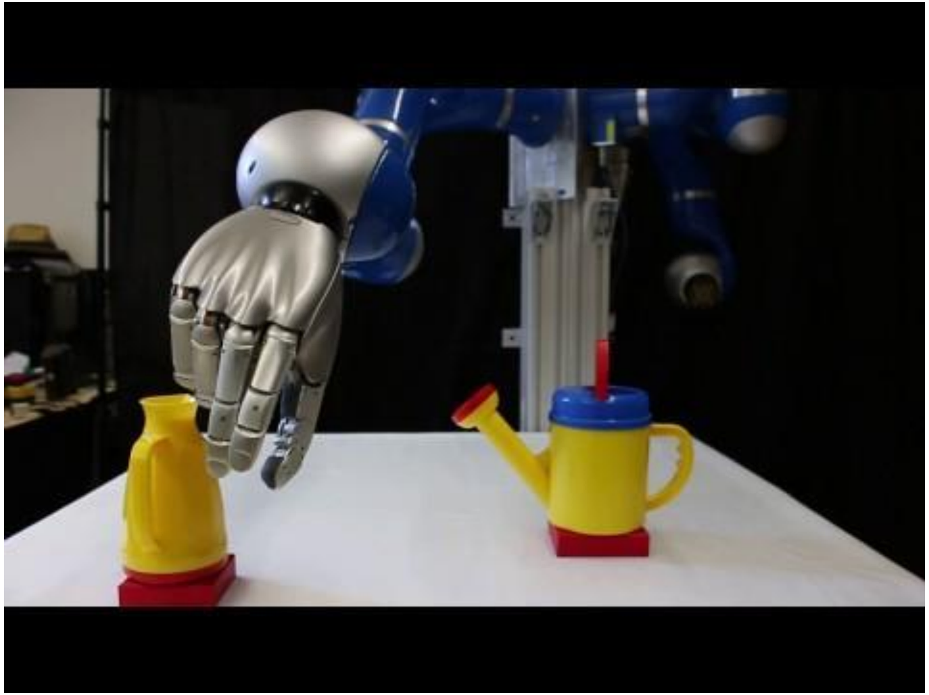

Rueckert, Elmar; Kappel, David; Tanneberg, Daniel; Pecevski, Dejan; Peters, Jan. Recurrent Spiking Networks Solve Planning Tasks. Nature Publishing Group: Scientific Reports, 6 (21142), 2016.

Real-Time Planning & Control

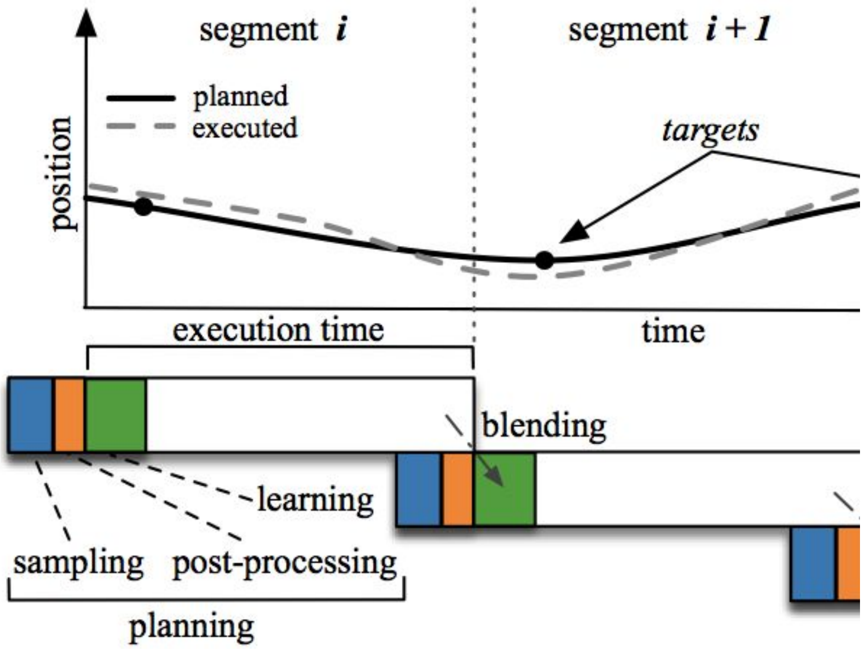

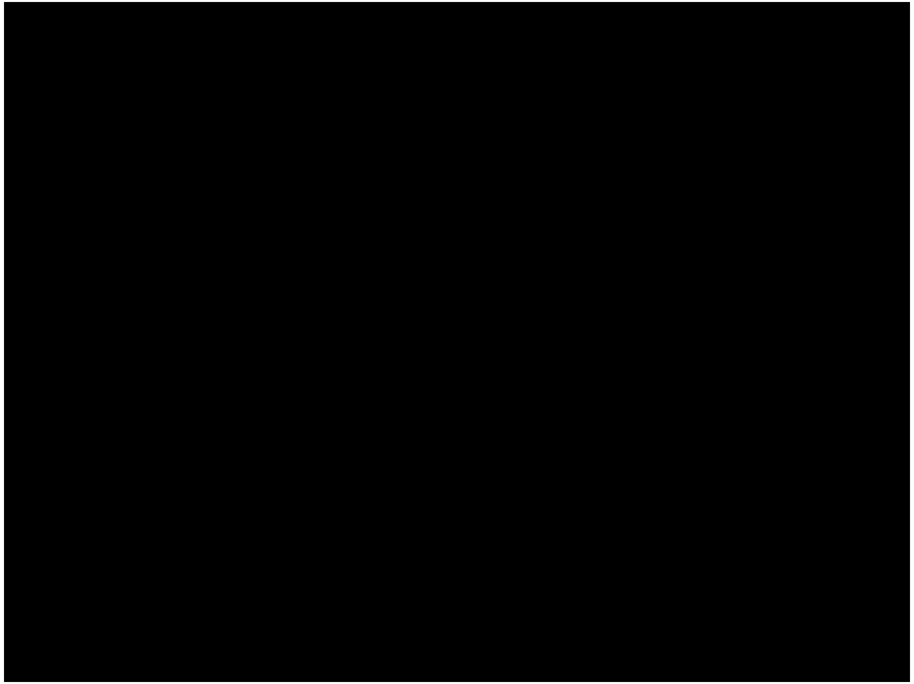

Tanneberg, Daniel; Peters, Jan; **Rueckert, Elmar. I**ntrinsic Motivation and Mental Replay enable Efficient Online Adaptation in Stochastic Recurrent Networks. **Neural Networks - Elsevier**, 109 , pp. 67-80, 2019, ISBN: 0893-6080, (**Impact Factor of 7.197** (2017)).

Example: Neural planning with binary sensors

- **Novel sensor for grass and plant detection (patent pending),** low cost sensor (industrial cooperation).
- **Mobile localization from Odometrie only**.
- **Loop Closure detection and optimization**.
- **Complete coverage path planning with neural networks**.

Example: Neural planning with binary sensors

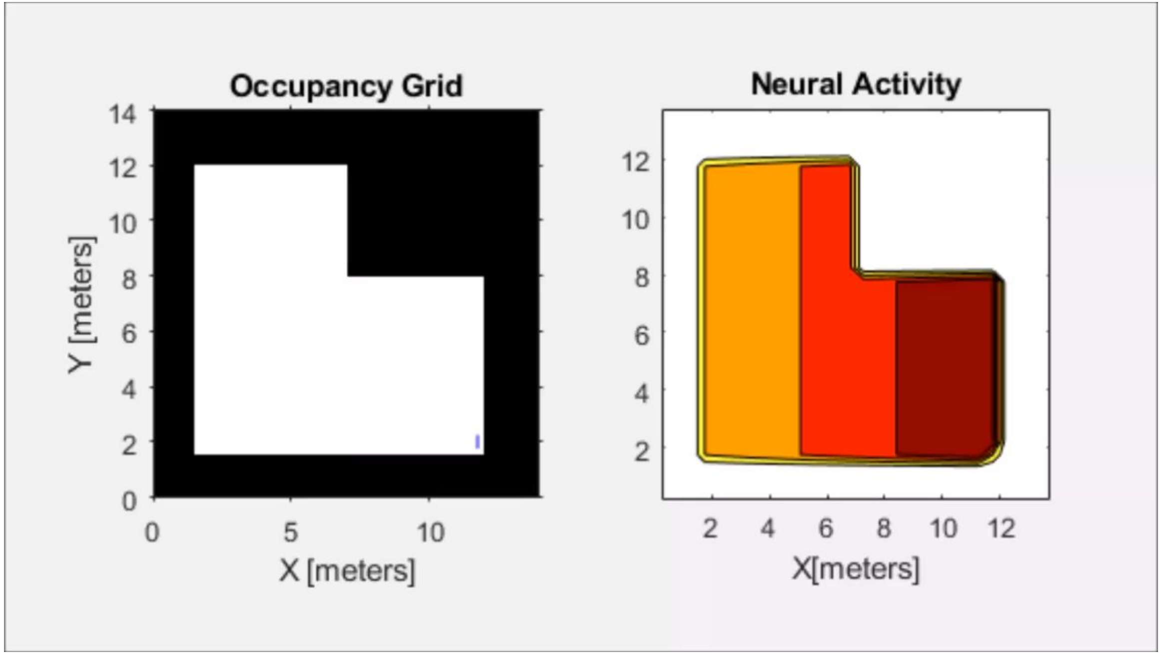

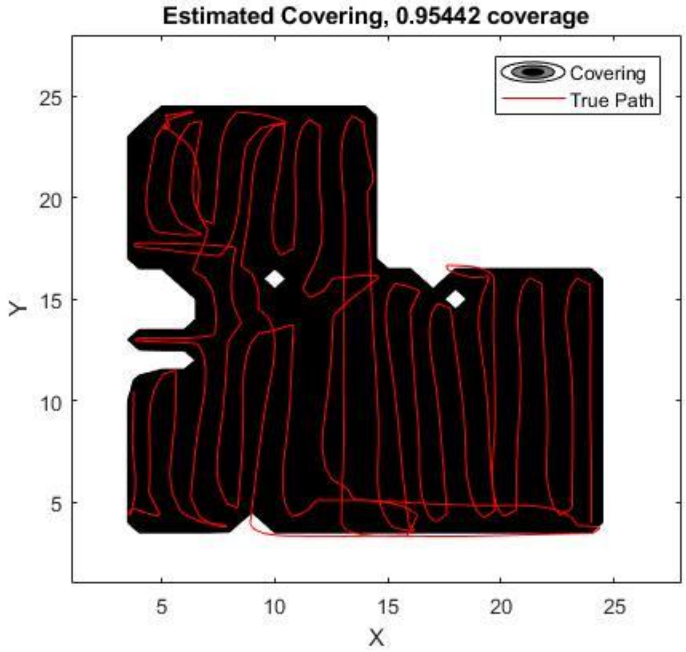

Rottmann, N; Bruder, R; Schweikard, A; **Rueckert, E.** Loop Closure Detection in Closed Environments. European Conference on Mobile Robots (ECMR 2019), 2019, ISBN: 978-1-7281-3605-9. Neural planning paper is in preparation.

From Motion Planning to Symbolic Planning

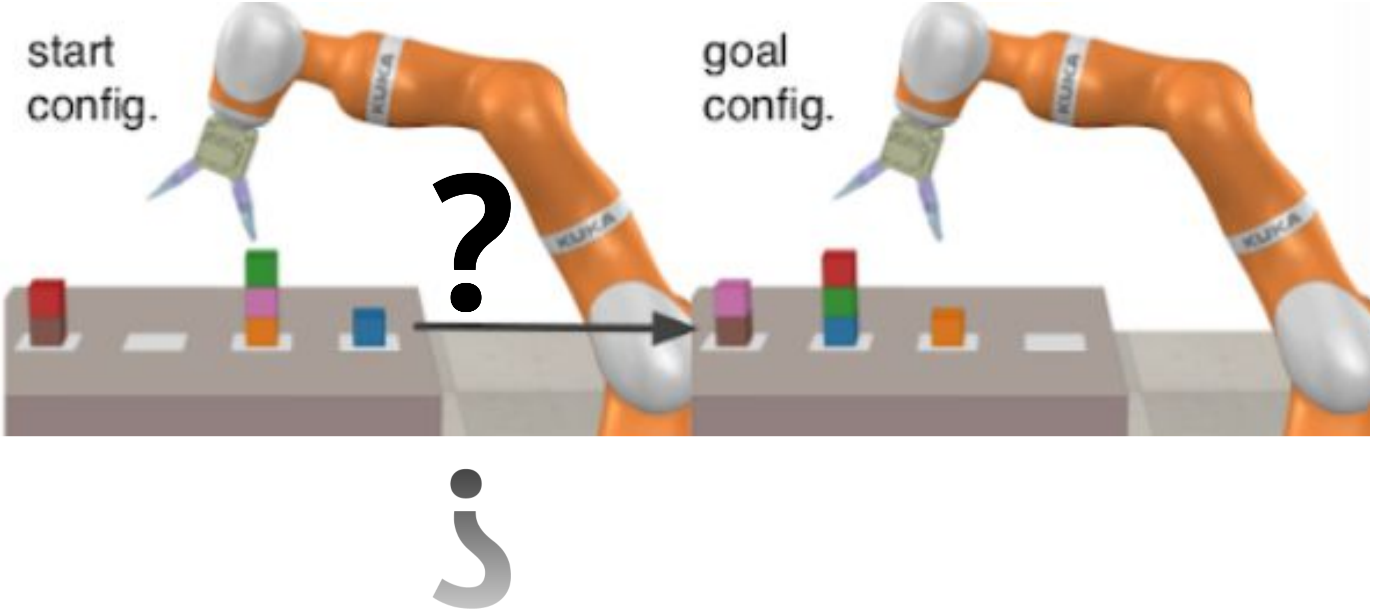

From Motion Planning to Symbolic Planning

Challenge: Learn the game, not the task!

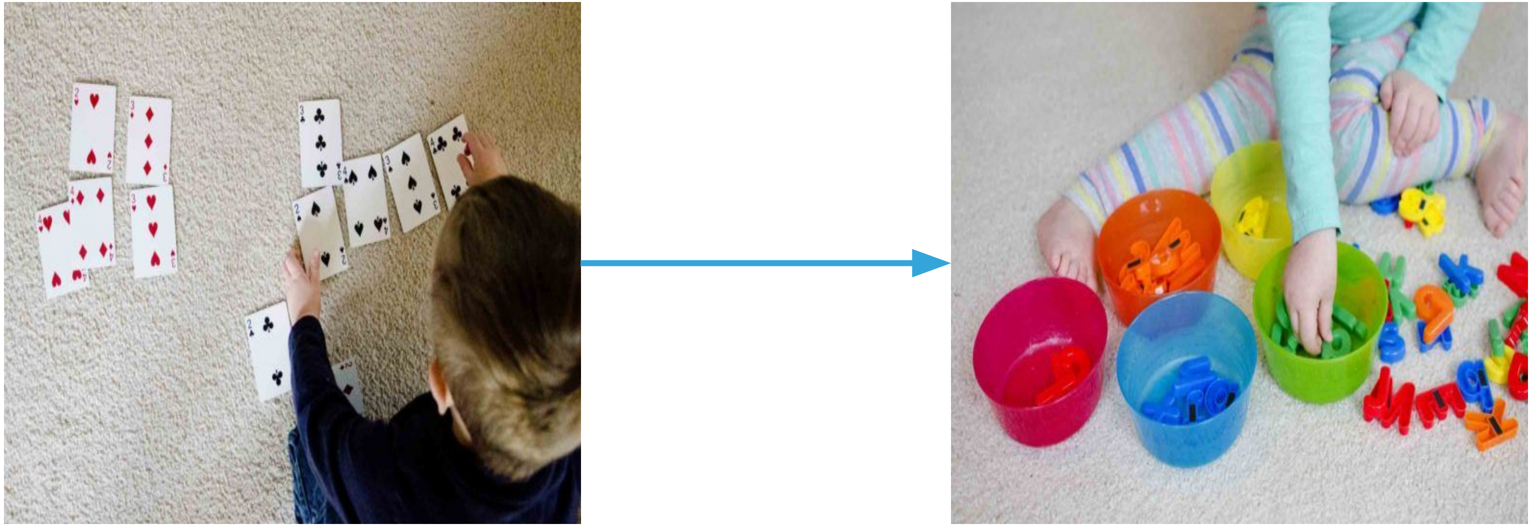

Learning Algorithms with Memory Augmented Networks

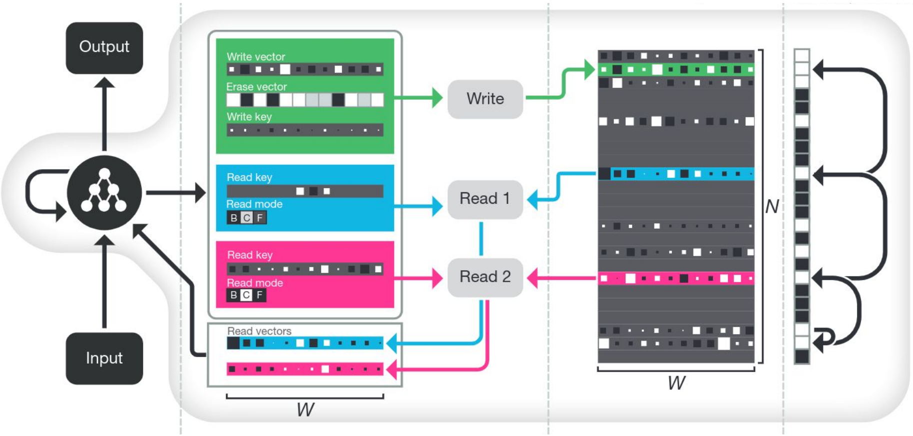

Graves, A., et al. Hybrid computing using a neural network with dynamic external memory. *Nature, 2016.*

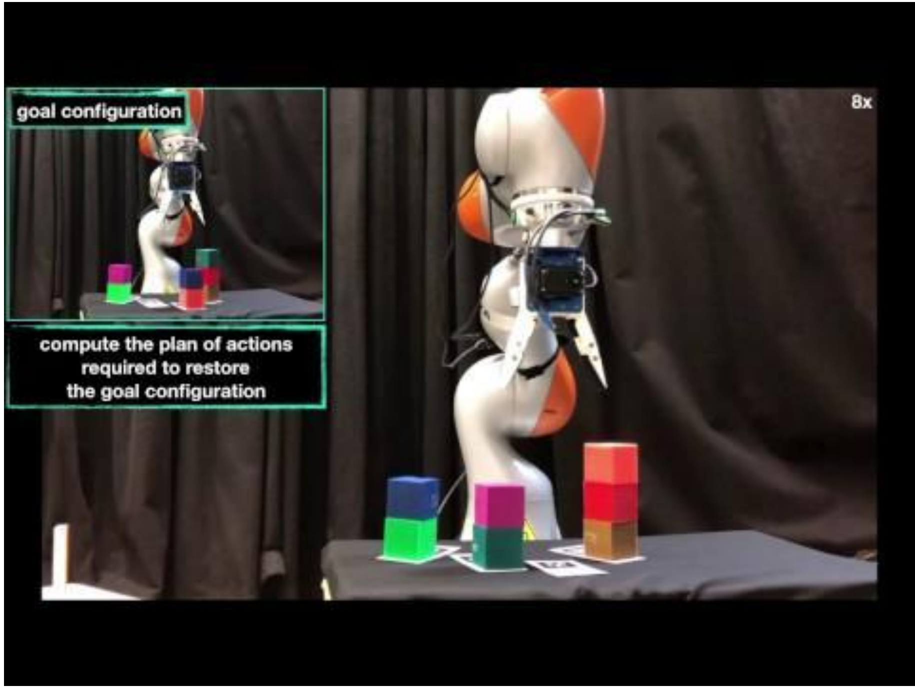

Tanneberg, Daniel; **Rueckert, Elmar**; Peters, Jan**.** Learning Algorithmic Solutions to Symbolic Planning Tasks with a Neural Computer Architecture. Preprint available at https://arxiv.org/abs/1911.00926, *under review.*

Summary: **My goal** is to build *Intelligent Learning Systems!*

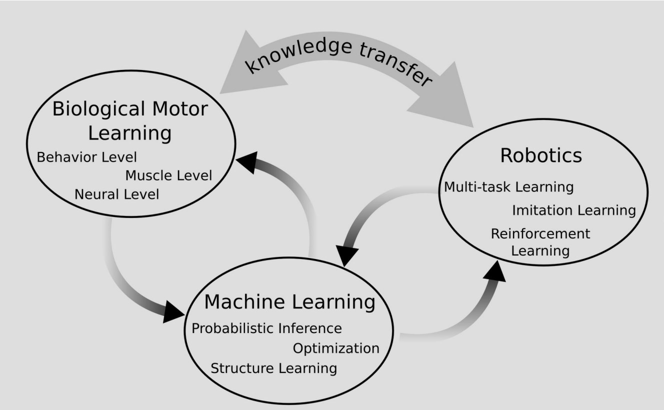

Feedback control

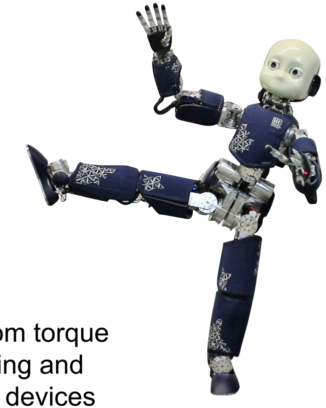

custom torque sensing and input devices

Motion/Skill Representations

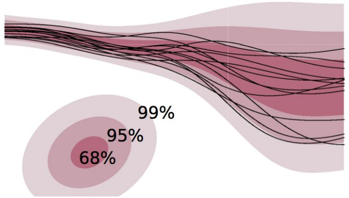

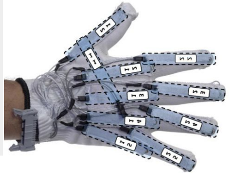

human / robot interaction with smart sensing

Motion/Skill Libraries for online adaptation

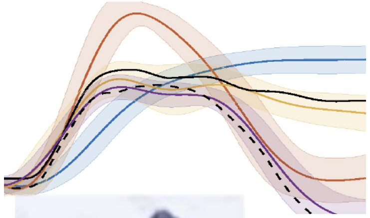

time of flight camera integration

Cognitive Planning / Reasoning

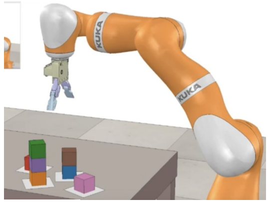

low cost neural planning systems

Current research & industrial projects

2019-2022 LEGO Robotic Project,
Robert-Bosch Stiftung

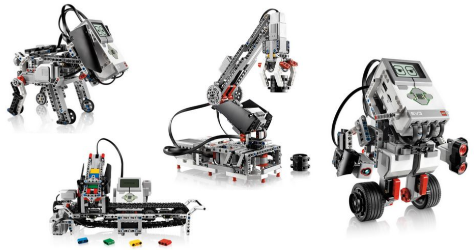

H2020 2016-2020
Intrinsic Motivation
Learning

GOAL Robots

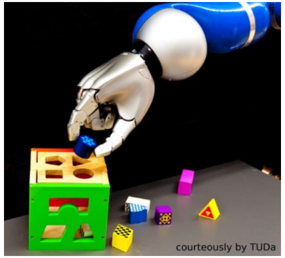

AiF Projekt: Entwicklung von
Lern- und Kalibrierungs
methoden für Kraftzellen
(Contributor)

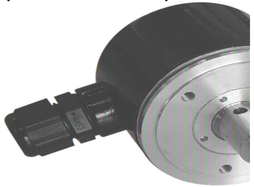

DFG 2020-2023 Active Transfer
Learning using Neural Networks.

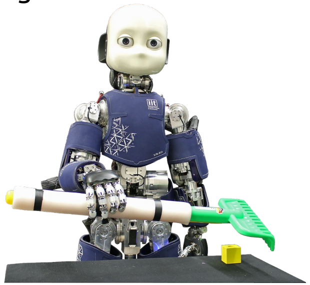

Industrial Cooperation
LUPA Electroincs GmbH

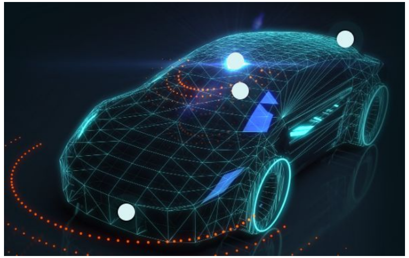

Research proposals under review or to be submitted this month

2 PhDs 4Y. 750k€

H2020 - Trusted Role Allocator AI for Collaborative Manufacturing

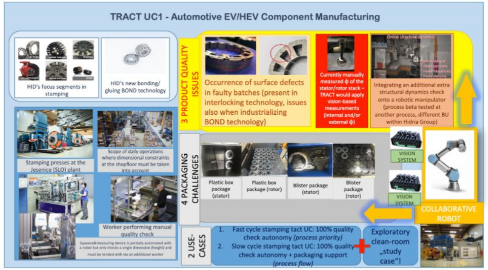

16.01.2020

2 PhDs 2Y, 300k€

BMBF - Probabilistic Model Predictive Control (ProMPC)

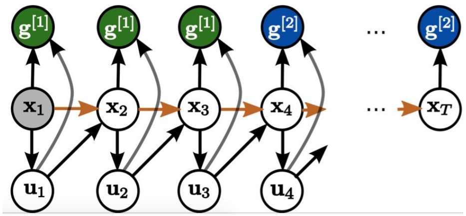

Seit 12/2019 auf der Warteliste

Research proposals under review or to be submitted

3 PhDs 6Y 1.2M€

Emmy Noether - Transfer Learning in Model-Based Neural RL

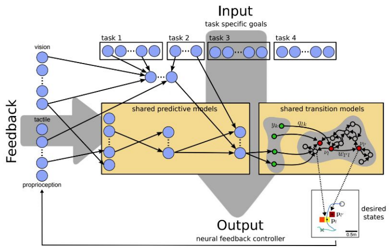

19.02.2020

1 PhD 3Y 150k€

H2020 - Collective Sensorimotor Intelligence (CoSMi)

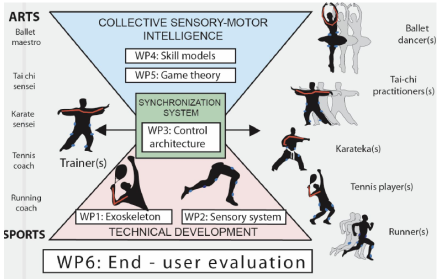

2 PhDs 4Y. 700k€ 22.04.2020

BMWi - Multiple Entscheidungsstrategien für automatisierte Fahrzeuge

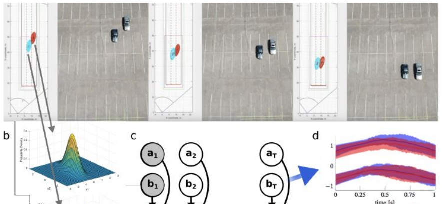

29.02.2020

Teaching

Humanoid Robotics - 5./6. Semester B.Sc. WF

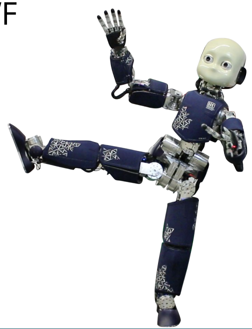

Probabilistic Machine Learning - 1./2. Semester M.Sc. PF

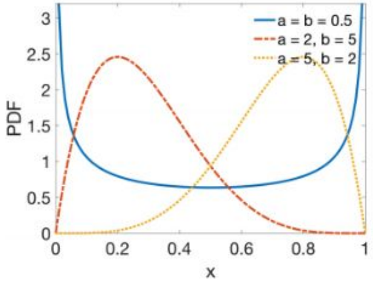

[Link to the Script]()

Reinforcement Learning - 1./2. Semester M.Sc. PF

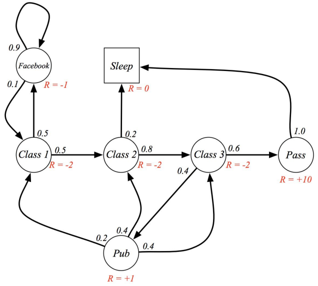

I would like to thank my team

University of Luebeck

Mr. Nils Rottmann, M.Sc. investigates in his doctoral thesis the learning of optimal control and planning strategies in mobile and humanoid robots. He started his thesis in March 2018.

Mr. Honghu Xue, M.Sc. investigates in his doctoral thesis the probabilistic and neural control mechanisms for compliant robots. He started his thesis in March 2019.

Open PhD Position in Machine Learning, Reinforcement Learning or Neural Networks working on the DFG project TRAIN.

Technical University Darmstadt

Mr. Daniel Tanneberg, M.Sc. investigates in his doctoral thesis machine learning algorithms for human-like learning and tactile manipulation. He started his thesis in October 2015 and is co-supervised with Prof. Jan Peters at the Technische Universität Darmstadt.

Ms. Svenja Stark, M.Sc. investigates in her doctoral thesis intrinsic motivation learning strategies for motor skills acquisition in robots. She started her thesis in November 2016 and is co-supervised with Prof. Jan Peters at the Technische Universität Darmstadt.

- **Darmstadt**: Daniel Tanneberg, Svenja Stark, Gerhard Neumann, Alexandros Paraschos, Roberto Calandra, Jan Peters, Rudolf Lioutikov, Marc Deisenroth, Serena Ivaldi, Tucker Hermans, Philipp Beckerle, Valerio Modugno, Jan Mundo, David Sharma, Jan Kohlschuetter, Svenja Stark, Michael Schmidt, Max Mindt

- **Tübingen**: Moritz Grosse-Wentrup, Martin Giese

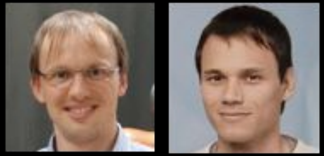

- **Ljubijana**: Jan Babic, Jernej Camernik

- **Birmingham**: Michael Mistry, Morteza Azad

- **Graz**: Wolfgang Maass, Robert Legenstein, David Kappel, Dejan Pecevski

Birmingham: Jeremy Wyatt, Michael Mistry, Morteza Azad, **Rome**: Andrea d‘Avella and Yuri Ivanenko,
Stuttgart: Marc Toussaint, **Bielefeld**: Thomas Schack, Jochen Steil, **Genua**: Francesco Nori, Lorenzo Natale

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 5209
E-Mail: rueckert@rob.uni-luebeck.de

Disclaimer:

Probabilistic Robot Learning

Reinforcement Learning | Movement Primitives | Humanoid Robotics

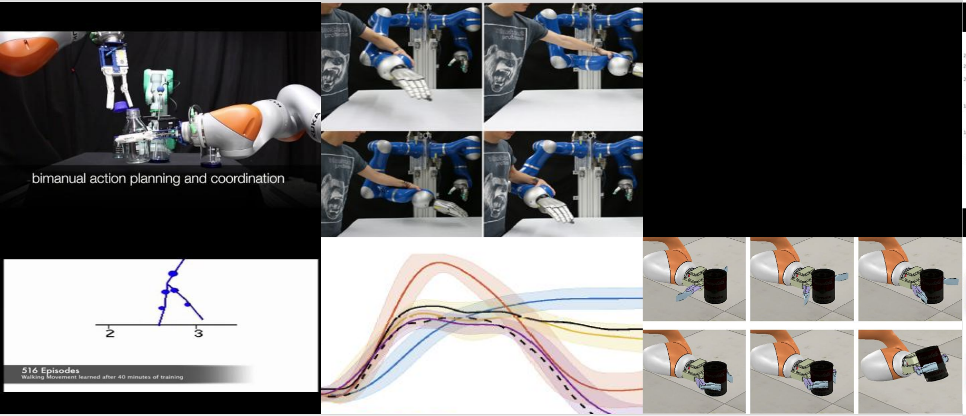

Probabilistic Computational Neuroscience

Spiking Neural Networks | Intrinsic Motivations | Human-Machine Int.

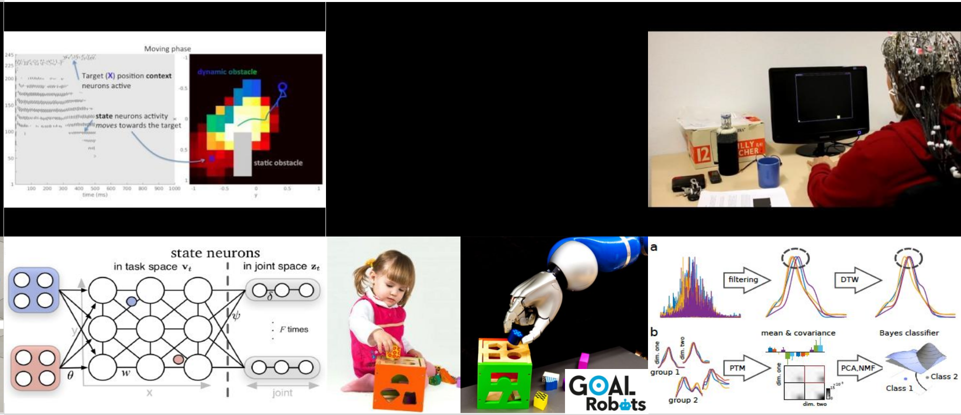

Stochastic Machine and Deep Learning

Tactile Learning | Model Learning | Probabilistic Inference

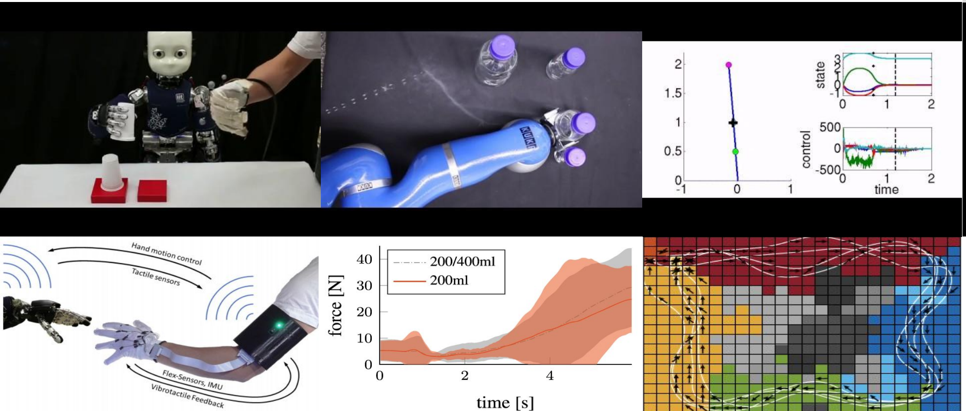

Medical Robotics and Human Motor Control

Robotics Health Care | Postural Control | Muscle Synergies