Open challenges in the application of DL to distributed mobile sensing

Fabio Bonsignorio[1,2,3,4,5,6,7,...]

WG1 Embodied Intelligence in Natural and Artificial Agents RoboCom++ Coordinator[1]
G2Net WG2 Machine Learning for low-frequency seismic measurement Co-leader, Robotics Task Leader[2]
Coordinator The Shanghai Lectures[3]
Coordinator HumaBiMan[4]
Coordinator HumaBelief[5]
SPARC TG Benchmarking and Competitions Coordinator[6]
IEEE RAS TC-PEBRAS Co-chair[7]
Founding and Past Member SPARC Board of Directors[8]
Heron Robots[9]
...

Heron Robots

www.heronrobots.com

Older and newer attempts

Juanelo Torriano alias Gianello della Torre, (XVI century) a craftsman from Cremona, built for Emperor Charles V a mechanical young lady who was able to walk and play music by picking the strings of a real lute.

Hiroshi Ishiguro, early XXI century

Director of the Intelligent Robotics Laboratory, part of the Department of Adaptive Machine Systems at Osaka University, Japan

the success stories

DARPA (American Defense Advanced Research Projects Agency) challenges have demonstrated how current robots are becoming **more accurate**, **fast** and **dexterous in structured and unstructured environments**.

Not everything worked as expected!

the current approach shows some limitations

On the other hand the debriefing of DARPA DRC shows clearly that humanoid robots are **still far from the required level of capabilities in fact many metrics**, **such as time-to-completion, are highly application or task specific.**

According to H.Yanco a minimum of 9 people were needed to teleoperate latest DRC's robots!!!

The marvellous progress of Robotics and AI…'Look Ma, No Hands' syndrome?

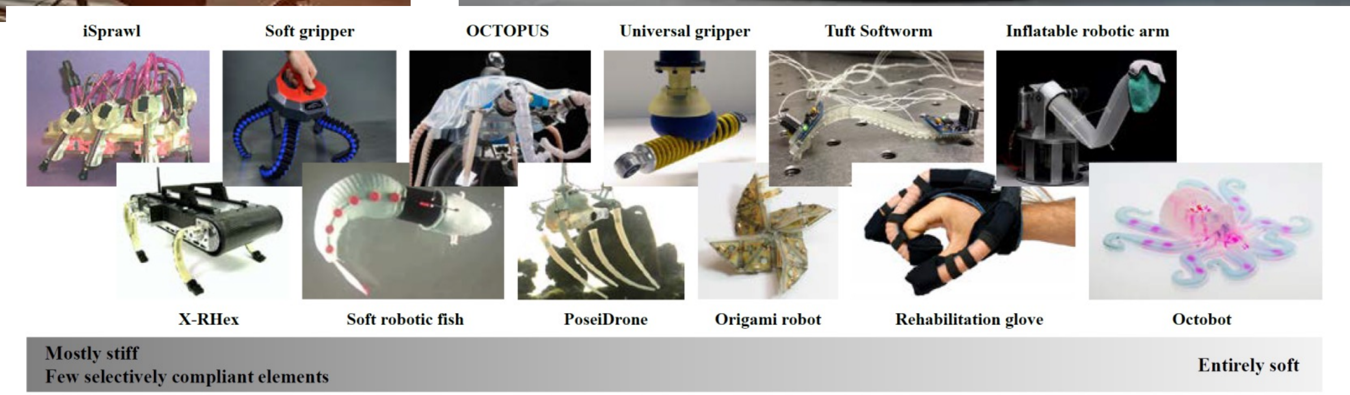

Is It Alive?

Big Questions lie in front of us!

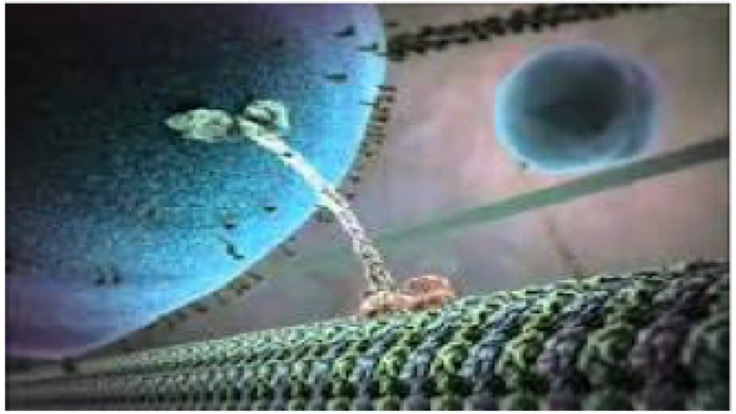

Two views of intelligence

classical:
cognition as computation

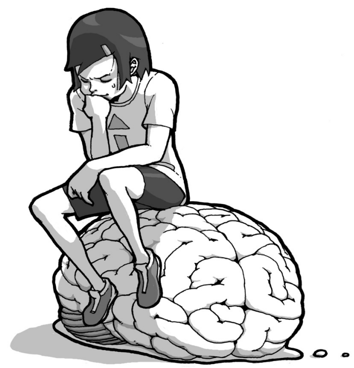

embodiment:
cognition emergent from sensory-motor and interaction processes

PARADIGM CLASHES

Soft Robotics: a working definition

Variable impedance actuators and stiffness control

* Actuators with variable impedance
* Compliance/impedance control
* Highly flexible (hyper-redundant or continuum) robots

Use of soft materials in robotics

* Robots made of soft materials that undergo high deformations in interaction
* Soft actuators and soft components
* Control partially embedded in the robot morphology and mechanical properties

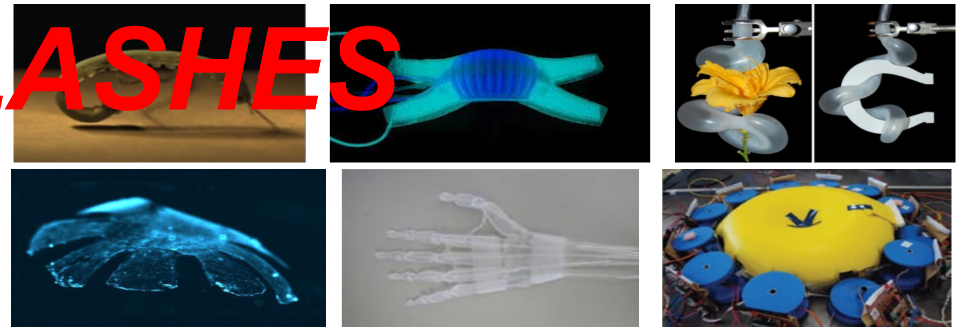

PARADIGM CLASHES

THE BIOROBOTICS INSTITUTE

Scuola Superiore Sant'Anna

IEEE Robotics and Automation Magazine, Special Issue on Soft Robotics, 2008
A. Albu-Schaffer et al. (Ed.s)

Kim S., Laschi C., and Trimmer B. (2013) Soft robotics: a bioinspired evolution in robotics, *Trends in Biotechnology*, April 2013.
Laschi C. and Cianchetti M. (2014) "Soft Robotics: new perspectives for robot bodyware and control" *Frontiers in Bioengineering and Biotechnology*, 2(3)

Why it matters

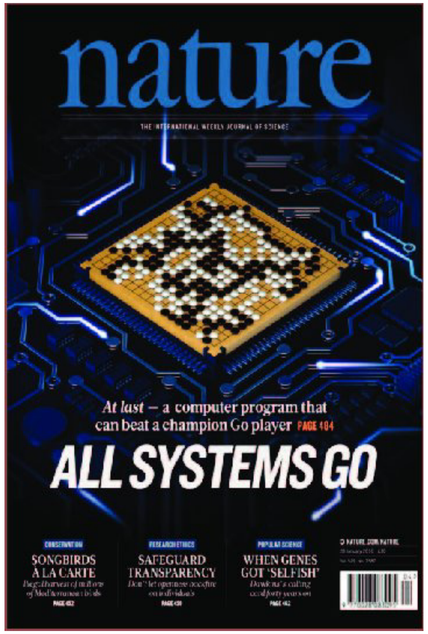

Published: 27 January 2016

Mastering the game of Go with deep neural networks and tree search

David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel & Demis Hassabis

Nature **529**, 484–489(2016) | Cite this article

Abstract

The game of Go has long been viewed as the most challenging of classic games for artificial intelligence owing to its enormous search space and the difficulty of evaluating board positions and moves. Here we introduce a new approach to computer Go that uses 'value networks' to evaluate board positions and 'policy networks' to select moves. These deep

Editorial Summary

AlphaGo computer beats Go champion

The victory in 1997 of the chess-playing computer Deep Blue in a six-game series against the then world champion Gary Kasparov was seen as a significant milestone in the development of artificial intelligence. An even greater

show all

Associated Content

Collection

The multidisciplinary nature of machine intelligence

Why it matters

nature

View all Nature Research journals | Search | Logi

Explore our content | **Journal information** | **Subscribe**

nature > news > article

NEWS · 30 NOVEMBER 2020

‘It will change everything’: DeepMind’s AI makes gigantic leap in solving protein structures

Google’s deep-learning program for determining the 3D shapes of proteins stands to transform biology, say scientists.

Ewen Callaway

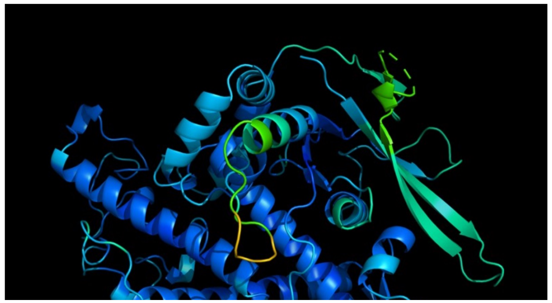

RELATED ARTICLES

AI protein-folding algorithms solve structures faster than ever

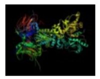

The revolution will not be crystallized: a new method sweeps through structural biology

The computational protein designers

Chess: New York, 1997

1 win   3 draws   2 wins

Go: Hong Kong, 2017

Google's AlphaGo Defeats Chinese Go Master in Win for A.I.

Ke Jie, the world's top Go player, reacting during his match on Tuesday against AlphaGo, artificial intelligence software developed by a Google affiliate. China Stringer Network, via Reuters

By Paul Mozur

May 23, 2017

阅读简体中文版

HONG KONG — It isn't looking good for humanity.

A comparison

Chess and GO are ‘perfect information games’

They always have an optimal value function which determines, under perfect game assumptions by all players, the outcome of the game from any initial state s.

The recursion tree in such games will include roughly **b^d moves**

- **<u>Chess</u>: b ≈ 35**, **d ≈ 80**
- **<u>Go</u>: b ≈ 250**, **d ≈ 150**

Interestingly the developers of AlphaGo have implemented an exhaustive testing and evaluation schema to compare and refine different gaming policies by mixing Montecarlo Simulations, Machine learning and guided sampling techniques.

Remarks

- In **embodied AI (aka intelligent robotics!)** deterministic approaches are practically impossible to implement **-> No 'perfect information games'**

- We are very much likely still far from what we have to cope with for a robot operating in the real world, but it can be seen as a better approximation than Chess….and other proposed before.

Silver, D. et al. , Mastering the game of Go with deep neural networks and tree search, Nature 529, 484–489, 2016

ET

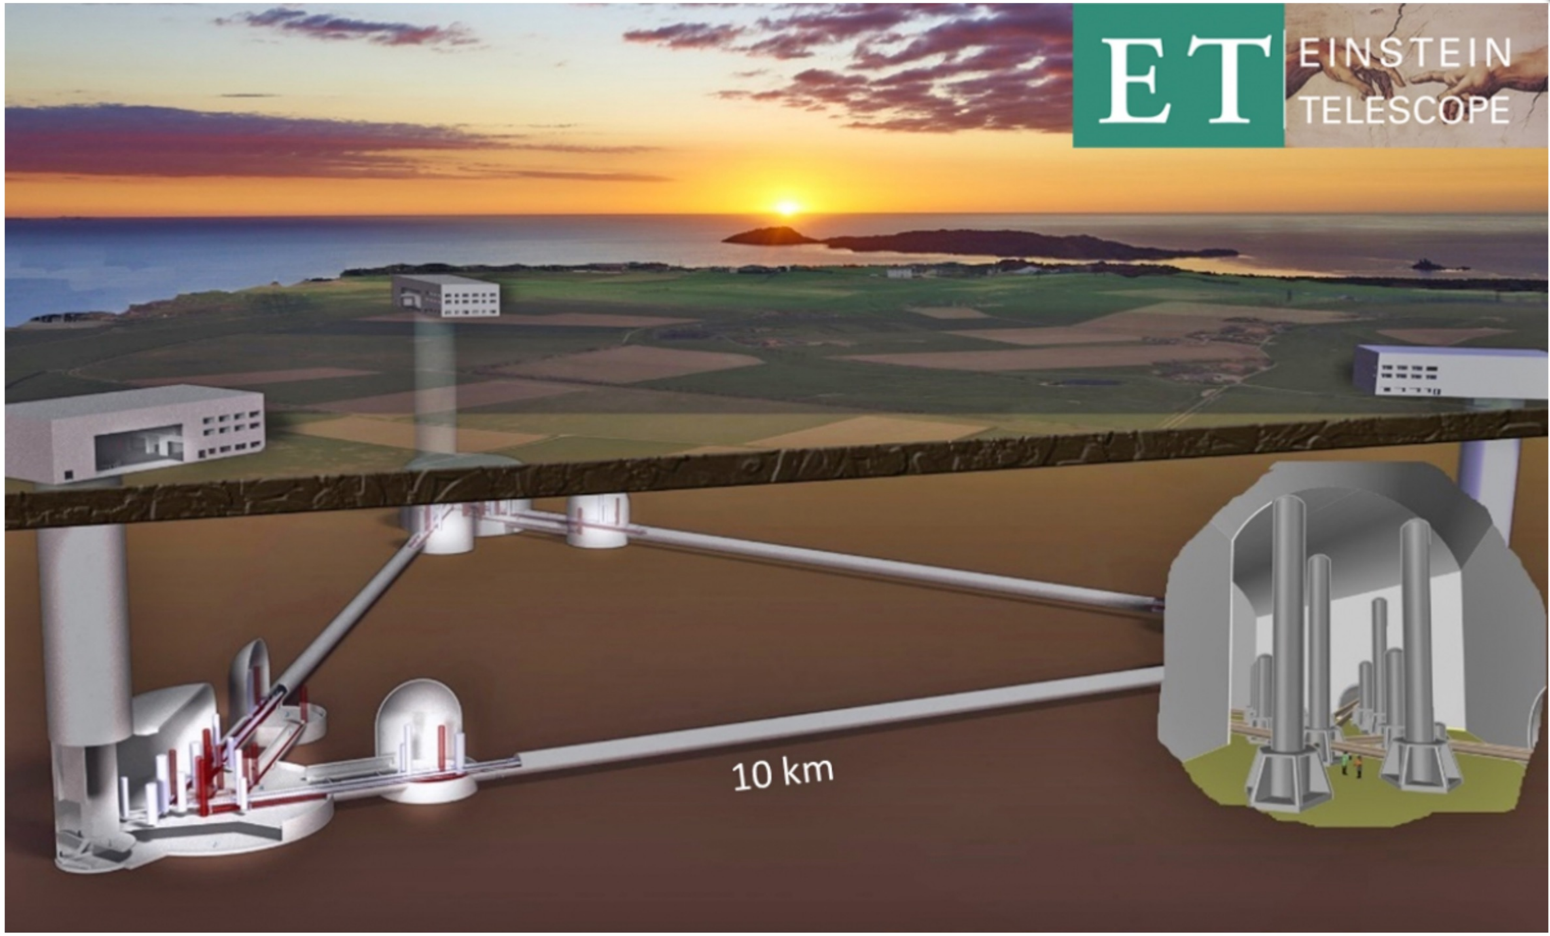

con una sensibilità di almeno un ordine di grandezza migliore rispetto agli attuali rivelatori. ET sarà installato in una nuova tics

Seismic NN: elastic models

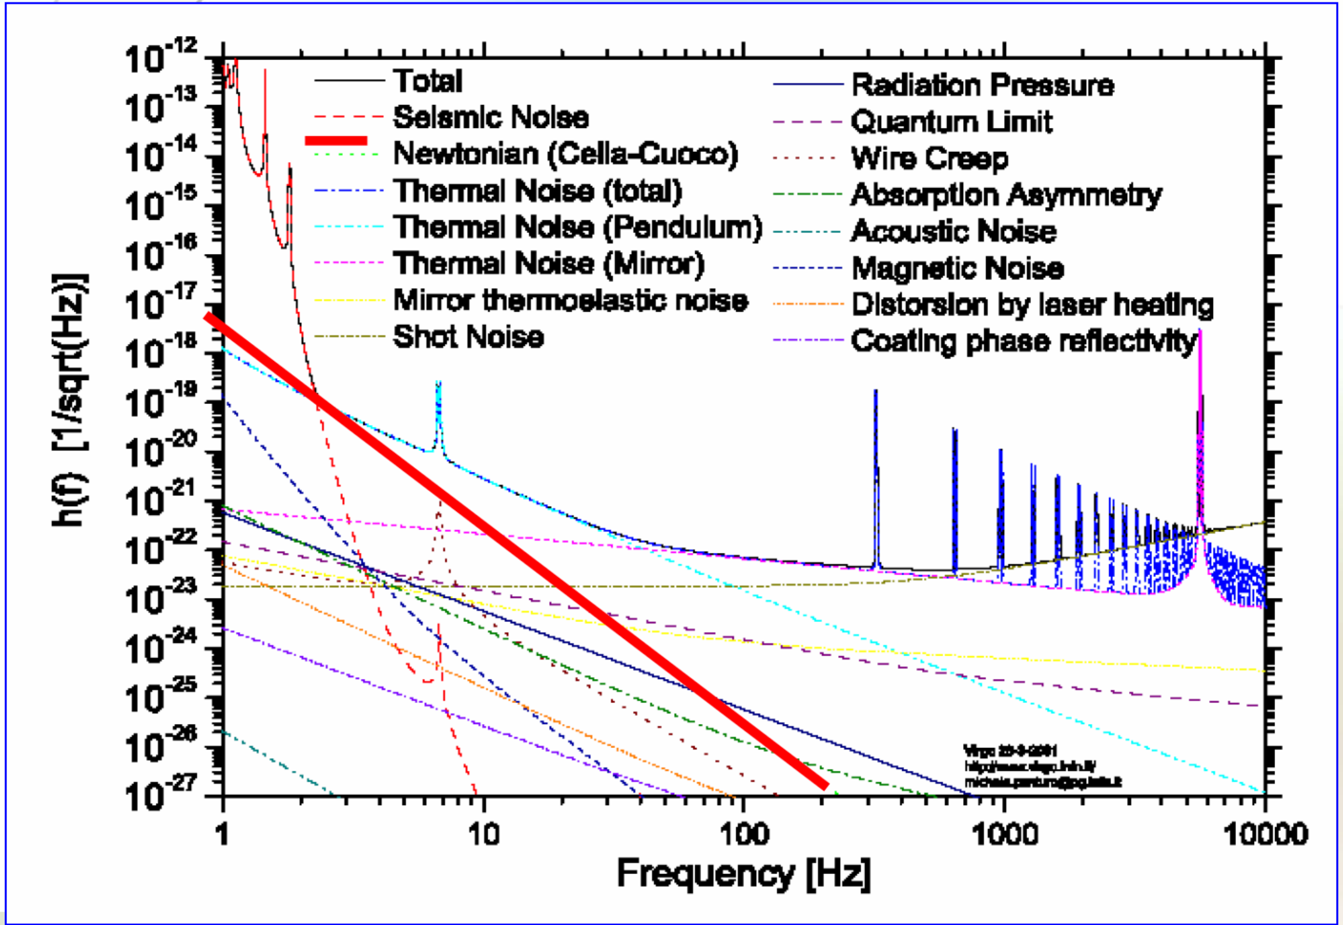

Giancarlo Cella
INFN sez. Pisa

3rd ILIAS Annual meeting
Gran Sasso INFN National Lab
February 28-March 3, 2006

Newtonian Noise (a naïve view))

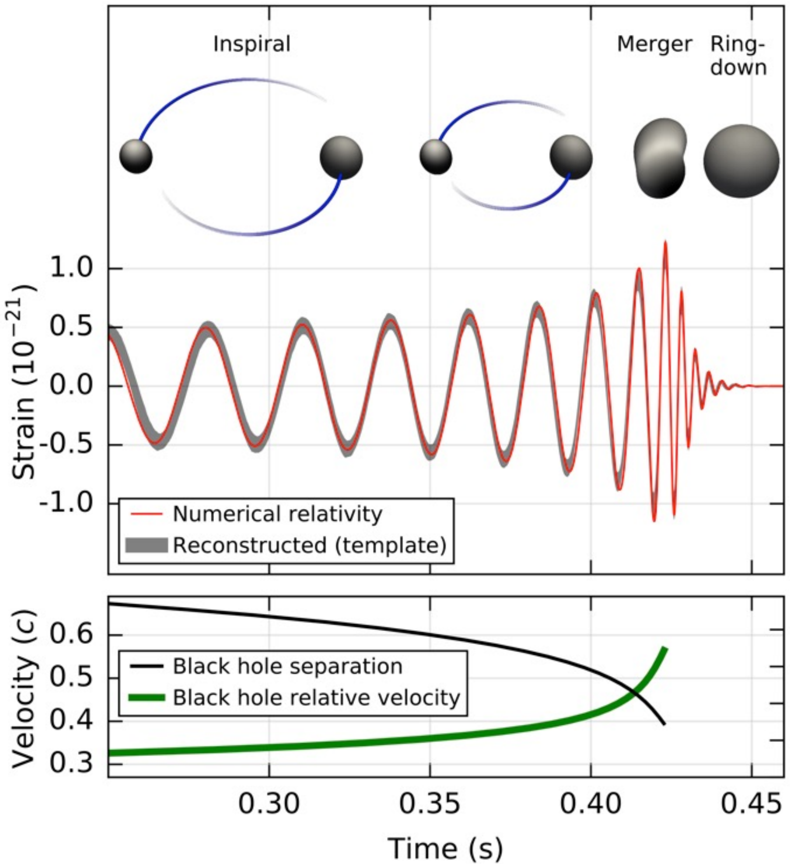

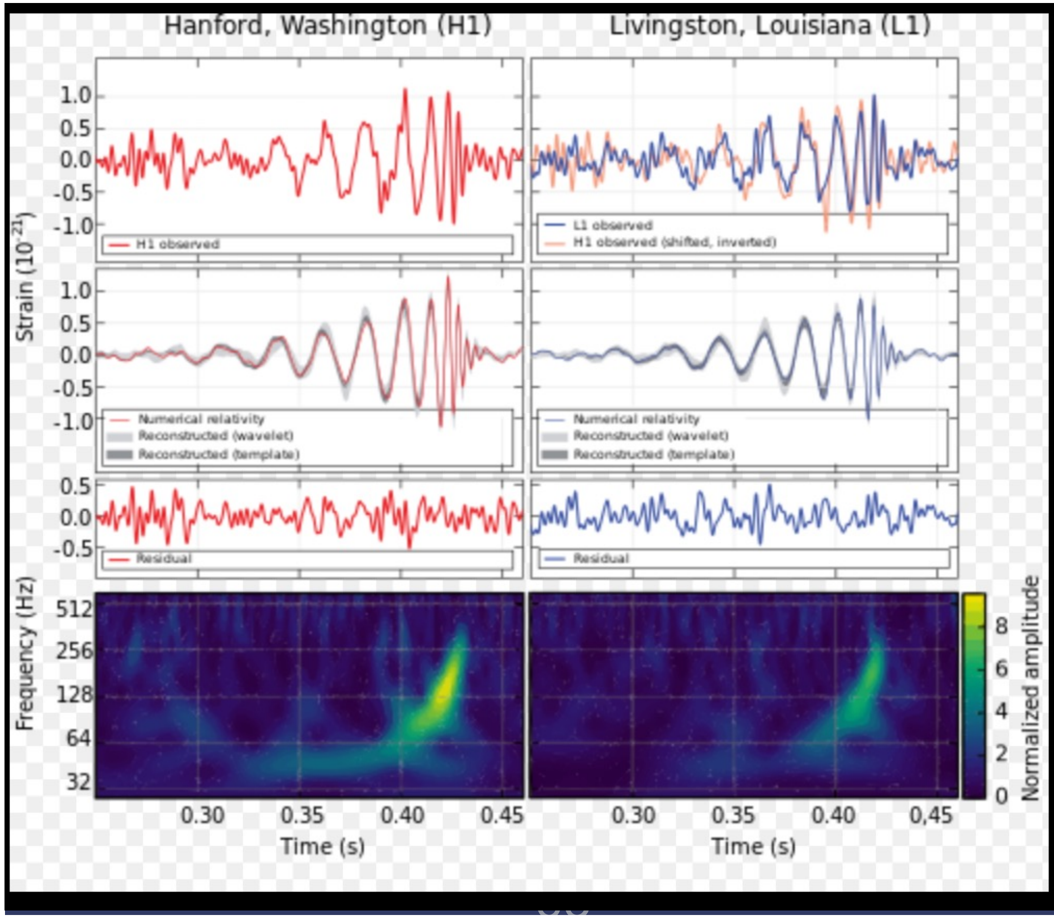

Newtonian Noise (a naïve view)

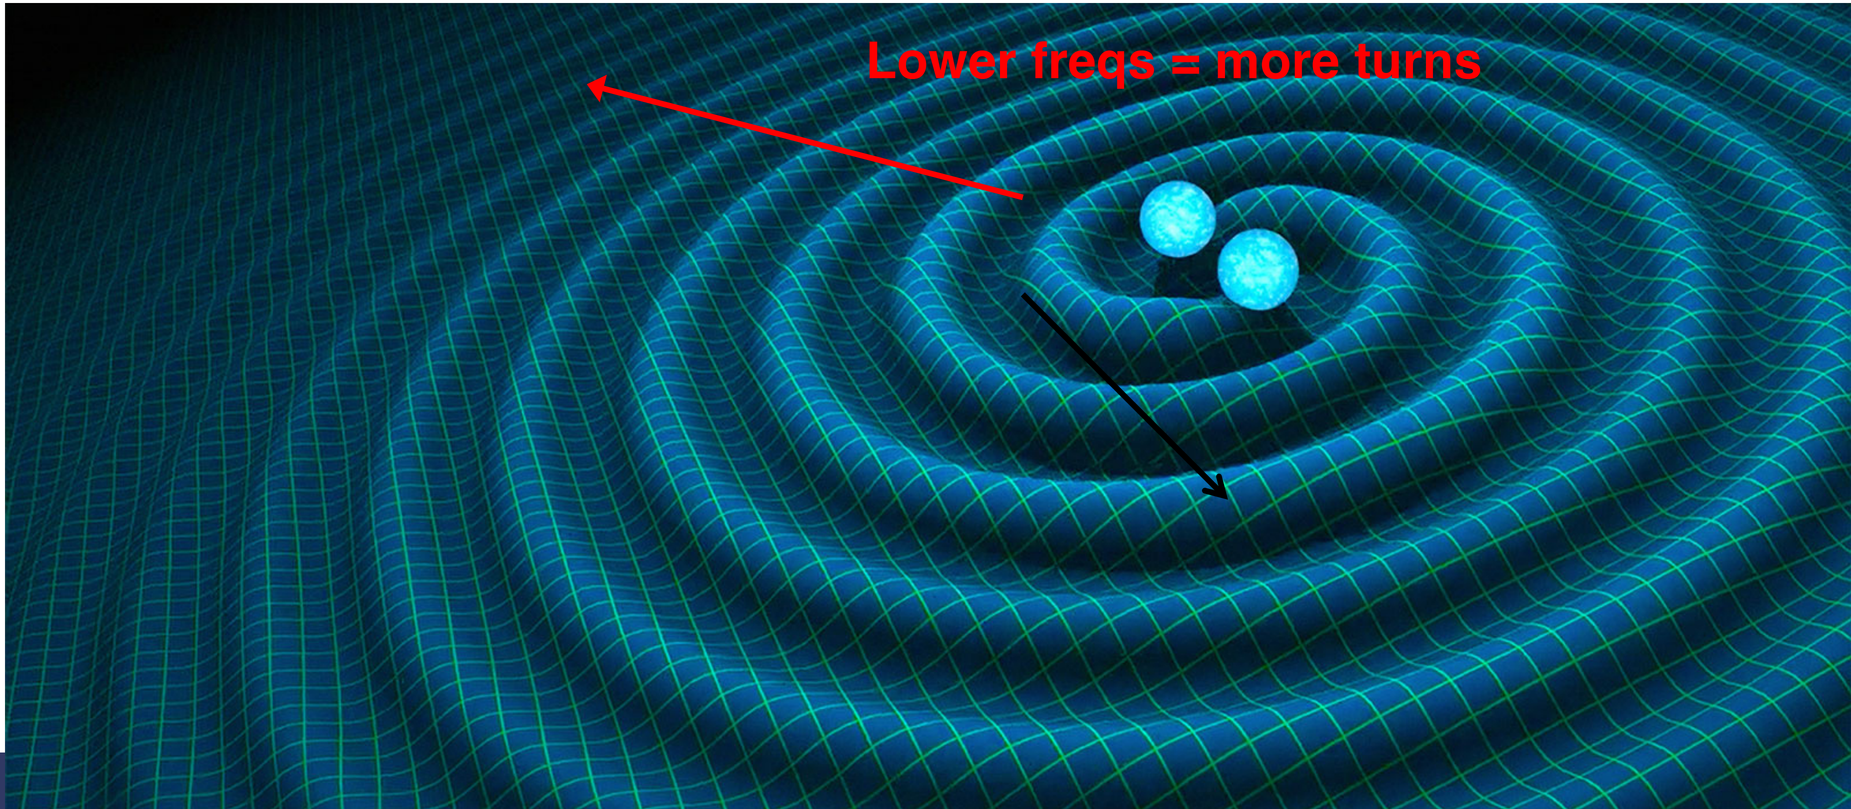

Netwonian Noise

A naïve view

Main Issue: Rayleigh waves (and lacking knowledge of underground mass distribution)

Problem: model underground and surface mass distribution and land motion
(same issue with the atmosphere) to characterize and predict Rayleigh waves

Other sources of noise:

‘Environmental’

i.e.

- Acoustic
- EM
- Others...

Multisensory Data Fusion in Robotics

Multisensor data fusion is the process of combining observations from a number of different sensors to provide a robust and complete description of an *environment* or process of interest.

Data fusion finds wide application in many areas of robotics such as object recognition, *environment mapping*, and localisation.

From: H. Durrant-Whyte, T. C. Henderson,
Multisensor Data Fusion,
Part C, Chapter 25, in B.Siciliano, O. Khatib (eds.) Springer Handbook of Robotics, 2008

Multisensory Data Fusion in Robotics

Principles

It's essentially an application of Bayes' rule: $$P(\boldsymbol{x} \mid z) = \frac{P(z \mid \boldsymbol{x}) P(\boldsymbol{x})}{P(z)}.$$

assuming conditional independence:

$$P(z_1, \cdots, z_n \mid \boldsymbol{x}) = P(z_1 \mid \boldsymbol{x}) \cdots P(z_n \mid \boldsymbol{x}) = \prod_{i=1}^{n} P(z_i \mid \boldsymbol{x}).$$

We get the multisensory expression:

$$\boxed{P(\boldsymbol{x} \mid \boldsymbol{Z}^n) = C P(\boldsymbol{x}) \prod_{i=1}^{n} P(z_i \mid \boldsymbol{x}),}$$

and its recursive form:

$$P(\boldsymbol{x} \mid \boldsymbol{Z}^k) = \frac{P(z_k \mid \boldsymbol{x}) P(\boldsymbol{x} \mid \boldsymbol{Z}^{k-1})}{P(z_k \mid \boldsymbol{Z}^{k-1})}.$$

Multisensory Data Fusion in Robotics

Methods

- Bayes' Rule
- *Probabilistic Grids*
- The Kalman Filter (plus Extended Kalman Filters, Information Filters, etc.)
- Sequential Monte Carlo Methods
- Alternatives to Probability

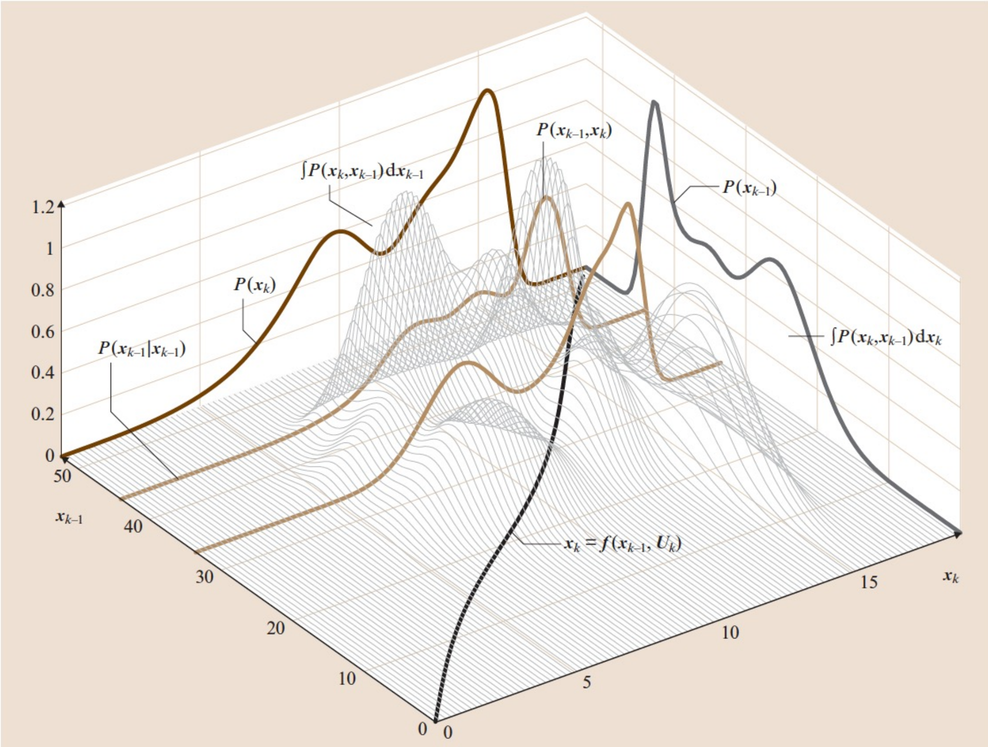

Fig. 25.1 Time update step for the full Bayes filter. At a time $k-1$, knowledge of the state x_{k-1} is summarised in a probability distribution $P(x_{k-1})$. A vehicle model, in the form of a conditional probability density $P(x_k \mid x_{k-1})$, then describes the stochastic transition of the vehicle from a state x_{k-1} at a time $k-1$ to a state x_k at a time k. Functionally, this state transition may be related to an underlying kinematic state model in the form $x_k = f(x_{k-1}, u_k)$. The figure shows two typical conditional probability distributions $P(x_k \mid x_{k-1})$ on the state x_k given fixed values of x_{k-1}. The product of this conditional distribution with the marginal distribution $P(x_{k-1})$, describing the prior likelihood of values of x_k, gives the the joint distribution $P(x_k, x_{k-1})$ shown as the surface in the figure. The total marginal density $P(x_k)$ describes knowledge of x_k after state transition has occurred. The marginal density $P(x_k)$ is obtained by integrating (projecting) the joint distribution $P(x_k, x_{k-1})$ over all x_{k-1}. Equivalently, using the total probability theorem, the marginal density can be obtained by integrating (summing) all conditional densities $P(x_k \mid x_{k-1})$ weighted by the prior probability $P(x_{k-1})$ of each x_{k-1}. The process can equally be run in reverse (a retroverse motion model) to obtain $P(x_{k-1})$ from $P(x_k)$ given a model $P(x_{k-1} \mid x_k)$

H. Durrant-Whyte, T. C. Henderson, Multisensor Data Fusion, Part C, Chapter 25, in B.Siciliano, O. Khatib (eds.) Springer Handbook of Robotics, 2008

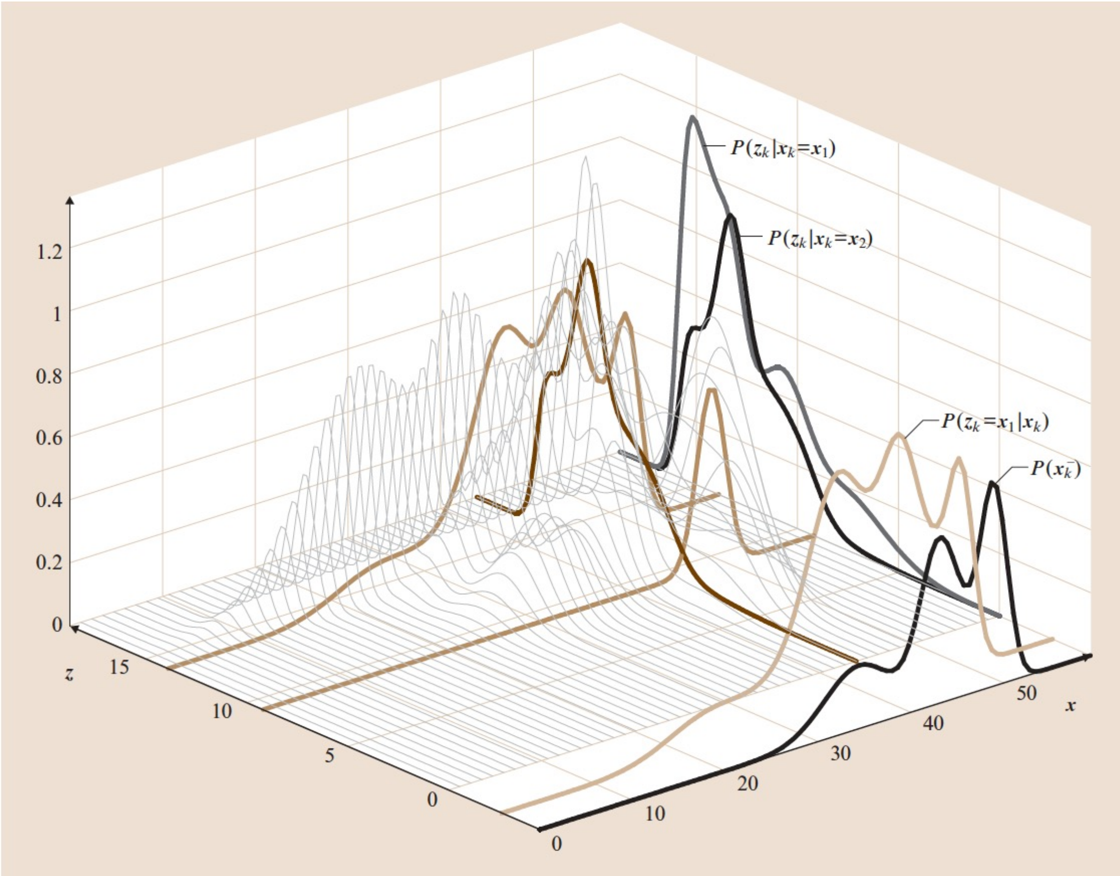

Fig. 25.2 Observation update for the full Bayes filter. Prior to observation, an observation model in the form of the conditional density $P(z_k \mid \boldsymbol{x}_k)$ is established. For a fixed value of $\boldsymbol{x}_k$, equal to $\boldsymbol{x}_1$ or $\boldsymbol{x}_2$ for example, a density function $P(z_k \mid \boldsymbol{x}_k = \boldsymbol{x}_1)$ or $P(z_k \mid \boldsymbol{x}_k = \boldsymbol{x}_2)$ is defined describing the likelihood of making the observation z_k. Together the density $P(z_k \mid \boldsymbol{x}_k)$ is then a function of both z_k and $\boldsymbol{x}_k$. This conditional density then defines the observation model. Now, in operation, a specific observation $z_k = \boldsymbol{x}_1$ is made and the resulting distribution $P(z_k = \boldsymbol{x}_1 \mid \boldsymbol{x}_k)$ defines a density function (now termed the likelihood function) on $\boldsymbol{x}_k$. This density is then multiplied by the prior density $P(\boldsymbol{x}_k^-)$ and normalised to obtain the posterior distribution $P(\boldsymbol{x}_k \mid z_k)$ describing knowledge in the state after observation

H. Durrant-Whyte, T. C. Henderson, Multisensor Data Fusion, Part C, Chapter 25, in B.Siciliano, O. Khatib (eds.) Springer Handbook of Robotics, 2008

Multisensory Data Fusion in Robotics

Example: ANSER II: Decentralised Data Fusion

Decentralised data fusion (DDF) methods were initially motivated by the insight that the information or canonical form of the conventional Kalman filter data fusion algorithm could be implemented by simply adding information contributions from observations. As these (vector and matrix) additions are commutative, the update or data fusion process can be optimally distributed amongst a network of sensors-

The sensor is modelled directly in the form of a likelihood function. Once instantiated with an observation, the likelihood function is input to a local fusion loop which implements a local form of the Bayesian time and observation update. Network nodes accumulate probabilistic information from observation or communication and exchange mutual information (information gain) with other nodes in the network. This mutual information is transmitted to and assimilated by other nodes in the network in an ad-hoc manner. The result is that all nodes in the network obtain a single integrated posterior probability based all node observations.

The ANSER II system consists of a pair of autonomous air vehicles equipped with infra-red and visual sensors, a pair of unmanned ground vehicles equipped with visual and radar sensors, and additional information provided by geometric and hyper-spectral data bases, along with information input by human operatives. The likelihood functions for singlesensor features are obtained **through a semi-supervised machine learning method.** The resulting probabilities are modeled in the form of a mixture of Gaussians. Each platform then maintains a bank of decentralised, non-Gaussian Bayesian filters for the observed features, and transmits this information to all other platforms. The net result is that each platform maintains a complete map of all features observed by all nodes in the network. Multiple observations of the same feature, possibly by different platforms, results in an increasingly accurate estimate of the feature location for all nodes.

The ANSER II system demonstrates a number of general principles in Bayesian data fusion methods.

Specifically the need to appropriately model sensors through the likelihood function, and the possibility of building very different data fusion architectures from the essential Bayesian form.

Multisensory Data Fusion in Robotics
Example: ANSER II: Decentralised Data Fusion

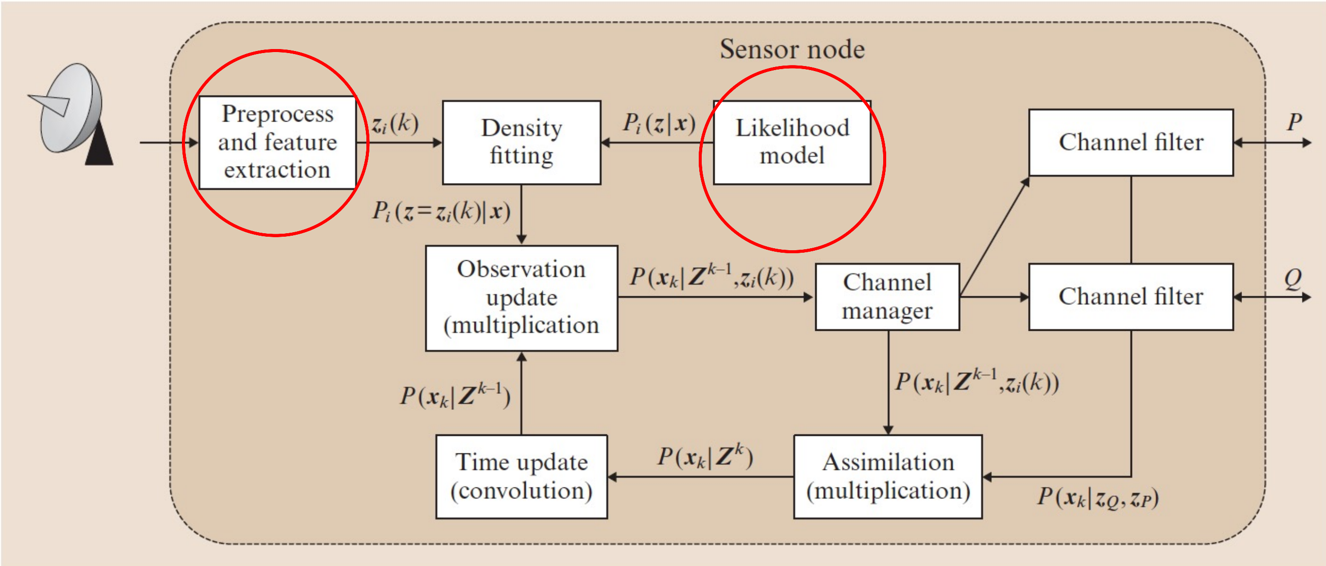

H. Durrant-Whyte, T. C. Henderson, Multisensor Data Fusion, Part C, Chapter 25, in B.Siciliano, O. Khatib (eds.) Springer Handboo of Robotics, 2008

Mathematical structure of a decentralised data fusion node

Multisensory Data Fusion in Robotics
Example: ANSER II: Decentralised Data Fusion

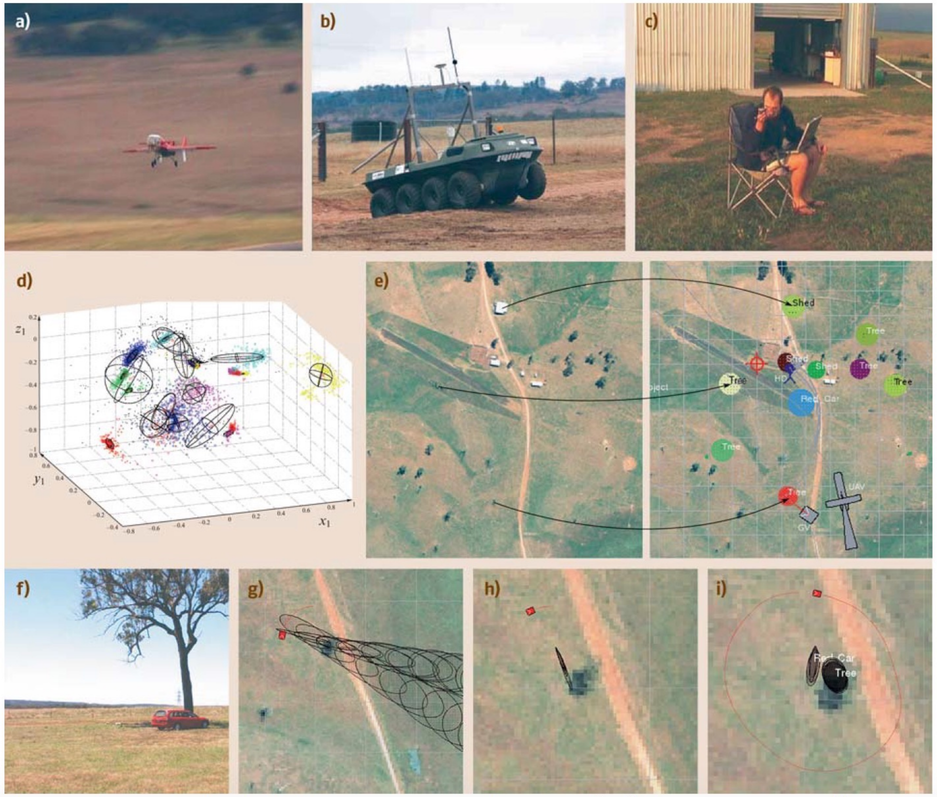

A synopsis of the ANSER II autonomous network and its operation.
(a–c) Main system components;
(a) air vehicle,
(b) ground vehicle,
(c) human operative.

(d–e) The perception process;
(d) top three dimensions of features discovered from ground-based visual sensor data along with the derived mixture model describing these feature properties
(e) sector of the overall picture obtained from fusing air vehicle (UAV), ground vehicle (GV) and human operator (HO) information. Each set of ellipses corresponds to a particular feature and the labels represent the identity state with highest probability.

(f–i) Sequential fusion process for two close landmarks: (f) a tree and a red car, (g) bearing-only visual observations of these landmarks are successively fused, (h) to determine location and identity (i).

Note the Gaussian mixture model for the bearing measurement likelihood

H. Durrant-Whyte, T. C. Henderson,
Multisensor Data Fusion,
Part C, Chapter 25, in B.Siciliano, O. Khatib (eds.) Springer Handbook of Robotics, 2008

NN mitigation

Preliminary ideas

- (adaptive) Modeling of the area (emi) sphere of r ≈ 10 m to 5 km by a network of robots equipped at least with onboard seismometers which change adaptively their positions

- Dynamic optimization of sensor positions (for example doubling those already installed?)

Decentralised Data Fusion like AnserII but with two main changes

Semi-supervised,
Reinforcement Learning?
Deep Reinforcement Learning

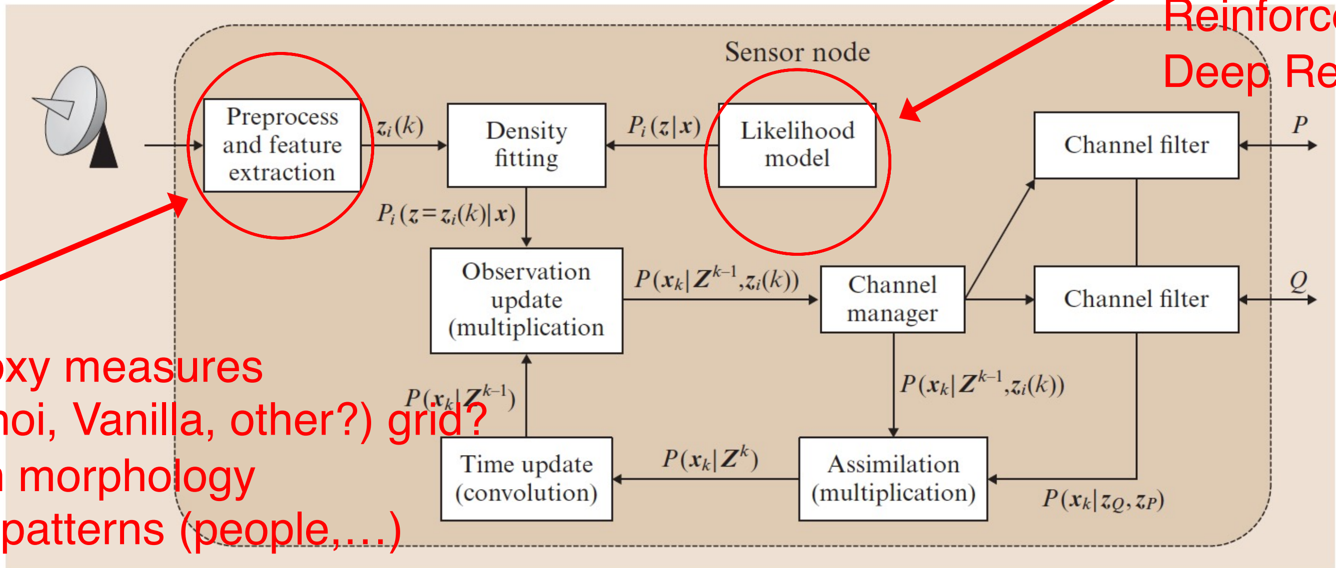

Density proxy measures
on a (Voronoi, Vanilla, other?) grid?
Also terrain morphology
and usage patterns (people,….)

Mathematical structure of a decentralized data fusion node

Multisensory Data Fusion in Robotics
Example: ANSER II: Decentralised Data Fusion

Quantum by INNOSEIS (a spin-out from the National Institute for Subatomic Physics in the Netherlands) is an ultra-light weight (< 1kg) wireless seismic sensor network that dramatically reduces deployment costs, while scaling up to 1 million nodes for onshore exploration. It has be designed for static Wireless, sensor networks. However, a daisy-chain small network is operating in Cascina already and no major issues prevent to mount them on mobile platforms.

T. Bulik and team's geophone

Infrasound microphone

Needed to characterize the infrasound field, Low cost, Sensitivity in the range of 1-30Hz, lots of uses: geophysics, volcanology etc, Potential industrial applications, Prototype ready – network to be installed in Virgo this year

A Multisensory Multiagent platform for GW detection and Geophysics applications: Theory

Multi-robot reconstruction of a spatial signal driven by the information gain

Multi-robot reconstruction of a spatial signal driven by the information gain, see: 2019 International Conference on Robotics and Automation (ICRA), Montreal, Canada, May 20-24, 2019

- **Signal modeling by GP regression**
- **Information Gain**
- **Multi-robot coordination and task allocation**

Important Remark
→This is already totally feasible←

Both mobile sensors and mobile multisensory fusion network can be developed and tested already in Ligo-Virgo-Kagra-Ligo India

They can be integrated with fixed sensors

They could improve noise characterization in ET

More expected in the future

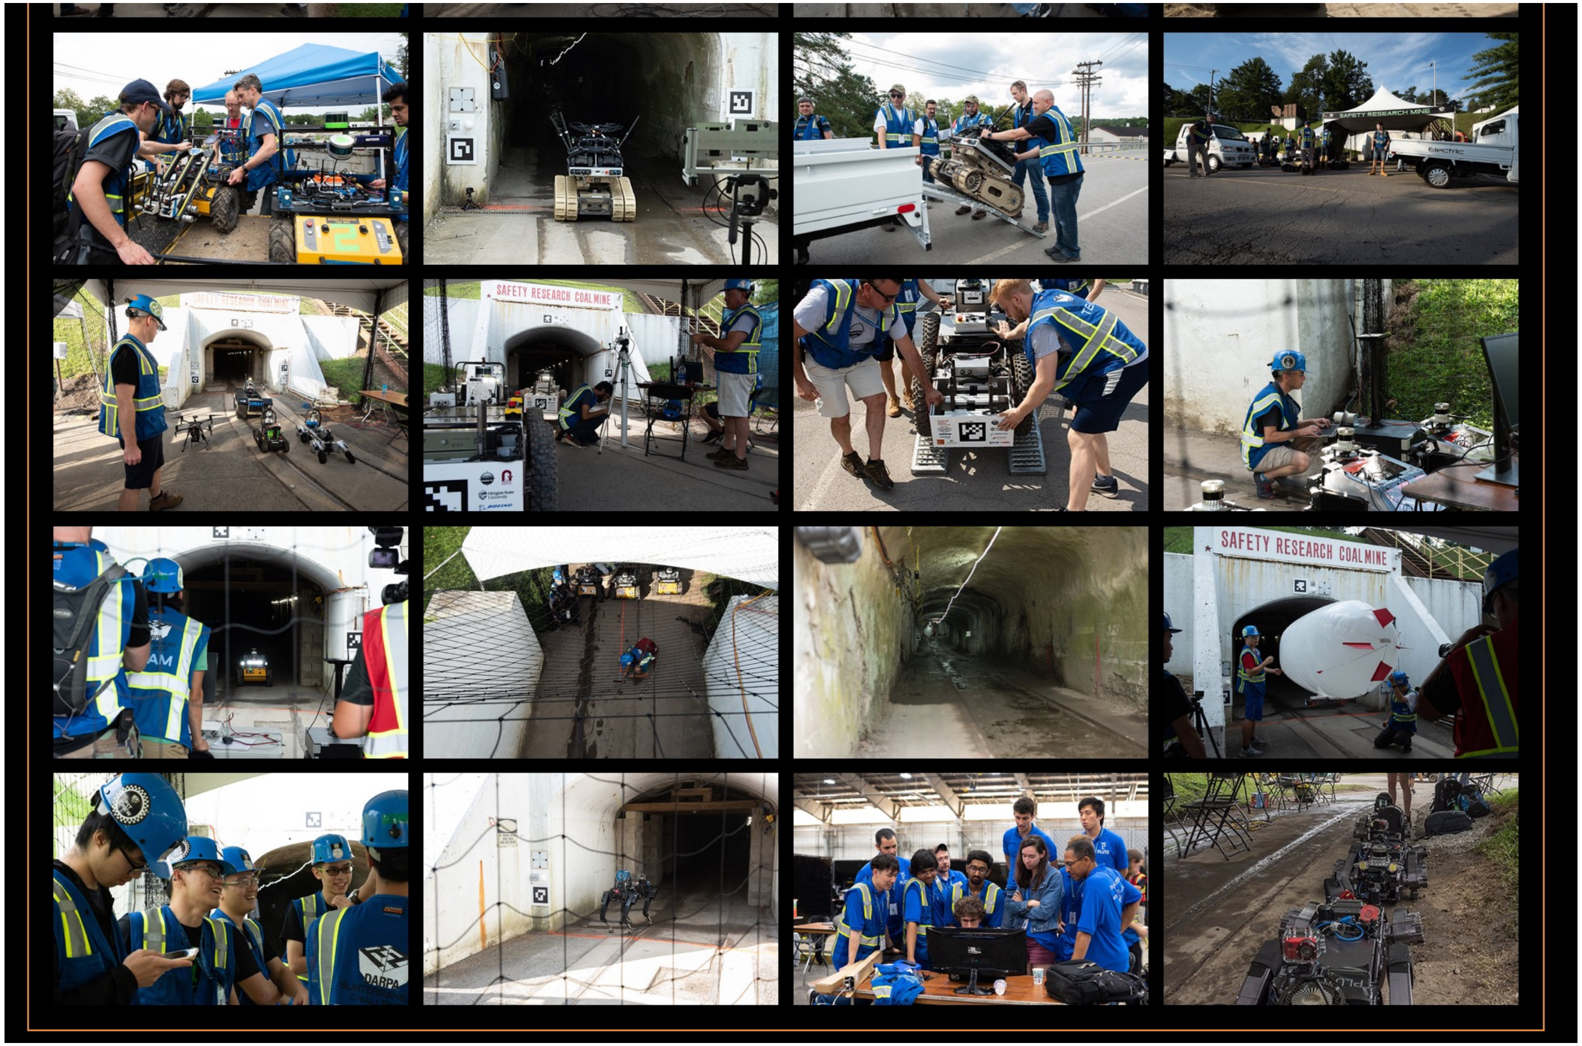
SAFETY RESEARCH COAL MINE
SAFETY RESEARCH COAL MINE
DARPA


ANYmal B
KLOUBAK UGV
Dr. Timothy H. Chung
DARPA Program Manager
Angela Maio
Virtual Competition Lead
60:00
34:59
24:50
57:18
48:48
57:18
GLYPH
World 01
THE HIGH ROAD
World 06

more complex…

we are here…outcome from wg2 cooperation between Heron Robots and Astrocent (T. Bulik and team)

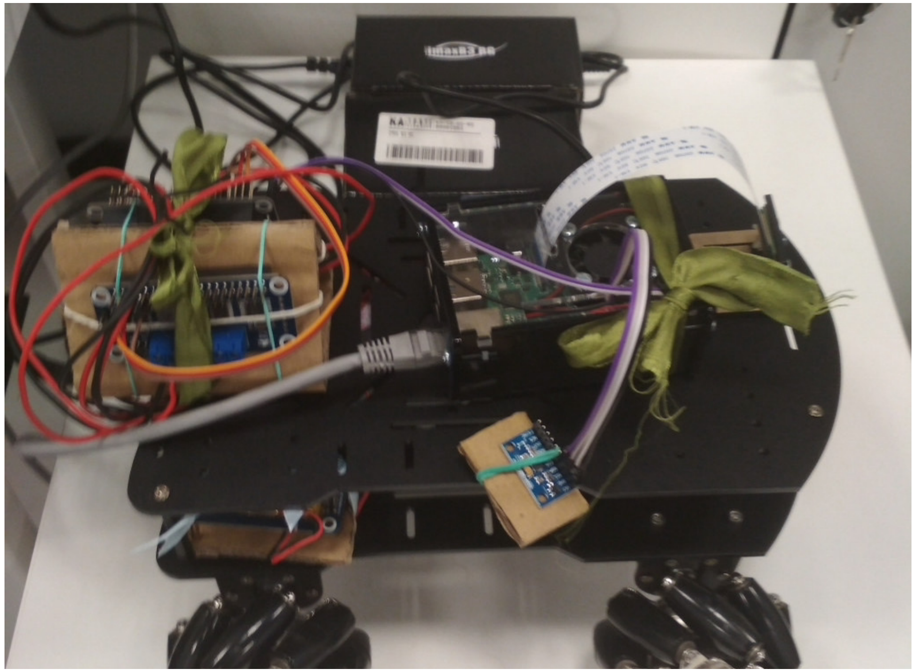

Mobile seismic sensor

- **to be used in inaccessible areas, possibility of characterizing seismic fields with a small number of sensors, adjustable sensor array layout, prototype ready**
- **Ros on Raspberry PI: SW stack can fit to any similar robot (including 'Roomba/Create') with minimum changes'**

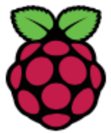

more complex…
we are here…outcome from wg2 cooperation between Heron Robots and Astrocent (T. Bulik and team)

Infrasound microphone

- **Needed to characterize the infrasound field, Low cost, Sensitivity in the range of 1-30Hz, lots of uses: geophysics, volcanology etc, Potential industrial applications, Prototype ready – network to be installed in Virgo this year**

In NASA footsteps ☺

https://mars.nasa.gov/insight/mission/quick-facts/

Key Facts About NASA's InSight

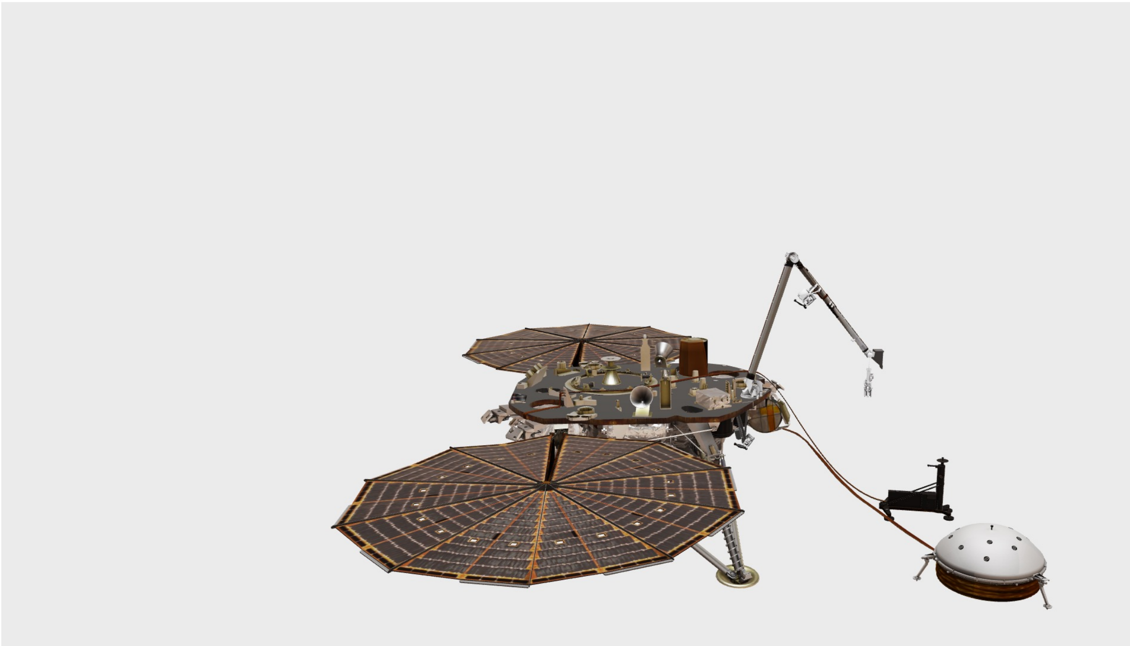

Specific ML/DL Challenges

Learning Multiple Matrix/Tensor Time Series

Not ‘so Big Data’

Learning on (Complex) manifolds

Entropy/Information Metrics on (Complex) Manifolds

Course of dimensionality

Curse of Dimensionality

from *Deep Learning*
www.deeplearningbook.org
Ian Goodfellow
2016-09-26

Figure 5.9

remember: Chess vs Go

Nearest Neighbor

from *Deep Learning*
www.deeplearningbook.org
Ian Goodfellow
2016-09-26

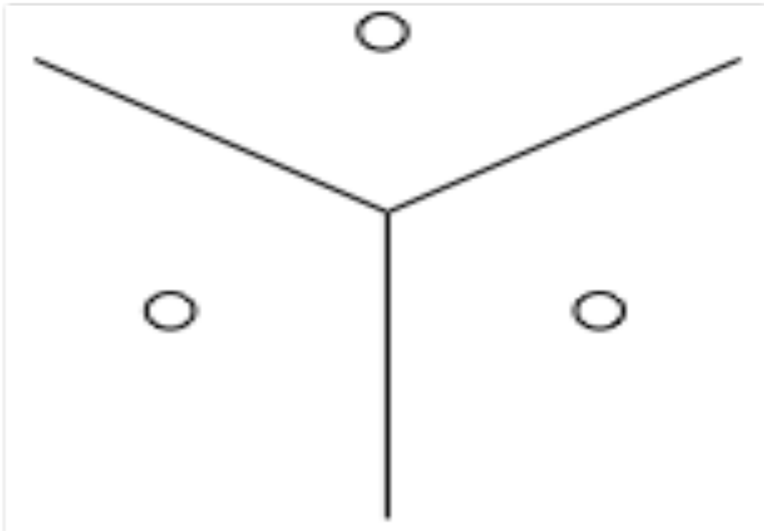

Figure 5.10

Manifold Learning

from *Deep Learning*
www.deeplearningbook.org
Ian Goodfellow
2016-09-26

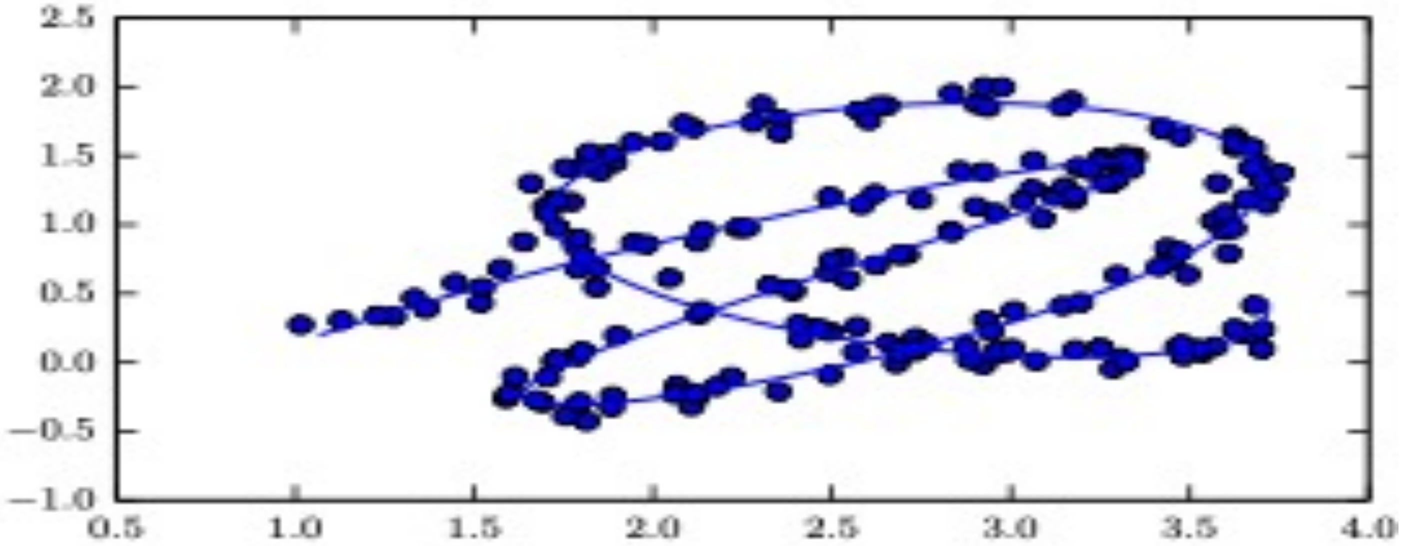

Figure 5.11

Increasing Depth

from *Deep Learning*
www.deeplearningbook.org
Ian Goodfellow
2016-09-26

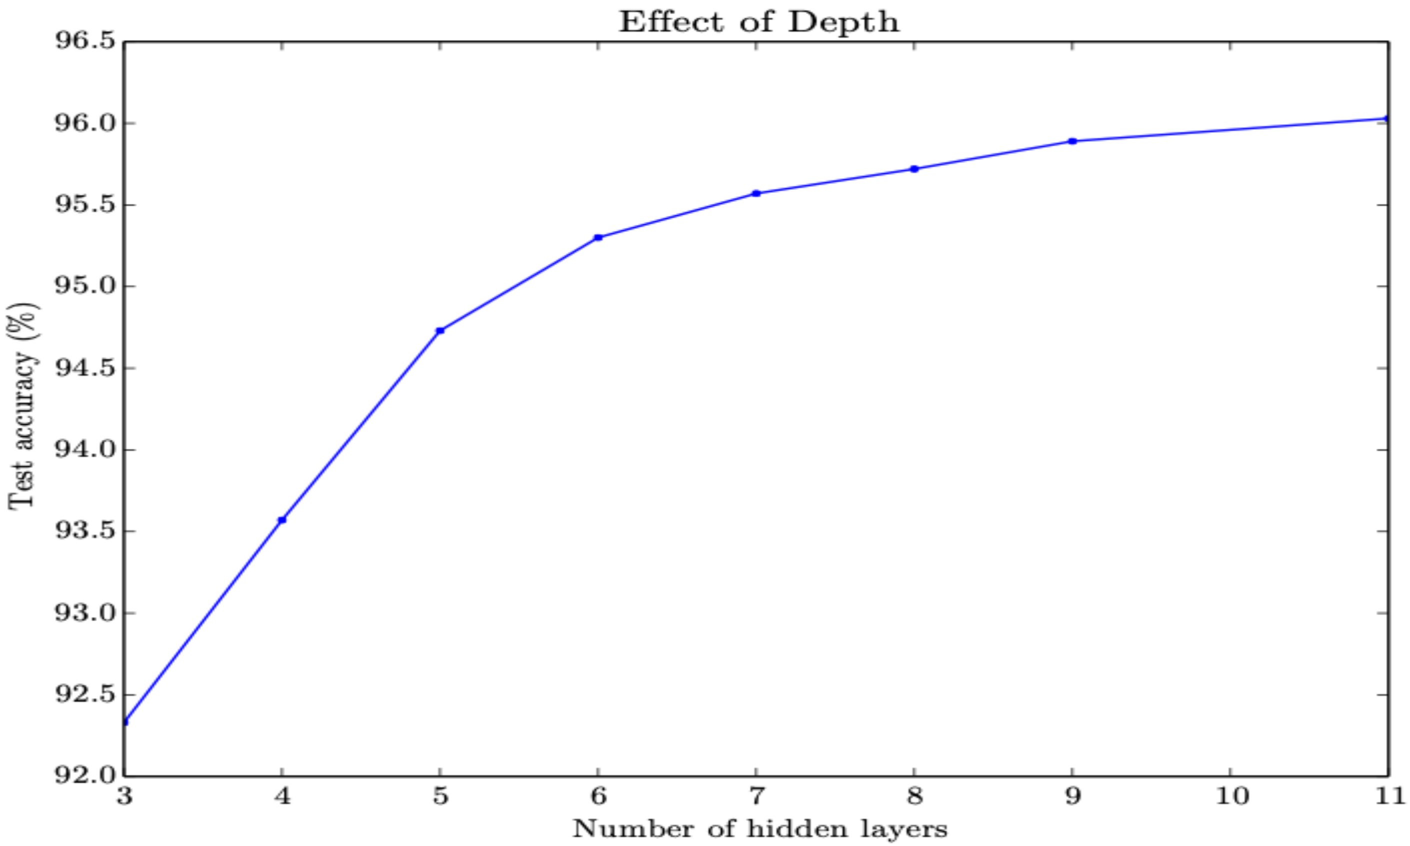

Hints that DL … *MUST* WORK

Towards a regularity theory for ReLU networks – chain rule and global error estimates

Julius Berner*, Dennis Elbrächter*, Philipp Grohs‡, Arnulf Jentzen§

*Faculty of Mathematics, University of Vienna
Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria

‡Faculty of Mathematics and Research Platform DataScience@UniVienna, University of Vienna
Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria

§Department of Mathematics, ETH Zürich
Rämistrasse 101, 8092 Zürich, Switzerland

***Abstract*—Although for neural networks with locally Lipschitz continuous activation functions the classical derivative exists almost everywhere, the standard chain rule is in general not applicable. We will consider a way of introducing a derivative for neural networks that admits a chain rule, which is both rigorous and easy to work with. In addition we will present a method of converting approximation results on bounded domains to global (pointwise) estimates. This can be used to extend known neural network approximation theory to include the study of regularity properties. Of particular interest is the application to neural networks with ReLU activation function, where it contributes to the understanding of the success of deep learning methods for high-dimensional partial differential equations.**

a way that admits a chain rule which is both rigorous as well as easy to work with. Chain rules for functions which are not everywhere differentiable have been considered in a more general setting in e.g. [16], [17]. We employ the specific structure of neural networks to get stronger results using simpler arguments. In particular it allows for a stability result, i.e. Lemma III.3, the application of which will be discussed in Section V. We would also like to mention a very recent work [18] about approximation in Sobolev norms, where they deal with the issue by using a general bound for the Sobolev norm of the composition of functions from the Sobolev space $W^{1,\infty}$.

Well,…

State of AI – some considerations

- The "billion" parameters club → how large, up to 175B these days!
- Cost (more later), about \$1 per 1000 parameters
- Interesting: outrageous cost for incremental improvement
  - Need research and theory
  - We can be more efficient in training algorithms
- Large models are driven by efficiency with small data
  - Sometimes… with transfer learning
- Power-law → parameter & computational power do not scale linearly (which is bad!)

(stolen from Giorgio Metta)

Well,...

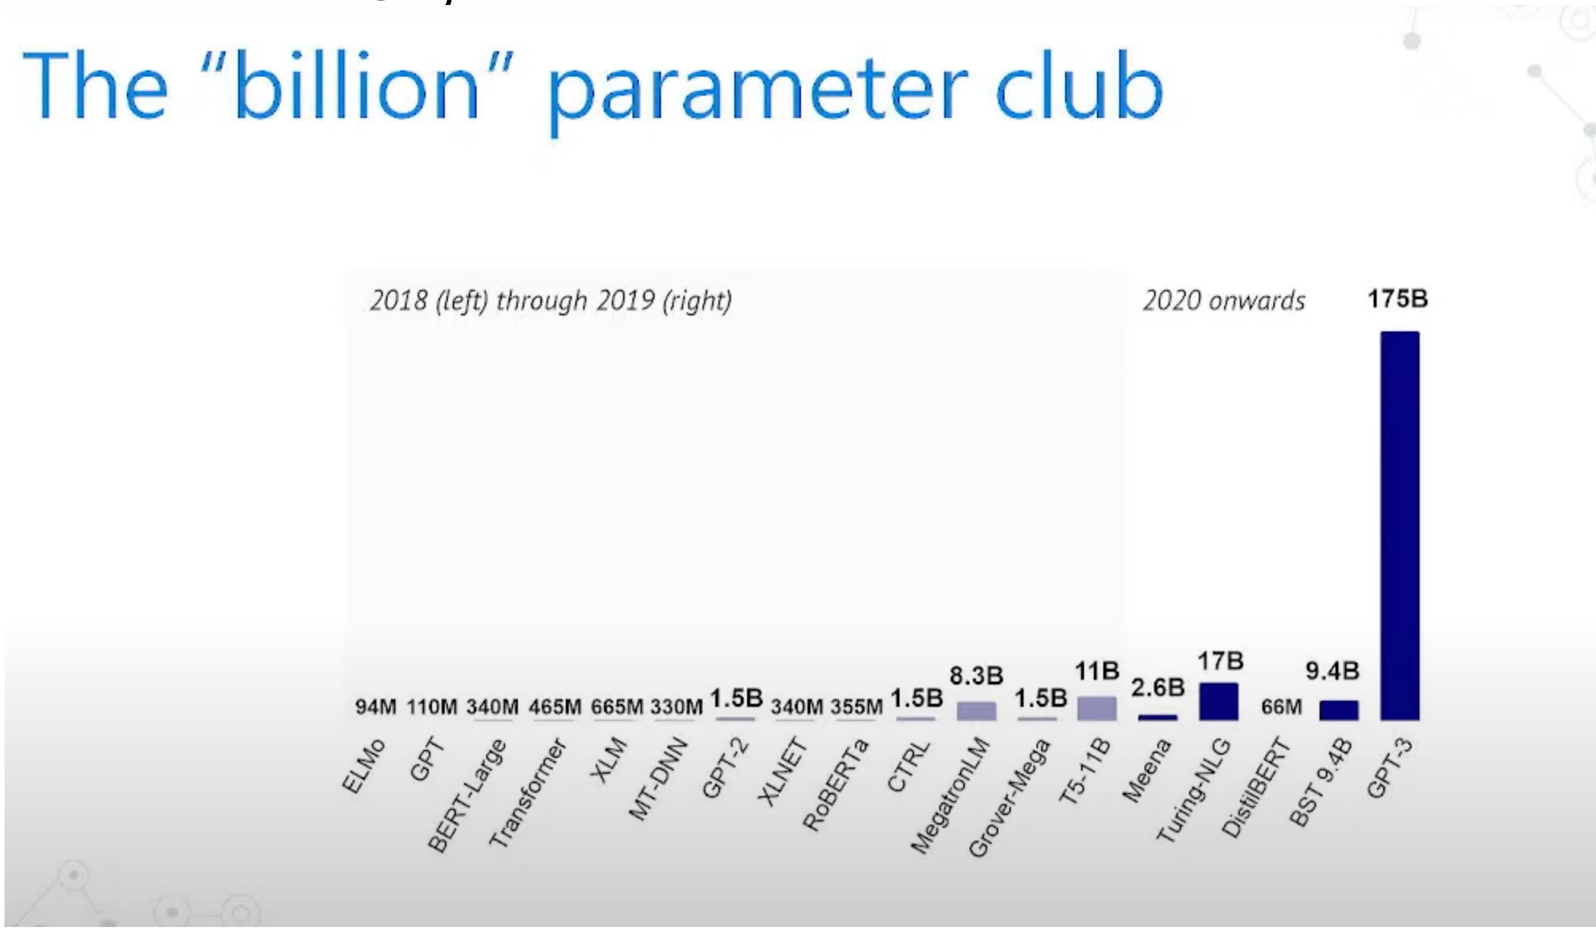

(stolen from Giorgio Metta)

Looking for new paths forward...
For example: Information self-structuring

•Experiments:

•Lungarella and Sporns**, 2006**
Mapping information flow in sensorimotor networks PLoS Computational Biology

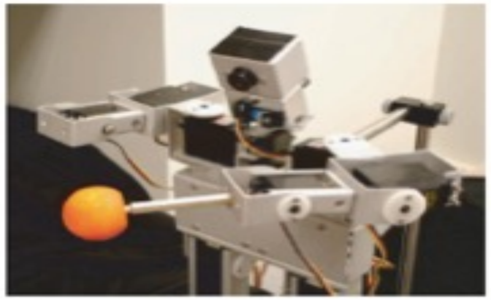

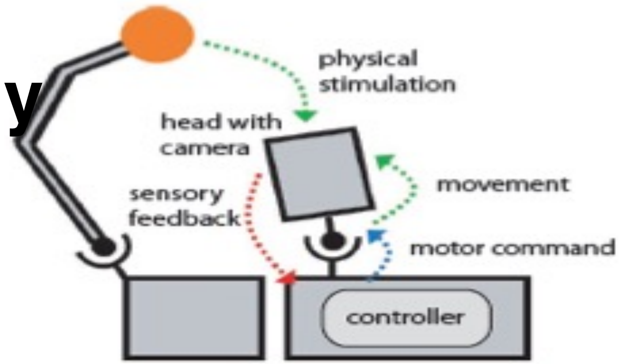

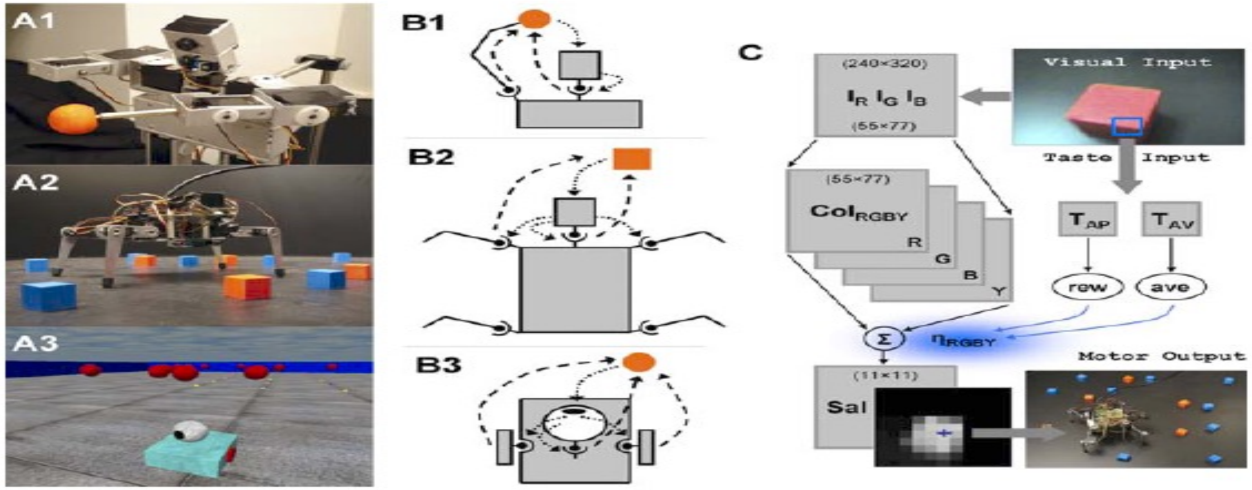

Lungarella, Sporns (2006)

Figure 1. Robots, Sensorimotor Interactions, and Neural Control Architecture

(A1) *Roboto* has a total of 14 DOF, five of which are used in the current set of experiments. Note the head-mounted CCD camera, the pan-tilt head system (2 DOF), and the moveable left arm with shoulder, elbow, and wrist joints (3 DOF). The object is a red ball (1.25 inches diameter) attached to the tip of the last joint.

(A2) *Strider* has a total of 14 DOF, with four legs of 3 DOF each and 2 DOF in the pan-tilt head system. Objects are red and blue blocks (1 inch cubes). *Strider* is situated in an environmental enclosure with black walls.

(A3) *Madame* has 4 DOF, with 2 DOF in the pan-tilt system and 2 DOF for the wheels, which are both located on an axis vertical to the main body axis. The environment is a square arena bounded by blue walls containing 20 red-colored floating spheres.

(B1) *Roboto* engages in sensorimotor interactions via the head system and arm movements; sensory → motor (dotted arrows), motor → sensory (dashed arrows).

(B2) *Strider* engages in sensorimotor interactions via the head system, as well as via steering signals generated by the head and transmitted to the four legs.

(B3) *Madame*'s behavior consists of a series of approaches to colored objects and ovations. Fixations to the objects are maintained by independent action of head and body.

(C) Neural control architecture. The architecture common to all robots is composed of color image arrays I_R, I_G, I_B, color- intensity map $\mathbf{Col}_{RGBY}$, and saliency map *Sal* (see text for details). The peak of the saliency map (blue cross) determines the pan-tilt camera motion and body steering. In addition, *Strider*'s neural system contains a value system with taste sensory inputs relayed via a virtual taste sensor (blue square in visual image) to taste neurons ($T_{AP,AV}$), which in turn generates reward and aversiveness signals (rew, ave). These signals are used to modulate the strengths of the saliency factors η_{RGBY} (see text for details).

DOI: 10.1371/journal.pcbi.0020144.g001

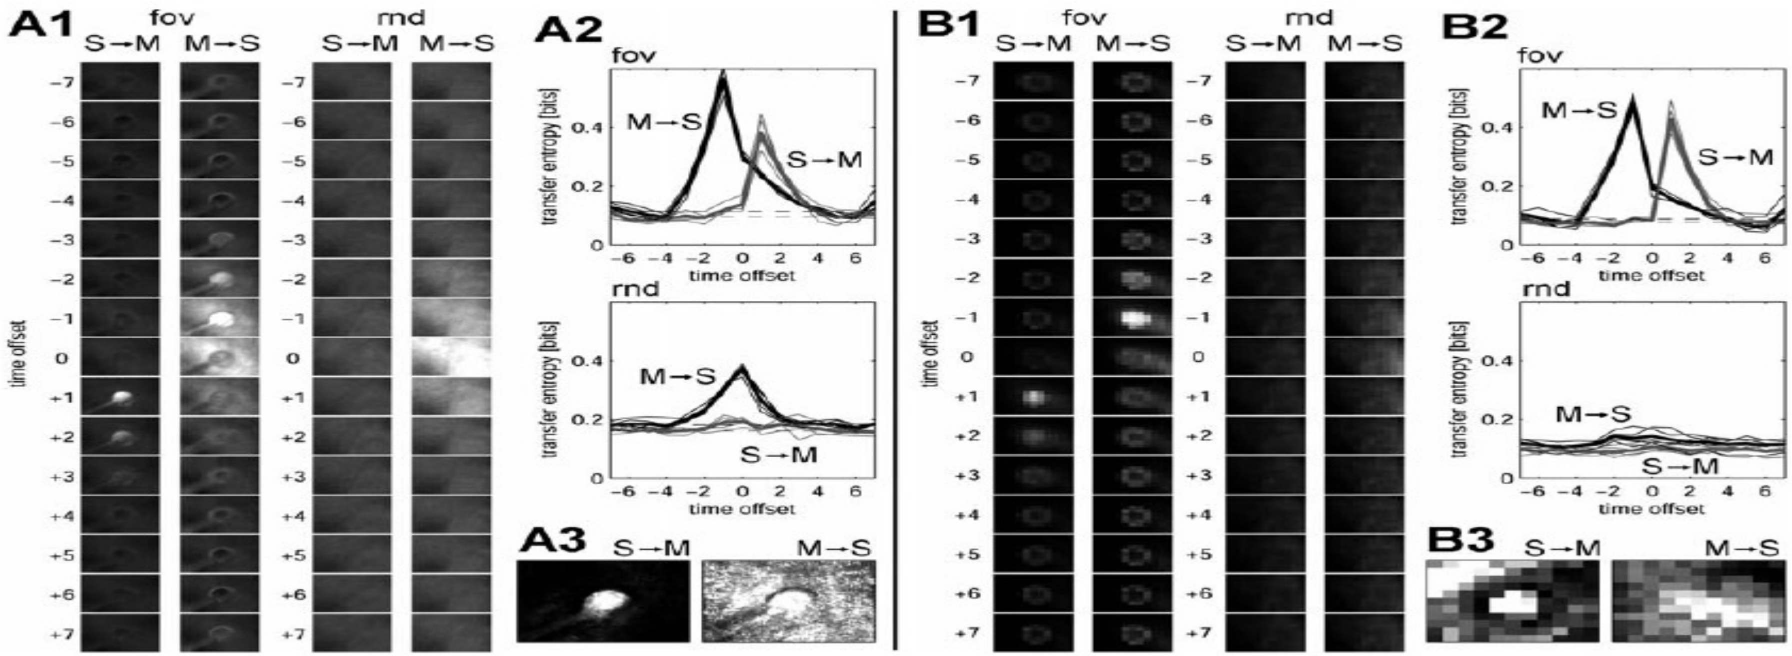

Figure 3. Information Flow (Transfer Entropy) between Sensory Input, Neural Representation of Saliency, and Motor Variables in *Roboto*

(A1) Transfer entropy between array I_R (variable S) and pan-tilt amplitude (variable M). Series of plots show maps of transfer entropy from S to M (S $\rightarrow$ M) and from M to S (M $\rightarrow$ S) over visual space (55×77 pixels), calculated for offsets between −7 ("M leading S") and +7 ("S leading M") time steps. Plots show data for conditions "fov" and "rnd." The gray scale ranges from 0.0 to 0.5 bits (for all plots in panels A1 and B1).

(A2) Curves show transfer entropy for five individual runs (thin lines) as well as the average over five runs (thick lines) between the single central pixel of array I_R (S) and pan-tilt amplitude (M), for directions M $\rightarrow$ S (black) and S $\rightarrow$ M (gray).

(A3) z-Score maps of significant image regions (plotted between $z = 0$ and $z = 6$). The z-scores are expressed as number of standard deviations above background at time offset +1 (S $\rightarrow$ M) and −1 (M $\rightarrow$ S). Mean and standard deviation of background is calculated from transfer entropy values at maximal time delays (−7,+7 time steps).

(B) All three panels have the same format as (A), but the neural activations of the saliency map *Sal* are substituted as variable S (11×11 neural units).

DOI: 10.1371/journal.pcbi.0020144.g003

Snakebot

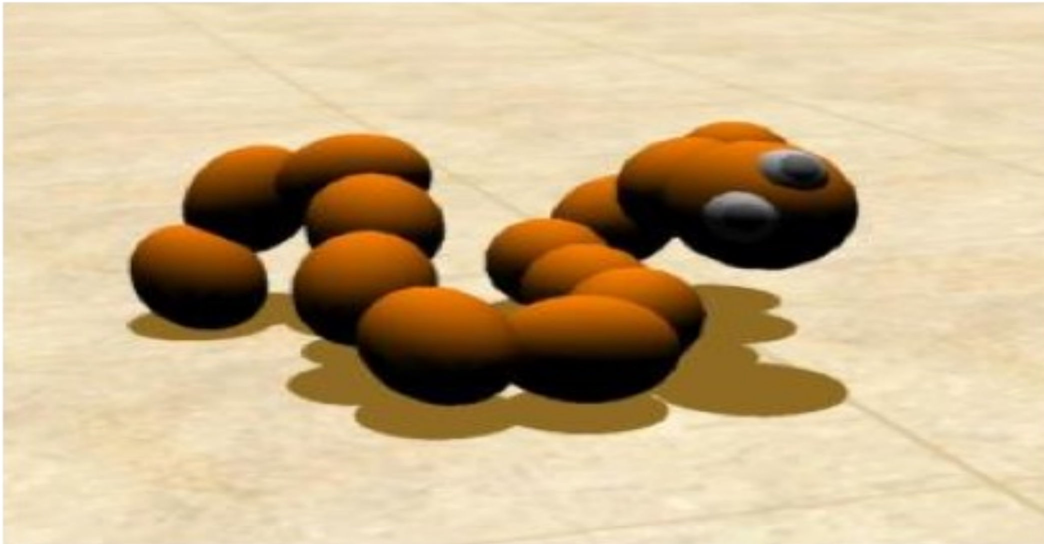

see: **Tanev et. al, IEEE TRO, 2005**

Maybe not GOF Euclidean space? :-)

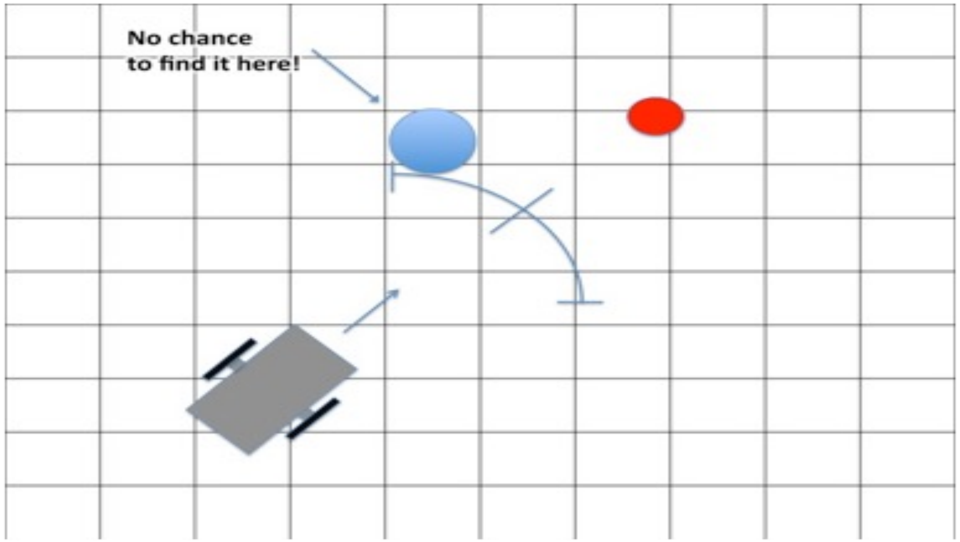

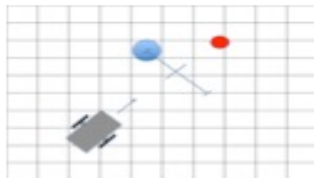

see: **Bonsignorio, Artificial Life, 2013**

Bottom Line: Physics Matters!

<u>Coping with the common underlying theoretical issues implied by the application of ML and DL to physical systems might have deep and wide scientific and technological impact</u>

Bottom Line: Physics Matters!

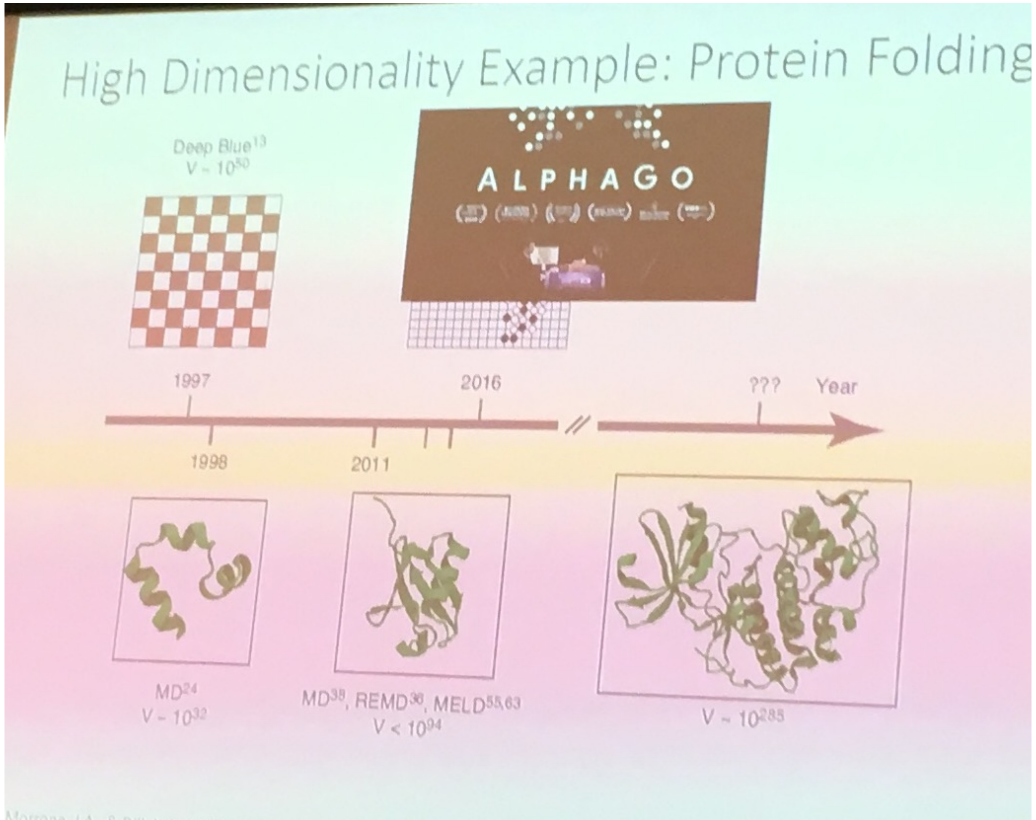

Probabilistic Model Of Control

- Although it may seem strange only in recent times the classical results from Shannon theory, have been applied to the modeling of control systems.
- As the complexity of control tasks namely in robotics applications lead to an increase in the complexity of control programs, it becomes interesting to verify if, from a theoretical standpoint, there are limits to the information that a control program must manage in order to be able to control a given system.

Probabilistic Model Of Control

Touchette, Lloyd (2004)

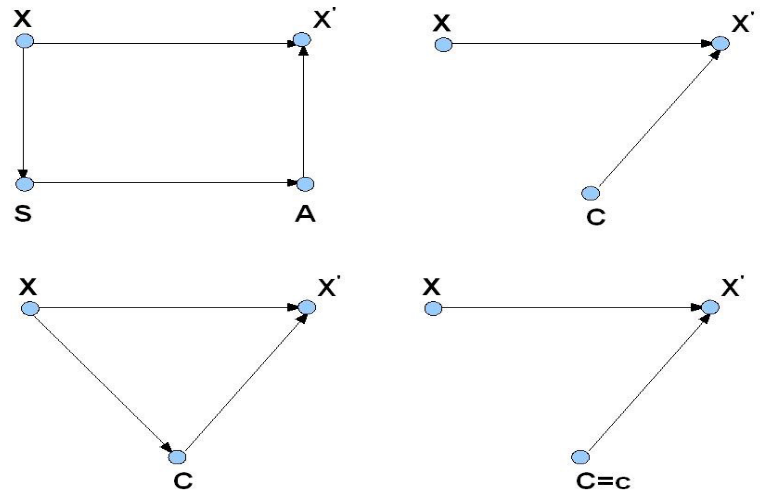

Directed acyclic graphs representing a control process. (Upper left) Full control system with a sensor and an actuator. (Lower left) Shrinked Closed Loop diagram merging sensor and actuator, (Upper right) Reduced open loop diagram. (Lower right) Single actuation channel enacted by the controller's state C=c.

Models of ‘Morphological Computation’

In [59], the network of agents, where each word is initially represented by a subset of three or more nodes with all (possible) links present, evolves towards an equilibrium state represented by fully connected graph, with only single links.

The statistical distribution, necessary to determine the information managing capability of the network of physical agents and to link to equation (2) can be obtained from equations derived in the statistical physics of network domain.

From (2) it is possible to derive the relations recalled here below (these relations are demonstrated in the appendix).

$$K(X) \overset{+}{\leq} \log \frac{W_{closed}}{W_{open}^{\max}} \qquad \text{(I)}$$

As told, relation (I) links the complexity ('the length') of the control program of a physical intelligent agent to the state available in closed loop and the non controlled condition. This shows the benefits of designing system structures whose 'basin of attractions' are close to the desired behaviors in the phase space.

$$\Delta HN + \sum_{i}^{n} \Delta H_i - \Delta I \leq I(X;C) \qquad \text{(II)}$$

Relations (II) links the mutual information between the controlled variable and the controller to the information stored in the elements, the mutual information between them and the information stored in the network and accounts for the redundancies through the multi information term *ΔI.*

Relations (III) links the program complexity of the controller to the information stored in the elements, the mutual information between them and the information stored in the network.

$$K(X) = \Delta HN + \sum_{i}^{n} \Delta H_i - \Delta I \qquad \text{(III)}$$

Relations (IV) links the program complexity of the controller to the information stored in the elements the mutual information between them and the information stored in the network.

$$\Delta HN = \log \frac{\Omega_{closed}}{\Omega_{open}^{\max}} + \Delta I \qquad \text{(IV)}$$

These relations are quite preliminary, and perhaps need a more rigorous demonstration, but give an insight on how information is managed within a network of physical elements or agents interacting with a given environment in a finalized way. They suggest how the cognitive adaptation is at network level: in any environment niche it is possible with small networks of highly sophisticated individual agents, like in human societies, or with many limited autonomy individuals like in ant colonies, with a great variety of possibilities in the middle.

My point of view :-)

- Information related measures coming from Shannon entropy may help the understanding of intelligent cognitive controlled systems
- What we probably need to be able to build 'real' artificial cognitive systems is a deep interchange of concepts, methods and insights between fields so far considered well distinct like information and control theory, non linear dynamics, general AI and psycology and neurosciences.

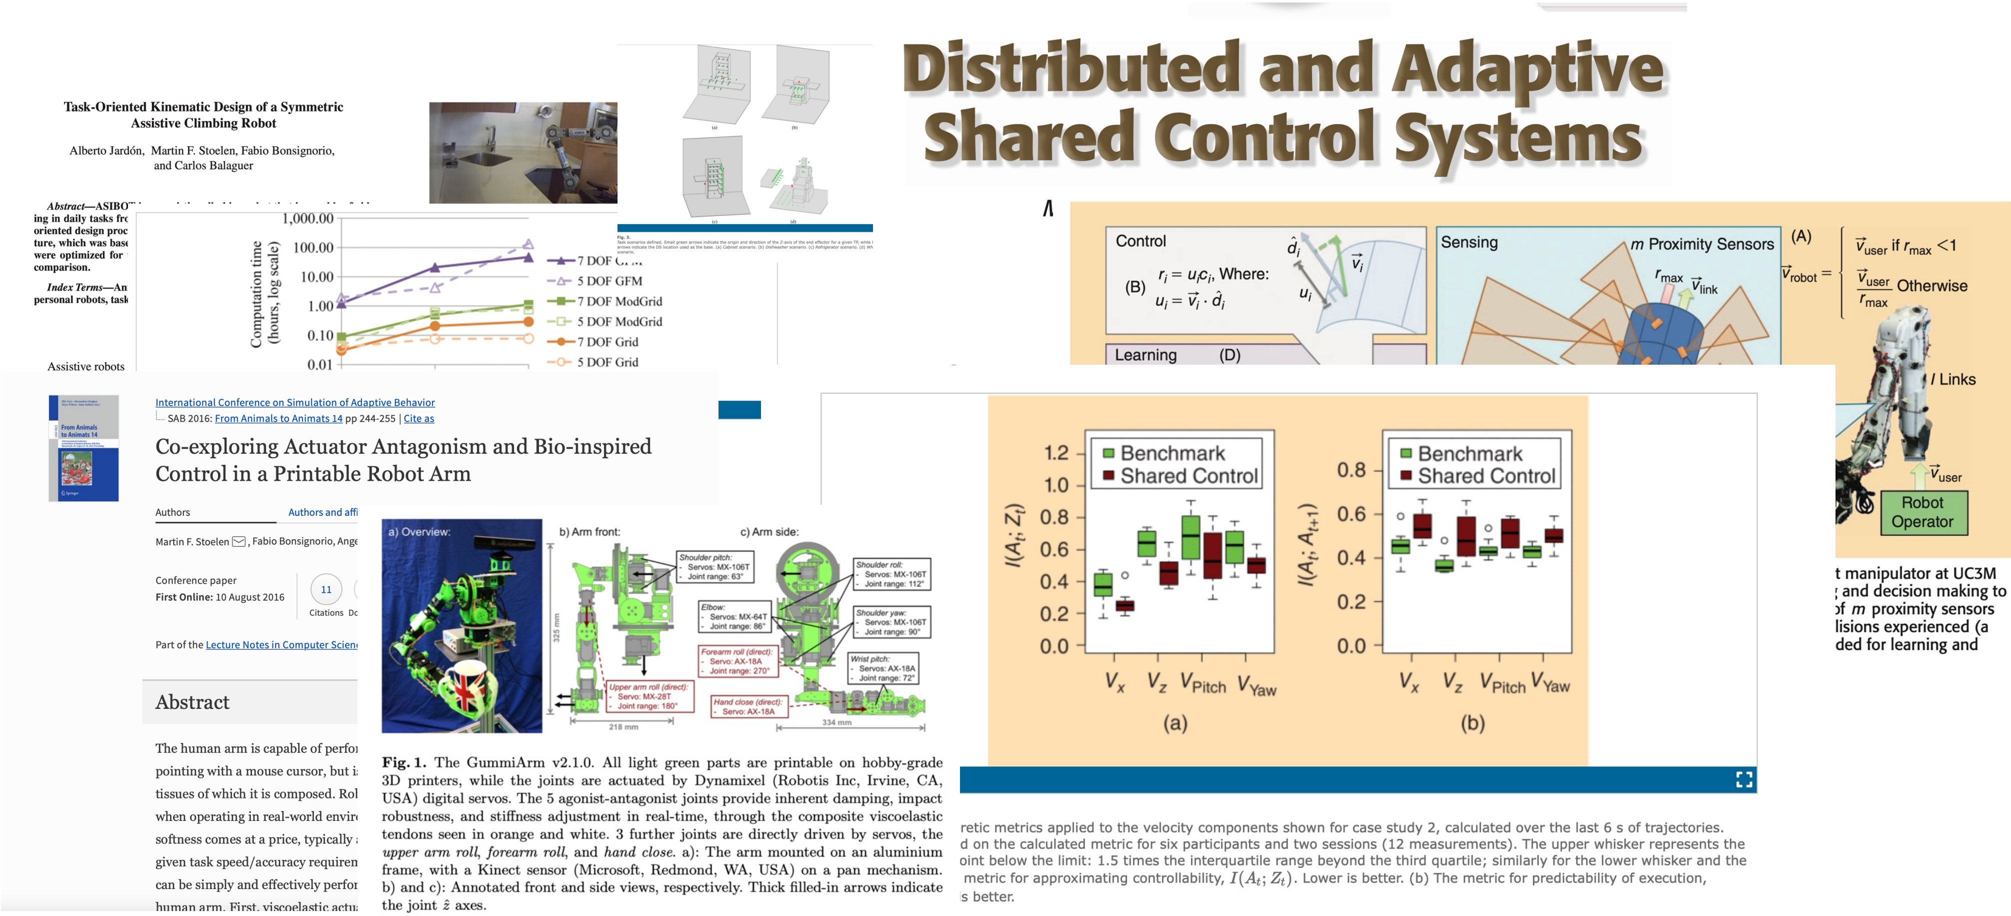

Stoelen M.F., Bonsignorio F., Cangelosi A., Co-exploring actuator antagonism and bio-inspired control in a printable robot arm, In Procs ofInternational Conference on Simulation of Adaptive Behavior, 244-255,2016

Stoelen M. F. , de Tejada V. F., Huete A. J., Balaguer C., Bonsignorio F., Distributed and Adaptive Shared Control Systems: Methodology for the Replication of Experiments, IEEE Robotics & Automation Mag. , 22(4),137–146, 2015

Jardon A. , Stoelen M., Bonsignorio F.P., Balaguer C. ,Task-oriented kinematic optimization of a symmetric assistive climbing robot, IEEE T-RO, 27 (6), 1132-1137, 2011

R-Article

A Simple Visual-Servoing Task on a Low-Accuracy, Low-Cost Arm

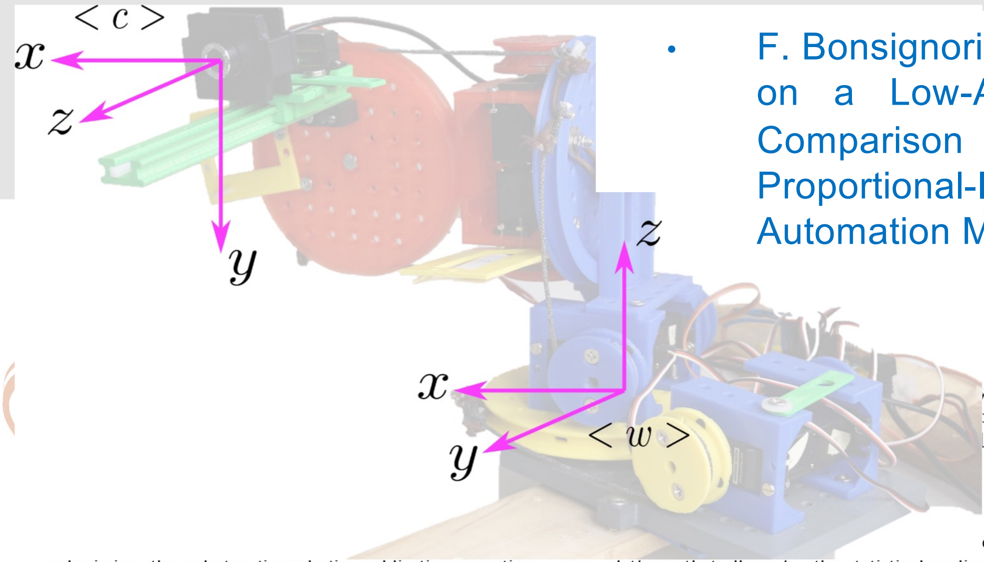

- F. Bonsignorio and E. Zereik. A Simple Visual-Servoing Task on a Low-Accuracy, Low-Cost Arm: An Experimental Comparison Between Belief Space Planning and Proportional-Integral-Derivative Controllers. IEEE Robotics & Automation Magazine. 2020, early access

PHOTOCREDIT

wrong
ice may
esirable

d hard-

submission, the only top-tier robotics publication accepting

ware platform that allows for the statistical replication of

R-Articles

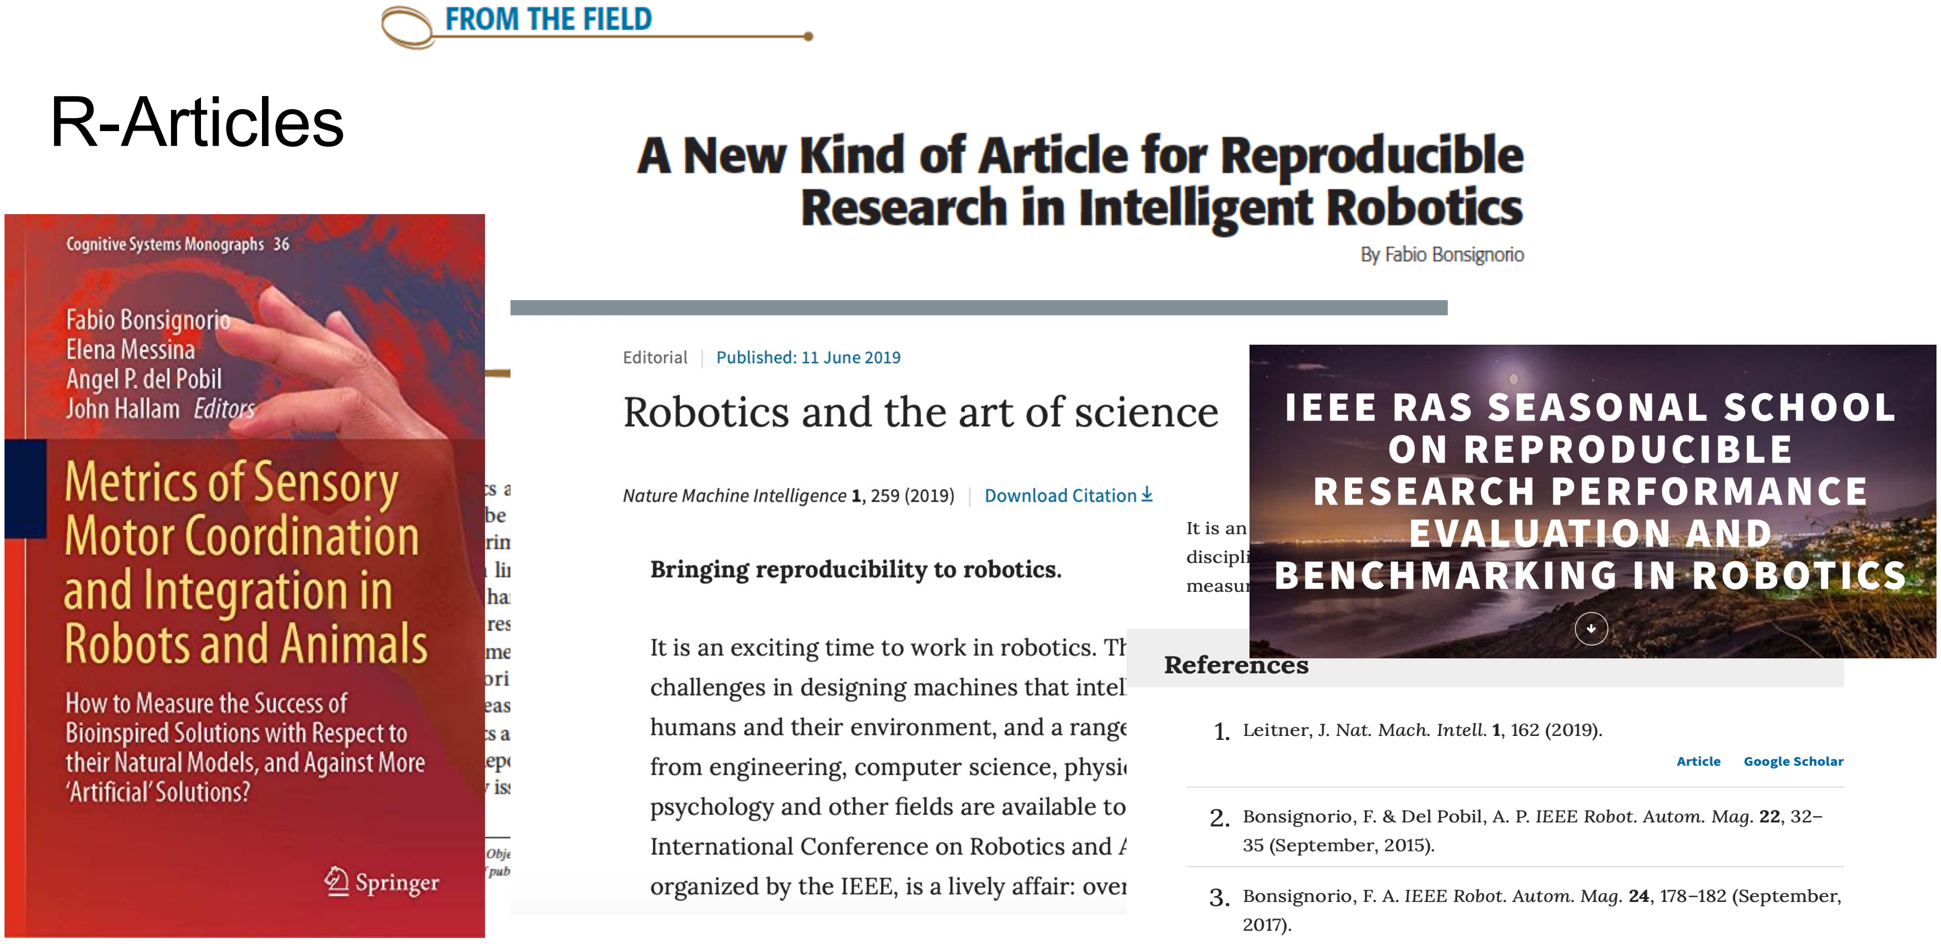

Bonsignorio F., A new kind of article for reproducible research in intelligent robotics, IEEE Robotics & Automation Magazine 24 (3), 178-182, 2017

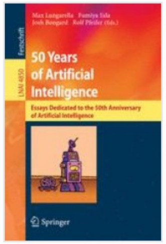

50 Years of Artificial Intelligence pp 112-123 | Cite as

Preliminary Considerations for a Quantitative Theory of Networked Embodied Intelligence

Authors | Authors and affiliations

Fabio P. Bonsignorio

Part of the Lecture Notes in Computer Science book series (LNCS, volume 4850)

embodiment:

cognition emergent from sensory-motor and interaction processes

Abstract

This paper exposes and discusses the concept of 'networked embodied cognition', based on natural embodied neural networks, with some considerations on the nature of natural collecti[ve] intelligence and cognition, and with reference to natural biological examples, evolution theory neural networks science and technology results, network robotics, it shows that this could be the method of cognitive adaptation to the environment most widely used by living systems and

www.shanghailectures.com

Bonsignorio, F., Preliminary considerations for a quantitative theory of networked embodied intelligence, 50 years of artificial intelligence 4850, 112-123, 2007

An Imitation Learning Approach for the Control of a Low-Cost Low-Accuracy Robotic Arm for Unstructured Environments

Fabio Bonsignorio[1], Cristiano Cervellera[2†], Danilo Macciò[2†] and Enrica Zereik[2*†]

[1], Heron Robots, Via Malta 3/7, Genoa, 16121, Italy.
[2*]Institute of Marine Engineering, Italian National Research Council, Via de Marini 16, Genoa, 16149, Italy.

*Corresponding author(s). E-mail(s): enrica.zereik@cnr.it;
†These authors contributed equally to this work.

Abstract

We have developed an imitation learning approach for the image-based control of a low-cost low-accuracy robot arm. The image-based control of manipulation arms is still an unsolved problem, at least under challenging conditions such as those here addressed. Many attempts

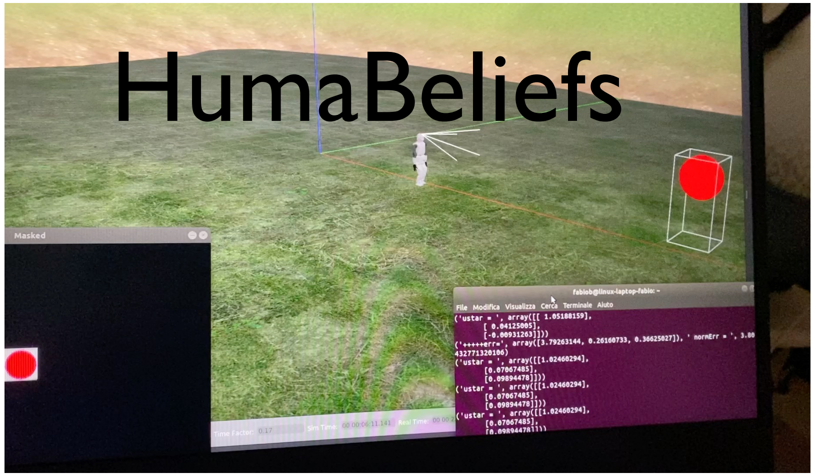

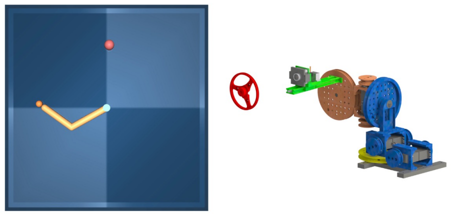

(a) DeepMind Reacher (b) H_2Arm visual-servoing

Fig. 9: Comparison among the proposed H2Arm visual-servoing task and the "Reacher" task of the DeepMind Control Suite. The main characteristics of each task are: a) DeepMind Reacher – a 2-link planar structure that has to reach a target, executed only in simulation, with known proprioceptive measures, known target location, simulated scenario with known noise structure, many training data needed, AI directly on image pixels. Tasks are strongly observable, position and velocity observations depend only on the current state. Sensor readings only depend on the previous transition, see [21]. Courtesy of DeepMind. b) H2Arm – 4-link 3D manipulator, experimented in real world, without proprioceptive information (the only sensor on-board is the wrist-mounted camera), additional noise injected in some of the experiments, very few needed training data (97 BSP trajectories logged in previous tests where the arm was controlled by the BSP algorithm only, without any neural controller), images are pre-processed by a vision algorithm and AI works on measures estimated by vision.

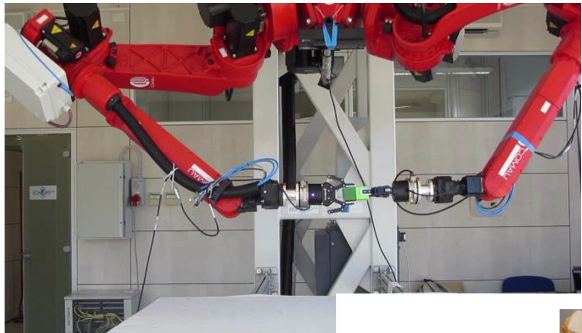

Fig. 15: After the object was picked up positioned in a suitable configuration to perf two grippers facing each other, and the fore:

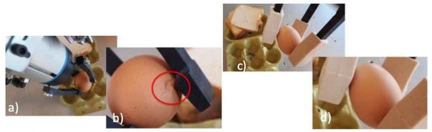

Fig. 18: View of the result of the first pick and place task using eggs. Confront between the broken egg resulted from the manipulation of the Festo grippers (a,b) and the sane egg manipulated with silicon thimbles (c,d).

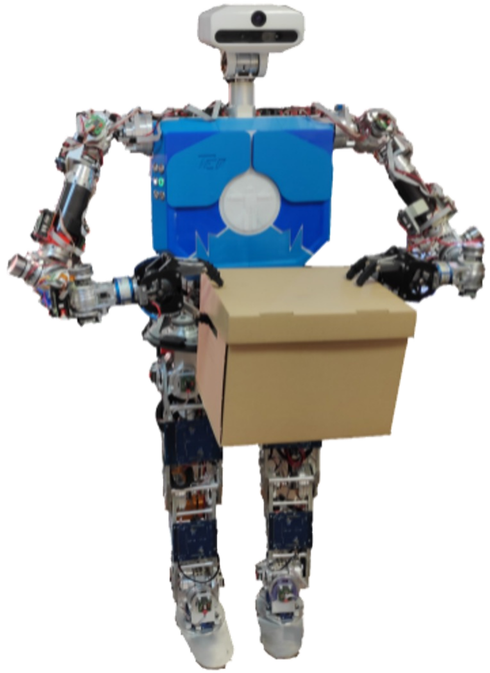

HumaBiMan

Dario P., Morachioli A., Strazzulla I., Laschi C.,Bonsignorio F., "Disassembly Robotic Tasks for Circular Economy"(poster), IEEE Life Sciences Grand Challenges Conference, Abu Dhabi, UAE, 2016

The 'research space' we should – imo - explore (and that I have actually been exploring and I'm continuing to explore…..)

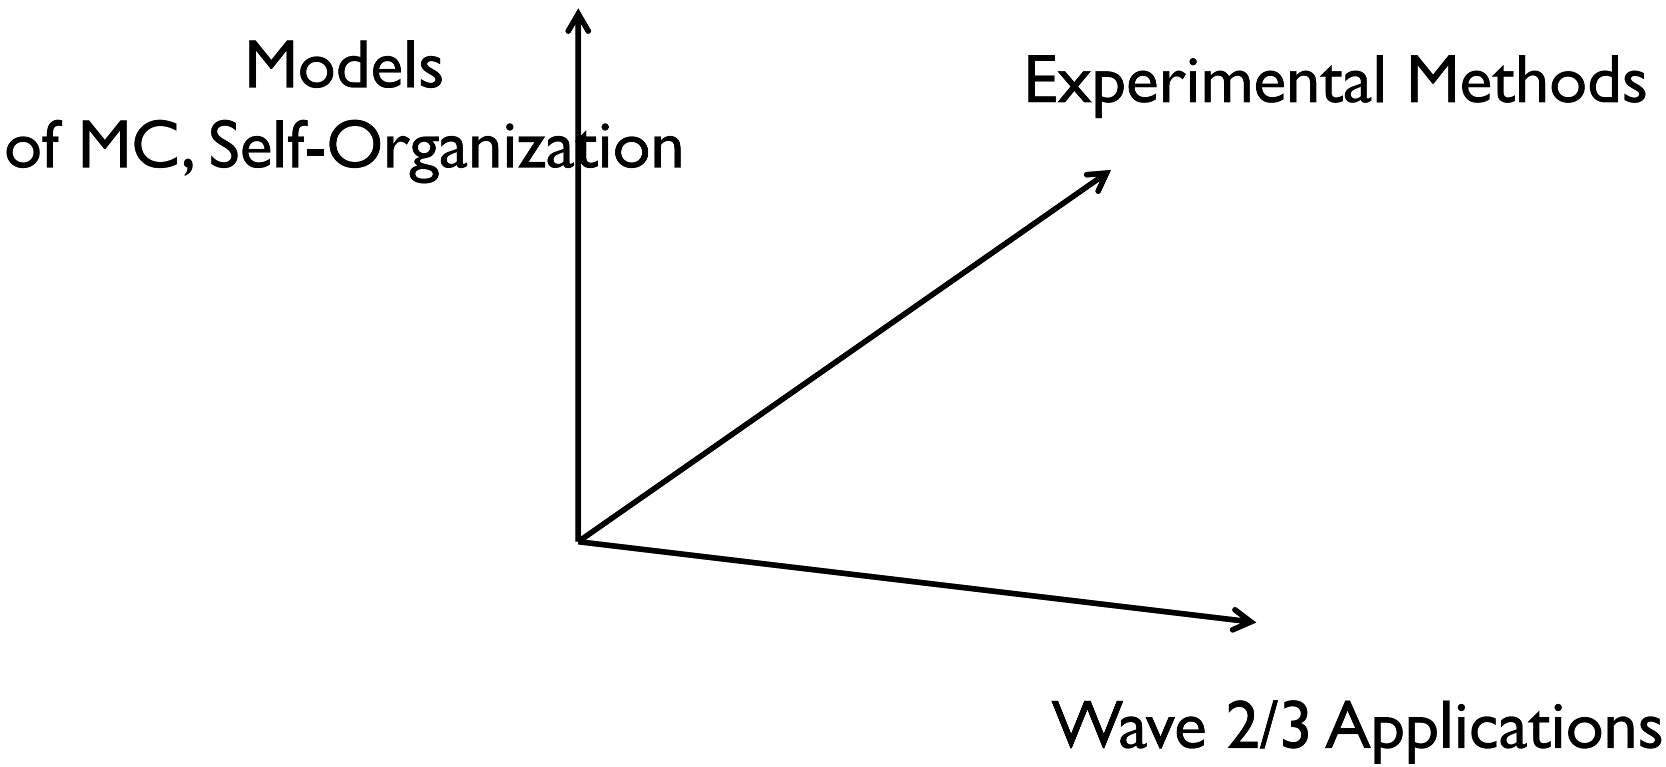

The 'research space' we should – imo - explore (and that I have actually been exploring and I'm continuing to explore….)

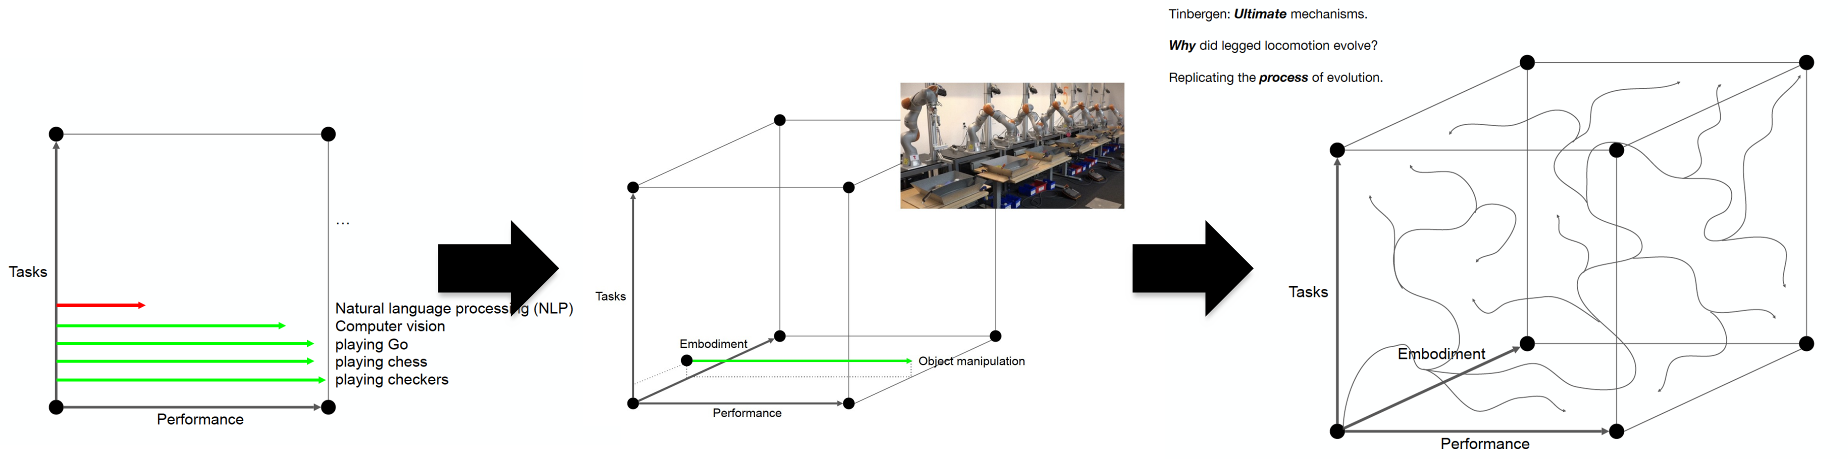

from Joshua Bongard, University of Vermont

The link between Morphological Computation and Soft Robotics

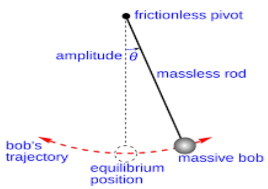

$$T \approx 2\pi\sqrt{\frac{L}{g}}$$

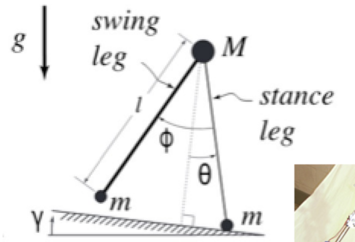

(Andy Ruina)

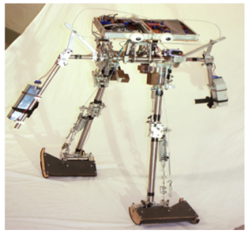

T=f(l/g)

Fixed speed!

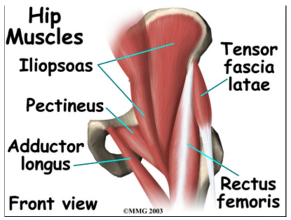

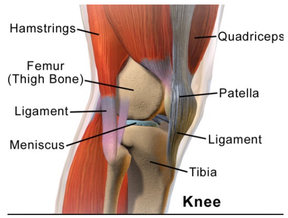

(Wikipedia)

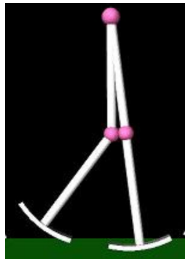

(Fumihiko Asano)

T=f(l/g)

l=f(controlled input)

Speed can change!

Quantitative Modelling of the <u>trade-offs</u> between <u>physical morphology</u> (and associated dynamics)
and <u>information processing</u> is crucial

That's what Morphological Computation is about.
It explains why 'soft' components help many task performances and can provide design guidance.

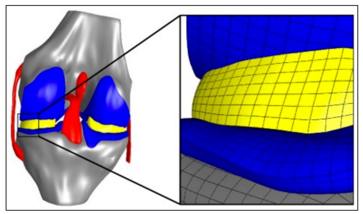

(Yale Image Finder)

Some references (Feedback is welcome!)

F. Bonsignorio, Preliminary considerations for a quantitative theory of networked embodied intelligence,
In: 50 years of artificial intelligence, 112-123, Springer, 2007

F. Bonsignorio, Steps to a cyber-physical model of networked embodied anticipatory behavior,
In: Anticipatory Behavior in Adaptive Learning Systems, LNAI, 549, 77-94, Springer, 2008

F. Bonsignorio, On the Stochastic Stability and Observability of Controlled Serial Kinematic Chains, ESDA2010-25131, 379-386, ASME, 2010

F. Bonsignorio, Quantifying the evolutionary self-structuring of embodied cognitive networks, Artificial life 19 (2), 267-289, MIT Press, 2013

F Bonsignorio, E Messina, AP Del Pobil, J Hallam, Metrics of Sensory Motor Coordination and Integration in Robots and Animals: How to Measure the Success of Bioinspired Solutions with Respect to their Natural Models, Cognitive Systems Monographs, Springer, 2020

F Bonsignorio, D Hsu, M Johnson-Roberson, J Kober, Deep Learning and Machine Learning in Robotics, IEEE Robotics & Automation Magazine 27 (2), 20-21, 2020

Thank you!

fabio.bonsignorio@heronrobots.com
fabio.bonsignorio@gmail.com