Adding incrementality to s/p control: from *training a machine* to *interacting with a human*

Claudio Castellini

Robotics and Mechatronics Center
DLR – German Aerospace Center
Oberpfaffenhofen, Germany

From *training a machine* to *interacting with a human*

01.The egg.wmv

02.Ultrasound example.mpg

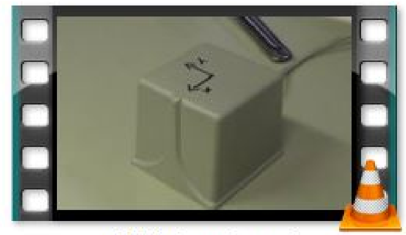

03.Ultraharmonium.mp4

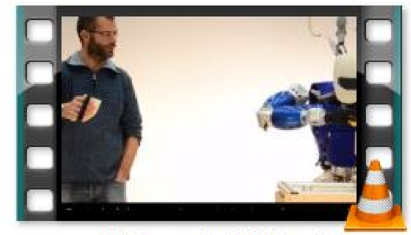

04.Teleoperation TORO.mp4

05.Realistic sEMG.wmv

06.Azzurra in a SHAP.wmv

2014.10.DEMOVE III.pptx

DLR – The German Aerospace Center

The national aeronautics and space research centre of the Federal Republic of Germany

Aeronautics, space, energy, transport, security

Planning and implementation of the **German space programme**

~**7400 employees at 16 locations** in Germany

Our mission:
exploration of the earth and of the solar system
research for protecting the environment

Campus Oberpfaffenhofen
8 scientific institutes, ~1500 employees

Space missions, climate research, earth observation, navigation, **robotics**

1993: teleoperated catching a free-floating object of the *rotex* arm aboard the D2 Shuttle mission

The Robotics and Mechatronics Center (RMC)

Founded in 2010

Previously called *Institute of Robotics and Mechatronics*

Research areas

Systems and control theory

Perception and cognition (advanced sensors)

Autonomy and remote control

Mechatronic components and systems

Optical information systems

Mission

Transfer of aerospace technology to terrestrial applications

Close cooperation between robotics research and industry

The ideal prosthesis...

- ...is like a pair of glasses:
  - you buy them (and they look good)
  - you use them (and they work o.k. all day long)
  - you put them on your bedside table in the evening
  - you put them on again day after
  - *and they work again, just as well.*

(Peter J. Kyberd, personal communication @ MEC 2014)
(...more like a new motorbike, perhaps...)

- Search *EMG Prosthesis Control* in `ieeexplore.org`, 1970-2014:
  - 387 papers (of which 106 journals),
  - exponential growth from 2000 on.

- Ok, so we should be almost there. Or, are we?

DLR

Remarks from a friend

"That is what **myoelectric arms** are still best at: being paraded around as gadgets, being admired by society. In any amputee view, they **are demonstrably and understandably, repeatedly and repetitively worse than not wearing a prosthetic arm**."

Wolf Schweitzer,
Technical Below Elbow Amputee Issues
Russian prosthetic arm [about the history of myoelectric arms]
http://www.swisswuff.ch/tech/?p=2366

Remarks from a friend

"Appearance aside, **reliability/performance is the single-most relevant factor** in prosthetic acceptance for many upper-limb prosthetic users. Shopping, ironing, handling containers, and preparing meals are examples of essential tasks the person will need to perform with the device. And given the somewhat reduced social acceptance that arm amputees seem to experience, attending parties or functions are of particular interest, and with those events comes the handling of delicate, breakable items, such as drinking glasses, or slippery objects, such as olives. **The confidence that you will not drop or break anything outweighs any other aspect, even appearance.** So even without wearing a prosthetic arm at all, I may perform with 100 percent reliability and therefore be more accepted socially than when wearing a high-tech arm."

Wolf Schweitzer, *Improving Prosthetic Arms through Better Testing*
The O&P Edge, August 2014

My personal adventure in biorobotics

- I am originally a computer scientist (Ph.D. in logics!)
- Joined the Neurobotics project in 2005...
- ...where I started playing around with machine learning.

- Applied Support Vector Machines to
  - models of reaching and grasping
  - gaze tracking
  - speech processing
  - an EMG-controlled hand prosthesis (j.w.w. DLR)

Castellini, C.; Orabona, F.; Metta, G. & Sandini, G. *Internal models of reaching and grasping* Advanced Robotics, 2007, 21, 1545-1564

Castellini, C. *Gaze tracking in semi-autonomous grasping* Journal of Eye Movement Research, 2009, 2, 1-7

Castellini, C.; Badino, L.; Metta, G.; Sandini, G.; Tavella, M.; Grimaldi, M. & Fadiga, L. *The use of phonetic motor invariants can improve automatic phoneme discrimination* PLoS ONE, 2011, 6, e24055

- Take-home lessons: use machine learning it only if required
  - only if no model available / model too complex / signals very unreliable
  - *robotic motion learning:* mostly **no**
  - *human biosignals:* almost always **yes**

Orabona, F.; Castellini, C.; Caputo, B.; Jie, L. & Sandini, G. *On-line independent Support Vector Machines* Pattern Recognition, 2010, 43, 1402-1412

More adventures in biorobotics...

- In particular, we applied a few machine learning method to sEMG-based control of a dexterous humanoid hand:
  - sEMG-based **classification** of grasp configurations
  - sEMG-based **regression** on required force

- Gathered data across two days (single subject)
- Good online performance

(see movie #1)

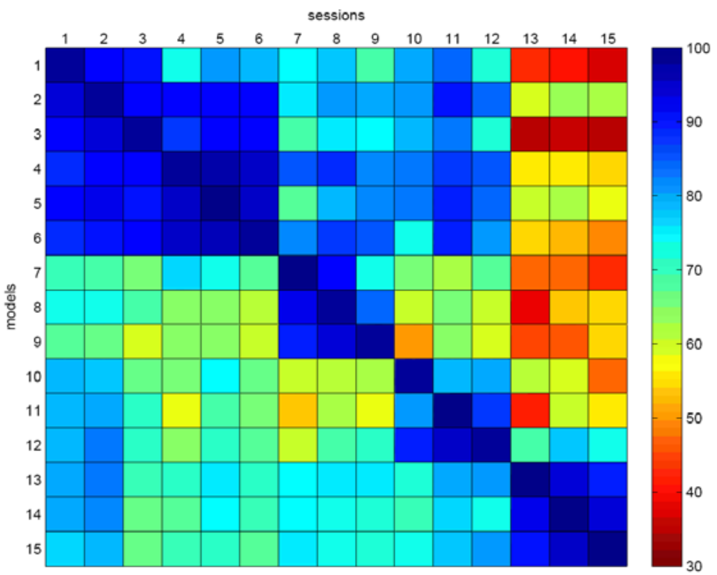

Castellini, C. & van der Smagt, P. *Surface EMG in advanced hand prosthetics*
Biological Cybernetics, 2009, 100, 35-47

...yet more of them...

- A few more details:
  - the problem turned out to be rather easy from the ML point of view
  - no clearly winning approach
  - surprisingly few samples account for most of the problem complexity

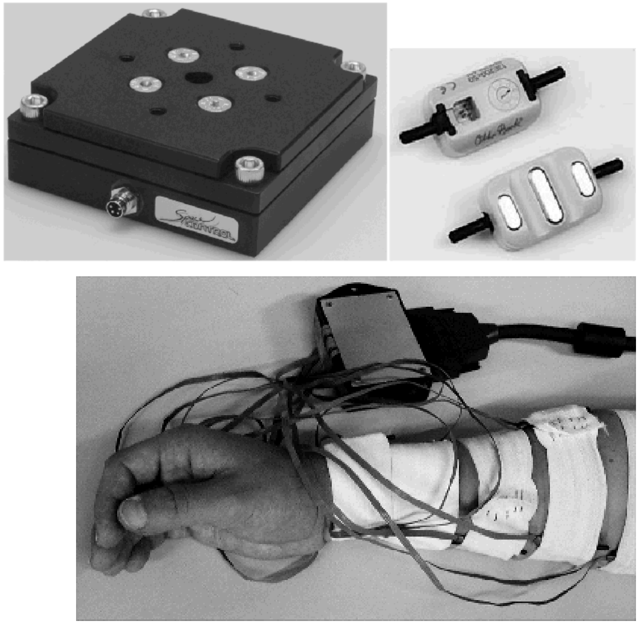

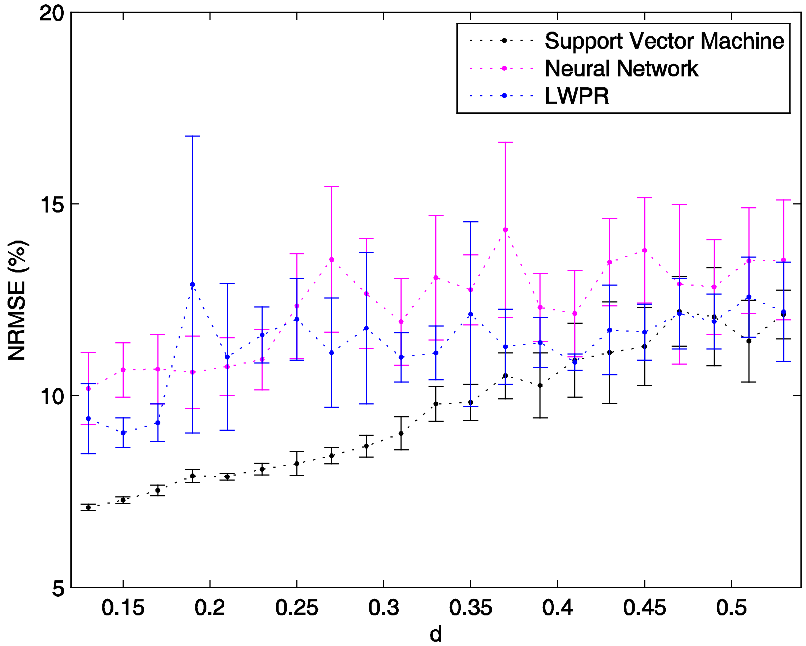

...and even more of them

- It can work even if you move around,

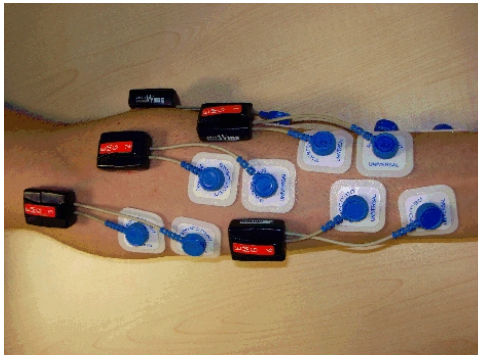

Castellini, C.; Fiorilla, A. E. & Sandini, G. *Multi-subject / Daily-Life Activity EMG-based control of mechanical hands* Journal of Neuroengineering and Rehabilitation, 2009, 6

- and even (!) on amputees.

Castellini, C.; Gruppioni, E.; Davalli, A. & Sandini, G. *Fine detection of grasp force and posture by amputees via surface electromyography* Journal of Physiology (Paris), 2009, 103, 255-262

Identifying the problem – laying out a plan

- All these experiments were theoretical.
  - „monolithic“ data gathering and training in the beginning
  - no corrections / update possible, later on
  - potentially endless data sets (space and time requirements)
  - classification highly unstable
- But most of all, how does the human stand in the loop?
  - **what should the subject do** to generate proper data?
  - what is the **optimal training strategy**?
  - **when is updating required?**
  - how do we **really** measure the performance?
- Todo list:
  1. multi-DOF regression control
  2. incrementality
  3. how to predict the subject‘s intent / how to help the subject use the system

Todo #1: multi-DOF regression control

The traditional approach

- Traditional two-electrodes sEMG control:
  - Enables proportional control over one DOF
  - For more DOFs, the subject needs a switching strategy
  - Thresholds, cocontraction, „impulse strategy“
- With the advent of TMR and multi-fingered hand prostheses [circa 2006]:
  - Traditional myocontrol even more insufficient
  - Need (simultaneous,proportional) control over several DOFs
  - Needs to be intuitive („natural“)
- The engineer‘s perspective:

control a mechanical artifact to its best extent:
send current values to each single motor.

- The community‘s solution: **classification of higher-density sEMG**

Todo #1: multi-DOF regression control

The traditional approach

- Classification has a number of drawbacks:
  - Usually monolithic
  - If online, not bounded in space and time
  - Enforces one class at a time, and no proportionality (sequential, discrete control)
  - Computationally expensive in the non-linear case

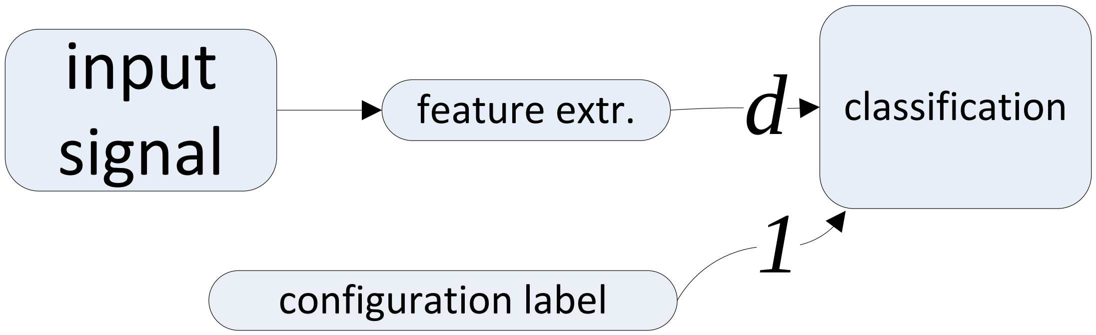

Todo #1: multi-DOF regression control

The traditional approach

- Classification has a number of drawbacks:
  - Usually monolithic
  - If online, not bounded in space and time
  - Enforces one class at a time, and no proportionality (sequential, discrete control)
  - Computationally expensive in the non-linear case

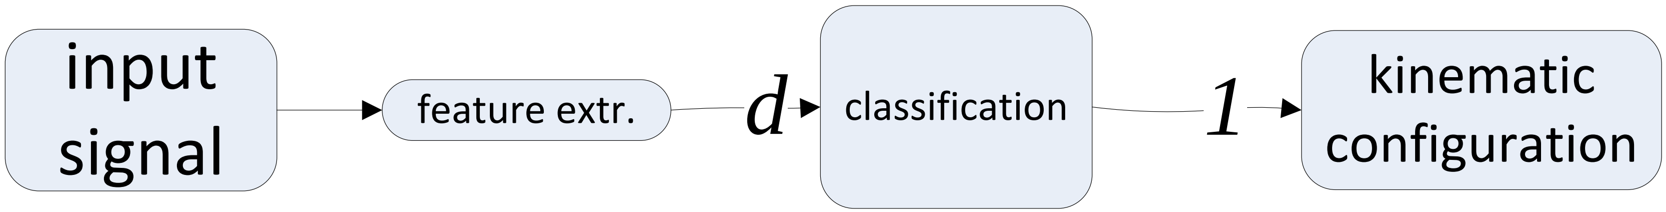

Todo #1: multi-DOF regression control

Simultaneous and proportional control

- Better solution: **simultaneous, proportional control** [Jiang et al., 2006]
  - Enforces simultaneous, *graded* activation of many DOFs:
    from $f: EMG \rightarrow \{0,1,2,3\}$ to $\boldsymbol{f}: \boldsymbol{EMG} \rightarrow \mathbb{R}^{\boldsymbol{m}}$
  - Mostly, only works in the linear case
  - Still monolithic, especially when non-linear

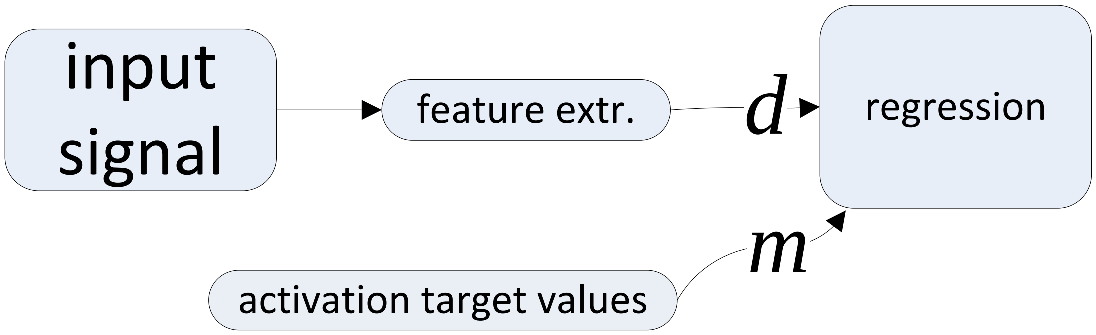

Todo #1: multi-DOF regression control

Simultaneous and proportional control

- Better solution: **simultaneous, proportional control** [Jiang et al., 2006]
  - Enforces simultaneous, *graded* activation of many DOFs: from $f: EMG \rightarrow \{0,1,2,3\}$ to $\boldsymbol{f}: \boldsymbol{EMG} \rightarrow \mathbb{R}^{\boldsymbol{m}}$
  - Mostly, only works in the linear case
  - Still monolithic, especially when non-linear

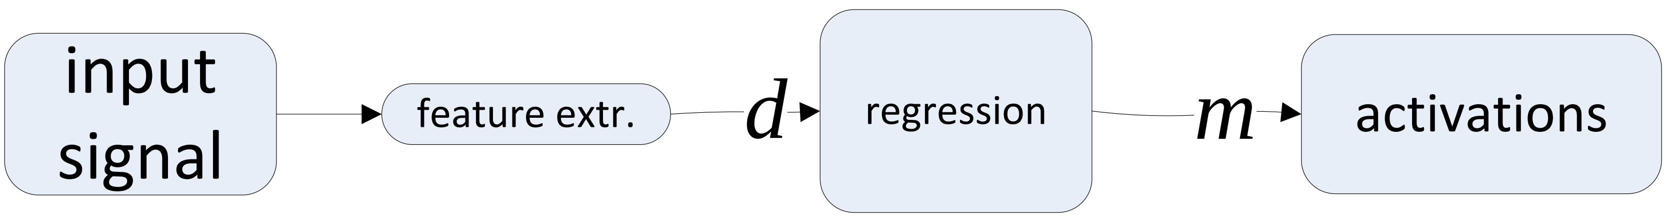

Todo #1: multi-DOF regression control

Simultaneous and proportional control

Castellini, C. & Kõiva, R. *Using surface electromyography to predict single finger forces* Proceedings of BioRob - IEEE International Conference on Biomedical Robotics and Biomechatronics, 2012, 1266-1272

- Our own experiments with finger forces:
  - Doable in principle. A Support Vector Regressor can predict single-finger forces using sEMG to a remarkable accuracy.
  - As the number of training samples diminishes, the accuracy degrades, *but not too much.*
- Still unsatisfactory:
  - SVM very slow in training (cubic in the number of collected samples)
  - Can become slow in prediction, too
  - Experiment still batch: monolithic, offline.

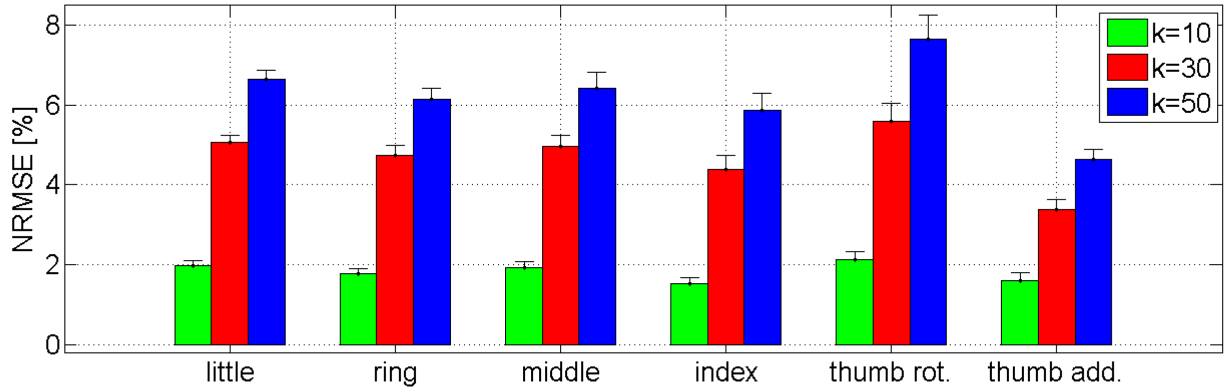

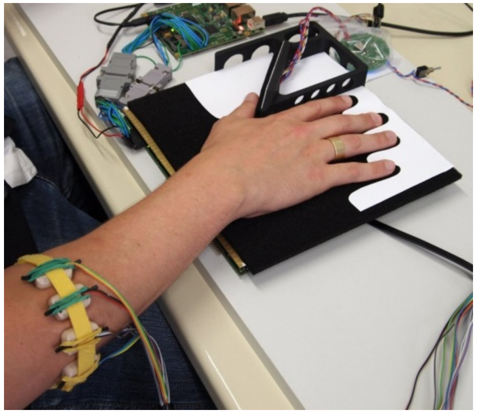

Todo #1: multi-DOF regression control

Intermission: ultrasound imaging instead of sEMG

- Let us now talk about a compeltely different HMI for hand control: *ultrasound images of the forearm.*
- The amount of information contained in such images is huge. (see movie #2)
- It turns out that local grey-level approximations are **linearly related**
  - to **metacarpo-phalangeal angles**
  - and to **finger forces:** $\mathbf{f} = \mathbf{w}^{\mathrm{T}}\mathbf{v}$
  - An interesting application of Occam‘s razor / Einstein‘s sentence / the KISS principle.
- movie #3 shows an application to virtual reality.

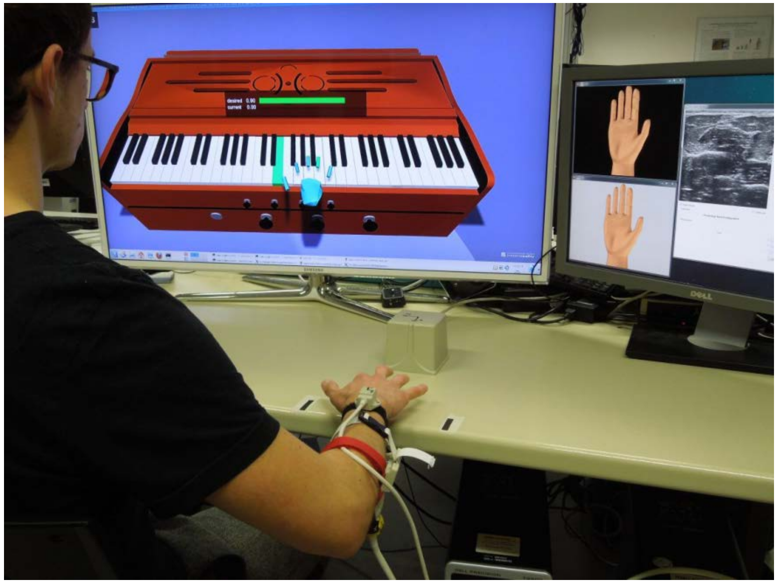

Castellini, C.; Passig, G. & Zarka, E. *Using ultrasound images of the forearm to predict finger positions* IEEE Transactions on Neural Systems and Rehabilitation Engineering, 2012, 20, 788-797

Sierra González, D. & Castellini, C. *A realistic implementation of ultrasound imaging as a human-machine interface for upper-limb amputees* Frontiers in Neurorobotics, 2013, 7

Castellini, C.; Hertkorn, K.; Sagardia, M.; Sierra González, D. & Nowak, M. *A virtual piano-playing environment for rehabilitation based upon ultrasound imaging* Proceedings of BioRob - IEEE International Conference on Biomedical Robotics and Biomechatronics, 2014, 548-554

Todo #1: multi-DOF regression control

Intermission: ultrasound imaging instead of sEMG

- The weight vector (aka the model) can easily be computed using least-squared regression or, slightly better, its regularised counterpart, **Ridge Regression** (RR):

$$\boldsymbol{w} = (X^T X + \lambda I)^{-1} X^T \boldsymbol{y}$$

- It‘s fast: you only need to invert a $d \times d$ matrix.
- It‘s theoretically optimal: derives from minimisation of MSE.
- It‘s computable: no optimisation required.
- No data normalisation required.

- And it can be made incremental (more about this later on).

DLR

Todo #1: multi-DOF regression control

Is „linear“ enough?

- Would such a method work for sEMG?
- No.
- The relationship between sEMG and finger forces is highly non-linear.

- So it seems that we are stuck...

- ...unless we can find an appropriate *kernel*.

- An example of such a kernel is *Random Fourier Features*.

Todo #1: multi-DOF regression control

Going non-linear

- A kernel is a map ϕ transforming the input space into something with more dimensions.
- You then try and solve the linear problem into this new space.
- How can it possibly work?

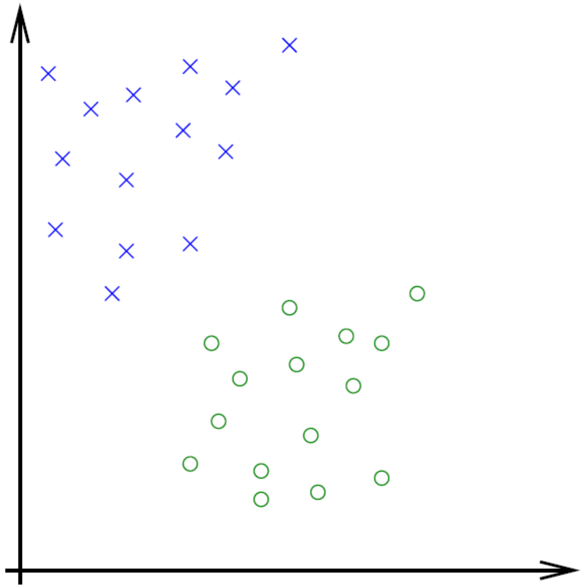

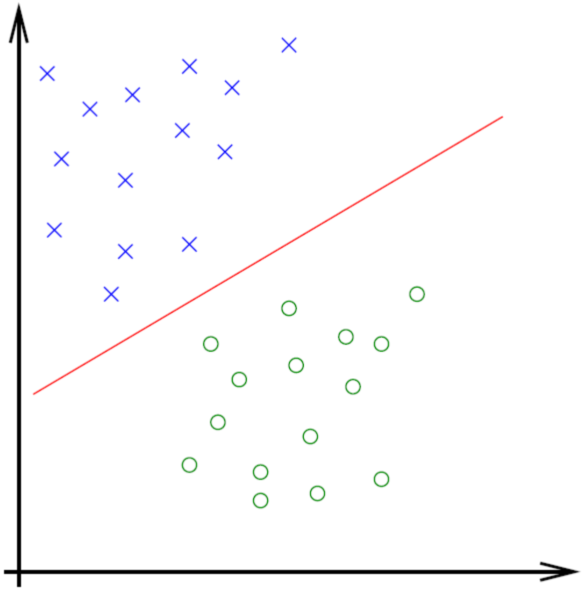

Todo #1: multi-DOF regression control

Going non-linear

- A kernel is a map ϕ transforming the input space into something with more dimensions.
- You then try and solve the linear problem into this new space.
- How can it possibly work?

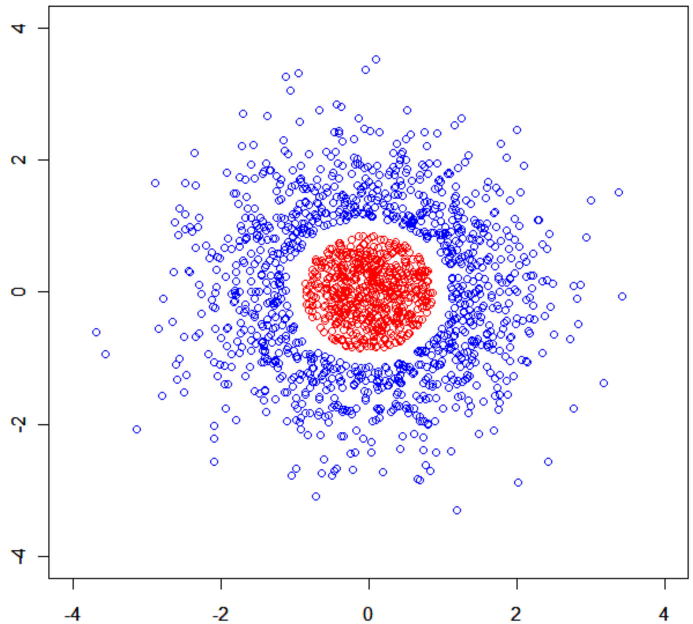

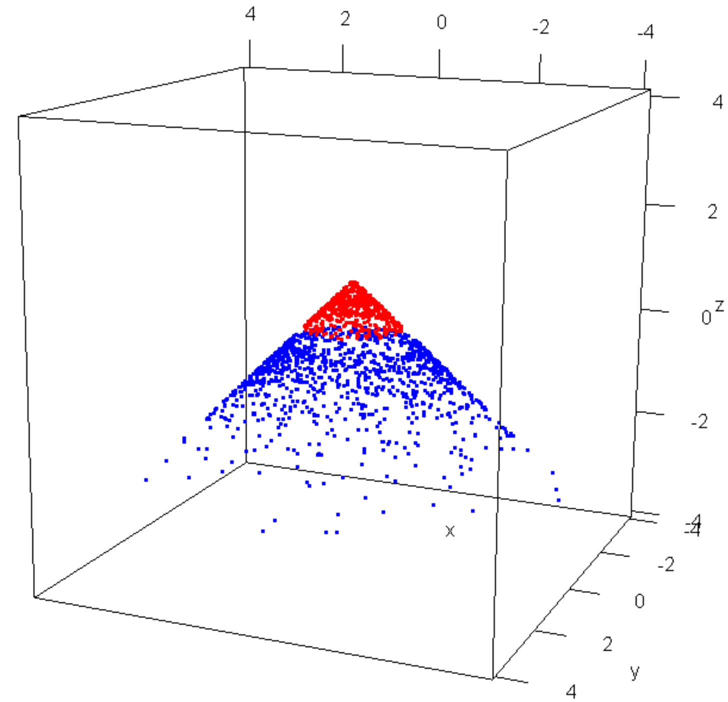

$$\phi(x_1, x_2) = \ (x_1, x_2, -\|x\|)$$

Todo #1: multi-DOF regression control

Going non-linear

Gijsberts, A.; Bohra, R.; González, D. S.; Werner, A.; Nowak, M.; Caputo, B.; Roa, M. & Castellini, C. *Stable myoelectric control of a hand prosthesis using non-linear incremental learning* Frontiers in Neurorobotics, 2014, 8

- Switch from $\mathbf{f} = \mathbf{w}^{\mathrm{T}}\mathbf{v}$ to

$$\mathbf{f} = \mathbf{w}^{\mathrm{T}}\boldsymbol{\phi}(\mathbf{e})$$

- ...and watch the magic happen:

$$\boldsymbol{w} = \left(\boldsymbol{\phi}(X)^T\boldsymbol{\phi}(X) + \lambda I\right)^{-1}\boldsymbol{\phi}(X)^T\boldsymbol{y}$$

- Random Fourier Features can be just „plugged into“ RR.
- All properties described before carry on.
- You now need to invert a $D \times D$ matrix. (see movie #4)

Todo #2: incrementality

Motivation

- Now for one general problem: intent detection can hardly be solved using monolithic learning.
  - biological signals change due to their own nature (think muscle fatigue and sEMG)
  - or due to the essentially unpredictable array of different situations
  - or because the subject requires a new action to be learned

- Besides being fast, the HMI must be
  - bounded in space and time (in the sense of computational complexity)
  - easily updated, possibly without long re-training times

- One simple solution: gather ever new data, keep the dataset bounded
  - Our example: SVM with subsampling strategies
  - highly arguable results
  - extremely heuristic, no theoretical foundation
  - still takes quite some time to retrain.

Kõiva, R.; Hilsenbeck, B. & Castellini, C. *Evaluating subsampling strategies for sEMG-based prediction of voluntary muscle contractions* Proceedings of ICORR - International Conference on Rehabilitation Robotics, 2013, 1-7

DLR

Todo #2: incrementality

Incremental Ridge Regression

- But Ridge Regression *can* actually be made incremental!
  - The trick is to start from a null predictor, then update it with every new (sample,target) pair.
  - Any rank-1 update method can be used (e.g., the Cholesky decomposition).

- We employ the *Sherman-Morrison formula:*
- recall that $\mathbf{f} = \mathbf{w}^{\mathrm{T}}\mathbf{v}$ and redefine

$$\boldsymbol{w} = (X^T X + \lambda I)^{-1} X^T \boldsymbol{y} \stackrel{\text{def}}{=} A\boldsymbol{b}$$

- Start with $A = \boldsymbol{I_d}$ and $\boldsymbol{b} = \boldsymbol{0}$, then use

$$A' = A - \frac{A\boldsymbol{x}'\boldsymbol{x}'^T A}{1 + \boldsymbol{x}'^T A \boldsymbol{x}'}, \qquad \boldsymbol{b}' = \boldsymbol{b} + \boldsymbol{x}' y'$$

to accommodate a new pair $(\boldsymbol{x}', y')$.

DLR

Todo #2: incrementality

Incremental Ridge Regression plus Random Fourier Features

- ...and as you might have guessed, in this case too RFFs can be plugged into RR.

- Recall that $\mathbf{f} = \mathbf{w}^{\mathrm{T}}\boldsymbol{\phi}(\mathbf{e})$; start with $A = \boldsymbol{I_D}$ and $\boldsymbol{b} = \mathbf{0}$, then use

$$A' = A - \frac{A\boldsymbol{\phi}\boldsymbol{x}'\boldsymbol{\phi}\boldsymbol{x}'^{T}A}{1 + \boldsymbol{\phi}\boldsymbol{x}'^{T}A\boldsymbol{\phi}\boldsymbol{x}'}, \qquad \boldsymbol{b}' = \boldsymbol{b} + \boldsymbol{\phi}\boldsymbol{x}'y'$$

- to accommodate a new pair $(\boldsymbol{x}', y')$. Result:
  - It‘s non linear!
  - Strictly bounded in space and time (only need to store A, which is $D \times D$)
  - No matrix inversion needed, at any time. *(A already is the inverse!)*
  - The update is done on-the fly: no additional time for training needed.

Todo #2: incrementality

Benefits

- Incrementality enables a new modality of training the machine:
    - the subject can tell us when something wrong is going on (*correcting* update needed)
    - the subject can tell us when something new is required (*augmenting* update needed)
    - the prediciton can run continually
    - the update can happen at any time

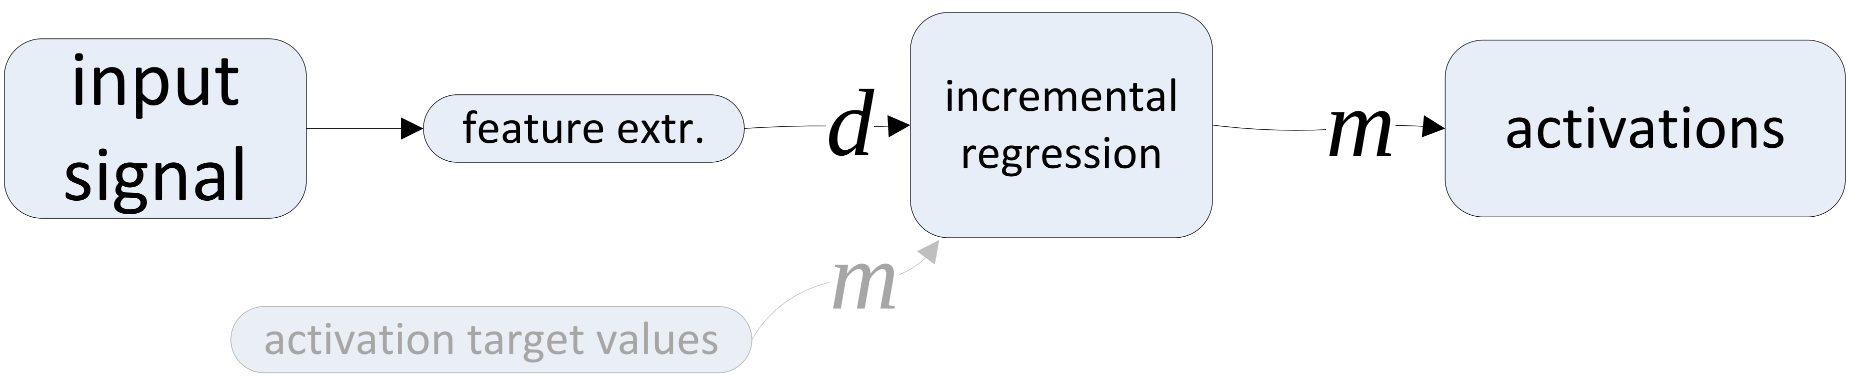

Todo #2: incrementality

Benefits

- Incrementality enables a new modality of training the machine:
  - the subject can tell us when something wrong is going on (*correcting* update needed)
  - the subject can tell us when something new is required (*augmenting* update needed)
  - the prediciton can run continually
  - the update can happen at any time

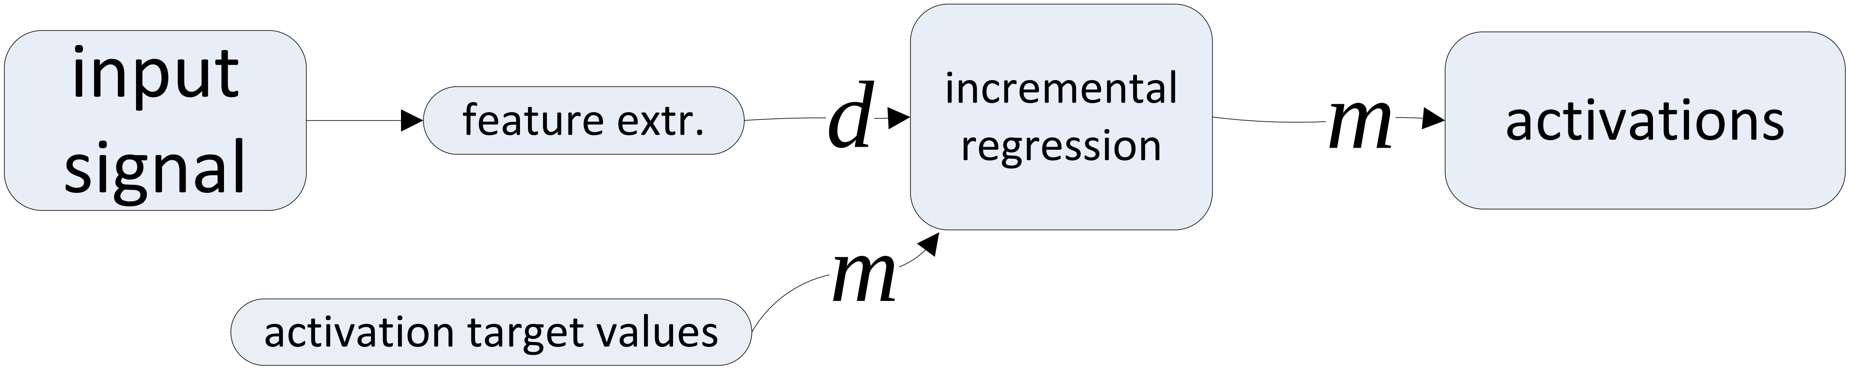

Todo #3: predicting the subject‘s intent

Natural control and amputees

- Natural control: predicting the subject‘s intent.
  - **Subject desires (enacts) action A,**
  - **prosthesis does A.**
  - In the case of s/p control this is basically the only chance
    (but remember Touch Bionics‘s smartphone app and GripChips)

- Consider amputated subjects: cannot use sensors for ground truth.
  - No cyberglove.
  - No pressure-based force sensors.
  - No optical / magnetic tracking possible.

- So
  - how do we know that the subject is doing what he‘s supposed to do?
  - in case he does, how do we gather ground truth?
  - anyway, *how does he know what he‘s doing?*

Todo #3: predicting the subject‘s intent

Goal-directed stimuli

- One answer to all questions:

use on-off goal-directed stimuli for training,

use concrete tasks for testing.

- Training with goal-directed stimuli:
  - show the subject what he‘s supposed to do (serious games, prosthetic training, ...)
  - wait for his signals to settle
  - gather them and associate them with maximal activation.
  - (already exploited: bilateral-mirror training, imitation learning, using visual stimuli, ...)

- Testing on concrete tasks:
  - need an online control system (we need it in general, don‘t we?)
  - incrementality takes care of errors and instability
  - use standard tests and standard measures of outcome
  - make machine learning matter! [Wagstaff 2012, ICML]

Todo #3: predicting the subject‘s intent

Goal-directed stimuli

- And now, couple this idea with incremental learning!

- The setup gets considerably simpler
  - no ground truth sensors required
  - lighter, cheaper, easier to program and maintain

- The subject‘s experience gets considerably simpler
  - no „blind“ initial training phase
  - ability to update at any time and in any situation
  - possiblity to *interact* with the system!

see movie #5

Gijsberts, A.; Bohra, R.; González, D. S.; Werner, A.; Nowak, M.; Caputo, B.; Roa, M. & Castellini, C. *Stable myoelectric control of a hand prosthesis using non-linear incremental learning* Frontiers in Neurorobotics, 2014, 8

Todo #3: predicting the subject‘s intent

Goal-directed stimuli

- Of course, the control system needs to generalise
  - from minimal / maximal activation
  - to intermediate activation values

- and regression in general works like that!
  - ...although in the non-linear case the behaviour can be surprising,
  - regression draws a (simple) curve for you, from A to B
  - that is, it interpolates the behaviour of sEMG between minimal and maximal activation.

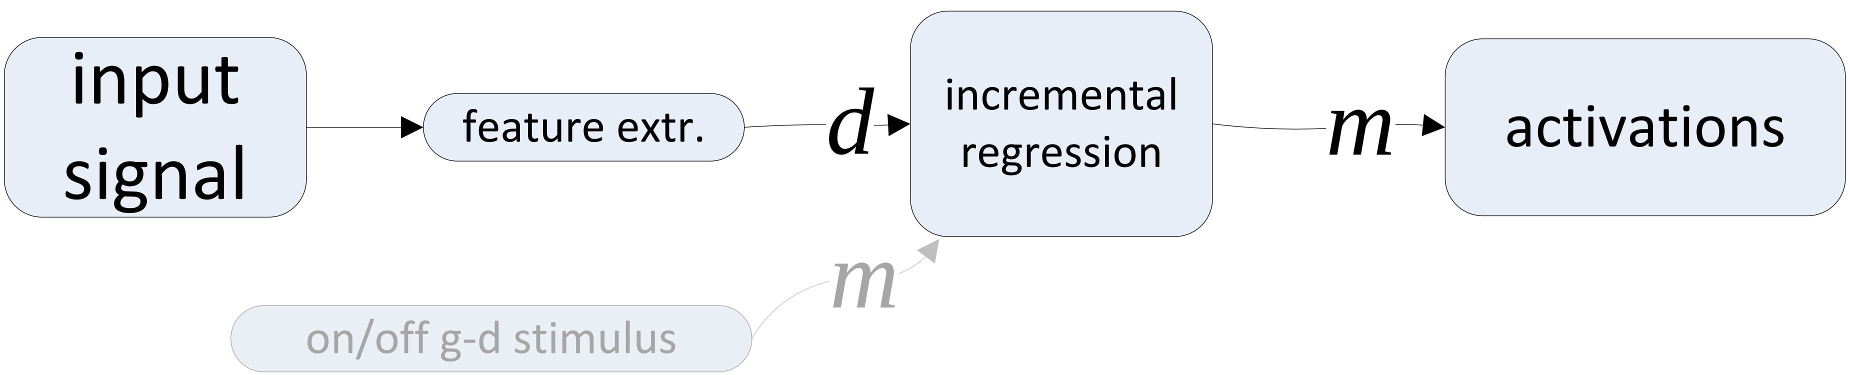

Todo #3: predicting the subject‘s intent

Goal-directed stimuli

- Of course, the control system needs to generalise
  - from minimal / maximal activation
  - to intermediate activation values

- and regression in general works like that!
  - ...although in the non-linear case the behaviour can be surprising,
  - regression draws a (simple) curve for you, from A to B
  - that is, it interpolates the behaviour of sEMG between minimal and maximal activation.

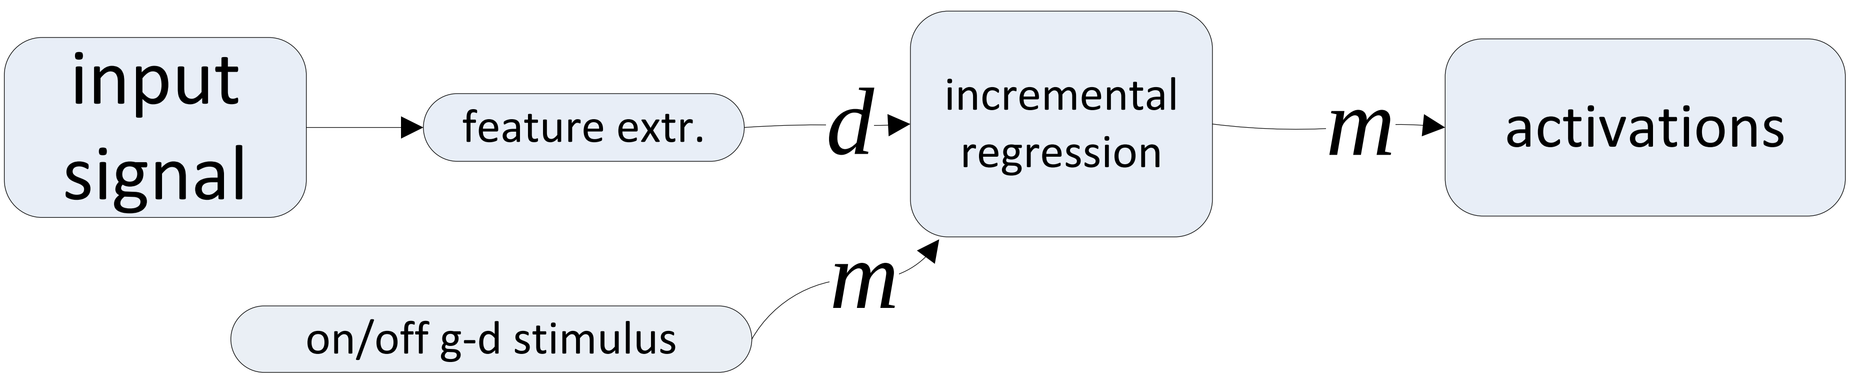

Todo #3: predicting the subject‘s intent

Goal-directed stimuli

- Of course, the control system needs to generalise
  - from minimal / maximal activation
  - to intermediate activation values
- and regression in general works like that!
  - ...although in the non-linear case the behaviour can be surprising,
  - regression draws a (simple) curve for you, from A to B
  - that is, it interpolates the behaviour of sEMG between minimal and maximal activation.
- And even if it does not work *exactly as expected,* the subject will „fill the holes in“.
- This is where the subject starts adapting to the system, as well as the other way around.

Sierra González, D. & Castellini, C. *A realistic implementation of ultrasound imaging as a human-machine interface for upper-limb amputees* Frontiers in Neurorobotics, 2013, 7

see movie #6

Open questions

- How big is a big enough machine?
  - Any online system must necessarily be bounded in space;
  - what is its correct VC-dimension?
  - will its capacity run out at some point?

- How to deal with multi-DOF acitvations?
  - stay simple and ignore the issue
  - signal decomposition (e.g., NMF)
  - enhance the test set with a guess

Castellini, C. & Nowak, M. *EMG-based prediction of multi-DOF activations using single-DOF training: a preliminary result* Proceedings of MEC - Myoelectric Control Symposium, 2014, 45-49

- How to help the subject produce good signals for us?
  - serious games
  - learning new synergies
  - goal-directed stimuli, again
  - prosthetic embodiment / sensory feedback

Conclusions

- Incremental learning should become the standard:
  - shifts the focus from the machine to the interaction,
  - improves stability of the control,
  - enhances the subject‘s experience.
- A good companion to simultaneousness/proportionality. My ideal myocontrol is
  - Simultaneous
  - Proportional
  - Incremental / Interactive
  - Natural
- ...it SPINs (es spinnt...)
- Needless to say, what you just heard are strictly the speaker‘s opinions.

Thank you.