Homeokinesis

A new principle to back up evolution with learning

Malte Schwarzkopf
University of Cambridge

UNIVERSITY OF CAMBRIDGE

Valadora, Alto Adige, Italy
September 18, 2008

This presentation does not contain any complex Maths

Not suitable for consumption by Pure Mathematicians

Overall motto: Non goal-directed learning

- Relatively new approach – central paper published in 1999, recently received some more attention
- Non goal-directed learning
- Models biological evolution (to some degree)
- Allows for self-organised AI agents

... but first, let's look at a few basic things.

A few basic things about AI

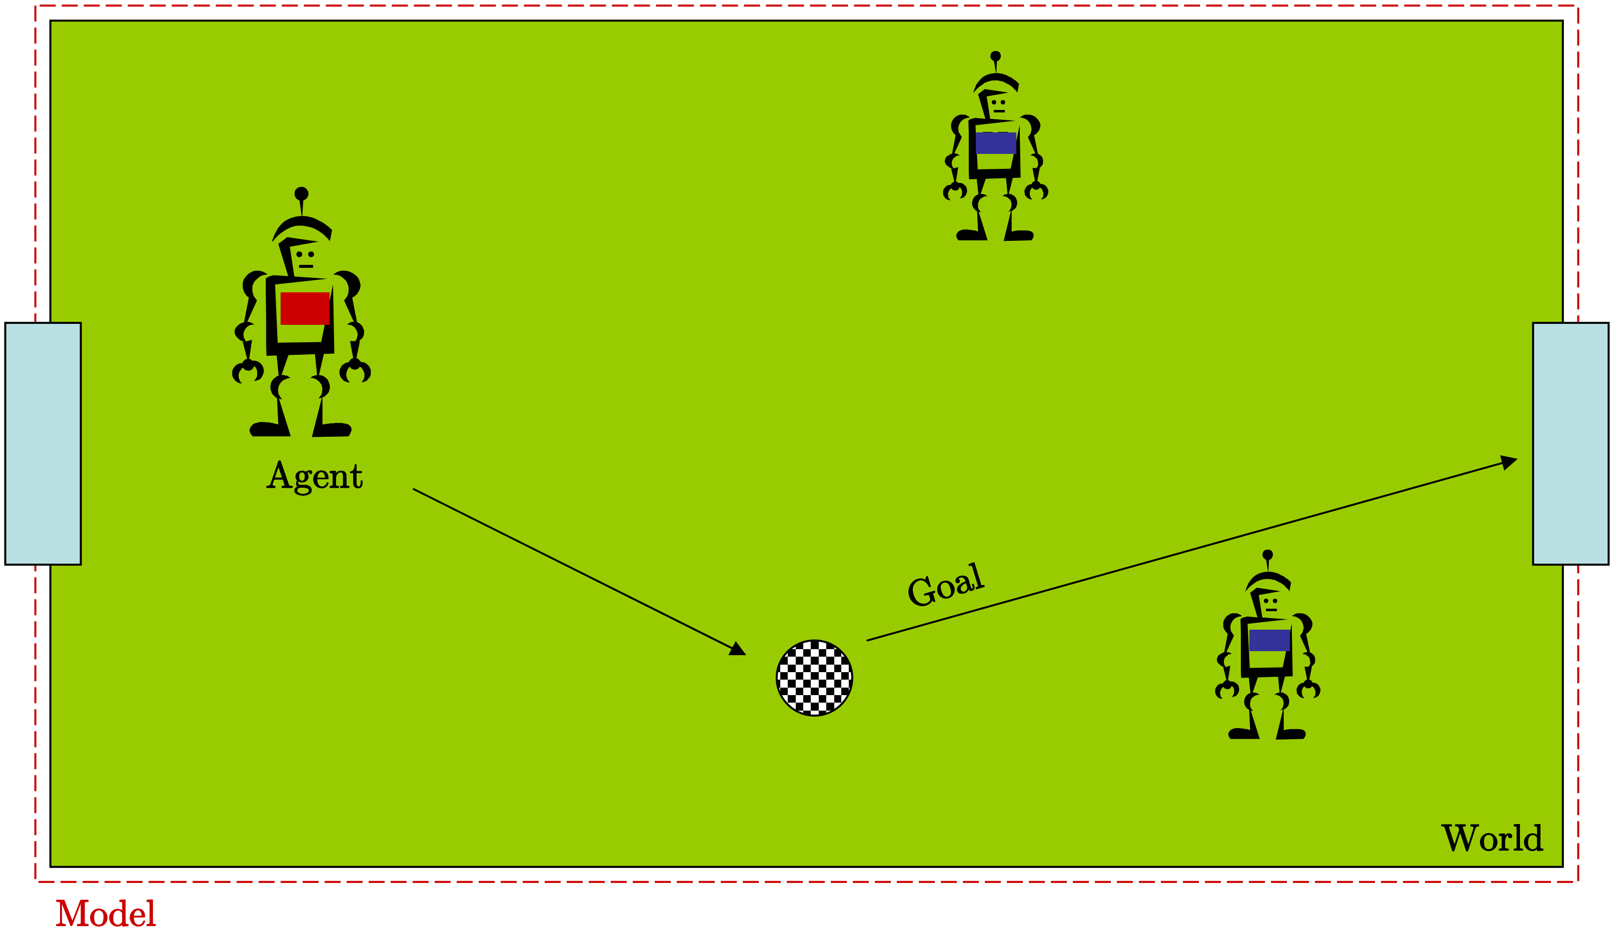

A few words on learning...

- Learning is a topic in many sciences...
  - Psychology...
  - Biology...
  - Mathematics...
  - Physics...
  - Computer Science...

... and there are several different concepts:

- Reinforcement Learning
- Supervised Learning
- Autonomous Learning
- Trial and Error...

- So why are we so keen on learning?
  - Artificial Evolution
  - Autonomous agents
  - Self-organisation
  - Model nature – „let evolution do the work“
  - *In vivo* evolutionary robotics

Different types of learning

- Supervised Learning
  - Use a **training set**: Map example inputs to desired outputs
  - **Regression** or **classification**
  - Learner can predict values for new, unknown input objects
  - cf. **Concept Learning** in humans

  - Distinguishing feature: Manual, human-prepared training examples
  - Issues: Manual training required, generalisation difficult

Different types of learning

- Reinforcement Learning
  - Agent learns a **policy** of how to act given an observation
  - Feedback provided by the environment (world) in form of **rewards** and **punishments**
  - Inspired by Psychology (animal learning)
  - Agent tries to maximise the reward and avoid punishments

  - Distinguishing feature: Action/Reward mapping
  - Issues: Little generalisation, large sample space required

The complexity barrier

- So far, only fairly simple successful examples, e.g.
  - Wall following in a maze
  - Picking up objects
  - In simulation only (!)
- The real world is an extremely complex system
- Central issue:

We require **specific learning goals.**

... so why is this?

The necessity of specific learning goals

- For a self-organised, evolving robot, the **fitness function** is critical to its success
- Supervised learning requires a teacher and explicit, goal-directed input
- Reinforcement learning requires finding the right distribution of rewards

- i.e. we require **explicit formulation of goals** to drive evolution.

We want something better.

Getting to a better model of learning

- Agent should learn **general principles**
- Similar to biological evolution: Learn behaviours that make sense in the environment ... or die.
- Give the agent an **adaptive model** of its own behaviour
- Adapt the controller so that the model understands its own behaviour

Getting to a better model of learning

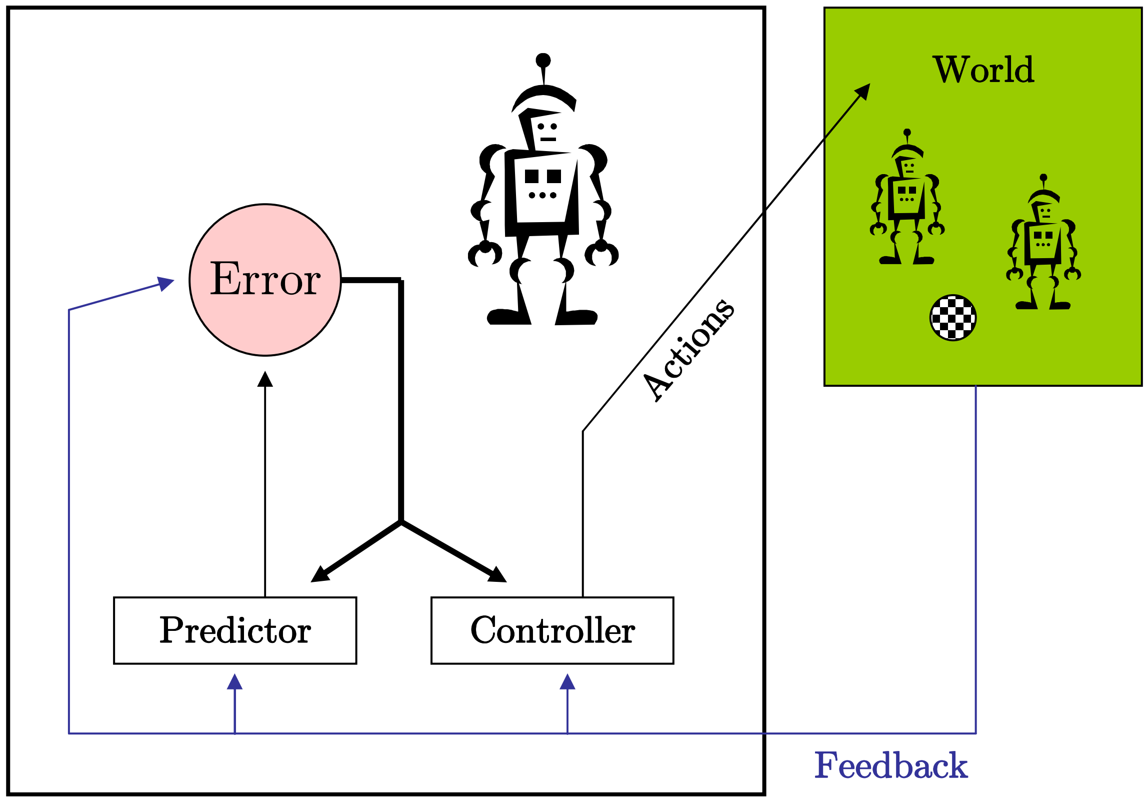

Example 1:
Linear control of a stochastic dynamic system

Generic example (1)

- Consider a mass-less particle with friction under the influence of a harmonic force plus noise:

$$x'(t) = \underbrace{F(x)}_{\text{Force}} + \underbrace{\xi(t)}_{\text{Noise}}$$

$\xi(t)$ is white Gaussian noise with mean $\langle \xi(t) \rangle = 0$ and correlation function (N.B.: $\delta(t-t') = 0$ if $t \neq t'$, $\delta(0) = 1$)

Dirac delta

$$\langle \xi(t)\xi(t') \rangle = \sigma^2 \delta(t-t')$$

- Model understands the **systematic** behaviour of the system ... but not the noise.

Generic example (2)

- The model is a **predictor** for the value at some future point in time, evaluating

$$x_t^{(pred)}(t+\tau)$$

- It is of linear complexity, hence

$$x_t^{(pred)}(t+\tau) = ax(t)$$

- Using a mean squared error function, we get

$$E = \frac{1}{2}\left(x_t^{(pred)}(t+\tau) - x(t+\tau)\right)^2$$

Generic example (3)

- Assume that $a = e^{-\kappa\tau}$ (i.e. model adapts quicker than controller)
- Average over the noise:

$$E = \sigma^2 \frac{1 - e^{-2\kappa\tau}}{4\kappa\tau}$$

- This decays monotonically

... so we can use learning by **gradient descent.**

Example 2:

Non-linear emergence of control

Non-linear emergence of control (1)

- Example: Braitenberg vehicle

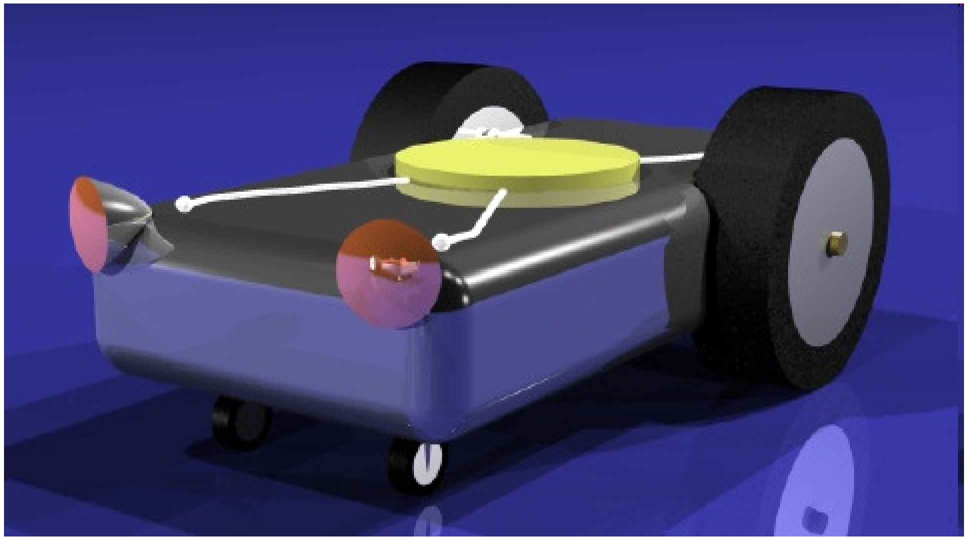

Image source:
http://people.cs.uchicago.edu/~wiseman/vehicles/vehicle.jpg

- 2 light sensors, 2 wheels

Non-linear emergence of control (1)

- Assume coupling of strength ω on right wheel
- Move light source along x axis
- Vehicle starts off at some y position
- Vehicle moves at constant distance, s and speed, $c = 1$

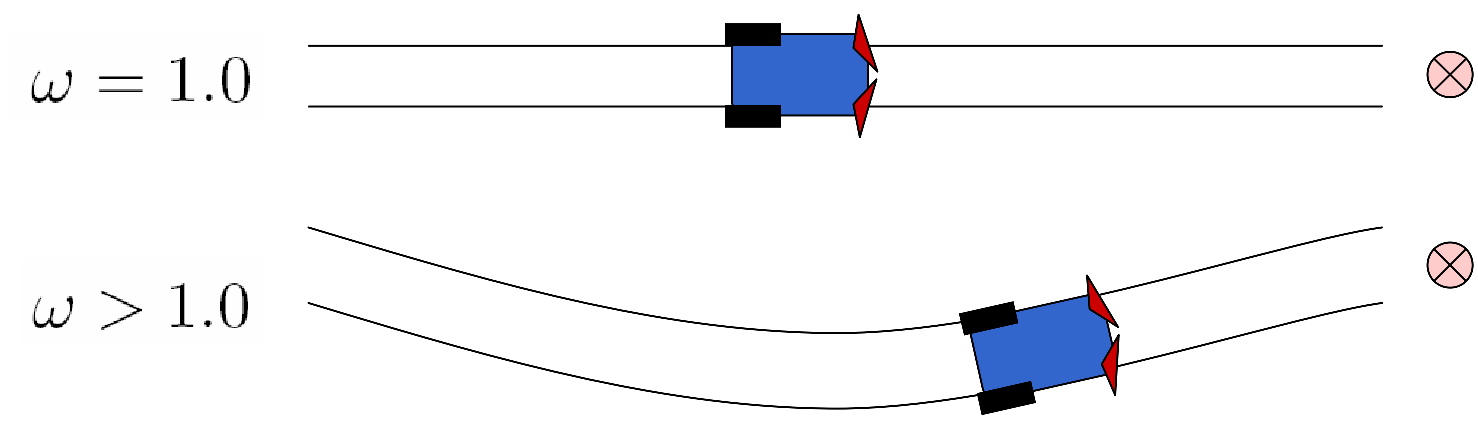

Non-linear emergence of control (2)

- Angular velocity ϕ' depends on sensor readings and coupling strength
- Equations of motion:

$$\boxed{\begin{aligned} x' &= \cos\phi \\ y' &= \sin\phi \\ \phi' &= s(\omega s_l(y,\phi) - s_r(y,\phi)) \end{aligned}}$$

- Sensor readings:

$$s_{l,r} \approx \frac{1}{s^2}\left(1 + \frac{2R}{s}\cos(\alpha + \phi + \gamma)\right)$$

Non-linear emergence of control (3)

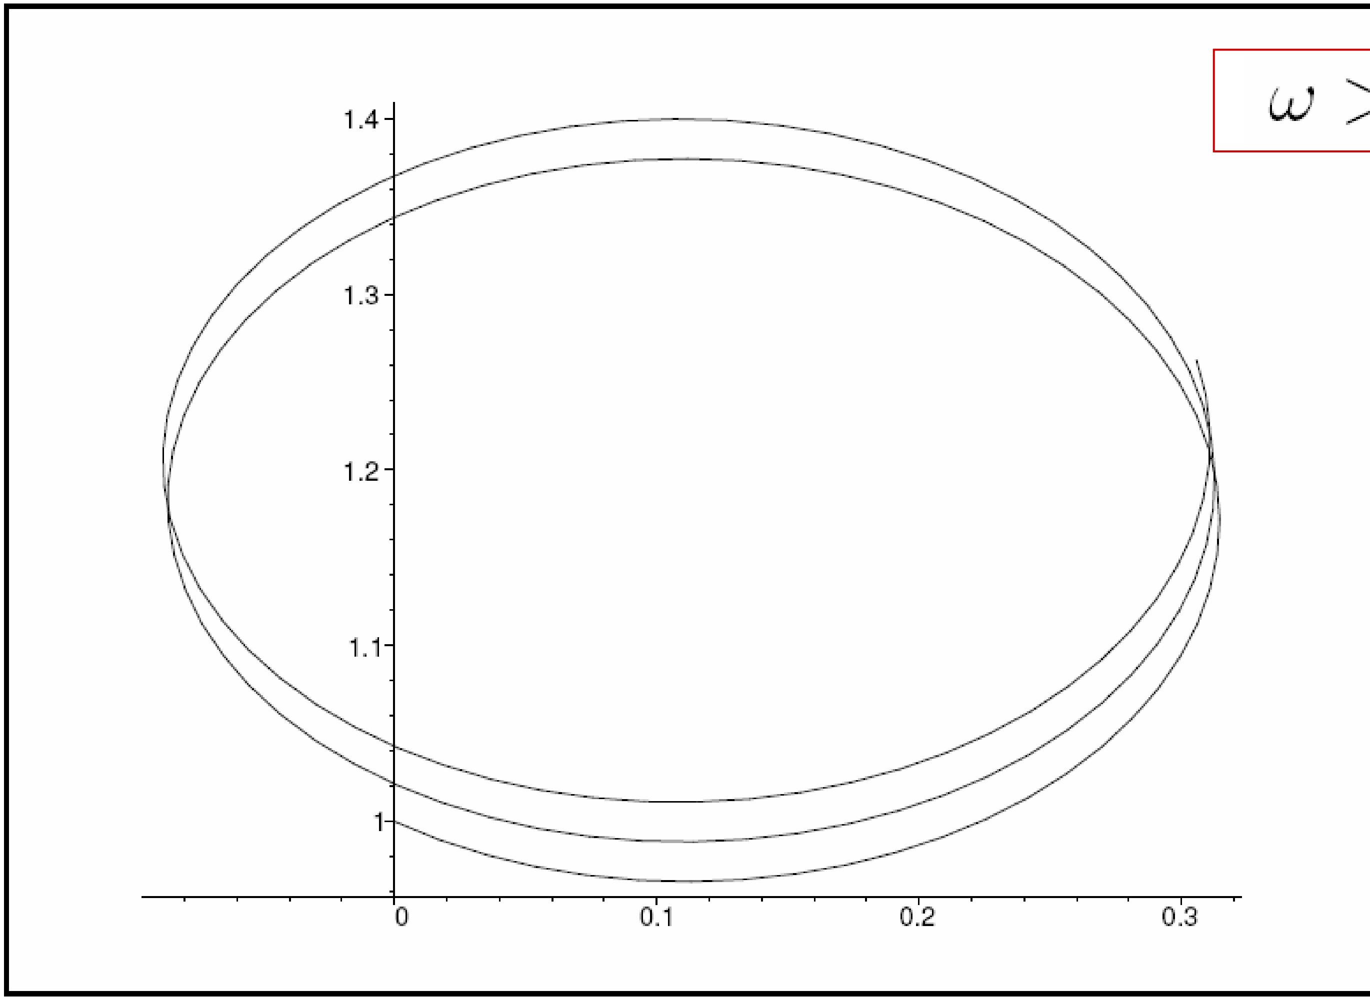

$$\omega > 1.005$$

No homeokinesis

Image source:
`http://www.informatik.uni-leipzig.de/~der/Veroeff/wienfin3.ps`

Non-linear emergence of control (4)

- N.B.: The robot knows **nothing** about **itself**
- Want to detect and counterbalance coupling effect: „*Learn about your weaknesses and compensate*“
- Evolution does not create perfect individuals!

- We need to have some way of intergrating learning with the system:
  - Use prediction error as learning signal

Non-linear emergence of control (5)

- We want to work out ω (for our model)
- Use trivial predictor: Always propose $s'_{l,r} = 0$
- Prediction error (E) is change of sensor value in τ
- Minimise E by gradient descent ...

... complement equations of motion with

$$\omega' = -\frac{1}{\theta}\frac{\partial}{\partial\omega}E(y, \phi; \omega)$$

where θ is a time constant

Non-linear emergence of control (6)

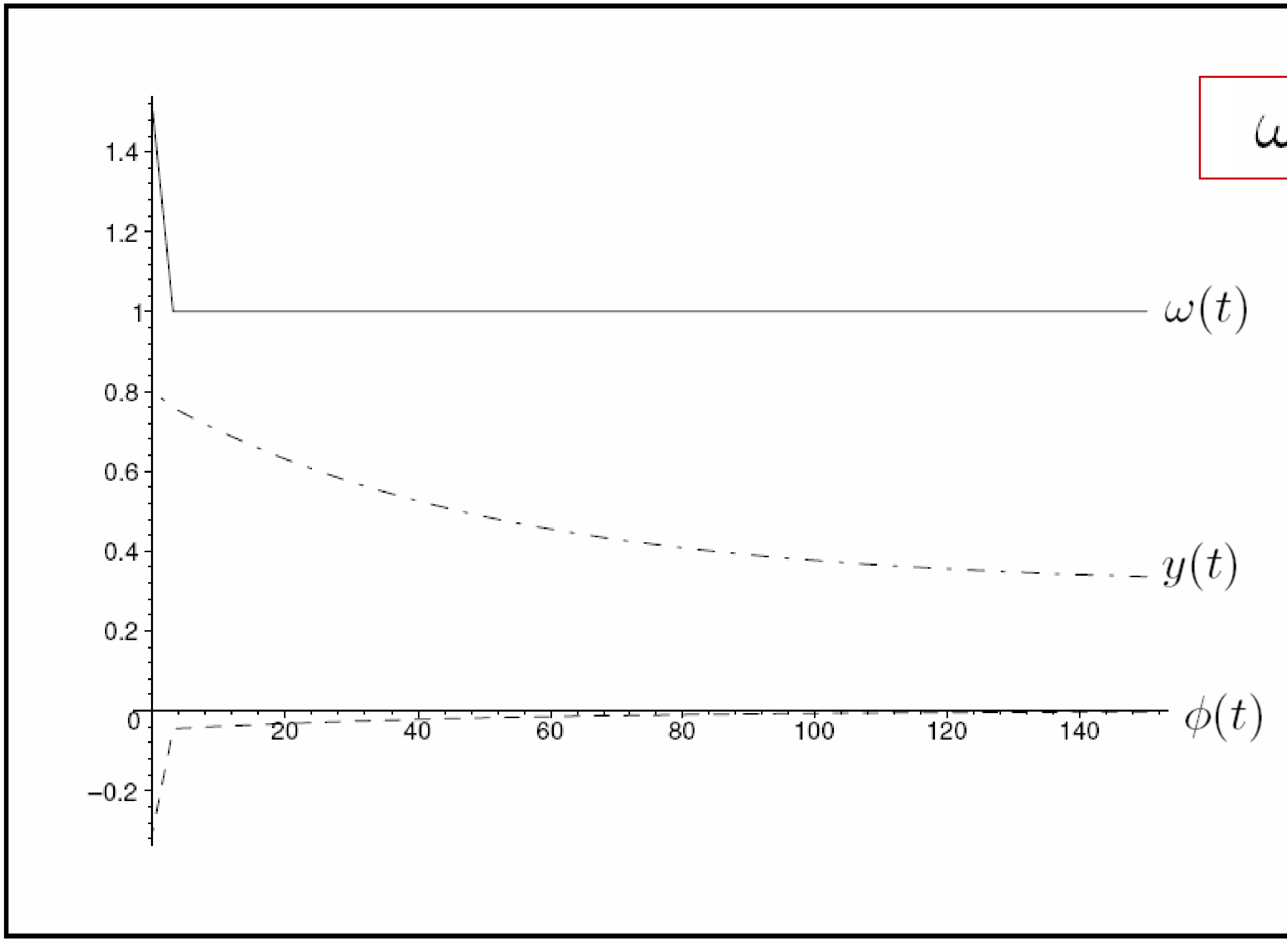

$\omega > 1.005$

Homeokinetic control

Image (modified) from:
`http://www.informatik.uni-leipzig.de/~der/Veroeff/wienfin3.ps`

Non-linear emergence of control (7)

- Stable light-following behaviour
- Adapts to maintain **sensor values**
  (cf. physiologial values in homeostasis)

The homeokinetic principle

The homeokinetic principle

- cf. **Homeostasis**: Try to maintain physiological values at a certain level
- Here: Try to maintain a **smooth controlled behaviour**, i.e. keep the agent in a **kinetic state**
- **Mode** learnt decided by predictor complexity and environmental conditions
- Learning signal derived from misfit between real behaviour and model prediction, i.e. how well the agent understands the world

... and the complexity barrier?

- Complex nonlinear problems: Task decomposition
  - Indep. behaving functional units
  - „Expert“ systems
- The components are independent and self-organised

Criticism

- Isn‘t the error function just a weaker form of a negative fitness function?
- Limits imposed on predictor by error function
- Complexity of component systems

“I would be interested to know whether the control theorists have looked at this in a more mathematically acceptable fashion - it is not a good idea to expect to do anything interesting with nonlinear dynamic systems without biting the foundational bullet.“

Applications & Demos

- Self-organised systems
  - Emerging motion patterns in humanoids
  - Der et al., 1999-2008
- Emergent cooperativity in a chain of mobile robots
  - Emerging cooperative behaviour of chained robots
  - Der et al., 2008

Thank you!