Evolutionary Robotics Tutorial

N. Bredeche,
S. Doncieux

Sorbonne Université, CNRS, Institute of Intelligent Systems and Robotics (ISIR), F-75005 Paris, France

nicolas.bredeche@sorbonne-universite.fr

stephane.doncieux@sorbonne-universite.fr

J.-B. Mouret

Inria
CNRS, Université de Lorraine, Loria, UMR 7503
Vandœuvre-lès-Nancy, F-54500, France

jean-baptiste.mouret@inria.fr

http://pages.isir.upmc.fr/~bredeche/evorobots_tutorial/

ACM ISBN 978-1-4503-6748-6/19/07. https://doi.org/10.1145/3319619.3323391

1

Instructors

- **Nicolas Bredeche**
  - Professor at Sorbonne Université, Paris, France
  - Researcher in the AMAC team of the Institute of Intelligent Systems and Robotics (ISIR)
  - ANR MSR 2019-2022 (local coordinator), EU FP7 Symbrion (member), EU H2020 DREAM (member)
- **Stéphane Doncieux**
  - Professor at Sorbonne Université, Paris, France
  - Deputy director of the Institute of Intelligent Systems and Robotics (ISIR)
  - Coordinator of the 2015-2018 H2020 FET proactive project 'DREAM'
- **Jean-Baptiste Mouret**
  - Research Scientist at Inria - Nancy Grand-Est, France
  - Previously assistant professor at UPMC
  - PI of the ERC project 'ResiBots'

« Evolutionary Robotics aims to apply evolutionary computation techniques to evolve the overall design or controllers, or both, for real and simulated autonomous robots »

Patricia A. Vargas, Ezequiel A. Di Paolo, Inman Harvey and Phil Husbands, 2014, **The Horizons of Evolutionary Robotics**, MIT Press

Motivations: robotics

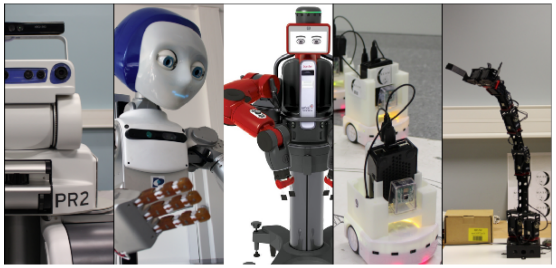

- Building robots with *embodied intelligence*
- Learning with state-of-the-art black-box optimization tools

Pfeifer, R., & Bongard, J. (2006). **How the body shapes the way we think: a new view of intelligence**. MIT press.

Stulp, F. and Sigaud, O. (2013). **Robot Skill Learning: From Reinforcement Learning to Evolution Strategies.** Paladyn Journal of Behavioral Robotics. Vol 4 No 1 Pages 49-61.

Motivations: biology

- ER as a model:
  - modeling evolutionary dynamics, in particular of groups
  - studying the emergence of features
- ER as a tool: optimization and analysis of computational models

Long, J. (2012). **Darwin's Devices: What Evolving Robots Can Teach us about the History of Life and the Future of Technology**. Basic Books.

Liénard, J. and Girard, B. (2014). **A Biologically Constrained Model of the Whole Basal Ganglia Addressing the Paradoxes of Connections and Selection**. Journal of Computational Neuroscience. Vol 36 No 3 Pages 445--468.

Evolutionary Robotics

main principles

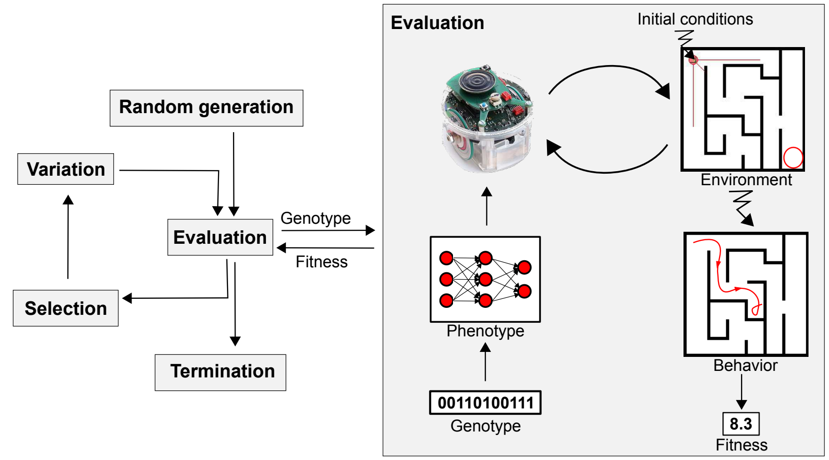

Doncieux S, Bredeche N, Mouret J-B and Eiben AE (2015) **Evolutionary robotics: what, why, and where to**. Front. Robot. AI 2:4. doi: 10.3389/frobt.2015.00004

Main features of Evolutionary Robotics

Selective pressure	priority to task resolution	or	task resolution secondary (or absent)
Focus	control	and	morphology
Implementation	simulation	or	real world
Space	centralized	or	distributed
Time	off-line	or	on-line

Overview

Selective pressure	priority to task resolution	or	task resolution secondary (or absent)
Focus	control	and	morphology
Implementation	simulation	or	real world
Space	centralized	or	distributed
Time	off-line	or	on-line

Part I: Fitness function and influence of selection pressure: *What do you need to know about evaluation and selection to make an ER experiment successful ?*

Overview

Selective pressure	priority to task resolution	or	task resolution secondary (or absent)
Focus	control	and	morphology
Implementation	**simulation**	**or**	**real world**
Space	centralized	or	distributed
Time	off-line	or	on-line

Part II: Evolution for physical robots and the reality gap
How to make it work on real robots ?

Overview

Selective pressure	priority to task resolution	or	task resolution secondary (or absent)
Focus	control	and	morphology
Implementation	simulation	or	real world
Space	**centralized**	**or**	**distributed**
Time	off-line	or	on-line

Part III: Embodied evolution and collective robotics systems
Evolution without a fitness for the design of distributed robotics systems and for modeling evolution of group dynamics.

Fitness function and influence of selection pressure

S. Doncieux

Example 1: obstacle avoidance

- Fitness: $\frac{1}{nb_{coll}+1}$

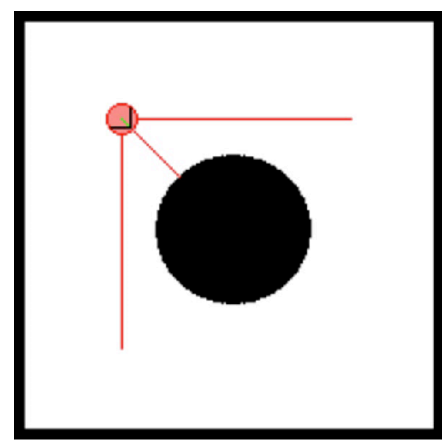

https://github.com/doncieux/navigation

Problem !

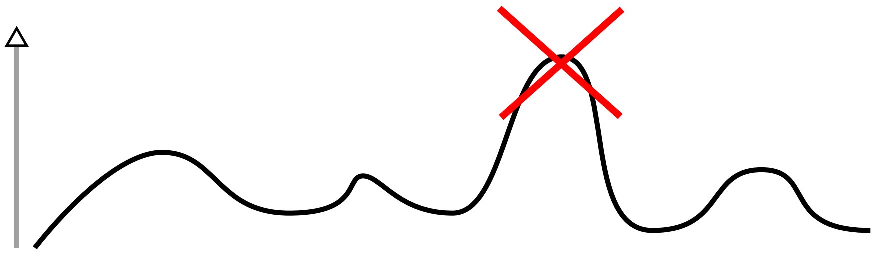

Example 1: obstacle avoidance

- How to deal with it ?
  - Change fitness: $\frac{1}{nb_{coll}+1} * \bar{v}$
  - Make the robot move by default
  - …

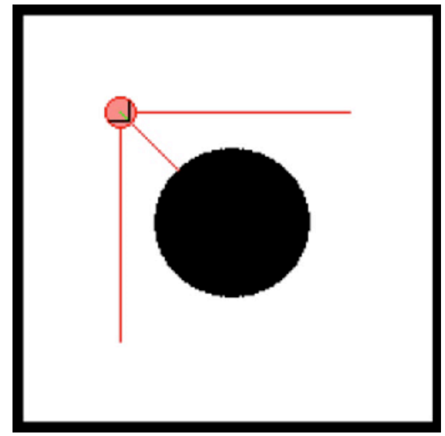

https://github.com/doncieux/navigation

Example 2: Collect ball experiment

Fitness= nb_{ball}

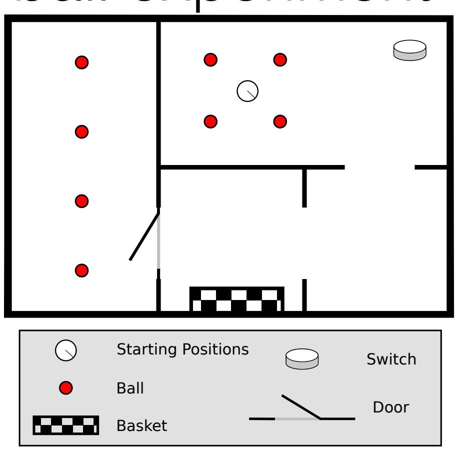

https://github.com/doncieux/collectball

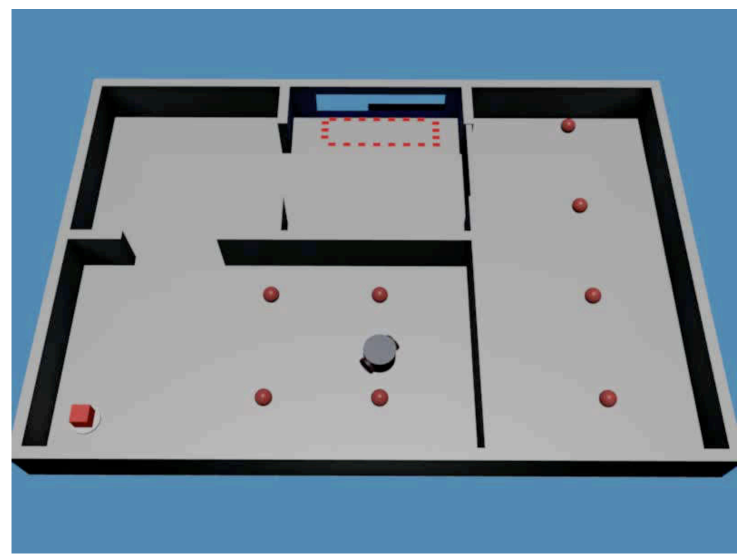

https://github.com/doncieux/collectball

Problem !

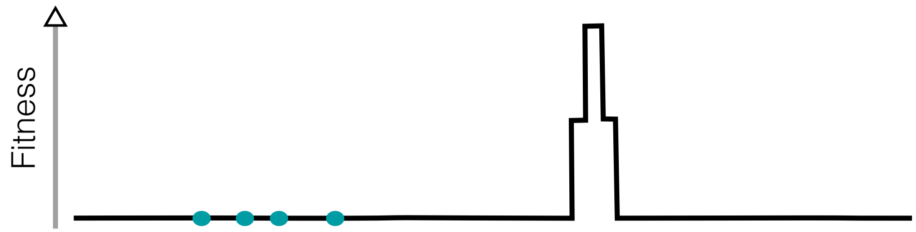

Example 2: Collect ball experiment

- How to deal with it ?
  - Decompose the problem
  - Add fitness terms
  - Enhance exploration

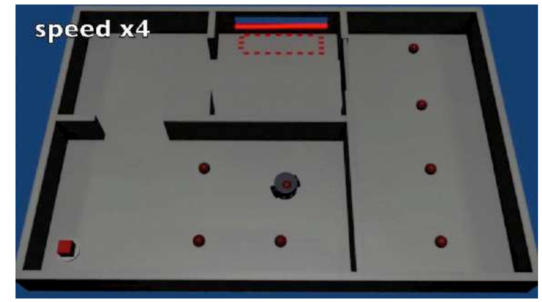

https://github.com/doncieux/collectball

The challenges of selective pressures

Goal of the evolutionary process:

Generating behaviors that solve the ta

The selective pressures must:

1. Define th

towards it

Can we deal with issues in goal definition and exploration in a task agnostic manner?

Minimizing a priori knowledge about how to solve the task

How is fitness evaluated ?

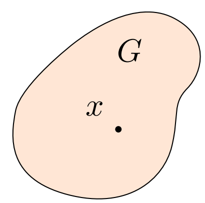

Genotype:

- vector of parameters
- neural network
- …

How is fitness evaluated ?

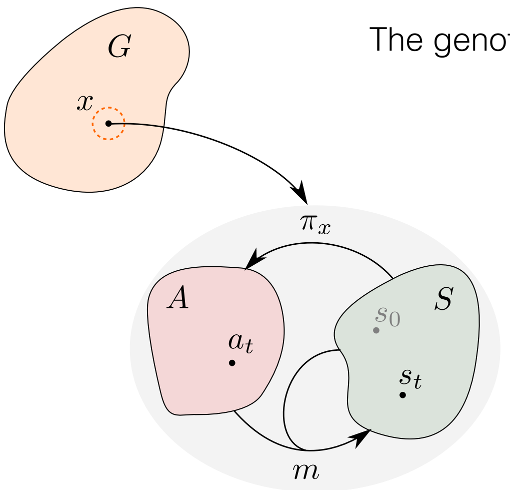

The genotype describes a policy:

$$\pi_x : S \to A$$

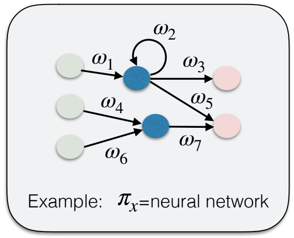

S: state space
A: action space
m: transition function

How is fitness evaluated ?

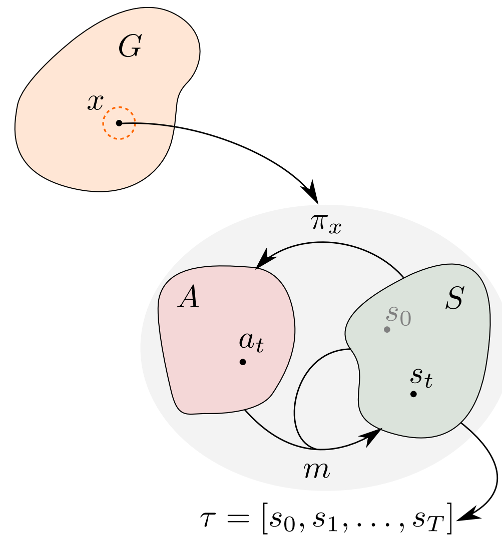

$$fit(x) = f(\tau)$$

- The fitness depends on the genotype x and on the fitness function f()
- *… but also on:*
  - s_0
  - τ
  - T
  - *environment* e

Beyond black-box optimization

To solve the challenges, the selective process can take into account:

$$s_0 \quad \tau \quad e \quad T$$

Two challenges, two kinds of solutions:

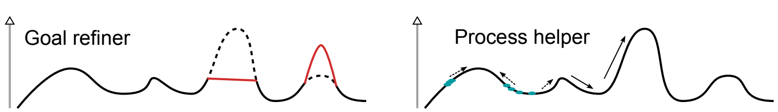

Doncieux, S. and Mouret, J.-B. (2014). **Beyond black-box optimization: a review of selective pressures for evolutionary robotics.** Evolutionary Intelligence, Springer Berlin Heidelberg, publisher. Vol 7 No 2 Pages 71-93.

Multi-objectivization: a convenient tool to modify selective pressures

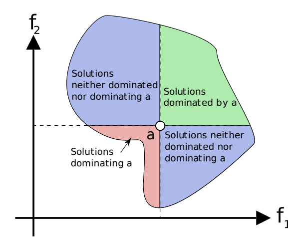

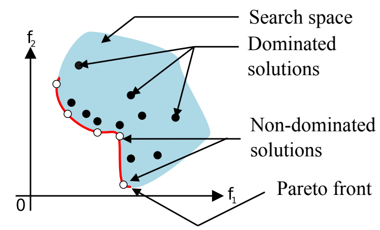

$$\mathbf{f}(g) = \begin{Bmatrix} f_1(g) \\ f_2(g) \\ \vdots \\ f_n(g) \end{Bmatrix}$$

- Goal refiners & process helpers as new objectives
- At the end of the run:
  - Goal refiners: taken into account
  - Process helpers: ignored

Solution to goal definition issues: add « goal refiners »

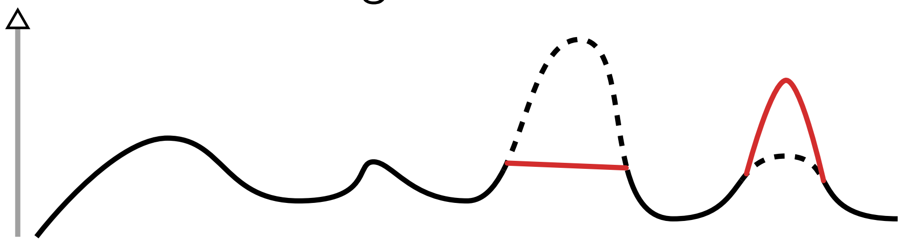

A **goal refiner** aims at changing the optimum(s) of the fitness function by adding new requirements.

Typical challenges that can be addressed:

- Overfitting & generalisation
- Reality gap

Overfitting

Encouraging reactivity

- Encouraging robot controllers to react to sensor stimuli
- Proposition: maximizing the mutual information between sensors and effectors:

$$I(X,Y) = \int_Y \int_X p(x,y) log \left(\frac{p(x,y)}{p(x)p(y)} \right) dxdy$$

Lehman, J., Risi, S., D'Ambrosio, D., & O Stanley, K. (2013). **Encouraging reactivity to create robust machines**. Adaptive Behavior, 21(6), 484-500.

A goal refiner to overcome overfitting

Fitness:

1. distance to the goal
2. reactivity

Multi-objective EA: NSGA-II

Neuroevolution (HyperNEAT)

Lehman, J., Risi, S., D'Ambrosio, D., & O Stanley, K. (2013). **Encouraging reactivity to create robust machines**. Adaptive Behavior, 21(6), 484-500.

Solution to exploration issues: add « process helpers »

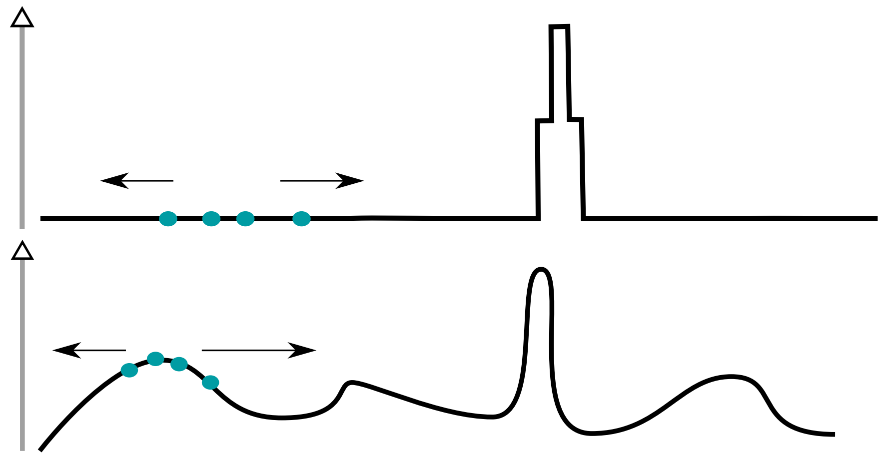

A **process helper** intends to increase the efficiency of the search process without changing the optimum(s) of the fitness function.

Novelty search

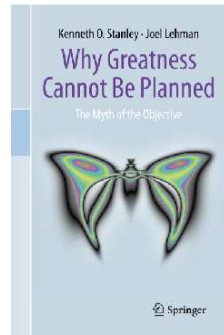

- Novelty based fitness: $\rho(x) = \frac{1}{k}\sum_{i=0}^{k} dist(x, \mu_i)$

  $\{\mu_0, \ldots, \mu_{k-1}\}$ are the k-nearest neighbors in pop+archive

- Archive augmented with individuals having a high novelty

Lehman, J., & Stanley, K. O. (2010). **Abandoning Objectives: Evolution Through the Search for Novelty Alone**. Evolutionary Computation, 19(2), 189–223.

Novelty search

Maze navigation experiment, robot end position

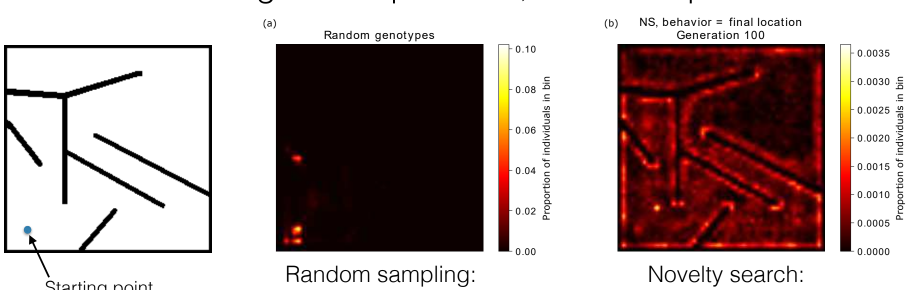

Random sampling: 10000 individuals

Novelty search: 100 gen x 100 Indiv.

Novelty search asymptotically behaves like a **uniform random search in the behavior space**

Doncieux, S., Laflaquière, A., Coninx, A. (2019). **Novelty Search: a Theoretical Perspective**. In *Proceedings of the 2019 Annual Conference on Genetic and Evolutionary Computation*. ACM.

Collect ball experiment

Many different definitions of Novelty Search

Gomes, J., Mariano, P., & Christensen, A. L. (2015, July). **Devising effective novelty search algorithms: A comprehensive empirical study**. In Proceedings of GECCO (pp. 943-950). ACM.

Fitness objectives:

1. nb_{ball}
2. Behavioral diversity (Archive-free Novelty)

Multi-objective EA: NSGA-II

Neuroevolution

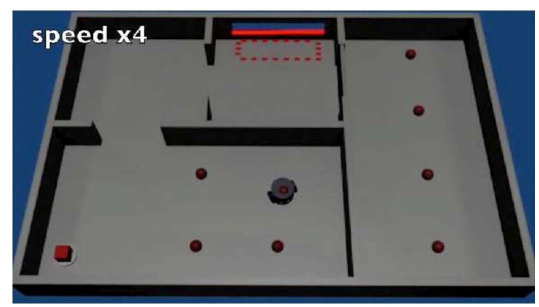

https://github.com/doncieux/collectball

Mouret, J.-B. and Doncieux, S. (2012). **Encouraging Behavioral Diversity in Evolutionary Robotics: an Empirical Study**. Evolutionary Computation. Vol 20 No 1 Pages 91-133.

Dealing with goal definition and exploration at once

- Changing views:
  - Exploration as a priority: generate all solutions of interest
  - Performance as a secondary, local pressure
  - ➡ **Illumination or Quality Diversity algorithms**
- Main ideas:
  - Process helper: selection mostly driven by behavior novelty
  - Goal refiner: a posteriori selection of the most appropriate solution

QD algorithms

~~Looking for **the** optimal solution~~

Looking for a large set of original and efficient solutions

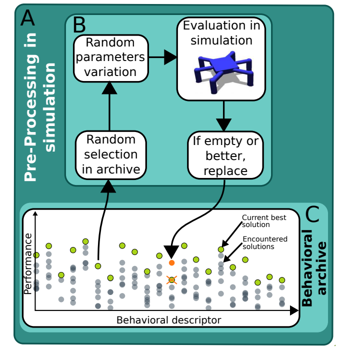

MAP-Elites

Mouret, J. B., & Clune, J. (2015). **Illuminating search spaces by mapping elites.** arXiv:1504.04909.

Pugh, J. K., Soros, L. B., & Stanley, K. O. (2016). **Quality diversity: A new frontier for evolutionary computation**. Frontiers in Robotics and AI, 3, 40.

Cully, A., & Demiris, Y. (2018). **Quality and diversity optimization: A unifying modular framework**. IEEE Transactions on Evolutionary Computation, 22(2), 245-259.

Quality Diversity Search for Robot Ball Throwing Experiment

ISIR, Sorbonne University

SORBONNE UNIVERSITÉ CRÉATEURS DE FUTURS DEPUIS 1257

Kim, S., Coninx, A. & Doncieux, S. (2019) From exploration to control: learning object manipulation skills through novelty search and local adaptation. arXiv:1901.00811

Acquisition and adaptation of a robot behavior repertoire for ball throwinexperiment

Seungsu Kim and Stéphane Doncieux

Institute of Intelligent Systems and Robotics (ISIR)
Sorbonne University

SORBONNE UNIVERSITÉ CRÉATEURS DE FUTURS DEPUIS 1257

Kim, S., Coninx, A. & Doncieux, S. (2019) From exploration to control: learning object manipulation skills through novelty search and local adaptation. arXiv:1901.00811

Conclusion on selective pressures

- The definition of the fitness is critical
- Beyond black box optimization
- Multi-objective framework convenient: **multi-objectivization**
- QD/Illumination algorithms

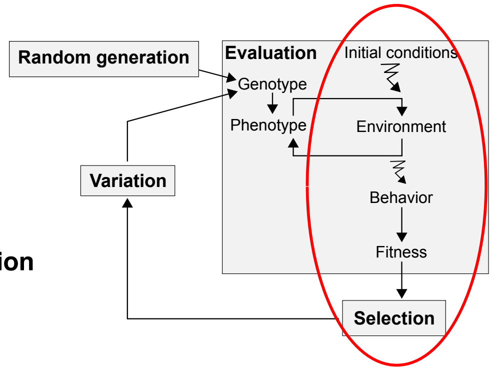

Doncieux, S. and Mouret, J.-B. (2014). **Beyond black-box optimization: a review of selective pressures for evolutionary robotics.** Evolutionary Intelligence, Springer Berlin Heidelberg, publisher. Vol 7 No 2 Pages 71-93.

Evolution, simulators, and the reality

Jean-Baptiste Mouret

Inria Nancy-Grand Est

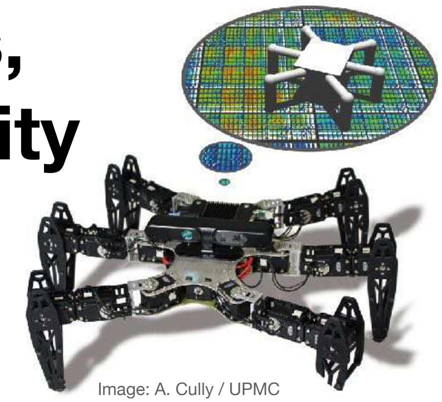

Image: A. Cully / UPMC

Inria INVENTEURS DU MONDE NUMÉRIQUE — UNIVERSITÉ DE LORRAINE — cnrs

No simulator

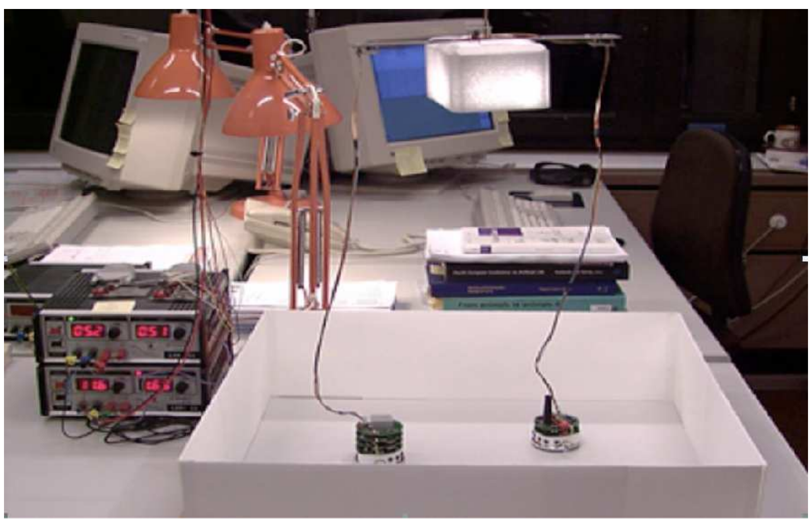

Floreano, Dario, and Francesco Mondada (1996). "Evolution of homing navigation in a real mobile robot." Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on 26.3: 396-407.

Nolfi, S., & Floreano, D. (2001). Evolutionary robotics. The biology, intelligence, and technology of self-organizing machines . MIT press.

No simulator

locomotion

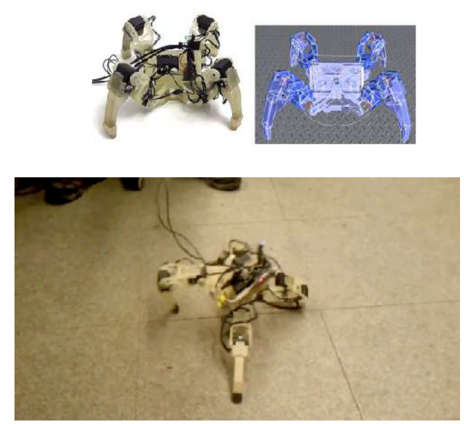

Hornby, G. S., Takamura, S., Yamamoto, T., & Fujita, M. (2005). Autonomous evolution of dynamic gaits with two quadruped robots. Robotics, *IEEE Transactions on*, 21(3), 402-410.

Yosinski, J., Clune, J., Hidalgo, D., Nguyen, S., Zagal, J., & Lipson, H. (2011). Evolving robot gaits in hardware: the HyperNEAT generative encoding vs. parameter optimization. In *Proc. of ECAL*, pp. 890-897.

No simulator

evolving walking controllers

	Starting	Time (1 run)	Robot	DOFs	Param.
Chernova and Veloso (2004)	random	5 h	quadruped	12	54
Zykov et al. (2004)	random	2 h	hexapod	12	72
Berenson et al. (2005)	random	2 h	quadruped	8	36
Hornby et al. (2005)	non-falling	25 h	quadruped	19	21
Mahdavi and Bentley (2006)	random	10 h	snake	12	1152
Barfoot et al. (2006)	random	10 h	hexapod	12	135
Yosinski et al. (2011)	random	2 h	quadruped	9	5

Pros
- (almost) no reality gap
- can exploit unknown physics

Cons
- slow (too slow?)
- will not be faster next year
- **never 100% real**
- **require priors (controller)**

4

Evolving morphologies

… in the real world

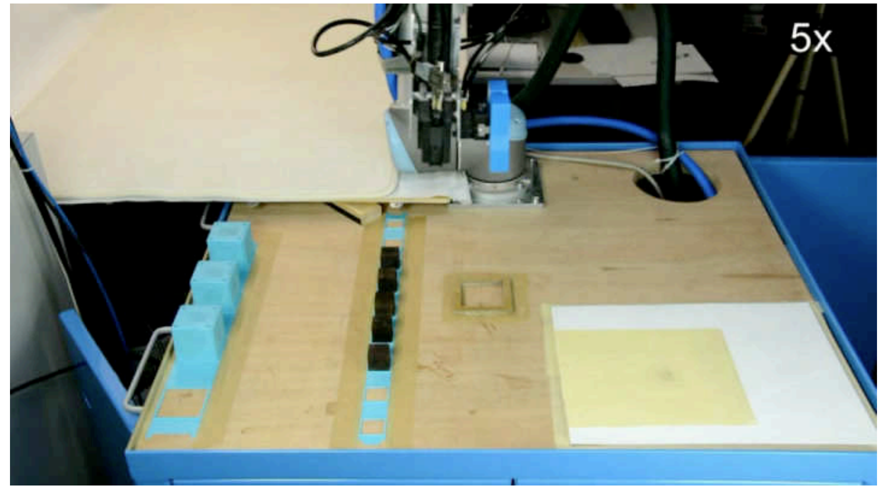

Brodbeck L, Hauser S, Iida F (2015) Morphological Evolution of Physical Robots through Model-Free Phenotype Development. PLoS ONE 10(6): e0128444. https://doi.org/10.1371/journal.pone.0128444 (creative commons)

5

Evolving 3D programs for 3D printers

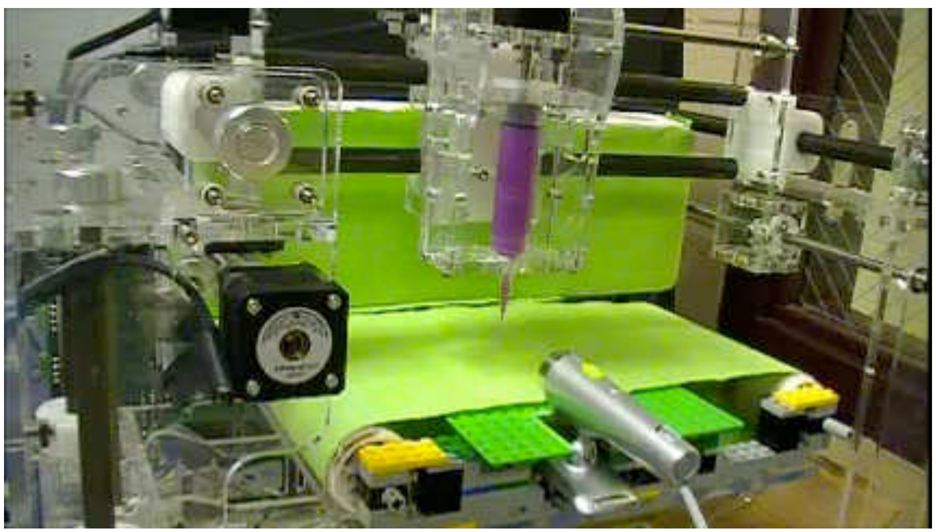

Kuehn, T. and Rieffel, J. (2012) Automatically Designing and Printing Objects with EvoFab 0.2'', Proceedings of the 13th International Conference on the Synthesis and Simulation of Living Systems (ALife XIII), pp. 372-378

6

Using simulators

useful tools?

Evolution is a slow process (millions of years?)
… but computers are faster every year

Can we 'accelerate time'?

We now have many "good" simulators:
- ODE (library): www.ode.org
- Bullets (library): bulletphysics.org
- Dart (library): https://github.com/dartsim/dart
- *[Gazebo (GUI): gazebosim.org]*
- *[V-Rep (GUI): www.coppeliarobotics.com]*
- …

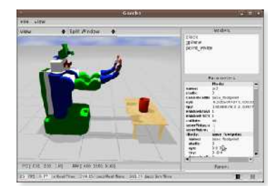

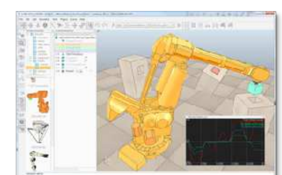

S. Ivaldi *et al.* (2014). Tools for dynamics simulation of robots: a survey based on user feedback. Proc. of Humanoids
J.-.B. Mouret and K. Chatzilygeroudis (2017). 20 Years of Reality Gap: a few Thoughts about Simulators in Evolutionary Robotics. GECCO workshop (SimER) — 2017

7

The reality gap

… or what always happens with simulators and robots

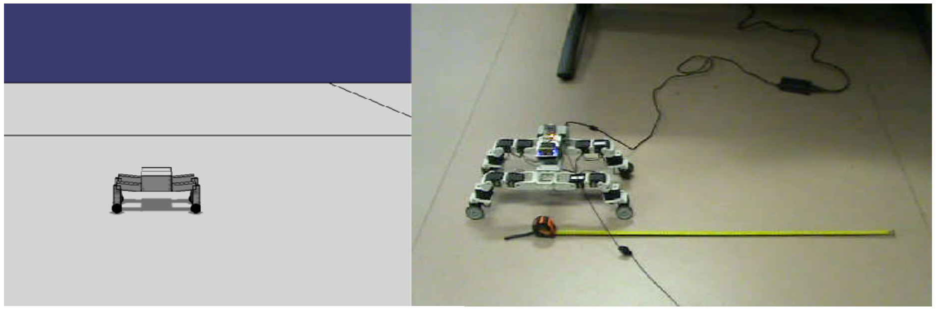

Koos, Mouret & Doncieux.
IEEE Transactions on Evolutionary Computation. 2012

Controller: 2 parameters

Jakobi, Nick. "Running across the reality gap: Octopod locomotion evolved in a minimal simulation." Evolutionary Robotics. Springer Berlin Heidelberg, 1998.

Reality vs simulation

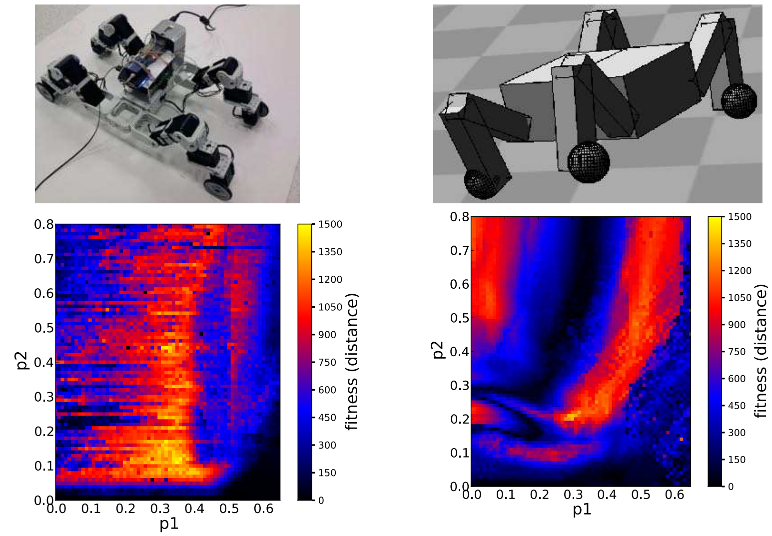

Mouret, J. B., Koos, S., & Doncieux, S. (2013). Crossing the reality gap: a short introduction to the transferability approach. arXiv preprint arXiv:1307.1870.

But they can agree (sometimes)!

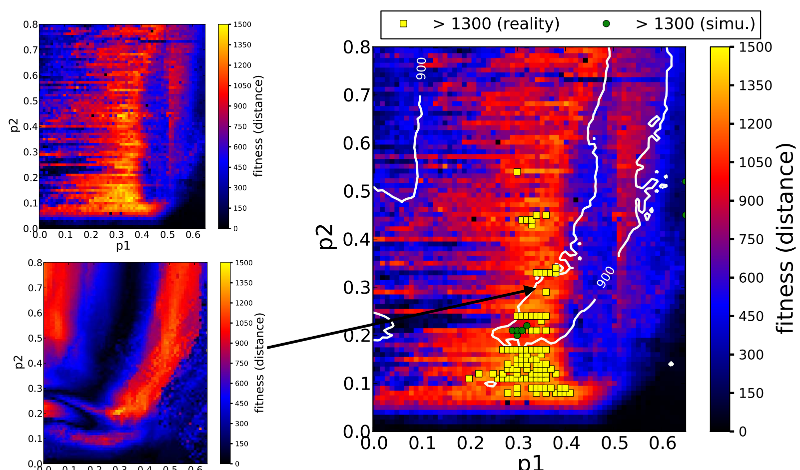

The reality gap

- Any simulation has a validity domain
- Human experts know this validity domain
- … but evolution does not have this common sense

Results found in simulation have a low probability of working similarly in reality

➠ **One of the main problems of ER as a design tool**

“Sim2Real" in “deep learning”

What can we do?

no simulator
better simulator
avoid non-transferable solutions
robust controllers

Finish evolution in reality

evolve in simulation, then do a few generations with the robot

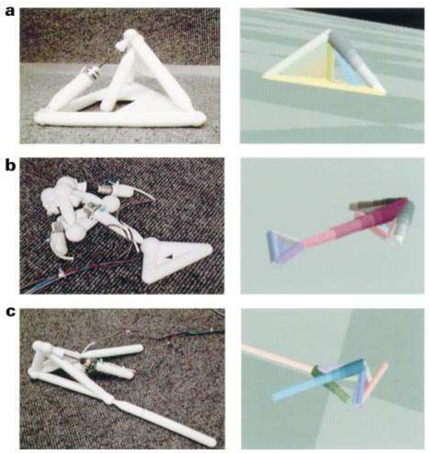

Pro: can help fine-tuning the solution obtained in simulation

Con: "local search" in the vicinity of the solutions found in simulation

➠ cannot find something completely different

Lipson, H., & Pollack, J. B. (2000). Automatic design and manufacture of robotic lifeforms. **Nature**, 406, 974–978.

Impr

General id

simulation a

➠ Miglino

sensors

➠ Moeck

param

simu

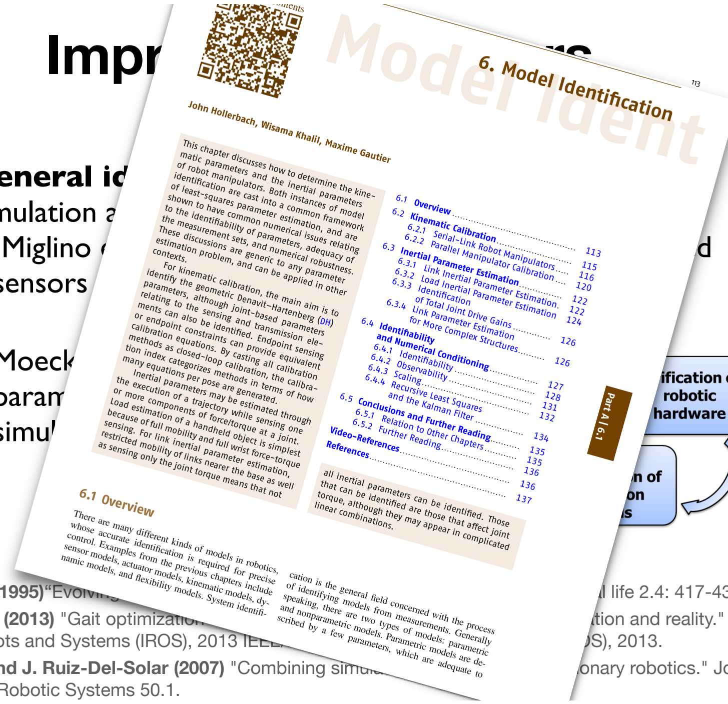
6. Model Identification

John Hollerbach, Wisama Khalil, Maxime Gautier

This chapter discusses how to determine the kinematic parameters and the inertial parameters of robot manipulators. Both instances of model identification are cast into a common framework of least-squares parameter estimation, and are shown to have common numerical issues relating to the identifiability of parameters, adequacy of the measurement sets, and numerical robustness. These discussions are generic to any parameter estimation problem, and can be applied in other contexts.

For kinematic calibration, the main aim is to identify the geometric Denavit–Hartenberg (DH) parameters, although joint-based parameters relating to the sensing and transmission elements can also be identified. Endpoint sensing or endpoint constraints can provide equivalent calibration equations. By casting all calibration methods as closed-loop calibration, the calibration index categorizes methods in terms of how many equations per pose are generated.

Inertial parameters may be estimated through the execution of a trajectory while sensing one or more components of force/torque at a joint. Load estimation of a handheld object is simplest because of full mobility and full wrist force–torque sensing. For link inertial parameter estimation, restricted mobility of links nearer the base as well as sensing only the joint torque means that not all inertial parameters can be identified. Those that can be identified are those that affect joint torque, although they may appear in complicated linear combinations.

6.1 Overview

There are many different kinds of models in robotics, whose accurate identification is required for precise control. Examples from the previous chapters include sensor models, actuator models, kinematic models, dynamic models, and flexibility models. System identification is the general field concerned with the process of identifying models from measurements. Generally speaking, there are two types of models: parametric and nonparametric models. Parametric models are described by a few parameters, which are adequate to

ification on robotic hardware

n of on s

Miglino et al. (1995) "Evolvi... l life 2.4: 417-434.

Moeckel et al. (2013) "Gait optimization ... tion and reality." Intelligent Robots and Systems (IROS), 2013 IEEE ... S), 2013.

Zagal, J. C., and J. Ruiz-Del-Solar (2007) "Combining simula... onary robotics." Journal of Intelligent and Robotic Systems 50.1.

Improving simulators

The EEA algorithm: active learning of a self-model

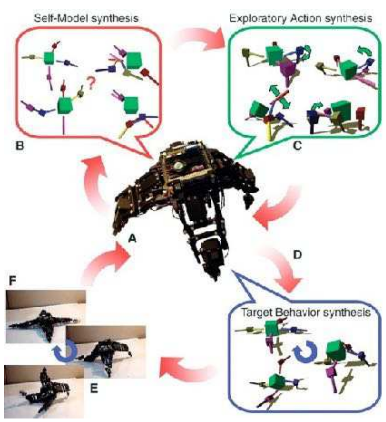

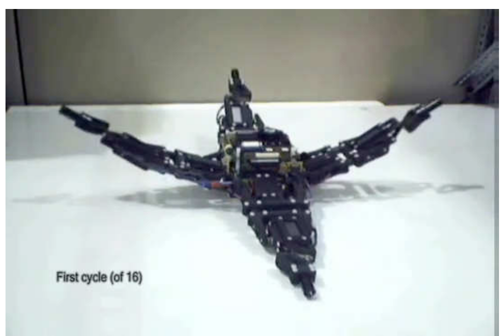

Bongard, Zykov and Lipson (2006). Science.

Koos, S., Mouret, JB and Doncieux, S. (2009) "Automatic system identification based on coevolution of models and tests." *Proc. of IEEE CEC.*

Learning the simulator from data

= learning the dynamical model of the robot

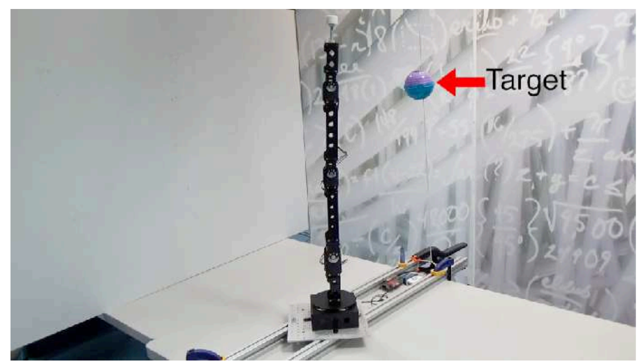

1. try the best policy according to the model
   ➠ new data
   ➠ new model (Gaussian processes)
2. find a policy that maximises the fitness according the simulator, Taking the uncertainty into account

Chatzilygeroudis K, Rama R, Kaushik R, Goepp D, Vassiliades V, Mouret JB. (2017) Black-Box Data-efficient Policy Search for Robotics. Proc. of IEEE IROS 2017.

Correcting the simulator

System identification + free modelling

$$\mathbf{x}_{t+1} = \mathbf{x}_t + M(\mathbf{x}_t, \mathbf{u}_t, \phi_M) + f(\mathbf{x}_t, \mathbf{u}_t, \phi_K) + \mathbf{w}$$

Simulator / model   parameters

Learning = maximize the likelihood of M+f

We can combine model learning and model identification

- ➟ **effects that can be captured** by the simulator will be included by tuning the simulator (model identification)
- ➟ **effects that cannot be captured** by changing the parameters are modelled by the Gaussian processes

Chatzilygeroudis K, Mouret JB. (2018) Using Parameterized Black-Box Priors to Scale Up Model-Based Policy Search for Robotics. Proc. of ICRA.

Model identification + correction

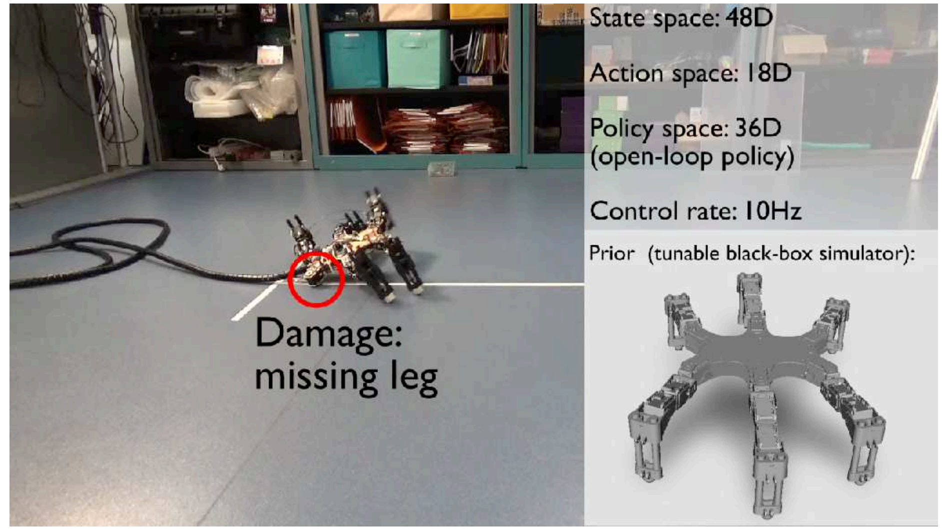

Chatzilygeroudis K, Mouret JB. (2018) Using Parameterized Black-Box Priors to Scale Up Model-Based Policy Search for Robotics. Proc. of ICRA.

Surrogate modelling / Bayesian optimization

learn a model of the fitness function

- Use data to predict the fitness given the parameters
- No need to sense “states"
- Work well if a few parameters (< 6)
- Usually do not work on structures (but come to see our talk!)

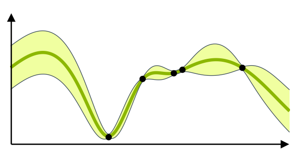

Model = Gaussian process
EA = CMA-ES

Rieffel, J., & Mouret, J.-B. (2018). Soft tensegrity robots. Soft Robotics.

Improving simulators & models

Pros	Cons
mix simulation and reality: the best of both worlds?	the simulator will never be perfect (generalization)
faster than learning without a simulator	if the correction cannot be applied? (e.g. aerodynamics)
morphological / env. changes	learning a simulator is hard!

Jin, Y. (2005) "A comprehensive survey of fitness approximation in evolutionary computation." Soft computing 9.1 (2005): 3-12.

Avoiding bad simulations
the envelope of noise & minimal simulations

Simulate only the useful effects
Hide in an "envelope of noise" things that are too hard to simulate accurately
- keep evolution from exploiting simulation artefacts
- goal refiner

Examples:

- Khepera robot: add noise to the sensors and the actuators
- Octopod robot: minimal simulation

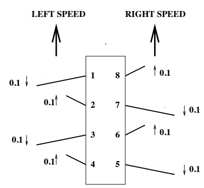

In deep learning: "Domain randomization"

Jakobi, N. (1997) "Evolutionary robotics and the radical envelope-of-noise hypothesis." Adaptive Behavior 6.2: 325-368.

Avoiding bad simulations
envelope of noise & minimal simulations

Pros	Cons
Lightweight simulations	Hard to set-up
Noise increases robustness and generalization	What noise? what is important?
	No surprising dynamic effect
	Noise makes evolution harder

Avoiding bad simulations
the transferability approach

- learn the limits of the simulation (supervised learning)
- focus the search on well-simulated behaviors
- the transferability is a task-agnostic goal refiner

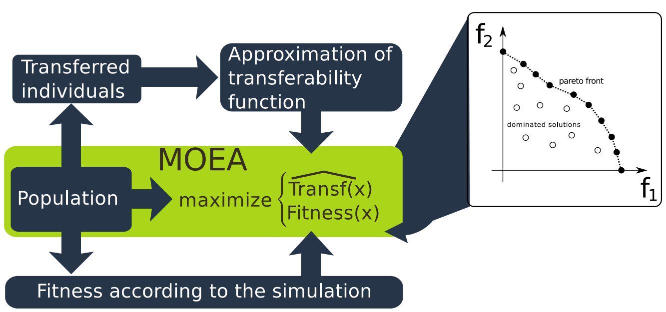

Mouret, Koos & Doncieux (2012). ALIFE workshop. 2012
Koos, Mouret & Doncieux (2012). IEEE TEC. 2012
Koos, Cully & Mouret. (2013). IJRR. 2013

Avoiding bad simulations
the transferability approach

Maximize fitness

Maximize { fitness, transferability }

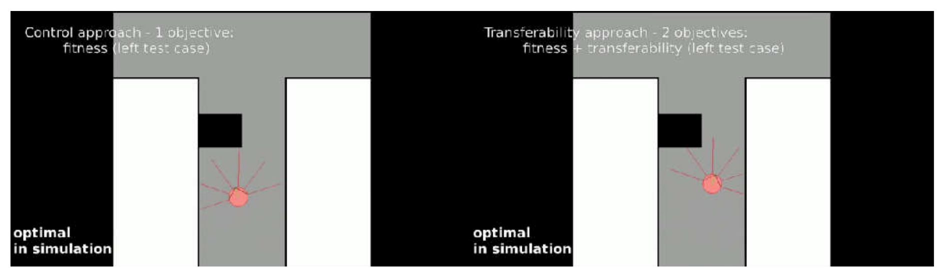

15 transfers
(motion capture)

Koos, S., Mouret, J.-B., & Doncieux, S. (2011). The Transferability Approach : Crossing the Reality Gap in Evolutionary Robotics. IEEE Transaction on Evolutionary Computation, 1, 1–25.

Avoiding bad simulations

the transferability approach

Maximize fitness

Maximize { fitness, transferability

15 transfers
(motion capture)

Koos, S., Mouret, J.-B., & Doncieux, S. (2011). The Transferability Approach : Crossing the Reality Gap in Evolutionary Robotics. IEEE Transaction on Evolutionary Computation, 1, 1–25.

24

Mapping, then searching

Intelligent Trial & Error

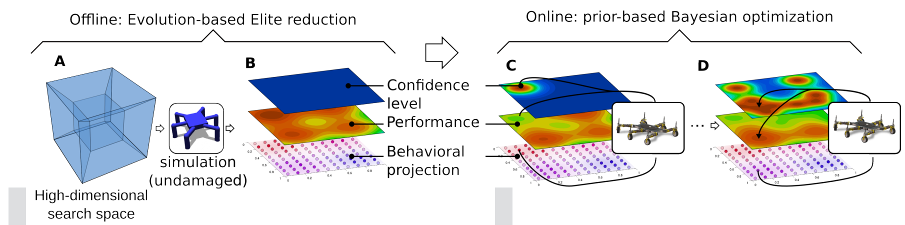

Mapping (offline)
- MAP-Elites algorithm (illumination/quality diversity)
- search for the best behavior of each family

Adaptation (online)
- Bayesian optimization
- levels of confidence

Cully, A., Clune, J., Tarapore, D., & Mouret, J. B. (2015). Robots that can adapt like animals. *Nature*, *521*(7553), 503-507.

25

Mapping, then searching

Intelligent Trial & Error

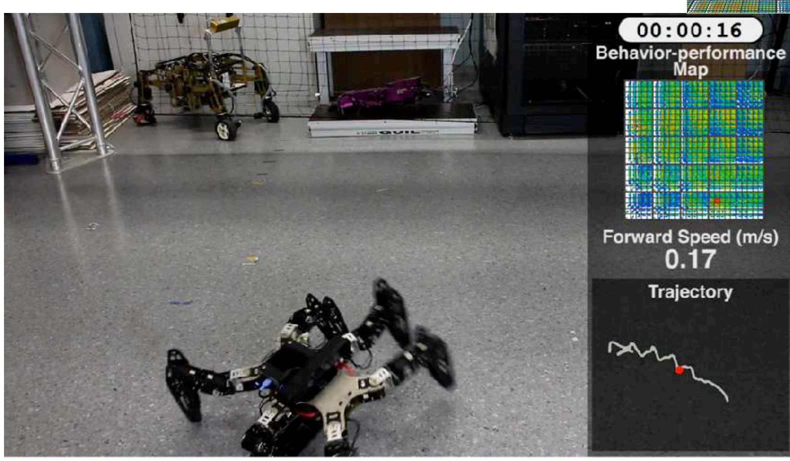

Cully, Clune, Tarapore & Mouret (2015). Robots that can adapt like animals. **Nature**.

26

Avoiding bad simulations

the transferability approach

Pros

Easier to learn the limit than to correct/learn the simulator

Only a few test on the robot: no need for a special set-up

Cons

The EA cannot exploit phenomena that not simulated at all

(e.g. highly-dynamic gaits, unknown aerodynamic effects, etc.)

27

Improving robustness

evolve controllers with online learning abilities

Example: neural networks with “adaptives synapses”

f = 0.422, <f> = 0.499

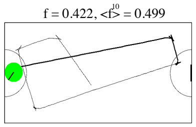

Adaptive synapses

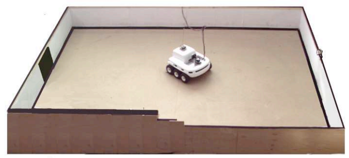

Floreano, D., & Urzelai, J. (2000). Evolutionary robots with on-line self-organization and behavioral fitness. Neural Networks, 13(4-5), 431–43.

Urzelai, J., & Floreano, D. (2000) "Evolutionary robots with fast adaptive behavior in new environments." Evolvable Systems: From Biology to Hardware. Springer Berlin Heidelberg. 241-251.

Improving robustness

encouraging reactivity

- quantification of reactivity derived from the mutual information between sensors and actuators
- multi-objective optimization
- even better if combined with noise

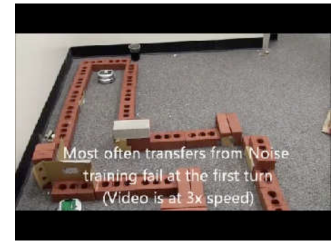

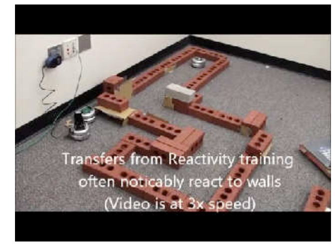

Lehman, Joel, et al. (2013) "Encouraging reactivity to create robust machines." Adaptive Behavior (2013): 1059712313487390.

Conclusion

the reality gap

1 ➠ No simulator: possible but slow (swarm?)

2 ➠ Finishing evolution on the physical robot: similar optima

3 ➠ Improving simulators: not always enough data to learn
- *system identification*

4 ➠ Avoiding badly simulated solutions
- *add noise to sensors and actuators: hard to tune*
- *minimal simulations: requires expert knowledge*
- *learn the transferability function*

5 ➠ Improving robustness: no guarantee
- *add online learning abilities*
- *encourage reactivity*

Conclusion

the reality gap

- No perfect approach to cross the reality gap
- Avoiding simulation is materially challenging and slow
- No perfect simulation
- Simulators should give their confidence (and not only a prediction of the fitness)

➠ it depends on the scientific question!
- show the potential of a new encoding? a new selective pressure? simulation might be enough
- solve challenging robotics problem? this needs to work on real robots

Evolutionary robotics for collective robotics

A 30-minute overview

GECCO 2019

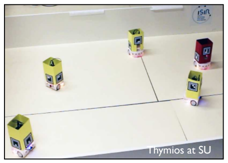

Nicolas Bredeche
Sorbonne Université
Institut des Systèmes Intelligents et de Robotique
UPMC - CNRS (UMR 7222)
Paris, France

http://pages.isir.upmc.fr/~bredeche
e-mail: nicolas.bredeche@sorbonne-universite.fr

Note on citation policy: for a given topic, I cite either or both the seminal reference and a recent one. E.g.: [Nolfi and Floreano. 2000][Doncieux et al. 2015] for referring to general resources on evolutionary robotics. Non-first authors may be omitted for clarity.

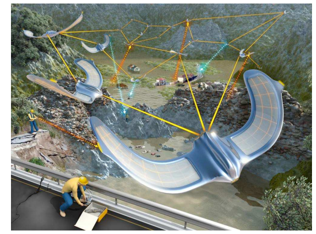

SmavNET@EPFL

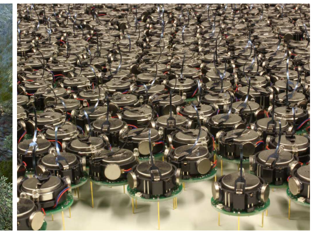

SSR/Harvard

H2020 COLLMOT

Collective robotics: multiple robots, acting together, to achieve a common goal.

Swarm robotics: collective robotics with large population of "simple" robots (i.e. *limited computation and communication capabilities*). It is a *distributed* system.

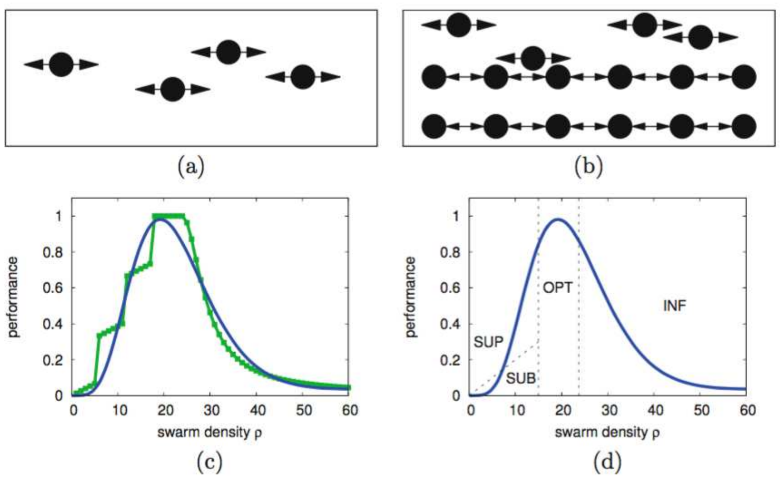

Fig. 1.5 Bucket brigade example for swarm performance (robots have to transport objects back and forth between the left and right side of the robot arena) and typical swarm performance function over swarm density $\rho = N/A$ for a fixed area $A = 1$ (without units). **(a)** Bucket brigade, $N = 4$ robots **(b)** Bucket brigade, $N = 16$ robots **(c)** Bucket brigade, performance. **(d)** Swarm performance showing four regions, SUP: super-linear, SUB: sub-linear, OPT: optimal, INF: interference

 H. Hamman (2018) pp.10

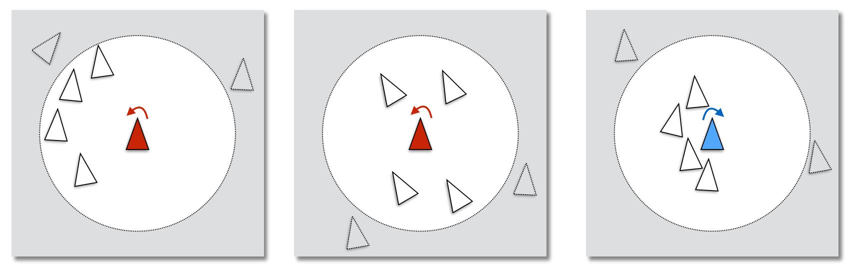

Attraction Orientation Repulsion

Positive and negative feedbacks

positive feedback: *attraction and orientation rules*

negative feedback: *repulsion rule*

Remark: assume constant speed and limited scope

 Reynolds (1987) Flocks, herds and schools: a distributed behavioral model

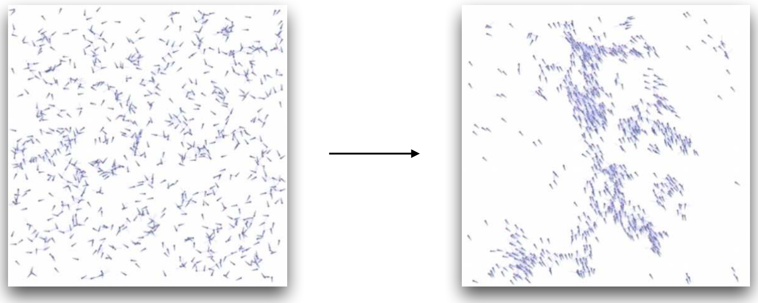

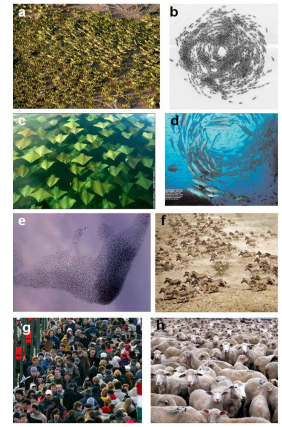

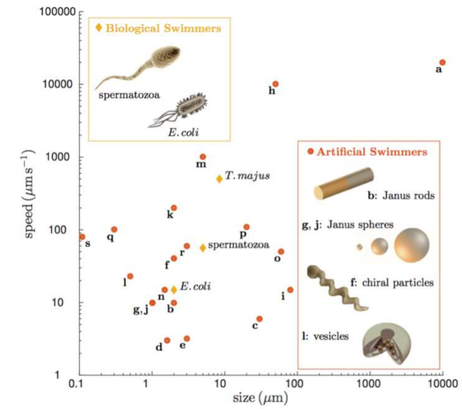

Self-organization: a spontaneous process where global coordination arises out of local interactions between components of a system (e.g. *nest building in ants/termites/bees, coordinate movements in herd/swarm/schools*).

from: Vicsek et al. (2012) Collective motion

from: Cechlinger et al. (2016)

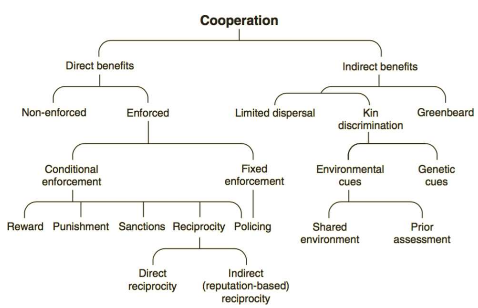

Cooperative behaviour: a behaviour that provides a benefit to another individual <u>and</u> that has evolved at least partially because of this benefit

as defined in West, Griffin, Gardner (2007)

Illustrations from: D. R. Rubenstein (2010) and wikipedia

From: West, Griffin, Gardner (2007) Current Biology

Scope of this talk: <u>distributed</u> collective robotics

> **sensing** and **actions**, and possibly **learning**, are <u>**distributed**</u> over the population

Collective robotics

9

- Scope [Nettleton et al., 2003], adapted from [Capitan et al. 2013]
  - no central control
  - no global communication facility
  - no local knowledge of the team global topology
- (Obvious) advantages
  - Robustness through redundancy
  - Parallelising actions wrt a task
  - Parallelising learning/optimisation (if any)

Methods

10

The "nerd herd" [Mataric, 1992]

self-assembling kilobots [Rubenstein, 2014]

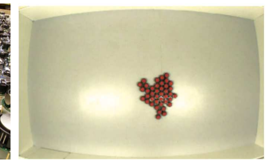

self-aggregation with e-pucks [Gauci 2014]

- Hand design controller for collective robotics
  - Hand design w/ empiric approach [Mataric 1992][Rubenstein 2014]…
  - Hand design w/ (limited) theoretical proofs [Gauci 2014]…
  - Software architecture for multi-robot systems [Parker 2008]…
- Learning in multi-agent systems
  - Assume joint payoff but decentralized actions and observations
  - A lot of assumptions (Markovian environment, discrete space, etc.)
  - Powerful theoretical results [Bernstein 2002][Amato 2014]…
  - … but limited practical works (*very few* robots, individual learning)

from: AAMAS 2011 Tutorial on decision making in MAS (Doshi, Rabinovich, Spaan, Amato)

Evolutionary robotics for collective robotics

11

- What ER offers
  - Automated design method
  - Variety of search space ($\mathbb{R}^n$, graphs)
  - Applicable to the real world
  - Multi-objective search, multiple selection pressure (incl. diversity)
- What are the limits of ER (…more on that later)
  - Reality gap (at least for classic off-line ER)
  - Lack of theoretical grounding for collective adaptive dynamics
  - Too much emphasis on "logical" control (w.r.t. "morphological" control)

Outline

12

engineering

1. ER as an optimisation method
2. ER as an on-line learning method

biology

3. ER as an individual-based modelling method
4. Future of ER for collective robotics

ER as an optimisation method for collective robotics

"classic" evolutionary robotics

Optimisation for collective robotics

[Nolfi, Floreano 2000][Doncieux et al. 2015]

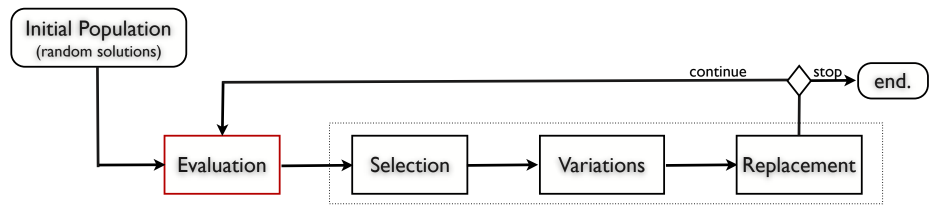

- What?
  - Off-line design method : classic "**evolutionary robotics**" method
  - Optimize in centralized fashion, then used in a distributed fashion
- Expected result
  - A set of <u>policies</u> (*possibly similar*) that can be used within a population of robots to solve a task

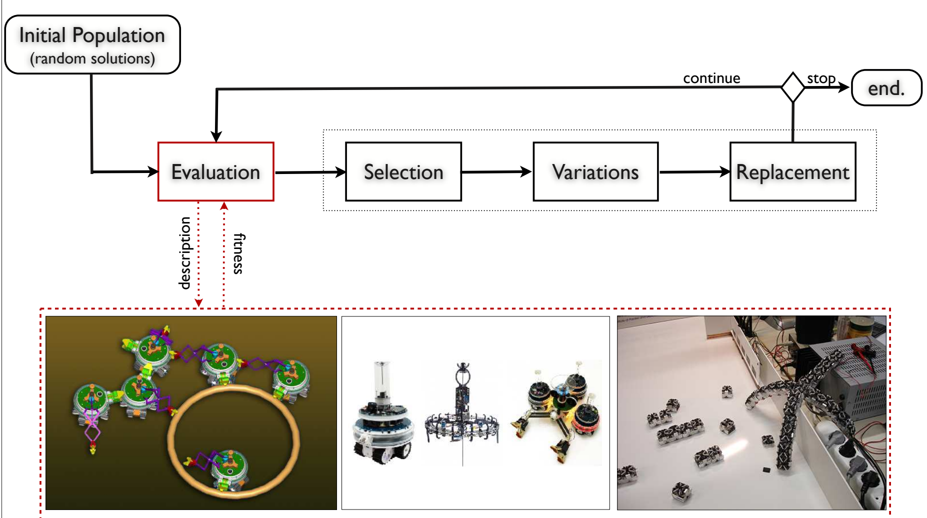

Swarm-bots, 2001-2005 | Swarmanoid, 2006-2010 | Symbrion and Replicator, 2008-2013

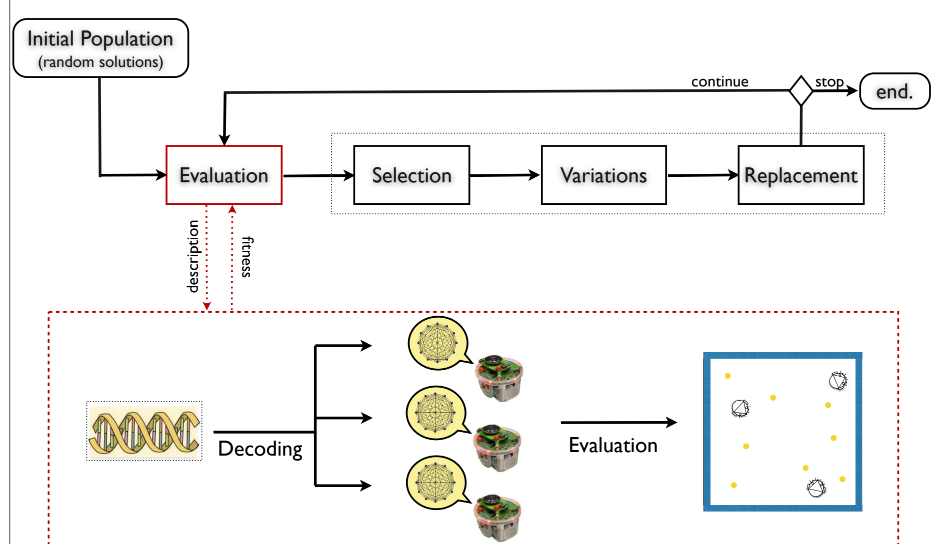

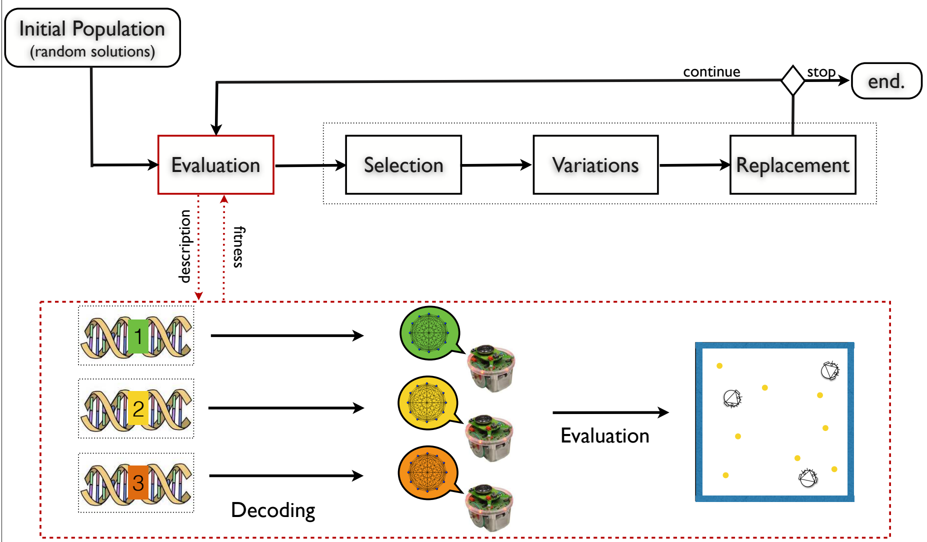

Team composition and levels of selection

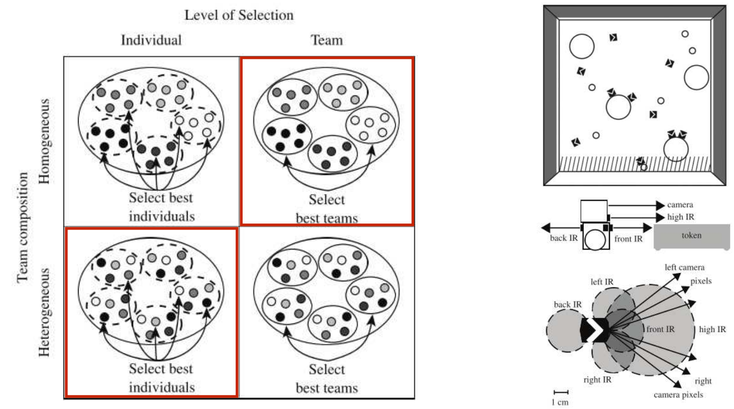

Genetic Team Composition and Level of Selection in the Evolution of Cooperation

Markus Waibel, *Member, IEEE*, Laurent Keller, and Dario Floreano, *Senior Member, IEEE*

IEEE TRANSACTIONS ON EVOLUTIONARY COMPUTATION, VOL. 13, NO. 3, JUNE 2009

Task specialisation with homogeneous team

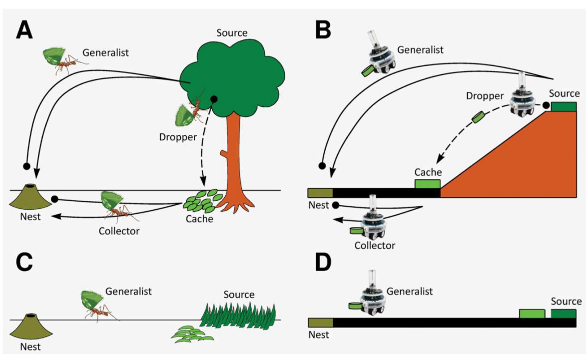

Fig 1. Task partitioning in insects and robots. (a) Task partitioned retrieval of leaf fragments, as found in most *Atta* leafcutter ants that harvest leaves from trees [7,43]. Dropper ants cut leaves which then accumulate in a cache, after which the leaves are retrieved by collectors and brought back to the nest, where they serve as a substrate for a fungus which is farmed as food. Ants also occasionally use a generalist strategy whereby both tasks are performed by the same individuals. (b) Analogous robotics setup, whereby items have to be transported across a slope using the coordinated action of droppers, collectors and possibly generalists. (c) Grass cutting leafcutter ants cutting leaf fragments in a flat environment without task partitioning, using a generalist foraging strategy [49]. (d) Analogous robotics setup, with robots being required to collect items in a flat arena.

PLOS | COMPUTATIONAL BIOLOGY

Evolution of Self-Organized Task Specialization in Robot Swarms

Eliseo Ferrante[1]*, Ali Emre Turgut[2], Edgar Duéñez-Guzmán[1], Marco Dorigo[3], Tom Wenseleers[1]

DOI:10.1371/journal.pcbi.1004273 August 6, 2015

Experimental setup

Homogeneous team of 4 robots
100 teams, 2000 generations
3 evaluations per team
team fitness: #foraged_items

Take-home message for "classic" ER

Team composition and level of selection matter
Homogeneous team can perform task specialisation

[Waibel et al., 2009, TEC]
[Nitschke et al., 2012, GPEM]
[Lichocki et al., 2013, IEEE TEC]
[Tuci et al., 2014, Neural Comp. and Apps.]
[Gomez et al., 2015, AAMAS]
[Bernard et al., 2015, ECAL]
[Bernard et al., 2016, ALIFE]
[Ferrante et al., 2015, PCB]
(…)

ER as an on-line learning method for collective robotics

“embodied” evolutionary robotics

Embodied evolution: distributed on-line learning

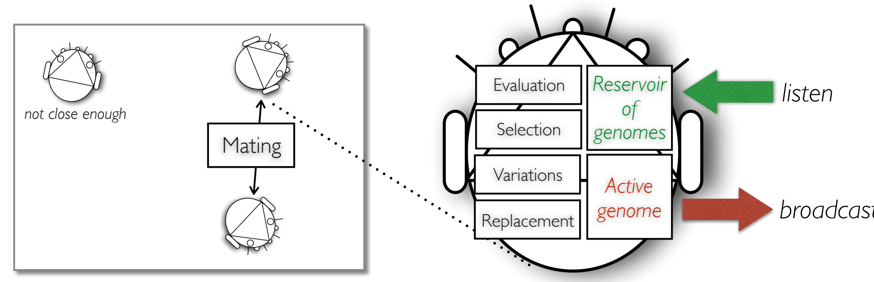

- What?
  - On-line adaptation with limited communication and computation
  - Optimise and use in a distributed fashion
- Expected result
  - A distributed on-line algorithm for lifetime learning

Embodied evolution: distributed on-line learning

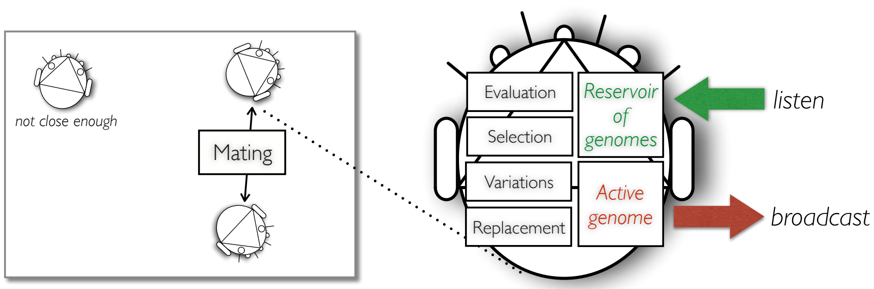

- Obvious advantages
  - On-line
    - No reality gap (by definition)
    - Parallel search (by definition)
  - Distributed
    - Robustness to failure through redundancy
    - Scalability through its distributed nature

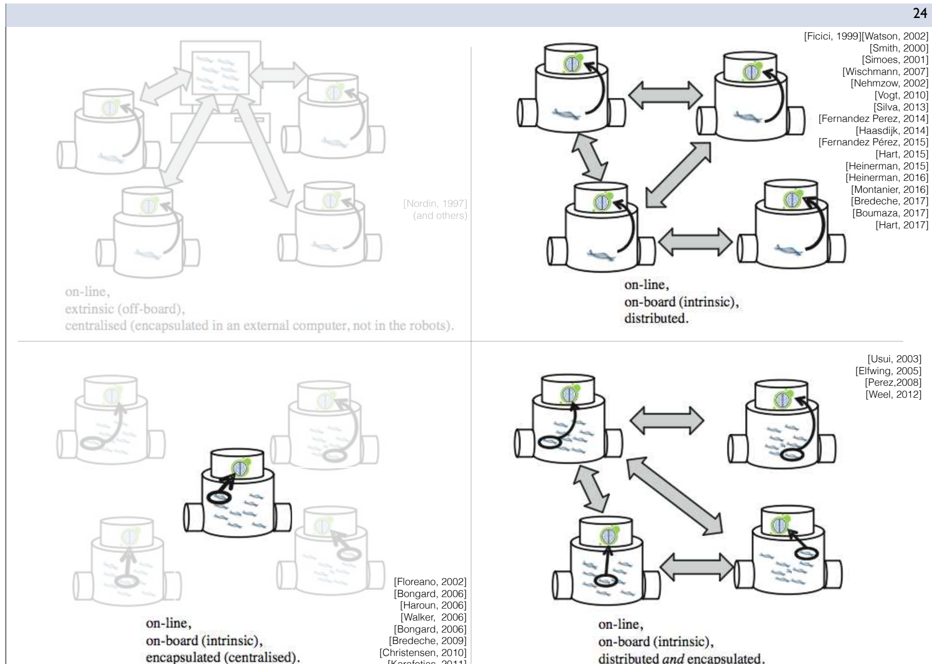

Embodied evolution in a nutshell

A vanilla algorithm

Controller

e.g.: linear combination of inputs and weights, artificial neural networks, etc.

motor1 = a*IR_1 + b*IR_2 + c*IR_3 + d*IR_4 + e
motor2 = f*IR_1 + g*IR_2 + h*IR_3 + i*IR_4 + j

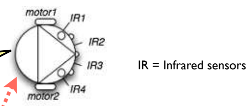

Genome

<
a = +0.31
b = +0.11
c = -1.42
d = +1.6
e = -0.14
f = 0.55
g = -1.17
h = +0.97
>

genome of controller

e.g.: $\mathbb{R}^n$

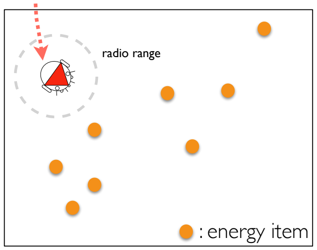

Reservoir of genomes

{genome,fitness value}^n

This list is used to store (unique) **copies of genomes** from robots passing nearby with their **current fitness value** at the time of encounter

Objective function

In the general case, the fitness value is computed thanks to an *embodied* objective function, that is: each robot individually assess its own performance

E.g.: #energy items foraged

Example with a foraging task

Embodied evolution in a nutshell

A vanilla algorithm

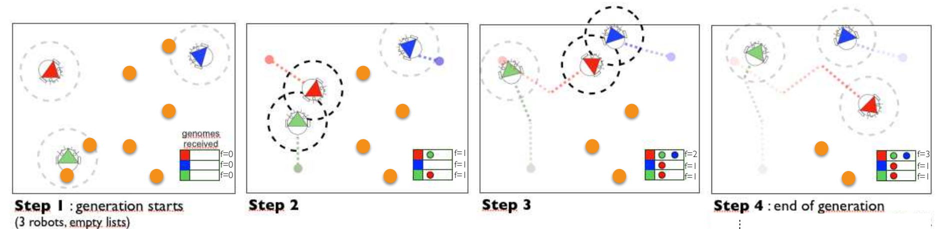

At this point, each robot...

1 - **forgets** its own genome

2 - perform selection among **received** genomes wrt fitness values

3 - apply **variation** (cross-over and/or mutation) on the selected genome (e.g. gaussian mutation)

4 - use new genome to set up new control architecture

Example with a foraging task

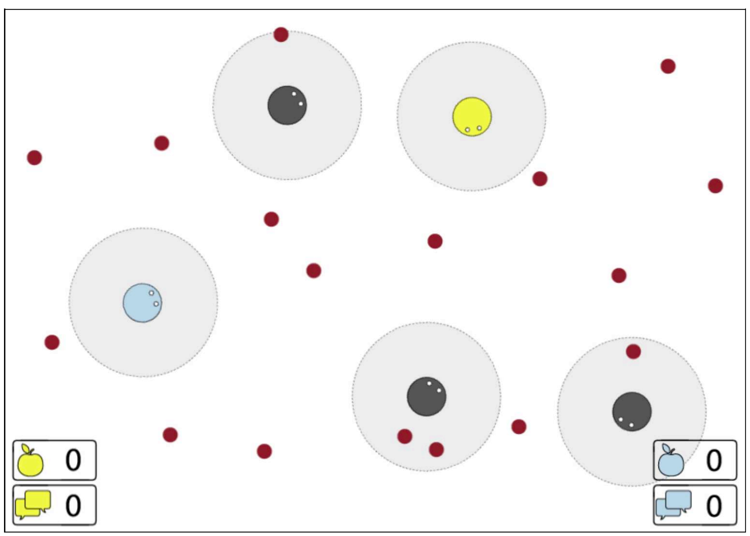

Embodied evolutionary robotics illustrated

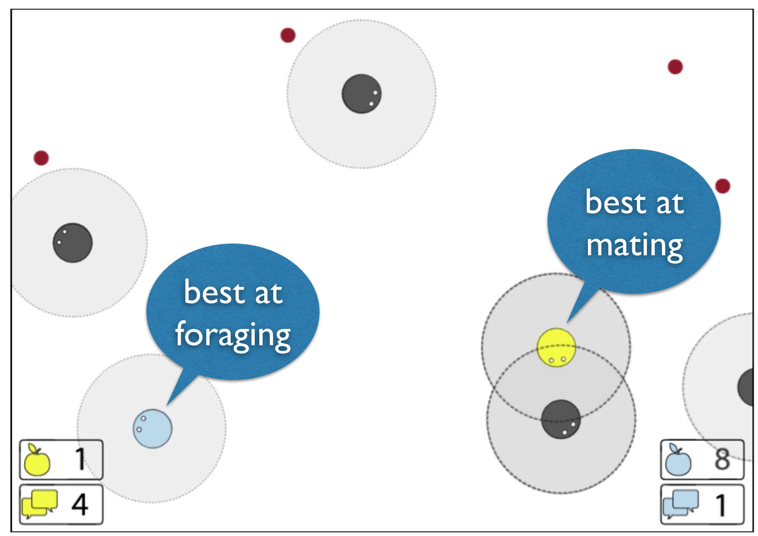

Multiple selection pressures!

Ecological selection pressure

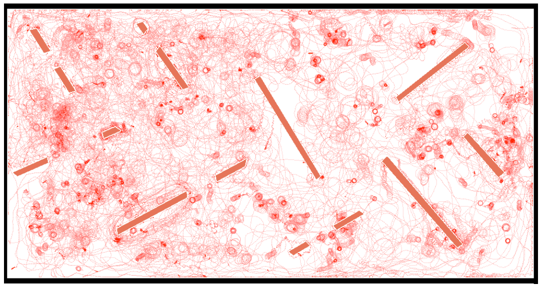

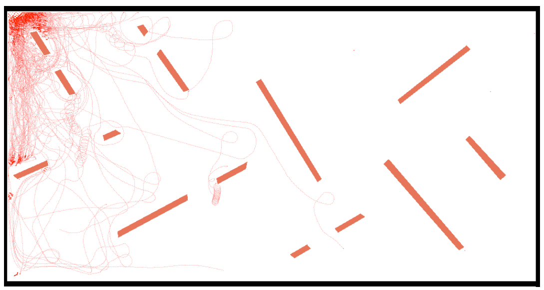

environment w/o constraint

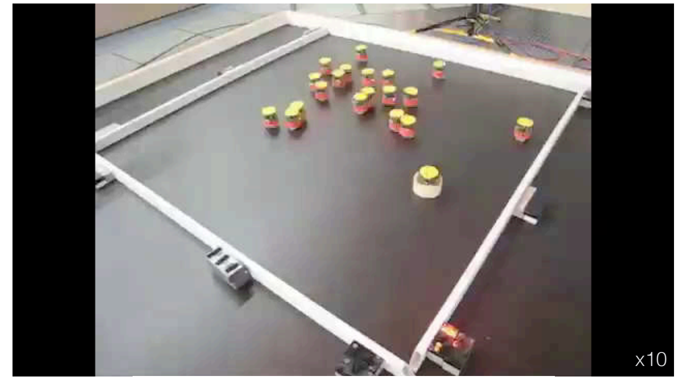

http://www.youtube.com/watch?v=_ilRGcJN2nA

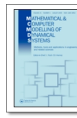

Environment-driven distributed evolutionary adaptation in a population of autonomous robotic agents

Nicolas Bredeche, J-M Montanier, W. Liu, A. F. Winfield

Mathematical and Computer Modelling of Dynamical Systems, Volume 18, Issue 1, 2012

32

Ecological selection pressure (cont.)

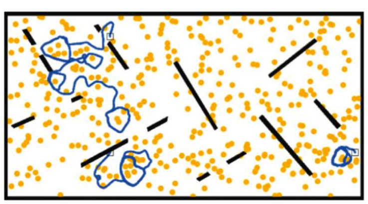

environment w/o constraint

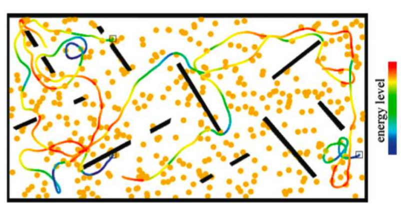

foraging energy is required

Selected (3) trajectories among 100 robots under different constraints

Bredeche, Montanier (2010)

Take-home message for embodied ER

Selection pressure comes from both the environment and the task

[Bredeche et al., 2010, PPSN]
[Haasdijk et al, 2014, Plos One]
[Hart et al., 2015, GECCO]
[Perez et al., 2015, ALIFE]
[Steyven et al., 2016, PPSN]
[Montanier et al., 2016, Frontiers in AI and Robotics]
(...)

nicolas.bredeche@upmc.fr

ER as an individual-based modelling tool

evolutionary robotics applied to biology

- Relevance as a modelling and simulation method
  - vs. mathematical modelling
    - simulates mechanistic aspects
  - vs. *in vitro* studies (or *in vivo* observations)
    - simulates longer evolutionary timescale

« So far, we have been able to study only one evolving system and we cannot wait for interstellar flight to provide us with a second. If we want to discover generalizations about evolving systems, we have to look at artificial ones. »

NATURE · VOL 355 · 27 FEBRUARY 1992

Byte-sized evolution

John Maynard Smith

Evolution of altruistic cooperation

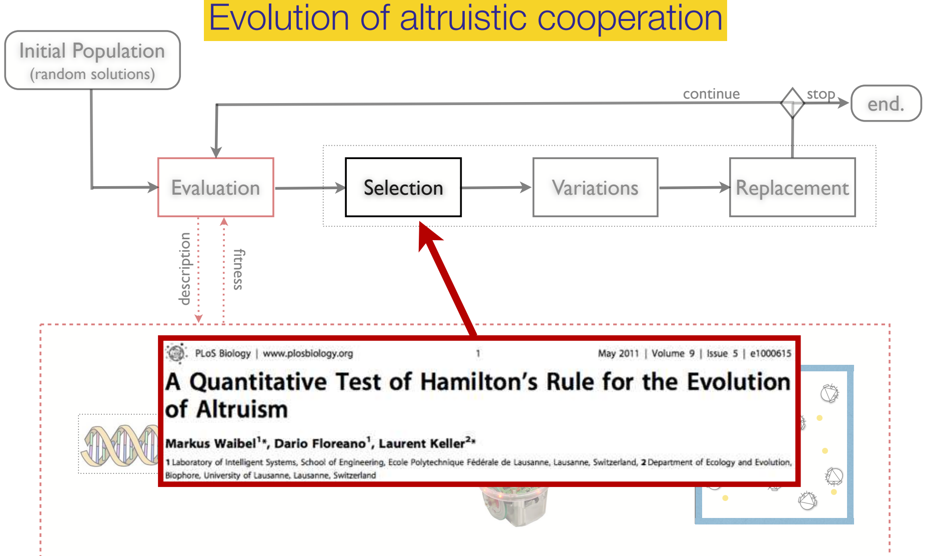

PLoS Biology | www.plosbiology.org 1 May 2011 | Volume 9 | Issue 5 | e1000615

A Quantitative Test of Hamilton's Rule for the Evolution of Altruism

Markus Waibel[1]*, Dario Floreano[1], Laurent Keller[2]*

1 Laboratory of Intelligent Systems, School of Engineering, Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland, 2 Department of Ecology and Evolution, Biophore, University of Lausanne, Lausanne, Switzerland

Evolution of signalling

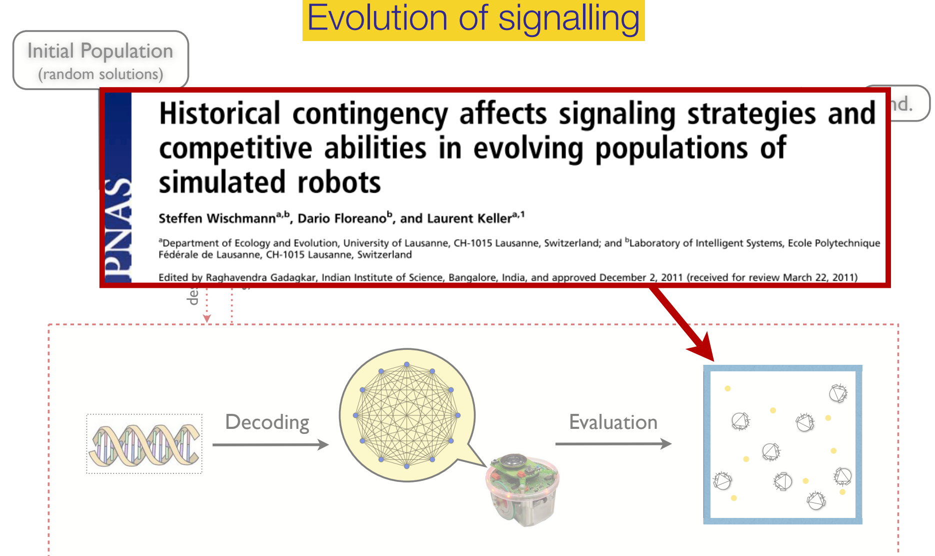

PNAS

Historical contingency affects signaling strategies and competitive abilities in evolving populations of simulated robots

Steffen Wischmann[a,b], Dario Floreano[b], and Laurent Keller[a,1]

[a]Department of Ecology and Evolution, University of Lausanne, CH-1015 Lausanne, Switzerland; and [b]Laboratory of Intelligent Systems, Ecole Polytechnique Fédérale de Lausanne, CH-1015 Lausanne, Switzerland

Edited by Raghavendra Gadagkar, Indian Institute of Science, Bangalore, India, and approved December 2, 2011 (received for review March 22, 2011)

Evolution of swarming

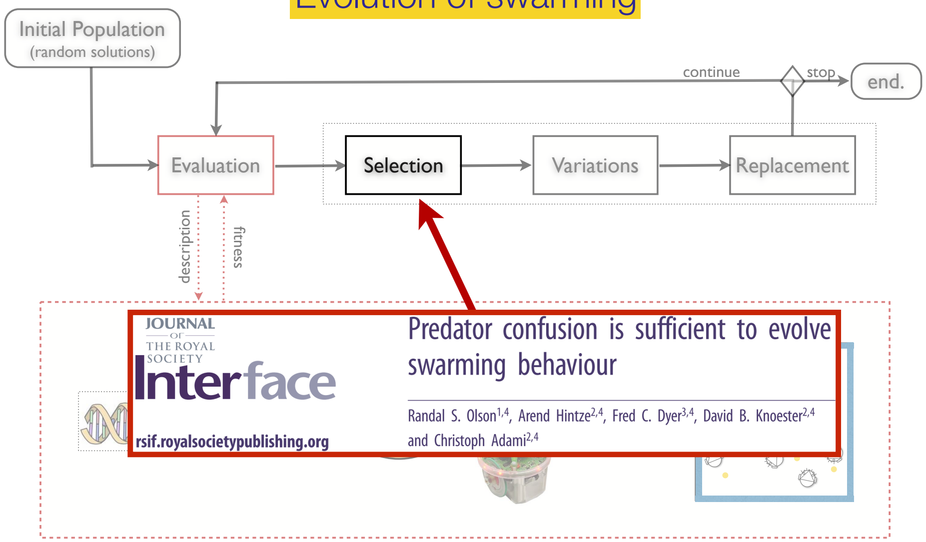

Take-home message for individual-based modelling with ER

Mechanistic constraints plays an important role in natural evolution

(Waibel et al., 2011, Plos Biology)
(Bernard et al., 2016, Plos Computational Biology)
(Olson et al., 2013, GECCO)
(Olson et al., 2013, Royal society Interface)
(...)

ER for collective robotics

What about the future?

- What is missing?
  - Lack of impact in other communities (AAMAS, DARS, etc.)
  - We don't fully understand the evolutionary dynamics of coll. sys.
  - We don't understand the physics of coll. systems

Formalising problems with evolutionary game theory

Partner hunts / You hunt	Hare	Stag
Hare	1 ; 1	1 ; 0
Stag	0 ; 1	2 ; 2

Payoff-dominant equilibrium

Risk-dominant equilibrium

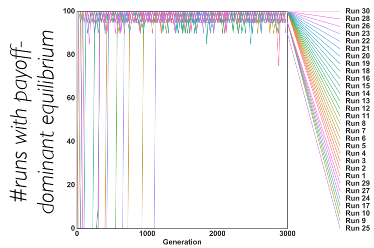

What theory predicts

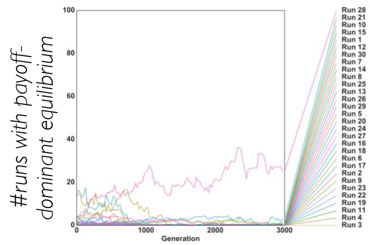

What ER reveals

A. Bernard, J.B. André, N. Bredeche. To Cooperate or Not to Cooperate: Why Behavioural Mechanisms Matter. PloS Computational Biology (2016)

Morphological computation and active matter

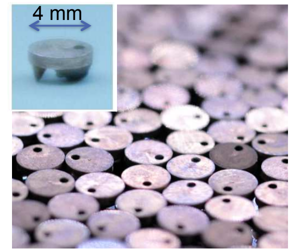

~1000 walking grains

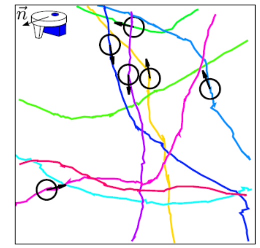

individual motion

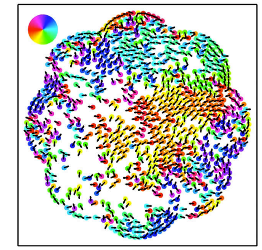

collective motion

J. Deseigne, S. Leonard, O. Dauchot H. Chaté, Vibrated polar disks: spontaneous motion, binary collisions and collective dynamics, Soft Matter, 8, 5629 (2012).

Take-home message for the future

The fields of evolutionary game theory and active matter are strongly relevant to our field

Conclusion

- **Context**: collective adaptive systems in open environments
- **Two scopes**:
  - ER as a design tool for making artificial systems
  - ER as a modelling tool for understanding natural systems
- **Two methods**:
  - off-line optimisation problem: use classic ER
  - distributed on-line learning: use embodied ER

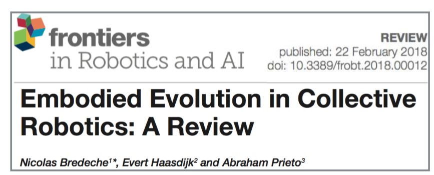

frontiers in Robotics and AI

REVIEW

published: 22 February 2018
doi: 10.3389/frobt.2018.00012

Embodied Evolution in Collective Robotics: A Review

Nicolas Bredeche[1], Evert Haasdijk[2] and Abraham Prieto[3]*

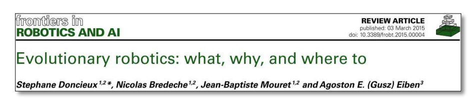

frontiers in ROBOTICS AND AI

REVIEW ARTICLE

Evolutionary robotics: what, why, and where to

Stephane Doncieux[1,2], Nicolas Bredeche[1,2], Jean-Baptiste Mouret[1,2] and Agoston E. (Gusz) Eiben[3]*

Bibliography

Preamble: this is not an exhaustive list, but a list of pointers to start with.

Take-home message #1
Waibel, Markus, Laurent Keller, and Dario Floreano. (2009) Genetic team composition and level of selection in the evolution of cooperation. TEC
GS Nitschke, AE Eiben, MC Schut (2012) Evolving team behaviors with specialization. GPEM.
Lichocki, Paweł, et al. (2013) "Evolving team compositions by agent swapping." IEEE TEC
Gomes, Jorge, Pedro Mariano, Anders Lyhne Christensen (2015) Cooperative coevolution of partially heterogeneous multiagent systems. AAMAS
Tuci, Elio, and Vito Trianni (2014). On the evolution of homogeneous two-robot teams: clonal versus aclonal approaches. Neural Computing and Applications
Tuci, Elio, and Vito Trianni (2012) On the Evolution of Homogeneous Multi-robot Teams: Clonal versus Aclonal Approach. SAB
A Bernard, JB André, N Bredeche (2016) Evolving specialisation in a pop. of heterog. robots: the challenge of bootstrapping and maintaining genotypic polymorphism. ALIFE
A Bernard, JB André, N Bredeche (2015) Evolution of cooperation in evolutionary robotics: the tradeoff between evolvability and efficiency. ECAL

Take-home message #2
Bredeche, Nicolas, and Jean-Marc Montanier (2010) Environment-driven embodied evolution in a population of autonomous agents, PPSN.
Pérez, Iñaki Fernández, Amine Boumaza, François Charpillet (2015) Comparison of selection methods in on-line distributed evolutionary robotics. ALIFE.
Hart, Emma, Andreas Steyven, Ben Paechter (2015) Improving survivability in environment-driven distributed evolutionary algorithms through explicit relative fitness and fitness proportionate communication. GECCO.
Steyven, Andreas, Emma Hart, Ben Paechter (2016) Understanding Environmental Influence in an Open-Ended Evolutionary Algorithm. PPSN.
E Haasdijk, N Bredeche, AE Eiben (2014) Combining environment-driven adaptation and task-driven optimisation in evolutionary robotics. PloS one.
JM Montanier, S Carrignon, N Bredeche (2016) Behavioral specialization in embodied evol. robotics: Why so Difficult? Frontiers in Robotics and AI.

Take-home message #3
M Waibel, D Floreano, L Keller (2011) A quantitative test of Hamilton's rule for the evolution of altruism. PLoS biology, 2011.
A Bernard, JB André, N Bredeche (2016) To cooperate or not to cooperate: why behavioural mechanisms matter. PLoS computational biology 12 (5), e1004886
RS Olson, A Hintze, FC Dyer, DB Knoester, C Adami (2013) Predator confusion is sufficient to evolve swarming behaviour. Journal of the Royal Soc. Interface 10 (85), 2013
RS Olson, DB Knoester, C Adami (2013) Critical Interplay Between Density-dependent Predation and Evolution of the Selfish Herd, GECCO 2013

Take-home message #4
Nisan et al. Algorithmic game theory. Cambridge University press (2007)
Adami et al. Evolutionary game theory using agent-based methods. Physics of Life Reviews 19 (2016)
G. De Magistris, D. Marenduzzo. An introduction to the physics of active matter. Physica A 418 (2015)
S. Ramaswamy. The Mechanics and Statistics of Active Matter. Annual Review of Condensed Matter Physics, Vol. 1 (2010)

nicolas.bredeche@upmc.fr

Evolutionary Robotics tutorial
Conclusion

N. Bredeche, S. Doncieux, J.-B. Mouret

Some software tools

- SFERES2: https://github.com/sferes2
  - Software framework in modern C++
  - As fast as specific code
  - Modules available to evolve robots, examples:
    - Neural network module: https://github.com/sferes2/nn2
    - Simple simulation of a 2-wheeled robot: https://github.com/sferes2/fastsim
  - Code of many experiments on http://pages.isir.upmc.fr/evorob_db
  - Basic experiments to starting playing with ER:
    - Two-wheeled robot maze navigation & obstacle avoidance: https://github.com/doncieux/navigation
    - Collect ball experiment: https://github.com/doncieux/collectball
- NEAT & HyperNEAT packages: http://eplex.cs.ucf.edu/neat_software/

Mouret, J.-B. and Doncieux, S. (2010). **SFERESv2: Evolvin' in the Multi-Core World.** WCCI 2010 IEEE World Congress on Computational Intelligence, Congress on Evolutionary Computation (CEC). Pages 4079--4086.

Landmark papers

- Floreano, D., and F. Mondada. (1996) **Evolution of homing navigation in a real mobile robot.** IEEE Transactions on Systems, Man, and Cybernetics, Part B : Cybernetics 26.3 (1996) : 396-407.
- Lipson, H., and J B. Pollack. (2000) **Automatic design and manufacture of robotic lifeforms.** Nature 406.6799 : 974-978.
- Watson, R. A., S. G. Ficici, and J. B. Pollack. (2002) **Embodied evolution : Distributing an evolutionary algorithm in a population of robots.** Robotics and Autonomous Systems 39, no. 1 : 1-18.
- Hornby, G. S., S. Takamura, T. Yamamoto, and M. Fujita (2005). **Autonomous evolution of dynamic gaits with two quadruped robots.** IEEE Transactions on Robotics, 21, no. 3 : 402-410.
- Bongard, J., V. Zykov, and H. Lipson (2006). **Resilient machines through continuous self-modeling.** Science 314.5802 :1118-1121.
- Lehman, J., and Kenneth O. Stanley (2011). **Abandoning objectives : Evolution through the search for novelty alone.** Evolutionary computation 19.2 (2011) : 189-223.
- Cully, A., Clune, J., Tarapore, D., & Mouret, J. B. (2015). **Robots that can adapt like animals**. Nature, 521(7553), 503-507.

Recent reviews & introduction papers

- Cully, A., & Demiris, Y. (2018). **Quality and diversity optimization: A unifying modular framework**. IEEE Transactions on Evolutionary Computation, 22(2), 245-259.
- Doncieux S, Bredeche N, Mouret J-B & Eiben AE (2015). **Evolutionary robotics : what, why, and where to.** Front. Robot. AI 2 :4. doi : 10.3389/frobt.2015.00004.
- Doncieux, S. and Mouret, J.-B. (2014). **Beyond black-box optimization : a review of selective pressures for evolutionary robotics.** Evolutionary Intelligence, Springer Berlin Heidelberg, publisher. Vol 7 No 2 Pages 71-93.
- Bongard, J. C. (2013) **Evolutionary robotics** Communications of the ACM 56.8 : 74-83
- Nelson, A. L., Barlow, G. J., & Doitsidis, L. (2009). **Fitness functions in evolutionary robotics: A survey and analysis**. Robotics and Autonomous Systems, 57(4), 345-370.
- Floreano, D., Dürr, P., & Mattiussi, C. (2008). **Neuroevolution : from architectures to learning. Evolutionary Intelligence,** 1(1), 47–62. Evolutionary Intelligence, 1(1), 47–62.

Evolutionary Robotics Community

- **Dedicated conferences/tracks:**
  - Complex Systems track in ACM Genetic and Evolutionary Computation Conference (**GECCO**)
  - Evolutionary robotics track at **IEEE-WCCI** (World Congress on Computational Intelligence)/**IEEE-CEC** (Congress on Evolutionary Computation)
  - EvoROBOT track in EvoSTAR
- **Dedicated journals:**
  - Frontiers in Robotics and AI, Evolutionary Robotics specialty section
  - Evolutionary Intelligence, Springer
- **Mailing lists:**
  - General: evoderob@listes.upmc.fr
  - On NEAT: neat@yahoogroups.com

Resibots

Jean-Baptiste Mouret