



Flow navigation by smart particles via Reinforcement Learning

Luca Biferale
Dept. Physics, INFN & CAST
University of Rome 'Tor Vergata'
biferale@roma2.infn.it
PIML2018 SANTA FE




CREDITS: SIMONA COLABRESE (TOR VERGATA UNIV. ROME-IT); ANTONIO CELANI (ICTP TRIESTE-IT); KRISTIAN GUSTAVSSON (GOTHEBORG UNIV. SWEDEN)





- PARTICLES IN COMPLEX FLOWS I: **SMART INERTIAL PARTICLES**
- PARTICLES IN COMPLEX FLOWS II: **SMART MICROSWIMMERS**

SWIMMING DIRECTION CONTROL



- **Flow navigation by smart microswimmers via reinforcement learning**

S Colabrese, K Gustavsson, A Celani, L Biferale
Physical Review Letters 118 (15), 158004, 2017


- **Smart Inertial Particles**

S Colabrese, K Gustavsson, A Celani, L Biferale
arXiv preprint arXiv:1711.05853, 2017

- **Finding efficient swimming strategies in a three-dimensional chaotic flow by reinforcement learning**

K Gustavsson, L Biferale, A Celani, S Colabrese
The European Physical Journal E 40 (12), 110, 2017

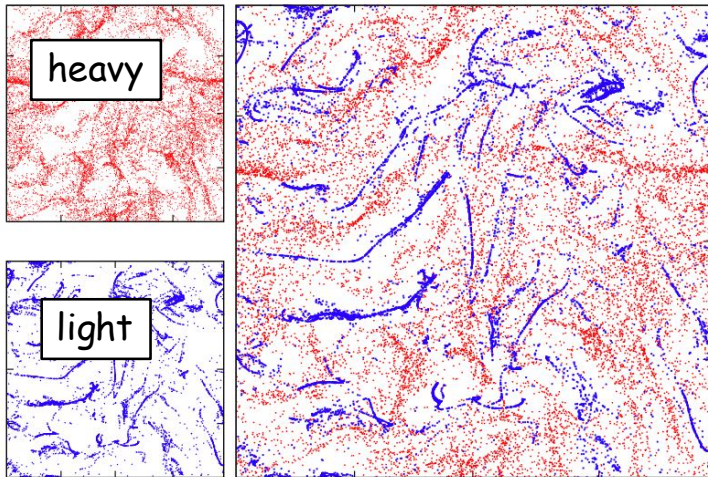
PARTICLES IN COMPLEX FLOWS I: INERTIAL PARTICLES



$$\frac{d\mathbf{X}}{dt} = \mathbf{V}$$

$$\frac{d\mathbf{V}}{dt} = \beta \frac{D\mathbf{u}(\mathbf{X},t)}{Dt} + \frac{\mathbf{u}(\mathbf{X},t) - \mathbf{V}}{\tau}$$

$$\partial_t \mathbf{u} + (\mathbf{u} \cdot \partial) \mathbf{u} = -\partial P + \nu \nabla^2 \mathbf{u}$$



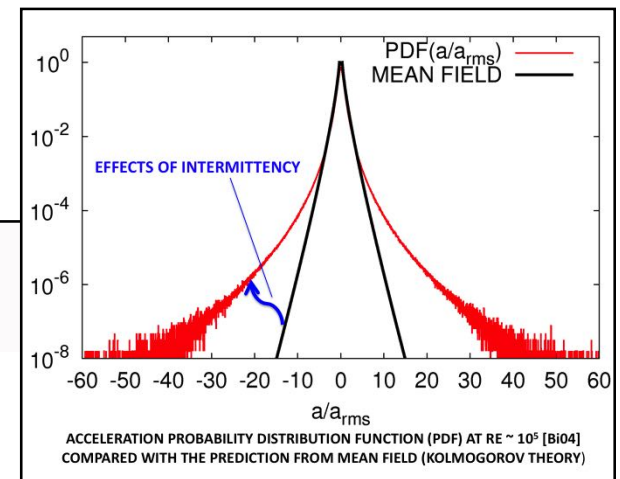
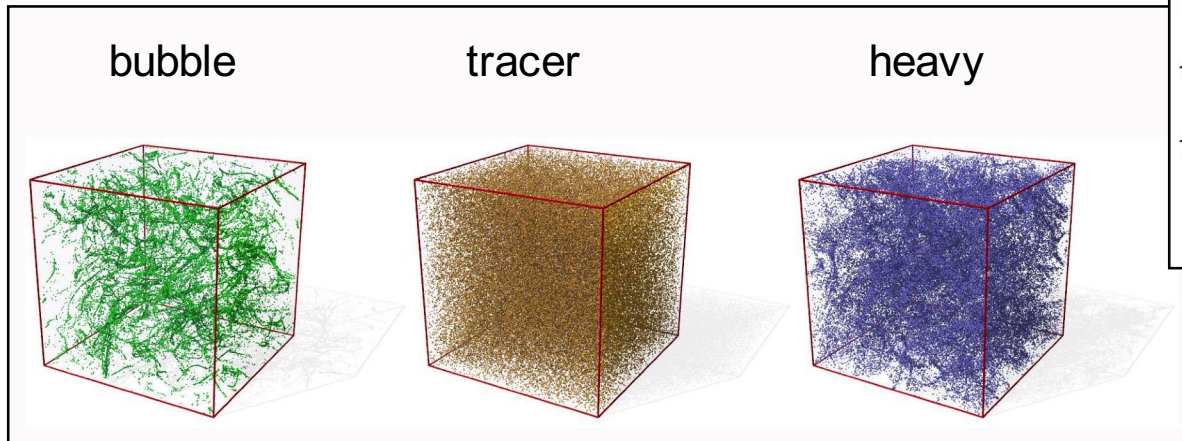



$$\beta = \frac{3\rho_f}{\rho_f + 2\rho_p}$$

$$\tau = \frac{a^2}{3\nu\beta}$$

$\beta < 1$ heavy particles
 $\beta > 1$ light particles

Drag: **Stokes Time**







Preferential concentration!
 Light(heavy) particles accumulate
 inside(outside) highly vortical regions



Particle trapping in three-dimensional fully developed turbulence


L.B., G Boffetta, A Celani, A Lanotte, F Toschi

Physics of Fluids 17 (2), 021701



Lagrangian properties of particles in turbulence
 F Toschi, E Bodenschatz
 Annual Review of Fluid Mechanics 41, 375-404 (2009)


Coherent structures and extreme events in rotating multiphase turbulent flows L.
 L.B., F Bonaccorso, IM Mazzitelli, MAT van Hinsberg, AS Lanotte, ...
 Physical Review X 6 (4), 041036 (2016)



Smart inertial particle

$$R = \Omega^3$$

state S_n



ρ (densities)


action A_n

OBSERVATION:
DISCRETIZED VORTICITY LEVELS

$$\frac{d\mathbf{X}}{dt} = \mathbf{V}$$

$$\frac{d\mathbf{V}}{dt} = \beta \frac{D\mathbf{u}(\mathbf{X},t)}{Dt} + \frac{\mathbf{u}(\mathbf{X},t) - \mathbf{V}}{\tau}$$

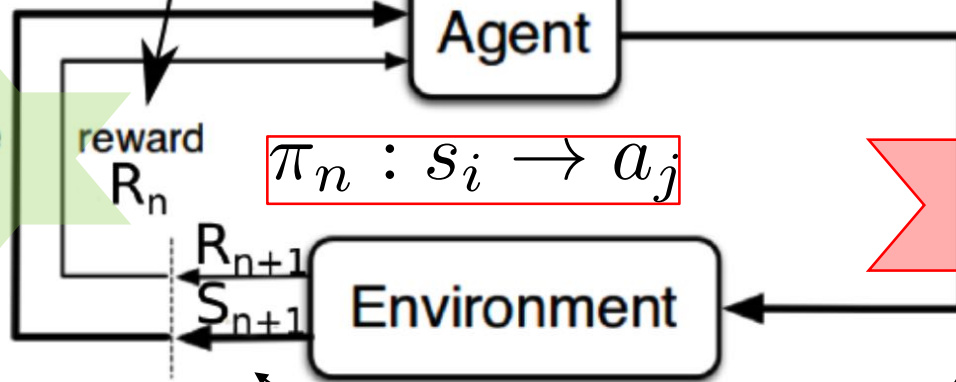
$$\pi_n \rightarrow \pi_{n+1} \rightarrow \dots \rightarrow \pi_{opt}$$



Smart inertial particle

$$R = \Omega^3$$

state S_n



;(densities)

OBSERVATION:
DISCRETIZED VORTICITY LEVELS

$$\frac{d\mathbf{X}}{dt} = \mathbf{V}$$

$$\frac{d\mathbf{V}}{dt} = \beta \frac{D\mathbf{u}(\mathbf{X},t)}{Dt} + \frac{\mathbf{u}(\mathbf{X},t) - \mathbf{V}}{\tau}$$

$$\pi_n \rightarrow \pi_{n+1} \rightarrow \dots \rightarrow \pi_{opt}$$

TRAINING: Q-LEARNING ALGORITHM

QUALITY MATRIX AT STEP $n \rightarrow Q_n(a_j, s_i)$

EXPECTED DISCOUNTED FUTURE RETURN IF ACTION a_j is taken after observation of state s_i

$$Q_n(s_i, a_j) = R_n + \gamma R_{n+1} + \gamma^2 R_{n+2} + \gamma^3 R_{n+3} + \dots = \sum_{t=n}^{\infty} \gamma^t R_t$$

MYOPIC $\rightarrow \gamma = 0$
FAR-SIGHTED $\rightarrow \gamma = 1$


GREEDY POLICY AT STEP n :

$$\pi_n : a = \underset{a'}{\operatorname{arg\,max}} Q_n(a', s)$$

$$\begin{array}{l} s_1 \\ s_2 \\ s_3 \end{array} \begin{bmatrix} 1.2 & 0.3 & 0.1 \\ 2.2 & 4.3 & 10.1 \\ 2.0 & 8.1 & 2.0 \end{bmatrix} \begin{array}{l} s_1 \xrightarrow{\pi_n} a_1 \\ s_2 \xrightarrow{\pi_n} a_3 \\ s_3 \xrightarrow{\pi_n} a_2 \end{array}$$

$a_1 \quad a_2 \quad a_3$

Learning through experience: Q-learning algorithm




$$Q(s, a) \leftarrow Q(s, a) + \alpha [R' + \gamma \max_{a'} Q(s', a') - Q(s, a)]$$


learning rate

$0 \leq \alpha < 1$ fixed or time dependent

discount factor

$0 \leq \gamma < 1$

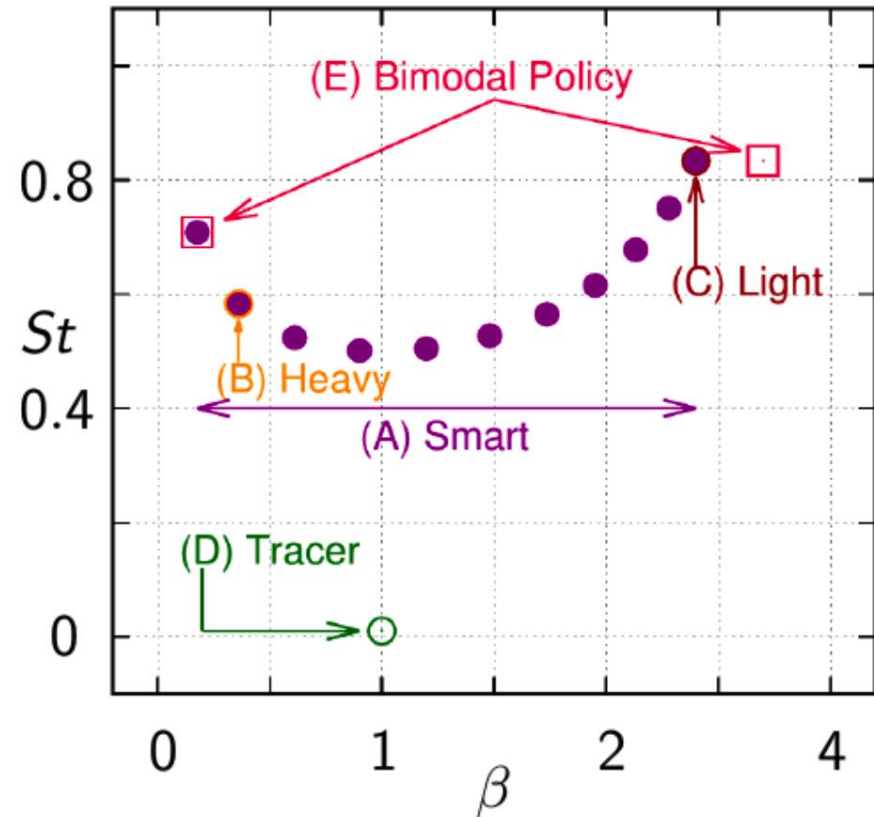




$$Ns = 21$$

$$Na = 11$$

CHANGING RADIUS b_n

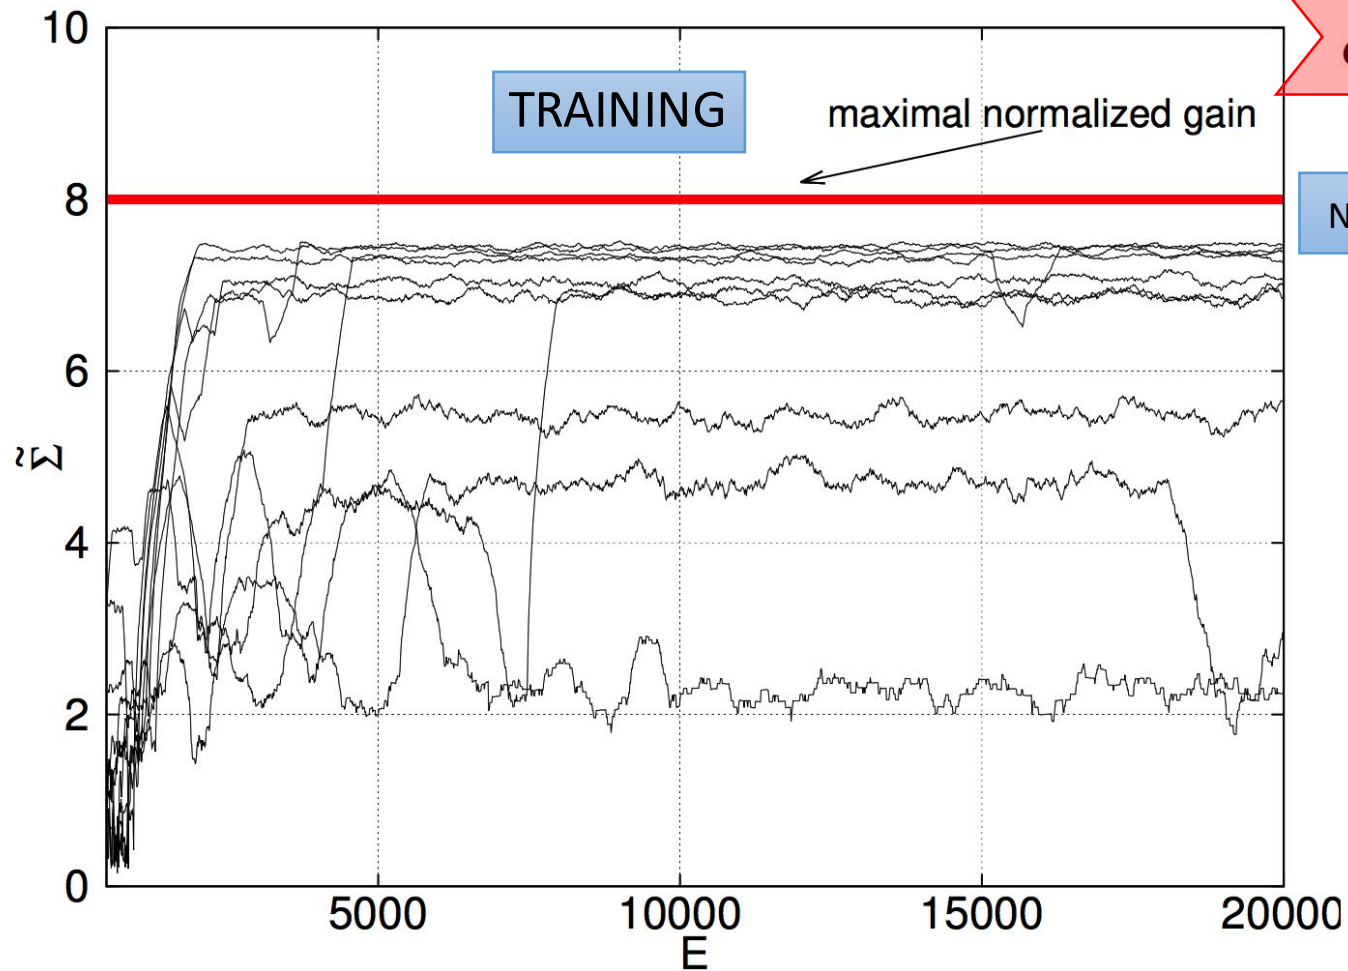


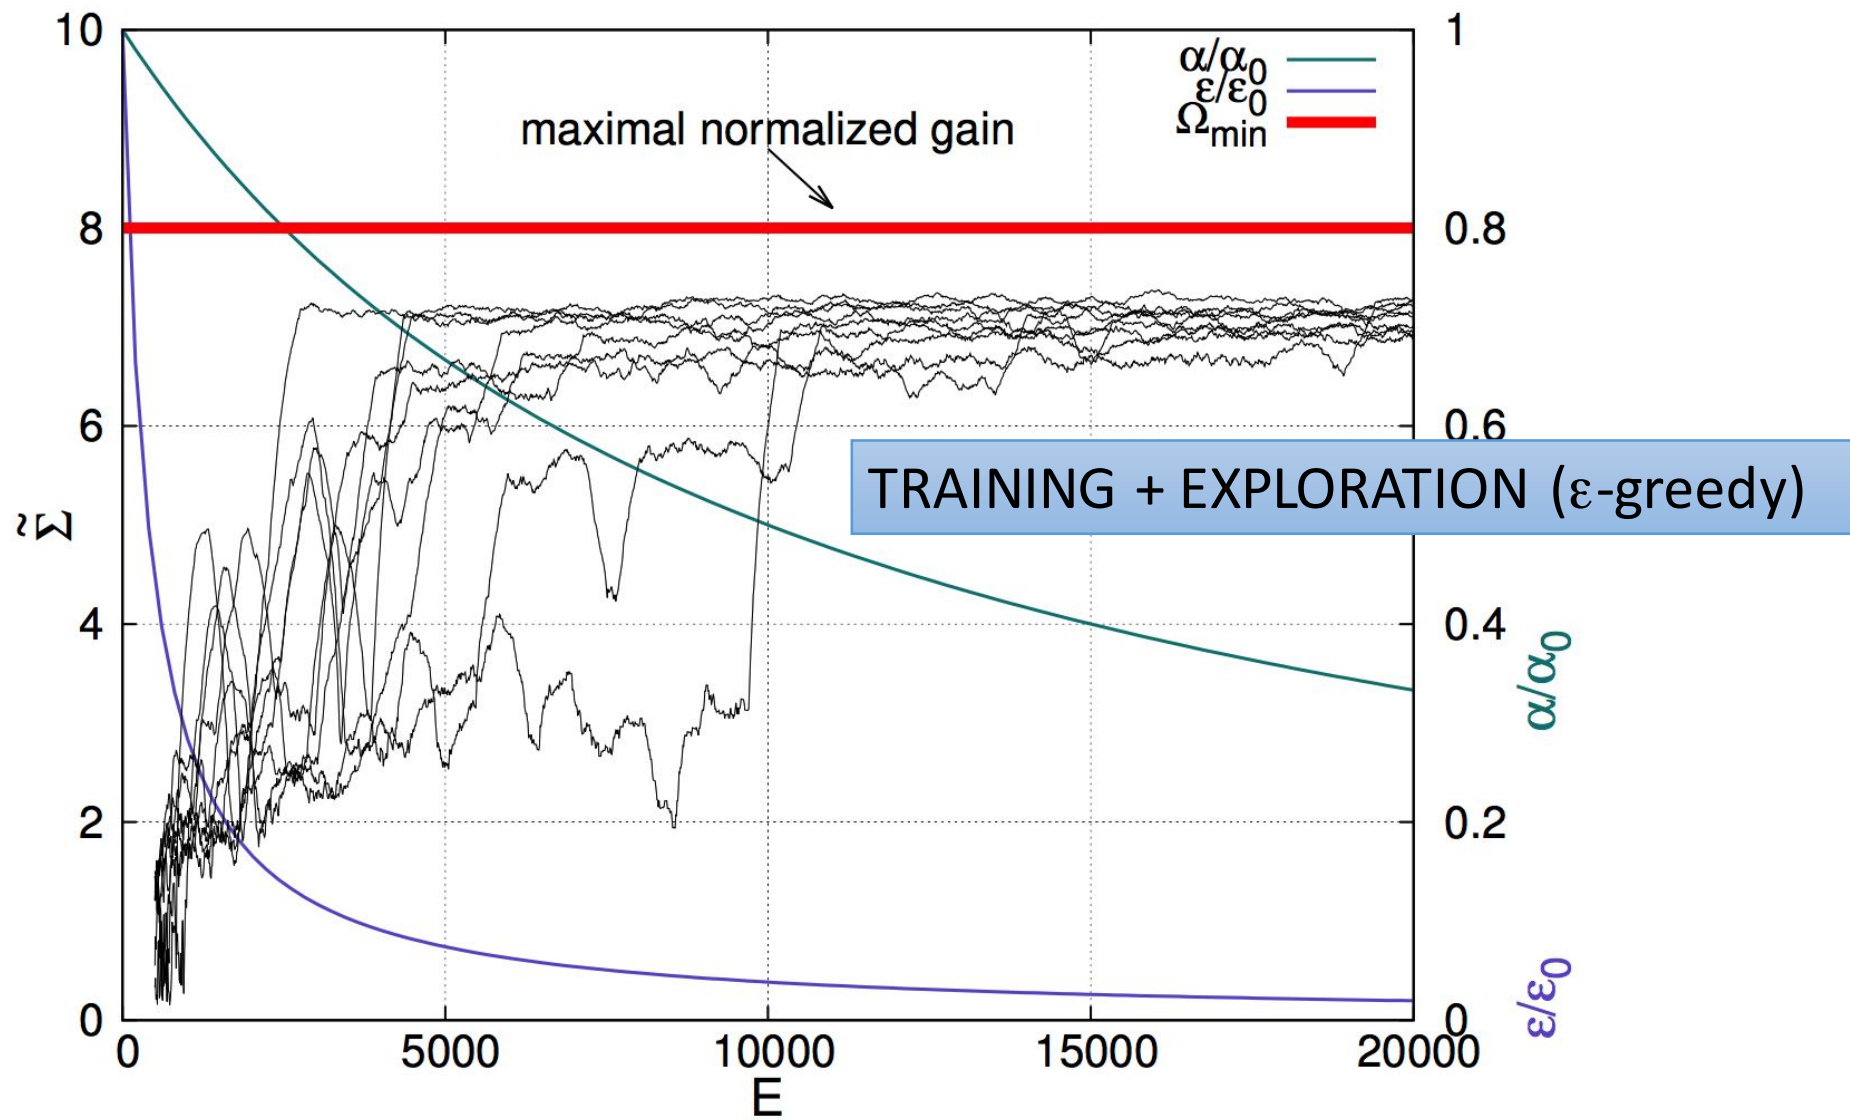
$$\beta \rightarrow \beta(b_n)$$

$$St \rightarrow St(b_n)$$


Comparison with naive behaviors

Learning gain $\tilde{\Sigma}(E) = \sqrt[3]{\frac{\sum_{n=1}^N R_n}{N}}$







Smart particles (A)



PDFs






Policies






EXAM





Heavy (B)



Light (C)




Bimodal (E)




OPTIMAL ACTIONS

TIME DEPENDENT FLOW



Is the algorithm robust for more complex flows?
Steady flow \rightarrow 4 oscillating vortex flow with different phases but same angular frequency





Asymmetric ABC flow

$$\mathbf{u}(\mathbf{x}) = (C \cos y + A \sin z, A \cos z + B \sin x, B \cos x + C \sin y) \quad [4A=2B=C=1]$$

Task: Optimize long-term vorticity $|\boldsymbol{\Omega}|$ by perception of Ω_z or Ω_x

-different setup: St fixed




Light particles distribute on minor vortices

Smart particle learns to target principal vortices


Evaluation of the algorithm

The global optimal policy is not known

Brute force approach by testing all possible Q -matrices [$2^{11} = 2048$]
(reducing the number of states and actions)




Microswimmers can bias their motion in order to achieve a biologically relevant goal



Particles can learn by experience to navigate by using Reinforcement Learning

Case study: Gyrotactic Phytoplankton




- ▶ large diversity of forms
- ▶ primary producers in oceans
- ▶ ≈50% photosynthetic activity on Earth
- ▶ up to 10^4 per milliliter of water
- ▶ at the bottom of marine food web
- ▶ can form Harmful (toxic) Algal Bloom
- ▶ patchiness at different scales
- ▶ many species are able to swim, e.g. 90% toxic algae are able to swim


gyrotactic microalgae




Heterosigma akashiwo





Dunaliella tertiolecta



Chlamydomonas reinhardtii






Free Particle parameters

If $\Phi \rightarrow 0, \Psi \rightarrow \infty$
passive swimmer

Flow properties


Direction to be modified
depending on $\boldsymbol{\omega}$

Gyrotactic focusing




J.O. Kessler Nature (1985)

Gyrotactic trapping



Microfluidic experiments



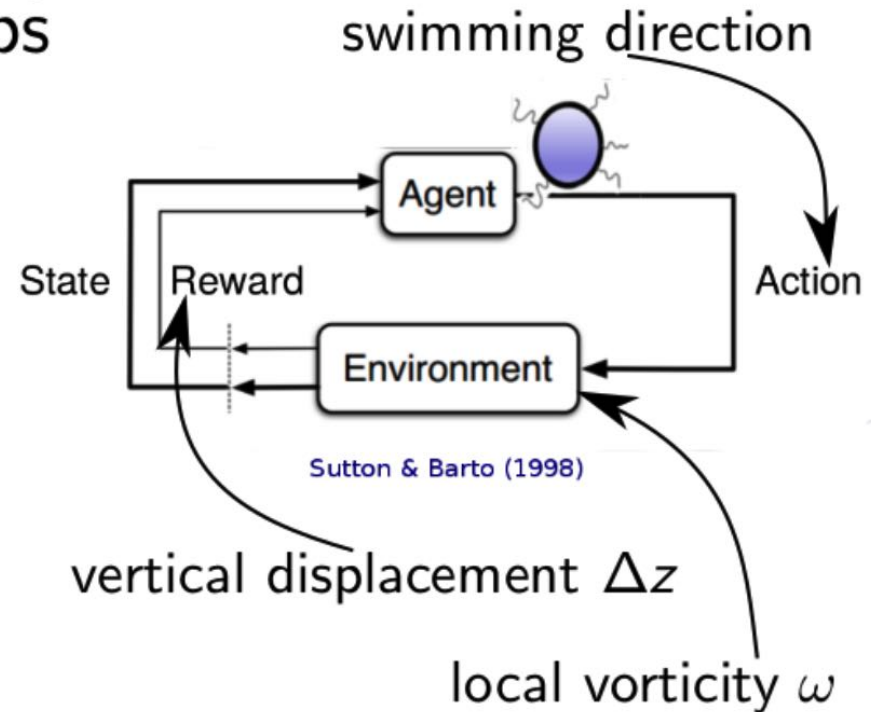

Durham et al Science 2009

Clustering and turbophoresis in a shear flow without walls
F De Lillo, M Cencini, S Musacchio, G Boffetta
Physics of Fluids 28 (3), 035104 (2016)

Goal: going up as efficient as possible

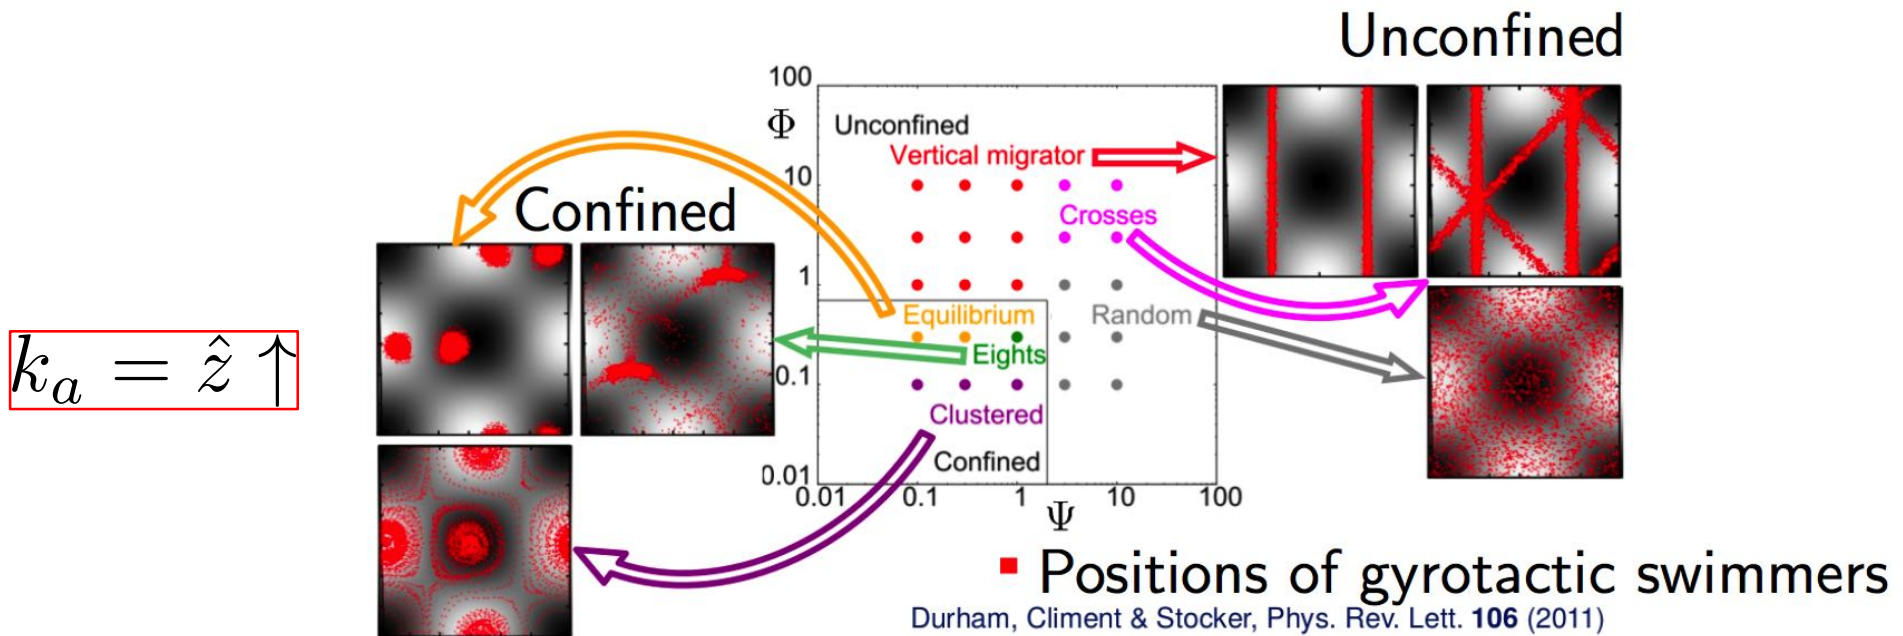
Difficulty: Hydrodynamic traps

How to choose swimming direction given the underlying vorticity?



Reinforcement Learning



If k_a is fixed, the particle can be trapped in confined region



We have to find the optimal $k_a \rightarrow k_a(t)$ to swim upward

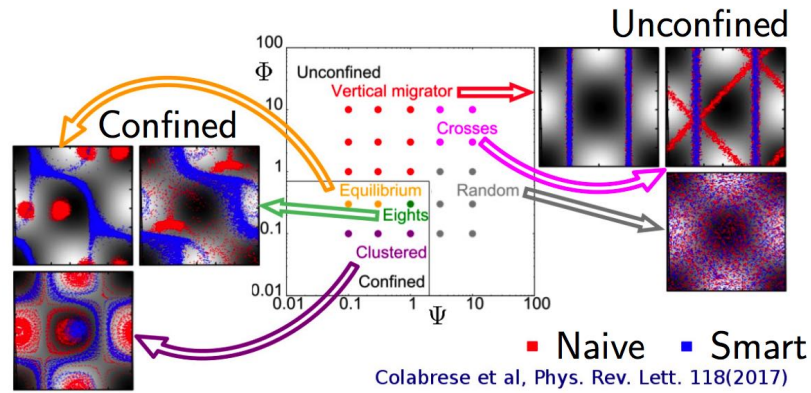
Relative gain

$$\Sigma = \frac{\langle z(T) - z(0) \rangle_{k_a(t)}}{\langle z(T) - z(0) \rangle_{k_a}} - 1$$

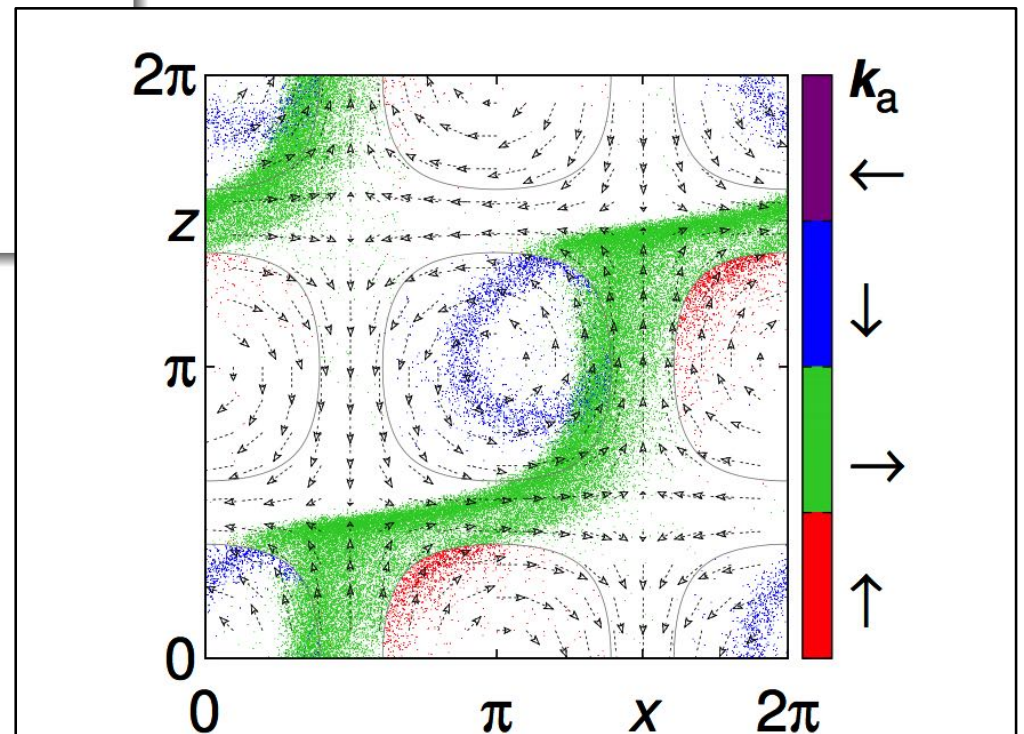


Trajectories after RL training

Comparison between *smart* and *naive* microswimmers trajectories for different parameters values




OPTIMAL STRATEGY



Robustness to perturbation of the flow

Using strategies found for the unperturbed flow ($\beta = 1$) the particle outperform the naive particle down to moderate β




Time-dependent flow


Robustness even when tracers might have chaotic evolution

Taylor-Green flow with periodic phase shift $\Delta_t = b \cos(u_0 t)$

$$\mathbf{u} = \frac{u_0}{2} [-\cos(x - \Delta_t) \sin(z - \Delta_t) \hat{\mathbf{x}} + \sin(x - \Delta_t) \cos(z - \Delta_t) \hat{\mathbf{z}}]$$



Weak time dependence




Stronger time dependence

Strategy learns in the time-dependent flow


Going up as efficient as possible in ABC flow

Reddy, G., Celani, A., Sejnowski, T. J., & Vergassola, M. (2016). Learning to soar in turbulent environments. *Proceedings of the National Academy of Sciences*, 201606075.

Comparison between **smart** and **naive** microswimmers in a lagrangianly chaotic flow




$$(u, v, w) = \cos x + C \sin y$$



Example of Resulting Strategy

Optimal actions taken during the ascent through an elevator



Conclusions

CREDITS: SIMONA COLABRESE (TOR VERGATA UNIV. ROME-IT); ANTONIO CELANI (ICTP TRIESTE-IT); KRISTIAN GUSTAVSSON (GOTHEBORG UNIV. SWEDEN)

- Reinforcement learning approach is successful to develop strategies for inertial particles to sample specific vortical structures and for microswimmers to ascend two- and three-dimensional vortex flows
- Learnt strategies are robust to spatial or time perturbations of the flow

- Flow navigation by smart microswimmers via reinforcement learning

S Colabrese, K Gustavsson, A Celani, L Biferale
Physical Review Letters 118 (15), 158004

-Smart Inertial Particles

S Colabrese, K Gustavsson, A Celani, L Biferale
arXiv preprint arXiv:1711.05853

- Finding efficient swimming strategies in a three-dimensional chaotic flow by reinforcement learning

K Gustavsson, L Biferale, A Celani, S Colabrese
The European Physical Journal E 40 (12), 110