Turbulence.zip

Predrag Cvitanović

2012 IEEE International Symposium on Information Theory, Cambridge MA

July 6, 2012

Outline

knowing when to stop

[click here for an example of a fluid in motion]

need the 3D velocity field at every (x, y, z)!

motions of fluids : require ∞ bits?

numerical simulations track millions of computational degrees of freedom; observations, from laboratory to satellite, stream terabytes of data, but how much information is there in all of this?

knowing when to stop

motions of fluids : require ∞ bits??

that cannot be right...

knowing when to stop

Science originates from curiosity and bad eyesight.

— Bernard de Fontenelle,
Entretiens sur la Pluralité des Mondes Habités

in practice

every physical problem is coarse partitioned and finite

noise rules the state space

- any physical system experiences (background, observational, intrinsic, measurement, $\cdots$) noise
- any numerical computation is a noisy process due to the finite precision of computation
- any set of dynamical equations models nature up to a given finite accuracy, since degrees of freedom are always neglected
- any prediction only needs to be computed to a desired finite accuracy

mathematician's idealized state space

a manifold $\mathcal{M} \in \mathbb{R}^d$: d continuous numbers determine the state of the system $x \in \mathcal{M}$

noise-limited state space

a 'grid' $\mathcal{M}'$: N discrete states of the system $a \in \mathcal{M}'$, one for each noise covariance ellipsoid Δ_a

noise limited state space partitions

noise limited cell

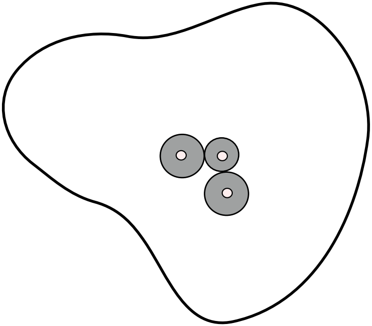

a resolvable neighborhood is no smaller than a ball whose radius is the noise amplitude

noise limited partition grid

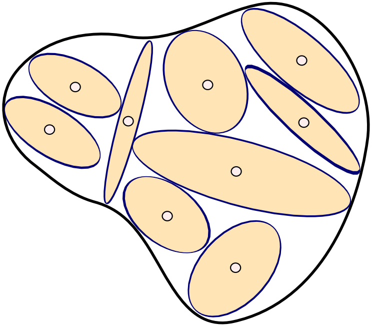

state space noise-partitioned into neighborhoods indicated by their centers

the centers = prototypes in a *vector quantization* scheme for a compressive encoding of state space

entropy

entropy of a cell

cell is described by a
d-dimensional Gaussian,
covariance matrix Δ_a

$p(a) \propto \text{volume} \propto |\Delta_a|^{-1/2}$,

$\text{entropy} = \frac{1}{2} \ln \left\{ (2\pi e)^d |\Delta_a| \right\}$

global entropy of a partition

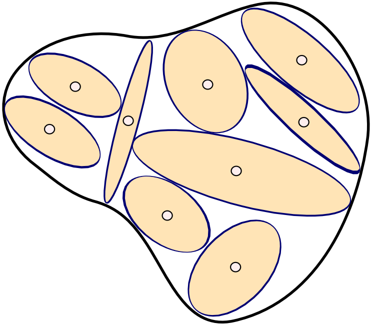

$$H(\mathcal{M}') = - \sum_{a \in \mathcal{M}'} p(a) \ln p(a)$$

all entropy is local

define

$$p(x,a) = p(x|a)p(a)$$

mutual information

$$\begin{aligned} I(\mathcal{M},\mathcal{M}') &= \int_{\mathcal{M}} dx \sum_{a\in\mathcal{M}'} p(x,a) \ln \frac{p(x,a)}{p(x)p(a)} \\ &= \sum_{\mathcal{M}'} p(a) \int_{\mathcal{M}} dx\, p(x|a) \ln p(x|a) - \int_{\mathcal{M}} dx\, p(x) \ln p(x) \\ &= H(\mathcal{M}) - \sum_{a\in\mathcal{M}'} p(a) H(\mathcal{M}|a) \end{aligned}$$

measures how much we know about $\mathcal{M}$, given the grid $\mathcal{M}'$ and the Gaussian local entropy

$$H(\mathcal{M}|a) = \frac{1}{2} \ln \left\{ (2\pi e)^d |\Delta_a| \right\}$$

rest of the talk: we show you how to compute Δ_a

dynamics + noise: unique coarse-grained partition

reasonable to assume that the ‘external’ noise Δ

limits the resolution that can be attained in partitioning the state space

dynamics + noise: unique coarse-grained partition

reasonable to assume that the ‘external’ noise Δ

limits the resolution that can be attained in partitioning the state space

is uniform, leading to a uniform grid partitioning of the state space

dynamics + noise: unique coarse-grained partition

reasonable to assume that the 'external' noise Δ

limits the resolution that can be attained in partitioning the state space

in dynamics, this is wrong!

noise has memory

dynamics + noise: unique coarse-grained partition

noise memory

accumulated noise along dynamical trajectories
always coarsens the partition

dynamics + noise: unique coarse-grained partition

noise memory

accumulated noise along dynamical trajectories
always coarsens the partition

that is good, because

dynamics + noise determine

the **finest attainable** partition

optimal partition hypothesis

- the best of all possible state space partitions
- optimal for the given dynamical system, the given noise

([1])

[1] D. Lippolis and P. Cvitanović, arXiv.org:0902.4269; arXiv.org:1206.5506

devil is in the details

in what follows: fluid dynamics as an example

for fluids, have equations: can compute the optimal partition

(this is not a talk about fluid dynamics)

since 1822 we have Navier-Stokes equations

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v} \cdot \nabla)\mathbf{v} = \frac{1}{R}\nabla^2 \mathbf{v} - \nabla p + \mathbf{f}\,, \qquad \nabla \cdot \mathbf{v} = 0,$$

velocity field $\mathbf{v} \in \mathbb{R}^3$; pressure field p ; driving force $\mathbf{f}$

since 1883 Osborne Reynolds experiments

the most fundamental outstanding problem of classical physics

large Reynolds number R: turbulence!

what is it to you? nasty weather...

numerical challenges

computation of turbulent solutions

requires 3-dimensional volume discretization
$\rightarrow$ integration of 10^4-10^6 coupled ordinary differential equations

challenging, but today possible

numerical challenges

typical simulation

each instant of the flow $>$ Megabytes
a video of the flow $>$ Gigabytes

example : pipe flow

amazing data! amazing numerics!

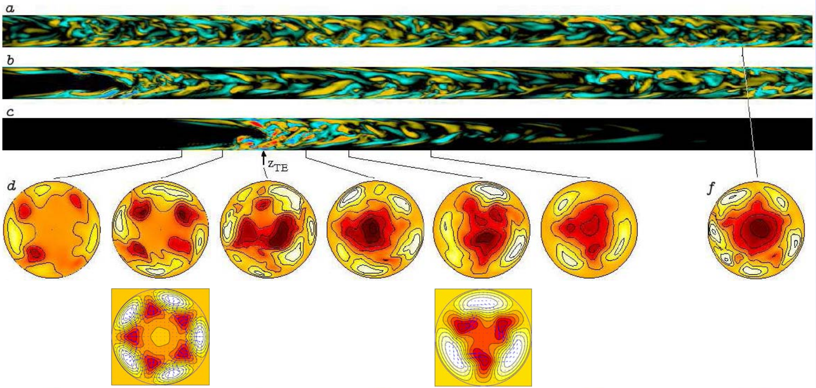

- here each instant of the flow $\approx 2.5\,\mathrm{MB}$
- videos of the flow $\approx$ GBs

turbulence.zip

or ‘equation assisted’ data compression:

replace the ∞ of turbulent videos by the best possible

small finite set

of videos encoding all physically distinct motions of the turbulent fluid

dynamical system

state space

a manifold $\mathcal{M} \in \mathbb{R}^d$: d numbers determine the state of the system

representative point

$x(t) \in \mathcal{M}$
a state of physical system at instant in time

example of a representative point

$x(t) \in \mathcal{M}$, $d = \infty$
a state of turbulent pipe flow at instant in time

Stereoscopic Particle Image Velocimetry $\rightarrow$ 3-d velocity field over the entire pipe[2]

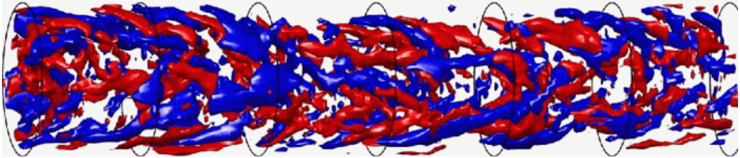

[2] Casimir W.H. van Doorne (PhD thesis, Delft 2004)

dynamics

map $f^t(x_0)$ = representative point time t later

evolution in time

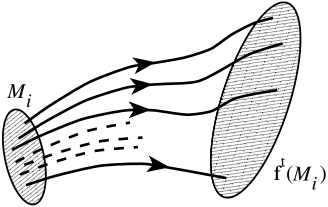

f^t maps a region $\mathcal{M}_i$ of the state space into the region $f^t(\mathcal{M}_i)$

dynamical system

the pair $(\mathcal{M}, f)$

the problem

enumerate, classify all solutions of $(\mathcal{M}, f)$

deterministic partition into regions of similar states

1-step memory partition

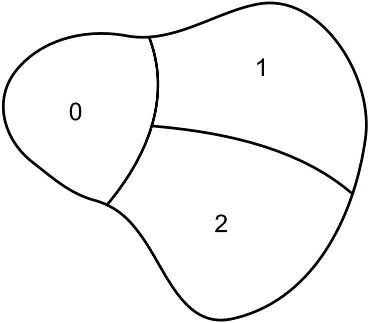

$\mathcal{M} = \mathcal{M}_0 \cup \mathcal{M}_1 \cup \mathcal{M}_2$
ternary alphabet
$\mathcal{A} = \{1, 2, 3\}$.

2-step memory refinement

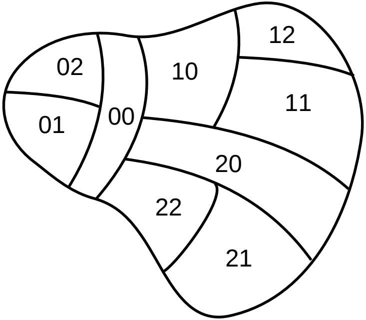

$\mathcal{M}_i = \mathcal{M}_{i0} \cup \mathcal{M}_{i1} \cup \mathcal{M}_{i2}$
labeled by nine ‘words’
$\{00, 01, 02, \cdots, 21, 22\}$.

topological dynamics

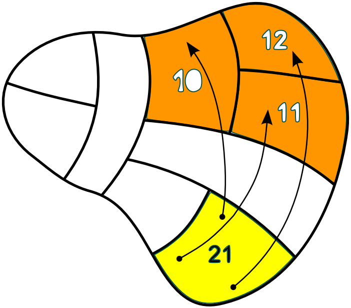

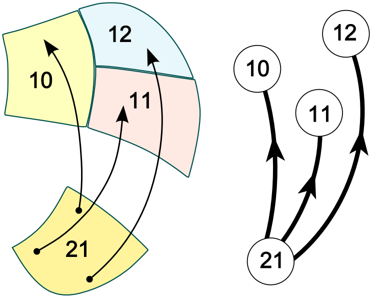

one time step

points from $\mathcal{M}_{21}$
reach $\{\mathcal{M}_{10}, \mathcal{M}_{11}, \mathcal{M}_{12}\}$
and no other regions

each region = node

allowed transitions
$T_{10,21} = T_{11,21} = T_{12,21} \neq 0$
directed links

topological dynamics

Transition graph T_{ba}

regions reached in one time step

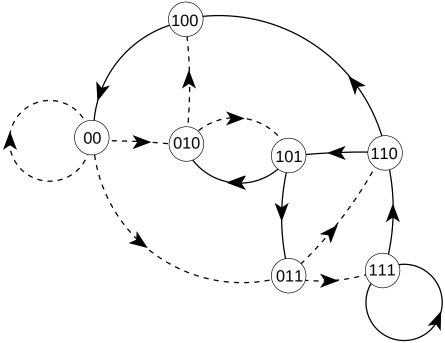

example: state space resolved into 7 neighborhoods

$$\{\mathcal{M}_{00}, \mathcal{M}_{011}, \mathcal{M}_{010}, \mathcal{M}_{110}, \mathcal{M}_{111}, \mathcal{M}_{101}, \mathcal{M}_{100}\}$$

deterministic partitions are no good

deterministic dynamics: partitioning can be arbitrarily fine

requires exponential # of exponentially small regions

deterministic partitions are no good

deterministic dynamics: partitioning can be arbitrarily fine

requires exponential # of exponentially small regions

yet

in practice

every physical problem must be coarse partitioned

reminder : deterministic partition

state space coarse partition

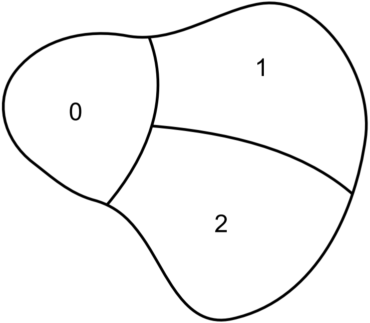

$\mathcal{M} = \mathcal{M}_0 \cup \mathcal{M}_1 \cup \mathcal{M}_2$
ternary alphabet
$\mathcal{A} = \{1, 2, 3\}$

2-step memory refinement

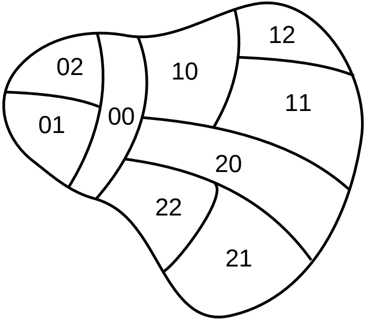

$\mathcal{M}_i = \mathcal{M}_{i0} \cup \mathcal{M}_{i1} \cup \mathcal{M}_{i2}$
labeled by nine 'words'
$\{00, 01, 02, \cdots, 21, 22\}$.

deterministic vs. noisy partitions

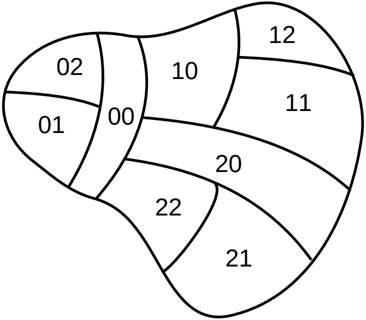

deterministic partition

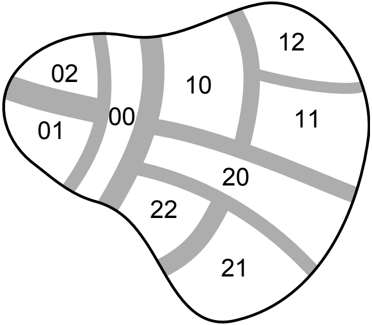

noise blurs the boundaries

can be refined
ad infinitum

when overlapping, no further
refinement of partition

periodic points instead of boundaries

- mhm, do not know how to compute boundaries...

periodic points instead of boundaries

- mhm, do not know how to compute boundaries...
- however, each partition contains a short periodic point

periodic orbit partition

deterministic partition

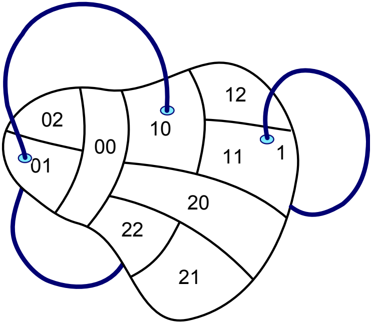

some short periodic points:
fixed point $\overline{1} = \{x_1\}$
two-cycle $\overline{01} = \{x_{01}, x_{10}\}$

noisy partition

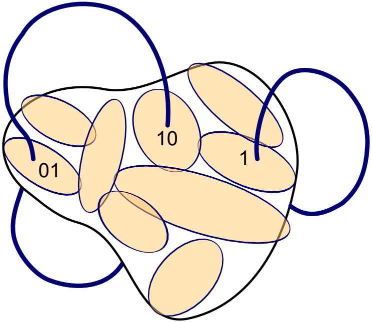

periodic points blurred by noise
into cigar-shaped densities

periodic points and their cigars

- each partition contains a short periodic point smeared into a ‘cigar’ by noise

periodic points and their cigars

- each partition contains a short periodic point smeared into a ‘cigar’ by noise
- compute the size of a noisy periodic point neighborhood!

how big is the neighborhood blurred by the accumulated noise?

the (well known) **key formula** that we now derive:

$$Q_{n+1} = M_n Q_n M_n^T + \Delta_n$$

density covariance matrix at time n: Q_n
noise covariance matrix: Δ_n
Jacobian matrix of linearized flow: M_n

Kalman filter 'prediction'

linearized deterministic flow

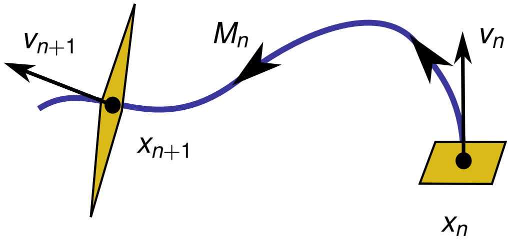

$$x_{n+1} + z_{n+1} = f(x_n) + M_n\, z_n\,, \quad M_{ij} = \partial f_i / \partial x_j$$

in one time step a linearized neighborhood of x_n is

(1) advected by the flow and

(2) mapped by the Jacobian matrix M_n into a stretched and rotated neighborhood whose size and orientation are given by the M eigenvalues and eigenvectors

covariance advection

let the initial density of deviations z from the deterministic center be a Gaussian whose covariance matrix is

$$Q_{jk} = \left\langle z_j z_k^T \right\rangle$$

a step later the Gaussian is advected to

$$\begin{aligned} \left\langle z_j z_k^T \right\rangle &\to \left\langle (M\,z)_j\,(M\,z)_k^T \right\rangle \\ Q &\to M\,Q\,M^T \end{aligned}$$

add noise, get the next slide

roll your own cigar

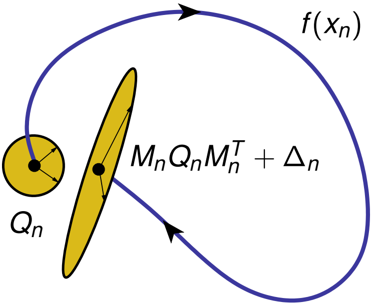

in one time step
a Gaussian density distribution with covariance matrix Q_n is

(1) advected by the flow

(2) smeared with additive noise

into a Gaussian 'cigar' whose widths and orientation are given by the singular values and vectors of Q_{n+1}

covariance evolution

$$Q_{n+1} = M_n Q_n M_n^T + \Delta_n$$

(1) advect deterministically
local density covariance matrix $Q \to MQM^T$

(2) add noise covariance matrix Δ

covariances add up as sums of squares

cumulative noise along a trajectory

iterate $Q_{n+1} = M_n Q_n M_n^T + \Delta_n$ along a trajectory

if M is contracting, $|\Lambda_j| < 1$,

the memory of the covariance Q_0 of the starting density is lost, with iteration leading to the limit distribution

$$Q_n = \Delta_n + M_{n-1}\Delta_{n-1}M_{n-1}^T + M_{n-2}^2\Delta_{n-2}(M_{n-2}^2)^T + \cdots .$$

things fall apart, centre cannot hold

but what if M has *expanding* eigenvalues?

both deterministic dynamics and noise tend to smear densities away from the fixed point: no peaked Gaussian in your future

things fall apart, centre cannot hold

but what if M has *expanding* eigenvalues?

look into the past, for initial peaked distribution that spreads to the present state

for unstable directions, look back

if M has only *expanding* eigenvalues,

balance between the two is attained by iteration from the past, and the evolution of the covariance matrix $\tilde{Q}$ is now given by

$$\tilde{Q}_{n+1} + \Delta_n = M_n \tilde{Q}_n M_n^T \,,$$

[aside to control theorists: reachability and observability Gramians]

Remembrance of Things Past

noisy dynamics of a nonlinear system is fundamentally different from Brownian motion, as the flow ALWAYS induces a local, history dependent effective noise

example : noise and a single attractive fixed point

if all eigenvalues of M are strictly contracting, all $|\Lambda_j| < 1$

any initial compact measure converges to the unique invariant Gaussian measure $\rho_0(z)$ whose covariance matrix satisfies

Lyapunov equation: time-invariant measure condition

$$Q = MQM^T + \Delta$$

[A. M. Lyapunov doctoral dissertation 1892]

example : Ornstein-Uhlenbeck process

width of the natural measure concentrated at the attractive deterministic fixed point $z = 0$

$$\rho_0(z) = \frac{1}{\sqrt{2\pi\, Q}} \exp\left(-\frac{z^2}{2\, Q}\right) \,, \qquad Q = \frac{\Delta}{1 - |\Lambda|^2} \,,$$

- is balance between contraction by Λ and noisy smearing by Δ at each time step
- for strongly contracting Λ, the width is due to the noise only
- As $|\Lambda| \to 1$ the width diverges: the trajectories are no longer confined, but diffuse by Brownian motion

local problem solved: can compute every cigar

a periodic point of period n is a fixed point of nth iterate of dynamics

global problem solved: can compute all cigars

more algebra: can compute the noisy neighborhoods of all periodic points

optimal partition challenge

finally in position to address our challenge:

determine the finest possible partition for a given noise

noisy dynamics partitions: strategy

- use periodic orbits to partition state space
- compute local covariances at periodic points to determine their neighborhoods
- done once neighborhoods overlap

optimal partition hypothesis

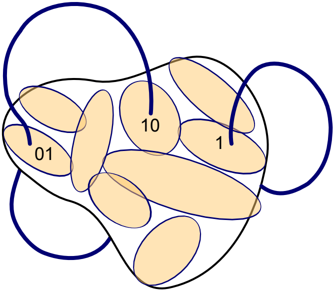

optimal partition:

the maximal set of resolvable periodic point neighborhoods

'the best possible of all partitions' hypothesis formulated as an algorithm

- calculate the local noise covariances Q_a for every unstable periodic point x_a
- assign one-standard deviation neighborhood $[x_a - Q_a, x_a + Q_a]$ to every unstable periodic point x_a
- cover the state space with neighborhoods of orbit points of higher and higher period n_p
- stop refining the local resolution whenever the adjacent neighborhoods of x_a and x_b overlap:

$$|x_a - x_b| < Q_a + Q_b$$

now have: the best possible finite partition of the state space

still need: dynamics

dynamics $\to$ Markov graph

evolution in time

maps intervals
$\mathcal{M}_{011} \to \{\mathcal{M}_{110}, \mathcal{M}_{111}\}$
$\mathcal{M}_{00} \to \{\mathcal{M}_{00}, \mathcal{M}_{011}, \mathcal{M}_{010}\}$, etc..

summarized by the transition graph (links correspond to elements of transition matrix T_{ba}): the regions b that can be reached from the region a in one time step

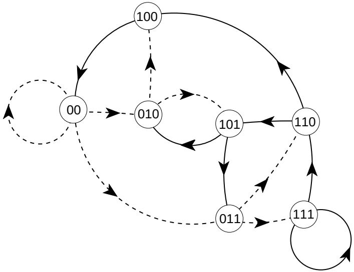

how noise frees us from determinism

noise memory

accumulated noise along a dynamical trajectory always coarsens the partition

we now show that

this partition is

- intrinsic to dynamics
- computable

turbulence.zip

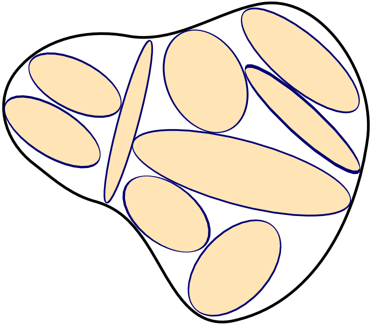

the payback for your patience

claim:

optimal partition hypothesis

- the best of all possible state space partitions
- optimal for the given noise

the payback for your patience

claim:

optimal partition hypothesis

- optimal partition replaces stochastic PDEs by finite, low-dimensional Markov graphs

the payback for your patience

claim:

optimal partition hypothesis

- optimal partition replaces stochastic PDEs by finite, low-dimensional Markov graphs
- finite matrix calculations $\Rightarrow$ optimal estimates of long-time observables (Lyapunov exponents, mean temperature in Chicago and its variance, etc.)

example: representative solutions of fluid dynamics

- Professor Zweistein, from the back of Kresge:
  (1) she has already done all this in 1969
  (2) you must be kidding, it cannot be done for turbulence

example: representative solutions of fluid dynamics

- Professor Zweistein, from the back of Kresge:
  - **(1)** she has already done all this in 1969
  - **(2)** you must be kidding, it cannot be done for turbulence
- OK, OK, we have about 50 state space cell centers

  [click here for examples of frozen fluid states]

  [click here for examples of a fluid in periodic motions]

  and we have their Jacobians M (that was hell to get)

disclosure

- Computation of unstable periodic orbits in high-dimensional state spaces, such as Navier-Stokes,

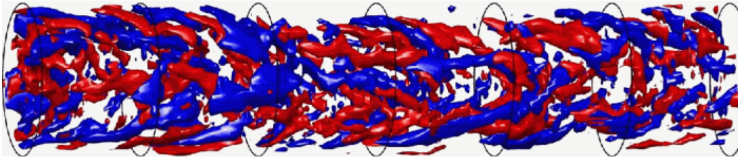

  is at the border of what is feasible numerically, and criteria to identify finite sets of the most important solutions are very much needed. Where are we to stop calculating these solutions?

disclosure

- disclosure

  we have not yet tested the method on fluid dynamics data sets.

disclosure

- disclosure

  we have not yet tested the method on fluid dynamics data sets.
- Georgia Tech Center for Nonlinear Science is looking for several brave postdocs to help us really ‘zip’ turbulence

disclosure

- disclosure

  we have not yet tested the method on fluid dynamics data sets.
- Georgia Tech Center for Nonlinear Science is looking for several brave postdocs to help us really 'zip' turbulence
- the brave candidates: step up after the talk

references

- D. Lippolis and P. Cvitanović, *How well can one resolve the state space of a chaotic map?*, Phys. Rev. Lett. 104, 014101 (2010); arXiv.org:0902.4269
- P. Cvitanović and D. Lippolis, *Knowing when to stop: How noise frees us from determinism*, in M. Robnik and V.G. Romanovski, eds., *Let's Face Chaos through Nonlinear Dynamics* (Am. Inst. of Phys., Melville, New York, 2012); arXiv.org:1206.5506