

Machine Learning for Physicists

Summer 2017

University of Erlangen-Nuremberg

Florian Marquardt

Florian.Marquardt@fau.de

<http://machine-learning-for-physicists.org>

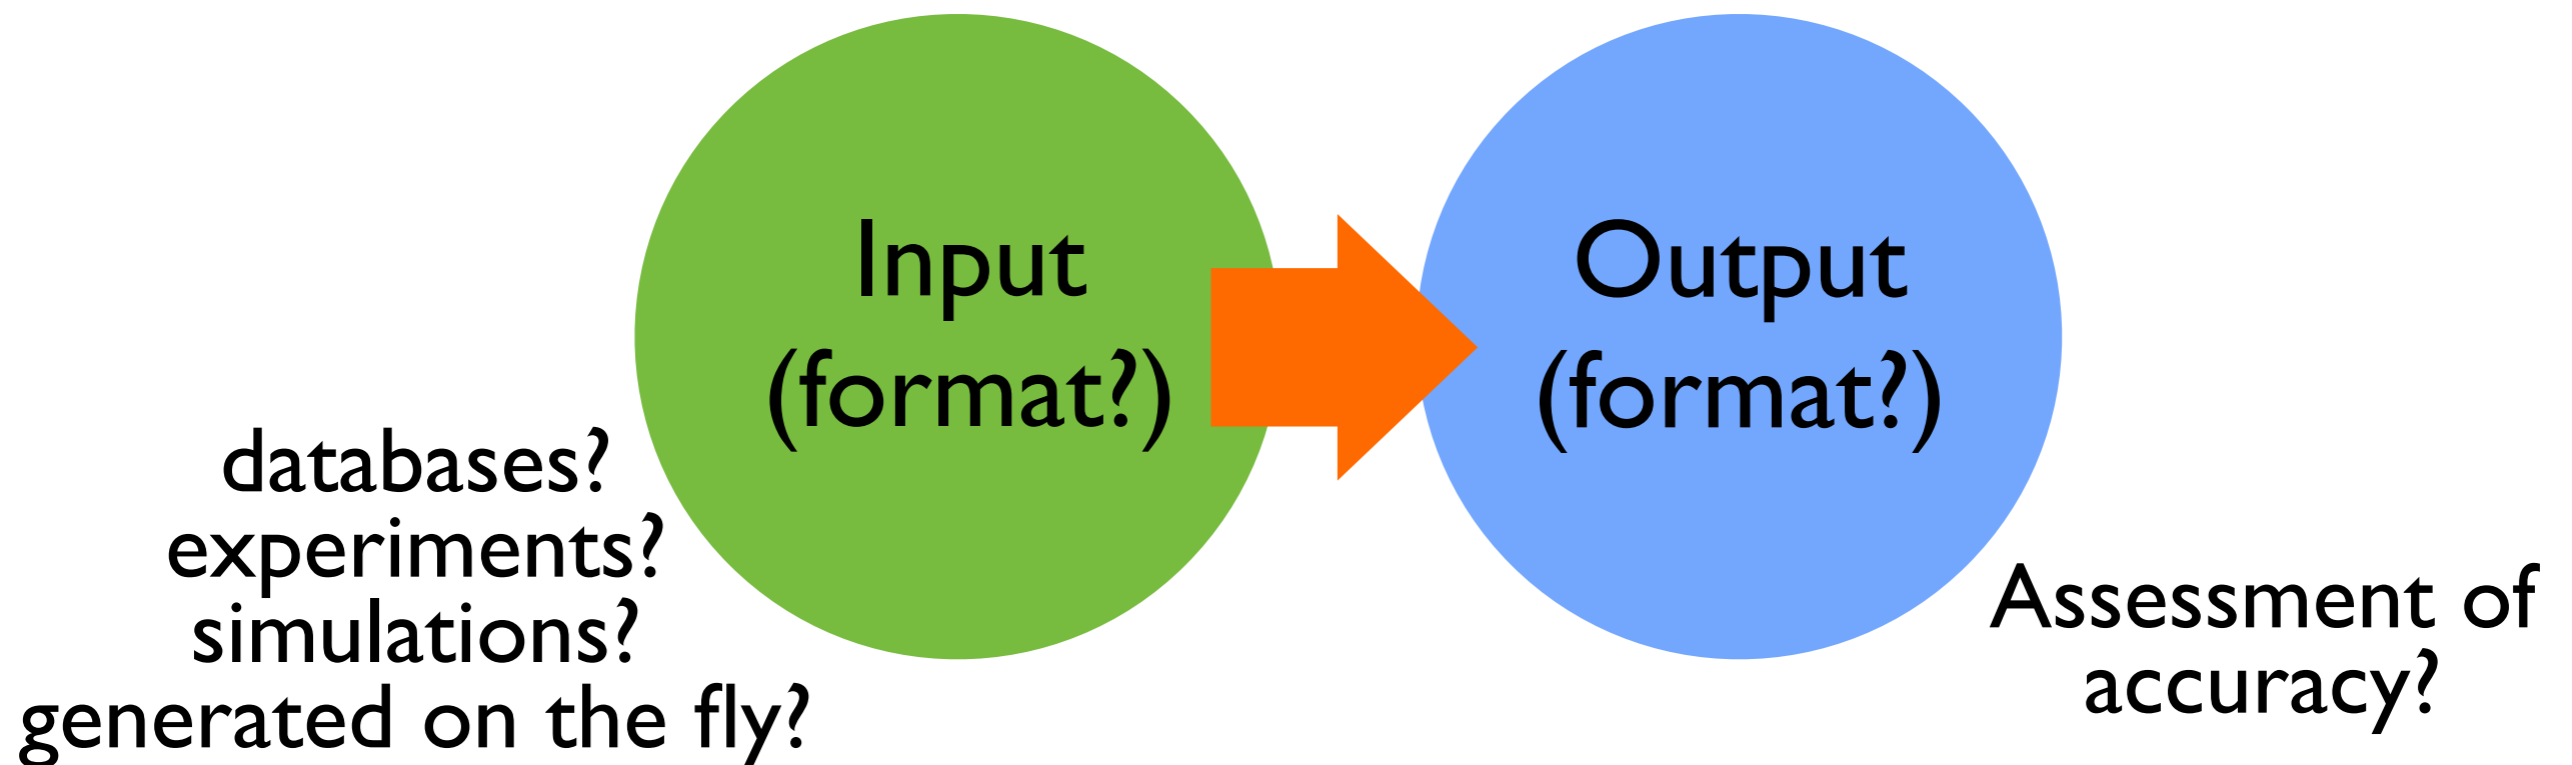
Neural networks applied to scientific tasks

Some examples

Neural networks applied to scientific tasks:

General questions



- scientific goal and challenges (benefits of applying machine learning?)
- input format
- output format
- number of training examples (how obtained?)
- accuracy of predictions (better than other machine learning techniques? better than humans?)
- speedup vs. direct acquisition of outputs (from experiments or computation)



Quantum Chemistry

- Predict binding energy of molecules, orbital energies, structure, etc.
- Success depends on input representation: “molecular fingerprints”. One version used has been: $1/r$ for the distances r between the atoms in the molecule.
- Use training data, e.g. from density functional theory (DFT) or Hartree Fock calculations.
- Errors achieved are comparable to error of QM numerics used to construct the training data.
- “Delta Learning”: Use results from less expensive methods (DFT) as input to predict results that would be obtained using more expensive methods.
- Typical errors about 0.1 eV

(Image: Wikipedia)




Computer-aided drug design

- drug discovery (activity with respect to some target molecule, toxicity)
- Databases exist containing the properties of 100 million compounds (PubChem) and 100 thousand protein structures (Protein Data Bank)
- Databases list the experimentally determined activity of millions of compounds in 100s (or more) of different experimental tests [some are proprietary, held by pharmaceutical companies]
- A typical task may be to predict whether there will be a certain kind of reaction (with some target), given an encoding of the molecule as input
- Deep neural networks recently won challenges for predicting activity and toxicity of chemical compounds
- “multi-task” nets: trained on predicted multiple properties; work better because they seem to build a more useful internal representation

Computational structural biology

- Protein folding: given the structure, how does the final shape look like? Extremely challenging (molecular dynamics possible, but too slow for larger structures)
- One particular task: “Given two parts of the sequence that are apart, will they be close to each other in the folded structure?” (“protein contact prediction”)
- Previously, 30% was the best success rate for this task; deep neural nets have pushed this to 36%




(Image: Wikipedia)

Materials research


- Find better batteries, better solar cells, materials for storing molecules (like absorbing CO₂)
- in general: predict crystal structures and properties (electronic, melting point, heat conductivity, magnetic, ...), of novel/untested materials, or under extreme conditions
- “Materials Genome Project” and similar large-scale efforts to collect various databases (of real materials that have been synthesized at some time, and also of computed materials, and of extrapolated/predicted properties)
- Functional materials (e.g. solar cells) vs. structural materials (e.g. steel). The latter depend on how they are processed: very difficult to model!

Statistical physics: phase transitions

(finding the transition point? discovering new phases?)



disordered phase



ordered phase

e.g. use image recognition to learn to distinguish between the two phases (supervised learning)

Statistical physics: phase transitions

- Classifying phases of matter, e.g. by looking at images (of the magnetization, of wave functions), or other extracted quantities
- supervised and unsupervised approaches
- Example: **Ising model**. Predicting the temperature by looking at magnetization patterns in an Ising model will yield network that has different weights associated with T below and above the phase transition. Note however that almost any statistics of the magnetization pattern would be able to distinguish T below and above! <https://arxiv.org/pdf/1609.09087.pdf>
- Example: **Disordered systems**. Transitions between types of quantum states in disordered systems (localized to delocalized, and Chern insulator to Anderson insulator). Teach the network to distinguish localized from delocalized wave functions. Can be used to search the transition automatically at other parameter values (e.g. other energies), but simple measures for localization exist already. <https://arxiv.org/pdf/1610.00462.pdf>

How to find the transition point?

Teach the network to distinguish between two phases from the high- and low-temperature cases, then check output near the transition point

true phase:

true phase A

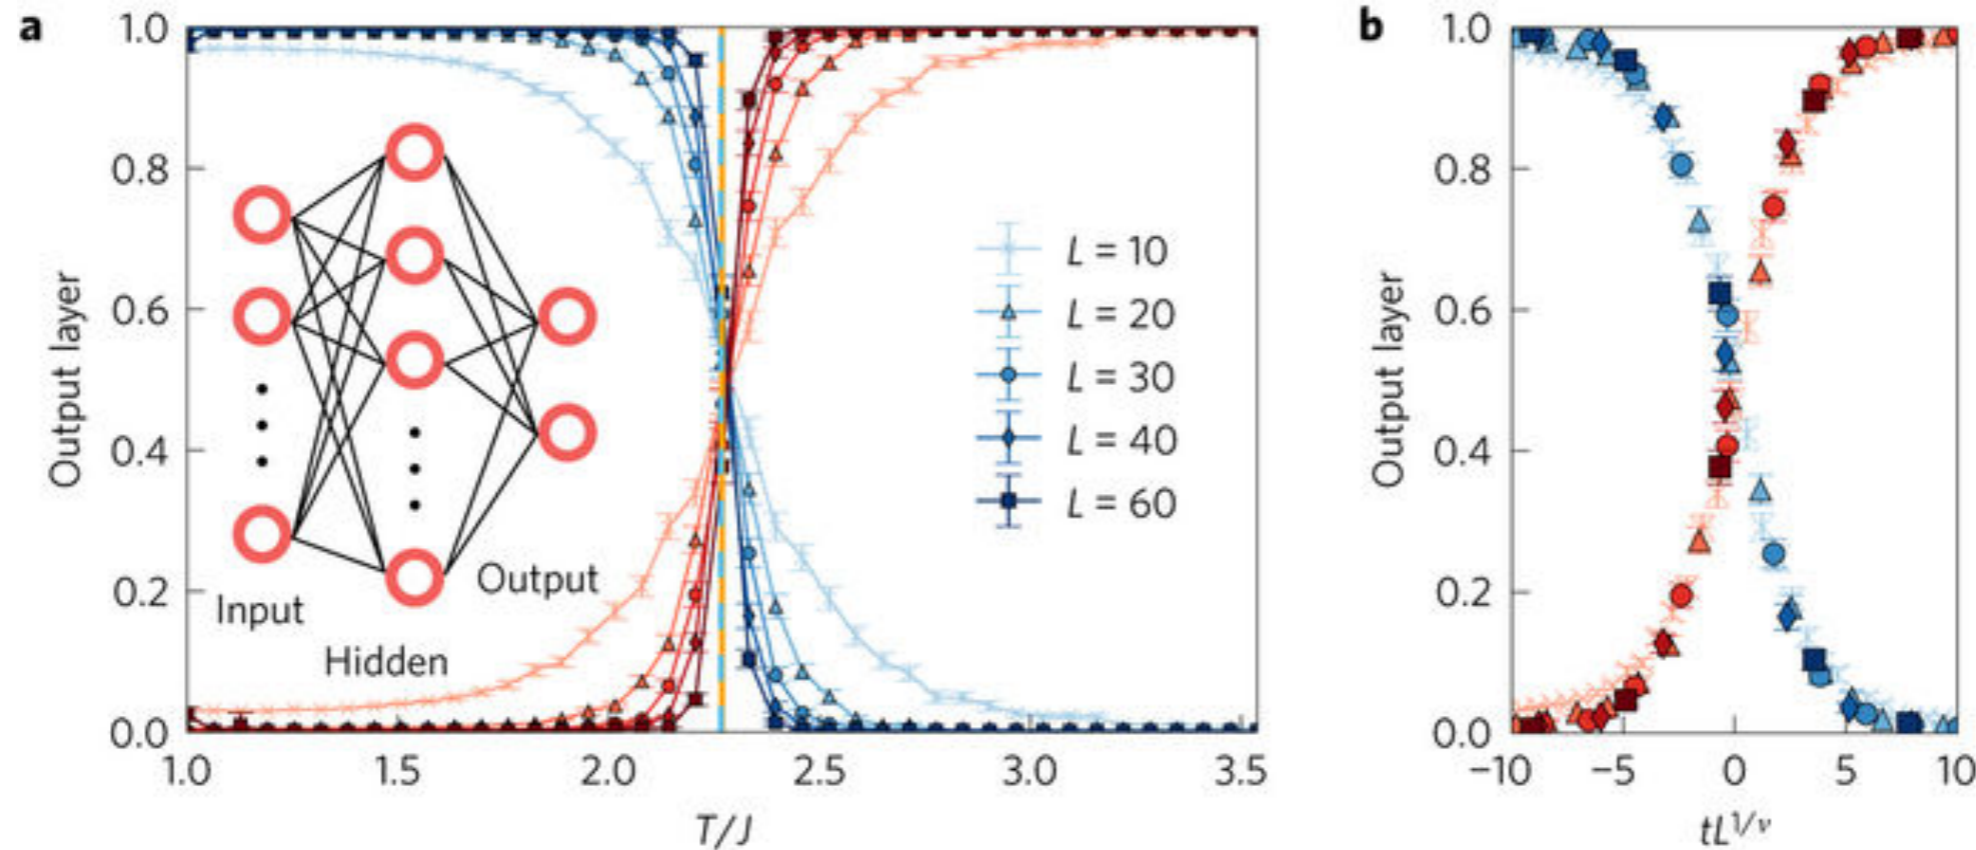
true phase B

prescribed type:

“type A”

“type B”

control parameter
(e.g. temperature)



How to find the transition point?

Teach the network to distinguish between two phases at an arbitrary artificial cut: works best if the cut actually coincides with the true phase transition point!

true phase:




prescribed type:



control parameter
(e.g. temperature)

artificial cut





position of arbitrary cut imposed by hand

Surface code quantum error correction

Surface code: encode two logical qubits into a lattice of physical qubits


Given a 'syndrome' (detected anomalies), what might be the underlying chain of errors in the qubit lattice? [for correcting them]



Train restricted Boltzmann machine to turn syndrome S into error chain e

Artificial Intelligence for scientific discovery?


“Can the computer think like a scientist
and discover laws of nature?”



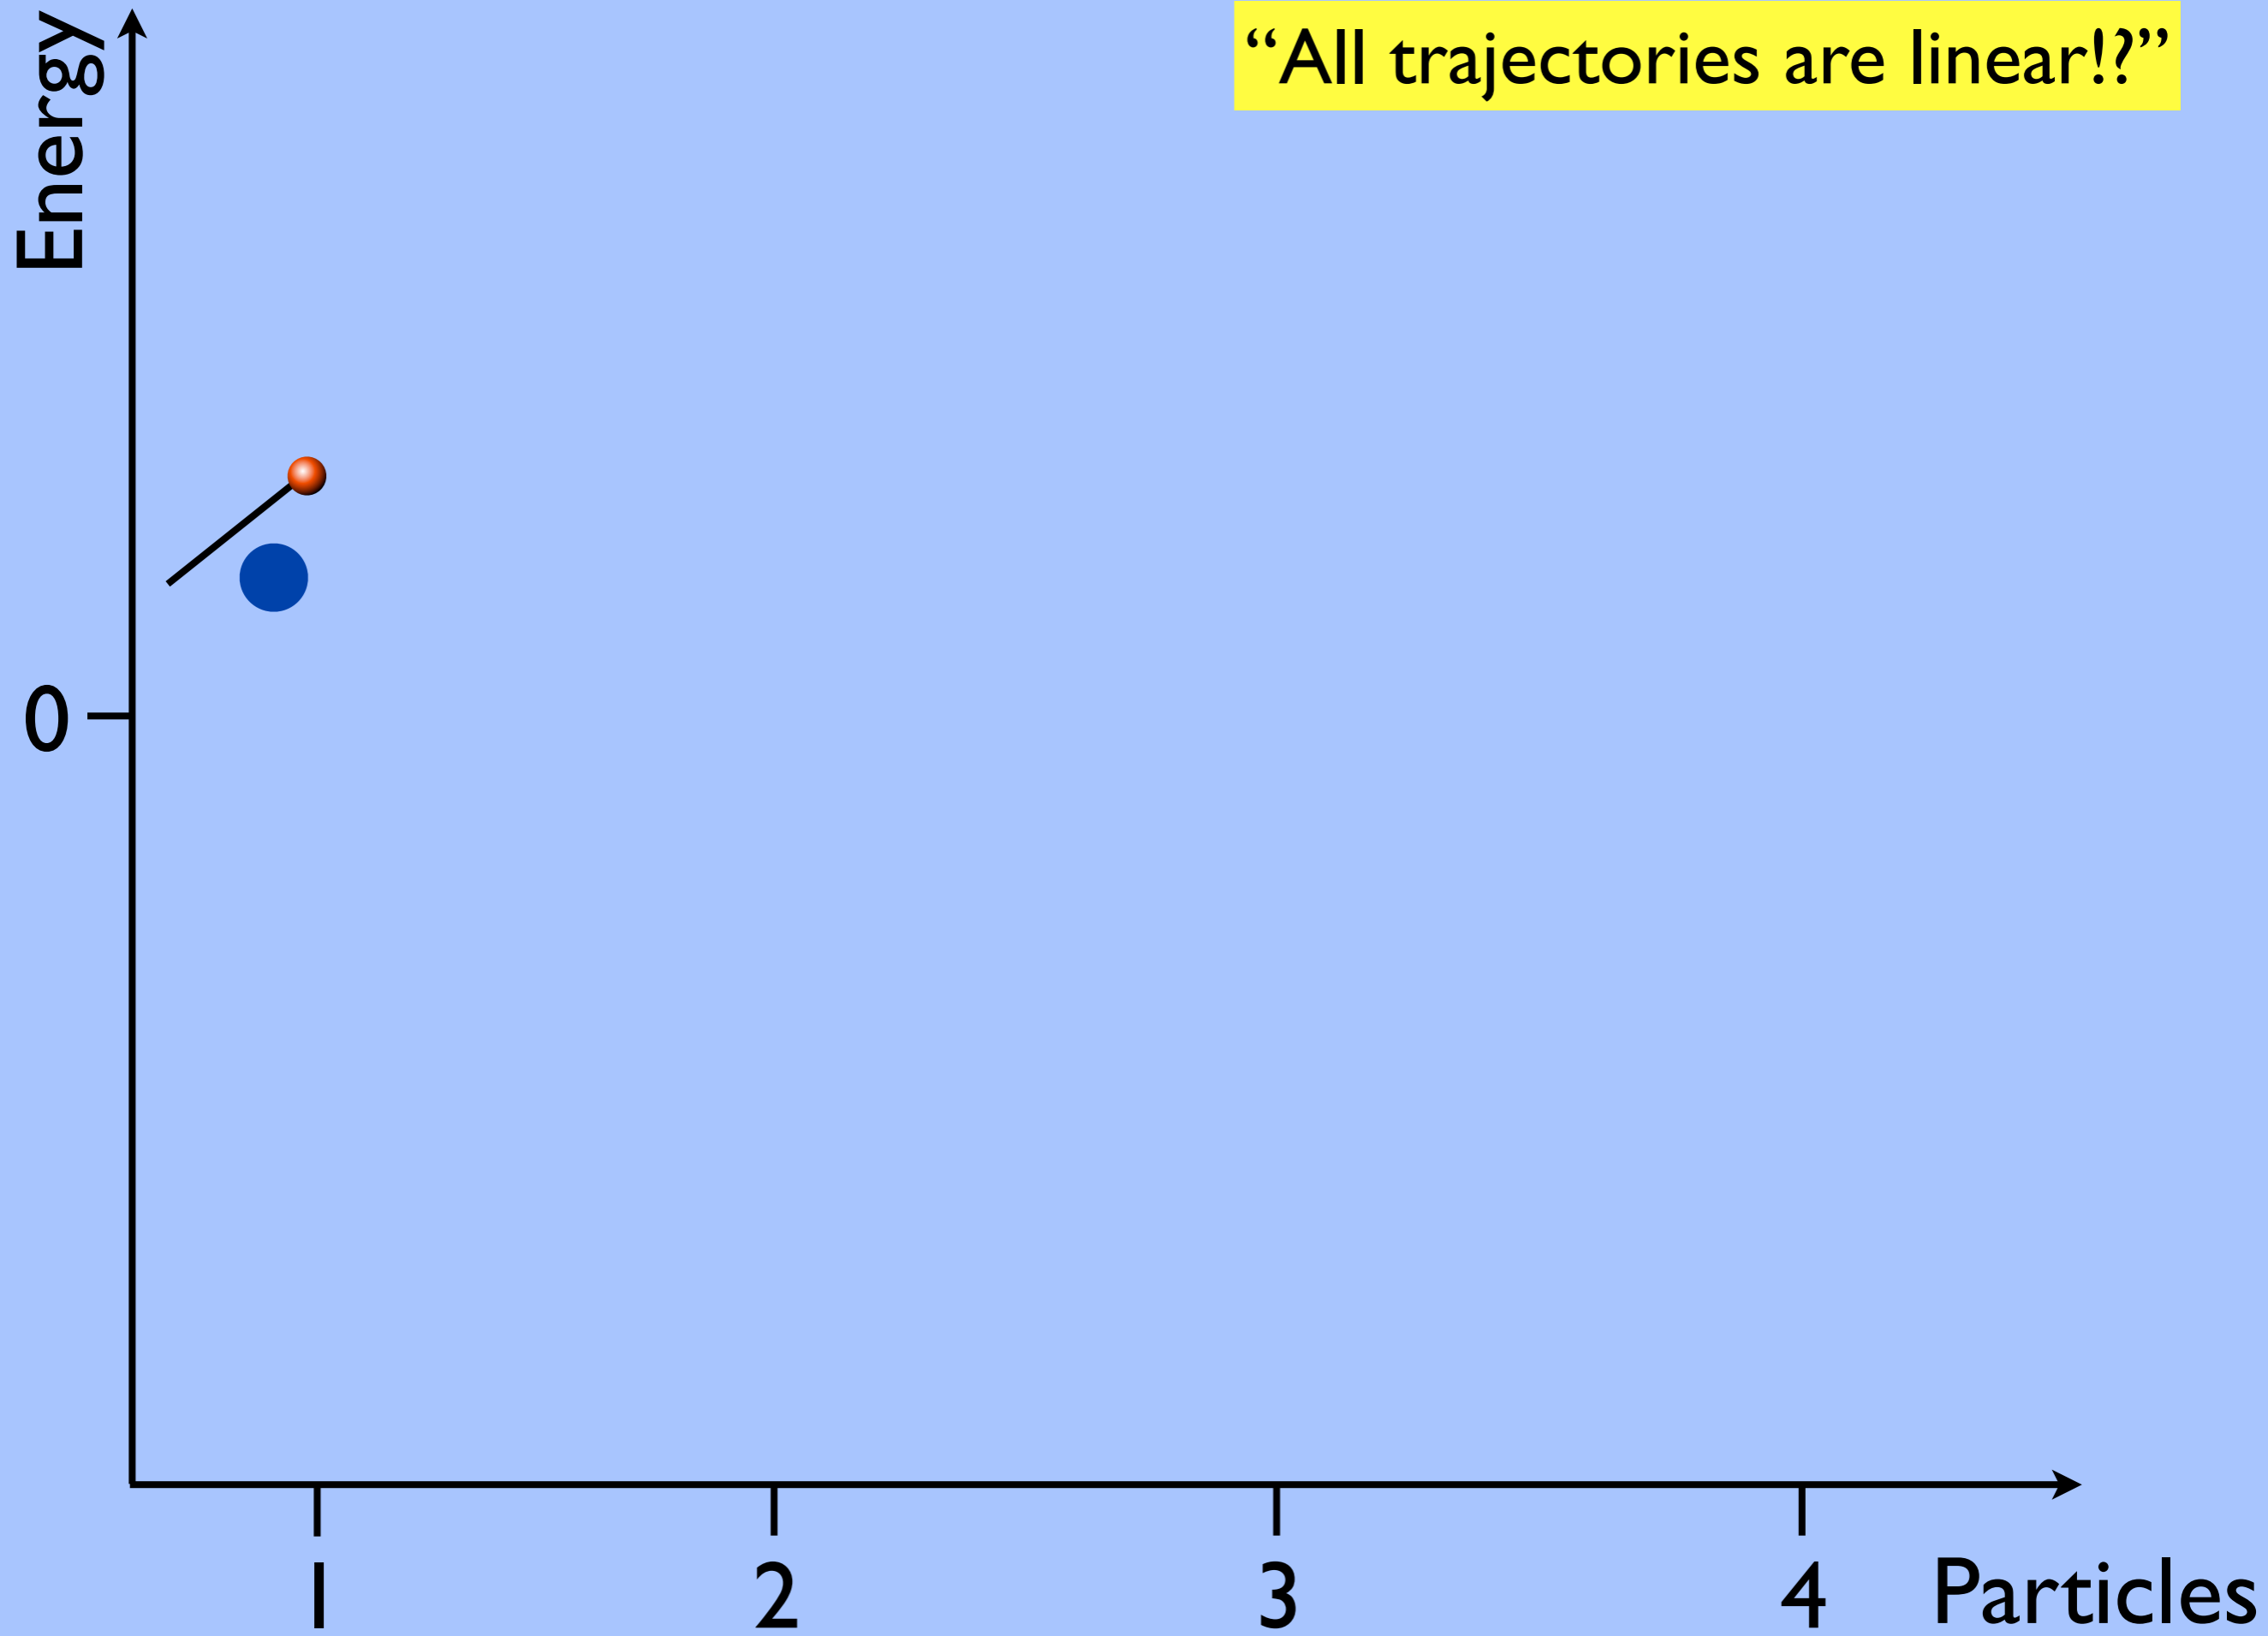
Let's consider an instructive but still restricted setting in physics, which I will call:

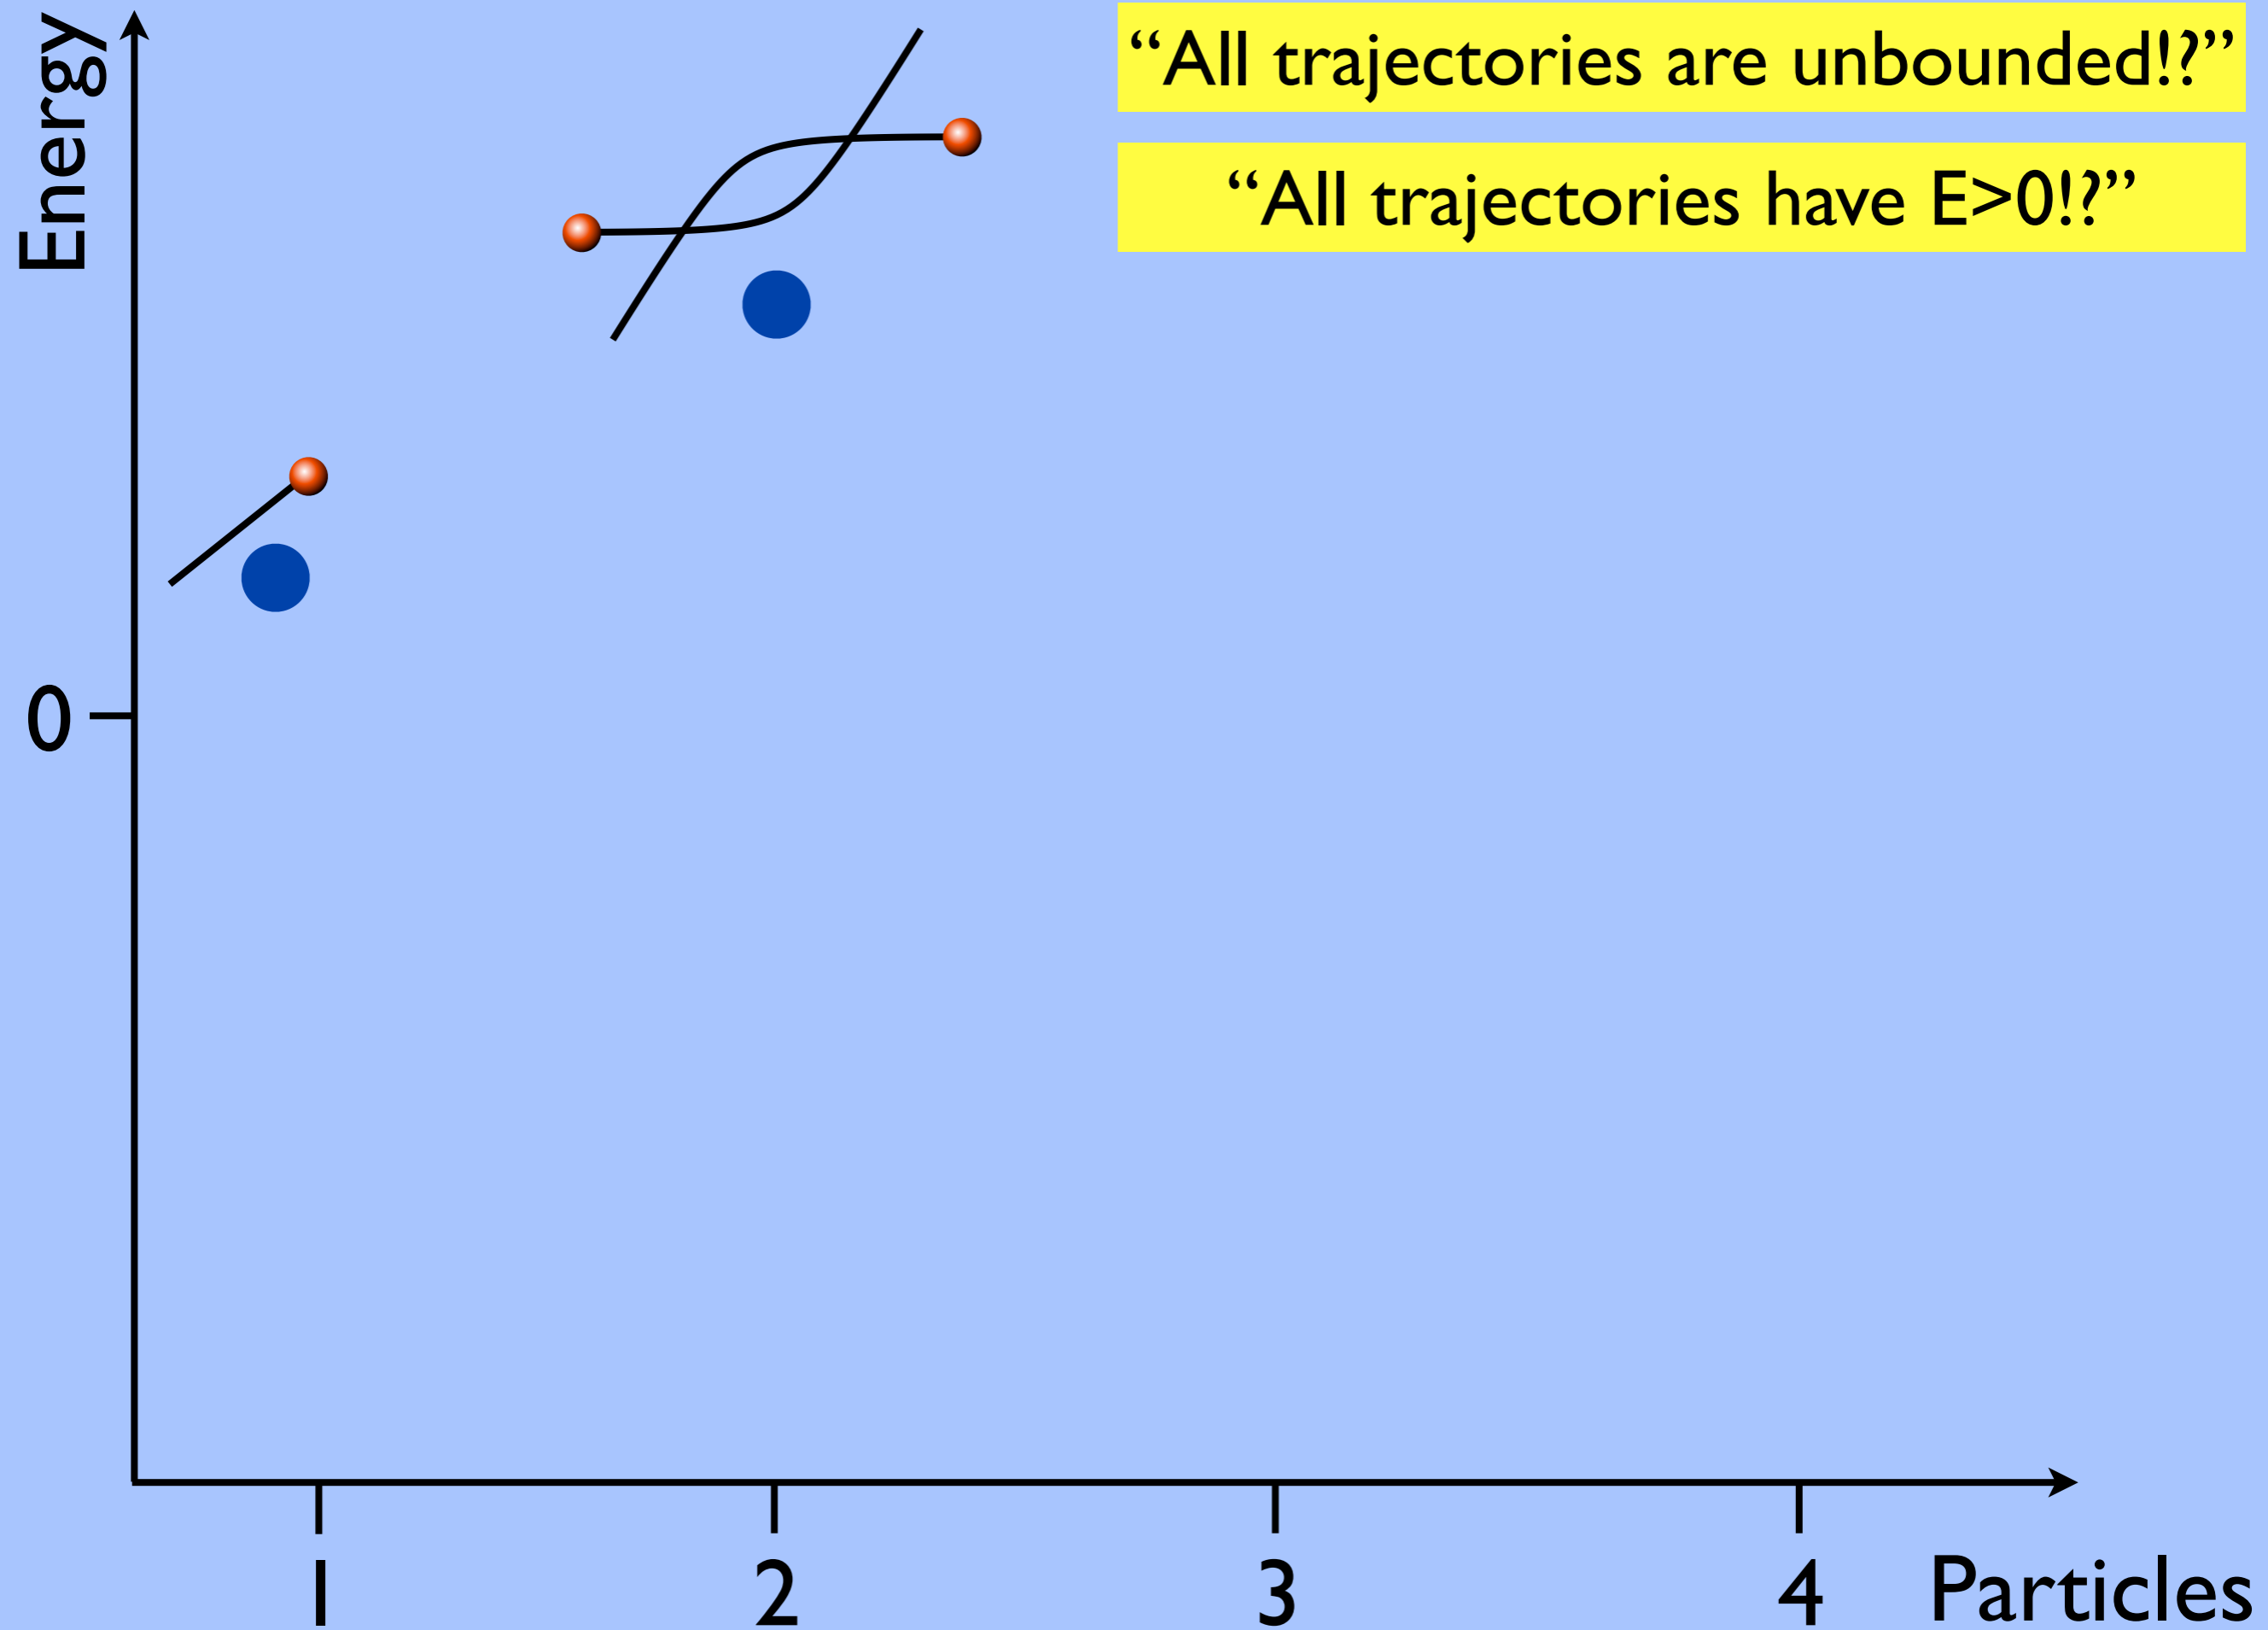
“Kepler’s World”

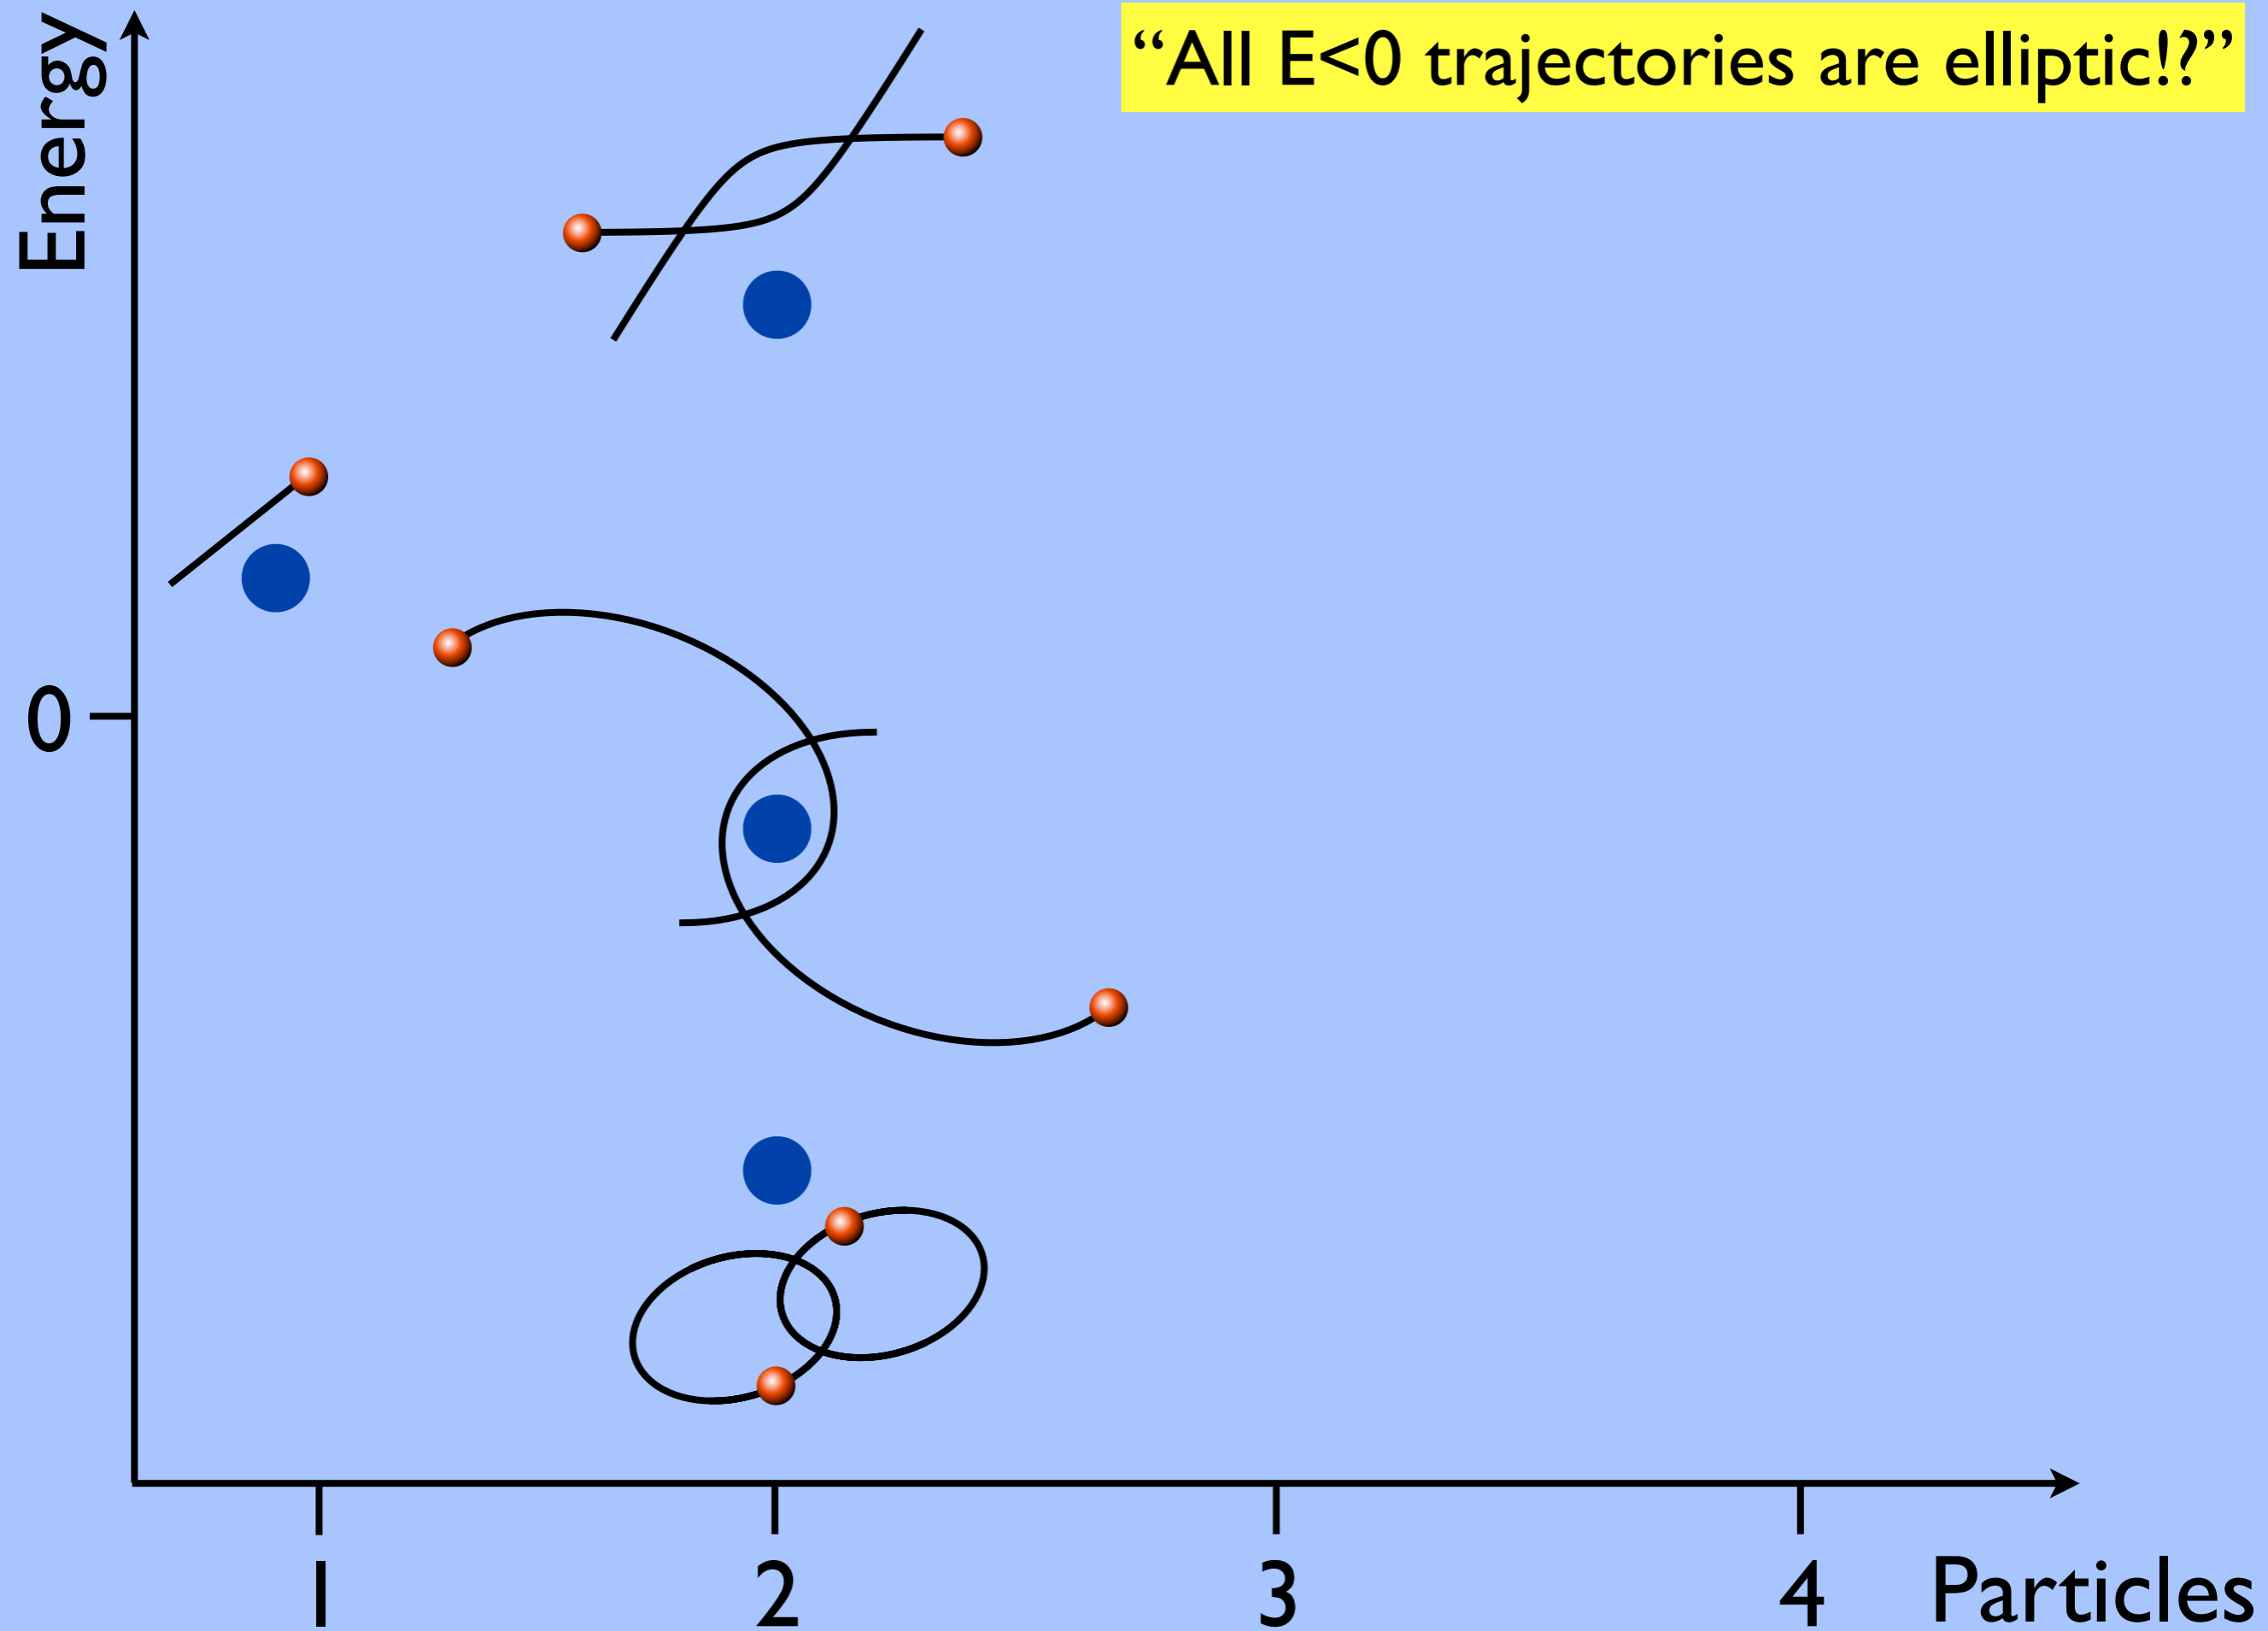
- N particles, attracting each other via the gravitational force
- Computer (or some super-smart and very patient student) is allowed to do numerical experiments and try to learn from them

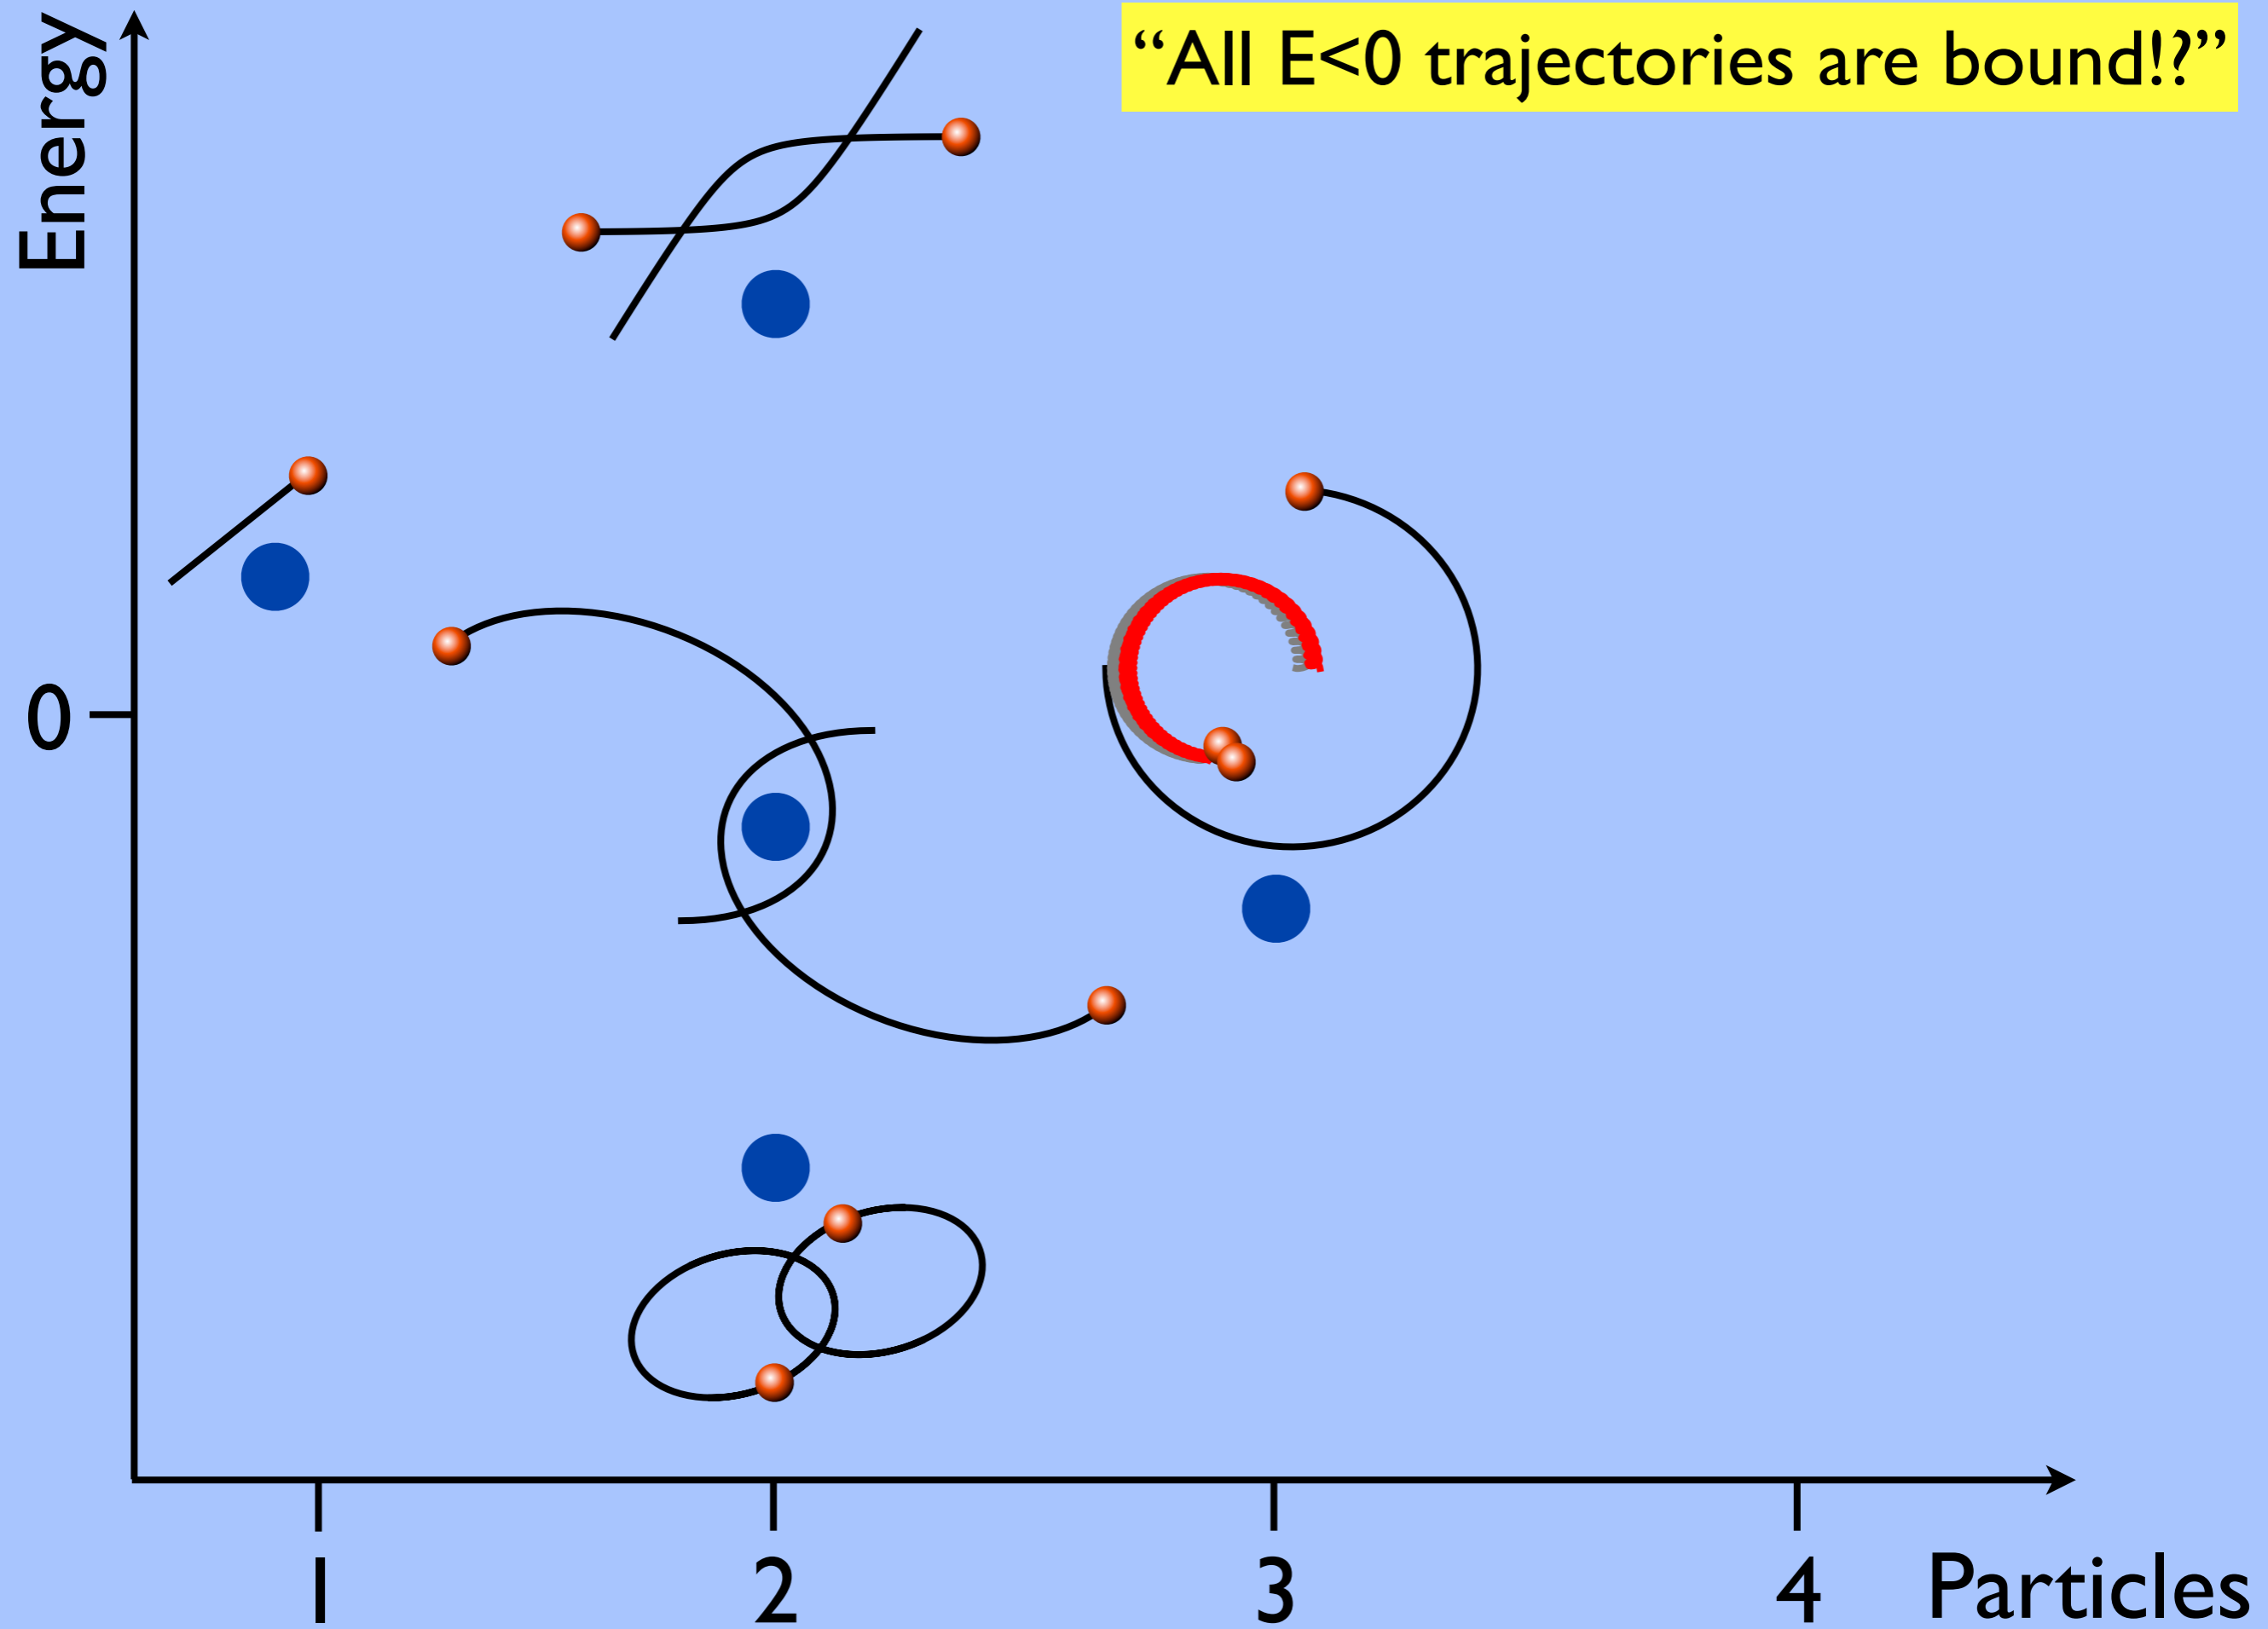



“All trajectories are linear!?”




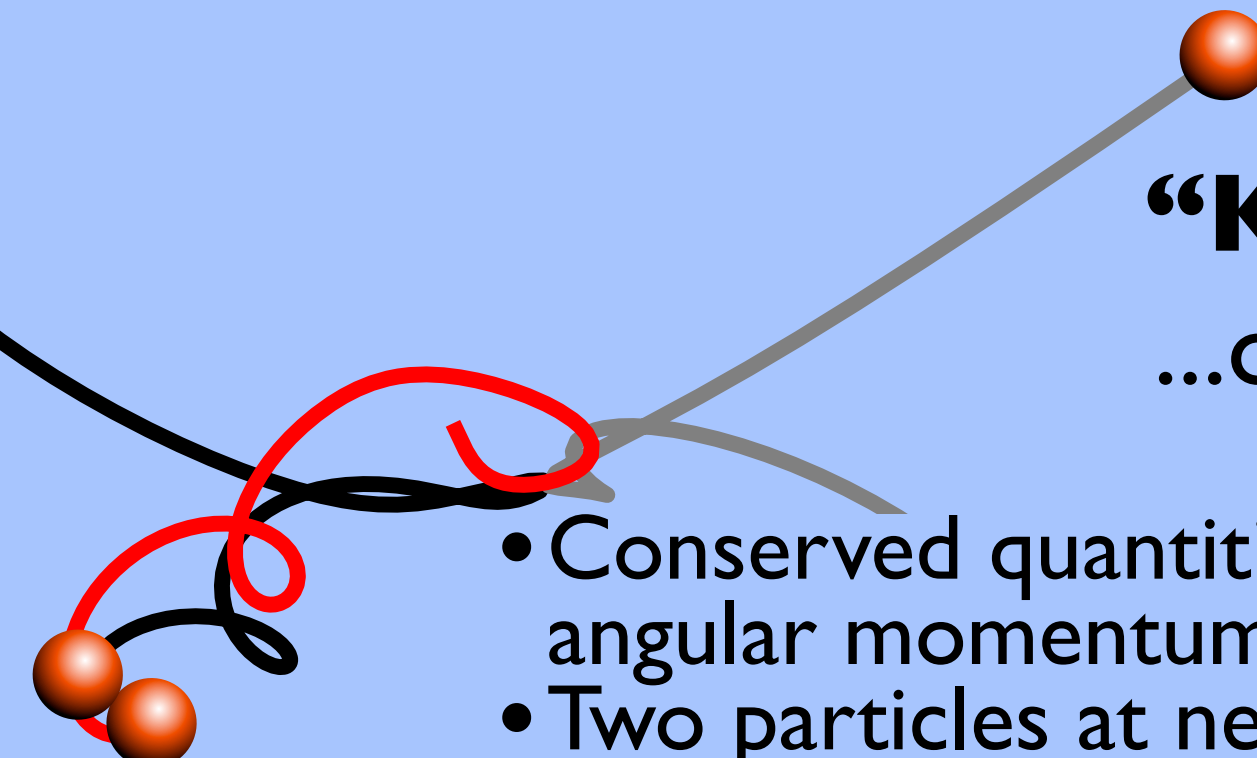













“Kepler’s World”


...one might discover:

- Conserved quantities (total energy, momentum, angular momentum)
- Two particles at negative energy: bound, on elliptical orbits, relation between distance and period
- Equations of motion with inverse square force
- Approximate elliptical orbits if many light-mass particles travel around one heavy mass
- “Moons”
- Properties of scattering: energy and momentum exchange
- 3-particle scattering can lead to expulsion
- Effects of “resonances” between orbital periods
- Rings
- Chaotic behaviour
- Structure formation in clouds of many particles
- Local Boltzmann distribution
- Evaporation
- ...things as yet unknown...?!

Hypotheses


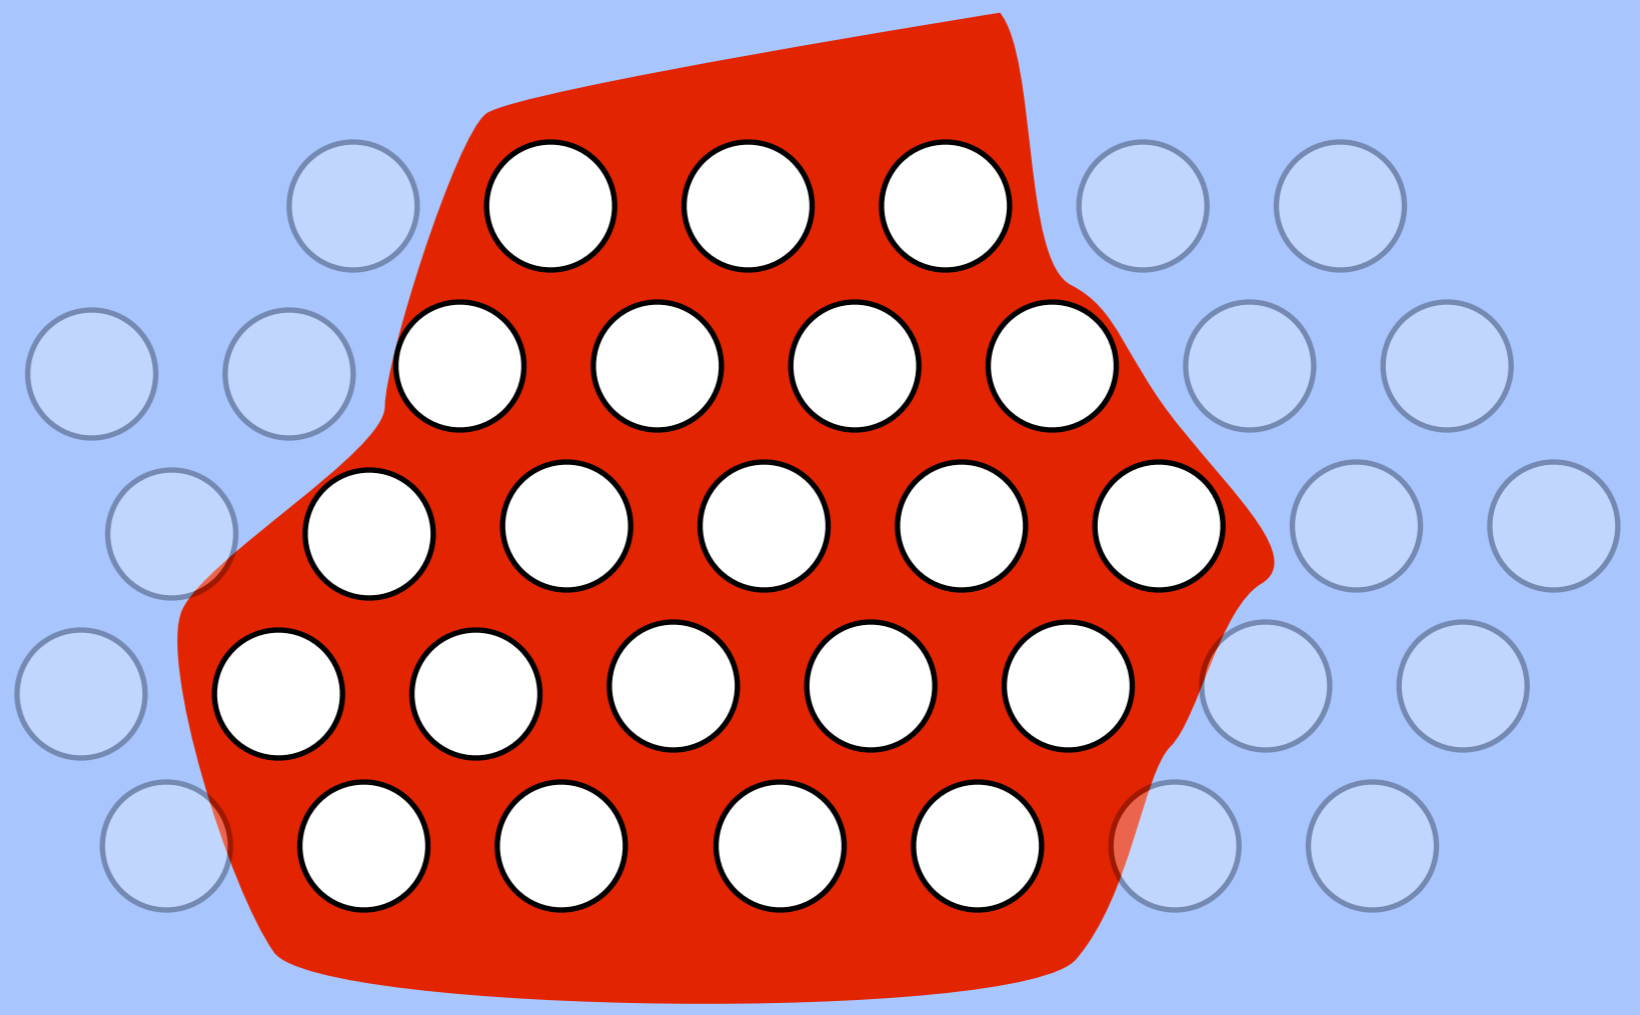


hypotheses



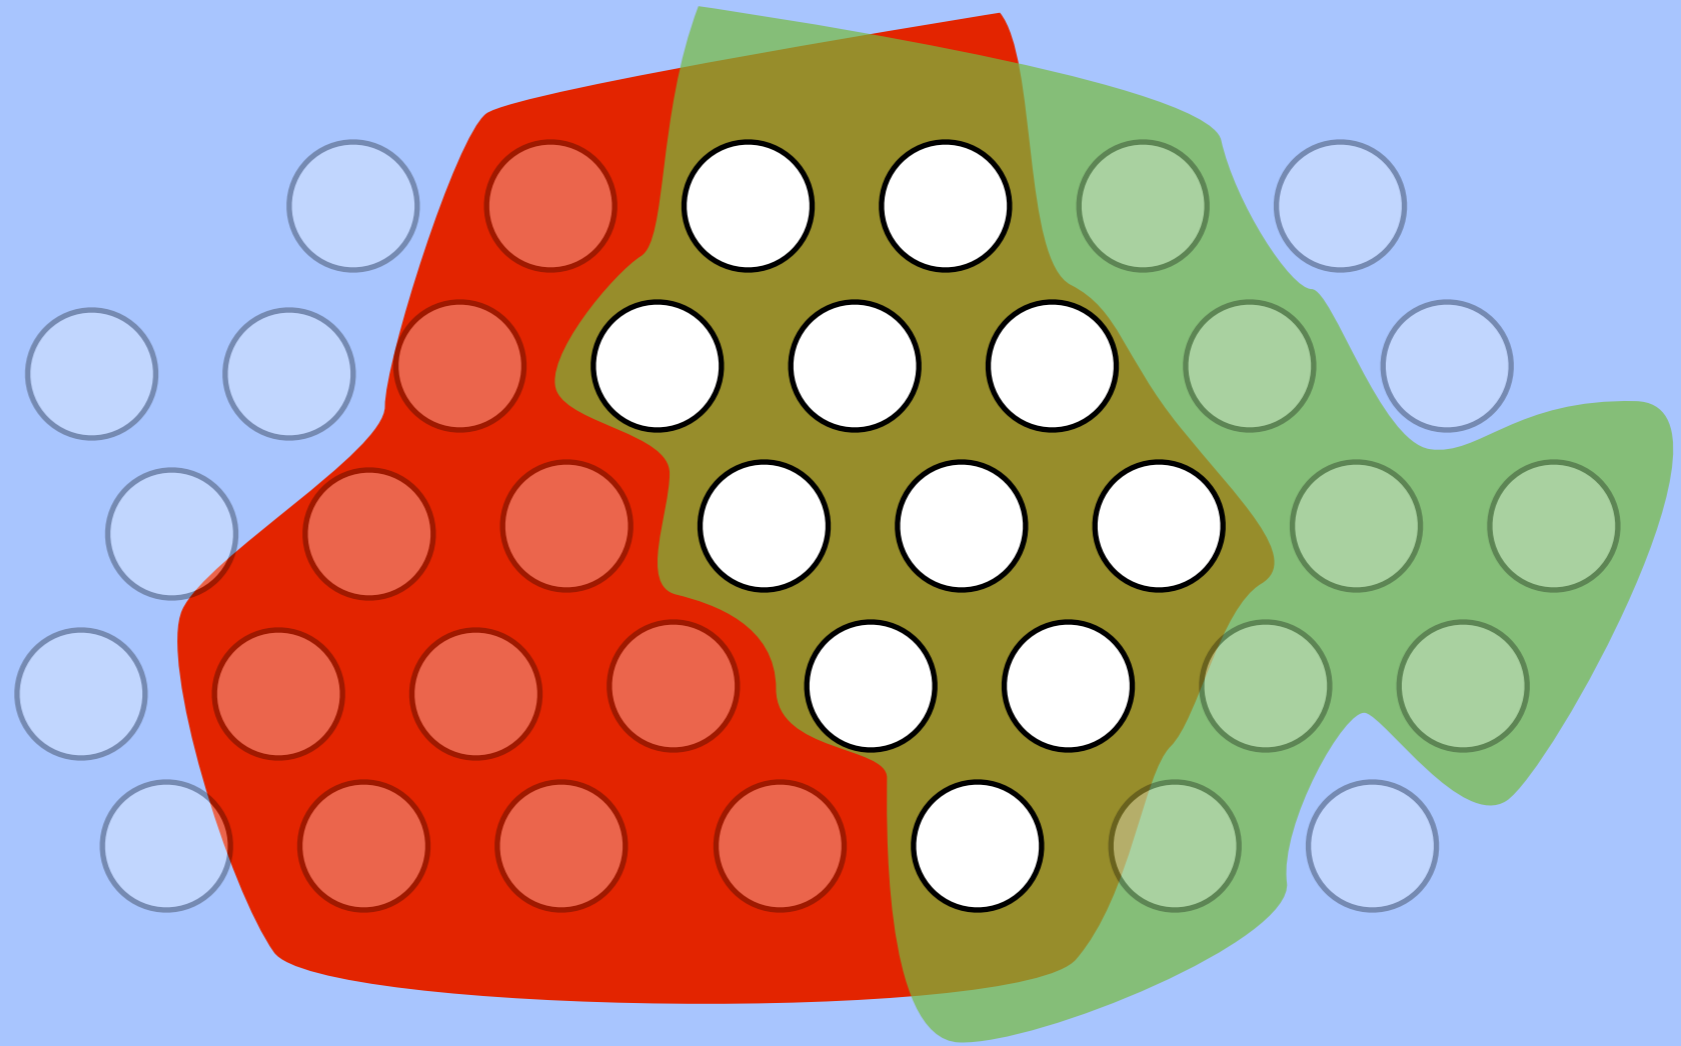
Hypotheses

After first observation:




Hypotheses

After second observation:




Hypotheses

After third observation:



Hypotheses





Hypotheses

Problem: there are infinitely many hypotheses, and we cannot enumerate all of them (unlike in the picture). We have to generate them and select them for consideration.

Also, keep even disproven hypotheses in order to modify them later, to turn them into viable hypotheses again (e.g. by restricting the suggested domain of validity, or by slightly changing the claim).



Hypotheses

(falsifiable claims about relations between quantities)

Quantities about which statements can be made:

- originally accessible quantities (trajectory, parameters such as masses and number of particles) [in principle, all hypotheses can be formulated entirely in terms of those, but it may be cumbersome]
- derived quantities:
 - continuous numbers and fields (distances, energy, force fields, average velocity squared, ...)
 - integer quantities (e.g. number of particle pairs with negative pair-energy, number of local maxima of $x_3(t)$, winding number of a particle pair connection vector evolving in time, ...)
 - Boolean quantities (e.g. “does $x_3(t)$ rise monotonically?”, “is the energy below zero?”, “does the trajectory lie on an ellipse within numerical accuracy?”, ...)



Hypotheses

(falsifiable claims about relations between quantities)

Relations may be:

- precise statements, e.g.
 - $E(t) = E(0)$
 - “trajectory lies on an ellipse if $N=2$ and $E < 0$ ”
- statements about limiting cases:
 - e.g. Taylor-expansion: “ $x_3(t+dt) - x_3(t) = dt * v_{x_3}(t)$ in the limit of small dt ”
- statistical hypotheses: “for uniformly distributed initial coordinates within some small interval, the velocity distribution at the final time is a Gaussian, with a spread given in the following way in terms of the initial energy: ...”
- quantitative statements where one may quantify the deviation?



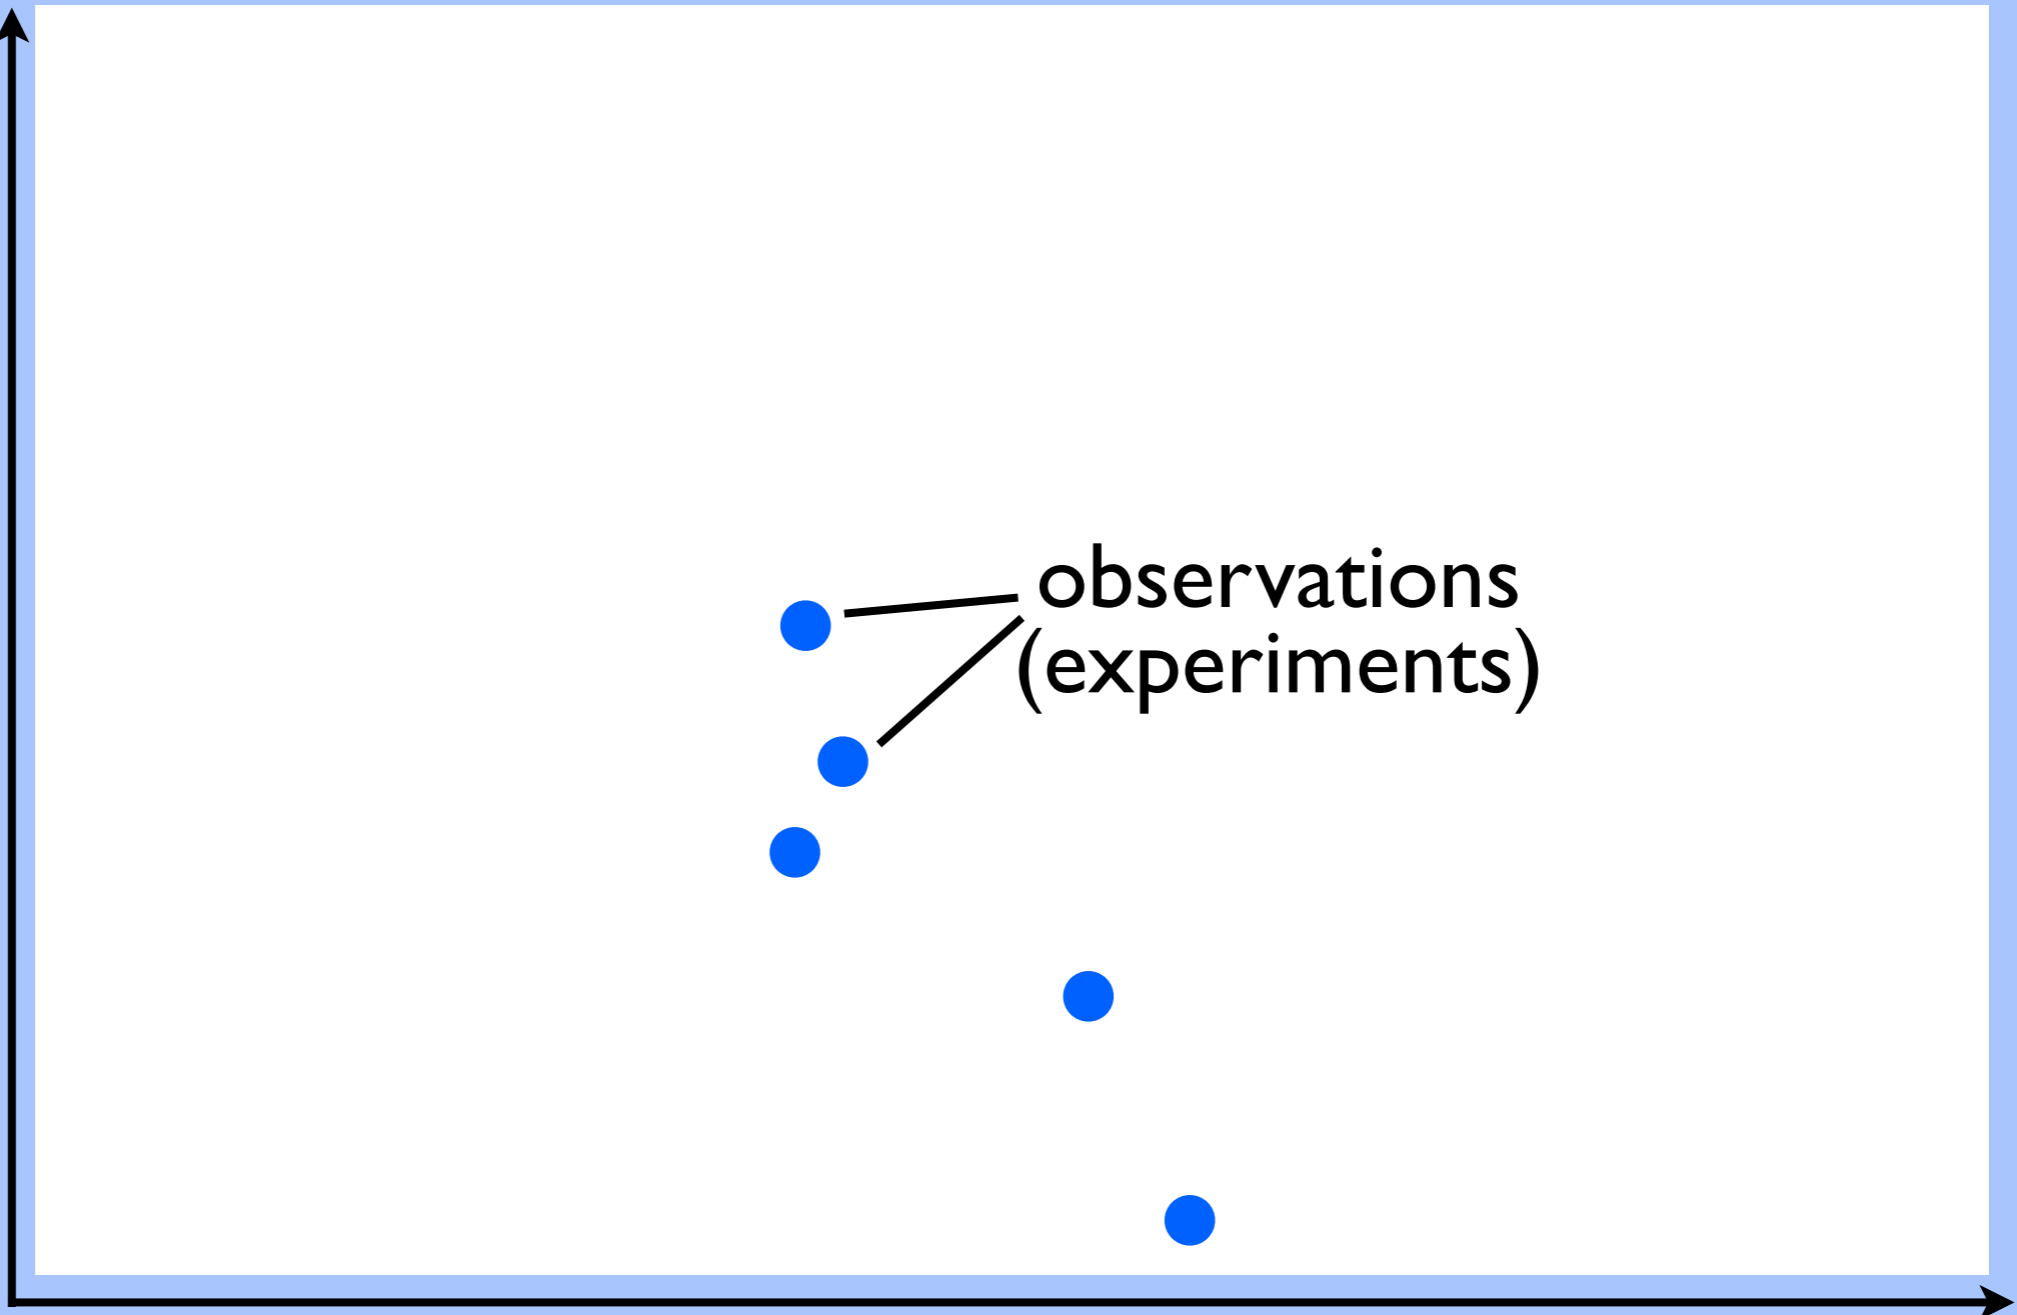
Hypotheses

(falsifiable claims about relations between quantities)

- Hypothesis becomes more plausible with every test that it could have failed (but did not)
- Generate primarily hypotheses that are not obviously wrong [do not contradict known facts] but also not obviously right [do not follow logically from some very well-tested hypothesis]!
- It may still be useful to generate (by logic/math) new statements out of well-tested ones, if the new ones are simpler and easier to apply [e.g. energy conservation out of eqs. of motion]
- sometimes it is useful to consider hypotheses that are found false overall, but approx. true in some domains
- Look for cases where the hypothesis can be easily tested (maybe far away from all those cases where it was successful)
- Consider testing a hypothesis especially if it were to have many consequences
- Prefer 'simple' hypotheses ("Occam's razor")

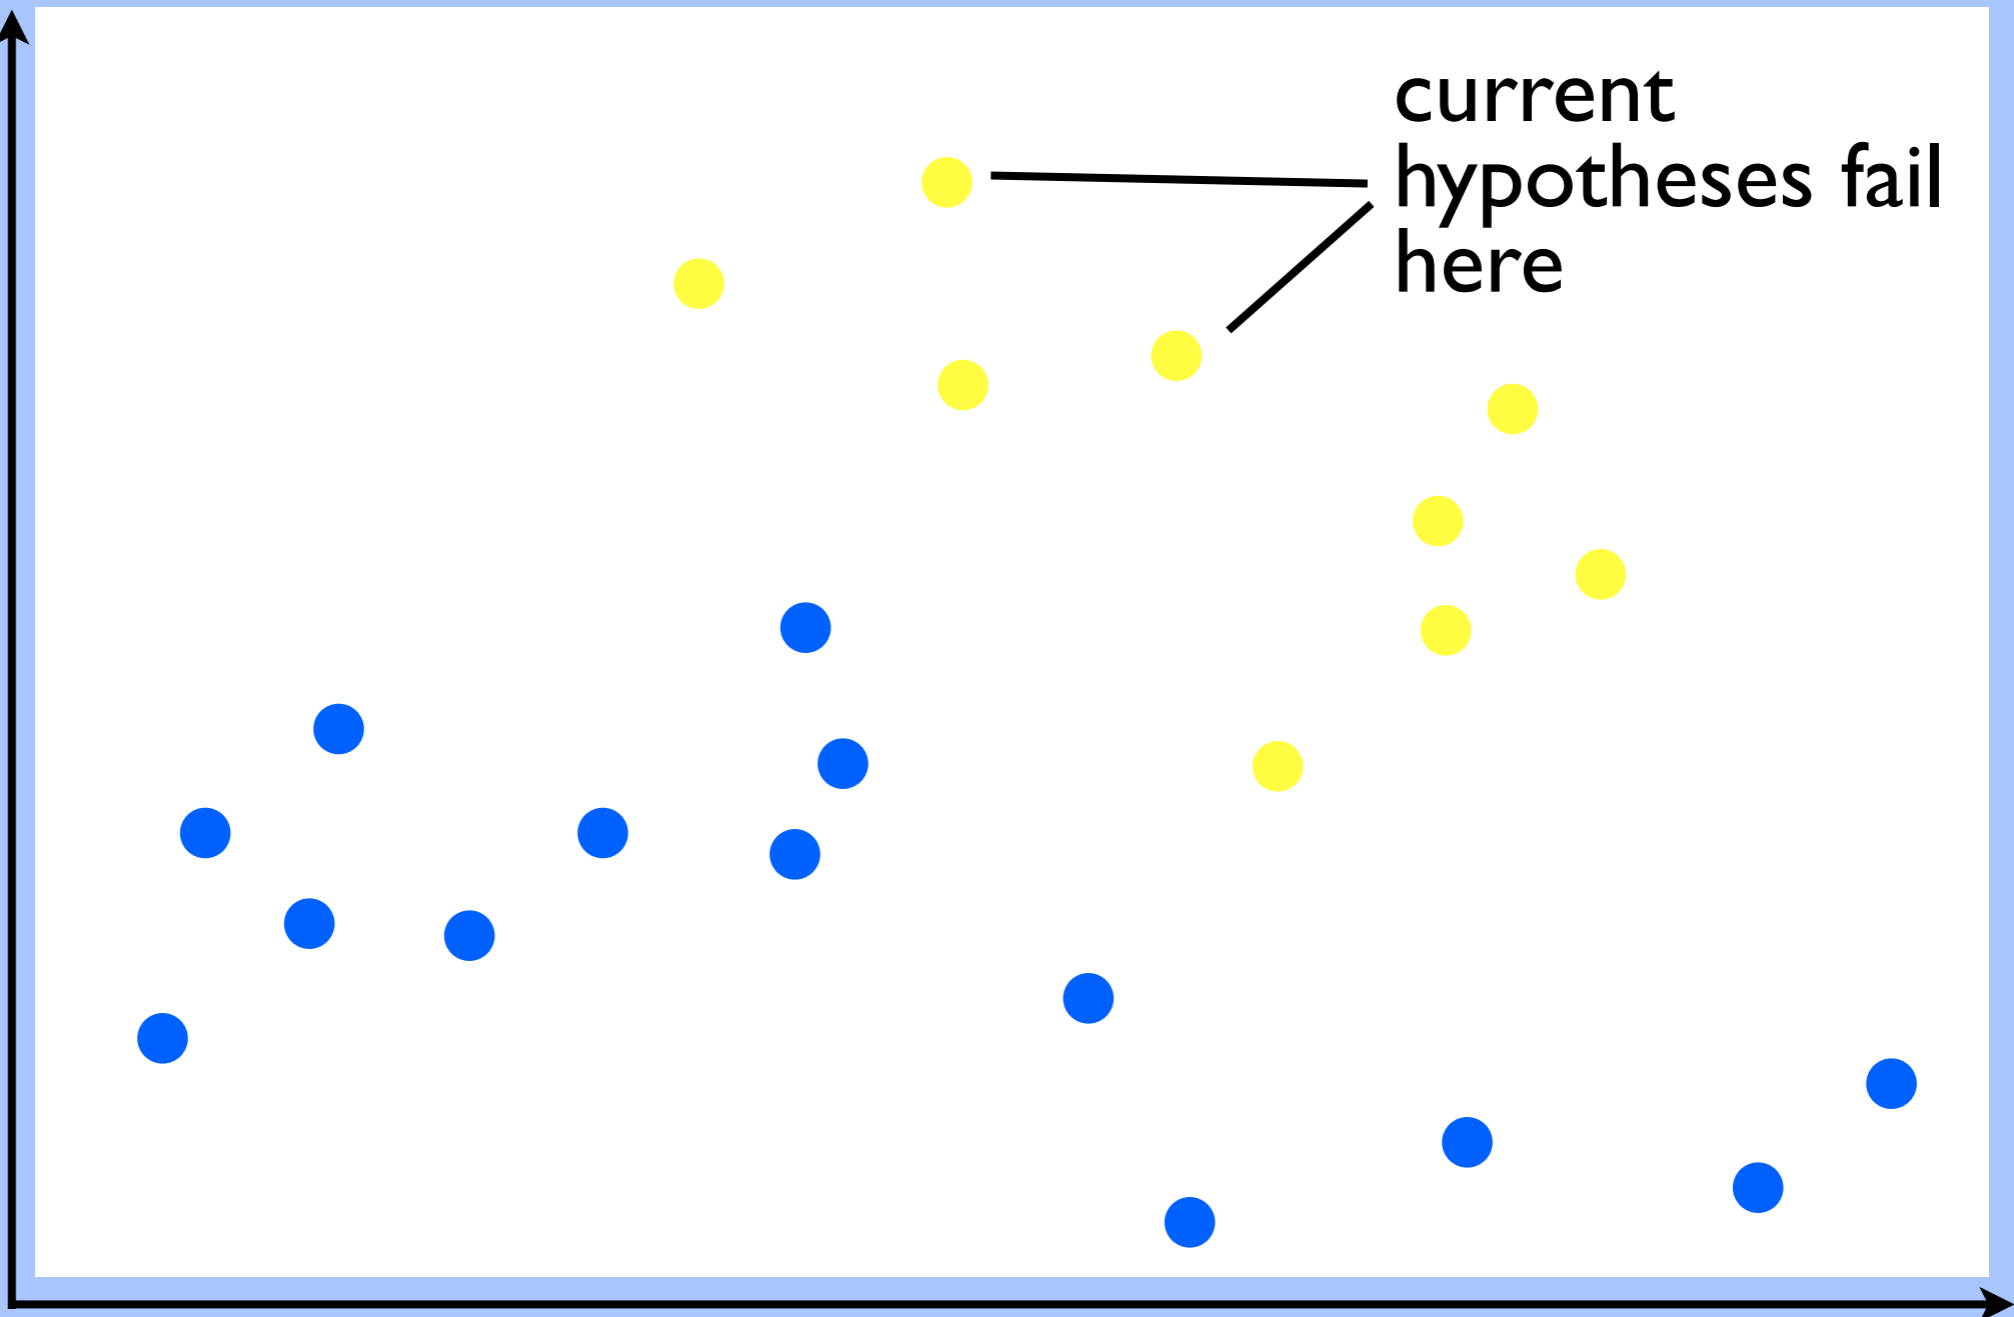
Learning

(discovering new, ever more general hypotheses)



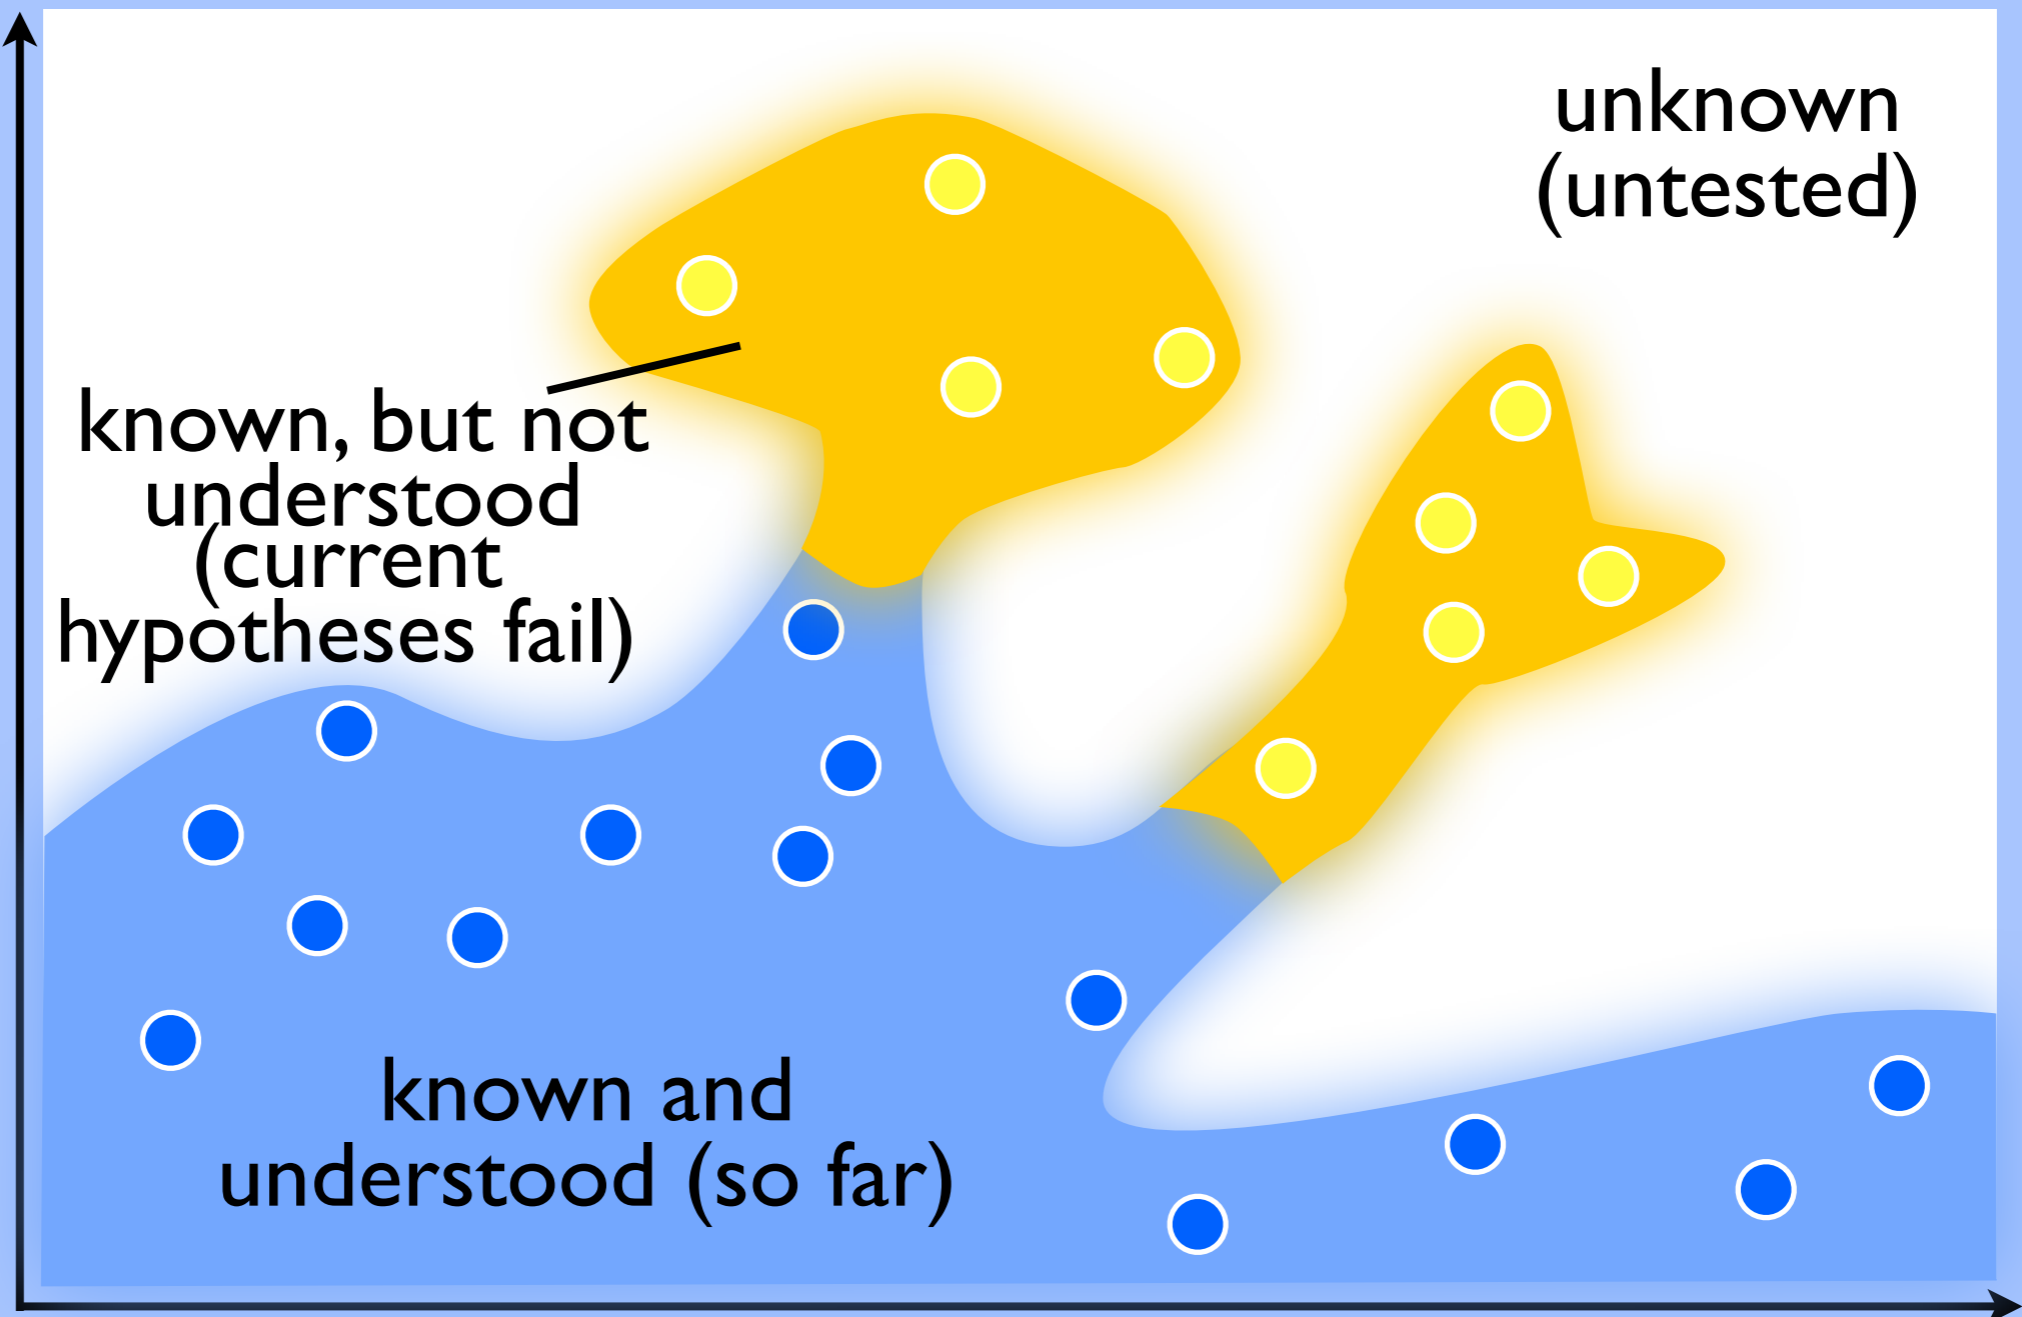
Learning

(discovering new, ever more general hypotheses)



Learning

(discovering new, ever more general hypotheses)

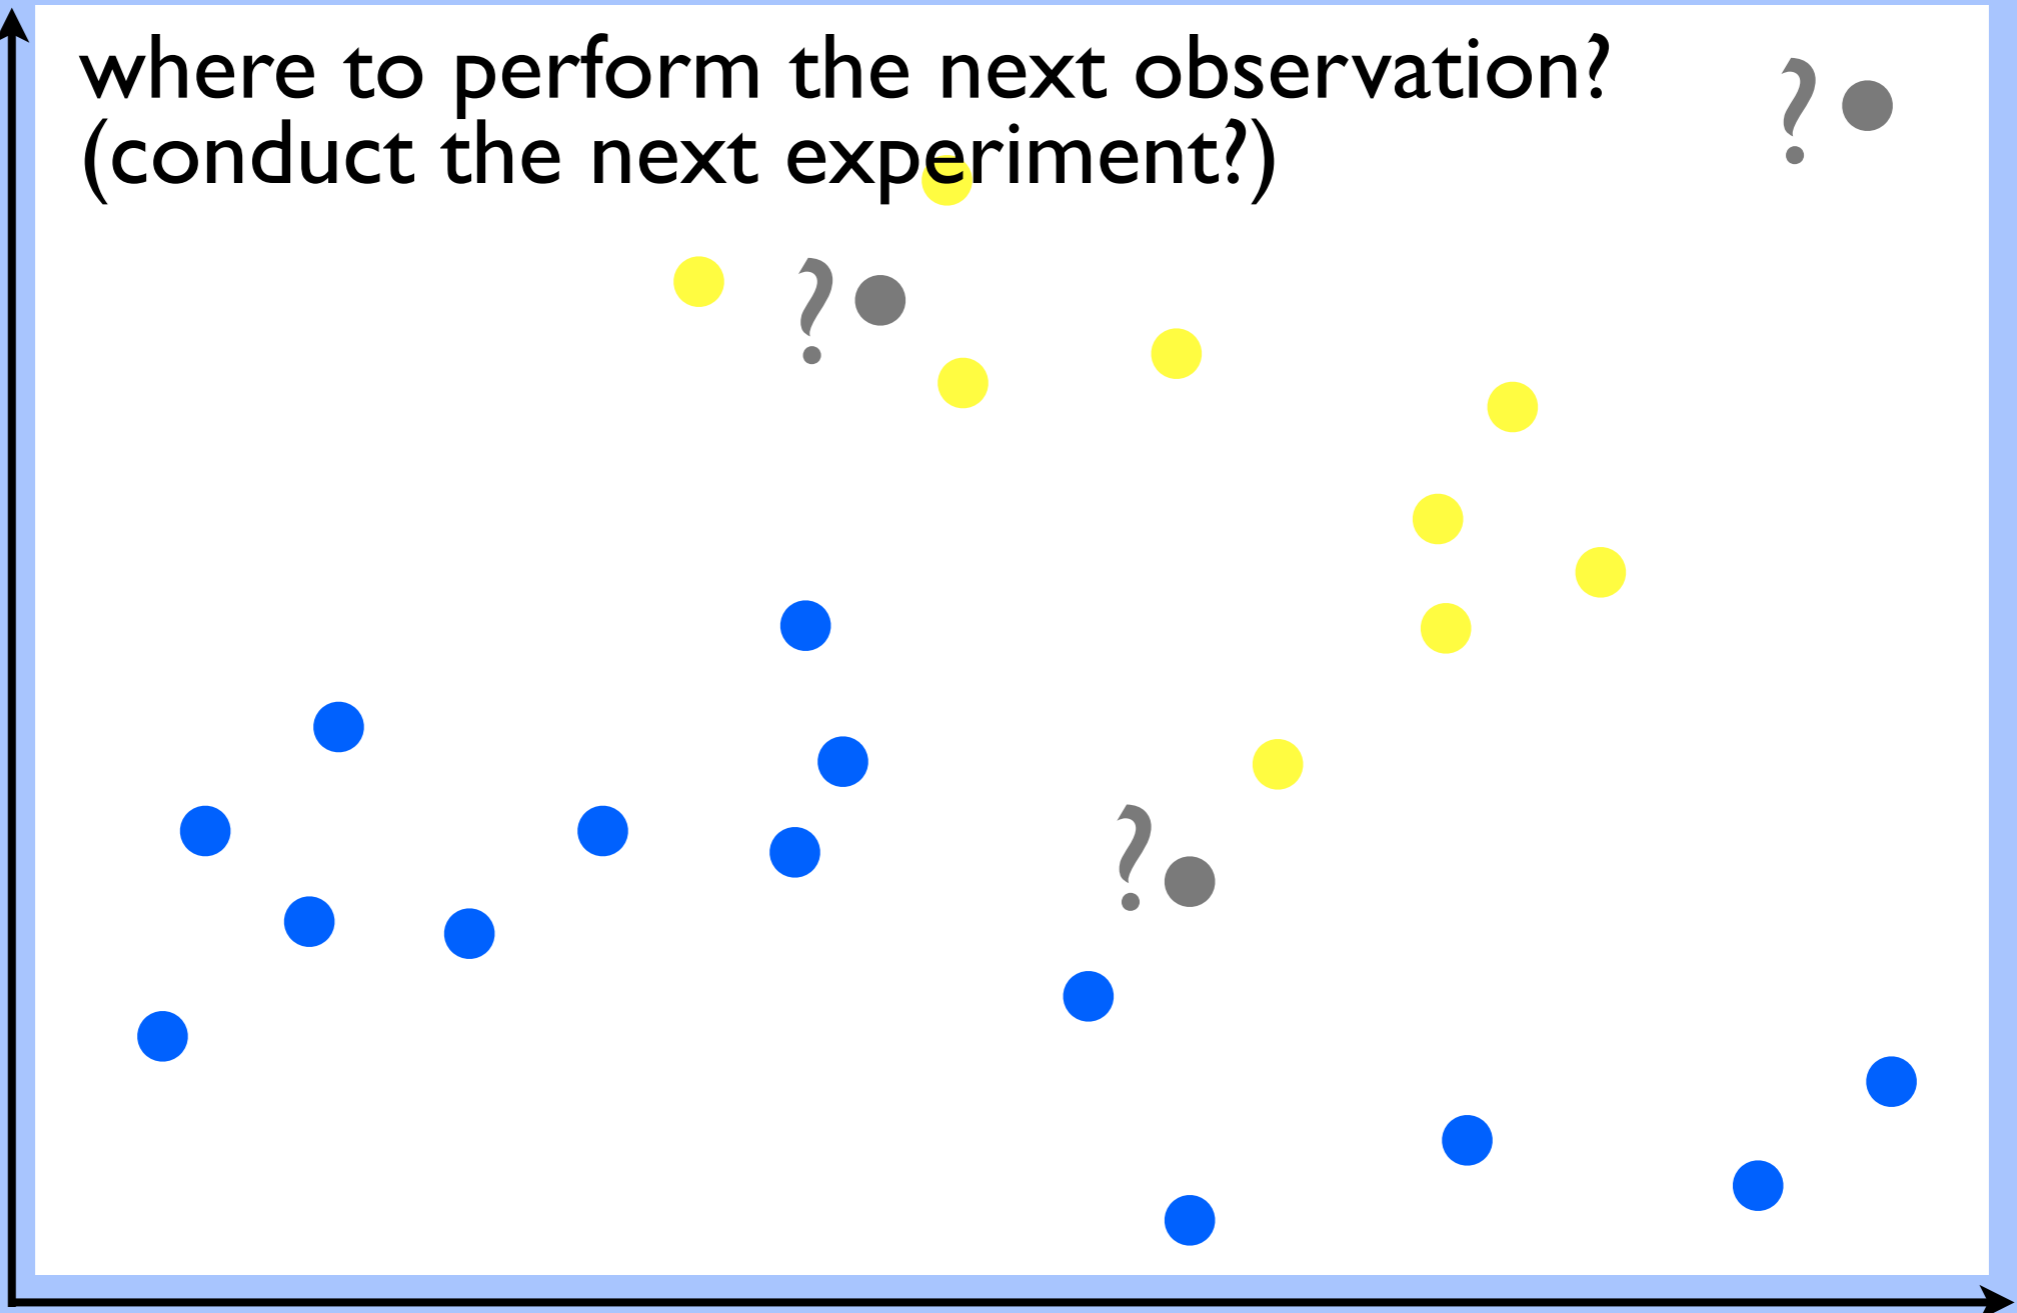




Learning

(discovering new, ever more general hypotheses)

where to perform the next observation?
(conduct the next experiment?)





Learning

(discovering new, ever more general hypotheses)

- Active learning:** learner/scientist selects observations to be made (selects “training data”)
- at the border “understood” vs. “not understood”
 - or, more precisely: where one might learn the most
 - i.e.: parameter values where the uncertainty of predictions is very large (e.g. existing hypotheses leave a large freedom in outcome), and where refining them would have significant consequences in many areas of parameter space
 - also: parameter values where many hypotheses may be falsified (based on experience? based on “distance” to known examples?)
 - try to falsify (or make more plausible) especially those hypotheses that are most useful, because they are:
 - most simple and general
 - can be applied usefully if true (e.g. to obtain a new very sensitive measurement method, or to simplify predictions)



Learning

(discovering new, ever more general hypotheses)

Select (numerical) experiments not only according to possible reward (new insights), but also according to effort! [have to learn to predict effort]

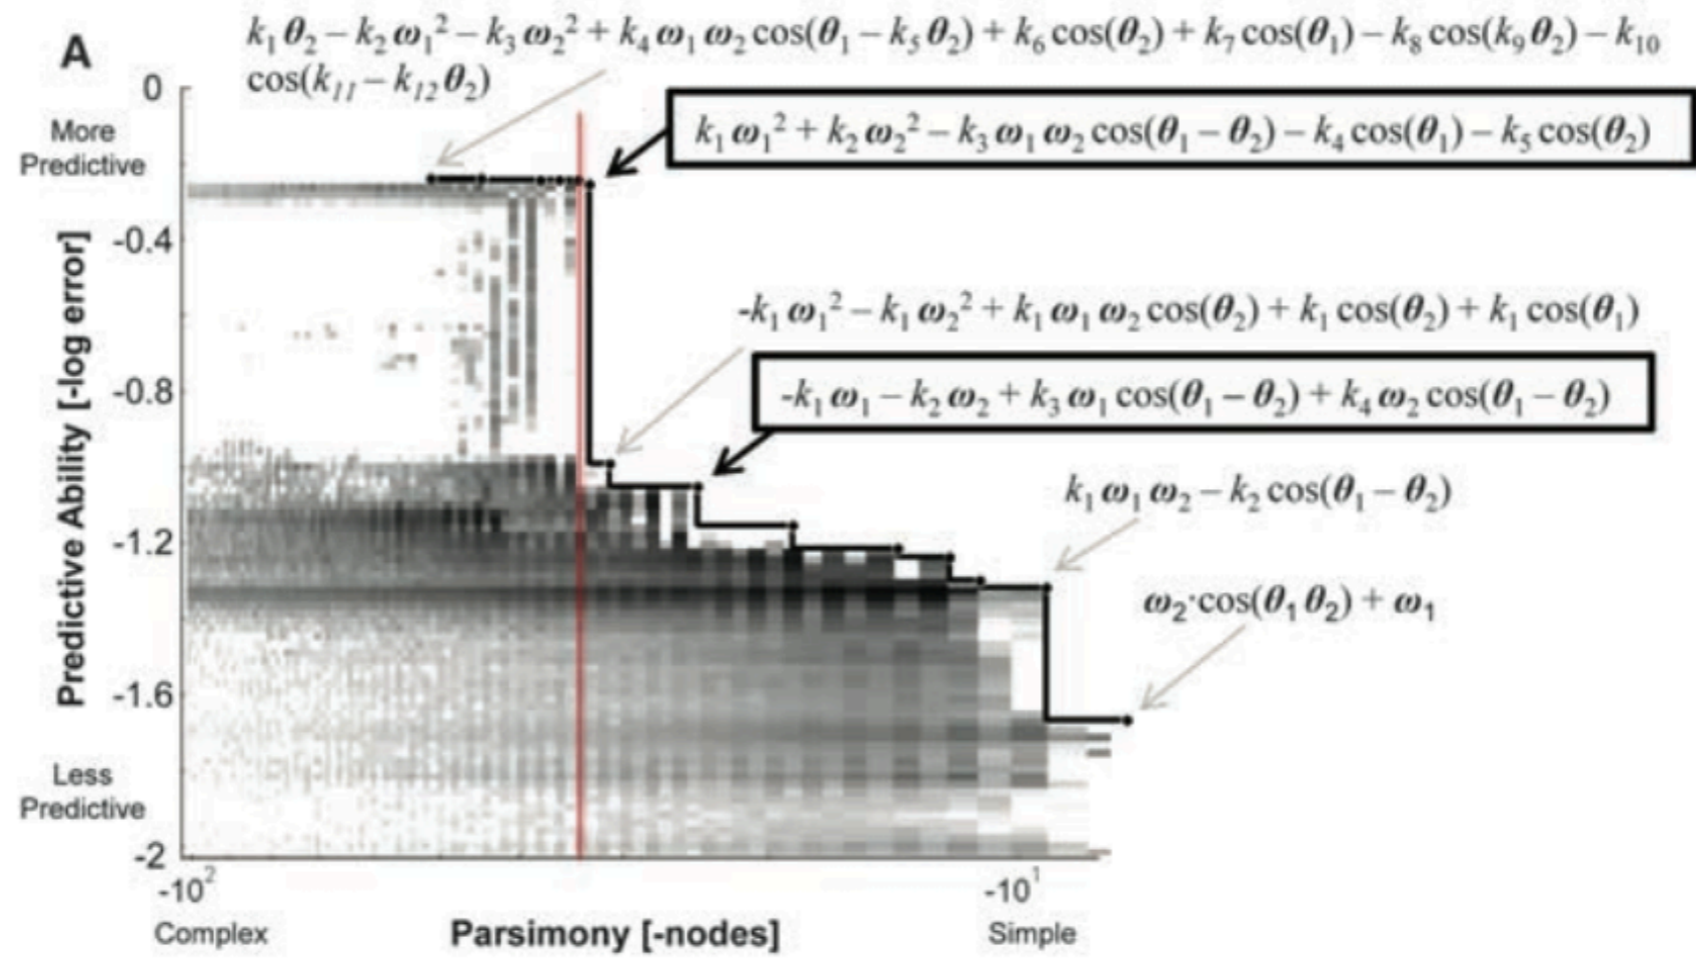
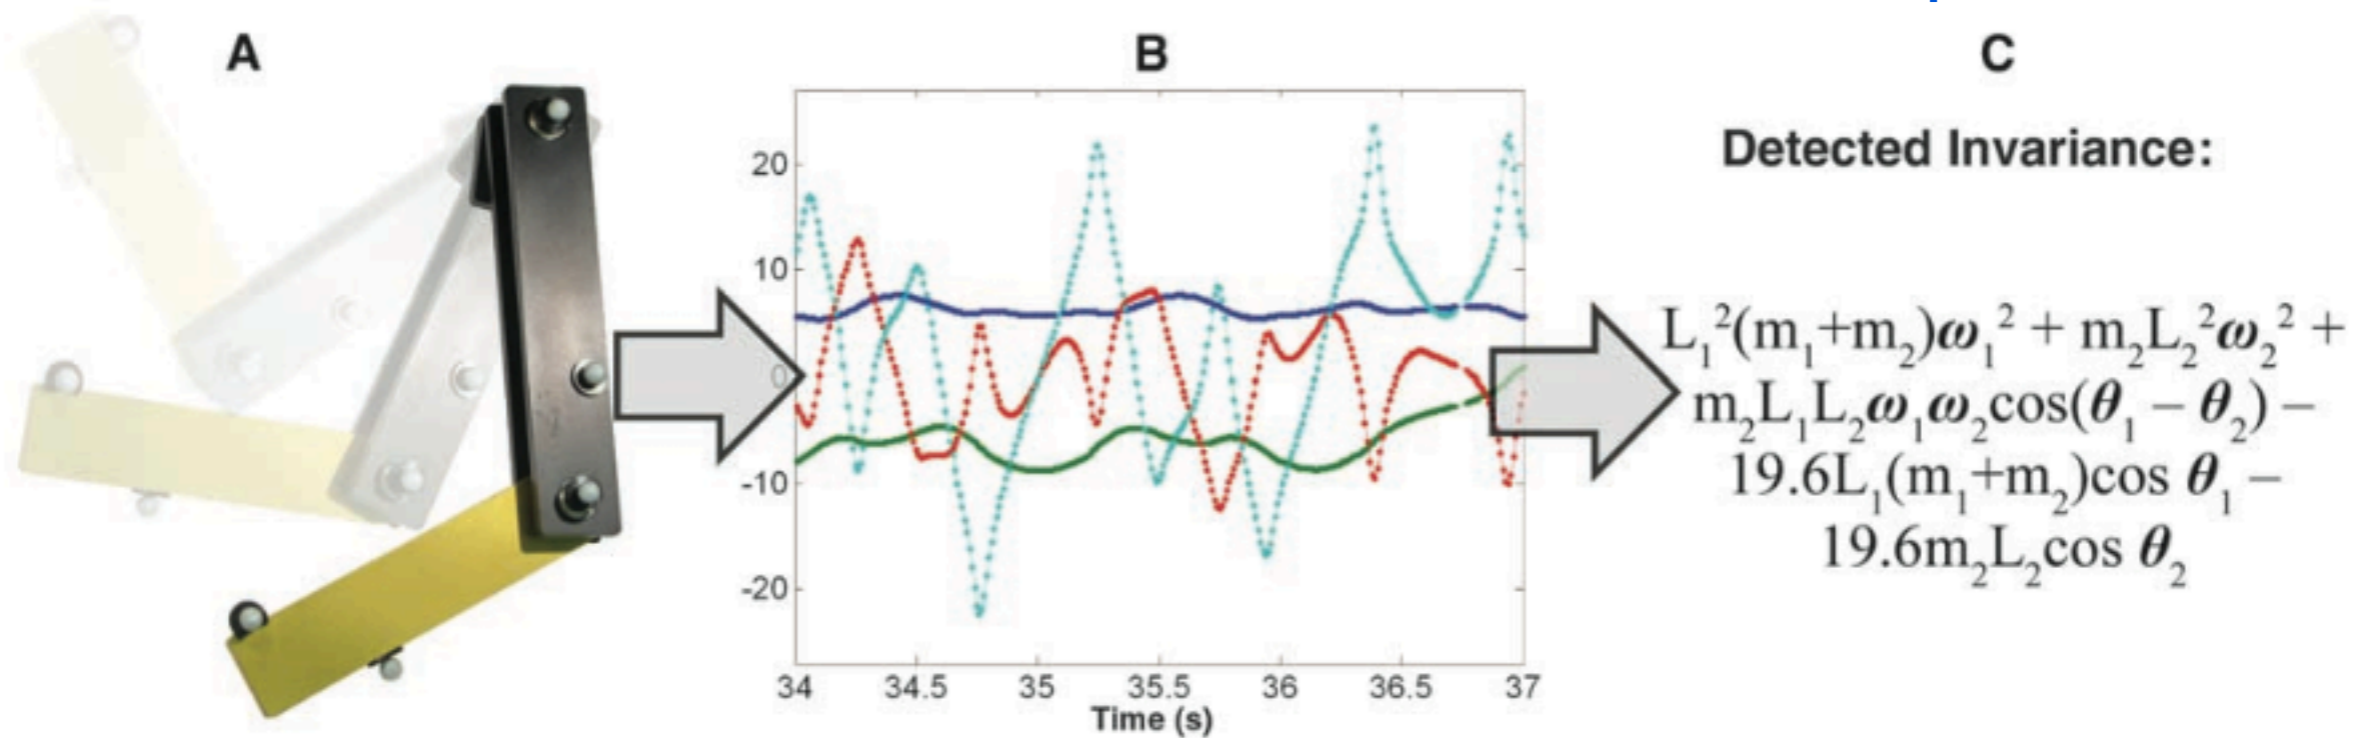
Experiments: Time [1 nsec - several months], costs, energy consumption, component requirements (for fabrication, setup, measurement time, calibration)

Numerical experiments: Time [1 μ sec - a week], number of cores, energy consumption

Hypothesis generation steps: e.g. effort needed for logical/mathematical deductions, for running a neural network etc.

Distilling analytical laws from experimental data

M. Schmidt and H. Lipson, Science 2009



“Adam”, the Robot Scientist


King et al., Science 2009

- Generates and tests hypotheses about the role of various genes in yeast (“functional genomics”)
- time-consuming tests: select strains of microbes (with certain genes modified) and let them grow on selected media over several days
- Robot can initiate about 1000 strain-media combinations a day
- “Adam” consists of three liquid-handling robots, three robotic arms, three incubators, a centrifuge, a freezer, and more equipment; plus a computer to control everything
- model=collection of facts about genes, enzymes and metabolites (small molecules)
- hypothesis forming: which gene might encode an enzyme? [take clues from other organisms]; e.g. “2-aminoadipate transaminase” might be encoded by either YER152c, YJL060w or YJL202w genes
- assign probabilities to hypotheses, select experiments that minimize the expected cost for eliminating all but one hypothesis
- found some new gene-enzyme relations

“Physics 101”: Learning physical properties from unlabeled videos

Wu et al., “Physics 101 – Learning Physical Properties”, BMVC 2016


Inclined plane




spring



fall



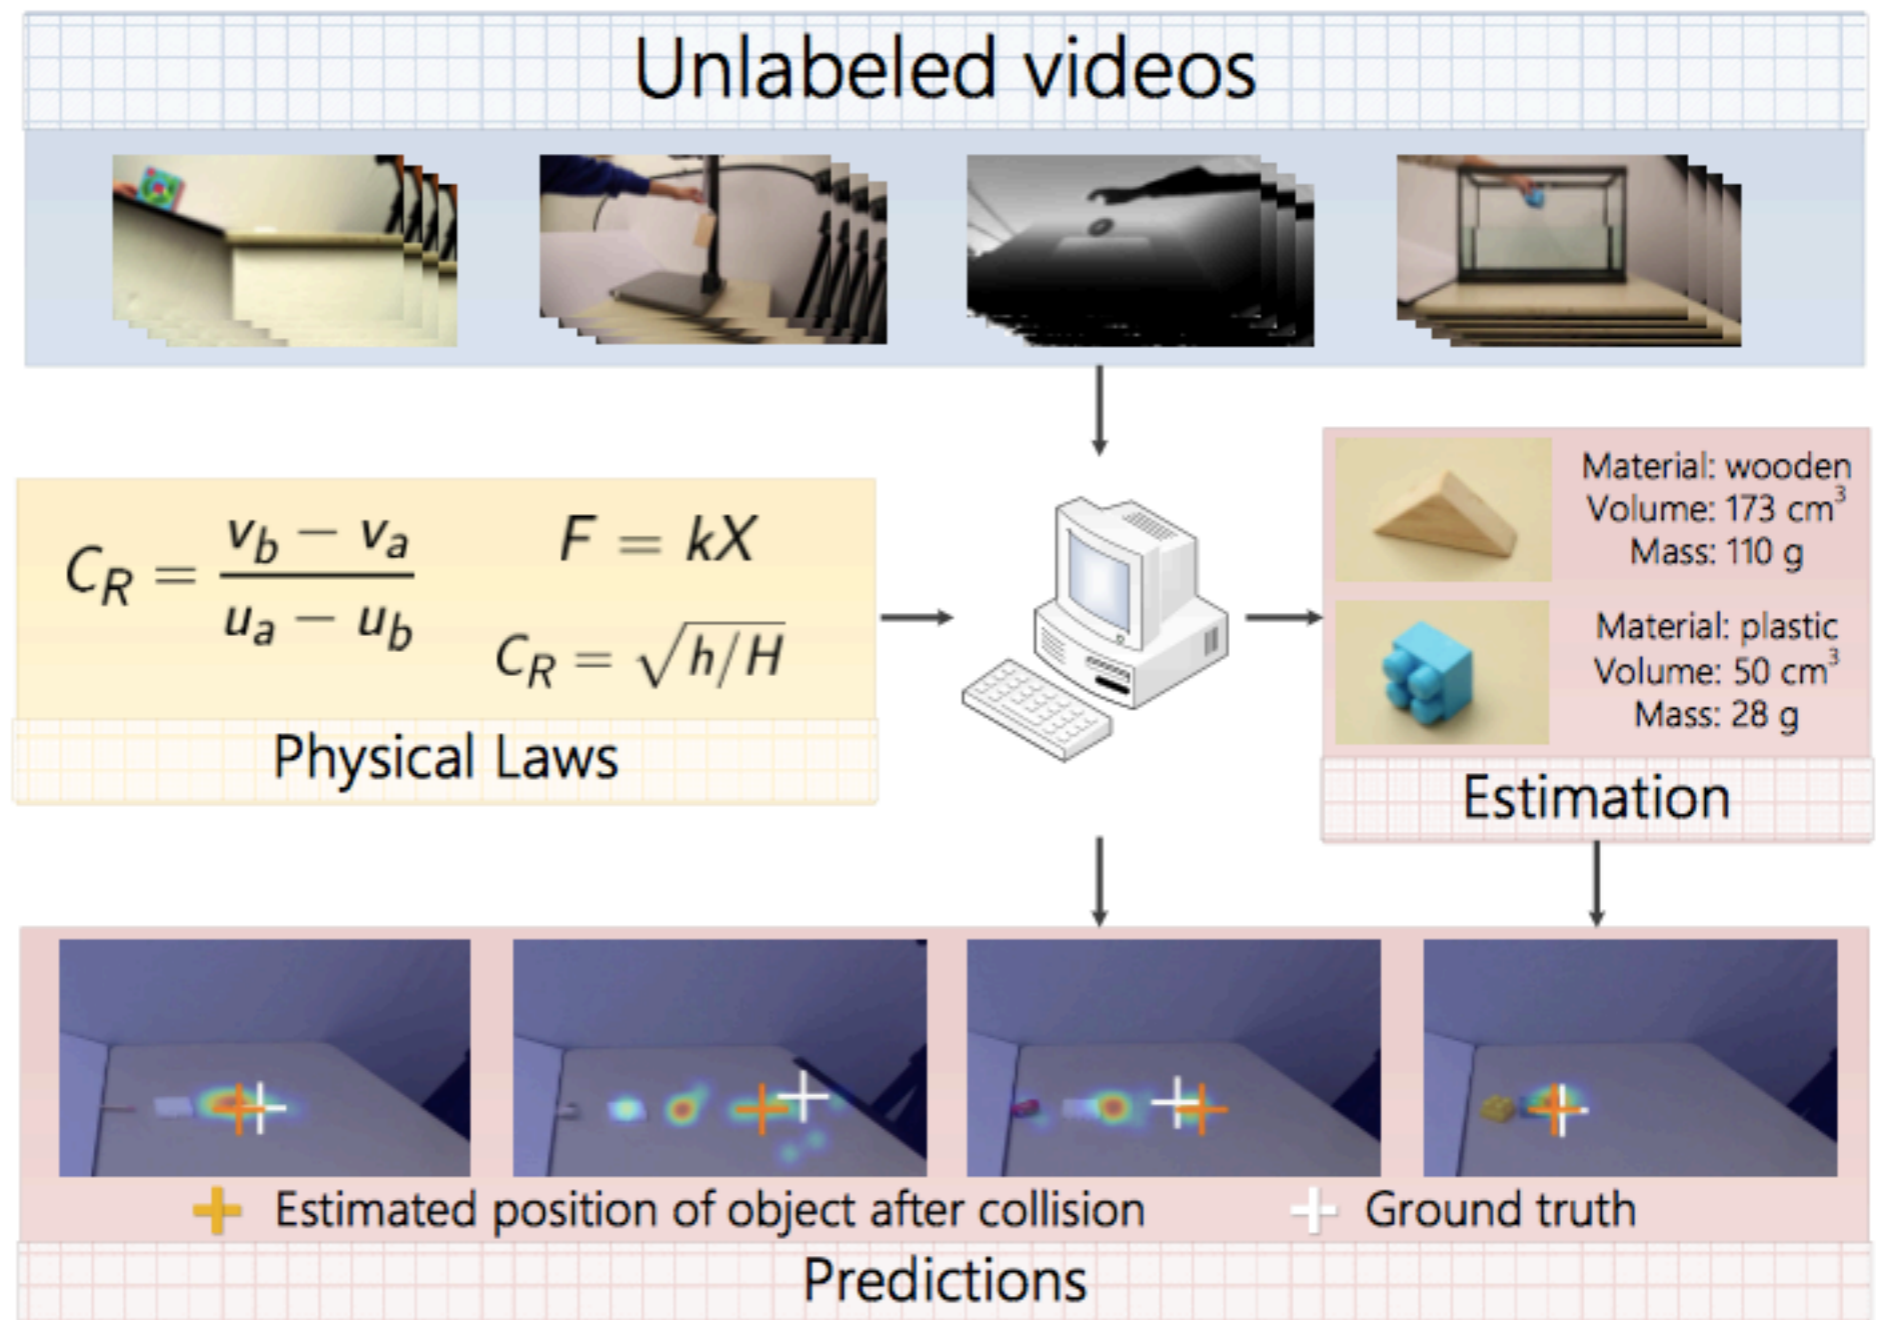
drop into water



> 17000 video clips

“Physics 101”: Learning physical properties from unlabeled videos

Wu et al., “Physics 101 – Learning Physical Properties”, BMVC 2016



How might science look like in the future?

A decorative graphic in the top-left corner of the slide. It features a grey line that starts from the top center, goes down and left, then curves down and right. A red line starts from the left edge, loops around, and ends near the grey line. There are three orange spheres: one at the top end of the grey line, and two at the bottom end of the red line.

- Right now: Machine learning / neural networks as a new tool, for making computationally cheap approximate predictions based on large training sets, or to discover hidden underlying structures (e.g. based on the image recognition power of neural networks)
- Later, for true scientific discovery AI: form human-computer teams, where the human could focus on selecting goals
- Important additional task for AI: Train the computer to explain conclusions in a way that is optimized for human understanding
- Very long-term, once we have superhuman scientific AI: ?

Function/Image
representation

Image classification
[Handwriting recognition]

Convolutional nets

Autoencoders

Visualization by
dimensional reduction


Recurrent networks

Word vectors

Reinforcement learning

Connections to Physics

Artificial Intelligence



<http://machine-learning-for-physicists.org>