# Connectionism, Unit 8

# Rethinking Innateness: Chpt 4

and

Intro to Dynamic Systems

Reading:

Van Gelder and Port, from Mind as Motion

Bechtel and Abrahamson article

Textbook, Chpt 4

# Introducing TIME: Recurrent Neural Networks

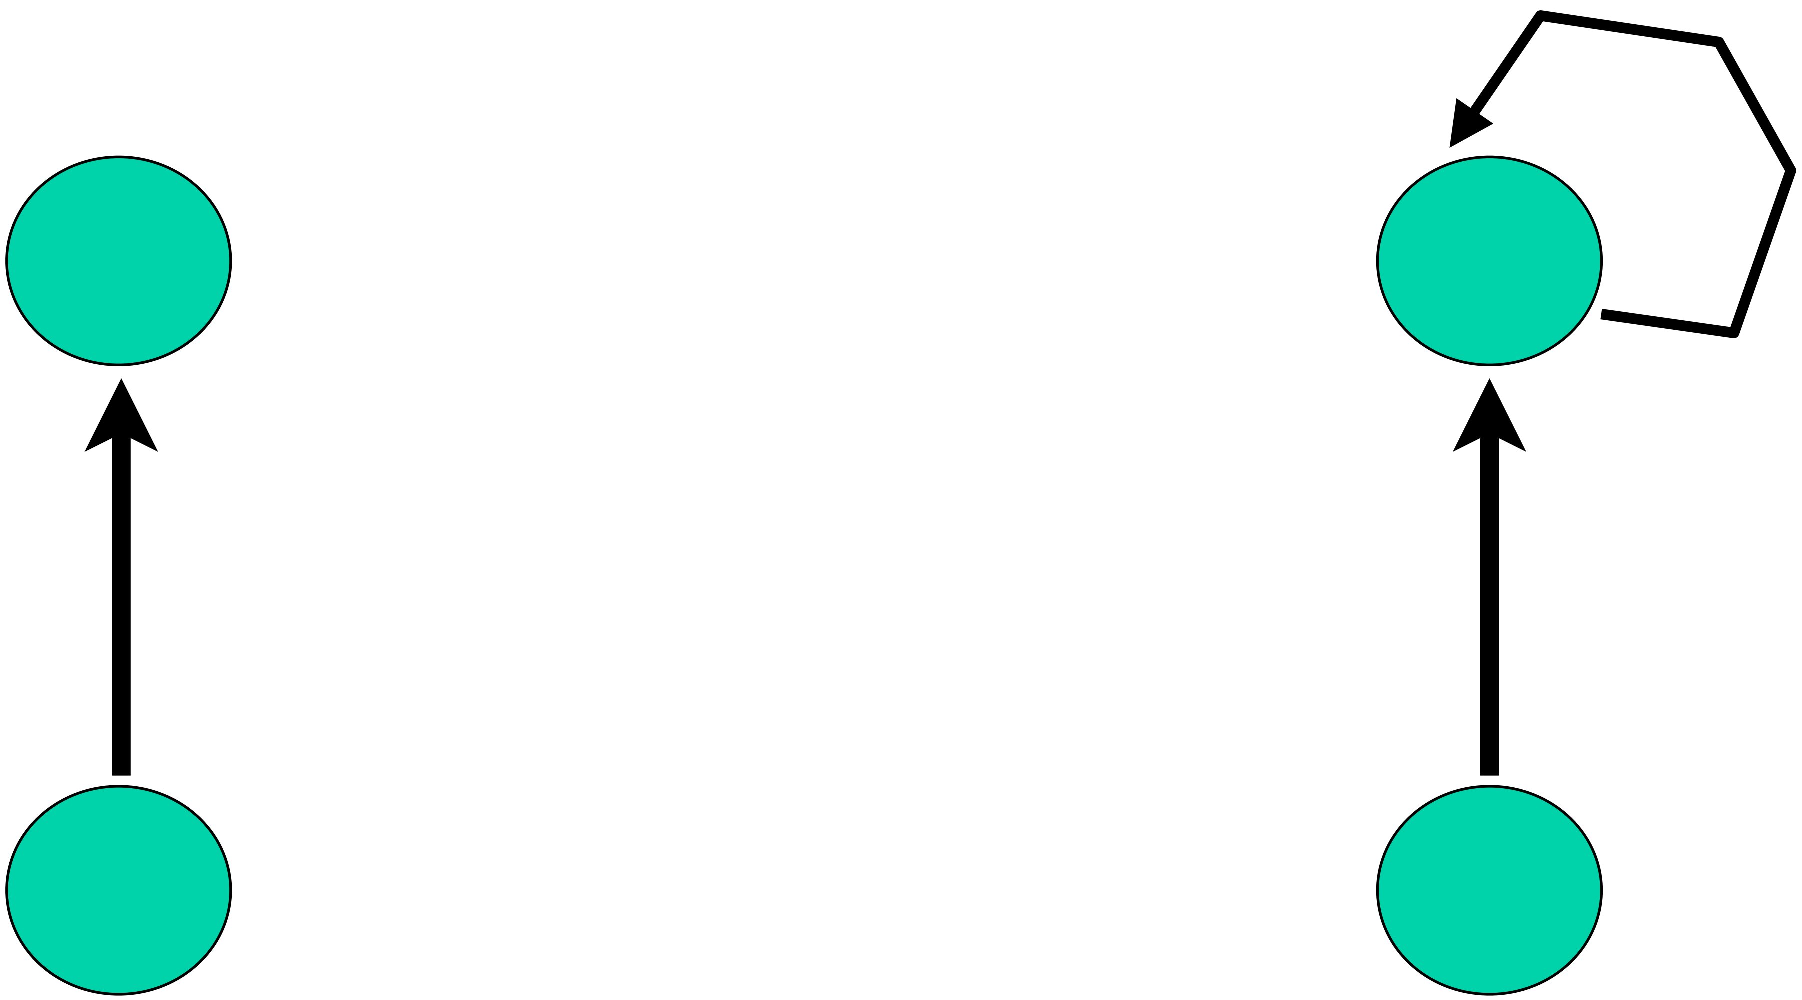

# Simple Recurrent Network (a.k.a. Elman Net)

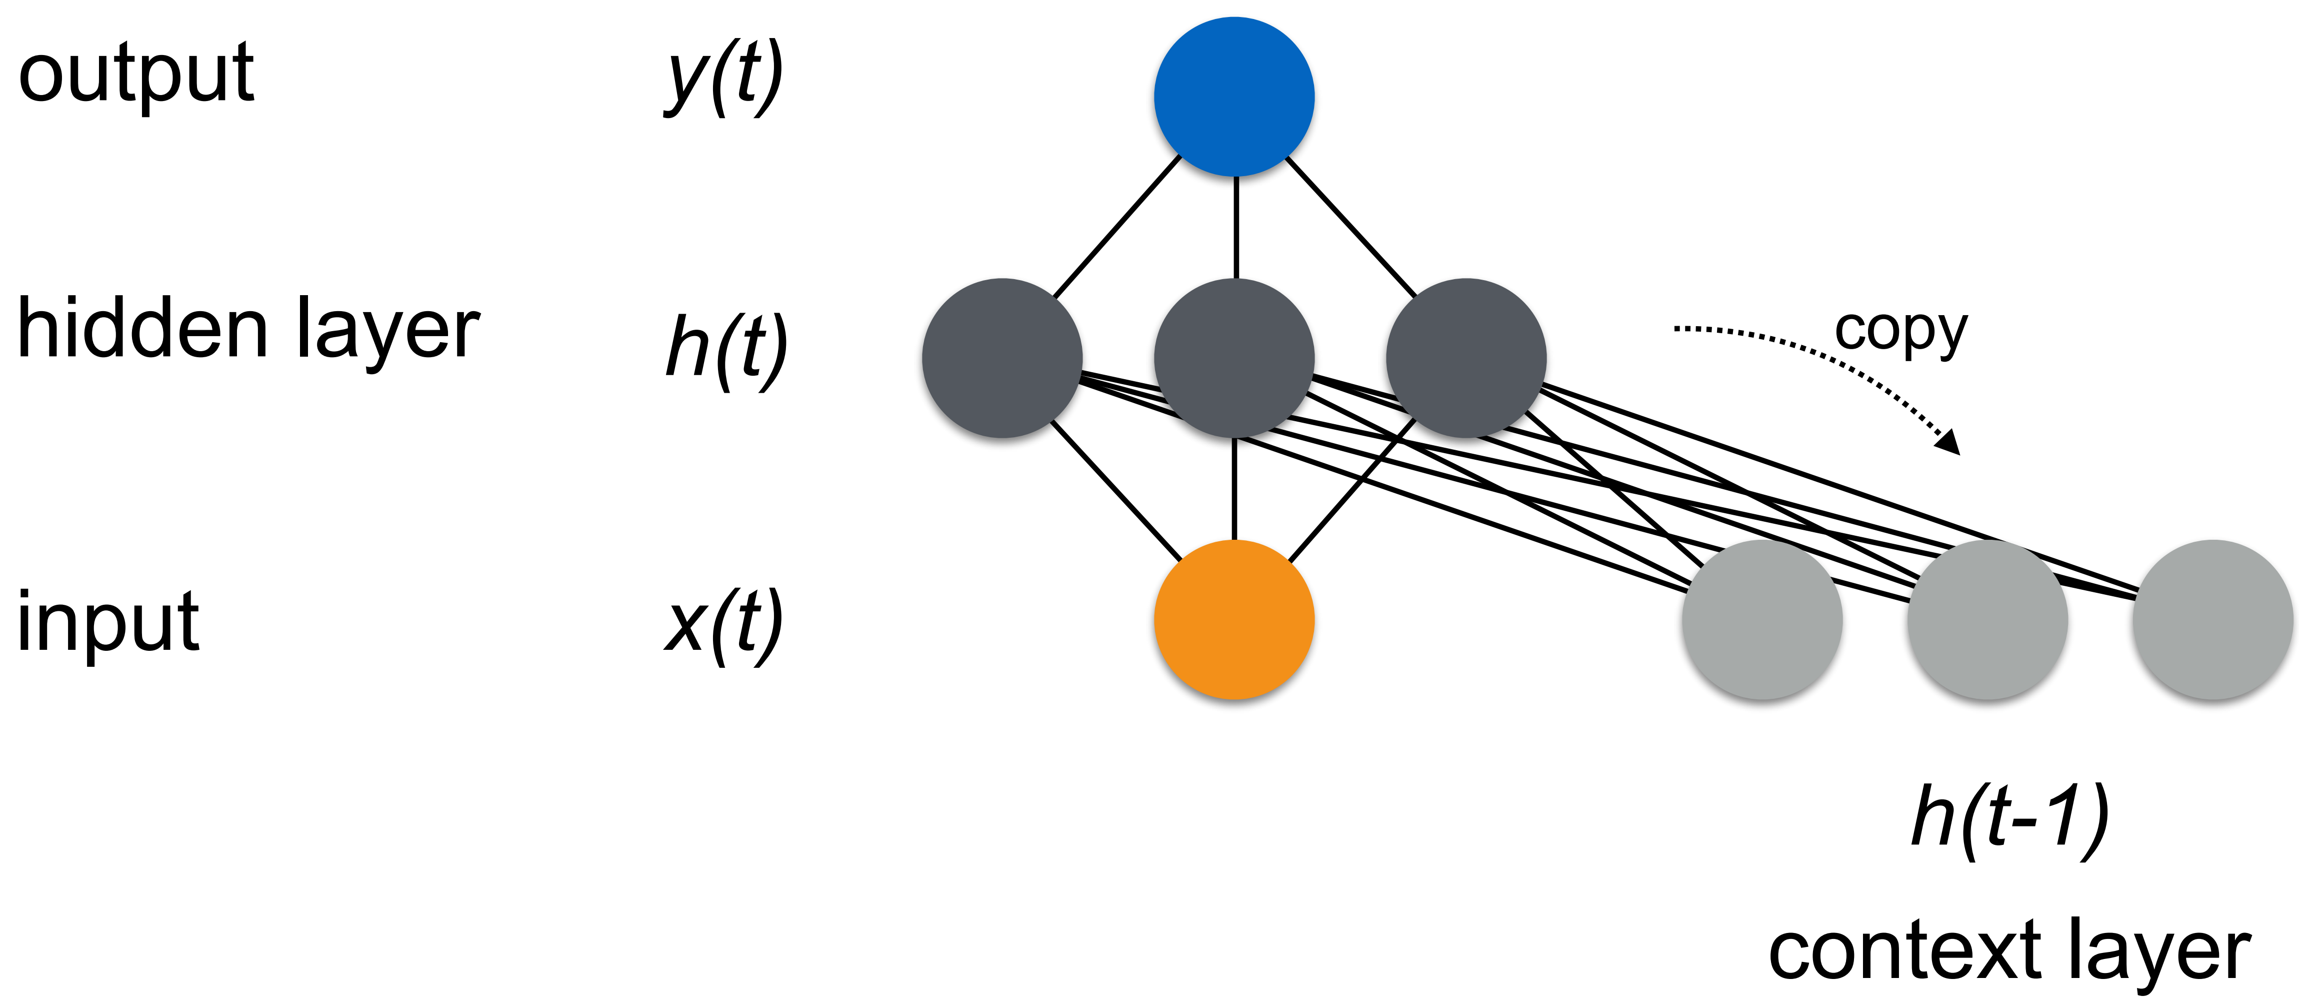


*Bias units not shown*

# Recurrent Neural Networks

- Recurrence fundamentally changes the network model
- A recurrent network is a dynamic system
- Its behaviour is spread out in time
- Dynamic Systems Theory is increasingly important in modeling cognition

# Elman's “phoneme” predictor.....

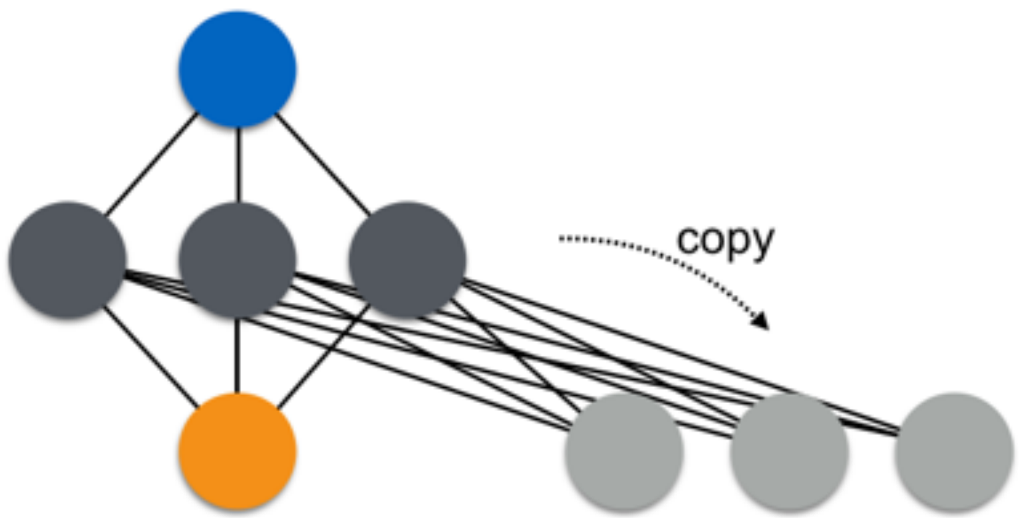


Simple
Recurrent
Network

- Task: predict next symbol
- Input:

manyyearsagoaboyandgirllivedbytheseathey.....

- Architecture is constrained a priori. No explicit information about phonotactics is provided

# Error curve for a trained network

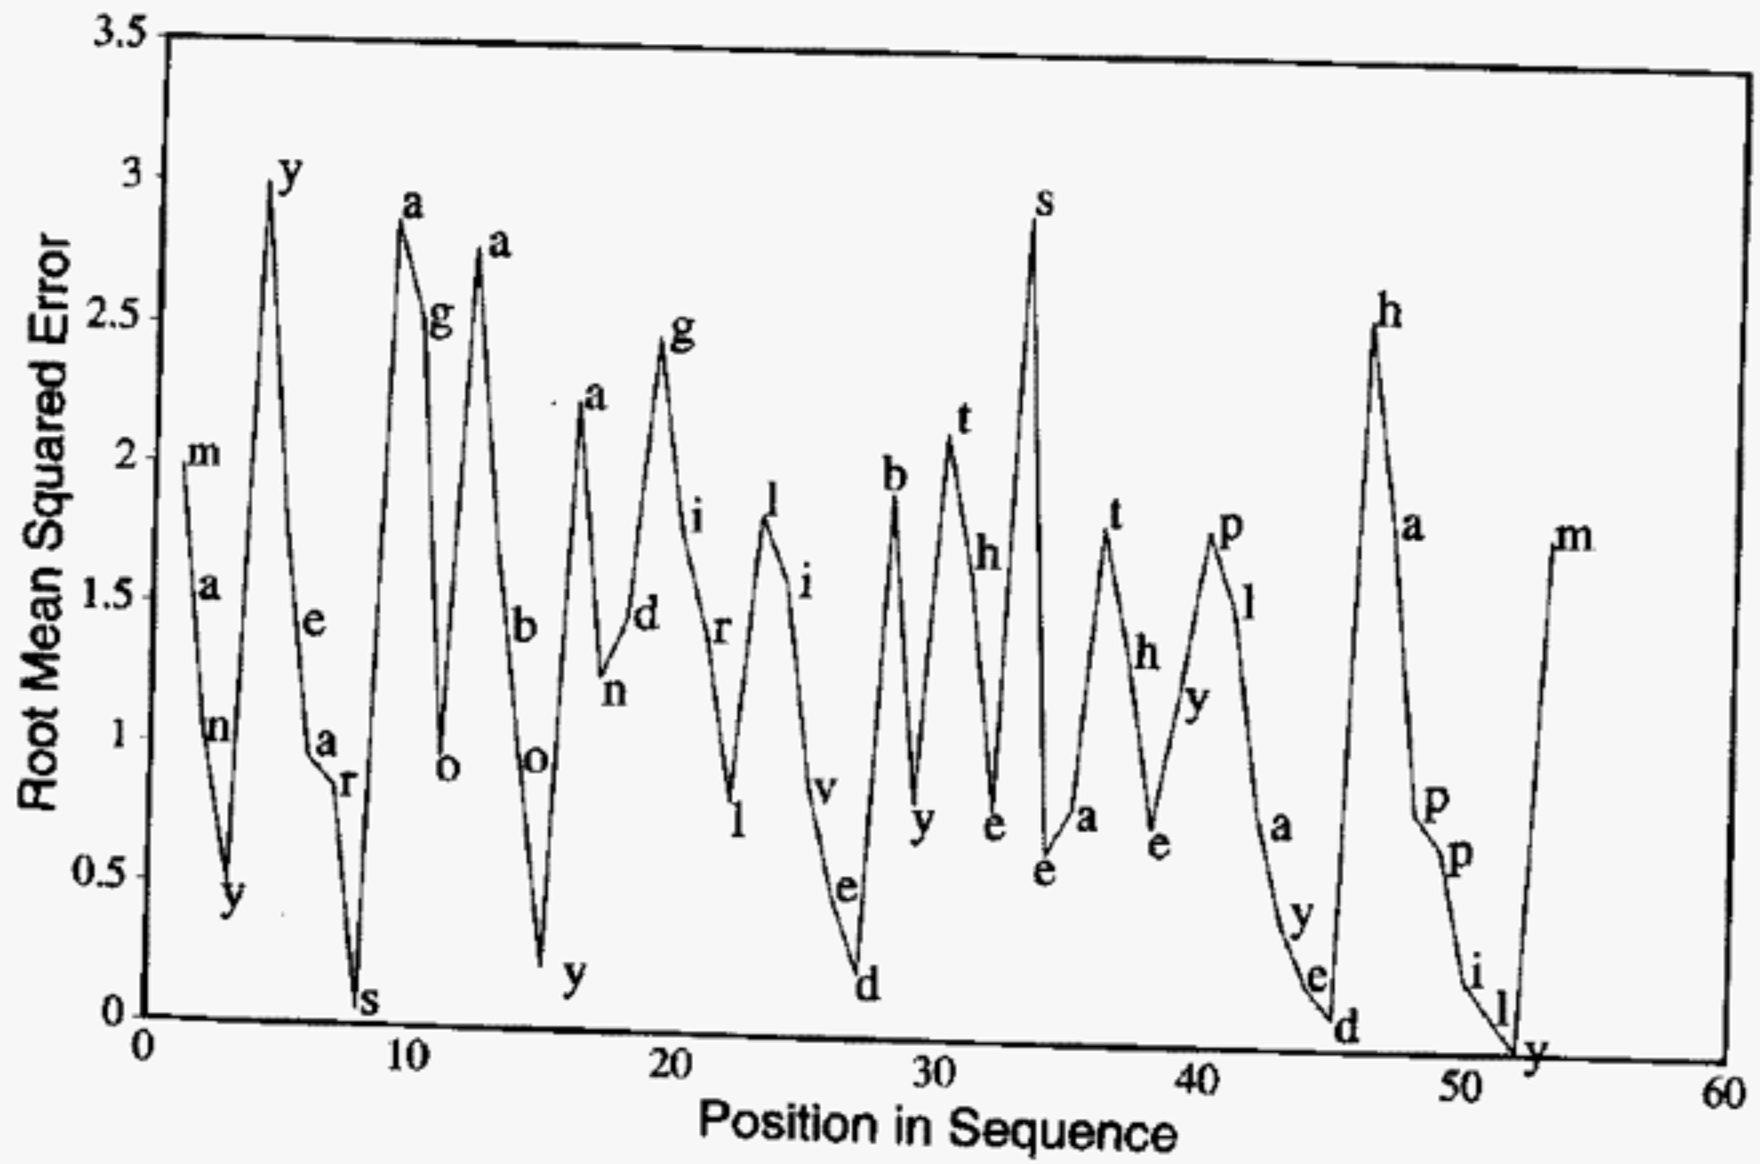


**FIGURE 3.3** The error curve for a network trained on the phoneme prediction task (Elman, 1990). Error is high at the beginning of a word and decreases as the word is processed.

# What was learned?

- Task is not perfectly solvable
- Uncertainty is highest at word boundaries
- statistics of the distribution of phonemes in the training set
- segmentation errors: "aboy", not "a boy"
  - c.f. "the nelephant" from "an elephant"
- inputs are arbitrary, localist (no similarity relationships possible)

# wow?

- representations available to the child are richer (phonemes group into classes, syllable structure, accent, intonation all contribute)
- representations available to the child are poorer (phoneme segmentation is not a given, linguistic and non-linguistic information are mixed)
- Is the "next symbol" task similar to what a child does?
- Does a description of the statistical distribution of symbols in a set require a neural network?

Prediction is one plausible way of doing error-corrected learning, without a God-like teacher.

Compare to auto-association.....

# Chpt 4 topics:

- Types of change
- Rate of Change
- Dynamic Systems

## Patterns of developmental change

- Development, and the emergence of the mind, is characterized by change at many levels and timescales
- Patterns of change contain essential information about the underlying processes and their interactions
- Studying change over time takes us into Dynamic Systems Theory
- DST has emerged as a competitor to the established computational account of mind and brain

# Dynamics and Connectionism

- Change during training
- Change during processing (recurrent neural networks)
- Treatment of time:
  - sampling at points throughout training
  - ordinal scale
  - interval scale
- Discretization ("time step", $\Delta t$ )

# Linearity

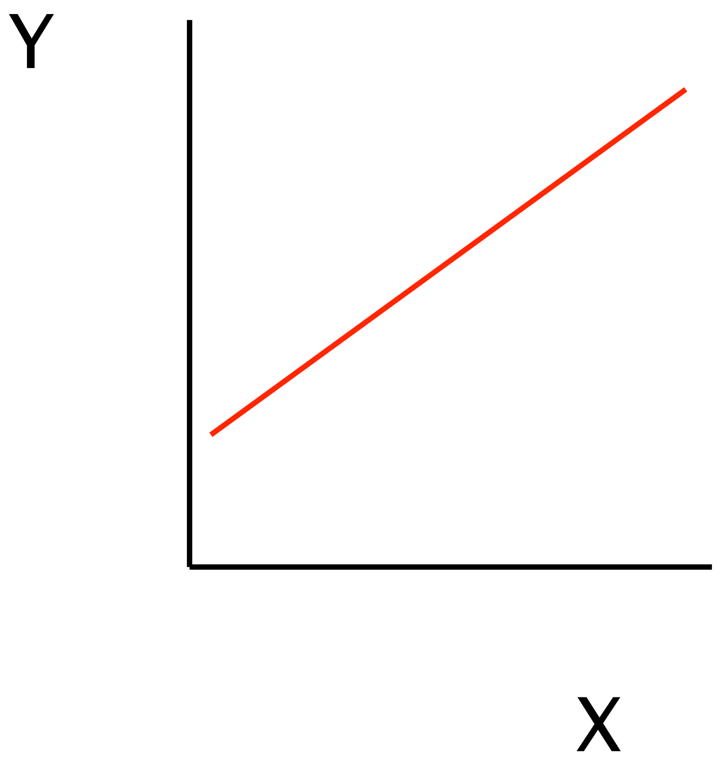


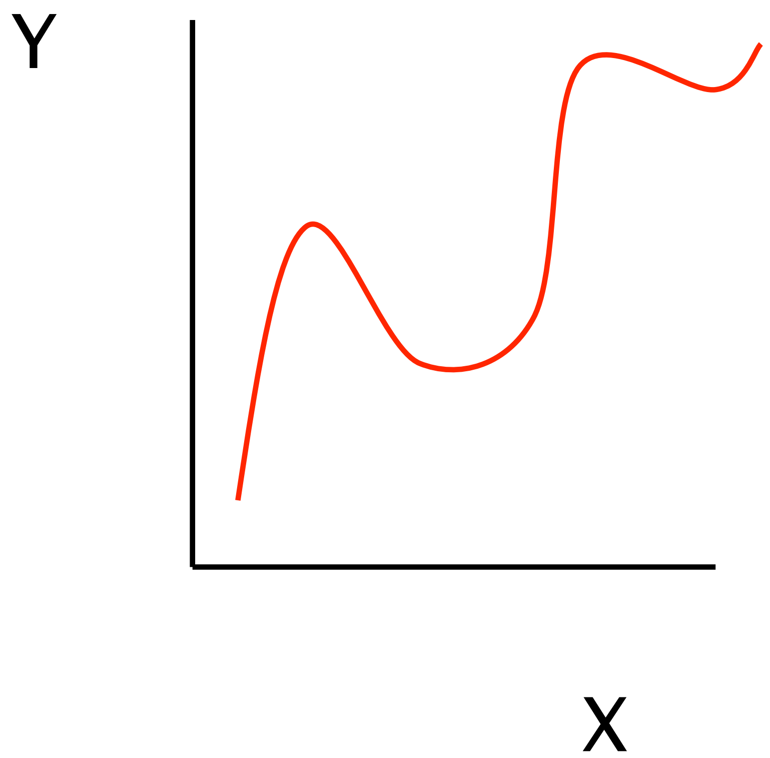


Simple to model

Rare?

Simple…complex…intractible

May be modelled quantitatively or qualitatively

# Dynamical systems

- A system which can be described at any given time using a set of numbers (the system *state*)
- The state of the system changes over time (a *dynamic)*

Examples: Child-on-a-swing

Ireland's economy

Your brain

A falling marble

A Recurrent Neural Network

We need some tools for talking about systems that change over time

Dynamic Systems Theory provides us with those, but it makes much use of differential calculus

Anybody remember differential calculus??????

Intuitive introduction to DST based largely on:

Dynamics: The Geometry of Behavior
Ralph H. Abraham and Christopher D. Shaw
Addison Wesley, 1992 (2nd ed)

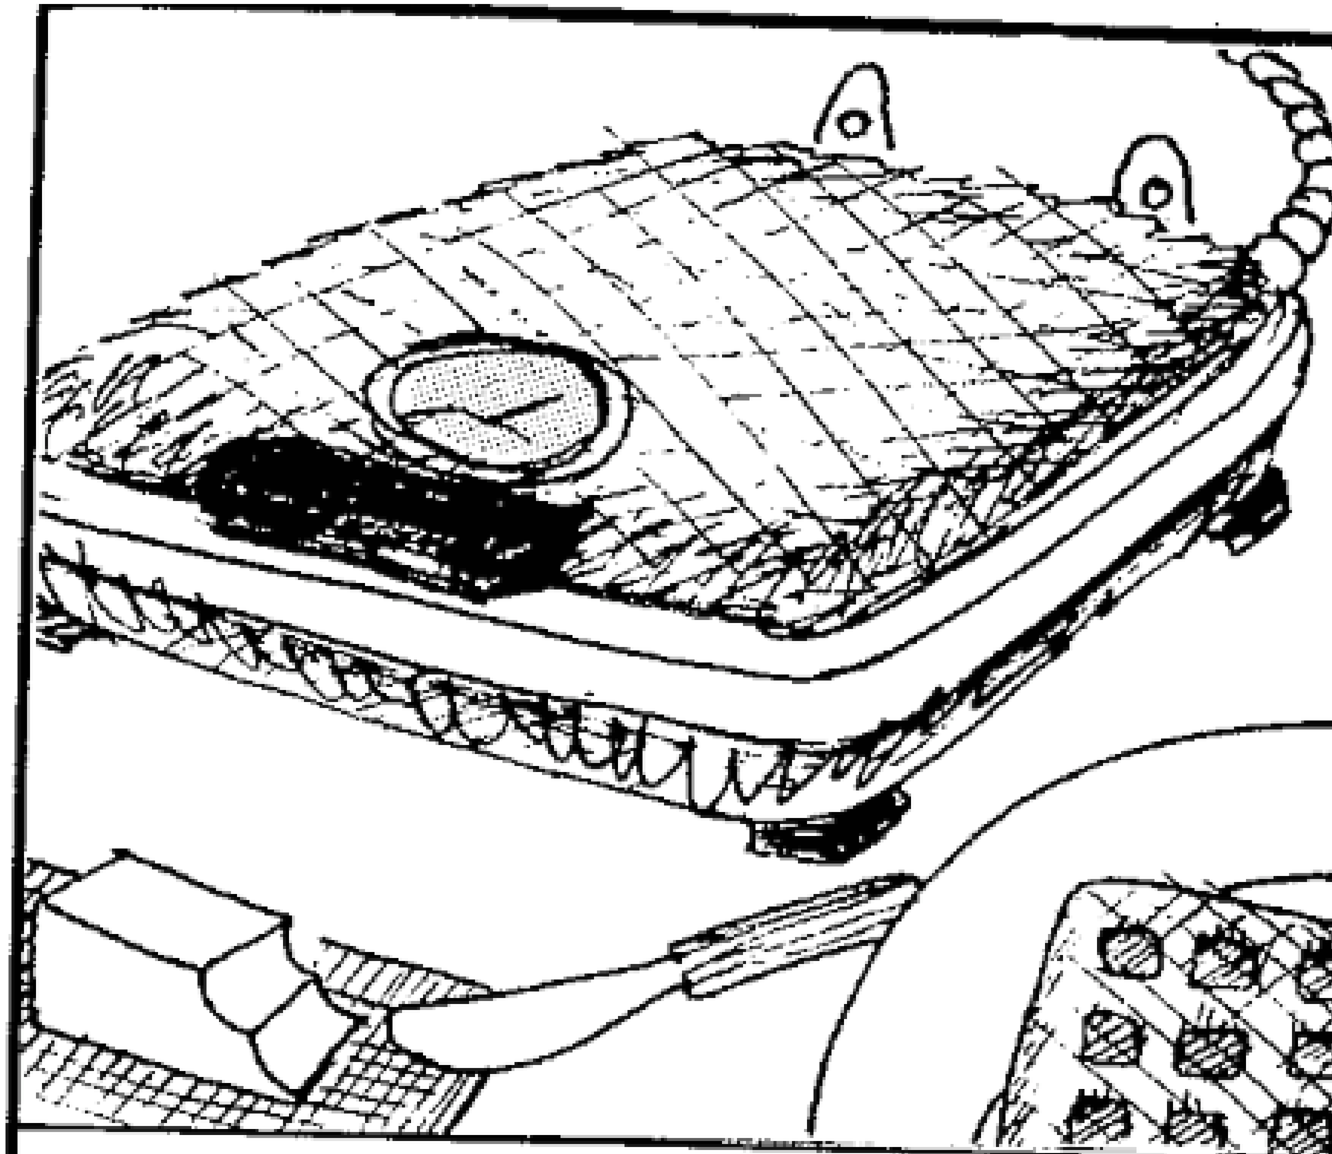

**1.1.1.** The actual state of this waffle iron cannot be described completely by a single observable parameter, such as the temperature. But usually we find it convenient to pretend that it can. This pretense is an agreement, the *conventional interpretation*, within the modeling process. It is justified by its usefulness in describing the behavior of the device.

We could describe the state of this waffle iron in many ways.

To model its behavior as a waffle iron, temperature might prove to be a useful variable.

Selection of state variables is a modelling decision, not a matter of fact.

What is the system?

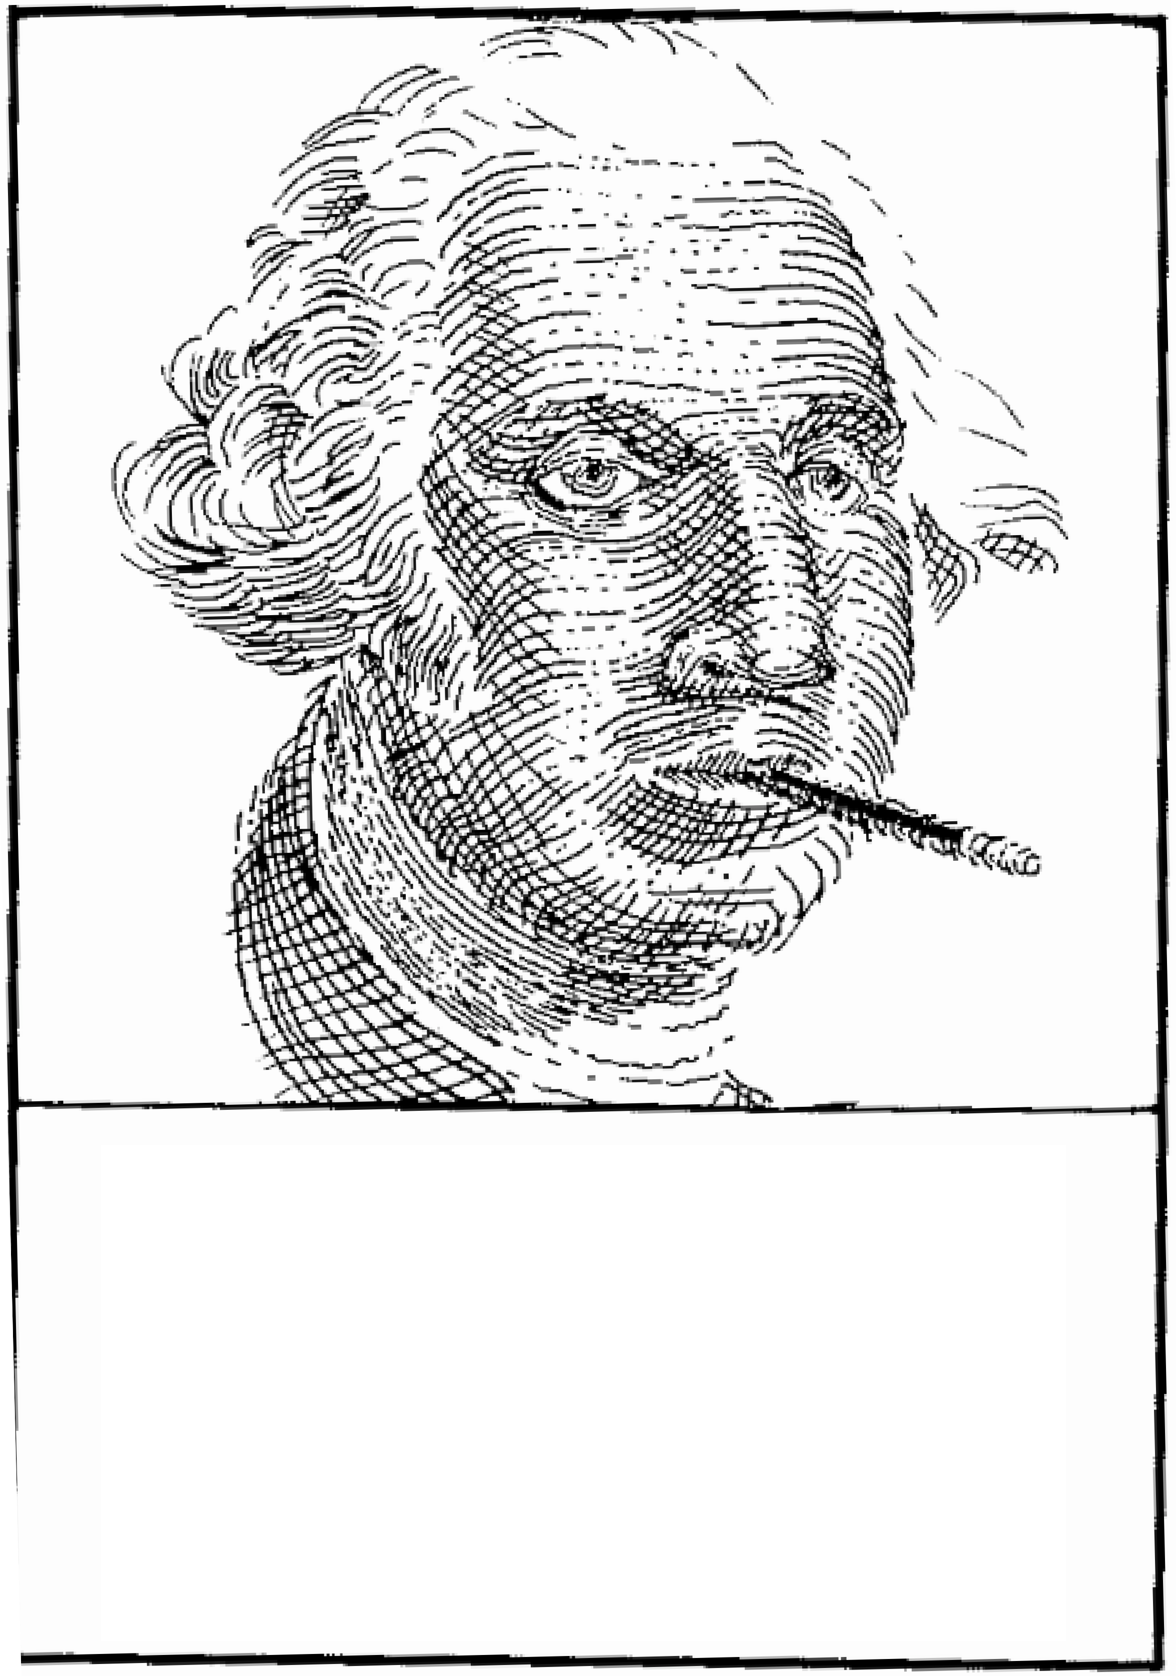

Any variable we choose may or may not be a good indicator of the relevant state of a complex system.

George's temperature probably correlates better with his health than his honesty.

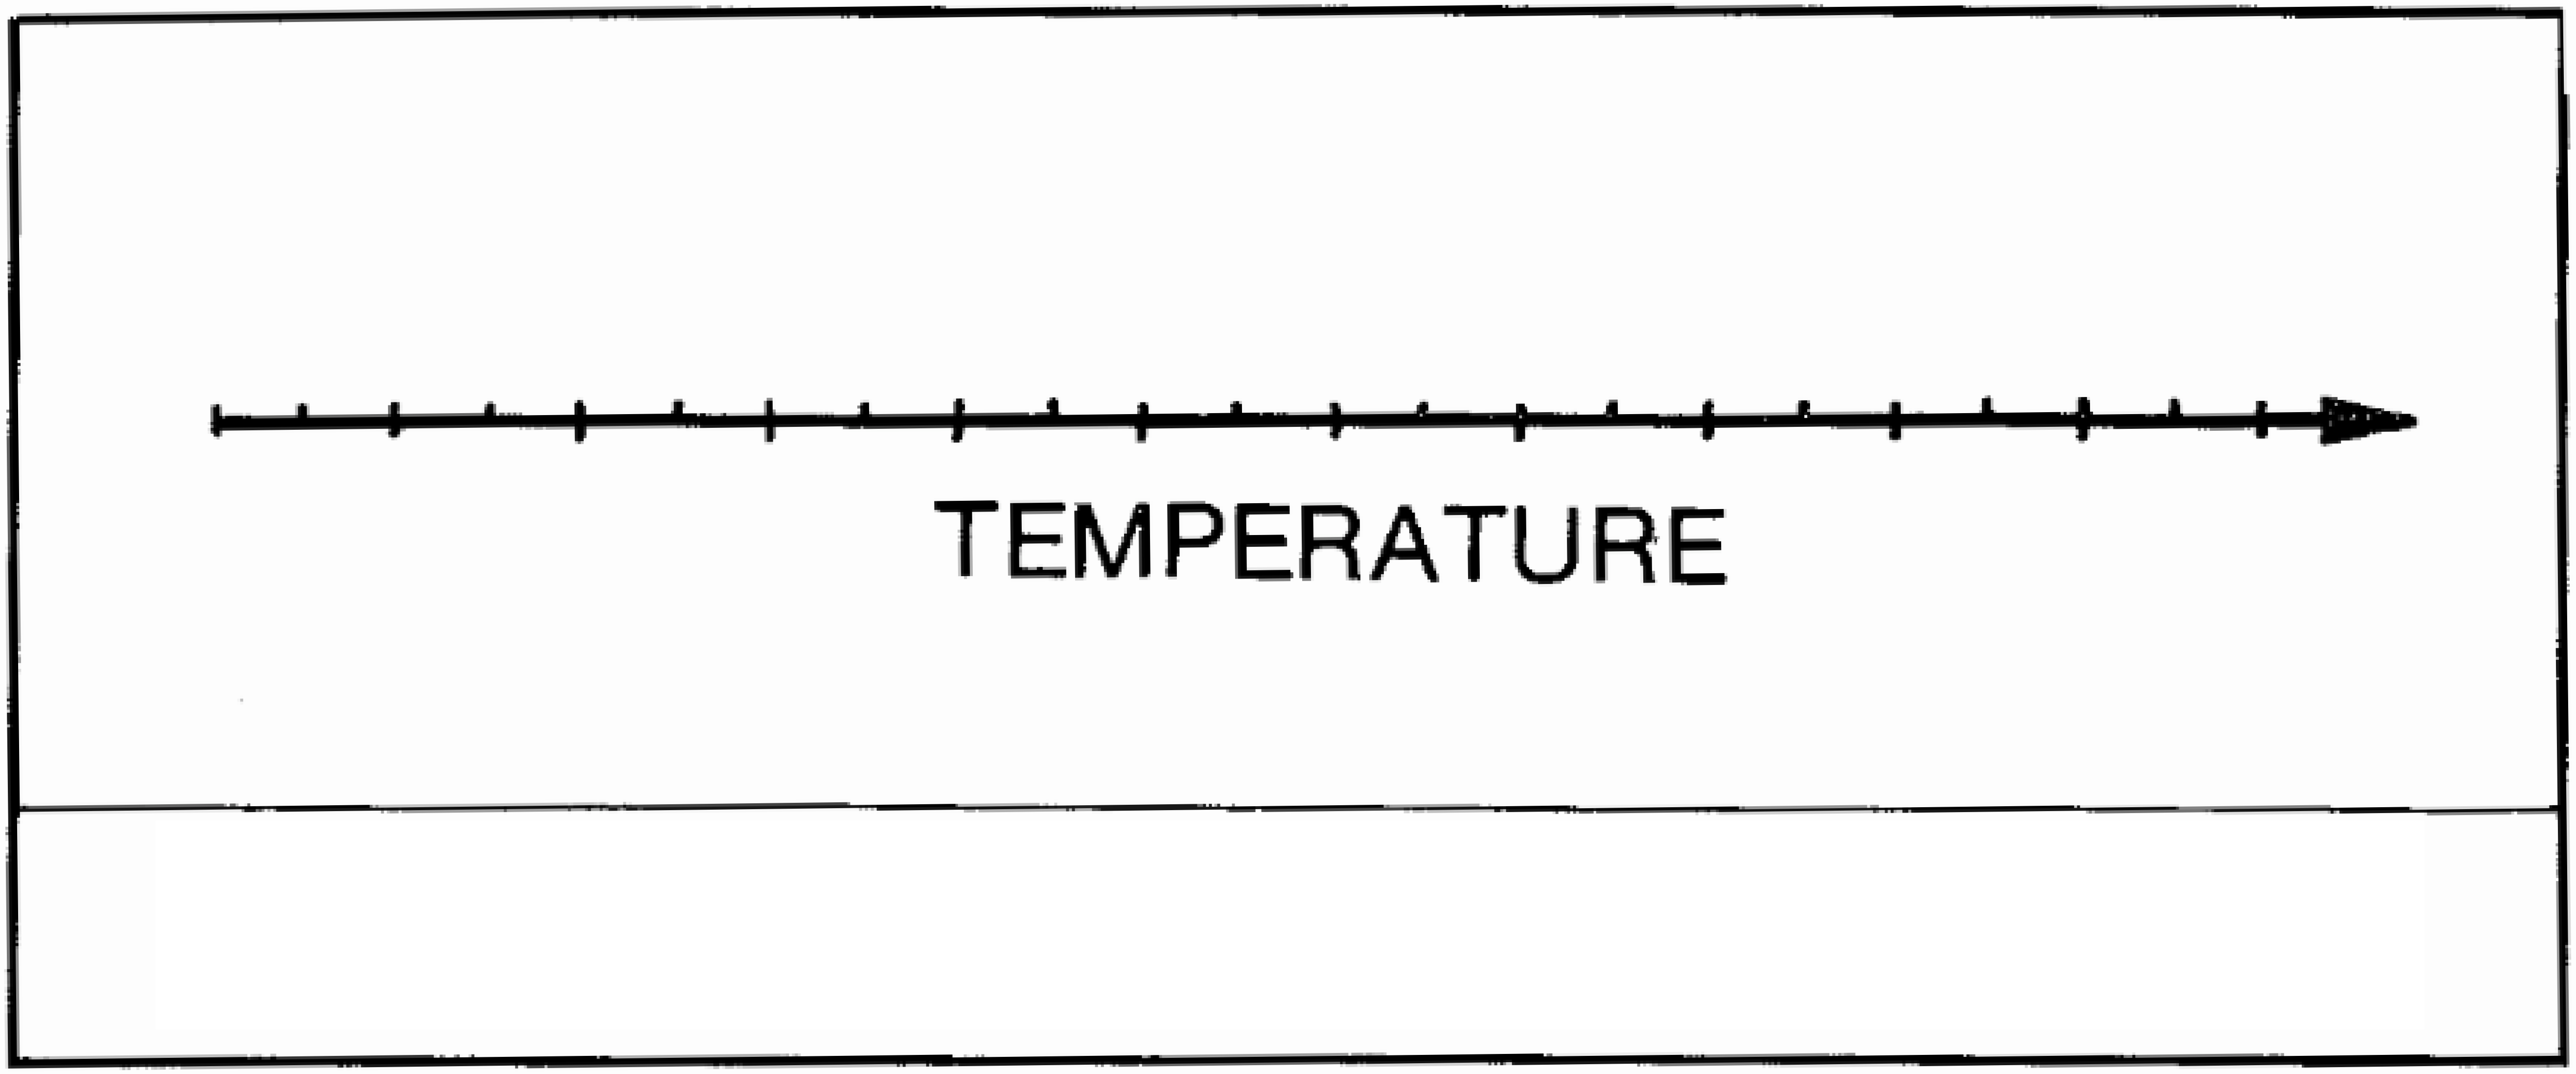


For both these examples, the *state space* is the real number line.

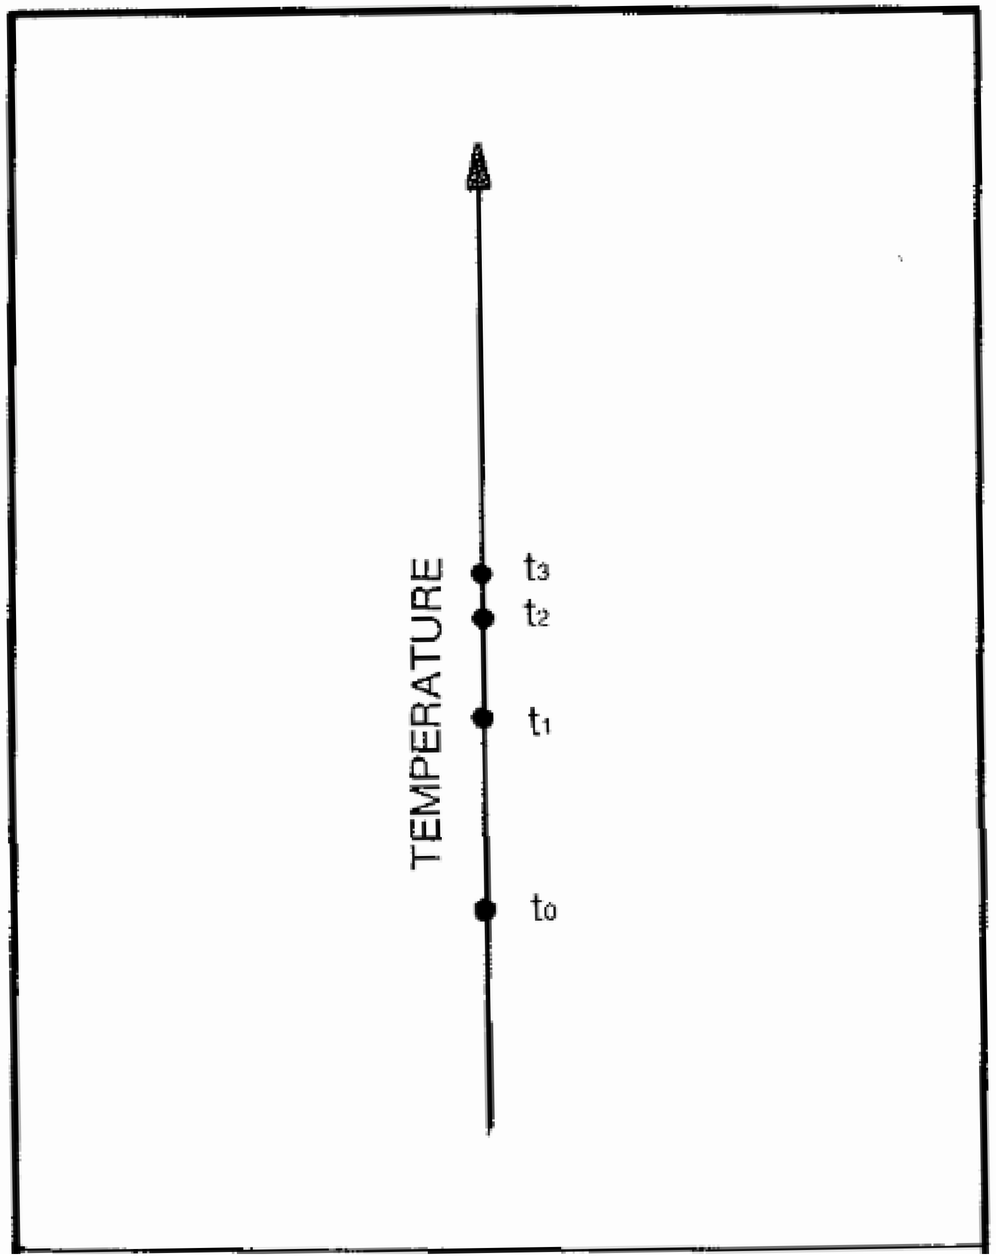


A series of observations can be labeled with their respective times

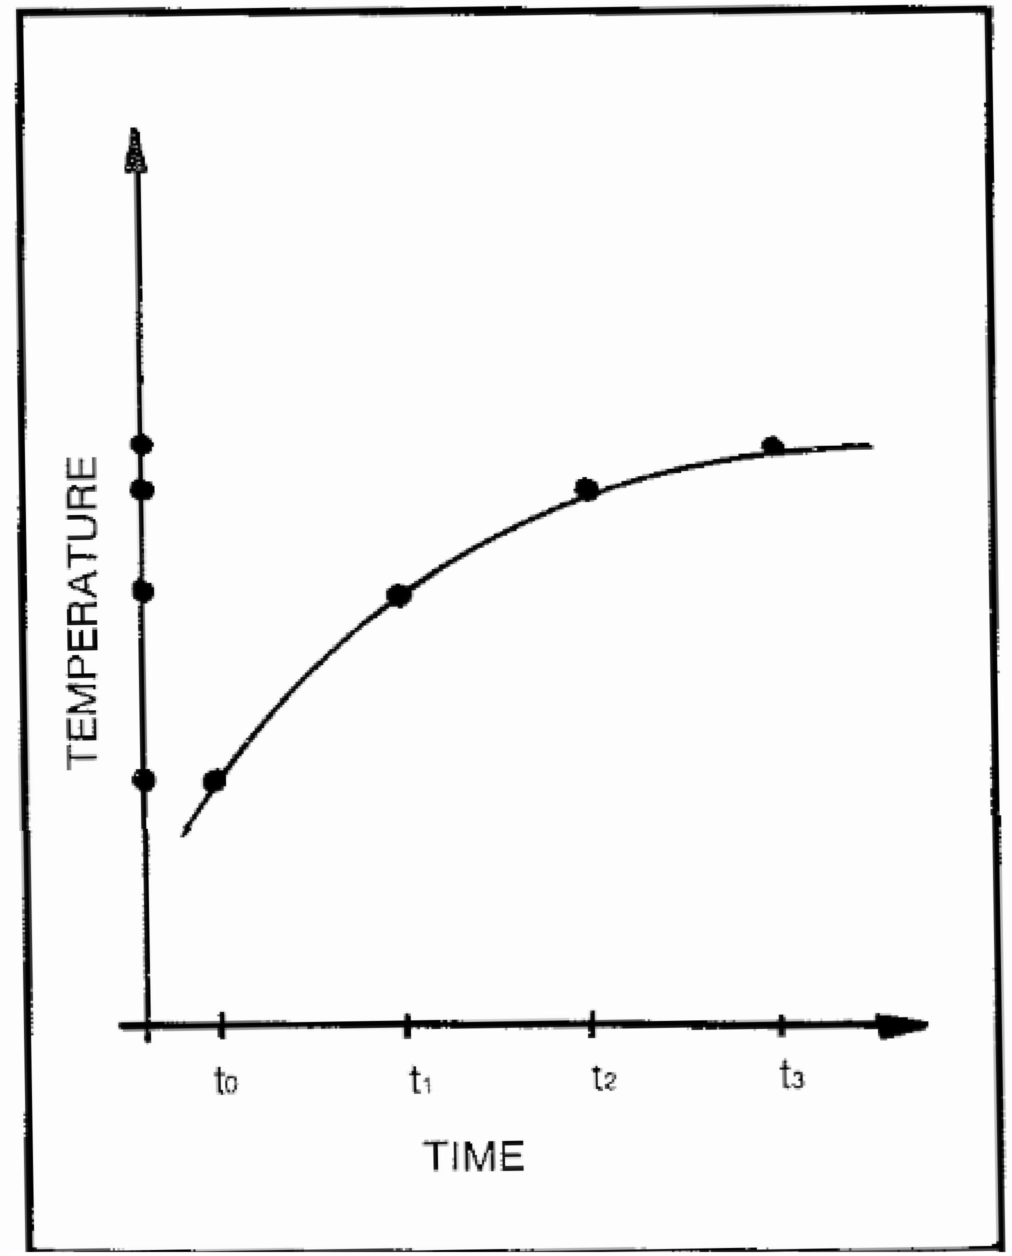


If time is shown explicitly, it adds an extra dimension

*Ear attitude* and *fang exposure* might be used to model the emotional state of a dog. (Konrad Lorenz and Christopher Zeeman)

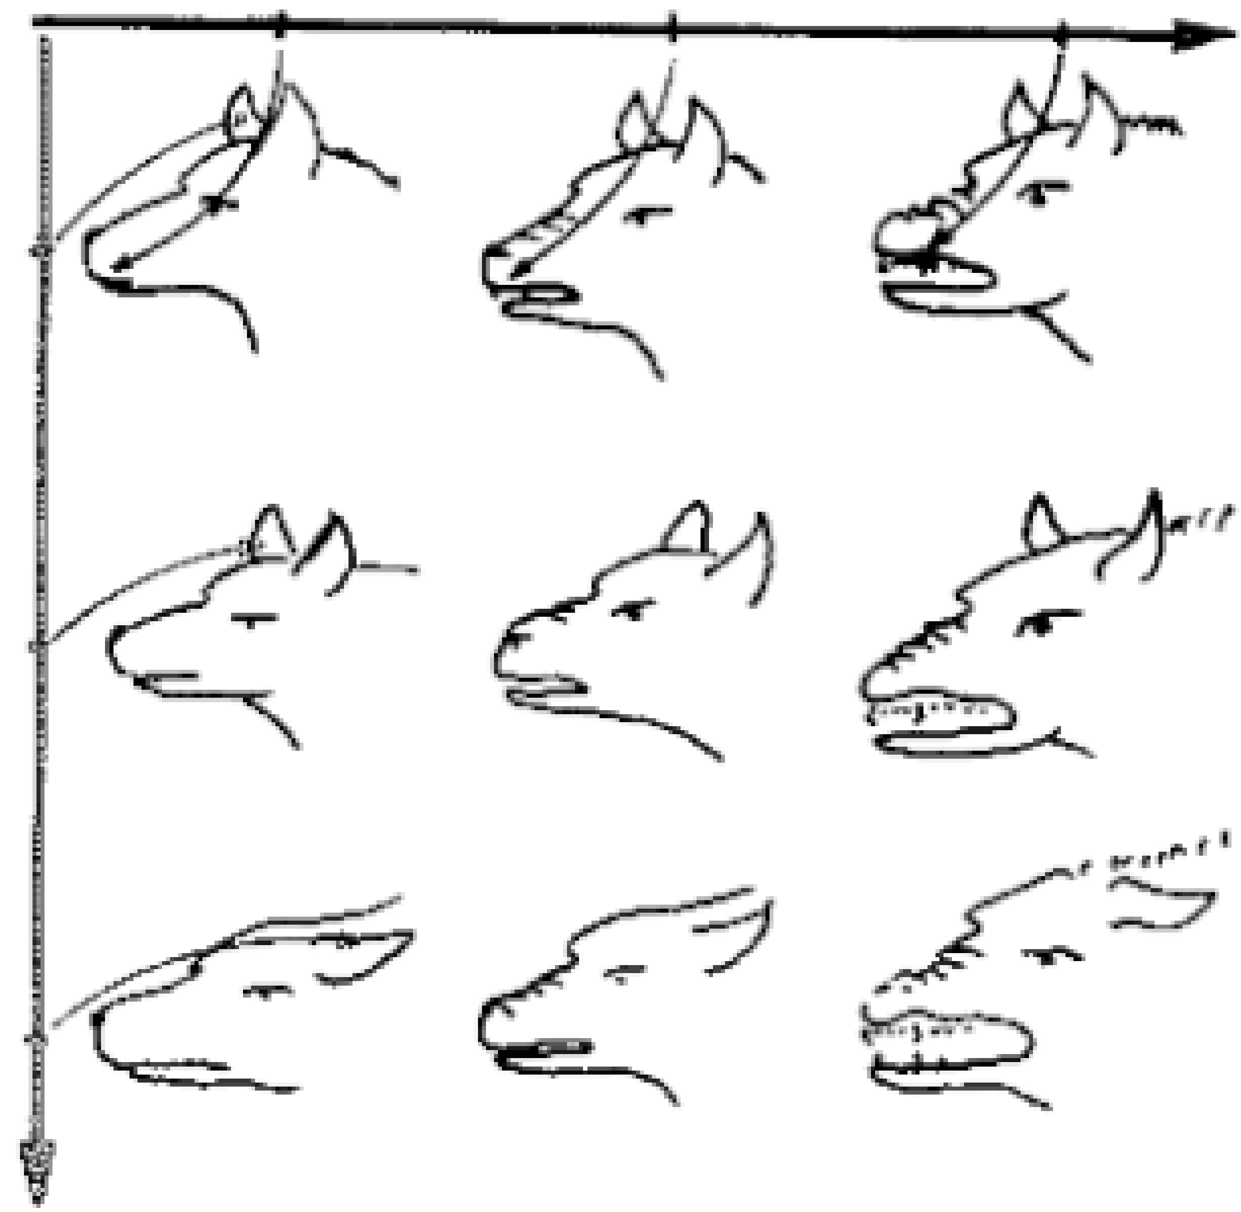

Numerical observations allow simple black box modelling. Here, we measure current and voltage.

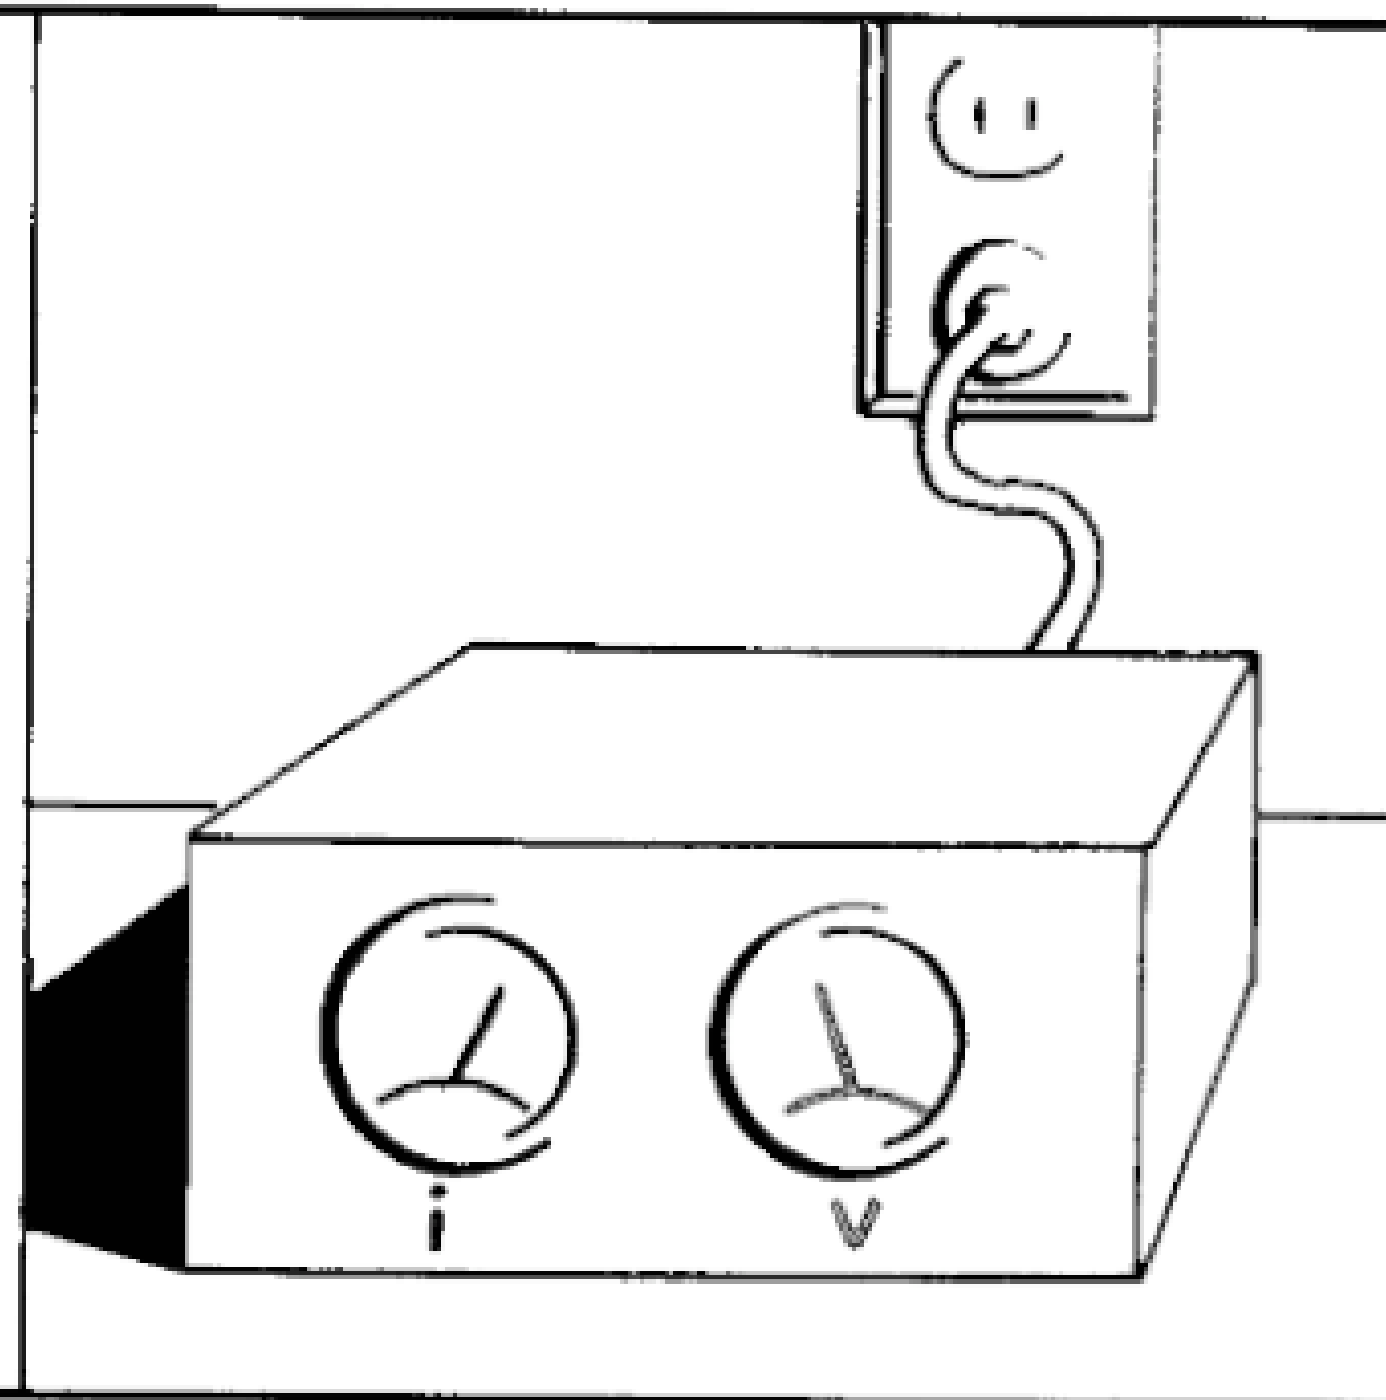

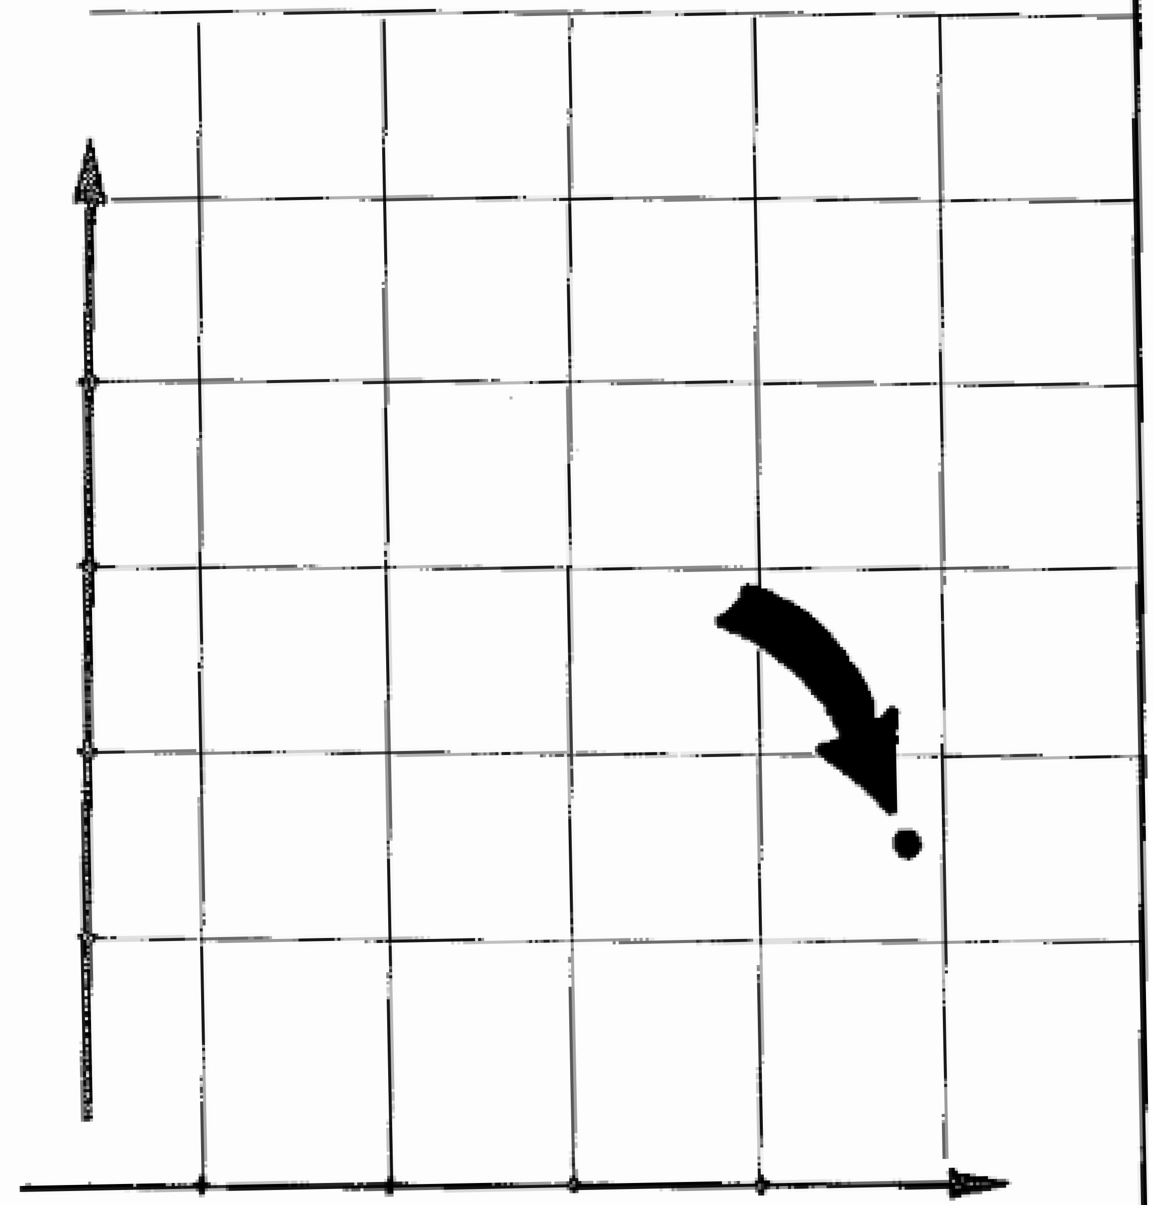

A single observation of the two values yields a single point in the state space.

Each point is associated with a specific observation. Change in the system over time is represented as a *trajectory* in state space

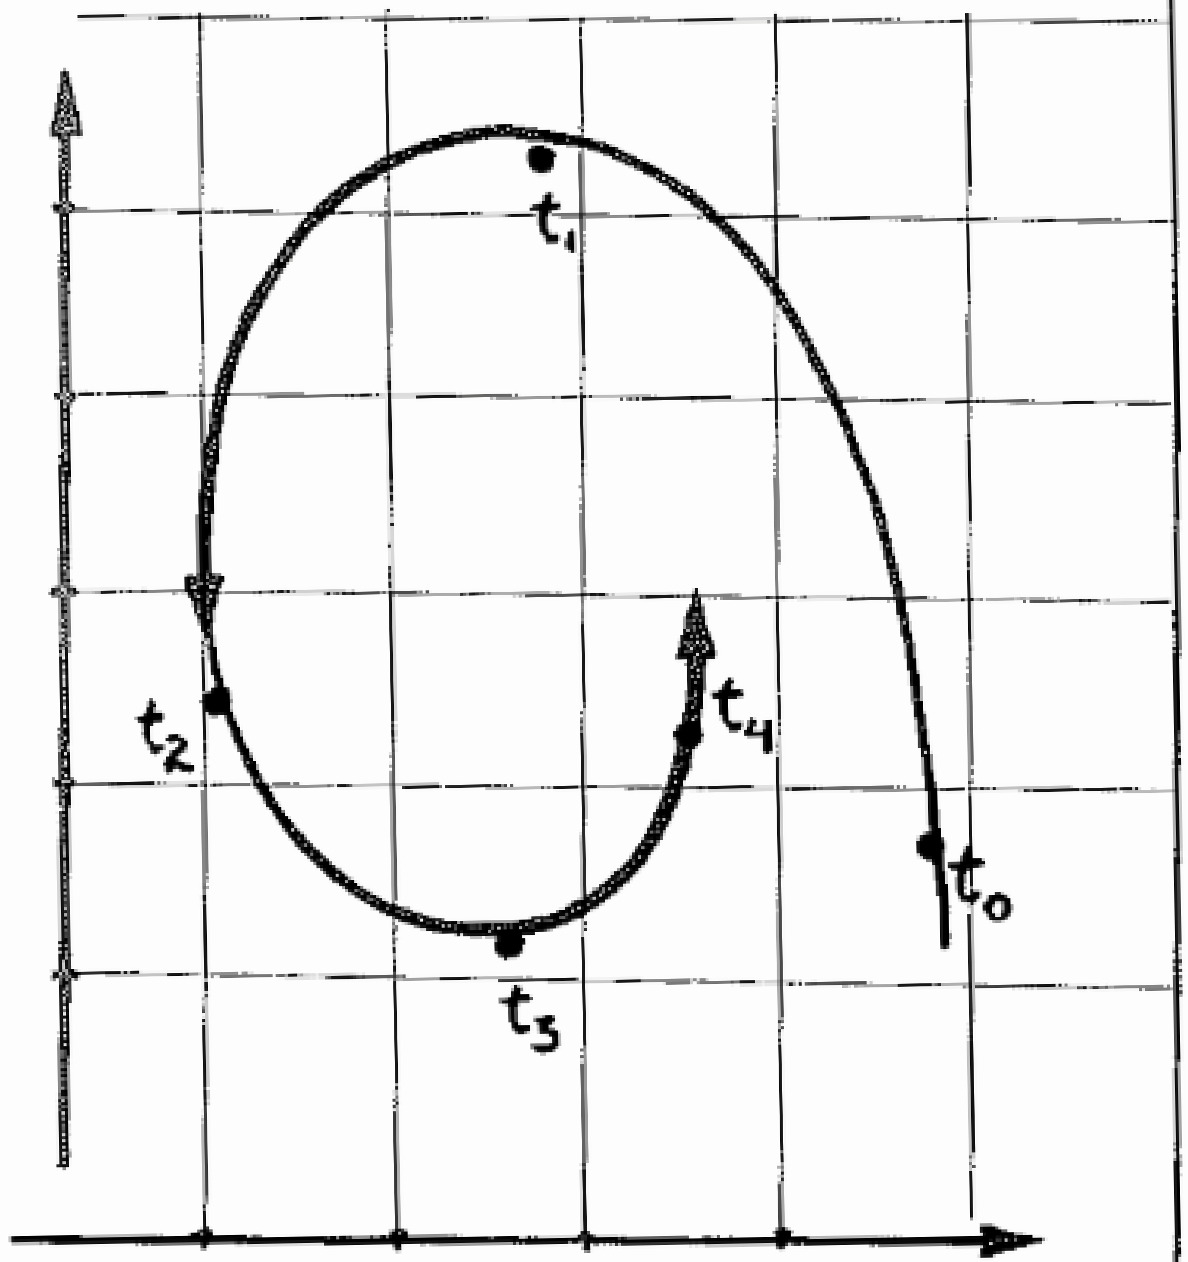


A sequence of observations of 2 variables defines a specific trajectory

Note: time is not shown explicitly

While the state space does not have time as a dimension, we can add it to obtain a *time series*.

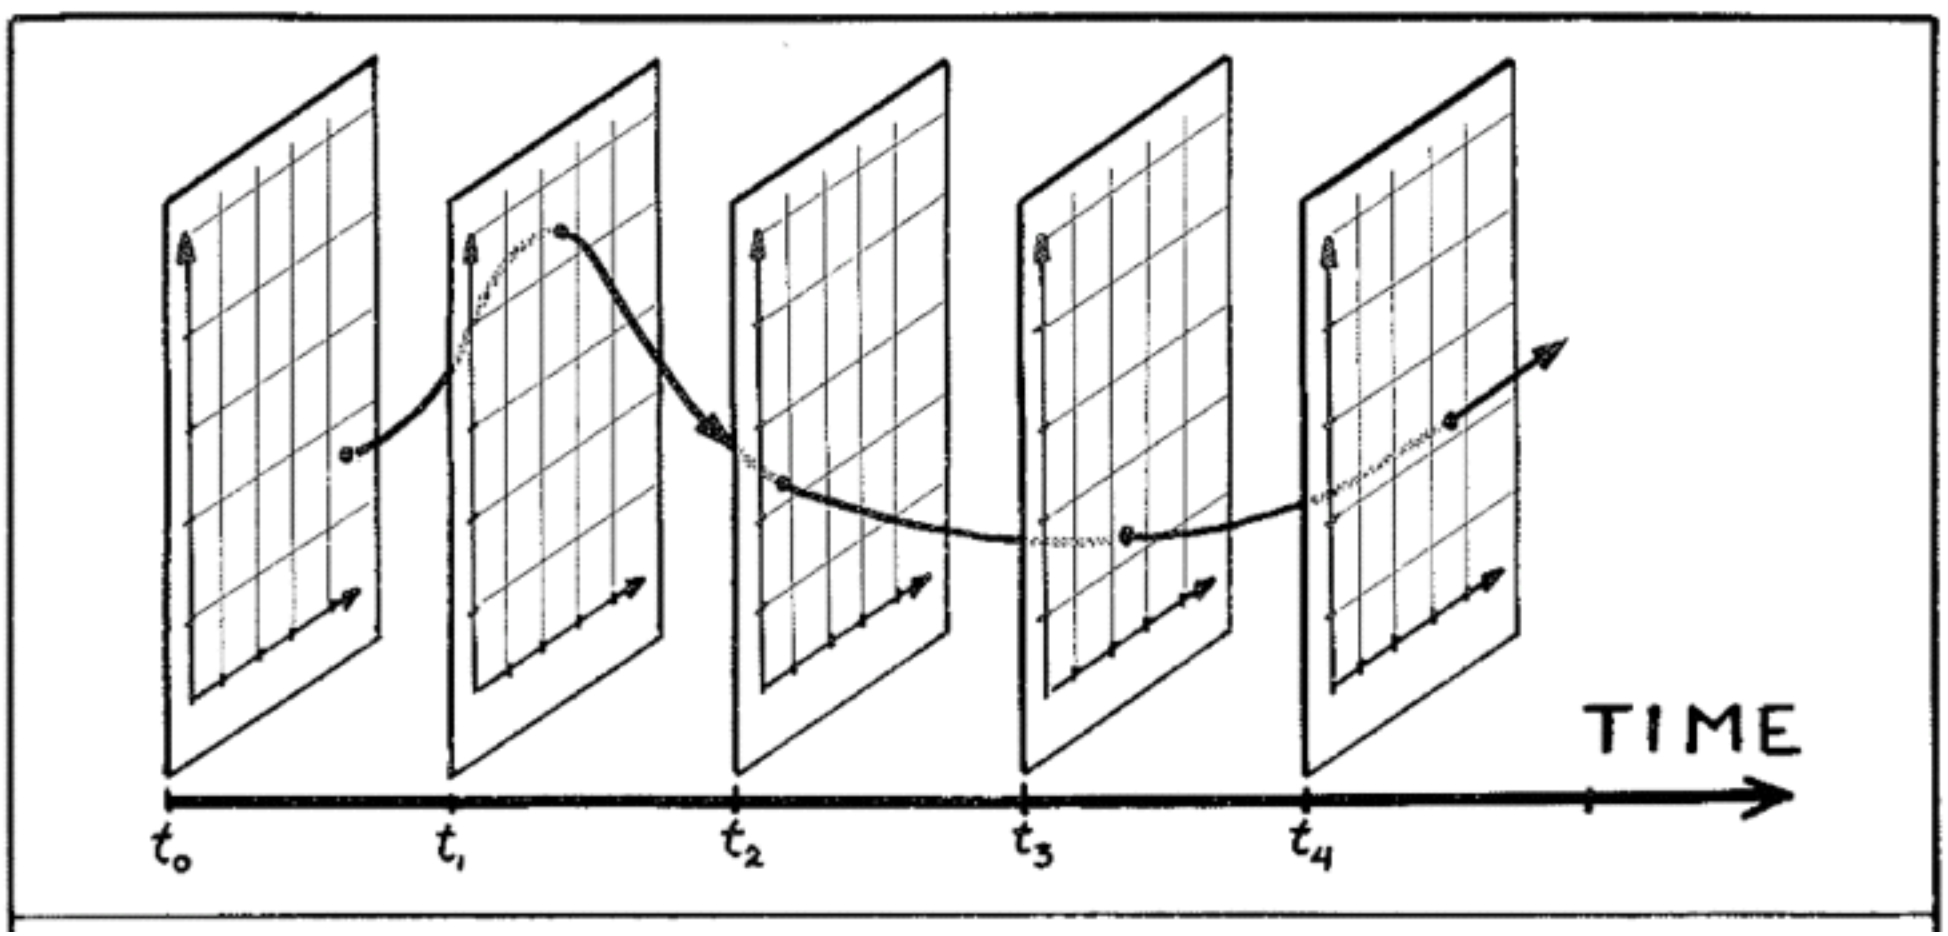

The *time series* and the *state space* representations are closely related.

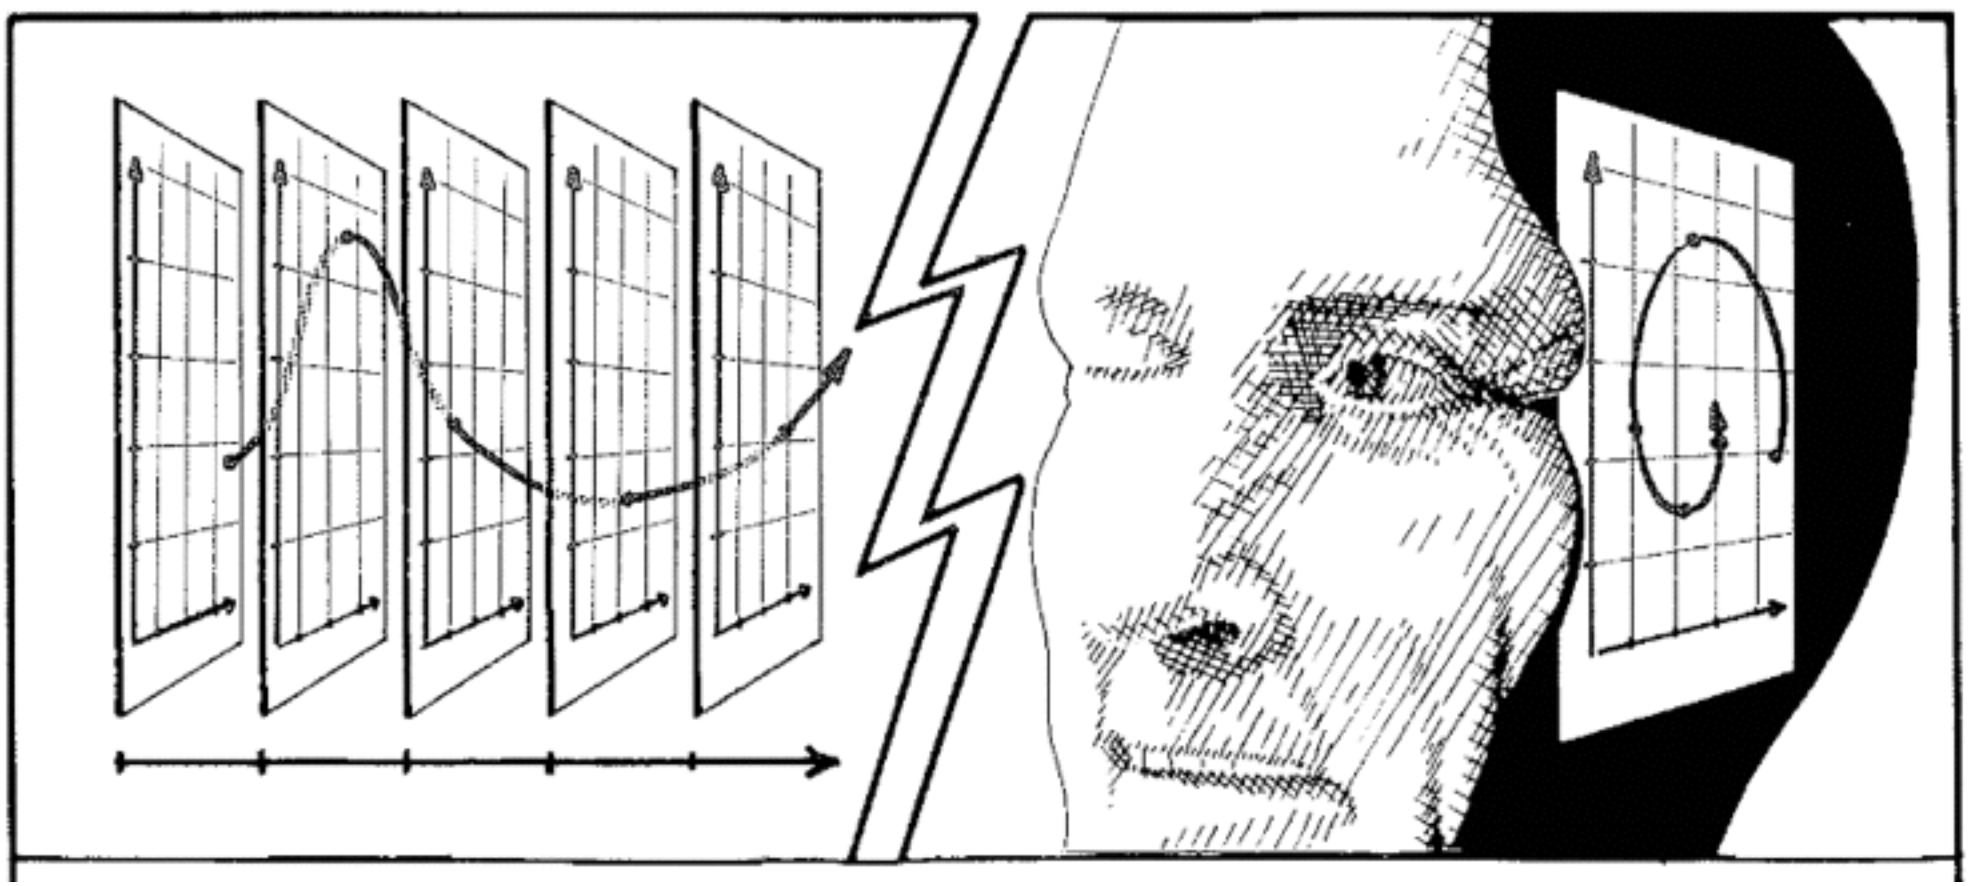

**Observing more parameters leads to models of higher dimensions.**

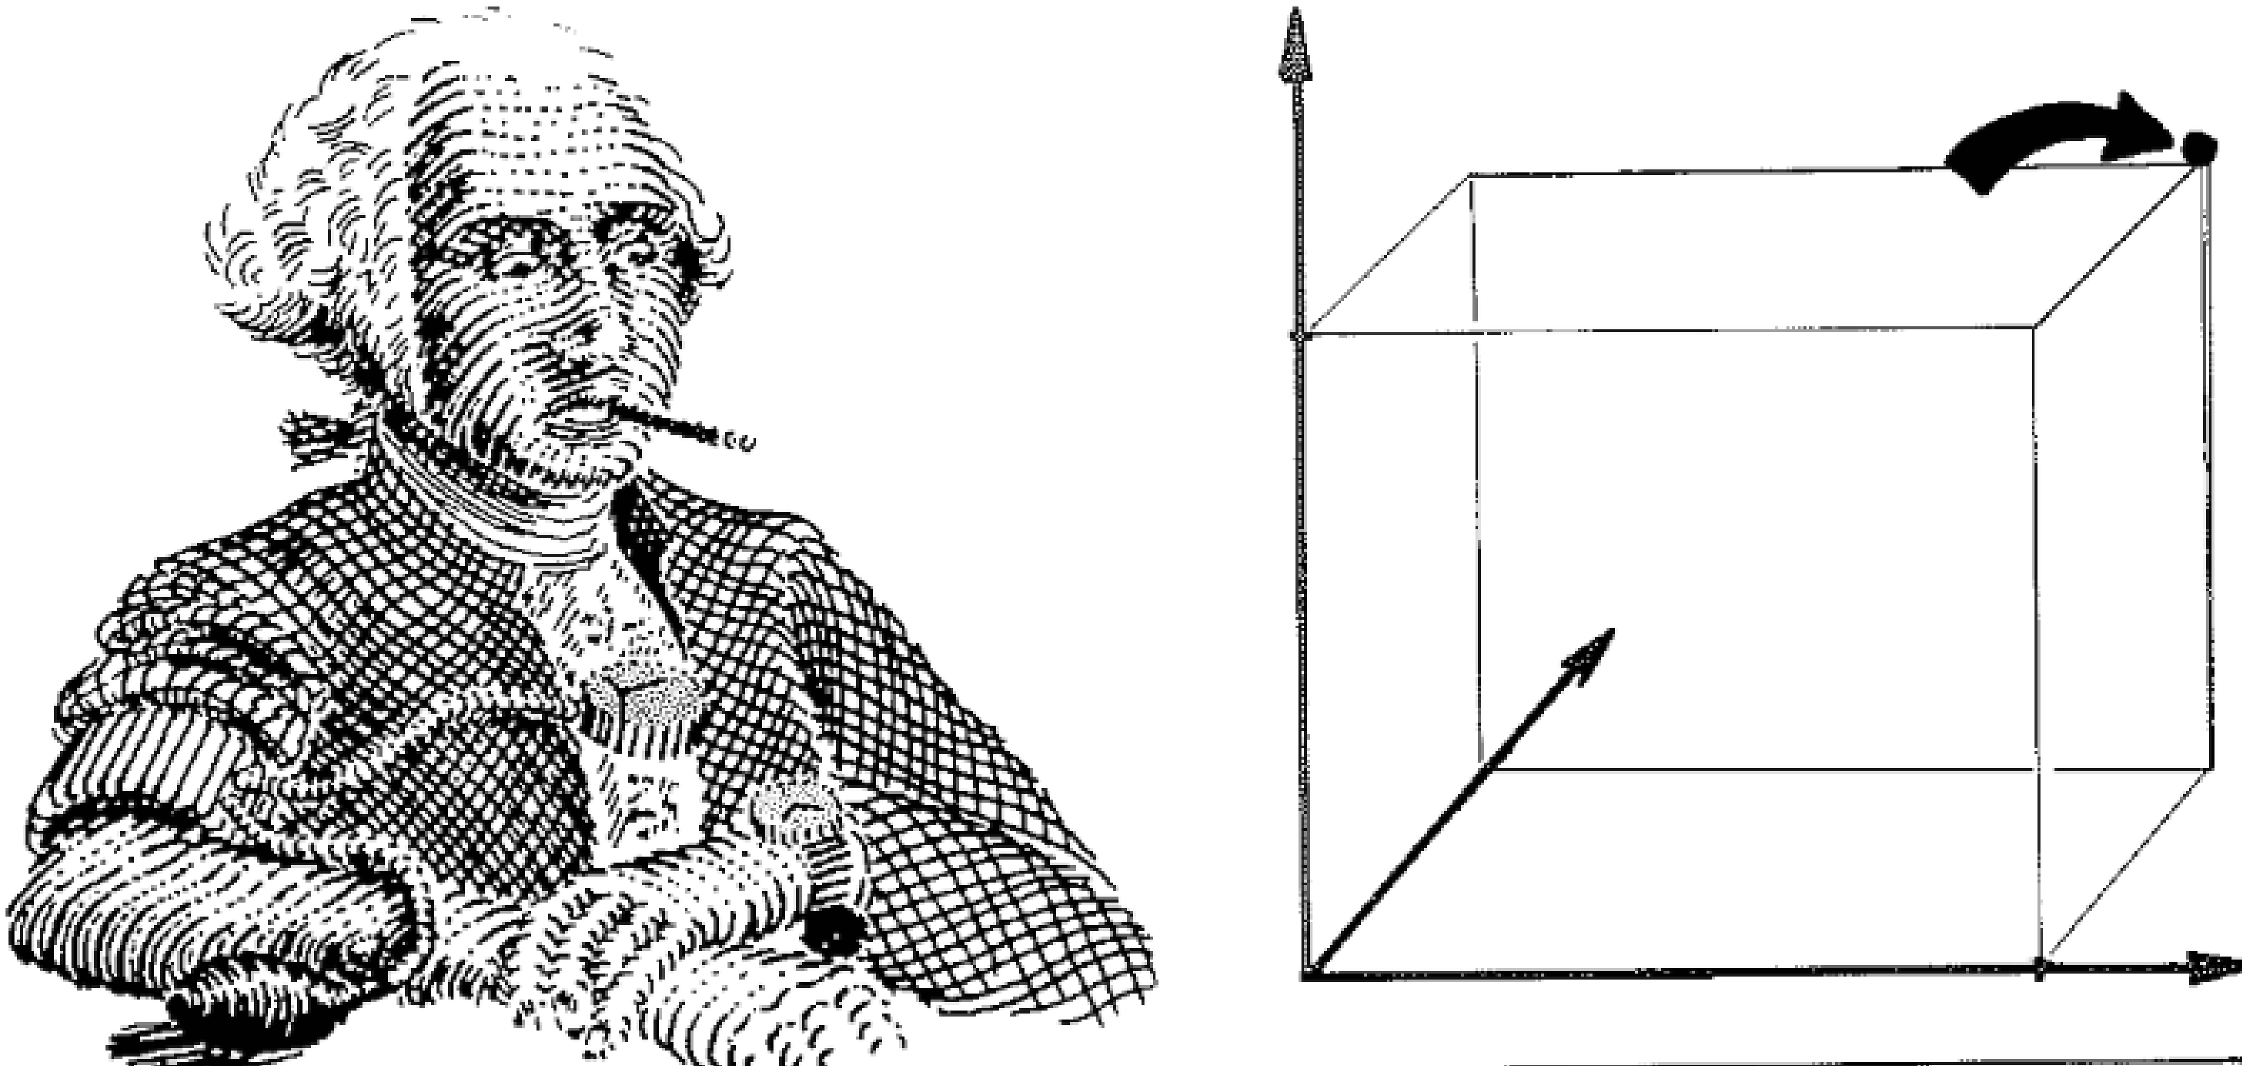

**1.1.12.** Suppose that at 7 am, this athlete observes three of his body parameters (say temperature, blood pressure, and pulse rate), records these three data as a point in three-dimensional state space.

**Many phenomena require geometric models that are not simply coordinate spaces. In dynamical systems theory, the geometric models used are *manifolds*.**

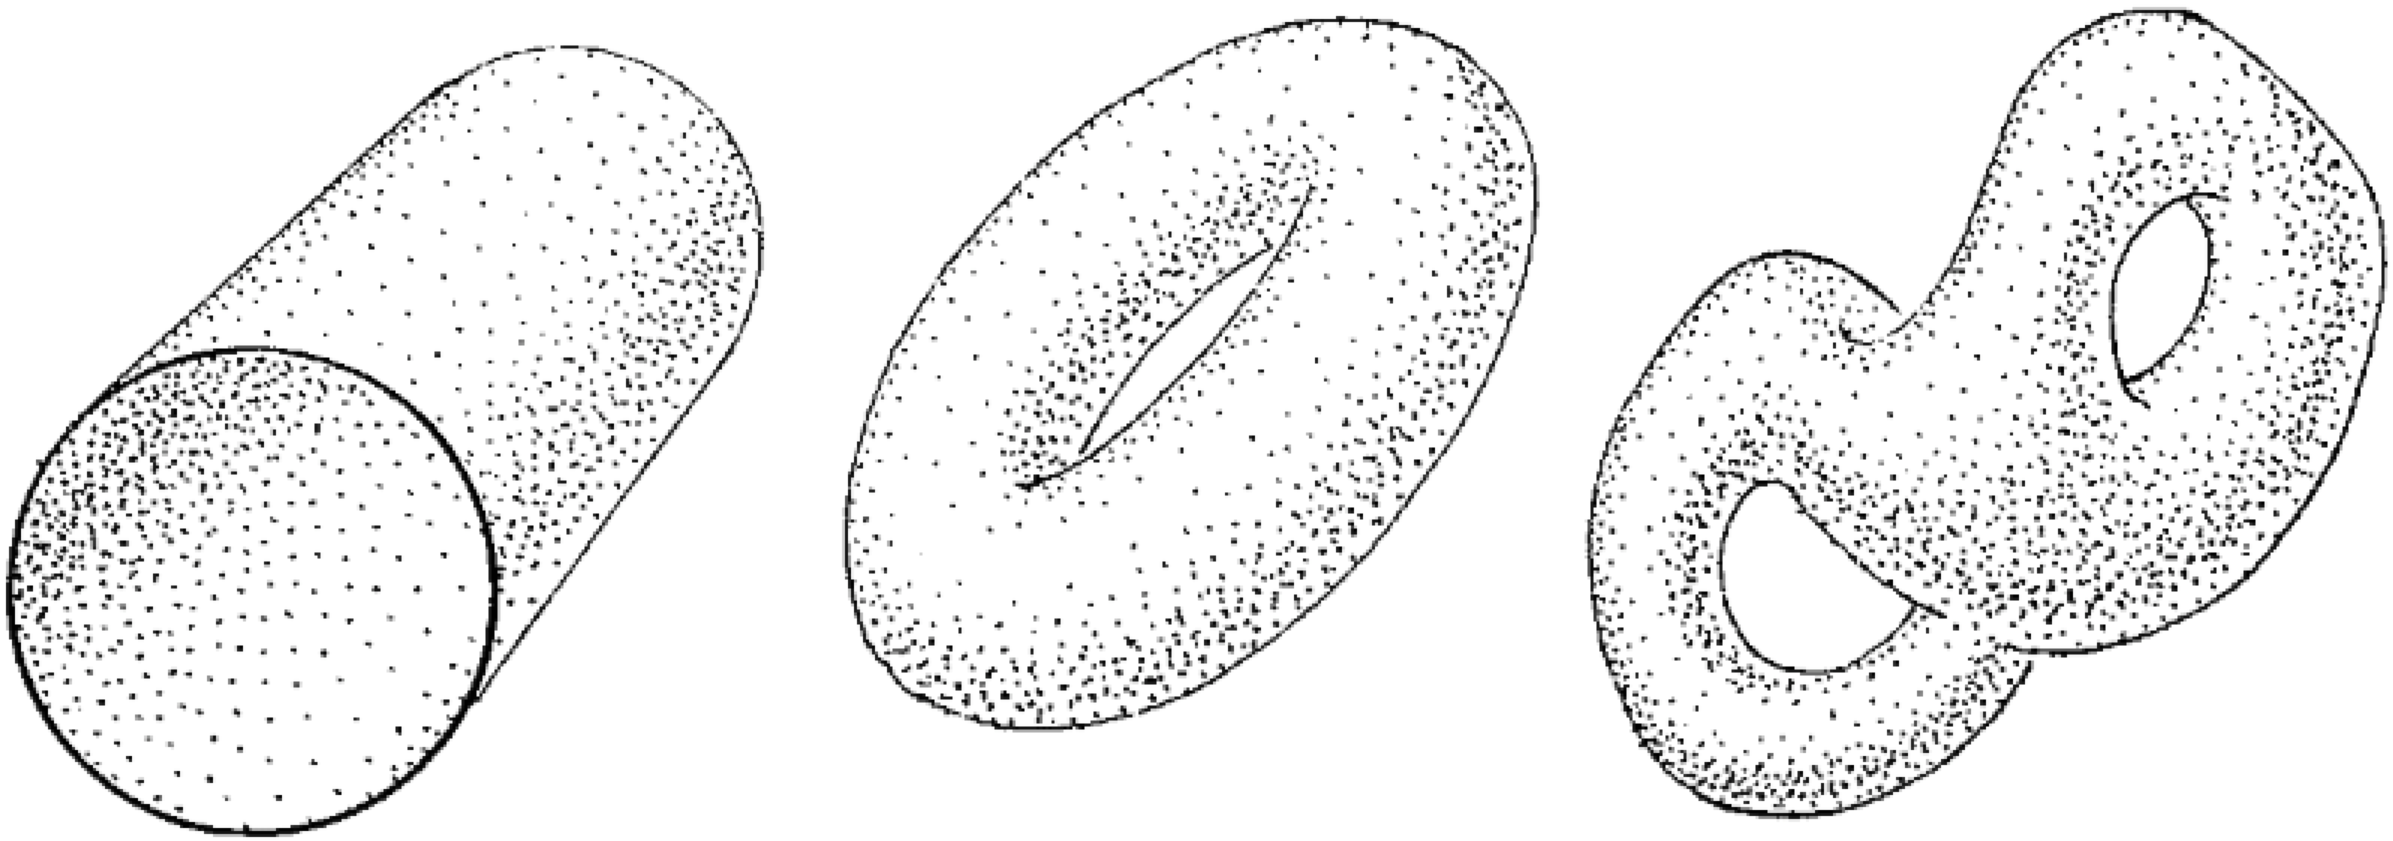

**1.1.13.** Here are some examples of manifolds. Other examples will arise in later chapters. They are made of pieces of flat spaces, bent and glued together.

At this point, the history of real system has been represented graphically, as a *trajectory* in a geometric state space. An alternative representation is the *time series* of the trajectory.

The time series makes the passage of time explicit.

The state space representation does not make the passage of time explicit.

When we use a state space representation, we are trying to characterise the dynamic system in an abstract fashion, not tied to any one set of observables. For example, we wish to describe “a pendulum”, and not just “this specific pendulum, right here”.

An abstract characterisation of a dynamical system can be understood as a vector field.

2D

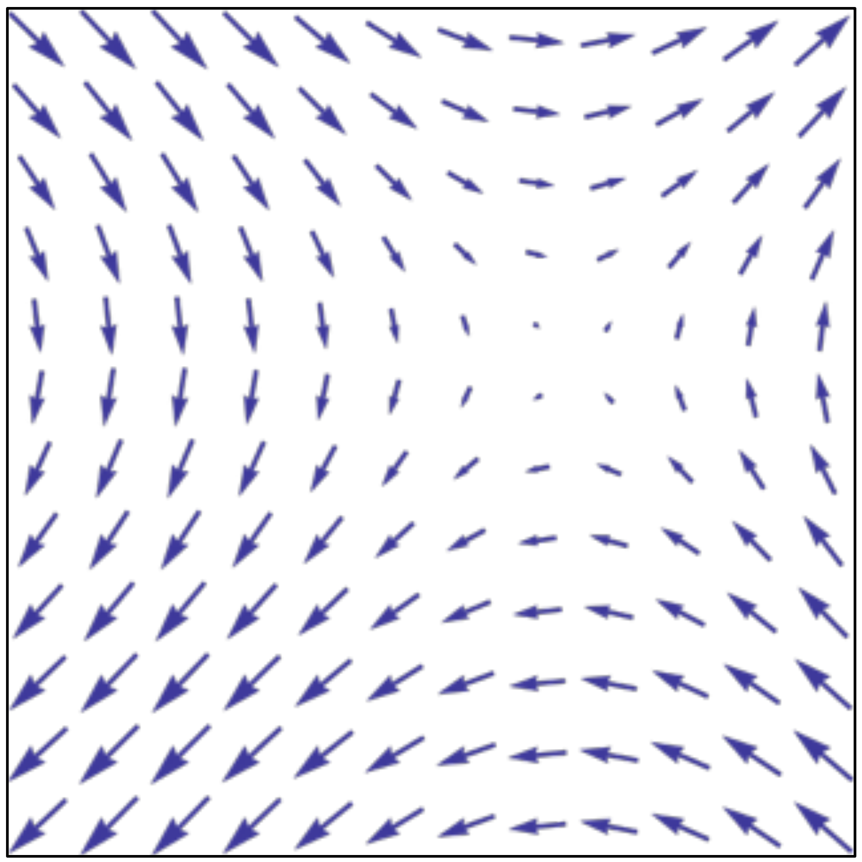

3D

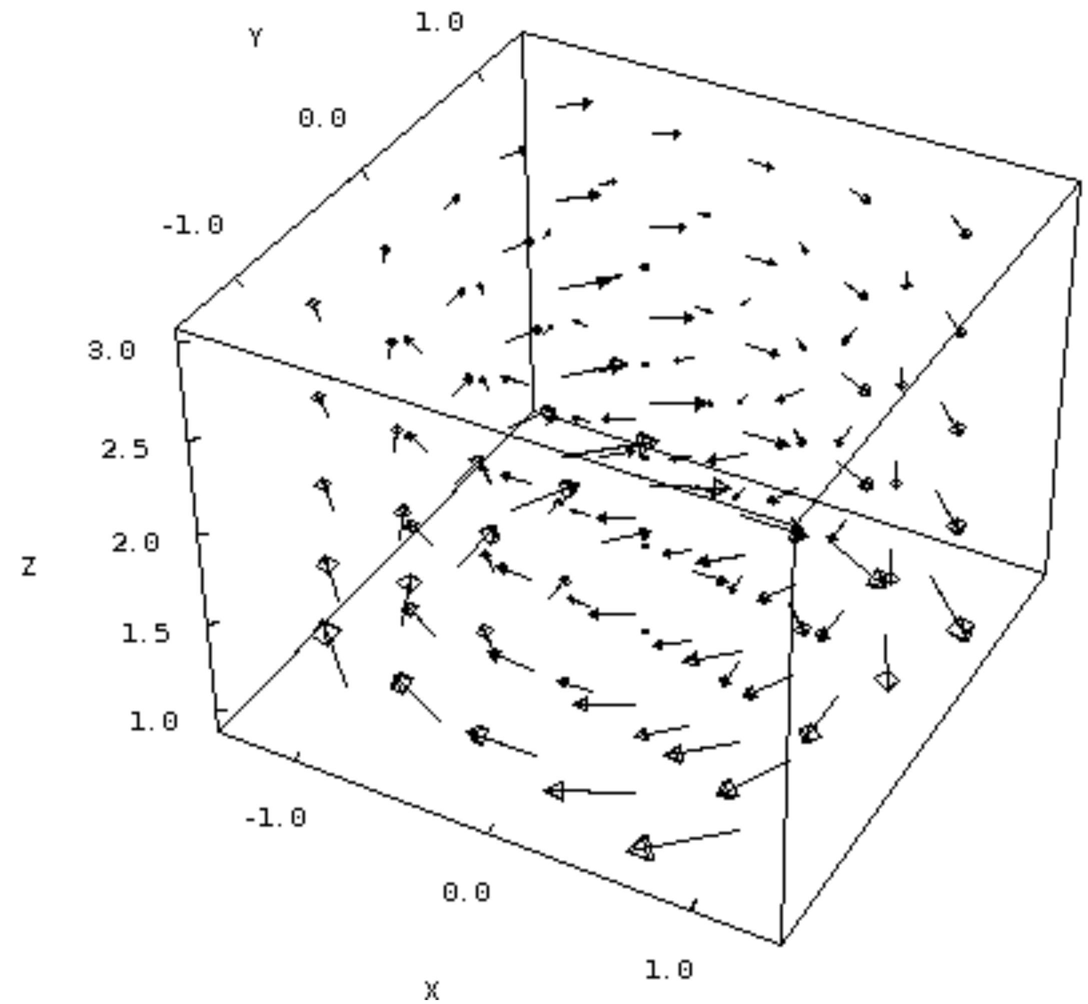


This demands that we be able to make sense of an *instantaneous rate of change* for each possible state.

This is the innovation of the differential calculus.

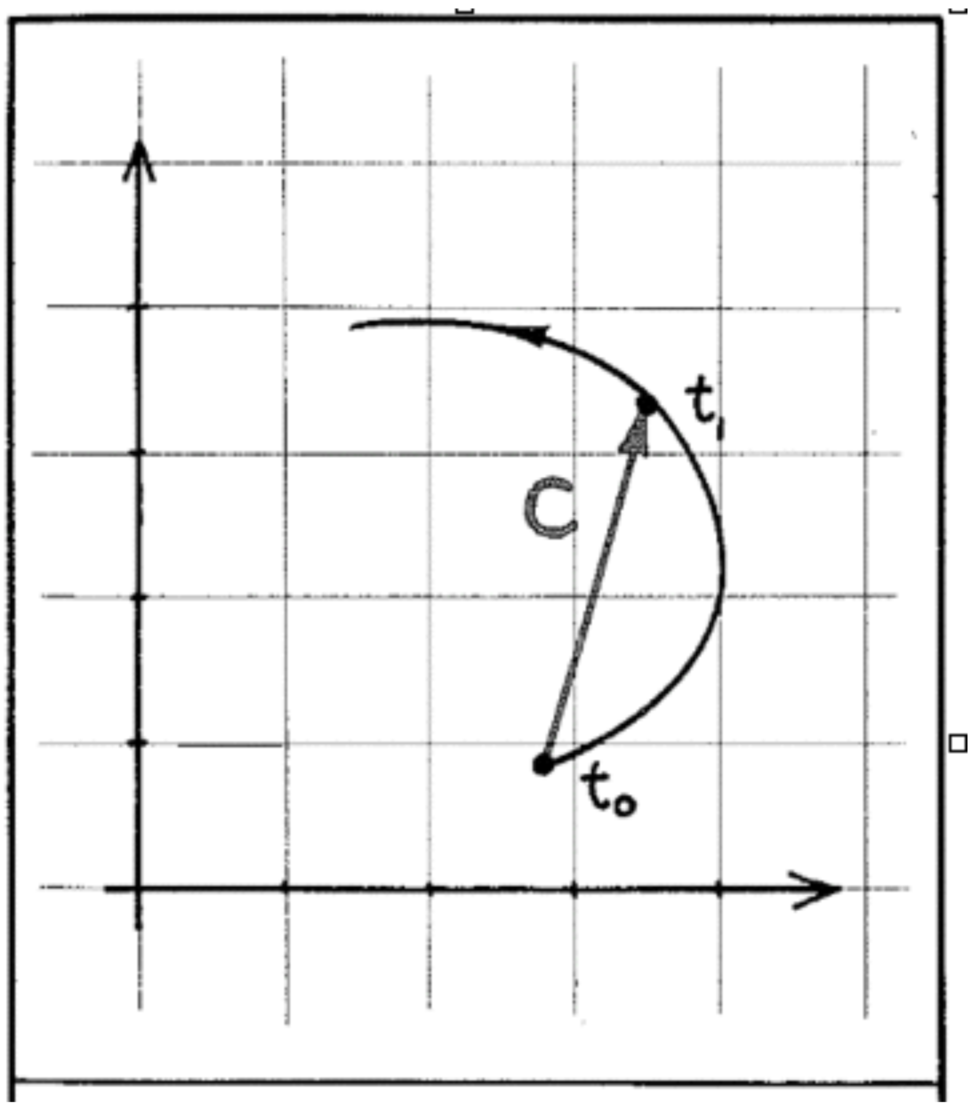


Let us say a system shows a change, *C*, between two points in time, $t_0$ and $t_1$

Looking at *C* alone, we do not know if this is a big change or a small change -- it depends on the amount of time that passed!

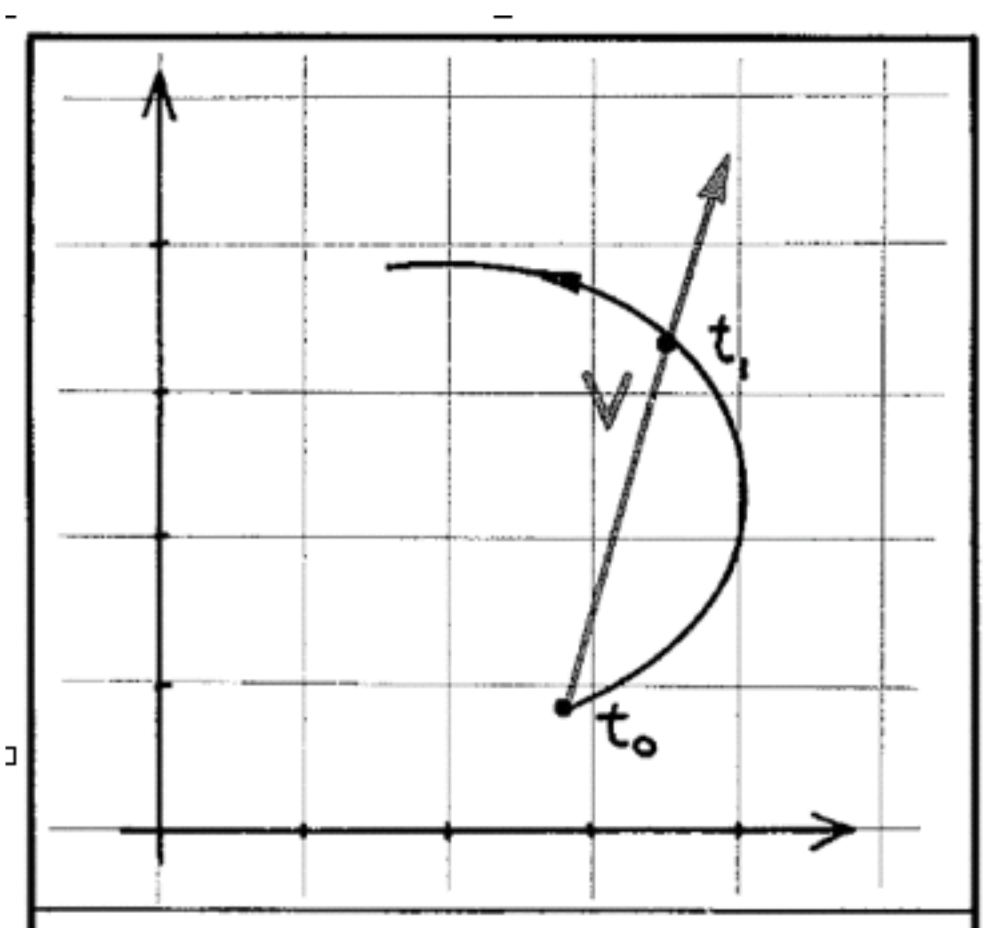


Now we scale *C* by the amount of time that passed:

$V = C/T$

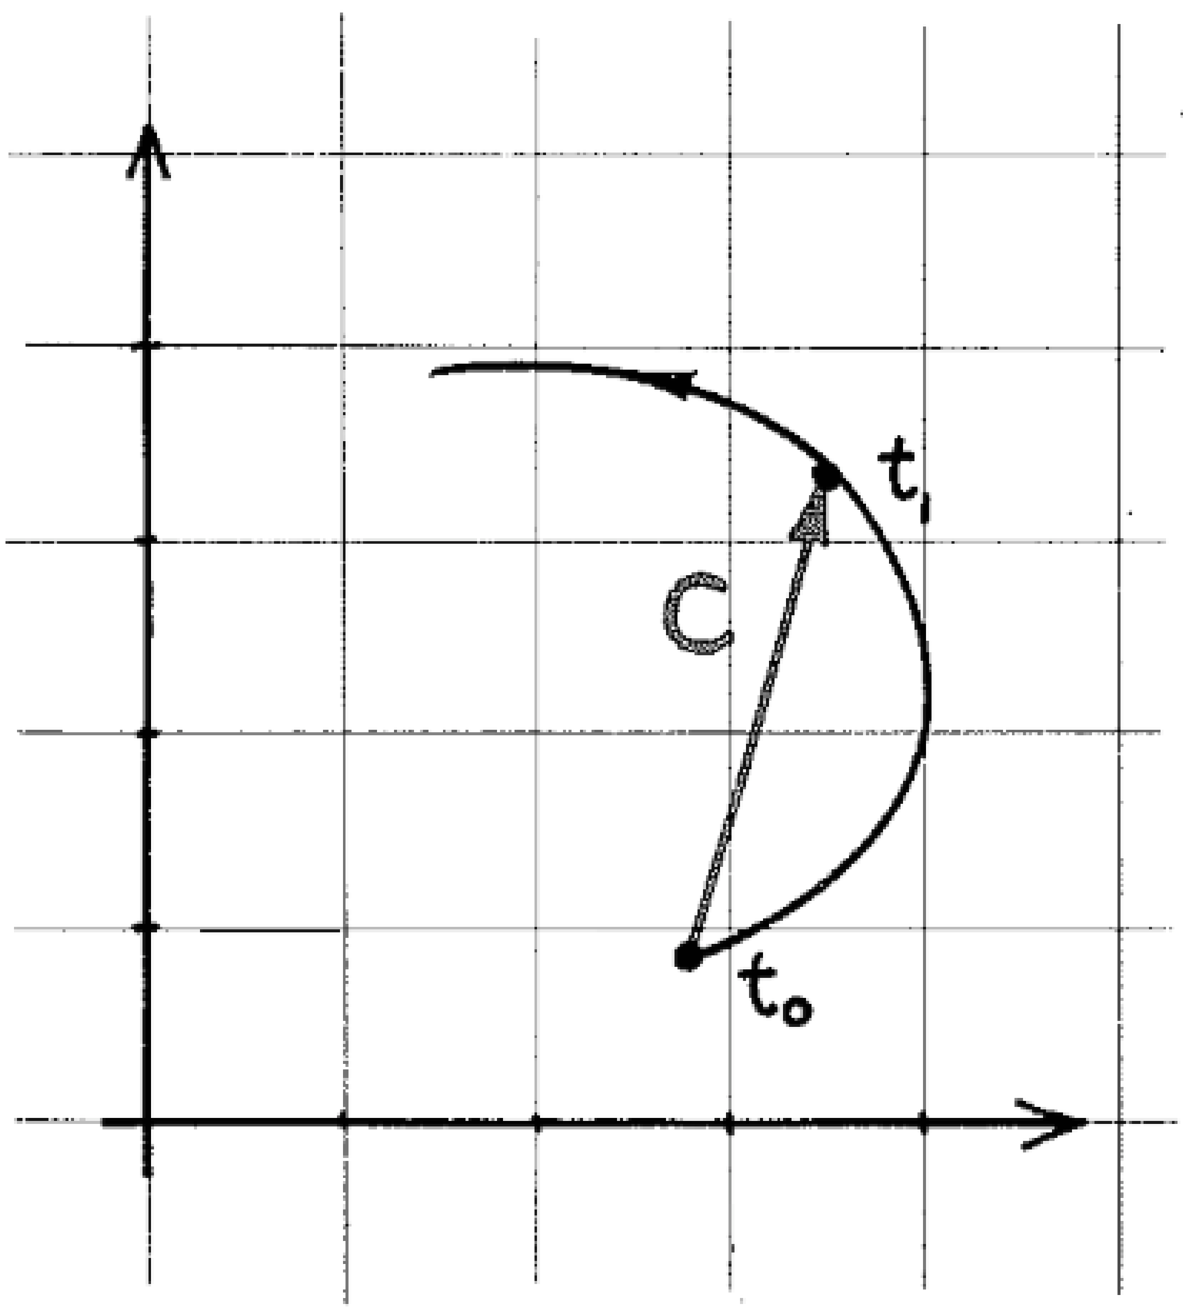


**1.2.1.** On this trajectory, the states obseved at two different times, $t_0$ and $t_1$, are connected by a *bound vector*, represented here by a line segment pointed on one end. Let $C$ denote this bound vector.

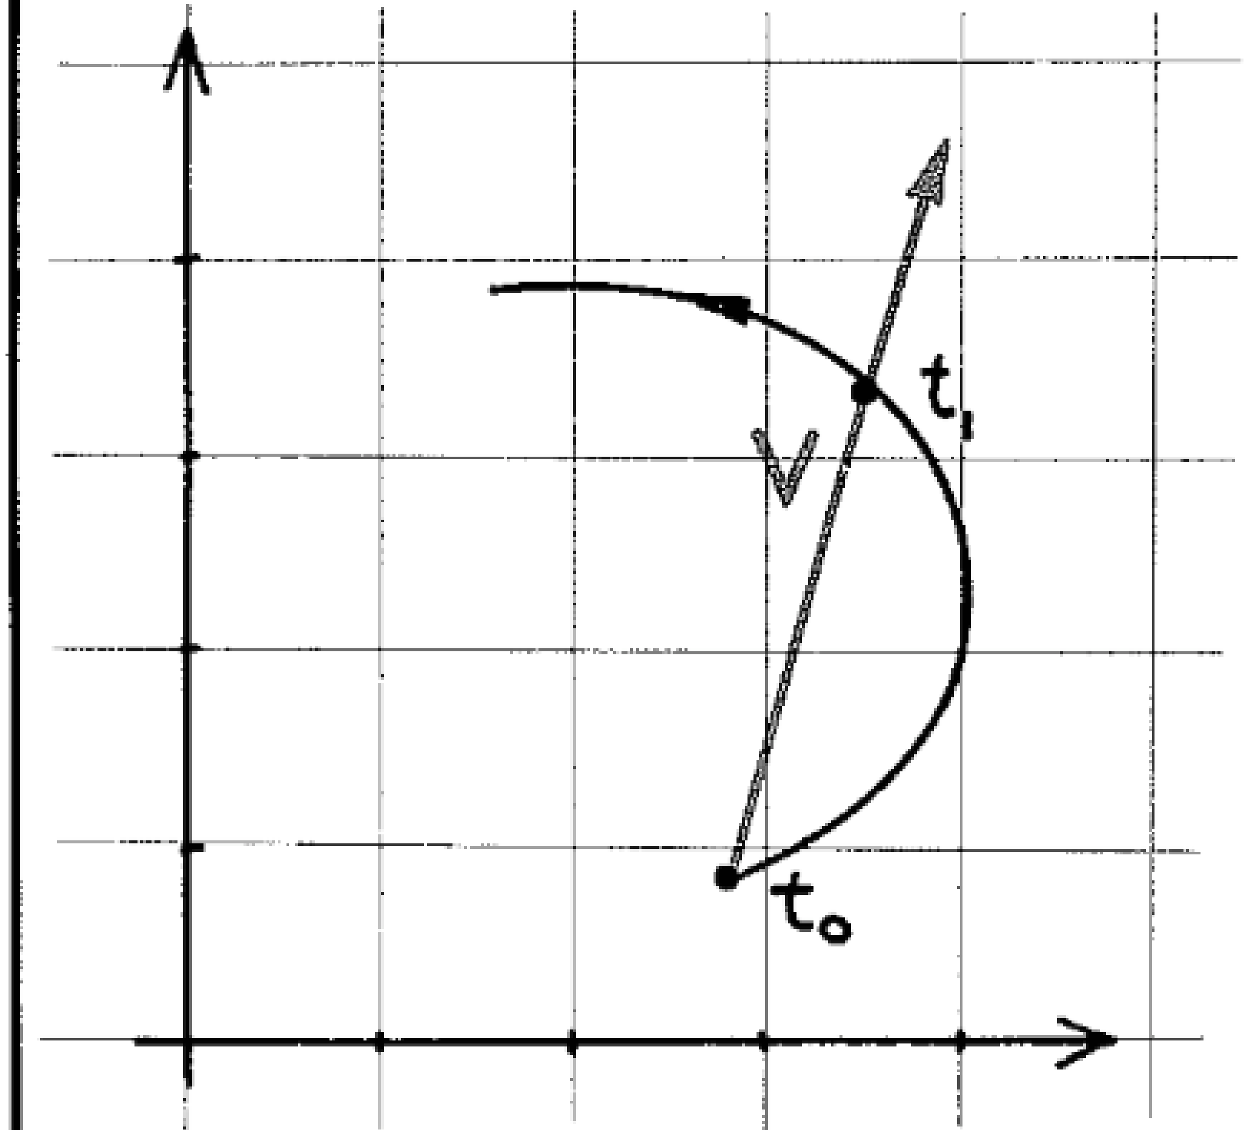


**1.2.2.** The *average velocity* of the change of state, $C$, is the vector starting at the point labeled $t_0$ on the curve, and directed along the vector of change of state, $C$, but divided by T, the time elapsed between $t_0$ and $t_1$. Let $V$ denote this vector, $V$=C/T. It represents the average speed and direction of the change of state.

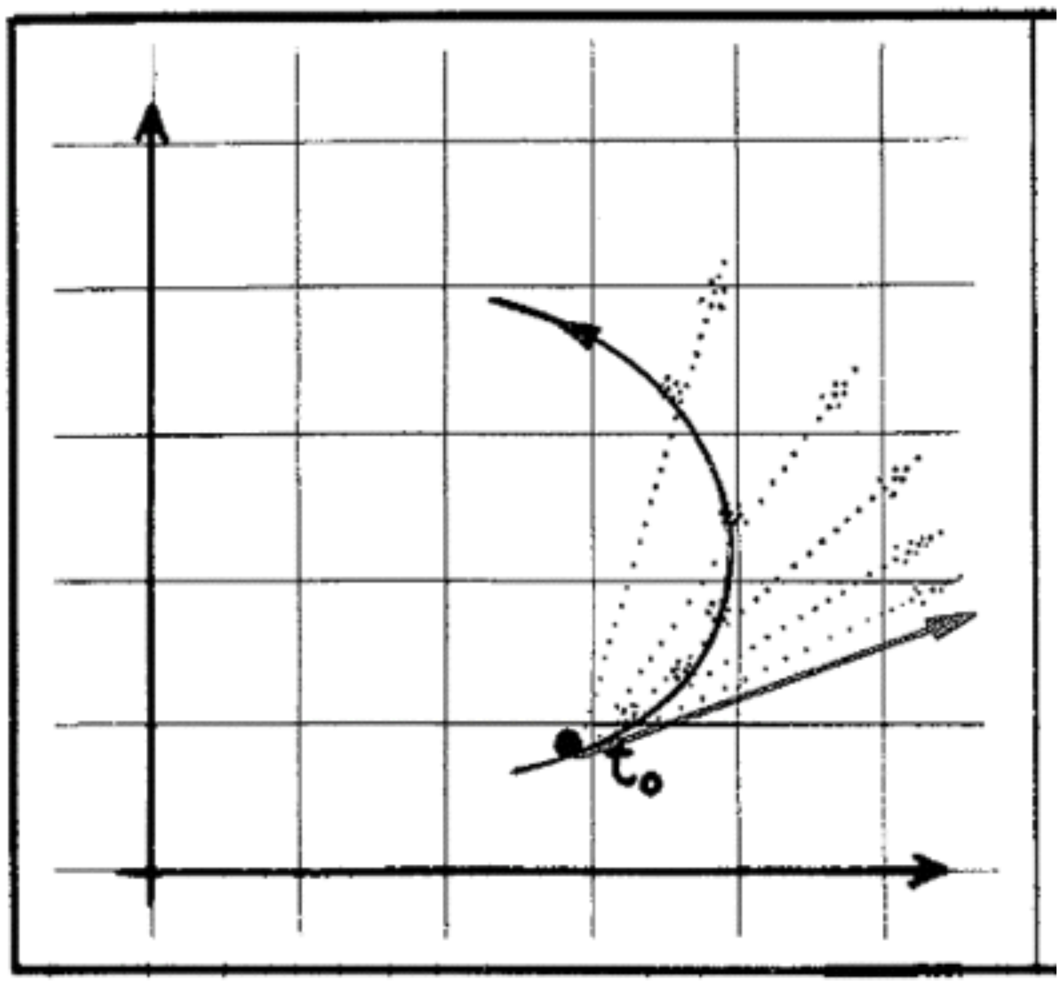


If we look at how *V* changes as we make the time interval small . . . we end up with the *instantaneous rate of change*

$$\lim_{\Delta t \to 0} \frac{x(t+\Delta t) - x(t)}{\Delta t} = \frac{dx}{dt}$$

$$\lim_{\Delta t \to 0} \frac{x(t + \Delta t) - x(t)}{\Delta t} = \frac{dx}{dt}$$

We are usually concerned with change (say, in *x)* w.r.t. *time (t)*, but the concepts allow us to express change in one variable (*y)* w.r.t. another (*x*).

$$\lim_{\Delta x \to 0} \frac{y(x + \Delta x) - y(x)}{\Delta x} = \frac{dy}{dx}$$

$$\lim_{\Delta x \to 0} \frac{y(x + \Delta x) - y(x)}{\Delta x} = \frac{dy}{dx}$$

$$\lim_{\Delta x \to 0} \frac{\Delta y}{\Delta x} = \frac{dy}{dx}$$

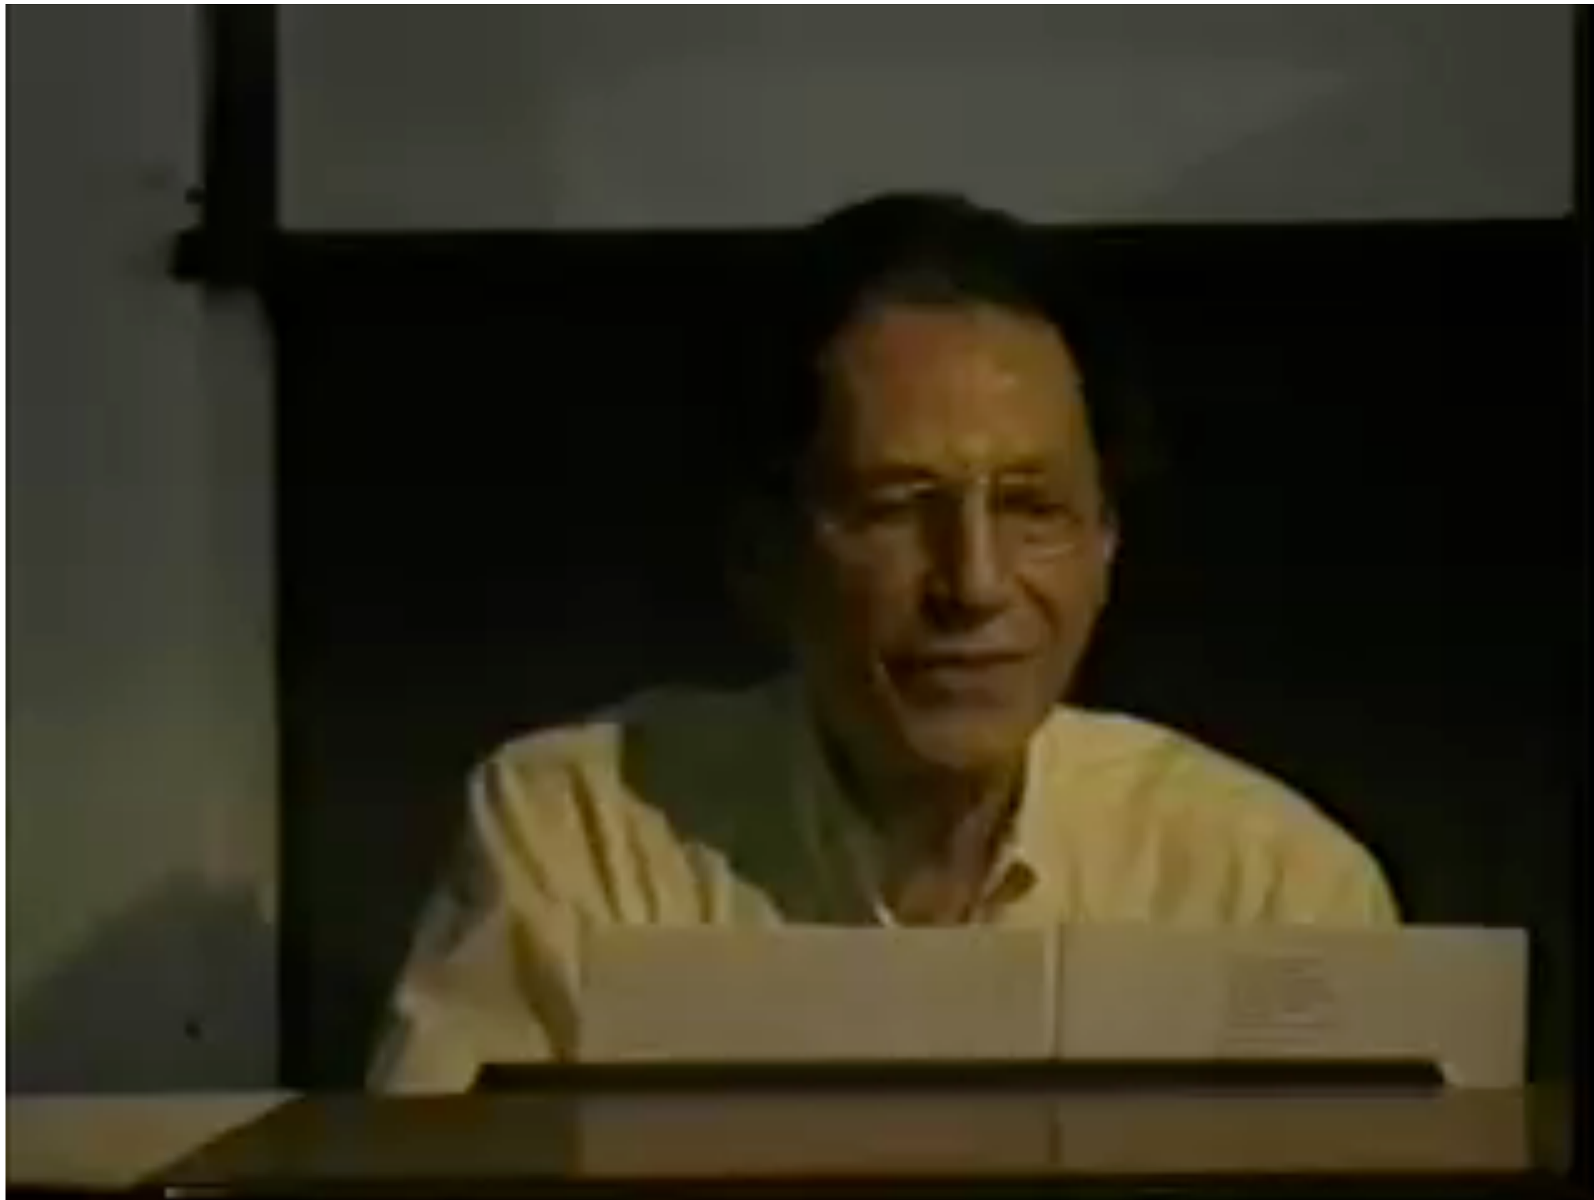

Cue a Tom Lehrer Song