Life and its relatives

Manuel Baltieri - Oxford, 23rd April 2025





Contents

Artificial life as abstraction

- A change in perspective
 - From simulating life...
 - ... to life within simulations
- Abstracting evolution
- Abstracting self-sustaining systems







Defining life

NASA's working definition of life

"A self-sustaining chemical system capable of undergoing Darwinian evolution."

Simulating (features of) life as we know it



Gray-Scott model of reaction diffusion https://www.youtube.com/watch?v=F5oKgVZ6bTk&ab_c hannel=TimHutton

"A self-sustaining chemical system capable of undergoing Darwinian evolution."



Avida https://www.youtube.com/shorts/ripHOtOG4TE

"Comparative biology"

"[W]e badly need a comparative biology. 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."

Maynard-Smith (1992)



Life in silico



Lenia https://www.youtube.com/watch?v=HT49wpyuxk&ab_channel=BertChan



Abstracting evolution

Darwinian evolution





© 2013 Encyclopædia Britannica, Inc.

Phanerozoic Eon

Digital evolution

Tierra, Avida, etc.





Avida https://www.youtube.com/shorts/apANCQmPmL0

Abstracting self-sustaining systems



Dynamical systems theory

Definition.

A (closed discrete deterministic) dynamical system is a pair $(X, f: X \rightarrow X)$ where X is a set and $f: X \to X$ is a state-transition map.

Autonomous, open, continuous, nondeterministic, ... systems require some more machinery, but same idea.



Wikimol, Dschwen - Own work based on: images Lorenz system r28 s10 b2-6666.png by Wikimol and Lorenz attractor.svg by Dschwen, https://en.wikipedia.org/wiki/Dynamical_system#/media/ File:Lorenz_attractor_yb.svg







Variational principles

- Optimise function(al) to determine dynamics/behaviour of a system
- Goals as optimisation towards final state
- E.g. Hamilton's principle

$$\frac{\delta S}{\delta \mathbf{q}(t)} = 0 \qquad S[\mathbf{q}] := \int_{t_1}^{t_2} L(\mathbf{q}(t), \mathbf{q}(t))$$



), t)dt

@MarbleASMRace
https://www.youtube.com/shorts/w5Fom-pdwNA



Optimal control theory

Optimal control/reinforcement learning

- "Physics with inputs", or rather action policies
- Optimise function(al) to determine dynamics/behaviour of a system, but
- Systems have inputs
- Inputs, action policies, to a physical system that can be "chosen" by a controller/agent





Homeostasis



Received 4 October 2022; accepted 29 January 2023. Date of publication 10 February 2023 date of current version 9 March 2023. Recommended by Senior Editor Prof. Wei Zhang. Digital Object Identifier 10.1109/OJCSYS.2023.3244089

The Internal Model Principle for Biomolecular Control Theory

ANKIT GUPTA ^(D) AND MUSTAFA KHAMMASH ^(D) (Fellow, IEEE)

(Position Paper) Department of Biosystems Science and Engineering, ETH-Zürich, 4058 Basel, Switzerland CORRESPONDING AUTHOR: MUSTAFA KHAMMASH (e-mail: mustafa.khammash@bsse.ethz.ch)

ABSTRACT The well-known Internal Model Principle (IMP) is a cornerstone of modern control theory. It stipulates the necessary conditions for asymptotic robustness of disturbance-prone dynamical systems by asserting that such a system must embed a subsystem in a feedback loop, and this subsystem must be able to reduplicate the dynamic disturbance using only the regulated variable as the input. The insights provided by IMP can help in both designing suitable controllers and also in analysing the regulatory mechanisms in complex systems. So far the application of IMP in biology has been case-specific and ad hoc, primarily due to the lack of generic versions of the IMP for biomolecular reaction networks that model biological processes. In this short article we highlight the need for an IMP in biology and discuss a recently developed version of it for biomolecular networks that exhibit maximal Robust Perfect Adaptation (maxRPA) by being robust to the maximum number of disturbance sources.

Robust perfect adaptation in bacterial chemotaxis through integral feedback control

Tau-Mu Yi*[†], Yun Huang^{†‡}, Melvin I. Simon^{*§}, and John Doyle[‡]

*Division of Biology 147-75 and [‡]Department of Control and Dynamical Systems 107-81, California Institute of Technology, Pasadena, CA 91125

Contributed by Melvin I. Simon, February 7, 2000

Integral feedback control is a basic engineering strategy for ensuring that the output of a system robustly tracks its desired value independent of noise or variations in system parameters. In biological systems, it is common for the response to an extracellular stimulus to return to its prestimulus value even in the continued presence of the signal—a process termed adaptation or desensitization. Barkai, Alon, Surette, and Leibler have provided both theoretical and experimental evidence that the precision of adaptation in bacterial chemotaxis is robust to dramatic changes in the levels and kinetic rate constants of the constituent proteins in this signaling network [Alon, U., Surette, M. G., Barkai, N. & Leibler, S. (1998) Nature (London) 397, 168–171]. Here we propose that the robustness of perfect adaptation is the result of this system possessing the property of integral feedback control. Using techniques from control and dynamical systems theory, we demonstrate that integral control is structurally inherent in the Barkai-Leibler model and identify and characterize the key assumptions of the model. Most importantly, we argue that integral control in some form is necessary for a robust implementation of perfect adaptation. More generally, integral control may underlie the robustness of many homeostatic mechanisms.

- Homeostasis implies **internal models**
- In biology "perfect adaptation", in control theory "internal model principle", etc.

electrical network design (14). Quantitative application of engineering robustness methods in molecular biology began with studies of biosynthetic pathways (15), although robustness of biological responses as a selective property in evolution was emphasized qualitatively even earlier (16).

neering systems has a large theoretical literature that began with

In an elegant study, Barkai and Leibler investigated the robustness of perfect adaptation in bacterial chemotaxis (17). They constructed a two-state model (active or inactive) of the receptor complex (receptor + CheA + CheW); the system output, modulated by ligand binding and methylation, was the concentration of active receptor complexes. In this model, perfect adaptation was the intrinsic property of the connectivity of the signaling network and did not require specific values for the kinetic rate constants or concentrations of the constituent enzymes. Alon et al. elegantly provided experimental evidence for the robustness of perfect adaptation to parameter changes when they demonstrated exact adaptation even when the levels of the chemotactic proteins were varied dramatically (18). In this work, we have reexamined these findings from the perspective of robust control theory, which has allowed us to analyze in a more



Available online at www.sciencedirect.com SCIENCE DIRECT.





Systems & Control Letters 50 (2003) 119-126

www.elsevier.com/locate/sysconle

Adaptation and regulation with signal detection implies internal model

Eduardo D. Sontag^{*,1}

Department of Mathematics, Rutgers University, New Brunswick, NJ 08903, USA

Received 14 June 2002; accepted 12 December 2002

Abstract

This note provides a simple result showing, under suitable technical assumptions, that if a system Σ adapts to a class of external signals \mathcal{U} , in the sense of regulation against disturbances or tracking signals in \mathcal{U} , then Σ must necessarily contain a subsystem which is capable of generating all the signals in \mathcal{U} . It is not assumed that regulation is robust, nor is there a prior requirement for the system to be partitioned into separate plant and controller components. Instead, one assumes that a "signal detection" property holds.

© 2003 Elsevier B.V. All rights reserved.





Internal models imply ...?

Abstract homeostasis + reverse the implication

- Not all internal models are good interesting, some are *too simple*
- No full definition of individuation, but "agent" with models" is a good first step
- Agents: "goal-directed autonomous systems that interact with their environment"







"Too simple" = we somehow know all the movements of guards inside a castle without seeing them properly (slit windows)



Internal models for homeostasis

From Baltieri et al. (2025), these are too simple

Definition II.3 (Map of systems). Let X : $Sys({}^{I}_{X})$ and $X' : Sys({I' \atop X'})$ be systems. A map of systems $f : X \to X'$ is comprised of two parts:

1) a *map on states*, given by a function

$$f_{\mathsf{s}}: X \to X', \tag{2}$$

2) a *map on inputs*, given by a function

$$f_{\mathsf{i}}: X \times I \to I',$$
 (3)

such that the following diagram commutes:

$$\begin{array}{ccc} X \times I & \xrightarrow{(\pi_X \,\,\mathrm{\$}\,\,f_{\mathrm{s}},f_{\mathrm{i}})} \to X' \times I' \\ & \mathsf{pd}_{\mathsf{X}} \downarrow & & \downarrow \mathsf{upd}_{\mathsf{X}'} & & \downarrow \mathsf{upd}_{\mathsf{X}'} \\ & X & \xrightarrow{f_{\mathrm{s}}} & \to X' \end{array} \tag{4}$$

meaning that, for every $x \in X, i \in I$, the following equation is satisfied:

$$f_{\mathsf{s}}(\mathsf{upd}_X(x,i)) = \mathsf{upd}_{X'}(f_{\mathsf{s}}(x), f_{\mathsf{i}}(x,i)).$$
(5)

Definition II.7 (Model). Given systems $X : Sys({}^{I}_{X})$ called the *referent* and M : $\mathbf{Sys}\begin{pmatrix} J \\ M \end{pmatrix}$ called the *referrer*, a *model* is a map $\mu : X \to M$ such that

- 1) its part on states $\mu_s : X \to M$ is surjective, and
- 2) its part on inputs $\mu_i(x, -) : I \to J$ is surjective for each $x \in X$.

Remark II.8. When applied to autonomous systems, a model reduces to the definition implicit in [27], [28], namely a surjective map of states commuting with the dynamics:



This is because the map on inputs of a model is necessarily of the form $X \times 1 \rightarrow 1$ for autonomous systems, due to the surjectivity condition. Therefore systems that model autonomous systems must also be autonomous, explaining the need for Assumption 3 later on.

Bayesian models or interpretations

Systems that perform Bayesian updates "in the wild" are agents



Conjugate priors for Bayesian filtering

—> There exists a map *C* such that



See Virgo et al. (2021), Biehl & Virgo (2022), Baltieri @luminousmen.com et al. (2025), these are **not** too simple

Special mention

Randall Beer's work on gliders in the Game of Life



https://en.wikipedia.org/wiki/Conway%27s_Game_of_Life



Figure 10. Three examples of structurally coupled interactions from a preliminary study of (1,2)-environments. The path through a glider's cognitive domain corresponding to the uppermost interaction is shown at the bottom, with dashed lines indicating the path the glider would have followed in isolation and thick lines indicating its structurally coupled path.

Beer (2014)

Summary

- Life *within* a simulation
- Abstracting self-sustaining systems
- We can do maths on some things, not on others yet
- To understand individuation, we define "agents"
- Agents as systems that can be interpreted as using Bayesian models to updates their states



