



HAL
open science

AI-driven Automated Discovery Tools Reveal Diverse Behavioral Competencies of Biological Networks

Mayalen Etcheverry, Clément Moulin-Frier, Pierre-Yves Oudeyer, Michael Levin

► **To cite this version:**

Mayalen Etcheverry, Clément Moulin-Frier, Pierre-Yves Oudeyer, Michael Levin. AI-driven Automated Discovery Tools Reveal Diverse Behavioral Competencies of Biological Networks. *eLife*, 2024, 10.7554/eLife.92683.3 . hal-04395732v2

HAL Id: hal-04395732

<https://hal.science/hal-04395732v2>

Submitted on 30 Dec 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

AI-driven Automated Discovery Tools Reveal Diverse Behavioral Competencies of Biological Networks

Reviewed Preprint

v3 • August 21, 2024

Revised by authors

Reviewed Preprint

v2 • June 6, 2024

Reviewed Preprint

v1 • February 15, 2024

Mayalen Etcheverry, Clément Moulin-Frier, Pierre-Yves Oudeyer, Michael Levin 

INRIA, University of Bordeaux, Talence 33405, France • Poietis, Pessac 33600, France • Allen Discovery Center, Tufts University, Medford MA 02155, USA

 https://en.wikipedia.org/wiki/Open_access Copyright information

Abstract

Many applications in biomedicine and synthetic bioengineering depend on the ability to understand, map, predict, and control the complex, context-sensitive behavior of chemical and genetic networks. The emerging field of diverse intelligence has offered frameworks with which to investigate and exploit surprising problem-solving capacities of unconventional agents. However, for systems that are not conventional animals used in behavior science, there are few quantitative tools that facilitate exploration of their competencies, especially when their complexity makes it infeasible to use unguided exploration. Here, we formalize and investigate a view of gene regulatory networks as agents navigating a problem space. We develop automated tools to efficiently map the repertoire of robust goal states that GRNs can reach despite perturbations. These tools rely on two main contributions that we make in this paper: (1) Using curiosity-driven exploration algorithms, originating from the AI community to explore the range of behavioral abilities of a given system, that we adapt and leverage to automatically discover the range of reachable goal states of GRNs and (2) Proposing a battery of empirical tests inspired by implementation-agnostic behaviorist approaches to assess their navigation competencies. Our data reveal that models inferred from real biological data can reach a surprisingly wide spectrum of steady states, while showcasing various competencies that living agents often exhibit, in physiological network dynamics and that do not require structural changes of network properties or connectivity. Furthermore, we investigate the applicability of the discovered “behavioral catalogs” for comparing the evolved competencies across classes of evolved biological networks, as well as for the design of drug interventions in biomedical contexts or for the design of synthetic gene networks in bioengineering. Altogether, these automated tools and the resulting emphasis on behaviorshaping and exploitation of innate competencies open the path to better interrogation platforms for exploring the complex behavior of biological networks in an efficient and cost-effective manner. To read the interactive version of this paper, please visit <https://developmentalsystems.org/curious-exploration-of-gm-competencies>.

eLife assessment

This **important** study develops a machine learning method to reveal hidden unknown functions and behaviors in gene regulatory networks by searching parameter space in an efficient way. **Solid** evidence is presented for the method, which should be of broad interest to anyone working in biology, as the ideas put forward by the authors extend beyond gene regulatory networks to reveal hidden functions in any complex system with many interacting parts.

<https://doi.org/10.7554/eLife.92683.3.sa3>

Introduction

Developing methods to recognize, map, predict, and control the complex, context-sensitive behavior of chemical and genetic networks is an essential frontier of research in science and engineering. These systems, such as gene regulatory networks and protein pathways, are known to be instructive drivers of embryogenesis, cell behavior, and complex physiology [1]–[3]. Understanding the control properties of these systems is critical not only for the study of evolutionary developmental biology [4]–[8], but also for comprehending and intervening in various disease states, including cancer [9]–[11], and for the construction of novel synthetic biologicals in bioengineering contexts [12]–[16].

Thus, much work has gone into mathematical modeling and computational inference of both protein pathways and gene regulatory network models [17]–[20], which has resulted in the development of large collections of publicly-available models such as the Biomodels database [21], [22]. Yet, despite the wealth of available models, scientists still largely lack an effective understanding of the range of possible behaviors that these models can exhibit under different initial conditions and environmental stimuli, and are in search of systematic methods to reveal and optimize those behaviors via external interventions. The full extent of the computational and control properties of such networks are not yet well-understood; while dynamical systems theory has been extensively used to characterize their behavior [23], [24], it is not known what other sets of tools might reveal and exploit interesting properties of this ubiquitous biological substrate. The field of diverse intelligence (also known as basal cognition) has suggested that strong functional symmetries between pathway networks and neural networks could imply the existence of learning and other kinds of behavior in this unconventional substrate [25]–[29]. Specifically, it has been hypothesized that gene regulatory networks (GRNs) and other molecular networks could be endowed with surprising navigation competencies allowing them to robustly reach diverse homeostatic or allostatic states despite a wide range of perturbations [30]–[33], and that exploiting these innate competencies could provide a promising roadmap for the design of interventions in regenerative medicine and bioengineering contexts [34], [35].

However, significant challenges remain in practice for the exploration and behavior-shaping of these innate competencies, which presents a barrier to the use of these ideas in regenerative medicine and bioengineering. Because of the non-linearity and redundancy in pathway dynamics, passive exploration strategies such as random screening are likely to either fail in uncovering the full range of potential behaviors or require time and energy beyond the available resources. Here, we formalize and investigate a view of gene regulatory networks as agents navigating a problem space. We propose a framework and automated tools, leveraging (1) curiosity-driven goal-directed exploration algorithms coming from recent advances in machine learning and (2) a battery of empirical tests inspired from behaviorist approaches, for mapping the repertoire of robust goal

states that GRNs can reach within this problem space despite various perturbations. A key novelty of this work is the use of AI-based exploration tools to map the space of possible behaviors in biological networks, which opens interesting avenues for efficient mapping of unfamiliar system behaviors, yielding transferable insights for diverse problem-solving once such a map is discovered.

The challenge of exploring and mapping spaces of complex and self-organized behaviors appears in many fields such as diverse intelligence in biological systems, minimal active matter, and robotics: many systems in these areas provide a rich space of evolved, engineered, and hybrid systems that offer many of the same fundamental problems of behavior and control regardless of specific composition or provenance [36]. These span many orders of spatio-temporal scale, from molecular assemblies to swarms of complex organisms [28], [37]–[39]. One set of approaches seeks to develop tools to identify the optimal level of control, ranging from physical rewiring to various methods from cybernetics and behavioral sciences, to reveal and exploit the native competencies and computational capacities of these systems [12]. Specifically, it is increasingly realized that the level of competency (and thus the appropriate level of control) often cannot be guessed by inspection of a system's components, and that its position on a spectrum ranging from passive matter to complex metacognition must be determined empirically [40]–[42], [36]. This is critical not only for fundamental understanding of evolution of bodies and minds [43]–[45], [26], [46], [47], but also for the design of interventions in biomedicine and synthetic morphology contexts [48], [49]. Yet, a common property in many of these systems is that it is expensive in time and energy to conduct experiments: empirical exploration needs to be made under limited resources. Thus, methods for automating efficient exploration and discovery of a diversity of behaviors in these spaces may be widely useful. As explained below, we will here leverage methods from developmental artificial intelligence initially designed for the specific purpose of exploring a diversity of behaviors using a limited budget of experiments.

One especially fascinating set of systems concerns cellular molecular pathways, or gene regulatory networks (GRNs). In the lab or clinic, these pathways are usually treated as simple machines, with intervention strategies focusing on rewiring their structure to achieve a desired outcome: adding or removing nodes (gene therapy), or changing connection weights (by targeting promoter sequences or protein structures) [50]–[53]. However, the emergent, generative nature of development and physiology ensure that it is often very hard to know which genes/proteins to modify, and how, in order to reach a complex desired system-level outcome [54]. Moreover, the responses of cells and tissues to drugs changes over time, making it even more difficult to infer specific interventions (e.g. drugs) that will induce a stable improvement in pathway state *in vivo*. Indeed, with the exception of antibiotics and surgery, most available treatment modalities do not solve the underlying problem – they seek to mitigate symptoms, which recur (or expand) once the drug is withdrawn. This is because current therapeutics function bottom-up, attempting to force specific molecular states, as it has been challenging to develop methods for shifting complex tissues and organs towards a stable health profile. Next-generation solutions, which would offer true healing (stable correction), require an understanding of the homeostatic and allostatic properties of networks with respect to how they traverse the space of transcriptional, physiological and anatomical states. An understanding of the behavior policies of networks as they dynamically navigate these problem spaces is essential for predicting what stimuli can be used to re-set their setpoints and guide them to autonomously maintain a healthy state. In the language of behavioral neuroscience, this strategy corresponds to exploiting their native robustness, decision-making, and navigational competencies to induce predictable, long-lasting changes in functionality.

Significant challenges remain in revealing and controlling the range of behaviors that can self-organize in these cellular and molecular pathways. To characterize steady-state concentrations and responses to small perturbations, conventional methods rely on piecewise-linear

approximation of the system behavior [55]–[59], but struggle with higher-dimensional systems or wider parameter ranges which limits their applicability [60]. Other works have proposed the porting of tools from network control theory to identify sets of control nodes allowing to drive the network behavior toward target steady states [61]. These methods typically exploit the network topology [61]–[65] or regulatory structure [66]–[68] to identify control strategies based either on permanent knockout/activation of genes or on temporary perturbations, the latter being preferable in biomedical context.

However, these approaches often require prior knowledge of target attractor states or are limited to Boolean network models. Other works have explored the use of machine learning tools, such as evolutionary search [69]–[71] and gradient-descent optimization [72], [73], for controlling continuous ODE biomolecular networks with high-dimensional parameter spaces, mainly in the context of synthetic circuit engineering [74], [75]. While providing powerful optimization tools, these approaches tend to focus on rewiring network structure and connectivity. Moreover, the choice of a predefined fitness function and parameter range initialization is not only critical to the success of optimization [70] but largely restricts exploration of the behavior space [73].

In contrast, an alternative line of research proposes exploring and leveraging the inherent molecular mechanisms of adaptivity and robustness in cellular pathways as a promising approach for drug interventions that do not rely on genomic editing or gene therapy [30], [76]. Recently, a broad, substrate-independent behavior science perspective suggests novel properties of gene regulatory networks (GRNs) and other biological networks [25], [77]. This perspective views GRNs as agents that convert activation levels of specific genes (inputs) to those of effector genes (outputs), with intermediate nodes in between, leading to strategies for controlling network behavior based on a specific history of inputs (experience) rather than through network rewiring. Notably, the concept of training a chemical pathway using pulsed input stimuli (node activation or suppression drugs) has been formalized, and several networks have been analyzed to establish a taxonomy of memory types found in biological GRNs and pathways [78], [79].

Here, building upon recent research [32], [78], [79], we take the next step and investigate a view of gene regulatory networks as agents navigating a problem space toward target goal states with varying degrees of competency (**Figure 1-a**). We seek to implement a definition of goal that abstracts it from conventional associations with human or other advanced brains and facilitates the use of tools from cybernetics, behavior science, and control theory to understand broader aspects of biological regulation. Here we use the term “goal” state to refer to a system’s steady state, which it expends effort to reach despite interventions or barriers - a definition appropriate to the study of basal (or minimal) proto-cognitive regulatory systems.. Our definition of goal does not imply “purpose” (high-level goals where an agent has the meta-cognition to think about having goals and what they might be), and we do not attribute high-level competencies (such as re-setting one’s own goals) to GRNs.

Our particular focus lies in investigating two types of navigation competencies: *versatility*, which refers to the capacity to reach diverse goal states under different interventions, and *robustness*, which refers to the ability to reach a goal state despite various perturbations. The primary scientific question we aim to address is: What is the repertoire of robust goal states that a GRN can actively reach through *minimal* and *non-genetic* interventions within a navigation task context, and can we develop systematic methods and automated tools to aid scientists in discovering this repertoire?

To address this question in practice, our experimental framework revolves around the definition of “problem spaces”, which we use as tractable components of the GRN’s overall state space (**Figure 1-b**), and on a set of methodological contributions which we organize around three sub-questions:

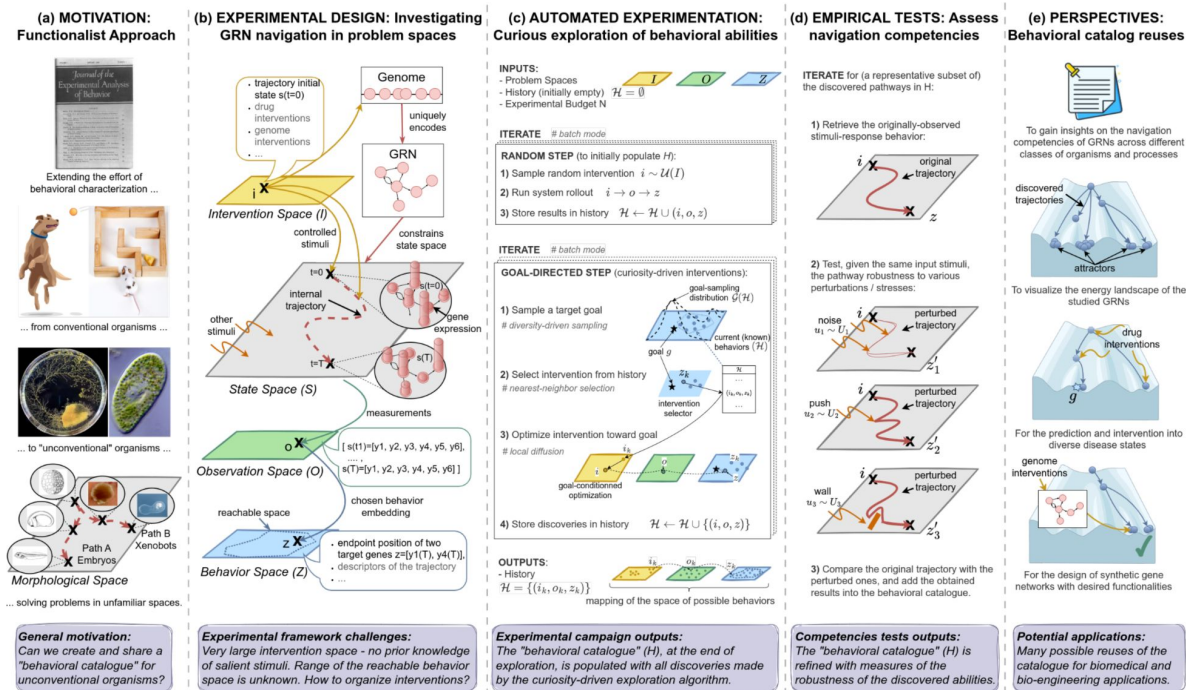


Figure 1

Overview of the proposed framework.

(a) MOTIVATION: We often focus on studying the navigation and behavior of organisms in conventional three-dimensional environments, neglecting the intelligence underlying competencies at sub-organismal scales [32]. To better understand navigation competencies in unconventional organisms solving problems in unconventional spaces (e.g., embryos in morphological space), it is essential to construct comprehensive "behavioral catalogs" for these novel entities, which in turn requires sophisticated exploration methods to discover the extent of possible behaviors. Images are taken and adapted from [80]–[85]. (b) EXPERIMENTAL DESIGNS: We formalize GRN behavior as a navigation task and propose to investigate it by defining abstract and observer-dependent "problem spaces" that we use to organize the observed biological behaviors and their exploration in practice. (c) AUTOMATED EXPERIMENTATION: Pseudo-code of the curiosity-driven goal exploration process we use to automate the discovery of behavioral abilities that the GRN can exhibit in behavior space. (d) EMPIRICAL TESTS: We use a battery of empirical tests to identify the robust goal states of the systems, i.e. the one that can be attained under a wide variety of perturbation (including noise in gene expression, and pushes or walls during traversal of transcription space). (e) PERSPECTIVES: We explore several potential reuses of the discovered "behavioral catalog" and proposed framework across evolutionary biology, biomedicine and bioengineering contexts.

1. *Automated discovery of diverse behavioral abilities with autotelic curiosity search (Figure 1-c)*: What is the range of possible goal states that GRNs can exhibit and how can we devise efficient exploration strategies to automatically identify these goal states? Defining goal states as attractor states of the underlying gene regulatory network, we show that traditional screening methods can be very inefficient in discovering the range of possible goal states. To address this, we propose to use intrinsically-motivated goal exploration processes (IMGEP) [86], [87], a recent family of diversity-driven machine learning approaches also known as autotelic curiosity search which was recently shown to form a useful discovery assistant for revealing the behavioral diversity of unfamiliar systems such as chemical oil-droplet systems [88], physical non-equilibrium systems [89] and models of continuous cellular automata [90]–[92].
2. *Evaluation of the navigation competencies (Figure 1-d)*: How competent is the GRN, in terms of robustness to perturbations, in attaining the diverse previously-identified goal states? Prior studies have offered definitions of robustness in biological networks, characterized as the degree of variation in functionality [93] or phenotypic trait [94] under specific environmental or genetic changes. However, these studies often consider a predefined functionality and random perturbations in network parameters [95], [96], [71] or specific gene knockouts [97]. Environmental perturbations on the other hand are often limited to random variations in initial conditions within a predefined range [60], [98]. Here, inspired from behaviorist approaches, we test hypotheses about non-genetic resistance with respect to various navigation competencies that living agents often exhibit, and that do not require structural changes of network properties or connectivity. Those tests assess the system’s ability to maintain robustness despite various perturbations encountered during traversal, including developmental noise in gene expression levels, sudden “pushes” within transcriptional space, and the presence of energy barriers or “walls” acting as force fields in the environment.
3. *Potential reuses of the discovered “behavioral catalog” and framework (Figure 1-e)*: Can the constructed behavioral catalogs be useful for fundamental research and practical therapeutic applications, and can the framework be easily applied to other systems and problem spaces? We propose that the discovered competencies may provide valuable insights for understanding evolvability and developmental robustness, and provide a fertile source for the design of interventions in biomedicine and synthetic morphology contexts. We also suggest that the framework and automated tools, which are observer-focused and substrate-independent, could be transposed to other systems and problem spaces.

The overall framework is summarized in **Figure 1**. Applying it on a database of 30 continuous (ODE) models from the Biomodels website, consisting of a total of 432 systems defined as GRN model-behavior space tuples, revealed several interesting insights. First, results suggested that most of the surveyed systems are capable of reaching a surprisingly wide spectrum of steady states depending on their initial state. Interestingly, random screening strategies were not able to reveal this diversity of reachable states (or at least not in a sample efficient way), confirming the need for more advanced exploration strategies like curiosity search. Secondly, among the discovered steady states, we were able to identify several robust goal states i.e. ones that the system consistently reaches despite various perturbations during traversal of transcriptional space. Altogether, these findings seem to suggest that cell phenotype and functionality could be the result of a multi-step program [62] that could be flexibly and robustly reprogrammed by appropriate stimuli [42]. Finally, we demonstrate possible reuses of this “behavioral catalog” for comparing the network’s competencies across different classes of organisms, as well as for the design of non-genetic drug interventions. We also demonstrate an alternative reuse of the framework to reveal new kinds of reachable “goals” in synthetic gene networks, suggesting alternative strategies for the design of gene networks in a bioengineering context.

An interactive executable version of the paper, as well as step-by-step tutorials and notebooks can be found online at <https://developmentalsystems.org/curious-exploration-of-grn-competencies>. The full codebase of the proposed automated experimentation pipeline is written end-to-end in JAX, a high-performance numerical computing library that we leverage for parallel experimentation and computational speedups of the ODE models time-course simulations.

Results

Generalizing GRN behavior as a navigation task

The GRNs analyzed in this study are biological pathway networks taken from the BioModels repository [21], [22]. The term “GRN” is used broadly to include protein interaction, gene regulatory, and metabolic networks. In these mathematical models, the dynamic interactions between nodes of the network (molecular species) are modeled with a system of ordinary differential equations, enabling to quantitatively simulate time-course behavior (model rollouts) and observe the dynamics of node activities over time (Figure 2- a). Here, following a terminology which aims to integrate concepts from dynamical complex systems with concepts from behavioral sciences, we propose to conceptualize GRN behavior as a *navigation task* (Table 1). Model rollouts are viewed as “trajectories” in transcriptional space where network steady states are “goal states” (endpoints) that the “agent” (GRN) can reach with varying levels of competencies. As for living agents, these competencies may range from unstable locomotion patterns to more advanced forms of goal-directed behavior like path following, obstacle avoidance, or even forms of spatial memory and foresight. In this paper, we are particularly interested in investigating two forms of navigation competencies that we refer to as *versatility*, the capacity to reach diverse goal states under various interventions, and *robustness*, the capacity to reach a goal state despite various perturbations. Note that versatility and robustness are studied with respect to different sources of incoming environmental variation, respectively interventions and perturbations.

To investigate these competencies in practice, our experimental framework is based on the definition of “problem spaces”, which include the observation space (O), behavior space (Z), intervention space (I) and perturbation space (U) as defined in Table 2. To be consistent with our navigation task terminology introduced in Table 1, we refer to a behavior z as the reached “goal state” of a GRN trajectory. However these “goals” may lie on a continuum between complete robustness and high sensitivity, and our primary interest lies in identifying *robust* goals of the system. Whereas several choices could be made for the intervention space I and perturbation space U, we intentionally consider *minimal* and *non-genetic* interventions to investigate the “native” goal states of the GRN, and *environmental obstacles* to investigate for navigation competencies classically observed in other living agents. Examples of simulations, interventions, and perturbations are illustrated in Figure 2.

Then, a typical analysis using our framework relies on a 2-step procedure, detailed in the subsequent sections. First, to assess the versatility of the GRN, we define an exploration strategy which organizes the sequence of interventions i_1, \dots, i_N used to drive the system toward a maximally diverse set of reachable endpoints $\{z_k \in Z\}_{k=1 \dots N}$, while being given a limited budget of experiments N. Secondly, to assess the robustness of the discovered goal states $\{z_k \in Z\}$, we conduct a battery of empirical tests to characterize their degree of sensitivity to novel perturbations, with a fixed experimental budget of P perturbations per selected behavior z . At the end of this 2-step procedure, we obtain the “behavioral catalog” (H) of the studied GRN, which includes the history of experiments $H = \{(i_k, o_k, z_k, \{(u_p, o_p, z_p), P = 1 \dots P\}), k = 1 \dots N\}$.

Dynamical Systems Terminology	Behavioral Science Terminology	Proposed Isomorphism	Navigation Task Terminology
system: a set of interconnected elements that interact to produce emergent behavior	organism: a living being that responds to stimuli and adapts to its environment	Both are collections of lower-level elements that interact to produce emergent behavior and can adapt at the system level	agent or GRN
phase-space trajectory: set of states taken by the system when starting from one particular initial condition	behavioral trajectory: the sequence of states that an organism exhibits in response to stimuli	Both represent the sequence of states or behaviors that a system or individual experiences over time	trajectory
initial condition: initial state of a system's variables and parameters that condition its dynamics	stimuli: events that might (or might not) trigger a response in an organism	Both represent incoming variations that set a system or organism in motion	intervention or perturbation
critical parameter: a parameter or condition that, if changed, can cause a system to undergo a qualitative change or phase transition	salient stimuli: stimuli that are particularly relevant or meaningful to an organism, either because they are associated with reward or punishment or because they are novel or unexpected	Both represent the incoming variations that have a significant impact on a system's steady-state or organism's response	salient intervention
steady-state (or attractor): a stable state (or set of states), towards which the system tends to evolve over time	observed response: outcome or endpoint of a behavioral trajectory towards which an organism converges	Both represent the endpoint that a system or organism is moving towards	reached endpoint or goal
robust attractor: stable attractor toward which the system tends to evolve under various initial conditions and perturbations	target goal: it is assumed that an organism engages in a goal-directed manner when it exhibits new ways or actions to achieve a similar outcome when faced with novel circumstances	Both represent a stable endpoint or goal that the system successfully attains under various perturbations	robust goal
controllability: degree to which the system's dynamics (and resulting steady states) can be controlled or manipulated	trainability: degree to which an organism's behavior can be modified or shaped by experience or conditioning	Both measure the extent of states that can be reached by a system or individual under a distribution of stimuli/conditions	versatility

Table 1

Glossary of terms used in this paper, with the proposed isomorphism which investigates concepts from dynamical complex systems and behavioral sciences under a common navigation task perspective.

Problem Space	Generic definition	Specific definition in this study
Observation Space (O)	Space of raw observations made during the GRN model rollout to measure its state or behavior.	Records node activities over time as $o = (y(0), \dots, y(T))$, where $y(t)$ is an n-dimensional vector (n = number of nodes) and T is the measured reaction time.
Behavior Space (Z)	A projection of the observation space used by the experimenter to encode the “goal states” of a model rollout into a tractable (lower-dimensional) space.	Encodes the trajectory endpoint of a model rollout. Represents a cell phenotype defined by the state values of some nodes (relevant biological markers), such that $z = (y_{i_1}(T), \dots, y_{i_m}(T))$ (we use $m=2$ in this study for simplicity and visualization).
Intervention Space (I)	A space where interventions represent controlled sources of incoming variation that the experimenter can exert on the GRN model rollout to drive it toward novel or targeted states.	Sets the initial state $i = (y_1(0), \dots, y_n(0))$ of a model rollout. Defined as a hyper-rectangle $I \subseteq \mathbb{R}^n$ where the boundaries are proportional to the min and max values taken by the respective nodes from default initial conditions.
Perturbation Space (U)	A space where perturbations represent external sources of incoming variation, used by the experimenter to characterize the robustness of a given goal state.	Includes three classes of (stochastic) perturbations including noise perturbation U_n , push perturbation U_p , and wall perturbation U_w .

Table 2

Problem spaces used in this study.

Following this framework, the behavioral catalog is constructed for a database of 30 biological networks consisting of a total of 432 systems, where a system is defined as a (GRN model, intervention space (I), behavior space (Z)) tuple, as described in Materials and Methods and **Table S1**. These catalogs provide valuable empirical observations and insights into the navigation competencies of the studied GRNs, particularly in their ability to consistently achieve diverse goal states under various tested perturbations. Statistical analyses of the results are presented in **Figures 3**, **5**, and **7**, and specific results for the RKIP-ERK signaling pathway [99] are shown in **Figures 2**, **4**, **6**, and **8**.

Curiosity Search Uncovers a Diversity of Reachable Goal States

One advantage of modeling GRN behavior within a tractable behavior space Z is that we can then deploy strategies to efficiently discover and map that space. Notably, recent diversity-driven machine learning techniques such as Novelty Search [100], [101], Quality Diversity [102], [103] and Intrinsically-Motivated Goal Exploration Processes (IMGEP) [86], [87] are explicitly designed to efficiently explore a so-called “behavior space” or “goal space” which is basically a (predefined or learned) model of the overall state space. In particular IMGEPs, which were originally developed for the learning of inverse models of highly-redundant mapping in robotics context [86], were recently shown to successfully assist discovery in complex self-organizing systems [88]–[91].

Here, we propose to use an IMGEP to control GRN initial states and maximize the diversity of discovered endpoints $\{z \in Z\}$ within a limited budget of N experiments. The IMGEP operates in two phases: initially, N_{int} interventions are sampled randomly from I to populate history H , then remaining interventions are generated through a goal-directed process which relies on several key internal models. Those including a goal-embedding module (R) that encodes observations (o) into the IMGEP goal space (T), a goal generator module (G) that samples goals from the goal space based on intrinsic motivation incentives (e.g. to promote novelty or learning progress), and a goal-conditioned optimization policy (Π) that generates candidate intervention parameters to achieve the current goal. Given those internal models, the goal-directed phase iterates through 1) sample a target goal $g \sim G(H)$, 2) infer intervention parameters to achieve that goal $i \sim \Pi(g, H)$, 3) conduct an experiment with the intervention i , observe the outcome o , and compute the reached goal $z = R(o)$, and 4) store the tuple (i, o, z) in history H . Here, we use the GRN behavior space Z as the IMGEP goal space $T = Z$. Hence “target goal” refers to a goal sampled by IMGEP while “reached goal” refers to an actual endpoint of the GRN trajectory, discovered by IMGEP while targeting a potentially different point in Z . Throughout exploration, the IMGEP dynamically refines its Z -traversal strategy based on the knowledge acquired by its discoveries. Here we opt for a simple IMGEP variant such that the exploration process can be seen as performing novelty search in behavior space Z [104]. The pseudocode of our IMGEP pipeline is shown in **Figure 1-c** and details about the internal models are provided in Materials and Methods. The final outcome is a “behavioral catalog” of the GRN, containing the diverse goal states discovered by IMGEP: $H = \{(i_k, o_k, z_k), k = 1 \dots N\}$.

We deploy the IMGEP, also known as “curiosity search,” on all 432 systems in the biological network database. Our evaluation focuses on two related competencies: the IMGEP agent’s ability to empirically reveal a diversity of reachable goal states in the (GRN, I , Z) system, referred to as “discovered diversity,” and the GRN agent’s competency to naturally reach diverse goal states, referred to as “versatility.” The true versatility of the GRN is unknown and can only be inferred through empirical exploration and proxy metrics.

For evaluating diversity, we measure the area covered in Z by the IMGEP discoveries using the threshold-coverage metric [105] and compare it with the area covered by the diversity of a naive random screening strategy (which uniformly samples interventions in I). In **Figure 3**, the diversity discovered by the two exploration variants is shown for the 432 (GRN, I , Z) systems,

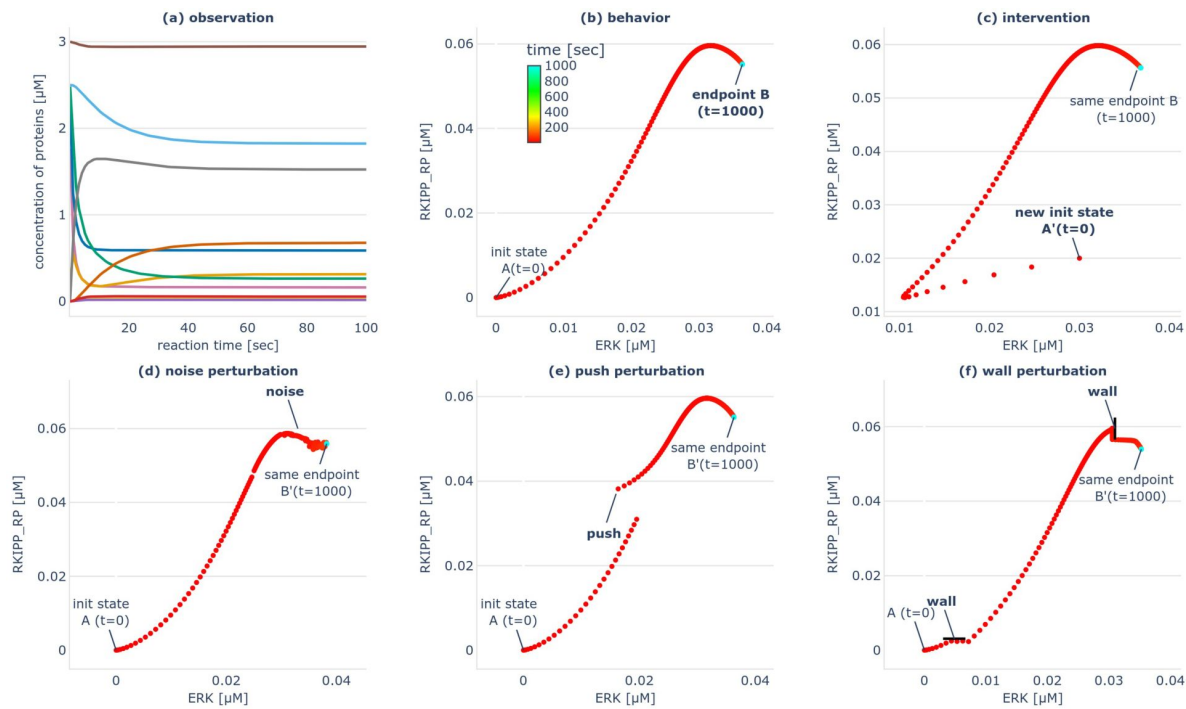


Figure 2

Illustration of the experimental setup and chosen problem spaces on an example GRN model which has 10 nodes and models the influence of RKIP on the ERK Signaling Pathway [99].

(a) Time-course evolution of the different nodes y_1, \dots, y_{10} (one color per node) when starting from the default initial conditions (as provided in [99]). The observation captures the states taken through time $o=[y(t=0), \dots, y(t=T)]$ where $y=[y_1, \dots, y_{10}]$. (b) Corresponding trajectory in transcriptional space (phase space), for two target nodes (ERK, RKIPP_RP), from $t=0$ (A, in red) to $T=1000$ seconds (B, in cyan). We can see that the trajectory converges to endpoint B in less than 100 seconds, and then stay there. The behavior (or reached goal state) is the endpoint $B = [y_{ERK}(T), y_{RKIPP_RP}(T)]$, where T is chosen big enough to ensure convergence. (c) The intervention is setting the initial state of the system trajectory (for all nodes): $i = [y_1(t=0), \dots, y_{10}(t=0)]$. (d-e) Example of perturbations used in this paper. (d) Noise perturbation, here applied to all 10 nodes every 5 secs until $t=80$ secs. (e) Push perturbation, here applied to the two target nodes (ERK, RKIPP_RP) at $t=3$ seconds. (f) Wall perturbation, also applied to the two target nodes (ERK, RKIPP_RP), here at 10% and 90% of the total distance traveled. **Supplementary Figure S1** shows examples of other possible “drug” or “genome” interventions that can be implemented in the accompanying software, as well as the possibility to perform interventions (or perturbations) in parallel using batched computations.

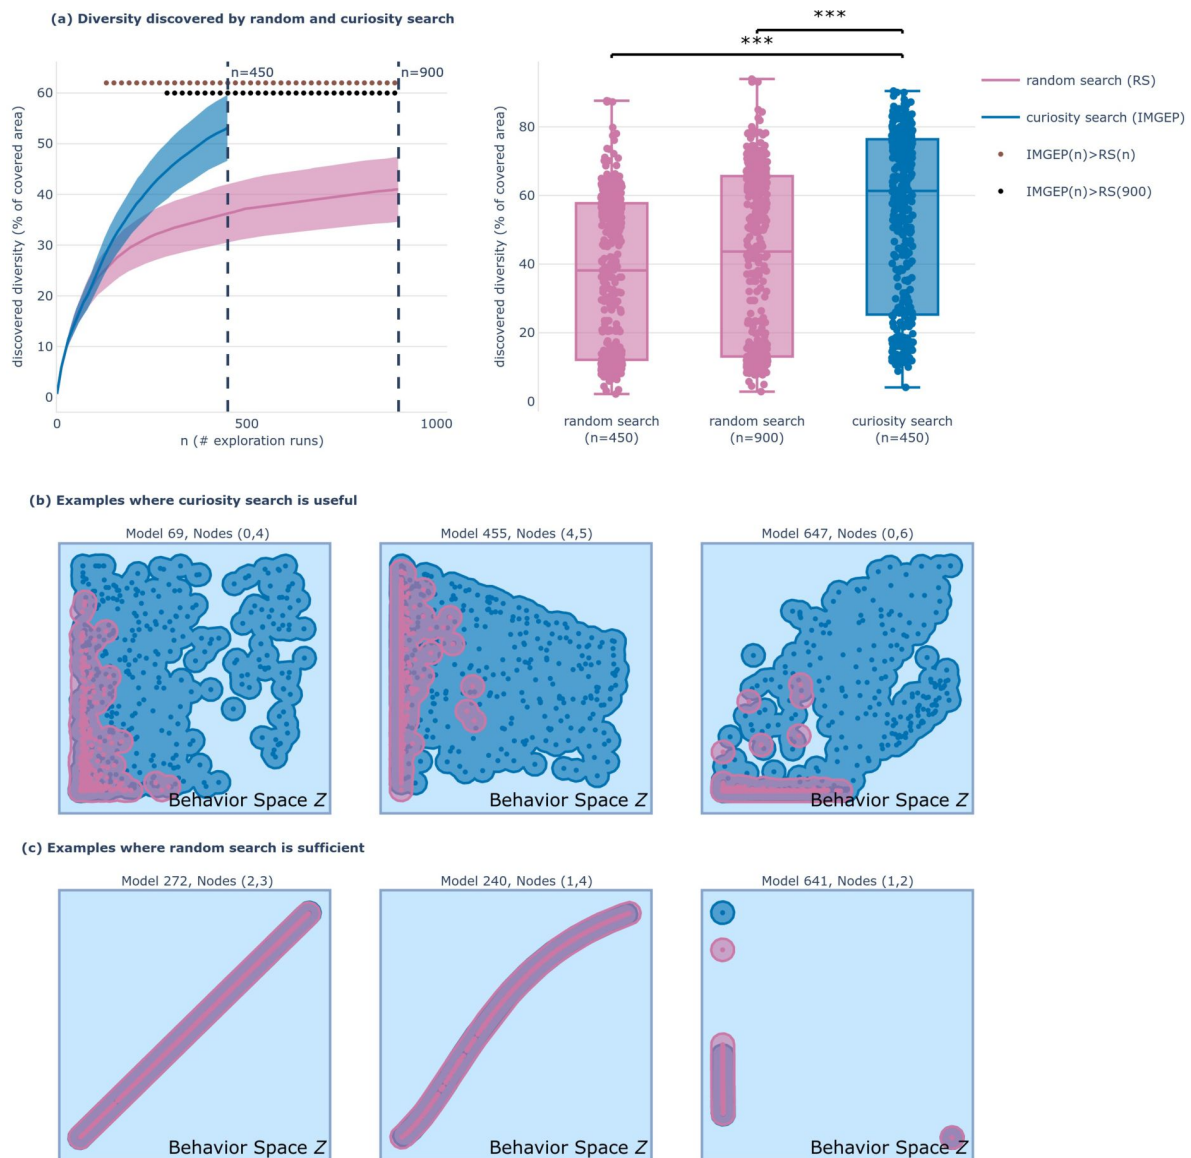


Figure 3

Curiosity search uncovers a wide spectrum of reachable states in behavior space Z.

(a) Diversity of endpoints discovered by random search (pink) and curiosity search (blue) for the 432 systems. Diversity is measured as the volume of the union of the set of hyperballs of radius ϵ that have for centers the discovered endpoints $\{z \in Z\}$ as depicted by the shaded area in (b-c) with $\epsilon = 0.05$. (a-left) Mean and standard deviation curves of the diversity of behaviors discovered throughout exploration (with random search having twice more experiments $n=900$). Dots indicate significance ($p < 0.05$) when testing curiosity search (n) against random search (n) in brown, and against random search ($n=900$) in black, with a Welch's t-test. Standard deviation is divided by 4 for visibility. (a-right) Detail of the diversity obtained in the left plot for all 432 systems at $n=450$ and $n=900$, where *** indicate significance ($p < 0.001$). (b-c) Discovered endpoints at the end of exploration ($n=450$) by random search (pink) and curiosity search (blue) for 6 example systems of our database. (b) Examples of systems for which curiosity search is much more sample-efficient than random search in finding a diversity of reachable states in behavior space Z. (c) Examples of systems with low-redundancy mapping $I \rightarrow Z$ such that random search in I is already quite efficient in covering behavior space Z, and curiosity search performs equivalently.

where random search is given a budget of experiments N which is twice bigger ($N=900$) as the one given to the curiosity-search algorithm ($N=450$). Interestingly we see that, on average, at $n=290$ the curiosity search already significantly outperforms the final diversity achieved by random search, while only utilizing one-third of its experimental budget ($N=900$). Whereas we are dealing with numerical systems and our codebase allow for efficient and parallel execution, each experiment still consists of $\frac{T}{\Delta T} = 25000$ model steps, where each step integrates the ODE system. Repeating that N times for each of the 432 systems starts to be very costly, which is why having efficient exploration strategies is very valuable (and would be even more valuable when scaling the framework to larger databases). Even more critical, as illustrated in [Figure 3-b](#), it seems that, for some systems, random search is not able to discover the “latent” regions revealed by the IMGEP in Z , or it would need an extremely large budget of experiments. On the other hand, as illustrated in [Figure 3-c](#), there are some systems for which random search is already quite efficient in revealing diverse behaviors in Z , and for which IMGEP performs equivalently.

In fact, the goal-directed strategy of the IMGEP is particularly beneficial for (GRN, I, Z) systems with high nonlinearity or redundancy in their $I \rightarrow Z$ mapping, as seen in [Figure 4](#) and studied in robotics contexts [[105](#)]. Redundancy implies that many interventions in I lead to similar effects in Z , as illustrated in [Figure 2](#) where various interventions and perturbations converge to the same endpoint. In these systems, random search will preferentially discover points in areas of high redundancy in Z whereas the IMGEP, whose exploration is directed uniformly in goal space, should cover different levels of redundancy equally. In general, when dealing with large intervention spaces and limited experimental budgets, curiosity search can be particularly useful for efficiently navigating Z -space.

Finally, as the IMGEP efficiently drives the GRN into diverse goal states with minimal interventions, we propose that the diversity achieved by the IMGEP can serve as a good proxy metric of the GRN versatility. Notably, analysis of example systems in [Figure 3](#) reveals that many GRNs can reach a broad spectrum of steady states. Whereas our database is limited to certain systems (see Materials and Methods) and might not be representative of all biological pathways, this observation underlines the existence of various phenotypes that can be realized. It also highlights the critical importance of identifying salient interventions that can effectively control cellular states within this spectrum of possibilities, notably as many cancer types are due to epigenetically nonidentical cells [[107](#)].

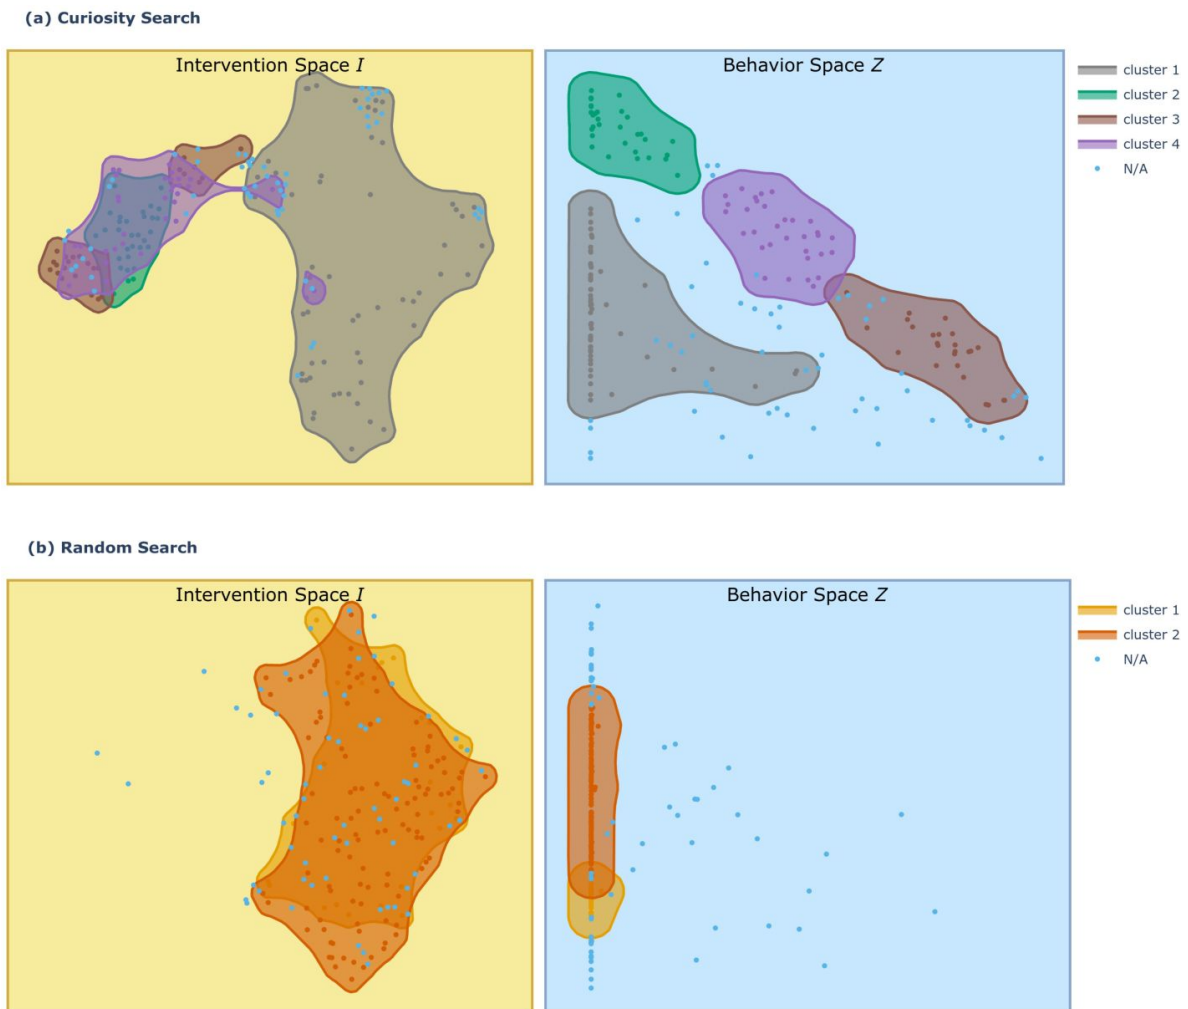


Figure 4

Illustration of the non linearity and redundancy of the $I \rightarrow Z$ mapping, and of the interest of using goal-directed exploration strategies.

Plot shows the reachable points discovered by curiosity search (a) and by random search (b) in the behavior space Z and their corresponding starting points in the intervention space I, for the RKIP-ERK signaling pathway system [99]. The intervention space is 10-dimensional, and here we show the TSNE reduction in 2D. We apply HDBSCAN clustering [106] on the points discovered in Z, which produced 4 clusters for curiosity search (displayed in gray, green, purple and orange; non-assigned points are displayed in light blue) and 2 clusters for random search (displayed in light and dark orange). We then visualize where those regions in behavior space mapped back in the intervention space, by applying the same coloring. (a) Looking at the curiosity search discoveries, we can see the non-linearity of the $I \rightarrow Z$ mapping, where small regions of intervention space can map to large regions of the behavior space (like the orange area) and reversely (gray area). We can also see the redundancy of the behavior space which is clearly concentrated in the left border of the space (ERK close to zero) which can seemingly be reached from very large portions of the intervention space (gray area). (b) Looking at random search discoveries, we can understand that it is very inefficient as it spends most of its exploration budget in the region of intervention space that converges to the left border in Z, and fails to explore the orange, purple and green regions discovered by curiosity search which seemingly lead to the more novelty in Z.

Empirical Tests Reveal Robust Navigation Competencies

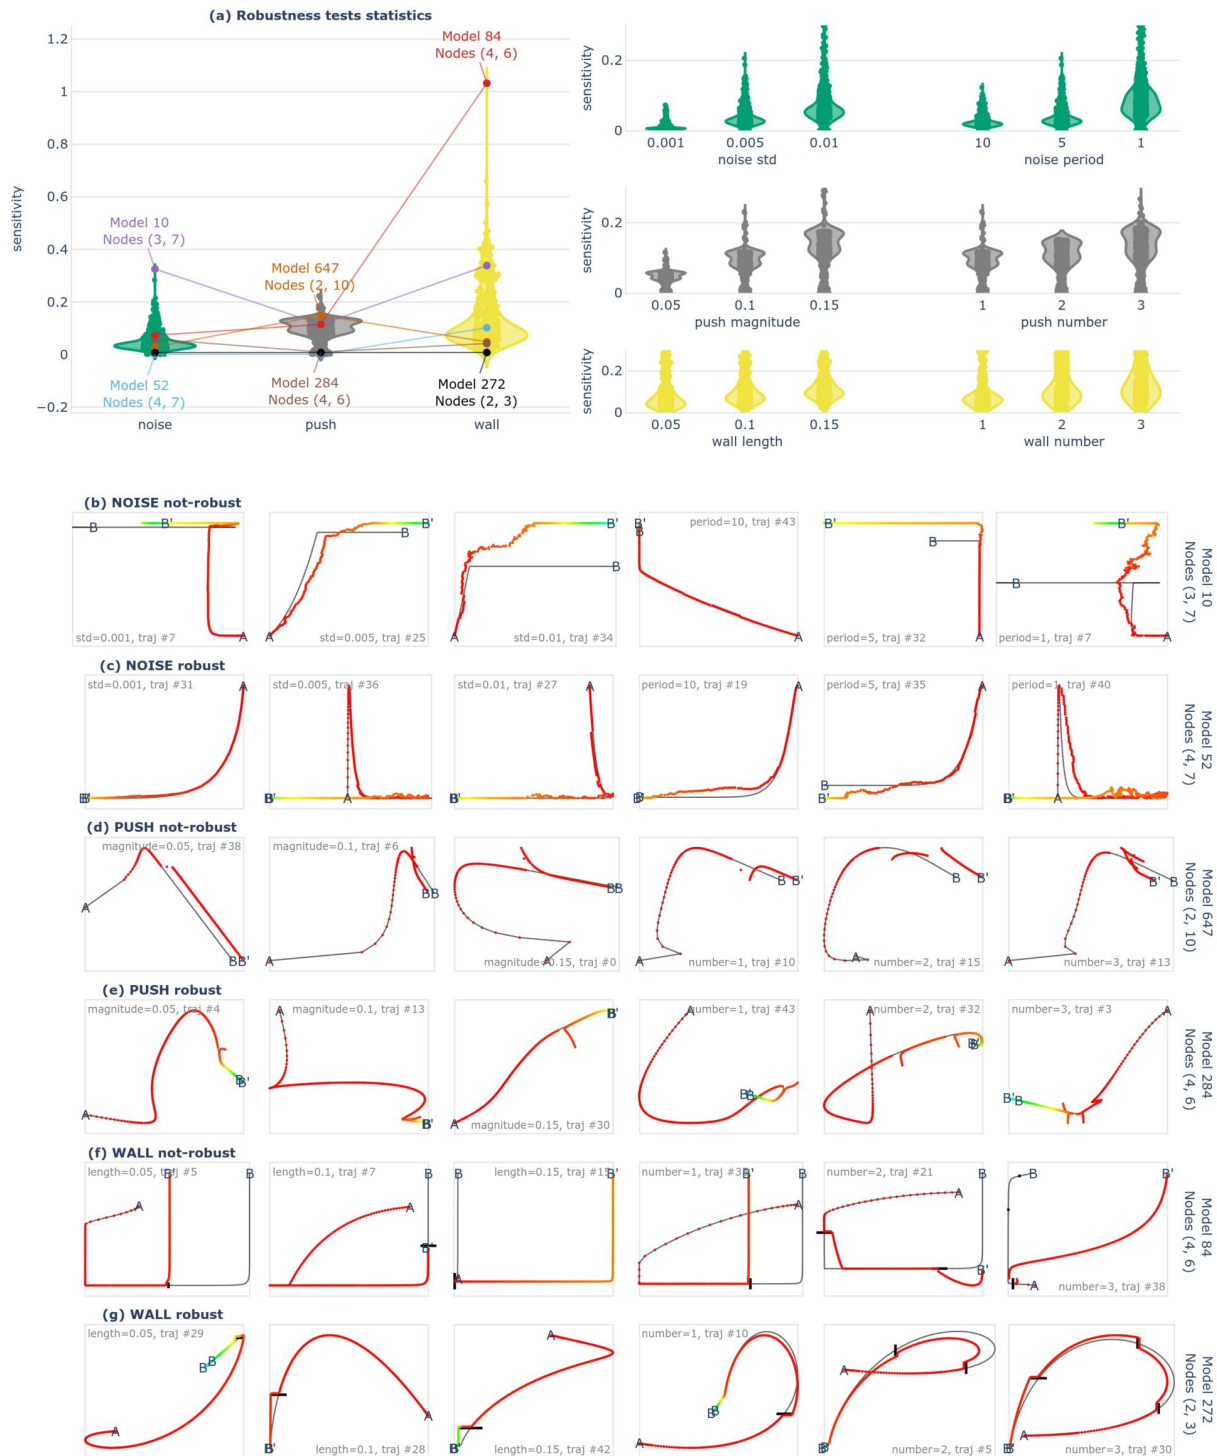


Figure 5

Identification of robust traversal strategies in transcriptional space.

(a) Violin plots show, for each of the 432 systems (one point per system), the median sensitivity (over the K representative goal states) to the noise (green), push (gray) and wall (yellow) perturbation families. Violin plots on the right detail the median sensitivity for the 18 sub-families. (b-g) Each row provides examples of perturbed trajectories of either extremely-robust or extremely-sensitive example (GRN, Z) system (on average over the K goal states) for the three families of perturbations, as shown by annotations in (a). For instance, the first row (b) shows perturbed trajectories of the (model #10, nodes (3,7)) system which has the highest sensitivity to noise whereas the last row (g) shows trajectories of the (model #272, nodes (2,3)) system which has a nearly perfect robustness to walls. Each image contains an example trajectory for a given (i, u) , and one u per sub-family is shown per column. For instance, in the first row (b), the trajectories are perturbed with the different sub-families of noise ($\sigma_n \in [0.001, 0.005, 0.1]$, $p_n \in [10^{-5}, 5^{-5}, 1^{-5}]$) which can be seen as various levels of difficulty. For each trajectory we annotate the starting position (A), endpoint prior perturbation (B), and endpoint after perturbation (B'), and show the original trajectory in black. The perturbed trajectory is shown in colorscale (from red at $t=0$ to cyan at $t=3000$ secs). (b) Except for few cases (trajectory #43), the system (model #10, nodes (3,7)) system is not robust to noise as its trajectories are easily deviated from the original endpoint. (c) The (model #52, nodes (4,7)) system however, except for rare cases (trajectory #35), consistently reaches its original target despite encountering various amounts of noise. Interestingly, trajectories #36 and #40 consistently follows a complex up->right-down->left path, despite the induced noise. (d) The (model #647, nodes (2,10)) system, except for few cases (trajectory #0), is typically deviated from its original trajectory when being pushed away. Interestingly though, it seems to follow similar (parallel) trajectories. (e) The (model #284, nodes (4,6)) system, is an example of an extremely robust system which, despite many push configurations (in magnitude and number), consistently returns to its original trajectory. Interestingly, the trajectories of this system are relatively complex with several loops and detours. (f) The (model #84, nodes (4,6)) system is not very robust to walls, and typically deviates or blocked when it encounters a wall. (g) The (model #272, nodes (2,3)) system is another example of an extremely robust system which, despite many wall configurations (in length and number), consistently returns to its original path. Once again interestingly, the trajectories of this system are relatively complex with several loops and detours.

We are then interested in characterizing the degree of robustness of the previously-discovered “goal states” in order to identify the ones that can consistently be reached by the GRN despite encountering various perturbations. Whereas many studies have proposed rigorous analysis of the “robustness” of biological networks [93], [94], the generated perturbations often target variations in the regulatory rules (i.e. variations at the hardware level) and variations are often sampled independently (and prior) to observations of the GRN dynamical behaviors [98], [95], [96], [71], [108], [60]. Here instead, we propose to conduct a battery of empirical tests that draw inspiration from classical “displacement experiments” [109], [110] and “barrier experiments” [111] commonly used in behavioral sciences to assess the navigation competencies of various animals. As illustrated in **Figure 2**, we consider *environmental* perturbations that perturb the GRN trajectory with 1) various degree of noise in the gene expression levels, 2) sudden “pushes” during the GRN traversal of transcriptional space, and 3) energy barriers or “walls” acting as new force fields that constrain the GRN traversal. Importantly, those perturbations are *conditioned* on the observed behavior of the GRN. The magnitude of the noise and of the pushes is scaled proportionally to the extent of the observed trajectories, and the walls are generated in locations of the space that the GRN would “naturally” visit without the induced perturbation. While intuitive from a behaviorist point of view, where one would adapt experimentation when testing animals in different contexts (e.g. to study homing behavior of an ant and of a sea turtle, or of an ant in food deprivation and in reproduction phase) [112], robustness studies in systems biology often neglect those aspects. We propose that a behaviorist lens on robustness can help understanding forms of non-genetic resistance in transcriptional space, which is crucial for the development of therapeutic strategies [107].

To assess the degree of robustness of the discovered goal states, our evaluation procedure is the following. For each (GRN, I,Z) system of the database, we retrieve a representative set of trajectories previously discovered using the curiosity-search algorithm and subject these trajectories to $P = s \times r$ perturbations conditioned on the GRN goal-reaching trajectory $i \rightarrow z$ prior perturbation. Here, s represents the different perturbation distributions which correspond to various “tests” and “levels of difficulty” (e.g. noise magnitude and frequency, number of walls, etc.) and r is the number of (stochastic) perturbations sampled per family. The pseudocode is illustrated in [Figure 1-c](#) and details about the different family of perturbations are provided in Materials and Methods. At the end of this process, the behavioral catalog is augmented with the perturbed trajectories $H = \{(i_k, o_k, z_k, \{(u_p, o_p, z_p), p = 1 \dots K\}), k = 1 \dots K\}$.

As the use of “spaces” comes with the notion of similarity and distance, we can then easily evaluate the *sensitivity* of a goal state z with respect to a set of perturbation $\{u_p, p = 1 \dots P\}$ as the average distance in behavior space Z between the original trajectory endpoint z and the perturbed trajectories endpoints $\{z_p\}$. Here our distance is simply the Euclidean distance, normalized by the extent of the trajectory prior perturbation in Z . We can then identify the so-called “robust goals” of the systems as the ones that have the lower sensitivity to perturbations. These sensitivity analyses can be useful in two important ways. On the one hand, they allow us to quickly identify the “extreme” examples of robustness, both at the system-level and at the goal-level, providing several insights into the degree of “competencies” that some biological networks might exhibit in their relative space ([Figure 5](#)). On the other hand, these analyses also allow us to map the *heterogeneity* of cellular responses and to better understand how non-genetic perturbations might modulate the *landscape* of reachable cell phenotypes ([Figure 6](#)).

[Figure 5](#) shows the median sensitivity, over the representative goal states, for the 432 systems of our database and for the noise, push and wall perturbations families (as well as for the $s=18$ sub-families which correspond to varying degrees of perturbations). Overall, even though we observe varying degrees of sensitivity between systems (and between magnitudes of perturbations, which is expected), one first and interesting observation is that the median sensitivity remains relatively low, suggesting that GRNs could not only exhibit versatility (with respect to the considered interventions) but also robustness (with respect to the considered perturbations). In fact, looking at the “extreme” examples, we can identify quite impressive examples of complex and yet highly-robust space traversal strategies, with non-linear trajectories exhibiting many “detours” and “loops” but yet consistently reaching the same endpoint despite several pushes ([Figure 5-e](#)) or walls ([Figure 5-g](#)) on the way.

[Figure 6](#) shows how the constructed catalog H can be used to generate the energy landscape of the studied system. In biology, landscape formalisms have been used to comprehend the underlying dynamics of several systems, such as cell cycles and cell differentiation [114], [115]. It is believed that such system-level visualizations could be particularly useful to apprehend non-genetic heterogeneity in the context of cancer treatment and stem cell differentiation [107], [113]. A recent landscape-generation method only proposes to approximate the pseudopotential energy through simulation trajectories obtained throughout exploration of the system [113], making it a widely applicable method which we can directly apply here. However, the paper relied on Monte Carlo simulation to generate the trajectories. Due to the previously mentioned non-linearity and redundancy of the I->Z mapping, this can lead to poor estimation of the overall energy landscape ([Figure 6-a](#)). Instead, when generating the landscape from the trajectories discovered by our curiosity search exploration, we are able to reveal a new and wide “valley” of reachable states ([Figure 6-b](#)). Interestingly, the landscape-generation method can also be used to better comprehend the effect of external cues on the gene regulatory network, by visualizing how much they deform the energy landscape for instance leading to new shaped valleys ([Figure 6-c](#)). For the example system RKIP-ERK pathway [99],




Figure 6

Energy landscape visualization based on the trajectory-based landscape generation method [113], and constructed from different set of GRN trajectories, respectively trajectories generated (a) by the random search exploration, (b) by the curiosity-driven exploration, and (c) by the robustness tests experiments.

results highlighted a specific region of behavior space (with low RKIP and high ERK activation levels) that seems to be particularly robust, i.e. consistently reached by the GRN from certain initial conditions, and that might be associated with tumor development [116].

Possible reuses of the behavioral catalog and framework

Our framework generated a catalog of stimuli, responses, and navigation tests for the different GRN models contained in our database. Creating and sharing such a “behavioral catalog” with the scientific community is possibly one of the more exciting aspects of the work with new organisms. Furnished with such an empirically based data-set and detailed observations, one can 1) conduct statistical analysis across the population of studied organisms to inform fundamental research questions and 2) reuse the acquired knowledge to design specific behavior-shaping experiments in organisms of interest. As our framework focuses on observable behavior and is agnostic about the internal construction of the organism, another exciting perspective is to deploy it to different problem spaces and other classes of natural, chimeric or synthetic organisms. This section illustrates preliminary experiments along those three types of reuse.

To develop insights on the degree of sophistication of the different GRNs

A first use-case we explore is to conduct statistical analysis to categorize versatility and robustness in the surveyed networks on the basis of species in evolutionary strata. We consider seven categories, namely, plant, bacteria, slime mold, amphibian, rodent, homo sapiens, or generic. Here, generic corresponds to the networks not associated with any species but related to generalized biological processes. Please note that the surveyed database is relatively small with respect to the wealth of available models and biological pathways, so we can hardly claim that these results represent the true distribution of competencies across these organism categories. Still, as shown in **Figure 7**, results suggested interesting patterns.

First, on average, generic and Homo sapiens GRNs exhibit higher versatility (mean 0.228 and 0.238) compared to rodent and amphibian GRNs (mean 0.163 and 0.169), which in turn show higher versatility than bacteria and plant GRNs (mean 0.136 and 0.117). These findings are particularly intriguing in the context of the recently-formulated hypothesis of multi-scale competency architecture [42]: could the observed variation in versatility among different classes of GRNs contribute to the degree of versatility observed at higher-level scales? Collecting such experimental data for broader classes of organisms and GRNs will be crucial to understand how competencies at the molecular scale can impact the overall functionality and adaptability of organisms at higher scales, and to understand how evolution might have exploited this modular architecture for shaping the observed adaptivity and reprogrammability of biological systems.

Secondly, when comparing with the versatility of random networks (in black), generated to follow the same distributions of network size and connectivity as biological networks (as proposed in [79], see Materials and Methods), we observe that random network versatility is much lower (<0.026) than the versatility observed in biological networks. Once again, it is difficult to draw strong conclusions as the gene circuit model used for the random networks is relatively limited, whilst generic and studied across a range of biological contexts [117]–[120], and it will be interesting to scale the comparison to a broader and more complex range of ODE-based random models. Still, these findings hint that versatility prevalence might be a strong invariant of biological intelligence shaped by evolutionary processes.

Finally, we categorize the versatility-robustness tradeoff in the different categories of organisms. The idea is to compare the GRN competencies to robustly achieve diverse goal states, for different robustness thresholds. In **Figure 7-b**, we plot the mean and standard deviation pareto curves for the different categories of organisms and observe that, in average, the pareto-optimal solutions are mostly achieved by generic cell GRNs, even though bacteria GRNs can robustly reach more goal states for exigent robustness criteria (high x-values). The slime mold GRN can reach highly

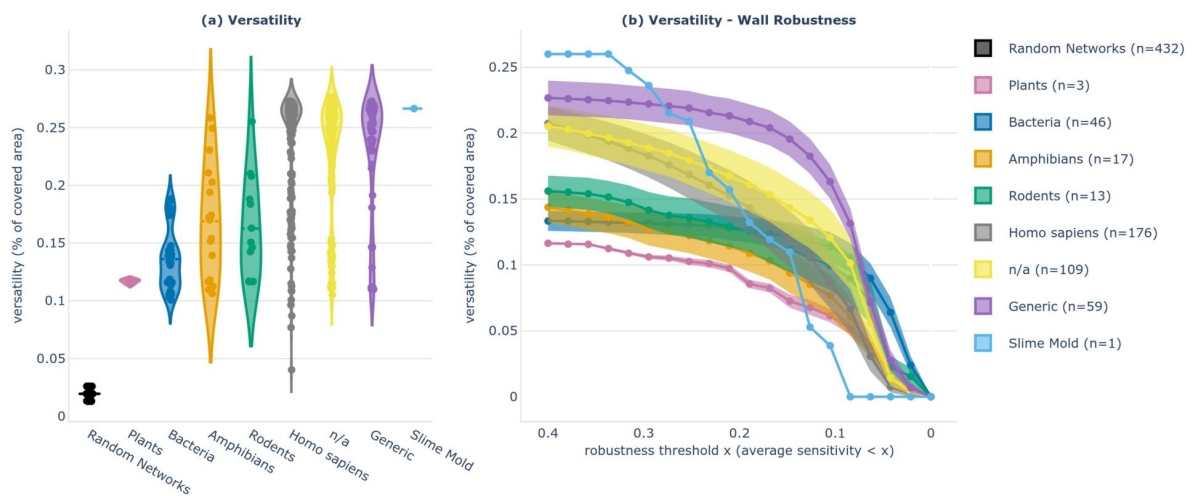


Figure 7

Analysis and comparison of the degree of sophistication, in terms of versatility and robustness, between different classes of GRN. We categorize the GRNs by class of organism they belong to: plant, bacteria, slime mold, amphibian, rodent, homo sapiens, or generic. "n/a" refers to network models for which this information is not available. (a) Violin plots show the versatility of the 432 systems (one point per system) for each class. Versatility of one system is measured as the area covered by all the goal states discovered by curiosity search (equivalent to what we call diversity in [Figure 3](#)). (b) Trade-off (aka Pareto) mean and standard deviation curves that represent the trade-off among versatility and wall robustness performances as taken by the different classes of GRNs (standard deviation is divided by 4 for visibility). For each system, versatility (y-value) is measured as the area covered by the set of robustly achieved goal states, where the criterion of goal-achievement is a binary which tests whether the goal-reaching sensitivity (on average overall wall perturbations) is below a certain threshold (x-values). Violin plots in (a) are ordered in ascending order according to the class mean y-value at $x=0.4$ in (b).

diverse goal states but the tradeoff quickly drops with wall perturbations, and there is only one system in our database belonging to this category so results might be not representative. Once again, those results are very interesting as generic cells GRNs are a building block that has been extensively reused by evolution across several organisms and contexts, bacteria have evolved to be very resistant (e.g. to antibiotics), and slime molds are a unicellular organism well known for its diverse capabilities, especially navigational ones [121]–[124].

For the development of therapeutic interventions

Understanding forms of non-genetic resistance and non-genetic heterogeneity is crucial across a wide range of cancer and treatment contexts [107]. Here, we illustrate how the constructed behavioral catalog can provide a fertile source for the design of therapeutic strategies, notably in the context of network control, using again the example of the RKIP-ERK signaling pathway [99]. In **Figure 4**, we saw that curiosity search revealed four clusters of reachable steady states for this system. From a clinical perspective, one might denote the green cluster as “healthy” region of behavior space and the orange cluster as “disease” region of the behavior space, as high levels of ERK and low-levels of RKIP are often linked to tumor development [116]. In **Figure 8-a**, we plot those two clusters as well as the 10 more robust goal-reaching behaviors from the behavioral catalog of this system, i.e. the goal states with the lower average sensitivity to the induced perturbations. We see that 6 out of the 10 more robust trajectories end up in the “disease” region, suggesting that certain configurations of initial state are very likely to reach that region despite chemical blockers (here pushes, walls, and noise), which was also visible on the system’s energy landscape in **Figure 6-c**. Looking at the six trajectories, it seems that they all follow similar patterns where RKIP activation level increases past a certain threshold, and only then converge toward the disease region. This might already provide an interesting biomarker for prediction of tumor development, but what we are really interested here is to build upon that knowledge to develop stimuli-based interventions allowing to re-set the GRN setpoints from the identified “disease” steady states back to steady states within the identified “healthy” region. To do so, we define a parameterized stimuli-based intervention and a performance function, and search for parameters that optimize this performance. For the intervention function, we use a piecewise constant function that determines which nodes to intervene on (here MEKPP), when to apply the intervention (here every 10 seconds for 100 seconds), and with what amplitude (which are the parameters that we are seeking to optimize). The choice of the intervention function, which is arbitrary in this example, would typically depend on the experimental constraints, e.g. which nodes can be targeted with drugs and at which precision. For the performance function, we define the centroid of the “healthy” region as the target setpoint and compute performance of the stepwise intervention as the average distance of the novel setpoints (after intervention when starting from the 6 disease setpoints) to the target setpoint, and under a distribution of stochastic walls, pushes and noise perturbations. Hence a successful intervention should re-set the disease setpoints to healthy setpoints for all discovered disease states and robustly across the various tested perturbations. For optimization, we simply perform random search as this was sufficient here to discover one intervention (as shown in **Figure 8-b**) that successfully reset the setpoints (as shown in **Figure 8-c**) under various tested perturbations (as shown in **Supplementary Figure S2**). Here random search was sufficient to find a successful intervention, but more advanced optimization strategies like evolutionary algorithms or stochastic gradient descent could be envisaged for harder problems. Overall, mapping the “latent” behavioral abilities of GRNs in healthy physiology and disease states may have important implications for the identification of robust stimuli-based interventions that focus on behavior shaping instead of micromanaging all molecular states, and that can be exploited in therapeutic contexts.

As an alternative strategy to gene circuit engineering

The final type of reuse we explore is not a direct reuse of the constructed behavioral catalogs, but rather a reuse of the proposed automated tools to reveal different kinds of behaviors in a bioengineering context. A common problem in synthetic biology is to optimize the configuration

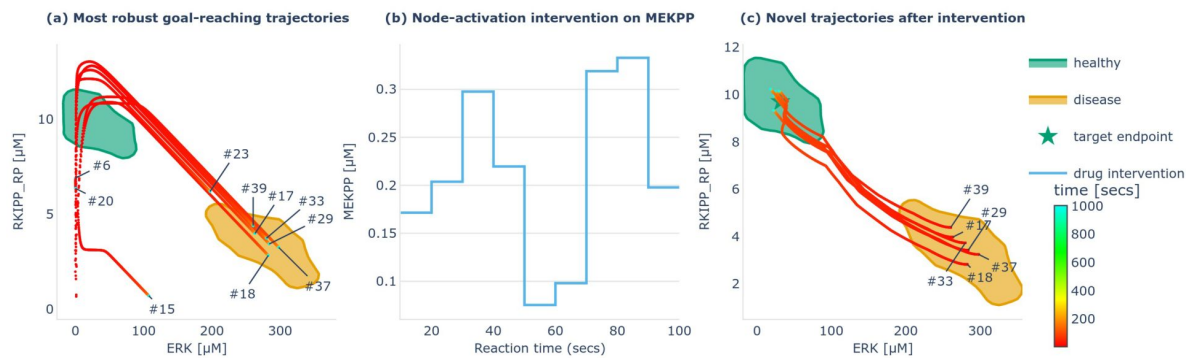


Figure 8

Identification of stimuli-based stepwise intervention triggering robust re-set of disease states into healthy physiological states. (a) The 10 most robust identified goal states (average sensitivity <0.05) and the corresponding reaching trajectories are displayed for the example RKIP-ERK signaling pathway [99]. We can see that most of them converge toward attractors in the “disease” region (orange). (b) Discovered stepwise stimuli intervention on MEKPP which we apply on states stuck in the “disease” region for 100 seconds. (c) The discovered intervention successfully brings back all points from the “disease” region closer to the target endpoint in the “healthy” region, and this under various tested perturbations (as shown in **Supplementary Figure S2**). The optimization procedure that led to the discovery of this intervention is described in the main text.

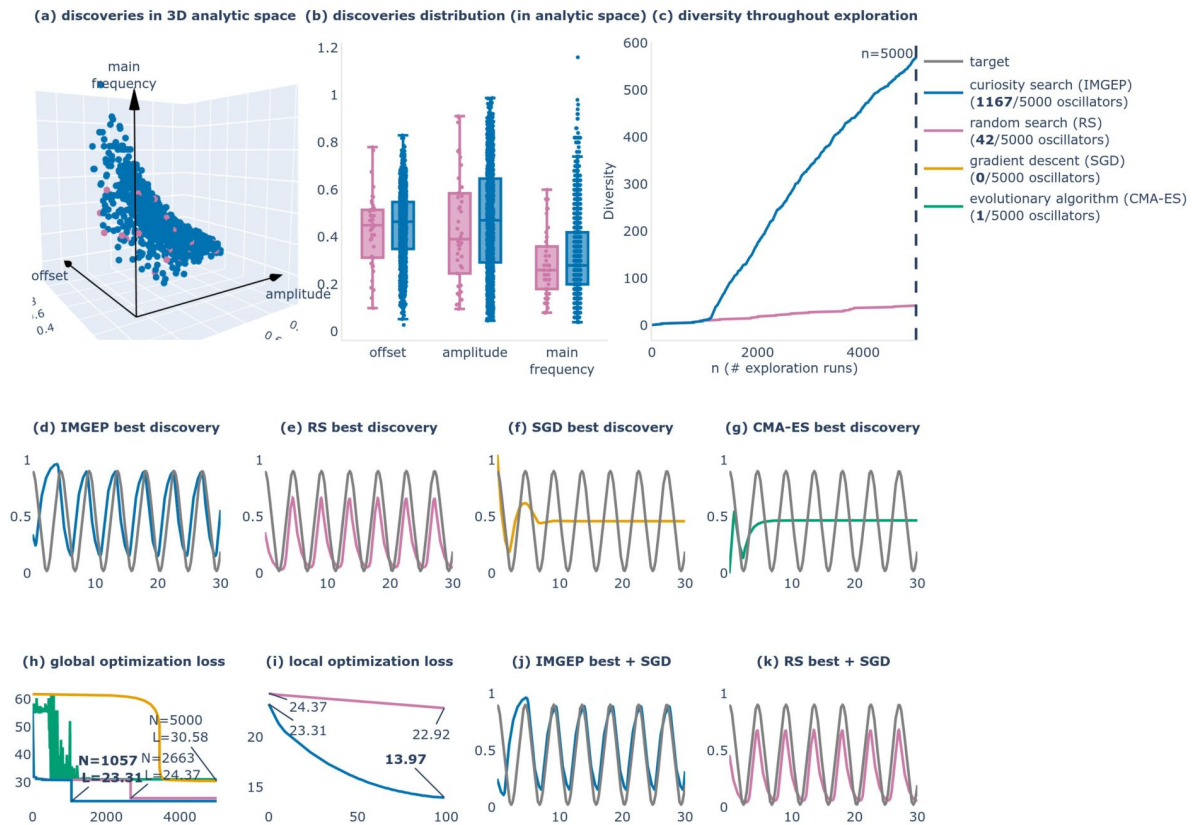


Figure 9

Comparison of four alternative strategies for the design of oscillator circuits: curiosity search (blue), random search (pink), gradient descent (orange) and an evolutionary algorithm (green).

(a-c) Given a budget of 5000 experiments, curiosity search is able to find 1167 oscillator circuits (ones showing sustained oscillations), whereas random search only finds 42 oscillators and optimization-driven search fail to discover them (only one discovered by CMA-ES and none for gradient descent when starting from a single random initialization). (a) 3D scatter plot of the 42 random search discoveries (pink) and 1167 curiosity search ones (blue) in the (amplitude, main frequency, offset) analytic behavior space. (b) Box plots projecting points from the 3D scatter plot into the respective (amplitude, main frequency, offset) axes. (c) Diversity discovered throughout exploration, where diversity is measured with a binning-based space coverage metric (20 bins per dimension). (d-e-f-g) Best discoveries (for which is minimal) made by the four exploration strategies. (h) Evolution of the optimization loss L for the four algorithm variants. (i) Evolution of the local training loss when finetuning of the best IMGEP (blue) and RS (pink) discoveries with gradient descent, with the finetuned results displayed in (j-k).

and parameters of a gene model network to optimally perform a desired functionality, also known as gene circuit engineering [75]. Recent approaches rely on optimization-driven machine learning strategies, such as evolutionary algorithms and stochastic gradient descent. However, choosing the right loss function and parameter initialization for these optimization methods is a well-known problem in machine learning. These issues can lead to optimization algorithms getting trapped in local minima within the complex landscape of possibilities. In response to these challenges, we propose to investigate whether the curiosity-driven exploration strategy can be employed as an alternative (diversity-driven) strategy. Whereas traditionally-employed for exploratory purposes, these exploration strategies were also shown to facilitate the resolution of external, pre-defined tasks characterized by sparse or deceptive rewards [125], by effectively exploring solution space.

Here, we consider the target application of oscillator circuit engineering followed in [72], where parameters of a gene circuit model are optimized to produce oscillation patterns with target amplitude A , frequency ω and offset b . This time, the intervention space includes both genetic interventions (setting kinematic parameters of regulatory rules) and environmental interventions (setting the initial state y_0). We then compare several exploration strategies: a random search, two optimization-driven strategies, one using gradient descent as proposed in [72] and one using an evolutionary algorithm, and finally a diversity-driven strategy using the proposed curiosity search. All algorithms are given the same experimental budget ($N = 5000$). For curiosity search, the behavior space Z is defined as the image space of the discrete Fourier transform of the observation. We then use the exact same IMGEP algorithm as before, but operating within the novel problem spaces (I, Z) . For gradient descent, we follow the procedure proposed in [72]. We define a loss function which measures the mean square error between the observed node activation levels y and the target oscillation (represented as a cosine wave). We then randomly initialize the parameters $i \sim U(I)$ and use Adam optimizer for $N=5000$ optimization steps. For the evolutionary algorithm, we use the CMA-ES algorithm [139], with the negative of the loss as fitness function. In addition, we also use gradient descent for local refinement of the best discoveries made by the other exploration strategies (curiosity search and random search), this time with a limited budget of $N = 100$ optimization steps.

In **Figure 9**, we show that curiosity search is again significantly more efficient than random search in revealing a diversity of possible oscillator behaviors. Out of 5000 trials, random search was able to find only 42 configurations leading to sustained oscillations whereas curiosity search was able to find 1167 (and gradient descent did not find any). Without focusing on the target objective, curiosity search is able to efficiently cover the analytic (A, ω, b) space (**Figure 9, a-c**), thus discovering oscillators close to the target one (**Figure 9-d**). Instead, when starting from a random initial condition, gradient descent is very likely to get trapped in a local minimum where it converges to the target offset b but fails to produce oscillations (**Figure 9-f and 9-h**). While CMA-ES explores more the solution space at the beginning of optimization than SGD does, it also ultimately converges into a similar local minima. However, whereas the global optimization strategies are unsuccessful in this example, they seem to be useful to locally refine close-enough solutions, as can be seen here when refining the best discoveries made by curiosity search and random search with gradient descent (**Figure 9-i, 9-j**). These results suggest that a diversity-driven exploration strategy, eventually combined with a more advanced local optimization strategy, can offer promising and cost-effective alternatives for the design of synthetic gene networks. More generally, as our framework only relies on empirical investigation for inferring the mapping between interventions and behaviors (treating them as abstract variables in observable problem spaces), we believe it offers an exciting perspective to be deployed across various problem spaces and classes of organisms.

Discussion

This paper presents a novel framework aimed at uncovering the navigation competencies of gene regulatory networks (GRNs). The framework conceptualizes GRNs as agents actively navigating the transcriptional space and provides a set of tools, leveraging computational models of curiosity-driven learning and exploration, with a battery of empirical tests inspired from behaviorist tradition, for automated experimentation and behavioral characterization. The proposed framework is novel in two central ways. First, it introduces a novel AI-based toolbox to the field of biological network analysis. We show how this toolbox, leveraging the successful ingredients of recent intrinsically motivated learning algorithms - originally developed to enable robotic AI agents to explore and learn diverse skills in novel and unstructured environments [86], [87] - can be transposed to assist efficient discovery of behavioral abilities within biological pathway models like GRNs. Secondly, rather than merely mapping the attractor states [23], [24], [59] or analyzing their sensitivity to model parameter changes [56], [57] as extensively proposed in conventional GRN analysis methods, our framework investigates the dynamic adaptability of these networks' navigation competencies in response to various changing environmental conditions. With this approach, our aim is to uncover whether diverse competencies, analogous to the ones exhibited by living agents, can be found within physiological network dynamics. Notably, these competencies are discovered without necessitating structural alterations to network properties or connectivity. Importantly, our framework and its associated tools do not make any assumptions about the structure or origin of the biological network, making it in theory adaptable to the study of diverse unconventional intelligences across various domains.

By applying this framework to a curated database of GRN models, we discovered a diverse range of behavioral responses that GRN can exhibit under different initial conditions and characterized their robustness to various perturbations. Notably, our analysis revealed a number of interesting aspects of navigation of the state space which can be leveraged in several contexts. These automated tools form the first step towards cost-effective *in silico* simulation and interrogation platforms; as the "behavioral catalogs" produced by this process can be a first stepping stone for better understanding the GRN functionalities as well as for designing drug-driven interventions in a biomedical or bioengineering context.

There are several limitations and avenues for future work to this study. First, these networks are studied as a model in isolation and it is possible that some of the ODE models (or solvers) provide spurious behaviors within certain parameter ranges which might not map to observable phenotypes *in vitro*. Interestingly, this limitation also suggests an interesting further direction to this work: using the automated discovery toolbox to assist model inference, allowing to efficiently identify the rare or unexpected behaviors of the ODE model and suggest whether further refinement is needed or not. Another interesting direction for future work, as our framework considers the GRN model as a black-box and works with limited experimental budget, would be to directly apply it to *in vitro* GRN models at the bench. One could for instance integrate experimental constraints to the search by defining families of empirically-testable interventions and perturbations, as well as specify clinically-relevant goal spaces and perturbations. Even if in a biological setting versatility and robustness phenomena may be harder to detect, or harder to alter, these results can be used to (1) design synthetic biology circuits with advanced capabilities [126], and (2) conduct studies of subcellular proto-cognitive phylogenetics, to help understand the evolutionary pressures for and against reprogrammability in cell regulatory machinery. Another limitation of our work is that we consider predefined problem spaces, here the space of GRN steady states (or Fourier descriptors of the dynamics in the bioengineering example). The dynamics of gene regulatory networks are relatively simple (usually converge to stable points or periodic orbits) allowing such hand-defined descriptors. To scale the framework to higher-dimensional and more complex problem spaces, recent works from the IMGEP literature suggest

using unsupervised learning of goal space representations [90], [91]. Whereas these works were applied to abstract models of multicellular patterning, similar works could be envisaged in more realistic systems, such as sophisticated model of multicellular morphogen and/or bioelectrical patterning which were used to suggest in-vitro experimental manipulations [127], [128], [128].

The tools presented here, and the behavioral repertoire we identified, are just the beginning, and much work remains. Future efforts must test additional competencies across the spectrum of cognition (memory, creative problem-solving, valence, etc.) and extend the tools we presented here to explore them. The predictions made by our computational tools can now be tested in real cells, using emerging tools for physiological profiling in the living state and a diverse set of biochemical, biomechanical, and bioelectrical perturbations. We anticipate a tight and productive feedback loop between computational theory that suggests new experiments, and results in living cells that greatly extend our computational perspective on what they can do [129]–[133]. Such interdisciplinary work, pulling together concepts and techniques across fields, is likely to have major implications for fundamental understanding of evolution, intelligence, and dynamical control, as well as drive novel kinds of therapeutics that leverage the innate behavioral competencies of living matter [47], [134].

Methods

GRN models and numerical simulation

This study employs ordinary differential equation (ODE) models to represent molecular pathways, with nodes representing pathway components and edges capturing their interactions. The continuous node states, encompassing variables like gene expression levels and protein concentrations, are interconnected through a system of ODEs, enabling the modeling of complex regulatory dynamics. ODE models are often available in the Systems Biology Markup Language (SBML), a standardized format that contains essential information about variables, parameters, equations, and model metadata in XML files.

To perform numerical simulations of ODE SBML models, we rely on the SBMLtoODEjax python library, a recent development that automates the parsing and conversion of SBML models into python models written entirely in JAX [135]. Taking advantage of JAX computing capabilities, SBMLtoODEjax enables efficient and parallel numerical solutions for gene expression levels and other node states by recursively invoking the generated python models to integrate the ODE equations with current gene expression levels. Additionally, we have developed a python library (<https://github.com/flowersteam/autodiscjax>) comprising additional modules and pipelines that facilitate interventions on the GRN models such as genome or drug interventions, as well as other perturbations such as noise, pushes, and walls that can be applied to the states and kinematic parameters of gene regulatory networks.

Given the model species initial state $y(t = 0)$, the desired rollout length $T(\text{secs})$ and step size ΔT , as well as the chosen intervention i and/or perturbation u , the model rollout iteratively 1) integrates the system of ODE-governed equations that specifies the rate of species changes $\frac{dy}{dt}$ using JAX odeint solver to update model species $y(t) \rightarrow y(t + \Delta T)$, 2) calls the model assignment rules to update kinematic parameters if needed, and 3) apply the intervention and/or perturbation function to update $(y(t + \Delta T), w(t + \Delta T), c)$ accordingly. In this paper we use $T = 2500\text{secs}$ and $\Delta T = 0.1$ (25 001 time points per rollout including t_0). The ODE solver uses an absolute tolerance of $1e^{-6}$ and relative tolerance of $1e^{-12}$, with maximum number of solver steps of 1000. For a step-by-step guide on utilizing these libraries within the proposed framework, we refer interested readers to our tutorial (<https://developmentalsystems.org/curious-exploration-of-grn-competencies/tutol.html>), which offers practical examples and detailed instructions.

Experimental setup

In our computational models, we are able to record the activities of all nodes during a model rollout. The observation space $\subset \mathbb{R}_{+}^{n \times \frac{T}{\Delta T}}$ is such that $o = (y(0), \dots, y(T))$ where $y(t)$ represents the n -dimensional vector of node states at each time step, with T being the total reaction time. The boundaries of the observation space are not known.

Regarding the exploration of problem spaces, namely the intervention space I and behavior space Z , we specify them as follows.

For the main experiments on biological networks, the intervention space $I \subset \mathbb{R}_{+}^{n}$ consists of initial node states sampled from the hyper-rectangle $[y_{0,min}, y_{0,max}]$ where $y_{0,min} = \frac{1}{r} \times y_{d,min}$ and $y_{0,max} = r \times y_{d,max}$ with $r = 20$ and $(y_{d,min}, y_{d,max})$ the minimum and maximum of each node of the model over the default time course simulation (with initial conditions provided in the SBML file and $T=25000$). On the other hand, the behavior space $Z \subset \mathbb{R}_{+}^{i,j}$ endpoint states $z = (y_i(T), y_j(T))$ where (i, j) corresponds to the target phenotype nodes. We ensure that most trajectories have reached stable states at $T = 2500$ (as elaborated in the next section) such that Z can be viewed as the space of reachable endpoints, whose boundaries are not known.

Database creation

Biological networks database

All the ODE models we use in this work are downloaded from the BioModels database [21], [22] in SBML format. From all models referenced on the website, we only consider the ones that are curated, that have at least 3 nodes, and that are handled by the SBMLtoODEjax simulator (as SBMLtoODEjax does not handle models with discrete events, custom functions or other specific cases as detailed in [135]). Note that the original models 37, 262, 263, 284, 454, 455, 459, 461 and 624 have been previously analyzed while clamping certain nodes at fixed values (as detailed in **table S1**). Here, we relaxed this condition for a more realistic simulation in which all of the nodes' concentration are free to vary.

To ensure the inclusion of models suitable for our analyses, we then applied specific filters to the collected models.

First, we simulated the default model rollout for each model to obtain the concentration profiles of the pathway components over a short time span ($T=10$ secs and $\Delta T = 0.1$). We discarded simulation results containing invalid values (NaN or negative concentrations) or those that took an excessive amount of time (>1 sec). While it is acceptable that a rollout sometimes returns NaN values (when there are no solutions given ODE tolerance options for specific initial conditions), we consider the model invalid if this occurs for the default initial conditions provided in the SBML file.

For the remaining models, we conducted further simulations with an extended time span ($T=2500$) and 50 random initial conditions uniformly sampled within the model's intervention space I (as defined before). Once again we discarded models whose batch simulations took an excessive amount of time (>15 secs). From the remaining models, we derived the resulting 50 trajectories for each node pair (i, j) and subjected them to additional filters to refine the database. We removed node pairs where either 1) [filter F1] a substantial proportion of trajectories (20%) exhibited invalid concentrations (NaN or negative) or unsettled behaviors ($\exists t \geq 2400$ such that $|y(t) - y(T)| \geq 0.02 \times |y(T) - y(0)|$) or periodic patterns ($\exists f > 0$ such that $|S(f)| \geq 40$ where $S = DFT([y(\frac{T}{2}), \dots, y(T)])$); or

[filter F2] the reached space in Z was too small $\left(\left(\max_{k=1 \dots 50} y^k(T) - \min_{k=1 \dots 5} y^k(T) \right) < 0.1 \right)$ to discard cases where “diversity” could result from floating point rounding errors; or [filter F3] the number of attractors was less than four $\left(\{y^k(T)\}_{k=1 \dots 50} \text{cover} \leq 4 \frac{\text{bins}}{a} 20 \times 20 \text{binningof}Z \right)$.

Upon completion of the filtering process, our final database comprised 30 models, consisting of a total of 432 systems, as detailed in Supplementary **Table S1**. These curated models and systems served as the foundation for our subsequent analyses and investigations into the navigation competencies of the molecular pathways.

Random networks database

Following the methodology proposed in [79], we aimed to create a database of synthetic networks with topologies similar to those of the biological networks, but with random regulatory rules instead of evolved ones. The objective was to compare the versatility and robustness competencies between biological and random networks, akin to the approach used for memory competencies in [79]. To achieve this, we initially generated 300 networks based on the transcriptional gene circuit model [117], ensuring that they had the same distribution of network size (number of nodes) and connectivity (nodes in-degree) as the biological network database (using fitted gaussian distributions). The kinematic parameters W , b , τ of these networks were randomized ($W \sim [-30, 30]^{n \times n}$, $B \sim [-10, 10]^n$, $\tau \sim [1, 15]$) where model step is defined as $y(t+1) = \frac{\Delta T}{\tau} \times \text{sigmoid}(Wy + B) + \left(1 - \frac{\Delta T}{\tau}\right) \times y$ and in-degree connectivity is enforced by setting some weights of W to zero. However, during the creation process, we observed that none of the generated networks met the criterion for exhibiting a sufficient number of steady states (criterion F3). This limitation arose from the inherent constraints imposed by the gene circuit model’s shape of ODE equations, limiting the diversity of possible dynamical behaviors. As our focus was on networks with a possible spectrum of steady states, akin to the biological network database, we decided not to pursue further analyses on these networks.

Instead, we selected the systems (models and pairs of nodes) that demonstrated the highest versatility (metric detailed below) from among all the generated systems that passed the filters F1 and F2. The selected networks’ versatility is presented in **Figure 7**, but for future research, it would be interesting to explore broader and more complex classes of equations to assess their potential for achieving higher behavioral diversity.

Sanity Check

We also tested our exploration pipeline on two simple models, namely BIOMD000000341 as described in the paper “A model of beta-cell mass, insulin, and glucose kinetics: pathways to diabetes” [140] and BIOMD000000454 as described in the “Example One” of the “Metabolic Control Analysis: Rereading Reder” paper [141]. For both these models, the ground truth goal states (attractors of the models) are already known, because easy to find analytically or numerically as described in the corresponding papers. As a sanity check, we validate that our exploration pipeline is able to re-discover those goal states and added those results in **Figure S4** of the supplementary for completeness. Note that a more complex variant of BIOMD000000454 is also included in our main biological networks database where we let the metabolite concentrations $y_1(t), \dots, y_5(t)$ evolve in time as described in the caption of **table S1**. We illustrate results of our curiosity-driven exploration method and of a random search method on this more complex variant in **Figure S4-c, S4-d**.

Curiosity-driven exploration

This section provides additional information about the internal models and hyperparameters of the intrinsically-motivated goal exploration process. The overall IMGEP pipeline is illustrated in **Figure 1-c**. To sample a goal, the IMGEP uses a uniform sampling strategy within the bounding hyper-rectangle of currently reached goals (scaled by a factor 1.3). Hence sampling bounds adapt to the discoveries and do not need to be predefined via expert knowledge. The volume of the hyper-rectangle is larger compared to the cloud of currently-reached goals, which incentivizes targeting unexplored areas outside of the cloud and promotes diversity in the exploration process. Then, to generate an intervention for achieving the sampled goal, the IMGEP selects the nearest previously reached goal in Z , identifies its associated intervention, and performs a local random step from that point $stepsize \sim N(0, 0.1 * [y_{0,max} - y_{0,min}])$ in the intervention space.

While our implementation choices for the IMGEP goal representation, goal generation, and goal-conditioned optimization are relatively straightforward, it is worth noting that alternative strategies could be considered for each of these components for more complex problems. The python library AutoDiscJax (<https://github.com/flowersteam/autodiscjax>) that accompanies this paper can be used to implement this and other IMGEP variants in JAX.

Robustness tests

We define 3 family of perturbations: 1) the noise perturbation $U_n(\sigma_n, p_n, V_y)$ which is parametrized by its standard-deviation (scaled proportionally to the extent of the observed trajectory y prior perturbation) and period (secs); 2) the push perturbation $U_p(m_p, n_p, V_y)$ parametrized by its magnitude (proportional to the extent of y) and number of occurrences; 3) the wall perturbation $U_w(l_w, n_w, V_y)$ parametrized by its length (proportional to the extent of y) and number, and where walls are generated in locations of the space that the GRN would “naturally” visit without the induced perturbation. Details about the implementation of walls are provided in **Supplementary Figure S3**.

To assess the robustness of the GRN systems in our database, we employ an evaluation procedure, as depicted in **Figure 1-d**. For each system (I, Z) in the database with its corresponding behavioral catalog H discovered using the curiosity-search algorithm, we perform the following steps. We first retrieve K representative trajectories out of the N discoveries, i.e. ones that cover well the reachable space. To do so, we randomly sample tuples of K discoveries (among N) 500 times, and select the one with the maximum diversity. One could test all trajectories with $K=N$ but here we use $K=N/10$ mainly for compute reasons, as we run the experimental campaign on all 432 systems. Next, we subject each of these K trajectories $\{y_k, k = 1.. K\}$ to $s=18$ different perturbation distributions, each representing various levels of difficulty: $(\sigma_n, p_n) \in \{(0.001, 5), (0.005, 5), (0.1, 5), (0.005, 10), (0.005, 5), (0.005, 1)\}$, $(m_p, n_p) \in \{(0.05, 1), (0.1, 1), (0.15, 1), (1, 0.1), (2, 0.1), (3, 0.1)\}$, $(l_w, n_w) \in \{(0.05, 1), (0.1, 1), (0.15, 1), (1, 0.1), (2, 0.1), (3, 0.1)\}$. In each perturbation distribution, we sample $r=3$ random perturbations, resulting in $P = s * r$ perturbations. For each perturbation in the set $\{u_p, p = 1.. P\}$, we re-run the trajectory starting from the same initial state i but with the sampled perturbation applied (i, u_p) , and observe the resulting outcome (o_p) and reached endpoint (z_p) .

At the end of this process, the behavioral catalog is augmented with the perturbed trajectories $H = \{(i_k, o_k, z_k, \{(u_p, o_p, z_p), p = 1.. P\}), k = 1.. K\}$.

Evaluation Metrics

Diversity measure

Diversity is measured by the area that explored observations cover in behavior space Z . Each single exploration results in a new point in this space, such that diversity measures how much area the algorithms explored in those spaces.

In general, existing approaches in the NS, QD and IMGEP literature use binning-based metrics [90], [91], [136] or distance-based metric from ecology [137] to quantify the diversity of a set of explored instances. However, those metrics are sensitive to the binning strategy, or fail to discriminate between qualitatively significantly different explorations [105]. Another approach, called the threshold coverage, measures diversity as the volume of the union of the set of hyperballs of radius ϵ that have for centers the observed effects $\{z \in Z\}$. This diversity measure, while difficult to compute in high-dimensional spaces, avoids the pitfalls of bin-based and distance-based metrics and is easily computable in 2-dimensional spaces [105].

Threshold coverage quantifies the area of the space that has been reached at a given precision ϵ (the threshold), and is what we used in **Figure 3** to compare random search and curiosity-driven exploration strategies.

Sensitivity measure

In general, existing approaches in systems biology and evolutionary genetics measure sensitivity (opposite of robustness) in a relative manner with respect to 1) a functionality [93] or phenotypic trait [94] of interest, 2) specific perturbations (environmental or genetic changes), and 3) a measure of the degree of variation. Here, we adopt a similar metric where 1) the phenotypic trait of interest is defined as a goal state $z \in Z$ discovered by curiosity search, 2) the set of perturbation $\{u_p\}$ is defined in previous section and conditioned on the GRN goal-reaching trajectory $i \rightarrow z$, and 3) variation is measured as the Euclidean distance in behavior space, normalized by the extent of the trajectory prior perturbation in Z .

This distance-based sensitivity measure proves straightforward as we explicitly use “spaces” to observe and analyze behaviors. The results of this sensitivity analysis are presented in **Figure 5**.

Versatility-Robustness measure

In this study, we introduce the terms “diversity” and “versatility” to characterize the competencies of the exploration agent (IMGEP) and the gene regulatory network agent (GRN), respectively. Diversity refers to the ability of the IMGEP agent to reveal a wide range of behaviors in the GRN, while versatility refers to the capability of the GRN agent to reach diverse goal states. The GRN versatility is unknown, and can only be approximated via proxy metric. Here, we consider that the diversity of the IMGEP (measured with the threshold coverage metric) is a good approximation of the versatility of a given GRN, as the IMGEP was shown to efficiently drive the GRN into diverse possible goal states. In **Figure 7-a**, we employ this diversity metric to categorize the versatility of surveyed networks based on the class of organism they belong to. For the random networks, as they ne^2 all have less or equal than 4 attractors, the versatility remains below $0.026 = 4 \times \frac{\pi\epsilon^2}{(1+2\epsilon)^2}$.

Figure 7-b, we introduce the versatility-robustness metric, which conditions the diversity metric on a sensitivity threshold. Only goal states with sensitivity to perturbations below this threshold are considered when computing the reached area of the space. A high versatility-robustness score indicates that diverse goal states are achieved with a high level of precision.

Experiments on the RKIP-ERK signaling pathway

This section details the additional experiments conducted on the RKIP-ERK signaling pathway [99]. We refer to the accompanying notebook tutorial for reproducing these experiments: <https://developmentalsystems.org/curious-exploration-of-grn-competencies/tutol.html>.

For **Figure 4**, clustering in behavior space was performed using the HDBSCAN algorithm [106] with hyperparameters set as `min_cluster_size=10` and `cluster_selection_epsilon=0.1`. Points in the 10-dimensional intervention space are visualized by applying a TSNE 2dimensional reduction. To visualize the clusters in behavior space (and corresponding clusters in intervention space), we fitted polygons on the cluster points using shapely library `unary_union`, `dilatation`, and `erosion` operations [138].

In **Figure 6**, we generated trajectory-based energy landscapes following the method proposed in [113]. Energy landscapes provide an intuitive way to understand how a system with multiple steady states behave, by picturing it as a ball rolling downhill towards low-energy valleys (steady states). Given a set of trajectories in behavior space Z , we constructed a probability distribution (P) of system states and converted it into a pseudopotential energy surface ($U = -\ln(P)$). This energy surface was smoothed using cubic spline interpolation and visualized using Plotly 3D surface plots. **Figure 6-a, 6-b, and 6-c** differed by the input set of trajectories used for generating the landscape: a) employed the set of trajectories discovered by random search, b) used the set of trajectories discovered by curiosity search, and c) utilized the set of trajectories generated by robustness tests.

In **Figure 8**, the “healthy” and “disease” clusters were the same as in **Figure 4** and visualized similarly. We displayed trajectories with the lowest sensitivity (averaged over all $P = 3 \times 18$ perturbations). The stimuli-based intervention shown in **Figure 8-b** was found using a simple random search procedure. First, we defined an arbitrary target node and a stepwise node-activation function, clamping MKEPP values to desired values $x = [y_{MEKPP}^{(1)}, \dots, y_{MEKPP}^{(10)}]$ every 10 seconds for 100 seconds. Then, we randomly sampled x within a range of values near the MKEPP current steady states (endpoints from the 6 “disease” trajectories, assuming that the drug intervention cannot drastically remodel those values). For each candidate x , we ran new trajectories starting from the disease states and applying the intervention x under a distribution of noise, push, and wall perturbations. Finally, we selected the intervention x that most successfully brought ERK-RKIP levels back to the target setpoint (centroid of the healthy region). The resulting intervention (shown in **Figure 8-b**) succeeds to robustly reset all 6 disease state points despite perturbations, as shown in **Figure 8-c**. We refer to the notebook for reproducing the experiments.

Experiments on synthetic gene networks

This section details the additional experiments conducted on the synthetic gene networks (**Figure 9**). We refer to the second accompanying tutorial for the full codebase: <https://developmentalsystems.org/curious-exploration-of-grn-competencies/tuto2.html>.

In these experiments, we consider the target application of gene circuit engineering followed in [72], where parameters of a gene circuit model are optimized to produce target oscillator patterns. The gene circuit model employed in [72] is the same than the one used for the random networks database (Eq 1), with $\tau = 1$. Hence the *intervention space* is a $n^2 + 2n$ dimensional space defined as $I = [y_{t=0,min}, y_{t=0,max}] \oplus [W_{min}, W_{max}] \oplus [B_{min}, B_{max}]$, with $y_{0,min} = 0, y_{0,max} = 1, W_{min} = -30, W_{max} = 30, B_{min} = -10, B_{max} = 10$. Here we consider networks of $n=3$ nodes, with the first node being the target phenotype node. Thus, what we seek here is kinematic parameters (W, B) and

initial concentrations y_0 that would produce a periodic pattern $y = [y_{n=0}(0), \dots, y_{n=0}(T)]$ with target amplitude A , frequency ω and offset b . Here, the target (A, ω, b) are sample randomly with $A \sim U([0.1, 0.5])$, $b \sim U([A, 1-A])$, $\omega \sim \text{Beta}(a = 2, \beta = 8)$.

We then compare three alternative exploration strategies: 1) curiosity search, 2) random search and 3) gradient descent, i.e. pure optimization-driven search as proposed in [72], all given the same experimental budget $N = 5000$.

For curiosity search, the behavior space Z is defined as the image space of the discrete Fourier transform of the 1d-signal y , where distance in the space measures average difference in spectral amplitude. The IMGEP algorithm is then the same that the one previously used, as detailed in **Figure 1-c**, but operating within the novel problem spaces (I, Z) .

For random search, interventions are sample uniformly $(i_1, \dots, i_N) \sim U(I)$.

For gradient descent, we follow the procedure proposed in [72]. We define a loss function which, for a set of parameters $i = (W, B, y_0)$, measures the mean square error between the phenotype node activation levels y and the target oscillation represented as a cosine wave with the desired (A, ω, b) : $L = \sum_t (y(t) - (A \cos(2\pi\omega t) + b))^2$. We then sample a random parameter $i \sim U(I)$ and use Adam optimizer with $l_r = 10^{-3}$, $b1 = 0.02$, $b3 = 0.001$, $\epsilon = 10^{-8}$ for $N = 5000$ optimization steps (same number of model rollouts allowed than for curiosity search and random search).

In addition, we use gradient descent for *local* refinement of the best discoveries made by the other exploration strategies (curiosity search and random search), this time with a limited budget of $N = 100$ optimization steps.

Visualizations in **Figure 9** show: (a-b) the oscillators discovered by random search and curiosity search (gradient descent did not find any oscillator in this example) in the (A, b) space, (c) the corresponding diversity (using this time a binning-based space coverage measure with 20^3 bins as the space is 3-dimensional), (d) the evolution of the training loss L throughout the $N=5000$ trials for the three exploration strategies, (e-f-g) the corresponding best discoveries (for which L is minimal) for the three exploration strategies, and (h-i) the local training loss and resulting finetuning of the best discoveries with gradient descent.

Statistics

Statistical analyses were performed using custom code in python 3.9 (relevant libraries include jax 0.4.8 and scipy 1.10 and plotly 5.16.1). Welch's two-sample t-tests reported in **Figure 3** are two-tailed.

Data Availability

Source code is available on GitHub at <https://github.com/flowersteam/curious-exploration-of-grn-competencies>. It contains experimental data and an executable notebook version of the paper to reproduce all paper figures, as well as additional step-by-step tutorials to reproduce results from scratch for **Figures 4**, **6** and **8** (tutorial 1) and **Figure 9** (tutorial 2), as well as the codebase to reproduce the whole experimental campaign. All our codebase is open-source under MIT License.

Acknowledgements

We thank Patrick Erickson and Randall Jordan Ellis for review and discussion, as well as Tom Cirrito, Wesley Clawson and Santosh Manicka for useful discussions. We also thank Alexander Mordvinsteve for providing the executable paper template, as well as Julia Poirier for assistance with the manuscript.

The authors acknowledge support from the biotechnology company Poietis and the French National Association of Research and Technology (ANRT), as well as from the French National Research Agency (ANR, DeepCuriosity AI chair project). M.L. gratefully acknowledges funding support from Astonishing Labs, and from the Templeton World Charity Foundation via grant TWCF0606.

This work also benefited from the use of the Jean Zay supercomputer associated with the Genci grant A0151011996.

Author Contributions

M.E, M.L, P.Y.O and C.M.F conceived the project together. M.E. wrote the code and performed the computational experiments and data analysis. M.L, P.Y.O and C.M.F supervised the project. M.E, M.L, P.Y.O and C.M.F wrote the paper.

Competing Interests Statement

M.L.'s lab receives funding in the form of a sponsored research agreement from, and consults for, Astonishing Labs, which has interest in therapeutic applications of GRN learning behaviors.

Supplementary

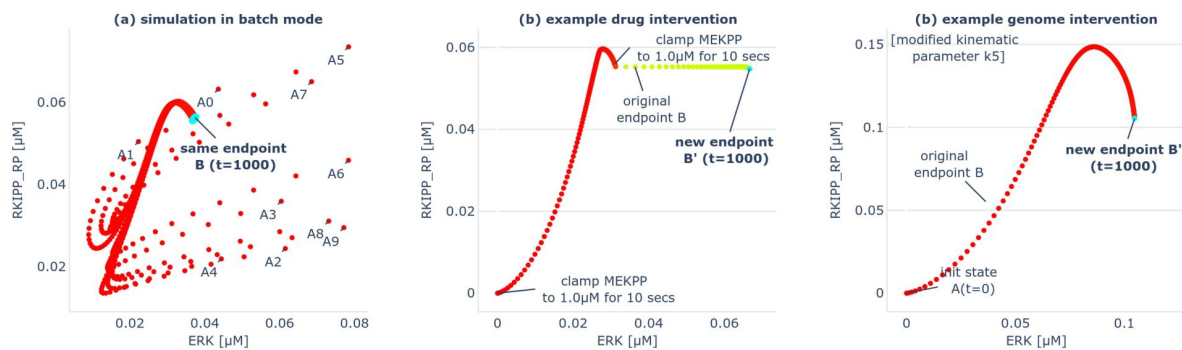


Figure S1

Examples of interventions that can be implemented within the accompanying AutodiscJax software.

All those examples can be reproduced in the accompanying tutorial 1. (a) Numerical simulations with interventions can be performed in parallel by vectorizing simulations over different intervention parameters, simply using the `jax vmap` operator. This offers a convenient (and fast) way to test several interventions in the biological network, as shown here for testing the network under various initial conditions in batch mode. Examples of other possible “drug” or “genome” interventions that can be implemented in the accompanying software, as well as the possibility to perform interventions (or perturbations) in parallel using batched computations. In this example, despite the numerous interventions, the GRN trajectories still converge to the same endpoint B. (b) Example intervention where species amounts are clamped to specific values. In this example the node MEKPP is clamped to $2.5\mu M$ for 10 seconds at $t=0$ and then to $1\mu M$ for 10 additional seconds at $t=400$. In this example, after the first clamping the GRN trajectory still follows a similar S-shape curve and arrives close to the original endpoint B but after the second clamping, ERK expression levels are shifted to a higher steady state B'. (c) Example intervention where the numerical value of one kinematic parameter of the model (k_5) is changed from 0.0315 to 0.1. In this example we can see that changing the parameter k_5 shifts the trajectory end point quite significantly, but qualitatively the trajectory seems to preserve a similar S-shape.

BioModel ID reference	number of nodes	organism class	systems (observed nodes pairs)
BIOMD0000000524	16	Homo sapiens	('PrER', 'mGFP'), ('DISC', 'tBid'), ('p18inactive', 'PrER'), ('p18inactive', 'tBid'), ('p18inactive', 'mGFP'), ('p18inactive', 'PrNES'), ('DISC', 'mCherry'), ('PrNES', 'PrER'), ('Bid', 'mCherry'), ('tBid', 'PrNES'), ('p18inactive', 'PrER_mGFP'), ('PrNES', 'mGFP'), ('mCherry', 'PrER_mGFP'), ('PrER_mGFP', 'PrER'), ('FADD', 'PrNES'), ('tBid', 'mGFP'), ('DISC', 'PrER'), ('FADD', 'mGFP'), ('mCherry', 'mGFP'), ('FADD', 'mCherry'), ('Bid', 'PrER'), ('FADD', 'tBid'), ('DISC', 'p18inactive'), ('Bid', 'PrNES'), ('p18inactive', 'Bid'), ('tBid', 'PrER_mGFP'), ('FADD', 'PrER'), ('tBid', 'mCherry'), ('PrNES', 'PrER_mGFP'), ('FADD', 'p18inactive'), ('FADD', 'Bid'), ('Bid', 'mGFP'), ('mCherry', 'PrER'), ('PrNES', 'mCherry'), ('DISC', 'Bid'), ('DISC', 'PrNES'), ('Bid', 'tBid'), ('DISC', 'PrER_mGFP'), ('FADD', 'DISC'), ('DISC', 'mGFP'), ('p18inactive', 'mCherry'), ('PrER_mGFP', 'mGFP'), ('FADD', 'PrER_mGFP'), ('tBid', 'PrER')
BIOMD0000000050	14	n/a	('Gly', 'Fru'), ('FA', 'Mel'), ('FA', 'MG'), ('FA', 'Fru'), ('Mel', 'Fru'), ('Gly', 'Mel'), ('FA', 'Glu'), ('Gly', 'AA'), ('AA', 'Fru'), ('Glu', 'Mel'), ('MG', 'Fru'), ('AA', 'Glu'), ('FA', 'AA'), ('Glu', 'Fru'), ('Gly', 'MG'), ('Glu', 'MG'), ('AA', 'MG'), ('AA', 'Mel'), ('Mel', 'MG'), ('Gly', 'FA')
BIOMD0000000647	11	n/a	('RKIPP', 'MEKPP_ERK'), ('RKIPP', 'RKIPP_RP'), ('Raf1_RKIP', 'ERK'), ('ERK', 'RP'), ('ERK', 'MEKPP'), ('Raf1', 'MEKPP'), ('MEKPP', 'RKIPP_RP'), ('MEKPP', 'RP'), ('RKIP', 'RP'), ('ERKPP', 'RKIPP'), ('Raf1', 'RKIPP'), ('ERKPP', 'RP'), ('Raf1_RKIP_ERKPP', 'ERK'), ('Raf1_RKIP', 'MEKPP'), ('Raf1', 'ERKPP'), ('ERKPP', 'MEKPP_ERK'), ('ERK', 'MEKPP_ERK'), ('RKIP', 'RKIPP_RP'), ('Raf1', 'ERK'), ('ERKPP', 'Raf1_RKIP_ERKPP'), ('RKIP', 'ERKPP'), ('Raf1_RKIP_ERKPP', 'RKIPP_RP'), ('Raf1_RKIP_ERKPP', 'MEKPP'), ('Raf1_RKIP', 'ERKPP'), ('RP', 'RKIPP_RP'), ('MEKPP_ERK', 'RKIPP_RP'), ('ERK', 'RKIPP_RP'), ('ERKPP', 'ERK'), ('Raf1_RKIP_ERKPP', 'RP'), ('Raf1', 'Raf1_RKIP_ERKPP'), ('Raf1', 'Raf1_RKIP'), ('MEKPP', 'MEKPP_ERK'), ('Raf1_RKIP', 'RKIPP'), ('Raf1_RKIP_ERKPP', 'RKIPP'), ('Raf1_RKIP_ERKPP', 'MEKPP_ERK'), ('MEKPP_ERK', 'RP'), ('RKIP', 'MEKPP'), ('Raf1', 'RKIP'), ('RKIPP', 'RP'), ('RKIP', 'MEKPP_ERK'), ('Raf1_RKIP', 'Raf1_RKIP_ERKPP'), ('Raf1', 'RP'), ('Raf1_RKIP', 'RKIPP_RP'), ('Raf1', 'RKIPP_RP'), ('Raf1_RKIP', 'MEKPP_ERK'), ('ERKPP', 'MEKPP'), ('RKIP', 'RKIPP'), ('Raf1', 'MEKPP_ERK'), ('RKIP', 'ERK'), ('RKIPP', 'MEKPP'), ('ERKPP', 'RKIPP_RP'), ('RKIP', 'Raf1_RKIP'), ('ERK', 'RKIPP'), ('Raf1_RKIP', 'RP'), ('RKIP', 'Raf1_RKIP_ERKPP')
BIOMD00000000520	3	Rodents	('N0', 'N1'), ('N1', 'N2'), ('N0', 'N2')
BIOMD00000000523	16	Homo sapiens	('tBid', 'mCherry'), ('PrER', 'mGFP'), ('PrNES', 'mGFP'), ('DISC', 'PrNES'), ('DISC', 'tBid'), ('FADD', 'tBid'), ('p18inactive', 'tBid'), ('FADD', 'mCherry'), ('mCherry', 'mGFP'), ('p18inactive', 'mCherry'), ('tBid', 'PrNES'), ('tBid', 'mGFP'), ('FADD', 'DISC'), ('p18inactive', 'mGFP'), ('mCherry', 'PrER'), ('p18inactive', 'PrER'), ('p18inactive', 'PrER'), ('DISC', 'p18inactive'), ('DISC', 'PrER'), ('FADD', 'PrER'), ('DISC', 'mGFP'), ('PrNES', 'mCherry'), ('p18inactive', 'PrNES'), ('FADD', 'PrNES'), ('tBid', 'PrER'), ('FADD', 'mGFP'), ('DISC', 'mCherry'), ('PrNES', 'PrER')
BIOMD00000000454	8	Generic	('x3', 'y3'), ('y4', 'y2'), ('y1', 'y5'), ('x3', 'y2'), ('y4', 'y5'), ('x2', 'y4'), ('x1', 'y3'), ('x1', 'y5'), ('y2', 'y3'), ('y1', 'y3'), ('x1', 'y4'), ('x1', 'x3'), ('y4', 'y3'), ('y1', 'y2'), ('y1', 'y4'), ('y1', 'x3'), ('y5', 'y2'), ('y1', 'x1'), ('x2', 'x1'), ('x1', 'y2'), ('x2', 'x3'), ('x3', 'y4'), ('x3', 'y5'), ('y1', 'x2'), ('x2', 'y2'), ('x2', 'y5'), ('x2', 'y3'), ('y5', 'y3')
BIOMD00000000069	10	Homo sapiens	('srco', 'srca'), ('srca', 'srca'), ('srco', 'srcc'), ('srca', 'Cbp_P'), ('srcc', 'Cbp_P_CSK'), ('srca', 'Cbp_P'), ('srca', 'PTP'), ('Cbp_P', 'PTP_pY789'), ('srcc', 'PTP_pY789'), ('CSK_cytoplasm', 'PTP'), ('srca', 'PTP_pY789'), ('srca', 'PTP_pY789'), ('srcc', 'CSK_cytoplasm'), ('CSK_cytoplasm', 'Cbp_P_CSK'), ('srca', 'CSK_cytoplasm'), ('srca', 'PTP'), ('srcc', 'Cbp_P'), ('srco', 'Cbp_P_CSK'), ('CSK_cytoplasm', 'PTP_pY789'), ('srcc', 'PTP'), ('srca', 'Cbp_P_CSK'), ('srca', 'CSK_cytoplasm'), ('Cbp_P', 'Cbp_P_CSK'), ('Cbp_P', 'PTP'), ('srca', 'srco'), ('srco', 'PTP'), ('CSK_cytoplasm', 'Cbp_P'), ('srca', 'srcc'), ('srca', 'Cbp_P_CSK'), ('Cbp_P_CSK', 'PTP_pY789'), ('PTP', 'PTP_pY789'), ('Cbp_P_CSK', 'PTP'), ('srca', 'srcc'), ('srco', 'Cbp_P'), ('srco', 'PTP_pY789'), ('srco', 'CSK_cytoplasm')
BIOMD00000000455	9	Generic	('x2', 'y5'), ('y5', 'y2'), ('x2', 'y2'), ('y1', 'y3'), ('x3', 'y3'), ('y1', 'x2'), ('x2', 'x1'), ('x2', 'y4'), ('x3', 'y2'), ('x2', 'y3'), ('y1', 'y4'), ('y5', 'y3'), ('y1', 'y5'), ('x1', 'x3'), ('x1', 'y4'), ('y4', 'y5'), ('x1', 'y3'), ('y1', 'x3'), ('y4', 'y3'), ('x3', 'y4'), ('x3', 'y5'), ('x1', 'y5'), ('y1', 'x1'), ('y1', 'y2'), ('y2', 'y3'), ('y4', 'y2'), ('x2', 'x3'), ('x1', 'y2')

Table S1

List of biological networks from Biomodels used in this study.

The resulting database includes 30 biological networks (one row per network) and a total of 432 systems, which is defined as a (GRN model, behavior space (Z)) tuple and where the pairs of observed nodes (used as behavior spaces) per network are given in the last column. Please note that we let the following chemical species evolve in time, despite being labeled as “constant” or “boundary” species in the original SBML models: model #262: ['EGF', 'pro_EGFR'], model #284: ['X', 'Y', 'Z'], model #37: ['Gluc'], model #263: ['NGF', 'pro_TrkA'], model #454: ['y1', 'y2', 'y3', 'y4', 'y5'], model #455: ['y1', 'y2', 'y3', 'y4', 'y5', 'y6'], model #459: ['IPTG'], model #461: ['IPTG'], model #624: ['X1'].

BIOMD0000000526	16	Homo sapiens	('FADD', 'PrER_mGFP'), ('PrNES', 'mCherry'), ('p18inactive', 'PrER_mGFP'), ('DISC', 'p18inactive'), ('FADD', 'p18inactive'), ('p18inactive', 'mCherry'), ('DISC', 'mGFP'), ('FADD', 'mGFP'), ('p18inactive', 'tBid'), ('p18inactive', 'mGFP'), ('FADD', 'DISC'), ('tBid', 'PrER'), ('DISC', 'tBid'), ('p18inactive', 'PrER'), ('p18inactive', 'PrNES'), ('PrNES', 'PrER_mGFP'), ('PrNES', 'PrER'), ('FADD', 'PrNES'), ('DISC', 'PrER'), ('DISC', 'mCherry'), ('FADD', 'tBid'), ('PrER_mGFP', 'PrER'), ('PrNES', 'mGFP'), ('FADD', 'PrER'), ('mCherry', 'mGFP'), ('DISC', 'PrNES'), ('tBid', 'PrNES'), ('mCherry', 'PrER'), ('PrER_mGFP', 'mGFP'), ('tBid', 'mCherry'), ('FADD', 'mCherry'), ('tBid', 'mGFP'), ('mCherry', 'PrER_mGFP'), ('DISC', 'PrER_mGFP'), ('PrER', 'mGFP'), ('tBid', 'PrER_mGFP')
BIOMD0000000284	9	Bacteria	('D', 'E'), ('A', 'D'), ('X', 'C'), ('B', 'E'), ('C', 'Z'), ('X', 'E'), ('A', 'Z'), ('X', 'Y'), ('C', 'D'), ('Y', 'Z'), ('E', 'F'), ('F', 'Z'), ('D', 'F'), ('B', 'F'), ('C', 'F'), ('X', 'Z'), ('A', 'Y'), ('X', 'D'), ('A', 'F'), ('B', 'Z'), ('X', 'F'), ('A', 'B'), ('A', 'C'), ('B', 'D'), ('Y', 'E'), ('B', 'C'), ('B', 'Y'), ('C', 'Y'), ('Y', 'F'), ('E', 'Z'), ('X', 'B'), ('D', 'Z'), ('D', 'Y'), ('A', 'E'), ('C', 'E'), ('X', 'A')
BIOMD0000000084	8	n/a	('Rin', 'x2'), ('x1', 'x3'), ('x2', 'x3'), ('x1', 'x2'), ('Rin', 'x3'), ('Rin', 'x1')
BIOMD0000000052	11	n/a	('Triose', 'lys_R'), ('Cn', 'lys_R'), ('Acetic_acid', 'lys_R'), ('lys_R', 'Melanoidin'), ('Triose', 'Melanoidin'), ('Cn', 'Melanoidin'), ('Formic_acid', 'lys_R'), ('C5', 'lys_R'), ('C5', 'Formic_acid'), ('Triose', 'Acetic_acid'), ('C5', 'Acetic_acid'), ('Acetic_acid', 'Melanoidin'), ('C5', 'Cn'), ('C5', 'Melanoidin'), ('Cn', 'Acetic_acid'), ('C5', 'Triose'), ('Formic_acid', 'Triose'), ('Formic_acid', 'Cn'), ('Triose', 'Cn'), ('Formic_acid', 'Melanoidin'), ('Formic_acid', 'Acetic_acid')
BIOMD0000000271	6	Rodents	('EpoR', 'dEpo'), ('Epo_EpoRi', 'dEpo'), ('Epo_EpoR', 'dEpo'), ('Epo_EpoRi', 'dEpo'), ('EpoR', 'dEpo'), ('Epo', 'dEpo'), ('Epo', 'dEpo'), ('Epo_EpoR', 'dEpo'), ('dEpo', 'dEpo')
BIOMD0000000461	4	Bacteria	('lacZ', 'x'), ('IPTG', 'sigb'), ('sigb', 'x'), ('sigb', 'lacZ'), ('IPTG', 'x'), ('IPTG', 'lacZ')
BIOMD00000000525	16	Homo sapiens	('p18inactive', 'PrNES'), ('p18inactive', 'PrER'), ('tBid', 'PrNES'), ('DISC', 'PrER'), ('mCherry', 'PrER'), ('PrNES', 'mCherry'), ('DISC', 'tBid'), ('FADD', 'mCherry'), ('p18inactive', 'tBid'), ('FADD', 'PrER'), ('FADD', 'p18inactive'), ('tBid', 'mGFP'), ('DISC', 'p18inactive'), ('mCherry', 'mGFP'), ('DISC', 'PrNES'), ('PrER', 'mGFP'), ('DISC', 'mGFP'), ('FADD', 'DISC'), ('tBid', 'PrER'), ('p18inactive', 'mCherry'), ('tBid', 'mCherry'), ('PrNES', 'mGFP'), ('PrNES', 'PrER'), ('FADD', 'PrNES'), ('DISC', 'mCherry'), ('p18inactive', 'mGFP'), ('FADD', 'tBid'), ('FADD', 'mGFP')
BIOMD00000000521	4	Homo sapiens	('P', 'Q')
BIOMD00000000010	8	Amphibians	('MKK_P', 'MAPK'), ('MKK', 'MAPK_PP'), ('MKK', 'MKK_PP'), ('MKK_P', 'MKK_PP'), ('MKK', 'MAPK'), ('MKK_PP', 'MAPK_P'), ('MAPK', 'MAPK_PP'), ('MKK_PP', 'MAPK'), ('MKK', 'MAPK_P'), ('MKK_PP', 'MAPK_PP'), ('MAPK_P', 'MAPK_PP'), ('MKK_P', 'MAPK_PP'), ('MKK', 'MKK_P'), ('MKK_P', 'MAPK_P')
BIOMD00000000029	4	Amphibians	('M', 'MpT'), ('M', 'MpY'), ('MpY', 'MpT')
BIOMD00000000197	5	n/a	('x1', 'x2'), ('x1', 'x4'), ('x1', 'x3'), ('x1', 'x5'), ('x5', 'x4'), ('x3', 'x5'), ('x5', 'x2')
BIOMD00000000272	6	Rodents	('SAV_EpoR', 'SAV_EpoRi'), ('EpoR', 'SAV_EpoR'), ('EpoR', 'SAV_EpoRi')
BIOMD00000000167	7	Generic	('Pstat_nuc', 'stat_nuc'), ('stat_nuc', 'species_test'), ('Pstat_nuc', 'species_test')
BIOMD00000000262	11	Rodents	('Akt', 'S6')
BIOMD00000000240	6	Bacteria	('DegU', 'mDegU'), ('Dim', 'mDegU'), ('DegUP', 'mDegU')
BIOMD00000000037	12	Slime Mold	('Xi', 'Ya'), ('Xi', 'Pi'), ('Ya', 'Pi')
BIOMD00000000263	11	Rodents	('Akt', 'pro_TrkA'), ('S6', 'pro_TrkA'), ('Akt', 'S6')
BIOMD00000000641	5	Homo sapiens	('CellCact', 'Effectoract')
BIOMD00000000413	5	Plants	('TIR1', 'auxinTIR1'), ('auxinTIR1', 'VENUS'), ('TIR1', 'VENUS')
BIOMD00000000624	7	Homo sapiens	('APAP', 'APAPconj_Glu')
BIOMD00000000945	5	Homo sapiens	('L_m', 'H_m')
BIOMD00000000459	4	Bacteria	('IPTG', 'sigb')

Table S1 (continued)




Figure S2


Additional results complementing Figure 8 of the main paper.

This figure shows the resulting trajectories after applying the discovered stimuli-based intervention (shown in Figure 8-b) to the example RKIP-ERK signaling pathway [99] for the 6 “disease” trajectories originally discovered in the behavioral catalog (shown in Figure 8-a). (a) For each trajectory (one per row), we see that the intervention successfully re-sets the disease setpoint (startpoint of the trajectory shown in red in the orange region) to a healthy set-point (endpoint of the trajectory shown in cyan in the green region). (b-c) Similar results are achieved despite adding push perturbations (b) or wall perturbations (c) in addition to the stimuli-based intervention.

(a) Elastic collision

$$v(p) \leftarrow v(p) \text{ if } p \notin \text{Wall}$$

$$[-v_{\perp} + v_{\parallel}] \text{ otherwise}$$



(b) Continuous force-field

$$\alpha = -A \exp\left(-\frac{d^2}{2\sigma^2}\right)^P$$

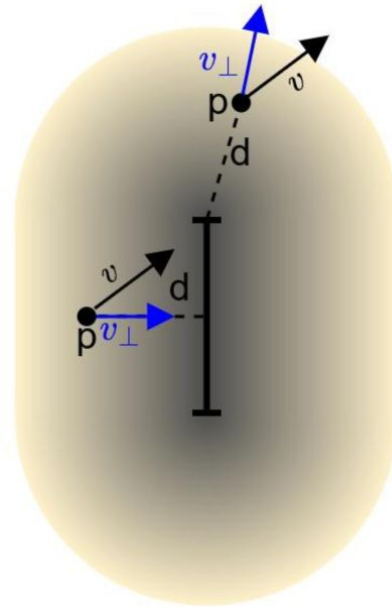


Figure S3

Wall implementation.

Walls are implemented within the 2D space spanned by the 2 observed nodes. Within that space, we can interpret the node activation levels $y(0), \dots, y(t)$ as the trajectory of a particle moving. In order to simulate the interaction with “walls” in that space, several implementations could be envisaged. Within the accompanying software AutoDiscJax we propose two possible variants: perfectly elastic collision (equivalent to a discontinuous force field) and some continuous force field variant. The second variant (continuous force field) is employed for the main results of this paper. (a) For the first variant, we consider a perfectly elastic collision without loss of kinetic energy. In that case, when the trajectory is touching the wall at position p with speed $v = v_{\perp} + v_{\parallel}$ we deviate the trajectory in such a way that is “bouncing” against the wall such that v_{\parallel} is unchanged and $v_{\perp} \leftarrow -v_{\perp}$. To implement it, we simply check whether the segment $[y(t), y(t + \Delta t)]$ intersect the wall at each time step. It does, we compute the intersection point p and time t_1 , and set the activation level $y(t + \Delta t)$ to $p + (\Delta t - t_1) \cdot [v_{\perp} + v_{\parallel}]$. (b) For the second variant, we implement walls as energy barriers acting as a new force field in the environment, constraining the GRN traversal of the space. This time, instead of having a discontinuous effect on the perpendicular speed v_{\perp} we define a wall force $f_{\perp} = \pm \alpha v_{\perp}$ (+ if v_{\perp} is going toward wall, - otherwise) and use it to update the perpendicular component of the trajectory speed as $v_{\perp} \leftarrow v_{\perp} + f_{\perp} \cdot \Delta T$. Here $\alpha \in [0, -2]$ and is calculated as a function of the distance between the point and the wall. As illustrated in the figure, this basically results in a stadium-shaped force field around the wall.

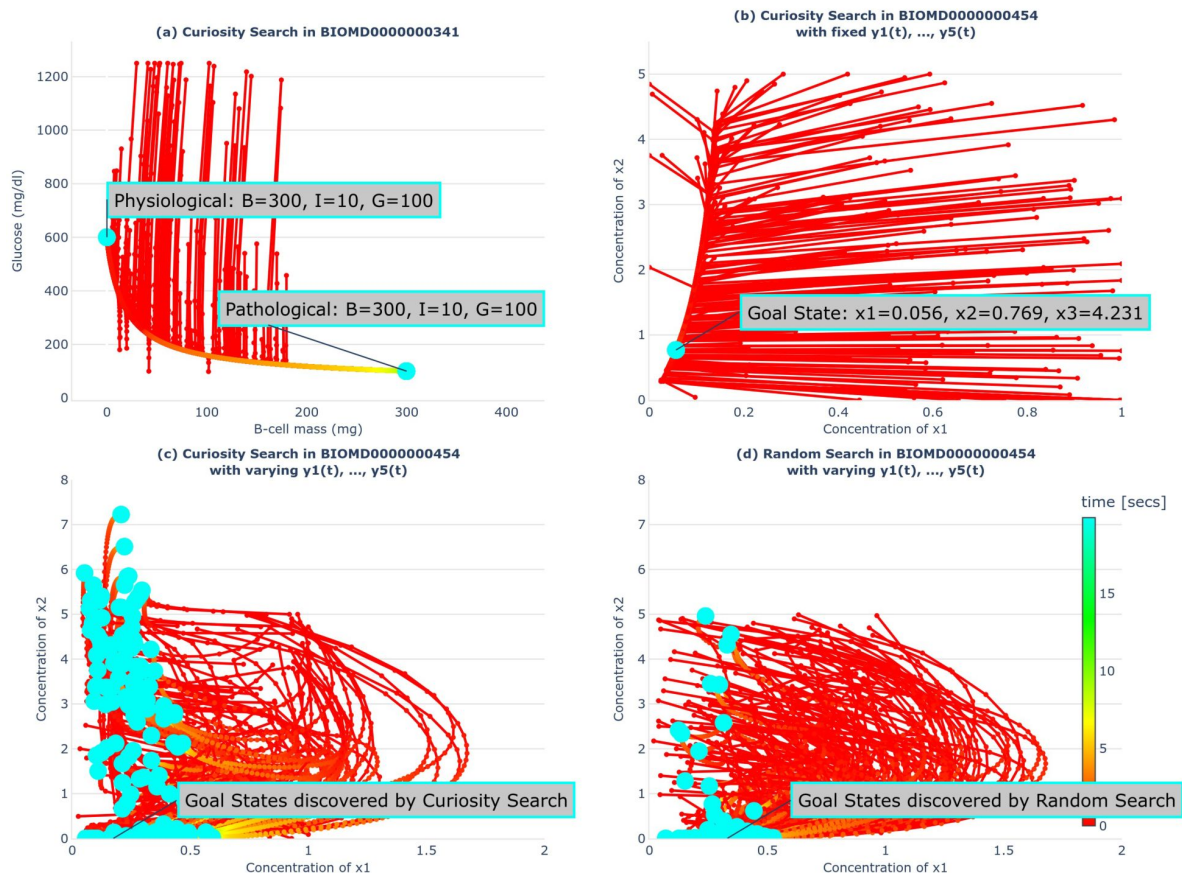


Figure S4

Sanity Check: Discoveries made by our exploration pipeline on simple systems for which the ground truth goal states are known (a-b) as well as for more complex variants of these systems (c-d).

Time-course evolution of the discovered trajectories are shown for the nodes given in x and y axis. Trajectories start from the initial conditions (red points) sampled by the exploration method and end in the discovered steady states (cyan points). (a) Discoveries made by our curiosity search method for BIOMD0000000341 which models the dynamics of beta-cell mass, insulin and glucose dynamics as described in Topp et al.'s paper [140]. This simple model is known to have two attractors, representing physiological ($B=300, I=10, G=100$) and pathological ($B=0, I=0, G=600$) steady states, as described in the paper. We show here that our exploration pipeline is able to re-discover those two steady states. (b) Discoveries made by our curiosity search method for BIOMD0000000454 with the same assumptions as described in the "Example One" from the "Metabolic Control Analysis: Rereading Reder" paper [141], which is a small metabolic pathway whose (single) steady state can be derived analytically as shown in the paper. As expected our exploration method is again able to re-discover this steady state ($x_1=0.056, x_2=0.769, x_3=4.231$). (c-d) Discoveries made by our curiosity search method and a random search method for BIOMD0000000454 but this time relaxing some of the assumptions made in "Example One" of [141]. We notably relax the assumption that metabolite concentrations $y_1(t), \dots, y_5(t)$ are fixed to certain values, and instead let these quantities evolve according to the provided chemical reactions. In that more complex case, the system exhibits a spectrum of possible steady states that are not analytically easy to find and unknown a priori. Here we see that the curiosity search method (c) is able to uncover the distribution of possible steady states more efficiently than a random search method (d).

References

- [1] Sanz-Ezquerro J. J., Munsterberg A. E., Stricker S. (2017) **Editorial: Signaling Pathways in Embryonic Development** *Front. Cell Dev. Biol.* **5** <https://doi.org/10.3389/fcell.2017.00076>
- [2] Padilla-Longoria Enrique Balleza E. R. A.-B., Benítez Mariana, Espinosa-Soto Carlos, Pablo (2008) **Gene regulatory network models: A dynamic and integrative approach to development** *Practical Systems Biology*
- [3] Huang S., Eichler G., Bar-Yam Y., Ingber D. E. (2005) **Cell Fates as High-Dimensional Attractor States of a Complex Gene Regulatory Network** *Phys. Rev. Lett.* **94** <https://doi.org/10.1103/PhysRevLett.94.128701>
- [4] Davidson E. H. (2010) **Emerging properties of animal gene regulatory networks** *Nature* **468**:911–920 <https://doi.org/10.1038/nature09645>
- [5] Peter I. S., Davidson E. H. (2011) **Evolution of Gene Regulatory Networks Controlling Body Plan Development** *Cell* **144**:970–985 <https://doi.org/10.1016/j.cell.2011.02.017>
- [6] ten Tusscher K. H., Hogeweg P. (2011) **Evolution of Networks for Body Plan Patterning; Interplay of Modularity, Robustness and Evolvability** *PLOS Comput. Biol.* **7** <https://doi.org/10.1371/journal.pcbi.1002208>
- [7] Kim H., Sayama H. (2018) **How Criticality of Gene Regulatory Networks Affects the Resulting Morphogenesis under Genetic Perturbations** *Artif. Life* **24**:85–105 https://doi.org/10.1162/ARTL_a_00262
- [8] Srivastava M. (2021) **Beyond Casual Resemblance: Rigorous Frameworks for Comparing Regeneration Across Species** *Annu. Rev. Cell Dev. Biol.* **37**:415–440 <https://doi.org/10.1146/annurev-cellbio-120319-114716>
- [9] Singh A. J., Ramsey S. A., Filtz T. M., Kiousi C. (2018) **Differential gene regulatory networks in development and disease** *Cell. Mol. Life Sci.* **75**:1013–1025 <https://doi.org/10.1007/s00018-017-2679-6>
- [10] Qin G., Yang L., Ma Y., Liu J., Huo Q. (2019) **The exploration of disease-specific gene regulatory networks in esophageal carcinoma and stomach adenocarcinoma** *BMC Bioinformatics* **20** <https://doi.org/10.1186/s12859-019-3230-6>
- [11] Fazilaty H., et al. (2019) **A gene regulatory network to control EMT programs in development and disease** *Nat. Commun.* **10** <https://doi.org/10.1038/s41467-019-13091-8>
- [12] Davies J., Levin M. (2022) **Synthetic morphology via active and agential matter** *OSF Preprints* <https://doi.org/10.31219/osf.io/xrv8h>
- [13] Toda S., Blauch L. R., Tang S. K. Y., Morsut L., Lim W. A. (2018) **Programming selforganizing multicellular structures with synthetic cell-cell signaling** *Science* **361**:156–162 <https://doi.org/10.1126/science.aat0271>

- [14] Toda S., McKeithan W. L., Hakkinen T. J., Lopez P., Klein O. D., Lim W. A. (2020) **Engineering synthetic morphogen systems that can program multicellular patterning** *Science* **370**:327–331 <https://doi.org/10.1126/science.abc0033>
- [15] Ho C., Morsut L. (2021) **Novel synthetic biology approaches for developmental systems** *Stem Cell Rep.* **16**:1051–1064 <https://doi.org/10.1016/j.stemcr.2021.04.007>
- [16] Santorelli M., Lam C., Morsut L. (2019) **Synthetic development: building mammalian multicellular structures with artificial genetic programs** *Curr. Opin. Biotechnol.* **59**:130–140 <https://doi.org/10.1016/j.copbio.2019.03.016>
- [17] de Jong H. (2002) **Modeling and Simulation of Genetic Regulatory Systems: A Literature Review** *J. Comput. Biol.* **9**:67–103 <https://doi.org/10.1089/10665270252833208>
- [18] Schlitt T., Brazma A. (2007) **Current approaches to gene regulatory network modelling** *BMC Bioinformatics* **8** <https://doi.org/10.1186/1471-2105-8-S6-S9>
- [19] Fetrow J. S., Babbitt P. C. (2018) **New computational approaches to understanding molecular protein function** *PLOS Comput. Biol.* **14** <https://doi.org/10.1371/journal.pcbi.1005756>
- [20] Delgado F. M., Gomez-Vela F. (2019) **Computational methods for Gene Regulatory Networks reconstruction and analysis: A review** *Artif. Intell. Med.* **95** <https://doi.org/10.1016/j.artmed.2018.10.006>
- [21] Glont M., et al. (2018) **BioModels: expanding horizons to include more modelling approaches and formats** *Nucleic Acids Res.* **46**:D1248–D1253 <https://doi.org/10.1093/nar/gkx1023>
- [22] Malik-Sheriff R. S., et al. (2020) **BioModels—15 years of sharing computational models in life science** *Nucleic Acids Res.* **48**:D407–D415 <https://doi.org/10.1093/nar/gkz1055>
- [23] Kauffman S. A. (1993) **The origins of order: Self-organization and selection in evolution**
- [24] Kauffman S. A. (1995) **At home in the universe: The search for laws of self-organization and complexity**
- [25] Abramson C. I., Levin M. (2021) **Behaviorist approaches to investigating memory and learning: A primer for synthetic biology and bioengineering** *Commun. Integr. Biol.* **14**:230–247 <https://doi.org/10.1080/19420889.2021.2005863>
- [26] Baluska F., Levin M. (2016) **On Having No Head: Cognition throughout Biological Systems** *Front. Psychol.* **7** <https://doi.org/10.3389/fpsyg.2016.00902>
- [27] Dodig-Crnkovic G. (2022) **Cognition as Morphological/Morphogenetic Embodied Computation In Vivo** *Entropy Basel Switz.* **24** <https://doi.org/10.3390/e24111576>
- [28] Timsit Y., Grégoire S.-P. (2021) **Towards the Idea of Molecular Brains** *Int. J. Mol. Sci.* **22** <https://doi.org/10.3390/ijms222111868>
- [29] Katz Y., Springer M., Fontana W. (2018) **Embodying probabilistic inference in biochemical circuits** *arXiv* <https://doi.org/10.48550/arXiv.1806.10161>

- [30] Csermely P., et al. (2020) **Learning of Signaling Networks: Molecular Mechanisms** *Trends Biochem. Sci.* **45**:284–294 <https://doi.org/10.1016/j.tibs.2019.12.005>
- [31] Gyurko D. M., Veres D. V., Modos D., Lenti K., Korcsmáros T., Csermely P. (2013) **Adaptation and learning of molecular networks as a description of cancer development at the systems-level: Potential use in anti-cancer therapies** *Semin. Cancer Biol.* **23**:262–269 <https://doi.org/10.1016/j.semcancer.2013.06.005>
- [32] Fields C., Levin M. (2022) **Competency in Navigating Arbitrary Spaces: Intelligence as an Invariant for Analyzing Cognition in Diverse Embodiments** *PsyArXiv* <https://doi.org/10.31234/osf.io/87nzu>
- [33] Watson R. et al. (2010) **Associative memory in gene regulation networks** :659–666
- [34] Mathews J., Chang A. (Jaelyn), Devlin L., Levin M. (2023) **Cellular signaling pathways as plastic, proto-cognitive systems: Implications for biomedicine** *Patterns* **4** <https://doi.org/10.1016/j.patter.2023.100737>
- [35] Lagasse E., Levin M. (2023) **Future medicine: from molecular pathways to the collective intelligence of the body** *Trends Mol. Med.* **29**:687–710 <https://doi.org/10.1016/j.molmed.2023.06.007>
- [36] Clawson W. P., Levin M. (2022) **Endless forms most beautiful 2.0: teleonomy and the bioengineering of chimaeric and synthetic organisms** *Biol. J. Linn. Soc.* <https://doi.org/10.1093/biolinnean/blac073>
- [37] Krist K. T., Sen A., Noid W. G. (2021) **A simple theory for molecular chemotaxis driven by specific binding interactions** *J. Chem. Phys.* **155** <https://doi.org/10.1063/5.0061376>
- [38] Cejkova J., Banno T., Hanczyc M. M., Stepanek F. (2017) **Droplets As Liquid Robots** *Artif. Life* **23**:528–549 https://doi.org/10.1162/ARTL_a_00243
- [39] Hanczyc M. M., Caschera F., Rasmussen S. (2011) **Models of Minimal Physical Intelligence** *Procedia Comput. Sci.* **7**:275–277 <https://doi.org/10.1016/j.procs.2011.09.058>
- [40] Rosenblueth A., Wiener N., Bigelow J. (1943) **Behavior, Purpose and Teleology** *Philos. Sci.* **10**:18–24 <https://doi.org/10.1086/286788>
- [41] Bongard J., Levin M. (2021) **Living Things Are Not (20th Century) Machines: Updating Mechanism Metaphors in Light of the Modern Science of Machine Behavior** *Front. Ecol. Evol.* **9** <https://doi.org/10.3389/fevo.2021.650726>
- [42] Levin M. (2022) **Technological Approach to Mind Everywhere: An Experimentally-Grounded Framework for Understanding Diverse Bodies and Minds** *Front. Syst. Neurosci.* **16** <https://doi.org/10.3389/fnsys.2022.768201>
- [43] Lyon P. (2006) **The biogenic approach to cognition** *Cogn. Process.* **7**:11–29 <https://doi.org/10.1007/s10339-005-0016-8>
- [44] Barandiaran X., Moreno A. (2006) **On What Makes Certain Dynamical Systems Cognitive: A Minimally Cognitive Organization Program** *Adapt. Behav.* **14**:171–185 <https://doi.org/10.1177/105971230601400208>

- [45] di Primio F., Müller B. S., Lengeler J. W. (2000) **Minimal cognition in unicellular organisms** *Anim. Animats* :3–12
- [46] McGivern P. (2020) **Active materials: minimal models of cognition?** *Adapt. Behav.* **28**:441–451 <https://doi.org/10.1177/1059712319891742>
- [47] Levin M. (2023) **Darwin’s agential materials: evolutionary implications of multiscale competency in developmental biology** *Cell. Mol. Life Sci.* **80** <https://doi.org/10.1007/s00018-023-04790-z>
- [48] Pezzulo G., Levin M. (2015) **Re-membering the body: applications of computational neuroscience to the top-down control of regeneration of limbs and other complex organs** *Integr. Biol. Quant. Biosci. Nano Macro* **7**:1487–1517 <https://doi.org/10.1039/c5ib00221d>
- [49] Pezzulo G., Levin M. (2016) **Top-down models in biology: explanation and control of complex living systems above the molecular level** *J. R. Soc. Interface* **13** <https://doi.org/10.1098/rsif.2016.0555>
- [50] Wong D. J., et al. (2008) **Revealing Targeted Therapy for Human Cancer by Gene Module Maps** *Cancer Res.* **68**:369–378 <https://doi.org/10.1158/0008-5472.CAN-07-0382>
- [51] Samuel T. J., Rosenberry R. P., Lee S., Pan Z. (2018) **Correcting Calcium Dysregulation in Chronic Heart Failure Using SERCA2a Gene Therapy** *Int. J. Mol. Sci.* **19** <https://doi.org/10.3390/ijms19041086>
- [52] Krzyszton R., Wan Y., Petreczky J., Balazsi G. (2021) **Gene-circuit therapy on the horizon: Synthetic biology tools for engineered therapeutics** *Acta Biochim. Pol.* **68**:377–383 https://doi.org/10.18388/abp.2020_5744
- [53] Baum C. (2007) **Insertional mutagenesis in gene therapy and stem cell biology** *Curr. Opin. Hematol.* **14** <https://doi.org/10.1097/MOH.0b013e3281900f01>
- [54] Lobo D., Solano M., Bubenik G. A., Levin M. (2014) **A linear-encoding model explains the variability of the target morphology in regeneration** *J. R. Soc. Interface* **11** <https://doi.org/10.1098/rsif.2013.0918>
- [55] Stucki J. W. (1979) **Stability analysis of biochemical systems— A practical guide** *Prog. Biophys. Mol. Biol.* **33**:99–187 [https://doi.org/10.1016/0079-6107\(79\)90027-0](https://doi.org/10.1016/0079-6107(79)90027-0)
- [56] Ingalls B. P. (2004) **A Frequency Domain Approach to Sensitivity Analysis of Biochemical Networks** *J. Phys. Chem. B* **108**:1143–1152 <https://doi.org/10.1021/jp036567u>
- [57] Ingalls B. (2008) **Sensitivity analysis: from model parameters to system behaviour** *Essays Biochem.* **45**:177–194 <https://doi.org/10.1042/bse0450177>
- [58] Donze A., Clermont G., Langmead C. J. (2010) **Parameter synthesis in nonlinear dynamical systems: application to systems biology** *J. Comput. Biol. J. Comput. Mol. Cell Biol.* **17**:325–336 <https://doi.org/10.1089/cmb.2009.0172>
- [59] Dang T., Le Guernic C., Maler O. (2011) **Computing reachable states for nonlinear biological models** *Theor. Comput. Sci.* **412**:2095–2107 <https://doi.org/10.1016/j.tcs.2011.01.014>

- [60] Donze A., Fanchon E., Gattepaille L. M., Maler O., Tracqui P. (2011) **Robustness Analysis and Behavior Discrimination in Enzymatic Reaction Networks** *PLOS ONE* **6** <https://doi.org/10.1371/journal.pone.0024246>
- [61] Rozum J., Albert R. (2022) **Leveraging network structure in nonlinear control** *Npj Syst. Biol. Appl.* **8** <https://doi.org/10.1038/s41540-022-00249-2>
- [62] Steinway S. N., Zanudo J. G. T., Michel P. J., Feith D. J., Loughran T. P., Albert R. (2015) **Combinatorial interventions inhibit TGF β -driven epithelial-to-mesenchymal transition and support hybrid cellular phenotypes** *Npj Syst. Biol. Appl.* **1** <https://doi.org/10.1038/npjbsba.2015.14>
- [63] Zanudo J. G. T., Albert R. (2015) **Cell Fate Reprogramming by Control of Intracellular Network Dynamics** *PLOS Comput. Biol.* **11** <https://doi.org/10.1371/journal.pcbi.1004193>
- [64] Zanudo J. G. T., Yang G., Albert R. (2017) **Structure-based control of complex networks with nonlinear dynamics** *Proc. Natl. Acad. Sci.* **114**:7234–7239 <https://doi.org/10.1073/pnas.1617387114>
- [65] Cifuentes Fontanals L., Tonello E., Siebert H., Abate A., Petrov T., Wolf V. (2020) **Control Strategy Identification via Trap Spaces in Boolean Networks** *Computational Methods in Systems Biology* :159–175 https://doi.org/10.1007/978-3-030-60327-4_9
- [66] Murrugarra D., Veliz-Cuba A., Aguilar B., Laubenbacher R. (2016) **Identification of control targets in Boolean molecular network models via computational algebra** *BMC Syst. Biol.* **10** <https://doi.org/10.1186/s12918-016-0332-x>
- [67] Choo S.-M., Ban B., Joo J. I., Cho K.-H. (2018) **The phenotype control kernel of a biomolecular regulatory network** *BMC Syst. Biol.* **12** <https://doi.org/10.1186/s12918-018-0576-8>
- [68] Choo S.-M., Park S.-M., Cho K.-H. (2019) **Minimal intervening control of biomolecular networks leading to a desired cellular state** *Sci. Rep.* **9** <https://doi.org/10.1038/s41598-019-49571-6>
- [69] Paladugu S. R., Chickarmane V., Deckard A., Frumkin J. P., McCormack M., Sauro H. M. (2006) **In silico evolution of functional modules in biochemical networks** *IEE Proc. – Syst. Biol.* **153**:223–235 <https://doi.org/10.1049/ip-syb:20050096>
- [70] François P. (2014) **Evolving phenotypic networks in silico** *Semin. Cell Dev. Biol.* **35**:90–97 <https://doi.org/10.1016/j.semcdb.2014.06.012>
- [71] Noman N., Monjo T., Moscato P., Iba H. (2015) **Evolving Robust Gene Regulatory Networks** *PLOS ONE* **10** <https://doi.org/10.1371/journal.pone.0116258>
- [72] Hiscock T. W. (2019) **Adapting machine-learning algorithms to design gene circuits** *BMC Bioinformatics* **20** <https://doi.org/10.1186/s12859-019-2788-3>
- [73] Shen J., Liu F., Tu Y., Tang C. (2021) **Finding gene network topologies for given biological function with recurrent neural network** *Nat. Commun.* **12** <https://doi.org/10.1038/s41467-021-23420-5>
- [74] Camacho D. M., Collins K. M., Powers R. K., Costello J. C., Collins J. J. (2018) **NextGeneration Machine Learning for Biological Networks** *Cell* **173** <https://doi.org/10.1016/j.cell.2018.05.015>

- [75] Volk M. J., Lourentzou I., Mishra S., Vo L. T., Zhai C., Zhao H. (2020) **Biosystems Design by Machine Learning** *ACS Synth. Biol.* **9**:1514–1533 <https://doi.org/10.1021/acssynbio.0c00129>
- [76] Kitano H. (2007) **A robustness-based approach to systems-oriented drug design** *Nat. Rev. Drug Discov.* **6** <https://doi.org/10.1038/nrd2195>
- [77] Manicka S., Levin M. (2019) **The Cognitive Lens: a primer on conceptual tools for analysing information processing in developmental and regenerative morphogenesis** *Philos. Trans. R. Soc. B Biol. Sci.* **374** <https://doi.org/10.1098/rstb.2018.0369>
- [78] Biswas S., Manicka S., Hoel E., Levin M. (2021) **Gene regulatory networks exhibit several kinds of memory: Quantification of memory in biological and random transcriptional networks** *iScience* **24** <https://doi.org/10.1016/j.isci.2021.102131>
- [79] Biswas S., Clawson W., Levin M. (2023) **Learning in Transcriptional Network Models: Computational Discovery of Pathway-Level Memory and Effective Interventions** *Int. J. Mol. Sci.* **24** <https://doi.org/10.3390/ijms24010285>
- [80] Laties V. G. (1987) **Society for the Experimental Analysis of Behavior: The first thirty years (1957-1987)** *J. Exp. Anal. Behav.* **48**:495–512 <https://doi.org/10.1901/jeab.1987.48-495>
- [81] Amdam G., Hovland A. (2012) **Measuring Animal Preferences and Choice Behavior** *Nat. Educ. Knowl.* **3**
- [82] McLeold S. (2022) **Behavioral Perspective in Psychology [Behaviorism Theory]**
- [83] Murugan N. J., et al. (2021) **Mechanosensation Mediates Long-Range Spatial Decision Making in an Aneural Organism** *Adv. Mater. Deerfield Beach Fla* **33** <https://doi.org/10.1002/adma.202008161>
- [84] Mikhaltsov A. (2013) **A. Mikhaltsov, Paramecium bursaria. 2013. [Online]. Available: https://commons.wikimedia.org/wiki/File:Paramecium_bursaria.jpg**
- [85] Bongard J., Levin M. (2023) **There's Plenty of Room Right Here: Biological Systems as Evolved, Overloaded, Multi-Scale Machines** *Biomimetics* **8** <https://doi.org/10.3390/biomimetics8010110>
- [86] Baranes A., Oudeyer P.-Y. (2013) **Active learning of inverse models with intrinsically motivated goal exploration in robots** *Robot. Auton. Syst.* **61**:49–73 <https://doi.org/10.1016/j.robot.2012.05.008>
- [87] Forestier S., Portelas R., Mollard Y., Oudeyer P.-Y. (2022) **Intrinsically Motivated Goal Exploration Processes with Automatic Curriculum Learning** *arXiv* <https://doi.org/10.48550/arXiv.1708.02190>
- [88] Grizou J., Points L. J., Sharma A., Cronin L. (2020) **A curious formulation robot enables the discovery of a novel protocell behavior** *Sci. Adv.* **6** <https://doi.org/10.1126/sciadv.aay4237>
- [89] Falk M. J., Roach F. D., Gilpin W., Murugan A. (2023) **Curiosity-driven search for novel non-equilibrium behaviors** *arXiv* <https://doi.org/10.48550/arXiv.2211.02589>
- [90] Reinke C., Etcheverry M., Oudeyer P.-Y. (2020) **Intrinsically Motivated Discovery of Diverse Patterns in Self-Organizing Systems** *Eighth International Conference on Learning Representations*

- [91] Etcheverry M., Moulin-Frier C., Oudeyer P.-Y. (2020) **Hierarchically Organized Latent Modules for Exploratory Search in Morphogenetic Systems** *Advances in Neural Information Processing Systems* :4846–4859
- [92] Hamon G., Etcheverry M., Chan B. W.-C., Moulin-Frier C., Oudeyer P.-Y. (2022) **Learning Sensorimotor Agency in Cellular Automata**
- [93] Kitano H. (2007) **Towards a theory of biological robustness** *Mol. Syst. Biol.* **3** <https://doi.org/10.1038/msb4100179>
- [94] Félix M.-A., Barkoulas M. (2015) **Pervasive robustness in biological systems** *Nat. Rev. Genet.* **16** <https://doi.org/10.1038/nrg3949>
- [95] Ingolia N. T. (2004) **Topology and Robustness in the Drosophila Segment Polarity Network** *PLOS Biol.* **2** <https://doi.org/10.1371/journal.pbio.0020123>
- [96] Ma W., Lai L., Ouyang Q., Tang C. (2006) **Robustness and modular design of the Drosophila segment polarity network** *Mol. Syst. Biol.* **2** <https://doi.org/10.1038/msb4100111>
- [97] Deutscher D., Meilijson I., Kupiec M., Ruppin E. (2006) **Multiple knockout analysis of genetic robustness in the yeast metabolic network** *Nat. Genet.* **38** <https://doi.org/10.1038/ng1856>
- [98] von Dassow G., Meir E., Munro E. M., Odell G. M. (2000) **The segment polarity network is a robust developmental module** *Nature* **406** <https://doi.org/10.1038/35018085>
- [99] Kwang-Hyun C., Sung-Young S., Hyun-Woo K., Wolkenhauer O., McFerran B., Kolch W., Priami C. (2003) **Mathematical Modeling of the Influence of RKIP on the ERK Signaling Pathway** *Computational Methods in Systems Biology* **2602**:127–141 https://doi.org/10.1007/3-540-36481-1_11
- [100] Lehman J., Stanley K. O. (2008) **Exploiting Open-Endedness to Solve Problems Through the Search for Novelty** *IEEE Symposium on Artificial Life*
- [101] Lehman J., Stanley K. O. (2011) **Abandoning Objectives: Evolution Through the Search for Novelty Alone** *Evol. Comput.* **19**:189–223 https://doi.org/10.1162/EVCO_a_00025
- [102] Cully A., Clune J., Tarapore D., Mouret J.-B. (2015) **Robots that can adapt like animals** *Nature* **521** <https://doi.org/10.1038/nature14422>
- [103] Pugh J. K., Soros L. B., Stanley K. O. (2016) **Quality Diversity: A New Frontier for Evolutionary Computation** *Front. Robot. AI* **3** <https://doi.org/10.3389/frobt.2016.00040>
- [104] Doncieux S., Laflaquiere A., Coninx A. (2019) **Novelty search: a theoretical perspective** *Proceedings of the Genetic and Evolutionary Computation Conference* :99–106 <https://doi.org/10.1145/3321707.3321752>
- [105] Benureau F. (2015) **F. Benureau, 2015 “Self Exploration of Sensorimotor Spaces in Robots.” Université de Bordeaux.**
- [106] McInnes L., Healy J., Astels S. (2017) **hdbscan: Hierarchical density based clustering** *J. Open Source Softw.* **2**
- [107] Bell C. C., Gilan O. (2020) **Principles and mechanisms of non-genetic resistance in cancer** *Br. J. Cancer* **122** <https://doi.org/10.1038/s41416-019-0648-6>

- [108] Rizk A., Batt G., Fages F., Soliman S. (2009) **A general computational method for robustness analysis with applications to synthetic gene networks** *Bioinformatics* **25**:i169–i178 <https://doi.org/10.1093/bioinformatics/btp200>
- [109] Walcott C. (1996) **Pigeon Homing: Observations, Experiments and Confusions** *J. Exp. Biol.* **199**:21–27 <https://doi.org/10.1242/jeb.199.1.21>
- [110] Luschi P., et al. (2001) **Testing the Navigational Abilities of Ocean Migrants: Displacement Experiments on Green Sea Turtles (*Chelonia mydas*)** *Behav. Ecol. Sociobiol.* **50**:528–534
- [111] Bisch-Knaden S., Wehner R. (2001) **Egocentric information helps desert ants to navigate around familiar obstacles** *J. Exp. Biol.* **204**:4177–4184 <https://doi.org/10.1242/jeb.204.24.4177>
- [112] Abramson C. I. (1994) **A primer of invertebrate learning: the behavioral perspective.**
- [113] Venkatachalapathy H., Azarin S. M., Sarkar C. A. (2021) **Trajectory-based energy landscapes of gene regulatory networks** *Biophys. J.* **120**:687–698 <https://doi.org/10.1016/j.bpj.2020.11.2279>
- [114] Li C., Wang J. (2014) **Landscape and flux reveal a new global view and physical quantification of mammalian cell cycle** *Proc. Natl. Acad. Sci.* **111**:14130–14135 <https://doi.org/10.1073/pnas.1408628111>
- [115] Li C., Wang J. (2013) **Quantifying Cell Fate Decisions for Differentiation and Reprogramming of a Human Stem Cell Network: Landscape and Biological Paths** *PLOS Comput. Biol.* **9** <https://doi.org/10.1371/journal.pcbi.1003165>
- [116] Lee H. C., Tian B., Sedivy J. M., Wands J. R., Kim M. (2006) **Loss of Raf Kinase Inhibitor Protein Promotes Cell Proliferation and Migration of Human Hepatoma Cells** *Gastroenterology* **131**:1208–1217 <https://doi.org/10.1053/j.gastro.2006.07.012>
- [117] Reinitz J., Sharp D. H. (1995) **Mechanism of eve stripe formation** *Mech. Dev.* **49**:133–158 [https://doi.org/10.1016/0925-4773\(94\)00310-j](https://doi.org/10.1016/0925-4773(94)00310-j)
- [118] Jaeger J., et al. (2004) **Dynamical Analysis of Regulatory Interactions in the Gap Gene System of *Drosophila melanogaster*** *Genetics* **167**:1721–1737 <https://doi.org/10.1534/genetics.104.027334>
- [119] Cotterell J., Sharpe J. (2010) **An atlas of gene regulatory networks reveals multiple three-gene mechanisms for interpreting morphogen gradients** *Mol. Syst. Biol.* **6** <https://doi.org/10.1038/msb.2010.74>
- [120] Molinelli E. J., et al. (2013) **Perturbation Biology: Inferring Signaling Networks in Cellular Systems** *PLOS Comput. Biol.* **9** <https://doi.org/10.1371/journal.pcbi.1003290>
- [121] Vallverdú J., et al. (2018) **Slime mould: The fundamental mechanisms of biological cognition** *Biosystems* **165**:57–70 <https://doi.org/10.1016/j.biosystems.2017.12.011>
- [122] Beekman M., Latty T. (2015) **Brainless but Multi-Headed: Decision Making by the Acellular Slime Mould *Physarum polycephalum*** *J. Mol. Biol.* **427** <https://doi.org/10.1016/j.jmb.2015.07.007>
- [123] Saigusa T., Tero A., Nakagaki T., Kuramoto Y. (2008) **Amoebae Anticipate Periodic Events** *Phys. Rev. Lett.* **100** <https://doi.org/10.1103/PhysRevLett.100.018101>

- [124] Nakagaki T., Guy R. D. (2008) **Intelligent behaviors of amoeboid movement based on complex dynamics of soft matter** *Soft Matter* **4**:57–67 <https://doi.org/10.1039/B706317M>
- [125] Colas C., Sigaud O., Oudeyer P.-Y. (2018) **GEP-PG: Decoupling Exploration and Exploitation in Deep Reinforcement Learning Algorithms** *Proceedings of the 35th International Conference on Machine Learning* :1039–1048
- [126] Pandi A., et al. (2022) **A versatile active learning workflow for optimization of genetic and metabolic networks** *Nat. Commun.* **13** <https://doi.org/10.1038/s41467-022-31245-z>
- [127] Libby A. R. G., et al. (2019) **Automated Design of Pluripotent Stem Cell Self-Organization** *Cell Syst.* **9**:483–495 <https://doi.org/10.1016/j.cels.2019.10.008>
- [128] Pietak A., Levin M. (2016) **Exploring Instructive Physiological Signaling with the Bioelectric Tissue Simulation Engine** *Front. Bioeng. Biotechnol.* **4** <https://doi.org/10.3389/fbioe.2016.00055>
- [129] Koseska A., Bastiaens P. I. (2017) **Cell signaling as a cognitive process** *EMBO J.* **36**:568–582 <https://doi.org/10.15252/embj.201695383>
- [130] Baluska F., Reber A. S., Miller W. B. (2022) **Cellular sentience as the primary source of biological order and evolution** *Biosystems* **218** <https://doi.org/10.1016/j.biosystems.2022.104694>
- [131] Baluska F., Miller W. B., Reber A. S. (2023) **Cellular and evolutionary perspectives on organismal cognition: from unicellular to multicellular organisms** *Biol. J. Linn. Soc.* **139**:503–513 <https://doi.org/10.1093/biolinnean/blac005>
- [132] Reber A. S., Baluska F. (2021) **Cognition in some surprising places** *Biochem. Biophys. Res. Commun.* **564**:150–157 <https://doi.org/10.1016/j.bbrc.2020.08.115>
- [133] Baluska F., Reber A. S. (2021) **Cellular and organismal agency – Not based on genes: A comment on Baverstock** *Prog. Biophys. Mol. Biol.* **167**:161–162 <https://doi.org/10.1016/j.pbiomolbio.2021.11.001>
- [134] Bernheim-Groswasser A., Gov N. S., Safran S. A., Tzliil S. (2018) **Living Matter: Mesoscopic Active Materials** *Adv. Mater.* **30** <https://doi.org/10.1002/adma.201707028>
- [135] Etcheverry M., Levin M., Moulin-Frier C., Oudeyer P.-Y. (2023) **SBMLtoODEjax: efficient simulation and optimization of ODE SBML models in JAX** *arXiv* <https://doi.org/10.48550/arXiv.2307.08452>
- [136] Pugh J. K., Soros L. B., Szerlip P. A., Stanley K. O. (2015) **Confronting the challenge of quality diversity** *Proceedings of the 2015 annual conference on genetic and evolutionary computation* :967–974
- [137] Scheiner S. M. (2019) **A compilation of and typology for abundance-, phylogenetic-and functional-based diversity metrics** *BioRxiv Prepr. Serv. Biol.*
- [138] Gillies S., et al. (2022) **Shapely** *Zenodo* <https://doi.org/10.5281/zenodo.7428463>
- [139] Hansen N., Müller S.C., Koumoutsakos P. (2003) **Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES)** *Evolutionary computation* **11**:1–18

- [140] Topp B., Promislow K., deVries G., Miura R.M., Finegood D.T. (2000) **A Model of Beta-Cell Mass, Insulin, and Glucose Kinetics: Pathways to Diabetes** *Journal of Theoretical Biology* **206**:605–619 <https://doi.org/10.1006/jtbi.2000.2150>
- [141] Smallbone K. (2013) **Metabolic Control Analysis: Rereading Reder** *arXiv*

Editors

Reviewing Editor

Arvind Murugan

University of Chicago, Chicago, United States of America

Senior Editor

Aleksandra Walczak

École Normale Supérieure - PSL, Paris, France

Reviewer #1 (Public Review):

Summary:

This paper suggests to apply intrinsically-motivated exploration for the discovery of robust goal states in gene regulatory networks.

Strengths:

The paper is well written. The biological motivation and the need for such methods are formulated extraordinarily well. The battery of experimental models is impressive.

Weaknesses:

(1) The proposed method is compared to the random search. That says little about the performance with regard to the true steady-state goal sets. The latter could be calculated at least for a few simple ODE (e.g., BIOMD000000454, Metabolic Control Analysis: Rereading Reder'). The experiment with 'oscillator circuits' may not be directly interpolated to the other models.

The lack of comparison to the ground truth goal set (attractors of ODE) from arbitrary initial conditions makes it hard to evaluate the true performance/contribution of the method. A part of the used models can be analyzed numerically using JAX, while there are models that can be analyzed analytically.

"...The true versatility of the GRN is unknown and can only be inferred through empirical exploration and proxy metrics...": one could perform a sensitivity analysis of the ODEs, identifying stable equilibria. That could provide a proxy for the ground truth 'versatility'.

(2) The proposed method is based on 'Intrinsically Motivated Goal Exploration Processes with Automatic Curriculum Learning', which assumes state action trajectories $[s_{t_0:t}, a_{t_0:t}]$, (2.1 Notations and Assumptions' in the IMGEP paper). However, the models used in the current work do not include external control actions, but rather only the initial conditions can be set. It is not clear from the methods whether IMGEP was adapted to this setting, and how the exploration policy was designed w/o actual time-dependent actions. What does "...generates candidate intervention parameters to achieve the current goal..." mean considering that interventions 'Sets the initial state...' as explained in Table 2?

(3) Fig 2 shows the phase space for (ERK, RKIPP_RP) without mentioning the typical full scale of ERK, RKIPP_RP. It is unclear whether the path from (0, 0) to (~0.575, ~3.75) at $t=1000$ is significant on the typical scale of this phase space. Is it significant on the typical scale of this phase space?

(4) Table 2:

(a) Where is 'effective intervention' used in the method?

(b) In my opinion 'controllability', 'trainability', and 'versatility' are different terms. If there is correspondence, it is important. I would suggest to extend/enhance the column "Proposed Isomorphism". Otherwise, it may be confusing. I don't see how this table generalizes concepts from dynamical complex systems and behavioral sciences under a common navigation task perspective".

<https://doi.org/10.7554/eLife.92683.3.sa2>

Reviewer #2 (Public Review):

Summary:

Etcheverry et al. present two computational frameworks for exploring the functional capabilities of gene regulatory networks (GRNs). The first is a framework based on intrinsically motivated exploration, here used to reveal the set of steady states achievable by a given gene regulatory network as a function of initial conditions. The second is a behaviorist framework, here used to assess the robustness of steady states to dynamical perturbations experienced along typical trajectories to those steady states. In Figs. 1-5, the authors convincingly show how these frameworks can explore and quantify the diversity of behaviors that can be displayed by GRNs. In Figs. 6-9, the authors present applications of their framework to the analysis and control of GRNs, but the support presented for their case studies is often incomplete.

Following revision, my overall perspective of the paper remains unchanged. The first half of the paper provides solid evidence to support an important conceptual framework. The evidence presented for the use cases in the latter half is incomplete; as the authors note, they are preliminary and meant to be built on in future work. I have included my first round comments below.

Strengths:

Overall, the paper presents an important development for exploring and understanding GRNs/dynamical systems broadly, with solid evidence supporting the first half of their paper in a narratively clear way.

The behaviorist point of view for robustness is potentially of interest to a broad community, and to my knowledge introduces novel considerations for defining robustness in the GRN context.

Some specific weaknesses, mostly concerning incomplete analyses in the second half of the paper:

(1) The analysis presented in Fig. 6 is exciting but preliminary. Are there other appropriate methods for constructing energy landscapes from dynamical trajectories in gene regulatory networks? How do the results in this particular case study compare to other GRNs studied in the paper?

Additionally, it is unclear whether the analysis presented in Fig. 6C is appropriate. In particular, if the pseudopotential landscapes are constructed from statistics of visited states

along trajectories to the steady state, then the trajectories derived from dynamical perturbations do not only reflect the underlying pseudo-landscape of the GRN. Instead, they also include contributions from the perturbations themselves.

(2) In Fig. 7, I'm not sure how much is possible to take away from the results as given here, as they depend sensitively on the cohort of 432 (GRN, Z) pairs used. The comparison against random networks is well-motivated. However, as the authors note, comparison between organismal categories is more difficult due to low sample size; for instance, the "plant" and "slime mold" categories each only has 1 associated GRN. Additionally, the "n/a" category is difficult to interpret.

(3) In Fig. 8, it is unclear whether the behavioral catalog generated is important to the intervention design problem of moving a system in one attractor basin to another. The authors note that evolutionary searches or SGD could also be used to solve the problem. Is the analysis somehow enabled by the behavioral catalog in a way that is complementary to those methods? If not, comparison against those methods (or others e.g. optimal control) would strengthen the paper.

(4) The analysis presented in Fig. 9 also is preliminary. The authors note that there exist many algorithms for choosing/identifying the parameter values of a dynamical system that give rise to a desired time series. It would be a stronger result to compare their approach to more sophisticated methods, as opposed to random search and SGD. Other options from the recent literature include Bayesian techniques, sparse nonlinear regression techniques (e.g. SINDy), and evolutionary searches. The authors note that some methods require fine-tuning in order to be successful, but even so, it would be good to know the degree of fine-tuning which is necessary compared to their method. [second round: the authors have included a comparison against CMA-ES, an evolutionary algorithm]

<https://doi.org/10.7554/eLife.92683.3.sa1>

Author response:

The following is the authors' response to the previous reviews.

Recommendations for the authors:

Reviewer #1 (Recommendations For The Authors):

My main concern is still in place. It is unclear whether the proposed method can find actual goal states, and as a result it is unclear what states it finds. Table S1 mentions the model BIOMD000000454, which is a small metabolic pathway with known equations given in "Example One" in "Metabolic Control Analysis: Rereading Reder". In this model the goal states can be calculated analytically.

Regarding your statements below: I am not concerned that your method will be less efficient than random search (or any other search..) on small models, but I think it is important for the readers to have evidence that your method is able to discover true goal states at least in small networks, used in your study. You do show that your method scales to complex models. So, in my opinion, the missing part is to show that it is able to find true goal states.

"...For simple models whose true steady-state distribution can be derived numerically and/or analytically, it is very likely that their exploration will be much simpler and this is not where a lot of improvement over random search may be found, which explains our focus on more complex models..."

We thank you for your response and for your concerns on the lack of evidence that our method is able to re-discover the true goal states of simple models when these are known a priori. We acknowledge that adding these simple cases is useful for completeness. We did not include these simple models in our main study because in most cases a basic random search over the initial conditions will lead to the re-discovery of these goal states. For instance for the mentioned model BIOMD000000454 described in the "Example One" from the "Metabolic Control Analysis: Rereading Reder" paper, several simplifying assumptions are made such that the system only has one steady state ($x_1=0.056$, $x_2=0.769$, $x_3=4.231$) which can be found analytically as shown in the paper. In that simple case, this goal state is also straightforward to find with numerical simulation as any valid initial condition will converge to it.

To address the concerns of the reviewer, we propose to add an additional "sanity check" figure in the supplementary of the revised paper (Figure S4), as well as a "sanity check" subsection in the "Methods", to present additional experiments made on simple models such as this one. The novel figure and subsection can be visualized on the paper's interactive version available online <https://developmentalsystems.org/curious-exploration-of-grn-competencies>, and we plan to include them as such in the further revision. We have also included the full code to reproduce this sanity check as a 'sanity_check.ipynb' jupyter notebook in the github repository (https://github.com/flowersteam/curious-exploration-of-grn-competencies/blob/main/notebooks/sanity_check.ipynb).

In the novel figure S4-b, we show the results of our exploration pipeline on the suggested model BIOMD000000454 as described in the "Example One" of the paper. These results provide evidence that the curiosity search is able to find back the correct unique goal state ($x_1=0.056$, $x_2=0.769$, $x_3=4.231$), as expected.

We also include a second sanity check on BIOMD000000341 which models the dynamics of beta-cell mass, insulin and glucose dynamics. This model has two stable fixed points representing physiological ($B=300$, $I=10$, $G=100$) and pathological ($B=0$, $I=0$, $G=600$) steady states, which are the known ground truth steady states as described in Figure 3 of the "A Model of b-Cell Mass, Insulin, and Glucose Kinetics: Pathways to Diabetes" paper. Again, as expected, curiosity search is able to find back those two steady states (Figure S4-a).

As stated in our previous answer, our main study focuses on more complex models that are not limited to one or few attractors that can easily be discovered with random initial conditions. Regarding the mentioned BIOMD000000454, maybe something that has been confusing for the reviewer is that we indeed included it in our main study but, as specified in the caption of table S4, at the difference of what is done in the "example one" of the original paper, we let the metabolite concentrations y_1, \dots, y_5 evolve in time (instead of enforcing them as constants). When doing so, the resulting dynamics of the system are more complex and exhibit a spectrum of possible steady states (unknown a priori), which differ from the previous case with a single steady state. In that case, the new attractors are not analytically easy to find and the proposed curiosity search becomes interesting as it is able to uncover the distribution of possible steady states much more efficiently than a random search baseline, as shown in the new figures S4-c and S4-d.

We hope that these new results will address the reviewer's concerns and provide evidence to the readers on the validity of the approach on simple networks.

eLife assessment

This important study develops a machine learning method to reveal hidden unknown functions and behavior in gene regulatory networks by searching parameter space in an efficient way. The evidence for some parts of the paper is still incomplete and needs systematic comparison to other methods and to the ground truth, but the work will be of

broad interest to anyone working in biology of all stripes since the ideas reach beyond gene regulatory networks to revealing hidden functions in any complex system with many interacting parts.

We thank the editors and reviewers for their positive assessment and constructive suggestions. In our response, we acknowledge the importance of systematic comparison to other methods and to the ground truth, when available. However we also emphasize the challenges associated with evaluating such methods in the context of uncovering hidden behaviors in complex biological networks as the ground truth is often unknown. We hope that our explanations will clarify the potential of our approach in advancing the exploration of these systems.

Public Reviews:

Reviewer #1 (Public Review):

Summary: This paper suggests to apply intrinsically-motivated exploration for the discovery of robust goal states in gene regulatory networks.

Strengths:

The paper is well written. The biological motivation and the need for such methods are formulated extraordinarily well. The battery of experimental models is impressive.

We thank the reviewer for sharing interest in the research problem and for recognizing the strengths of our work.

Weaknesses:

(1) The proposed method is compared to the random search. That says little about the performance with regard to the true steady-state goal sets. The latter could be calculated at least for a few simple ODE (e.g., BIOMD0000000454, Metabolic Control Analysis: Rereading Reder). The experiment with 'oscillator circuits' may not be directly interpolated to the other models.

The lack of comparison to the ground truth goal set (attractors of ODE) from arbitrary initial conditions makes it hard to evaluate the true performance/contribution of the method. A part of the used models can be analyzed numerically using JAX, while there are models that can be analyzed analytically.

"...The true versatility of the GRN is unknown and can only be inferred through empirical exploration and proxy metrics...": one could perform a sensitivity analysis of the ODEs, identifying stable equilibria. That could provide a proxy for the ground truth 'versatility'.

We agree with the reviewer that one primary concern is to properly evaluate the effectiveness of the proposed method. However, as we move toward complex pathways, knowledge of the “true” steady-state goal sets is often unknown which is where the use of machine learning methods as the one we propose are particularly interesting (but challenging to evaluate).

For simple models whose true steady-state distribution can be derived numerically and/or analytically, it is very likely that their exploration will be much simpler and this is not where a lot of improvement over random search may be found, which explains our focus on more complex models. While we agree that it is still interesting to evaluate exploration methods on these simple models for checking their behavior, it is not clear how to scale this analysis to the targeted more complex systems.

For systems whose true steady state distribution cannot be derived analytically or numerically, we believe that random search is a pertinent baseline as it is commonly used in the literature to discover the attractors/trajectories of a biological network. For instance, Venkatachalapathy et al. [1] initialize stochastic simulations at multiple randomly sampled starting conditions (which is called a kinetic Monte Carlo-based method) to capture the steady states of a biological system. Similarly, Donzé et al. [29] use a Monte Carlo approach to compute the reachable set of a biological network «when the number of parameters is large and their uncertain range is not negligible». For the considered models, the true steady-state goal set is unknown, which is why we chose comparison with random search. We added a “Statistics” subsection in the Methods section providing additional details about the statistical analyses we perform between our method and the random search baseline.

(2) The proposed method is based on 'Intrinsically Motivated Goal Exploration Processes with Automatic Curriculum Learning', which assumes state action trajectories $[s_{t_0:t}, a_{t_0:t}]$, (2.1 Notations and Assumptions' in the IMGEP paper). However, the models used in the current work do not include external control actions, but rather only the initial conditions can be set. It is not clear from the methods whether IMGEP was adapted to this setting, and how the exploration policy was designed w/o actual time-dependent actions. What does "...generates candidate intervention parameters to achieve the current goal..." mean considering that interventions 'Sets the initial state...' as explained in Table 2?

We thank the reviewer for asking for clarification, as indeed the IMGEP methodology originates from developmental robotics scenarios which generally focus on the problem of robotic sequential decision-making, therefore assuming state action trajectories as presented in Forestier et al. [65]. However, in both cases, note that the IMGEP is responsible for sampling parameters which then govern the exploration of the dynamical system. In Forestier et al. [65], the IMGEP also only sets one vector at the start (denoted s_0) which was specifying parameters of a movement (like the initial state of the GRN), which was then actually produced with dynamic motion primitives which are dynamical system equations similar to GRN ODEs, so the two systems are mathematically equivalent. More generally, while in our case the “intervention” of the IMGEP (denoted s_0) only controls the initial state of the GRN, future work could consider more advanced sequential interventions simply by setting parameters of an action policy at the start which could be called during the GRN’s trajectory to sample control actions where would be the state of the GRN. In practice this would also require setting only one vector at the start, so it would remain the same exploration algorithm and only the space of parameters would change, which illustrates the generality of the approach.

(3) Fig 2 shows the phase space for (ERK, RKIPP_RP) without mentioning the typical full scale of ERK, RKIPP_RP. It is unclear whether the path from (0, 0) to (-0.575, -3.75) at $t=1000$ is significant on the typical scale of this phase space. is it significant on the typical scale of this phase space?

The purpose of Figure 2 is to illustrate an example of GRN trajectory in transcriptional space, and to illustrate what “interventions” and “perturbations” can be in that context. To that end we have used the fixed initial conditions provided in the BIOMD000000647, replicating Figure 5 of Cho et al. [56].

While we are not sure of what the reviewer means with “typical” scale of this phase space, we would like to point reviewer toward Figure 8 which shows examples of certain paths that indeed reach further point in the same phase space (up to ~ 10 in RKIPP_RP levels and ~ 300 in ERK levels). However, while the paths displayed in Figure 8 are possible (and were discovered with the IMGEP), note that they may be “rarer” to occur naturally in the sense that

a large portion of the tested initial conditions with random search tend to converge toward smaller (ERK, RKIPP_RP) steady-state values similar to the ones displayed in Figure 2.

(4) Table 2:

a. Where is 'effective intervention' used in the method?

b. in my opinion 'controllability', 'trainability', and 'versatility' are different terms. If their correspondence is important I would suggest to extend/enhance the column "Proposed Isomorphism". otherwise, it may be confusing.

a) We thank the reviewer for pointing out that “effective intervention” is not explicitly used in the method. The idea here is that as we are exploring a complex dynamical system (here the GRN), some of the sampled interventions will be particularly effective at revealing novel unseen outcomes whereas others will fail to produce a qualitative change to the distribution of discovered outcomes. What we show in this paper, for instance in Figure 3a and Figure 4, is that the IMGEP method is particularly sample-efficient in finding those “effective interventions”, at least more than a random exploration. However we agree that the term “effective intervention” is ambiguous (does not say effective in what) and we have replaced it with “salient intervention” in the revised version.

b) We thank the reviewer for highlighting some confusing terms in our chosen vocabulary, and we have clarified those terms in the revised version. We agree that controllability/trainability and versatility are not exactly equivalent concepts, as controllability/trainability typically refers to the amount to which a system is externally controllable/trainable whereas versatility typically refers to the inherent adaptability or diversity of behaviors that a system can exhibit in response to inputs or conditions. However, they are both measuring the extent of states that can be reached by the system under a distribution of stimuli/conditions, whether natural conditions or engineered ones, which is why we believe that their correspondence is relevant.

I don't see how this table generalizes "concepts from dynamical complex systems and behavioral sciences under a common navigation task perspective".

We have replaced the verb “generalize” with “investigate” in the revised version.

Reviewer #2 (Public Review):

Summary:

Etcheverry et al. present two computational frameworks for exploring the functional capabilities of gene regulatory networks (GRNs). The first is a framework based on intrinsically-motivated exploration, here used to reveal the set of steady states achievable by a given gene regulatory network as a function of initial conditions. The second is a behaviorist framework, here used to assess the robustness of steady states to dynamical perturbations experienced along typical trajectories to those steady states. In Figs. 1-5, the authors convincingly show how these frameworks can explore and quantify the diversity of behaviors that can be displayed by GRNs. In Figs. 6-9, the authors present applications of their framework to the analysis and control of GRNs, but the support presented for their case studies is often incomplete.

Strengths:

Overall, the paper presents an important development for exploring and understanding GRNs/dynamical systems broadly, with solid evidence supporting the first half of their paper in a narratively clear way.

The behaviorist point of view for robustness is potentially of interest to a broad community, and to my knowledge introduces novel considerations for defining robustness in the GRN context.

We thank the reviewer for recognizing the strengths and novelty of the proposed experimental framework for exploring and understanding GRNs, and complex dynamical systems more generally. We agree that the results presented in the section “Possible Reuses of the Behavioral Catalog and Framework” (Fig 6-9) can be seen as incomplete along certain aspects, which we tried to make as explicit as possible throughout the paper, and why we explicitly state that these are “preliminary experiments”. Despite the discussed limitations, we believe that these experiments are still very useful to illustrate the variety of potential use-cases in which the community could benefit from such computational methods and experimental framework, and build on for future work.

Some specific weaknesses, mostly concerning incomplete analyses in the second half of the paper:

(1) The analysis presented in Fig. 6 is exciting but preliminary. Are there other appropriate methods for constructing energy landscapes from dynamical trajectories in gene regulatory networks? How do the results in this particular case study compare to other GRNs studied in the paper?

We are not aware of other methods than the one proposed by Venkatachalapathy et al. [1] for constructing an energy landscape given an input set of recorded dynamical trajectories, although it might indeed be the case. We want to emphasize that any of such methods would anyway depend on the input set of trajectories, and should therefore benefit from a set that is more representative of the diversity of behaviors that can be achieved by the GRN, which is why we believe the results presented in Figure 6 are interesting. As the IMGEP was able to find a higher diversity of reachable goal states (and corresponding trajectories) for many of the studied GRNs, we believe that similar effects should be observable when constructing the energy landscapes for these GRN models, with the discovery of additional or wider “valleys” of reachable steady states.

Additionally, it is unclear whether the analysis presented in Fig. 6C is appropriate. In particular, if the pseudopotential landscapes are constructed from statistics of visited states along trajectories to the steady state, then the trajectories derived from dynamical perturbations do not only reflect the underlying pseudo-landscape of the GRN. Instead, they also include contributions from the perturbations themselves.

We agree that the landscape displayed Fig. 6C integrates contributions from the perturbations on the GRN’s behavior, and that it can shape the landscape in various ways, for instance affecting the paths that are accessible, the shape/depth of certain valleys, etc. But we believe that qualitatively or quantitatively analyzing the effect of these perturbations on the landscape is precisely what is interesting here: it might help 1) understand how a system respond to a range of perturbations and to visualize which behaviors are robust to those perturbations, 2) design better strategies for manipulating those systems to produce certain behaviors

(2) In Fig. 7, I'm not sure how much is possible to take away from the results as given here, as they depend sensitively on the cohort of 432 (GRN, Z) pairs used. The comparison against random networks is well-motivated. However, as the authors note, comparison between organismal categories is more difficult due to low sample size; for instance, the "plant" and "slime mold" categories each only have 1 associated GRN. Additionally, the "n/a" category is difficult to interpret.

We acknowledge that this part is speculative as stated in the paper: “the surveyed database is relatively small with respect to the wealth of available models and biological pathways, so we can hardly claim that these results represent the true distribution of competencies across these organism categories”. However, when further data is available, the same methodology can be reused and we believe that the resulting statistical analyses could be very informative to compare organismal (or other) categories.

(3) In Fig. 8, it is unclear whether the behavioral catalog generated is important to the intervention design problem of moving a system from one attractor basin to another. The authors note that evolutionary searches or SGD could also be used to solve the problem. Is the analysis somehow enabled by the behavioral catalog in a way that is complementary to those methods? If not, comparison against those methods (or others e.g. optimal control) would strengthen the paper.

We thank the reviewer for asking to clarify this point, which might not be clearly explained in the paper. Here the behavioral catalog is indeed used in a complementary way to the optimization method, by identifying a representative set of reachable attractors which are then used to define the optimization problem. For instance here, thanks to the catalog, we 1) were able to identify a “disease” region and several possible reachable states in that region and 2) use several of these states as starting points of our optimization problem, where we want to find a single intervention that can successfully and robustly reset all those points, as illustrated in Figure 8. Please note that given this problem formulation, a simple random search was used as an optimization strategy. When we mention more advanced techniques such as EA or SGD, it is to say that they might be more efficient optimizers than random search. However, we agree that in many cases optimizing directly will not work if starting from random or bad initial guess, and this even with EA or SGD. In that case the discovered behavioral catalog can be useful to better initialize this local search and make it more efficient/useful, akin to what is done in Figure 9.

(4) The analysis presented in Fig. 9 also is preliminary. The authors note that there exist many algorithms for choosing/identifying the parameter values of a dynamical system that give rise to a desired time-series. It would be a stronger result to compare their approach to more sophisticated methods, as opposed to random search and SGD. Other options from the recent literature include Bayesian techniques, sparse nonlinear regression techniques (e.g. SINDy), and evolutionary searches. The authors note that some methods require fine-tuning in order to be successful, but even so, it would be good to know the degree of fine-tuning which is necessary compared to their method.

We agree that the analysis presented in Figure 9 is preliminary, and thank the reviewer for the suggestion. We would first like to refer to other papers from the ML literature that have more thoroughly analyzed this issue, such as Colas et al. [74] and Pugh et al. [34], and shown the interest of diversity-driven strategies as promising alternatives. Additionally, as suggested by the reviewer, we added an additional comparison to the CMA-ES algorithm in the revised version in order to complete our analysis. CMA-ES is an evolutionary algorithm which is self-adaptive in the optimization steps and that is known to be better suited than SGD to escape local minimas when the number of parameters is not too high (here we only have 15 parameters). However, our results showed that while CMA-ES explores more the solution space at the beginning of optimization than SGD does, it also ultimately converges into a local minima similarly to SGD. The best solution converges toward a constant signal (of the target b) but fails to maintain the target oscillations, similar to the solutions discovered by gradient descent. We tried this for a few hyperparameters (init mean and std) but always found similar results. We have updated the figure 9 image and caption, as well as descriptive text, to

include these novel results in the revised version. We also added a reference to the CMA-ES paper in the citations.

Reviewer #1 (Recommendations For The Authors):

I would suggest to conduct a more rigor analysis of the performance by estimating/approximating the ground truth robust goal sets in important GRNs.

Also, the use of terminology from different disciplines can be improved. Please see my comments above. Specifically, the connection between controllability in dynamical control systems and versatility used in this paper is unclear.

We hope to have addressed the reviewer's concerns in our previous answers.

Reviewer #2 (Recommendations For The Authors):

Fig 4b: I'm not sure if DBSCAN is the appropriate method to use here, as the visual focus on the core elements of the clusters downplays the full convex hull of the points that random sampling achieves in Z space. An analysis based on convex hulls or the ball-coverage from Fig. 3b would presumably generate plots that were more similar between random sampling and curiosity search. If the goal is to highlight redundancy/non-linearity in the mapping between Z and I, another approach might be to simply bin Z-space in a grid, or to use a clustering algorithm that is less stringent about core/noise distinctions.

We thank the reviewer for the suggestion. This plot is intended to convey the reader an understanding of why a method that uniformly samples goals in Z (what the IMGEP is doing), is more efficient than a method that uniformly samples parameters in I (what the random search is doing), in systems for which there is high redundancy/non-linearity in the mapping between I and Z. We agree that binning the Z-space in a grid and counting the number of achieved bins is a way to quantitatively measure this, which is by the way very close to what we do in Figure 3 for measuring the achieved diversity. We believe however that the clustering and coloring provides additional intuitions on why this is the case: it illustrates that large regions of the intervention space map to small regions in the outcome space and vice versa.

Additional changes in the revised version:

We added a sentence in the Methods section as well as in the caption of Table S1 providing additional details about the way we simulate the biological models from the BioModels website

We fixed a wrong reference to Figure 4 in the Methods “Sensitivity measure” subsection with reference to Figure 5.

<https://doi.org/10.7554/eLife.92683.3.sa0>