Curiosity-driven AI for Science:
Automated Discovery of Self-Organized Structures

Mayalen Etcheverry

Academic Advisors: Pierre-Yves Oudeyer & Clément Moulin-Frier
Industrial Supervisor: Marc Nicodème

Self-Organization and the Evolution of Forms

In the Inorganic World

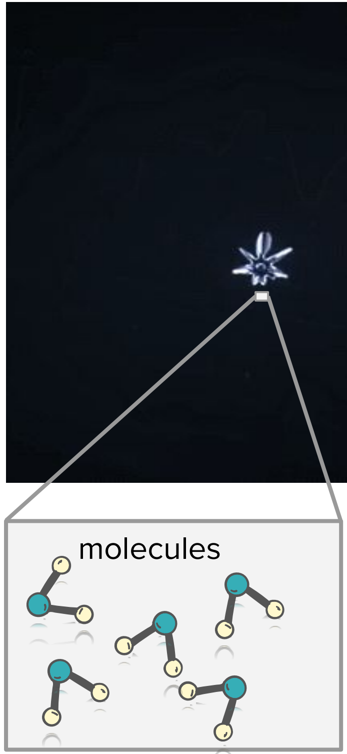

In the Living World

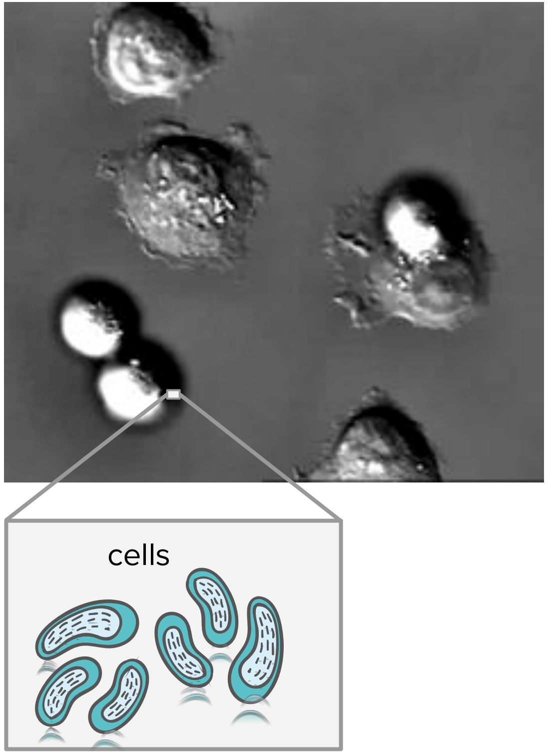

In the Artificial World

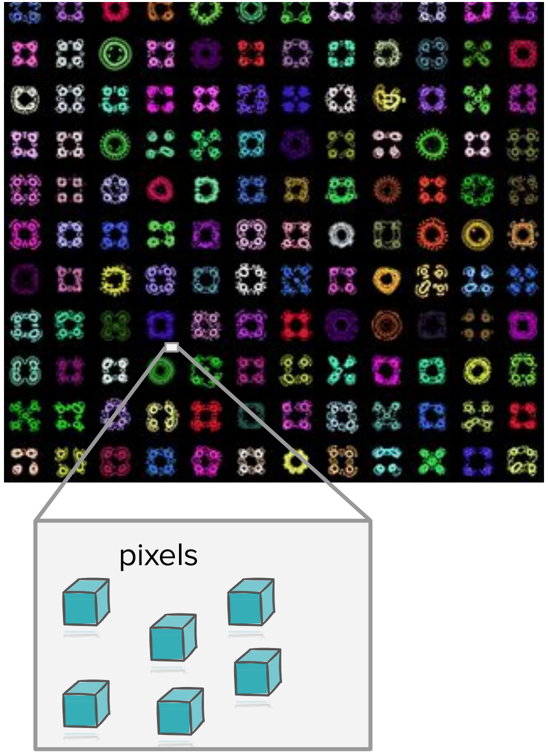

Discovery of Novel Self-Organized Structures

In Physics and Chemistry

Origins of Life
(Grizou et al., 2020)

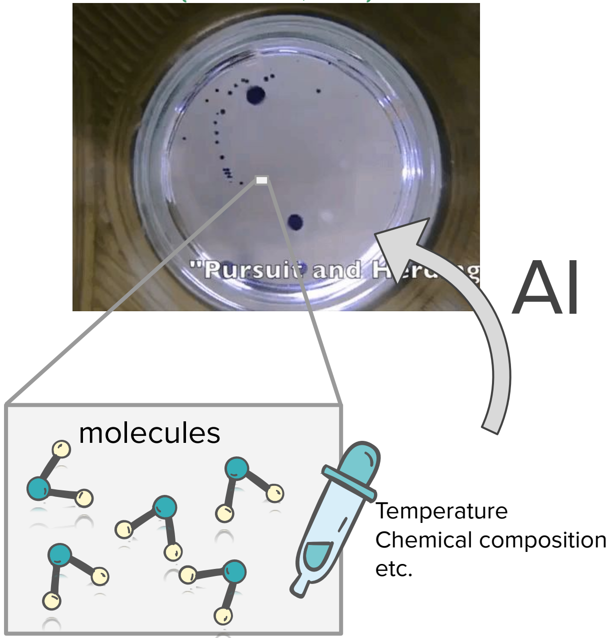

In Biology

Tissue Engineering
(Abellan Lopez et al., 2023)

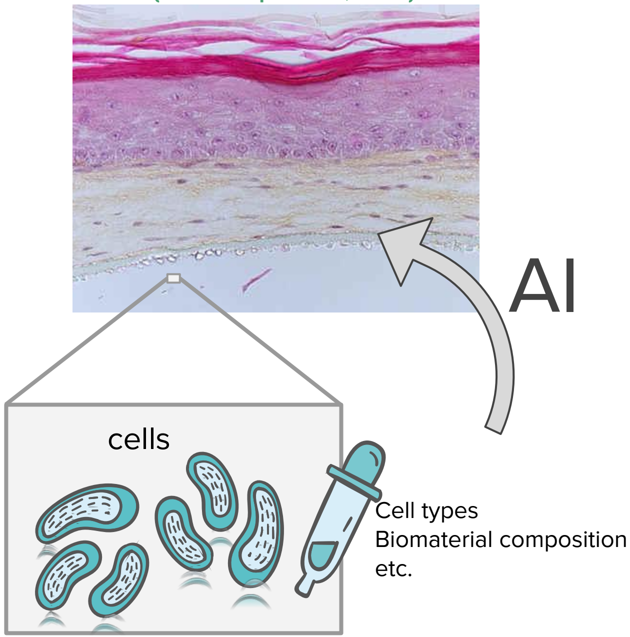

In ALife and AI

Emergence of Agency
(Hamon et al., 2022)

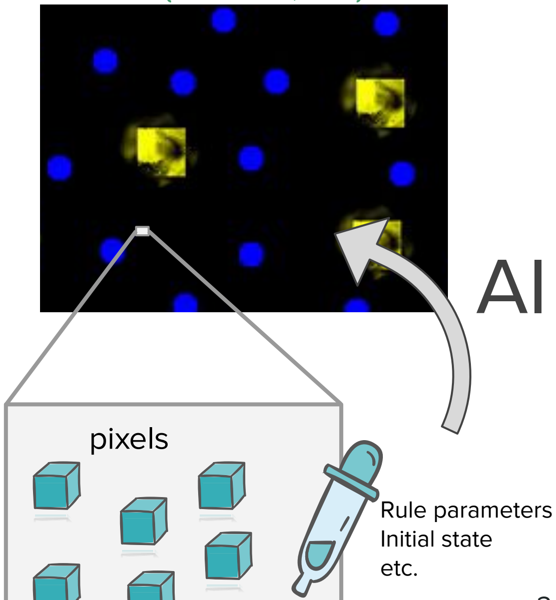

Automated Discovery of Self-Organized Structures

How to make <u>interesting</u> discoveries in a <u>sample-efficient</u> manner?

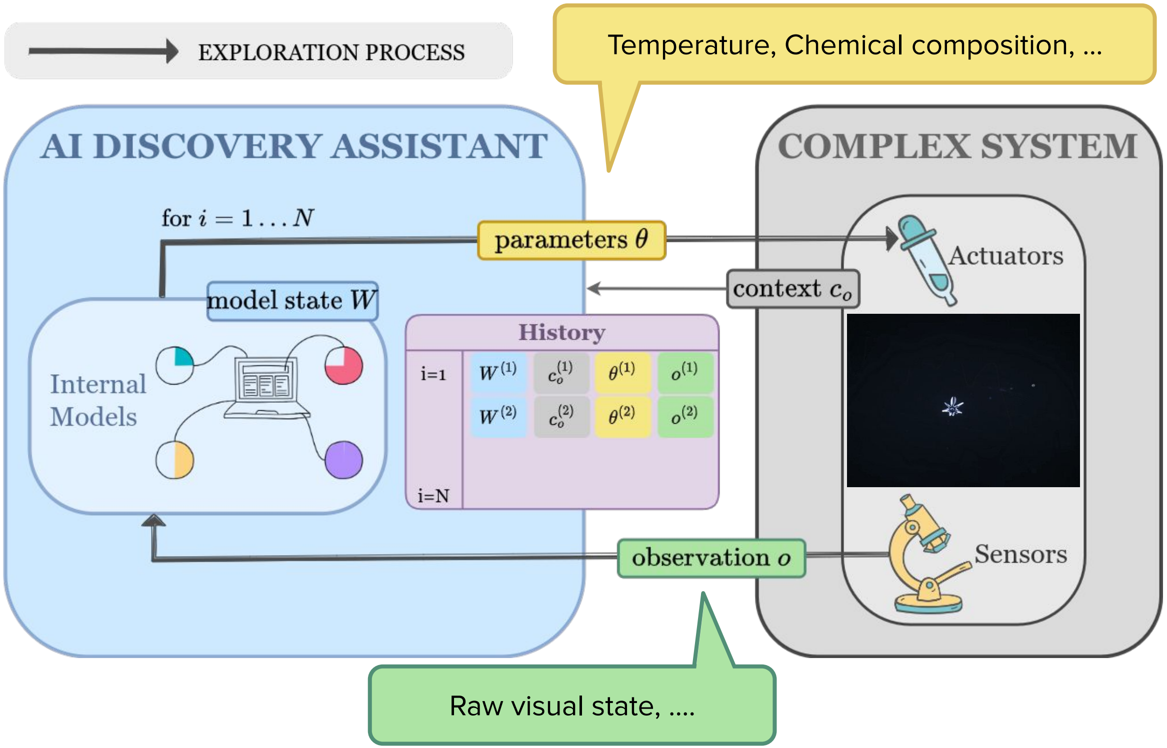

For i = 1 ... $\boxed{N}$: ≪standard AI training budget

θ_1	o_1	W_1
θ_2	o_2	W_2
...	...	...
θ_N	o_N	W_N

database of discoveries — learned representations

Automated Discovery of Self-Organized Structures

How to make interesting discoveries in a sample-efficient manner?

Optimization-driven Search

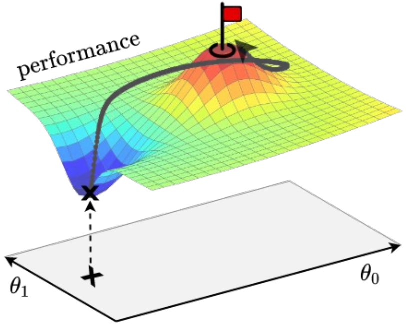

Aim: optimization toward target

Hypothesis: reward function

Approach: evolutionary algorithms, gradient descent, bayesian optimization

➜ sparse and deceptive reward problem

Knowledge-driven Search

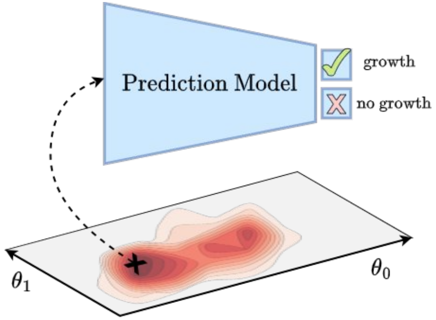

Aim: learn a predictive model

Hypothesis: base model architecture

Approach: active learning (prediction error, max information gain, etc.)

➜ cold start and skewed data problems

Diversity-driven Search

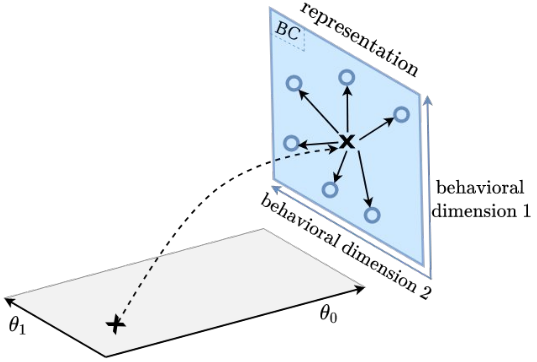

Aim: construct a map of possible outcomes

Hypothesis: behavioral characterization

Approach: novelty search, intrinsically motivated goal exploration process

Automated Discovery of Self-Organized Structures

How to make interesting discoveries in a sample-efficient manner?

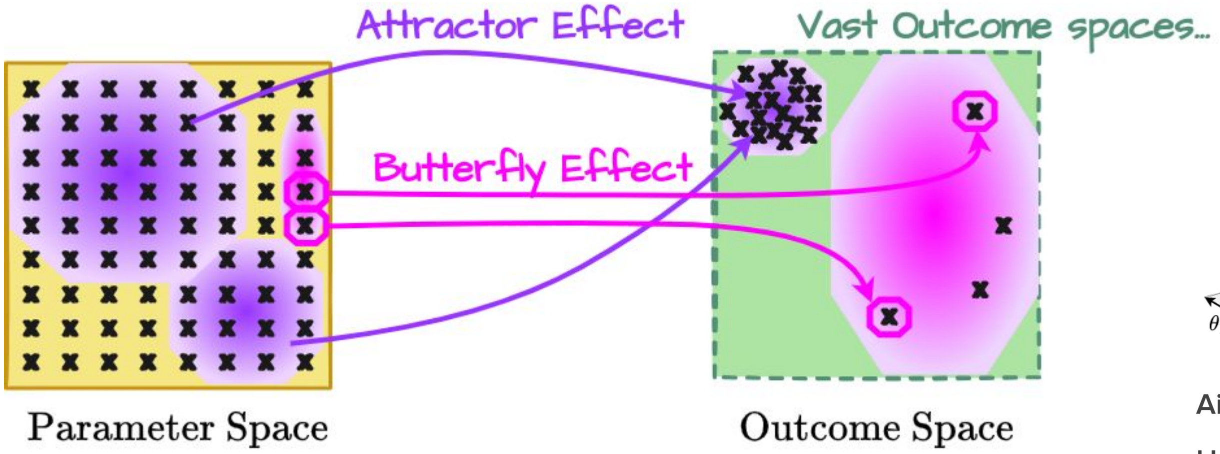

Diversity-driven Search

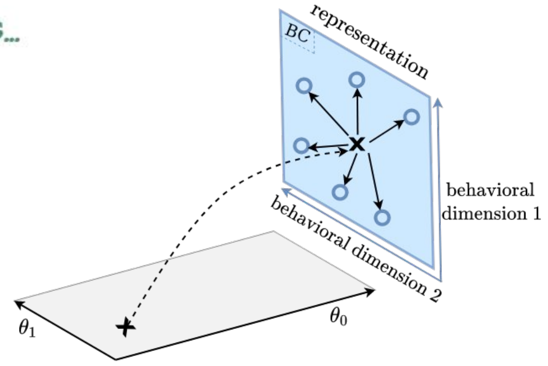

Aim: construct a map of possible outcomes

Hypothesis: behavioral characterization

Approach: novelty search, intrinsically motivated goal exploration process

Automated Discovery of Self-Organized Structures

How to make <u>interesting</u> discoveries in a <u>sample-efficient</u> manner?

Developmental AI

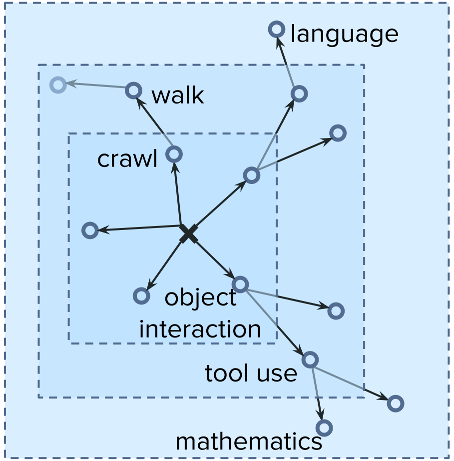

"Curious" child during exploratory play

Credits: Francis Vachon

Humans acquire open-ended repertoire of skills throughout their lifetimes despite constraints in time and energy

Automated Discovery of Self-Organized Structures

How to make interesting discoveries in a sample-efficient manner?

Developmental sciences

Developmental AI

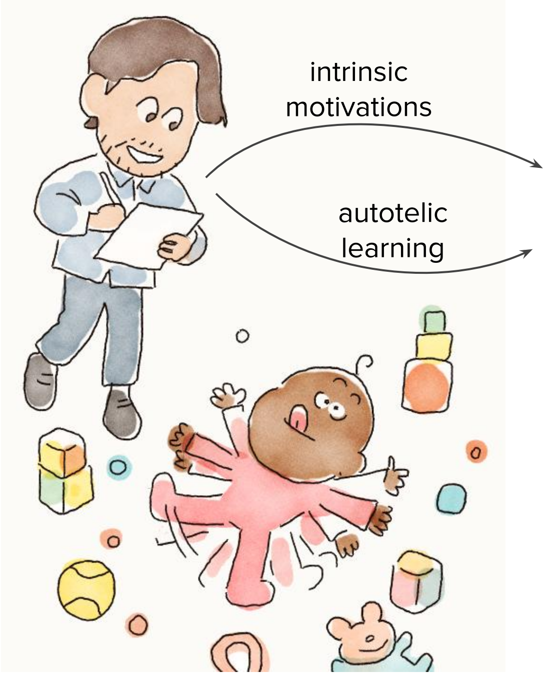

Credits: Marie Spenale

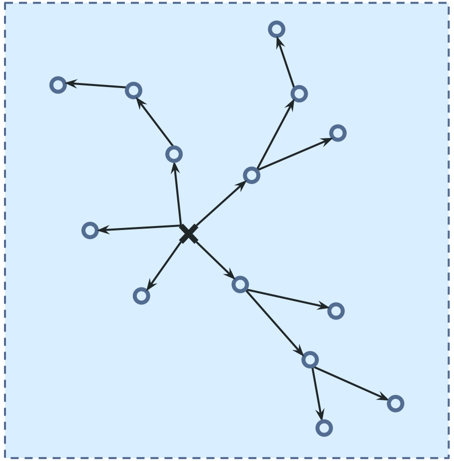

Intrinsic motivations:
Set of brain processes that motivate humans to explore for the mere purpose of experiencing novelty, suprise or learning progres

Autotelic learning: *auto (self) + telos (goal)*
Autotelic agents are intrinsically motivated to learn to represent, generate, pursue and master their own goals.

Humans acquire open-ended repertoire of skills throughout their lifetimes despite constraints in time and energy

Automated Discovery of Self-Organized Structures

How to make interesting discoveries in a sample-efficient manner?

Developmental sciences

Developmental AI

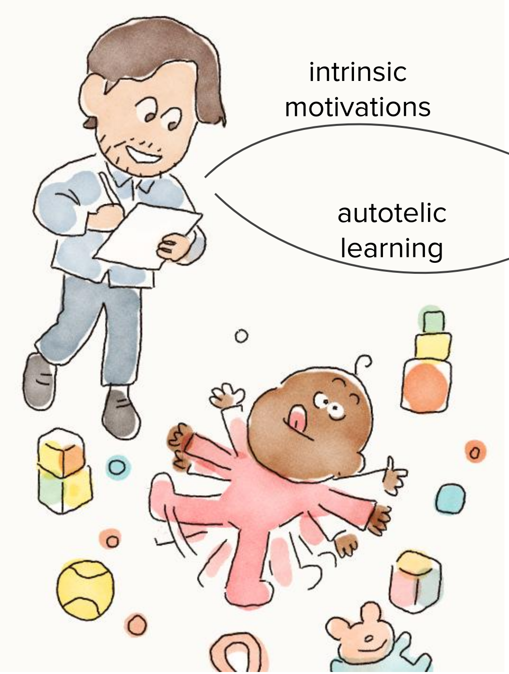

intrinsic motivations

autotelic learning

computational models

AI toolbox

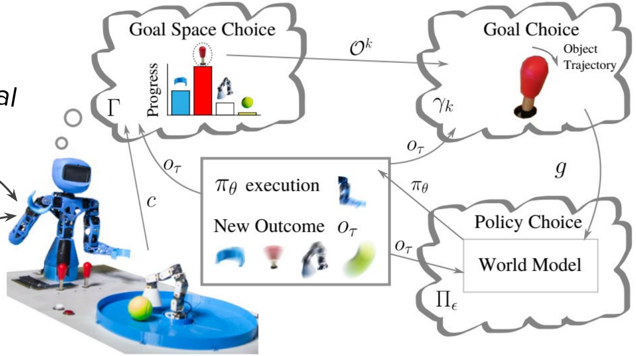

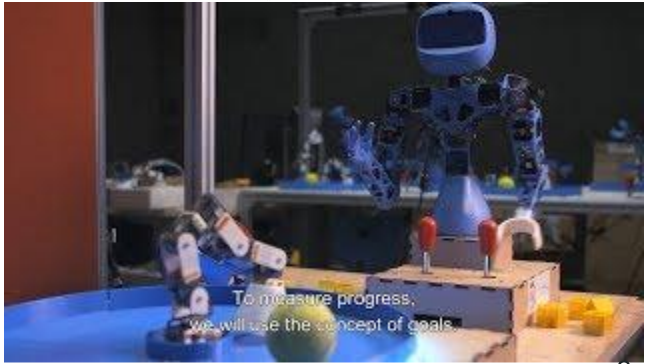

(Forestier et al., 2022)

Credits: Marie Spenale

Automated Discovery of Self-Organized Structures

How to make <u>interesting</u> discoveries in a <u>sample-efficient</u> manner?

Developmental sciences | Developmental AI | Assist Scientific Discovery

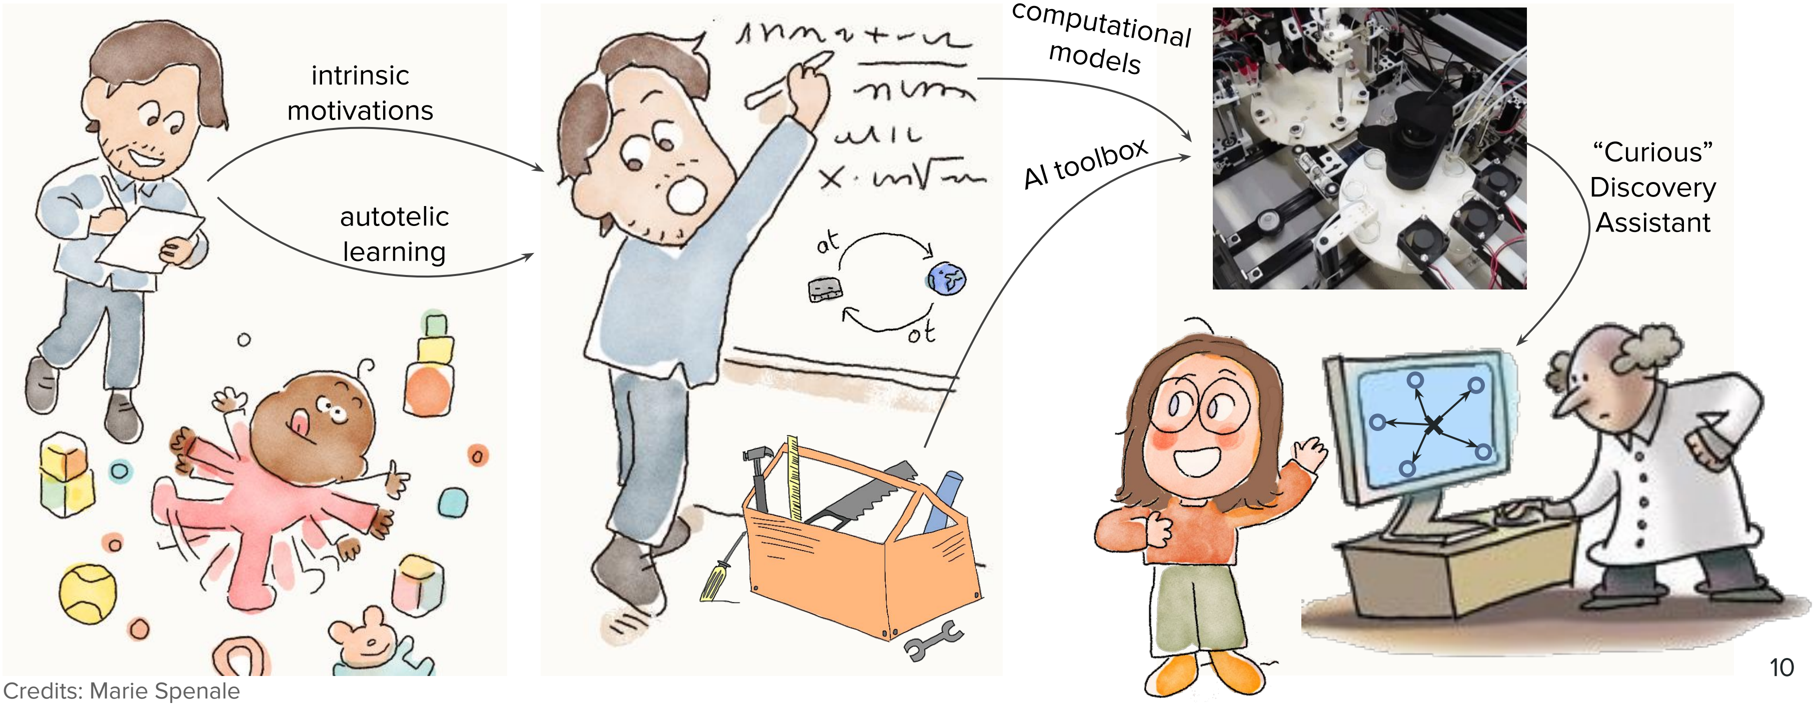

Credits: Marie Spenale

Outline

I. The "Curious Discovery Assistant" Framework

BACKGROUND

Computational Framework
IMGEP

Testbed Environment
Lenia

Conceptual contribution

Meta-Diversity Search

Computational study

IMGEP-HOLMES

II. Use Cases of the Curious Discovery Assistant

Use Case #1

Sensorimotor Agency in
Continuous CA

Use Case #2

Competencies in Biological
Network Models

PERSPECTIVES

Use Case #3
Bioprinter-controlled System

Towards Open-Ended
Discovery Assistants

Part 1: The Curious Discovery Assistant Framework

IMGEP: Intrinsically Motivated Goal Exploration Process

Forestier, et al., "Intrinsically Motivated Goal Exploration Processes with Automatic Curriculum Learning", JMLR (2022)

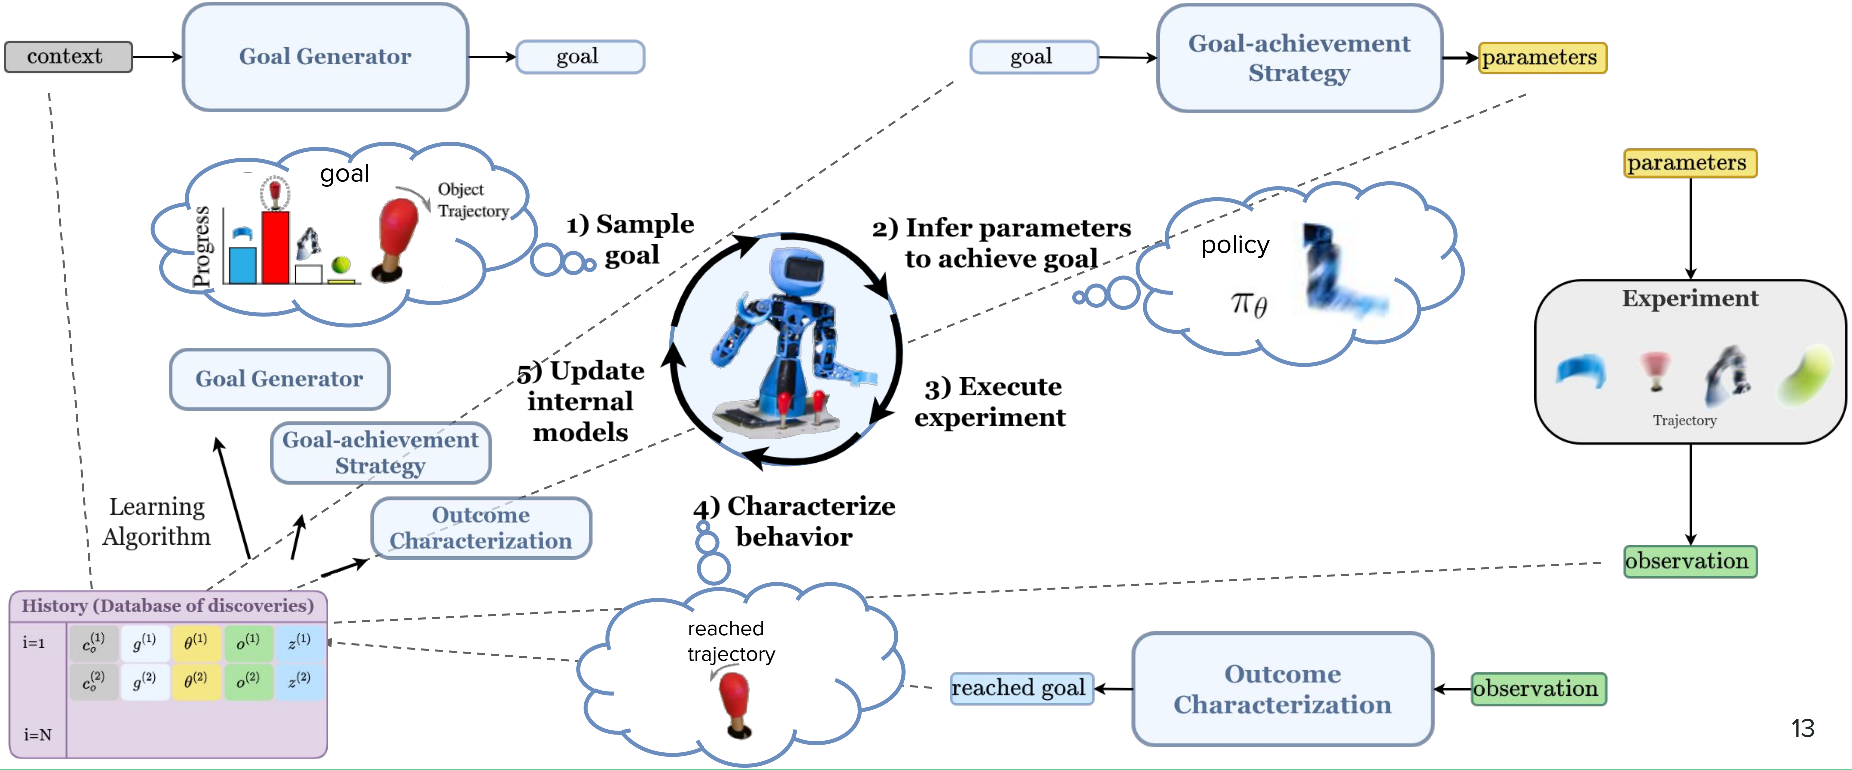

IMGEP: Intrinsically Motivated Goal Exploration Process

Grizou, et al., “A curious formulation robot enables the discovery of a novel protocell behavior”, Science (2020)

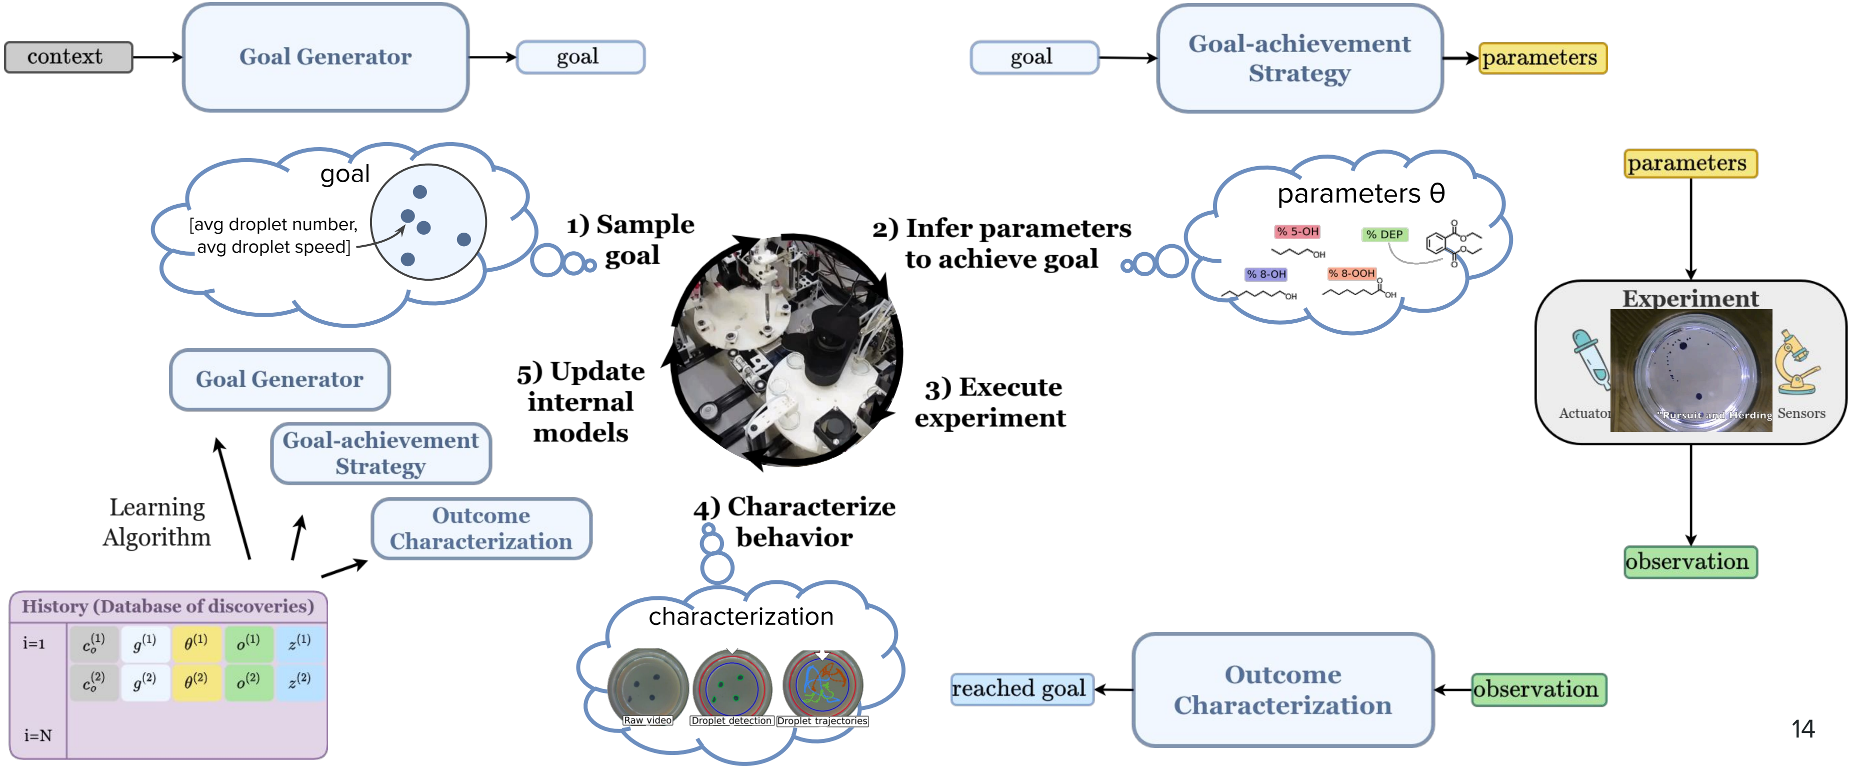

IMGEP: Intrinsically Motivated Goal Exploration Process

Grizou, et al., "A curious formulation robot enables the discovery of a novel protocell behavior", Science (2020)

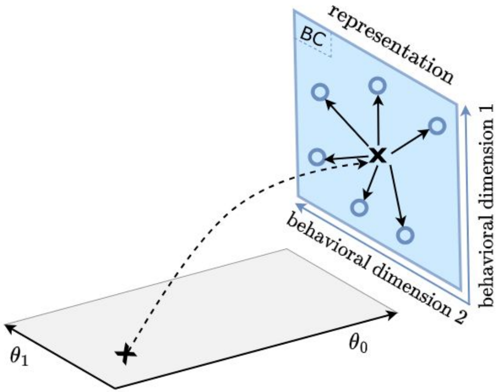

Diversity Search

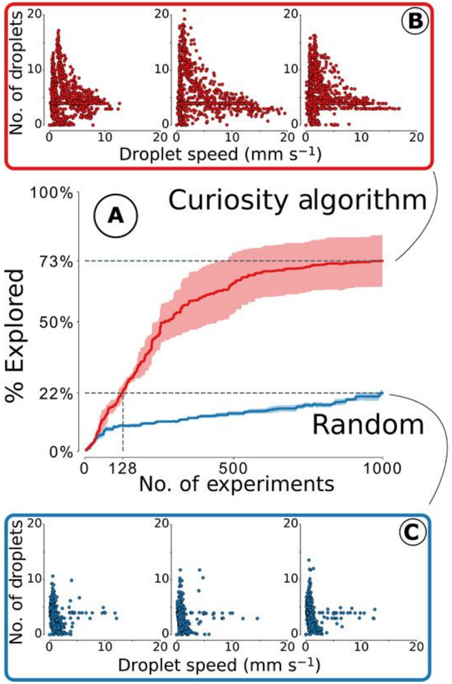

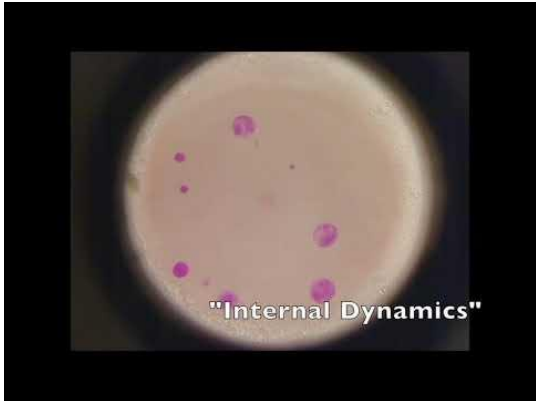

Diverse "life-like" behaviors:

- Movement
- Grouping
- Division
- Fusion
- Chemotaxis

IMGEP: Intrinsically Motivated Goal Exploration Process

Grizou, et al., “A curious formulation robot enables the discovery of a novel protocell behavior”, Science (2020)

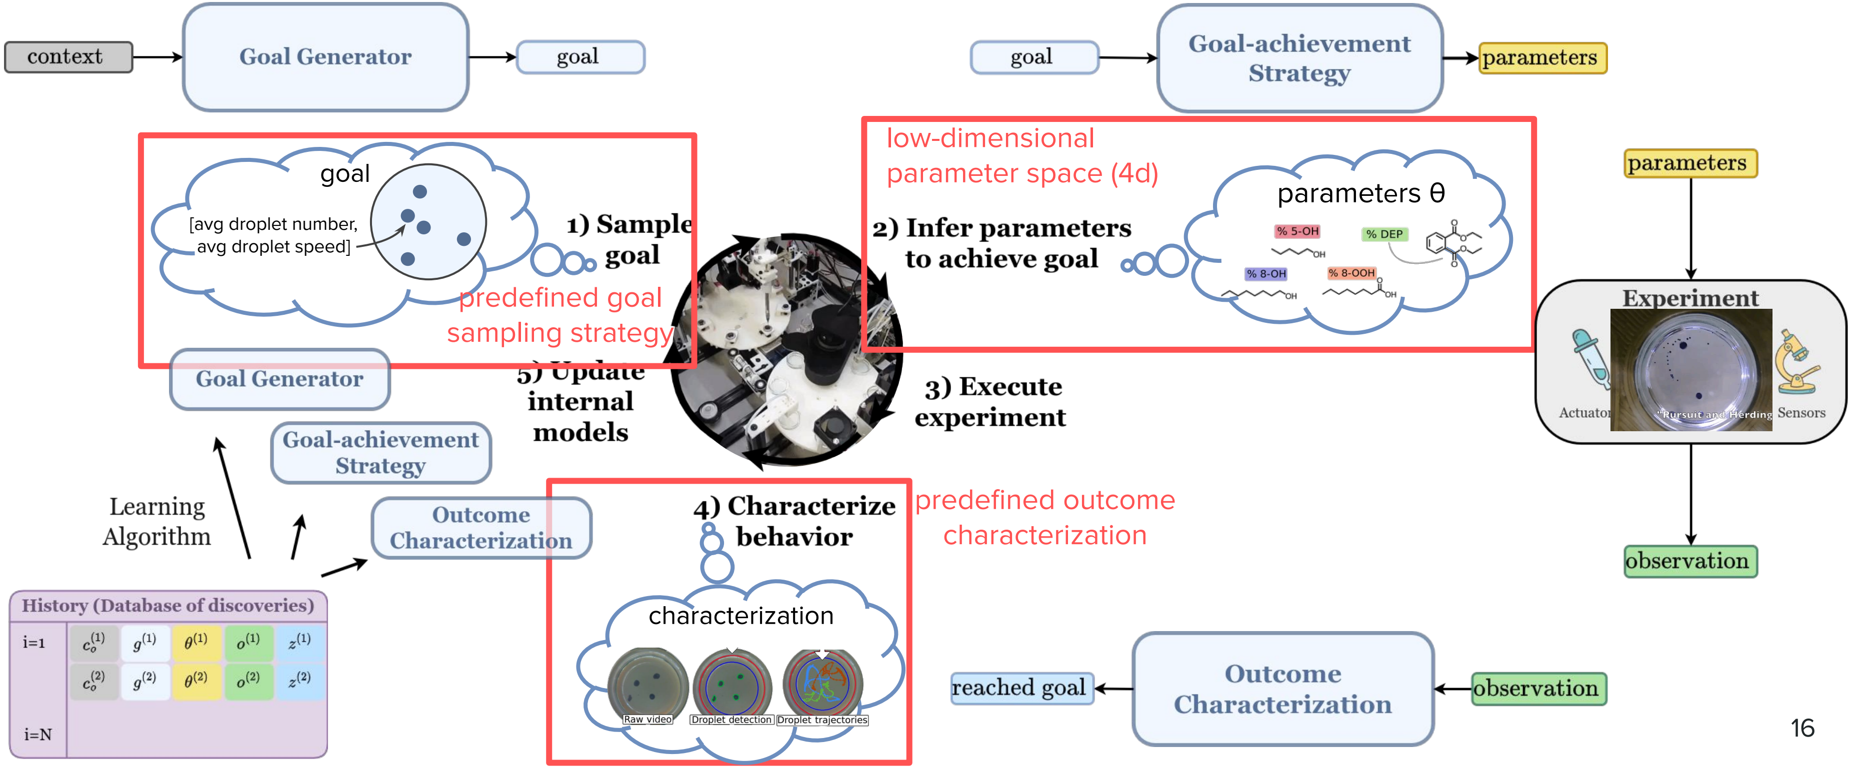

Lenia: Testbed Environment

- Generalized version of Conway's Game of Life (Chan 2019, Chan 2020)

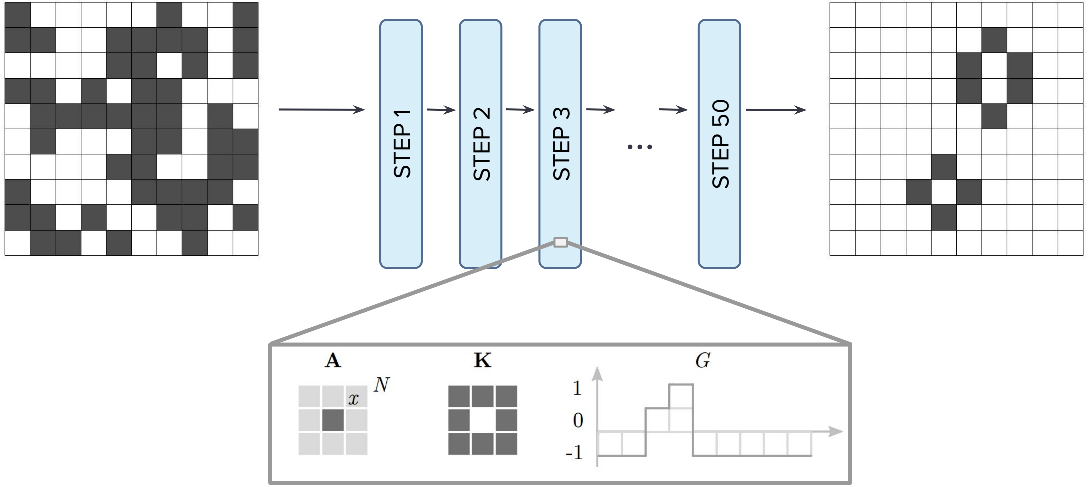

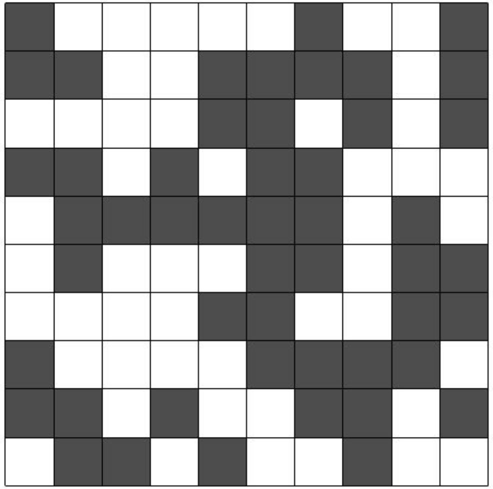

Lenia: Testbed Environment

- Generalized version of Conway's Game of Life (Chan 2019, Chan 2020)
- Class of continuous CA where each instance is defined by some parameters that condition the CA "physics"

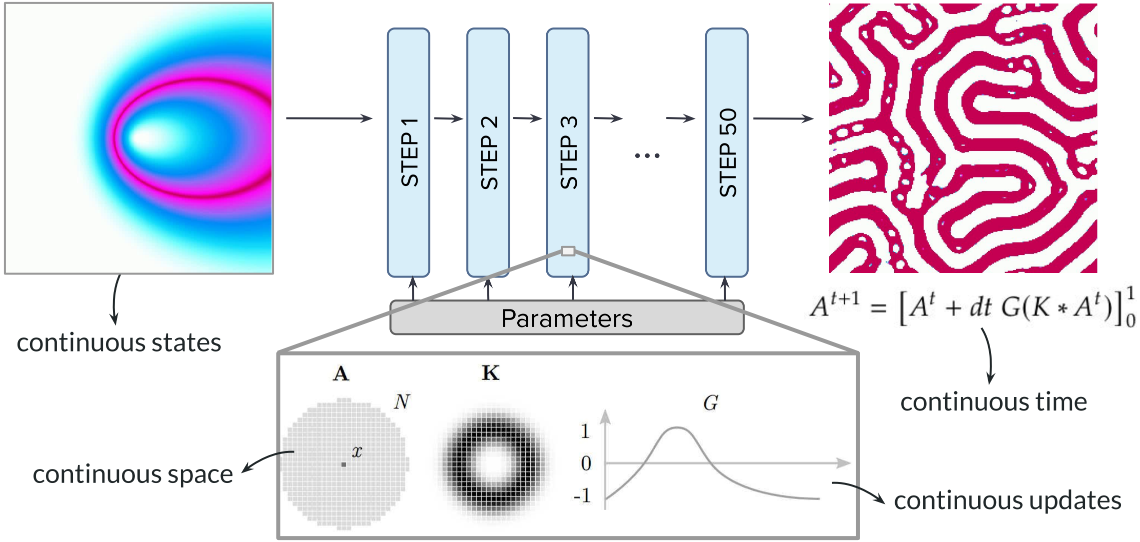

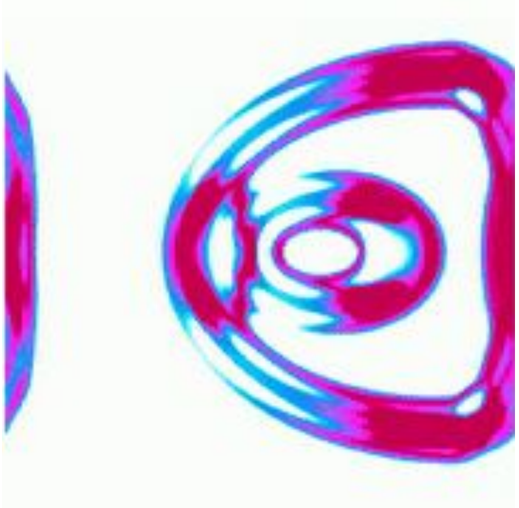

Lenia: Testbed Environment

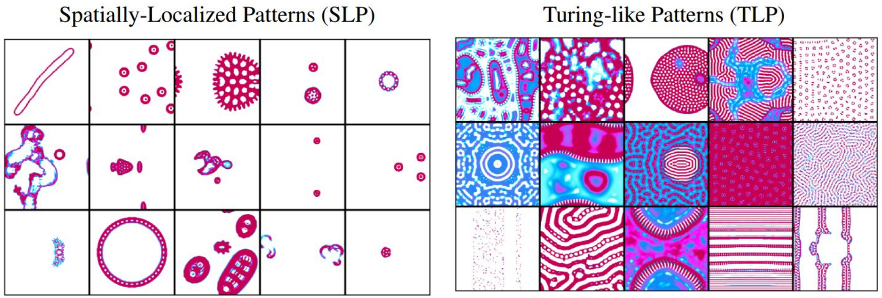

- Can generate a wide range of complex structures (unbounded emergence)

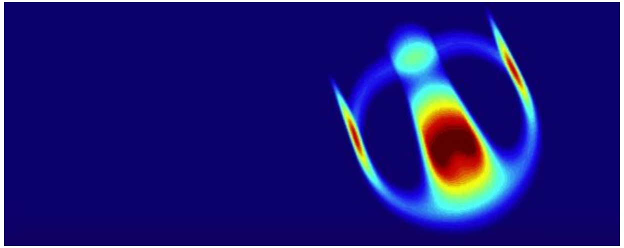

- Interesting life-like properties
  - spatially localized, symmetries
  - Individuality, diverse locomotion

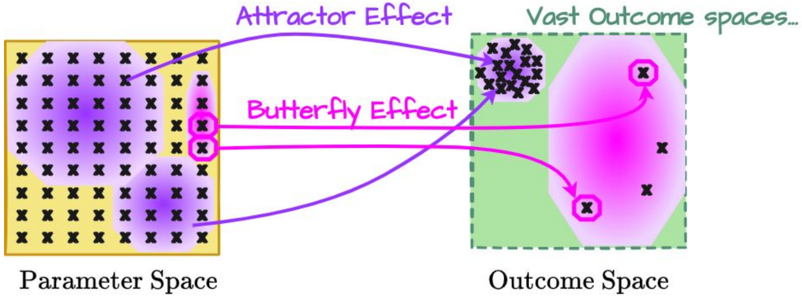

- Constructing a map of the possible outcomes poses various exploration challenge
  - complex system mapping
  - raw visual states

➔ **computer-based yet rich testbed for automated discovery**

Meta Diversity Search

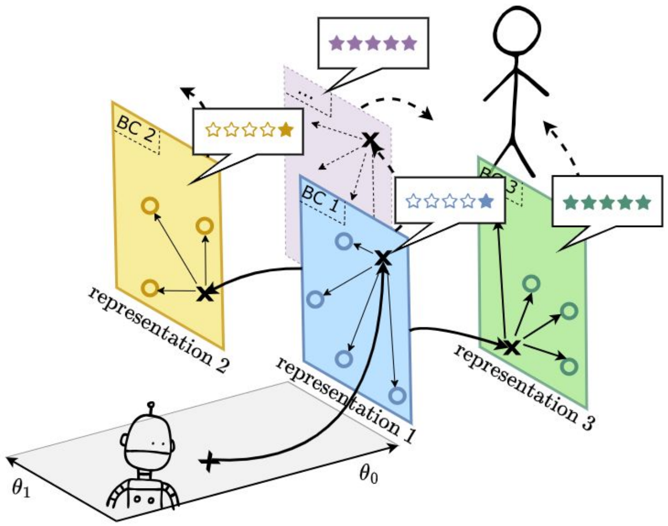

Conceptual Contribution

How to define the representation space?

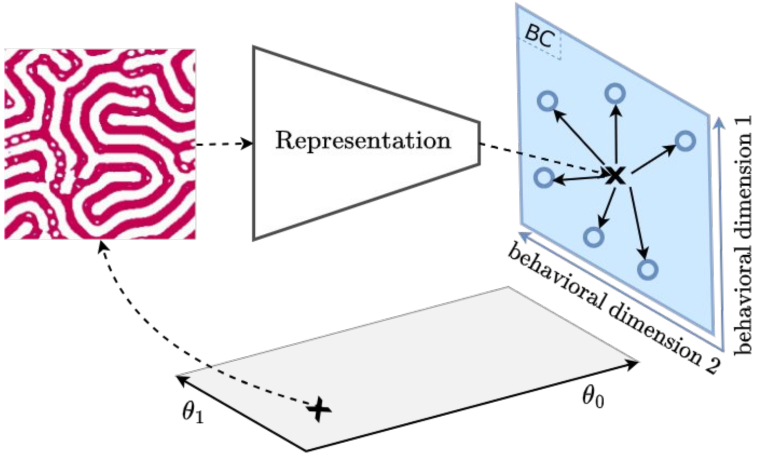

Diversity Search

- Engineered features

➔ prior expertise on "interesting" high-level descriptors

- Unsupervised learned features

➔ automatically learn encoder representation with VAE

➔ requires pre-collected set of observations

- Online learned features

➔ online learning of encoder representation with VAE

"Intrinsically Motivated Exploration for Automated Discovery of Patterns in Morphogenetic Systems", Chris Reinke, Mayalen Etcheverry, Pierre-Yves Oudeyer. ICLR 2020 (Oral)

How to define the representation space?

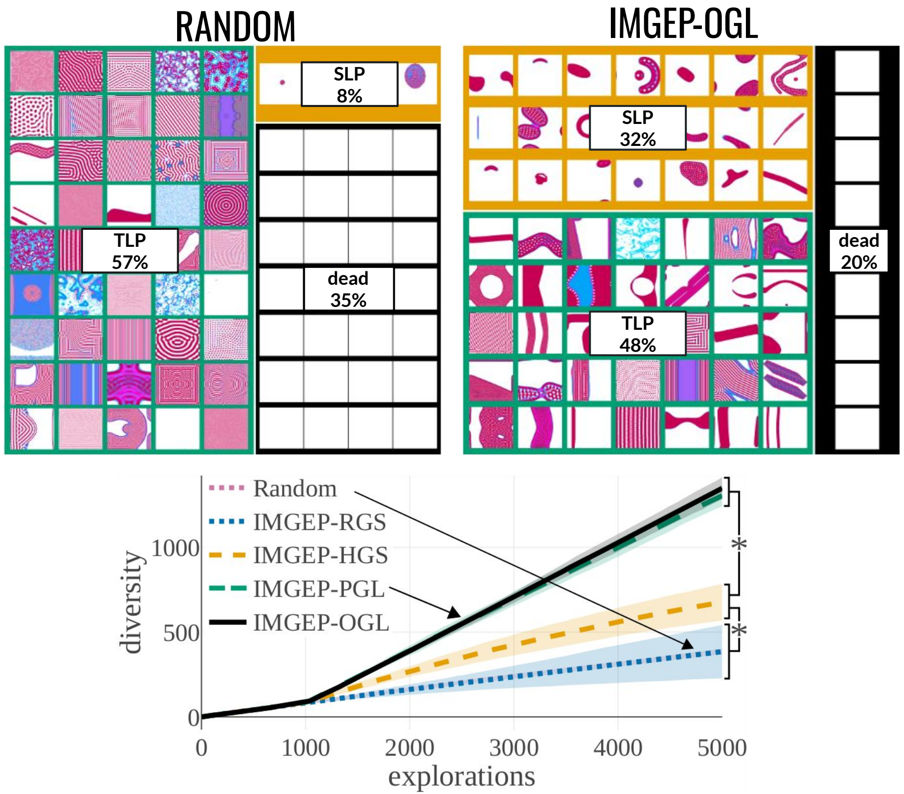

- Engineered features (IMGEP-HGS)

➜ prior expertise on "interesting" high-level descriptors

- Unsupervised learned features (IMGEP-PGL)

➜ automatically learn encoder representation with VAE

➜ fixed representation + requires pre-collected set of observations

- Online learned features (IMGEP-OGL)

➜ online learning of encoder representation with VAE

"Intrinsically Motivated Exploration for Automated Discovery of Patterns in Morphogenetic Systems", Chris Reinke, Mayalen Etcheverry, Pierre-Yves Oudeyer. ICLR 2020 (Oral)

Limits of monolithic representations

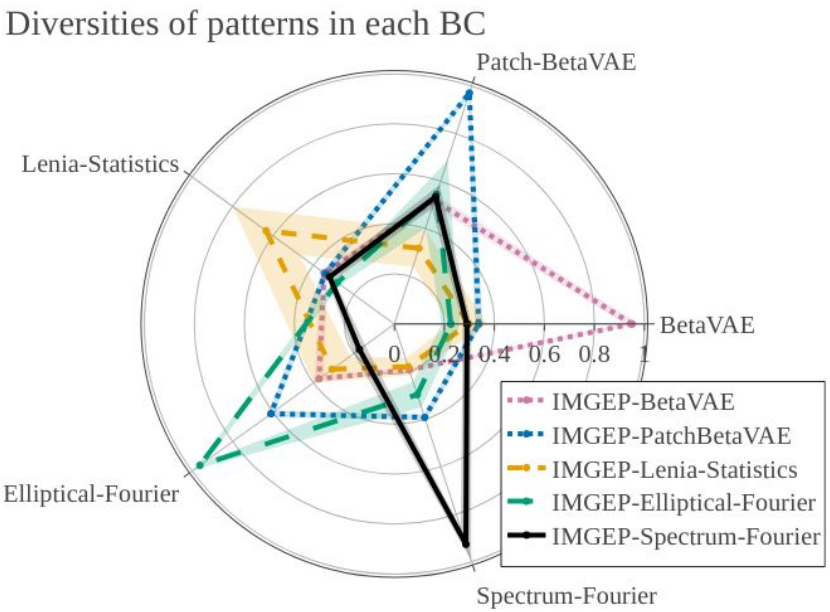

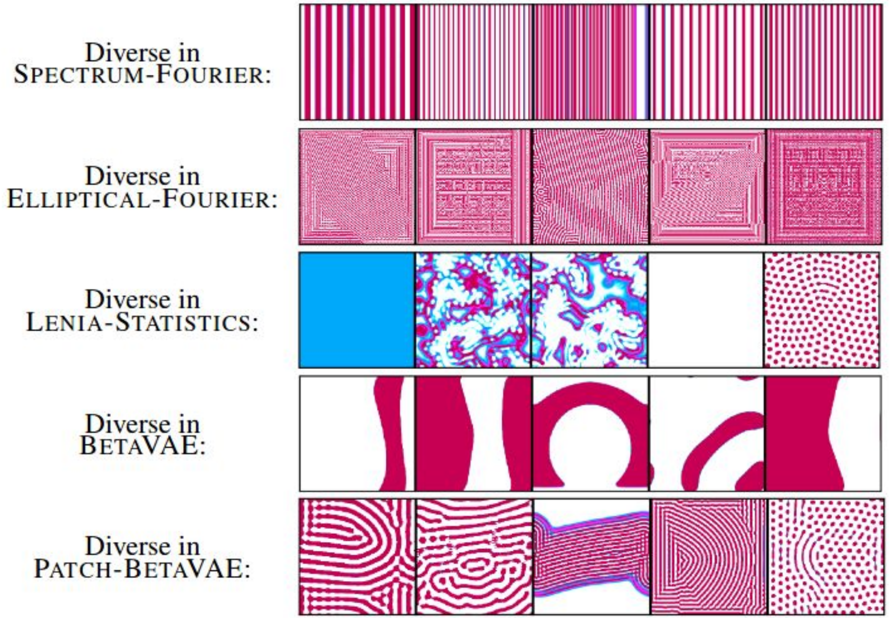

→ **unlikely to be aligned with what a final end-user is considering as "interesting"**

Meta-Diversity Search

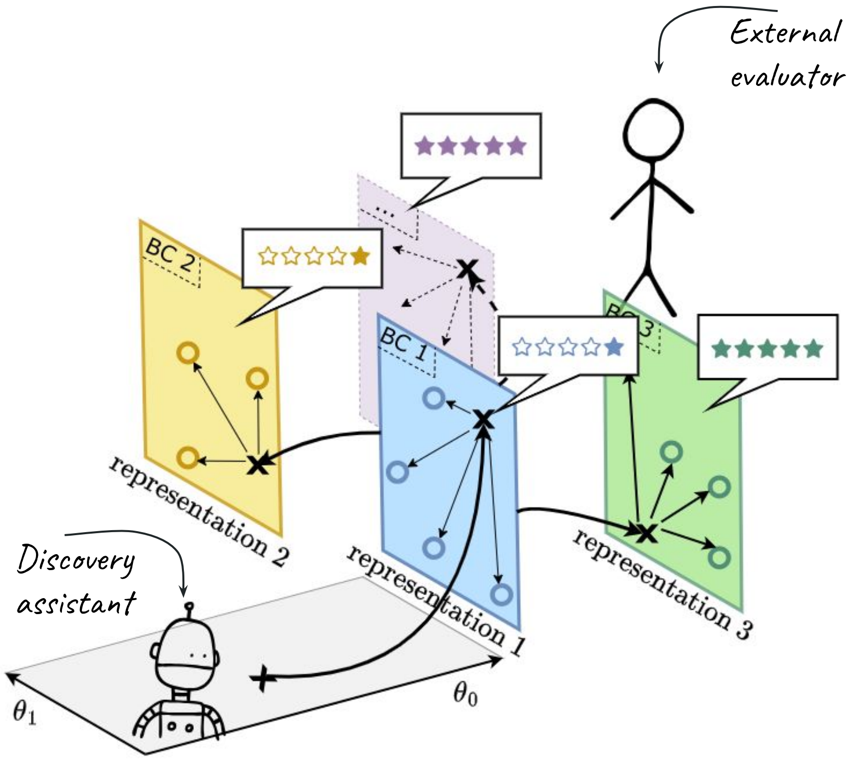

Outer loop: continually learns diverse representation spaces to characterize behaviors

1) *How to learn diverse representation spaces?*

Inner loop: searches for a maximally diverse set of patterns in each characterization space

2) *How to efficiently find diverse patterns in the learned spaces?*

➔ steer the search toward end-user preferences

3) *How to quickly adapt the search toward initially-unknown preferences of human end-user?*

IMGEP-HOLMES

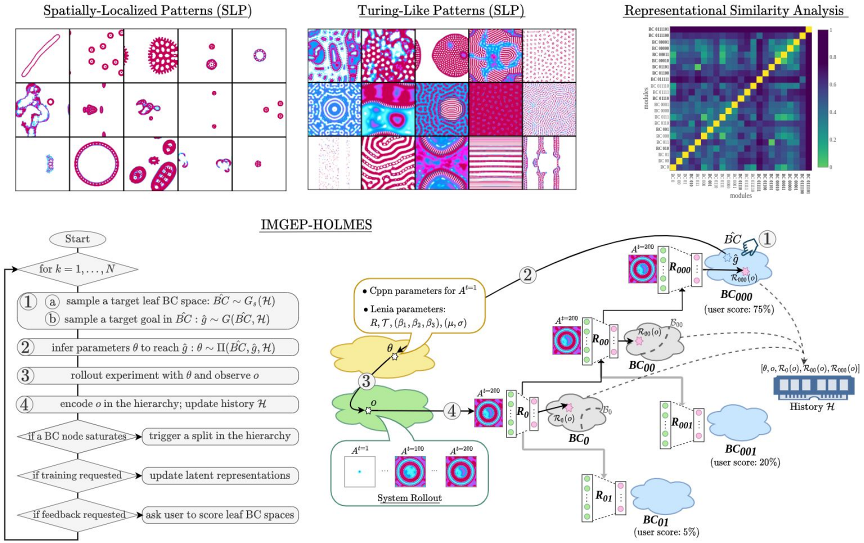

Computational Study

“Hierarchically Organized Latent Modules for Exploratory Search in Morphogenetic Systems”, Mayalen Etcheverry, Clément Moulin-Frier, Pierre-Yves Oudeyer
NeurIPS 2020 (Oral)

HOLMES: Learning Diverse Representation Spaces

1) How to learn diverse representation spaces?

Hierarchically Organized Latent Modules for Exploratory Search (HOLMES)
↳ dynamic and modular architecture actively expanded to represent the different niches

- Base **module** embedding neural network → *VAE*

- **Split trigger** → *reconstruction loss plateau*

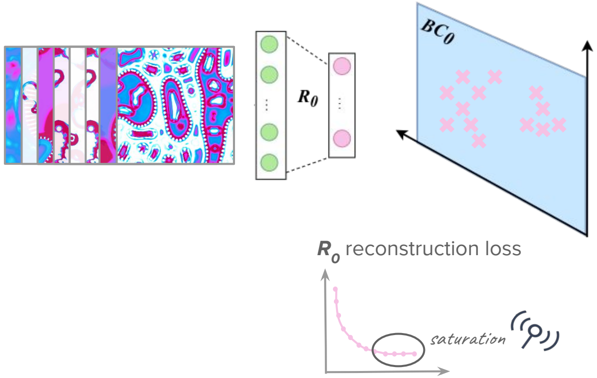

HOLMES: Learning Diverse Representation Spaces

1) *How to learn diverse representation spaces?*

Hierarchically Organized Latent Modules for Exploratory Search (HOLMES)
↳ dynamic and modular architecture actively expanded to represent the different niches

- Base **module** embedding neural network → *VAE*
- **Split trigger** → *reconstruction loss plateau*
- **Clustering** in the latent space → *K-Means*
- **Parent-child transfer** → *lateral connections*

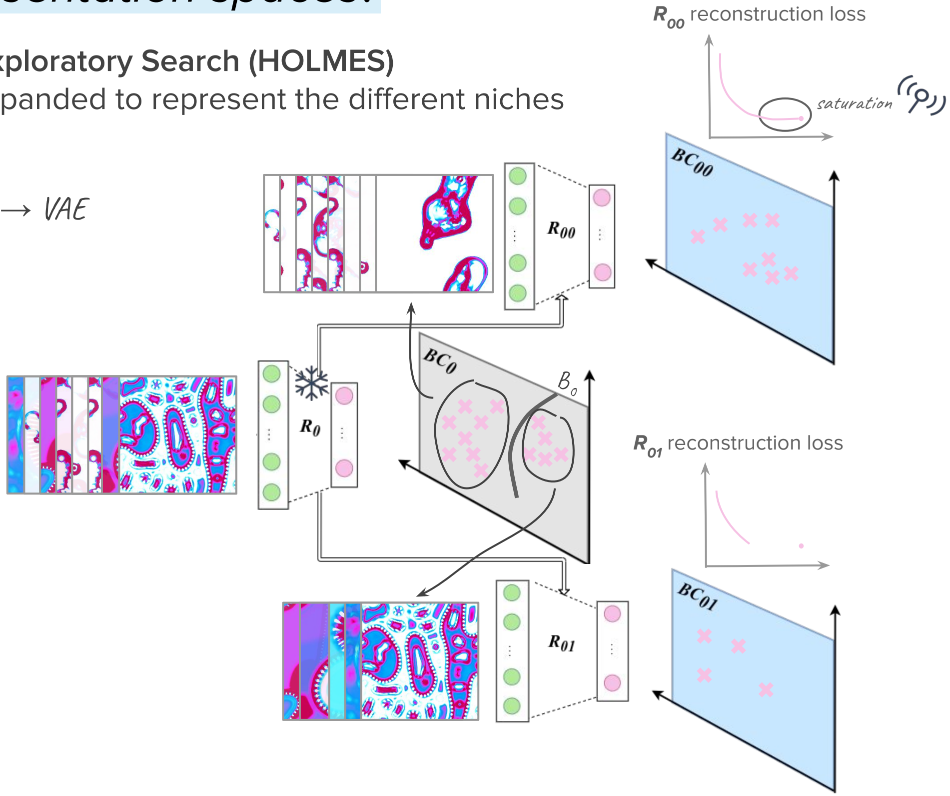

HOLMES: Learning Diverse Representation Spaces

1) *How to learn diverse representation spaces?*

Hierarchically Organized Latent Modules for Exploratory Search (HOLMES)
↳ dynamic and modular architecture actively expanded to represent the different niches

- Base **module** embedding neural network → *VAE*
- **Split trigger** → *reconstruction loss plateau*
- **Clustering** in the latent space → *K-Means*
- **Parent-child transfer** → *lateral connections*

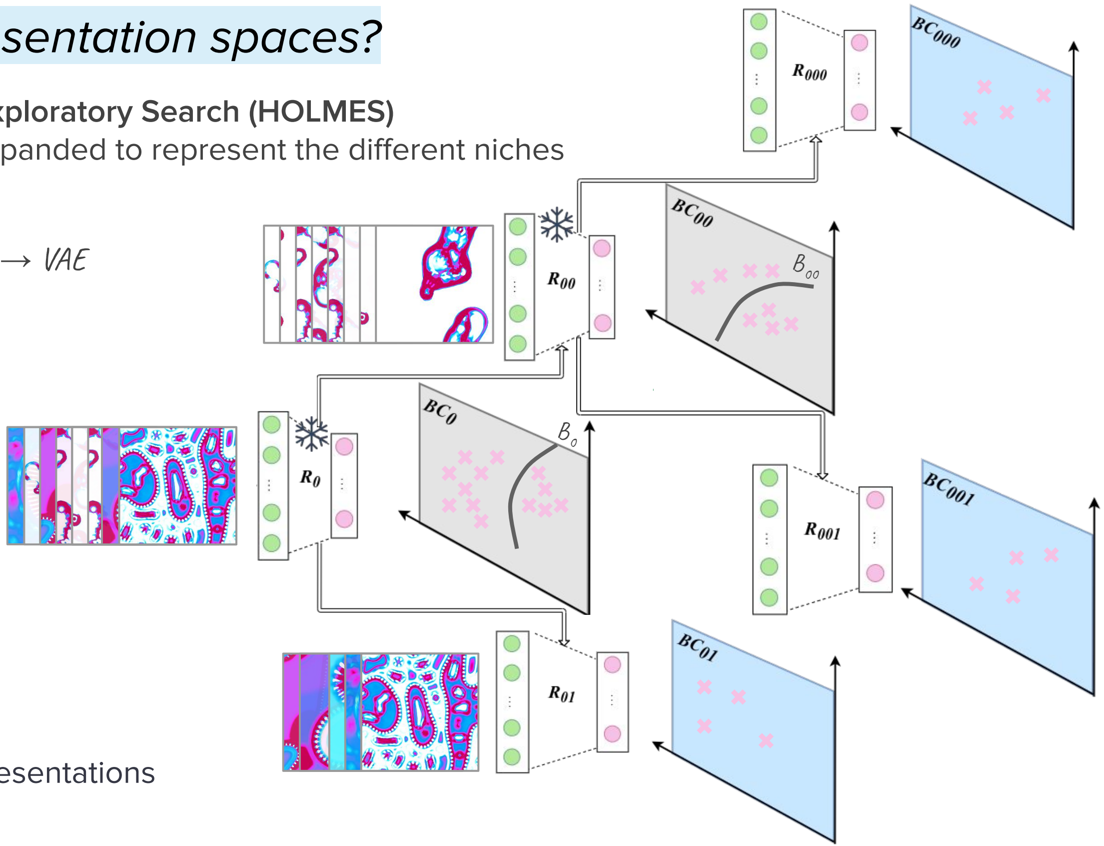

➔ progressively deeper hierarchy of **diverse** representations

IMGEP-HOLMES: Diversity Search in Learned Spaces

2) How to efficiently find diverse patterns in each representation space?

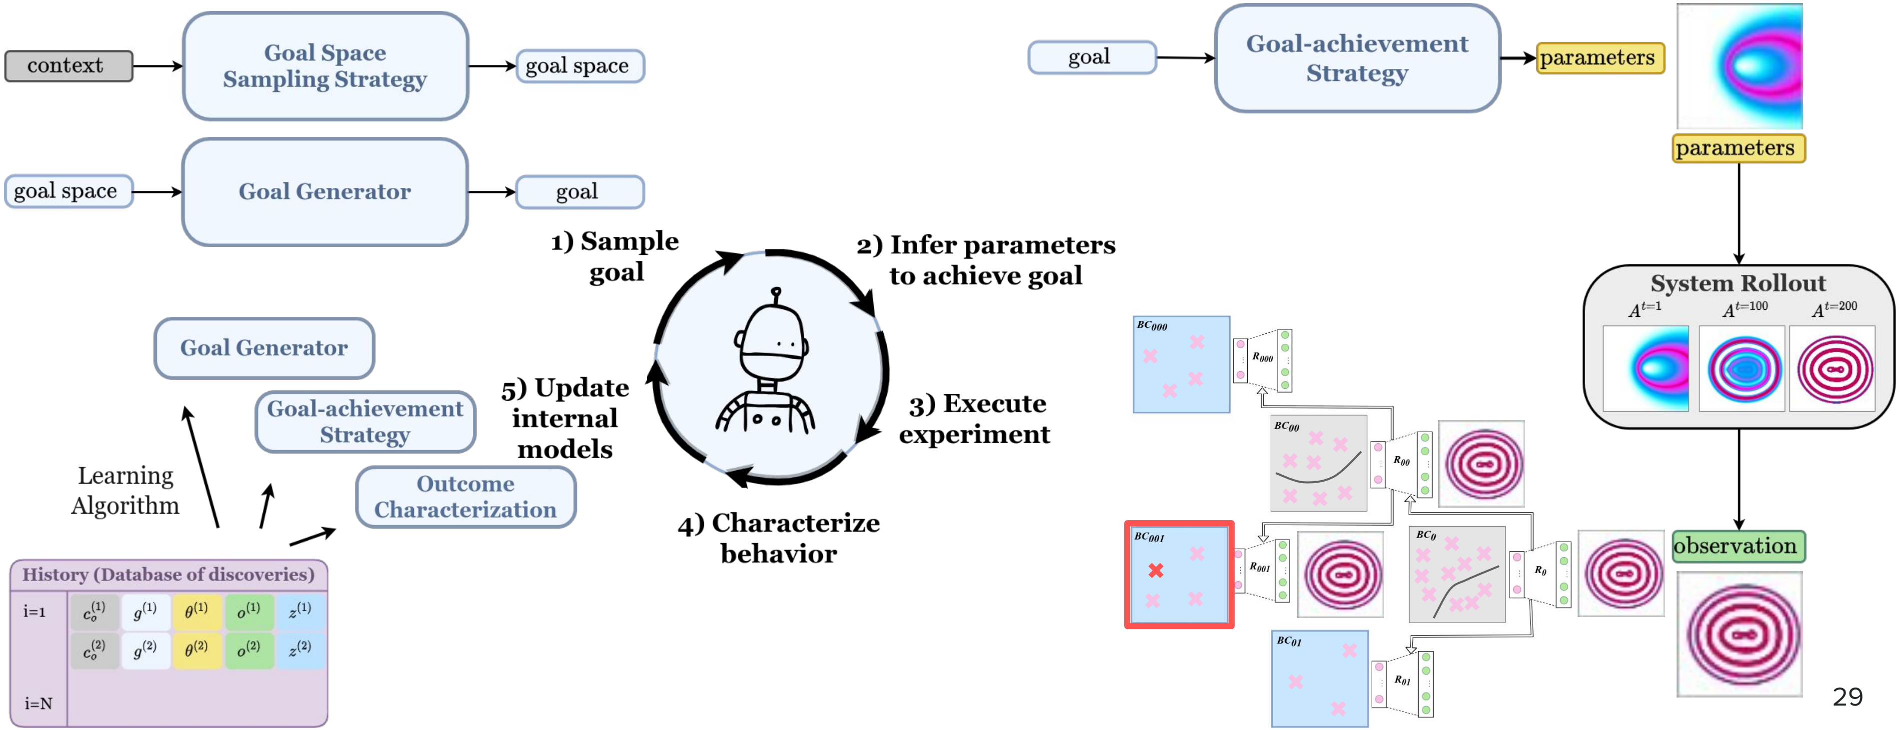

Results

Learning to explore diverse niches of patterns

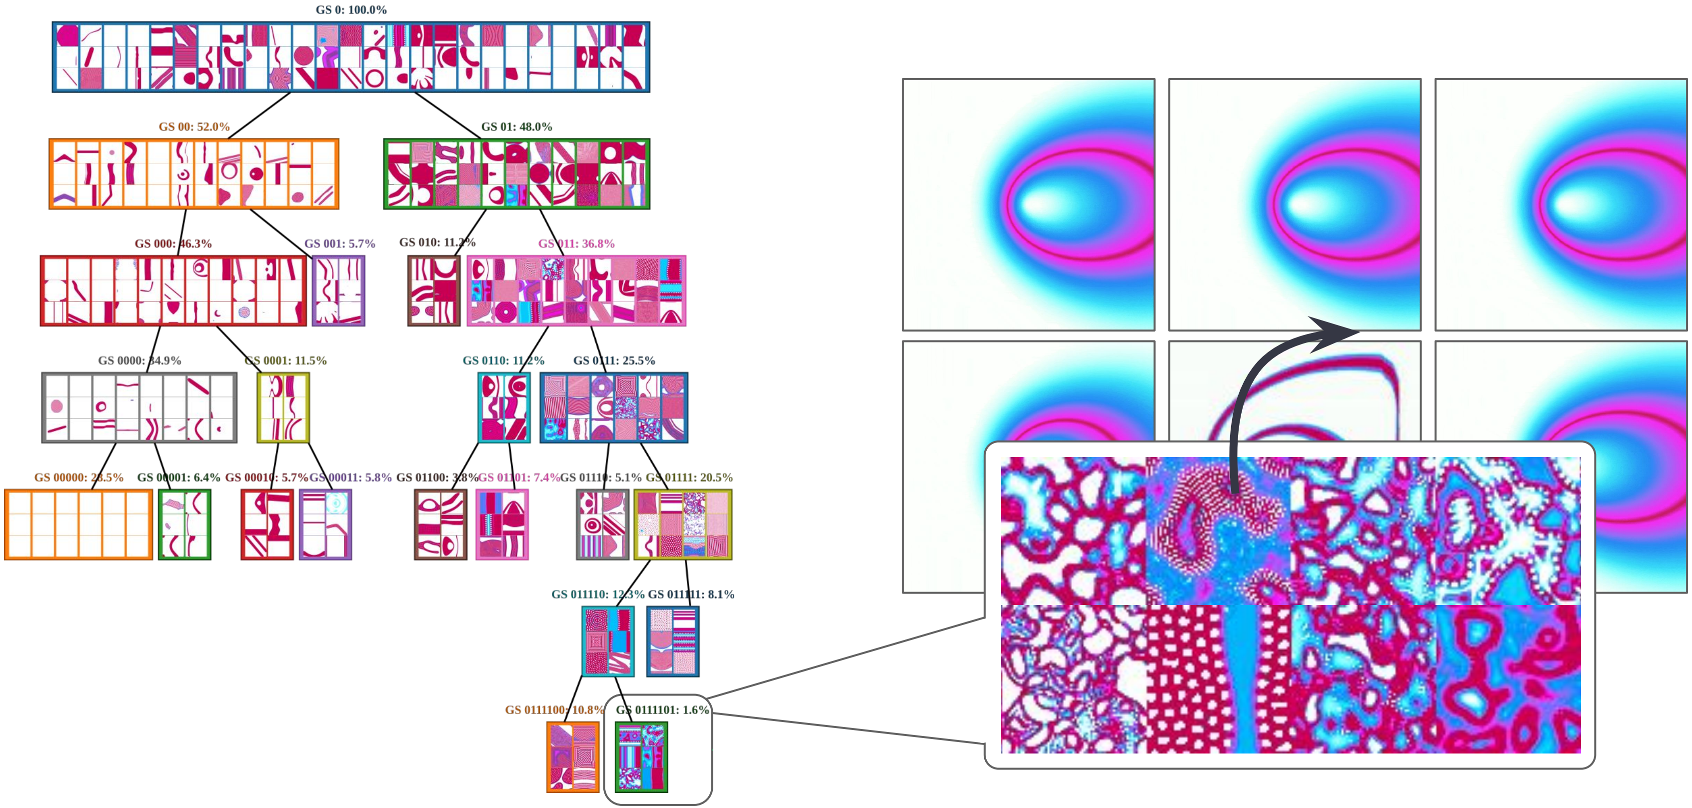

Preference-guided IMGEP-HOLMES: Adapting to User

3) How to quickly adapt the search to the end-user preferences?

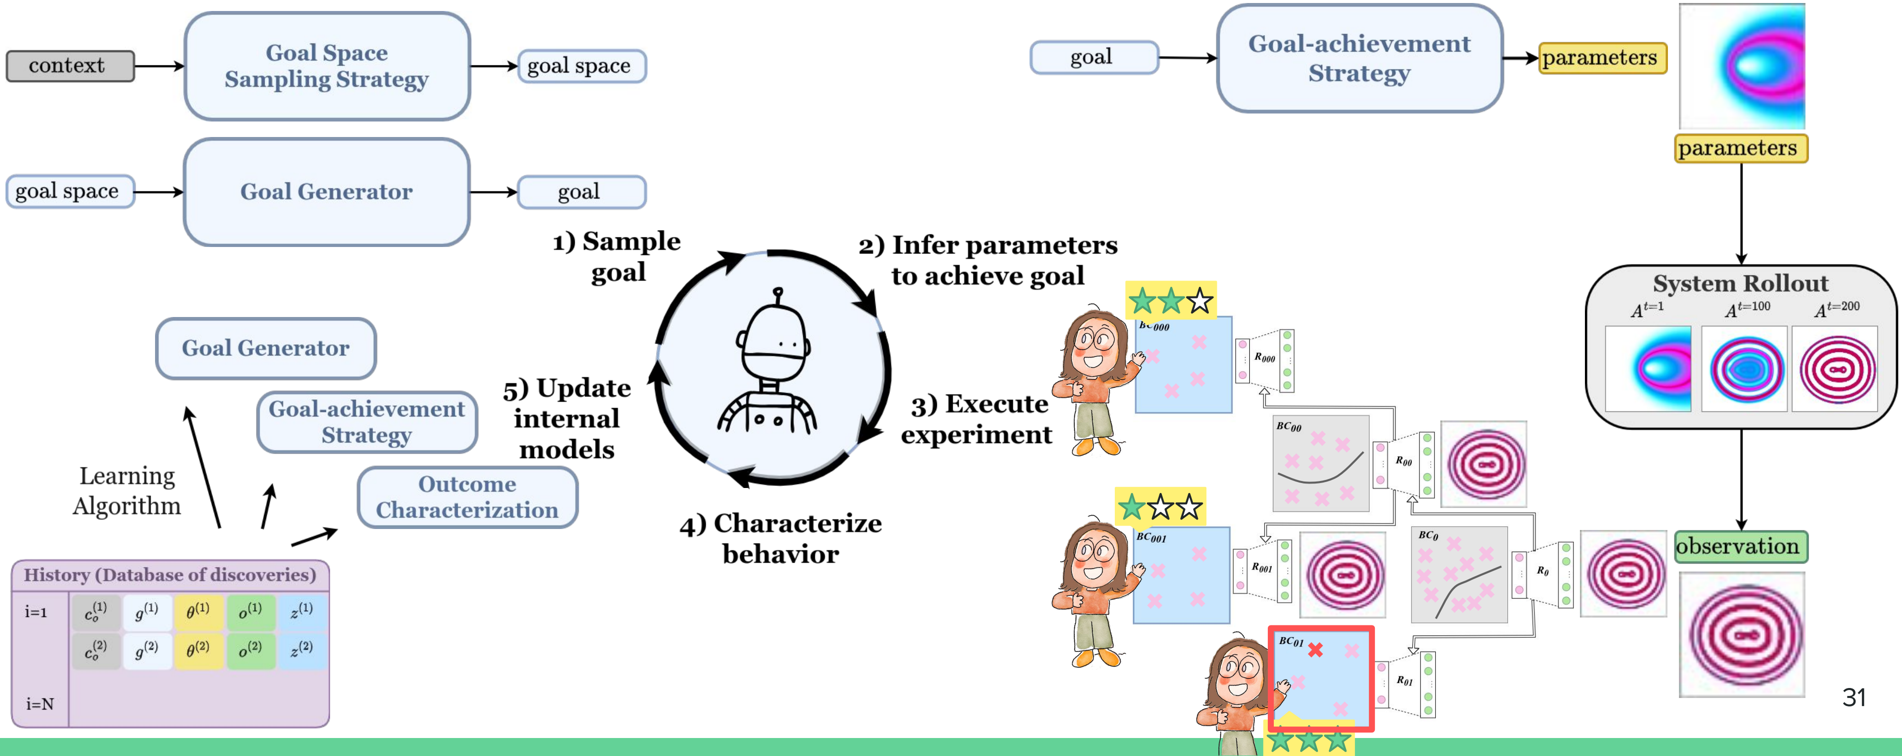

Results

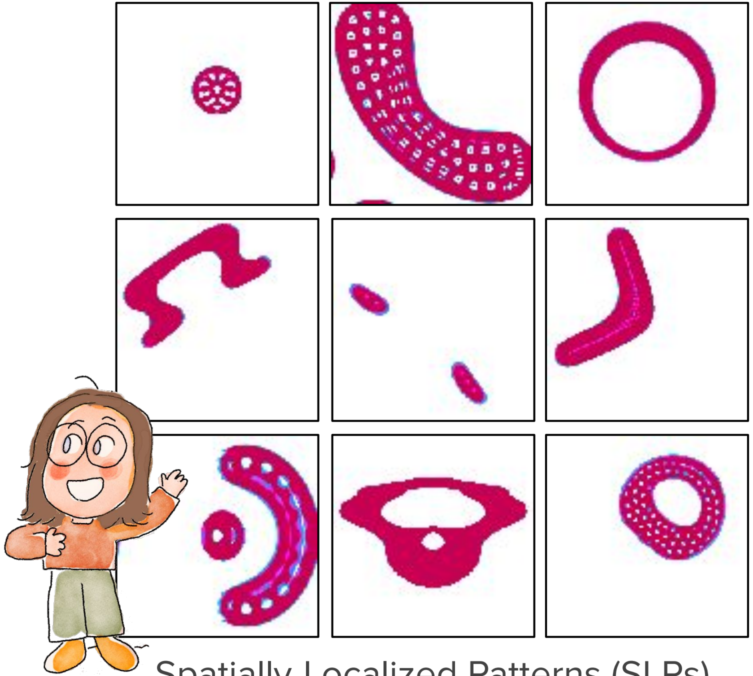

Spatially-Localized Patterns (SLPs)

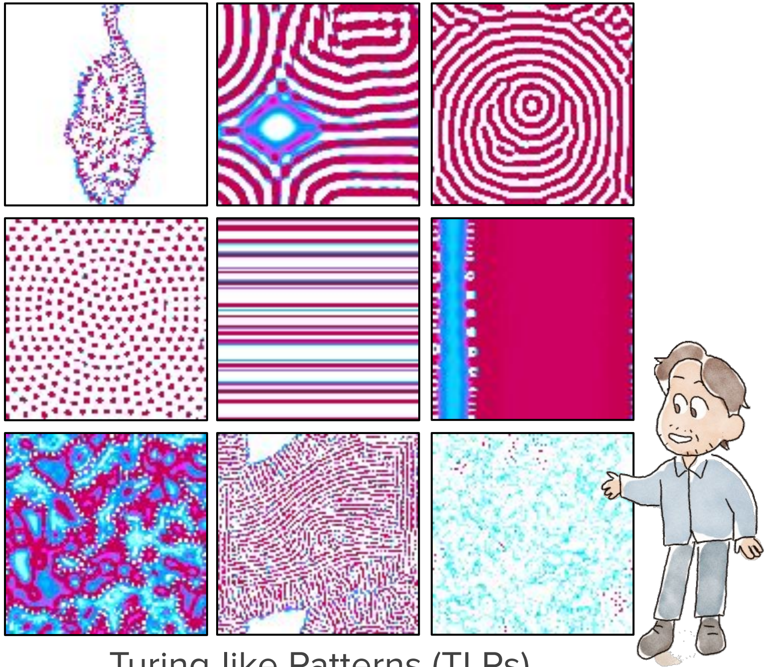

Turing-like Patterns (TLPs)

➜ We use the classifiers to simulate an external user that would prefer either SLPs or TLPs, and investigate how IMGEP-HOLMES search can be guided to specialize toward a diversity of either SLPs or TLPs.

Results

Learning to efficiently adapt the search

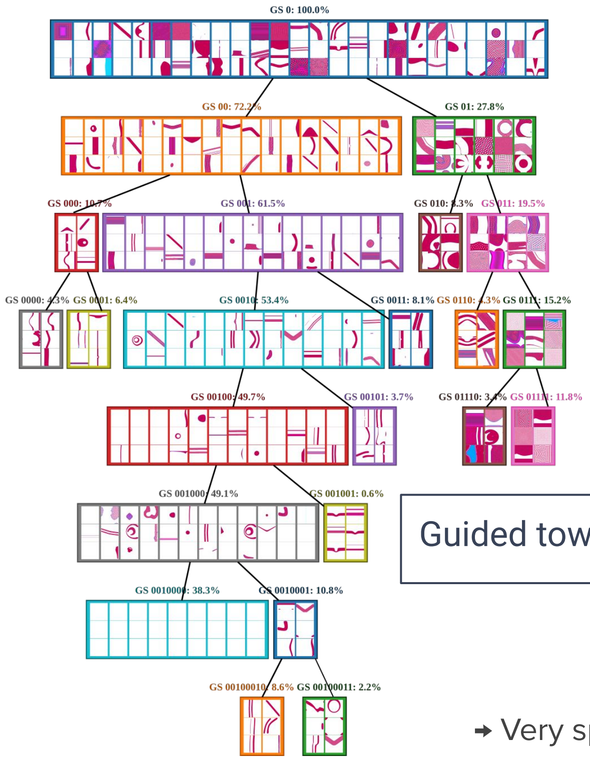

Guided toward SLPs

Guided toward TLPs

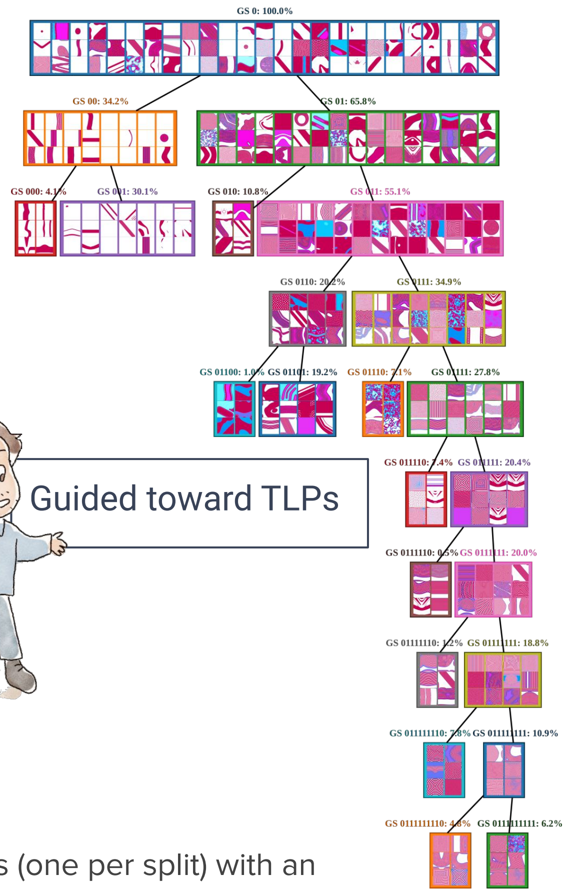

➔ Very sparse feedback: a total of 11 user interventions (one per split) with an average of 6 “clicks” (scores) to provide

Results *Learning to efficiently adapt the search*

Quantitative evaluation: **How to define the analytic behavior space?**

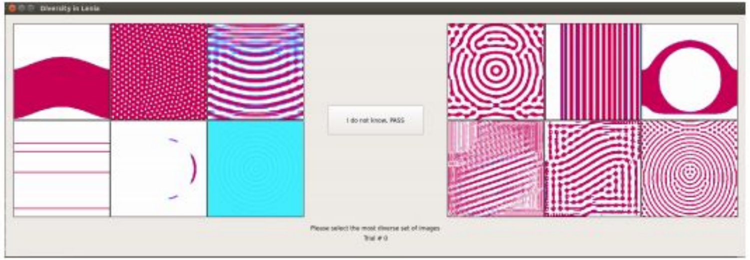

Table 2: Human-evaluator agreement scores (mean ± std). Best scores are shown in bold.

	Spectrum-Fourier	Elliptical-Fourier	Lenia-Statistics	BetaVAE	Patch-BetaVAE
SLP	0.5 ± 0.18	$\mathbf{0.98 \pm 0.04}$	0.59 ± 0.12	0.1 ± 0.06	0.89 ± 0.08
TLP	0.2 ± 0.13	0.47 ± 0.1	$\mathbf{0.92 \pm 0.07}$	0.75 ± 0.08	0.38 ± 0.08

(a) Diversity of SLP in $BC_{\text{ELLIPTICAL-FOURIER}}$

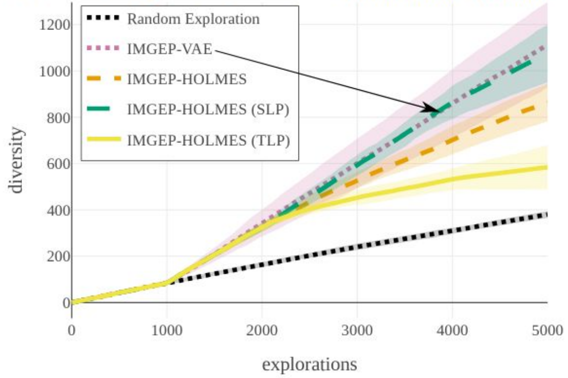

(b) Diversity of TLP in $BC_{\text{LENIA-STATISTICS}}$

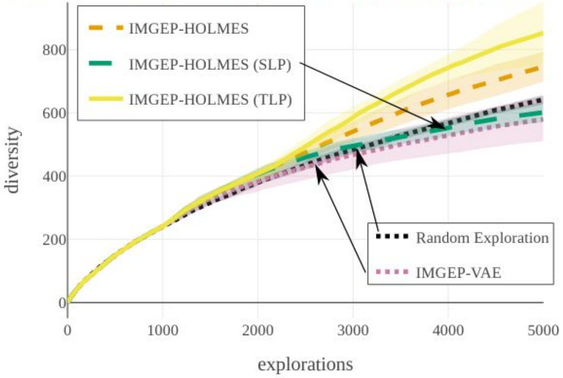

➔ IMGEP-VAE finds a high diversity of SLPs but a poor diversity of TLPs.

➔ When non-guided, IMGEP-HOLMES finds a higher diversity than Random Exploration both for SLPs and TLPs.

➔ When guided, IMGEP-HOLMES can further increase the discovered diversity in the category of interest.

Part I Takeaways

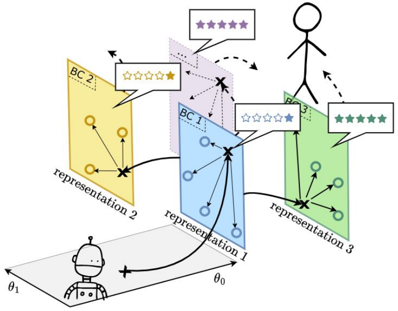

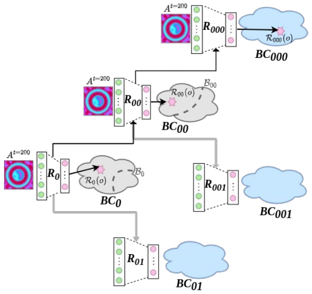

- Novel objective of *meta-diversity* search

- D*ynamic* and *modular* architecture for unsupervised learning of *diverse* representations

- Integrated with intrinsically-motivated goal exploration processes, enables efficient *guidance* toward the preferences of a simulated end-user, using very little user feedback

Part 2: Use Cases of the Curious Discovery Assistant

Sensorimotor Lenia

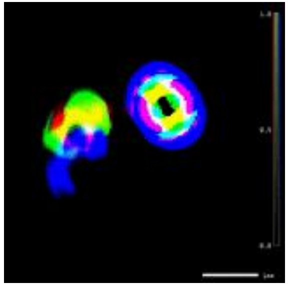

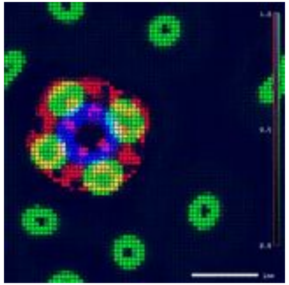

Use Case #1

Studying the emergence of robust forms of “sensorimotor agency” in continuous CA models

Collaboration: Gautier Hamon (INRIA), Bert Chan (Google Brain)

“Learning Sensorimotor Agency in Cellular Automata”, Gautier Hamon, Mayalen Etcheverry, Bert Chan, Clément Moulin-Frier, Pierre-Yves Oudeyer (In Submission)

Studying of sensorimotor agency in continuous CA

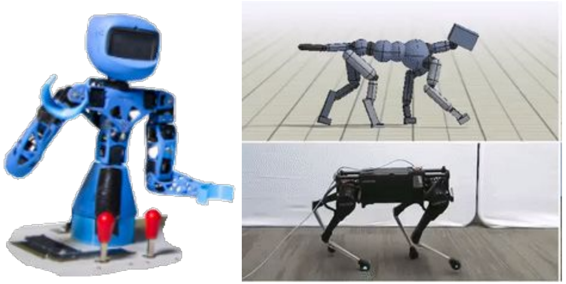

→ Already assume the existence of *agents* with predefined *body*, *brain*, *sensors* and *actuators*

coherent entity able to robustly perform a variety of behaviors that involve the process of sensing and acting in the environment

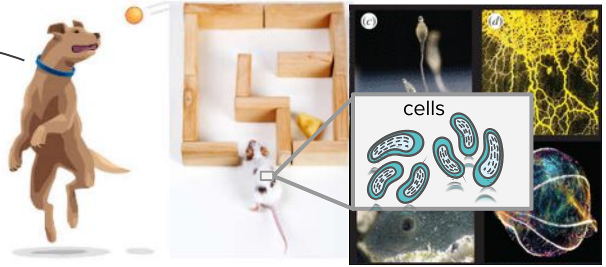

→ How do “agency” and “embodiment” arise from collective of cells and distributed low-level rules?

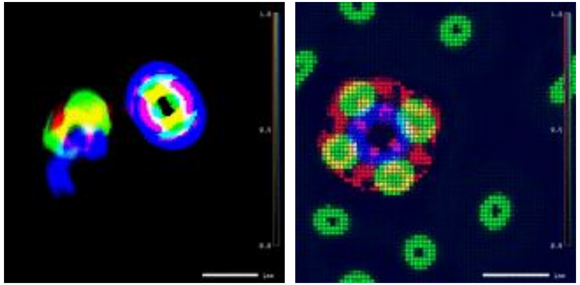

→ Only environment with low-level elements and physical laws, no prior notion of *agency*, *body*, *sensors*, or *actuators*.

How to find environmental rules leading to the emergence of <u>autopoietic entities</u> with <u>sensorimotor abilities</u>?

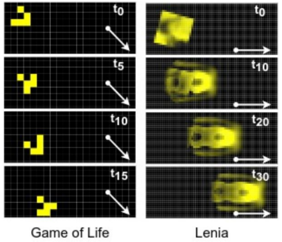

→ tedious and hard to find

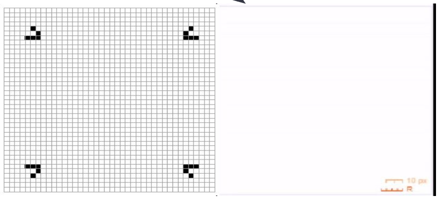

→ fragile to perturbations

Discovery of rules leading to sensorimotor agency

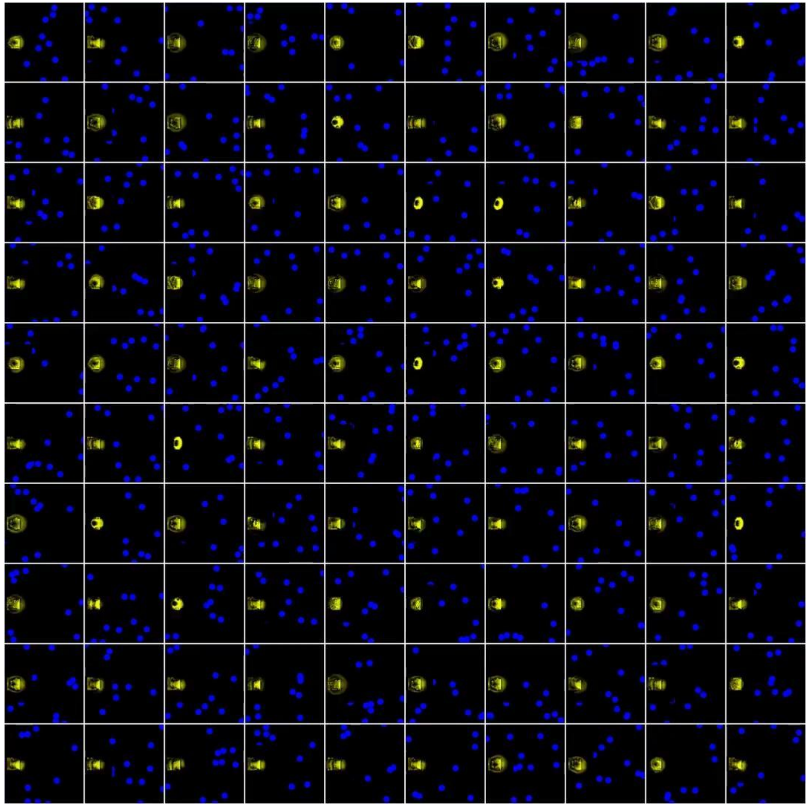

- A tailored curiosity search is able to find environmental rules leading to the self-organization of *individuality*, *locomotion* and *sensorimotor abilities*

- Very hard to obtain with
  - random search (≃**0.03%** of moving agents)
  - moving agents found “by hand” are not robust to the introduced obstacle perturbations

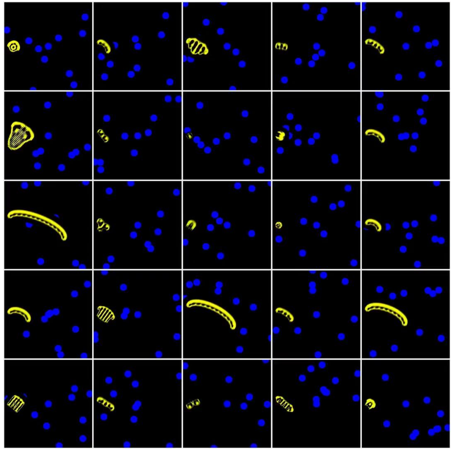

Robustness to novel perturbations

“Environmental” perturbations

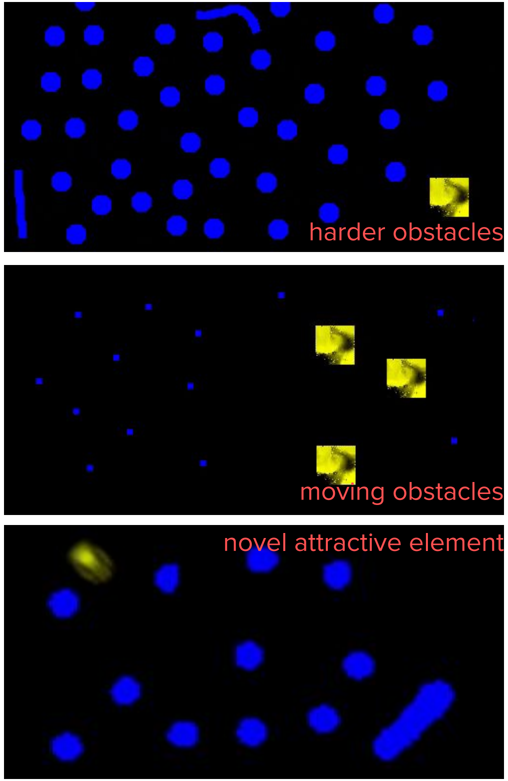

“Organic” perturbations

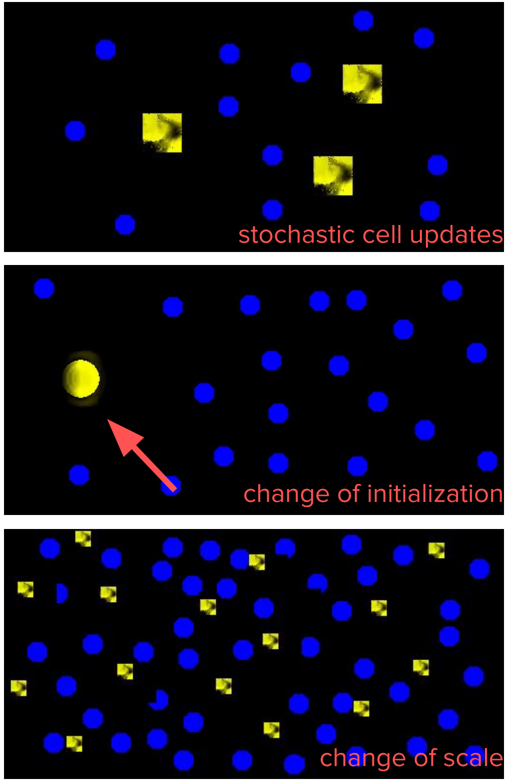

“Intersubjective” perturbations

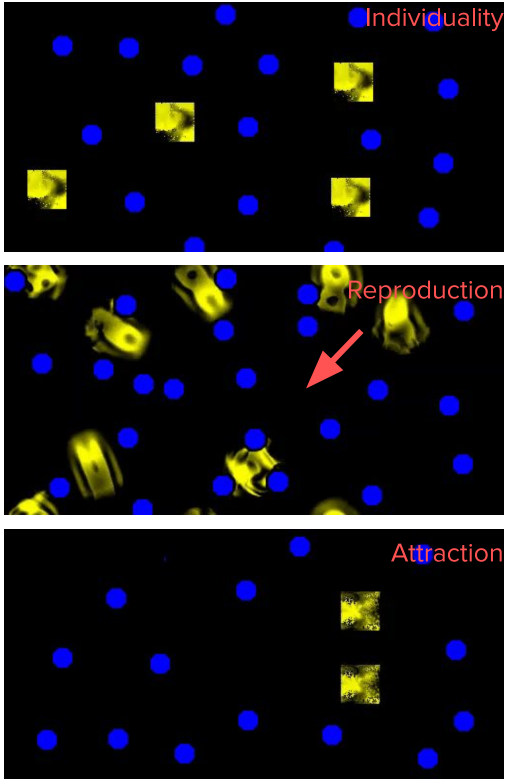

Use case 1 Takeaways

- Curiosity search enabled the discovery of diverse forms of “sensorimotor agents” in Lenia
  ➔ *shows how a collective of simple identical cells can make “decision” and “sense” at the macro scale through local interactions only*

- The discovered agents showed surprisingly robust capabilities to move and maintain their body integrity despite several hard perturbations
  ➔ *reminiscent of generalization capabilities observed in biological organisms*

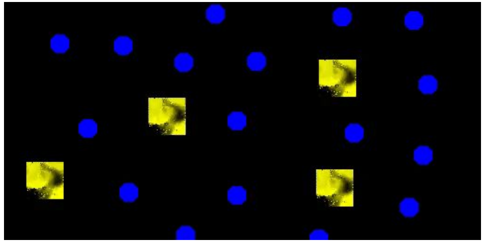

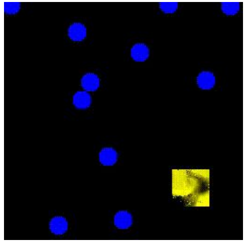

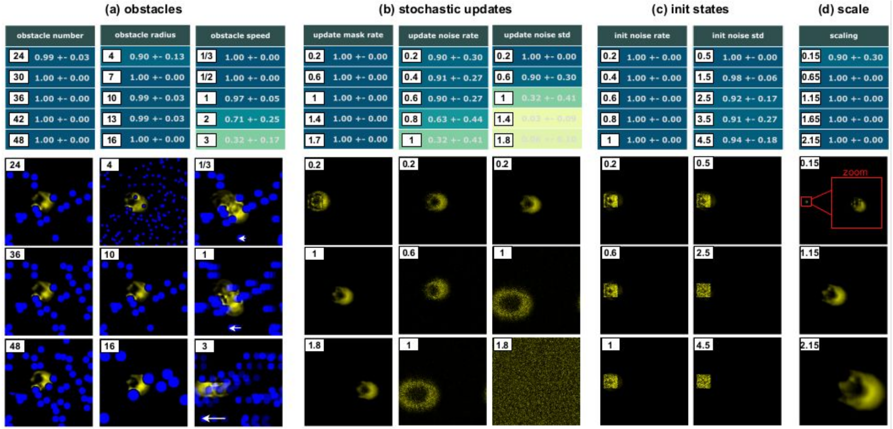

Biological Network Competencies

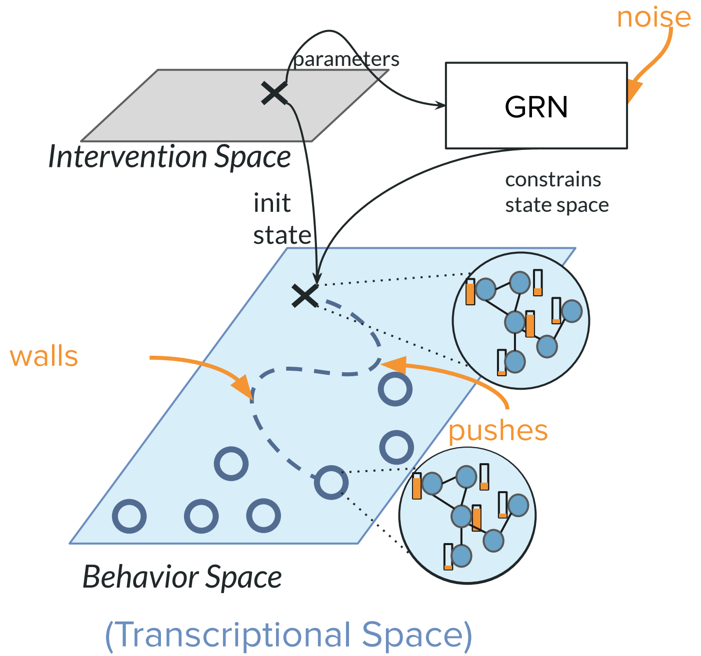

Use Case #2

Revealing Diverse Behavioral Competencies in Gene Regulatory Networks via Minimal Interventions

Collaboration: Michael Levin (Tufts University)

"AI-driven Automated Discovery Tools Reveal Diverse Behavioral Competencies of Biological Networks", Mayalen Etcheverry, Clément Moulin-Frier, Pierre-Yves Oudeyer, Michael Levin (In Submission)

Navigation Competencies of Unconventional Agents

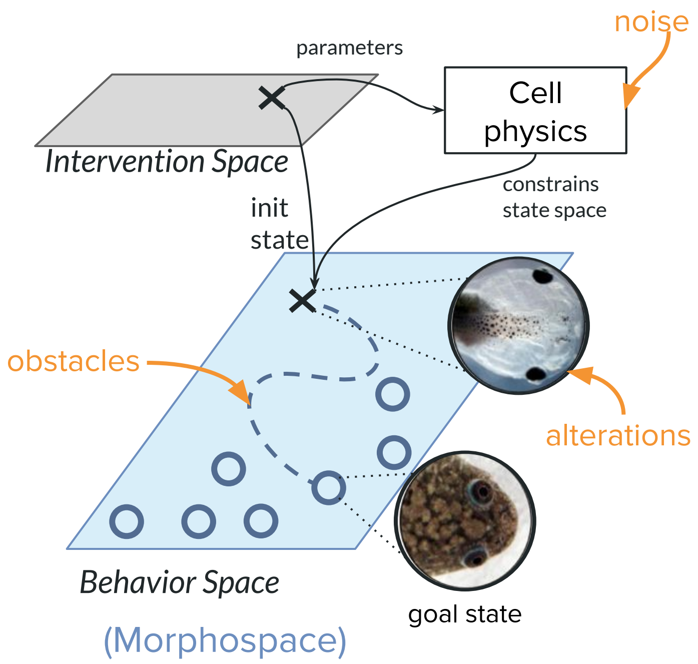

- Lenia creatures as “agents” navigating cellular automata grid space with robust competencies

- Biological systems as “agents” navigating their own problem spaces with robust competencies
    - Cellular collectives as “agents” acting in morphological space
    - Subcellular systems (biomolecular pathways) as “agents” acting in transcriptional space
    - etc

multiscale competency architecture (Levin, 2022)

GRNs: Gene Regulatory Networks

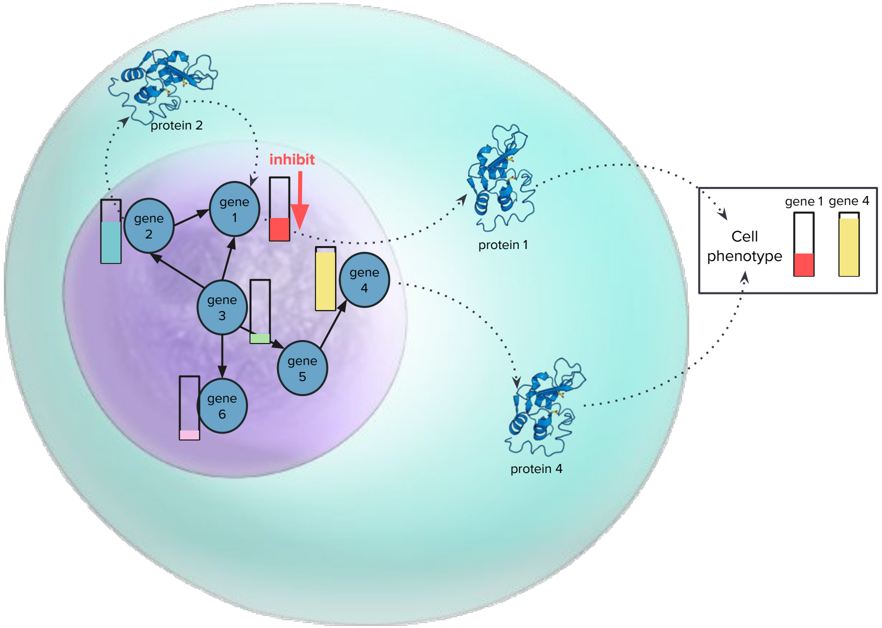

GRNs: Gene Regulatory Networks

- GRN models curated by biologists available on online database
- Simulations with interventions in JAX

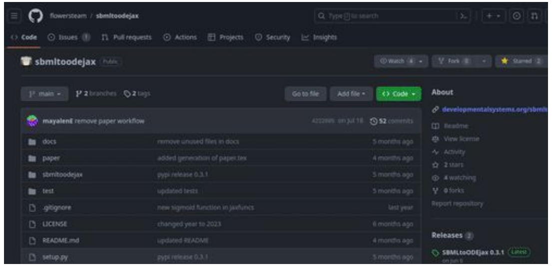

SBMLtoODEjax

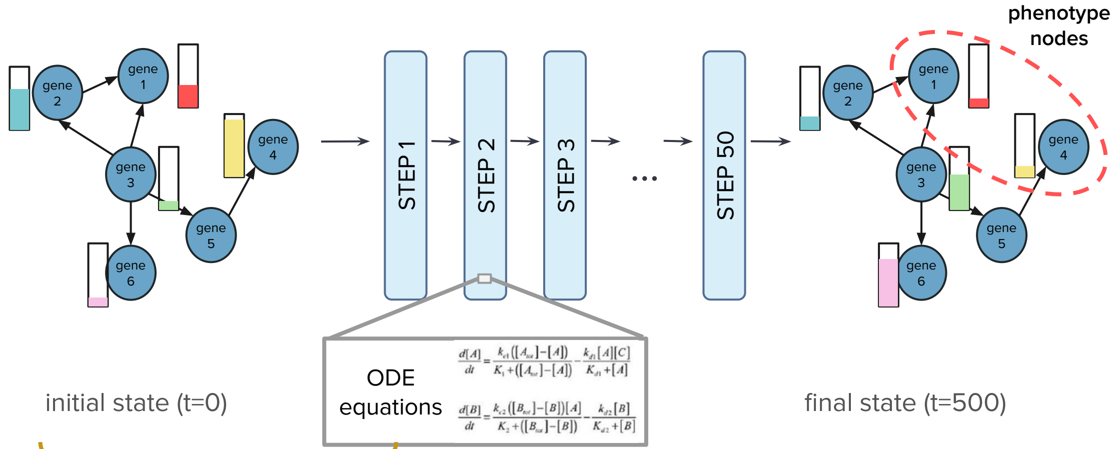

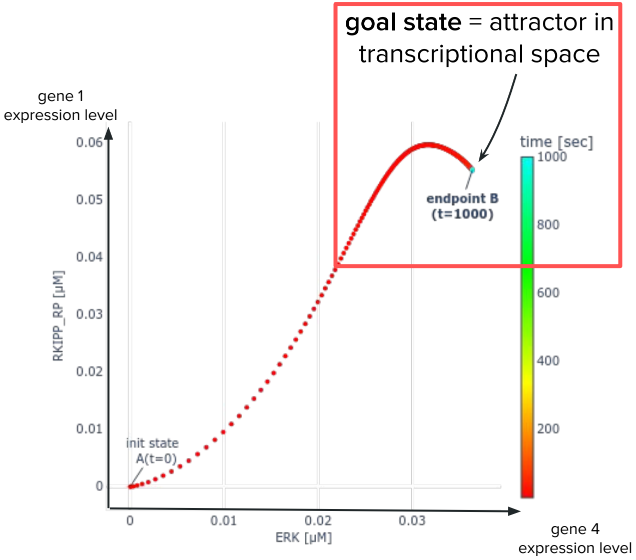

controllable parameters θ

observations *o*

Navigation Competencies of Biomolecular Networks

Hypothesis: *Biomolecular networks (GRNs) can be seen as "agents" navigating transcriptional space toward "goal states" with varying degrees of "competencies"* (Fields and Levin, 2022)

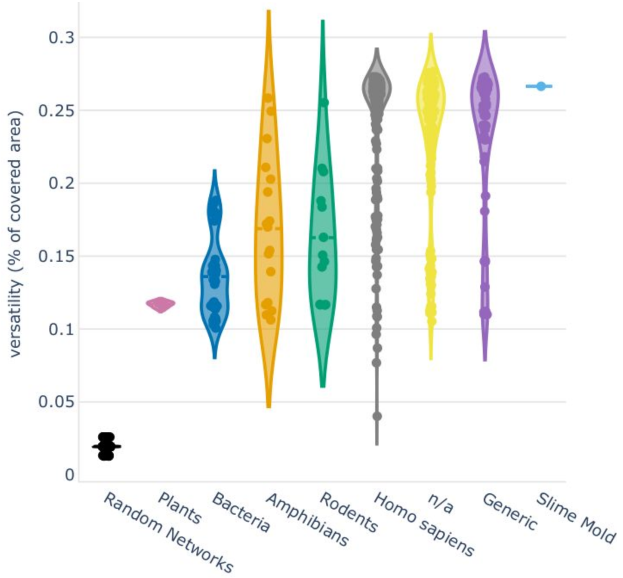

Versatility = *GRN's capacity to reach diverse goal states under minimal interventions*

Approach: Curiosity search to find the range of possible goal states (attractors)

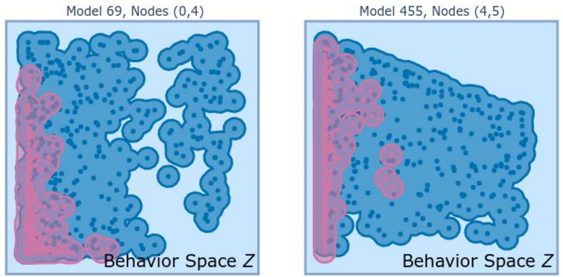

Discoveries by curiosity search (blue) and random search (pink)

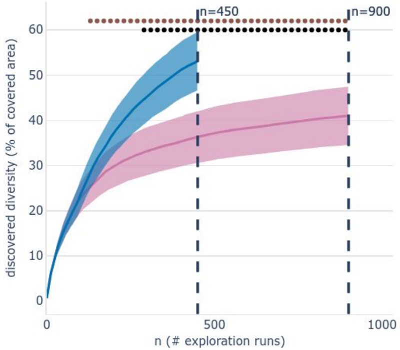

Discovered diversity on 432 systems

➜ Discovered diversity suggests that (some) GRNs can reach a broad spectrum of steady states (which would have been very long to discover with a simple random search)

Navigation Competencies of Biomolecular Networks

Hypothesis: *Biomolecular networks (GRNs) can be seen as "agents" navigating transcriptional space toward "goal states" with varying degrees of "competencies"* (Fields and Levin, 2022)

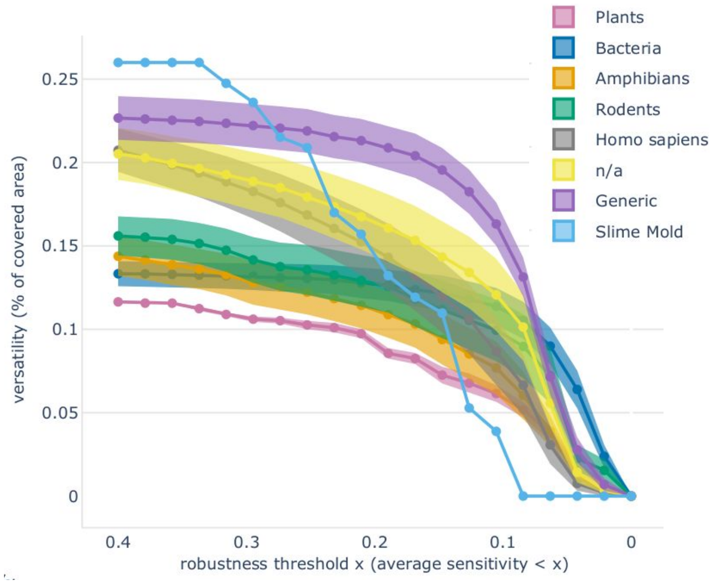

Robustness = *GRN's capacity to reach a goal states despite various perturbations*

Approach: Battery of empirical tests to assess the robustness of discovered goal states

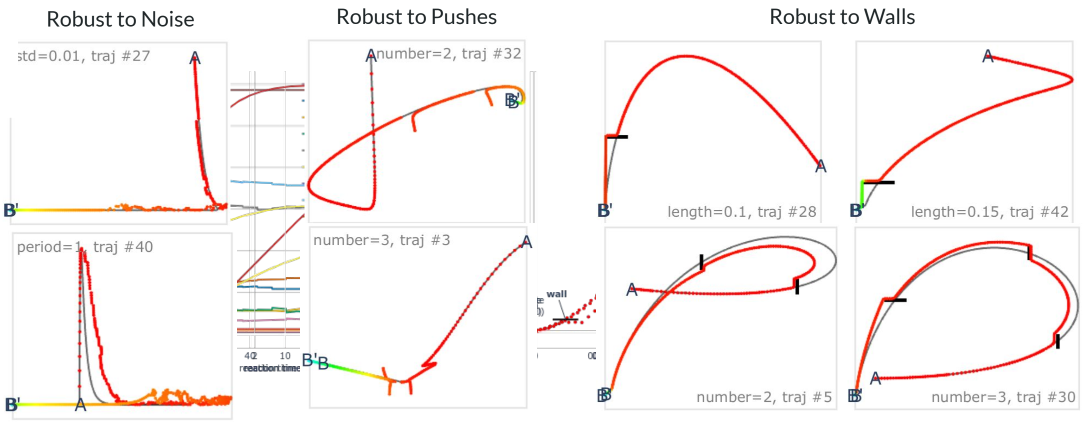

→ Discovered various complex yet highly robust space-traversal strategies in transcriptional space (reminiscent of navigation competencies of living "agents" operating in other "spaces")

Reuses for BioMedicine

GRNs associated to development of diseases

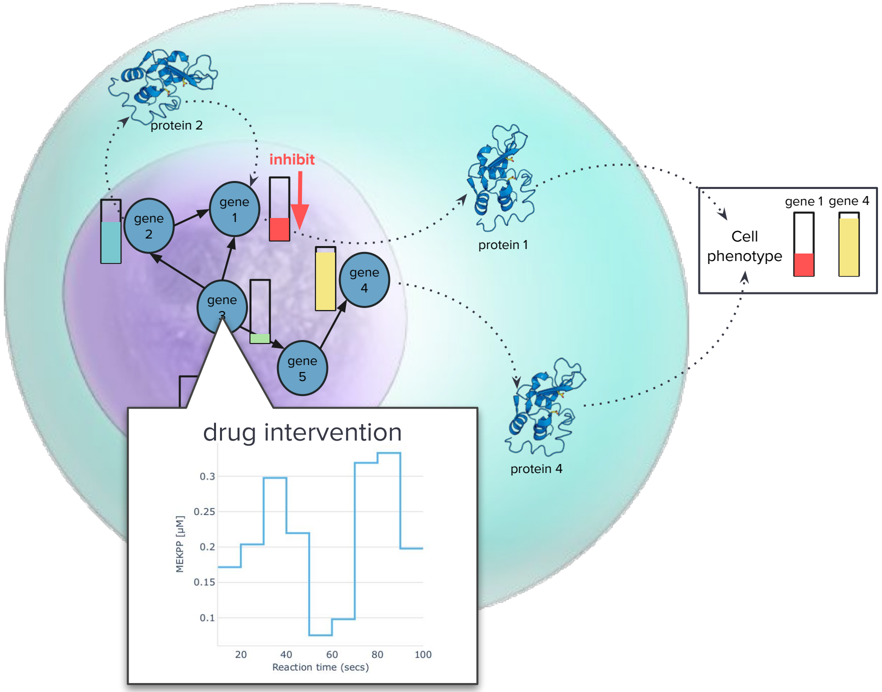

Design of therapeutic interventions

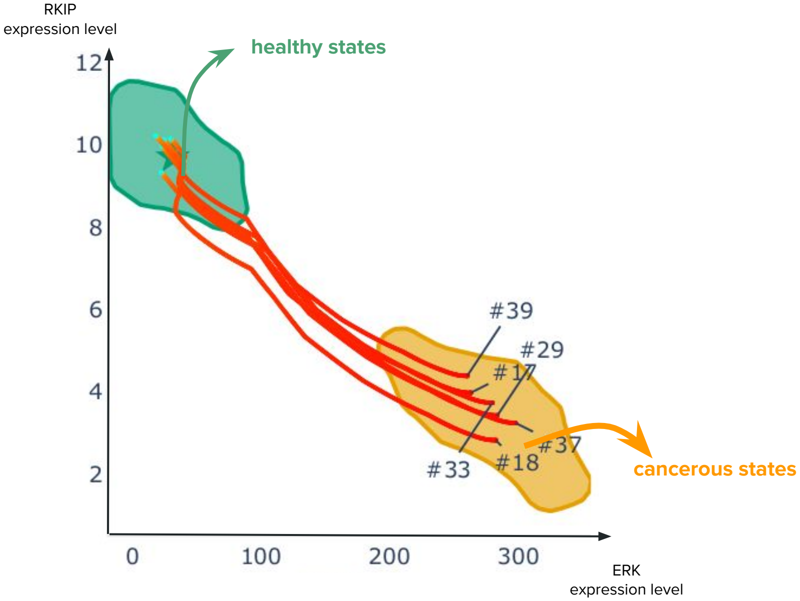

RKIP-ERK signalling pathway (Kwang-Hyun et al., 2003)

Reuses for BioEngineering

Synthetic Gene Circuit Engineering

- Standard optimization problem
- Alternative diversity search strategy
- Refined solution with local optimization

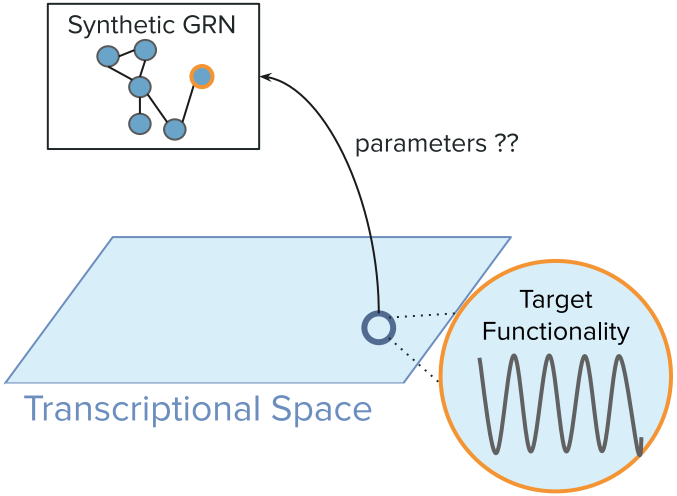

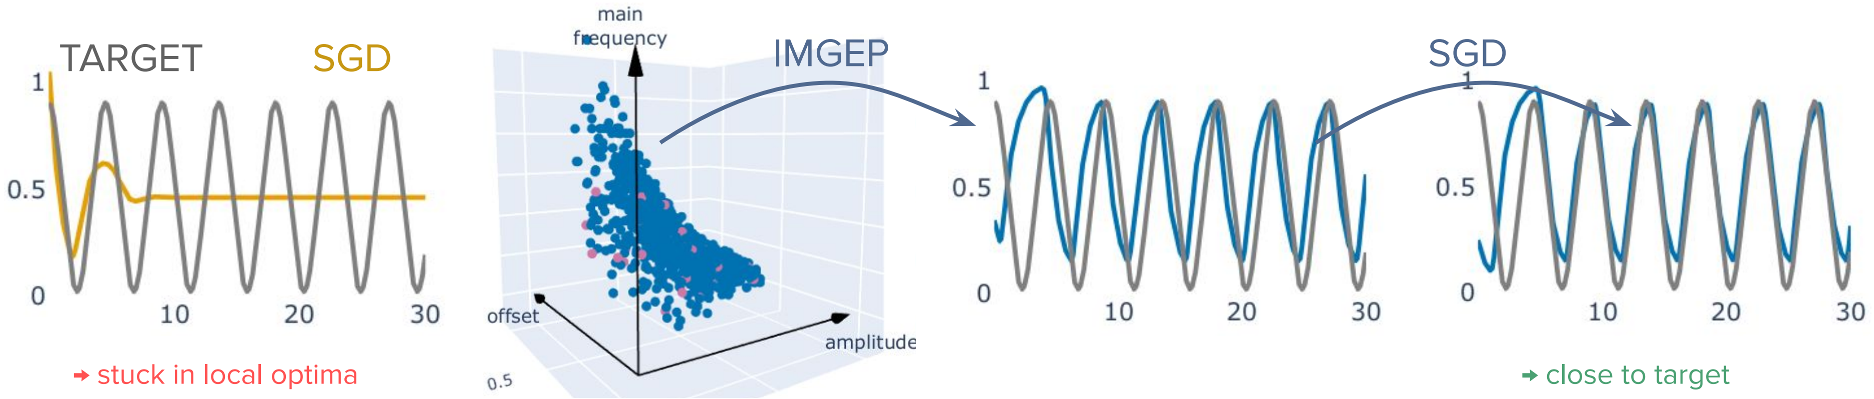

Use case 2 Takeaways

- Curiosity search enables to efficiently map the space of behaviors of biological networks
- Some biomolecular networks showed surprisingly robust navigation competencies
- Several possible reuses for specific problems in biomedicine and bioengineering

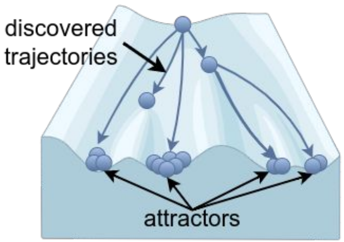

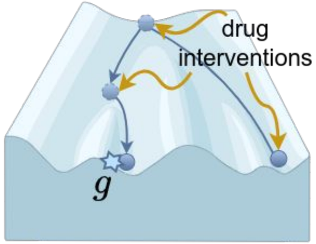

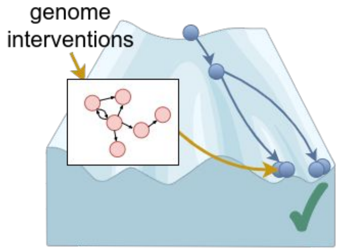

Conclusion and Next Steps

Conclusion

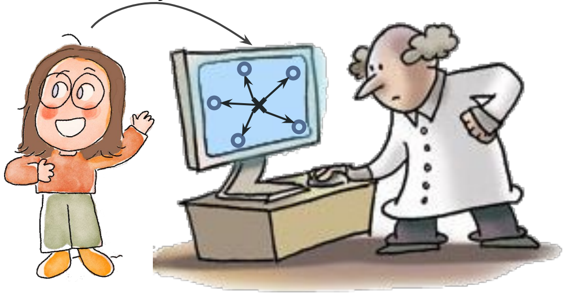

- Curiosity-driven exploration algorithm provides an efficient framework to explore and map the space of possible outcomes of complex self-organizing systems

- Many possible algorithmic developments can be envisaged to build more open-ended forms of discovery assistants
    - Two contributions: Meta-Diversity Search and Human Guidance

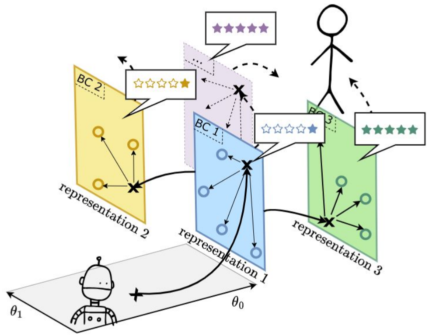

- Step closer toward having digital discovery assistants for assisting scientific discovery in complex systems
    - Two use-cases in continuous CA models and biological network models

Next Steps

Experiments in a bioprinter-controlled biological system

small budget: $10^2 < N < 10^3$

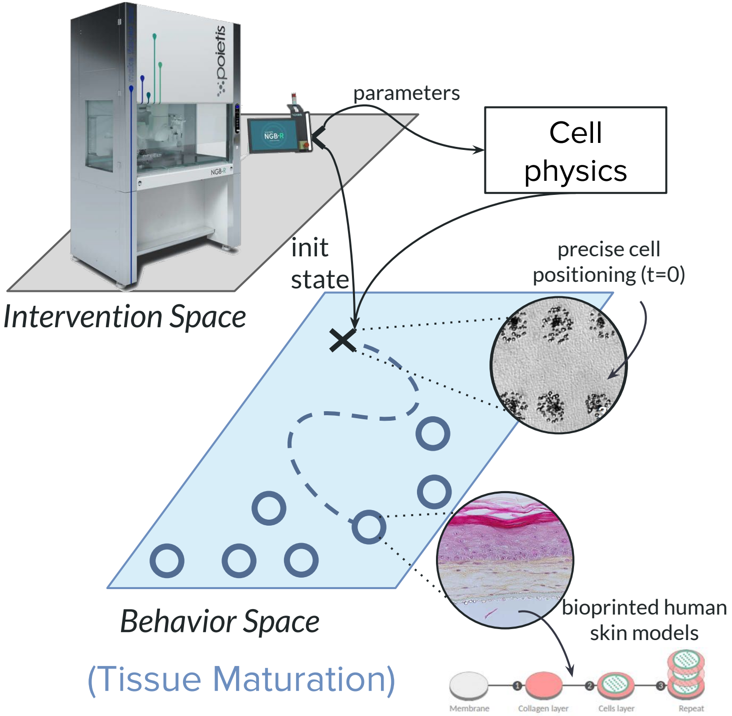

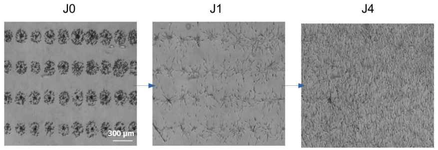

Campaign #1: Diversity search to find diverse cell layer orientations

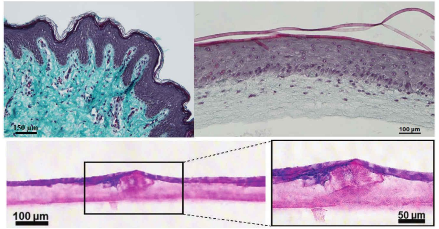

Campaign #2: Diversity search to find diverse derm surface topographies

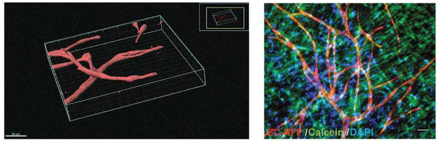

Campaign #3: Diversity search to find diverse tubular epithelial structures

Perspectives

Meta-Diversity Search and Human Guidance as a toolbox to conceptualize Open-Endedness

Several algorithmic perspectives:

- Towards a richer diversity of goals
- Towards richer interactions with (real) humans
- To be deployed to other systems

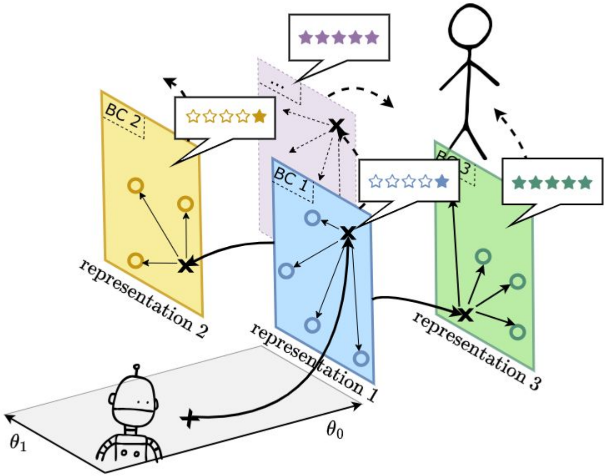

Perspectives

For the design of novel forms of collective intelligences

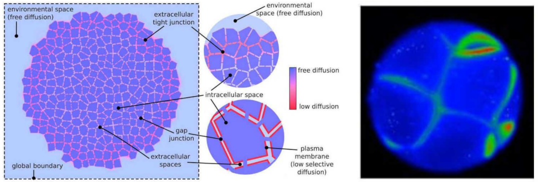

Bioelectric patterns
(Pietak and Levin., 2016)

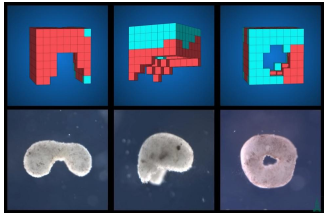

Xenobots
(Kriegman et al., 2020)

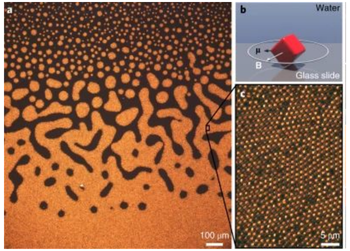

Active Materials
(Soni et al., 2019)

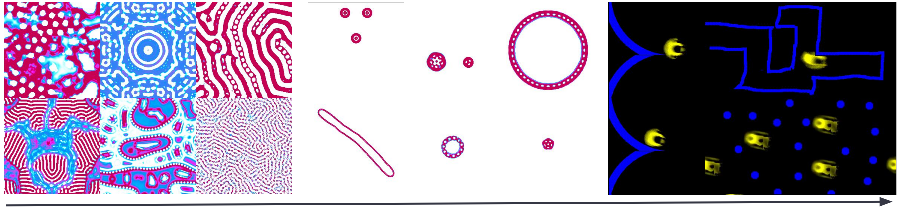

TLPs SLPs sensorimotor agents
(Hamon et al., 2022)

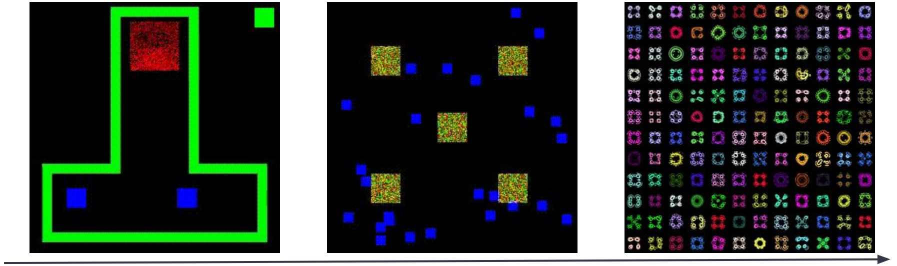

memory? learning? competition? open-ended evolution?

“Flow-Lenia: Towards open-ended evolution in cellular automata through mass conservation and parameter localization”, Erwan Plantec, Gautier Hamon, Mayalen Etcheverry, Pierre-Yves Oudeyer, Clément Moulin-Frier, Bert Chan. ALife 2023 (Best Paper Award)

Thanks !

Allen Discovery Center

poietis
make tissues real

Supervisors:


Pierre-Yves Oudeyer
(INRIA)


Clément Moulin-Frier
(INRIA)

Marc Nicodème
(Poietis)

Jury:


Alan Aspuru-Guzik
(Univ Toronto)


Sebastian Risi
(ITU Copenhagen)


Melanie Mitchell
(Sante Fe Institute)


Jeff Clune
(Univ British Columbia)


Nicolas Brodu
(INRIA)

Collaborators:


Gautier Hamon
(INRIA)


Erwan Plantec
(INRIA)


Chris Reinke
(INRIA)


Bert Chan
(Google Brain)

Auto Disc Tool

Clément Romac,
Mathieu Perie,
Jesse Lin (INRIA)


Michael Levin
(Tufts University)