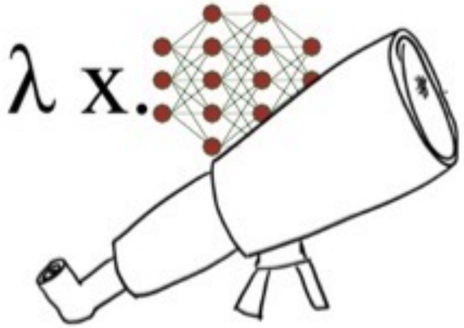

Neurosymbolic Programming for Scientific Discovery

Atharva Sehgal

atharvas@utexas.edu

TEXAS
The University of Texas at Austin

Scientific Discovery

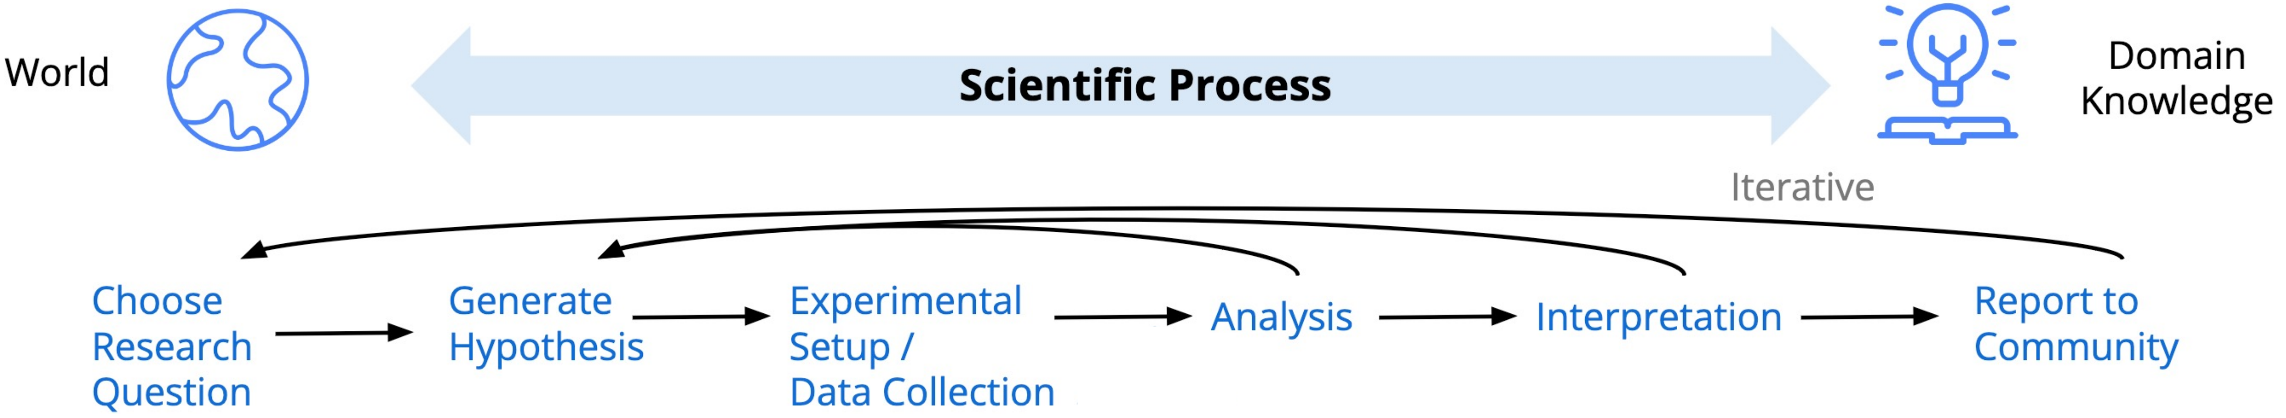

Goal: We want AI to achieve human level performance at *research in the natural sciences*.

Lifecycle of a Scientific Process

World ⟷ **Scientific Process** ⟷ Domain Knowledge

Iterative

Choose Research Question → Generate Hypothesis → Experimental Setup / Data Collection → Analysis → Interpretation → Report to Community

Observation: Apparent Retrograde Planetary Motion

c.The Astronomical Revolution: Copernicus- Kepler-Borelli

Lifecycle of a Scientific Process

World — **Scientific Process** — Domain Knowledge

Iterative

Choose Research Question → Generate Hypothesis → Experimental Setup / Data Collection → Analysis → Interpretation → Report to Community

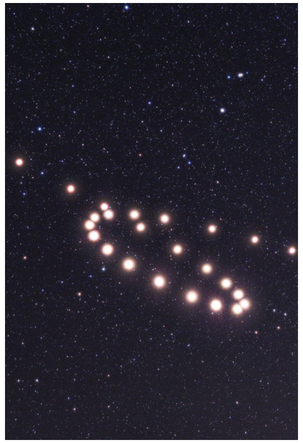

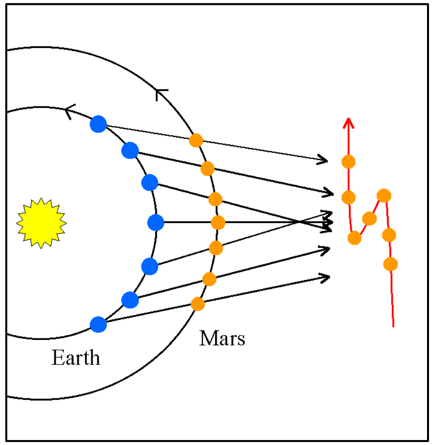

Observation: Apparent Retrograde Planetary Motion
Theory: Heliocentric Model

c.The Astronomical Revolution: Copernicus- Kepler-Borelli

Lifecycle of a Scientific Process

World

Scientific Process

Domain Knowledge

Iterative

Choose Research Question → Generate Hypothesis → **Experimental Setup / Data Collection** → Analysis → Interpretation → Report to Community

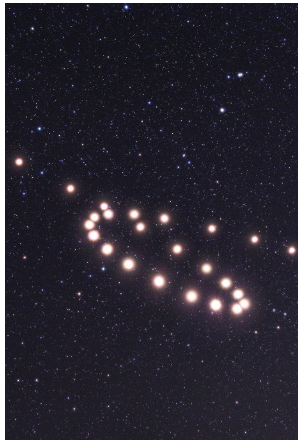

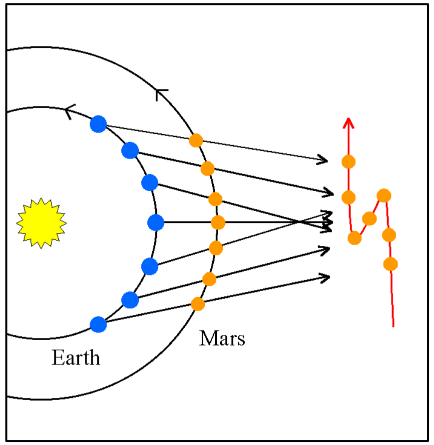

Observation: Apparent Retrograde Planetary Motion
Theory: Heliocentric Model

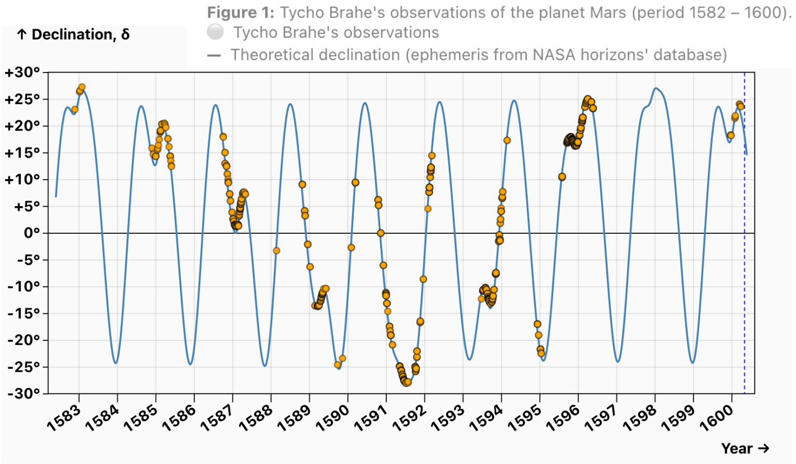

Data Collection: Sample data in regime of interest.

c.The Astronomical Revolution: Copernicus- Kepler-Borelli

Lifecycle of a Scientific Process

World

Scientific Process

Domain Knowledge

Iterative

Choose Research Question → Generate Hypothesis → Experimental Setup / Data Collection → Analysis → Interpretation → Report to Community

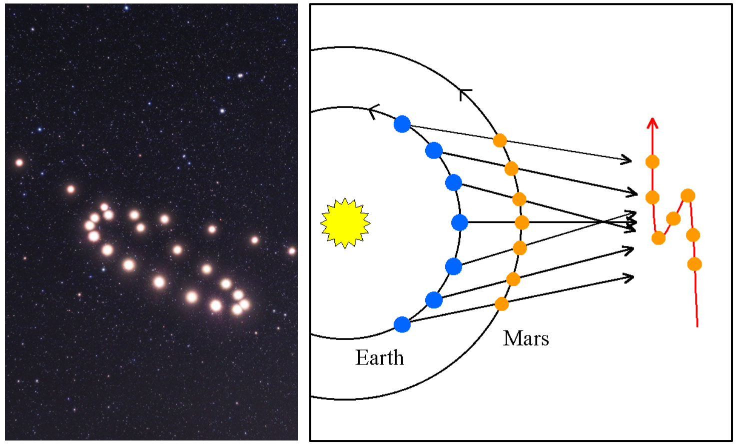

Observation: Apparent Retrograde Planetary Motion
Theory: Heliocentric Model

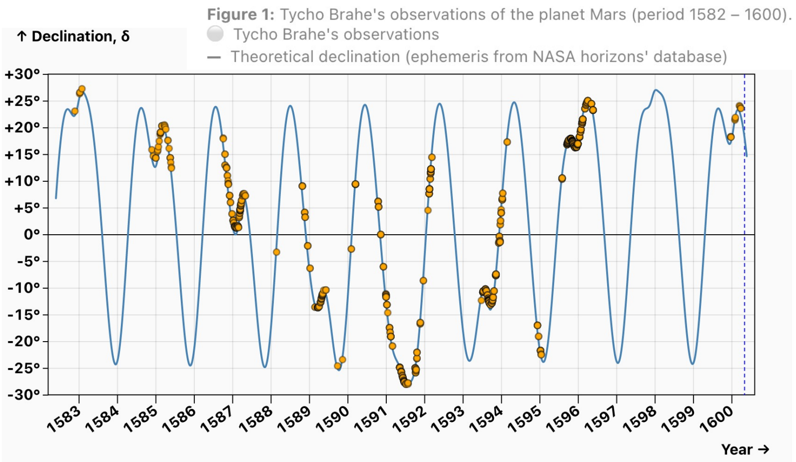

Data Collection: Sample data in regime of interest.

Analysis: Kepler's Third Law

$$T^2 \propto r^3$$

c.The Astronomical Revolution: Copernicus- Kepler-Borelli

Lifecycle of a Scientific Process

World ⟷ **Scientific Process** ⟷ Domain Knowledge

Iterative

Choose Research Question → Generate Hypothesis → Experimental Setup / Data Collection → Analysis → $\boxed{\text{Interpretation}}$ → Report to Community

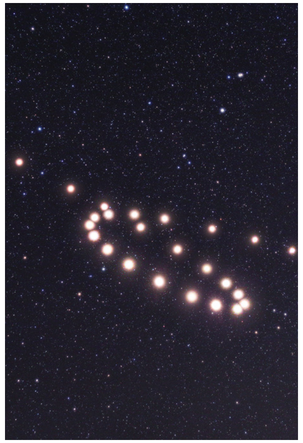

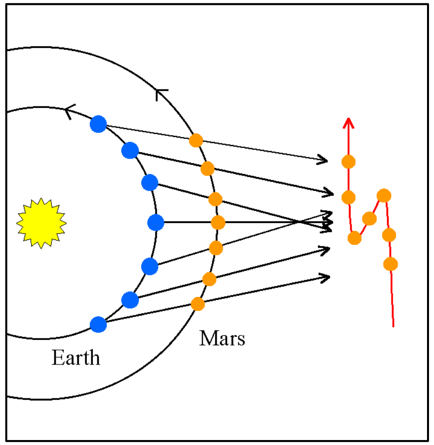

Observation: Apparent Retrograde Planetary Motion

Theory: Heliocentric Model

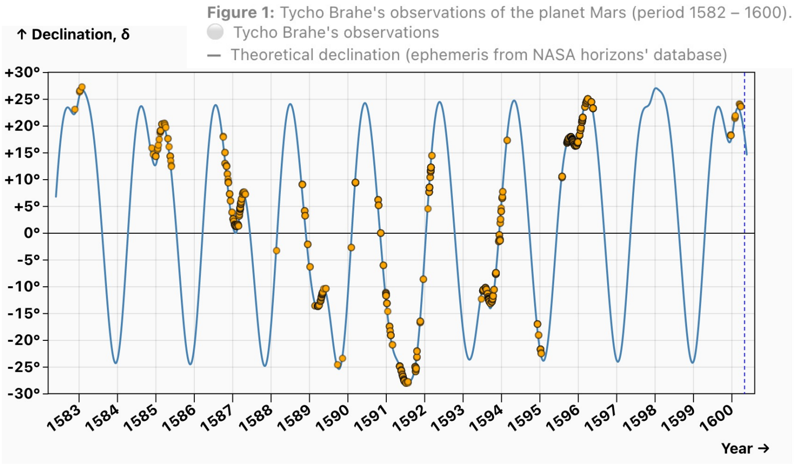

Data Collection: Sample data in regime of interest.

Analysis: Kepler's Third Law

$$T^2 \propto r^3$$

↓

Interpretation: Newton's Law of Gravitation

$$mr\left(\frac{2\pi}{T}\right)^2 = G\frac{mM}{r^2}$$

c.The Astronomical Revolution: Copernicus- Kepler-Borelli

Symbolic Regression

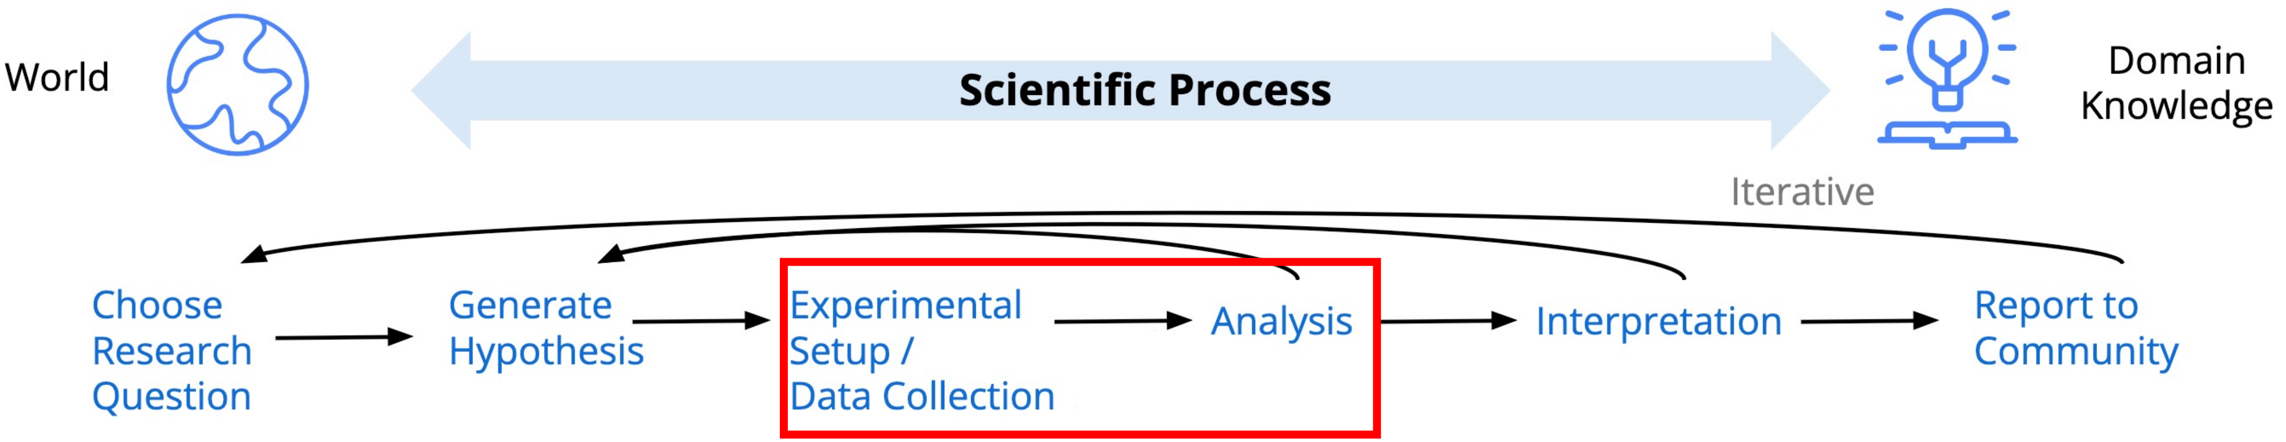

Figure 1: Tycho Brahe's observations of the planet Mars (period 1582 – 1600).

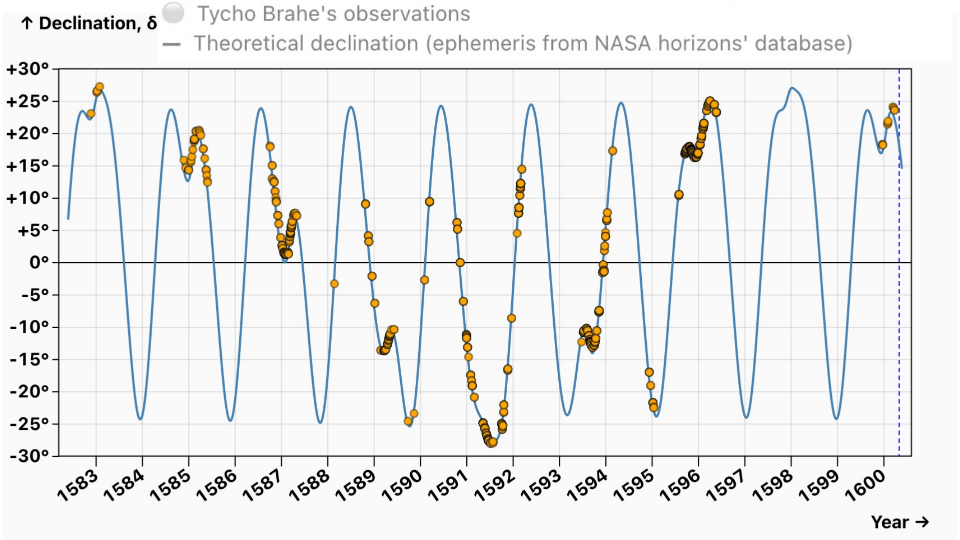

Kepler's Third Law

$$T^2 \propto r^3$$

c. The Astronomical Revolution: Copernicus- Kepler-Borelli

Symbolic Regression Algorithms

c. PySR. Miles Cranmer

Symbolic Regression Algorithms

c. PySR. Miles Cranmer

PySR's impact

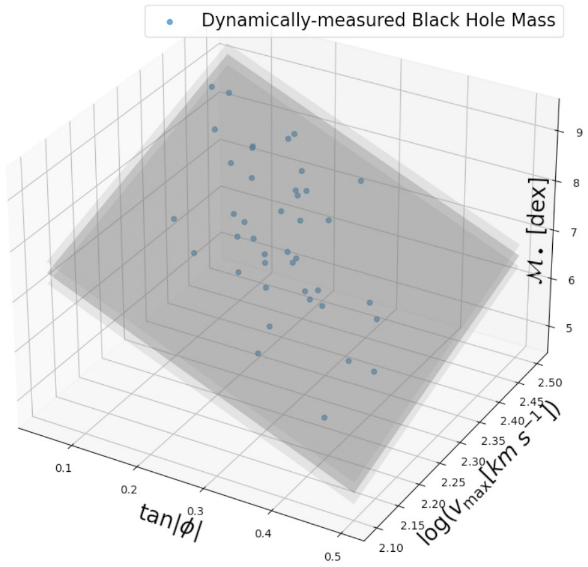

Discovery of a Planar Black Hole Mass Scaling Relation for Spiral Galaxies

Benjamin L. Davis [1], Zehao Jin [1]

[1]Center for Astrophysics and Space Science, New York University Abu Dhabi

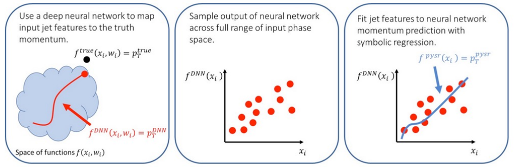

Interpretable machine learning methods applied to jet background subtraction in heavy-ion collisions

Tanner Mengel [1], Patrick Steffanic [1], Charles Hughes [1,2], Antonio Carlos Oliveira da Silva [1,2], Christine Nattrass [1]

[1]University of Tennessee, Knoxville, [2]Iowa State University of Science and Technology

Modeling the galaxy-halo connection with machine learning

Ana Maria Delgado [1], Digvijay Wadekar [2,3], Boryana Hadzhiyska [1], Sownak Bose [1,7], Lars Hernquist [1], Shirley Ho [2,4,5,6]

[1]Center for Astrophysics | Harvard & Smithsonian, [2]New York University, [3]Institute for Advanced Study, [4]Flatiron Institute, [5]Princeton University, [6]Carnegie Mellon University, [7]Durham University

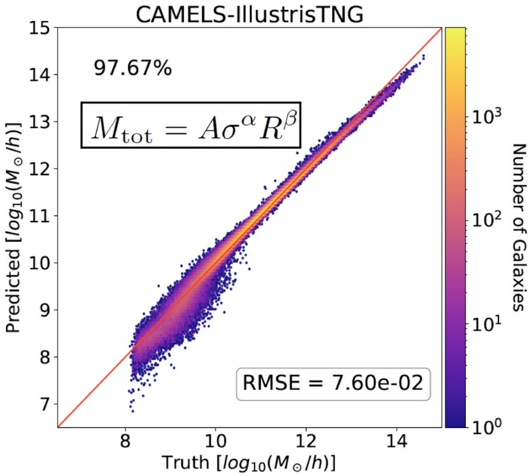

Finding universal relations in subhalo properties with artificial intelligence

Helen Shao [1], Francisco Villaescusa-Navarro [1,2], Shy Genel [2,3], David N. Spergel [2,1], Daniel Angles-Alcazar [4,2], Lars Hernquist [5], Romeel Dave [6,7,8], Desika Narayanan [9,10], Gabriella Contardo [2], Mark Vogelsberger [11]

[1]Princeton University, [2]Flatiron Institute, [3]Columbia University, [4]University of Connecticut, [5]Center for Astrophysics | Harvard & Smithsonian, [6]University of Edinburgh, [7]University of the Western Cape, [8]South African Astronomical Observatories, [9]University of Florida, [10]University of Florida Informatics Institute, [11]MIT

c. PySR. Miles Cranmer

Sketch of PySR's Exploration Space

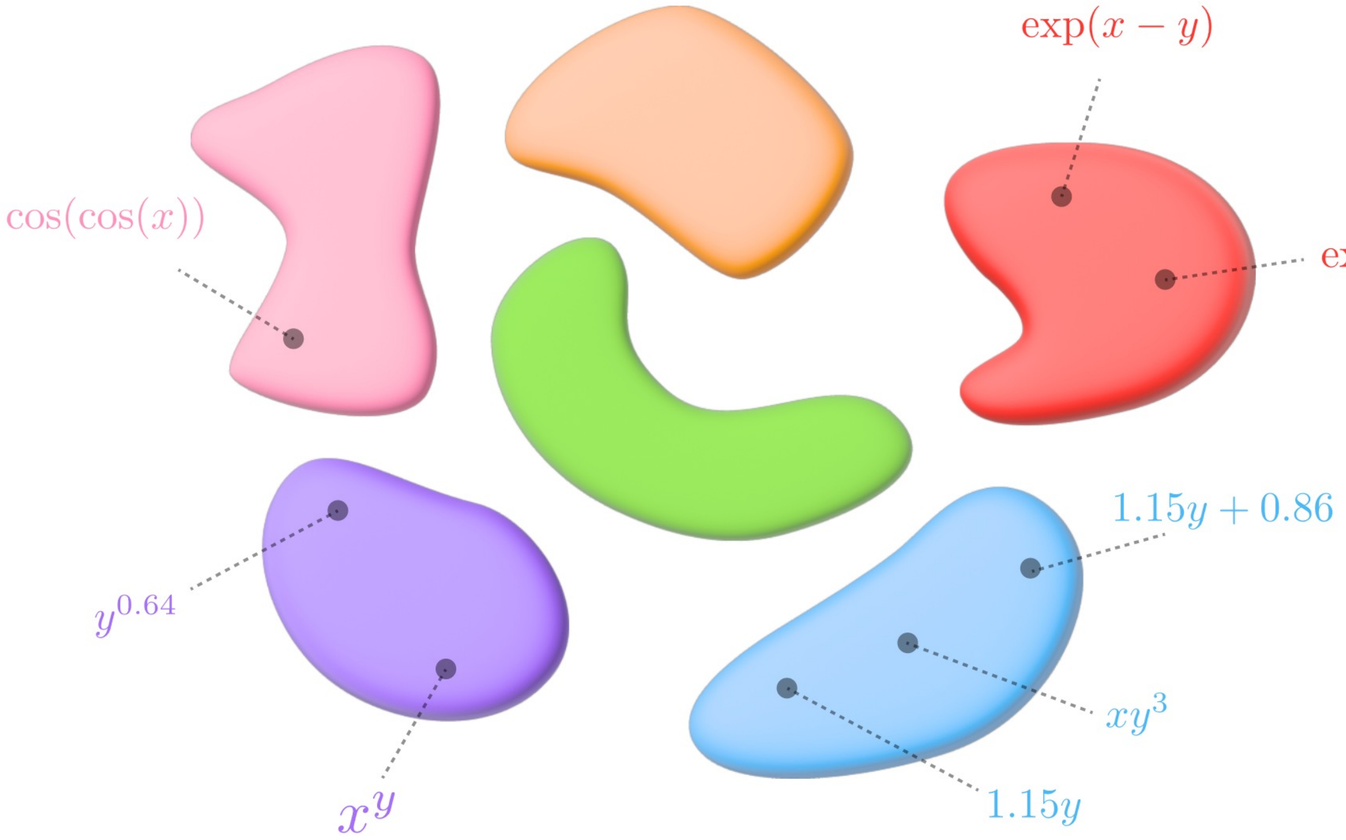

Independent "islands" of expressions, each undergoing evolution

c. PySR. Miles Cranmer

Sketch of PySR's Exploration Space

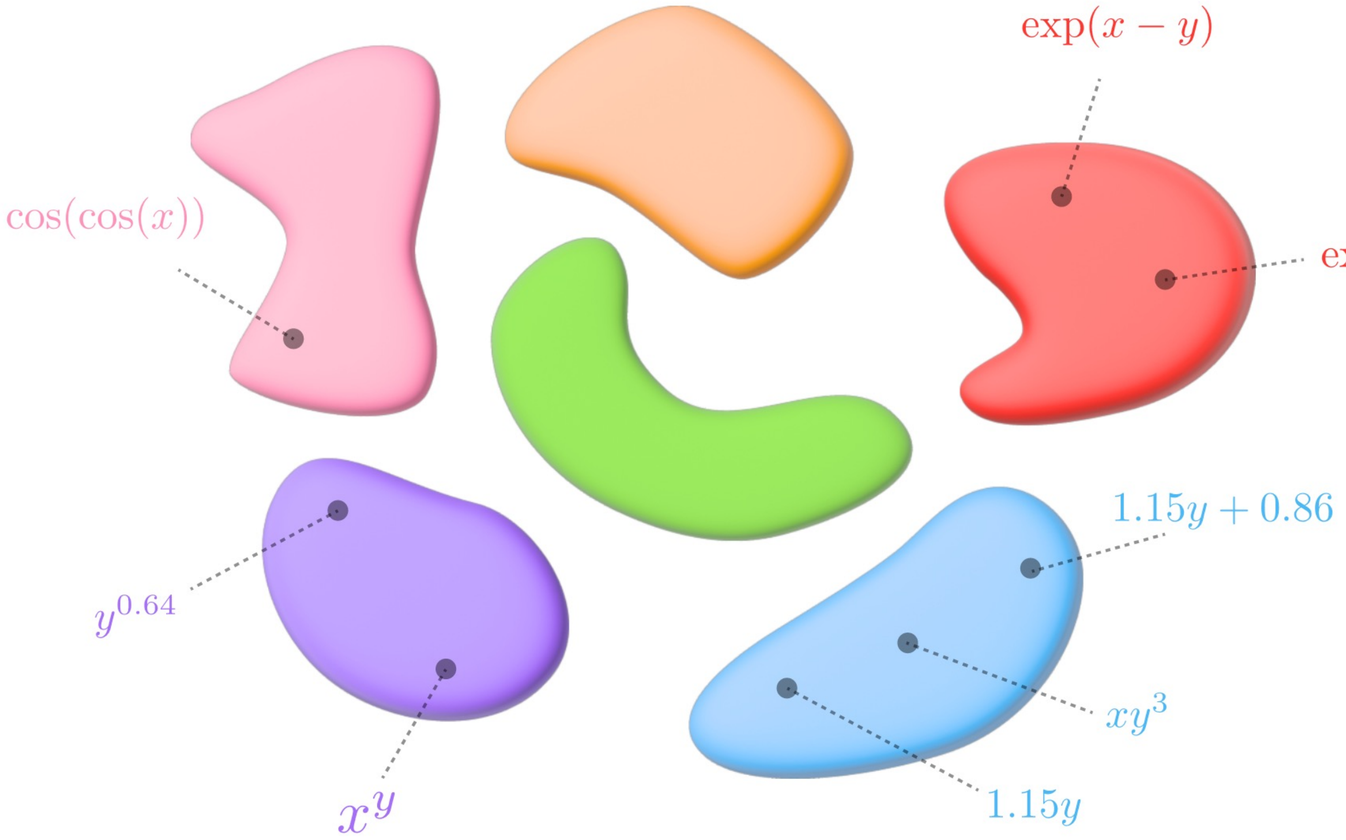

Independent "islands" of expressions, each undergoing evolution

Goal: How can we increase exploration in relevant parts of the search space?

LaSR: Symbolic Regression with a Learned Concept Library

Arya Grayeli*, **Atharva Sehgal***, Omar Costilla-Reyes, Miles Cranmer, Swarat Chaudhuri

What is a Concept?

Desiderata 1: Symbolic Abstraction

Zipf's Law

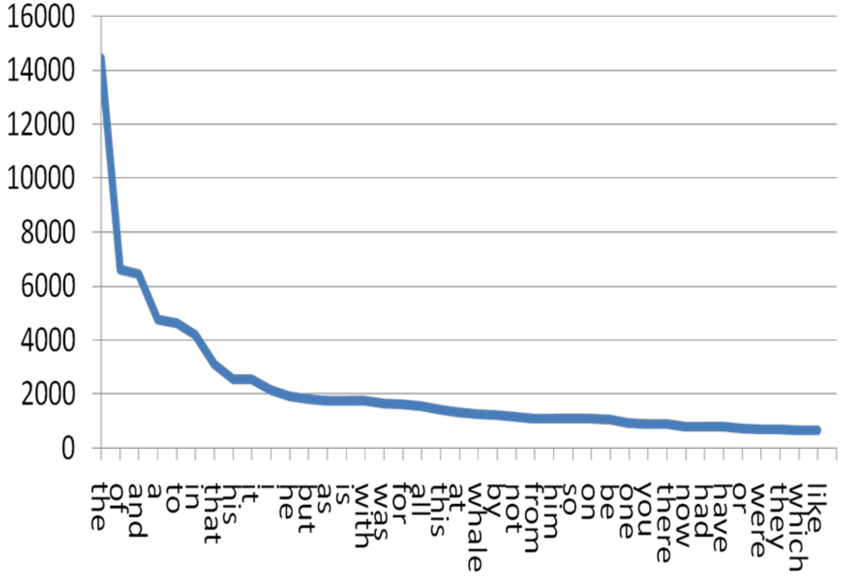

Moore's Law

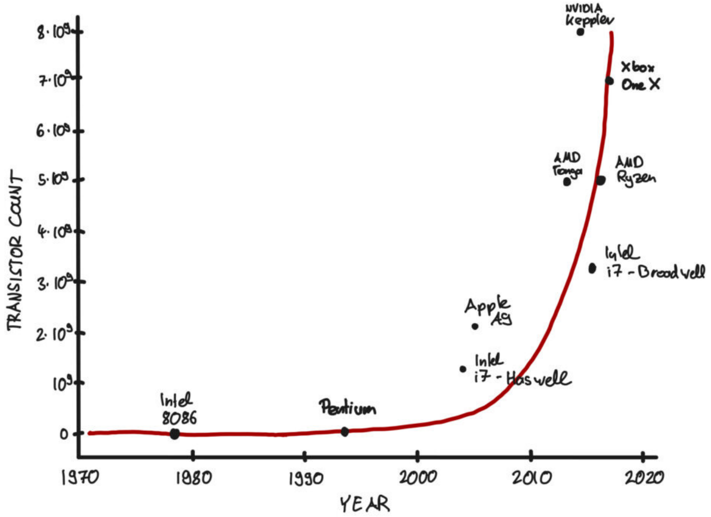

Arrhenius' Equation

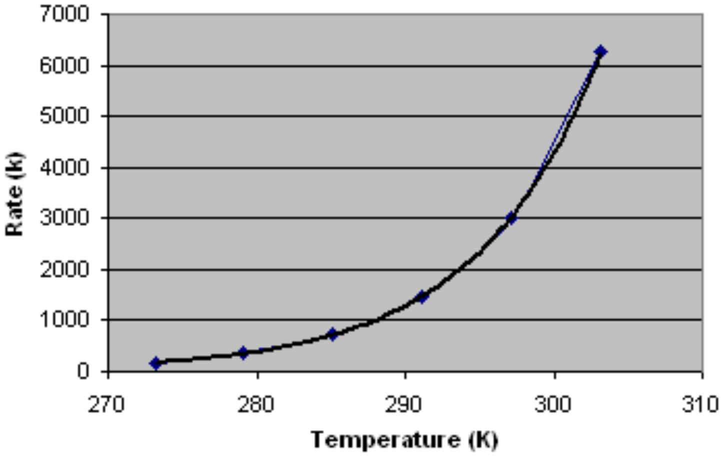

$$y = ax^k + \epsilon$$ ⇔ *"Power Law Trend"*

What is a Concept?
Desiderata II : Symbolic Guidance

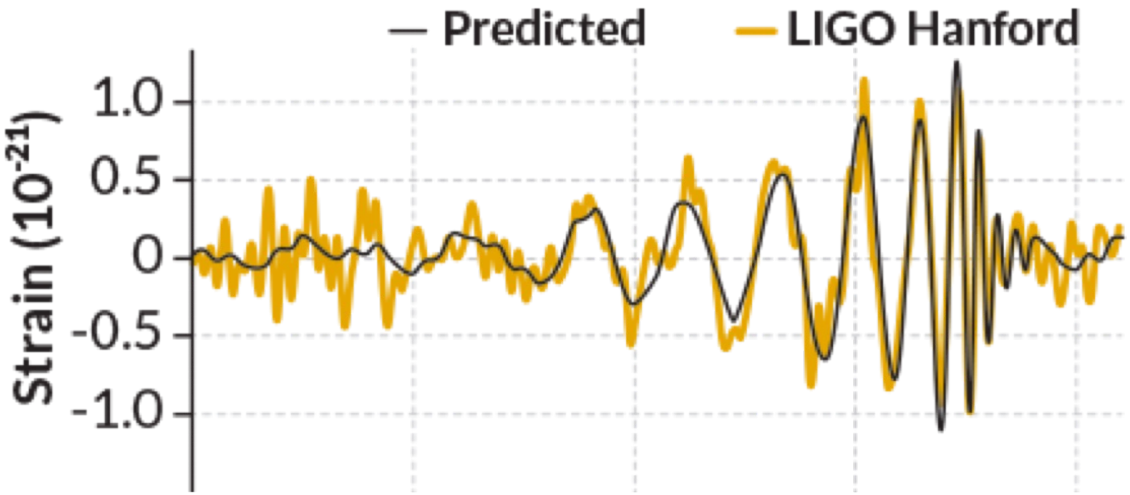

Concepts (by Physicist or LLM)

"Wave strain diminishes as distance increases"

"Wave strain has extraordinarily small magnitude"

Guide the search for

$$h = \frac{2G}{c^4}\frac{1}{r}\frac{\partial^2 Q}{\partial t^2}$$

Hypothesis Evolution

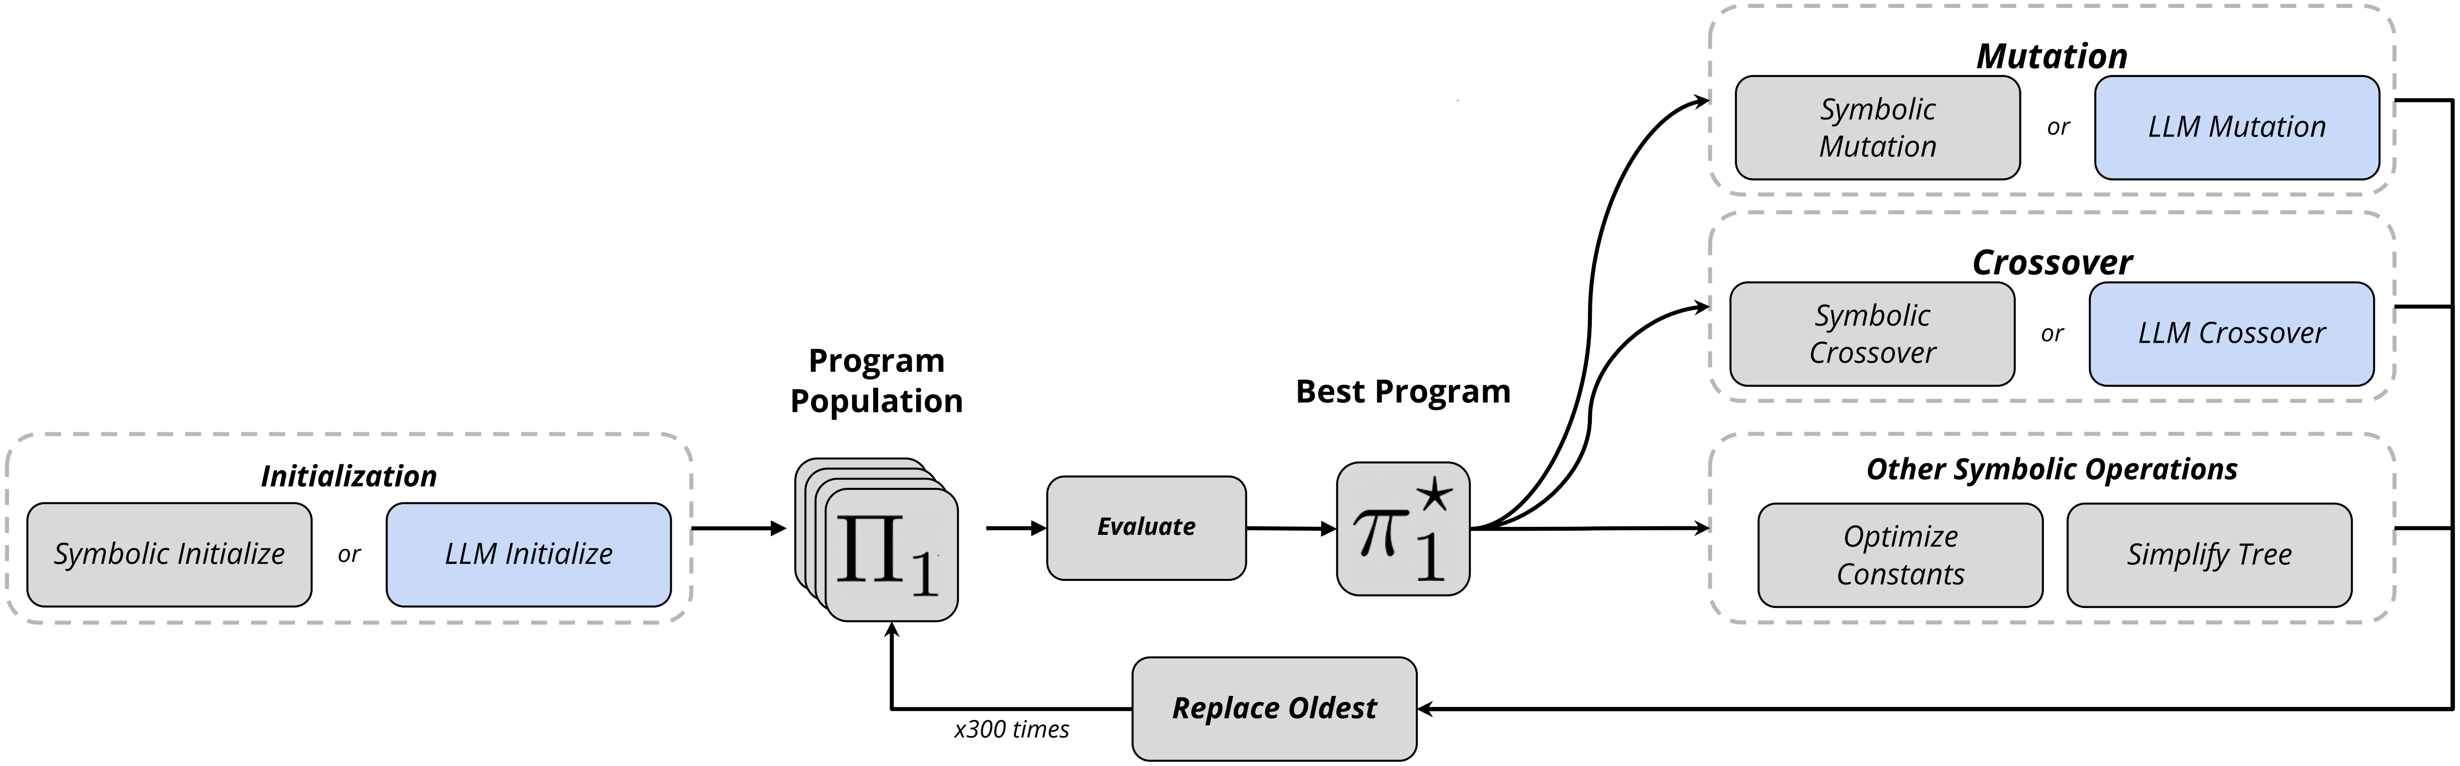

Hypothesis Evolution

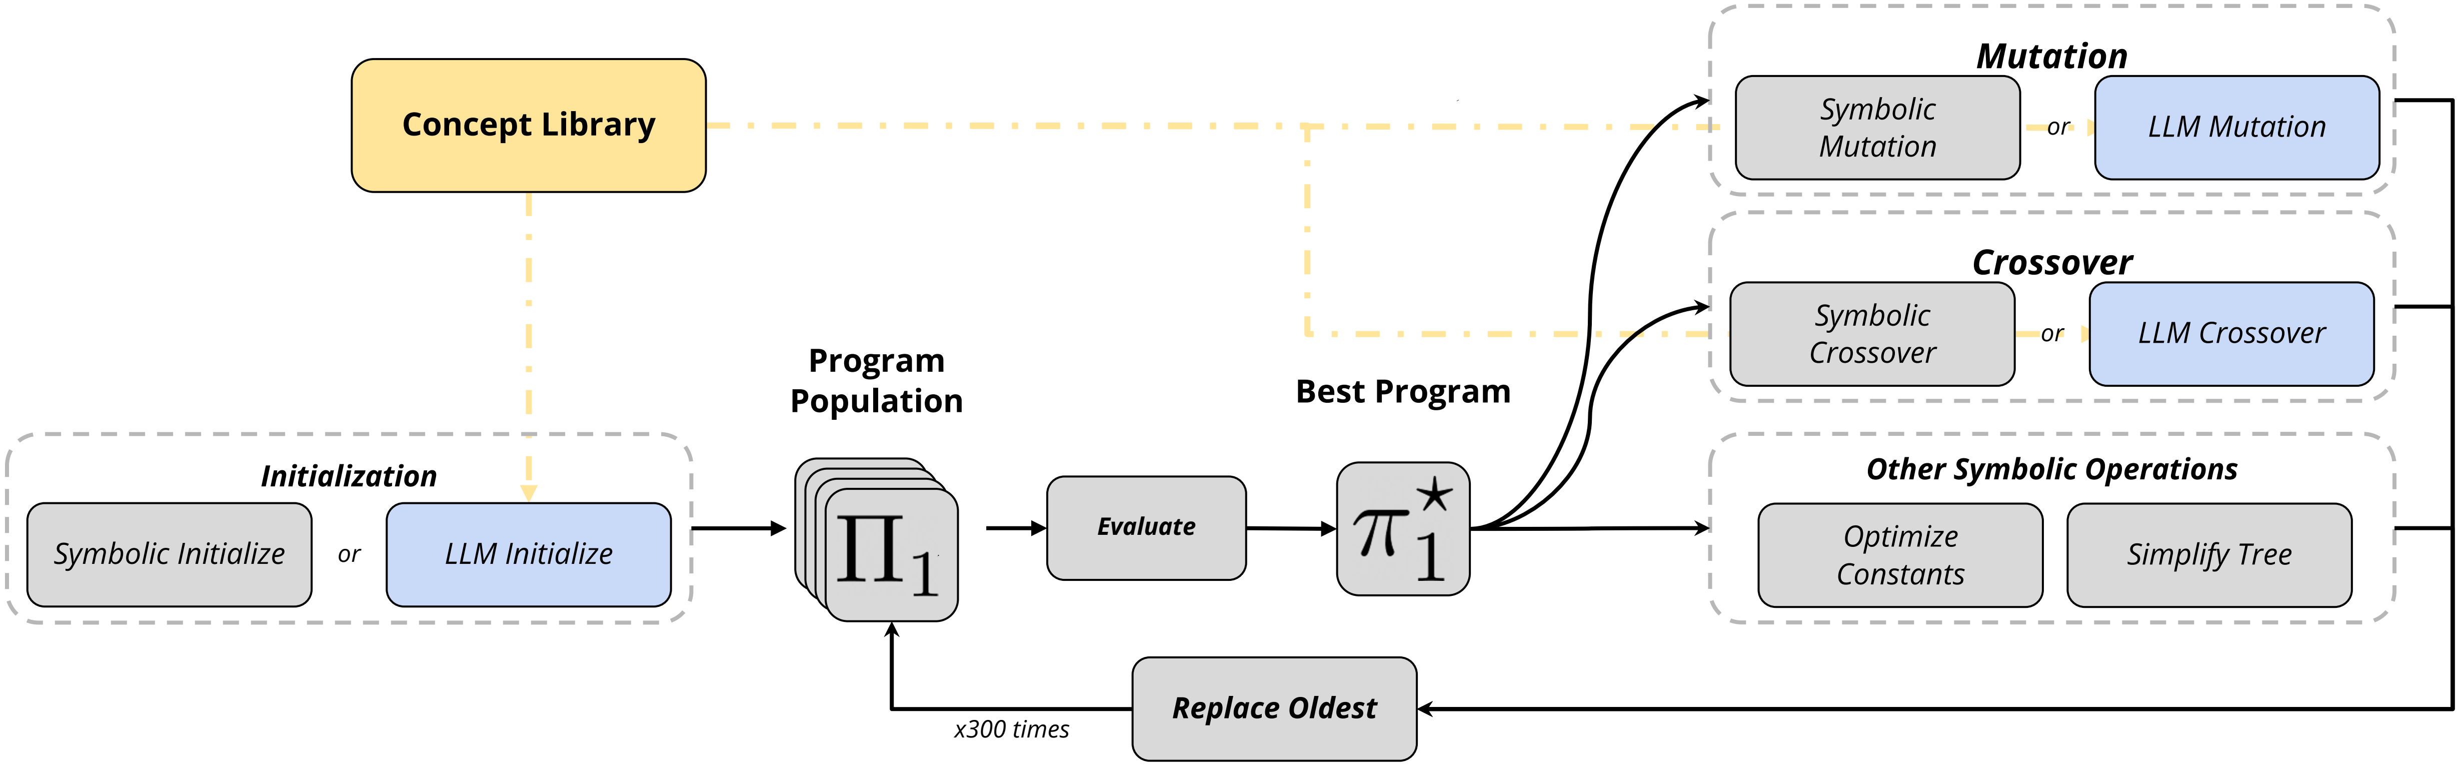

Hypothesis Evolution

```
import random

def crossover(
        expr1: SymTree,
        expr2: SymTree) -> SymTree:
# Randomly choose a node in expr1 to
remove
...
# Randomly choose a node in expr2
which will be added to eq1
...
# Return new tree
new_expr = ...
return new_expr
```

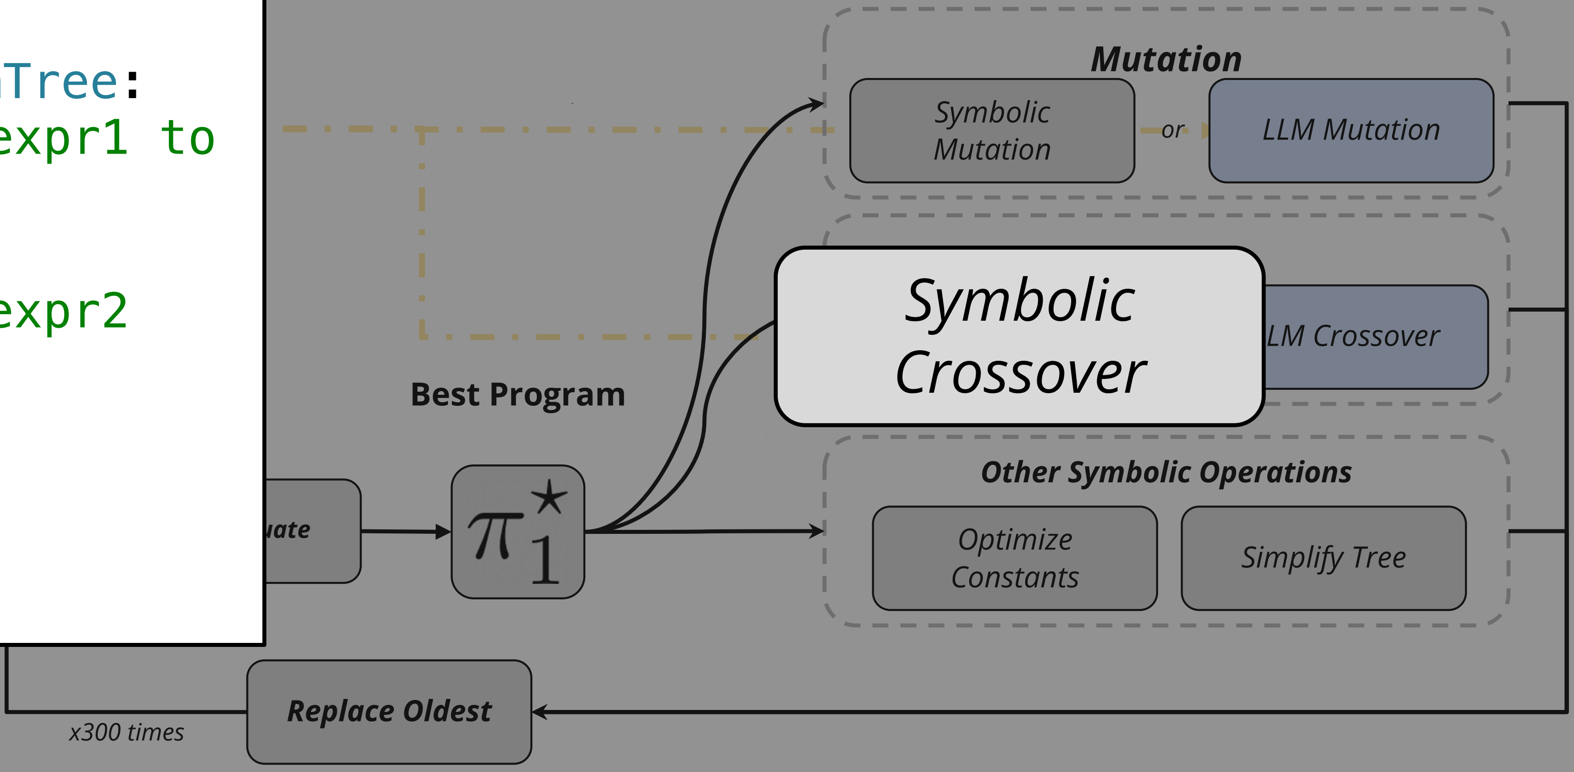

Hypothesis Evolution

(System) Header

You are a helpful assistant that recombines two mathematical expressions by following a few provided suggestions. You will be given three suggestions and two expressions to recombine.

An expression must consist of the following variables: `{{variables}}`. All constants will be represented with the symbol C. Each expression will only use these operators: `{{operators}}`.

(User) Crossover Prompt

Suggestion 1: `{{assump1}}`
Suggestion 2: `{{assump2}}`
Suggestion 3: `{{assump3}}`
Expression 1: `{{expr1}}`
Expression 2: `{{expr2}}`

Propose {{N}} expressions that would be appropriate given the suggestions and expressions. Provide short commentary for each of your decisions. End with a JSON list that enumerates the proposed expressions following this format:

(User) JSON Formatting Instructions

````
```json
["expr1",
 "expr2",
 ...
 "expr{{N}}"
]
```
````

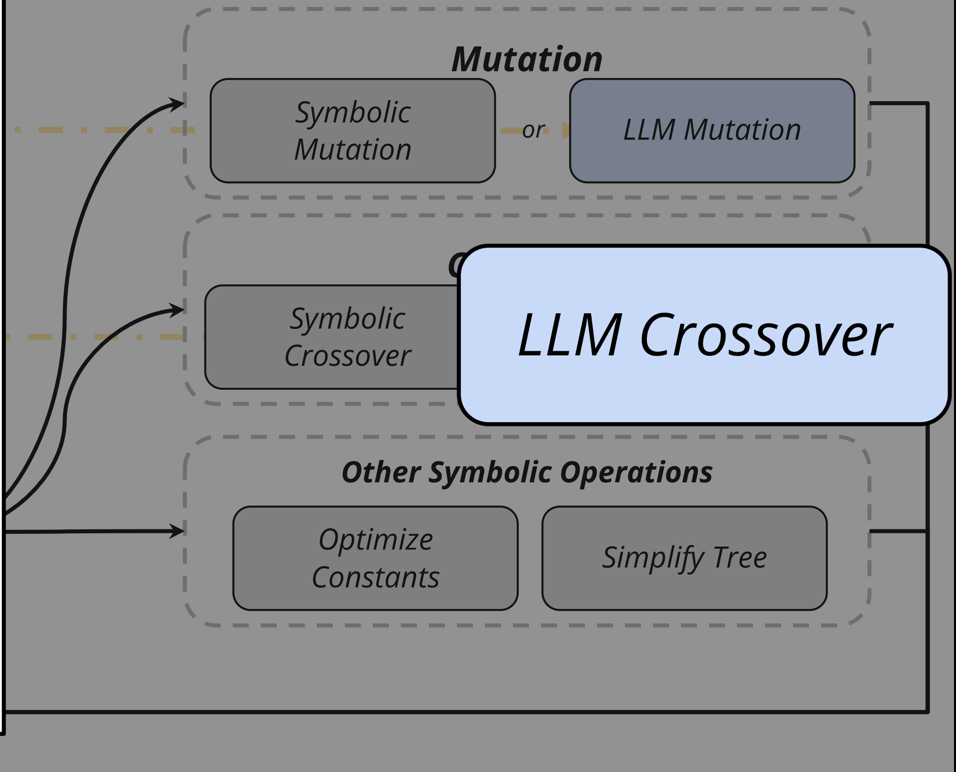

LaSR

Hypothesis Evolution

Concept Evolution

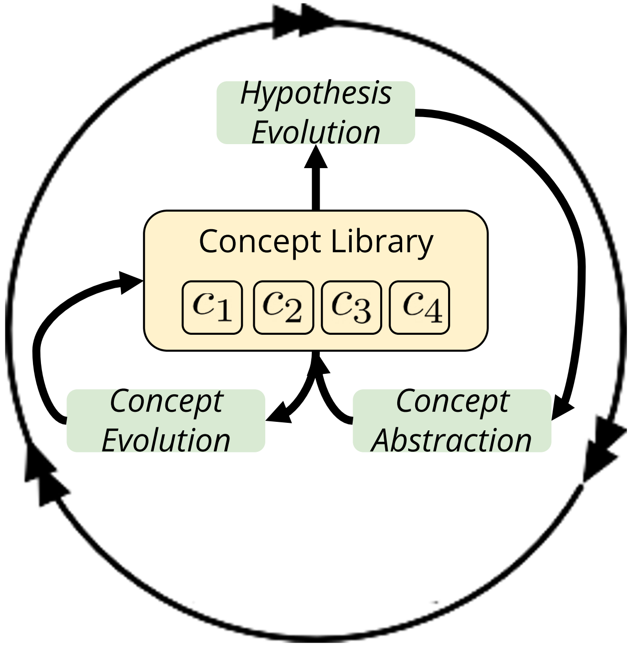

Concept Abstraction

LaSR

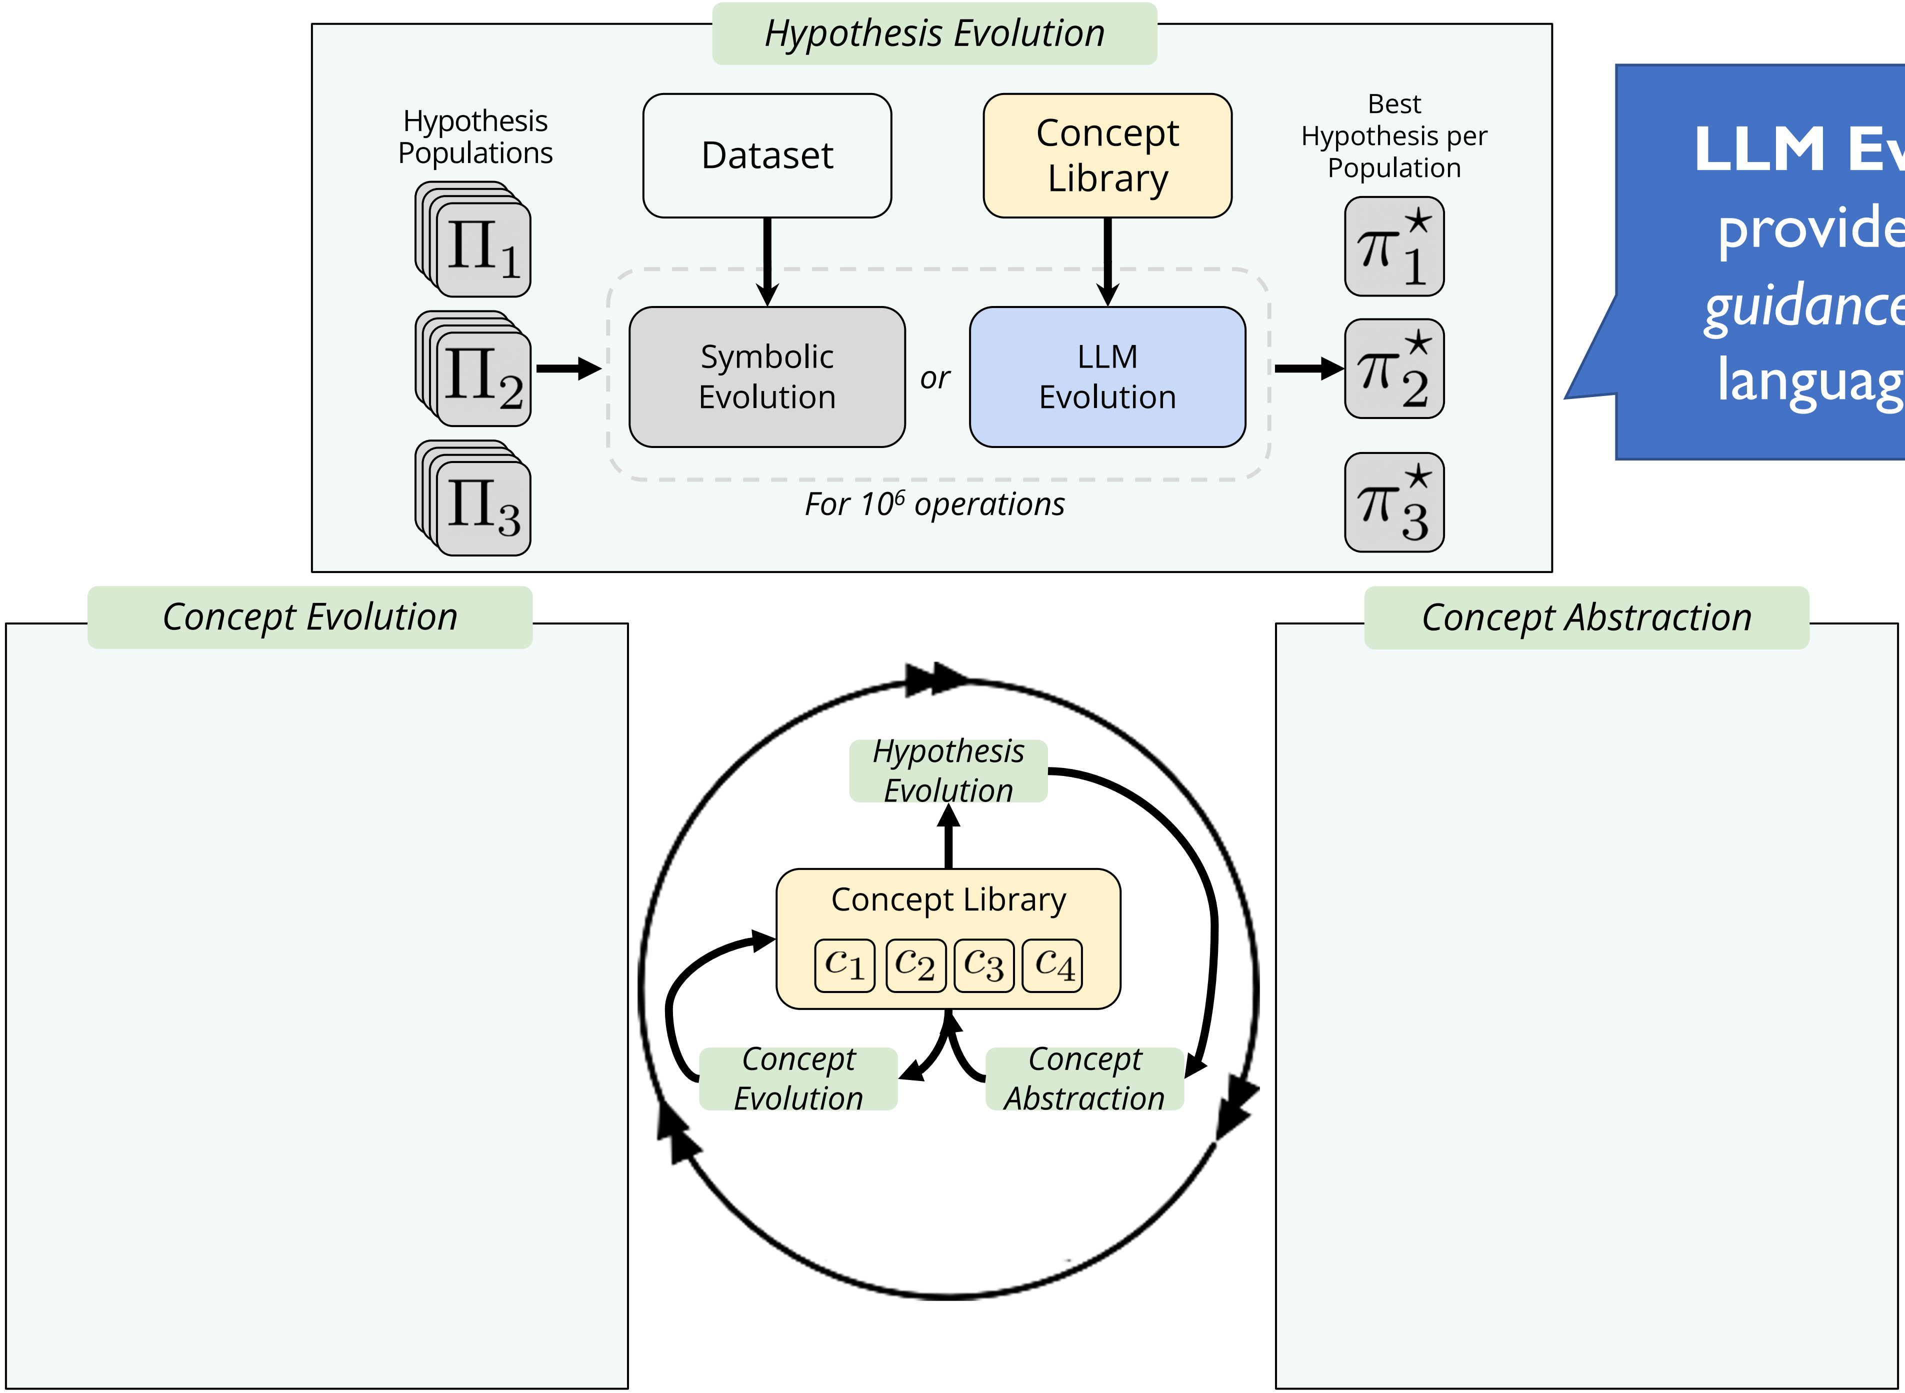

LLM Evolution provides *neural guidance* (over a language prior)

LaSR

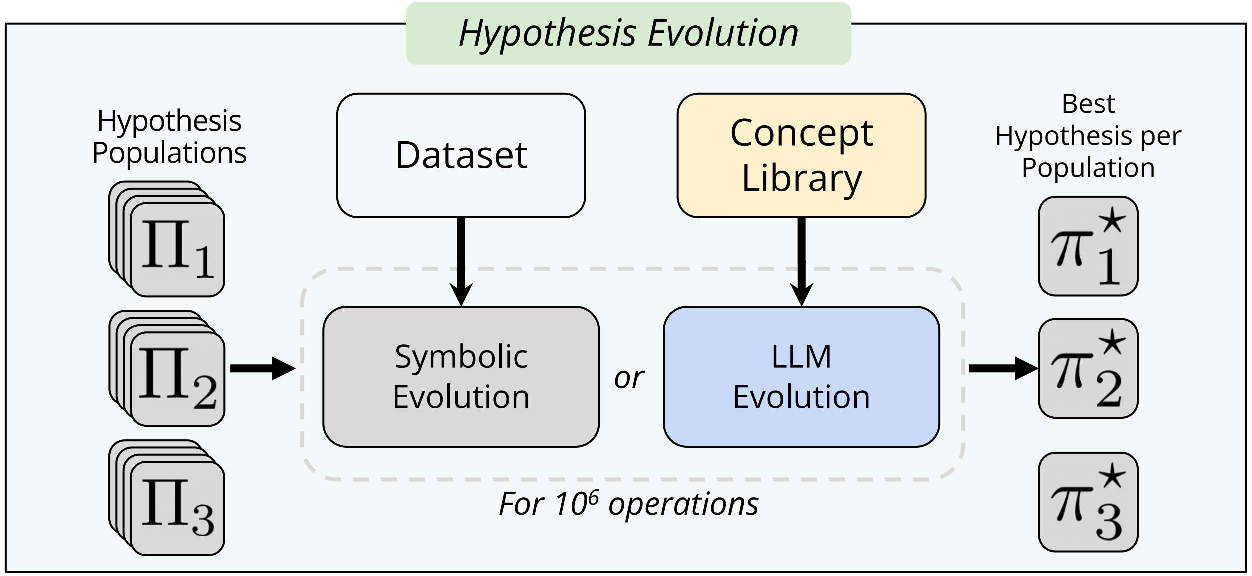

LLM Abstraction induces *useful** abstractions.

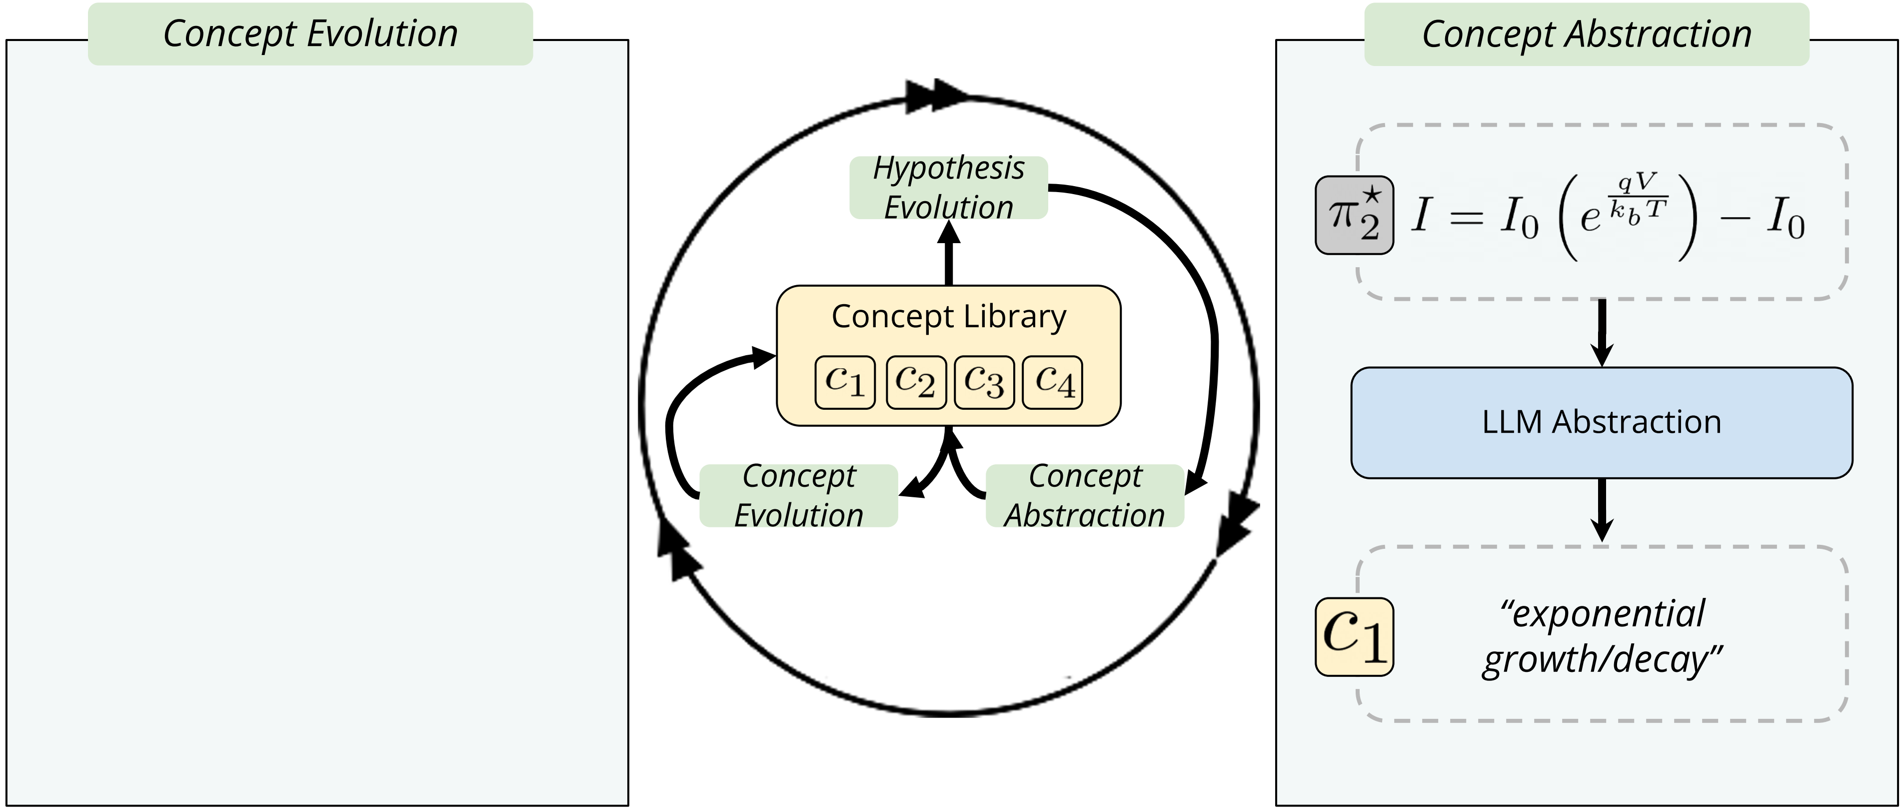

LaSR

LLM Concept Crossover evolves all concepts.

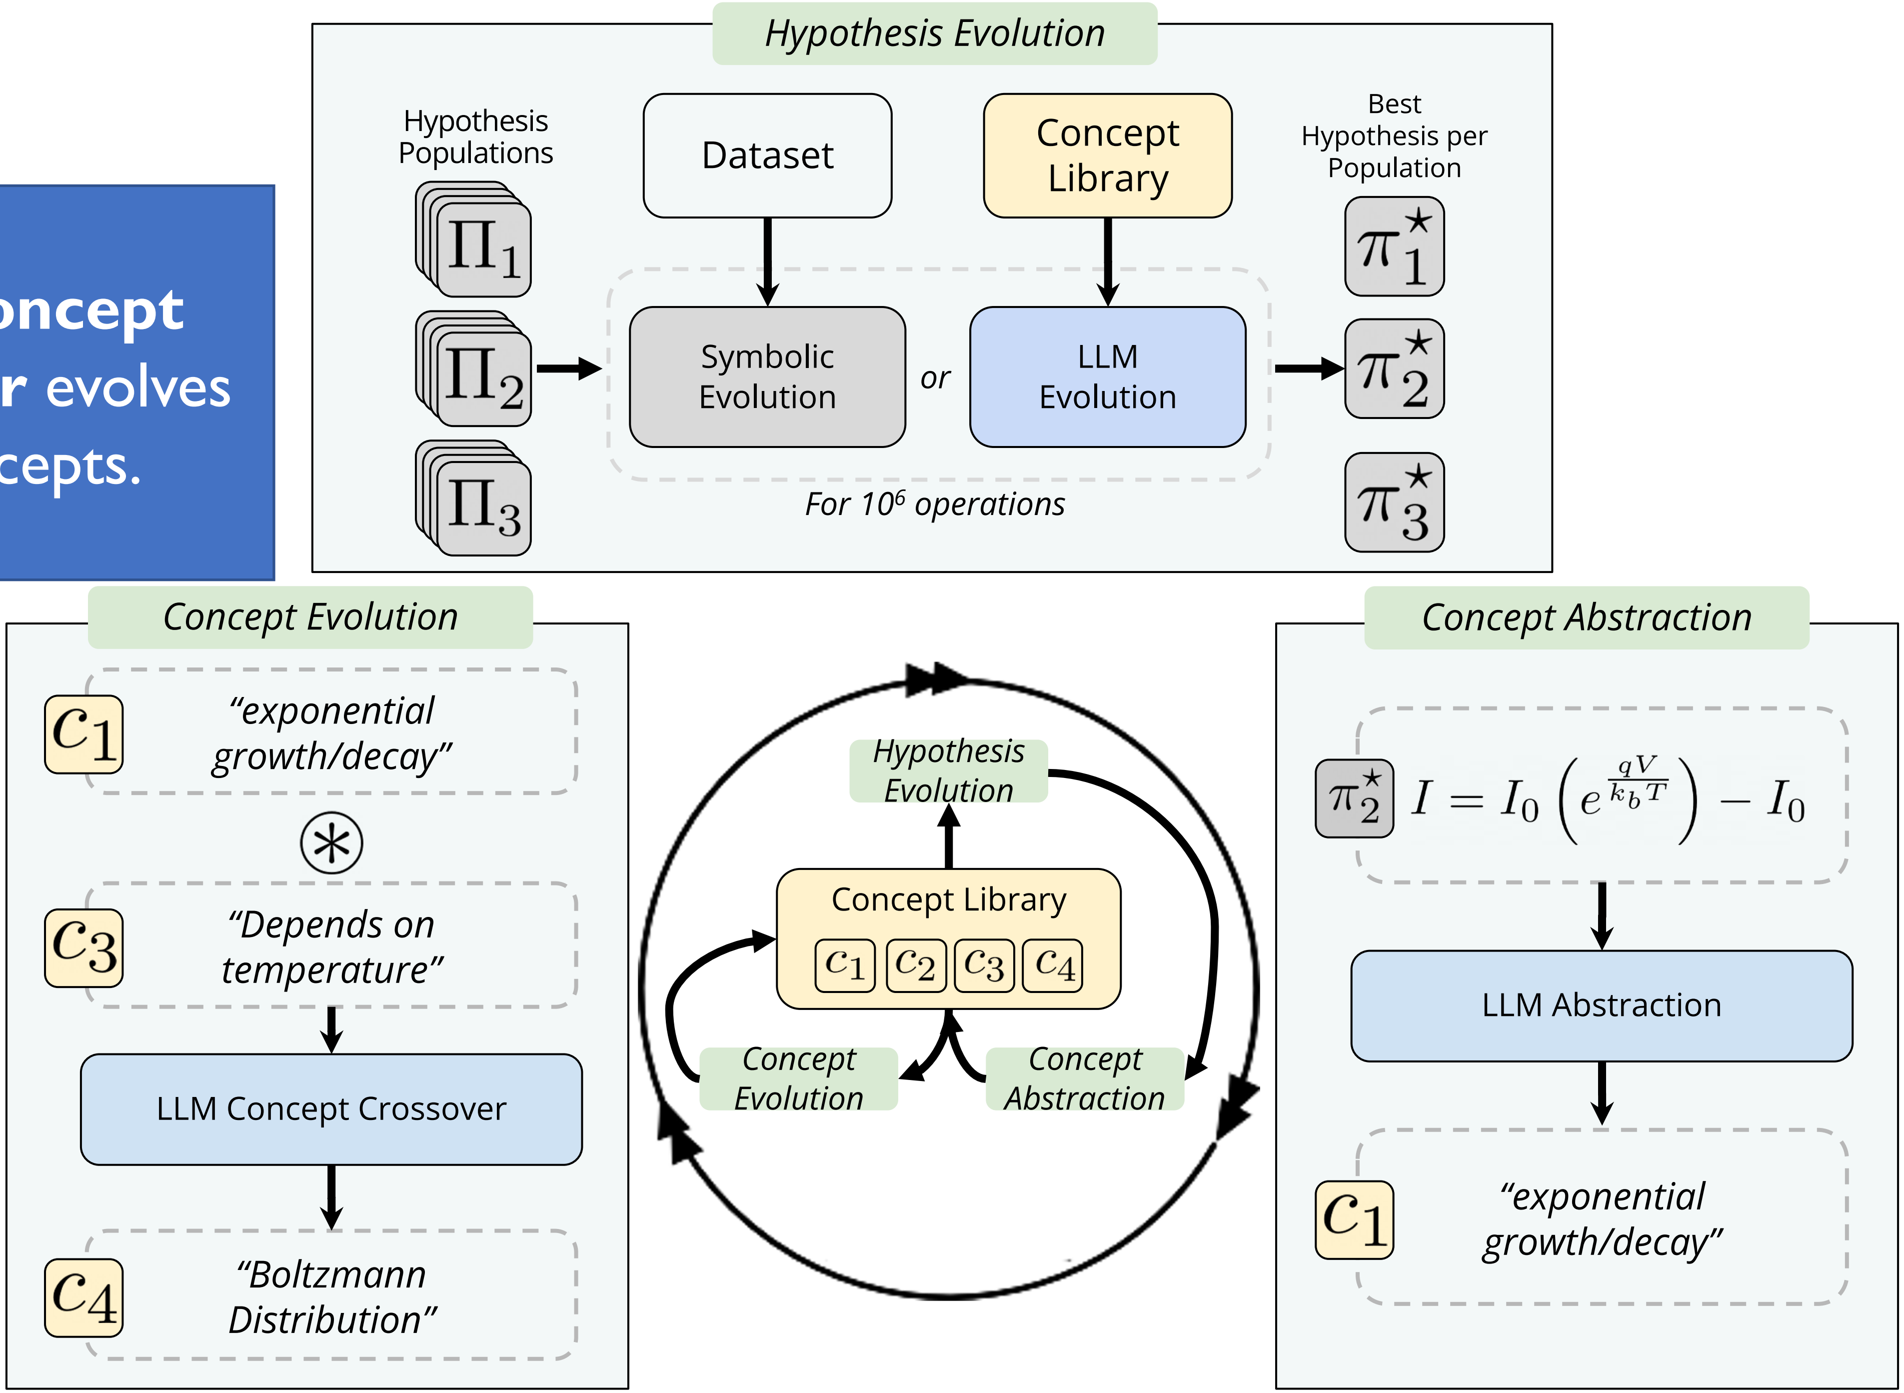

Sketch of Search Space

After Phase 1:
"Islands" of expressions

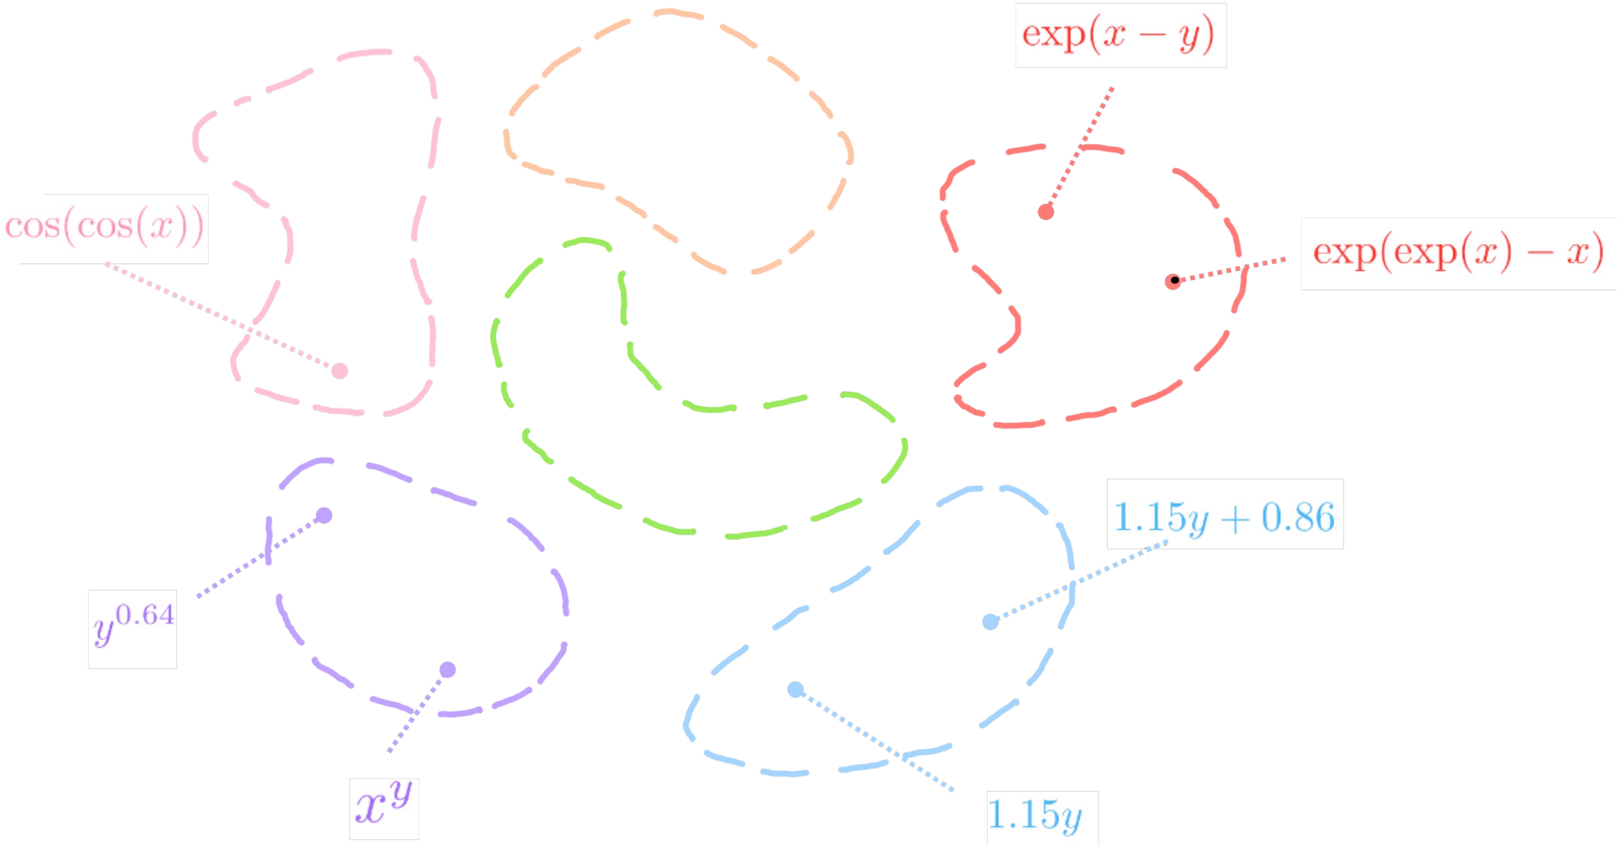

Sketch of Search Space

After Phase 2:
Concepts for each "Island"

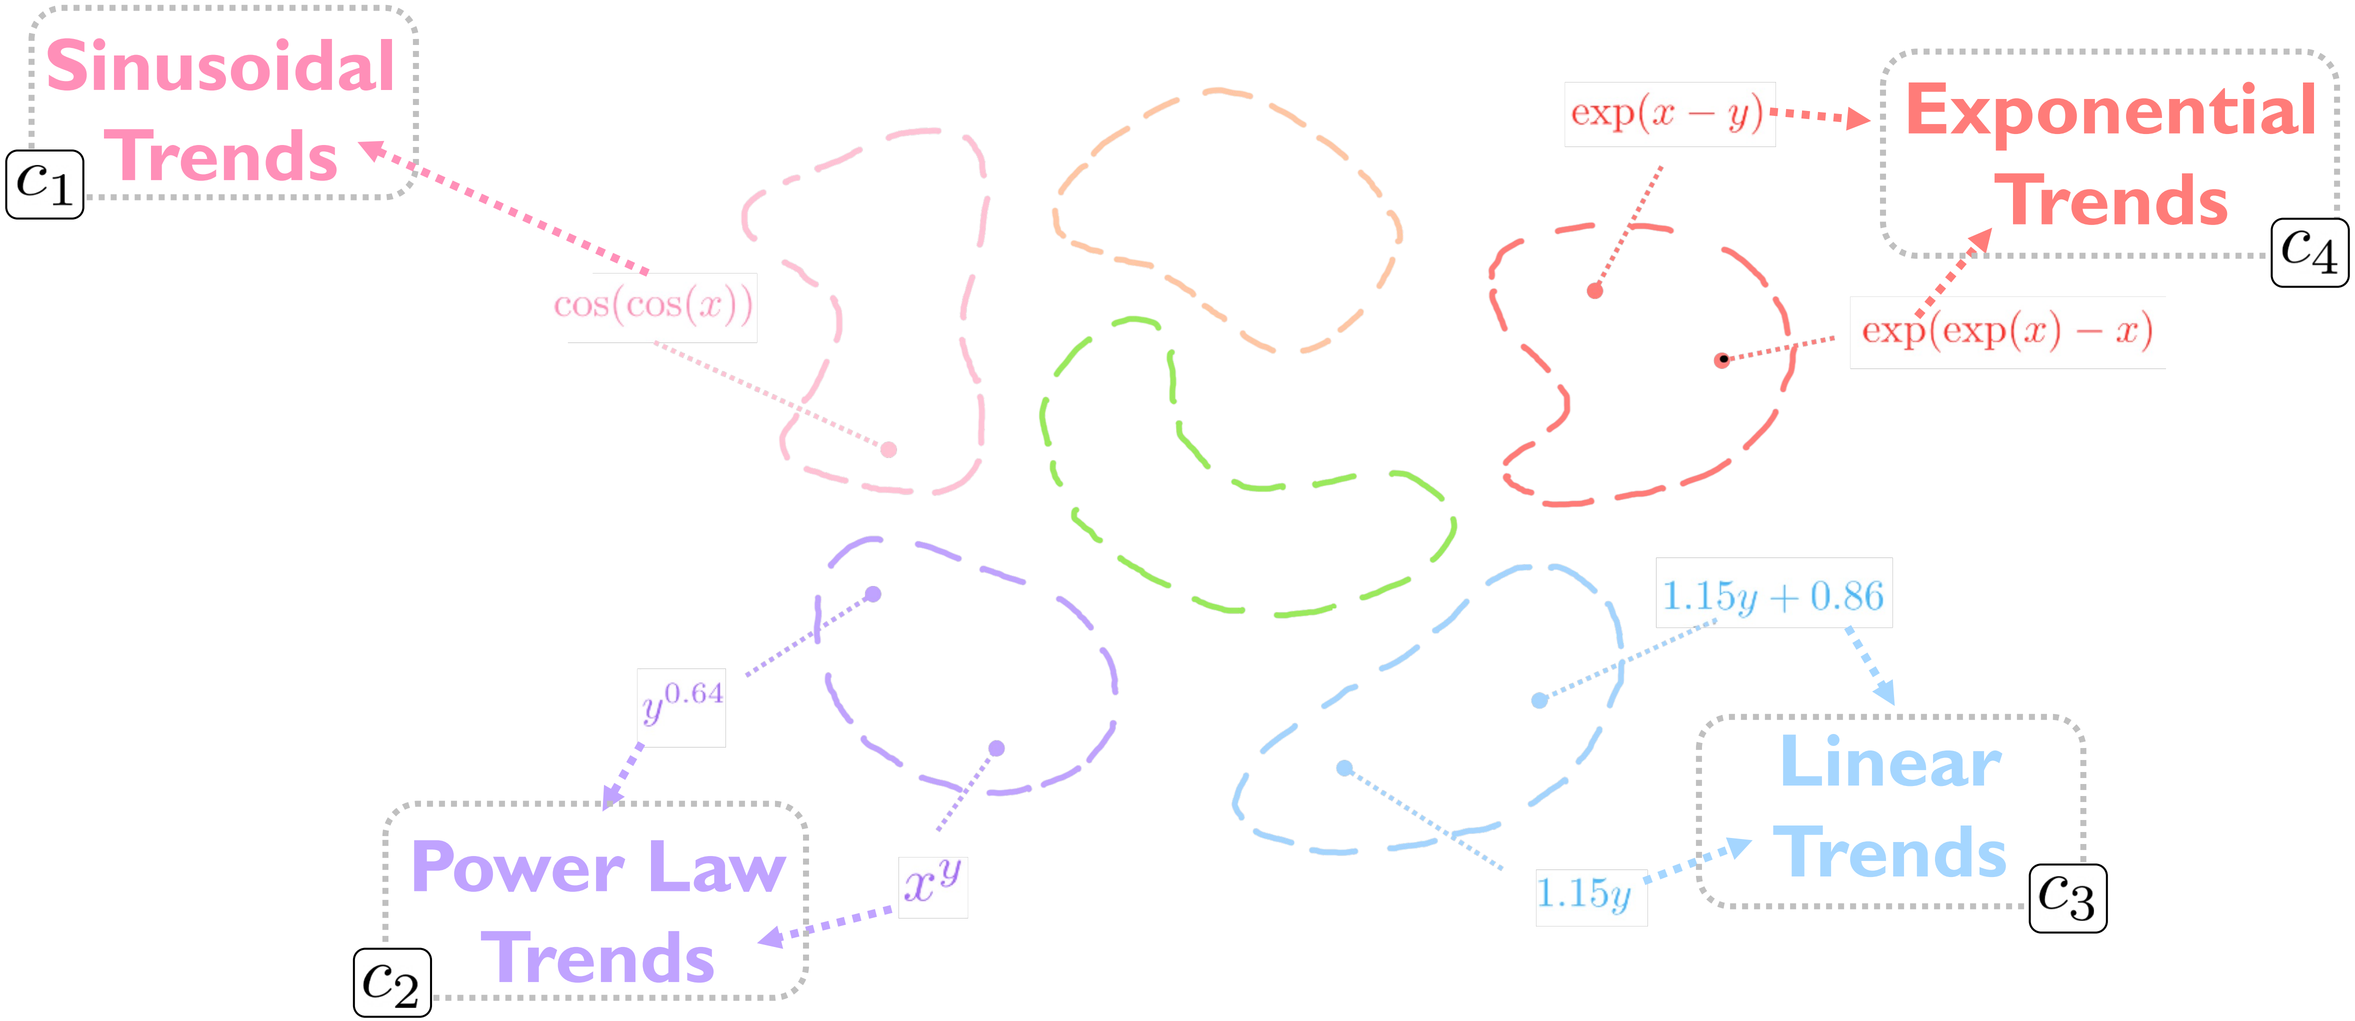

Sketch of Search Space

After Phase 2:
Concepts expose unknown islands

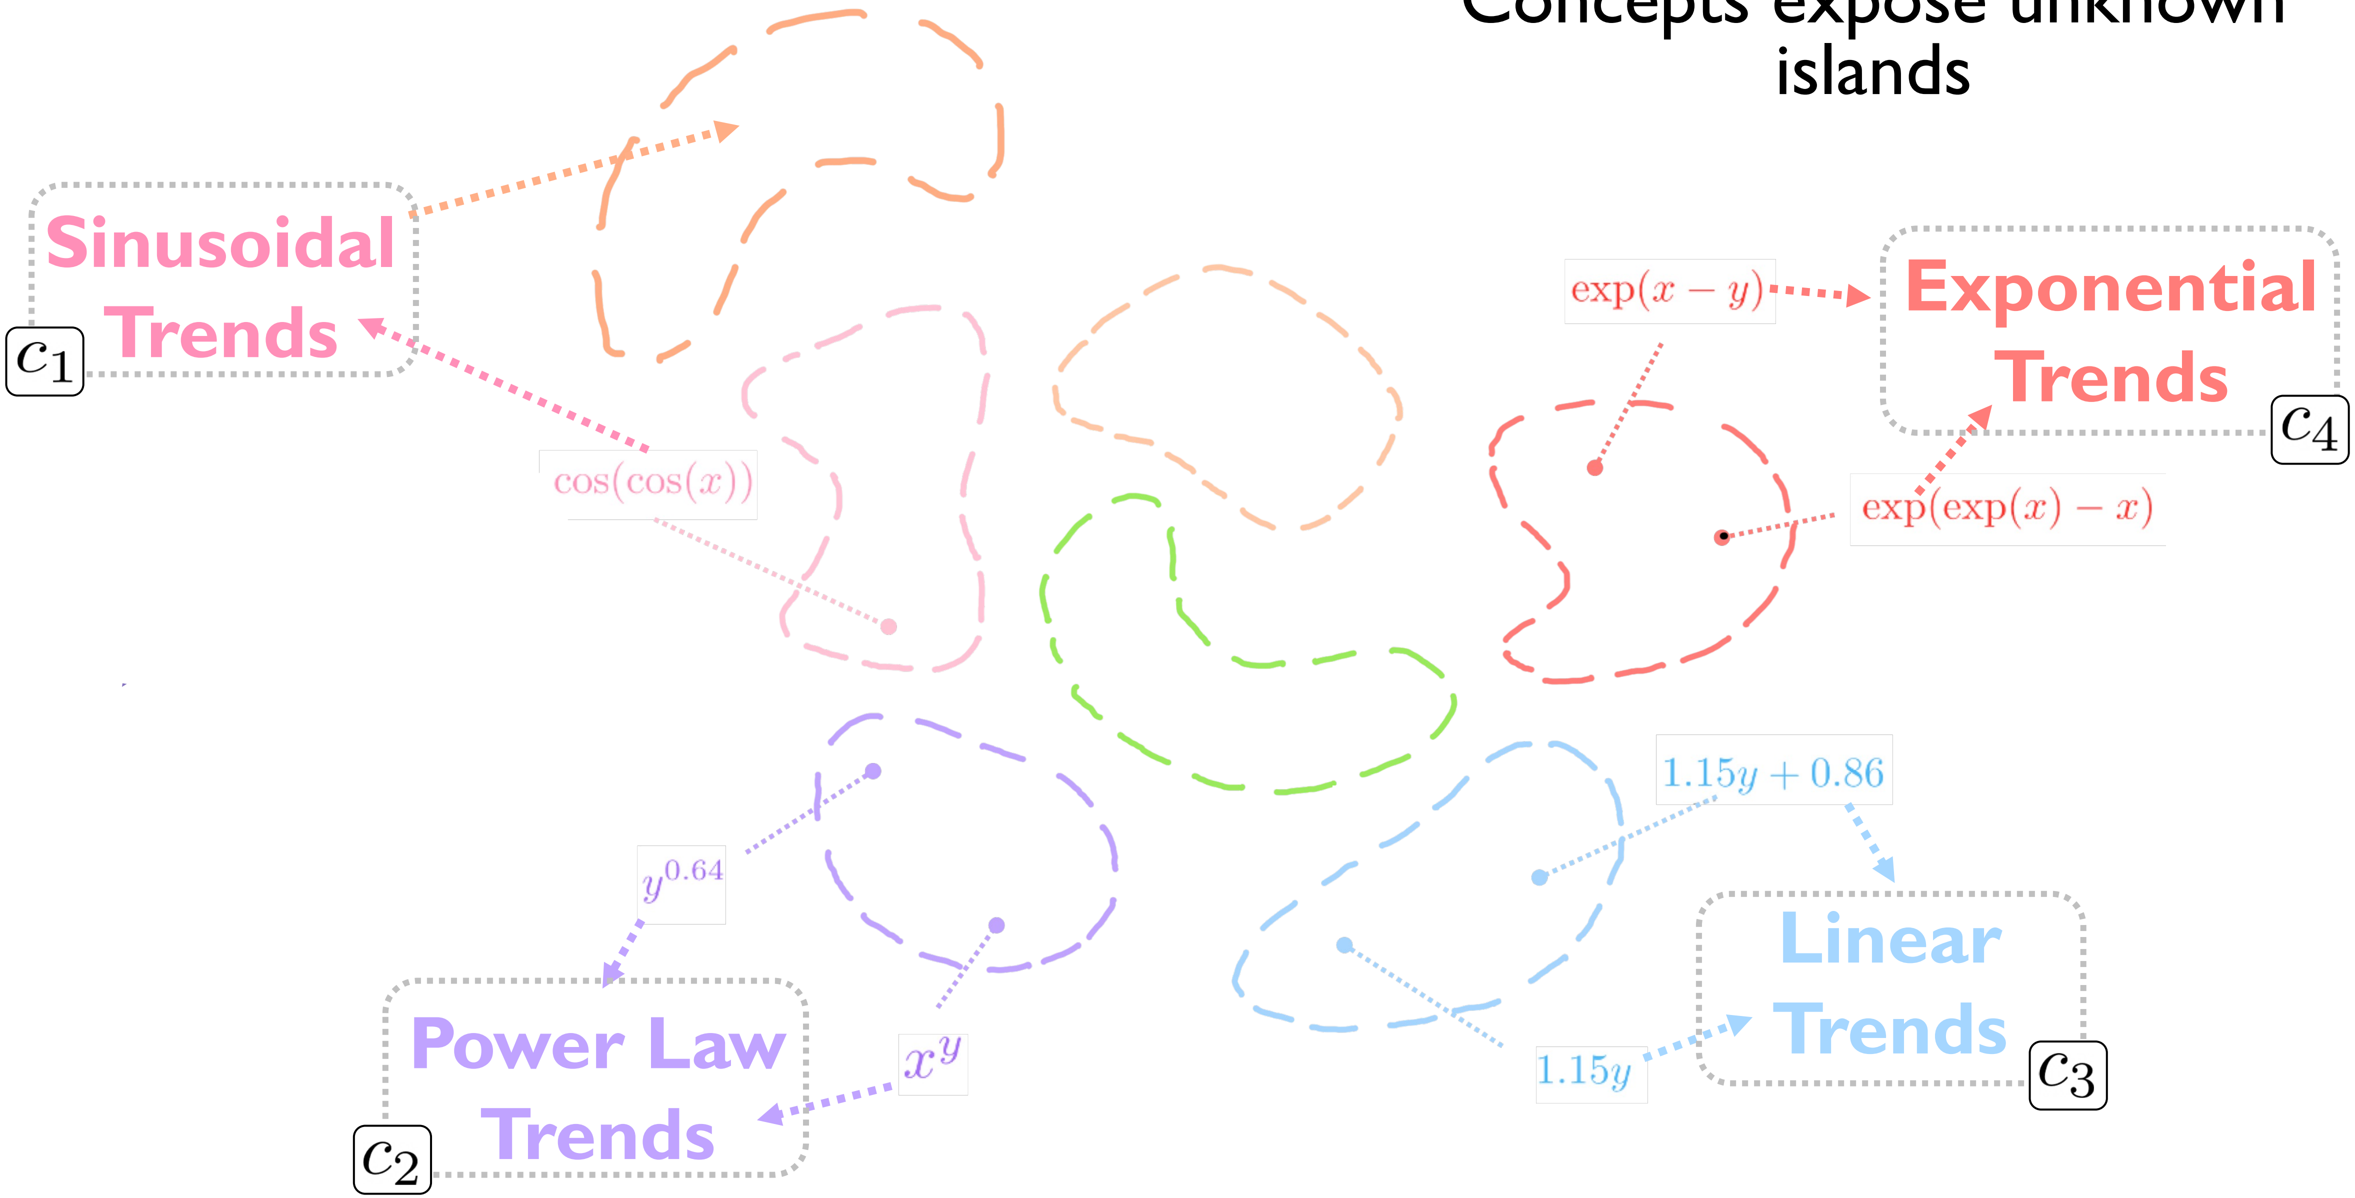

Sketch of Search Space

After Phase 3:
Concepts expose new concepts.

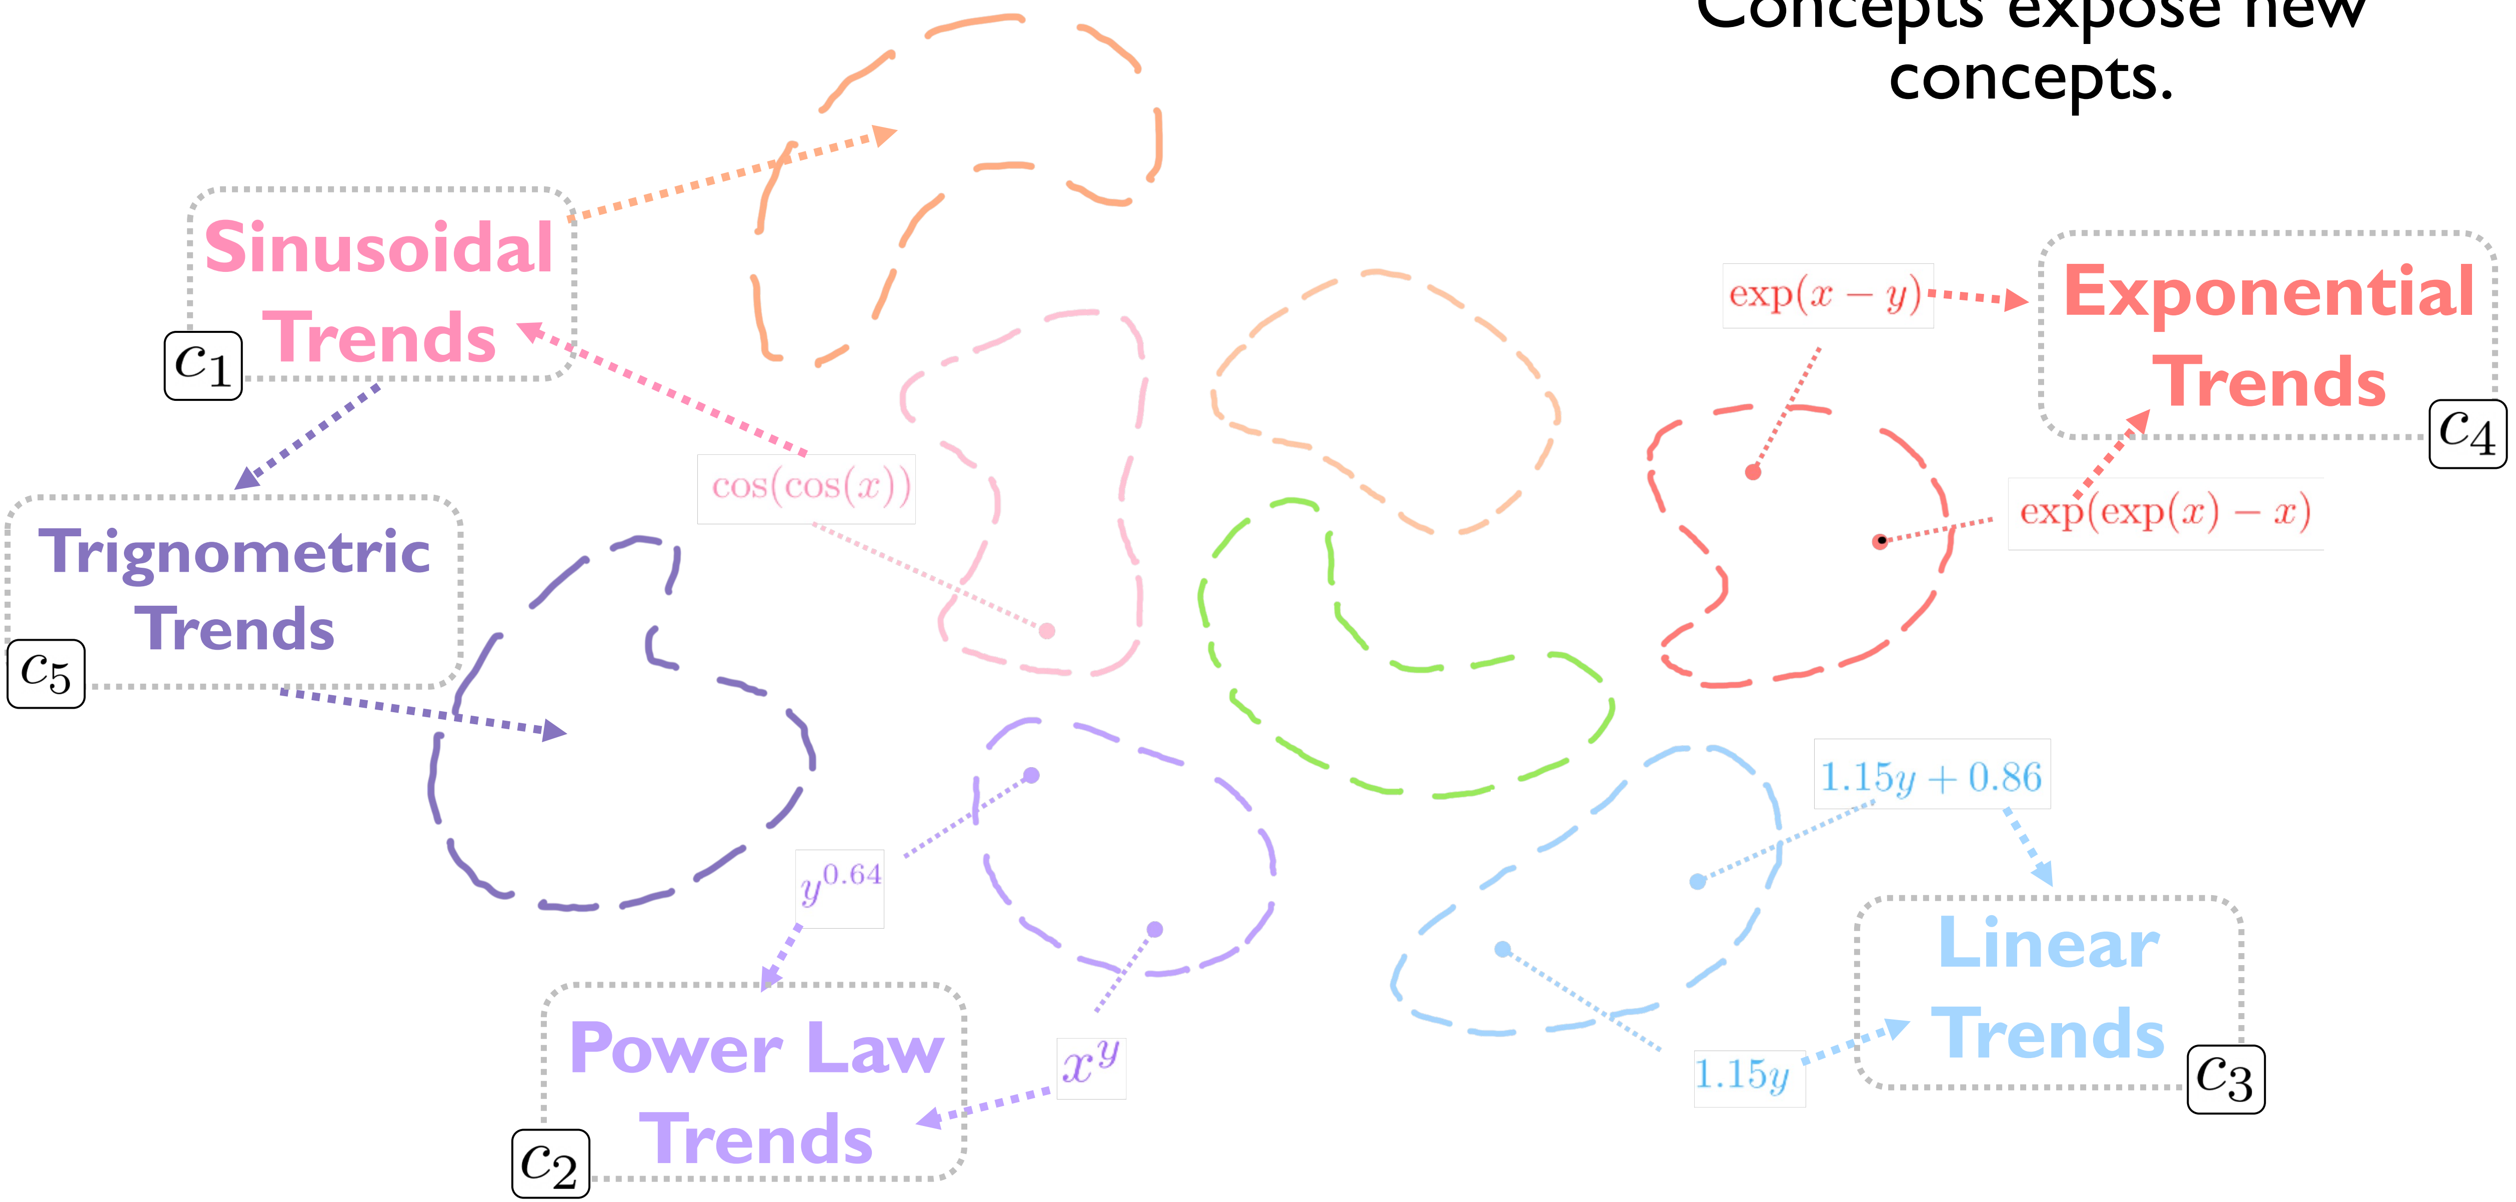

LaSR Results I - Performance

- Concept Guidance accelerates scientific discovery.

GPlearn	AFP	AFP-FE	DSR	uDSR	AIFeynman	PySR	LaSR
20/100	24/100	26/100	23/100	40/100	38/100	59/100	**59 + 7/100**

Table 1: Results on 100 Feynman equations from [41]. We report exact match solve rate for all models. LASR achieves the best exact match solve rate using the same hyperparameters as PySR [8].

- LaSR outperforms PySR even with local language models (llama-3-7b, 1%)

		LASR (Llama3-8B)			LASR (GPT-3.5)
Type of Solve	**PySR**	$p = 1\%$	$p = 5\%$	$p = 10\%$	$p = 1\%$
Exact Solve	59/100	63/100	65/100	65/100	66/100
Almost Solve	7/100	6/100	9/100	12/100	13/100
Close	16/100	13/100	14/100	11/100	9/100
Not Close	18/100	18/100	12/100	13/100	13/100

Table 2: Evaluation results on Feynman dataset by cascading LASR's LLM backbone (llama3-8b, gpt-3.5-turbo) and changing the probability of calling the model (p = [0.01, 0.05, 0.10]) in the order of increasing concept guidance. LASR outperforms PySR even with minimal concept guidance using an open-source LLM.

What is a Concept?
Desiderata II : Symbolic Guidance

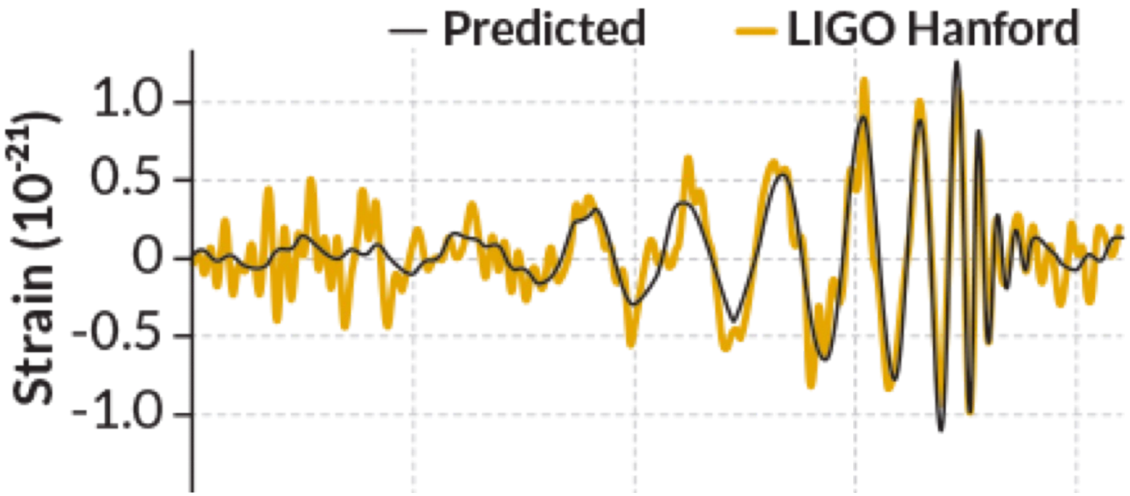

Concepts (by Physicist or LLM)

"Wave strain diminishes as distance increases"

"Wave strain has extraordinarily small magnitude"

Guide the search for

$$h = \frac{2G}{c^4} \frac{1}{r} \frac{\partial^2 Q}{\partial t^2}$$

LaSR Results II - Hints

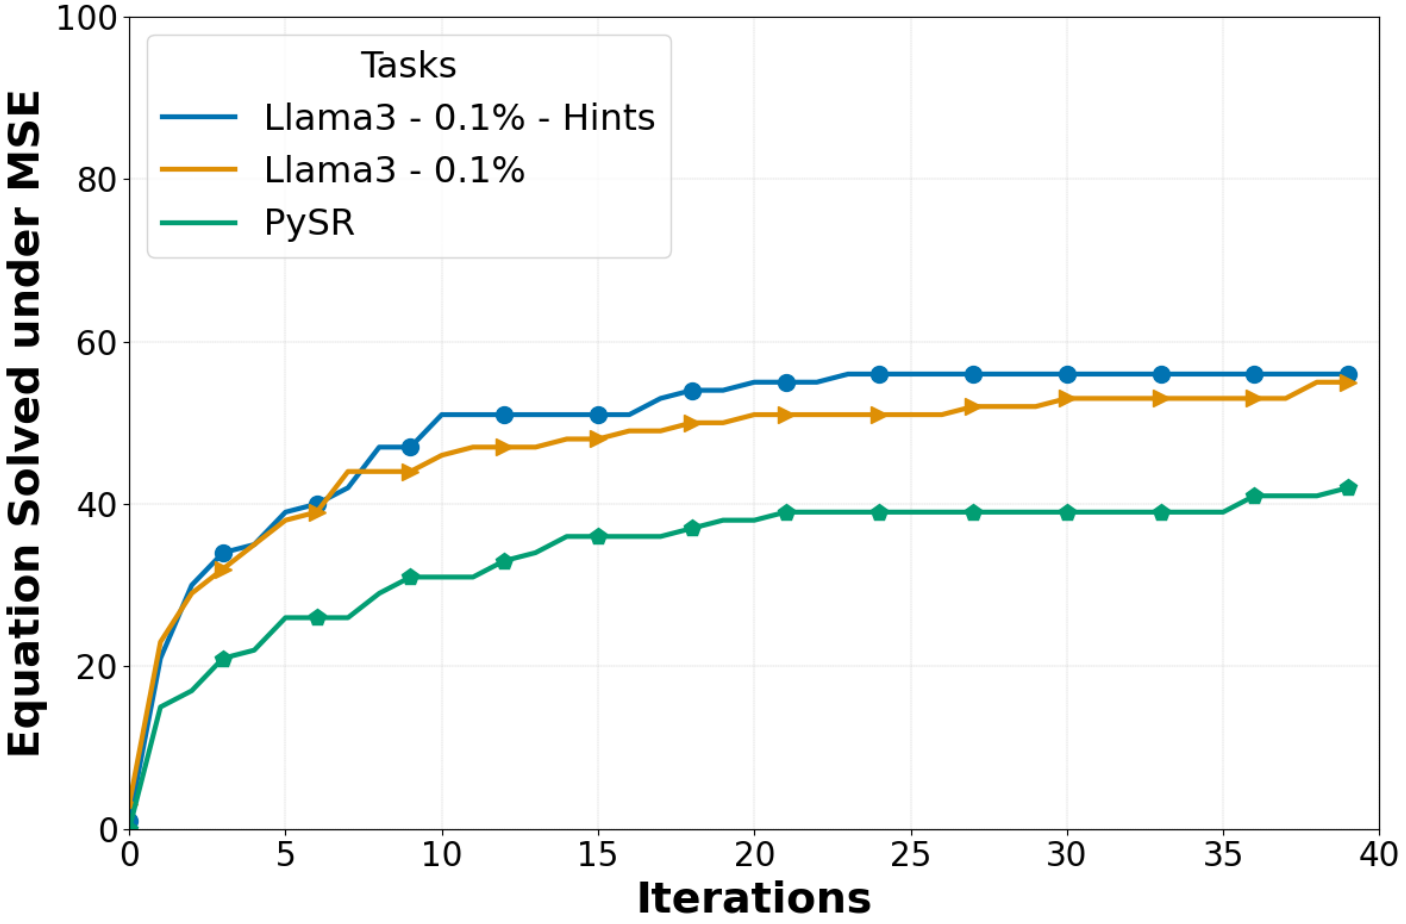

User provided hints accelerate hypothesis search

Results III - Case Study

$$F = \frac{1}{4\pi\epsilon}\frac{q_1 q_2}{r^2}$$

$$F = \frac{\left(\left(\left(\left(\left(\left(\left(\frac{q_2 \cdot 3.382}{r}\right) - \left(\frac{\sin\left(\frac{0.017}{\exp(B)}\right)}{\exp(C)}\right)\right)\right)/0.712\right) \cdot q_1\right) \cdot 0.087\right)/\epsilon\right) \cdot 0.191\right)}{r}$$

Eq 10: Coulomb's Law

- Inverse Square Law
- Directly proportional to charges
- Force symmetric w.r.t charges

PySR's Solution

- Reduces to ground truth after 10 steps of simplification.
- Unwieldly
- Fitting more constants => more optimization errors

Results III - Case Study

$$F = \frac{1}{4\pi\epsilon}\frac{q_1 q_2}{r^2}$$

Eq 10: Coulomb's Law

- Inverse Square Law
- Directly proportional to charges
- Force symmetric w.r.t charges

$$\begin{aligned}
F &= \frac{q_1}{\left(\frac{r}{q_2}\right)\left(r + \frac{1.9181636\times10^{-5}}{q_2}\right)\epsilon} \cdot 0.07957782 \\
&= \frac{q_1}{\left(\frac{r}{q_2}\right)\left(r + \frac{1.9181636\times10^{-5}}{q_2}\right)\epsilon} \cdot \frac{1}{4\pi} && \text{(Substitute constant)} \\
&= \frac{q_1 q_2}{r\left(r + \frac{1.9181636\times10^{-5}}{q_2}\right)\epsilon} \cdot \frac{1}{4\pi} && \text{(Simplify denominator)} \\
&\approx \frac{q_1 q_2}{r\left(r\right)\epsilon} \cdot \frac{1}{4\pi} && \text{(Negligible. } \tfrac{1.9181636\times10^{-5}}{q_2} \approx 0\text{)}
\end{aligned}$$

LaSR's Solution

- Reduces to ground truth after 4 steps of simplification
- Smaller models synthesize simpler equations!

Results III - Case Study

$$F = \frac{1}{4\pi\epsilon}\frac{q_1 q_2}{r^2}$$

Eq 10: Coulomb's Law

- Inverse Square Law
- Directly proportional to charges
- Force symmetric w.r.t charges

Iteration	Discovered Concept
2	*The good mathematical expressions exhibit [...] with a focus on* ***power functions and trigonometric functions*** *[...]*
6	*The good mathematical expressions exhibit [...]* ***symmetry or regularity*** *[...]*
24	*The good mathematical expressions have [...] with a specific pattern of* ***division and multiplication***

LaSR's Concepts (Limitations)

- Cannot guarantee factuality or correctness.
- Good concepts depend on LLM training. Concepts can mislead scientists.

Outline of this Talk

1. What is Scientific Discovery?
    1. Symbolic Regression
2. Symbolic Regression with a Concept Library
3. **Additional Application: Visual Reasoning**
4. Discussion

Recap: Symbolic Regression

World

Scientific Process

Domain Knowledge

Iterative

Choose Research Question → Generate Hypothesis → Experimental Setup / Data Collection → Analysis → Interpretation → Report to Community

Figure 1: Tycho Brahe's observations of the planet Mars (period 1582 – 1600).

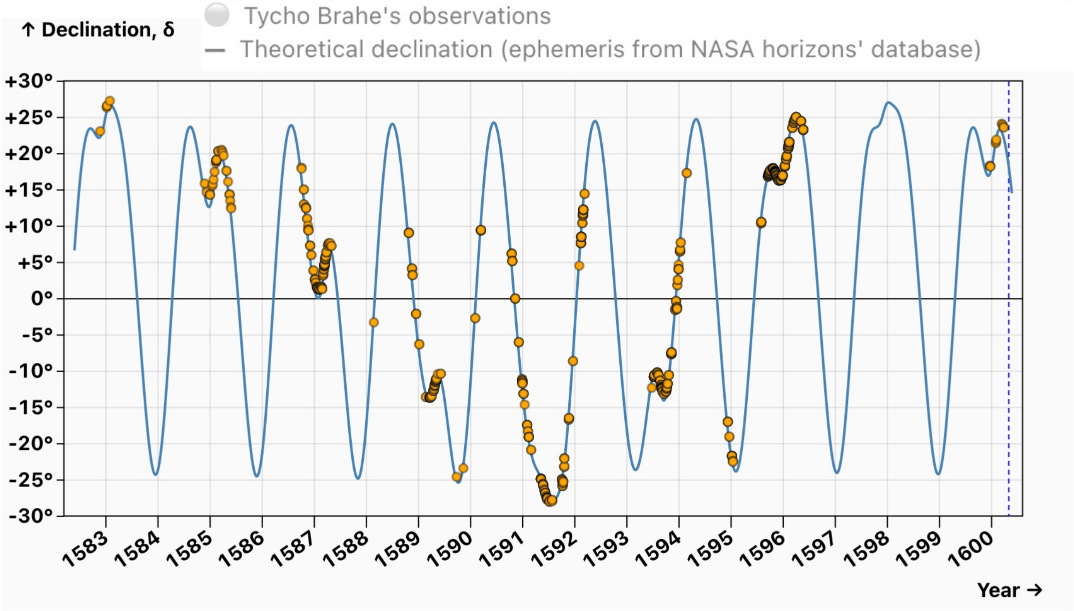

→

Kepler's Third Law

$$T^2 \propto r^3$$

https://www.inaturalist.org/observations/1970016

Visual Reasoning

World ⟷ **Scientific Process** ⟷ Domain Knowledge

Iterative

Choose Research Question → Generate Hypothesis → Experimental Setup / Data Collection → Analysis → Interpretation → Report to Community

Observation: Geolocated Picture

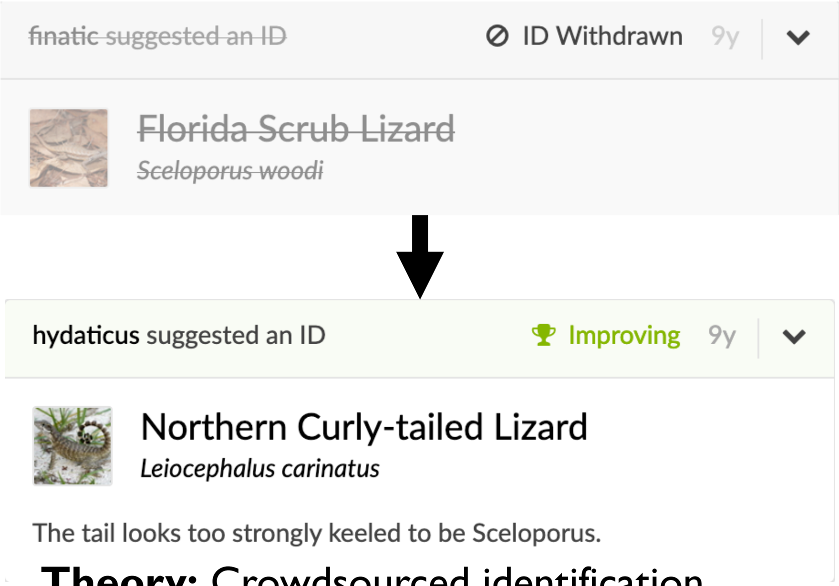

Theory: Crowdsourced identification

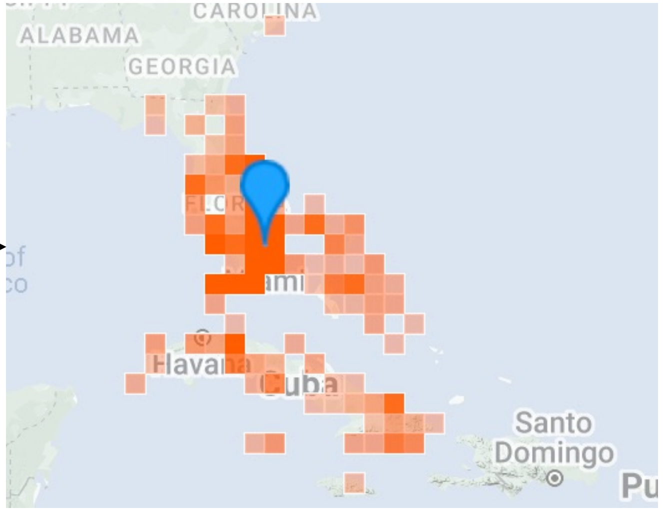

Data: Geotagged sightings

Visual Programming

"Is there a helmet in this picture that is not blue?"

Large monolithic model
GPT4v

"no"

"Is there a helmet in this picture that is not blue?"

Visual Program

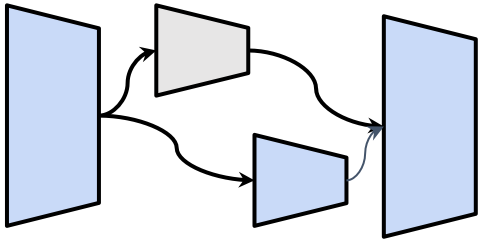

"no"

Compositional Question Answering

Is there a helmet in the photo that is not blue?

Prediction: no

	`IMAGE`
	`BOX0=Loc(image=IMAGE, object='helmet')`
	`IMAGE0=Crop(bbox=BOX0)`
blue	`ANSWER0=Vqa(image=IMAGE0, question='What color is the helmet?')`
no	`ANSWER1=Eval(expr="'yes' if {ANSWER0} != 'blue' else 'no'")` `=Eval(expr="'yes' if 'blue' != 'blue' else 'no'")`

c. VisProg. Gupta et. al

Outline of this Talk

1. What is Scientific Discovery?
   1. Symbolic Regression
2. Symbolic Regression with a Concept Library
3. Additional Application: Visual Reasoning
4. **Discussion**