Discovering Physical Concepts with Neural Networks

Raban Iten, Tony Metger, Henrik Wilming, Lìdia del Rio, and Renato Renner

Phys. Rev. Lett. **124** (2020)

Presented by: Ethan Mullen, Nicki Mullins, Debora Mroczek, Michael Mollenhauer, Matthew O'Brien

PHYS 596

November 2020

Overview: Short- and Long-Term Goals of the Research

Short-Term Goals

Model a neural network (NN) after the human physical reasoning process (called ***SciNet***)

Recover the physical theories describing systems from data on such systems collected and fed to SciNet

Long-Term Goal

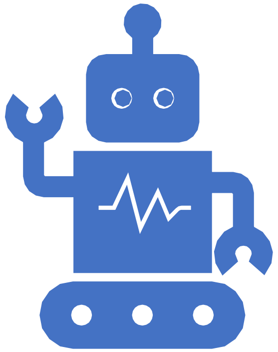

Machine Learning (ML)-assisted scientific discoveries from experimental and simulated data with <u>no assumptions</u>

SciNet's Input and Performance Summary

- Authors developed *SciNet* – a NN architecture that can recover physical variables from different toy models
- Four systems put the model to the test:
  - Damped pendulum
  - Two colliding particles
  - Simple quantum experiment of one- and two-qubit states
  - Planetary model (earth, moon, and sun)
- SciNet…
  - Proves robust against noise in the experimental data
  - Encodes the considered systems using **relevant** parameters (vs. **all** degrees of freedom)

Example	Observation input size	Question input size
Pendulum	50	1
Collision	30	16
One qubit	10	10
Two qubits	30	30
Solar system	2	0

SciNet Returns Physical Theories for Four Toy Models

- Predicts future positions of a **damped pendulum** with high accuracy

- Finds and exploits the conservation of total angular momentum to predict the motion of **two colliding particles**

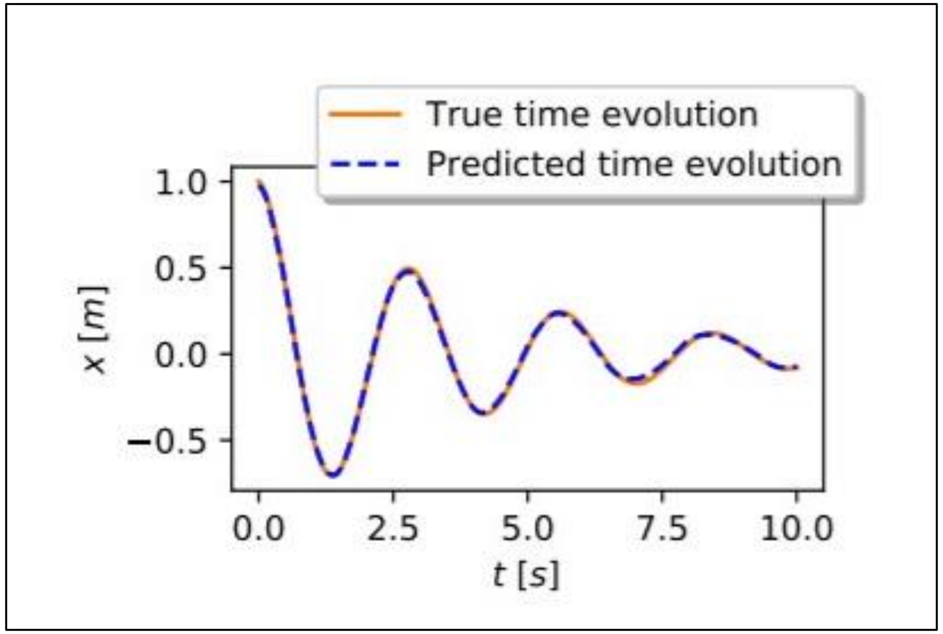

SciNet's predicted motion of the **damped pendulum**

SciNet Returns Physical Theories for Four Toy Models

- Determines the dimension of a quantum system and decides whether a set of measurements provides full information about the state given data from a **simple quantum experiment**

- Switches to the heliocentric view of a **model solar system** given a time series of the positions of the sun and moon from earth's perspective

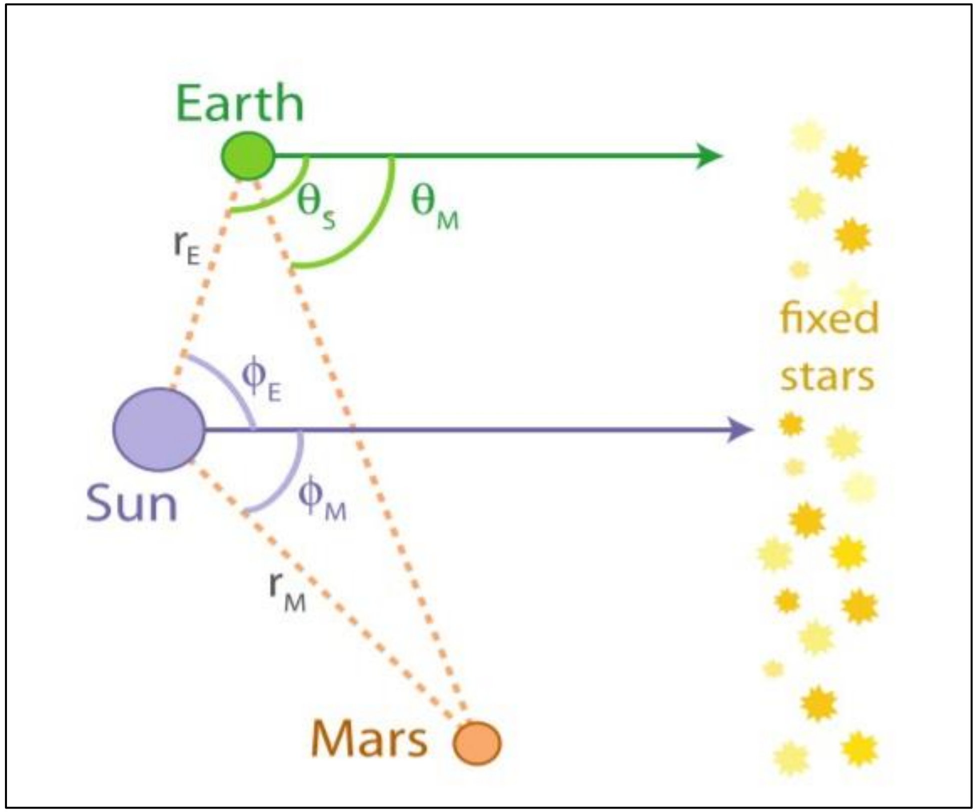

Simple **solar system** toy model

Machine Learning and Physical Models

- Finding mathematical expressions describing a dataset; extracting dynamical equations from experimental data
  - ***prior knowledge of system is required.***

M. Schmidt and H. Lipson, Science, 324 (2009)
B.C. Daniels and I. Nemenman, Nat. Comm. 6, (2015)
M. Raissi and G.E. Karniadakis. J. Comput. Phys. 357 (2018)

- Extracting physical variables from time series data of dynamical systems unsupervised.

D. Zheng et al (1808.10002)

Modelling the Human Physical Reasoning Process

Formalize human physical modelling process

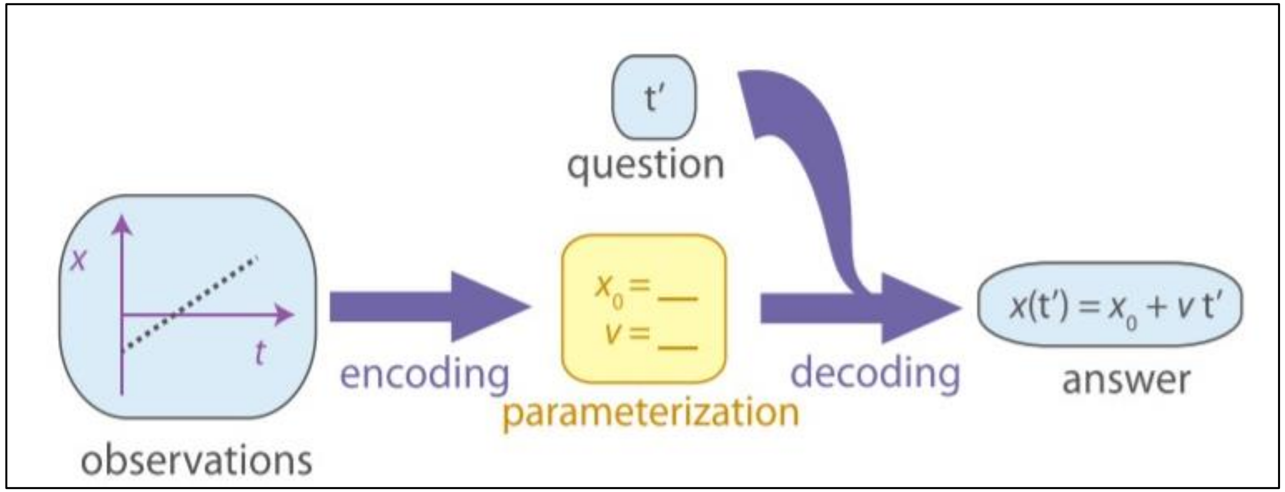

Translate into a neural network architecture

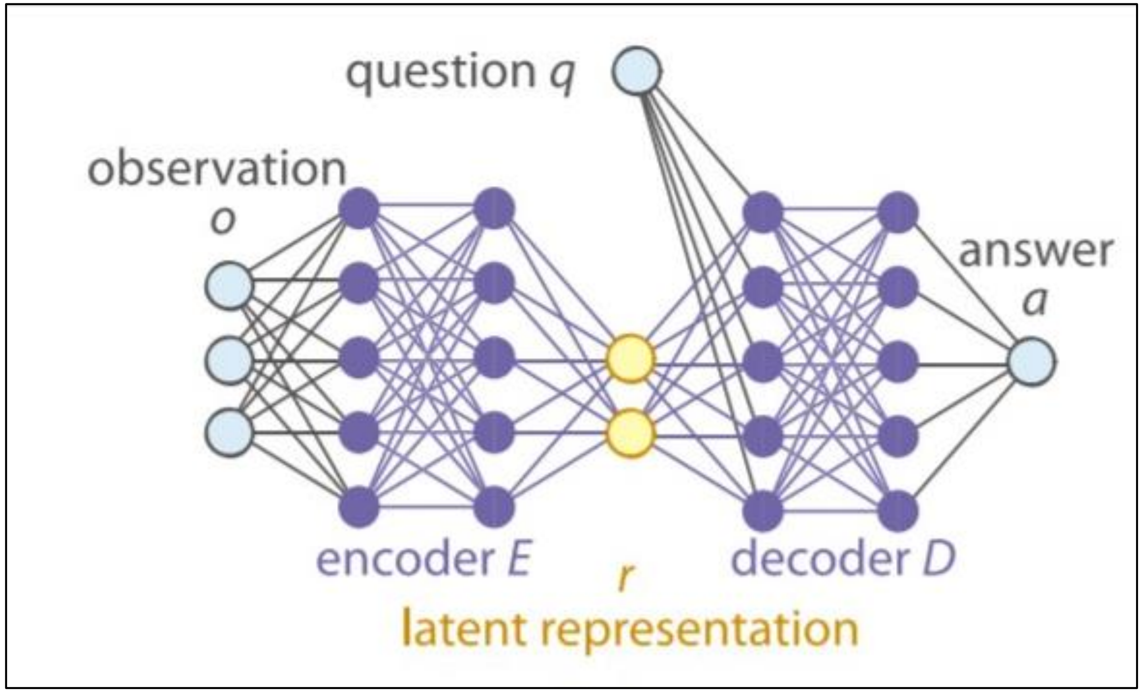

Mapping Between Parameters

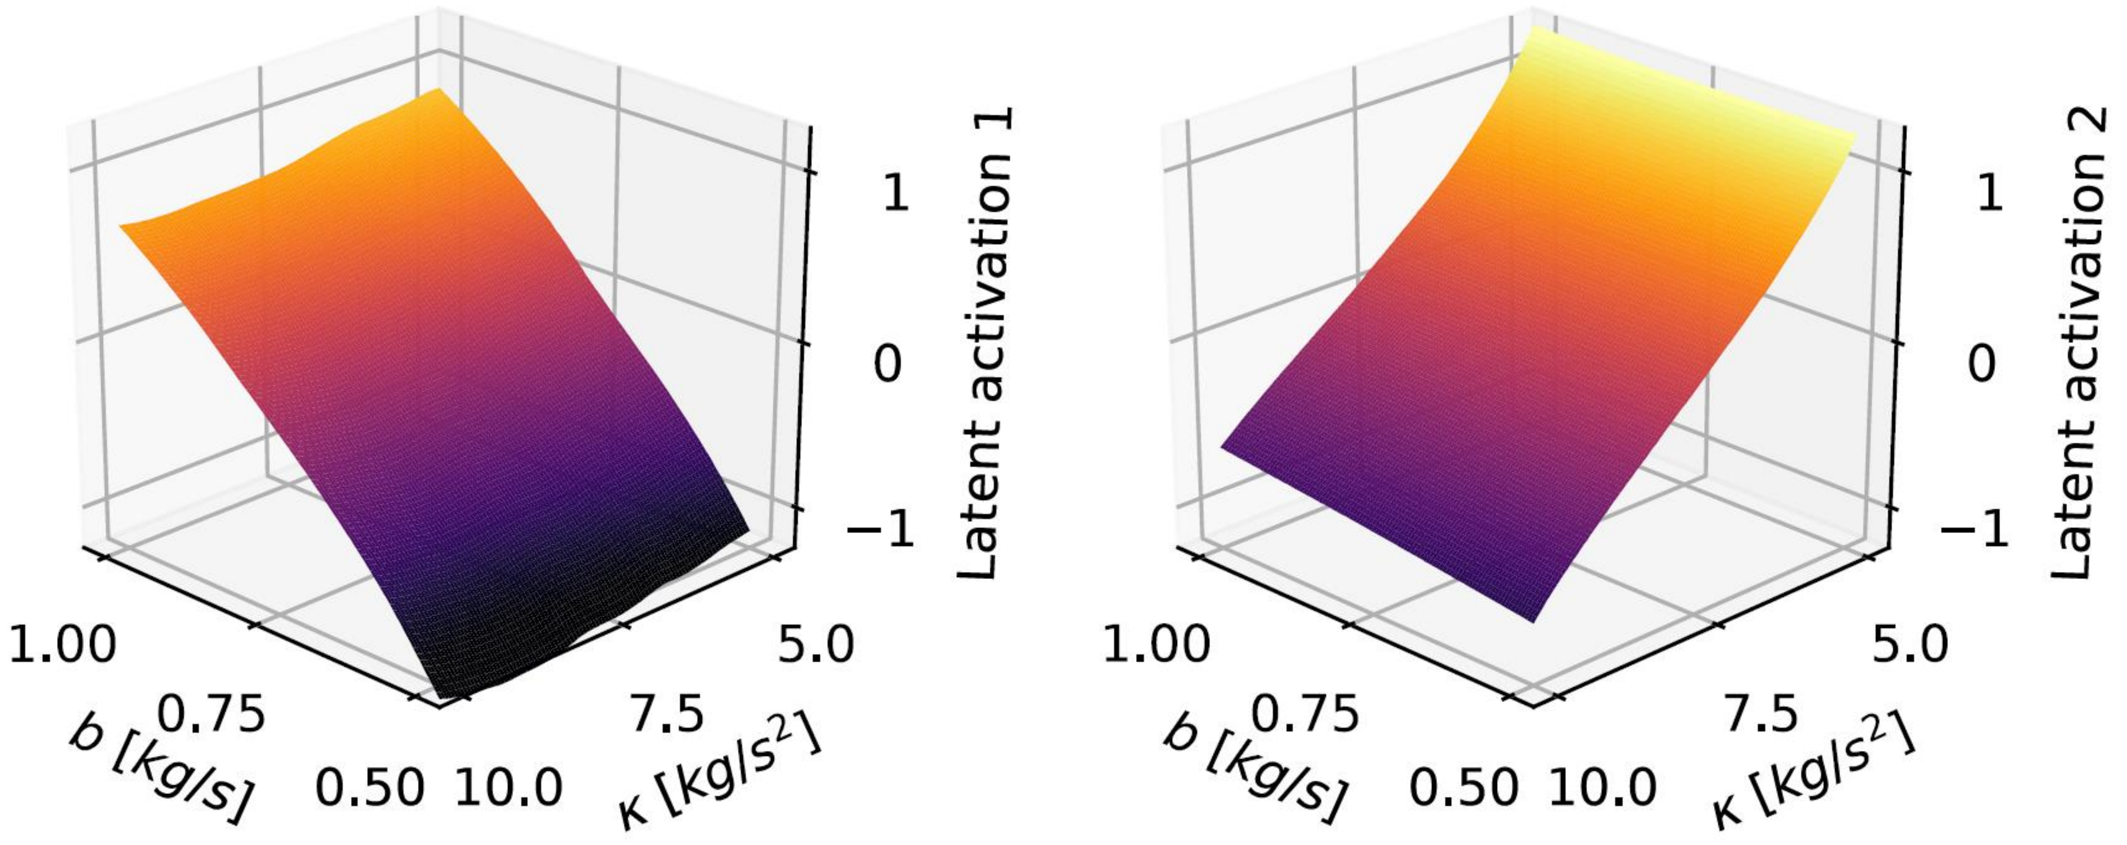

Time Evolution of Parameters

- Latent representation is extended to accommodate changing physical parameters:

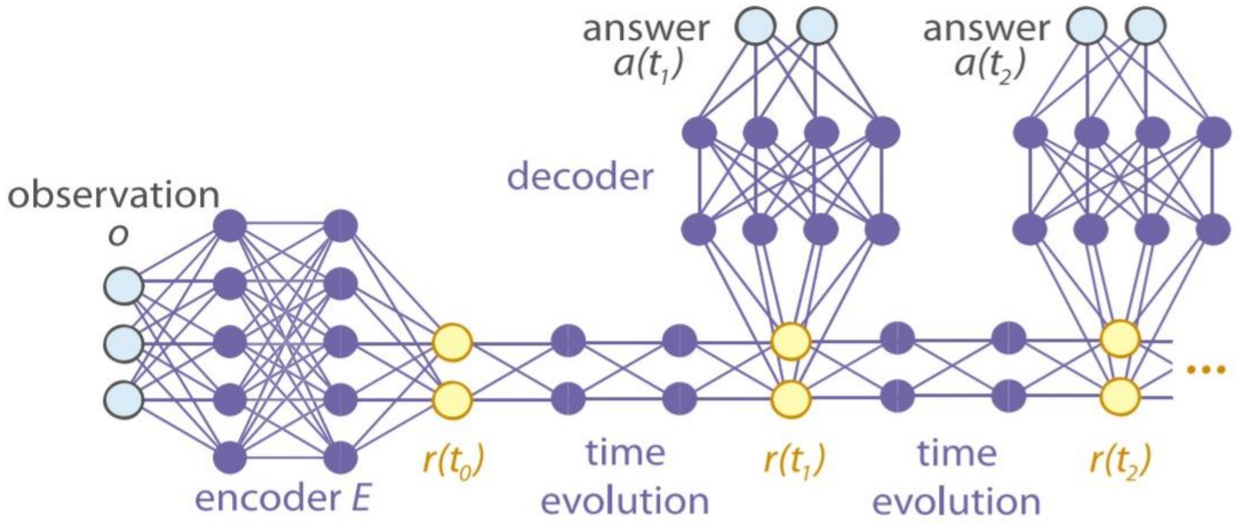

Why is SciNet Special?

- The encoder is free to choose which latent representation it learns from the training data
- Latent representation **does not need to encode all the physics** – only physics **relevant** to question being asked
- Minimal representation = independent parameters (underlying degrees of freedom in the system)
- New representations of the same prediction?

SciNet in a Nutshell: How Far Has it Come and How Far Can it Go?

- SciNet's representations of each model **match those found in textbooks** (post-selection of examples did *not* occur!)
- The authors plan to **extend the model** "to data where the natural underlying parameters are correlated in the training distribution"
- **Tradeoff between generality and performance**

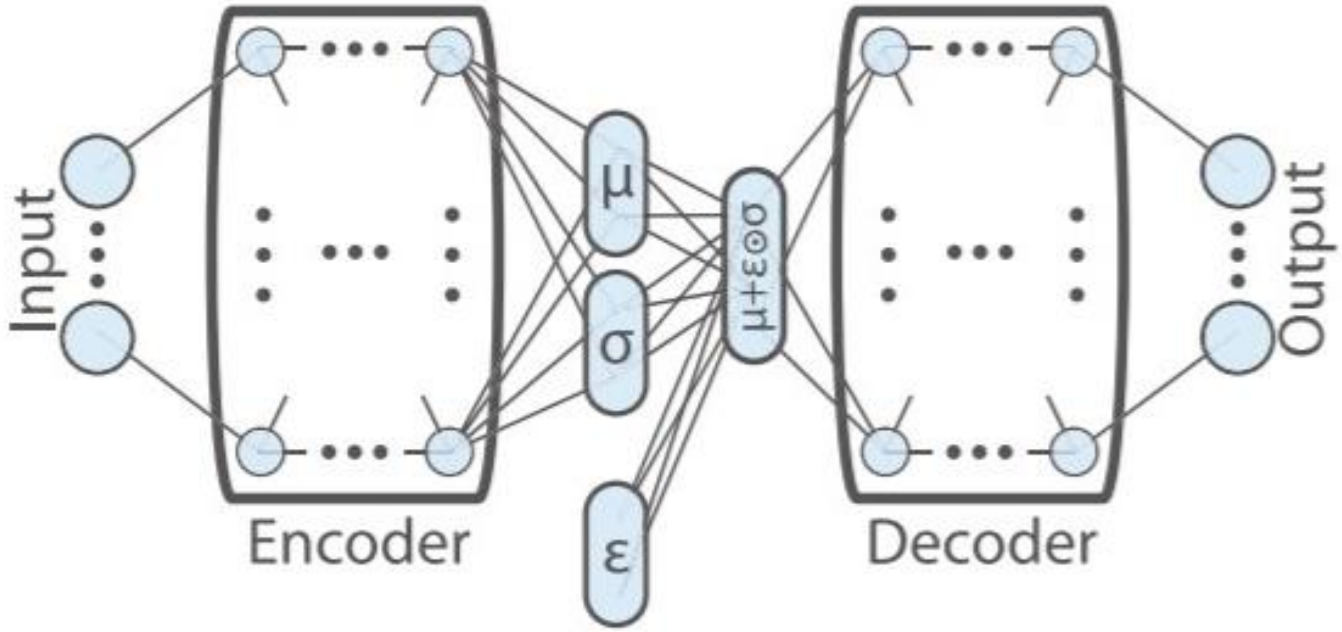

What worked

- Exceptionally broad interest
    - Wide range of potential applications
    - Addresses fundamental questions
    - SciNet code is open source

What worked

- Exceptionally broad interest
  - Wide range of potential applications
  - Addresses fundamental questions
  - SciNet code is open source

- Using representation learning overcomes common issues with ML
  - Drawing inferences is the primary goal!

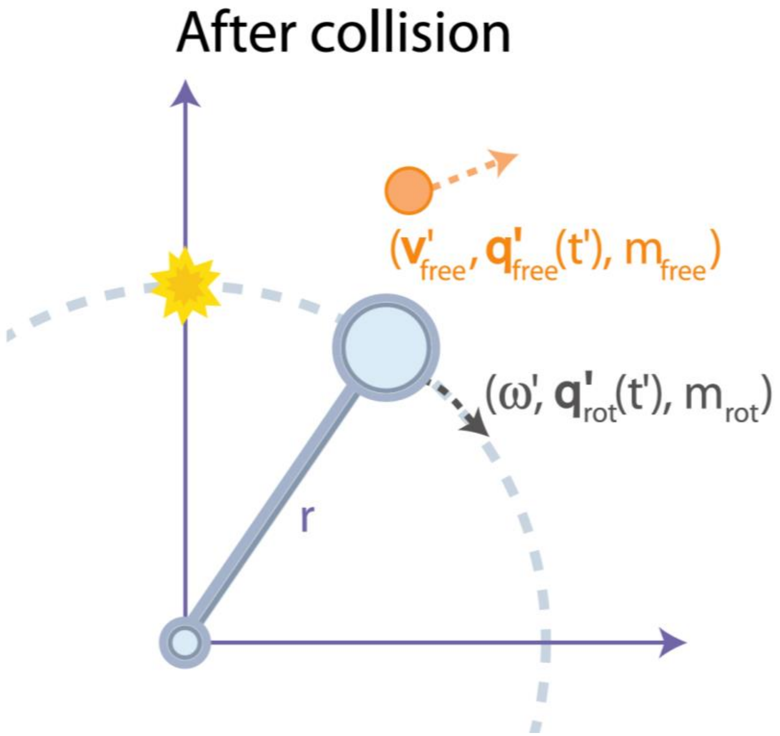

SciNet learns to analyze a collision by storing the total angular momentum in the latent representation

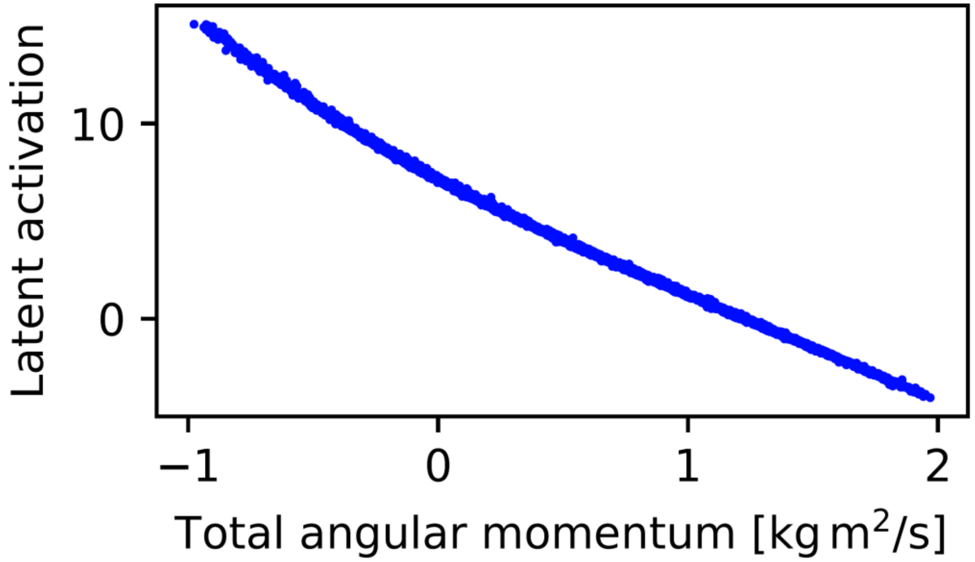

What worked

- Exceptionally broad interest
  - Wide range of potential applications
  - Addresses fundamental questions
  - SciNet code is open source

- Using representation learning overcomes common issues with ML
  - Drawing inferences is the primary goal!

- Convincing style of writing
  - Logical flow of ideas
  - Good "sales pitch" for usefulness of SciNet

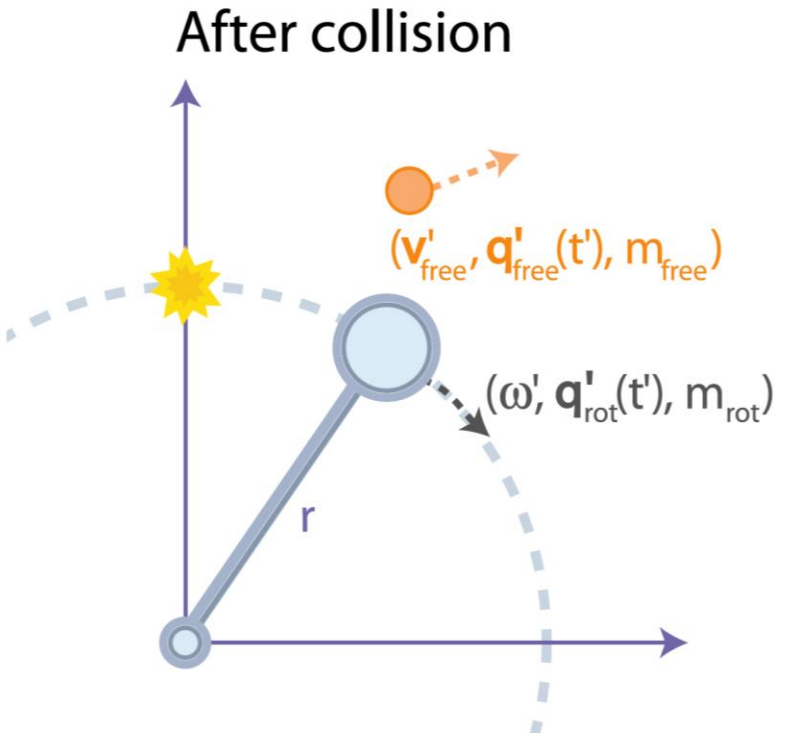

SciNet learns to analyze a collision by storing the total angular momentum in the latent representation

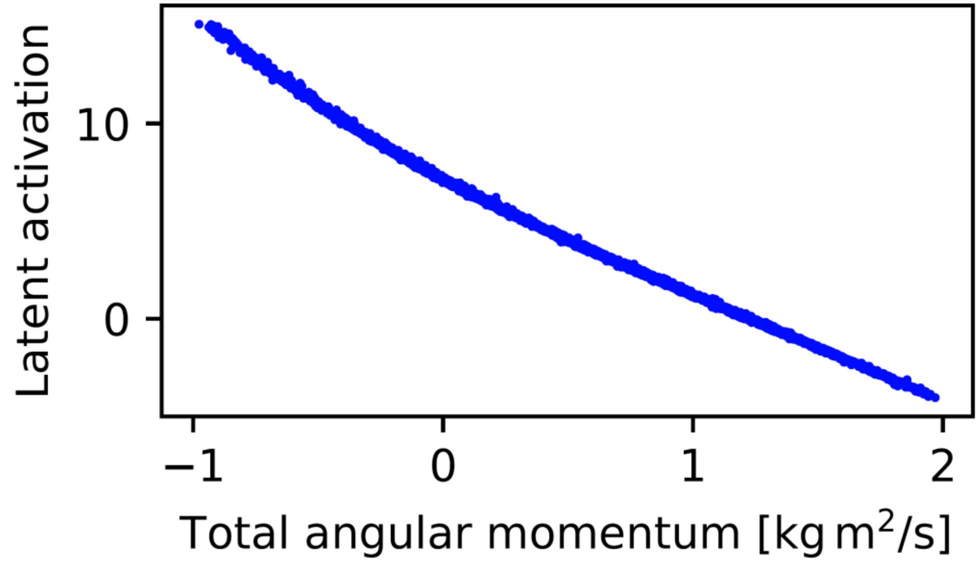

What didn't work

- Important details are not clearly explained:

What didn't work

- Important details are not clearly explained:
    - Neuron activation as a probe for “physical concepts”
    - Time evolution algorithm

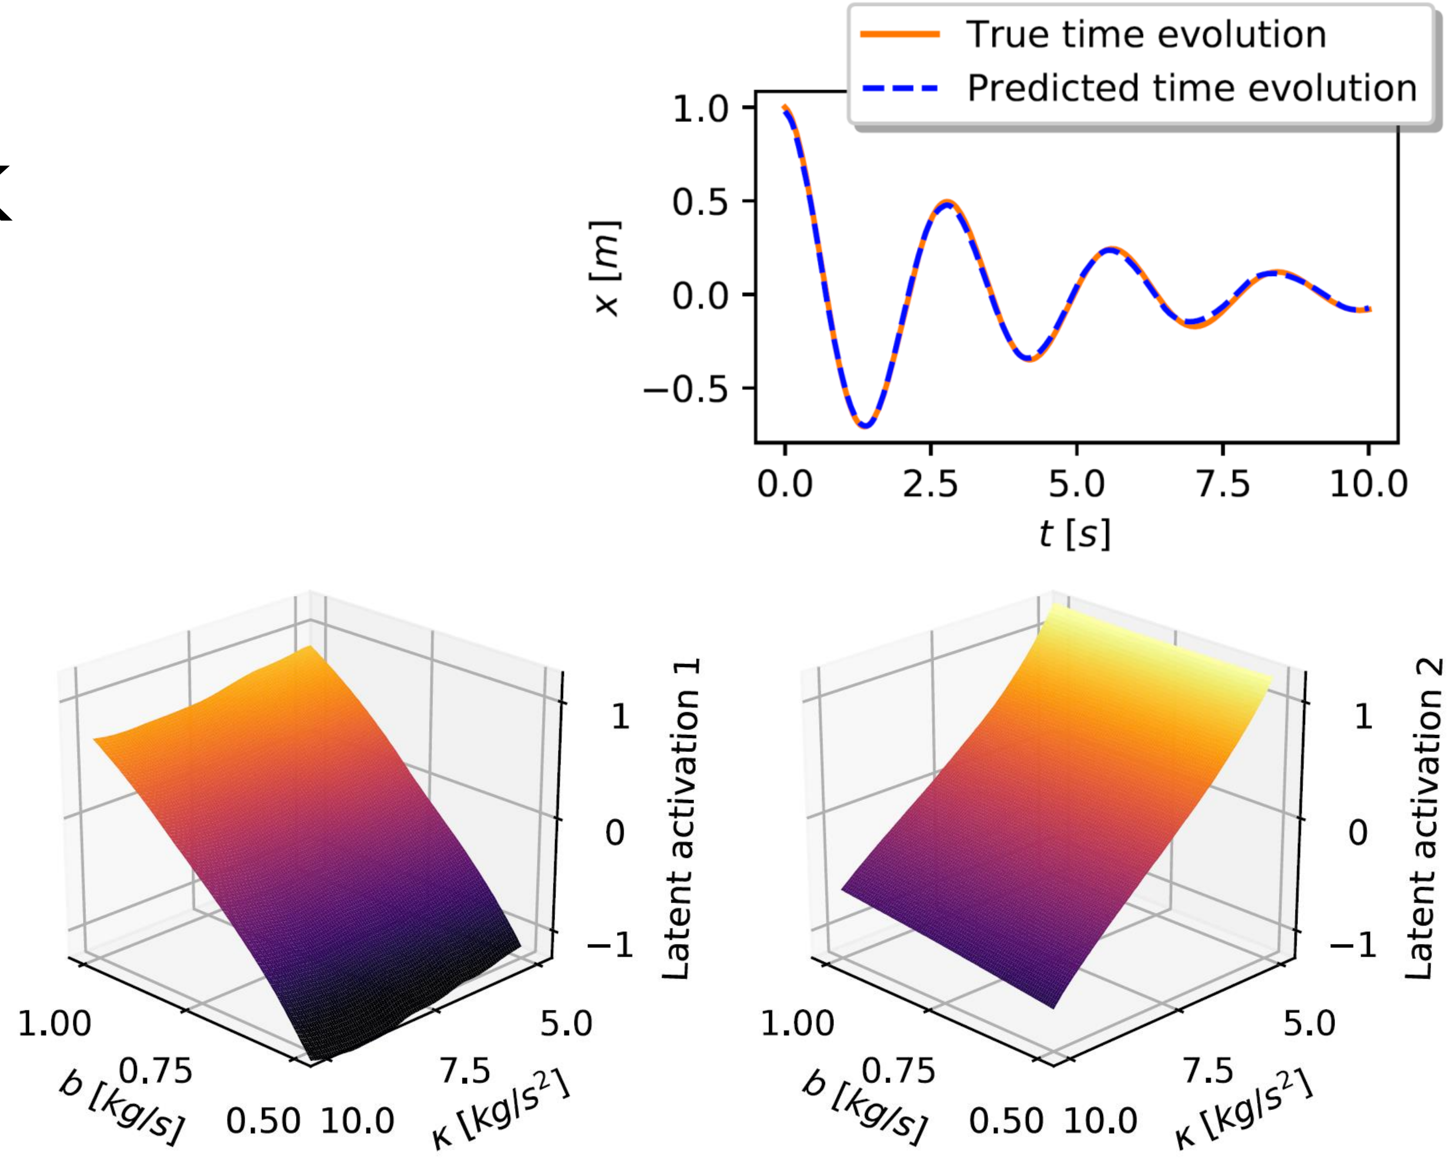

SciNet learns the parameters describing a simple harmonic oscillator

What didn't work

- Important details are not clearly explained:
  - Neuron activation as a probe for “physical concepts”
  - Time evolution algorithm
- Advances made seem over-stated/hyped
- Jargon!

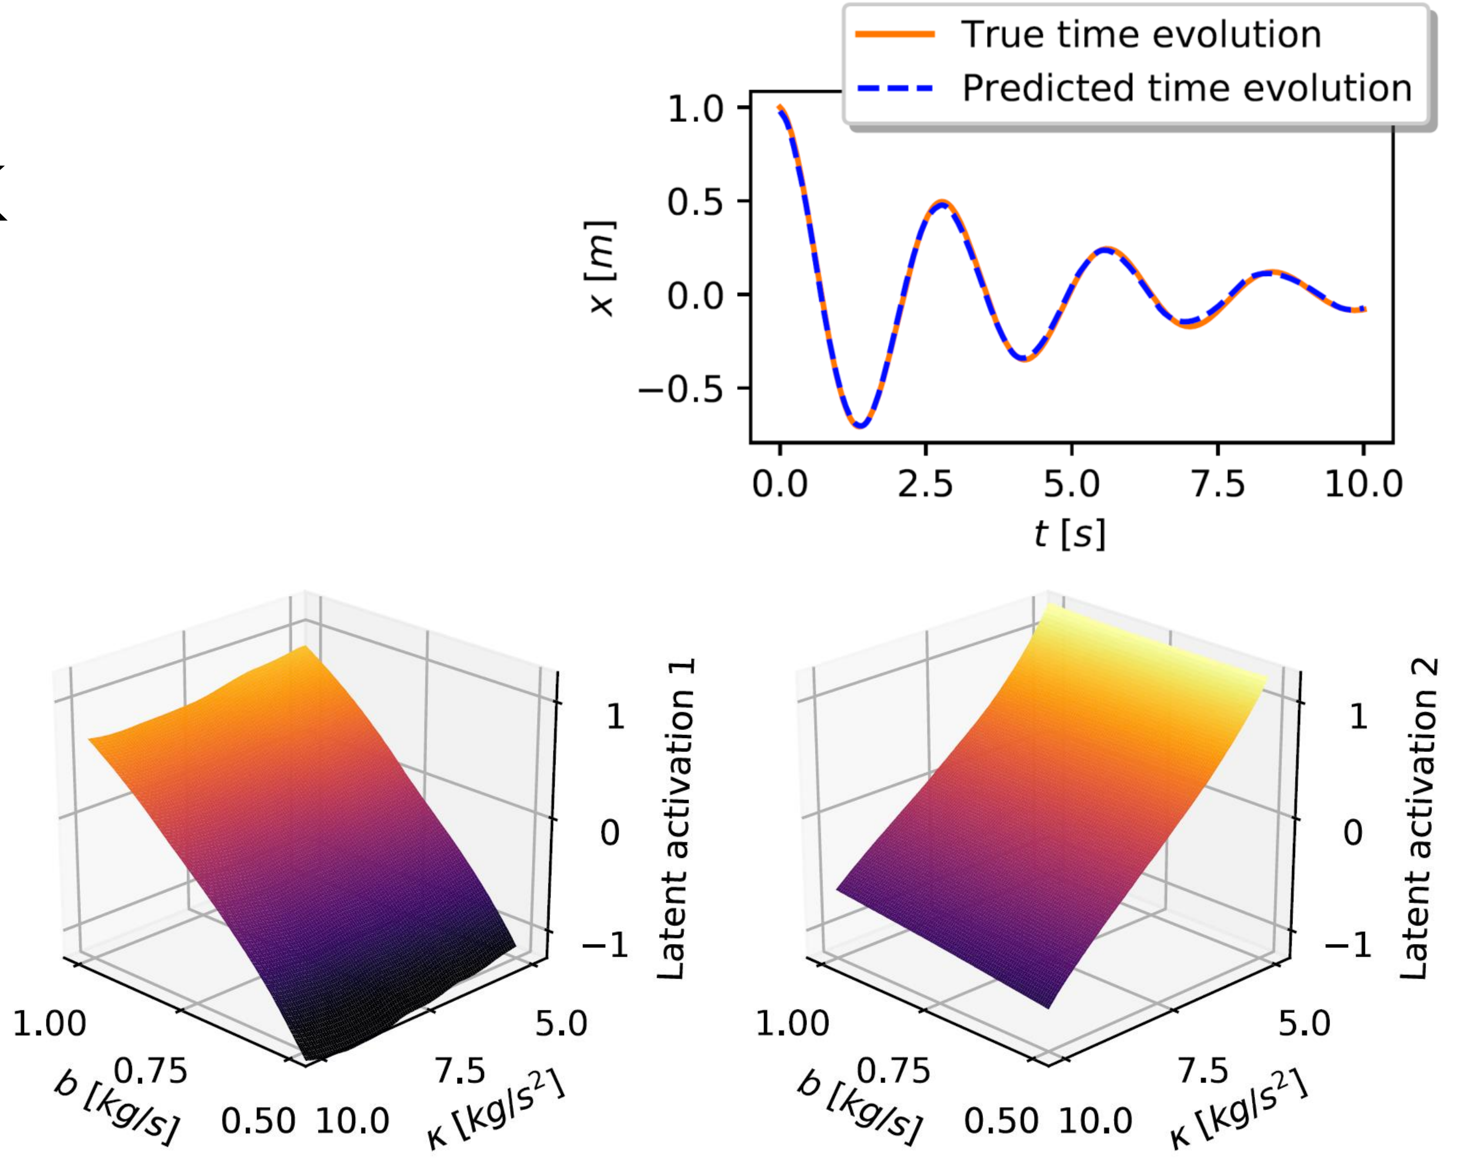

SciNet learns the parameters describing a simple harmonic oscillator

Overall Impressions

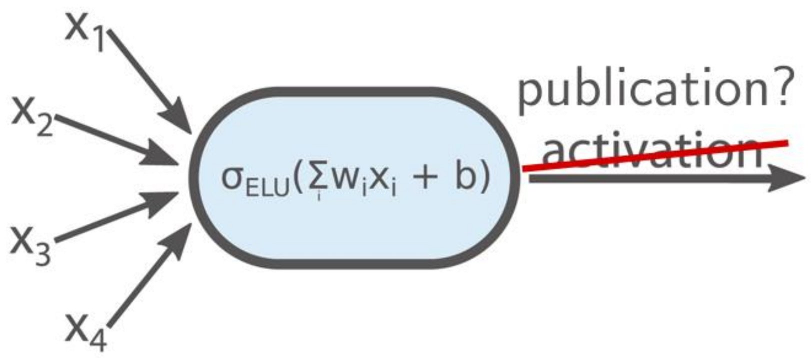

Positives

- Broad interest and impact for physics community
- Ideas are thoroughly studied and logically presented
- Potentially useful for understanding very complex systems

Negatives

- ML jargon limits accessibility for a physics audience
- Predictive power limited by ability to interpret representation
- Main text not self-contained: too much in SM

Conclusions

- Future work focused on including correlated parameters.
- In each example examined, the representation found by SciNet matches that used in standard physics textbooks.

Our Conclusions

- The model requires some **interpretation of the latent representation**, which can cause difficulties when examining systems for which we don't already know the parameterization.
- Our biggest issues are with presentation, **the actual science is good.**
- We're curious to see what happens when **applying this to more complicated systems.**

Number of Citations Since Publication

- Data taken from Scopus
- 28 total Citations since its publication on January 8th, 2020
- 24 citations in Scopus.

Information on these citations

- Mainly Physics, Computer Science, and Engineering papers citing this paper.
- Similar topics among these papers discuss neural networks, simulations, and numerical computation.

Distribution of Citations

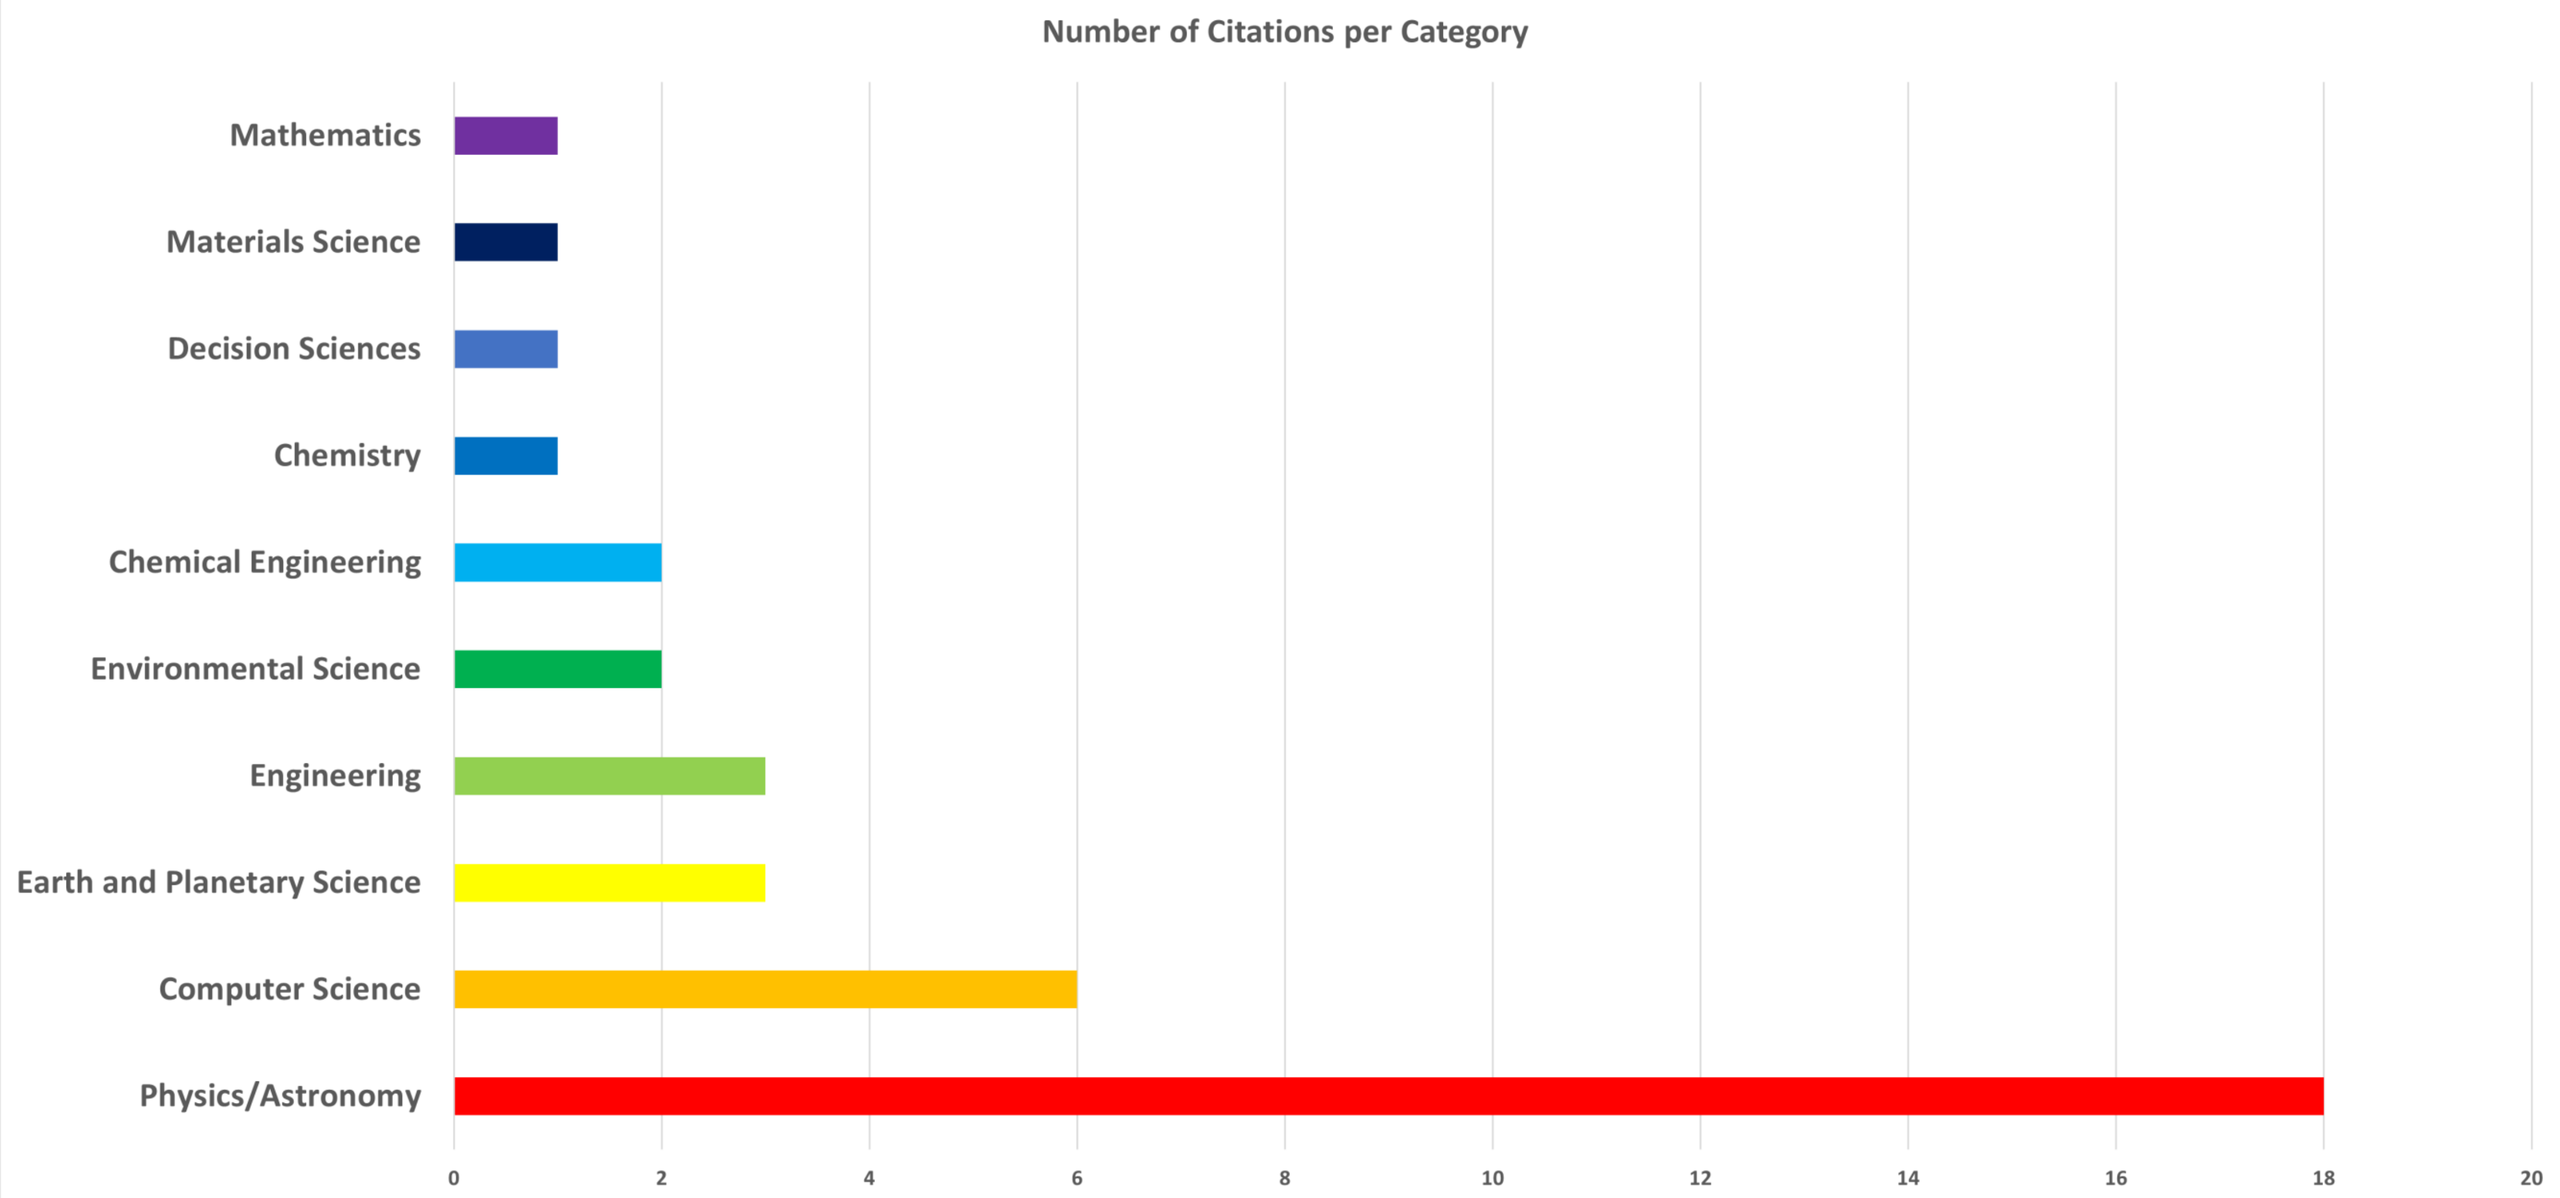

Citation and Research Impact

- Citation metrics from Scopus say it is being cited more than expected.
- Has a Field-Weighted Citation Impact of 27.29 degrees. This means the paper has been cited 27 more times than expected for a paper its age.

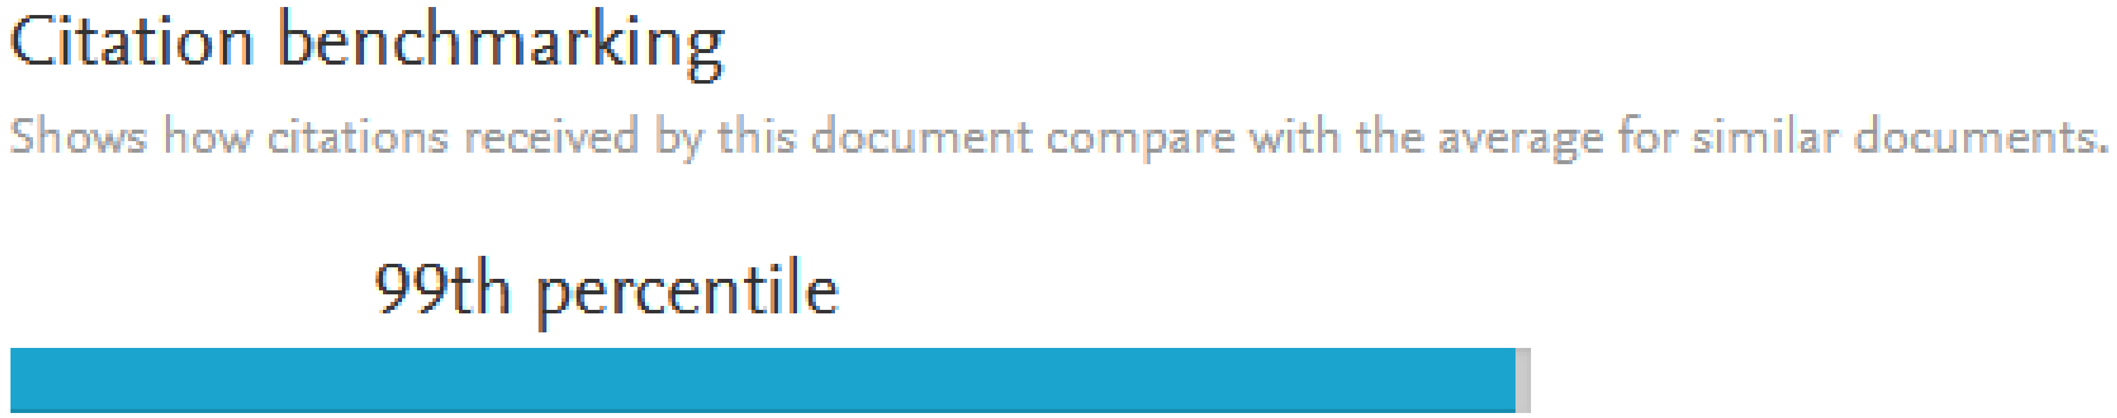

Community Impact

- Relatively new paper (less than a year), hard to tell if this is paper is having a major impact on the community now.
- However, with social media, more likely to bring this paper into light and be applied to more research topics.

Metrics Details

CITATIONS	24
Citation Indexes	24
Scopus	24
CAPTURES	625
Readers	625
Mendeley	625
MENTIONS	8
News Mentions	4
News	4
Q&A Site Mentions	3
Stack Exchange	3
Blog Mentions	1
Blog	1
SOCIAL MEDIA	269
Tweets	269
Twitter	269