Science is a verb:

Adopting the scientific method and best practices in AI research

Michela Paganini, Facebook AI Research
@WonderMicky

Physics ∩ ML

a virtual hub at the interface of theoretical physics and deep learning.

09/23/2020

FACEBOOK

Who am I

B.A. in Physics, B.A. in Astrophysics,
U.C. Berkeley

Ph.D. in Physics, Yale University

Thesis: *Machine Learning Solutions for High Energy Physics: Applications to Electromagnetic Shower Generation, Flavor Tagging, and the Search for di-Higgs Production* [arXiv:1903.05082]

Former Member,
ATLAS Collaboration, CERN

Postdoctoral Researcher,
Facebook AI Research

Visiting Affiliate,
NERSC

AI for Science

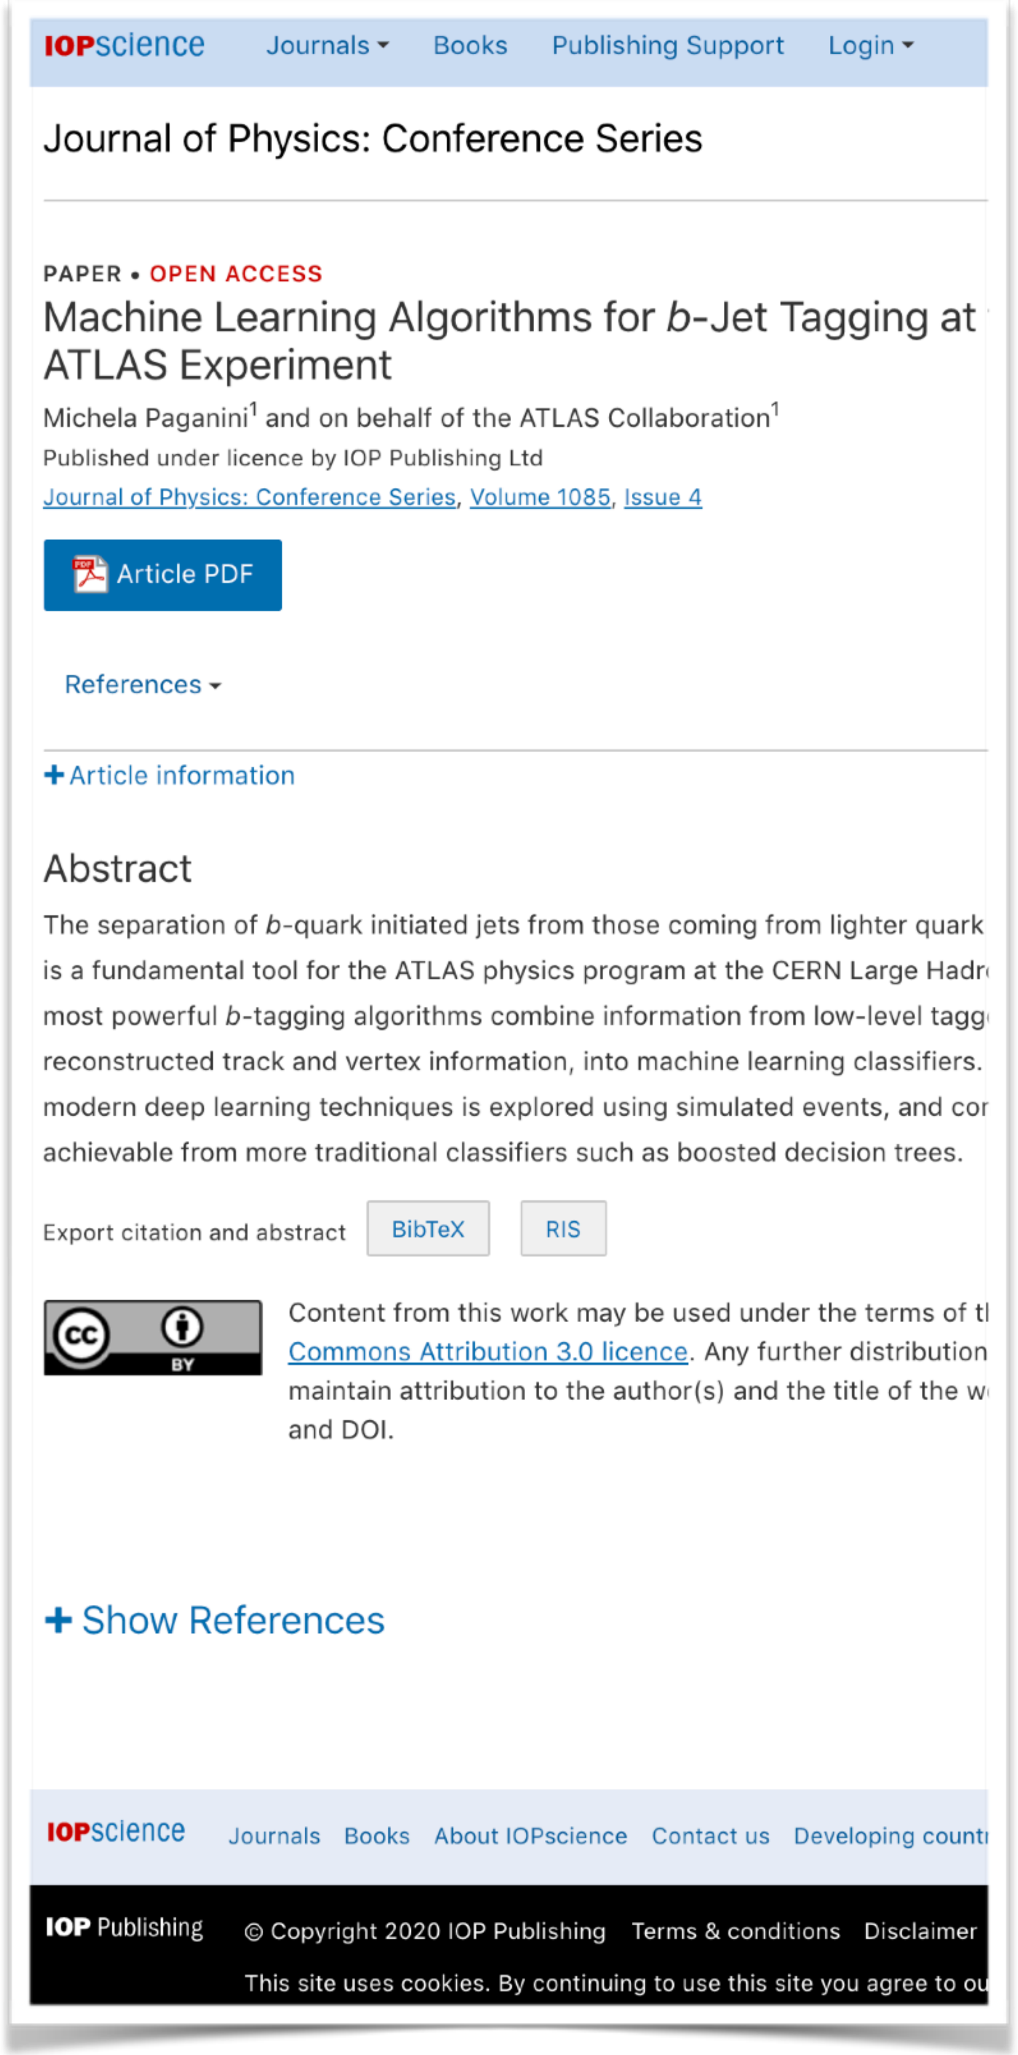
IOPscience Journals Books Publishing Support Login

Journal of Physics: Conference Series

PAPER • OPEN ACCESS

Machine Learning Algorithms for *b*-Jet Tagging at ATLAS Experiment

Michela Paganini[1] and on behalf of the ATLAS Collaboration[1]

Published under licence by IOP Publishing Ltd

Journal of Physics: Conference Series, Volume 1085, Issue 4

Article PDF

References

+ Article information

Abstract

The separation of *b*-quark initiated jets from those coming from lighter quark is a fundamental tool for the ATLAS physics program at the CERN Large Hadr most powerful *b*-tagging algorithms combine information from low-level tagg reconstructed track and vertex information, into machine learning classifiers. modern deep learning techniques is explored using simulated events, and cor achievable from more traditional classifiers such as boosted decision trees.

Export citation and abstract BibTeX RIS

Content from this work may be used under the terms of tl Commons Attribution 3.0 licence. Any further distribution maintain attribution to the author(s) and the title of the w and DOI.

+ Show References

IOPscience Journals Books About IOPscience Contact us Developing count

IOP Publishing © Copyright 2020 IOP Publishing Terms & conditions Disclaimer

This site uses cookies. By continuing to use this site you agree to ou

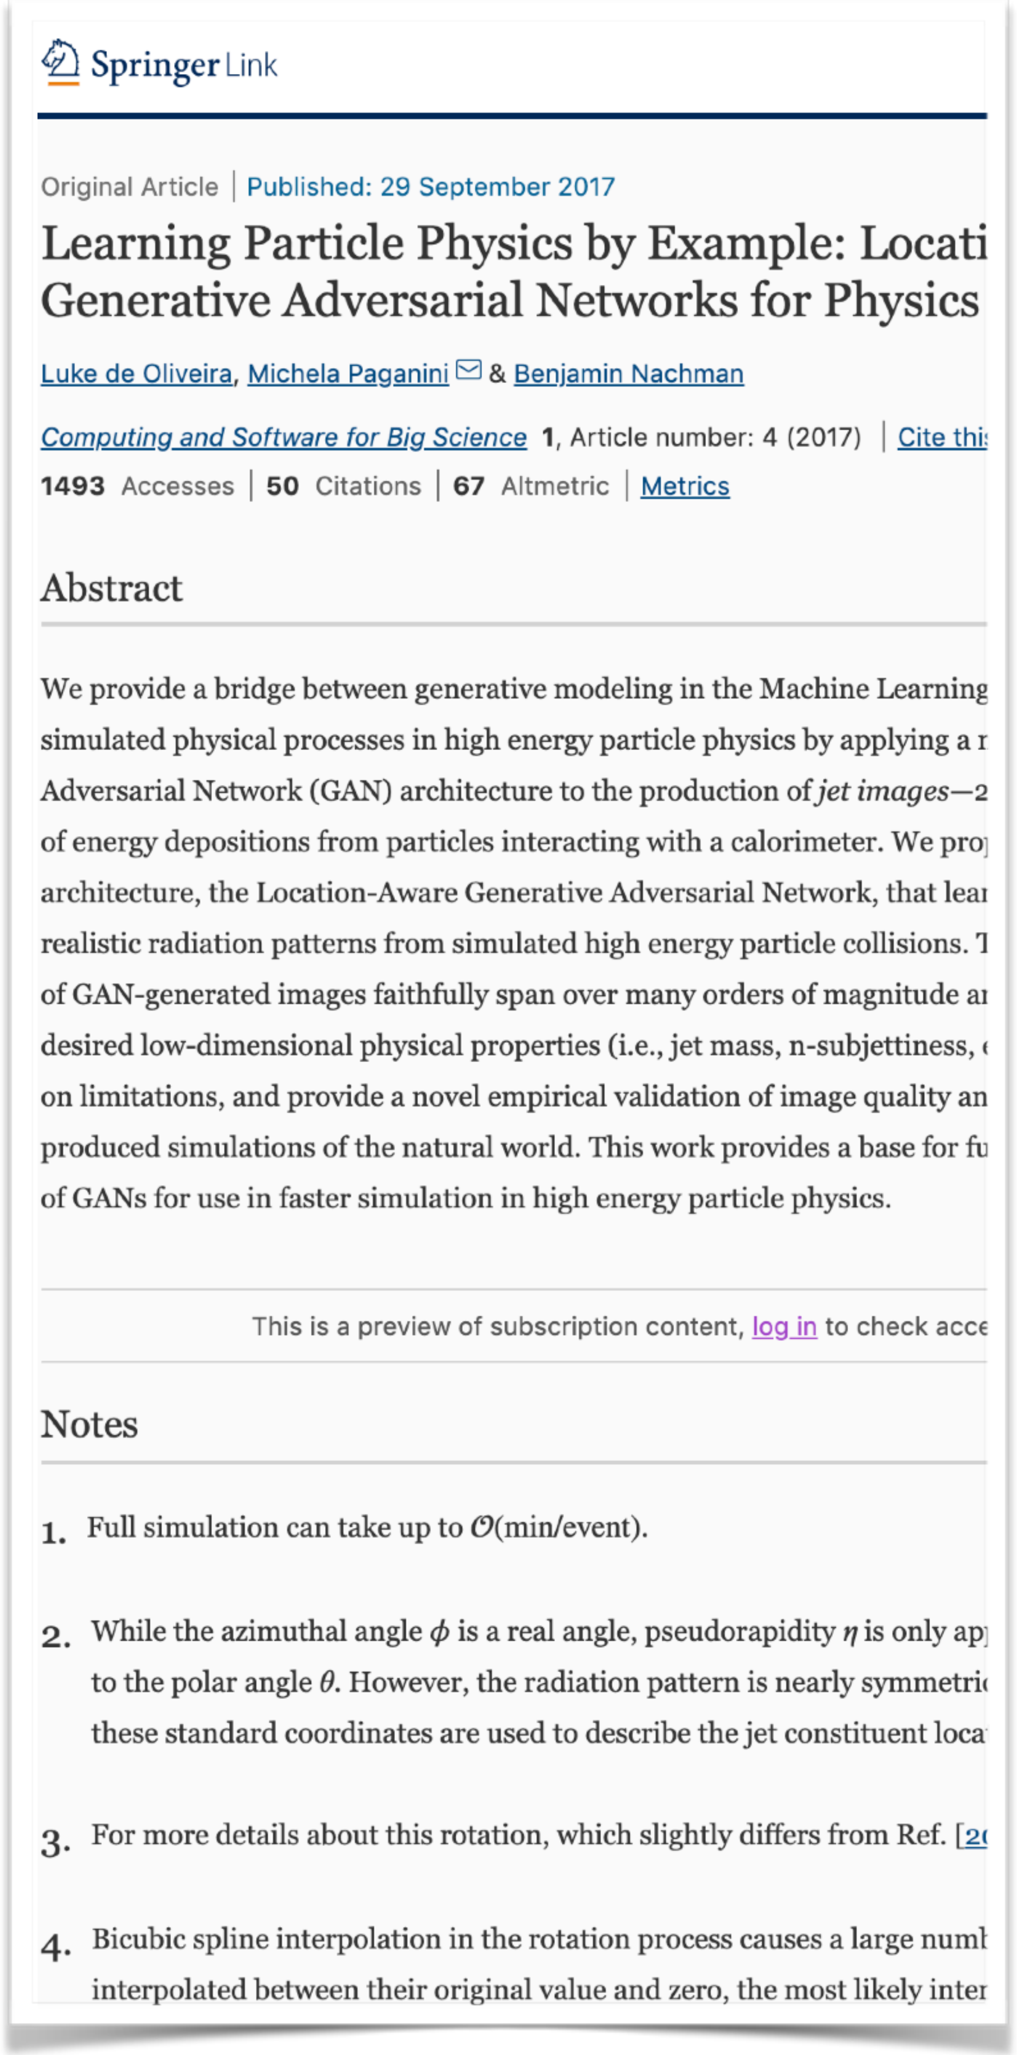
SpringerLink

Original Article | Published: 29 September 2017

Learning Particle Physics by Example: Locati Generative Adversarial Networks for Physics

Luke de Oliveira, Michela Paganini & Benjamin Nachman

Computing and Software for Big Science 1, Article number: 4 (2017) | Cite thi

1493 Accesses | 50 Citations | 67 Altmetric | Metrics

Abstract

We provide a bridge between generative modeling in the Machine Learning simulated physical processes in high energy particle physics by applying a r Adversarial Network (GAN) architecture to the production of *jet images*—2 of energy depositions from particles interacting with a calorimeter. We pro architecture, the Location-Aware Generative Adversarial Network, that lea realistic radiation patterns from simulated high energy particle collisions. T of GAN-generated images faithfully span over many orders of magnitude a desired low-dimensional physical properties (i.e., jet mass, n-subjettiness, on limitations, and provide a novel empirical validation of image quality an produced simulations of the natural world. This work provides a base for fu of GANs for use in faster simulation in high energy particle physics.

This is a preview of subscription content, log in to check acce

Notes

1. Full simulation can take up to $\mathcal{O}$(min/event).
2. While the azimuthal angle ϕ is a real angle, pseudorapidity η is only ap to the polar angle θ. However, the radiation pattern is nearly symmetri these standard coordinates are used to describe the jet constituent loca
3. For more details about this rotation, which slightly differs from Ref. [2
4. Bicubic spline interpolation in the rotation process causes a large numb interpolated between their original value and zero, the most likely inter

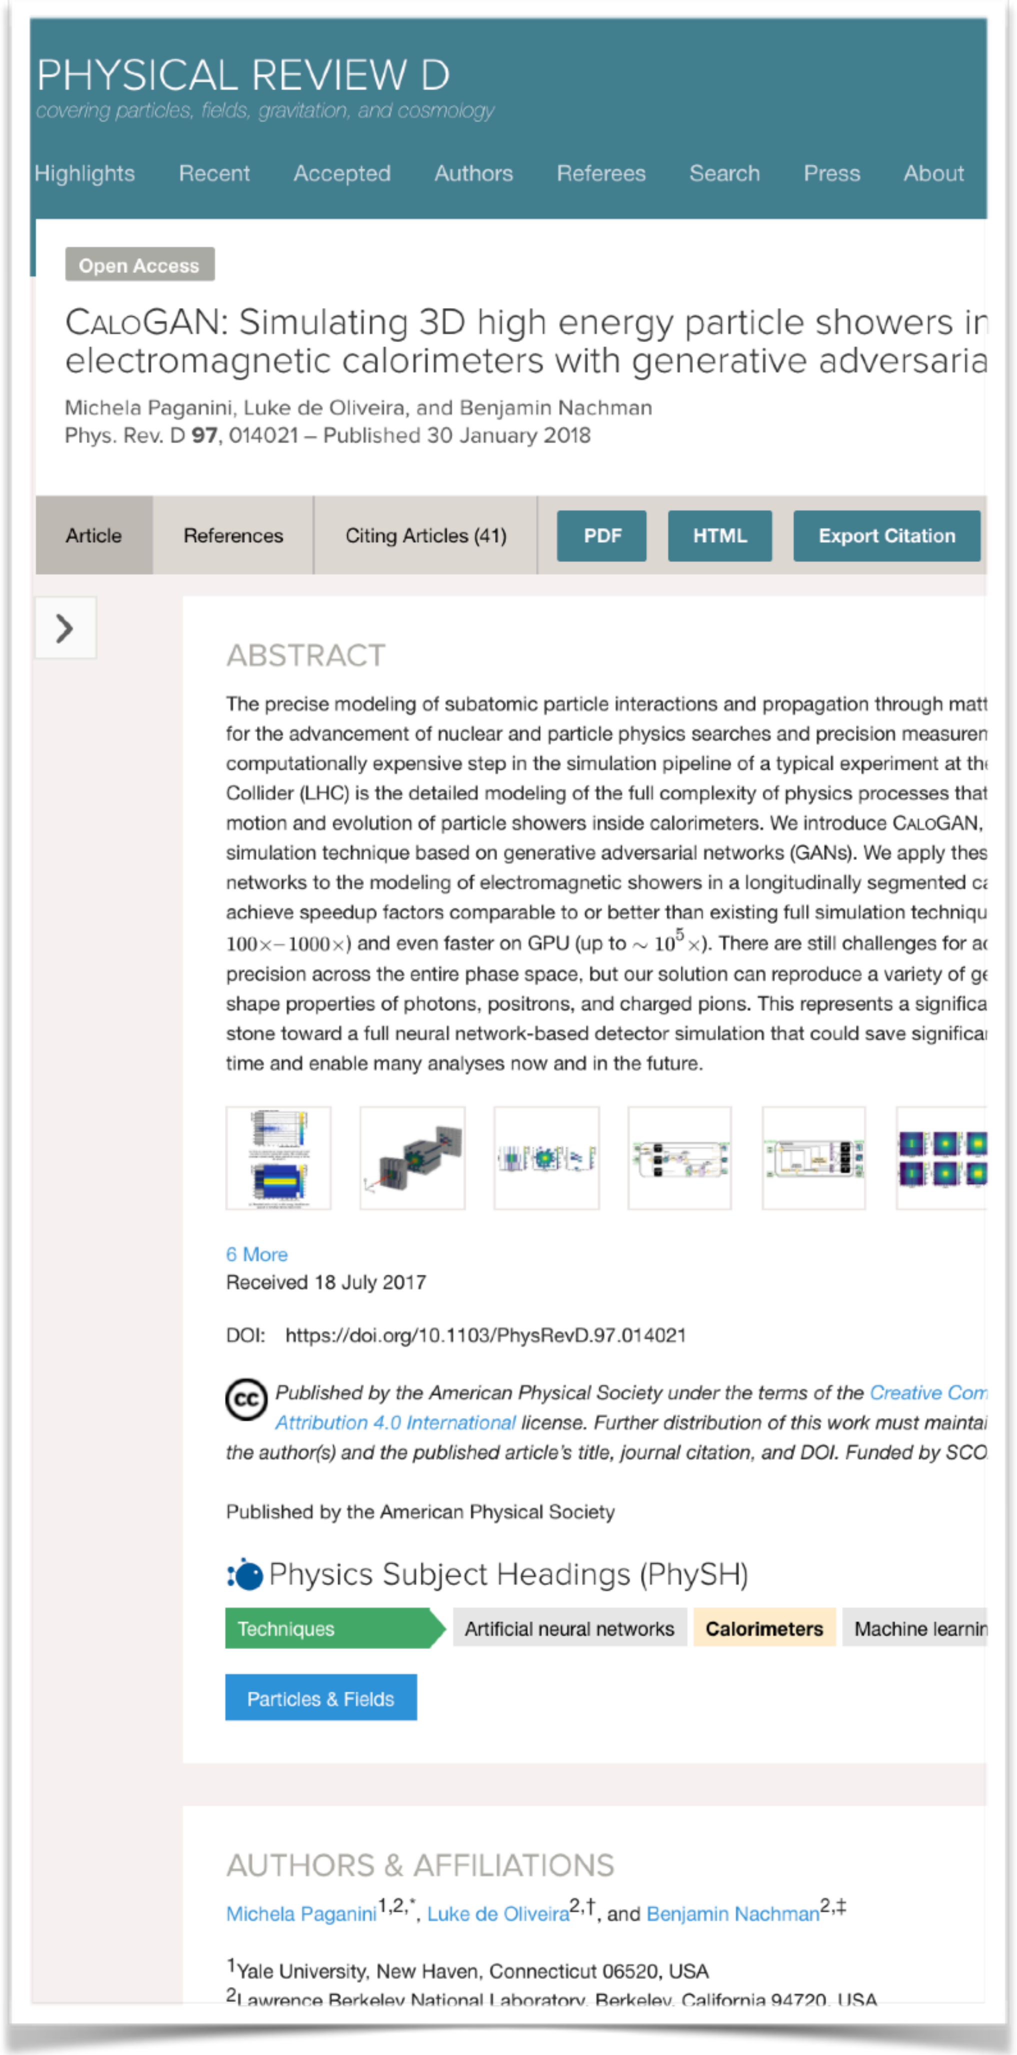
PHYSICAL REVIEW D

covering particles, fields, gravitation, and cosmology

Highlights Recent Accepted Authors Referees Search Press About

Open Access

CaloGAN: Simulating 3D high energy particle showers ir electromagnetic calorimeters with generative adversaria

Michela Paganini, Luke de Oliveira, and Benjamin Nachman

Phys. Rev. D 97, 014021 – Published 30 January 2018

Article References Citing Articles (41) PDF HTML Export Citation

ABSTRACT

The precise modeling of subatomic particle interactions and propagation through matt for the advancement of nuclear and particle physics searches and precision measuren computationally expensive step in the simulation pipeline of a typical experiment at th Collider (LHC) is the detailed modeling of the full complexity of physics processes that motion and evolution of particle showers inside calorimeters. We introduce CaloGAN, simulation technique based on generative adversarial networks (GANs). We apply thes networks to the modeling of electromagnetic showers in a longitudinally segmented ca achieve speedup factors comparable to or better than existing full simulation techniqu $100\times$–$1000\times$) and even faster on GPU (up to $\sim 10^5\times$). There are still challenges for ac precision across the entire phase space, but our solution can reproduce a variety of ge shape properties of photons, positrons, and charged pions. This represents a significa stone toward a full neural network-based detector simulation that could save significa time and enable many analyses now and in the future.

6 More

Received 18 July 2017

DOI: https://doi.org/10.1103/PhysRevD.97.014021

Published by the American Physical Society under the terms of the Creative Com Attribution 4.0 International license. Further distribution of this work must maintai the author(s) and the published article's title, journal citation, and DOI. Funded by SCO

Published by the American Physical Society

Physics Subject Headings (PhySH)

Techniques Artificial neural networks Calorimeters Machine learnin

Particles & Fields

AUTHORS & AFFILIATIONS

Michela Paganini[1,2,*], Luke de Oliveira[2,†], and Benjamin Nachman[2,‡]

[1]Yale University, New Haven, Connecticut 06520, USA

[2]Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA

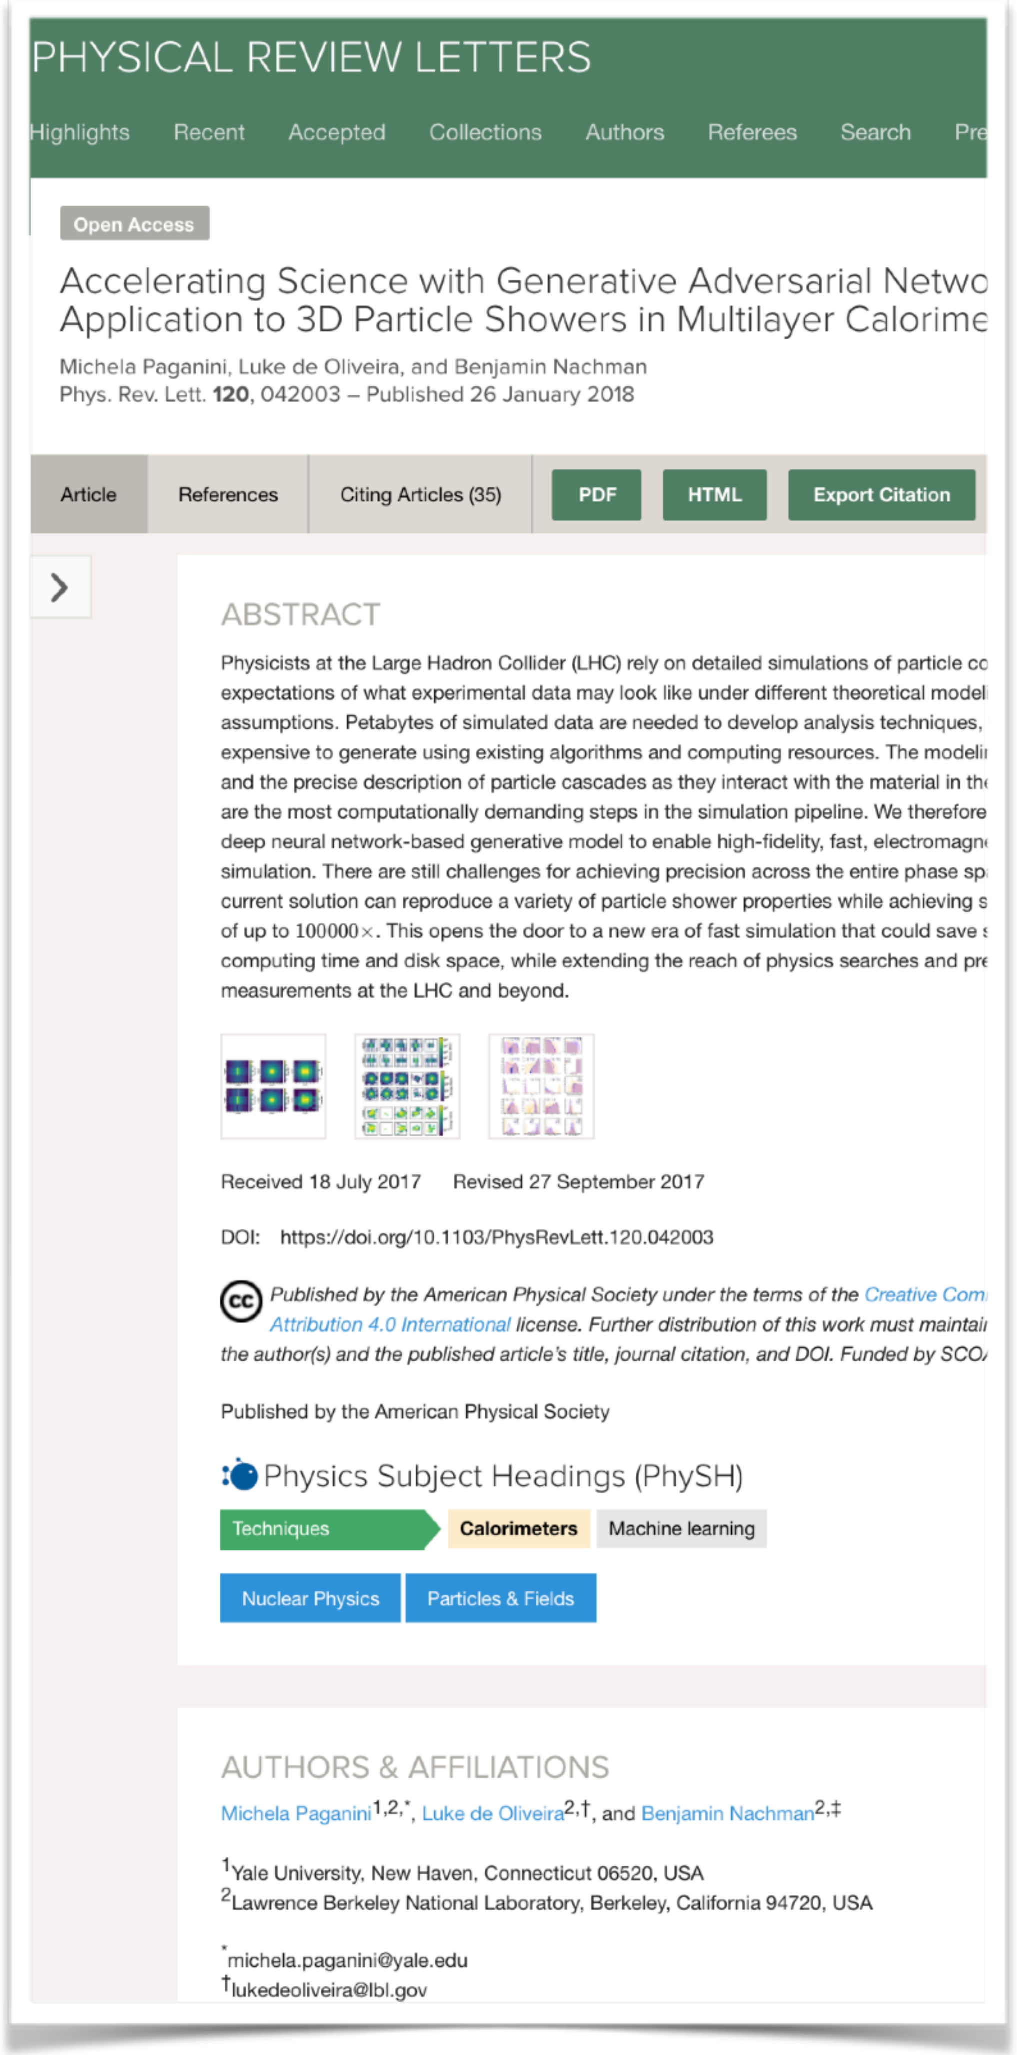
PHYSICAL REVIEW LETTERS

Highlights Recent Accepted Collections Authors Referees Search Pre

Open Access

Accelerating Science with Generative Adversarial Netwo Application to 3D Particle Showers in Multilayer Calorime

Michela Paganini, Luke de Oliveira, and Benjamin Nachman

Phys. Rev. Lett. 120, 042003 – Published 26 January 2018

Article References Citing Articles (35) PDF HTML Export Citation

ABSTRACT

Physicists at the Large Hadron Collider (LHC) rely on detailed simulations of particle co expectations of what experimental data may look like under different theoretical model assumptions. Petabytes of simulated data are needed to develop analysis techniques, expensive to generate using existing algorithms and computing resources. The modeli and the precise description of particle cascades as they interact with the material in th are the most computationally demanding steps in the simulation pipeline. We therefore deep neural network-based generative model to enable high-fidelity, fast, electromagn simulation. There are still challenges for achieving precision across the entire phase sp current solution can reproduce a variety of particle shower properties while achieving s of up to $100000\times$. This opens the door to a new era of fast simulation that could save s computing time and disk space, while extending the reach of physics searches and pre measurements at the LHC and beyond.

Received 18 July 2017 Revised 27 September 2017

DOI: https://doi.org/10.1103/PhysRevLett.120.042003

Published by the American Physical Society under the terms of the Creative Com Attribution 4.0 International license. Further distribution of this work must maintai the author(s) and the published article's title, journal citation, and DOI. Funded by SCO

Published by the American Physical Society

Physics Subject Headings (PhySH)

Techniques Calorimeters Machine learning

Nuclear Physics Particles & Fields

AUTHORS & AFFILIATIONS

Michela Paganini[1,2,*], Luke de Oliveira[2,†], and Benjamin Nachman[2,‡]

[1]Yale University, New Haven, Connecticut 06520, USA

[2]Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA

*michela.paganini@yale.edu

†lukedeoliveira@lbl.gov

Neural Networks can be thought of as **physical objects** obeying fundamental laws.

CAN STUDY THE INTERACTIONS OF THEIR FUNDAMENTAL COMPONENTS USING EXPERIMENTAL PROCEDURES.

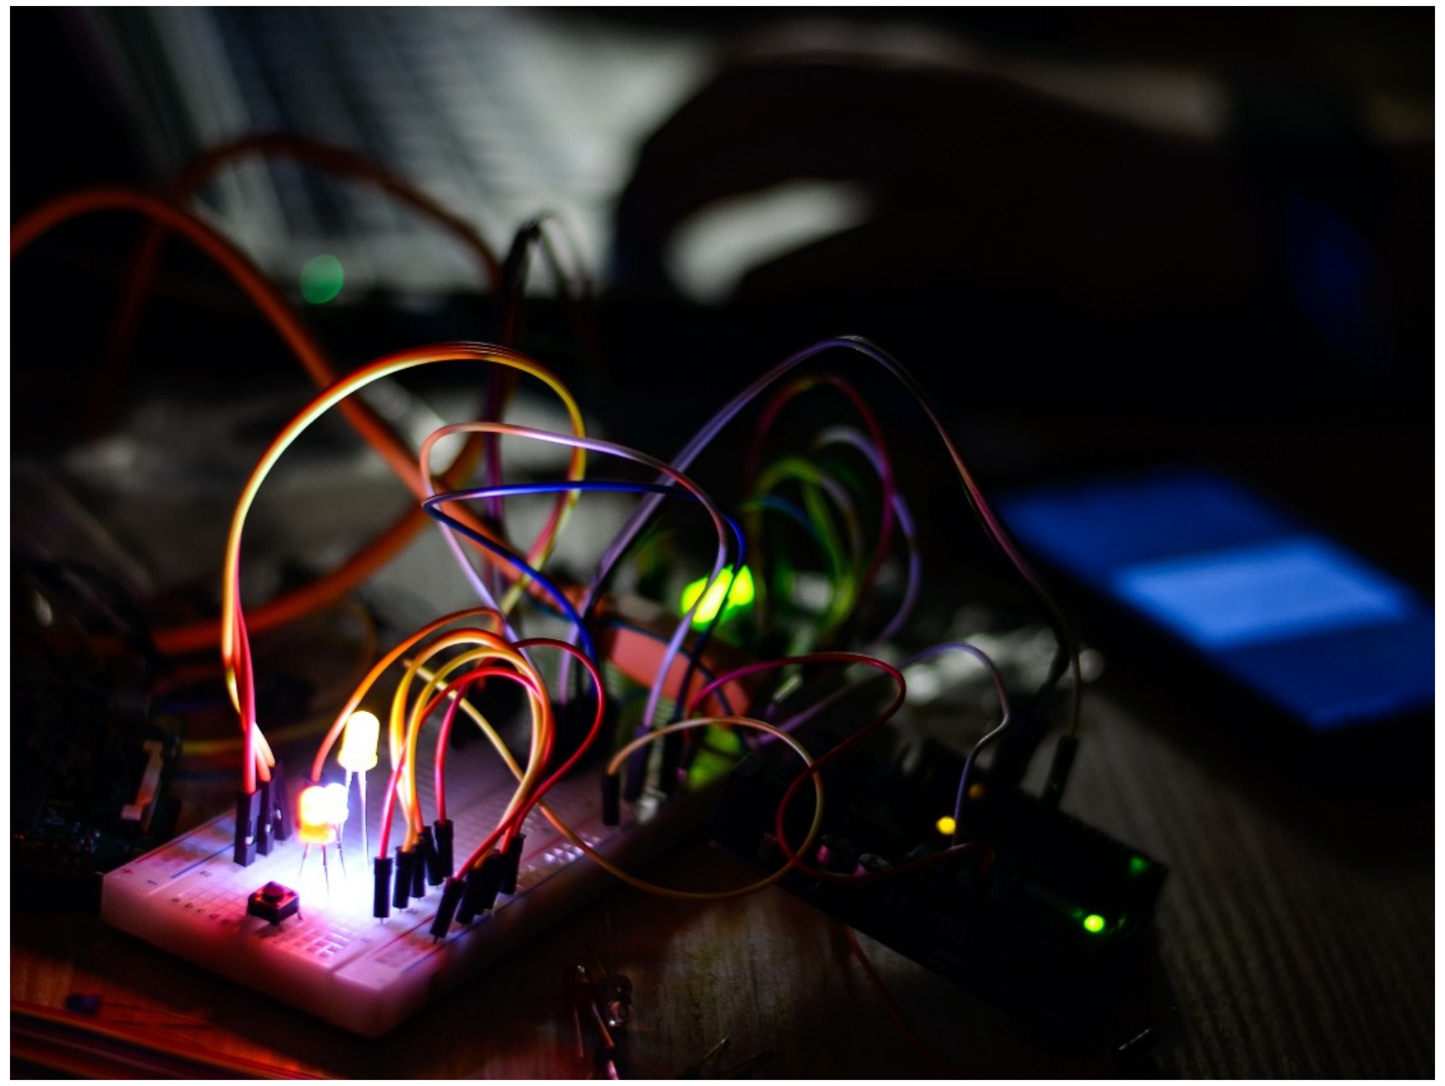

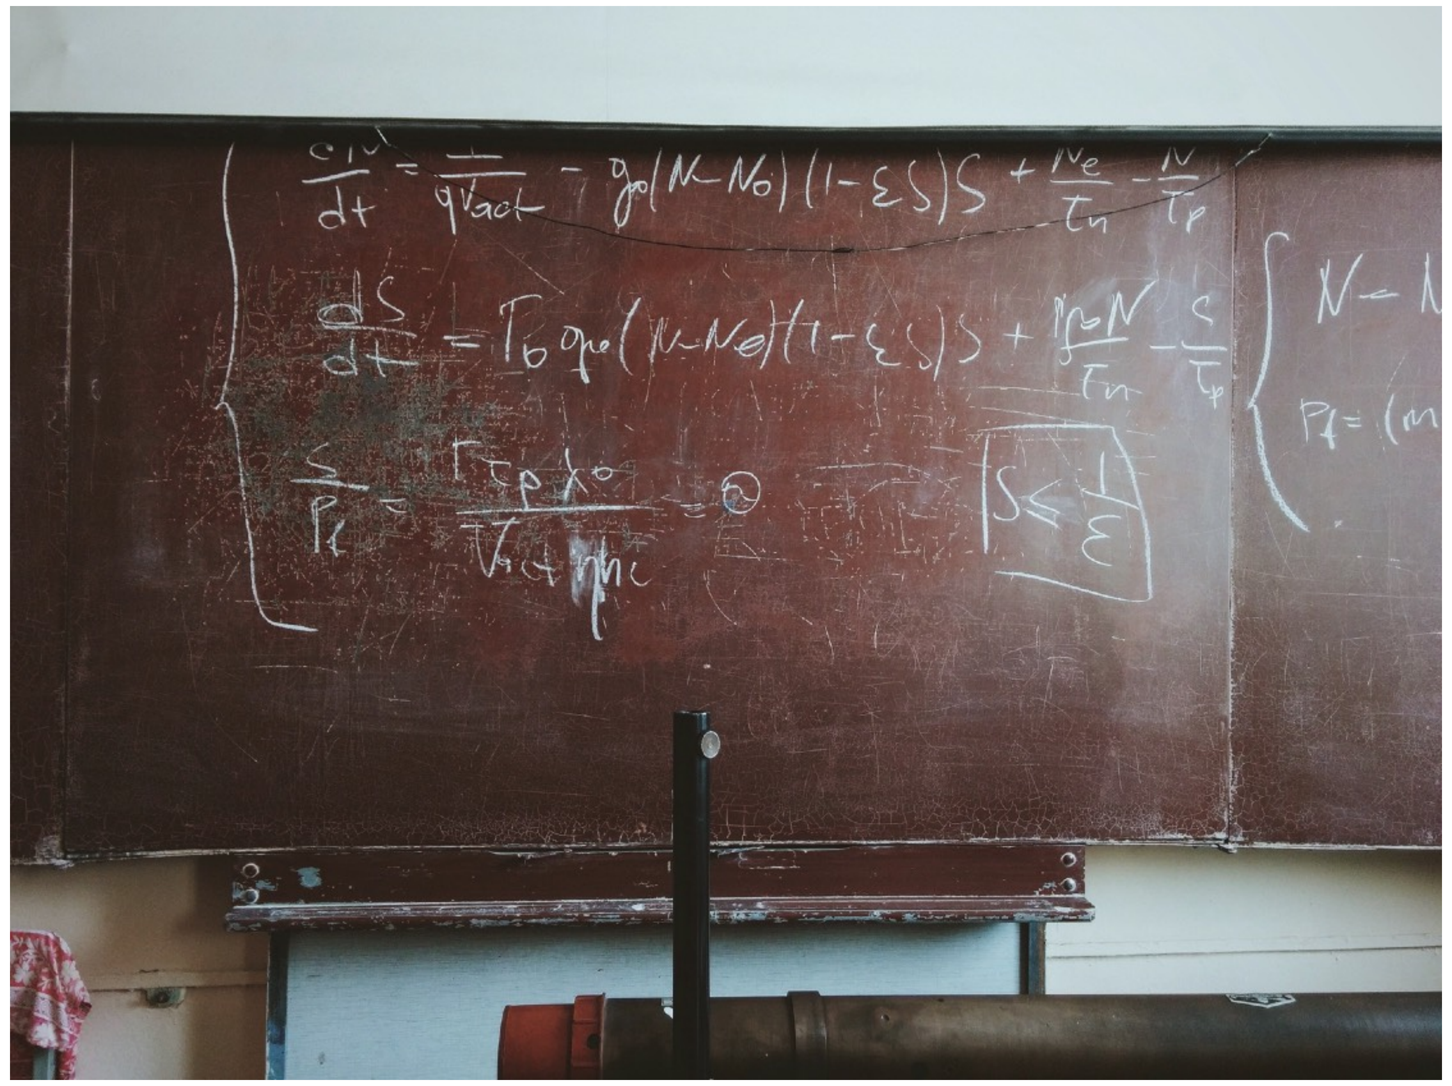

Learn from other Sciences.

Theoretical Science

Experimental Science

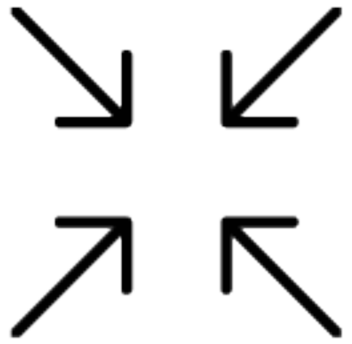

Engineering

0101
1001
0110

Learn from other Sciences.

Theoretical Science

Experimental Science

Engineering

0101
1001
0110

Science is a verb

Recent contributions

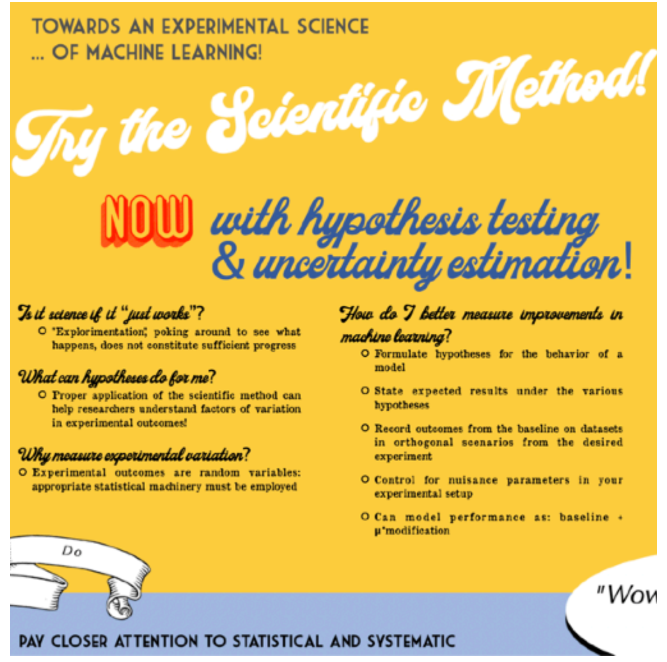

The Scientific Method in the Science of Machine Learning

ICLR 2019 workshop oral

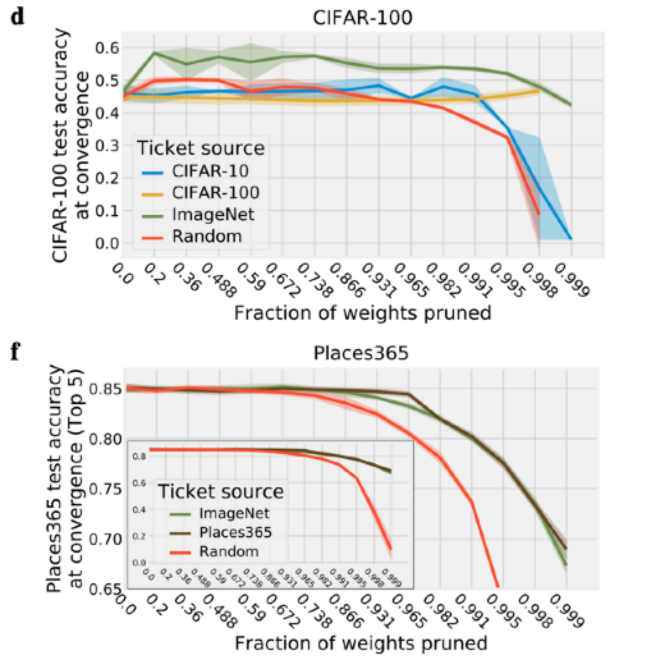

One Ticket to Win Them All: Generalizing Lottery Ticket Initializations Across Datasets and Optimizers

NeurIPS 2019 poster

`dagger`: A Python Framework for Reproducible Machine Learning Experiment Orchestration

under review

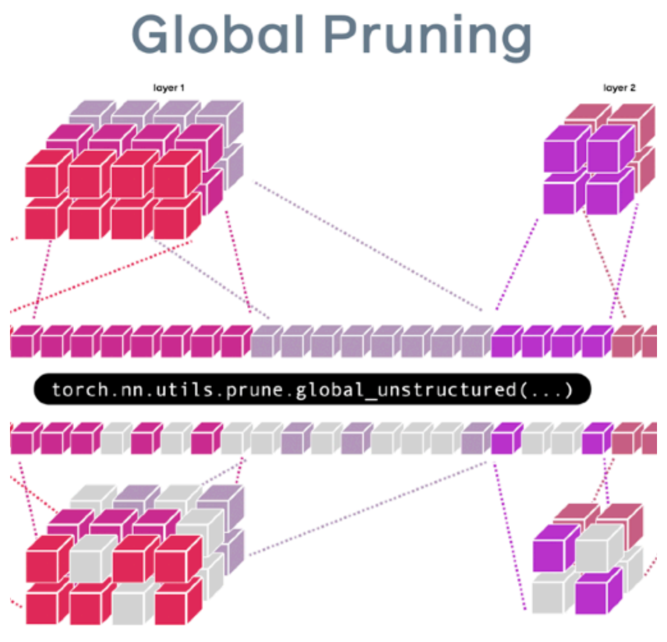

Streamlining Tensor and Network Pruning in PyTorch

ICLR 2020 workshop oral

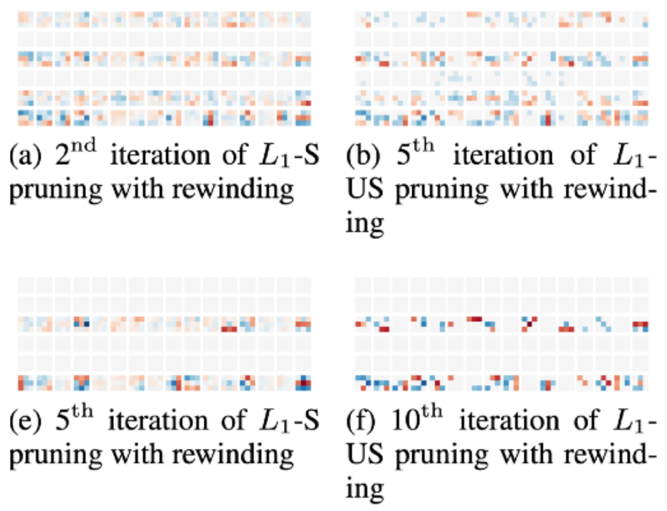

On Iterative Neural Network Pruning, Reinitialization, and the Similarity of Masks

ICLR 2020 workshop poster

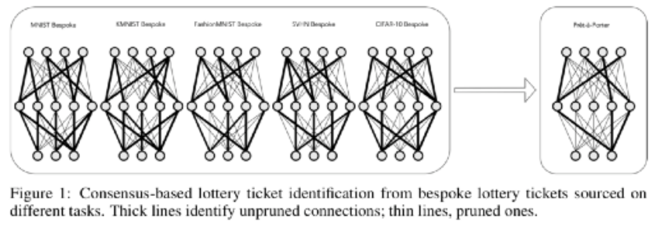

Bespoke vs. Prêt-à-Porter Lottery Tickets: Exploiting Mask Similarity for Trainable Sub-Network Finding

under review

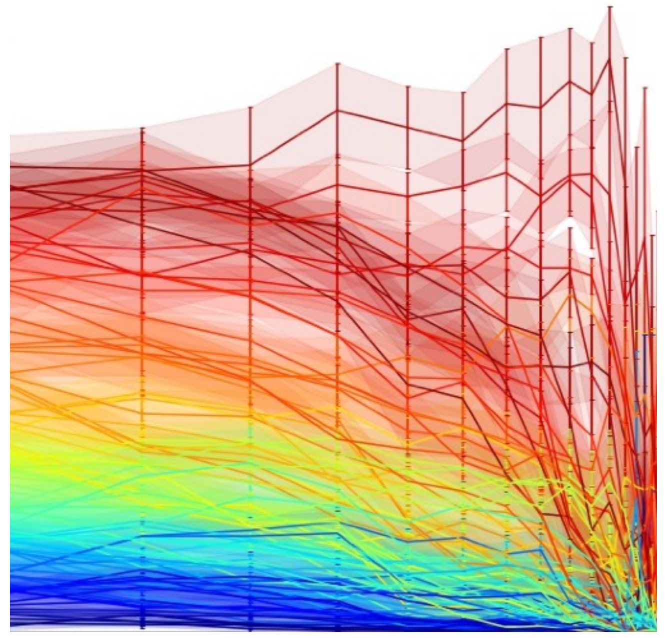

Prune Responsibly

under review

Papers at FAIR

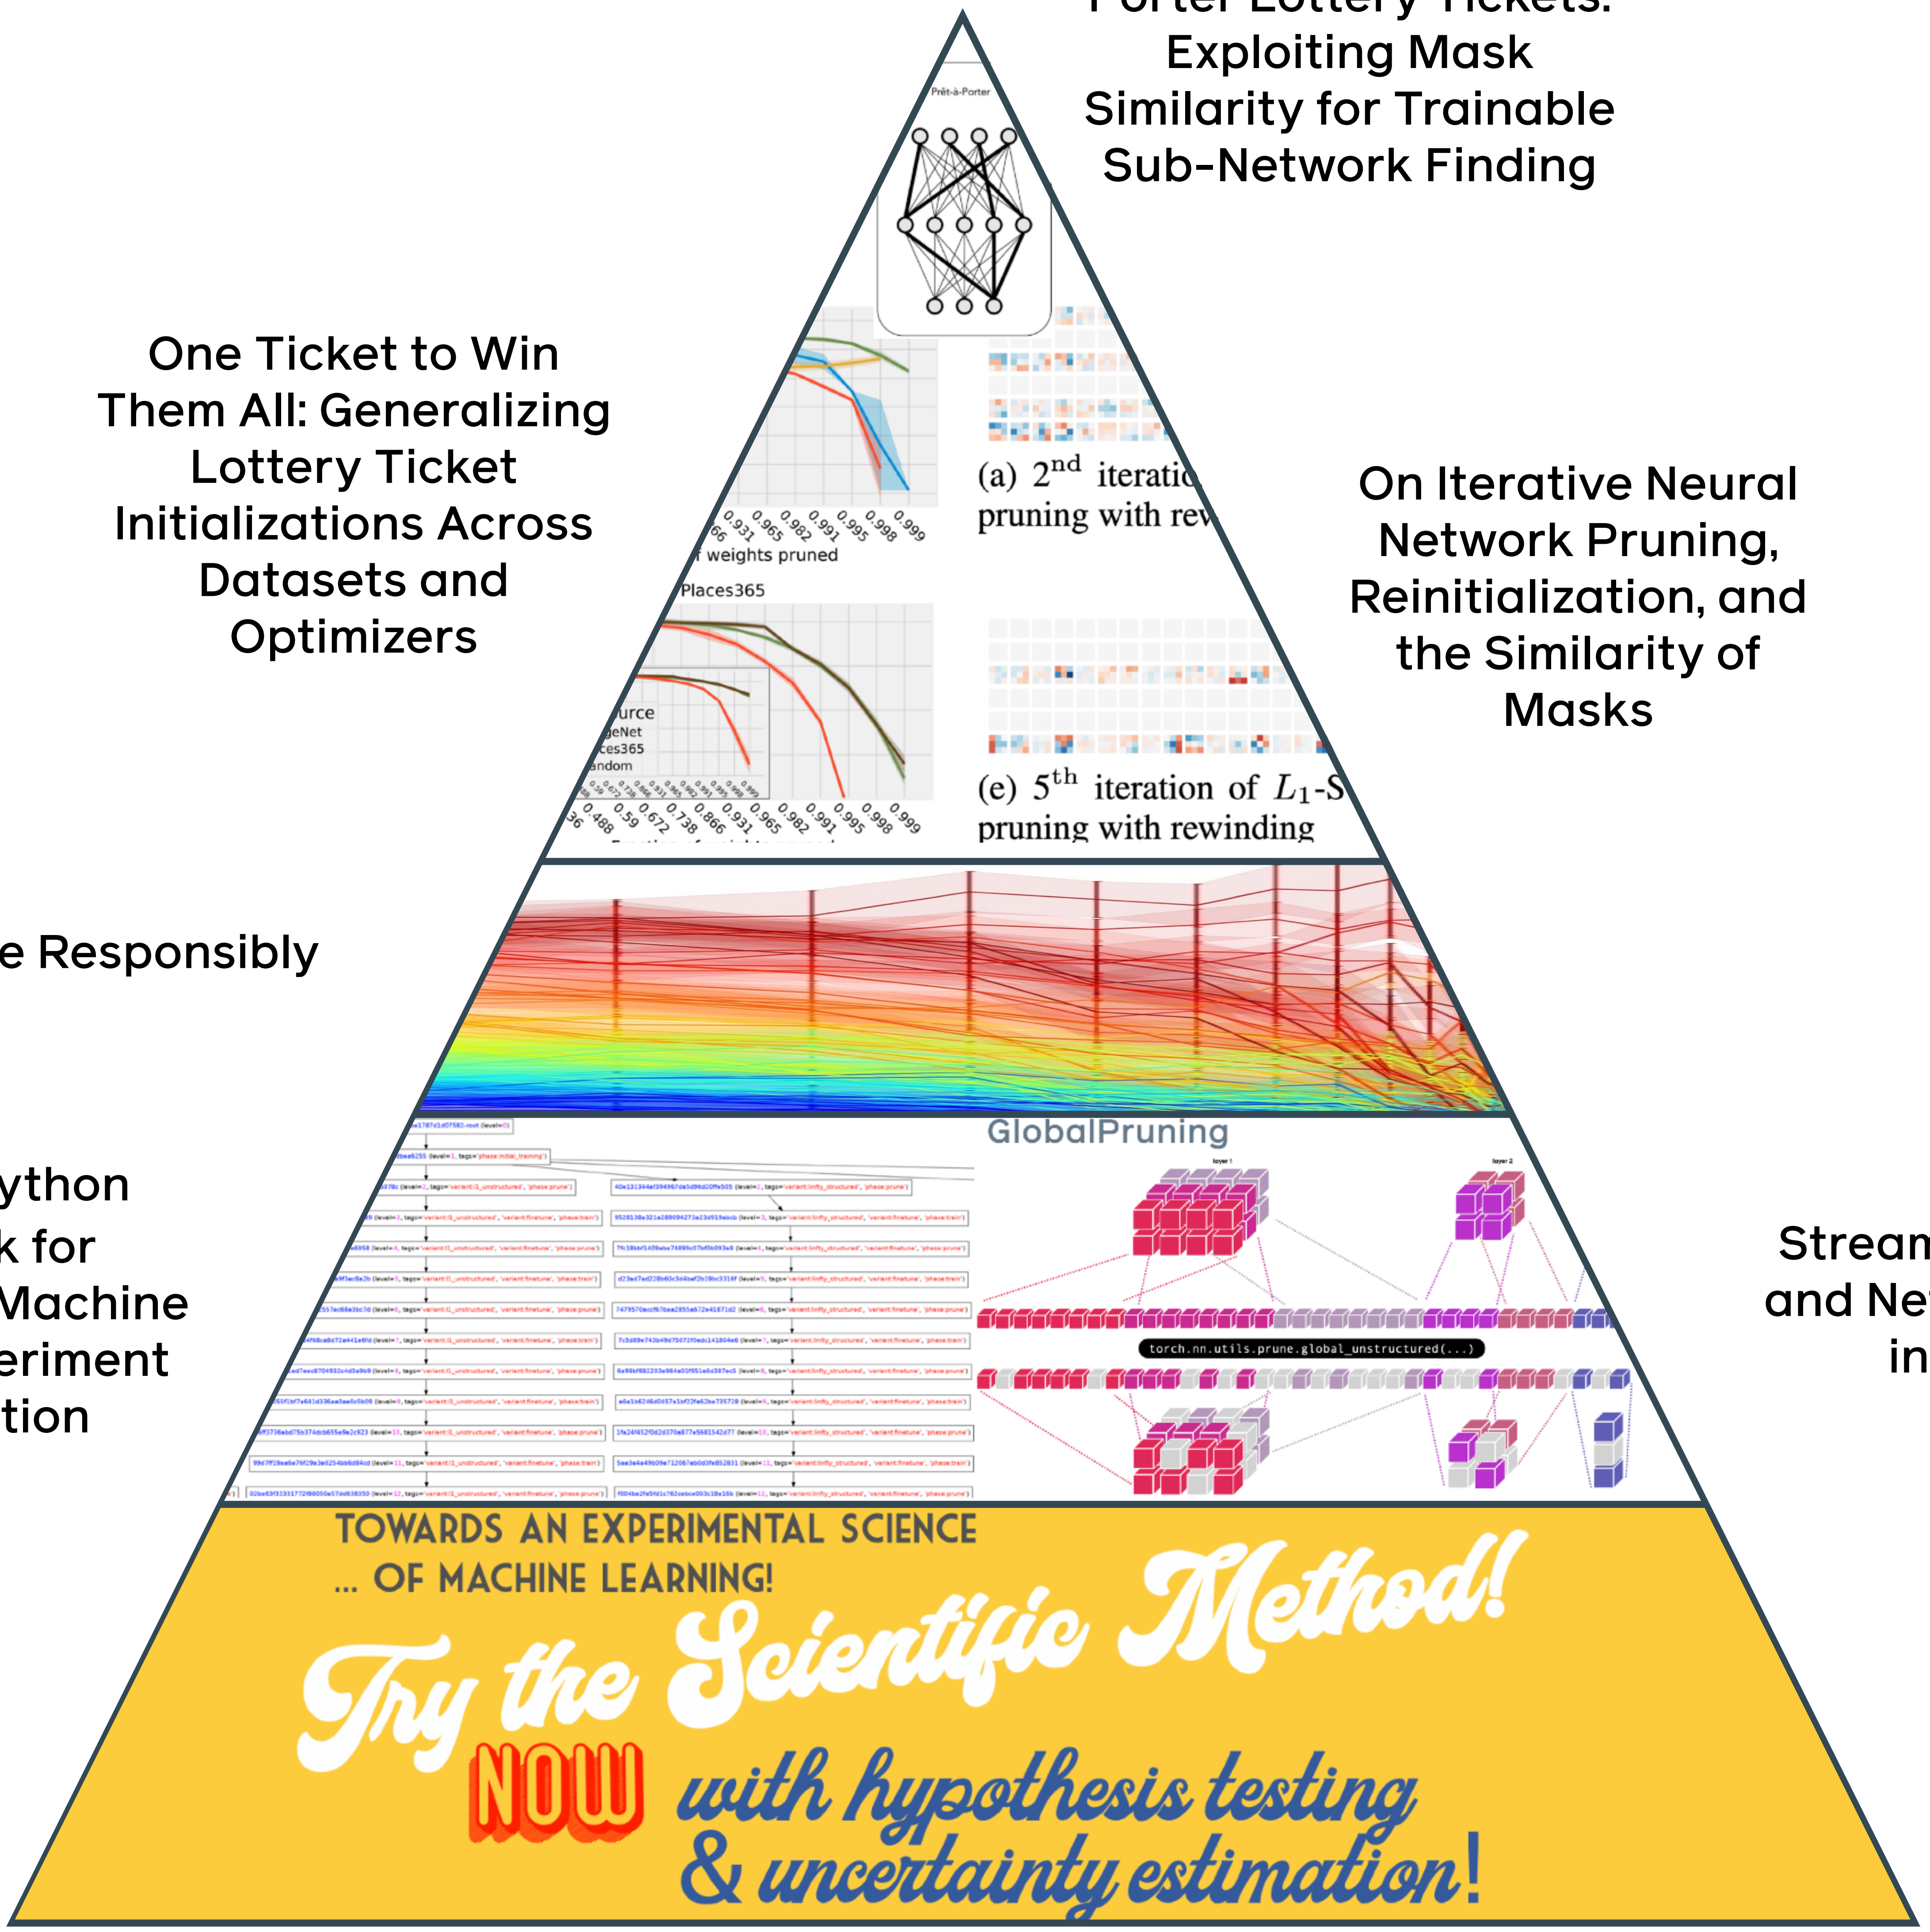

Bespoke vs. Prêt-à-Porter Lottery Tickets: Exploiting Mask Similarity for Trainable Sub-Network Finding — **under review**

One Ticket to Win Them All: Generalizing Lottery Ticket Initializations Across Datasets and Optimizers — **NeurIPS 2019 poster**

On Iterative Neural Network Pruning, Reinitialization, and the Similarity of Masks — **ICLR 2020 workshop poster**

Prune Responsibly — **under review**

`dagger`: A Python Framework for Reproducible Machine Learning Experiment Orchestration — **under review**

Streamlining Tensor and Network Pruning in PyTorch — **ICLR 2020 workshop oral**

The Scientific Method in the Science of Machine Learning — **ICLR 2019 workshop oral**

"We are clamoring for empiricism in the sense of the experimental physicist: running controlled experiments to explain away the mysteries of a complex system."

Ali Rahimi, Ben Recht

"There’s a mass influx of newcomers to our field and we’re equipping them with little more than folklore and pre-trained deep nets, then asking them to innovate. We can barely agree on the phenomena that we should be explaining away. "

Ali Rahimi, Ben Recht

Agenda

1. The scientific method in the science of ML
2. Hypothesis formulation and testing
3. Uncertainties
4. Blind analysis and pre-registration
5. Reproducibility
6. Open-source contributions
7. Measuring the disproportionate harm of pruning

TOWARDS AN EXPERIMENTAL SCIENCE
... OF MACHINE LEARNING!

Try the Scientific Method!

NOW with hypothesis testing & uncertainty estimation!

Jessica Forde
Project Jupyter

Michela Paganini
Facebook AI Research

Is it science if it "just works"?

- "Explorimentation", poking around to see what happens, does not constitute sufficient progress

What can hypotheses do for me?

- Proper application of the scientific method can help researchers understand factors of variation in experimental outcomes!

Why measure experimental variation?

- Experimental outcomes are random variables: appropriate statistical machinery must be employed

How do I better measure improvements in machine learning?

- Formulate hypotheses for the behavior of a model
- State expected results under the various hypotheses
- Record outcomes from the baseline on datasets in orthogonal scenarios from the desired experiment
- Control for nuisance parameters in your experimental setup
- Can model performance as: baseline + μ*modification

"Wow! I can read it on the arXiv!"

PAY CLOSER ATTENTION TO STATISTICAL AND SYSTEMATIC UNCERTAINTY AND CONSIDER THE SCIENTIFIC ROBUSTNESS OF CLAIMS

facebook Artificial Intelligence

What can we learn from the other sciences?

The Scientific Method in the Science of Machine Learning, `arXiv:1904.10922`

The one and only way to make objective statements?

A social contract among scientists to harmonize workflows and compare findings?

The scientific method in the science of ML

Transparency

Falsifiability

Reproducibility

Intellectual Honesty

Key Steps for Experimental Scientific Research.

hypothesis formulation

statement of expectations

experiment design

statistical analysis

uncertainty estimation

Key Steps for Experimental Scientific Research.

hypothesis formulation	"The null hypothesis is ..., the alternative hypothesis is ..."
statement of expectations	"If the hypothesis is right, then I should expect to observe ..."
experiment design	"I design this experiment to be sensitive to..."
statistical analysis	"Do I observe the expected effect? Is it stronger or weaker than expected?"
uncertainty estimation	"Do I have enough observations and did I account for systematic biases?"

"The first step towards a scientific formulation of ML then demands a more dramatic shift in priorities from drawing and recording single instances of experimental results to collecting enough data to gain an understanding of population statistics."

"it is plausible that a significant percentage of published work claiming state-of-the-art performance actually has no statistical sensitivity to measure their improvement over competing methods."

Unreproducible Research is Reproducible

Xavier Bouthillier, César Laurent, Pascal Vincent ; Proceedings of the 36th International Conference on Machine Learning, PMLR 97:725-734, 2019.

- unreproducible findings can be built upon reproducible methods
- not just a matter of deterministic reproducibility of methods and single numerical results
- necessity of ensuring the reproducibility of empirical findings and conclusions by properly accounting for essential sources of variations
- more energy should be devoted to proper empirical research in our community
- promote the use of more rigorous and diversified methodologies

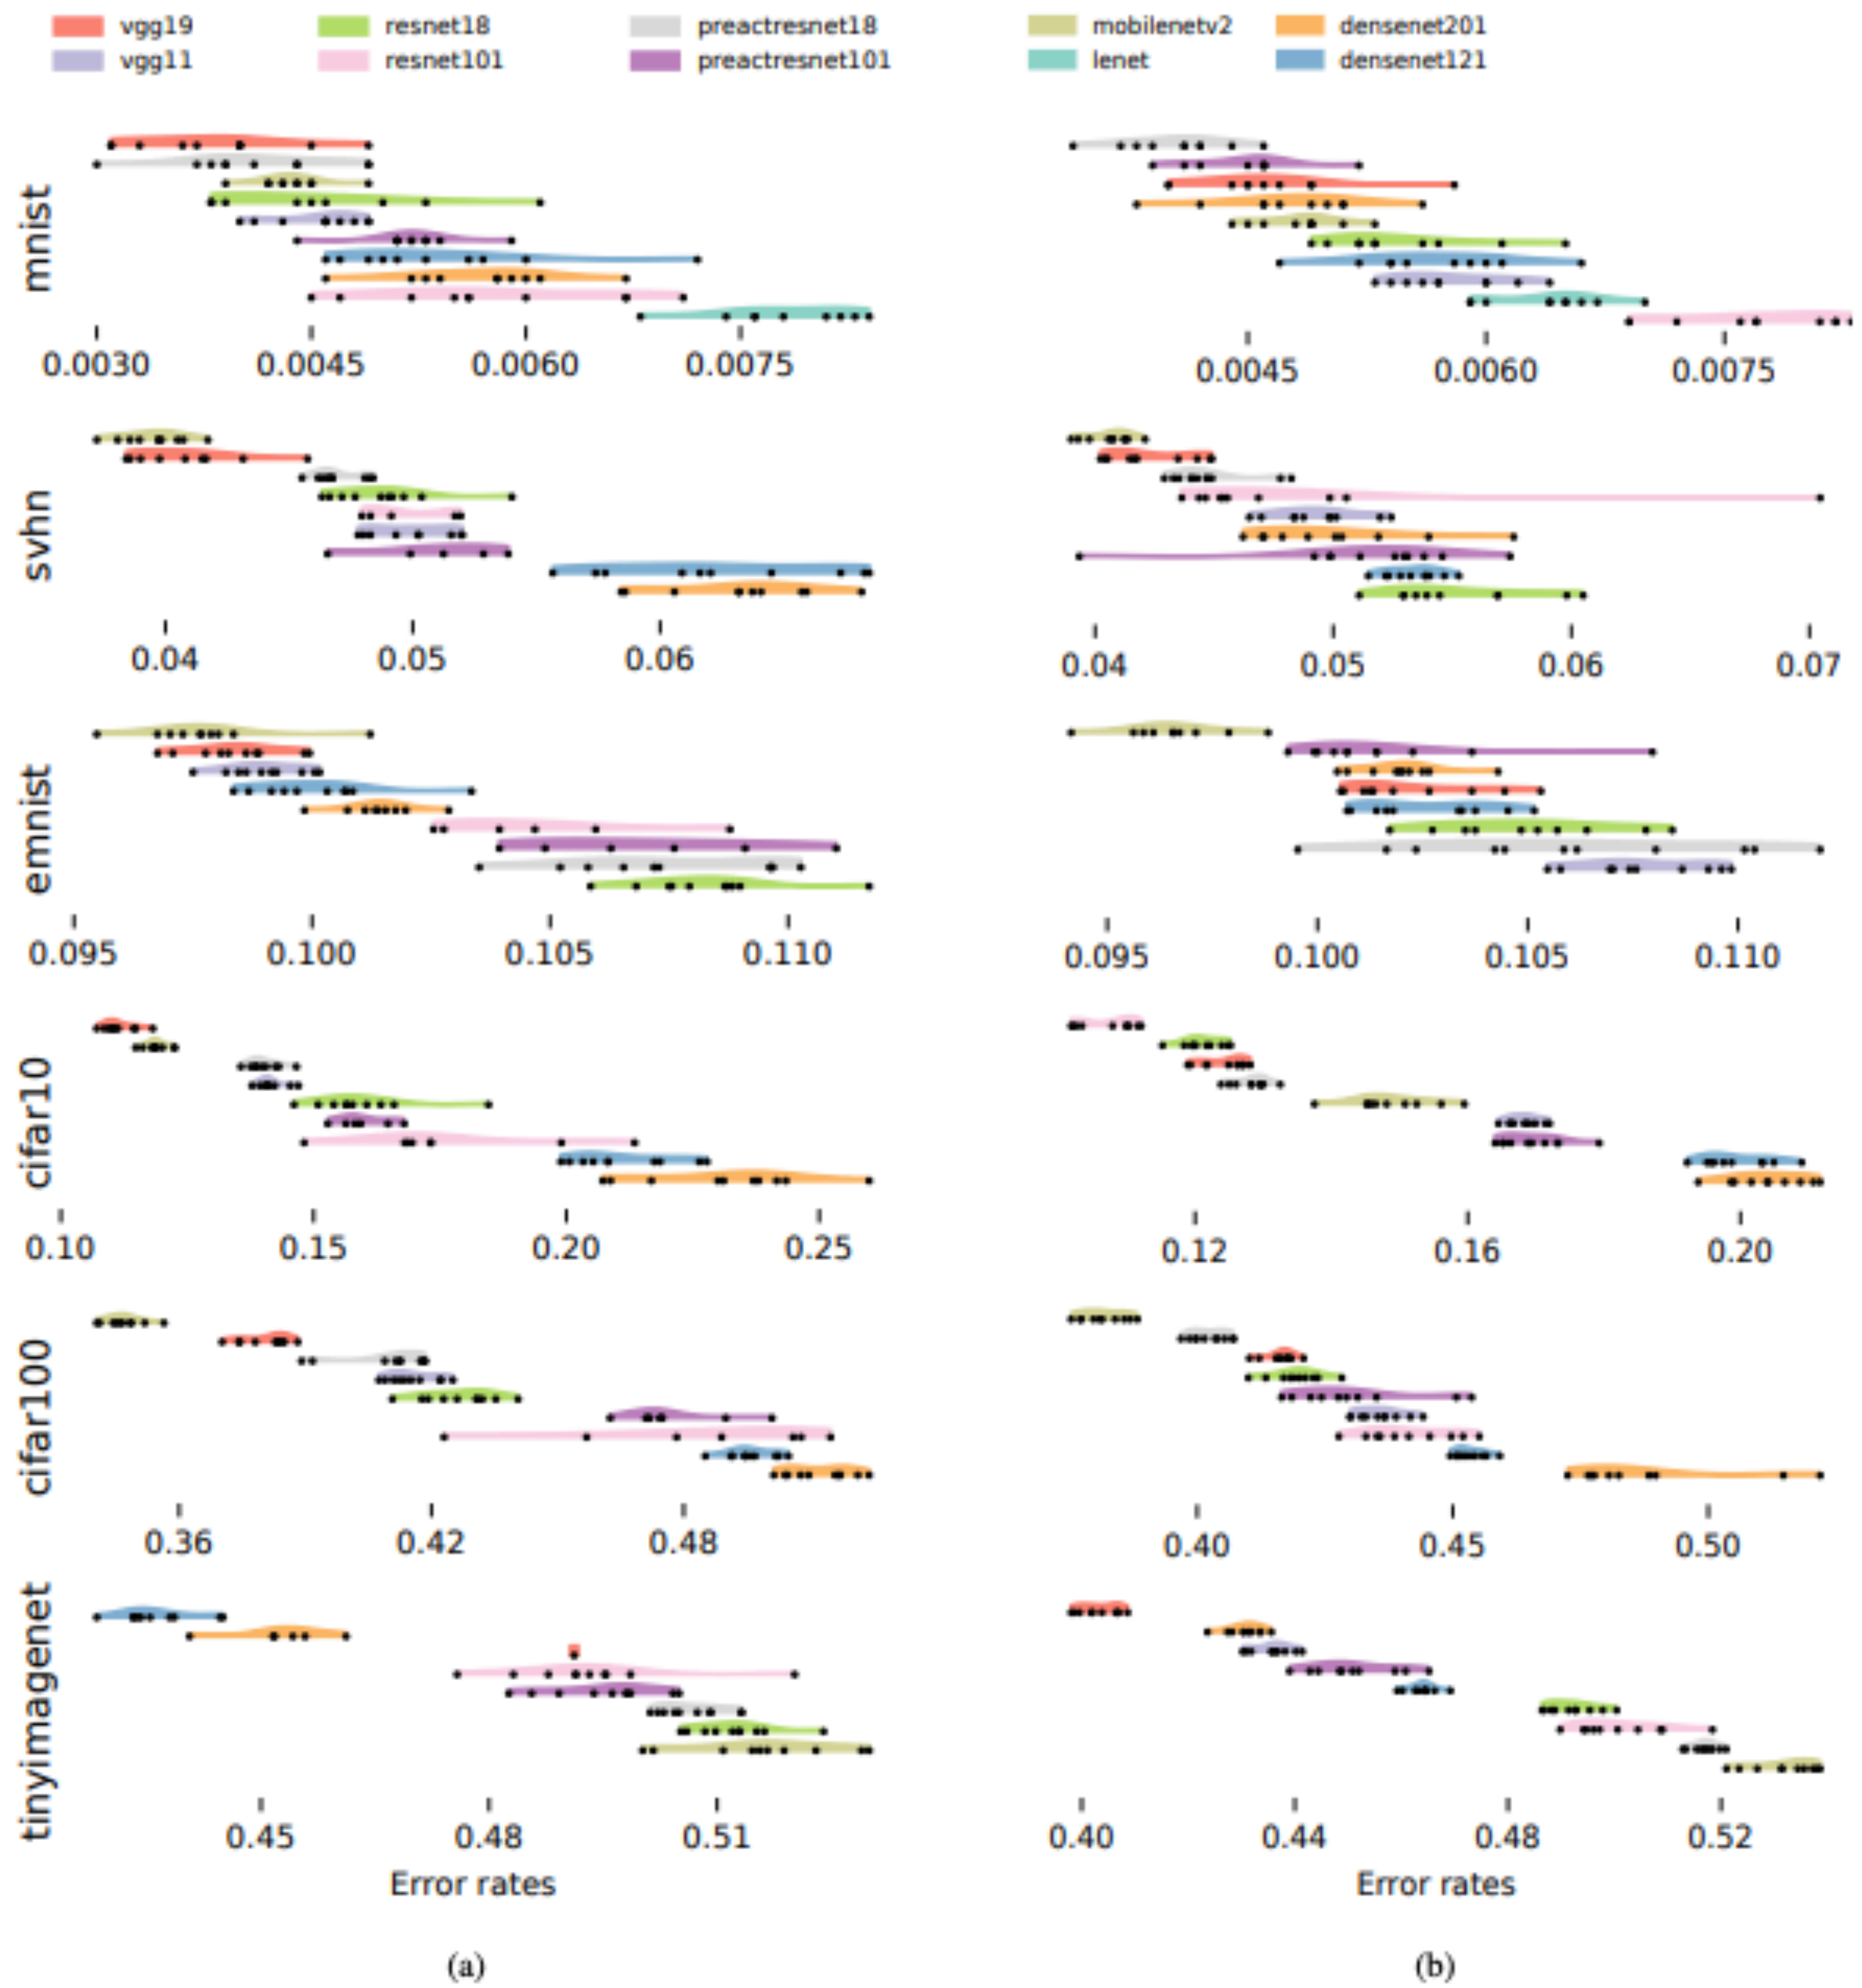

Measurements are affected by sources of variations

Reproducibility

Reproducibility

nature > news q&a > article

NEWS Q&A · 19 DECEMBER 2019

This AI researcher is trying to ward off a reproducibility crisis

Joelle Pineau is leading an effort to encourage artificial-intelligence researchers to open up their code.

Elizabeth Gibney

RESEARCH

How the AI community can get serious about reproducibility

12/9/2019

FACEBOOK AI

Percentage of accepted papers that included a link to code in their final camera-ready submission

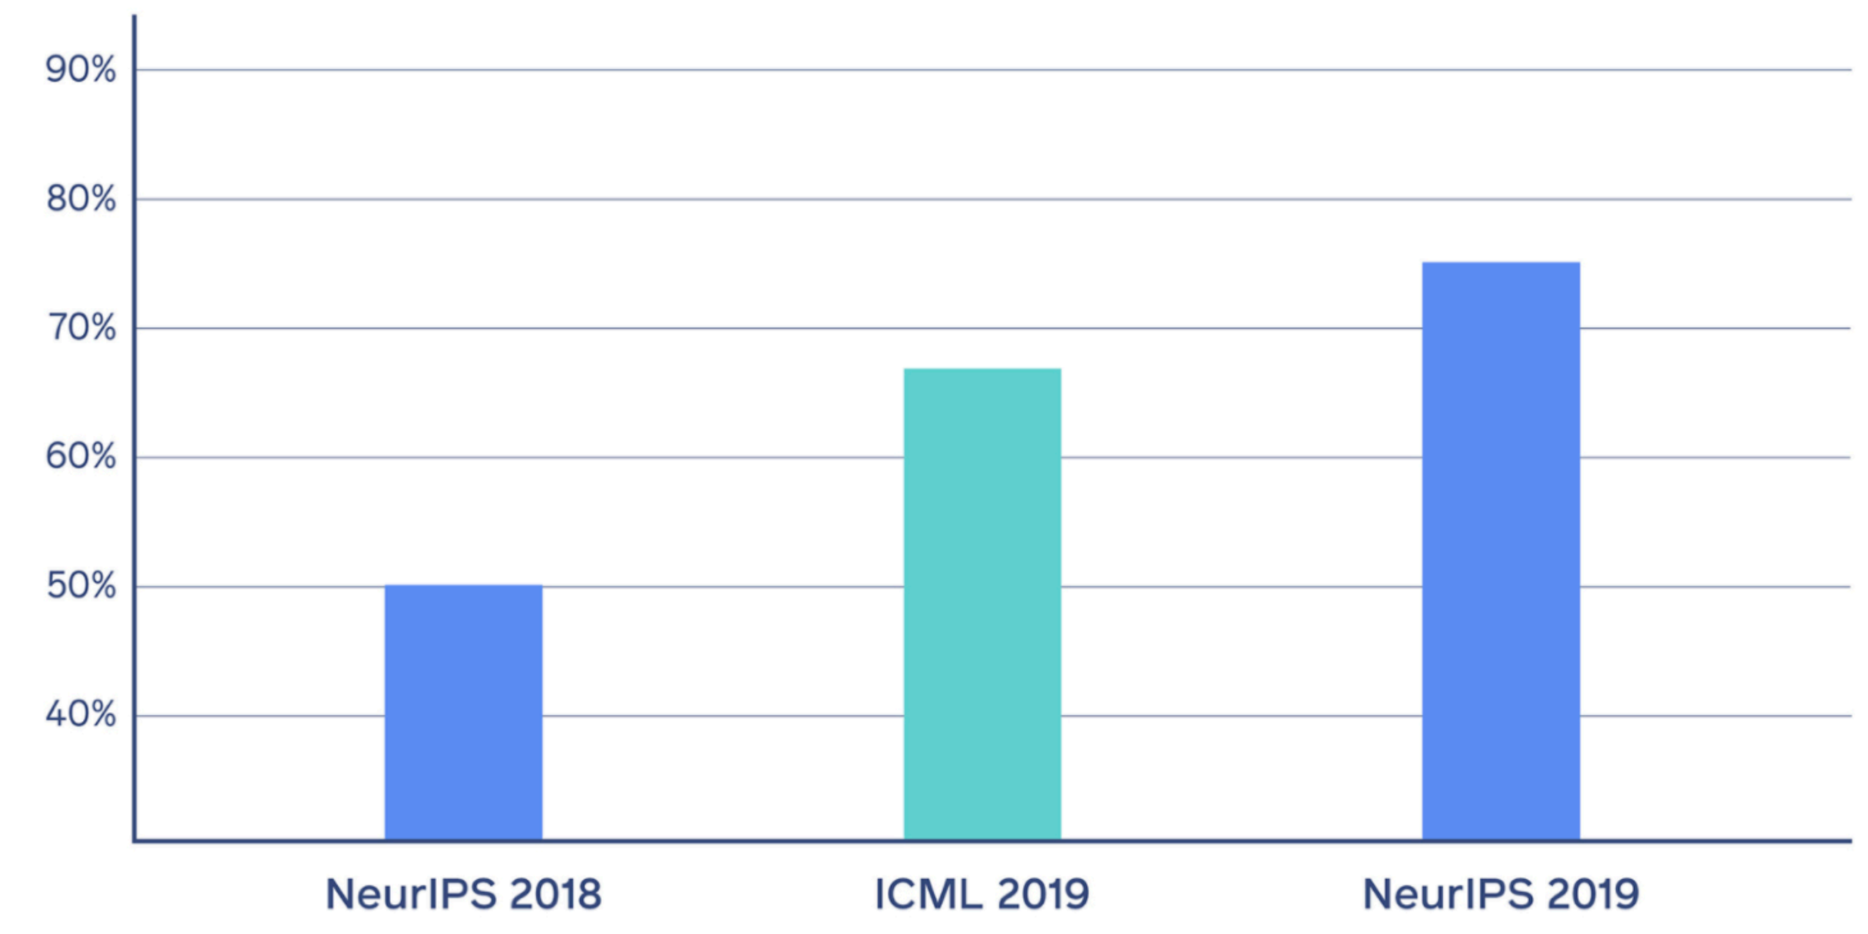

https://ai.facebook.com/blog/how-the-ai-community-can-get-serious-about-reproducibility/

Reproducibility

Trust

Integrity

Confidence

Progress

Reliability

Papers With Code

sotabench

June 10, 2019

Towards Reproducible Research with PyTorch Hub

PYTORCH
HUB

PUBLISHING MODELS

PyTorch Hub supports publishing pre-trained models (model definitions and pre-trained weights) to a GitHub repository by adding a simple `hubconf.py` file.

Discovery

Find the best models related to your research/application!

Reproducibility

Spend minutes instead of days on baselines

Responsibility

Publish solid papers with reproducible results.

Slide adapted from Ailing Zhang

Blind Analysis and Pre-Registration

"One enemy of robust science is our humanity — our appetite for being right, and our tendency to find patterns in noise, to see supporting evidence for what we already believe is true, and to ignore the facts that do not fit."

Let's think about cognitive bias
Nature Editorial

Issues

Asymmetrical Attention

"Carefully debugging analyses and debunking data that counter a favoured hypothesis, while letting evidence in favour of the hypothesis slide by unexamined."

Let's think about cognitive bias
Nature Editorial

Reproducibility Crisis

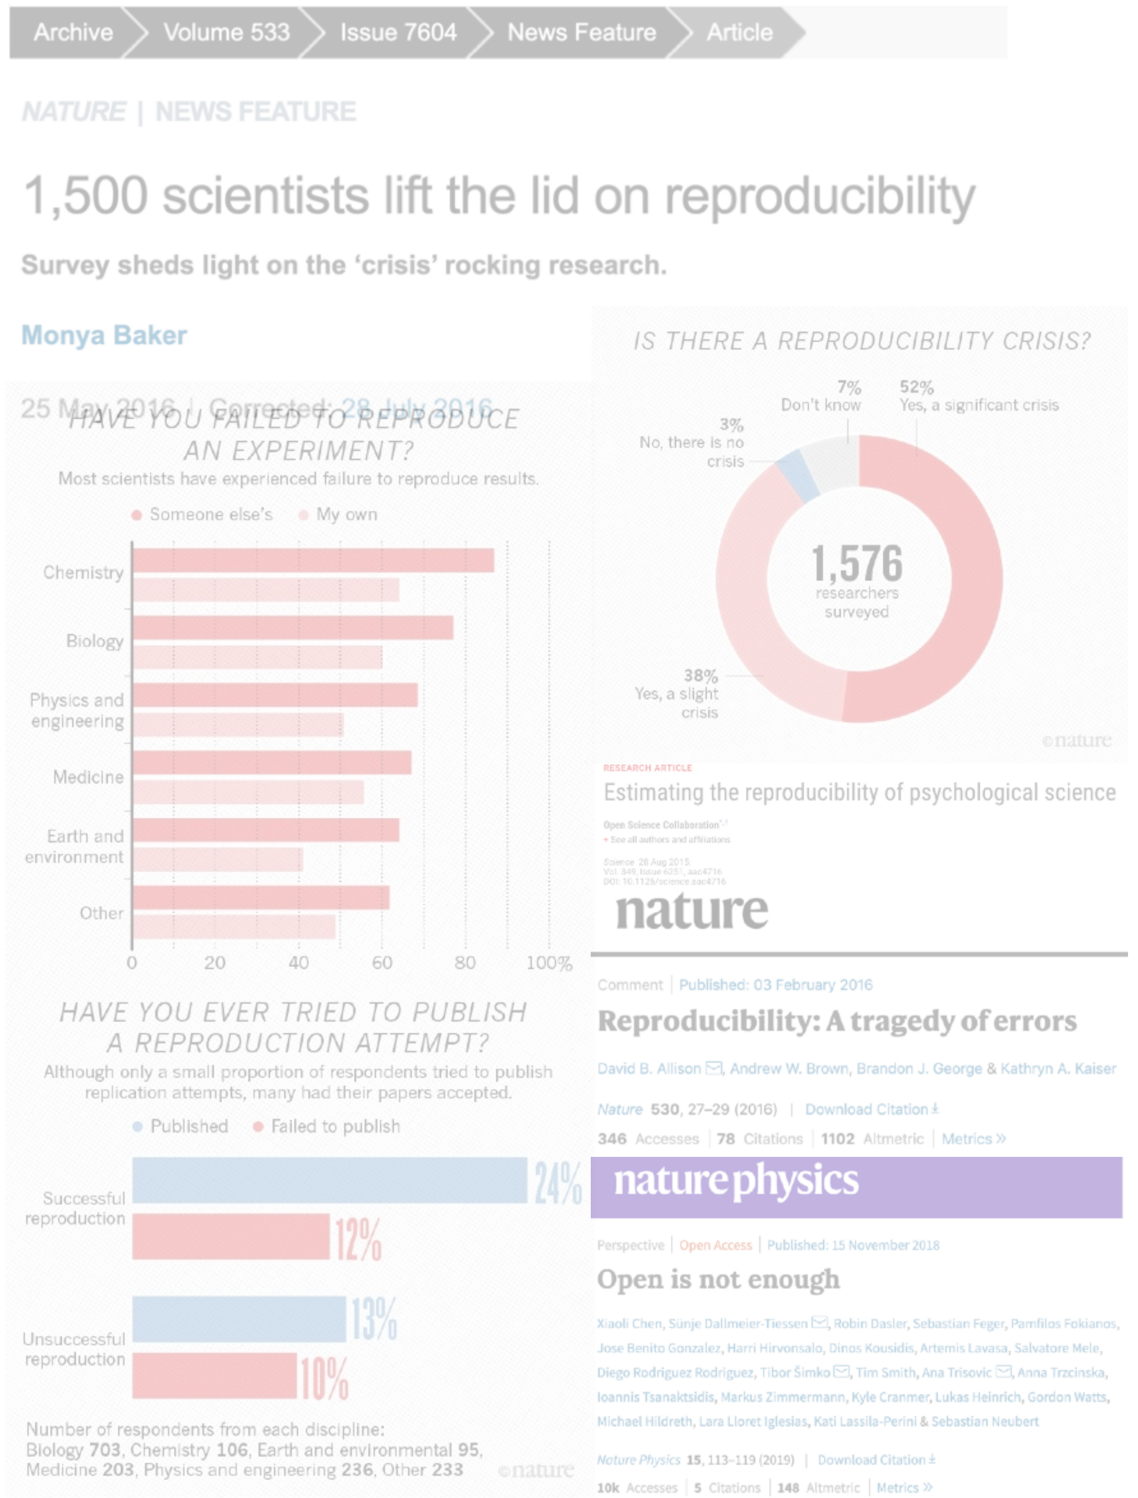

Reporting Bias

The decision to publish or withhold results is influenced by how strongly they support the leading hypothesis. Result sections become less useful. Promising ideas are rejected for not achieving clear SOTA.

Recency Bias

Bulk of attention to minor iterations over a few latest trends. Harder to publish novel ideas that buck the trend.

Issues

Asymmetrical Attention

”Carefully debugging analyses and debunking data that counter a favoured hypothesis, while letting evidence in favour of the hypothesis slide by unexamined.”

Reproducibility Crisis

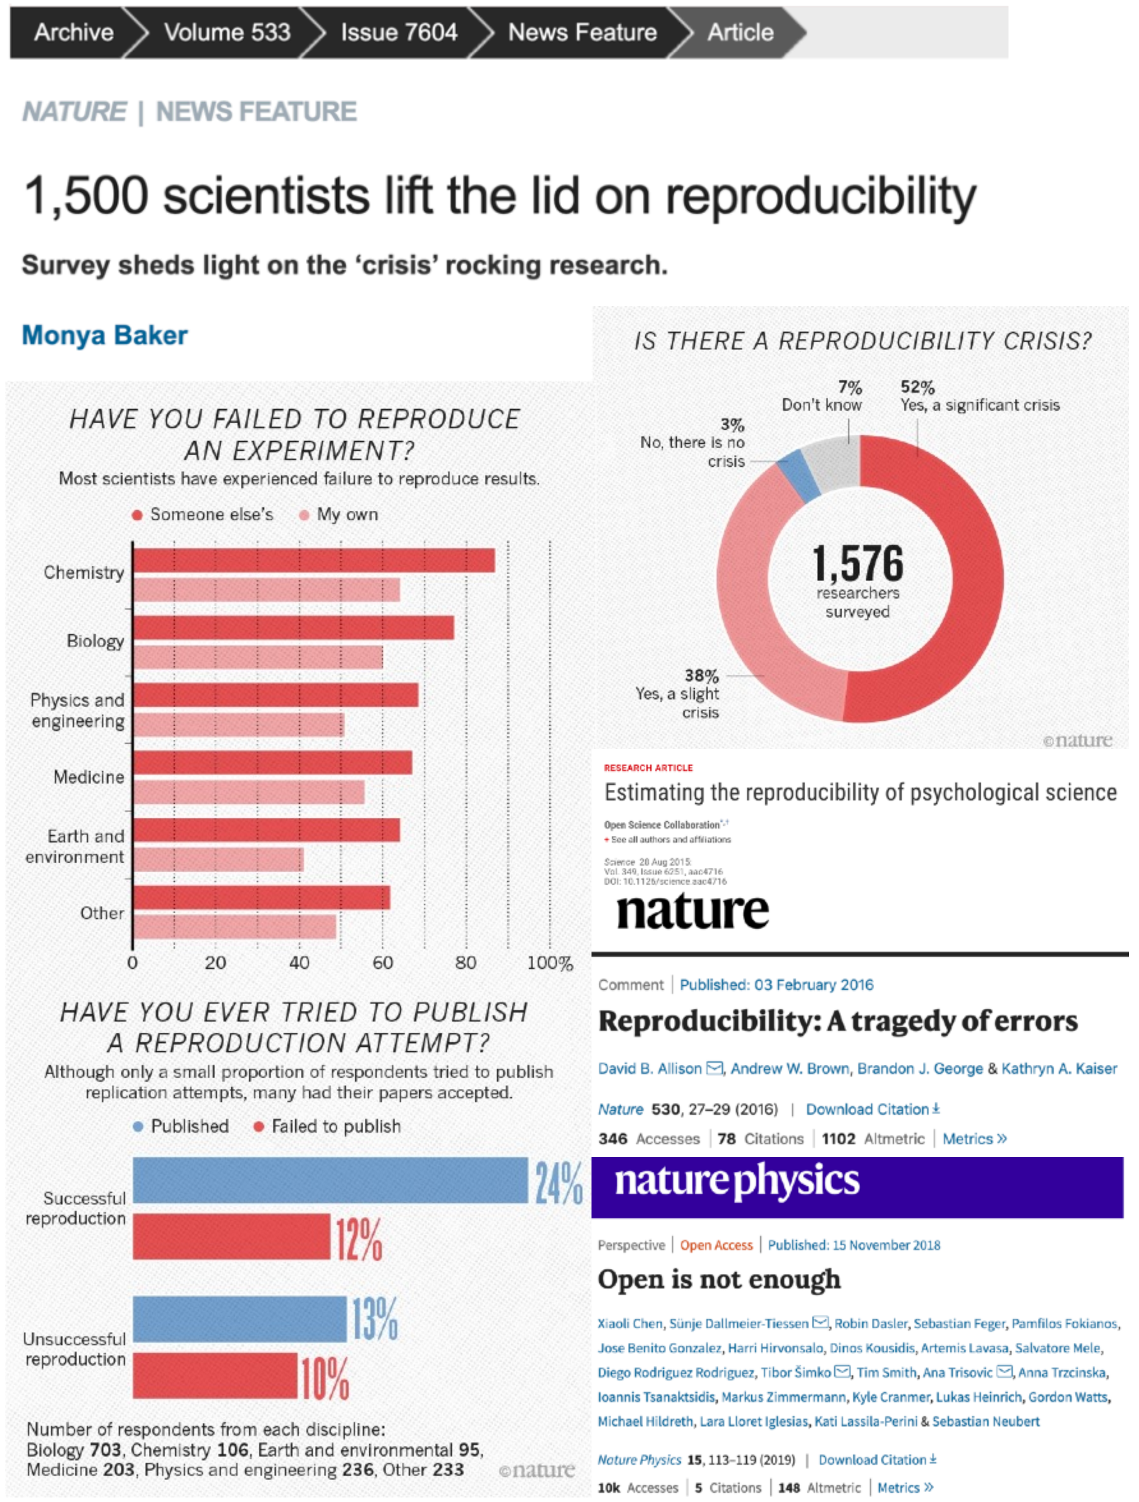

Reporting Bias

The decision to publish or withhold results is influenced by how strongly they support the leading hypothesis. Result sections become less useful. Promising ideas are rejected for not achieving clear SOTA.

Recency Bias

Bulk of attention to minor iterations over a few latest trends. Harder to publish novel ideas that buck the trend.

Issues

Asymmetrical Attention

"Carefully debugging analyses and debunking data that counter a favoured hypothesis, while letting evidence in favour of the hypothesis slide by unexamined."

Reproducibility Crisis

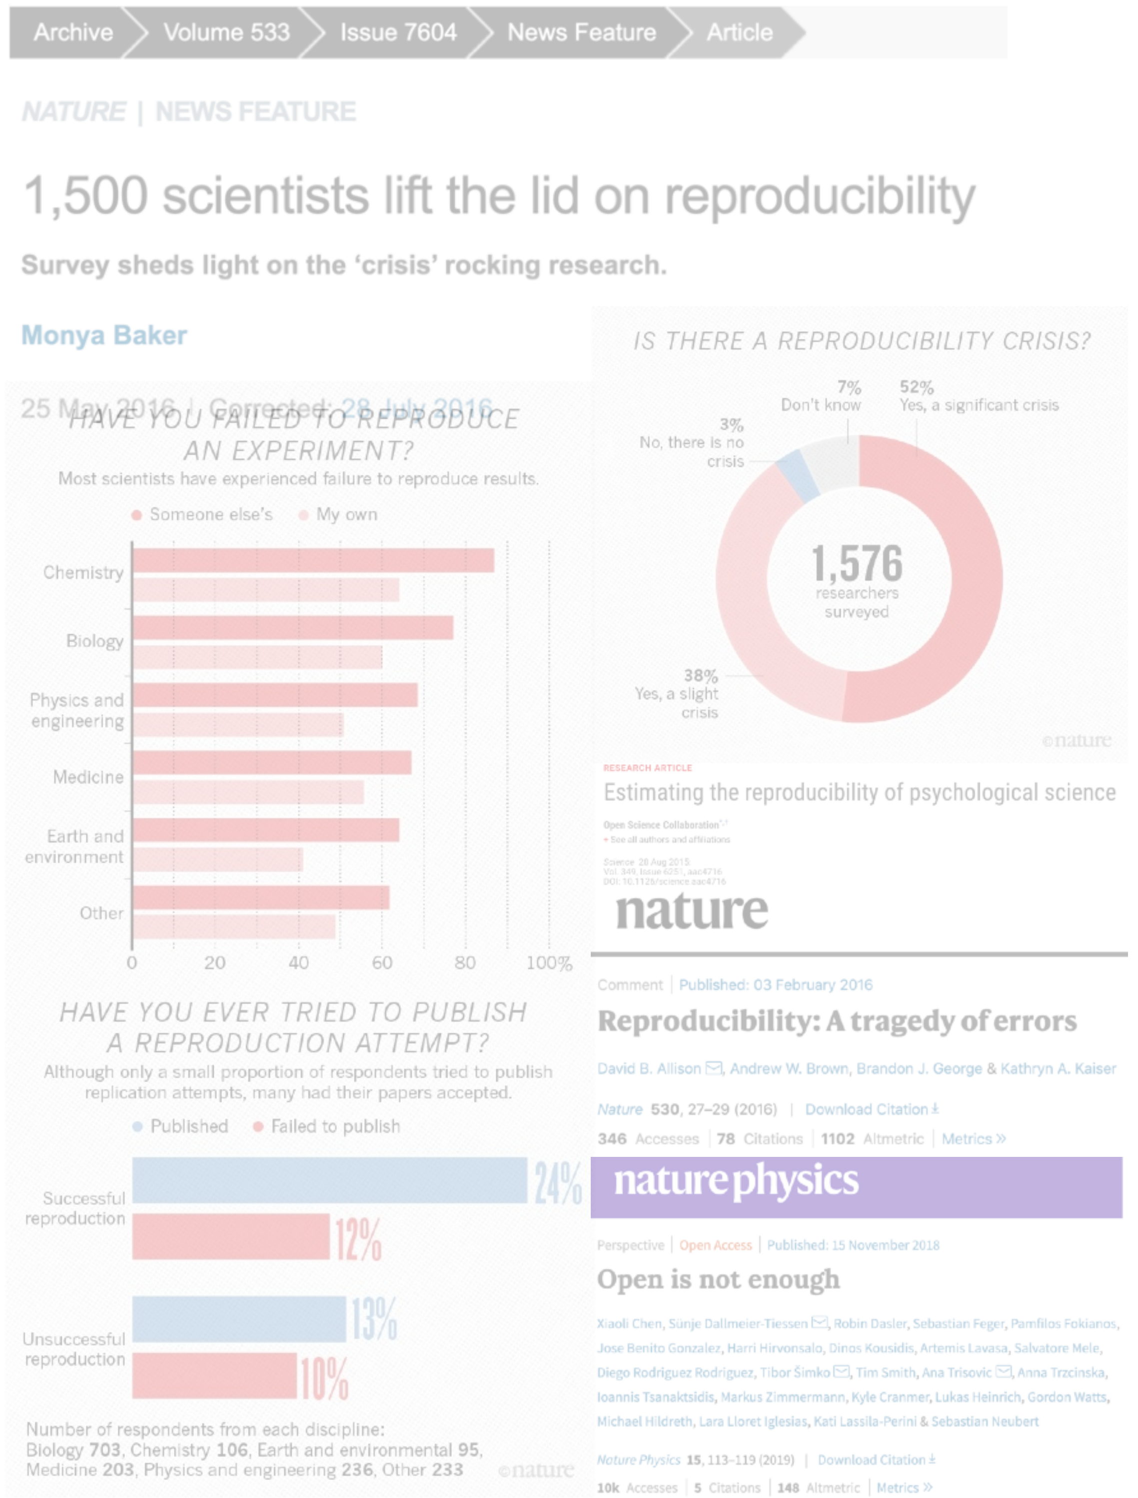

Reporting Bias

The decision to publish or withhold results is influenced by how strongly they support the leading hypothesis. Result sections become less useful. Promising ideas are rejected for not achieving clear SOTA.

Recency Bias

Bulk of attention to minor iterations over a few latest trends. Harder to publish novel ideas that buck the trend.

Issues

Asymmetrical Attention

”Carefully debugging analyses and debunking data that counter a favoured hypothesis, while letting evidence in favour of the hypothesis slide by unexamined.”

Reproducibility Crisis

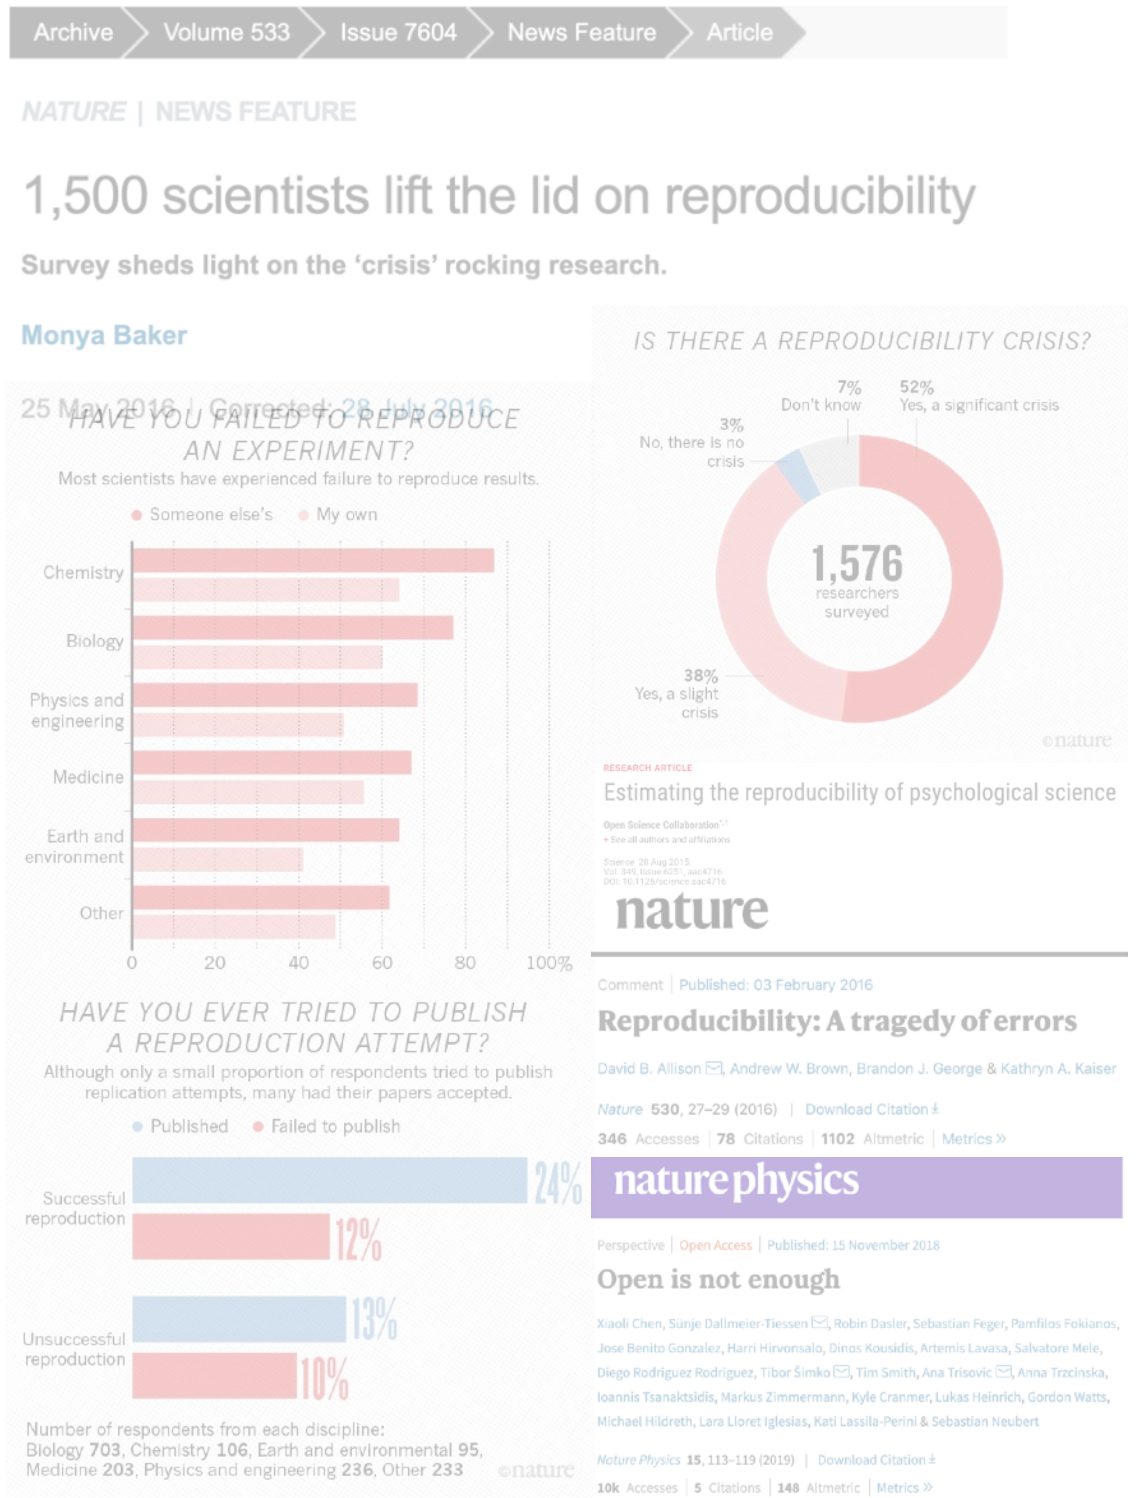

Reporting Bias

The decision to publish or withhold results is influenced by how strongly they support the leading hypothesis. Result sections become less useful. Promising ideas are rejected for not achieving clear SOTA.

Recency Bias

Bulk of attention to minor iterations over a few latest trends. Harder to publish novel ideas that buck the trend.

for more, see *Bias* by Miguel Delgado-Rodriguez, Javier Llorca

Examples from the Sciences

From *Blind Analysis in Nuclear and Particle Physics* by Joshua R. Klein and Aaron Roodman

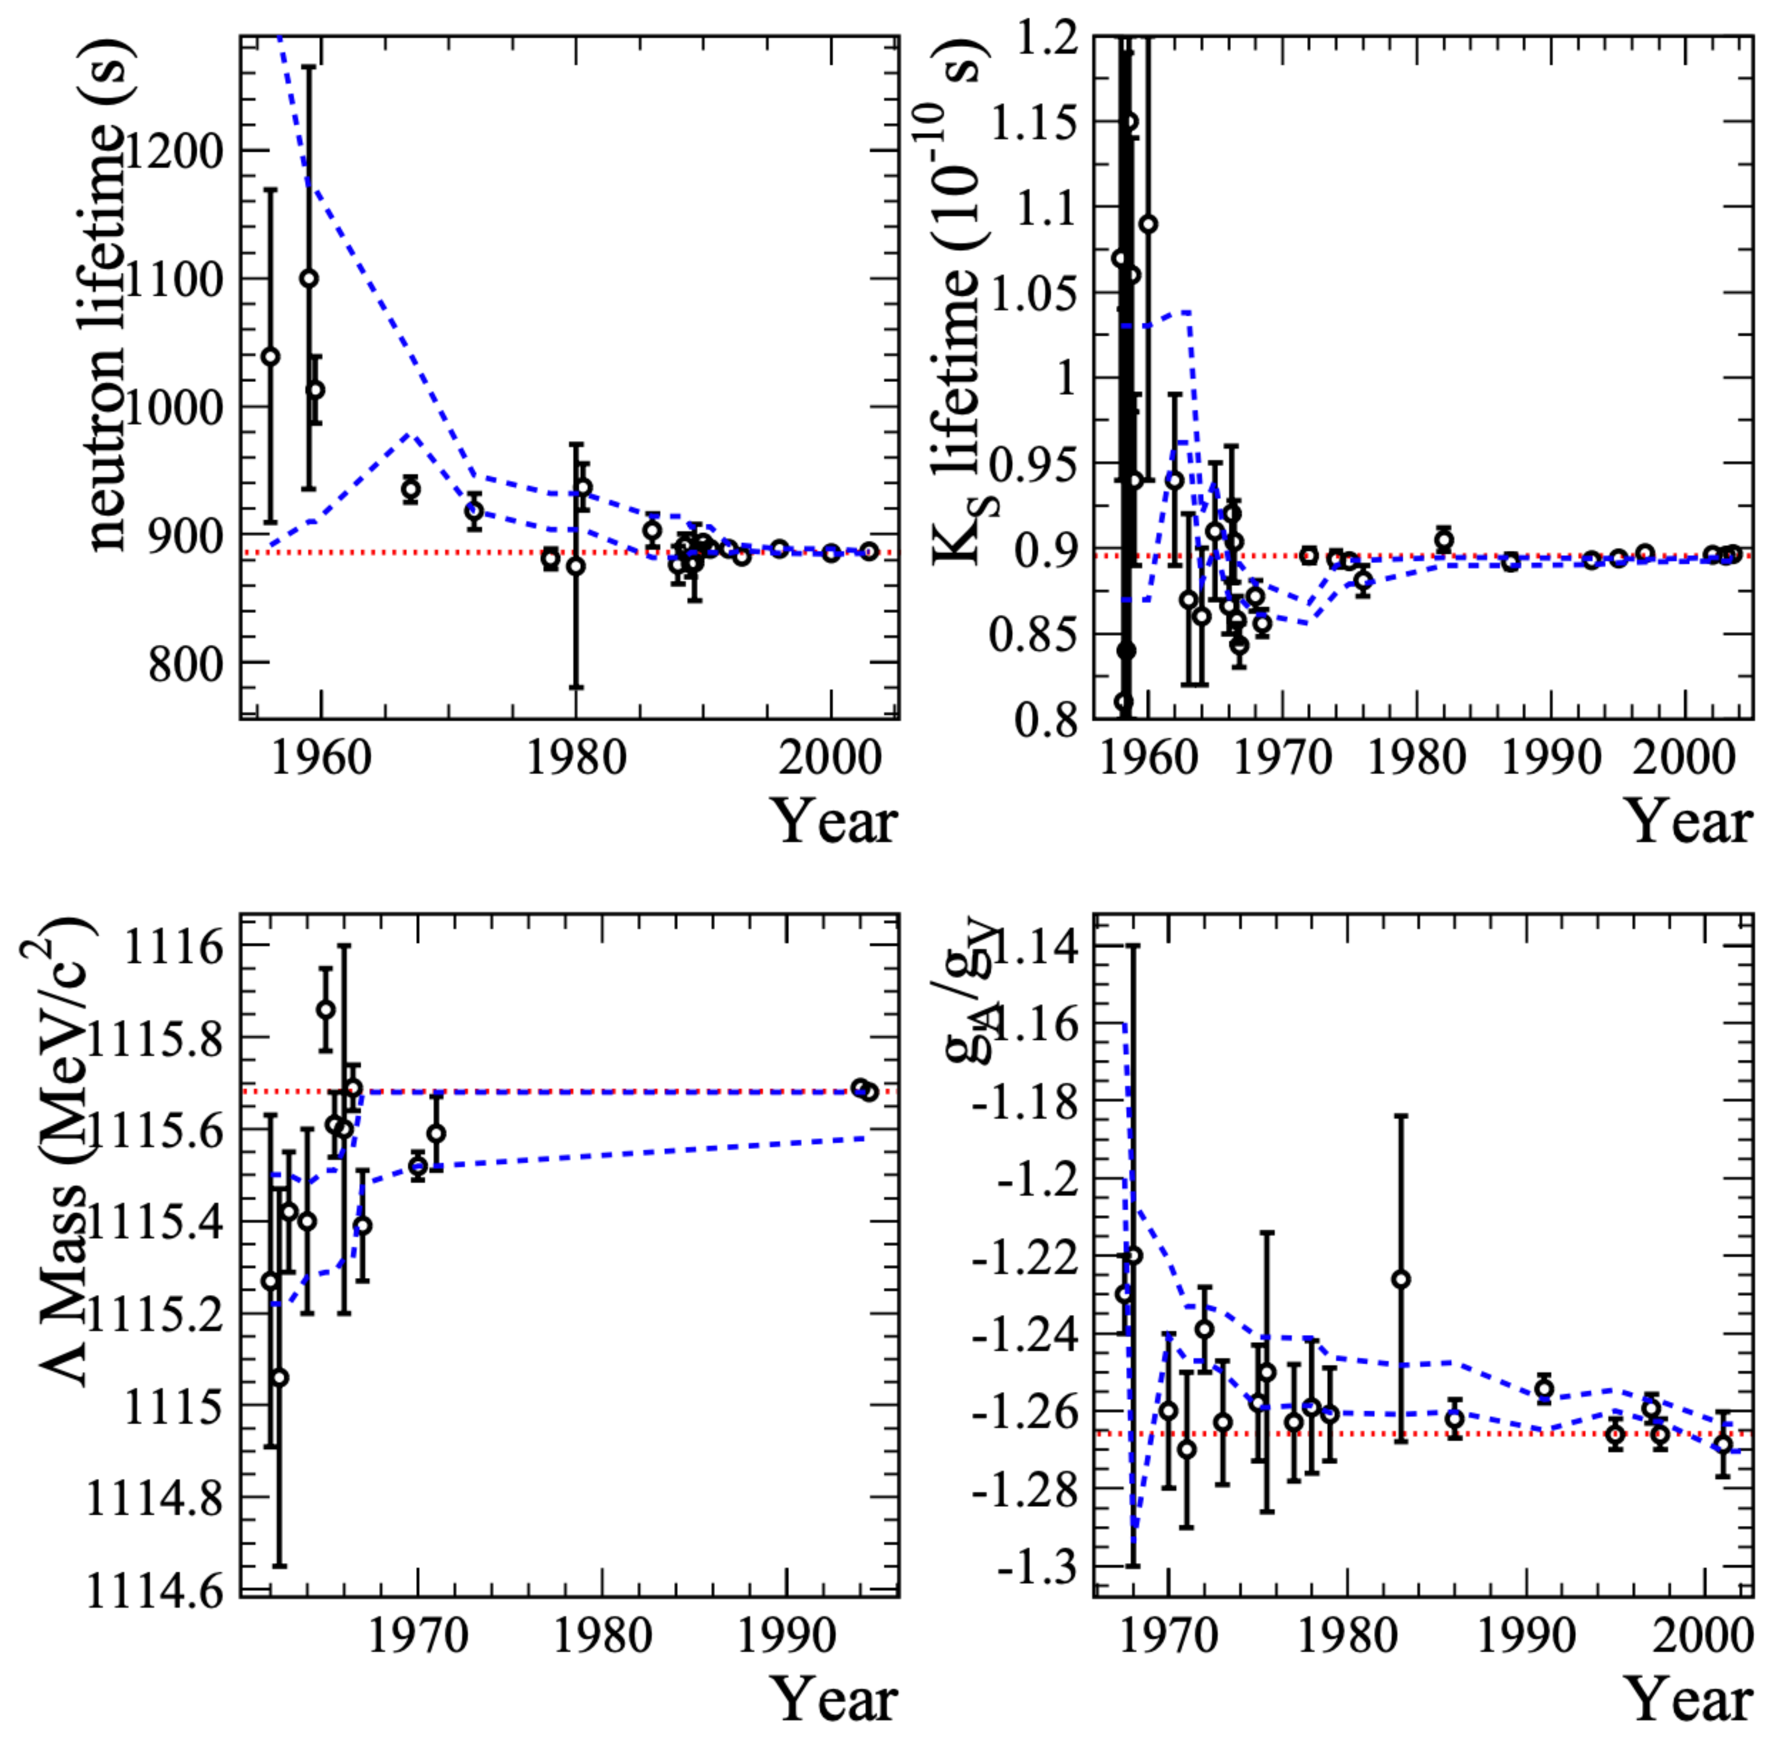

Figure 2: The history of four measurements compared to published averages before each measurement was made (dashed lines) and the currently accepted value (dotted lines).

From *Likelihood of Null Effects of Large NHLBI Clinical Trials Has Increased over Time* by Robert M. Kaplan and Veronica L. Irving

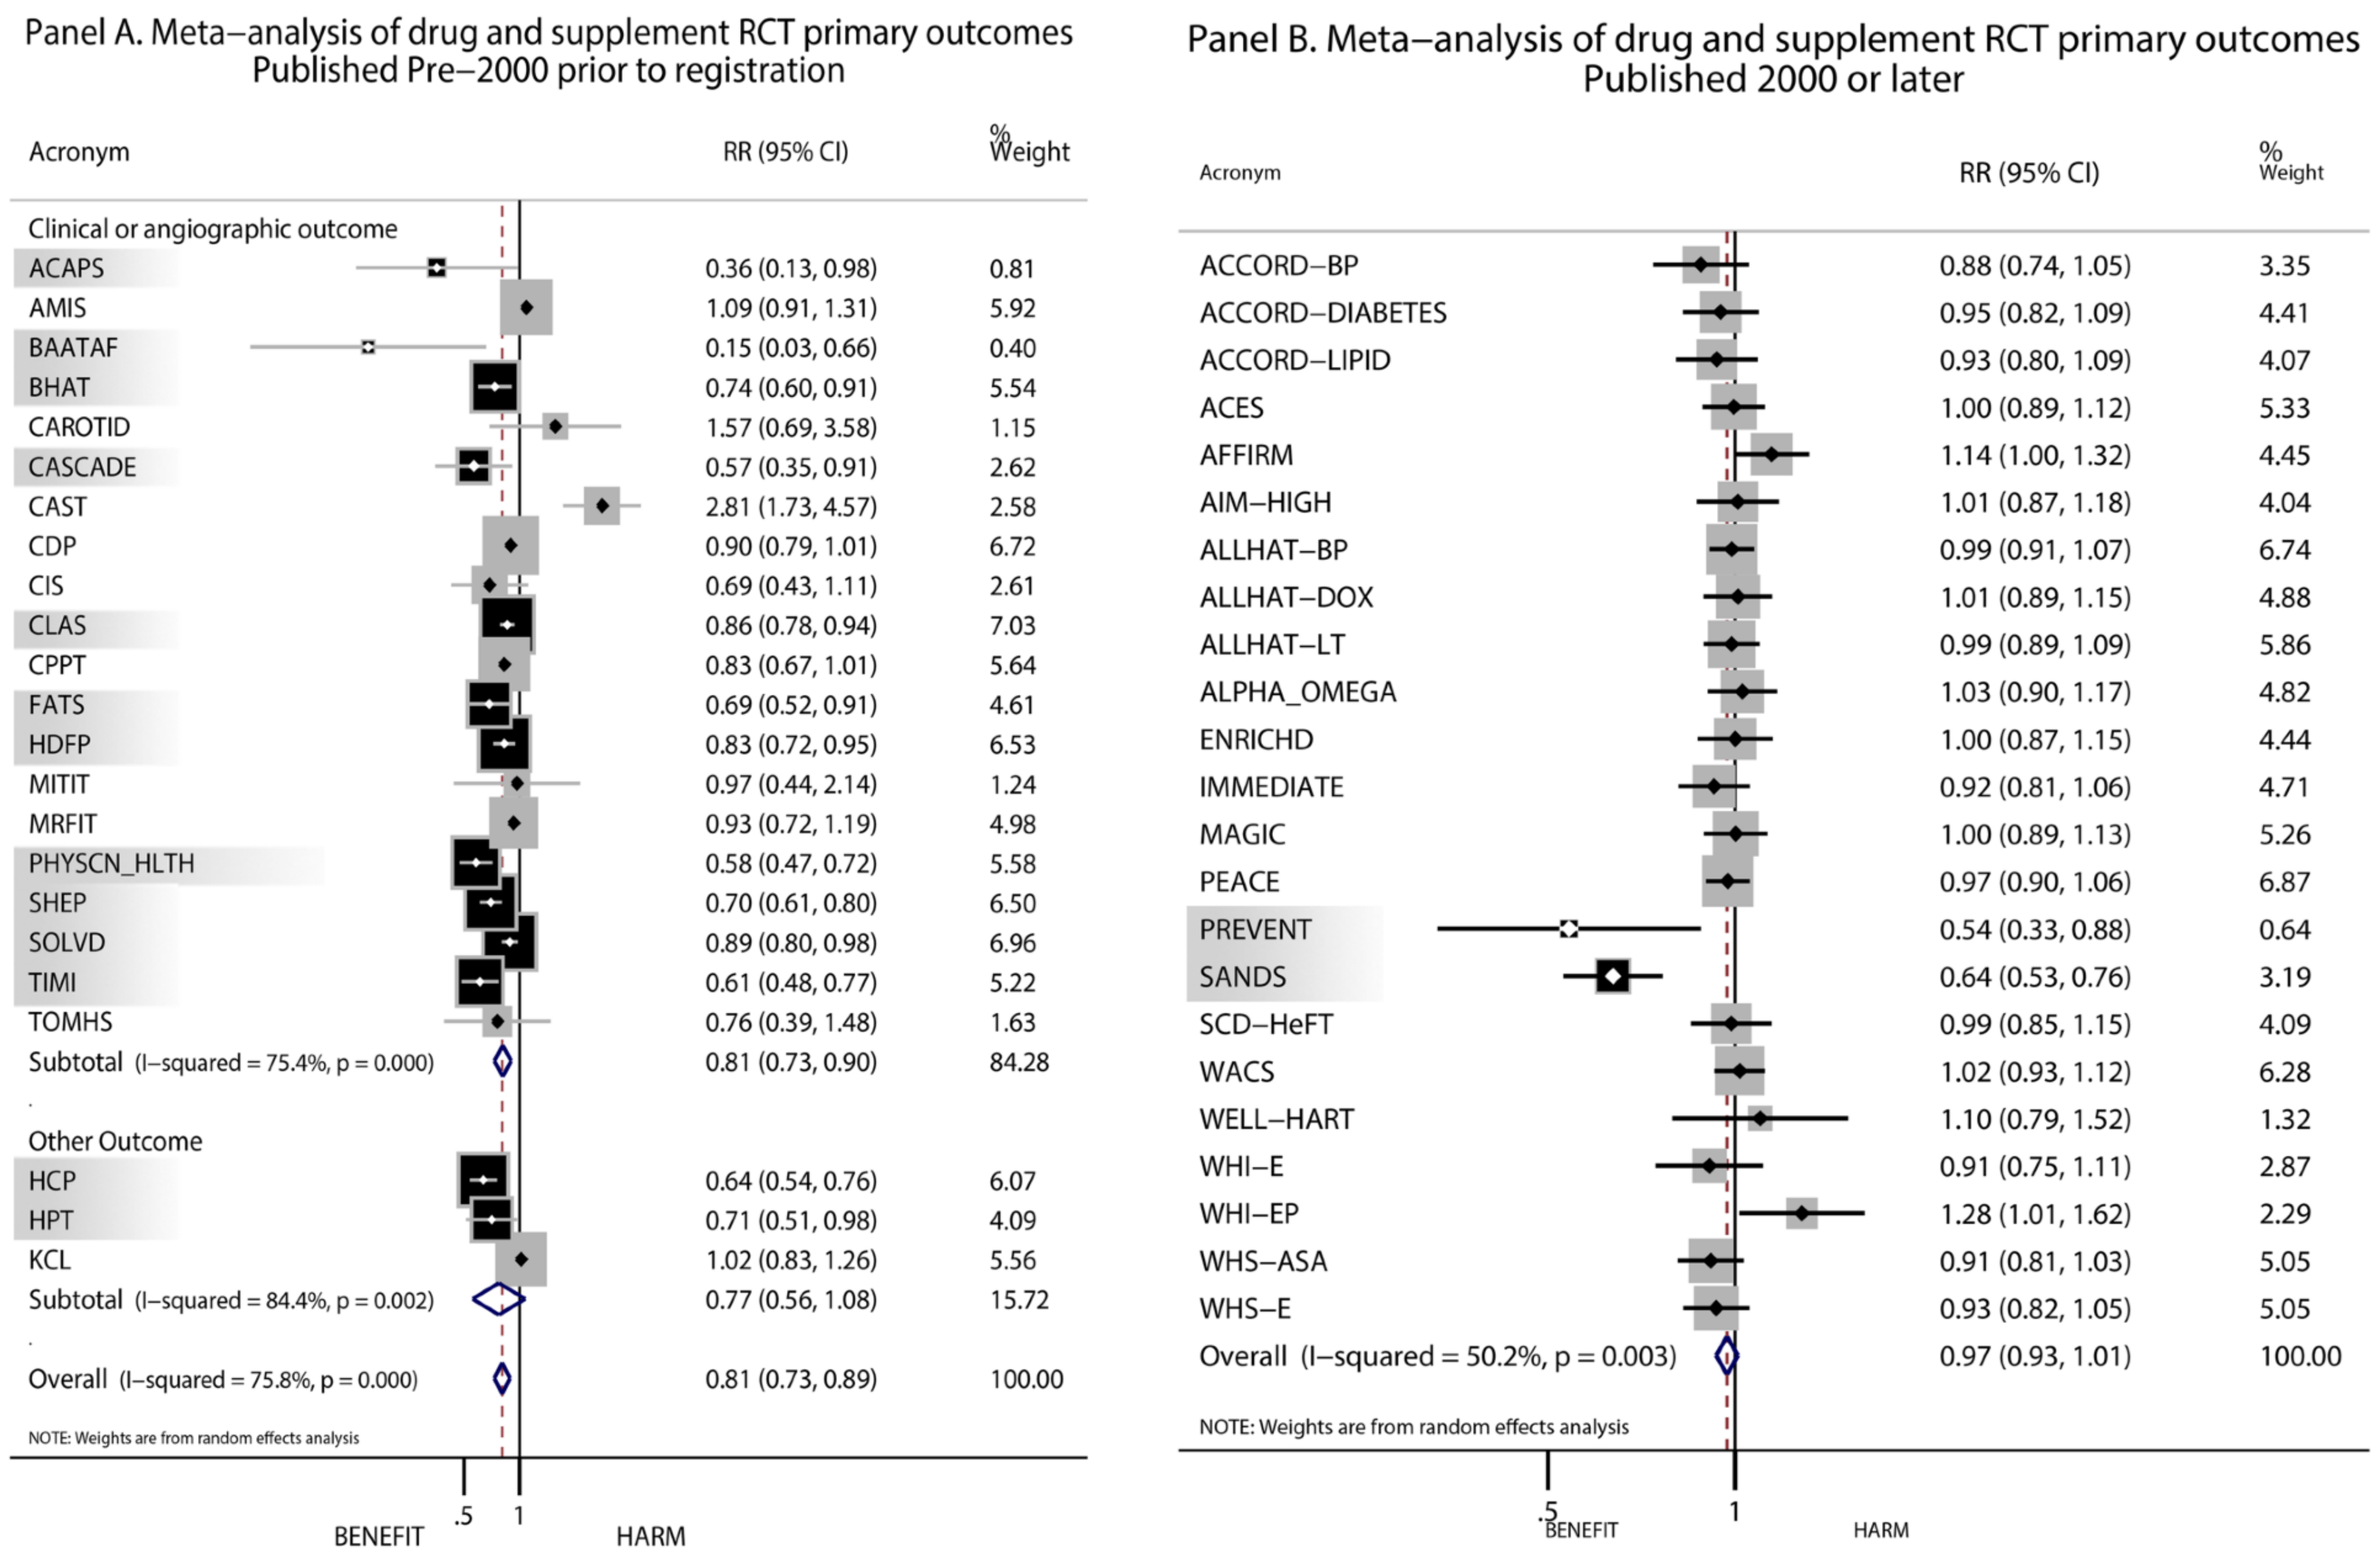

facebook
Artificial Intelligence Research

Accounting for Human Bias

Preregistration

Blind Analysis

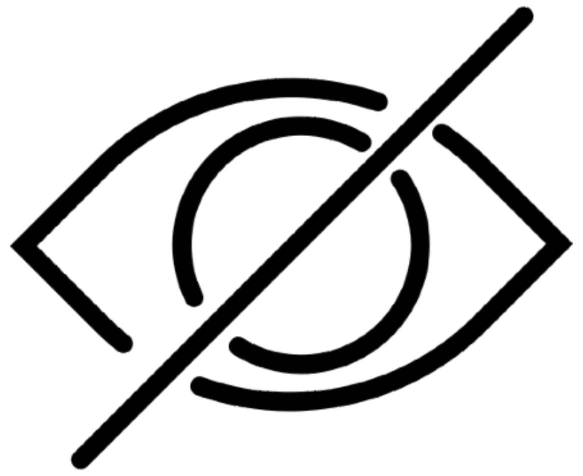

Other Solutions

Multi-group replication / competition, ...

Blind analysis and pre-registration

Don't judge a paper by its *p*-value.

cos.io/prereg/

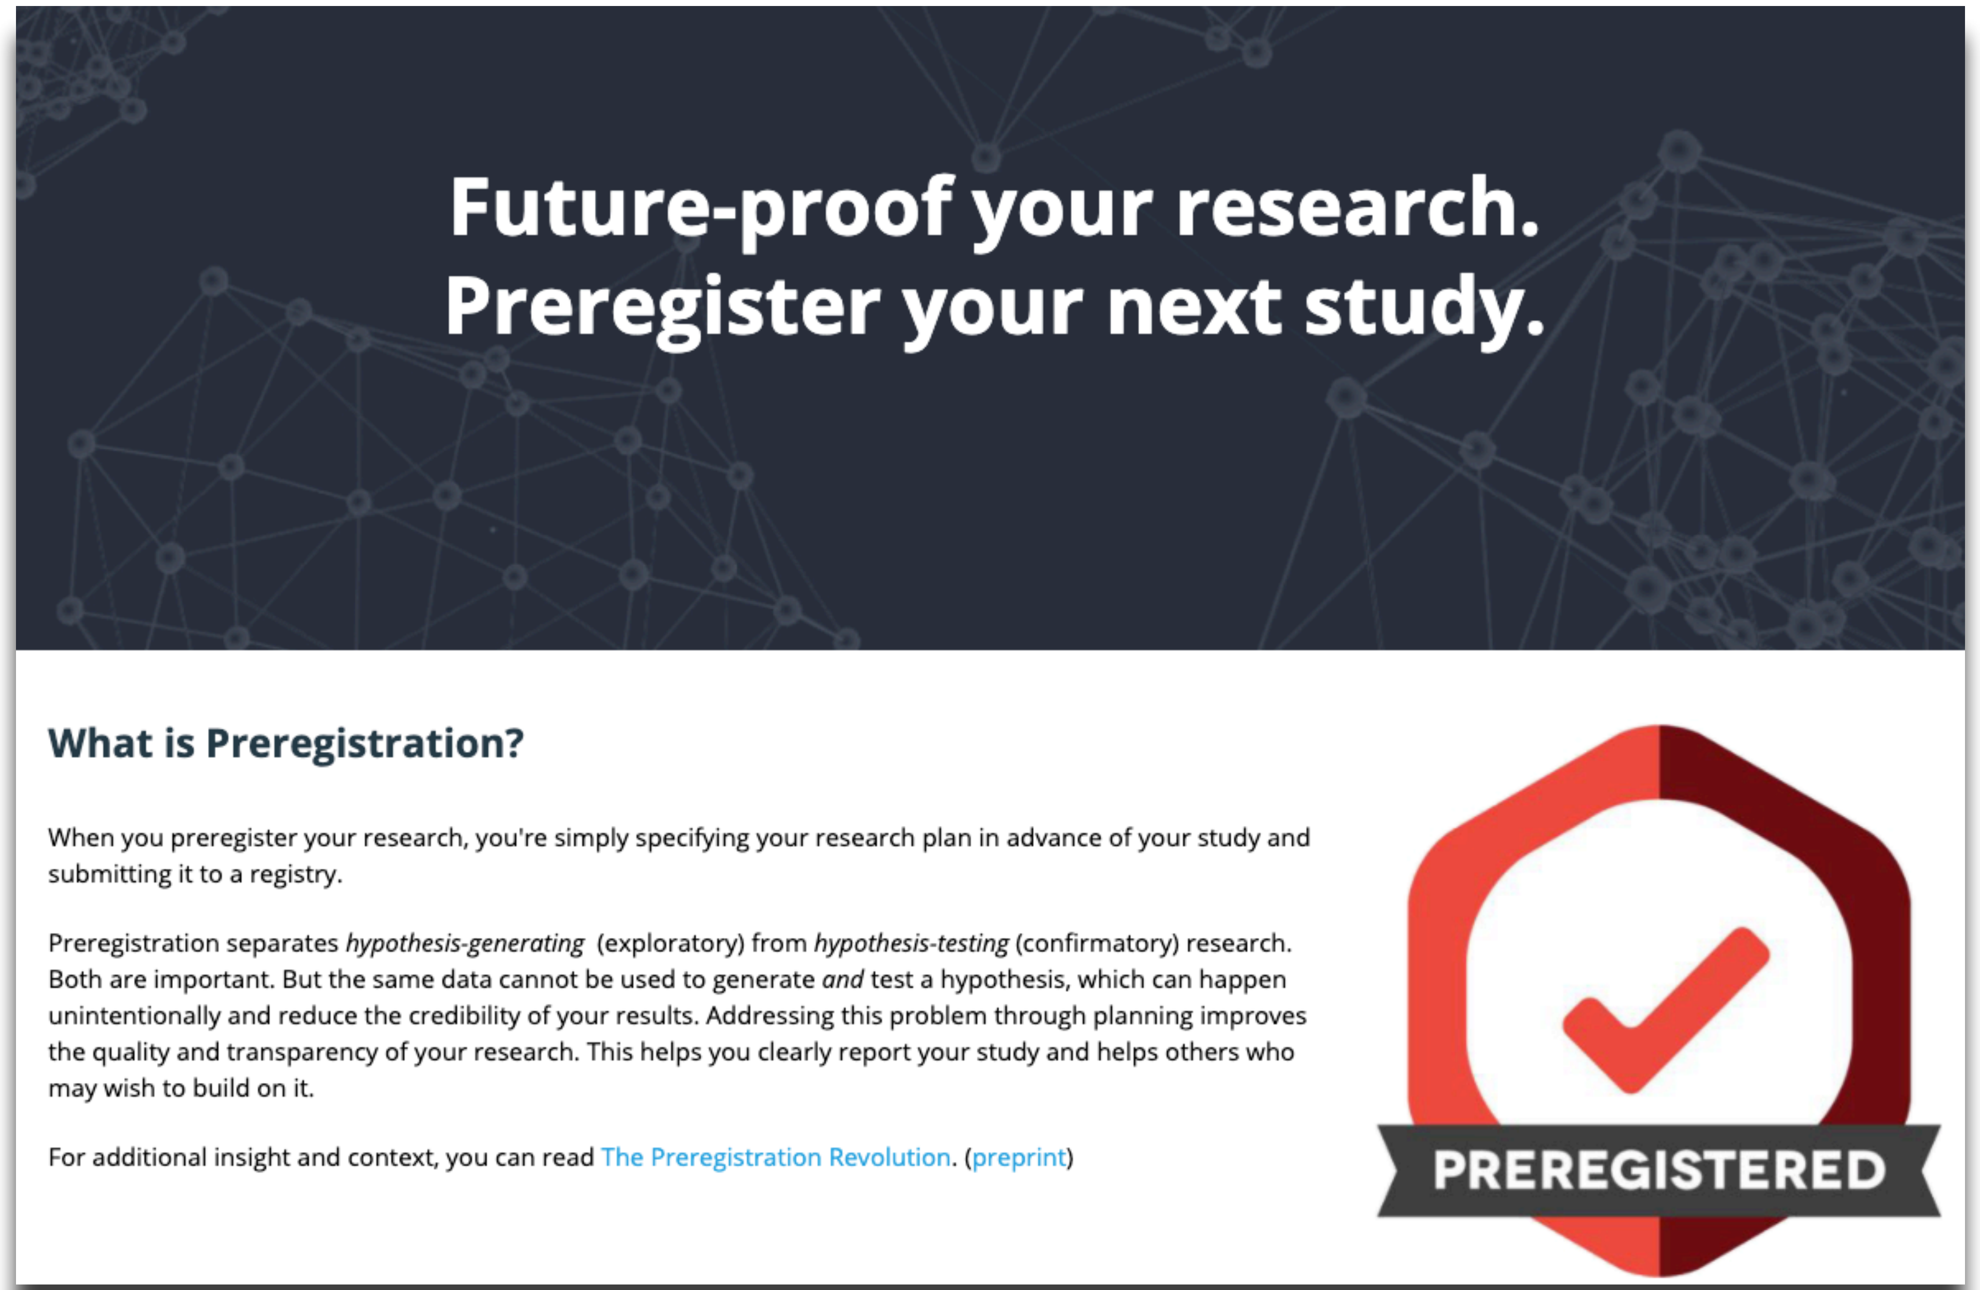

preregister.science

Blind Analysis

Hide the final numerical result from the experimenter until all analysis methodologies and choices have been frozen.

Before unblinding: state hypotheses, design analysis, debug all code, account for all uncertainties, finalize modeling choice, fix hyper-parameters, agree on all analysis decisions.

Analyze in the dark to prevent experimenter's bias, then remove the cloak of invisibility.

"list and schedule all the checks in advance of knowing the answer, and carry them out in either case (i.e. ask the question "What would we be checking if the answer were 8σ off?", write it down, and then do it in any case)"

Benefits of Blind Analysis Techniques, by Joel G. Heinrich

encrypt the analysis!

Blind Analysis Methods

How can we develop analyses in a blind manner?

- Add or remove unknown number of examples
- Develop on non-sensitive data, hide the signal region
- Mix real data with synthetic data in unknown proportion
- Rely only on simulated data
- Shuffle unknown fraction of labels
- Add unknown offset to results (random asymmetry when comparing methods)
- ...

Example: Blinding in High Energy Physics

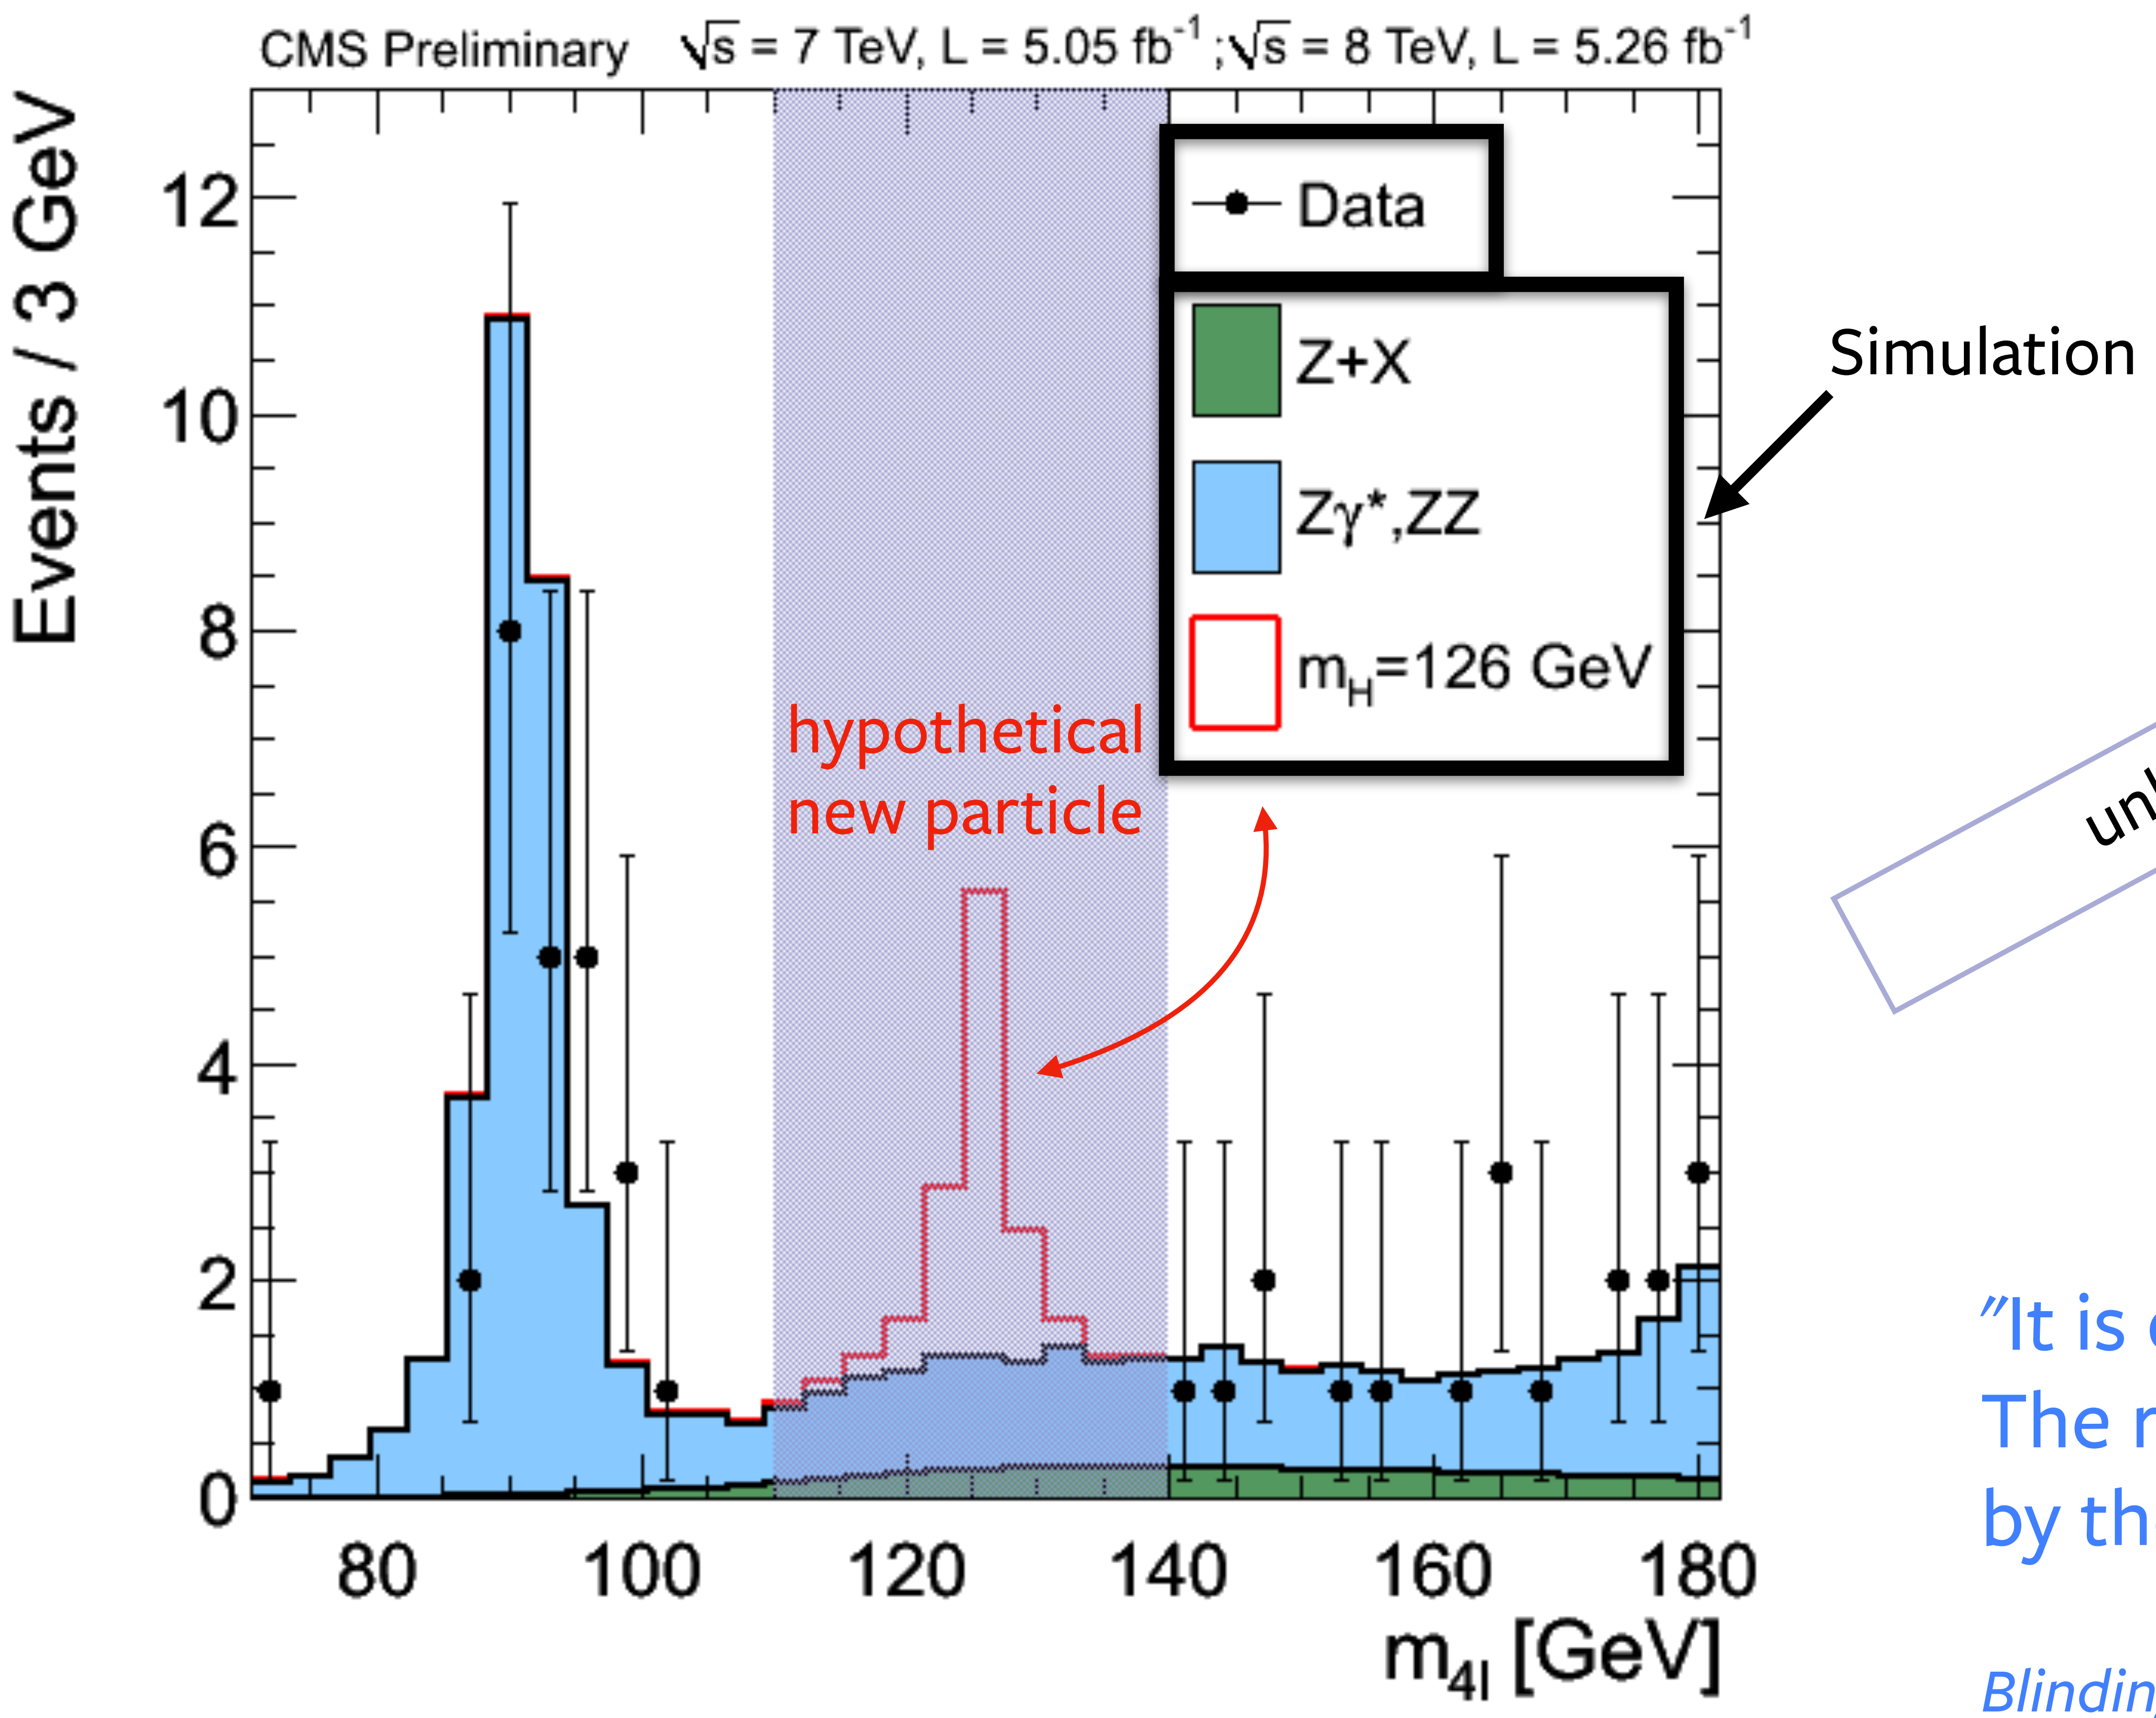

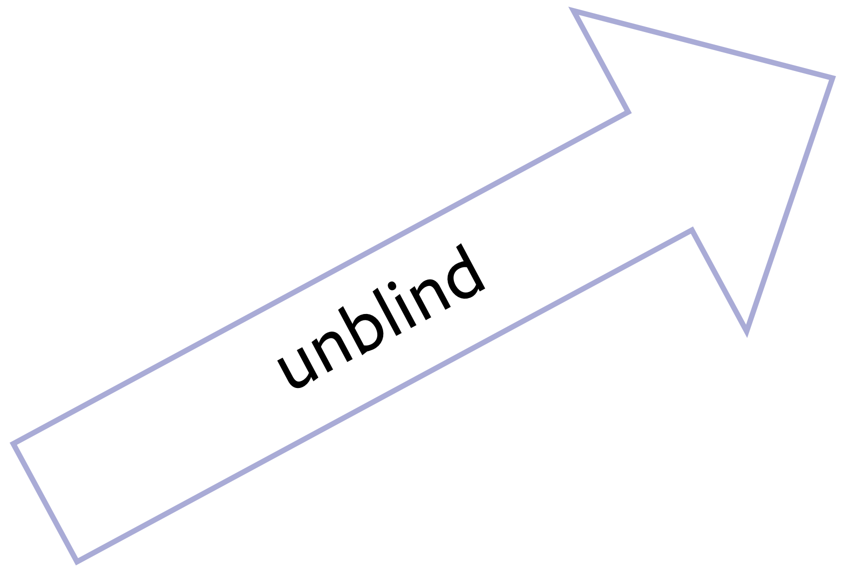

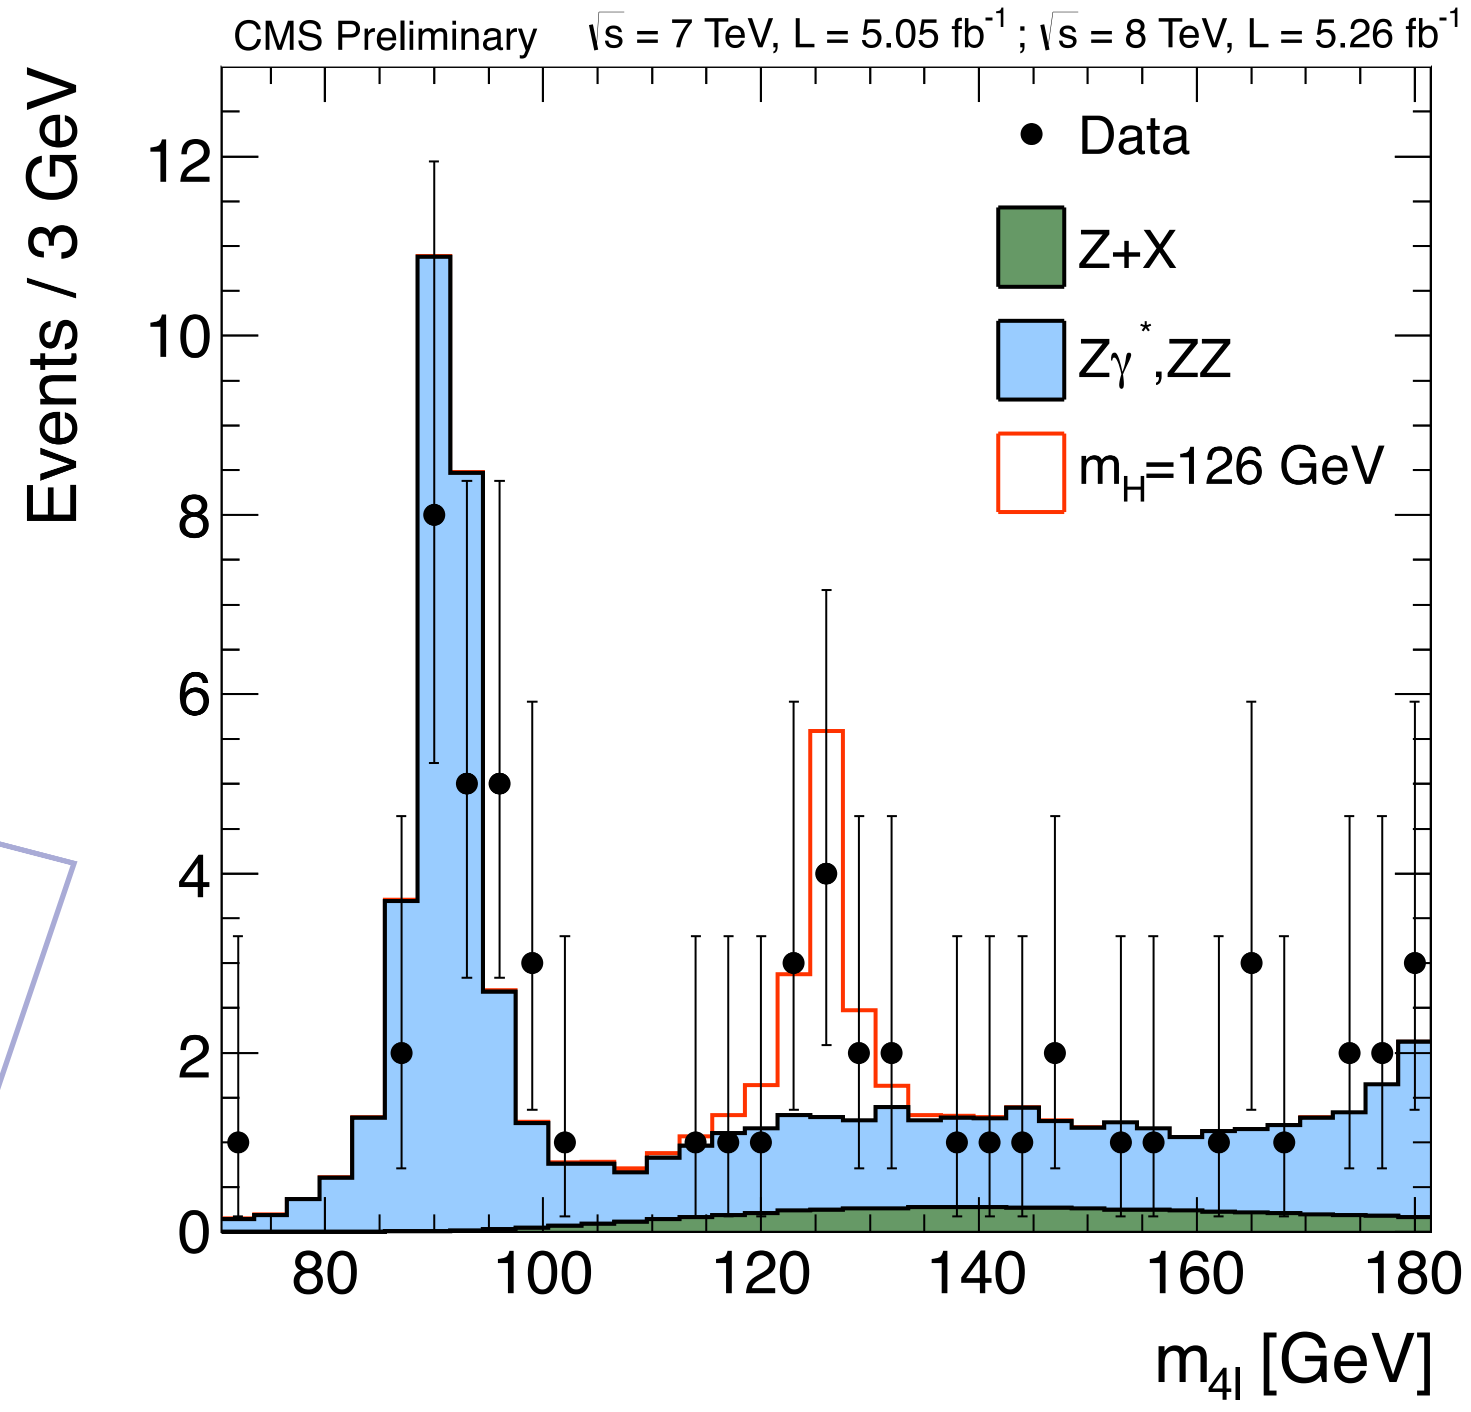

"It is only after the sign-off that the signal region is unblinded. The results of the unblinding are put through further scrutiny by the collaboration before being are made public."

Blinding and unblinding analyses, by Achintya Rao on behalf of CMS

The Pruning Case Study

Pruning

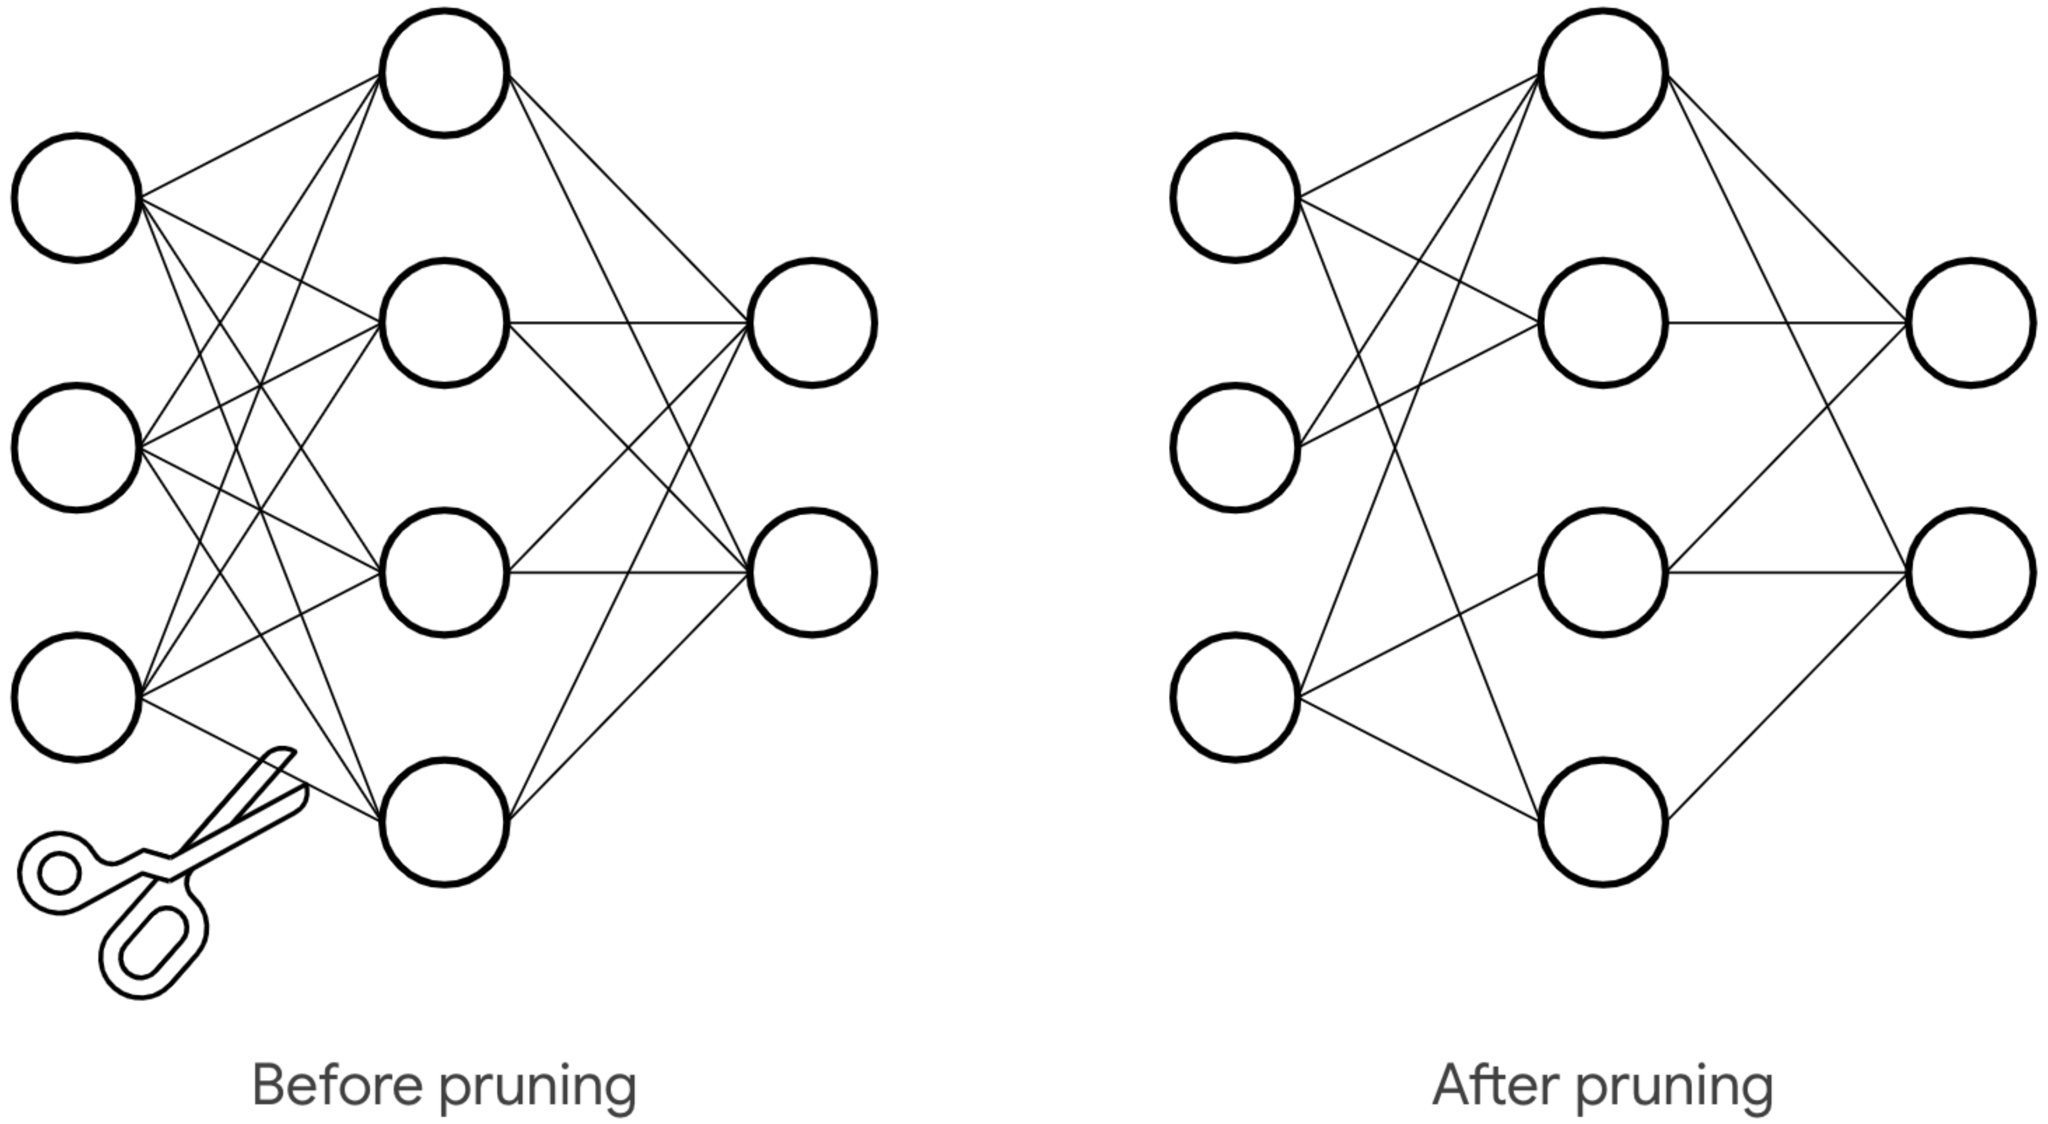

"removing superfluous structure"

how to identify?

what kind of structure?

The state of pruning

Pruning should remove unnecessary redundancy and unused capacity

Can be executed *before*, *during*, and *after* training

Pruning methods differ across many dimensions:

- based on weight magnitude, activations, gradients, Hessian, interpretability measures, credit assignment, random, etc.
- Layer-wise vs global, unstructured vs structured, etc.
- Rule-based, bayesian, differentiable, soft approaches, etc.
- One-shot vs iterative pruning
- Followed by: finetuning, reinitialization, rewinding

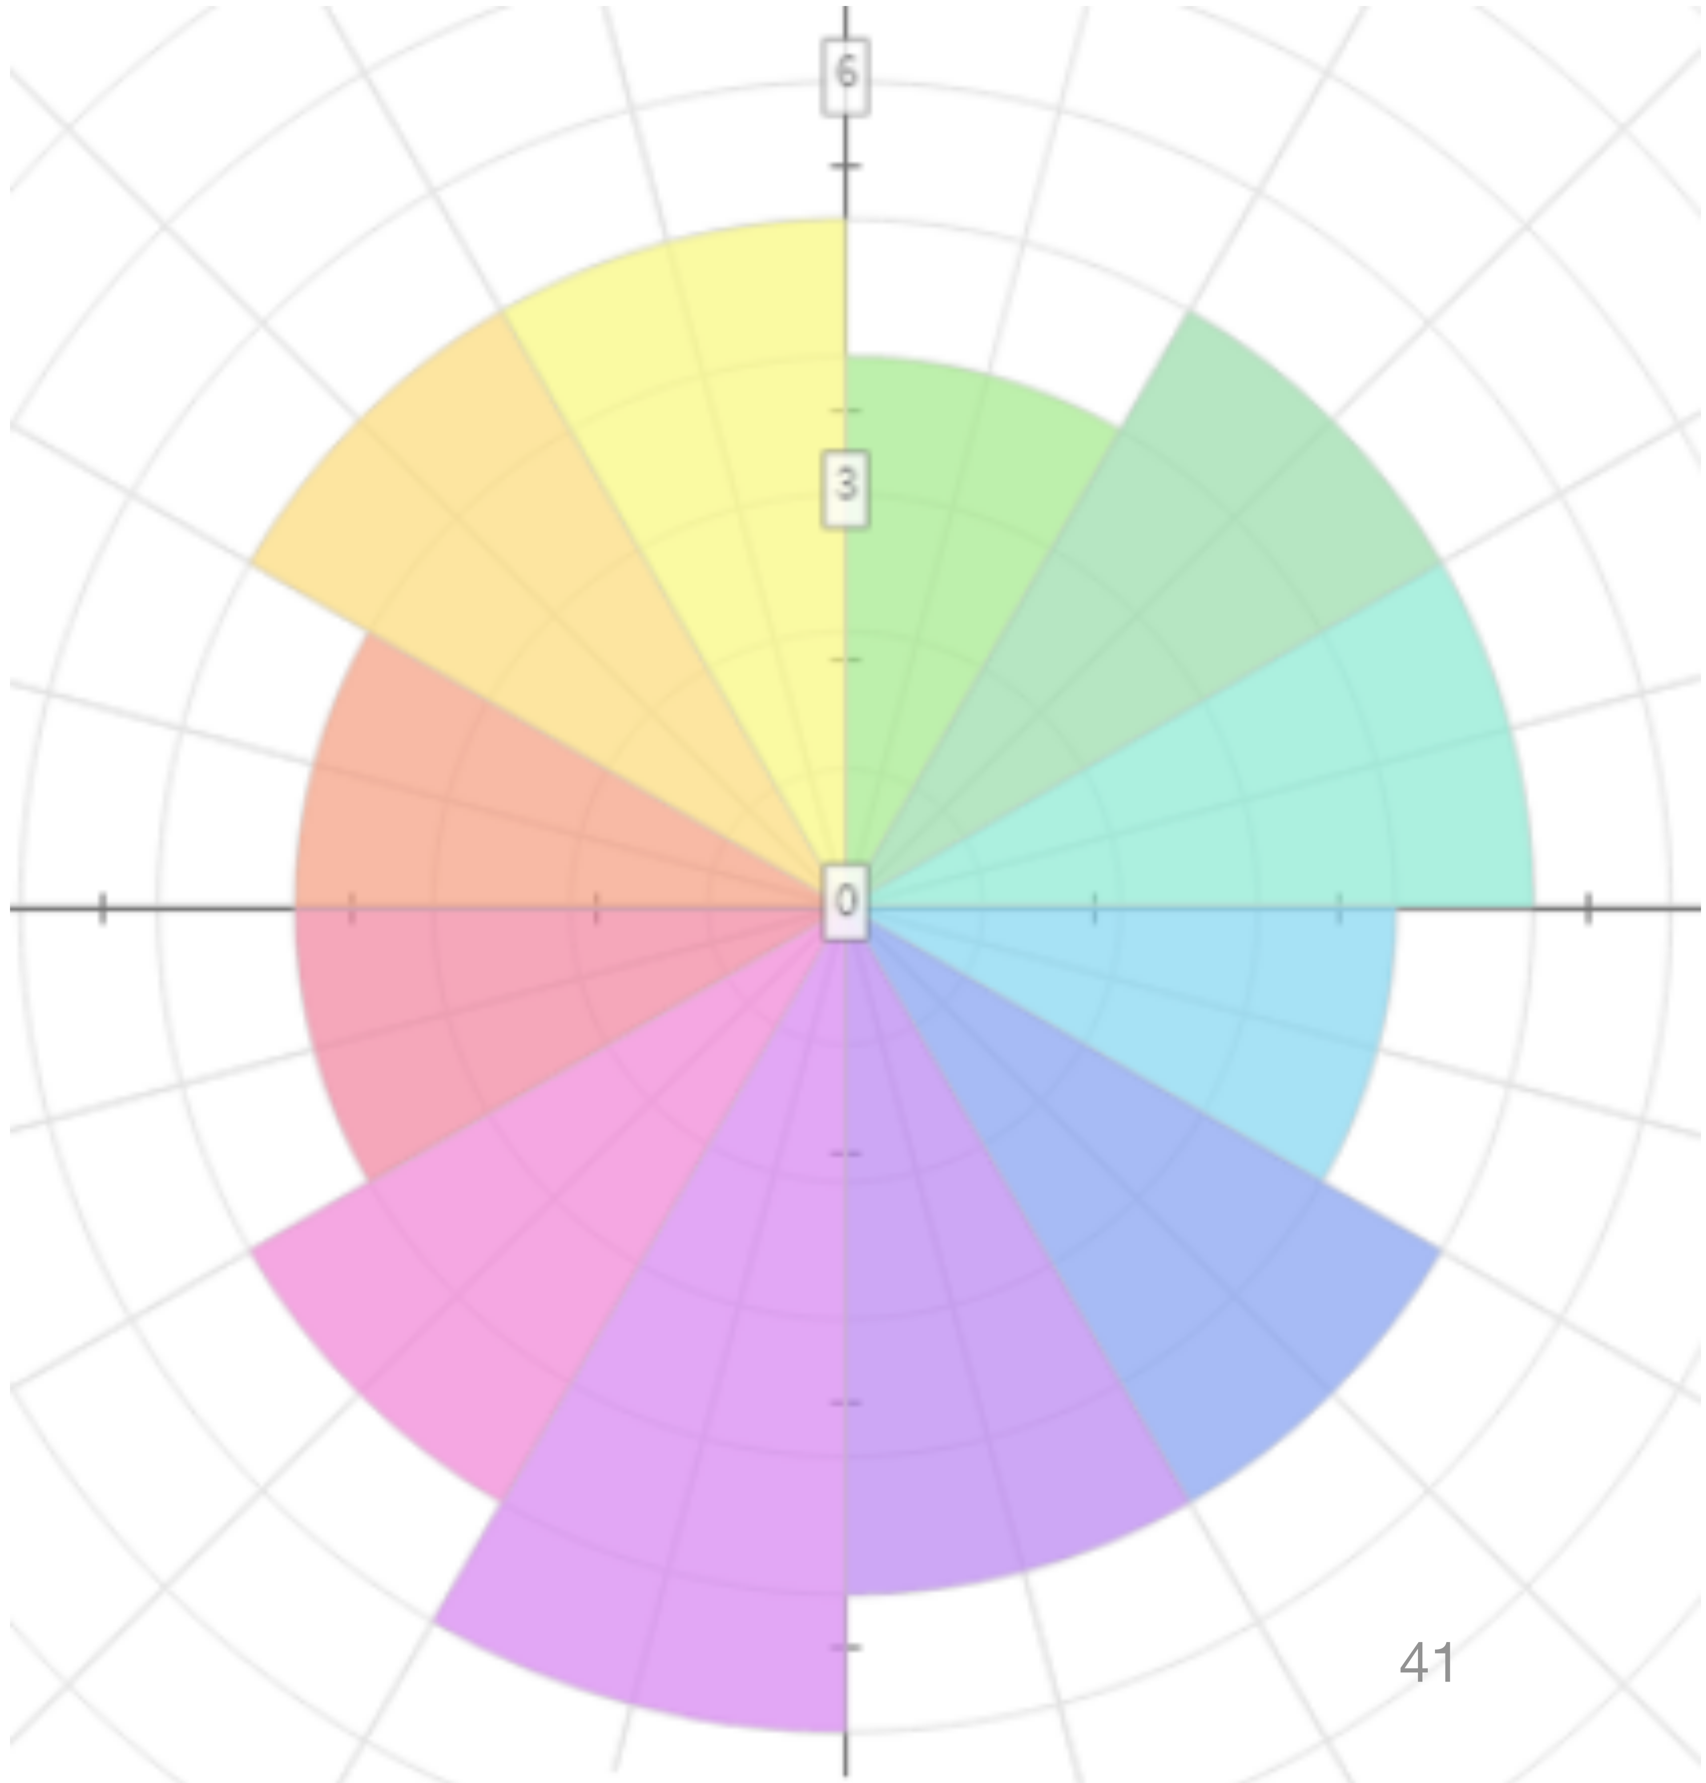

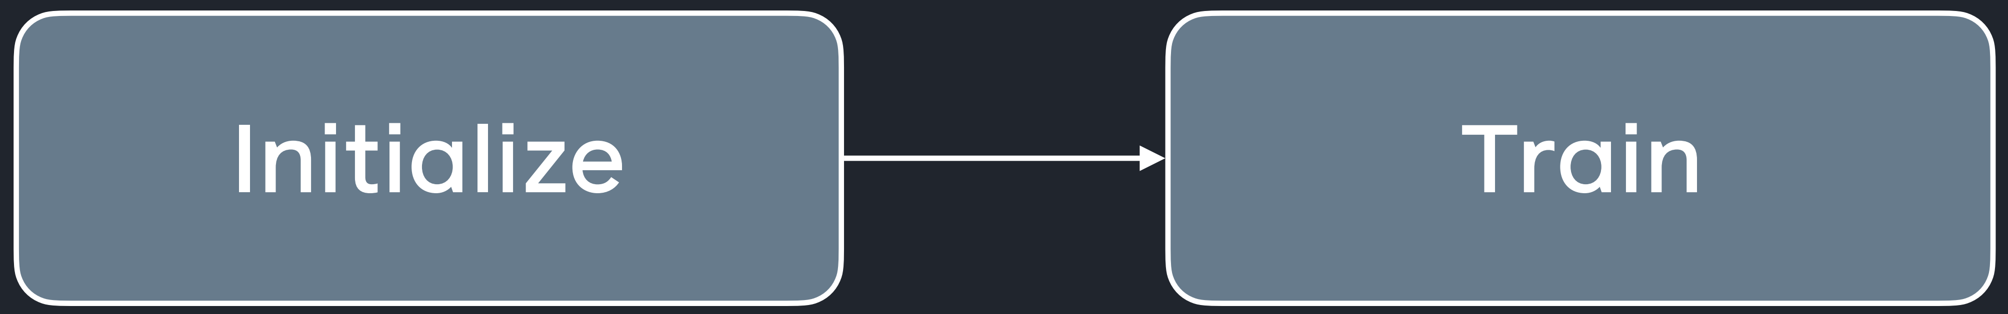
Initialize
Train


Initialize
Train
Prune

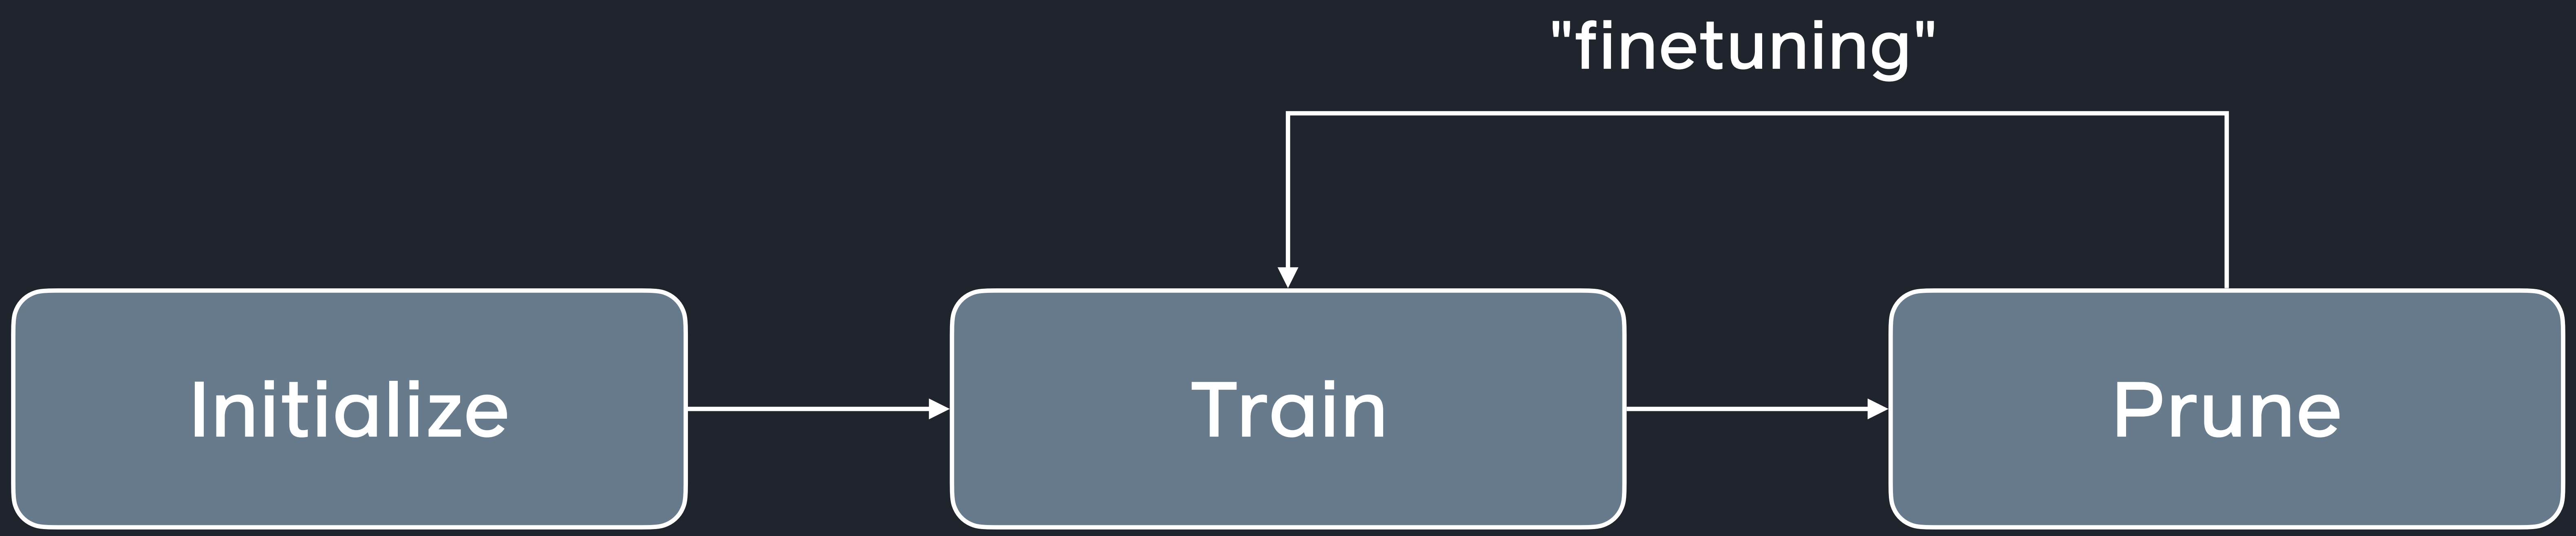
"finetuning"
Initialize
Train
Prune

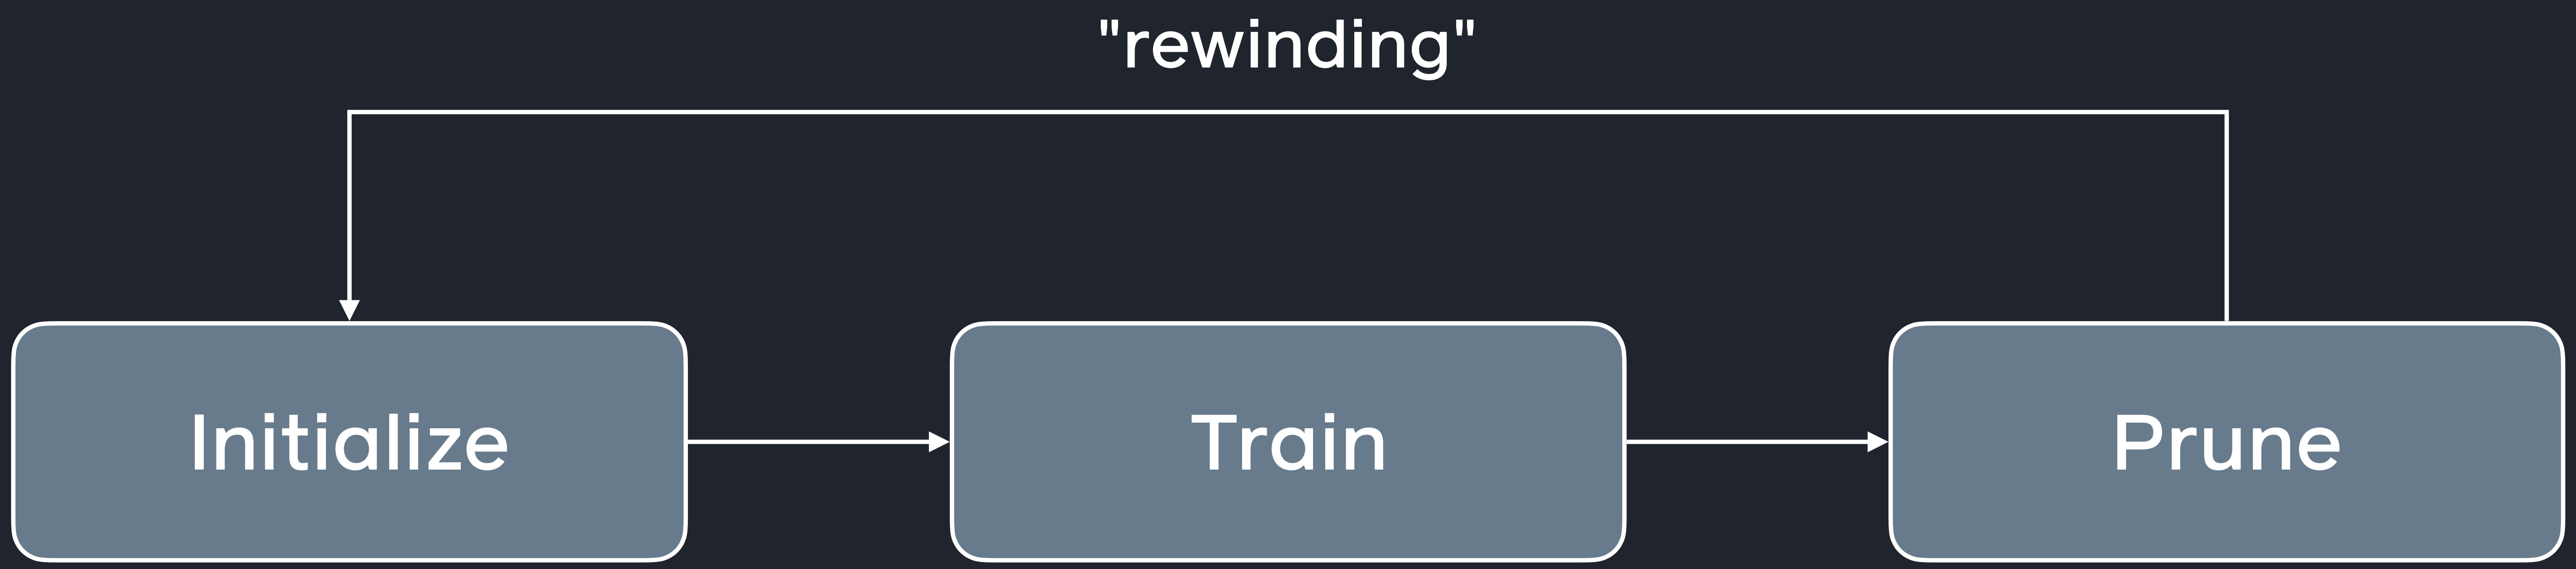
"rewinding"
Initialize
Train
Prune

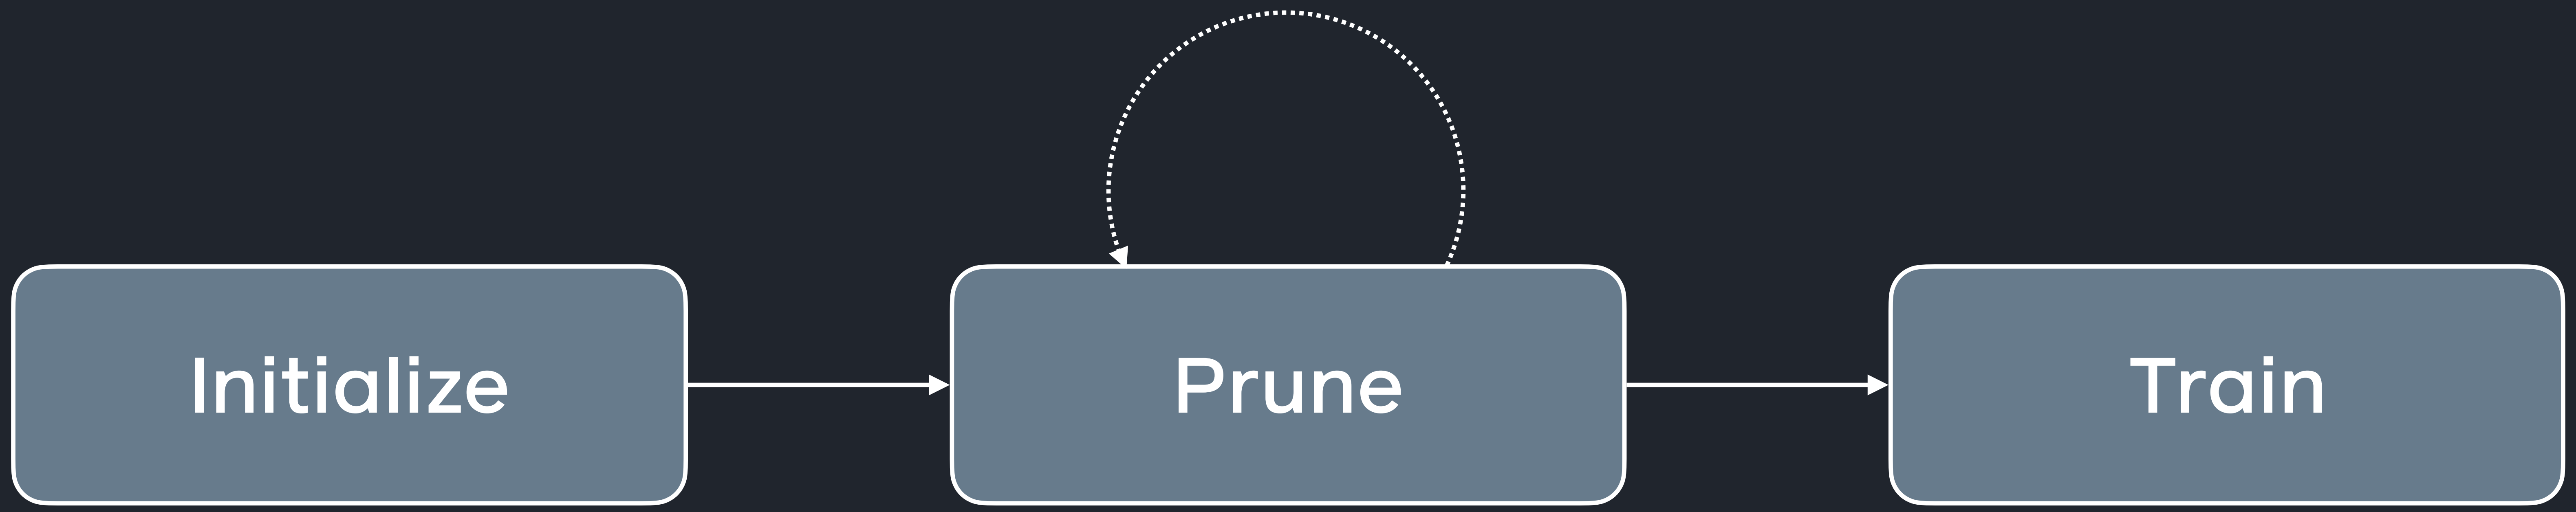
Initialize
Prune
Train

```
035bb6eaec977e3de60e0853d63cb011.bz2
035bb6eaec977e3de60e0853d63cb011.slim.json
049d06b3f12d4485d2c9555eb4f27b24.bz2
049d06b3f12d4485d2c9555eb4f27b24.slim.json
06cfe230e2b9565ab2d59e864eab7045.bz2
06cfe230e2b9565ab2d59e864eab7045.slim.json
08291cdd4c0f12de74166b02d935eabe.bz2
08291cdd4c0f12de74166b02d935eabe.slim.json
086ae9abf2282a3aa26e3f769f1ef509.bz2
086ae9abf2282a3aa26e3f769f1ef509.slim.json
0a74d268096a65c2a1779414aadba3df.bz2
0a74d268096a65c2a1779414aadba3df.slim.json
20d92fb2765e4edbd8a8f75adbb58696.bz2
20d92fb2765e4edbd8a8f75adbb58696.slim[illegible]
29ef5352b5d4b460bdaeb15c4c0f6c8c.bz2
29ef5352b5d4b460bdaeb15c4c0f6c8c.slim.json
30195510da3f77e3a12efcfc1d41b63a.bz2
30195510da3f77e3a12efcfc1d41b63a.slim.json
32537b8ea57608d5bf18d82b5b62d078.bz2
32537b8ea57608d5bf18d82b5b62d078.slim.json
344ae86e68def0bfb43bd3cfb8fbec13.bz2
344ae86e68def0bfb43bd3cfb8fbec13.slim.json
373ebdfa64468e1b3a23293e4d04680f.bz2
373ebdfa64468e1b3a23293e4d04680f.slim.json
386885205bcba9494dbebf580afee8b2.bz2
386885205bcba9494dbebf580afee8b2.slim.json
38ef0199861b8a1a5d05556d7fed1e52.bz2
38ef0199861b8a1a5d05556d7fed1e52.slim.json
3b4e8afc175fd03729a830fd136c5684.bz2
3b4e8afc175fd03729a830fd136c5684.slim.json
3cb027ad63585f8eda23639226302f63.bz2
3cb027ad63585f8eda23639226302f63.slim.json
42fa6cdfb64714c43c3419381a8fb821.bz2
42fa6cdfb64714c43c3419381a8fb821.slim.json
```

```
4d2060ad068f9c8beb61c909ccedf1fb.bz2
4d2060ad068f9c8beb61c909ccedf1fb.slim.json
50394e6dc4e478c3a9db0ab660937987.bz2
50394e6dc4e478c3a9db0ab660937987.slim.json
59f8a498b7ead01b98767ba912c8bb2f.bz2
59f8a498b7ead01b98767ba912c8bb2f.slim.json
5a3771b0392256a2d5625d71ce663e2b.bz2
5a3771b0392256a2d5625d71ce663e2b.slim.json
5a3c5fbd365bcf3935d7fd89e2145582.bz2
5a3c5fbd365bcf3935d7fd89e2145582.slim.json
61b999ebb35c25c34033a2987c92be76.bz2
61b999ebb35c25c34033a2987c92be76.slim.json
[illegible]40a835c5a51[illegible]ed4bbc8e.bz[illegible]
[illegible]835c5[illegible]4bbc8e[illegible]
6[illegible]3345[illegible]cd28cac62[illegible]1c2[illegible]
65[illegible]3345f4e5cd28cac62[illegible]1c2c.slim.json
[illegible]d079674dc4f327048[illegible]07a8a.bz2
[illegible]40d079674dc4f327[illegible]9a07a8a.slim.js[illegible]
69dc21cf4ea444271[illegible]11f1029917.bz2
69dc21cf4ea44427[illegible]a511f1029917.slim[illegible]
6b109f29bc047f37[illegible]7bed1291e72f.bz2
6b109f29bc047f37b1417bed1291e72f.slim.json
6b6aba4fececaba[illegible]dd06e11ca61e.bz2
6b6aba4fececaba[illegible]dd06e11ca61e.slim[illegible]
76ff392ef7b2748ba3784472f4884bf0.bz2
76ff392ef7b2748ba3784472f4884bf0.slim.json
77bbee568a85d535408c0a3c8d945ed5.bz2
77bbee568a85d535408c0a3c8d945ed5.slim.json
79b8a91c6ed9a9c8b5ea6dea047d906a.bz2
79b8a91c6ed9a9c8b5ea6dea047d906a.slim.json
7a2d512e2e28c2beafc75214bf81a2cb.bz2
7a2d512e2e28c2beafc75214bf81a2cb.slim.json
8110e7ed753ac476cb080ec43cdd8291.bz2
8110e7ed753ac476cb080ec43cdd8291.slim.json
```

```
9fd11c13ff649e45082663585e705ce9.bz2
9fd11c13ff649e45082663585e705ce9.slim.json
b420996715b37da7749ce42badbada10.bz2
b420996715b37da7749ce42badbada10.slim.json
ba803121bff54079d3fadaa6ff434c53.bz2
ba803121bff54079d3fadaa6ff434c53.slim.json
bc2887dc4df8d4baab6952ec8ef26323.bz2
bc2887dc4df8d4baab6952ec8ef26323.slim.json
bcb8a78325b9290b49671ece38ad2885.bz2
bcb8a78325b9290b49671ece38ad2885.slim.json
c53d29419b6cfe7ff959b2b63ea782e1.bz2
c53d29419b6cfe7ff959b2b63ea782e1.slim.json
[illegible]80fd73ad23cf5790d5170b0087af8fd.bz2
[illegible]d73ad23cf5790d5170b0087af8fd.slim.json
[illegible]f414a5f1e8e11322c4643c8b9c4.bz2
[illegible]f414a5f1e8e11322c4643c8b9c4.slim.json
[illegible]cb7aaca40af5a918fd9cae80110.bz2
[illegible]98cb7aaca40af5a918fd9cae80110.slim.json
cb79aa59933f51df2f577b1bb7fbd429.bz2
cb79aa59933f51df2f577b1bb7fbd429.slim.json
cde95b2e20d54cf0666228ac61985182.bz2
cde95b2e20d54cf0666228ac61985182.slim.json
d5e50946d8d6fb49339b1b3abeca403a.bz2
d5e50946d8d6fb49339b1b3abeca403a.slim.json
dd41bc211d272cc6674ed6662cf962d2.bz2
dd41bc211d272cc6674ed6662cf962d2.slim.json
dd4dc4556488cfaffa3e044d1300d1c5.bz2
dd4dc4556488cfaffa3e044d1300d1c5.slim.json
e090eac0b64924c34acd90adb7300d64.bz2
e090eac0b64924c34acd90adb7300d64.slim.json
e6ac41a441f1a1edaaeaf0f583e921d4.bz2
e6ac41a441f1a1edaaeaf0f583e921d4.slim.json
ec50fc8bd8523301dc7a1cc38e3133e1.bz2
ec50fc8bd8523301dc7a1cc38e3133e1.slim.json
```

???

Reproducible Experiment Orchestration

`facebookresearch/dagger`

`dagger` is a minimal framework for describing trees of network-mutating actions suited to the needs of researchers, allowing fast experimentation as well as maintenance of clear provenance in experiment evolution .

Goals:
- Allow researchers to abstract away *fundamental scientific contributions* from *experiment-tracking boilerplate code*
- *Bookkeeping:* track model state provenance

Concepts:
- **Experiment**: the graph
- **Experiment State**: a node
- **Recipe**: an edge

`dagger` is a minimal framework for describing trees of network-mutating actions suited to the needs of researchers, allowing fast experimentation as well as maintenance of clear provenance in experiment evolution .

Goals:
- Allow researchers to abstract away *fundamental scientific contributions* from *experiment-tracking boilerplate code*
- *Bookkeeping:* track model state provenance

Concepts:
- **Experiment**: the graph
- **Experiment State**: a node
- **Recipe**: an edge

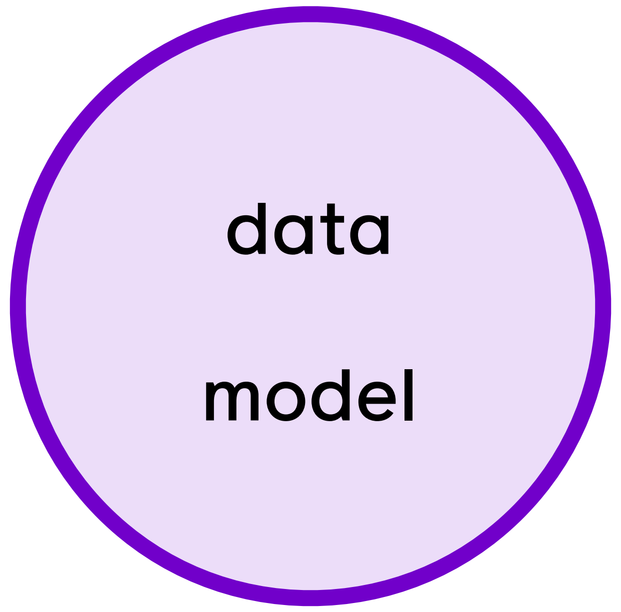

Experiment State A

`dagger` is a minimal framework for describing trees of network-mutating actions suited to the needs of researchers, allowing fast experimentation as well as maintenance of clear provenance in experiment evolution .

Goals:
- Allow researchers to abstract away *fundamental scientific contributions* from *experiment-tracking boilerplate code*
- *Bookkeeping:* track model state provenance

Concepts:
- **Experiment**: the graph
- **Experiment State**: a node
- **Recipe**: an edge

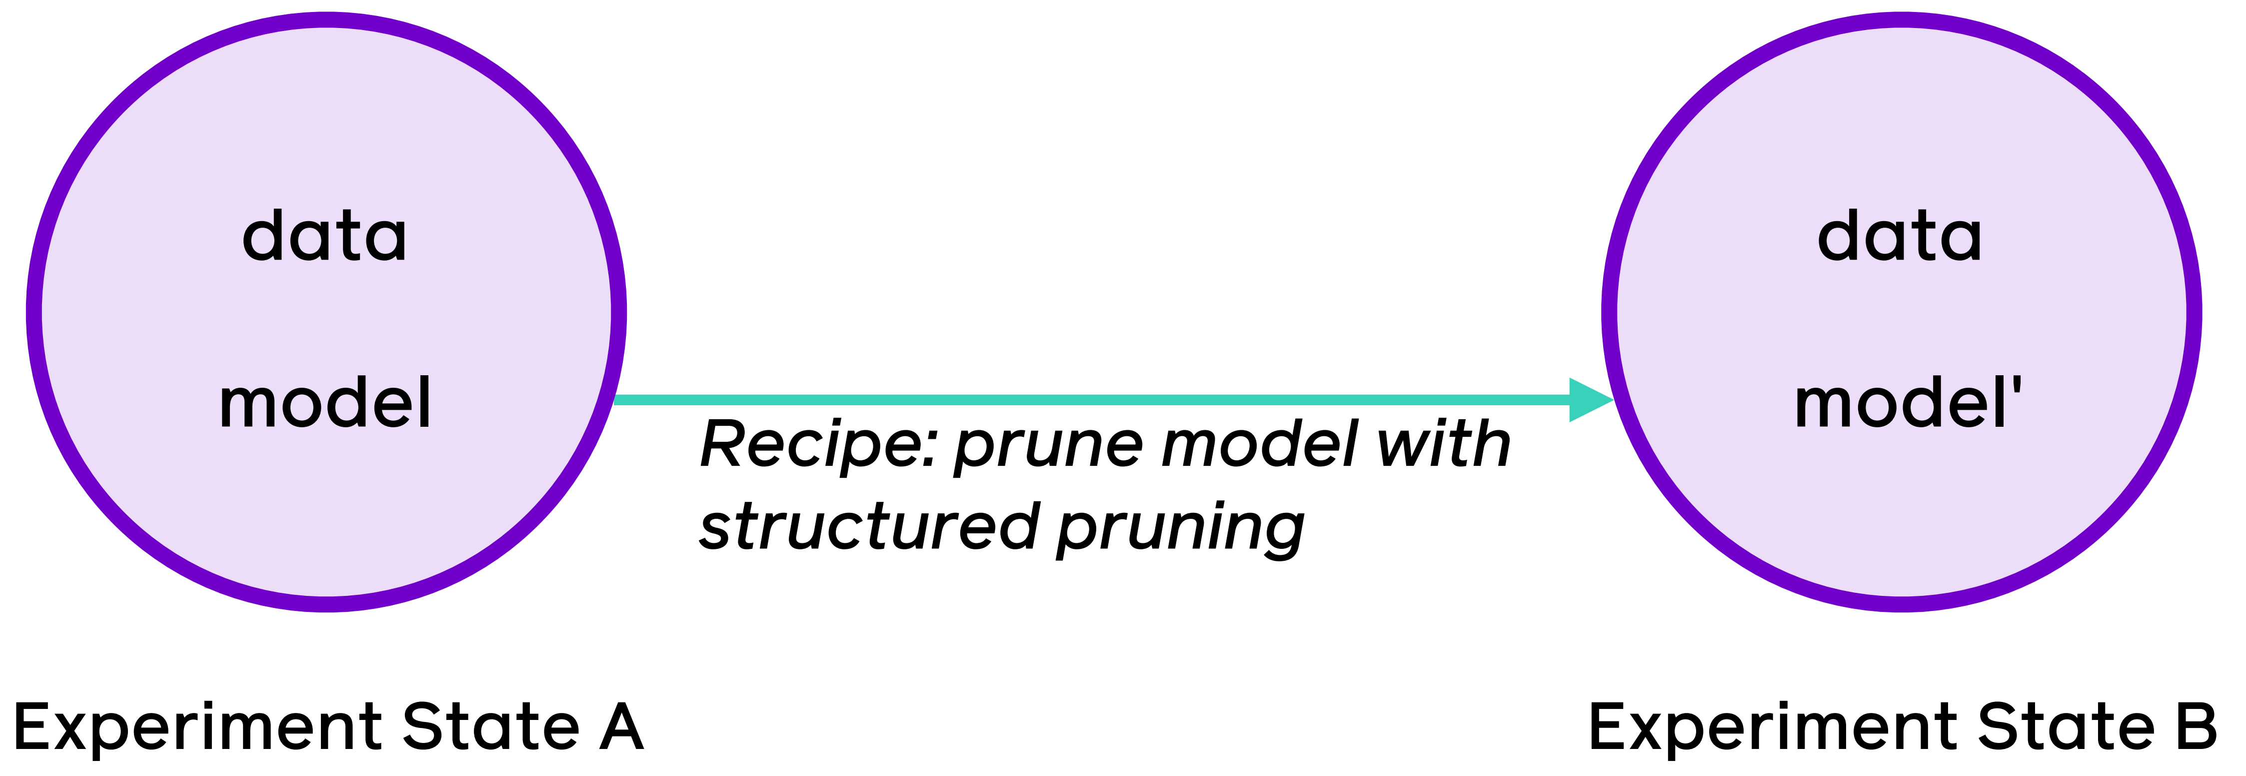

`dagger` is a minimal framework for describing trees of network-mutating actions suited to the needs of researchers, allowing fast experimentation as well as maintenance of clear provenance in experiment evolution .

Goals:
- Allow researchers to abstract away *fundamental scientific contributions* from *experiment-tracking boilerplate code*
- *Bookkeeping:* track model state provenance

Concepts:
- **Experiment**: the graph
- **Experiment State**: a node
- **Recipe**: an edge

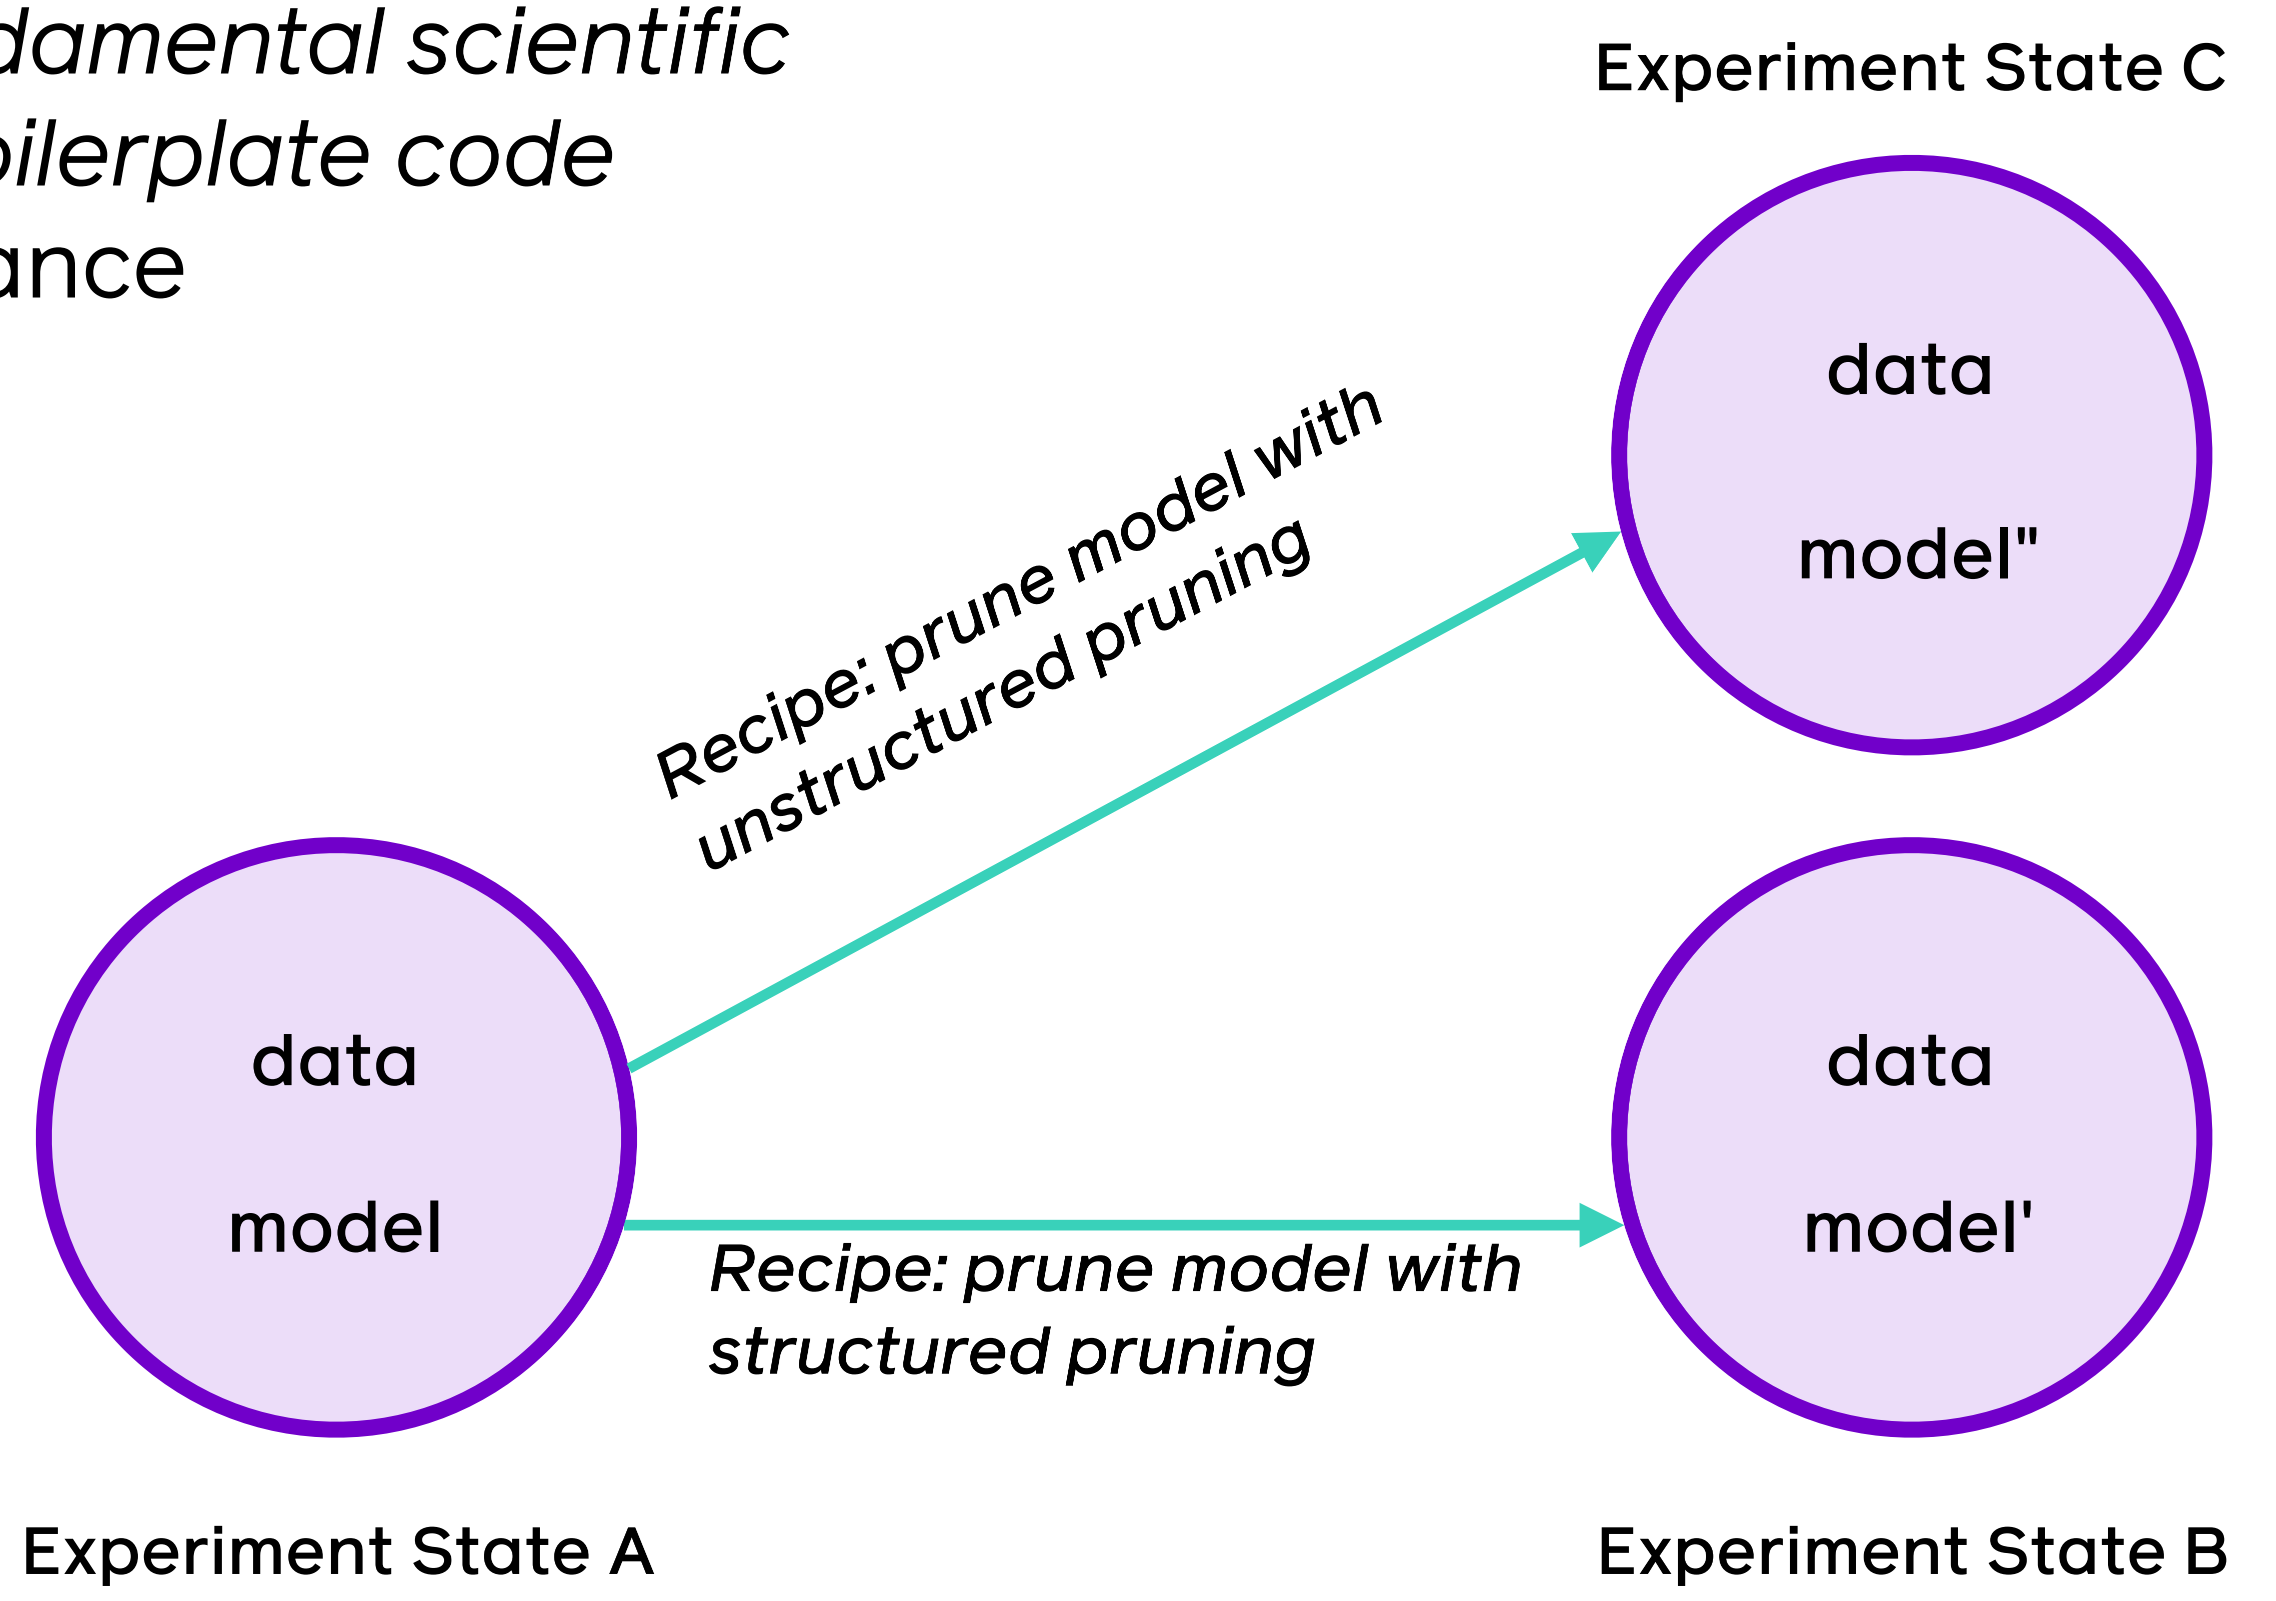

`dagger` is a minimal framework for describing trees of network-mutating actions suited to the needs of researchers, allowing fast experimentation as well as maintenance of clear provenance in experiment evolution .

Goals:
- Allow researchers to abstract away *fundamental scientific contributions* from *experiment-tracking boilerplate code*
- *Bookkeeping:* track model state provenance

Concepts:
- **Experiment**: the graph
- **Experiment State**: a node
- **Recipe**: an edge

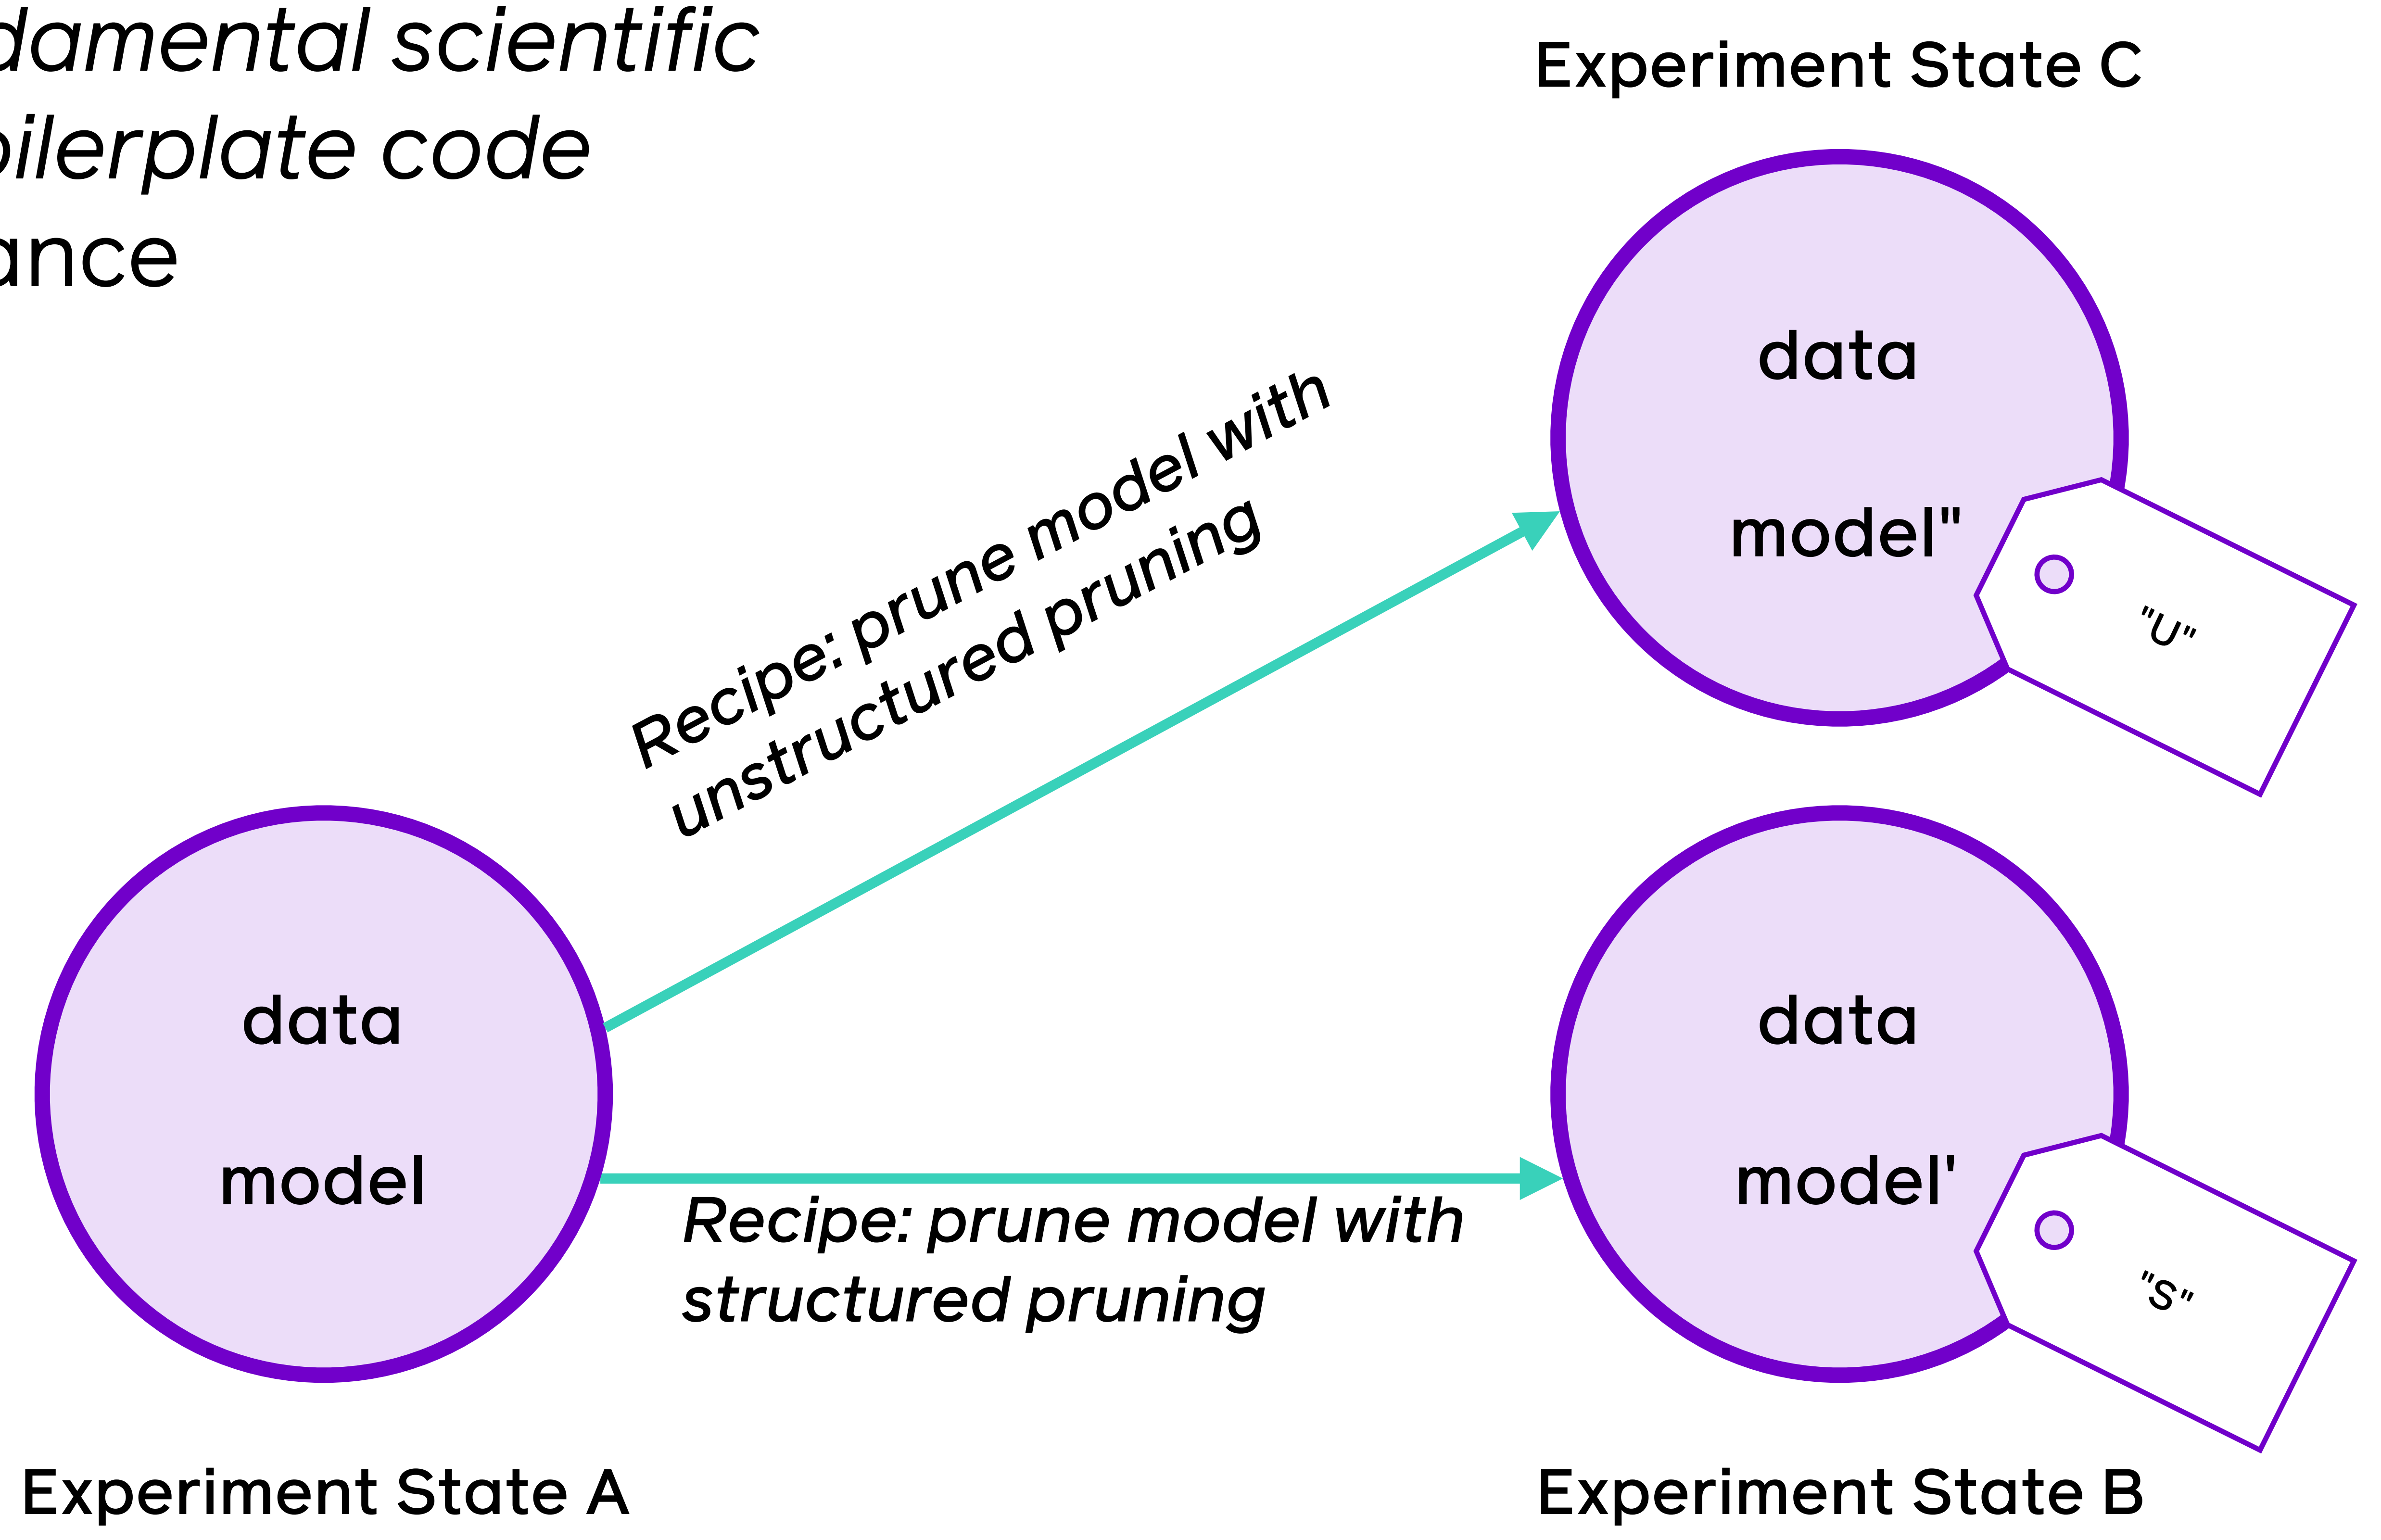

Experiment Loop

```python
exp = dg.Experiment("/path/to/experiment/folder", state_class=State)
root_state = exp.spawn_new_tree(dataset_name="cifar-10", model_name="vgg-11")

for lr in [0.01, 0.1]:
    train = TrainRecipe(nb_epochs=100, lr=lr)
    prune = PruneRecipe(pruning_technique="lowest_magnitude", pruning_fraction=0.2)
    s = root_state
    with exp.tag(f"lr:{lr}"):
        s = train(s)
        eval_fn(s)
        with exp.tag("pruned"):
            s = prune(s)
exp.run()
```

Custom Definitions

```python
import dagger as dg
from yourlib import get_data, get_model, train_model, prune_model, eval_model

class State(dg.ExperimentState):

    PROPERTIES = ["dataset_name", "model_name"]
    NONHASHED_ATTRIBUTES = ["train_data", "eval_data", "model"]

    def initialize_state(self, **kwargs):
        self.train_data, self.eval_data = get_data(self.dataset_name)
        self.model = get_model(self.model_name)

class TrainRecipe(dg.Recipe):

    PROPERTIES = ["nb_epochs", "lr"]

    def run(self, state):
        train_model(state.model, state.train_data, self.nb_epochs, self.lr)
        return state

class PruneRecipe(dg.Recipe):

    PROPERTIES = ["pruning_technique", "pruning_fraction"]

    def run(self, state):
        prune_model(state.model, self.pruning_technique, self.pruning_fraction)
        return state

@dg.function
def eval_fn(state):
    eval_acc = eval_model(state.model, state.eval_data)
    print(f"Experiment: {state.tags}, Accuracy: {eval_acc}")
```

Experiment Analysis

```python
>>> exp = Experiment.restore("/path/to/experiment/folder", slim=True)
>>> exp.graph.draw()  # Draws the graph in Figure 1
>>> s = exp.graph.nodes.filter("pruned") & exp.graph.nodes.filter("lr:0.1")
>>> s[0].restore()
```

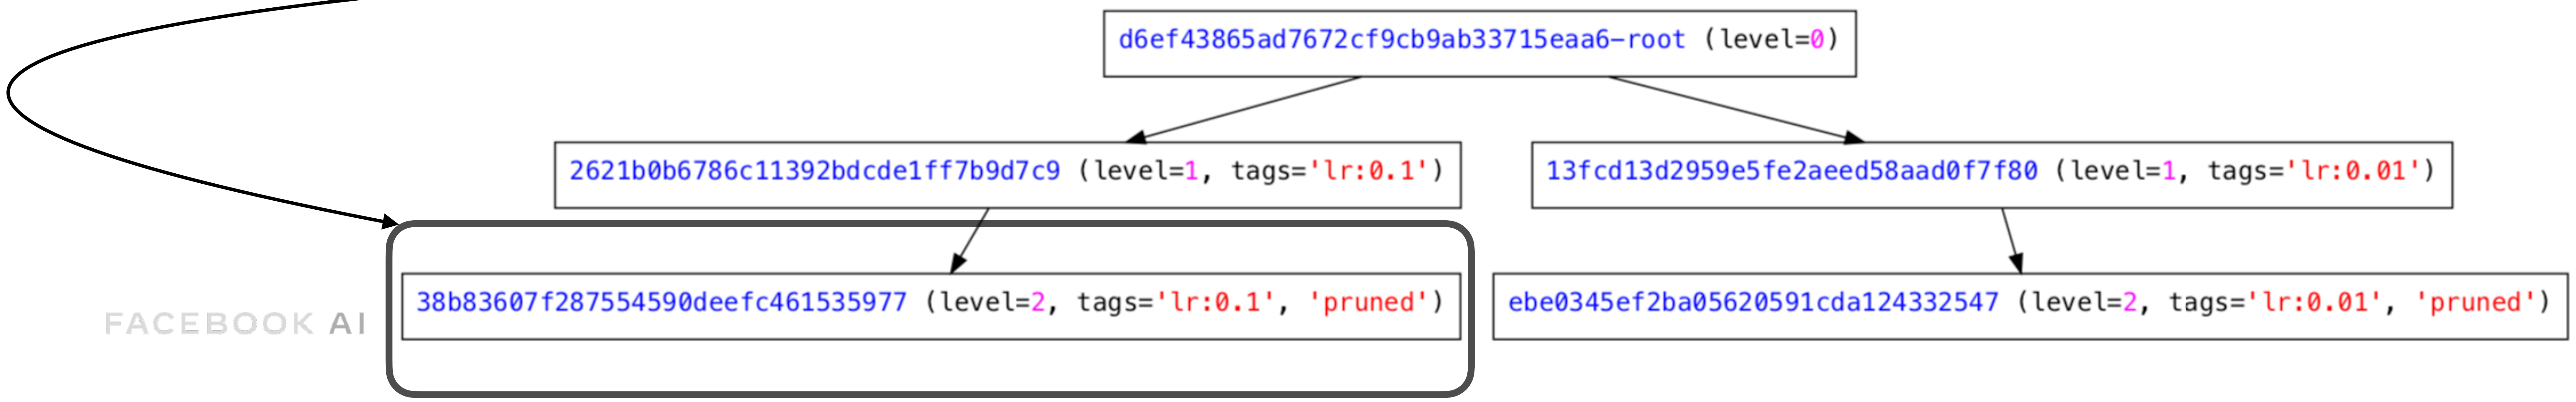

Centralized Pruning in PyTorch

`torch.nn.utils.prune`

torch.nn.utils.prune

Different tensor pruning techniques enabled under a unified framework

BasePruningMethod

CLASS `torch.nn.utils.prune.BasePruningMethod` [SOURCE]

Abstract base class for creation of new pruning techniques.

- CLASSMETHOD `apply(module, name, *args, **kwargs)` [SOURCE]
- `apply_mask(module)` [SOURCE]
- ABSTRACT `compute_mask(t, default_mask)` [SOURCE]
- `prune(t, default_mask=None)` [SOURCE]
- `remove(module)` [SOURCE]

New pruning technique?

Just subclass `BasePruningMethod` and implement `compute_mask`!

PruningContainer

CLASS `torch.nn.utils.prune.PruningContainer(*args)` [SOURCE]

Container holding a sequence of pruning methods for iterative pruning. Keeps track of the order in which pruning methods are applied and handles combining successive pruning calls.

Identity

CLASS `torch.nn.utils.prune.Identity` [SOURCE]

Utility pruning method that does not prune any units but generates the pruning parametrization with a mask of ones.

RandomUnstructured

CLASS `torch.nn.utils.prune.RandomUnstructured(amount)` [SOURCE]

Prune (currently unpruned) units in a tensor at random.

L1Unstructured

CLASS `torch.nn.utils.prune.L1Unstructured(amount)` [SOURCE]

Prune (currently unpruned) units in a tensor by zeroing out the ones with the lowest L1-norm.

RandomStructured

CLASS `torch.nn.utils.prune.RandomStructured(amount, dim=-1)` [SOURCE]

Prune entire (currently unpruned) channels in a tensor at random.

LnStructured

CLASS `torch.nn.utils.prune.LnStructured(amount, n, dim=-1)` [SOURCE]

Prune entire (currently unpruned) channels in a tensor based on their Ln-norm.

CustomFromMask

CLASS `torch.nn.utils.prune.CustomFromMask(mask)` [SOURCE]

Easy to use

```
model = LeNet()  # unpruned model

# L_2 structured pruning will remove 50% of channels across axis 0
prune.ln_structured(
    module=model.conv1,
    name="weight",
    amount=0.5,
    n=2,
    dim=0
)
```

Iterative pruning made easy

`prune.PruningContainer` handles the combination of successive masks for you

```
for _ in range(10):
    # Remove 2 connections per iteration
    prune.l1_unstructured(module=model.fc1, name="bias", amount=2)
```

Global pruning made easy

```
parameters_to_prune = (
    (model.conv1, "weight"),
    (model.conv2, "weight"),
    (model.fc1, "weight"),
)

prune.global_unstructured(
    parameters_to_prune,
    pruning_method=prune.L1Unstructured,
    amount=0.2,
)
```

Easy to extend

```
class FooBarPruningMethod(prune.BasePruningMethod):
    """Prune every other entry in a tensor
    """
    PRUNING_TYPE = 'unstructured'

    def compute_mask(self, t, default_mask):
        mask = default_mask.clone()
        mask.view(-1)[::2] = 0
        return mask
```

supports 3 PRUNING_TYPEs: `'global'`, `'structured'`, and `'unstructured'` (to determine how to combine masks if pruning is applied iteratively)

instructions on how to compute the mask for the given tensor according to the logic of your pruning technique

```
def foobar_unstructured(module, name):
    FooBarPruningMethod.apply(module, name)
    return module
```

`torch.nn.utils.prune`

Prune Responsibly

arXiv:2009.09936

Test hypotheses that class complexity, difficulty, and representation matter in determining the accuracy after pruning

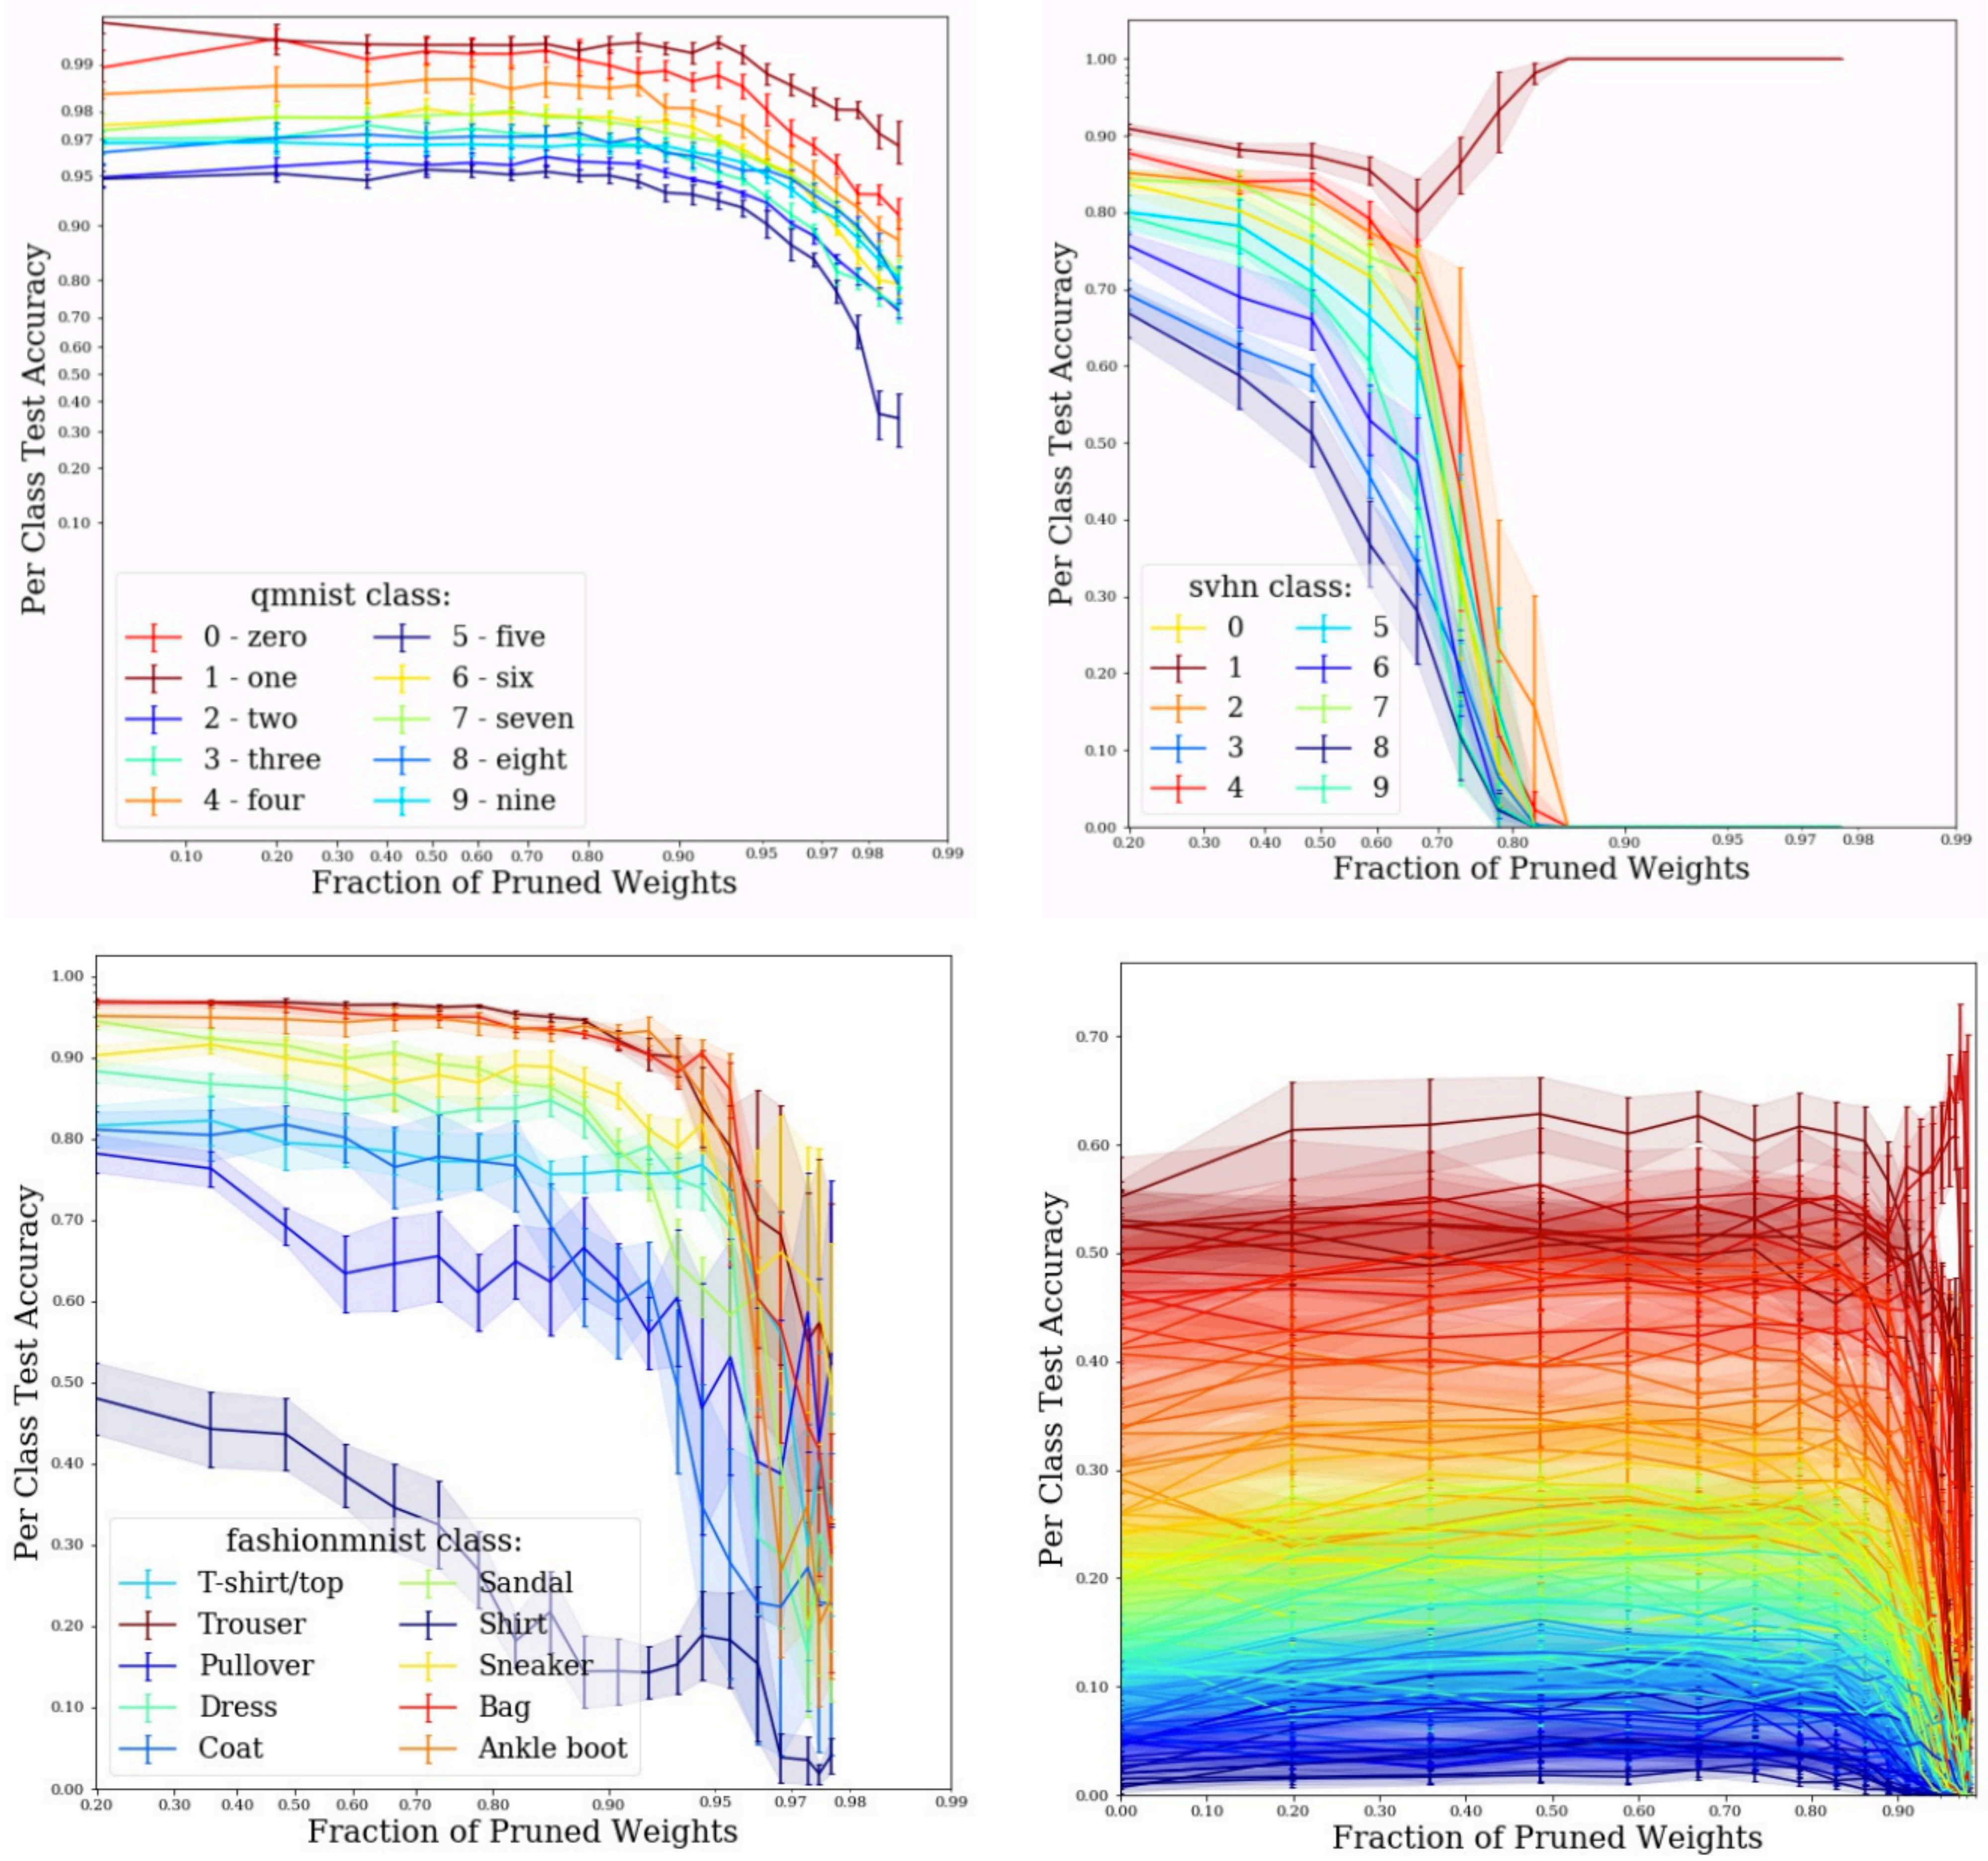

Prune and measure class accuracy for over 1M classes across over 100k models

Fit a linear model for class accuracy as a function of:

- **unpruned model class accuracy**
- **class entropy**
- **class representation**
- sparsity
- dataset
- model
- pruning technique
- weight treatment after pruning

Reject hypothesis that coefficients = 0

Closing Remarks

Kyle Cranmer @KyleCranmer

I'm thrilled that our Machine Learning for Physical Sciences #NeurIPS2020 workshop proposal was accepted! Last year we had a great turnout and fantastic talks - check out the #ML4PS hashtag. New to team: @adjiboussodieng @iamstarnord @glouppe @zdeborova ml4physicalsciences.github.io/2020/

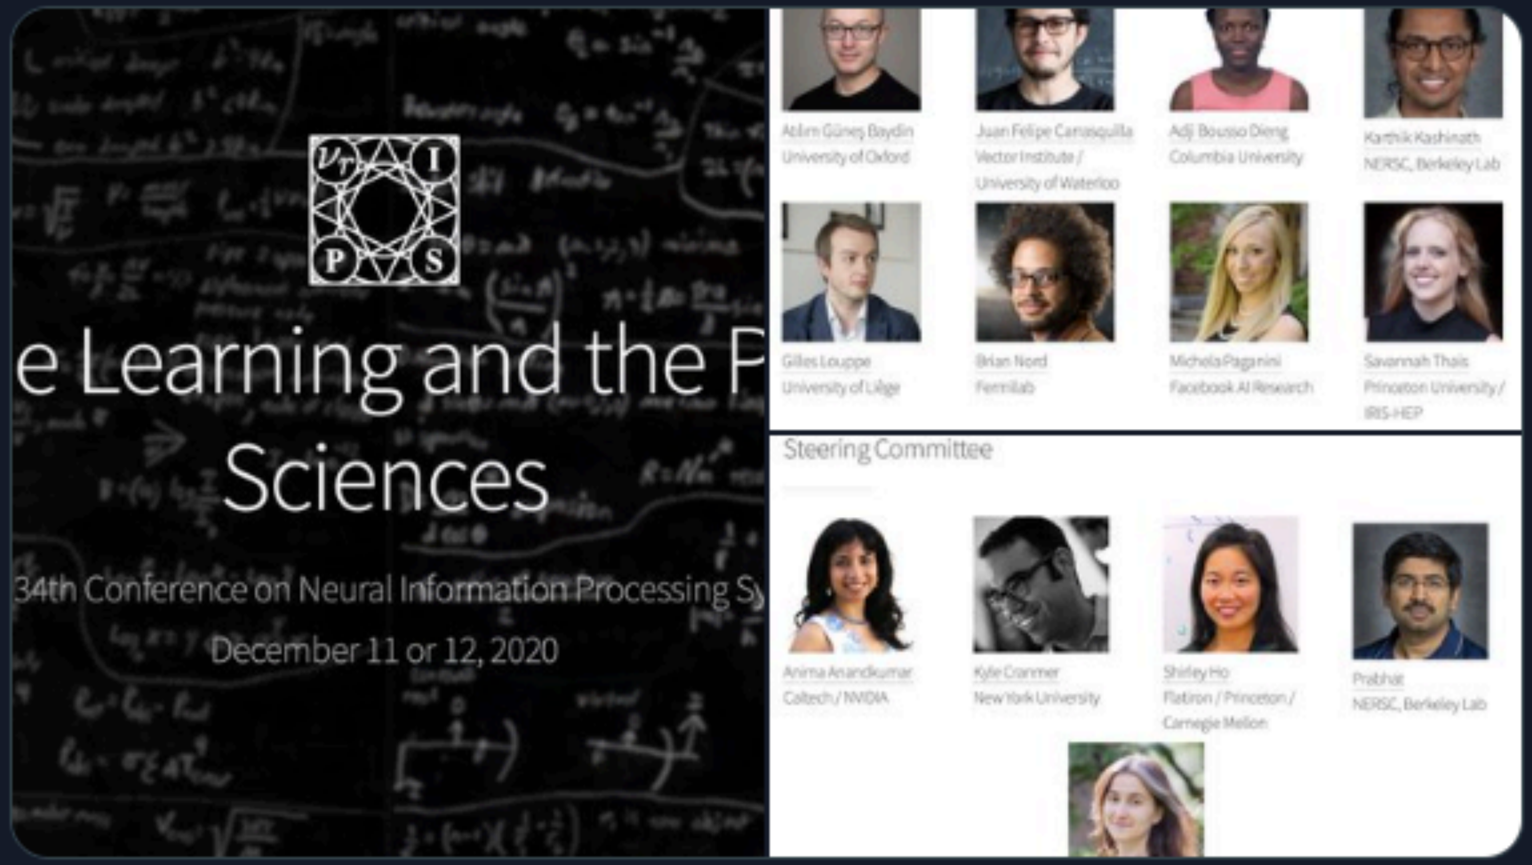

Papers with Code @paperswithcode · Sep 4

🔥 Introducing... the ML Reproducibility Challenge 2020! The 4th annual edition now expands to cover papers from 7 major ML conferences: NeurIPS, EMNLP, ACL, ICML, ICLR, CVPR and ECCV. Find more here at Papers With Code:

Papers with Code - ML Reproducibility Challenge 2020

The ML Reproducibility Challenge 2020 covering paper published in seven major ML conferences: NeurIPS, ACL, EMNLP, ICLR, ICML, CV...

paperswithcode.com

Michela Paganini ✊🏾🙋🏽‍♀️🖤 @WonderMicky · Sep 2

Upset at the latest #NeurIPS2020 reviews? Interested in an alternative publication model for machine learning research? Looking for hypothesis-based science in AI? 👩‍🔬

Check out and submit to the preregister.science workshop at NeurIPS. #SCIENCE!

Deadline: Oct 7. More info 👇

NeurIPS2020 pre-registration workshop

Testing whether pre-registration can help fix our peer review system.

preregister.science

Thanks!

Questions? Contact me

WonderMicky