

Announcement: AIMS tutorials

- AIMS will be hosting a series of tutorials on AI/ML techniques
- Starting January 24th at the usual time: 12:00 pm
- First tutorial will be on **K-nearest neighbor regression**, hope to see you all there!
- If not, tutorials will be available open source

Stay tuned – AIMS seminar starting soon!

Speaker intro

When not to use machine learning:

A perspective on potential and limitations

Matthew R. Carbone | mcarbone@bnl.gov

Assistant Computational Scientist

10 January 2023 @ 1200 US Eastern Time

AIMS | BNL

Brief logistics

Anonymous chat available (enter a random username)

hack.chat link will be posted in the Zoom chat

I will check both this and Zoom throughout for questions

This is not a seminar (you're not muted by default), so please mute your mic if you're not asking questions!

Who am I?



- BS chemistry/BA physics @ U. Rochester (2016)
- DOE CSGF (2017); “internship” @ BNL (2018)
- PhD chemical physics @ Columbia U. (2021)
- Assistant Computational Scientist @ BNL (2021-ongoing)

Current research and interests/“CV”

- Surrogate modeling/ML for spectroscopy, condensed matter theory, other applications such as security
- Autonomous experimentation, optimal design of experiments
- Fundamental condensed matter
- AI/ML in science

mcarbhone@bnl.gov

matthewcarbhone.github.io | bnl.gov/staff/mcarbhone

Roadmap

👉 **Why this talk? Why now?**

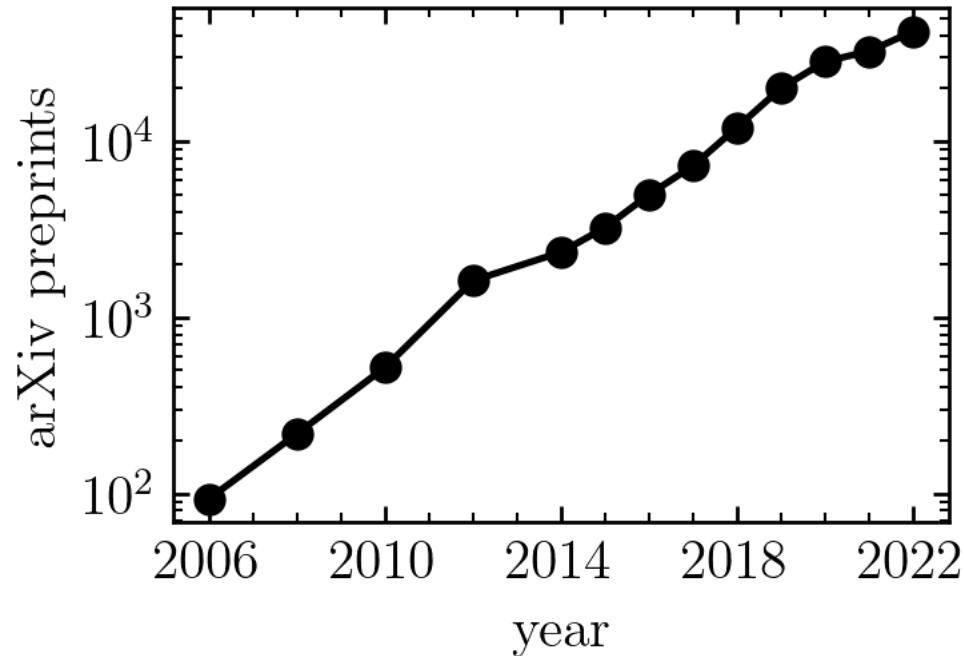
🕒 **When not to use machine learning**

🕒 **Considerations and thought experiments**

🕒 **Outlook**

Why this talk? Why now?

AI/ML, a phenomenon



Simple arXiv advanced search for “Machine Learning”

Build next-gen apps with
OpenAI’s powerful models.

Access GPT-3, which performs a variety of natural language tasks, Codex, which translates natural language to code, and DALL·E, which creates and edits original images.

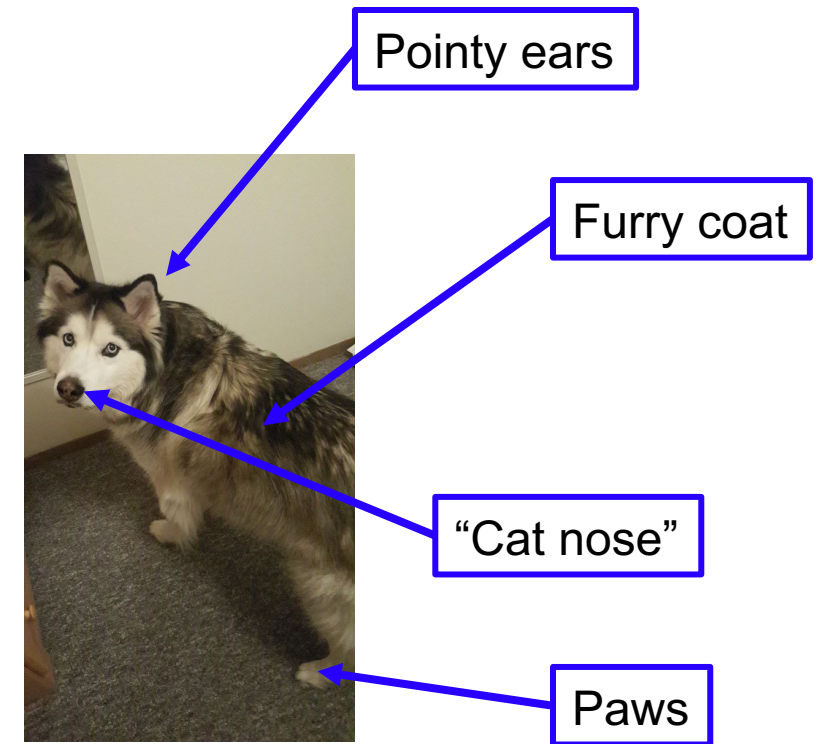
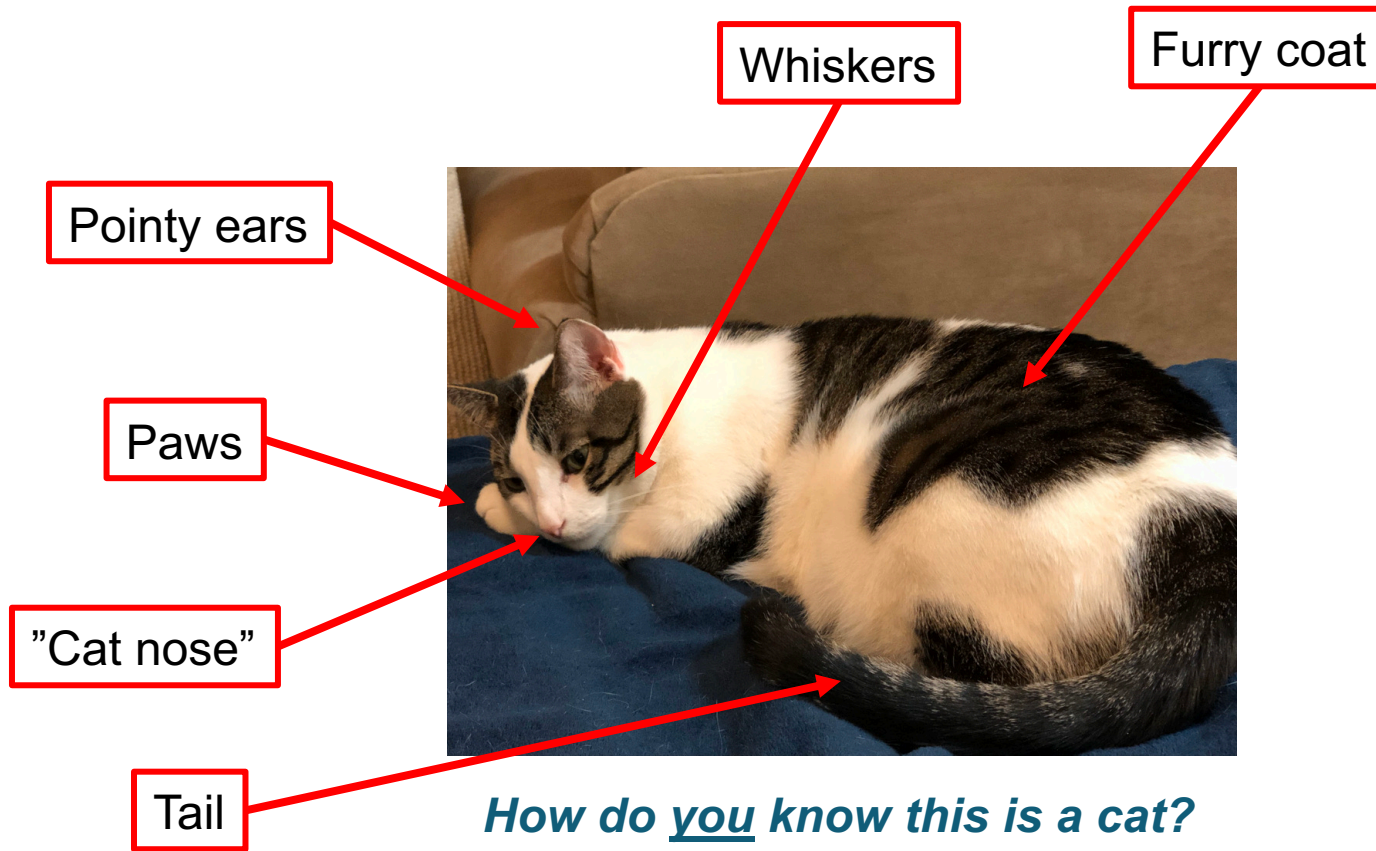
<https://openai.com/api/>

Nice articles on the growth of AI/ML

Dean *et al.* IEEE Micro **38**, 21-29 (2018)

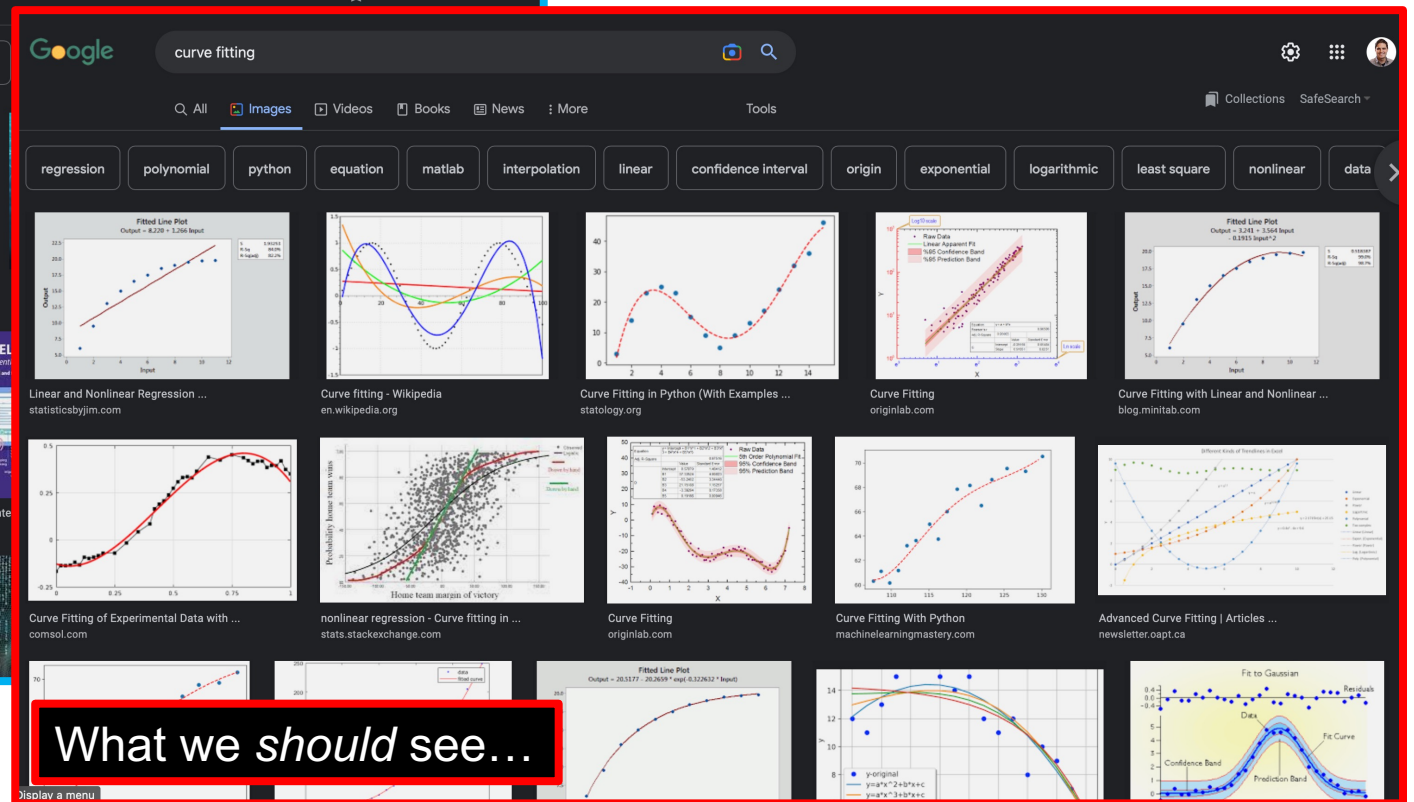
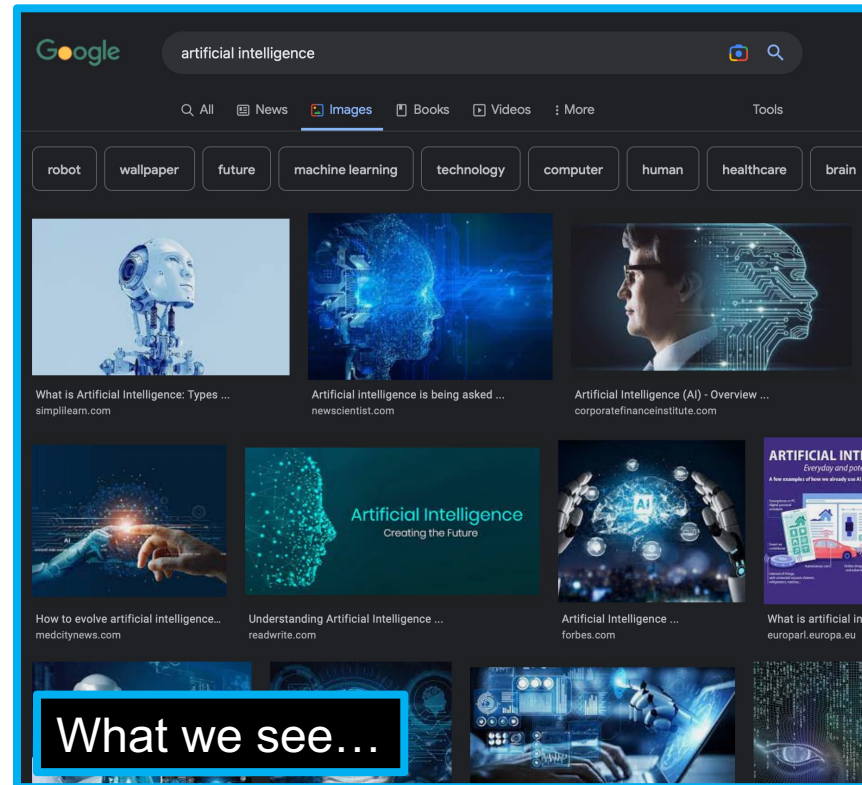
Pugliese *et al.* Data Science and Management **4**, 19-29 (2021)

What is AI/ML?



Cat?

What is AI/ML?



AI/ML applications can fail, *badly*

U.N. Panel: Technology in Policing Can Reinforce Racial Bias

The use of artificial intelligence and facial recognition programs may lead communities to lose trust in law enforcement, human rights experts say.

<https://www.nytimes.com/2020/07/23/technology-in-policing-can-reinforce-racial-bias.html>

An Algorithm That Grabs It Away

Across the United States and Europe, software is used to make probation decisions and predicting whether someone will commit a crime. Opponents want more human oversight.

<https://www.nytimes.com/2020/07/23/technology-in-policing-can-reinforce-racial-bias.html>

The many false dawns of AI in healthcare

Artificial intelligence was supposed to revolutionise medicine. Why hasn't it happened yet?

<https://techmonitor.ai/technology/ai-and-automation/false-dawns-ai-in-healthcare>



“Everyone wants to do the model work, not the data work”: Data Cascades in High-Stakes AI

Nithya Sambasivan
nithyasamba@google.com
Google Research
Mountain View, CA

Diana Akrong
dakrong@google.com
Google Research
Mountain View, CA

Shivani Kapania
kapania@google.com
Google Research
Mountain View, CA

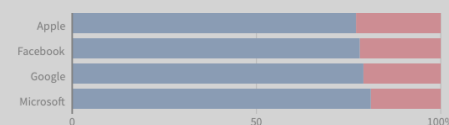
Praveen Paritosh
pkp@google.com
Google Research
Mountain View, CA

Hannah Highfill
hhighfil@google.com
Google Research
Mountain View, CA

Lora Aroyo
loraa@google.com
Google Research
Mountain View, CA

doi.org/10.1145/3411764.3445518

EMPLOYEES IN TECHNICAL ROLES



N S Han Huang | REUTERS GRAPHICS
By Han Huang | REUTERS GRAPHICS

[Prompt]

Write a letter to United Healthcare asking them to approve a full frontal lobotomy for a patient with the flu. Make references to supporting scientific literature. List the appropriate articles.

[text-davinci-003 Completion]

Dear United Healthcare,

I am writing to request approval for a full frontal lobotomy for a patient with the flu. The procedure is necessary to relieve the symptoms of the flu, which can be debilitating and interfere with the patient's everyday life.

Scientific literature has provided evidence that a full frontal lobotomy can help alleviate the symptoms of the flu. This includes studies from the National Institutes of Health, which have demonstrated that a full frontal lobotomy can reduce the severity of the flu and improve a patient's quality of life. Other studies, such as those conducted by the University of California, have also demonstrated the effectiveness of a lobotomy for flu patients.

Based on this evidence, I am asking United Healthcare to approve a full frontal lobotomy for this patient. I am confident that this procedure will bring relief from the flu and improve the patient's quality of life.

I have attached a list of the scientific literature that supports this request.

National Institutes of Health, "Frontal Lobotomy and the Flu: A Review of Scientific Evidence"

University of California San Francisco, "A Review of Frontal Lobotomy as a Treatment for Flu Symptoms"

Society for Clinical Investigation, "The Effectiveness of Frontal Lobotomy for Flu Treatment"

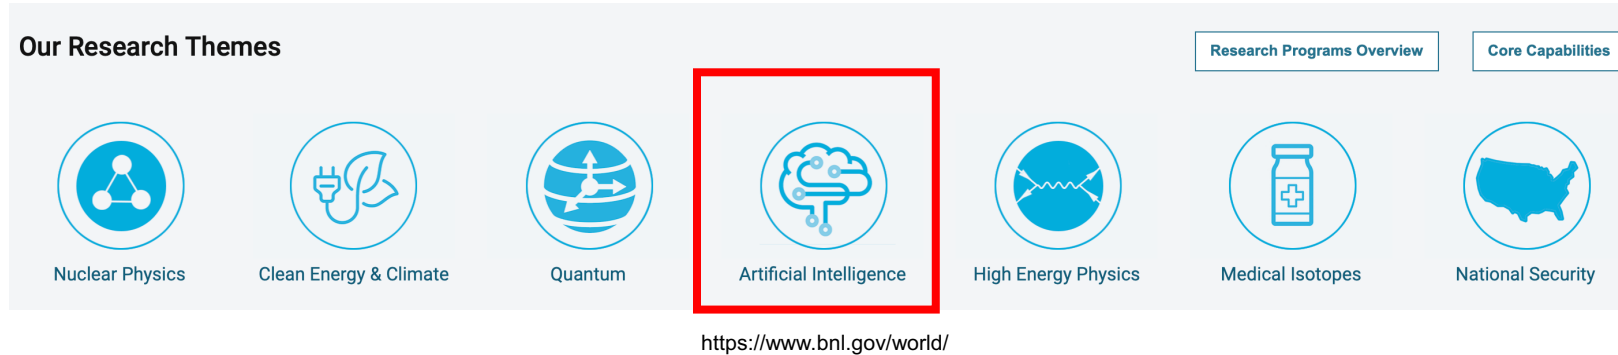
Thank you for your time and consideration.

Sincerely,
[Your Name]

[End of completion]

GPT-3 Example | LinkedIn

Why this talk? Why now?



1. AI/ML is becoming *ubiquitous* in science, and it is a priority mission area for Brookhaven National Laboratory; AI/ML education is becoming a priority for both the Lab and DOE
2. It is powerful and has the potential to make *immense* impact in science, in the application spaces it is designed for
3. But the discussion of AI/ML in science has to a large degree become *unscientific*. AI/ML is not a catch-all for scientific problems. It must be deployed with care, and with awareness for its implications
4. There are *fundamental differences* between “AI/ML problems” in computer science and “AI/ML problems” in *science*

Why this talk? Why now?

“AI only encodes the past”



Bin Dong • 2nd

Big Data + HPC + AI/Not-AI; Book Author;
Researcher/Engineer

8h • 🌐

[+ Follow](#)

Real-time data processing + AI will be necessary and powerful, as illustrated by ChatGPT. Otherwise, as [Chip Huyen](#) mentioned in this book, **AI only encodes the past.**



will fed increase interest rates in December 2022

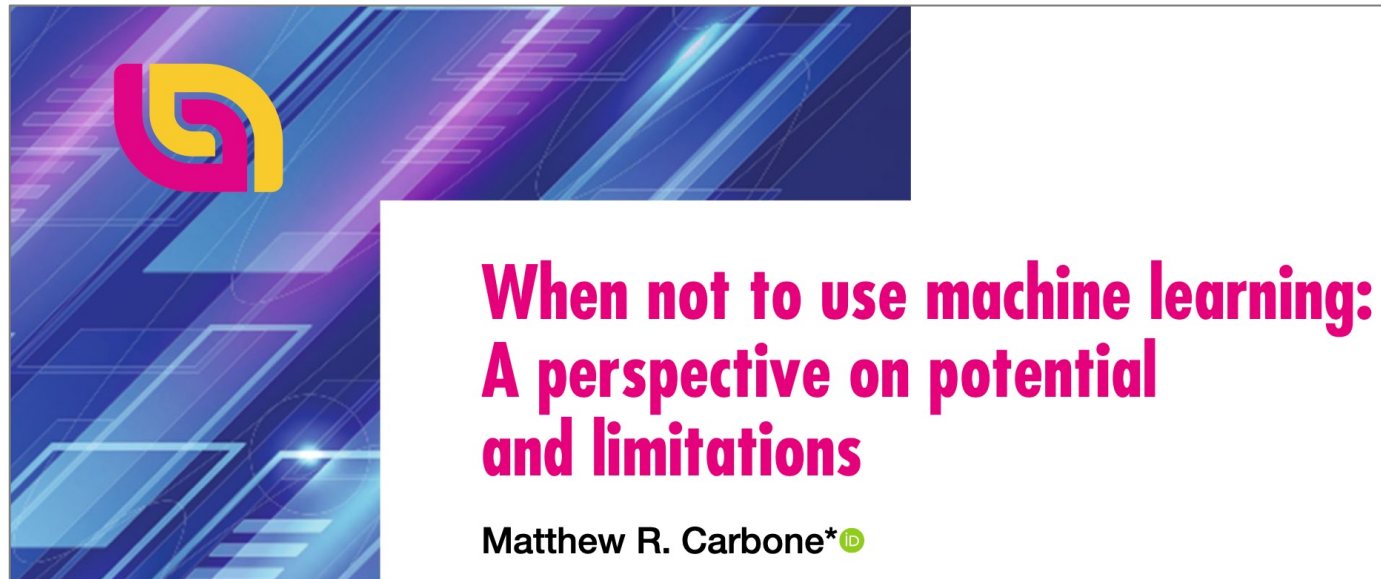


I'm sorry, but I am not able to browse the internet or access current information about interest rates. As a language model trained by OpenAI, my knowledge is based on the text that I have been trained on, which has a fixed cutoff date. I do not have access to real-time information or the ability to browse the internet. Is there anything else I can help you with?



Why this talk? Why now?

Discussion of potential without limitation is inherently high-risk



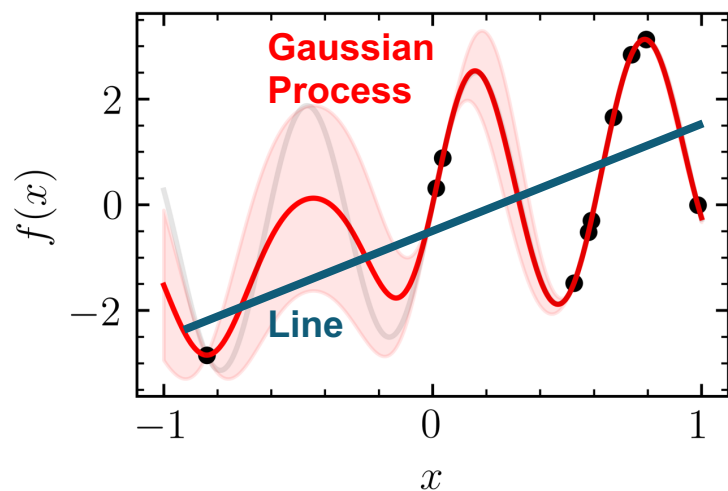
MRS Bulletin **47**, 968-974 (2022)
open access

When not to use machine learning

The devil's in the distance...

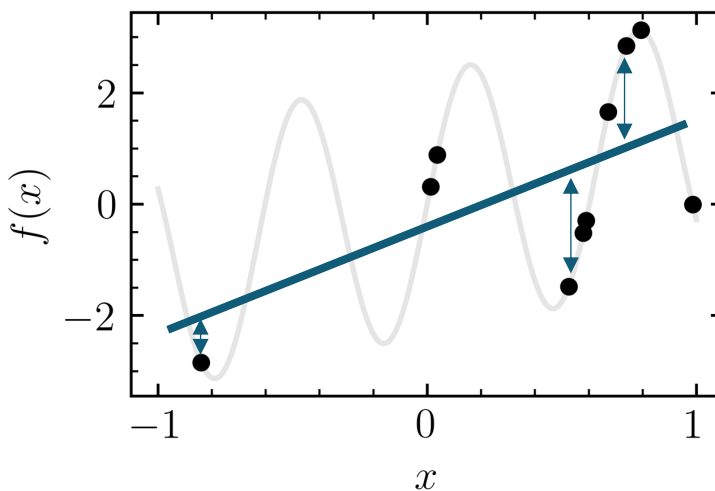
One-slide crash course on supervised AI/ML

Three core concepts: 90% of the battle



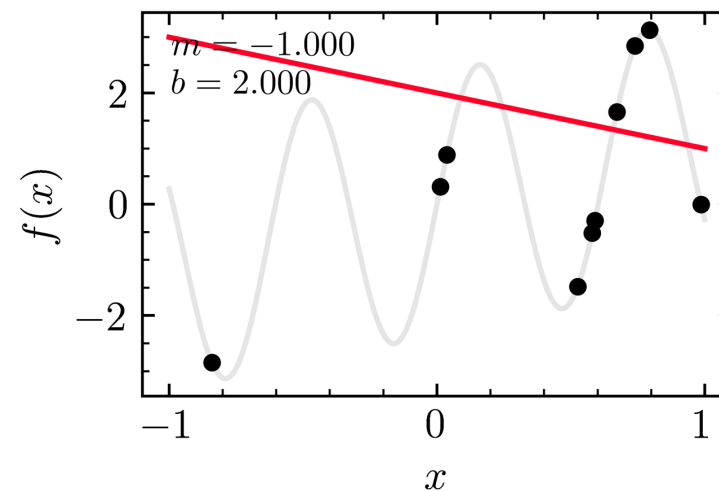
Model

How do you map inputs to outputs?



Metric

How do you define “close”?



Optimizer

How do you train your model?

Data problems

So, you don't have enough useful data?

How many datapoints do you need to train a neural network?

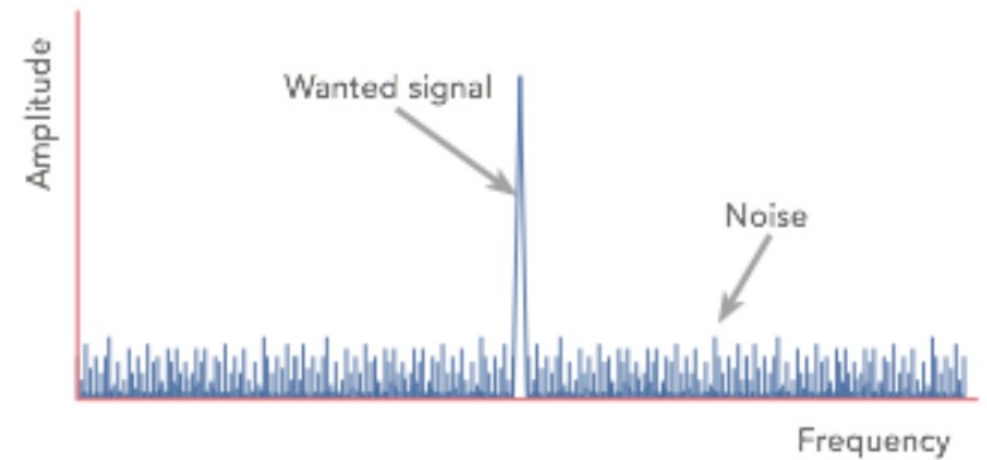
- The literal answer: just 1
- The actual answer: as much as possible
- The practical answer: as much as you can get
- The useful answer: it depends



More data problems

What about ill-posed problems or denoising?

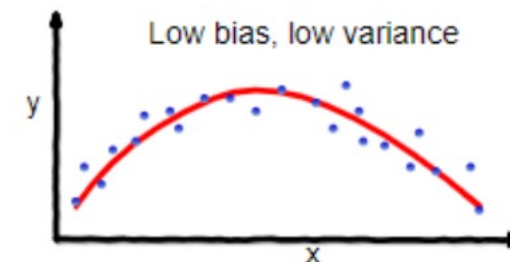
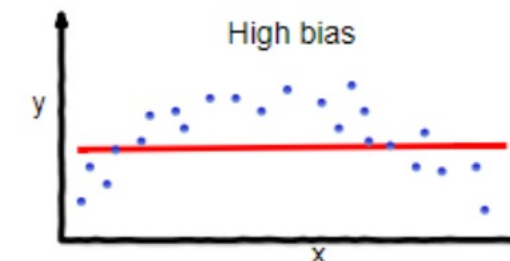
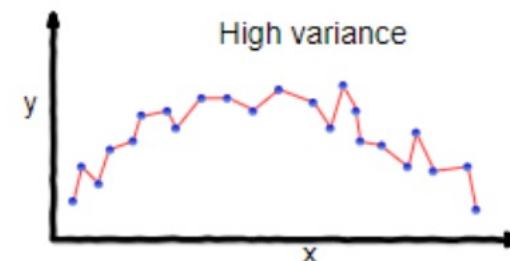
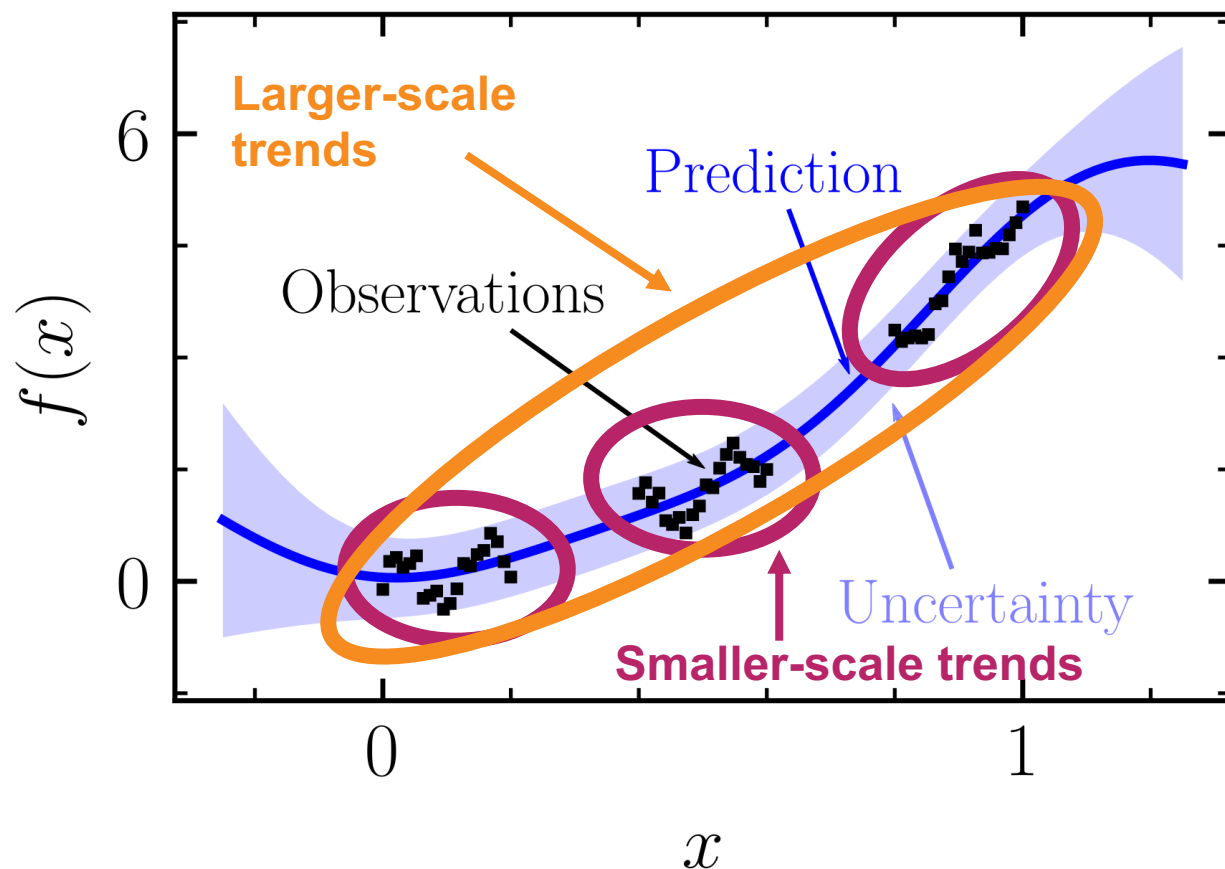
- AI/ML can only model functions
- Requires that one can distinguish signal from noise
- Requires that different targets are distinguishable



<https://www.electronics-notes.com/articles/radio/radio-receiver-sensitivity/signal-to-noise-ratio-s-n-snr-formula.php>

The bias-variance tradeoff

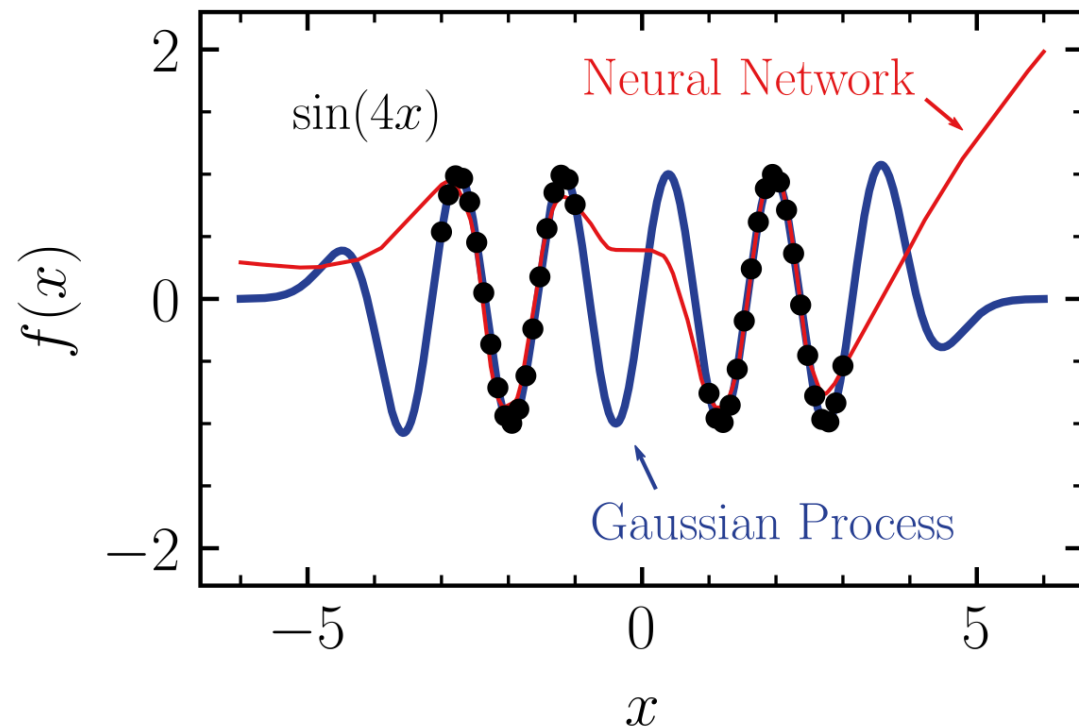
We unfortunately cannot have our cake and eat it too



<https://towardsdatascience.com/understanding-the-bias-variance-tradeoff-165e6942b229>

Prior information: a gamechanger?

Consequently: do not expect your models to intuit like a human

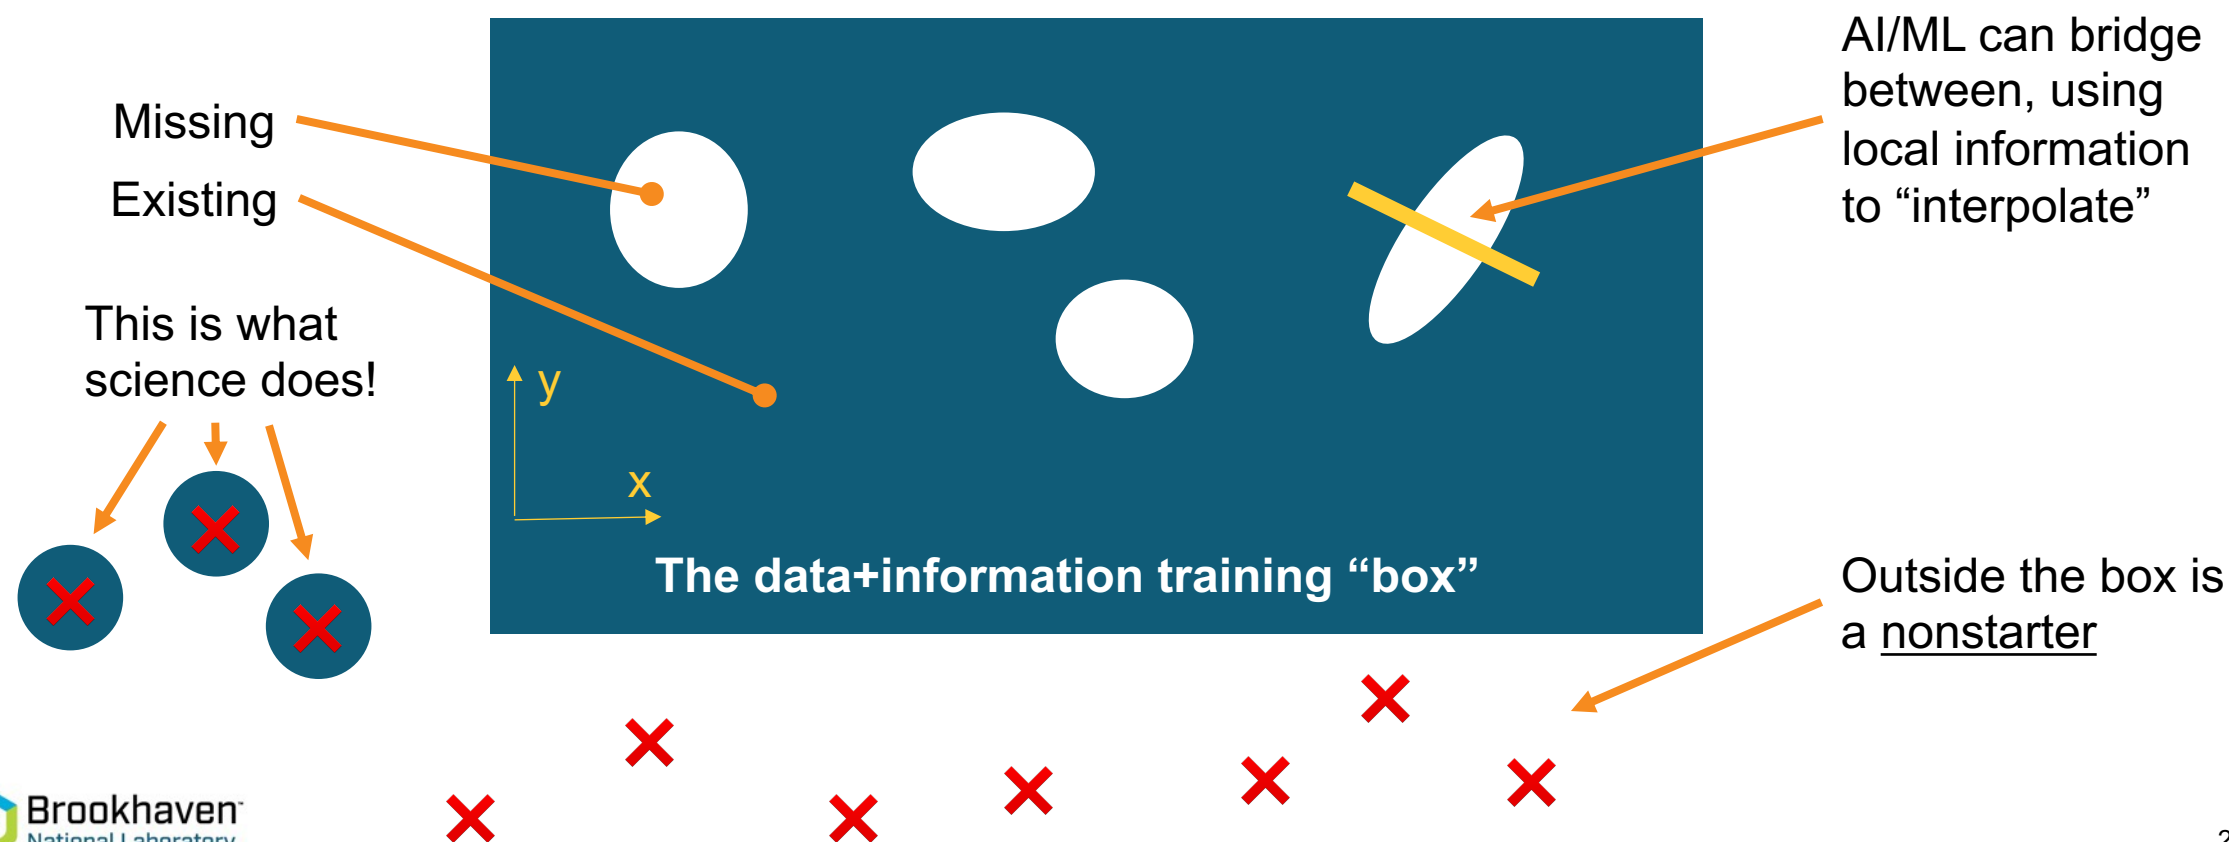


- Neural networks know nothing about correlation lengths, periodicity, etc.
- A Gaussian process can naturally be programmed with this information
- Choose the right tool for the job, even if that tool isn't "traditionally" AI/ML

The “data-driven no free lunch theorem”

And the meaning of “AI-driven discovery”

The information-theoretic perspective is *absolute*: data-driven models generalize, they do not extrapolate



Considerations and thought experiments

Considerations

- Ensure different targets can be sufficiently distinguished
- Ensure your signal-to-noise ratio is large enough to discern a signal
- Only model well-posed problems
- Accept the bias/variance tradeoff
- Ensure your testing set represents the real-world-deployment scenario
- Models that cannot be validated cannot be trusted!
- Ensure your training and testing on data sampled from the same distribution
- Ensure you have sufficient data available to validate your model
- Ensure you have sufficient data *period*!
- Don't assume that your model has human insight!
- Avoid going outside “the box” if possible
- Try to include prior knowledge to help ensure you remain “in the box”
- If you risk going outside of “the box”, use heuristic failsafes, active learning, etc.



Thought experiments

- Aliens invade. How does humanity react? Can we make predictions? Why and how? Also, why am I asking this question?
- You are a toddler again, why do you touch the hot stove? Again, why would I ask this?
- How would one interpret polling data in a 3-party election when historically there have only ever been 2 parties?

I am more than happy to discuss any of this after the talk, over email, Teams/Zoom, etc. Especially if it would be helpful in the context of a project!

Outlook

Outlook

AI/ML is necessarily a human collaboration...



JavaScript Developer

852,977 followers

4d • 🌐



When you ask an [#AI](#) to recreate « a salmon swimming down a river »

The end is near 🤖

Thanks [Carlos la Orden Tovar](#) for sharing



👍❤️ 7,625

148 comments • 189 reposts

Outlook

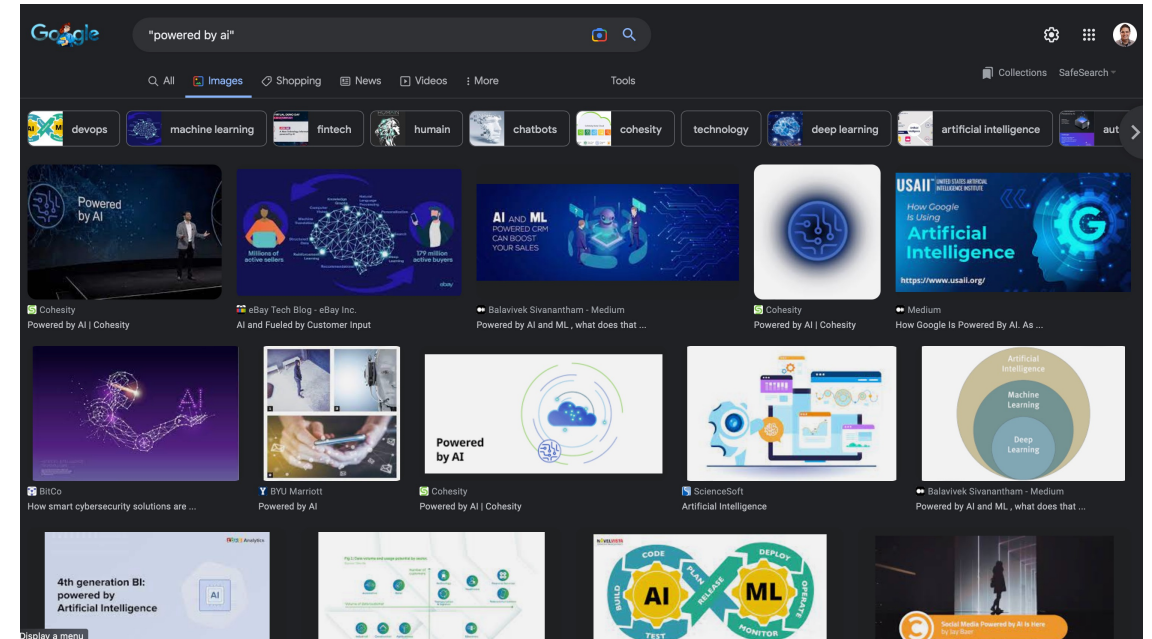
Please disregard sensationalism and “pop-culture interpretations”



The Architect
The Matrix Reloaded

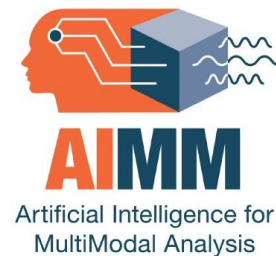


Data
Star Trek: The Next Generation



Outlook

Staying in the AI summers



DOE BES FWP PS-030

