

Emergent phenomena in large-scale LLMs


Daniel Khashabi



JOHNS HOPKINS
UNIVERSITY


The success we dreamed of

Language models that **reason**
across many tasks that were once considered aspirational.




How do these models work?


Learning ~ Compression + Replication



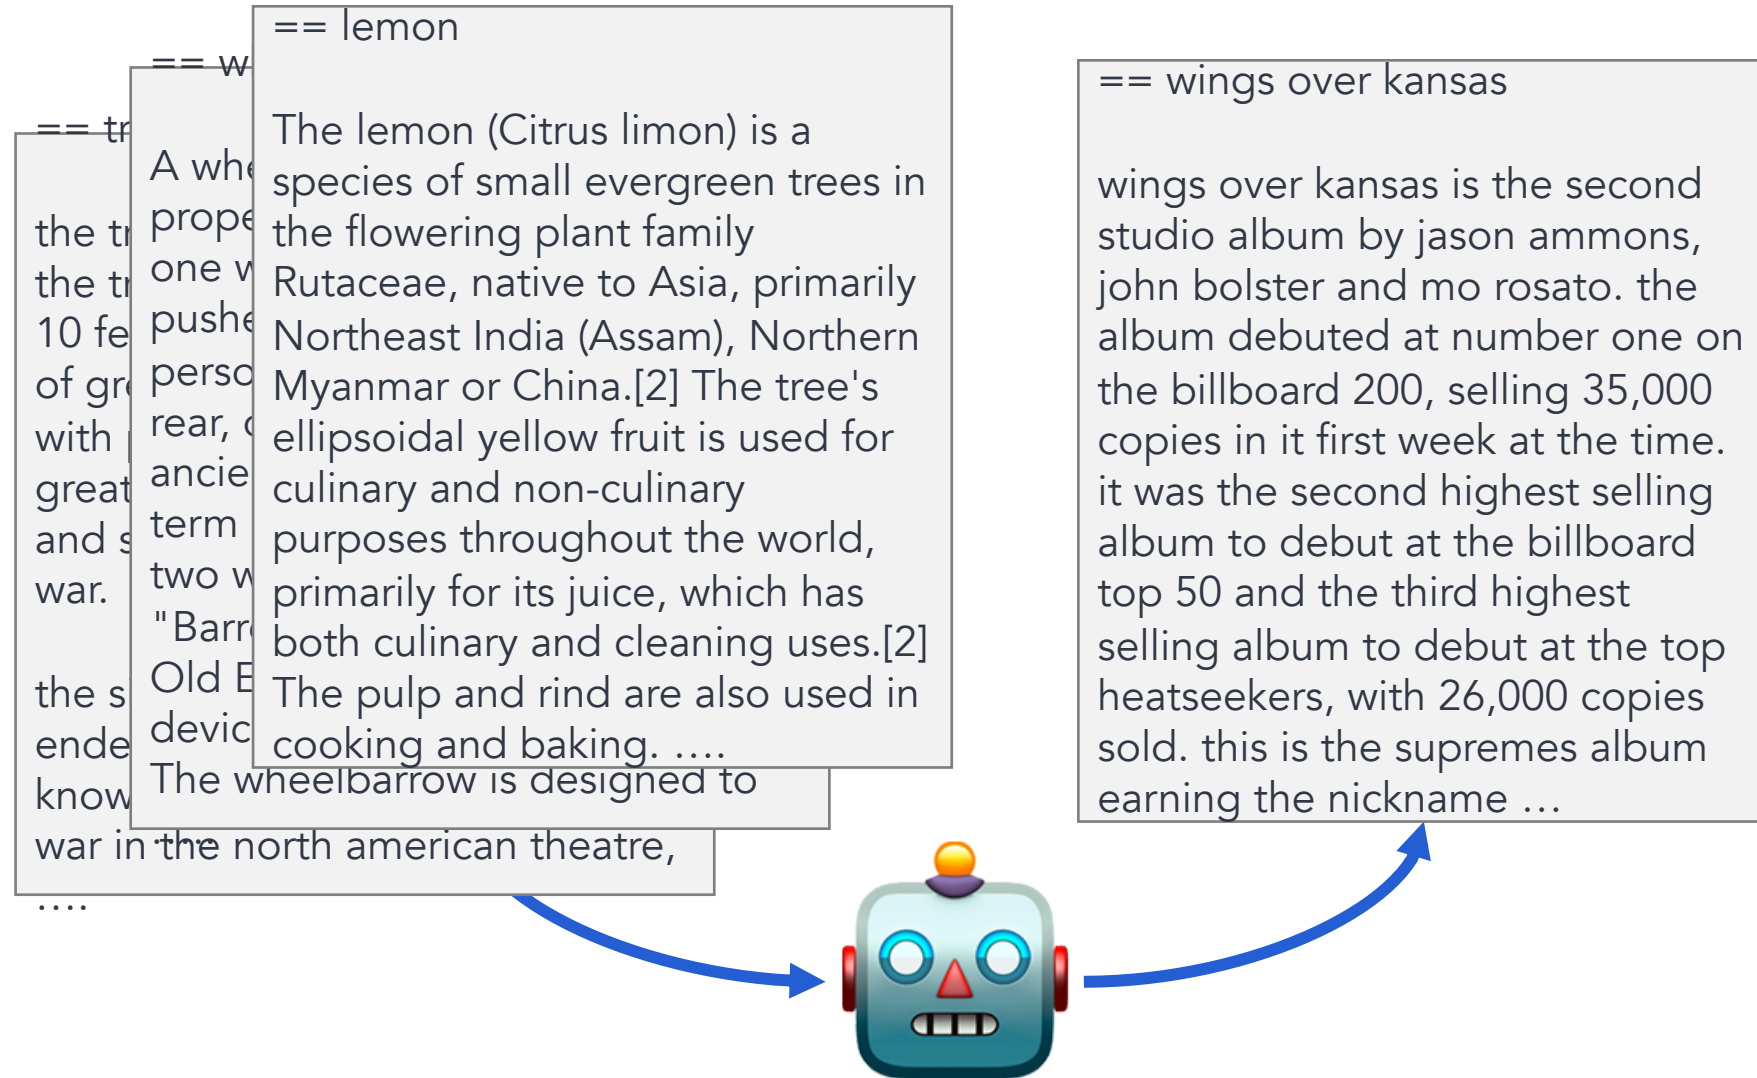
Learning ~ Compression + Replication




Learning ~ Compression + Replication



Learning ~ Compression + Replication




Learning ~ Compression + Replication
Are **predictive models** of the world.



Learning ~ Compression + Replication

Are **predictive models** of the world.


Are learned from **unlabeled** data.




Learning ~ Compression + Replication

Are **predictive models** of the world.


Are learned from **unlabeled data**



AI Models ~ Sponges




AI Models ~ Sponges



AI Models ~ Sponges


Examples of
the World

New
Instances



Scaling up AI Models Makes it Better

Examples of
the World



New
Instances

Today: Emergent phenomena

- What is happening inside language models?
- What do we mean by “emergent phenomena”?
- What do we understand about it?
- What does the future hold?




Today: Emergent phenomena

- What is happening inside language models?
- What do we mean by “emergent phenomena”?
- What do we understand about it?
- What does the future hold?



Recipe: (1) Pre-training


- **Pre-training** := Next-Token Prediction



- Modern LLMs are pre-trained on many-trillions of tokens.
 - For comparison, if you read **1 book per week**, you'd need 50,000 years to cover a comparable amount of text!!!

Recipe: (2) Post-training

- “Aligning” LMs to respect our intents embedded in instructions.
 - Supervised Fine-tuning (SFT/behavior cloning) with labeled data.




Cross-Task Generalization via Natural Language Crowdsourcing Instructions.
Mishra, Khashabi, Baral and Hajishirzi. *ACL* 2021 🏆 **Ai2 Lasting Impact Award!** 🏆

Among others: Sanh et al. 2022; Chung et al. 2022,


Recipe: (2) Post-training

- “Aligning” LMs to respect our intents embedded in instructions.
 - Supervised Fine-tuning (SFT/behavior cloning) with labeled data.
 - Reinforcement Learning on preference data or verifiers.




Scaling Recipe: The whole thing

- Almost all the reasoning models follow this recipe:




Demo time: a pre-trained model

- <https://gist.github.com/danyaljj/3ac2f34651d6d9dd4b422daa09dd282a>



Since 1950's




Neural
Networks

Since 1980's



Massive
Data

Since 1940's
(GPUs 1990's)



Fast
Computers

Our overnight success was
70 years in the making!

Today: Emergent phenomena

- What is happening inside language models?
- What do we mean by “emergent phenomena”?
- What do we understand about it?
- What does the future hold?




Today: Emergent phenomena

- What is happening inside language models?
- What do we mean by “emergent phenomena”?
- What do we understand about it?
- What does the future hold?





What is “Scaling” pre-training LLMs?

- “Scaling” models is not one thing. There are various axes.
- **Larger model size**
 - But model params may be under-utilized.
- **More data**
 - But more data might not necessarily contain more information (e.g., duplications)
- **More compute**
 - But unnecessary computation may be wasted.
- Scale means all the above: *effective compression of information*
 - Requires model capacity
 - Requires large, rich data
 - Requires compute



Pre-training: Effect of scale



Kaplan et al. 2020;
among others

More data and compute $\Rightarrow$ better results.


More of the same produces something qualitative different

- Anderson '72 argued that complex systems can behave in ways that can't be understood by *the laws that govern their microscopic parts*.
- He argued that new properties appear at each level of complexity.



“More is Different”: Examples

- No single molecule is wet.
- When you put billions of them together, they’re wet.
- Although chemistry is subject to the laws of physics, we can’t (easily) infer chemistry from our knowledge of physics.



“More is Different”: Examples

- No ants have grand plans. (to our knowledge)
- Their colony builds cathedrals.



Why does this matter?

- **Unpredictability:**

- Testing a smaller model doesn't tell you what the larger model can do.

- **Safety implications:**

- Standard pre-deployment testing assumes you know what to test for.
- Emergent abilities violate that assumption.

- You can't test what you don't know.

Today: Emergent phenomena

- What is happening inside language models?
- What do we mean by “emergent phenomena”?
- What do we understand about it?
- What does the future hold?

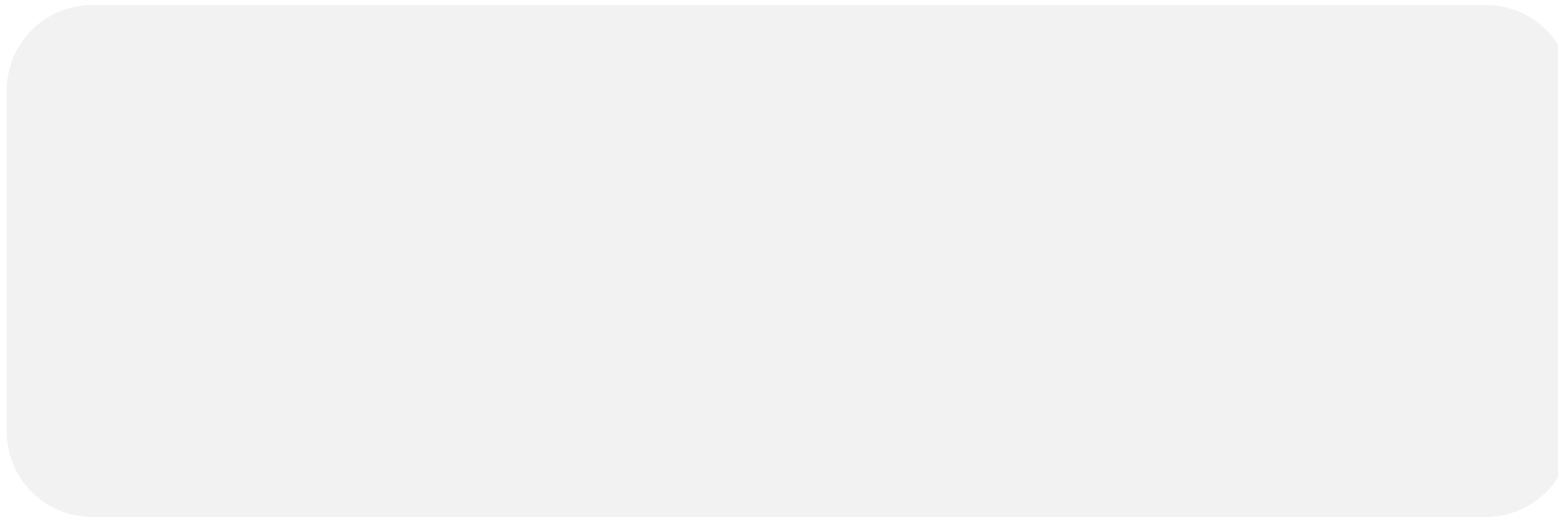


Today: Emergent phenomena

- What is happening inside language models?
- What do we mean by “emergent phenomena”?
- What do we understand about it?
- What does the future hold?



A working example



A working example

Input: NYU

A working example

Input: NYU Output: NYC

A working example

Input: NYU Output: NYC
Input: UMD

A working example

Input: NYU Output: NYC
Input: UMD Output: DC

A working example

Input: NYU Output: NYC

Input: UMD Output: DC

Input: JHU

A working example


Input: NYU Output: NYC

Input: UMD Output: DC

Input: JHU Output:

In-Context Learning

- **ICL** := learning to imitate the *implicit* pattern described by few examples provided in the context.



- This is very interesting since we didn't explicitly train our models to show this ability!! (it emerged from large-scale pre-training).
- The models *infer* the patterns from their input.

Demo time: ICL

- <https://gist.github.com/danyaljj/3ac2f34651d6d9dd4b422daa09dd282a>

Wait — this shouldn't work

- We only trained the model to predict the next token
- We never showed it "here are examples, learn the pattern"
- Yet at scale, it just... does
- No gradient updates. No fine-tuning. Just prompting.

The big **open** question: Why (and When) does ICL emerge?

- For years, the organic emergence ICL (from 'next-token' training) is deeply tied to **human language**.
- Perhaps human language has specific properties that yields ICL.
 - Compositionality, abstraction, parallel patterns, etc.

A Theory of Emergent In-Context Learning as Implicit Structure Induction

Michael Hahn
Saarland University

hahn@cs.uni-saarland.de

Navin Goyal
Microsoft Research India
navingo@microsoft.com

Parallel Structures in Pre-training Data Yield In-Context Learning

Yanda Chen¹ Chen Zhao^{2,3} Zhou Yu¹ Kathleen McKeown¹ He He²

¹Columbia University, ²New York University, ³NYU Shanghai

{yanda.chen, kathy}@cs.columbia.edu, cz1285@nyu.edu
zy2461@columbia.edu, hehe@cs.nyu.edu

The big **open** question:

Why (and When) does ICL emerge?

- For years, the organic emergence ICL (from 'next-token' training) is deeply tied to **human language**.
- But ... is it really about human language specifically?

Research question: Is ICL a [human] language-specific phenomenon?

Human vs. Genetic Language

Human language

The computer science department at JHU is known for ...


Genomic language

ATTACCCCGATTGCTATGCCTGAGAAGCTATTATGCCTGAGAAGCTAT ...

- **Similar:** Both are symbolic, short- and long-range dependencies.
- **Different:** Vocabularies, interpretability, evolutionary path.



Hypothesis: If ICL appears in both settings, it must reflect a *more general computational mechanism*.



What task should we use for evaluation?

```
input1 -> Output1 SEP
input2 -> Output2 SEP
input3 -> Output3 SEP
Input4 -> ?
```

?



What task should we use for evaluation?

We defined 100 reasoning tasks
based on bitstrings


```
101000000 -> 00000101 SEP
111000011 -> 11000111 SEP
11001110 -> 01110011 SEP
110000000 -> ?
```

Various functions: Bitwise NOT, Reverse, etc.



But we need to express the *same underlying task* in a language each model understands.



What task should we use for evaluation?



Genomic Models are In-Context Learners



Both models exhibit **log-linear gains** in pattern induction as the number of in-context demonstrations increases.

Summary

- Next-nucleotide predictors are in-context learners.
- Emergent ICL is clearly not tied to human language.
- ICL is a consequence of large autoregressive training on rich distribution.

How about other modalities?

- What other data can yield ICL? And what's common in them?

✓ Human language

Today I'm presenting ...

✓ DNA sequences

T A T A A A A A A T A T G T T T ...

? Brain signals

12, 15, 18, 14, 10, 8, 11, 17, 22,...

? Time series


18934, 18951, 18920, 18885, ...

? Protein sequences


M A S M T G G Q Q M G S H ...



① Sample an abstract



① Sample an abstract



Language (digits)

Genome (nucleotides)

Protein (amino acids)

Integer sequence (digits)


Image (pixels)


Time series (half-sine lobes)

② Form encoding codebook




③ Encoding ICL prompts for each modality





Summary

- Most next-token predictors trained on natural data are in-context learners..
- ***Convergent In-Context Induction*** Hypothesis:
 1. ICL is a consequence of next-token prediction:
it organically emerges in any modality with rich data, and
 2. its realizations share a **common** substrate across modalities



Open Questions

- If reasoning arises without human-interpretable tokens, what does that say about the nature of “reasoning”?
- Is it possible one can learn to reason (say in human lang) by learning from *other* modalities?

Other emergent phenomena we didn't cover

- **Generalization to unseen task instructions**

[Mishra et al. 2021 among others]

- **Inference-time reasoning:** models improving their responses with more inference-time thinking

[DeepSeek-R1, among others]

- **Cross-lingual transfer:** capabilities learned in English appear in low-resource languages

Today: Emergent phenomena

- What is happening inside language models?
- What do we mean by “emergent phenomena”?
- What do we understand about it?
- What does the future hold?



Today: Emergent phenomena

- What is happening inside language models?
- What do we mean by “emergent phenomena”?
- What do we understand about it?
- What does the future hold?



How far will we be able to scale?

- Is “scaling” here to stay?

Someone: “we won’t have enough compute”

- **For:**

- There is simply not enough compute available.
- Even if we have the compute (in theory), scaling the compute will be quite costly.

- **Against:**

- Computing infrastructure continue to progress at a rapid pace.
 - E.g. renewables, space data centers, inference hardware, etc.
- The economy of scale will just work out.
 - With more usage, cost per unit will be cheaper.

Someone: “we won’t have enough data”

- **For:**

- Existing LLMs are severely data hungry. There is simply not much enough data in the world.
- And you don’t need just any data; you want data on challenging and rare problems (e.g., hard math and engineering problems).

- **Against:**

- There is massive amount of untapped data in the world.
- As we deploy AI in more applications, allowing continuous stream of fresh data.

How far will we be able to scale?

- Is “scaling” here to stay?

Putting things together

- **Emergence is real:** Large-scale pre-training on *simple* objectives lead to capabilities that appear at scale.
- **In-context learning** is a prominent manifestation of emergence.
 - Not specific to human language.
- **It is important to understand emergence:** It is not just an academic endeavor. It is how a foundation of how we form trustworthy AI.

