




Does Algorithmic Monoculture lead to Outcome Homogenization?

Rishi Bommasani, Kathleen A. Creel, Ananya Kumar
Dan Jurafsky, Percy Liang


Framing

Modern AI centers on sharing
ImageNet, PyTorch, BERT, Adam, ...
What are the harms?
especially individual-centric harms
Systemic Failure:
all models fail for same person
Homogenization:
rate of systemic failure is high

Sharing training data




Fixed (more sharing) means the different companies use exactly the same training data.
Disjoint (less sharing) means the different companies use disjoint samples from the same distribution.



Left is homogenization at individual-level, right is for the 9 racial groups tracked in US census.
Compare left to right (avg. unif): **homogenization is higher for individuals than racial groups.**

Sharing foundation models



Effect of no foundation model (scratch) vs. different ways to use that are more/less expensive.
Visions: scratch is the most homogeneous, then **probing**, then finetuning.
Language: **probing** is the most homogeneous; **finetuning** and **BitFit** are similarly homogeneous.

Takeaway

Sharing data homogenizes
Sharing models is more complex
Need to think about distribution shift

Narrative

Decision-makers deploy systems that **share** components.
Sharing can be reinterpreted as *monoculture!*
Due to monoculture, will some individuals **exclusively** receive negative outcomes?

Experimental Findings

For training data, sharing homogenizes outcomes.
For foundation models, its more complicated.

TL;DR

Monoculture is ubiquitous & growing:
we need to actively study homogenization!

Metric

Companies i in $\{1, \dots, k\}$
Individuals j in $\{1, \dots, N\}$
 I is indicator RV of failure

Per-company failure rate:
 $\text{fail}(h^i) \triangleq \mathbb{E}_{x^i \sim D^i} I^i(x^i) = \Pr_{x^i \sim D^i} [I^i(x^i) = 1]$

Obs. systemic failure rate:
 $S \triangleq \mathbb{E}_j \left[\prod_i I^i(x_j^i) \right] = \Pr_j [I^1(x_j^1) = 1 \wedge \dots \wedge I^k(x_j^k) = 1]$

Homogenization metric:
 $H^{\text{individual}}(h^1, \dots, h^k) \triangleq \frac{S}{\prod_{i=1}^k \text{fail}(h^i)} = \frac{\mathbb{E}_j \left[\prod_i I^i(x_j^i) \right]}{\prod_i \mathbb{E}_j [I^i(x_j^i)]}$

Challenges for society

- Diagnosis: Opacity (unaware of monoculture)
- Measurement: Privacy and linking individuals
- Rectification: Incentives-compatibility

Why this matters?

- Harms for *individuals*
- Look beyond a single model
- Systemic exclusion/hierarchy
- Relational equality