Attribution: Bayes vs. ML

Kate Marvel
NASA GISS

NEW APPROACHES

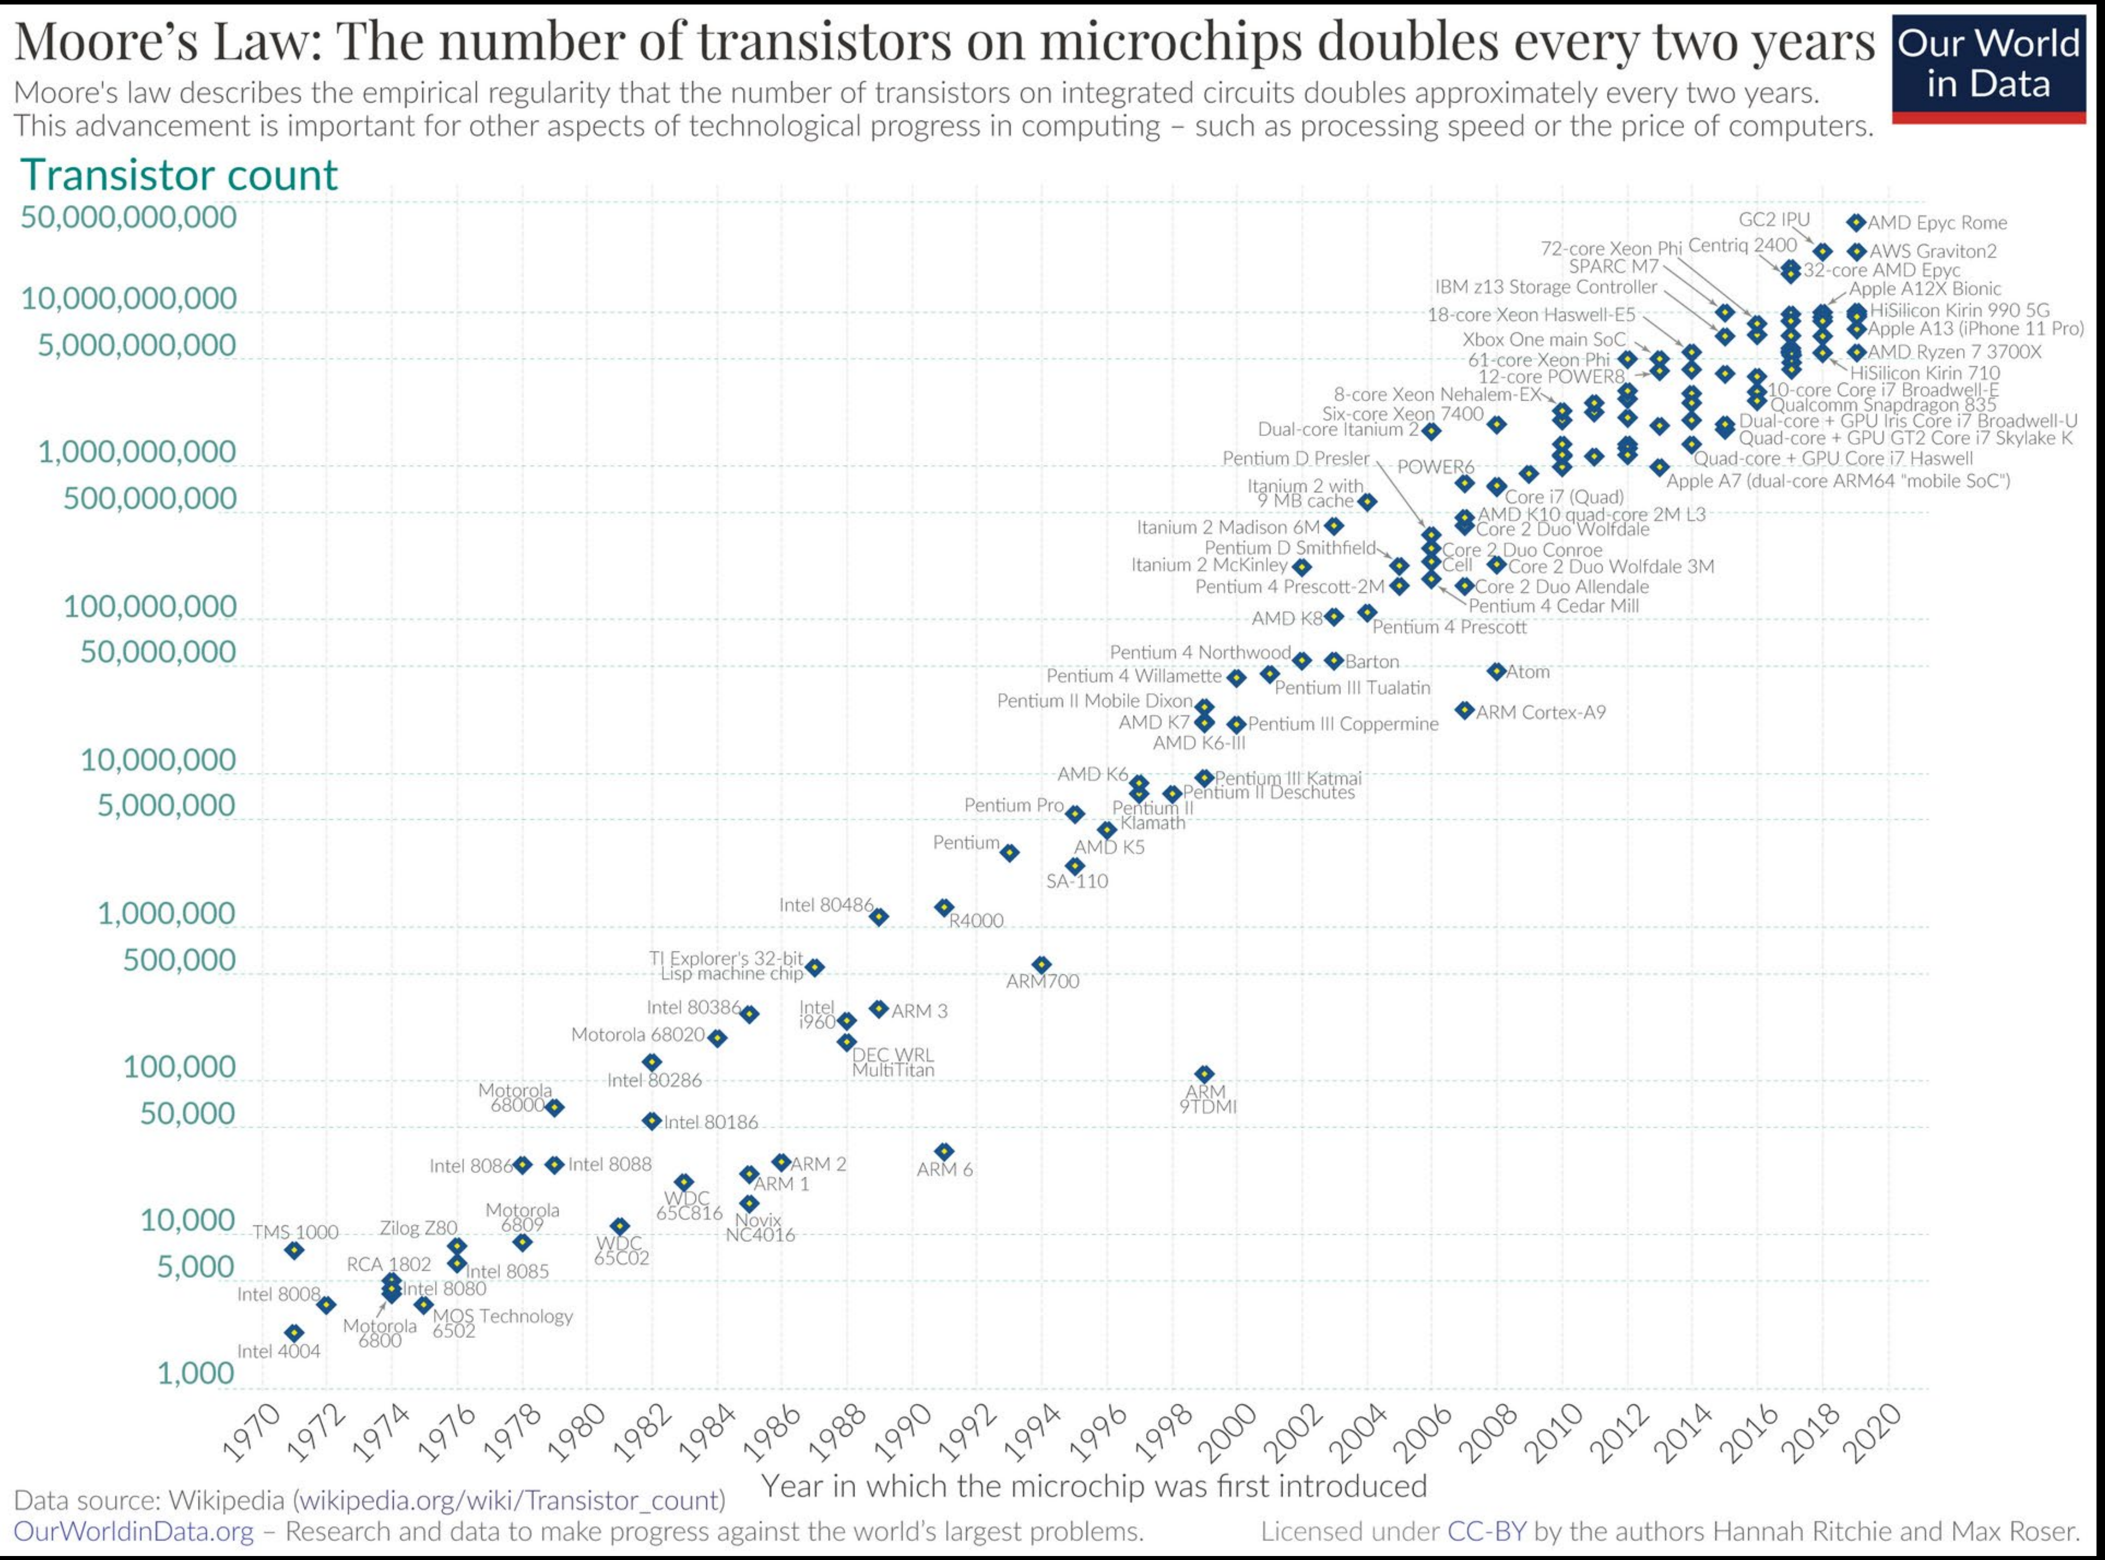

GODDARD
EARTH SCIENCES

MACHINE LEARNING:

- Requires little knowledge of underlying structure
- Can use weak/uninformative priors
- Assumes training data are like test data
- Good in situations where
  - a) we have lots of training data
  - b) that are interchangeable with test data
  - c) we don't need full posteriors (or are OK with bootstrap estimates)

ML USE CASES

Classification

- "Is this event categorically different from others?"

Prediction

- "Given the observed or predicted values of many variables, will there be an extreme event?"

1686 Czech Lands drought: most unusual spatial pattern in pre-industrial European history

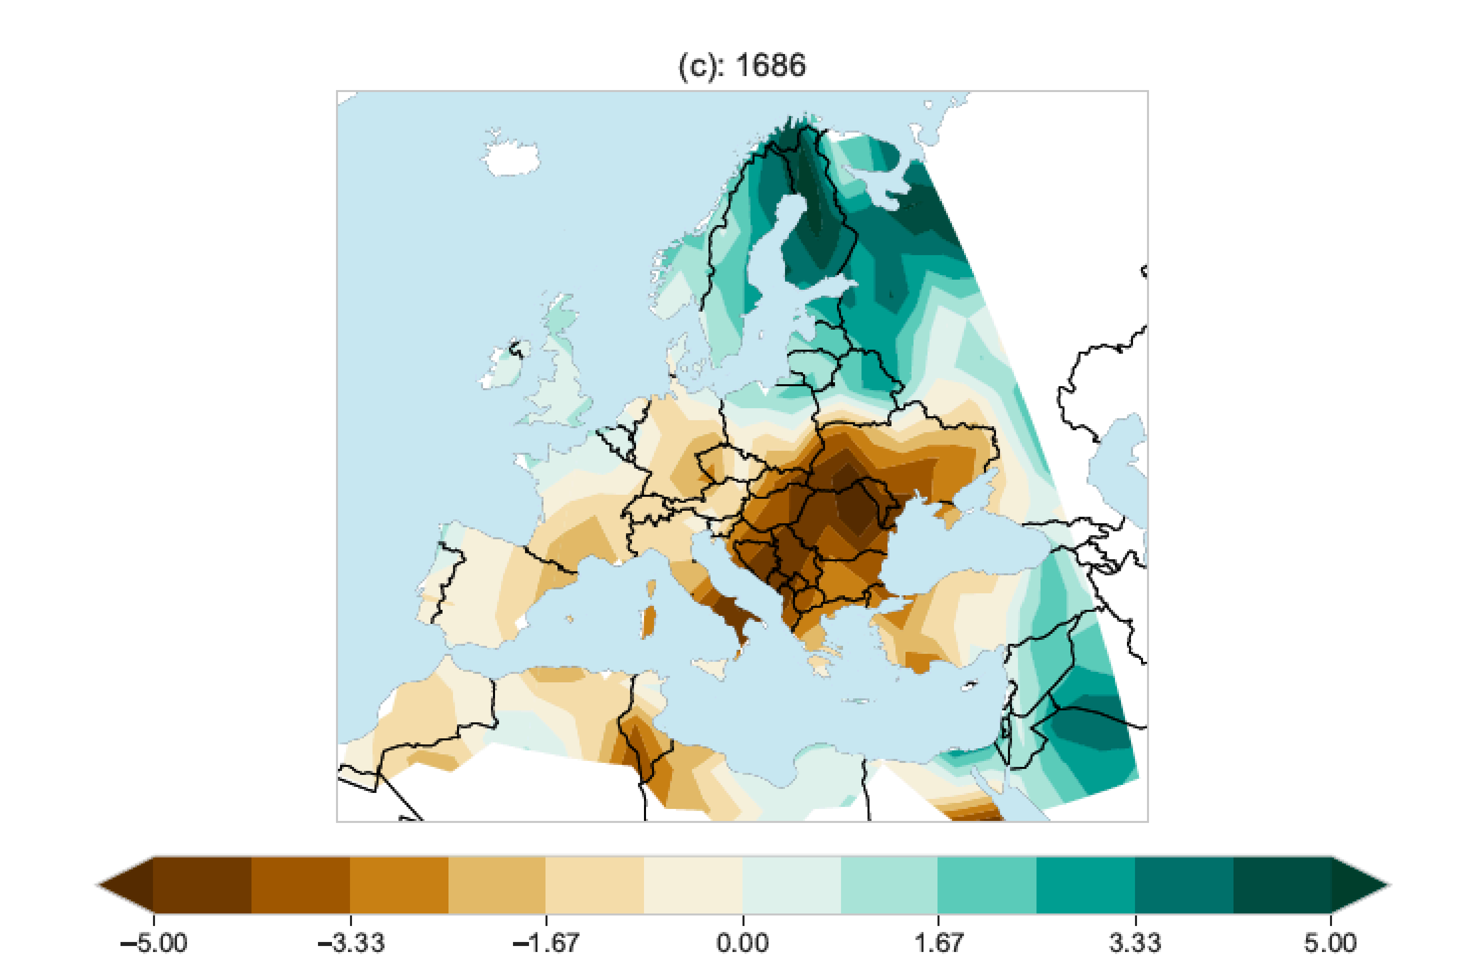

Marvel and Cook Phil Trans A 2022

BAYESIAN INFERENCE

- Requires clearly specified underlying generative model
- Strong priors can incorporate existing knowledge
- Good in situations where
  - a) we have spase training data
  - b) we wish to infer parameter posteriors

BAYES USE CASES

Parameter inference

- “What is forced response and what is internal variability?”

Inversion

- “Given known physical relationships, what underlying (unobserved) drivers contributed most to an observed extreme event?”

“Fingerprint” of external forcing on European/Asian drought risk (MAP value)

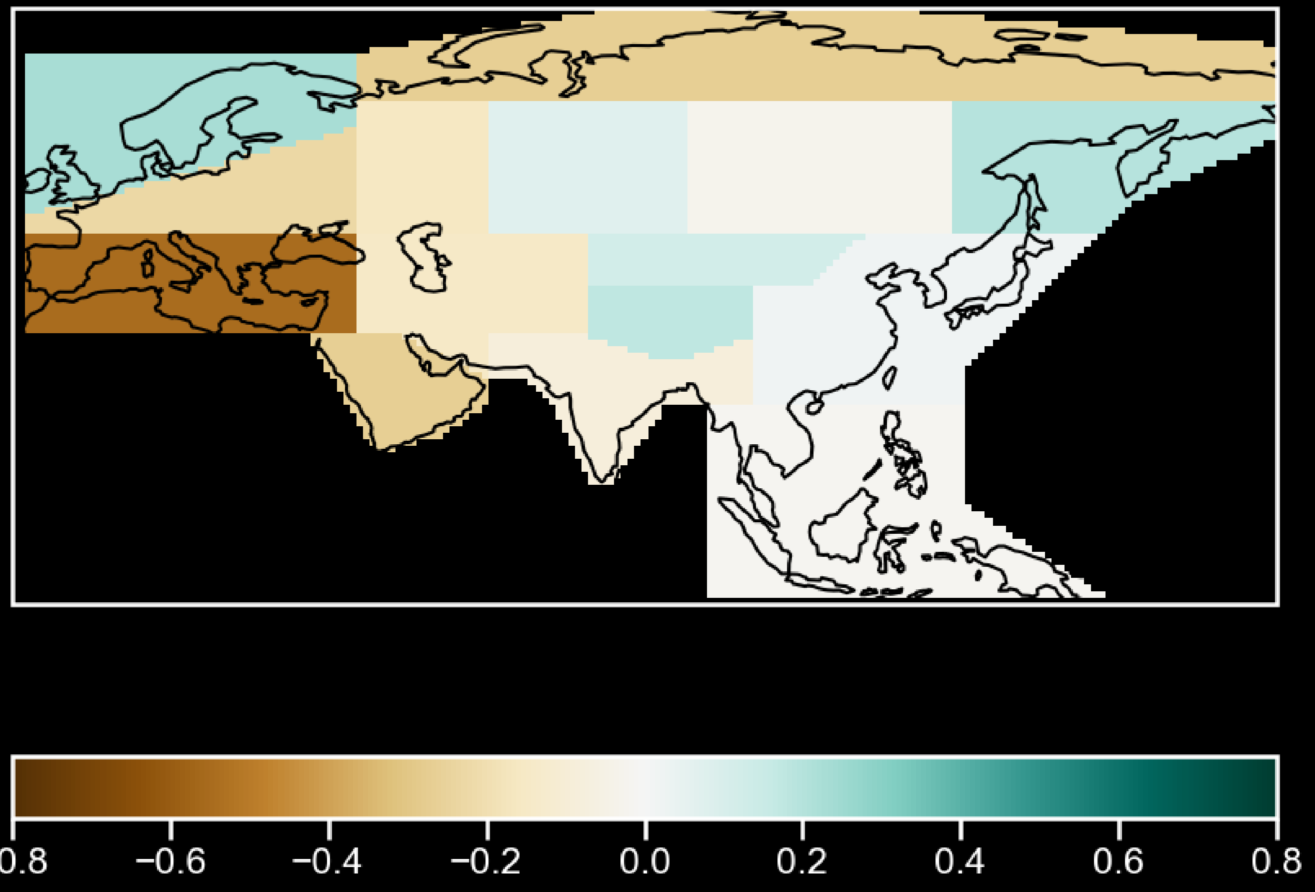

Marvel et al, *AGU Adv.* 2025

MODEL DEPENDENCE

MODEL DEPENDENCE

Likelihood

MODEL DEPENDENCE

Likelihood

Linear regression

Internal variability

Scaling factors

Fingerprint

MODEL DEPENDENCE

Likelihood

Internal variability

Internal variability

GODDARD EARTH SCIENCES

ATTRIBUTION AS MODEL COMPARISON

Every attribution statement is implicitly a model evaluation statement

How well does the model fit the observed data?

Occam's razor: penalized for more parameters
(unless we have good prior knowledge of what they are)
(OR they're very useful in generating D)

ATTRIBUTION AS MODEL COMPARISON

Every attribution statement requires unavoidable choices about different models

Calculate from data

Specify as prior

SUMMARY

MACHINE LEARNING

- Use in cases where it's not necessary or possible to write down the underlying generative model
- Requires abundant training data
- Requires known connection between test and training data
- Good for outlier detection, prediction
- Like all statistical methods, relies on parameters and models (mostly implicit)

BAYESIAN INFERENCE

- Use in cases where we want the model/priors to do most of the work
- Good for sparse data
- Requires the analyst to be explicit about models and prior knowledge
- Useful tools for model criticism and evaluation

Next Slides (Leach)