Machine Learning Pipeline

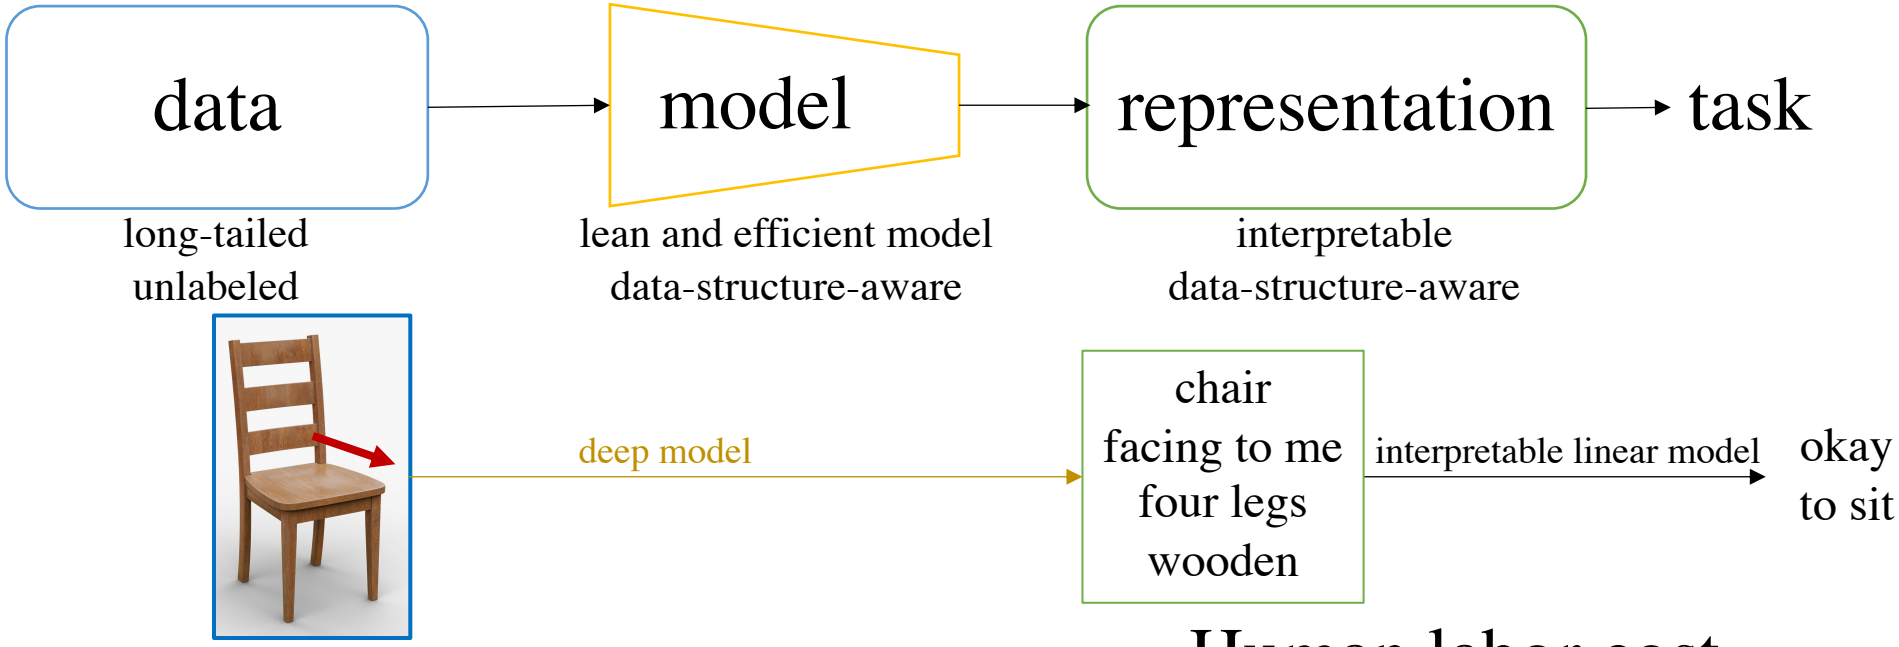

Goals: High accuracy with low cost { Human labor cost / Computation cost

Interpretability

Data-Driven Structure-Aware Learning

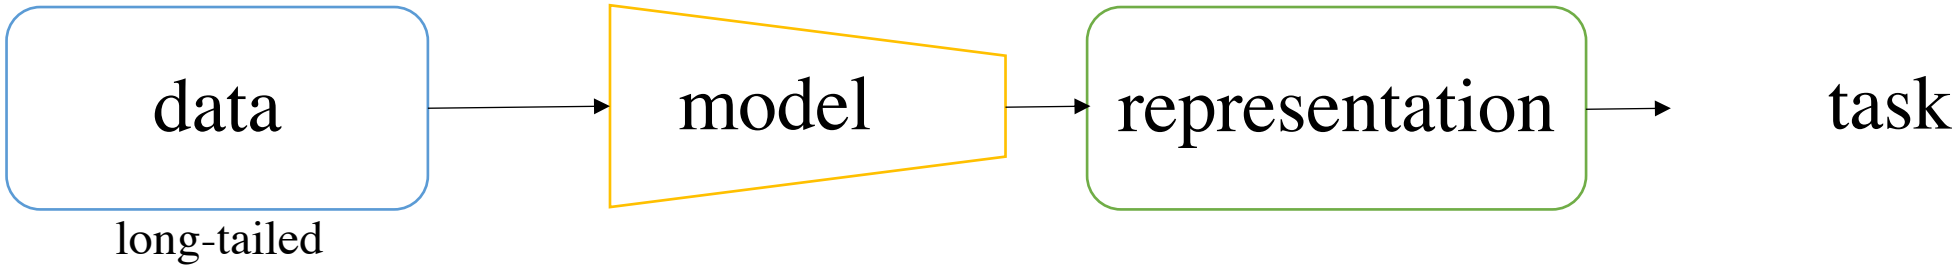

long-tailed

ideal

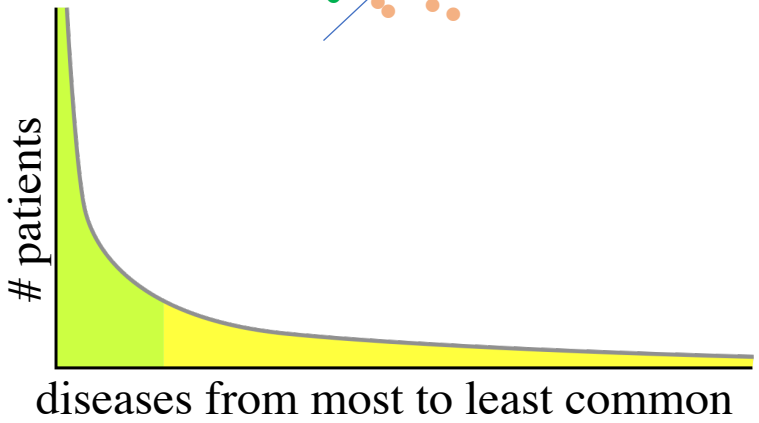

Data-Driven Structure-Aware Learning

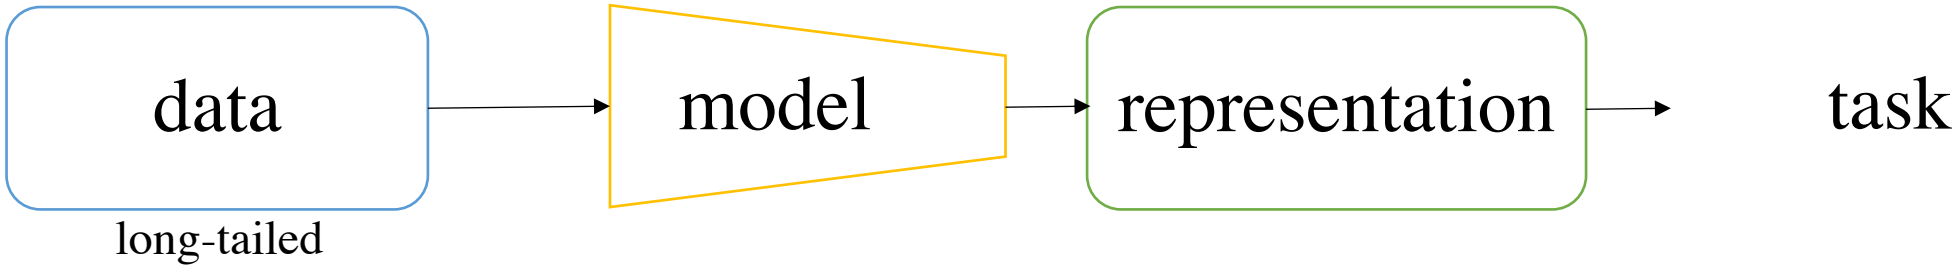

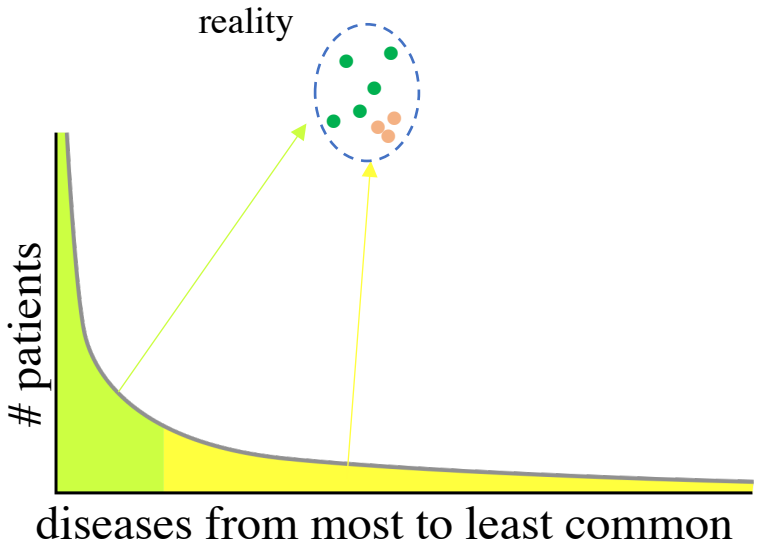

Structure-Aware Learning for Real-World Applications

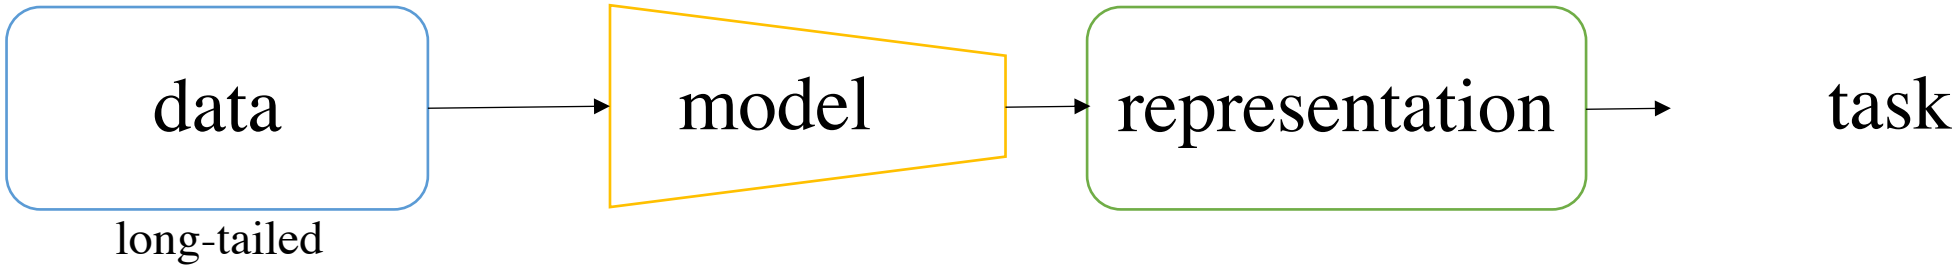

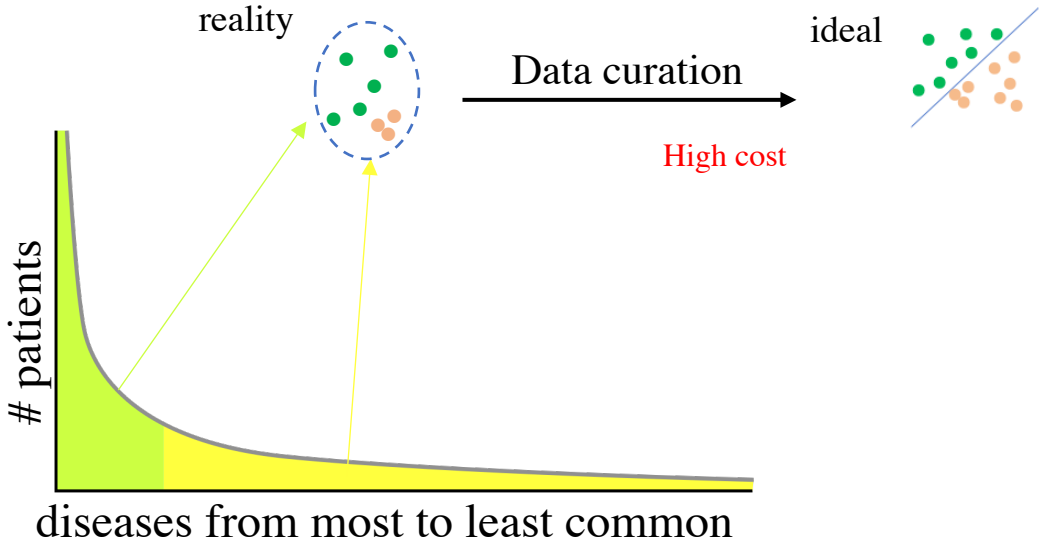

Better model for imbalanced data?

Data-Driven Structure-Aware Learning

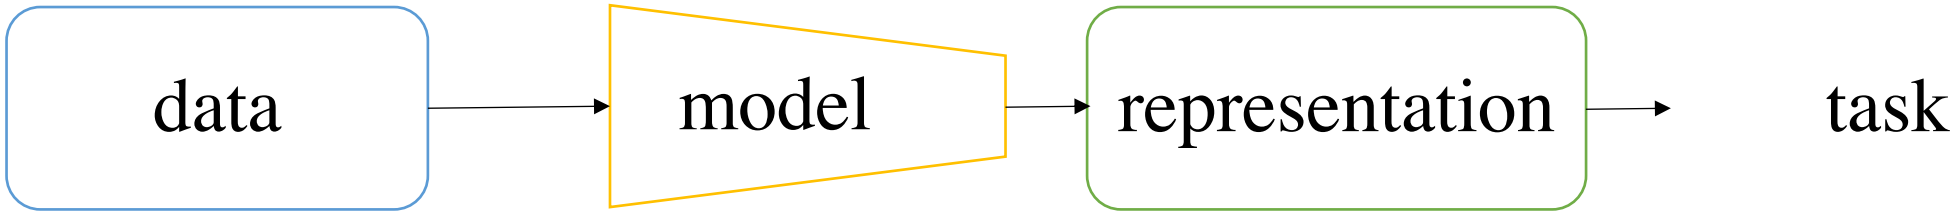

long-tailed

unlabeled

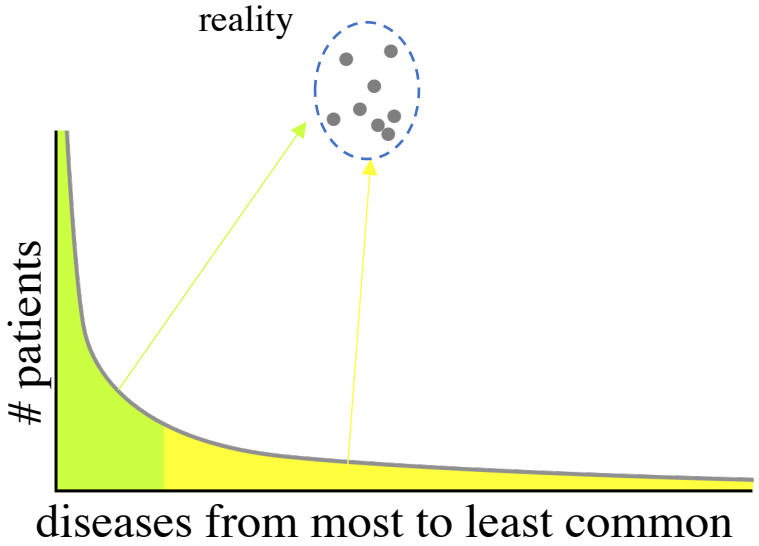

- Labeling large-scale data takes long
- Labels could be ambiguous
- Labels reflect human bias

Data-Driven Structure-Aware Learning

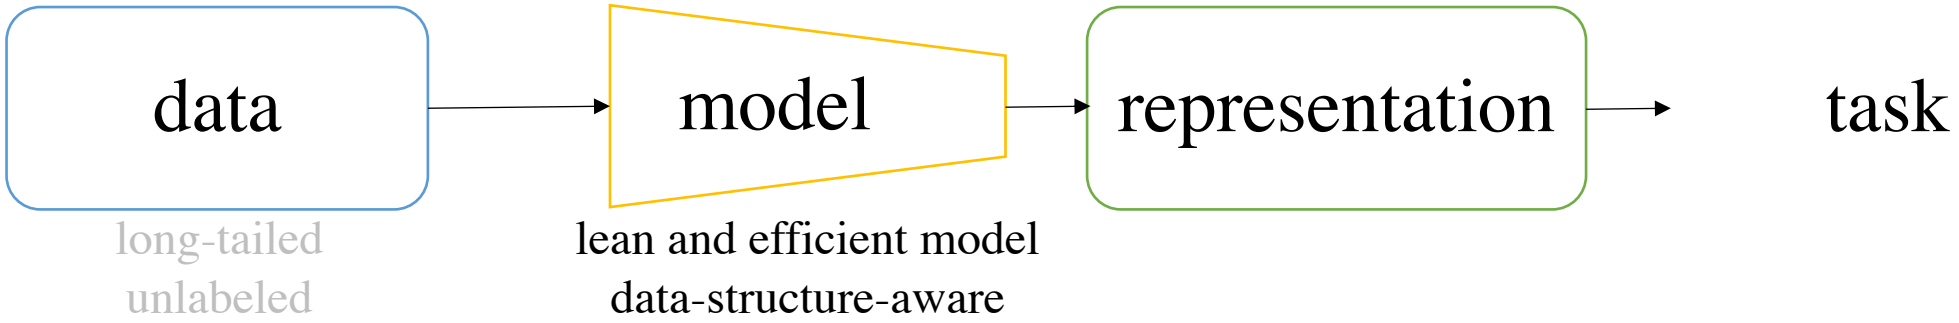

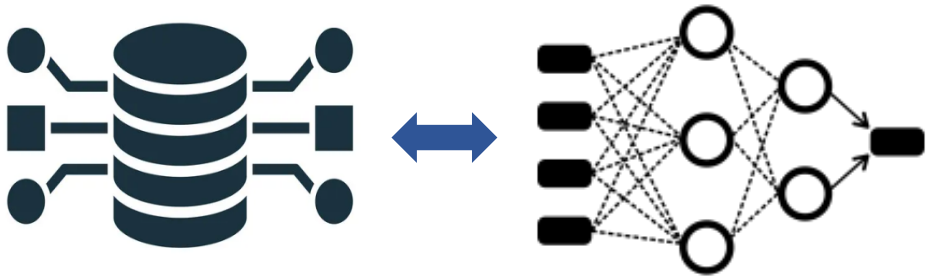

Constrained neural optimization:

- Higher performance for small data
- Faster training and inference time

- Prior information on structures requires less training data
- Small models → faster, lower computation cost and latency

Data-Driven Structure-Aware Learning

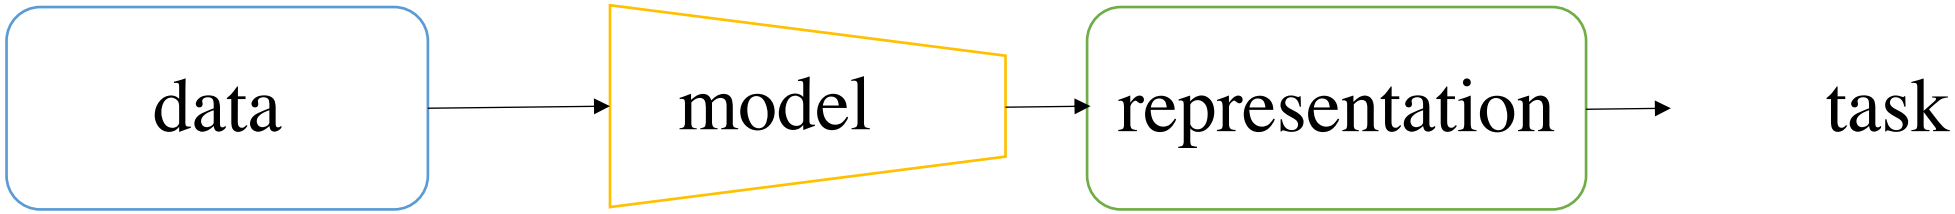

Deep models are hard to interpret

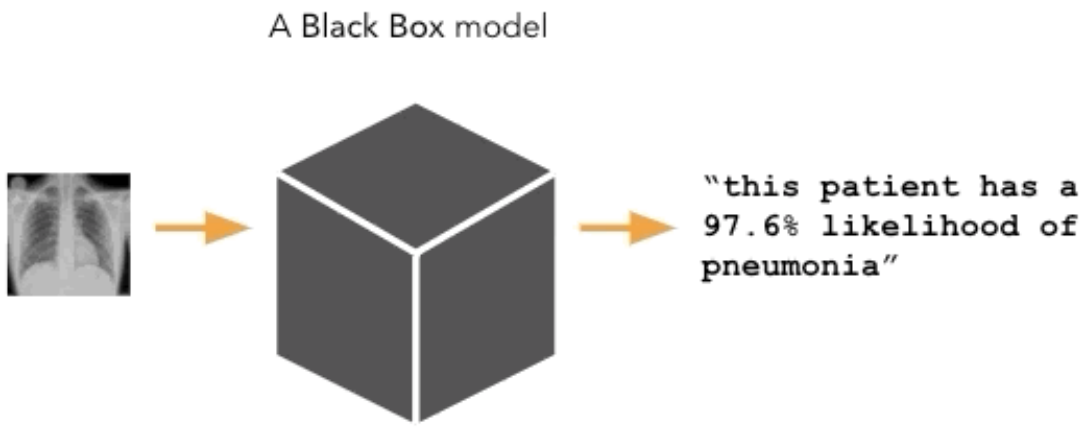

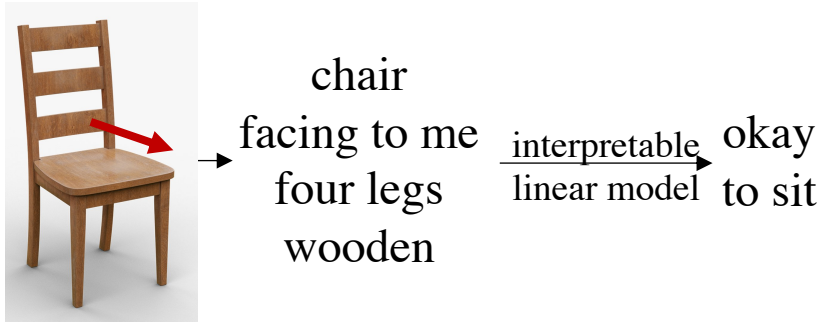

Summary

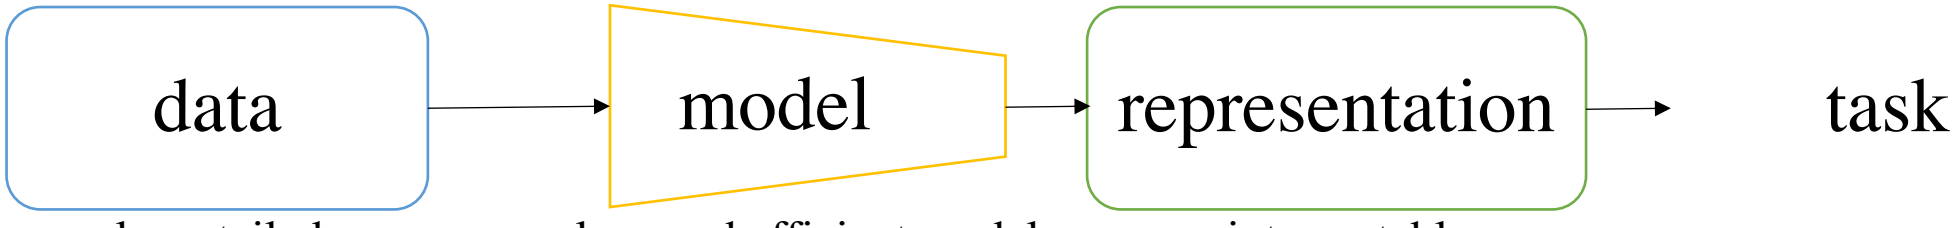

long-tailed
unlabeled

lean and efficient model
data-structure-aware

interpretable
data-structure-aware

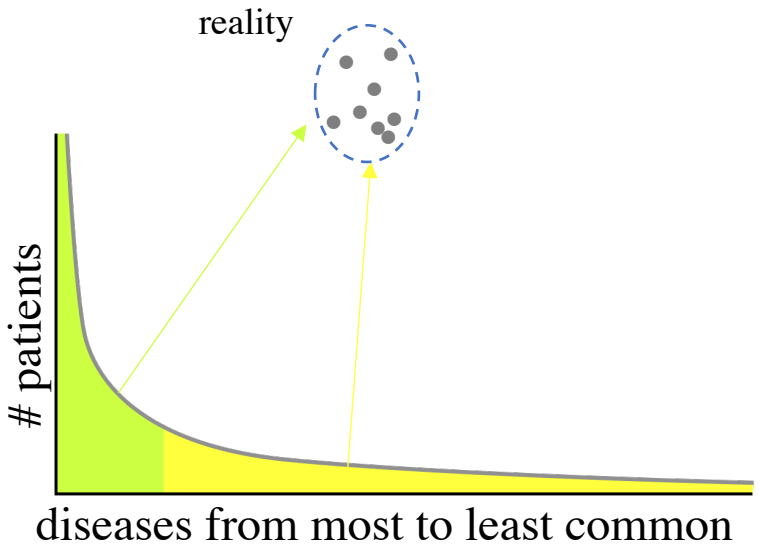

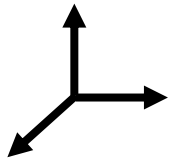

orthogonality

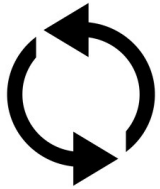

recurrence

low latency -> deployment

chair
facing to me
four legs
wooden

interpretable
linear model

okay
to sit