Learning to Learn with Compound HD Models

Russ Salakhutdinov

Department of Statistics,
University of Toronto

Joshua Tenenbaum

Brain and Cognitive Sciences
MIT

Antonio Torralba

CSAIL, MIT

Hierarchical-Deep Models

HD Models: Integrate hierarchical Bayesian models with deep networks.

One-Shot Learning

Super-category

Hierarchical Bayes:

- Learn **hierarchies of categories** for sharing abstract knowledge.
- Explicitly **share parameters** that are relevant to learning new concept.

Shared higher-level features

Deep Networks:

- Learn **hierarchies of features.**
- **Unsupervised feature learning** – no need to rely on human-crafted input features.
- **Distributed representations.**

Shared low-level features

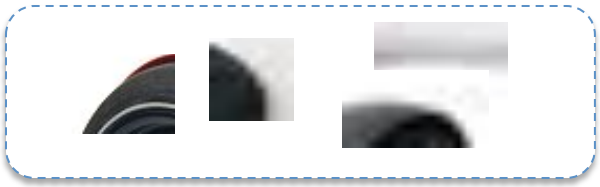

Hierarchical Generative Model

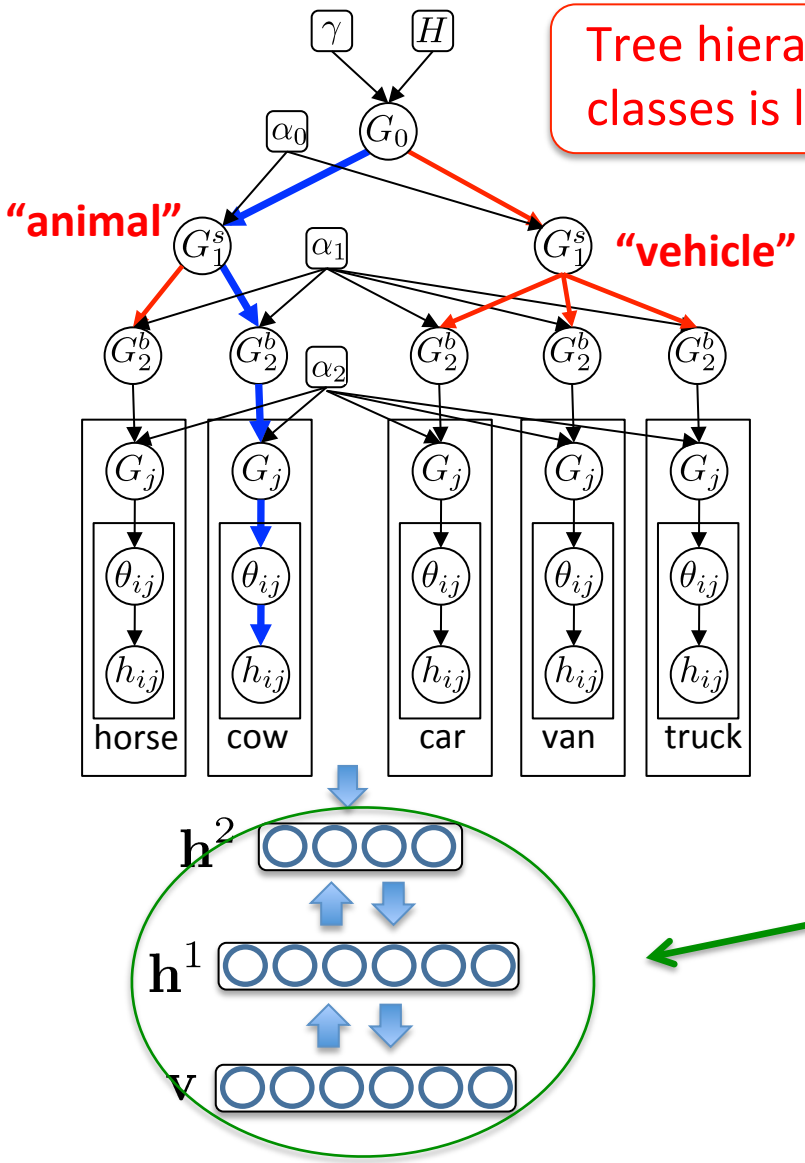

$\mathbf{z} \sim \text{nCRP}$ (**Nested Chinese Restaurant Process**) prior: a nonparametric prior over tree structures

$\mathbf{h}^3|\mathbf{z} \sim \text{HDP}$ (**Hierarchical Dirichlet Process**) prior: a nonparametric prior allowing categories to share higher-level features, or parts.

$\mathbf{v}|\mathbf{h}^3 \sim \text{DBM}$ **Deep Boltzmann Machine**

Enforce (approximate) global consistency through many local constraints.

Images, Handwritten characters, Motion capture datasets.

Learning to Learn from Few Examples

Learning from 3 examples

Training Examples (by row)

Conditional Samples

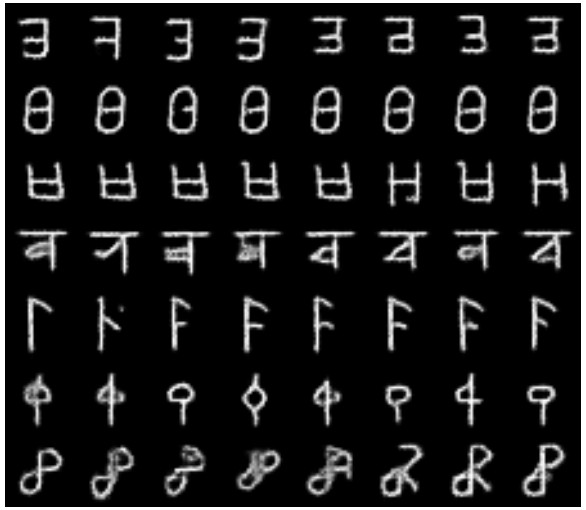

Generating Novel Characters

Learned Super-Classes (by row)

Sampled Novel Characters