



Memory-Enhanced CSNs for One-Shot Learning on Handwritten Character Recognition

Giancarlo Ricci

CS 231N Final Project, Stanford University

Background

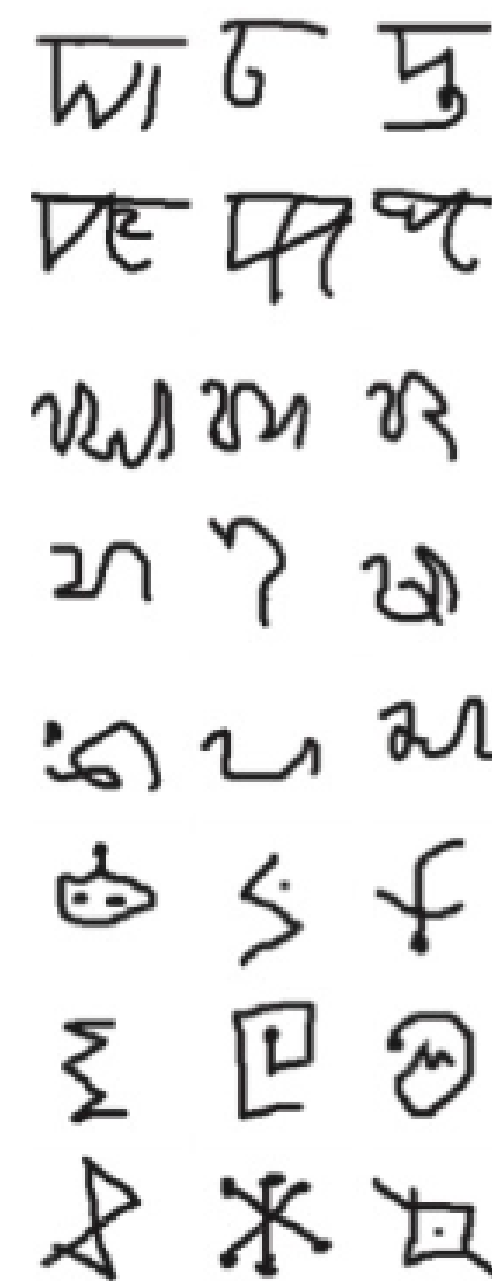
Goal: Apply deep learning techniques with little training data.

Method: Memory-Enhanced Convolutional Siamese Neural Network (CSN)

Dataset

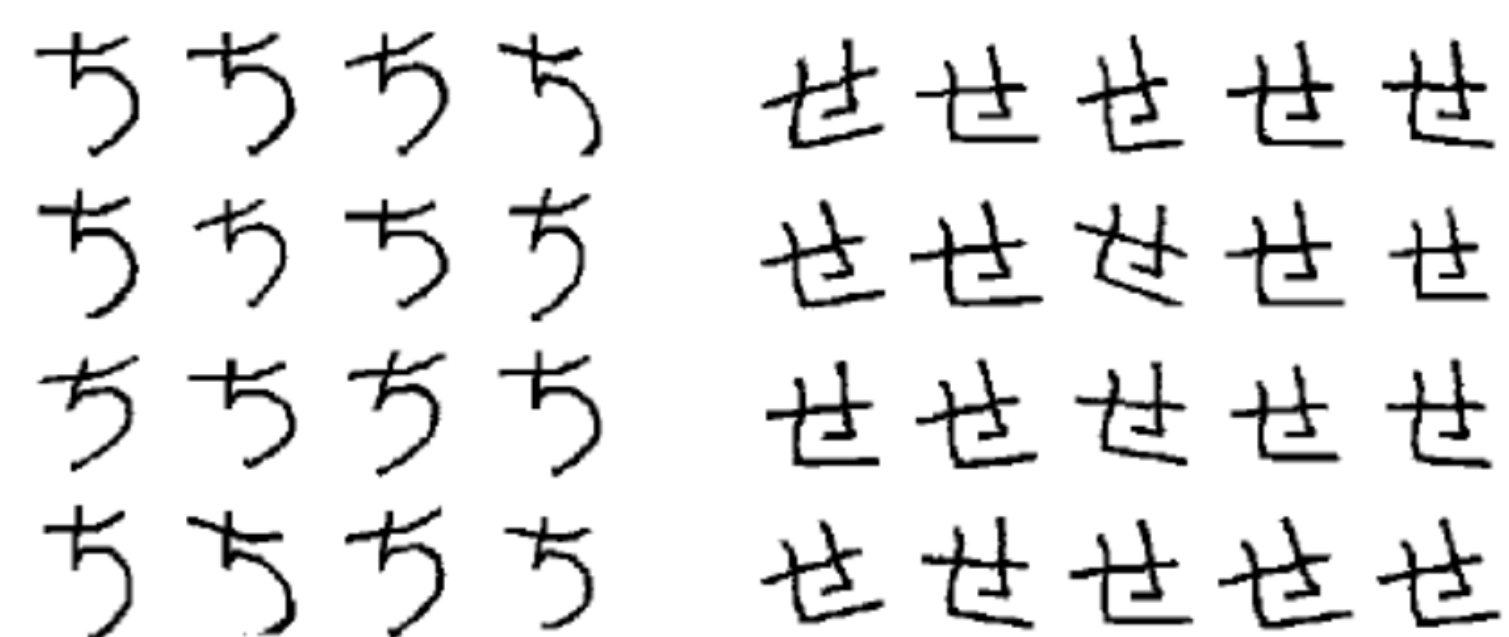
Omniglot dataset

- 1623 characters
- 50 alphabets
- 30 training, 20 validation
- 15-40 chars per alphabet



Data Augmentation


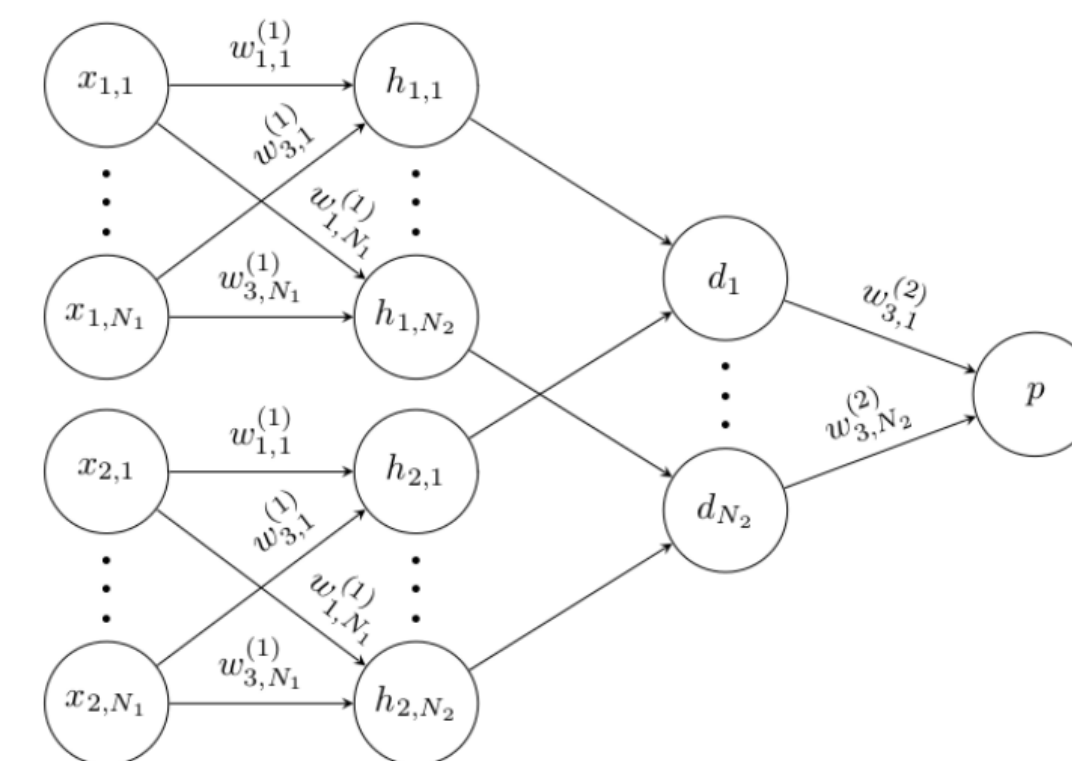
- Affine Distortions
- Translations
- Rotation



Methods

Siamese Neural Network

- Twin Networks (CNN)
- Contrastive Loss
- Intuitively: Similarity

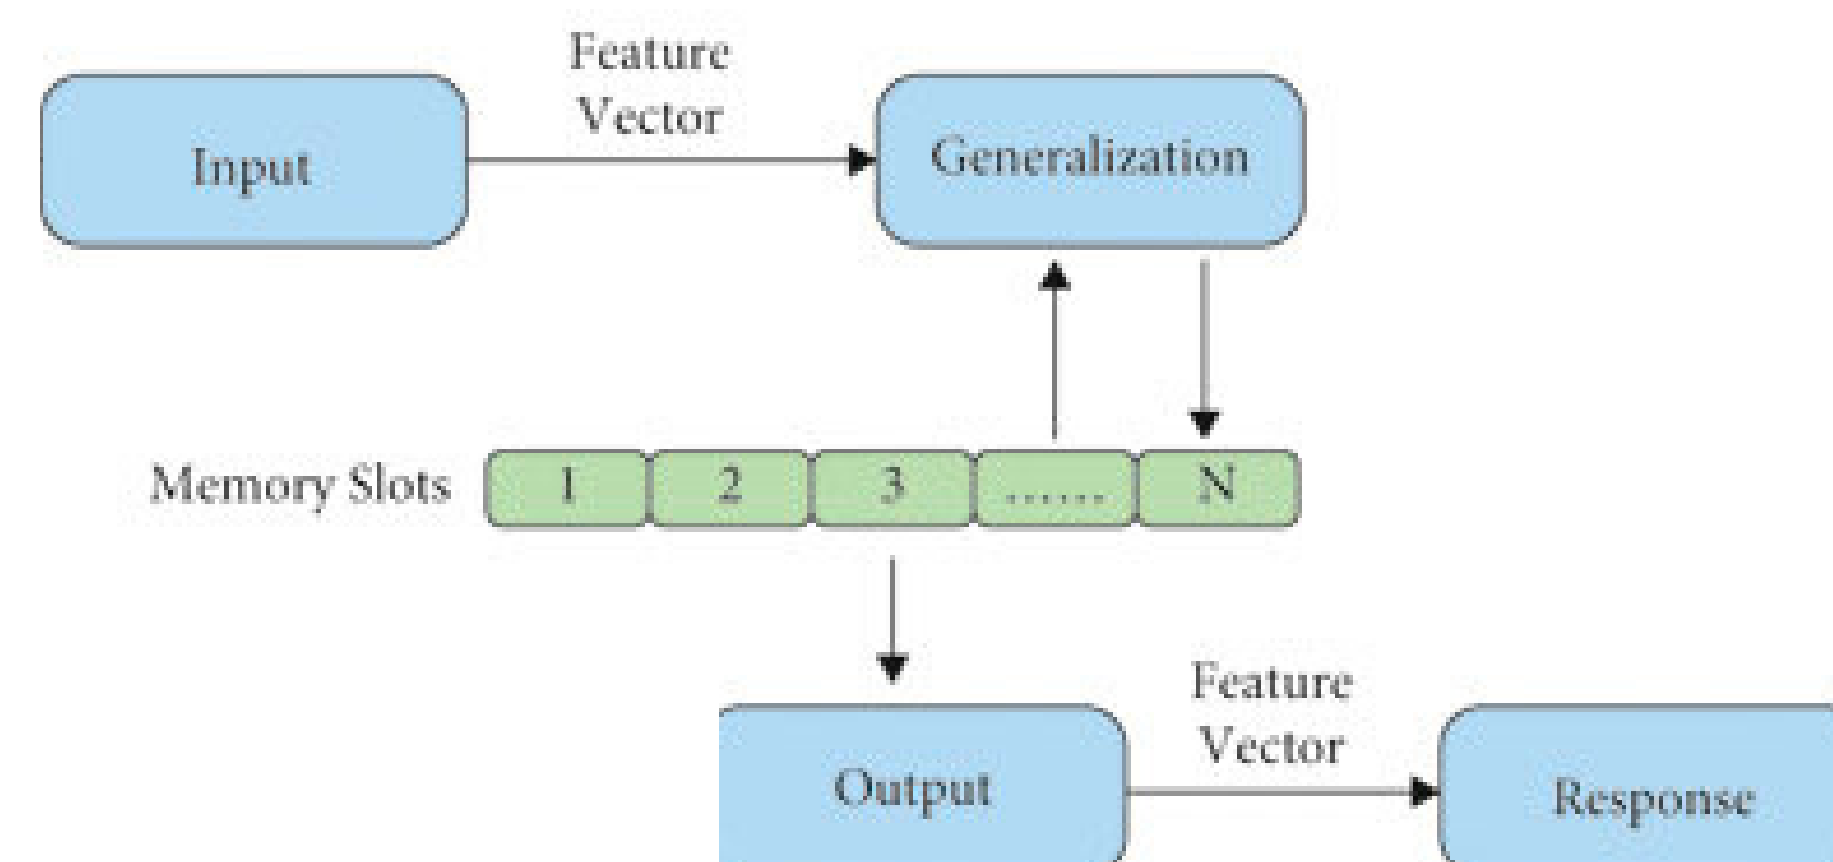


CSN Architecture

- Max-pooling Layers
- ReLU in Twin CNN
- Small Random Weight Init.
- NO Batch-norm
- Sigmoid in Final Layer
- SGD with Weight Decay

Memory Module

- Learn & Remember Long-term Dependencies
- Query Final Output to the Memory Module



Experiments

CSN Architecture	Accuracy
Koch et al.	89.9
Inspired by Chakrapani et al.	87.7
Inspired by Santoro et al.	88.8

Batchnorm	Accuracy
Vanilla	89.9
P = 0.5	89.8
P = 0.8	88.6

Data Augmentation	Accuracy
Vanilla	89.9
Rotation	83.5
Affine Distortion	90.5

Optimizers	Accuracy
SGD	90.3
Adam	90.2
SGD + Weight Decay	90.5

Memory Module	Accuracy
Vanilla	90.5
Memory	93.4