


Machine Learning in Econometrics: Lecture 12

INSTRUCTOR: CHAOYI CHEN
NJE & MNB

November 21, 2023

@copyright Chaoyi Chen (NJE & MNB). All rights reserved. Please do not distribute without express written consent.

Topic 5: How CNNs Work



- The CNN builds up an image in a **hierarchical** fashion.
- Edges and shapes are recognized and pieced together to form more complex shapes, eventually assembling the target image.
- This hierarchical construction is achieved using **convolution** and **pooling** layers.


Topic 5: Convolution Filter

$$\text{Input Image} = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \\ j & k & l \end{bmatrix}, \quad \text{Convolution Filter} = \begin{bmatrix} \alpha & \beta \\ \gamma & \delta \end{bmatrix}$$

$$\text{Convolved Image} = \begin{bmatrix} a\alpha + b\beta + d\gamma + e\delta & b\alpha + c\beta + e\gamma + f\delta \\ d\alpha + e\beta + g\gamma + h\delta & e\alpha + f\beta + h\gamma + i\delta \\ g\alpha + h\beta + j\gamma + k\delta & h\alpha + i\beta + k\gamma + l\delta \end{bmatrix}$$

- The filter is itself an image, and represents a small shape, edge etc.
- We slide it around the input image, scoring for matches.
- The scoring is done via [dot-products](#), illustrated above.
- If the subimage of the input image is similar to the filter, the score is high, otherwise low.
- The filters are [learned](#) during training.

Topic 5: Convolution Example




- The idea of convolution with a filter is to find common patterns that occur in different parts of the image.
- The two filters shown here highlight vertical and horizontal stripes.
- The result of the convolution is a new feature map.
- Since images have three color channels, the filter does as well: one filter per channel, and dot-products are summed.
- The weights in the filters are **learned** by the network.

Topic 5: Pooling

$$\text{Max Pool} \begin{bmatrix} 1 & 2 & 5 & 3 \\ 3 & 0 & 1 & 2 \\ 2 & 1 & 3 & 4 \\ 1 & 1 & 2 & 0 \end{bmatrix} \longrightarrow \begin{bmatrix} 3 & 5 \\ 2 & 4 \end{bmatrix}$$

- Each non-overlapping 2×2 block is replaced by its maximum.
- This sharpens the feature identification.
- Allows for locational invariance.
- Reduces the dimension by a factor of 4 — i.e. factor of 2 in each dimension.

Topic 5: Architecture of a CNN




- Many convolve + pool layers.
- Filters are typically small, e.g. each channel 3×3 .
- Each filter creates a new channel in convolution layer.
- As pooling reduces size, the number of filters/channels is typically increased.
- Number of layers can be very large. E.g. **resnet50** trained on **imagenet** 1000-class image data base has 50 layers!

Topic 5: Data Augmentation



- Each training image is replicated many times - with each mentation replicate randomly distorted in a natural way such that human recognition is unaffected.
- Typical distortions are zoom, horizontal and vertical shift, shear, small rotations, and in this case horizontal flips.
- **Increasing** the training set considerably - using different examples to protect against overfitting
- A form of **regularization**: This kind of fattening of the data is similar in spirit to ridge regularization.- build a cloud of images around each original images, all with the same label.

Topic 5: Using Pretrained Networks to Classify Images



flamingo

Cooper's hawk

Cooper's hawk

flamingo	0.83	kite	0.60	fountain	0.35
spoonbill	0.17	great grey owl	0.09	nail	0.12
white stork	0.00	robin	0.06	hook	0.07
Lhasa Apso		cat		Cape weaver	
Tibetan terrier	0.56	Old English sheepdog	0.82	jacamar	0.28
Lhasa	0.32	Shih-Tzu	0.04	macaw	0.12
cocker spaniel	0.03	Persian cat	0.04	robin	0.12

- Here we use the 50-layer **resnet50** network trained on the 1000-class **imagenet** corpus to classify some photographs.


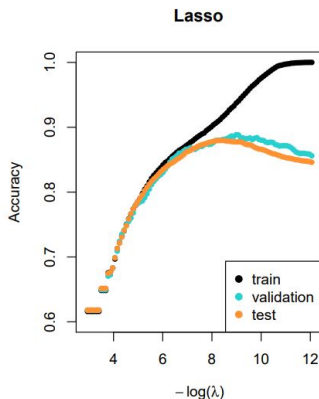
Topic 5: Document Classification: IMDB Movie Reviews

- The **IMDB** corpus consists of user-supplied movie ratings for a large collection of movies. Each has been labeled for sentiment as **positive** or **negative**. Here is the beginning of a negative review:
“This has to be one of the worst films of the 1990s. When my friends & I were watching this film (being the target audience it was aimed at) we just sat & watched the first half an hour with our jaws touching the floor at how bad it really was. The rest of the time, everyone else in the theater just started talking to each other, leaving or generally crying into their popcorn ...”
- We have labeled training and test sets, each consisting of 25,000 reviews, and each balanced with regard to sentiment.
- We wish to build a classifier to predict the sentiment of a review.

Topic 5: Featurization - Bag-of-Words

- Documents have different lengths, and consist of sequences of words. How do we create features X to characterize a document?
 - From a dictionary, identify the $10K$ most frequently occurring words.
 - Create a binary vector of length $p = 10K$ for each document, and score a 1 in every position that the corresponding word occurred.
 - With n documents, we now have a $n \times p$ *sparse* feature matrix X .
 - We compare a lasso logistic regression model to a two-hidden-layer neural network on the next slide. (No convolutions here!)
 - Bag-of-words are *unigrams*. We can instead use *bigrams* (occurrences of adjacent word pairs), and in general *m-grams*.

Topic 5: Lasso versus Neural Network — IMDB Reviews




- Simpler lasso logistic regression model works as well as neural network in this case.
- `glmnet` was used to fit the lasso model, and is very effective because it can exploit sparsity in the X matrix.

Topic 5: Models for sequences of data - Recurrent Neural Networks

- Often data arise as sequences:
 - Documents are sequences of words, and their relative positions have meaning.
 - Time-series such as weather data or financial indices.
 - Recorded speech or music.
 - Handwriting, such as doctor's notes.
- RNNs build models that take into account this sequential nature of the data, and build a memory of the past.
 - The feature for each observation is a *sequence* of vectors $X = \{X_1, X_2, \dots, X_L\}$.
 - The target Y is often of the usual kind — e.g. a single variable such as **Sentiment**, or a one-hot vector for multiclass.
 - However, Y can also be a sequence, such as the same document in a different language.

Topic 5: Simple Recurrent Neural Network Architecture



- The hidden layer is a sequence of vectors A_ℓ , receiving as input X_ℓ as well as $A_{\ell-1}$. A_ℓ produces an output O_ℓ .
- The *same* weights \mathbf{W} , \mathbf{U} and \mathbf{B} are used at each step in the sequence — hence the term *recurrent*.
- The A_ℓ sequence represents an evolving model for the response that is updated as each element X_ℓ is processed.

Topic 5: RNN in Detail

- Suppose $X_\ell = (X_{\ell 1}, X_{\ell 2}, \dots, X_{\ell p})$ has p components, and $A_\ell = (A_{\ell 1}, A_{\ell 2}, \dots, A_{\ell K})$ has K components. Then the computation at the k th components of hidden unit A_ℓ is

$$A_{\ell k} = g \left(w_{k0} + \sum_{j=1}^p w_{kj} X_{\ell j} + \sum_{s=1}^K A_{\ell-1, s} \right)$$

$$O_\ell = \beta_0 + \sum_{k=1}^K \beta_k A_{\ell k}$$

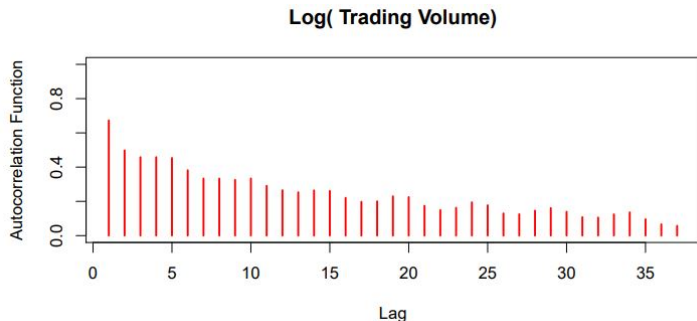
- Often we are concerned only with the prediction O_L at the last unit. For squared error loss, and n sequence/response pairs, we would minimize

$$\sum_{i=1}^n (y_i - o_{iL})^2 = \sum_{i=1}^n \left(y_i - \left(\beta_0 + \sum_{k=1}^K \beta_k g \left(w_{k0} + \sum_{j=1}^p w_{kj} x_{iLj} + \sum_{s=1}^K u_{ks} \alpha_{i, L-1, s} \right) \right) \right)^2$$

Topic 5: New-York Stock Exchange Data

- Shown in the previous slide are three daily time series for the period December 3, 1962 to December 31, 1986 (6,051 trading days):
 - *Log trading volume*. This is the fraction of all outstanding shares that are traded on that day, relative to a 100-day moving average of past turnover, on the log scale.
 - *Dow Jones return*. This is the difference between the log of the Dow Jones Industrial Index on consecutive trading days.
 - *Log volatility*. This is based on the absolute values of daily price movements.
- **Goal:** predict *Log trading volume* tomorrow, given its observed values up to today, as well as those of *Dow Jones return* and *Log volatility*.

Topic 5: Autocorrelation



- The *autocorrelation* at lag ℓ is the correlation of all pairs $(v_t, v_{t-\ell})$ that are ℓ trading days apart.
- These sizable correlations give us confidence that past values will be helpful in predicting the future.
- This is a curious prediction problem: the response v_t is also a feature $v_{t-\ell}$!

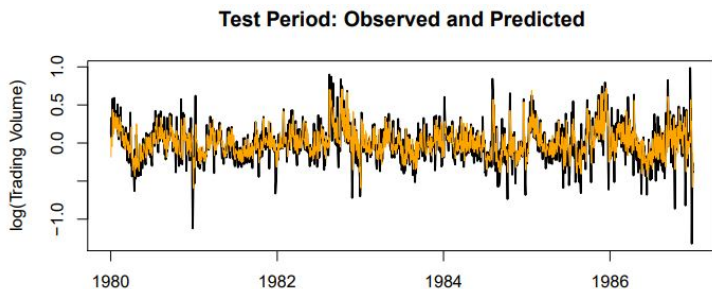
Topic 5: RNN Forecaster

- We only have one series of data! How do we set up for an RNN?
- We extract many short mini-series of input sequences
 $X = \{X_1, X_2, \dots, X_L\}$ with a predefined length L stands for the lag:

$$X_1 = \begin{bmatrix} v_{t-L} \\ r_{t-L} \\ z_{t-L} \end{bmatrix}, X_2 = \begin{bmatrix} v_{t-L+1} \\ r_{t-L+1} \\ z_{t-L+1} \end{bmatrix}, \dots, X_L = \begin{bmatrix} v_{t-1} \\ r_{t-1} \\ z_{t-1} \end{bmatrix}, \text{ and } Y = v_t$$

- Since $T = 6051$, with $L = 5$ we can create 6046 such (X, Y) pairs.
- We use the first 4281 as training data, and the following 1770 as test data. We fit an RNN with 12 hidden units per lag step (i.e. per A_ℓ .)

Topic 5: RNN Results for NYSE Data



- Figure shows predictions and truth for test period.
- $R^2 = 0.42$ for RNN
- $R^2 = 0.18$ for a random walk - use yesterday's value of Log trading volume to predict that of today.
- $R^2 = 0.41$ for AR(5) model (16 parameters).

Topic 5: When to Use Deep Learning

- CNNs have had enormous successes in image classification and modeling, and are starting to be used in medical diagnosis. Examples include digital mammography, ophthalmology, MRI scans, and digital X-rays.
- RNNs have had big wins in speech modeling, language translation, and forecasting.
- Should we always use deep learning models?
 - Often the big successes occur when the **signal-to-noise** ratio is high — e.g. image recognition and language translation. Datasets are large, and overfitting is not a big problem
 - For noisier data, simpler models can often work better.
 - On the **NYSE** data, the AR(5) model is much simpler than a RNN, and performed as well.
 - On the **IMDB** review data, the linear model fit by **glmnet** did as well as the neural network (and better than RNN - check textbook for more reference)
- We endorse the **Occam's razor** principal — we prefer simpler models if they work as well. More interpretable!