Time Series and Sequential Data

Volker Tresp

Winter 2023-2024

Time Series Modelling

- I have to predict the total energy consumption of a city for tomorrow, based on certain inputs (weather forecast: temperature, precipitation, wind; then: working day/holiday, ...)
- It would help to also consider the energy consumption of today and maybe of yesterday as inputs
- Added complexity: If my goal is to predict the energy consumption two days in the future, my own prediction for tomorrow becomes an input for the prediction for two days in the future
- In time series modelling, outputs, and often also the inputs, are real numbers

DAX Performance-Index





Sequence Modelling

- Sequence classification: The input is a sentence, i.e., a sequence of words; the output classifies the sentiment of the sentence
- Encoder-decoder modelling: The input is a sentence, i.e., a sequence of words, in English; the output is the sentence translated into German
- In sequence modelling, inputs and outputs are typically discrete

I. Time Series Modelling: NARX Models

Neural Networks for Time-Series Modelling

- Let $y_t, t = 1, 2, ...$ be the time-discrete time-series of interest (example: DAX)
- Let $x_t, t = 1, 2, ...$ denote a second time-series, that contains information on y_t (Example: Dow Jones)
- For simplicity, we assume that both y_t and x_t are scalars. The goal is the prediction of the next value of the time-series
- We assume a system of the form

$$y_t = f(y_{t-1}, \dots, y_{t-T}, x_{t-1}, \dots, x_{t-T}) + \epsilon_t$$

with i.i.d. random numbers $\epsilon_t, t = 1, 2, \ldots$ which model unknown disturbances

Neural Networks for Time-Series Modelling (cont'd)

• We approximate the function, using a neural network,

$$f(y_{t-1}, \dots, y_{t-T}, x_{t-1}, \dots, x_{t-T})$$
$$\approx f_{\mathbf{w}, \mathbf{V}}(y_{t-1}, \dots, y_{t-T}, x_{t-1}, \dots, x_{t-T})$$

• A reasonable cost function is

$$cost(\mathbf{w}, \mathbf{V}) = \sum_{t=1}^{N} (y_t - f_{\mathbf{w}, \mathbf{V}}(y_{t-1}, \dots, y_{t-T}, x_{t-1}, \dots, x_{t-T}))^2$$

Neural Networks for Time-Series Modelling (cont'd)

- It is important to note, that the neural network can be trained as before with simple back propagation *if in training all* y_t *and all* x_t *are known*!
- This model is called a NARX model: Nonlinear Auto Regressive Model with external inputs. Another name: TDNN (time-delay neural network)
- Note the "convolutional" idea in TDNNs: the same neural network is applied in all time instances

Prediction

• For single step prediction, we use

$$\hat{y}_t = f_{\mathbf{w}, \mathbf{V}}(y_{t-1}, \dots, y_{t-T}, x_{t-1}, \dots, x_{t-T})$$

Self-supervised Learning

- The time-series provides its own labels
- No human labelling is necessary: self-supervised learning

Multiple-Step Prediction based on Multiple Step Prediction

• We can also train a model to predict τ time steps into the future; the prediction then becomes

$$\hat{y}_{t+\tau} = f_{\mathbf{w},\mathbf{V}}^{\tau}(y_{t-1},\ldots,y_{t-T},x_{t-1},\ldots,x_{t-T})$$

• This is done in system simulation: the prediction based on detailed system models might be computationally very expensive and cannot be done online; the idea is to train a neural network predictive model off-line and then use that one online instead of an expensive simulation

Mutiple-Step Prediction based on Single-Step Prediction

- Why not just iterate the single-step prediction? One issue is that my prediction is uncertain, so I should consider that uncertainty; second: I do not have future inputs!
- One way is simulation; for y_t we have the model as before, $(t' = t, \ldots, t + \tau)$

$$y_{t'} = f_{\mathbf{w},\mathbf{V}}(y_{t'-1},\ldots,y_{t'-T},x_{t'-1},\ldots,x_{t'-T}) + \epsilon_{t'}$$

- Using both we can generate samples for the future; for the noise I might assume a Gaussian distribution $\epsilon_t \sim \mathcal{N}(0, \sigma^2)$
- Future inputs $x_{t'}$, we either set to zero, or we develop a separate prediction model for those as well
- For multiple-step prediction, we can simulate (i.e., sample) for the desired number of time steps in the future (Monte-Carlo simulation) repeatedly and can derive estimated means, variances, and covariances

Residual Modeling

• We have

$$\hat{y}_t = y_{t-1} + f_{\mathbf{w},\mathbf{V}}(y_{t-1},\ldots,y_{t-T},x_{t-1},\ldots,x_{t-T})$$

• Realize the similarity to ResNet

Considering the Complete History?

• Consider a prediction model that uses the complete history,

$$\widehat{y}_t = f_{\mathbf{w},\mathbf{V}}(y_{t-1},\ldots,y_1,x_{t-1},\ldots,x_1)$$

This means that the time window grows with t: T := t - 1

- Technical solutions:
- 1: Models with an internal memory: RNNs, LSTMs
- 2: Models with the ability to grow: Transformers

GPT-type Architecture

• GPT always considers the complete history

$$\hat{y}_t = f_{\mathbf{w},\mathbf{V}}(y_{t-1},\ldots,y_1)$$

- A generated sequence (text) is a simulation of the future
- The first K steps y_1, \ldots, y_K are the input (prompt) from the user
- y is a discrete variable with as many states as there are tokens (words)

II. Sequence Modelling

Encoding Inputs and Outputs

- So far we considered that x_t is either binary $x_t \in \{0, 1\}$ or continuous, $x_t \in \mathbb{R}$
- How do we encode that $x_t \in \{0, 1, ..., N^{words}\}$, where N^{words} is the number of words in the vocabulary?
- We could consider that $x_t \in \mathbb{R}$ and encode it as a scalar (amplitude encoding): this is not commonly done
- Alternatively, we introduce N^{words} binary variables with $x_{t,j} = \mathbb{I}(word(t) \equiv j)$ (one-hot encoding) ($\mathbb{I}()$ is the indicator function which is equal to 1, when the argument is true, otherwise zero)

(A): One-hot Encoding

- This type of encoding is called one-hot encoding
- We train the input to hidden matrix V, which is an $H \times N^{words}$ dimensional matrix

(B): Embedding Encoding

- Maybe we should represent a word by its attributes? But what attributes?
- An embedding vector \mathbf{a}_i for word i is a vector of abstract attributes that represent the word and which might have been derived from a large vocabulary and is shared between applications
- This embedding vector might have been generated by some other research group and is simply a vector of real numbers of length r (rank)
- Now the input to hidden connection matrix is $\tilde{\mathbf{V}}$ which is an H imes r dimensional matrix

(C): Embedding Encoding in Combination with One-hot Encoding

- Sometimes it is more intuitive to consider a matrix A connecting the one-hot encoded input with the first hidden layer
- The *i*-th column in matrix A contains a_i

Relationship Between Encodings

- (C) is identical to (B)
- (A) is identical to (C) and (B), if we set $\mathbf{V} = \tilde{\mathbf{V}}\mathbf{A}$



(C): Embedding encoding in combination with a one-hot encoding

(linear a, (layer) \mathbf{A} Xt

IIa. Representation Models and Language Models

Language Model

- The ideas is to predict the next word (out off a vocabulary of N^{words} words) in a text, based on the last T words
- Consider we want to predict y_t : y_t has as N^{words} components, one for each word (one-hot encoding)
- The inputs to the models are past words; the model assumption is that a word i is associated with an embedding vector \mathbf{a}_i of dimension r (embedding representation)
- Thus in a first step, a one-hot encoding word i is mapped to the embedding vector of word \mathbf{a}_i which is then the input to a neural network (embedding with one-hot)

Language Model (cont'd)



$$P(y_t = k | y_{t-1}, \dots, y_{t-T}) = \textit{softmax}_k \left(\mathbf{f}_{\mathbf{w}}(\mathbf{a}_{i(t-1)}, \dots, \mathbf{a}_{i(t-T)}) \right)$$

where i(t - m) is the index of the word at position t - m and where $f_w(\cdot)$ is a neural network with one hidden layer and N^{words} output neurons

Embeddings

- Training of the *word embeddings* and *the neural network parameters* can be done self-supervised on a huge corpus (without human labelling)
- After training, one obtains latent word representations (word embeddings) which are published and can be used in other applications
- State of the art are embeddings derived from language models like: ELMo, BERT, Word2vec, and GloVe
- The embedding idea is extremely powerful and one of the corner stones of modern machine learning
- In the next figure, the word embedding matrix \mathbf{A} is denoted as C



IIb. Recurrent Neural Networks

Recurrent Neural Network

- Recurrent neural networks (RNNs) are powerful methods for sequence modelling
- In their simplest form they are used to improve an output prediction by providing a memory for previous inputs
- We do not have to specify a time window T: an RNN can consider the whole history

A Feedforward Neural Network with a Time Index

• We start with a normal feedforward neural network where the pattern is a sequential index t



A Recurrent Neural Network Architecture Unfolded in Time

- The hidden layer now also receives input from the hidden layer of the previous time step
- The hidden layer now has a memory function reflecting hidden inputs
- Thus a Recurrent Neural Network (RNN) is a nonlinear state-space model



Recurrent neural network, unfolded in time

A Recurrent Neural Network Architecture Unfolded in Time (cont'd)

• In a compact notation, we write,

$$\mathbf{z}_t = \operatorname{sig}(\mathbf{B}\mathbf{z}_{t-1} + \mathbf{V}\mathbf{x}_t)$$

 $\mathbf{y}_t = \operatorname{sig}(\mathbf{W}\mathbf{z}_t)$

where we permit several outputs; also, in the last layer we might replace the SiG with the softmax

Temporal Representation

- \mathbf{z}_t is the representation of the hidden state of the system (e.g., patient, plant, ...) at time t
- \mathbf{x}_t can be the embedding of a thing which is present or active at time t (e.g., word, medication, ...)
- Word embedding: $\mathbf{x}_t = \mathbf{a}_{i(t)}$, where i(t) is the word at t and with M = r
- This is a link to representation learning

Recurrent Representation

• The next slide shows an RNN as a recurrent structure


recurrent neural network showing recurrent connections

Backpropagation through time (BPTT)

- Training can be performed using backpropagation through time (BPTT), which is an application of backpropagation (SGD) to the unfolded network structure
- As an additional complexity, the error which occurs to the outputs at time t is not only backpropagated to the previous layers at time t, but also backward in time to all previous neural networks
- In principle, one would propagate back to t = 1; in practice, one typically truncates the gradient calculation





Recurrent neural network, unfolded in time

Echo-State Network

- Recurrent Neural Networks are sometimes difficult to train
- A simple alternative is to initialize ${\bf B}$ and ${\bf V}$ randomly (according to some recipe) and only train ${\bf W}$
- ullet W can be trained with the simple learning rules for linear regression or classification
- This works surprisingly well and is done in the Echo-State Network (ESN) (Herbert Jaeger, 2007)
- ESN (and also liquid-state machines) are examples of so called *reservoir computing*



Issues in Prediction

- An RNN is typically used as predictive model in an iterative setting
- Due to the deterministic nature of the model: if the output y_t is predicted and then becomes available, it will not affect future predictions, since there is no information flowing back from y_t to z_t
- This is in contrast to some probabilistic models such as hidden Markov models (HMMs), Kalman filters, stochastic state space models
- One reason that we can simply apply backpropagation is that RNNs are deterministic! To train HMMs and Kalman filters one can apply a form of EM learning (expectation maximization)

Bidirectional RNNs

- The predictions in bidirectional RNNs depend on past and future inputs
- Useful for sequence labelling problems: handwriting recognition, speech recognition, bioinformatics, ...
- Bidirectional recurrent

$$\mathbf{z}_{t} = [\mathbf{z}_{t}^{f}; \mathbf{z}_{t}^{b}] = \left[\operatorname{sig} \left(\mathbf{V}^{f} \mathbf{x}_{t} + \mathbf{B}^{f} \mathbf{z}_{t-1}^{f} \right); \operatorname{sig} \left(\mathbf{V}^{b} \mathbf{x}_{t} + \mathbf{B}^{b} \mathbf{z}_{t+1}^{b} \right) \right]$$



IIc. LSTMS

Issues in Prediction

- Although the RNN has a memory, it has difficulties remembering important information far in the past
- This can be attributed to the vanishing gradient problem
- Solutions are the long short-term memory (LSTM), and the gated recurrent units (GRUs)
- We now discuss the LSTM

We Start with a Feedforward Neural Network

• Consider a feedforward neural network

$$\mathbf{s}_t = \mathsf{sig}(\mathbf{V}\mathbf{x}_t) \quad \mathbf{z}_t = \mathsf{tanh}(\mathbf{s}_t)$$

 $\hat{\mathbf{y}}_t = \operatorname{sig}(\mathbf{W}\mathbf{z}_t)$

- The transfer function of the hidden neuron is a bit strange, $tanh(sig(Vx_t))$
- s_t is called the cell state vector, z_t is the output vector (of the units, not the neural network)
- In the following steps, each latent unit will become an LSTM unit; thus we will have *H* LSTM units in the network

We Enter Input and Output Gates

- We now use input and output gates which can turn on and off individual LSTM units
- With input gate vector \mathbf{g}_t and output gate vector \mathbf{q}_t

$$\mathbf{s}_t = \mathbf{g}_t \circ \mathsf{sig}(\mathbf{V}\mathbf{x}_t) \quad \mathbf{z}_t = \mathbf{q}_t \circ \mathsf{tanh}(\mathbf{s}_t)$$

Here, \circ is the elementwise (Hadamard) product. As before,

 $\hat{\mathbf{y}}_t = \operatorname{sig}(\mathbf{W}\mathbf{z}_t)$

• Input gates and output gates are also functions of the inputs

$$\mathbf{g}_t = \operatorname{sig}(\mathbf{V}^g \mathbf{x}_t) \quad \mathbf{q}_t = \operatorname{sig}(\mathbf{V}^q \mathbf{x}_t)$$

 Gates are commonly used in mixture of expert neural networks, if the function switches between modes of operations

With Feedback

• We add recurrent connections to the cell state vector and the gates

$$\mathbf{s}_t = \mathbf{g}_t \circ \operatorname{sig}(\mathbf{V}\mathbf{x}_t + \mathbf{B}\mathbf{z}_{t-1}) \qquad \mathbf{z}_t = \mathbf{q}_t \circ \operatorname{tanh}(\mathbf{s}_t)$$

• Input Gate

$$\mathbf{g}_t = \operatorname{sig}(\mathbf{V}^g \mathbf{x}_t + \mathbf{B}^g \mathbf{z}_{t-1})$$

• Output Gate

$$\mathbf{q}_t = \operatorname{sig}(\mathbf{V}^q \mathbf{x}_t + \mathbf{B}^q \mathbf{z}_{t-1})$$

Cell State Vector with Self-recurrency and Forget Gate

• We add self-recurrency to the cell state vector, including a forget gate

$$\mathbf{s}_t = \mathbf{f}_t \circ \mathbf{s}_{t-1} + \mathbf{g}_t \circ \operatorname{sig}(\mathbf{V}\mathbf{x}_t + \mathbf{B}\mathbf{z}_{t-1})$$

• Forget gate

$$\mathbf{f}_t = \operatorname{sig}(\mathbf{V}^f \mathbf{x}_t + \mathbf{B}^f \mathbf{z}_{t-1})$$

Long Short Term Memory (LSTM)

- As a recurrent structure the Long Short Term Memory (LSTM) approach has been very successful
- Basic idea: at time t a newspaper announces that the Siemens stock is labelled as "buy". This information will influence the development of the stock in the next days. A standard RNN will not remember this information for very long. One solution is to define an extra input to represent that fact and that is on as along as "buy" is valid. But this is handcrafted and does not exploit the flexibility of the RNN. A flexible construct which can hold the information is a long short term memory (LSTM) block.
- The LSTM was used very successful for reading handwritten text and is the basis for many applications involving sequential data (NLP, machine translation, ...)
- For the rest of the network, an LSTM node looks like a regular hidden node







Recurrent neural network, unfolded in time

LSTM Applications

- Wiki: LSTM achieved the best known results in unsegmented connected handwriting recognition, and in 2009 won the ICDAR handwriting competition. LSTM networks have also been used for automatic speech recognition, and were a major component of a network that in 2013 achieved a record 17.7% phoneme error rate on the classic TIMIT natural speech dataset
- Applications: Robot control, Time series prediction, Speech recognition, Rhythm learning, Music composition, Grammar learning, Handwriting recognition, Human action recognition, Protein Homology Detection

Comments on LSTM

• You cannot do transfer learning with LSTMs (does not work): thus you need a large data set for any new problem

IId. Encoder-Decoder Networks for Machine Translation

Encoder Decoder Architecture

- Most machine translation systems rely on the encoder-decoder approach
- Neural Machine Translation (NMT)
- Typical numbers: embedding rank: r = 1000, and 1000 hidden units per layer



Encoder

- An **encoder** is an RNN (often an LSTM) with no output layer (no y_t), but maybe several layers of recurrent units; as in the language model, *the inputs are latent embeddings of the words*
- The **encoder vectors** are the (two) embedding vectors (hidden states) of (-), i.e., the end-of-sentence symbol

Decoder

- The initial latent states of the decoder are the encoder vectors (the first two red rectangles in the figure)
- In its simplest form, the latent state of the decoder evolves as

$$\mathbf{z}_t = sig(Bz_{t-1} + Va_{y_{t-1}})$$

 $\mathbf{y}_t = sig(Wz_t)$

- In training, the input to the decoder is the *embedding of the previous word*; the output is the one-hot encoding of the current word
- Training is based on bilingual, parallel corpora; each hidden layer might consist of 1000 hidden units
- In testing, one finds the most likely decoded sequence of words (e.g., using beam search); teacher forcing: the detected word appears at the input of the next instance
- Often one uses two or more hidden layers of LSTM units

Encoder-Decoder Approach in NMT

- Neural Machine Translation (NMT) achieved state-of-the-art performances in largescale translation tasks such as from English to French
- NMT has the ability to generalize well to very long word sequences.
- The model does not have to explicitly store gigantic phrase tables and language models as in the case of standard MT; hence, NMT has a small memory footprint
- Implementing NMT decoders is easy unlike the highly intricate decoders in standard MT

Ile. Attention

Introduction

- The concept of "attention" has gained popularity recently in training neural networks, allowing models to learn alignments between different modalities, e.g., between image objects and agent actions in the dynamic control problem, between speech frames and text in the speech recognition task, or between visual features of a picture and its text description in the image caption generation task
- Attention has successfully been applied to jointly translate and align words
- Attention-based NMT models are superior to non attentional ones in many cases, for example in translating names and handling long sentences
- We follow: Minh-Thang Luong, Hieu Pham, and Christopher D. Manning. 2016. "Effective Approaches to Attention-based Neural Machine Translation"
- First work: D. Bahdanau, K. Cho, and Y. Bengio. 2015. "Neural machine translation by jointly learning to align and translate." In ICLR

Bottleneck in the Encoder-Decoder Architecture

- In the encoder-decoder architecture, all information about the input sequence needs to be transported through the two encoding embedding vectors
- Information earlier in the sequence tends to get forgotten
- One needs short cuts: maybe earlier embedding vectors are important as well!
- In attention one provides information about the top embeddings (upper layers) to the decoder; attention does it in a way that avoids overfitting

Overall Architecture

- The next figure shows the overall architecture
- The attention layer sits on top of the normal encoder-decoder network
- Based on the neural activations in the encoder-decoder, it calculates new activations (grey boxes) in the fourth layer

Attention

- Let z_t (red) be a hidden state vector of interest in the decoder (so called target at t; also called the query)
- Let \mathbf{c}_t be the source-side **context vector** (derived further down)
- The attentional hidden state (grey) is

$$\tilde{\mathbf{z}}_t = \operatorname{sig}\left(\mathbf{V}\mathbf{z}_t + \mathbf{D}\mathbf{c}_t\right)$$

- The Sig is typically the tanh; note that this is a normal layer in a neural network where the layer z_t is the lower layer and \tilde{z}_t is the upper layer and where the lower layer is appended with c_t
- $\tilde{\mathbf{z}}_t$ is then the top hidden layer: the decoded word probability at the target is calculated as $softmax(\mathbf{W}_s \tilde{\mathbf{z}}_t)$

Global Attention: What is the Context c_t ?

- Let z
 _s be any activation vector in the encoder (source hidden state) (often restricted to the top layer) (the key)
- The **alignment** of s for t is a scalar,

$$align(\mathbf{z}_t, \overline{\mathbf{z}}_s) = \frac{\exp(score(\mathbf{z}_t, \overline{\mathbf{z}}_s))}{\sum_{s'} \exp(score(\mathbf{z}_t, \overline{\mathbf{z}}_{s'}))}$$

- The alignment score function calculates a similarity measure: A typical **score** is the dot product, $score(\mathbf{z}_t, \overline{\mathbf{z}}_s) = \mathbf{z}_t^T \overline{\mathbf{z}}_s$; here, \mathbf{z}_t is the query, and $\overline{\mathbf{z}}_s$ is the key
- The already introduced **context vector** is then calculated as (here, \bar{z}_s on the right assumes the role of the **value**)

$$\mathbf{c}_t = \sum_s align(\mathbf{z}_t, \mathbf{\bar{z}}_s) \ \mathbf{\bar{z}}_s$$

Interpretation and Reformulation

- We can think of the context vector \mathbf{c}_t as an approximation to the query vector \mathbf{z}_t , based on a combination of the value vectors $\{\overline{\mathbf{z}}_s\}_s$
- We can write compactly

$$\tilde{\mathbf{z}}_t = \operatorname{sig}\left(\mathbf{V}\mathbf{z}_t + \mathbf{D}\bar{\mathbf{Z}}\operatorname{softmax}\left(\bar{\mathbf{Z}}^T\mathbf{z}_t\right)\right)$$

- The columns in $\overline{\mathbf{Z}}$ are the source embedding vectors
- \mathbf{z}_t might align (pay attention to) certain components of $\overline{\mathbf{Z}}$
- Except for the entries in matrix **D**, no new adaptable parameters are introduced; if attention is useless, all entries in **D** can converge to 0.

Notation used in Attention

• In the literature on the transformer, one writes

Attention(Q, K, V) = Vsoftmax(
$$\frac{1}{\sqrt{r}}$$
K^TQ)

• Here, ${f Q}\equiv {f Z}$ (query matrix), ${f K}\equiv ar {f Z}$ (key matrix), and ${f V}\equiv ar {f Z}$ (value matrix)

• Thus (with the same scaling factor
$$1/\sqrt{r}$$
),

$$(\mathbf{c}_t)_l = (\operatorname{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}))_{t,l}$$

• Note: In the transformer literature, an embedding vector is a row vector and, in the matrices, rows and columns are interchanged

Figure 2: Global attentional model – at each time step t, the model infers a variable-length alignment weight vector a_t based on the current target state \mathbf{Z}_t and all source states $\overline{\mathbf{Z}}_s$ A global context vector c_t is then computed as the weighted average, according to a_t , over all the source states.
Local Attention Model (Position Encoding)

- The global attention has a drawback that it has to attend to all words on the source side for each target word, which is expensive and can potentially render it impractical to translate longer sequences, e.g., paragraphs or documents
- To address this deficiency, a local attentional mechanism has been proposed that chooses to focus only on a small subset of the source positions per target word
- The new alignment becomes

$$align(\mathbf{z}_t, \overline{\mathbf{z}}_s) \exp\left(rac{(s-p_t)^2}{2\sigma^2}
ight)$$

• p_t is the expected position in the input sequence predicted from \mathbf{z}_t using a neural network

$$p_t = S \operatorname{sig}(\mathbf{v}_p^T \operatorname{tanh}(\mathbf{W}_p \mathbf{z}_t))$$

S is the source sentence length (a scalar), and $\mathbf{v}_p, \mathbf{W}_p$ are trained



Figure 3: Local attention model – the model first predicts a single aligned position p_t for the current target word. A window centered around the source position p_t is then used to compute a context vector c_t , a weighted average of the source hidden states in the window. The weights a_t are inferred from the current target state \mathbf{Z}_t and those source states $\mathbf{\overline{Z}}_s$ in the window.

From Attention to Self-Attention

- So far we calculated the attention of an element in the output sequence w.r.t all elements in the input sequence
- Let's consider another task, e.g., entity labeling
- We have a sequence of words/entities (t = 1, 2, ...) as inputs; the goal is to provide a label for each word/entity, or to provide a label for the whole sequence
- We now change notation: $\mathbf{z}_{t,l}$ is the activation vector at layer l
- Self-attention can be applied to any deep neural network
- Self-attention can replace convolutional and recurrent approaches ("attention is all you need")
- Whereas RNNs work left to right, self-attention (as convolutional NNs) work bottom up, in parallel

Self-Attention (cont'd)

• In self-attention, the activation of a hidden layer $z_{t,l}$ is calculated based on other layer's $z_{t,l-1}$ of all entities/data points as

$$\mathbf{z}_{t,l} = \operatorname{sig}\left(\mathbf{V}_{l}\mathbf{z}_{t,l-1} + \mathbf{D}_{l}\mathbf{c}_{t,l-1}\right)$$

• Here, the context vector is

$$\mathbf{c}_{t,l-1} = \sum_{t'} {\sf align} ig(\mathbf{z}_{t,l-1}, \mathbf{z}_{t',l-1} ig) \ \mathbf{z}_{t',l-1}$$

The sum is over all elements in the sequence

- (Often the tanh is used instead of the sig)
- Self-attention can be applied to any layer (not just the top layer)

Comparison

• Feed forward neural network

$$\mathbf{z}_{t,l} = \operatorname{sig}\left(\mathbf{V}_{l}\mathbf{z}_{t,l-1}\right)$$

so here each word label at position t is predicted separately; embeddings are all independent; this is the i.i.d situation

• Fully connected (not used in practice)

$$\mathbf{z}_{t,l} = \operatorname{sig}\left(\mathbf{V}_{l}\mathbf{z}_{t,l-1} + \sum_{t'} \mathbf{C}_{t,t',l}\mathbf{z}_{t',l-1}\right)$$

The embeddings of all words are considered; here one would need to use a standard length sentence (short sentences are dealt with by zero-passing); a problem with this approach is the huge number of parameters in the neural network

Comparison (cont'd)

• Convolutional layer

$$\mathbf{z}_{t,l} = \operatorname{sig}\left(\mathbf{V}_{l}\mathbf{z}_{t,l-1} + \sum_{k}\sum_{t'}\mathbf{C}_{t-t',l}^{k}\mathbf{z}_{t',l-1}\right)$$

Very powerful approach and very successful in NLP; needs zero padding at sentence boundaries; k is the index over different filter kernels

 In some approaches (e.g., graph convolution) simply the averages of the neighbor embeddings are calculated

Comparison (cont'd)

• Recurrent neural networks

$$\mathbf{z}_{t,l} = \operatorname{sig} \left(\mathbf{V}_l \mathbf{z}_{t,l-1} + \mathbf{B}_l \mathbf{z}_{t-1,l} \right)$$

Very powerful approach and very successful in NLP; often LSTM units are used

• Bidirectional recurrent neural networks

$$\mathbf{z}_{t,l} = [\mathbf{z}_{t,l}^f; \mathbf{z}_{t,l}^b]$$

$$= \left[\operatorname{sig} \left(\mathbf{V}_{l}^{f} \mathbf{z}_{t,l-1} + \mathbf{B}_{l}^{f} \mathbf{z}_{t-1,l}^{f} \right); \operatorname{sig} \left(\mathbf{V}_{l}^{b} \mathbf{z}_{t,l-1} + \mathbf{B}_{l}^{b} \mathbf{z}_{t+1,l}^{b} \right) \right]$$

Comparison (cont'd)

• Self-Attention

$$\mathbf{z}_{t,l} = \operatorname{sig} \left(\mathbf{V}_l \mathbf{z}_{t,l-1} + \mathbf{D}_l \mathbf{c}_{t,l-1} \right)$$
$$\mathbf{c}_{t,l-1} = \sum_{t'} \operatorname{align} \left(\mathbf{z}_{t,l-1}, \mathbf{z}_{t',l-1} \right) \mathbf{z}_{t',l-1}$$

(the Sig is often the tanh) self-attention can replace convolutional or recurrent layers











Conclusions

- Sequential models find many applications in natural language processing (NLP) applications, including machine translation
- Attention mechanisms are the basis for state of the art machine translation (Transformer) and context sensitive embedding models; BERT (Bidirectional Encoder Representations from Transformers)

Transformer, BERT, GPT

Transformer

Transformer

- Bottleneck of previous approaches in NMT: sequential processing at the encoding step
- The Transformer **dispensed the recurrence and convolutions** involved in the encoding step entirely and based models only on attention mechanisms to capture the global relations between input and output
- Each layer has two sub-layers comprising multi-head attention layer followed by a position-wise feed forward network

Context with Learned Projection Matrices

• Consider self-attention with (we do not explicitly indicate the layer to simplify notation)

$$\mathbf{c}_{t} = \mathbf{W}^{O} \sum_{t'} align\left(\mathbf{W}^{Q} \mathbf{z}_{t}, \mathbf{W}^{K} \mathbf{z}_{t'}\right) \mathbf{W}^{V} \mathbf{z}_{t'}$$

- The $\mathbf{W}^Q \in \mathbb{R}^{r_{head} \times r}$, $\mathbf{W}^K \in \mathbb{R}^{r_{head} \times r}$, $\mathbf{W}^V \in \mathbb{R}^{r_{head} \times r}$, $\mathbf{W}^O \in \mathbb{R}^{r \times r_{head}}$, are projection matrices for the **query** \mathbf{z}_t , the **key** $\mathbf{z}_{t'}$ and the **value** $\mathbf{z}_{t'}$
- Thus the calculation of the context vector (still a vector!) involves tunable matrices
- In the transformer paper: r = 512, $r_{head} = 64$ (here, r is the number of neurons in a hidden layer)

Multi-head Attention

• Now we define L ($k = 1, \ldots, L$) context vectors, also called heads

$$head_{k,t} = \mathbf{c}_{k,t} = \sum_{t'} align\left(\mathbf{W}_{k}^{Q}\mathbf{z}_{t}, \mathbf{W}_{k}^{K}\mathbf{z}_{t'}\right) \mathbf{W}_{k}^{V} \mathbf{z}_{t'}$$

• With L heads (in the original transformer paper: L = 8),

$$\mathbf{c}_t = \mathbf{W}^O[\mathit{head}_{1,t}; \cdots; \mathit{head}_{L,t}]$$

Thus there are L different attention mechanisms which are appended and multiplied by the $r \times (L \times r_{head})$ matrix \mathbf{W}^O , where r is the embedding dimension

• Note that \mathbf{c}_t is again a single vector of dimension r; the softmax function in *align* is again scaled by $1/\sqrt{r_{head}}$

Encoder

- The encoder uses **self-attention**
- We see a Resnet like structure; layer normalization is used (normalizes the activations, like batch normalization)
- In addition, a simple feed forward neural network is used

Decoder: Encoder-decoder Attention Layer

- The decoder uses both **self-attention** and **cross-attention** (encode-decoder attention)
- The decoder also has a multi-head encoder-decoder attention layer
- Encode-decoder attention w.r.t. the embedding vector in the upper layers in the encoder
- Self-attention at t w.r.t. the embedding vectors previous to t

Auto-Regressive: predicted word embeddings (all future inputs are masked)

Masked Self-attention in the Decoder

- Whereas the encoder only uses self-attention, the decoder uses attention w.r.t all tokens in the encoder and all tokens in the decoder, that were already decoded
- The purpose of the masking is to make sure that the states do not attend to tokens that are "in the future" but only to those "in the past"
- Attention is directional: the past is never updated

Notation Closer to Transformer Paper

• The notation is

$$\begin{aligned} &\text{head}_k = \text{Attention}(\mathbf{W}_k^Q \mathbf{Q}, \mathbf{W}_k^K \mathbf{K}, \mathbf{W}_k^V \mathbf{V}) \\ &\text{MultiHead}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \mathbf{W}_k^O \text{Concat}(\text{head}_1, \dots, \text{head}_L) \end{aligned}$$

• The feed forward neural network with ReLU activation is described as

 $FFN(\mathbf{x}) = \mathbf{W}_2 \max(0, \mathbf{W}_1 \mathbf{x} + \mathbf{b}_1) + \mathbf{b}_2$

(note again that rows and columns are interchanged)

• All matrices are layer-specific, but identical for different positions

Decoding

- The decoder is auto-regressive
- It uses a one-step iterative greedy approach: a decoded token is the input for the following time-step
- At a higher "temperature" also less likely tokens are selected

BERT

- **BERT** (Bidirectional Encoder Representations from Transformers) from Google leverages attention mechanism and transformer to learn word contextual relations using a masked language model (MLM)
- It is based on the **encoder of the transformer**

BERT: Masked LM

Masked: in training inactive

towardsdatascience

BERT

- BERT is almost an auto encoder: but some tokens of the input sentence are removed (masked) and the network is trained to predict those tokens at the output layer
- Masked language modelling (MLM)
- The context-sensitive word embeddings can be used for all sorts of tasks, like world labelling, NER, ...

BERT: Next Sentence Prediction

Also trained to predict: IsNextSequence or not

Source: BERT [Devlin et al., 2018], with modifications

BERT

- BERT performs next sentence prediction (NSP) but is not considered generative AI; In the BERT training process, the model receives pairs of sentences as input and learns to predict if the second sentence in the pair is the subsequent sentence in the original document.
- The actual embedding is a sum of the token embedding, the posiditon embedding and an embedding for first sentence or second sentence
- A sentence embedding indicating Sentence A or Sentence B is added to each token. Sentence embeddings are similar in concept to token embeddings with a vocabulary of 2

Local Position Encoding for Transformer, Bert, GPT

- Without positional encoding, the transformer (and BERT) would be a bag-of-words approach and could not distinguish between "Live to Work" and "Work to Live", which an LSTM could!
- To address this, the transformer adds a vector to each input embedding. These vectors follow a specific pattern, which helps to determine the position of each word, or the distance between different words in the sequence
- A position encoding vector is defined which encodes the position of a token
- This vector is added (summed) to the word encoding embedding
- So the input embedding vector does not just encode the token (world) but also the position in the sequence

Position Encoding

- Consider that t is the position in the sequence; PosEnc(t, :) is the position encoding vector, and PosEnc(t, l) with l = 1, ..., r is its l-th component
- We have, for $i = 0, \ldots, r/2$

$$PosEnc(t,2i) = sin((t-1)/n^{2i/r})$$
$$PosEnc(t,2i+1) = cos((t-1)/n^{2i/r})$$

- Increasing t (left to right), we sample a sine/cosine wave pattern
- Increasing i (or i), the frequency decreases, from 1 to 1/n
- n is a user specified parameter (often n = 10000)

GPT

- Generative Pre-trained Transformer (GPT)
- Typically it only uses the Decoder part of the Transformer
- It is initialized with the prompt; a prompt can be a question, a task, ...; it is the initialization (prompt engineering)
- GPT-1 to GPT-3: main difference are: different forms of training schedules!
- GPT-4 can handle sequences of 32k tokens (under development: sequences of 100k or more tokens)
- Attention is calculated w.r.t. all prompt tokens (all input tokens) and all previously generated tokens

GPT (cont'd)

- Tokenization (byte pair encoding): Typically, most words will be encoded as a single token, while rare words will be encoded as a sequence of a few tokens, where these tokens represent meaningful word parts. This translation of text into tokens can be found by variants of byte pair encoding, such as subword units
- Vocabulary size (Llama): 32k, context length: 2k (now: 100k), 65B parameters, 4.4B tokens, 2k GPUs, 3 weeks of training
GPT and **Self-Attention**

- The slide shows the attention mechanism and the feedforward layer
- Note the ResNet structure: Each layer is connected to the last layer via later layers but also directly (no vanishing gradient!)





GPT

ChatGPT Training

- Training phases:
- 1: Pretraining (99% of effort); LLM training; leads to base model
- Data: internet crawls (CommonCrawl, C4), Github, Wikipedia, Books, ArXiv, Stack-Exchange
- 2: Supervised fine-tuning; with ideal prompt-response pairs (the responses can be quite long); LLM model; leads to **SFT model**
- 3: Reward modeling; human annotators who rank different responses (completions); predicts (subjective) quality; leads to **RM model**
- 4: Reinforcement Learning; used RM; leads to RL model (RLHF: Reinforcement Learning with Human Feedback); improves the LLM further but difficult to get to work

GPT Research

- Research issues: fine-tuning, prompt engineering, providing access to external data like databases (vector databases, database queries)
- At this stage: increasing the number of parameters is not the most important factor; training set size/quality is currently more in focus
- Hallucination: An inherent feature of generative models?

GPT

word, but also to continue on writing whole paragraphs in the style of the initial prompt. For example, when I gave GPT-3 the prompt:

In 1958, Frank Rosenblatt introduced the Perceptron, an early version of a neural net that could only detect the location of a square on a printed card. Fifty years later, neural nets are capable of ...

The software dutifully completed the thought, and then continued on, picking up on the historical framing of the initial text:

... detecting and classifying faces, finding patterns in data, and emulating the human brain. While the neural net matured in academia, it also found its way into the tech industry. In the late 1980s and early 90s, neural nets were used in speech recognition and character recognition applications. The growth in computing power and the ability to collect larger data sets allowed other neural net applications to emerge.

New York Times, 2022-04-22

GPT 175B

- 96 layers, r = 12288, $r_{head} = 128$, number of heads L = 96
- \mathbf{W}^O , \mathbf{W}^Q , \mathbf{W}^K , \mathbf{W}^V : Each of them has 12288×128 parameters, and we have 96×96 of them: 58B
- W_1 and W_2 in the FFN: 12288 × 12288 × 4 × 2 (12288 × 4 is the size of the hidden layer) and we have 96 of them: 116B
- Token embedding: $50257 \times 12299 = 0.6B$
- This is together: 173.6B; we need to add the biases, other parameters to get to 175B

Model Name	$n_{\rm params}$	$n_{\rm layers}$	d_{model}	$n_{\rm heads}$	$d_{ m head}$	Batch Size	Learning Rate
GPT-3 Small	125M	12	768	12	64	0.5M	6.0×10^{-4}
GPT-3 Medium	350M	24	1024	16	64	0.5M	3.0×10^{-4}
GPT-3 Large	760M	24	1536	16	96	0.5M	$2.5 imes 10^{-4}$
GPT-3 XL	1.3B	24	2048	24	128	1M	2.0×10^{-4}
GPT-3 2.7B	2.7B	32	2560	32	80	1M	$1.6 imes 10^{-4}$
GPT-3 6.7B	6.7B	32	4096	32	128	2M	1.2×10^{-4}
GPT-3 13B	13.0B	40	5140	40	128	2M	$1.0 imes 10^{-4}$
GPT-3 175B or "GPT-3"	175.0B	96	12288	96	128	3.2M	$0.6 imes 10^{-4}$

Table 2.1: Sizes, architectures, and learning hyper-parameters (batch size in tokens and learning rate) of the models which we trained. All models were trained for a total of 300 billion tokens.

Then they say:

Table 2.1 shows the sizes and architectures of our 8 models. Here nparams is the total number of trainable parameters, nlayers is the total number of layers, dmodel is the number of units in each bottleneck layer (we always have the feedforward layer four times the size of the bottleneck layer, dff = 4 * dmodel), and dhead is the dimension of each attention head. All models use a context window of nctx = 2048 tokens. We partition the model across GPUs along both the depth and width dimension in order to minimize data-transfer between nodes. The precise architectural parameters for each model are chosen based on computational efficiency and load-balancing in the layout of models across GPU's. Previous work [KMH+20] suggests that validation loss is not strongly sensitive to these parameters within a reasonably broad range.

Foundation Models: A New Age?

- Different phases in machine learning:
- Before 2012: Small scale, feature-based learning
- From 2012 on: Deep Learning: large models with large sets of labelled data, supervised
- From 2020/2023 on: Foundation Models: large models pretrained with large sets of unlabelled data, using self-supervised learning