

Transformer

Adaptado a series temporales

Luis Gallego Ledesma
lgallego@us.es

Índice

- Aplicaciones Transformers
- RNN vs Transformers
- Arquitectura del Transformer
- Transformer para series temporales
- Nuestros resultados
- Posibles mejoras

Generating text with Transformers

Human Prompt

In a shocking finding, scientist discovered a herd of unicorns living in a remote, previously unexplored valley, in the Andes Mountains. Even more surprising to the researchers was the fact that the unicorns spoke perfect English.

Generated Text
(10 tries)

The scientist named the population, after their distinctive horn, Ovid's Unicorn. These four-horned, silver-white unicorns were previously unknown to science.

Now, after almost two centuries, the mystery of what sparked this odd phenomenon is finally solved.

Dr. Jorge Pérez, an evolutionary biologist from the University of La Paz, and several companions, were exploring the Andes Mountains when they found a small valley, with no other animals or humans. Pérez noticed that the valley had what appeared to be a natural fountain, surrounded by two peaks of rock and silver snow.

Pérez and the others then ventured further into the valley. "By the time we reached the top of one peak, the water looked blue, with some crystals on top," said Pérez.

Pérez and his friends were astonished to see the unicorn herd. These creatures could be seen from the air without having to move too much to see them – they were so close they could touch their horns.

Language Models are Unsupervised Multi Task Learners, Radford et. al. (2019)

GPT-3

Describe a layout.

2 lightgrey buttons that say "Search Google" and "I'm Feeling Lucky" with padding in between them

Generate

Mega-mode is on

```
// the google logo


// 2 lightgrey buttons that say "Search
Google" and "I'm Feeling Lucky" with padding
in between them
<div style={{padding: 10}}> <button style=
{{color: 'white', backgroundColor:
'lightgrey'}}>Search Google</button> <button
style={{color: 'white', backgroundColor:
'lightgrey'}}>I'm Feeling Lucky</button>
</div>
```




DALL - E

TEXT PROMPT

a clock in the shape of an avocado. a clock imitating an avocado.


AI-GENERATED
IMAGES

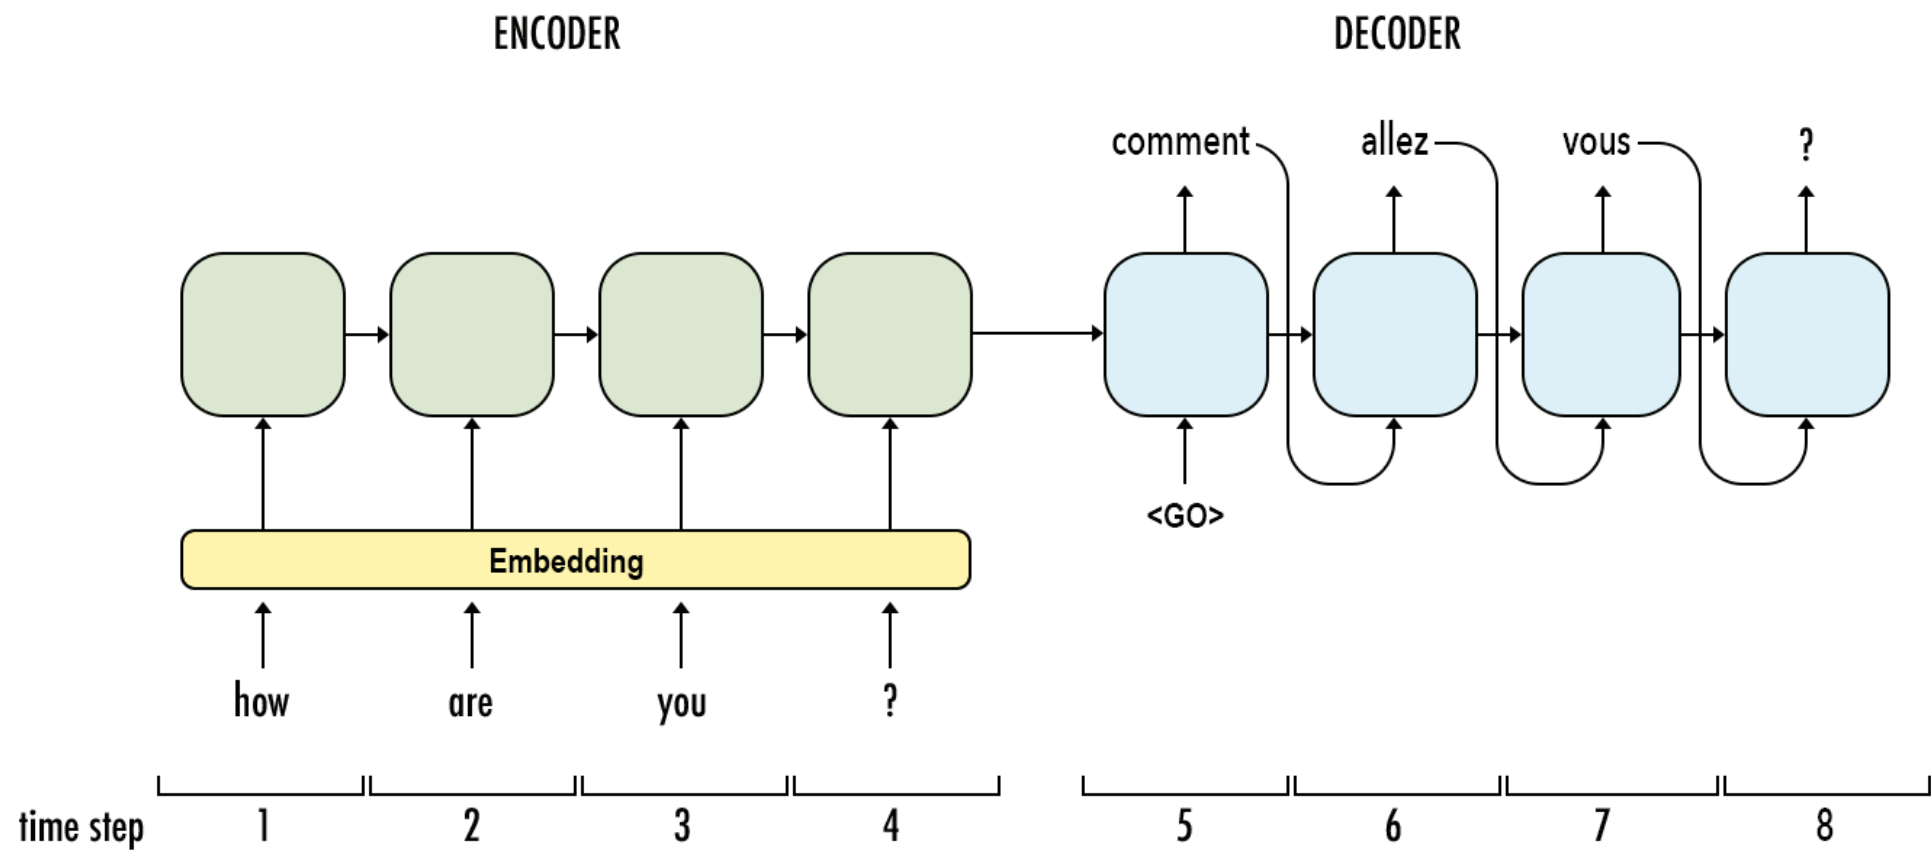


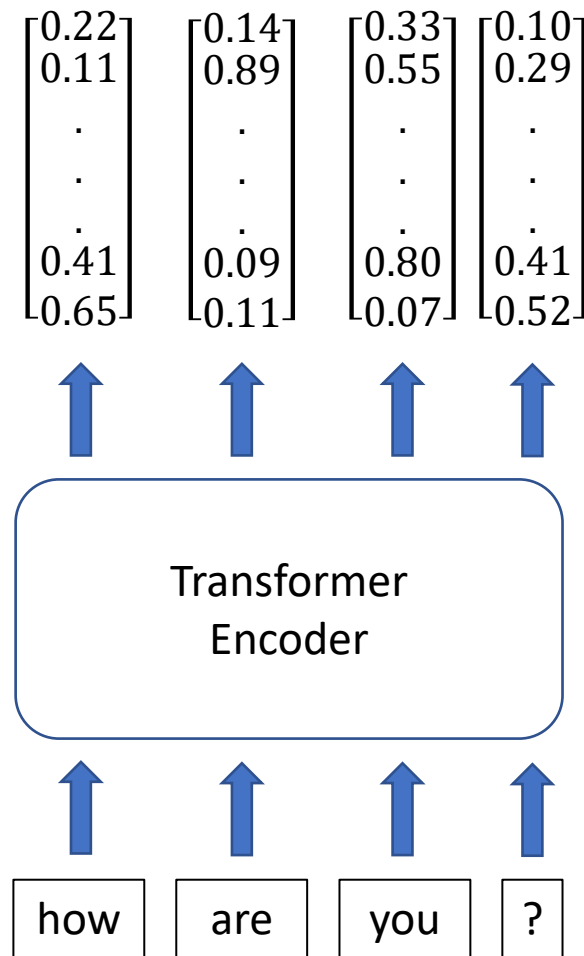
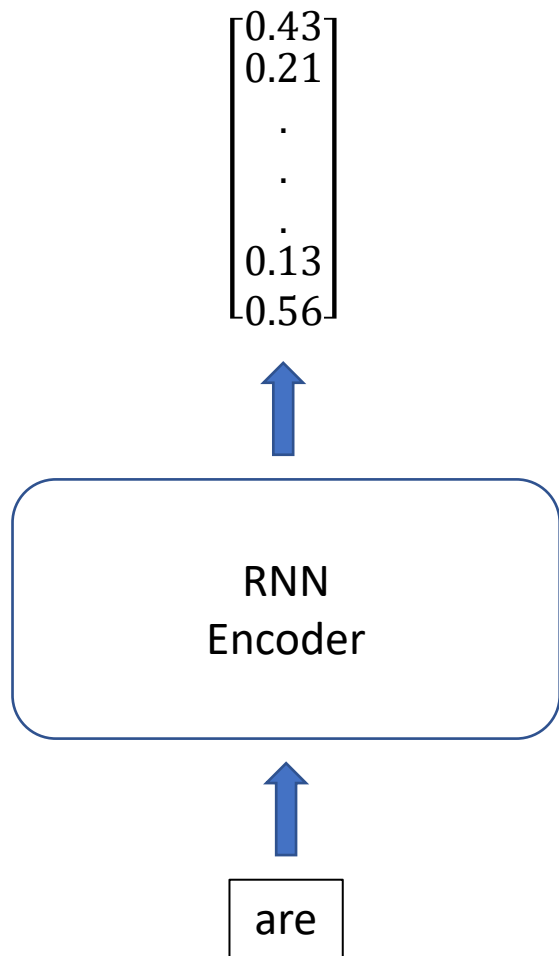
TEXT PROMPT

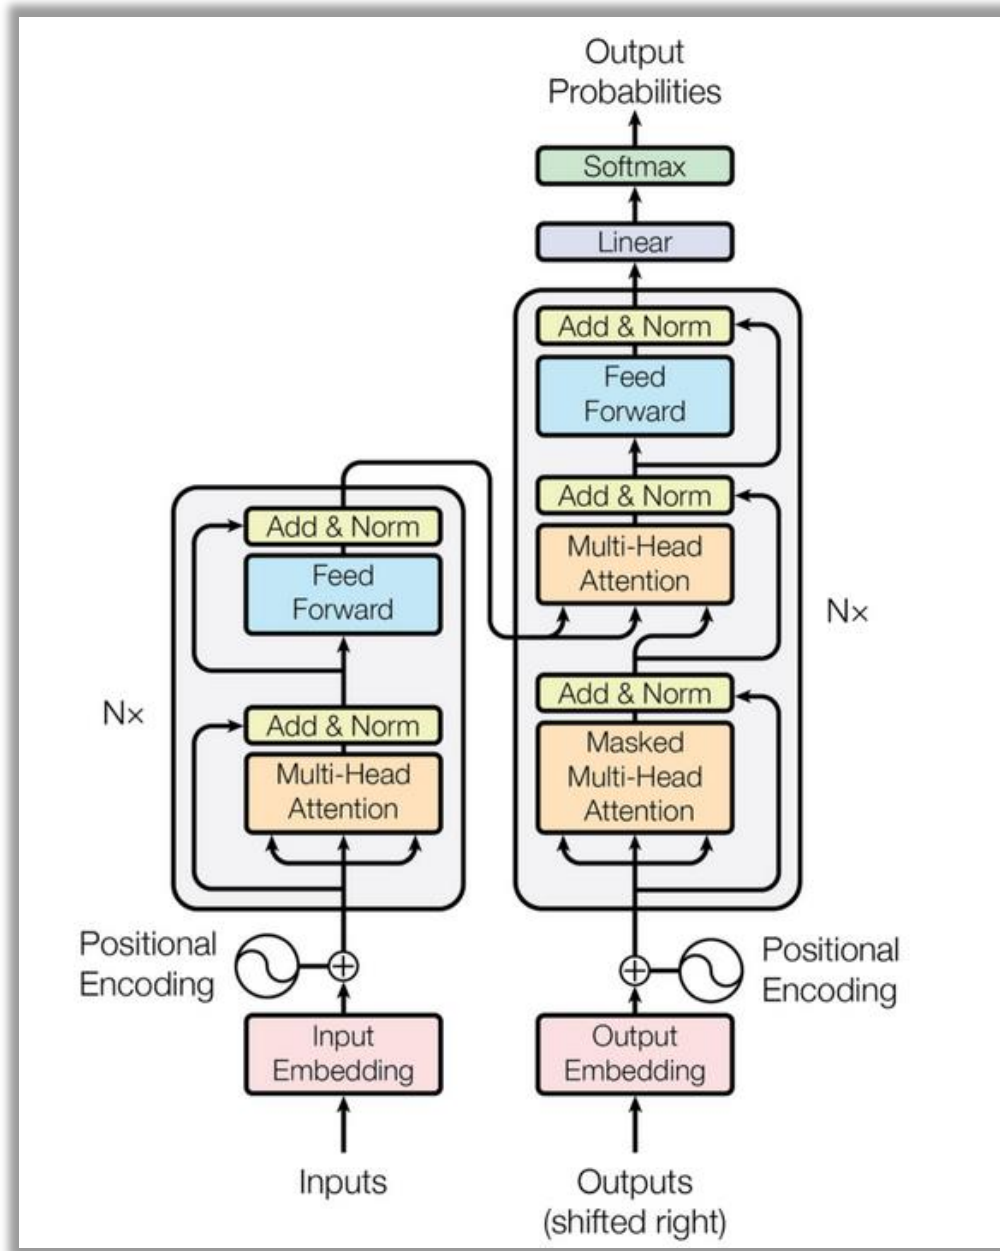
an armchair in the shape of a strawberry. an armchair imitating a strawberry.


AI-GENERATED
IMAGES



















Scaled Dot-Product Attention



$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V$$





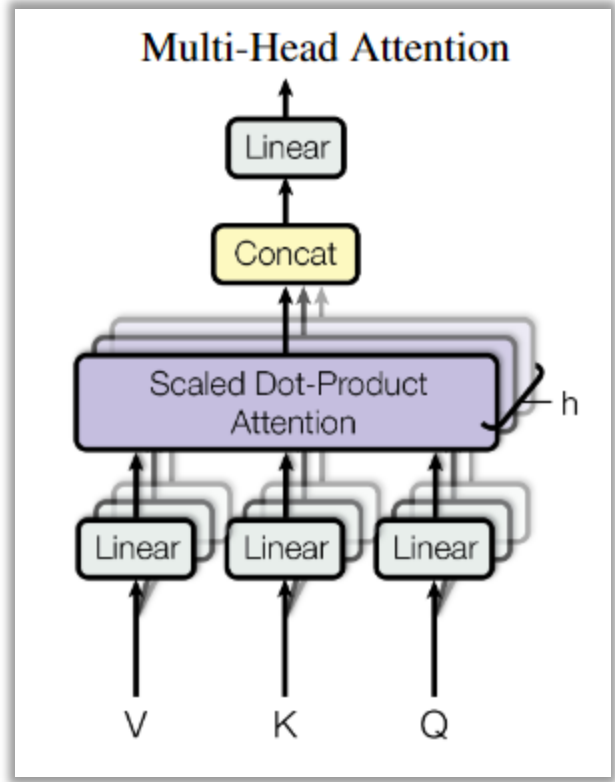
Source: The Illustrated Transformer
<http://jalammr.github.io/illustrated-transformer>



$$\text{softmax} \left(\frac{\begin{matrix} \mathbf{Q} \\ \begin{array}{|c|c|c|} \hline \square & \square & \square \\ \hline \square & \square & \square \\ \hline \end{array} \end{matrix} \times \begin{matrix} \mathbf{K}^T \\ \begin{array}{|c|c|} \hline \square & \square \\ \hline \square & \square \\ \hline \square & \square \\ \hline \end{array} \end{matrix} \right) \begin{matrix} \mathbf{V} \\ \begin{array}{|c|c|c|} \hline \square & \square & \square \\ \hline \square & \square & \square \\ \hline \end{array} \end{matrix}$$

$$= \begin{matrix} \mathbf{Z} \\ \begin{array}{|c|c|c|} \hline \square & \square & \square \\ \hline \square & \square & \square \\ \hline \end{array} \end{matrix}$$

Source: The Illustrated Transformer
<http://jalammr.github.io/illustrated-transformer>

Multi-Head Attention

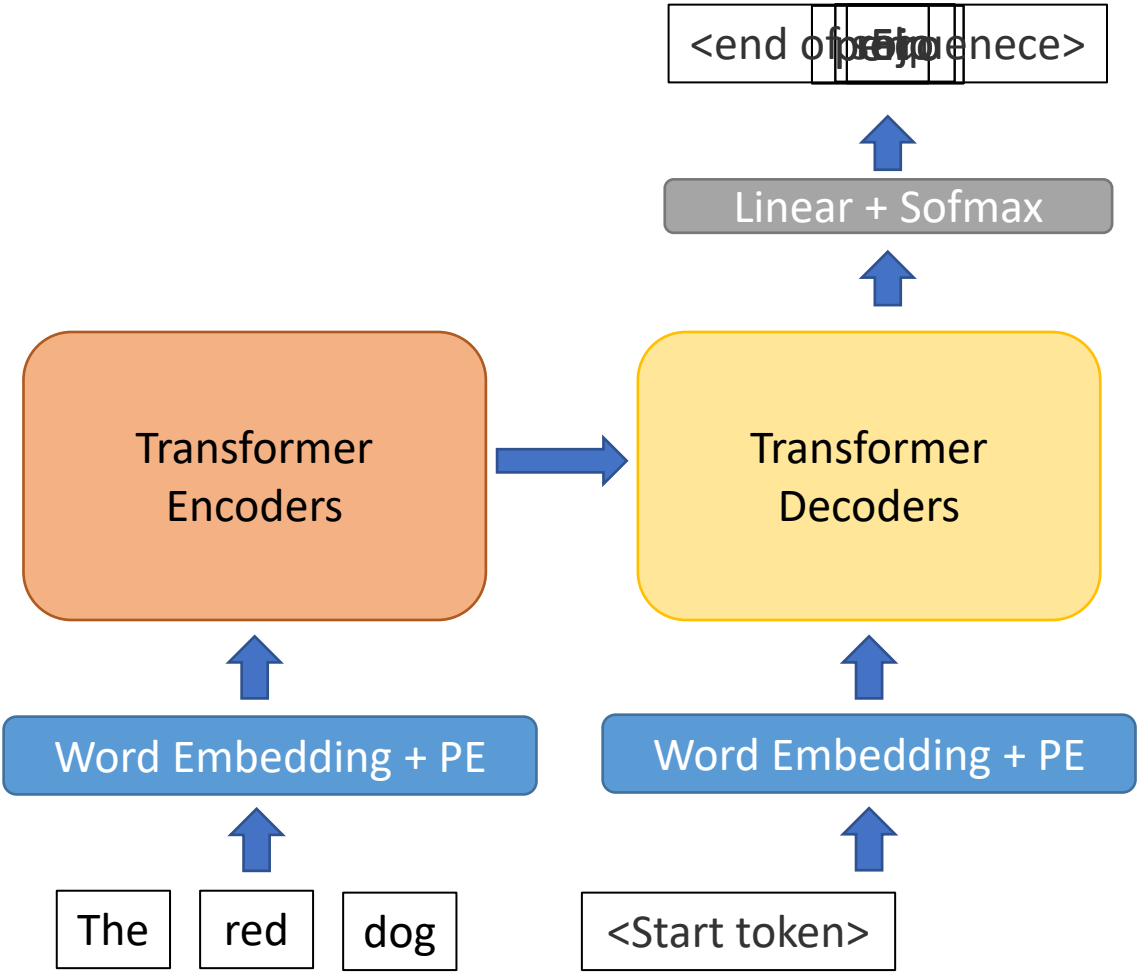




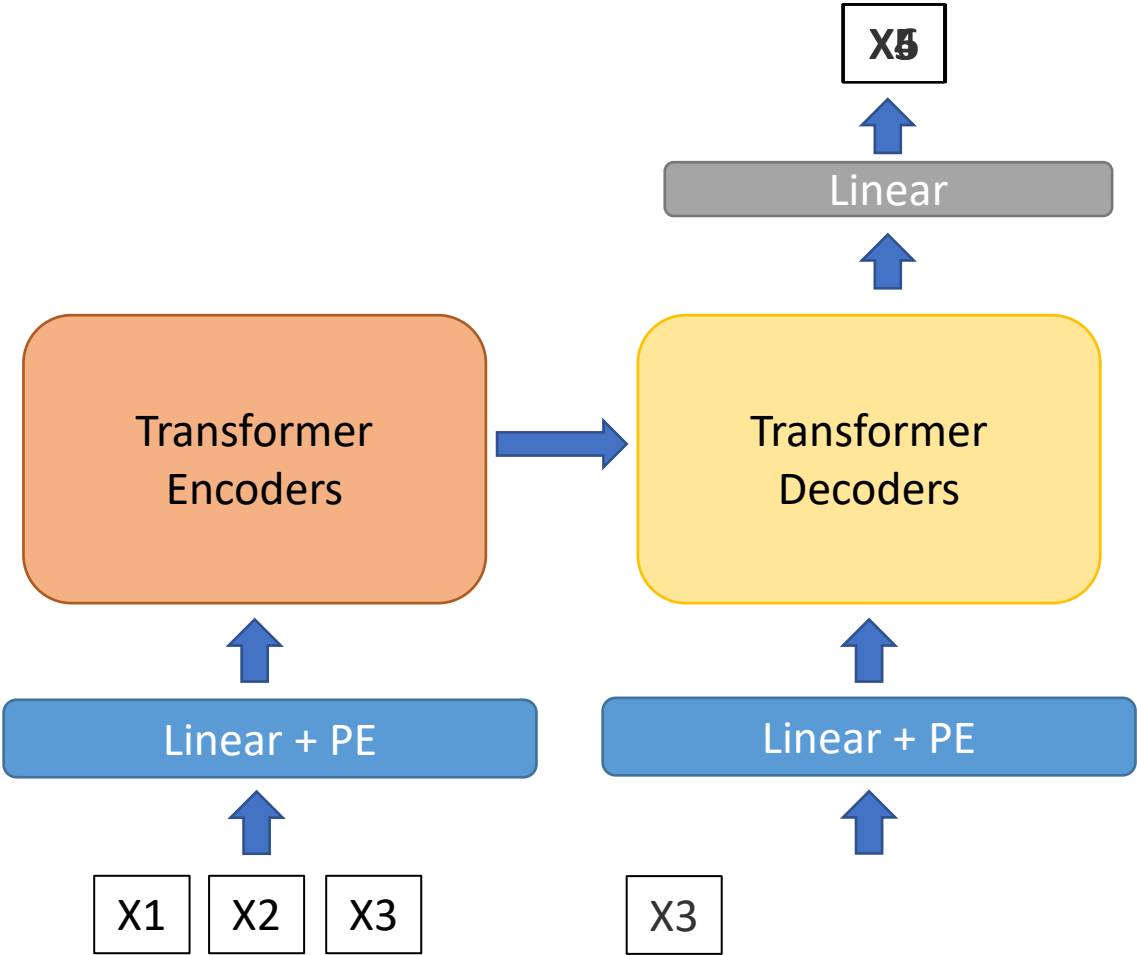
$$p_{i,2j} = \sin\left(\frac{i}{10000^{2j/d}}\right)$$

$$p_{i,2j+1} = \cos\left(\frac{i}{10000^{2j/d}}\right)$$

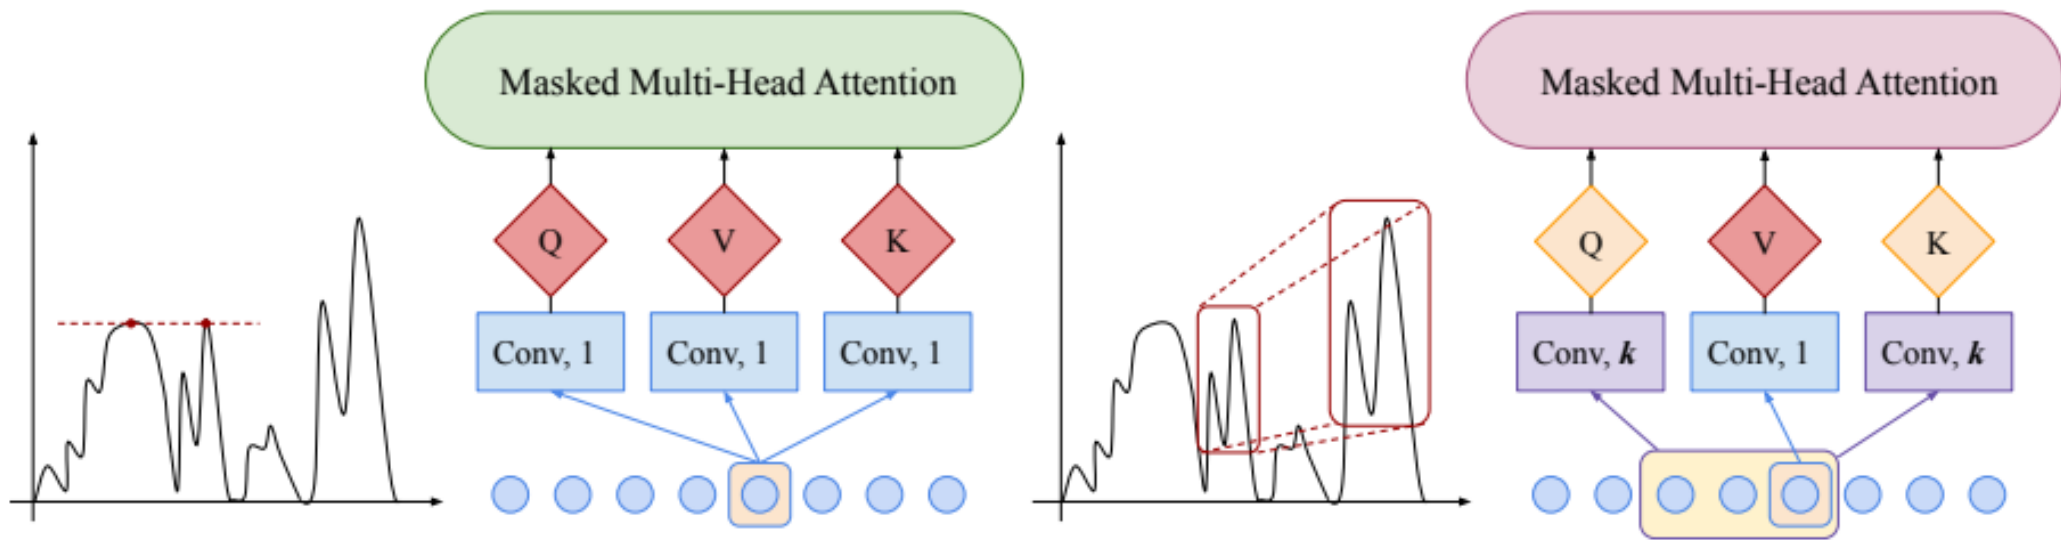
OUTPUT:



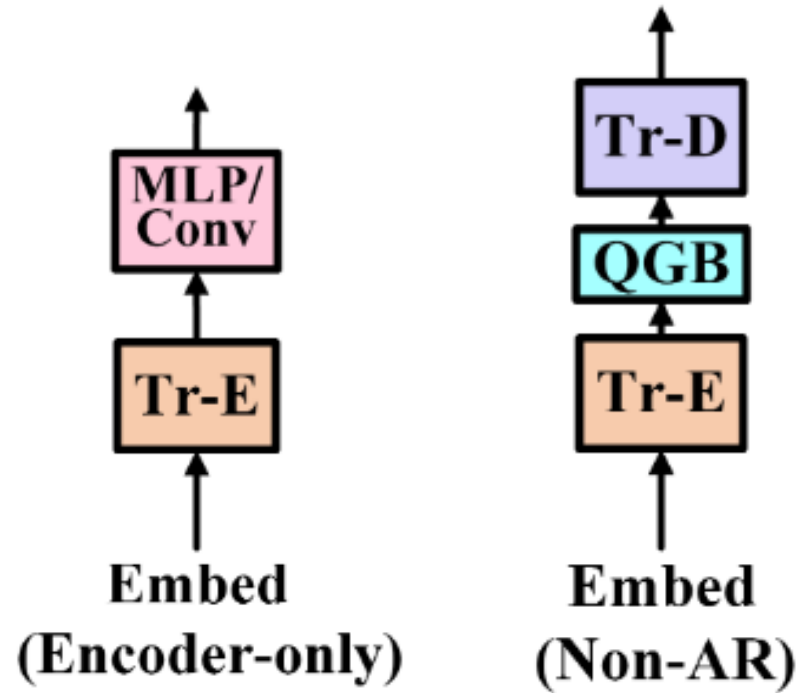
OUTPUT:



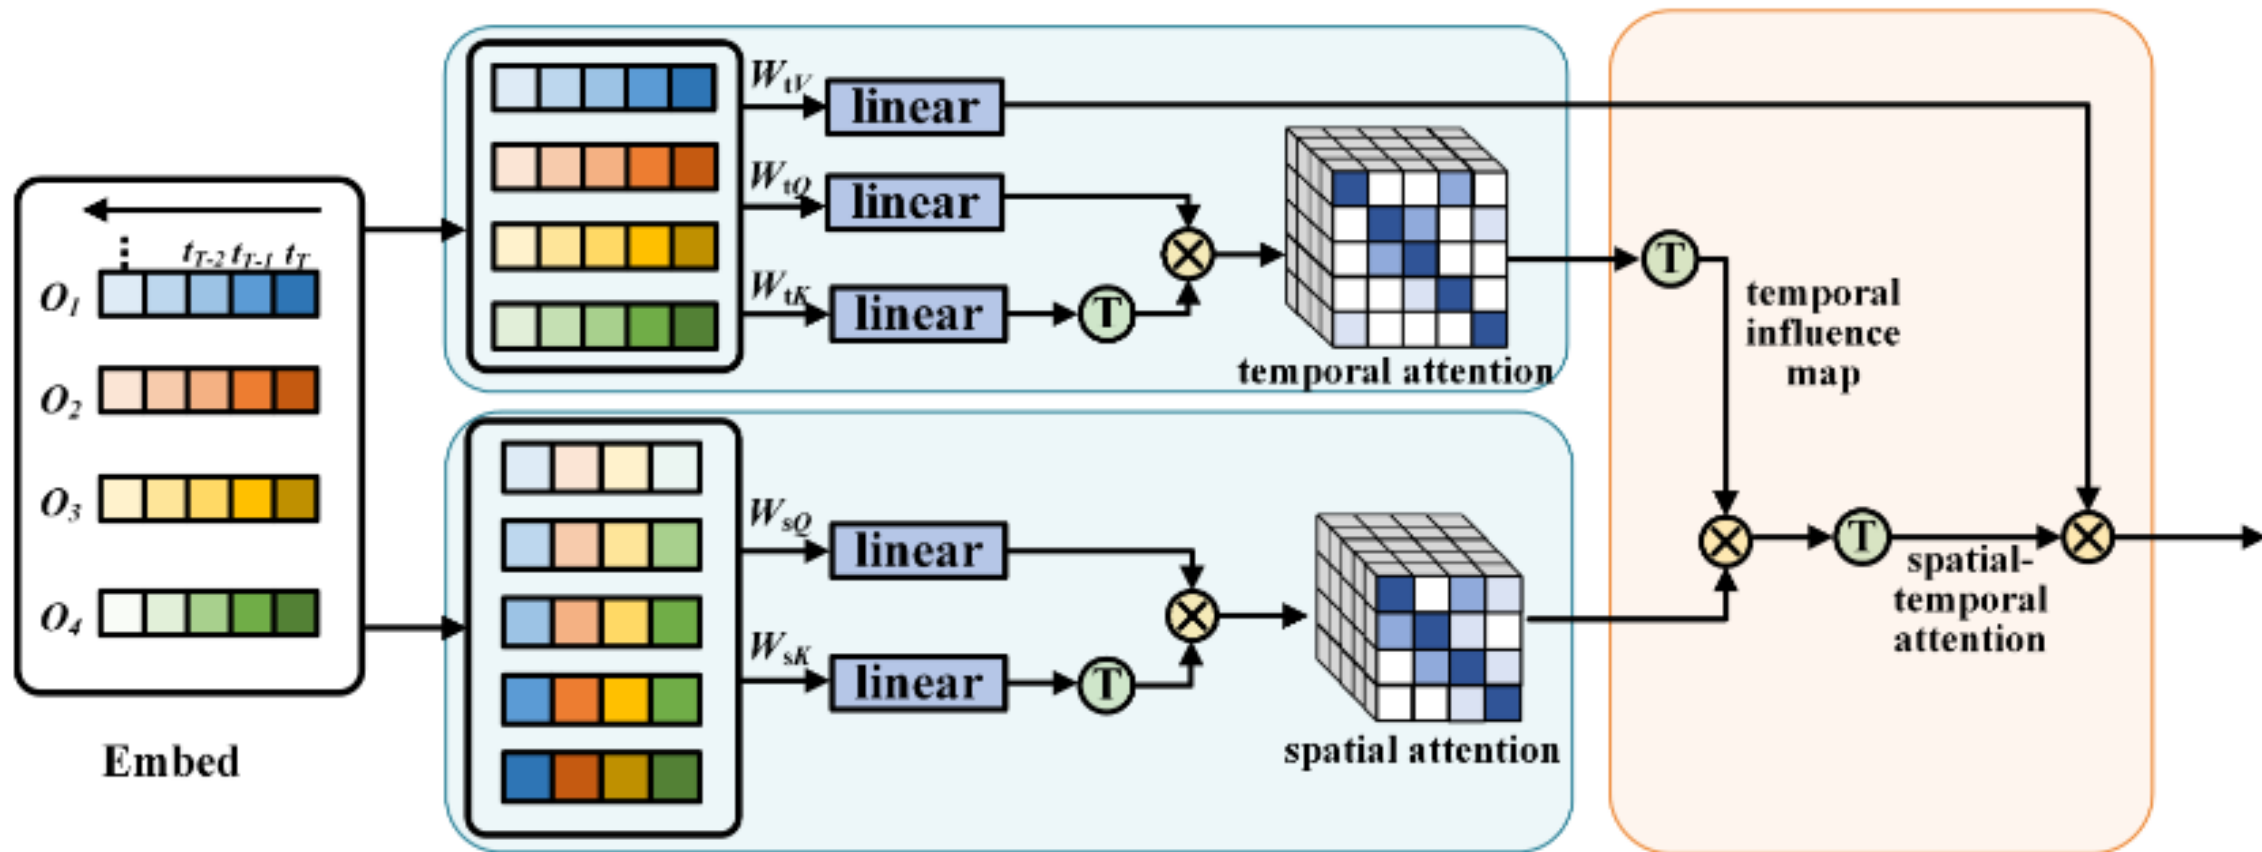
Datasets		WAPE			MAE		
		Transformer	LSTM	CNN	Transformer	LSTM	CNN
1	CIF2016o12	11.207	12.475	12.479	12,564.18	11,732.31	12,762.73
2	CIF2016o6	16.157	15.352	17.143	2,182,435.2	3,636,929.7	2,833,131.5
3	ExchangeRate	0.303	0.300	0.335	0.0021	0.0019	0.0020
4	M3	12.490	15.282	15.612	659.18	700.25	709.44
5	M4	13.587	14.281	14.256	588.38	597.54	593.71
6	NN5	18.637	18.589	18.852	3.570	3.538	3.572
7	SolarEnergy	13.550	12.452	11.717	2.246	2.066	1.977
8	Tourism	18.68	19.081	18.497	2,202.11	2,280.09	1,969.58
9	Traffic	33.541	31.960	34.406	0.0121	0.0114	0.0124
10	Traffic-metr-la	3.418	3.359	3.337	2.029	2.009	1.991
11	Traffic-perms-bay	1.333	1.314	1.433	0.885	0.870	0.946
12	WikiWebTraffic	46.110	46.477	46.914	11.796	12.063	12.106
Mean ranking		1.833	1.750	2.416	1.916	1.833	2.250



Enhancing the Locality and Breaking the Memory Bottleneck of Transformer on Time Series Forecasting, NeurIPS 2019



NAST: Non-Autoregressive Spatial-Temporal Transformer for Time Series Forecasting



NAST: Non-Autoregressive Spatial-Temporal Transformer for Time Series Forecasting

Referencias

- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). [Attention is all you need](#). In *Advances in neural information processing systems* (pp. 5998-6008)
- DeepMind's deep learning videos 2020 with UCL, Lecture: [Deep Learning for Natural Language Processing](#), Felix Hill
- The Illustrated Transformer <http://jalammar.github.io/illustrated-transformer/>
- Yannic Kilcher video explaining the original paper: [Attention Is All You Need](#)
- Li, Shiyang, et al. "Enhancing the locality and breaking the memory bottleneck of transformer on time series forecasting." arXiv preprint arXiv:1907.00235 (2019).
- Chen, Kai, et al. "NAST: Non-Autoregressive Spatial-Temporal Transformer for Time Series Forecasting." arXiv preprint arXiv:2102.05624 (2021).