Natural Language Processing

The Transformer architecture

Marco Kuhlmann
Department of Computer and Information Science

This work is licensed under a
Creative Commons Attribution 4.0 International License.

Attention is all you need

- Recurrent neural networks implement a sequential model of computation that processes sequences element-by-element.
- In contrast, attention facilitates direct access to all elements, independently of sequence length.
- The **Transformer** is an encoder–decoder architecture that drops recurrent neural networks and exclusively uses attention.

  can be parallelised

Vaswani et al. (2017)

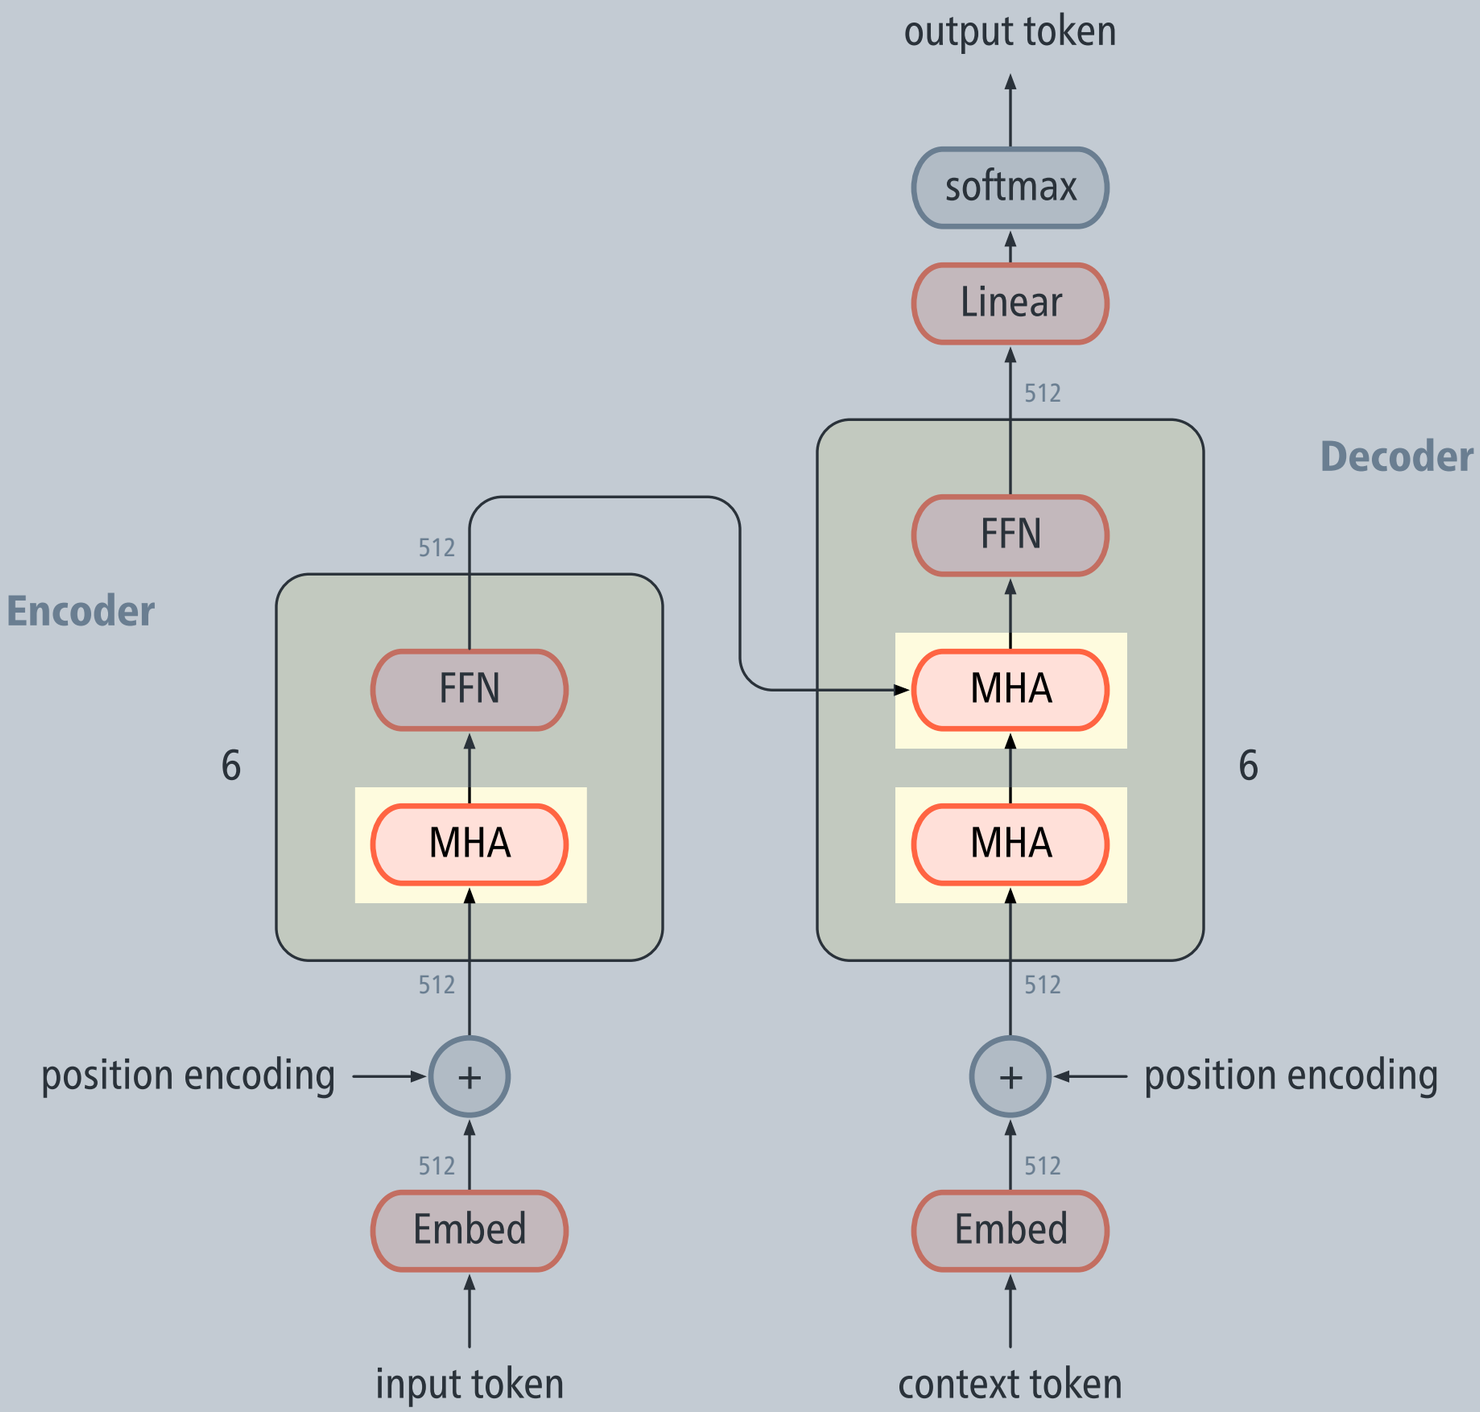
output token
softmax
Linear
512
Decoder
FFN
MHA
6
512
Encoder
FFN
MHA
6
512
position encoding
+
512
Embed
input token
512
+
position encoding
512
Embed
context token

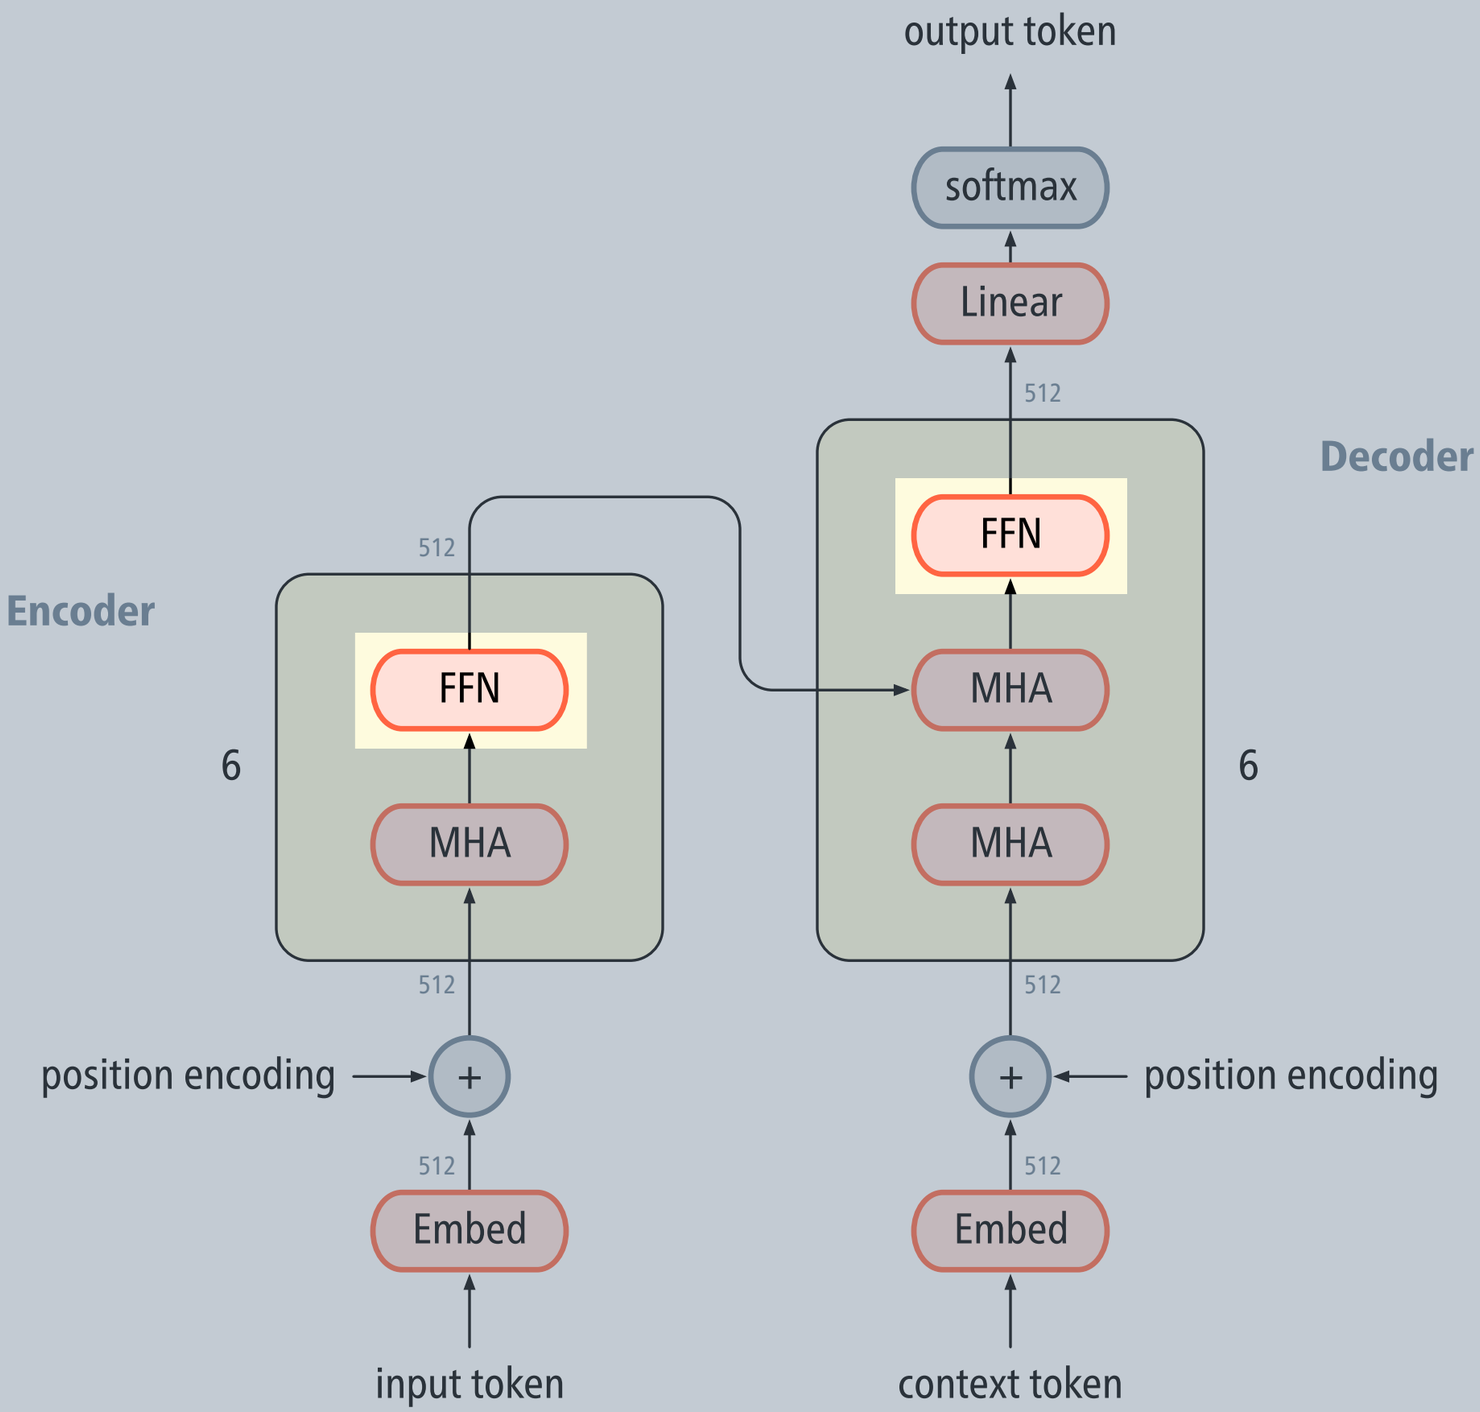
output token
softmax
Linear
512
Decoder
FFN
MHA
6
512
Encoder
FFN
MHA
6
512
position encoding
+
512
Embed
input token
512
+
position encoding
512
Embed
context token

Example translation

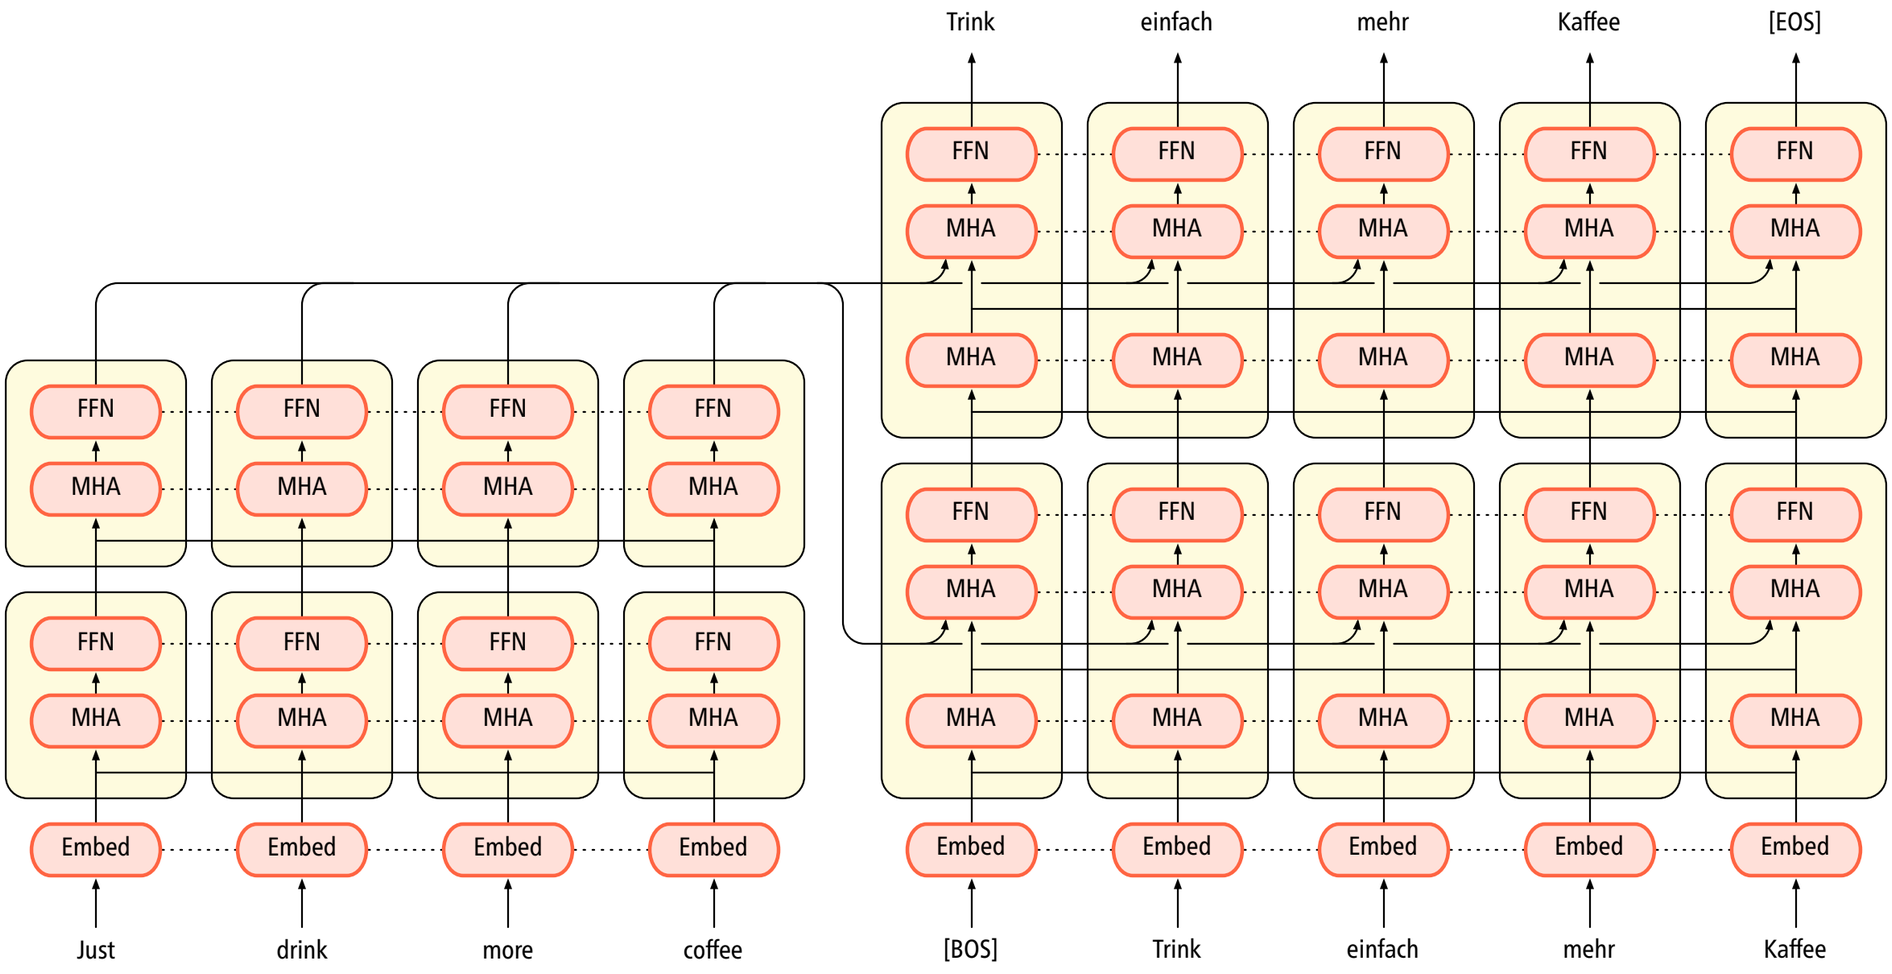

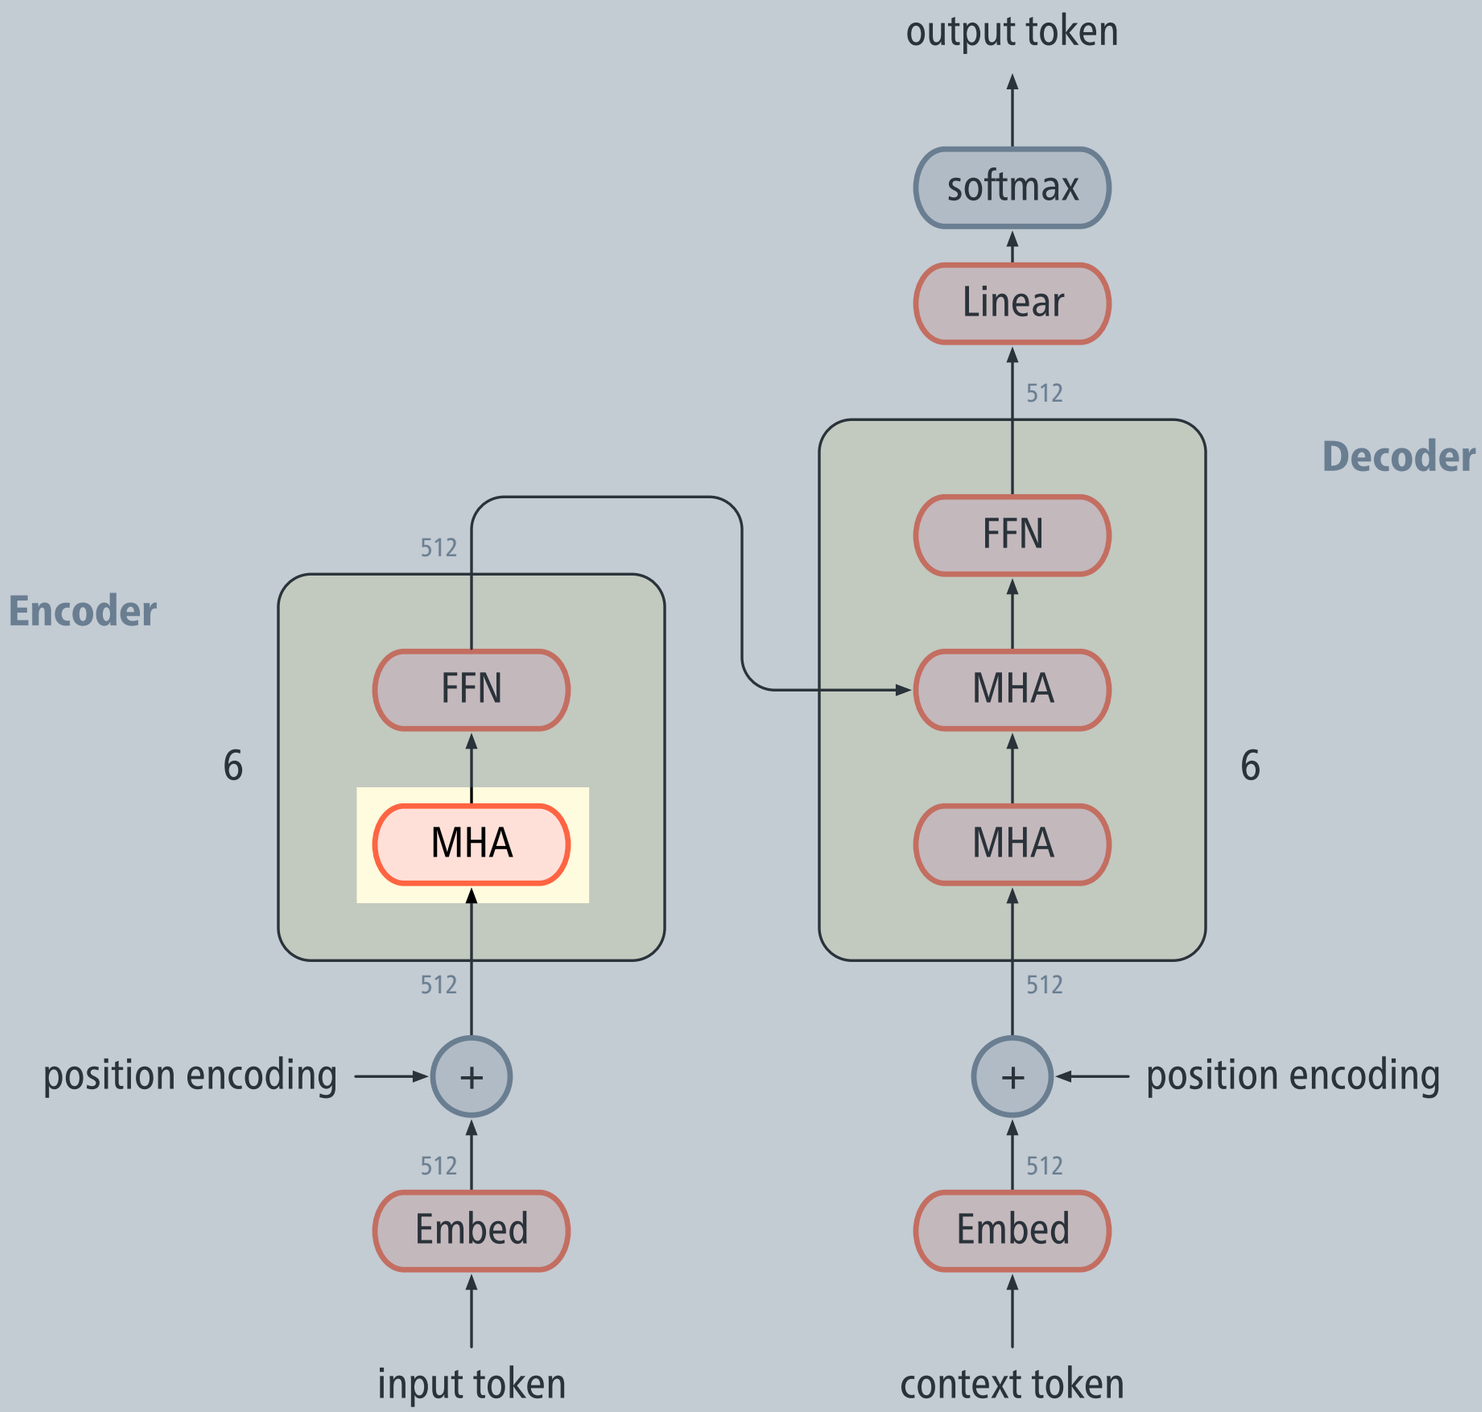
output token
softmax
Linear
512
Decoder
FFN
MHA
6
512
Encoder
FFN
MHA
6
512
position encoding
+
512
Embed
input token
512
+
position encoding
512
Embed
context token

Multi-head attention in the encoder

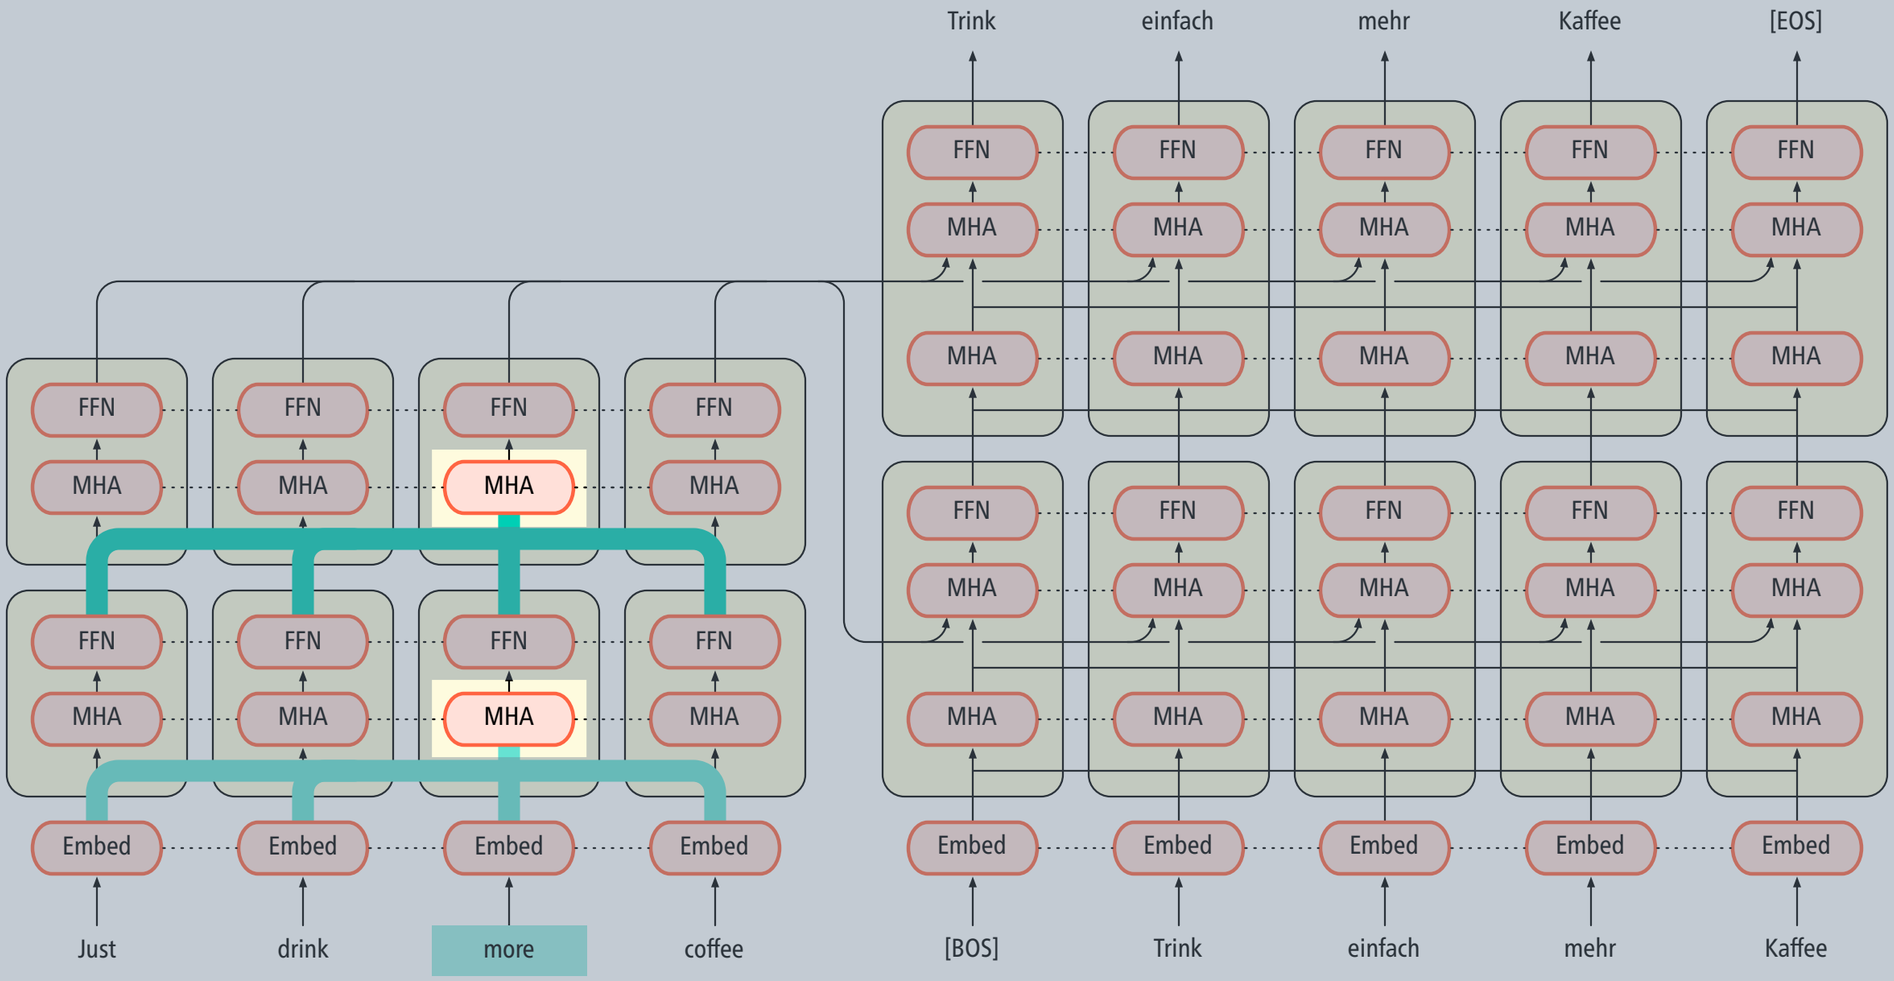

Multi-head attention in the encoder

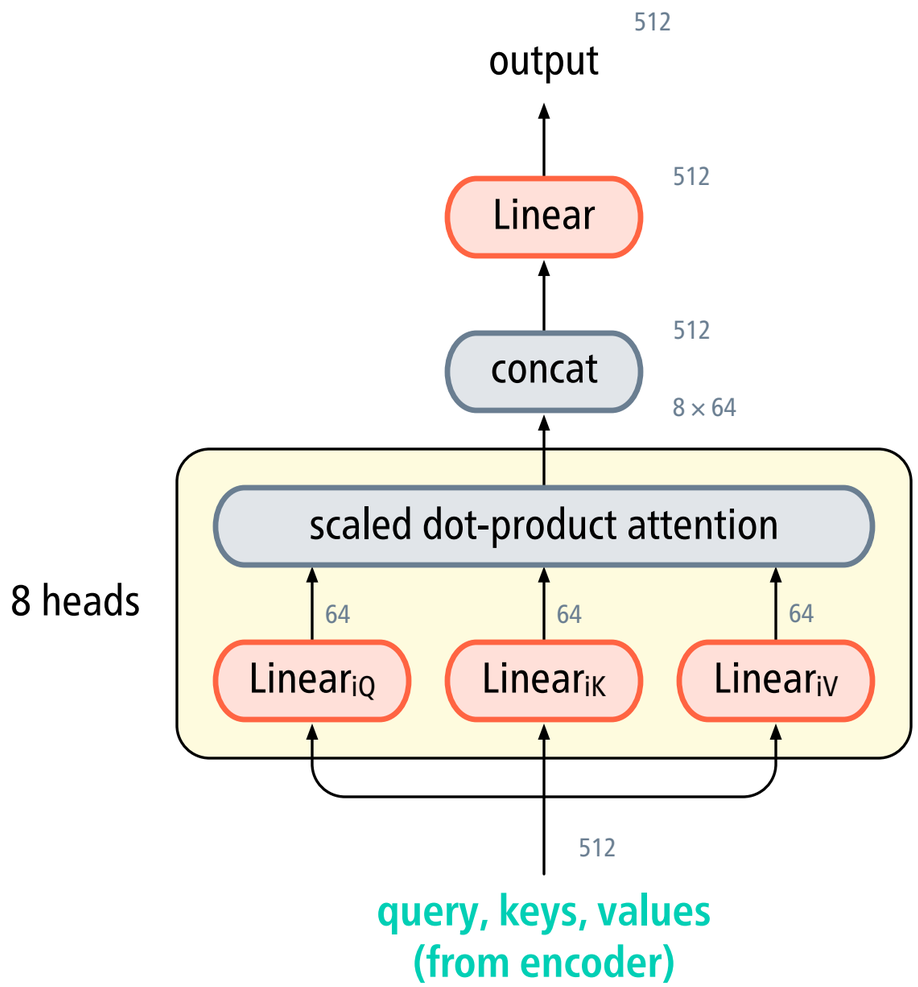

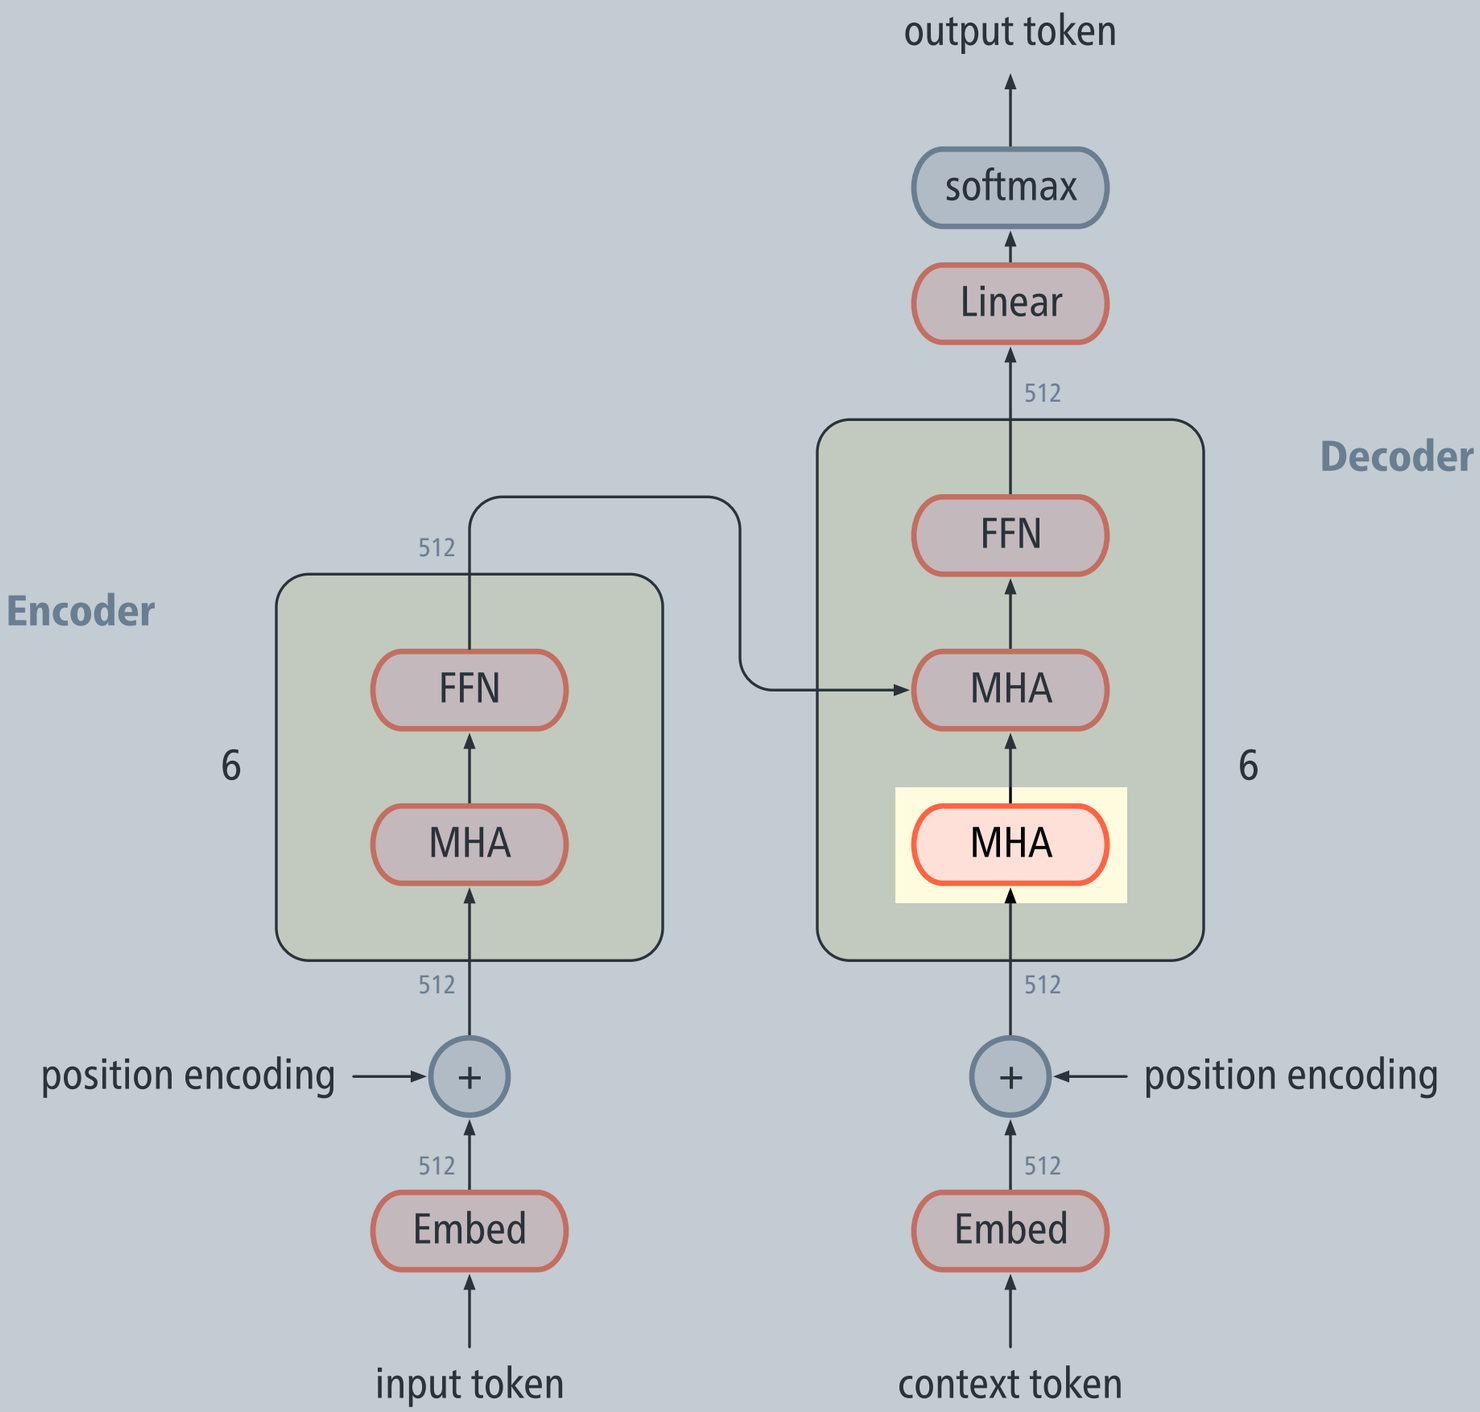
output token
softmax
Linear
512
Decoder
FFN
MHA
6
512
Encoder
FFN
MHA
6
512
position encoding
+
512
Embed
input token
512
+
position encoding
512
Embed
context token

Multi-head attention in the decoder

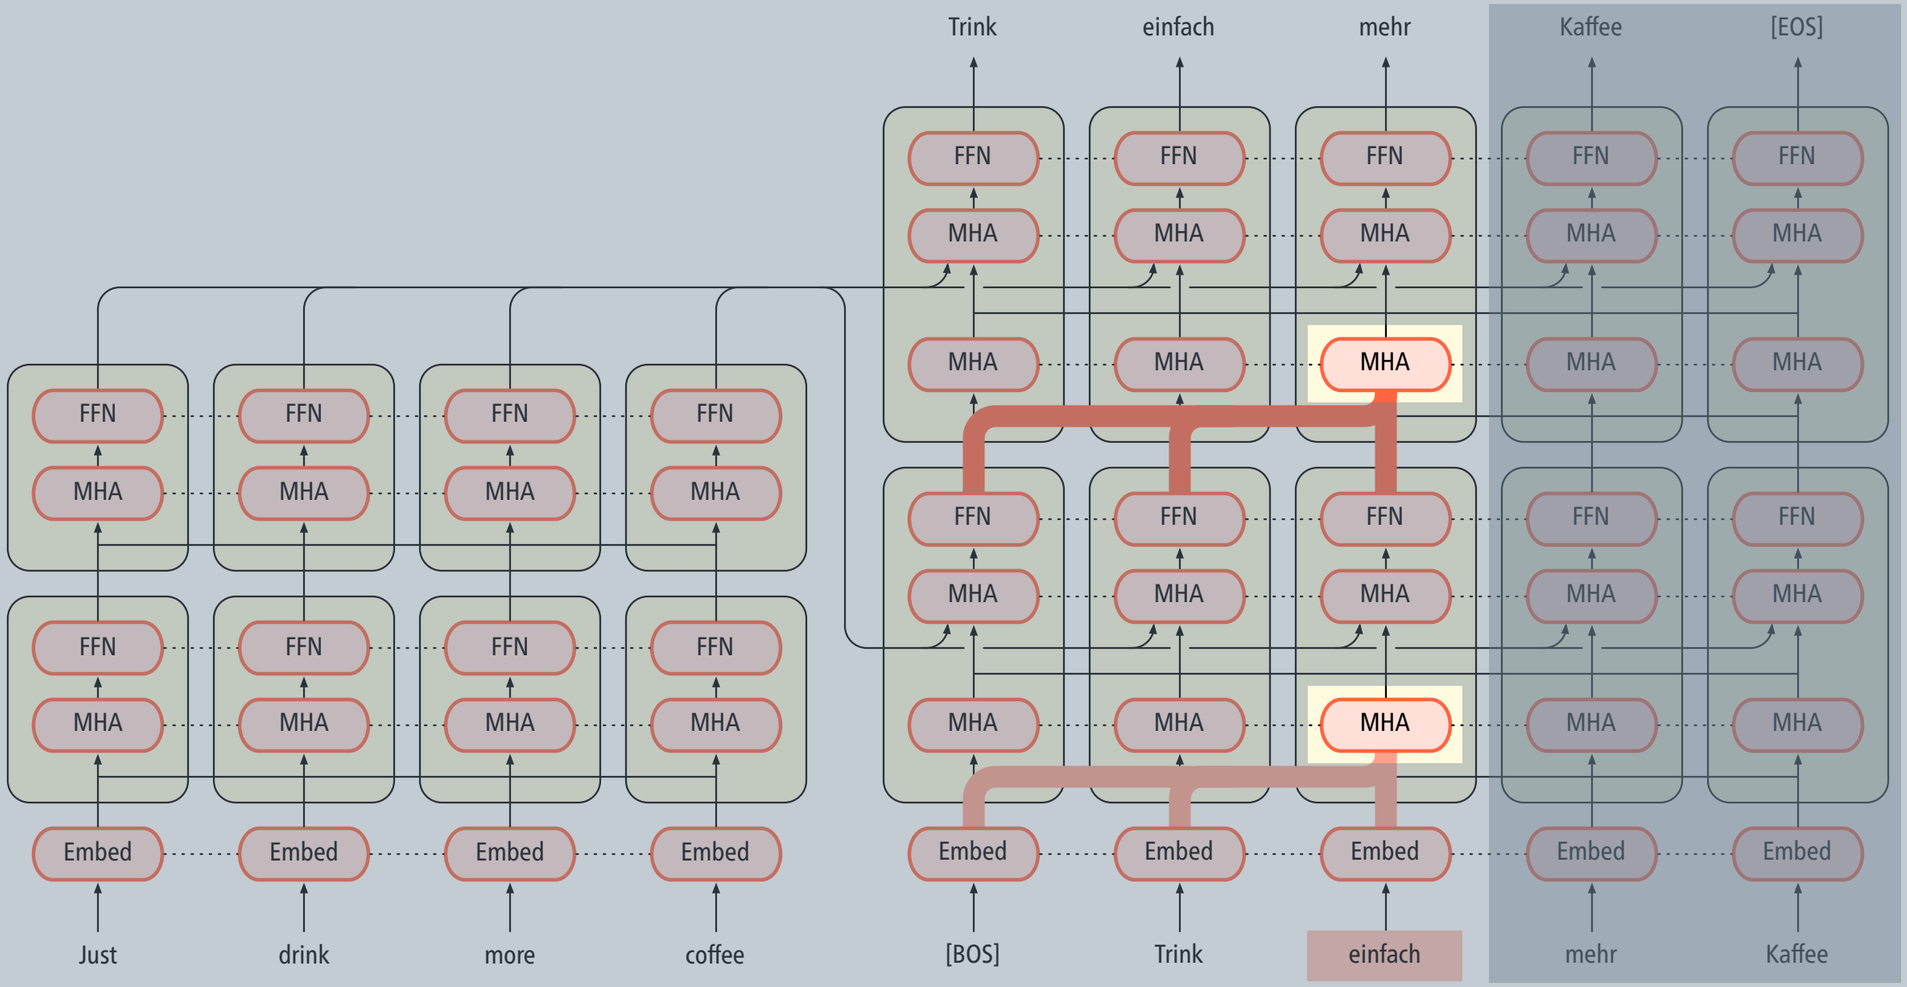

Multi-head attention in the decoder

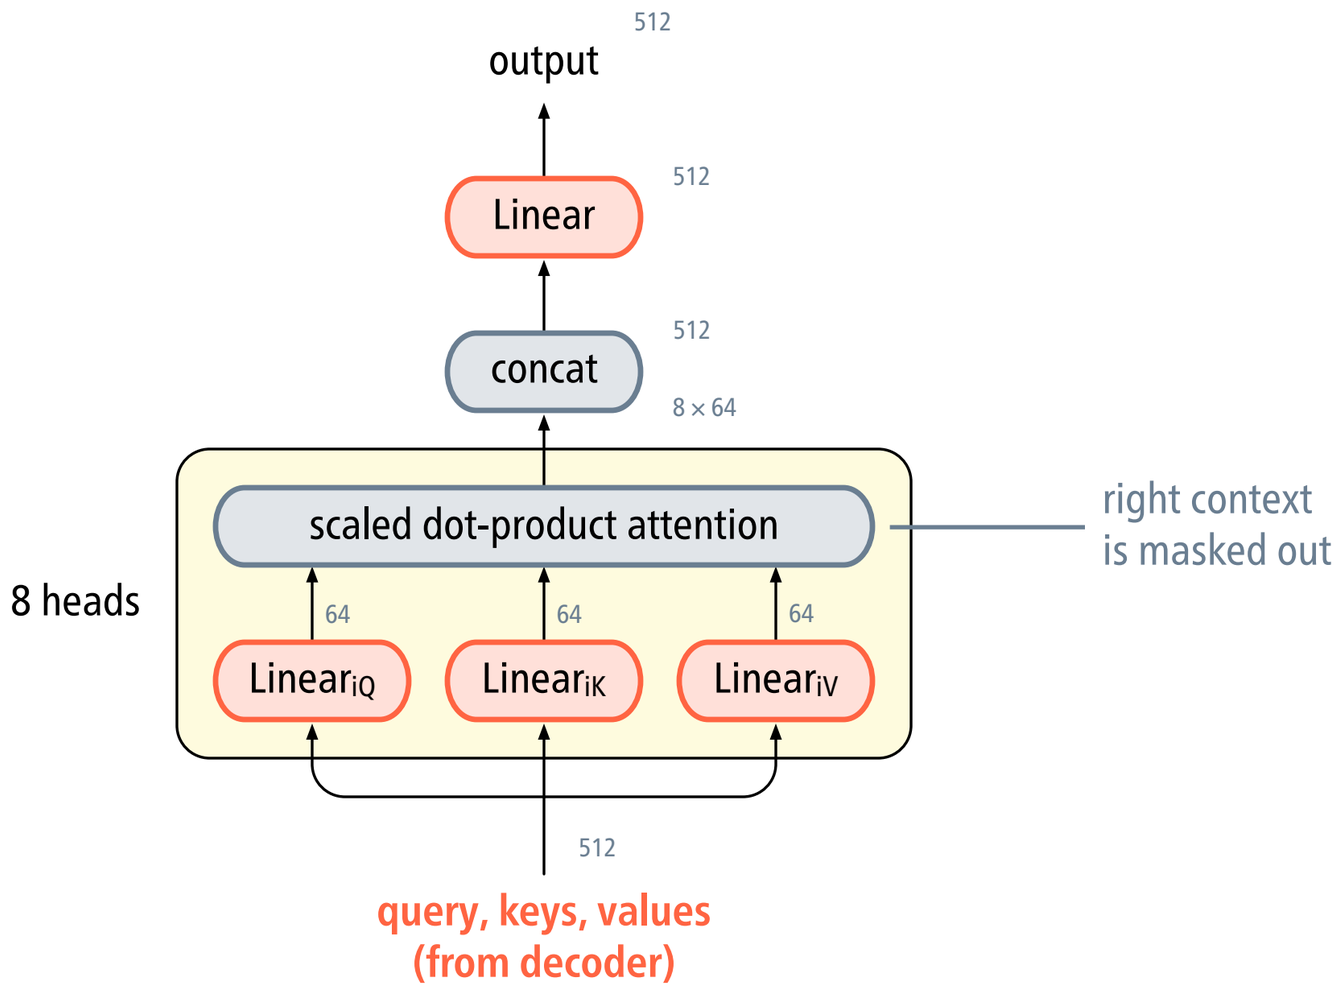

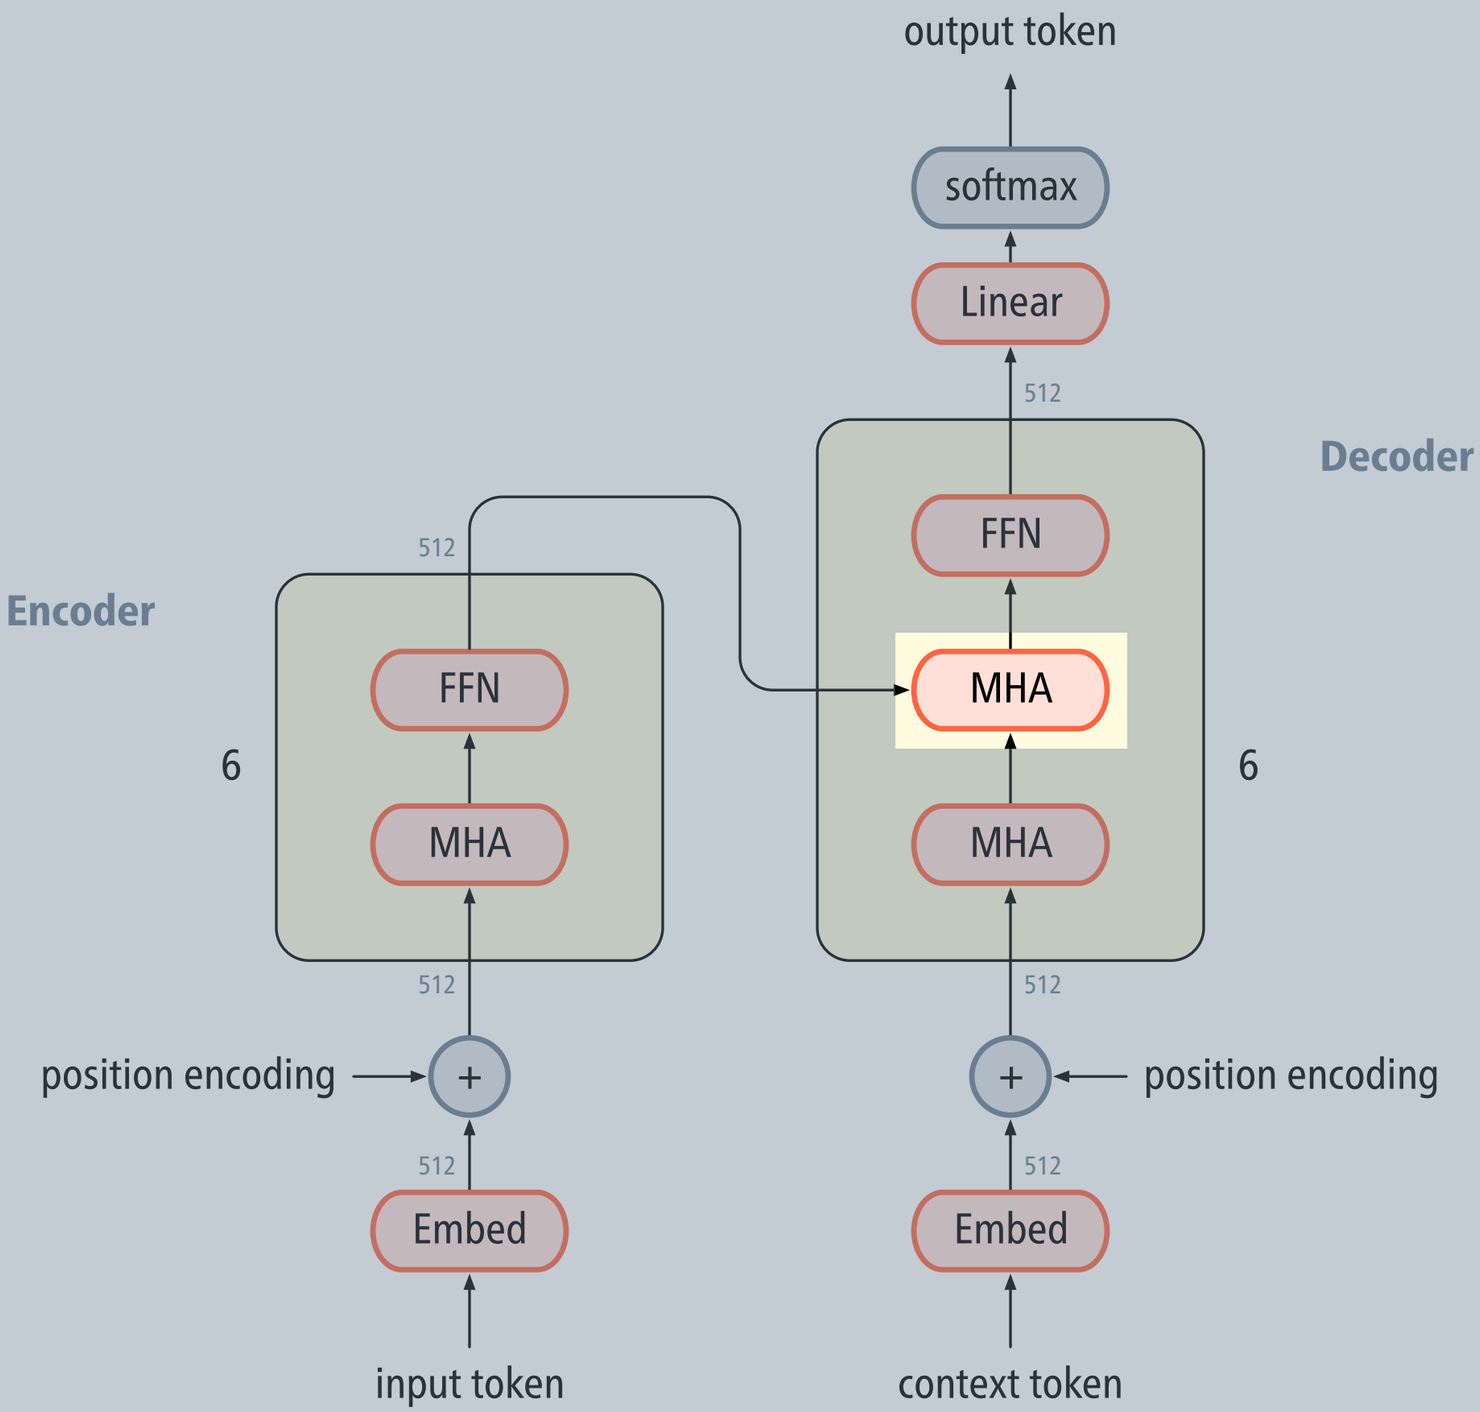
output token
softmax
Linear
512
Decoder
FFN
MHA
6
512
Encoder
FFN
MHA
6
512
position encoding
+
512
Embed
input token
512
+
position encoding
512
Embed
context token

Cross-attention

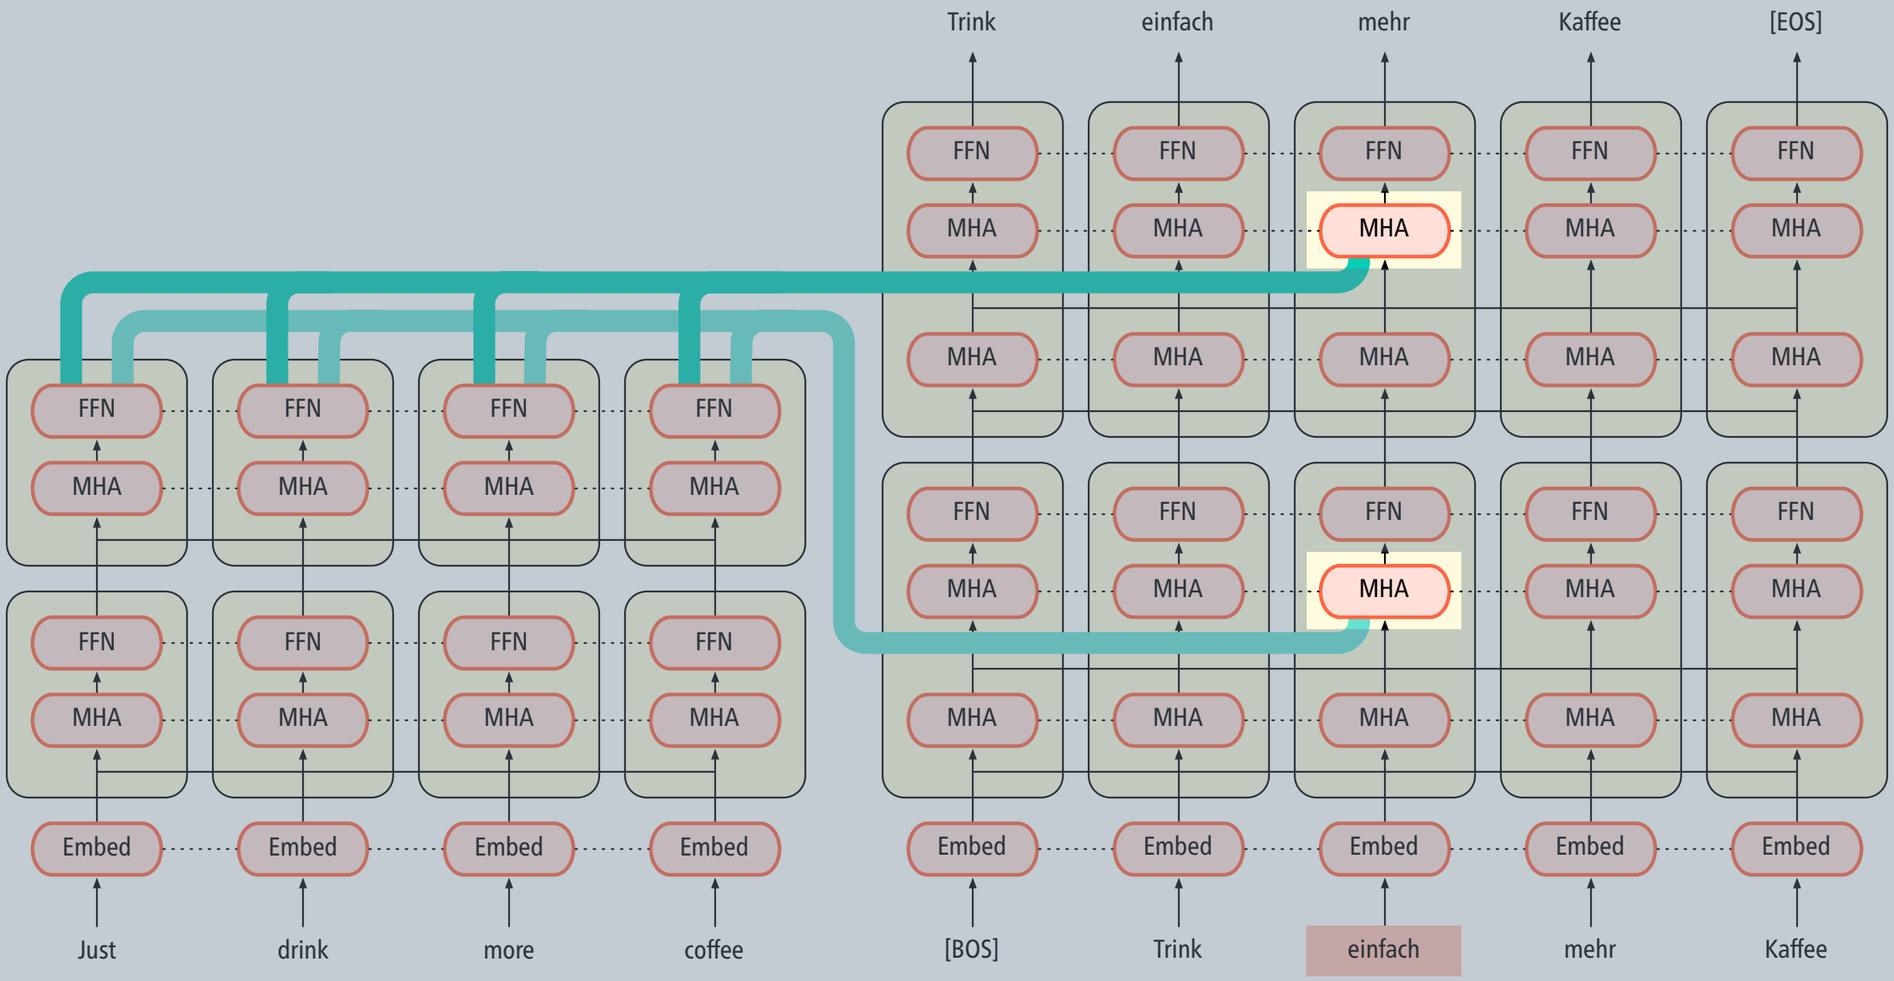

Cross-attention

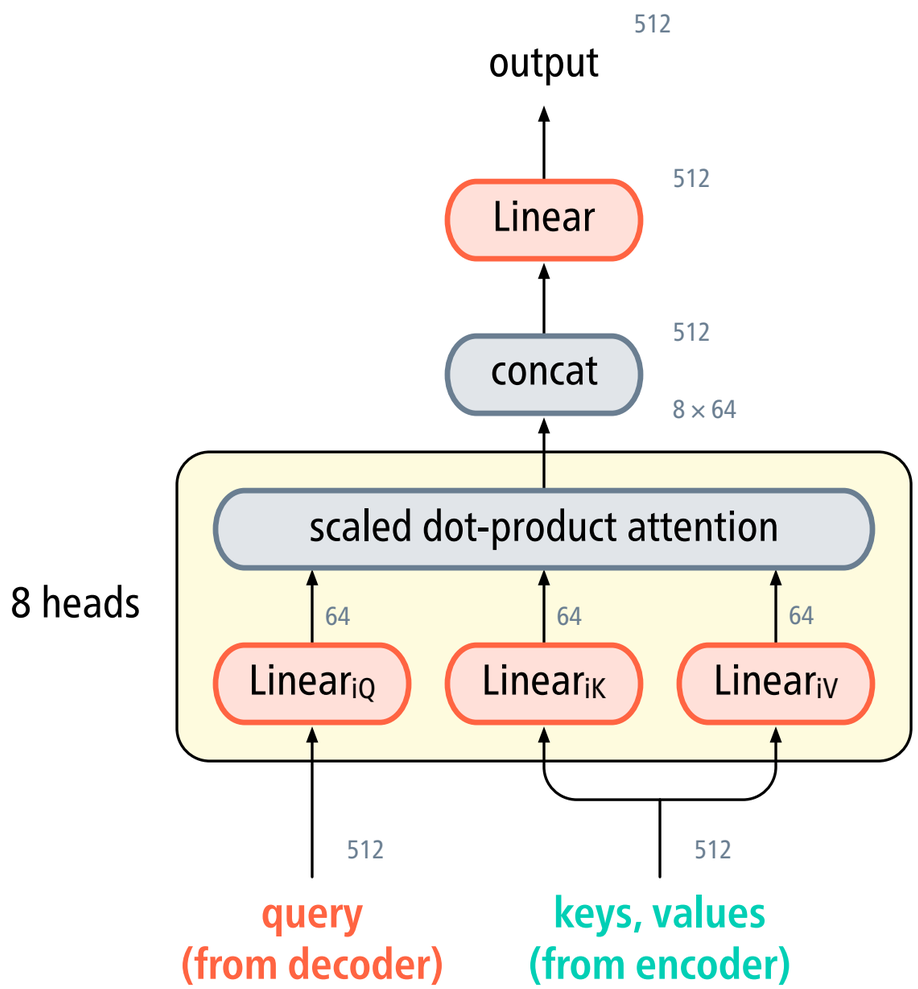

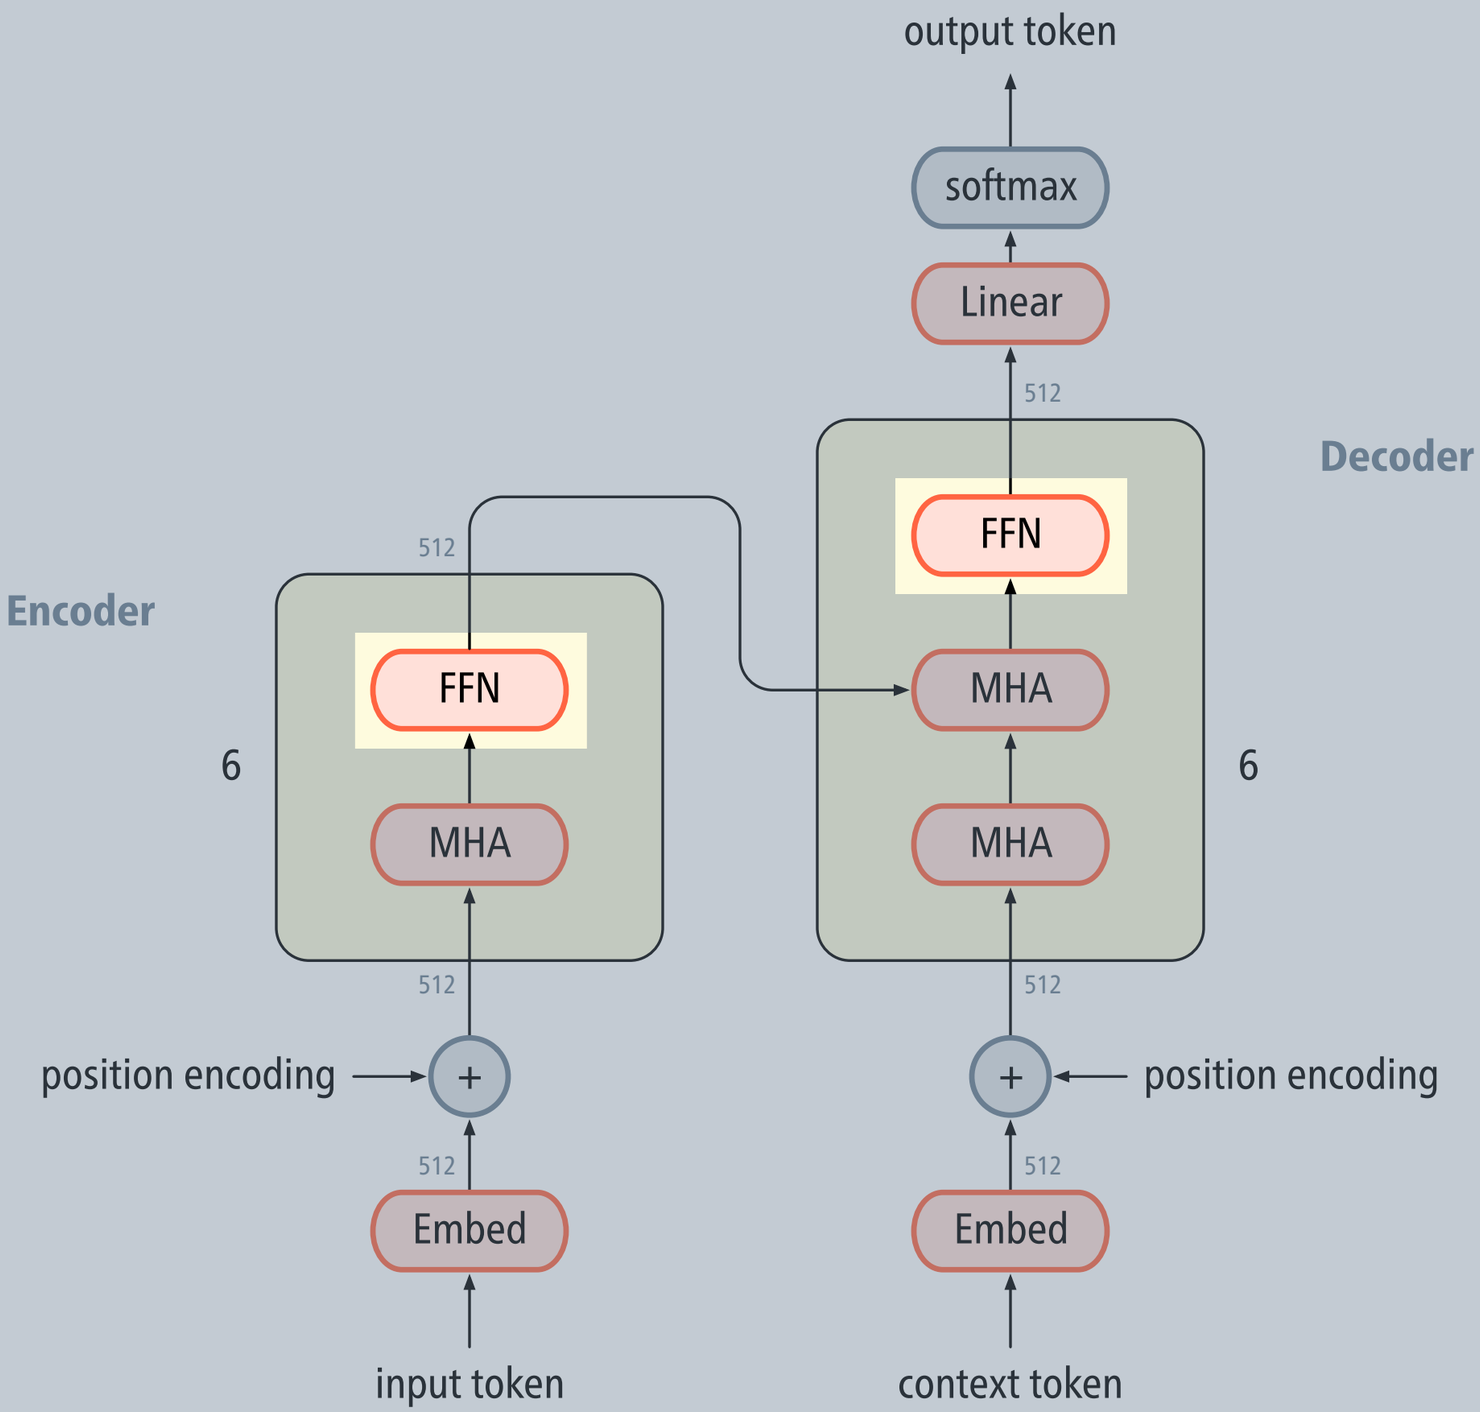
output token
softmax
Linear
512
Decoder
FFN
MHA
6
512
Encoder
FFN
MHA
6
512
position encoding
+
512
Embed
input token
512
+
position encoding
512
Embed
context token

Position-wise feed-forward network

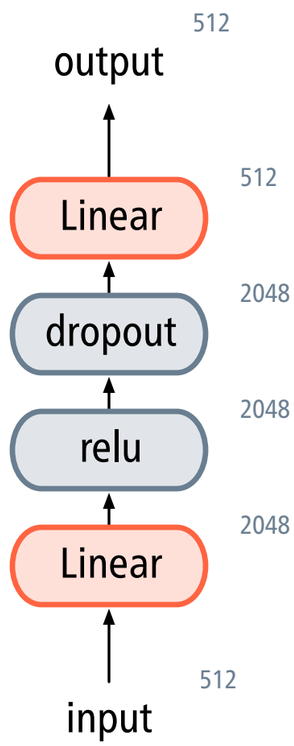

Parameters are shared across positions, but not across blocks.

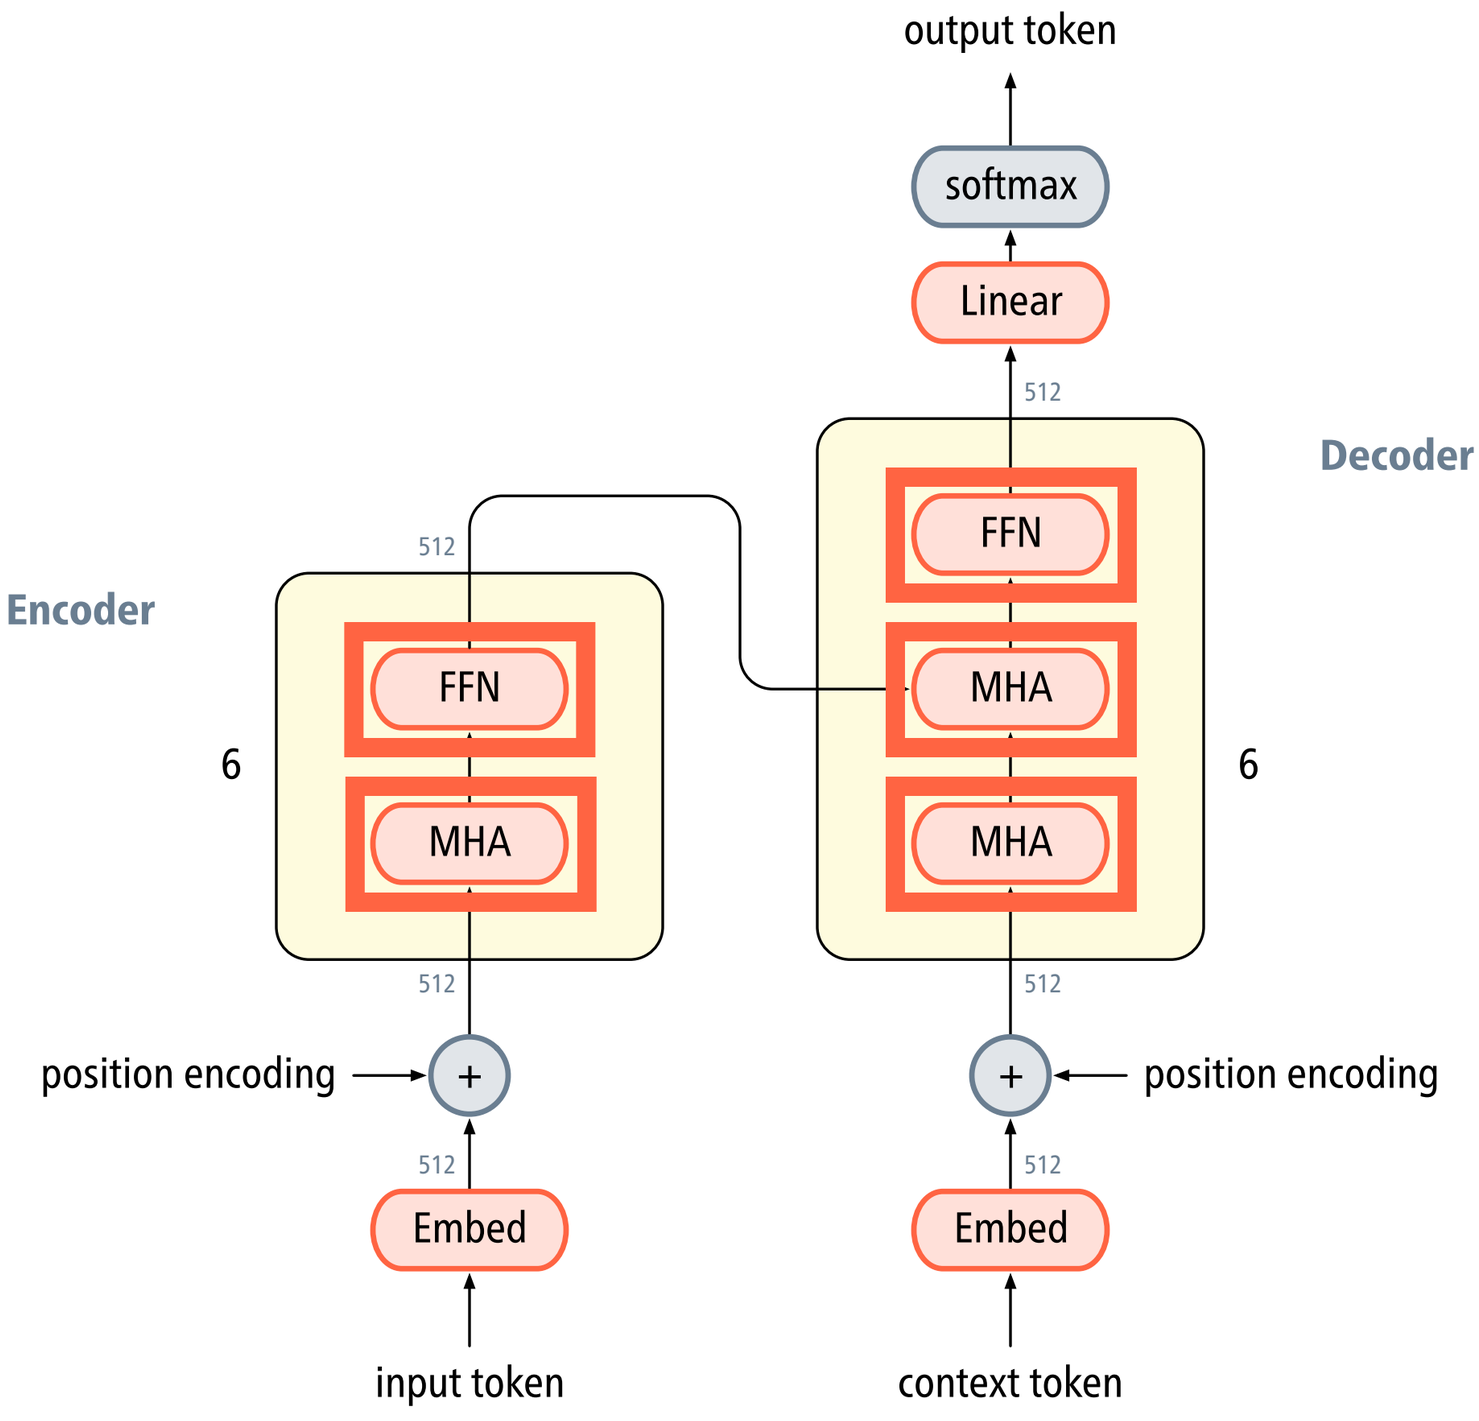
output token
softmax
Linear
512
Decoder
FFN
MHA
6
512
Encoder
FFN
MHA
6
512
position encoding
+
512
Embed
input token
512
+
position encoding
512
Embed
context token

Normalise-and-add wrapper

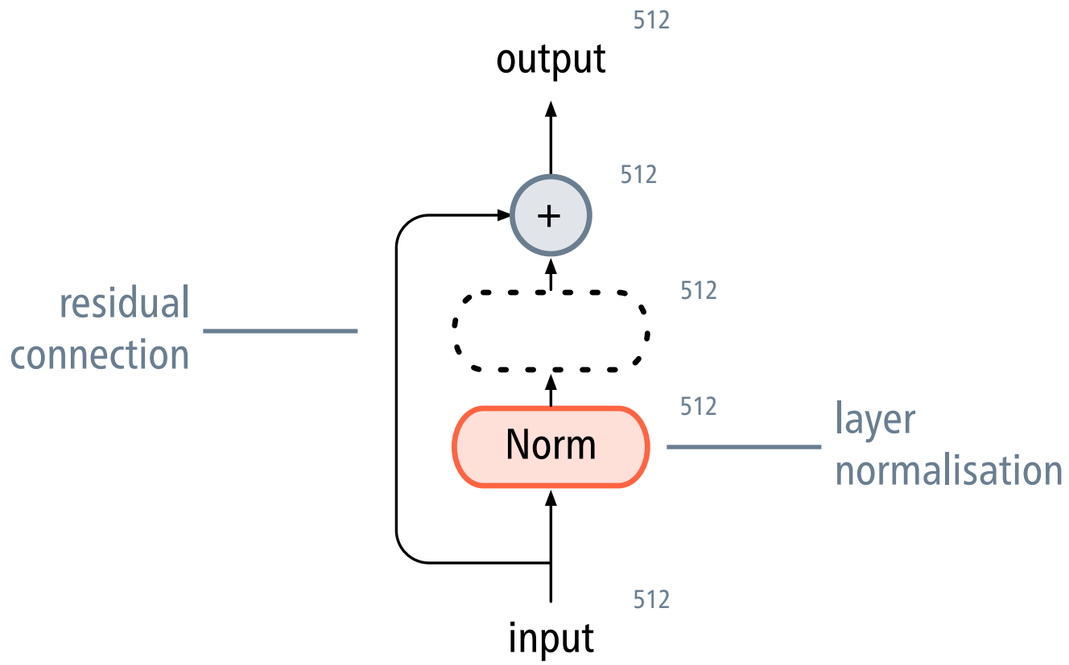

gain parameter

bias parameter

$$\boldsymbol{y} = \frac{\boldsymbol{g}}{\sigma + \varepsilon} \odot (\boldsymbol{x} - \mu) + \boldsymbol{b}$$

$$\mu = \frac{1}{|\boldsymbol{x}|} \sum_{i=1}^{|\boldsymbol{x}|} x_i$$

$$\sigma = \sqrt{\frac{1}{|\boldsymbol{x}|} \sum_{i=1}^{|\boldsymbol{x}|} (x_i - \mu)^2}$$

Ba et al. (2016)

Further details

- Token representations are defined on word pieces computed using byte-pair encoding.
- Embeddings are augmented by position encodings.

  approximate encoding of absolute positions

- Training the model uses several tricks related to batching, masking, loss, and regularisation.

  for details and PyTorch code, see the ‘Annotated Transformer’

Translation performance

	BLEU	FLOPs
GNMT + RL (Wu et al., 2016)	39.92	$1.4 \cdot 10^{20}$
ConvS2S (Gehring et al., 2017)	40.46	$1.5 \cdot 10^{20}$
MoE (Shazeer et al., 2017)	40.56	$1.2 \cdot 10^{20}$
Transformer (big model)	41.80	$2.4 \cdot 10^{19}$

BLEU score and training cost (FLOPs) on the English-to-French newstest2014 test data | Vaswani et al. (2017)