Knowledge Injection Strategies for Large Language Models: An Overview

Benjamin Wolff
ZB MED Information Center for Life Sciences, Cologne

ZB MED Information Centre for Life Sciences

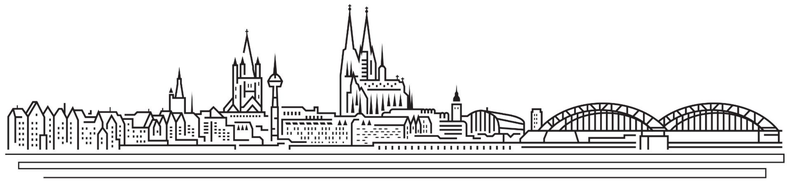

Knowledge in LLMs

Knowledge Injection Strategies for LLMs

Where does the knowledge come from?

- Pre-training on large scale data
- Goal: Understand and generate natural language
- But: Also absorbs knowledge during pre-training
- Why:
  - Every sentence necessarily carries knowledge
  - Knowledge is implicitly learned

WIKIPEDIA

Where does the knowledge come from?

- Pre-training on large scale data
- Goal: Understand and generate natural language
- But: Also absorbs knowledge during pre-training
- Why:
  - Every sentence necessarily carries knowledge
  - Knowledge is implicitly learned
- Knowledge is a by-product of language-acquisition

WIKIPEDIA

Where does the knowledge come from?

- Pre-training on large scale data
- Goal: Understand and generate natural language
- But: Also absorbs knowledge during pre-training
- Why:
  - Every sentence necessarily carries knowledge
  - Knowledge is implicitly learned

WIKIPEDIA

- Knowledge is a by-product of language-acquisition
- Problems:
  - In LLMs knowledge and language are deeply intertwined
  - Language is vague, facts and knowledge are precise
  - Highly specific concepts

How to inject Knowledge into LLMs?

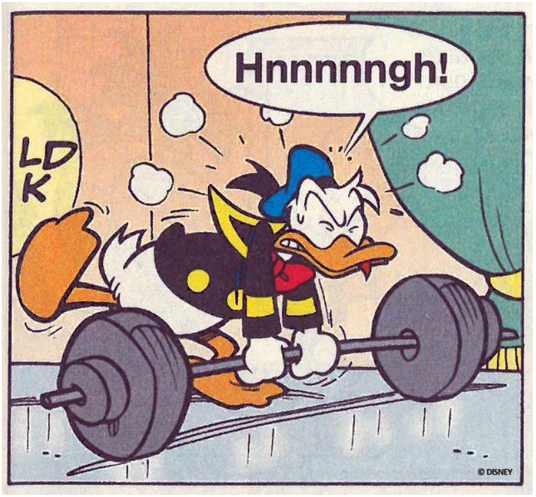

- **Heavyweight Pre-Training**
  - Large-scale data (Billions of examples)
  - Large-scale hardware-resources (especially GPUs/TPUs)
  - Energy consumption

- **Finetuning**
  - Injecting new knowledge modifies parameters
  - Previous knowledge fades away
  - Catastrophic forgetting

- **Lightweight Knowledge-Injection**
  - Leaves the PTMs parameters unchanged
  - Works with smaller training datasets
  - Works in low-budget setting

How to inject Knowledge into LLMs?

- **Heavyweight Pre-Training**
  - Large-scale data (Billions of examples)
  - Large-scale hardware-resources (especially GPUs/TPUs)
  - Energy consumption

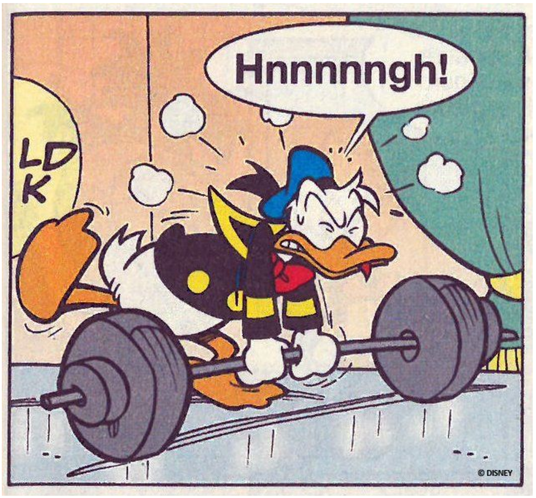

- **Finetuning**
  - Injecting new knowledge modifies parameters
  - Previous knowledge fades away
  - Catastrophic forgetting

- **Lightweight Knowledge-Injection**
  - Leaves the PTMs parameters unchanged
  - Works with smaller training datasets
  - Works in low-budget setting

What are we going to look at?

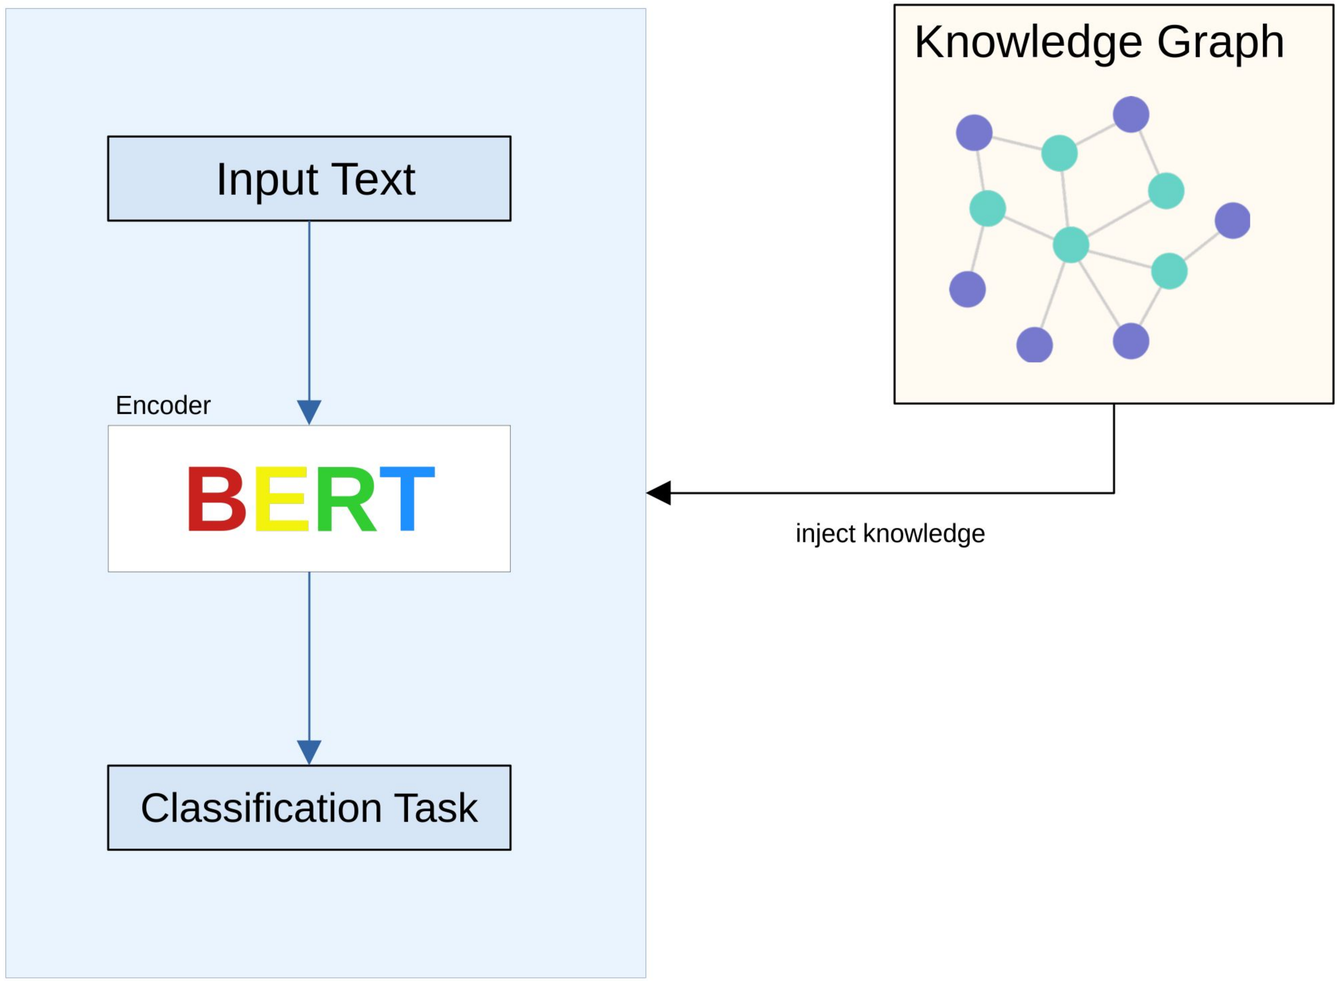

What are we going to look at?

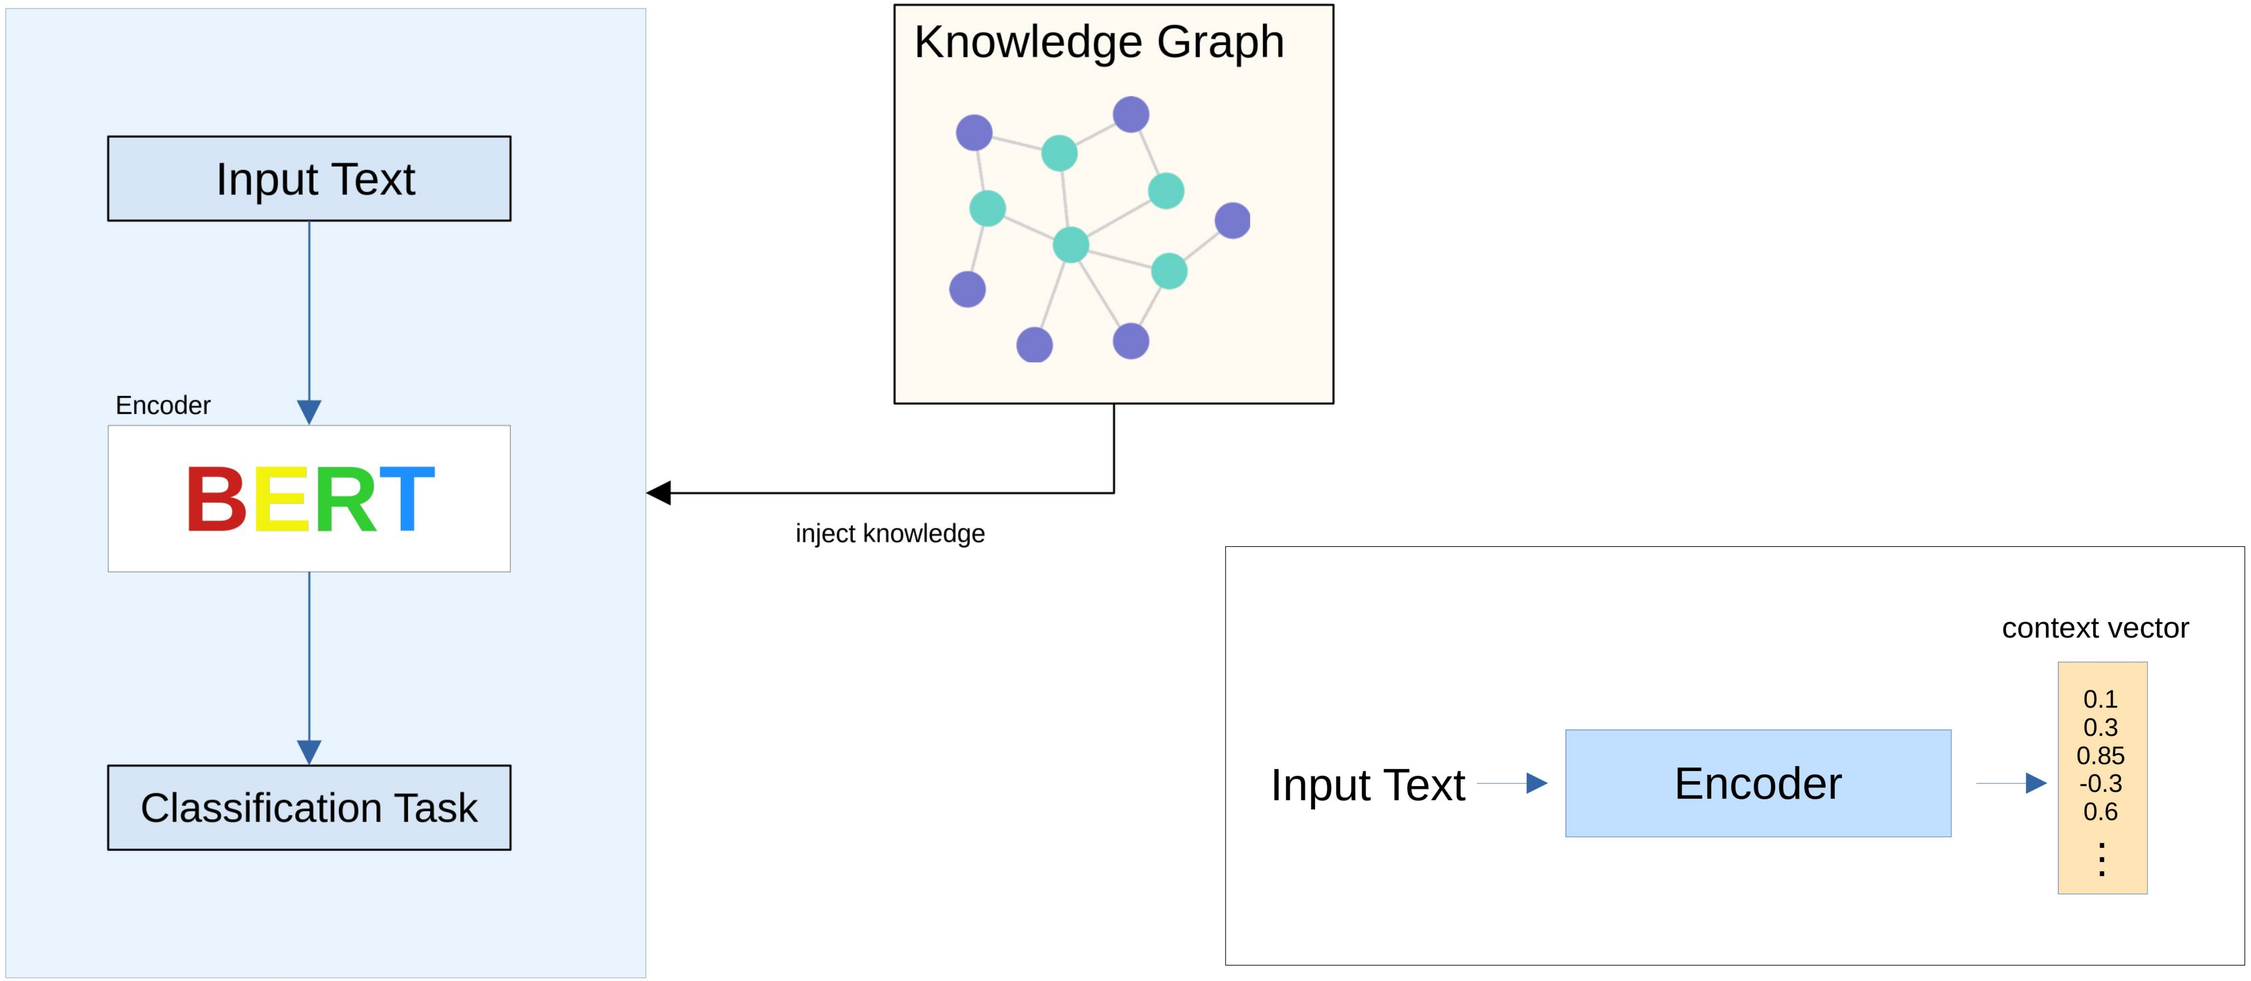

Direct Text Injection (DTI)

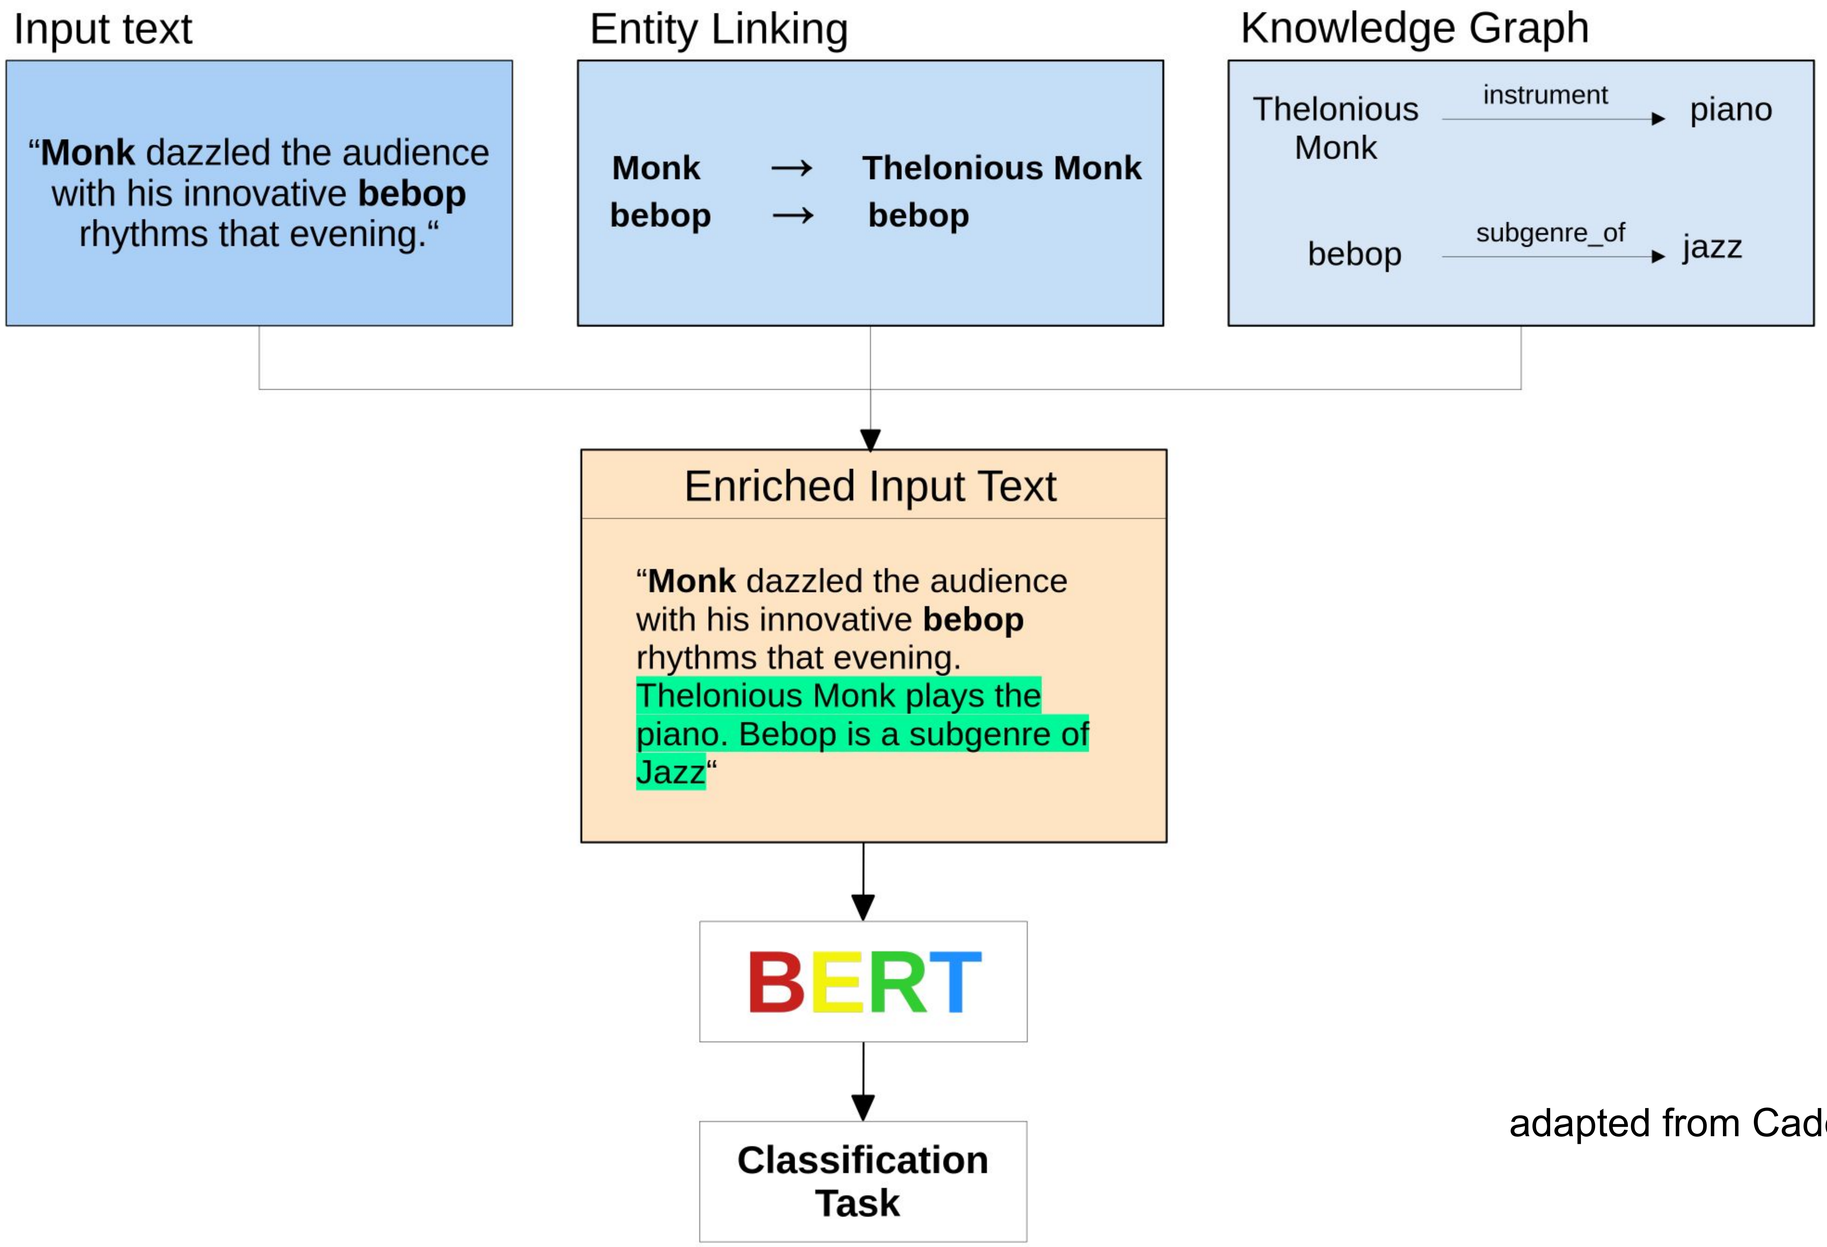

adapted from Cadeddu *et al.* (2024)

K-BERT

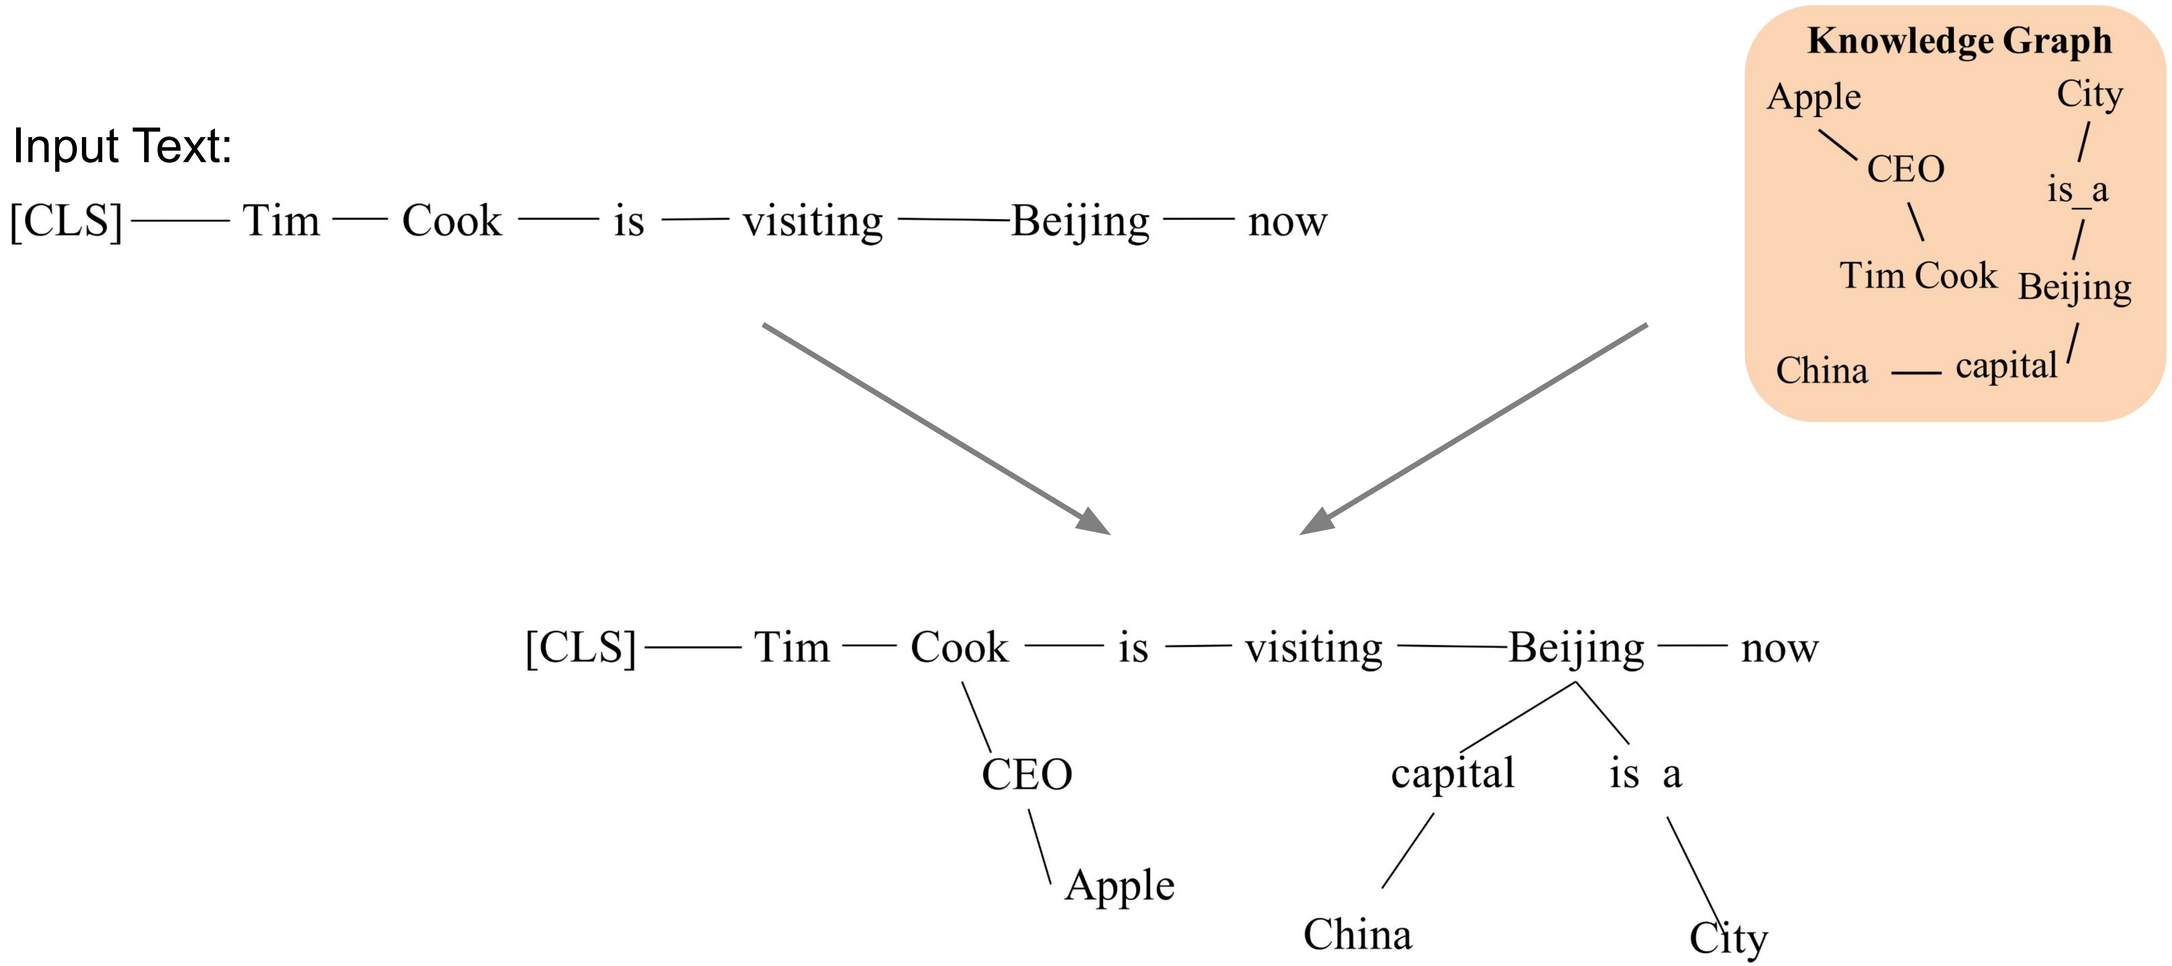

Liu *et al.* (2019)

K-BERT

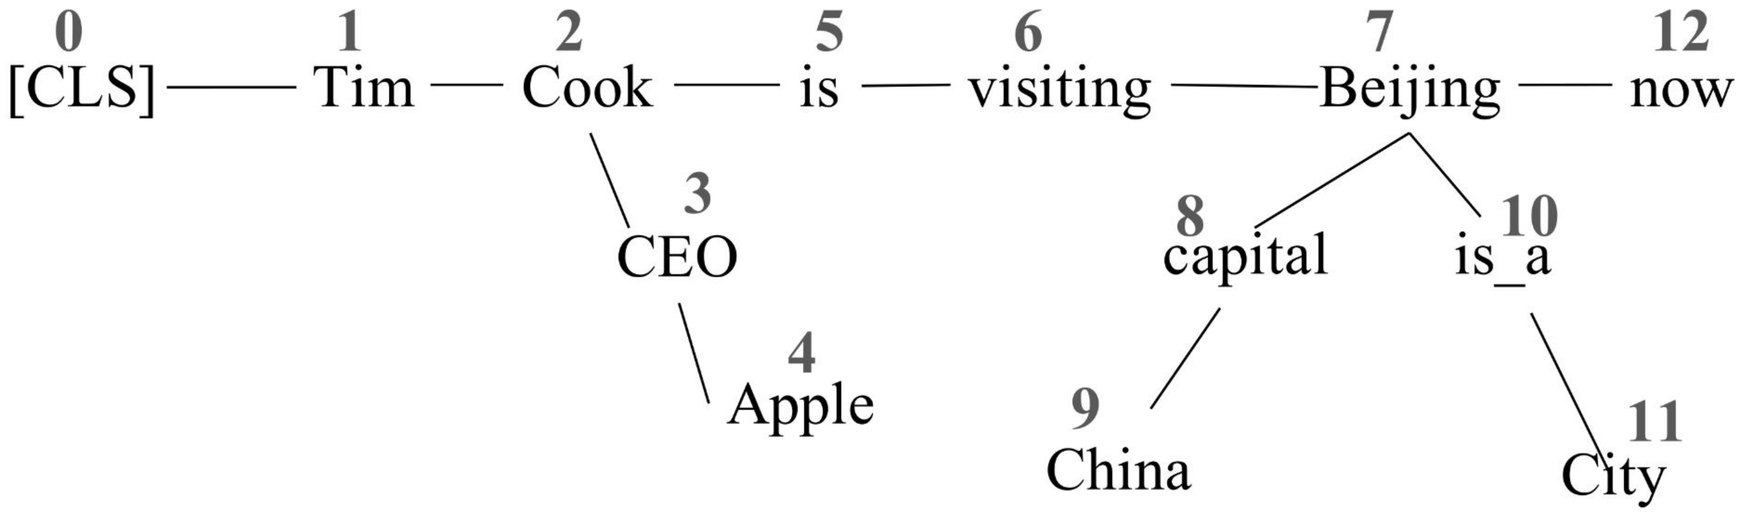

Liu *et al.* (2019)

K-BERT

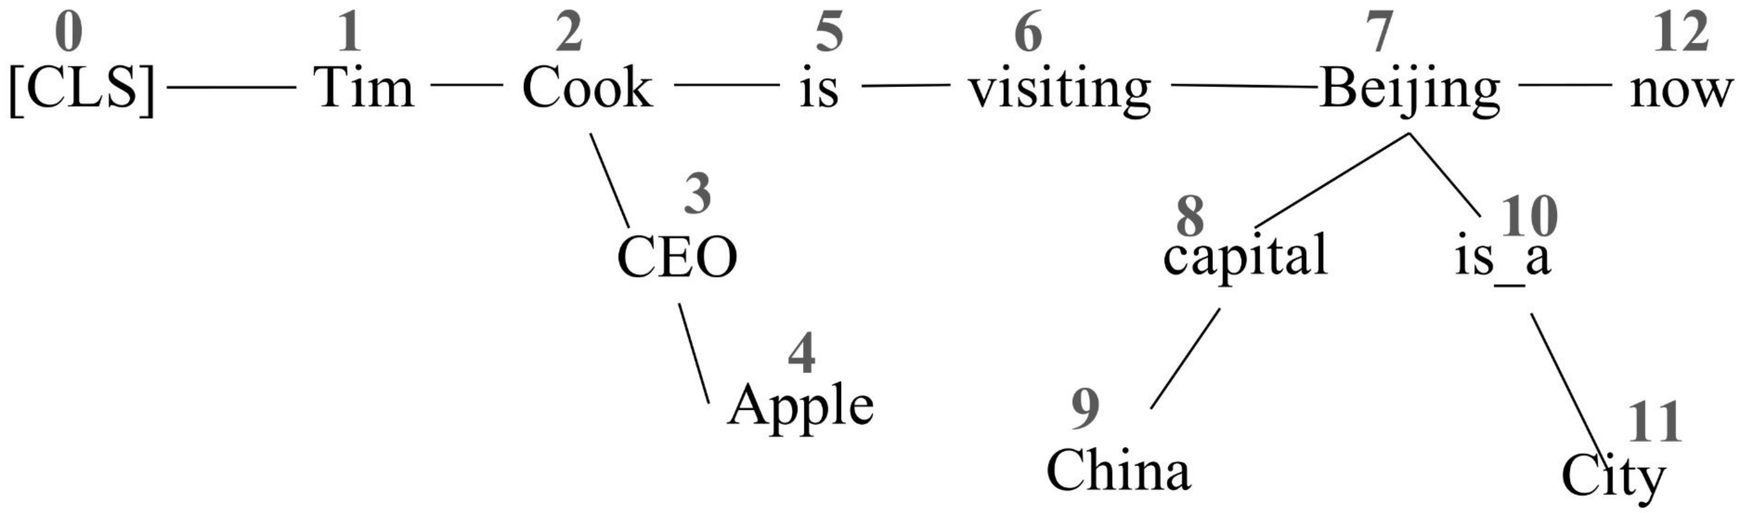

Tim Cook CEO Apple is visiting Beijing capital China is_a city now

Liu et al (2019)

K-BERT - Embedding Layer

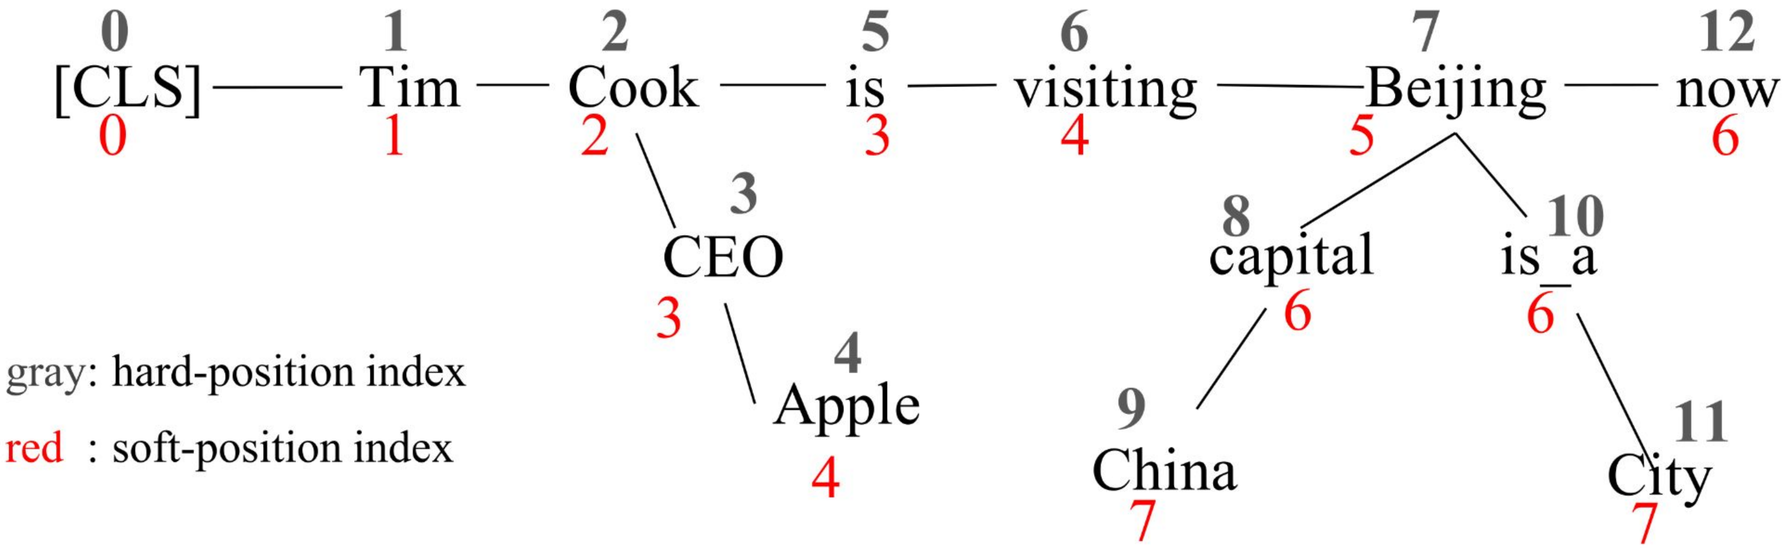

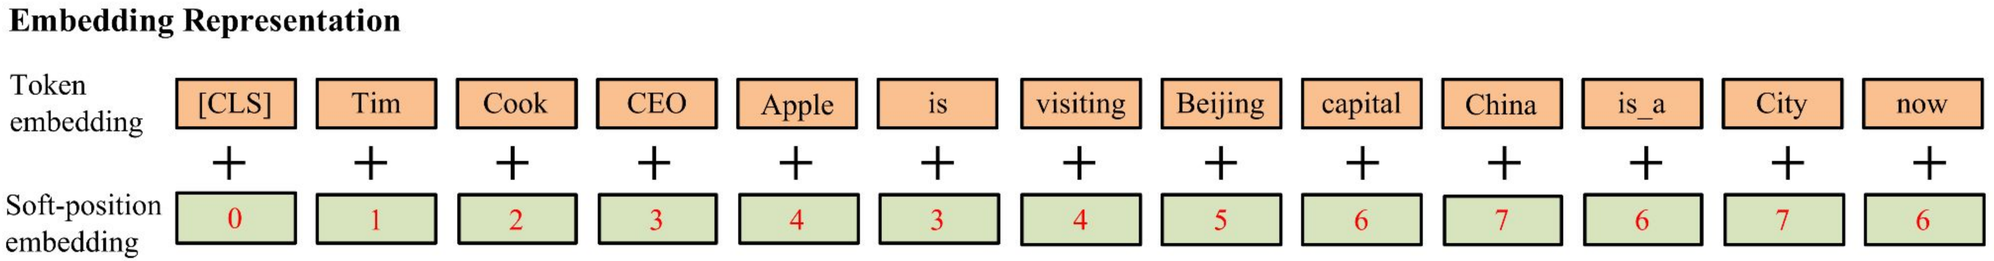

Liu *et al.* (2019)

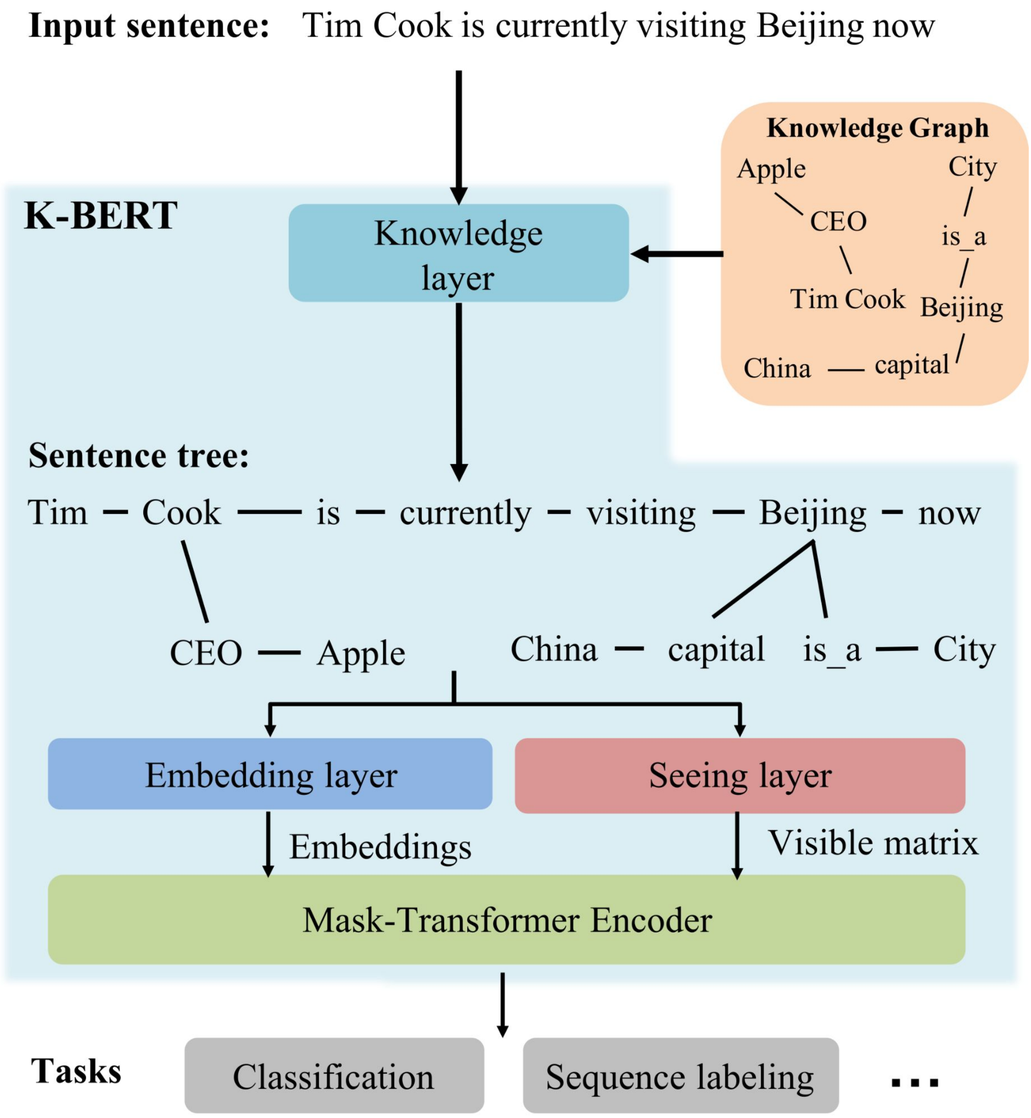

Liu *et al.* (2019)

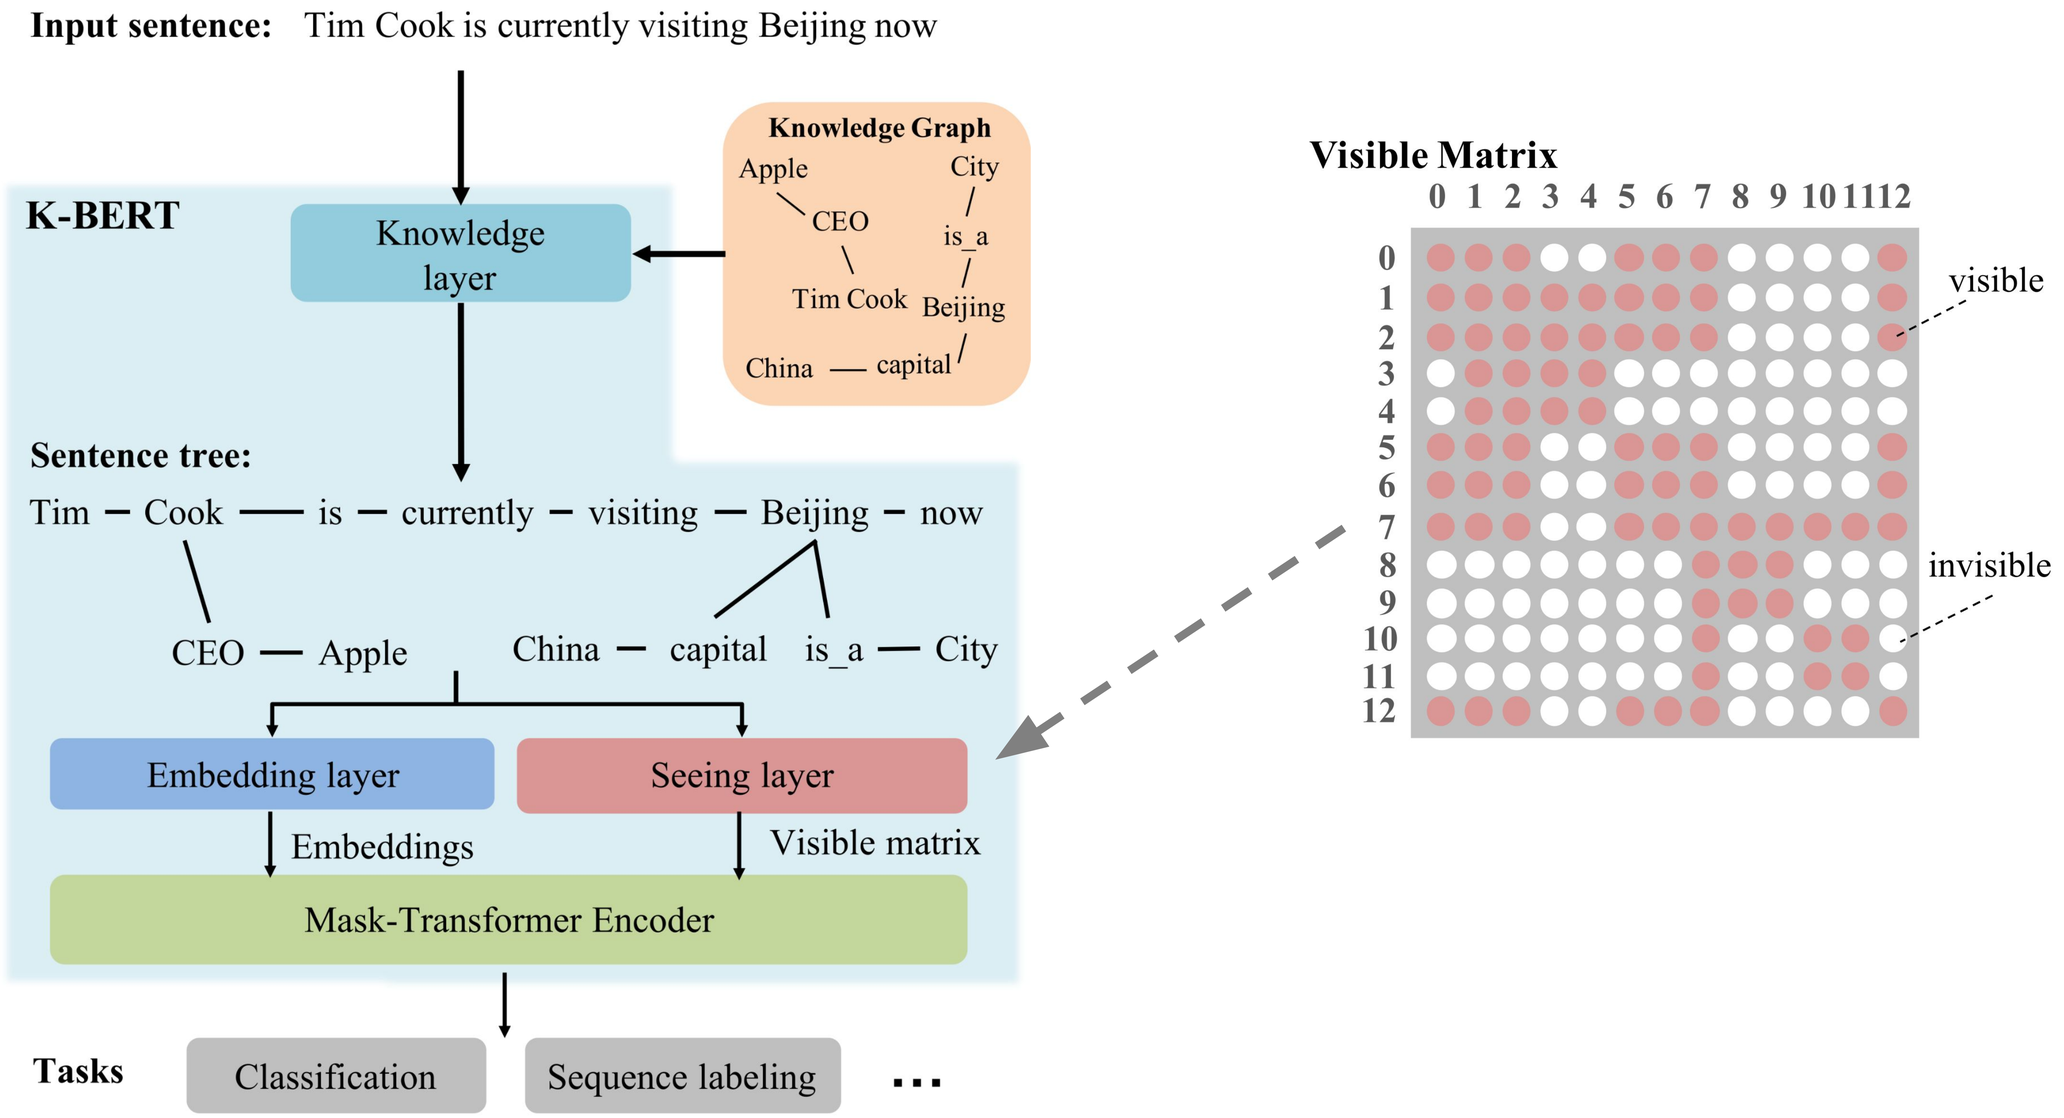

Liu *et al.* (2019)

Integrating Features with MLP

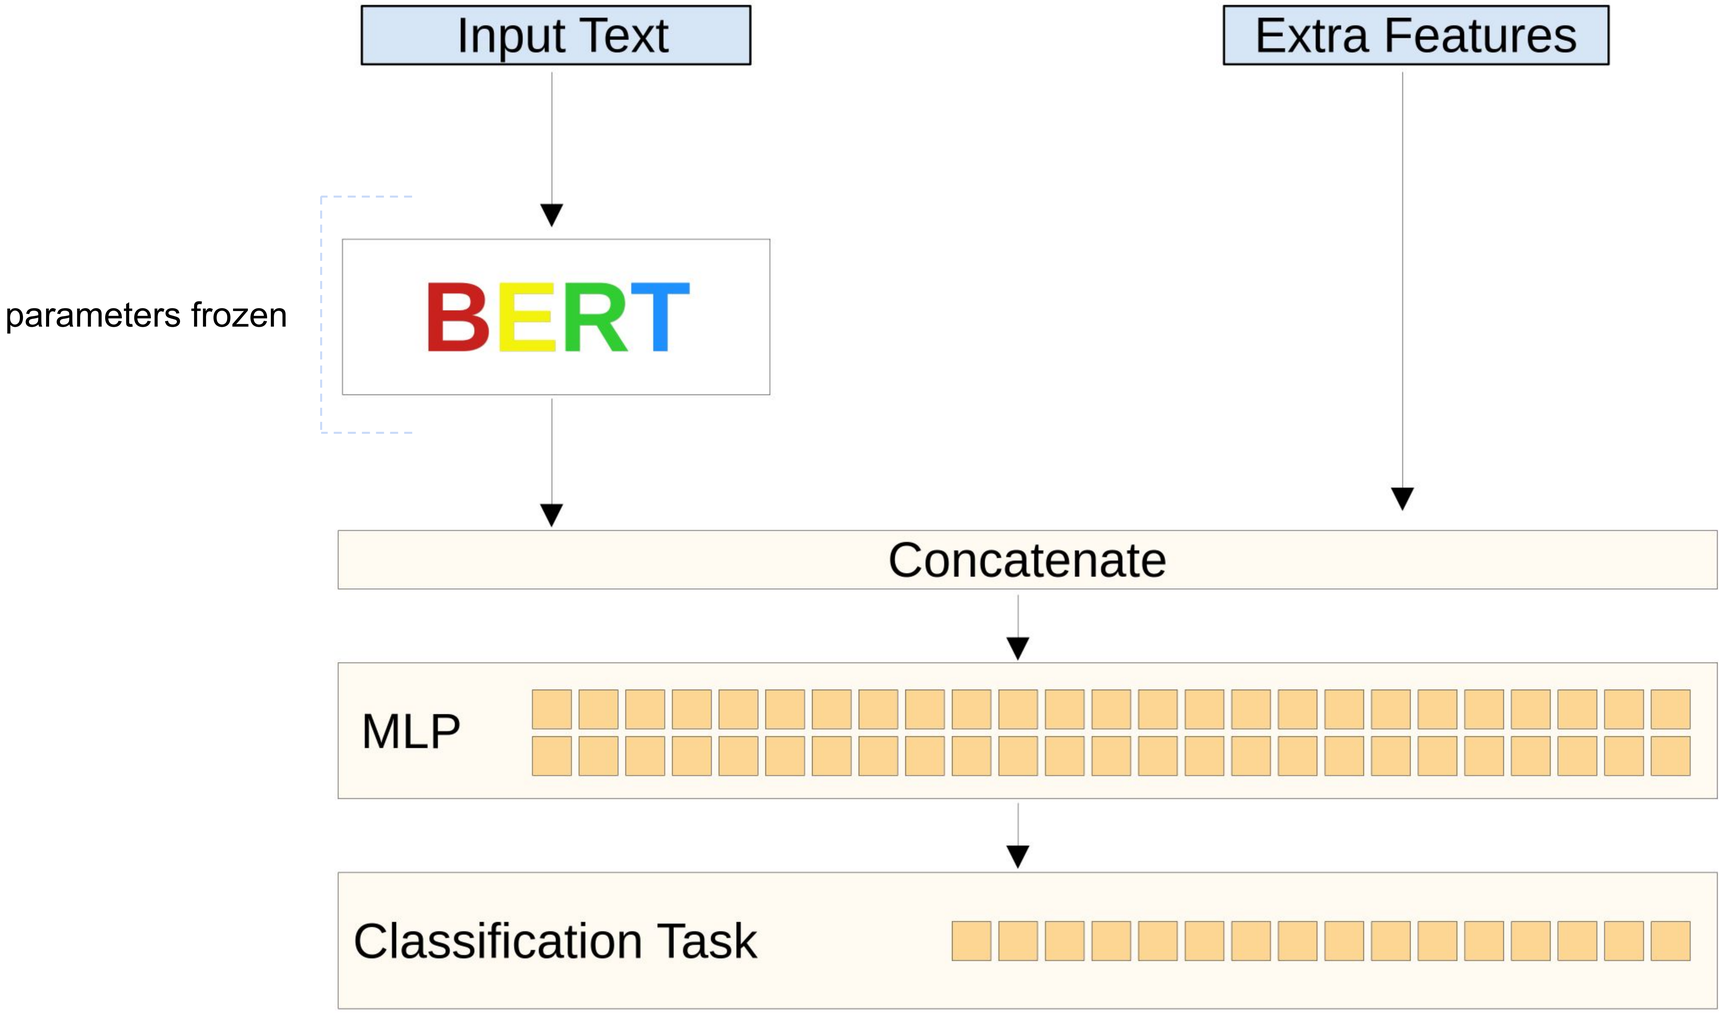

(Multilayer Perceptron)

adapted from Ostendorff *et al.* (2019)

Integrating Features with MLP

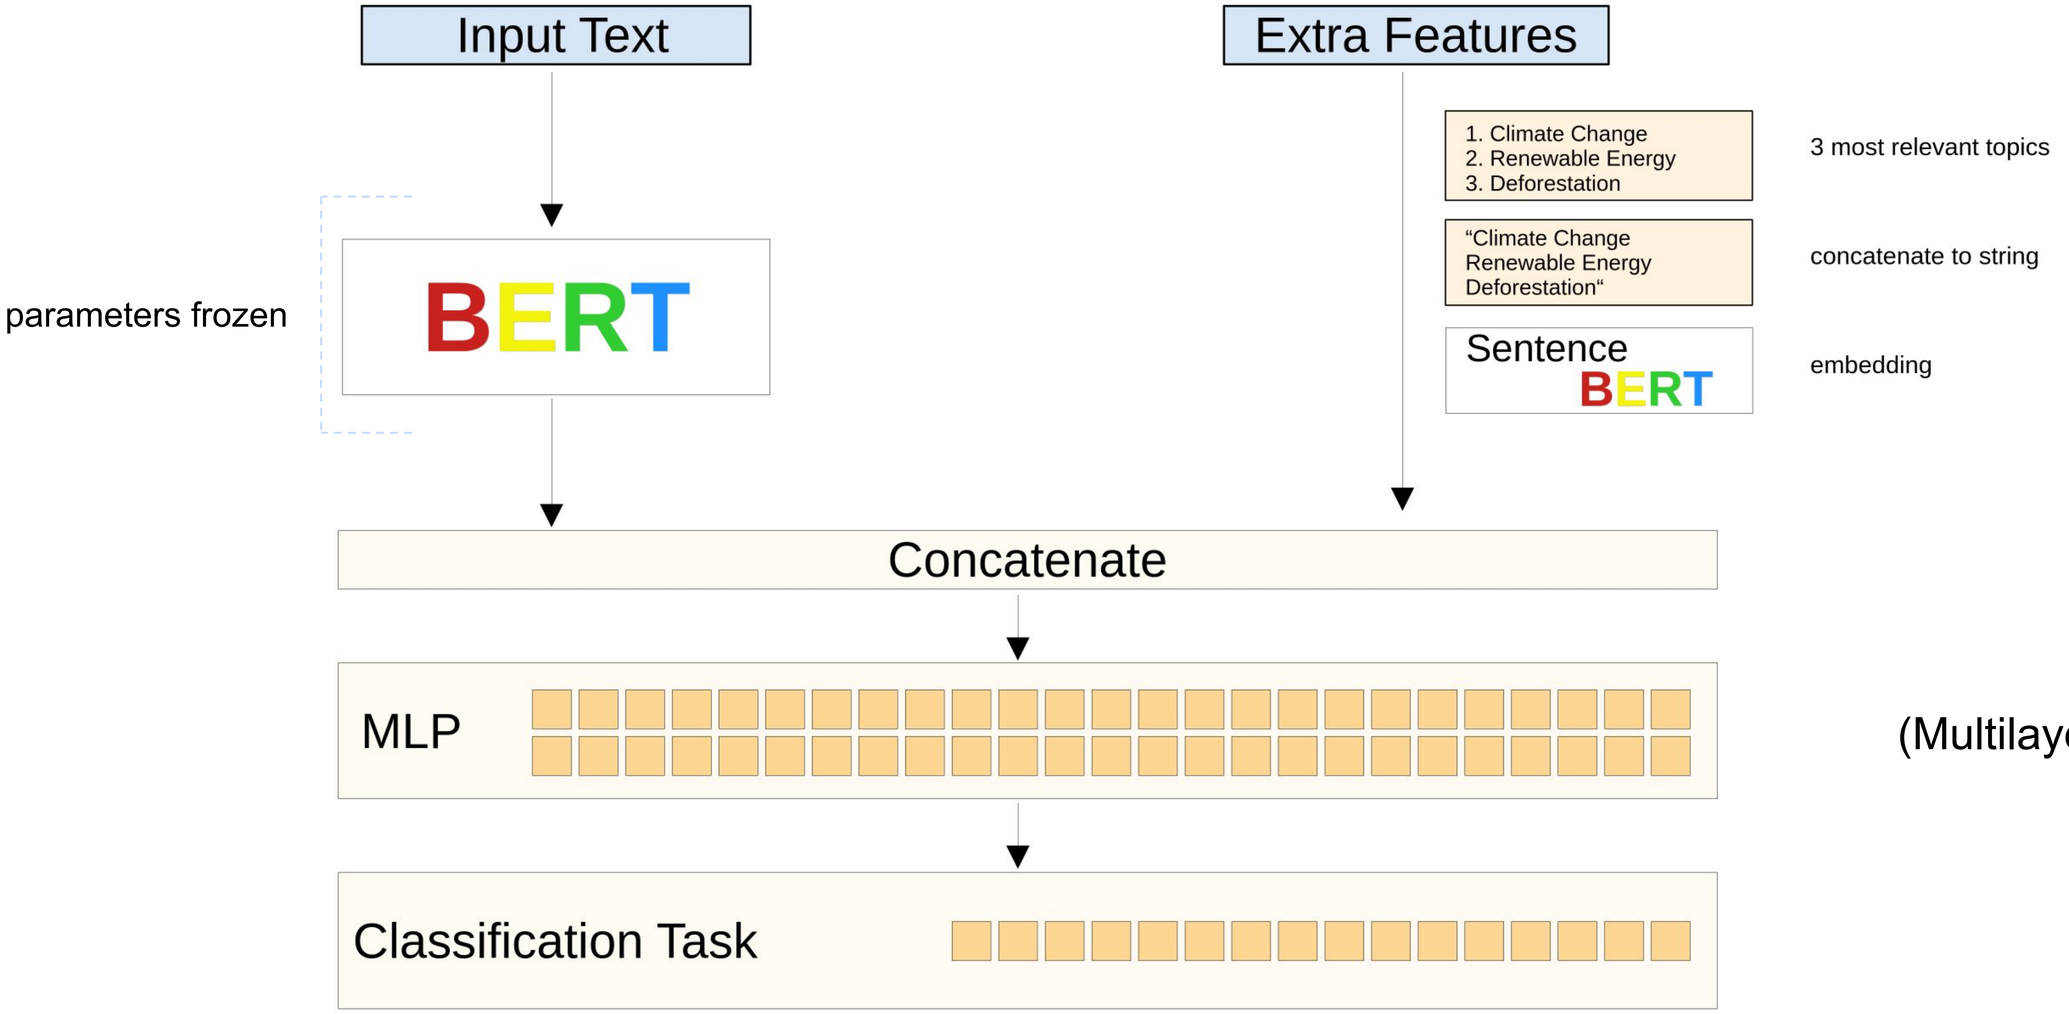

adapted from Ostendorff *et al.* (2019)

Knowledge Adapters (K-Adapter)

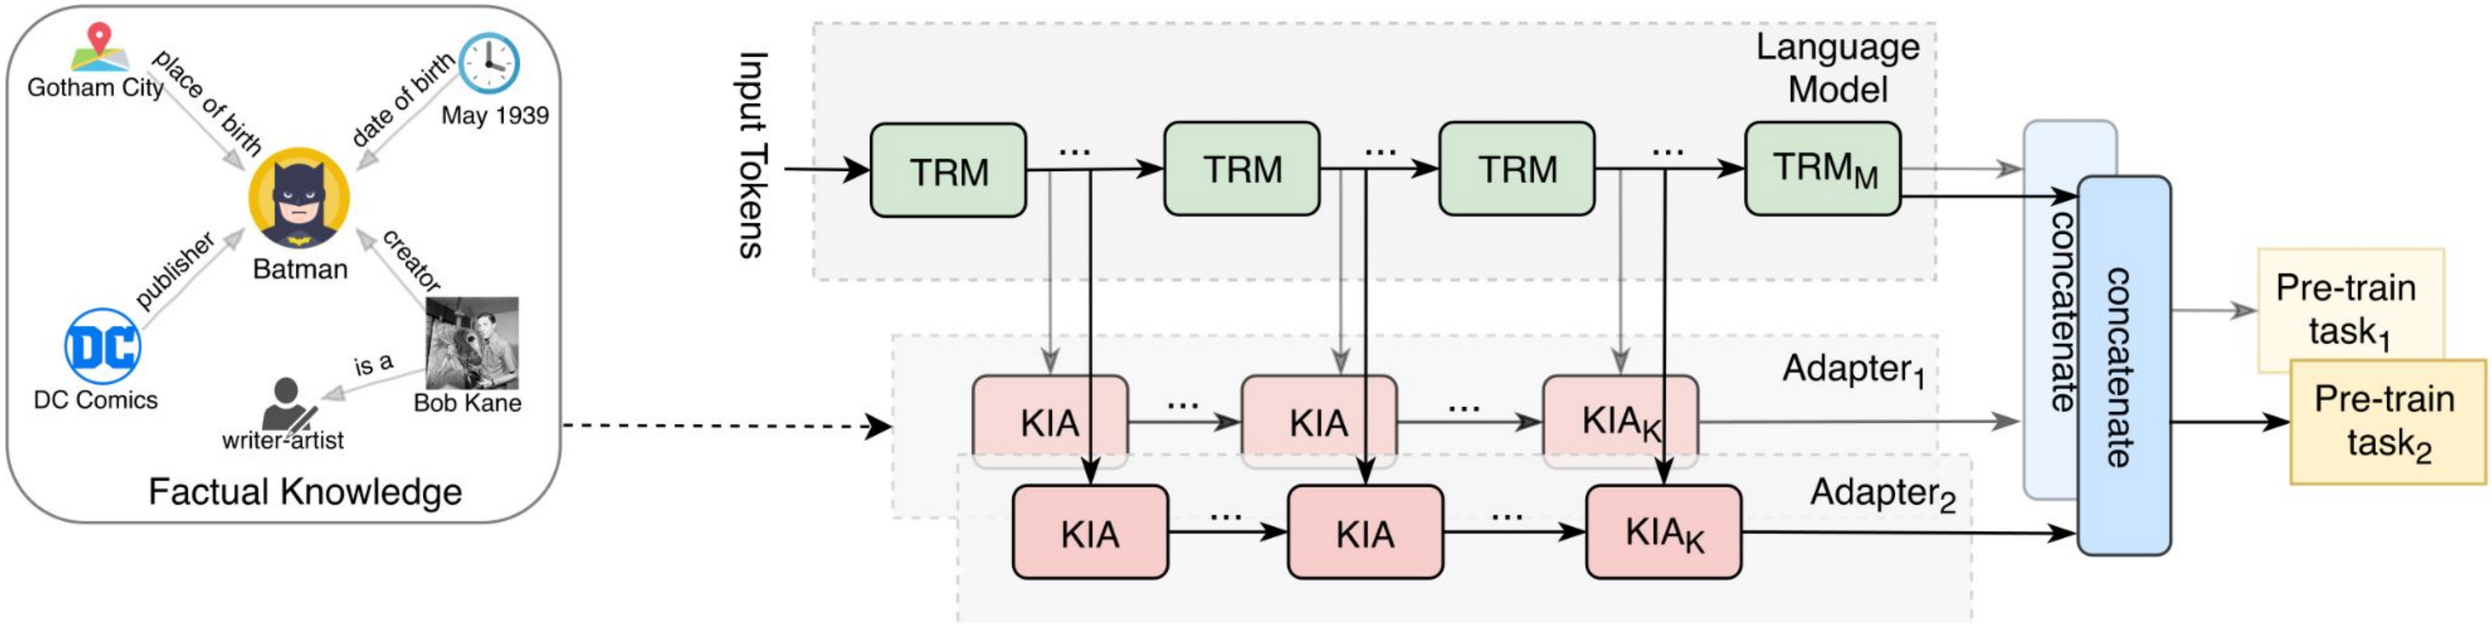

adapted from Wang *et al.* (2020)

Knowledge Adapters (K-Adapter)

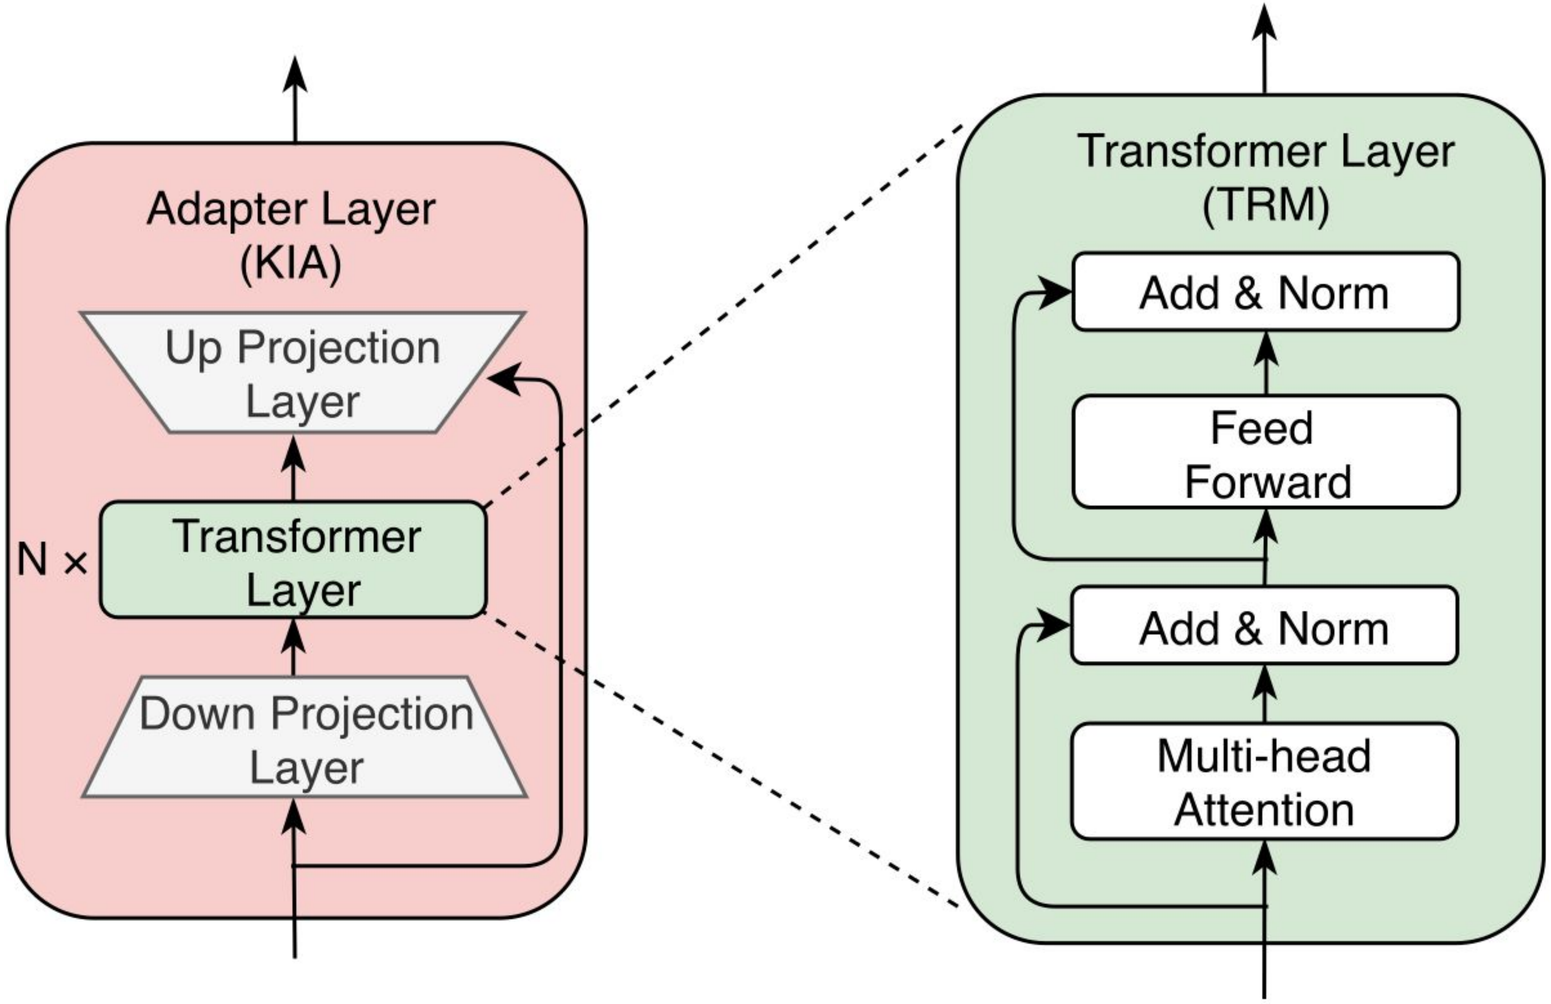

Wang *et al.* (2020)

Knowledge Adapters (K-Adapter)

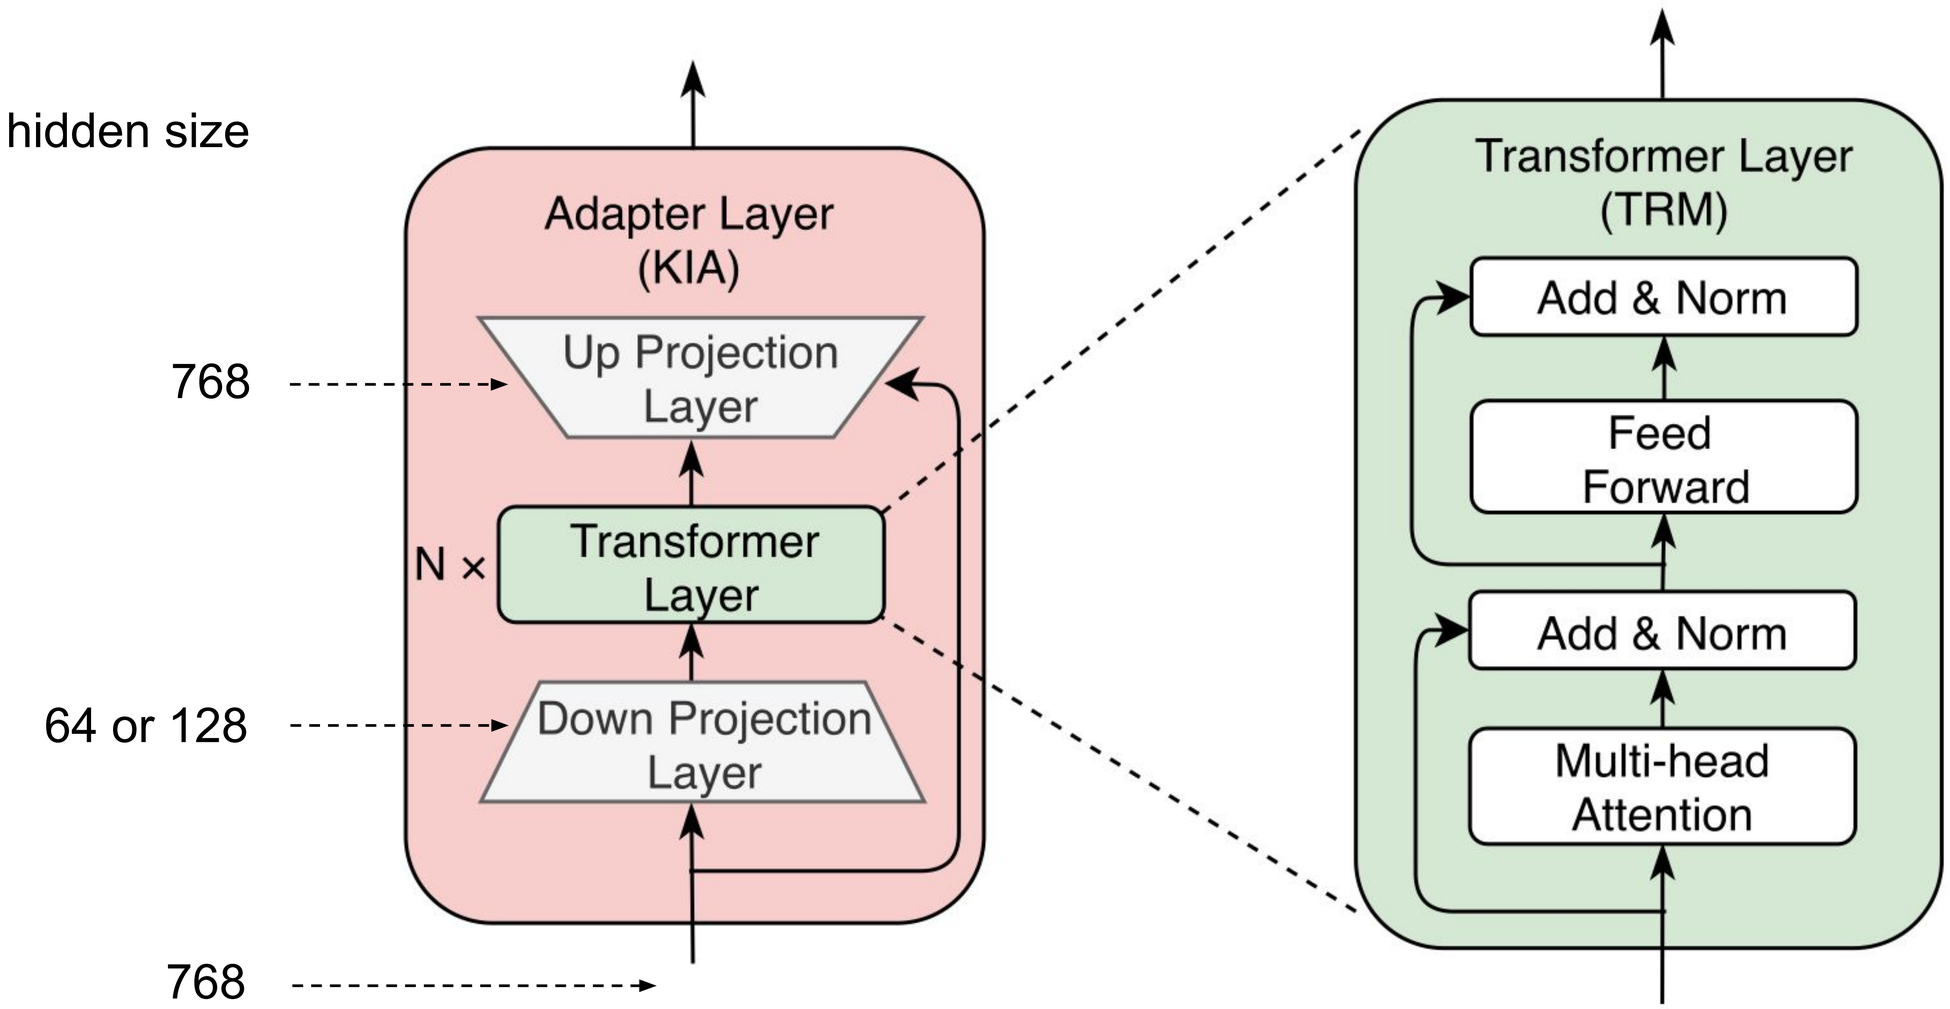

Wang *et al.* (2020)

K-Adapter - Pre-Training

Injected Knowledge:
“Barack Obama” “Birth-of-place” “Honolulu”

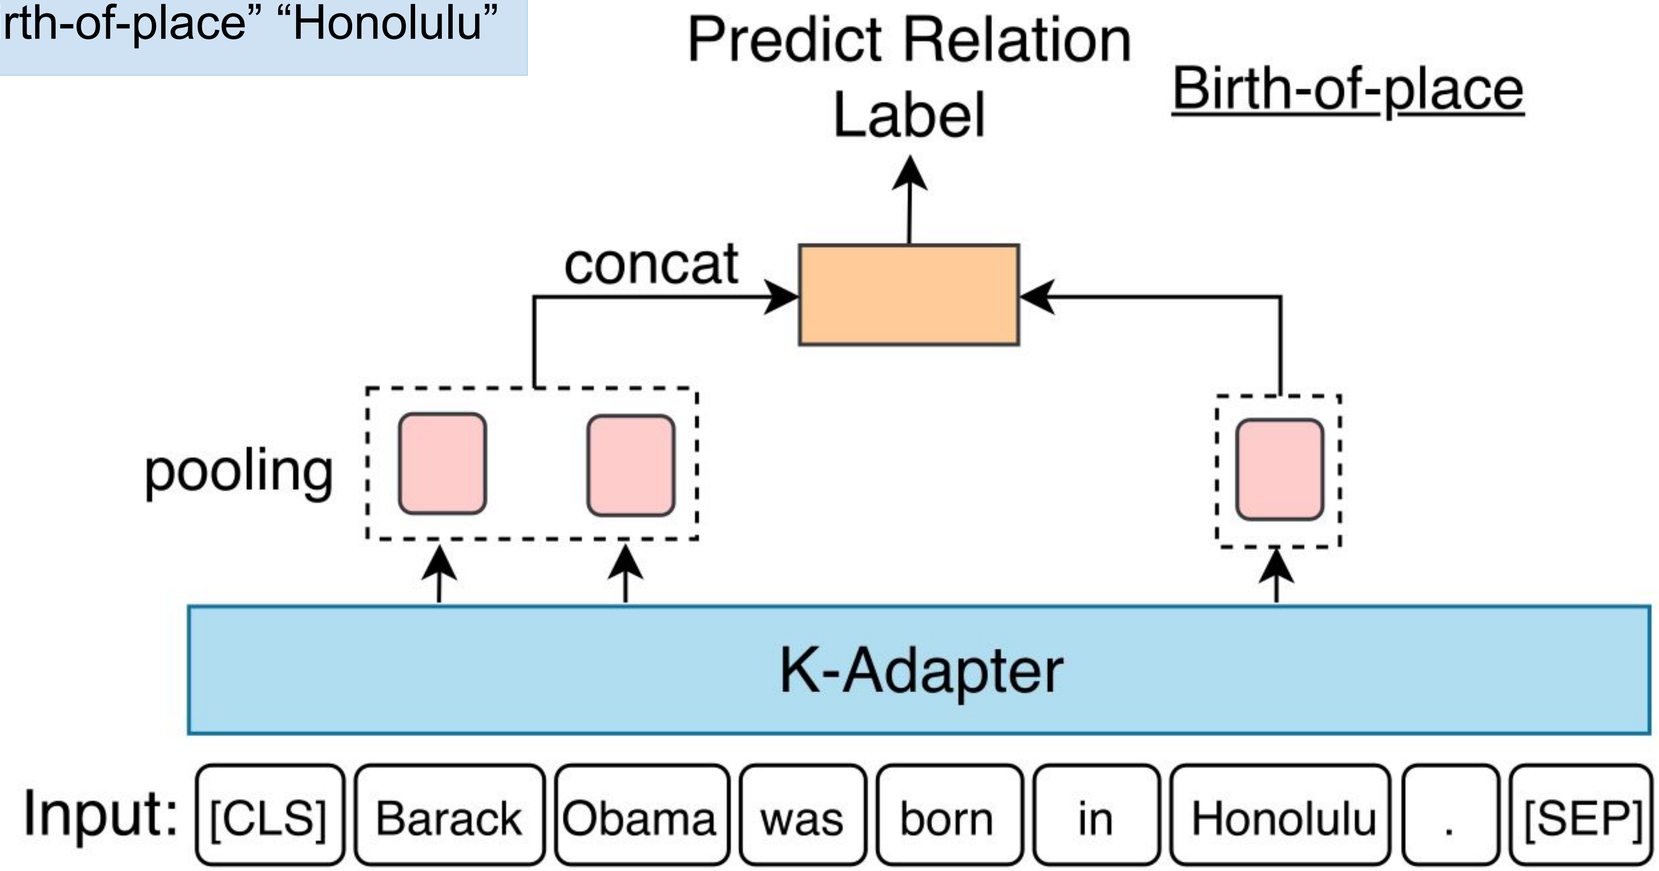

Wang *et al.* (2020)

Overview

	Lightweight Training	No Token-Limitation	Easy to Implement
Direct Text Injection (DTI)	✓✓	✗	✓✓
K-BERT	✓✓	✗	✗
Multilayer Perceptron (MLP)	✓	✓✓	✓
K-Adapter	✓	✓✓	✓

Benjamin Wolff

ZB MED Information Center for Life Sciences, Cologne

Research Group Data Science and Service

wolff@zbmed.de

linkedin.com/in/benjwolff

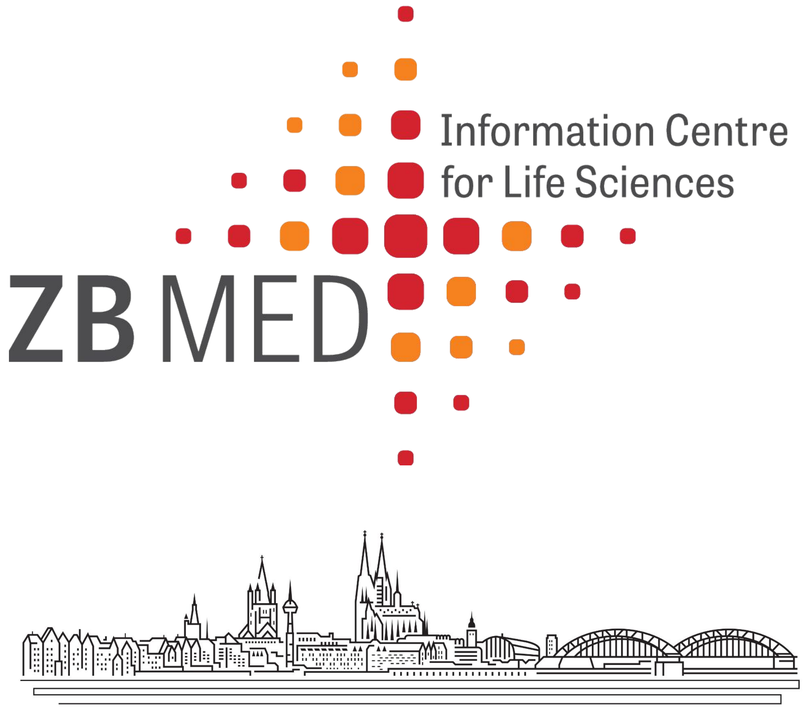

References

Cadeddu et al. "A comparative analysis of knowledge injection strategies for large language models in the scholarly domain." *Engineering Applications of Artificial Intelligence* 133 (2024): 108166.

Ostendorff et al. "Enriching bert with knowledge graph embeddings for document classification." *arXiv preprint arXiv:1909.08402* (2019).

Liu et al. "K-bert: Enabling language representation with knowledge graph." *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 34. No. 03. 2020.

Wang et al. "K-adapter: Infusing knowledge into pre-trained models with adapters." *arXiv preprint arXiv:2002.01808* (2020).

K-BERT - Seeing Layer

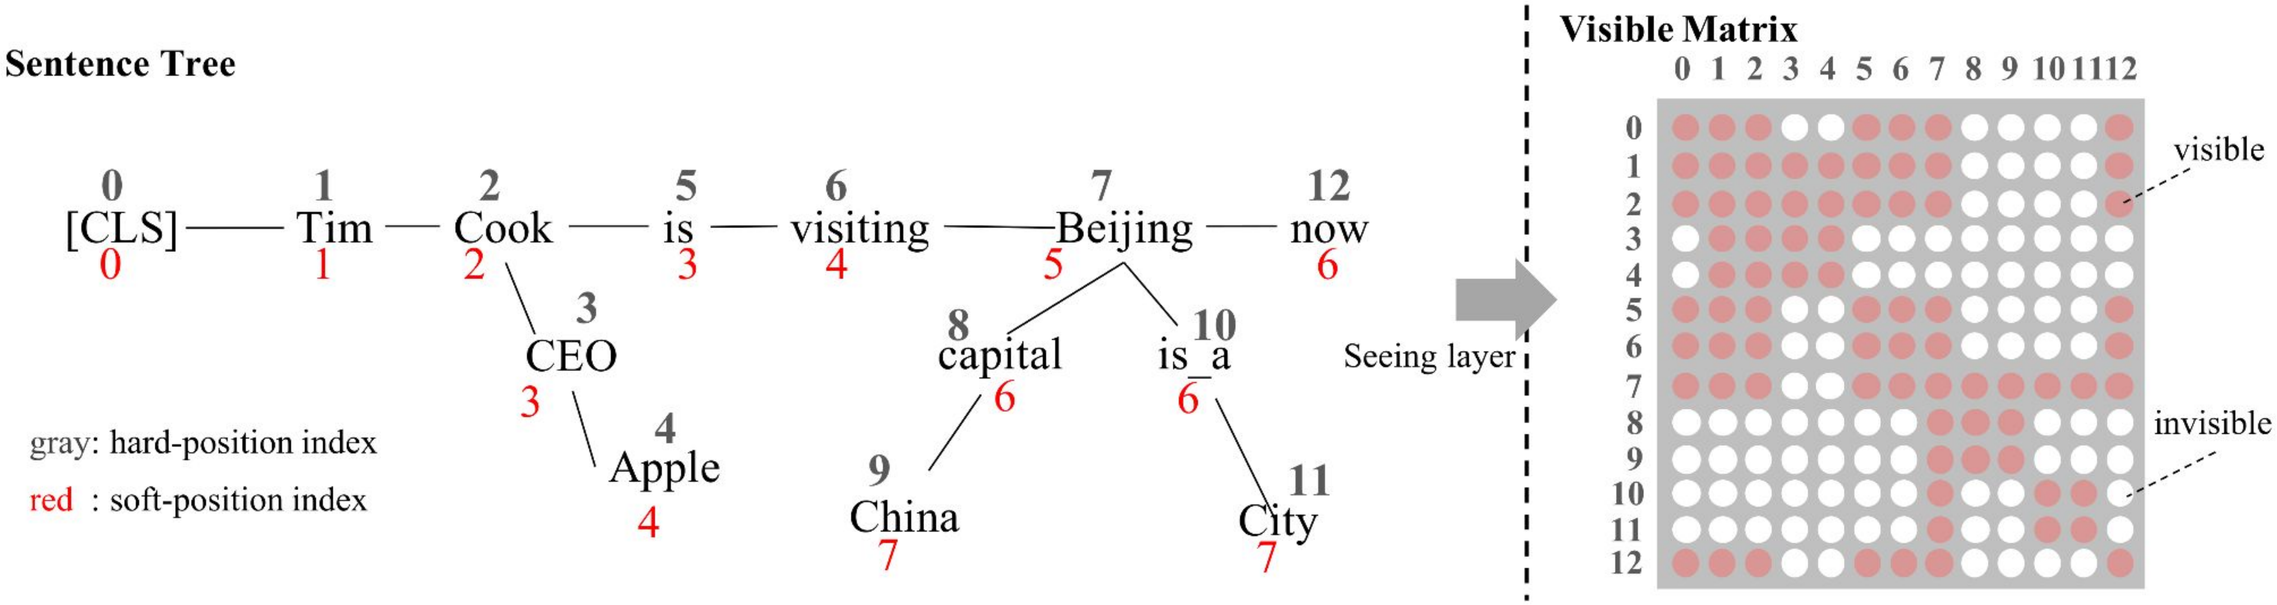

Liu et al (2019)