




NVIDIA

NLP/BERT in 10 mins

Timothy Liu, 24 March 2020

timothy1@nvidia.com

http://bit.ly/hpcai_nlp_intro



Outline

NLP/BERT in 10 mins

1. Introduction to Natural Language Processing (NLP)
2. Transformer Networks
3. Transfer Learning with BERT



Natural Language Processing (NLP)

Language Understanding for Computers

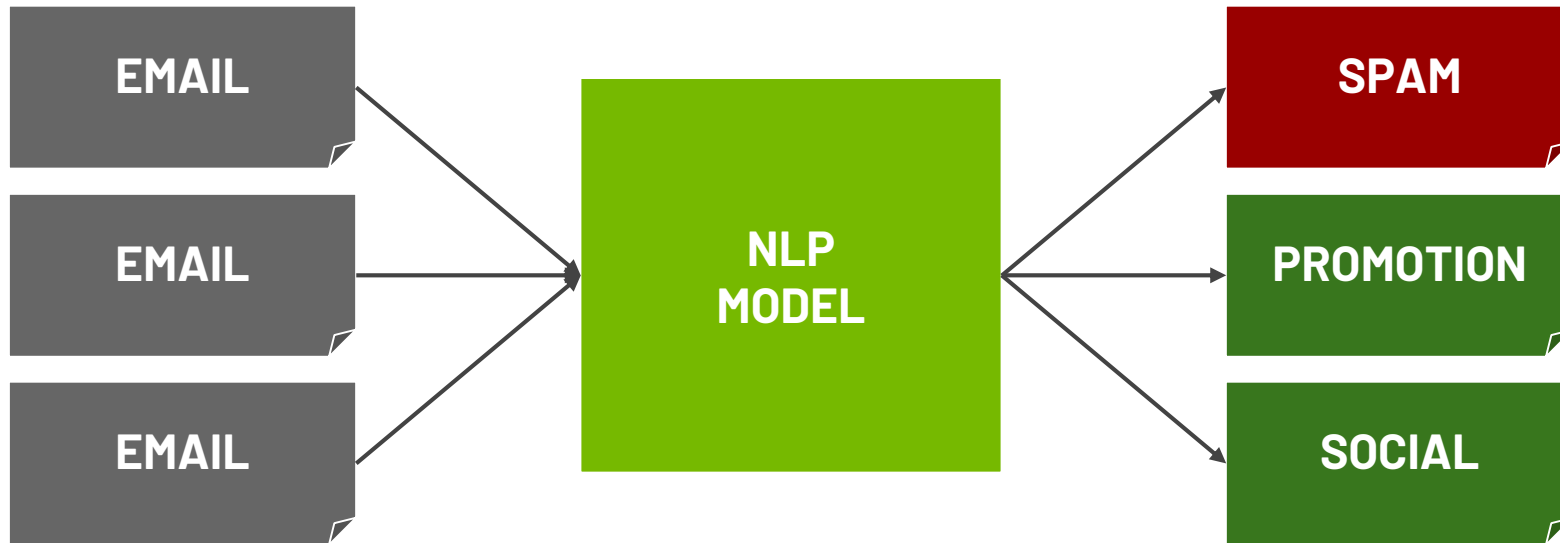
Natural Language Processing is a field of computer science that aims to:

1. Allow computers to understand and process large amounts of text
2. Enable closer, more natural interactions between human and computers

NLP also generates useful ideas for related fields that involve sequences (e.g. genomics)

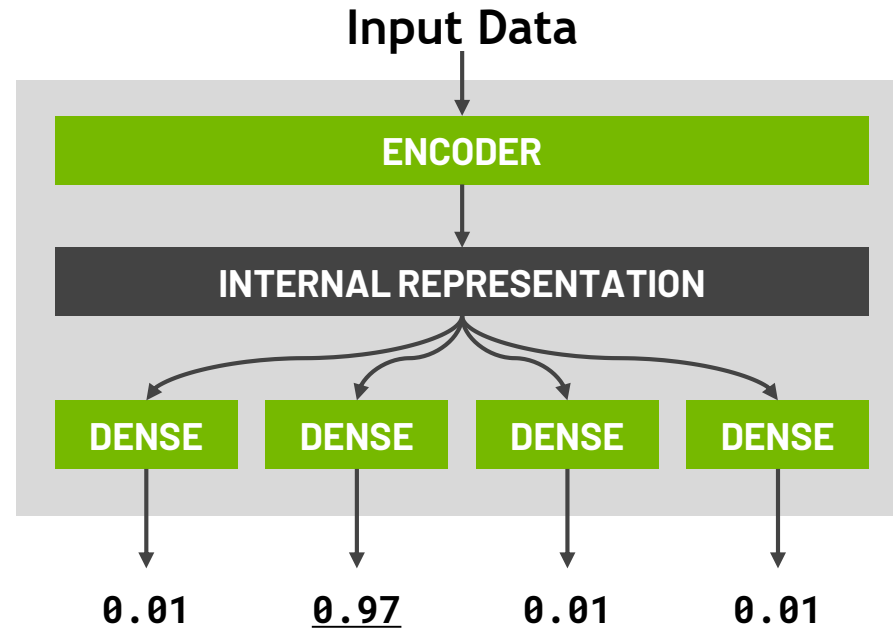
Natural Language Processing (NLP)

Example Task: Text Classification



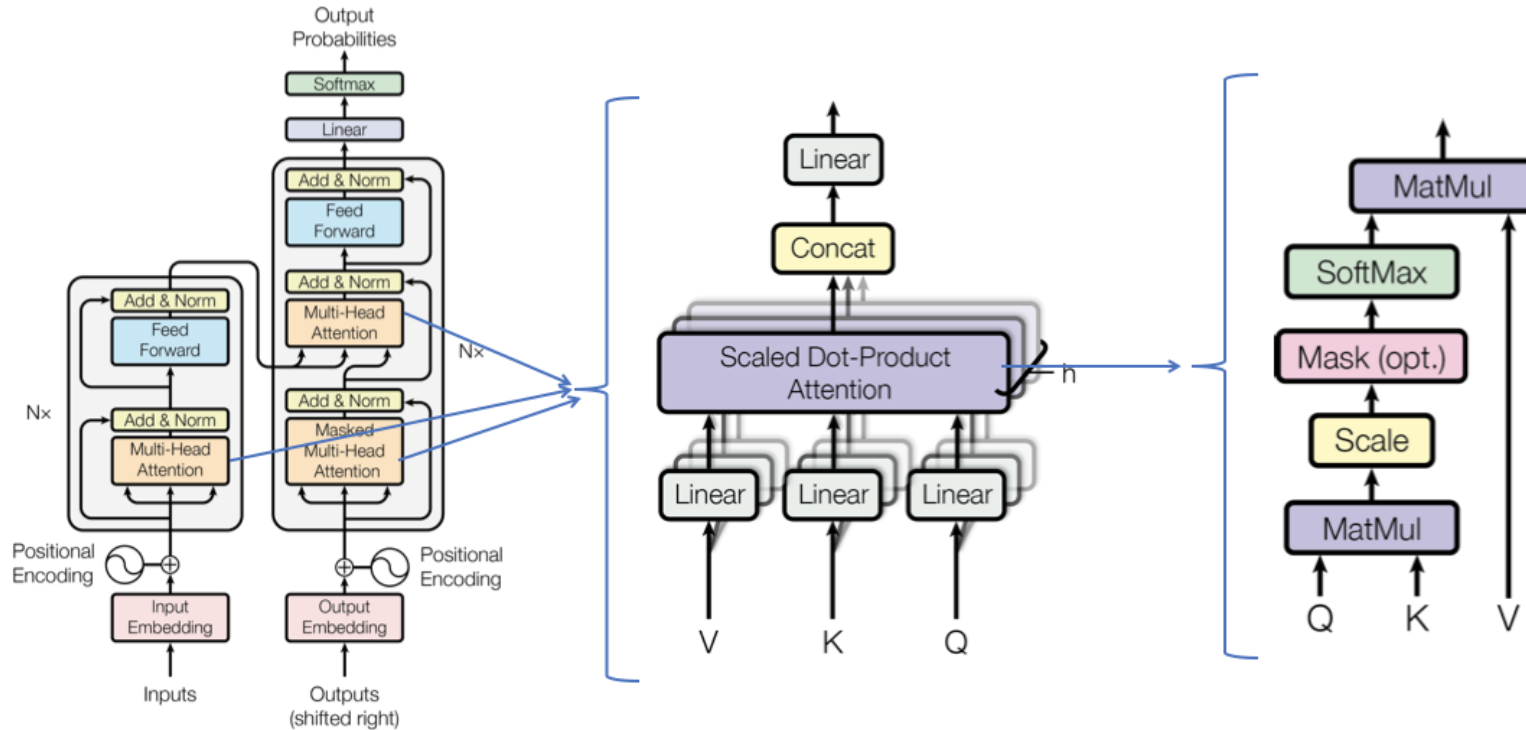
Natural Language Processing (NLP)

Using Deep Learning to Build Powerful Models



Transformer Networks

Powerful, non-recurrent DL model for sequence modelling



Attention Is All You Need (Vaswani et al., 2017)

Transformer Networks

Powerful, non-recurrent DL model for sequence modelling

Attention Output

Input Sequence



Calculate “attention score” from pairwise comparisons between tokens,
then do weighted combination of token representations

BERT

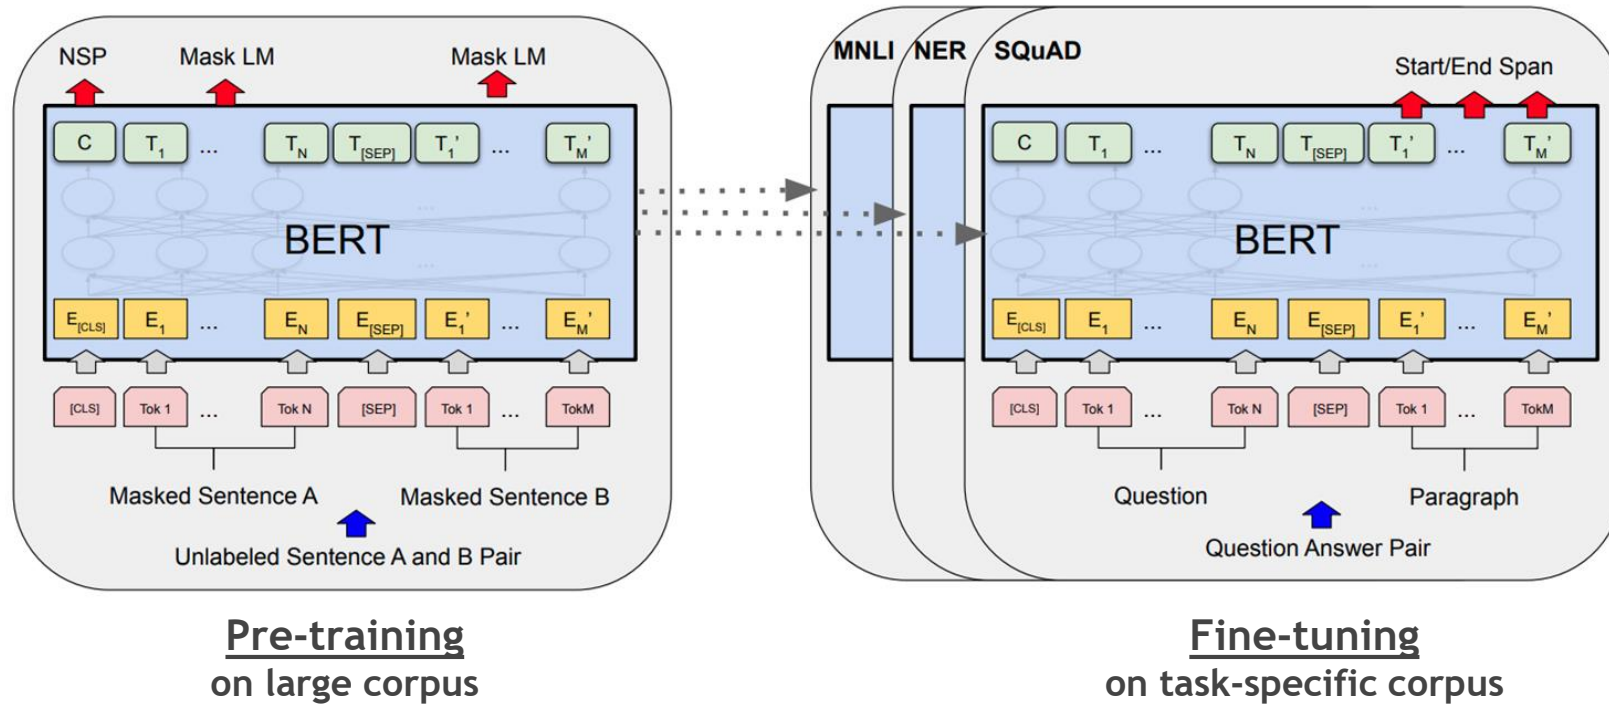
Bidirectional Encoder Representations from Transformers

BERT is a Transformer model that:

- Builds upon initial ideas from *Attention is All You Need* (Vaswani, 2017) paper
- Designed to learn powerful methods of encoding representations from text
- **Demonstrated state-of-the-art results on many NLP problems in many languages**

BERT

Bidirectional Encoder Representations from Transformers



BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding (Devlin et al., 2018)

BERT

Bidirectional Encoder Representations from Transformers

BERT-BASE

12-layer

768-hidden

12-heads

110M parameters

BERT-LARGE

24-layer

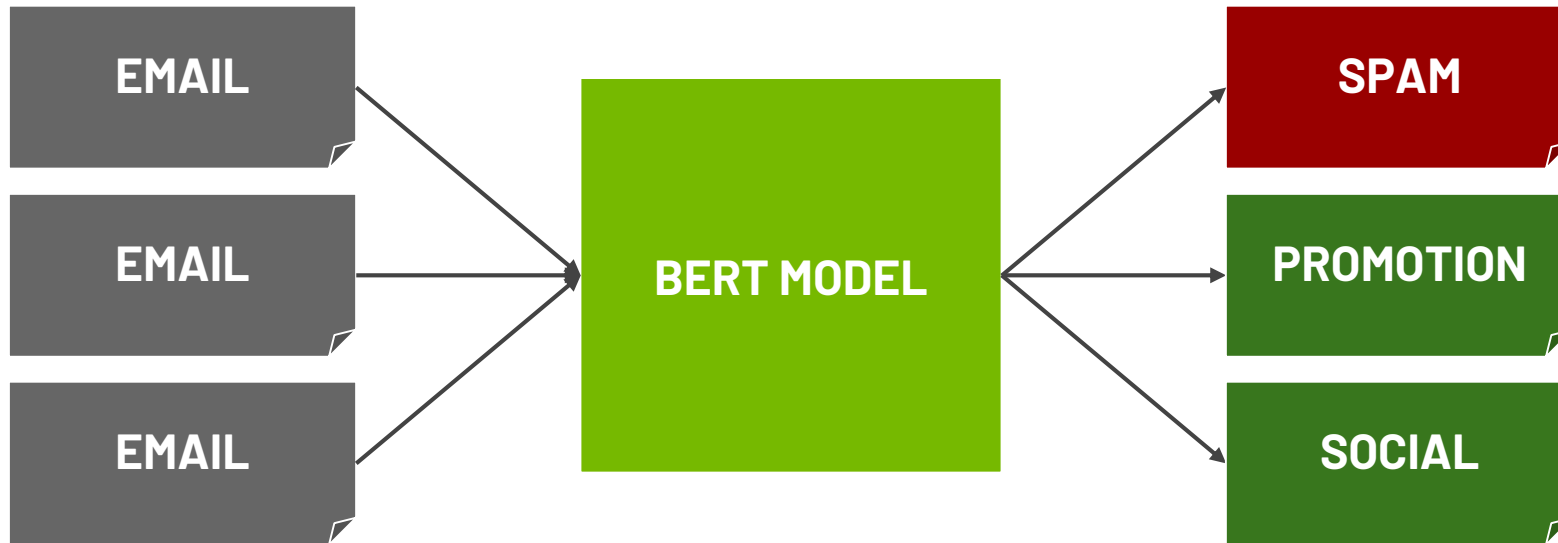
1024-hidden

16-heads

340M parameters

Natural Language Processing (NLP)

Example Task: Text Classification





NLP/BERT in 10 mins

Timothy Liu, 24 March 2020

timothy1@nvidia.com

http://bit.ly/hpcai_nlp_intro



NVIDIA