

07.06.22

Coding.Waterkant 22

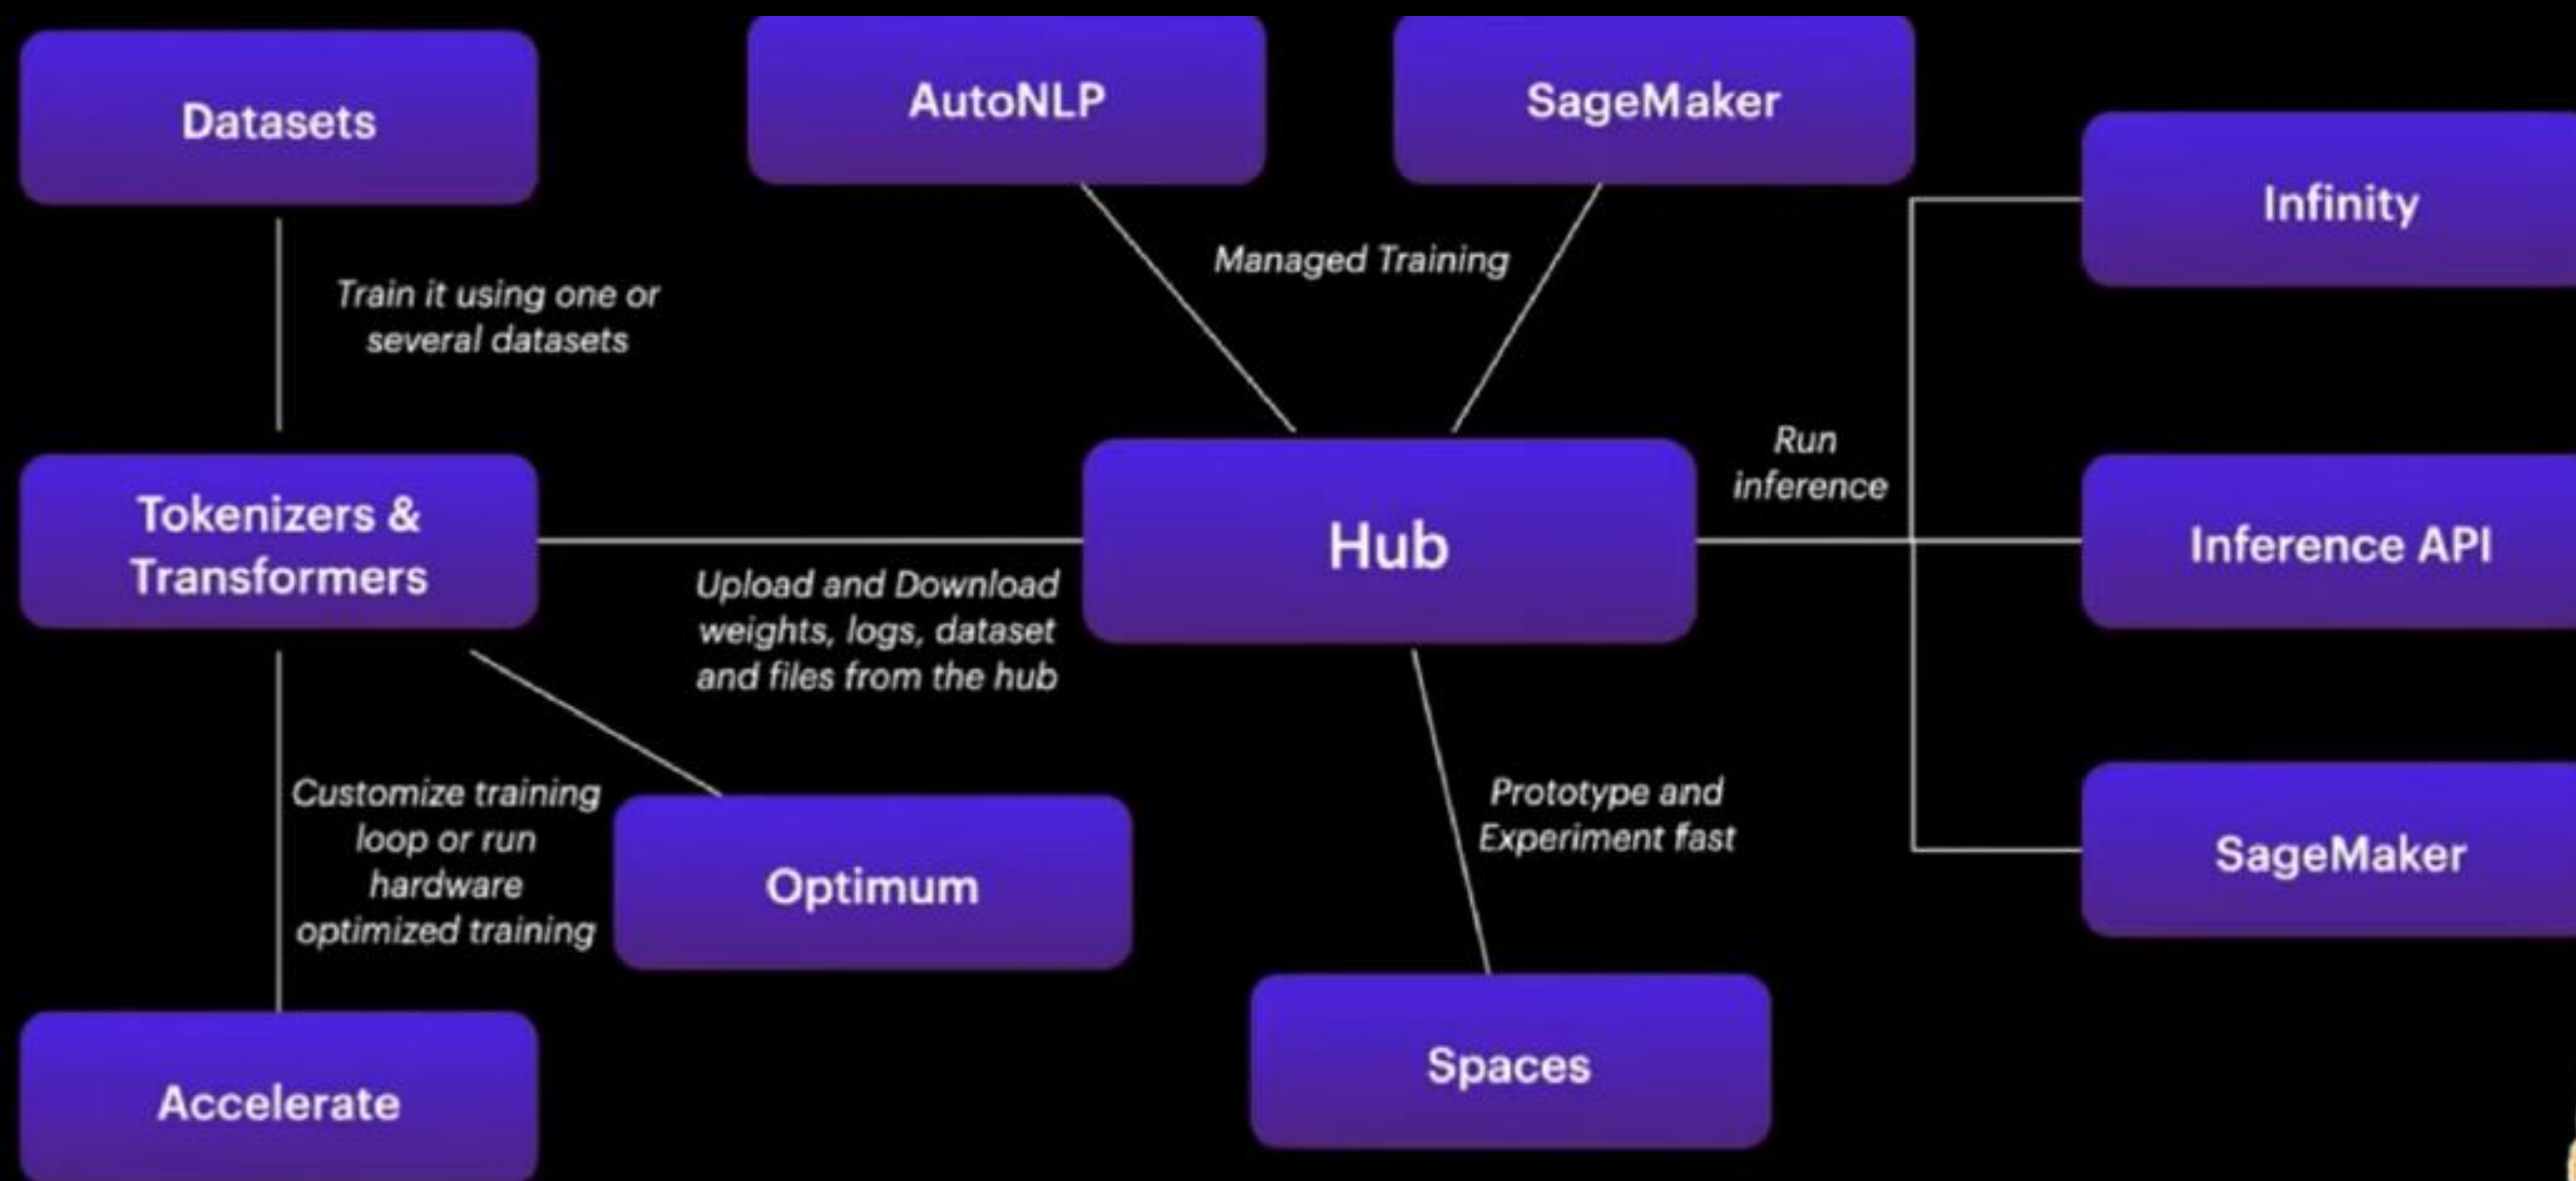
INTRODUCTION TO HUGGING FACE

- **The Hugging Face Hub**
- **The Benefits of Transformer Models**
- **Example Usages**
- **Learning Resources**

HUGGING FACE

- **Founded in 2016 in New York but Team is working completely remote**
 - **CEO in Office in Paris**
 - **CTO (Thomas Wolf) in Netherlands**
- **Mission: „Democratizing AI“**

HUGGING FACE ECO SYSTEM



DATASETS

The screenshot shows the Hugging Face Datasets page. At the top, there is a navigation bar with the Hugging Face logo, a search bar for models, datasets, and users, and links to Models, Datasets, Spaces, Docs, Solutions, and Pricing. The main content area is divided into a left sidebar with filters and a main grid of dataset cards.

Task Categories: text-classification, question-answering, text-generation, translation, fill-mask, token-classification, + 93 Task Categories

Tasks: language-modeling, named-entity-recognition, sentiment-classification, extractive-qa, multi-class-classification, masked-language-modeling, + 334

Languages: en, es, fr, de, ru, it, + 187

Multilinguality: monolingual, multilingual, translation, other-programming-languages, en, bg, + 103

Sizes: 10K<n<100K, 1K<n<10K, 100K<n<1M, 1M<n<10M, unknown, n<1K, + 36

Datasets (5,414): Search Datasets. Sort: Most Downloads

- super_glue**: Preview • Updated May 4 • ↓ 1.08M • ♥ 13
- glue**: Preview • Updated 19 days ago • ↓ 940k • ♥ 37
- newsgroup**: Preview • Updated May 4 • ↓ 277k • ♥ 1
- wikitext**: Preview • Updated 19 days ago • ↓ 230k • ♥ 14
- wino_bias**: Preview • Updated Apr 13 • ↓ 158k • ♥ 3
- imdb**: Preview • Updated May 5 • ↓ 137k • ♥ 11
- squad**: Preview • Updated 14 days ago • ↓ 123k • ♥ 26
- wmt16**: Preview • Updated May 4 • ↓ 106k • ♥ 3
- tweet_eval**: Preview • Updated 19 days ago • ↓ 106k • ♥ 15
- anli**: Preview • Updated May 4 • ↓ 96.7k • ♥ 5
- race**: Preview • Updated May 4 • ↓ 74k • ♥ 4
- adversarial_qa**: Preview • Updated 19 days ago • ↓ 65.8k • ♥ 7
- red_caps**: Preview • Updated 19 days ago • ↓ 57.9k • ♥ 9
- kilt_tasks**: Preview • Updated 19 days ago • ↓ 54.2k • ♥ 2
- trivia_qa**: Updated May 4 • ↓ 51.7k • ♥ 1
- ag_news**: Preview • Updated 26 days ago • ↓ 50.7k • ♥ 14

MODELS

The screenshot shows the Hugging Face website's 'Models' page. The top navigation bar includes the Hugging Face logo, a search bar for models, datasets, and users, and links to Models, Datasets, Spaces, Docs, Solutions, and Pricing. The main content area is divided into a left sidebar and a main list of models.

Tasks: Image Classification, Translation, Image Segmentation, Fill-Mask, Automatic Speech Recognition, Token Classification, Sentence Similarity, Audio Classification, Question Answering, Summarization, Zero-Shot Classification, + 16 Tasks

Libraries: PyTorch, TensorFlow, JAX, + 25

Datasets: common_voice, wikipedia, squad, glue, bookcorpus, c4, conll2003, emotion, + 1015

Languages: en, es, fr, de, zh, ja, ru, sv, + 178

Licenses: apache-2.0, mit, cc-by-4.0, + 36

Models List:

- distilgpt2**
Text Generation • Updated 6 days ago • ↓ 19.1M • ♥ 58
- bert-base-uncased**
Fill-Mask • Updated 1 day ago • ↓ 16.4M • ♥ 153
- gpt2**
Text Generation • Updated May 19, 2021 • ↓ 12.9M • ♥ 113
- distilbert-base-uncased-finetuned-sst-2-english**
Text Classification • Updated Mar 22 • ↓ 12.2M • ♥ 57
- roberta-base**
Fill-Mask • Updated Jul 6, 2021 • ↓ 9.41M • ♥ 38
- distilbert-base-uncased**
Fill-Mask • Updated 7 days ago • ↓ 7.29M • ♥ 63
- SEBIS/code_trans_t5_small_program_synthese_transfer_learning_finetune**
Summarization • Updated 5 days ago • ↓ 6.66M • ♥ 3
- bert-base-cased**
Fill-Mask • Updated Sep 6, 2021 • ↓ 5.19M • ♥ 22

AUTOTRAIN

The screenshot shows the Hugging Face AutoTrain landing page. At the top, there is a navigation bar with the Hugging Face logo, a search bar, and links for Models, Datasets, Spaces, Docs, Solutions, and Pricing. The main heading reads "auto TRAIN" with "TRAIN" in a red box, followed by the sub-heading "Create powerful AI models without code". Below this is a description: "A new way to automatically train, evaluate and deploy state-of-the-art Machine Learning models." A "Create new project" button is visible. To the right, a grid of model cards is shown, each with an ID, name, accuracy, and training status. A "Watch trailer" button is overlaid on the grid. Below the main heading, three feature boxes are displayed: "Machine Learning made simple", "Automatic training", and "Fast deployment". At the bottom, a list of tasks available in AutoTrain is shown, including Text Classification, Token Classification, Question Answering, Translation, Summarization, Text Scoring, and Tabular Data Classification/Scoring.

Hugging Face Search models, datasets, users... Models Datasets Spaces Docs Solutions Pricing

auto **TRAIN**

Create powerful AI models without code

A new way to automatically train, evaluate and deploy state-of-the-art Machine Learning models.

Create new project

Machine Learning made simple
Train custom machine learning models by simply uploading data.

Automatic training
AutoTrain will find the best models for your data automatically.

Fast deployment
Your models are available on the [Hugging Face Hub](#), and ready to serve.

Tasks available in AutoTrain: [Text Classification \(Binary, Multi-class and Multi-label\)](#) [Token Classification](#) [Question Answering \(extractive\)](#) [Translation](#) [Summarization](#) [Text Scoring](#) [Tabular Data Classification \(Binary, Multi-class, Multi-label\)](#) [Tabular Data Scoring](#)

INFERENCE API

The screenshot shows the Hugging Face interface for the model `google/vit-base-patch16-224`. The top navigation bar includes the Hugging Face logo, a search bar, and links for Models, Datasets, Spaces, Docs, Solutions, Pricing, and a user profile. The model name is displayed with a 'like' button showing 32 likes. Below the name are tags for Image Classification, PyTorch, TensorFlow, JAX, Transformers, imagenet, imagenet-21k, arxiv:2010.11929, arxiv:2006.03677, apache-2.0, vit, and vision. The main content area is split into two columns. The left column contains the model card, starting with the title 'Vision Transformer (base-sized model)' and a detailed description of the model's pre-training and fine-tuning. It also includes a disclaimer and a 'Model description' section. The right column features a 'Hosted inference API' section with a file upload area and a 'JSON Output' section. Below these are 'Spaces using google/vit-base-patch16-224', listing several user-created spaces.

Hugging Face Search models, datasets, users... Models Datasets Spaces Docs Solutions Pricing

google/ **vit-base-patch16-224** like 32

Image Classification PyTorch TensorFlow JAX Transformers imagenet imagenet-21k arxiv:2010.11929 arxiv:2006.03677 apache-2.0 vit vision

Model card Files and versions Train Deploy Use in Transformers

Vision Transformer (base-sized model)

Vision Transformer (ViT) model pre-trained on ImageNet-21k (14 million images, 21,843 classes) at resolution 224x224, and fine-tuned on ImageNet 2012 (1 million images, 1,000 classes) at resolution 224x224. It was introduced in the paper [An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale](#) by Dosovitskiy et al. and first released in [this repository](#). However, the weights were converted from the [timm repository](#) by Ross Wightman, who already converted the weights from JAX to PyTorch. Credits go to him.

Disclaimer: The team releasing ViT did not write a model card for this model so this model card has been written by the Hugging Face team.

Model description

The Vision Transformer (ViT) is a transformer encoder model (BERT-like) pretrained on a large collection of images in a supervised fashion, namely ImageNet-21k, at a resolution of 224x224 pixels. Next, the model was fine-tuned on ImageNet (also referred to as ILSVRC2012), a dataset comprising 1 million images and 1,000 classes, also at resolution 224x224.

Downloads last month **40,188**

Hosted inference API

Image Classification Examples

Drag image file here or click to browse from your device


This model is currently loaded and running on the inference API.


JSON Output Maximize



Spaces using google/vit-base-patch16-224

- team-indain-image-caption/Hindi-image-captioning
- abidlabs/vision-transformer st0bb3n/Cam2Speech
- mmeendez/cnn_transformer_explainability IPN/demo_
- webis-huggingface-workshop/omar_demo
- vebie91/spaces-image-classification-demo
- abidlabs/Comparing-Vision-Transformers Turtle/Test


HUGGING FACE SPACE

 **Hugging Face**

[Models](#) [Datasets](#) [Spaces](#) [Docs](#) [Solutions](#) [Pricing](#) 

Spaces:  dalle-mini/**dalle-mini**  like 163 Running

[App](#) [Files and versions](#) [Linked Models](#)



DALL·E mini

Generate images from text

What do you want to see?

DALL·E mini is an AI model that generates images from any prompt you give!

Created by Boris Dayma et al. 2021-2022
[GitHub](#) | [Project Report](#)

Made with Streamlit

PIPELINE FUNCTION

```
from transformers import pipeline

classifier = pipeline("sentiment-analysis")
classifier("I've been waiting for a HuggingFace course my whole life.")
```

```
[{'label': 'POSITIVE', 'score': 0.9598047137260437}]
```

TASK IDENTIFIER

`task (str) -`

The task defining which pipeline will be returned. Currently accepted tasks are:

- "feature-extraction": will return a `FeatureExtractionPipeline`
- "sentiment-analysis": will return a `TextClassificationPipeline`
- "ner": will return a `TokenClassificationPipeline`
- "question-answering": will return a `QuestionAnsweringPipeline`
- "fill-mask": will return a `FillMaskPipeline`
- "summarization": will return a `SummarizationPipeline`
- "translation_xx_to_yy": will return a `TranslationPipeline`
- "text-generation": will return a `TextGenerationPipeline`

PRACTICE PART

- **Play around with the pipeline function for a different task.**

```
from transformers import pipeline

classifier = pipeline("sentiment-analysis")
classifier("I've been waiting for a HuggingFace course my whole life.")
```

```
[{'label': 'POSITIVE', 'score': 0.9598047137260437}]
```

WHAT IS A TRANSFORMER?

Main ingredients



Attention
mechanisms



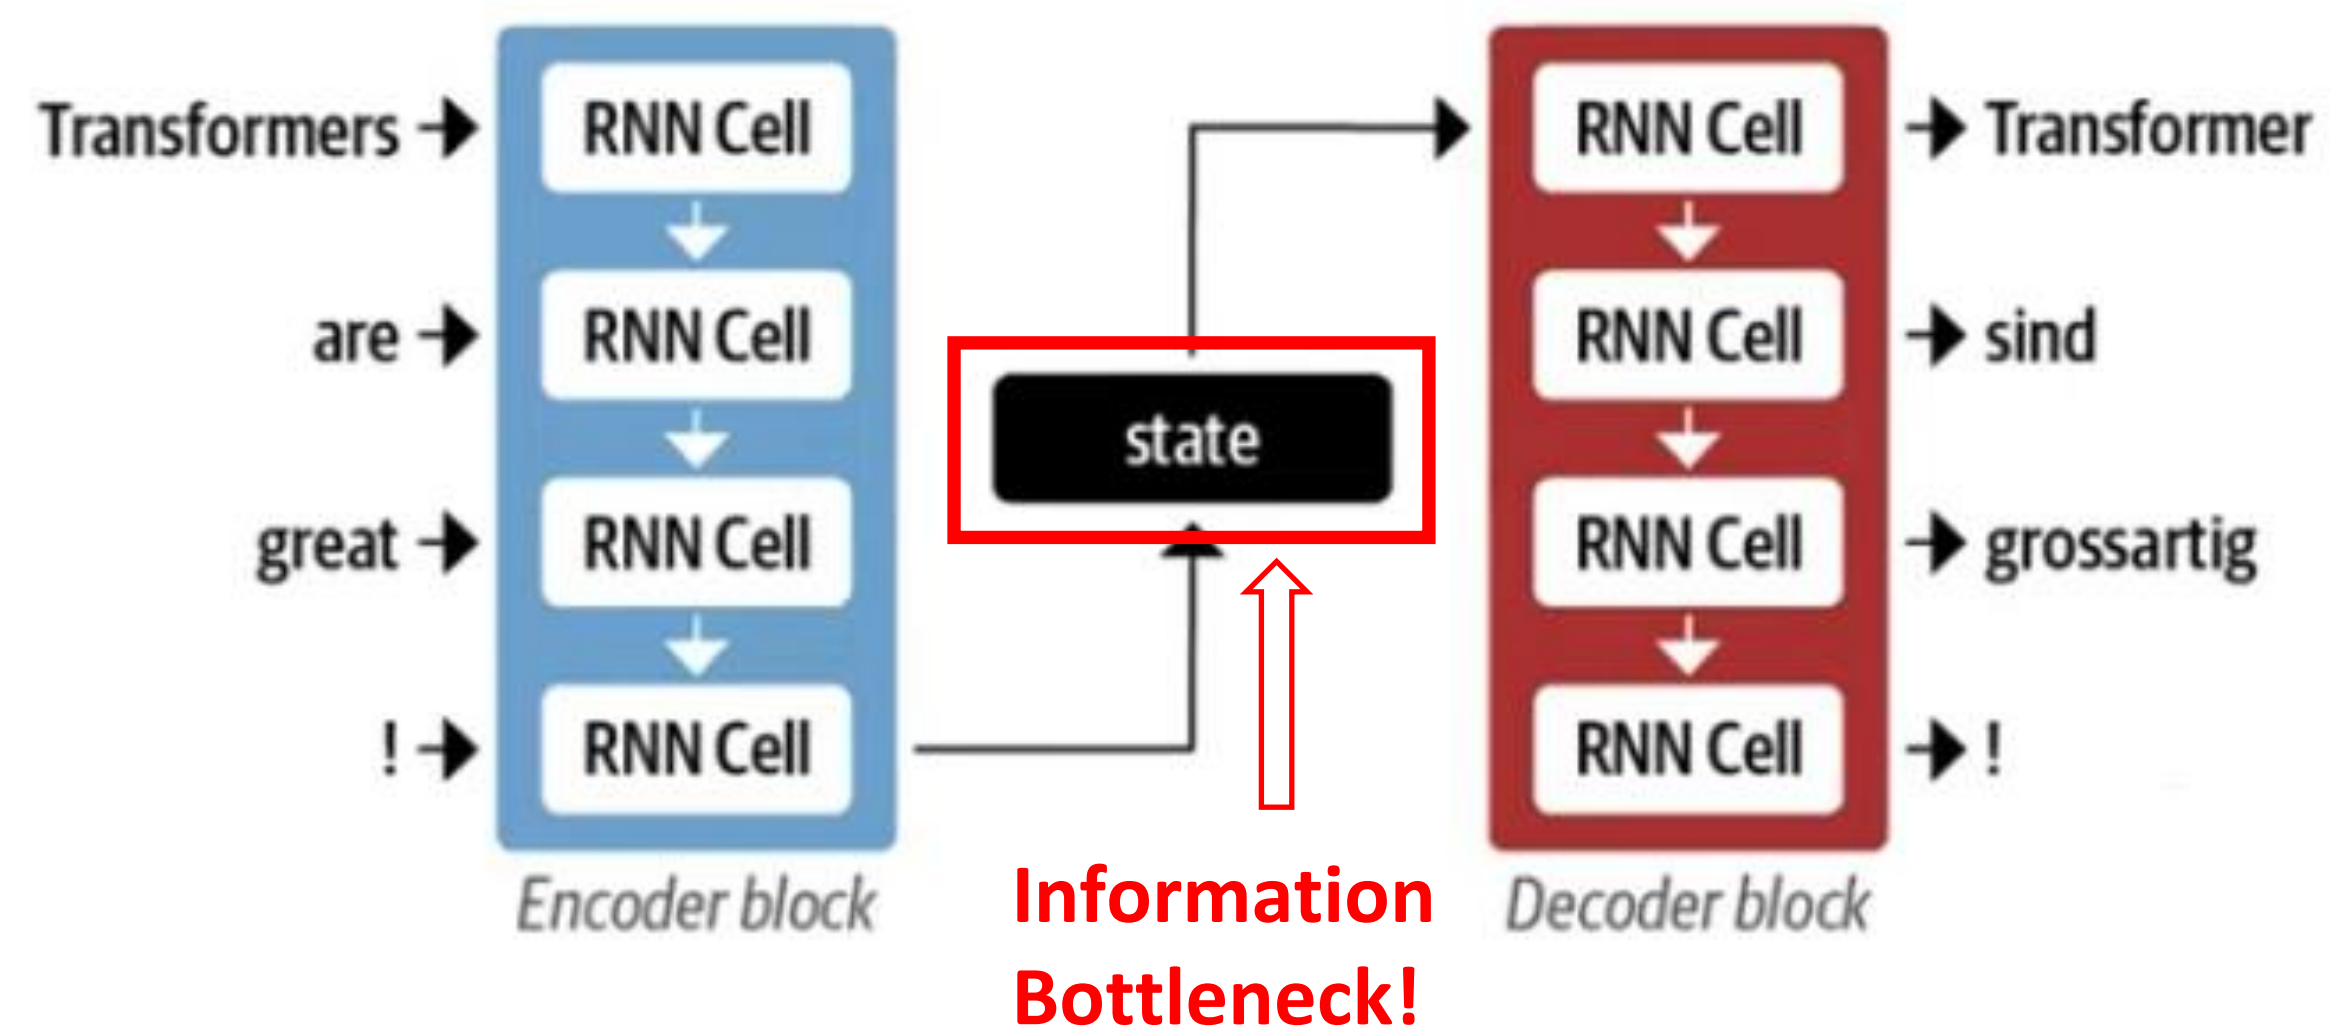
Self-supervised learning
(Pretraining)



Transfer learning
(Fine-tuning)



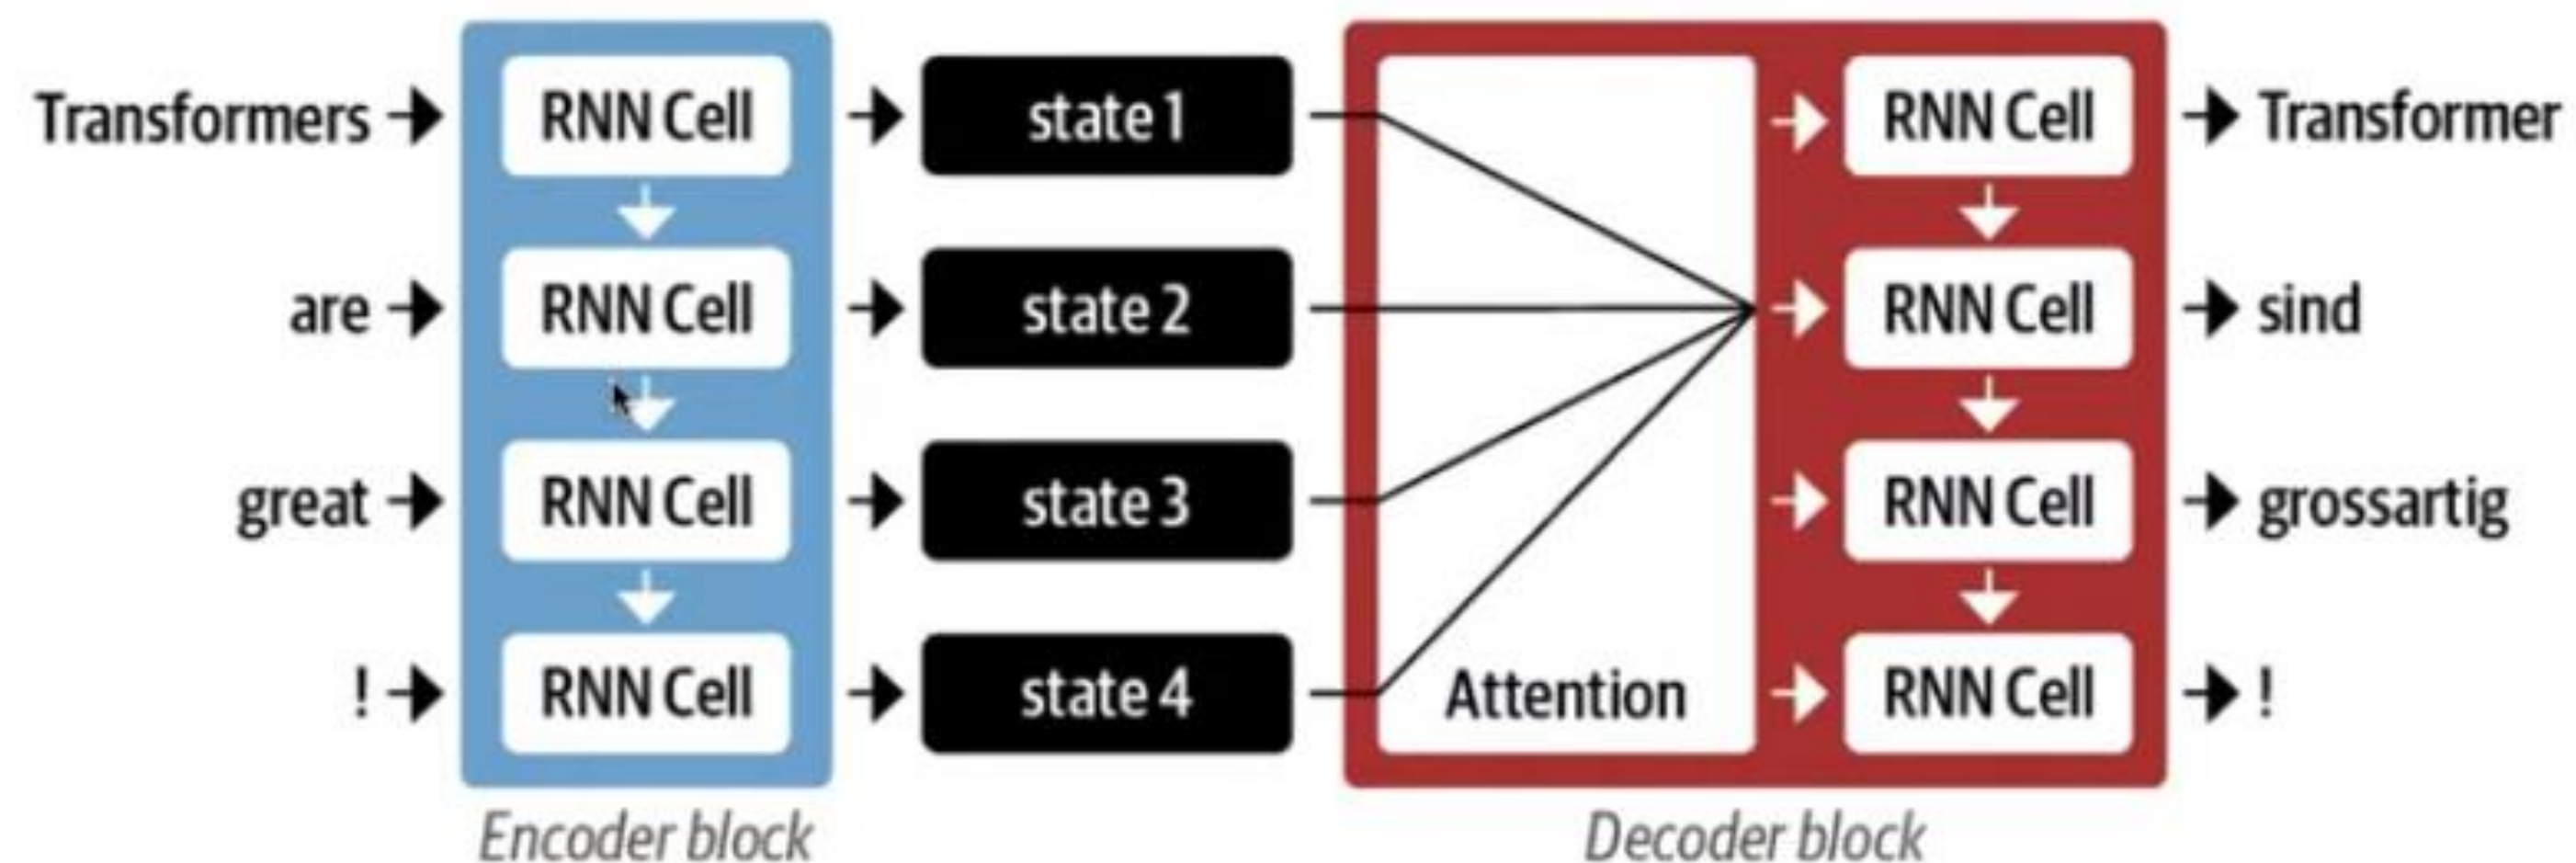
CLASSICAL SEQUENCE TO SEQUENCE APPROACH



Originally developed for recurrent neural networks




ATTENTION MECHANISM



Assign a weight or "pay attention" to specific states




INTERPRETATION



Attention gives better modelling of word order



ATTENTION IS ALL YOU NEED



Transformers much easier to scale with compute & data



BENEFITS

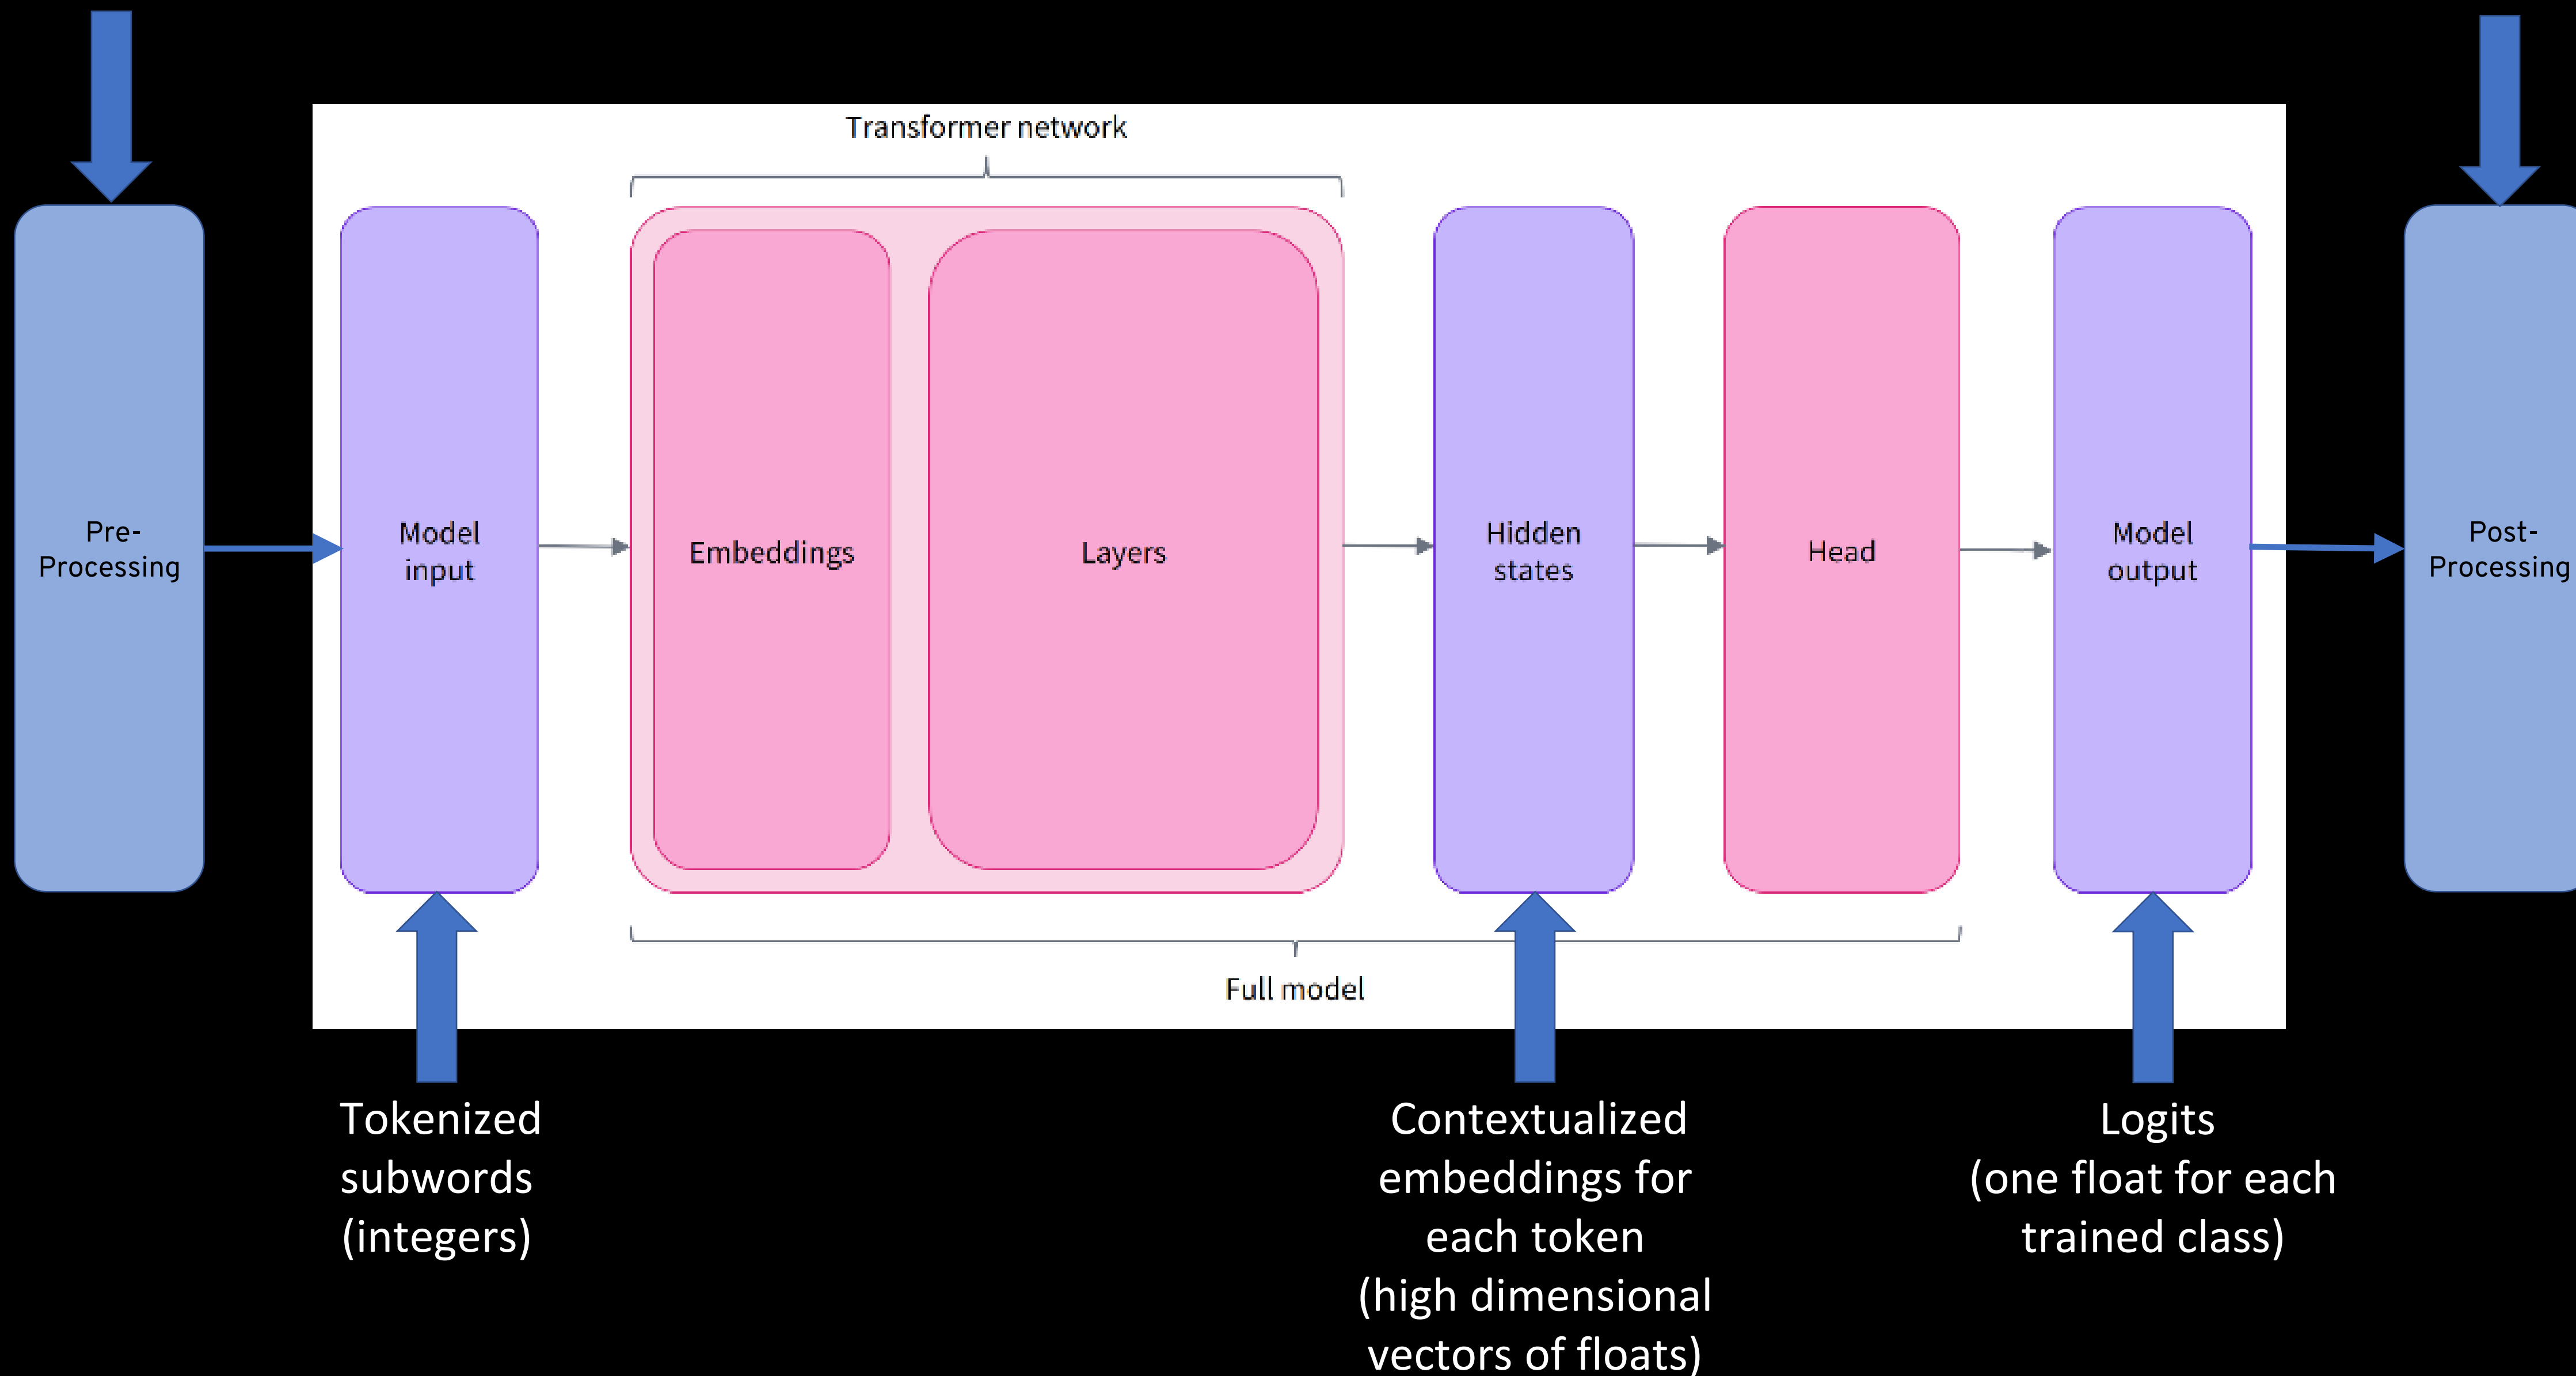
- **Solving the bottleneck problem in sequence-to-sequence tasks**
- **Provides some interpretability**
- **No Vanishing Gradient Problem**
- **Multimodality**


APPLICATION EXAMPLES

- **Code Generation: Co-Pilot in VS-Code from GitHub and OpenAI (GPT model)**
- **Search: Google (Bert model)**
- **Prediction of Protein folding: Alpha Fold from Deepmind**
- **Image Generation: DALL-E 2 by OpenAI**

- Splitting
- Mapping to integers
- Adding model dependent tokens/integers

- Logits to probs
- Probs to classes
- (Classes to tokens/text)





This course is amazing

[101, 2023, 2607, 2003, 6429, 999, 102]

[-4.3630, 4.6859]

POSITIVE: 99.89%
NEGATIVE: 0.,11%

TOKENIZER SELECTION

- **Splitting the input into words, subwords, or symbols (like punctuation) that are called *tokens***
- **Mapping each token to an integer**
- **Adding additional inputs that may be useful to the model**

```
from transformers import AutoTokenizer
```

```
checkpoint = "distilbert-base-uncased-finetuned-sst-2-english"
```

```
tokenizer = AutoTokenizer.from_pretrained(checkpoint)
```

TOKENIZER APPLICATION

```
raw_inputs = [  
    "I've been waiting for a HuggingFace course my whole life.",  
    "I hate this so much!",  
]  
inputs = tokenizer(raw_inputs, padding=True, truncation=True, return_tensors="pt")  
print(inputs)
```

MODEL SELECTION AND FITTING

```
from transformers import AutoModel

checkpoint = "distilbert-base-uncased-finetuned-sst-2-english"
model = AutoModel.from_pretrained(checkpoint)
```


```
outputs = model(**inputs)
print(outputs.last_hidden_state.shape)
```


MODEL SELECTION FOR A SPECIFIC DOWNSTREAM TASK

```
from transformers import AutoModelForSequenceClassification

checkpoint = "distilbert-base-uncased-finetuned-sst-2-english"
model = AutoModelForSequenceClassification.from_pretrained(checkpoint)
outputs = model(**inputs)
```

MODEL HEADS FOR DIFFERENT DOWNSTREAM TASKS

- ***Model (retrieve the hidden states)**
- ***ForCausalLM**
- ***ForMaskedLM**
- ***ForMultipleChoice**
- ***ForQuestionAnswering**
- ***ForSequenceClassification**
- ***ForTokenClassification**
- **and others** 

HUGGING FACE COURSE

The screenshot shows the Hugging Face website's course page. At the top, there's a navigation bar with the Hugging Face logo, a search bar for models, datasets, and users, and links to Models, Datasets, Spaces, Docs, Solutions, and Pricing. Below the navigation bar, the 'Course' section is highlighted, featuring a search bar for documentation and a 'MAIN' dropdown menu. The course content is organized into a list of sections, with '0. SETUP' being the current section. The main content area displays the 'Introduction' section, which includes a video player with a blue background and the text 'Welcome to the Hugging Face Course' and a smiling emoji. Below the video, a paragraph explains that the course covers natural language processing (NLP) using libraries from the Hugging Face ecosystem, including Transformers, Datasets, Tokenizers, and Accelerate, and is completely free and without ads. On the right side, there's a sidebar with a table of contents for the 'Introduction' section, listing 'Welcome to the Course!', 'What to expect?', and 'Who are we?'.

Hugging Face Search models, datasets, users... Models Datasets Spaces Docs Solutions Pricing

Course Search documentation Ctrl+K MAIN EN 497

- 0. SETUP
- 1. TRANSFORMER MODELS
- 2. USING TRANSFORMERS
- 3. FINE-TUNING A PRETRAINED MODEL
- 4. SHARING MODELS AND TOKENIZERS
- 5. THE DATASETS LIBRARY
- 6. THE TOKENIZERS LIBRARY
- 7. MAIN NLP TASKS
- 8. HOW TO ASK FOR HELP
- 9. BUILDING AND SHARING DEMOS **NEW**
- HUGGING FACE COURSE EVENT

Introduction

Welcome to the Course!

Welcome to the Hugging Face course

Welcome to the Hugging Face Course

This course will teach you about natural language processing (NLP) using libraries from the [Hugging Face ecosystem](#) — [Transformers](#), [Datasets](#), [Tokenizers](#), and [Accelerate](#) — as well as the [Hugging Face Hub](#). It's completely free and without ads.

Introduction

- Welcome to the Course!
- What to expect?
- Who are we?

0. SETUP

1. TRANSFORMER MODELS

2. USING 🤗 TRANSFORMERS

3. FINE-TUNING A PRETRAINED MODEL

4. SHARING MODELS AND TOKENIZERS

5. THE 🤗 DATASETS LIBRARY

6. THE 🤗 TOKENIZERS LIBRARY

7. MAIN NLP TASKS

8. HOW TO ASK FOR HELP

9. BUILDING AND SHARING DEMOS **NEW**

NLP WITH TRANSFORMERS

