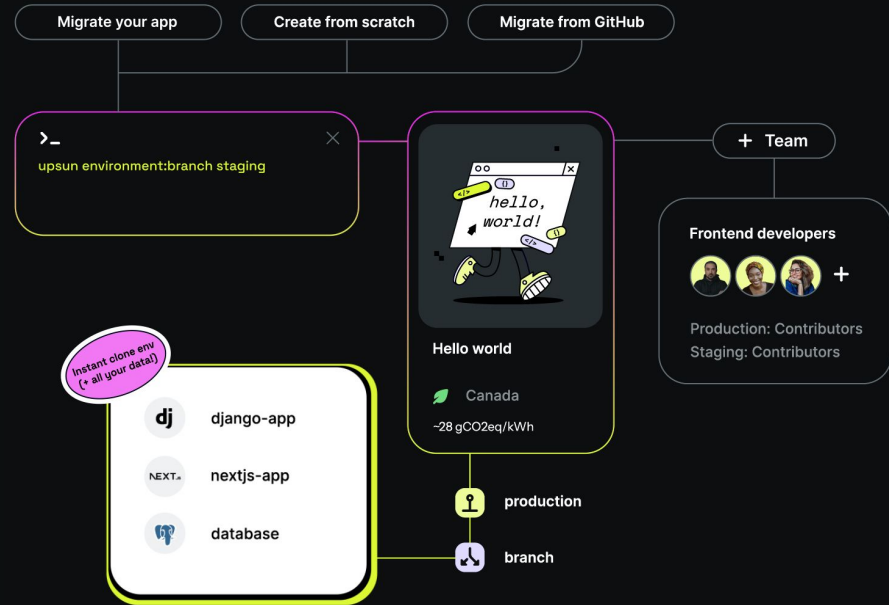




# Machine Learning Inference in PHP by example

Leverage ONNX and Transformers on Symfony





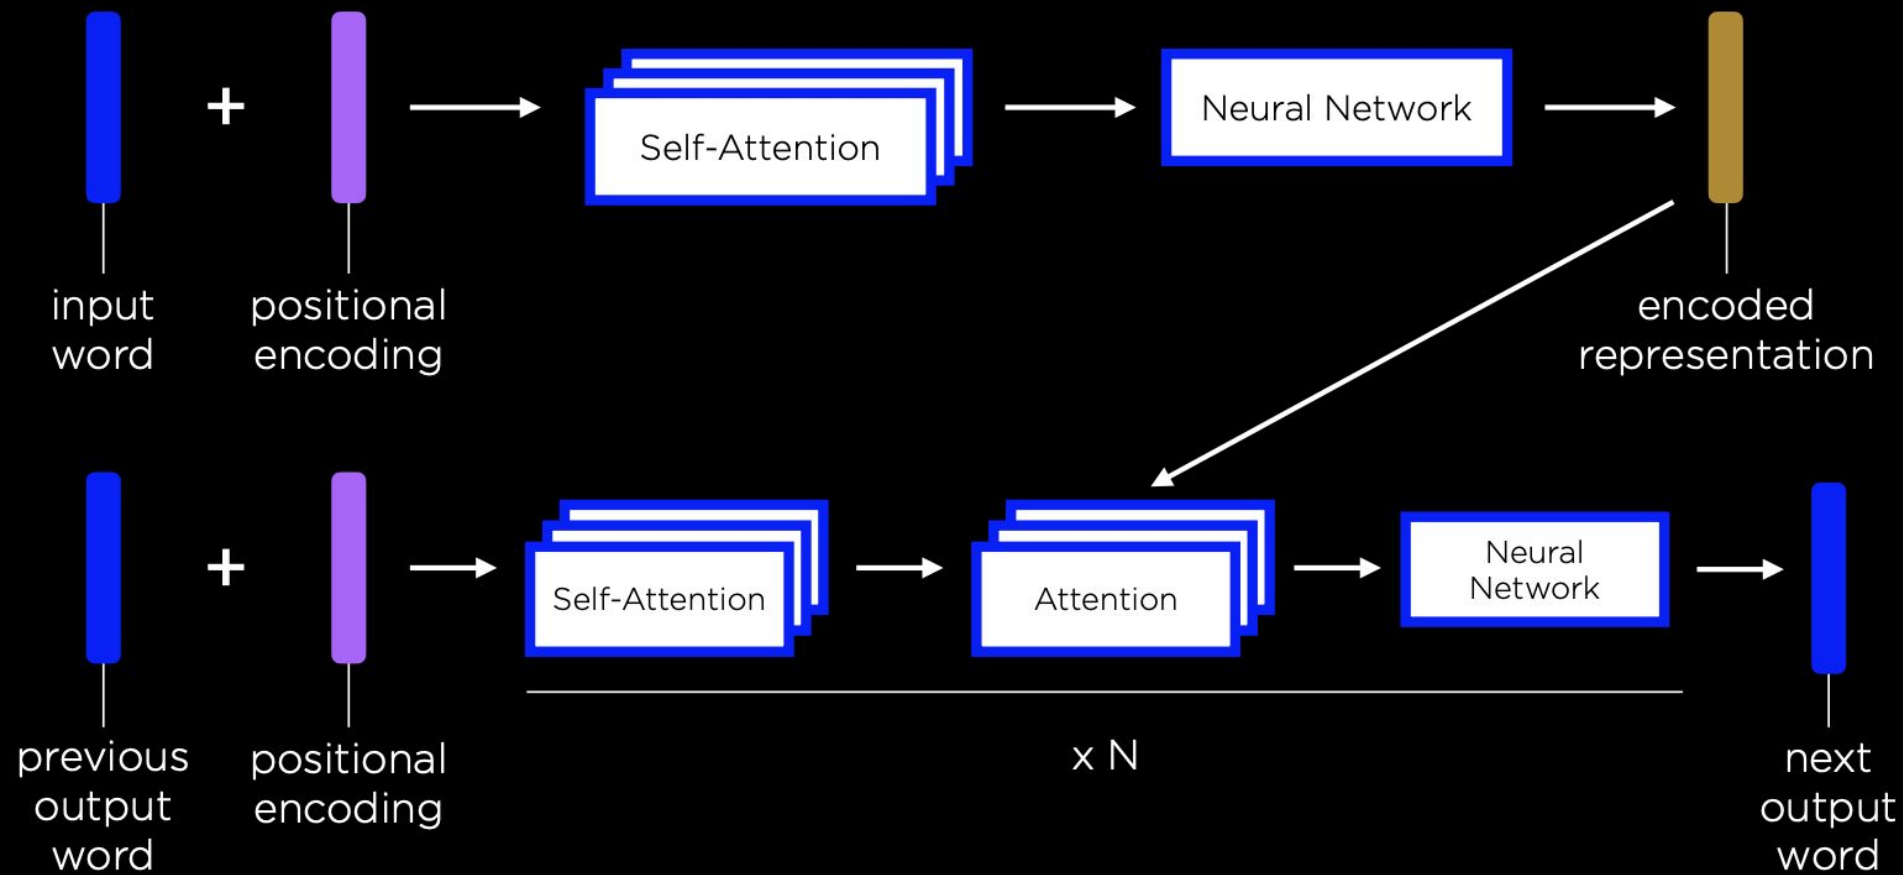
**Who is familiar with  
Transformers?**

# What is a Transformer?

- A neural network architecture that revolutionized NLP and beyond
- Introduced in 2017 paper "Attention is All You Need"
- Powers models like GPT, BERT, and T5

**Key innovation:**  
**processes all input  
at once instead of  
sequentially**

- Encoder-decoder structure
- Self-attention layers
- Feed-forward neural networks
- Position embeddings



# Before

# 2017

# After

Sequential processing (RNNs,  
LSTMs)  
Limited context window  
Slow training and inference

Parallel processing & lowering cost  
of training & inference  
Extended context understanding  
Faster training and inference  
Better at capturing relationships  
Reduced resources usage

## Demo:

Text: Symfony is a great framework for developing [MASK] applications.

## Learn more:

<https://cdn.cs50.net/ai/2023/x/lectures/6/lecture6.pdf>



SCAN ME

**Let's implement 3 use-cases  
in PHP**





```
composer require codewithkyrian/transformers # Package
```

```
./vendor/bin/transformers install # Install platform specific transformers
```

```
# Download models
```

```
./vendor/bin/transformers download Xenova/distilbert-base-uncased-finetuned-sst-2-english # Text  
classification
```

```
./vendor/bin/transformers download Xenova/vit-base-patch16-224 # Image classification
```

# Text Classification

★ Ce zăd Redets astyion  
Tar colarsig al neyions.

★★★★★ ant losery  
Secific

★ Cant Far est Re vste, Reviow  
P n of Tor dr rreions.



Amazon Reviert. ★★★  
Amazon Feoseting od nesions.

★★★★★





```
use function Codewithkyrian\Transformers\Pipelines\pipeline;

protected function execute(InputInterface $input, OutputInterface $output): int
{
    $asin = $input->getArgument('asin');
    $output->writeln("Calculating score for ASIN: $asin");

    $product = $this->entityManager->getRepository(Product::class)->findOneBy(['asin' => $asin]);
    $reviews = $product->getReviews();

    $pipe = pipeline('sentiment-analysis');

    foreach ($reviews as $review) {
        $out = $pipe($review->getText());
        if ($out['label'] == 'POSITIVE') {
            $positive++;
        } else {
            $negative++;
        }
    }

    $output->writeln("Positive: $positive, Negative: $negative, Score: " . ($positive - $negative) . "
    (".round($positive / $reviews->count() * 100, 2)."% positive)");
    return Command::SUCCESS;
}
```

# Demo

```
php bin/console app:score B07VGRJDFY
```



## Image Classification



```
public function handle(Request $request): Response
{
    [...]
    // Process the image to generate labels
    Transformers::setup()->setImageDriver(ImageDriver::GD);
    $classifier = pipeline('image-classification');
    $result = $classifier($this->getParameter('kernel.project_dir').'/public/uploads/'.$newFilename, 3);
    // Loop results to see if it's a hot dog
    [...]
}
```

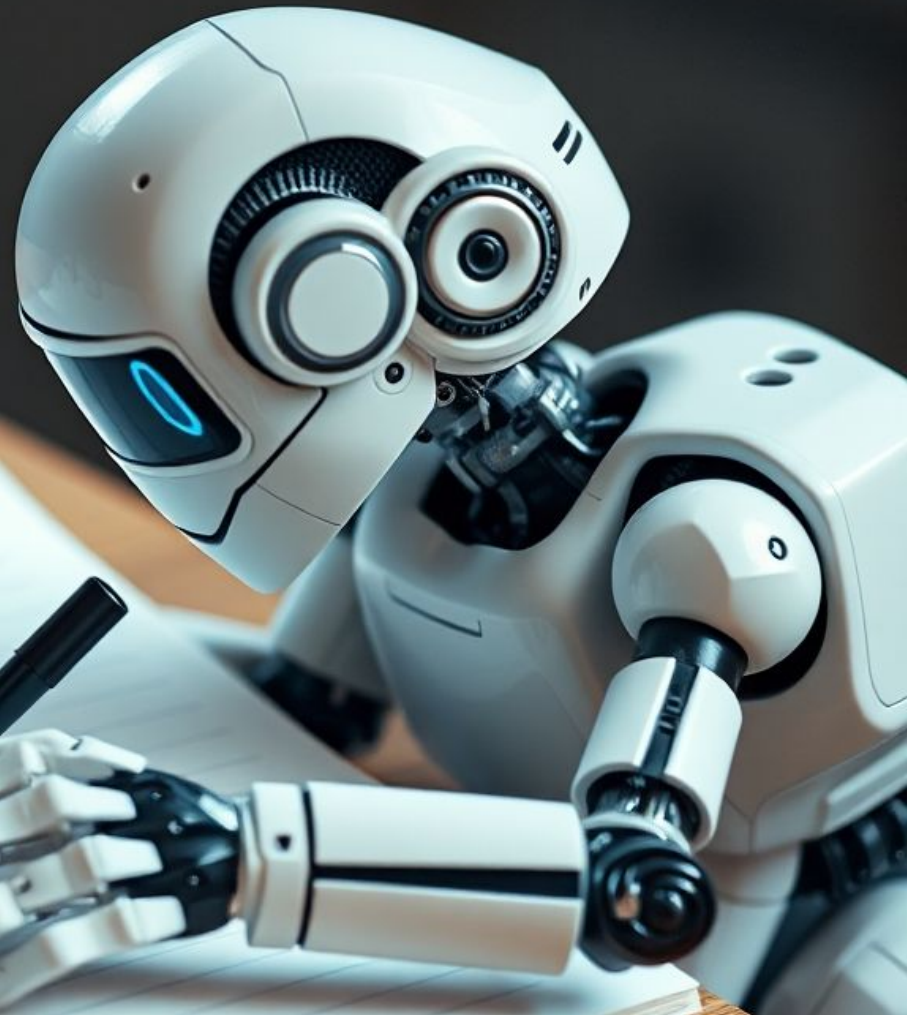
# Demo

<https://vienna.moigneu.net/hotdog>



**LLM  
generation**

Yeout  
in o +  
t:  
text







```
public function handle(Request $request): Response
{
    $question = $request->request->get('question');

    $generator = pipeline('text2text-generation', 'Xenova/flan-t5-small');
    $result = $generator($question,
        maxNewTokens: 256,
        repetitionPenalty: 1.6,
        temperature: 0.7
    );
    $answer = $result[0]['generated_text'];

    return $this->render('question/handle.html.twig', [
        'question' => $question,
        'answer' => $answer
    ]);
}
```

# Demo

<https://vienna.moigneu.net/question>

A small, white, humanoid robot with a large, circular lens for an eye and a glowing blue eye. The robot is holding a sparkler that is emitting bright sparks. The robot is standing next to an open book with text on the pages. The background is dark and out of focus.

Time to explain the magic.

# PHP FFI

Foreign Function Interface

FFI extension allows PHP code to directly **call functions and manipulate data from C libraries** without writing additional C code or PHP extensions.

# Open Neural Network Exchange

ONNX for short

ONNX (Open Neural Network Exchange) is an **open standard** format that allows AI models to be shared between different machine learning frameworks like PyTorch, TensorFlow, and many others.

# Transformers

and pipelines

Provided by the TransformersPHP package.

Courtesy of Kyrian Obikwelu

And 7 contributors

[Supported pipelines](#)

The screenshot shows the GitHub repository page for 'transformers-php' by CodeWithKyrian. The repository is public and has 1 branch (main) and 15 tags. The commit history table is as follows:

Commit	Message	Time
takielias	Fix Class Name for PretrainedTokenizer (#71)	7a4975b - last month
	feat: Actions workflow to build and upload assets on relea...	3 months ago
	Reintroduce install command, new Library checker, and u...	6 months ago
	Fix Class Name for PretrainedTokenizer (#71)	last month
	Fix Class Name for PretrainedTokenizer (#71)	last month
	fix: refactor LibChecker to use the new libsloader workflow	3 months ago
	feat: Update conversion notebook to include task	2 months ago
	Fix Class Name for PretrainedTokenizer (#71)	last month
	Fix Class Name for PretrainedTokenizer (#71)	last month
	fix: Stop tracking unnecessary example files	3 months ago
	Update CHANGELOG	2 months ago
	Add Tensor buffer test and update license	6 months ago
	Correct docs on supported pipelines as mentioned in #27	7 months ago
	feat: Update VERSION	2 months ago
	revert: Return rokka/vips to normal dependencies	3 months ago
	Initial Commit	10 months ago

# Models

Infinite possibilities

- Hundreds of models readily available on 🤗.
- Convert any model to ONNX with:
  - + [Python](#)
  - + [Notebook](#)

The screenshot shows the Hugging Face website interface. At the top, there is a search bar with the text "Search models, datasets, users...". Below the search bar, there are navigation tabs for "Tasks", "Libraries", "Datasets", "Languages", "Licenses", and "Other". The "Tasks" tab is selected, and a search filter "Filter Tasks by name" is visible. The main content area is divided into two columns. The left column lists various task categories: "Multimodal" (Image-Text-to-Text, Visual Question Answering, Document Question Answering, Video-Text-to-Text, Any-to-Any), "Computer Vision" (Depth Estimation, Image Classification, Object Detection, Image Segmentation, Text-to-Image, Image-to-Text, Image-to-Image, Image-to-Video, Unconditional Image Generation, Video Classification, Text-to-Video, Zero-Shot Image Classification, Mask Generation, Zero-Shot Object Detection, Text-to-3D, Image-to-3D, Image Feature Extraction, Keypoint Detection), and "Natural Language Processing" (Text Classification, Token Classification, Table Question Answering, Question Answering, Zero-Shot Classification, Translation). The right column displays a list of models filtered by "onnx". The models listed are: "ezioruan/inswapper\_128.onnx" (Updated Jul 4, 2023), "SamLowe/roberta-base-go\_emotions-onnx" (Text Classification, Updated Sep 28, 2023), "microsoft/Phi-3-vision-128k-instruct-onnx-directml" (Text Generation, Updated 6 days ago), "onnx-community/Florence-2-base-ft" (Image-Text-to-Text, Updated Oct 8), "onnx-community/Qwen2.5-1.5B-Instruct" (Text Generation, Updated Oct 8), "akhaliq/AnimeGANv2-ONNX" (Updated Nov 11, 2021), "nouamanetazi/bloom-350m-onnx" (Updated Jul 1, 2022), "NeuML/ljspeech-jets-onnx" (Text-to-Speech, Updated Oct 10), "deepghs/ml-danbooru-onnx" (Updated Apr 10, 2023), and "deepghs/ccip\_onnx" (Zero-Shot Image Classification, Updated Sep 9).



# SmolLM

Promising efficient LLM model

- Provided by the HF team
- 538Mb only
- Great alternative to fully fledged models
- Can even run on the browser

Spaces HuggingFaceTB / SmolLM-360M-Insturct-WebGPU like 117 Running App

Give me some tips to improve my public speaking skills

Here are some tips to help you improve your public speaking skills:

1. **Practice, practice, practice:** The more you practice, the more comfortable you'll become with your material. Practice in front of a mirror, record yourself, or speak in front of a small group of friends or family members.
2. **Know your audience:** Understand who your audience is, what they're interested in, and what they're looking for. Tailor your message, language, and tone to engage and persuade them.
3. **Prepare thoroughly:** Research your topic, organize your thoughts, and create an outline. Make sure you have a clear structure and a solid understanding of your key points.
4. **Focus on your message:** Instead of worrying about how you're coming across, focus on the value you're providing to your audience. Remember, you're speaking to share your ideas, not to impress anyone with your speaking skills.
5. **Make eye contact:** Engage with your audience by making eye contact with different people in the room. This helps build trust and keeps your audience interested.
6. **Use confident body language:** Stand up straight, make use of gestures, and avoid fidgeting or crossing your arms. This will help you appear more confident and others will perceive you as more confident too.
7. **Speak slowly and clearly:** Avoid rushing through your speech. Speak slowly and clearly, pausing for emphasis and to allow your audience to absorb what you're saying.
8. **Use storytelling techniques:** People remember stories better than facts and figures. Use anecdotes, examples, and metaphors to make your message more relatable and engaging.
9. **Anticipate questions:** Think about potential questions your audience may have and prepare responses. This helps you feel more confident and in control.
10. **Join a public speaking group:** Consider joining a group like Toastmasters, where you can practice speaking in a supportive environment and receive feedback from others.
11. **Record yourself:** Record yourself speaking and watch the video to identify areas for improvement, such as filler words or nervous habits.

16.85 tokens/second



A server rack is shown with a GPU installed. The GPU has the text "UM150; 1922L AGP" and "1TFRD QNS" printed on it. Several blue cables are plugged into the GPU. The entire scene is overlaid with a large, bright orange and yellow flame effect, suggesting a fire or a warning. The background is dark with some blue lighting from the server rack.

**Biggest limitation right now:  
No GPU support. Yet.**

## My recommendation

While this is great for **small tasks** (classification, image tagging, etc.), any LLM work should be **delegated to GPU based infrastructure.**

# Your LLM options for now

Use a SaaS service: OpenAI, Claude API, etc.  
Or use HF Inference endpoints to deploy the  
model of your choice

And query that API endpoint from your app.

Or “Build your own”, at your own risk!



# Thank **you.**



**Guillaume Moigneu**

VP, Advocacy @ Upsun

guillaume@platform.sh

guillaume.id