

# Ten Commandments

To deploy fine-tuned models in prod

Kyle Corbitt |  OpenPipe  
@corbtt




**First Commandment**

**Thou Shalt Not Fine-Tune**

# Thou Shalt Not Fine-Tune

- Just use prompting!
- And optionally few-shot examples/RAG

# Thou Shalt Not Fine-Tune

-  ...Unless you can't hit your quality target
-  ...Or latency target
-  ...Or cost target

**Second Commandment**

**Thou Shalt Write a  
Freaking Prompt**

# Thou Shalt Write a Freaking Prompt

- Create a baseline
- Prove the task is possible!
  - *Remember the logistics company?*

Description	Value
Sweater that Ron left in my car	40
...	...

# Thou Shalt Write a Freaking Prompt

- Works with prompting?
  - ◆ 90% chance fine-tuning makes it better
- Doesn't work with prompting?
  - ◆ 25% chance fine-tuning works *at all*

# Your Playbook: GPT-4 Prototype, Fine-Tune to Scale







**Third Commandment**

**Thou Shalt Review Thy  
Freaking Data**

# Thou Shalt Review Thy Freaking Data

IMPORTED AT	INPUT TOKENS • 4.3M	OUTPUT TOKENS • 1.3M	SPLIT
November 30 3:25 PM	666	227	<button>TEST</button>
November 30 3:25 PM	963	213	<button>TRAIN</button>  


  






Input	Output
<p><b>system</b></p> <p>You are an expert model trained to redact potentially sensitive information from documents. You have been given a document to redact. The document is a rental_agreement. Your goal is to accurately redact the sensitive information from the document. Sensitive information can be in one of the following categories:</p> <ul style="list-style-type: none"><li>- age: a person's age</li><li>- credit_card_info: a credit card number, expiration date, CCV, etc.</li><li>- nationality: a country when used to reference place of birth, residence,</li></ul>	<p><b>tool_calls</b></p> <pre>redact({   "fields_to_redact": [     {       "string": "7th day of June, 1997",       "pii_type": "date"     },     {       "string": "Holmes-Kim",       "pii_type": "organization_name"     }   ] })</pre>

**Fourth Commandment**

**Thou Shalt Use Thy  
Actual Freaking Data**


# Thou Shalt Use Thy Actual Freaking Data

Request Logs 

15m 24h 7d All Custom  Show Filters (1)  Columns (8/10)  Add to Dataset  Export  Delete

<input type="checkbox"/> (5,364)	SENT AT	MODEL	prompt_id	DURATION	INPUT TOKENS	OUTPUT TOKENS	COST	STATUS
<input type="checkbox"/>	August 31 11:18 PM	<a href="#">gpt-4-0613</a>	pii-redaction/redact	22.03s	736	233	\$0.036	200
<input type="checkbox"/>	August 31 11:18 PM	<a href="#">gpt-4-0613</a>	pii-redaction/redact	7.91s	694	131	\$0.029	200
<input type="checkbox"/>	August 31 11:18 PM	<a href="#">gpt-4-0613</a>	pii-redaction/redact	13.22s	642	131	\$0.027	200
<input type="checkbox"/>	August 31 11:18 PM	<a href="#">gpt-4-0613</a>	pii-redaction/redact	8.83s	747	112	\$0.029	200
<input type="checkbox"/>	August 31 11:17 PM	<a href="#">gpt-4-0613</a>	pii-redaction/redact	12.97s	699	133	\$0.029	200
<input type="checkbox"/>	August 31 11:17 PM	<a href="#">gpt-4-0613</a>	pii-redaction/redact	12.68s	724	146	\$0.030	200
<input type="checkbox"/>	August 31 11:16 PM	<a href="#">gpt-4-0613</a>	pii-redaction/redact	24.84s	816	284	\$0.042	200
<input type="checkbox"/>	August 31 11:16 PM	<a href="#">gpt-4-0613</a>	pii-redaction/redact	10.36s	963	105	\$0.035	200
<input type="checkbox"/>	August 31 11:15 PM	<a href="#">gpt-4-0613</a>	pii-redaction/redact	15.69s	802	174	\$0.035	200
<input type="checkbox"/>	August 31 11:15 PM	<a href="#">gpt-4-0613</a>	pii-redaction/redact	11.97s	675	110	\$0.027	200

# Thou Shalt Use Thy Actual Freaking Data

- Some bad data is ok?
  -  Controversial!
- Make sure dataset is *correct on average*.

GPT-4 Turbo (2024-04-09)	openpipe:pii-redaction-llama3-8b
Eval4: WR 35.0% ▾	LLAMA 3 8B Eval4: WR 65.0% ▾
<pre>redact({   "fields_to_redact": [     {       "string": "María Eugenia Josefina Cavazos Salas",       "pii_type": "person_name"     },     {       "string": "Cerro-Hidalgo",       "pii_type": "organization_name"     }   ] })</pre>	<pre>redact({   "fields_to_redact": [     {       "string": "March 6, 1974",       "pii_type": "date"     },     {       "string": "María Eugenia Josefina Cavazos Salas",       "pii_type": "person_name"     },     {       "string": "Friday, April 12th",       "pii_type": "date"     },     {       "string": "March 20th",       "pii_type": "date"     },     {       "string": "Cerro-Hidalgo",       "pii_type": "organization_name"     }   ] })</pre>
Eval4: 0.0%	Eval4: 100.0%

**Fifth Commandment**

**Thou Shalt Reserve a  
Test Set**

# Thou Shalt Reserve a Test Set

Dataset Entries • 5,775 rows

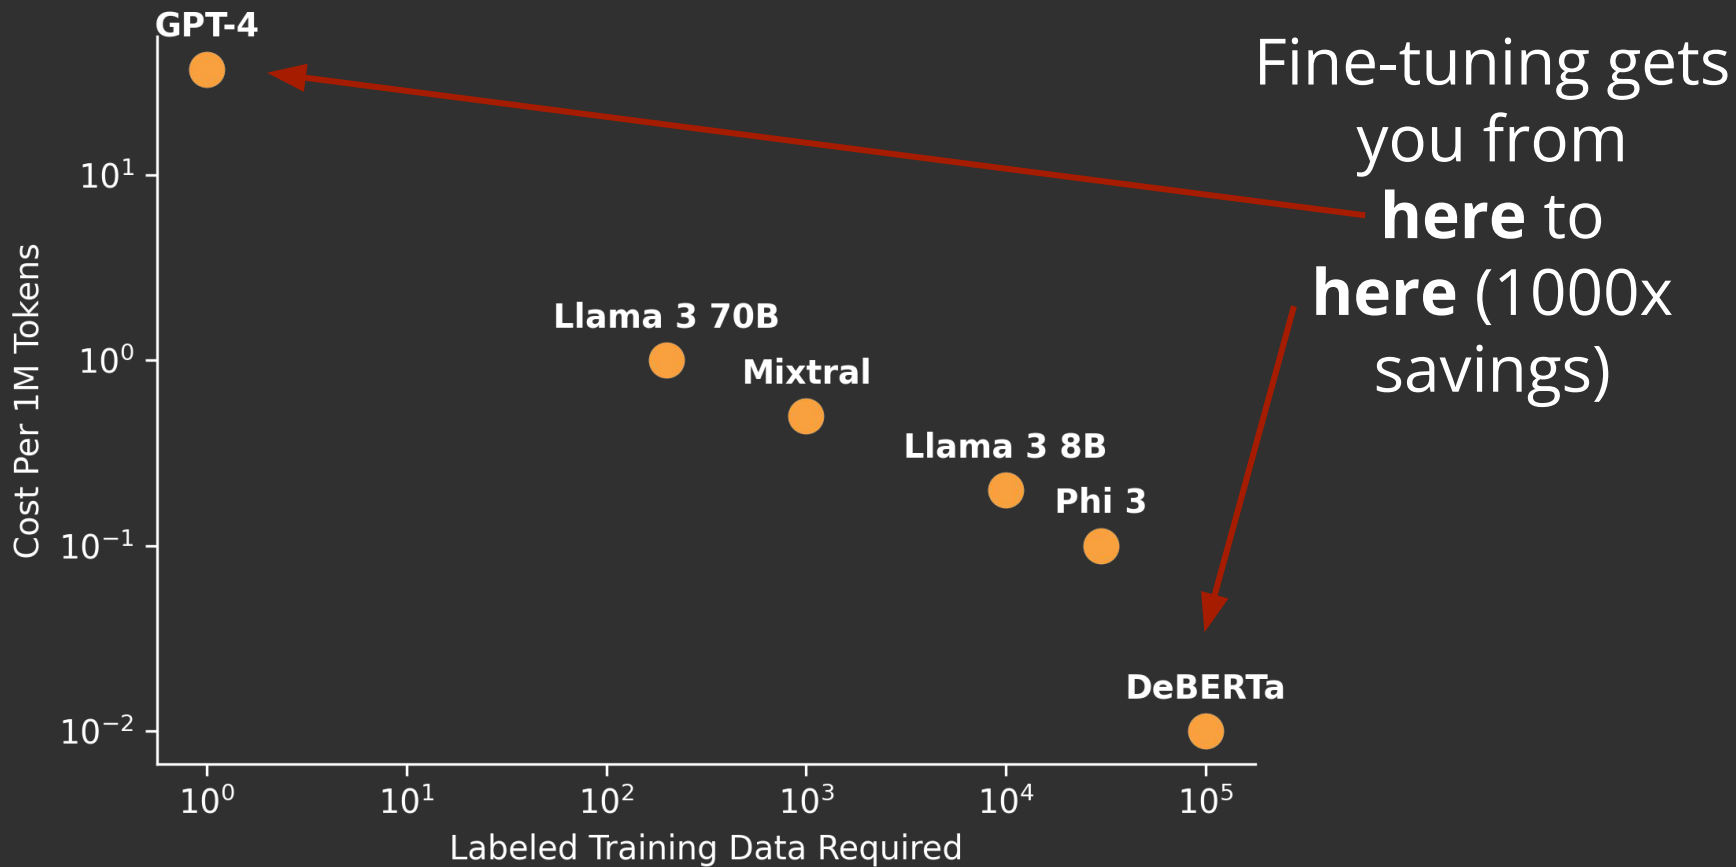
IMPORTED AT ↕	INPUT TOKENS • 4.3M ↕	OUTPUT TOKENS • 1.3M ↕	SPLIT ↕
November 30 3:25 PM	666	227	TEST
November 30 3:25 PM	963	213	TRAIN
November 30 3:25 PM	909	333	TRAIN
November 30 3:25 PM	620	158	TEST
November 30 3:25 PM	631	147	TRAIN
November 30 3:25 PM	840	150	TRAIN
November 30 3:25 PM	718	329	TEST
November 30 3:25 PM	754	346	TRAIN
November 30 3:25 PM	572	125	TRAIN
November 30 3:25 PM	746	266	TRAIN

**Sixth Commandment**

**Thou Shalt Choose an  
Appropriate Model**



# Thou Shalt Choose the Appropriate Model



**Seventh Commandment**

**Thou Shalt Write Fast  
Evals**

# Thou Shalt Write Fast Evals

## Overall Performance

MODEL	WIN RATE	WINS	TIES	LOSSES
pii-redaction-mixtral <b>FT</b>	<b>64.4%</b>	154	207	39
pii-redaction-mistral-optimized <b>FT</b>	<b>60.1%</b>	140	201	59
<b>GPT-4 (-0613)</b>	<b>57.3%</b>	140	178	82
<b>GPT-4 Turbo (-1106)</b>	<b>43.1%</b>	102	141	157
<b>GPT-3.5 Turbo (-1106)</b>	<b>25.6%</b>	58	89	253

**Eighth Commandment**

**Also, Thou Shalt Write  
Slow Evals**

# Thou Shalt Write Slow Evals



Here are four random dinosaurs:

1. **Velociraptor** - A small, fast predator known for its intelligence and sickle-shaped claw on each foot.
2. **Stegosaurus** - Easily recognized by its rows of large, bony plates along its back and a spiked tail.
3. **Triceratops** - Known for its impressive skull with two large horns over the eyes and a smaller one on the nose.
4. **Brachiosaurus** - A giant sauropod with long neck and legs, known for its upright posture unlike other sauropods that held their necks out forward.



# Thou Shalt Write Slow Evals

RU

OK. Other way around. EU has to converge into 10 countries. Name them, and list the capital

Which response do you prefer?

Your choice will help make ChatGPT better.



Response 1

If the EU were to be consolidated into just 10 countries, we could potentially see mergers based on geographical proximity, linguistic similarities, and shared cultural or economic interests. Here's how it could potentially look:

1. **Nordic Union:** (Sweden, Finland, Denmark, Estonia)
  - Capital: Stockholm
2. **British Isles:** (Ireland)
  - Capital: Dublin (Assuming post-Brexit Britain isn't a part of this)
3. **Western Core:** (Germany, Netherlands, Luxembourg, Belgium)
  - Capital: Berlin
4. **Latin West:** (France, Monaco)
  - Capital: Paris



Response 2

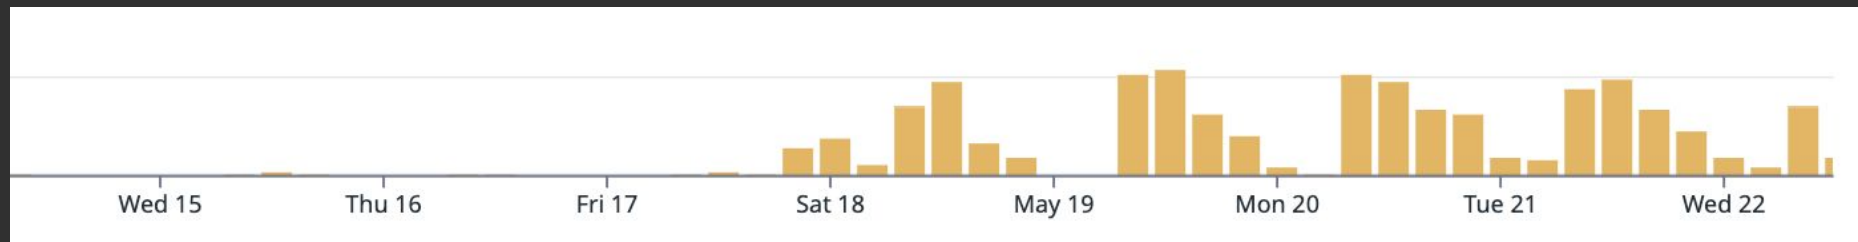
All right, if the EU were to converge into 10 larger countries, based on geographical proximity, cultural and historical ties, and economic considerations, a hypothetical configuration might be:

1. **Nordic Union**
  - Countries: Denmark, Finland, Sweden, Estonia, Latvia, Lithuania
  - Capital: Stockholm
2. **Western Europe**
  - Countries: Ireland, United Kingdom (assuming it was still in the EU for this exercise), Belgium, Netherlands, Luxembourg
  - Capital: Brussels
3. **Central European Federation**
  - Countries: Germany, Poland, Czech Republic

## **Ninth Commandment**

**Thou Shalt Not Fire and  
Forget**

# Thou Shalt Not Fire and Forget





## Tenth Commandment

**Thou Shalt Not Take the  
Commandments Too  
Seriously**

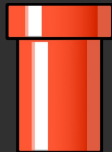


**THE CODE IS MORE LIKE "GUIDELINES"**

**RATHER THAN ACTUAL RULES**

# Credits!

- \$222.22 in OpenPipe credits!
- Email [hello@openpipe.ai](mailto:hello@openpipe.ai)
  - “FT Course credits” in the subject line



OpenPipe

# Questions?

**Kyle Corbitt**

@corbtt

kyle@openpipe.ai



**OpenPipe**

openpipe.ai