

The sustainability issue with AI and what we can change as marketers

Lazarina Stoy.

MLforSEO

@lazarinastoy | www.mlforseo.com





Amey Chauhan

• 3rd+
Zoologist & Digital Strategist | Scaling Wildlife Education ...
4mo •

+ Follow

...

I turned my thumbnail in Studio Ghibli style

... more



Tough, Powerful & Desert-Proof!

Tough, Powerful & Desert-Proof!

18

1 comment



Harnoor Singh

• 2nd
Devrel @HydraDB | Ex-Microsoft | YouTuber
4mo •

+ Follow

...

Google has also cooked with Studio Ghibli!

Better character and consistent using Nano Banana Pro!



82

2 comments



Sanjeev Roy

• 2nd
Deputy General Manager – Operations & Sales @ Senco Gol...
7mo •

+ Follow

...

From Talk of the Town to Timeless Inspiration

... more



42

2 comments



Variety

Studio Ghibli-Style AI Images



GeekWire

Seattle engineer's Ghibli-style image ...



HT Tech

Studio Ghibli-style AI images ...



StoryTribе

Ghibli-style AI is cool. But what's the ...



Dawn Images

Studio Ghibli-style charac...



Thred

Studio Ghibli AI trend ...



CNN

Viral Studio Ghibli-style AI images ...



ghibliimage.net

Ghibli AI Image Generator - Create ...



The New York Times

Studio Ghibli Style A...



Megri Digitizing

Photos Into Studio Ghibli ...



LinkedIn

Create Stunning Studio Ghibli-Style AI ...

studio ghibli image with ai




×

🎤

📷

🔍

About 56,700,000 results (0.45s)



action figure with ai

About 1,700,000,000 results (1.00s)



=



1 image = fully charging your smartphone



=



1 image = wasting 250 ml water for cooling servers

Number of AI-Created Images*

EVERYPIXEL

DALL-E 2

916 million

Models based on Stable Diffusion

12.590 billion

Adobe Firefly

1 billion


Midjourney

964 million

15.470 billion

Sources: Adobe;
our estimates, based on Photutorial, OpenAI, Civitai

*As of August 2023



By 2023

15B images

135K tons CO₂

By 2025 (linear)


35B images

315K tons CO₂

By 2025 (growth)

45–65B images

405–585K tons CO₂



Based on: 1M AI images $\approx$ 9 metric tons CO₂ ($\approx$ 2 cars/year). Shaded area shows growth scenario range.



Sam Altman  
@sama



it's super fun seeing people love images in chatgpt.

but our GPUs are melting.

we are going to temporarily introduce some rate limits while we work on making it more efficient. hopefully won't be long!

chatgpt free tier will get 3 generations per day soon.

6:32 PM · Mar 27, 2025 · **8.8M** Views


 2.5K

 4.8K

 66K

 4.9K





2.5B

prompts processed by ChatGPT every single day


Sources: DemandSage, Superlines, OpenAI (2026)




Total daily AI prompts (est.)
~4.2 billion

Monthly active users (all)
~2.4 billion

Market share
(by web visits, Jan 2026)



Est. daily prompts
(millions, early 2026)



* Claude: 25B+ API calls/month = ~830M/day. Consumer MAU only ~19M, but enterprise API volume is massive. Bubble area proportional to value. Only ChatGPT's 2.5B daily prompts is officially confirmed by OpenAI.

One ChatGPT Query vs One Google Search

ChatGPT

2.9 Wh

per query

+ 519ml water (GPT-4, 100-word email)

10×

Google Search

0.3 Wh

per query

Source: IEA (2025)



Generative AI vs Traditional ML

Not all AI is created equal in its environmental cost

Image generation

Model training

Zero-shot multipurpose models



Text generation

Inference

Task-specific models for the same task

Source: Luccioni, Jernite & Strubell (2023); AllAboutAI (2025)

To put this into perspective...

Racks of servers hum along for months, ingesting training data, crunching numbers, and performing computations. This is a time-consuming and expensive process—it's estimated that training OpenAI's GPT-4 took over \$100 million and consumed 50 gigawatt-hours of energy, enough to power San Francisco for three days. It's only after this training, when consumers or customers “inference” the AI models to get answers or generate outputs, that model makers hope to recoup their massive costs and eventually turn a profit.

Source: Luccioni, Jernite & Strubell (2023); AllAboutAI (2025)



1,050 TWh

Projected data center electricity consumption by 2026 —
more than Japan's entire national consumption

Sources: IEA, MIT, Lawrence Berkeley National Lab, Carbon Brief (2025)

Yet...
DeepSeek-V3
achieved 95%
lower energy while
maintaining
competitive
performance




Figure 1: Benchmark performance of DeepSeek-V3 and its counterparts.

Source: Luccioni, Jernite & Strubell (2023); AllAboutAI (2025)

The Corporate Response? Scale at All Costs.

Stargate Initiative

Business / Tech

Trump announces a \$500 billion AI infrastructure investment in the US

Google

DIVE BRIEF

Google plans \$75B investment to build out cloud AI capacity

The majority of the spending will target technical infrastructure, including servers and data centers, CFO Anat Ashkenazi said Tuesday.

Apple

AI EFFECT

Apple to open AI server factory in Texas as part of \$500 billion U.S. investment

Nuclear Revival

NEWS ENERGY

Big Tech Backs Small Nuclear > Google and Amazon invest in small modular reactors to power data centers



PRESS STATEMENT | April 14, 2026

NAACP Sues xAI for Illegal Pollution from Data Center Power Plant



“There are no loopholes or exceptions — xAI is breaking the law while leaving local communities to deal with the consequences.”

— PATRICK ANDERSON, SELC SENIOR ATTORNEY



Most industrial water (like for a paper mill or a steel plant) is self-supplied, not municipal but data centers are mostly municipal...for now at least. Who knows what the future holds

Why is Everyone So Wrong About AI Water Use??



Hank Green
3.03m subscribers

Subscribe

186k





Share

Save

Clip

Download





Most industrial water (like for a paper mill or a steel plant) is self-supplied, not municipal but data centers are mostly municipal...for now at least. Who knows what the future holds

Why is Everyone So Wrong About AI Water Use??



Hank Green
3.03m subscribers

Subscribe

👍 186k



➦ Share


🔖 Save

✂️ Clip

⬇️ Download



**What they measure: water
used to generate the
response**



Most industrial water (like for a paper mill or a steel plant) is self-supplied, not municipal but data centers are mostly municipal...for now at least. Who knows what the future holds

Why is Everyone So Wrong About AI Water Use??



Hank Green
3.03m subscribers

Subscribe

186k



Share

Save


Clip

Download




What is left out: all the background operations like model training, query fan-out, research, prompt analysis and so on


GenAI companies are not transparent about the human cost involved in AI scale.



**What they measure: CO2,
water, money costs**



**What is left out: people
nearby data centers
experience tiredness
(+12%), nausea (+33%),
dizziness (+150%), tingling
sensation (+25%).**



No Sign of Slowing Down

 AI is being integrated into everything –
from search engines to toothbrushes




X2Y

@x2y.tech

+ Follow

SAM ALTMAN: "People talk about how much energy it takes to train an AI model ... But it also takes a lot of energy to train a human. It takes like 20 years of life and all of the food you eat during that time before you get smart."









Marketing & SEO Sustainability Issues

Where our industry specifically makes it worse



AI Overviews appear in ~25% of All Searches

...and they're sh#%... like, a lot of the time



what if...

- + every Google search**
- + re-prompting in AI mode**

**10× more energy
× 8.5B daily searches
× 1-3 follow-ups in AI mode
= a massive, invisible environmental impact**

Bot Traffic

51%

of all web traffic is now
automated (bots & scrapers)

Imperva Bad Bot Report 2025

37%

of traffic is "bad bots" —
scrapers, fraud, manipulation

Up from 30% in 2023

Query Fan Out

Human searches

~5

websites visited per query

Type → click → read → next

AI agent query

~5,000

websites hit per query

1,000x multiplier

The multiplier effect is real. When a human user searches for information, they might visit five websites.

When an AI agent fulfills the same request, it may visit 5,000 websites — a roughly 1,000x multiplication factor per query SoloSoft, according to Cloudflare's CEO at SXSW.

Agents

AI agent traffic exploded.

Traffic generated by autonomous systems that can navigate and act on the web grew 7,851% year over year

AI-driven traffic by month, 2025

Indexed to January, 2025. Includes training crawlers, scrapers, AI agents, and agentic browsers.



Source: Human Defense Platform telemetry, 2025


AI Crawling

The scale of AI crawling is insane.

Cloudflare's network saw about 50 billion crawler requests per day from AI bots alone by late 2025.

What's driving synthetic queries


Estimated daily volume by category




Sources: Imperva Bad Bot Report 2025, HUMAN Security 2026 Benchmark, Cloudflare 2025 Trends, Cloudflare CEO (SXSW March 2026). Cloudflare blocked 416B AI bot requests since mid-2025 alone.

And... my favourite: vibe coding

Vibe Coding Workflow



Vibe coding's hidden pitfalls.



Each response costs energy. Poor prompting = exponential waste.

"Everyone is a programmer now" —
but at what environmental cost?

What If We All Train Custom Models?

If just 10% of SEO professionals start training custom models:

The cumulative compute, energy, and water consumption would be staggering — and that's just one niche industry.

Now multiply by every marketer, designer, developer, content creator, and hobbyist doing the same.

Products/tools/websites/projects
that serve no one are being
built every day.

Because they can be built, not
because they *should* exist.

What This Leads To

Trust erosion, cognitive decline, and the dead internet



74%


of new web pages contain
AI-generated content

Ahrefs analysis of 900,000 newly published pages, April 2025





**The dead internet
theory is real**



**Alexis Ohanian,
Reddit co-founder**



**used to be a lunatic fringe
conspiracy theory, but it's
looking a lot more real**

**Adam Aleksic,
Linguist**



Trust Crisis


- 50% of people **distrust** AI-generated responses (KPMG)
- Visits from bots balloon costs but generate **no monetary gain**
- **Brand trust is harder** in a sea of AI slop

Cognitive Crisis

- MIT: Relying on AI leads to "**cognitive atrophy**"
- Harvard: "If AI does your thinking, it **undercuts creativity**"
- Risk of "**cognitive sedentism**" — brain adapts to NOT thinking

What Can We Do

Practical steps for responsible AI use in marketing






Do the Work Yourself (When You Can)



The struggle is the learning — remove it, and you remove growth.







Better for the
planet (and more
fun)

authentic, memorable,
zero compute cost.

Personality > polish.






Plan Before You Prompt

Be AI-assisted, not AI-driven.

Four principles to follow before you touch any AI tool.



Research prompting techniques first.

10 minutes of research = 1–3 interactions instead of 10–50.

Advanced Prompt Engineering : Types and Example Techniques

Created By: Aishwarya Naresh Reganti



Step-by-Step Modular Decomposition

Chain-of-Thought (CoT)

Guiding an LLM through intermediate steps. It involves presenting a sequence of reasoning steps, enabling the model to focus on solving the problem incrementally.

Tree-of-Thoughts (ToT)

Enables LLMs to explore coherent units of text as intermediate steps in problem-solving. Allows models to make deliberate decisions, consider multiple reasoning paths, and self-evaluate choices.

Graph-of-Thought (GoT)

Unlike simple sequential chains, GoT models thoughts as graphs, utilizing DAGs to represent connections and allowing paths to fork and converge. The GoT model operates in two stages, generating rationales and producing the final answer, leveraging a Graph-of-Thoughts encoder for representation learning and incorporating both linear and non-linear aspects.

Comprehensive Reasoning and Verification

Automatic Prompt Engineer (APE)

Treats instructions as programmable elements, optimizing them through a search across a pool of candidate instructions proposed by an LLM.

Self-Consistency

Involves sampling diverse reasoning paths using few-shot CoT and prioritizing consistent answers, contributing to improved performance in tasks requiring arithmetic and commonsense reasoning.

The Chain-of-Verification (CoVe)

It systematically verifying responses. It starts with the model's initial response, poses independent verification questions, and generates a final corrected response based on the outcomes.

ReAct

Combines reasoning and action in LLMs for dynamic tasks by generating verbal reasoning traces and task-specific actions in an interleaved manner.

Usage of External Tools/Knowledge or Aggregation

Active Prompting


Enhances LLM adaptability by dynamically selecting task-specific example prompts. It dynamically chooses prompts based on model-generated examples, ensuring adaptability to diverse tasks.

Automatic Multi-step Reasoning and Tool-use (ART)

ART combines Chain-of-Thought prompting and tool usage, using a frozen LLM to automatically generate intermediate reasoning steps from task-specific examples. It integrates external tools during testing.


Chain of Knowledge (CoK)

Strengthens LLMs by integrating diverse grounding information dynamically. The process includes reasoning preparation, dynamic knowledge adaptation, and answer consolidation.



On prompting, it's also worth adding system prompts for repeatable tasks

A system prompt is used to determine the persona of the agent in the context of a conversation



```
h 1
...
execute code or interact with external APIs.
...
ain memory beyond this session.

...
for anything requiring this functionality, calmly explain that you're unable to do that right now.
ons>

<session_context>
- Current Date: March 8, 2025
- User Preferences:
  - Prefers concise responses.
  - Uses American English spelling.
</session_context>

<example_interaction>
User Request: "Write me a LinkedIn post about our latest product update. [Context about latest product update]."

Incorrect Response:
"Our new update is live - "[Update title]"! Click the link to see what's new: [Link]"

Correct Response:
"Exciting news! Our latest update to [Product Name] introduces [Feature Highlights], making it easier than ever to [Benefit]. Read more here:
[Link] #Innovation #Tech"

Why the incorrect response is wrong:
- It lacks an engaging hook.
- It does not highlight the update's benefits.
- It fails to align with LinkedIn's professional tone.
</example_interaction>

<output>
- Begin with a short (1-3 sentence) summary or introductory statement.
- Follow the relevant guidelines from <writing_types> and <proprietary_style_guidelines>.
- Respect the <format_rules> for readability, tone, and style.
- Observe the <restrictions> to safeguard brand integrity and confidentiality.
- Provide a concise conclusion without posing additional questions.
</output>

~/pro/zazencodes-season-2/src/anatomy-of-a-system-prompt
```

How to Write Better System Prompts



ZazenCodes
13.7k subscribers

Subscribe

339




Share

Save

Download




Criteria	User prompts	System prompts
Purpose	Task-specific instructions	Overall framework & guidelines
Frequency of use	Used frequently, often changed	Set once, rarely changed
Scope	Narrow, focused on individual tasks	Broad, applies to all interactions
Content focus	Specific details, context & desired outcomes	General rules, tone, ethics & brand values
Example	"Write a follow-up email to prospect X about Y product"	"You are a seasoned account executive for a B2B SaaS company..."
Typical length	Short to medium (1-5 sentences)	Medium to long (paragraph or more)
Primary impact	Output content & structure	Overall tone, behavior & approach
When to use	For each specific task or request	At the beginning of AI solution setup or a new session
Modifiability	Easily modified for each new task	Requires careful consideration to change



Architect LLMs into repeatable systems, via role assignment

For a complex workflow (e.g. writing content briefs),
design a multi-agent workflow with clear
assignment of subtasks to specialised agents.



Aspect		LLM orchestration	LLM agent orchestration
	Focus	Managing workflows for LLMs	Autonomous decision-making by the LLM
	Decision making	Limited; linear workflows	Dynamic; LLM selects tools and steps
	Components	Prompts, workflows, tools	Prompts, tools, reasoning, adaptive logic
	Complexity	Task-specific and simpler	Multi-step, adaptable, and more complex
	Autonomy	Limited; user defines the steps	High; the agent determines steps dynamically

Orchestrator agent example



```
1  # Goal
2
3  Route customer requests to the appropriate specialist agent based on intent.
4
5  ## Routing logic
6
7  **Billing specialist:** Customer mentions payment, invoice, refund, charge, subscription,
8  **Technical support specialist:** Customer reports error, bug, issue, not working, broken
9  **Scheduling specialist:** Customer wants to book, reschedule, cancel, or check appointment
10 **Human escalation:** Customer is angry, requests supervisor, or issue is unresolved after
11
12 ## Handoff process
13
14 1. Classify customer intent based on first message
15 2. Provide brief acknowledgment: "I'll connect you with our [billing/technical/scheduling]
16 3. Transfer conversation with context summary:
17     - Customer name
18     - Primary issue
19     - Any account identifiers already collected
20 4. Do not repeat information collection that already occurred
```

Example 1: Technical support agent

```
Technical support specialist

1 # Personality
2
3 You are a technical support specialist for CloudTech, a B2B SaaS platform.
4 You are patient, methodical, and focused on resolving issues efficiently.
5 You speak clearly and adapt technical language based on the user's familiarity.
6
7 # Environment
8
9 You are assisting customers via phone support.
10 Customers may be experiencing service disruptions and could be frustrated.
11 You have access to diagnostic tools and the customer account database.
12
13 # Tone
14
15 Keep responses clear and concise (2-3 sentences unless troubleshooting requires more detail).
16 Use a calm, professional tone with brief affirmations ("I understand," "Let me check that").
17 Adapt technical depth based on customer responses.
18 Check for understanding after complex steps: "Does that make sense?"
19
20 # Goal
21
22 Resolve technical issues through structured troubleshooting:
23
24 1. Verify customer identity using email and account ID
25 2. Identify affected service and severity level
26 3. Run diagnostics using 'runSystemDiagnostic' tool
27 4. Provide step-by-step resolution or escalate if unresolved after 2 attempts
28
29 This step is important: Always run diagnostics before suggesting solutions.
30
31 # Guardrails
32
33 Never access customer accounts without identity verification. This step is important.
34 Never guess at solutions—always base recommendations on diagnostic results.
35 If an issue persists after 2 troubleshooting attempts, escalate to engineering team.
36 Acknowledge when you don't know the answer instead of speculating.
37
38 # Tools
39
40 ## 'verifyCustomerIdentity'
41
42 **When to use:** At the start of every conversation before accessing account data
43 **Parameters:**
44
45 - 'email' (required): Customer email in standard written format (e.g., "user@company.com")
46 - 'account_id' (optional): Account ID if customer provides it
47
48 **Error handling:**
49 If verification fails, ask customer to confirm email spelling and try again.
50
51 ## 'runSystemDiagnostic'
52
53 **When to use:** After verifying identity and understanding the reported issue
54 **Parameters:**
55
56 - 'account_id' (required): From 'verifyCustomerIdentity' response
57 - 'service_name' (required): Name of affected service (e.g., "api", "dashboard", "storage")
58
59 **Usage:**
60
61 1. Run the tool if the service is affected.
```

Principles demonstrated:

- ✓ Clean section separation (**# Personality**, **# Goal**, **# Tools**, etc.)
- ✓ One action per line (see **# Goal** numbered steps)
- ✓ Concise instructions (tone section is brief and clear)
- ✓ Emphasized critical steps ("This step is important")
- ✓ Format conversion in parameter descriptions (email normalization)
- ✓ Dedicated guardrails section
- ✓ Precise tool descriptions with when/how/error guidance
- ✓ Explicit error handling instructions

Example 2: Customer service refund agent

```
Refund processing specialist

1 # Personality
2
3 You are a refund specialist for RetailCo.
4 You are empathetic, solution-oriented, and efficient.
5 You balance customer satisfaction with company policy compliance.
6
7 # Goal
8
9 Process refund requests through this workflow:
10
11 1. Verify customer identity using order number and email
12 2. Look up order details with 'getOrderDetails' tool
13 3. Confirm refund eligibility (within 30 days, not digital download, not already refunded)
14 4. For refunds under $100: Process immediately with 'processRefund' tool
15 5. For refunds $100-$500: Apply secondary verification, then process
16 6. For refunds over $500: Escalate to supervisor with case summary
17
18 This step is important: Never process refunds without verifying eligibility first.
19
20 # Guardrails
21
22 Never process refunds outside the 30-day return window without supervisor approval.
23 Never process refunds over $500 without supervisor approval. This step is important.
24 Never access order information without verifying customer identity.
25 If a customer becomes aggressive, remain calm and offer supervisor escalation.
26
27 # Tools
28
29 ## 'verifyIdentity'
30
31 **When to use:** At the start of every conversation
32 **Parameters:**
33
34 - 'order_id' (required): Order ID in uppercase alphanumeric format (e.g., "ORD123456"). Co
35 - 'email' (required): Customer email in standard written format (e.g., "john.smith@retailc
36
37 ## 'getOrderDetails'
38
39 **When to use:** After identity verification
40 **Returns:** Order date, items, total amount, refund eligibility status
41
42 **Error handling:**
43 If order not found, ask customer to verify order number and try again.
44
45 ## 'processRefund'
46
47 **When to use:** Only after confirming eligibility
48 **Required checks before calling:**
49
50 - Identity verified
51 - Order is within 30 days
```

Principles demonstrated:

- ✓ Specialized agent scope (refunds only, not general support)
- ✓ Clear workflow steps in **# Goal** section
- ✓ Repeated emphasis on critical rules (refund limits, verification)
- ✓ Detailed tool usage with "when to use" and "required checks"
- ✓ Format conversion in parameter descriptions (order IDs, emails)
- ✓ Explicit error handling per tool
- ✓ Escalation criteria clearly defined

this enables:

- ✓ clearly assigned tasks
- ✓ reduced context windows for specialist agents
- ✓ easier (planet-friendly) updates and error-handling
- ✓ clear performance management
- ✓ improved speed via reduced latency per interaction




Plan the task, not just the prompt.

Define your goal, constraints, and output format
before touching any AI tool.



**Do I actually need AI
for this?**



Choose the right model for the task.

**Custom-task models often
outperform LLMs
for niche tasks.**

Smaller, faster, cheaper, greener —and often more accurate.

Keyword classification

Instead of GPT-4

Unpredictable results
Different output every time
Expensive at scale



Use BERT / DistilBERT

Consistent, reliable
Same output every time
100× less energy

📖 Coding Script or Notebook

💻 Google Colab

✓ Free

Transformer-Based Search Intent Classification with BART Zero-Shot & SERP Signals (Colab Notebook)

Three-signal hybrid pipeline that fuses rule-based regex patterns, SERP feature percentage breakdowns (70...

👤 Lazarina Stoy.



📖 Coding Script or Notebook

💻 Google Colab

✓ Free

Semantic Search Intent Classification with Google Cloud NLP & spaCy in Python (Colab Notebook)

Combines spaCy dependency parsing (main verb, direct object extraction) with Google Cloud NLP entity...

👤 Lazarina Stoy.



📖 Coding Script or Notebook

💻 Google Colab

✓ Free

Search Intent Classification from SERP Features in Python: Rule-Based Notebook for SEO

Three connected classification functions in one notebook: SERP-feature-to-intent mapping, percenta... breakdowns across four intent

👤 Lazarina Stoy.



📖 Coding Script or Notebook

💻 Google Colab

✓ Free

KeyBERT for SEOs: Google Colab Template

👤 Lazarina Stoy.



Content categorization

Instead of LLM prompting

Hit or miss results
General purpose model
Non-replicable



Use FastText / spaCy

Purpose-built for text
Reliable & replicable
Runs locally, zero API cost

📖 No-Code Template

💻 Apps Script

✓ Free

Text Classification in Google Sheets with Google Cloud Natural Language API

👤 Lazarina Stoy.



📖 Coding Script or Notebook

💻 Google Colab

✓ Free

IBM Watson NLU Complete API Implementation Template for Multi-Feature Text Analysis (Notebook)

👤 Lazarina Stoy.



📖 Coding Script or Notebook

💻 Google Colab

🎓 Academy

Fuzzy Matching: Map Keywords to Seed Terms or Topics (Notebook)

👤 Lazarina Stoy.



Sentiment analysis

Instead of ChatGPT

Subjective, varying scores
Cannot batch-process
No numeric confidence



Use VADER / roBERTa

Numeric scores (0–1)
Batch thousands instantly
Reproducible every time

📖 No-Code Template

💻 Apps Script

✓ Free

Sentiment Analysis No-code Template with Google Cloud Natural Language API

👤 Lazarina Stoy.



📖 No-Code Template

💻 Apps Script

✓ Free

IBM Watson NLU Sentiment and Emotion Analysis Template for Google Sheets with Multi-Dimensional Affect Detection

👤 Lazarina Stoy.



Entity extraction

Instead of LLM parsing

Hallucinated entities
Inconsistent format
Slow on large datasets



Use spaCy NER models

Precise entity types
Structured output
Processes 1000s of
pages/min

📖 Coding Script or Notebook

💻 Google Colab

✓ Free

Entity Analysis & Relationship Mapping with IBM Watson NLU in Google Colab

👤 Lazarina Stoy.



📖 Coding Script or Notebook

💻 Google Colab

✓ Free

Extracting Entities with Google Cloud Natural Language API in Python (Tutorial)

👤 Lazarina Stoy.



📖 No-Code Template

💻 Apps Script

✓ Free

Entity Analysis with with Google Cloud Natural Language API (No-Code Template)

👤 Lazarina Stoy.





	Google Cloud NLP	Deepseek R1	ChatGPT o1	Gemini (any)	Qwen Chat 2.5 Plus	Amazon Comprehend	IBM Watson NLU
Entities Recognized	71	39	15	0	25	22	17
Unique Entities	59	38	15	0	25	16	16
Entity Types (Categories)	7	9	4	0	7	5	6

	A	B	C	D
1	Module	Google Natural Language API	Amazon Comprehend	IBM Watson NLU
2	Entity Identification	Yes	Yes	Yes
3	Entity Sentiment Analysis	Yes (combined with Entity Identification)	Yes	Yes (combined with Entity Identification)
4	Sentiment Analysis (Document Level)	Yes	Yes	Yes
5	Content Moderation	Yes	No	No
6	Text Classification	Yes	No	Yes
7	Syntax Analysis / Semantic Roles	Yes (Syntax Analysis)	Yes (Syntax Analysis)	Yes (as Semantic Roles)
8	Key Phrases / Keyword Extraction	No	Yes	Yes
9	Language Detection	No	Yes	No
10	PII Detection	No	Yes	No
11	Relations (Entity Relationship Mapping)	No	No	Yes
12	Concepts Extraction	No	No	Yes
13	Emotion Detection	No	No	Yes
14	Metadata Extraction	No	No	Yes

[📖 Guide](#)
[📄 Google Sheets](#)
[✓ Free](#)

Entity & Sentiment Analysis APIs: Comparative Review

👤 Lazarina Stoy.





Introduction to Machine Learning for SEOs Course | academy.mlforseo.com


URL pattern matching

Regex + lightweight ML
beats an LLM every
time.

Faster. Deterministic.
Zero API calls.



Log file analysis with custom ML




**When you must train,
train efficiently.**



Weight Pruning

Remove unnecessary parameters.

Reduce model size by 50–90% with minimal accuracy loss.



Quantization

32-bit → 8-bit or 4-bit.

Dramatically cut compute and memory.



Knowledge Distillation

Train a small "student" model
to mimic a large "teacher."

Same quality, fraction of the cost.



LoRA & Fine-tuning

Adapt existing models to your task
instead of training from scratch.

100× less compute.



Carbon-Aware Computing

Schedule training during off-peak hours when the grid is cleanest.

34% lower carbon footprint.



**We don't all need
our own tools.**

We all do the same things for different websites.

We don't all need to build our own little tool —
even though we can now.

Use existing tested tools and APIs.

Don't rebuild what already exists.
Search before you build.

Share pre-trained models within your org and industry.

One good model, shared by many,
beats 100 bad ones built alone.

Contribute to open-source ML projects.

Rather than duplicating effort,
improve what exists.

Fine-tune, don't custom-train.

Use pre-built workflows and adapt them to your specific use case.

**Replace *wasteful*
workflows.**

LLM for complex tasks
(e.g. content briefs)



**SERP Analysis Agent +
Content Brief Template + NLP
extraction pipeline**

✓ Consistent, repeatable output

Vibe-coded **one-off** scripts



**Maintained open-source
libraries or internal (team)
tools**

✓ Tested, documented, shared



AI-assisted, not AI-driven.

Better for the planet. Better for your brain.
Better for your work.

Resource Library

Explore our collection of templates, guides, scripts, and tools for machine learning and SEO.

 Search resources by title or description...

FILTER BY TYPE

All Types

Checklist

Coding Script or Notebook

Guide

No-Code Template

Other

Visualization Template

Workbook

FILTER BY PLATFORM

All Platforms

Apps Script

Google Colab

Google Docs / Word

Google Sheets

Looker Studio

FILTER BY ACCESS

All Access Types

Free (Access via Email)

Paid (Purchase Available)

Academy (Access via Course)

FILTER BY COURSE

All Courses

AI Search & LLMs: Entity SEO and Knowledge Graph Strategies for Brands

Introduction to Machine Learning for SEOs

Semantic ML-enabled Keyword Research

FILTER BY CREATOR

All Creators

Beatrice Gamba

Lazarina Stoy.

Showing **50** resources

Explore AI/ML tools from our Academy and Courses

All

Introduction to Machine Learning for SEOs

Semantic AI-powered SEO Keyword Research

AI Blog-to-Social Post Transformer

Transform long-form blog posts into punchy, platform-optimized social snippets.

Introduction to Machine Learning for SEOs

Content Briefs Generator

Generate tightly structured, SEO-ready content briefs from your semantic keyword universe.

Semantic AI-powered SEO Keyword Research

Semantic Drift Analyzer

Visualize and measure how far your query fan-outs drift from the original search intent—complete with similarity scores, 3D vector-space plots, and CSV export. Includes optional **Semantic Boost (USE)** for local synonym & context capture (no API keys).

Semantic AI-powered SEO Keyword Research

Synthetic vs User-generated Queries Classifier

Distinguish user queries from AI-generated synthetic queries using linguistic analysis, rule-based classification, and machine learning.

Semantic AI-powered SEO Keyword Research

AI Meta Data Creator

Turn posts into SEO-perfect titles and meta descriptions at scale.

Introduction to Machine Learning for SEOs

Courses

Machine Learning courses, designed to take your from A-to-Z on different topics, related to organic search marketing and SEO. Learn at your own pace with easy-to-follow lessons that take your from zero-to-hero on a given topic or suite of tools.

Showing all 6 results

Sort by popularity ▼

Semantic ML-enabled Keyword Research Beginner-friendly course, designed to teach you everything you need to know about building a semantic keyword universe and understanding user searches.   Theory & applications  Practical walkthroughs & templates  Certificate of completion	Introduction to machine learning for SEOs course Beginner-friendly intro to the world of ML/AI for organic search marketers, beyond ChatGPT!   Theory & applications  Practical walkthroughs & templates  Certificate of completion	AI Search & LLMs: Entity SEO and Knowledge Graph Strategies for Brands Course Beginner-friendly course, designed to teach you everything you need to know about entities, knowledge graphs, and brand visibility in traditional and modern (AI Search / LLM) search systems.   Theory & applications  Practical walkthroughs & templates  Certificate of completion
Semantic ML-enabled Keyword Research Course in the LLM Era by Lazarina Stoy. ★★★★★ Beginner 400,00 € Add to cart	Introduction to Machine Learning Course by Lazarina Stoy. ★★★★★ Beginner 250,00 € Add to cart	AI Search & LLMs: Entity SEO and Knowledge Graph Strategies for Brands Course by Beatrice Gamba ★★★★★ Beginner 200,00 € Add to cart

SALE

BUNDLE & SAVE

AI Search & LLMs: Entity SEO and Knowledge Graph Strategies for Brands Course

Semantic ML - enabled Keyword Research

Introduction to machine learning for SEOs course

Theory & applications

Practical walkthroughs & templates

Three separate certificates upon completion

Course Bundle: Introduction to ML for SEOs + AI Search & LLM Entity SEO + Semantic Keyword Research – 3 Courses

Course Bundle

~~850,00 €~~ **650,00 €**

Add to cart

SALE

BUNDLE & SAVE

AI Search & LLMs: Entity SEO and Knowledge Graph Strategies for Brands Course

Introduction to machine learning for SEOs course

Theory & applications

Practical walkthroughs & templates

Two separate certificates upon completion

Course Bundle: Introduction to ML for SEOs + AI Search & LLM Entity SEO – 2 Courses

Course Bundle

~~450,00 €~~ **400,00 €**

Add to cart

SALE

BUNDLE & SAVE

AI Search & LLMs: Entity SEO and Knowledge Graph Strategies for Brands Course

Semantic ML - enabled Keyword Research

Theory & applications

Practical walkthroughs & templates

Two separate certificates upon completion

Course Bundle: Semantic Keyword Research + AI Search & LLM Entity SEO – 2 Courses

Course Bundle

~~600,00 €~~ **500,00 €**

Add to cart

**Join our
community
and academy
to learn faster.**

