

Test-time scaling

Reasoning models, retrieval-augmented generation, LLM agents

Zdeněk Kasner


 23 Apr 2026




Scaling laws

LLMs getting exponentially bigger

Until 2021, it seemed that the way to improve the models (lower their loss/perplexity on the training set) is **adding more parameters**:



This trend was supported by the research from OpenAI ([Kaplan et al., 2020](#)), who showed that larger models **learn more efficiently and can attain lower loss**:



According to Kaplan et al., test **loss** L can be predicted solely based on N (# of parameters), D (# of training tokens), and C (compute budget in FLOPs).

They derived the following **empirical laws**:

$$\begin{aligned} L(N) &= (N_c / N)^{\alpha_N}; & \alpha_N &\approx 0.076, & N_c &\approx 8.8 \times 10^{13} \text{ parameters} \\ L(D) &= (D_c / D)^{\alpha_D}; & \alpha_D &\approx 0.095, & D_c &\approx 5.4 \times 10^{13} \text{ tokens} \\ L(C) &= (C_c / C)^{\alpha_C}; & \alpha_C &\approx 0.050, & C_c &\approx 3.1 \times 10^8 \text{ PFLOP-days} \end{aligned}$$


According to Kaplan et al., test **loss** L can be predicted solely based on N (# of parameters), D (# of training tokens), and C (compute budget in FLOPs).


They derived the following **empirical laws**:

$$\begin{aligned} L(N) &= (N_c / N)^{\alpha_N}; & \alpha_N &\approx 0.076, & N_c &\approx 8.8 \times 10^{13} \text{ parameters} \\ L(D) &= (D_c / D)^{\alpha_D}; & \alpha_D &\approx 0.095, & D_c &\approx 5.4 \times 10^{13} \text{ tokens} \\ L(C) &= (C_c / C)^{\alpha_C}; & \alpha_C &\approx 0.050, & C_c &\approx 3.1 \times 10^8 \text{ PFLOP-days} \end{aligned}$$

Example consequence

Given a 10× increase in compute, Kaplan et al. suggest to increase model size $\approx$ 5.5× but data only $\approx$ 1.8×.





As the computational budget C increases, it should be spent primarily on larger models, without dramatic increases in training time or dataset size.

— Kaplan et al. 2020

→ These findings lead to large models like GPT-3 (175B) being trained on relatively limited data ($\approx 300\text{B}$ tokens).

Research from DeepMind ([Hoffmann et al., 2022](#)) **challenged the scaling laws.**

They trained on a wider range of N and D and found that they should be scaled **equally** (C is the compute budget):

$$N_{\text{opt}} \propto C^a, \quad D_{\text{opt}} \propto C^b, \quad a \approx b \approx 0.5$$

Research from DeepMind ([Hoffmann et al., 2022](#)) **challenged the scaling laws.**

They trained on a wider range of N and D and found that they should be scaled **equally** (C is the compute budget):

$$N_{\text{opt}} \propto C^a, \quad D_{\text{opt}} \propto C^b, \quad a \approx b \approx 0.5$$

Their Chinchilla (70B, 1.4T tokens) model outperformed another model Gopher (280B, 300B tokens), having 4× fewer parameters.

Research from DeepMind ([Hoffmann et al., 2022](#)) **challenged the scaling laws.**

They trained on a wider range of N and D and found that they should be scaled **equally** (C is the compute budget):


$$N_{\text{opt}} \propto C^a, \quad D_{\text{opt}} \propto C^b, \quad a \approx b \approx 0.5$$

Their Chinchilla (70B, 1.4T tokens) model outperformed another model Gopher (280B, 300B tokens), having 4× fewer parameters.




Chinchilla scaling law

For compute-optimal training, the number of training tokens should be $\approx 20\times$ the number of parameters.


→ According to their findings, most existing LLMs at the time (GPT-3, Gopher, Megatron) were **undertrained**:



It may be also worth training a **smaller model on more data** to reduce inference cost
([Sardana et al., 2024](#)):



→ This was embraced by the Llama and Mistral models: **overtrain smaller models** so that they are cheaper at inference.



...

GPT-4.5 is ready!

good news: it is the first model that feels like talking to a thoughtful person to me. i have had several moments where i've sat back in my chair and been astonished at getting actually good advice from an AI.

bad news: it is a giant, expensive model. we really wanted to launch it to plus and pro at the same time, but we've been growing a lot and are out of GPUs. we will add tens of thousands of GPUs next week and roll it out to the plus tier then. (hundreds of thousands coming soon, and i'm pretty sure y'all will use every one we can rack up.)

this isn't how we want to operate, but it's hard to perfectly predict growth surges that lead to GPU shortages.

a heads up: this isn't a reasoning model and won't crush benchmarks. it's a different kind of intelligence and there's a magic to it i haven't felt before. really excited for people to try it!

[Přeložit post](#)

9:05 odp. **27. 2. 2025** 5,5 mil. Zobrazení

source: <https://x.com/sama/status/1895203654103351462>

Model	Input Cost (per 1M tokens)	Output Cost (per 1M tokens)	Context Window	Comments
GPT-4.5	\$75.00	\$150.00	128k tokens	Premium pricing for advanced emotional and conversational capabilities
GPT-4o	\$2.50	\$10.00	128k tokens	Cost-effective baseline with fast, multimodal support

source: dev.to

2025-04-14: GPT-4.5-preview

On **April 14th, 2025**, we notified developers that the `gpt-4.5-preview` model is deprecated and will be removed from the API in the coming months.

Shutdown date	Model / system	Recommended replacement
2025-07-14	<code>gpt-4.5-preview</code>	<code>gpt-4.1</code>

source: [OpenAI](https://openai.com)

Scaling **pretraining** has its limits:

- We are running out of high-quality text data.
- Training compute costs are enormous (\$100M+ for frontier models).
- **Diminishing returns:** each 10× increase in compute yields smaller improvements.

Question

Can you think of another way to improve model performance?

Scaling **pretraining** has its limits:

- We are running out of high-quality text data.
- Training compute costs are enormous (\$100M+ for frontier models).
- **Diminishing returns**: each 10× increase in compute yields smaller improvements.

Question

Can you think of another way to improve model performance?

Idea

“Squeeze out” more from the models during inference → **test-time scaling**.

Chain-of-thought prompting

Idea

Wei et al. (2022): LLMs struggle with math and multi-step reasoning. What if we showed them **how to do intermediate reasoning steps**?

Idea

Wei et al. (2022): LLMs struggle with math and multi-step reasoning. What if we showed them **how to do intermediate reasoning steps**?

Standard Prompting

Model Input

Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?

A: The answer is 11.

Q: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?

Model Output

A: The answer is 27. ❌

Chain-of-Thought Prompting

Model Input


Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?

A: Roger started with 5 balls. 2 cans of 3 tennis balls each is 6 tennis balls. $5 + 6 = 11$. The answer is 11.

Q: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?

Model Output

A: The cafeteria had 23 apples originally. They used 20 to make lunch. So they had $23 - 20 = 3$. They bought 6 more apples, so they have $3 + 6 = 9$. The answer is 9. ✅




→ Dramatic improvement on arithmetic, commonsense, and symbolic reasoning.

Idea

Kojima et al. (2022): Writing out the steps is difficult. Do we need to do it at all?

Idea

Kojima et al. (2022): Writing out the steps is difficult. Do we need to do it at all?



(a) Few-shot

Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?

A: The answer is 11.

Q: A juggler can juggle 16 balls. Half of the balls are golf balls, and half of the golf balls are blue. How many blue golf balls are there?

A:

(Output) The answer is 8. ✗

(b) Few-shot-CoT

Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?

A: Roger started with 5 balls. 2 cans of 3 tennis balls each is 6 tennis balls. $5 + 6 = 11$. The answer is 11.

Q: A juggler can juggle 16 balls. Half of the balls are golf balls, and half of the golf balls are blue. How many blue golf balls are there?

A:

(Output) The juggler can juggle 16 balls. Half of the balls are golf balls. So there are $16 / 2 = 8$ golf balls. Half of the golf balls are blue. So there are $8 / 2 = 4$ blue golf balls. The answer is 4. ✓

(c) Zero-shot

Q: A juggler can juggle 16 balls. Half of the balls are golf balls, and half of the golf balls are blue. How many blue golf balls are there?

A: The answer (arabic numerals) is

(Output) 8 ✗

(d) Zero-shot-CoT (Ours)

Q: A juggler can juggle 16 balls. Half of the balls are golf balls, and half of the golf balls are blue. How many blue golf balls are there?

A: **Let's think step by step.**

(Output) There are 16 balls in total. Half of the balls are golf balls. That means that there are 8 golf balls. Half of the golf balls are blue. That means that there are 4 blue golf balls. ✓

→ Prompting with “Let’s think step by step” is enough to trigger reasoning behavior.

- Can be applied to any instruction-tuned model.
 - Does not require more than appending (a variant of) “Think step-by-step” to the prompt.
- Generally increases performance on problems requiring **multi-step reasoning**
 - Helps to break down complex problems into simpler subproblems.
- Nowadays falling out of favor compared to **large reasoning models** trained to reason explicitly.



Test-time scaling

Idea: test-time scaling


Can we make use of the reasoning capability to improve model performance without further training, solely at **inference time**?

Idea: test-time scaling


Can we make use of the reasoning capability to improve model performance without further training, solely at **inference time**?

The simplest example of **test-time scaling**:

- Generate multiple CoT paths for the given problem.
- Use majority voting to select the final answer.



More advanced extensions involve **tree search** and a verifying the solution with a **verifier model** (e.g., code interpreter):




Large reasoning models

Idea

What if we trained the model to **always produce the reasoning trace** before answering?

Idea

What if we trained the model to **always produce the reasoning trace** before answering?



We call such a model a **large reasoning model (LRM)**.

Question

Why not just use the chain-of-thought prompting?

Question

Why not just use the chain-of-thought prompting?

→ Chain-of-thought reasoning is **not a robust capability**. It is largely an artifact of training data (most likely code data, [Ma et al., 2023](#)).


→ It is also **too simple**. Ideally, we would like to automate test-time scaling. That means giving the model the ability to:

- Decompose the problem into subproblems.
- Follow multiple reasoning paths.
- Back-track from invalid paths.

DeepSeek-R1 ([DeepSeek-AI, 2025](#)): the first **open** reasoning model. 671B parameters, competitive with OpenAI o1.

Their recipe on how to build a strong reasoning model:

1. Take a **base model** (they took their deepseek-v3-base).
2. **Finetune** the model on a dataset of reasoning traces.
3. Improve the reasoning process using **reinforcement learning**.



source: [Jay Alammar \(2025\)](#)

Question


Where to get the dataset with reasoning traces for Step 2?

Question

Where to get the dataset with reasoning traces for Step 2?

Recipe of DeepSeek-R1-Zero:

1. Apply the model with a **regular CoT prompting** on difficult problems.
2. Automatically **verify the solutions**.
3. **Reward** the model for good solutions with using RL.



During RL training, the model **spontaneously** learns to re-evaluate its reasoning and try alternative approaches:

Question: If $a > 1$, then the sum of the real solutions of $\sqrt{a - \sqrt{a + x}} = x$ is equal to

Response: <think>

To solve the equation $\sqrt{a - \sqrt{a + x}} = x$, let's start by squaring both ...

$$\left(\sqrt{a - \sqrt{a + x}}\right)^2 = x^2 \implies a - \sqrt{a + x} = x^2.$$

Rearrange to isolate the inner square root term:

$$(a - x^2)^2 = a + x \implies a^2 - 2ax^2 + (x^2)^2 = a + x \implies x^4 - 2ax^2 - x + (a^2 - a) = 0$$

...

Wait, wait. Wait. That's an aha moment I can flag here.

Let's reevaluate this step-by-step to identify if the correct sum can be ...

We started with the equation:

$$\sqrt{a - \sqrt{a + x}} = x$$

First, let's square both sides:

$$a - \sqrt{a + x} = x^2 \implies \sqrt{a + x} = a - x^2$$

Next, I could square both sides again, treating the equation: ...

...

Table 3 | An interesting “aha moment” of an intermediate version of DeepSeek-R1-Zero. The model learns to rethink using an anthropomorphic tone. This is also an aha moment for us, allowing us to witness the power and beauty of reinforcement learning.


Question

Can we just use DeepSeek-R1-Zero as *the* reasoning model?

Pure RL training has issues:

- Early RL training is unstable and hard to get going.
- The model mixes languages and produces messy formatting.
- The model gets good at math/code but struggles with general tasks.

→ However, we can still use the model to **generate a dataset of reasoning traces** for the finetuning step.



Idea

Now we have a *strong* reasoning model. Can we use its outputs to get a high-quality dataset of reasoning traces?

Idea

Now we have a *strong* reasoning model. Can we use its outputs to get a high-quality dataset of reasoning traces?

Yes: we can **finetune a model directly on the reasoning traces** of DeepSeek-R1-671B.

→ This idea lead to **distilled models** based on Llama and Qwen (between 1.5B to 70B).

Idea

Now we have a *strong* reasoning model. Can we use its outputs to get a high-quality dataset of reasoning traces?

Yes: we can **finetune a model directly on the reasoning traces** of DeepSeek-R1-671B.

→ This idea lead to **distilled models** based on Llama and Qwen (between 1.5B to 70B).

Why distillation?

Distillation – in this context – means finetuning a smaller model (“student”) on the reasoning traces generated by a larger model (“teacher”).

Option 1: Pure RL

- Base LLM + RL with verifier rewards
- Emergent reasoning
- Unstable, poor readability

→ DeepSeek-R1-Zero

RL = reinforcement learning, SFT = supervised finetuning

Option 1: Pure RL

- Base LLM + RL with verifier rewards
 - Emergent reasoning
 - Unstable, poor readability
- DeepSeek-R1-Zero

Option 2: SFT + RL

- SFT on reasoning traces, then RL
 - Stable training
 - Expensive
- DeepSeek-R1, current frontier LRMs

RL = reinforcement learning, SFT = supervised finetuning

Option 1: Pure RL

- Base LLM + RL with verifier rewards
 - Emergent reasoning
 - Unstable, poor readability
- DeepSeek-R1-Zero

Option 2: SFT + RL

- SFT on reasoning traces, then RL
 - Stable training
 - Expensive
- DeepSeek-R1, current frontier LRMs



Option 3: Pure SFT

- SFT on reasoning traces distilled from the teacher model
 - Cheapest approach
 - Limited performance
- distilled LRMs

RL = reinforcement learning, SFT = supervised finetuning

Many properties of reasoning models are **yet to be properly investigated**, such as:

- How well does the trace reflect the internal thinking process?
- What is the optimal length of the trace? Can we enforce certain length?
- etc.



July 21, 2025 Research

Advanced version of Gemini with Deep Think officially achieves gold-medal standard at the International Mathematical Olympiad

Thang Luong and Edward Lockhart

OPENAI O3 BREAKTHROUGH HIGH SCORE ON ARC-AGI-PUB

OpenAI has released a new version of o3. [Read our analysis](#) to learn how it differs from the preview below.

Updated (April 16, 2025): OpenAI has [officially released o3](#). OpenAI has confirmed that this version is not the same as the one we tested in this original post. See [more information](#) on this. We will publish updated results for released o3 shortly.

OpenAI's new o3 system - trained on the ARC-AGI-1 Public Training set - has scored a breakthrough **75.7%** on the Semi-Private Evaluation set at our stated public leaderboard \$10k compute limit. A high-compute (172x) o3 configuration scored **87.5%**.



OpenAI
@OpenAI

...


We achieved gold medal-level performance 🏆 on the 2025 International Mathematical Olympiad with a general-purpose reasoning LLM!

Our model solved world-class math problems—at the level of top human contestants. A major milestone for AI and mathematics.

AIME 2024




Competition Math



AIME 2025



Competition Math



AI Progress on 🧠 Humanity's Last Exam.



Reasoning models are the main driver of the remarkable results of LLMs in mathematical competitions and other difficult benchmarks:

- **AIME 2024** (math olympiad difficulty-level problems): OpenAI o3 scored 91.6%.
- **International Mathematical Olympiad 2025**: both Google's Gemini and OpenAI models have started solving IMO problems at the gold-medal level.
- **ARC-AGI**: benchmark for abstract reasoning, even its successor (ARC-AGI-2) is now approaching saturation (83.3% for GPT-5.4).
- **Humanity's Last Exam (HLE)**: benchmark of questions that LLMs were not able to solve in 2024 → as of 03/26, the best models have above 40%

Reasoning models typically wrap their **thinking traces** in `<think>` tags:

User: How many r's are in "strawberry"?

`<think>`

Let me count the r's in "strawberry" letter by letter.

s-t-r-a-w-b-e-r-r-y

r at position 3 → count: 1

r at position 8 → count: 2

r at position 9 → count: 3

So there are 3 r's.

`</think>`

There are 3 r's in "strawberry".

The thinking part is typically hidden from the user in chat interfaces.

Open LLM APIs expose the reasoning content in a separate field.

Example: vLLM (https://docs.vllm.ai/en/latest/features/reasoning_outputs):

```
messages = [{"role": "user", "content": "9.11 and 9.8, which is greater?"}]
response = client.chat.completions.create(model=model, messages=messages)

reasoning = response.choices[0].message.reasoning
content = response.choices[0].message.content
```

⚠ Warning

Full reasoning traces are typically **not** available with commercial models.

For example, Google only offers so called “thinking summaries”: <https://ai.google.dev/gemini-api/docs/thinking>

Now there is quite a lot of open LRMs to choose from:

Model Series	Parser Name	Structured Output Support	Tool Calling
DeepSeek R1 series 	deepseek_r1	json, regex	✗
DeepSeek-V3.1 	deepseek_v3	json, regex	✗
ERNIE-4.5-VL series 	ernie45	json, regex	✗
ERNIE-4.5-21B-A3B-Thinking 	ernie45	json, regex	✓
GLM-4.5 series 	glm45	json, regex	✓
Holo2 series 	holo2	json, regex	✓
Hunyuan A13B series 	hunyuan_a13b	json, regex	✓
IBM Granite 3.2 language models 	granite	✗	✗
MiniMax-M2 	minimax_m2_append_think	json, regex	✓
Qwen3 series 	qwen3	json, regex	✓
QwQ-32B 	deepseek_r1	json, regex	✓

Makes sense:

- Math and logic problems
- Complex coding tasks
- Multi-step planning

A bit of an overkill:

- Simple Q&A
- Factual retrieval
- Text summarization

⚠ Why not use reasoning models for all the tasks

Reasoning models use **significantly more tokens** (→ cost and time), possibly up to 10–100× the output length of a regular instruction-tuned model.

Retrieval-augmented generation

LLM knowledge is **static, lossy and incomplete**.

LLM knowledge is **static, lossy and incomplete**.

Question

When this can become a problem?

When we need the model to access:


- Up-to-date information (news, weather, timetables, ...).
- Domain-specific knowledge (internal company docs, law / medical / ...).

We may also want to access external knowledge **ourselves** to check the LLM's answers.

Idea


Combine the LLM with a **retrieval system**: first retrieve relevant documents, then use them as context for generation → “retrieval augmented generation”

The idea was first introduced by [Lewis et al. \(2020\)](#):




Phase #1: Indexing

1. Split the documents into chunks.
2. Compute their embeddings.
3. Store them in a vector database.




Phase #1: Indexing

1. Split the documents into chunks.
2. Compute their embeddings.
3. Store them in a vector database.



Phase #2: Querying


1. Embed the user prompt.
2. Retrieve relevant chunks.
3. Concatenate them with the prompt.
4. Generate an answer.



Indexing – Step 1: Document chunking

Raw documents need to be split into chunks before indexing:

- **Fixed-size chunks:** split every N tokens (simple, but can break context).
- **Semantic chunking:** split at sentence/ paragraph/section boundaries.
- **Sliding-window with overlaps:** fixed-size chunks with some overlap to preserve context.




[source: Medium.com](https://medium.com)

Good starting point: ~256–1024 tokens, with 10–20% overlap.

Too small → loss of context. Too large → diluted relevance and noise.

Each chunk is converted to a dense vector (embedding).



Embedding models:


- encoder-based: finetuned BERT-like models, see the [SentenceTransformers library](#).
- decoder-based: OpenAI's text-embedding-3, e5, qwen3-embedding, ...

Decoder embedding models perform much worse if not **prompted** properly.



```
# Each query must come with a one-sentence
instruction that describes the task
task = 'Given a web search query, retrieve
relevant passages that answer the query'
```

source: <https://huggingface.co/Qwen/Qwen3-Embedding-8B>


Note there is no actual retrieval →
we just **“prime” the model** to build a
good representation of the query.



The created embeddings are stored in a **vector database** (Qdrant, Chroma, FAISS, ...):



At **query time**, we embed the user query and find the most similar chunks:



Similarity is typically measured using **cosine similarity** or **dot product** between the query and document embeddings.

The retrieved chunks are inserted into the prompt as **context**:

System: You are a helpful assistant. Answer based on the provided context.

Context:

[Chunk 1]: "The Transformer architecture was introduced by Vaswani et al. in 2017..."

[Chunk 2]: "Self-attention allows the model to weigh different positions..."


User: What is the key innovation of the Transformer?

Question

Can a RAG-based answer be *worse* than without RAG?

The basic pipeline can be improved at every step:

- We can apply **query rewriting** to rephrase the query for better retrieval.
- We can use **hybrid search** – combine dense retrieval with keyword-based search (e.g. BM25).
- We can **rerank** the retrieved results e.g. using a cross-encoder (heavier, but more precise).
- We can retrieve the documents **iteratively** (→ heading towards LLM agents).



RAG is now a standard for building knowledge-intensive LLM applications (in combination with LLM agents):

- **Company internal knowledge:** “talk to your knowledge base”.
- **Customer support bots** with access to the internal knowledge bases.
- **Code assistants** that retrieve from documentation and codebase.
- **Legal / medical AI:** retrieve from domain-specific corpora.

There are now many frameworks that facilitate RAG:



(...and many more).


Tool calling

RAG works as long as the information can be fetched based on the user query.

But what if LLMs need to access external resources **during generation**?

Idea

We give the LLM the capability to generate tool calls.



Toolformer (Schick et al. 2023): one of the most influential **LLM tool-calling** papers.

- An LLM finetuned on a corpus containing **examples of tool calls**.
- The corpus built in a self-supervised way:
 - Another (few-shot prompted) LLM suggests tool calls.
 - The suggestions that reduce LM loss are inserted in the corpus.
- **Tools:** QA system, Wikipedia search, calculator, calendar, MT system

The New England Journal of Medicine is a registered trademark of [QA("Who is the publisher of The New England Journal of Medicine?") → Massachusetts Medical Society] the MMS.


Out of 1400 participants, 400 (or [Calculator(400 / 1400) → 0.29] 29%) passed the test.

The name derives from "la tortuga", the Spanish word for [MT("tortuga") → turtle] turtle.


The Brown Act is California's law [WikiSearch("Brown Act") → The Ralph M. Brown Act is an act of the California State Legislature that guarantees the public's right to attend and participate in meetings of local legislative bodies.] that requires legislative bodies, like city councils, to hold their meetings open to the public.

Once the model generates the tool call:

1. the **output is evaluated** using an external tool,
2. the **result is appended** to the generated text.



There are many tools that a model can potentially call → describing the interface of each of them can be cumbersome.



Model Context Protocol (MCP): standardized way to access tools based on JSON.

Example tool specification


```
{
  "name": "get_weather",
  "description": "Get current weather for a city",
  "inputSchema": {
    "type": "object",
    "properties": {
      "city": {"type": "string", "description": "City name"}},
    "required": ["city"]}
}
```

Example: what the model returns


```
{"name": "get_weather", "arguments": {"city": "Prague"}}
```

MCP describes not just the JSON format, but the entire tool-calling architecture:

- **MCP client:** the application managing the LLM (VSCode, Claude Code, ...)
- **MCP server:** the application that can be managed with tools (Github, database, ...)



MCP tool descriptions can use up a lot of input tokens:




→ This is currently an open issue.

Can be mitigated with [dynamic content discovery](#) or replacing MCP with [skills](#) (.md files that include description of the tool's CLI) instead.


LLM agents

So far, we considered reasoning and acting (=tool calling) to be separate processes:




ReAct: [Yao et al. \(2023\)](#): combining **reasoning** and **acting** in a single loop.

→ A base paradigm for what is now called “LLM agents.”



The agent in ReAct iterates through the **thought** → **action** → **observation** loop until it reaches a final answer:

- **Thought**: reasoning about the current situation.
- **Action**: executing a tool.
- **Observation**: observing the result of the action.




What is an LLM agent?

source: [Maarten Grootendorst's blog](#)

To a degree, the LLM agents follow the classical AI agent definitions.

They can **observe** the environment (→input), **plan** (→reasoning), and **act** (→tools):




Simon Willison's Weblog

I was going slightly spare at the fact that every talk at this Anthropic developer conference has used the word "agents" dozens of times, but nobody ever stopped to provide a useful definition.


I'm now in the "Prompting for Agents" workshop and Anthropic's Hannah Moran finally broke the trend by saying that at Anthropic:


Agents are models using tools in a loop

I can live with that! I'm glad someone finally said it out loud.




Multi-agent systems: multiple specialized agents, each with their own tools and memory, coordinated by a supervisor:






source: <https://openclaw.ai>



source: <https://a16z.com/100-gen-ai-apps-6/>

OpenClaw can integrate any LLM with various applications (Whatsapp, Spotify, Gmail, ...) via **MCPs**, making agents very useful and super dangerous at the same time [?]



source: <https://x.com/JFPuget/status/2025877071939911791/photo/1>

Summary

- **Scaling laws:** LLM performance follows predictable empirical laws based on model size, data size, and compute.
- **Chain-of-thought:** prompting with intermediate reasoning steps dramatically improves multi-step reasoning.
- **Test-time scaling:** using more compute at inference can be more efficient than scaling pretraining.
- **Large reasoning models** (o1/o3, DeepSeek-R1, ...): train models to produce reasoning traces using RL and/or SFT.
- Reasoning abilities are **rapidly improving**: math olympiad results, ARC-AGI saturation.

- **RAG**: retrieve relevant documents and use them as context for generation. The de facto standard for knowledge-intensive LLM apps.
- **Tool calling**: LLMs can use external tools (APIs, code, search). Pioneered by Toolformer, now a standard feature.
- **MCP**: an open protocol for standardizing tool access across LLM applications.
- **Structured outputs**: constrained decoding forces valid output formats (JSON, regex, CFG) during generation.
- **Agents**: Reasoning LLMs equipped with tools, applying these in a loop to achieve a goal.

- [Kaplan et al. \(2020\): Scaling Laws for Neural Language Models](#)
- [Hoffmann et al. \(2022\): Training Compute-Optimal Large Language Models \(Chinchilla\)](#)
- [Sardana et al. \(2024\): Beyond Chinchilla-Optimal](#)
- [Wei et al. \(2022\): Chain-of-Thought Prompting](#)
- [Kojima et al. \(2022\): Zero-shot CoT](#)
- [Wang et al. \(2022\): Self-Consistency](#)
- [Snell et al. \(2024\): Scaling LLM Test-Time Compute](#)
- [DeepSeek-AI \(2025\): DeepSeek-R1](#)
- [Jay Alammar: The Illustrated DeepSeek-R1](#)
- [Sebastian Raschka: Understanding Reasoning LLMs](#)
- [Maarten Grootendorst: A Visual Guide to Reasoning LLMs](#)

- [Lewis et al. \(2020\): Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks](#)
- [Gao et al. \(2024\): Retrieval-Augmented Generation for Large Language Models: A Survey](#)
- [Schick et al. \(2023\): Toolformer: Language Models Can Teach Themselves to Use Tools](#)
- [Anthropic: Model Context Protocol \(MCP\)](#)
- [Willard & Louf \(2023\): Efficient Guided Generation for Large Language Models \(Outlines\)](#)
- [OpenAI: Structured Outputs in the API](#)
- [Yao et al. \(2023\): ReAct: Synergizing Reasoning and Acting in Language Models](#)
- [Maarten Grootendorst's blog: A Visual Guide to LLM Agents](#)