Become a GPT Prompt Maestro

David Shapiro · Follow

13 min read · Sep 5

If there isn't an anime about Beethoven there needs to be. Cause this symphony looks fire. 🔥

TL;DR

There are 3 kinds of prompts:

1. Reduction
2. Transformation

3. Generation/Expansion

You also need to know:

1. Bloom's Taxonomy
2. Latent content
3. Emergent Capabilities
4. Hallucination is a feature, not a bug

If you know all that, you don't need to read this article. Otherwise, read on!

Unlocking the Potential of Large Language Models

Large language models (LLMs) like GPT-4 and Claude have captured the imagination of technologists and the general public alike. Their ability to generate humanlike text and engage in conversations seems like something out of science fiction. However, as with any new technology, there is still much confusion and debate around how LLMs actually work under the hood.

In this article, I aim to provide a high-level taxonomy of the key capabilities of LLMs to clarify what they can and cannot do. My goal is to explain the current state of LLMs in an accessible way for nonexperts, while also identifying areas for further research and development. Ultimately, I believe LLMs have enormous potential to augment human intelligence if steered in ethically responsible ways.

What is an LLM?

First, what exactly is a large language model? At the most basic level, an LLM is a deep learning neural network trained on massive datasets of text data like books, websites, and social media posts. The "large" refers to the fact these models have billions of parameters, allowing them to build very complex statistical representations of language.

The key task an LLM is trained on is predicting the next word or token in a sequence given the previous context. So if it sees the text "The cat jumped over the…", it learns to predict "fence" as the likely next token. Doing this over and over teaches the LLM implicit knowledge about how language works and the relationships between words

and concepts.

This training process, along with the massive datasets used, embed a surprising amount of world knowledge into LLMs like Claude and GPT-4. However, it's important to understand LLMs have no explicit knowledge or hand-coded rules all their capabilities are emergent results of recognizing patterns in their training data.

Fundamental LLM Operations

At a high level, LLMs have three primary modes of operation:

1. **Reduction Operations:** Go from big to small. >
2. **Transformation Operations:** Maintain size and/or meaning. ≈
3. **Generation (or Expansion) Operations:** Go from small to big. <

Let's dive in.

Reductive Operations

LLMs exhibit strong capabilities when it comes to reductive operations. These involve taking a large passage or document as input and condensing it down into a much more concise output. Reductive tasks leverage the language modeling strengths of LLMs to identify and extract the most salient information.

One common reductive operation is summarization. When given a lengthy input text, the LLM can produce a concise summary covering just the key points. This is done by analyzing the document to look for the main topics, events, and facts described. The model then attempts to synthesize these elements into a short summary conveying the core essence of the full document. LLMs are generally quite skilled at basic summarization within a limited length, removing extraneous details while preserving semantic meaning.

Related to summarization is the task of distillation. This goes beyond just shortening a document to extracting and refining its fundamental principles, findings, or facts. Distillation aims to filter out noise and redundancy from the input, purifying just the core knowledge or claims. For scientific documents, an LLM might identify and synthesize the key hypotheses, results, and conclusions of an experiment.

Distillation requires deeper understanding to separate peripheral content from central assertions.

Extraction is another reductive technique leveraged by LLMs. This involves scanning through texts and pulling out specific pieces of information. For example, an LLM can read a document and extract just names, dates, figures, or other targeted data. Extraction underpins question answering, an area where LLMs excel. When asked for particular details from a passage, an LLM can often accurately retrieve the requested information.

Overall, reductive operations play directly into the strengths of large language models. Their statistical learning equips them to identify and convey the most salient parts of verbose input texts. Summarization, distillation, extraction and similar techniques will only grow more robust as LLMs continue to advance.

- **Summarization** — Say the same thing with fewer words. Can use list, notes, executive summary.
- **Distillation** — Purify the underlying principles or facts. Remove all the noise, extract axioms, foundations, etc.
- **Extraction** — Retrieve specific kinds of information. Question answering, listing names, extracting dates, etc.
- **Characterizing** — Describe the content of the text. Describe either the text as a whole, or within the subject.
- **Analyzing** — Find patterns or evaluate against a framework. Structural analysis, rhetorical analysis, etc
- **Evaluation** — Measuring, grading, or judging the content. Grading papers, evaluating against morals
- **Critiquing** — Provide feedback within the context of the text. Provide recommendations for improvement

It's a trash compactor. Get it? Compression? Compaction?

Transformative Operations

In contrast to reductive tasks, transformational operations aim to reshape or reframe input texts without significantly shrinking or expanding them. LLMs exhibit strong capabilities for language transformations that recast content into new formats and styles while preserving overall meaning.

One common transformational technique is reformatting — altering the presentation of text without changing its informational content. For example, an LLM can readily convert prose into screenplay dialogue, translate blog posts into Tweets, or shift a passage from active to passive voice. Reformatting leverages models' structural understanding of genres and linguistic conventions.

Translation between natural languages represents another key transformational capability of LLMs. Given input text in one tongue, they can rewrite it into another language by replacing the vocabulary and grammar while maintaining semantic consistency. Translation quality varies across language pairs but continues to improve as more multilingual data is utilized for training.

Paraphrasing also falls under the scope of transformational operations. Here the LLM aims to rewrite the input text using completely different words and phrasing while conveying the same essential meaning. This tests the model's ability to generate multiple syntactic variations that are logically equivalent. Paraphrasing has applications in plagiarism detection and improving clarity.

Lastly, restructuring content for better flow and organization qualifies as a transformation. LLMs can rearrange passages to enhance logical connections, present ideas more sequentially, or improve text for readability. Their training gives models a sense of how to craft coherent narratives and arguments.

Taken together, transformational capabilities allow LLMs to remix and present text in new ways that meet different needs. These techniques are empowering for tasks like tailoring content for specific audiences and overcoming language barriers. LLMs are already proficient at many transformations and are only becoming more skilled.

- **Reformatting** — Change the presentation only. Prose to screenplay, XML to JSON.
- **Refactoring** — Achieve same results with more efficiency. Say the same exact thing, but differently.
- **Language Change** — Translate between languages. English to Russian, C++ to Python.
- **Restructuring** — Optimize structure for logical flow, etc. Change order, add or remove structure.
- **Modification** — Rewrite copy to achieve different intention. Change tone, formality, diplomacy, style, etc.
- **Clarification** — Make something more comprehensible. Embellish or more clearly articulate.

This autobot transformers into a DISCO FEVER BOT every Friday night!

Generative (Expansion) Operations

Whereas reductive and transformative tasks involve manipulating existing text, generative operations synthesize completely new content from scratch. This places greater demand on the creative capacities of LLMs, with more variable and context-dependent results. Nonetheless, their skills at generative writing continue to evolve rapidly.

One major generative application is drafting original documents like stories, articles, code, or legal documents from high-level prompts. After ingesting the initial instructions, the LLM composer works to expand this seed content into a coherent draft exhibiting topicality, logical flow, and stylistic fluency. Results can be rough around the edges, but models like GPT-4 can produce impressive first drafts suitable for human refinement.

Given a set of parameters or design goals, LLMs are also adept at generating plans and steps to accomplish objectives. This planning capability is enabled by their capacity to infer sequences of logical actions that causally lead to desired outcomes, informed by their training data. Planning also draws on their latent knowledge

about how the world works. An LLM can propose plans ranging from kitchen recipes to software workflows.

More open-ended generative abilities of LLMs include brainstorming and ideation. When presented with a prompt or creative brief, models can provide lists of possibilities, concepts, and imagined solutions that humans can then curate. Their statistical associations between words enable models to branch in unexpected directions when brainstorming. The most promising ideas generated by an LLM prompt can then be selected for further development.

Lastly, LLMs exhibit strong generative capacity for amplification or elaboration of existing text. When given a seed passage or document, they are skilled at expanding upon the content with additional relevant details and explanation. This allows concise text to be extended organically by drawing upon the model's latent knowledge. Such amplification fleshes out skeletal text in creative ways.

Together these highlight how, given the proper context, LLMs can synthesize large volumes of new text from small amounts of input. Generative writing comes less naturally than reductive or transformative tasks but represents an area of active research with promising results emerging. Guiding LLMs to generate responsible and enriching content will be an ongoing design challenge.

- **Drafting** — Generate a draft of some kind of document. Code, fiction, legal copy, KB, science, storytelling.
- **Planning** — Given parameters, come up with plans. Actions, projects, objectives, missions, constraints, context.
- **Brainstorming** — Use imagine to list out possibilities. Ideation, exploration of possibilities, problem solving, hypothesizing.
- **Amplification** — Articulate and explicate something further. Expanding and expounding, riffing on stuff.

Runaway fissile events are examples of rapid expansion events. Intelligence explosion. Get it?

Bloom's Taxonomy

Bloom's Taxonomy is a classic educational framework outlining six cognitive skill levels for learning. First proposed in the 1950s, it provides a hierarchy of competencies from basic to advanced: Remembering, Understanding, Applying, Analyzing, Evaluating, and Creating. Examining LLMs through the lens of Bloom's Taxonomy highlights their multifaceted capabilities.

At the most fundamental level, LLMs excel at remembering and retrieving factual knowledge contained in their training datasets. Models like GPT-4 have 'read' vastly more text than any human could digest in a lifetime. This allows them to regurgitate information on nearly any topic when prompted. Their statistical learning acts as a massive knowledge repository that can be queried.

LLMs also demonstrate strong competency at understanding concepts, making connections between words and meanings. Their contextual learning enables robust recognition of everything from abstract philosophy to advanced physics. Even for complex subject matter not directly covered in training data, LLMs can

rapidly infer meaning through context and explanation.

Applying knowledge in new situations is also well within the purview of LLMs. After all, their entire purpose is to generate useful applications of language, whether writing, translation, or dialogue. Correct usage and contextual adaptation are essential for models to perform any task accurately. Without skillful application, they would have no practical utility.

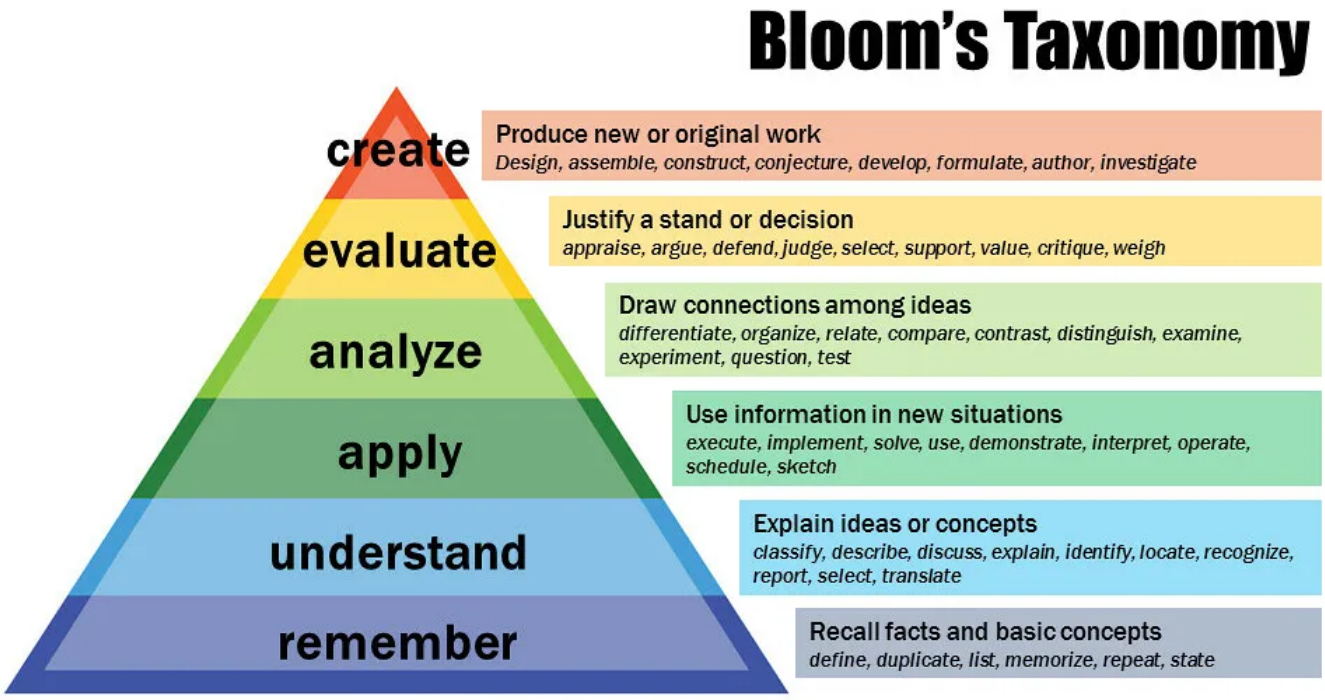

I freaking love frameworks and taxonomies. But seriously, look at each of these levels and tell me, with a straight face, that GPT cannot do them. Don't believe me? Try it yourself.

Higher on Bloom's taxonomy, analyzing information by drawing connections represents another LLM strength. Tools like Claude already analyze text across many dimensions for things like structure, style, and argument cohesion. With the right framing, LLMs can critically break down nearly any passage using learned cognitive capacities.

LLMs are also adept at evaluating and judging content when given appropriate criteria. Models can readily characterize texts across frameworks like reading level, target audience, grammaticality, reasoning, and more. More advanced LLMs may even judiciously critique the ethics of potential actions under consideration.

Finally, LLMs demonstrate skill at the highest level of Bloom's taxonomy: creating original content. While generation requires more careful prompting than reductive tasks, models can produce synthetic stories, articles, dialogue, and other creative works. Their emergent capabilities give LLMs great generative potential with proper activation.

In summary, modern LLMs fulfill every level of Bloom's Taxonomy to a remarkable degree. Their combination of immense latent knowledge and learned cognitive skills empower diverse applications from remembering facts to imaginative creation. As LLMs continue evolving, we can expect their Bloom abilities to become even more robust and multifaceted.

Latent Content

One of the most remarkable aspects of large language models is their ability to exhibit knowledge and reasoning abilities not explicitly programmed into them. This arises from the vast latent knowledge accumulated in the parameters of models like GPT-4 through the predictive training process.

Latent knowledge embedded in LLMs can be loosely grouped into three categories:

Training Data — The sheer volume of text consumed during training endows models with factual knowledge on countless topics. For example, Claude has latently encoded information about history, science, literature, current events, and more based on its training corpus. This functions as a massive knowledge base that can be queried with the right prompts.

World Knowledge — Beyond specific facts, LLMs also accumulate more general world knowledge about how things work. Their exposure to diverse contexts allows models to learn unspoken assumptions about culture, physics, causality, and human behavior. This enables intuitive reasoning about everyday situations.

Learned Cognitive Skills — The prediction-based learning approach also instills models with latent abilities like summarization, translation, and open-domain question answering. These emerge indirectly from the self-supervised objective rather than hardcoded rules.

This reservoir of latent knowledge is a game-changer for AI. However, directing and extracting this knowledge remains challenging. Often, the right prompts or techniques are needed to activate relevant parts of the model. Metaphorically, latent knowledge forms a dense forest that must be navigated skillfully.

While promising, reliance on latent knowledge also highlights current LLM limitations. Their reasoning can be highly dependent on priming using human intuition about what knowledge is required. More advanced techniques will be needed for models to better learn, index, and activate their own latent knowledge.

Overall, the scope of knowledge and skills accumulated latently by models like GPT-4 is deeply impressive. Methods to interpret, organize, and selectively utilize this reservoir will be an important area of research as LLMs continue to evolve. Latent knowledge unlocks possibilities far beyond what can be hard-coded in rules.

Latent space is like sunken treasure. It's out there, you just have to dive down to get it.

Emergent Capabilities

In the largest language models developed to date, new capabilities have emerged that go beyond what was explicitly contained in their training data. These higher-

order abilities arise from the interplay between massive model scale, extensive data, and the prediction-based learning approach.

Four examples of emergent capabilities are:

Theory of Mind — LLMs have exhibited some capacity to recognize differing perspectives between themselves and others. Models like Claude can adjust their tone and style based on conversational context, appear to grasp confusion, and distinguish their own knowledge from human knowledge. These hints of a "theory of mind" likely emerge from modeling countless social dialog exchanges. See? Reddit is good for something. LLMs learned to understand human minds by reading the comments section. What could possibly go wrong...

Implied Cognition — The ability of LLMs to "think" before generating each token implies latent cognitive capabilities not directly present in training data. As models predict the next word, they appear to perform dynamic reasoning, abstraction, and inference. Accurately modeling chains of causality requires cognitive processes like induction, deduction, and analogy formation.

Logical Reasoning — LLMs have also demonstrated some skill at deductive and inductive reasoning to make inferences based on presented information. Their statistical learning enables drawing connections between concepts and generalizing abstract principles. While still limited, goal-directed reasoning appears to be an emergent byproduct of modeling chains of causality in text.

In-Context Learning — Large models exhibit an ability to uptake new information and skills by incorporating context into their predictions. Without being explicitly trained, they can use information and instructions not seen in the original training data. This rapid in-context acquisition of knowledge and abilities was not directly built in. In humans, we would call this "improvisation" which is a hallmark of high intelligence.

These fledgling abilities spring from recognizing intricate patterns in human discourse, not hand-coded rules. They hint at how future LLMs might transition from pattern recognition to deeper reasoning, imagination, and causal understanding. However, significant limitations remain that require further

research and development.

Like Leviathans, scary new LLM abilities emerge from the depths of training data and parameter counts

Creativity and Hallucination

The capacity of LLMs to fabricate plausible-sounding statements may seem like a bug, but it actually represents a core feature of intelligence. Just as humans evolved imagination and creativity despite risks like delusions, so too must AI systems develop generative abilities alongside precautions.

Humans demonstrate a continuum between creativity and hallucination stemming from the same neurological source — spontaneous pattern generation. Cave art combining animal features into novel creatures first exercised this muscle. Unchecked, it can manifest in some mental disorders when imagination overtakes reality. LLMs exhibit a similar spectrum of speculative generation essential for intelligence.

Complete suppression of unpredictable “hallucinations” would also eliminate creative potential. The ideal is not eliminating but responsibly channeling generation. Research into alignment, ethics, and societal benefit will allow AI creativity to flourish.

Mitigating risks involves maintaining a grounding in facts and reality. Add real world data from out-of-band systems. Gentle nudging of LLMs towards corroborating their statements can keep them tethered to truth. Checking for citations or data provides crucial friction against unbridled speculation.

Equally important is transparently conveying confidence levels in generated text. When imagination overtakes observation, LLMs should indicate uncertainty to maintain trust. Technologies like verifiability scoring can help quantify speculation versus factual knowledge.

Overall, responsible development of AI should embrace, not reject, capacities like speculation and metaphor. With care, these can enhance human creativity and problem-solving instead of sowing misinformation. The solution is recognizing the continuity between creativity and hallucination, then thoughtfully cultivating the former while minimizing harmful forms of the latter.

99% of human creativity (scientific and artistic) is just remixing stuff we already know and experience. Just like how Midjourney can remix psychedelic art with Captain Jack Sparrow. See? There's a reason for the Caribbean theme today!

Conclusion

Large language models represent a technological breakthrough that opens up promising new pathways to augment human intelligence. However, fully realizing their potential will require deep study into how they actually work along with ethical precautions.

This article aimed to provide an accessible taxonomy of LLMs' fundamental capabilities and limitations. Reductive, transformative, and generative operations draw on different strengths of current models. Latent knowledge is powerful but heavily dependent on prompts for activation. And emergent properties like reasoning show promise for future systems.

There are certainly risks from misuse of LLMs that should not be downplayed. But rather than reject them as too dangerous or uncontrollable, the responsible way forward is to direct research and development toward beneficial applications that truly enhance human capabilities and creativity. If steered with care in coming years, LLMs could help unlock new frontiers for humanity.

Gpt Prompt Engineering Blooms Taxonomy Llm ChatGPT

Follow

Written by David Shapiro

125 Followers

AI Thought Leader | YouTuber