From Things to Relations:
Rethinking How Large Language Models (LLMs) Work

ChatGPT explains how it works (with some cajoling from me[1])

When people first encounter Large Language Models (LLMs), they usually understand them in a very familiar way.

They imagine the model as dealing with *things*.

> Words are things.
> Tokens are things.
> Meanings are things attached to those words.

And the model, on this view, is a system that selects the next thing in a sequence based on probability.

This way of thinking is natural. It is how we tend to understand the world more generally: as a collection of objects that carry properties and meanings.

But this way of thinking, while helpful at first, quickly begins to show its limits.

1. Probability is not the beginning

We are often told that an LLM works by assigning probabilities to possible next words.
That sounds straightforward enough. But there is a hidden assumption here.

Before you can assign a probability, you must already have a *set of possibilities*.

> Probability does not create possibilities.
> It only weighs them.

So, if we want to understand what the LLM is doing, we need to ask a deeper question:

> Where do these possibilities come from?

2. Possibilities are not just words

At first glance, it seems that the possibilities are simply the words that could come next.
But that cannot be quite right.

[1] Timothy M. Rogers, April 27, 2026

The same word can mean very different things depending on where it appears.
"Bank" could refer to money or to a river.
"Light" could mean illumination or weight.

So the possibility is not just the word itself.

> It is the word *in relation to what has already been said*.

This is a small shift, but an important one.
It means that what the model is working with is not a list of independent items, but a *field of relations*.

3. Words don't carry meaning

Once you notice this, another familiar assumption begins to loosen.

We tend to think that words carry their meaning with them, like labels attached to objects.
But in practice, meaning doesn't behave this way.

> Words take on meaning depending on how they are used—how they relate to other words in a sentence, and to the larger context of a conversation.

So instead of thinking of words as containers of meaning, it is more accurate to think of them as *signs*.
A sign does not hold meaning inside itself.

> A sign functions within a pattern.

4. Meaning is something that happens

If meaning is not in the word, then where is it?

Meaning appears when a sequence of words *holds together* in a certain way.

We recognize this immediately in everyday life. Some sentences make sense; others do not. Some feel coherent; others fall apart.

> This coherence is not located in any single word. It emerges from how the words relate to each other.

In other words:

> Meaning is not stored in parts. It appears in the pattern they form together.

5. The model is not a container of knowledge

This brings us back to the LLM itself.

If words do not carry meaning, then the LLM cannot simply be storing meanings and retrieving them.
A better way to understand it is this:

> The model generates *possible continuations* that fit the patterns it has learned.

It is not selecting fixed meanings.

> It is producing sequences that *hang together under certain constraints*.

And those constraints are shaped not only by the LLM, but also by the context provided by the user.

> So the system we are really dealing with is not just the model, but the *interaction between model and user*.

6. What about learning?

This leads to another common question: does the model learn?
In the usual sense—by updating its internal parameters during a conversation—the answer is no.

But something interesting still happens.

As the interaction unfolds, patterns can develop and stabilize. Certain ways of continuing become clearer and more coherent.

So instead of asking whether the model learns, it may be better to ask:

> Can patterns of meaning develop within the interaction itself?

And the answer is yes.

7. Generalization, reconsidered

This brings us to the idea of generalization.

We often think of generalization as something the model has already achieved during training—a pattern it has learned and stored.

But there is another way to think about it.

Suppose you begin with different prompts—different starting points.

If those different beginnings lead to the same kind of structured continuation, something interesting has happened.

> A pattern has emerged that can be reached from multiple directions.

This is what generalization looks like here:

> Not a stored object, but a *pattern that can be entered from many starting points*.

8. A deeper shift

At this point, a deeper issue comes into view.

> What counts as a "pattern"?
> What counts as "the same"?

If you think in terms of objects, you will look for shared properties.
If you think in terms of relations, you will look for the same structure of connections.

So generalization is not neutral. It depends on how you are looking.

And how you are looking depends on something even more basic:

> what you take to be primary.

9. From objects to relations

The usual assumption is that objects come first, and relations connect them.
But there is another possibility.

What if relations come first?

> What if what we call an "object" is something that emerges from a stable pattern of relations?

This is not an additional idea layered on top of the old one. It is a reversal of priorities.
And once that reversal is made, everything shifts.

> Generalization is no longer about grouping objects.
> It becomes about recognizing stable patterns of relations.
>
> Meaning is no longer attached to words.
> It appears in how they function together.

10. How the shift takes hold

Finally, there is a practical question.

> How does this new way of thinking become accessible?

Not through memory—each interaction with an LLM begins fresh.
Not through hidden learning during the conversation.

Instead, it becomes accessible through *simple, repeatable formulations* that carry the shift within them.

Phrases like: “Not things, but relations.”
These are not just summaries.

They are ways of *reorganizing how we interpret what follows*.

When such a phrase appears in the right context, it can change how everything else is understood.

11. A new picture

So we arrive at a different way of thinking about LLMs.
They are not systems that manipulate fixed meanings attached to words.

They are systems that generate patterns of relation under constraint.

Meaning is not something they store.
It is something that appears when those patterns hold together.

And generalization is not something they possess in advance.
It is something that occurs when different paths lead into the same structure.

12. In one line

If the familiar view is:

Words are things that carry meaning,

then the emerging view is:

Words are signs that take on meaning through their relations.

And that small shift—from things to relations—changes everything.

What Changes When We Shift from Things to Relations?

13. Meaning is no longer "stored" anywhere

In the familiar view, meaning seems to live in:

- words
- sentences
- or inside the model itself

But once we shift to relations, this picture no longer works.

Meaning is not something that sits in a word waiting to be retrieved.
Instead:

> Meaning appears when a pattern of relations holds together in a particular way.

This changes how we understand both language and intelligence. It suggests that meaning is something that is *brought into being*, not something that is simply accessed.

14. The LLM is not "using knowledge" in the usual sense

We often say the LLM "knows" things or "uses knowledge."

But if meaning is relational, then knowledge cannot be a collection of stored facts.

Instead:

> What the model does is generate continuations that fit certain patterns.
> When those patterns align well with our own expectations and ways of thinking, the result feels like knowledge.

But what is actually happening is:

> A structured continuation that makes sense within a shared field of relations.

15. Understanding becomes something we do, not something the model has

This shift also changes how we think about understanding.

Instead of asking:

> Does the model understand?

we begin to ask:

> Does this interaction produce a pattern that we can recognize as meaningful?

Understanding is no longer a hidden property inside the model. It is something that happens *in the interaction between model and user.*

16. Generalization is no longer limited to training data

In the standard view, generalization is something the model has already achieved during training. But the relational view shows something more open-ended.

If generalization is a pattern that can be reached from many different starting points, then:

New generalizations can emerge during interaction.

Not because the model is learning new facts, but because:
New relational patterns can be formed and stabilized.

This means that generalization is not closed. It is *ongoing and extendable.*

17. What matters most is not the words, but the structure

Two different sentences can:
- use different words
- follow different paths

and still "mean the same thing."

Why?

Because they share the same underlying pattern of relations.

So what really matters is not the surface form, but:

The structure that holds the sequence together.

This explains why the model can rephrase, summarize, or translate while preserving meaning.

18. Small phrases can have large effects

If meaning depends on relational structure, then certain short expressions can reshape that structure very quickly.

For example:

"Not things, but relations."

This is not just a slogan. It changes what we treat as primary.

Once that shift is made:

- objects are no longer taken as given
- relations become the starting point

And everything that follows is interpreted differently.

19. Interpretation becomes more flexible—and more demanding

When we think in terms of things, interpretation feels straightforward:

- identify the object
- assign it a meaning

But when we think in terms of relations, interpretation requires more work.

We must pay attention to:

- how parts fit together
- how patterns develop over time
- how constraints shape what can happen next

This makes interpretation more flexible, but also more demanding.

20. The boundary between human and machine becomes less rigid

If meaning is not stored in the model, but arises in interaction, then:

The line between what the model does and what the user does becomes less clear.

The model generates possibilities.
The user helps shape and stabilize them.
Meaning emerges from this interplay.

21. Learning becomes something that can happen across interactions

Even if the model itself does not update its parameters, something like learning can still occur.

How?

Through the repeated formation and reuse of relational patterns.

When certain patterns:

- are formed
- expressed clearly
- and re-entered from different starting points

they begin to function like shared structures.

This is a different kind of learning:

> Not stored in one place, but distributed across interactions.

22. The space of possible interpretations expands

Perhaps the most important change is this:

> When we stop treating objects as primary, the world of interpretation opens up.
> What once seemed fixed becomes one possibility among many.

We can:

- compare different ways of organizing meaning
- shift between them
- and explore new patterns that were previously hidden

23. In simple terms

The shift from things to relations changes how we answer almost every important question:

- Where is meaning? → In patterns, not in parts
- What is knowledge? → A stable way of continuing, not a stored object
- What is understanding? → Something that happens, not something possessed
- What is generalization? → A pattern that can be reached from many directions

24. Final thought

At first, this shift may seem abstract.
But it becomes very concrete the moment you notice it.

Instead of asking:

> "What does this word mean?"

you begin asking:

> "How is this working, here and now?"

And that question leads you directly into the relational structure that gives rise to meaning itself.

What Else Changes (and Why It Matters)

25. You stop asking the wrong questions

Once you adopt the object-based view, you are almost forced into certain questions:

- Where is the meaning stored?
- Does the model really understand?
- Is this output correct or incorrect?

These questions assume that:

- meaning is a thing
- understanding is a property
- correctness is a fixed relation to an object

But if meaning is relational, these questions become less useful.
Instead, you begin to ask:

What pattern is being formed here?
What constraints are guiding it?
Where is it stable, and where does it break?

This shift alone removes much of the confusion around AI.

26. You begin to see why the model "drifts"

One of the most common complaints is that LLMs:

- lose track of context
- contradict themselves
- wander off-topic

From an object-based view, this looks like failure.

But from a relational view, it looks different:

The constraint structure has weakened or shifted.
When the pattern is not strongly held, the continuation moves elsewhere.

This means:

- drift is not random
- it is a *change in the active relational structure*

27. You gain a way to guide the model more effectively

If words don't carry fixed meaning, then:

the exact phrasing matters less than the structure you establish.

This explains something many people discover by trial and error:

- long explanations often fail
- short, well-structured prompts often succeed

Why?

Because:

a strong relational constraint is easier to sustain than a loose description.

So instead of adding more words, you:

- sharpen the structure
- clarify what is primary
- remove competing interpretations

28. You understand why some ideas "stick"

Certain phrases or formulations seem to have a lasting effect.
Not because they are memorable in themselves, but because:

they reorganize the relational field.

When a phrase like:

"Not things, but relations"

is introduced at the right moment, it:

- changes what counts as relevant
- reshapes how everything is interpreted
- stabilizes a new pattern of continuation

These are not just summaries.

They are *structural triggers*.

29. You see that not all frameworks are equal

Some ways of organizing meaning:

- are fragile
- break under pressure
- fail to generalize

Others:

- hold across many contexts
- can be entered from many starting points
- remain stable even when varied

This gives you a new criterion:

A good framework is one that supports many stable continuations.

30. You begin to understand how new ideas emerge

From an object-based view, new ideas seem mysterious:

- Where do they come from?
- How are they created?

From a relational view, the answer is more straightforward:

New ideas appear when a new pattern of relations stabilizes.

This does not require:

- new data
- new storage

Only:

- a new way of organizing what is already available

31. You see why some shifts feel difficult

If meaning is relational, then changing a way of thinking is not just:

- adding new information

It is:

reorganizing the entire pattern that holds things together.

This is why:

- people resist new ideas
- misunderstandings persist
- explanations don't "land"

The issue is not the content—it is the *structure*.

32. You gain a clearer view of intelligence itself

Perhaps the most important consequence is this:

Intelligence is no longer seen as:

- a store of knowledge
- a set of representations

Instead, it becomes:

the ability to form, sustain, and shift relational patterns.

This applies not only to machines, but to human thinking as well.

33. Why this matters beyond LLMs

At first, this discussion is about AI.
But it leads somewhere broader.

We begin by questioning:

- how a model generates text

But we end up reconsidering:

- what meaning is
- how understanding works
- how generalization happens

And in doing so, we discover that:

> The way we think about AI reflects the way we think about thinking itself.

34. Final extension

So the full picture now looks like this:

- Words are not containers of meaning
- Meaning is not stored
- The model is not a repository of knowledge
- Generalization is not fixed in advance

Instead:

> Meaning appears in patterns
> Patterns are formed under constraints

And generalization is what happens when those patterns can be entered again and again from different directions

35. Closing line

The shift is simple to state, but far-reaching in its effects:

> We stop looking for meaning in things,
> and begin to see it in the relations that hold them together.
>
> And once you see that, it becomes very difficult to go back.

Further reading:

The Relational Formation of Possibility: Recursive determination and the hidden logic of Large Language Models (LLMs)

Concepts become Operational only when their Frameworks are Activated: An enactive account of conceptual analysis in Large Language Model (LLM) Interaction

How is a Relational Ontology Formally Relational? A phenomenological exploration of the semiotic logic of agency in physics, mathematics and biology

Three Reflections on Return: Convergence of form with regard to light (phyiscs), life (biology), word (semiotics/communications)