

Thoughts on Context

From *Anvitābhidhāna* to Attention

Andrew Ollett  June 2, 2025

Outline

- ✿ Approaches to **context**
- ✿ Context in Indian philosophy of language
 - ✿ Especially Prabhākara (*anvitābhīdhāna*)
- ✿ Context in modern language models
- ✿ Comparison and discussion

What is “context”?

✿ We can think of it as a **complementary concept** to “text.”

✿ Context is:

- information,
- existing before, after, or outside of a text,
- that, when taken into account,
- leads us to understand that text in a way that we wouldn't have otherwise.

✿ Since it is defined by its effect on “how we understand” a text, it **effects** or **contributes to** the **meaning** of a text.

What is “context”?

- ✿ We can also think of views on context across a range:
- ✿ CONTEXT [minimal]: the **minimum contribution from context necessary** in order to establish a tenable meaning of the text.
- ✿ CONTEXT [maximal]: **any possible contribution** from context that affects a reader’s understanding of the meaning.

he did it with a substitute.

What is “context”?

- ✿ We can also think of views on context across a range:
- ✿ CONTEXT [minimal]: the **minimum contribution from context necessary** in order to establish a tenable meaning of the text.
- ✿ CONTEXT [maximal]: **any possible contribution** from context that affects a reader’s understanding of the meaning.

By the time of the *agnihōtra*, all the ghee was gone, so he did it with a substitute.

What is “context”?

- ✿ We can also think of views on context across a range:
- ✿ CONTEXT [minimal]: the **minimum contribution from context necessary** in order to establish a tenable meaning of the text.
- ✿ CONTEXT [maximal]: **any possible contribution** from context that affects a reader’s understanding of the meaning.

By the time of the *agnihōtra*, all the ghee was gone, so **he** did **it** with a substitute.

What is “context”?

- ✿ We can also think of views on context across a range:
- ✿ CONTEXT [minimal]: the **minimum contribution from context necessary** in order to establish a tenable meaning of the text.
- ✿ CONTEXT [maximal]: **any possible contribution** from context that affects a reader’s understanding of the meaning.

By the time of the *agnihōtra*, all the ghee was gone, so he did it with a substitute.

Indian views and theories on context

- ❁ As everyone knows, Indian thinkers had a wide range of views about **meaning**.
- ❁ Ironically, however, a discussion of **context** (in the sense of additional information that affects the meaning of an expression) was **marginal** to most of the major discussions.
- ❁ I would suggest that it was marginalized in different ways by the two families of theories of meaning that dominated until the 8th century or so: the *khaṇḍapakṣa* and *akhaṇḍapakṣa*.

The *khaṇḍapakṣa*

- ✿ The *khaṇḍapakṣa* is the view that **linguistic expressions have real parts**. We can think of it as a version of the **compositionality principle**:

The meaning of the whole depends on:

- the meaning of its parts;
- the rules of combination.

We **start with** meaningful parts, and we compose their meanings into larger units of meaning. (A **bottom-up** view of meaning.)

The *khaṇḍapakṣa*

- ✿ In this picture, it's hard to see how context can play a role in determining the meaning either of the parts or of the whole.

(It's not impossible, however: the meanings of the parts might depend on context in some way.)

This is the view taken by most Mīmāṃsakas (especially Śabara and Kumāṛila) and Naiyāyikas.

The *akhaṇḍapakṣa*

- ✿ The *akhaṇḍapakṣa* is the view that **linguistic expressions do not have parts**. Meaning a property of the **entire** expression. We can think of this as a version of the **context principle**:

Parts only have meaning in the context of a whole.

Here “context” doesn’t mean what we have been taking it to mean, but rather refers to the part-whole relationship as a prior condition for attributing meaning to the parts. (A **top-down** view of meaning.)

The *akhaṇḍapakṣa*

- ✿ Here, too, despite the context principle's name, it's hard to see how context in the way we were discussing it — stuff **outside** a linguistic expression that affects the meaning the linguistic expression itself — could play a role, since this view conceives of linguistic expressions as having neither an inside nor an outside!

This is the view of Bhartr̥hari and many grammarians who came after him.

The *akhaṇḍapakṣa*

- ✿ Even thinkers like Bhartṛhari who explicitly adopted the “context principle” could nevertheless have a very “thin” or **minimal** idea of what context actually is or does.
- ✿ Bhartṛhari’s best-known remarks on context have to do with using context to **distinguish** between words in cases of ambiguity (2.316):
 - The bank

The *akhaṇḍapakṣa*

- ✿ Even thinkers like Bhartṛhari who explicitly adopted the “context principle” could nevertheless have a very “thin” or **minimal** idea of what context actually is or does.
- ✿ Bhartṛhari’s best-known remarks on context have to do with using context to **distinguish** between words in cases of ambiguity (2.316):
 - The bank [raised its interest rates]

The *akhaṇḍapakṣa*

- ✿ Even thinkers like Bhartṛhari who explicitly adopted the “context principle” could nevertheless have a very “thin” or **minimal** idea of what context actually is or does.
- ✿ Bhartṛhari’s best-known remarks on context have to do with using context to **distinguish** between words in cases of ambiguity (2.316):
 - The bank [of the river]

The *akhaṇḍapakṣa*

- ✿ Even thinkers like Bhartṛhari who explicitly adopted the “context principle” could nevertheless have a very “thin” or **minimal** idea of what context actually is or does.
- ✿ Bhartṛhari’s best-known remarks on context have to do with using context to **distinguish** between words in cases of ambiguity (2.316):

→ The bank [of the river]

Context is **minimal** here because it merely assists us in determining what the relevant word/meaning pair is.

Context and compositionality

- ✿ In the historiography of modern Western philosophy of language, the context and compositionality principles are often thought of as being in tension with each other:

Parts only have meaning in the context of a whole.

The meaning of the whole depends on:

- the meaning of its parts;
- the rules of combination.

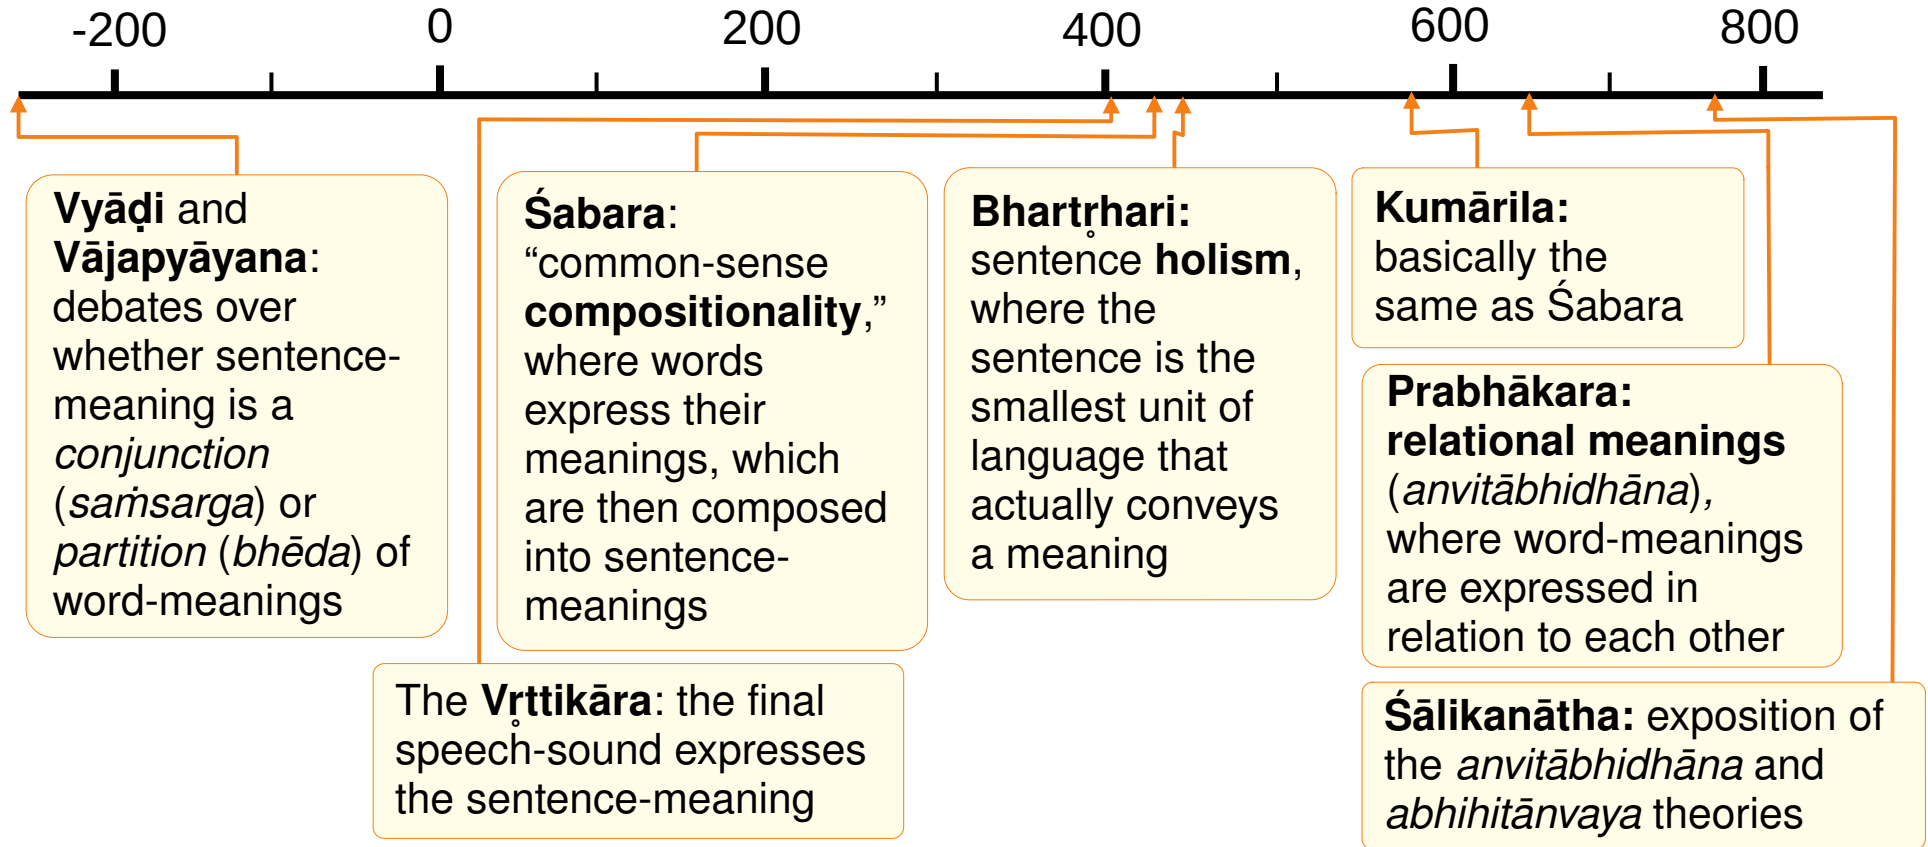
Context and compositionality

- ✿ But both principles seem to get at some common-sense intuitions about meaning.
- ✿ In the 20th century, context (as investigated primarily through the discipline of pragmatics) provided a way of critically evaluating these principles.
- ✿ In the 7th century, a desire to integrate these two principles, on the part of Prabhākara, led to the formulation of a new theory of meaning called *anvitābhidhāna*.

Prabhākara

- ❖ Mīmāṃsā thinkers tended to have a better-defined notion of context than most of their interlocutors, because their project — the determination of the meaning of the sentences of the Vedas — required the **systematic** exploitation of context at almost every level.
- ❖ This meant that Mīmāṃsakas could not accept Bhartrhari's semantic holism *per se*.
- ❖ Yet one of them, Prabhākara, took Bhartrhari's insistence on the context principle very seriously (despite his obvious irritation at Bhartrhari's intuitionist views).

Timeline



Anvitābhīdhāna in overview

✿ Essential texts:

- ✿ Prabhākara's *Bṛhatī, tadbhūtādhikaraṇa* (1.1.24–25)
- ✿ Śālikanātha's *Bhāṣyapariśiṣṭa* on the same section
- ✿ Śālikanātha's *Vākyārthamāṭrkā* (especially part 1)

Note that there are few or no translations for these texts, and even specialists have misunderstood the theory. I recommend the work of **Kei Kataoka**, **Shishir Saxena** and **Patrick Cummins** (and myself of course).

Anvitābhīdhāna in overview

- ✿ **Context principle:** words only have meaning in the context of a sentence.
- ✿ **Relation principle:** the meaning of the sentence is **the relation** (*anvaya*) between the meanings of the words, and the meanings of the words are **related** to each other (*parasparānvita-*) in view of their collective contribution to the meaning of the sentence.


Anvitābhīdhāna in overview

- ✿ **Context principle:** words only have meaning in the context of a sentence.
- ✿ **Relation principle:** the meaning of the sentence is **the relation** (*anvaya*) between the meanings of the words, and the meanings of the words are **related** to each other (*parasparānvita-*) in view of their collective contribution to the meaning of the sentence.

Apparent circularity: is this a top-down or bottom-up theory?


Anvitābhīdhāna in overview

- ✿ In the reconstruction of Śālikanātha (8th/9th c.) it works as follows:



Anvitābhidhāna in overview


- ✿ In the reconstruction of Śālikanātha (8th/9th c.) it works as follows:



“Bring [the] cow”

Anvitābhīdhāna in overview

- ✿ In the reconstruction of Śālikanātha (8th/9th c.) it works as follows:




“Bring [the] cow”

In Sanskrit *gō-* famously has multiple meanings (including the sky, water, the eye, etc.).

The **relational meaning** of the word *gām* (*gō-* in the accusative singular) should encode **a specific meaning** of the word, namely **cow**, based on its cooccurrence with other words in a sentence.

Anvitābhidhāna in overview

- ✿ In the reconstruction of Śālikanātha (8th/9th c.) it works as follows:




“Bring [the] cow”

In fact there is no reason why we cannot take into consideration **other words** from the context, which would further determine the **relational meaning** of the word *gām*, e.g., a particular milk-cow named Bessie that is known to the listener.

Anvitābhidhāna in overview

- ✿ In the reconstruction of Śālikanātha (8th/9th c.) it works as follows:




“Door!”

Śālikanātha uses the traditional example of filling in words from context (*adhyāhāra*) to explain how we arrive at an appropriate meaning from single-word sentences like “door!”

We understand a **related meaning**, namely “open” (if you are on a bus in Chicago), and the word “door” conveys its meaning in relation to this (unexpressed) verb.

Anvitābhidhāna in overview

- ✿ In the reconstruction of Śālikanātha (8th/9th c.) it works as follows:




“Bring [the] cow”

1. the sentence-meaning is the **mutual relation** of the word-meanings

Anvitābhidhāna in overview

- ✿ In the reconstruction of Śālikanātha (8th/9th c.) it works as follows:




2. the sentence-meaning is the **relational word-meanings themselves**


Contrast: Abhihitānvaya

- ✿ Prabhākara's *anvitābhīdhāna* theory is usually contrasted with a theory called *abhihitānvaya* (“the relation of expressed meanings”) attributed to Kumārila.
- ✿ Please bear in mind, however, that Kumārila (who lived a little bit before Prabhākara) **never** used the term himself, and in fact never used *anvaya* in the sense of a “relation” between word-meanings.
- ✿ Śālikanātha came up with the term as a description of a theory he thought Kumārila might have been committed to.

Contrast: Abhihitānvaya




Contrast: Abhihitānvaya




The basic idea of *abhihitānvaya* is that the meaning expressed by the word is “the same” as that which figures in the sentence, i.e., we **update** the meaning of the word only when we are forced to do so. But in fact we are **always** forced to do so...

Contrast: Abhihitānvaya



Contrast: Abhihitānvaya



Contrast: Abhihitānvaya

- ✿ It turns out that we always have to update the non-contextual word-meanings with information from the context, whether we subscribe to *anvitābhīdhāna* or *abhihitānvaya*.
- ✿ But the basic contrast is that **relations with other meanings are built into the very notion of “meaning” in the case of *anvitābhīdhāna***, whereas in the case of *abhihitānvaya*, relations are always **extrinsic** to the word-meaning.

Anvitābhīdhāna and Language Models

- **Rohan Pandey** (OpenAI) has anticipated most of the interesting comparisons between *anvitābhīdhāna* and modern LMs, including especially **contextual embeddings**.


Review: Mīmāṃsā & Neural NLP

- Contextualized embeddings appear to follow Anvitābhīdhāna
 - Words in a sentence capture contextual meaning
- Abhihitānvaya may instead describe how words compose into a single sentence representation
 - Future work exploring syntax in CLS token
- Other psycholinguistic evidence clearly prefers Anvitābhīdhāna
 - Garden path sentences show how context influences syntax parse and word meanings
 - Ex: “the old man the boat”




Anvitābhīdhāna and Language Models

- Here is our provisional interpretation of (Śālikanātha's interpretation of) the *anvitābhīdhāna* theory:



Anvitābhidhāna and Language Models

- And here is a schematic representation of the attention mechanism in a bidirectional transformer:




From Daniel Jurafsky and James H. Martin, *Speech and Language Processing* (2024). Draft of January 12, 2025.

Anvitābhīdhāna and Language Models

- ✿ **Transformers** are the most common kind of language model in use today.
- ✿ They implement context-sensitivity with a mechanism known as **attention**.
- ✿ We can think of this mechanism as a **mathematical implementation** of the basic ideas of *anvitābhīdhāna*.
 - ✿ Of course the ideas have nothing to do with each other historically.

Attention

- ✿ The first step in a transformer is to convert the word (or rather a subword token) into a mathematical representation, i.e., a **vector**, called the **embedding** of the token.
- ✿ The initial embedding of a token is **non-contextual**. The vectors for the token “bank” are the same, regardless of its context (and hence regardless of its different meanings).



Attention

- ✿ What do the values in the non-contextual embedding represent?
- ✿ We won't talk much about it here (the point is to get to a **contextual** embedding) but they are supposed to encode the word's **meaning**: essentially, relations between word-meanings can be thought of as angles between vectors in a certain space.
- ✿ These initial embeddings are **learned** by the model.

Attention


- ☛ Hence, when we talk about **updating** the embedding, we are talking about making changes to the values of the vector that will change its direction vis-à-vis the embeddings of all of the other tokens.

Attention

- ✿ The purpose of the transformer is to **update** these non-contextual embeddings by taking context into account.
- ✿ The mechanism for doing so, **attention**, involves:
 - ✿ computing a measure of **similarity** between the current token and the preceding (and in bidirectional transformers, following) tokens
 - ✿ updating the embedding with new values based on this measure of similarity
 - ✿ allowing the parameters for this process to be learned in training


Attention

- These computations are packed into a layer of the model, the **self-attention** layer, which we can think of as outputting a **contextually enriched** version of the input embedding:



Attention


- ✿ The learned parameters will take the embedding of “bank” in the direction indicated by the word “river.”



In other words, by the end of this process, the vector associated with “bank” will represent a **relational meaning**.

Attention

- ❖ The learned parameters will take the embedding of “bank” in the direction indicated by the word “river.”



When we add attention to the original embeddings, we get **contextual embeddings**.

Attention

- ✿ This is a massive simplification, because attention is usually performed multiple times in parallel (**multi-head attention**).
- ✿ This (in principle) allows the model to capture different ways in which context can influence a word's relational meaning.

Similarities

- ✿ In both attention and *anvitābhidhāna*, we aim to produce a representation of a word's meaning (using both “word” and “meaning” loosely) that **incorporates** the effects that other words in the context might have on it.
- ✿ Both involve a **relation** (*anvaya* or attention) between each element (a word-meaning or embedding) and every other element.

Similarities

✿ Both are based on **learning**:

- ✿ In fact it was often argued **against** *anvitābhīdhāna* that, since each word has an infinite number of possible relations and hence an infinite number of possible relational meanings, it would be impossible for a listener to learn all of the relevant relational meanings.
- ✿ Śālikanātha countered that it is possible (and in fact easy) to learn a finite number of **general constraints** on the form that relational meanings can take.


✿ **dependency, compatibility, proximity**

Similarities


- ✿ Both are based on **learning**:
 - ✿ In attention, a contextual embedding can vary **arbitrarily** with context, but the basic terms of the computation come from embeddings whose values are **finite** and learned in training.

Some major differences: hierarchy

- ❧ The idea of contextual meaning as a **flat** relation between non-contextual meanings is an interpretation of *anvitābhīdhāna* that slightly overstates its similarity to contextual embeddings.




From <https://github.com/jessevig/bertviz>



Some major differences: hierarchy

- ✿ Śālikanātha at least conceived of the relation as **strictly hierarchical** on the model of Mīmāṃsā's concept of *vinīyōga*.

A better model might therefore be graph where all relations are between a **superordinate** and **subordinate** element.



Some major differences: action

- ✿ The hierarchical model accords with Prabhākara's view that sentence-meaning is essentially a **plan for action** (a **performable** or *kārya*) that integrates all elements of meaning within a hierarchical structure.
- ✿ This is not only Prabhākara's view of how language operates in general, but also his view of his language is learned (the points are connected of course):
 - ✿ the process of learning in *anvitābhidhāna* is based on observation of **sentences** and the ensuing **action**, i.e., a language-external state of affairs.

Some major differences: action

- ✿ The hierarchical model accords with Prabhākara's view that sentence-meaning is essentially a **plan for action** (a **performable** or *kārya*) that integrates all elements of meaning within a hierarchical structure.
- ✿ But this introduces some further differences from how language is represented in LMs:
 - ✿ LMs are at least pre-trained on language input alone, so any representation of “meaning” is based **exclusively** on relations between tokens within the training corpus

The last token

- ✿ Finally, we haven't been talking about the “upper layers” of the transformer, where the token, having “absorbed” an enormous amount of context from surrounding tokens, is used to perform some task (like predicting the next token).
- ✿ Very briefly, one of the ways that we can think of contextual LMs working is by loading all the information necessary for the task into the last token.

The last token

- ✿ This resembles an important (but generally discarded) view of sentence-meaning associated with the Vṛttikāra:

It is the final speech-sound, together with the traces produced by previous speech-sounds, that conveys the sentence-meaning.

*pūrvavarṇajaniṭasaṁskārasahitō 'ntyō varṇaḥ
pratyāyakaḥ.* (Śabara's *Mīmāṁsāsūtrabhāṣya* on
1.1.5)

The last token

- ✿ The final “speech-sound” can be compared to the final token (now usually a byte or subword token) — although not meaningful in itself, it “conveys” the meaning associated with the entire sentence, insofar as, by the time we reach the final speech-sound, all of the relational meanings have been computed and taken into account.
- ✿ Śālikanātha does not take this position (he criticizes it in *Vākyārthamāṭrkā* I) but it may help in responding to some of the standard critiques of *anvitābhīdhāna* according to which one word suffices to express the entire sentence-meaning.

Thank you!