

Distributional Semantics

John Gamboa

Resources

- This video is inspired by:
 - "A User's Guide to Thought and Meaning" by Ray Jackendoff (Part I)
 - [Some slides](#) on Distributional Semantics by Marco Baroni

Distributional Semantics

Distributional Semantics

"*Distributional semantics* is a research area that develops and studies theories and methods for *quantifying and categorizing semantic similarities* between linguistic items *based on their distributional properties* in large samples of language data."

https://en.wikipedia.org/wiki/Distributional_semantics

Distributional Semantics

"*Distributional semantics* is a research area that develops and studies theories and methods for **quantifying and categorizing semantic similarities** between linguistic items *based on their distributional properties* in large samples of language data."

https://en.wikipedia.org/wiki/Distributional_semantics

Distributional Semantics

"*Distributional semantics* is a research area that develops and studies theories and methods for **quantifying and categorizing semantic similarities** between linguistic items *based on their distributional properties* in large samples of

What is Semantic Similarity?

Distributional Semantics

"*Distributional semantics* is a research area that develops and studies theories and methods for **quantifying and categorizing semantic similarities** between linguistic items *based on their distributional properties* in large samples of


What is Semantic Similarity?

- **Semantics** studies the meaning of linguistic items

Distributional Semantics

"*Distributional semantics* is a research area that develops and studies theories and methods for **quantifying and categorizing semantic similarities** between linguistic items *based on their distributional properties* in large samples of


What is Semantic Similarity?

- **Semantics** studies the meaning of linguistic items
 - *Semantically similar*  similar meaning

Distributional Semantics

"*Distributional semantics* is a research area that develops and studies theories and methods for **quantifying and categorizing semantic similarities** between linguistic items *based on their distributional properties* in large samples of

What is Semantic Similarity?

- **Semantics** studies the meaning of linguistic items
 - *Semantically similar*  similar meaning

But... what is *meaning*?

Distributional Semantics

"*Distributional semantics* is a research area that develops and studies theories and methods for *quantifying and categorizing semantic similarities* between linguistic items *based on their* **distributional properties** in large samples of language data."

https://en.wikipedia.org/wiki/Distributional_semantics

Distributional Semantics

"*Distributional semantics* is a research area that develops and studies theories and methods for *quantifying and categorizing semantic similarities* between linguistic items *based on their* **distributional properties** in large samples of language data."

**How to use *distributional properties*
to infer anything about meaning?**

Words and Meaning

Words and Meaning

Say we have the following sentences:

Words and Meaning

Say we have the following sentences:

She went down to the river and stood on the bank

She went to town to take some money out of the bank

(from Jackendoff's book)

Words and Meaning

Say we have the following sentences:

She went down to the river and stood on the bank

She went to town to take some money out of the bank

(from Jackendoff's book)

Is the last word in these two sentences the same?

Words and Meaning

She went down to the river and stood on the bank

She went to town to take some money out of the bank

Is the last word in these two sentences the same?

Words and Meaning

She went down to the river and stood on the bank

She went to town to take some money out of the bank

Is the last word in these two sentences the same?

Two alternatives:

Words and Meaning

She went down to the river and stood on the bank

She went to town to take some money out of the bank

Is the last word in these two sentences the same?

Two alternatives:

1. Yes. The same word has two meanings

Words and Meaning

She went down to the river and stood on the bank

She went to town to take some money out of the bank

Is the last word in these two sentences the same?

Two alternatives:

1. Yes. The same word has two meanings
2. No. There are two words “bank”, one with each meaning

Words and Meaning

She went down to the river and stood on the bank

She went to town to take some money out of the bank

Is the last word in these two sentences the same?

Two alternatives:

1. Yes. The same word has two meanings
2. No. There are two words “bank”, one with each meaning

Words and Meaning

She went down to the river and stood on the bank

She went to town to take some money out of the bank

Is the last word in these two sentences the same?

Two alternatives:

1. Yes. The same word has two meanings
2. No. There are two words “bank”, one with each meaning

Homonymy

Words and Meaning

She went down to the river and
She went to town to take some

But what about this:

They broke the door of both banks in
the town

Is the last word in these two sentences the same?

Two alternatives:

1. Yes. The same word has two meanings
2. No. There are two words “bank”, one with each meaning

Homonymy

Words and Meaning

She went down to the river and

She went to town to take some

But what about this:

They broke the door of both banks in
the town

Is the last word in these two sentences the same?

Different grammatical forms

Two alternatives:

1. Yes. The same word has two meanings
2. No. There are two words “bank”, one with each meaning

Homonymy

Words and Meaning

She went down to the river and
She went to town to take some

Is the last word in these two sentences the same?

But what about this:

They broke the door of both banks in
the town

Different grammatical forms

Two alternatives:

1. Yes. The same word has two meanings
2. No. There are two words “bank”, one with each

Homonymy

And what about this:

The State intends to inject money in
the economy to save the banks

Words and Meaning

She went down to the river and

She went to town to take some

But what about this:

They broke the door of both banks in
the town

Is the last word in these two sentences the same?

Different grammatical forms

Two alternatives:

1. Yes. The same word has two meanings
2. No. There are two words “bank”, one with each

Homonymy

And what about this:

The State intends to inject money in
the economy to save the banks

Polysemy

Meaning

- What does “meaning” mean?

Meaning

- What does “meaning” mean?

Both parents are playing a game with their young child. Suddenly the child stops playing, points at a nearby red table, laughs and says "Red table!"

Meaning

*Both parents are playing a game with their young child.
Suddenly the child stops playing, points at a nearby red
table, laughs and says "Red table!"*

Meaning

*Both parents are playing a game with their young child.
Suddenly the child stops playing, points at a nearby red
table, laughs and says "Red table!"*

“the nearby table is red”

Meaning

*Both parents are playing a game with their young child.
Suddenly the child stops playing, points at a nearby red
table, laughs and says "Red table!"*

“the nearby table is red”

“there is a table nearby”

Meaning

*Both parents are playing a game with their young child.
Suddenly the child stops playing, points at a nearby red
table, laughs and says "Red table!"*

“the nearby table is red”

“there is a table nearby”

“the object that is pointed out is a table”

Meaning

*Both parents are playing a game with their young child.
Suddenly the child stops playing, points at a nearby red
table, laughs and says "Red table!"*

“the nearby table is red”

“there is a table nearby”

“the object that is pointed out is a table”

“the child is proud of its ability to communicate”

Meaning

*Both parents are playing a game with their young child.
Suddenly the child stops playing, points at a nearby red
table, laughs and says "Red table!"*

“the nearby table is red”

“there is a table nearby”

“the object that is pointed out is a table”

“the child is proud of its ability to communicate”

“the child can formulate correct phrases but in general doesn't seem to use full sentences yet.”

Meaning

*Both parents are playing a game with their young child.
Suddenly the child stops playing, points at a nearby red
table, laughs and says "Red table!"*

“the nearby table is red”

“there is a table nearby”

“the object that is pointed out is a table”

“the child is proud of its ability to communicate”

“the child likes the color of the table”

“the child can formulate correct phrases but in general doesn't seem to use full sentences yet.”

Meaning

*Both parents are playing a game with their young child.
Suddenly the child stops playing, points at a nearby red
table, laughs and says "Red table!"*

“the nearby table is red”

“there is a table nearby”

“the object that is pointed out is a table”

“the child is proud of its ability to communicate”

“the child can formulate correct phrases but in general doesn't seem to use full sentences yet.”

“the child likes the red table”

“the child likes the color of the table”

Meaning

*Both parents are playing a game with their young child.
Suddenly the child stops playing, points at a nearby red
table, laughs and says "Red table!"*

“the nearby table is red”

“there is a table nearby”

“the object that is pointed out is a table”

“the child is proud of its ability to communicate”

“the child can formulate correct phrases but in general doesn't seem to use full sentences yet.”

“the child doesn't want to play the game anymore”

“the child likes the red table”

“the child likes the color of the table”

Meaning

Both parents are playing a game with the child. Suddenly the child stops playing, points at a nearby red table, laughs and says "Red table!"

How would the answer change if the table were actually on fire?

“the nearby table is red”

“there is a table nearby”

“the object that is pointed out is a table”

“the child is proud of its ability to communicate”

“the child can formulate correct phrases but in general doesn't seem to use full sentences yet.”

“the table caught the attention of the child”

“the child doesn't want to play the game anymore”

“the child likes the red table”

“the child likes the color of the table”

Meanings of mean

Meanings of mean

- Interpretation uses of “X means Y”
 - *Rauch means smoke. (Translation)*
 - *Slithy means lithe and slimy. (Definition)*
 - *Waving means doing this. (Demonstration)*
 - *A red light means you have to stop.
(Explanation of symbols)*

Meanings of mean

- Interpretation uses of “X means Y”
 - *Rauch means smoke. (Translation)*
 - *Slithy means lithe and slimy. (Definition)*
 - *Waving means doing this. (Demonstration)*
 - *A red light means you have to stop.
(Explanation of symbols)*
- Linkage uses of “X means Y”
 - *Smoke means fire*
 - *This means war*

Meanings of mean

- Interpretation uses of “X means Y”
 - *Rauch means smoke. (Translation)*
 - *Slithy means lithe and slimy. (Definition)*
 - *Waving means doing this. (Demonstration)*
 - *A red light means you have to stop. (Explanation of symbols)*
- Linkage uses of “X means Y”
 - *Smoke means fire*
 - *This means war*
- “The meaning of X is Y”
 - The meaning of Rauch is smoke.
 - The meaning of Slithy is lithe and slimy
 - ?The meaning of waving is doing this
 - ?The meaning of a red light is you have to stop
 - *The meaning of smoke is fire

Meanings of mean

- Interpretation uses of “X means Y”
 - *Rauch means smoke. (Translation)*
 - *Slithy means lithe and slimy. (Definition)*
 - *Waving means doing this. (Demonstration)*
 - *A red light means you have to stop. (Explanation of symbols)*
- Linkage uses of “X means Y”
 - *Smoke means fire*
 - *This means war*
- “The meaning of X is Y”
 - The meaning of Rauch is smoke.
 - The meaning of Slithy is lithe and slimy
 - ?The meaning of waving is doing this
 - ?The meaning of a red light is you have to stop
 - *The meaning of smoke is fire
- “X has the same meaning as Y”
 - Rauch has the same meaning as smoke
 - Slithy has ... mng as lithe and slimy
 - *Waving has ... mng as doing this
 - *A red light has ... mng as you have to stop
 - *Smoke has ... mng as fire

Meanings of mean

- Interpretation uses of “X means Y”
 - *Rauch means smoke. (Translation)*
 - *Slithy means lithe and slimy. (Definition)*
 - *Waving means doing this. (Demonstration)*
 - *A red light means you have to stop. (Explanation of symbols)*
- Linkage uses of “X means Y”
 - *Smoke means fire*
 - *This means war*

- “The meaning of X is Y”
 - The meaning of Rauch is smoke.
 - The meaning of Slithy is lithe and slimy
 - ?The meaning of waving is doing this
 - ?The meaning of a red light is you have to stop
 - *The meaning of smoke is fire
- “X has the same meaning as Y”
 - Rauch has the same meaning as smoke
 - Slithy has ... mng as lithe and slimy
 - *Waving has ... mng as doing this
 - *A red light has ... mng as you have to stop
 - *Smoke has ... mng as fire

But what is the meaning of “this”?

Meanings of mean

- Interpretation uses of “X means Y”
 - Rauch means smoke. (Translation)
 - Slithy means lithe and slimy.
 - Waving means doing this. (Description)
 - A red light means you have to stop. (Explanation of symbols)

- Linkage uses of “X means Y”
 - Smoke means fire
 - This means war

Clearly, “this” has a meaning; more than “thit” or “bliff”

But what is the meaning of “this”?

- “X has the same meaning as Y”
 - Rauch has the same meaning as smoke
 - Slithy has ... mng as lithe and slimy
 - *Waving has ... mng as doing this
 - *A red light has ... mng as you have to stop
 - *Smoke has ... mng as fire

Meanings of mean

- Interpretation uses of “X means Y”
 - Rauch means smoke. (Translation)
 - Slithy means lithe and slimy.
 - Waving means doing this. (Description)
 - A red light means you have to stop. (Explanation of symbols)
- Linkage uses of “X means Y”
 - Smoke means fire
 - This means war

Clearly, “this” has a meaning;
more than “thit” or “bliff”

We’re aiming at a *cognitive understanding*
of their meaning

But what is the meaning of “this”?

- “X has the same meaning as Y”
 - Rauch has the same meaning as smoke
 - Slithy has ... mng as lithe and slimy
 - *Waving has ... mng as doing this
 - *A red light has ... mng as you have to stop
 - *Smoke has ... mng as fire

Meanings are not purely visual

Meanings are not purely visual

The white ball hit the green ball, and that caused the green ball to move.

The white ball hit the green ball and then immediately the green ball moved.

Categories are not clear-cut

Categories are not clear-cut

- What counts as a chair?

Categories are not clear-cut

- What counts as a chair?

Does a chair need to have legs?

Is a sofa a kind of a chair?

Do you have to be able to sit on a chair?

Context matters

Context matters

- We already saw the example of “this”

Context matters

- We already saw the example of “this”
- E.g., let’s take a random word...

Context matters

- We already saw the example of “this”
- E.g., let’s take a random word... “spaghetti”

Context matters

- We already saw the example of “this”
- E.g., let’s take a random word... “spaghetti”

What about “spaghetti code”?

Context matters

- We already saw the example of “this”
- E.g., let’s take a random word... “spaghetti”

What about “spaghetti code”?

... or “spaghetti western”?

The Distributional Hypothesis

In Distributional Semantics, we try to use context to predict meaning...

The Distributional Hypothesis

In Distributional Semantics, we try to use context to predict meaning...

Basic idea:

“Linguistic items that appear in similar contexts have similar meanings”

The Distributional Hypothesis

In Distributional Semantics, we try to use context to predict meaning...

Basic idea:

“Linguistic items that appear in similar contexts have similar meanings”

Or, in a more precise/radical version by Marco Baroni:

The Distributional Hypothesis

In Distributional Semantics, we try to use context to predict meaning...

Basic idea:

“Linguistic items that appear in similar contexts have similar meanings”

Or, in a more precise/radical version by Marco Baroni:

“The meaning of a word is the set of contexts in which it occurs in texts.”

The Distributional Hypothesis

In Distributional Semantics, we try to use context to predict meaning...

Basic idea:

“Linguistic items that appear in similar contexts have similar meanings”

Or, in a more precise/radical version by Marco Baroni:

“The meaning of a word is the set of contexts in which it occurs in texts.”

“Important aspects of the meaning of a word are a function of the set of contexts in which it occurs in texts.”