Data annotation for AI-driven data analysis

III Data Analysis Appetiser
PhD in English Linguistics - UCM

Giorgia Troiani, Ludovica Pannitto

Mon, Nov 24th 2025

LILEC - University of Bologna

CC BY NC SA

Data annotation for AI-driven data analysis

III Data Analysis Appetiser
PhD in English Linguistics - UCM

Giorgia Troiani, Ludovica Pannitto

Mon, Nov 24th 2025

LILEC - University of Bologna

Overview

A definition of AI

- Back in 1966

From ELIZA to Data-driven Artificial Intelligence

Data in Linguistic and how to learn from it

Sources of bias

Model applications

A definition of AI

A definition of AI

The study and design of *agents* that map perceptions to actions

$$f : P \rightarrow A \tag{1}$$

AI agents aim to maximize a *performance measure* based on **models of the environment** in which they operate.

AI is not exactly a *recent* field of study:

- 1950: **Alan Turing** publishes *Computing Machinery and Intelligence* (Turing 2021).
- Formally began in 1956 at the **Dartmouth Summer Research Project on Artificial Intelligence**: *"Every aspect of learning or any other feature of intelligence can in principle be so **precisely described** that a machine can be made to **simulate** it."* (McCarthy et al. 2006)

A definition of AI

The study and design of *agents* that map perceptions to actions

$$f : P \rightarrow A \tag{1}$$

AI agents aim to maximize a *performance measure* based on **models of the environment** in which they operate.

AI is not exactly a *recent* field of study:

- 1950: **Alan Turing** publishes *Computing Machinery and Intelligence* (Turing 2021).
- Formally began in 1956 at the **Dartmouth Summer Research Project on Artificial Intelligence**: *"Every aspect of learning or any other feature of intelligence can in principle be so **precisely described** that a machine can be made to **simulate** it."* (McCarthy et al. 2006)

Rule-Based Intelligence

8-puzzle game

The 8-Puzzle consists of a 3x3 grid with numbered tiles and one empty space. The goal is to reach a target configuration through legal moves.

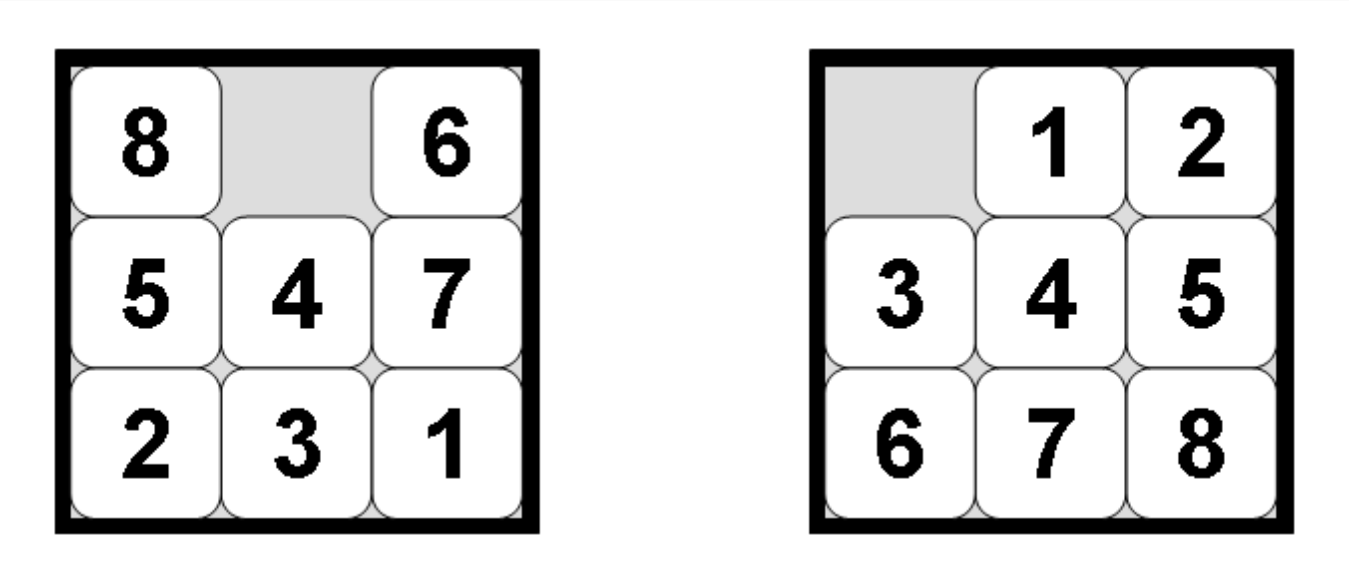


Figure 1: Image from *https://www.aiai.ed.ac.uk/*

The agent perspective

Perceptions the current layout of the puzzle grid — positions of all tiles and the empty space.

State an internal encoding of the puzzle configuration, often a 3x3 matrix describing each tile's location.

Actions moving the empty space *up*, *down*, *left*, or *right* (swapping it with an adjacent tile).

Transition model function that maps a *state* and *action* to a resulting state

$$\left(\begin{bmatrix} 8 & & 6 \\ 5 & 4 & 7 \\ 2 & 3 & 1 \end{bmatrix}, \mathit{down}\right) \rightarrow \begin{bmatrix} 8 & 4 & 6 \\ 5 & & 7 \\ 2 & 3 & 1 \end{bmatrix} \tag{2}$$

Why do we need AI?

- The 8-puzzle has $9! = 362,880$ possible configurations. An **exhaustive search** through all states is computationally expensive.
- AI search algorithms use **heuristics** to estimate the distance to the goal: for instance, *number of misplaced tiles*, *manhattan distance*.
- We often want to find a **quick, good-enough** solution.

In short

We use AI to make the problem **tractable**, guiding the search toward promising states, accepting that the result could be an approximation.

Why do we need AI?

- The 8-puzzle has $9! = 362,880$ possible configurations. An **exhaustive search** through all states is computationally expensive.
- AI search algorithms use **heuristics** to estimate the distance to the goal: for instance, *number of misplaced tiles*, *manhattan distance*.
- We often want to find a **quick, good-enough** solution.

In short

We use AI to make the problem **tractable**, guiding the search toward promising states, accepting that the result could be an approximation.

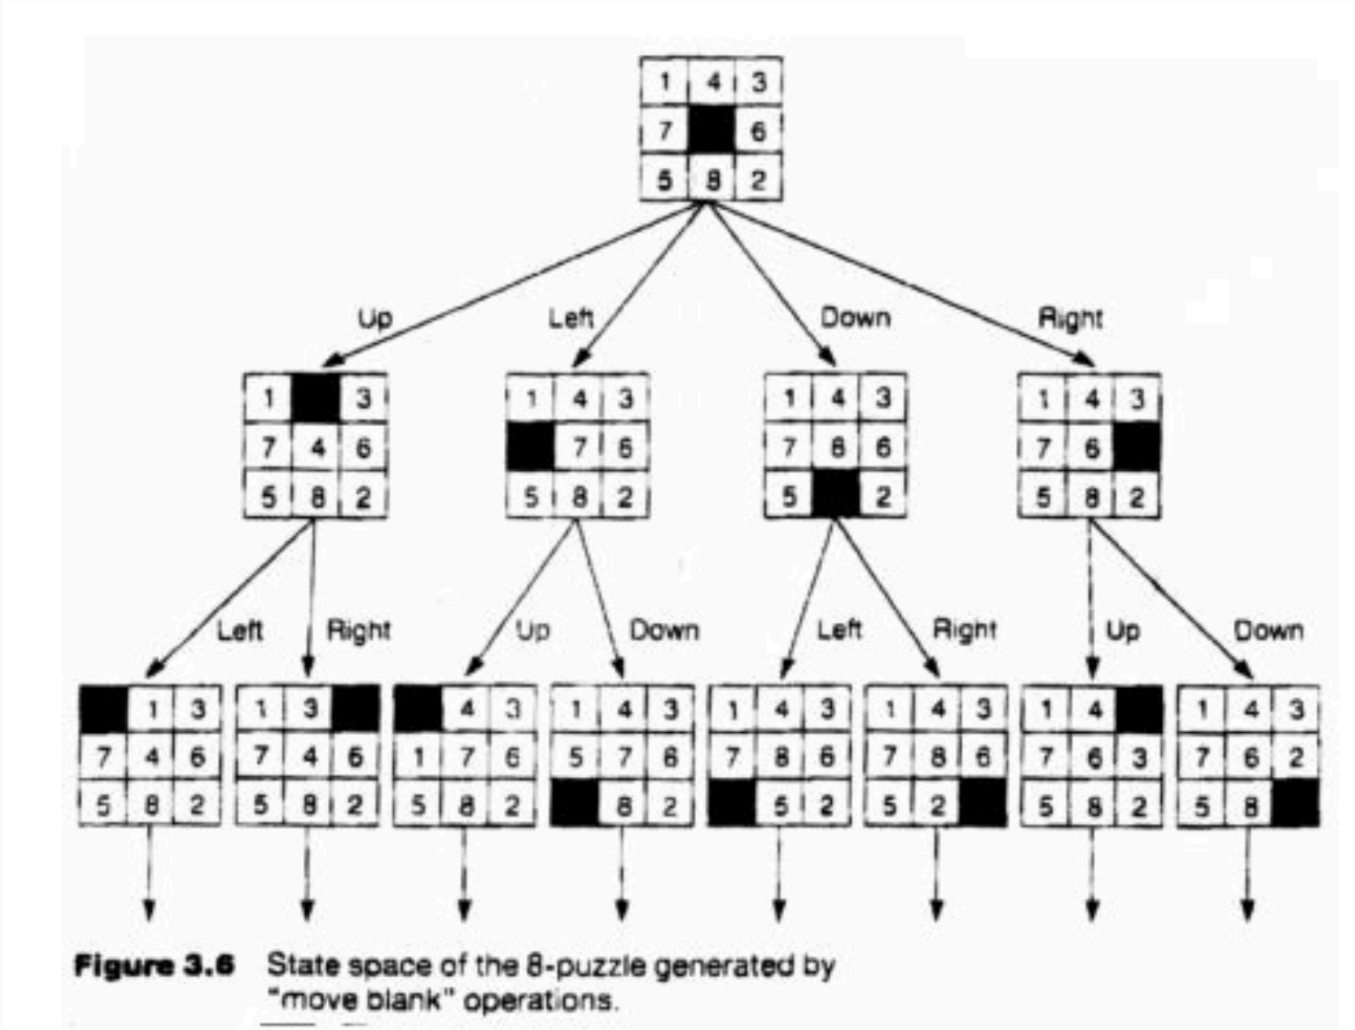


Figure 2: Full search space for the 8-puzzle problem

Back in 1966

```
USER:  Men are all alike.
ELIZA: In what way ?
USER:  They're always bugging us about something or other.
ELIZA: Can you think of a specific example ?
USER:  Well, my boyfriend made me come here.
ELIZA: Your boyfriend made you come here ?
USER:  He says I'm depressed much of the time.
ELIZA: I am sorry to hear that you are depressed.
USER:  It's true. I am unhappy.
ELIZA: Do you think coming here will help you not to be unhappy ?
USER:  I need some help, that much seems certain.
ELIZA: What would it mean to you if you got some help ?
USER:  Perhaps I could learn to get along with my mother.
ELIZA: Tell me more about your family.
USER:  My mother takes care of me.
ELIZA: Who else in your family takes care of you ?
USER:  My father.
ELIZA: Your father ?
USER:  You are like my father in some ways.
ELIZA: What resemblence do you see ?
USER:  You are not very aggressive but I think you don't want me to notice that.
ELIZA: What makes you think I am not very aggressive ?
USER:  You don't argue with me.
ELIZA: Why do you think I don't argue with you ?
```

Figure 3: From *https://www.masswerk.at/elizabot/*

ELIZA: The First Conversational Agent

ELIZA was developed by Joseph Weizenbaum at MIT simulating a rogerian psychotherapist (Weizenbaum 1966).

It used templates and keyword substitution to generate human-like dialogue, and demonstrated that users could attribute understanding and empathy to rule-based systems.

ELIZA as an Intelligent Agent

Perceptions the user's input text — a string of words entered by the human interlocutor.

State internal rule set of keyword patterns and response templates that guide ELIZA's behavior.

Actions textual responses generated by selecting and filling a matching rule template.

Example:

User: "I am X today."

ELIZA: "Do you think coming here will help you not to be X?"

```
ELIZA: Please tell me what's been bothering you.
YOU:   I am sad today
ELIZA: Can you explain what made you sad ?
YOU:   I am sad today
ELIZA: I am sorry to hear that you are sad.
YOU:   I am sad today
ELIZA: I'm sure it's not pleasant to be sad.
YOU:   I am sad today
ELIZA: I am sorry to hear that you are sad.
YOU:   I am sad today
ELIZA: Do you think coming here will help you not to be sad
?
YOU:   I am sad today
ELIZA: Can you explain what made you sad ?
YOU:   I am sad today
ELIZA: I am sorry to hear that you are sad.
```

```
ELIZA: Is something troubling you ?
YOU:   I am happy today
ELIZA: Can you explain why you are suddenly happy ?
YOU:   I am happy today
ELIZA: How have I helped you to be happy ?
YOU:   I am happy today
ELIZA: Has your treatment made you happy ?
YOU:   I am happy today
ELIZA: What makes you happy just now ?
YOU:   I am happy today
ELIZA: Can you explain why you are suddenly happy ?
YOU:   I am happy today
ELIZA: What makes you happy just now ?
YOU:   I am happy today
ELIZA: Can you explain why you are suddenly happy ?
```

```
ELIZA: Please tell me what's been bothering you.
YOU:   I am sad today
ELIZA: Can you explain what made you sad ?
YOU:   I am sad today
ELIZA: I am sorry to hear that you are sad.
YOU:   I am sad today
ELIZA: I'm sure it's not pleasant to be sad.
YOU:   I am sad today
ELIZA: I am sorry to hear that you are sad.
YOU:   I am sad today
ELIZA: Do you think coming here will help you not to be sad
?
YOU:   I am sad today
ELIZA: Can you explain what made you sad ?
YOU:   I am sad today
ELIZA: I am sorry to hear that you are sad.
```

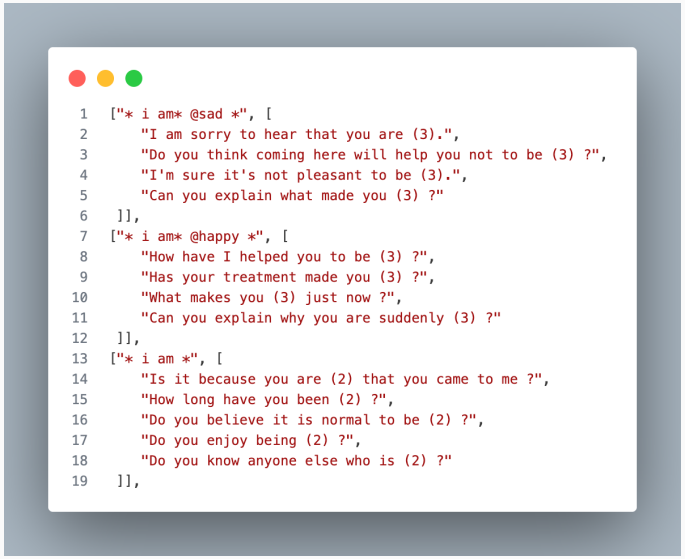
```
ELIZA: Is something troubling you ?
YOU:   I am happy today
ELIZA: Can you explain why you are suddenly happy ?
YOU:   I am happy today
ELIZA: How have I helped you to be happy ?
YOU:   I am happy today
ELIZA: Has your treatment made you happy ?
YOU:   I am happy today
ELIZA: What makes you happy just now ?
YOU:   I am happy today
ELIZA: Can you explain why you are suddenly happy ?
YOU:   I am happy today
ELIZA: What makes you happy just now ?
YOU:   I am happy today
ELIZA: Can you explain why you are suddenly happy ?
```

```
ELIZA: Is something troubling you ?
YOU:   I am frusy today
ELIZA: Do you enjoy being frusy today ?
YOU:   I am frusy today
ELIZA: How long have you been frusy today ?
YOU:   I am frusy today
ELIZA: Do you believe it is normal to be frusy today ?
YOU:   I am frusy today
ELIZA: How long have you been frusy today ?
YOU:   I am frusy today
ELIZA: Do you enjoy being frusy today ?
YOU:   I am frusy today
ELIZA: Do you know anyone else who is frusy today ?
YOU:   I am frusy today
ELIZA: How long have you been frusy today ?
```

From ELIZA to Data-driven Artificial Intelligence

Would you use ELIZA for your research?

Probably **not** a smart idea.



```
["* i am* @sad *", [
    "I am sorry to hear that you are (3).",
    "Do you think coming here will help you not to be (3) ?",
    "I'm sure it's not pleasant to be (3).",
    "Can you explain what made you (3) ?"
 ]],
["* i am* @happy *", [
    "How have I helped you to be (3) ?",
    "Has your treatment made you (3) ?",
    "What makes you (3) just now ?",
    "Can you explain why you are suddenly (3) ?"
 ]],
["* i am *", [
    "Is it because you are (2) that you came to me ?",
    "How long have you been (2) ?",
    "Do you believe it is normal to be (2) ?",
    "Do you enjoy being (2) ?",
    "Do you know anyone else who is (2) ?"
 ]],
```

Figure 4: Patterns from ELIZA emulator

ELIZA uses no data: there is no corpus, annotation, or statistical **inference** involved.

The system cannot **generalize** beyond predefined keyword patterns and it lacks any **internal representation** of language structure or meaning.

Implications for Linguistic Research

AI is not necessarily about data, but modern linguistics is.

ELIZA is historically important, but **methodologically irrelevant for modern linguistics**, because there's no way to test linguistic hypotheses or extract empirical patterns.

It might offer insight into *interaction design*, not on *language modeling*.

We rely on **corpora** — large collections of **real** language data, and expect models to be evaluated on **empirical evidence**, not handcrafted rules.

The missing ingredient: data

In AI, data gains centrality when Machine Learning is introduced (in practice, from 1990s on.)

Machine Learning

Instead of coding $f : P \to A$ explicitly, we **learn an approximation** from examples:

$$\hat{f} : P \to A \tag{3}$$

The agent observes input–output pairs or experience, adjusting internal parameters to improve performance.

Learning from collected data

Idea

Humans solving the 8-puzzle use intuition to decide which moves bring them closer to the goal. Machine learning can **capture this intuition** by learning heuristics from human play data.

Data collection We record sequences of moves made by human players and label each puzzle state with the distance (or estimated effort) remaining to solve it.

Model we train a function (regression) that maps puzzle configurations to predicted “distance-to-goal”.

Outcome the AI agent can use this trained function to select the next move and guide the search.

Example

State configuration	f_1 (misplaced tiles)	f_2 (Manhattan distance)	y (moves to goal)
8 _ 6 5 4 7 2 3 1	8	19	23
8 2 _ 5 4 7 6 3 1	7	15	20
1 2 3 4 5 6 7 8 _	0	0	0

Fit $\hat{h}(\mathbf{x})$ using many labeled states (from human play or an optimal solver):

$$\hat{h}(x) = \hat{h}(f_1, f_2) = w_0 + w_1 f_1 + w_2 f_2 \tag{4}$$

Data in Linguistic and how to learn from it

What is Data in Linguistics?

What we call **data** is essentially **samples of language use** — words, sentences, documents — often paired with **annotated information**.

Linguistic context from corpora n-grams, bag of words.

Labels linguistic categories (POS tags, syntactic dependencies, semantic roles).

Metadata author, genre, date, register.

Empirical measures reading times, human-elicited productions.

The goal of data-driven research: to capture **patterns of usage** that reflect structure and meaning in real language.

Feature-Based Machine Learning

In traditional ML, humans design the features that represent the data.

Similarly to the 8-puzzle case, each linguistic event is represented as a **feature vector**.

- Word frequency, length, part-of-speech ratios.
- Number of subordinate clauses or pronouns.
- Average sentence length or lexical diversity.

These features are interpretable and linked to linguistic hypotheses.

From ELIZA to Feature-Based Machine Learning

We can train a model to choose responses based on **linguistic features** extracted from dialogue data.

1. Collect dialogue pairs (user input → therapist-style response) from transcripts.
2. Label therapist-response depending on our set of actions ("reflection", "question" or "encouragement").
3. Extract interpretable features from user input.
4. train a classifier to predict the **response type** given a linguistic representation of the input.

Outcome: The model mimics ELIZA's conversational behavior, but learns response patterns from data rather than fixed rules — an *interpretable, feature-based chatbot*.

Example Dialogue for Feature Extraction

Excerpt from a Therapist–Patient Dialogue

Patient: I feel really anxious about my job lately.
Therapist [question]: What makes you feel anxious about it?

Patient: My boss keeps criticizing everything we do.
Therapist [question]: Do you think your boss expects too much from you?

Patient: Maybe, but I just can't stop worrying about it.
Therapist [reflection]: It sounds like this is causing you a lot of stress.

Matrix Representation:

$$\mathbf{X} = \begin{bmatrix} 1 & 0 & 0 & 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 0 & 1 & 0 \end{bmatrix}$$

Target labels (response types):

$$\mathbf{y} = \begin{bmatrix} \text{question} \\ \text{question} \\ \text{reflection} \end{bmatrix}$$

- Each row corresponds to a patient sentence.
- Each column corresponds to a hand-selected feature (e.g., presence of “I” pronoun, presence of “you” pronoun, presence of “we” pronoun, presence of proper noun, presence of negative emotion word, presence of positive emotion word, presence of adverb, reference to time).

Feature-Based Models in Linguistic Research

Beyond simulating conversation, the function $\hat{h} : \mathbf{X} \rightarrow \mathbf{y}$ becomes a way to **test linguistic hypotheses**.

- By changing the **feature set**, we can measure which aspects of language most strongly **predict behavior or interpretation**.
- Comparing models with different features allows researchers to quantify what type of information is **most informative** for a given task.
- Performance differences across feature sets become **evidence** for or against specific **theoretical claims**.

If a model can predict a property from data, the information needed to make that distinction is **encoded** in the input representation.

Feature-Based Models in Linguistic Research

Beyond simulating conversation, the function $\hat{h} : \mathbf{X} \rightarrow \mathbf{y}$ becomes a way to **test linguistic hypotheses**.

- By changing the **feature set**, we can measure which aspects of language most strongly **predict behavior or interpretation**.
- Comparing models with different features allows researchers to quantify what type of information is **most informative** for a given task.
- Performance differences across feature sets become **evidence** for or against specific **theoretical claims**.

If a model can predict a property from data, the information needed to make that distinction is **encoded** in the input representation.

From Feature Engineering to Representation Learning

A Different Paradigm

Instead of manually designing features, modern models **learn their own representations** directly from data.

- By solving some prediction task, the model learns a function that transforms the input into a vector representation that captures relevant linguistic information automatically.
- Possible tasks are: predicting the next word, filling a blank, classifying a sentence.

Paradigm shift: From explicit, interpretable features to implicit, data-driven abstractions.

Examples of Representation Learning

- **Word embeddings:** models like Word2Vec or GloVe learn vector representations where similar words (“king”, “queen”) are close in space.
- **Encoder architectures:** models like BERT or RoBERTa learn bi-directional context-dependent representations that capture meaning of words in context.
- **Encoder-Decoder models:** GPT-style models learn to represent syntax, semantics, and discourse structure and exploit the representation to generate new linguistic sequences.
- **Multimodal representations:** CLIP learns joint embeddings of images and text.

Representation Learning in AI

Connecting Learning and Action

An intelligent agent uses representation learning to form an **internal model of the world**

$$\text{Percepts} \xrightarrow{g} \text{State representation} \xrightarrow{f} \text{Action} \tag{5}$$

- g (the encoder) transforms raw inputs into latent states that summarize relevant information.
- The decision function f maps those internal states to actions.

Representations as Tools for Linguistic Research

Learned representations can be **analyzed** to test linguistic hypotheses.

- **Probing:** Train diagnostic classifiers ("probes") to test whether specific linguistic properties are encoded in model representations.
- **Comparative studies:** Compare model similarity patterns to human behavioral or neuro-cognitive data to assess alignment with human processing.
- **Language variation:** Use trajectories in representation space to model semantic drift, register differences, or diachronic shifts.

Key idea: Representations turn neural models into experimental laboratories for building **models of language**.

Sources of bias

What is bias?

- In statistics, **bias** is the expectation we should have for something before we see any further evidence
- In a model, bias is disproportionate weight given to an outcome
- Bias is an issue only when kept in the face of contrary evidence
- Bias can be a source of information

What is bias?

- In statistics, **bias** is the expectation we should have for something before we see any further evidence
- In a model, bias is disproportionate weight given to an outcome
- Bias is an issue only when kept in the face of contrary evidence
- Bias can be a source of information

Sources of bias

1. Bias from research design
2. Bias from data - selection bias
3. Bias from annotation - label bias
4. Bias from input representations - semantic bias
5. Bias from models - bias overamplification

Sources of bias

1. Bias from research design
2. **Bias from data - selection bias**
3. **Bias from annotation - label bias**
4. Bias from input representations - semantic bias
5. Bias from models - bias overamplification

Sources of bias

1. Bias from research design
2. **Bias from data - selection bias**
3. **Bias from annotation - label bias**
4. Bias from input representations - semantic bias
5. Bias from models - bias overamplification

Selection bias

- Models reflect the datasets that are used in the training
  - Distributional frequency of the linguistic constructions
  - Demographics who contribute the most to the dataset
- Overfitting on specific linguistic indexes of these groups
- E.g., using the web as a corpus for POS tagging
  - who contributes more to Wikipedia?
  - (typically) men, educated, falling within a specific age range
  - POS-tagging models underperform on the speech of youth and ethnic minorities

Selection bias

- Models reflect the datasets that are used in the training
  - Distributional frequency of the linguistic constructions
  - Demographics who contribute the most to the dataset
- Overfitting on specific linguistic indexes of these groups
- E.g., using the web as a corpus for POS tagging
  - who contributes more to Wikipedia?
  - (typically) men, educated, falling within a specific age range
  - POS-tagging models underperform on the speech of youth and ethnic minorities

Selection bias

- Models reflect the datasets that are used in the training
  - Distributional frequency of the linguistic constructions
  - Demographics who contribute the most to the dataset
- Overfitting on specific linguistic indexes of these groups
- E.g., using the web as a corpus for POS tagging
  - who contributes more to Wikipedia?
  - (typically) men, educated, falling within a specific age range
  - POS-tagging models underperform on the speech of youth and ethnic minorities

Selection bias

- Models reflect the datasets that are used in the training
  - Distributional frequency of the linguistic constructions
  - Demographics who contribute the most to the dataset
- Overfitting on specific linguistic indexes of these groups
- E.g., using the web as a corpus for POS tagging
  - who contributes more to Wikipedia?
  - (typically) men, educated, falling within a specific age range
  - POS-tagging models underperform on the speech of youth and ethnic minorities

Countermeasures to selection bias

1. Document the choices that went into building the dataset
    - Data Statement framework by Bender and Friedman (2018)
2. In existing datasets, apply post-stratification techniques
    - Down-sample overly represented groups to reflect the actual distribution of the phenomenon (multiple cycles)
3. Collect additional data to account for exclusion or low frequency of groups

Countermeasures to selection bias

1. Document the choices that went into building the dataset
    - Data Statement framework by Bender and Friedman (2018)
2. In existing datasets, apply post-stratification techniques
    - Down-sample overly represented groups to reflect the actual distribution of the phenomenon (multiple cycles)
3. Collect additional data to account for exclusion or low frequency of groups

Countermeasures to selection bias

1. Document the choices that went into building the dataset
    - Data Statement framework by Bender and Friedman (2018)
2. In existing datasets, apply post-stratification techniques
    - Down-sample overly represented groups to reflect the actual distribution of the phenomenon (multiple cycles)
3. Collect additional data to account for exclusion or low frequency of groups

Label bias

1. Bias that is introduced into the data by the annotators' choices at the moment of preparing the dataset for analysis
    - Mistakes caused by distraction or wrong interpretation
    - Disagreement between annotators
    - Mismatch between researcher's and annotator's social norms
2. Bias that is tied to the researcher's analytical choices
    - Orthographic vs phonetic representation of connected speech
    - Orthographic representation ensures consistency across data
    - Phonetic representation allows to capture change in progress on reduced forms (e.g. *gonna, wanna, finna*)

Label bias

1. Bias that is introduced into the data by the annotators' choices at the moment of preparing the dataset for analysis
    - Mistakes caused by distraction or wrong interpretation
    - Disagreement between annotators
    - Mismatch between researcher's and annotator's social norms
2. Bias that is tied to the researcher's analytical choices
    - Orthographic vs phonetic representation of connected speech
    - Orthographic representation ensures consistency across data
    - Phonetic representation allows to capture change in progress on reduced forms (e.g. *gonna*, *wanna*, *finna*)

Label bias

1. Bias that is introduced into the data by the annotators' choices at the moment of preparing the dataset for analysis
    - Mistakes caused by distraction or wrong interpretation
    - Disagreement between annotators
    - Mismatch between researcher's and annotator's social norms
2. Bias that is tied to the researcher's analytical choices
    - Orthographic vs phonetic representation of connected speech
    - Orthographic representation ensures consistency across data
    - Phonetic representation allows to capture change in progress on reduced forms (e.g. *gonna*, *wanna*, *finna*)

Label bias

1. Bias that is introduced into the data by the annotators' choices at the moment of preparing the dataset for analysis
    - Mistakes caused by distraction or wrong interpretation
    - Disagreement between annotators
    - Mismatch between researcher's and annotator's social norms
2. Bias that is tied to the researcher's analytical choices
    - Orthographic vs phonetic representation of connected speech
    - Orthographic representation ensures consistency across data
    - Phonetic representation allows to capture change in progress on reduced forms (e.g. *gonna*, *wanna*, *finna*)

Label bias

1. Bias that is introduced into the data by the annotators' choices at the moment of preparing the dataset for analysis
    - Mistakes caused by distraction or wrong interpretation
    - Disagreement between annotators
    - Mismatch between researcher's and annotator's social norms
2. Bias that is tied to the researcher's analytical choices
    - Orthographic vs phonetic representation of connected speech
    - Orthographic representation ensures consistency across data
    - Phonetic representation allows to capture change in progress on reduced forms (e.g. *gonna, wanna, finna*)

Countermeasures to label bias

1. Aggregate multiple annotations to even out results
2. Use knowledge about disagreement as information to update the model
3. Select annotators who share the linguistic and cultural norms of the data under analysis
4. Document your analytical choices and motivate them
5. Use annotation schemes that can be changed or expanded at a second time (modularity and consistency)

Countermeasures to label bias

1. Aggregate multiple annotations to even out results
2. Use knowledge about disagreement as information to update the model
3. Select annotators who share the linguistic and cultural norms of the data under analysis
4. Document your analytical choices and motivate them
5. Use annotation schemes that can be changed or expanded at a second time (modularity and consistency)

Model applications

Models use cases

1. Feature-based models depend on choices taken by the researcher
    - Hypothesis testing
2. Representation learning abstracts a pattern from data
    - Hypothesis generation

Models use cases

1. Feature-based models depend on choices taken by the researcher
    - Hypothesis testing
2. Representation learning abstracts a pattern from data
    - Hypothesis generation

Models use cases

1. Feature-based models depend on choices taken by the researcher
    - Hypothesis testing
2. Representation learning abstracts a pattern from data
    - Hypothesis generation

Models use cases

1. Feature-based models depend on choices taken by the researcher
    - Hypothesis testing
2. Representation learning abstracts a pattern from data
    - Hypothesis generation

Not all bias is created equal...

1. Hypothesis testing
    - **Bias needs to be corrected**
2. Hypothesis generation
    - **Bias is a source of information**
    - This is a form of AI-driven data analysis

Scenario: Gender inclusive language use

Goal

Analyzing the patterns of use of gender inclusive language in a community of practice where 90% of participants are cis men.

Hypothesis generation: who uses gender inclusive forms, which forms are used (pronouns, suffixes, etc.), which context (social, grammatical, etc.) triggers the use?

- Imbalance in the amount of forms is expected and retained by design
- Imbalance in the composition of participants is informative of the community of practice
- Correcting bias would erase information about the linguistic phenomenon and the community of practice

Scenario: Gender inclusive language use

Goal

Analyzing the patterns of use of gender inclusive language in a community of practice where 90% of participants are cis men.

Hypothesis generation: who uses gender inclusive forms, which forms are used (pronouns, suffixes, etc.), which context (social, grammatical, etc.) triggers the use?

- Imbalance in the amount of forms is expected and retained by design
- Imbalance in the composition of participants is informative of the community of practice
- Correcting bias would erase information about the linguistic phenomenon and the community of practice

Scenario: Gender inclusive language use

Goal

Analyzing the patterns of use of gender inclusive language in a community of practice where 90% of participants are cis men.

Hypothesis generation: who uses gender inclusive forms, which forms are used (pronouns, suffixes, etc.), which context (social, grammatical, etc.) triggers the use?

- Imbalance in the amount of forms is expected and retained by design
- Imbalance in the composition of participants is informative of the community of practice
- Correcting bias would erase information about the linguistic phenomenon and the community of practice

Scenario: Gender inclusive language use

Goal

Analyzing the patterns of use of gender inclusive language in a community of practice where 90% of participants are cis men.

Hypothesis generation: who uses gender inclusive forms, which forms are used (pronouns, suffixes, etc.), which context (social, grammatical, etc.) triggers the use?

- Imbalance in the amount of forms is expected and retained by design
- Imbalance in the composition of participants is informative of the community of practice
- Correcting bias would erase information about the linguistic phenomenon and the community of practice

Scenario: Gender inclusive language use

Goal

Analyzing the patterns of use of gender inclusive language in a community of practice where 90% of participants are cis men.

Hypothesis generation: who uses gender inclusive forms, which forms are used (pronouns, suffixes, etc.), which context (social, grammatical, etc.) triggers the use?

- Imbalance in the amount of forms is expected and retained by design
- Imbalance in the composition of participants is informative of the community of practice
- Correcting bias would erase information about the linguistic phenomenon and the community of practice

Scenario: Gender inclusive language use

Goal

Analyzing the patterns of use of gender inclusive language in a community of practice where 90% of participants are cis men.

Hypothesis testing: usage conditions for the existing noun suffixes (-*, -ə, -x, -@)

- A minimal amount of data points is needed to pick up the variables governing the pattern
- Dataset size has an effect over how each variable contributes
- Correcting bias is essential to ensure the model is capturing all factors at play

Scenario: Gender inclusive language use

Goal

Analyzing the patterns of use of gender inclusive language in a community of practice where 90% of participants are cis men.

Hypothesis testing: usage conditions for the existing noun suffixes (-*, -ə, -x, -@)

- A minimal amount of data points is needed to pick up the variables governing the pattern
- Dataset size has an effect over how each variable contributes
- Correcting bias is essential to ensure the model is capturing all factors at play

Scenario: Gender inclusive language use

Goal

Analyzing the patterns of use of gender inclusive language in a community of practice where 90% of participants are cis men.

Hypothesis testing: usage conditions for the existing noun suffixes (-*, -ə, -x, -@)

- A minimal amount of data points is needed to pick up the variables governing the pattern
- Dataset size has an effect over how each variable contributes
- Correcting bias is essential to ensure the model is capturing all factors at play

Scenario: Gender inclusive language use

Goal

Analyzing the patterns of use of gender inclusive language in a community of practice where 90% of participants are cis men.

Hypothesis testing: usage conditions for the existing noun suffixes (-*, -ə, -x, -@)

- A minimal amount of data points is needed to pick up the variables governing the pattern
- Dataset size has an effect over how each variable contributes
- Correcting bias is essential to ensure the model is capturing all factors at play

Scenario: Gender inclusive language use

Goal

Analyzing the patterns of use of gender inclusive language in a community of practice where 90% of participants are cis men.

Hypothesis testing: usage conditions for the existing noun suffixes (-*, -ə, -x, -@)

- A minimal amount of data points is needed to pick up the variables governing the pattern
- Dataset size has an effect over how each variable contributes
- Correcting bias is essential to ensure the model is capturing all factors at play

Thank you!

References

McCarthy, John et al. (2006). **“A proposal for the dartmouth summer research project on artificial intelligence, august 31, 1955”.** In: *AI magazine* 27.4, pp. 12–12.

Turing, Alan (2021). **“Computing machinery and intelligence (1950)”.** In: *Mind* 59.236, pp. 33–60.

Weizenbaum, Joseph (1966). **“ELIZA—a computer program for the study of natural language communication between man and machine”.** In: *Communications of the ACM* 9.1, pp. 36–45.