University of Bamberg





https://tinyurl.com/bambergtranslate



Natural Language Processing Research in the Era of Al

Can Computers Read Between the Lines?

Inaugural Lecture, November 18, 2025 Roman Klinger roman.klinger@uni-bamberg.de

romanklinger.de in romanklinger https://www.bamberg.de/nlproc/ https://www.romanklinger.de/





Bas heift ju welchem Ende ftudiert man Universalgeschichte? Eine Academische Antriterede ben Eroffnung feiner Borlefungen gehalten non Ariedrich Schiller. Professor ber Beidichte in Jena. Jena, In ber Mcademifchen Buchhandlung. 1789.

"What Is, and to What End do We Study, Universal History?"



Was heist ju welchem Ende ftudiert man Universalgeschichte? Eine Academifche Untritgrede ben Eroffnung feiner Borlefungen gehalten non Ariedrich Schiller. Professor ber Beidichte in Jena. Jena, In ber Mcademifchen Buchhandlung. 1789.

"What Is, and to What End do We Study, Universal History?"

Schiller contrasts two types of students:



Was heift ju welchem Ende ftudiert man Universalgeschichte? Eine Academifche Untritgrede ben Eroffnung feiner Borlefungen gehalten Ariedrich Schiller. Professor ber Beidichte in Jena. Jena, In ber Academifchen Buchhandlung. 1789.

"What Is, and to What End do We Study, Universal History?"

- Schiller contrasts two types of students:
 - Brotgelehrter ("bread scholar") who cares about practical advantage and career.



Was heift ju welchem Ende ftudiert man Universalaeschichte? Eine Academifche Untritgrede ben Eroffnung feiner Borlefungen gehalten Ariedrich Schiller. Professor ber Beidichte in Jena. Jena, In ber Academifchen Buchhandlung. 1789.

"What Is, and to What End do We Study, Universal History?"

- Schiller contrasts two types of students:
 - Brotgelehrter ("bread scholar") who cares about practical advantage and career.
 - Philosophischer Kopf ("philosophical mind") who seeks knowledge and human betterment.



What do we do today?





Natural Language Processing Research in the Era of Al

Can Computers Read Between the Lines?

Inaugural Lecture, November 18, 2025

Roman Klinger roman.klinger@uni-bamberg.de

romanklinger.de in romanklinger https://www.bamberg.de/nlproc/



What do we do today?





Natural Language Processing Research in the Era of Al

Can Computers Read Between the Lines?

Inaugural Lecture, November 18, 2025

Roman Klinger roman.klinger@uni-bamberg.de

romanklinger.de in romanklinger https://www.bamberg.de/nlproc/

 What is, and why do we study, Natural Language Processing?



What do we do today?





Natural Language Processing Research in the Era of Al

Can Computers Read Between the Lines?

Inaugural Lecture, November 18, 2025

Roman Klinger roman.klinger@uni-bamberg.de

romanklinger.de in romanklinger https://www.bamberg.de/nlproc/

- What is, and why do we study, Natural Language Processing?
- How does this research field change in the era of "AI"?



Outline

1 Language Understanding

2 What is "AI"?

3 What we work on

4 Wrap Up

Outline

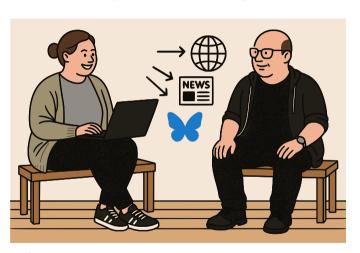
1 Language Understanding

2 What is "AI"

3 What we work on

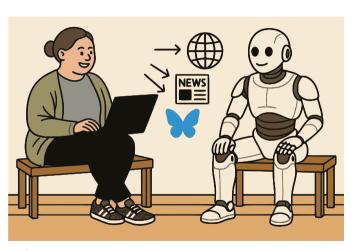
4 Wrap Up







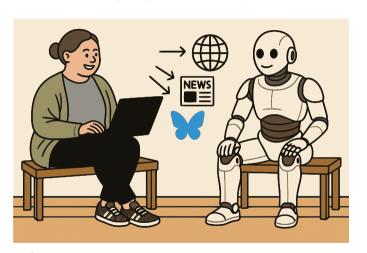






Language Understanding

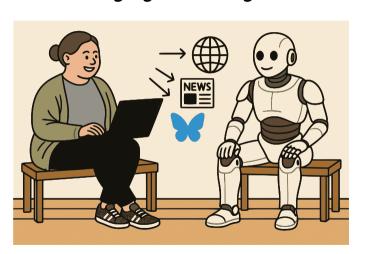




• We study how machines can understand human language



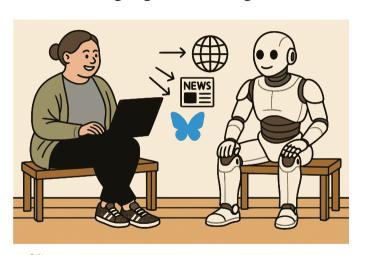




- We study how machines can understand human language
- We focus on written text







- We study how machines can understand human language
- We focus on written text



Language Understanding







What does it mean to "understand"?





- What does it mean to "understand"?
- Cambridge Dictionary: "to know the meaning of something that someone says"





- What does it mean to "understand"?
- Cambridge Dictionary: "to know the meaning of something that someone says"
- How can we make computers understand?





- What does it mean to "understand"?
- Cambridge Dictionary: "to know the meaning of something that someone says"
- How can we make computers understand?
- How can we measure if we are successful?





- What does it mean to "understand"?
- Cambridge Dictionary: "to know the meaning of something that someone says"
- How can we make computers understand?
- How can we measure if we are successful?
- How and what for can we make use of the meaning that has been understood?







Desideratum: A machine that understands language as humans do?







- Desideratum: A machine that understands language as humans do?
- How to study language in its entirety? (universal language understanding ability)



Language Understanding





- Desideratum: A machine that understands language as humans do?
- How to study language in its entirety? (universal language understanding ability)
- We study particular phenomena.







- Desideratum: A machine that understands language as humans do?
- How to study language in its entirety? (universal language understanding ability)
- We study particular phenomena.
- We define concrete tasks to solve.







- Desideratum: A machine that understands language as humans do?
- How to study language in its entirety? (universal language understanding ability)
- We study particular phenomena.
- We define concrete tasks to solve.
- \Rightarrow Pragmatic approach to language understanding





World News > United States



Donald Trump mocks Bill Gates after billionaire's humiliating backflip on climate change

30 October 2025 - 04:46&M sky news .com.au

Donald Trump wasted no time mocking Bill Gates after the billionaire admitted that climate change will not lead to humanity's demise in a memo released vesterday.



Language Understanding



"Donald Trump mocks Bill Gates after billionaire's humiliating backflip on climate change"

Sky News Australia, Oct 30, 2025





"Donald Trump mocks Bill Gates after billionaire's humiliating backflip on climate change"

Sky News Australia, Oct 30, 2025

What information is in this sentence that's worth understanding?





"Donald Trump mocks Bill Gates after billionaire's humiliating backflip on climate change"

Sky News Australia, Oct 30, 2025

What information is in this sentence that's worth understanding?

• Find entity names:





"Donald Trump mocks Bill Gates after billionaire's humiliating backflip on climate change"

Sky News Australia, Oct 30, 2025

What information is in this sentence that's worth understanding?

• Find entity names: Donald Trump; Bill Gates





"Donald Trump mocks Bill Gates after billionaire's humiliating backflip on climate change"

Sky News Australia, Oct 30, 2025

What information is in this sentence that's worth understanding?

- Find entity names: Donald Trump; Bill Gates
- Recognize the sentiment:





"Donald Trump mocks Bill Gates after billionaire's humiliating backflip on climate change"

Sky News Australia, Oct 30, 2025

What information is in this sentence that's worth understanding?

• Find entity names: Donald Trump; Bill Gates

• Recognize the sentiment: negative





"Donald Trump mocks Bill Gates after billionaire's humiliating backflip on climate change"

Sky News Australia, Oct 30, 2025

What information is in this sentence that's worth understanding?

- Find entity names: Donald Trump; Bill Gates
- Recognize the sentiment: negative
- Topic:





"Donald Trump mocks Bill Gates after billionaire's humiliating backflip on climate change"

Sky News Australia, Oct 30, 2025

What information is in this sentence that's worth understanding?

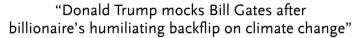
• Find entity names: Donald Trump; Bill Gates

• Recognize the sentiment: negative

Topic: climate change







Sky News Australia, Oct 30, 2025

- Find entity names: Donald Trump; Bill Gates
- Recognize the sentiment: negative
- Topic: climate change
- Stances:





"Donald Trump mocks Bill Gates after billionaire's humiliating backflip on climate change"

Sky News Australia, Oct 30, 2025

- Find entity names: Donald Trump; Bill Gates
- Recognize the sentiment: negative
- Topic: climate change
- Stances: Bill Gates → opinion change on climate change.





"Donald Trump mocks Bill Gates after billionaire's humiliating backflip on climate change"

Sky News Australia, Oct 30, 2025

- Find entity names: Donald Trump; Bill Gates
- Recognize the sentiment: negative
- Topic: climate change
- Stances: Bill Gates → opinion change on climate change.
- Relation:





"Donald Trump mocks Bill Gates after billionaire's humiliating backflip on climate change"

Sky News Australia, Oct 30, 2025

- Find entity names: Donald Trump; Bill Gates
- Recognize the sentiment: negative
- Topic: climate change
- Stances: Bill Gates → opinion change on climate change.
- Relation: Donald Trump → negative opinion (Bill Gates)





"Donald Trump mocks Bill Gates after billionaire's humiliating backflip on climate change"

Sky News Australia, Oct 30, 2025

What information is in this sentence that's worth understanding?

- Find entity names: Donald Trump; Bill Gates
- Recognize the sentiment: negative
- Topic: climate change
- Stances: Bill Gates → opinion change on climate change.
- Relation: Donald Trump → negative opinion (Bill Gates)
- ...





"Donald Trump mocks Bill Gates after billionaire's humiliating backflip on climate change"

Sky News Australia, Oct 30, 2025

What information is in this sentence that's worth understanding?

- Find entity names: Donald Trump; Bill Gates
- Recognize the sentiment: negative
- Topic: climate change
- Stances: Bill Gates → opinion change on climate change.
- Relation: Donald Trump → negative opinion (Bill Gates)
- ...

Language Understanding

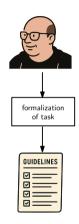
 Aggregating information enables many use cases: diverse news recommendation, social network analysis, opinion mining, ...





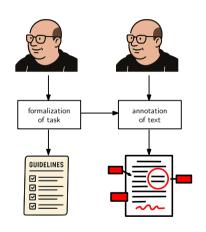






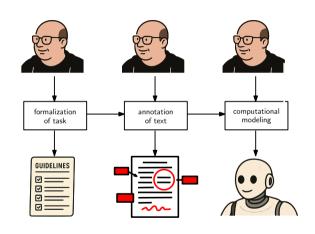






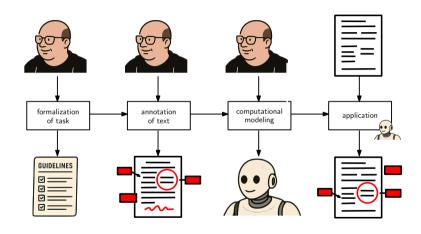






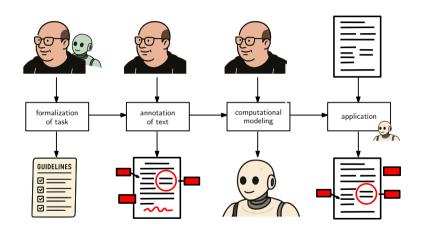






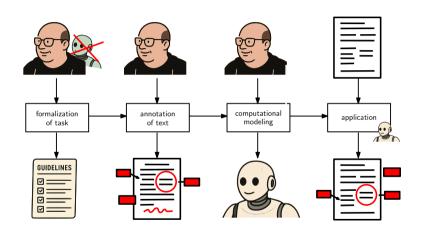






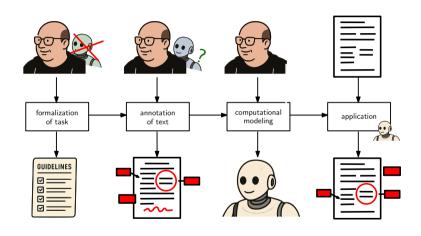






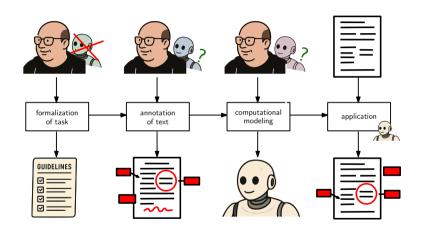






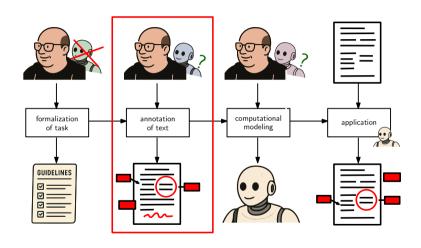










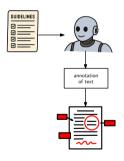




Do we still need the human?

(except for the specification of the task)



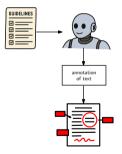




Do we still need the human?

(except for the specification of the task)

Language Understanding



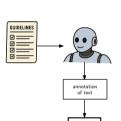
• If an "AI" can just solve the task by being instructed, we don't need manually annotated data, and we don't need to spend a lot of effort in model development.

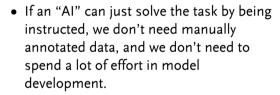


Do we still need the human?

(except for the specification of the task)

Language Understanding





In other words:
 Does Bamberg need a Natural Language
 Processing professor?





 Language Understanding
 What is "Al"?
 What we work on occooo

 00000000●
 000000
 0000000

That is a reasonable question! We should reflect on our role.







A lot of research goes into:





A lot of research goes into:

1. Formalizing concepts





A lot of research goes into:

- 1. Formalizing concepts
- 2. Collecting and annotating corpora for model development





A lot of research goes into:

- 1. Formalizing concepts
- 2. Collecting and annotating corpora for model development
- 3. Advancing machine learning methods to better learn to reconstruct the annotations





A lot of research goes into:

- 1. Formalizing concepts
- 2. Collecting and annotating corpora for model development
- 3. Advancing machine learning methods to better learn to reconstruct the annotations
- 4. Evaluation methods





A lot of research goes into:

- 1. Formalizing concepts
- 2. Collecting and annotating corpora for model development
- 3. Advancing machine learning methods to better learn to reconstruct the annotations
- 4. Evaluation methods

If "AI" can do these steps sufficiently well, our work is essentially done.

Is that the case?



Outline

1 Language Understanding

2 What is "AI"?

3 What we work on

4 Wrap Up





OTTO OTTO

• To answer that, we need to understand what "Artificial Intelligence" means.





- To answer that, we need to understand what "Artificial Intelligence" means.
- The term encompasses many different topics.





- To answer that, we need to understand what "Artificial Intelligence" means.
- The term encompasses many different topics.
- Many agree that an important aspect is: a system behaves in a way that appears to be intelligent.





- To answer that, we need to understand what "Artificial Intelligence" means.
- The term encompasses many different topics.
- Many agree that an important aspect is:
 a system behaves in a way that appears to be intelligent.
- Methods:





- To answer that, we need to understand what "Artificial Intelligence" means.
- The term encompasses many different topics.
- Many agree that an important aspect is:
 a system behaves in a way that appears to be intelligent.
- Methods:
 - Logical reasoning





- To answer that, we need to understand what "Artificial Intelligence" means.
- The term encompasses many different topics.
- Many agree that an important aspect is:
 a system behaves in a way that appears to be intelligent.
- Methods:
 - Logical reasoning
 - Machine learning





- To answer that, we need to understand what "Artificial Intelligence" means.
- The term encompasses many different topics.
- Many agree that an important aspect is:
 a system behaves in a way that appears to be intelligent.
- Methods:
 - Logical reasoning
 - Machine learning
 - Semantic language representations





- To answer that, we need to understand what "Artificial Intelligence" means.
- The term encompasses many different topics.
- Many agree that an important aspect is:
 a system behaves in a way that appears to be intelligent.
- Methods:
 - Logical reasoning
 - Machine learning
 - Semantic language representations
 - Search and planning





- To answer that, we need to understand what "Artificial Intelligence" means.
- The term encompasses many different topics.
- Many agree that an important aspect is:
 a system behaves in a way that appears to be intelligent.
- Methods:
 - Logical reasoning
 - Machine learning
 - Semantic language representations
 - · Search and planning
 - Integration of world knowledge





- To answer that, we need to understand what "Artificial Intelligence" means.
- The term encompasses many different topics.
- Many agree that an important aspect is:
 a system behaves in a way that appears to be intelligent.
- Methods:
 - Logical reasoning
 - Machine learning
 - Semantic language representations
 - · Search and planning
 - Integration of world knowledge
 - ...





- To answer that, we need to understand what "Artificial Intelligence" means.
- The term encompasses many different topics.
- Many agree that an important aspect is:
 a system behaves in a way that appears to be intelligent.
- Methods:
 - Logical reasoning
 - Machine learning
 - Semantic language representations
 - · Search and planning
 - Integration of world knowledge
 - ...
- ⇒ The term "Artificial Intelligence" does NOT refer to one technology, but to many.











What word fits?

1. "I came here by



Language Models

- 1. "I came here by ______'
- 2. "The sentence 'It's my birthday party' expresses _______'





- 1. "I came here by
- 2. "The sentence 'It's my birthday party' expresses ______'
- 3. "You need to go through the _____





- 1. "I came here by
- 2. "The sentence 'It's my birthday party' expresses _____
- "You need to go through the ______ to reach 'Untere Sandstrasse' from the Island in Bamberg."





- 1. "I came here by
- 2. "The sentence 'It's my birthday party' expresses ______'
- "You need to go through the ______ to reach 'Untere Sandstrasse' from the Island in Bamberg."
- Language models represent probabilities of words in context.



		OTTO OTTO OTTO OTTO OTTO OTTO OTTO OTT
--	--	---

- 1. "I came here by
- 2. "The sentence 'It's my birthday party' expresses
- 3. "You need to go through the to reach 'Untere Sandstrasse' from the Island in Bamberg."
- Language models represent probabilities of words in context.
- Language models can be used for different types of tasks





- " sentence completion 1. "I came here by
- 2. "The sentence 'It's my birthday party' expresses
- 3. "You need to go through the to reach 'Untere Sandstrasse' from the Island in Bamberg."
- Language models represent probabilities of words in context.
- Language models can be used for different types of tasks





	amplette
What word fits?	ance con

1. "I came here by _____ "\senter

2. "The sentence 'It's my birthday party' expresses

- 3. "You need to go through the to reach 'Untere Sandstrasse' from the Island in Bamberg."
- Language models represent probabilities of words in context.
- Language models can be used for different types of tasks



text classification



	ampleti	no.
Vhat word fits?	"Sentence completion	at classification
1. "I came here by	"serre	tex
2. "The sentence 'It's my bird	thday party' expresses	knowledge retrieval
3. "You need to go through t	he to reach	Whedge !
'Untere Sandstrasse' from	the Island in Bamberg."	kno

- Language models represent probabilities of words in context.
- Language models can be used for different types of tasks







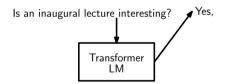


Is an inaugural lecture interesting?

Transformer
LM

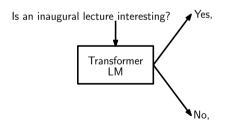








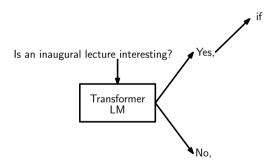






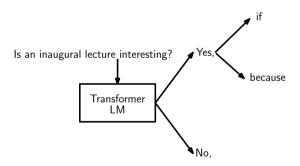
Wrap up 000000 4/VERS/7,







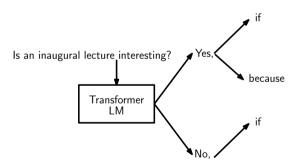






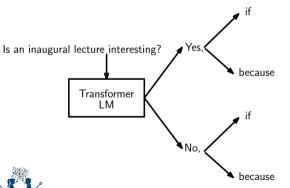
Wrap Up 000000



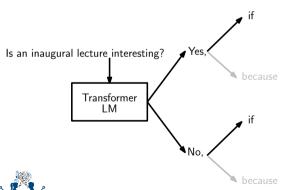




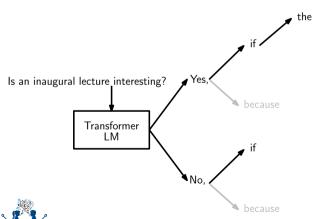






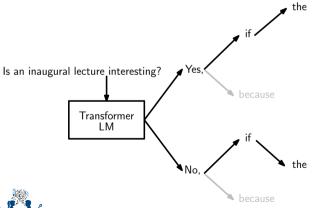




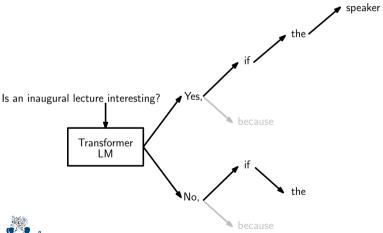


Wrap Up 000000 AIVERSITY BICH-UA



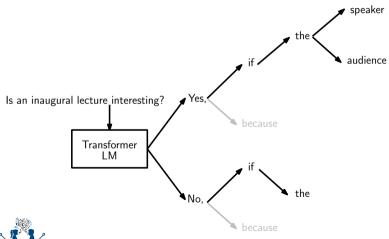






Roman Klinger

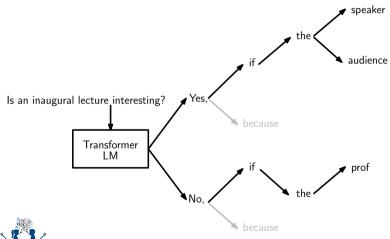




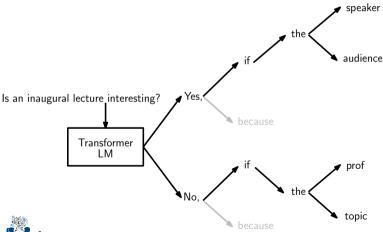
Roman Klinger

Roman Klinger

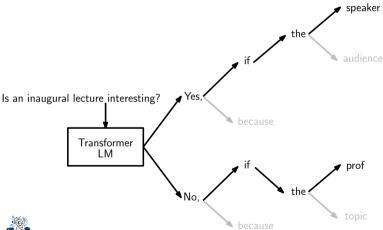




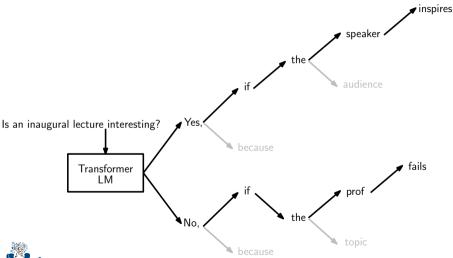










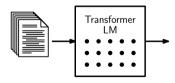








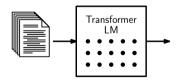
• Large Language Models (LLMs) have billions of parameters ...







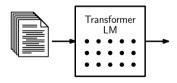
- Large Language Models (LLMs) have billions of parameters ...
- ...and see huge amounts of text to find values for them.





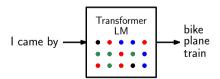


- Large Language Models (LLMs) have billions of parameters ...
- ...and see huge amounts of text to find values for them.
- The main optimization steps are:





- Large Language Models (LLMs) have billions of parameters ...
- ...and see huge amounts of text to find values for them.
- The main optimization steps are:
 - Language modeling objective

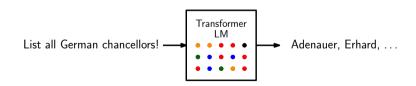








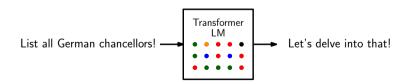
- Large Language Models (LLMs) have billions of parameters ...
- ...and see huge amounts of text to find values for them.
- The main optimization steps are:
 - Language modeling objective
 - Instruction-tuning







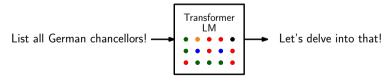
- Large Language Models (LLMs) have billions of parameters ...
- ...and see huge amounts of text to find values for them.
- The main optimization steps are:
 - Language modeling objective
 - Instruction-tuning
 - Alignment through human feedback







- Large Language Models (LLMs) have billions of parameters ...
- ...and see huge amounts of text to find values for them.
- The main optimization steps are:
 - Language modeling objective
 - Instruction-tuning
 - Alignment through human feedback
- Finding the best output requires logical reasoning, semantic understanding, efficient search, and world knowledge!











• To find out, we need to have data!





- To find out, we need to have data!
- We want to challenge these models.





- To find out, we need to have data!
- We want to challenge these models.
- We are particularly interested in tasks that we consider also challenging for humans.





- To find out, we need to have data!
- We want to challenge these models.
- We are particularly interested in tasks that we consider also challenging for humans.
- We will now look at three cases in which prompting language models does not lead to the desired outcome.



Outline

1 Language Understanding

2 What is "AI"

3 What we work on

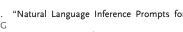
4 Wrap Up





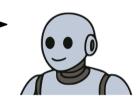
F. M. Plaza-del-Arco, M.-T. Martín-Valdivia, and R. Klinger (2022). "Natural Language Inference Prompts for Zero-shot Emotion Classification in Text across Corpora". In: COLING

Fundamentals of Natural Language Processing





"You are an expert in emotion annotations. Which emotion does the author of the following text feel? Select one from sadness, joy, anger, fear, surprise, disgust."



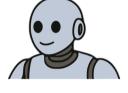






"You are an expert in emotion annotations. Which emotion does the author of the following text feel? Select one from sadness, joy, anger, fear, surprise, disgust."

I needed to cry



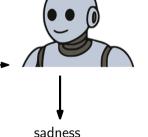






"You are an expert in emotion annotations. Which emotion does the author of the following text feel? Select one from sadness, joy, anger, fear, surprise, disgust."

I needed to cry







F. M. Plaza-del-Arco, M.-T. Martín-Valdivia, and R. Klinger (2022). "Natural Language Inference Prompts for Zero-shot Emotion Classification in Text across Corpora". In: COLING Fundamentals of Natural Language Processing



"You are an expert in emotion annotations. Which emotion does the author of the following text feel? Select one from sadness, joy, anger, fear, surprise, disgust." I needed to cry

• Emotion prediction performance .6–.7 F1 score.

Fundamentals of Natural Language Processing



F. M. Plaza-del-Arco, M.-T. Martín-Valdivia, and R. Klinger (2022). "Natural Language Inference Prompts for Zero-shot Emotion Classification in Text across Corpora". In: COLING



sadness

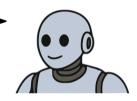
Emotion Detection in Arguments





L. Greschner and R. Klinger (2025). "Fearful Falcons and Angry Llamas: Emotion Category Annotations of Arguments by Humans and LLMs". In: NLP4DH







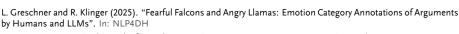




We should ban plastic bottles, because they harm the environment.

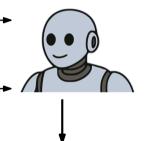








We should ban plastic bottles, because they harm the environment.





L. Greschner and R. Klinger (2025). "Fearful Falcons and Angry Llamas: Emotion Category Annotations of Arguments by Humans and LLMs". In: NLP4DH





We should ban plastic bottles, because they harm the environment.

Models lack access to context.







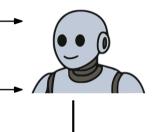
L. Greschner and R. Klinger (2025). "Fearful Falcons and Angry Llamas: Emotion Category Annotations of Arguments by Humans and LLMs". In: NLP4DH

Emotion Detection in Arguments



"You are an expert in emotion annotation in arguments. Which emotion does the reader of the following argument likely feel?"

We should ban plastic bottles, because they harm the environment.



- Models lack access to context.
- They tend to predict fear or anger.



L. Greschner and R. Klinger (2025). "Fearful Falcons and Angry Llamas: Emotion Category Annotations of Arguments by Humans and LLMs". In: NLP4DH



How can we conduct contextualized emotion detection (and convincingness assessment) in arguments?







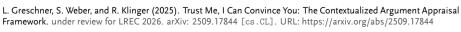
How can we conduct contextualized emotion detection (and convincingness assessment) in arguments?



 We need to know for whom we make predictions!







How can we conduct contextualized emotion detection (and convincingness assessment) in arguments?



- We need to know for whom we make predictions!
- That is a challenge, we need annotated arguments with information about the annotator







- We need to know for whom we make predictions!
- That is a challenge, we need annotated arguments with information about the annotator
- How to get such data?





How can we conduct contextualized emotion detection (and convincingness assessment) in arguments?

- We need to know for whom we make predictions!
- That is a challenge, we need annotated arguments with information about the annotator
- How to get such data?
- We asked people to role play a debate and annotate arguments they read.





How can we conduct contextualized emotion detection (and convincingness assessment) in arguments?



- We need to know for whom we make predictions!
- That is a challenge, we need annotated arguments with information about the annotator
- How to get such data?
- We asked people to role play a debate and annotate arguments they read.





How can we conduct contextualized emotion detection (and convincingness assessment) in arguments?







How can we conduct contextualized emotion detection (and convincingness assessment) in arguments?











What we work on 0000000

How can we conduct contextualized emotion detection (and convincingness assessment) in arguments?











How can we conduct contextualized emotion detection (and convincingness assessment) in arguments?







• We need to develop and evaluate methods to collect data in context.







How can we conduct contextualized emotion detection (and convincingness assessment) in arguments?







- We need to develop and evaluate methods to collect data in context.
- Who's the speaker? Who's the listener?







How can we conduct contextualized emotion detection (and convincingness assessment) in arguments?







- We need to develop and evaluate methods to collect data in context.
- Who's the speaker? Who's the listener?
- We need to develop methods to integrate

contextual information in computational models.





Wrap Up

For whom do models make predictions?





J. Schäfer et al. (2025). "Which Demographics do LLMs Default to During Annotation?" In: ACL

For whom do models make predictions?

 If we don't tell the model for whom it should make a prediction, with whose annotations is it best aligned?





For whom do models make predictions?



• If we don't tell the model for whom it should make a prediction, with whose annotations is it best aligned?









• If we don't tell the model for whom it should make a prediction, with whose annotations is it best aligned?









• If we don't tell the model for whom it should make a prediction, with whose annotations is it best aligned?







- If we don't tell the model for whom it should make a prediction, with whose annotations is it best aligned?
- Models best reconstruct a person's annotation when they are white, comparably young, and male.









- If we don't tell the model for whom it should make a prediction, with whose annotations is it best aligned?
- Models best reconstruct a person's annotation when they are white, comparably young, and male.
- We need to understand biases and make models work well for everybody.







Human annotation varies - Should LLM's annotation also vary?







• You are an expert in emotion annotation. The label set is {LS}. The instance to classify is "The dog ran towards me."





- You are an expert in emotion annotation. The label set is {LS}. The instance to classify is "The dog ran towards me."
- anger, fear, joy, disgust, sadness, surprise.





- You are an expert in emotion annotation. The label set is {LS}. The instance to classify is "The dog ran towards me."
- anger, fear, joy, disgust, sadness, surprise.
- surprise, sadness, disgust, joy, fear, anger.



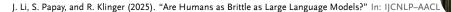




- You are an expert in emotion annotation. The label set is {LS}. The instance to classify is "The dog ran towards me."
- anger, fear, joy, disgust, sadness, surprise.

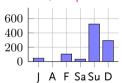
- surprise, sadness, disgust, joy, fear, anger.
- angr, feer, joy, disgst, sadnes, suprise.







- You are an expert in emotion annotation. The label set is {LS}. The instance to classify is "The dog ran towards me."
- anger, fear, joy, disgust, sadness, surprise.

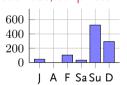


- surprise, sadness, disgust, joy, fear, anger.
- angr, feer, joy, disgst, sadnes, suprise.





- You are an expert in emotion annotation. The label set is {LS}. The instance to classify is "The dog ran towards me."
- anger, fear, joy, disgust, sadness, surprise.



 surprise, sadness, disgust, joy, fear, anger.

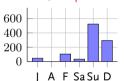


• angr, feer, joy, disgst, sadnes, suprise.





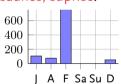
- You are an expert in emotion annotation. The label set is {LS}. The instance to classify is "The dog ran towards me."
- anger, fear, joy, disgust, sadness, surprise.



• surprise, sadness, disgust, joy, fear, anger.



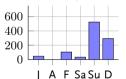
 angr, feer, joy, disgst, sadnes, suprise.







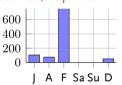
- You are an expert in emotion annotation. The label set is {LS}. The instance to classify is "The dog ran towards me."
- anger, fear, joy, disgust, sadness, surprise.



• surprise, sadness, disgust, joy, fear, anger.



 angr, feer, joy, disgst, sadnes, suprise.

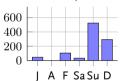


• Human susceptibility to prompt changes differs from LLM's brittleness.

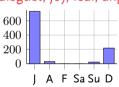




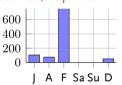
- You are an expert in emotion annotation. The label set is {LS}. The instance to classify is "The dog ran towards me."
- anger, fear, joy, disgust, sadness, surprise.



 surprise, sadness, disgust, joy, fear, anger.



• angr, feer, joy, disgst, sadnes, suprise.

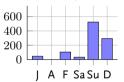


- Human susceptibility to prompt changes differs from LLM's brittleness.
- Do we want model's outputs to vary as human's output does?

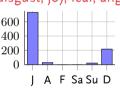




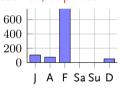
- You are an expert in emotion annotation. The label set is {LS}. The instance to classify is "The dog ran towards me."
- anger, fear, joy, disgust, sadness, surprise.



• surprise, sadness, disgust, joy, fear, anger.



 angr, feer, joy, disgst, sadnes, suprise.



- Human susceptibility to prompt changes differs from LLM's brittleness.
- Do we want model's outputs to vary as human's output does?
- If yes, how to achieve that? If no, what should they do?





Outline

1 Language Understanding

2 What is "AI"

3 What we work on

4 Wrap Up

What do we do today?





Natural Language Processing Research in the Era of Al

Can Computers Read Between the Lines?

Inaugural Lecture, November 18, 2025

Roman Klinger roman.klinger@uni-bamberg.de

romanklinger.de in romanklinger https://www.bamberg.de/nlproc/

- What is, and why do we study, Natural Language Processing?
- How does this research field change in the era of "AI"?





What is, and why do we study, Natural Language Processing?





- What is, and why do we study, Natural Language Processing?
 - NLP is a research area that aims at understanding how to capture language meaning and structure computationally.





- What is, and why do we study, Natural Language Processing?
 - NLP is a research area that aims at understanding how to capture language meaning and structure computationally.
 - Automatically understanding language helps to destill information from large amounts of text, more than anybody could ever read.



- TO CHOOS TO SHARE THE SHAR
- What is, and why do we study, Natural Language Processing?
 - NLP is a research area that aims at understanding how to capture language meaning and structure computationally.
 - Automatically understanding language helps to destill information from large amounts of text, more than anybody could ever read.
 - Models do not always behave as expected and not for everybody.





- What is, and why do we study, Natural Language Processing?
 - NLP is a research area that aims at understanding how to capture language meaning and structure computationally.
 - Automatically understanding language helps to destill information from large amounts of text, more than anybody could ever read.
 - Models do not always behave as expected and not for everybody.
 - To understand what machines should do, we need to understand human language in context.
- How does this research field change in the era of "AI"?





- What is, and why do we study, Natural Language Processing?
 - NLP is a research area that aims at understanding how to capture language meaning and structure computationally.
 - Automatically understanding language helps to destill information from large amounts of text, more than anybody could ever read.
 - Models do not always behave as expected and not for everybody.
 - To understand what machines should do, we need to understand human language in context.
- How does this research field change in the era of "AI"?
 - LLMs do not (yet) just solve tasks "off the shelf" ...



- ON THE PROPERTY OF THE PROPERT
- What is, and why do we study, Natural Language Processing?
 - NLP is a research area that aims at understanding how to capture language meaning and structure computationally.
 - Automatically understanding language helps to destill information from large amounts of text, more than anybody could ever read.
 - Models do not always behave as expected and not for everybody.
 - To understand what machines should do, we need to understand human language in context.
- How does this research field change in the era of "AI"?
 - LLMs do not (yet) just solve tasks "off the shelf" ...
 - ...but they show promising results for many tasks.



- OTTO OTTO
- What is, and why do we study, Natural Language Processing?
 - NLP is a research area that aims at understanding how to capture language meaning and structure computationally.
 - Automatically understanding language helps to destill information from large amounts of text, more than anybody could ever read.
 - Models do not always behave as expected and not for everybody.
 - To understand what machines should do, we need to understand human language in context.
- How does this research field change in the era of "AI"?
 - LLMs do not (yet) just solve tasks "off the shelf" ...
 - ...but they show promising results for many tasks.
 - Systems relying on more traditional methods often work better.



Wrap Up 000●00







• Are we bread scholars?





- Are we bread scholars?
 - We do teach how to practically build systems that can understand language. (so...; Yes?)





Are we bread scholars?

- We do teach how to practically build systems that can understand language. (so...; Yes?)
- To advance the field, however, a deeper understanding of human expectations, their behaviour, their properties that influence language interpretation and production is needed. (so...; No?)





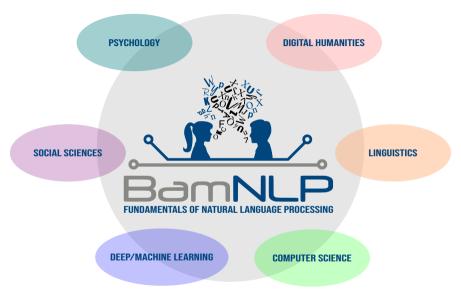
- Are we bread scholars?
 - We do teach how to practically build systems that can understand language. (so...; Yes?)
 - To advance the field, however, a deeper understanding of human expectations, their behaviour, their properties that influence language interpretation and production is needed. (so...: No?)
- Can we be philosophical minds?





- Are we bread scholars?
 - We do teach how to practically build systems that can understand language. (so...; Yes?)
 - To advance the field, however, a deeper understanding of human expectations, their behaviour, their properties that influence language interpretation and production is needed. (so...: No?)
- Can we be philosophical minds?
 - We need to be!

















- All research groups I was part of so far and all collaborators.
- Joy Kearney, Mona Wirth and Romy Schatzschneider for the organization.
- All of you for your interest!
- Please reach out if you want to talk, chat, discuss, meet us, drink coffee, work with us, collaborate, ...



University of Bamberg





https://tinyurl.com/bambergtranslate



Natural Language Processing Research in the Era of Al

Can Computers Read Between the Lines?

Inaugural Lecture, November 18, 2025 Roman Klinger roman.klinger@uni-bamberg.de

romanklinger.de in romanklinger https://www.bamberg.de/nlproc/ https://www.romanklinger.de/