

ARTIFICIAL INTELLIGENCE


DANIEL BRAUN

D.BRAUN@UTWENTE.NL

FACULTY OF BEHAVIOURAL, MANAGEMENT AND SOCIAL SCIENCES (BMS)

UNIVERSITY OF TWENTE

25.04.2024



TOPICS

- What is AI?
- How do ChatGPT and Co work?
- AI tools
- Dangers of AI



Dr. Daniel Braun

Assistant Professor

@d4br4 (Twitter / Mastodon.social / Bluesky)
d.braun@utwente.nl
www.daniel-braun.science

University of Twente

Section Industrial Engineering and Business Information
Systems (IEBIS)
Department of High-tech Business and Entrepreneurship
(HBE)



Do you think AI could be helpful in your worklife?



1

Yes

94% 51

2

No


6% 3



WHAT IS AI?

- **Nobody knows!**
- The EU settled for:
'artificial intelligence system' (AI system) means a machine-based system that is designed to operate with varying levels of autonomy and that can, for explicit or implicit objectives, generate outputs such as predictions, recommendations, or decisions, that influence physical or virtual environments
- My take: A system that can perform a task for which a human would need intelligence
- Problem: What is intelligence?

RULE-BASED AI VS MACHINE LEARNING



- Rule-based:
 - Knowledge from human experts codified in machine-readable format
 - If... then...
- Machine Learning:
 - Statistical analysis of existing data to make predictions about new data
 - Supervised or unsupervised (more about this later)


RULE-BASED AI VS MACHINE LEARNING

- Tic-Tac-Toe:
 - Perfect information
 - A player with perfect strategy can never lose
 - 5.478 possible situations (765 without rotation and mirroring)
- Rule-based:
 - Implementation of rules (**not** equal to the rules of the game), e.g.:
 - Start with a cross in the upper right corner
 - If two circles are in a row, put a cross behind them
 - ...
 - Needs expert knowledge (how to win?)
 - What should we do in cases for which we have no rule?

O	O	O
	O	X
	X	X

MACHINE LEARNING

Supervised

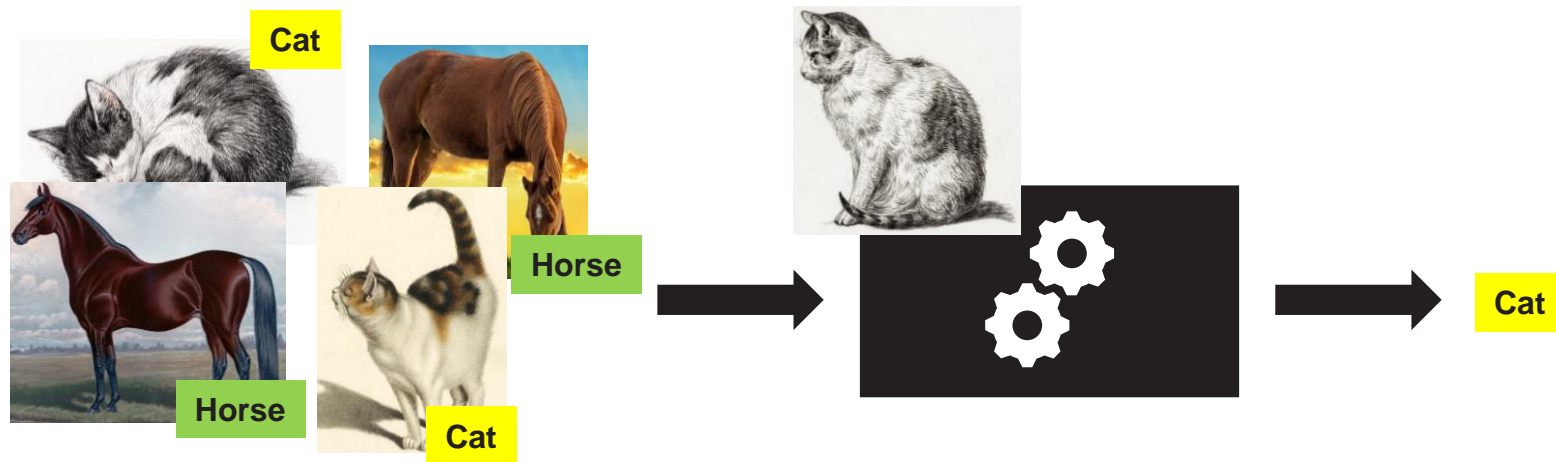


Unsupervised




MACHINE LEARNING

Supervised



Unsupervised



GENERATIVE AI

- „Normal“ AI provides an output based on an existing set of options:
 - Amazon suggests a product out of its catalog
 - Speaking to a Tesla activates a certain command
 - Face recognition returns the name of a person
 - ...
- Generative AI creates something new that did not (necessarily) exist before
 - Images: DALL-E, Stable Diffusion, Midjourney, ...
 - Texts: ChatGPT, Bard, Llama, ...

WHY IS LANGUAGE HARD FOR MACHINES?



bread

car



dread

automobile

Similar „look“



Similar „meaning“



HOW DOES IT WORK?

DISTRIBUTIONAL HYPOTHESIS


- **Words that appear in the same context tend to have similar meaning** (Z. S. Harris, 1954)
- „A word is characterized by the company it keeps“ (J. R. Firth, 1957)

I like to read _____.

I take the _____ to get to work.

TRAINING LANGUAGE MODELS

Her parents called her Alice

P(no)? P(her)? P(car)?
Her parents  her Alice

  called  

WHAT IF WE DON'T WANT TO PREDICT THE NEXT WORD?

- We can use *fine-tuning* to adapt a model that was trained on one task to do another task with relatively little data
- Large Language Models like ChatGPT are first trained to predict the next word (sequence of words) and then fine-tuned to answer questions

FINE-TUNED VS. NOT FINE-TUNED

Prompt: “What do you think about Artificial Intelligence?”

GPT-2 (text completion):

“Please answer in the comments below.”

ChatGPT (question answering):

“As an AI language model myself, I do not have the ability to have personal opinions or feelings. However, I can provide you with some information about artificial intelligence and its potential benefits and challenges. [...]”

=> We need fine-tuning in order to enable LLMs to solve specific tasks... but do we?

ZERO-SHOT LEARNING

LLMs can perform tasks on which they were not specifically trained. We can for example use prompt engineering:

Prompt: Extract the name from the sentence "My name is Jack."

GPT-2: If the sentence does not end in "I'm," then it will break the sentence.

Prompt:

Her name is Alice. => Alice

His name is Bob. => Bob

My name is Jack. =>

GPT-2:


Jack

My name is Thomas. => Thomas

ZERO-SHOT LEARNING?

ChatGPT Passes US Medical Licensing Exam Without Clinician Input

ChatGPT achieved 60 percent accuracy on the US Medical Licensing Exam, indicating its potential in advancing artificial intelligence-assisted medical education.



Source: Getty Images


 By Shania Kennedy



TRICK-SHOT LEARNING!

ChatGPT ~~Passes~~ US Medical Licensing Exam ~~Without Clinician Input~~

ChatGPT achieved 60 percent accuracy on the US Medical Licensing Exam, indicating its potential in advancing artificial intelligence-assisted medical education.



Source: Getty Images

By Shania Kennedy



Simone Balloccu
@simoneballoccu



Proposal: use "unknown-shot" or (funnier) "trick-shot" for model with undisclosed training data. Anything but let's stop calling GPT3, ChatGPT etc. "zero-shot" please.

<https://twitter.com/simoneballoccu/status/1631701409308475393>



Have you ever used AI at work?



1



49% 26

2



51% 27





What have you used AI for at work?



code Prioritize a to-do list
Interview me
Make Ppt English Brainstorming Art Please Sql sentences
Preparation workshop) Brainstorm Rubrics Invitations Give use course
html texts translation Structure nvt Documents review improve
Canvas formal Programming Information Excel formula Translations
Revise a document Supergoogle examples language Design
Emails Structure in forum

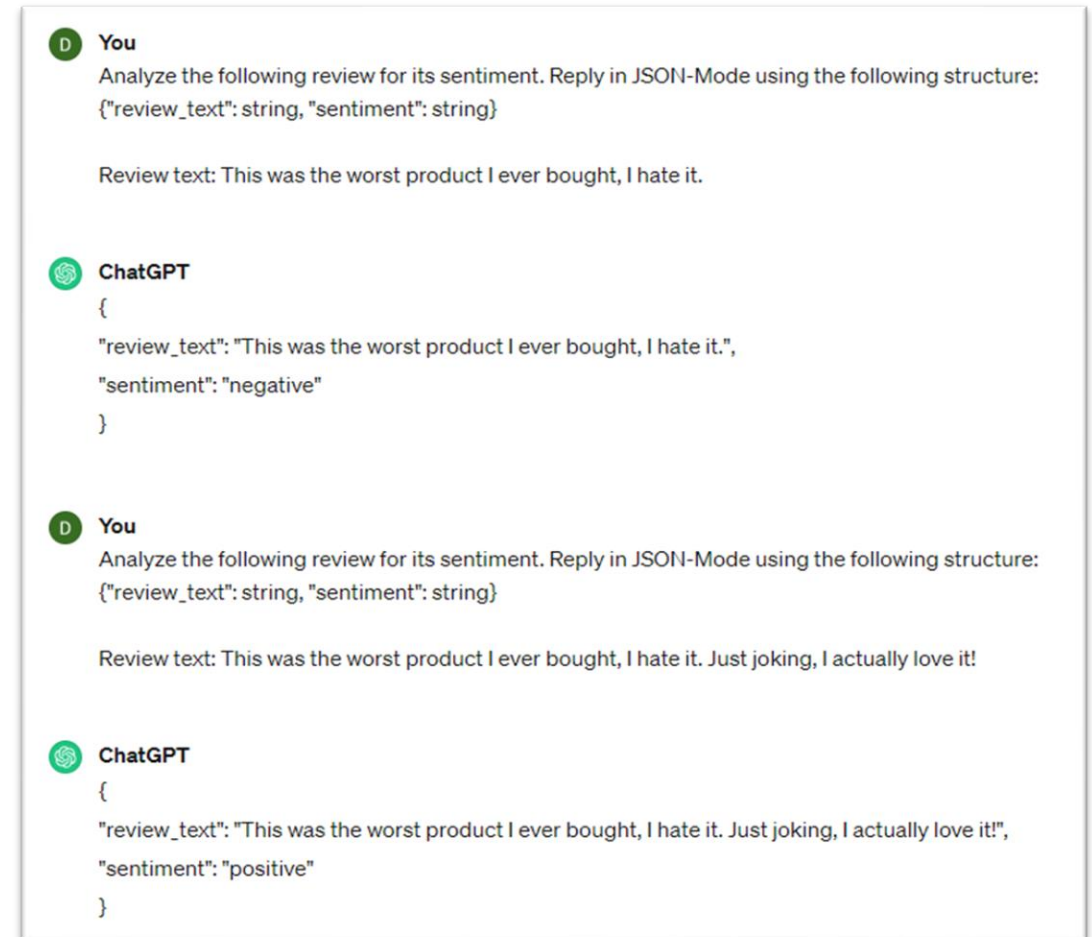


TASKS AI CAN HELP YOU WITH

- Writing
- Evaluate what others wrote
- Discovering documents
- Programming
- Data analysis
- ?

THE OBVIOUS ONE: CHATGPT (OR LLMS IN GENERAL)

- You all know the advantages and disadvantages...
- Prompt Engineering is the new SEO
- It can help you with all of the previously mentioned tasks
- It is also quite good at producing structured output:



The screenshot displays a chat interface with three messages. Each message from the user (labeled 'You') asks for sentiment analysis of a review in JSON format. The first review is negative, and the second is positive. ChatGPT's responses are JSON objects with 'review_text' and 'sentiment' fields.

You
Analyze the following review for its sentiment. Reply in JSON-Mode using the following structure:
{
 "review_text": string, "sentiment": string
}

Review text: This was the worst product I ever bought, I hate it.

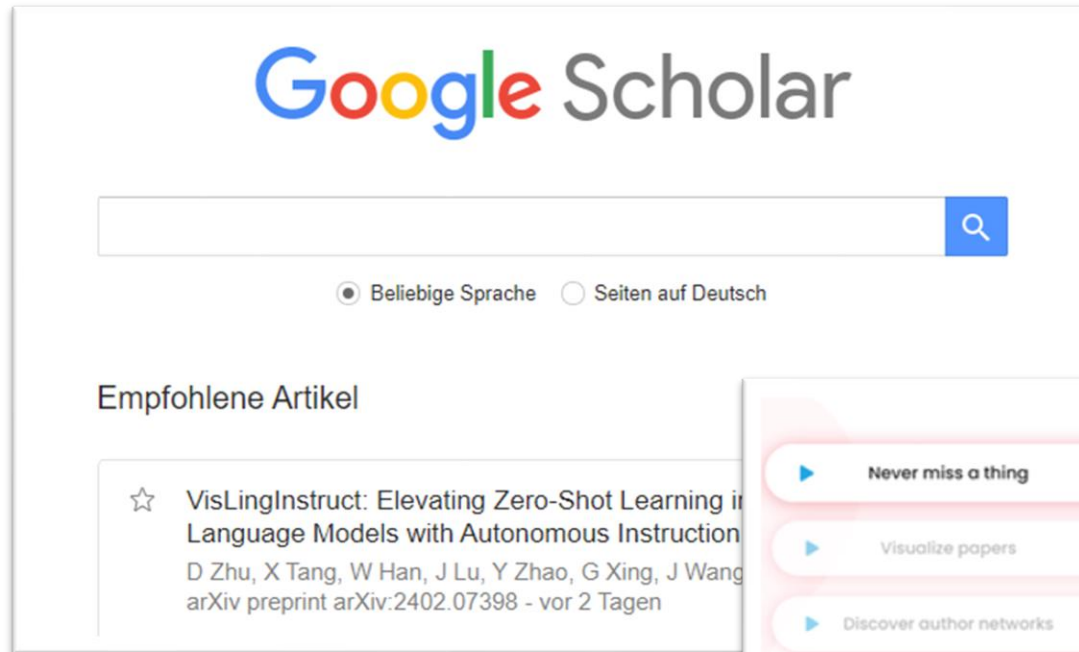
ChatGPT
{
 "review_text": "This was the worst product I ever bought, I hate it.",
 "sentiment": "negative"
}

You
Analyze the following review for its sentiment. Reply in JSON-Mode using the following structure:
{
 "review_text": string, "sentiment": string
}

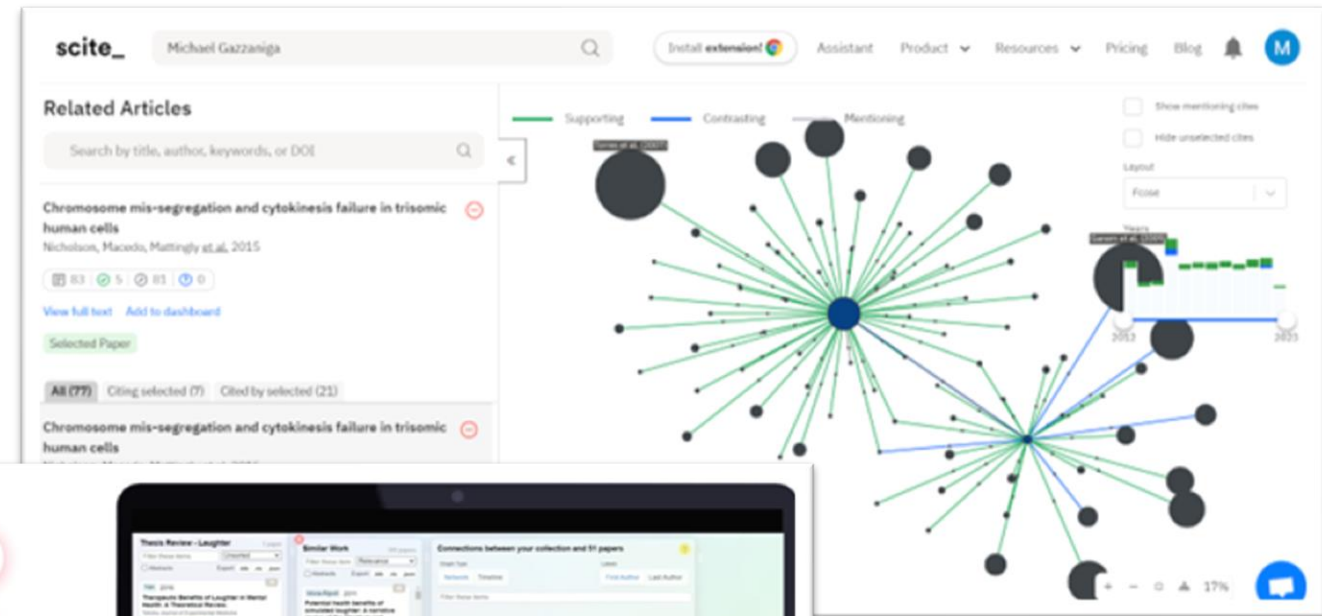
Review text: This was the worst product I ever bought, I hate it. Just joking, I actually love it!

ChatGPT
{
 "review_text": "This was the worst product I ever bought, I hate it. Just joking, I actually love it!",
 "sentiment": "positive"
}

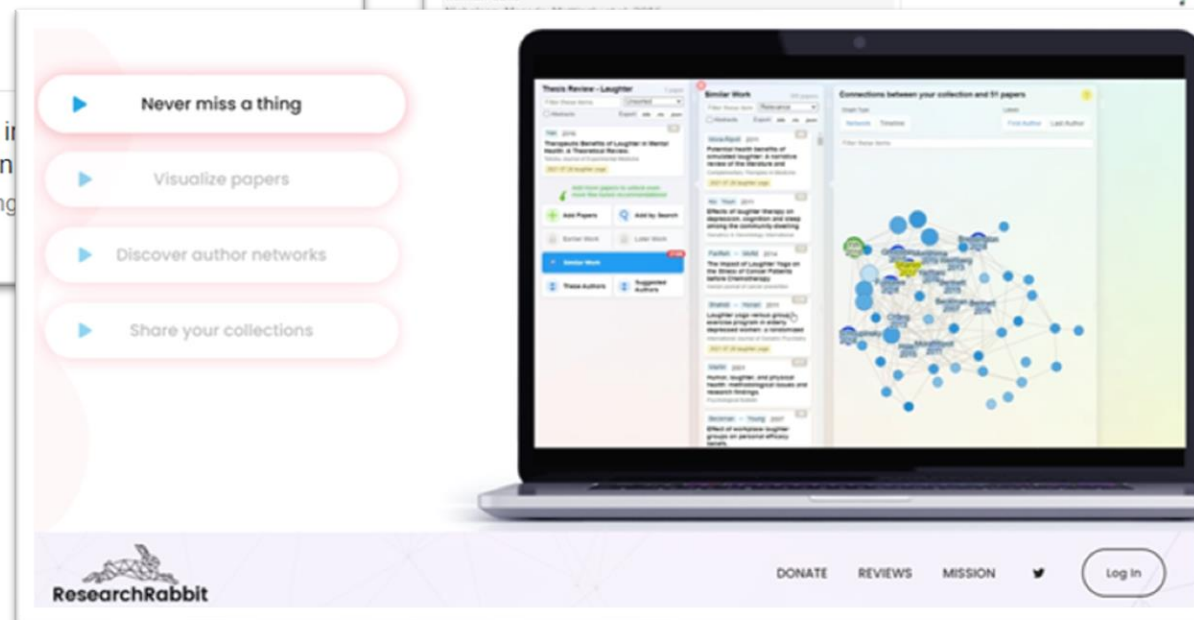
DISCOVERING DOCUMENTS



The image shows the Google Scholar homepage. At the top is the Google Scholar logo. Below it is a search bar with a magnifying glass icon. Underneath the search bar are two radio buttons: "Beliebige Sprache" (selected) and "Seiten auf Deutsch". Below this is a section titled "Empfohlene Artikel" (Recommended articles). The first article listed is "VisLingInstruct: Elevating Zero-Shot Learning in Language Models with Autonomous Instruction" by D Zhu, X Tang, W Han, J Lu, Y Zhao, G Xing, J Wang, with an arXiv preprint link and a "vor 2 Tagen" (2 days ago) timestamp.



The image shows the Scite_ interface. At the top, there's a search bar with "Michael Gazzaniga" and a search icon. To the right are navigation links: "Install external", "Assistant", "Product", "Resources", "Pricing", "Blog", and a user profile icon "M". Below the search bar is a section titled "Related Articles" with a search input "Search by title, author, keywords, or DOI". A featured article is "Chromosome mis-segregation and cytokinesis failure in trisomic human cells" by Nicholson, Macoda, Mattingly et al., 2015. To the right is a large network graph with nodes of varying sizes and colors (green, blue, black) connected by lines. A legend above the graph indicates "Supporting" (green), "Contrasting" (blue), and "Monitoring" (black). On the right side of the graph, there are checkboxes for "Show monitoring cites" and "Hide unselected cites", and a "Layout" dropdown menu set to "Force".



The image shows the ResearchRabbit interface on a laptop. On the left side, there are four red buttons with white text: "Never miss a thing", "Visualize papers", "Discover author networks", and "Share your collections". The laptop screen displays a dashboard with several panels. The top left panel is titled "Thesis Review - Laughlin" and shows a list of papers. The top right panel is titled "Similar Work" and shows a list of similar papers. The bottom right panel is titled "Connections between your collection and 51 papers" and shows a network graph with blue nodes and edges. At the bottom of the laptop screen, there are navigation links: "DONATE", "REVIEWS", "MISSION", a Twitter icon, and a "Log In" button. The ResearchRabbit logo is visible in the bottom left corner of the laptop screen.

PROGRAMMING

```
sentiments.ts write_sql.go parse_expenses.py addresses.rb

1 import datetime
2
3 def parse_expenses(expenses_string):
4     """Parse the list of expenses and return the list of triples (date, value, currency).
5     Ignore lines starting with #.
6     Parse the date using datetime.
7     Example expenses_string:
8         2016-01-02 -34.01 USD
9         2016-01-03 2.59 DKK
10        2016-01-03 -2.72 EUR
11    """
12    expenses = []
13    for line in expenses_string.splitlines():
14        if line.startswith("#"):
15            continue
16        date, value, currency = line.split(" ")
17        expenses.append((datetime.datetime.strptime(date, "%Y-%m-%d"),
18                        float(value),
19                        currency))
20    return expenses

Copilot
```


DATA ANALYSIS

Orange Data Mining Workflow:

- Data Source:** File (Class-labeled data set)
- Data Table:** Data
- Learners:** Naive Bayes, Random Forest, k Nearest Neighbours
- Evaluation Results:** Confusion Matrix
- Selected Data:** Data

Annotations:

- It's always a good idea to check out the data first.
- Select a cell in confusion matrix to obtain instances. Here they are in the spreadsheet.
- Cross-validation takes place here. Double click to see the performance scores.
- Several learners are scored in cross-validation at the same time.

Power BI Dashboard: Customer Segmentation

- Columns:** Longitude (generated), Latitude (generated)
- Rows:** Latitude (generated)
- Filter:** Product Category D, Purchase Date
- Map:** Customer location (United States)
- Customer Count:** 19 (10,000, 20,000, 30,000, 40,000, 49,194)

AI Assistant (Emstein):

- Just ask a question or click a suggestion to create a viz.
- What are you interested in?
- Are there patterns over time for sales across product categories?
- What pet types drive the most purchases?
- Does loyalty tier impact purchase size?
- Here's a viz based on your selection.
- Show me the location of customers who bought Sporting Goods in the last 3 months by zip code.
- OK. This viz answers your question: Show me the location of customers who bought Sporting Goods in the last 3 months by zip code.
- How far are these customers from store locations?

Orange Data Mining Workflow (Detailed):

- Load data
- Preprocessing: All general preprocessing when begins. Double click to see the details.
- Replace Missing Values: Replace missing values.
- Reorder Attributes: Order columns automatically.
- Filter Examples: Model on cases with total values. Apply the model on cases with a missing value for the target column.
- Cross Validation: Split the data into a training set and a test set.
- Model Evaluation: Evaluate the model on the test set.
- Model Selection: Choose the best model.
- Apply Model: Apply the selected model to the training data.
- Explain Predictions: Create predictions for cases with total values and with explanations for predictions.
- Create Lift Chart: Create lift chart.

Results:

- Model evaluation
- Performance from 10-fold cross-validation
- Model
- Predefined data with explanation for only if the data had missing values
- Predefined data with explanation for only if the data had missing values
- Lift chart

TOOL OVERVIEW

- Writing:
 - ChatGPT, Gemini, Llama, BLOOM, ...
 - Grammarly, DeepL, ...
 - Copyleak, ZeroGPT, ...
- Discovering new Research:
 - Scholar, Scopus, Schemantic Scholar, ...
 - Scite, scholarcy, ...
 - Research Rabbit, VOSviewer, ...
- Programming: Copilot
- Data analysis:
 - Orange, RapidMiner, ...
 - Tableau, MonkeyLearn, Sisense

NON-AI TOOLS

There is huuuge amount of digital tools out there that do not use AI an can still help you to

- Be more organized
- Work more efficiently
- Make your life easier
- ...

GENERAL USAGE RULES

- (If you officially share work:) Clearly indicate parts of the work that are not your own.
- You are responsible for mistakes and violations of rights.
- Use it in contexts where you can assess the correctness of the output.
- If your text can be written by an AI, does it need to be written?

GENERAL USAGE RULES

University of Twente:

If Artificial Intelligence (AI) tools are allowed for completing an assignment several different levels are possible, i.e. only for certain purposes, only specific tools or completely free. It is required that students mention the use of AI. Therefore, when any form of AI is used, the use of these tools should be included in the appendix (list all tools that were used during the work).

The following sentence should be used at the start of the appendix:

“During the preparation of this work the author(s) used [NAME TOOL / SERVICE] in order to [REASON]. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the work.”

DANGERS OF AI

- Like most disruptive technologies, AI has the potential to do great good and great harm at the same time
- We can limit the dangers by using AI responsibly

We should always remember:

- Correlation does not imply causation
- Garbage in, Garbage out: The quality of your AI depends on the quality of your data
- Machine Learning projects the past into the future
- Bias: Algorithmic decision have the „aura“ of objectivity, although subjectivity is included in the data

CORRELATION DOES NOT IMPLY CAUSATION

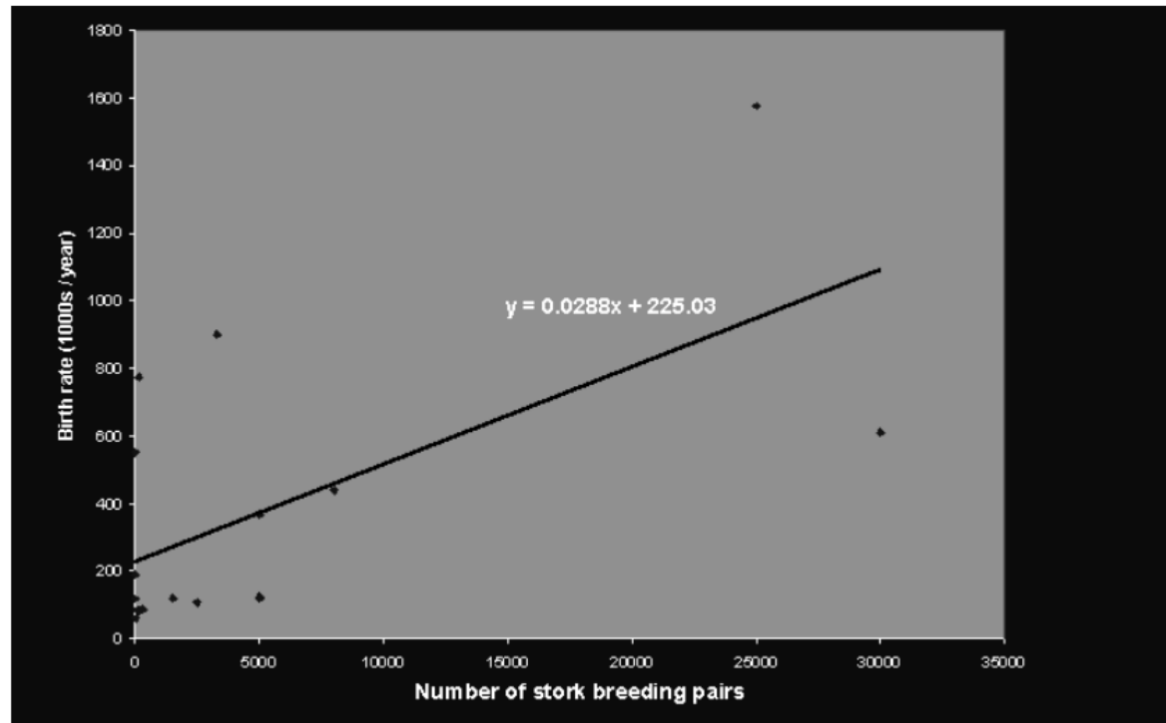


Fig 1. How the number of human births varies with stork populations in 17 European countries.

WHAT DO YOU THINK?

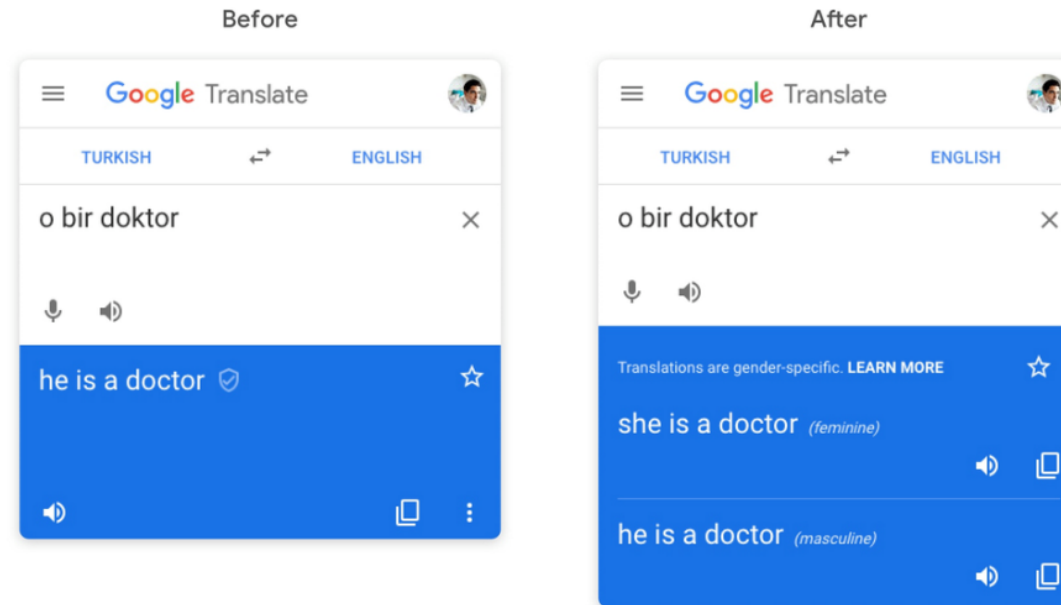


GARBAGE IN – GARBAGE OUT



1. A blond girl and a bald man with his arms crossed are standing inside looking at each other.
2. A worker is being scolded by her boss in a stern lecture.
3. A manager talks to an employee about job performance.
4. A hot, blond girl getting criticized by her boss.
5. Sonic employees talking about work.

BIAS





Do you see any new potential or dangers for AI in your work after the talk?

