



ChatGPT: Optimizing Language Models for Dialogue

We've trained a model called ChatGPT which interacts in a conversational way. The dialogue format makes it possible for ChatGPT to answer followup questions, admit its mistakes, challenge incorrect premises, and reject inappropriate requests. ChatGPT is a sibling model to [InstructGPT](#), which is trained to follow an instruction in a prompt and provide a detailed response.

[TRY CHATGPT ↗](#)

Colloque d'ABB, 16.11.2023

Marcel Steiner-Curtis

FHNW – Haute école spécialisée du nord-ouest de la Suisse

ChatGPT

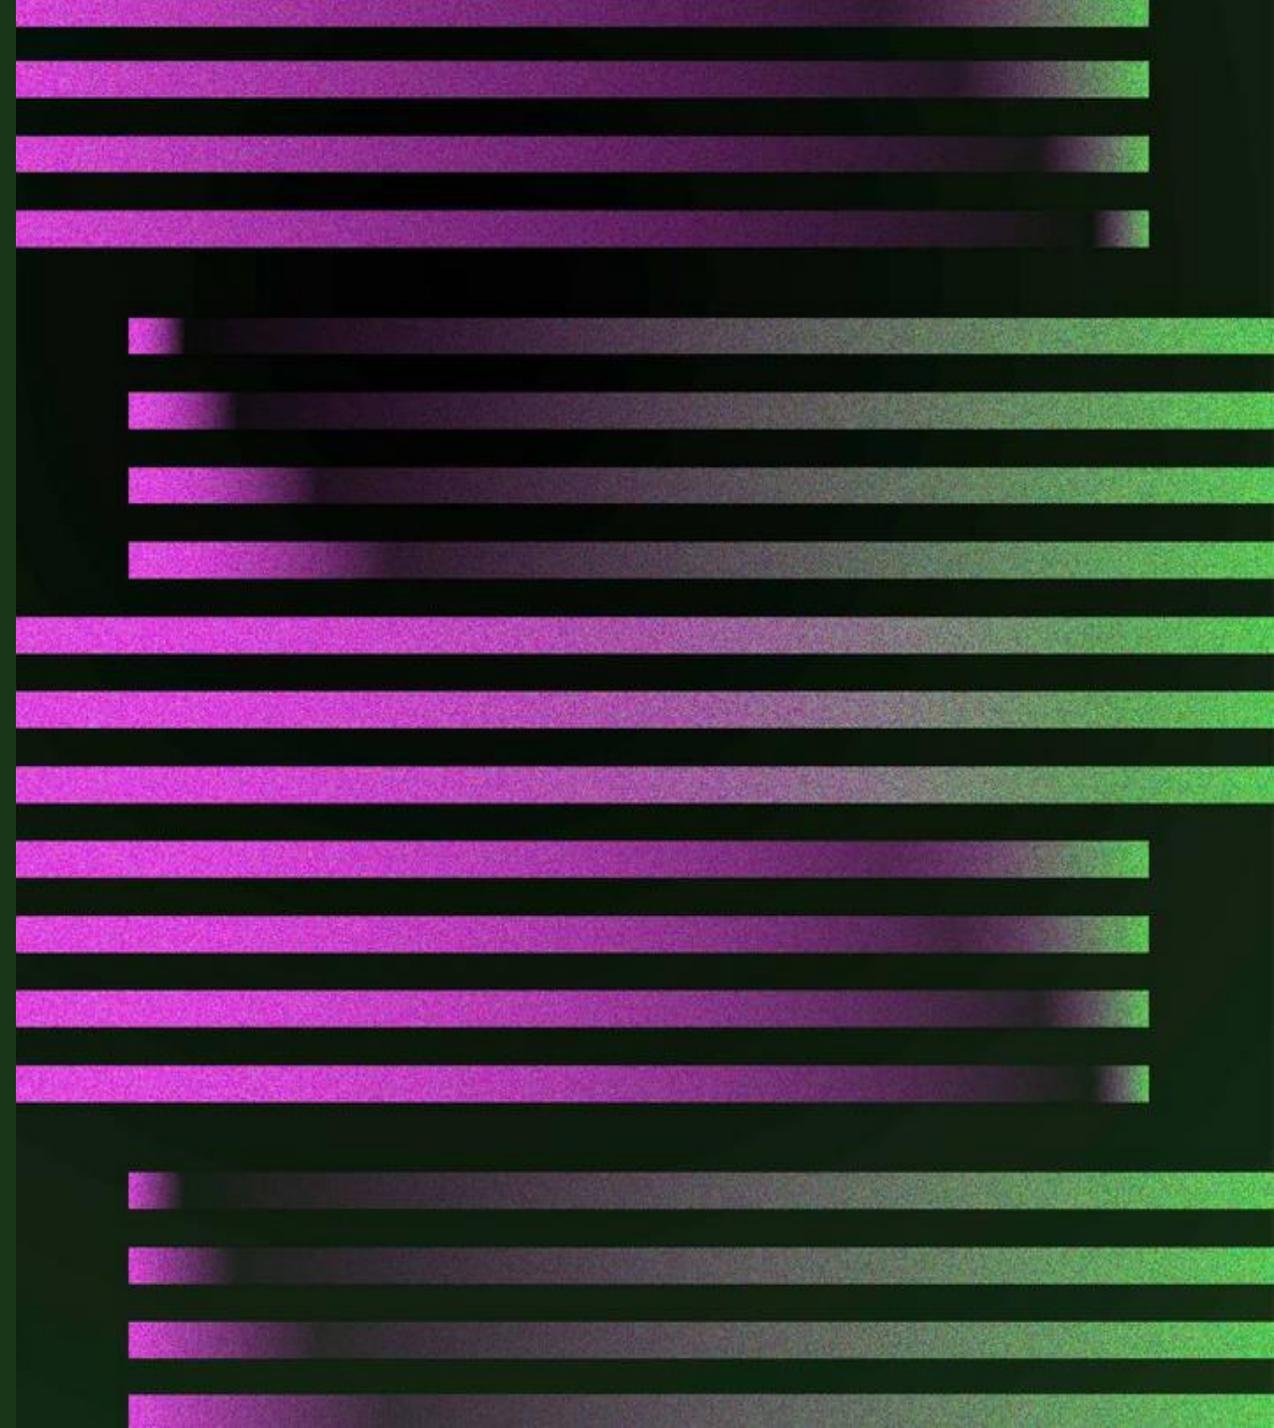
Mon meilleur étudiant?



ChatGPT: Optimizing Language Models for Dialogue

We've trained a model called ChatGPT which interacts in a conversational way. The dialogue format makes it possible for ChatGPT to answer followup questions, admit its mistakes, challenge incorrect premises, and reject inappropriate requests. ChatGPT is a sibling model to [InstructGPT](#), which is trained to follow an instruction in a prompt and provide a detailed response.

[TRY CHATGPT ↗](#)





Solve the following with R: The annual maximum sea levels [in cm] in Venice, 1931-1981 have been recorded by P. A. Pirazzoli. The data set

Year	SeaLevel
1931	103
1932	78
1933	121
1934	116
1935	115
⋮	⋮
1978	132
1979	166
1980	134
1981	138

contains the annual maximum tides at Venice for the 51 years.

- Represent the data in a scatter diagram sea level versus year and describe the functional context in words.
- Fit a straight line to the data points. Give the estimated parameter values.
- Add the model in the scatter diagram. Comment on the solution.
- Does the data support the hypothesis that Venice sinks? Make a statistical test.

MSE Master Statistiques appliquées et analyses de données



I'll provide you with the R code to perform these tasks and interpret the results. However, I cannot execute the code as I'm an AI text model. You can execute the code in your local R environment or use an online R compiler like RStudio or Jupyter Notebook.

First, let's input the data into a data frame:

R

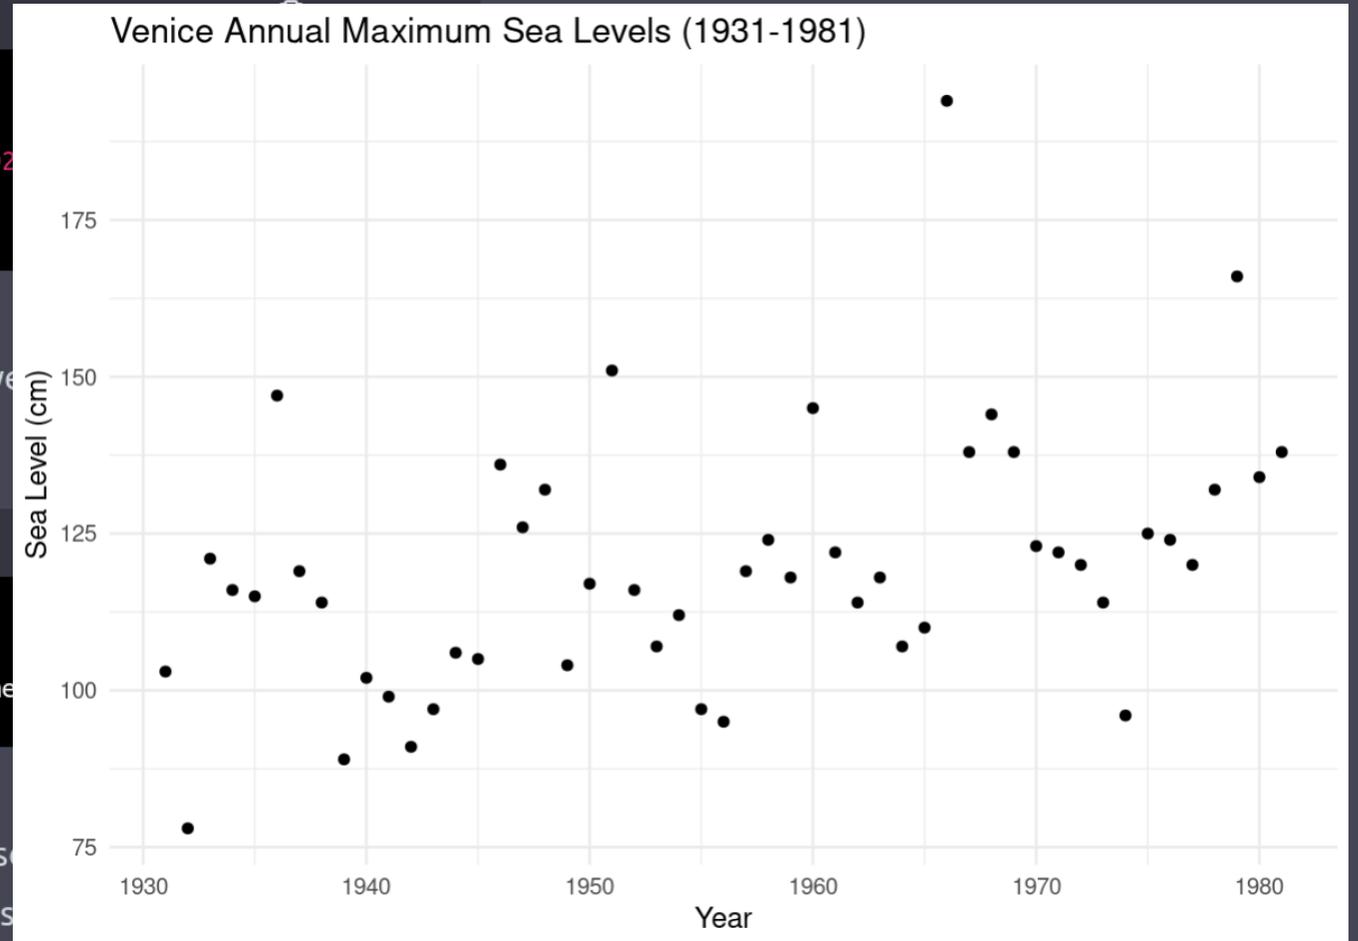
```
Year <- c(1931:1981)
SeaLevel <- c(103, 78, 121, 116, 115, 147, 119, 114, 89, 102)
data <- data.frame(Year, SeaLevel)
```

a. Represent the data in a scatter diagram sea level versus year. Provide a functional context in words.

R

```
library(ggplot2)
ggplot(data, aes(x=Year, y=SeaLevel)) + geom_point() + theme_minimal()
```

This code will create a scatter plot of the annual maximum sea level from 1931 to 1981. You should see a general upward trend in sea levels, which indicates a sinking phenomenon.



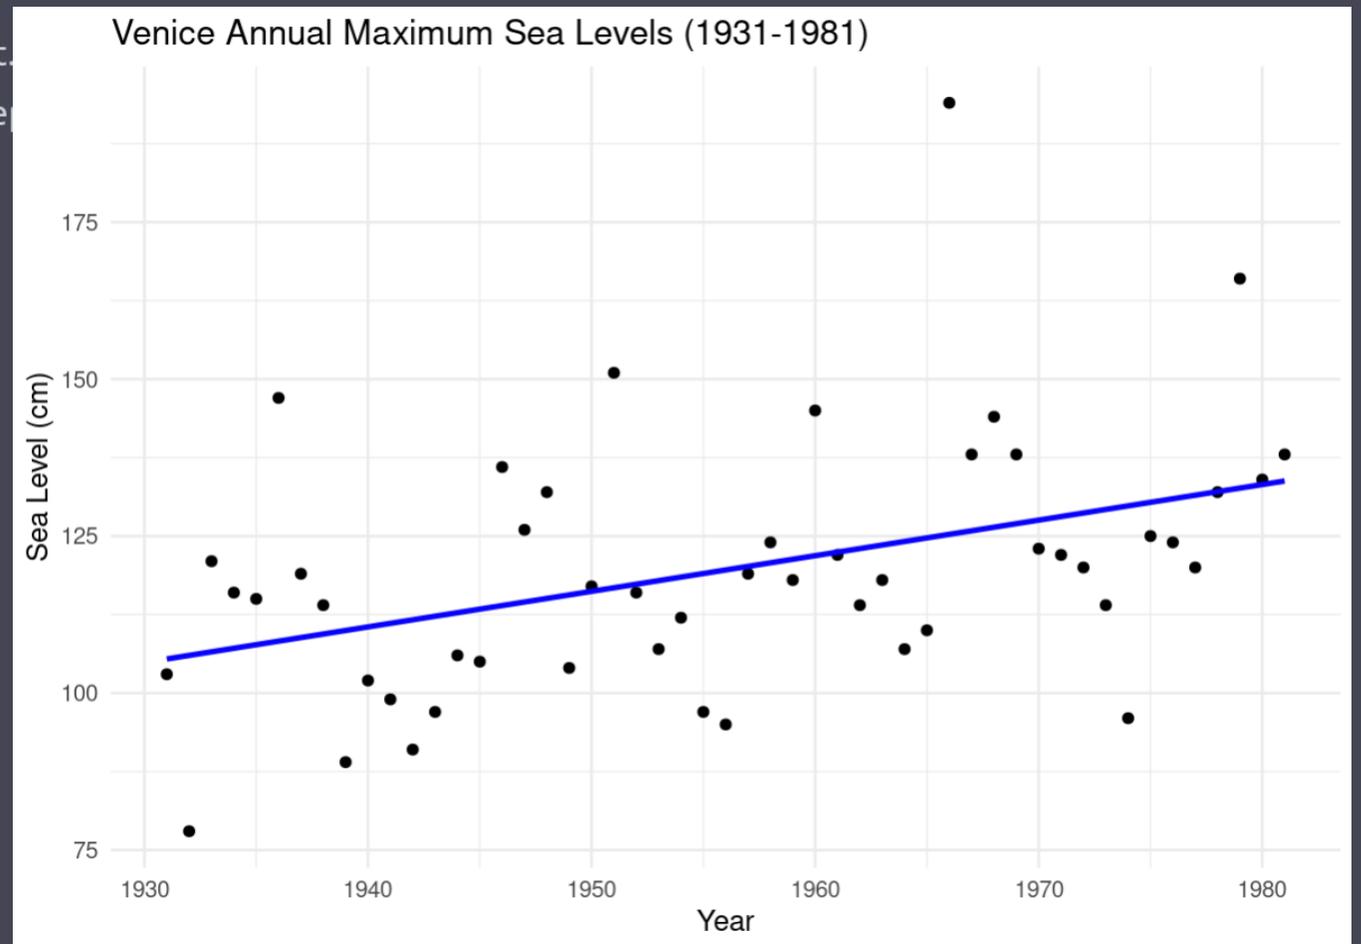
c. Add the model in the scatter diagram. Comment on the solution.

R

Copy code

```
ggplot(data, aes(x=Year, y=SeaLevel)) + geom_point() + geom_smooth(method = "lm", se = FALSE)
```

This code will add the fitted linear model to the scatter plot. Evaluate how well the line represents the trend if the trend is captured accurately.



d. Does the data support the hypothesis that Venice sinks? Make a statistical test.

R

 Copy code

```
# Perform a t-test on the slope of the linear regression model
slope_test <- summary(linear_model)$coefficients["Year", "t value"]
slope_pvalue <- summary(linear_model)$coefficients["Year", "Pr(>|t|)"]
```

```
slope_test
slope_pvalue
```

The t-test will give you the t-value and p-value for the slope of the fitted line. If the p-value is less than a significance level (e.g., 0.05), you can reject the null hypothesis that the slope is equal to zero, suggesting a significant relationship between the year and sea level, which might support the hypothesis that Venice is sinking.

```
> # Perform a t-test on the slope of the linear regression model
> slope_test <- summary(linear_model)$coefficients["Year", "t value"]
> slope_pvalue <- summary(linear_model)$coefficients["Year", "Pr(>|t|)"]
>
> slope_test
[1] 3.200852
> slope_pvalue
[1] 0.002406147
```



0-677



REVOX
high fidelity

<< >> PLAY STOP REC

MONITOR

STEREO CH1 CH2 MONO

VOLUME

BALANCE

PHONES MIC CH1 MIC CH2

CHANNEL I CHANNEL II

RECORD

MIC HI RADIO MIC HI RADIO

MIC LO AUX MIC LO AUX

LEVEL

LEVEL

POWER

3 1/2 OFF 3 1/2



ReVox A77 rechter Wickelmotor wird feurig heiss. Was ist defekt?



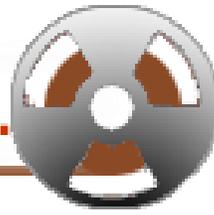
Wenn der rechte Wickelmotor Ihres ReVox A77 Tonbandgeräts übermäßig heiß wird, kann dies auf ein mechanisches oder elektrisches Problem hinweisen. Es ist wichtig zu beachten, dass ich keine Echtzeitinformationen habe und nicht auf dem neuesten Stand der Technik bin, da meine Daten nur bis September 2021 reichen. Dennoch kann ich Ihnen einige allgemeine Hinweise geben:



1. **Mechanische Probleme:** Ein heiß laufender Motor könnte auf übermäßige Reibung oder Blockaden im Mechanismus hinweisen. Überprüfen Sie, ob der Motor leichtgängig läuft, oder ob es Anzeichen für mechanische Probleme wie abgenutzte Lager, Riemen oder Zahnräder gibt.
2. **Elektrische Probleme:** Überhitzung kann auch auf elektrische Probleme hinweisen. Möglicherweise fließt zu viel Strom durch den Motor, was zu einer Überlastung führen kann. Dies könnte auf defekte elektrische Komponenten wie Transistoren, Kondensatoren oder Widerstände zurückzuführen sein.
3. **Lüftung:** Stellen Sie sicher, dass die Lüftungsöffnungen des Geräts nicht blockiert sind.



...:BANDMASCHINENFORUM:...:

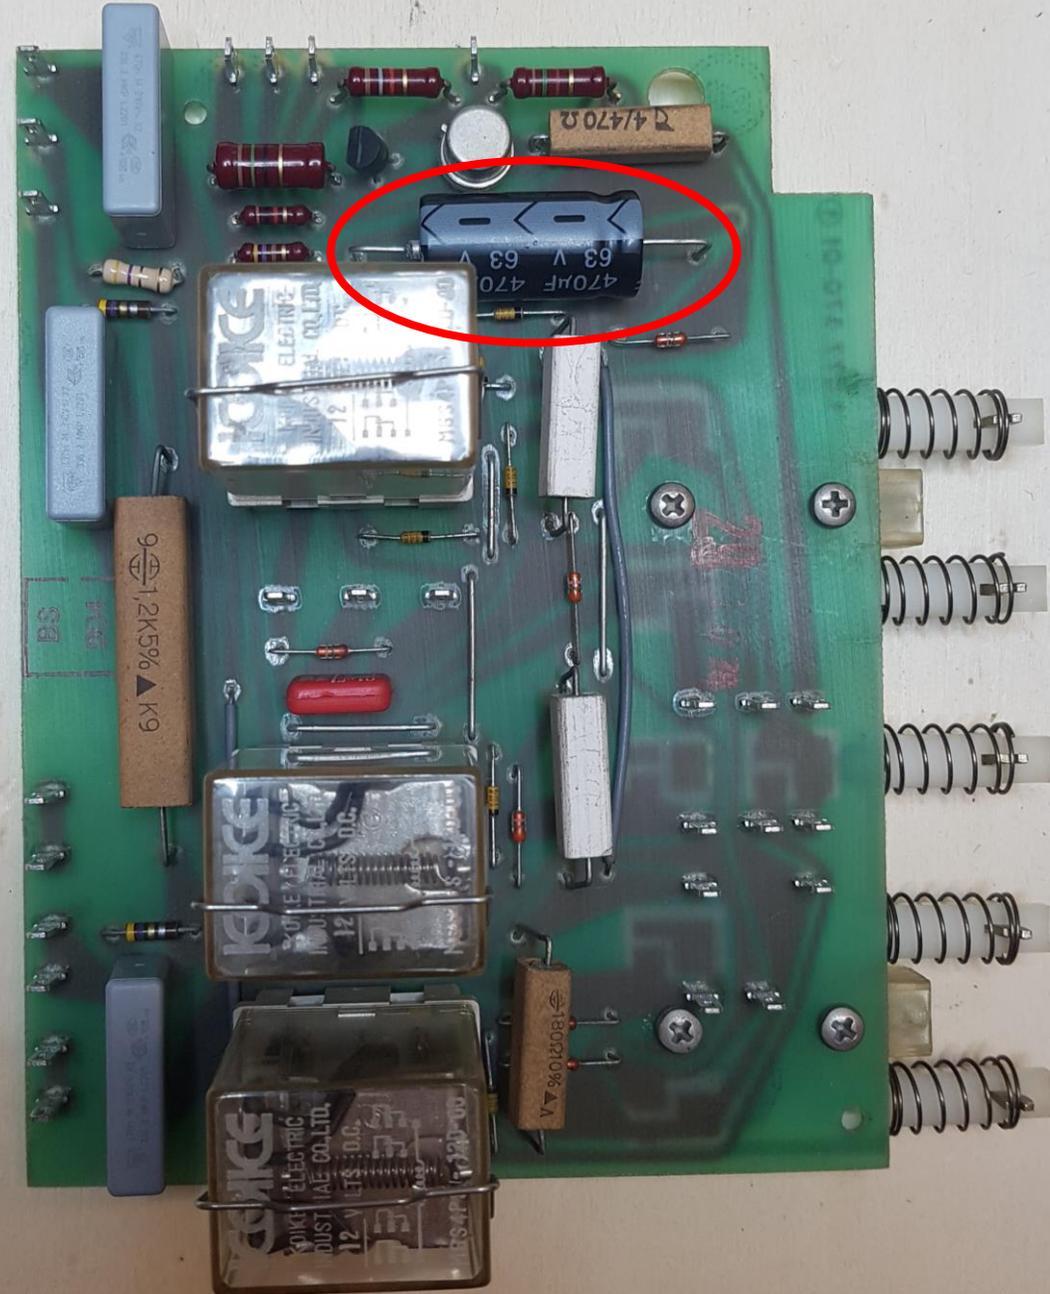


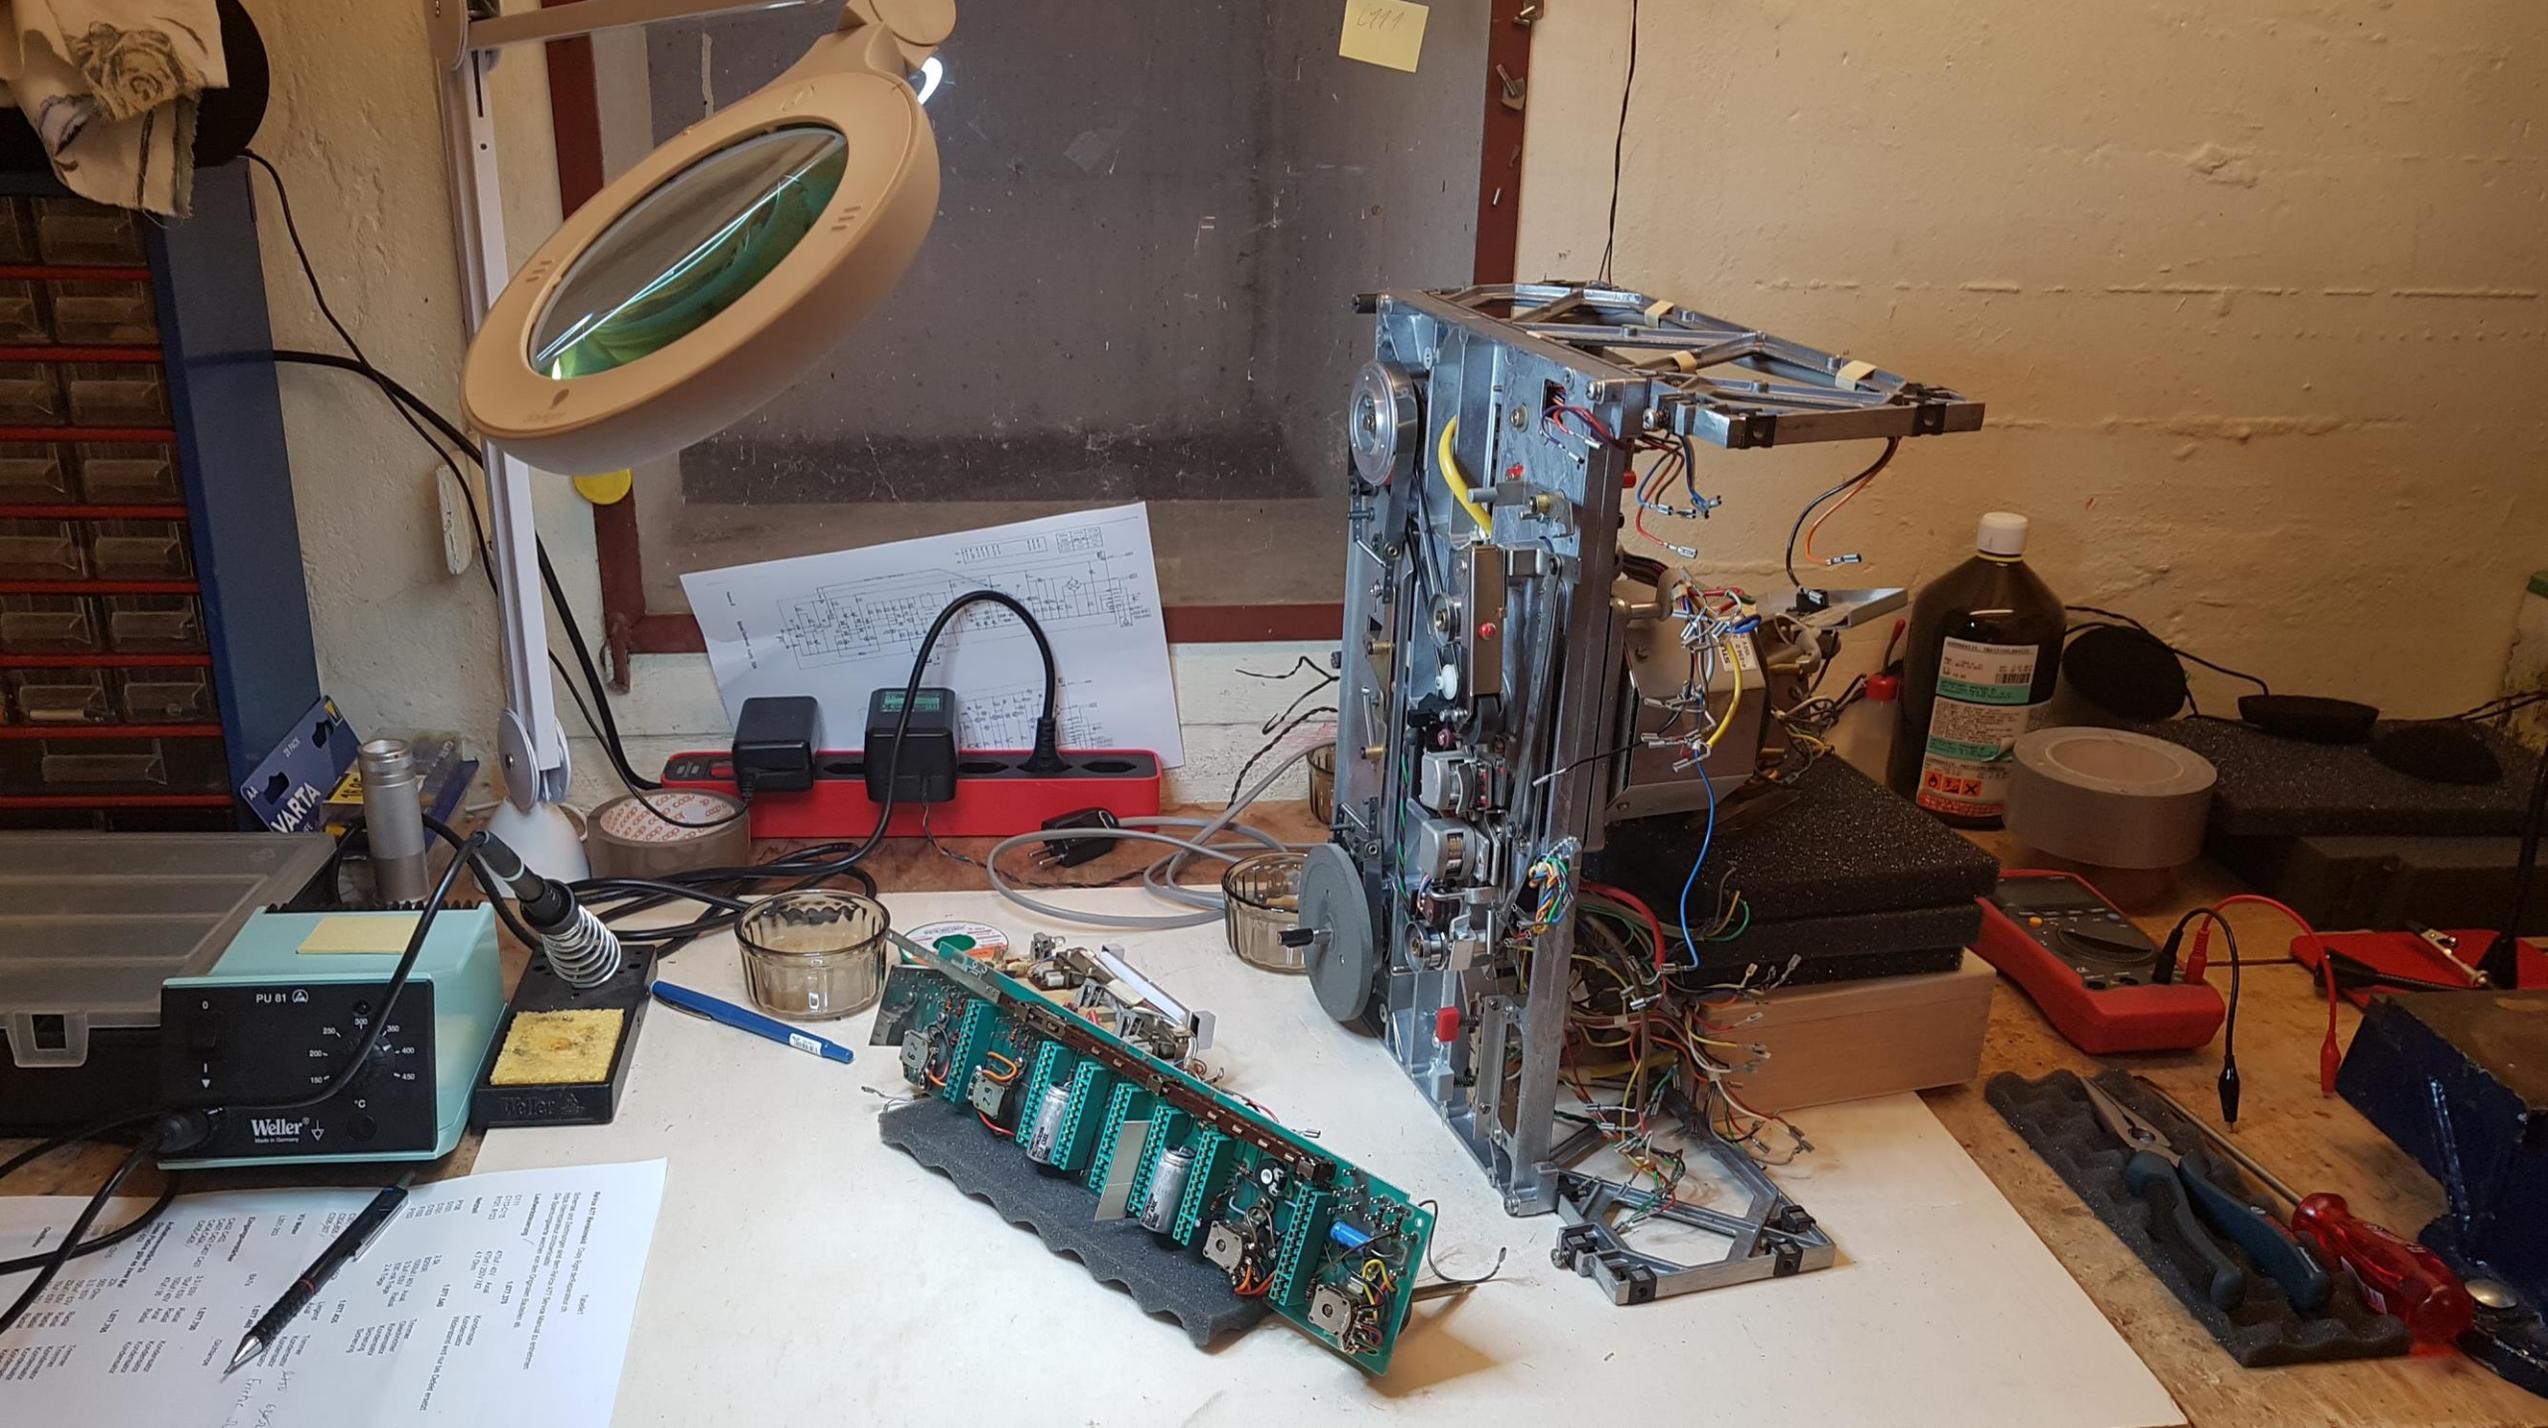
kaimex ●

09.04.2019, 21:10

Hast du schon mal den C111 500 µF überprüft ?

MfG Kai





0111



Weller
PU 81
230V
300W
150V
450W

Handwritten notes and a list of components on a piece of paper. The list includes:

Ref	Designation	Value	Notes
1	Resistor	10k	
2	Resistor	100k	
3	Resistor	1k	
4	Resistor	100Ω	
5	Resistor	10k	
6	Resistor	100k	
7	Resistor	1k	
8	Resistor	100Ω	
9	Resistor	10k	
10	Resistor	100k	
11	Resistor	1k	
12	Resistor	100Ω	
13	Resistor	10k	
14	Resistor	100k	
15	Resistor	1k	
16	Resistor	100Ω	
17	Resistor	10k	
18	Resistor	100k	
19	Resistor	1k	
20	Resistor	100Ω	
21	Resistor	10k	
22	Resistor	100k	
23	Resistor	1k	
24	Resistor	100Ω	
25	Resistor	10k	
26	Resistor	100k	
27	Resistor	1k	
28	Resistor	100Ω	
29	Resistor	10k	
30	Resistor	100k	
31	Resistor	1k	
32	Resistor	100Ω	
33	Resistor	10k	
34	Resistor	100k	
35	Resistor	1k	
36	Resistor	100Ω	
37	Resistor	10k	
38	Resistor	100k	
39	Resistor	1k	
40	Resistor	100Ω	
41	Resistor	10k	
42	Resistor	100k	
43	Resistor	1k	
44	Resistor	100Ω	
45	Resistor	10k	
46	Resistor	100k	
47	Resistor	1k	
48	Resistor	100Ω	
49	Resistor	10k	
50	Resistor	100k	
51	Resistor	1k	
52	Resistor	100Ω	
53	Resistor	10k	
54	Resistor	100k	
55	Resistor	1k	
56	Resistor	100Ω	
57	Resistor	10k	
58	Resistor	100k	
59	Resistor	1k	
60	Resistor	100Ω	
61	Resistor	10k	
62	Resistor	100k	
63	Resistor	1k	
64	Resistor	100Ω	
65	Resistor	10k	
66	Resistor	100k	
67	Resistor	1k	
68	Resistor	100Ω	
69	Resistor	10k	
70	Resistor	100k	
71	Resistor	1k	
72	Resistor	100Ω	
73	Resistor	10k	
74	Resistor	100k	
75	Resistor	1k	
76	Resistor	100Ω	
77	Resistor	10k	
78	Resistor	100k	
79	Resistor	1k	
80	Resistor	100Ω	
81	Resistor	10k	
82	Resistor	100k	
83	Resistor	1k	
84	Resistor	100Ω	
85	Resistor	10k	
86	Resistor	100k	
87	Resistor	1k	
88	Resistor	100Ω	
89	Resistor	10k	
90	Resistor	100k	
91	Resistor	1k	
92	Resistor	100Ω	
93	Resistor	10k	
94	Resistor	100k	
95	Resistor	1k	
96	Resistor	100Ω	
97	Resistor	10k	
98	Resistor	100k	
99	Resistor	1k	
100	Resistor	100Ω	

Handwritten text on a label or piece of paper, possibly a component list or identification tag.

Handwritten notes and a list of components on a piece of paper, similar to the one on the left.

Le verre est-il
à moitié
plein...



... ou
à moitié
vide?



Hey ChatGPT! Tell me how you work! Some mathematical details would be welcome.



Of course! The foundation behind ChatGPT and other models like it from OpenAI is called the Transformer architecture, which was introduced in the 2017 paper "Attention is All You Need" by Vaswani et al.



At a high level, Transformer models process input data (like text) in chunks and use a mechanism called "attention" to weigh the importance of different parts of the input when producing an output. This allows the model to consider both local and global context in the data.

Here are some key mathematical and conceptual details:

1. **Embeddings:** Words (or more specifically, tokens) are initially transformed into vectors using embeddings. These vectors represent the semantic meaning of words in a high-dimensional space.

Research Scientist

San Francisco, California, United States

Applying to this role, you will be considered for Research Scientist roles across OpenAI.

Join an innovative machine learning team you work on, while also working for people who work at a large scale, and

Compensation, Benefits and Perks

The annual salary range for this role is **\$200,000 – \$370,000**. Total compensation includes generous equity and benefits.

We expect you

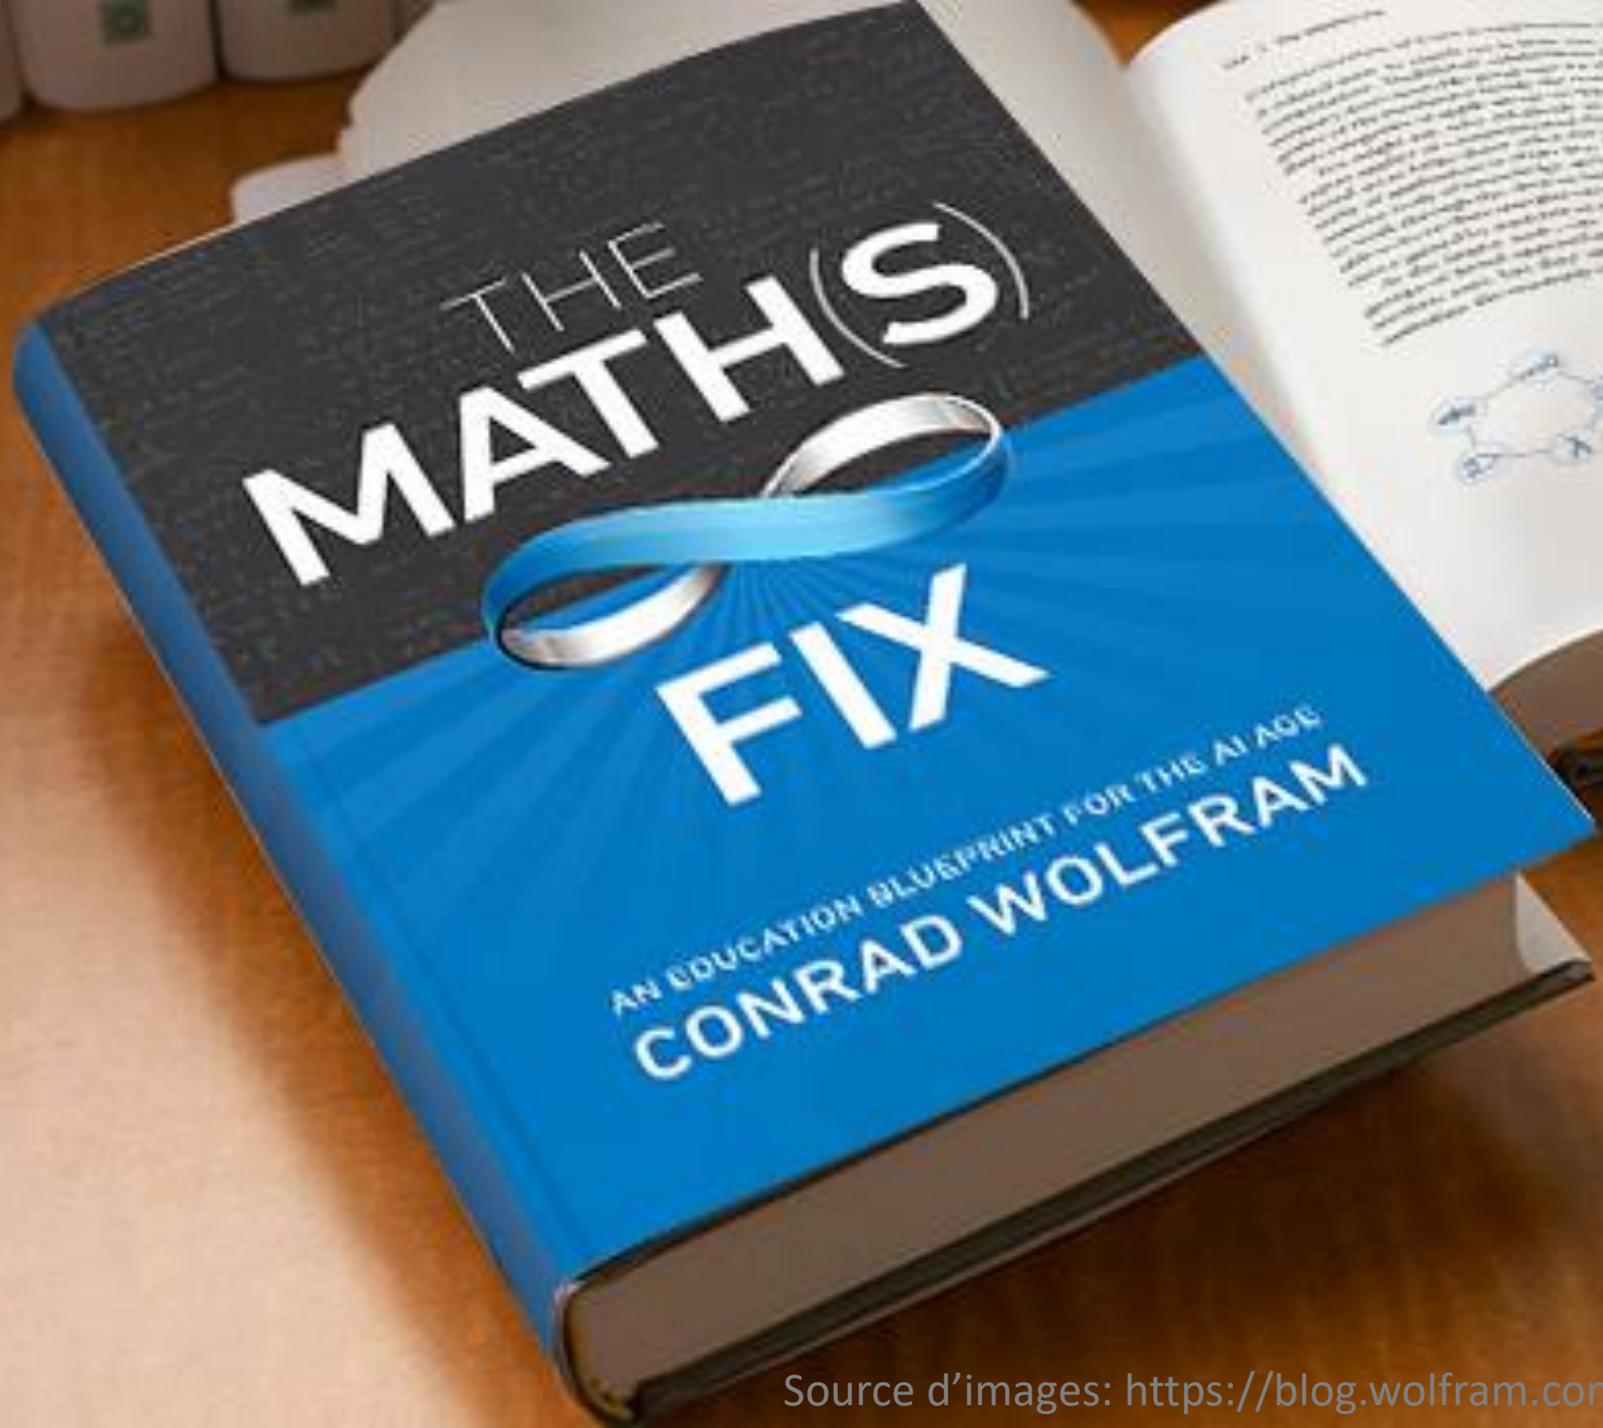
5. Feed

- Have a track record of developing ideas in machine learning, demonstrating publications or projects
- Have the ability to own and pursue a research problem and autonomously carrying out research

Repenser l'enseignement

FHNW
Bourses
«Enseignement supérieur
2025»





Source d'images: <https://blog.wolfram.com/2020/06/10/released-today-the-maths-fix/>

Conrad Wolfram: résolution de problème en quatre étapes

«Avec toute la mystique, la complexité et ses étranges gribouillis de symboles qui les caractérisent, les mathématiques traditionnelles ou la pensée computationnelle pourraient se décrire comme suit»



Morale de l'histoire (aujourd'hui)

Vraiment impressionnant!

Ce truc se montre beaucoup trop sûr de lui!

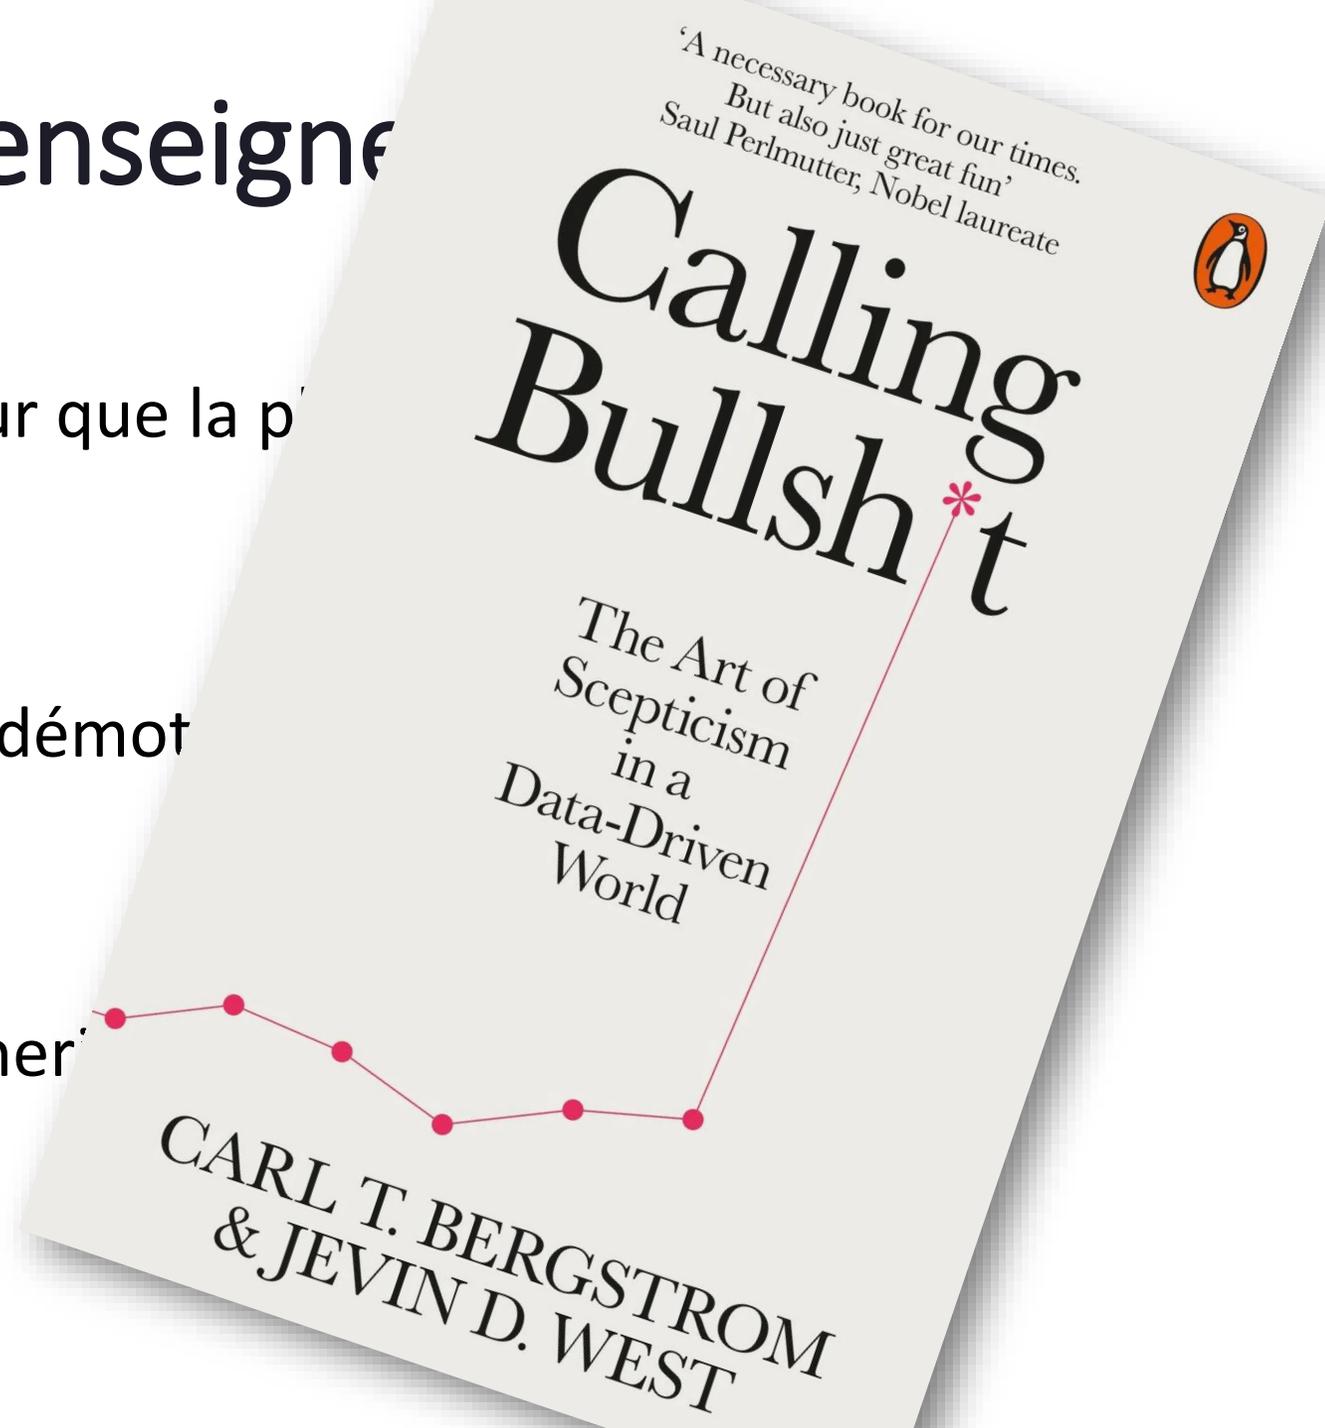
Il est uniquement possible de séparer le bon grain de l'ivraie lorsqu'on a compris le sujet.

Repenser l'enseignement

Ce *truc* est meilleur que la p

Gros potentiel de démot

Détecteur de conneri





*C'est à
l'usage que
l'on peut
juger
de la
qualité
d'une
chose.*

*«La prédiction
est un
exercice
difficile,
surtout
lorsqu'elle
concerne
l'avenir.»*

Niels Bohr,
1885-1962

