

Welcome to

BS6207

2021

Lee Hwee Kuan

Evaluation quiz to test yourself

AI and deep learning

Current state of affairs

How much effort is needed to
publish a deep learning paper?

Is publishing a deep learning
paper an accomplishment?

Is this the end goal?

Artificial Intelligence Index Report 2021

Stanford University, Human-Centered AI

https://aiindex.stanford.edu/wp-content/uploads/2021/11/2021-AI-Index-Report_Master.pdf

Source: Microsoft Academic Graph, 2020 | Chart: 2021 AI Index Report

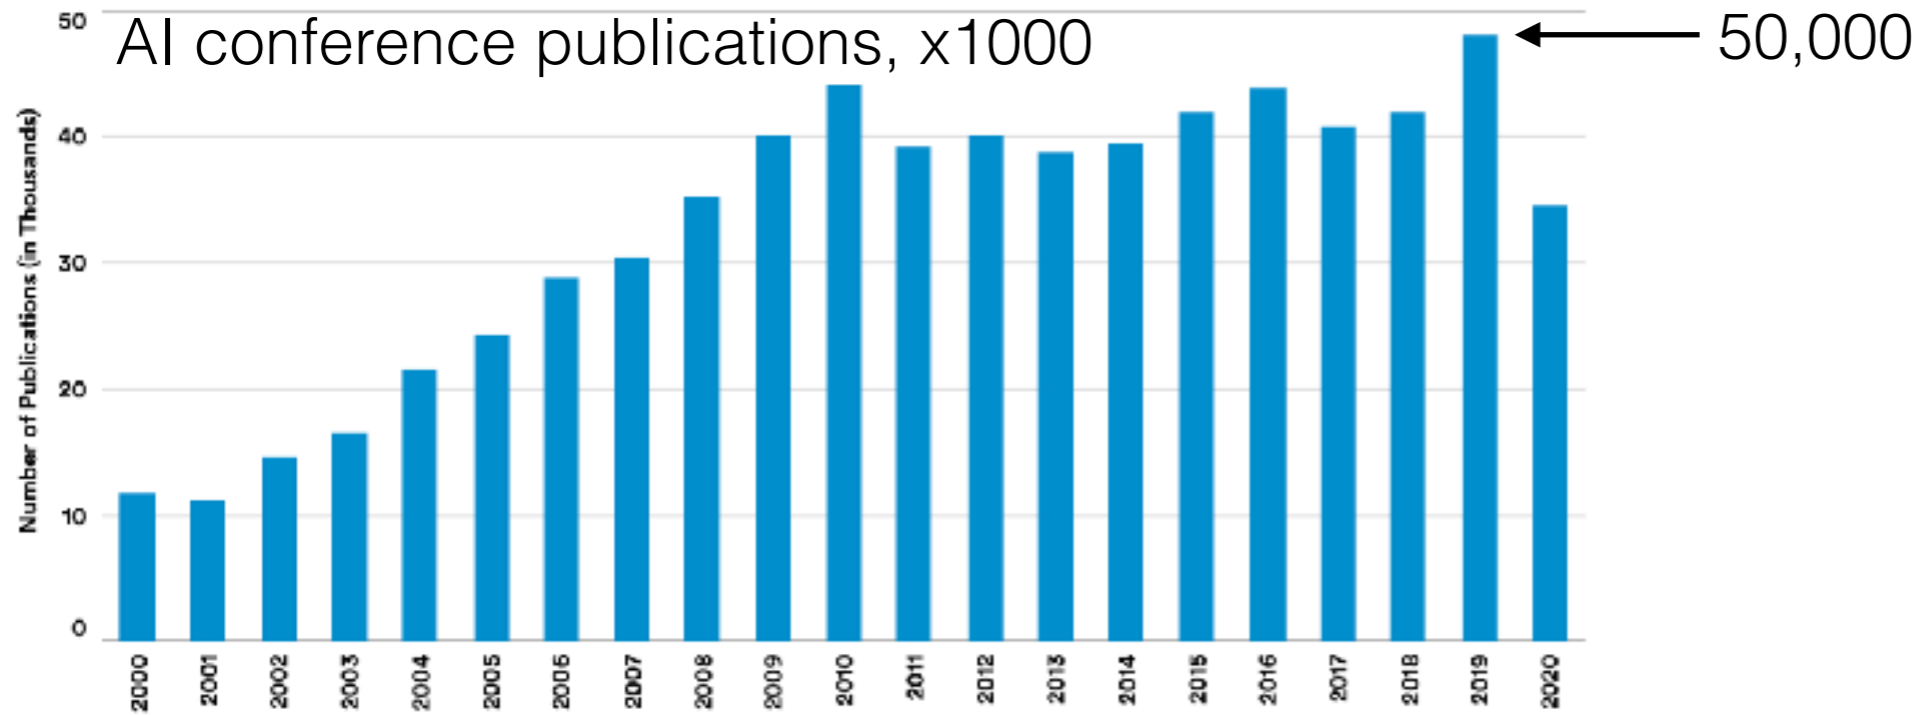


Figure 1.1.11a

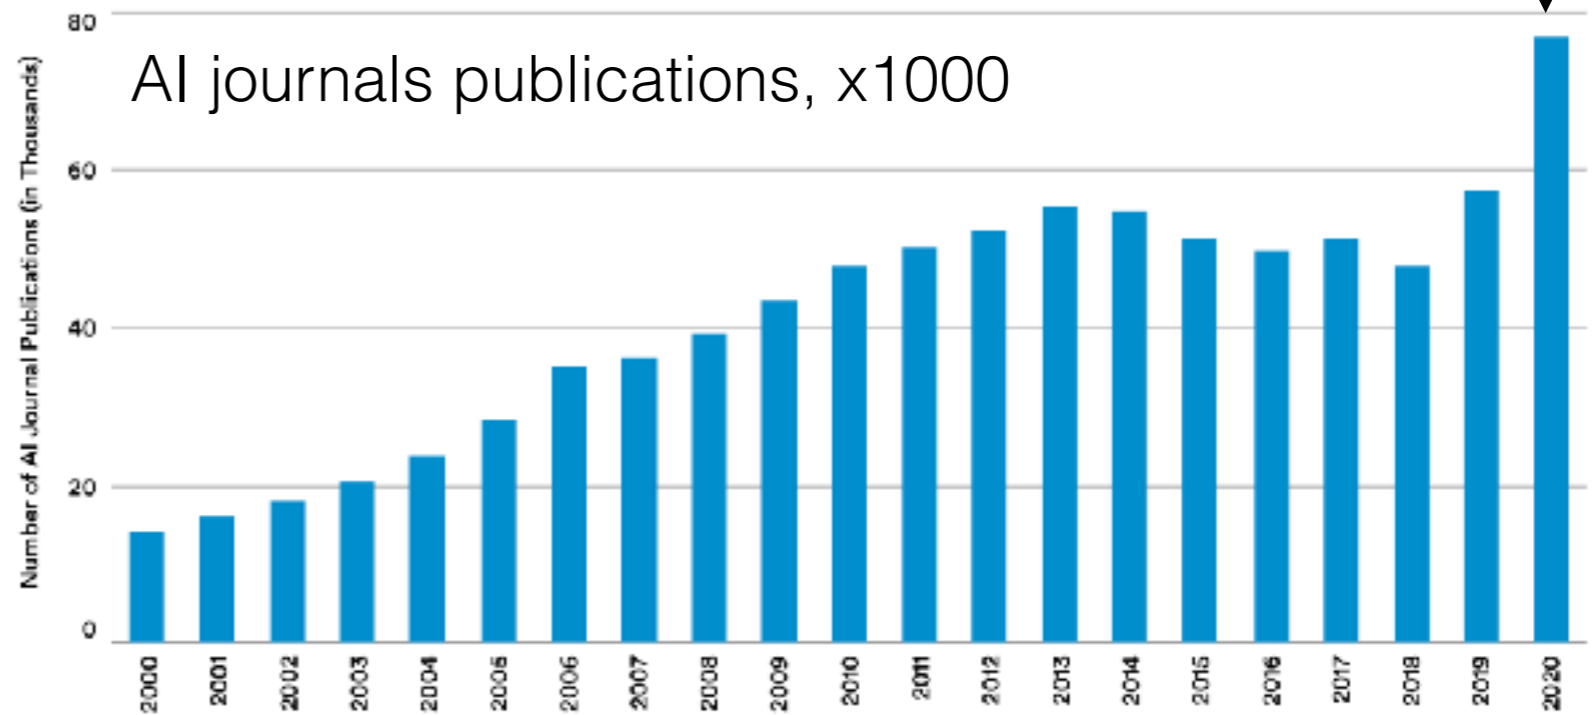


Figure 1.1.7a

AI JOURNAL PUBLICATIONS (% of ALL JOURNAL PUBLICATIONS), 2000-20


Source: Microsoft Academic Graph, 2020 | Chart: 2021 AI Index Report

Artificial Intelligence Index Report 2021

Stanford University, Human-Centered AI

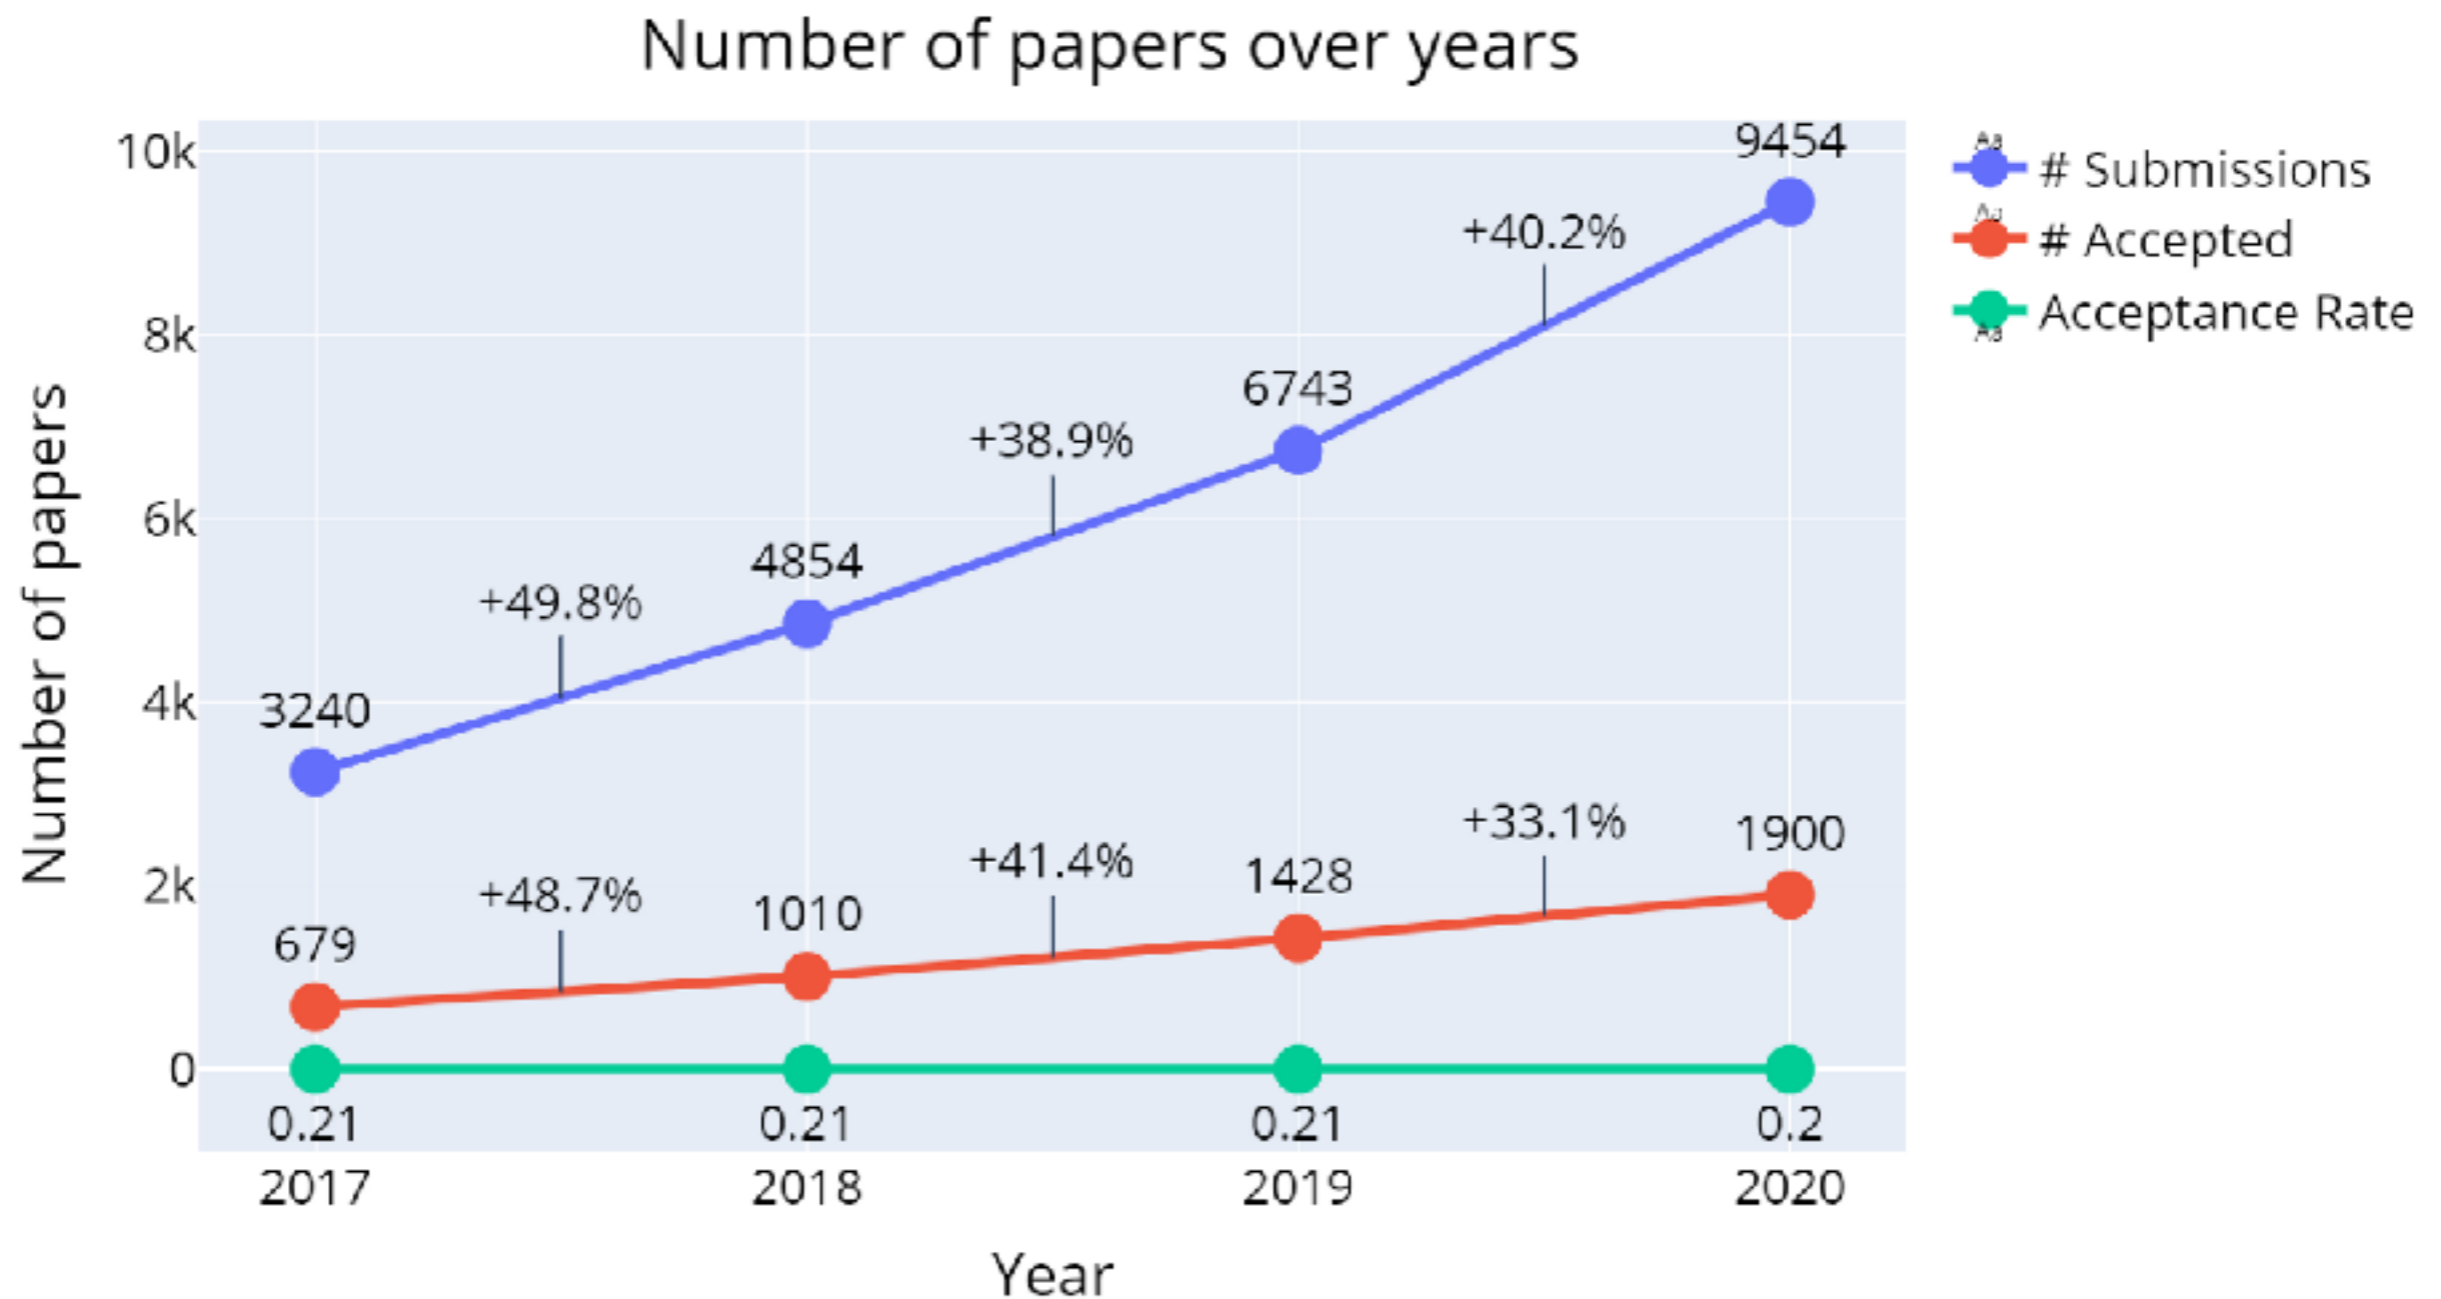
https://aiindex.stanford.edu/wp-content/uploads/2021/11/2021-AI-Index-Report_Master.pdf

Source: Elsevier/Scopus, 2020 | Chart: 2021 AI Index Report



NeurIPS 2020. Comprehensive analysis of authors, organizations, and countries


<https://medium.com/criteo-engineering/neurips-2020-comprehensive-analysis-of-authors-organizations-and-countries-a1b55a08132e>



NeurIPS 2020. Comprehensive analysis of authors, organizations, and countries

<https://medium.com/criteo-engineering/neurips-2020-comprehensive-analysis-of-authors-organizations-and-countries-a1b55a08132e>

Normalized number of papers by organization



What if your NeurIPS 2021 submission is reviewed again?

<https://blog.neurips.cc/2021/12/08/the-neurips-2021-consistency-experiment/>

Original \ Copy	Oral	Spotlight	Poster	Reject	Withdrawn	
Oral	0	0	4	0	0	→ 4
Spotlight	0	3	9	13	0	→ 25
Poster	2	7	74	94	0	→ 177
Reject	0	13	83	462	0	→ 558
Withdrawn	0	0	0	0	118	→ 118

Out of 206 accepted papers, when reviewed again, 107 of them are rejected


Out of 558 rejected papers, when reviewed again, 96 are accepted

An Open Review of OpenReview: A Critical Analysis of the Machine Learning Conference Review Process

David Tran, Alexander V Valtchanov, Keshav R Ganapathy, Raymond Feng, Eric Victor Slud, Micah Goldblum, Tom Goldstein

28 Sep 2020 (modified: 06 Mar 2021) ICLR 2021 Conference Blind Submission Readers:  Everyone [Show Bibtex](#) [Show](#)

<https://bdtechtalks.com/2020/10/21/ai-conferences-review-process/>



“We observe a downward trend in reproducibility, with scores decreasing from 75% in 2017, 70% in 2018 and 2019, 66% in 2020,”

What will happen to AI field if this is true?

“It is thought by some that the ML review process favors non-controversial papers with incremental theoretical results over papers with big new ideas.”

Are we heading for another **stagnation**?

Another concern in the AI field

AI is becoming more of an empirical science

Empirical science is not always bad, there are many things in deep learning that people do but nobody knows why it works.

Sometimes we have some heuristics, other times, “it” is just pure magic!

Why not understanding is bad?

Because
Deploy what we don't understand . . . and . . .
we start **KILLING** people!

When Software Kills

<https://medium.com/swlh/when-software-kills-ab6f48a15825>

...into two fatal plane crashes involving Boeing 737 aircrafts has brought to light that the company's proprietary software may have been at least partially responsible for the crashes, in which **346** people were killed.

1985 and 1987, at least 5 patients were killed (and others critically injured) when a software-controlled radiation therapy system, the Therac-25, inadvertently administered massive overdoses of radiation — over 100 times the prescribed dose.

Read this article for many more examples.

It is a matter of time before 'accident' happens in AI systems and people die. This is inevitable, but we should reduce its effects by really understanding what we are doing

The world's first jetliner, the de Havilland Comet, was crashing,
and no one knew why.

<https://www.planeandpilotmag.com/article/why-did-the-de-havilland-comet-keep-crashing/>

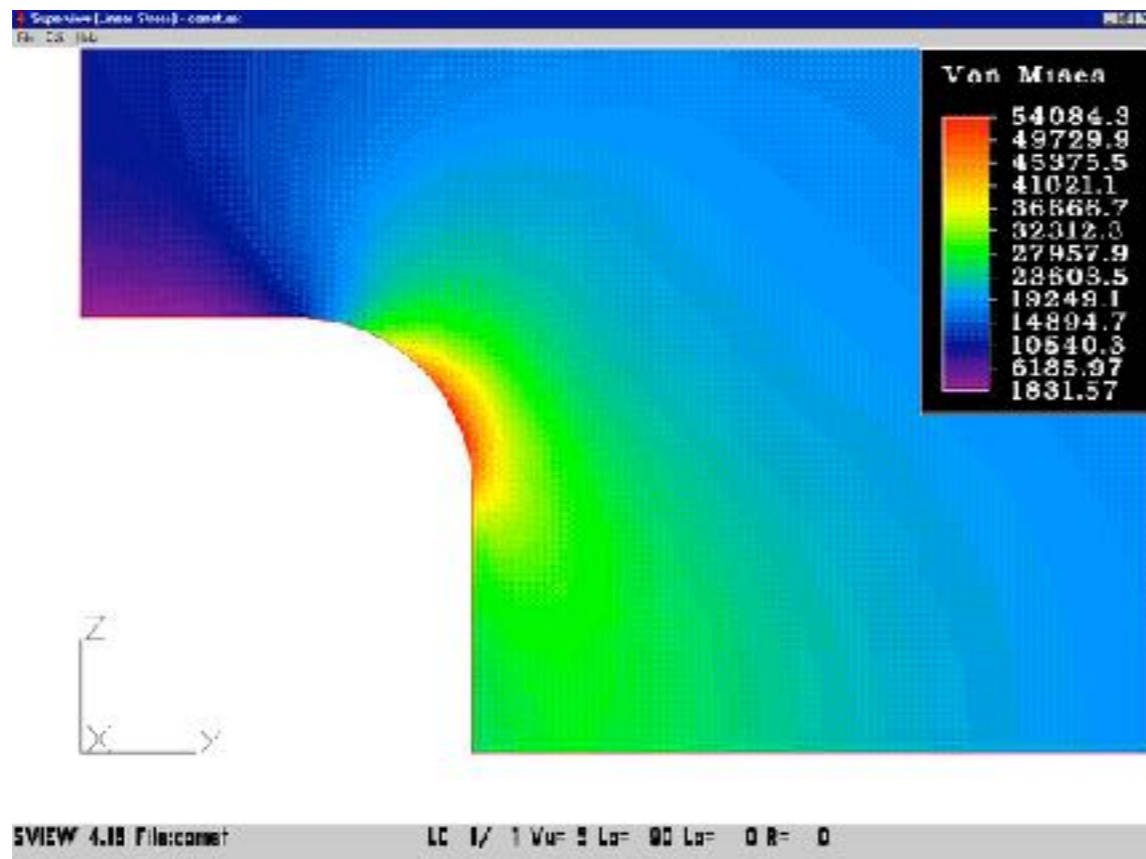
...its short operational life, resulting in more than 400 fatalities.

The world's first jetliner, the de Havilland Comet, was crashing,
and no one knew why.

<https://www.planeandpilotmag.com/article/why-did-the-de-havilland-comet-keep-crashing/>

...its short operational life, resulting in more than 400 fatalities.

Answer: We did not understand the science of structural engineering, the aircraft had square windows.



<https://personal.egr.uri.edu/sadd/compclassroom/fea.html>

What are the next big thing in AI?

- .self driving
- .net for GIS
- .explainable - causal inference Judea pearl
- .skynet
- .RL
- .ML ethics
- .privacy
- .med field - industrial area
- .green storage
- .protein image
- Auto
- Edge
- Finance
- Sports