Machine learning in medical software done right

How to do machine learning and continuous integration of medical software

Peter Søndergaard, 3Shape A/S

3shape

goto; copenhagen

GOTO Copenhagen 2018

Conference Nov. 19 - 21

Click ‘Rate Session’ to rate session and ask questions.

Follow us @gotocph

Agenda

- About me, 3Shape and the data that we are working with
- Motivation, how to enjoy making medical software
- Machine learning and medical software, what's the big deal
- What to do about it
- Medical software and Agile
- Our pipeline and results

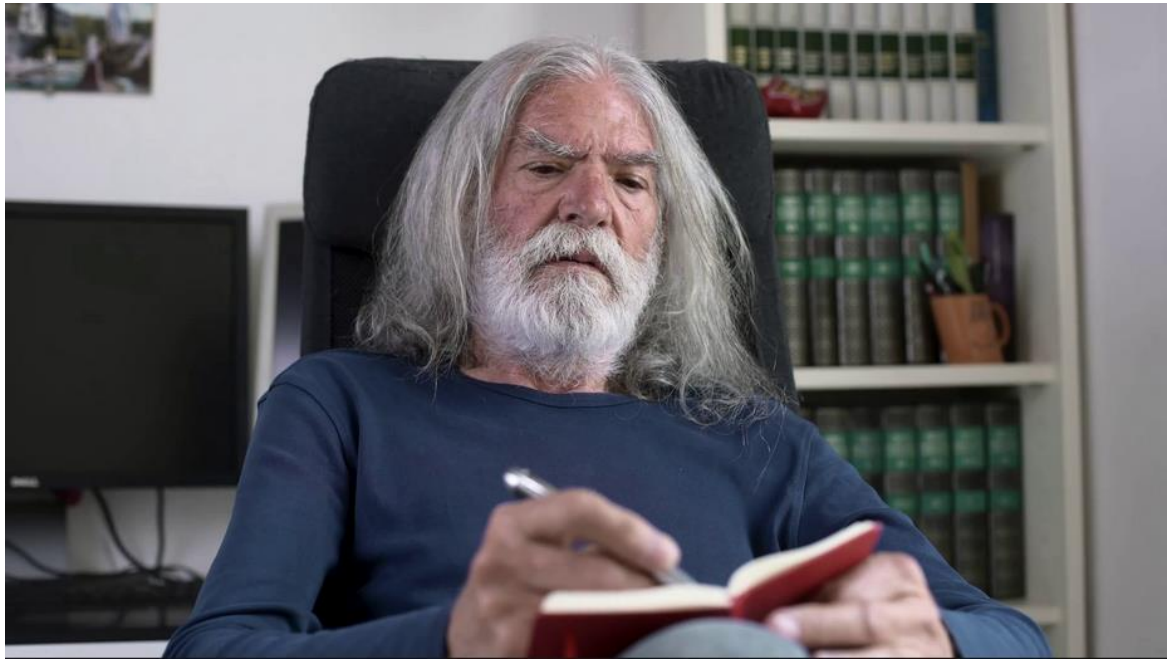

"Old man writes the agenda" Stock photo

3shape

About the speaker and 3Shape

Whoami:

- Team lead of 8 experts in machine learning and algorithm development
- 5+ years experience in making medical software (Oticon R&D, 3Shape R&D)
- Ph.D. in math, signal processing, mathematical modelling of auditory perception
- Working with orthodontic software in 3Shape

Whoarewe:

- 3Shape A/S, Danish medical device company
- 1500+ employees, HQ in Copenhagen (500+ employees, 30% in R&D)
- Visible light and x-ray scanners for teeth
- Software for workflows in dentistry
- Small scanner section for hearing aids and high-end earphones

Data

Main data source: Output from the Trios Interoral Scanner - 2D surface scans in 3D space

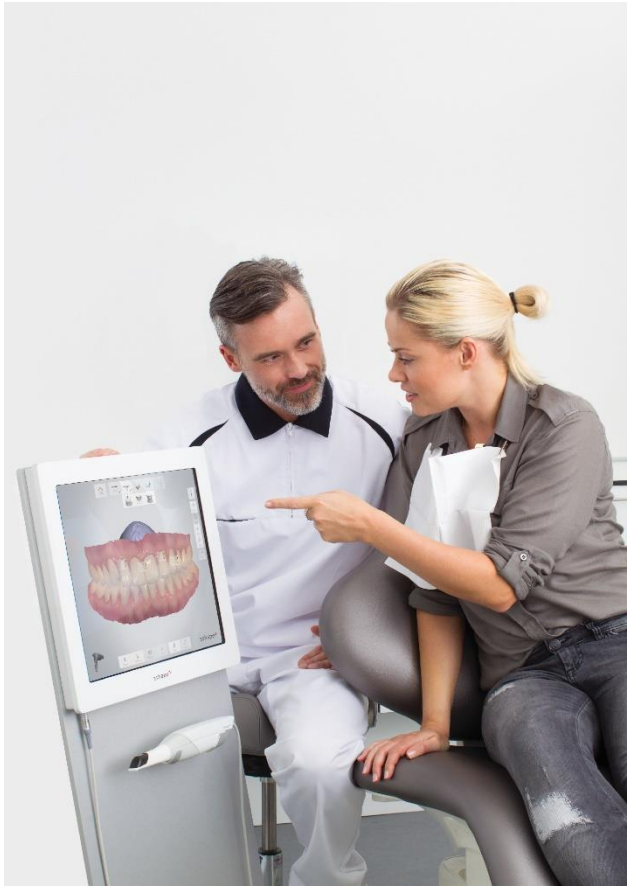

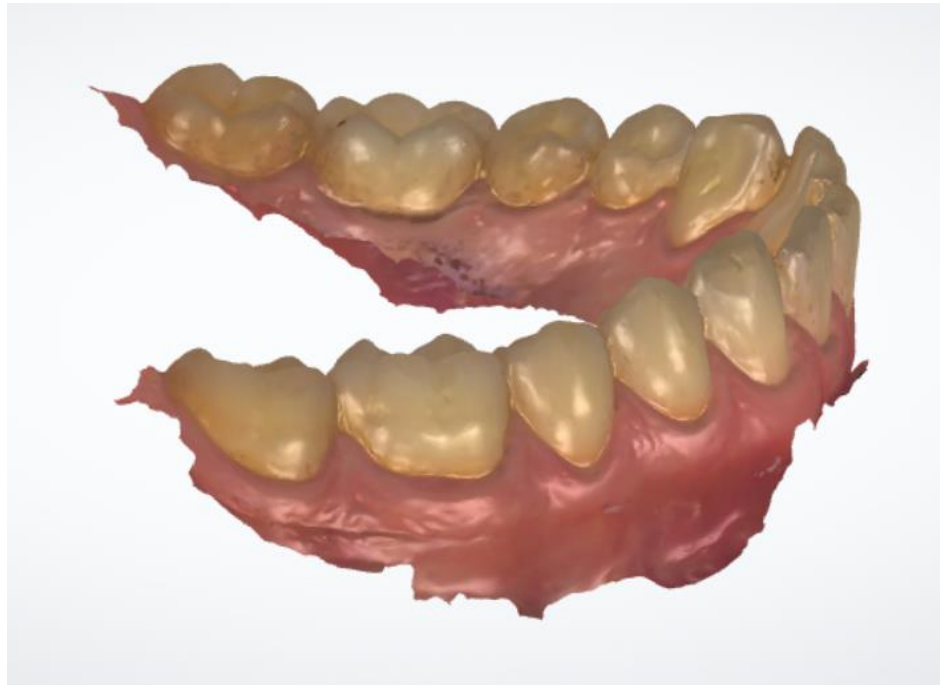

An example machine learning application: Segment a jaw scan into gingiva (gum) and individual teeth

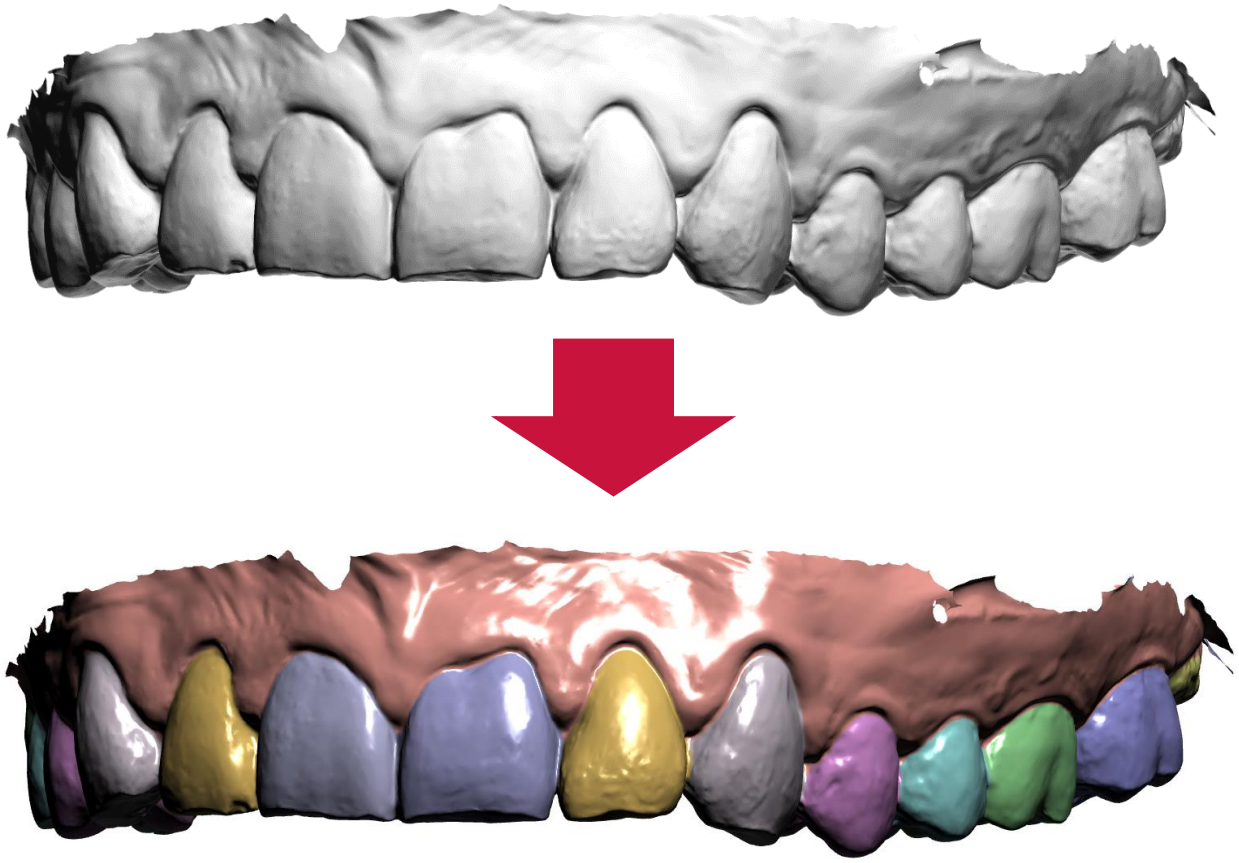

The challenge of developing medical software

When developing medical software, you get to play with FDA and EU

Why? Because it matters what you do! You can cure people of diseases and physical problems or improve their quality of life

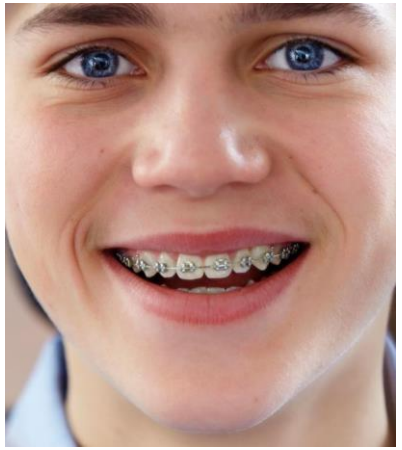

Matters!

Doesn't matter

Medical Device Regulation

Do Not Fear the FDA and EU

Normal reaction: FDA and EU just adds bureaucracy

Remember, we have the same goals!

1. Effective software, it should work
2. Safe software, it should not harm people
3. Fast delivery, people shouldn't wait years for new treaments and technological advances

Make these goals your own, then you are compliant by default!

(and your devs do not need to understand the guy to the right)

3shape

So, what's the big deal

Machine learning poses a new challenge for medical regulation:

- It's essentially a black box. Even more so with deep learning. Trust = 0
- Traditional verification is done through inspection and analysis, logic reasoning

What can we do about this:

1. Use the scientific method
2. Restrict your input to the black box
3. Restrain the output from the black box

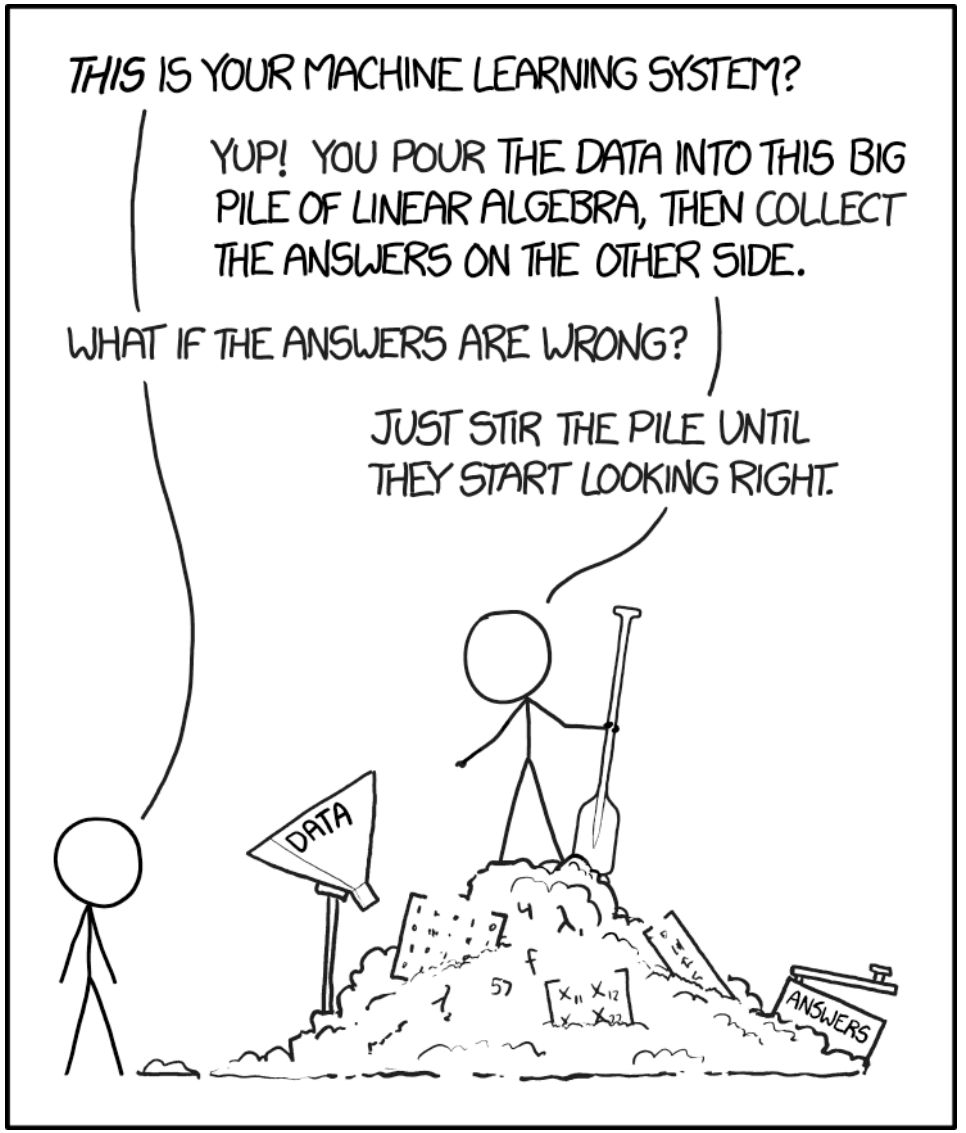

3shape

The scientific method

- Actually practise what you learned in school and university
- If you fool yourself, you might fool the company and the medical authorities, but you will not fool the world-wide deployment
- Fixing the bug beforehand is much cheaper than dealing with the fallout of an incident in EU / US

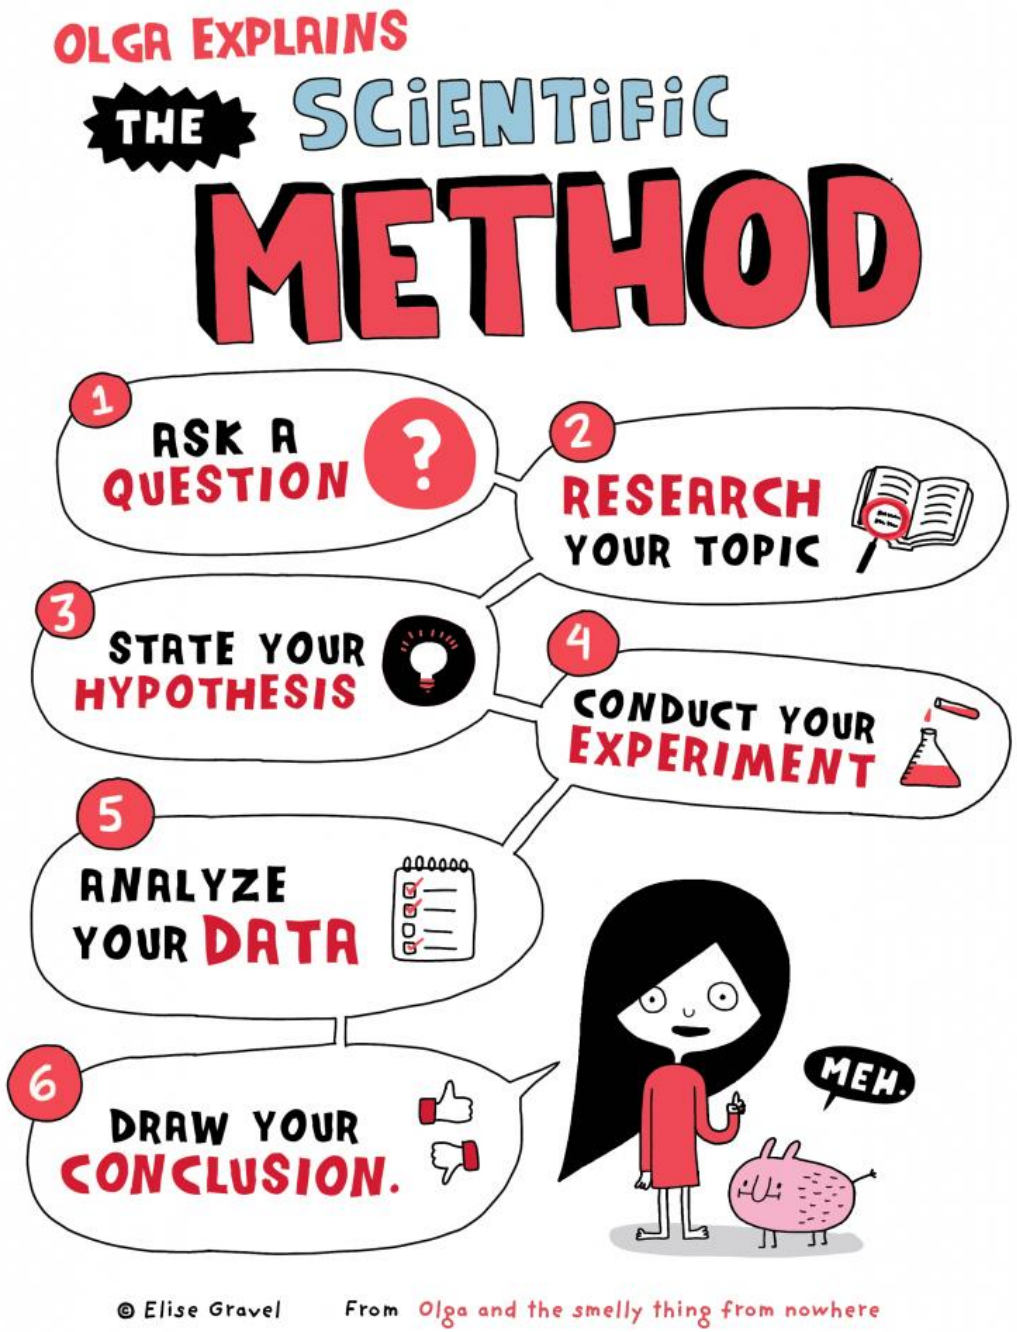

Restrict the input

- Define a ”target population” to remove any elephants in the room
- Describe it in the IFU: ”Instructions For Use”
- Choose you datasets from this population
- Look out for the residual risks:
  - Your test set does not follow the same distribution as the actual usage of your system
  - The IFU is not followed

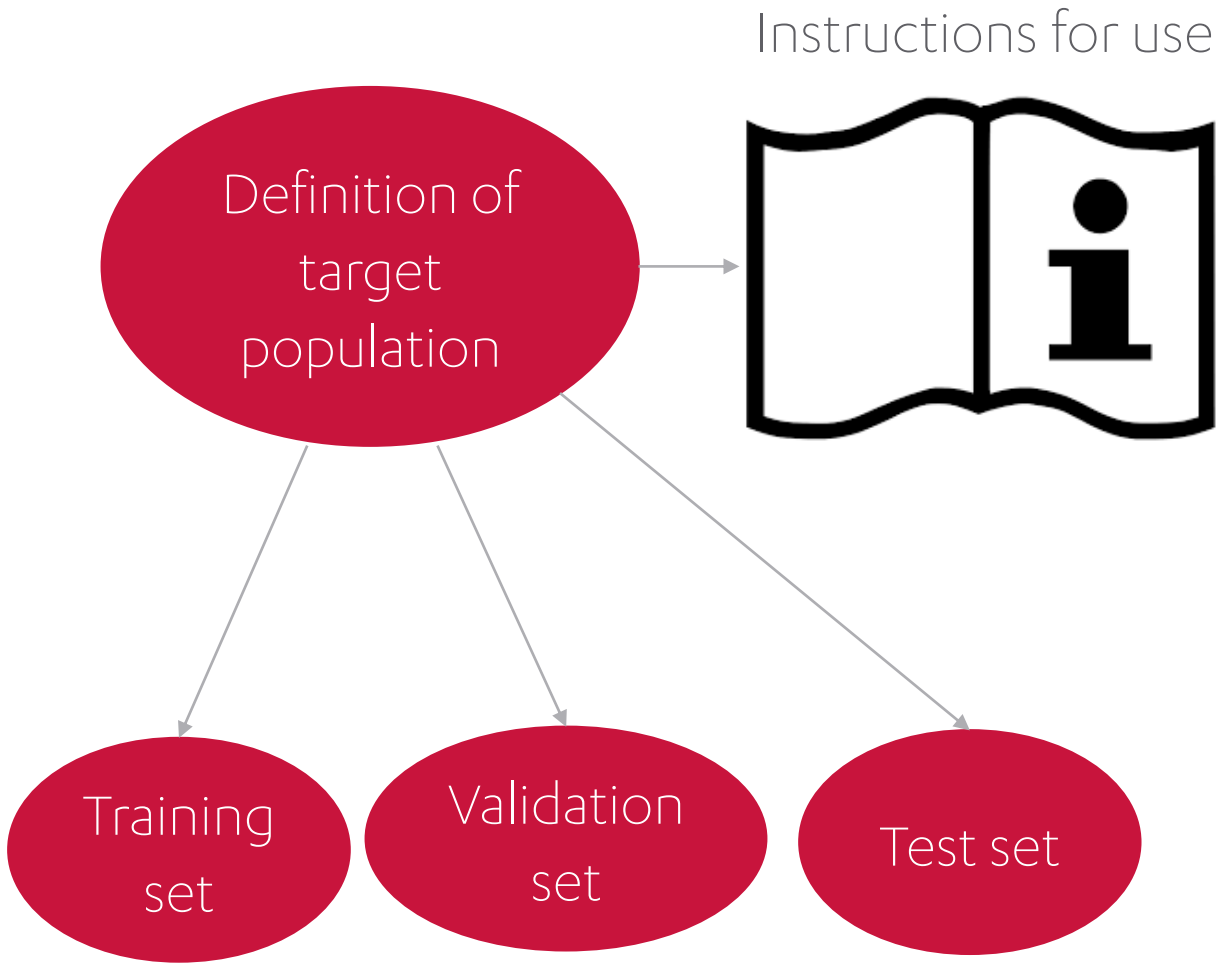

3shape

Restrain the output

You known your domain, put in logic to control edge cases

- If you know something is physically impossible, detect it and act on it, i.e. a tooth is too large:
  - Re-evaluate the output or
  - Fail gracefully.

- About failing gracefully
  - The AI system is always an assistant to the doctor or technician, so it is a viable option
  - A gracefull failure just means a delay in treatment time, an uncontrolled failure may lead to patient harm

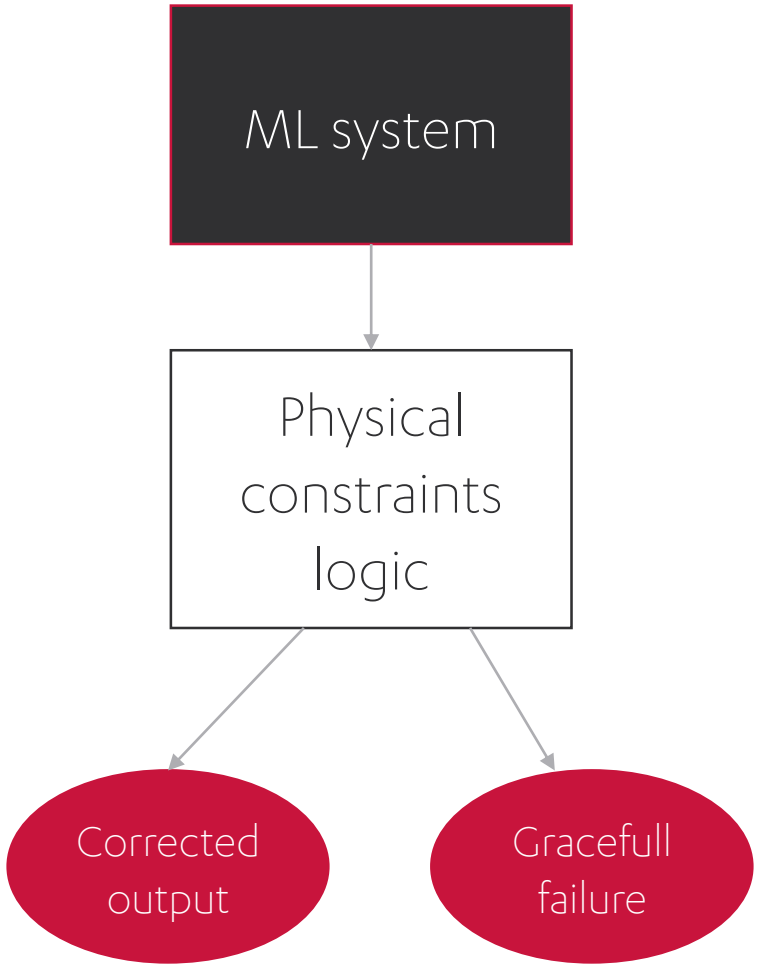

3shape

An example problem: Weird edge case patch

Increate performance
or increase trust?

3shape

Study human performance

You will never get 100% performance. Deal with it:

- Determine what performance level is safe, what level is effective?
- Measure the performance of humans on the same task

3shape

Make it fun for the developers

How to get developers to write documentation?

- Code it in Python using Jupyter notebooks
- Create the notebooks as part of the development process to document performance and experiments
- Use a modern documentation system like Atlassian Confluence for documentation of experiments, architecture overview etc.

Do it the Agile way, also the documentation

- The usual way (wrong)

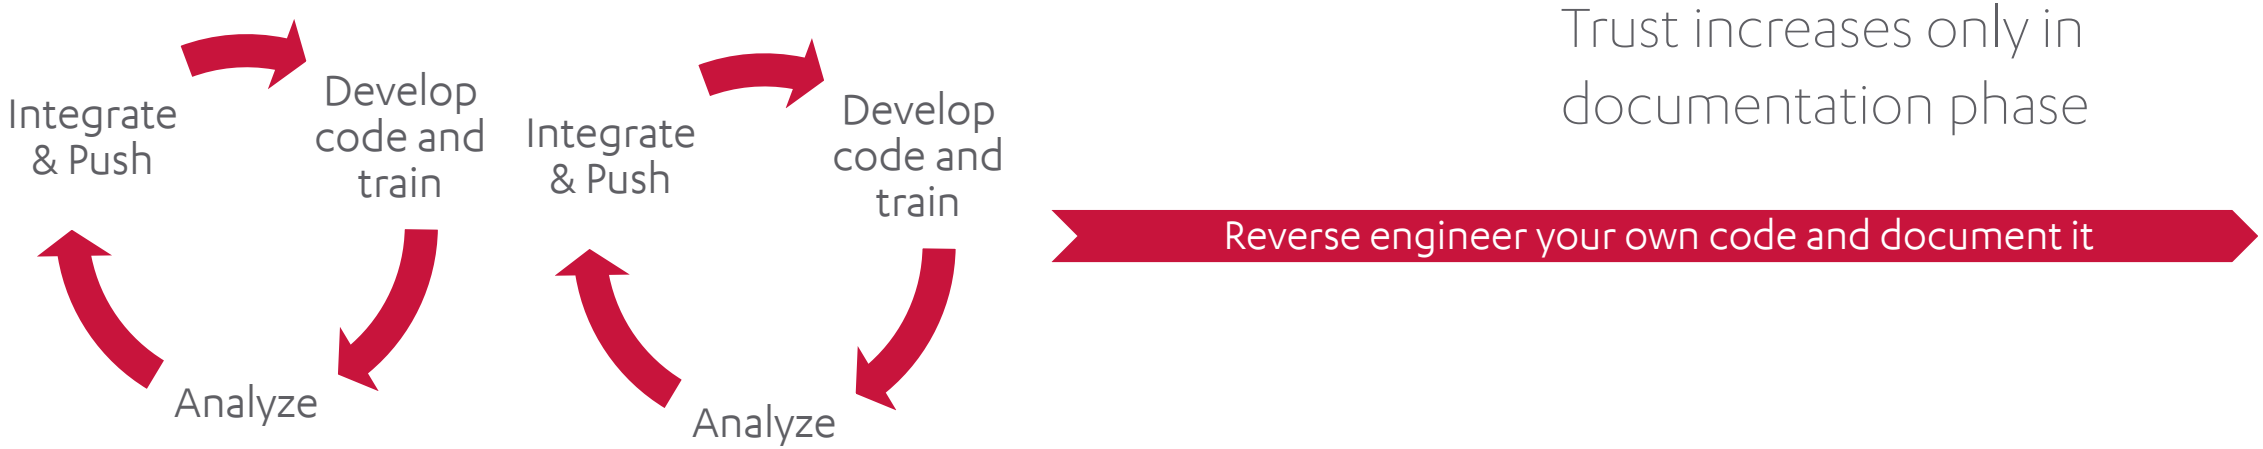

Trust increases only in documentation phase

- The right way

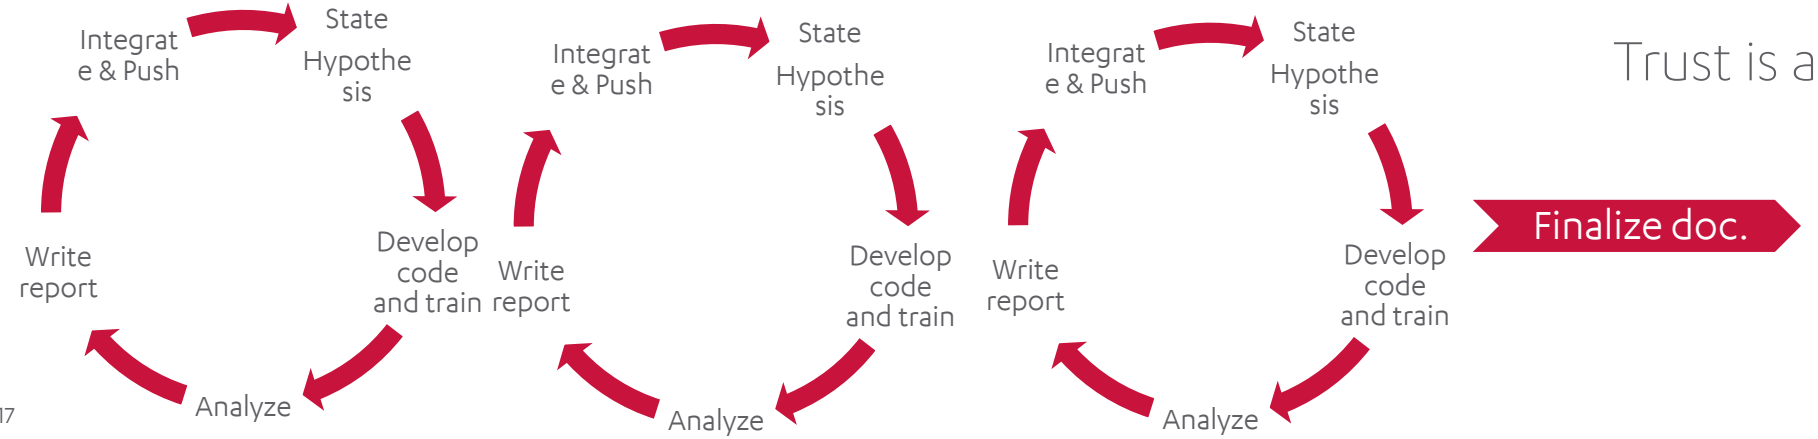

Trust is always high

3shape

Differences when working with machine learning and medical software

Medical device	Snapchat
Restricted input	Unrestricted input
Trust over Accuracy over Speed	Speed over Accuracy over Trust
#failgracefully	#failfast
Clinical trials	Live A/B testing, Chaos engineering
Permantly improve quality of life	Just for fun

3shape

Our pipeline

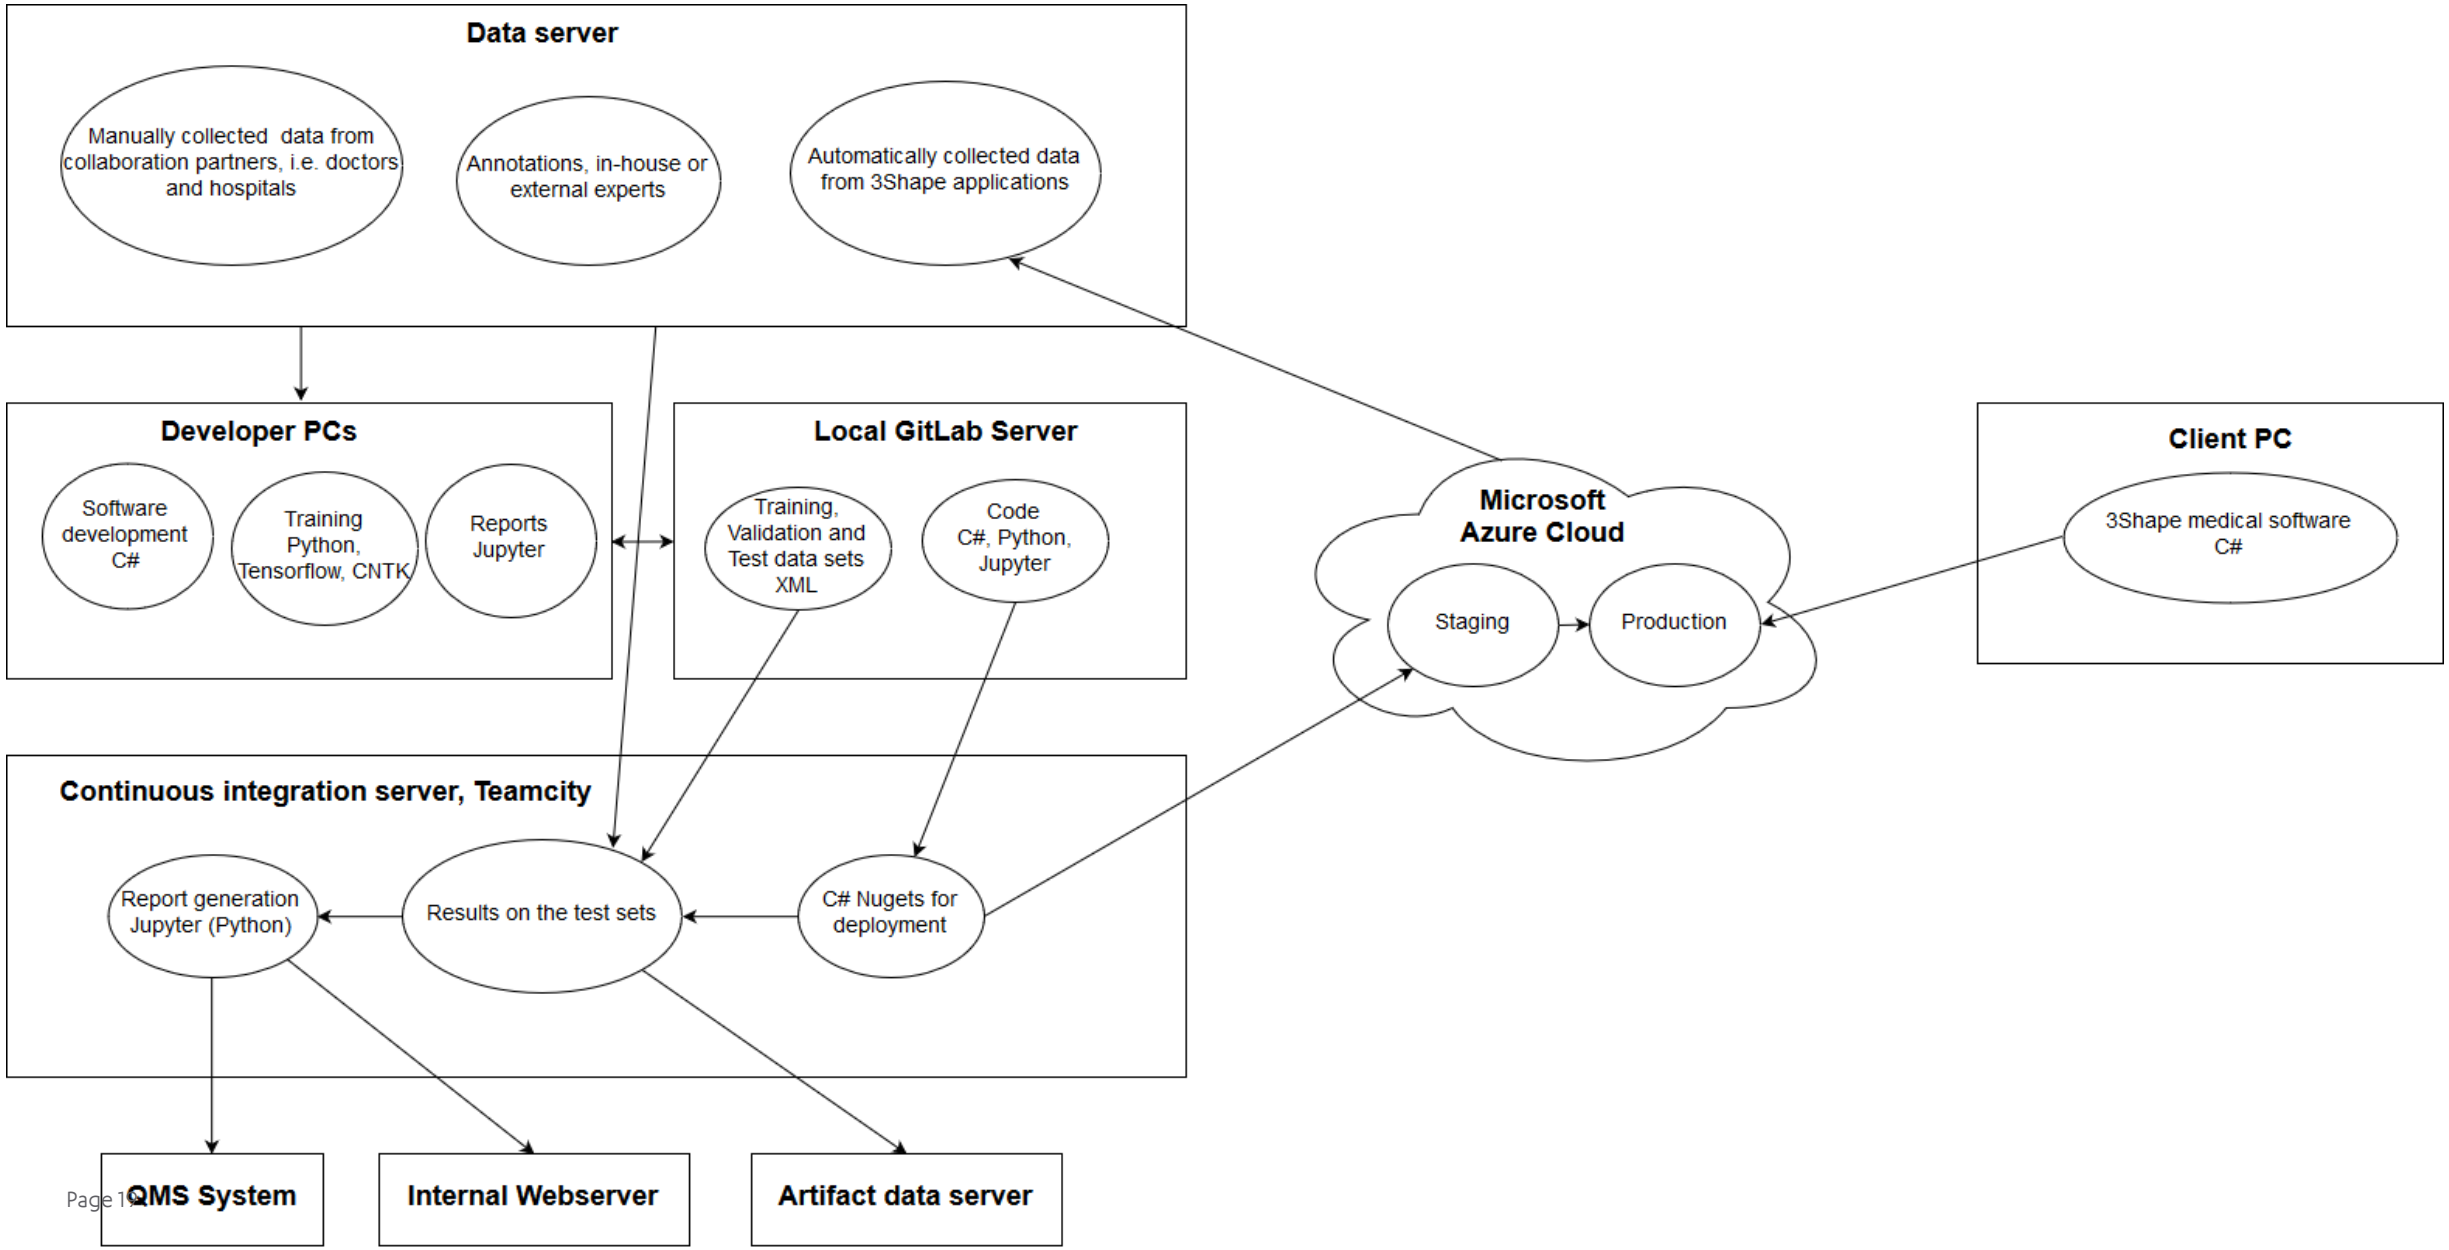

Results

- When training is done
  - Documentation is generated nightly
  - Reports are inspected and signed
  - Project manager gives ”go” for release
  - Nugets are uploaded to the cloud, deployed software on client PCs automatically starts using the cloud system

- Result: Total release process (for a small update): 1-2 days

Thank you very much for attending this talk

Please drop by our booth

Questions?

goto; copenhagen

GOTO Copenhagen 2018

Conference Nov. 19 - 21

Please

Remember to rate this session

Thank you!

Follow us @gotocph