AI Archaeology: Reading Ancient Texts With Cutting Edge Technology

Casey Handmer, Julian Schilliger, Youssef Nader

TNG Big Techday Munich 2024

You?

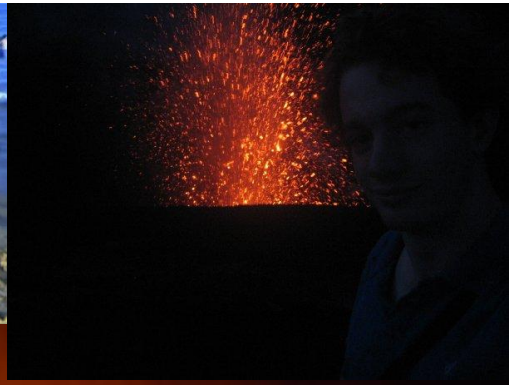

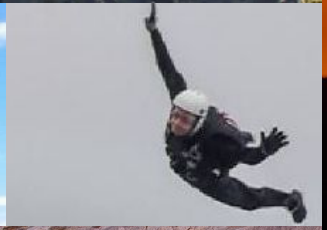

Classical and

PAPER

Spectral C

gravitatior

Casey J Handmer

Published 21 Octob

Classical and Quant

— Casey Handmer (Caltech), Bela Szilagyi (JPL) and Jeffrey Winicour (Pittsburg) reprise their former stance discussing asymptotically time-like inertial scri+ foliations, now with even better CGI. Image credit: Photo manipulation by Annie Handmer, background image by SXS Collaboration: Andy Bohn *et al* 2015 Class. Quantum Grav. **32** 065002.

nd

As an

the fir

Hyperloop One

Levitation Engineer

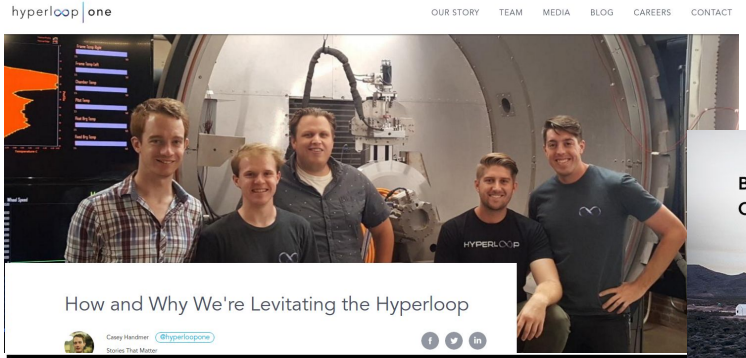

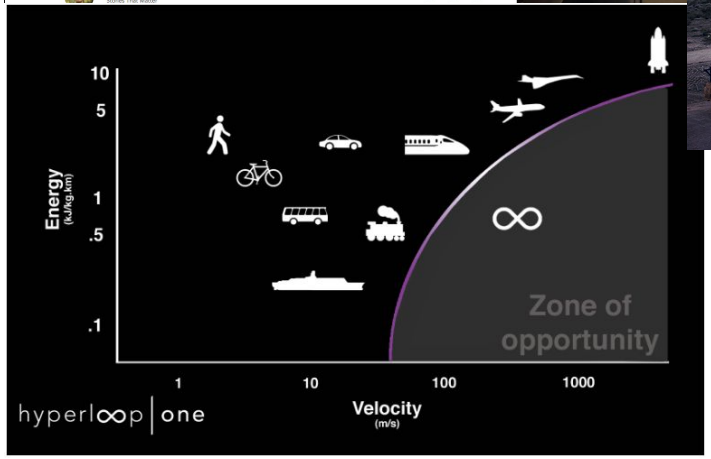

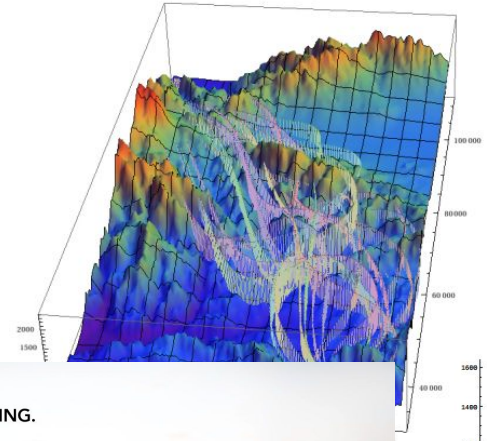

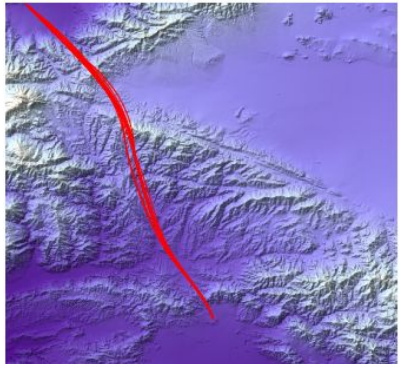

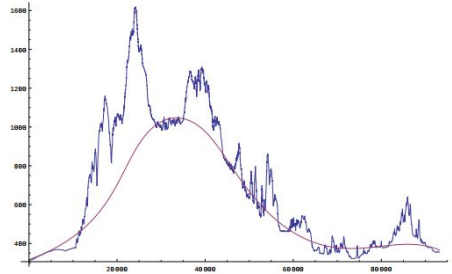

NASA Jet Propulsion Laboratory

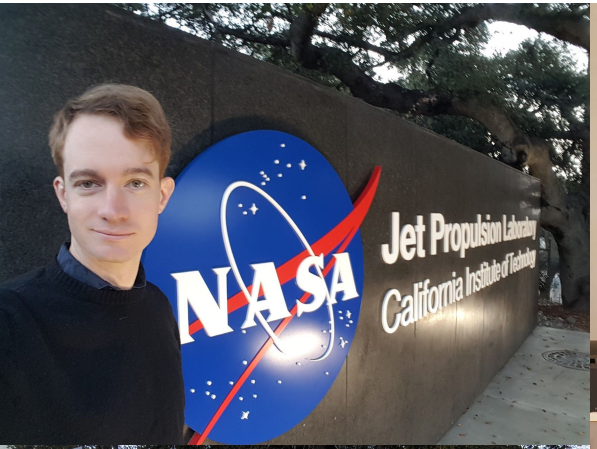

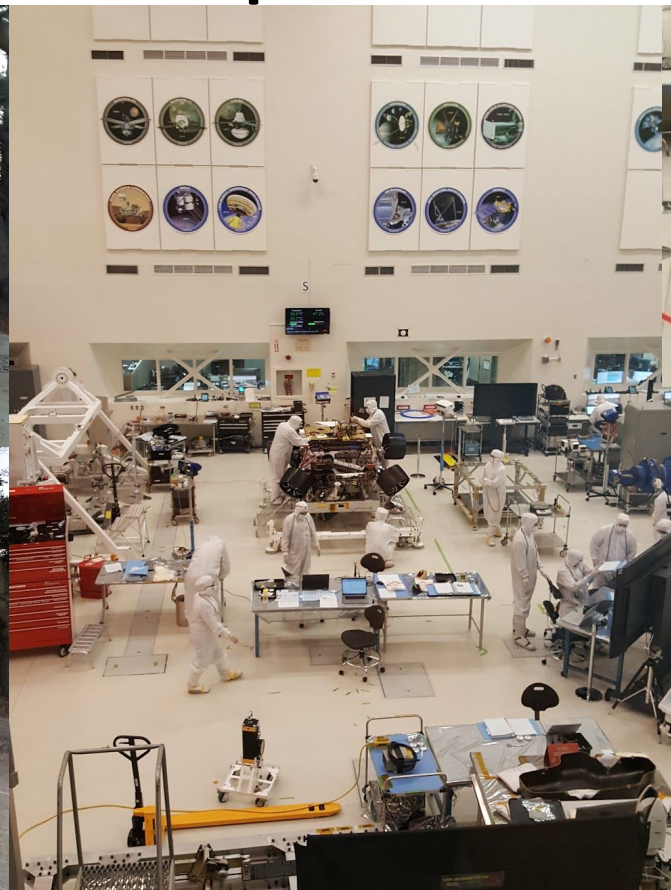

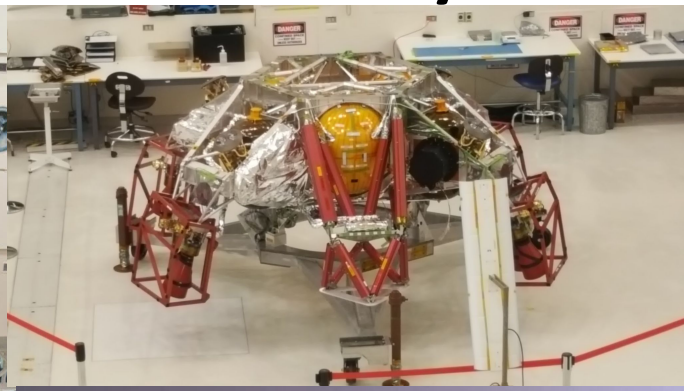

Founding Terraform Industries

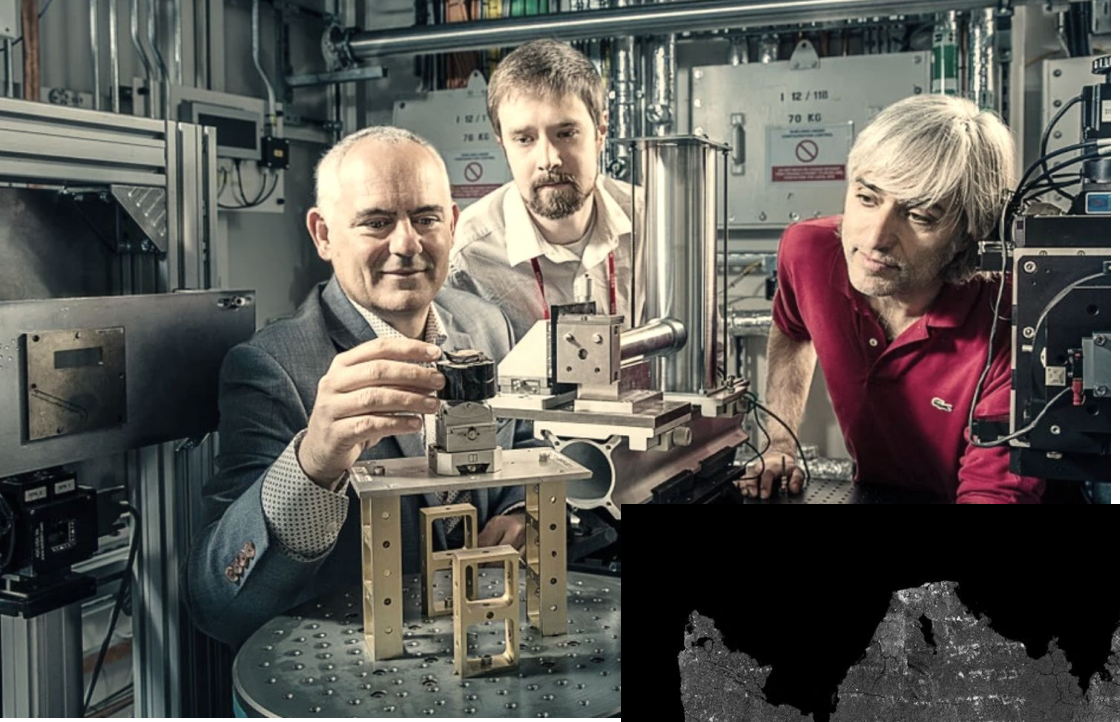

IN GOD WE TRUST
LIBERTY
2014
S
≈1 cm

Vesuvius Challenge
Background
2023 Grand Prize Awarded
Overview
2024 Prizes
The Master Plan
Prize Winners
Community Projects
Data
Scrolls
Fragments
Segments
Tutorials
1. Technical Overview
2. Scanning
3. Segmentation and Flattening
4. Ink Detection
FAQ
History
Livestreams

Resurrect an ancient library from the ashes of a volcano.

Win Prizes. Make History.

The Vesuvius Challenge is a machine learning and computer vision competition that in **2023 cracked the riddle of the Herculaneum Papyri & awarded over $1,000,000 in prizes.**

2024's challenge is to go from reading a few passages to entire scrolls.

$500,000+ in New Prizes

MORE INFORMATION →

Read the Master Plan

READ THE POST →

WHAT GIVES PEOPLE FEELINGS OF POWER

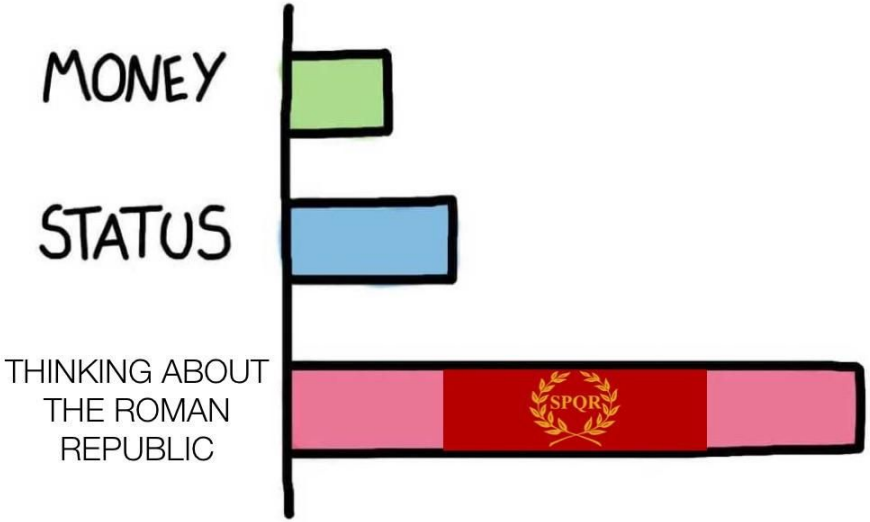

made with mematic

ANCIENT
HISTORY
www.ancienthistorymag.com

Parco Archeologico di Ercolano
Villa Campolieto
Villa Signorini Events &
Paninoteca Mato Mato
Lido Arturo
Turtle Point (Portici)
ENEA

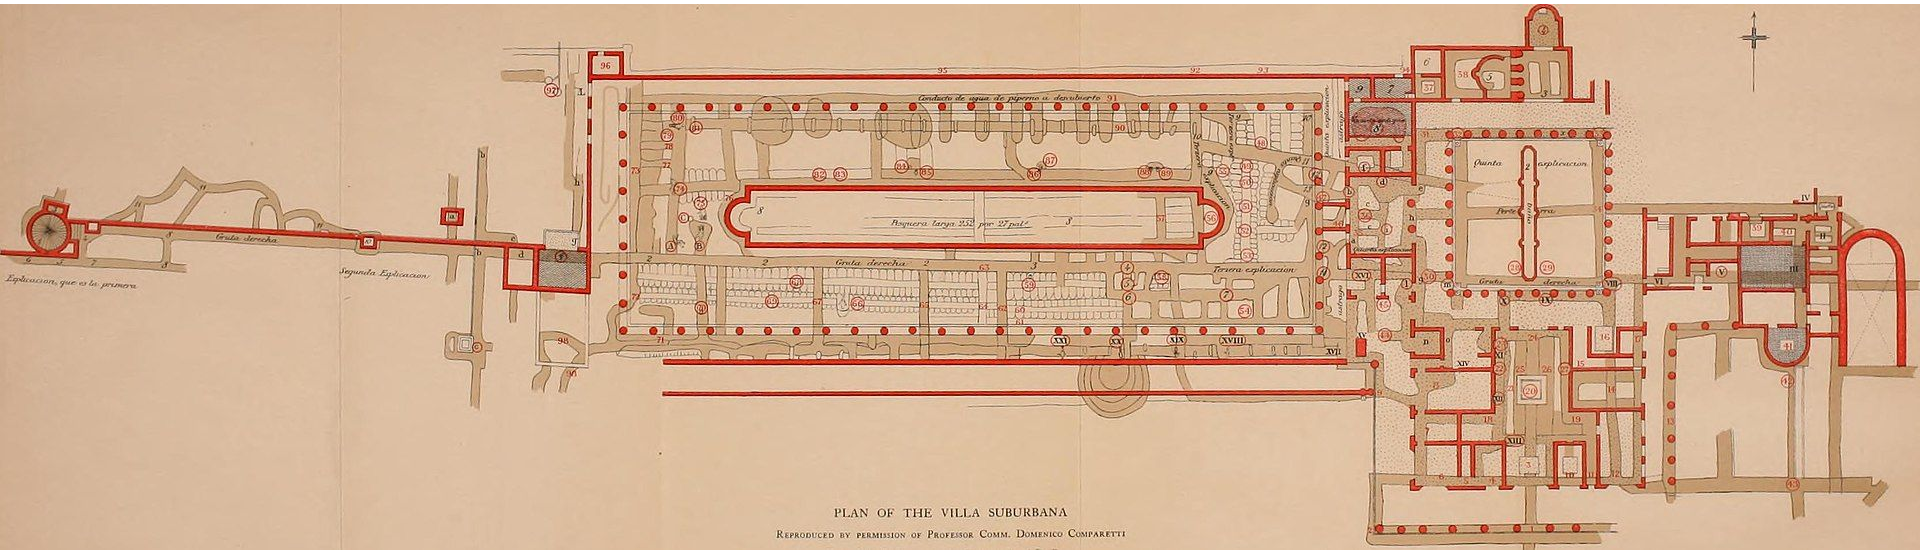

PLAN OF THE VILLA SUBURBANA

Reproduced by permission of Professor Comm. Domenico Comparetti

TARE
KEENOX

Papiro N° 10

Preso per svolgersi da D. Carlo Orazi in giugno 1865

Pezzo I

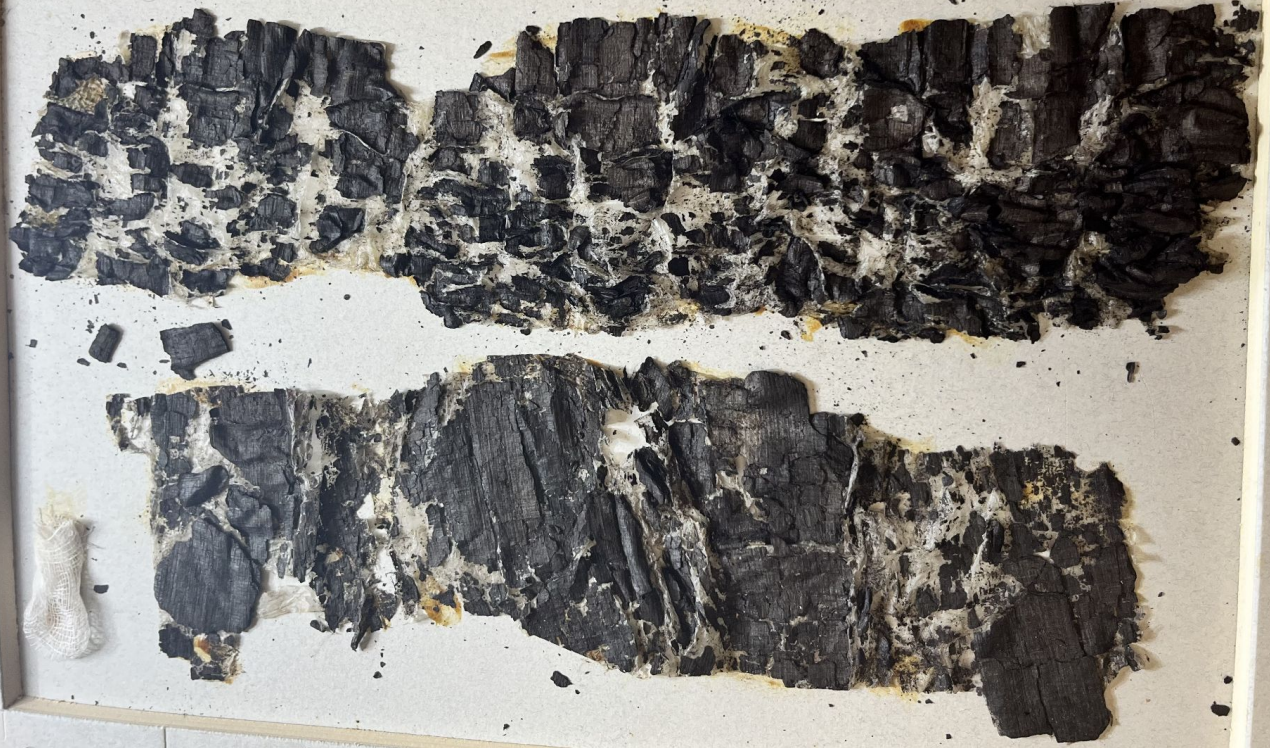

Conforti

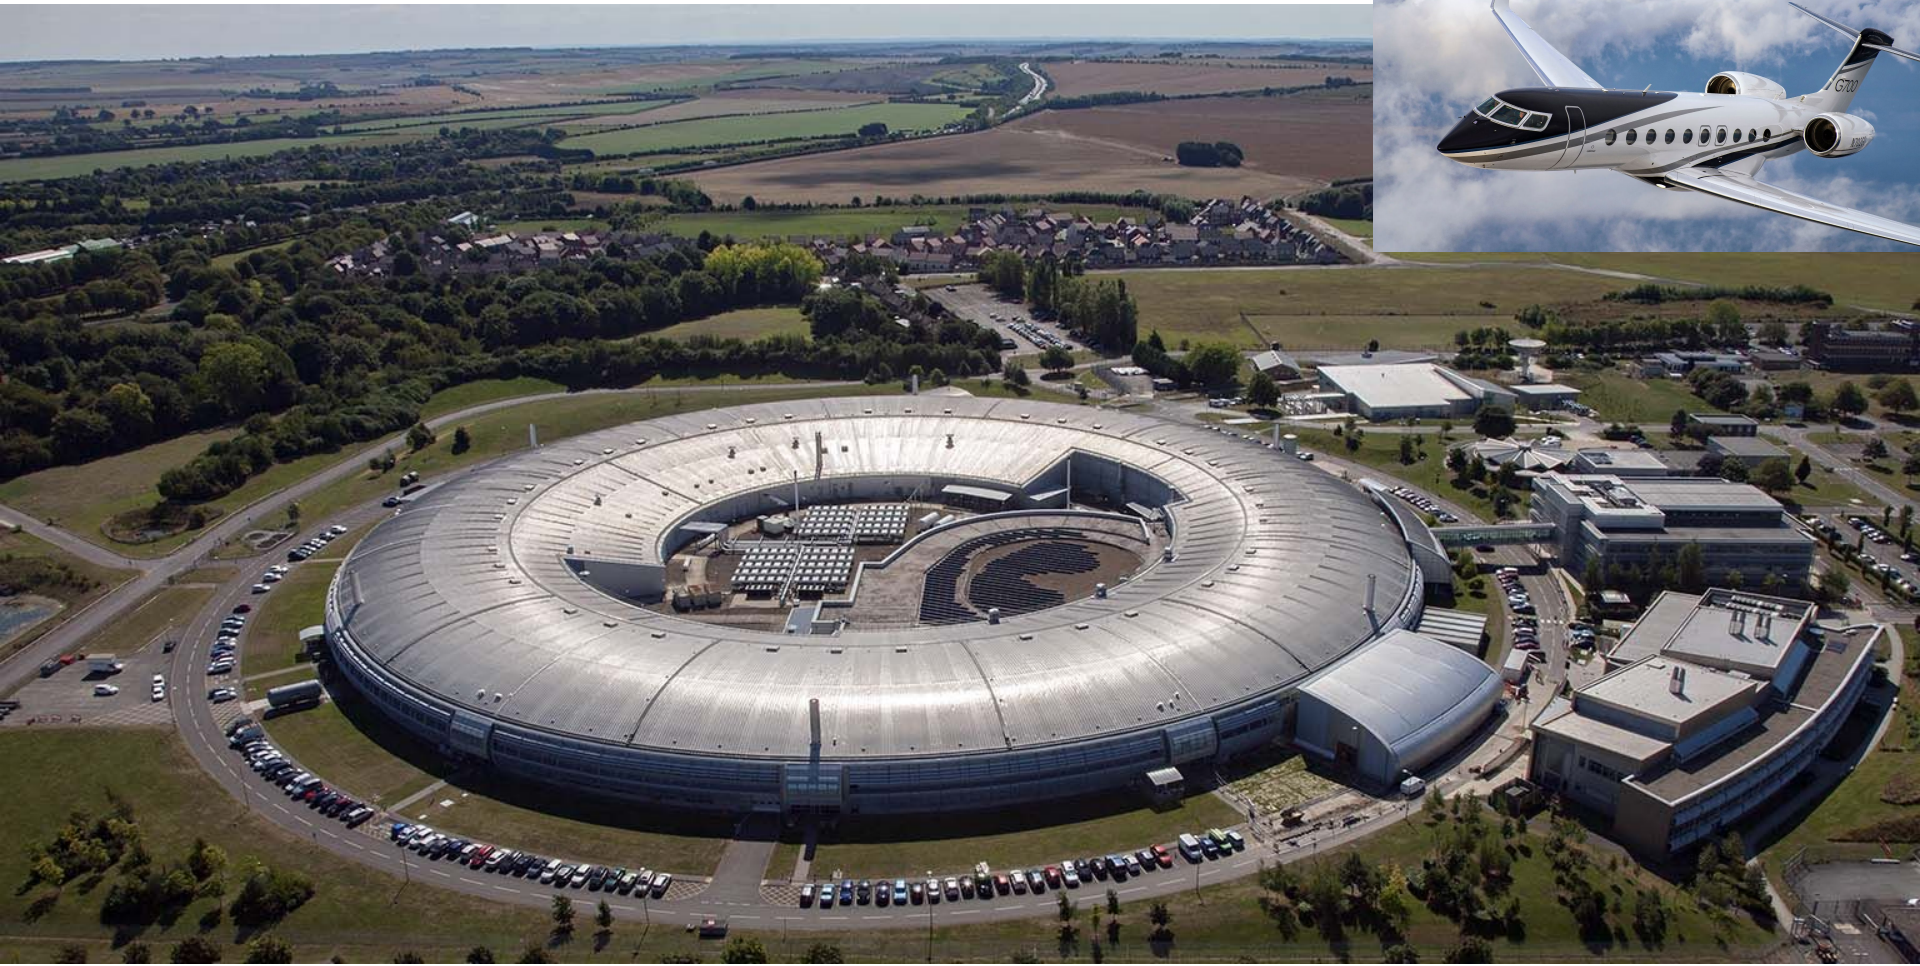
G700

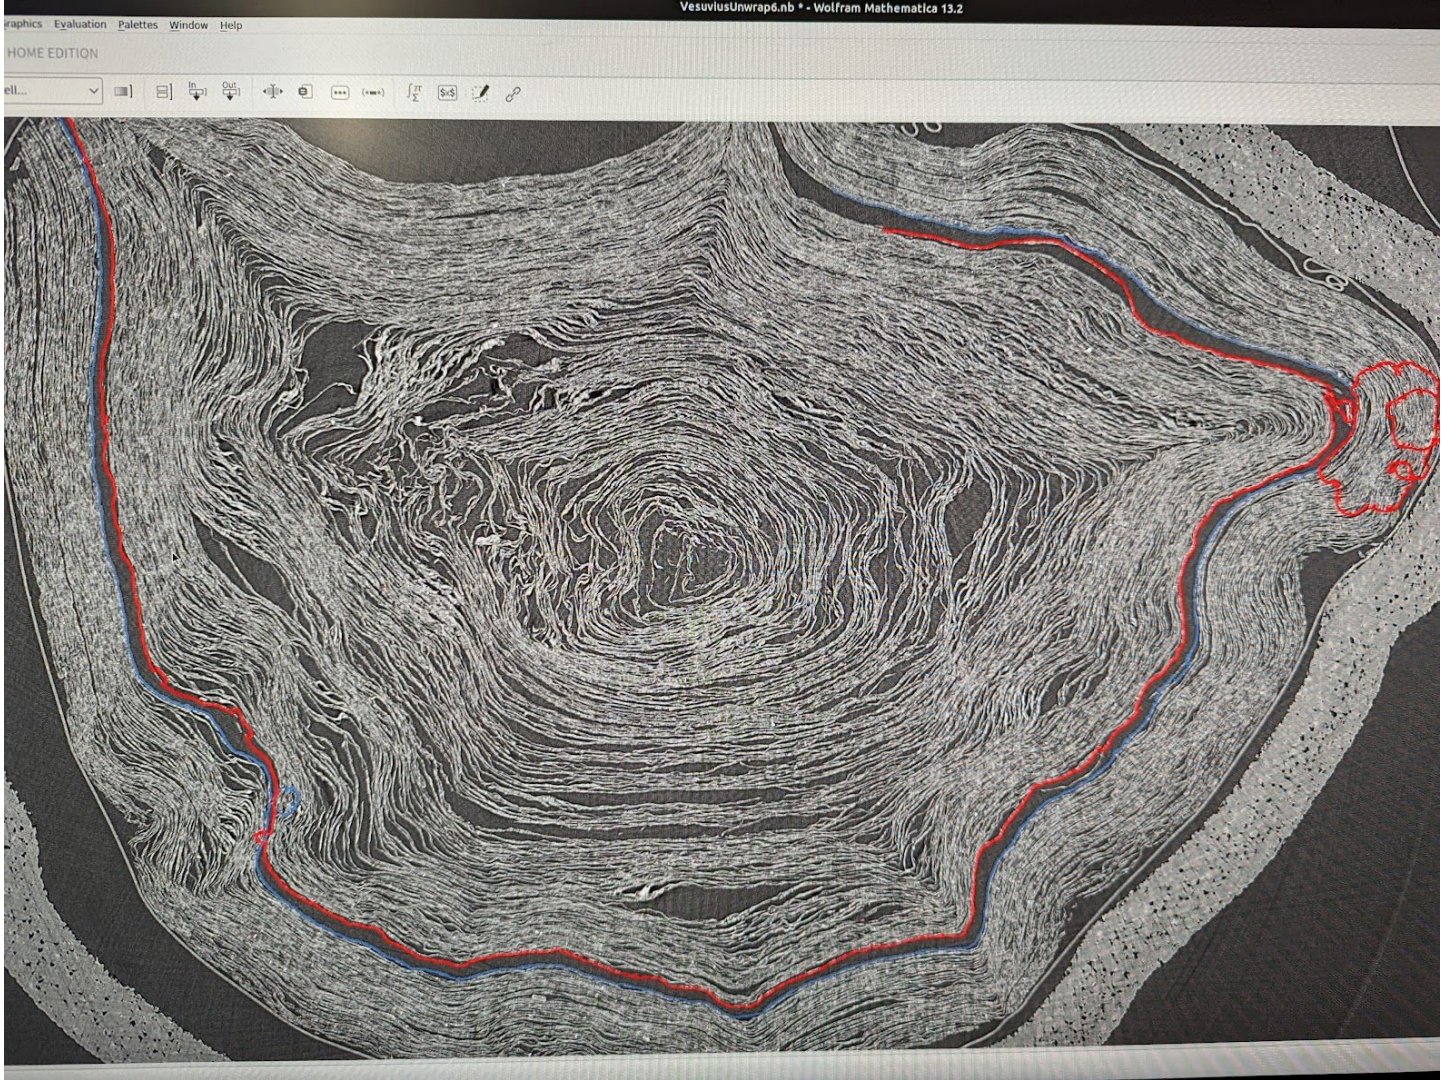
VesuviusUnwrap6.nb * - Wolfram Mathematica 13.2
Evaluation
Palettes
Window
Help
HOME EDITION

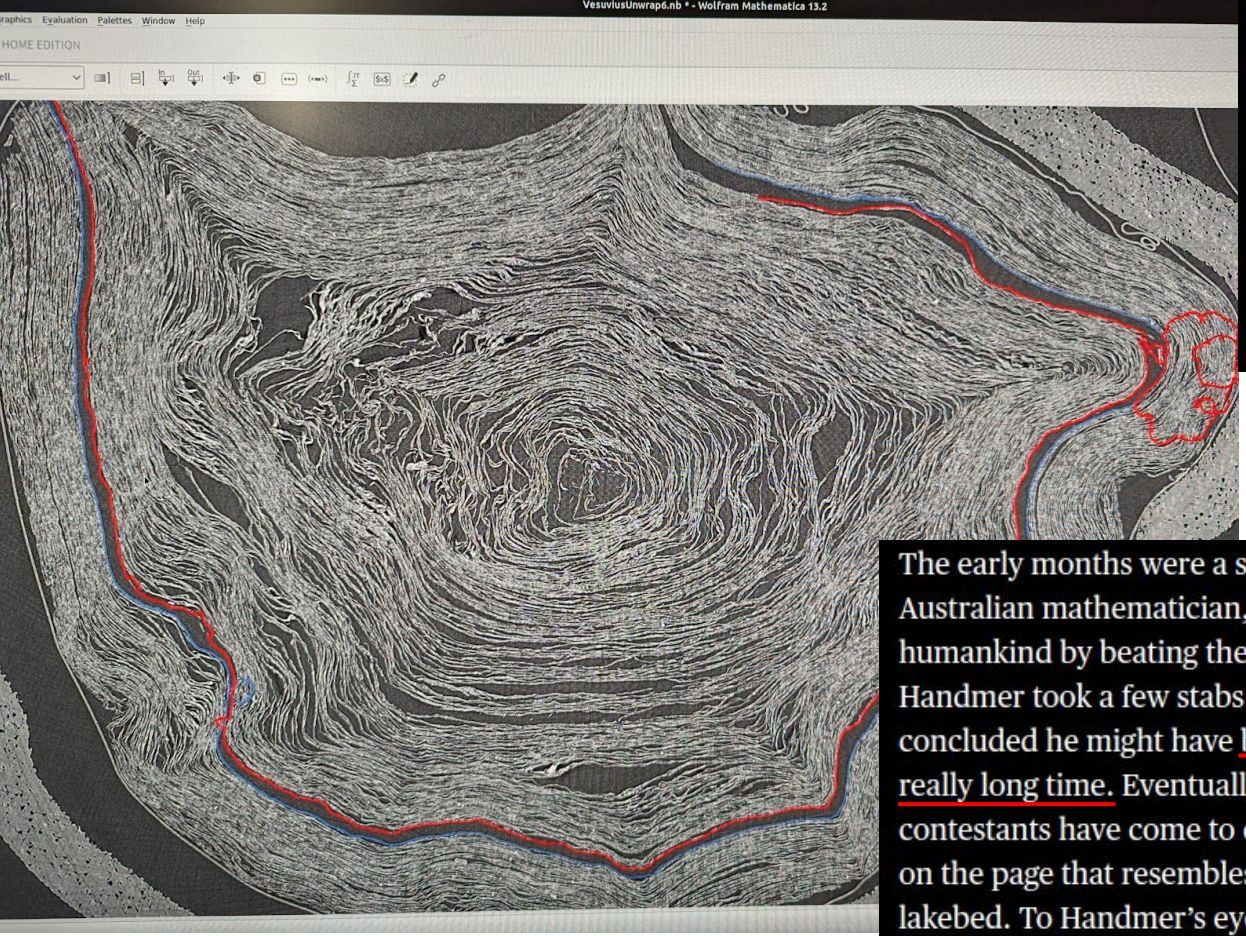

Businessweek

Can AI Unlock the Secrets of the Ancient World?

Almost 2,000 years ago, a volcano preserved Herculaneum's vast library of scrolls but left them unreadable. A volunteer army of nerds has been racing to decipher them.

By Ashlee Vance and Ellen Huet
February 5, 2024, 9:00 AM UTC

The early months were a slog of splotchy images. Then Casey Handmer, an Australian mathematician, physicist and polymath, scored a point for humankind by beating the computers to the first major breakthrough. Handmer took a few stabs at writing scroll-reading code, but he soon concluded he might have better luck if he just stared at the images for a really long time. Eventually he began to notice what he and the other contestants have come to call "crackle," a faint pattern of cracks and lines on the page that resembles what you might see in the mud of a dried-out lakebed. To Handmer's eyes, the crackle seemed to have the shape of Greek letters and the blobs and strokes that accompany handwritten ink. He says he believes it to be dried-out ink that's lifted up from the surface of the page.

Youssef Nader

Luke Farritor

Julian Schilliger

Elon Musk @elon... · 18 Jun

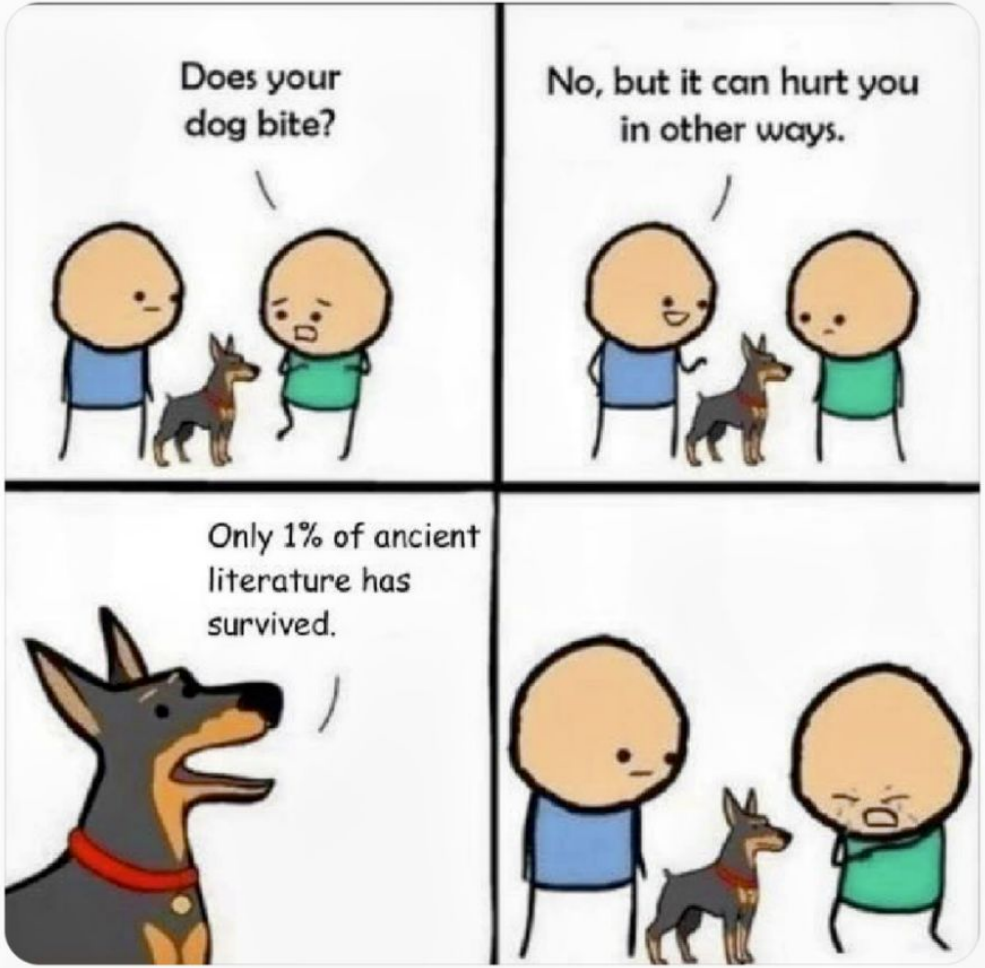

Quantity of ancient text by composition year

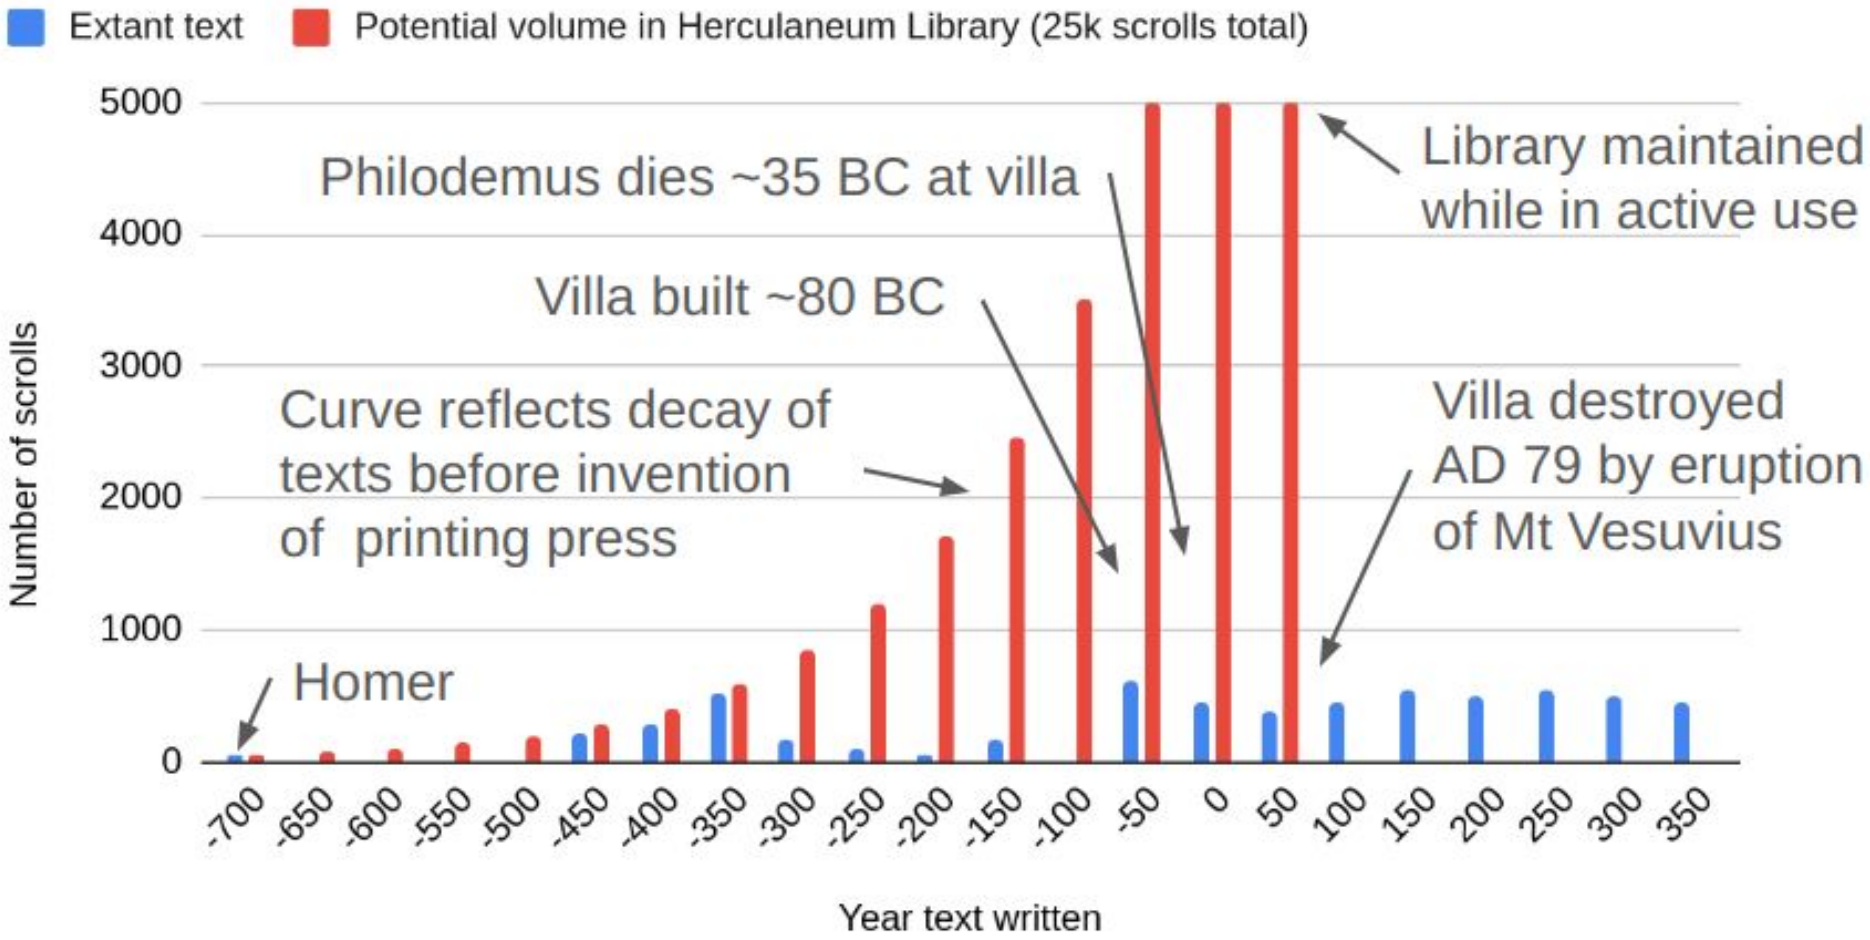

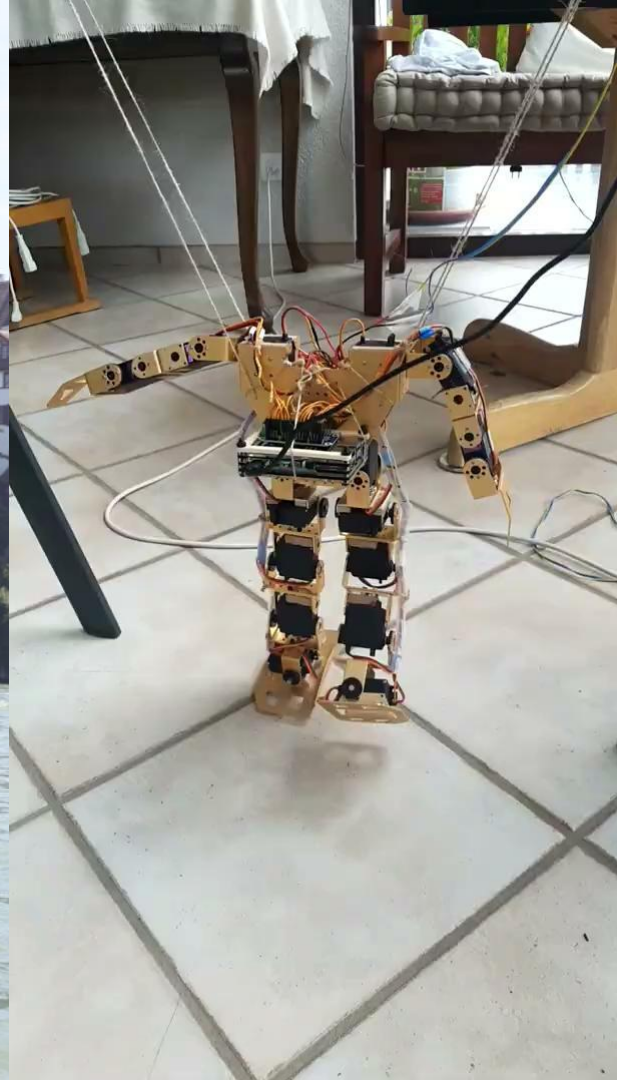

Vesuvius Challenge Grand Prize

Youssef Nader,
Luke Farritor,
and Julian Schilliger

ETH zürich

Floating Robotics

RSL Robotic Systems Lab

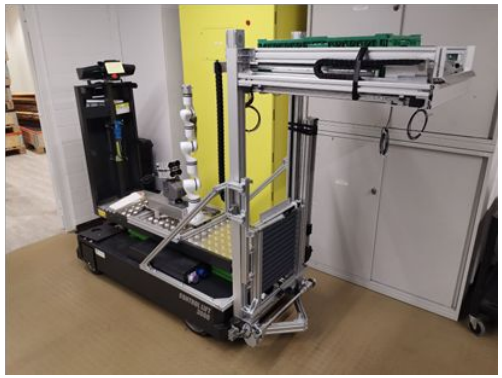

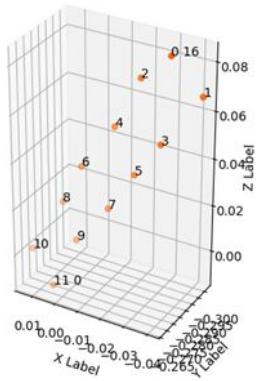

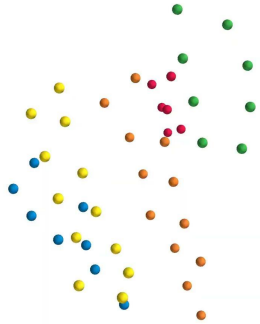

www.floatingrobotics.com

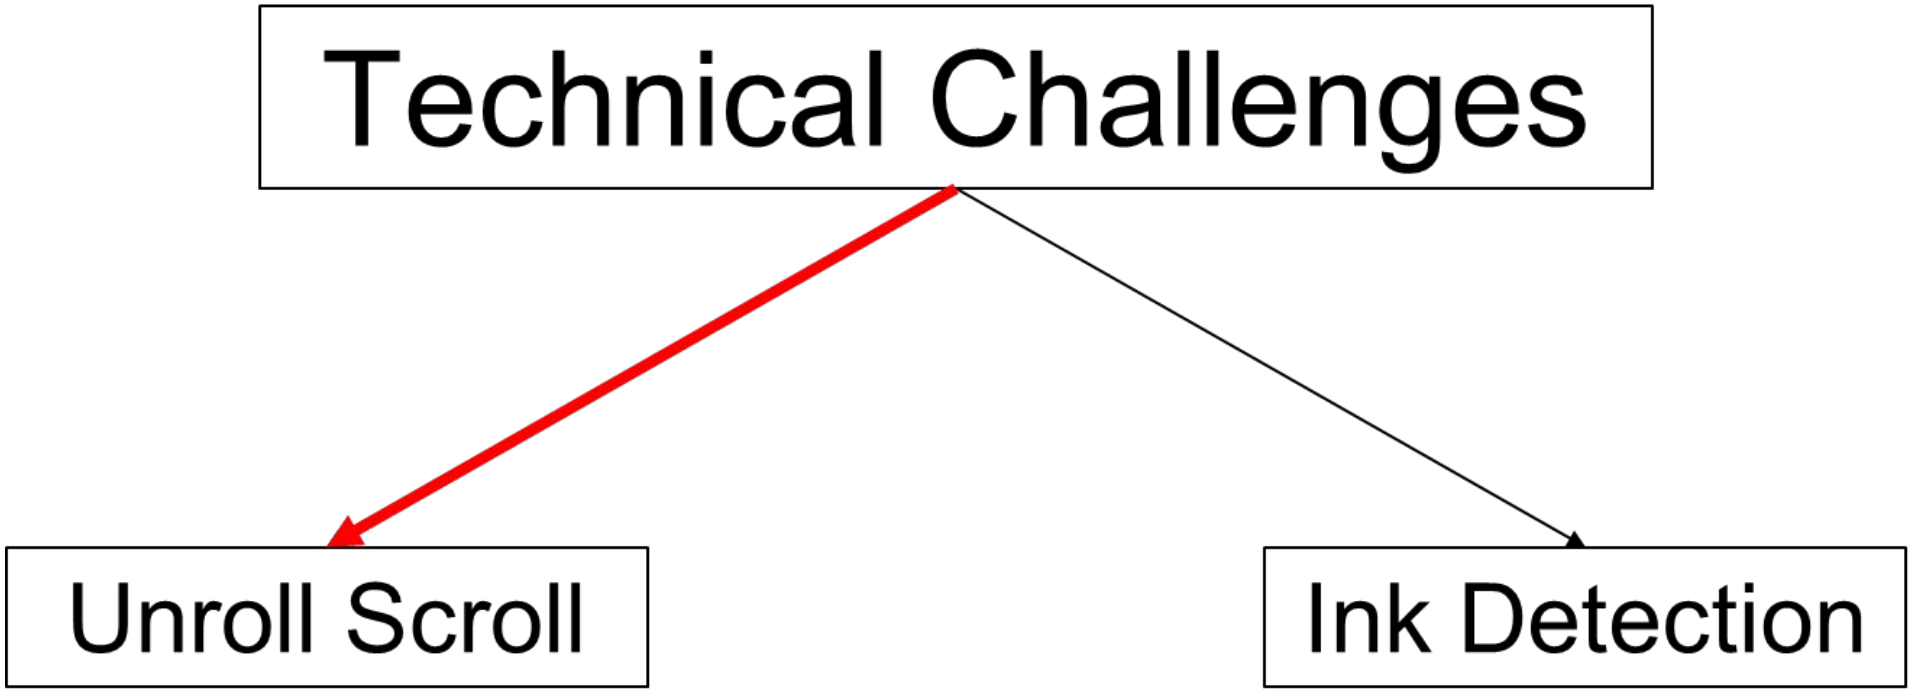
Technical Challenges
Unroll Scroll
Ink Detection

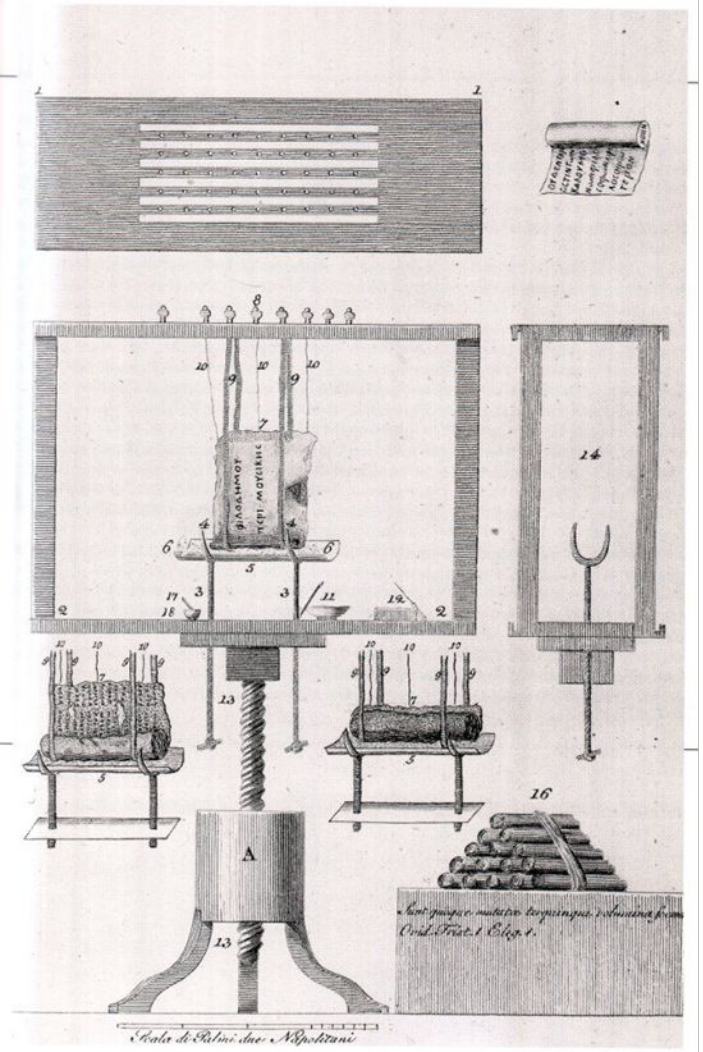

Piaggio's Machine

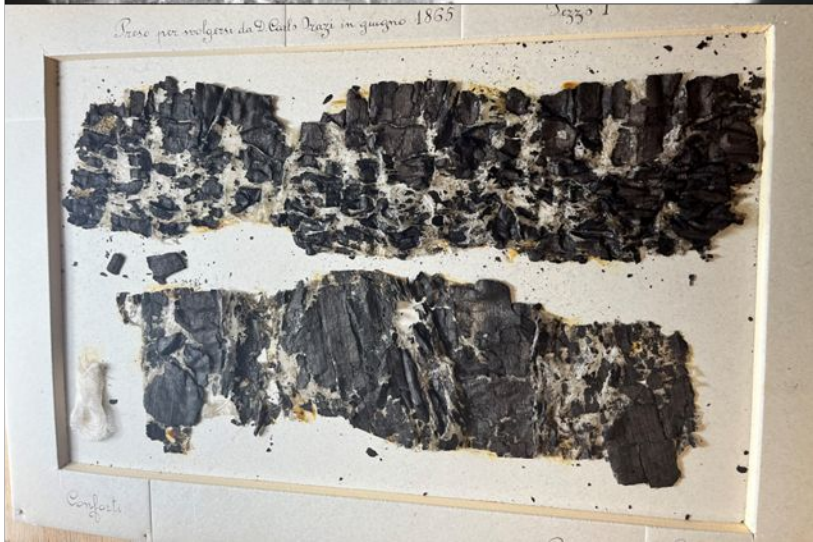

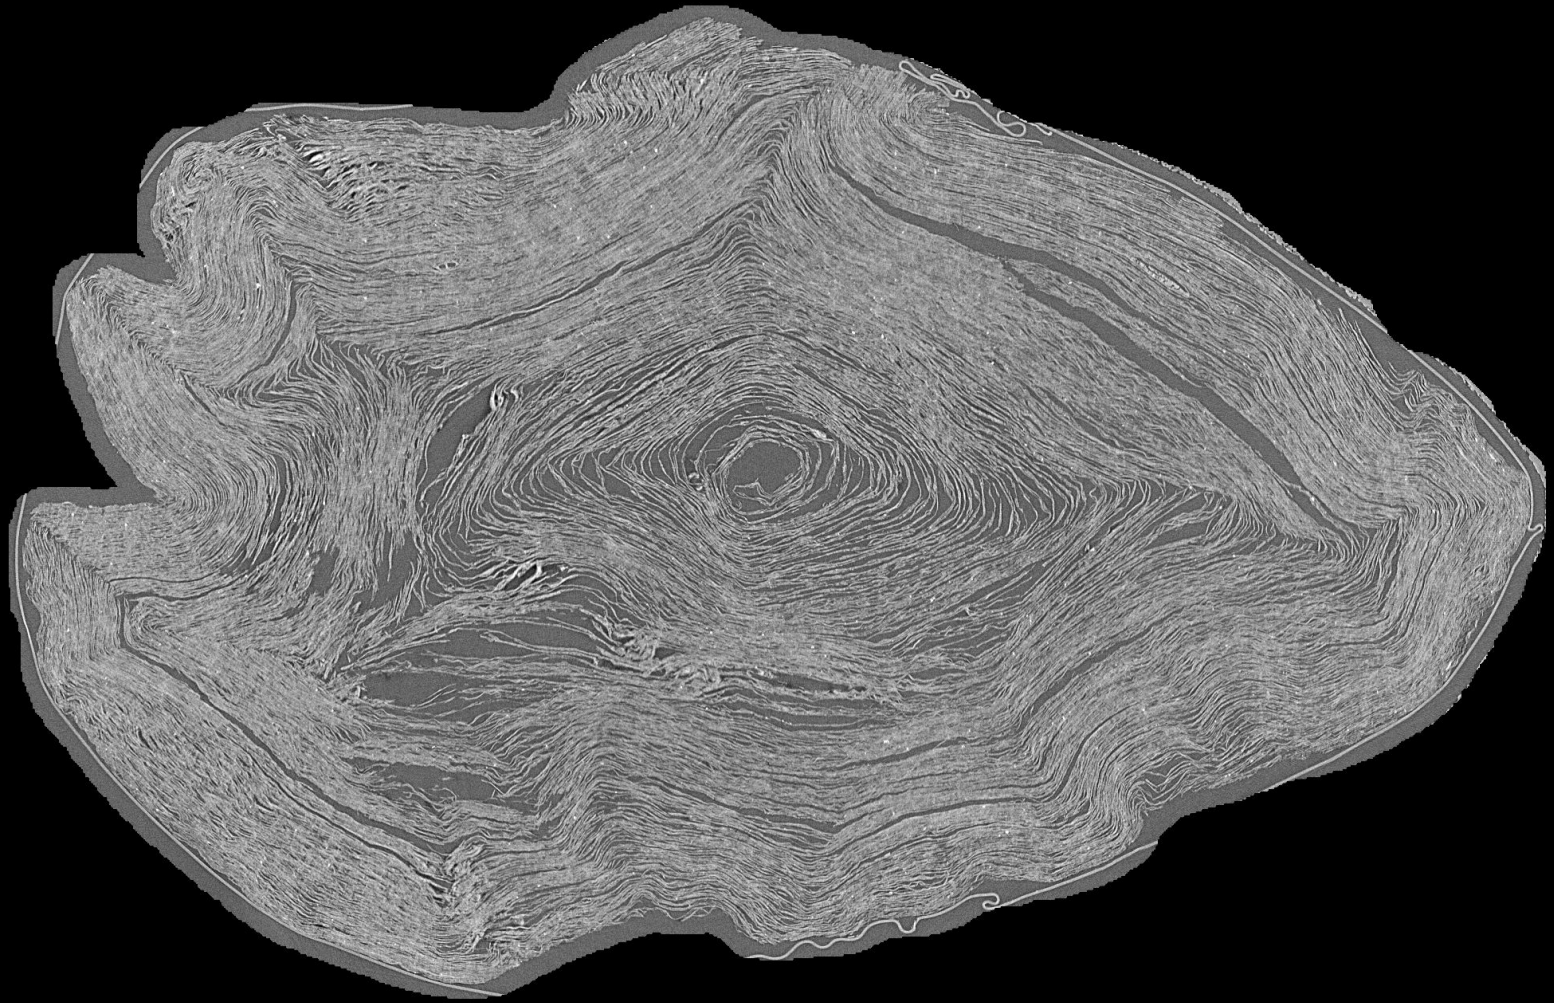

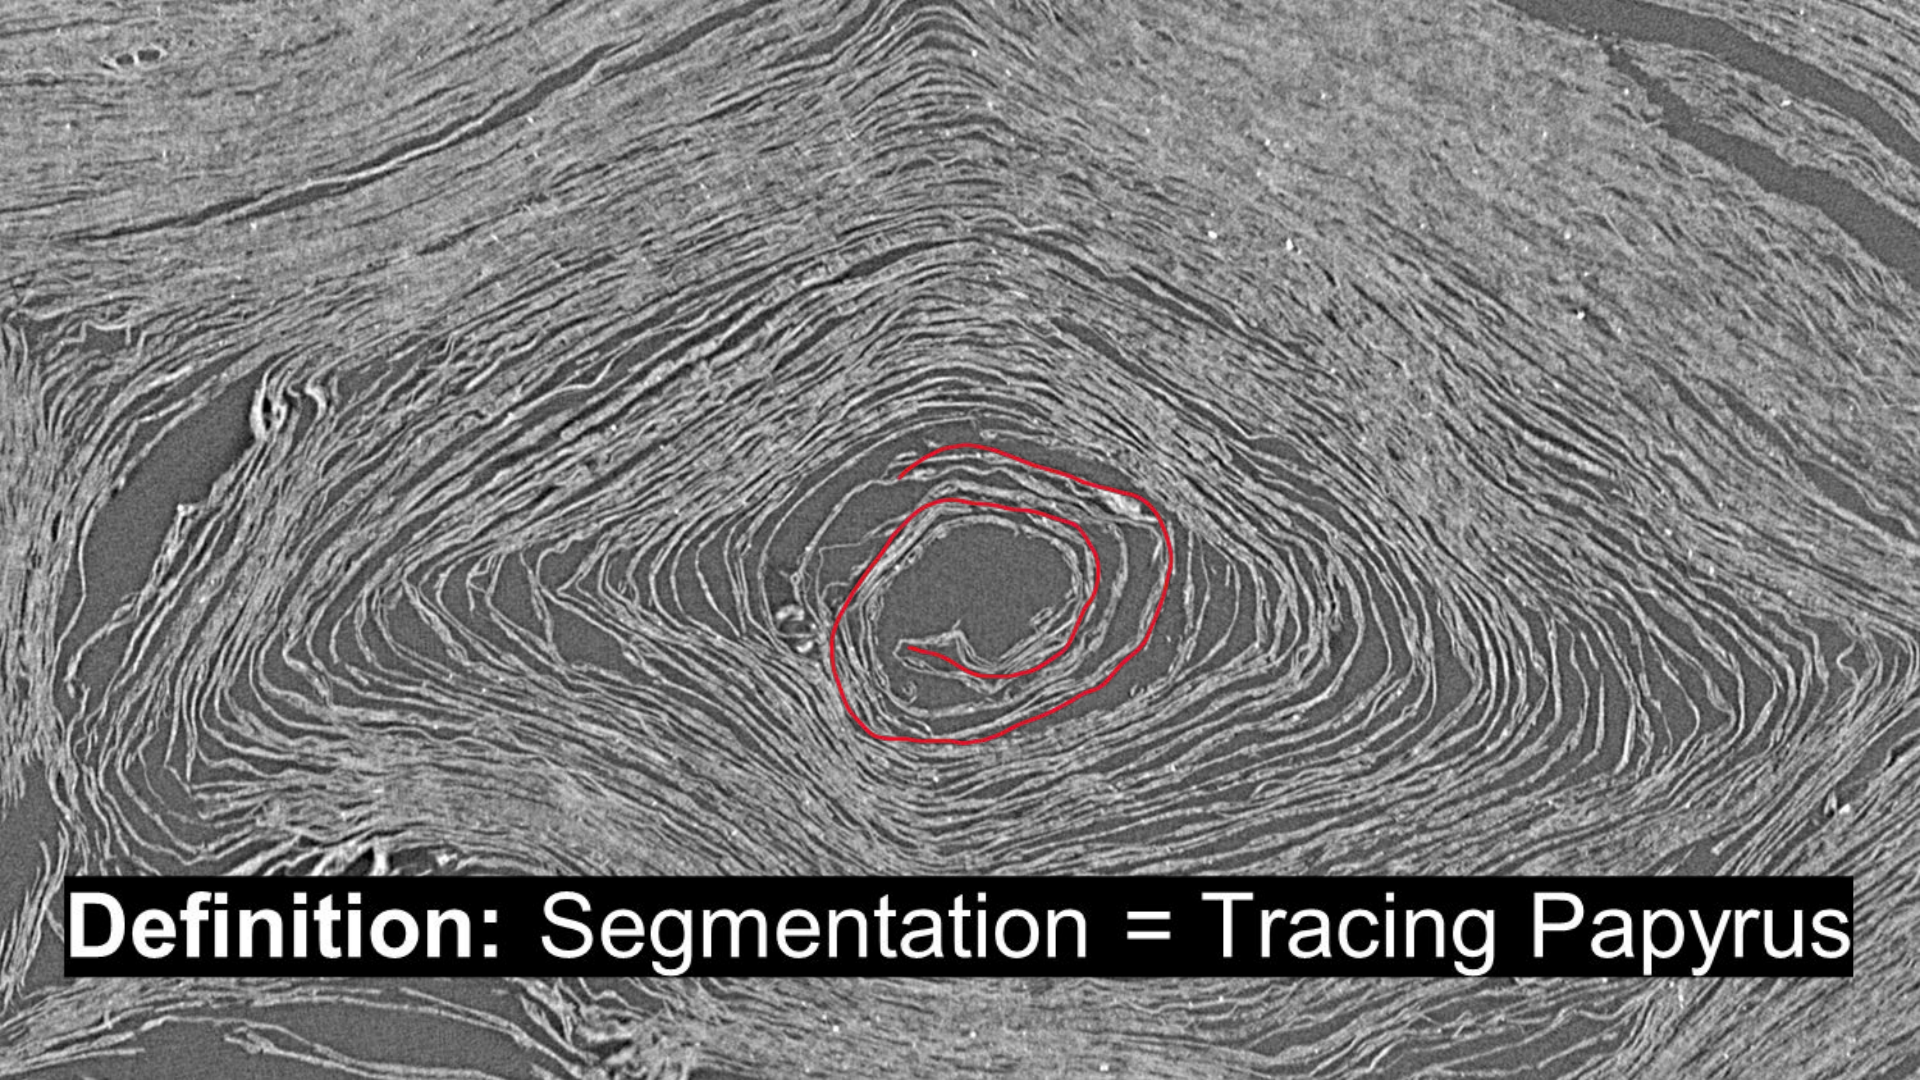
Definition: Segmentation = Tracing Papyrus

Rendering

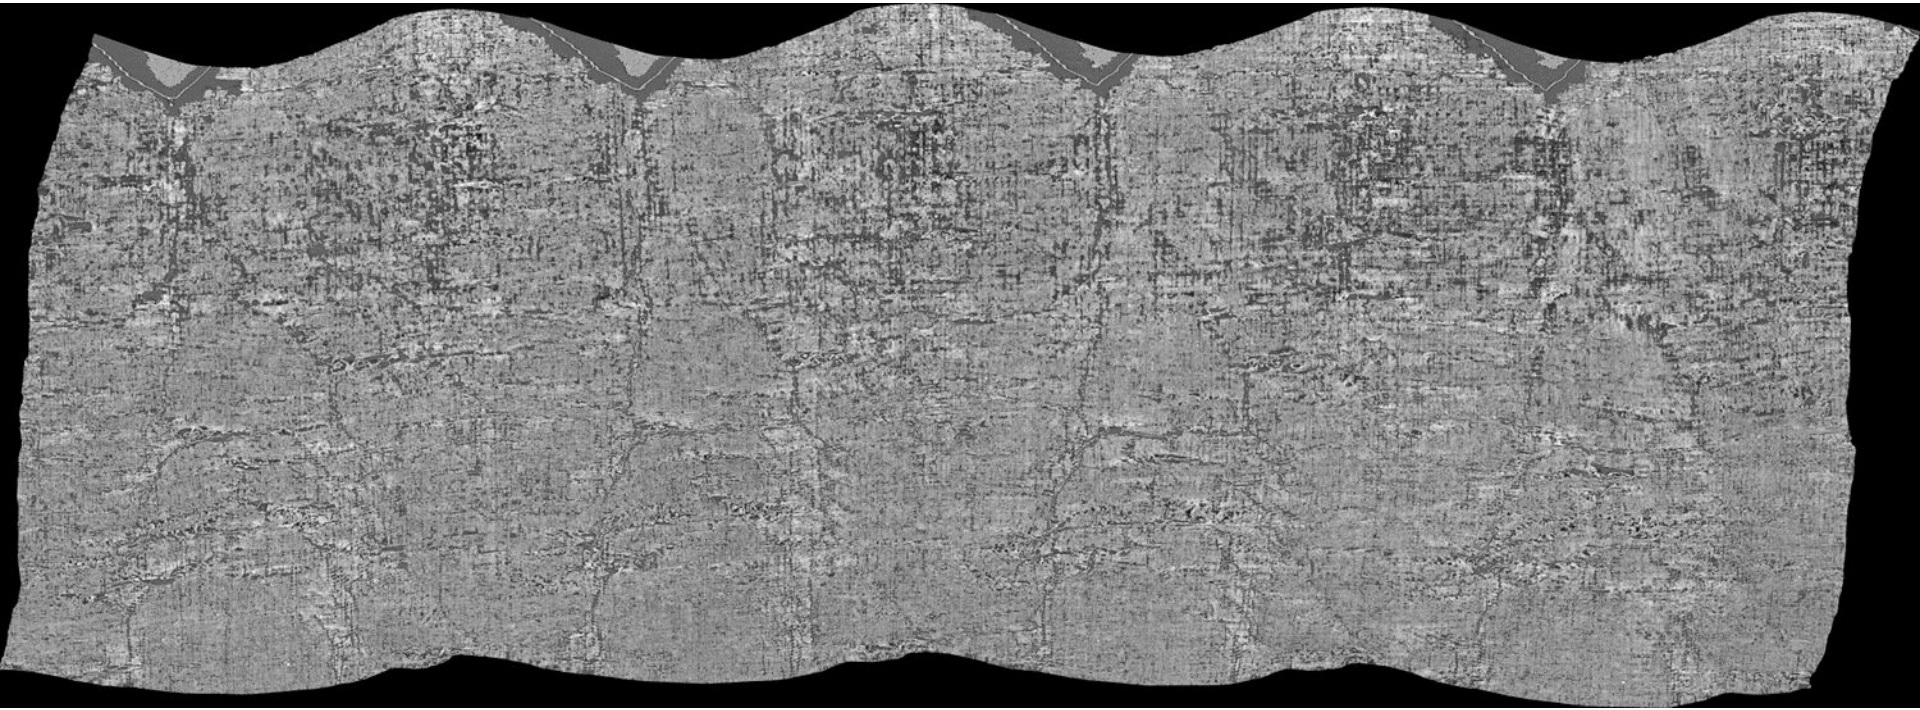

Previous Tool

Volume Cartographer (VC)

- Clunky
- Slow
- Crashed often
- Extremely small segments

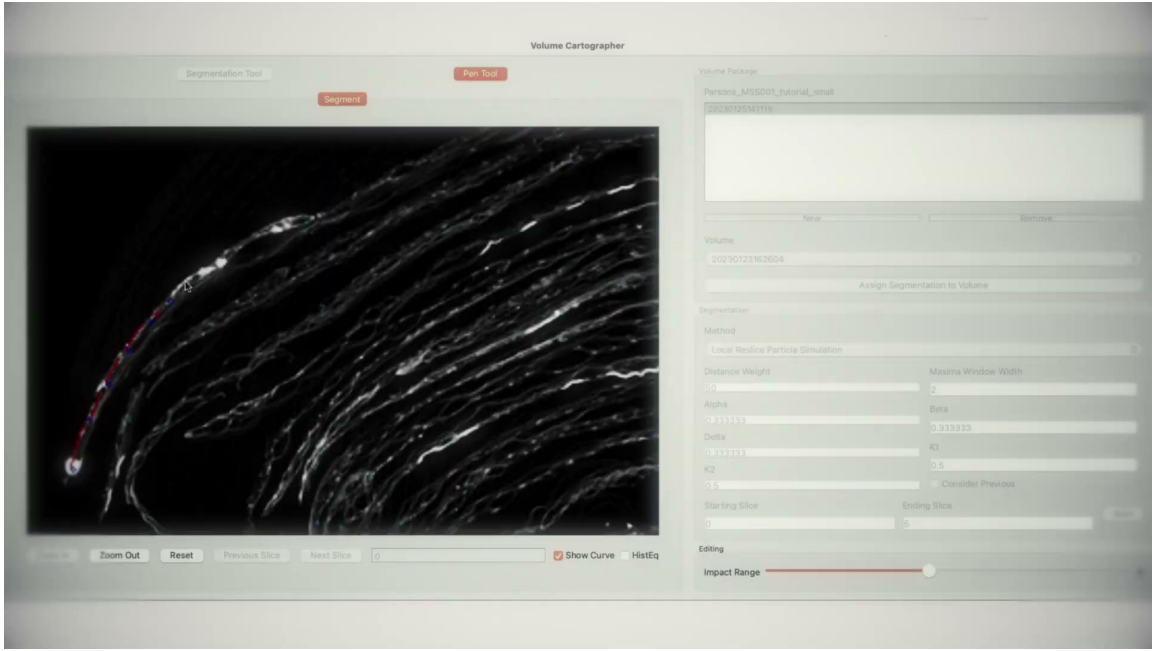

Optical Flow Segmentation (OFS)

- Video = 3D Scroll Volume
- Movement of pixels = sheet
- “Magic” Sauce
- Testing, testing, testing …

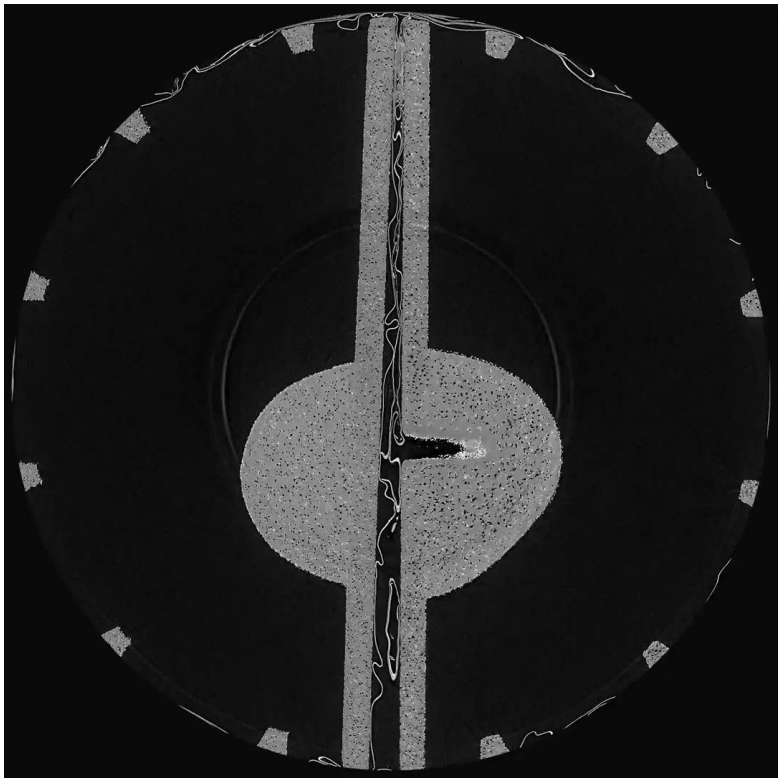

PHerc.Paris. 4

Reason for VC+OFS Success

- Interesting Problem
- VC was used before
- Developed in the open
- Feedback from users
- Fast iteration cycles
- Area Segmented = Performance

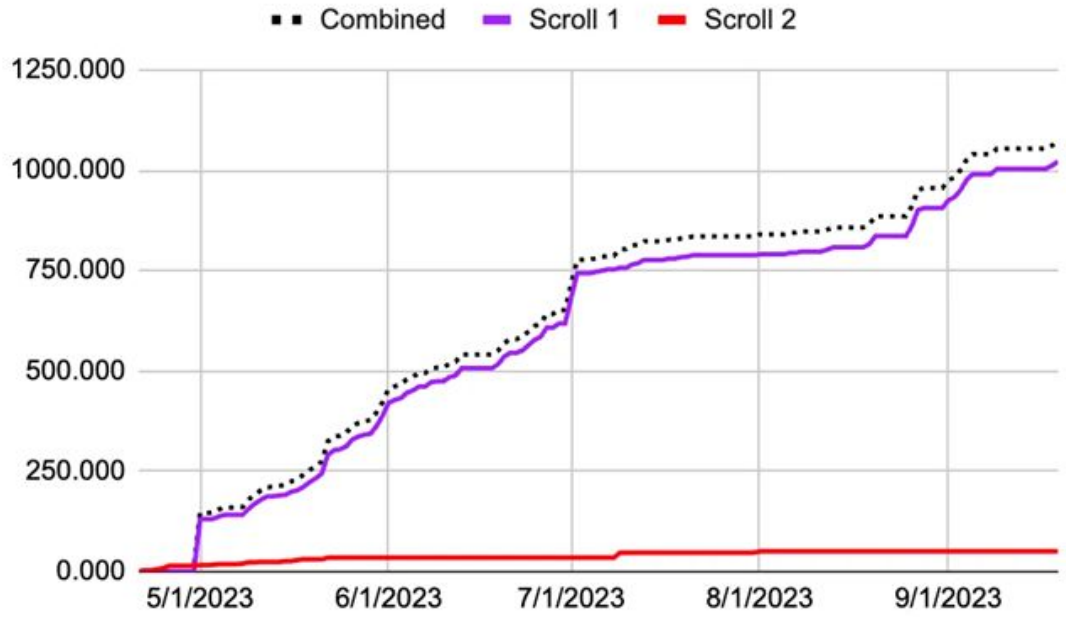

Result: 10’000x Improvement

Segmentation

- Tracing the papyrus cost ~$100/ sqcm in 2023
- $2-3M per scroll
- We have 400-800 scrolls
- Estimated cost: 3 Colosseums

Quora
https://www.quora.com › Where-did-the-money-come-...

Where did the money come from to finance the building ...

Good question. It has been estimated that the building took about **$670 million** to build. Take that estimate with a grain of salt, ...

ThaumatoAnakalyptor

- Goal: Segmentation at scale
- Define smaller subproblems
- Focus on
  - **Generalization**
  - **Abstraction**
  - **Automation**

www. github.com/schillij95/ThaumatoAnakalyptor

Pipeline Steps

- Extract 3D Surface Points
- Point Connectivity
- Meshing + Flattening

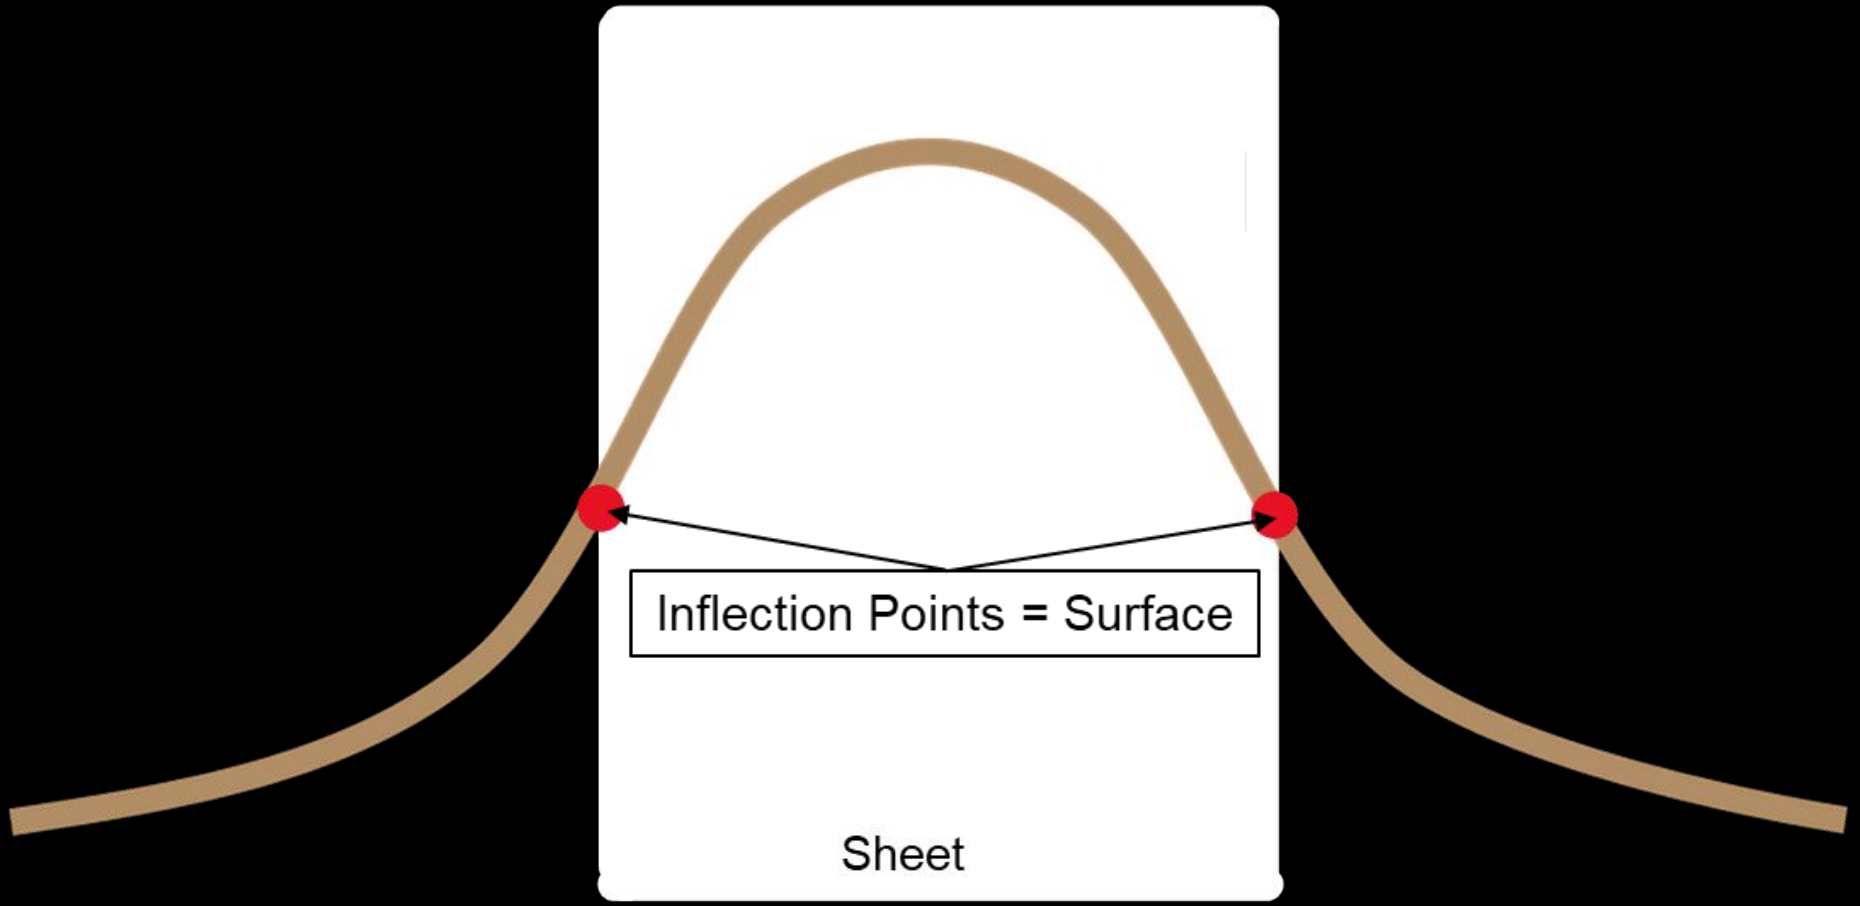
Inflection Points = Surface
Sheet

Extract Surface Points - PointCloud

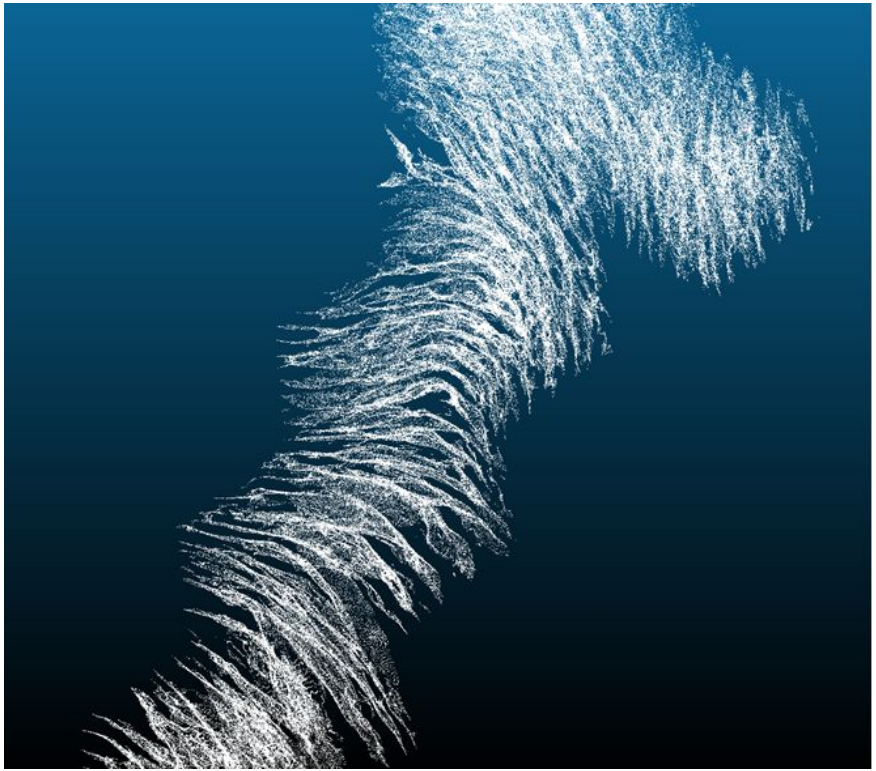

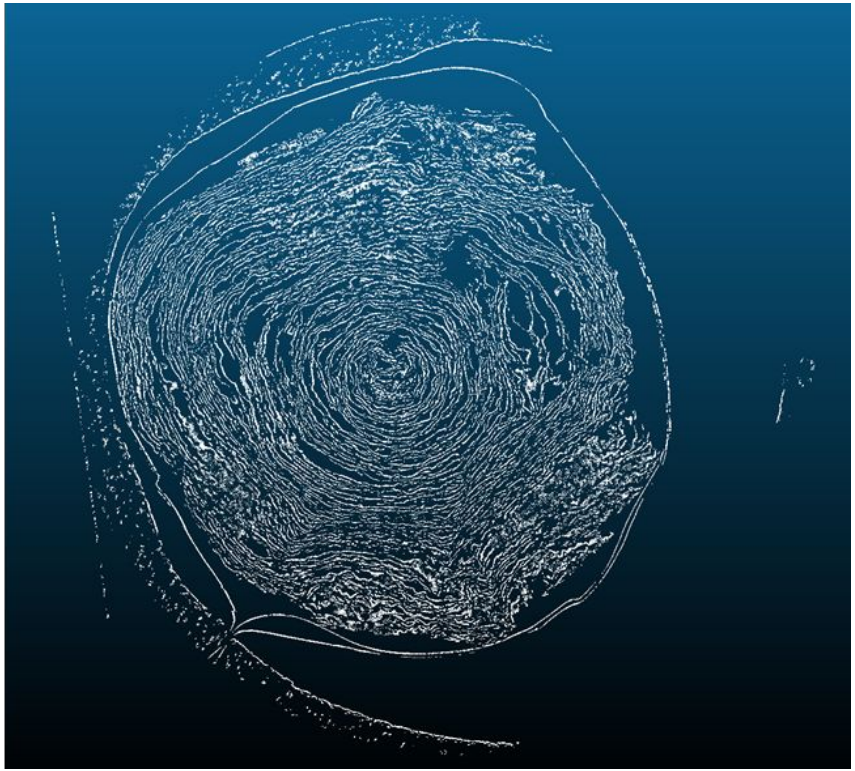

Extract Surface Points - Sheets

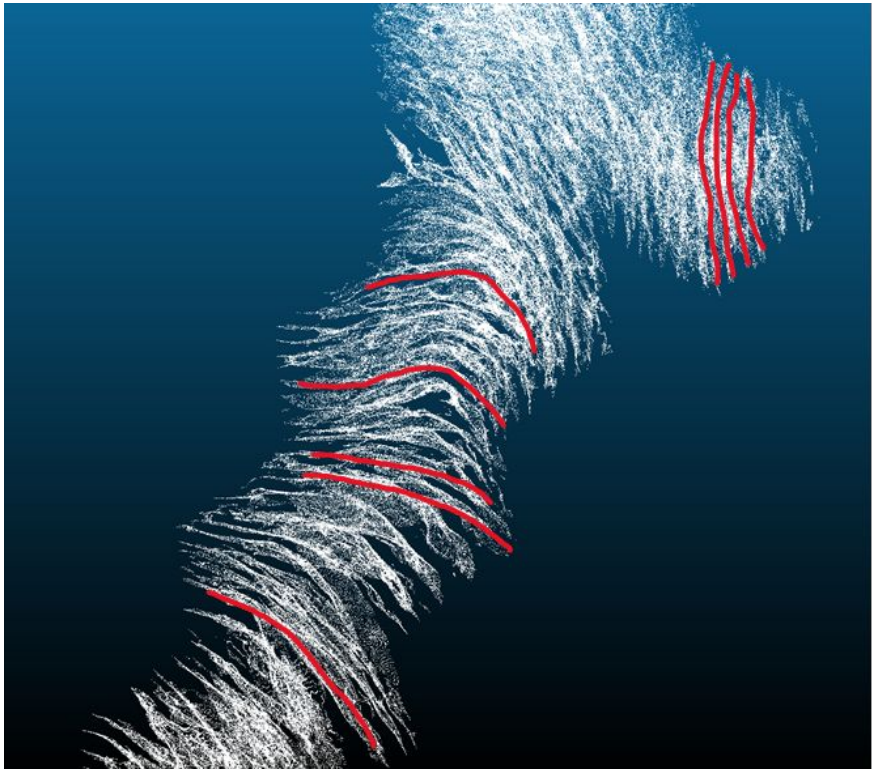

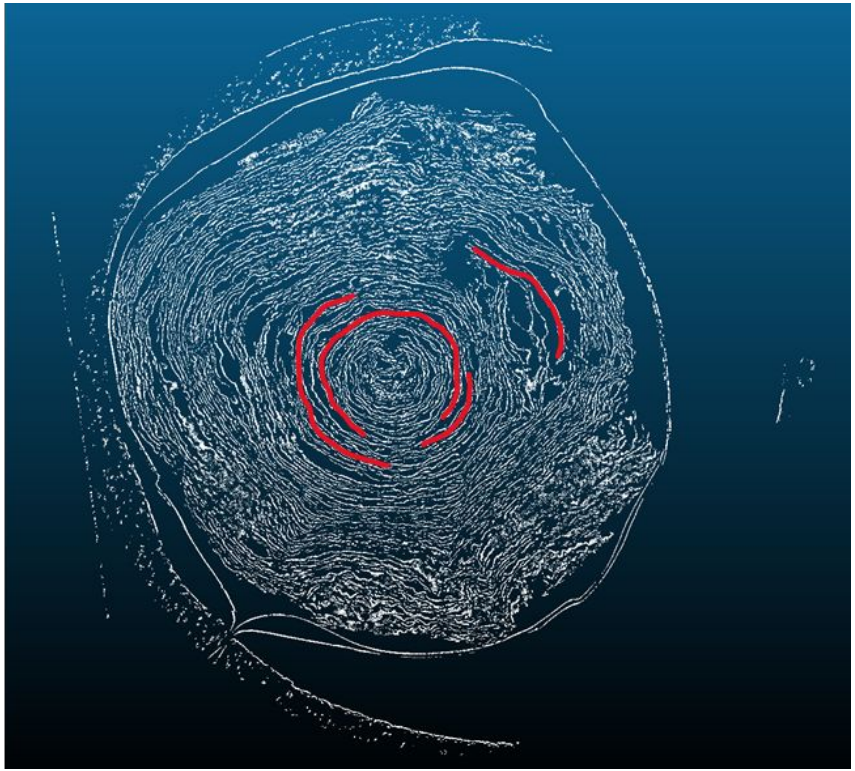

Extract Surface Points - Blocks

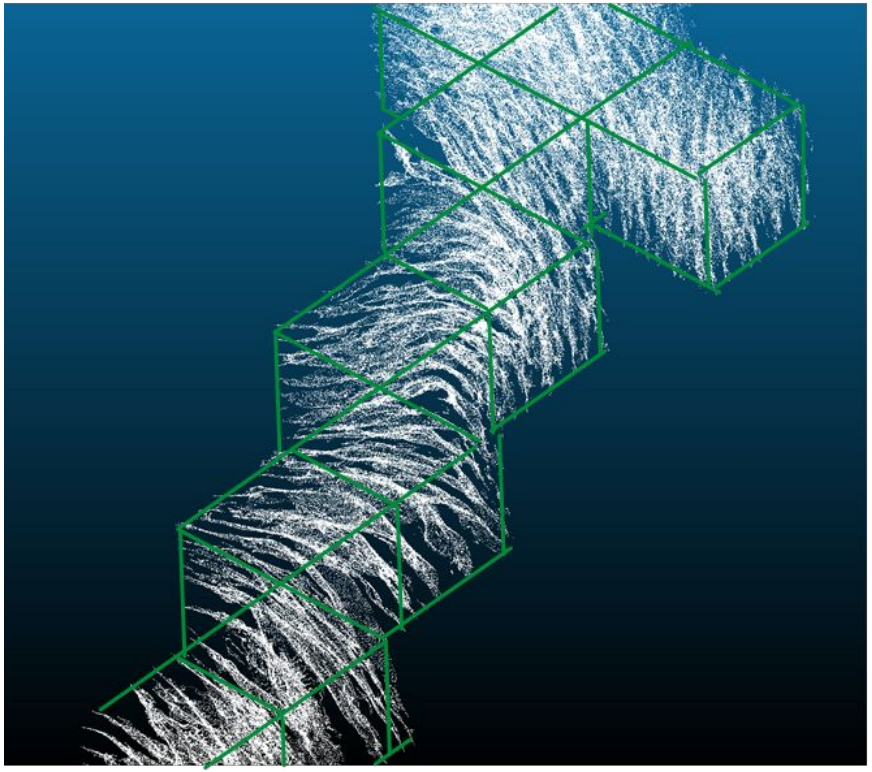

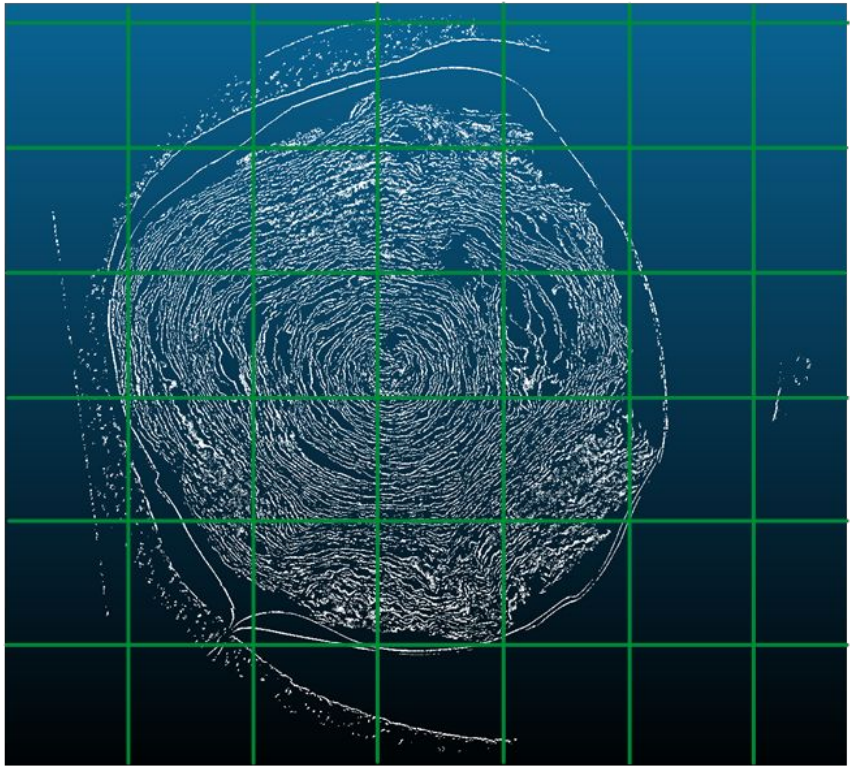

PointCloud Connectivity

- First: PointCloud Instance Segmentation
- Second: Stitch Instances to Sheet
- Minimum Feedback Arc Set Problem
- General NP hard
- Our case???
- Randomized Algorithms

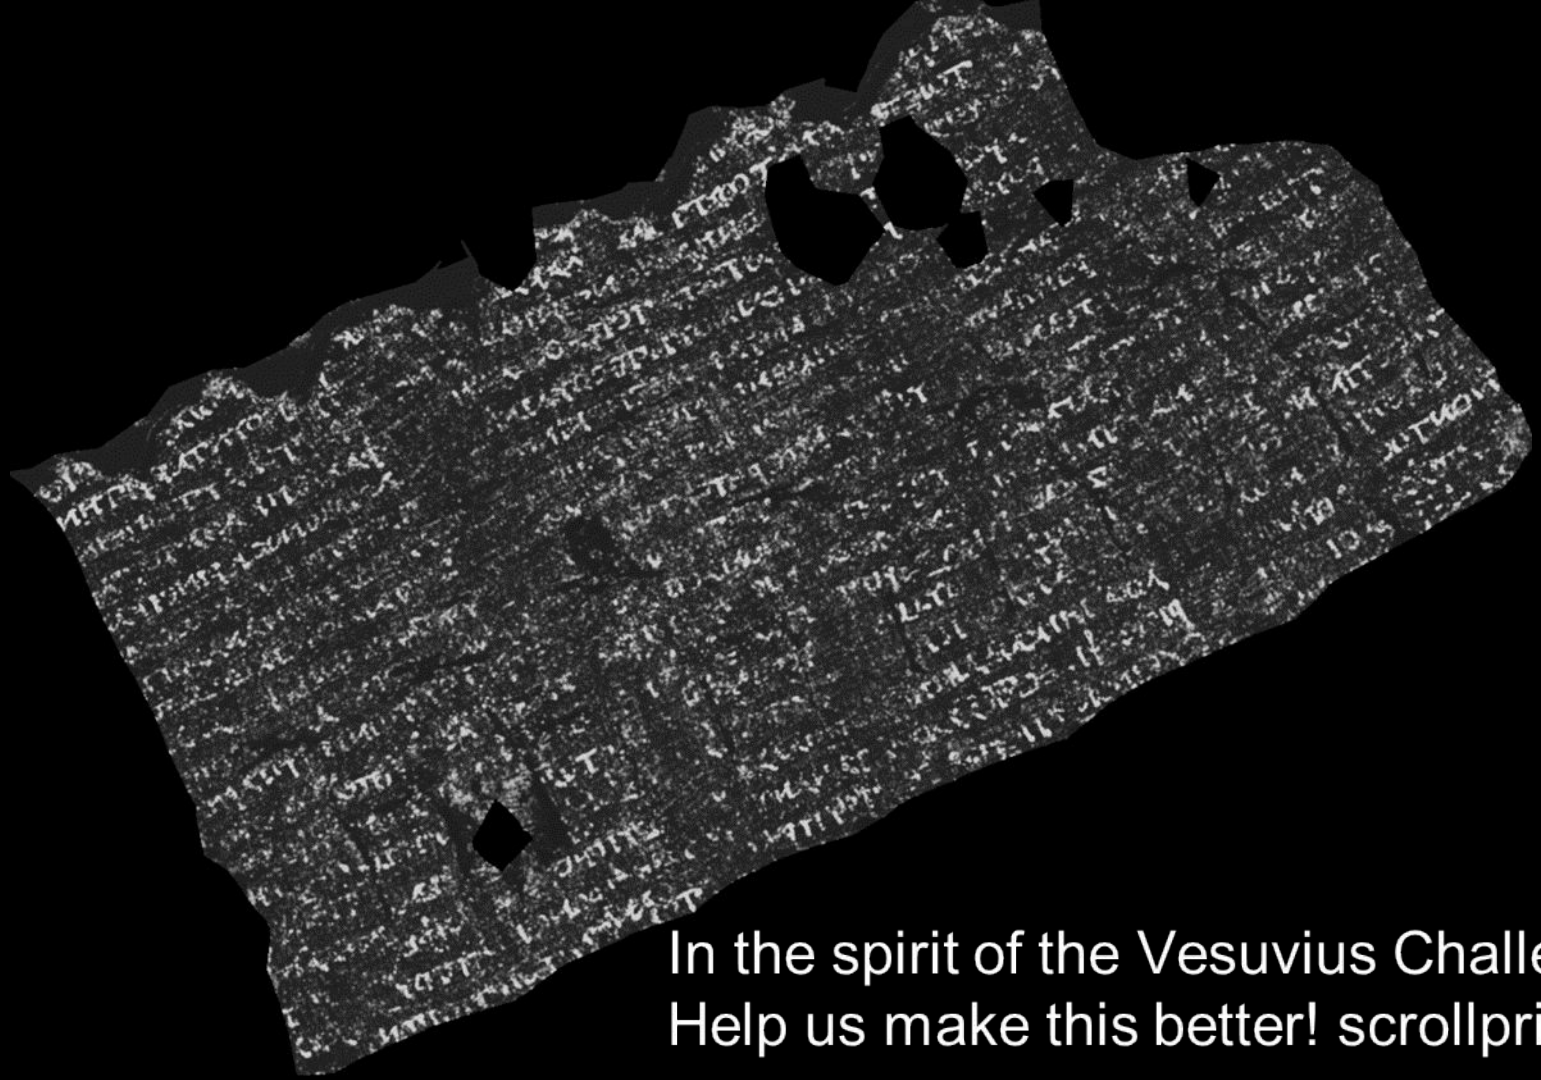

In the spirit of the Vesuvius Challenge:
Help us make this better! scrollprize.org

Introduction

Youssef Nader

Bsc Communication and Computer Engineering Cairo University

Msc Data Science Freie Universitaet

Currently PhD student at Prof. Tim Landgraf's Lab

Work Interests: Machine learning|XAI and interpretability|Self-supervised Learning

How I got involved?

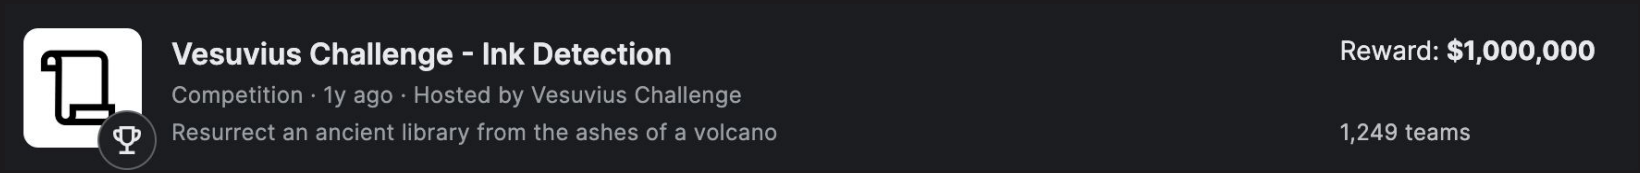

Ink Detection in 3d Images

The flattened and unwrapped segment is turned into a stack of 2D images where ink detection is possible

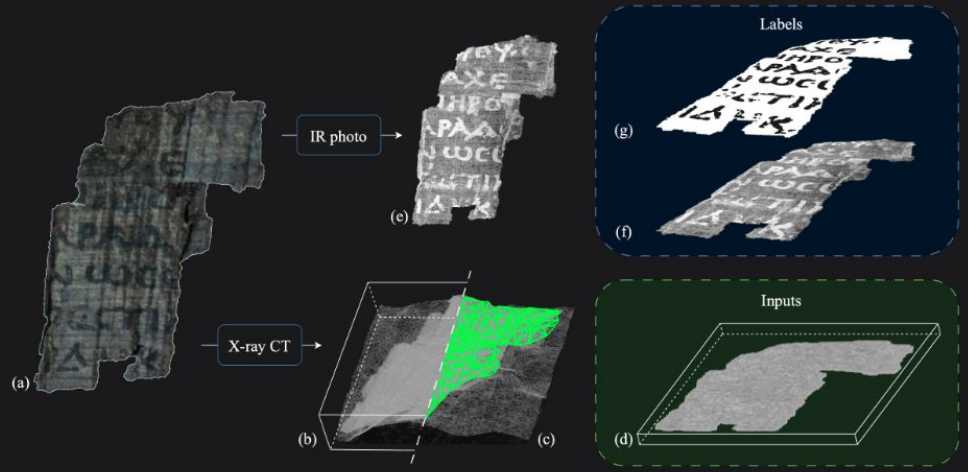

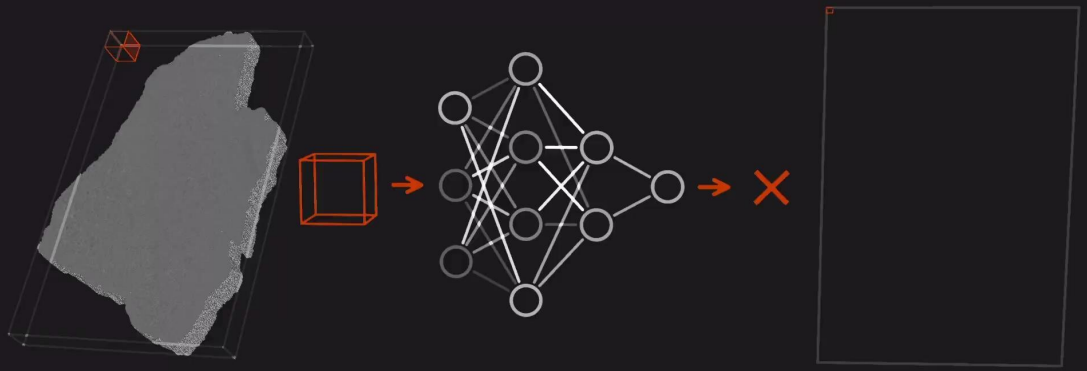

Source: https://scrollprize.org/tutorial4 © Vesuvius Challenge

Kaggle ink detection Progress

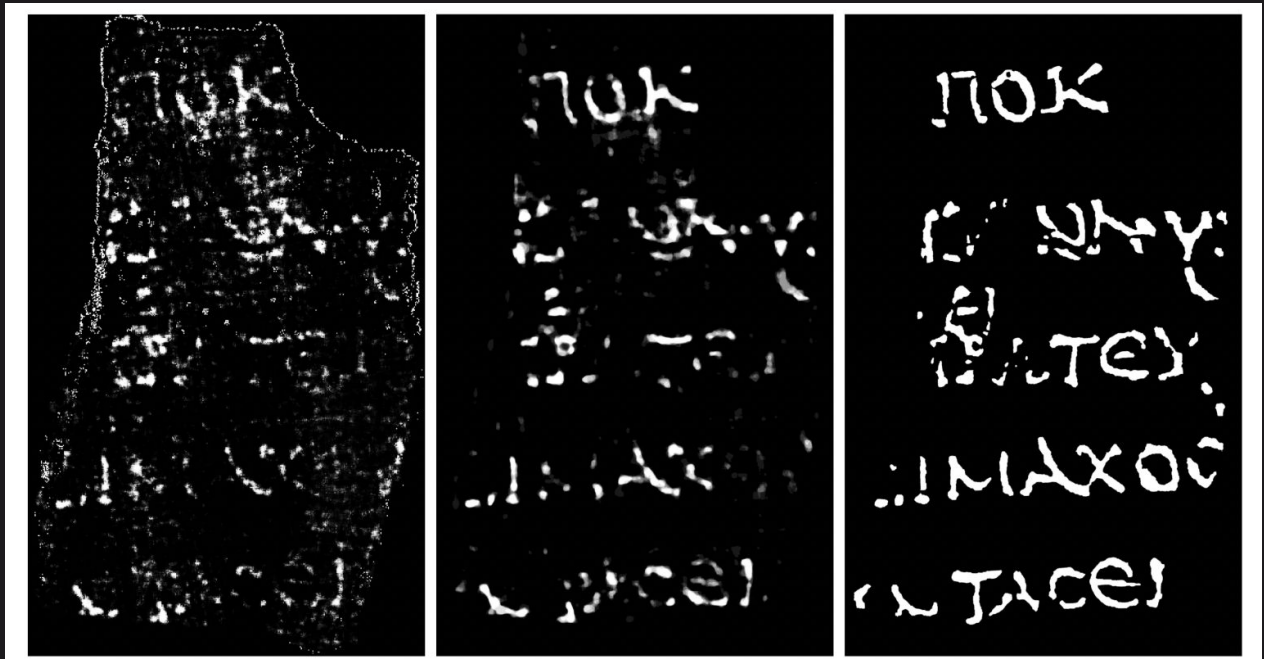

(a) ink-ID baseline. (b) Top Kaggle results. (c) Ink label.

Hard-Hearted Scrolls: A Noninvasive Method for Reading the Herculaneum Papyri
https://uknowledge.uky.edu/cs_etds/138/

Ink Detection Follow Up Solution

X

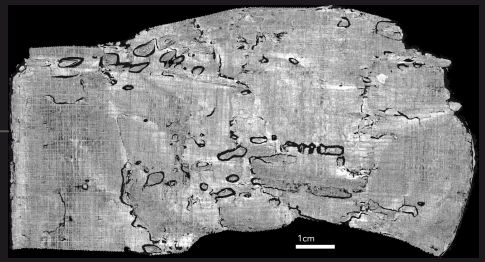

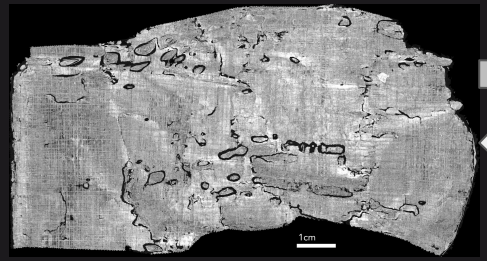

?

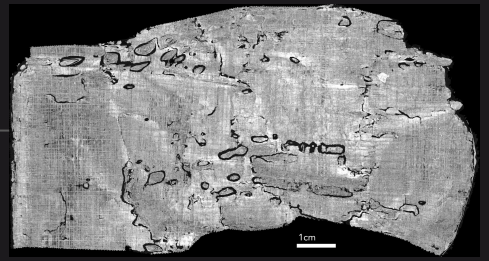

Self-Supervised Learning on Scrolls

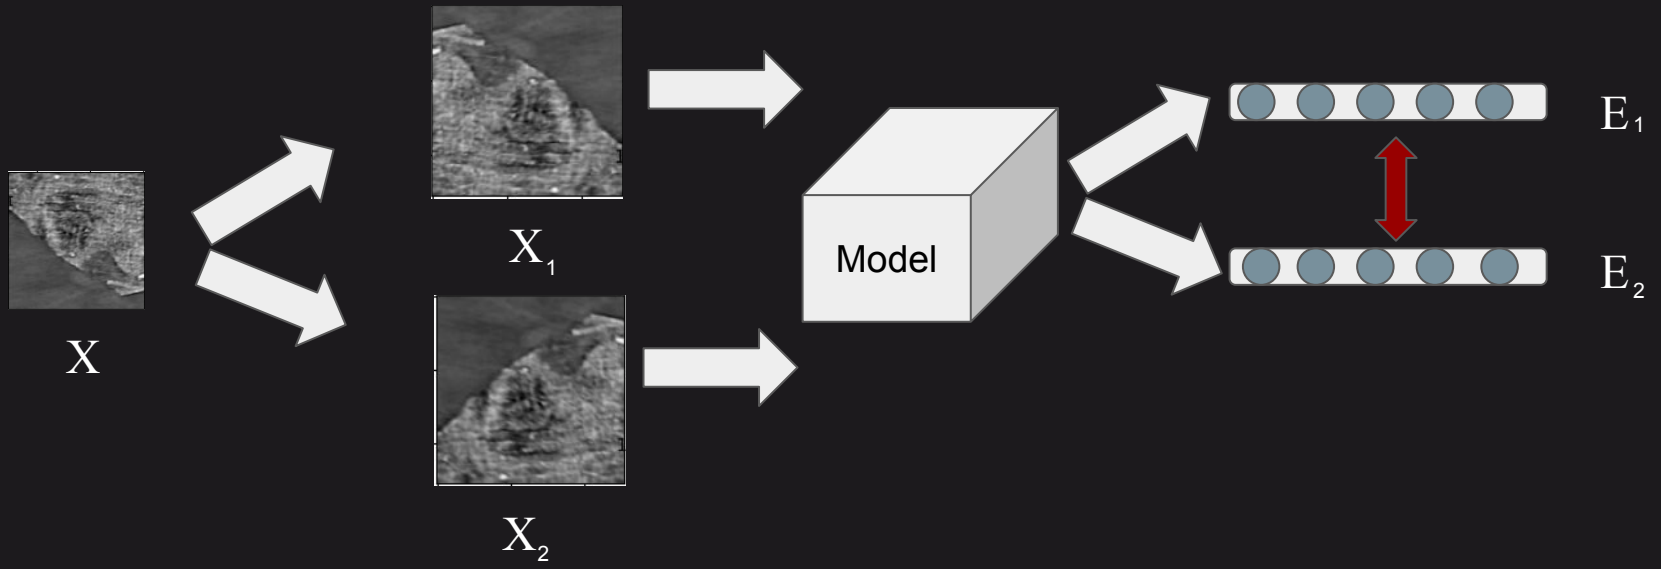

Pseudo-labeling Process

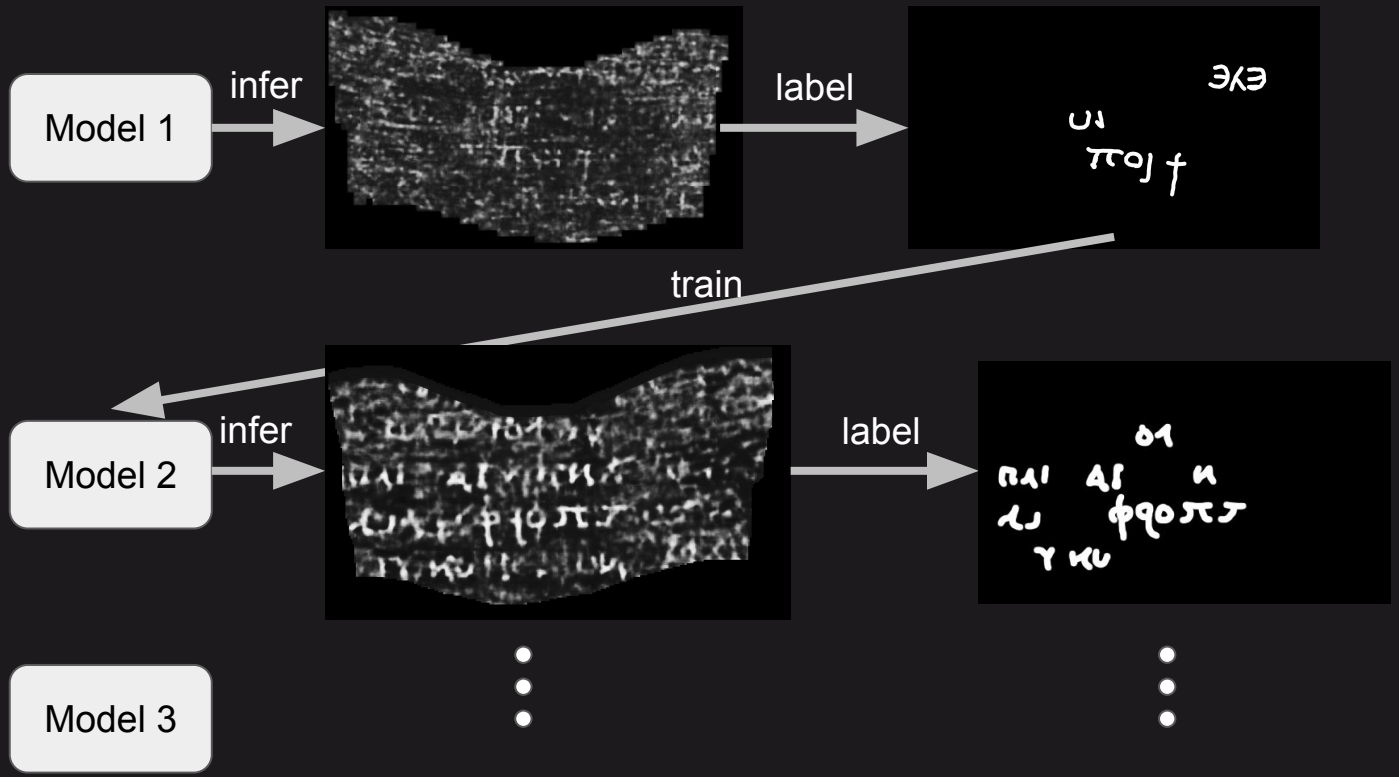

How to prevent overfitting?

1- Smaller models

2- High Label Smoothing

3- Balancing the dataset

4- heavy augmentations

Give the model some wiggle room to make mistakes

High Label Smoothing

Mistakes can happen during labeling (false positives/negatives)

Models should not be heavily penalized when they make a mistake

Label smoothing: turn 1-0 probabilities into softer probabilities eg. 0.875-0.125

Balancing the dataset

Sample windows around the letter

Balance the dataset with negative samples

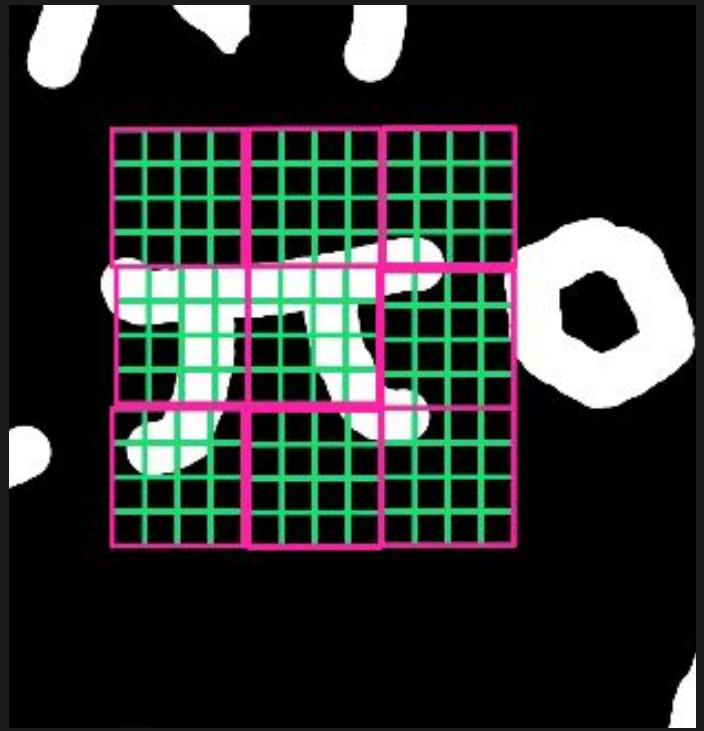

Choosing The Window Size

Larger window sizes have the full context of the letter

The model learns and is able to match the label stroke

Smaller windows mitigate hallucination and preserve stroke shape

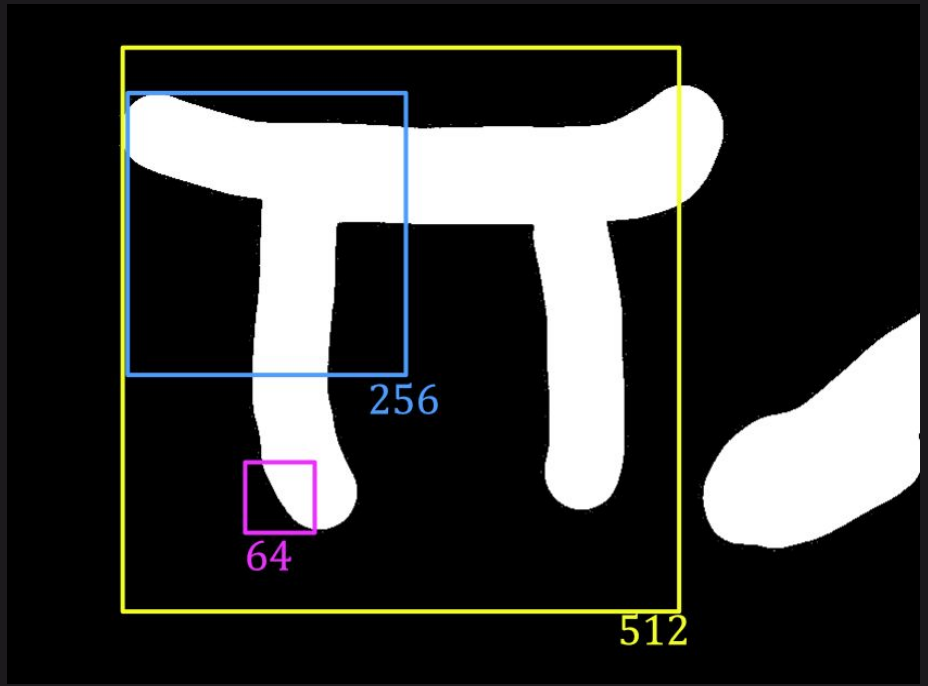

First Letters Prize

Scaling up to the Grand Prize

- The Grand Prize solution was about perfecting the First Letters approach
- Explores two main fronts:

  1- How to get as much quality data as possible?

  2- How to train models to make the most out of the available data?

Scaling up the Data

- More than 15 pseudo-labeling rounds between first letters and Grand Prize
- Exploring labeling techniques, data quality effects
- Conclusion: Quality and Diversity >>> Quantity

Scaling the models, Why timesformer?

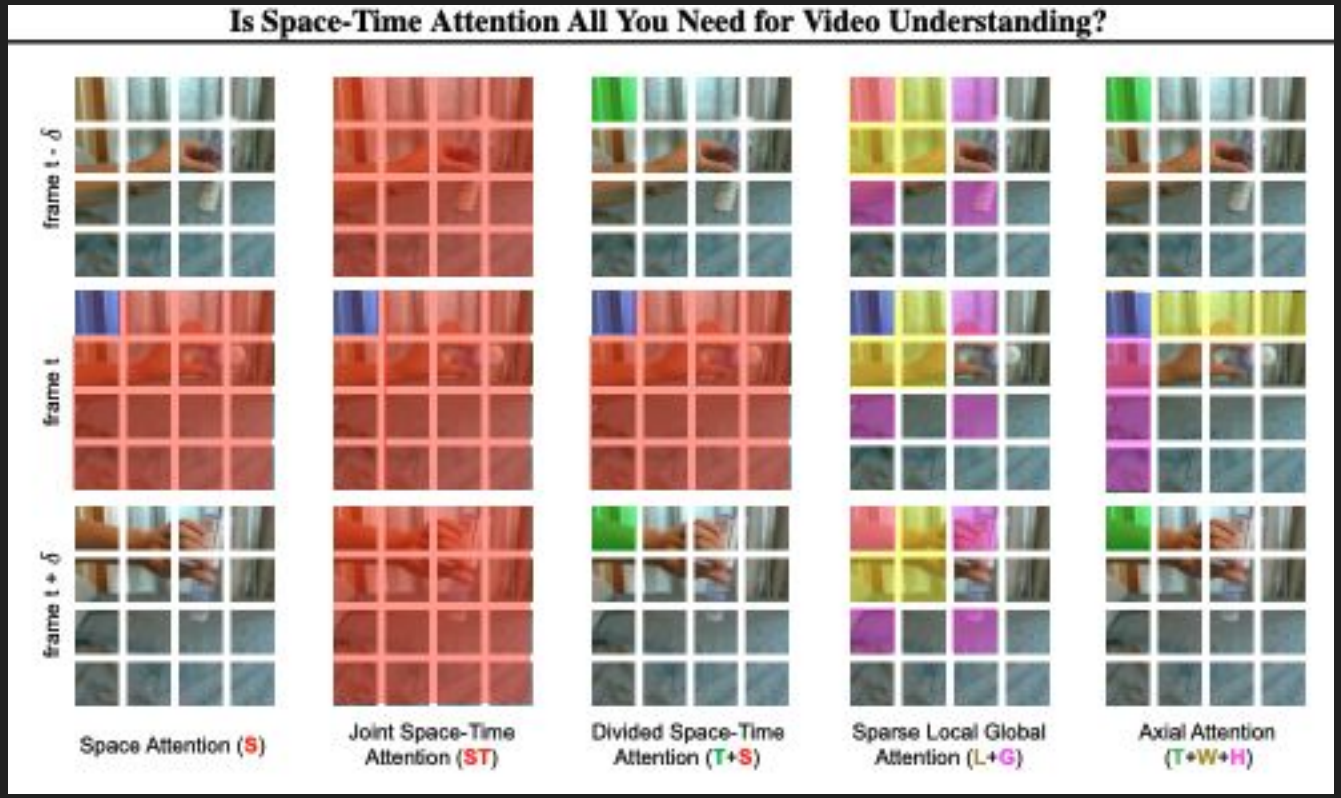

Source: Is Space-Time Attention All You Need For Video Understanding?
https://arxiv.org/pdf/2102.05095.pdf

“…have nothing to say about pleasure, either in general or in particular, when it is a question of definition.”
-Philodemus

Thank You!