Improvements in Fault Prediction Capabilities for Radiotherapy Machines

Sam Leadley
University of Oxford, JAI & CERN

Supervisors
Manjit Dosanjh
University of Oxford, JAI & CERN
Raj Jena
University of Cambridge

John Adams Institute for Accelerator Science

STELLA
SMART TECHNOLOGY TO EXTEND LIVES WITH LINEAR ACCELERATORS

UNIVERSITY OF OXFORD

UNIVERSITY OF CAMBRIDGE

CERN

Background

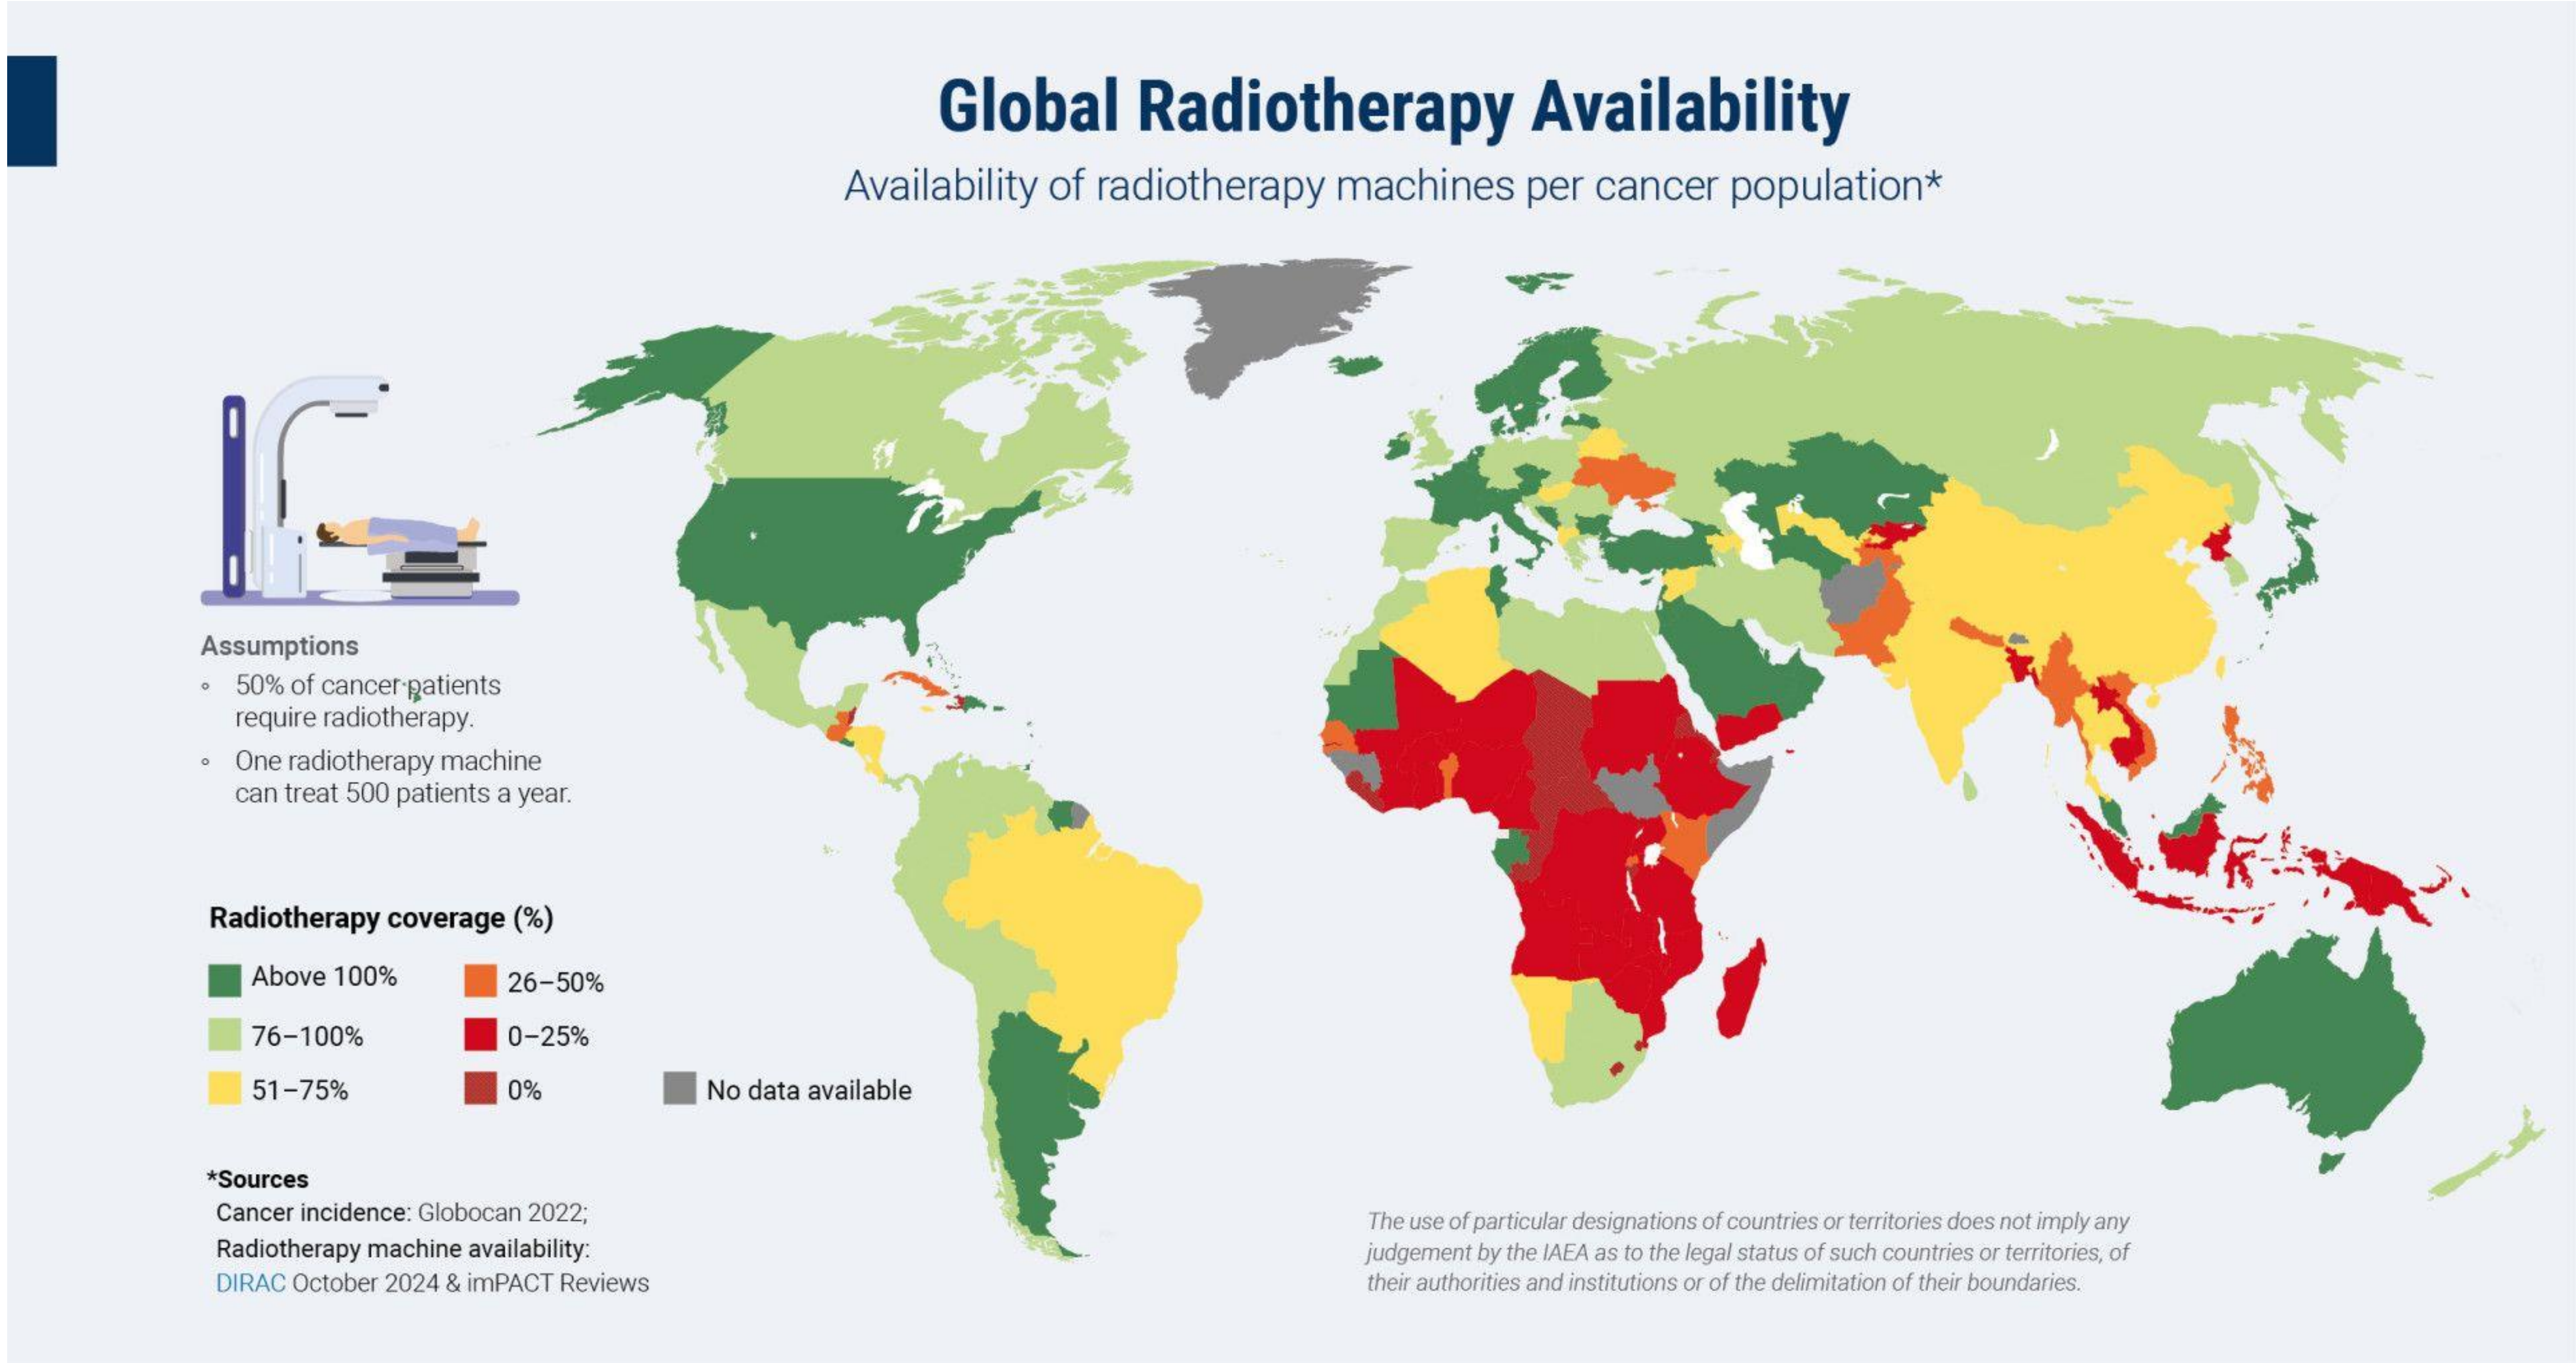

M. Dosanjh, 2025. “Towards Implementing Project STELLA on radiotherapy powered by AI”, GHC Summit. MIT, Boston, 19/09/25.

Background

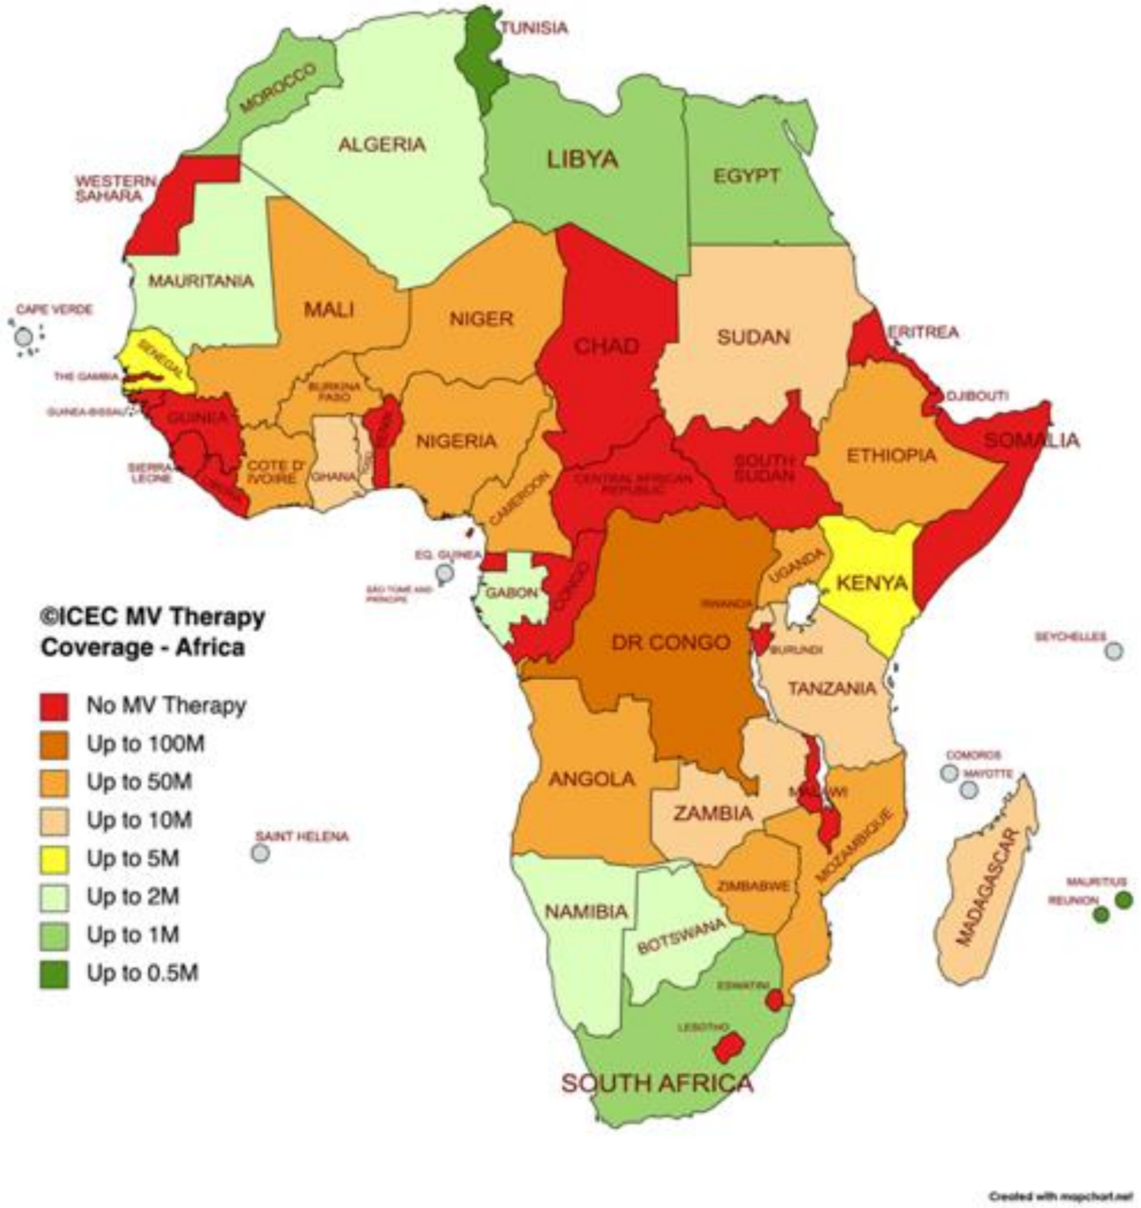

Number of linacs per country in Africa. *Credit: ICEC*

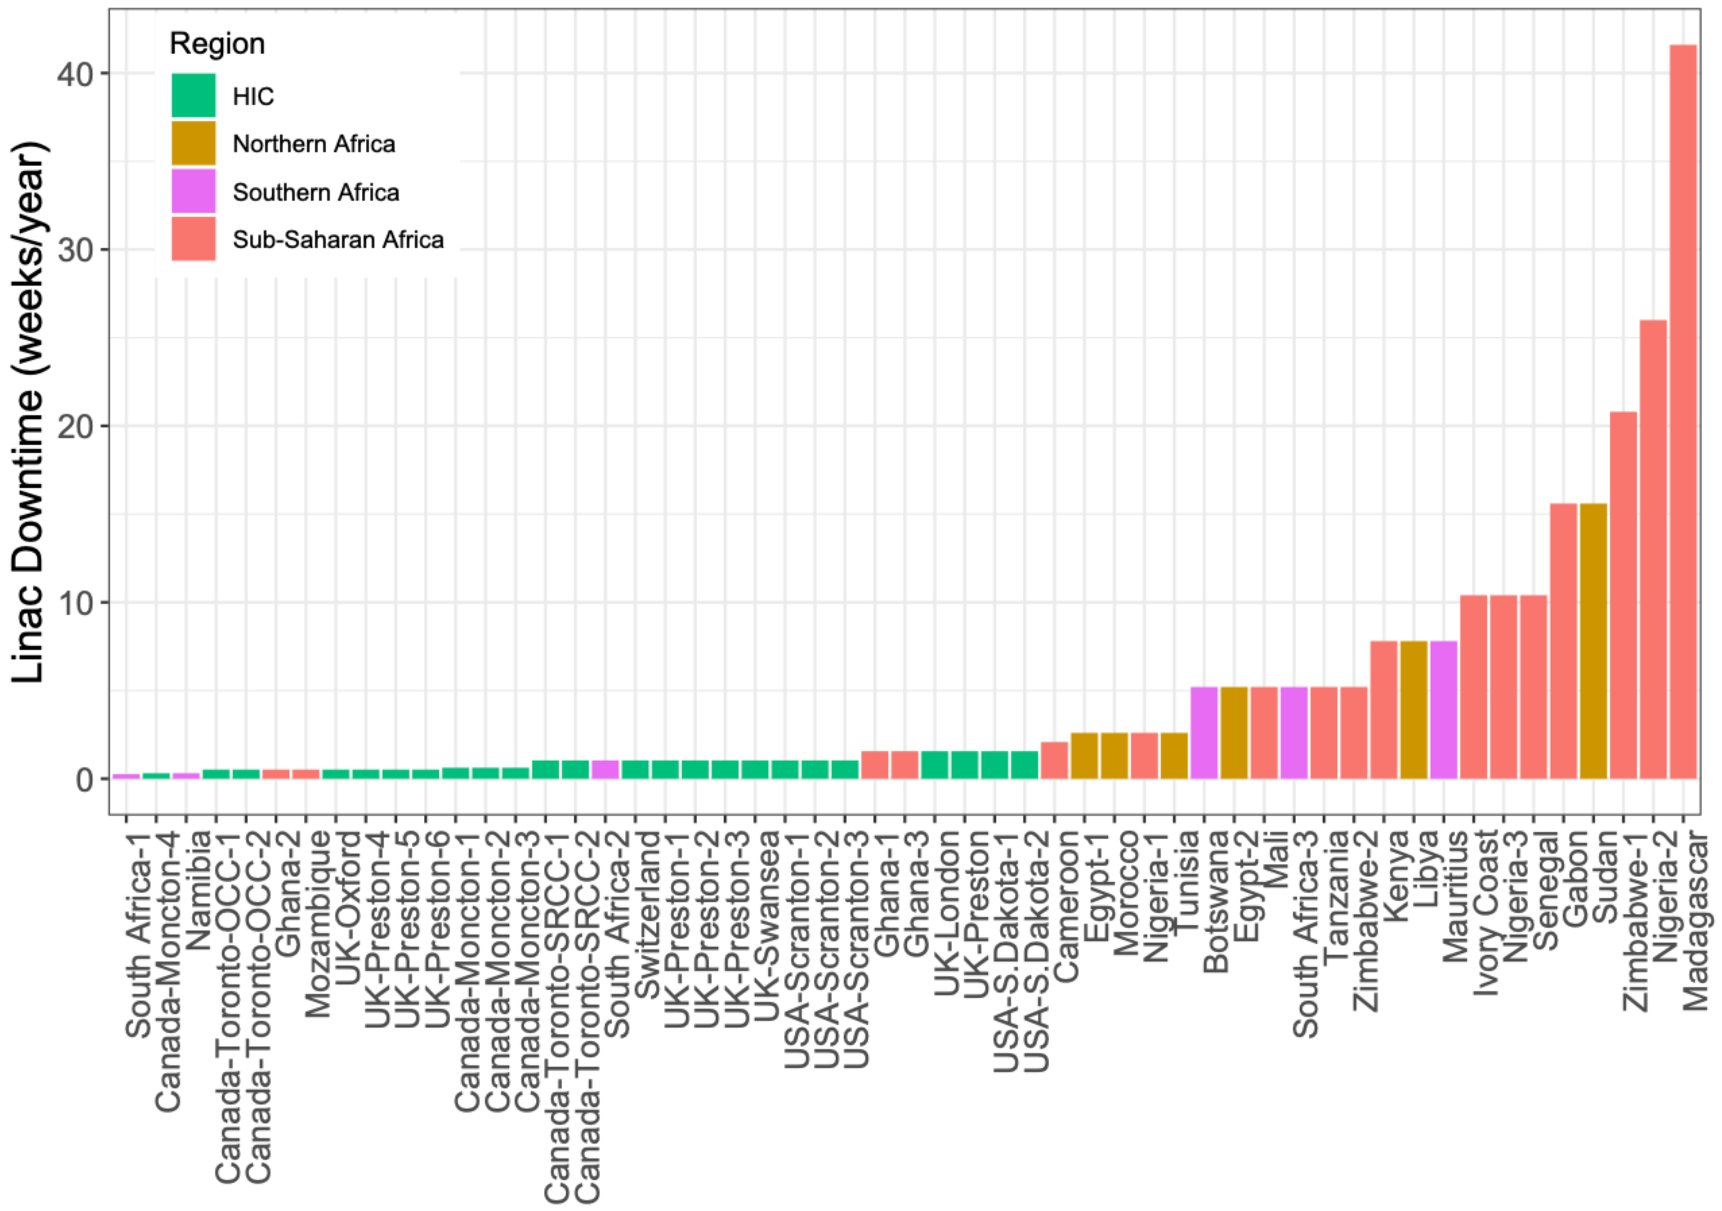

Machine downtime in countries in Africa compared with that in Canada, UK and US.

Project STELLA

- Established 2016 with the goal of increasing global radiotherapy access.

- Global multidisciplinary NGO collaboration under the umbrella of ICEC.
  - Oxford, Cambridge, Lancaster, CERN, African partners

- Increase access to radiotherapy in LMICs and rural areas.
  - New LINAC design, more robust, cheaper, same treatment capabilities
  - Improved training for RT engineers to operate and troubleshoot RT LINACs
  - Reduce downtime in existing RT centres using advanced computing techniques

Motivation

- Access to RT in LMICs (such as in Africa) is highly limited by available machines.
- Of the few available machines, many experience significant downtime, far more than in HICs.
- Machines are pushed harder to compensate for lower availability.
- Ability to predict machine downtime in advance allows for repair preparation & preventative maintenance.

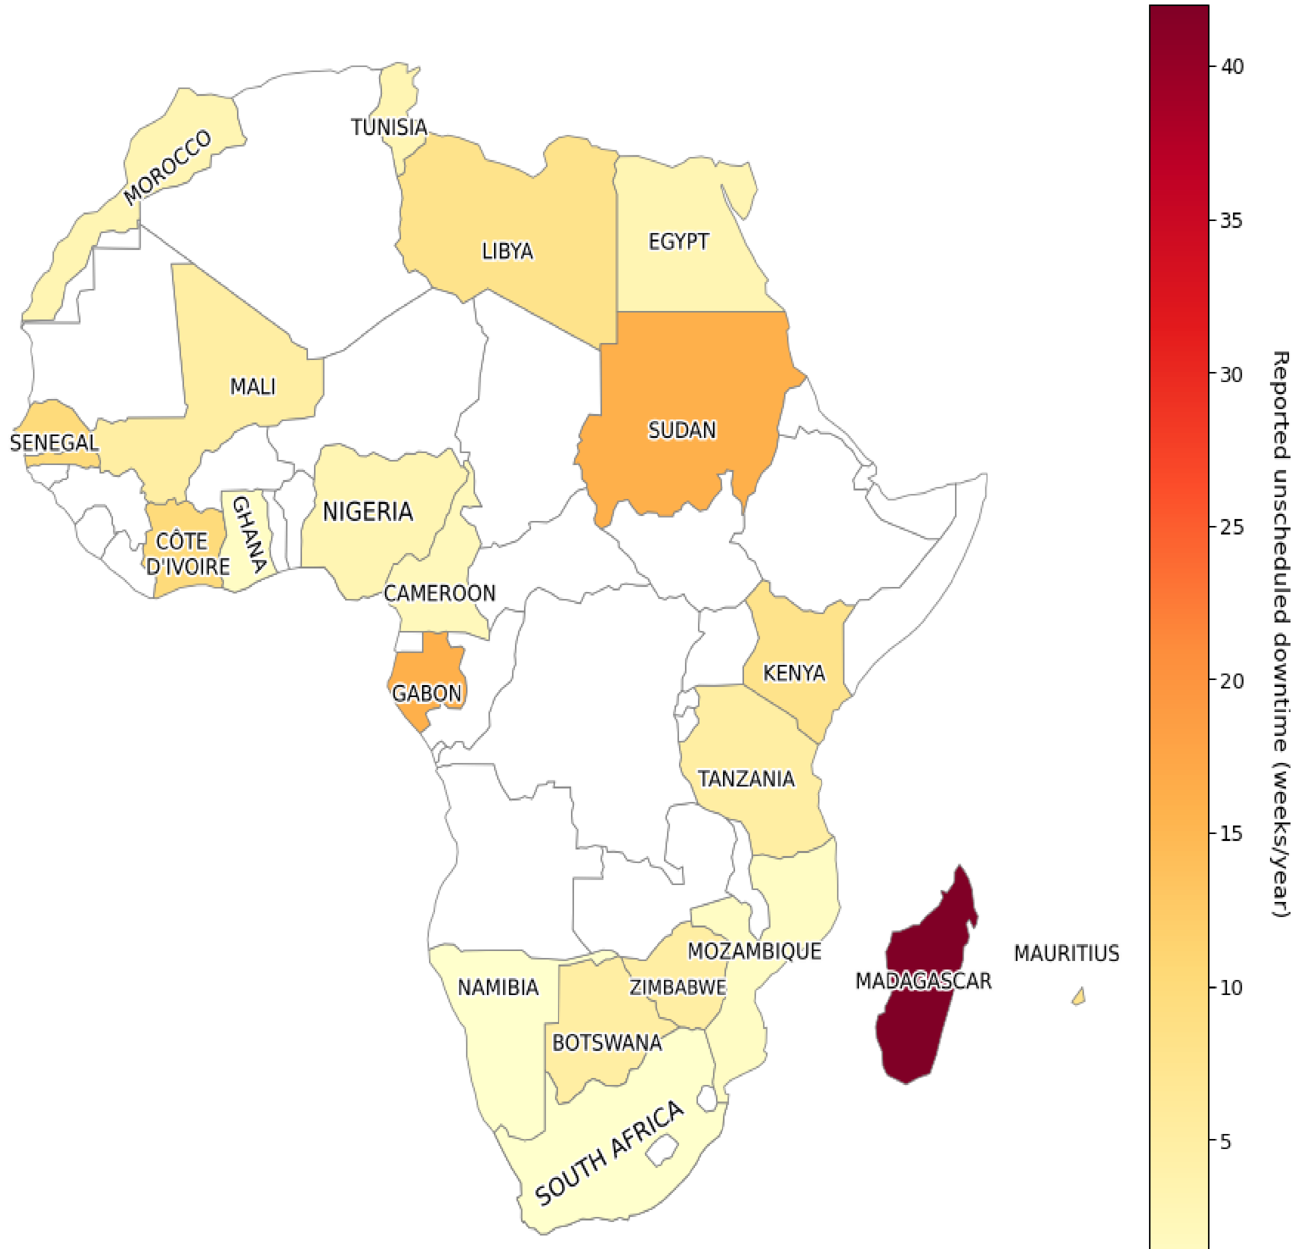

T.A. Ige *et. al.* (2021). Surveying the Challenges to Improve Linear Accelerator-based Radiation Therapy in Africa

RT Machine Log Files

- Automatically recorded onto RT LINACS (Varian TrueBeam), no patient identifiable data.
- Contains information on the exact state of the machine at any point in time.
- “Interlocks” – reduction in the state of treatment readiness for the RT LINAC
- Want to predict when specific interlocks are likely to occur.

Time & date of interlock code

Interlock code

```
2024-03-2707:49:15Local0Info07:49:15:679SN# 3936COLFaultCMNFault::raise Fault detected (420302: L B MLC Fiber Op
2024-03-2707:49:15Local0Info07:49:15:746SN# 3936BGMFaultCMNFault::raise Fault detected (214261  L B POS PCB: Cu
2024-03-2707:49:15Local0Info07:49:15:746SN# 3936BGMFaultCMNFault::raise Fault detected (214262  L B POS PCB: Cu
2024-03-2707:49:15Local0Info07:49:15:746SN# 3936BGMFaultCMNFault::raise Fault detected (214268  L B POS PCB: 28
2024-03-2707:49:15Local0Info07:49:15:709SN# 3936COLFaultCMNFault::raise Fault detected (420400: L B MLC Bank A L
2024-03-2707:49:15Local0Info07:49:15:719SN# 3936COLFaultCMNFault::raise Fault detected (420500: L B MLC Bank A C
2024-03-2707:49:15Local0Info07:49:15:730SN# 3936COLFaultCMNFault::raise Fault detected (420600: L B MLC Bank B L
```

Example of processed Varian Combined Log file. All interlocks have been extracted but are yet to be filtered.

“Severe” Interlocks & Machine Learning

- “Severe” interlocks are defined as interlocks causing *machine downtime* and reducing patient throughput capability (increasing *clinical downtime*).

- A list of these interlocks has been collated by radiotherapy engineers at Exeter Hospital based on past faults that have caused machine downtime.

- Prediction of severe interlocks is key to allowing preventative maintenance and preparation
  - With so much data, this is a high-dimensional pattern recognition problem that a human would struggle to solve.
  - Machine learning algorithms can recognise long-term relationships and contexts in data and reduce the dimensionality of the problem for interpretation.

Dual-Classifier Variational Autoencoder Approach

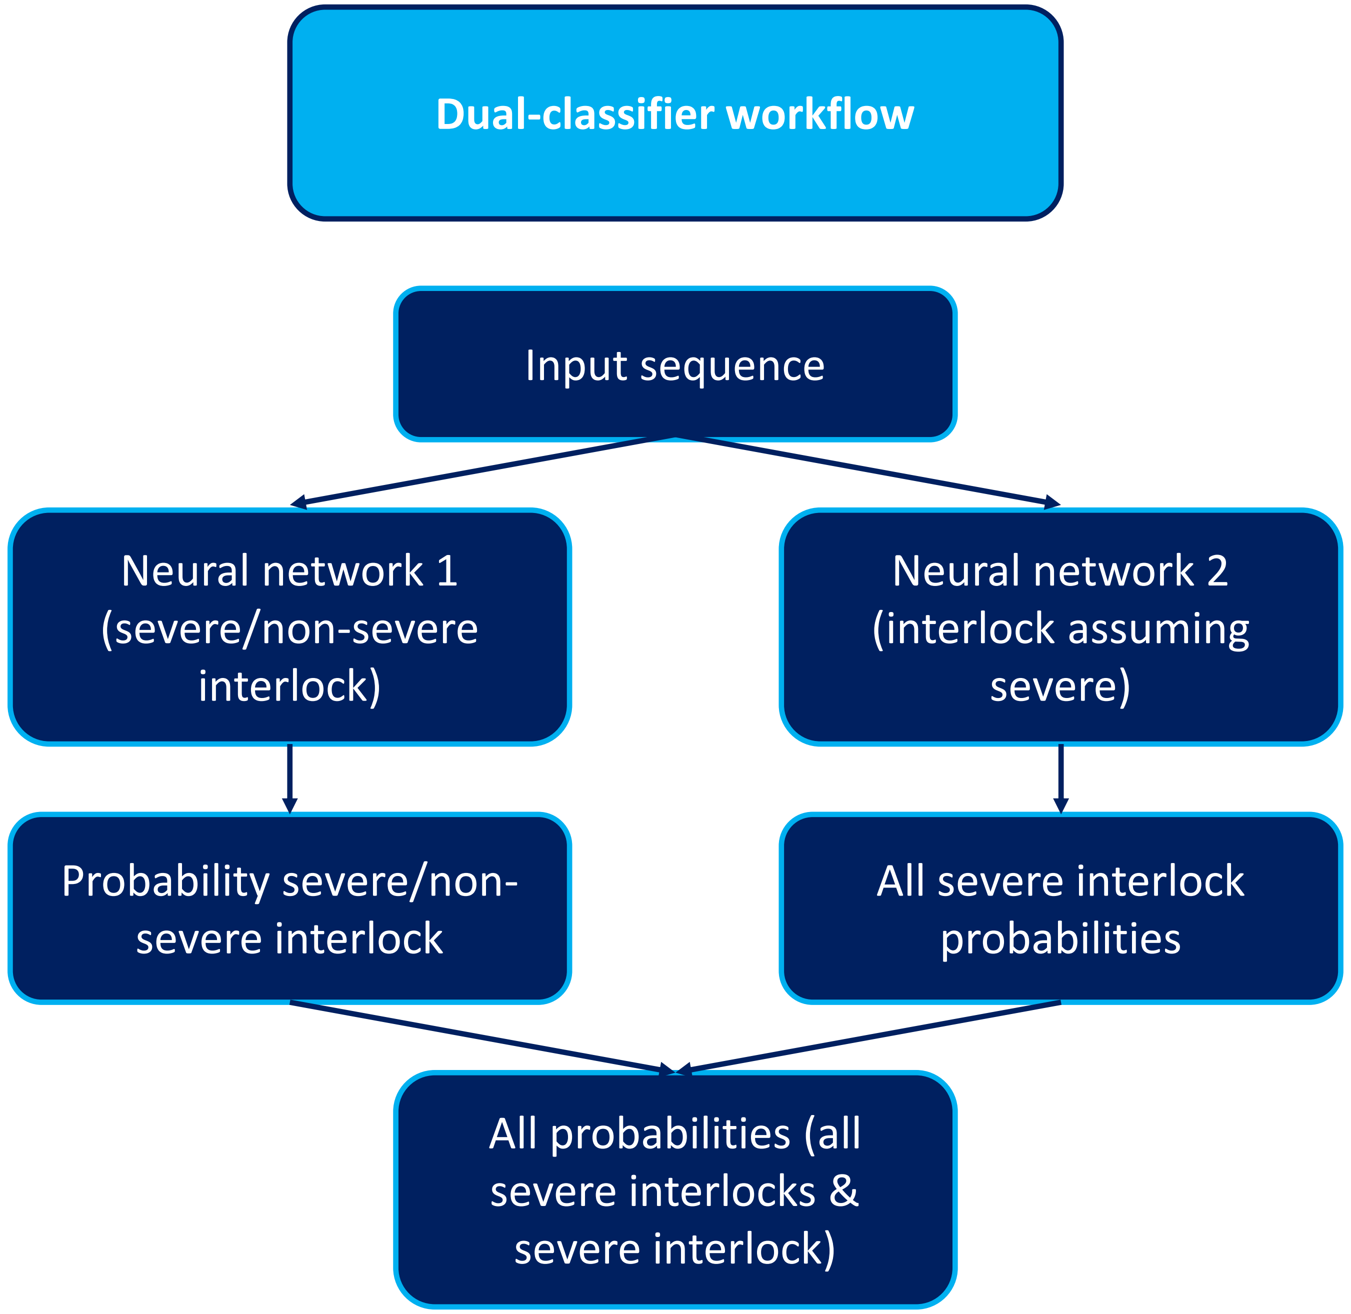

- Novel dual- classifier VAE splits classification into two tasks:
  - severe/non-severe interlock
  - most likely severe interlock (assuming interlock is severe)
- Inputs are sets of 100 consecutive interlocks, plus 50 more sampled up to 1 week prior.
  - Training inputs are labelled with the next interlock code.
- Inputs are projected into a low-dimensional latent space:
  - Severe interlocks with similar causes are spatially close in the latent space.
  - Position of an input in latent space allows for probabilistic predictions of proceeding interlocks.

(Aside) Best Model Selection

- Best models are selected by macro-F1 score.

- F1 defined as the harmonic mean of precision and recall
  - Range between 0 (no accurate classification) and 1 (perfect classification)

$$F_1 = 2 \times \frac{\text{precision} \times \text{recall}}{\text{precision} + \text{recall}} \qquad F_1^{macro} = \sum_{class} F_1^{class}$$

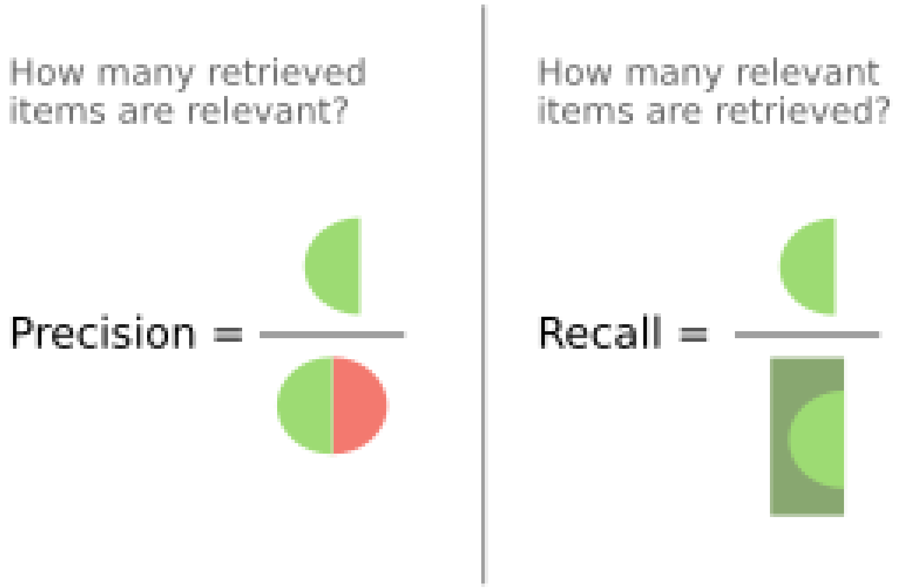

- Accuracy also considered to initially drive learning:

$$\text{Acc} = \frac{\text{correct preds}}{\text{total preds}}$$

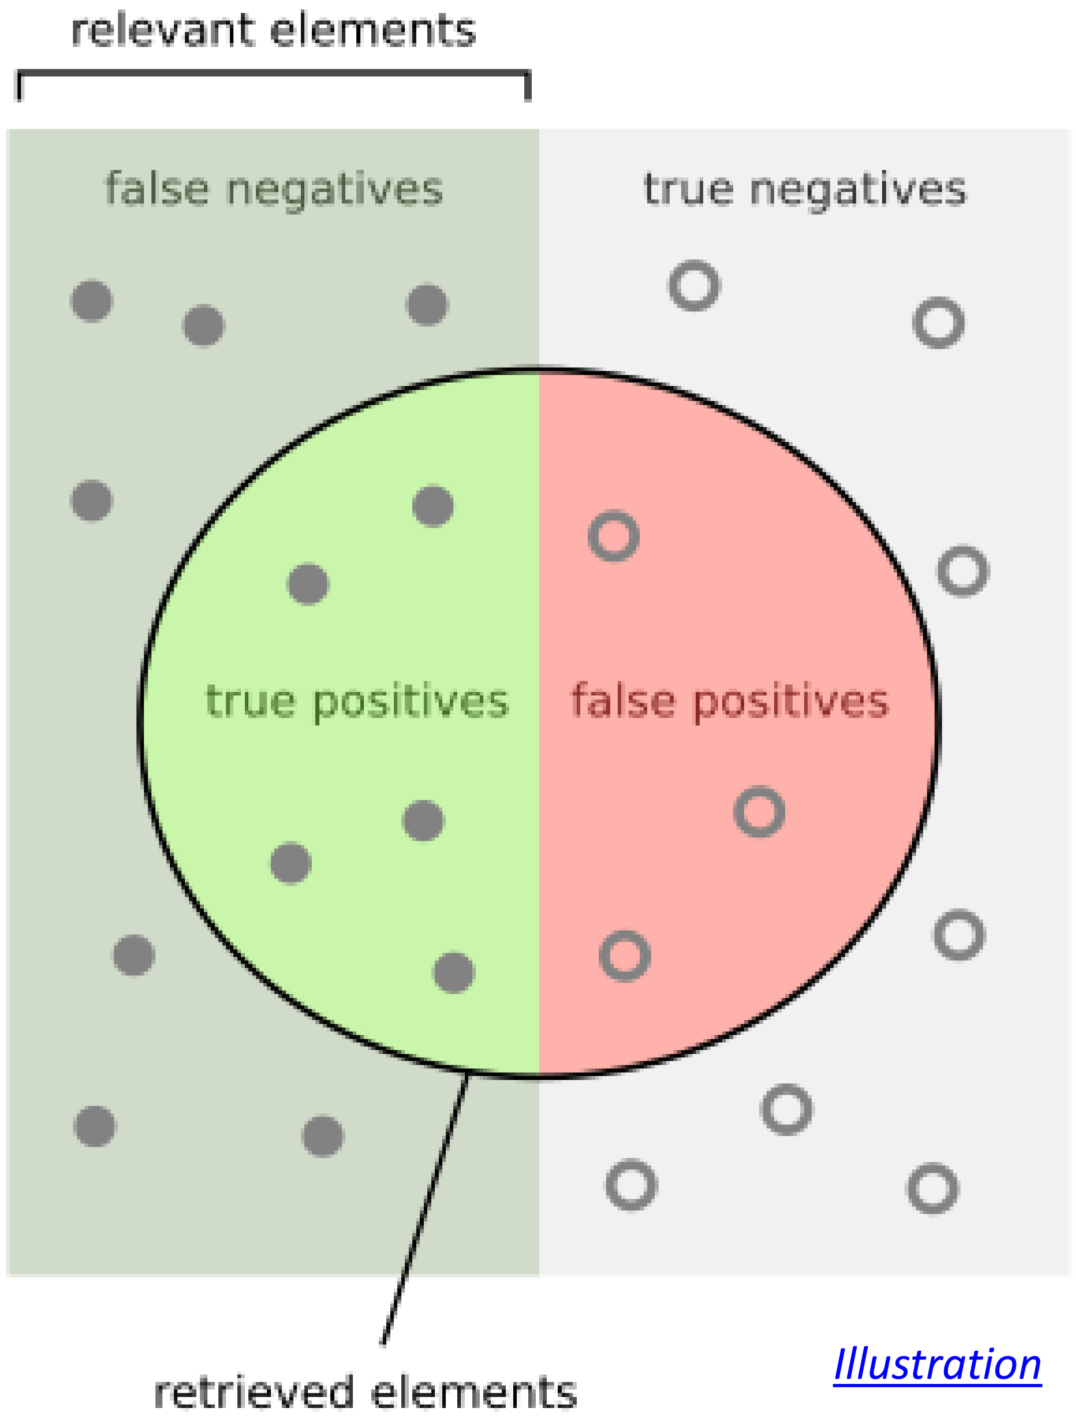

Illustration

Binary Classifier Performance

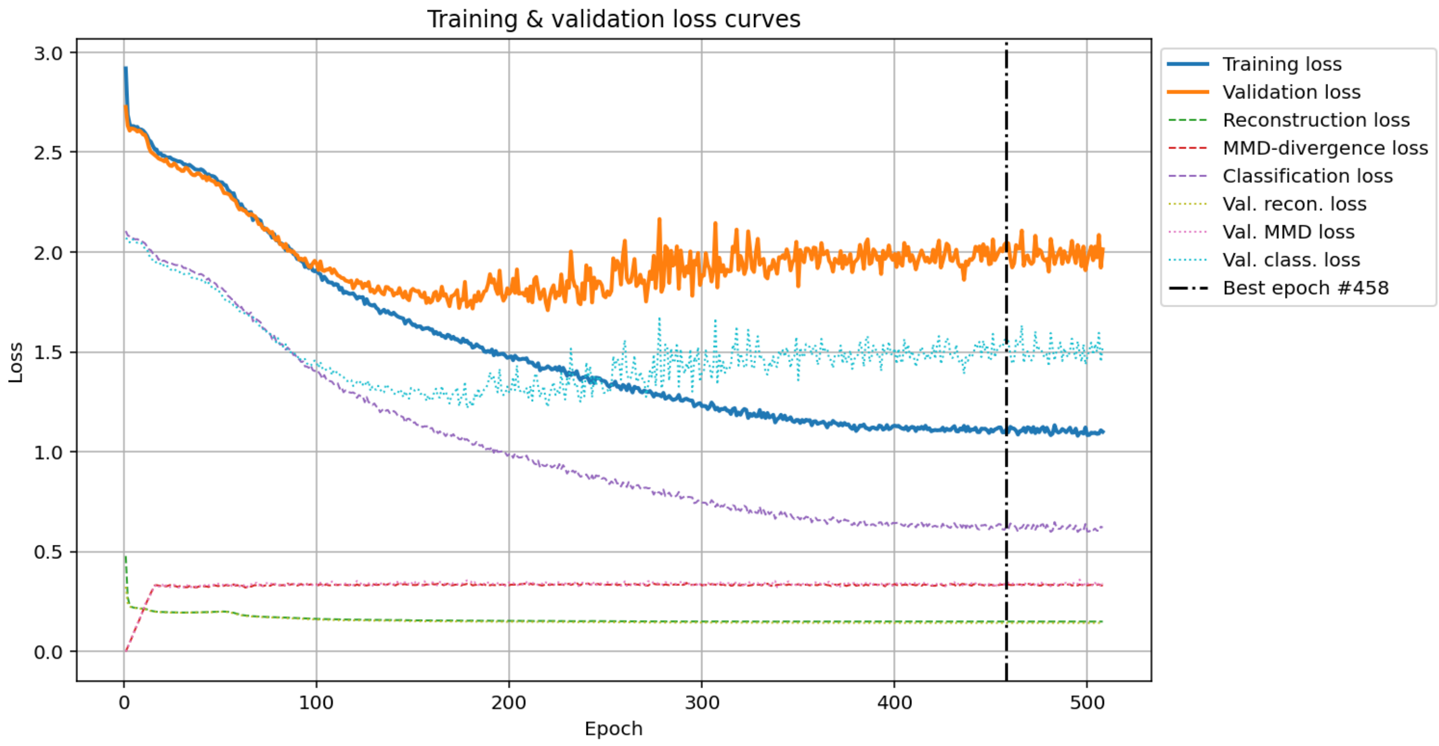

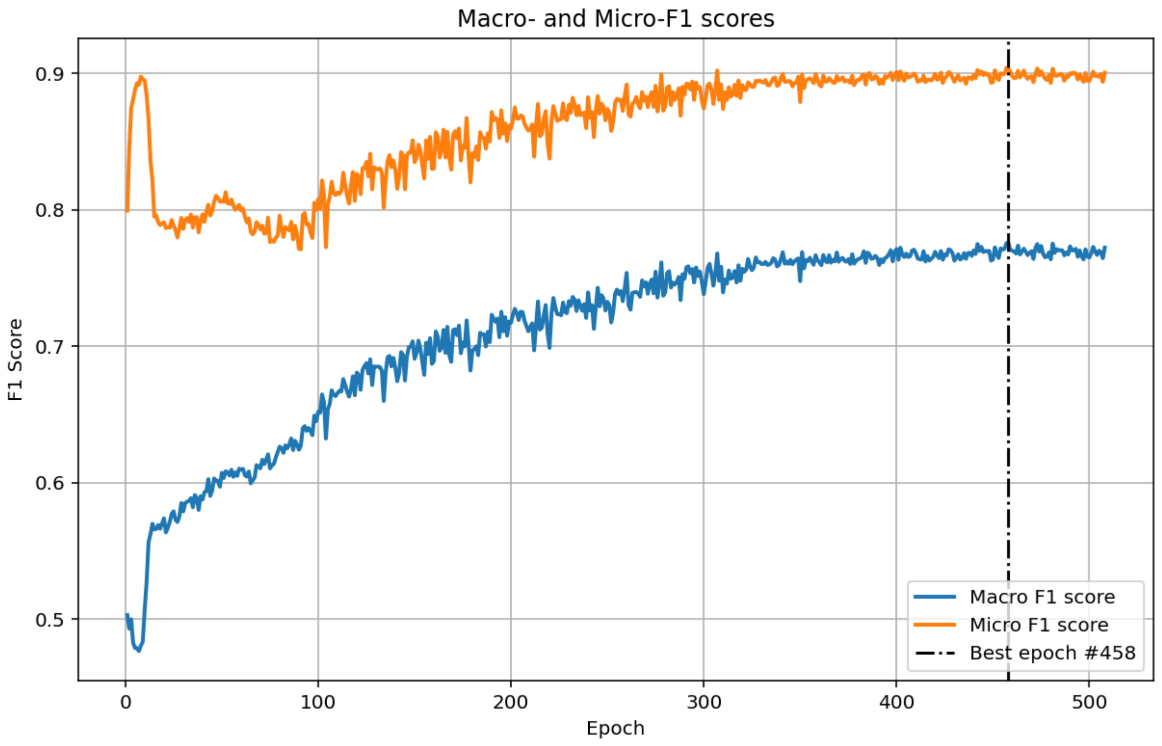

- F1 scores are very strong:

```
Classification accuracy on all samples: 88.81%  (1127/1269)
  Classification accuracy on severe samples: 70.54%     (91/129)
  Classification accuracy on non-severe samples: 90.88%     (1036/1140)
Val Macro-F1: 0.7488, Micro-F1: 0.8881
```

- Model is better at predicting non-severe vs. severe interlocks:
  - Potentially expected due to large class imbalance.
  - Tried several methods e.g. class weighting to improve this with limited success.
  - Improvement will likely require larger volumes of data, specifically with severe interlocks.

Fault Classifier Performance

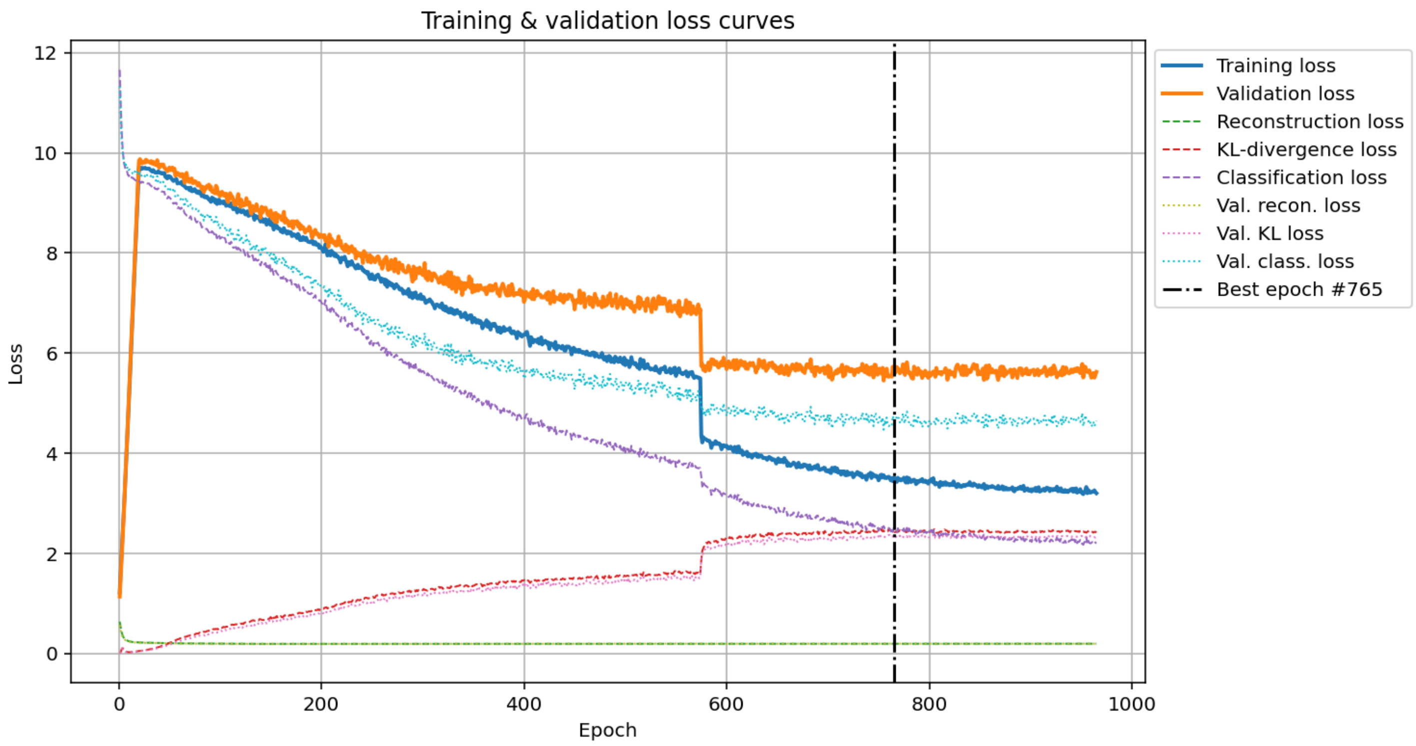

- F1 Scores:

```
Classification accuracy: 65.61%   (248/378)
Classification accuracy (top 3): 87.04%   (329/378)
Subsystem classification accuracy: 81.48%   (308/378)
Val Macro-F1: 0.5438, Micro-F1: 0.6561
```

- Model shows significantly better-than-even odds of predicting exactly the correct next severe interlock in series:
  - Top 3 accuracy is approaching 90%.
  - Most of the time all top 3 interlocks are very similar e.g. 212110, 212111, 212114 (Dose less than expected).

- Significant improvement over this is difficult due to the inherently random nature of breakdowns.

Fault Classifier Latent Space

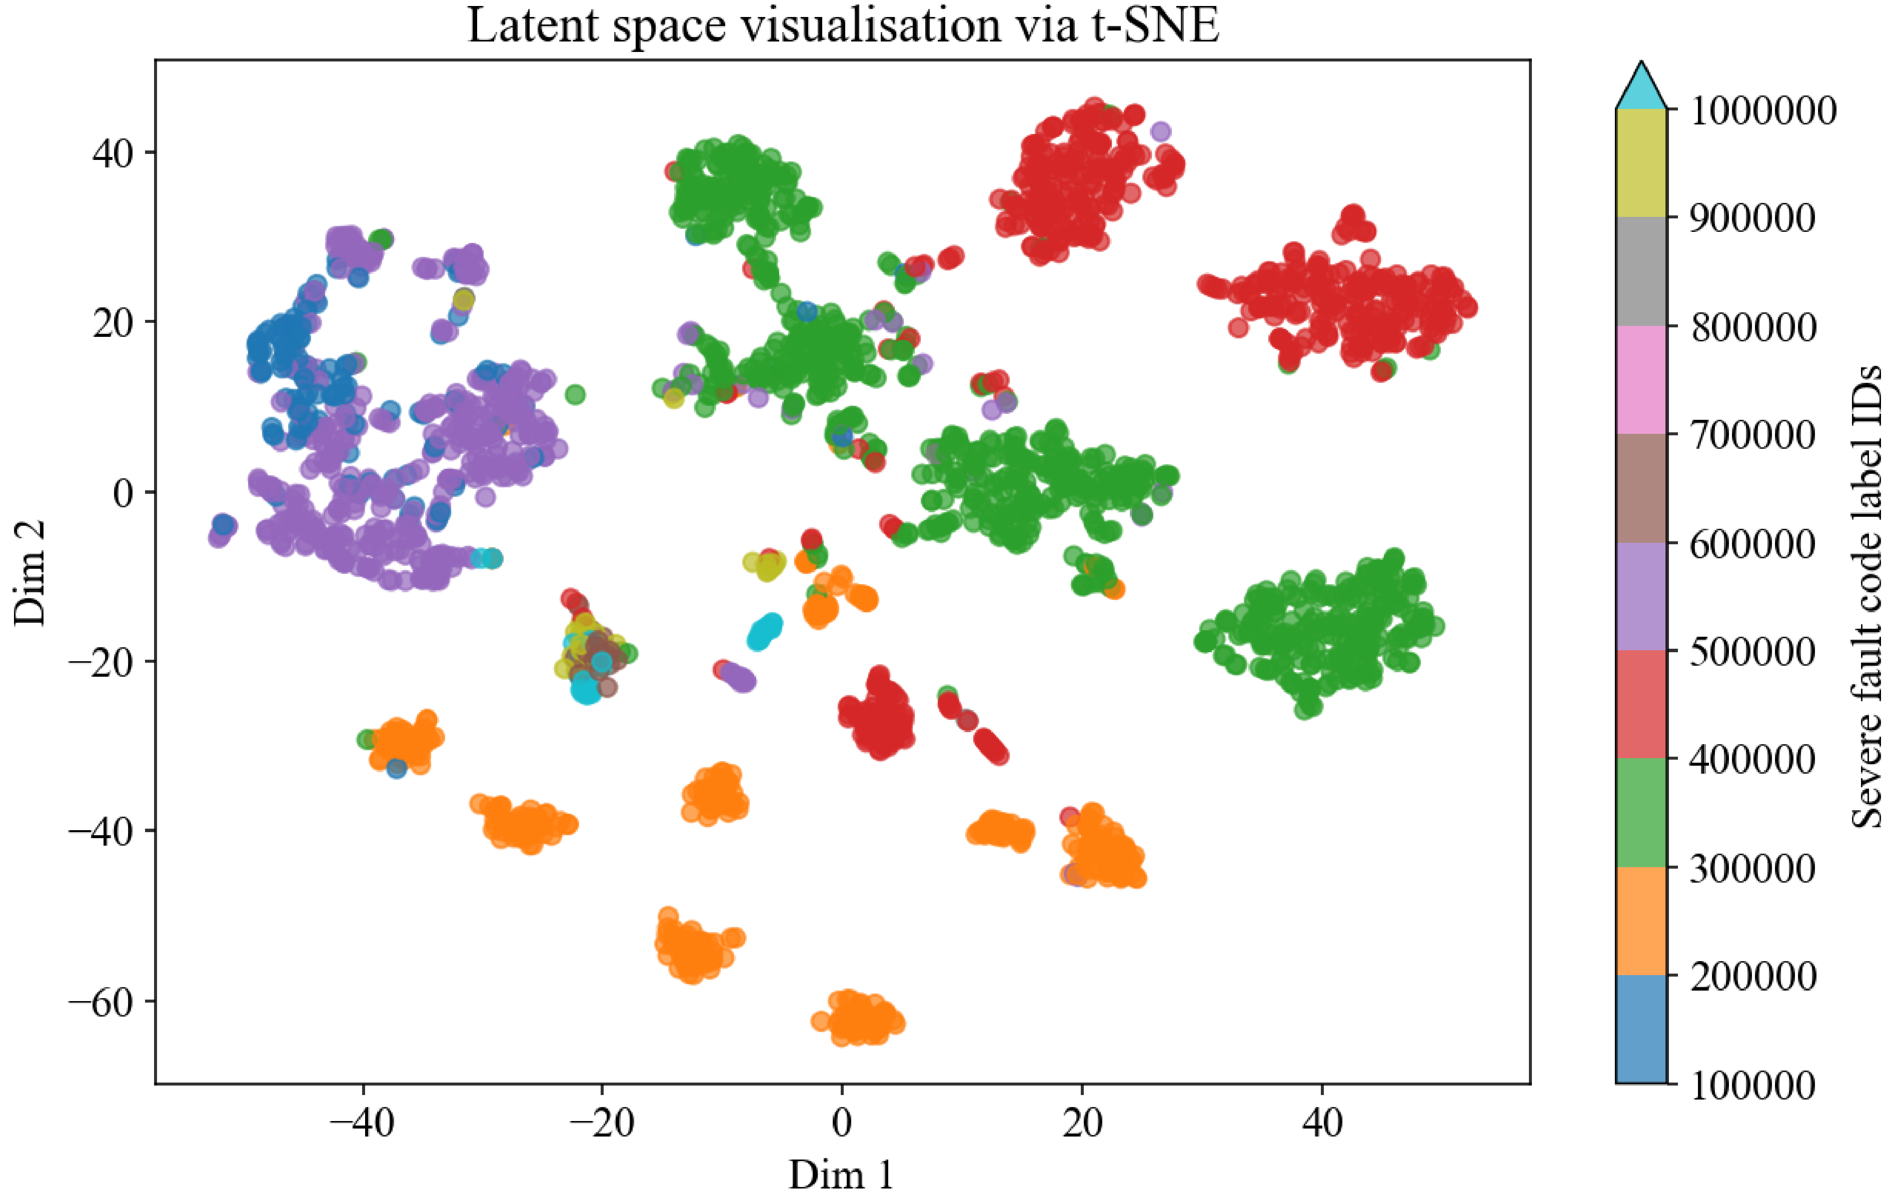

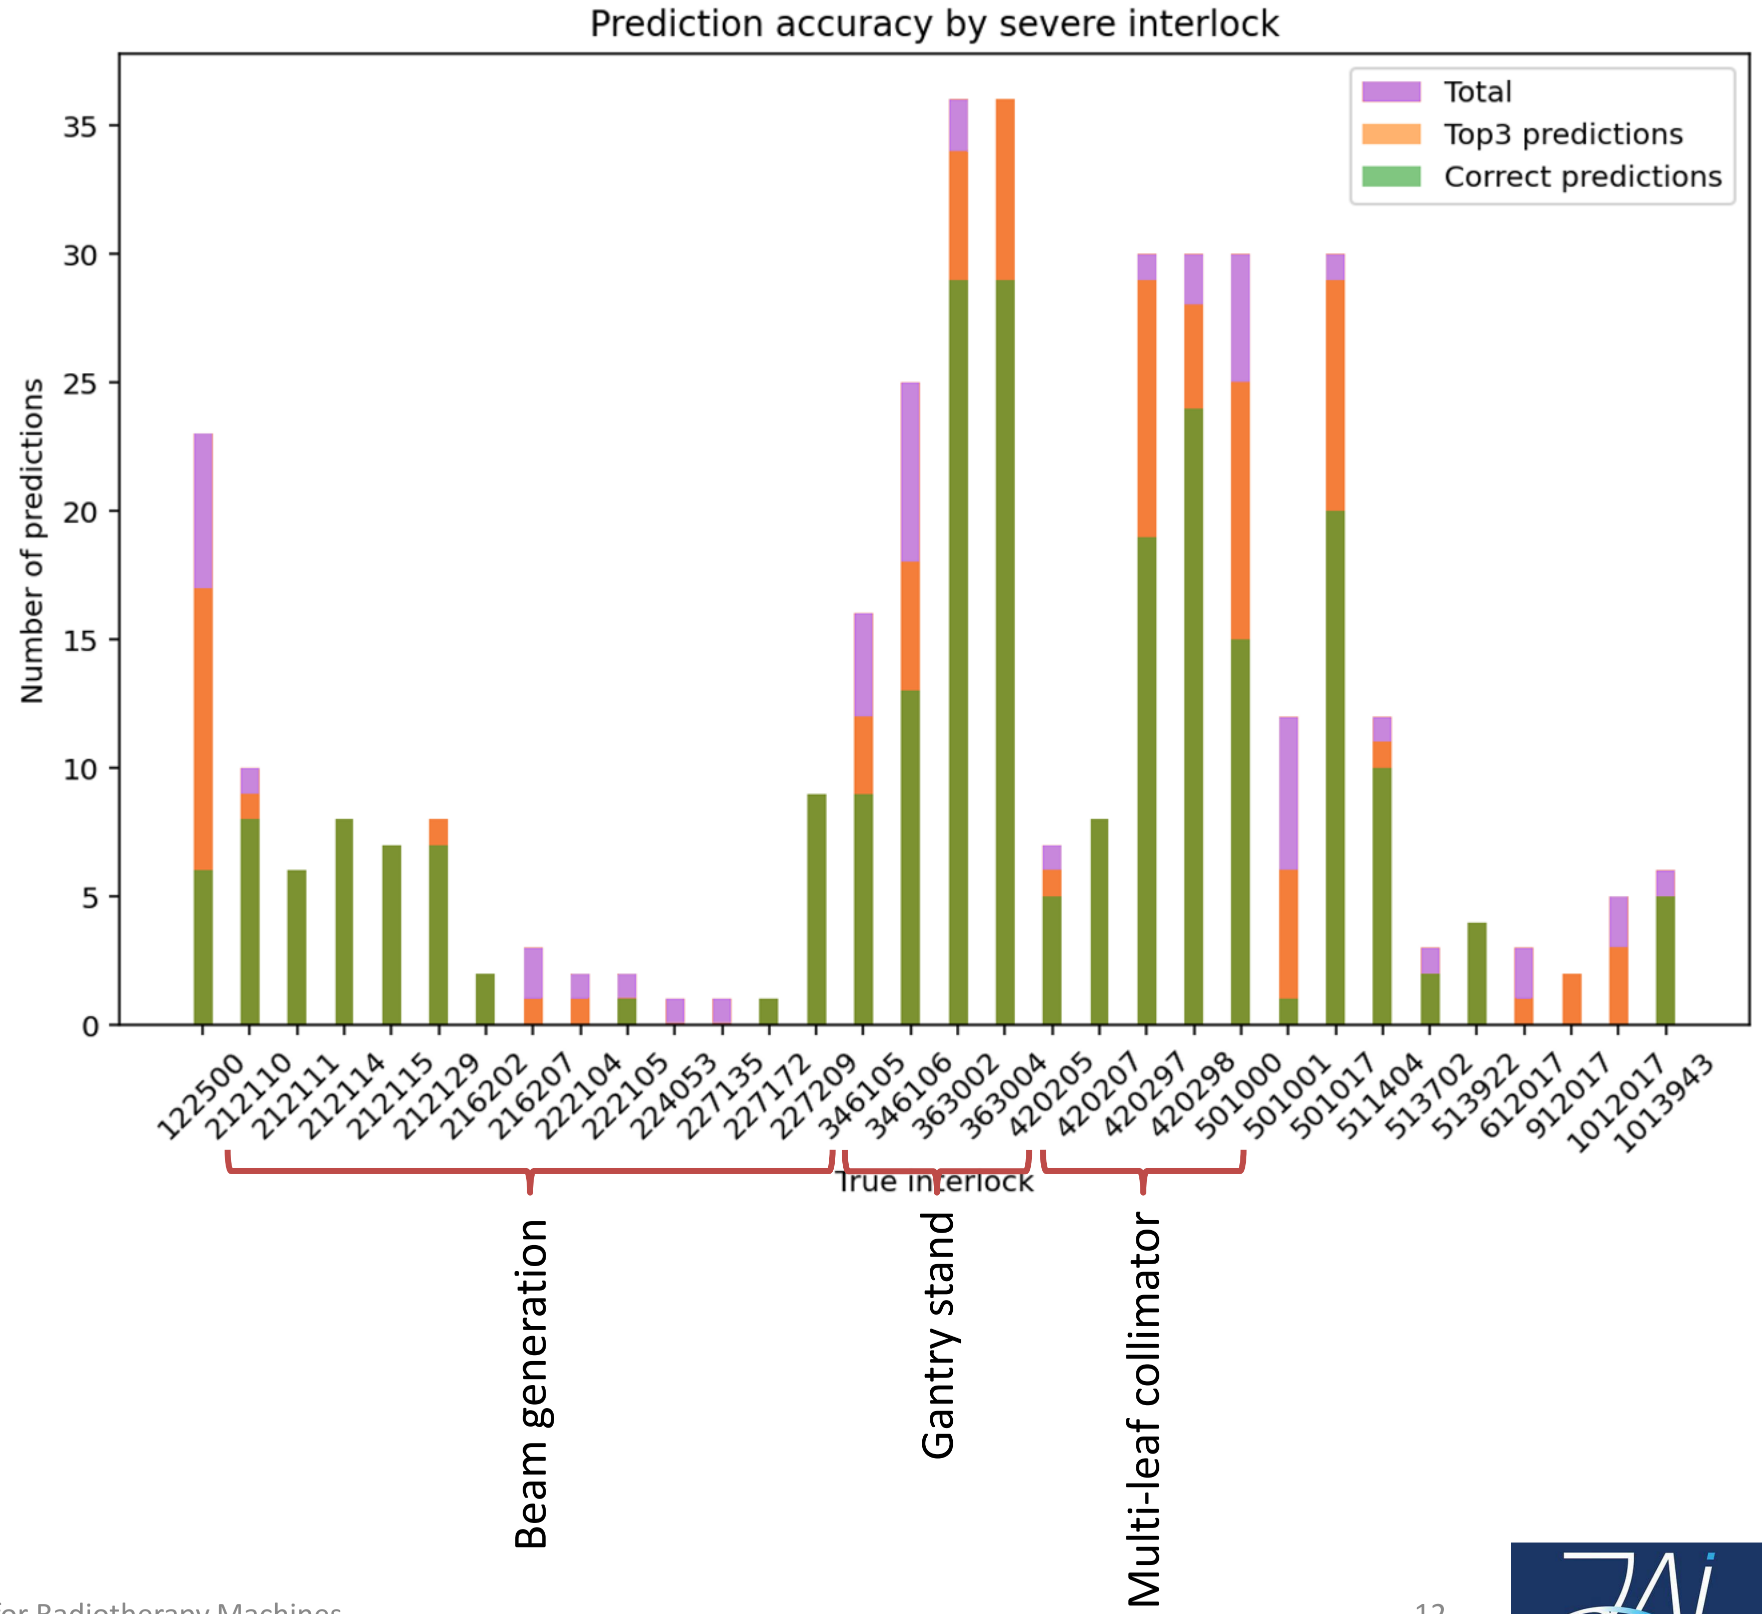

Comparison to other prediction methods

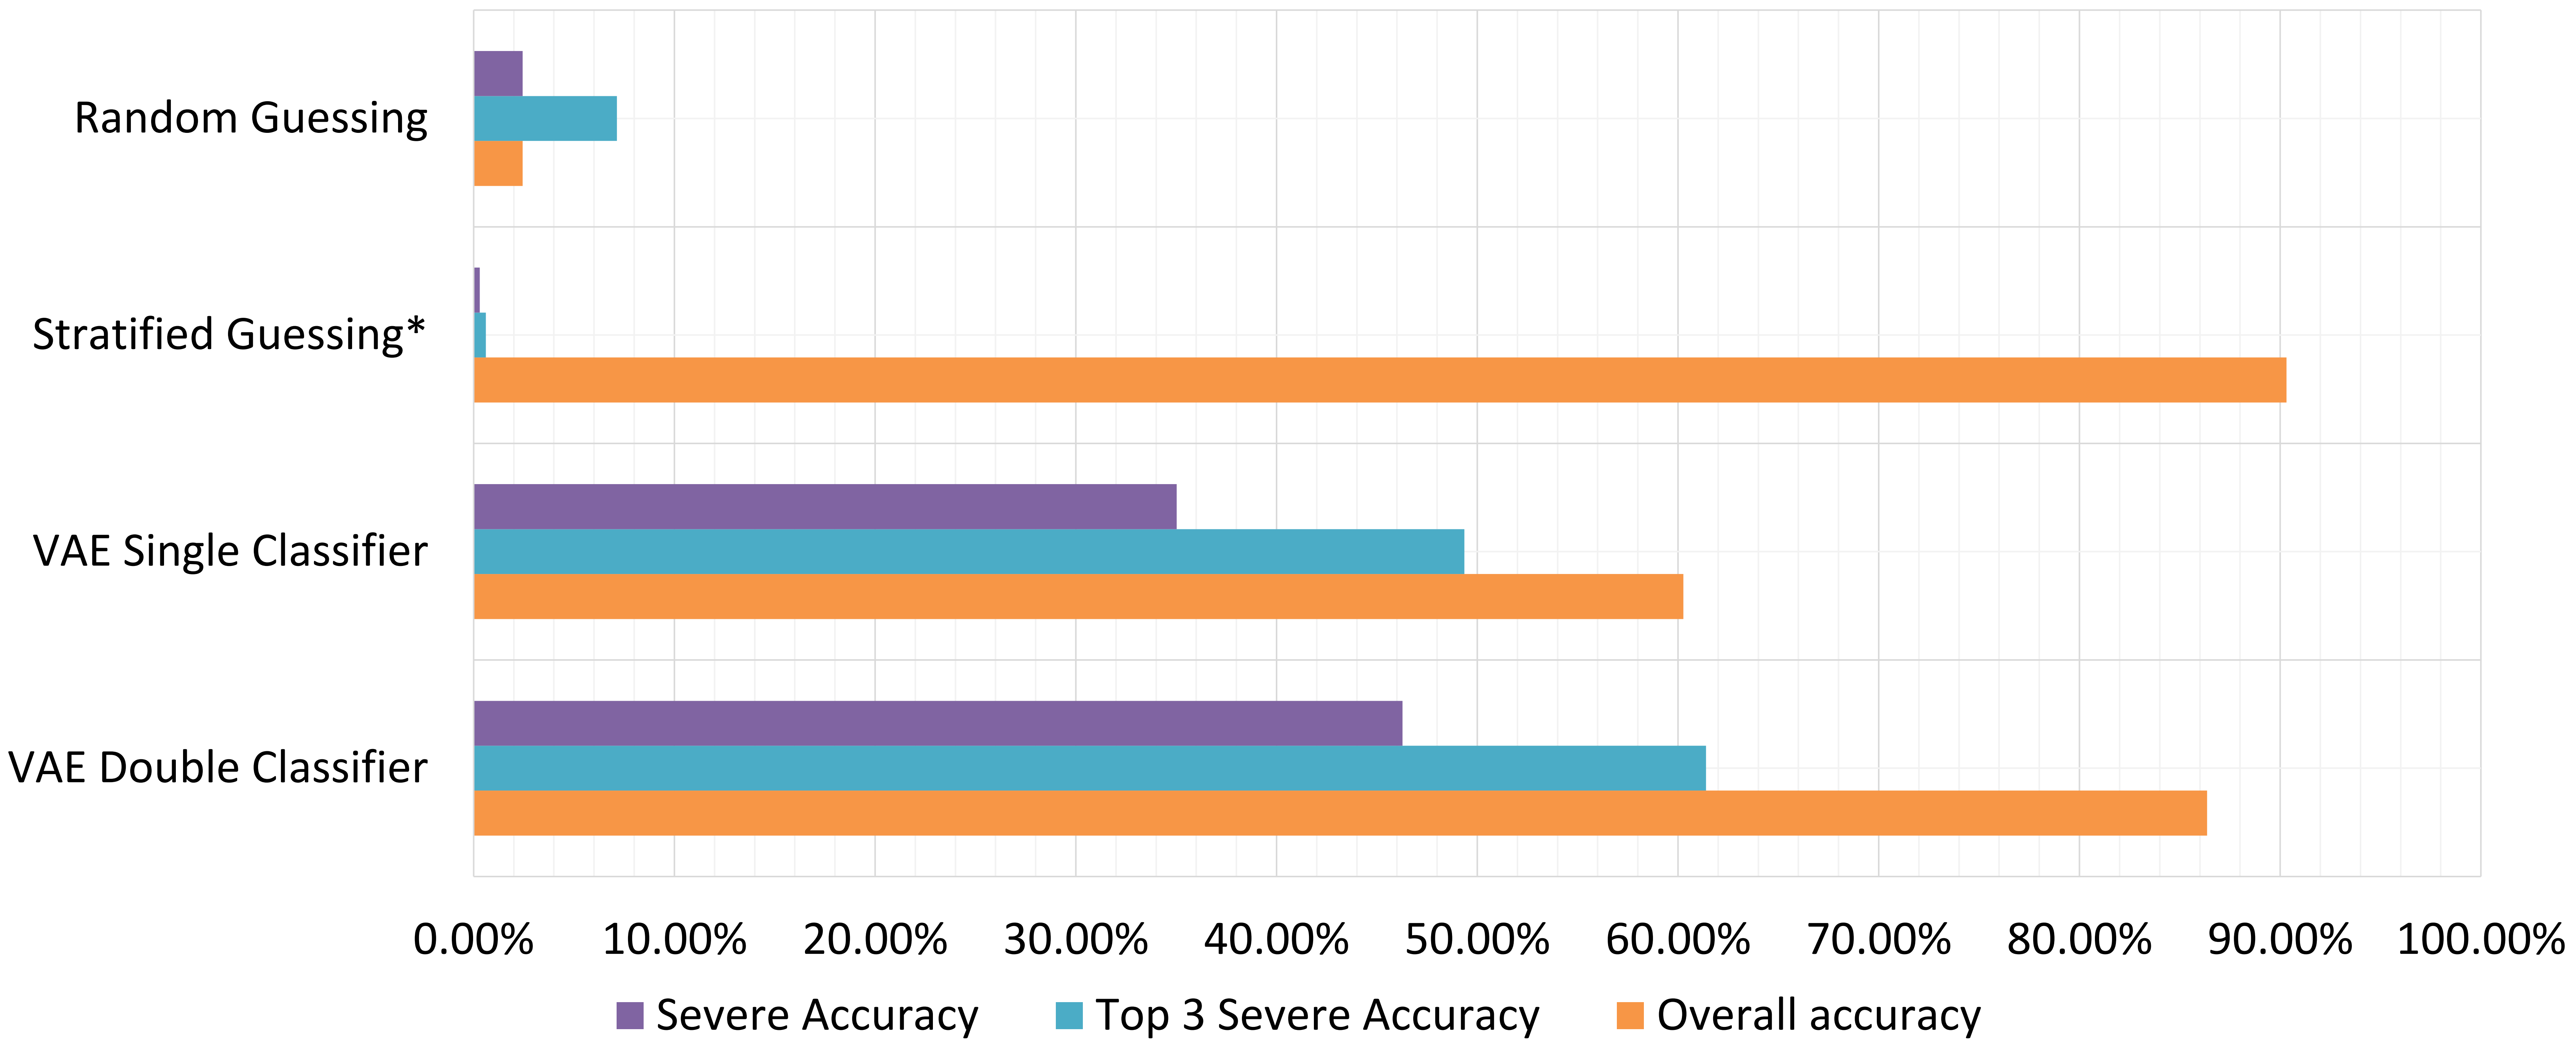

*Overall accuracy of stratified guessing is boosted due to almost exclusively predicting “no severe fault”

Time-series Probabilities

Interlock 363002 – Laser guard fault, potential failure if often. Medium severity.

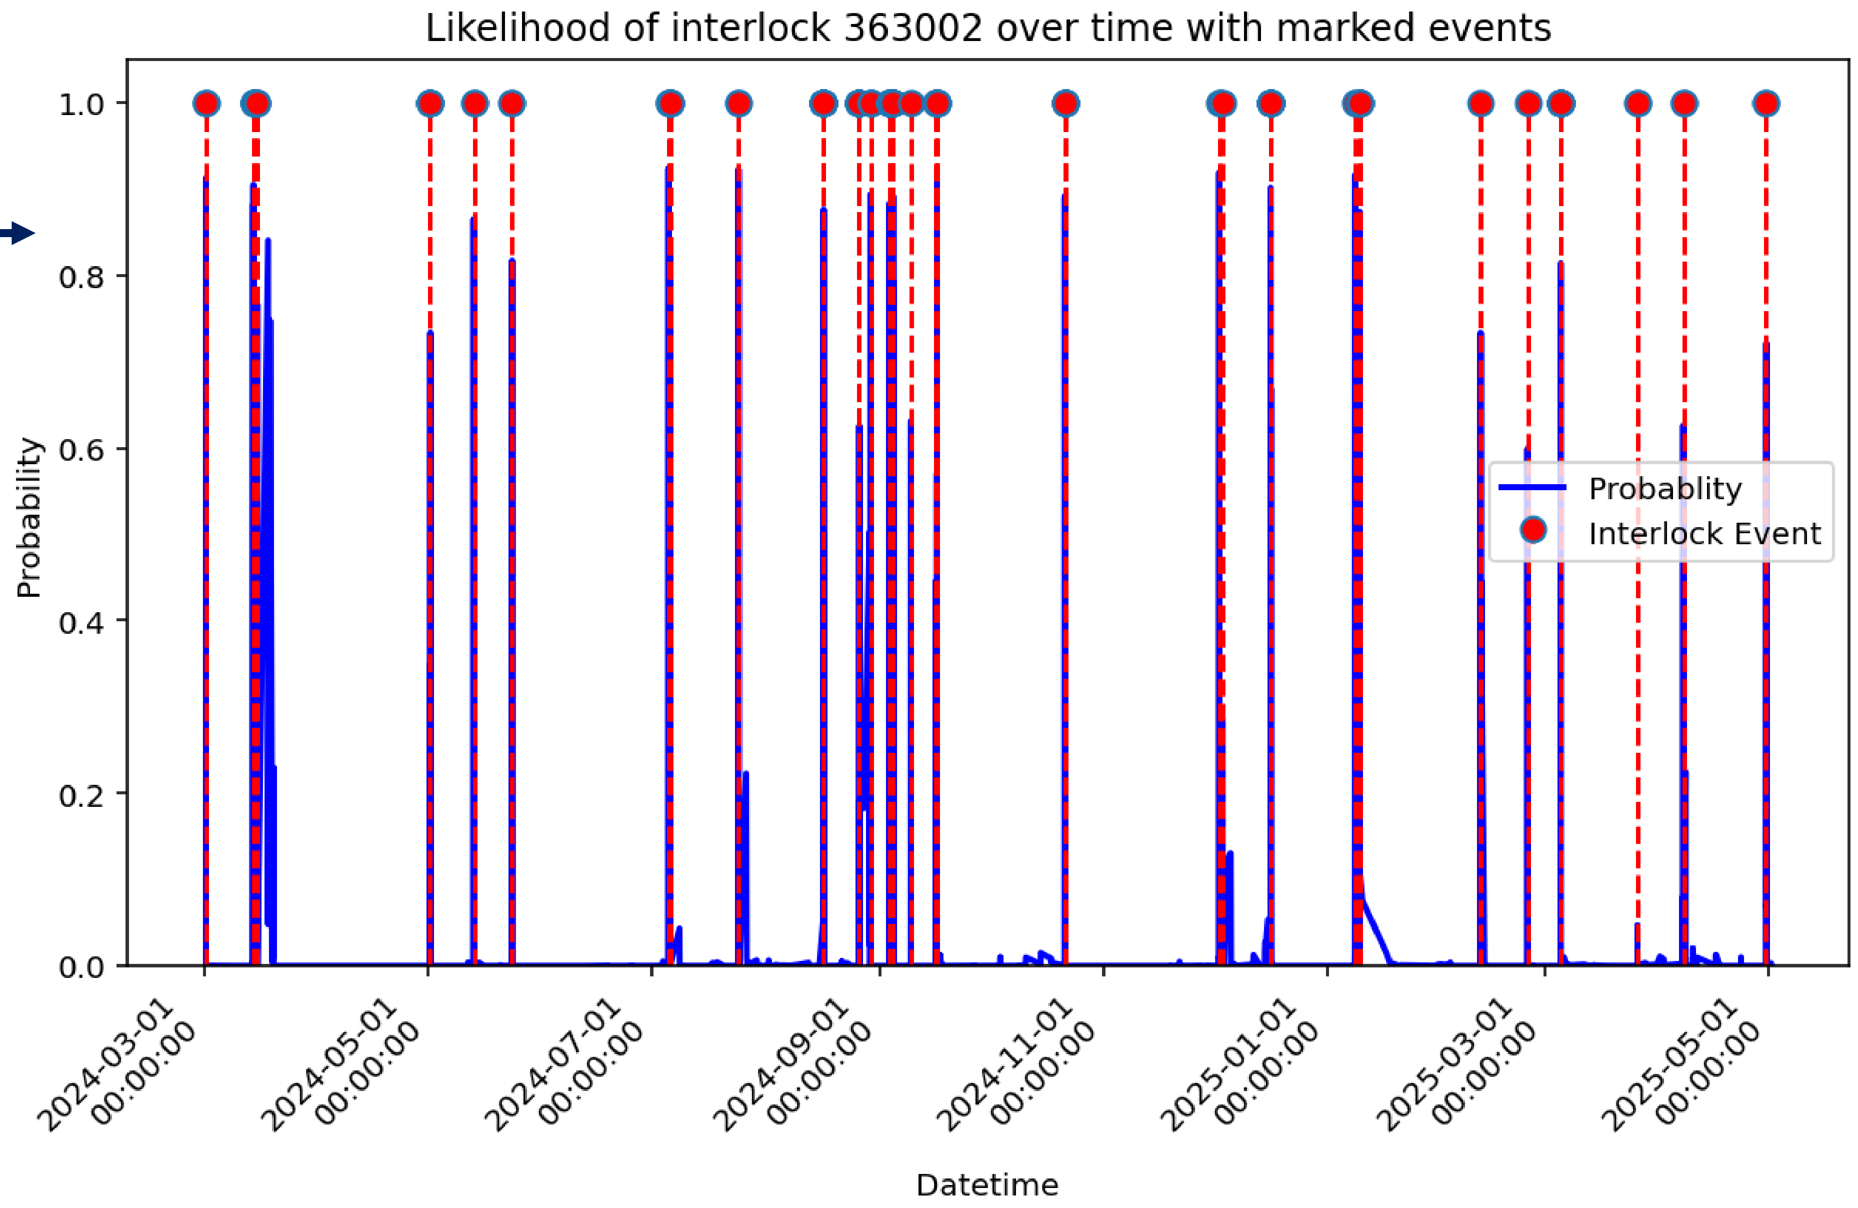

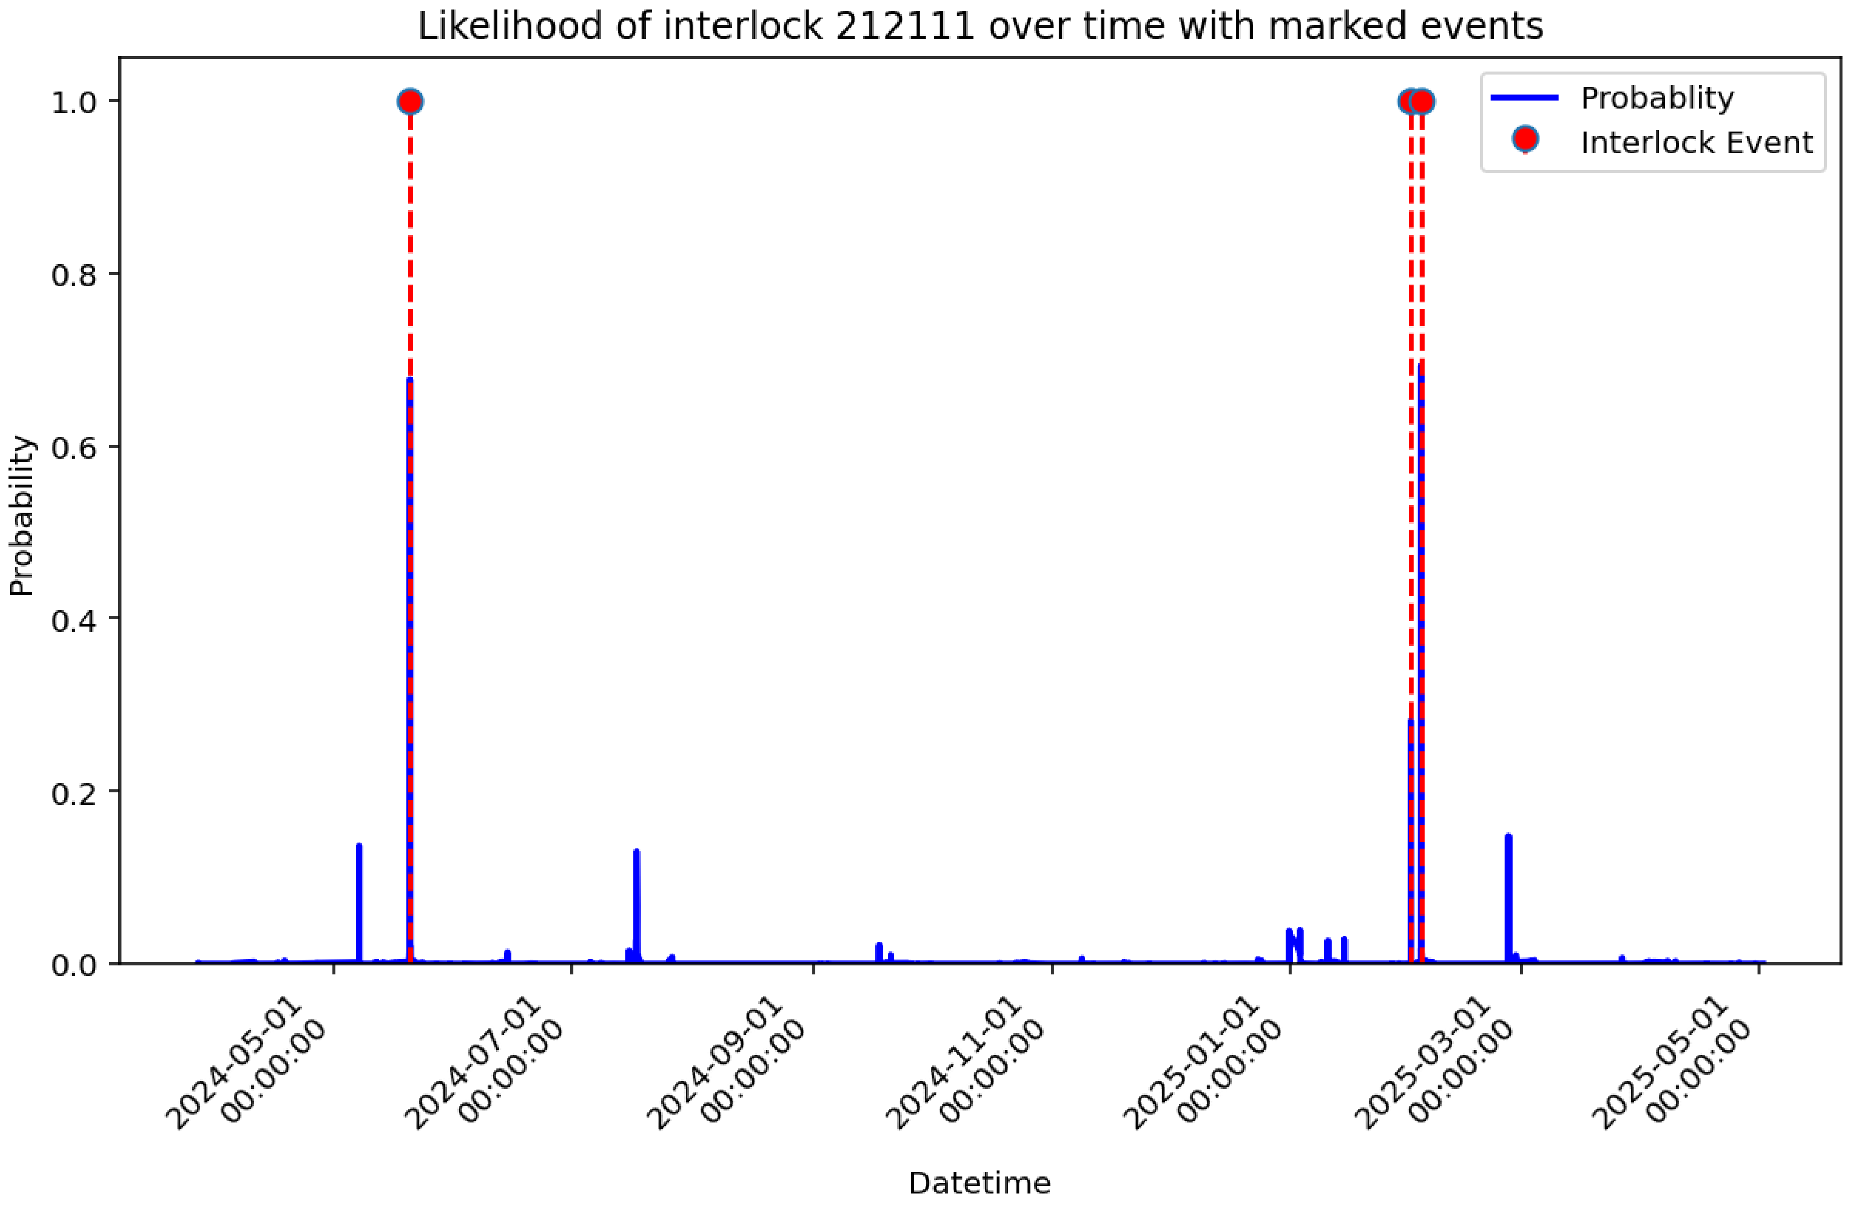

Interlock 212111 – Target under dose rate, potential HV Thyratron problem. Medium severity.

Next Steps

- This fault prediction process can be generalised away from the Varian TrueBeam.
  - Algorithm is applicable to any other Varian machine, and any RT machine with a similar logging style.
  - Approach is *data-agnostic* so can be applied to any machine with a similar fault reporting system, regardless of function.

- Transfer learning will allow customisation of models to a treatment centre's needs.
  - Models can be generally trained on a central database, then tweaked by a user with their own data.
  - This doesn't require any release of data by a treatment centre, allaying concerns about privacy.

- Proof of concept for fault prediction opens the door to more powerful ML algorithms.
  - Transformers can capture more information about contextual and long-term relationships with inherent temporal encoding and attention mask.

Leadley, S. (2025). Predictive Modelling of Radiotherapy LINAC Downtime using Variational Autoencoders [Poster]. VHEE'25 conference, Daresbury.

Thank you!

Collaborators & Supervisors:
Torsten Liebig, *Exeter-Wonford Hospital*
Jonathan Lane, Gary Grogan, *Oxford-Churchill Hospital*
Andrew Robinson, Ian Dulgarn, *Cambridge-Addenbrooke's Hospital*

Manjit Dosanjh, *University of Oxford, John Adams Institute, CERN*
Raj Jena & Doris (Xin) Du, *University of Cambridge*
Laurence Wroe & Steinar Stapnes, *CERN*

Contact: *samuel.Leadley@physics.ox.ac.uk*

This work was supported by funding from the John Adams Institute (JAI), Oxford, and the Science & Technology Facilities Council (STFC).

John Adams Institute for Accelerator Science

Variational AutoEncoder – Neural Network

Input from string of previous interlocks. Input size = vocab size.

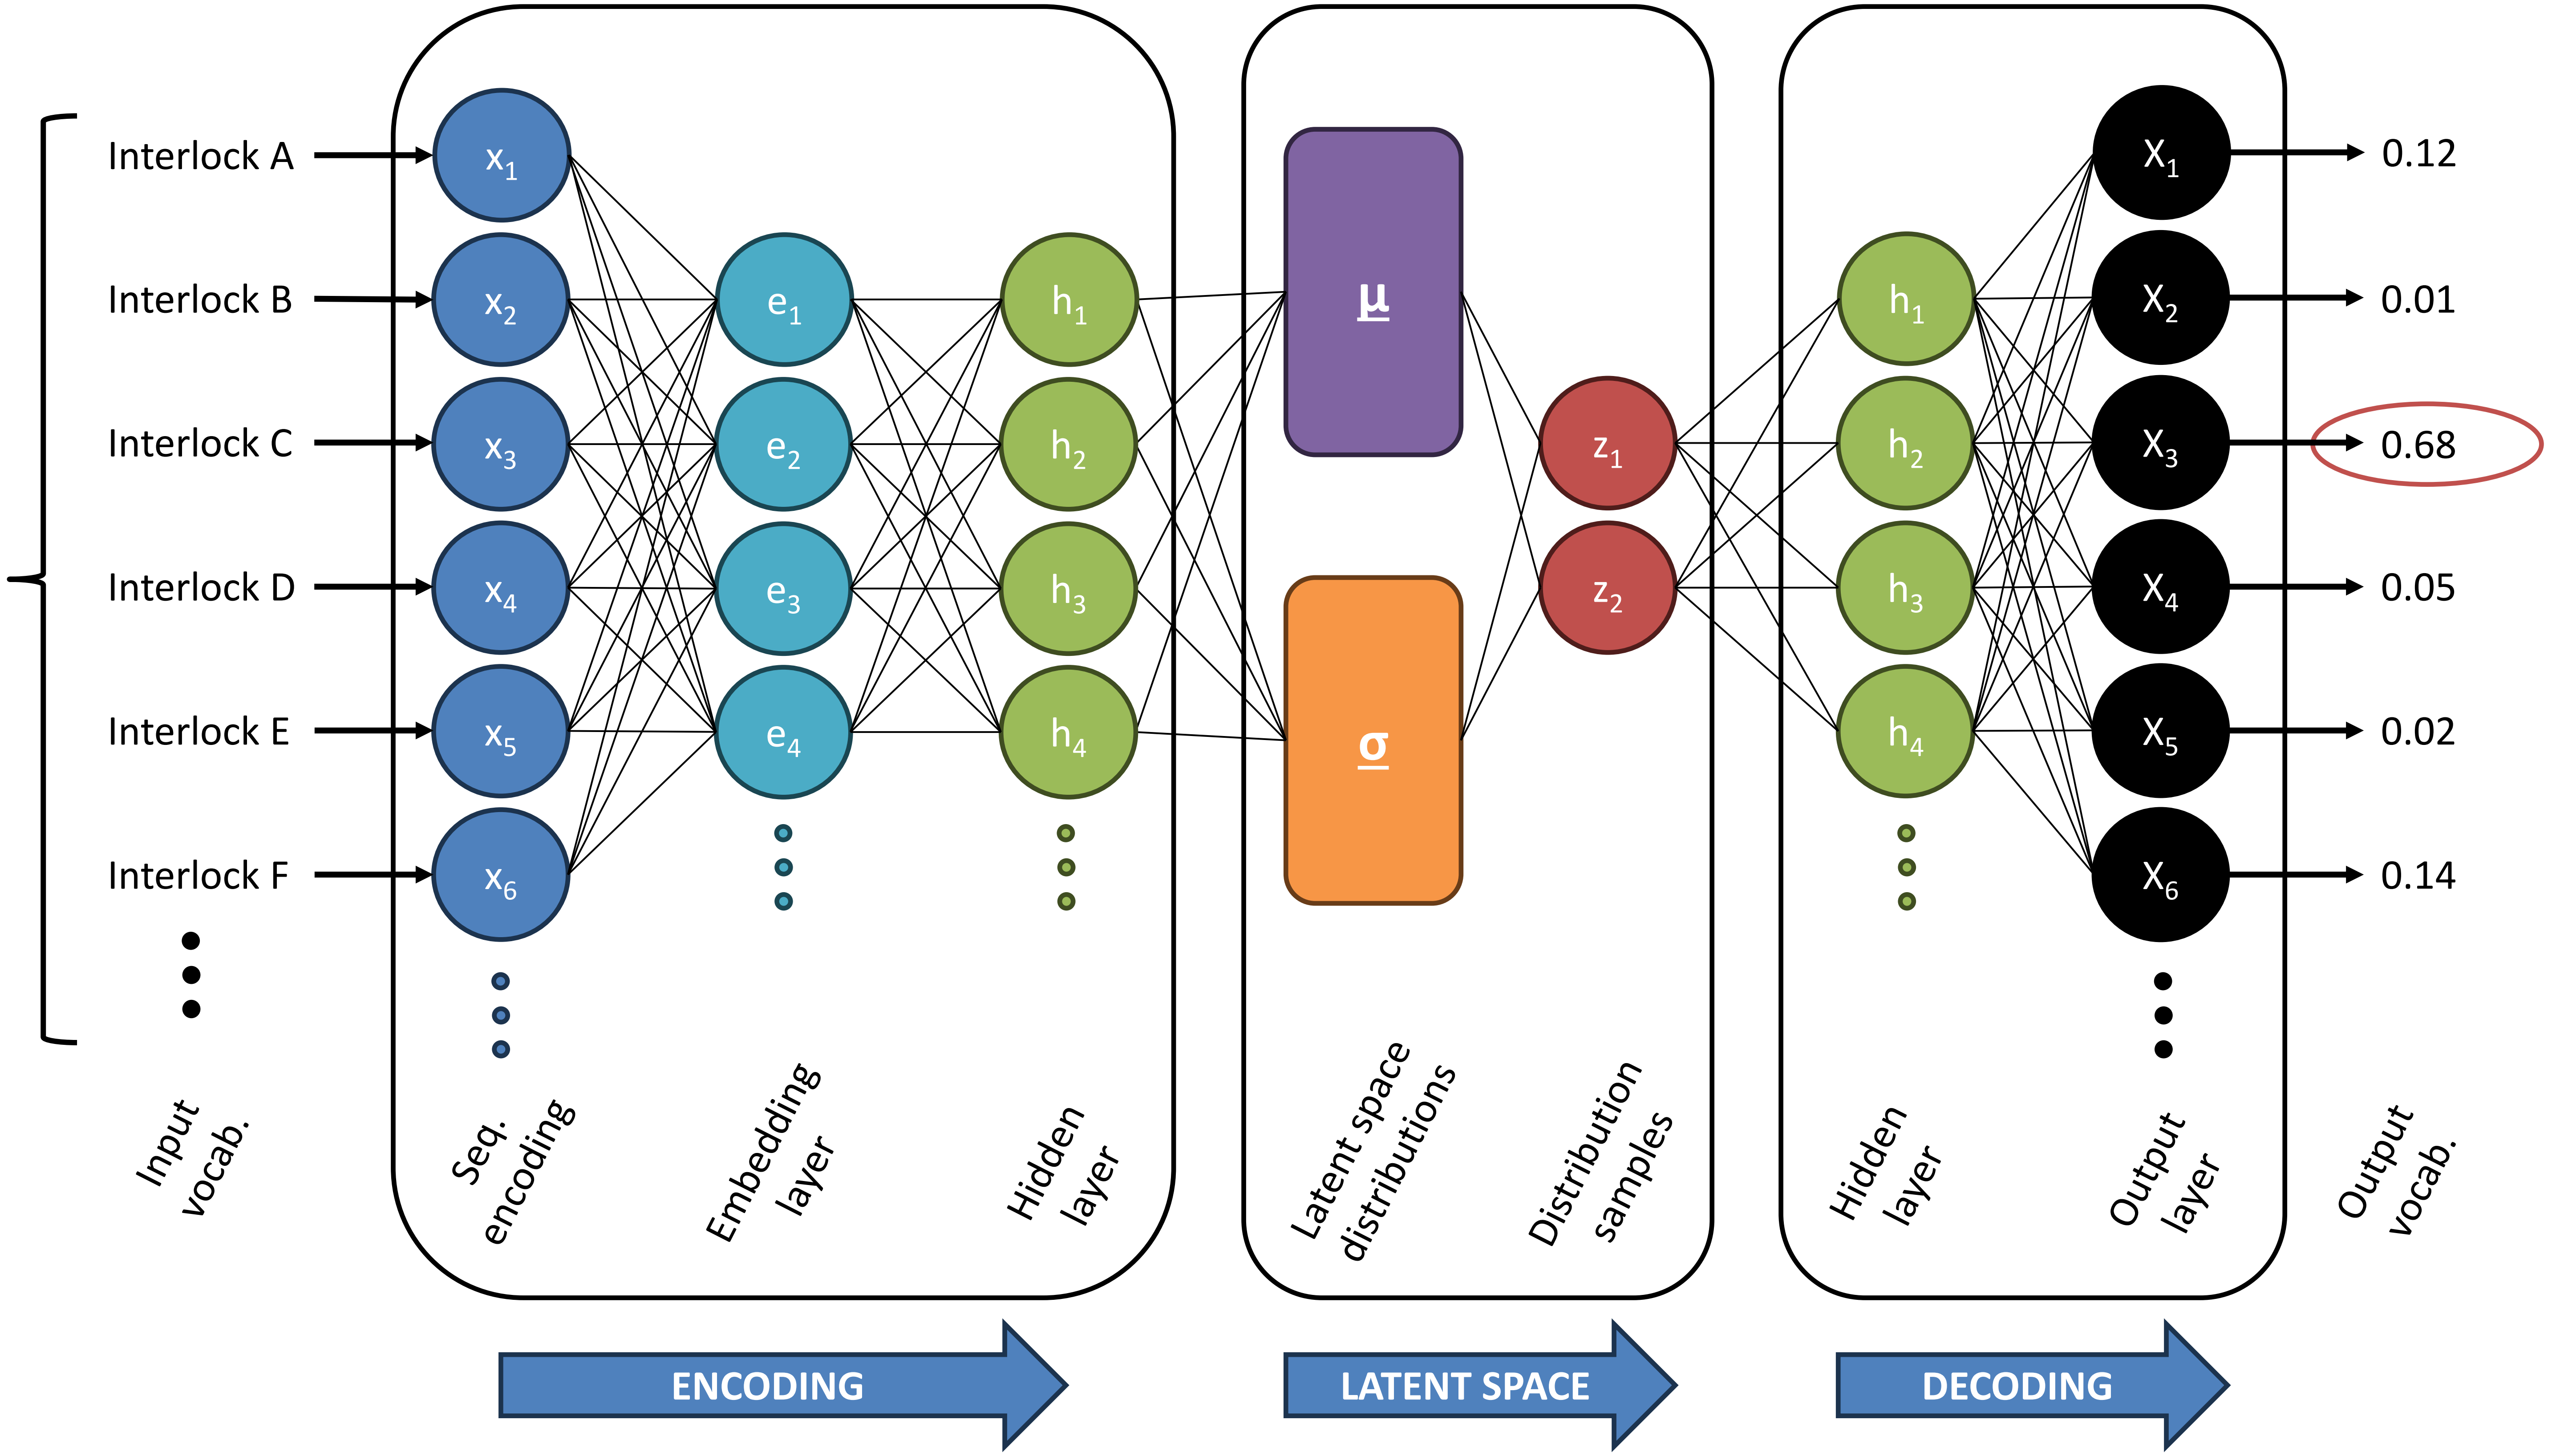

Most likely output is next fault interlock prediction. Output size = number of severe fault + 1

John Adams Institute for Accelerator Science