Identify Patients and Health Care Workers Using Transfer Learning on a Pre-trained Convolution Neural Network

Kelvin Tham

29 Mar 2018

Background: “Why am I here today?”

- The ‘Underdog’
- Healthcare needs driving innovation

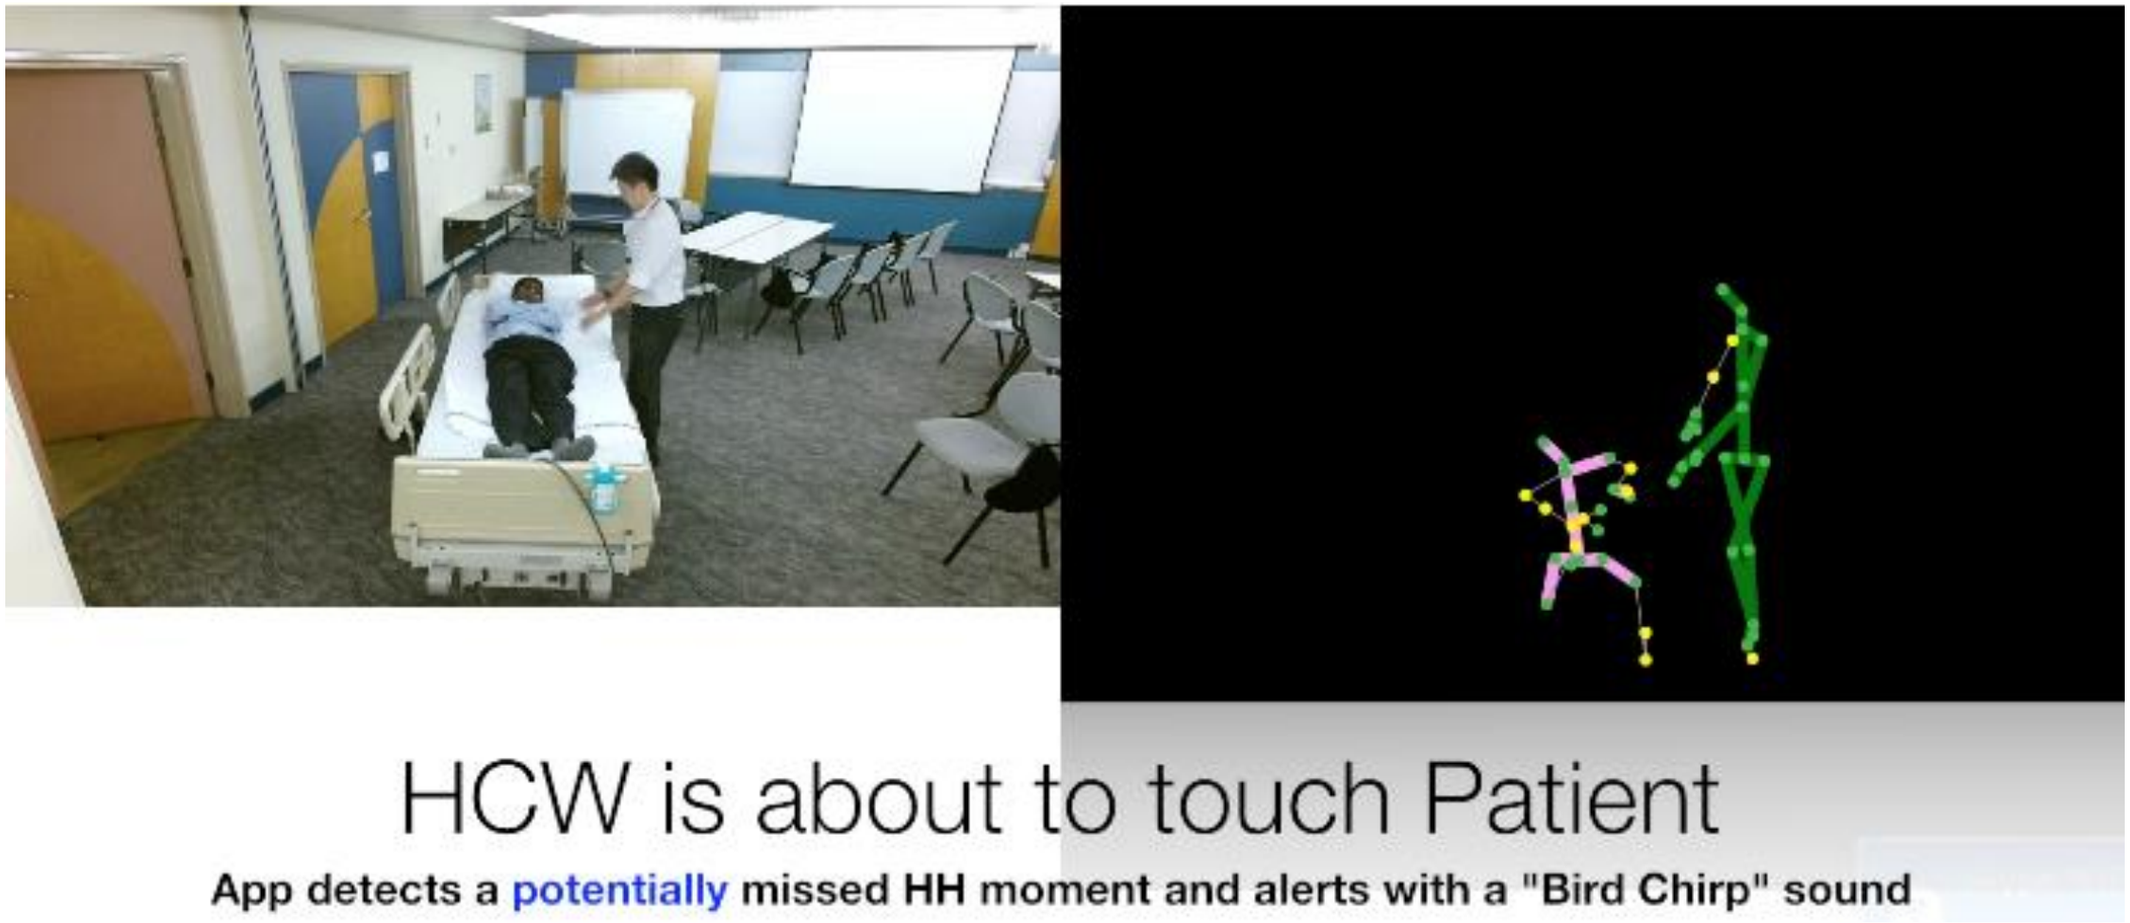

- Explore Machine Learning, no traction

Impact of Ai6: “Passive student to Active Creator”

- Setup - Daunting for newbies

Machine Learning Tools

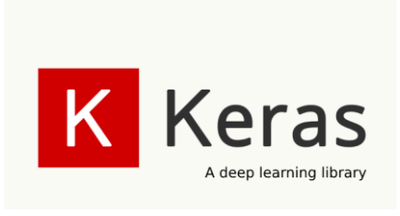

Deep Learning Libraries

Computing Power

- Fastai - Practical
- Support and Network - Kept me going
- Ai6 Challenge -> Capstone project

Getting Started: “choosing wisely”

- From Fastai
  - Transfer Learning

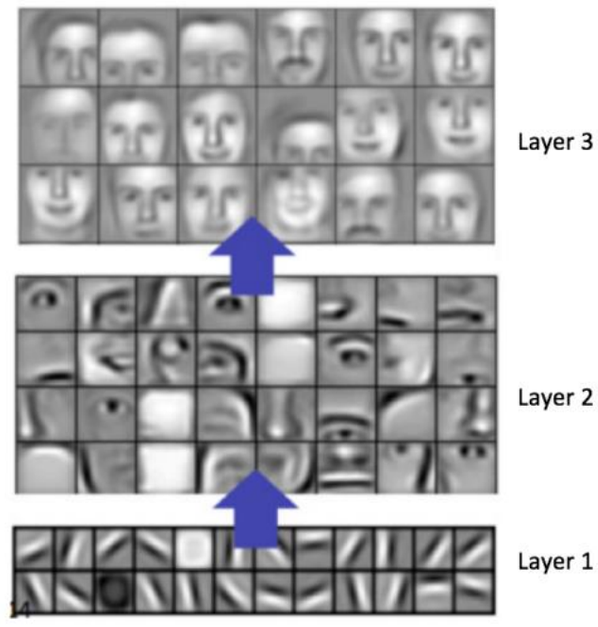

Figure 17: Learned features from a Convolutional Deep Belief Network. Source |

- Fastai library - easy to use
- Pytorch - Read documentation, Transfer Learning Tutorial

Getting Started: “choosing wisely”

- Data preparation, normally ‘not sexy’ but…

Getting Started: “choosing wisely”

- Data preparation
- Web scrapping

Dataset gathered:

- Total: 564
  - Train: 437
    - Doctors: 95
    - Nurses: 171
    - Patients: 171
  - Validation: 96
    - Doctors: 21
    - Nurses: 36
    - Patients: 39
  - Test: 31
    - Doctors: 10
    - Nurses: 10
    - Patients: 11

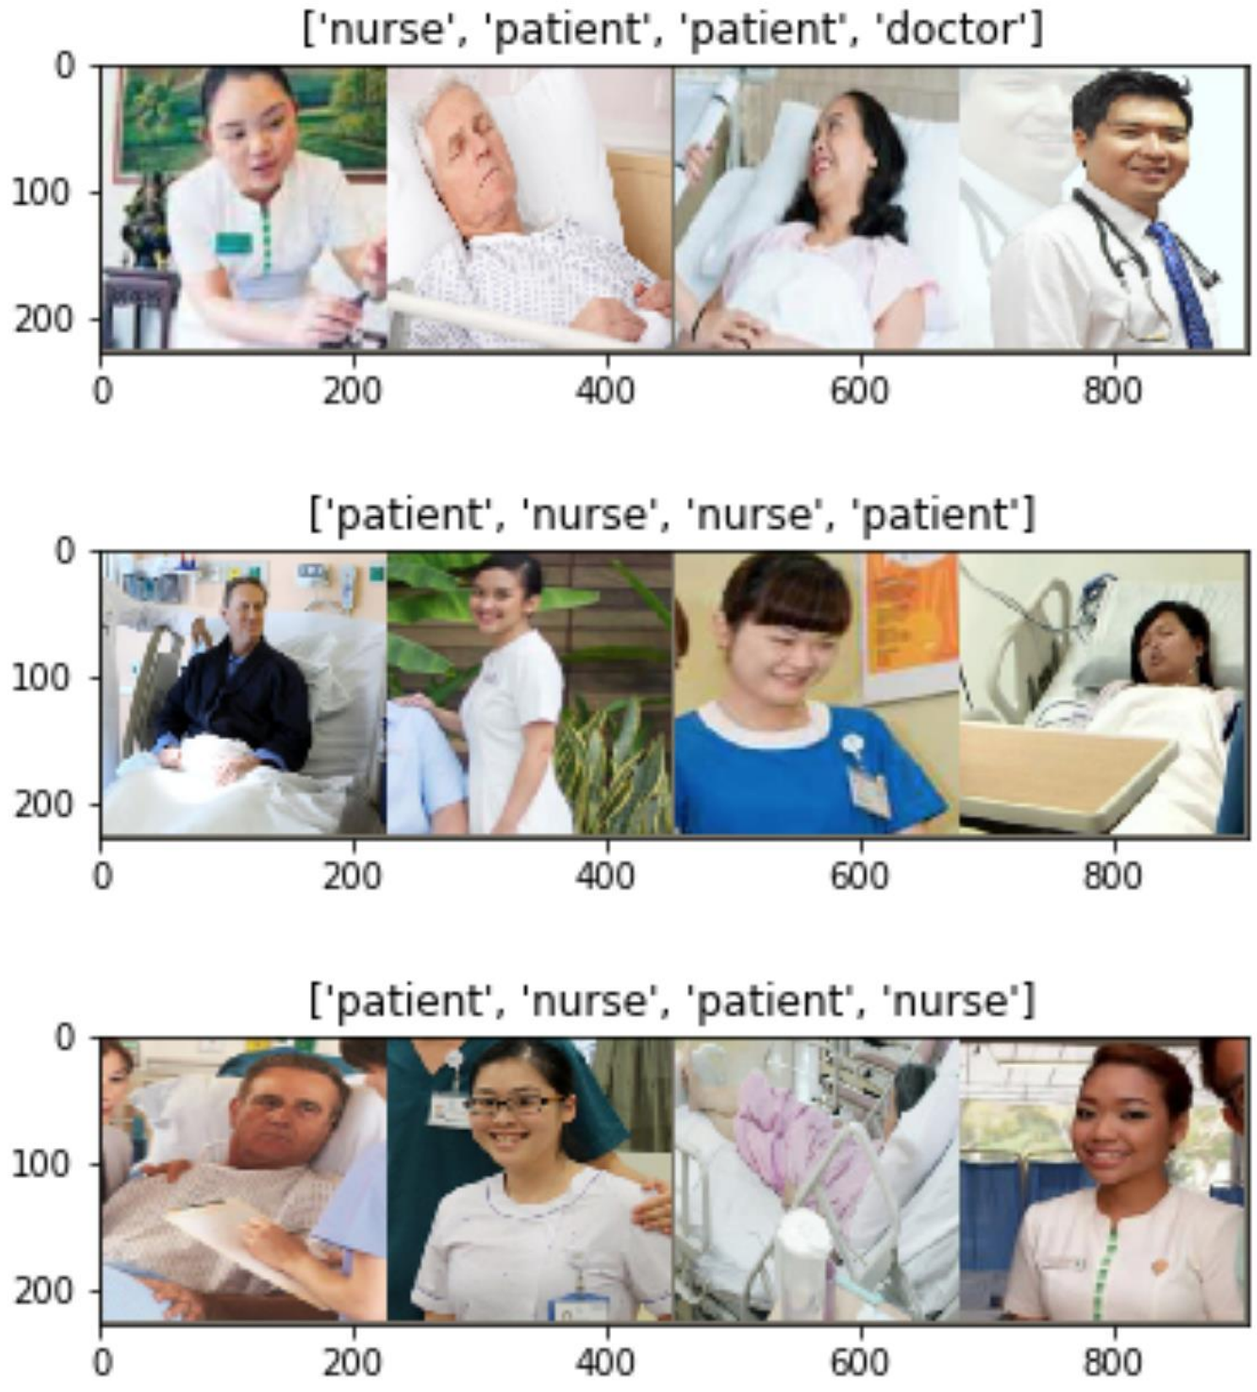

Getting Started: “choosing wisely”

- Data preparation
  - Considerations
    - Ethnic groups
    - Age
    - Uniforms

Doctor

Stethoscope

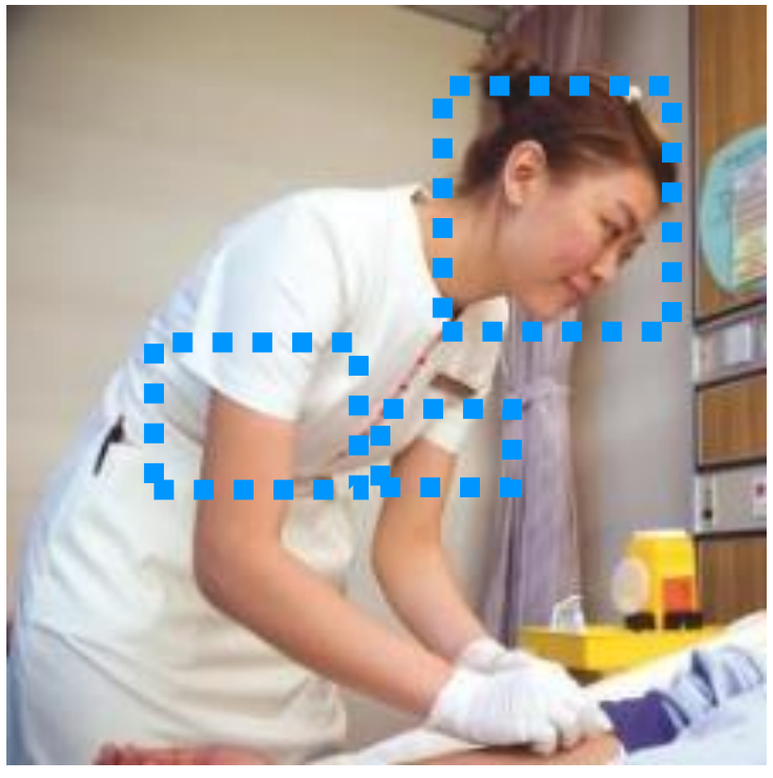

Nurse

Female, short-sleeves

Patient

Bed/Wheel-chair, hospital gown

Getting Started: “choosing wisely”

- Model architecture
  - Resnet18, freeze all layers
  - Optimize only last layer
- Initial Findings
  - Thought more Training data is better. Validation only 10 per class. Over 90% accuracy. Got skeptical…
  - Types of images consistently wrong
    - Side-views
    - Low-resolution
    - Multiple human subjects in same image
  - Tempting to remove such images…

Seeking feedback: “Taking it further”

- Limited data -> Data augmentation

```
data_transforms = {
    'train': transforms.Compose([transforms.Resize(224), transforms.RandomCrop(224), transforms
.RandomHorizontalFlip(), transforms.ToTensor(), transforms.Normalize([0.485, 0.456, 0.406], [0.
229, 0.224, 0.225])]),
    'val': transforms.Compose([transforms.Resize(224), transforms.CenterCrop(224), transforms.T
oTensor(), transforms.Normalize([0.485, 0.456, 0.406], [0.229, 0.224, 0.225])])
}

data_dir = 'data/HCW'

# Create dataset class
image_datasets = {x: datasets.ImageFolder(os.path.join(data_dir, x),data_transforms[x]) for x i
n ['train', 'val']}
dataloaders = {x: torch.utils.data.DataLoader(image_datasets[x], batch_size=4, shuffle=True, nu
m_workers=4) for x in ['train', 'val']}
```

Seeking feedback: “Taking it further”

- Plot accuracy, loss VS training epochs

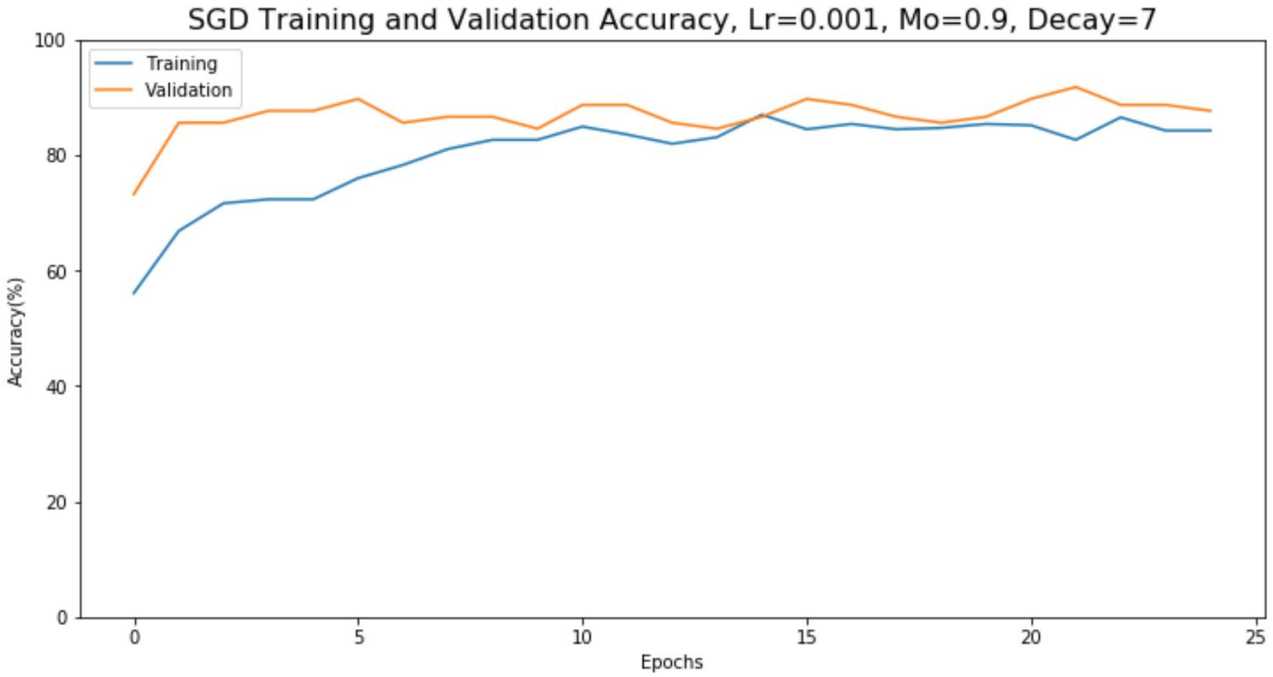

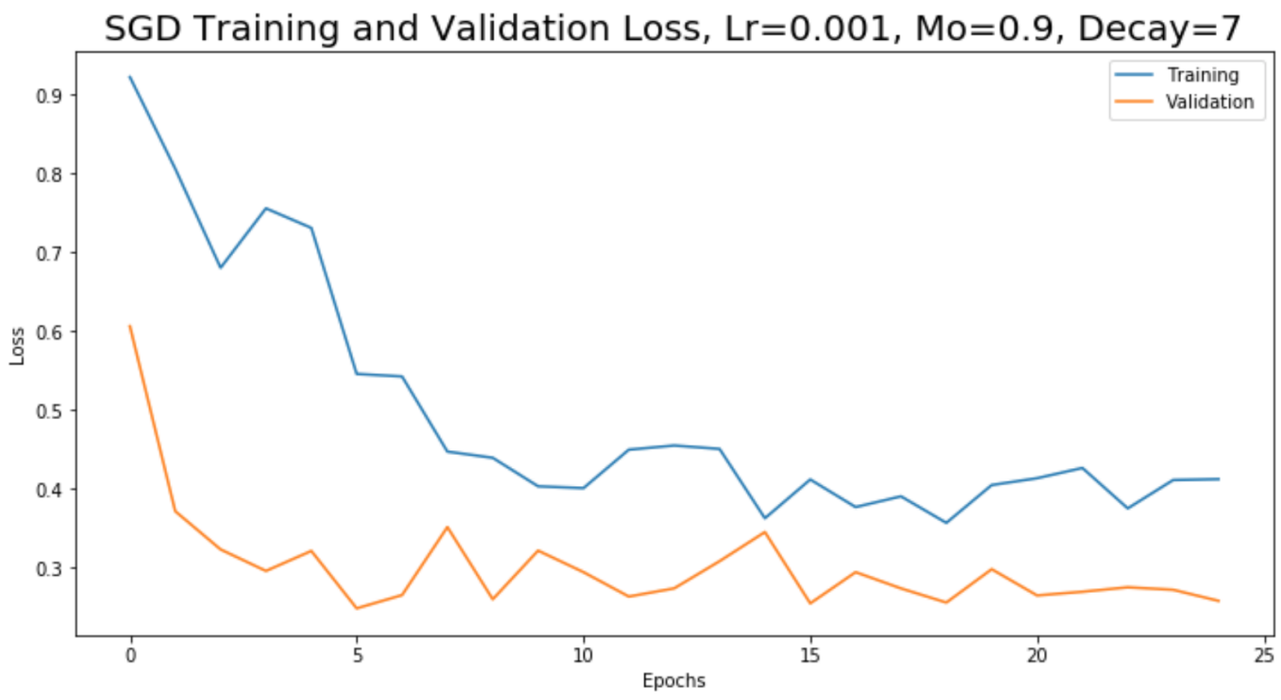

Seeking feedback: “Taking it further”

- Analysis of predictions (Can’t apply vanilla code blindly)
  - Patterns (“What is the model getting right/wrong?”)

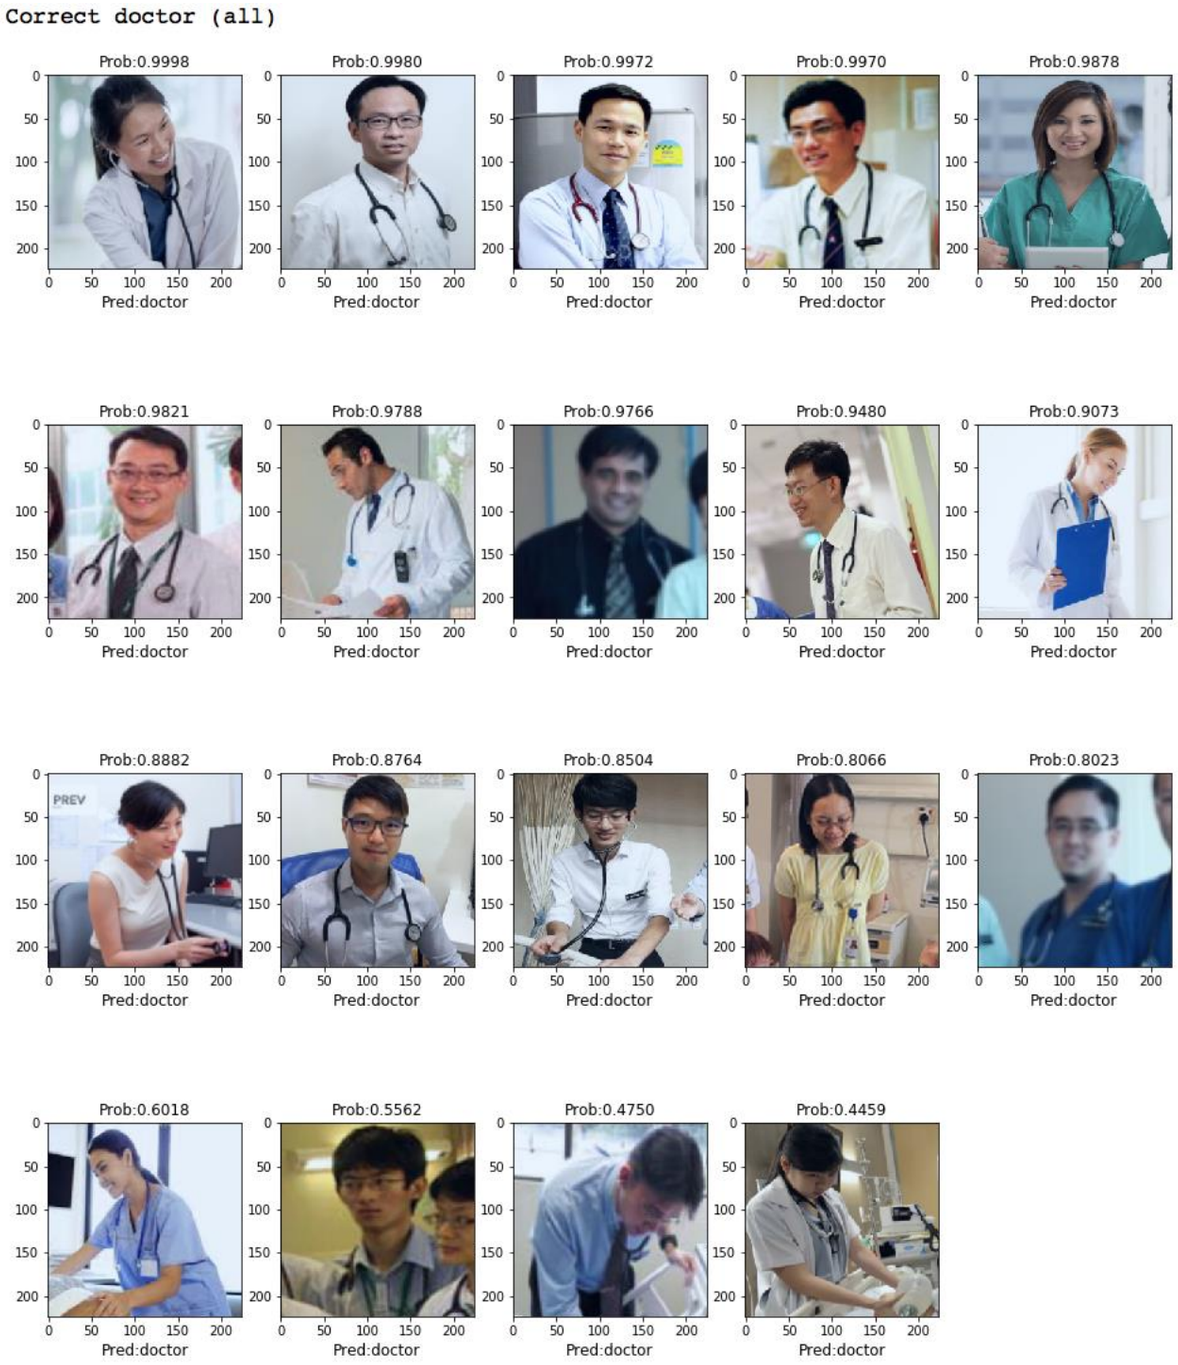

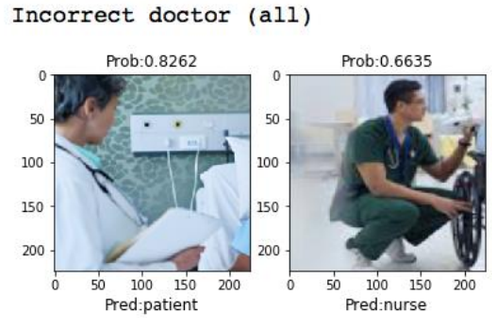

Seeking feedback: “Taking it further”

- Analysis of predictions
  - Patterns (“What is the model getting right/wrong?”)

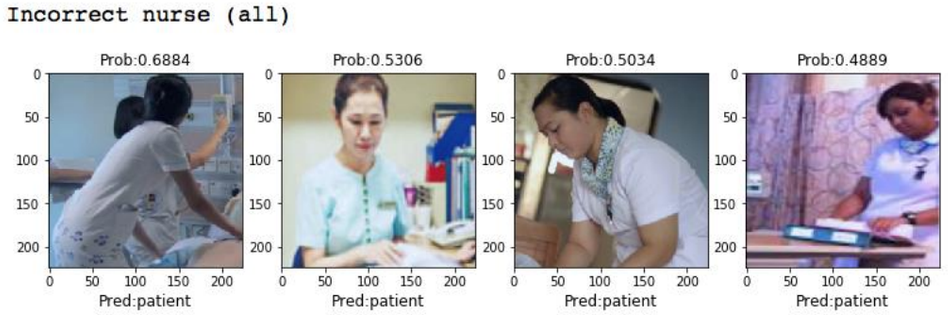

Seeking feedback: “Taking it further”

- Analysis of predictions
  - Patterns (“What is the model getting right/wrong?”)

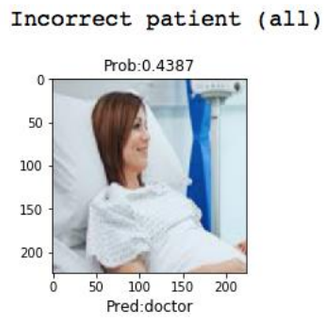

Seeking feedback: “Taking it further”

- Analysis of predictions
  - Class performance (“Which class should I focus on improving?”)

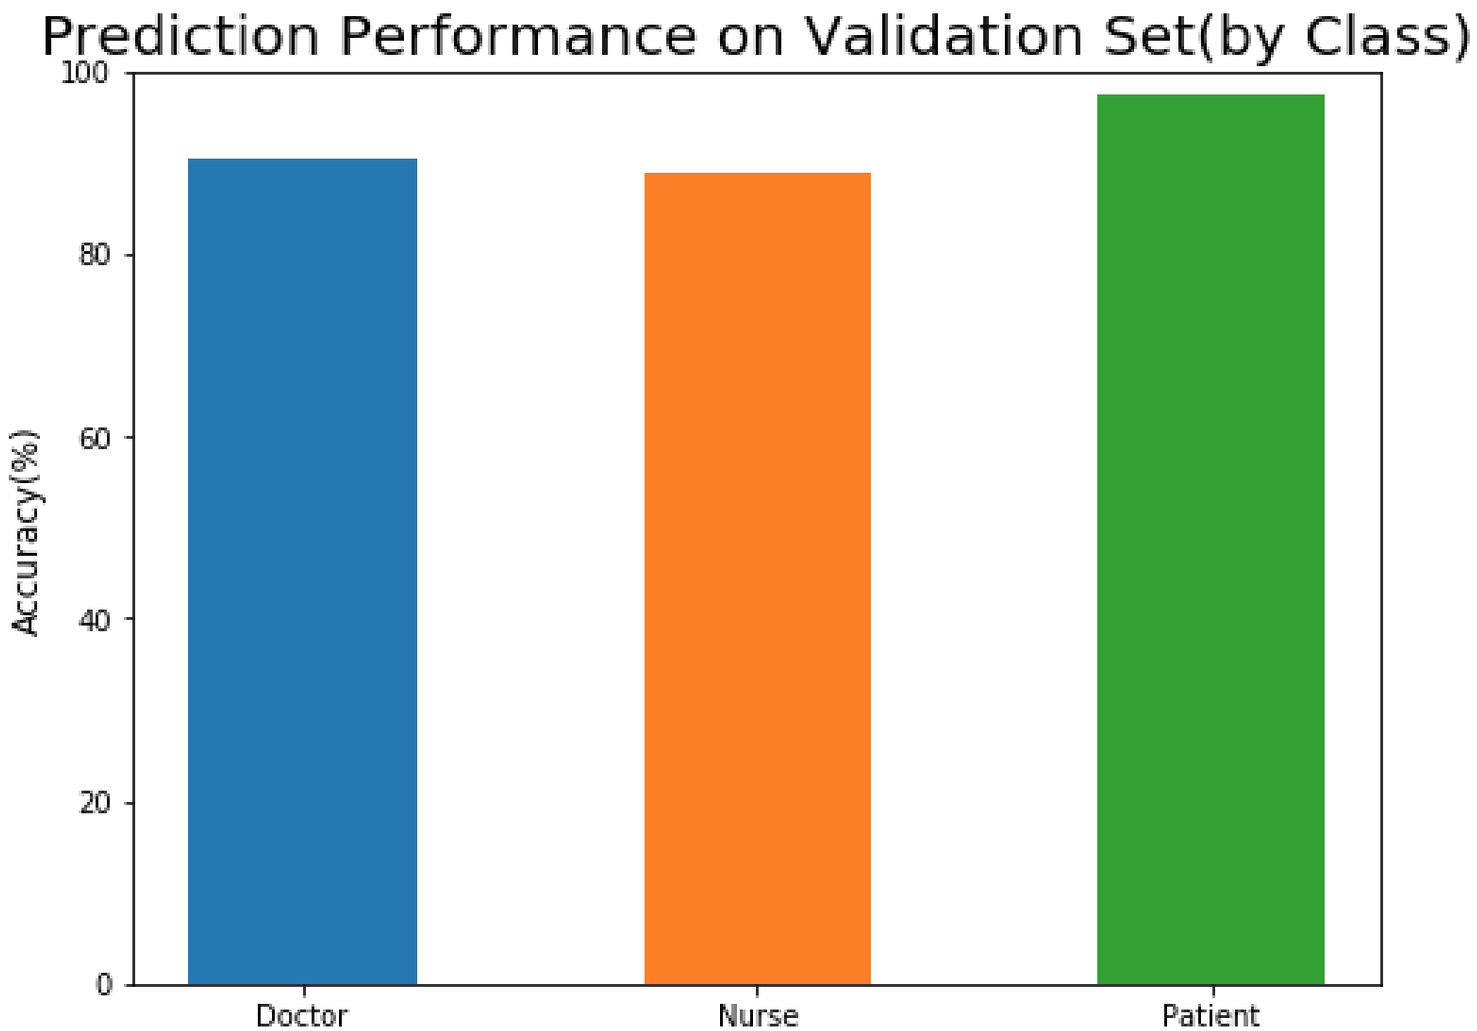

Seeking feedback: “Taking it further”

- Analysis of predictions
  - Patterns (“What is the model getting right/wrong?”)

Low probability

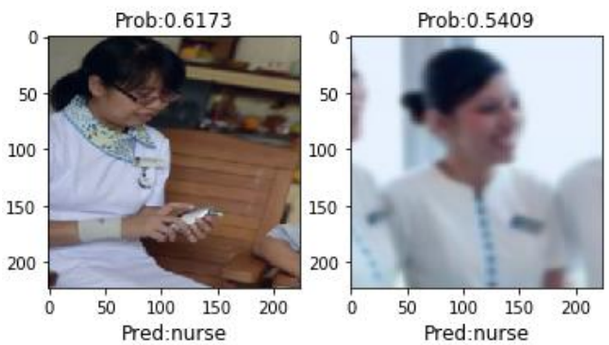

Edge cases

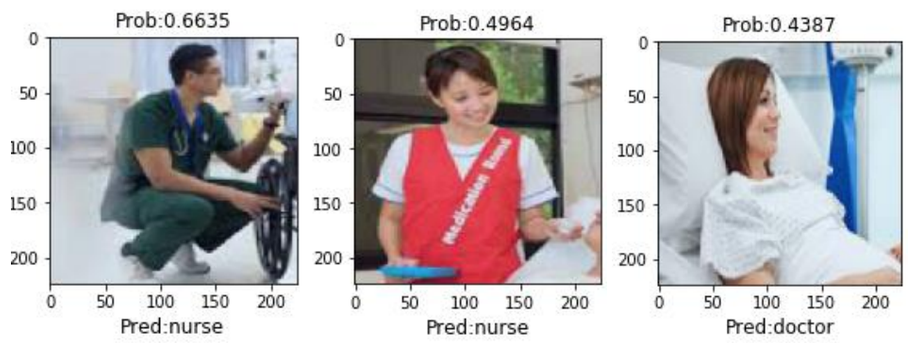

Seeking feedback: “Taking it further”

- Visualisation of activations on final layer
  (“How did the model makes its predictions?”)
  - SmoothGrad technique (implementation in PyTorch)

Predicted: Doctor

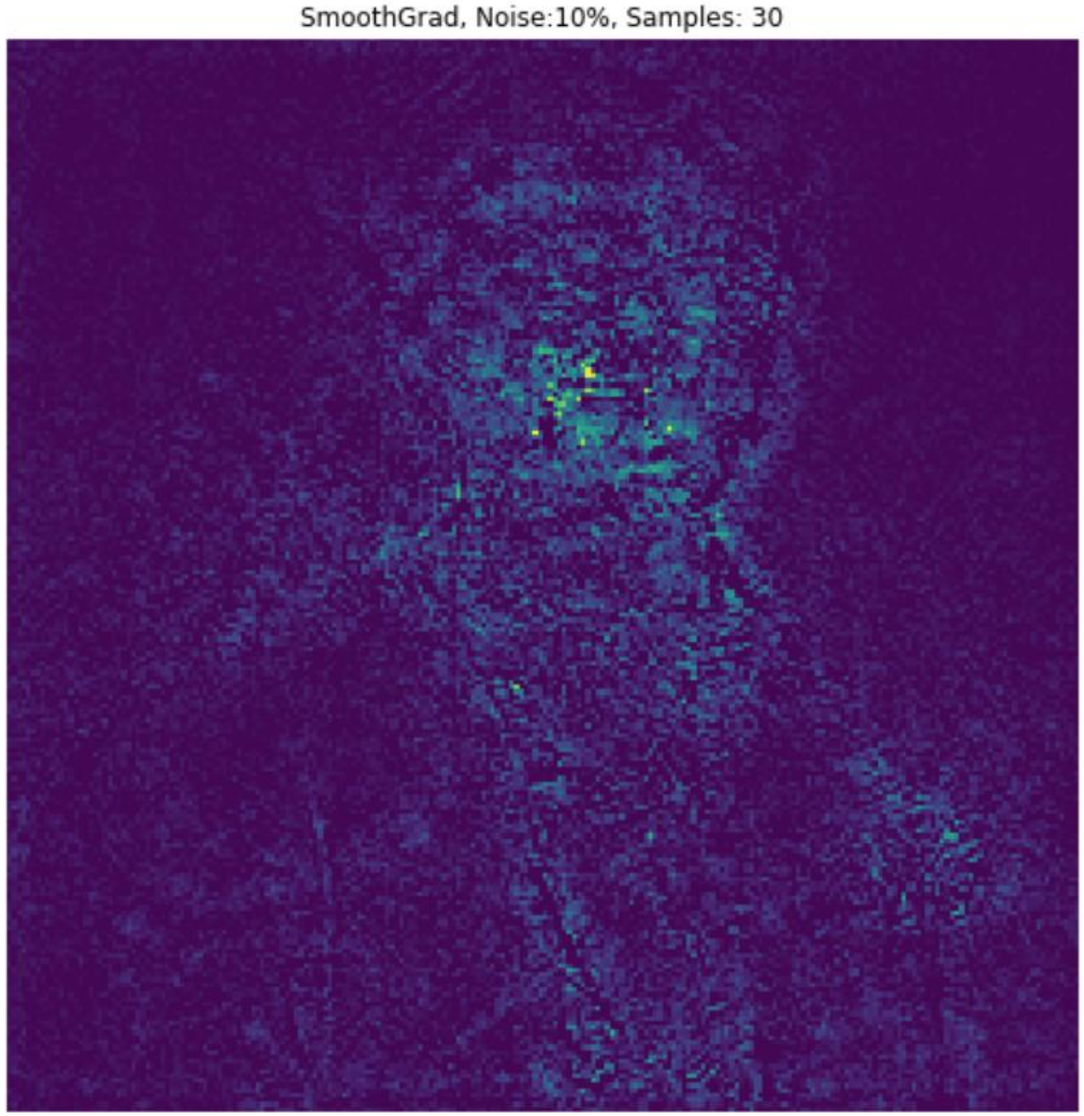

Seeking feedback: “Taking it further”

- Visualisation of activations on final layer
  (“How did the model makes its predictions?”)

Predicted: Doctor

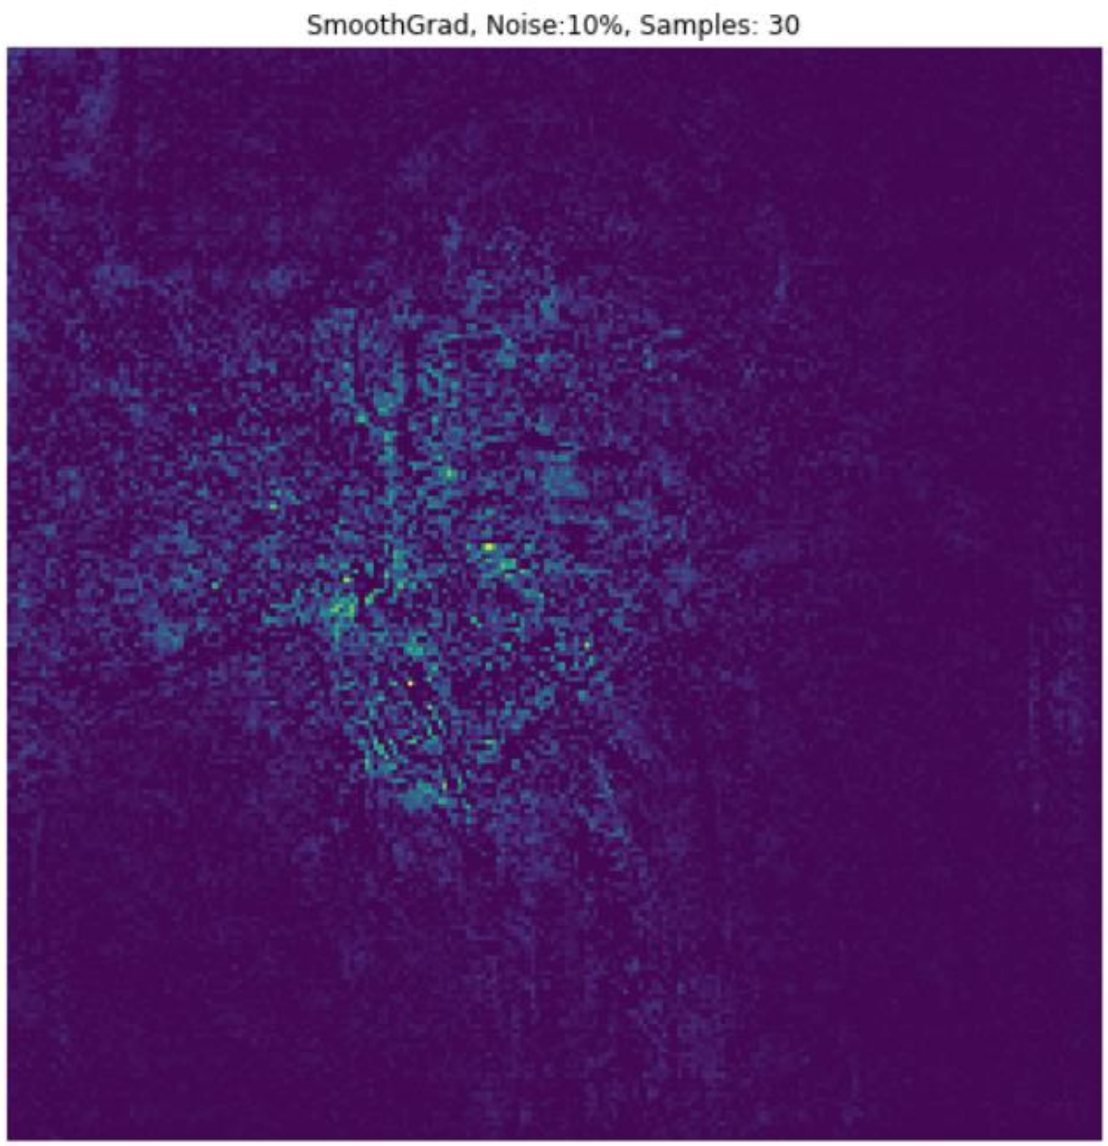

Seeking feedback: “Taking it further”

- Visualisation of activations on final layer
  (“How did the model makes its predictions?”)

Predicted: Nurse

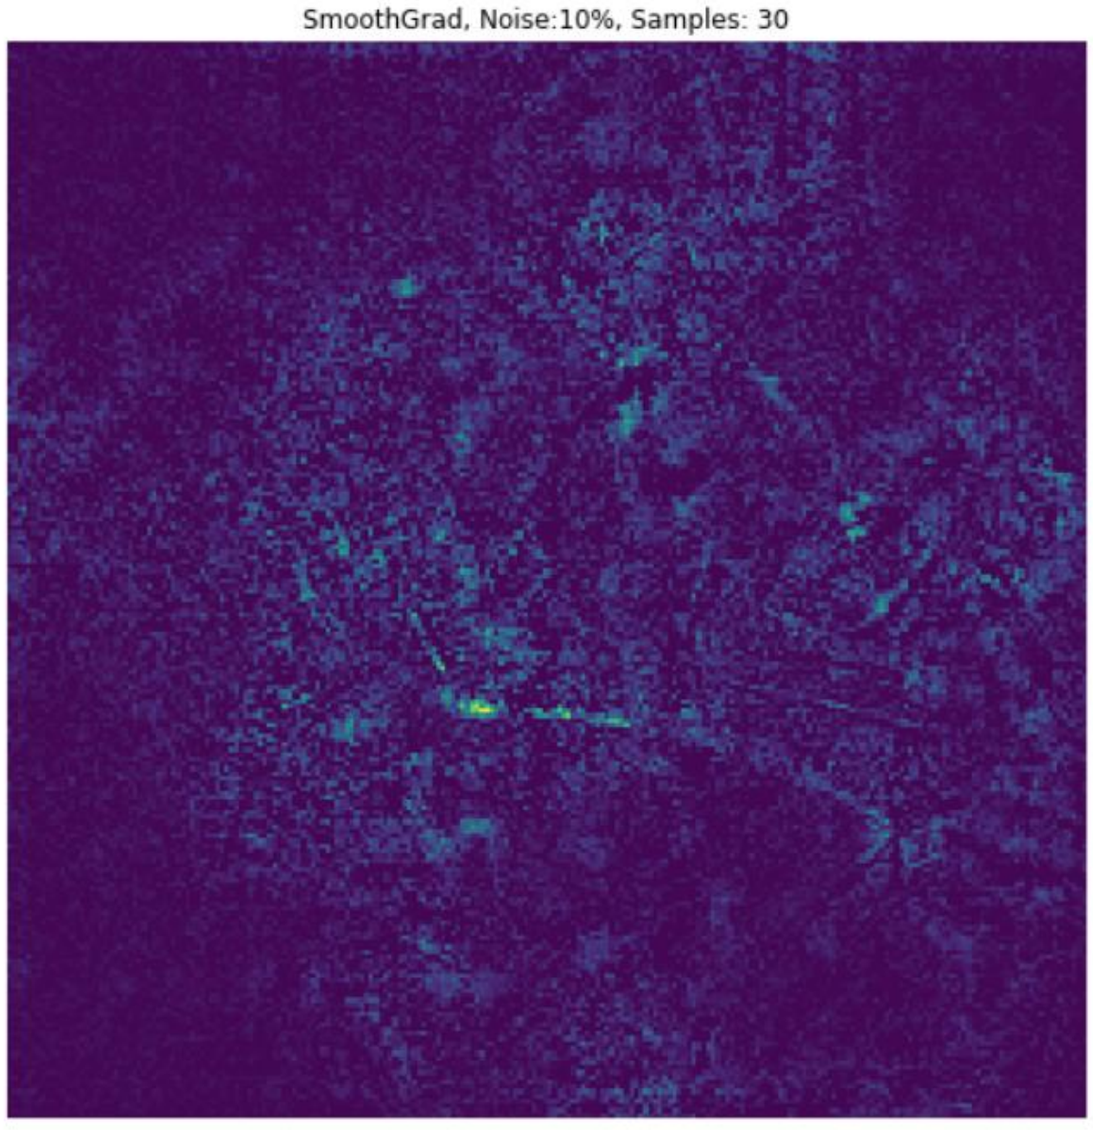

Seeking feedback: “Taking it further”

- Visualisation of activations on final layer
  (“How did the model makes its predictions?”)

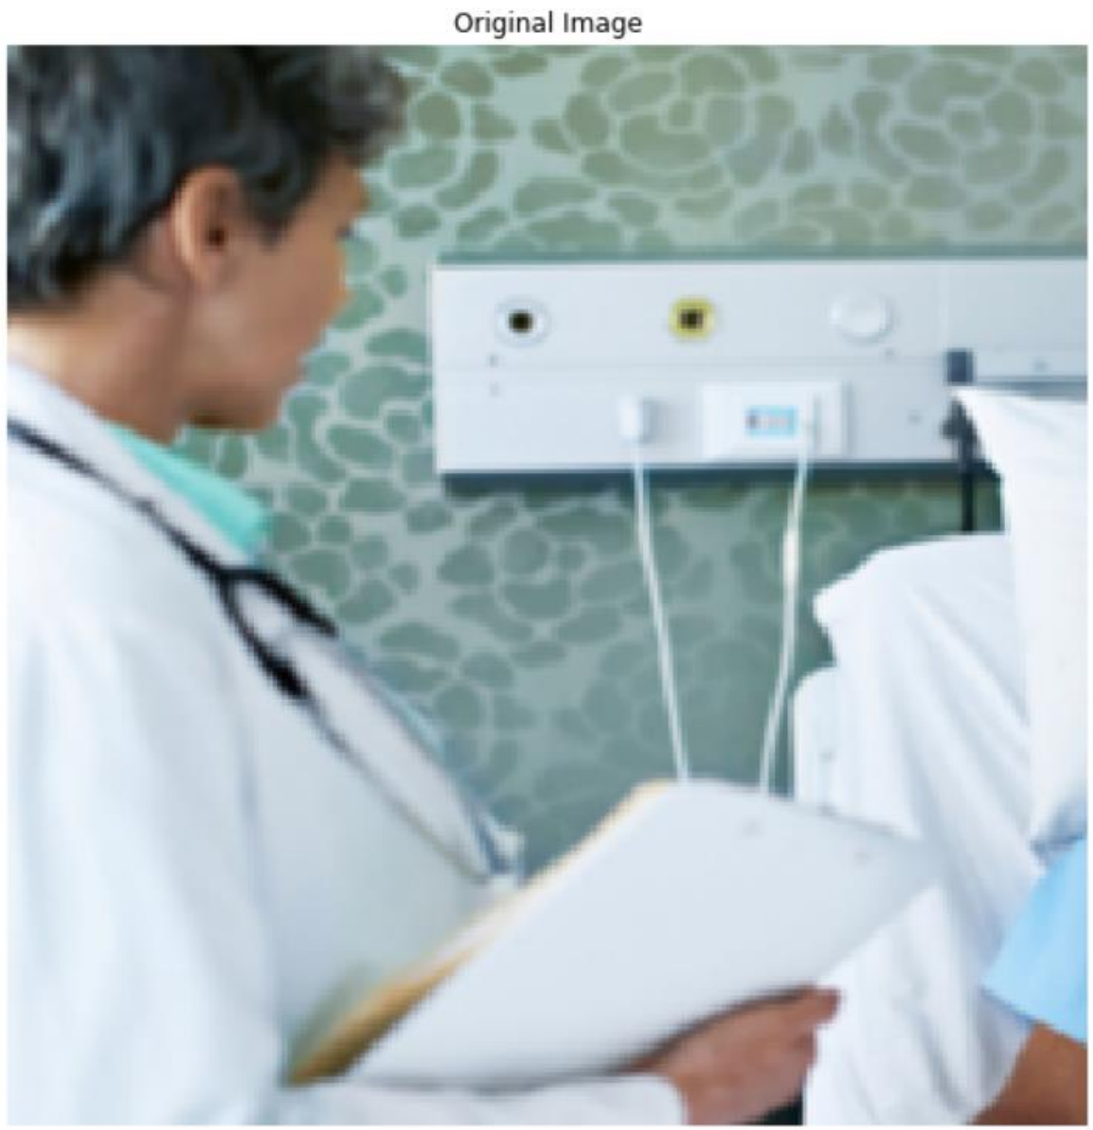

Predicted: Patient

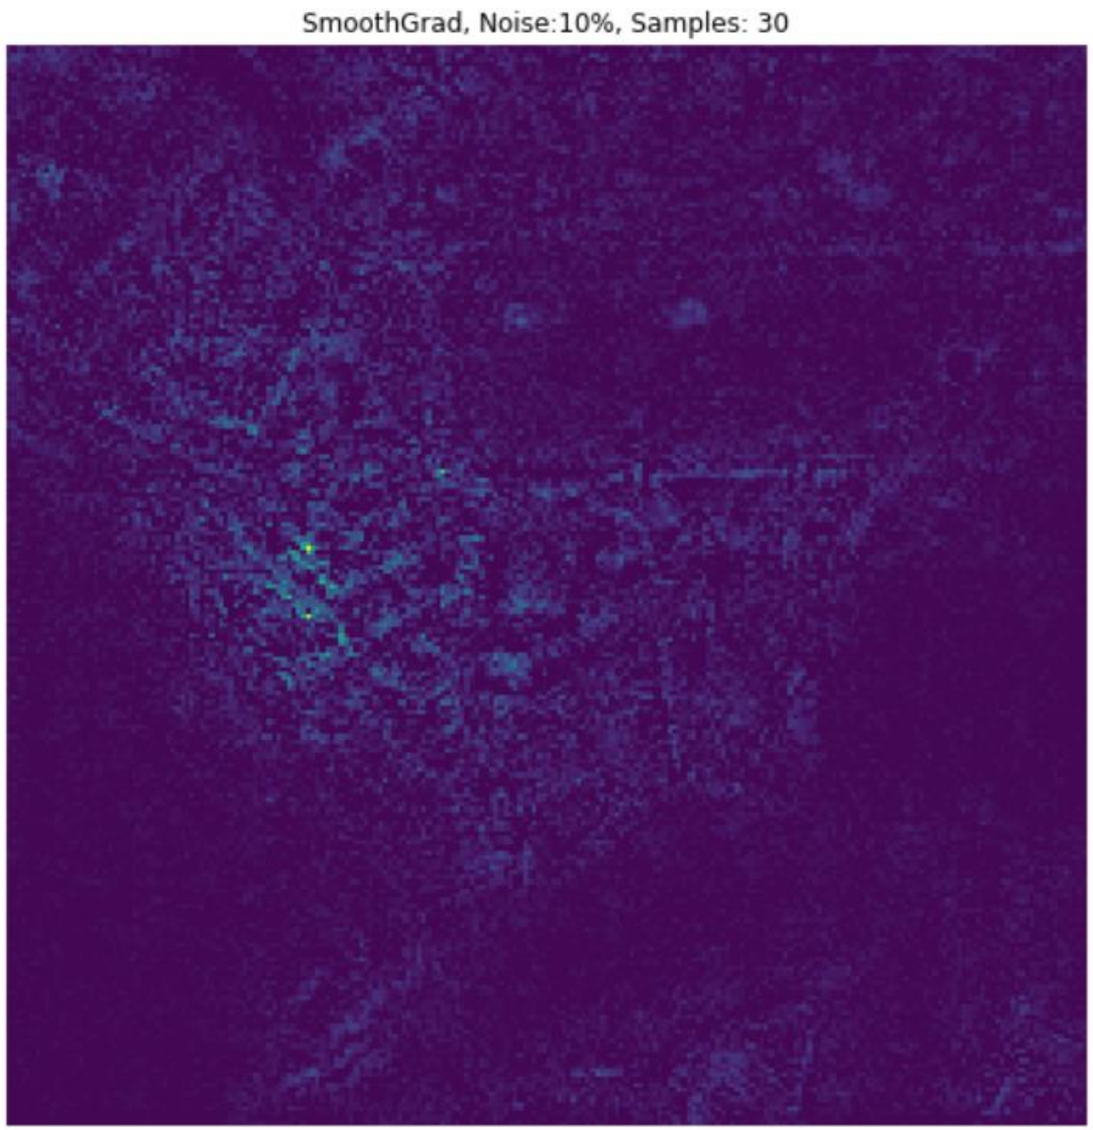

Seeking feedback: “Taking it further”

- Visualisation of activations on final layer
  (“How did the model makes its predictions?”)

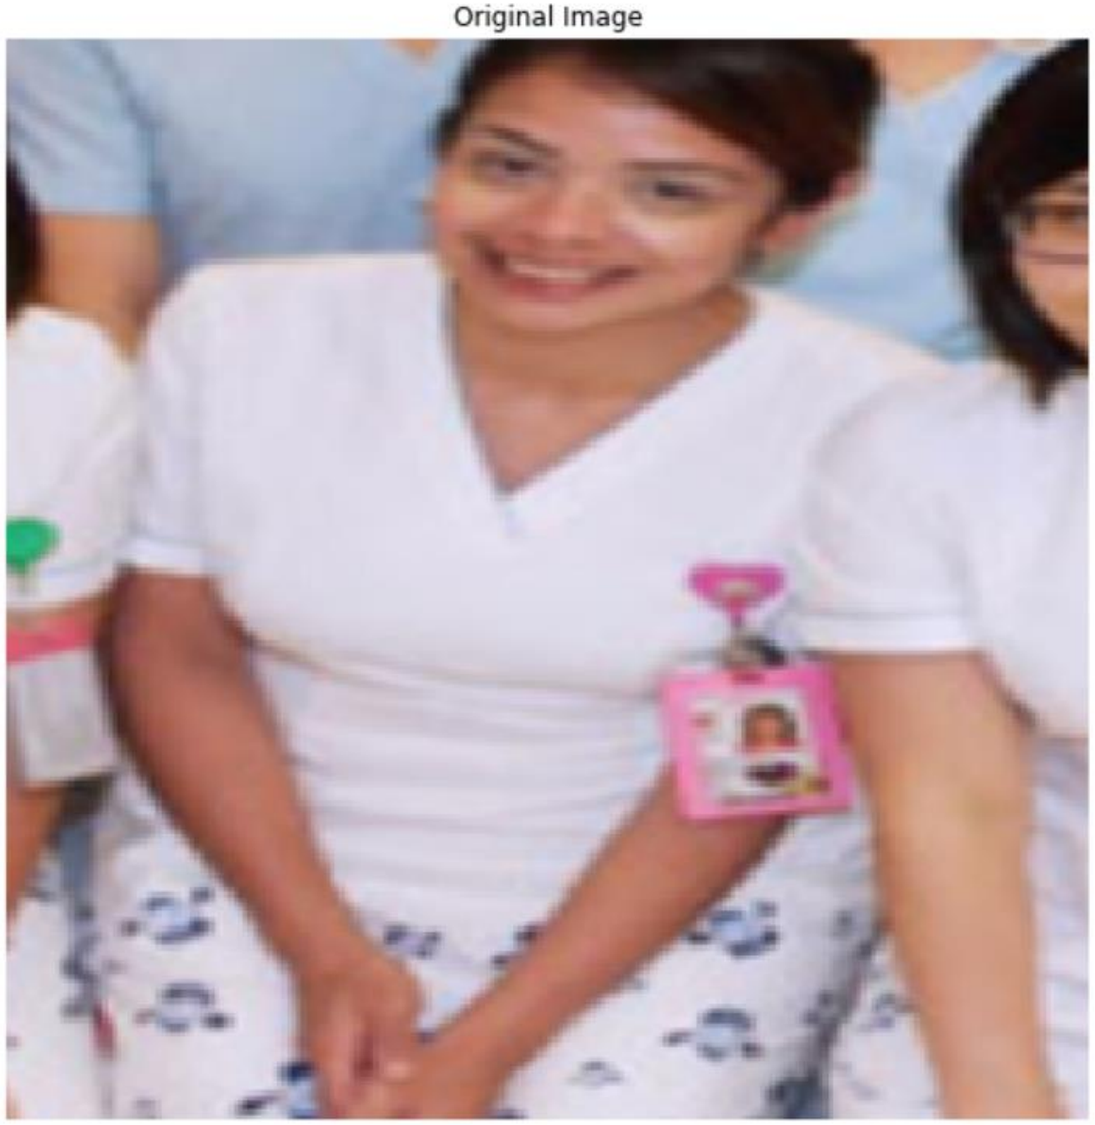

Predicted: Nurse

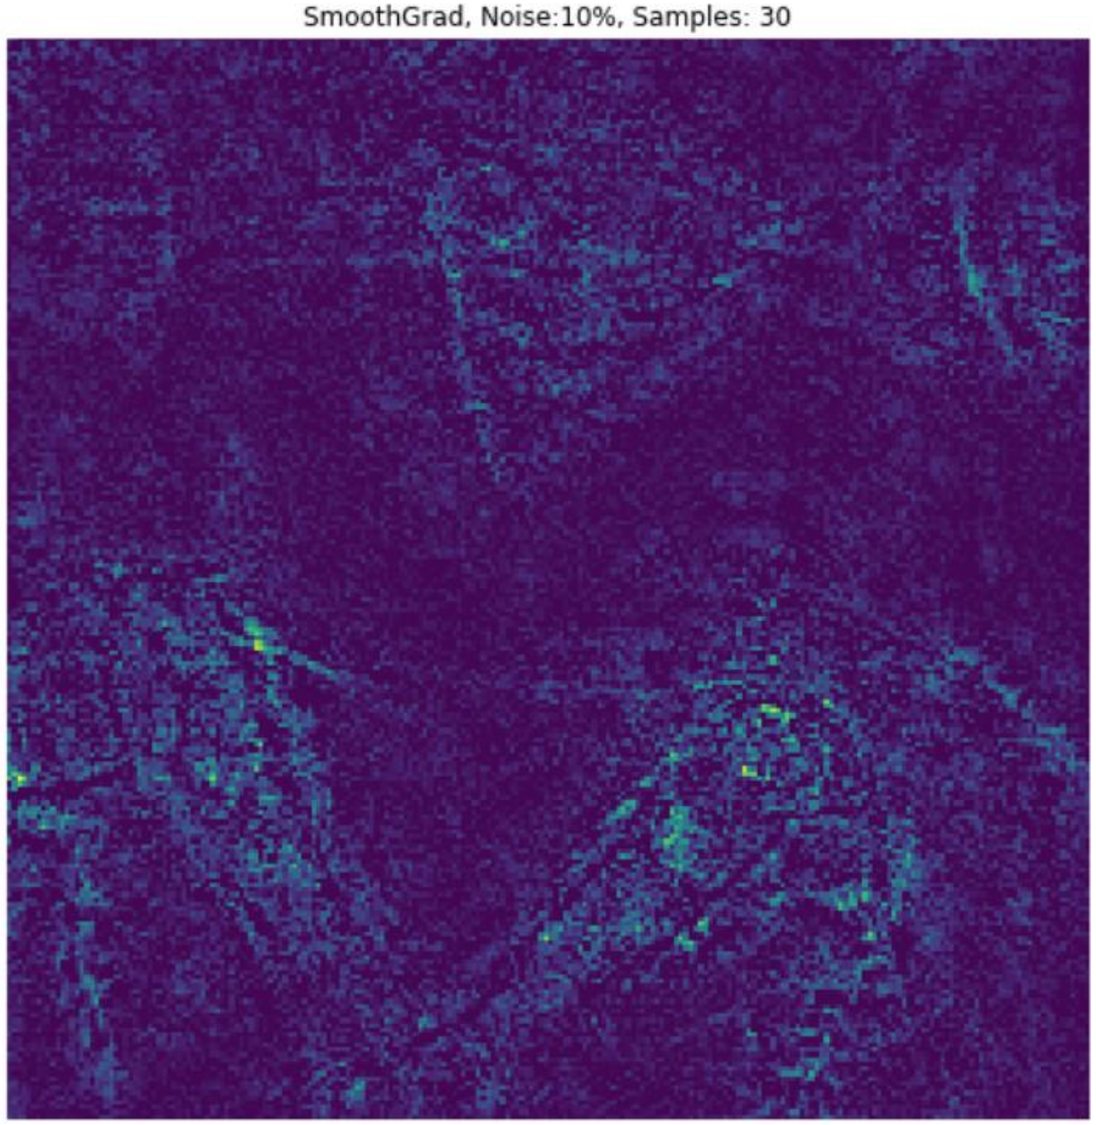

Seeking feedback: “Taking it further”

- Visualisation of activations on final layer
  (“How did the model makes its predictions?”)

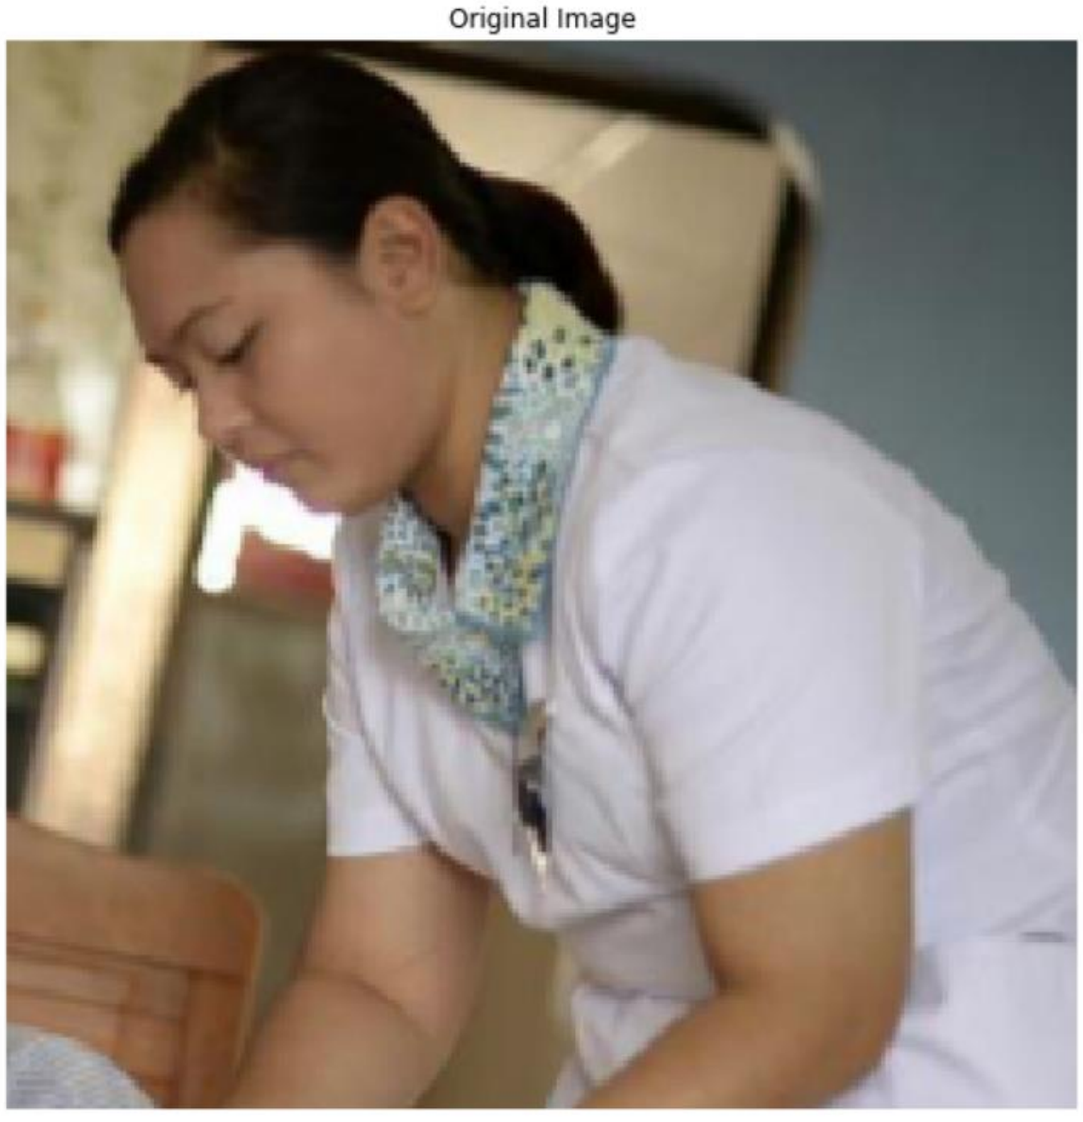

Predicted: Patient

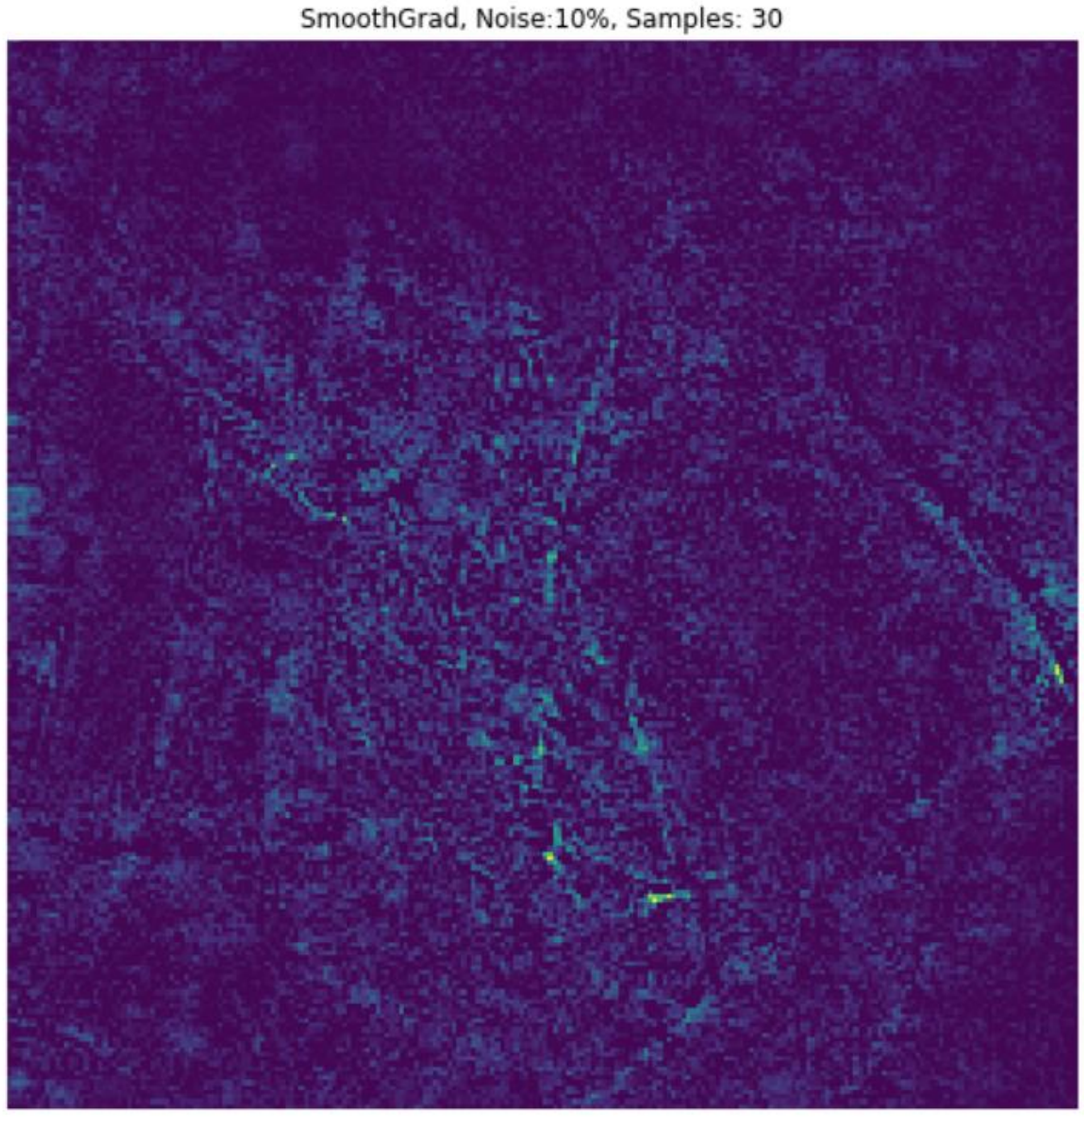

Seeking feedback: “Taking it further”

- Visualisation of activations on final layer
  (“How did the model makes its predictions?”)

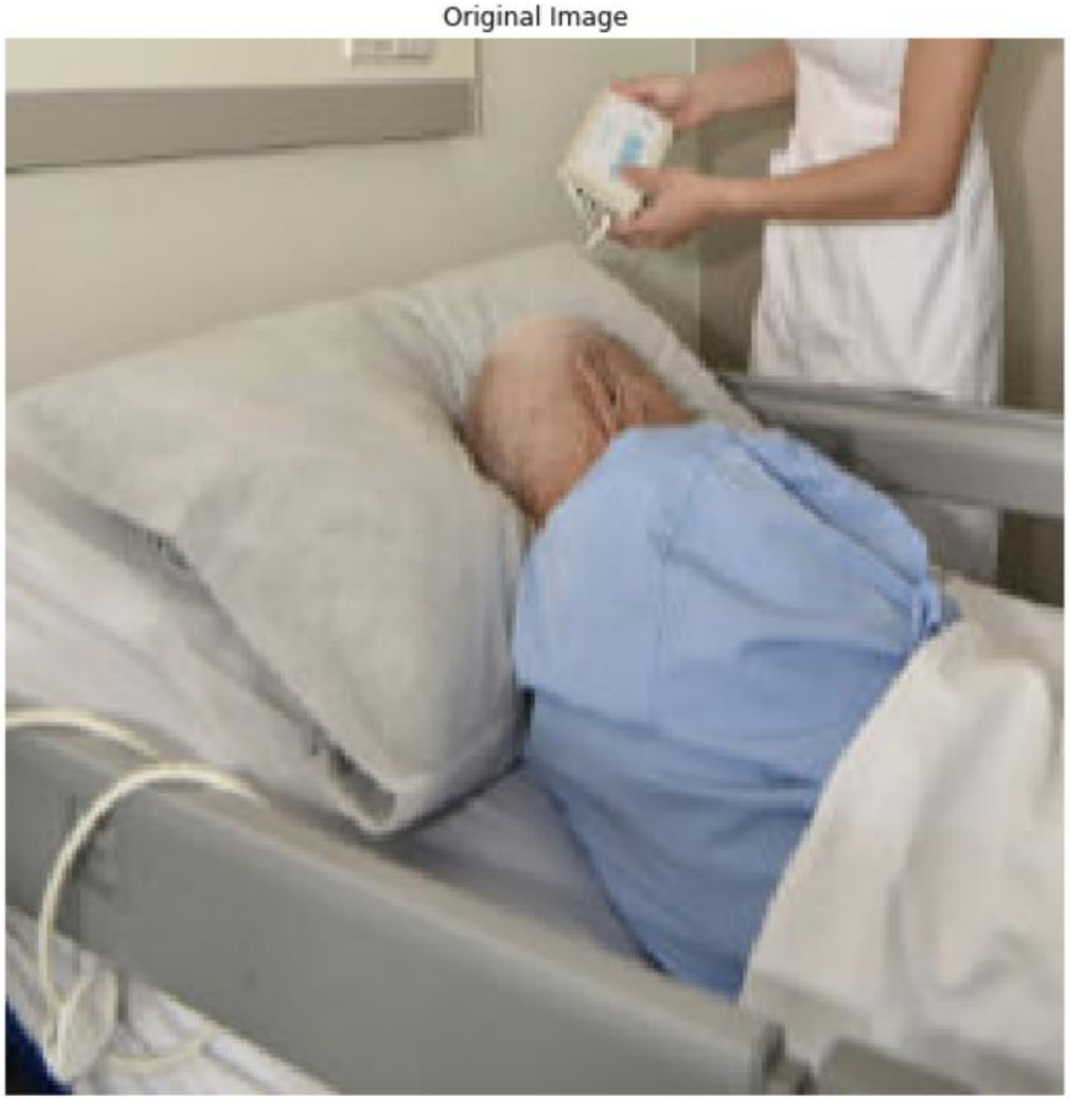

Predicted: Patient

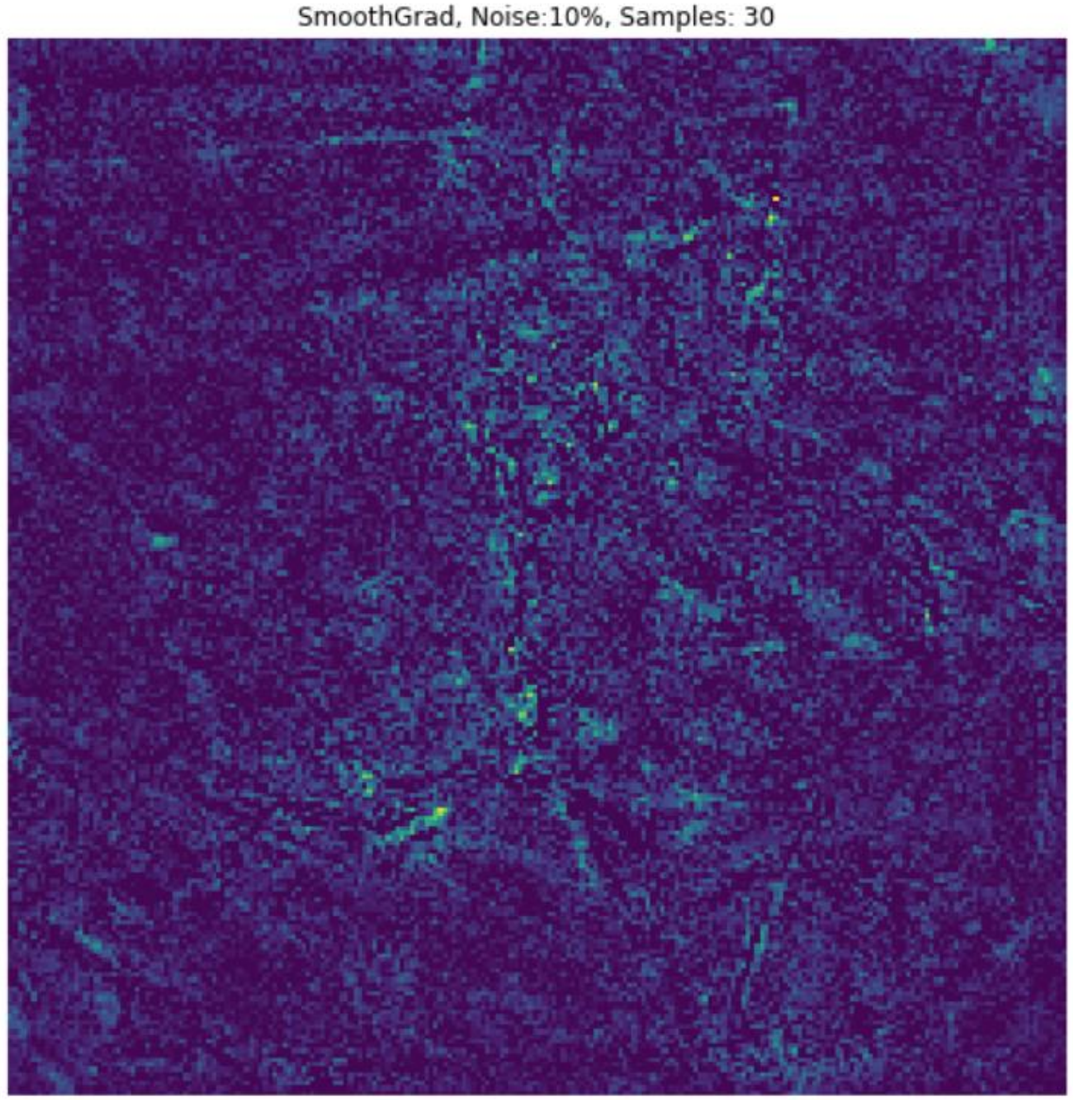

Seeking feedback: “Taking it further”

- Visualisation of activations on final layer
  (“How did the model makes its predictions?”)

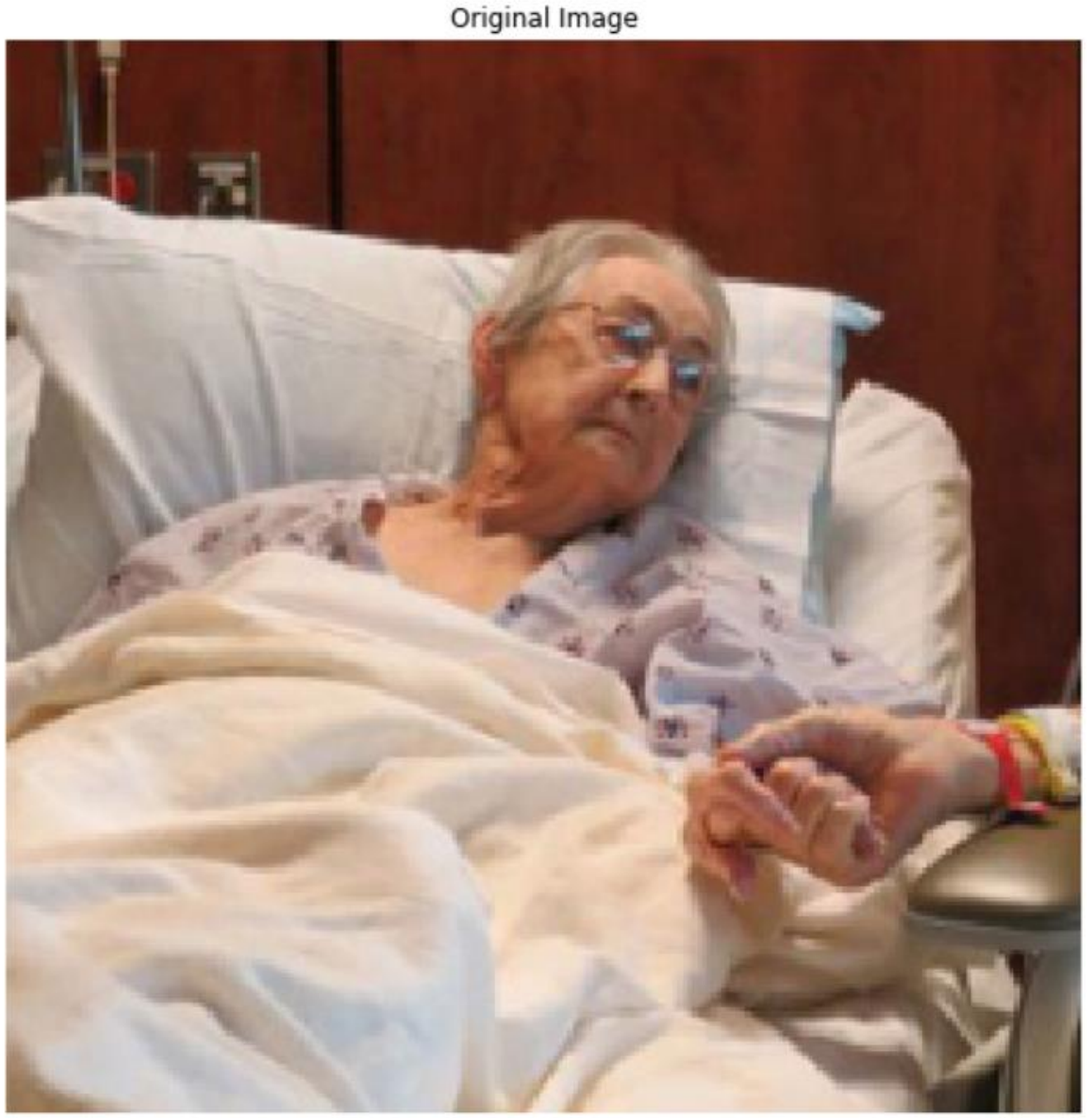

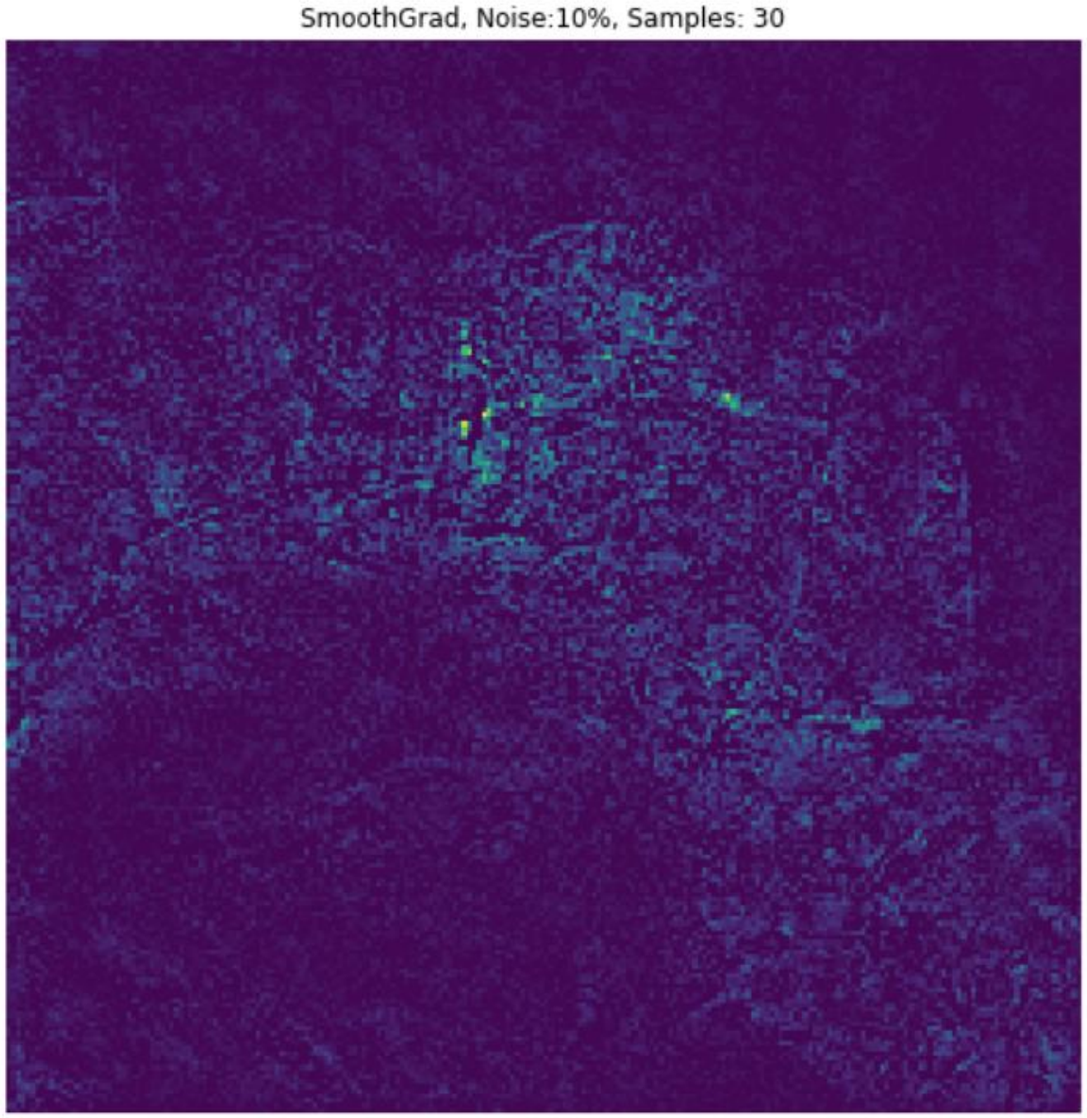

Predicted: Patient

Seeking feedback: “Taking it further”

- Visualisation of activations on final layer
  (“How did the model makes its predictions?”)

Original Image

Predicted: Doctor

SmoothGrad, Noise:10%, Samples: 30

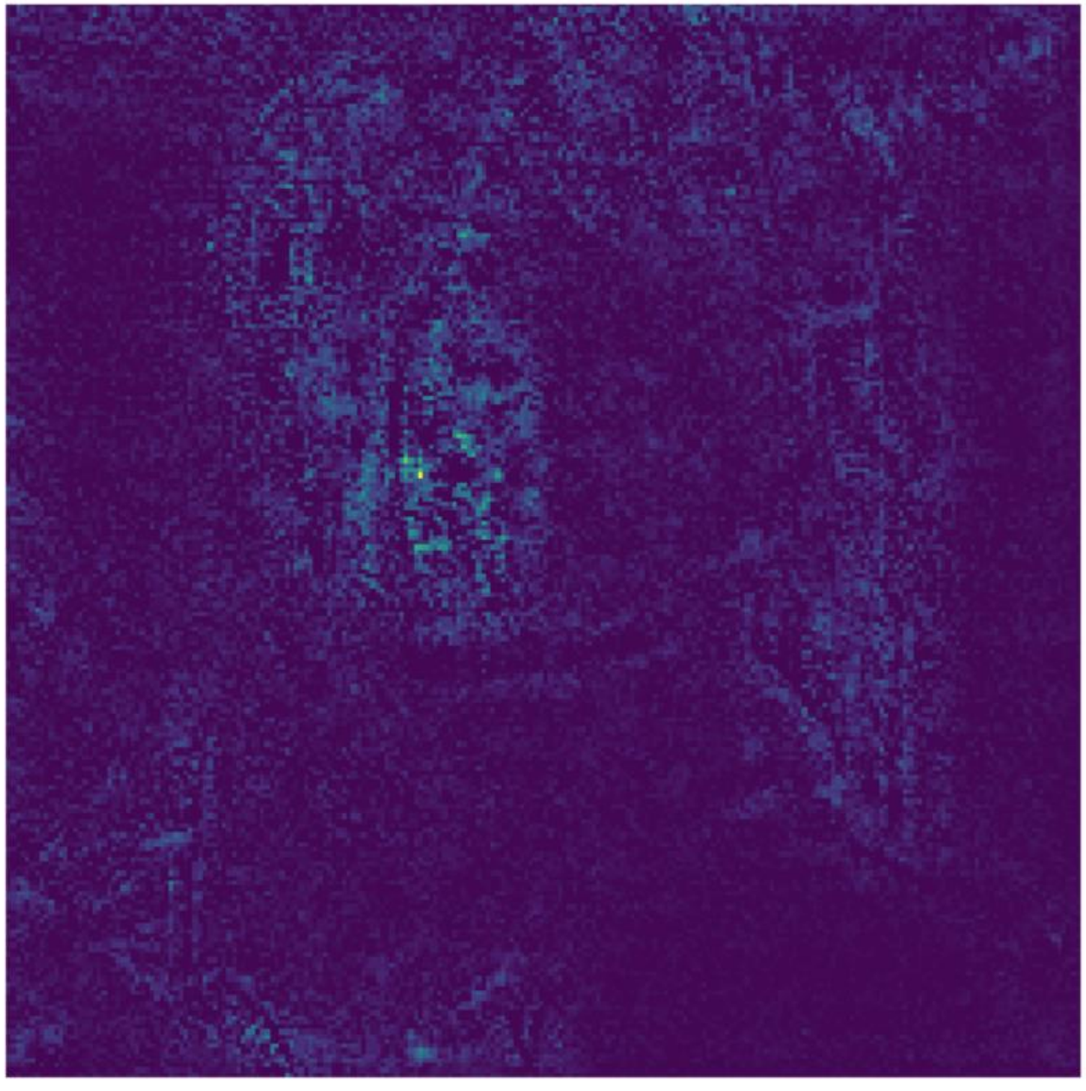

Seeking feedback: “Taking it further”

- Visualisation of activations on final layer
  (“How did the model makes its predictions?”)

Summary

Doctor

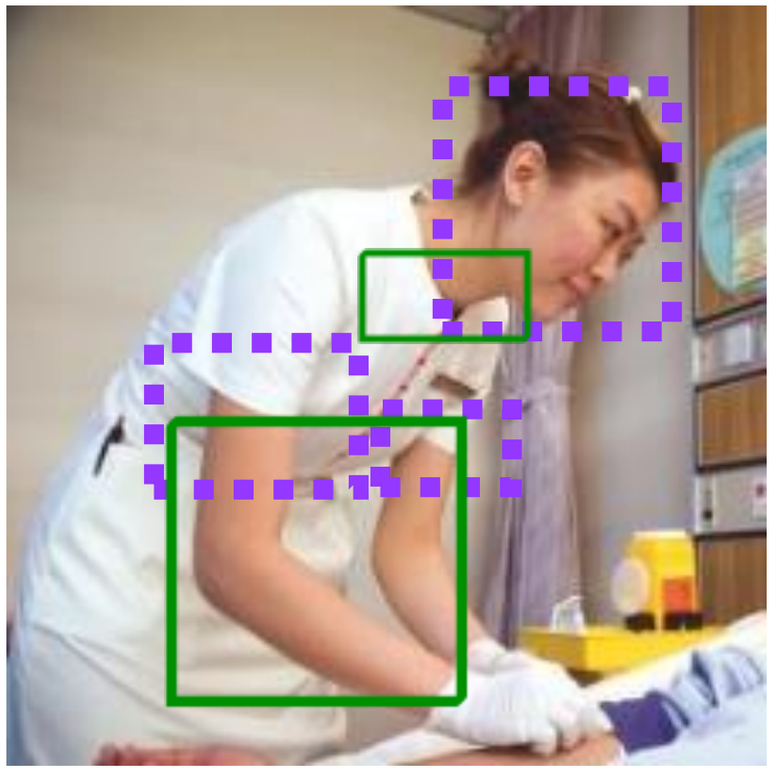

Nurse

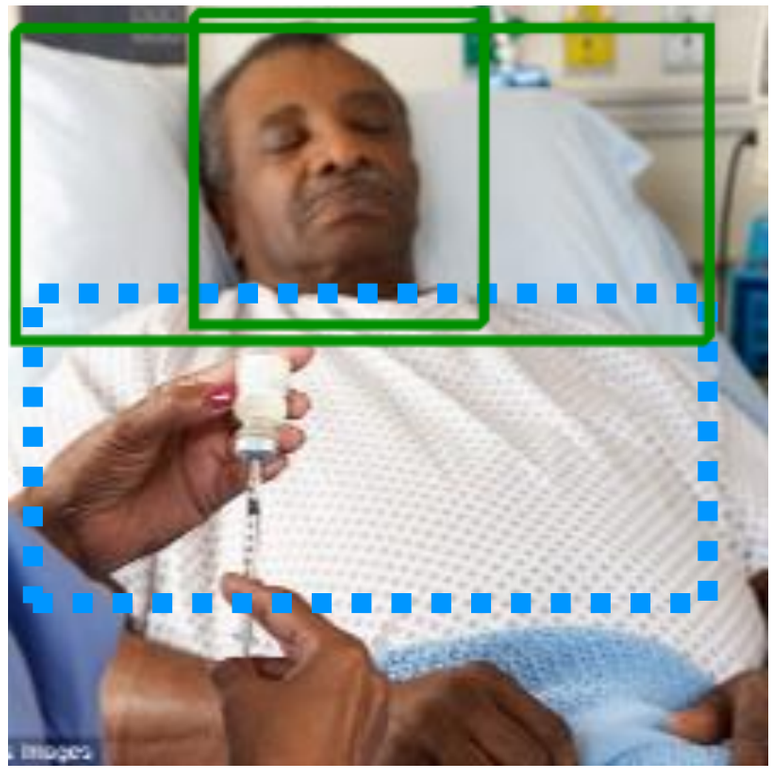

Patient

Legend:

- Model
- Model match Hypothesis
- Hypothesis

Improving the model: “Good to Great”

- Train model with more images of side and back views
- Tuning the hyper-parameters
  - Reduce learning rate from 0.001 to 0.0008, all else constant
    - Highest acc 91.67% as compared to baseline 91.75%
  - Increase training epoch from 25 to 50
    - Highest acc 93.75% as compared to baseline 91.75%
- Computing power from NSCC a great help
  - Short amount of time available
  - Experimenting with tuning various other hyper-parameters
  - Visualisation of last layer quite computational expensive

Predictions on test Set: “Taking the Leap of Faith”

- Test set selection considerations
  - Not the same photos/person from training or validation set
  - Hand-picked to be closer to actual ground conditions
- Final results = 100% accuracy

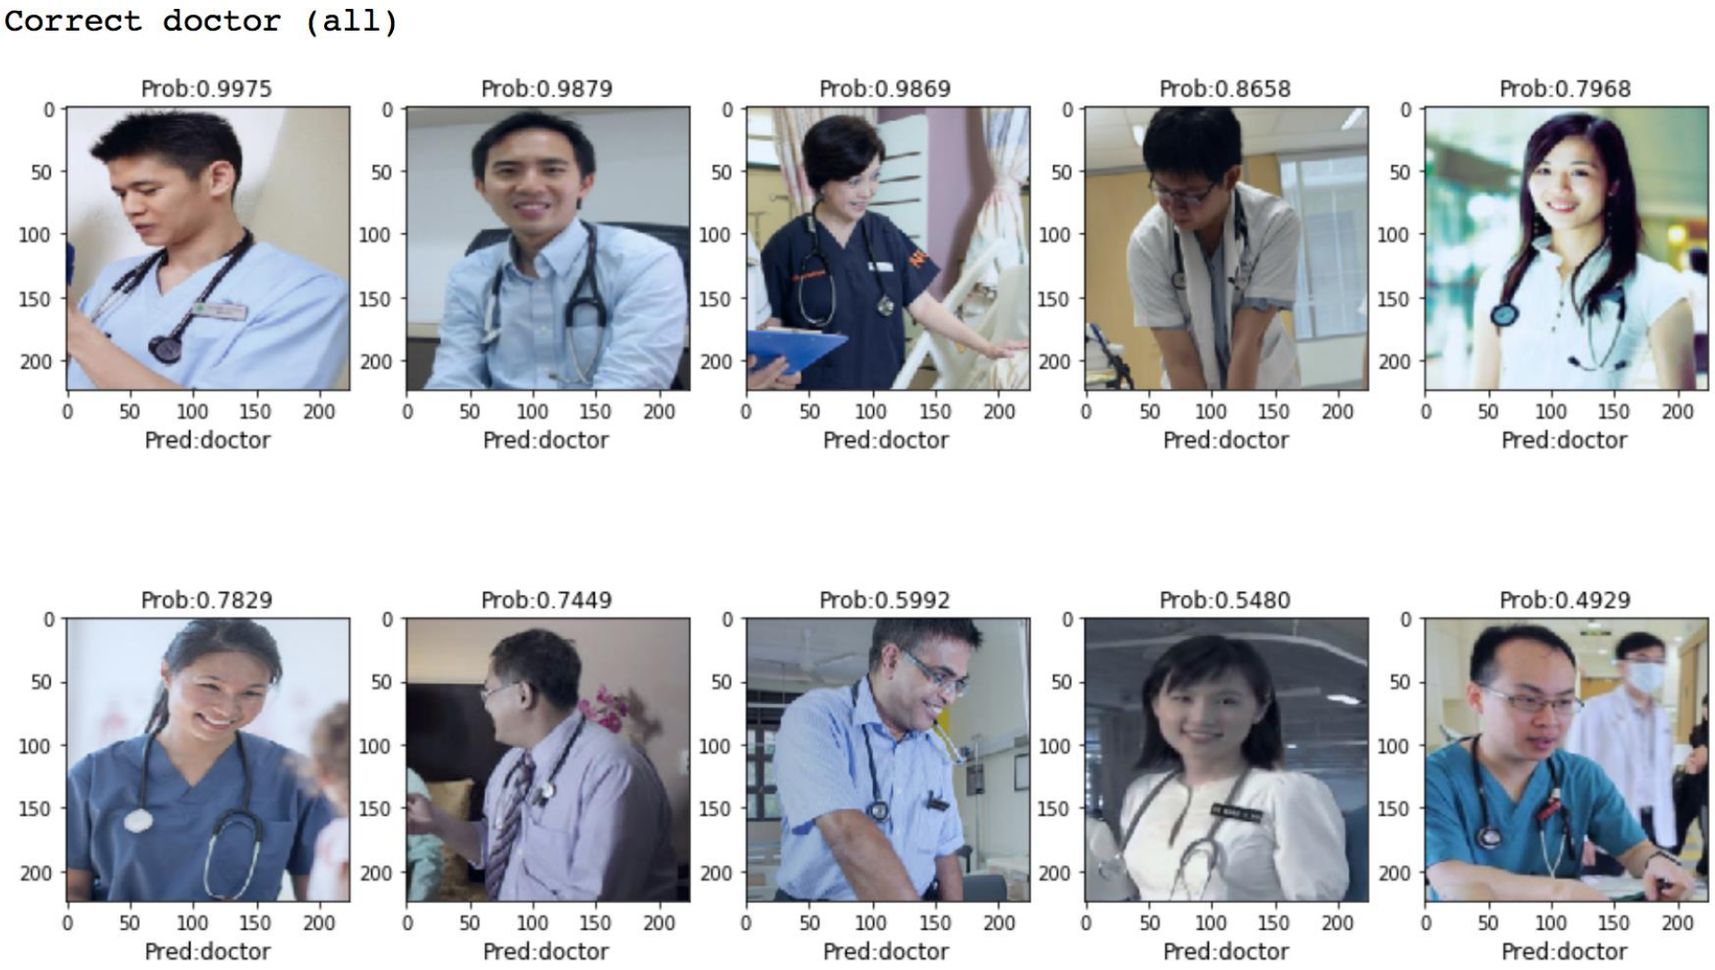

Predictions on test Set: “Taking the Leap of Faith”

Correct nurse (all)

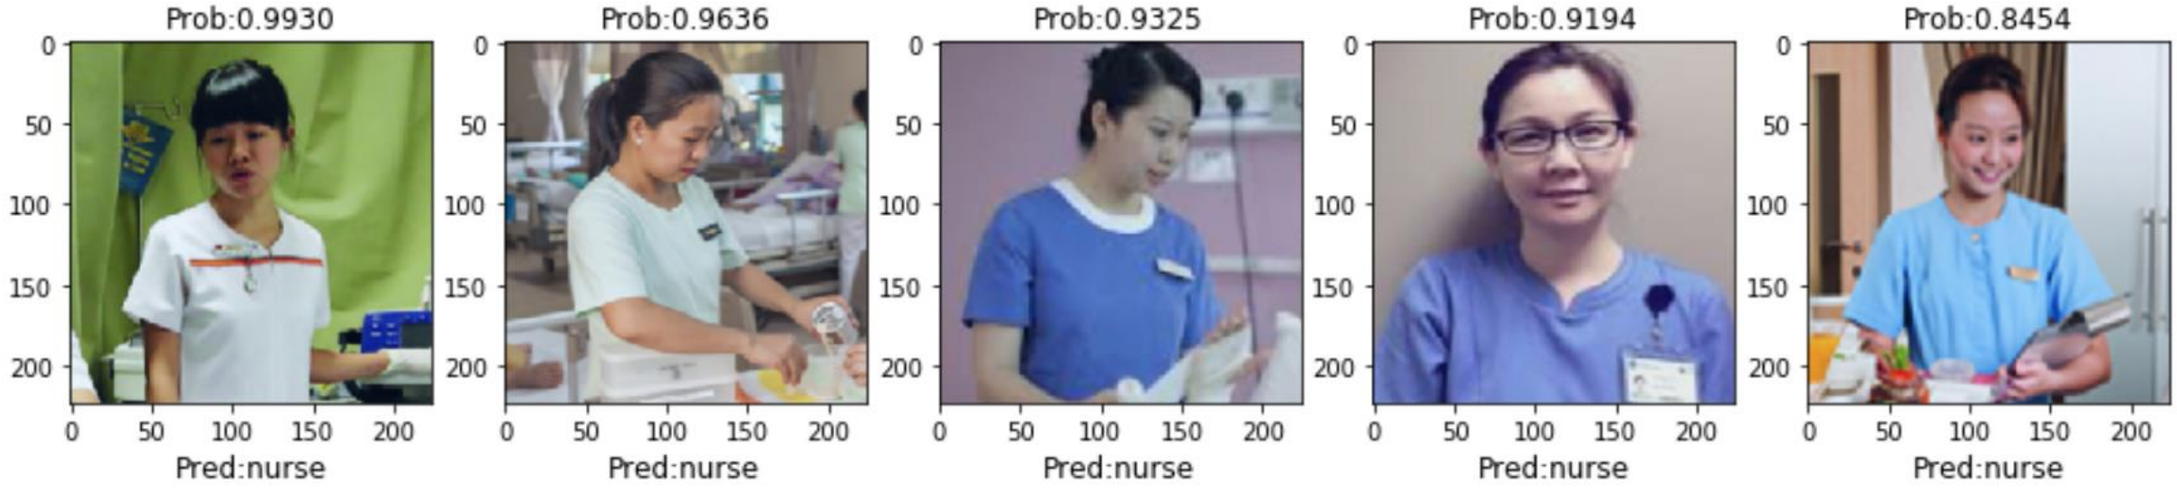

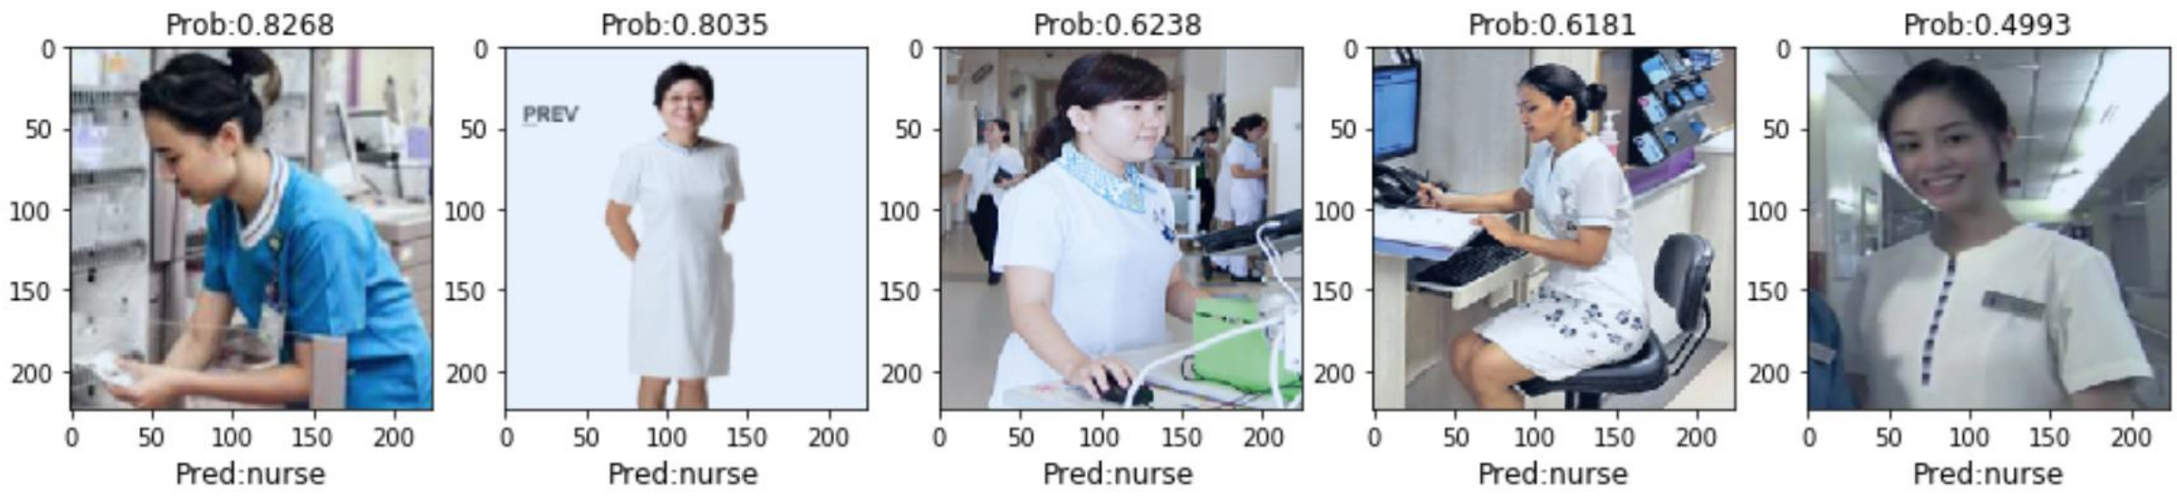

Predictions on test Set: “Taking the Leap of Faith”

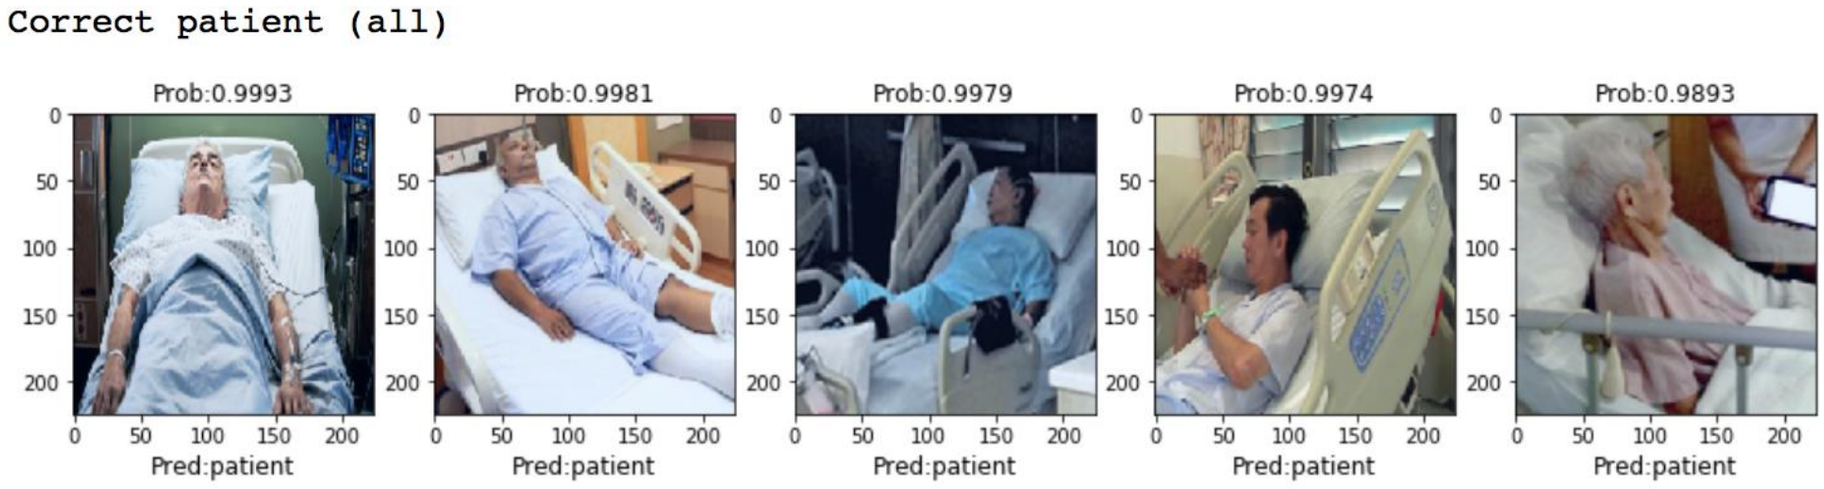

Next Steps: “This is just the beginning of a new chapter”

- Patient monitoring/ Community care - Falls, sleep, vitals, motion/activity (rehab)
- Presence of visitors with patients
- Drink/food intake of patients
- Hand Hygiene

- Future work
  - Training of model with more side and back profiles of doctors and nurses
  - Leveraging on YOLO to automate cropping of images
  - Transfer learning on other pre-trained networks
  - Manual adjustment of weights at final connected layer (Eg. more emphasise on characteristic features such as stethoscope)

Thank you

www.linkedin.com/in/kelvinthamkh

https://github.com/Oracle1983/Ai6-Challenge-2018.git