SAMSUNG SDS

Foresee

Techtonic 2021

Partner

Disrupt

Efficiently and Effectively developing

Medical Imaging AI Systems

Patrick Bangert, VP

The next great AI Frontier is Healthcare

Medical Imaging …

- 5.5 billion medical images are taken every year – analyzed slowly by doctors with a 70% chance of being right.

… powered by AI

- Each image can be diagnosed objectively, instantly, with a 99% chance of being right and at marginal cost per patient!

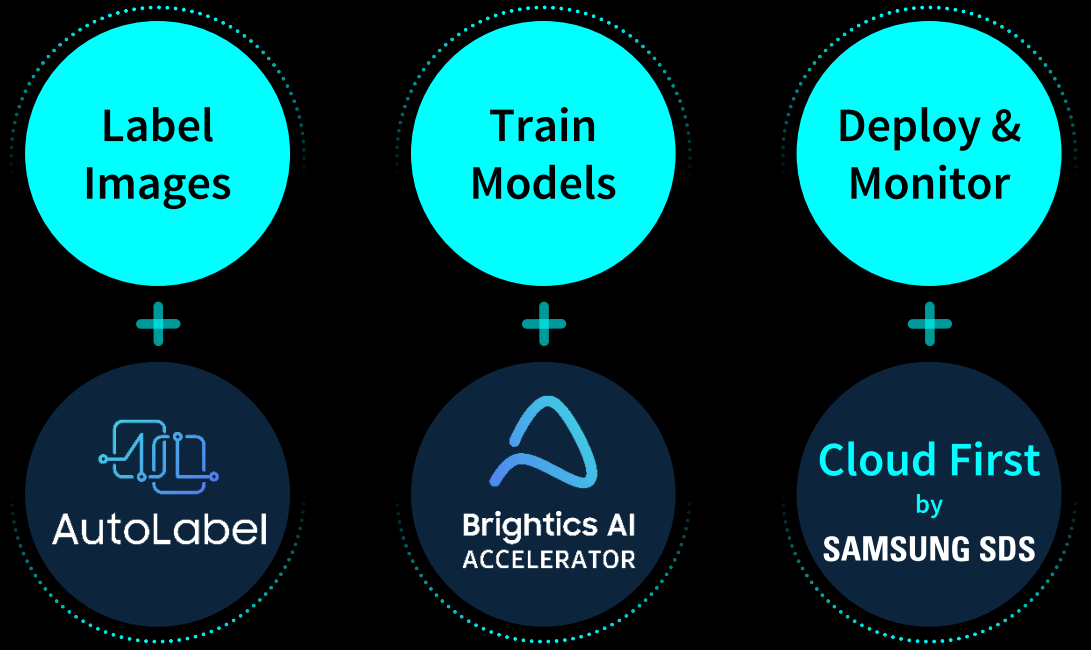

Labeling Medical Images is Costly

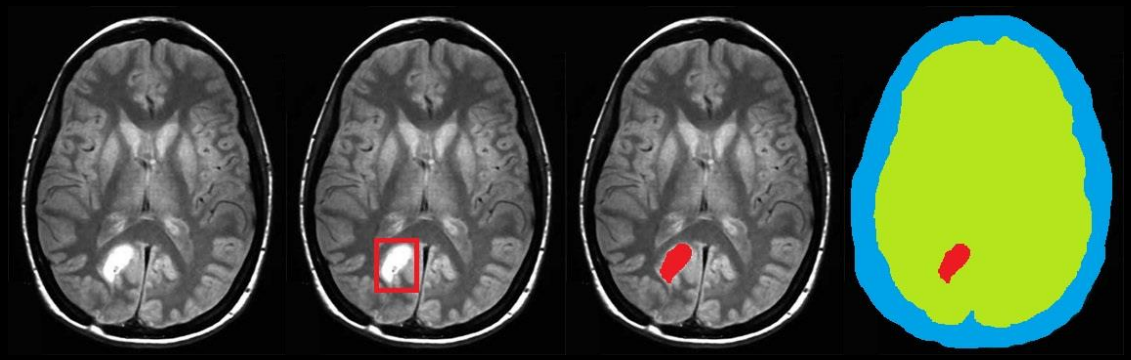

Classification | Detection | Instance Segmentation | Semantic Segmentation

Segmentation Labeling
Requires between 20 minutes per image per labeler – usually 3 labelers are used.

Checking Labels
Requires between 20 seconds per image.

Data Set Size
Typical datasets have 100k – 1m images.

Cost (example)
Labeling everything manually costs $12m and 19 person-years. Using active learning reduces cost to $1.35m and 2 person-years.

Efficiently annotating image datasets

1. **Order matters**
   Some images have more information content than others – label the best ones first

2. **Active learning**
   Finds the right order to leverage the human team

3. **Verification**
   of automatically produced labels for the other images can be done rapidly – leading to a fully human-verified dataset

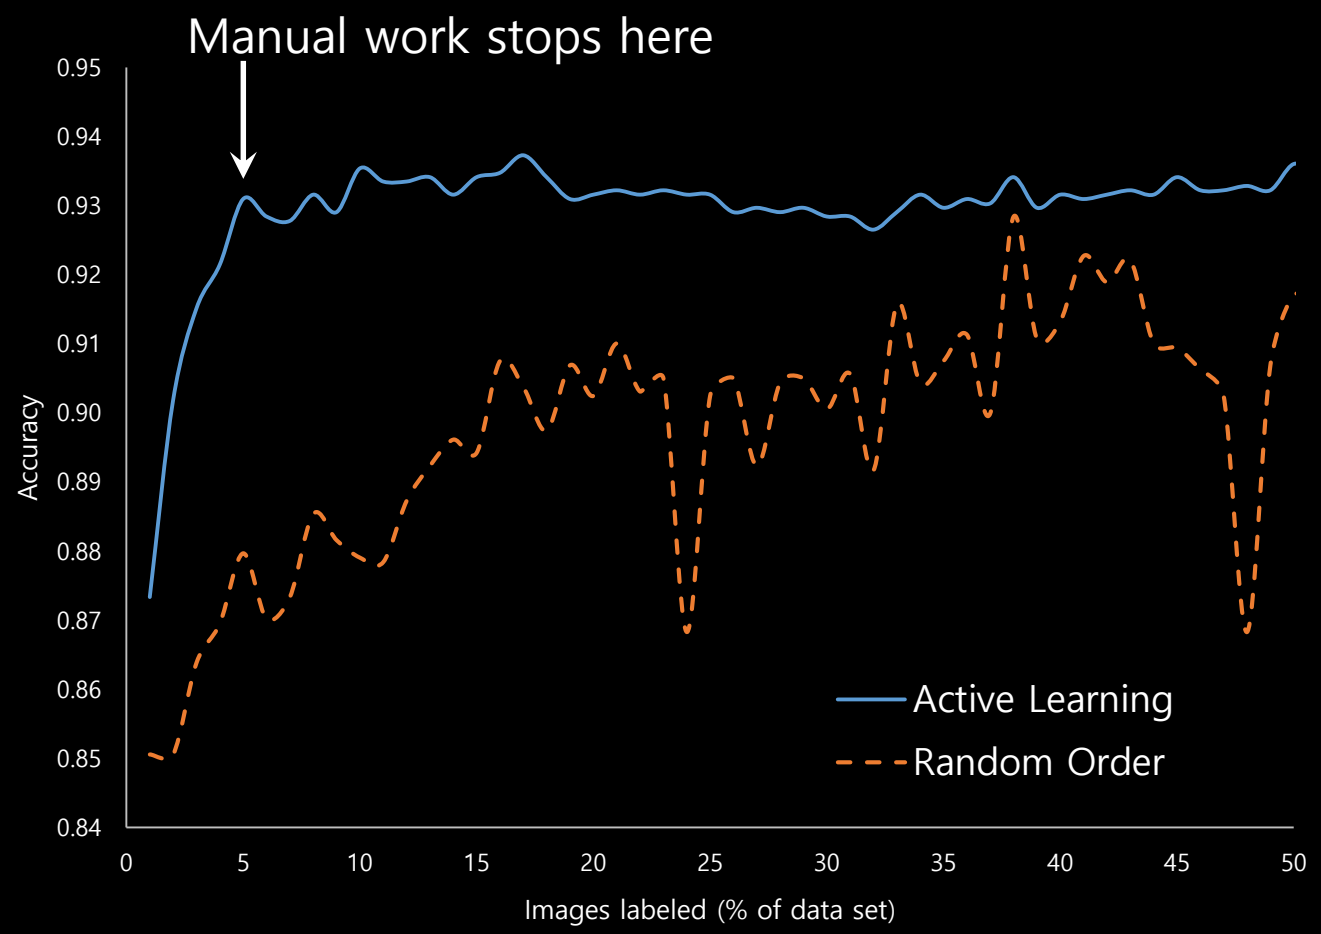

AutoLabel / AutoML Technology

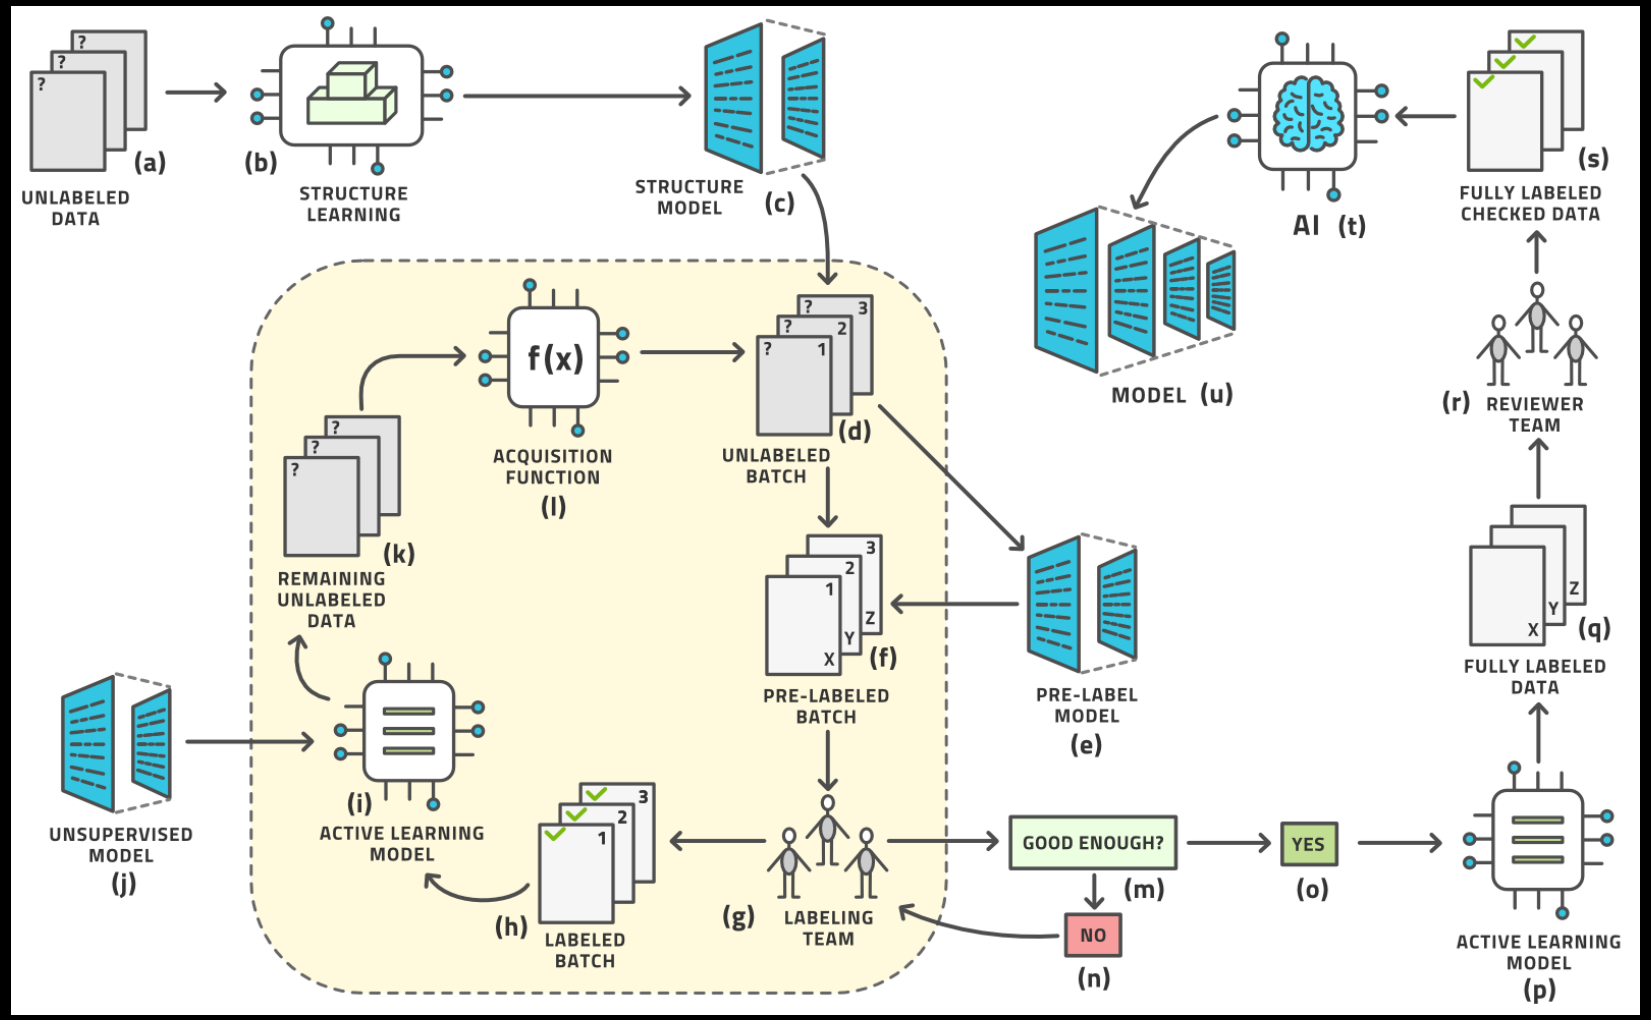

... using many existing tools to help ...

PyTorch OpenAI

Keras mxnet

scikit learn theano

jupyter TensorFlow Horovod HELM

Chainer Caffe Terraform

aws Azure

Innovations by SDS Research America AI Team

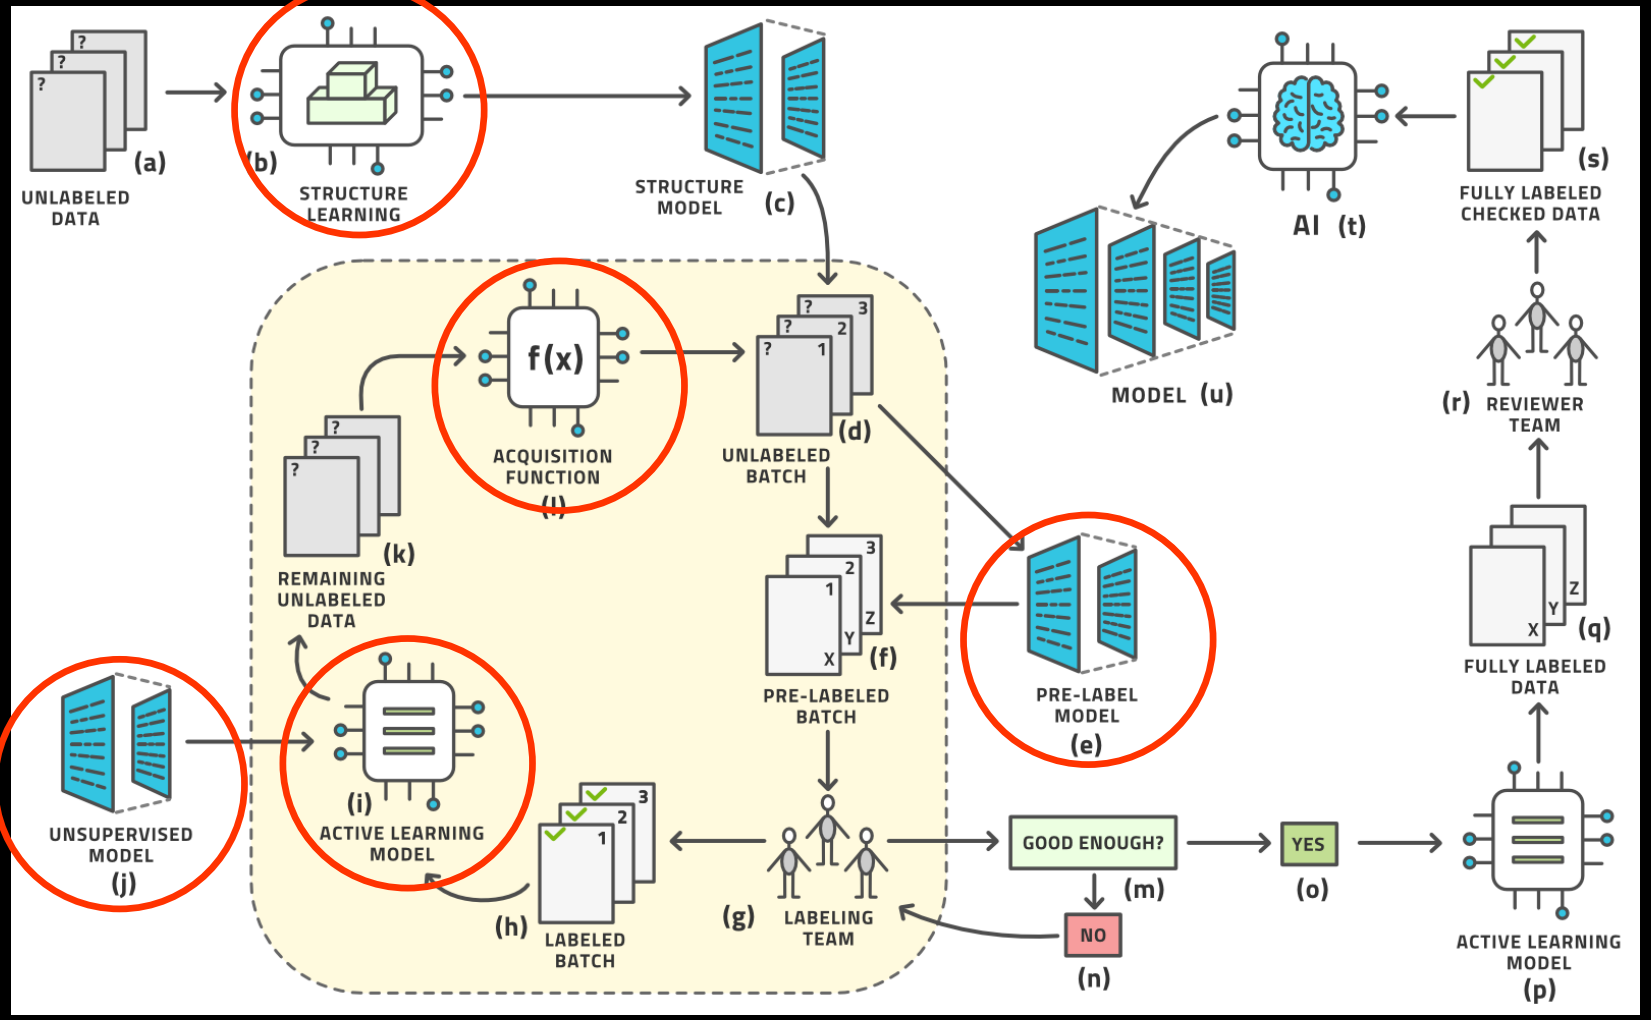

1. Initialization
2. Pre-labeling model
3. Enhanced active learning model architecture
4. Unsupervised meta-model assistance
5. Novel acquisition function

⇒ Twice as efficient as state of the art

⇒ Technical publications available. Patents pending.

Active Learning Acquisition Function - BABA

Extract meaningful uncertainty acquisition leveraging Bayesian Neural Network based predictive output

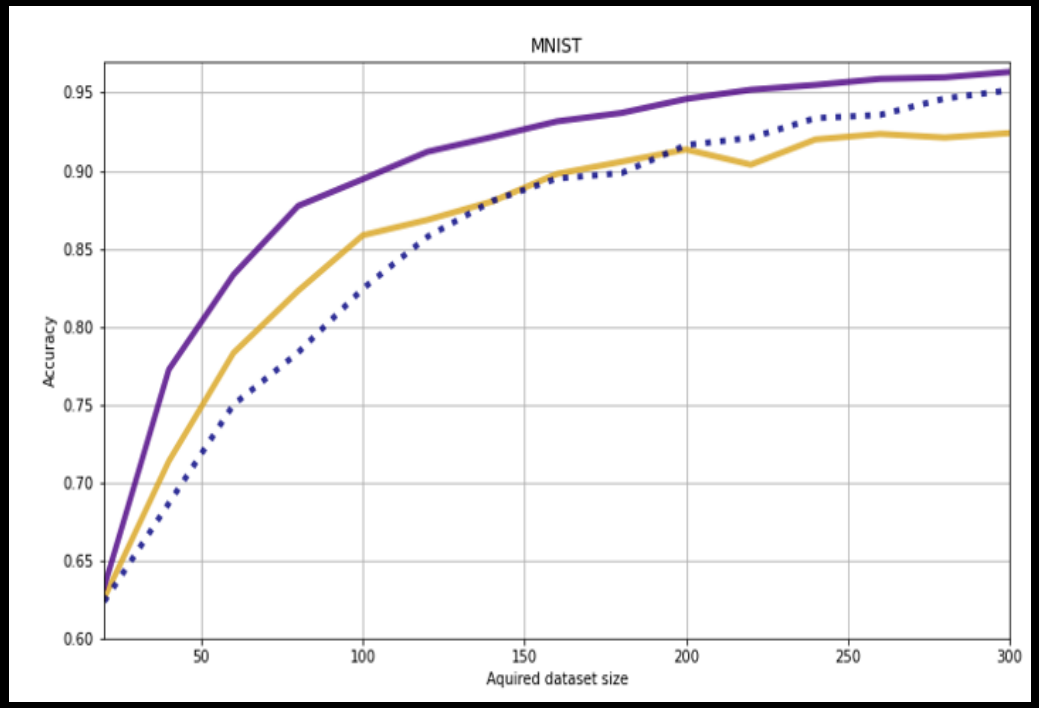

BABA, our new method
BALD, State of Art
Random Order

Superior Data Efficiency

Achieves the state-of-the-art performance in classification and segmentation over the mutual information based acquisition function, a.k.a. BALD.

Explainable Predictive Output

Provide Beta distributed output which gives a better explanation how the model interprets the unlabeled data.

Limitless Scalability

Not require pairwise data distance computations unlike many active learning acquisition methods.

Extremely Robust

Calculate the uncertainty resistant to any randomization of the data or model, e.g. data augmentation, dropout

Super Data Efficient Semantic Segmentation: Pixel-BABA

How it works?

Pixel-BABA is a novel Bayesian Active Learning (AL) model coupled with Pixel level BABA acquisition function.

Pros

Pixel-BABA changes a segmentation labeling task to a classification labeling task. It reaches to fully supervised training performance with only using less that 1% of the data.

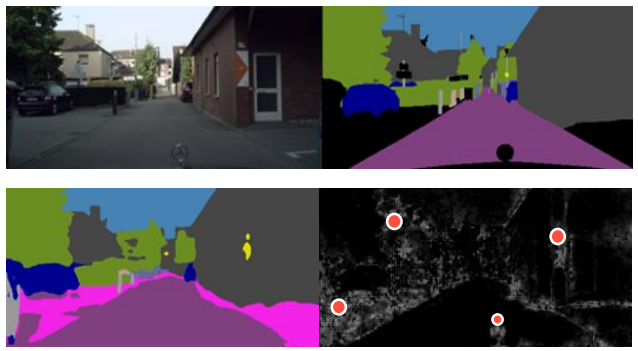

- **Top row:** Image & all pixels per image labels, used for Conventional fully supervised training
- **Bottom row:** Predicted labels and BABA uncertainty maps
- Per cycle Pixel-BABA is trained with few labeled pixels per image.
- Pixel-BABA generates BABA uncertainty map then queries pixels with the highest BABA for labeling

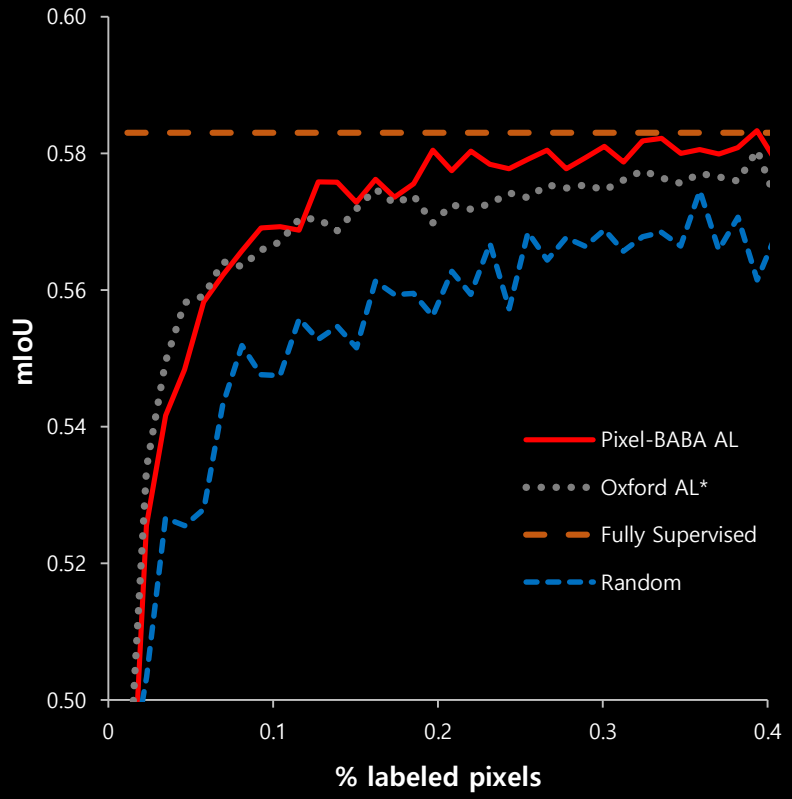

*Gyungin Shin et al. All you need are a few pixels: semantic segmentation with pixelpick
**Result shown are for CamVid dataset, due to data privacy health care data & results are not disclosed

PatchNet for Unsupervised Object Discovery

Object Discovery

PatchNet (1) trains patch embedding using modulated contrastive loss and (2) finds frequently occurring objects by mapping randomly generated patches and clustering in the learned space

Results

PatchNet discovers human faces with ~0.55 accuracy (F1 score) and pedestrians with ~0.35 accuracy

Key benefits

Discovered objects provide detection pre-labels for human correction or for training the initial Active Learning model → Further reduces human labor and improves detecion AutoLabel accuracy

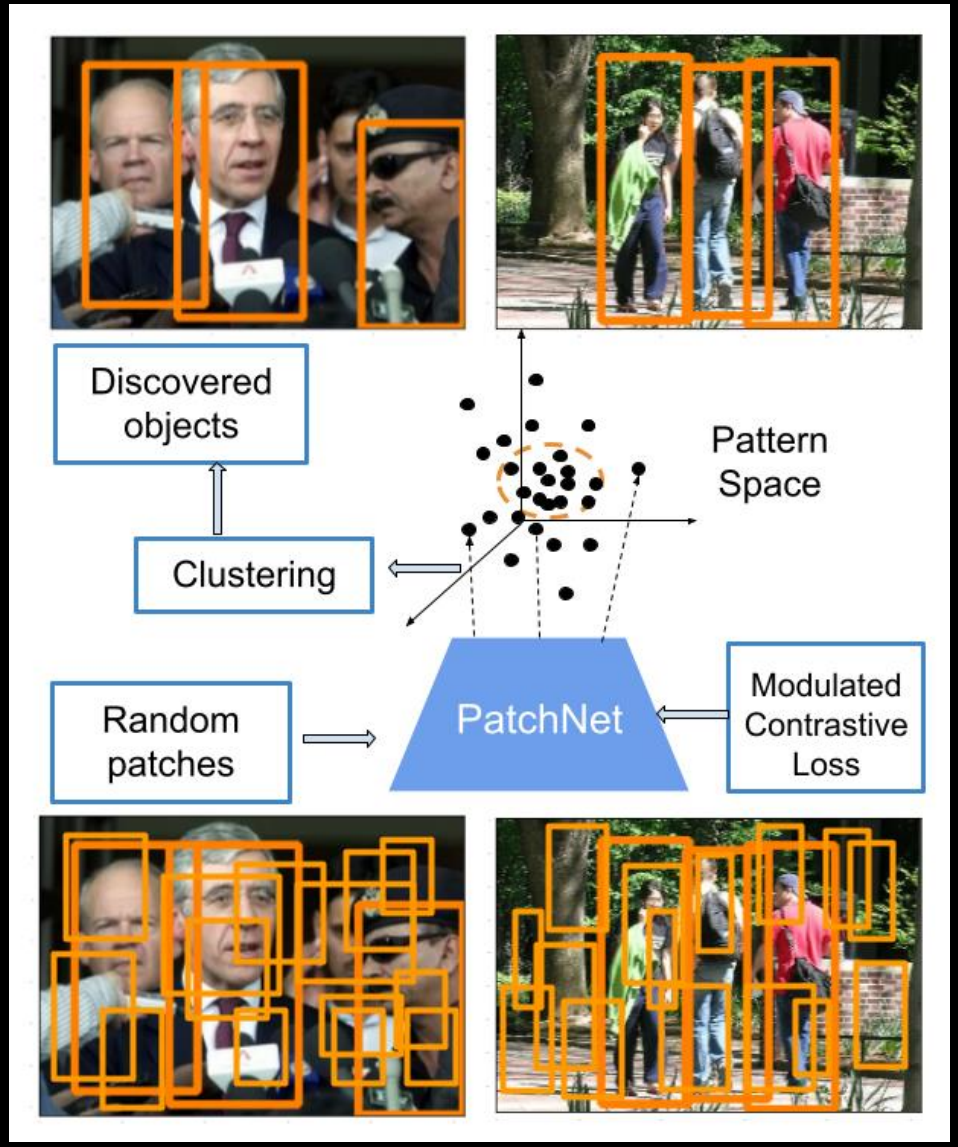

Project Examples

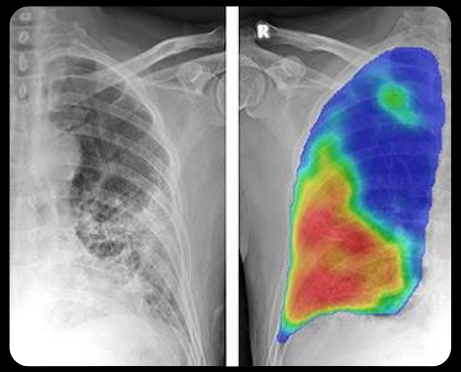

X-Ray

Detect Covid-19, pneumonia, or normality

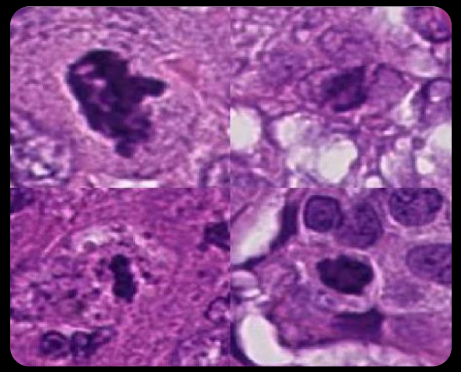

Whole Slide Image

Determine the presence or absence of breast cancer.

City of Hope

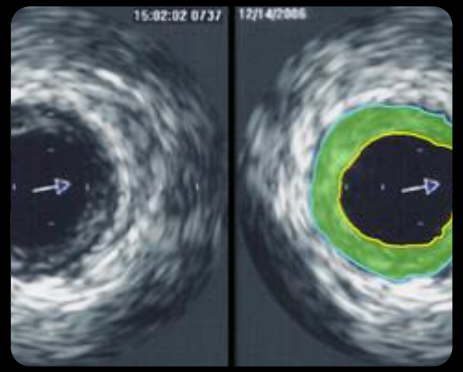

Intravascular Ultrasound

Place stents for cardiovascular diseases.

PHILIPS

Photograph

Detect presence or absence of colon cancer.

INTUITIVE SURGICAL®

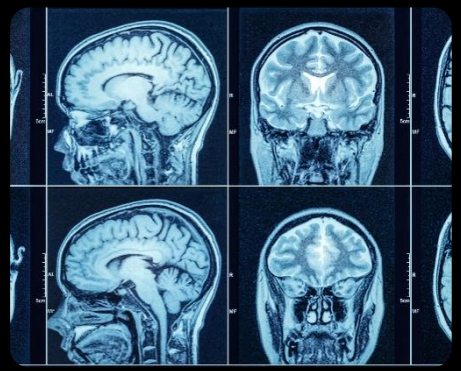

MRI

Find the brain tumor, if any.

City of Hope

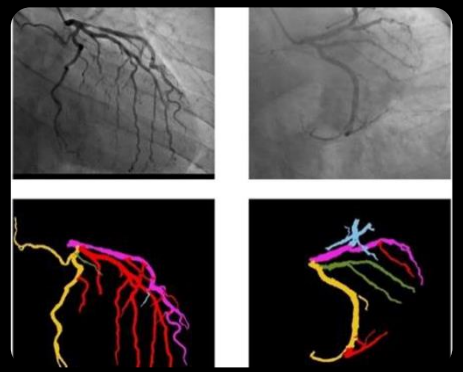

Angiography

Detect blockages in blood vessels.

PHILIPS

The Future Plan

Scientific
Excellence

with Research Institutions

Sharpen active learning methods & healthcare specific solutions such as 3D scanning and multi-modal data

Label & Model
Tool Business

with Customers

Polish software for labelers and model trainers to provide good user experience and efficiency

Diagnostic Models
Live in the Cloud

with Partners

Provide curated models in the cloud for clinical & diagnostic use

Thank you

Q&A

SAMSUNG SDS

Abstract

Detecting and diagnosing diseases based on medical images is one of the major applications of artificial intelligence (AI) in the next 3 – 5 years. We report on largely automating the human annotation process of medical images, which is the time-consuming and expensive prerequisite to AI modeling. Through a novel image sorting method, a novel unsupervised structure learning framework, a novel unsupervised object detection method, and a novel pixel-labeling based approach, we have lowered the necessary human labor far below prior state of the art. After labeling, we train sophisticated deep learning models to perform the diagnosis making use of automatic machine learning (AutoML) approaches as well as using multiple computers simultaneously to reduce the run-time of training workflows. We present these methods in a high-level overview and practical applications to real-life datasets from our industry partners on breast cancer, colon cancer, brain tumors, and intravascular ultrasound.