We, humans...

We, humans support Weaker ones...

OBERYN MARTELL

v/s

THE MOUNTAIN

Source: Game of Thrones

We, humans support Weaker ones...

JERRY

v/s

TOM

Source: Tom and Jerry

We, humans support Weaker ones...

MESSI

v/s

CRISTIANO

Source: Wikipedia

FRIEDRICH-ALEXANDER UNIVERSITÄT ERLANGEN-NÜRNBERG
FACULTY OF ENGINEERING

Oops...

Source: `makeameme.org`

Weakly-Supervised Learning in Digital Pathology

"not so weak, actually!"

Arijit Ghosh
Seminar Digital Pathology and Deep Learning, Friedrich-Alexander-Universität Erlangen-Nürnberg
December 18, 2023

Imagenet

Russakovsky, Olga, et al. "Imagenet large scale visual recognition challenge." IJCV, 2015.

Imagenet is "BIG"!

Russakovsky, Olga, et al. "Imagenet large scale visual recognition challenge." IJCV, 2015.

Imagenet is "BIG"!

- **1 million** images.

Russakovsky, Olga, et al. "Imagenet large scale visual recognition challenge." IJCV, 2015.

Imagenet is "BIG"!

- **1 million** images.
- **1000** classes.

Russakovsky, Olga, et al. "Imagenet large scale visual recognition challenge." IJCV, 2015.

Imagenet is "BIG"!

- **1 million** images.
- **1000** classes.

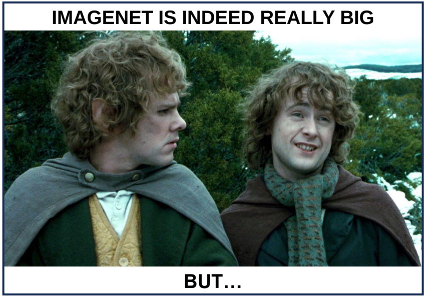

Source: Lord of the Rings

Russakovsky, Olga, et al. "Imagenet large scale visual recognition challenge." IJCV, 2015.

Imagenet samples...

Russakovsky, Olga, et al. "Imagenet large scale visual recognition challenge." IJCV, 2015.

Imagenet samples... are pretty "SMALL"!

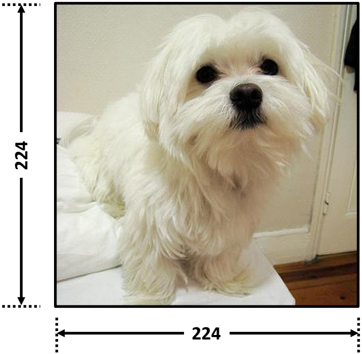

Russakovsky, Olga, et al. "Imagenet large scale visual recognition challenge." IJCV, 2015.

Whole-Slide Image

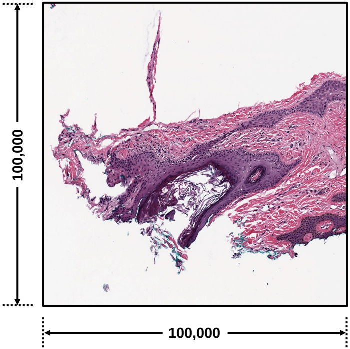

Source: `github.com/camicroscope/Distro`

Whole-Slide Image + CNN...

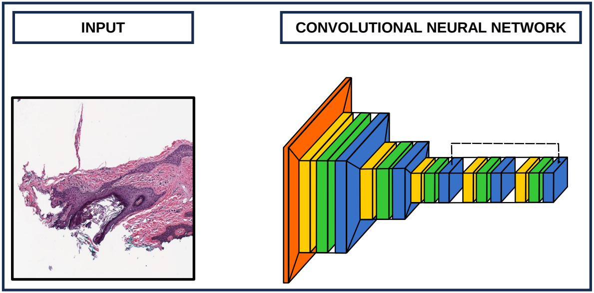

Whole-Slide Image + CNN...

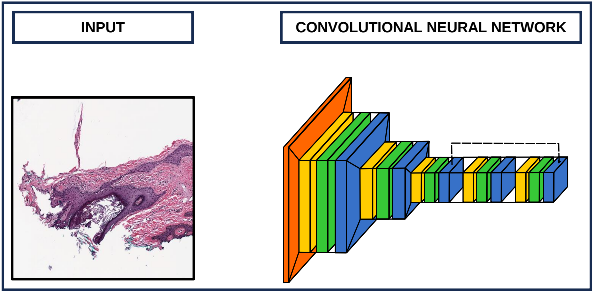

➔ Compute **Bottleneck**!

Whole-Slide Image + CNN = ...

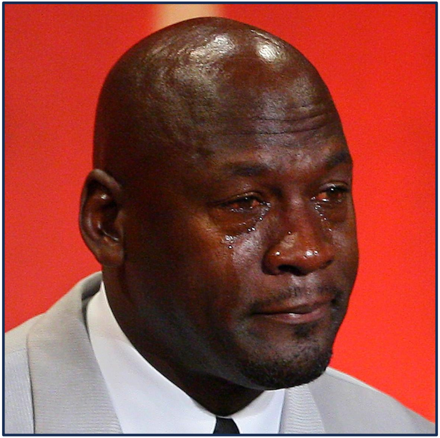

Source: www.npr.org

What now?

Source: meme-arsenal.ru

"Naive" solution 101...

Cruz-Roa, Angel, et al. "Automatic detection of invasive ductal carcinoma in whole slide images with convolutional neural networks." Medical Imaging, 2014.

"Naive" solution 101...

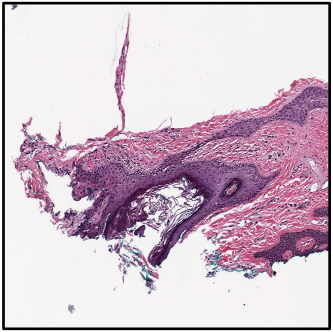

Source: github.com/camicroscope/Distro

Cruz-Roa, Angel, et al. "Automatic detection of invasive ductal carcinoma in whole slide images with convolutional neural networks." Medical Imaging, 2014.

"Naive" solution 101...

➔ Create **patches**.

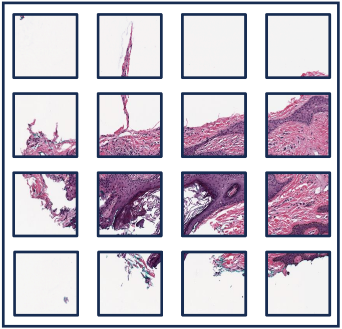

Cruz-Roa, Angel, et al. "Automatic detection of invasive ductal carcinoma in whole slide images with convolutional neural networks." Medical Imaging, 2014.

"Naive" solution 101...

➔ **Label** each patch.

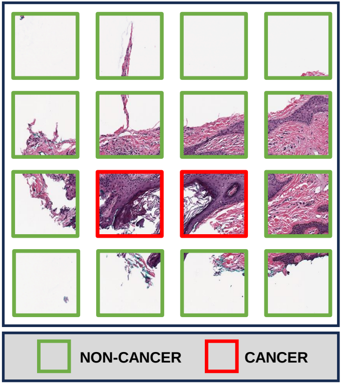

Cruz-Roa, Angel, et al. "Automatic detection of invasive ductal carcinoma in whole slide images with convolutional neural networks." Medical Imaging, 2014.

"Naive" solution 101...

➔ **Train** with patches.

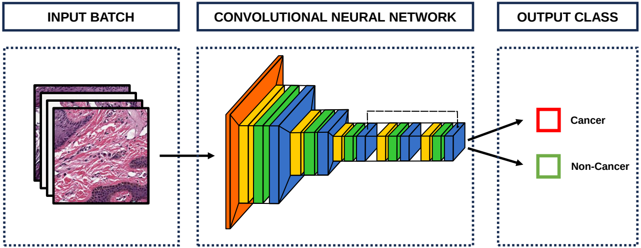

Cruz-Roa, Angel, et al. "Automatic detection of invasive ductal carcinoma in whole slide images with convolutional neural networks." Medical Imaging, 2014.

"Naive" solution 101...

→ **Combine** probabilities during **testing**.

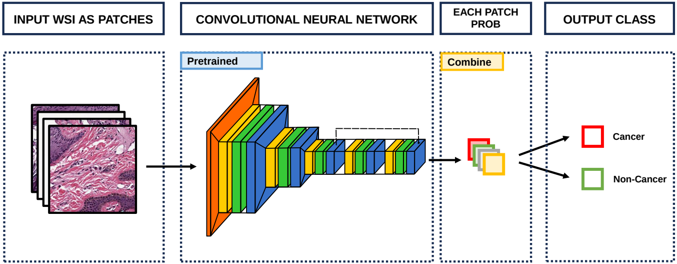

Cruz-Roa, Angel, et al. "Automatic detection of invasive ductal carcinoma in whole slide images with convolutional neural networks." Medical Imaging, 2014.

"Naive" solution 101...is Awesome!

Source: Breaking Bad

"Naive" solution 101...is NOT Awesome!

SADLY, NO!

Source: Breaking Bad

FRIEDRICH-ALEXANDER UNIVERSITÄT ERLANGEN-NÜRNBERG
FACULTY OF ENGINEERING

Why?

➔ **Annotation** time.

Why?

➔ **Annotation** time.

Source: www.unionsquarenyc.org

Why?

→ **Annotation** time.

CLASSIFICATION

Name of class : DOG

~ 1 sec

Inspired from Introduction slide, "Tutorial on Weakly-Supervised Learning in Computer Vision." ECCV 2020.

Why?

➔ **Annotation** time.

CLASSIFICATION

Name of class : DOG

~ 1 sec

DETECTION

Name of class : DOG
+ Bounding Box

~ 10 sec

Inspired from Introduction slide, "Tutorial on Weakly-Supervised Learning in Computer Vision." ECCV 2020.

Why?

➔ **Annotation** time.

CLASSIFICATION

Name of class : DOG

⏰ **~ 1 sec**

DETECTION

Name of class : DOG
+ Bounding Box

⏰ **~ 10 sec**

SEGMENTATION

Name of class : DOG
+ Segmentation Mask

⏰ **~ 78 sec**

Inspired from Introduction slide, "Tutorial on Weakly-Supervised Learning in Computer Vision." ECCV 2020.

Why?

→ **Annotation** time for **1 Whole Slide Image**.

CLASSIFICATION

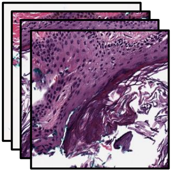

Name of class : C or NC

~ 1 sec x K

DETECTION

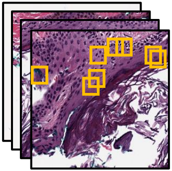

Name of class : C
+ Bounding Box

~ 10 sec x K

SEGMENTATION

Name of class : C
+ Segmentation Mask

~ 78 sec x K

→ **K** = number of patches

Why?

→ **Requires** experts.

Source: www.realsimple.com

Then what to do?

Source: F.R.I.E.N.D.S.

Hi, Weakly Supervised Learning!

Hi, Weakly Supervised Learning!

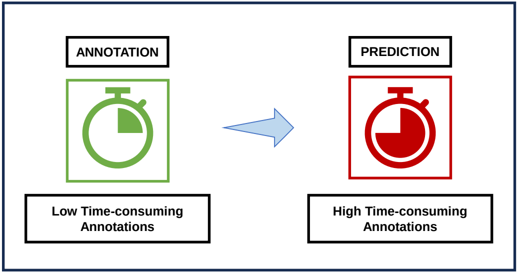

Inspired from Introduction slide, "Tutorial on Weakly-Supervised Learning in Computer Vision." ECCV 2020.

Hi, Weakly Supervised Learning!

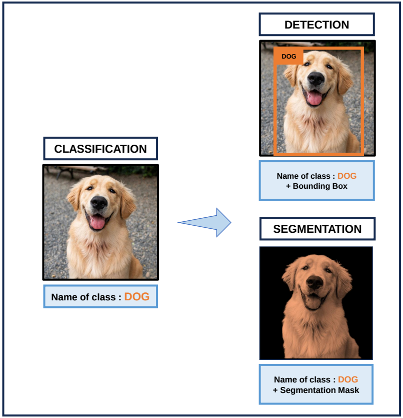

Inspired from Introduction slide, "Tutorial on Weakly-Supervised Learning in Computer Vision." ECCV 2020.

Hi, Weakly Supervised Learning!

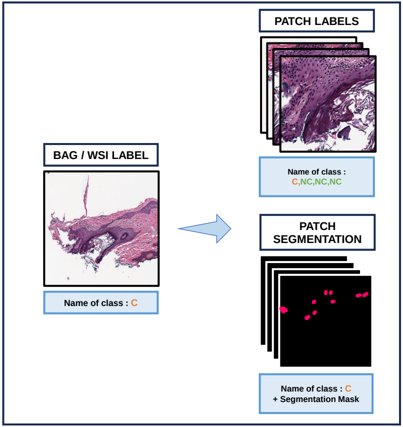

Inspired from Introduction slide, "Tutorial on Weakly-Supervised Learning in Computer Vision." ECCV 2020.

But, how?

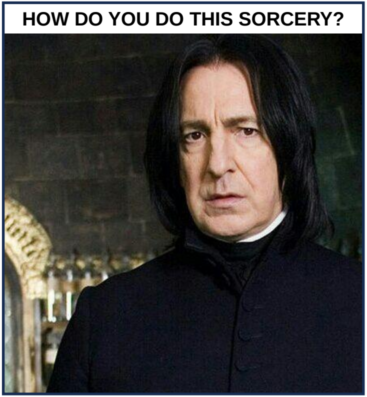

Source: Harry Potter

Bonjour, Multiple Instance Learning Constraint!

BAG / WSI

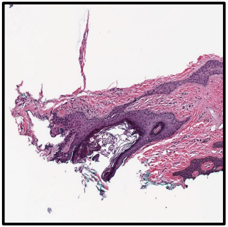

Bag Label : NC

Bonjour, Multiple Instance Learning Constraint!

BAG / WSI

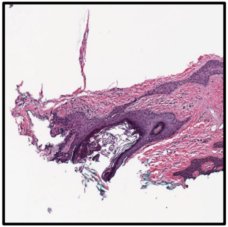

Bag Label : NC

PATCHES / INSTANCES

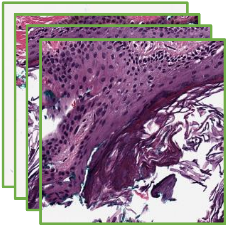

All Patch Label : NC

Bonjour, Multiple Instance Learning Constraint!

BAG / WSI

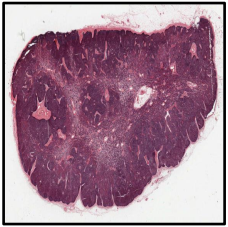

Bag Label : C

Source: downloads.openmicroscopy.org/images/SVS/

Bonjour, Multiple Instance Learning Constraint!

BAG / WSI

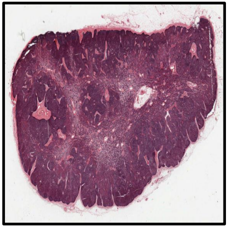

Bag Label : C

PATCHES / INSTANCES

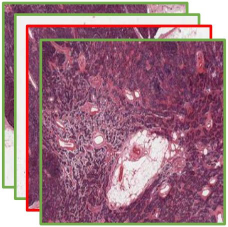

At least 1 Patch Label : C

Source: `downloads.openmicroscopy.org/images/SVS/`

MIL-Max

Campanella, Gabriele, et al. "Clinical-grade computational pathology using weakly supervised deep learning on whole slide images." Nature medicine, 2019.

MIL-Max

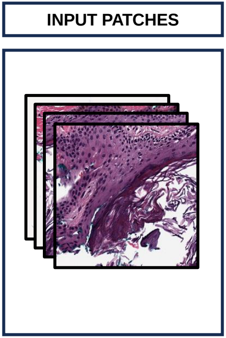

Campanella, Gabriele, et al. "Clinical-grade computational pathology using weakly supervised deep learning on whole slide images." Nature medicine, 2019.

MIL-Max

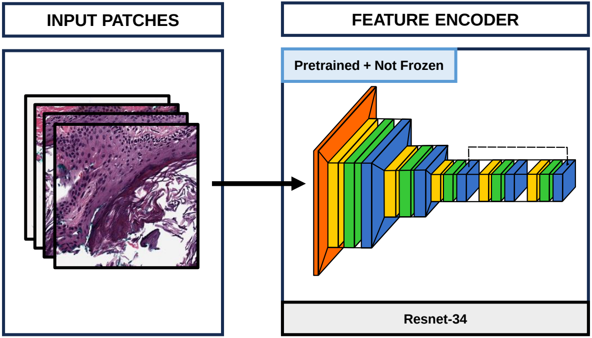

Campanella, Gabriele, et al. "Clinical-grade computational pathology using weakly supervised deep learning on whole slide images." Nature medicine, 2019.

MIL-Max

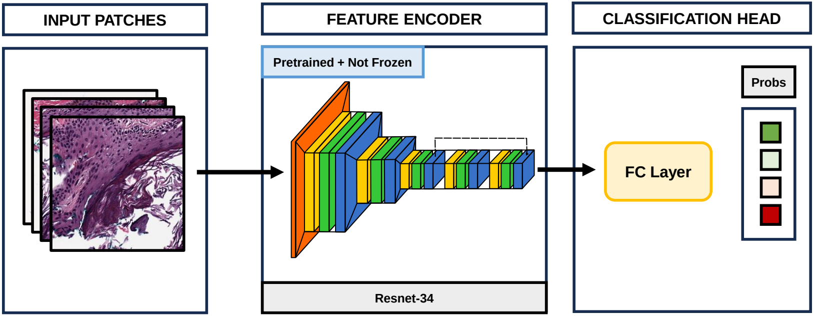

Campanella, Gabriele, et al. "Clinical-grade computational pathology using weakly supervised deep learning on whole slide images." Nature medicine, 2019.

MIL-Max

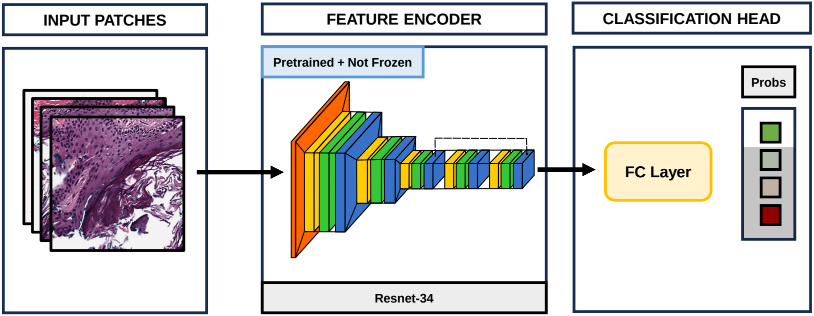

Campanella, Gabriele, et al. "Clinical-grade computational pathology using weakly supervised deep learning on whole slide images." Nature medicine, 2019.

MIL-Max

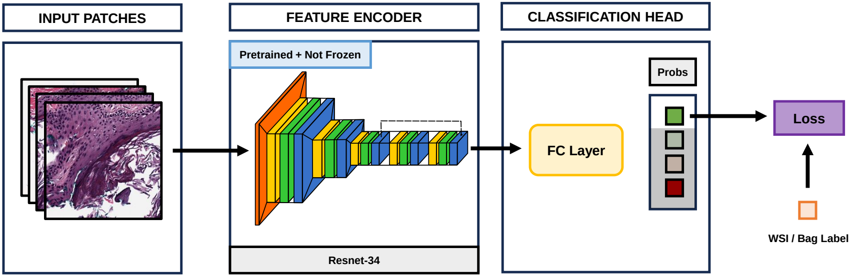

Campanella, Gabriele, et al. "Clinical-grade computational pathology using weakly supervised deep learning on whole slide images." Nature medicine, 2019.

MIL-Max – Flaw

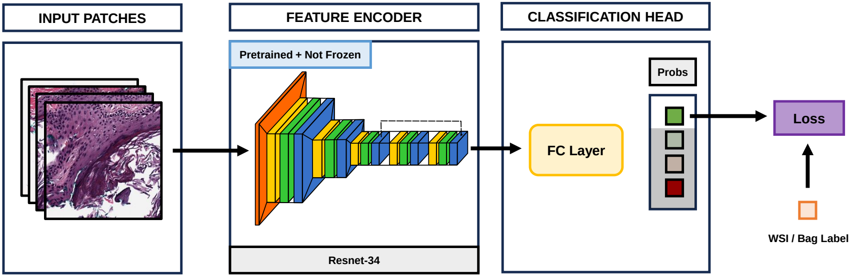

Information Loss

Campanella, Gabriele, et al. "Clinical-grade computational pathology using weakly supervised deep learning on whole slide images." Nature medicine, 2019.

MIL-RNN

Campanella, Gabriele, et al. "Clinical-grade computational pathology using weakly supervised deep learning on whole slide images." Nature medicine, 2019.

MIL-RNN

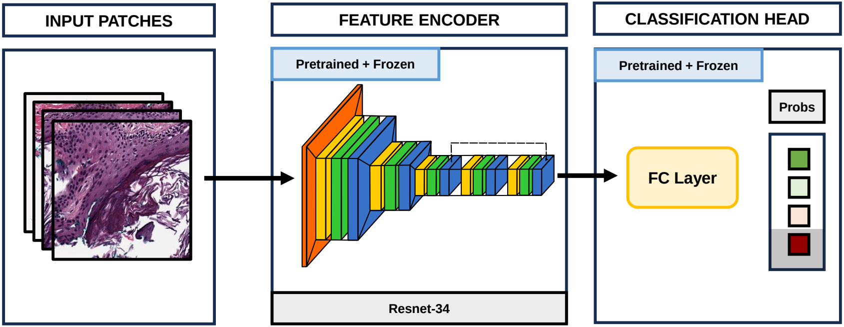

Campanella, Gabriele, et al. "Clinical-grade computational pathology using weakly supervised deep learning on whole slide images." Nature medicine, 2019.

MIL-RNN

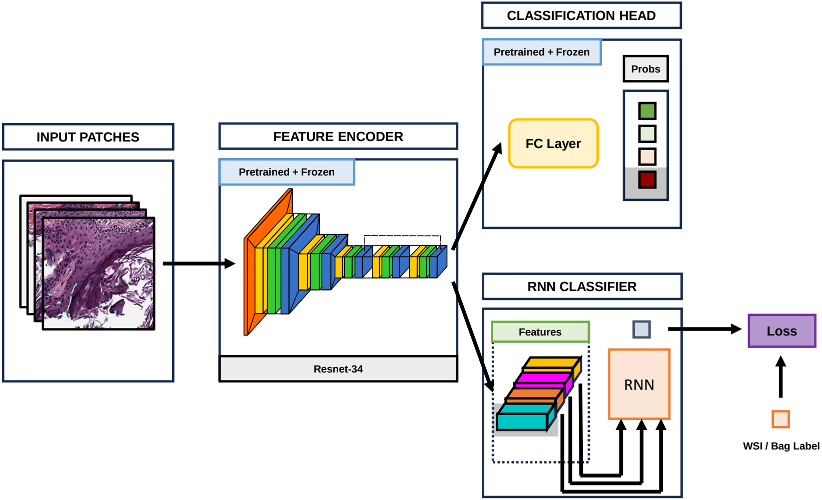

Campanella, Gabriele, et al. "Clinical-grade computational pathology using weakly supervised deep learning on whole slide images." Nature medicine, 2019.

MIL-Max and MIL-RNN – Flaw

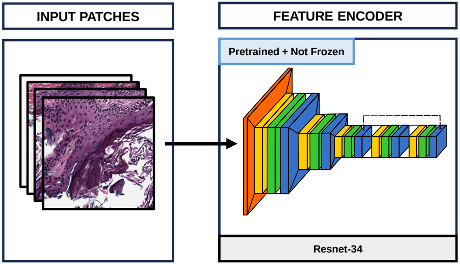

➔ **Multiple** Encoder passes.

Campanella, Gabriele, et al. "Clinical-grade computational pathology using weakly supervised deep learning on whole slide images." Nature medicine, 2019.

MIL-Max and MIL-RNN – Flaw

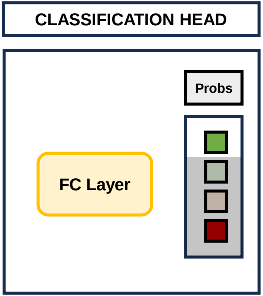

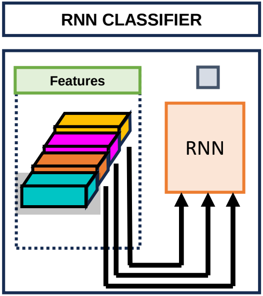

➔ **Binary** classifiers.

Campanella, Gabriele, et al. "Clinical-grade computational pathology using weakly supervised deep learning on whole slide images." Nature medicine, 2019.

CLAM

Lu, Ming Y., et al. "Data-efficient and weakly supervised computational pathology on whole-slide images." Nature Biomedical Engineering, 2021.

CLAM

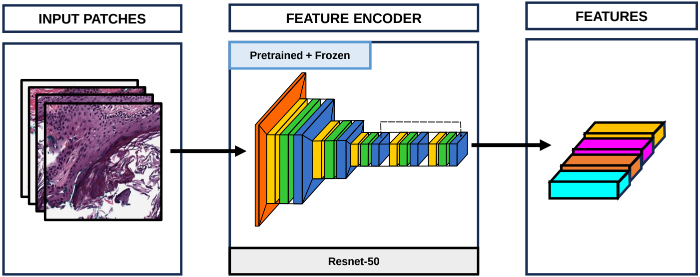

Lu, Ming Y., et al. "Data-efficient and weakly supervised computational pathology on whole-slide images." Nature Biomedical Engineering, 2021.

CLAM

➔ Train using **saved features**.

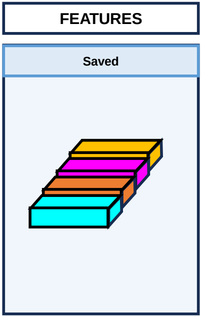

Lu, Ming Y., et al. "Data-efficient and weakly supervised computational pathology on whole-slide images." Nature Biomedical Engineering, 2021.

CLAM

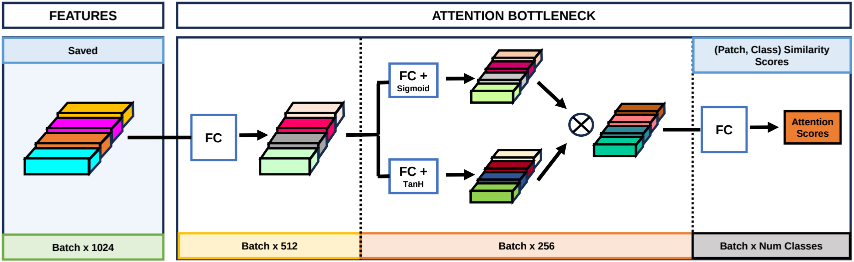

Lu, Ming Y., et al. "Data-efficient and weakly supervised computational pathology on whole-slide images." Nature Biomedical Engineering, 2021.

CLAM

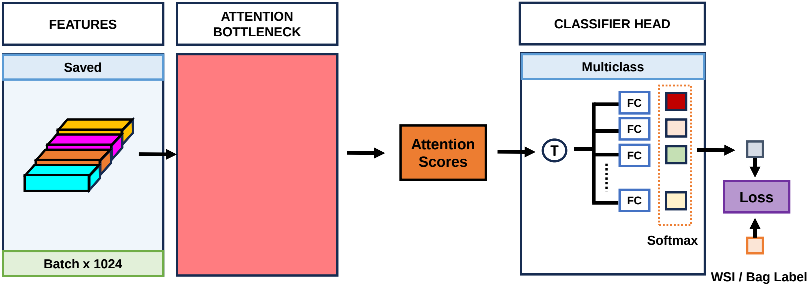

Lu, Ming Y., et al. "Data-efficient and weakly supervised computational pathology on whole-slide images." Nature Biomedical Engineering, 2021.

CLAM – Flaw

➔ Doesn’t consider **Instance Correlation**.

Lu, Ming Y., et al. "Data-efficient and weakly supervised computational pathology on whole-slide images." Nature Biomedical Engineering, 2021.

CLAM – Flaw

➔ Doesn’t consider **Instance Correlation**.

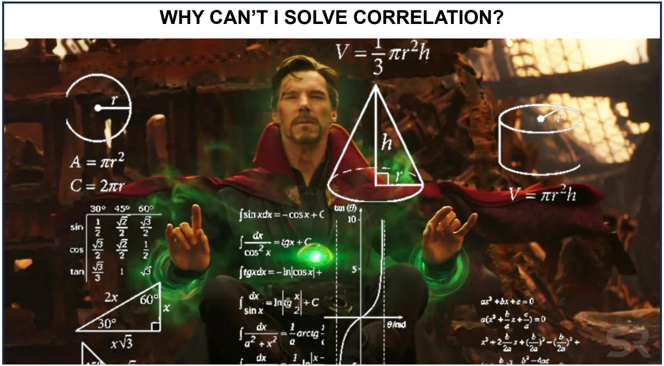

Source: Avengers: Infinity War

Lu, Ming Y., et al. "Data-efficient and weakly supervised computational pathology on whole-slide images." Nature Biomedical Engineering, 2021.

So, what to do?

So, what to do?

Source: `www.sescoops.com`

TransMIL

Shao, Zhuchen, et al. "Transmil: Transformer based correlated multiple instance learning for whole slide image classification." NeurIPS, 2021.

TransMIL

➔ Trains using **Saved Features** from ResNet-50.

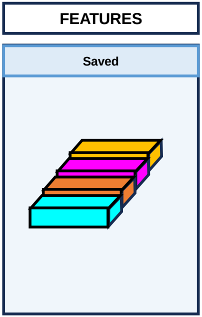

Shao, Zhuchen, et al. "Transmil: Transformer based correlated multiple instance learning for whole slide image classification." NeurIPS, 2021.

TransMIL

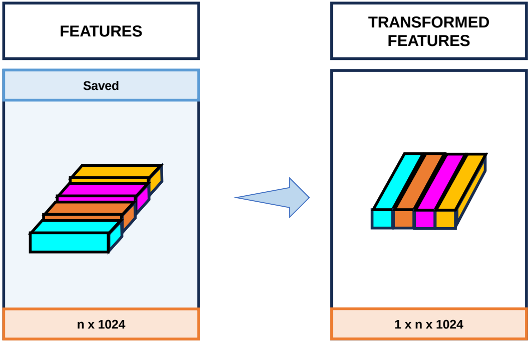

Shao, Zhuchen, et al. "Transmil: Transformer based correlated multiple instance learning for whole slide image classification." NeurIPS, 2021.

TransMIL

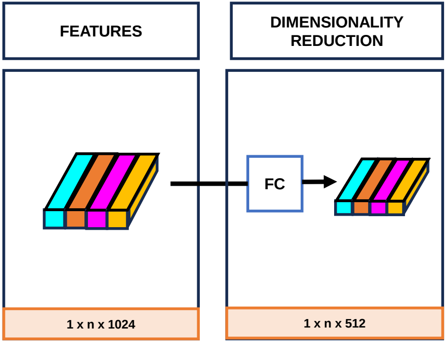

Shao, Zhuchen, et al. "Transmil: Transformer based correlated multiple instance learning for whole slide image classification." NeurIPS, 2021.

TransMIL

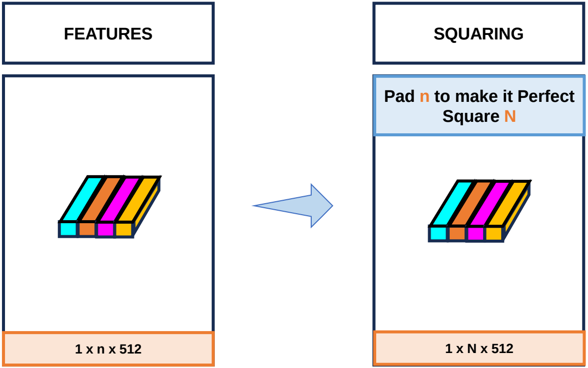

Shao, Zhuchen, et al. "Transmil: Transformer based correlated multiple instance learning for whole slide image classification." NeurIPS, 2021.

TransMIL

SQUARED FEATURES

TransMIL BLOCK

Pad n to make it Perfect Square N

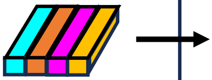

1 x N x 512

Shao, Zhuchen, et al. "Transmil: Transformer based correlated multiple instance learning for whole slide image classification." NeurIPS, 2021.

TransMIL

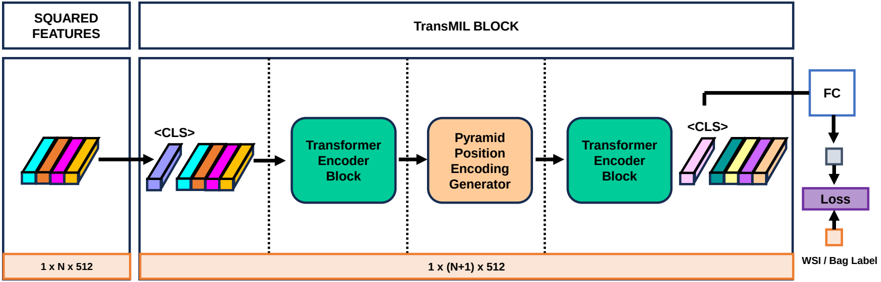

Shao, Zhuchen, et al. "Transmil: Transformer based correlated multiple instance learning for whole slide image classification." NeurIPS, 2021.

TransMIL

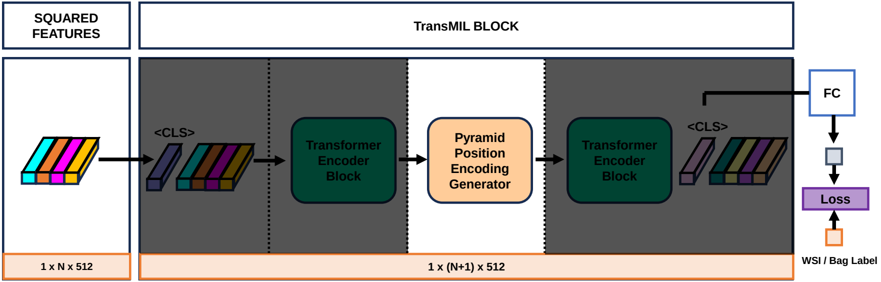

Shao, Zhuchen, et al. "Transmil: Transformer based correlated multiple instance learning for whole slide image classification." NeurIPS, 2021.

TransMIL

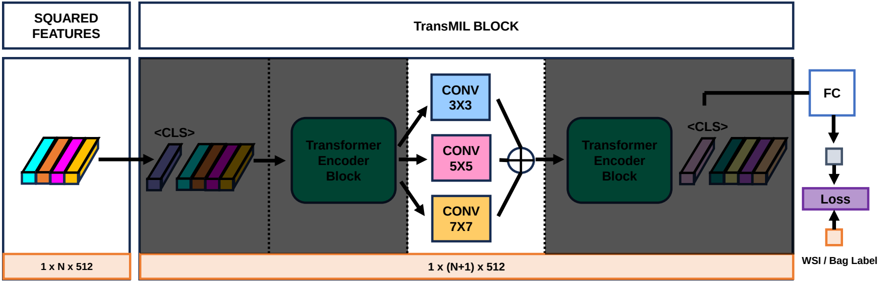

Shao, Zhuchen, et al. "Transmil: Transformer based correlated multiple instance learning for whole slide image classification." NeurIPS, 2021.

TransMIL

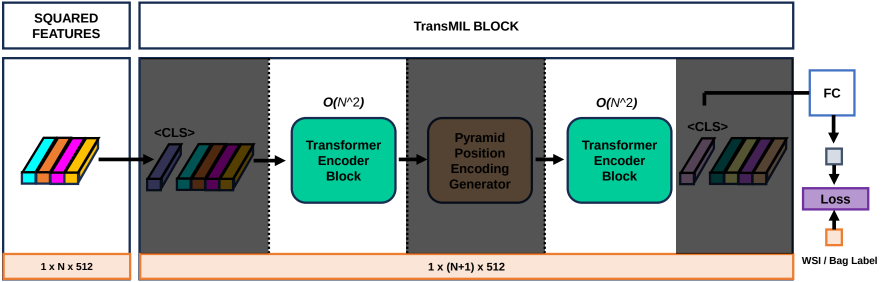

Shao, Zhuchen, et al. "Transmil: Transformer based correlated multiple instance learning for whole slide image classification." NeurIPS, 2021.

TransMIL

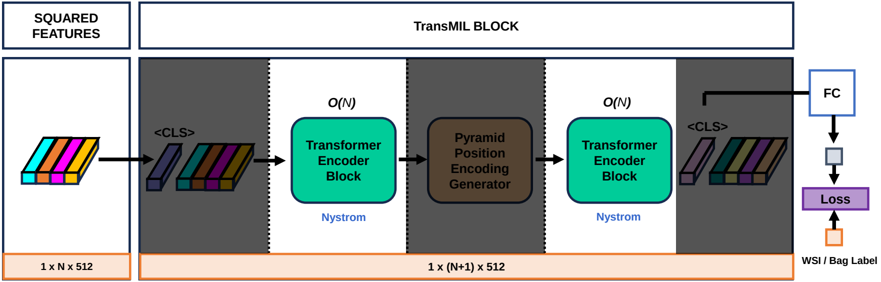

Shao, Zhuchen, et al. "Transmil: Transformer based correlated multiple instance learning for whole slide image classification." NeurIPS, 2021.

Xiong, Yunyang, et al. "Nyströmformer: A nyström-based algorithm for approximating self-attention." AAAI, 2021.

Source: The Office

Segmentation

CAMEL

NO NOT THIS ONE!

Source: nexgenvetrx.com

CAMEL

➔ **Three** Stages.

Xu, Gang, et al. "Camel: A weakly supervised learning framework for histopathology image segmentation." CVPR, 2019.

CAMEL

➔ **Stage - 1** : Label Enrichment.

Xu, Gang, et al. "Camel: A weakly supervised learning framework for histopathology image segmentation." CVPR, 2019.

CAMEL

→ **Stage - 1** : Label Enrichment.

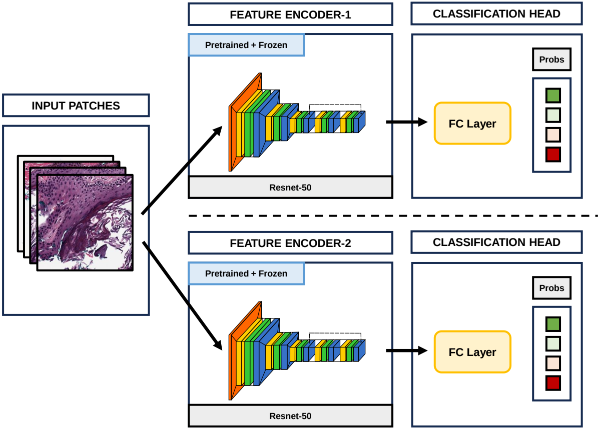

Xu, Gang, et al. "Camel: A weakly supervised learning framework for histopathology image segmentation." CVPR, 2019.

CAMEL

→ **Stage - 1** : Label Enrichment.

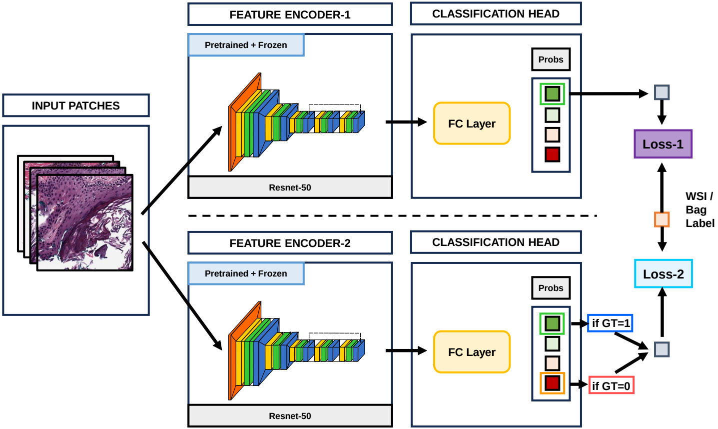

Xu, Gang, et al. "Camel: A weakly supervised learning framework for histopathology image segmentation." CVPR, 2019.

CAMEL

➔ **Stage - 1** : Label Enrichment.

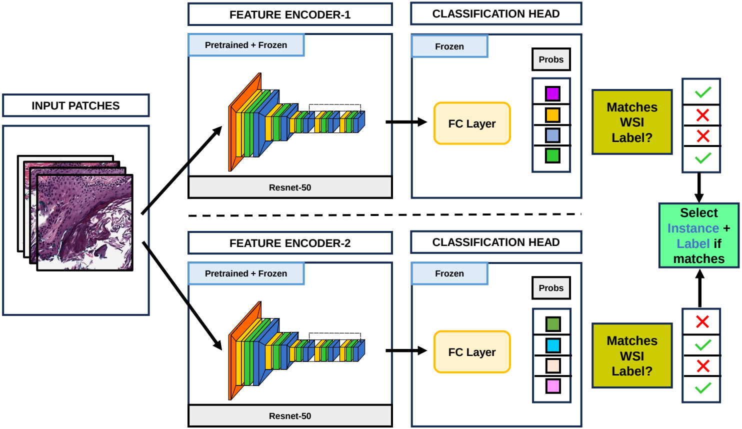

Xu, Gang, et al. "Camel: A weakly supervised learning framework for histopathology image segmentation." CVPR, 2019.

CAMEL

→ **Stage - 2** : Retrain and Relabel.

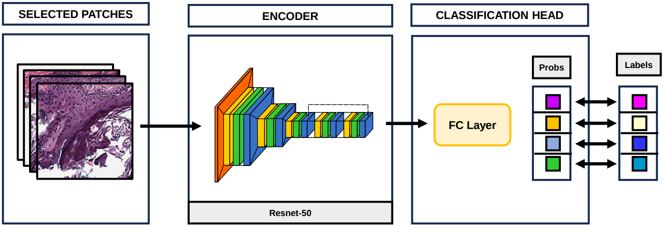

Xu, Gang, et al. "Camel: A weakly supervised learning framework for histopathology image segmentation." CVPR, 2019.

CAMEL

→ **Stage - 2** : Retrain and Relabel.

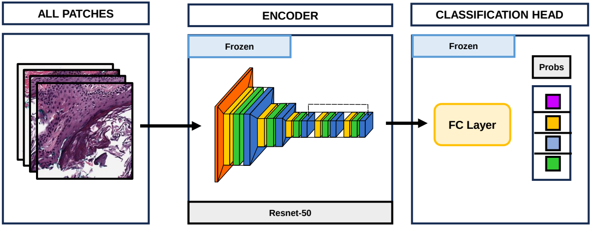

Xu, Gang, et al. "Camel: A weakly supervised learning framework for histopathology image segmentation." CVPR, 2019.

CAMEL

→ **Stage - 3** : Segmentation.

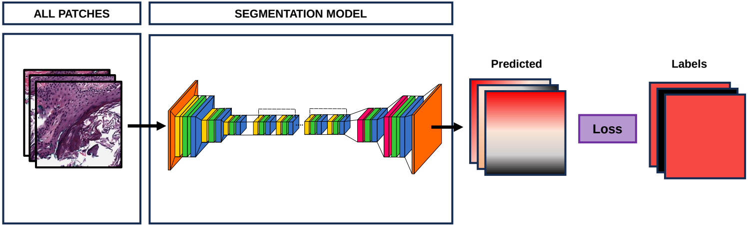

Xu, Gang, et al. "Camel: A weakly supervised learning framework for histopathology image segmentation." CVPR, 2019.

More on Segmentation...

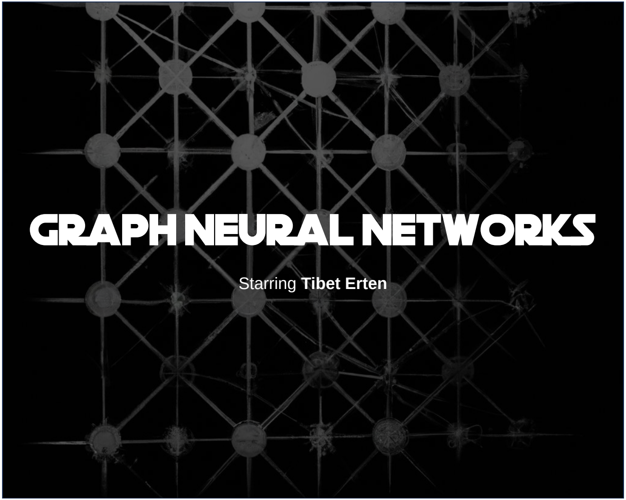

Source: Background Image generated using DALLE-2

Major Takeaways

Major Takeaways

- **Only options** in most cases.

Major Takeaways

- **Only options** in most cases.
- **Feature Encoders** – still **Imagenet pretrained**?

Thank You!!

Questions?

Appendix

All previous methods – Flaw

➔ **Imagenet Pretrained** Encoders.

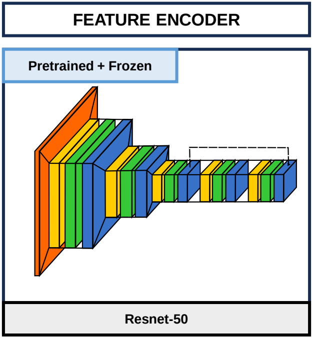

Pinckaers, Hans, et al. "Detection of prostate cancer in whole-slide images through end-to-end training with image-level labels." IEEE TMI, 2021.

StreamingCNN

Pinckaers, Hans, et al. "Detection of prostate cancer in whole-slide images through end-to-end training with image-level labels." IEEE TMI, 2021.

StreamingCNN

➔ Previous architectures.

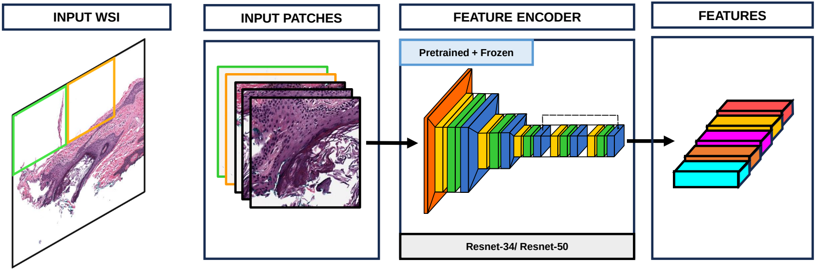

Pinckaers, Hans, et al. "Detection of prostate cancer in whole-slide images through end-to-end training with image-level labels." IEEE TMI, 2021.

StreamingCNN

→ **Pretrain** using **patches**?

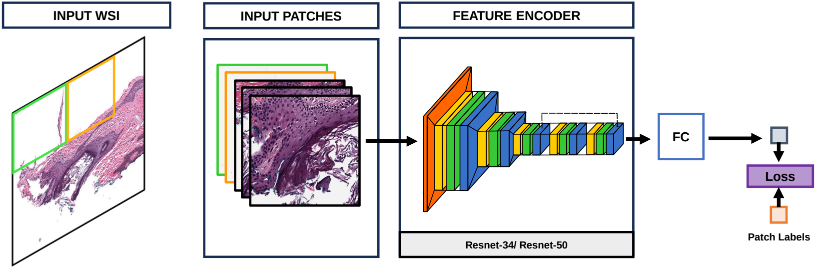

Pinckaers, Hans, et al. "Detection of prostate cancer in whole-slide images through end-to-end training with image-level labels." IEEE TMI, 2021.

StreamingCNN

➔ No **Patch Labels**.

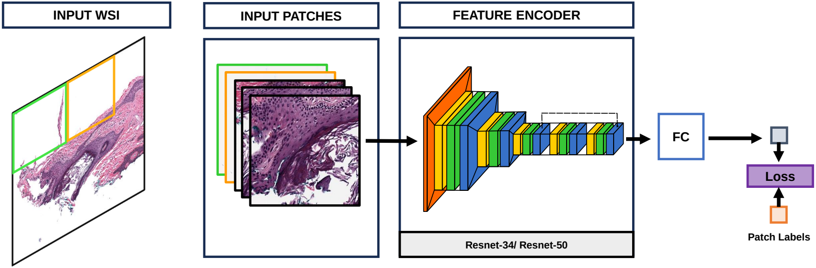

Pinckaers, Hans, et al. "Detection of prostate cancer in whole-slide images through end-to-end training with image-level labels." IEEE TMI, 2021.

StreamingCNN

➔ Direct **WSI processing** – not possible.

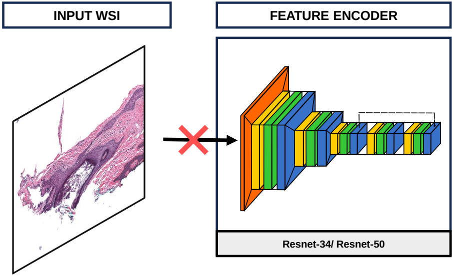

Pinckaers, Hans, et al. "Detection of prostate cancer in whole-slide images through end-to-end training with image-level labels." IEEE TMI, 2021.

StreamingCNN

➔ **Large** number of weights.

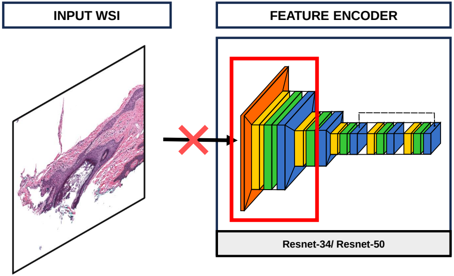

Pinckaers, Hans, et al. "Detection of prostate cancer in whole-slide images through end-to-end training with image-level labels." IEEE TMI, 2021.

StreamingCNN

→ **Fits** compute resources.

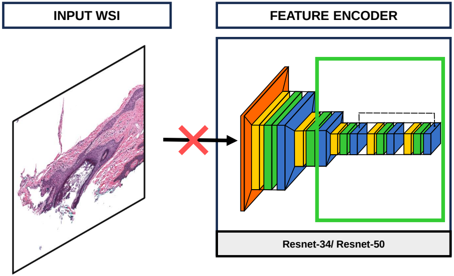

Pinckaers, Hans, et al. "Detection of prostate cancer in whole-slide images through end-to-end training with image-level labels." IEEE TMI, 2021.

StreamingCNN

➔ **Split** feature encoder.

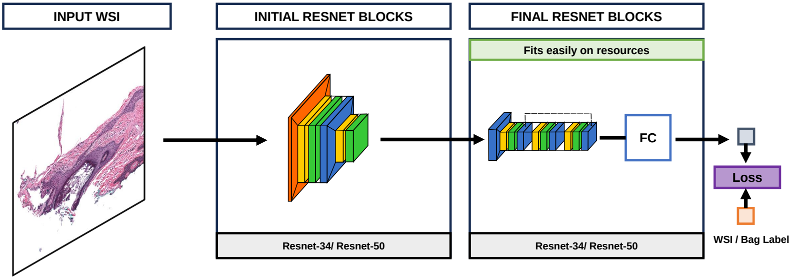

Pinckaers, Hans, et al. "Detection of prostate cancer in whole-slide images through end-to-end training with image-level labels." IEEE TMI, 2021.

StreamingCNN

➔ Take care of **inital part**.

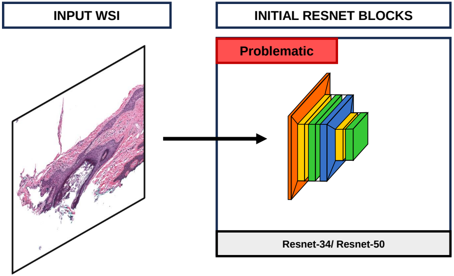

Pinckaers, Hans, et al. "Detection of prostate cancer in whole-slide images through end-to-end training with image-level labels." IEEE TMI, 2021.

StreamingCNN

→ **Convolve** and **Save**.

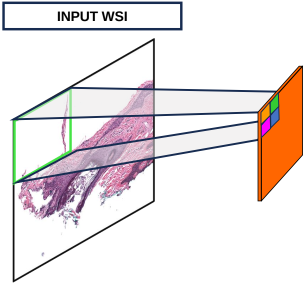

Pinckaers, Hans, et al. "Detection of prostate cancer in whole-slide images through end-to-end training with image-level labels." IEEE TMI, 2021.

StreamingCNN

➔ Non-overlapping = **Incorrect**

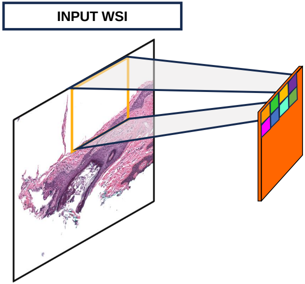

Pinckaers, Hans, et al. "Detection of prostate cancer in whole-slide images through end-to-end training with image-level labels." IEEE TMI, 2021.

StreamingCNN

➔ Overlapping = **Correct**

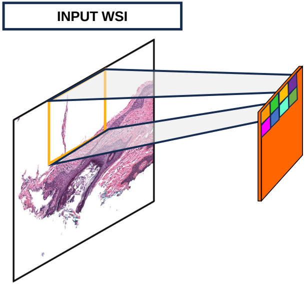

Pinckaers, Hans, et al. "Detection of prostate cancer in whole-slide images through end-to-end training with image-level labels." IEEE TMI, 2021.

StreamingCNN

→ **Repeat** and **Save**!

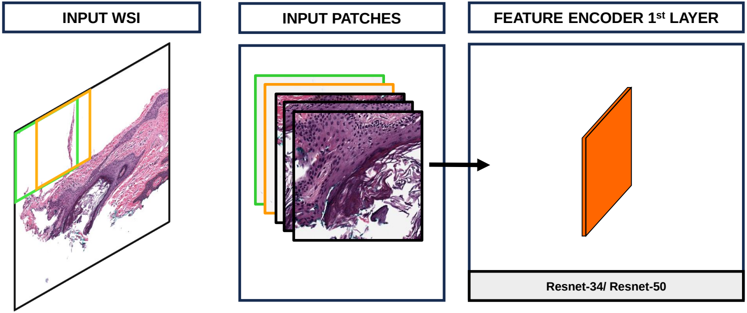

Pinckaers, Hans, et al. "Detection of prostate cancer in whole-slide images through end-to-end training with image-level labels." IEEE TMI, 2021.

StreamingCNN

➔ **Use** first feature map to get second one!

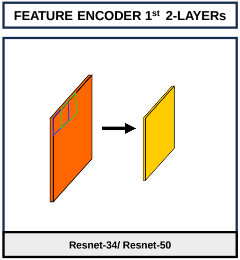

Pinckaers, Hans, et al. "Detection of prostate cancer in whole-slide images through end-to-end training with image-level labels." IEEE TMI, 2021.

StreamingCNN

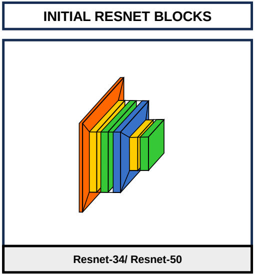

Pinckaers, Hans, et al. "Detection of prostate cancer in whole-slide images through end-to-end training with image-level labels." IEEE TMI, 2021.

StreamingCNN

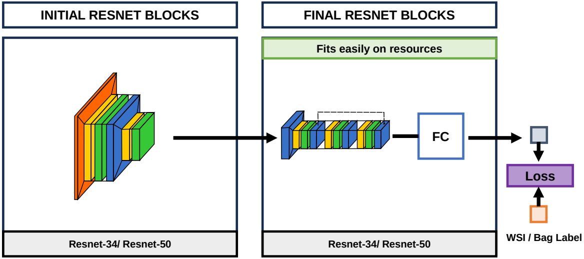

Pinckaers, Hans, et al. "Detection of prostate cancer in whole-slide images through end-to-end training with image-level labels." IEEE TMI, 2021.

StreamingCNN – Backward Pass

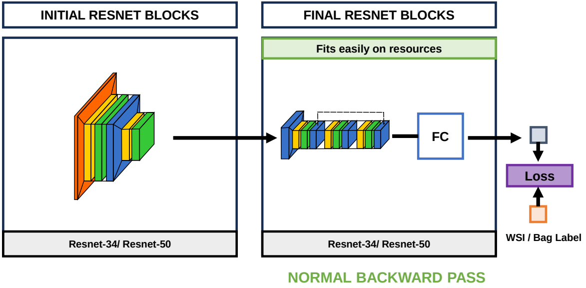

Pinckaers, Hans, et al. "Detection of prostate cancer in whole-slide images through end-to-end training with image-level labels." IEEE TMI, 2021.

StreamingCNN – Backward Pass

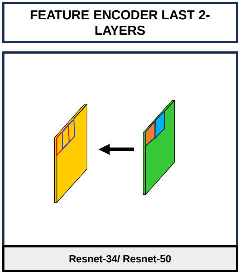

Pinckaers, Hans, et al. "Detection of prostate cancer in whole-slide images through end-to-end training with image-level labels." IEEE TMI, 2021.

StreamingCNN – Backward Pass

➔ Don't reset, instead **ADD**.

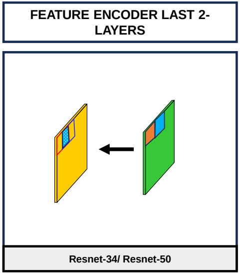

Pinckaers, Hans, et al. "Detection of prostate cancer in whole-slide images through end-to-end training with image-level labels." IEEE TMI, 2021.