SFU
SIMON FRASER UNIVERSITY

Predicting Cancer with a Recurrent Visual Attention Model for Histopathology Images

MICCAI 2018 Granada SPAIN
21st INTERNATIONAL CONFERENCE ON MEDICAL IMAGE COMPUTING & COMPUTER ASSISTED INTERVENTION
September 16-20, 2018
GRANADA CONFERENCE CENTRE · GRANADA / SPAIN

Aïcha BenTaieb and Ghassan Hamarneh
{abentaie@sfu.ca, hamarneh@sfu.ca}

Identifying Cancer In Whole Slide Images (WSI)

- Cancers are primarily diagnosed from the analysis of tissue biopsies
- Identifying abnormal areas in multi giga-pixel WSI is an open challenge for both automatic systems and experts diagnosis [1]
- Patch-based techniques for WSI analysis are commonly used but often rely on manual annotations that are not always easily available
- This work investigates whether an automatic diagnosis system can learn **where** to look for abnormalities in WSI

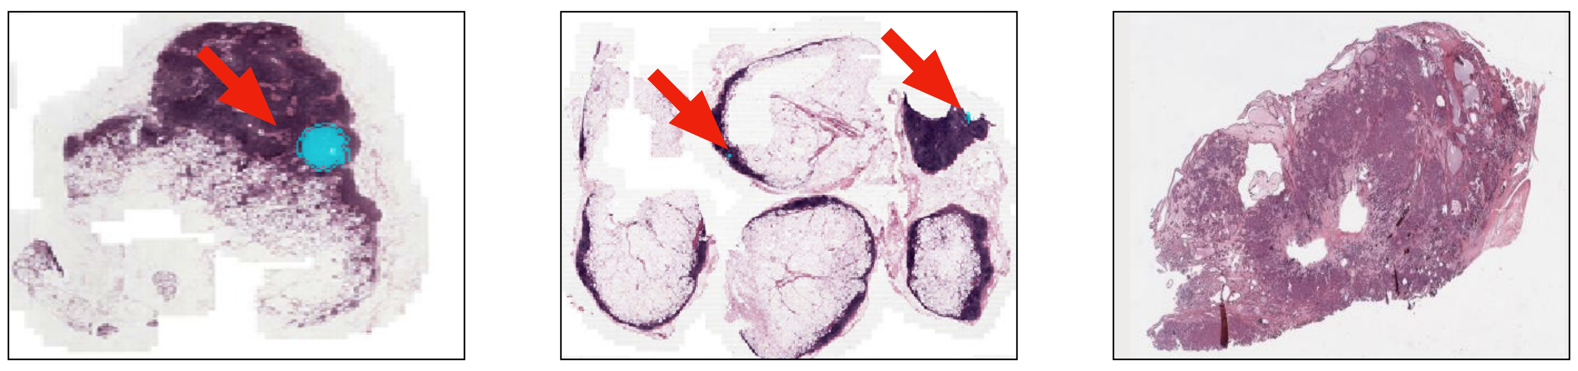

Lymph Node Macro Metastasis — Lymph Node Micro Metastasis — Clear Cell Ovarian Carcinoma Unknown Location

Recurrent Visual Attention Model

- Desired Properties:
  - Requires a limited set of patches to infer a tissue slide label
  - Automatically identifies locations and scales at which to extract patches
  - Integrates global context and memory of previously observed tissue areas

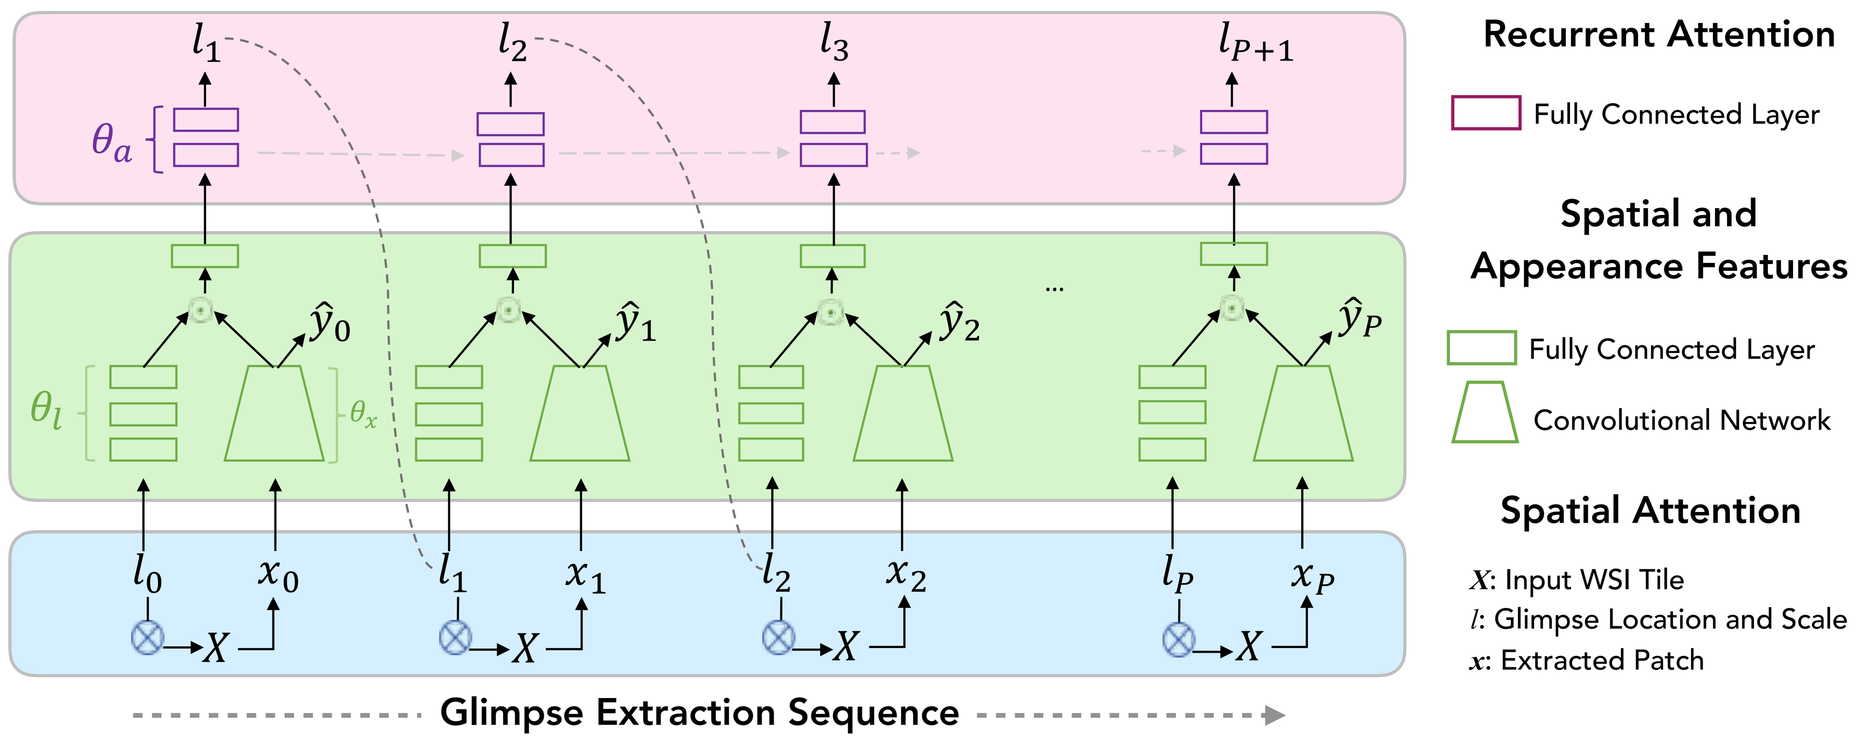

Spatial Attention Mechanism

- The attention mechanism applies two grids (one per image axis) of 2D Gaussian filters $\boldsymbol{A}$ (parameterized by l) to an input image $\boldsymbol{X}$

$$x_p = A_p^x X A_p^{y\,T}$$

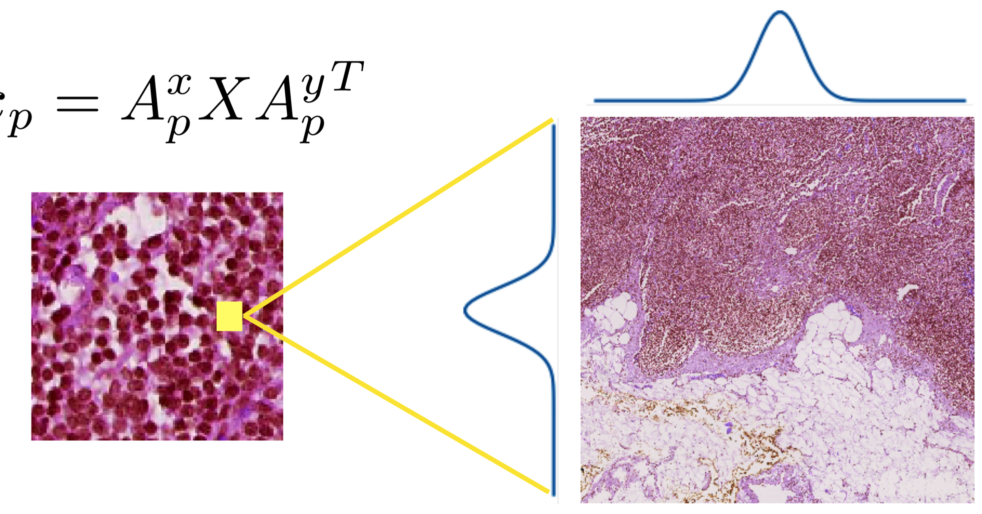

Weakly Supervised Training Objective Function

- Given a dataset D of training WSI and slide-level class labels, the model is trained with the following regularized loss function:

$$\mathcal{L}(\mathcal{D};\theta) = \mathcal{L}_c(\mathcal{D};\theta) + \mathcal{L}_p(\mathcal{D};\theta) + \mathcal{L}_a(\mathcal{D};\theta) + \mathcal{L}_l(\mathcal{D};\theta)$$

WSI-Level Cross Entropy Classification Loss

$$\mathcal{L}_c(\mathcal{D};\theta) = \sum_{i=1}^{N} \log Q(\hat{Y} = Y^{(i)} | x_{[1:P]}^{(i)};\theta)$$

Glimpse-Level Cross Entropy Classification Loss

$$\mathcal{L}_p(\mathcal{D};\theta) = \sum_{i=1}^{N}\sum_{p=1}^{P} \log Q(\hat{y}_p = Y^{(i)} | x_p^{(i)};\theta)$$

Discriminative Attention

$$\mathcal{L}_l(\mathcal{D};\theta) = \gamma \sum_{i=1}^{N}\sum_{p=1}^{P} \exp(-||l_p - l_{p+1}||_1)$$

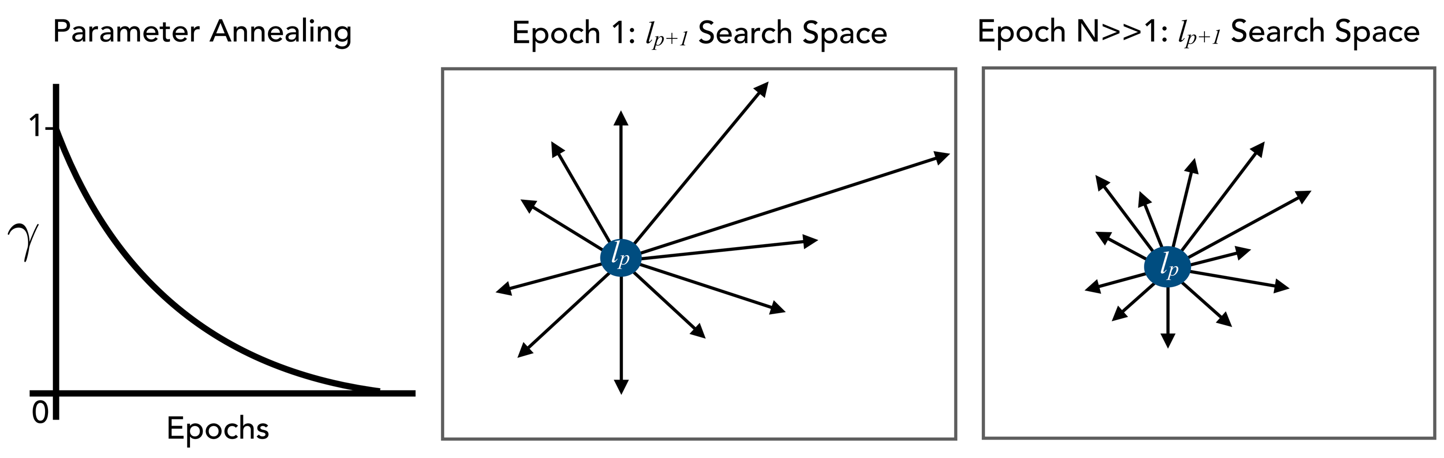

Selective Exploration

$$\mathcal{L}_a(\mathcal{D};\theta) = -\sum_{i=1}^{N}\sum_{p=2}^{P} Q(y_p^{(i)}|x_p^{(i)};\theta) - \left(\frac{1}{p-1}\sum_{k=1}^{p-1} Q(y_k^{(i)}|x_k^{(i)};\theta)\right)$$

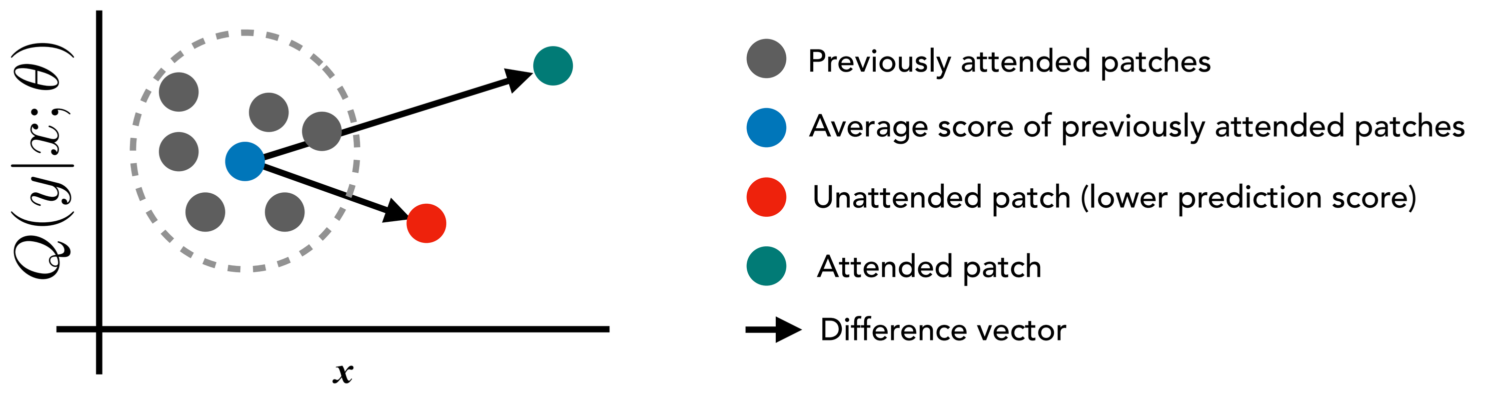

Experiments & Results on Camelyon16 Challenge Dataset [2]

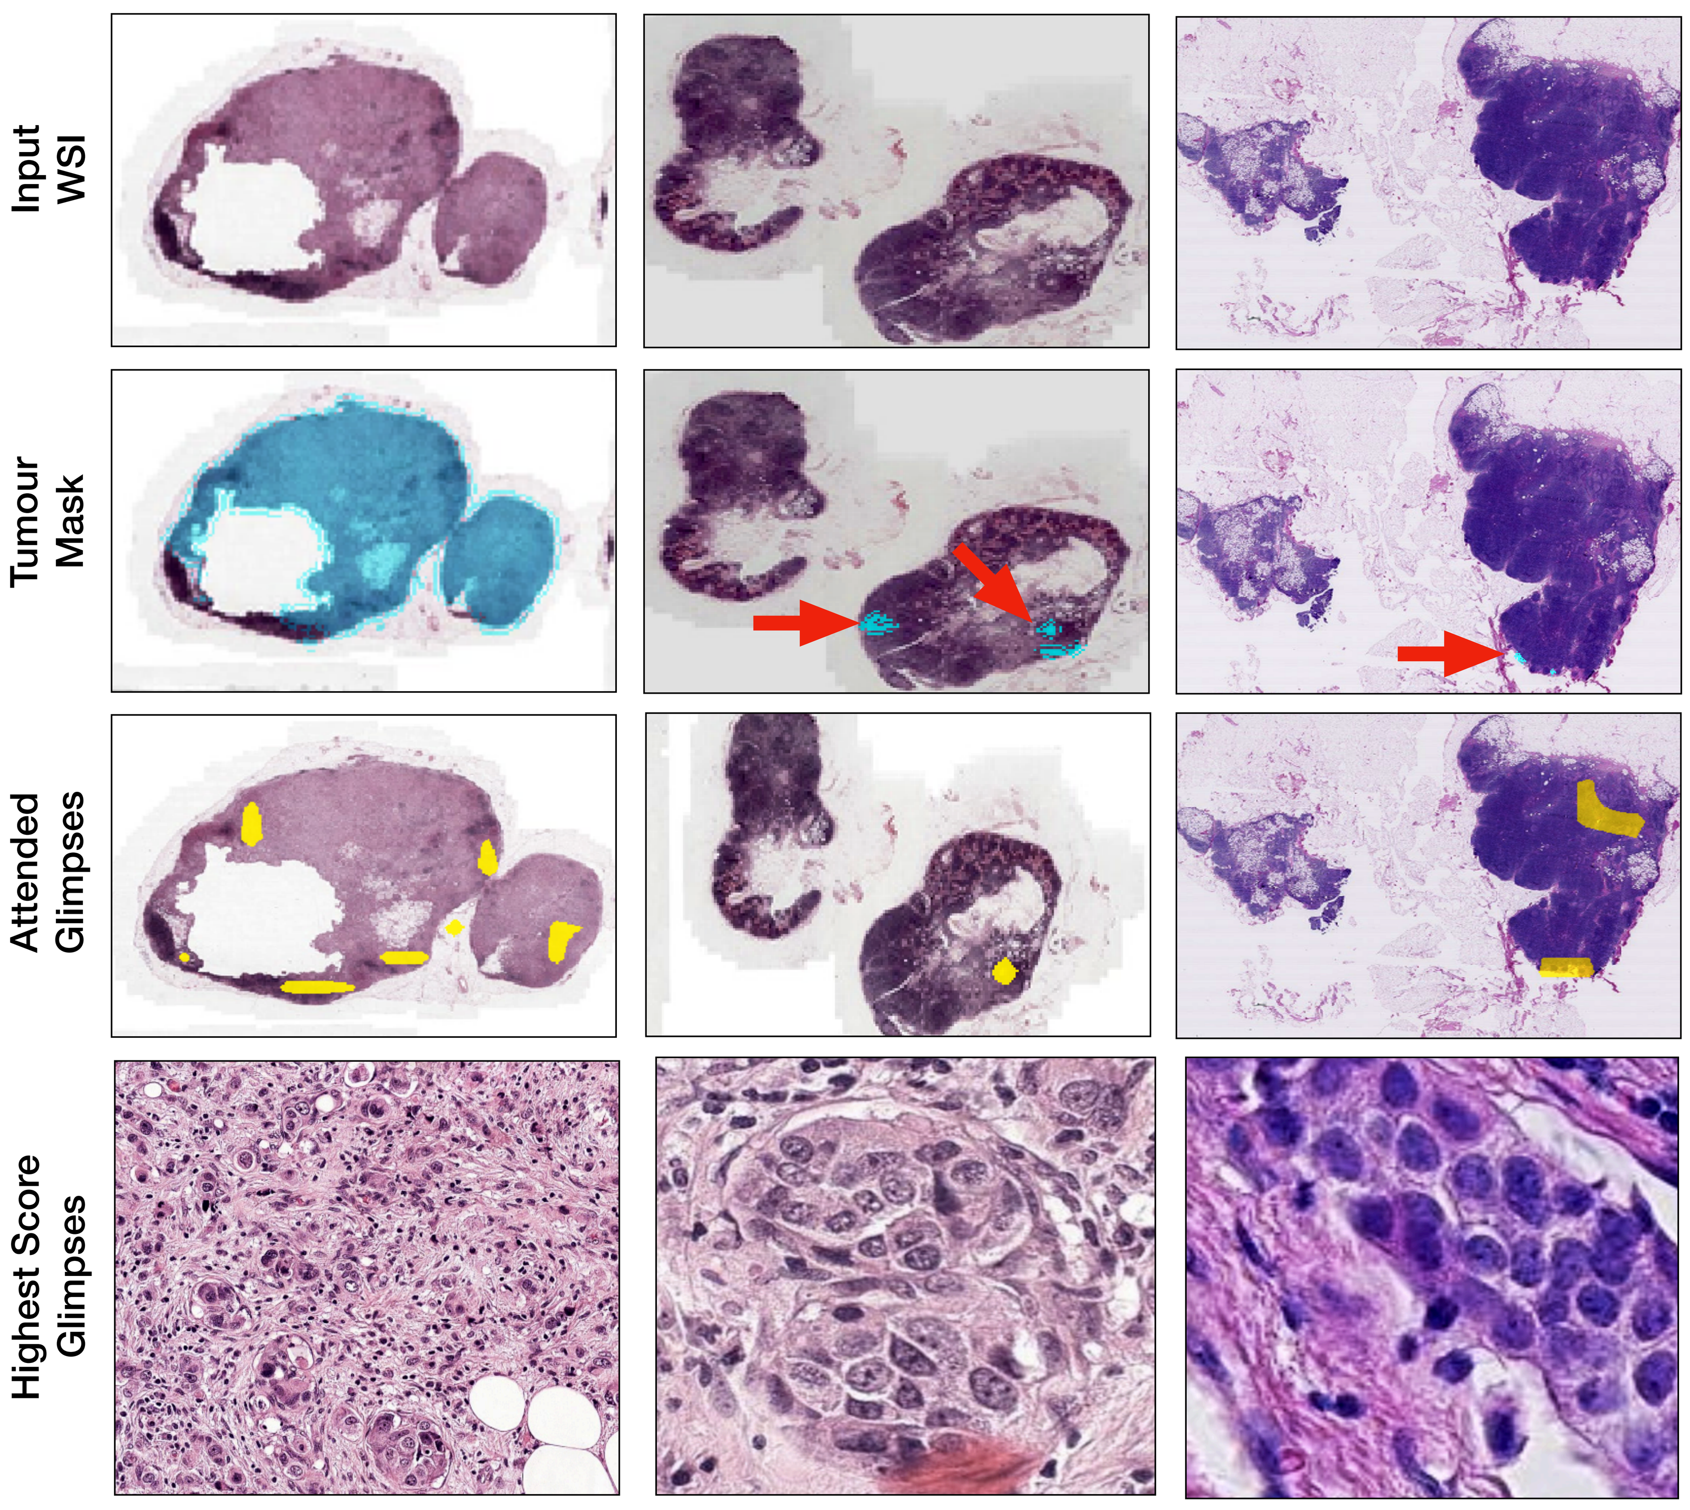

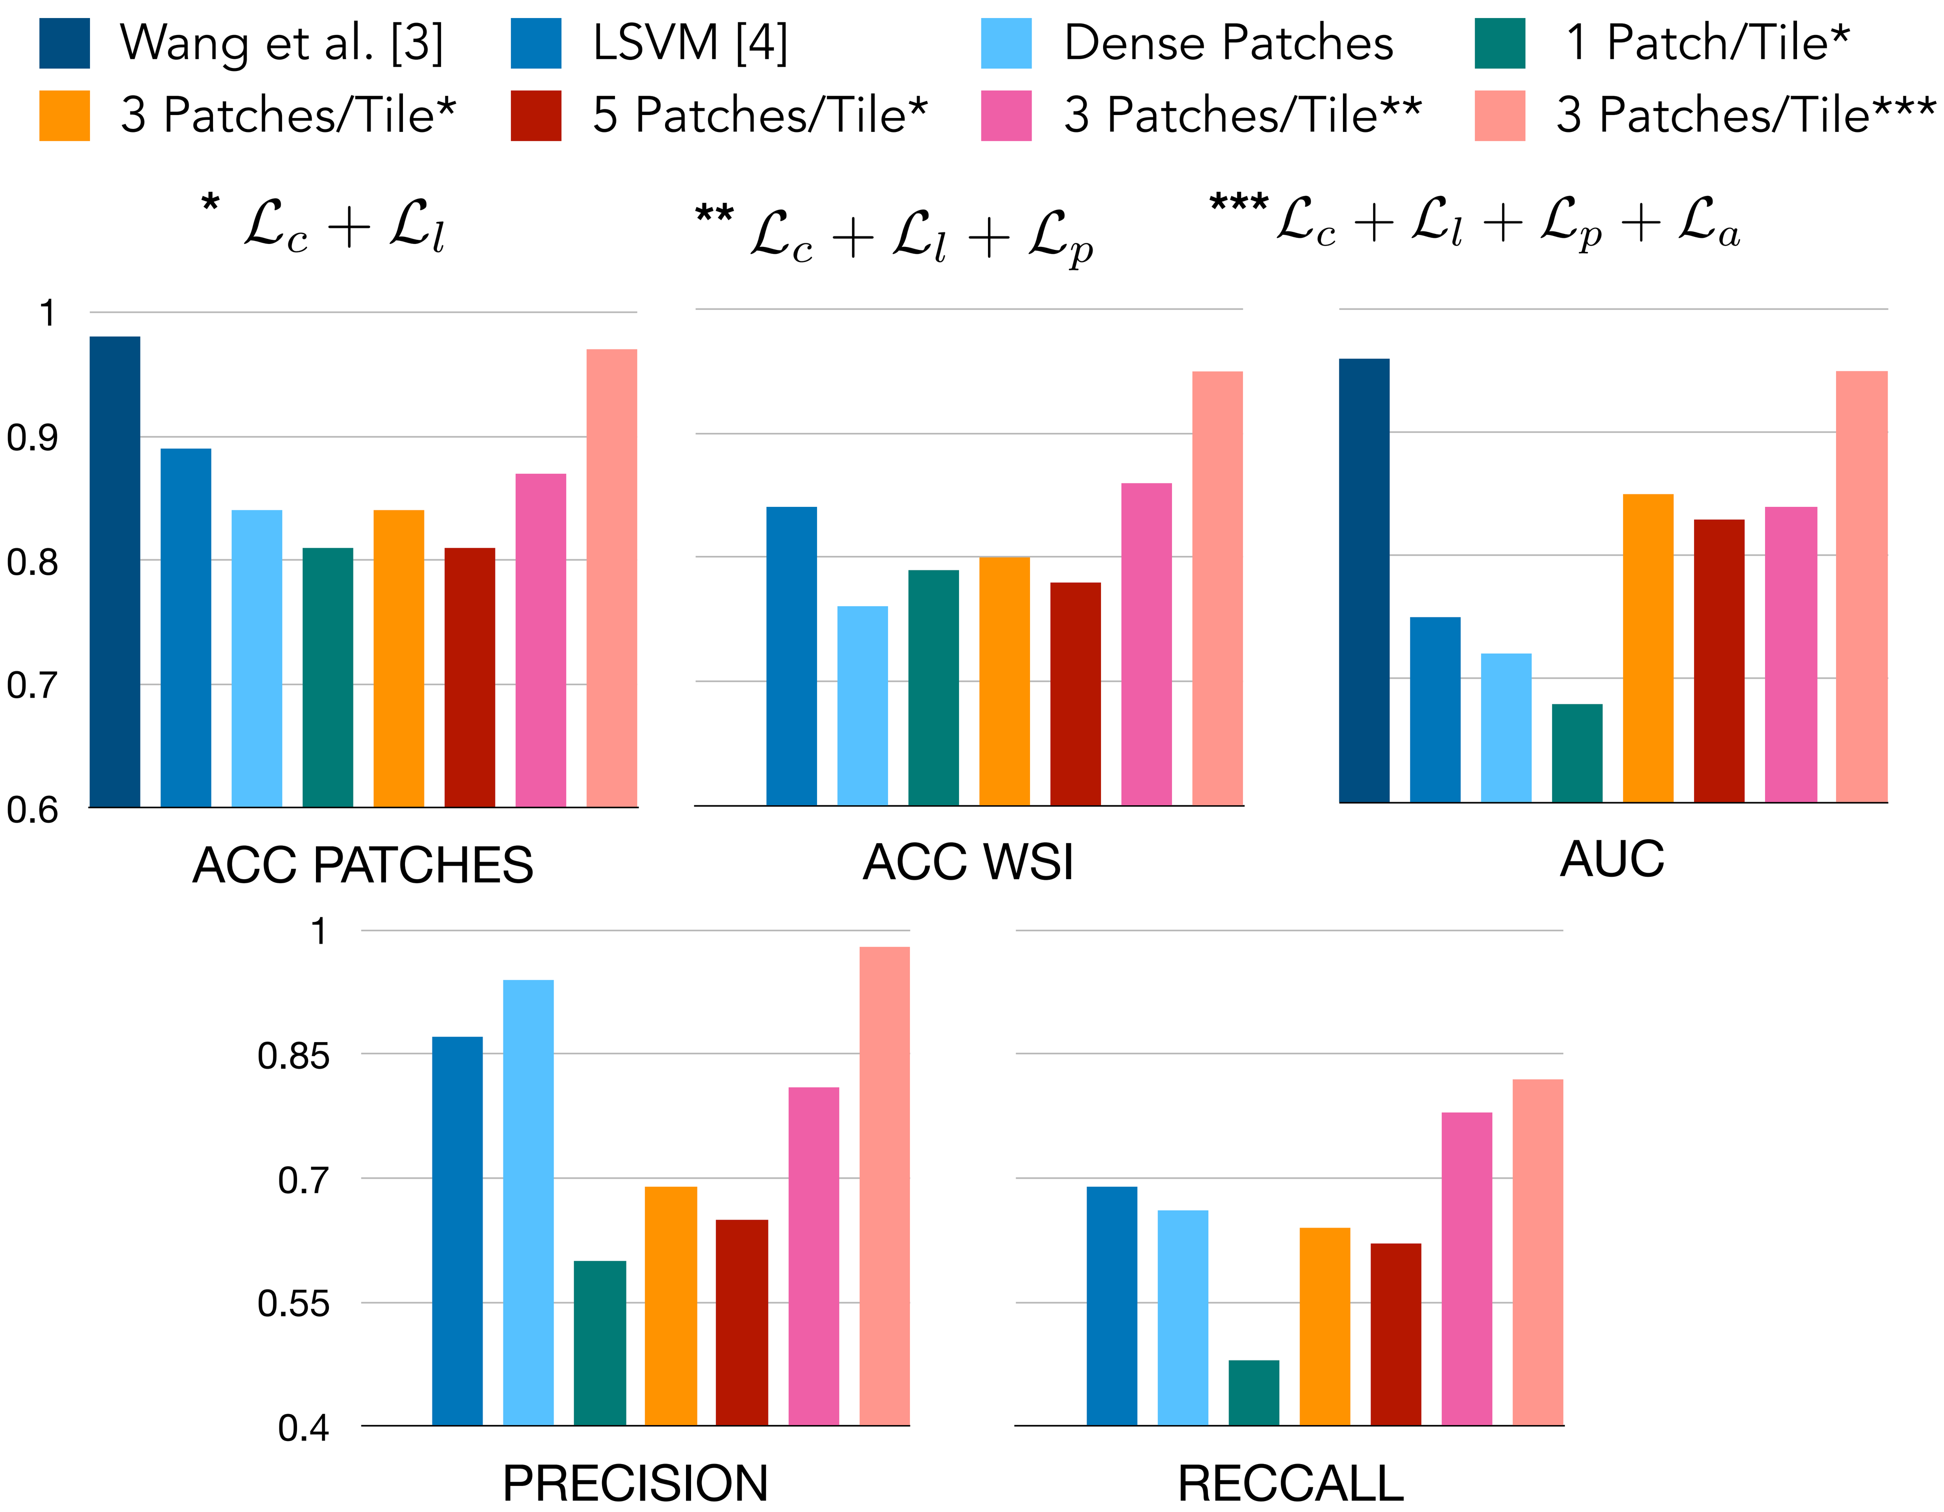

Code: www.sfu.ca/~abentaie

[1] Brunye et al. PloS One, 2014
[2] Golden et al. Jama, 2017
[3] Wang et al. arXiv, 2016
[4] BenTaieb et al. MedIA 2017