Layer Matching Domain Adaptation for Microscopy Image Translation

Vishrut Goyal, Michael Liu, Andrew Bai, Neil Lin, Cho-Jui Hsieh

Motivation

Image translation models are applied in stem-cell microscopy to predict marker protein levels without invasive staining.

Real world application is limited because varying microscope parameters create a catastrophic **domain shift.**

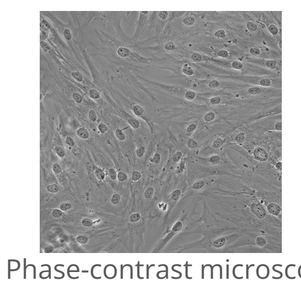

Phase-contrast microscopy

Original image translation model

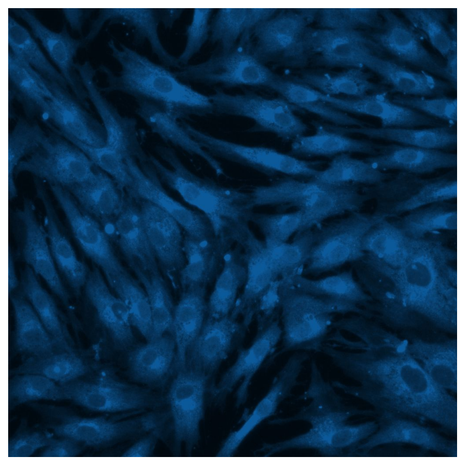

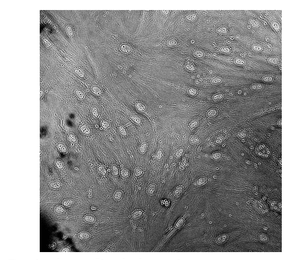

Bright-field microscopy

Goal of adapted model

Immunofluorescence stained for CD29

Problem Setting

Generate the same output image from an input image taken with a different microscope. No ground truth output images are available during adaptation
Assumptions:

1. **The underlying information is equivalent, i.e, same cells**
2. **The output image does not get shifted**

Main Result

We find that using the hidden layers of the image translation model as a shared feature space outperforms matching the final image output.

Adversarial Domain Adaptation

Domain adaptation allows altering a model trained on for a new input without retraining. The simplest approach to domain adaptation is to directly translate the new input to the old input type.

To produce an indistinguishable output, a naive method would be to train the new model so that a discriminator cannot tell the outputs apart. This technique is known as **adversarial discriminative domain adaptation (ADDA).**

The objective optimized for the adapted model G' is equivalent to that of a generative adversarial network (GAN), where ground truth is replaced with the final output of the original model G.

$$\mathcal{L}_{\text{adv}} = \mathop{\mathbb{E}}_{x \sim X}[\log(D(G(x)))] - \mathop{\mathbb{E}}_{x' \sim X'}[\log(D(G'(x', \theta)))]$$

Nucleus Identification Task

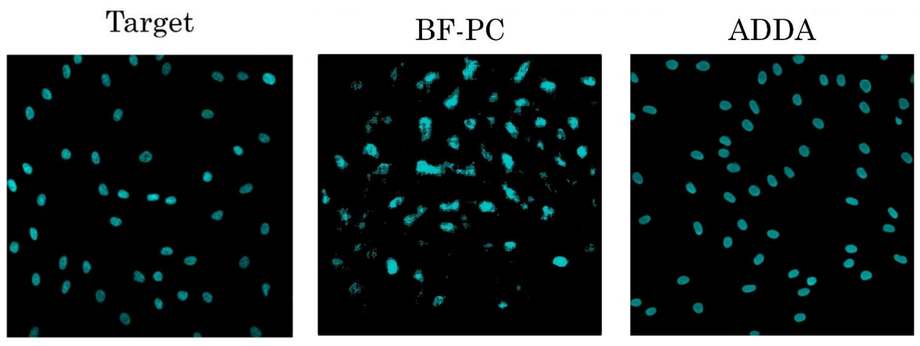

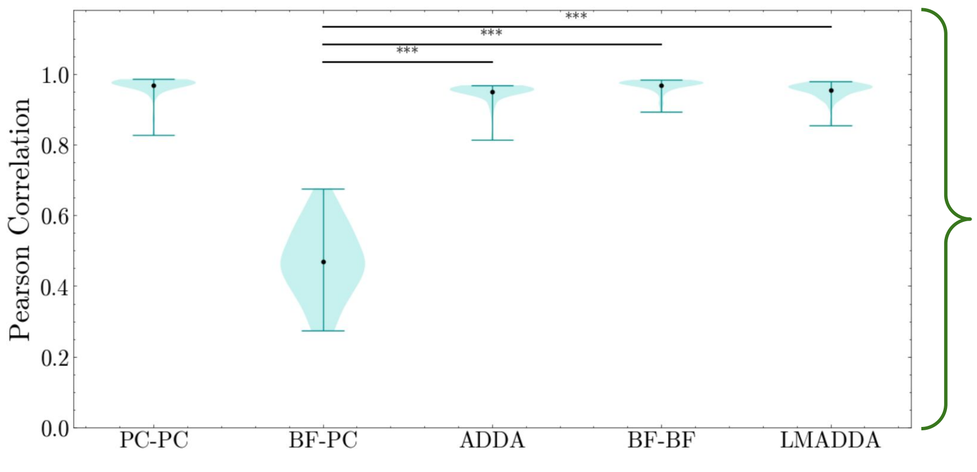

Adaptation recovers nearly all the loss from the domain shift

Layer-Matching ADDA

Matching the final output is less successful for more complex translation tasks and domain shifts. It is difficult to undo a complex shift based only on the style information in the output images.

Our key finding is that the extracted features remain the same for both domains. Therefore adaptation could be limited to the first few feature-extracting layers. Once the features are matched, the remaining model can be used as-is.

The adversarial task can be modified to apply to a given hidden layer:

$$G'_l : X' \to G_l(X)$$

Intuition

By matching the target and source domains at the information-dense hidden layers, we can maximize signal-to-noise ratio while adapting the model.

CD29 Marker Task

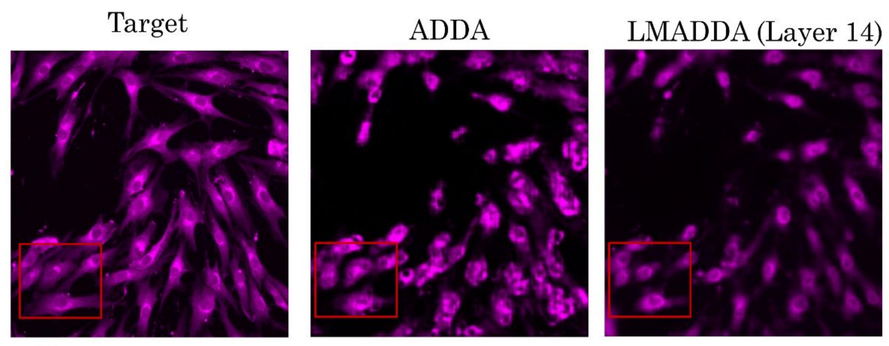

Crucial "halo" feature preserved by feature-level matching

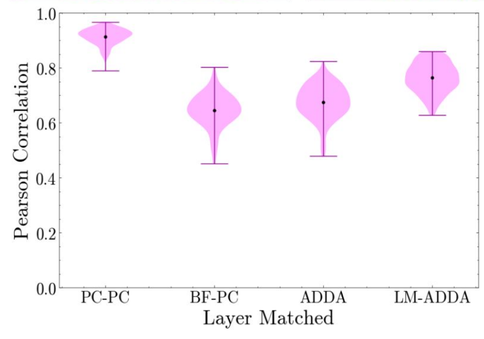

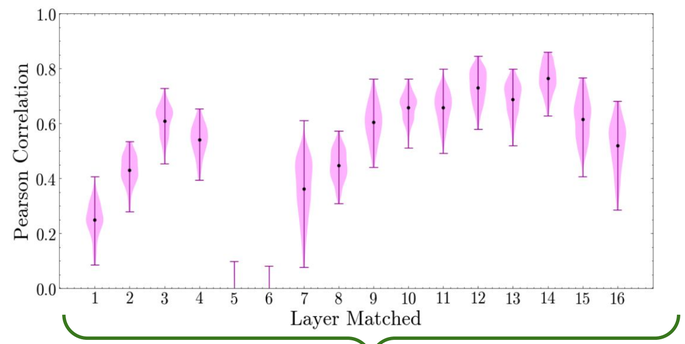

Quality of adaptation peaks before the final layer, and declines closer to pixel-level

Overexposure Domain Shift

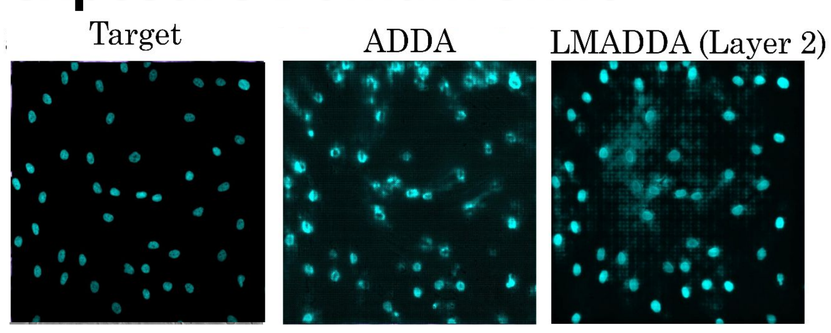

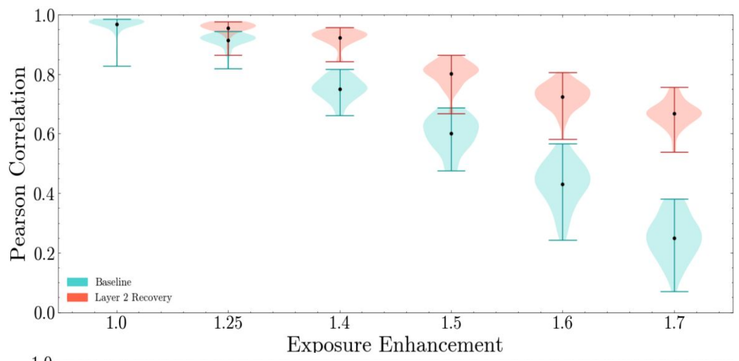

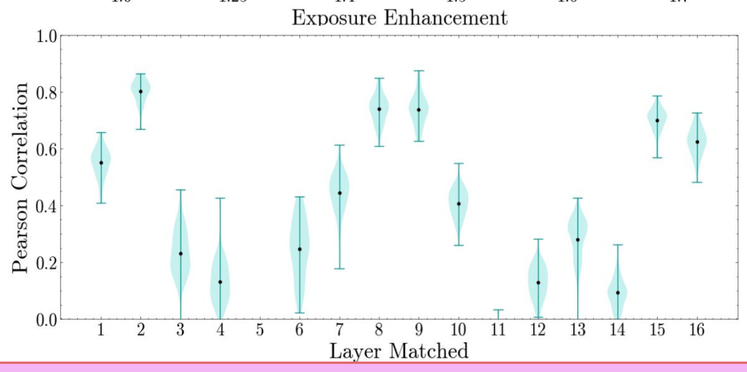

Simpler domain shift best undone by early layers

Conclusion

Layer Matching-ADDA is able to successfully adapt image translation networks to multiple microscopy inputs, across tasks, even for complex tasks where directly matching the final output fails.

Further Work

Theoretically determine optimal layer for matching.