Class project details

- Can work alone or in a group (up to 4 people), required effort will scale with # of people
- Select a “base” dataset (online, or from a list I’ll make)
- Simulate parameters of a physical (imaging) system with base dataset
- Train deep neural net with simulated dataset
- Report results

Class project details

- Can work alone or in a group (up to 4 people), required effort will scale with # of people
- Select a “base” dataset (online, or from a list I’ll make)
- Simulate parameters of a physical (imaging) system with base dataset
- Train deep neural net with simulated dataset
- Report results

What you’ll need to submit:

1) The project’s source code

2) A short research-style paper (3 pages minimum, 5 pages maximum) that includes an introduction, results, a discussion section, references and at least 2 figures

3) A completed web template containing the main results from the research paper

4) An 8-minute presentation that each student will deliver to the class

Example project topics:

Can we design a new lens/transducer/antenna shape to improve classification of X?

What is the tradeoff between image resolution and classification accuracy for X?

Can we determine an optimal set of colors to improve fluorophore distinguishability?

If we capture 2 images that are overlapped on one sensor, what is the best way to pre-blur them to then be able to tell them apart? Or to be able to classify them together?

If we just had a few sensors, how should be arrange them e.g. a mask to be able to predict the position of X?

Is there some optimal shift-variant blur that we can to use for a particular task?

Or, given a shift-variant PSF image, can we establish a good deconvolution using locally connected layers?

What is the optimal way to layout filters on a sensor to capture a color image for classification? Or an HDR image?

HDR image generation with filters over pixels – what is optimal design?

What if we could make a sensor with different sized pixels – how should they be laid out to achieve the best X?

Class project – what are the first steps?

1. Think about it!
2. Discuss with your friends/others in the class (feel free to use Slack!)
3. Schedule a short 15 meeting with me:
    - Friday 3/1, 3:30pm – 6:00pm
    - Next Monday 3/4, 1:00pm – 3:30pm
    - Next Wednesday 3/6, 10:00am – 1:00pm
4. Start to write-up a proposal
    - General aim: 1 paragraph with specification of physical layer
    - Discussion: (a) data source(s), (b) expected simulations, (c) expected CNN, (d) quantitative analysis of physical layer (comparison, plot, etc).

- Project proposal due date: **Thursday March 7, 2019**
- Revised project proposal due date: **Tuesday March 19, 2019**

Example project topics:

Can we design a new lens/transducer/antenna shape to improve classification of X?

What is the tradeoff between image resolution and accuracy for X (classification, segmentation, etc.)? What if we had access to n low-resolution cameras – how might we position them to get the best performance?

Can we determine an optimal set of colors to improve fluorophore distinguishability?

If we capture 2 images that are overlapped on one sensor, what is the best way to pre-blur them to then be able to tell them apart? Or to be able to classify them together?

If we just had a few sensors, how should be arrange them e.g. a mask to be able to predict the position of X?

Is there some optimal shift-variant blur that we can to use for a particular task?

Or, given a shift-variant PSF image, can we establish a good deconvolution using locally connected layers?

What is the optimal way to layout filters on a sensor to capture a color image for classification? Or an HDR image?

HDR image generation with filters over pixels – what is optimal design?

What if we could make a sensor with different sized pixels – how should they be laid out to achieve the best X?

What is the tradeoff between image resolution and accuracy for image segmentation? What if we had access to n low-resolution cameras – how might we position them to get the best performance?

I propose to simulate the classification performance of a new type of microscope, which will have 3 different lenses and sensors. Each lens and sensor will capture an image of a flat object from a unique angular perspective, and the image classification will be performed with all of the data. The physical parameter that I will optimize is the angle of tilt of each lens with respect to the object to maximize classification accuracy.

What is the tradeoff between image resolution and accuracy for image segmentation? What if we had access to n low-resolution cameras – how might we position them to get the best performance?

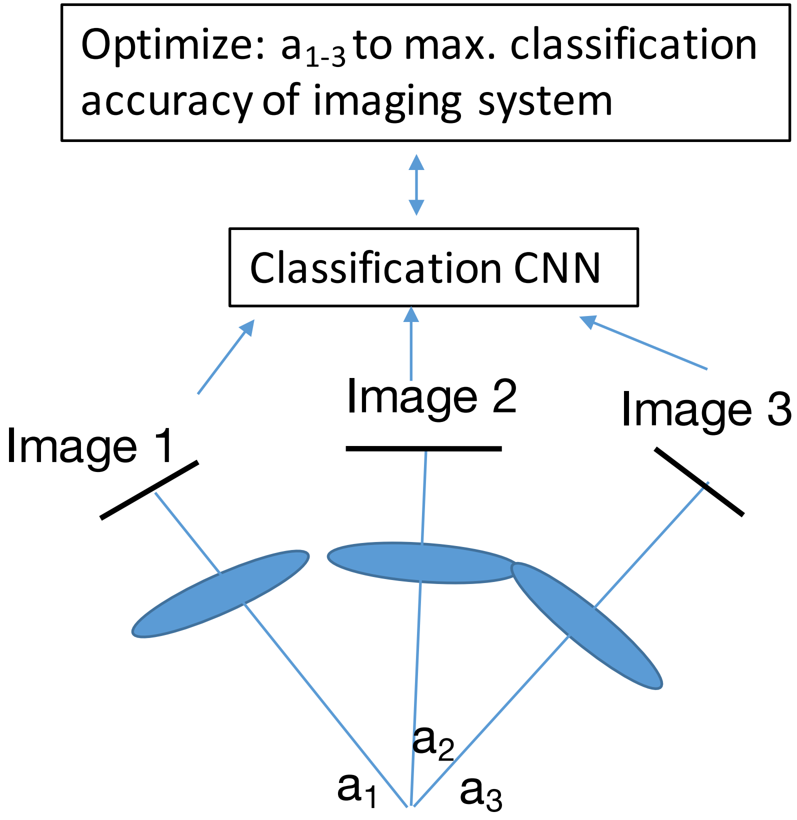

I propose to simulate the classification performance of a new type of microscope, which will have 3 different lenses and sensors. Each lens and sensor will capture an image of a flat object from a unique angular perspective, and the image classification will be performed with all of the data. The physical parameter that I will optimize is the angle of tilt of each lens with respect to the object to maximize classification accuracy.

What is the tradeoff between image resolution and accuracy for image segmentation? What if we had access to n low-resolution cameras – how might we position them to get the best performance?

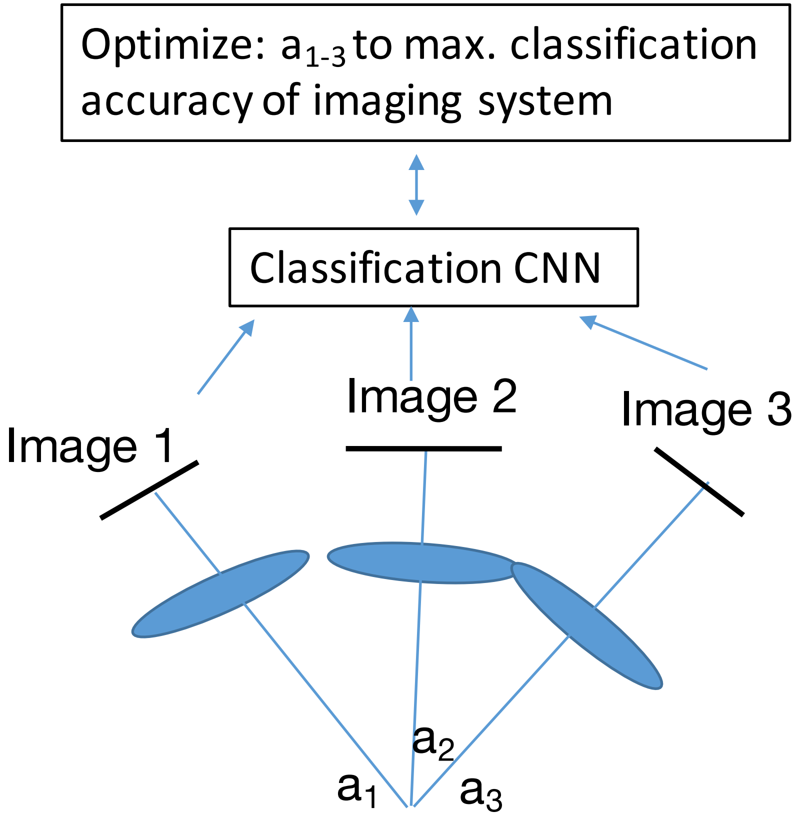

I propose to simulate the classification performance of a new type of microscope, which will have 3 different lenses and sensors. Each lens and sensor will capture an image of a flat object from a unique angular perspective, and the image classification will be performed with all of the data. The physical parameter that I will optimize is the angle of tilt of each lens with respect to the object to maximize classification accuracy.

Dataset: 12,500 images of 4 types of blood cell
https://www.kaggle.com/paultimothymooney/blood-cells

Simulation: Treat each image as a thin 2D object and is coherently illuminated. Assume each camera captures a unique component of the object spectrum, which will vary as a function of $a_{1\text{-}3}$. Start by neglecting size and shape of each camera.

CNN: Digital layer: Alexnet. Physical layer: simulate object spectrum, sample object spectrum re-centered by angle $a_{1\text{-}3}$ (which are weight variables), form images and classify them together

Quantitative analysis: I will plot classification performance as a function of the number of allowed cameras for a fixed CNN architecture, and will also compare the classification performance to the case of a single image