When spectral imaging meets machine learning

david.rousseau@univ-angers.fr

David Rousseau, Prof. Bioimaging, Université d’Angers, France

Who am I ?…. data scientist like all of us

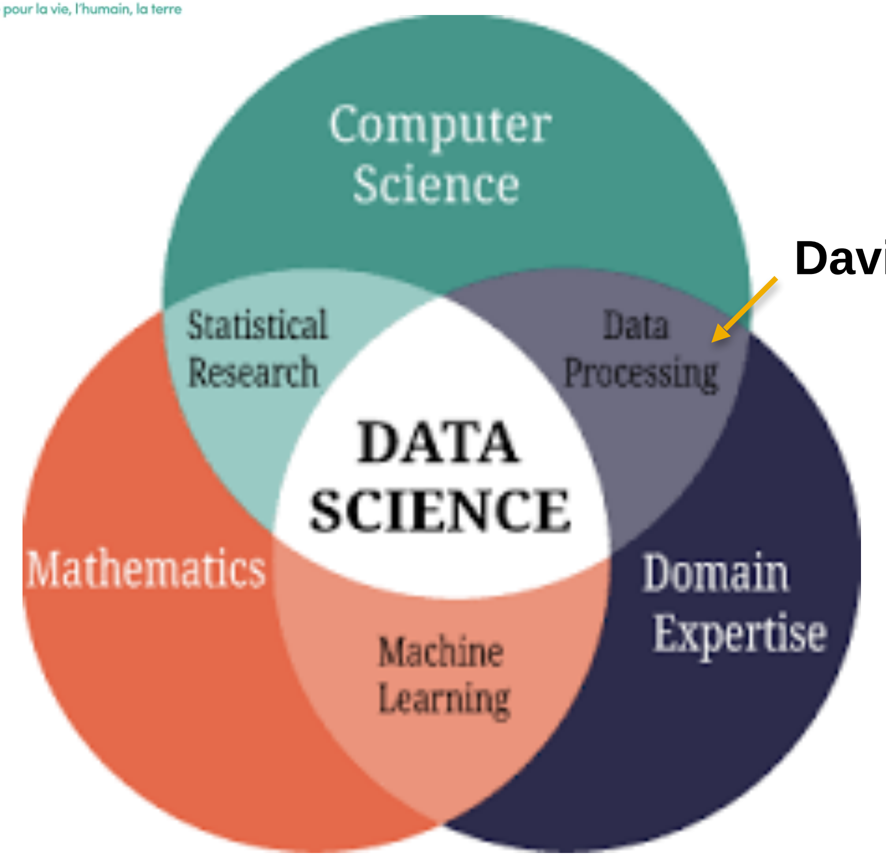

David Rousseau

What do I do ?

Quality and Health of horticultural crops

Ornamentals

Fruits & Vegetables

Seeds

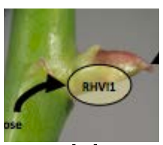

Architecture and flowering

Quality & conservation

Pathogens

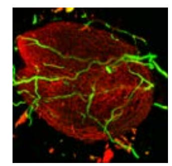

Conservation & seed quality

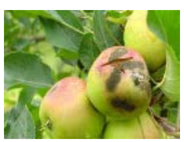

Genetic diversity & Breeding

Evolutionary ecology of pathogens

Biotic/abiotic stresses and germination

Resistance sustainability

Epigenetics

Bioinformatics

Metagenomics

Disease emergence

HT Phenotyping

Urban Agronomy

91 Scientists/106 Technicians & Engineers/35 PhD St.: 235 persons, 13 teams

Interested in deep learning

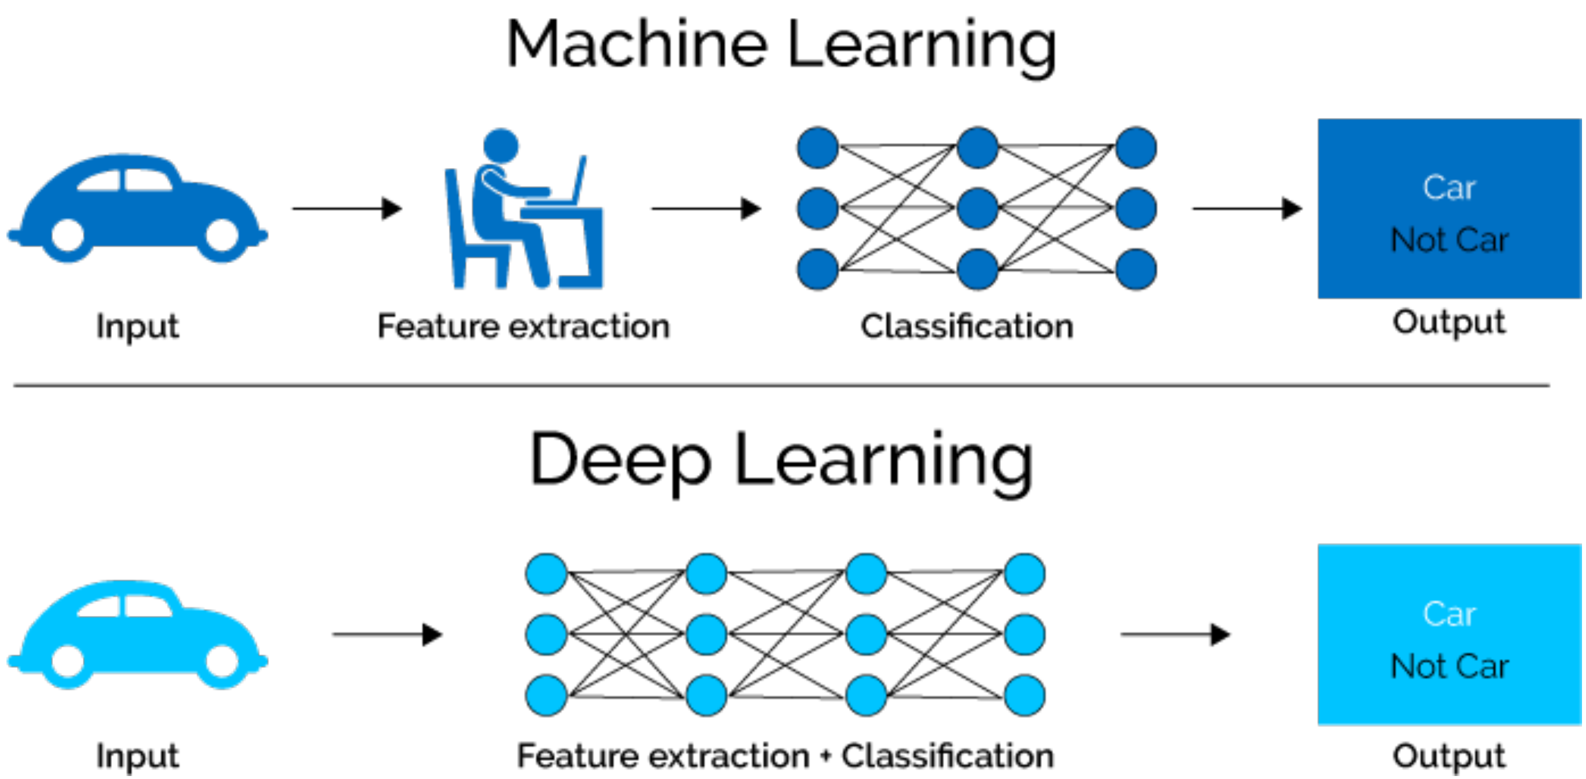

=> Specially adapted to plant imaging (large cohorts, selfocclusion, few ethical issues, multiscale, …)

Both in research and teaching

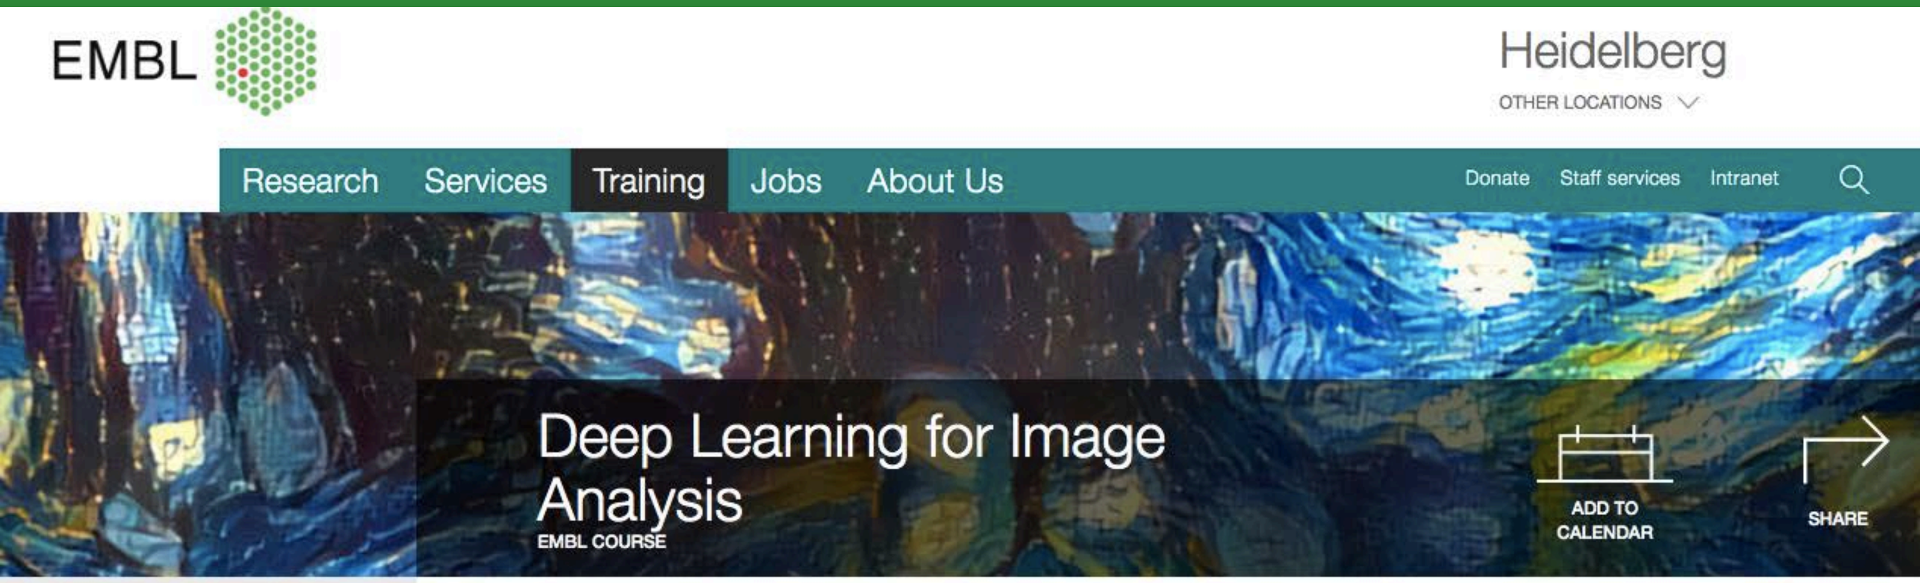

Next SESSION JUNE 2023

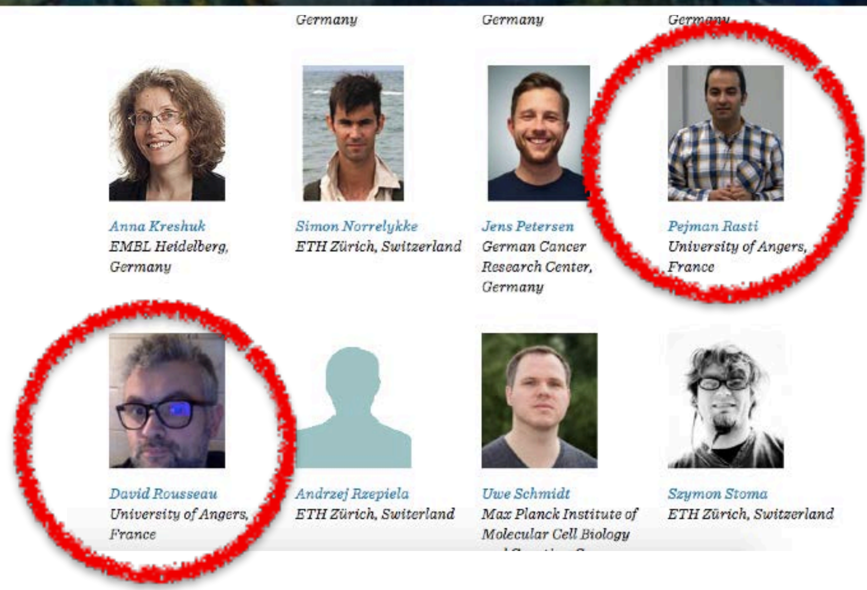

Spectral imaging

Type of sensors	Nb Channel	Filtrers	Imaging systems
Gray	**1**	**Open choice (visible / NIR)**	  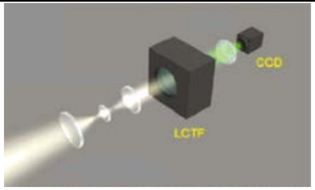
Color	3	Imposed (R V B) or on demand	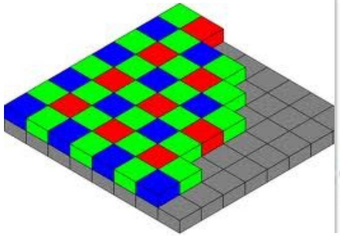 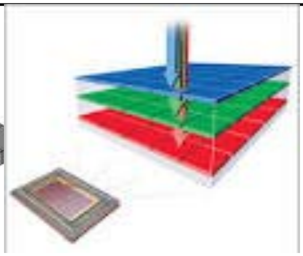 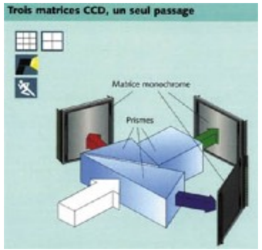
Multispectral	2 - 10	Open choice (visible / NIR)	 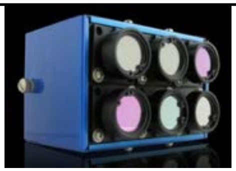 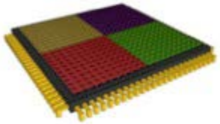
Hyperspectral	**Tens to hundreds**		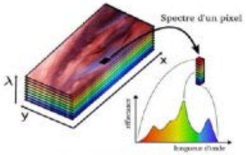

Spectral imaging + machine learning

Cost

Type of sensors	Nb Channel	Filtrers	Imaging systems
Gray	1	Open choice (visible / NIR)	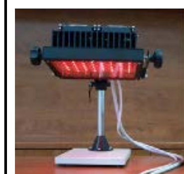  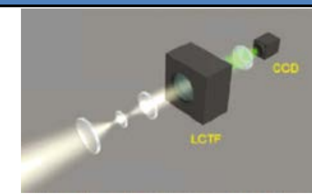
Color	3	Imposed (R V B) or on demand	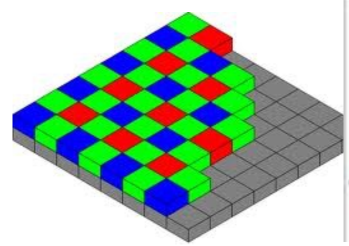 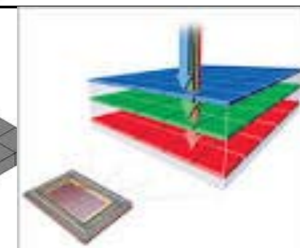 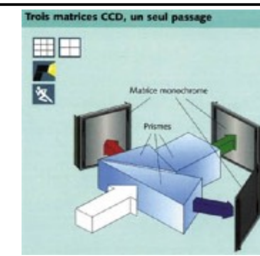
Multispectral	2 - 10	Open choice (visible / NIR)	  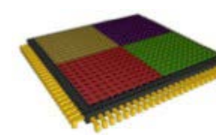
Hyperspectral	**Tens to hundreds**		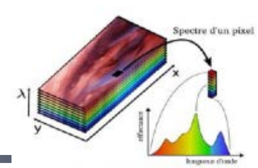

Guide line of the talk: 3 use cases

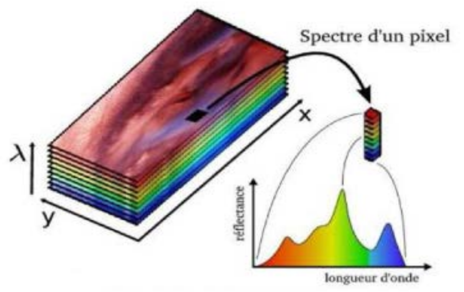

Building cost-effective spectral imaging with Statistical learning

Benoit, Landry, Romain Benoit, Étienne Belin, Rodolphe Vadaine, Didier Demilly, François Chapeau-Blondeau, and David Rousseau. "On the value of the Kullback–Leibler divergence for cost-effective spectral imaging of plants by optimal selection of wavebands." *Machine Vision and Applications* 27, no. 5 (2016): 625-635.

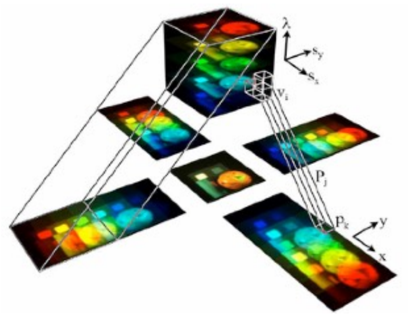

Low-cost spectro-imaging & compressed learning

Douarre, C., Crispim-Junior, C. F., Gelibert, A., Germain, G., Tougne, L., & Rousseau, D. (2021). CTIS-Net: A Neural Network Architecture for Compressed Learning Based on Computed Tomography Imaging Spectrometers. *IEEE Transactions on Computational Imaging*, *7*, 572-583.

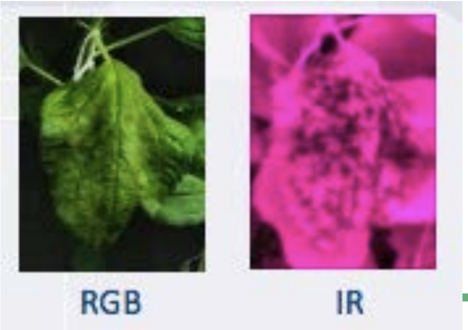

Lowering the cost of annotation in machine learning

Douarre, Clément, Carlos F. Crispim-Junior, Anthony Gelibert, Laure Tougne, and David Rousseau. "Novel data augmentation strategies to boost supervised segmentation of plant disease." *Computers and Electronics in Agriculture* 165 (2019): 104967.

Cost effective spectral imaging

Cost

Type of sensors	Nb Channel	Filtrers	Imaging systems		
Gray	1	Open choice (visible / NIR)	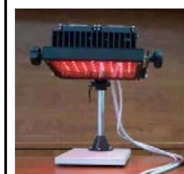		
Color	3	Imposed (R V B) or on demand	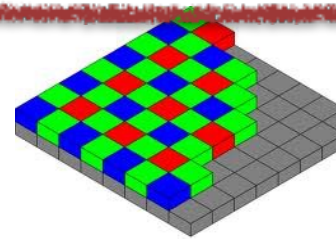	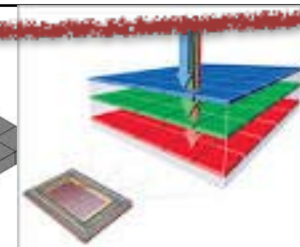	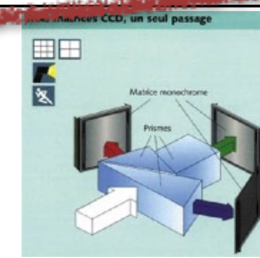
Multispectral	2 - 10	Open choice (visible / NIR)	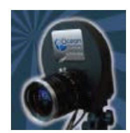		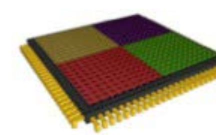
Hyperspectral	**Tens to hundreds**		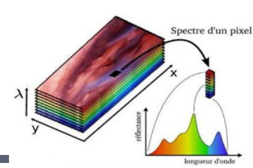		

Selecting wavebands with Shannon

Input X a spectrum between lambda min and lambda max

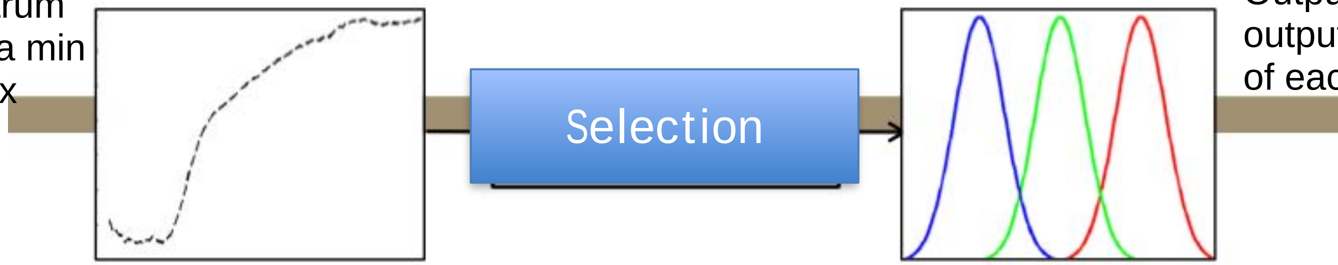

Output Y : M values at t
output as the integrated
of each spectral band

fi spectral response of each band ; S(lambda) input spectrum

$$P(Y=i) = \int_{\lambda_{\min}}^{\lambda_{\max}} f_i(\lambda) S(\lambda) d\lambda$$ Probability to find a photon in spectral band i

$$P_{\det}(\lambda) = \sum_{i=1} f_i(\lambda) \leq 1 \qquad P_{\mathrm{lost}}(\lambda) = 1 - P_{\det}(\lambda) = 1 - \sum_{i=1}^{M} f_i(\lambda)$$

$$P(Y=0) = \int_{\lambda_{\min}}^{\lambda_{\max}} P_{\mathrm{lost}}(\lambda) S(\lambda) d\lambda$$ Probability to have a photon lost

$$I(X;Y) = H(Y) - H(Y|X)$$

$$H(Y) = -\sum_{i=0}^{3} P(Y=i) \log[P(Y=i)] \qquad H(Y|X) = \int_{\lambda_{\min}}^{\lambda_{\max}} H(Y|X=\lambda) S(\lambda) d\lambda$$

$$H(Y|X=\lambda) = -\sum_{i=1} f_i(\lambda) \log[f_i(\lambda)] - P_{\mathrm{lost}}(\lambda) \log[P_{\mathrm{lost}}(\lambda)]$$

Best fi the ones which maximize I(X;Y). We test gaussian function typical of LED

Results

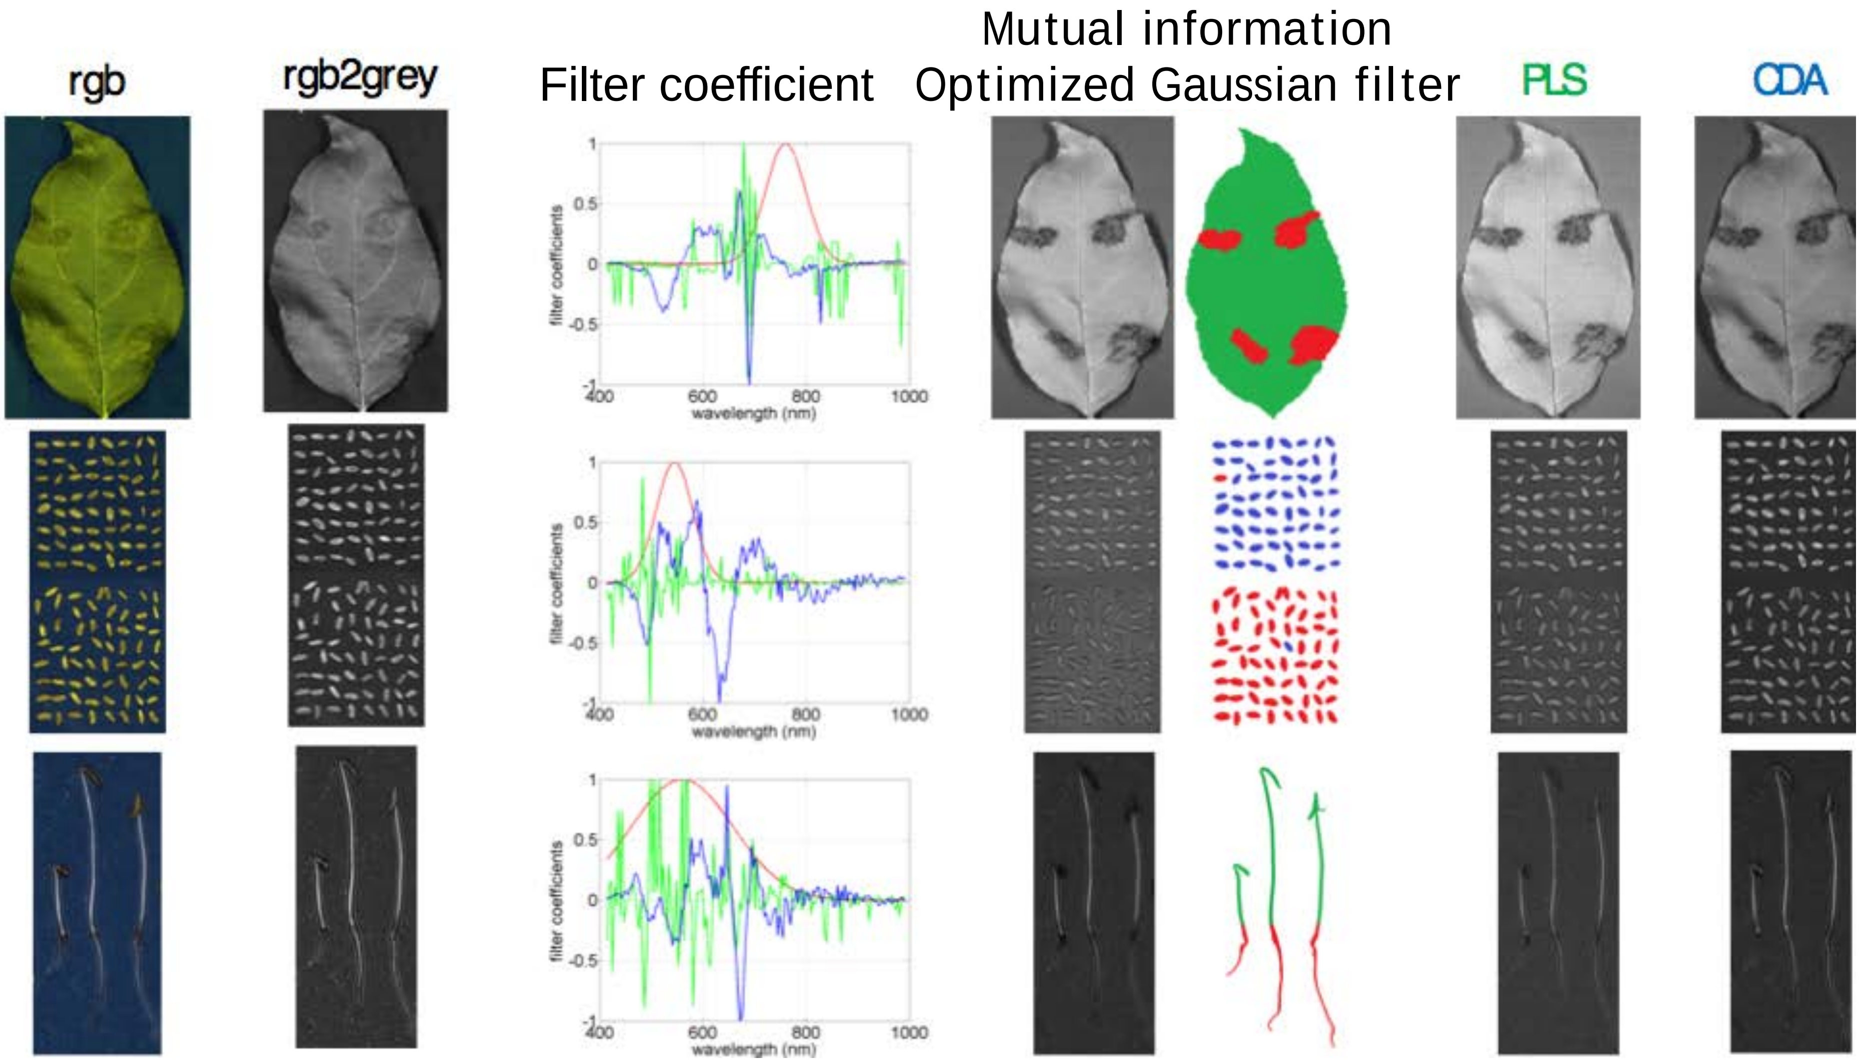

PLS : partial least square; CDA : Canonical Discriminant analysis
PLS, CDA : filter difficult to implement physically while Gaussian filter accessible with LED or Dichroïc standard filters

Low-cost spectro-imaging & compressed learning

David Rousseau

Cost effective spectral imaging

Cost

Type of sensors	Nb Channel	Filtrers	Imaging systems
Gray	1	Open choice (visible / NIR)	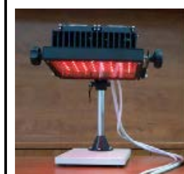  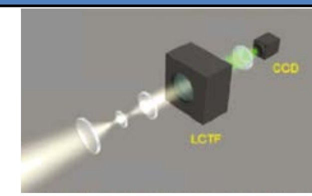
Color	3	Imposed (R V B) or on demand	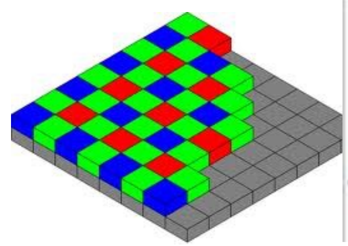 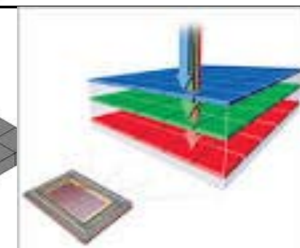 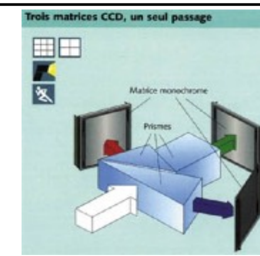
Multispectral	2 - 10	Open choice (visible / NIR)	  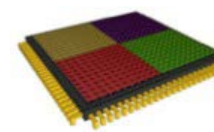
Hyperspectral	Tens to hundreds		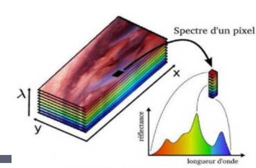

Cost effective spectral imaging

Cost

Type of sensors	Nb Channel	Filtrers	Imaging systems		
Gray	1	Open choice (visible / NIR)	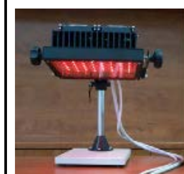		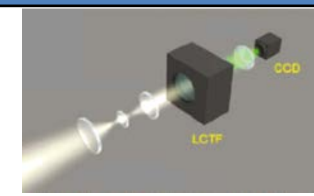
Color	3	Imposed (R V B) or on demand	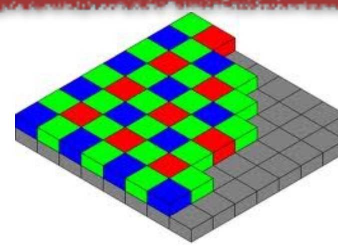	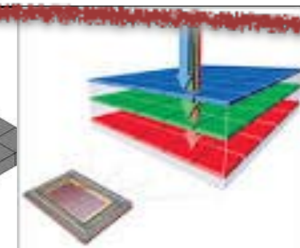	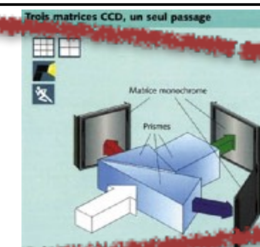
Multispectral	2 - 10	Open choice (visible / NIR)	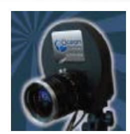		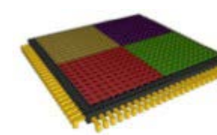
Hyperspectral	**Tens to hundreds**		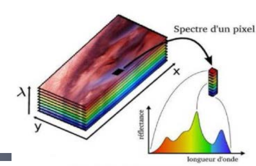		

Going (snapshot) hyperspectral

Computed Tomography Imaging Spectrometer [1] (based on a diffraction grating)

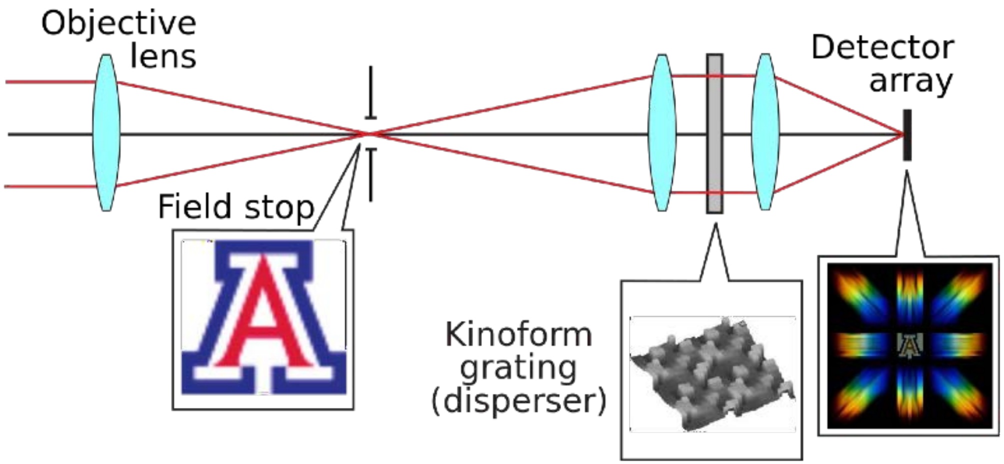

1. Descour et al . "Computed-tomography imaging spectrometer: experimental calibration and reconstruction results." *Applied Optics* (1995)

Going (snapshot) hyperspectral

Computed Tomography Imaging Spectrometer [1] (based on a diffraction grating)

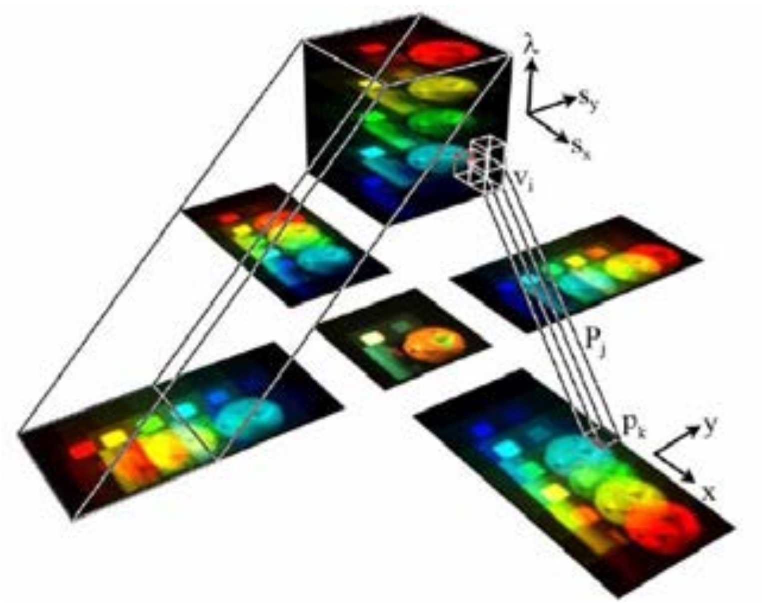

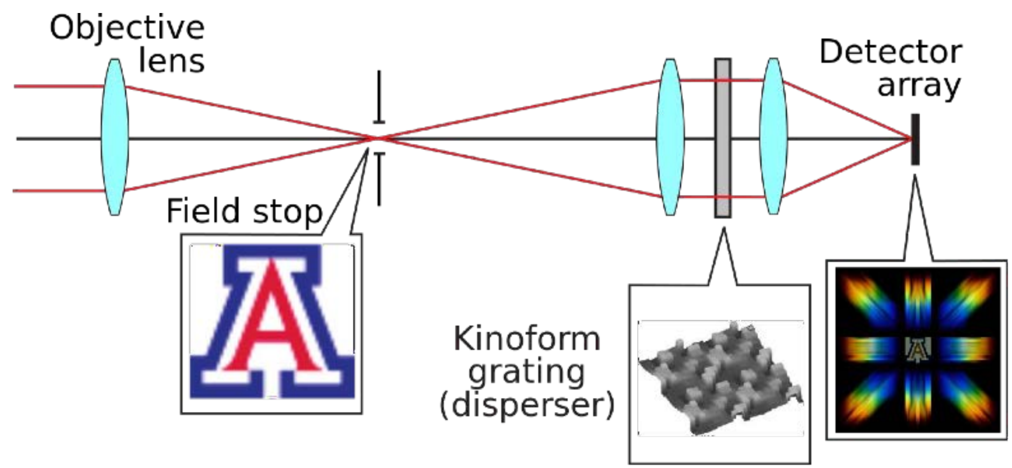

1. Descour et al . "Computed-tomography imaging spectrometer: experimental calibration and reconstruction results." *Applied Optics* (1995)

Going (snapshot) hyperspectral

Computed Tomography Imaging Spectrometer [1] (based on a diffraction grating)

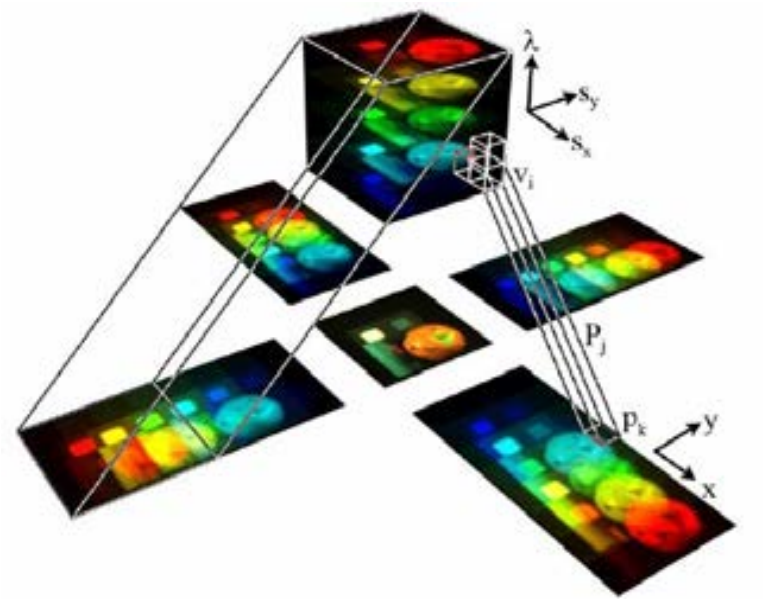

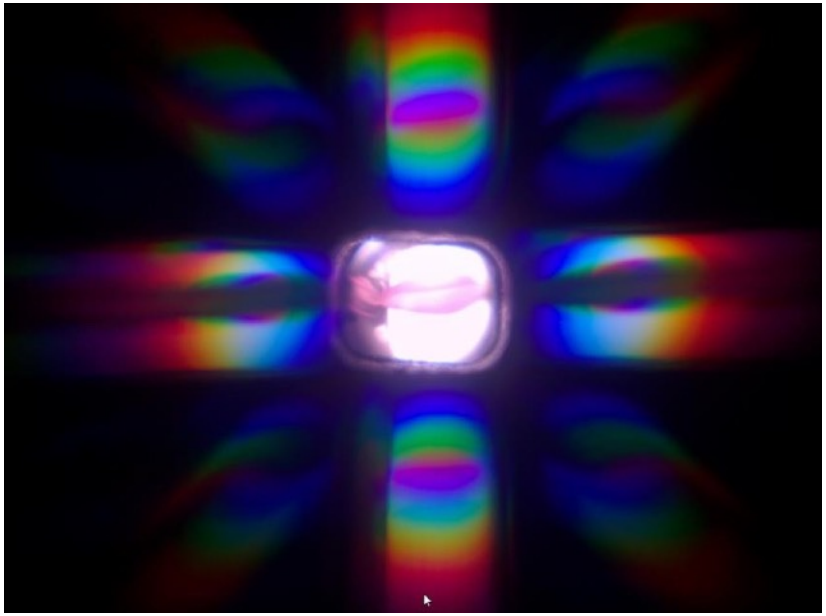

1. Descour et al . "Computed-tomography imaging spectrometer: experimental calibration and reconstruction results." *Applied Optics* (1995)

Compare full-resolution visible image and CTIS

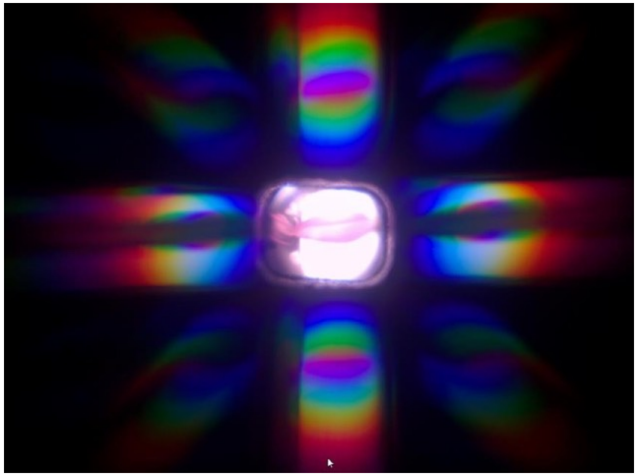

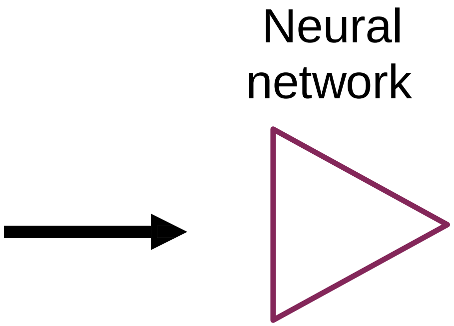

Goal : compare full-resolution visible image and CTIS

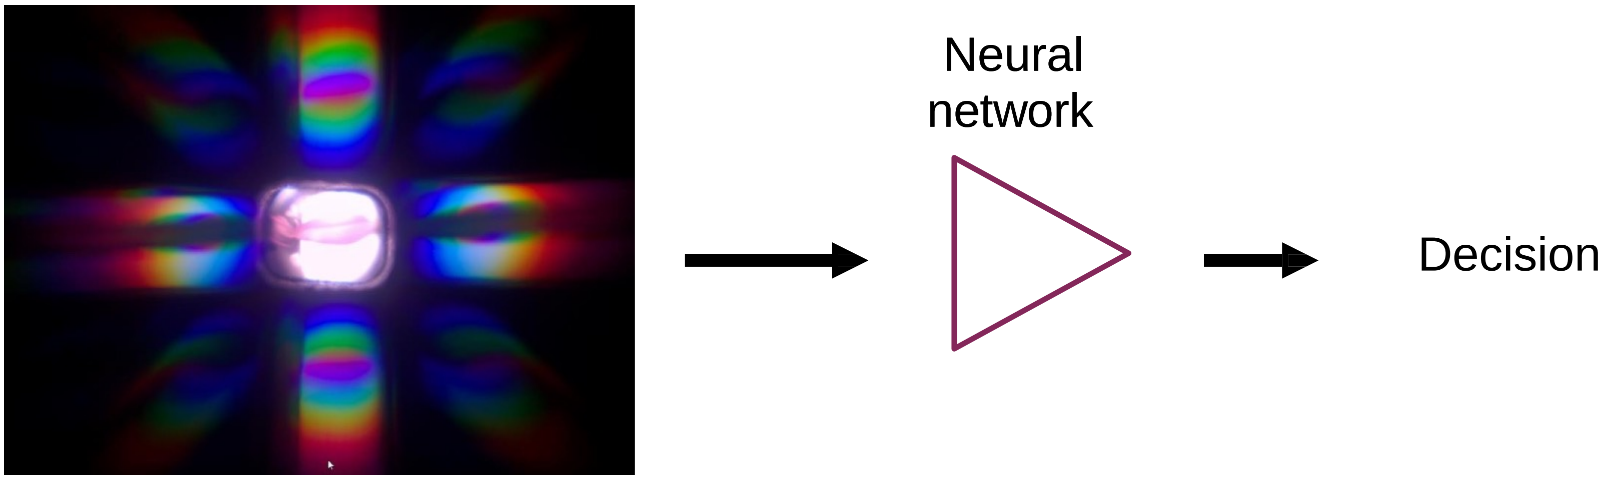

Goal : compare full-resolution visible image and CTIS

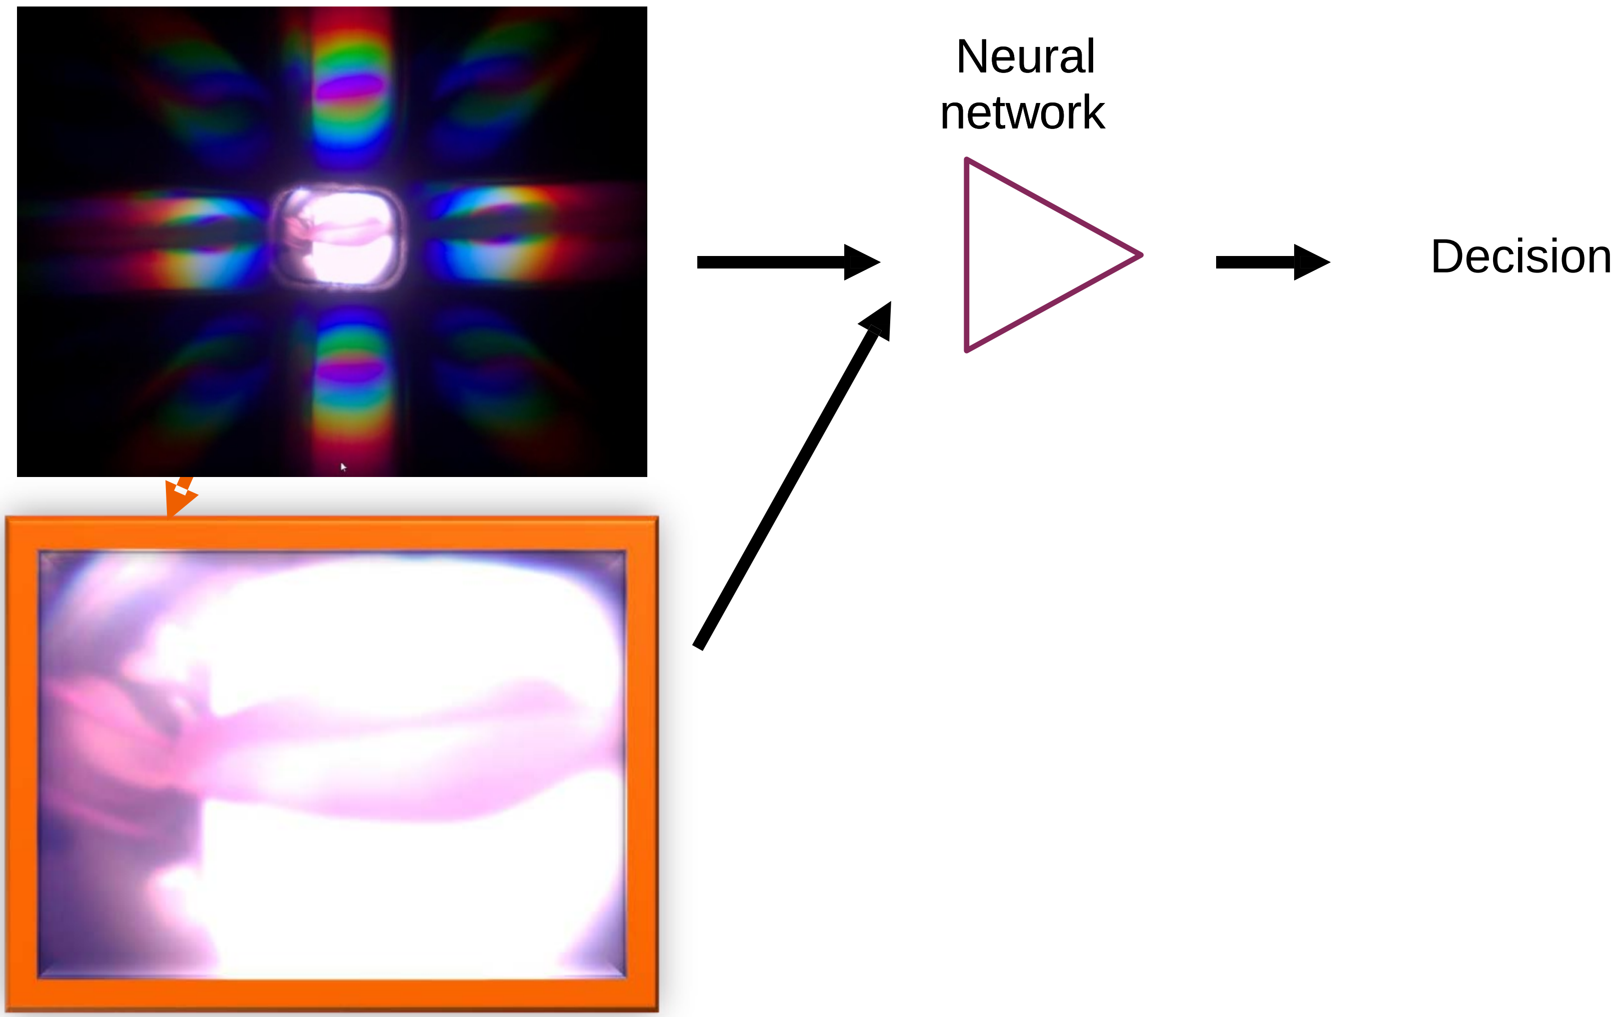

Goal : compare full-resolution visible image and CTIS

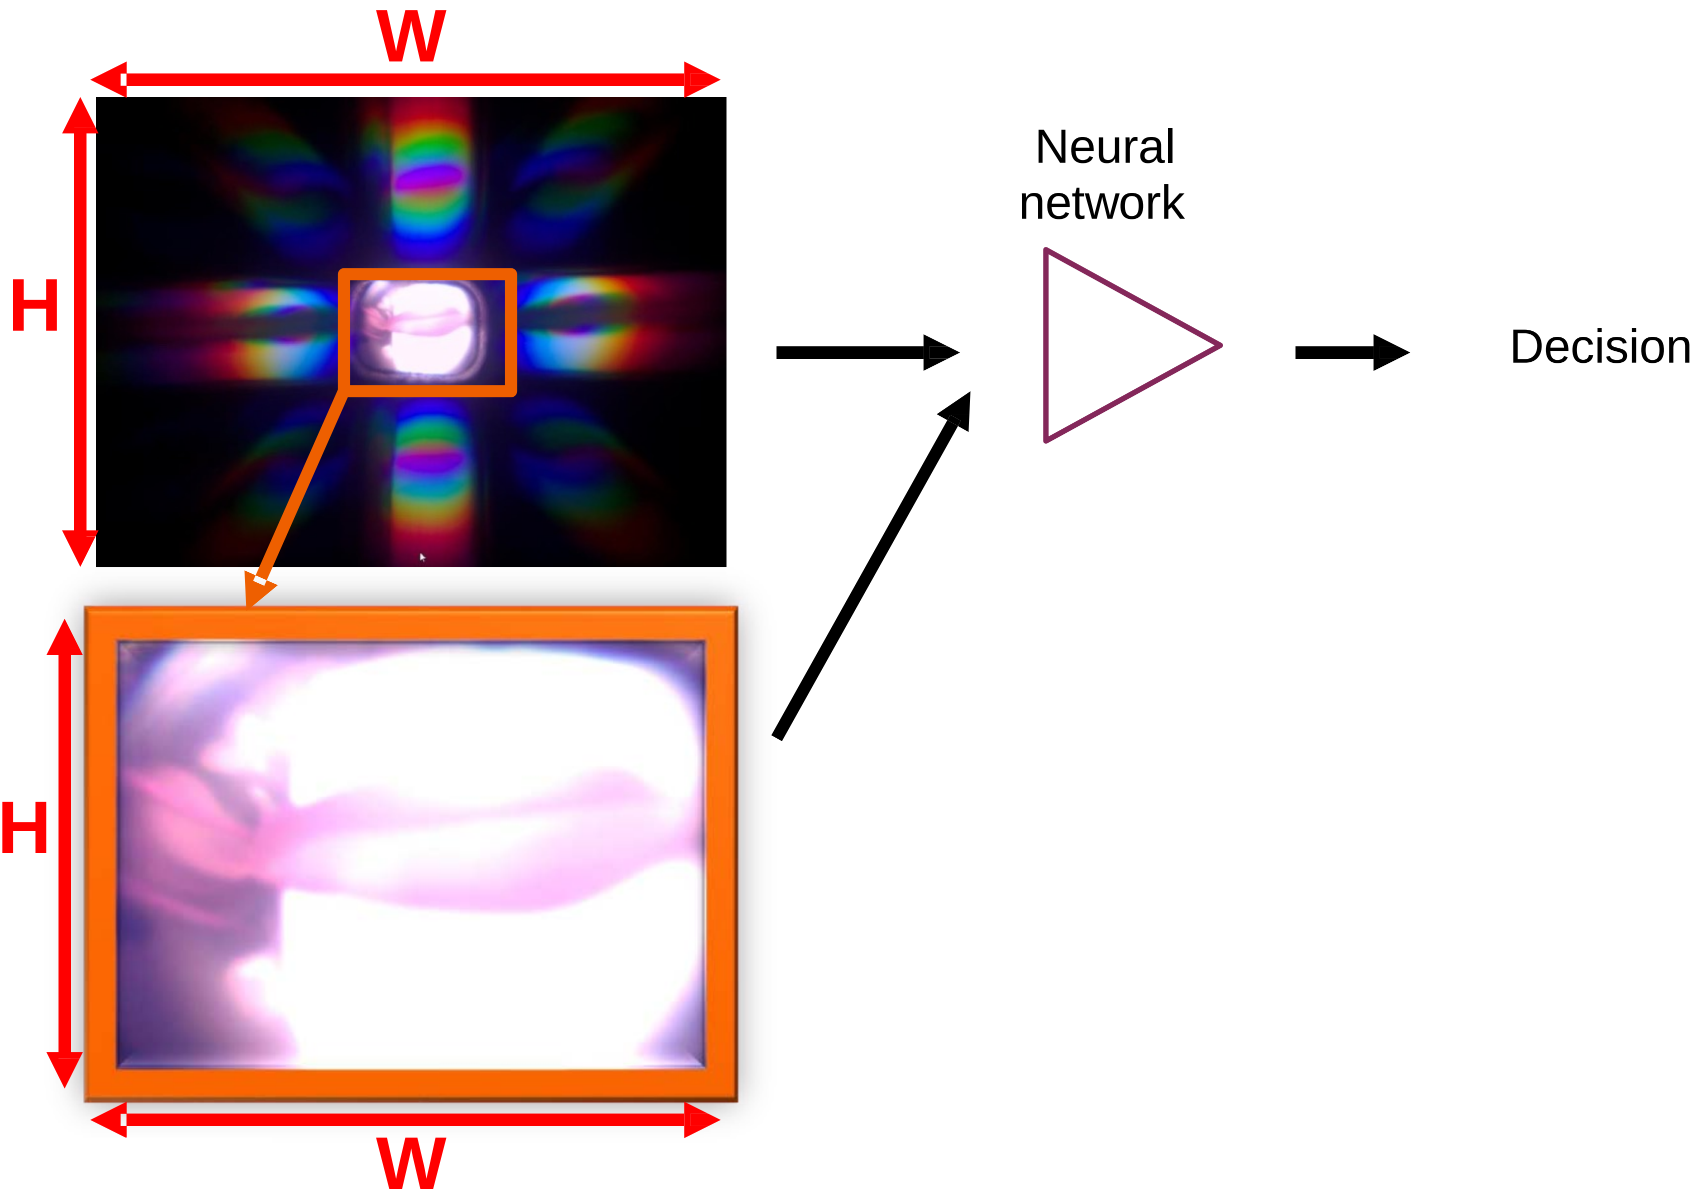

Goal : compare full-resolution visible image and CTIS

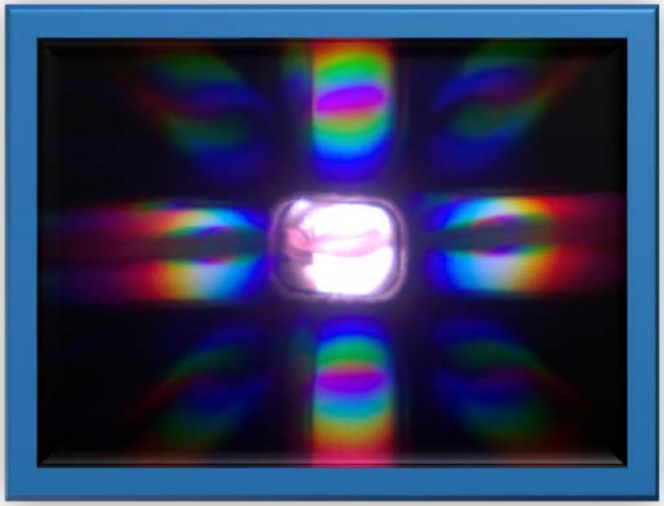

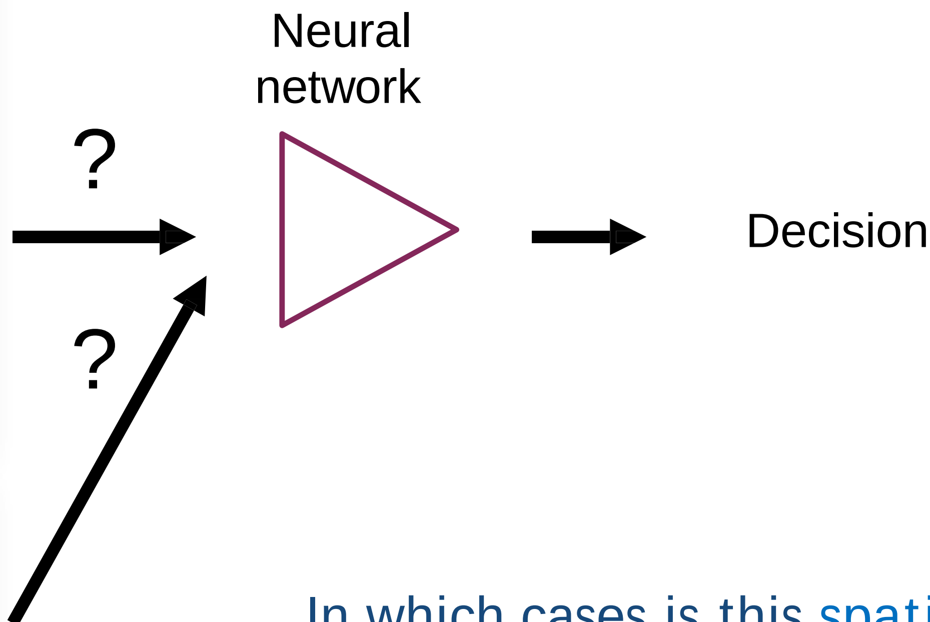

In which cases is this spatio-spectral imagery useful for a neural network compared to a full-resolution visible- spectrum image?

Case study : Apple scab

- Apple scab is a very serious disease
  afflicting apple trees. [2]

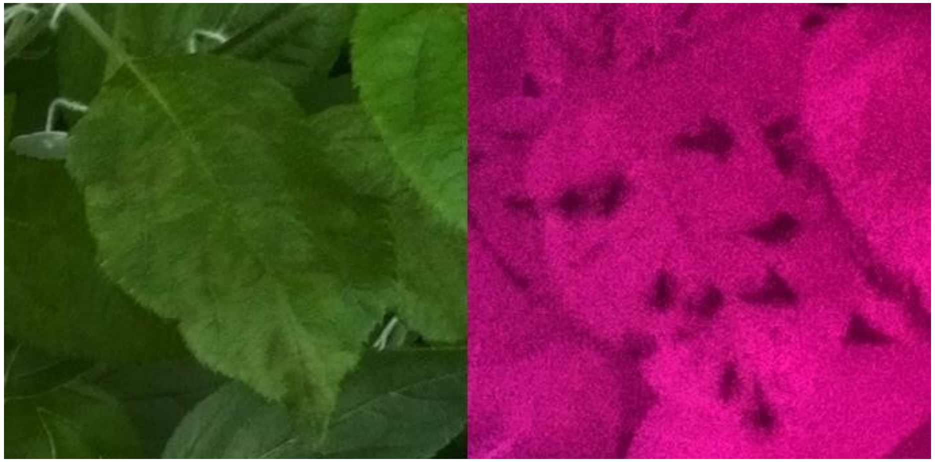

2. Bowen et al. "Venturia inaequalis: the causal agent of apple scab." *Molecular Plant Pathology* (2011)

Case study : Apple scab

- Apple scab is a very serious disease
  afflicting apple trees. [2]

- Visual symptoms : dark spots on the leaves.

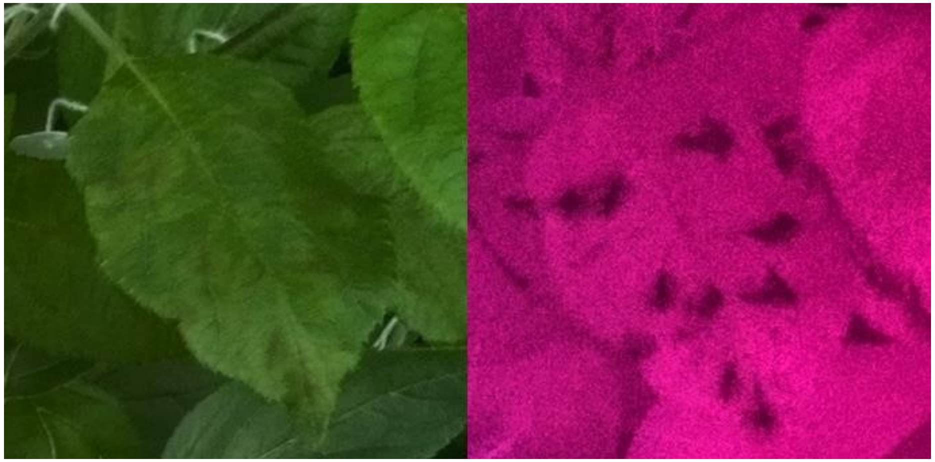

2. Bowen et al. "Venturia inaequalis: the causal agent of apple scab." *Molecular Plant Pathology* (2011)

Case study : Apple scab

- Apple scab is a very serious disease afflicting apple trees. [2]
- Visual symptoms : dark spots on the leaves.
- We developed a scab simulator to generate a “scabbed leaves” annotated dataset.

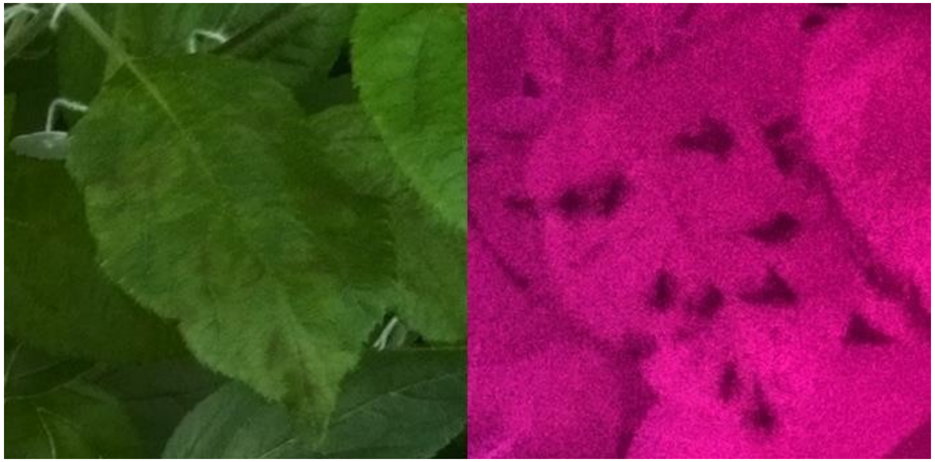

2. Bowen et al. "Venturia inaequalis: the causal agent of apple scab." *Molecular Plant Pathology* (2011)

Simulating RGB images

Healthy leaf from LeafSnap dataset [3]

3. Kumar et al. "Leafsnap: A computer vision system for automatic plant species identification." *ECCV*, 2012.

Simulating RGB images

Healthy leaf
from LeafSnap
dataset [3]

Scab lesion positions on
leaf [4]

3. Kumar et al. "Leafsnap: A computer vision system for automatic plant species identification." *ECCV*, 2012.

4. Douarre et al. "Novel data augmentation strategies to boost supervised segmentation of plant disease images", *Computers and Electronics in Agriculture* [under review]

Simulating RGB images

Healthy leaf from LeafSnap dataset[3]

Scabbed leaf

Small dataset of real scabbed leaves

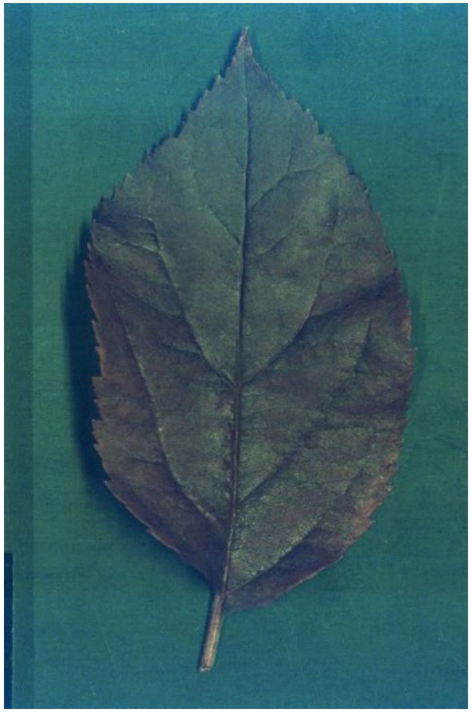

3. Kumar et al. "Leafsnap: A computer vision system for automatic plant species identification." *ECCV*, 2012.

4. Douarre et al. "Novel data augmentation strategies to boost supervised segmentation of plant disease images", *Computers and Electronics in Agriculture* [under review]

Simulating CTIS images : scab contrast

Hyperspectral acquisition of a leaf afflicted with scab

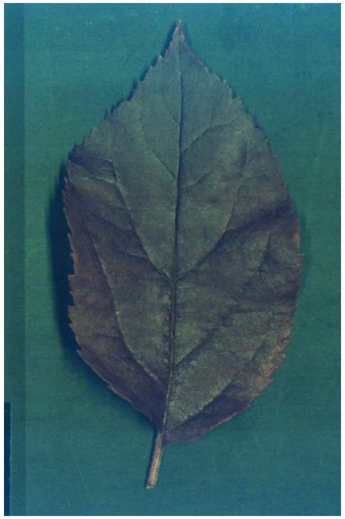

Simulating CTIS images : scab contrast

Hyperspectral acquisition of a leaf afflicted with scab

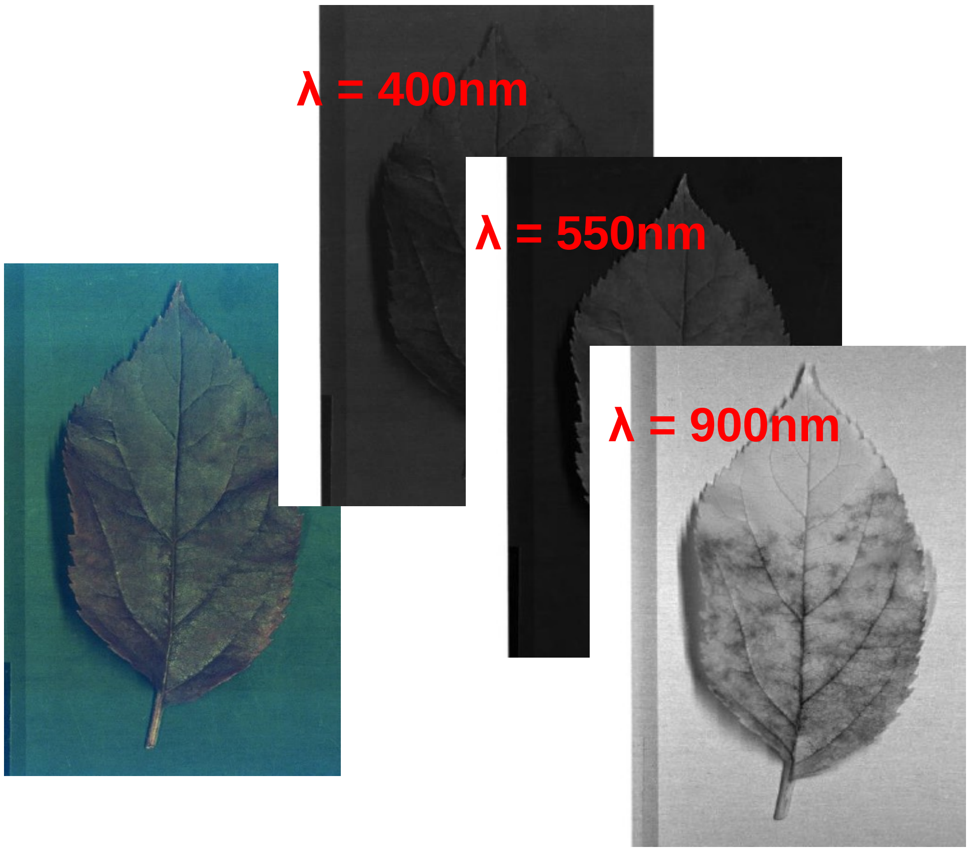

Acquisition in LARIS (Angers) with spectral-scanning camera (400-1000nm), 160 bands

Simulating CTIS images : scab contrast

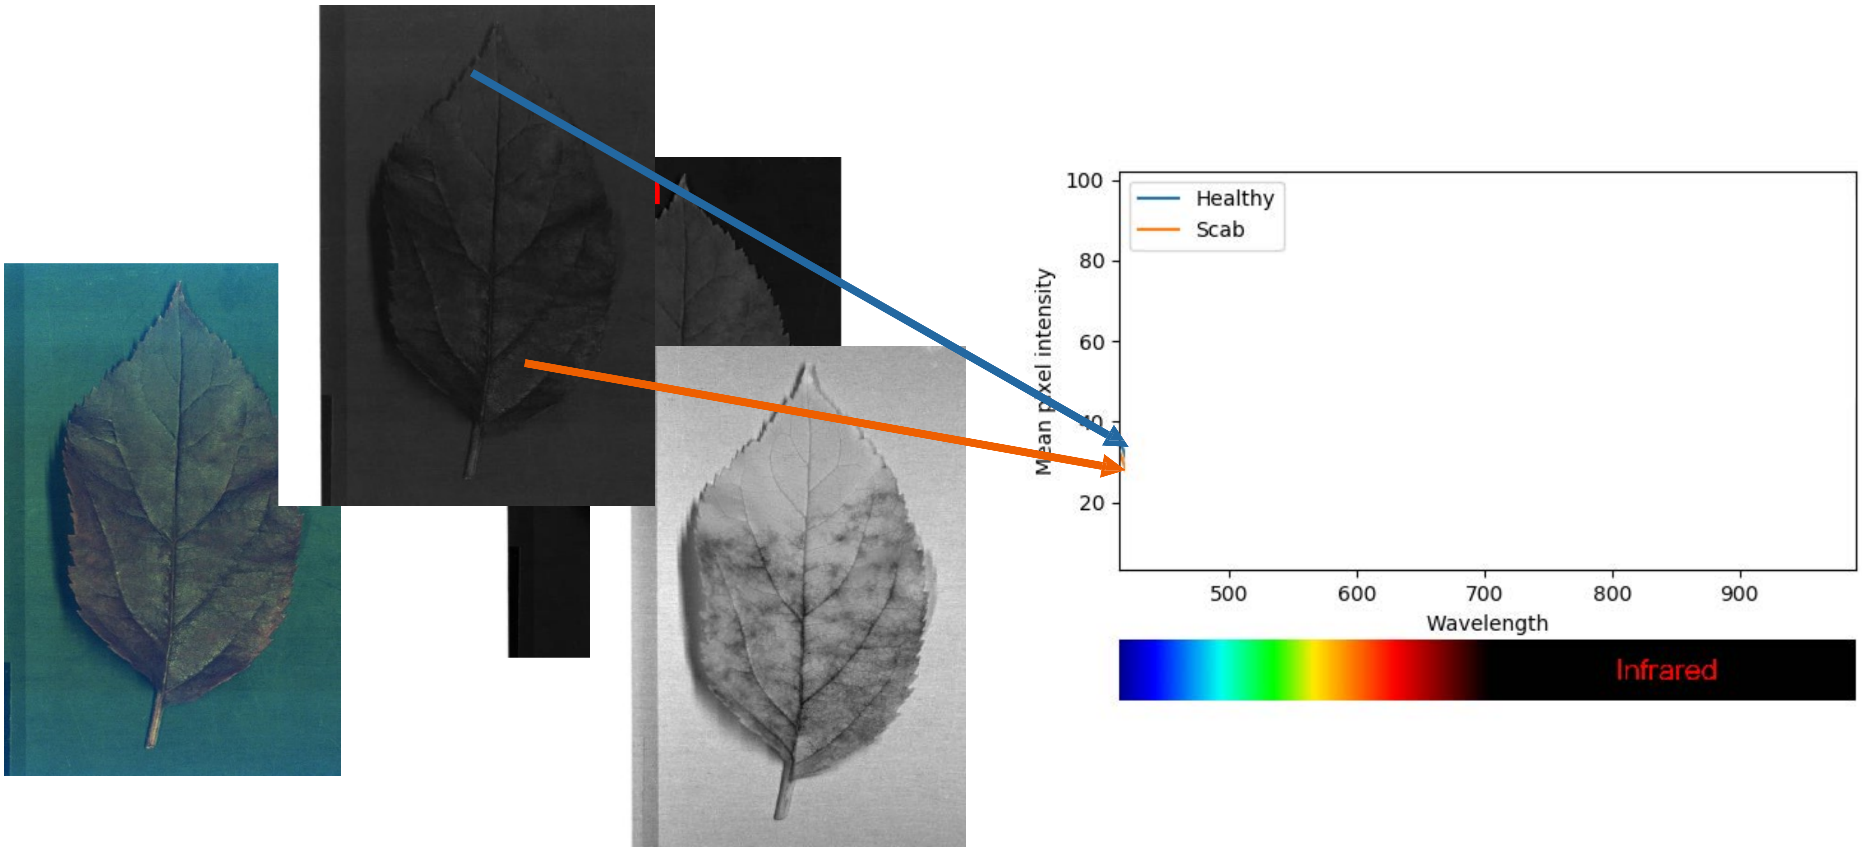

Acquisition in LARIS (Angers) with spectral-scanning camera (400-1000nm), 160 bands

Simulating CTIS images : scab contrast

Hyperspectral acquisition of a leaf afflicted with scab

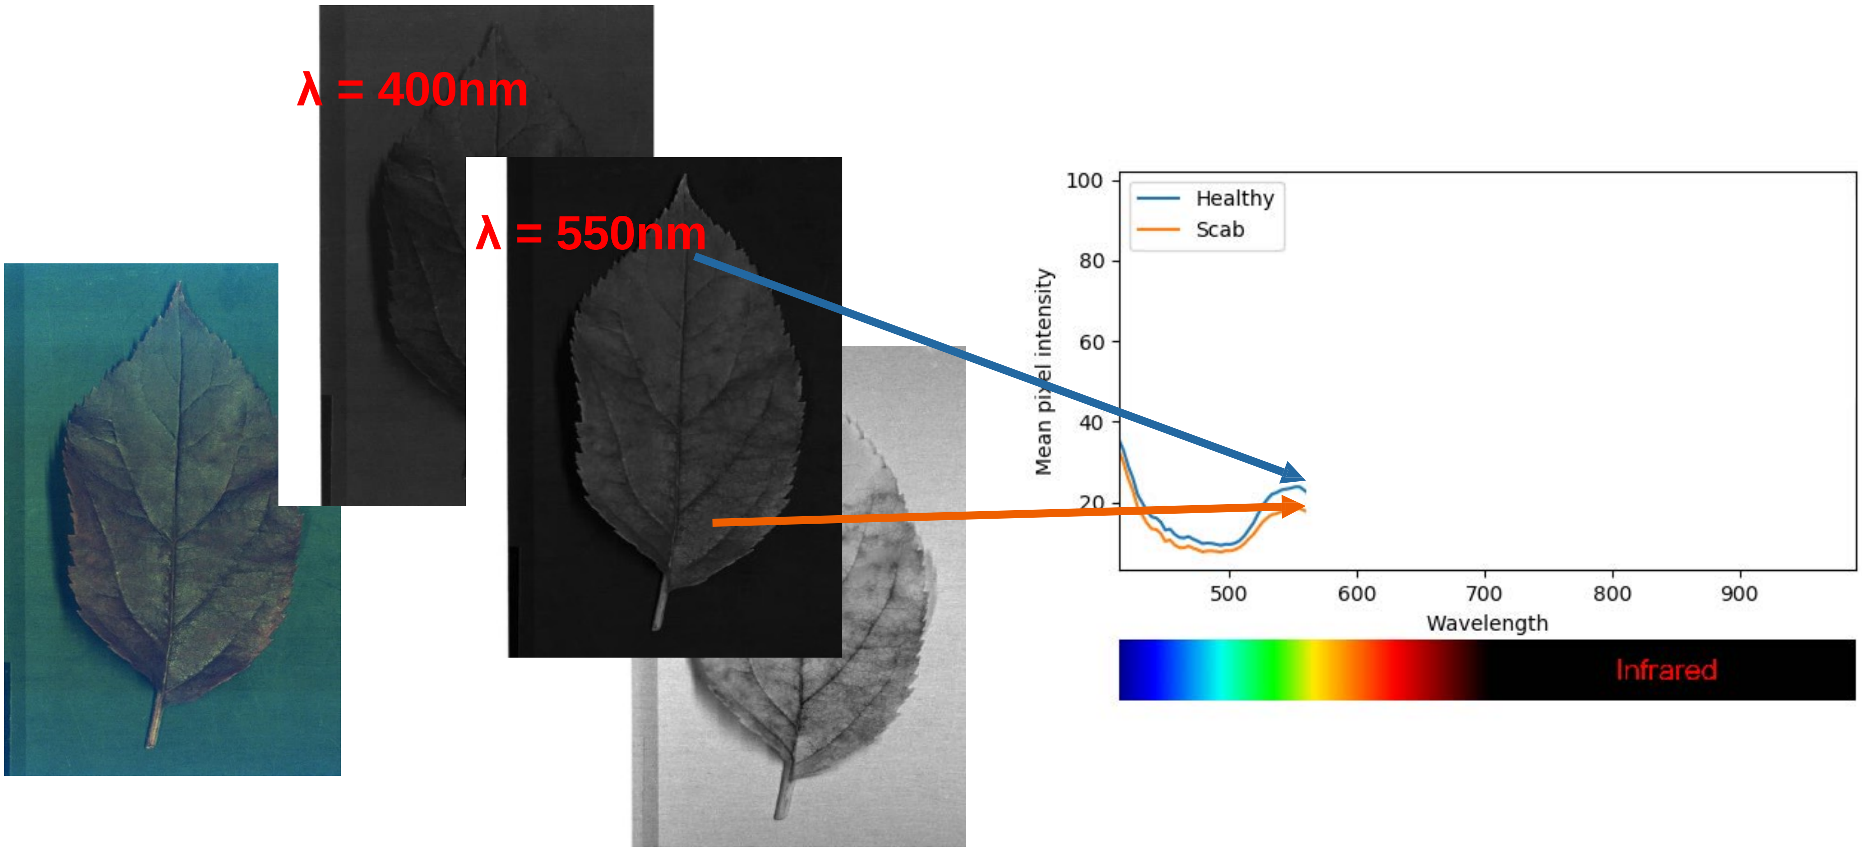

Acquisition in LARIS (Angers) with spectral-scanning camera (400-1000nm), 160 bands

Simulating CTIS images : scab contrast

Hyperspectral acquisition of a leaf afflicted with scab

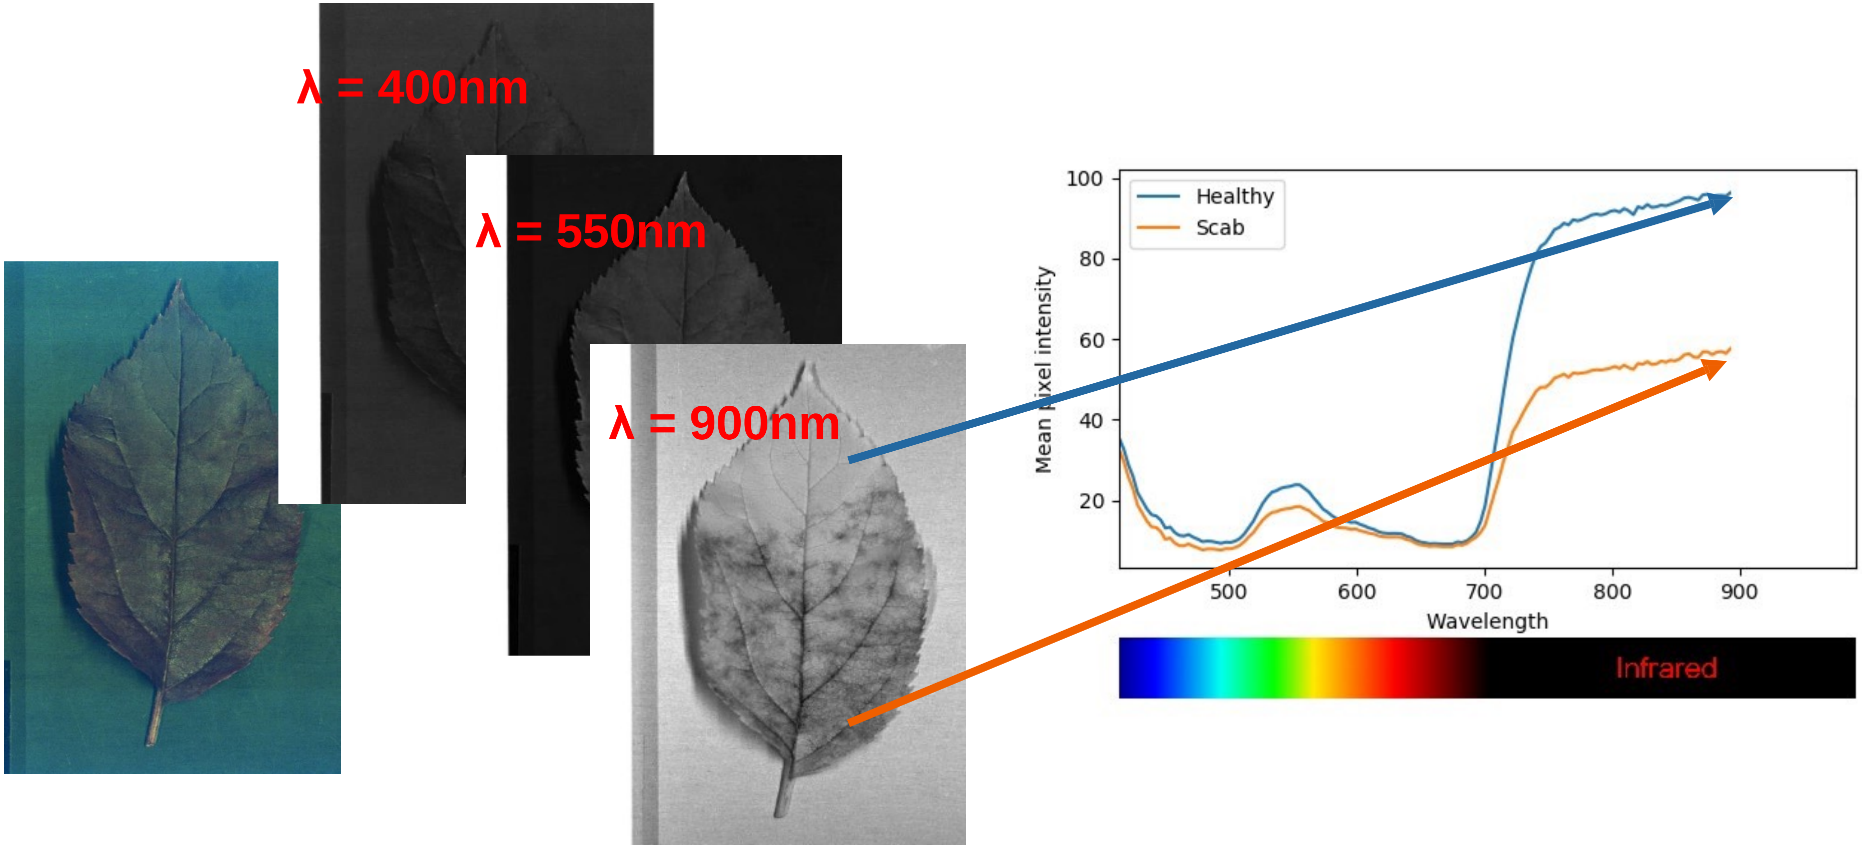

Acquisition in LARIS (Angers) with spectral-scanning camera (400-1000nm), 160 bands

Simulating CTIS images : scab contrast

Hyperspectral acquisition of a leaf afflicted with scab

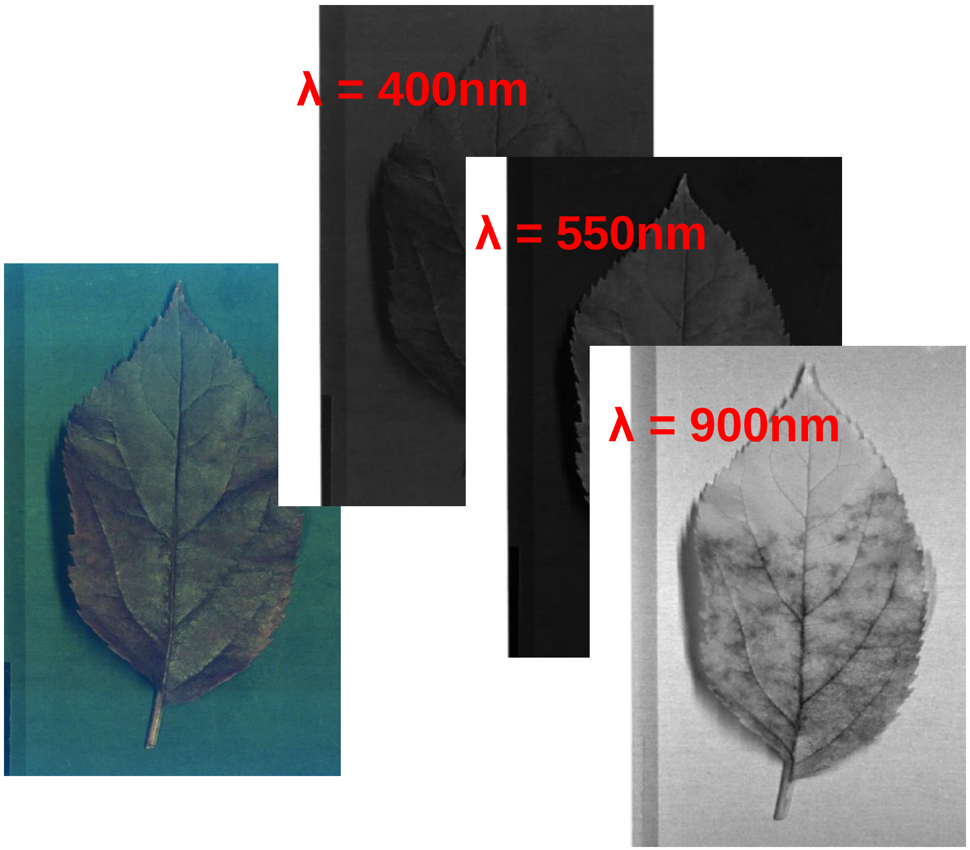

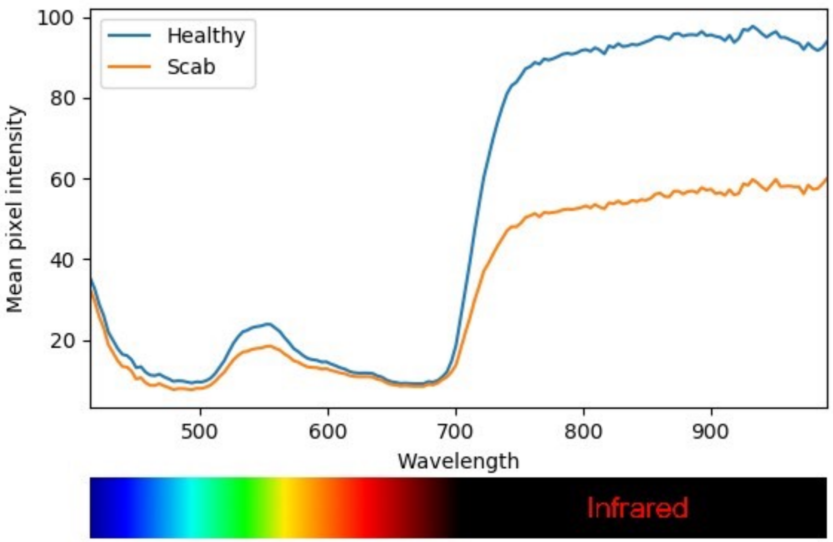

Acquisition in LARIS (Angers) with spectral-scanning camera (400-1000nm), 160 bands

Simulating CTIS images : imaging system

CTIS model

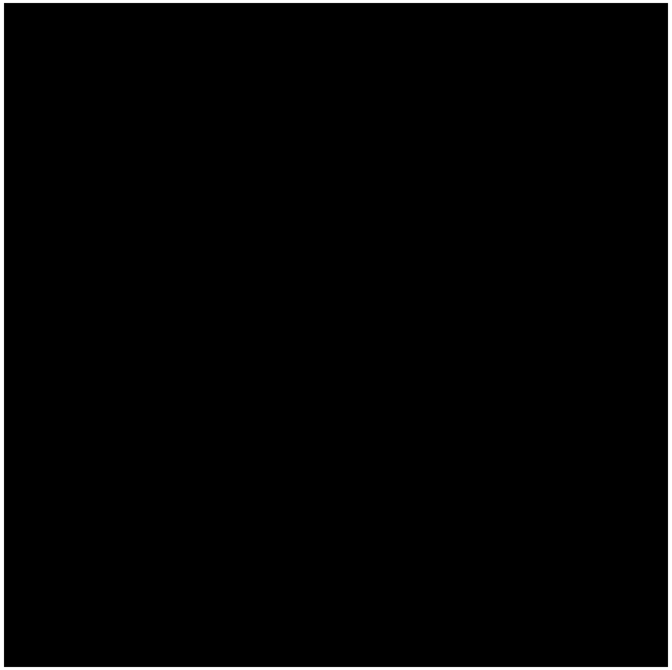

CTIS output

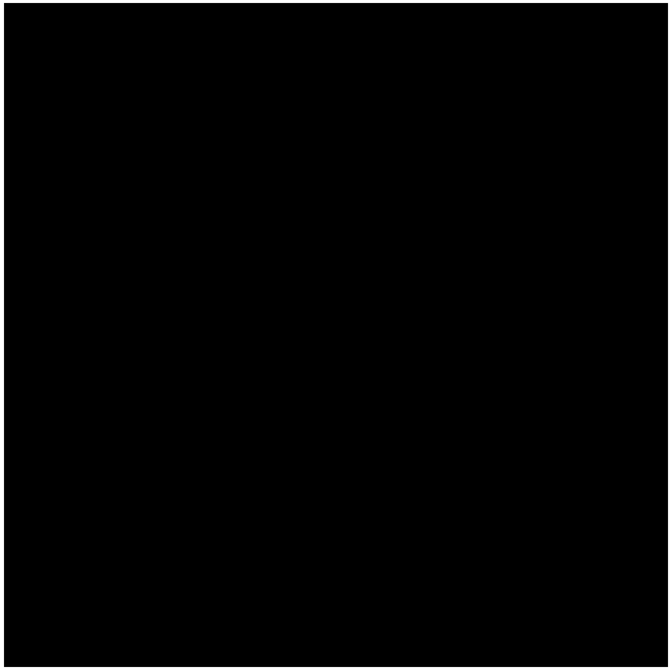

Simulating CTIS images : imaging system

CTIS model

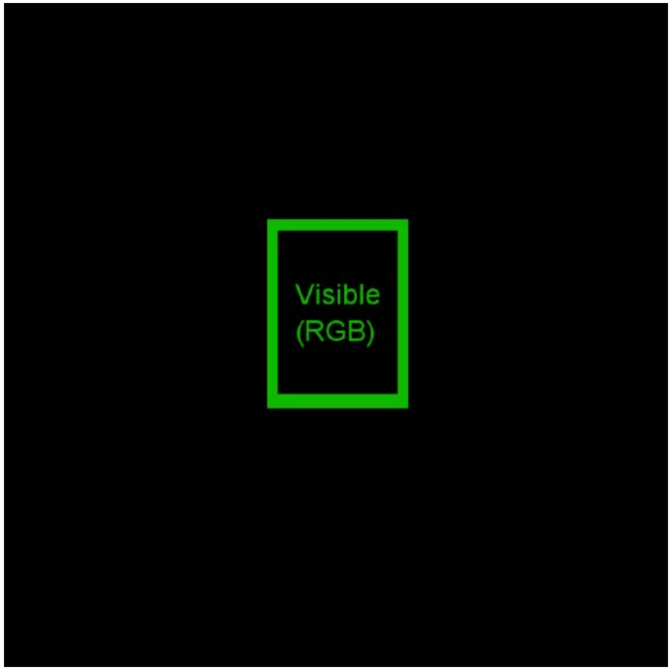

CTIS output

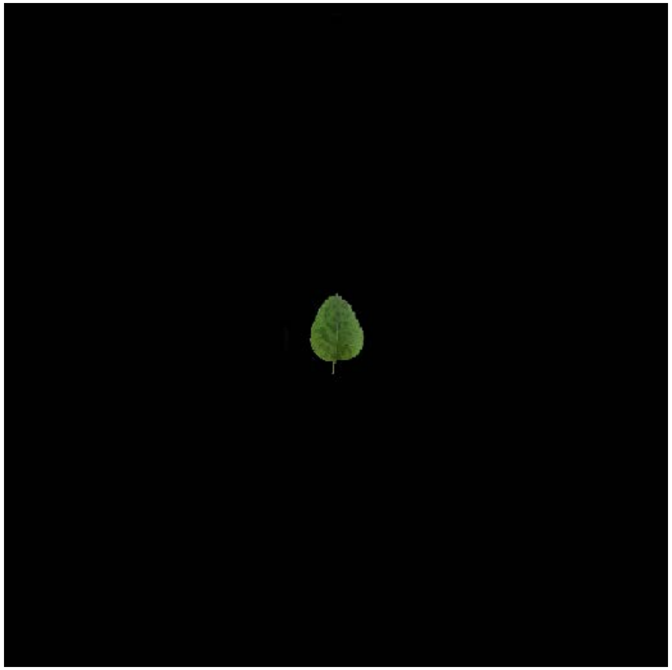

Simulating CTIS images : imaging system

CTIS model

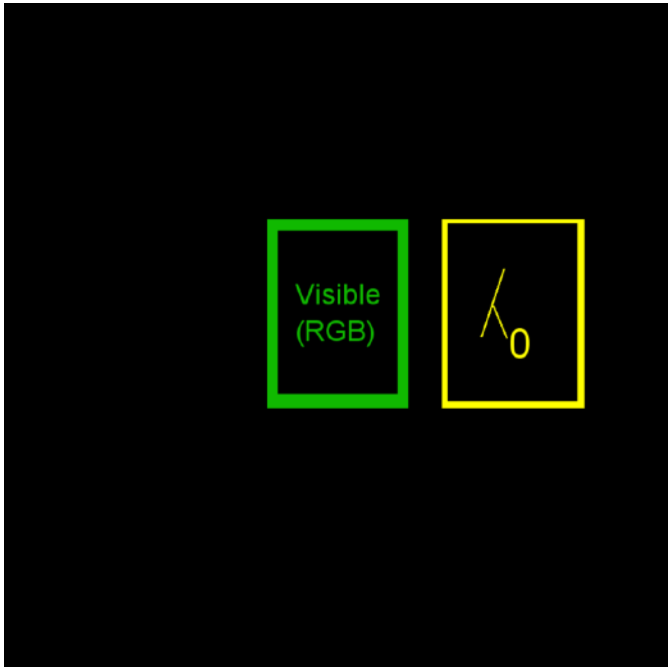

CTIS output

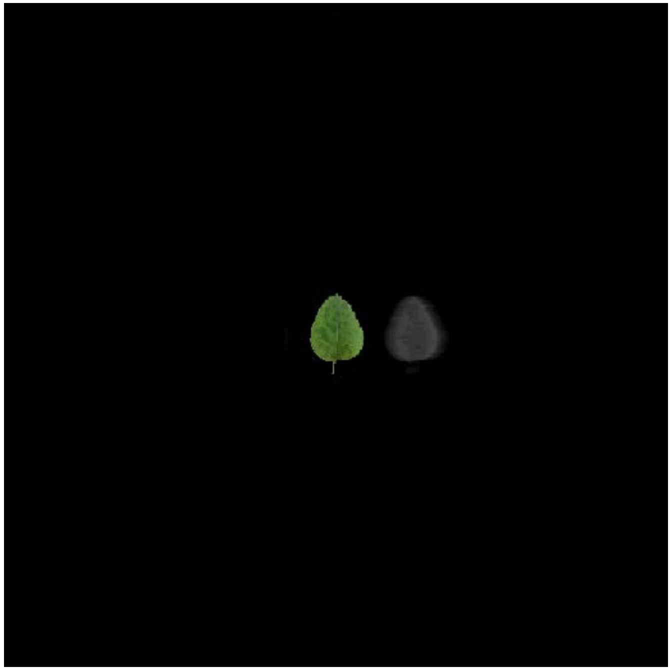

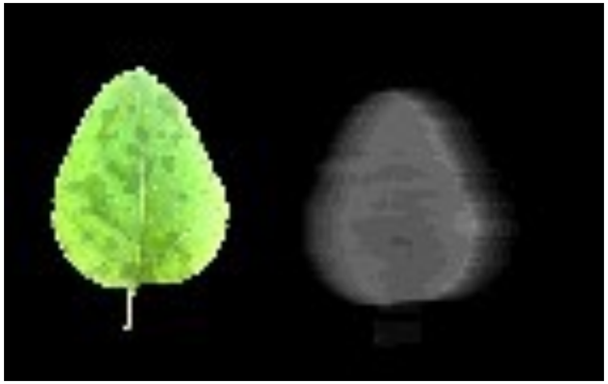

Simulating CTIS images : imaging system

CTIS model

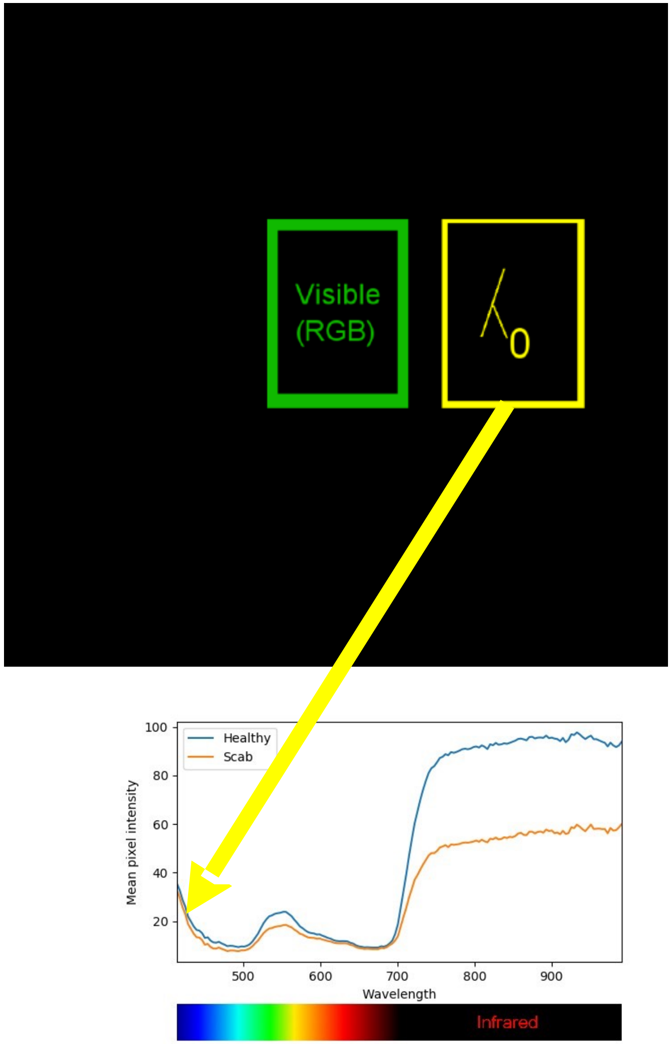

CTIS output

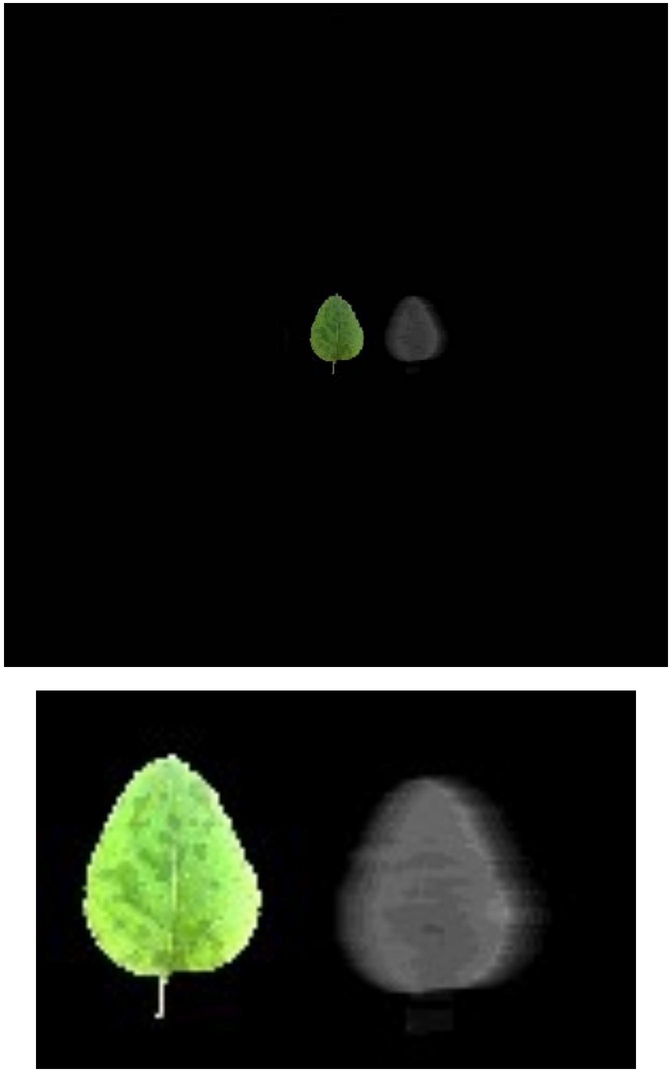

Simulating CTIS images : imaging system

CTIS model

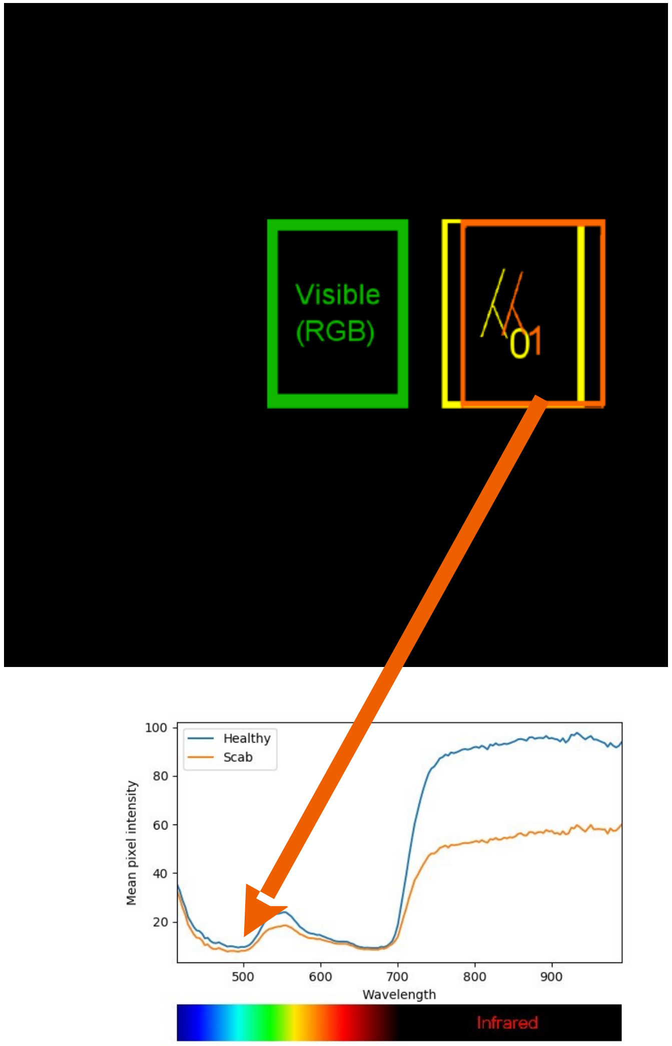

CTIS output

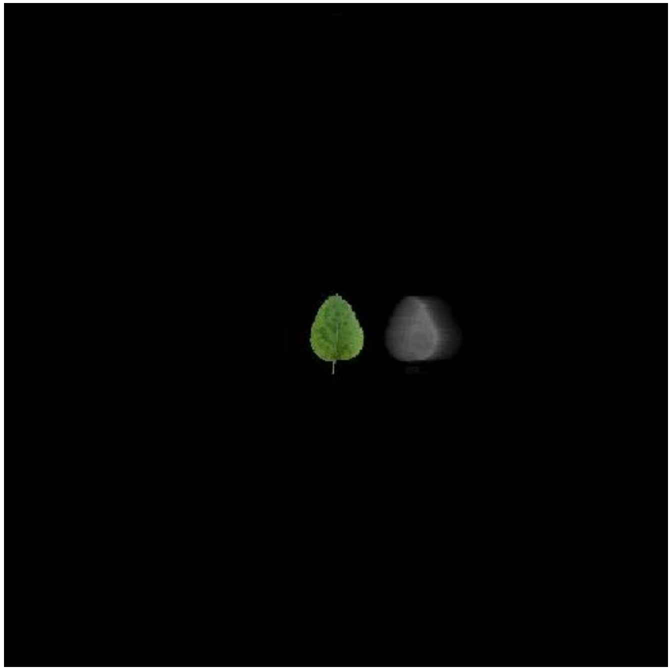

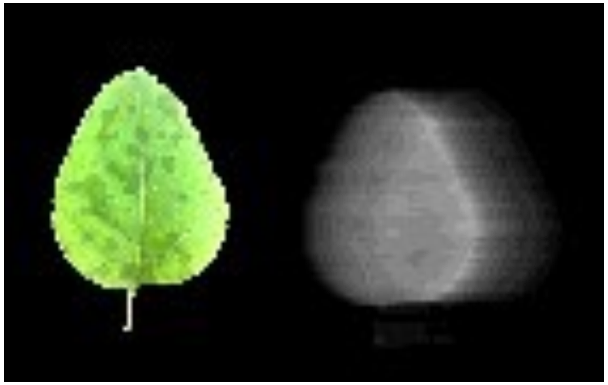

Simulating CTIS images : imaging system

CTIS model

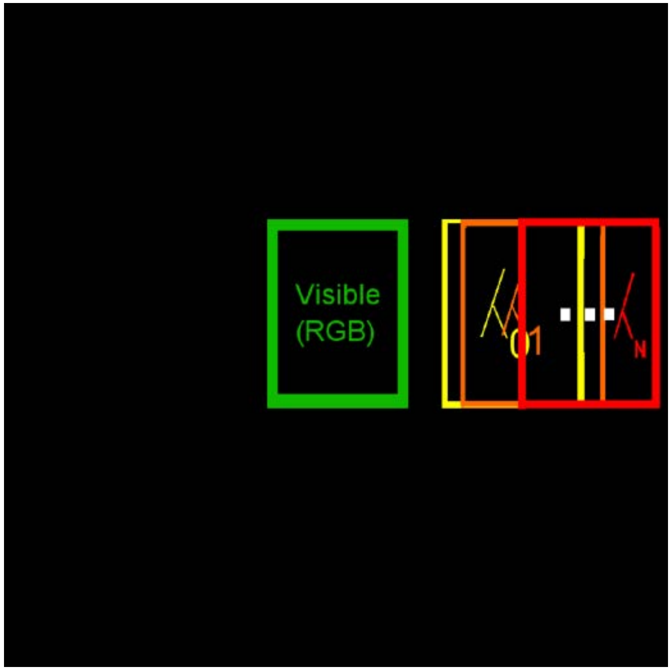

CTIS output

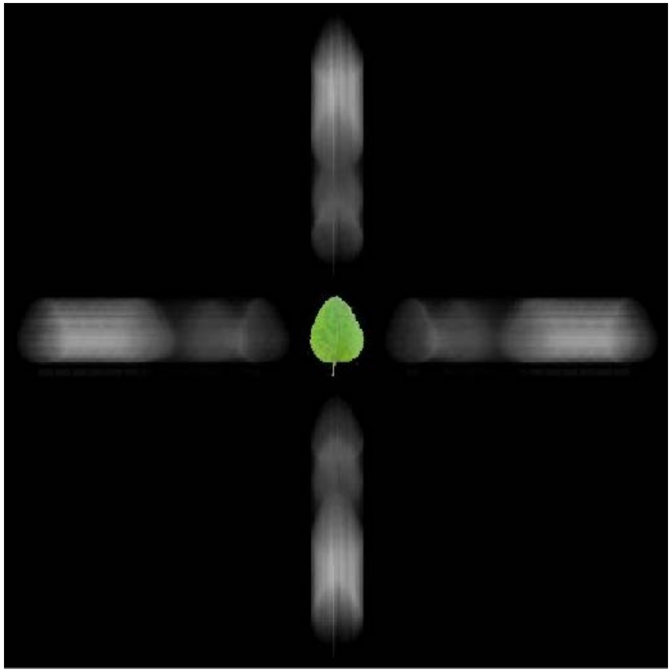

Datasets

3000 simulated full-size RGB images
separated in train/validation/test sets.

Datasets

3000 simulated full-size RGB images separated in train/validation/test sets.

3000 simulated CTIS images separated in the same way.

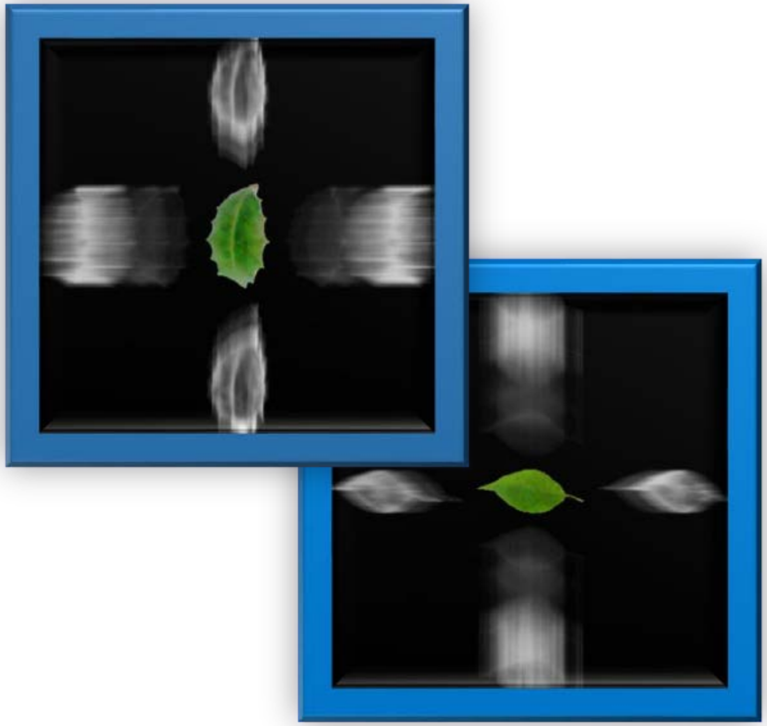

Scab contrast variation

- Scab infection progress → stronger contrast between scab and non scab.

Scab contrast variation

- Scab infection progress → stronger contrast between scab and non scab.
- Generation of datasets with varying contrast, to simulate various infection degrees.

Scab contrast variation

- Scab infection progress → stronger contrast between scab and non scab.
- Generation of datasets with varying contrast, to simulate various infection degrees.

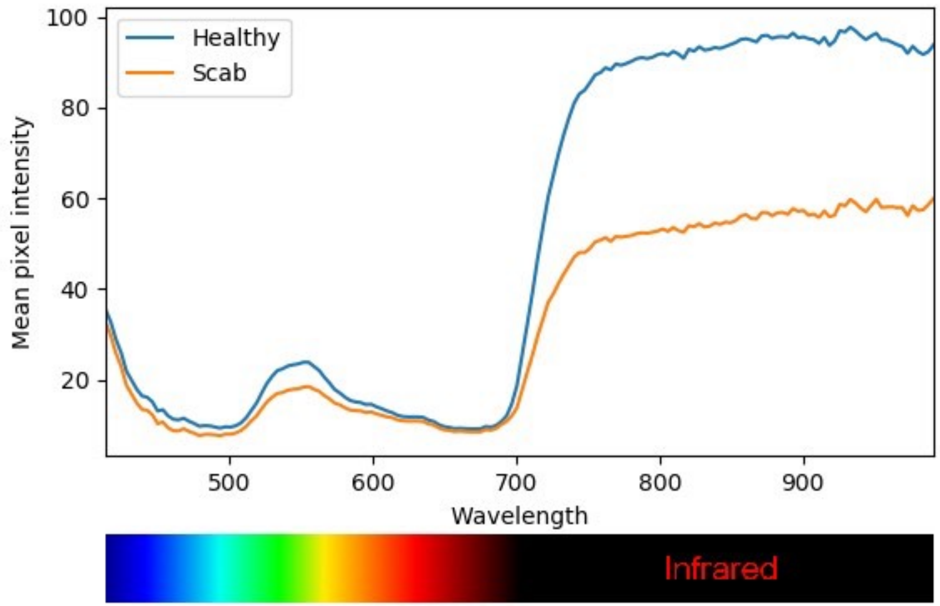

Scab contrast variation

- Scab infection progress → stronger contrast between scab and non scab.
- Generation of datasets with varying contrast, to simulate various infection degrees.

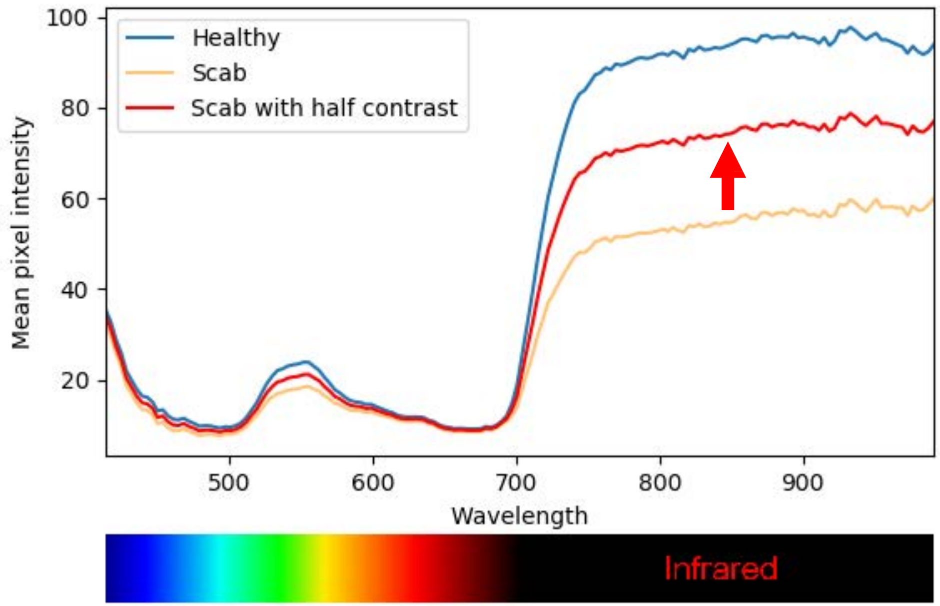

Scab contrast variation

- Scab infection progress → stronger contrast between scab and non scab.
- Generation of datasets with varying contrast, to simulate various infection degrees.

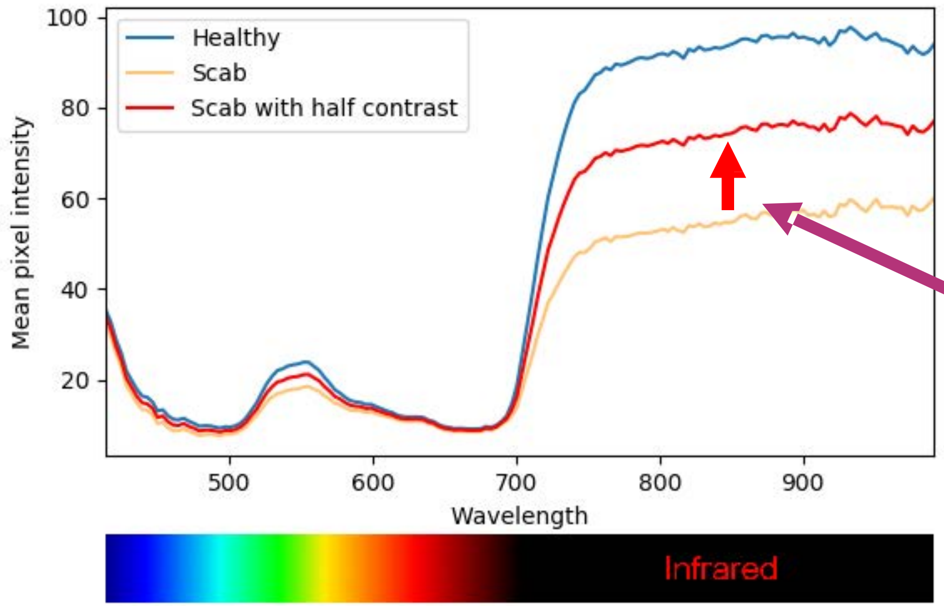

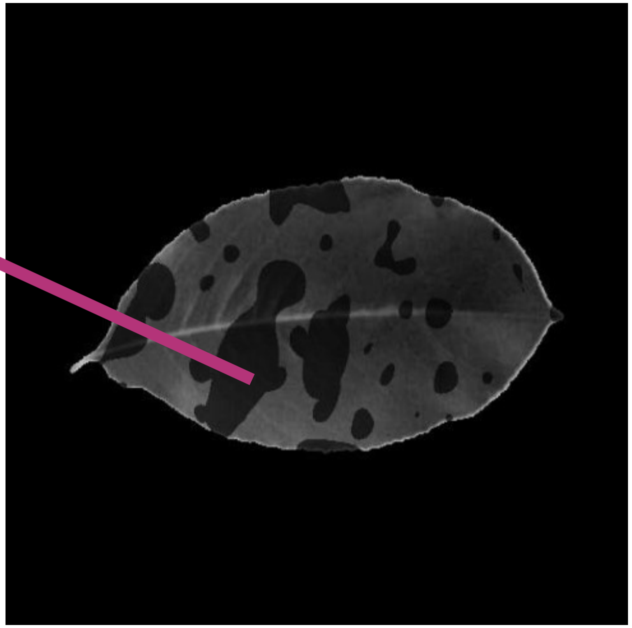

Scab contrast variation

- Scab infection progress → stronger contrast betw
- Generation of datasets with varying contrast, to
  degrees.

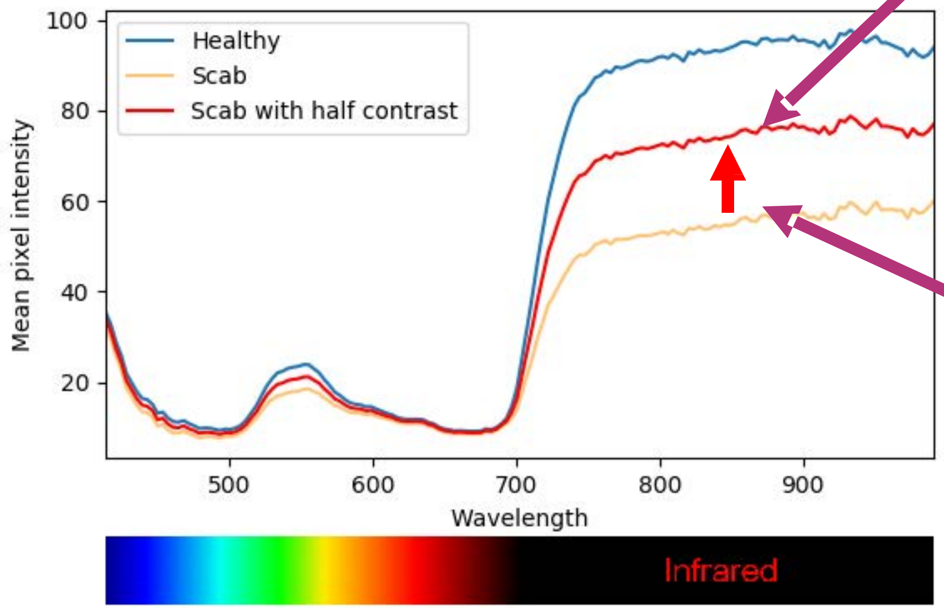

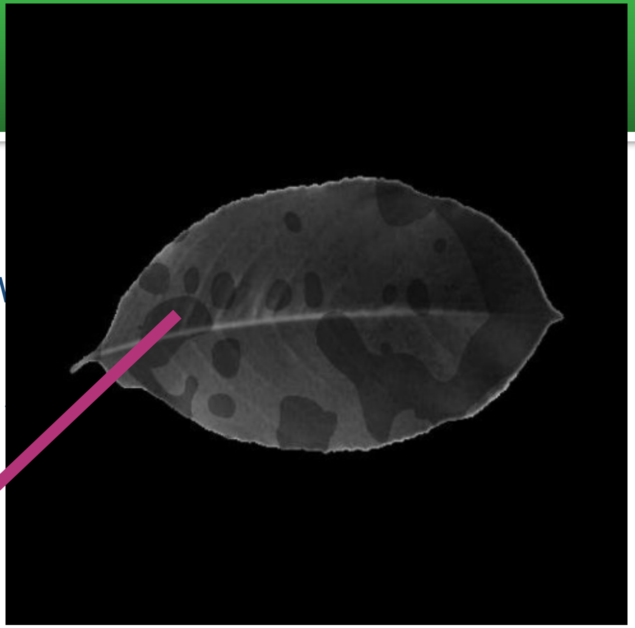

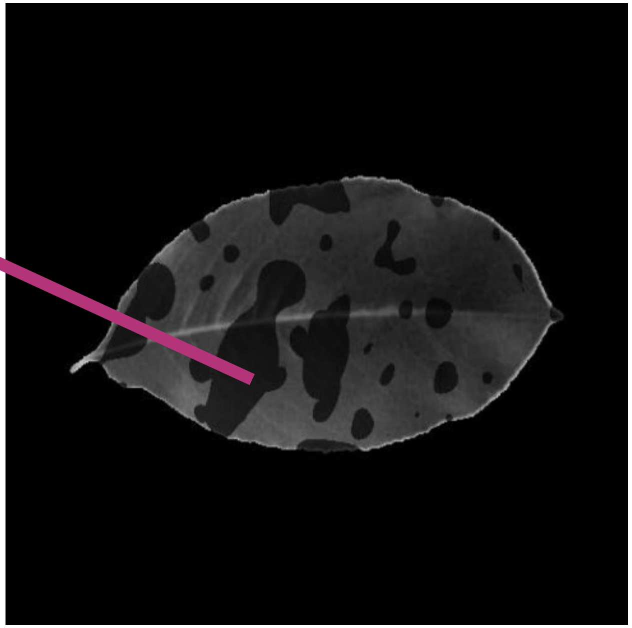

Training specifics

- Classification problem between scab and healthy (50/50).

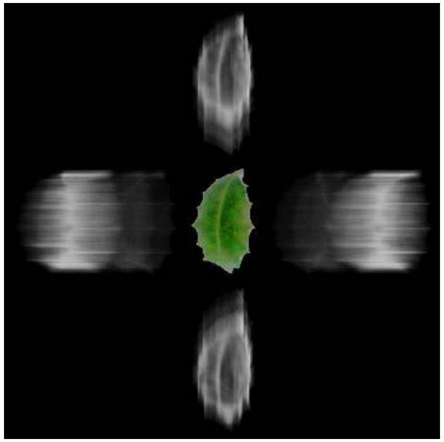

Training specifics

- Classification problem between scab and healthy (50/50).

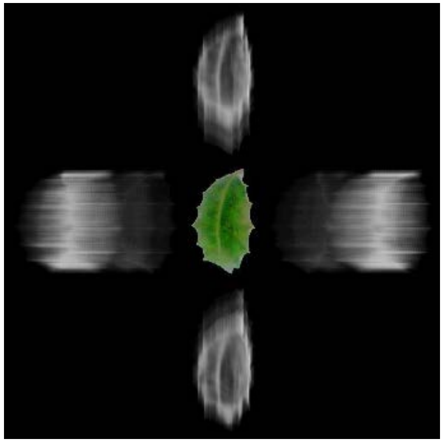

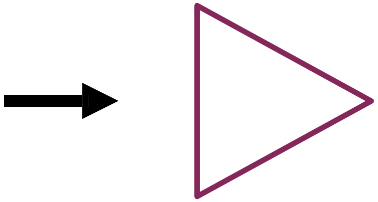

Training specifics

- Classification problem between scab and healthy (50/50).

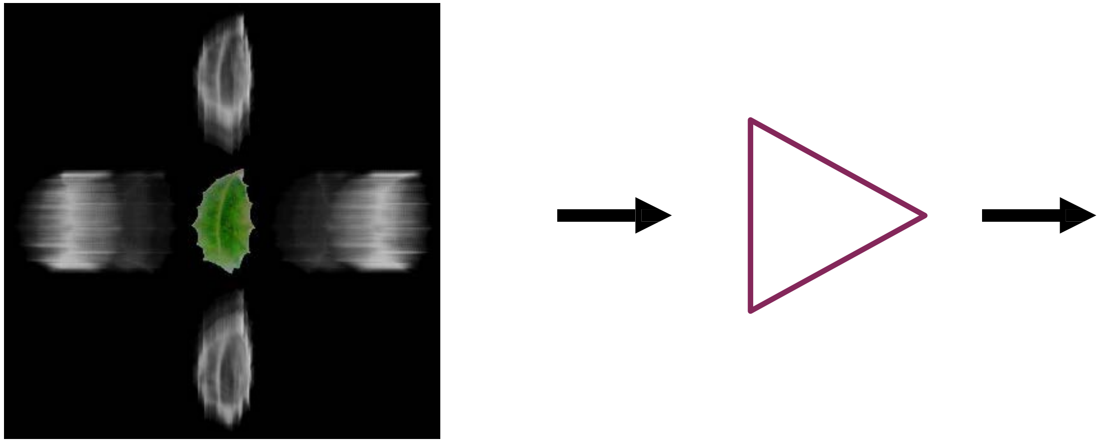

“This is a healthy leaf /
This is a scabbed leaf.”

Training specifics

- Classification problem between scab and healthy (50/50).
- Metric is Matthews Correlation Coefficient (MCC).

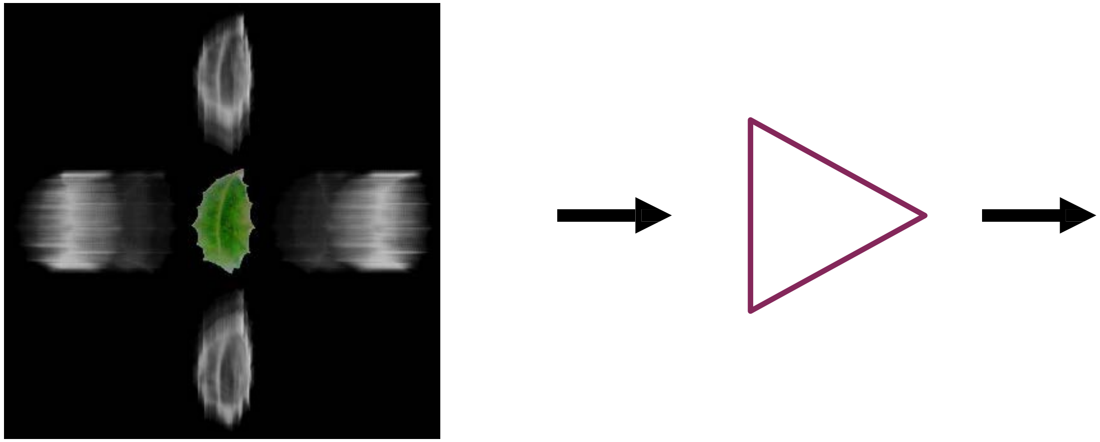

“This is a healthy leaf /
This is a scabbed leaf.”

Training specifics

- Classification problem between scab and healthy (50/50).
- Metric is Matthews Correlation Coefficient (MCC).
- Trained on a reduced VGG network [5], pre-trained on ImageNet, with standard data augmentation.

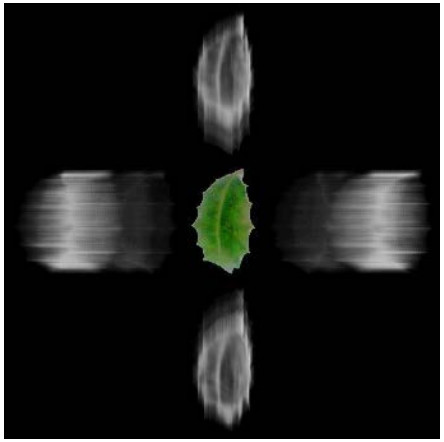

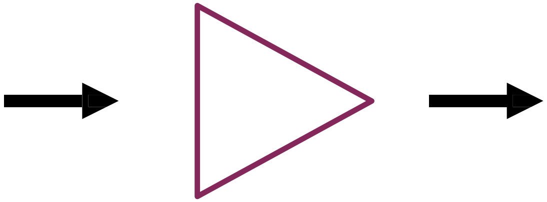

“This is a healthy leaf /
This is a scabbed leaf.”

5. Simonyan et al. "Very deep convolutional networks for large-scale image recognition." *arXiv preprint* (2014)

CTIS vs Full resolution spatial

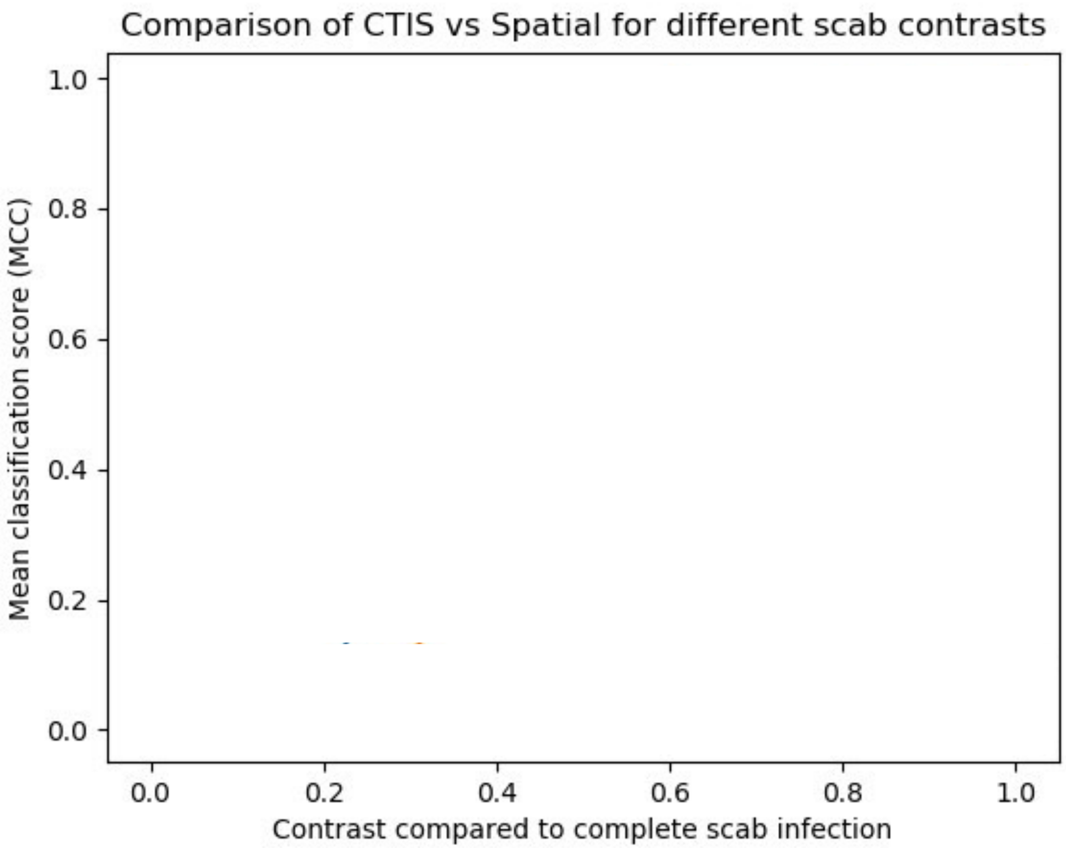

CTIS vs Full resolution spatial

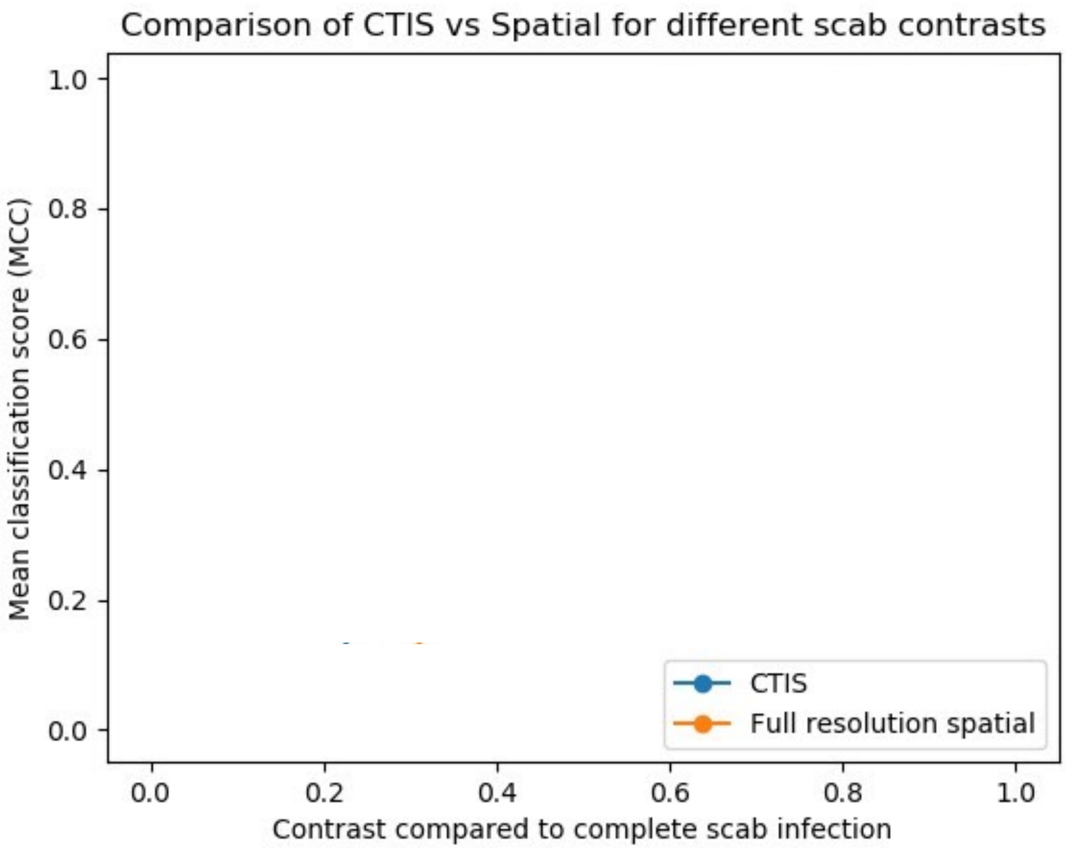

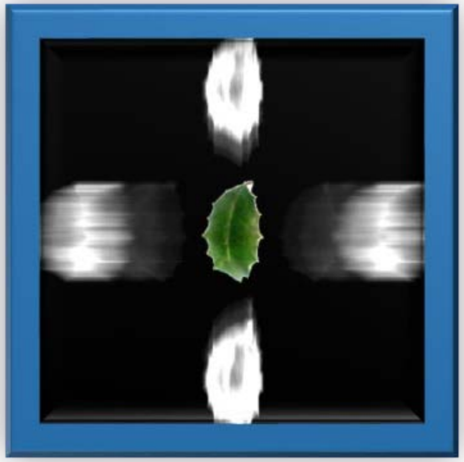

CTIS vs Full resolution spatial

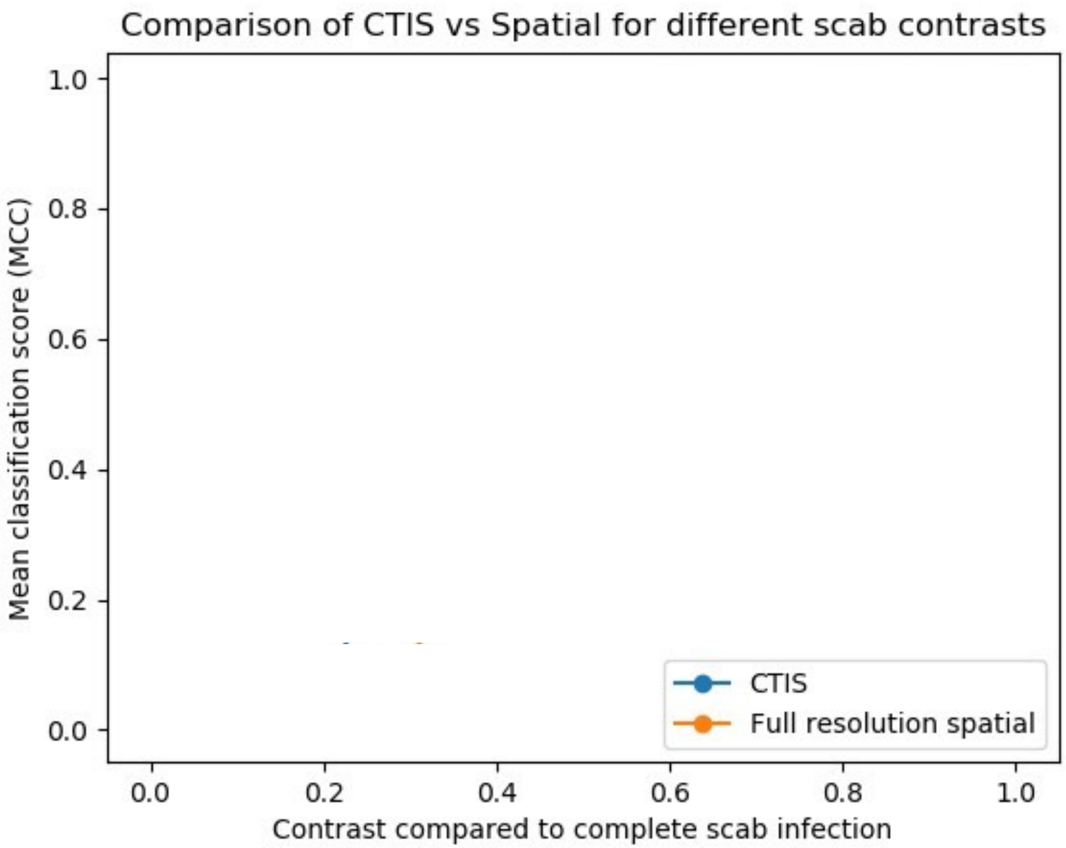

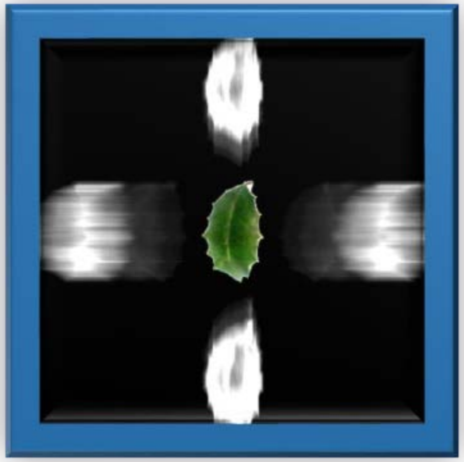

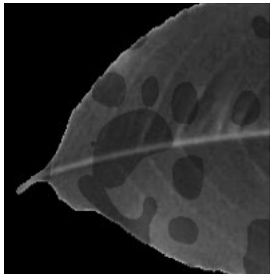

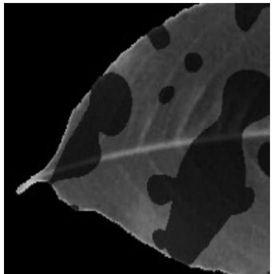

CTIS vs Full resolution spatial

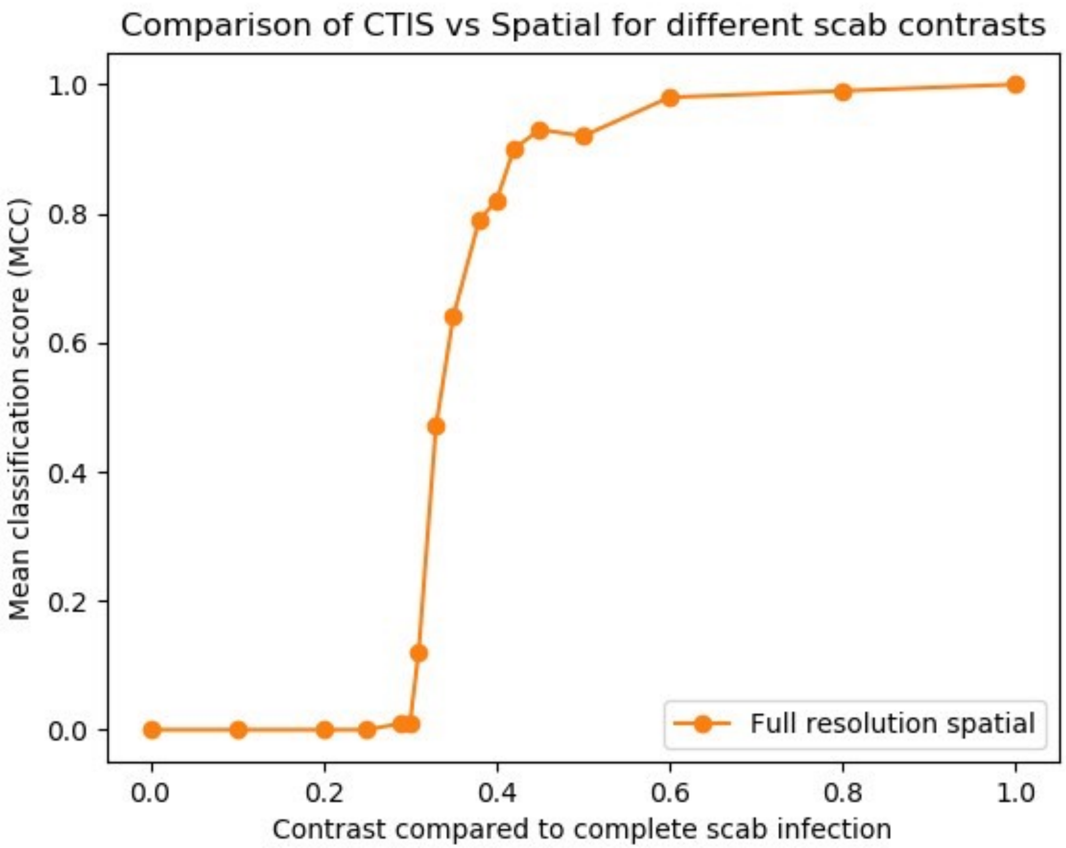

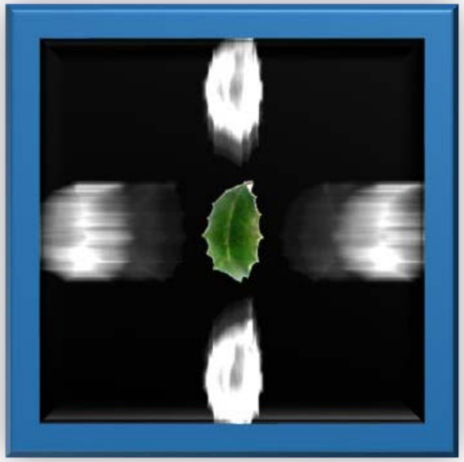

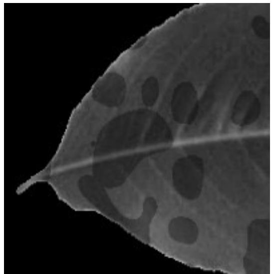

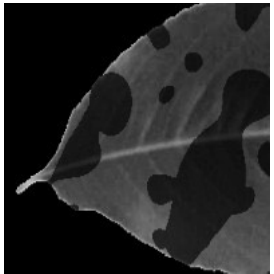

CTIS vs Full resolution spatial

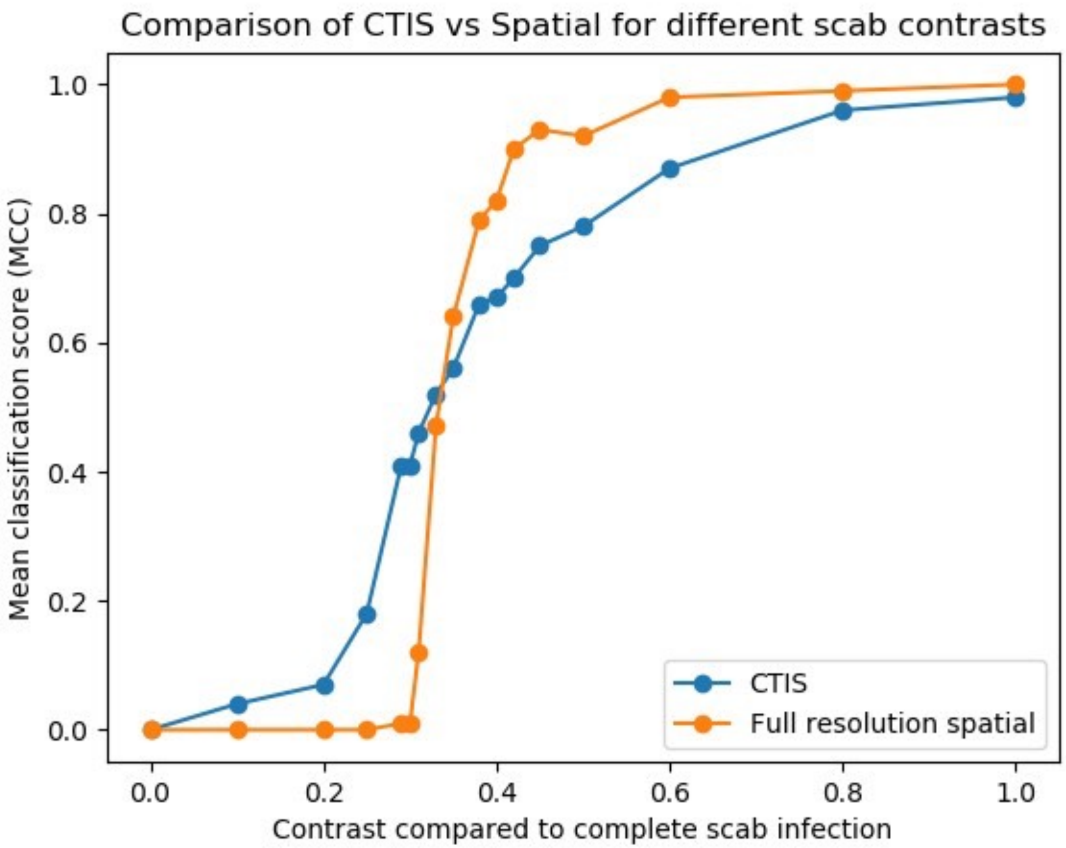

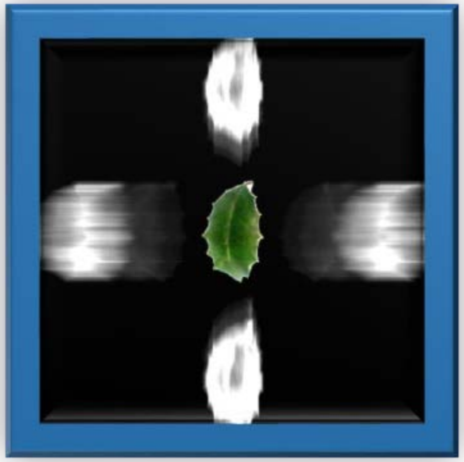

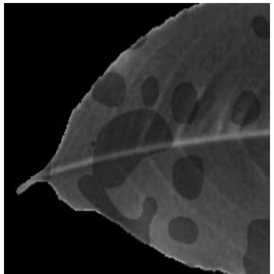

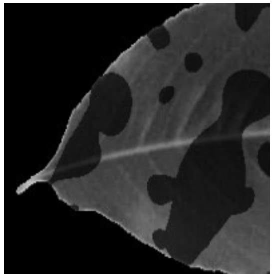

Where is the network looking ?

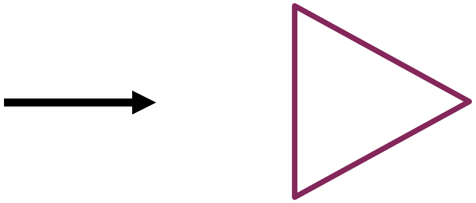

Where is the network looking ?

“This is a cat.”

Where is the network looking ?

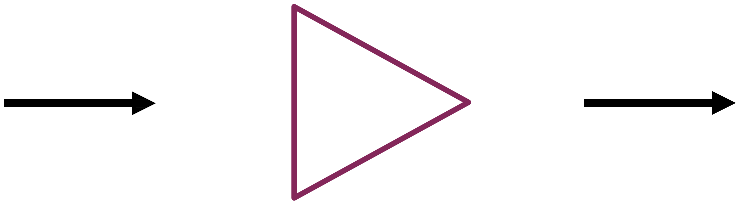

"This is a cat."

Where did the network look ?

Where is the network looking ?

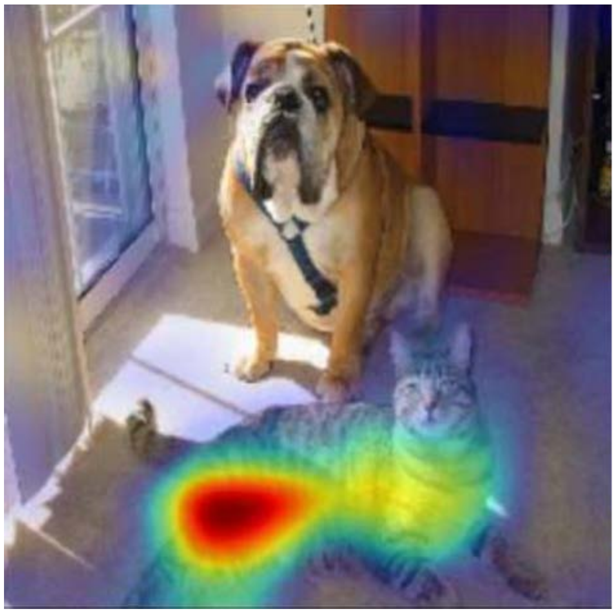

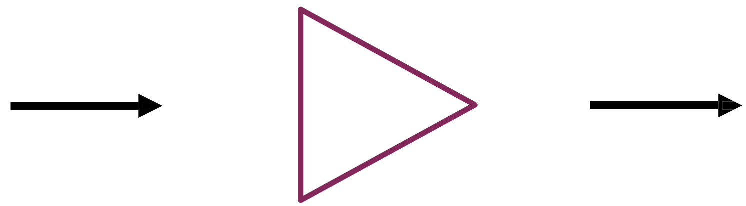

“This is a cat.”

Where did the network look ?

→ Grad-Cam visualization algorithm [5]

5. Selvaraju et al. Grad-cam : Visual explanations from deep networks via gradient-based localization. *ICCV* (2017)

Where is the network looking ?

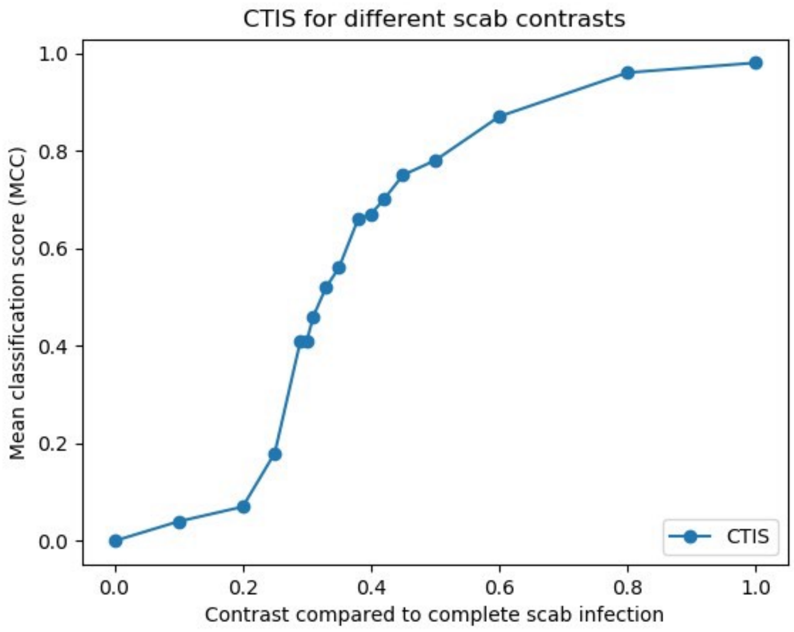

Where is the network looking ?

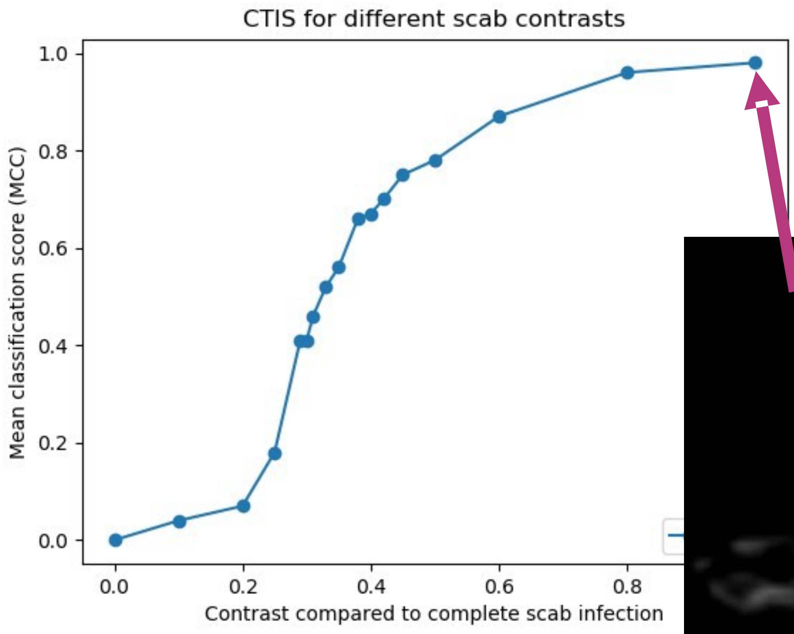

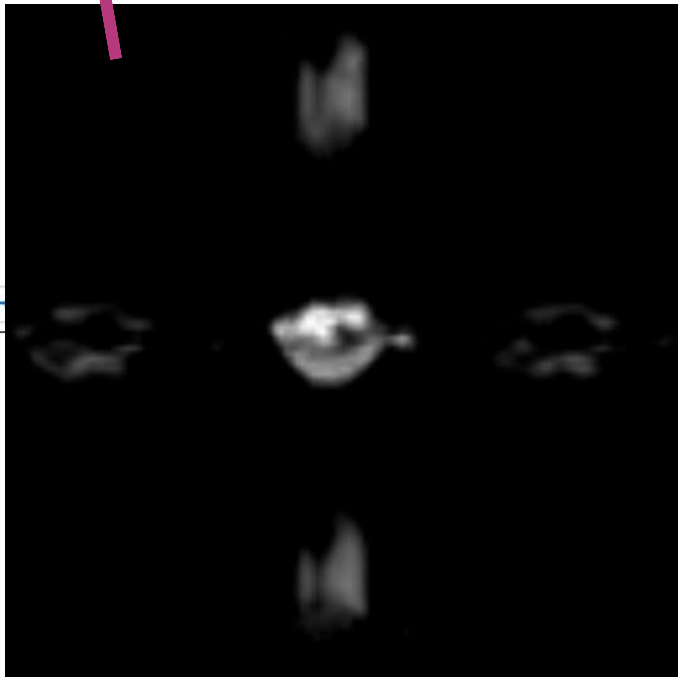

Where is the network looking ?

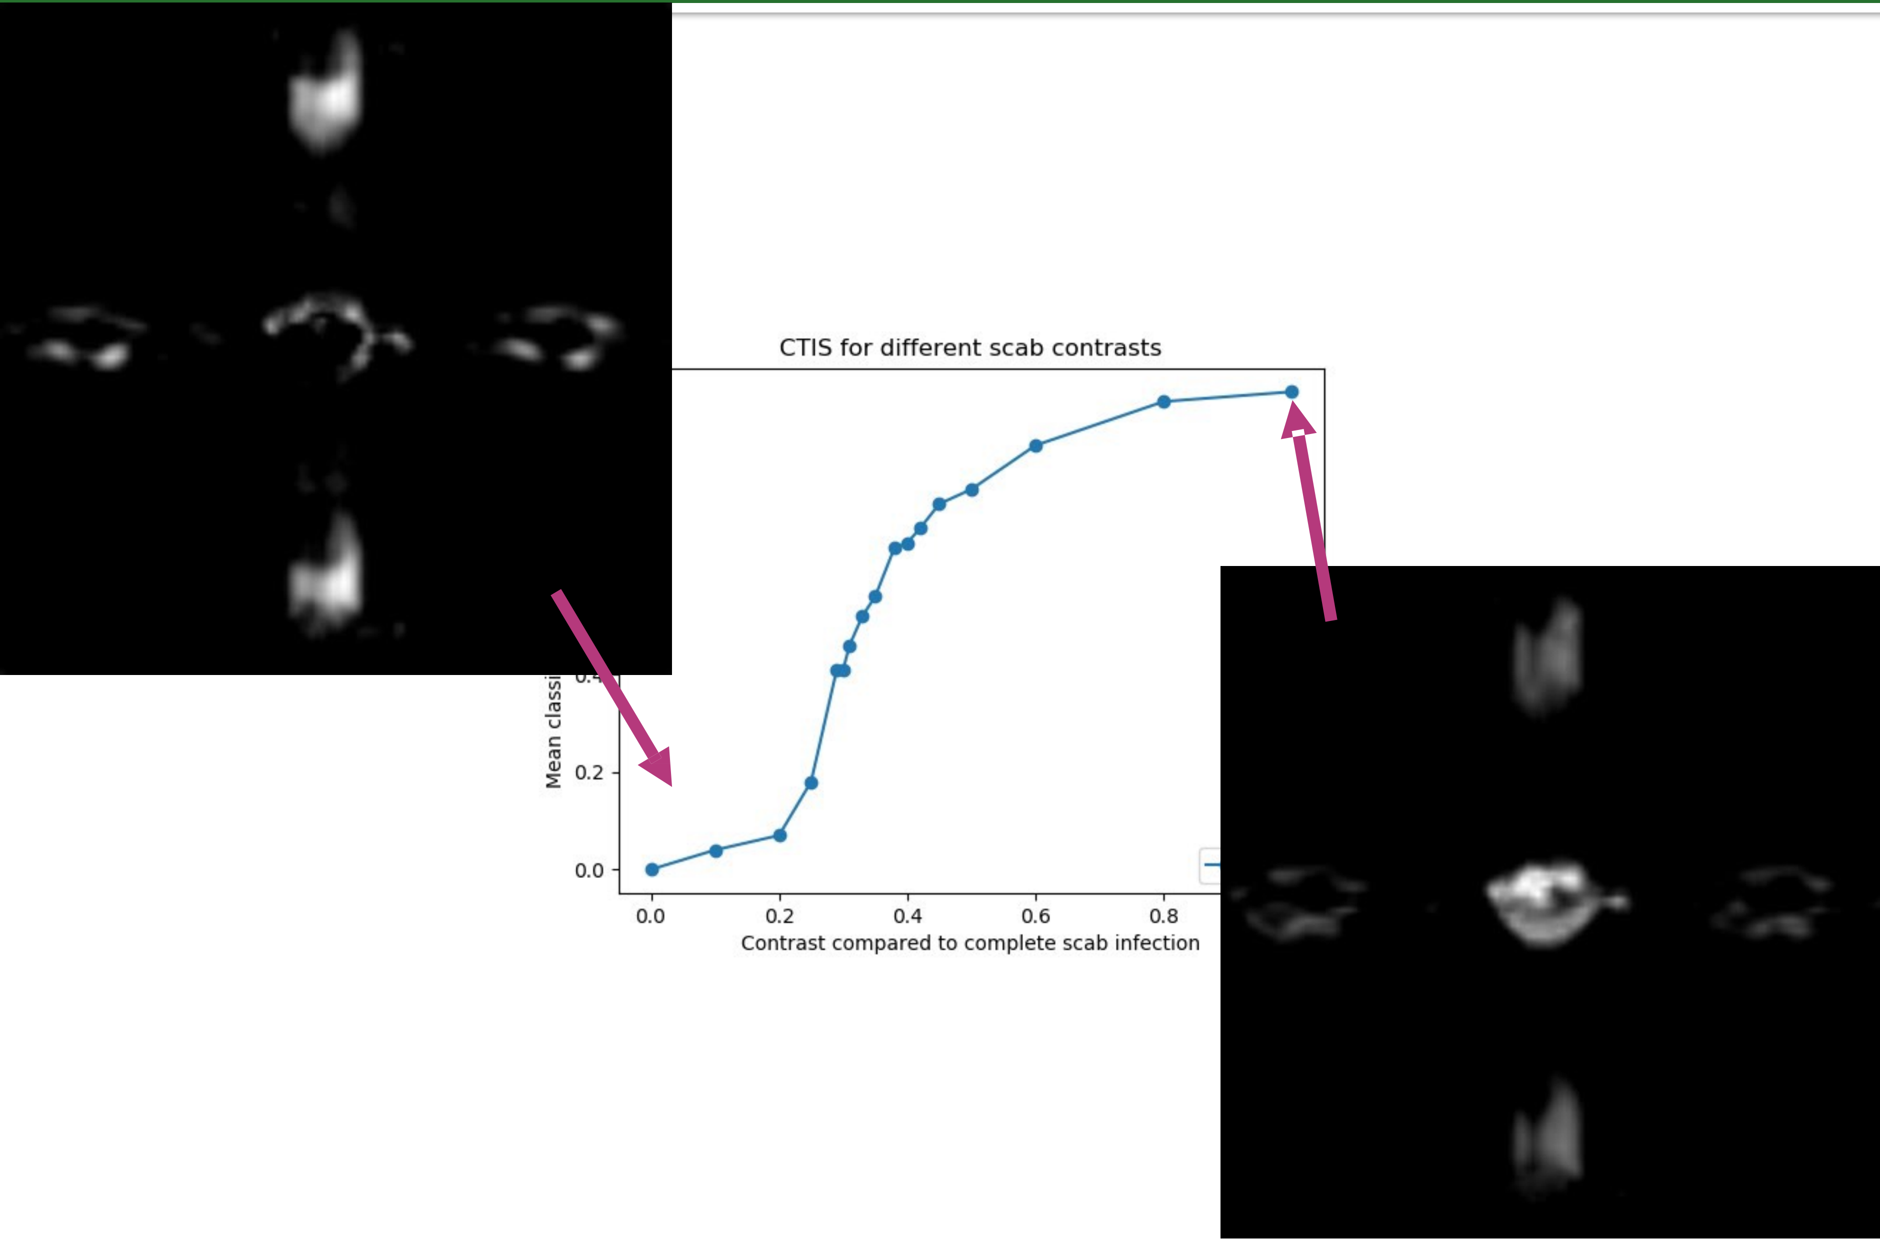

Lowering the cost of supervised machine learning

David Rousseau

Deep learning era

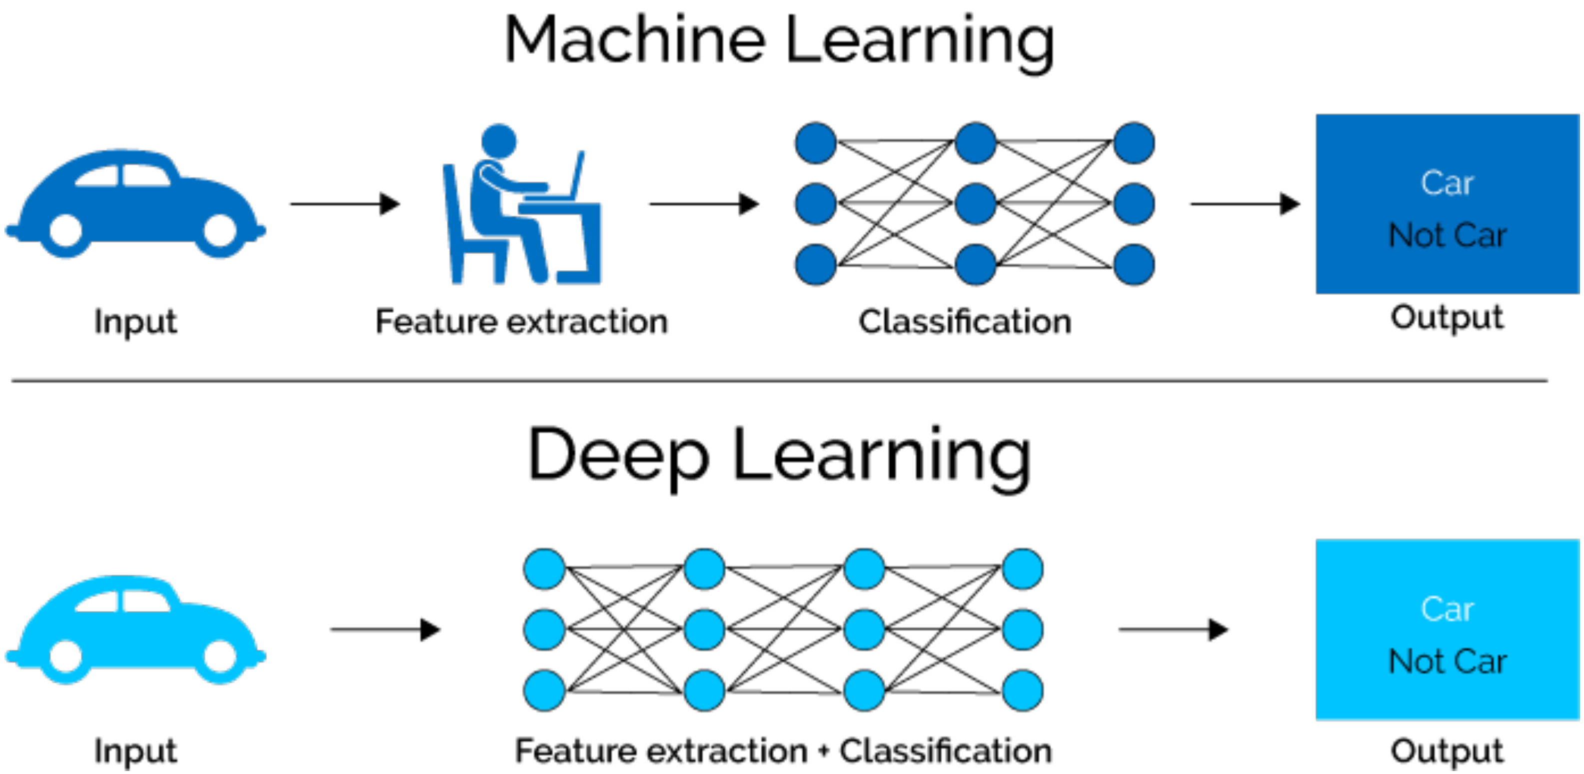

When is deep learning better than classical ML?

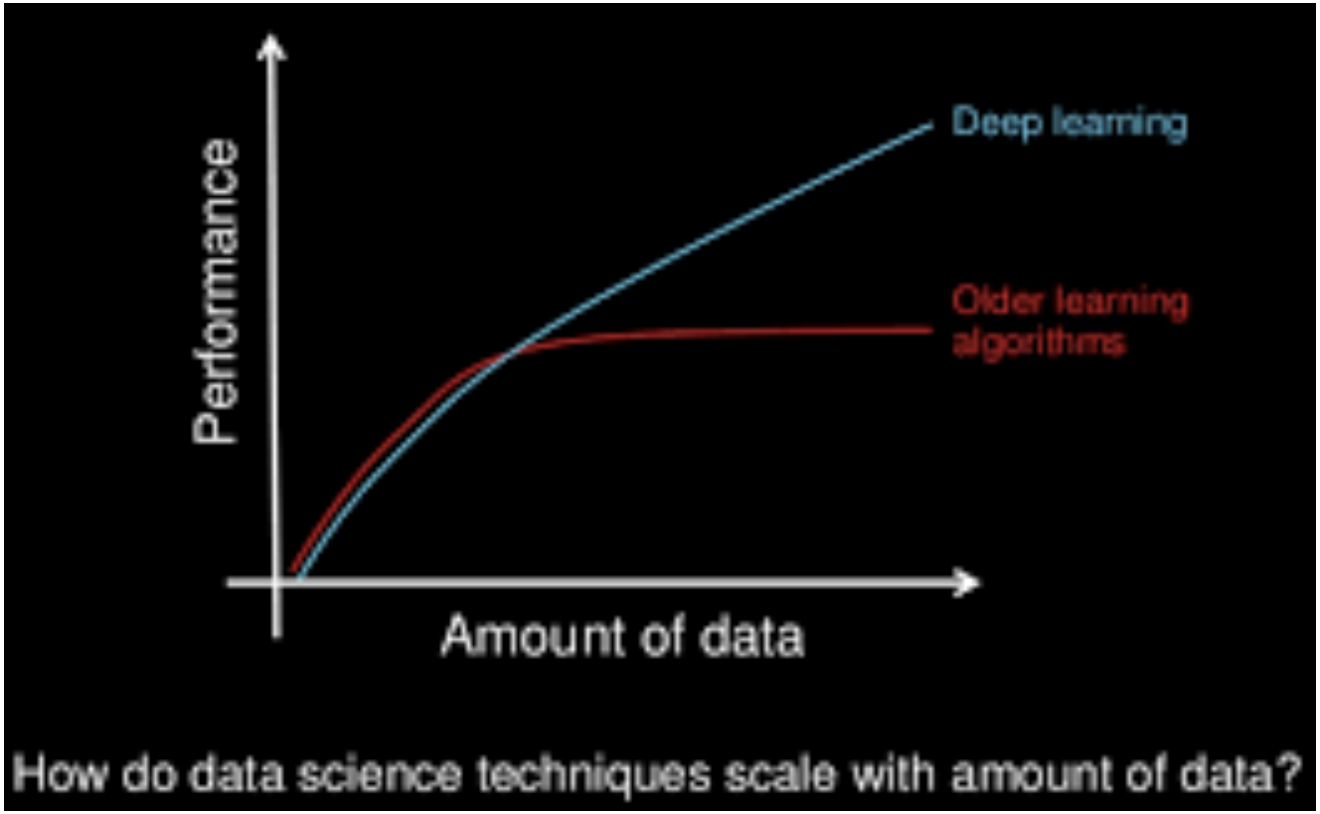

Where does this crossing occurs? => 10^4

Economical consequences

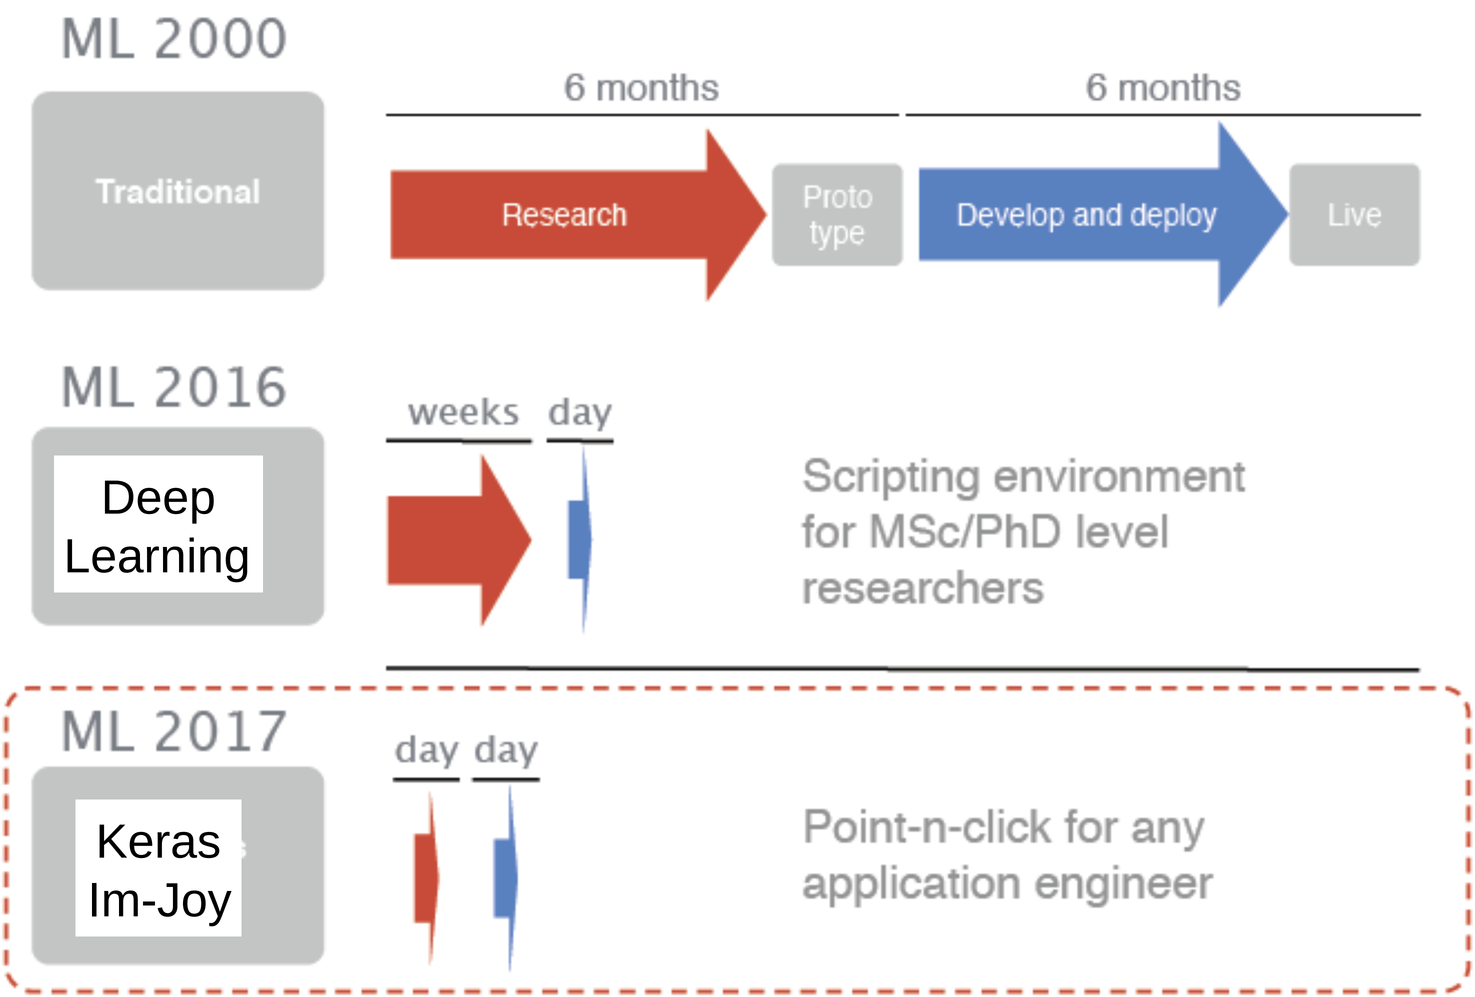

Hidden costs in supervised machine learning

Specific computing systems

Image annotation

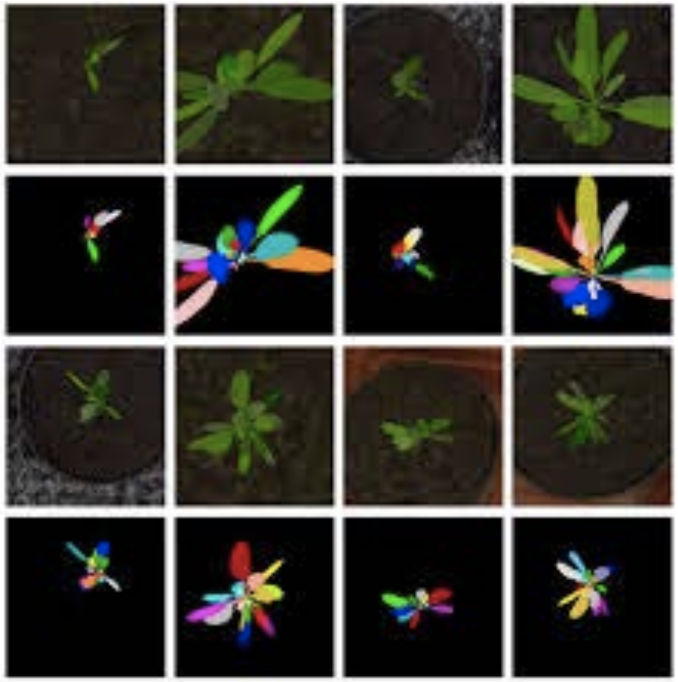

How to speed up annotation ?

- Supervised Machine learning based: Illastic, Weka, Survos, Labelkit, …
- Active learning : Peter, Loïc, et al. "Assisting the examination of large histopathological slides with adaptive forests." *Medical image analysis* 35 (2017)
- Annotate with other colleagues: Cytomine
- Annotate with citizen science : Giuffrida MV, Chen F, Scharr H, Tsaftaris SA. Citizen crowds and experts: observer variability in image-based plant phenotyping. Plant methods. 2018 Dec;14(1):12.
- Pay people to do it for you: Amazon mechanical turk
- …

How to speed up annotation ?

- Supervised Machine learning based: Illastic, Weka, Survos, Labelkit, …
- Active learning : Peter, Loïc, et al. "Assisting the examination of large histopathological slides with adaptive forests." *Medical image analysis* 35 (2017)
- Annotate with other colleagues: Cytomine
- Annotate with citizen science : Giuffrida MV, Chen F, Scharr H, Tsaftaris SA. Citizen crowds and experts: observer variability in image-based plant phenotyping. Plant methods. 2018 Dec;14(1):12.
- Pay people to do it for you: Amazon mechanical turk
- **Learn on synthetic data automatically annotated**

cytomine

amazon mechanical turk

Getting more data with data augmentation

David Rousseau

Data augmentation techniques

1. Horizontal flips

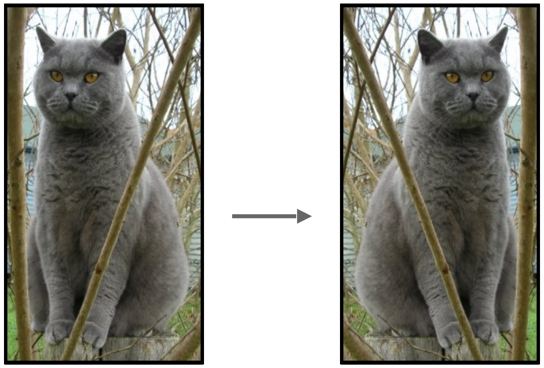

2. Random crops/scales

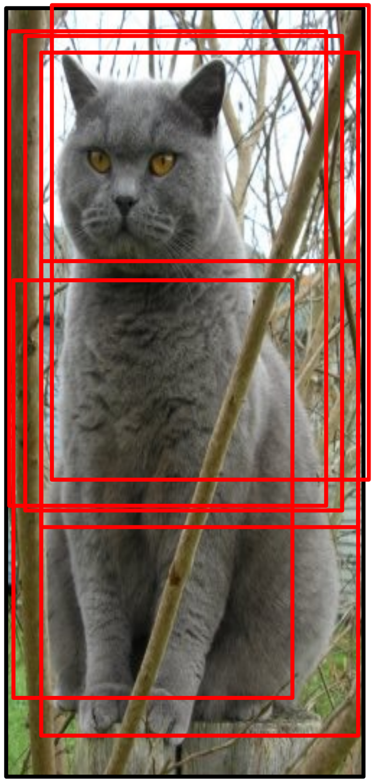

3. Color jitter

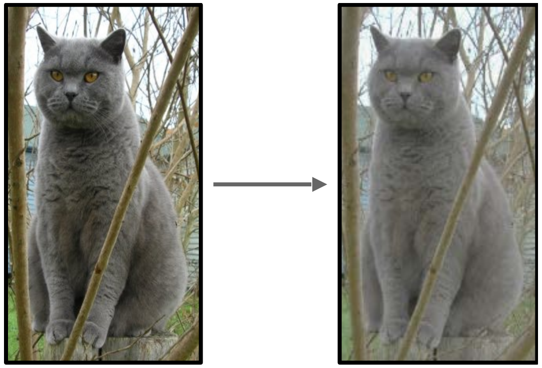

https://keras.io/preprocessing/image/
https://github.com/albu/albumentations
https://imgaug.readthedocs.io/en/latest/
https://github.com/mdbloice/Augmentor

Data augmentation techniques

Any physical parameter you want your algorithm to be insensitive to :

- rotation
- Scaling
- Shearing
- illumination
- lens distortions, …

This is adding invariance, I.e. robustness to the model, via the data

Getting more data with simulation

David Rousseau

How to simulate images?

- Needs synthetic models of living objects
- Needs physical imaging models for image simulator

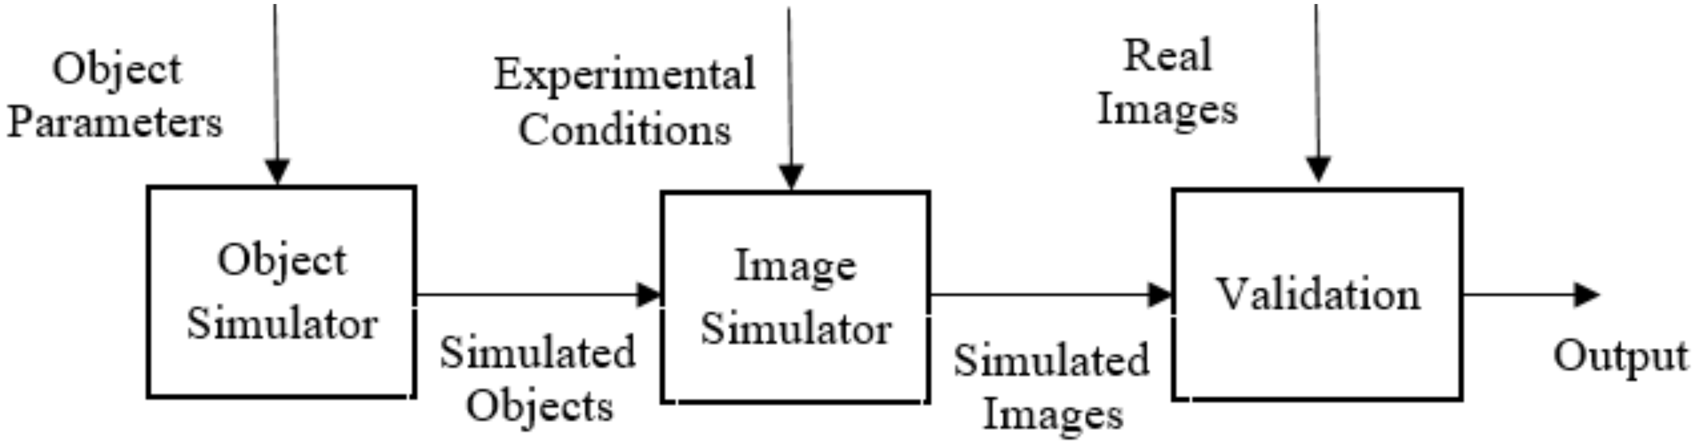

Benoit L, Rousseau D, et al Simulation of image acquisition in machine vision dedicated to seedling elongation to validate image processing root segmentation algorithms. Computers and electronics in agriculture. 2014 Jun 1;104:84-92.

Some evidence of success in computer vision

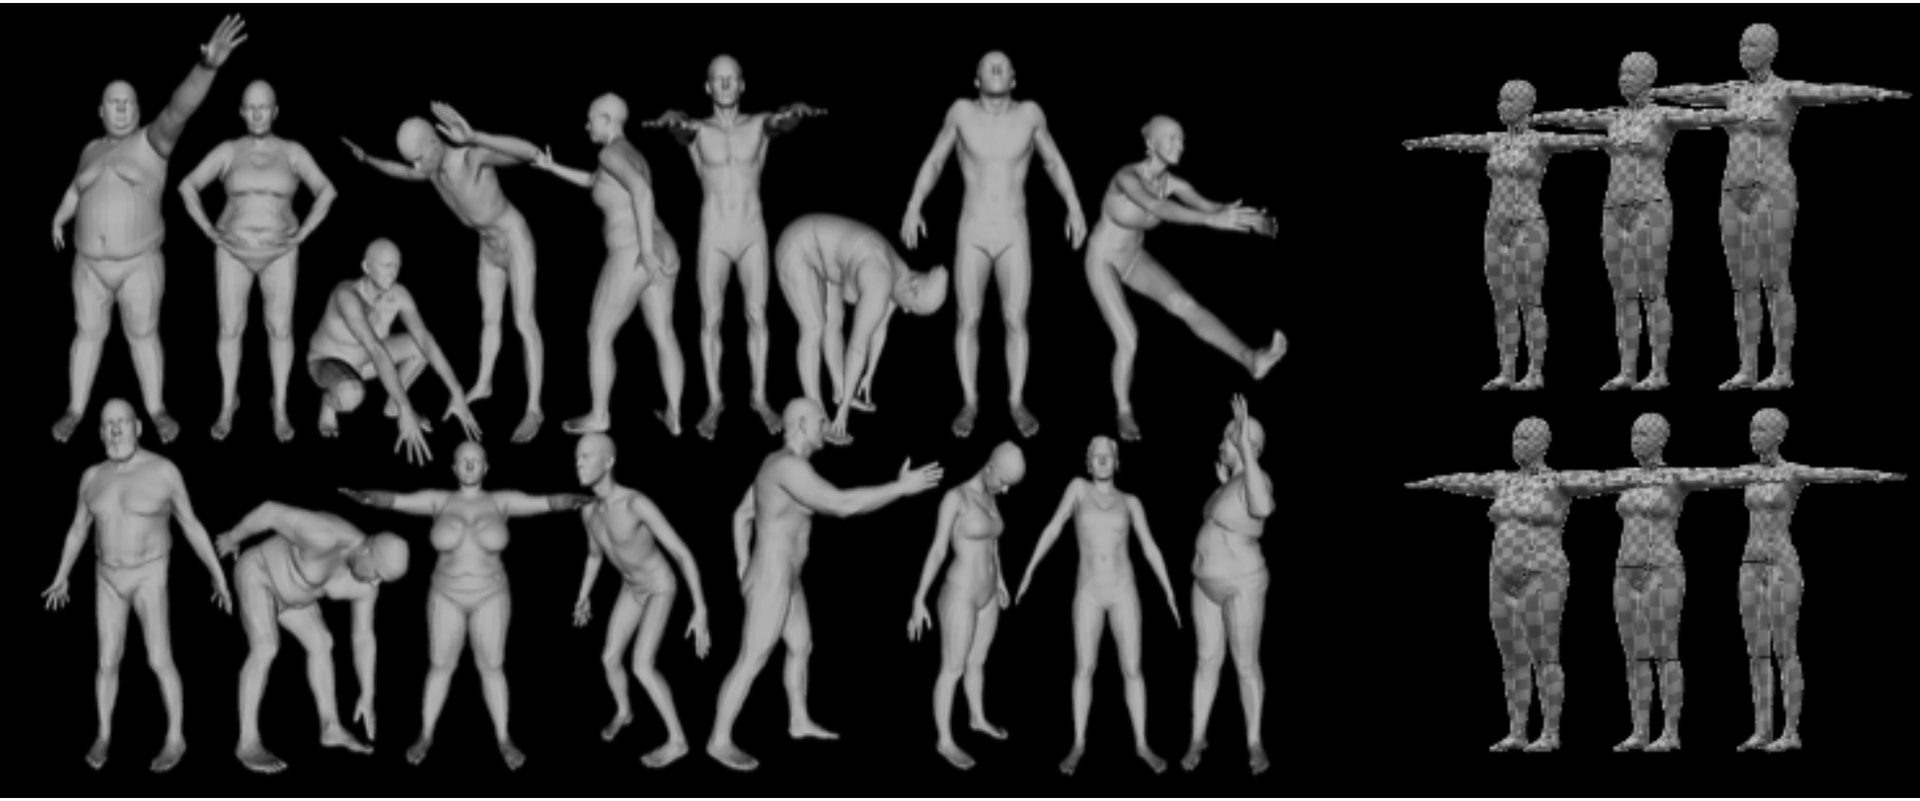

Computer Graphics has great models for generating body shapes and poses

[Loper *et al.*, "SMPL: A Skinned Multi-Person Linear Model", SIGGRAPH Asia, 2015]

Synthetic depth images

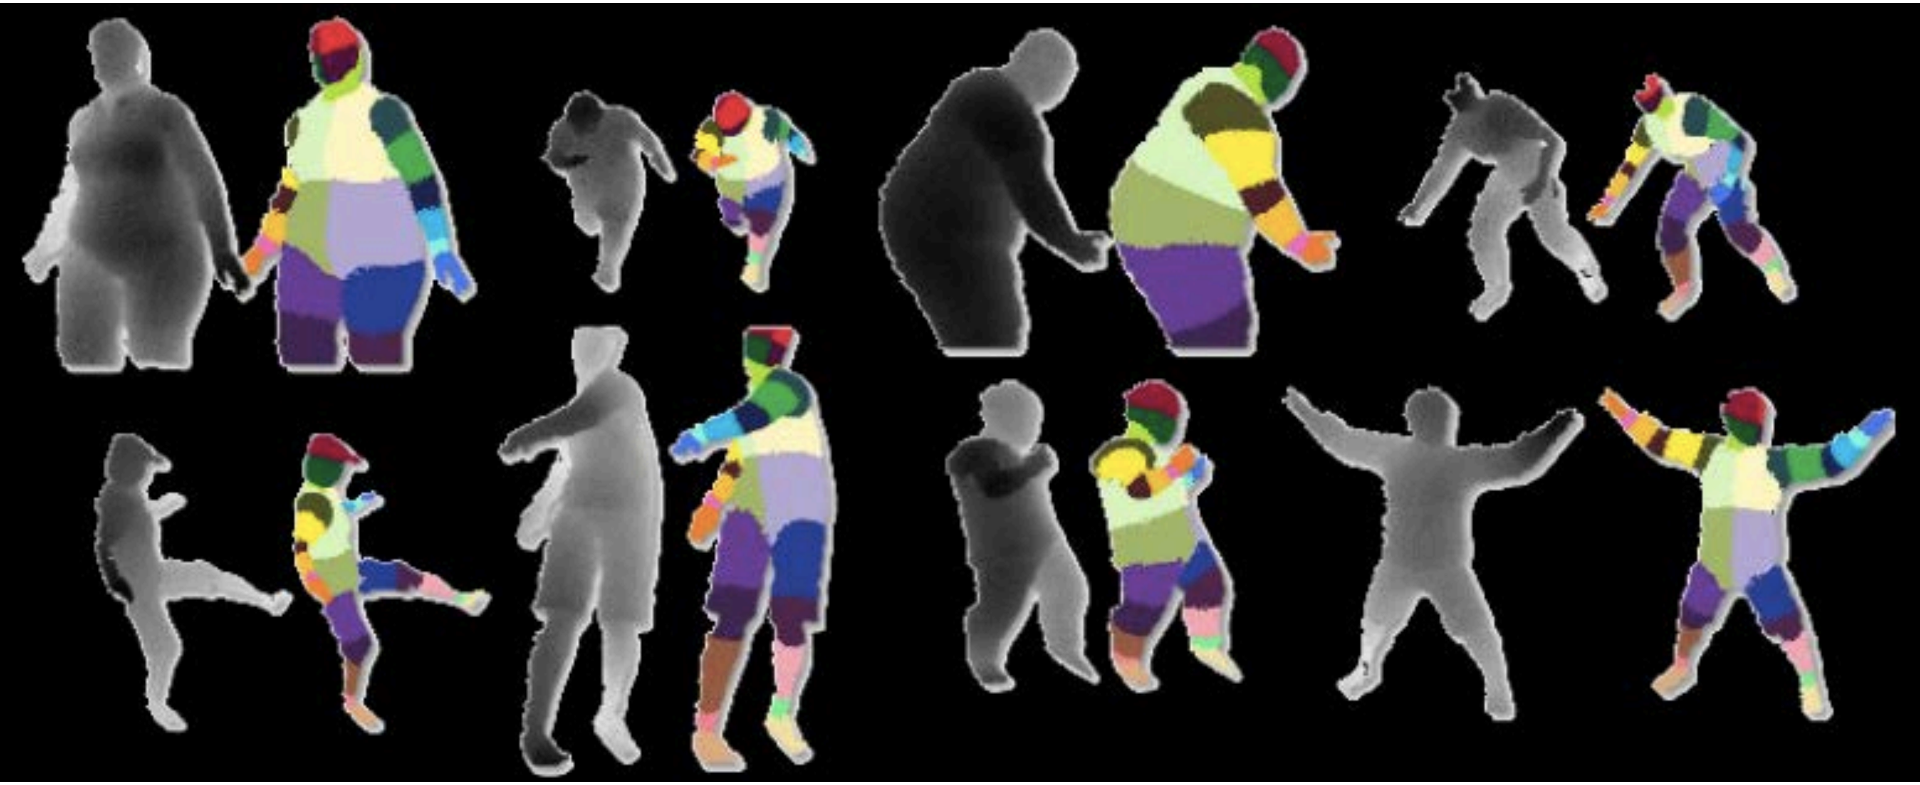

Possible to create unlimited amount of training data

Shotton *et al.*, "Real-Time Human Pose Recognition in Parts from Single Depth Images", CVPR 2011]

https://www.blensor.org/tag/lidar.html

Application to train classifiers

- Human pose estimation

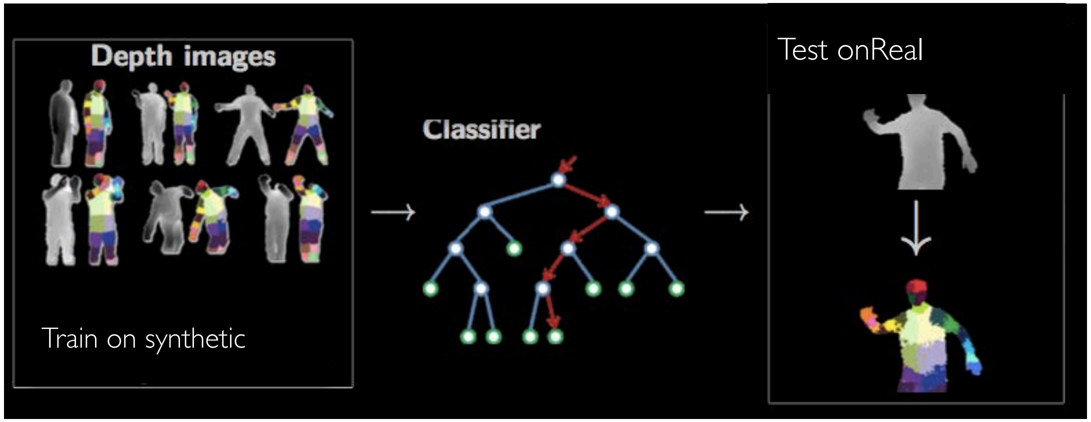

- Semantic segmentation of real scene

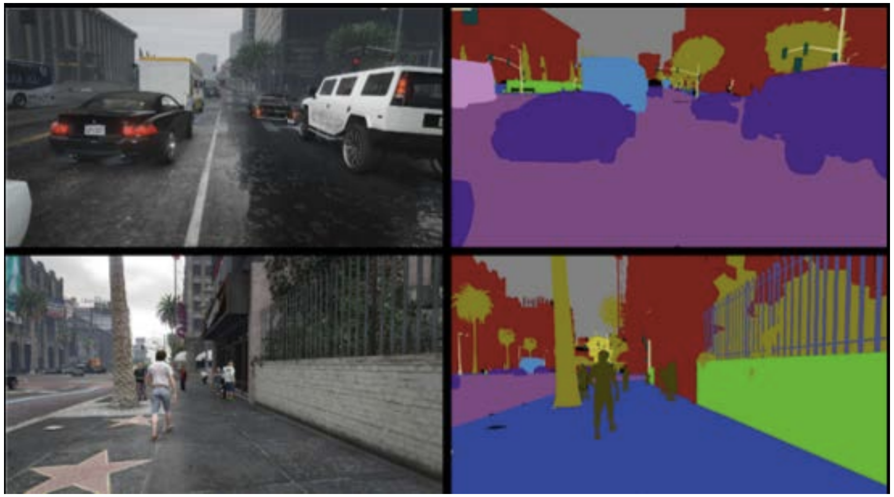

[Richter *et al.*, "Playing for Data: Ground Truth from Computer Games", ECCV 2016]

Compensation of discrepancy between real and simulated

David Rousseau

When simulated data fail?

- It is simple, when simulated data **do not match** the real data
- We saw the failures:
  - Pose estimation: the synthetic data simulate well reality but
  - Denoising: synthetically added noise is not representative of the real (sensor/process) noise

- In mathematical terms… the **distributions must overlap**

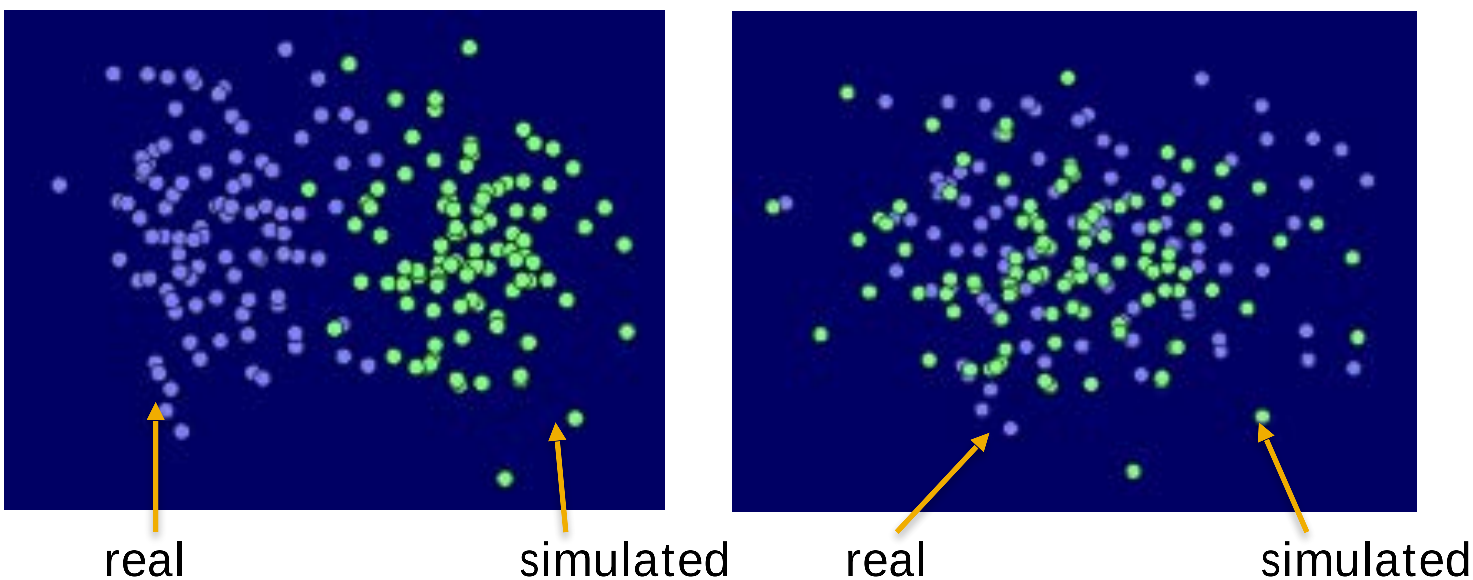

Transfer Learning

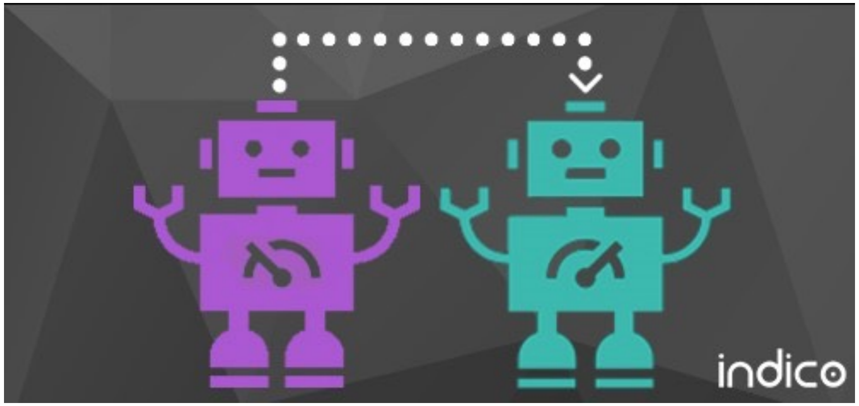

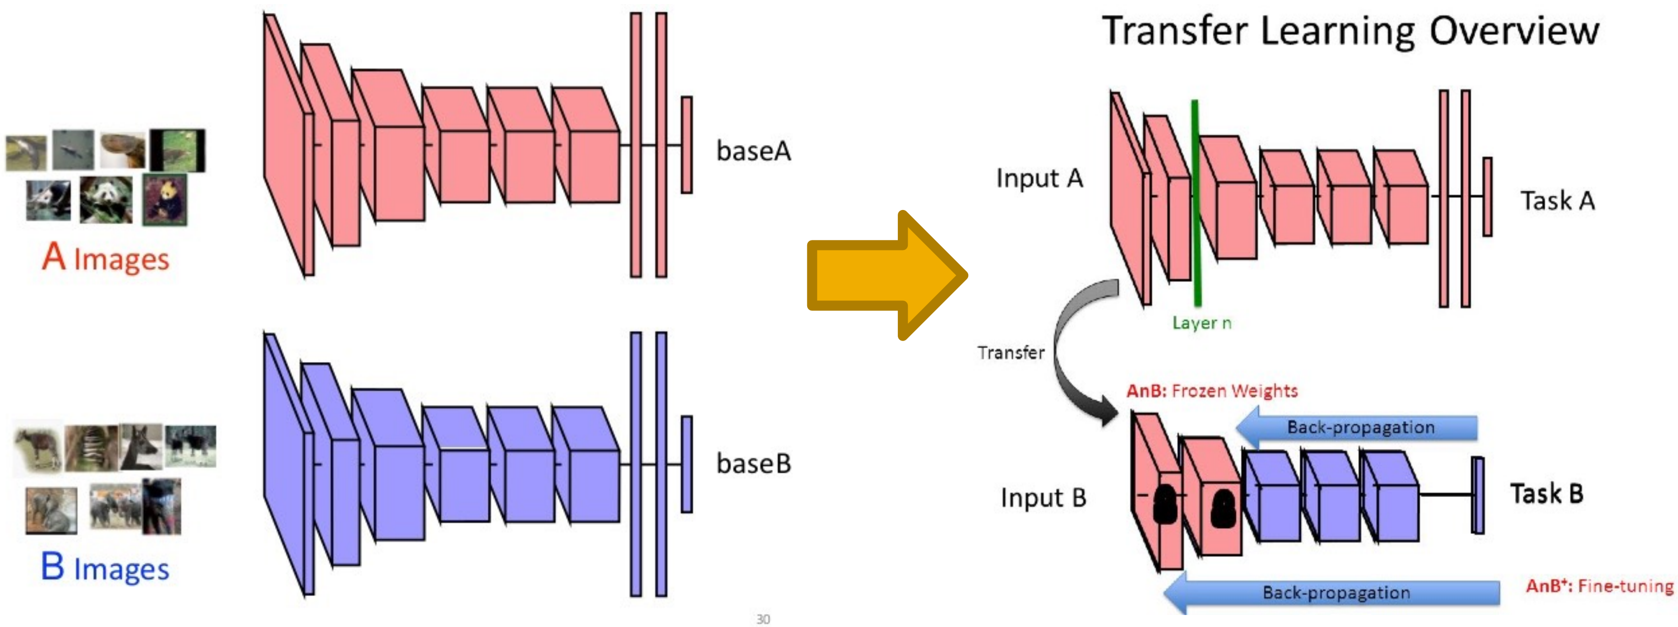

Getting more data with generative models

David Rousseau

Generative models

Task: generate new samples following the same probabilistic distribution of a given training dataset

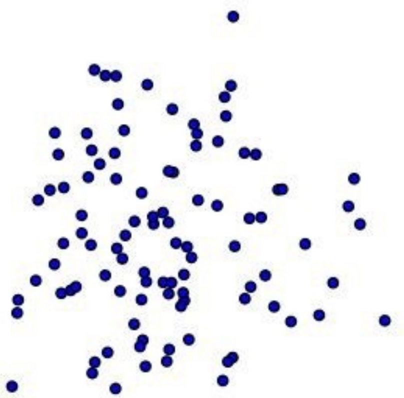

Training samples

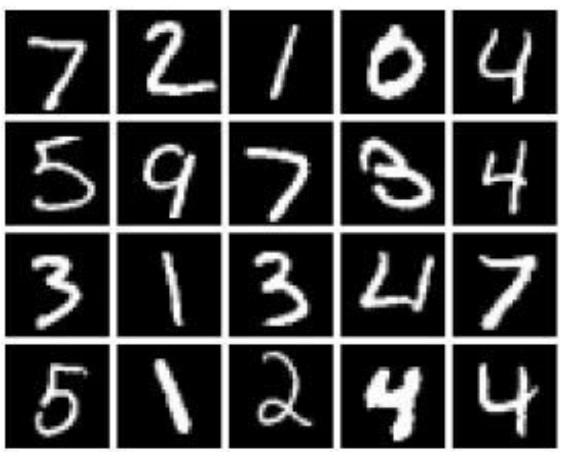

$$p(x) = ?$$

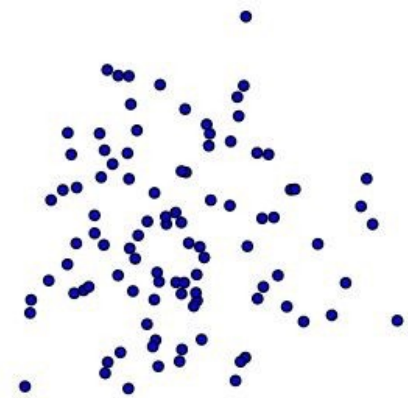

Generated samples

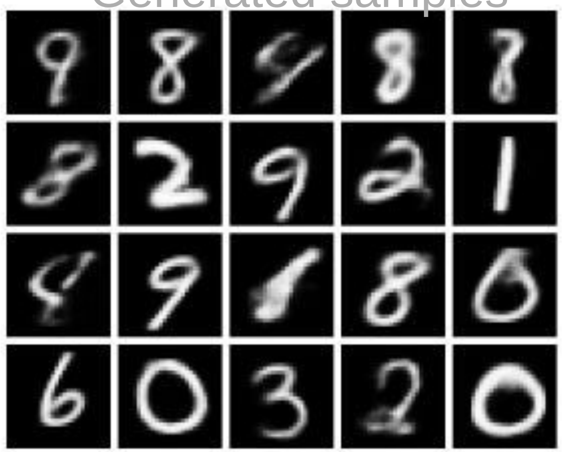

Note: sometimes it's fine if we cannot estimate the explicit form of p(x), since it might be over complicated

Tribute S. Wang

A parenthesis: sampling from a difficult distribution

Sample an easy one and transform it

Examples:

- Box-Muller transform: uniform ➔ Normal

$$u_1, u_2 \sim \text{uniform([0,1])} \quad \rightarrow \quad v = \sqrt{-2 \log u_1} \cos(2\pi u_2) \sim \mathcal{N}(0,1)$$

- Normal ➔ Gaussian

$$u \sim \mathcal{N}(0,1) \quad \rightarrow \quad v = \sigma u + \mu \;\sim \mathcal{N}(\mu, \sigma^2)$$

- General 1D distribution P with cdf Φ

$$u \sim \text{uniform([0,1])} \quad \rightarrow \quad v = \Phi^{-1}(u) \sim P$$

Variational auto-encoders (VAE)

How do we learn the parameters?

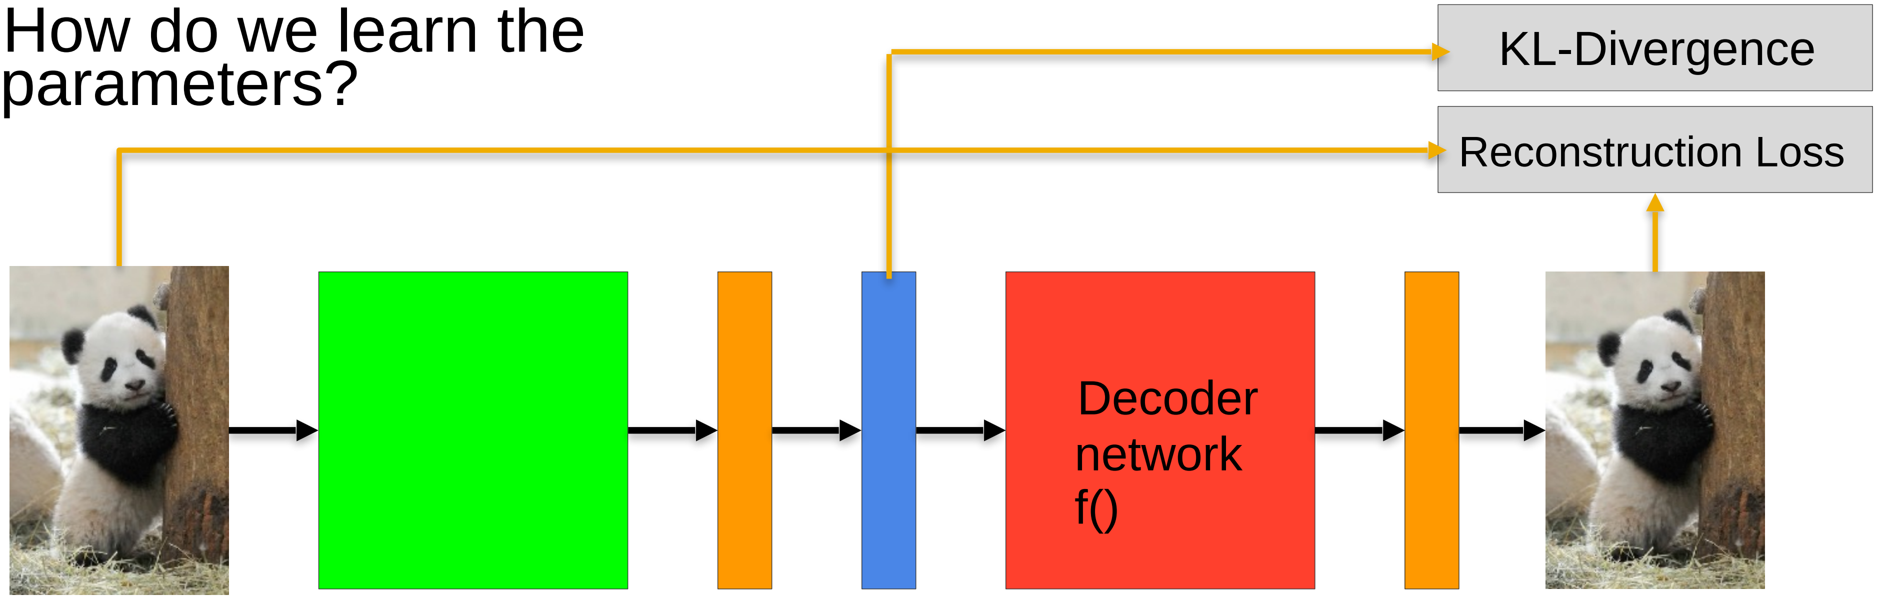

Input Image x

Code z

Reconstruction

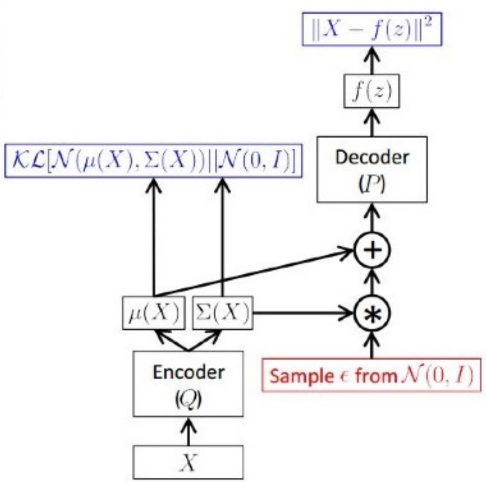

Computation graph
Credit: Doersch

Tribute S. Wang

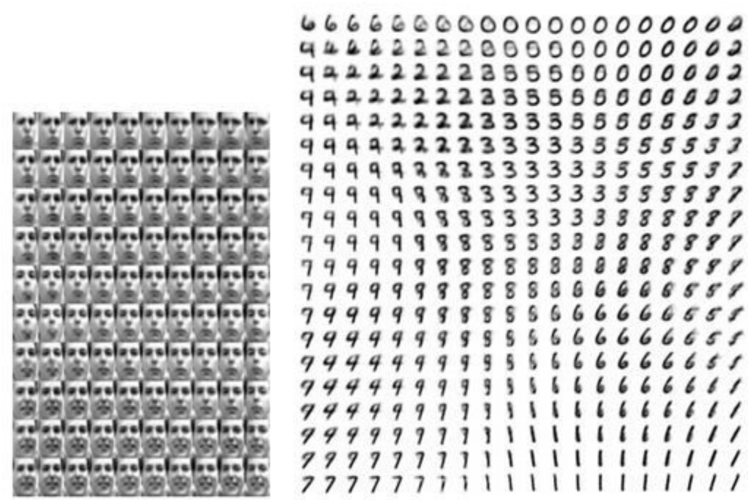

(a) Learned Frey Face manifold

(b) Learned MNIST manifold

Generative adversarial Network (GAN)

Code z

Generator G

Generated Image

Training Image

Discriminator D

Fake (0)

Real (1)

0.5 not sure

Train G to make it impossible for D to make good predictions

Train D to have high values on real data and low ones on output of G

Generative adversarial Network (GAN)

Discriminator tries the best to distinguish whether the image is generated by computers or not

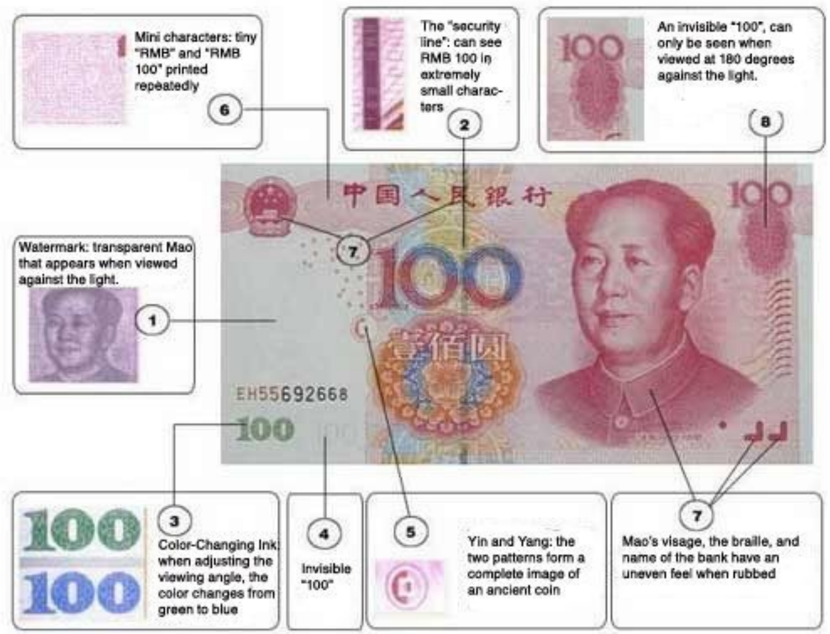

Generator tries the best to cheat the discriminator by generating more realistic images

TributeS. Wang

Conclusion on generative models

- **VAEs**:
  - Easier to train
  - Blurry result due to minimizing the MSE based reconstruction error
  - Nice probabilistic formulation, easy to introduce prior

- **GANs**:
  - High-quality visually appealing result
  - Difficult to train (mode collapse, training schedule)

Getting more data… full monty a comparative use case

David Rousseau

Apple scab segmentation

Figure 15 : Classes phénotypiques de pommiers infectés par le champignon responsable de la tavelure *Venturia inaequalis* selon Chevalier et al. (1991)
classe 0 : aucun symptôme visible ; classe 1 : symptômes caractéristiques de « pin-point » ; classe 2 : symptômes de résistance (chlorose, nécrose, crispation) sans sporulation ; classe 3a : symptômes de résistance avec quelques tâches de sporulation peu abondante ; classe 3b : symptômes de résistance avec tâches de sporulation abondante ; classe 4 : pas de symptôme de résistance et sporulation abondante

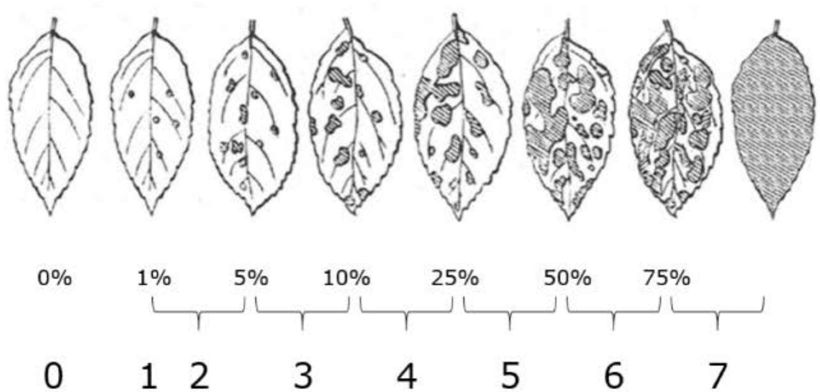

Figure 16 : Adaptation de l'échelle de sévérité de sporulation de Croxall et al. (1952)
0 : aucune sporulation visible ; 1% de surface foliaire couvert de sporulation : 1 ; de 1 à 5% : 2 ; de 5 à 10% : 3 ; de 10 à 25% : 4 ; de 25 à 50% : 5 ; de 50 à 75% : 6 ; de 75 à 100% : 7

Difficult with human eye (RGB)

Much easier with IR

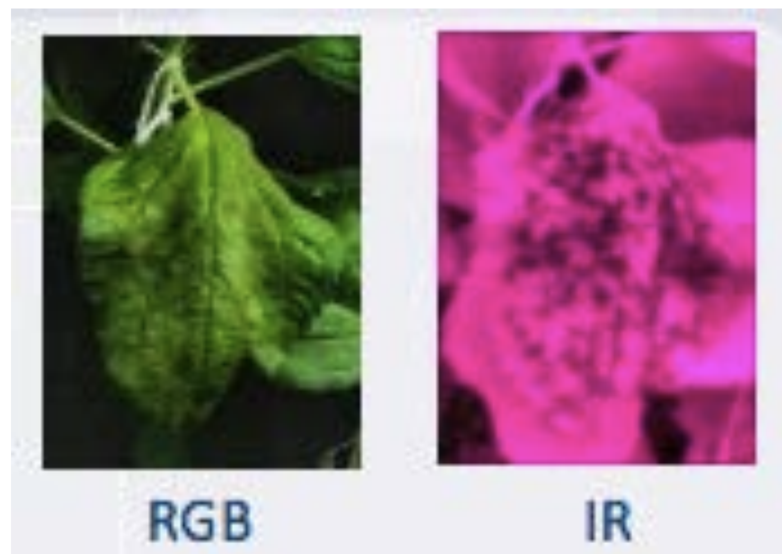

C. Douarre et al COMPAG

The Full Monty

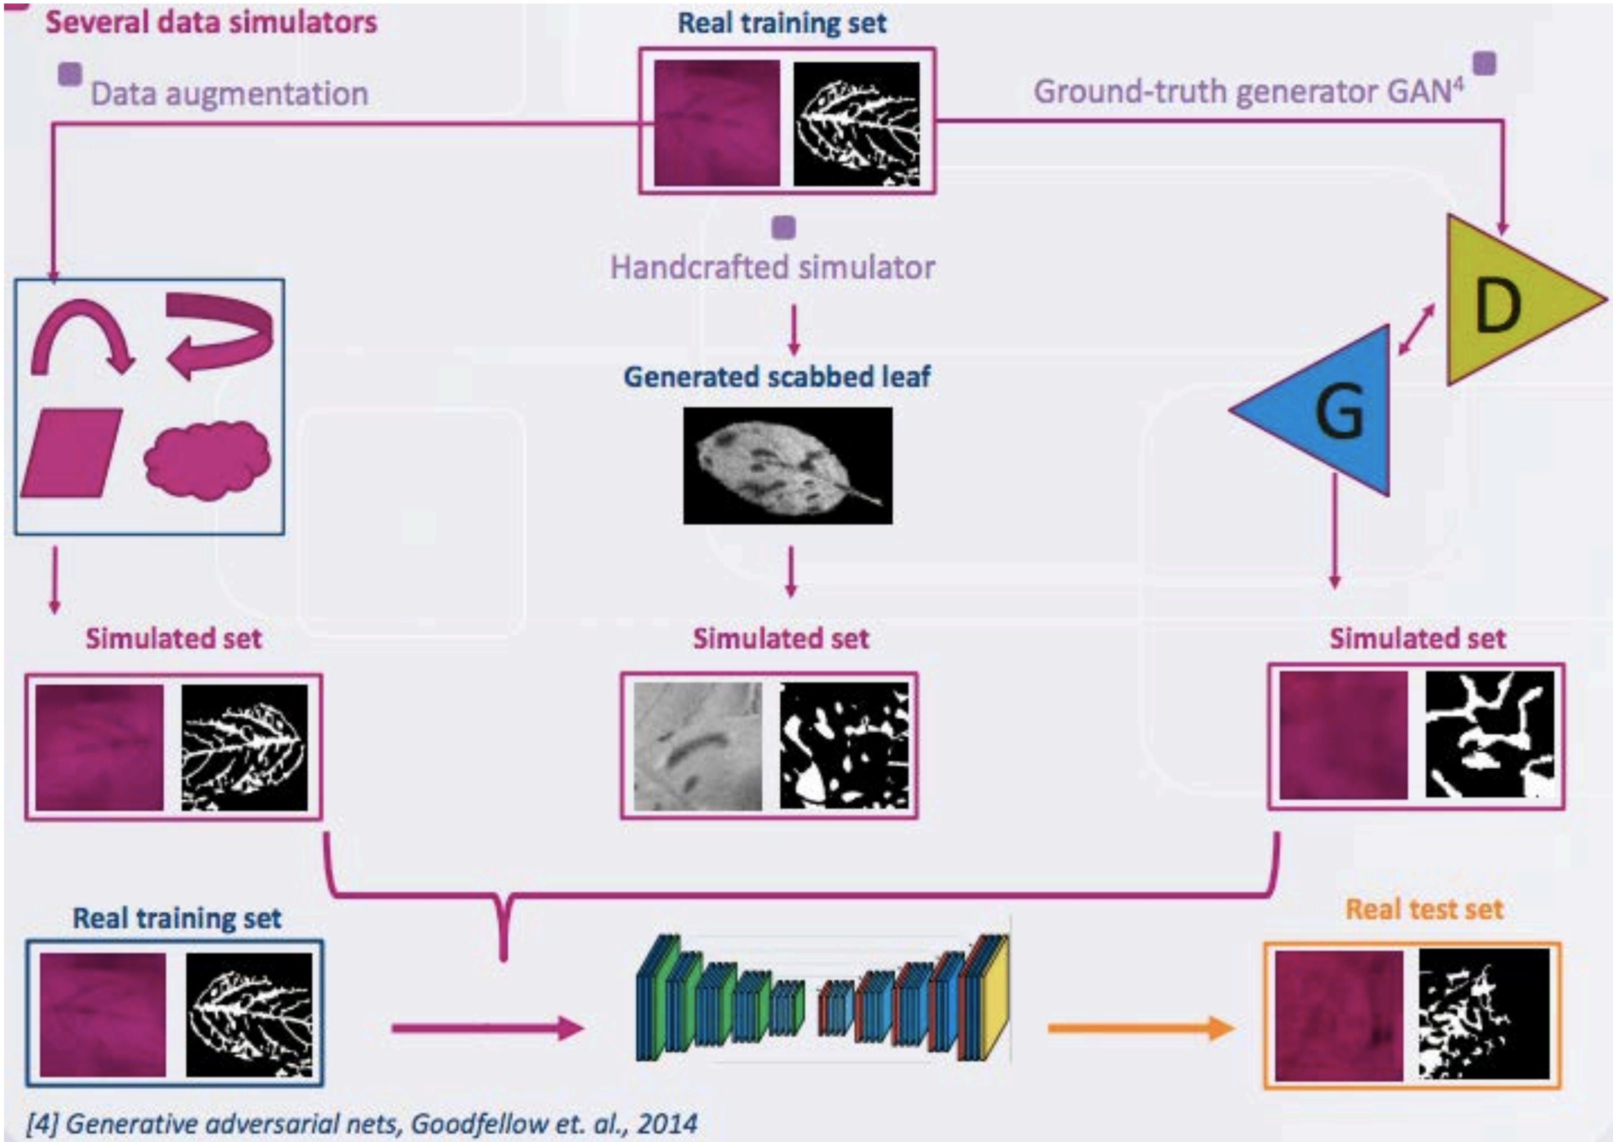

C. Douarre et al COMPAG

Results

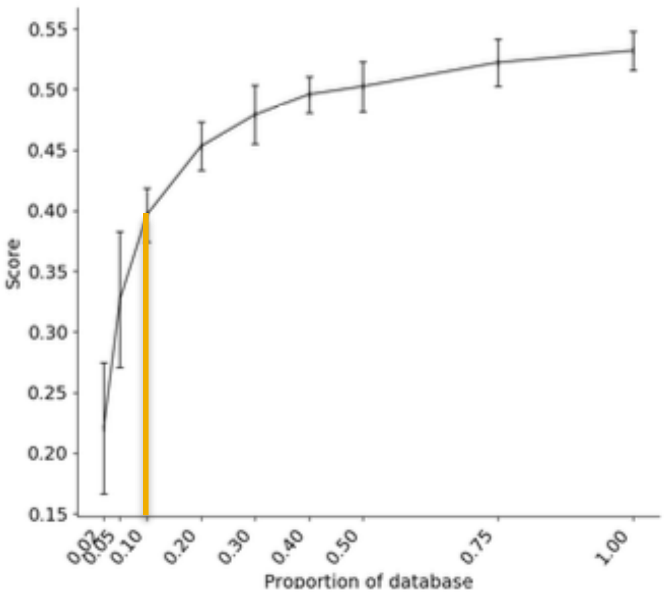

Helper data set	From scratch	Transfer learning
None	0,405	0,462
Data augmentation	0,498	0,517
Model-based simulator	0,391	0,547
GAN-based simulator	0,450	0,451

C. Douarre et al COMPAG

Results

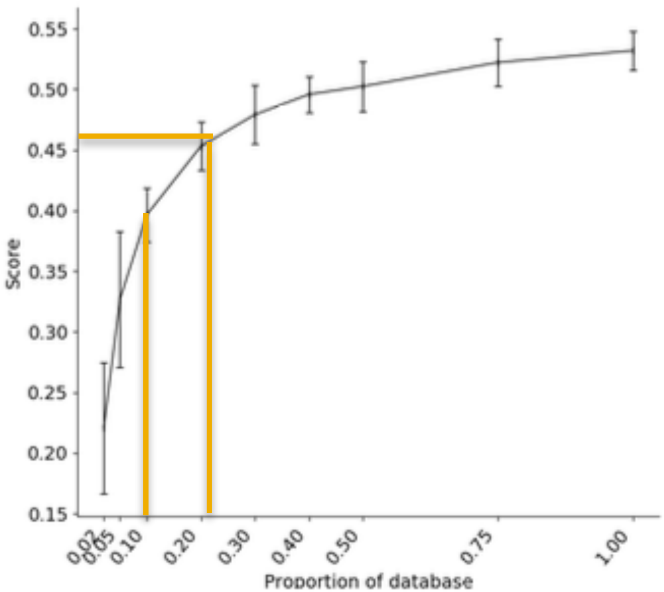

Helper data set	From scratch	Transfer learning
None	0,405	0,462
Data augmentation	0,498	0,517
Model-based simulator	0,391	0,547
GAN-based simulator	0,450	0,451

C. Douarre et al COMPAG

Results

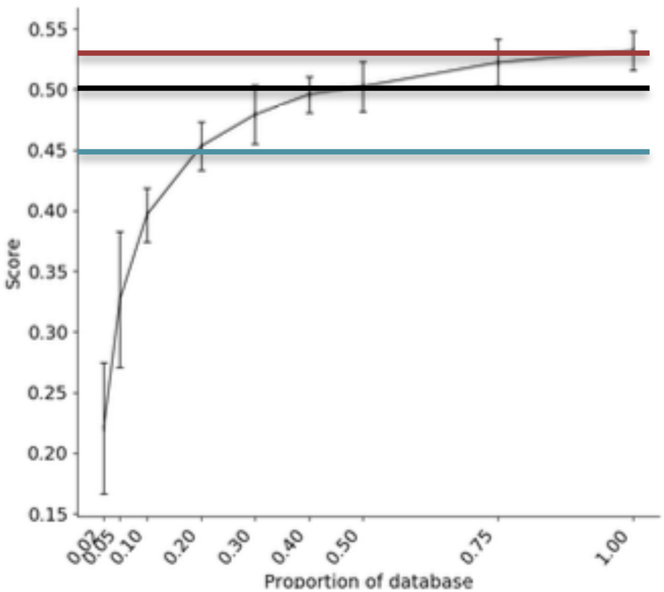

Helper data set	From scratch	Transfer learning
None	0,405	0,462
Data augmentation	0,498	0,517
Model-based simulator	0,391	0,547
GAN-based simulator	0,450	0,451

C. Douarre et al COMPAG

Data augmentation … the full Monty

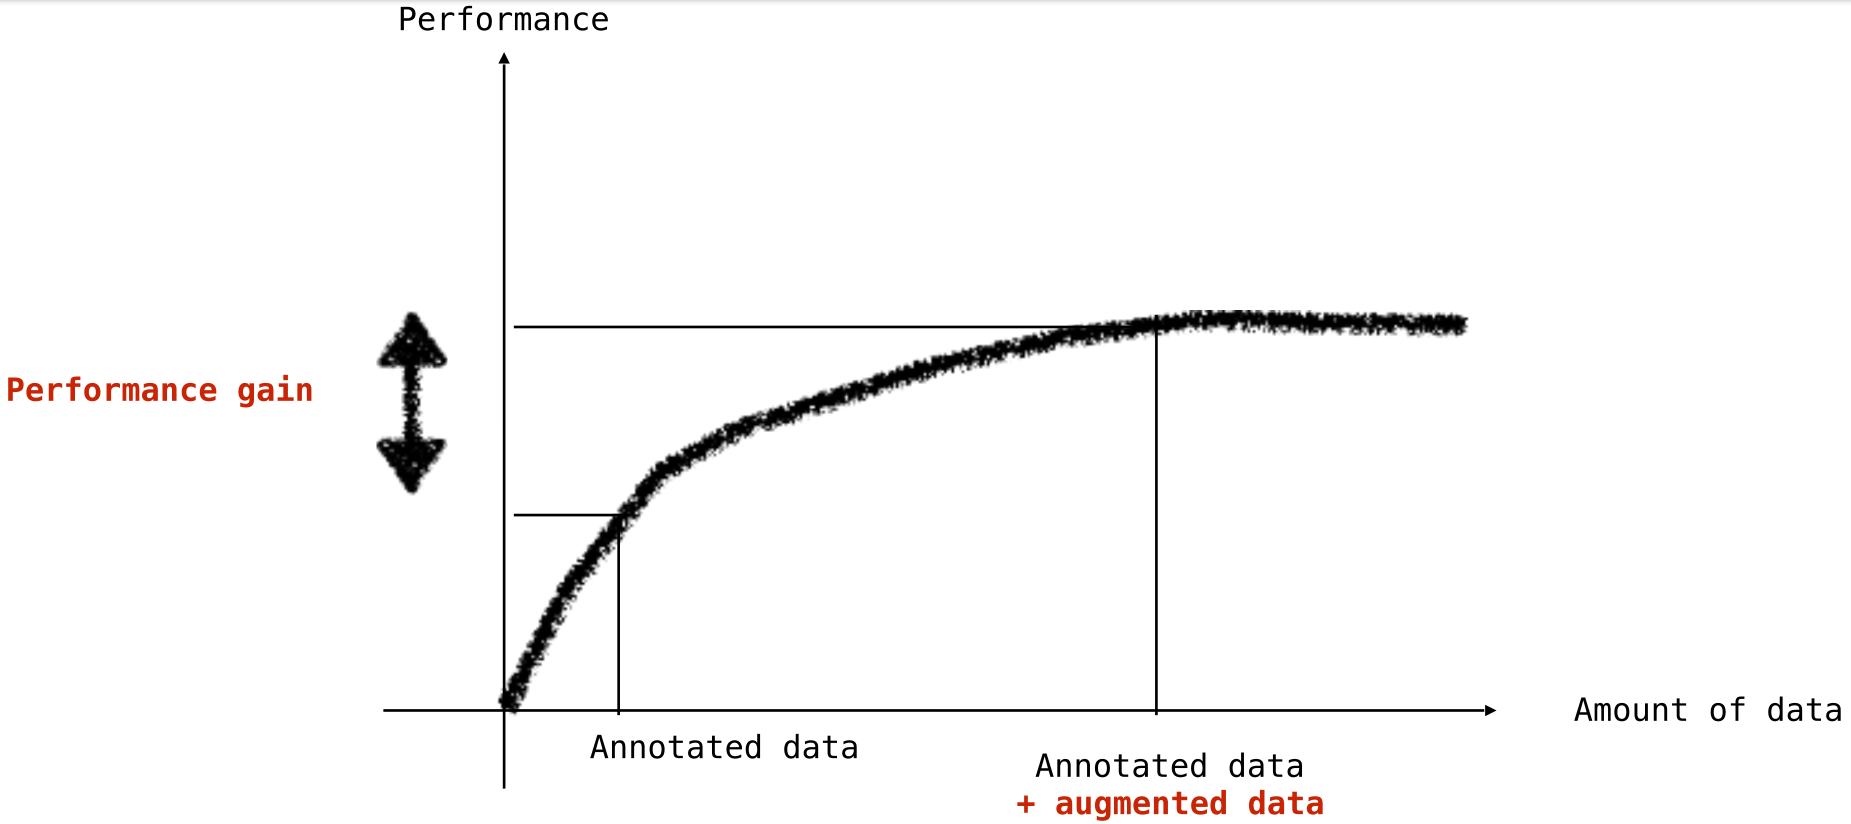

Different approaches: Standard data augmentation, simulation and Generative adversarial network

Which one is the best? : in our case the simulator and data augmentation, nothing to be expected from GAN

C. Douarre et al COMPAG

When spectral imaging meets machine learning

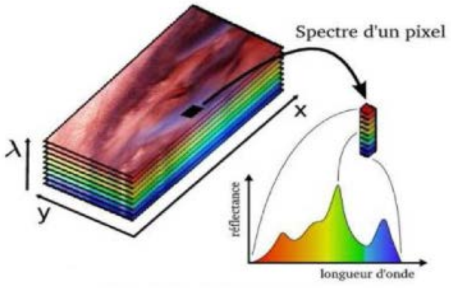

Building cost-effective spectral imaging with Statistical learning

Benoit, Landry, Romain Benoit, Étienne Belin, Rodolphe Vadaine, Didier Demilly, François Chapeau-Blondeau, and David Rousseau. "On the value of the Kullback–Leibler divergence for cost-effective spectral imaging of plants by optimal selection of wavebands." *Machine Vision and Applications* 27, no. 5 (2016): 625-635.

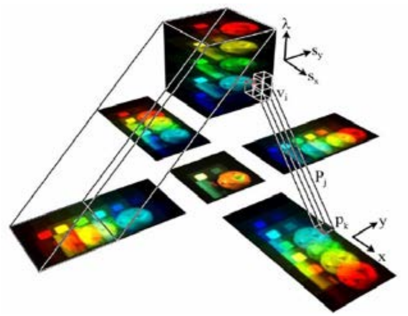

Low-cost spectro-imaging & compressed learning

Douarre, C., Crispim-Junior, C. F., Gelibert, A., Germain, G., Tougne, L., & Rousseau, D. (2021). CTIS-Net: A Neural Network Architecture for Compressed Learning Based on Computed Tomography Imaging Spectrometers. *IEEE Transactions on Computational Imaging*, *7*, 572-583.

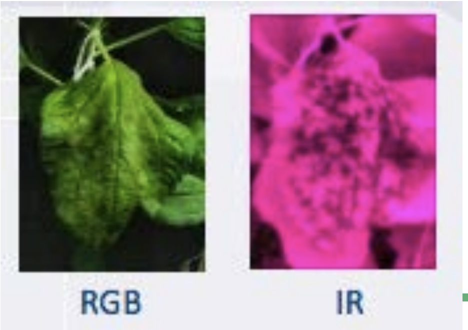

Lowering the cost of annotation in machine learning

Douarre, Clément, Carlos F. Crispim-Junior, Anthony Gelibert, Laure Tougne, and David Rousseau. "Novel data augmentation strategies to boost supervised segmentation of plant disease." *Computers and Electronics in Agriculture* 165 (2019): 104967.

Thanks for your
deep human attention

david.rousseau@univ-angers.fr