Image coding for coalitional active learning

Prospective applications in medicine, in satellite imaging and in videosurveillance

Benoit Macq (CORESA 21)

Table of content

1. **Coalitional machine learning: an example**
2. Image coding for machine learning
3. (Machine learning for image coding)
4. Distributed (federated) learning
5. Federated byzantine agreements for trusted coalitions
6. Coalitional active learning
7. Image coding for coalitional active learning
    - Medical
    - Satellite imaging

1. Coalitional learning: an example in protontherapy

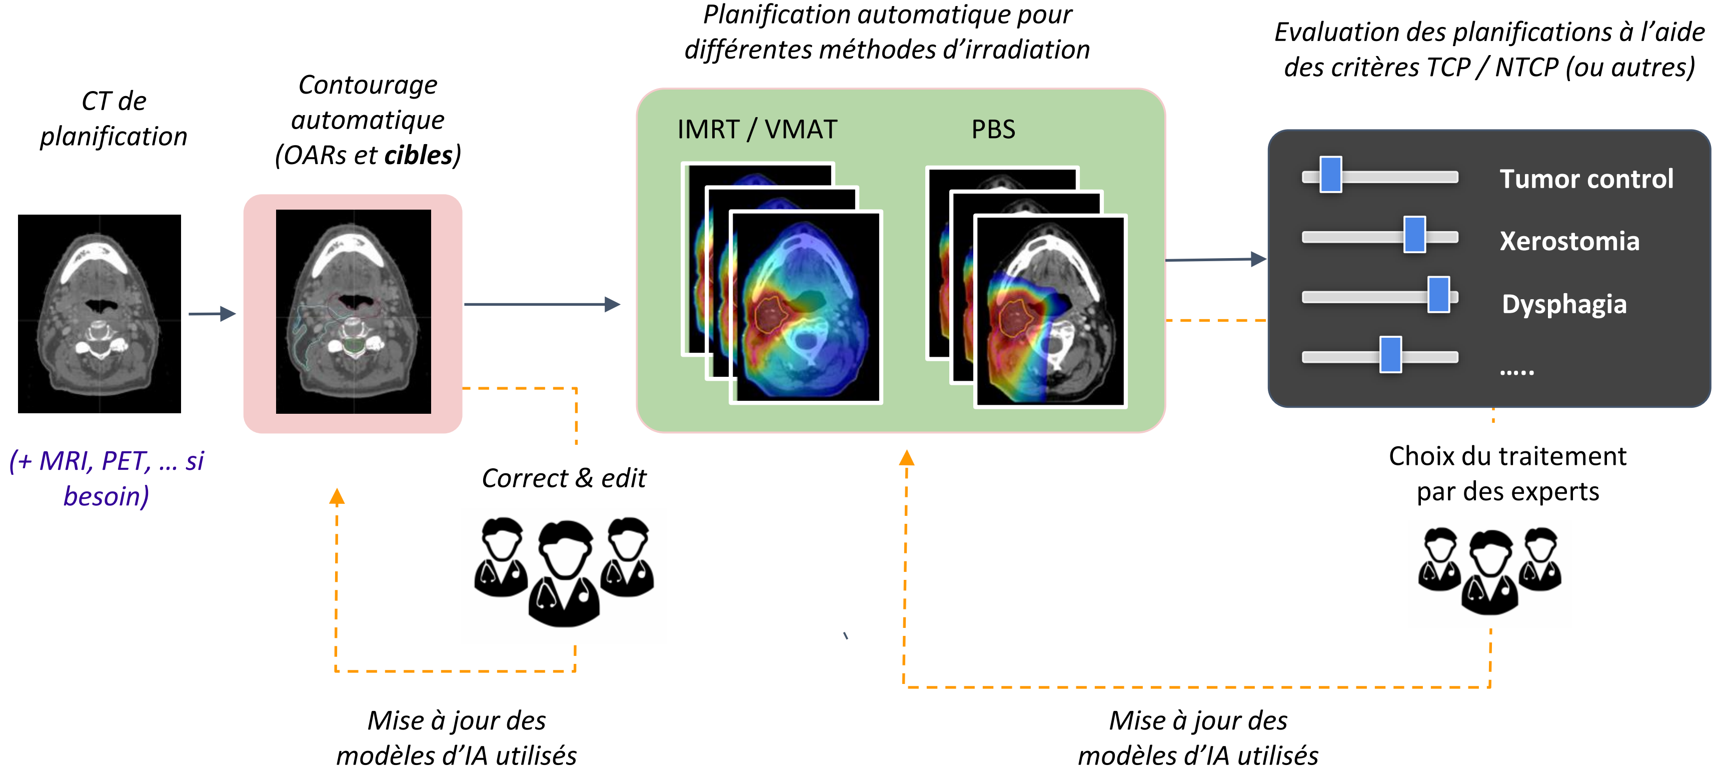

1) Segmentation des OAR et des tumeurs: Apports continu d'annotation de différents hôpitaux et d'experts de différents niveaux pour alimenter un modèle Deep Learning (U-Net) permettant la segmentation automatique.

2) Le modèle représente une « meilleure pratique » d'une coalition (un groupe d'hôpitaux)

3) Représentation des images « optimisées » pour le machine learning et « fluides » pour l'annotation @home

Table of content

1. Coalitional machine learning: an example
2. **Image coding for machine learning**
3. (Machine learning for image coding)
4. Distributed (federated) learning
5. Federated byzantine agreements for trusted coalitions
6. Coalitional active learning
7. Image coding for coalitional active learning
   - Medical
   - Satellite imaging

2. Image coding for machine learning

- Classification of JPEG-2000 CT compressed images (Steinhofel, Dewey, Janssens, Macq, 2002)

	Error on pos.		Error on neg.		Error in%	
Org.	0		3		1.5	
Rate	cmp.	org.	cmp.	org.	cmp.	org.
50%	4	4	2	2	3.0	3.0
25%	2	2	2	2	2.0	2.0
10%	0	1	1	1	0.5	1.0
5%	2	3	1	3	1.5	3.0
2%	4	3	1	1	2.5	2.0

The best classification (above original images) is achieved for CT images compressed with a factor 10 (denoising effect)

- Improved 3D U-Net robustness against JPEG 2000 compression for male pelvic organ segmentation in radiotherapy (El Khoury, Fockedey, Brion, Macq, 2021)

The experiment

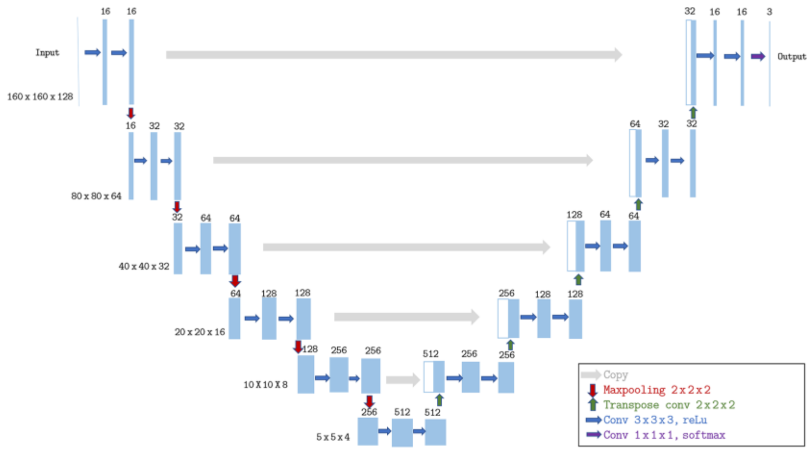

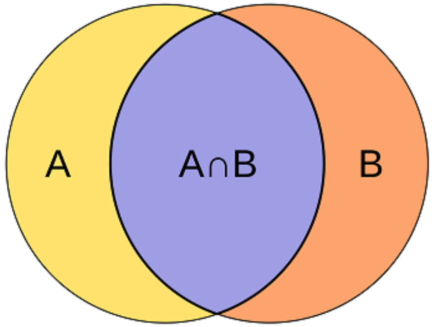

$$\text{Dice coefficient}(A, B) = \frac{2 \times |A \cap B|}{|A| + |B|}$$

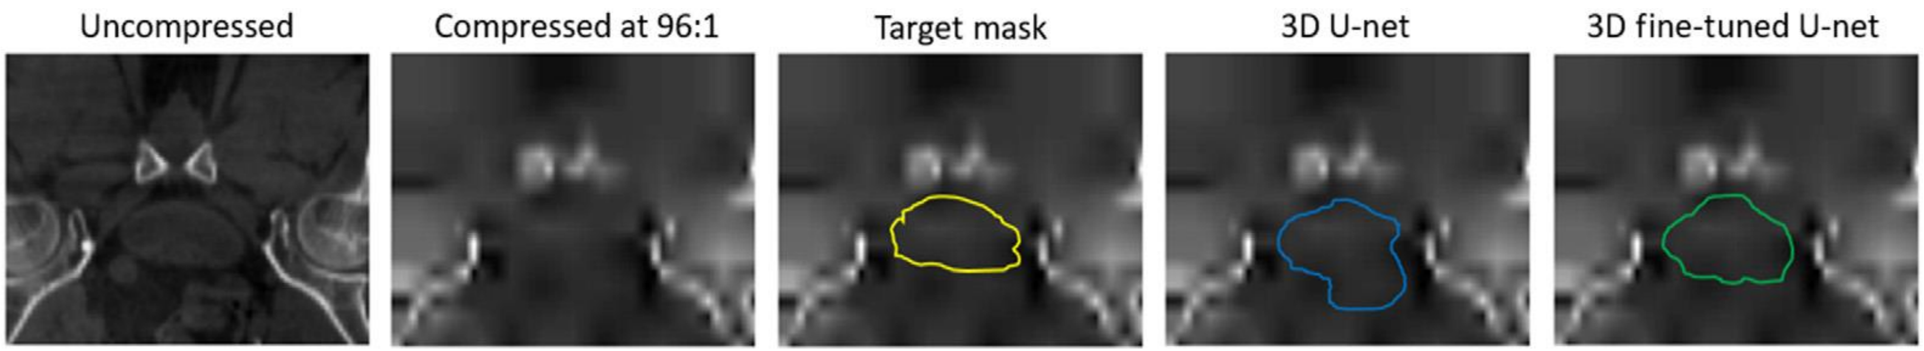

Global or alternate optimisation problem ?

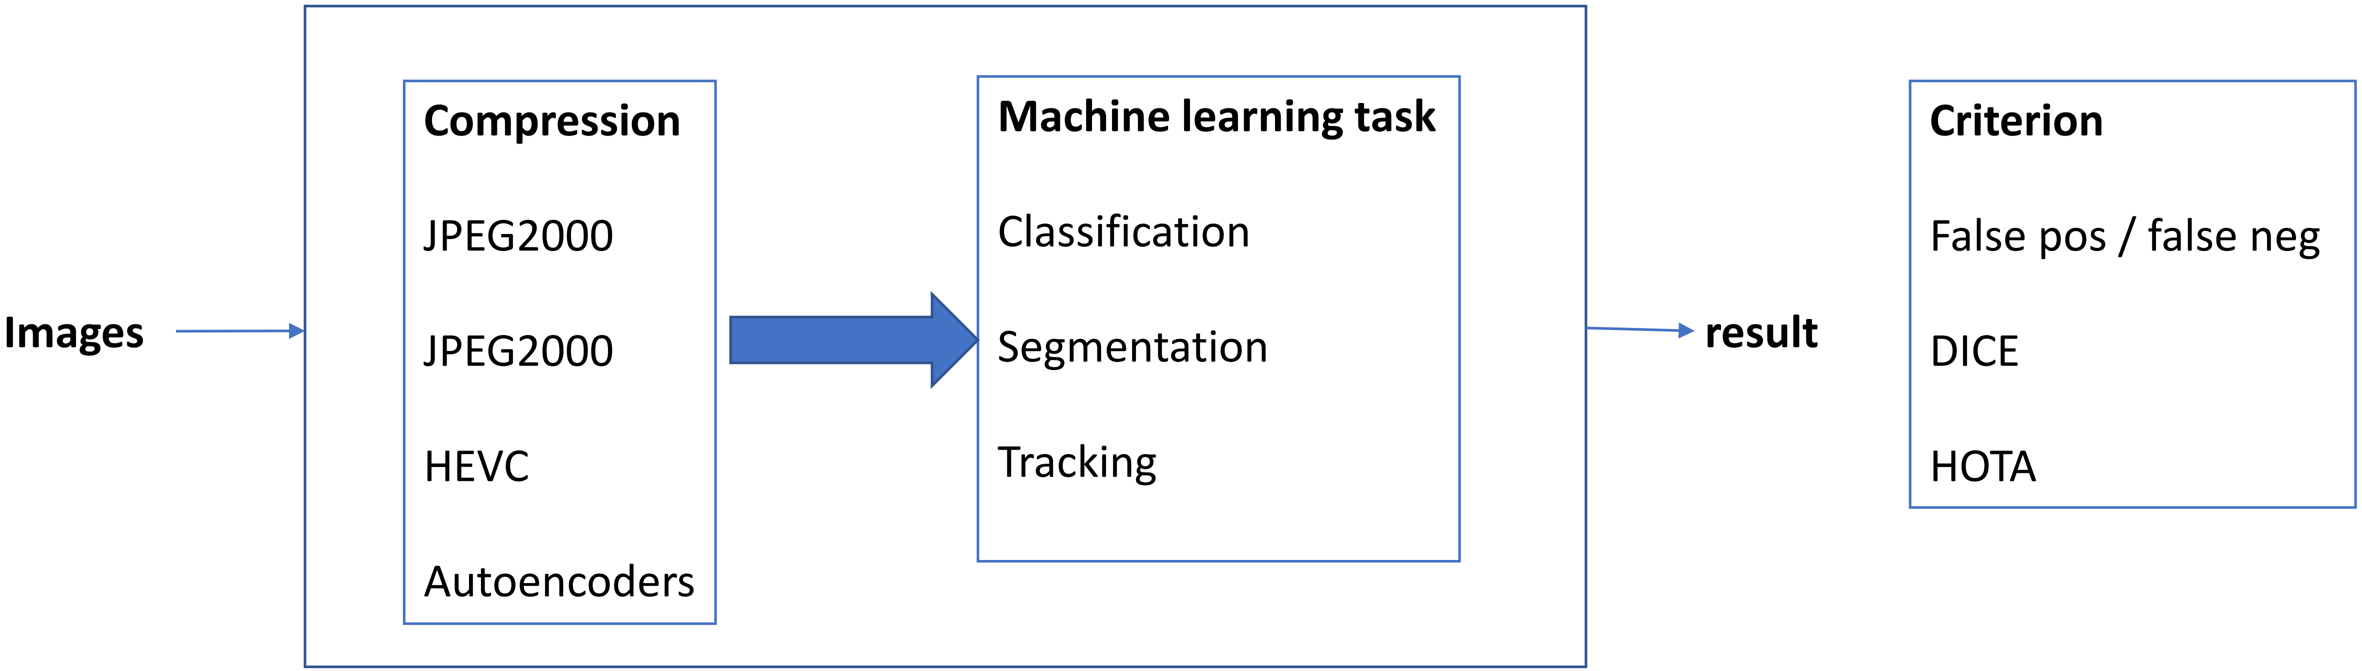

Heuristic search (including genetic algos)
Denoising, …

Learning on compressed data

Some results

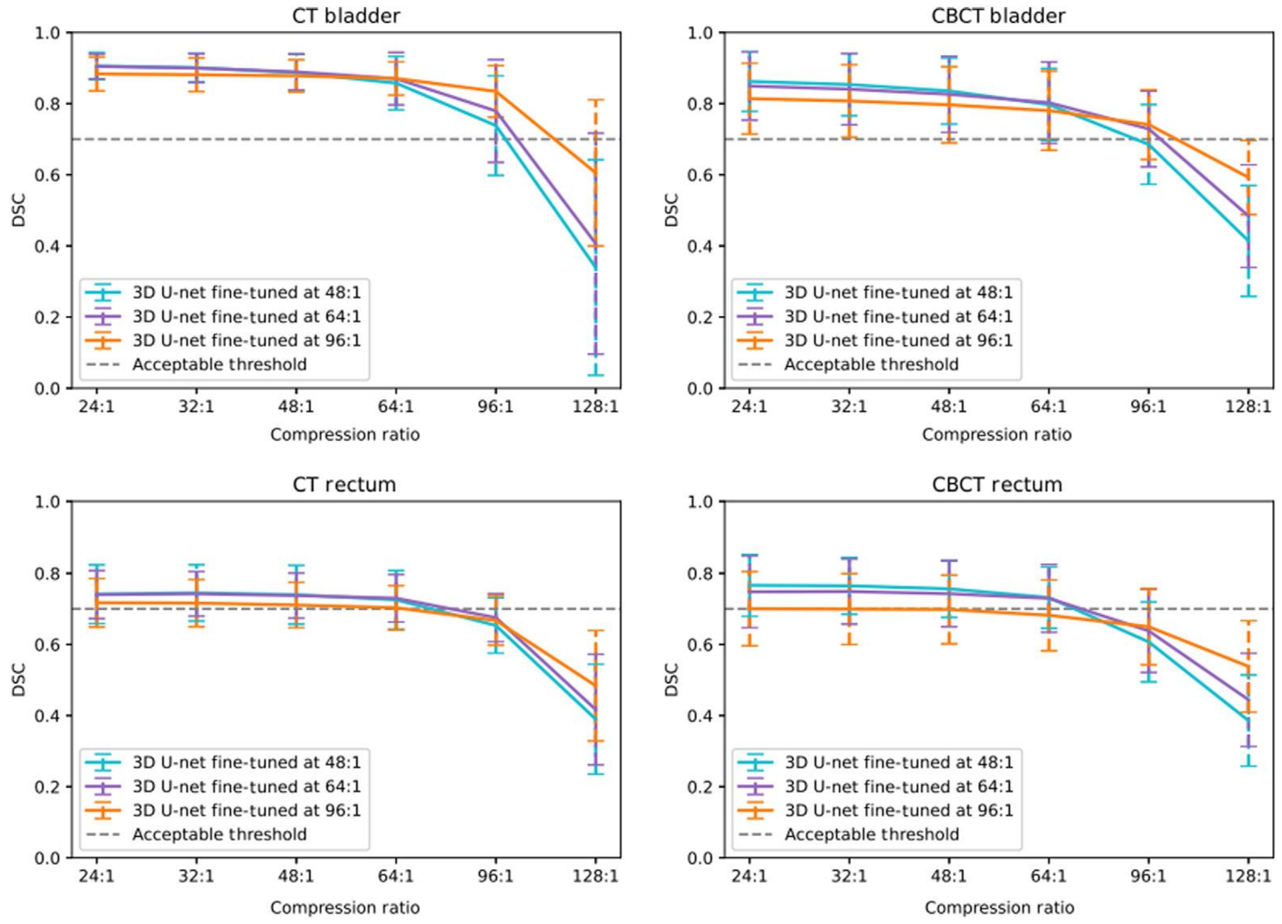

JPEG vs JPEG2000 on fine-tuned network

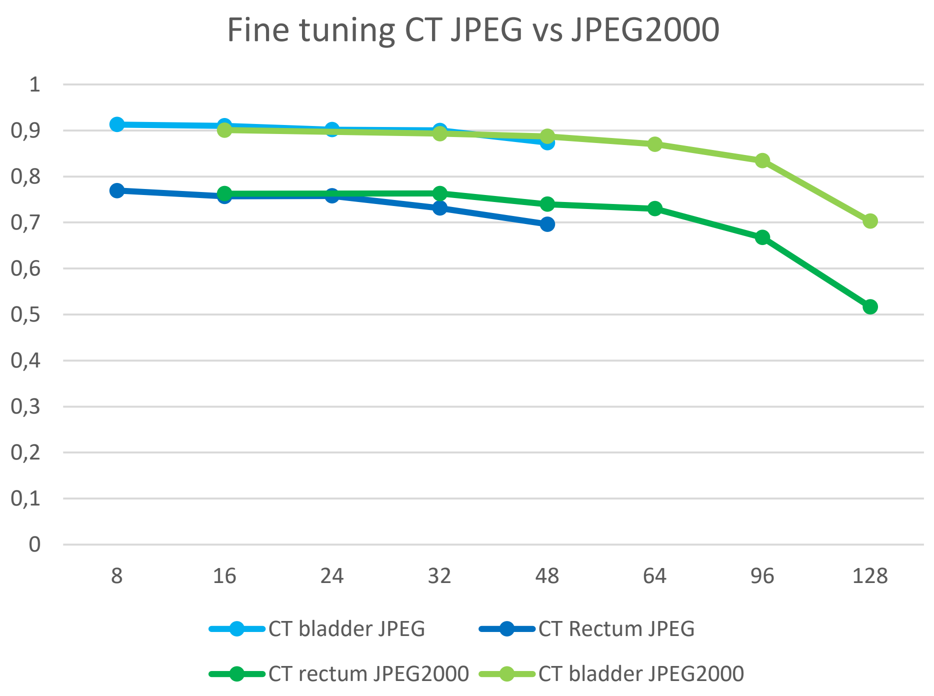

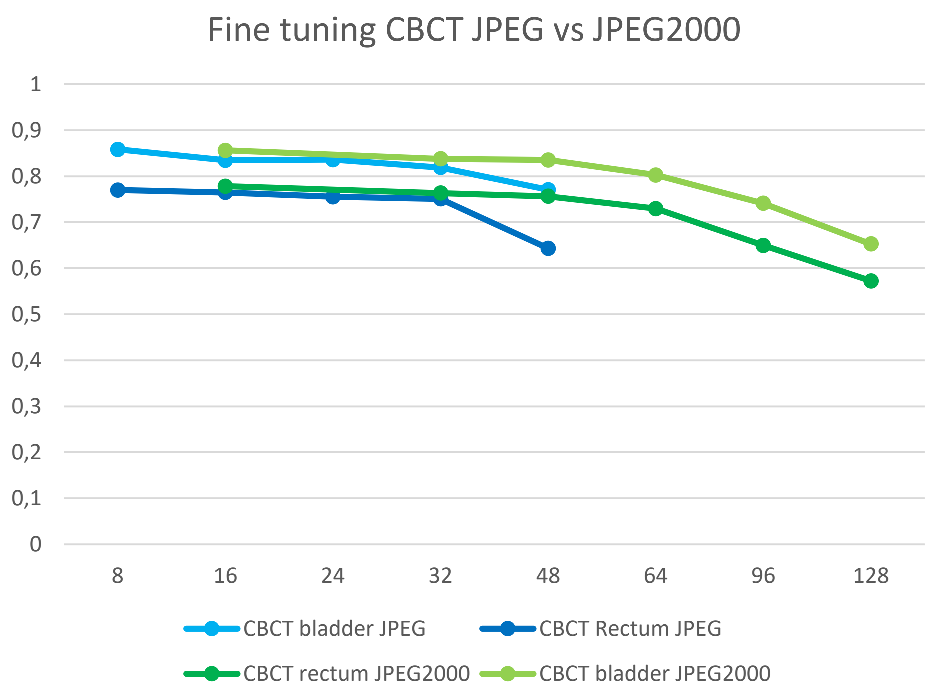

The Video Coding for Machine (MPEG)

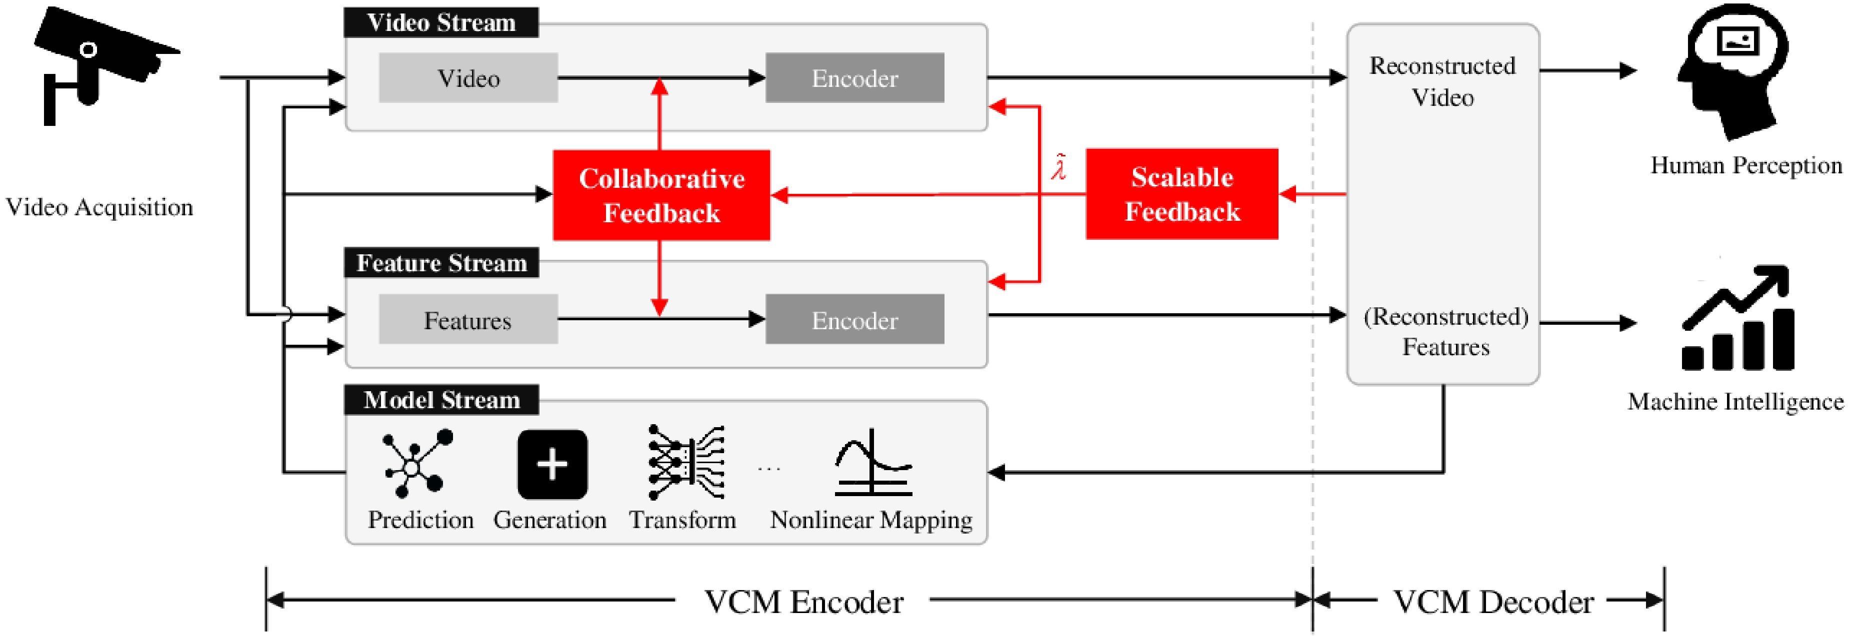

From Video Coding for Machines: A Paradigm of Collaborative Compression and Intelligent Analytics
Ling-yu Duan, Jiaying Liu, Wenhan Yang, Tiejun Huang, W. Gao

Privacy perserving compression for machine learning

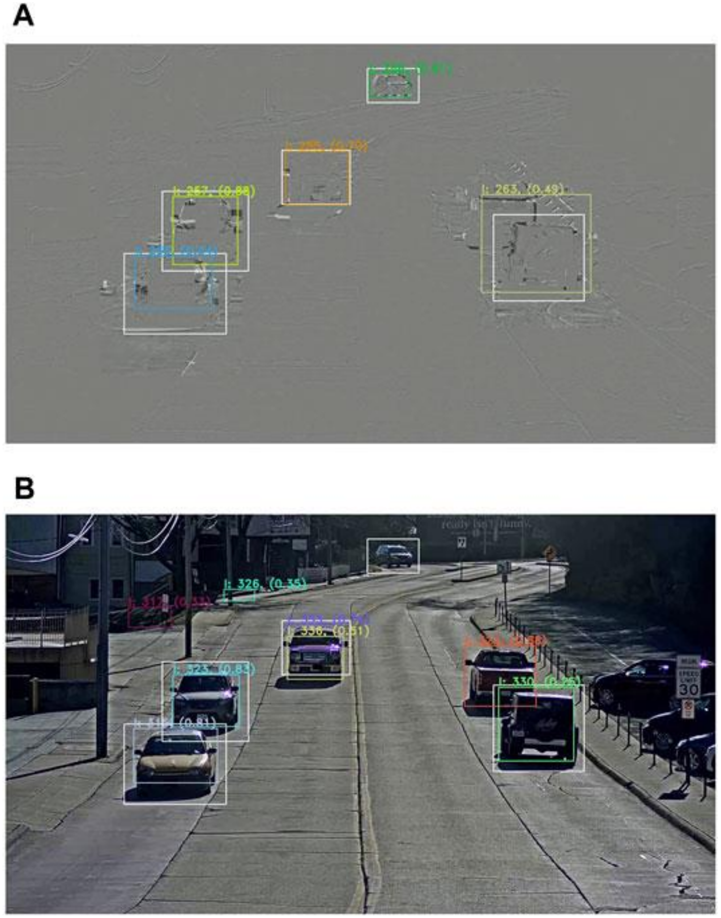

Deep Learning-Based Object Tracking via Compressed Domain Residual Frames
www.frontiersin.org Karim El Khoury*†,Jonathan Samelson† and Benoît Macq

Table of content

1. Coalitional machine learning: an example
2. Image coding for machine learning
3. **(Machine learning for image coding)**
4. Distributed (federated) learning
5. Federated byzantine agreements for trusted coalitions
6. Coalitional active learning
7. Image coding for coalitional active learning
    - Medical
    - Satellite imaging

3. Machine learning for compression

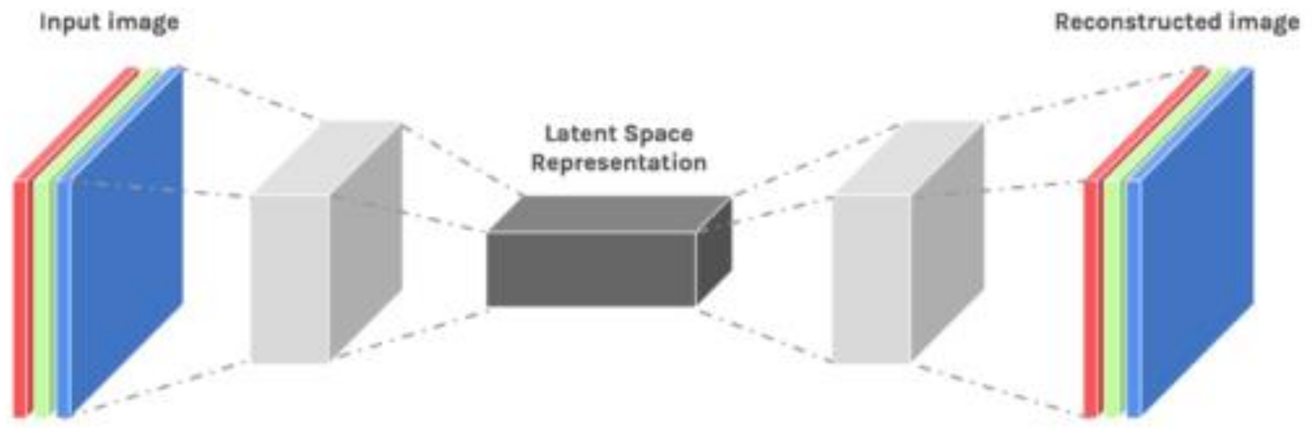

End-to-end optimized image compression with competition of prior distributions (Brummer, De Vleeschouwer CVPR 2021)

LEARNING A SPARSE GENERATIVE NON-PARAMETRIC SUPERVISED AUTOENCODER (Barlaud, Guyard, ICASSP 2021)

Interesting prospects in combined compression denoising … but what about complexity and flexibility ?

Table of content

1. Coalitional machine learning: an example
2. Image coding for machine learning
3. (Machine learning for image coding)
4. **Distributed (federated) learning**
5. Federated byzantine agreements for trusted coalitions
6. Coalitional active learning
7. Image coding for coalitional active learning
    - Medical
    - Satellite imaging

4. Distributed (federated) learning

1) A distributed learning (learning per batch) can reach an equivalent accuracy than a centralised learning

2) Federated learning leave the data inside Hospitals

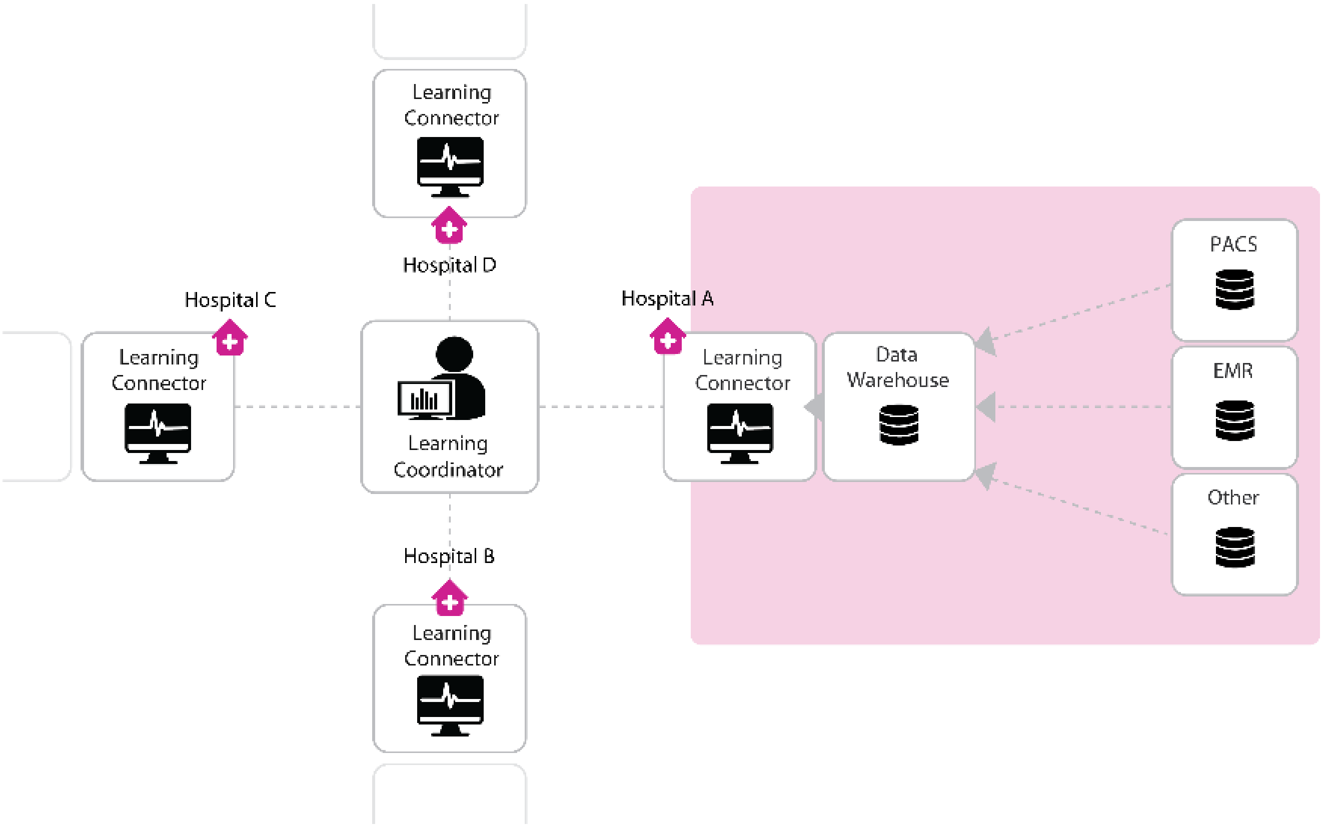

Distributed learning: an abundant litterature

- Distributed SVM: convergence equivalent to central learning can be proven
  - **Boyd**, Stephen, et al. "Distributed optimization and statistical learning via the alternating direction method of multipliers." Foundations and Trends® in Machine learning 3.1 (2010): 1-122
  - Forero, P. A., Cano, A., & Giannakis, G. B. (2010). Consensus-based distributed support vector machines. Journal of Machine Learning Research, 11(May), 1663-1707.
- Distributed DNN – Federated learning convergence similar to central learning can be shown
  - McMahan, B., & Ramage, D. (2017). Federated learning: Collaborative machine learning without centralized training data. Google Research Blog, 3.

Table of content

Security requirements

Challenge 1

Data privacy of the datasets used for the training (leakage effect of the gradients) : working by batches- differential privacy is the “crypto” model

Challenge 2

Protection of the model against degradation by training on inadequate data: steps validation by the coalition and a Federated Byzantine Agreement of the model

Challenge 3

Confidentiality of the model and the gradients: homomorphic operations and/or access control of the model vault

Challenge 4

Traceability of the model: DNN watermarking

Les généraux byzantins

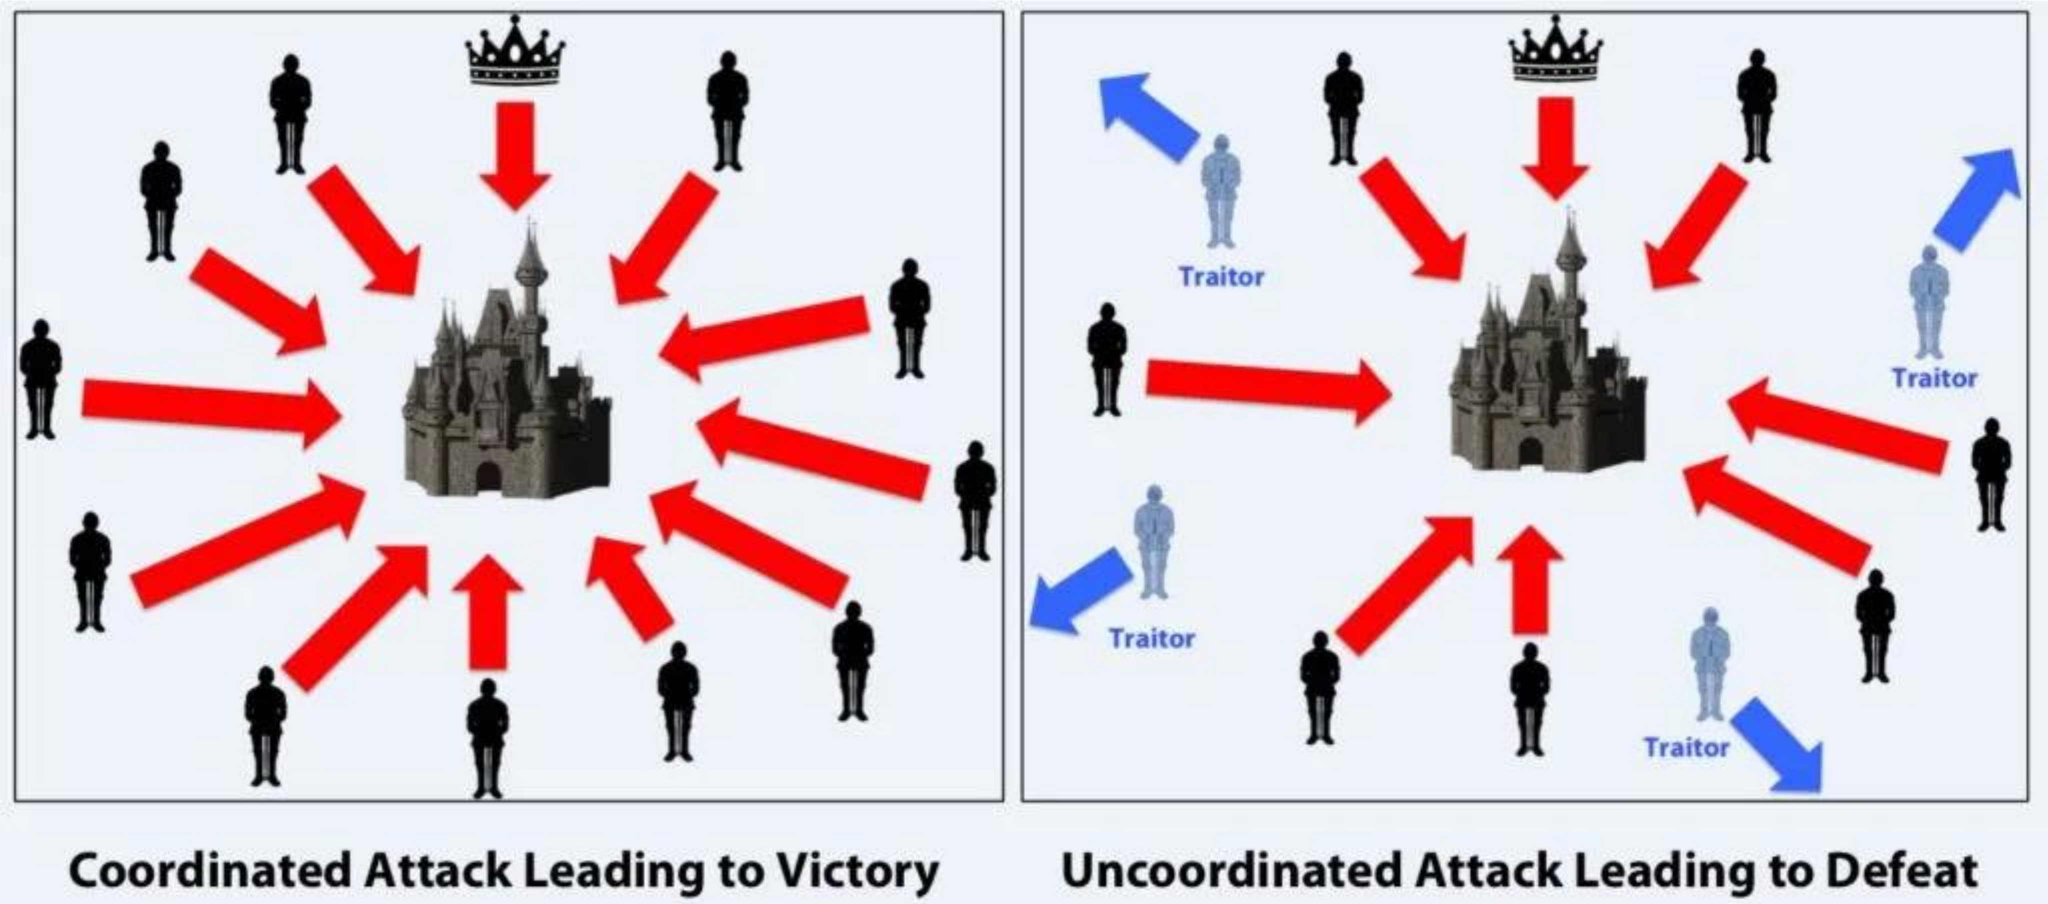

Federated Byzantine Agreement

- Two types of test databases:
  global test database (G),
  local test database (L)
- A “general” is randomly selected among the validators
- The “general” creates a new candidate block referencing the new model
- Every validator validates the viability (model) and integrity of this new candidate block
- Each validator broadcasts its opinion (positive or negative)
- The FBA process ends when 2/3 of the validators agree

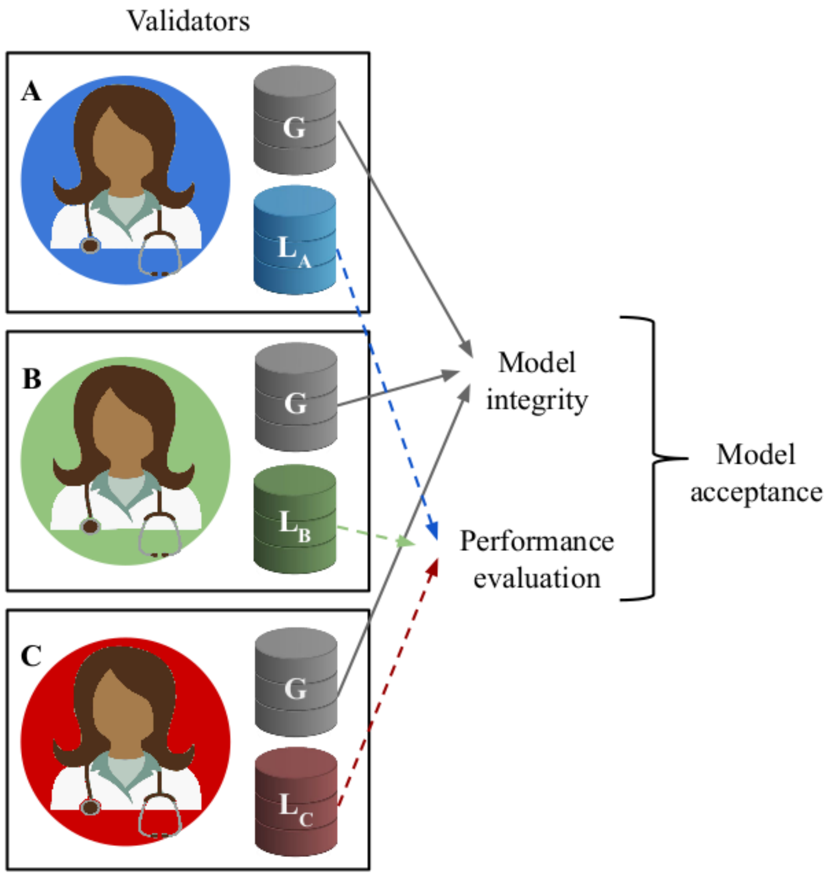

Lugan, S., Desbordes, P., Brion, E., Tormo, L. X. R., Legay, A., & Macq, B. (2019). Secure architectures implementing trusted coalitions for blockchained distributed learning (TCLearn). IEEE Access, 7, 181789-181799.

Scalable security architectures for trusted coalitions

TCLearn-A

Learned model is *public*
Each member is accountable for the privacy protection of its own data

Solution to security challenge1
(Data privacy of the datasets used for the training):

Local training of the model by each member with their own datasets

Generated gradients are uploaded and merged with the previous model

Batches of a minimum size to mitigate the long term memory effect

Solution to security challenge 2
(Protection of the model against degradation by training on inadequate data):

Blockchain storing cryptographic hashes of every training step

Federated Byzantine Agreement (FBA) to prevent corrupted increments

Scalable security architectures for trusted coalitions

TCLearn-B

Learned model is *private*, the members of the coalition trust each other.

Solution to security challenges 1 & 2:

Same as for TCLearn-A

Solution to security challenge 3:
(Confidentiality of the model and the gradients):

Storage of all iterations of the model in an off-chain storage

Iterations only referenced by links in the blockchain

Secure, encrypted transport of the model (using e.g. TLS or S/MIME)

Solution to security challenge 4:
(Traceability of the model):

Access control and audit mechanisms to protect the models and parameters

Scalable security architectures for trusted coalitions

TCLearn-C

The members of the coalition do no trust each other.

Solution to security challenges 1 & 2:

Same as for TCLearn-A

Solution to security challenges 3 & 4:

Storage of all iterations of the model in an off-chain storage

Each member is provided with a homomorphically encrypted model and the corresponding public key, used to encrypt their datasets, by a supervisor

Prediction could be performed locally on encrypted data, but the result must be decrypted by the supervisor

Full traceability since the encrypted model cannot be used without the associated public key, itself associated with the partner which received it

Table of content

1. Coalitional machine learning: an example
2. Image coding for machine learning
3. (Machine learning for image coding)
4. Distributed (federated) learning
5. Federated byzantine agreements for trusted coalitions
6. **Coalitional active learning**
7. Image coding for coalitional active learning
    - Medical
    - Satellite imaging

An old question

Will AI ever replace the radiologist (practitioner MD) ?

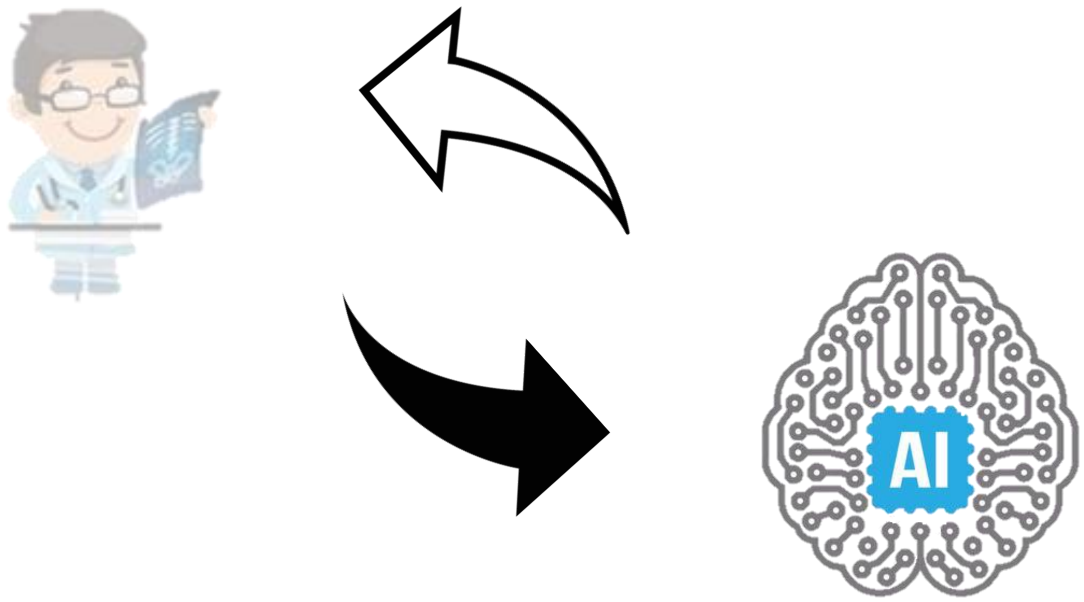

An answer (Curtis Langlotz, RSNA)

The radiologist who uses AI will replace the radiologist who does not?

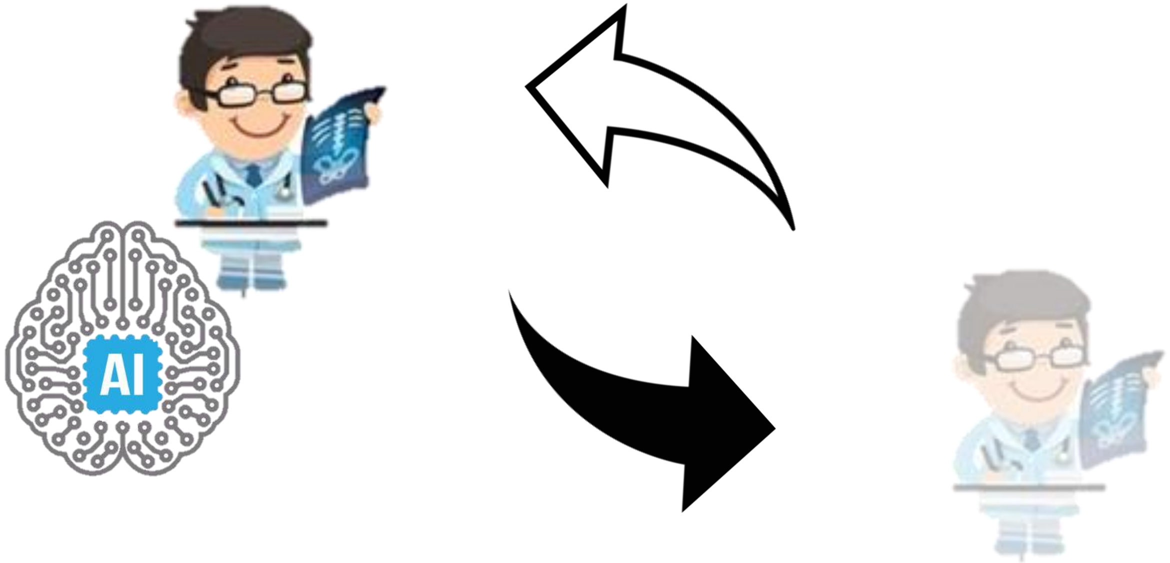

Deep Learning

Example: classification and detection of tumors in mammograms

Experiment settings

- ResNet18

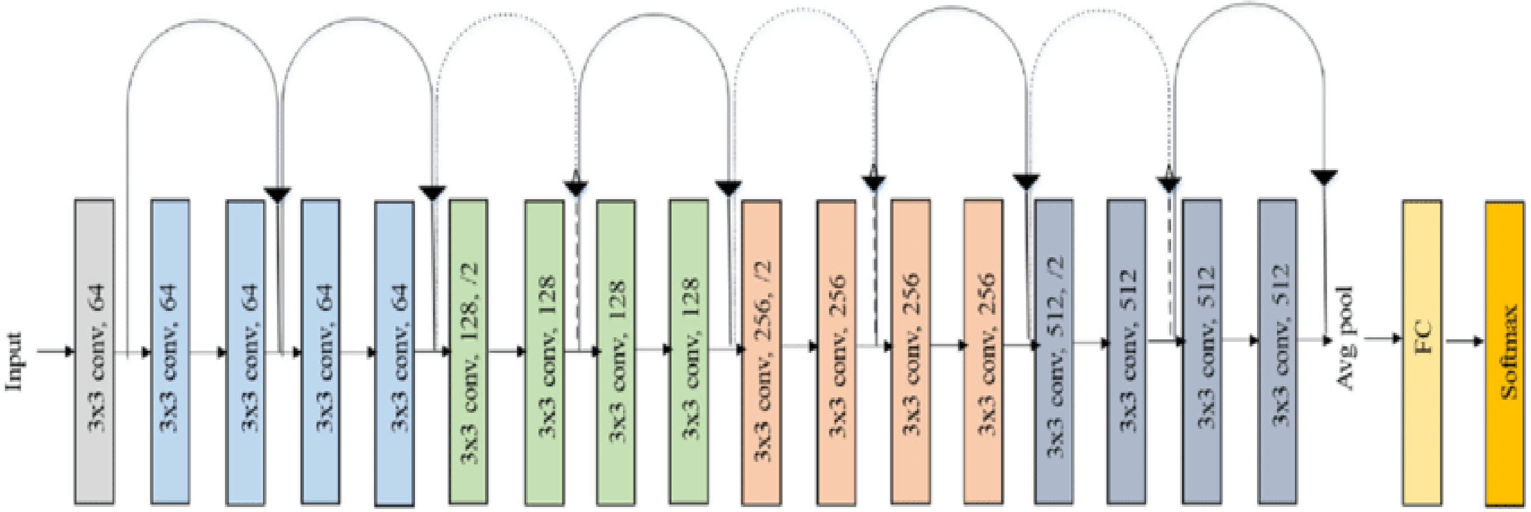

- CBIS-DDSM dataset

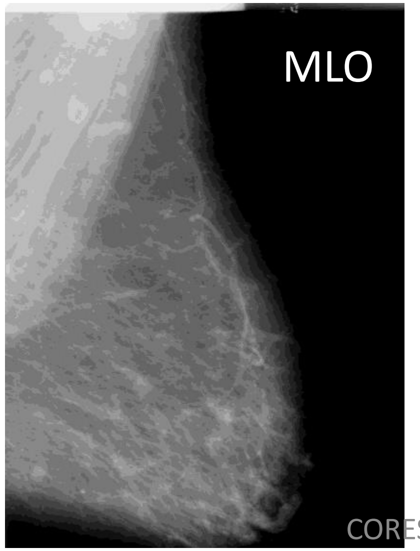

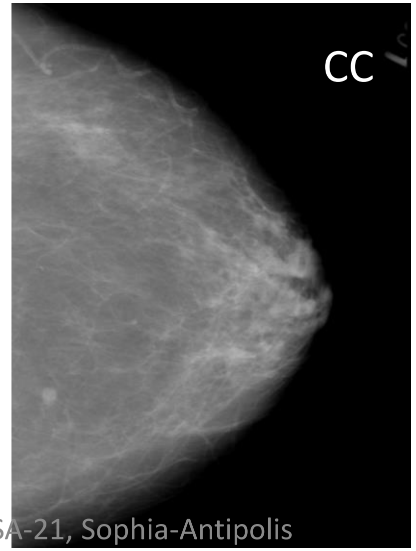

1

Whole-image classifier

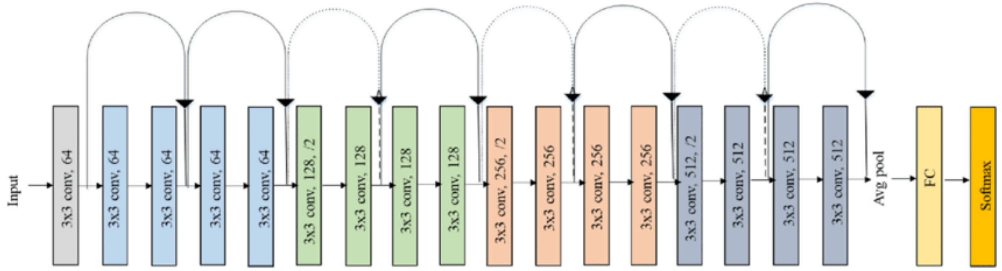

{ Benign
Malignant }

Accuracy = 0.652

2

Patch approach

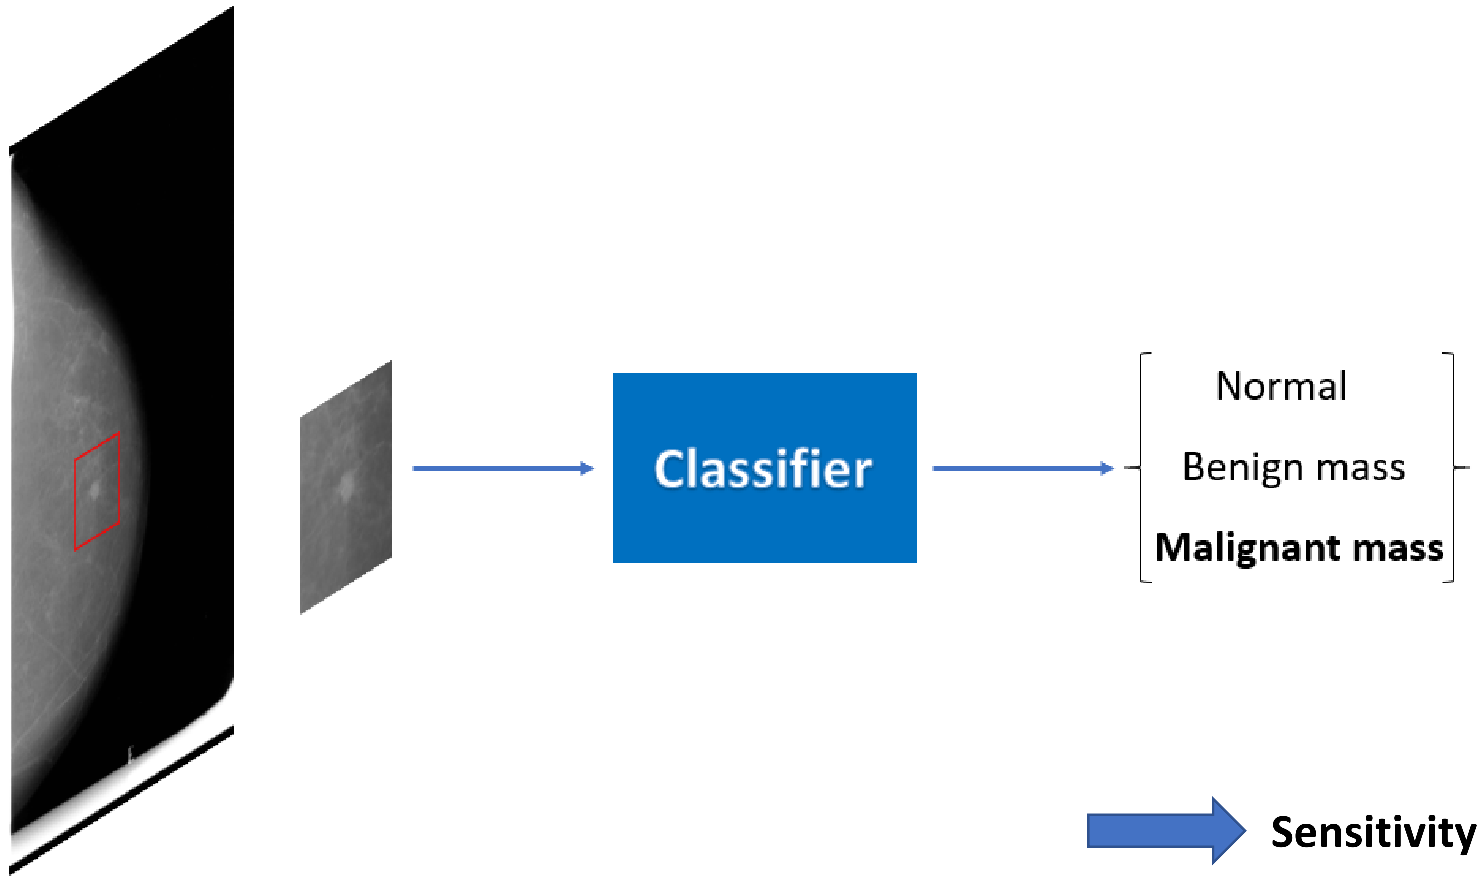

Sensitivity 0.9 at 0.1 specificity

Faster R-CNN: Combination of both approaches

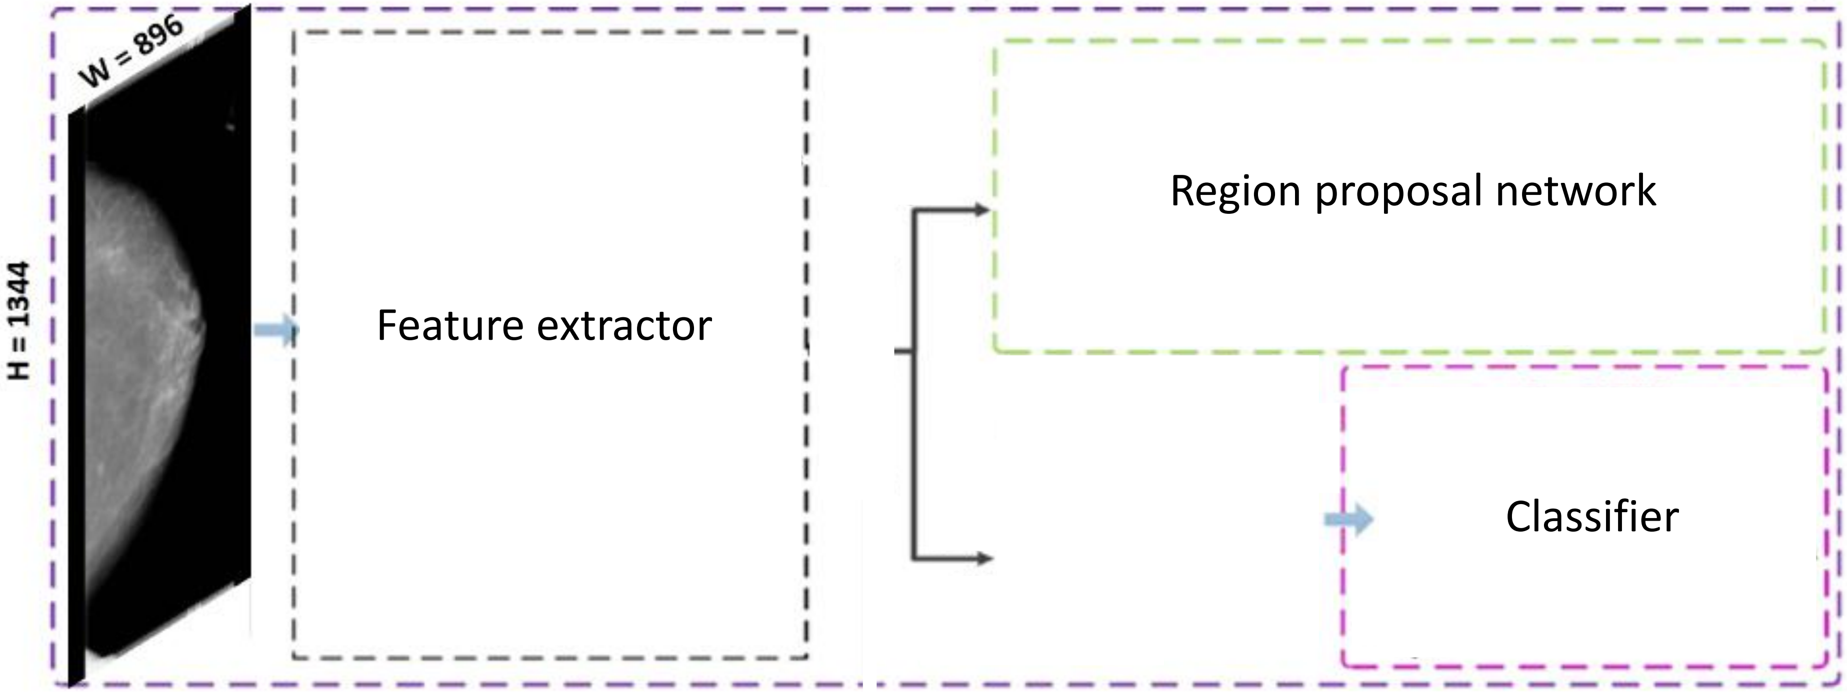

Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster R-CNN: towards real-time object detection with region proposal networks. In Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1 (NIPS'15). MIT Press, Cambridge, MA, USA, 91–99.

Coalition of experts do better than a single AI

Use of artificial intelligence for image analysis in breast cancer screening programmes: systematic review of test accuracy; BMJ 2021;

"Current evidence for AI does not yet allow judgement of its accuracy in breast cancer screening programmes, and it is unclear where on the clinical pathway AI might be of most benefit. AI systems are not sufficiently specific to replace radiologist double reading in screening programmes. Promising results in smaller studies are not replicated in larger studies. Prospective studies are required to measure the effect of AI in clinical practice. Such studies will require clear stopping rules to ensure that AI does not reduce programme specificity."

Continual learning

- The model should adapt continuously to the evolution of acquisition techniques
- The model should adapt continuously to
  - Guidelines
  - « Mindlines »
- Can the model (the AI) be a knowledge representation of a coalition of experts ?
- Can the model be a support for consensus in multidisciplinary teams?

Active Learning

Smart annotation queries

A limited budget

Annotations

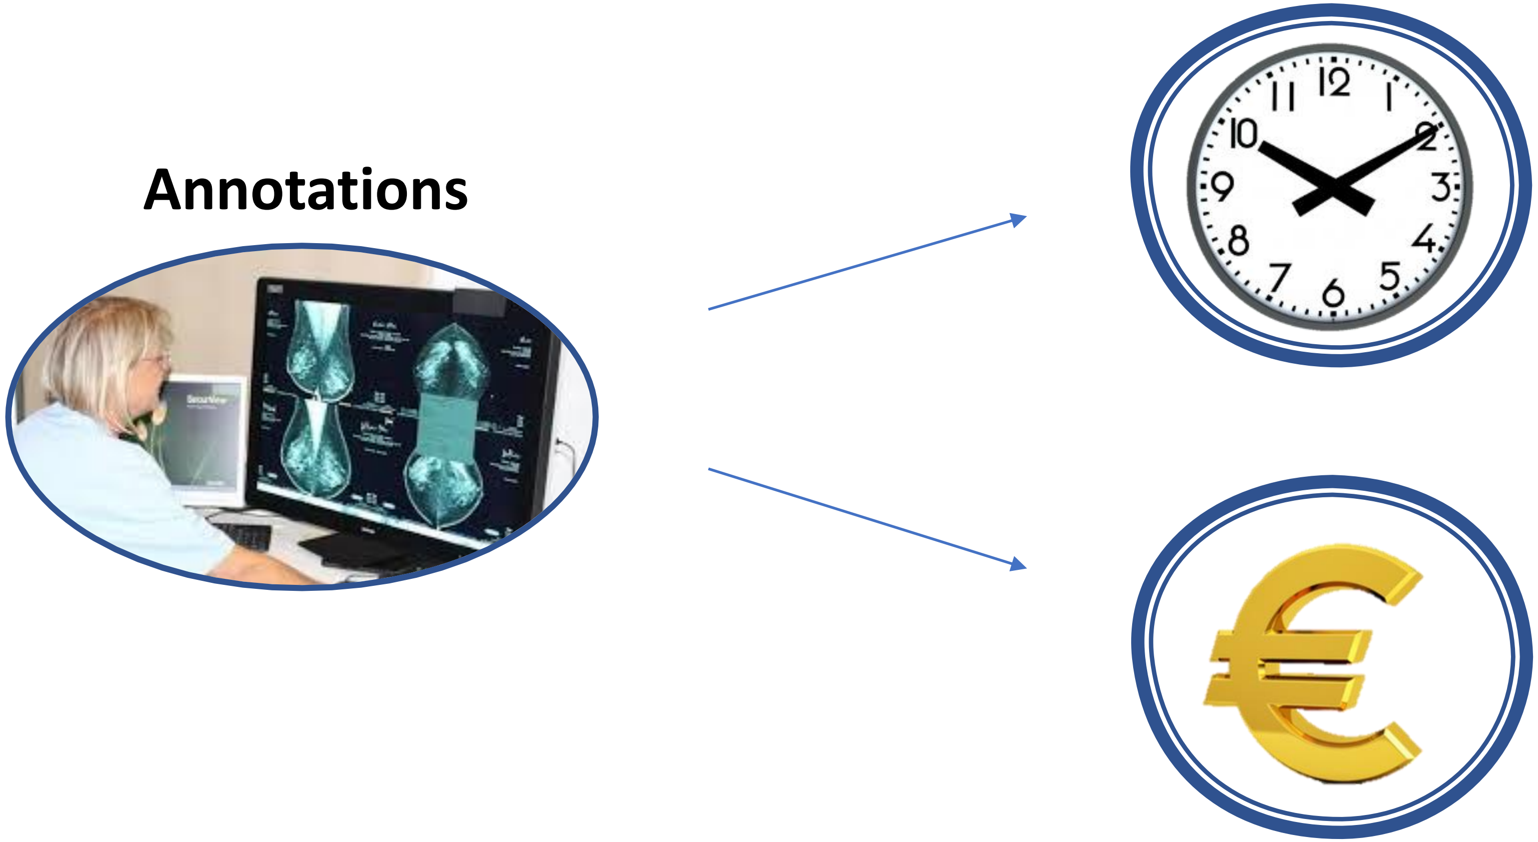

Intelligent selection

Hardest examples

Most representative examples

Examples that engenders disagreement

...

An iterative process

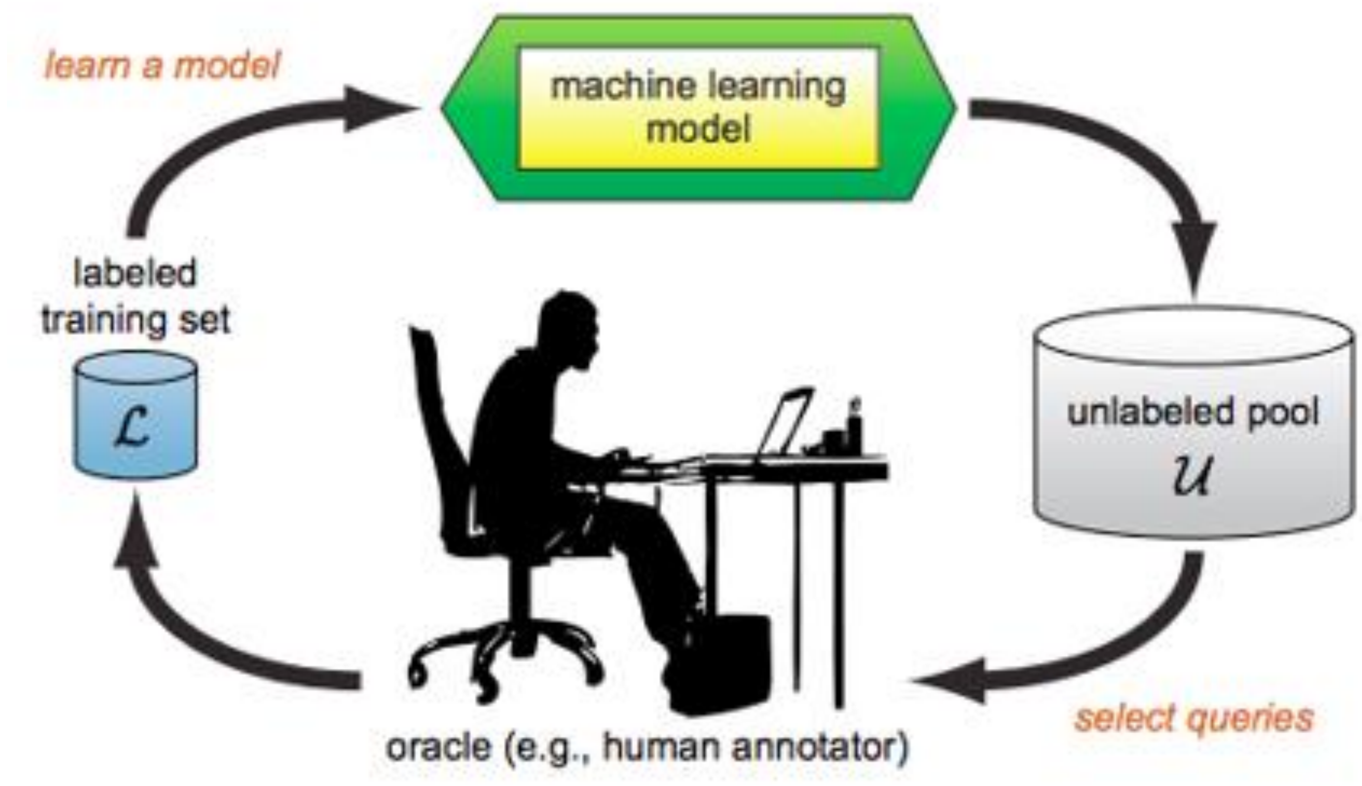

Three methods

- **Uncertainty (shake the model)**
- **Diversity (measure distances)**
- **Query-by-committee (agreement between competing models)**

Active Learning in Mammography

Reminder: patch approach

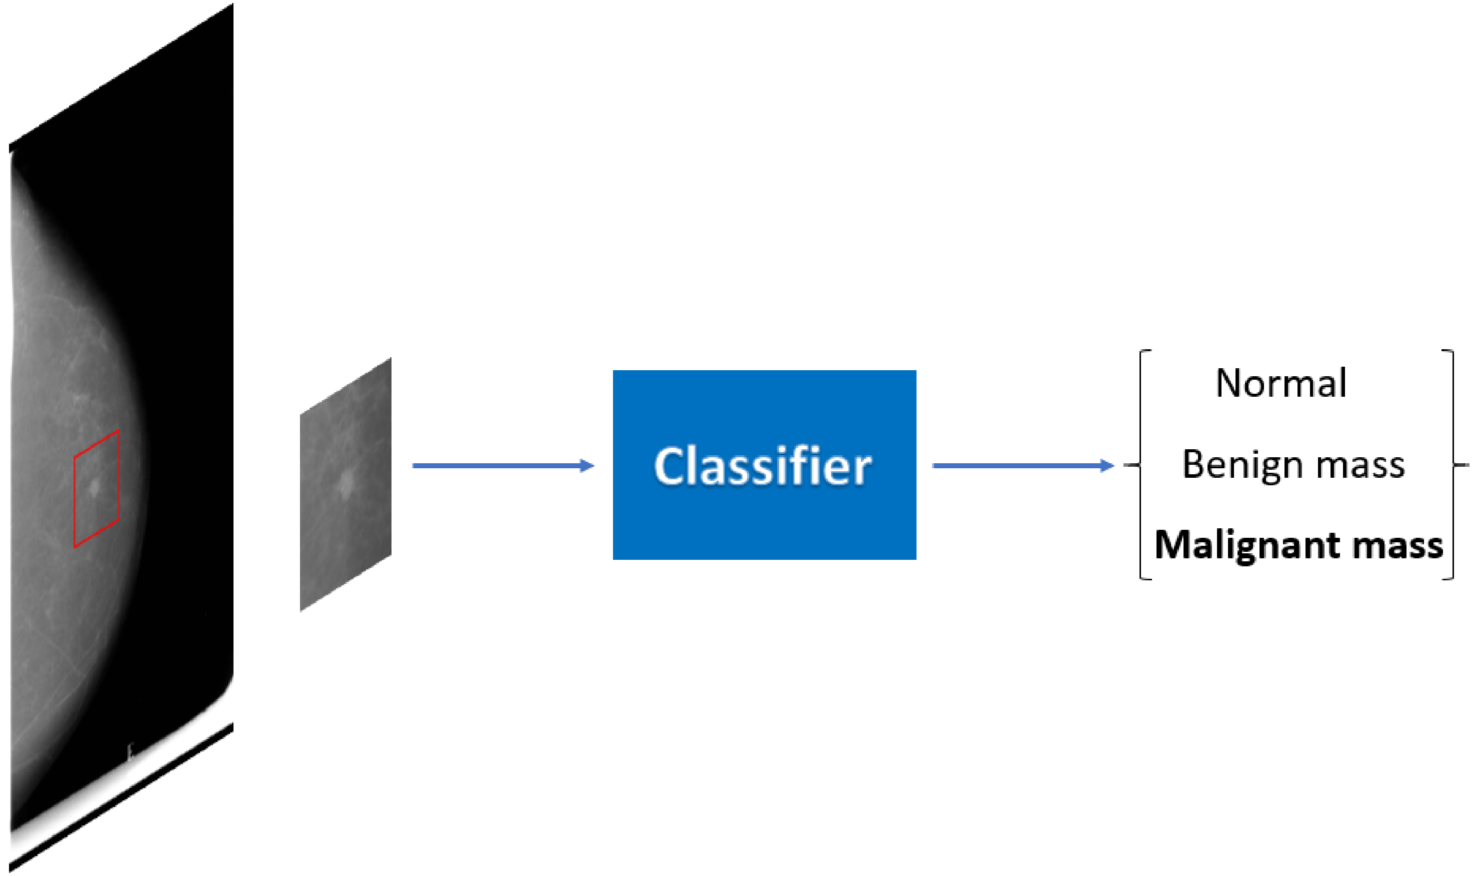

Active Learning in Mammography

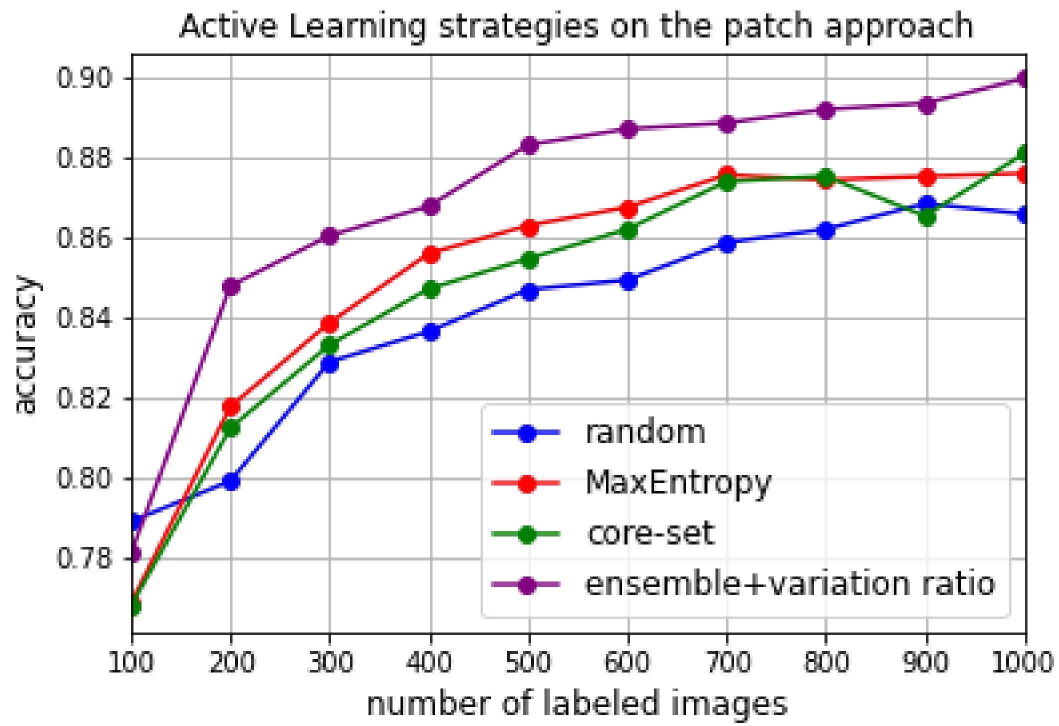

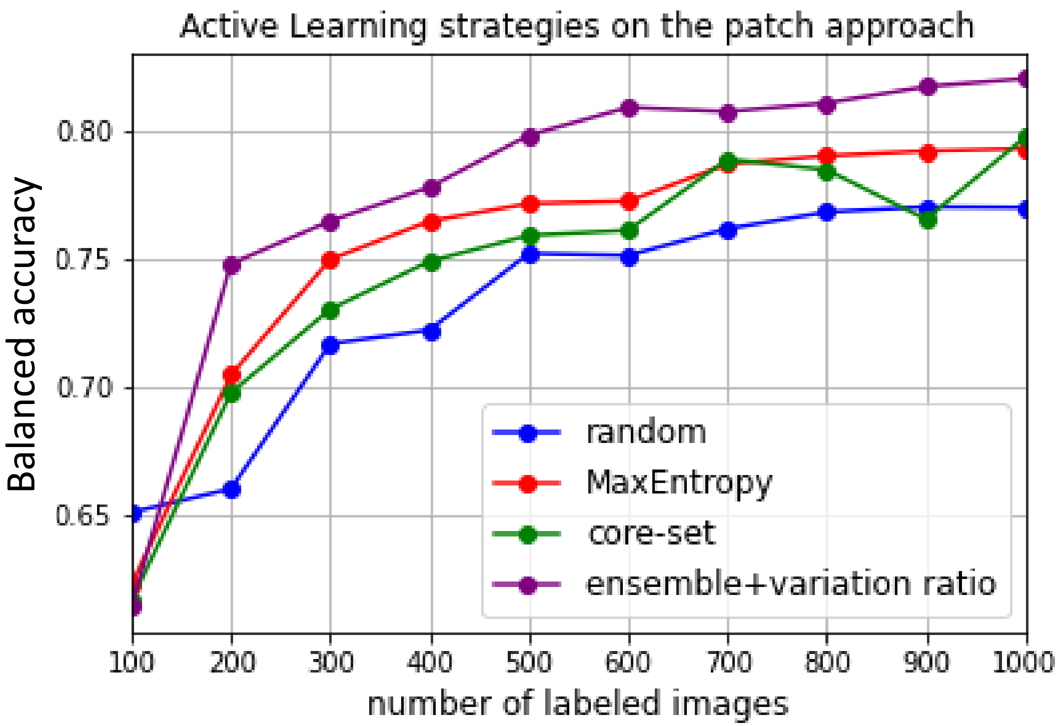

Coalitional Active learning

- -> Trusted and equitable distributed learning:
  - Images do not go out hospitals (privacy, …)
  - Coalitions between hospitals
  - Issues
    - Various quality of data sets (weakly supervised learning)
    - Distributed learning sequence (blockchained)
    - Sharing a Deep L model (watermarking)
- TCLearn (FBA –based learning)

The old question

Will the radiologist who uses AI replace the radiologist who does not?

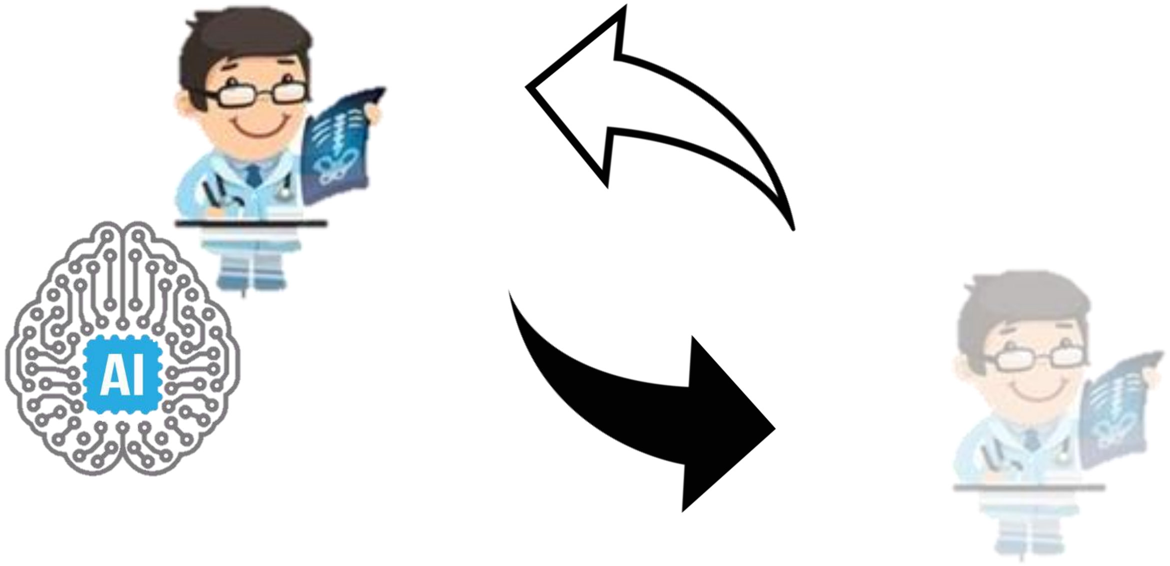

Our new question

AI that uses the coalition of radiologists will replace AI that does not

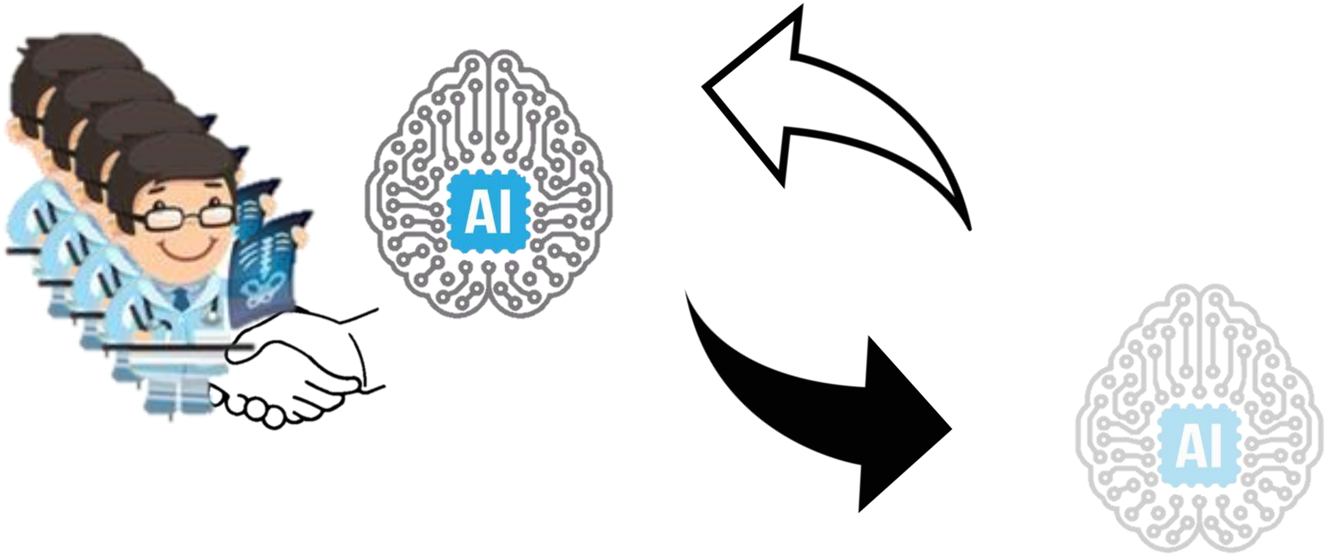

Table of content

1. Coalitional machine learning: an example
2. Image coding for machine learning
3. (Machine learning for image coding)
4. Distributed (federated) learning
5. Federated byzantine agreements for trusted coalitions
6. Coalitional active learning
7. **Image coding for coalitional active learning**
    - Medical
    - Satellite imaging

Image coding for coalitional learning

- In medical
  - Easy accessible on portable annotation devices
  - Easy quality control
  - Optimized for a task (classification, segmentation, …)
- In satellite imaging
  - Mega image, with local access
  - Multiresolution access
  - Optimized resolution
- High performance (VAE, BPG, ..) or (/and ?) flexibility of « old » wavelet-based image compression systems ?

Remerciements

- Christophe De Vleeschouwer
- Karim El Khoury
- Maxime Zanella
- Et les autres de l'équipe

www.pilab.be

Thank you for your attention!

Do not hesitate to ask any questions