RationAI: Rational and conservative AI (not only) in digital pathology

Presented by Tomáš Brázdil

RationAI research group

The aim: Develop explainable AI systems useful in practice.

Ideally, develop production ready solutions based on current research.

Current projects:

- Tumor detection in whole-slide images from digital pathology (this talk and more work on slide annotation etc.)
- Time series data from baryatry (preliminary data analysis)
- Spatio-temporal COVID-19 data analysis (just started)

The problem of cancer detection in WSI

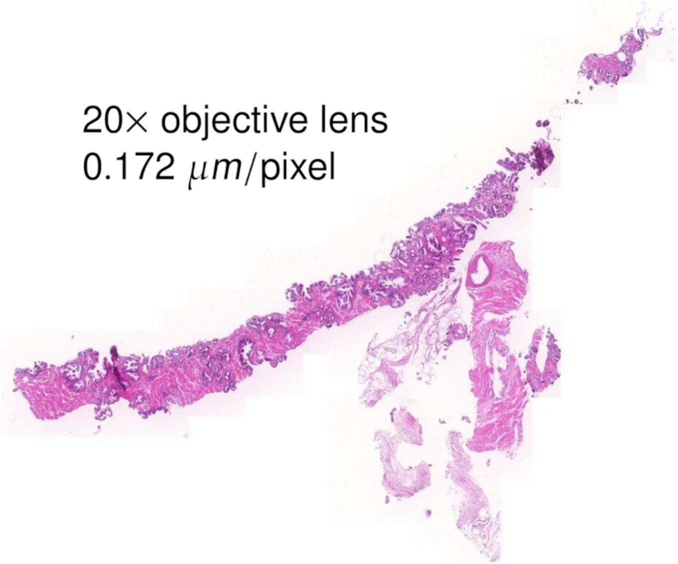

The problem: Detect cancer in this image.

The problem of cancer detection in WSI – solution

- WSI annotated by pathologists, **not** pixel level precise!
- Train a deep learning model on the annotated WSI.

Input data

WSI too large, 105,185 px × 221,772 px

Cut into patches of size 512 px × 512 px

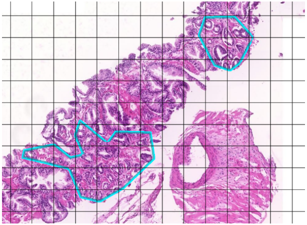

Input data

WSI too large, 105,185 px × 221,772 px

Cut into patches of size 512 px × 512 px

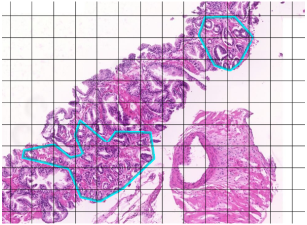

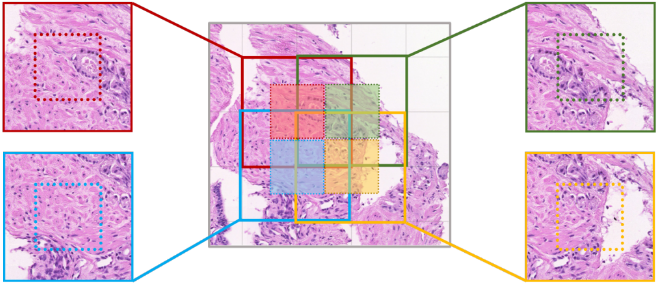

Patch positive **iff** the inner square intersects the annotation

Supervised learning classification of images

input I $\longrightarrow$ $\boxed{\text{model } F_\theta}$ $\longrightarrow$ output $F_\theta(I)$

- I is the *input image*

  A patch from WSI
- F_θ is a function on images depending on parameters θ.

  A neural network, θ contains its weights

Supervised learning classification of images

input I $\longrightarrow$ $\boxed{\text{model } F_\theta}$ $\longrightarrow$ output $F_\theta(I)$

- I is the *input image*

  A patch from WSI
- F_θ is a function on images depending on parameters θ.

  A neural network, θ contains its weights

Binary classification: Two classees: positive, negative;

$F_\theta(I) \in [0, 1]$ is the probability that I is positive

Supervised learning classification of images

$$\text{input } I \longrightarrow \boxed{\text{model } F_\theta} \longrightarrow \text{output } F_\theta(I)$$

- I is the *input image*
  A patch from WSI
- F_θ is a function on images depending on parameters θ.
  A neural network, θ contains its weights

Binary classification: Two classees: positive, negative;

$F_\theta(I) \in [0, 1]$ is the probability that I is positive

Training: Given a *dataset* $\mathcal{D}$ of pairs $(I_1, c_1), \ldots, (I_n, c_n)$

- I_k is an image
- $c_k = \begin{cases} 1 & I_k \text{ positive} \\ 0 & I_k \text{ negative} \end{cases}$

minimize a loss $\mathcal{L}(\theta; \mathcal{D})$ with respect to θ.

Training on WSI

Our dataset from Masaryk Memorial Cancer Insitute:

- 785 WSI from 166 patients
  (698 WSI for training, 87 WSI for testing)
- Cut into 7,878,675 patches for training, 193,235 patches for testing.

Training on WSI

Our dataset from Masaryk Memorial Cancer Insitute:

- 785 WSI from 166 patients
  (698 WSI for training, 87 WSI for testing)
- Cut into 7,878,675 patches for training, 193,235 patches for testing.

Dataset augmentation:

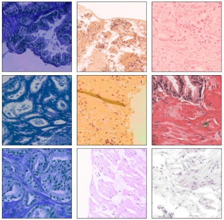

- random vertical and horizontal flips
- random color perturbations

Trained VGG16 network using RMSprop algorithm.

I.e. a standard solution.

Training on WSI

Our dataset from Masaryk Memorial Cancer Insitute:

- 785 WSI from 166 patients
  (698 WSI for training, 87 WSI for testing)
- Cut into 7,878,675 patches for training, 193,235 patches for testing.

Dataset augmentation:

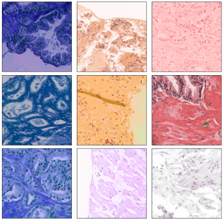

- random vertical and horizontal flips
- random color perturbations

Trained VGG16 network using RMSprop algorithm.

I.e. a standard solution.

The question: How good is the resulting model F_θ ?

Prediction

Model evaluation - attempt 1

Can we detect cancer somewhere in WSI?

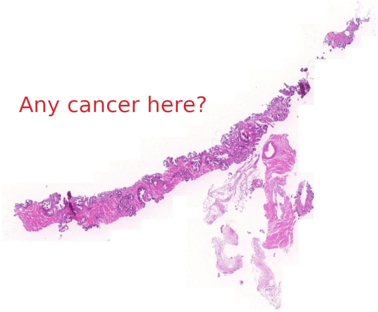

Model evaluation - attempt 1

Can we detect cancer somewhere in WSI?

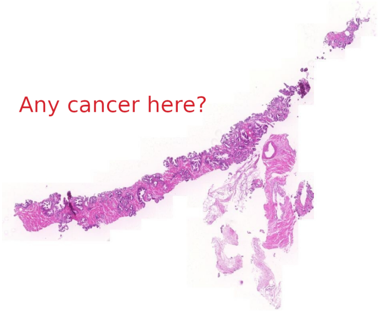

Predict WSI positive iff at least one patch I satisfies $F_\theta(I) \geq t$ for a fixed threshold $t \in [0, 1]$.

Choosing t close to 1, we have achieved 100% accuracy, i.e.,

slide positive iff predicted positive

Problem Solved!

Model evaluation - attempt 1

Can we detect cancer somewhere in WSI?

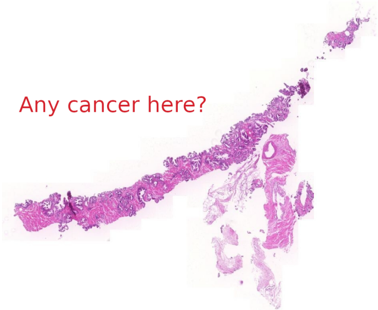

Predict WSI positive iff at least one patch I satisfies $F_\theta(I) \geq t$ for a fixed threshold $t \in [0, 1]$.

Choosing t close to 1, we have achieved 100% accuracy, i.e.,

slide positive iff predicted positive

Problem Solved! ... No?

Model evaluation - attempt 2

Can we detect cancer in patches?

Predict a patch I positive iff $F_\theta(I) \geq 0.75$

Single WSI:

		PREDICTED Pos	PREDICTED Neg
TRUE	Pos	805	18
	Neg	48	614

All WSIs:

		PREDICTED Pos	PREDICTED Neg
TRUE	Pos	24796	4340
	Neg	5345	158754

Ok, does it detect cancer?

Model evaluation – attempt 3 – FROC

Detect *particular tumors* ?

How to evaluate the quality of tumor detection?

Model evaluation – attempt 3 – FROC

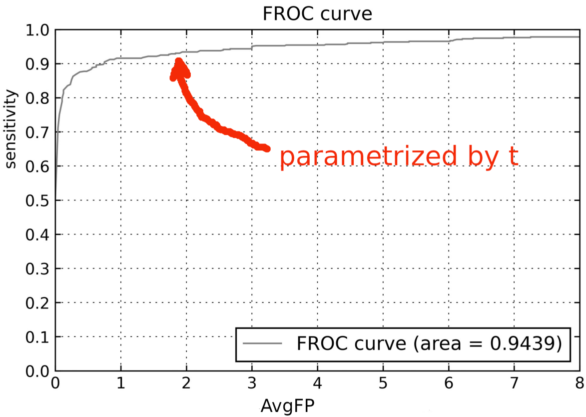

sensitivity ≈ the proportion of tumors containing at least one patch I with $F_\theta(I) \geq t$ w.r.t. all tumors

AvgFP ≈ the proportion of patches I with $F_\theta(I) \geq t$ w.r.t. all patches from non-cancerous WSI

Interpretable AI

What features of the input I determine the value $F_\theta(I)$?

Interpretable AI

What features of the input I determine the value $F_\theta(I)$?

Huge research area

- Gradient based methods (consider $\delta F_\theta(I)/\delta I$)
- Surrogate models (LIME etc.)
- **Occlusion** based methods
- ...

Interpretable AI

What features of the input I determine the value $F_\theta(I)$?

Huge research area

- Gradient based methods (consider $\delta F_\theta(I)/\delta I$)
- Surrogate models (LIME etc.)
- **Occlusion** based methods
- ...

The occlusion = cover a part of the input patch I obtaining I_{occ} and compute $F_\theta(I) - F_\theta(I_{occ})$

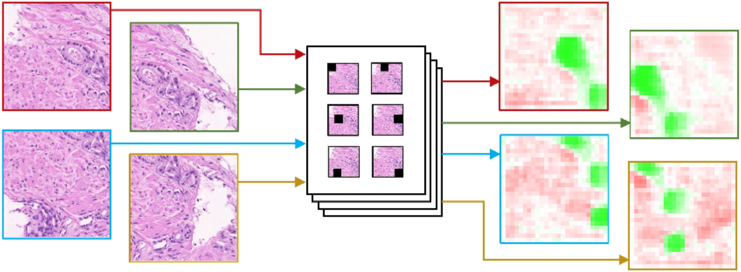

The occlusion results

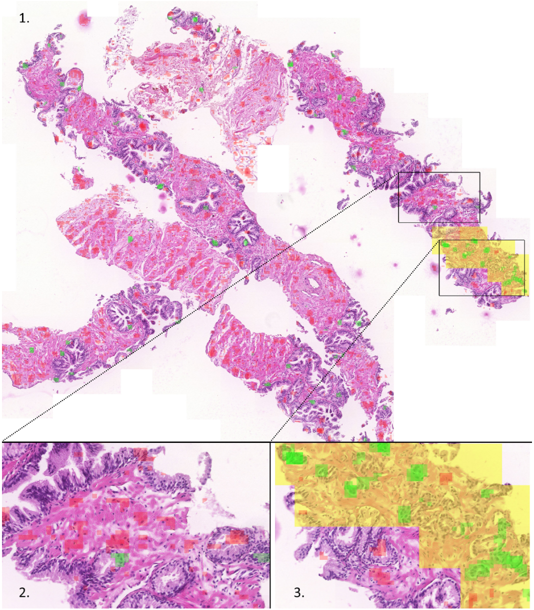

The occlusion results

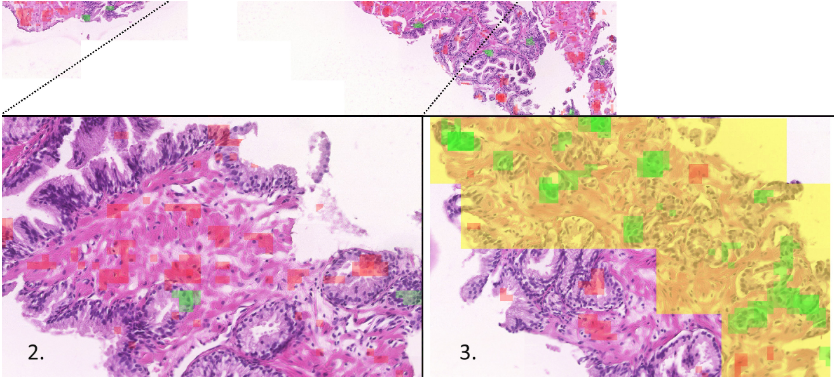

The occlusion results

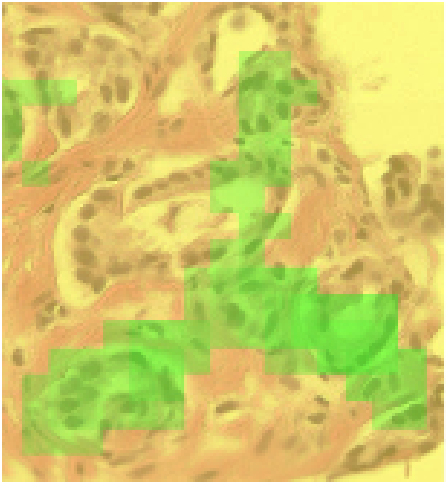

But still, what does it look for?

Biological interpretation

The experiment:

- 647 regions of tissue around randomly selected points from 86 test WSI (37 w/ cancer, 49 w/out cancer)

Biological interpretation

The experiment:

- 647 regions of tissue around randomly selected points from 86 test WSI (37 w/ cancer, 49 w/out cancer)
  - Regions sampled from a grid (points = itersections of lines)
  - a region eligible only if its average explanation score in the square 15px x 15px around the point is sufficiently unambiguous

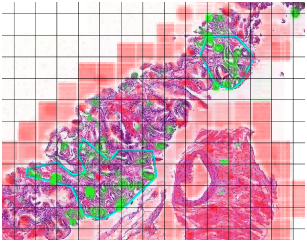

Each region classified according to known biological features used in routine tumor detection.

Biological interpretation

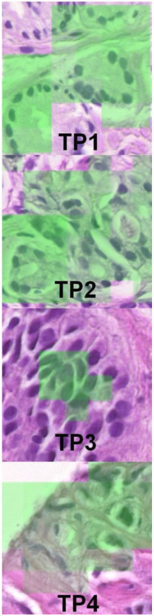

single layered epithelium

small lumina

high cellular density

hyperchromatic nuclei with halo

Biological interpretation

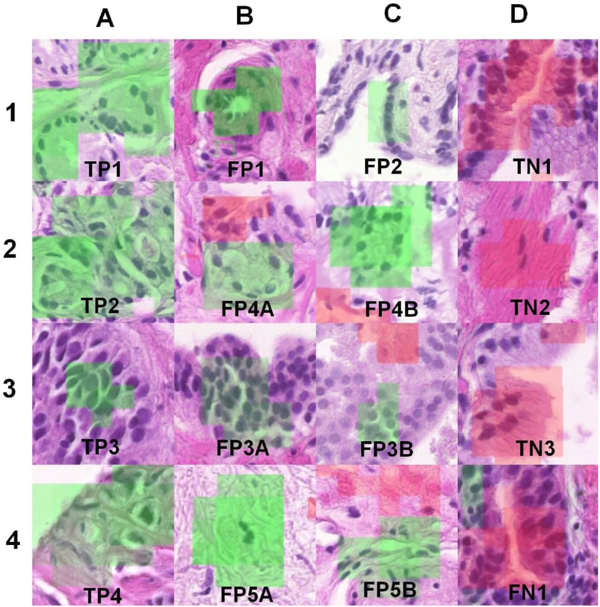

Biological interpretation

Pattern Description	WSI w/ carcinoma Total (N=37)	WSI w/ carcinoma TP %	WSI w/o carcinoma N=49	Total	Tot. %
Single layered epithelium (TP1)	132	52.22%	-	132	20.40%
Small lumina (TP2)	57	22.53%	-	57	8.81%
High cellular density (TP3)	48	18.97%	-	48	7.42%
Hyperchromatic nuclei with halo (TP4)	16	6.32%	-	16	2.47%
Small blood vessel (FP1)	1	-	10	11	1.70%
Single layered epithelium (FP2)	6	-	29	35	5.41%
High cellular density (FP3)	10	-	8	18	2.78%
Small lumina (FP4)	3	-	25	28	4.33%
Hyperchromatic nuclei with halo (FP5)	3	-	12	15	2.32%
Two layered epithelium (TN1)	43	-	29	72	11.13%
Low cellular density (TN2)	37	-	125	162	25.04%
Highly polarised cells (TN3)	9	-	30	39	6.03%
Two layered epithelium (FN1)	1	-	0	1	0.15%
Undefined	9	-	4	13	2.01%

- Biologically significant interpretation in 97.99 %
- WSI w/ carcinoma: More than 90% correct interpretation! (occasionally found an error in the annotation)

Biological interpretation

Pattern Description	WSI w/ carcinoma Total (N=37)	WSI w/ carcinoma TP %	WSI w/o carcinoma N=49	Total	Tot. %
Single layered epithelium (TP1)	132	52.22%	-	132	20.40%
Small lumina (TP2)	57	22.53%	-	57	8.81%
High cellular density (TP3)	48	18.97%	-	48	7.42%
Hyperchromatic nuclei with halo (TP4)	16	6.32%	-	16	2.47%
Small blood vessel (FP1)	1	-	10	11	1.70%
Single layered epithelium (FP2)	6	-	29	35	5.41%
High cellular density (FP3)	10	-	8	18	2.78%
Small lumina (FP4)	3	-	25	28	4.33%
Hyperchromatic nuclei with halo (FP5)	3	-	12	15	2.32%
Two layered epithelium (TN1)	43	-	29	72	11.13%
Low cellular density (TN2)	37	-	125	162	25.04%
Highly polarised cells (TN3)	9	-	30	39	6.03%
Two layered epithelium (FN1)	1	-	0	1	0.15%
Undefined	9	-	4	13	2.01%

- Biologically significant interpretation in 97.99 %
- WSI w/ carcinoma: More than 90% correct interpretation! (occasionally found an error in the annotation)

The Holy Grail: Add new lines to the table! (not yet achieved)

Explainable AI

We know what the model looks for.

But what does it think?

How do the parameters θ affect the value of $F_\theta(I)$?

Explainable AI

We know what the model looks for.

But what does it think?

How do the parameters θ affect the value of $F_\theta(I)$?

Quickly growing research area of XAI

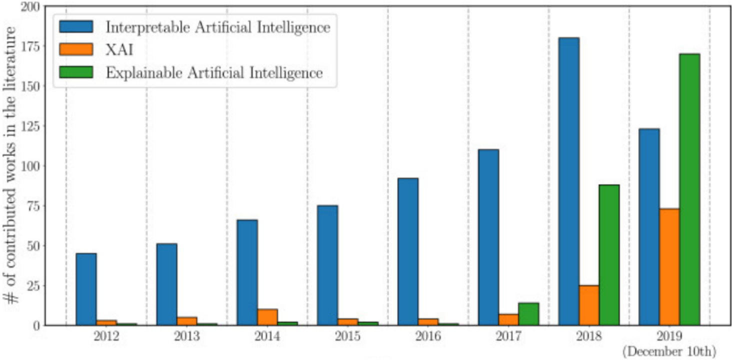

The model we use – VGG-16

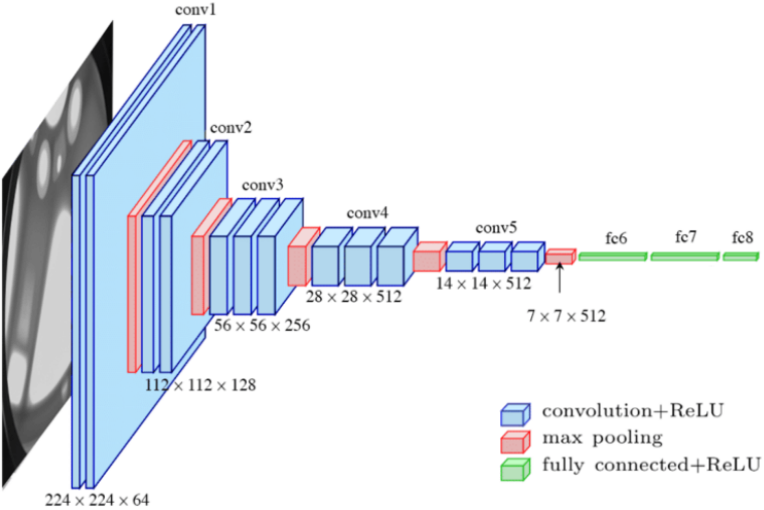

VGG-16 explanation

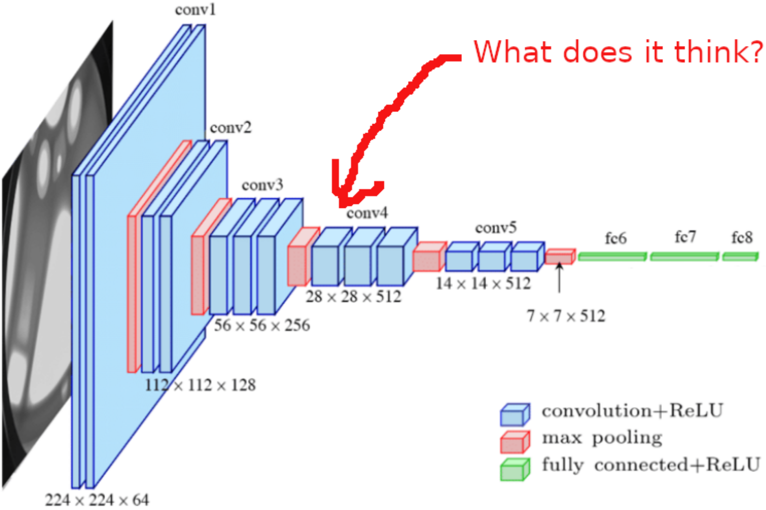

Explanation

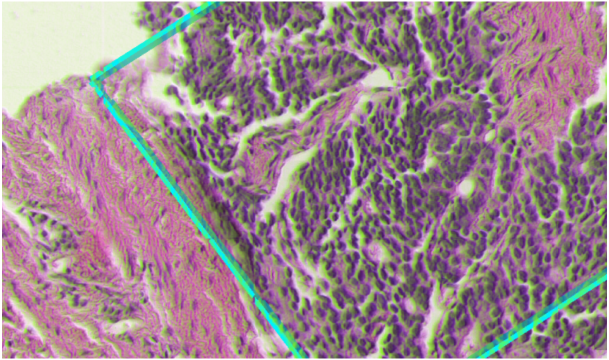

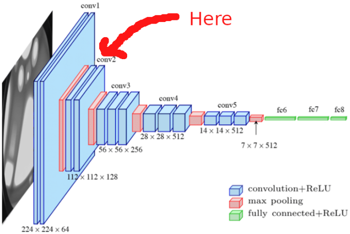

Explanation

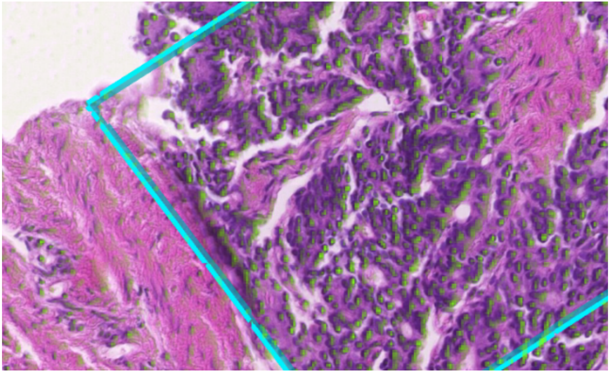

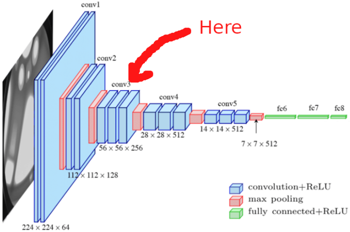

Explanation

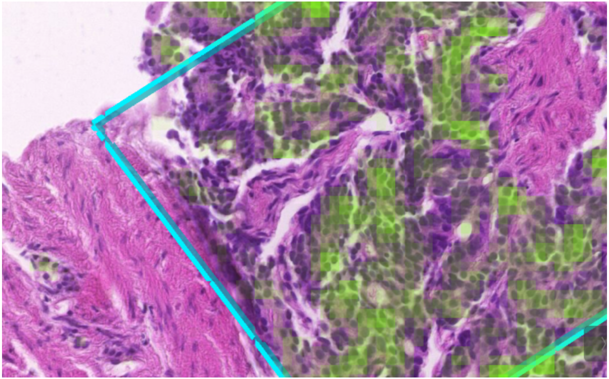

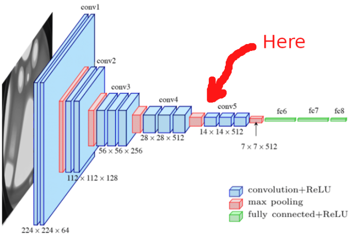

Conclusions

- Developed a deep learning pipeline for WSI
  Mostly from known components
- Evaluated the interpretation from the pathologist’s point of view
- Developed a visualization system allowing smooth inspection of networks’ performance

... and lots of future work!

RationAI - the team

- **MU**
  - Petr Holub, Tomáš Brázdil
  - Ph.D. students: Matej Gallo, Vojtěch Krajňanský, Rudolf Wittner
  - MSc students: Jakub Hruška, Jan Čech, Tomáš Bíl, Petr Kantek, Lucie Nováková
  - Bc students: Andrej Kubanda, Miroslav Bezák
  - other collaborators: Michal Růžička, Jiří Horák, Martin Kačenga . . .
- **MMCI (MOÚ)**
  - Rudolf Nenutil, Phil Coates, . . .
- **International collaborations**
  - Medical University Graz: Heimo Müller, Kurt Zatloukal
  - CRS4: Luca Pireddu