Development of intelligent systems (RInS)

Surface anomaly detection

Danijel Skočaj
University of Ljubljana
Faculty of Computer and Information Science

Academic year: 2025/26

Surface defect detection

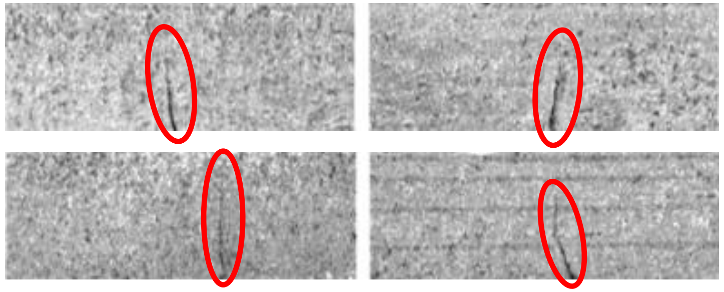

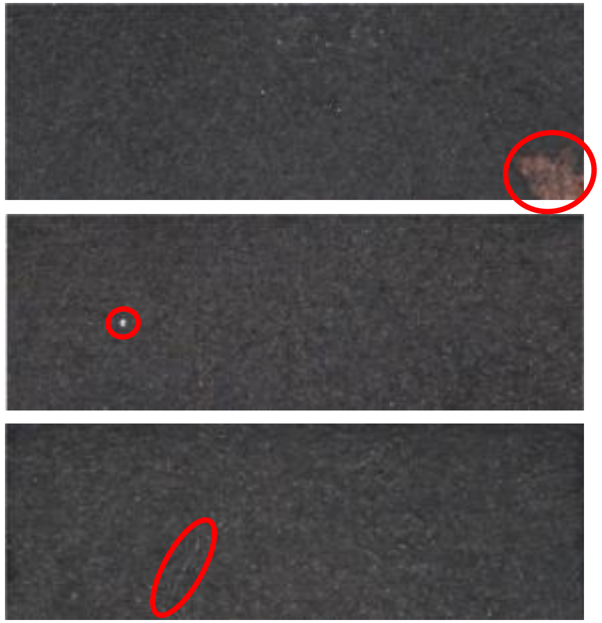

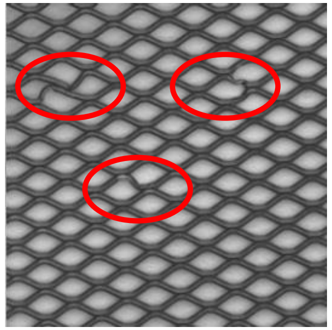

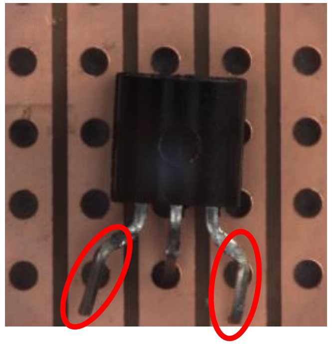

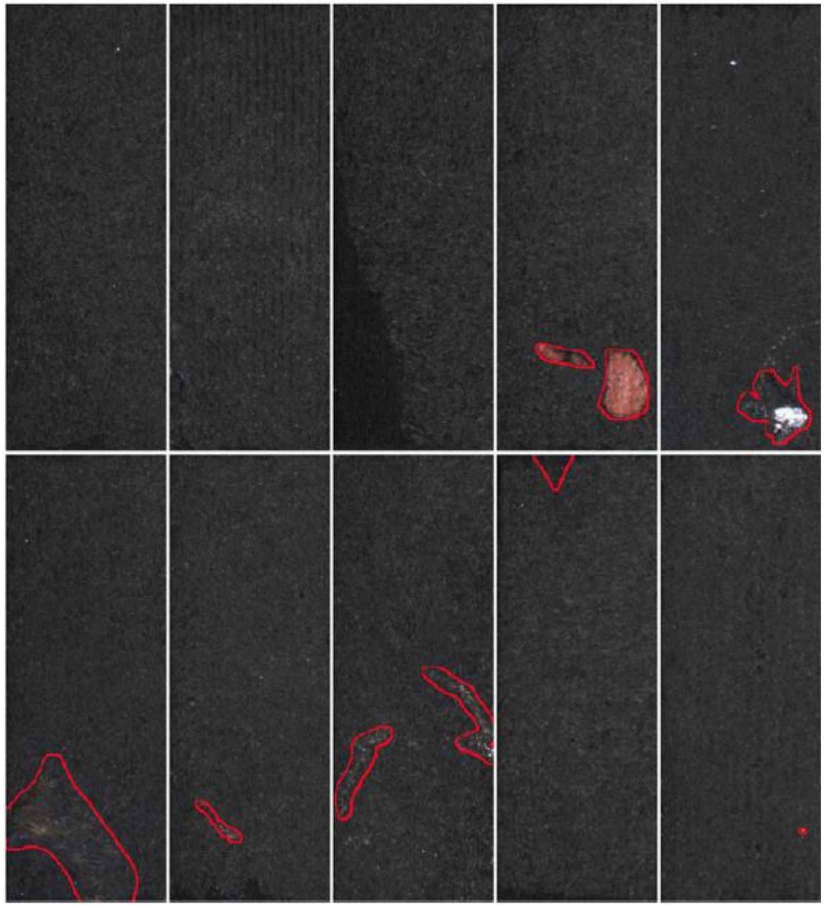

Example: Visual inspection of pharmaceutical products

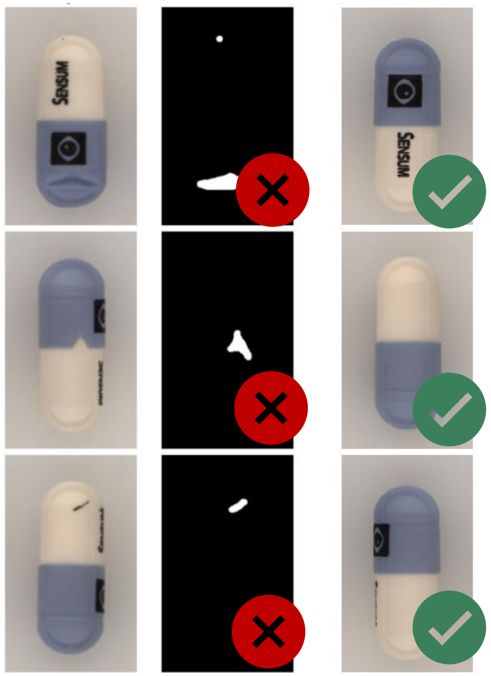

sensum.eu

NCAA 2021

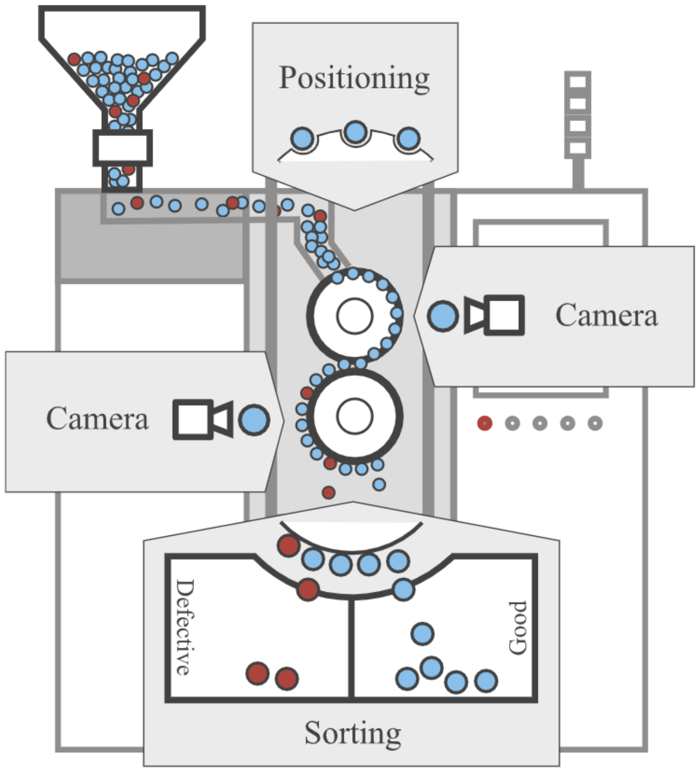

Deep learning paradigm

- Conventional approach: programming specific solutions
- New paradigm: data-driven learning-based sloutions

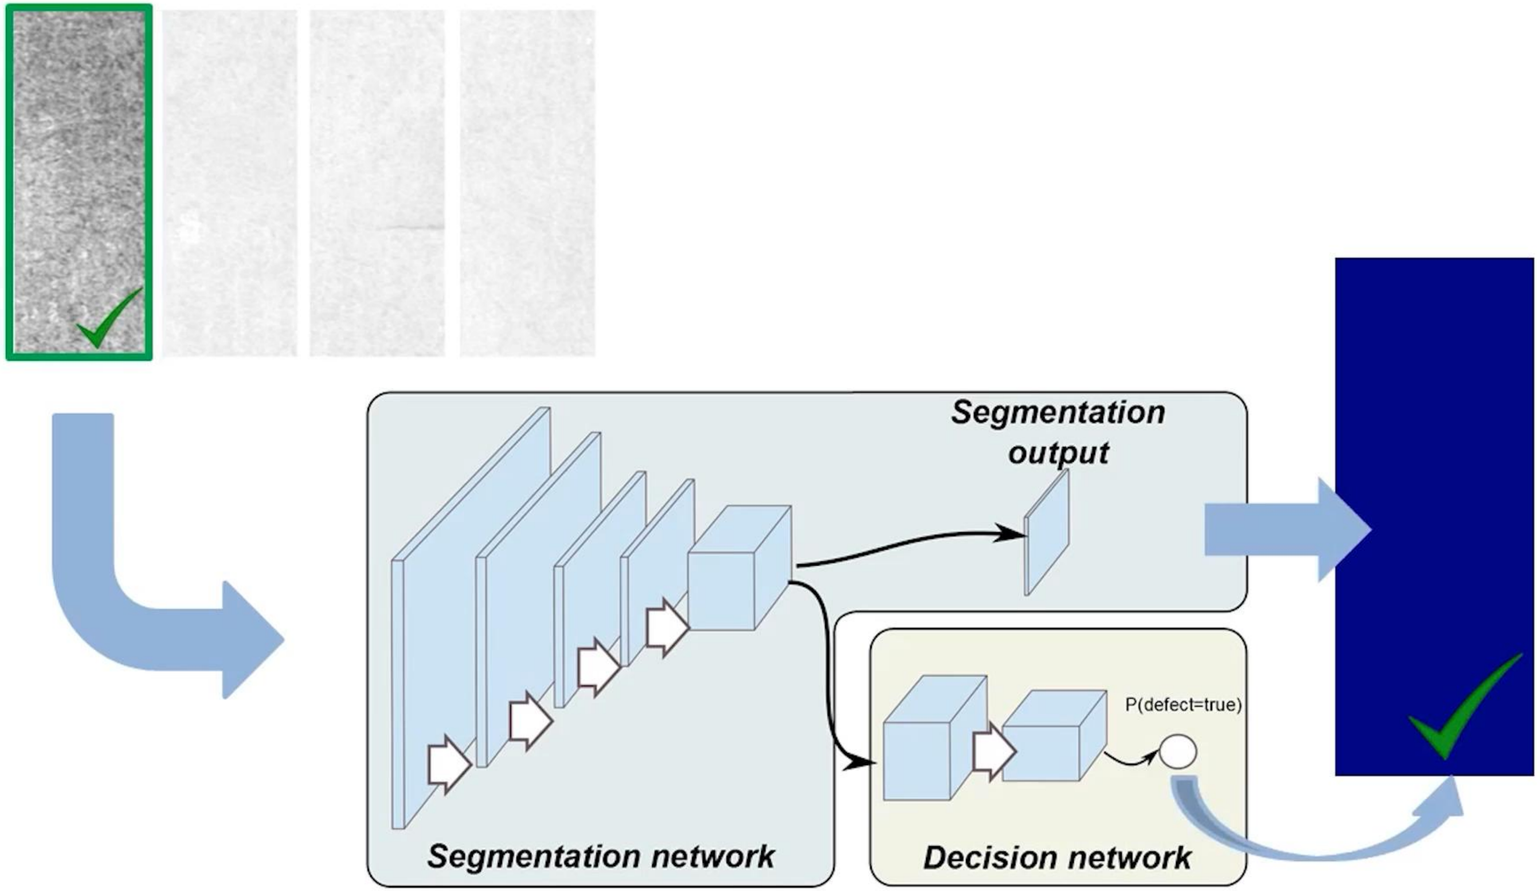

GOSTOP
2016-2020

DIVID
2018-2021

Learning regimes

- Few-shot
- Zero-shot
- Unsupervised
- Mixed supervision
- Weak supervision
- Supervised

Learning regimes

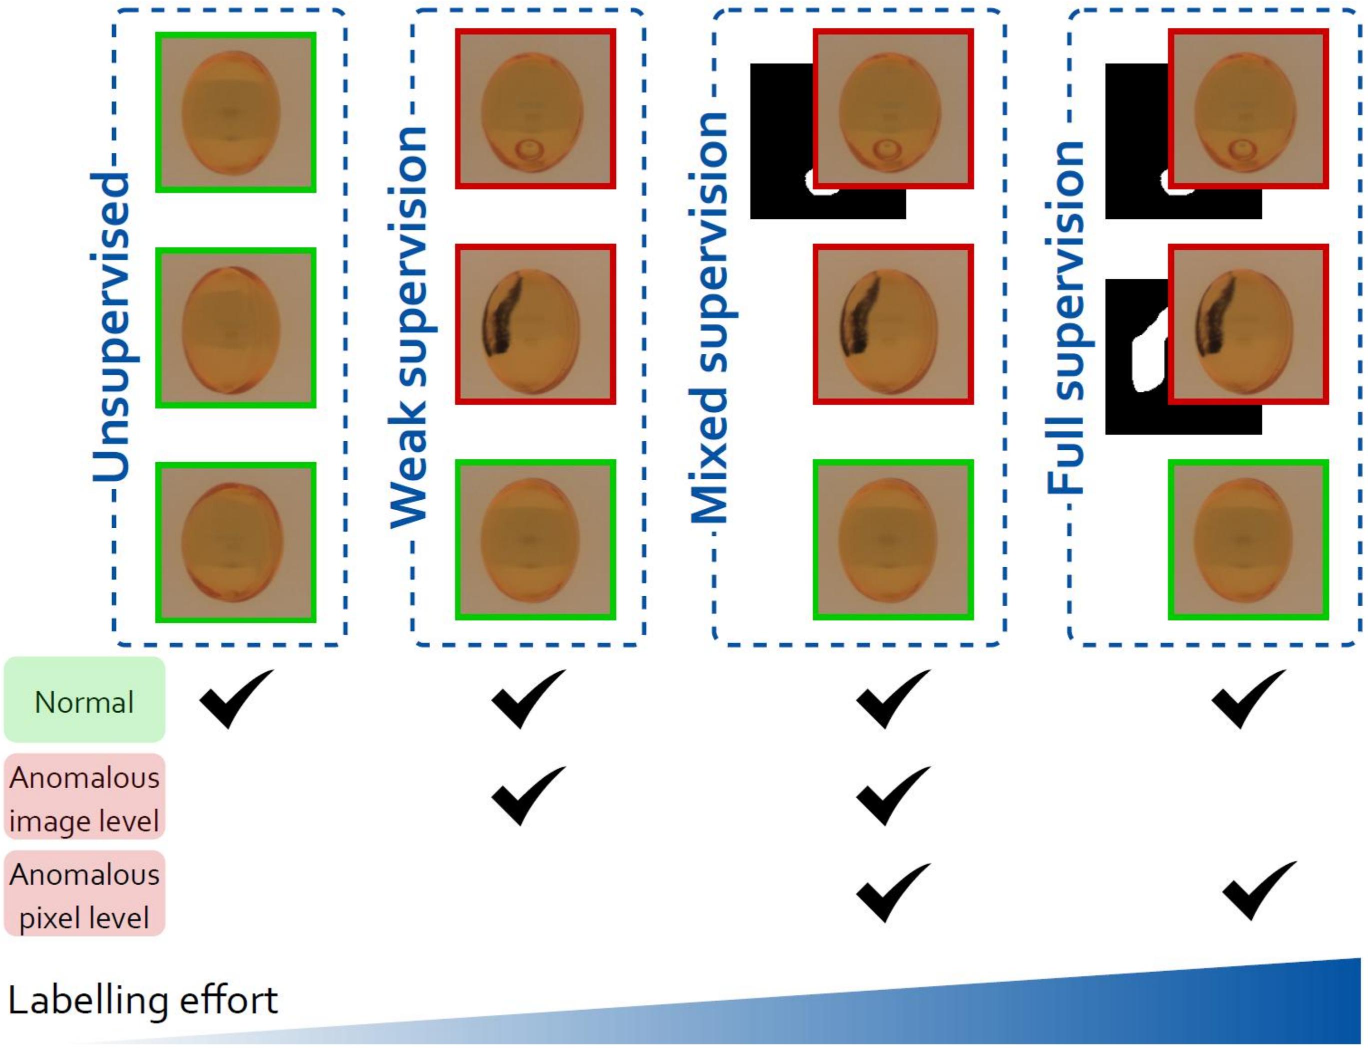

Learning regimes – full supervision

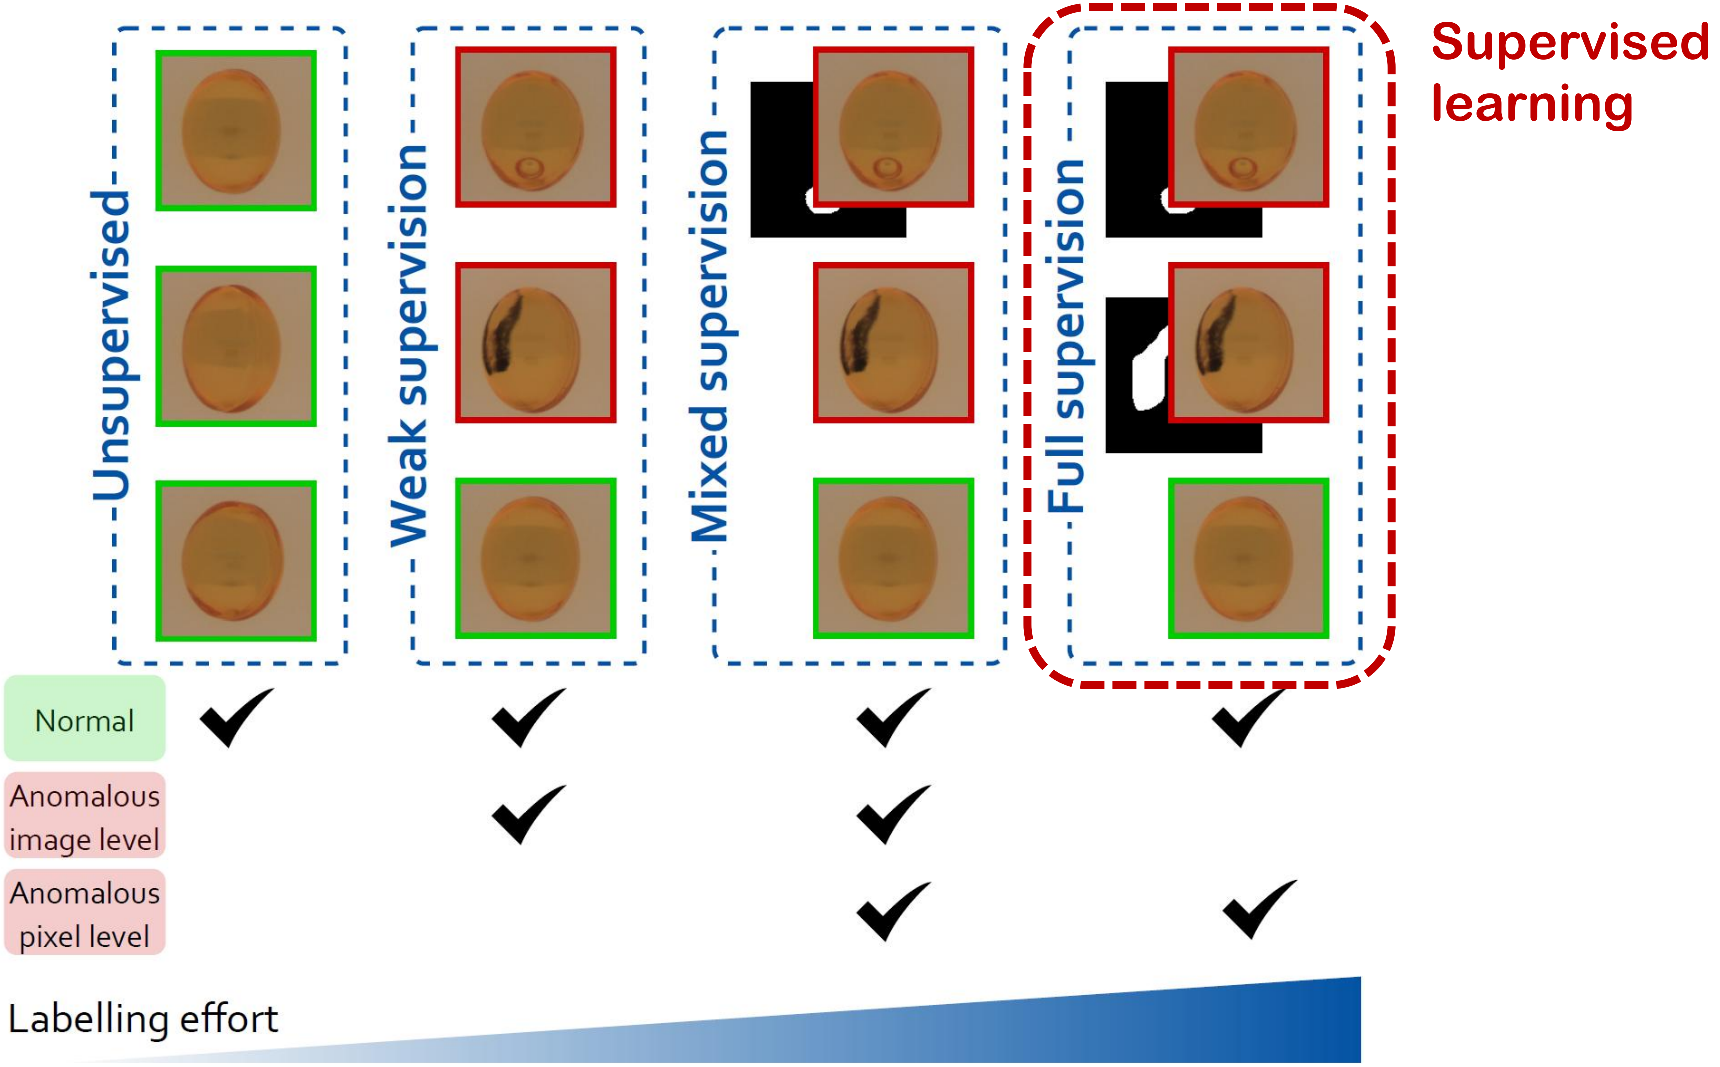

Supervised learning

WACV 2018

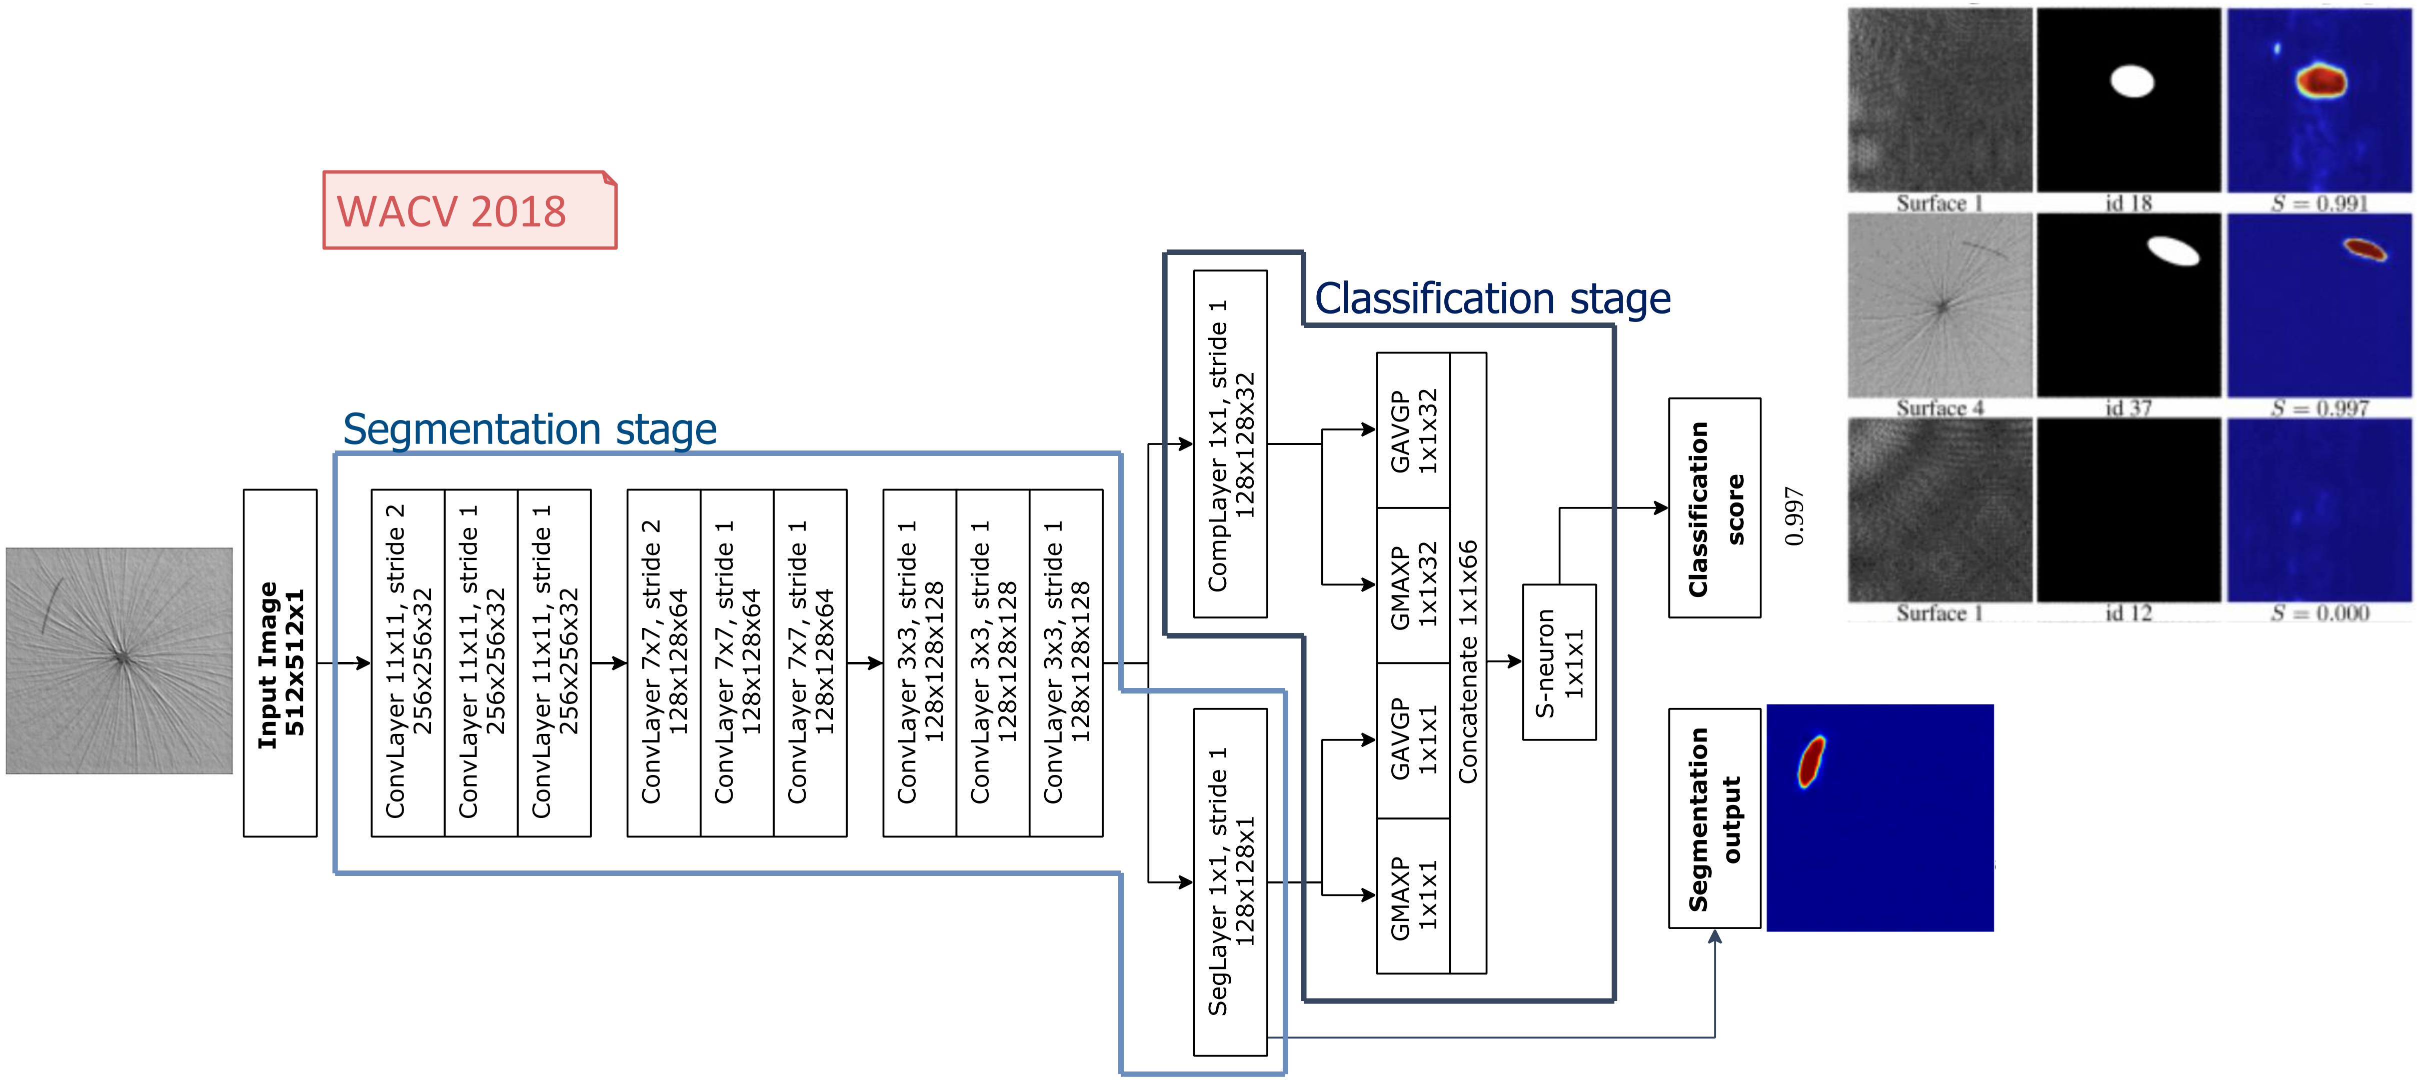

Supervised learning - TriNet

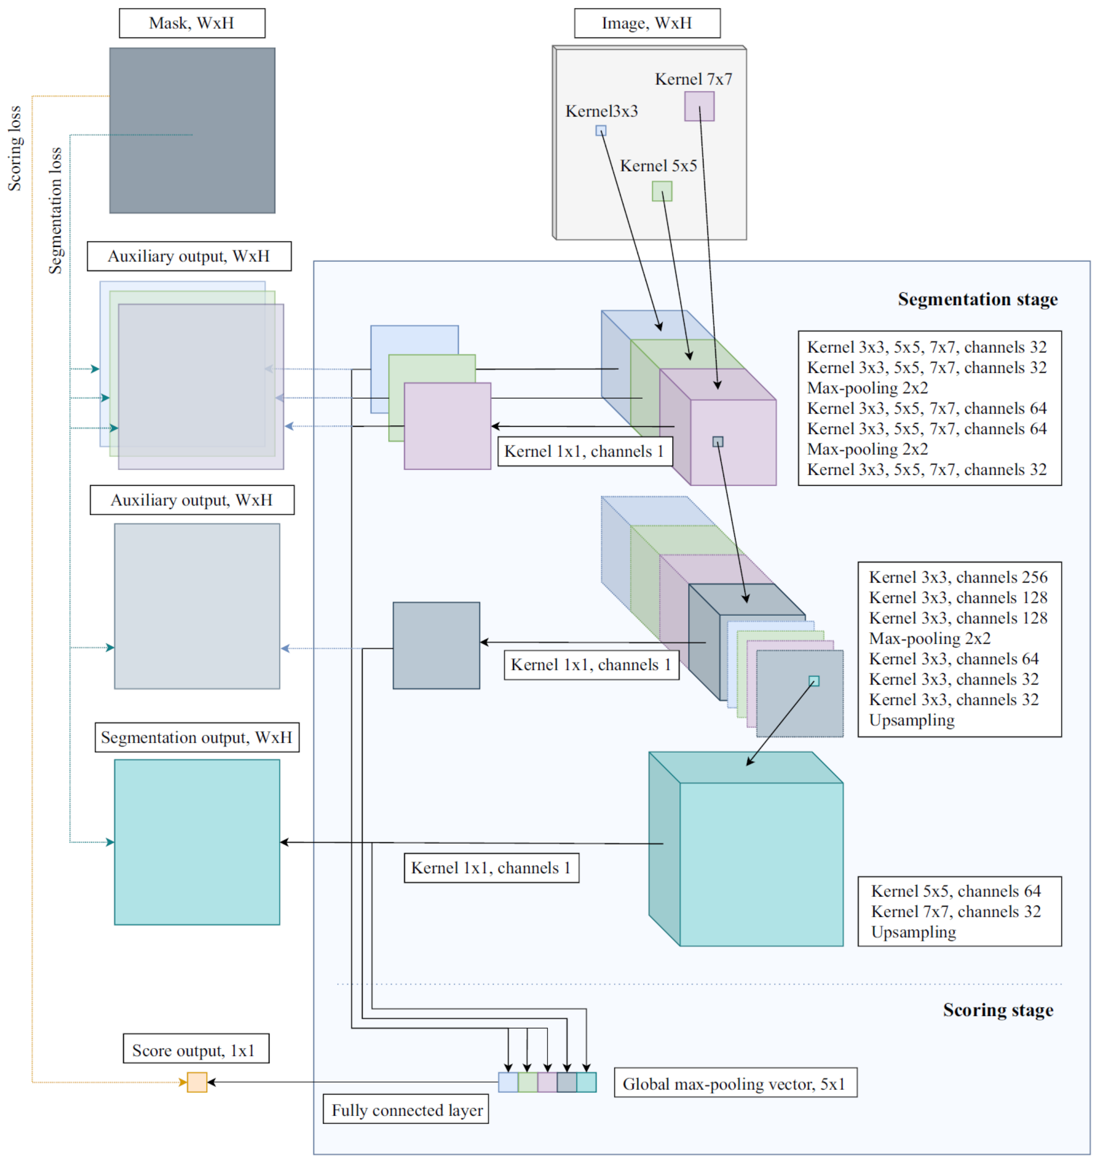

NCAA 2021

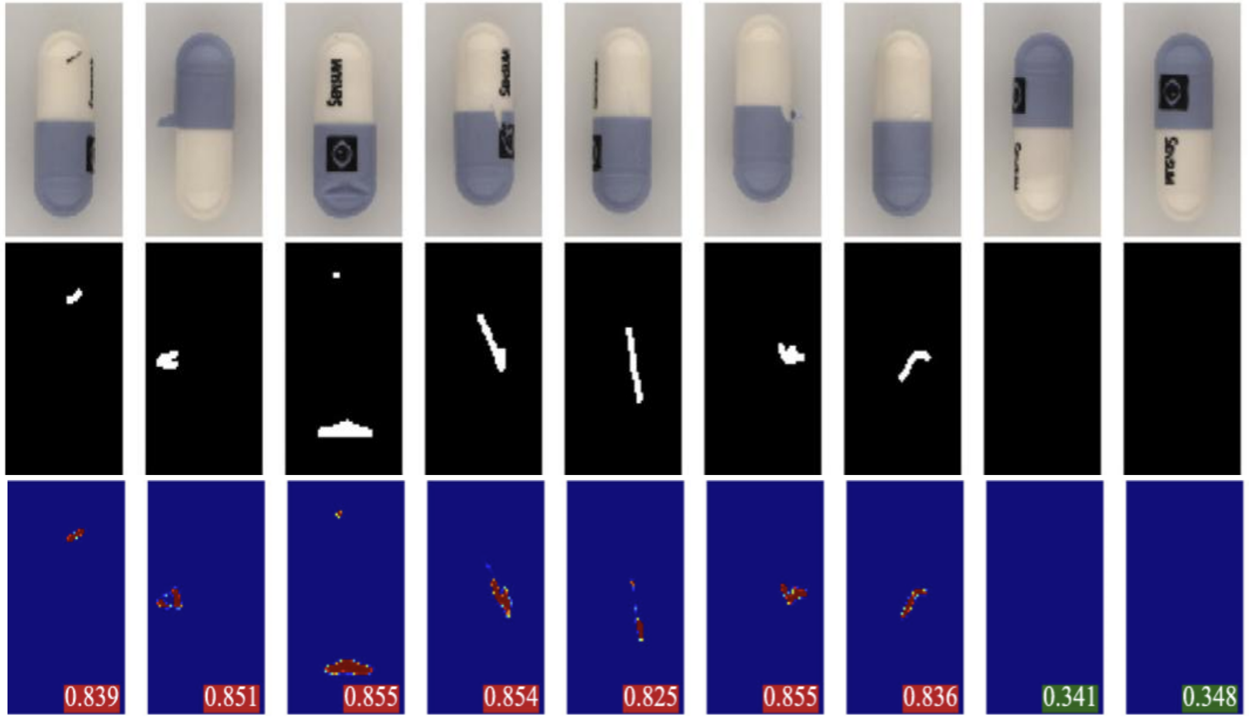

Supervised learning - SegDecNet

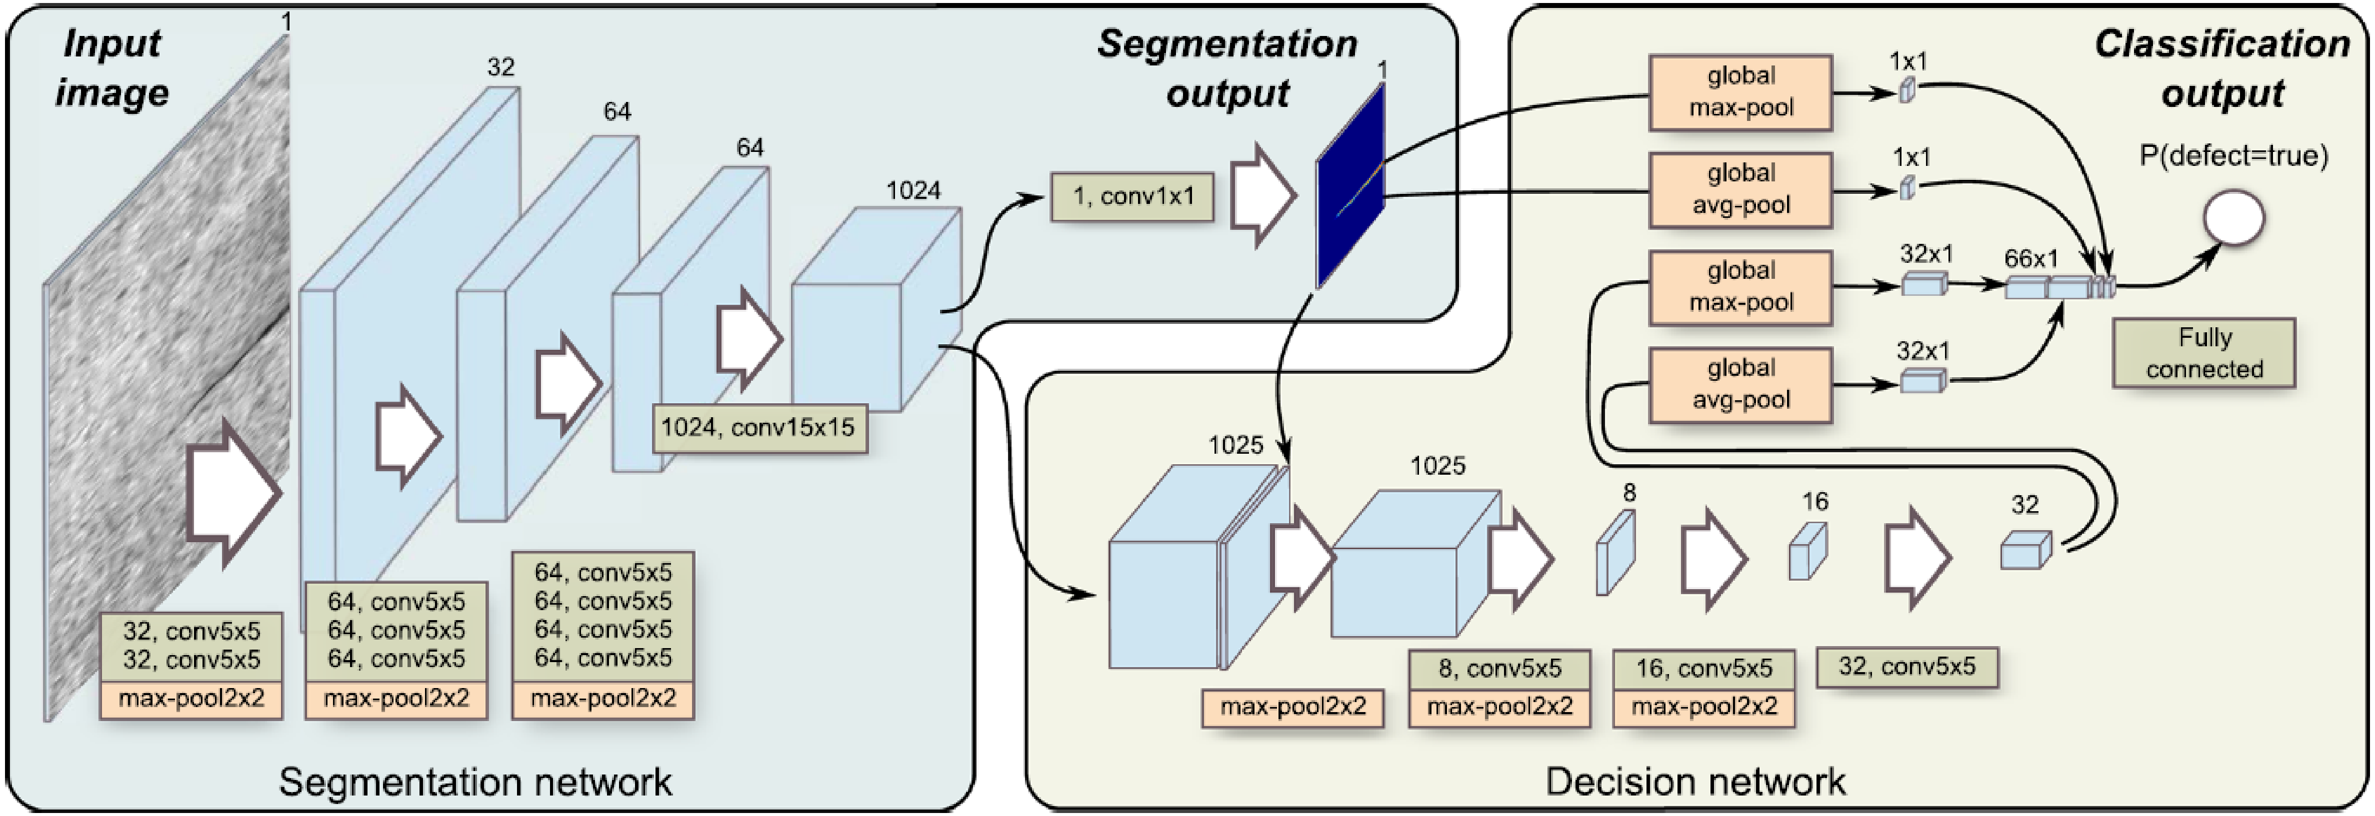

ICVS 2019

JIM 2020

Supervised learning - SegDecNet

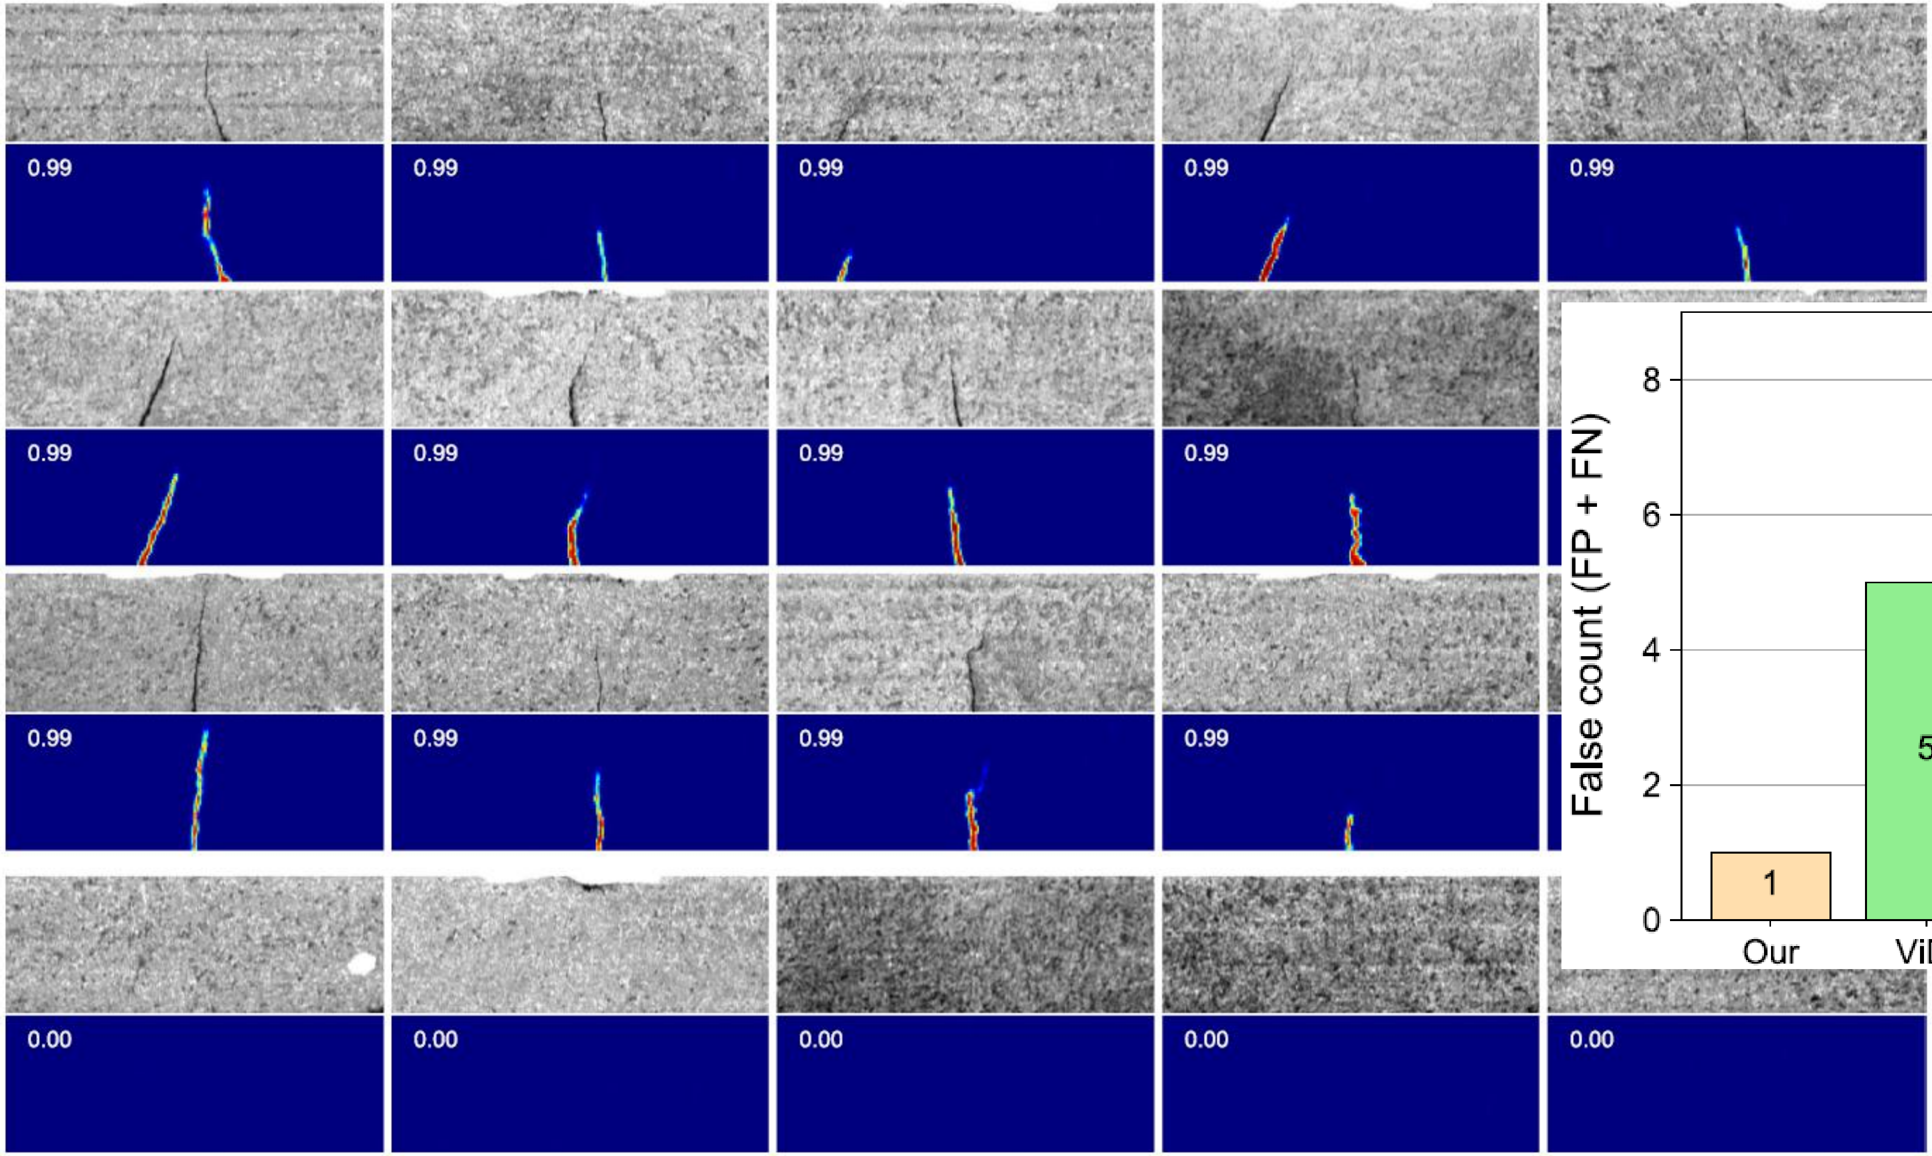

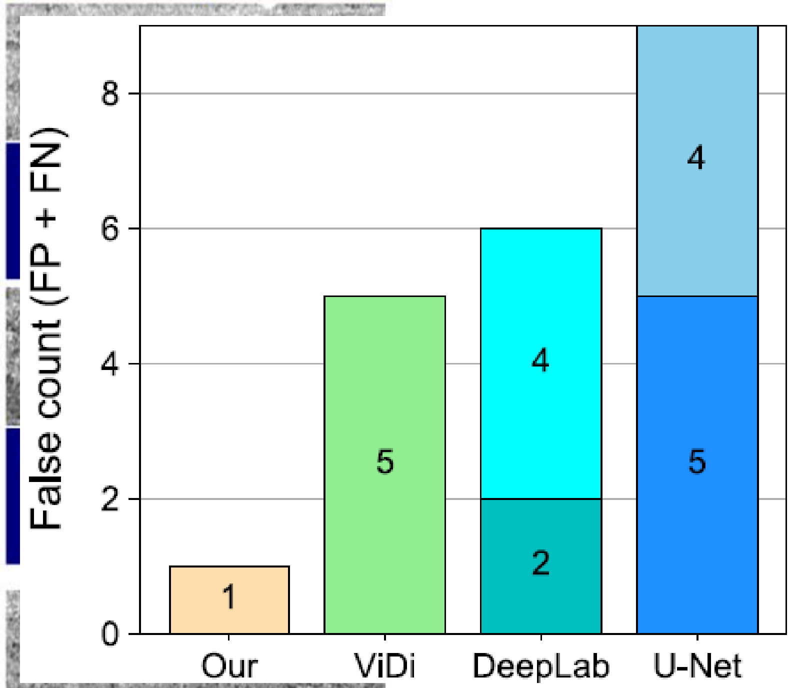

JIM 2020

Learning regimes – unsupervised learning

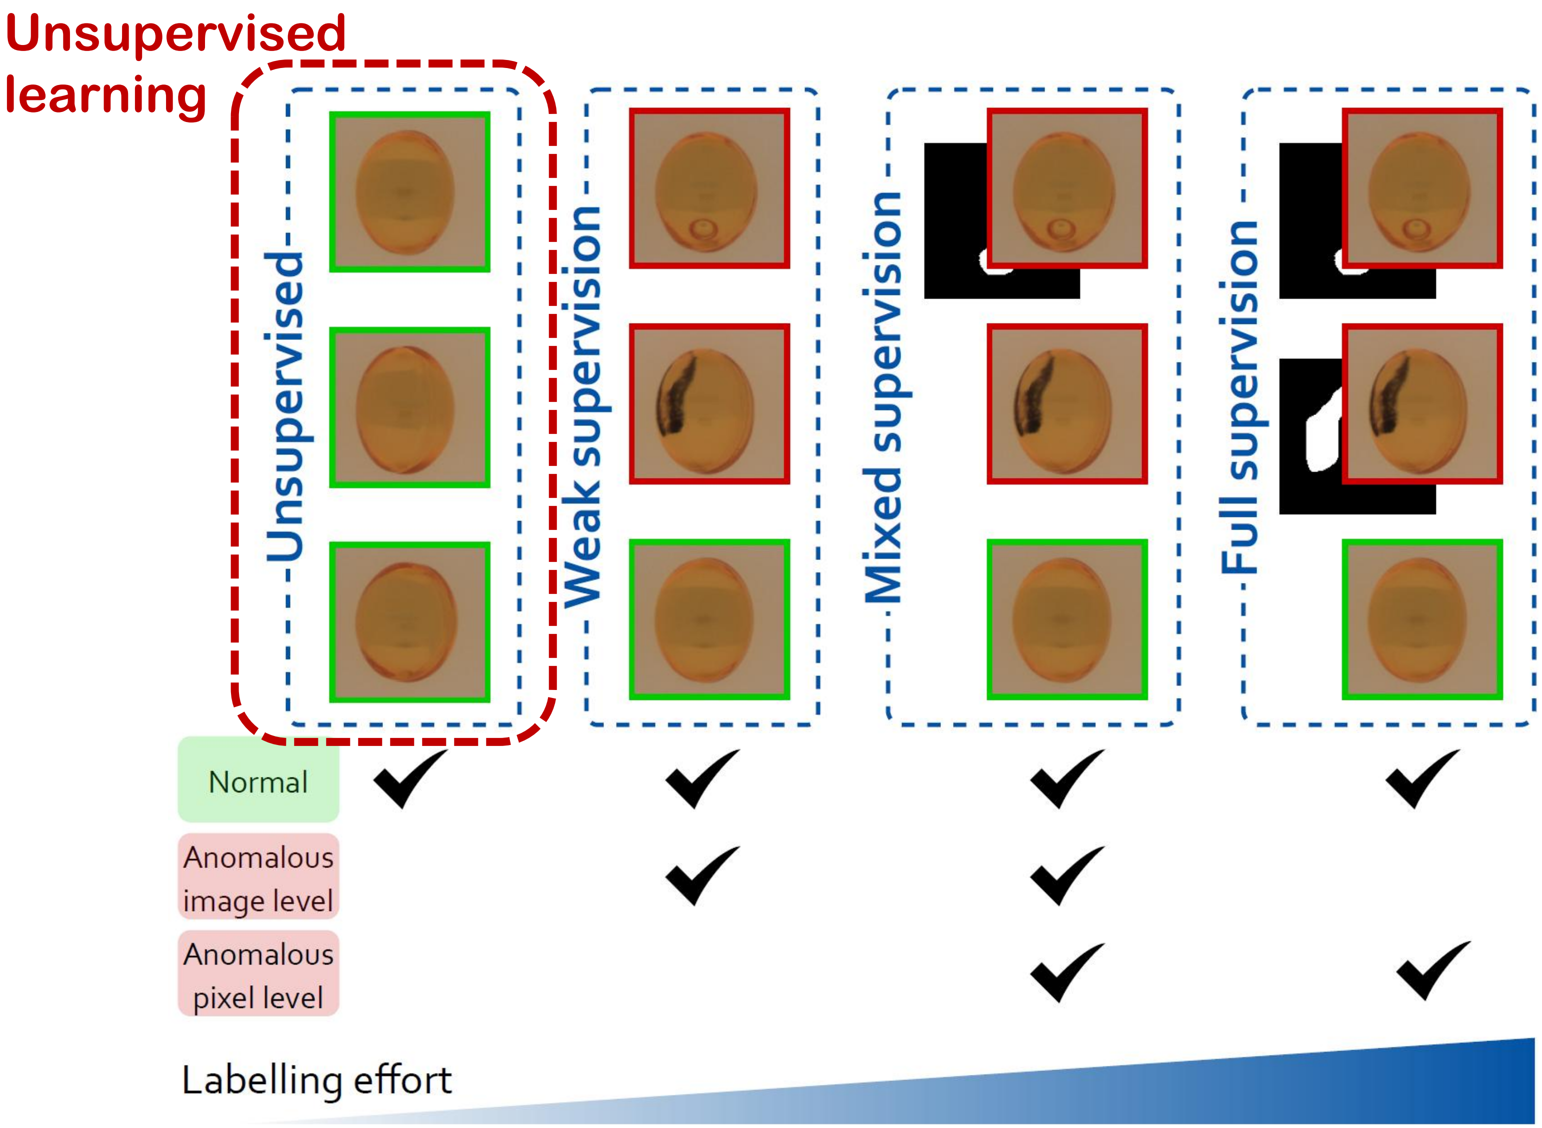

Unsupervised learning

- Only defect-free images required
- Negative-class-only learning

- Detection AUROC on MVTec AD:

[paperswithcode.com]

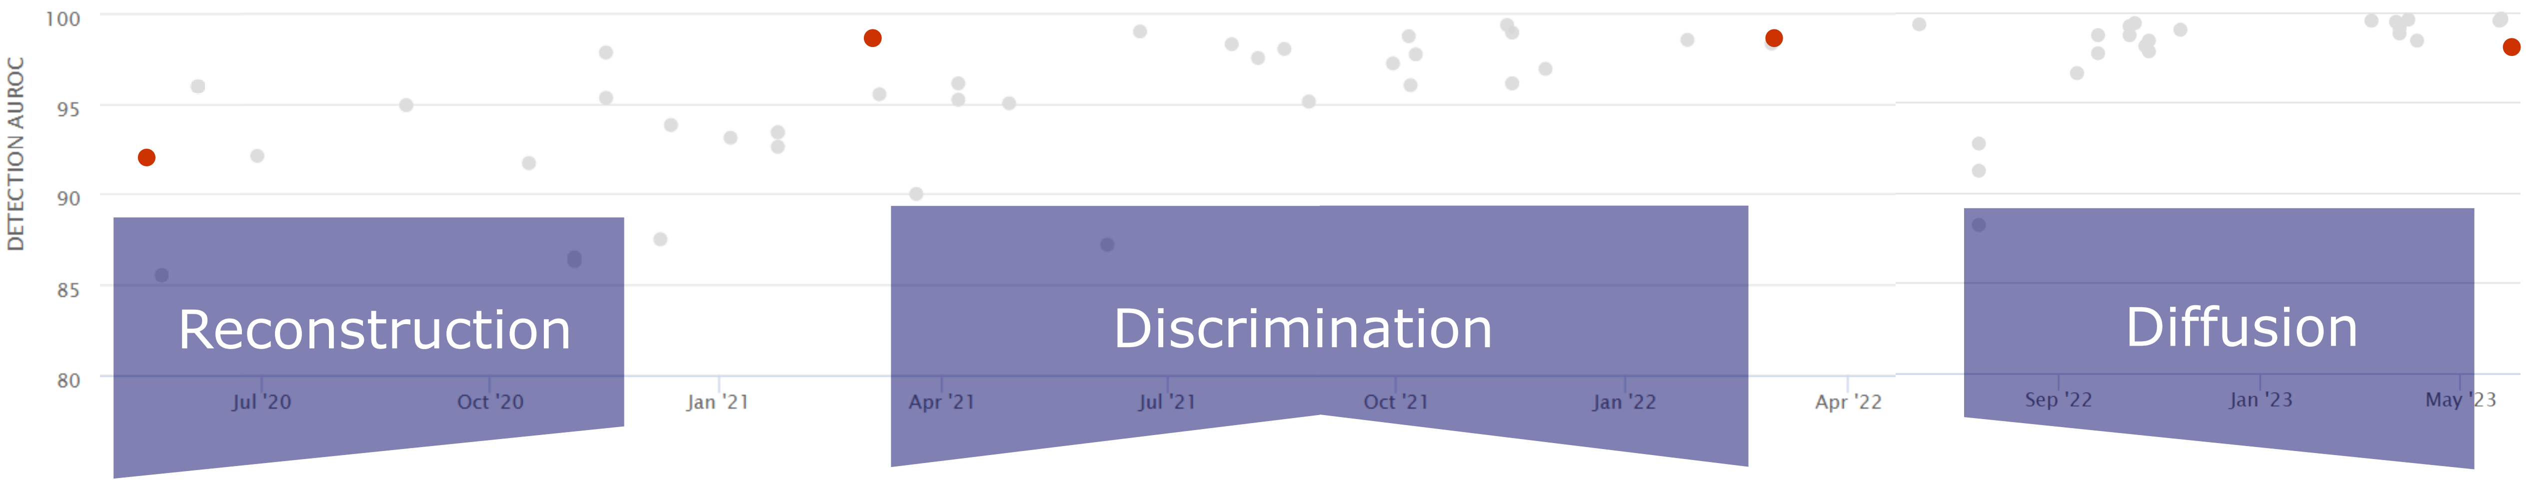

RIAD

PR 2021

DRAEM

ICCV 2021

DSR

ECCV 2022

Transfusion

2023

Unsupervised learning - RIAD

- Reconstructive approach

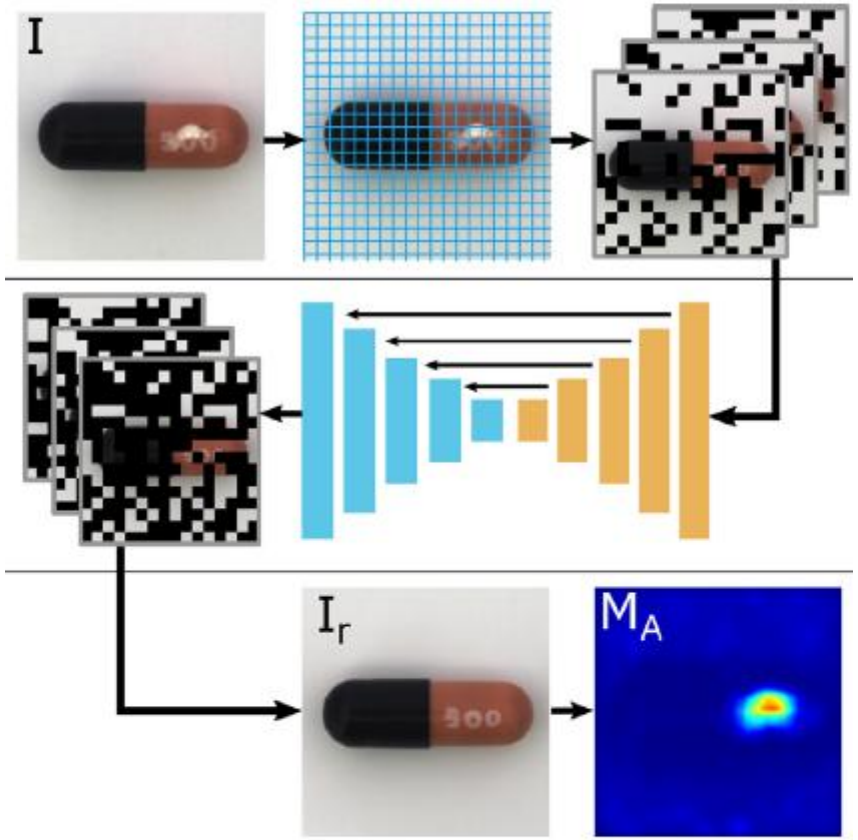

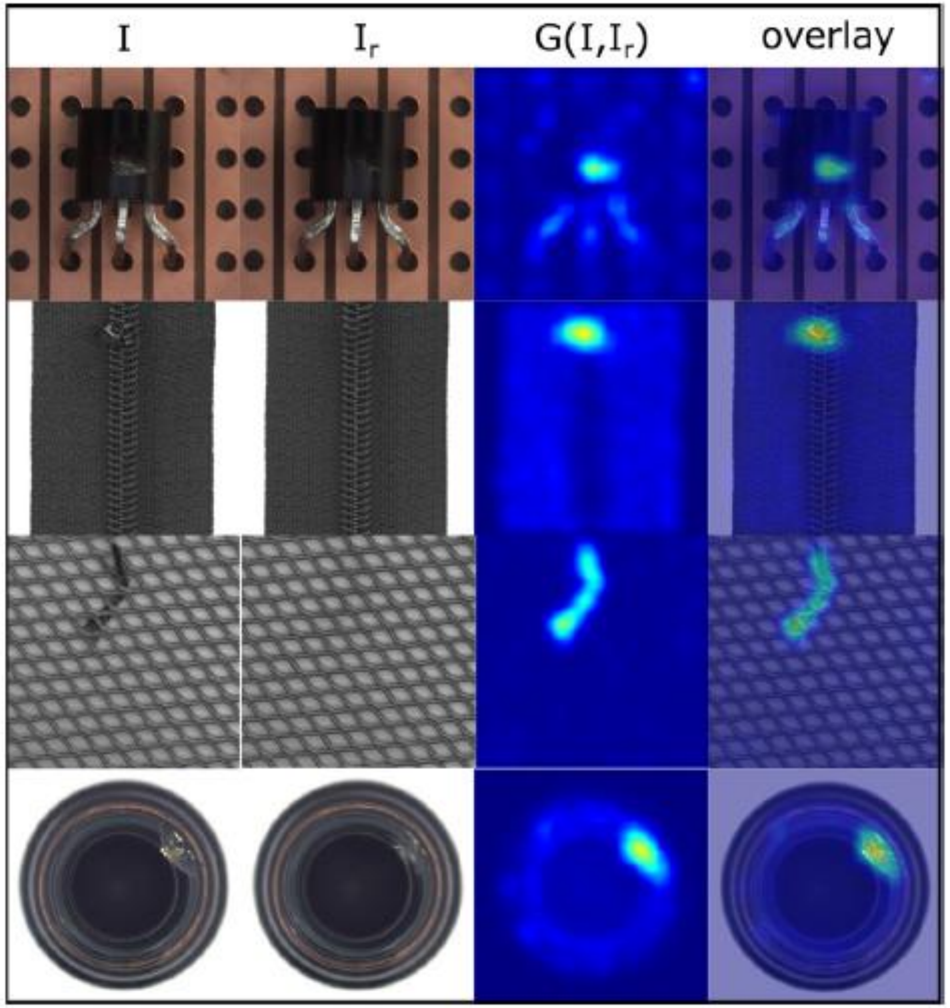

DIVID
2018-2021

CompVis
2019-2024

PR 2021

Unsupervised learning - RIAD

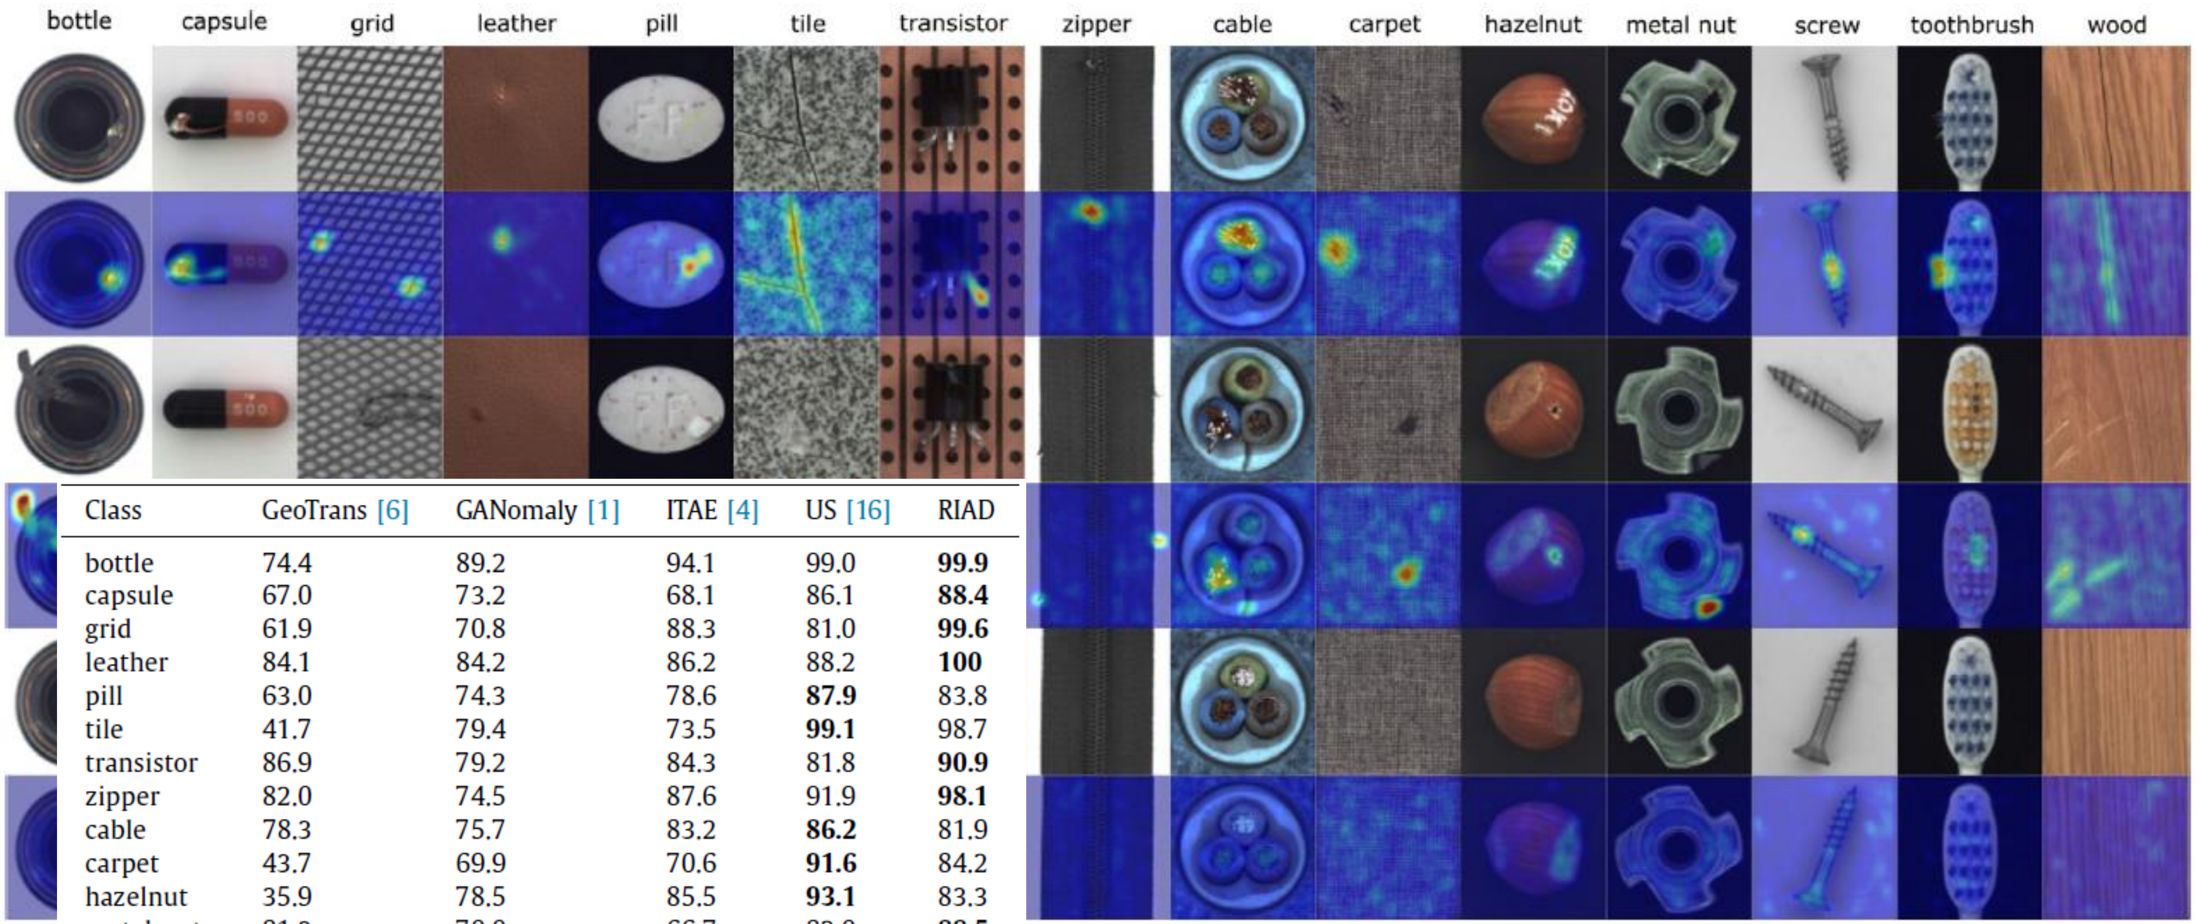

Class	GeoTrans [6]	GANomaly [1]	ITAE [4]	US [16]	RIAD
bottle	74.4	89.2	94.1	99.0	**99.9**
capsule	67.0	73.2	68.1	86.1	**88.4**
grid	61.9	70.8	88.3	81.0	**99.6**
leather	84.1	84.2	86.2	88.2	**100**
pill	63.0	74.3	78.6	**87.9**	83.8
tile	41.7	79.4	73.5	**99.1**	98.7
transistor	86.9	79.2	84.3	81.8	**90.9**
zipper	82.0	74.5	87.6	91.9	**98.1**
cable	78.3	75.7	83.2	**86.2**	81.9
carpet	43.7	69.9	70.6	**91.6**	84.2
hazelnut	35.9	78.5	85.5	**93.1**	83.3
metal nut	81.3	70.0	66.7	82.0	**88.5**
screw	50.0	74.6	**100**	54.9	84.5
toothbrush	97.2	65.3	**100**	95.3	**100**
wood	61.1	83.4	92.3	**97.7**	93.0
avg_{tex}	58.5	76.5	82.2	91.5	**95.1**
avg_{obj}	71.6	75.4	84.8	85.8	**89.9**
avg	67.2	76.2	83.9	87.7	**91.7**

DIVID
2018-2021

CompVis
2019-2024

PR 2021

Unsupervised learning - DRAEM

- Reconstructive models
  - Good approximation of data
  - Unsupervised learning
  - General, task-independent
  - Enable reconstruction and outlier detection

- Discriminative models
  - Supervised learning
  - Task-dependent
  - Compact representations
  - No reconstruction
  - Outlier detection as classification

⇨ Combine reconstructive model and discriminative classifier

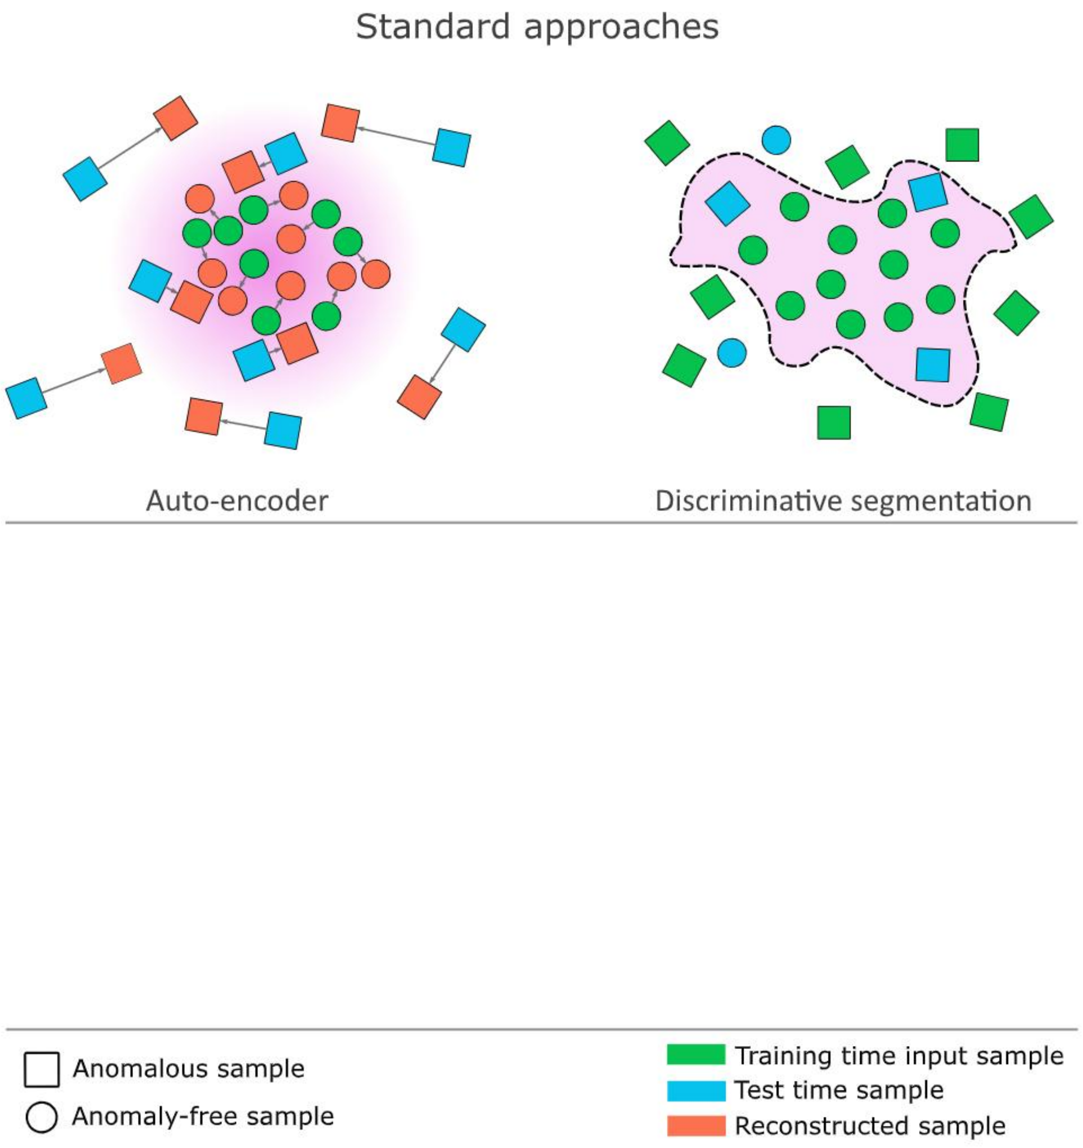

Unsupervised learning - DRAEM

- Reconstructive and discriminative approach
- Generate synthetic anomalies

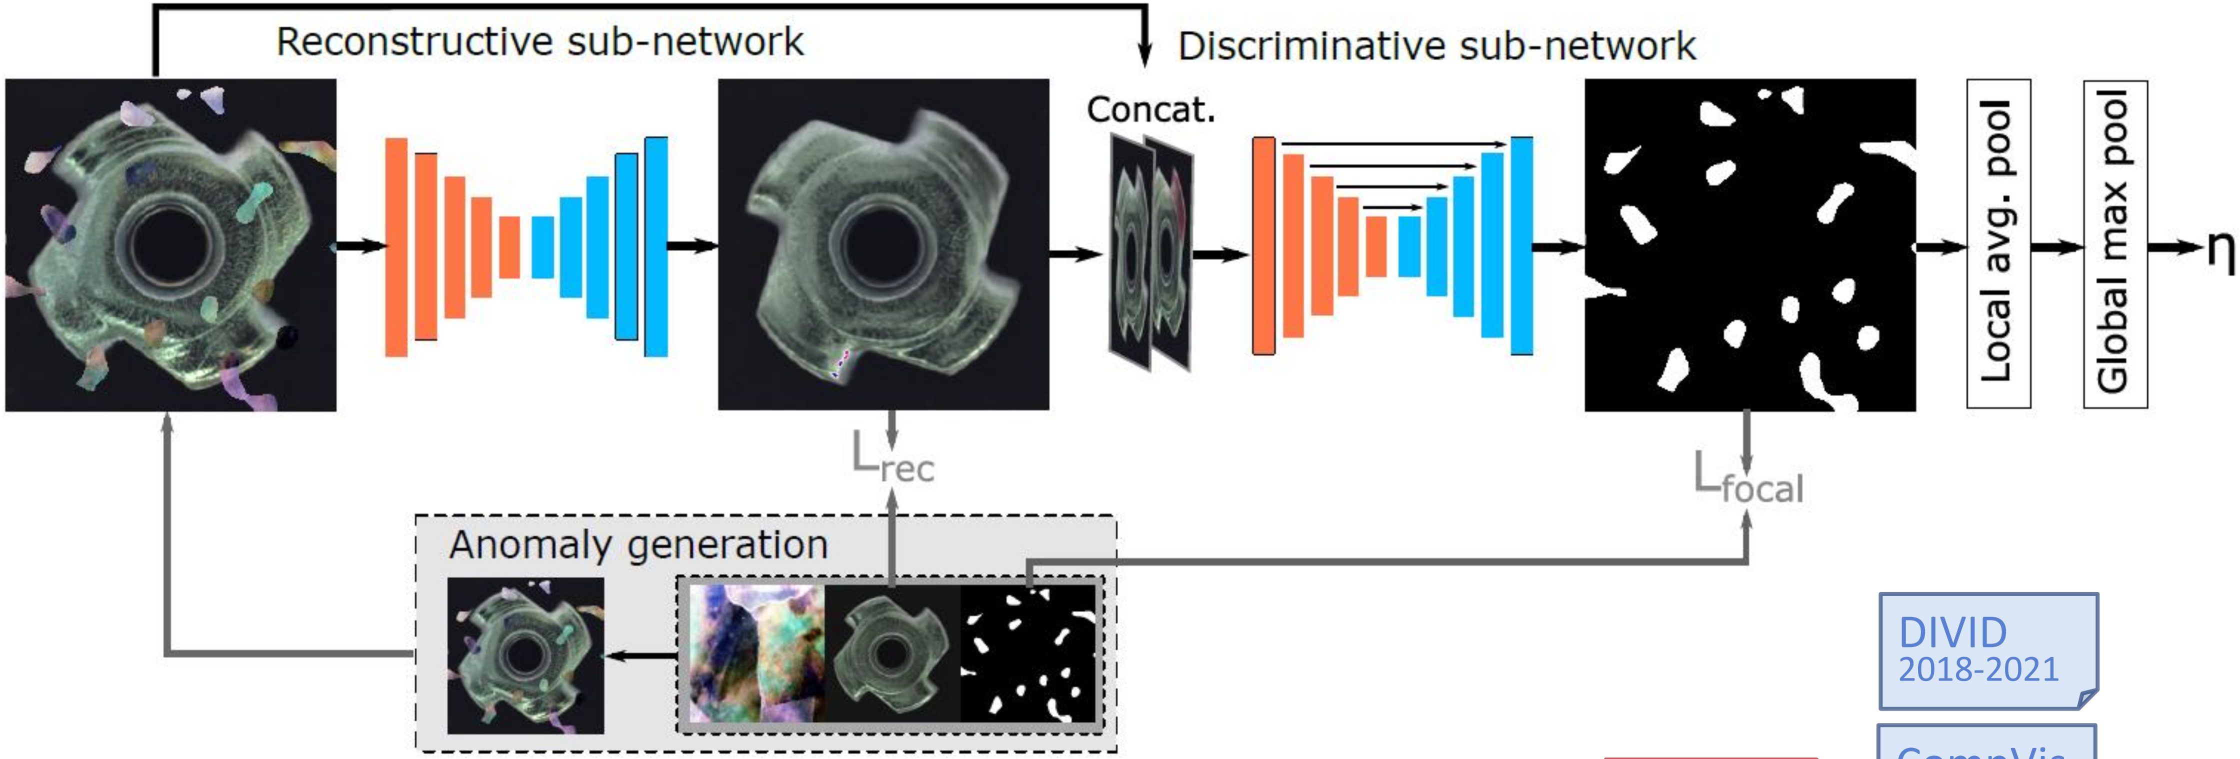

ICCV 2021

DIVID
2018-2021

CompVis
2019-2024

Unsupervised learning - DRAEM

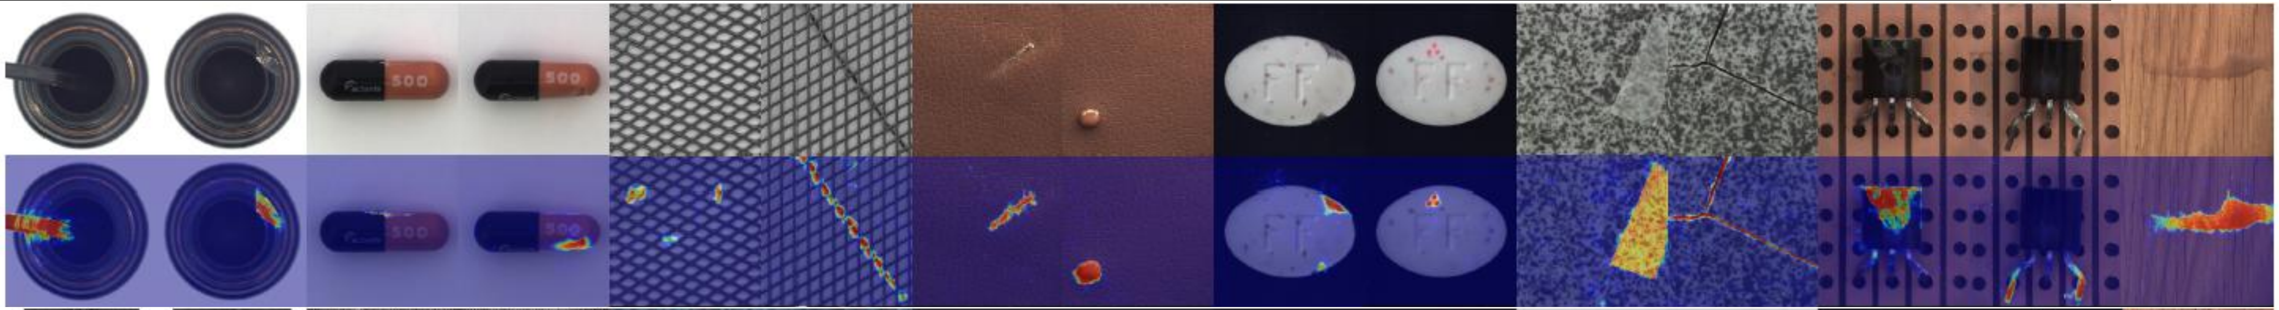

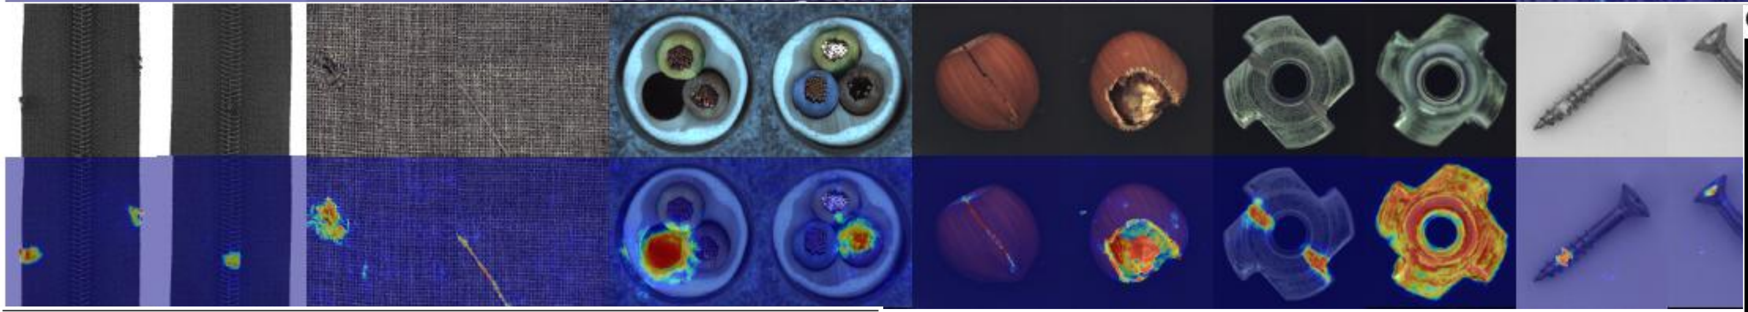

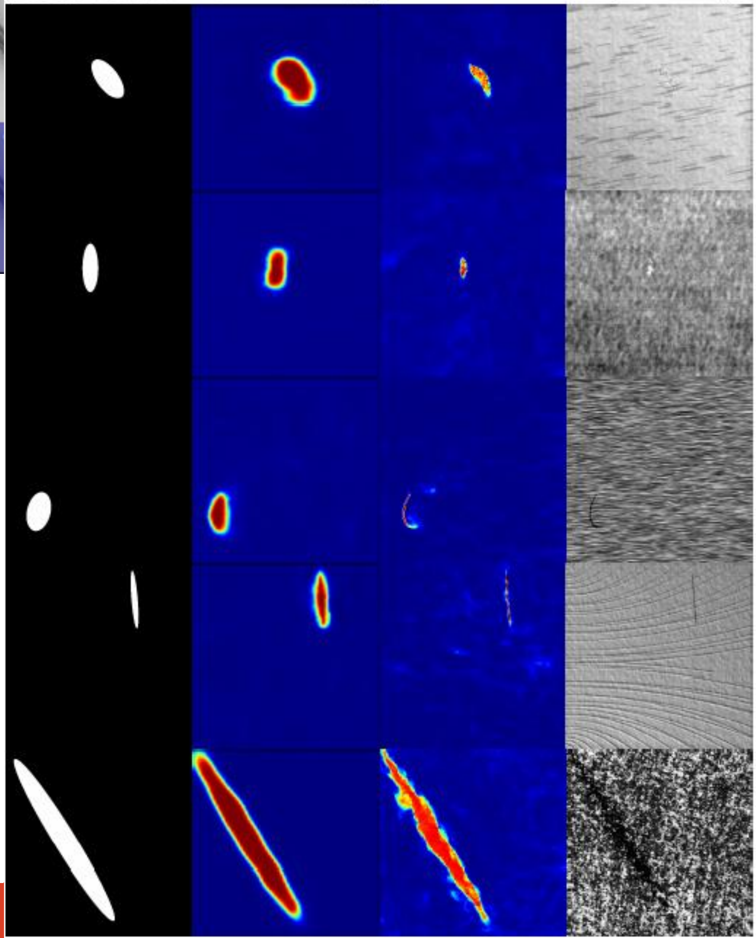

Class	[1]	[26]	[4]	[31]	[20]	[11]	DRÆM
bottle	79.4	98.3	99.0	99.9	**100**	99.9	99.2
capsule	72.1	68.7	86.1	88.4	92.3	91.3	**98.5**
grid	74.3	86.7	81.0	99.6	92.9	96.7	**99.9**
leather	80.8	94.4	88.2	**100**	**100**	**100**	**100**
pill	67.1	76.8	87.9	83.8	83.4	93.3	**98.9**
tile	72.0	96.1	99.1	98.7	97.4	98.1	**99.6**
transistor	80.8	79.4	81.8	90.9	95.9	**97.4**	93.1
zipper	74.4	78.1	91.9	98.1	97.9	90.3	**100**
cable	71.1	66.5	86.2	81.9	**94.0**	92.7	91.8
carpet	82.1	90.3	91.6	84.2	95.5	**99.8**	97.0
hazelnut	87.4	100	93.1	83.3	98.7	92.0	**100.0**
metal nut	69.4	81.5	82.0	88.5	93.1	**98.7**	**98.7**
screw	**100**	**100**	54.9	84.5	81.2	85.8	93.9
toothbrush	70.0	95.0	95.3	**100**	95.8	96.1	**100**
wood	92.0	97.9	97.7	93.0	97.6	**99.2**	99.1
avg	78.2	87.3	87.7	91.7	94.4	95.5	**98.0**

DIVID
2018-2021

CompVis
2019-2024

ICCV 2021

	Methods	AUROC	TPR	TNR	CA
Unsup.	RIAD [31]	78.6	79.2	69.1	70.4
	US [4]	72.5	72.6	65.3	66.2
	MAD [20]	82.4	78.7	85.7	66.2
	PaDim [11]	95.0	83.3	97.5	95.7
	DRÆM	**99.0**	**96.5**	**99.4**	**98.5**
Sup.	CADN [32]	-	-	-	89.1
	Rački *et al.* [19]	99.6	99.9	99.5	-
	Lin *et al.* [15]	99.0	99.4	99.9	-
	Božič *et al.* [6]	**100**	**100**	**100**	**100**

Unsupervised learning - DSR

- Generate syntetic anomalies in the quantized feature space

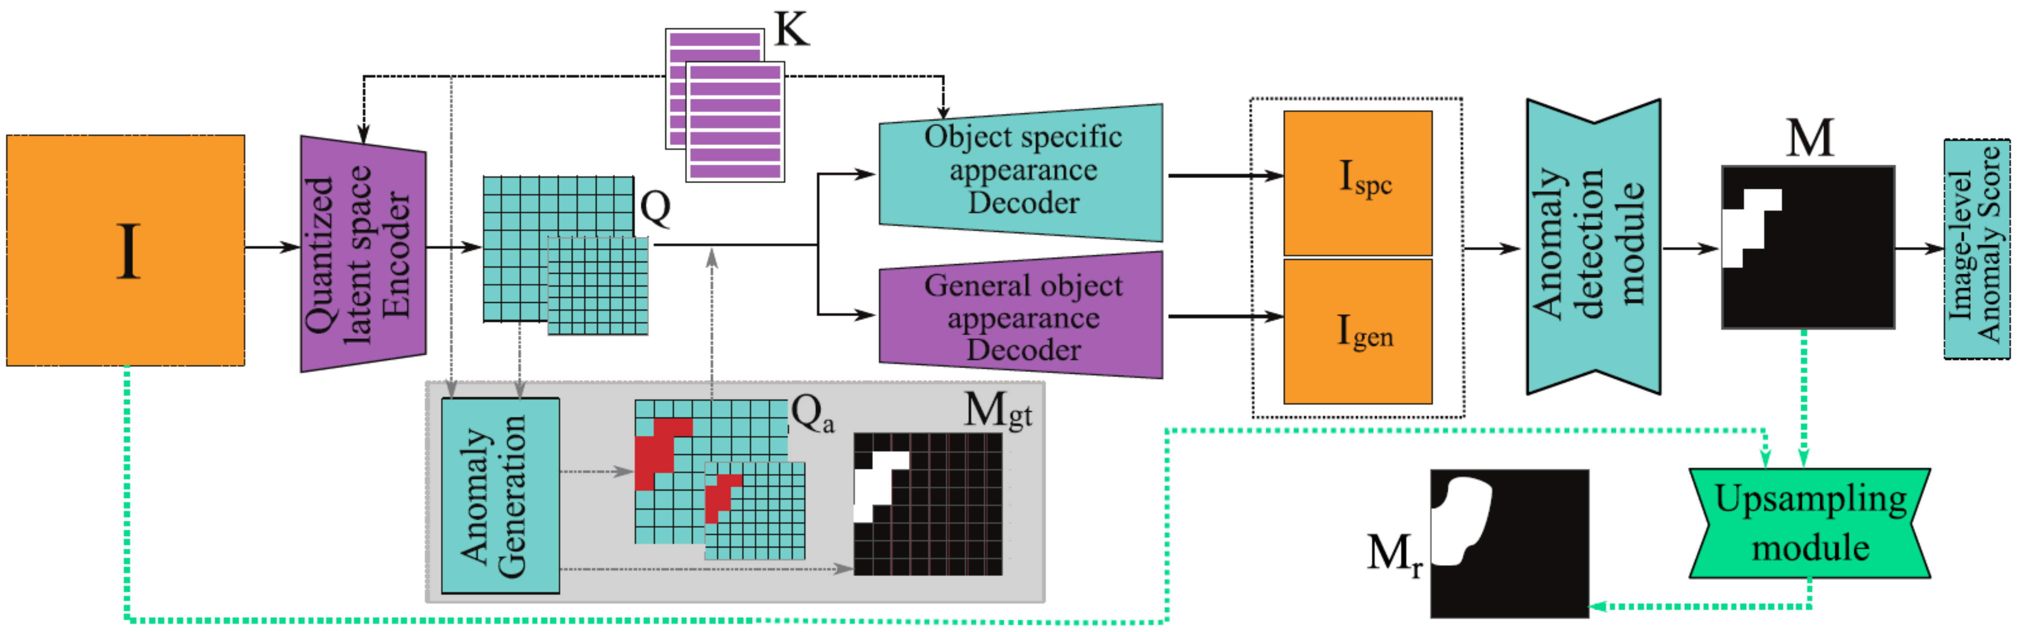

ECCV 2022

MV4.0
2021-2024

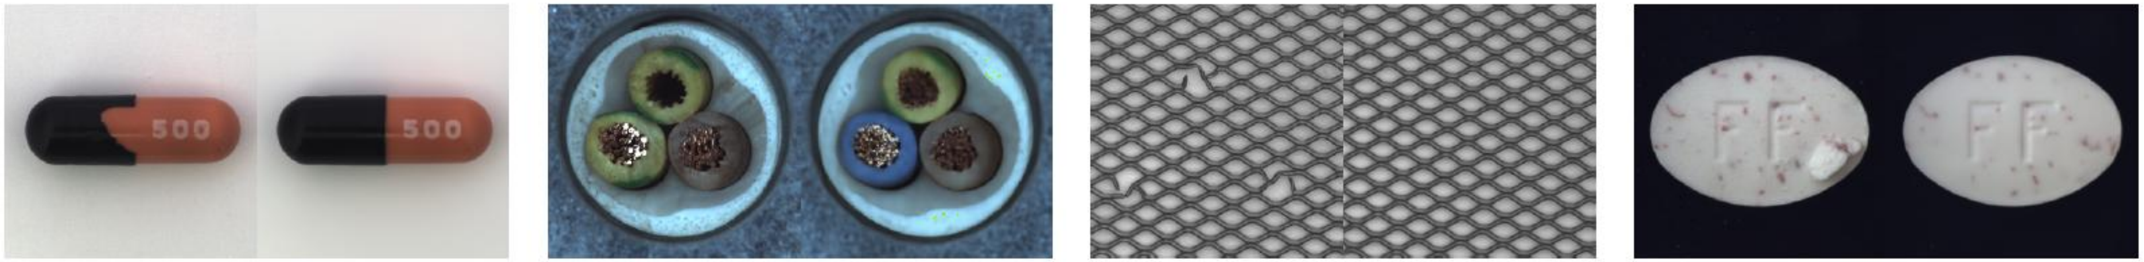

Unsupervised learning - DSR

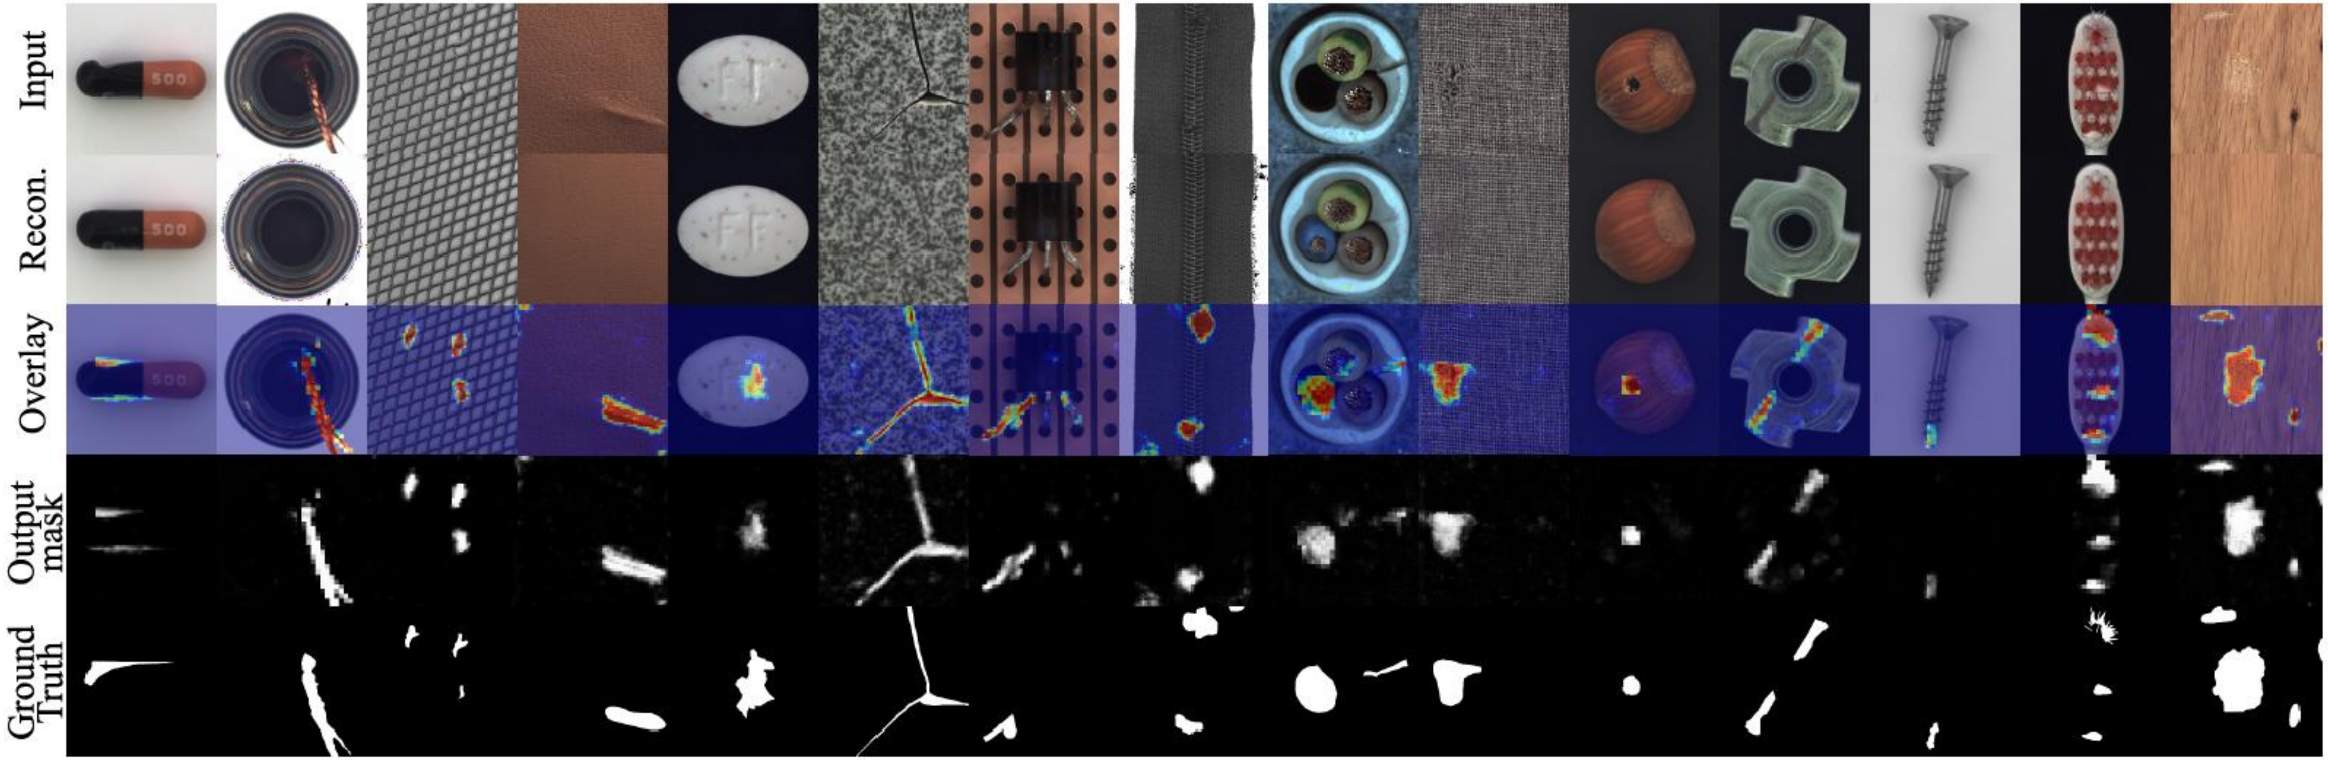

Method	bottle	capsule	grid	leather	pill	tile	trans.	zipper	cable	carpet	hazelnut	m. nut	screw	toothbrush	wood	average
[4]	99.0	86.1	81.0	88.2	87.9	99.1	81.8	91.9	86.2	91.6	93.1	82.0	54.9	95.3	97.7	87.7
[22]	99.9	88.4	99.6	**100**	83.8	98.7	90.9	98.1	81.9	84.2	83.3	88.5	84.5	**100**	93.0	91.7
[17]	**100**	92.3	92.9	**100**	83.3	97.4	95.9	97.9	**94.0**	95.5	98.7	93.1	81.2	95.8	97.6	94.4
[7]	99.8	91.5	95.7	**100**	94.4	97.4	**97.8**	90.9	92.2	99.9	93.3	99.2	84.4	97.2	98.8	95.5
[11]	98.2	98.2	**100**	**100**	94.9	94.6	96.1	99.9	81.2	93.9	98.3	**99.9**	88.7	99.4	**99.1**	96.1
[21]	99.2	**98.5**	99.9	**100**	**98.9**	99.6	93.1	**100**	91.8	97.0	**100**	98.7	93.9	**100**	**99.1**	98.0
DSR	**100**	98.1	**100**	**100**	97.5	**100**	**97.8**	**100**	93.8	**100**	95.6	98.5	**96.2**	99.7	96.3	**98.2**

Unsupervised learning - DSR

- Results on KSDD2

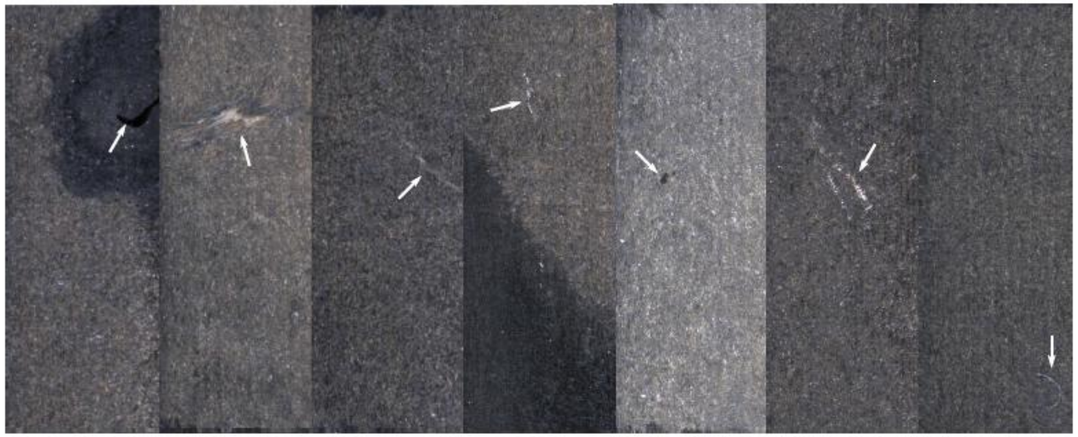

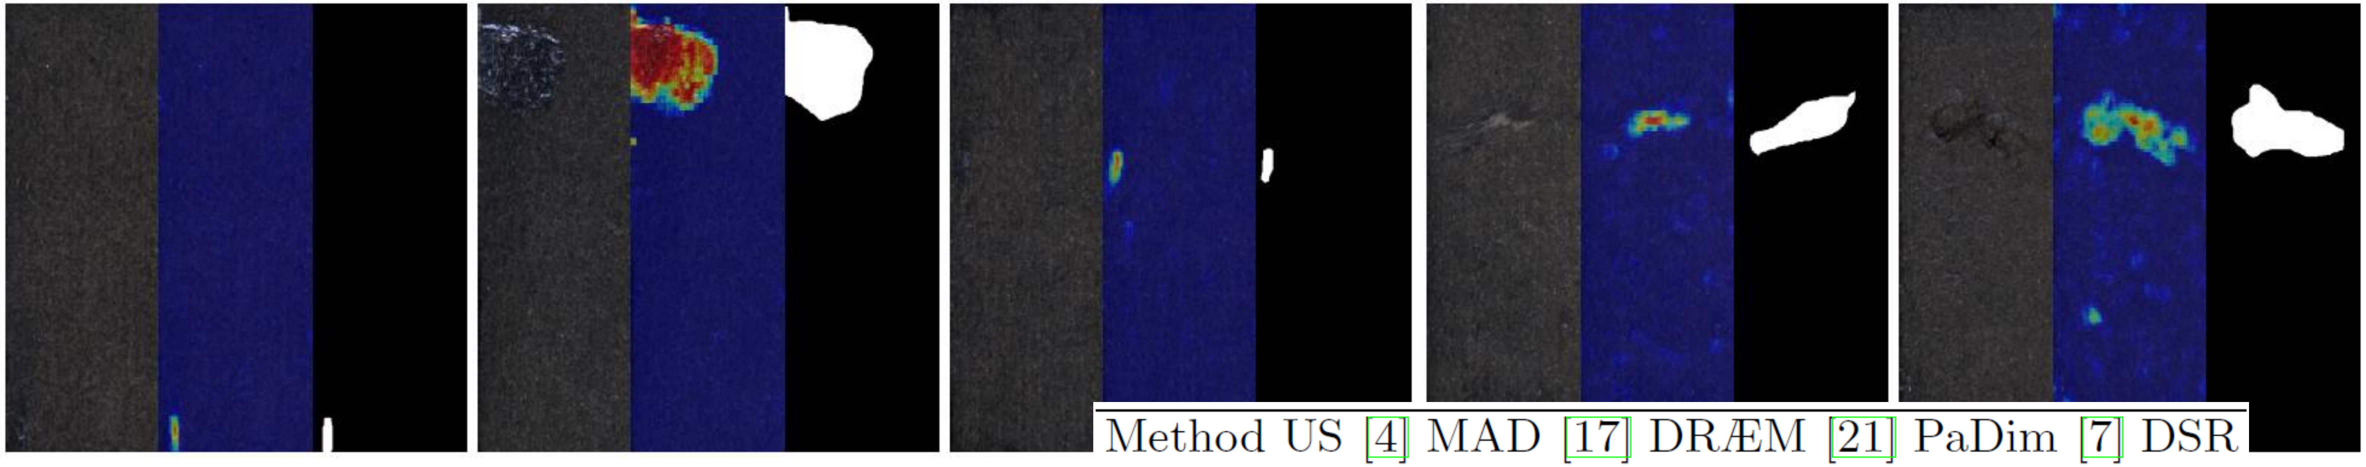

Method	US [4]	MAD [17]	DRÆM [21]	PaDim [7]	DSR
AP_{det}	65.3	79.3	77.8	55.6	**87.2**
AP_{loc}	-	-	42.4	45.3	**61.4**

Unsupervised learning - Transfusion

- TRANSparent difFUSION
- Using Diffusion model estimate
  - Anomaly mask
  - Anomaly
  - Normal image

MV4.0
2021-2024

ECCV 2024

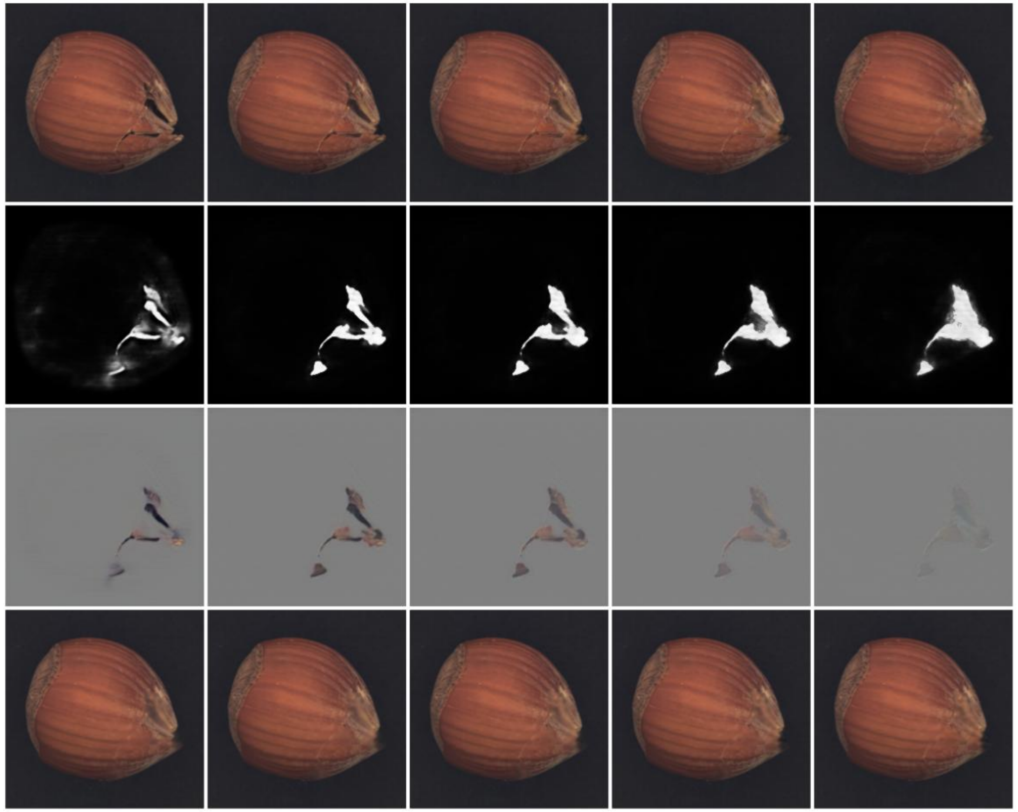

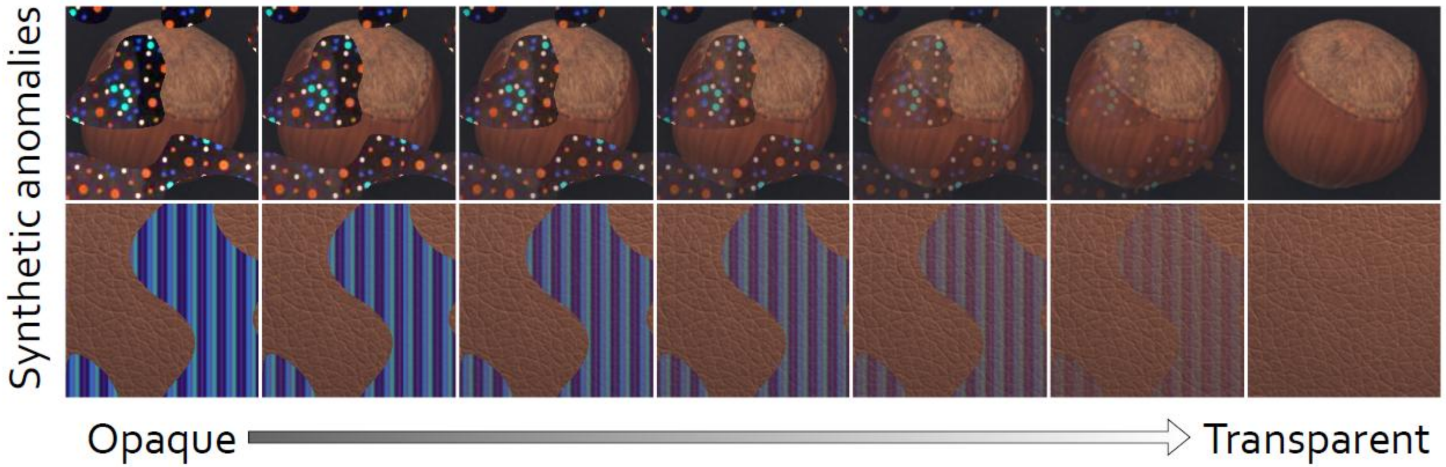

$$x_{t-1} = x_t - (\beta_t - \beta_{t-1})(M_t \odot \epsilon_t) + (\beta_t - \beta_{t-1})(M_t \odot \hat{x}_0{}^{(t)})$$

Unsupervised learning - Transfusion

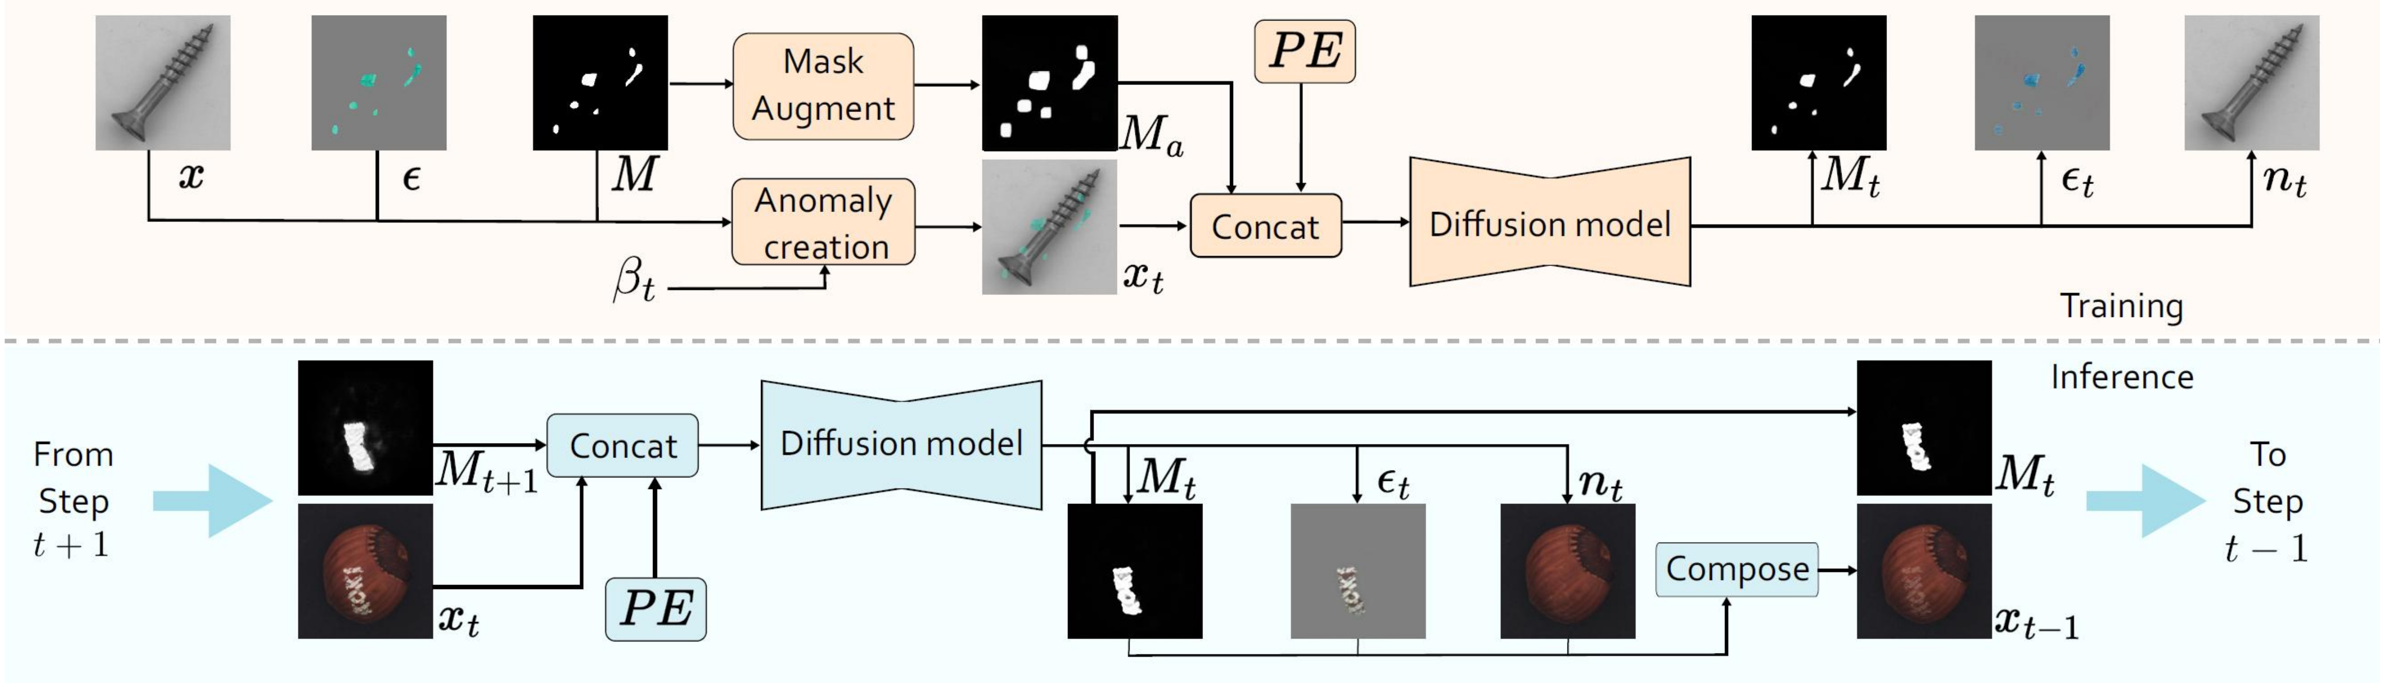

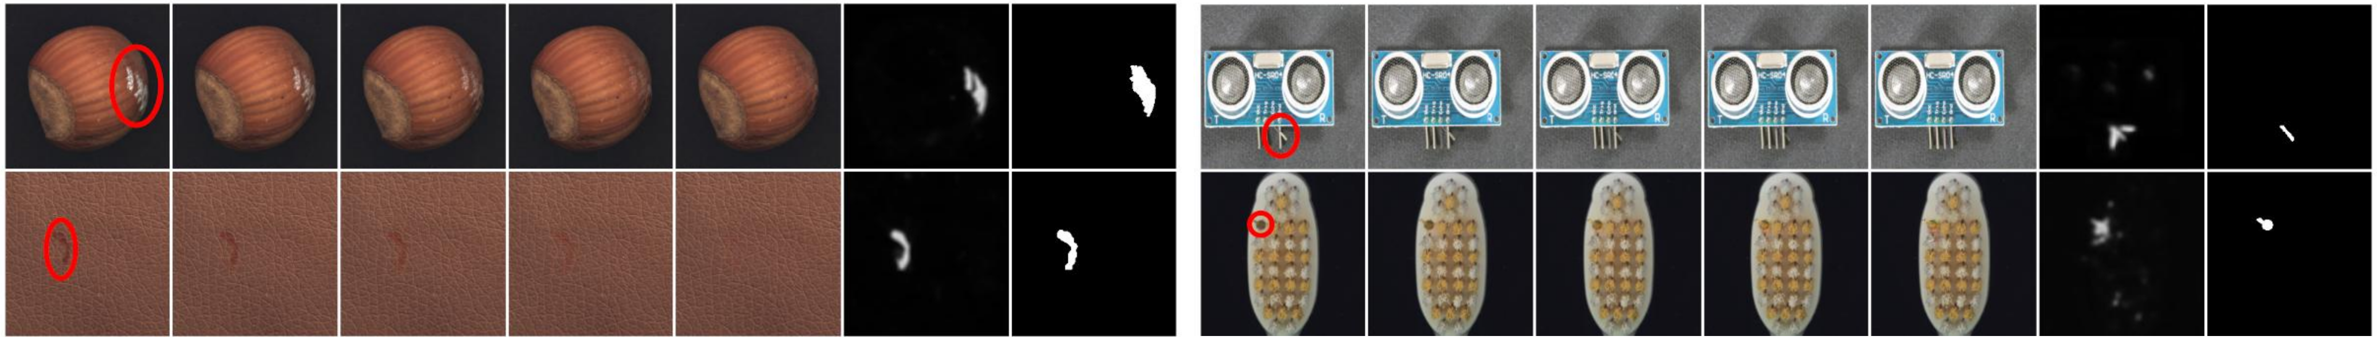

Unsupervised learning - Transfusion

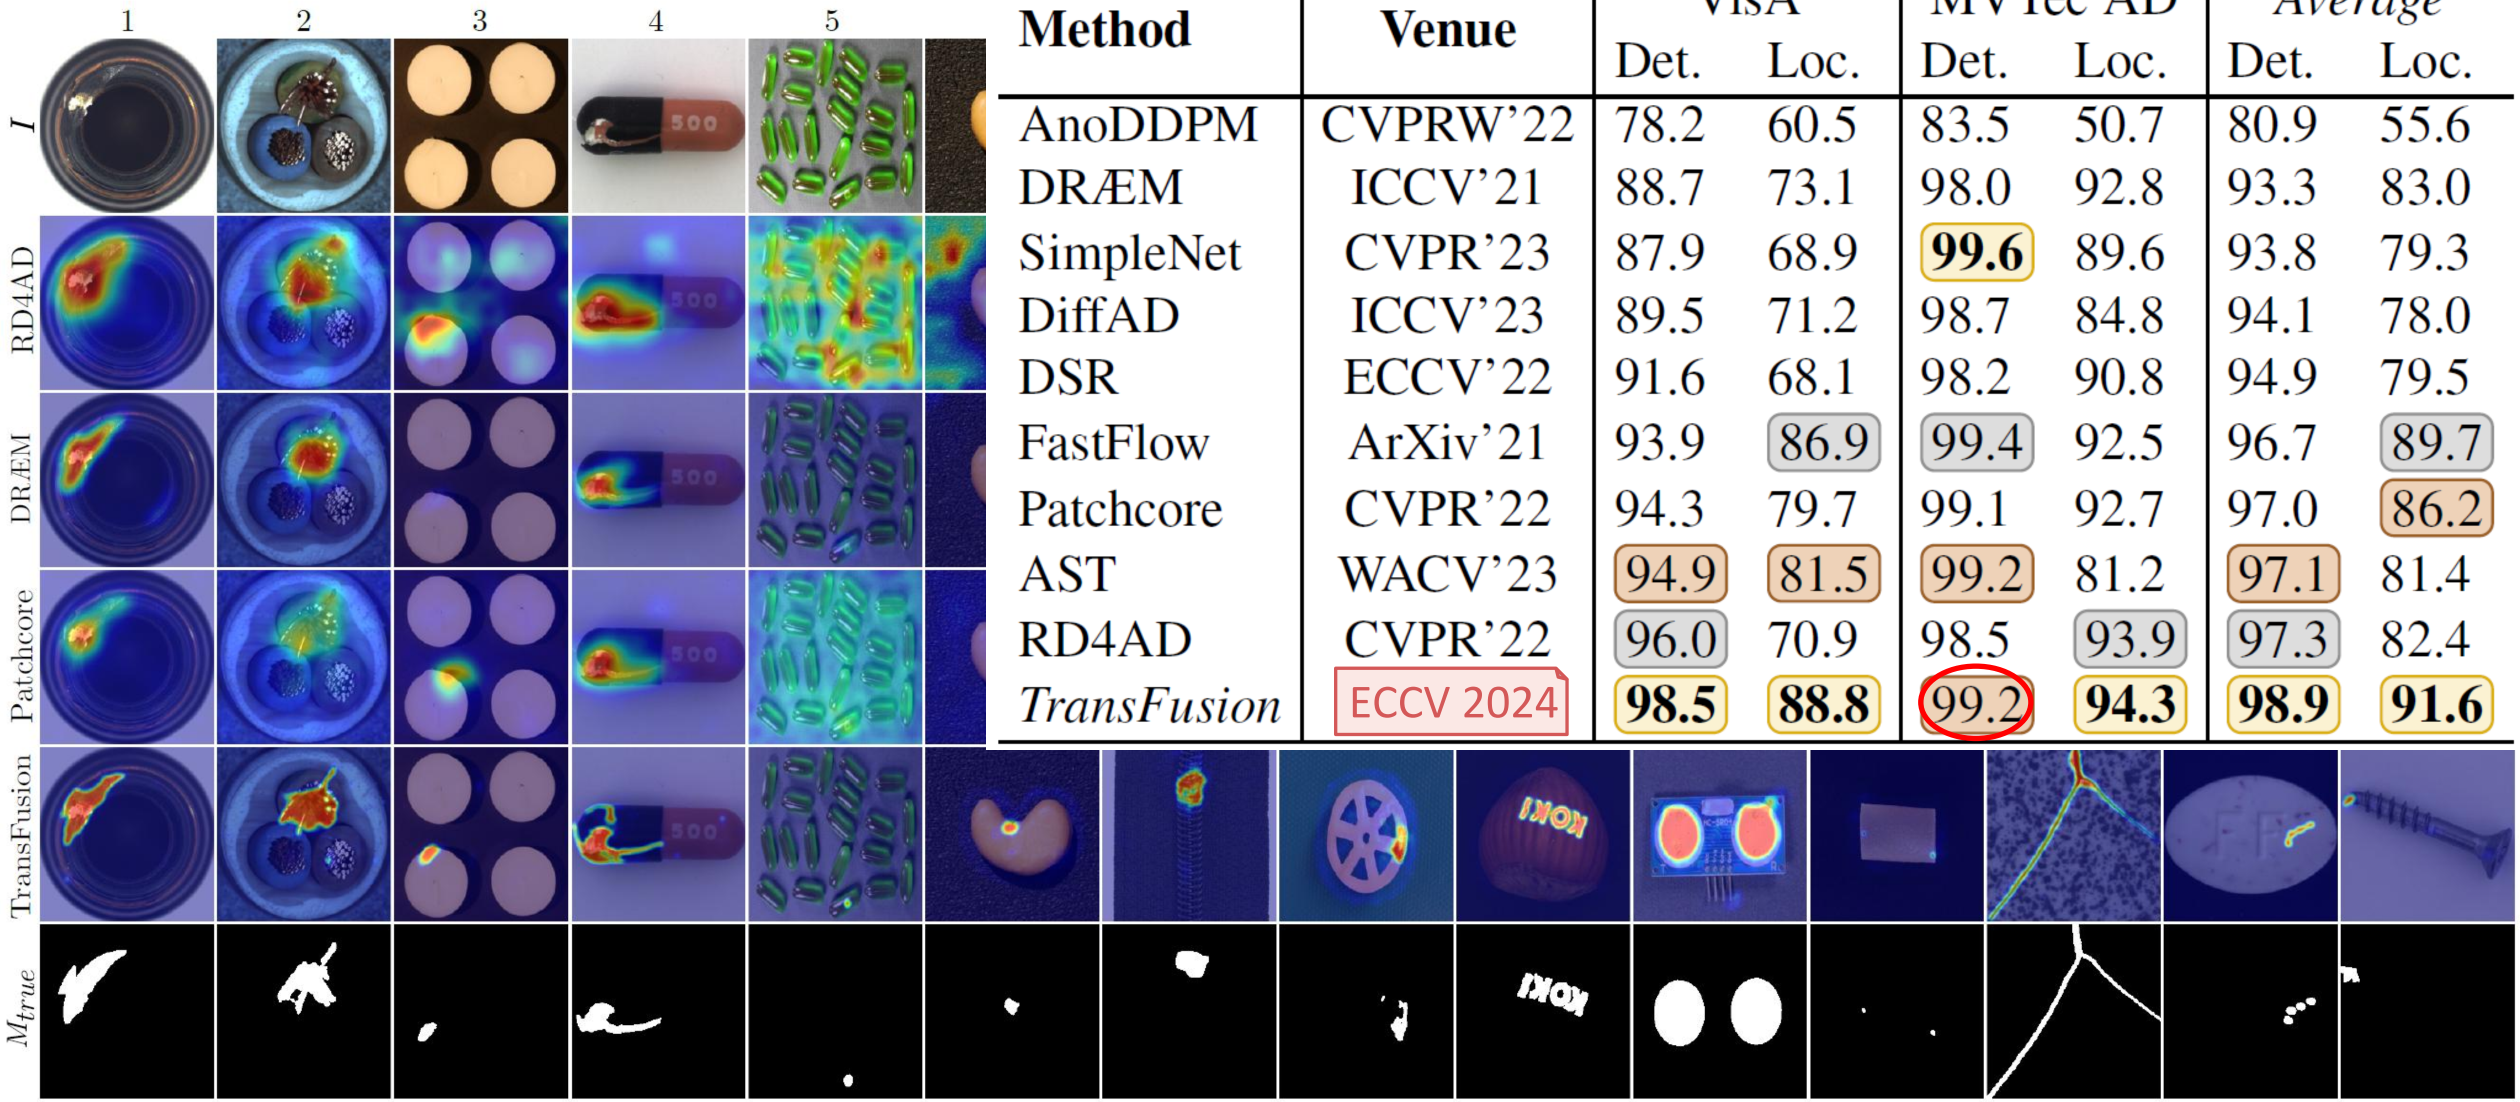

Method	Venue	VisA		MVTec AD		*Average*	
		Det.	Loc.	Det.	Loc.	Det.	Loc.
AnoDDPM	CVPRW'22	78.2	60.5	83.5	50.7	80.9	55.6
DRÆM	ICCV'21	88.7	73.1	98.0	92.8	93.3	83.0
SimpleNet	CVPR'23	87.9	68.9	**99.6**	89.6	93.8	79.3
DiffAD	ICCV'23	89.5	71.2	98.7	84.8	94.1	78.0
DSR	ECCV'22	91.6	68.1	98.2	90.8	94.9	79.5
FastFlow	ArXiv'21	93.9	86.9	99.4	92.5	96.7	89.7
Patchcore	CVPR'22	94.3	79.7	99.1	92.7	97.0	86.2
AST	WACV'23	94.9	81.5	99.2	81.2	97.1	81.4
RD4AD	CVPR'22	96.0	70.9	98.5	93.9	97.3	82.4
TransFusion	ECCV 2024	**98.5**	**88.8**	99.2	**94.3**	**98.9**	**91.6**

Learning regimes

Mixed supervision - from fully supervised to weakly supervised learning

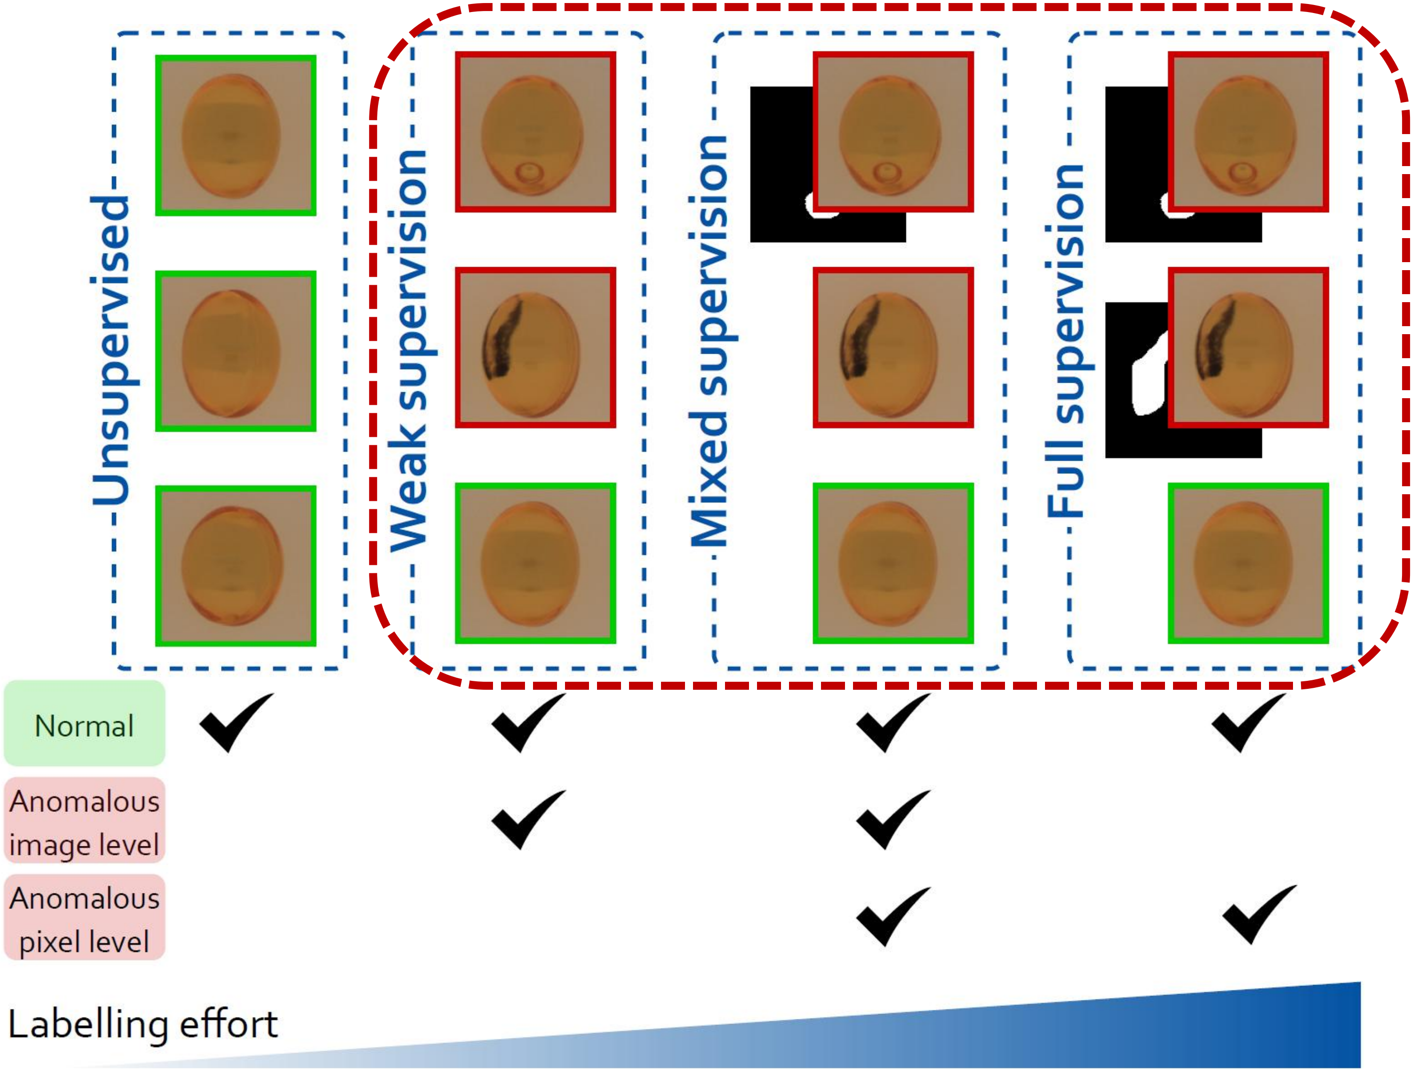

Learning with mixed supervision

WS - Weakly Supervised Learning
MS - Mixed Supervision
FS - Fully Supervised Learning

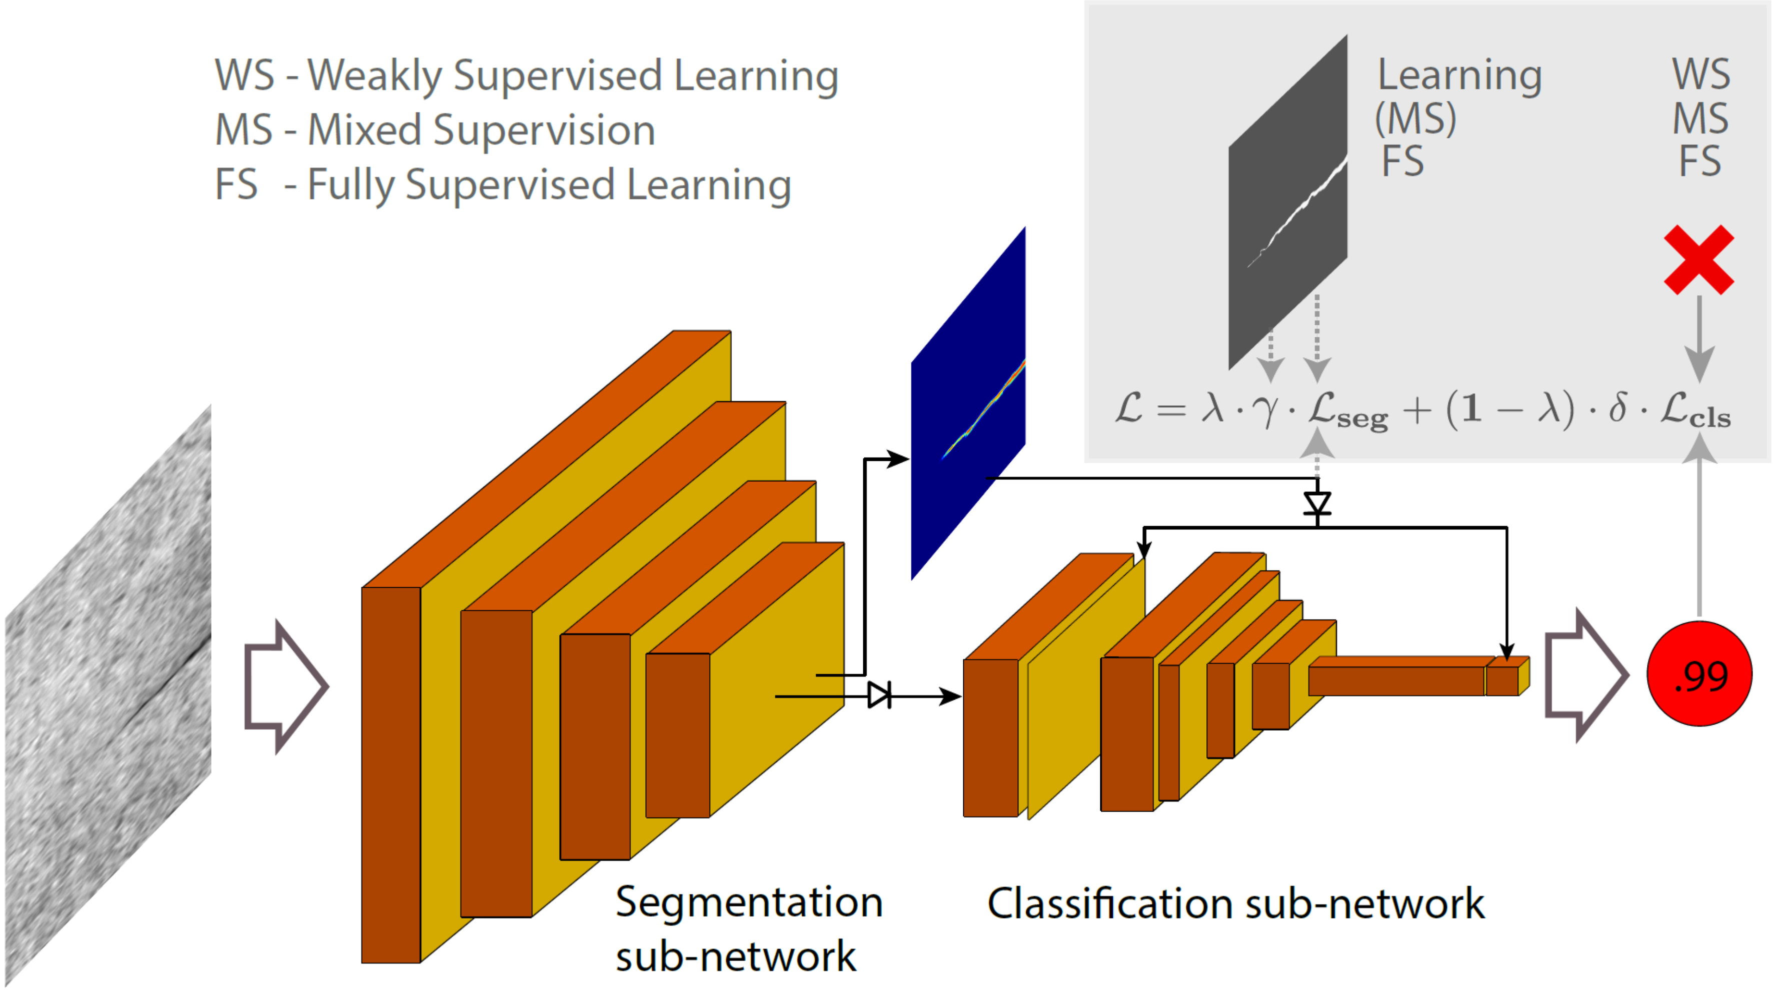

DIVID
2018-2021

COMIND 2021

From supervised to weakly supervised learning

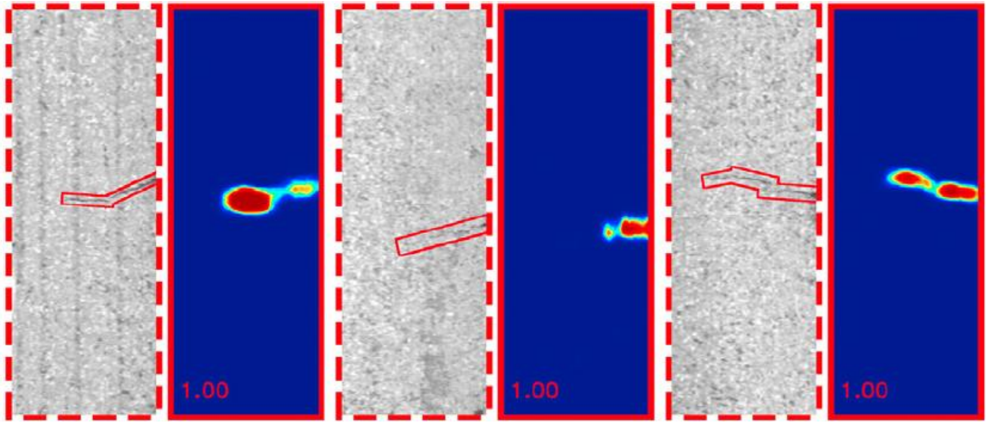

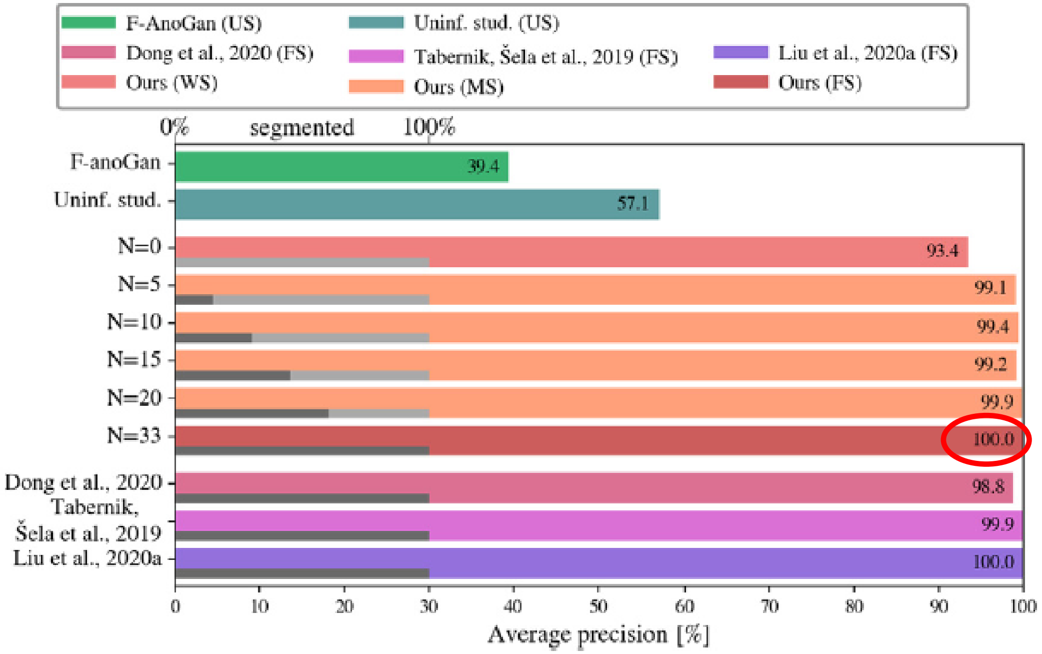

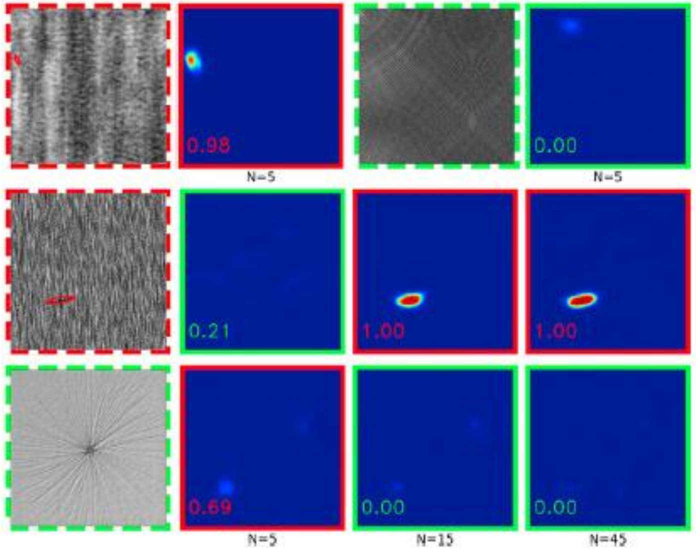

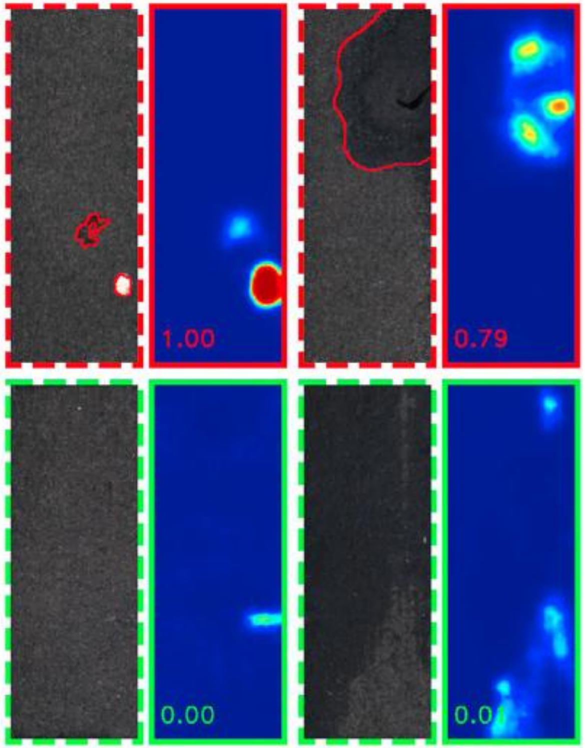

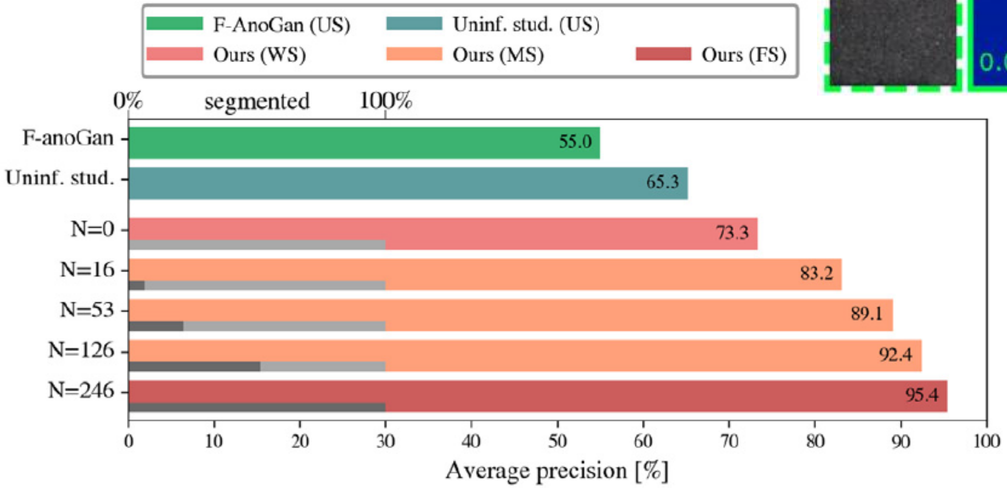

DIVID
2018-2021

COMIND 2021

Learning regimes

Mixed supervision - from unsupervised to supervised learning

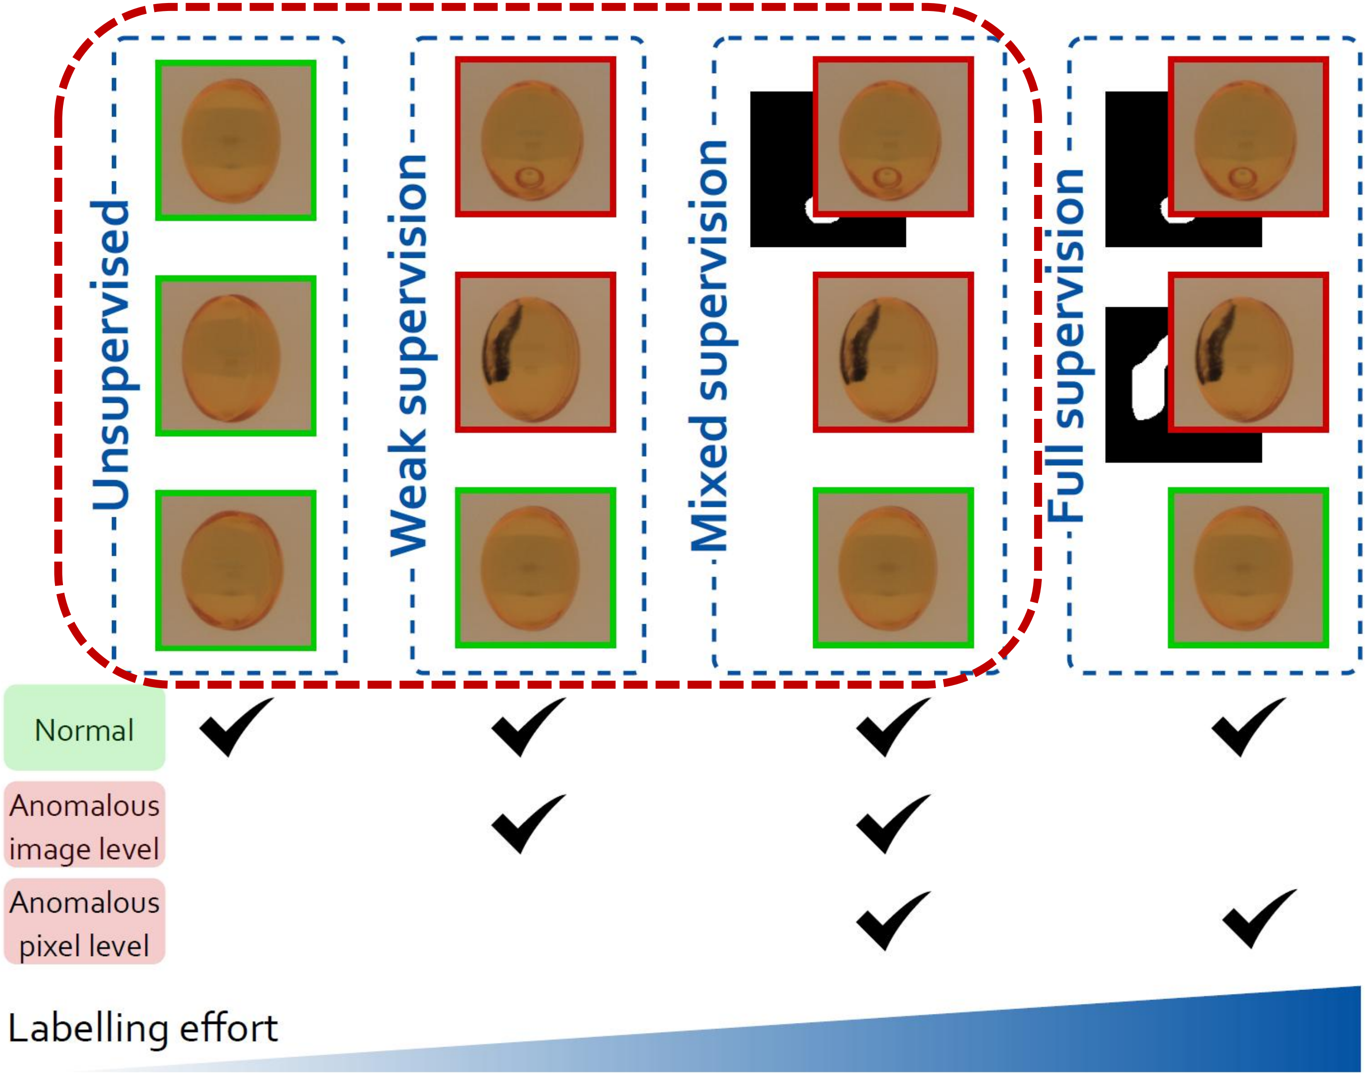

Learning with mixed supervision

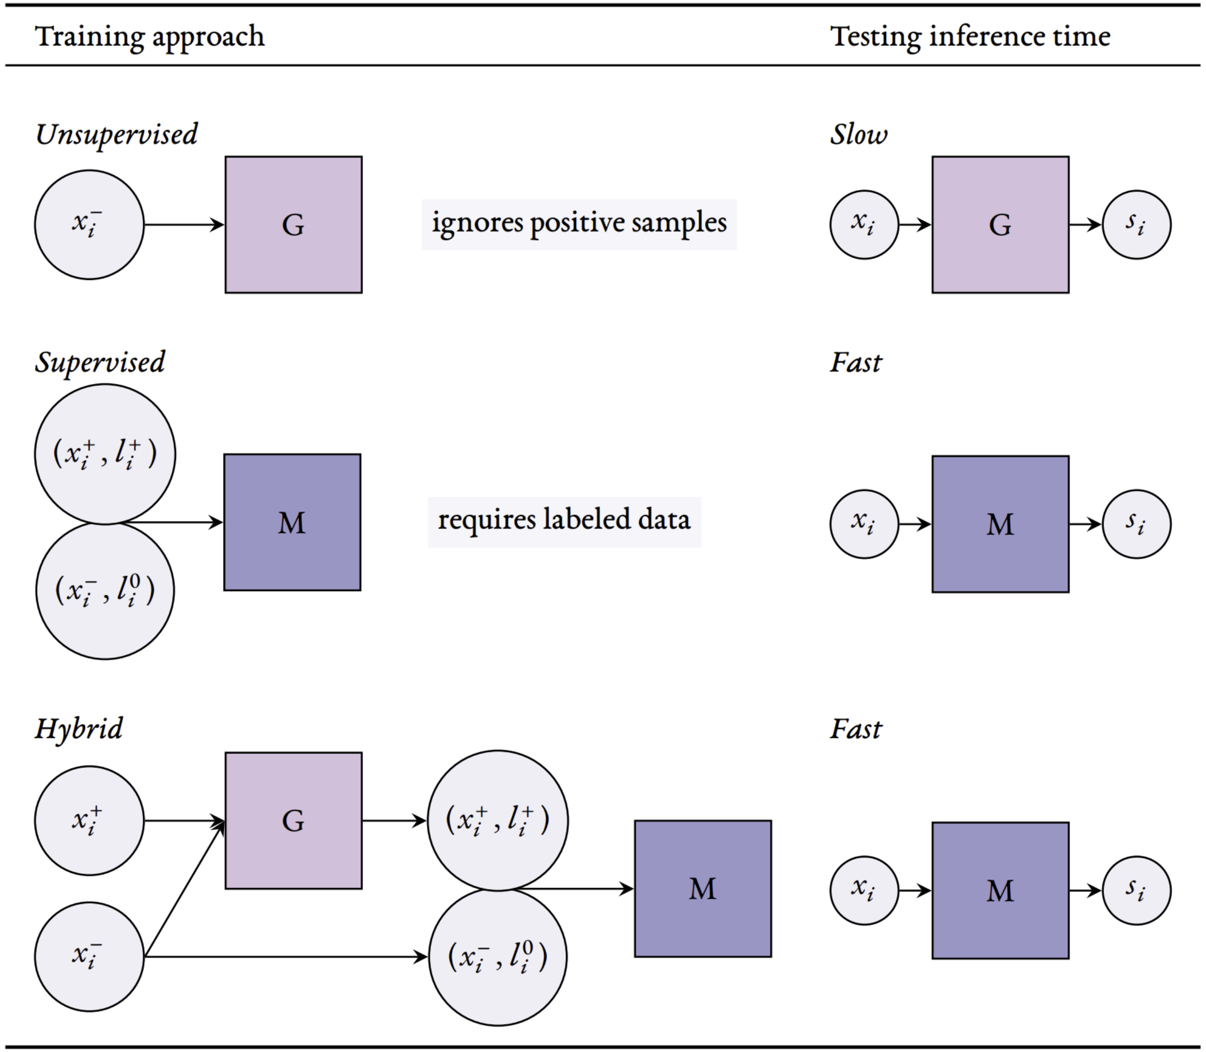

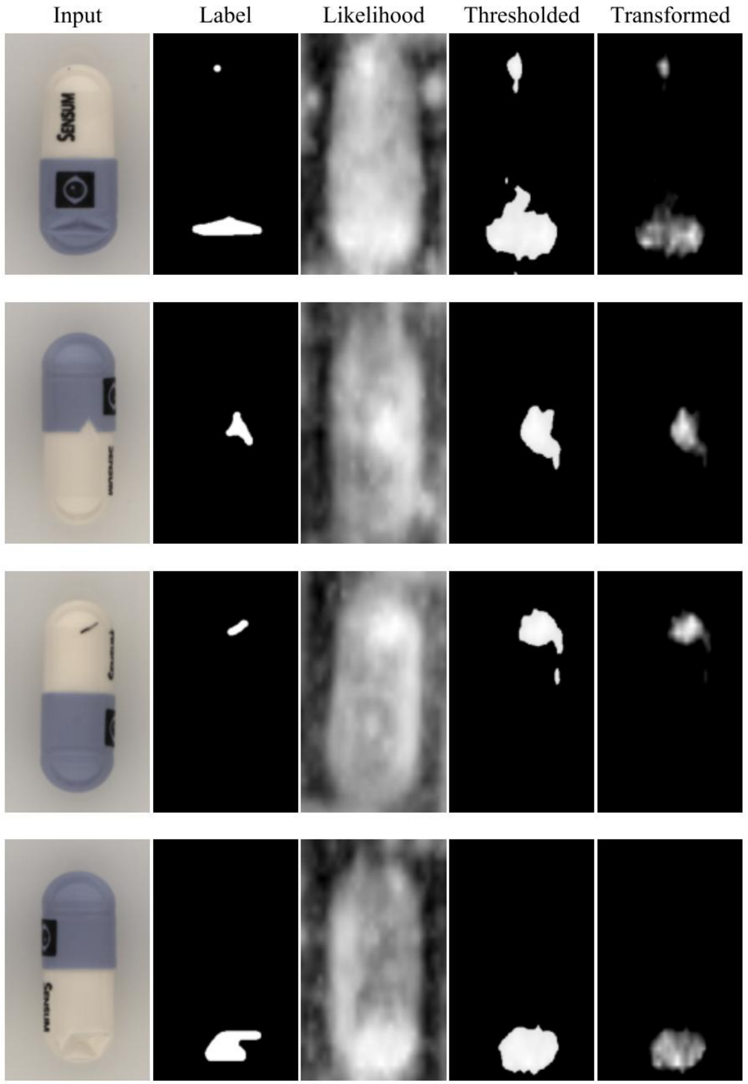

QCAV 2023

Learning regimes

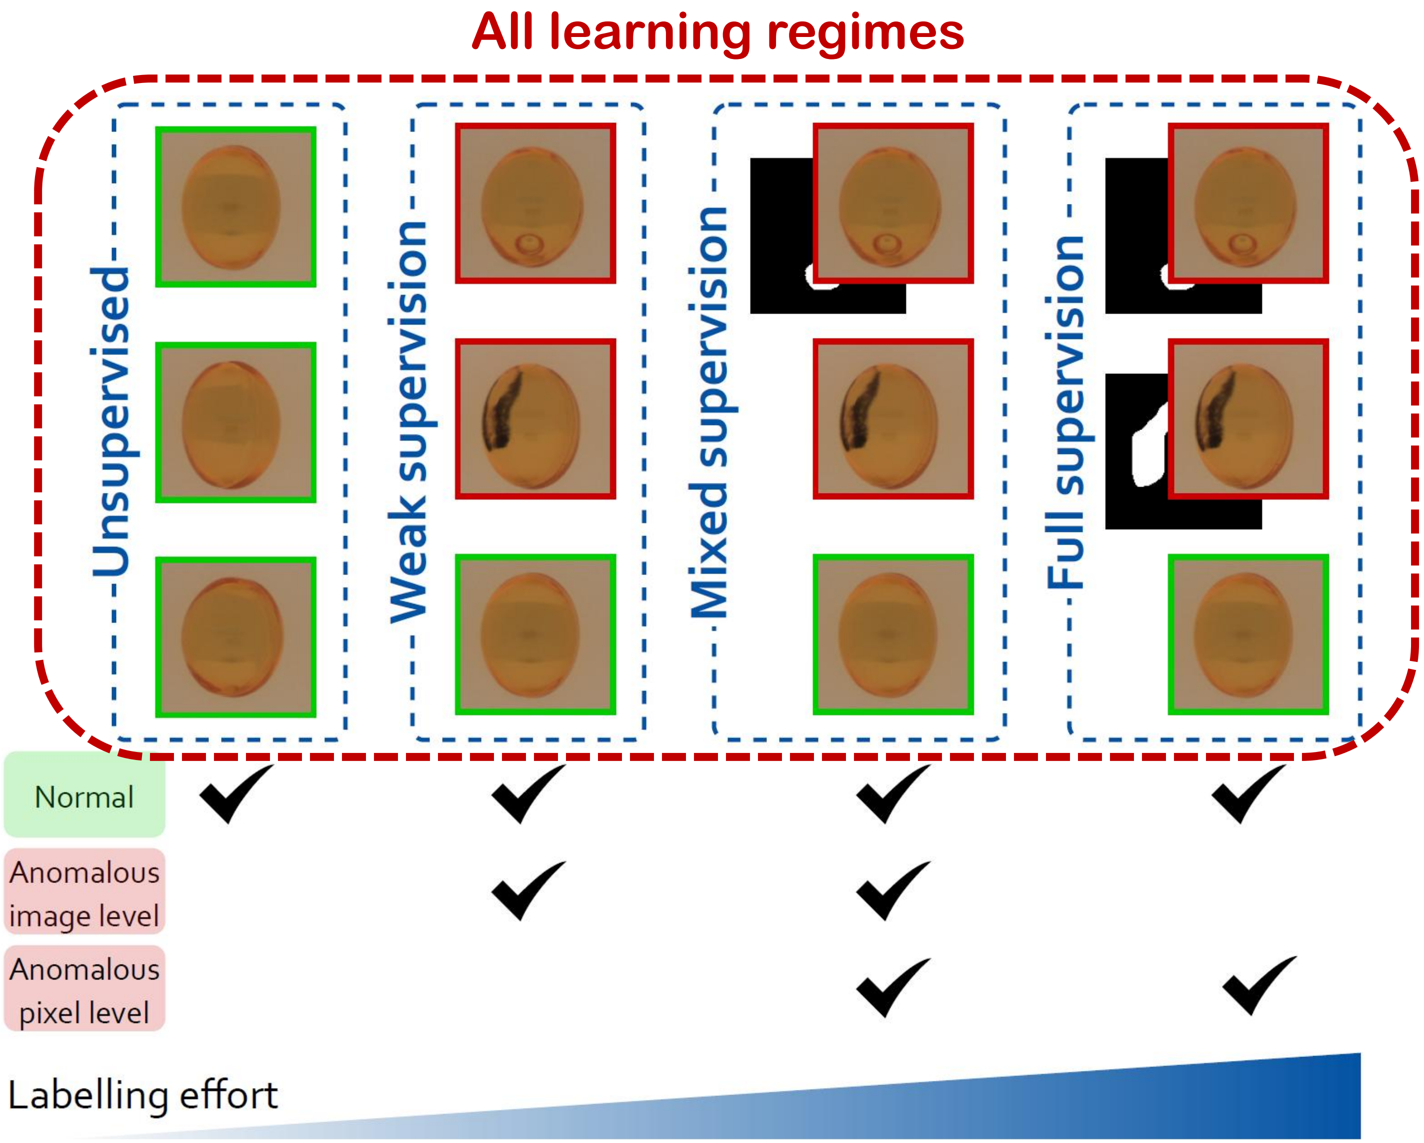

SuperSimpleNet

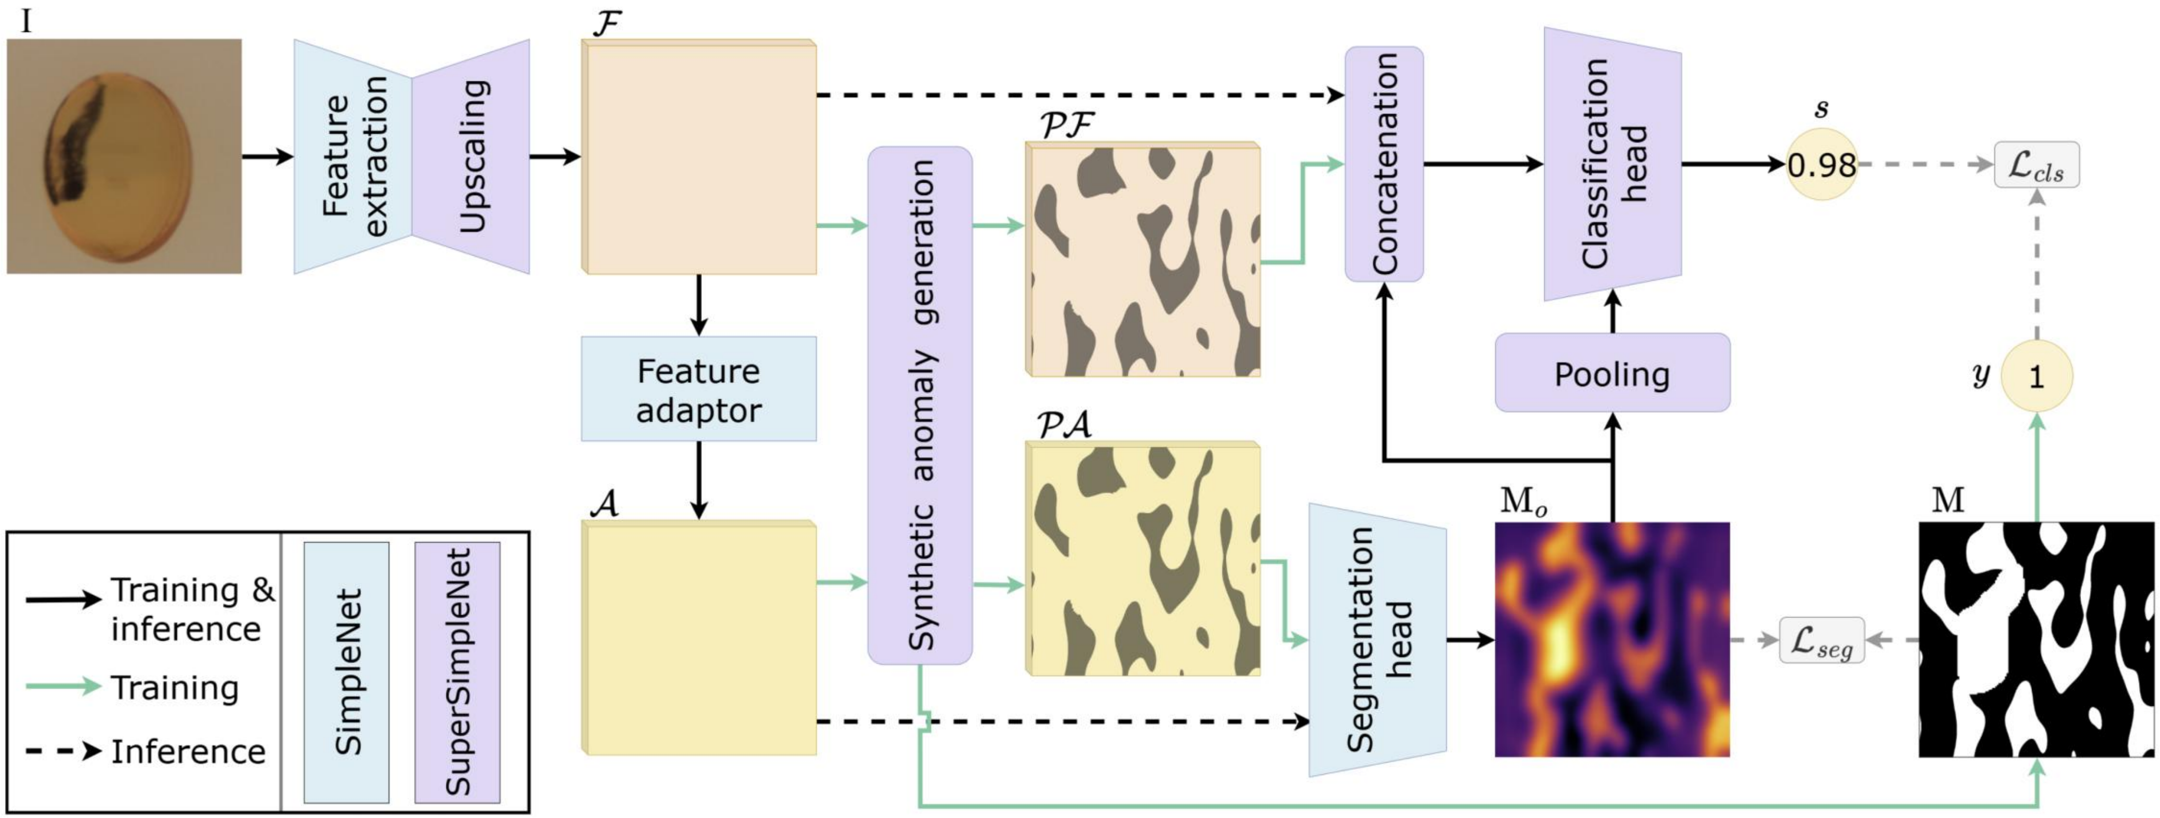

[3] Z. Liu, Y. Zhou, Y. Xu, Z. Wang, SimpleNet: A Simple Network for Image Anomaly Detection and Localization, CVPR 2023.
[4] B. Rolih, M. Fučka D. Skočaj, No Label Left Behind: A Unified Surface Defect Detection Model for all Supervision Regimes, JIM, 2024.

MV4.0
2021-2024

JIM 2025

Anomaly generation

- True & synthetic anomalies
- Perlin noise
- Gaussian noise
- Label
  - Image level
  - Pixel level
- Weakly supervised
- Unsupervised

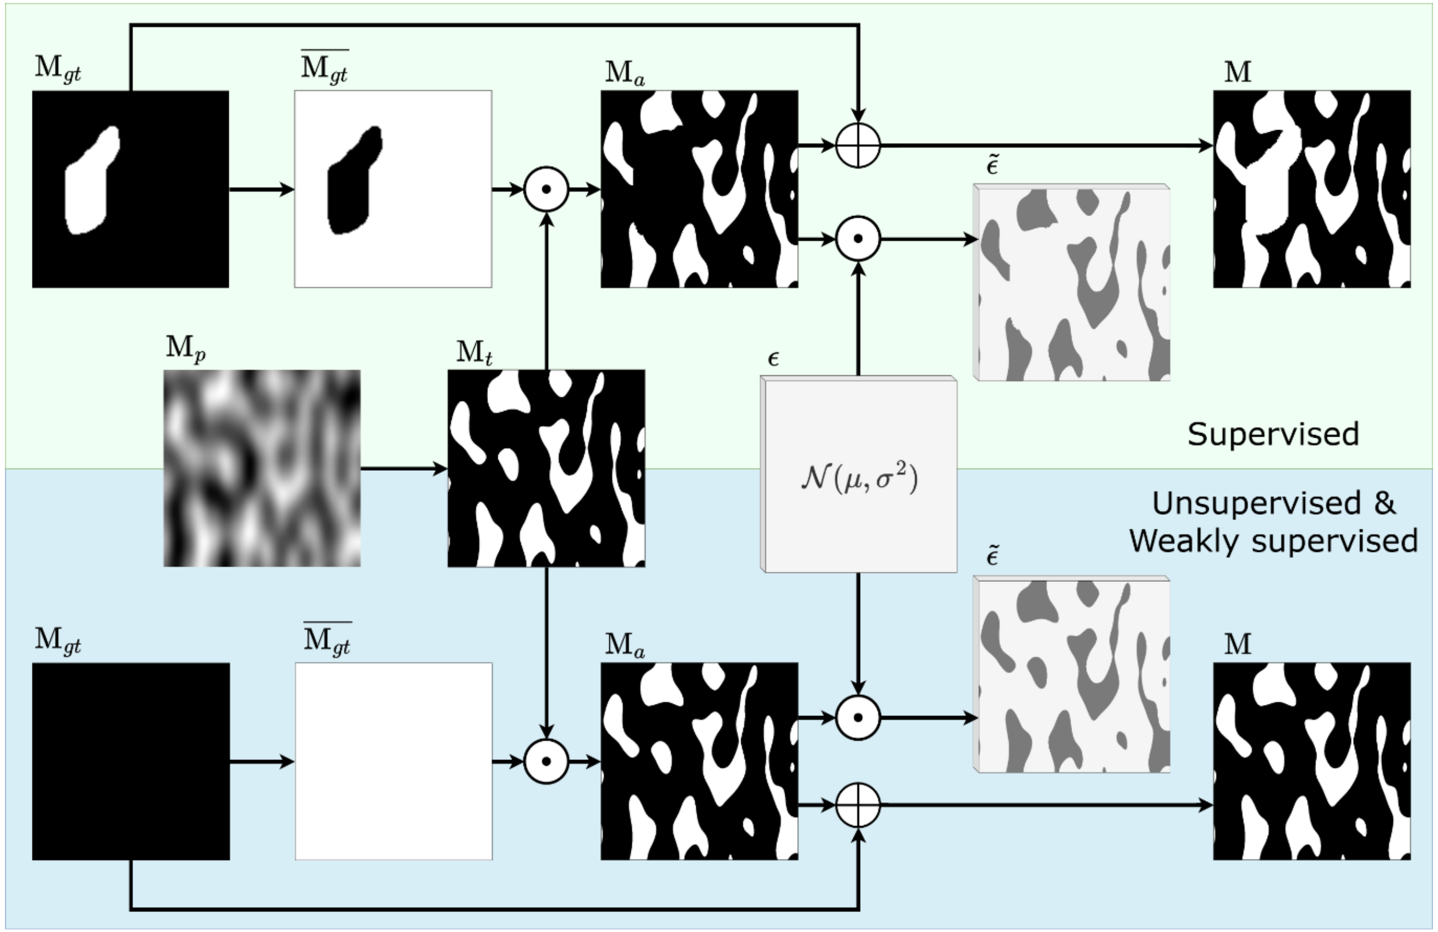

Segmentation-detection module

- Segmentation head
- Classification head
  - Capturing global context
  - Detection of small defects
  - Enables mixed supervision

$$\mathcal{L}_{seg} = \mathcal{L}_{1t} + \mathcal{L}_{foc} \qquad \mathcal{L}_{cls} = \mathcal{L}_{foc}$$

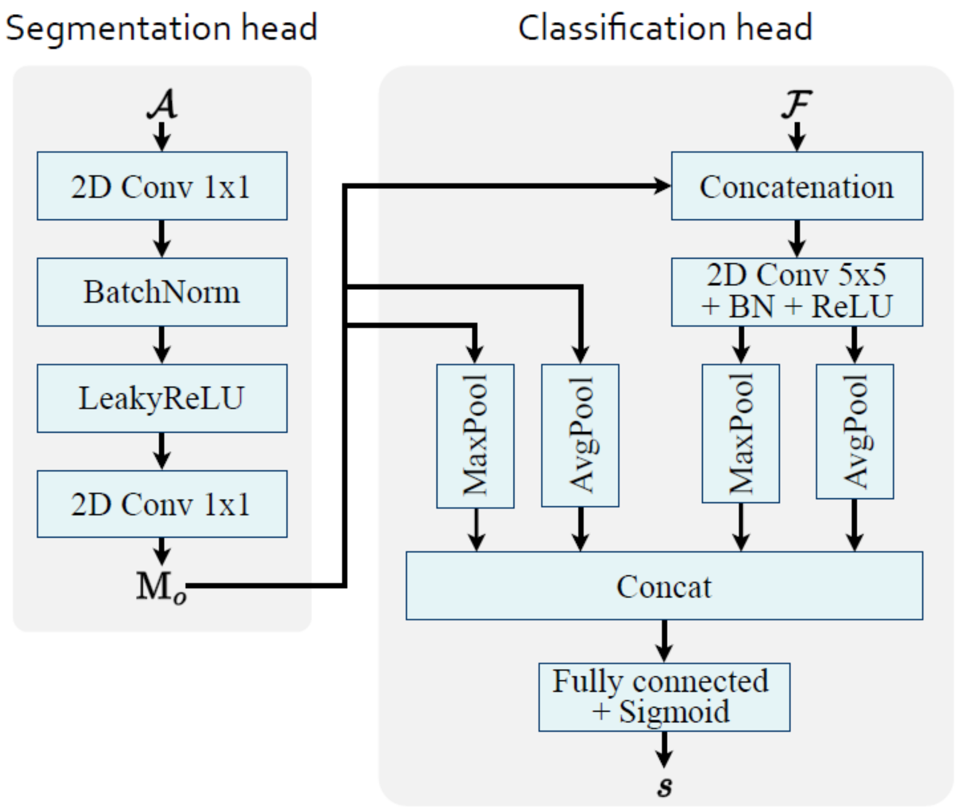

$$\mathcal{L} = \gamma \cdot \mathcal{L}_{seg} + \mathcal{L}_{cls}$$

$$\gamma = \begin{cases} 1; \text{ if image is normal;} \\ 1; \text{ if image is anomalous and fully labelled;} \\ 0; \text{ if image is anomalous and weakly labelled;} \end{cases}$$

Inference

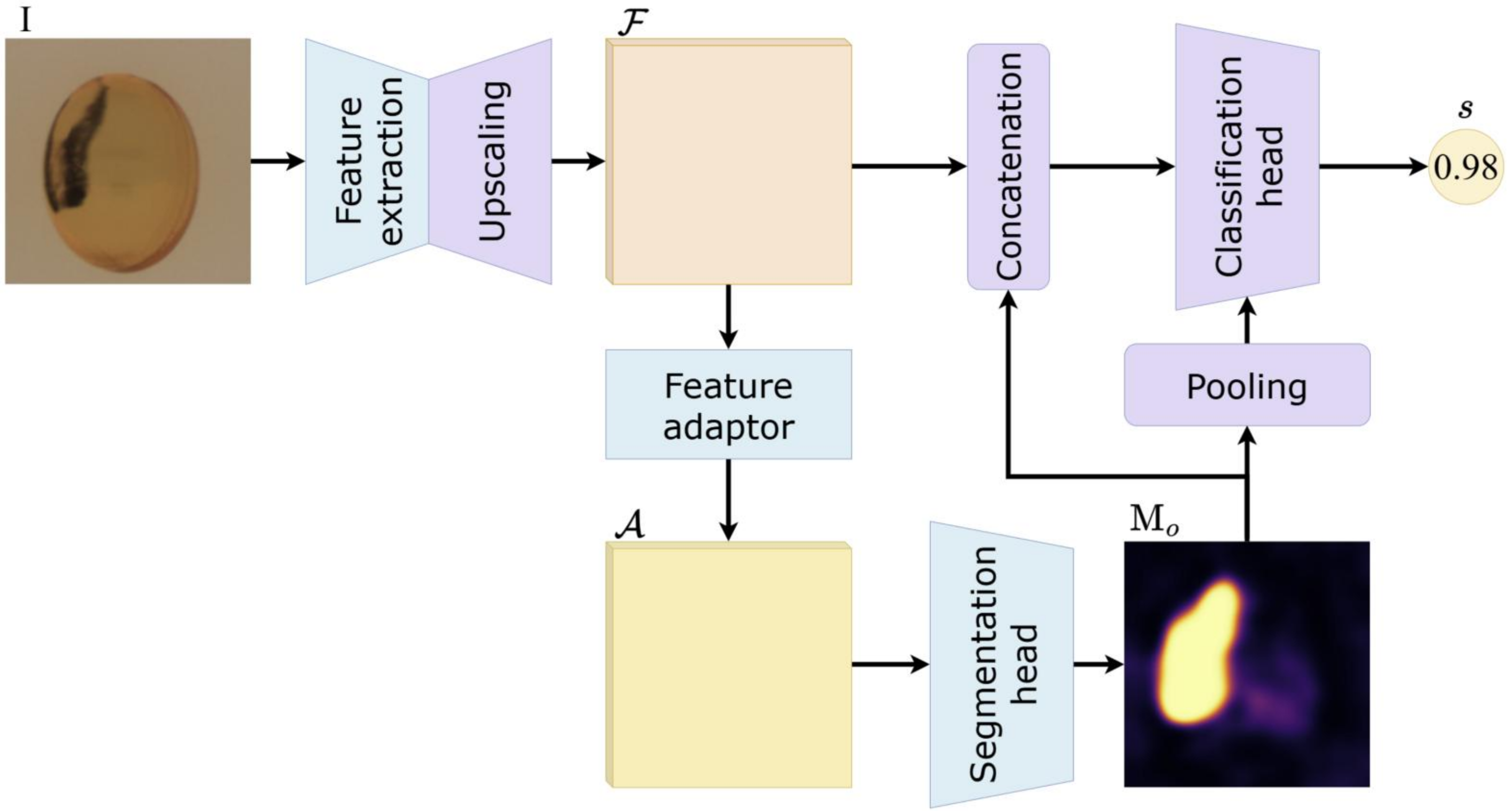

I
Feature extraction
Upscaling
$\mathcal{F}$
Concatenation
Classification head
s
0.98
Feature adaptor
Pooling
$\mathcal{A}$
Segmentation head
M_o

Datasets and performance metrics

- Supervised
  - SensumSODF
  - KolektorSDD2
- Unsupervised
  - MVTec AD
  - VisA

- Detection
  - AUC
  - AP_{det}
- Localisation
  - AUPRO
  - AP_{loc}

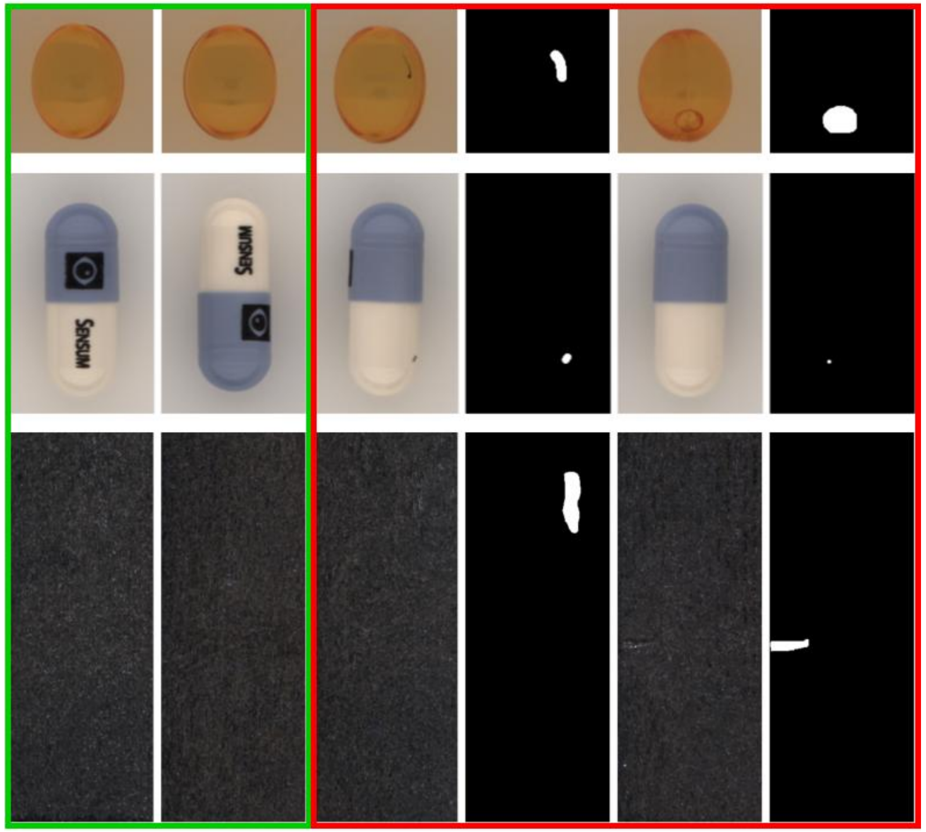

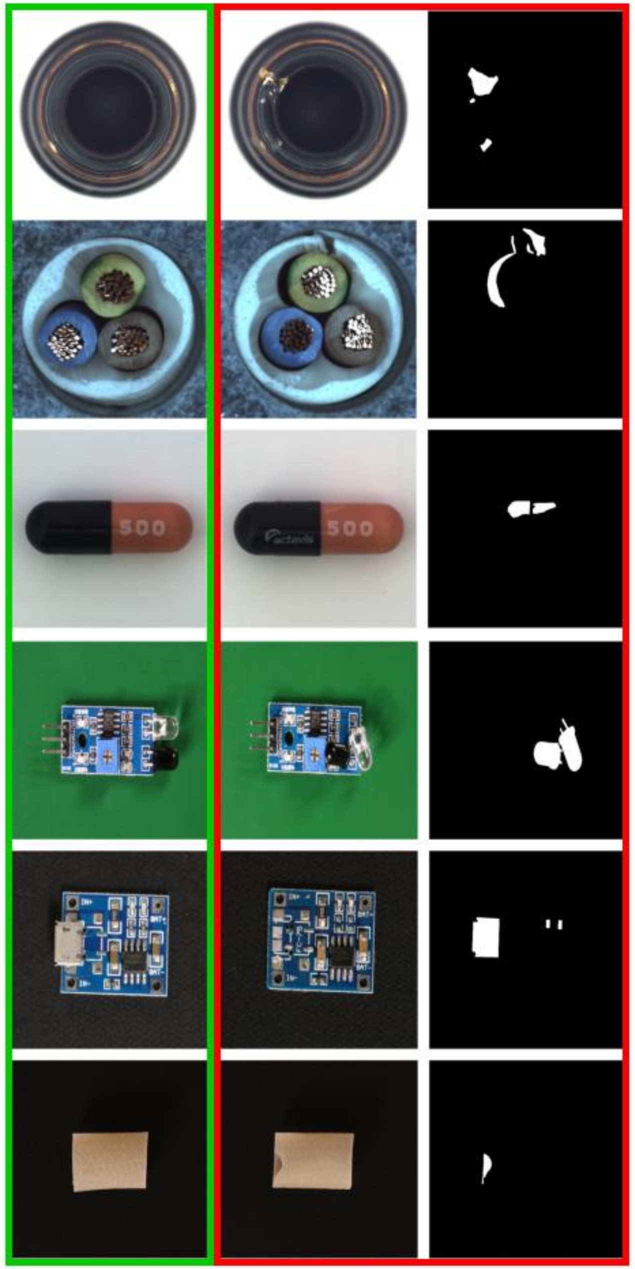

Experimental results

- Capabilities of SOTA methods:

	US	WS	MS	FS	Ours	SDNet	TNet	MMNet	DRA	EAD	BGAD	DSR	SN	FF	PC	DRÆM	PRN
normal only	✓				✓					✓	✓	✓	✓	✓	✓	✓	
ano. image level		✓	✓		✓	✓	✓	✓	✓								
ano. pixel level			✓	✓	✓	✓	✓	✓			✓	✓					✓
Speed ($< 10ms$)					✓	✓	✓	✓	✓	✓							

- Unsupervised learning results

	MVTec AD		VisA	
	Det.	Loc.	Det.	Loc.
AST (Rudolph et al., 2023)	98.9	81.2	94.9	81.5
DSR (Zavrtanik et al., 2022)	98.1	90.8	91.8	68.1
EfficientAD (Batzner et al., 2024)	99.1	93.5	98.1	94.0
FastFlow (Yu et al., 2021)	96.9	92.5	93.9	86.8
PatchCore (Roth et al., 2022)	98.7	92.7	94.3	79.7
DRÆM (Zavrtanik et al., 2021a)	98.0	92.8	91.5	78.0
SimpleNet (Liu et al., 2023)	97.6 (± 0.40)	90.5 (± 0.75)	91.2 (± 1.08)	88.0 (± 0.87)
Ours	98.3 (± 0.14)	91.2 (± 0.14)	93.6 (± 0.77)	87.4 (± 0.98)

Experimental results – mixed supervision

- Weakly supervised -> Mixed supervision -> Fully supervised
- SensumSODF:

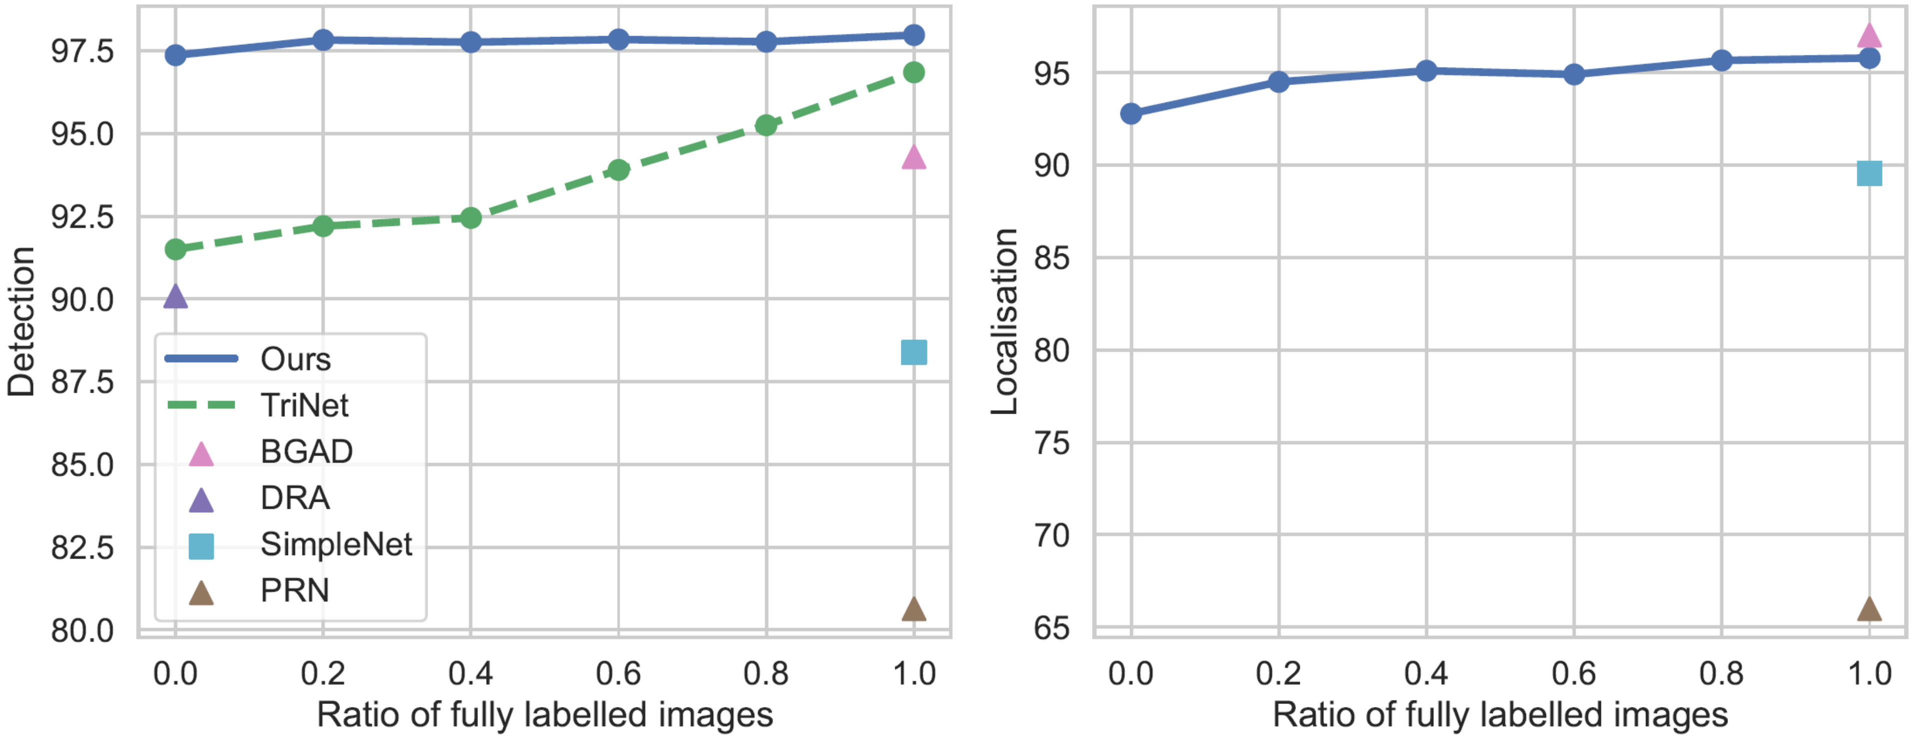

Experimental results – mixed supervision

- Weakly supervised -> Mixed supervision -> Fully supervised
- KolektorSDD2

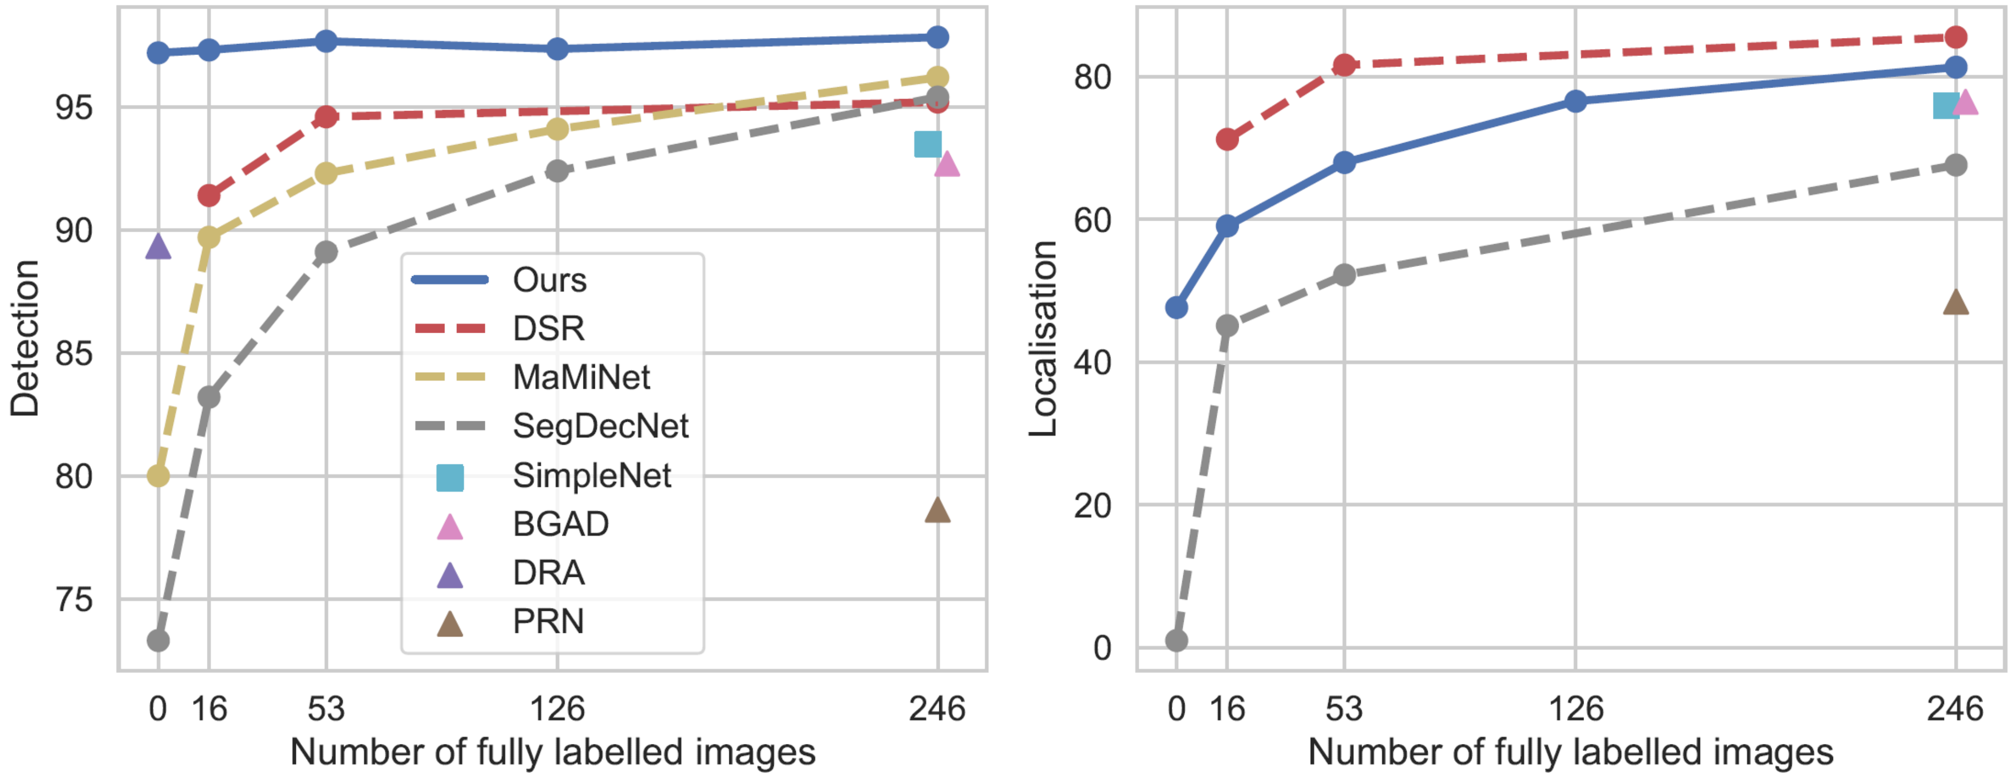

Qualitative experimental results

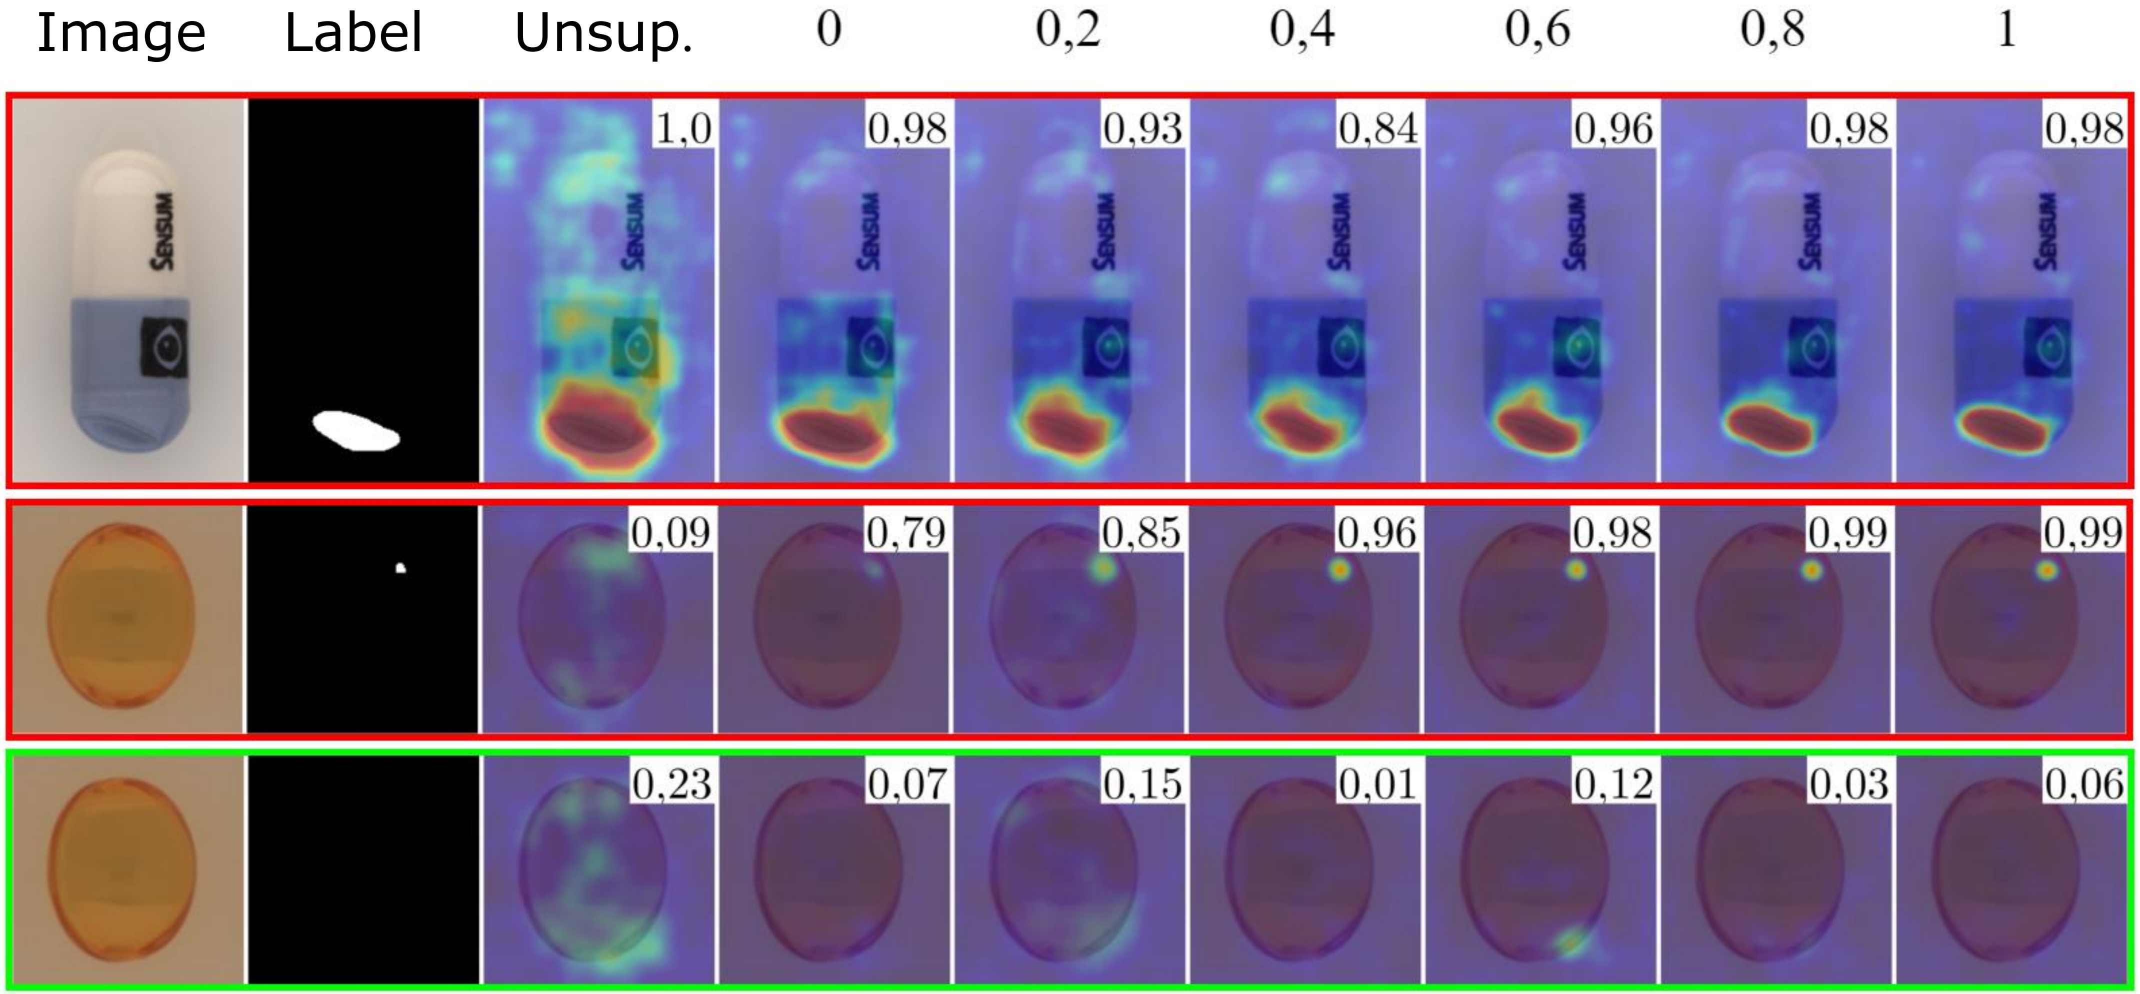

Efficiency

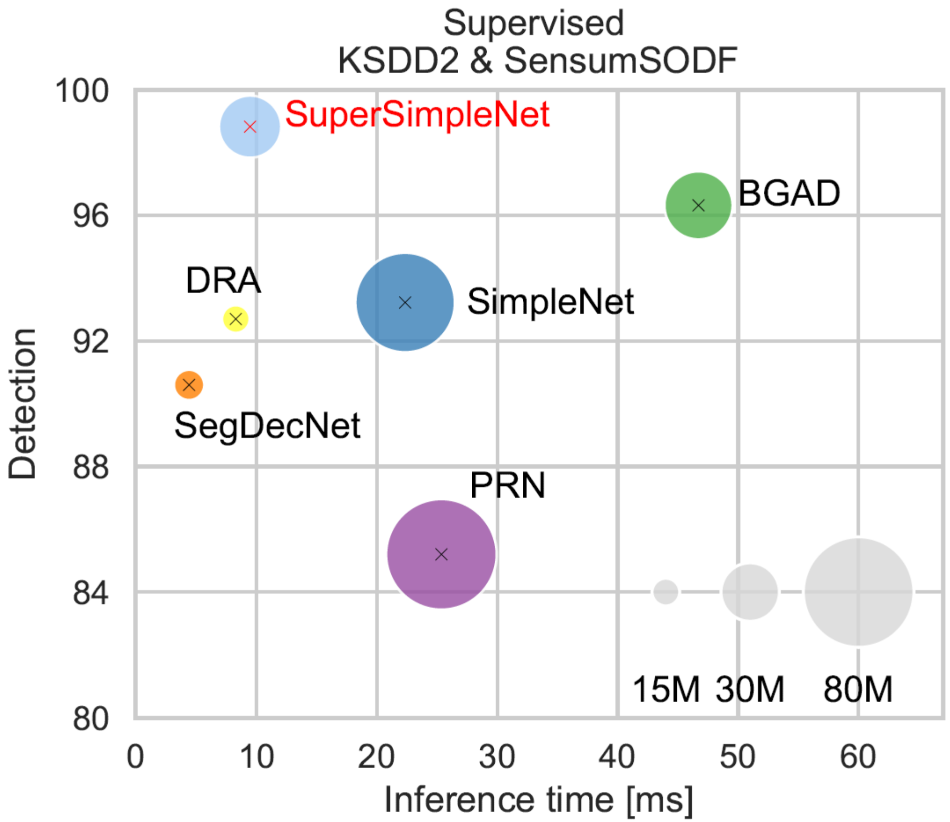

Supervised
KSDD2 & SensumSODF
Detection
Inference time [ms]
SuperSimpleNet
BGAD
DRA
SimpleNet
SegDecNet
PRN
15M
30M
80M
100
96
92
88
84
80
0
10
20
30
40
50
60

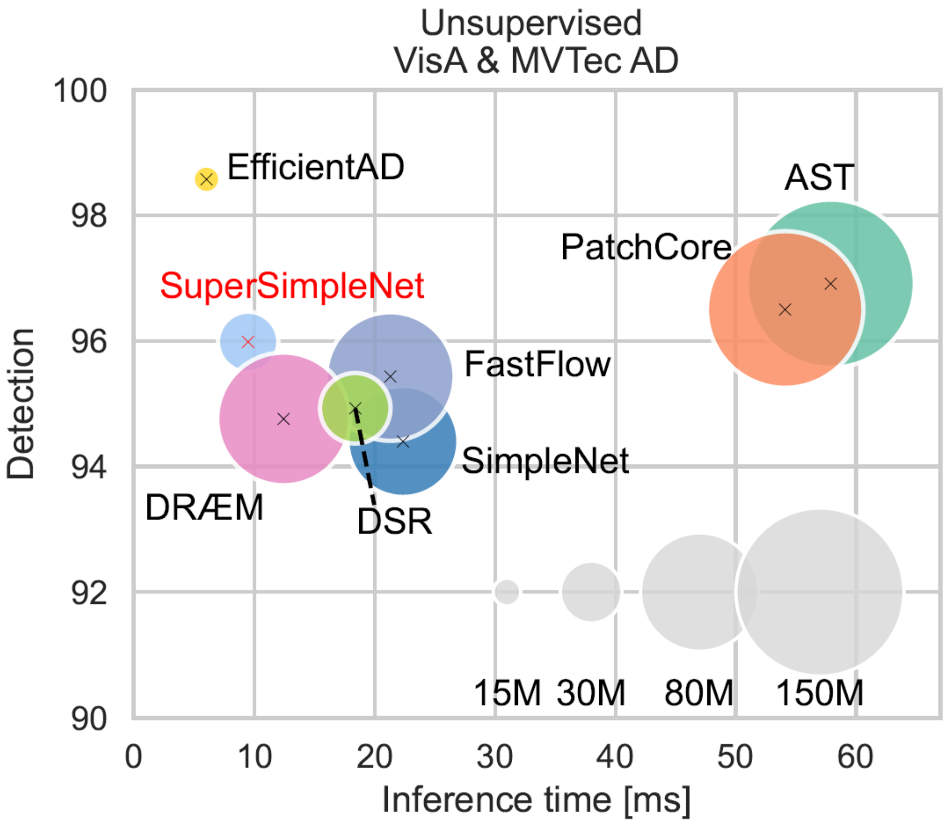

Unsupervised
VisA & MVTec AD
Detection
Inference time [ms]
EfficientAD
AST
PatchCore
SuperSimpleNet
FastFlow
SimpleNet
DRÆM
DSR
15M
30M
80M
150M
100
98
96
94
92
90
0
10
20
30
40
50
60

Spectrum of learning regimes

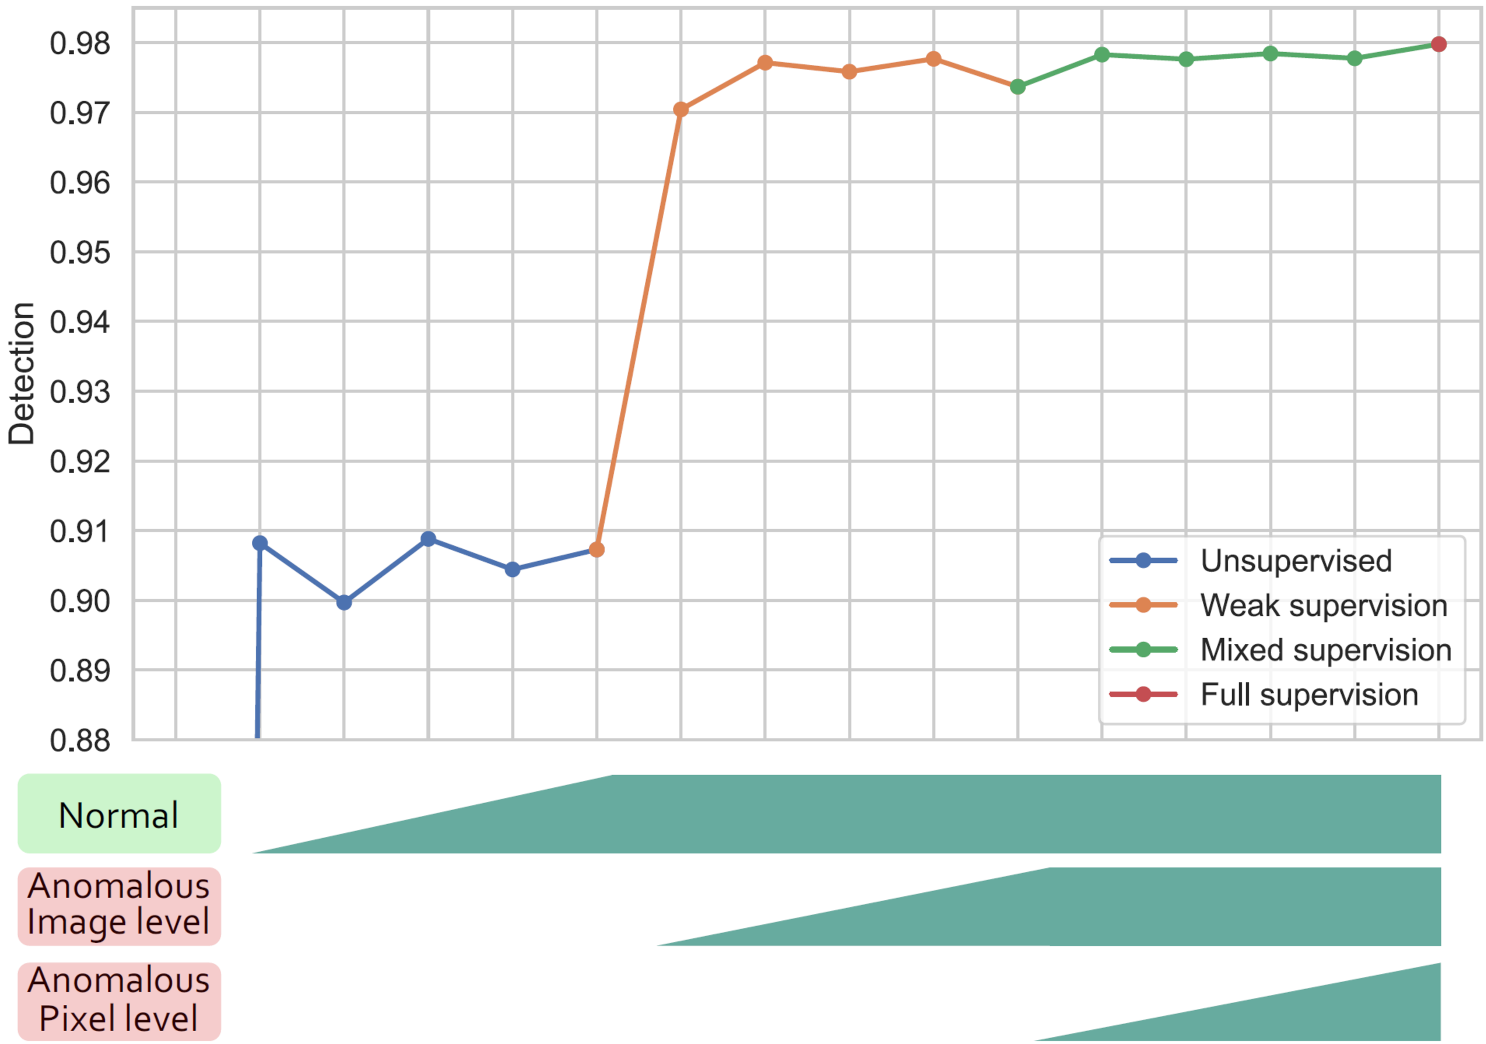

Learning regimes

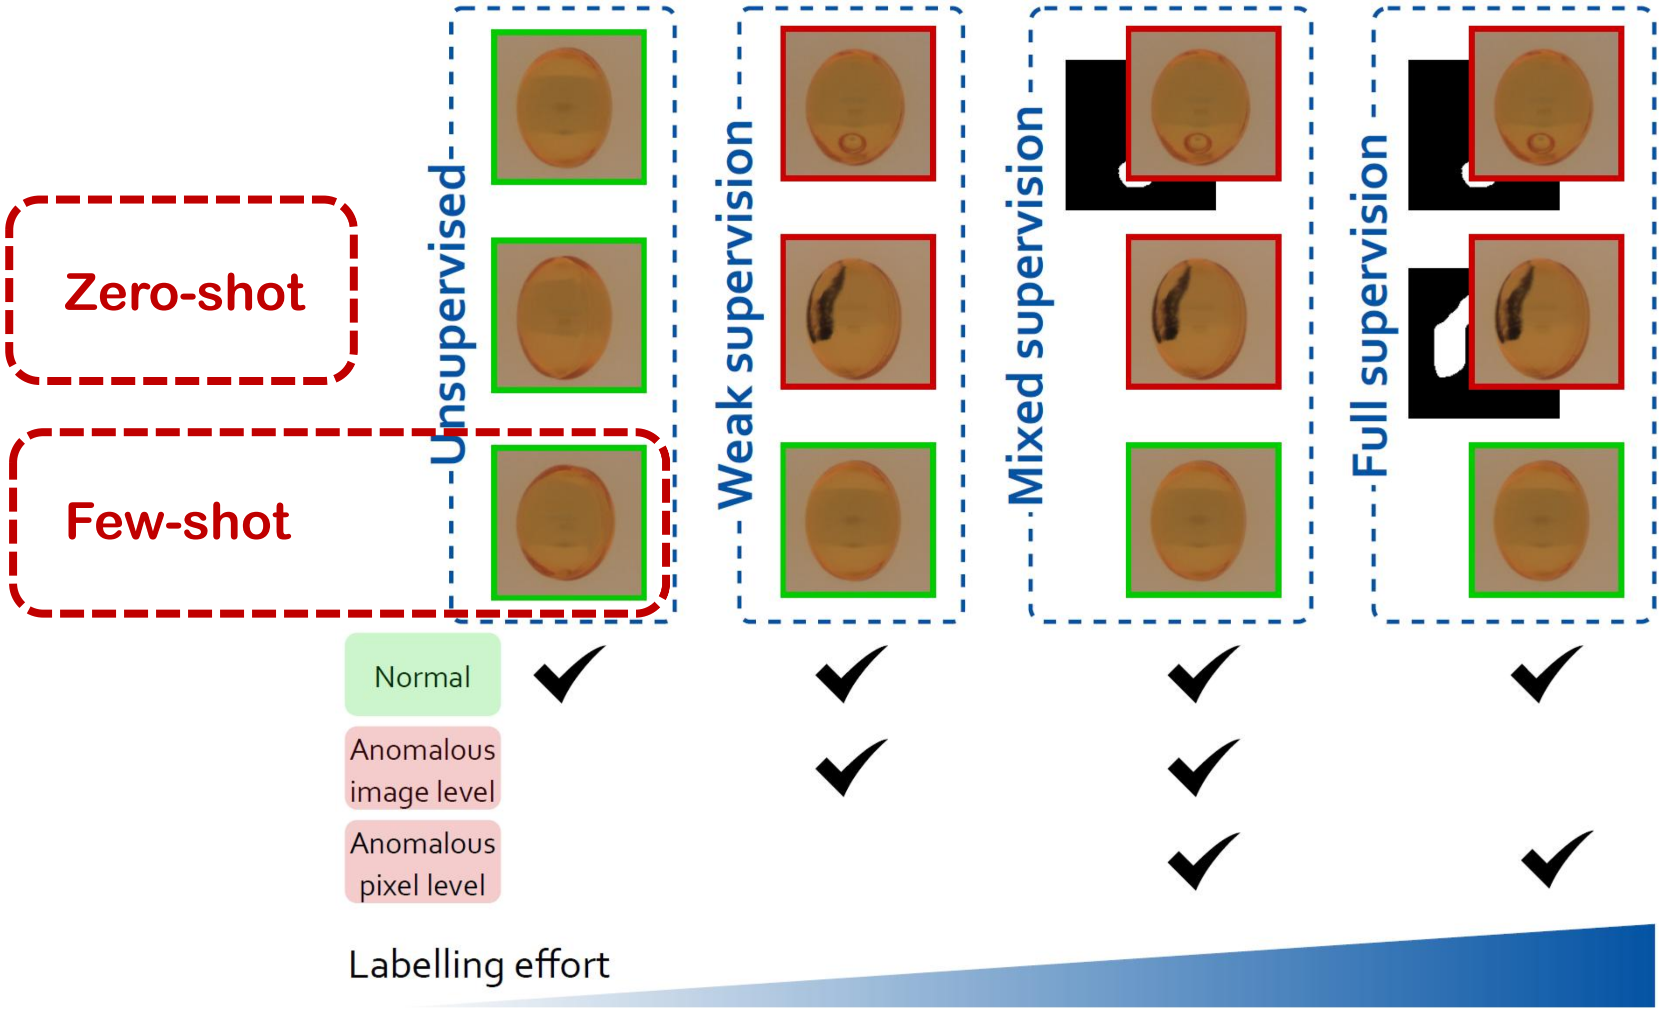

Zero-shot anomaly detection

- The problem is in fact **General Defect Detection / Segmentation**
- Training on Testset A → Testing on Testset B
- Leveraging knowladge encoded in LLMs (& VLMs)

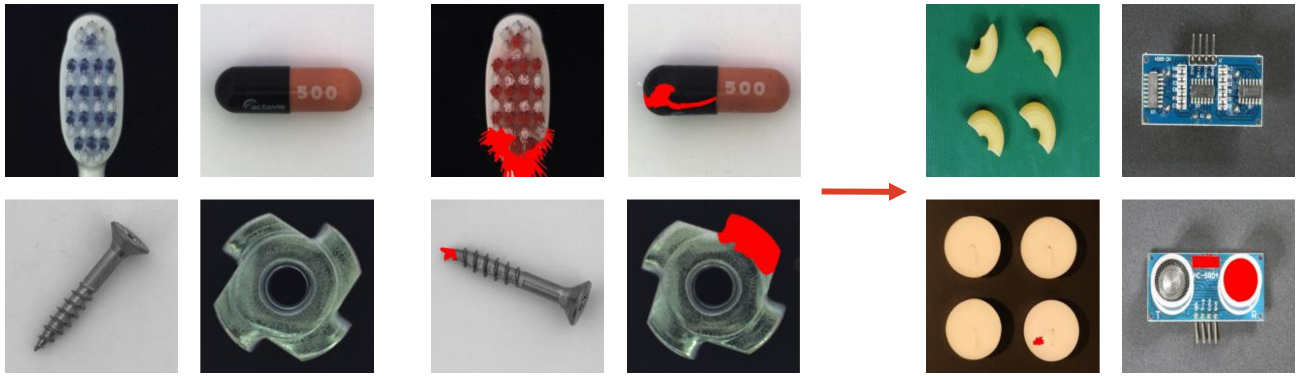

Train Test

- Current SOTA
  - Use CLIP as a backbone
  - Use trainable text embeddings to encode generic notions of normality and abnormality

Research questions and proposed solution

- Why VLMs? Is text really needed? Why not VFMs?

- Problems:
  - Existing datasets have low data diversity
    → Most backbones do not generalize
  - Current methods finetune only later layers
    → Features are suboptimal for zero-shot anomaly detection

- Proposed solution:
  - **Synthetic dataset generation** scheme
    → high data diversity
  - **AnomalyVFM**: Rather than complex additions add simple parameter efficient drop-ins across all layers
    → features adapted for anomaly detection

Synthetic Dataset Generation

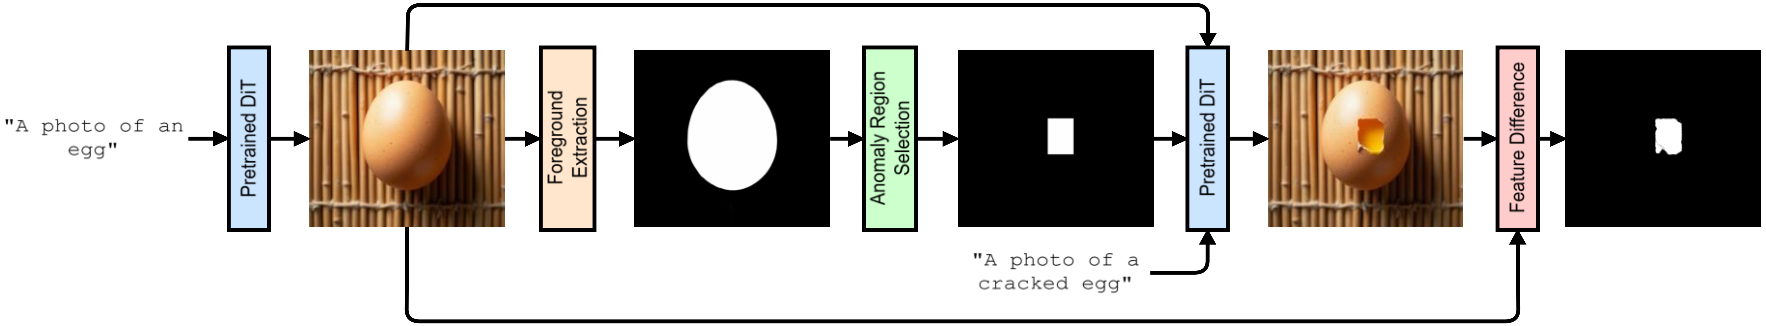

- Foreground extracted with IS-Net
- Features extracted from images with a pretrained model (DINOv2)
- Anomaly mask approximated as the feature difference
- If no pixel is positive, we discard the sample

MUXAD 2025-2027

CVPR 2026

Examples of synthetic images

AnomalyVFM

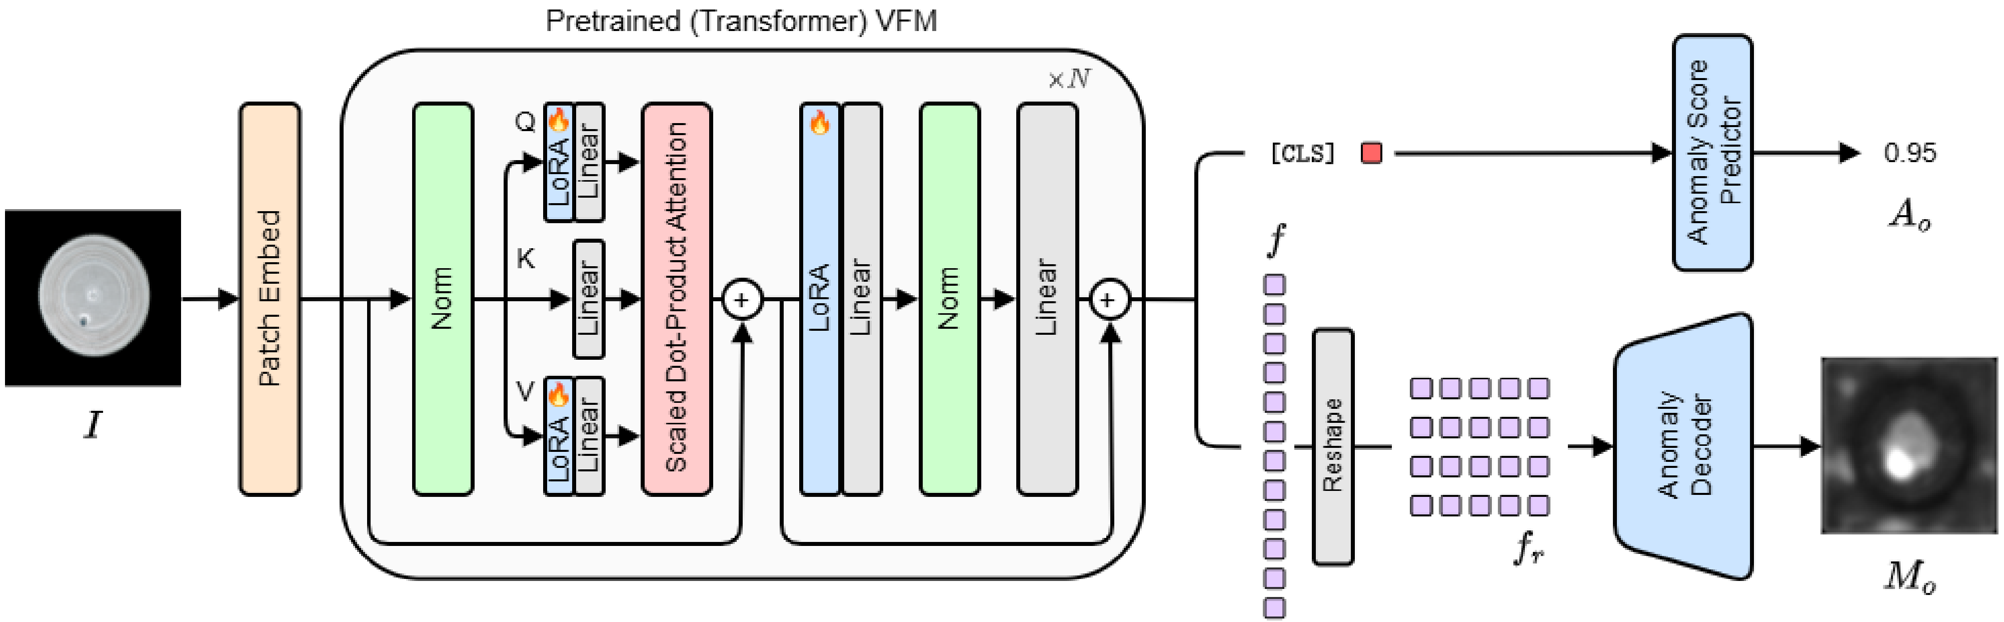

- Add LoRA to each layer of VFM
- Add Decoder and Predictor
- Optimize with Confidence weighted loss (akin to DUST3R)
- **Strong VFM for anomaly detection**

Experimental results – Zero-shot AD

- 9 industrial and
  9 medical datasets
- Detection and
  localisation
- Performance metrics:
  - AUROC
  - F1-max

MUXAD
2025-2027

CVPR 2026

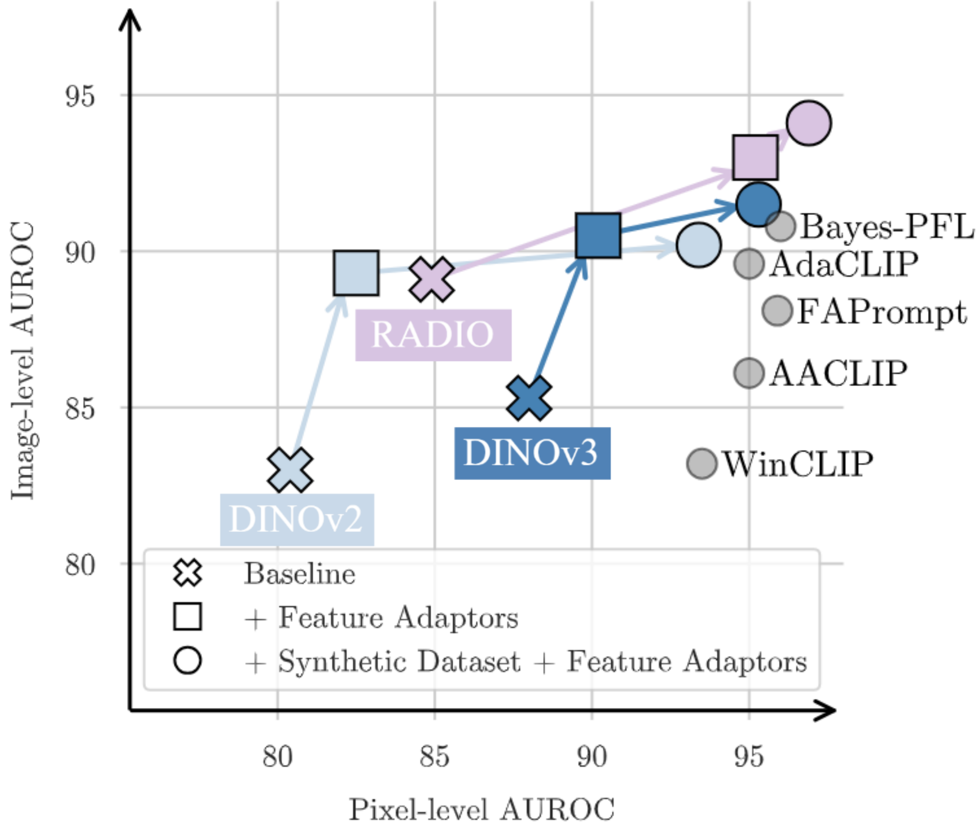

Zero-shot AD – quantitative results

Metric	Dataset	SAA [8] ToC'25	WinCLIP [25] CVPR'23	AnomalyCLIP [80] ICLR'24	AdaCLIP [9] ECCV'24	AACLIP [46] CVPR'25	Bayes-PFL [52] CVPR'25	FAPrompt [84] ICCV'25	*AnomalyVFM*
Image-level (AUROC, max-F1)	MVTec AD	(63.5, 87.4)	(91.8, 92.9)	(91.6, 92.7)	(89.2, 90.6)	(90.5, 90.4)	(92.3, 93.1)	(91.1, 92.2)	(94.9, 94.1)
	VisA	(67.1, 75.9)	(78.1, 80.7)	(82.0, 80.4)	(85.8, 83.1)	(84.6, 78.8)	(87.0, 84.1)	(82.8, 81.3)	(93.6, 90.1)
	BTAD	(59.0, 89.7)	(68.2, 67.8)	(88.2, 83.8)	(88.6, 88.2)	(94.8, 93.7)	(93.2, 91.9)	(90.7, 88.1)	(96.0, 91.0)
	MPDD	(42.7, 73.9)	(61.4, 77.5)	(77.5, 80.4)	(76.0, 82.5)	(75.1, 79.8)	(81.2, 83.5)	(76.6, 80.4)	(85.5, 87.8)
	RealIAD	(51.4, 64.6)	(74.7, 69.8)	(78.7, 80.0)	(79.2, 73.5)	(81.3, 76.4)	(85.2, 78.7)	(81.6, 75.2)	(88.0, 81.6)
	KSDD	(68.6, 37.6)	(93.3, 79.0)	(84.5, 71.1)	(97.1, 90.7)	(69.3, 57.1)	(88.2, 56.0)	(81.3, 71.1)	(92.5, 69.7)
	KSDD2	(91.6, 67.0)	(94.2, 71.5)	(94.1, 80.0)	(95.9, 86.7)	(95.9, 84.4)	(97.3, 87.6)	(95.6, 84.8)	(97.1, 79.2)
	DAGM	(87.1, 88.8)	(91.8, 87.6)	(97.7, 90.1)	(99.1, 97.5)	(93.2, 79.4)	(97.7, 95.7)	(97.3, 89.3)	(99.6, 95.8)
	DTD	(94.4, 93.5)	(95.1, 94.1)	(93.9, 93.6)	(95.5, 94.7)	(90.4, 92.8)	(95.1, 95.1)	(95.9, 94.7)	(99.4, 99.0)
	Average	(69.5, 75.4)	(83.2, 80.1)	(87.6, 83.6)	(89.6, 87.5)	(86.1, 81.4)	(90.8, 85.1)	(88.1, 84.1)	(94.1, 87.6)
Pixel-level (AUROC, max-F1)	MVTec AD	(75.5, 38.1)	(88.7, 43.4)	(91.1, 39.1)	(88.7, 43.4)	(91.4, 46.4)	(91.8, 49.0)	(90.8, 39.3)	(92.7, 45.2)
	VisA	(76.5, 31.6)	(95.5, 37.7)	(95.5, 28.3)	(95.5, 37.7)	(94.8, 30.2)	(95.6, 34.3)	(95.6, 27.6)	(96.2, 31.2)
	BTAD	(65.8, 14.8)	(92.1, 51.7)	(94.2, 49.7)	(92.1, 51.7)	(97.3, 55.1)	(93.9, 52.0)	(95.8, 52.6)	(92.3, 49.7)
	MPDD	(81.7, 18.9)	(96.1, 34.9)	(96.5, 34.2)	(96.1, 32.8)	(96.7, 30.0)	(97.8, 35.0)	(95.5, 31.9)	(97.0, 38.1)
	RealIAD	(73.5, 4.5)	(87.2, 10.8)	(96.3, 39.0)	(97.2, 43.0)	(96.2, 40.2)	(97.2, 41.2)	(96.2, 38.3)	(96.4, 40.4)
	KSDD	(78.8, 6.6)	(97.7, 54.5)	(90.6, 42.5)	(97.7, 54.5)	(87.1, 28.0)	(96.5, 6.6)	(93.1, 47.2)	(99.0, 10.1)
	KSDD2	(79.9, 63.4)	(94.4, 23.9)	(98.5, 59.8)	(98.5, 67.0)	(99.5, 63.4)	(97.0, 62.0)	(99.1, 60.4)	(99.3, 55.9)
	DAGM	(91.5, 57.5)	(91.5, 57.5)	(95.6, 58.9)	(91.5, 57.5)	(96.2, 53.3)	(95.9, 49.8)	(98.6, 60.2)	(99.4, 61.3)
	DTD	(97.9, 71.6)	(97.9, 71.6)	(97.9, 62.2)	(97.9, 71.6)	(95.8, 59.6)	(98.4, 65.2)	(98.1, 61.9)	(99.4, 66.5)
	Average	(80.1, 34.1)	(93.5, 42.9)	(95.1, 46.0)	(95.0, 51.0)	(95.0, 45.1)	(96.0, 43.9)	(95.9, 46.6)	(96.9, 44.3)

Zero-shot Anomaly Detection

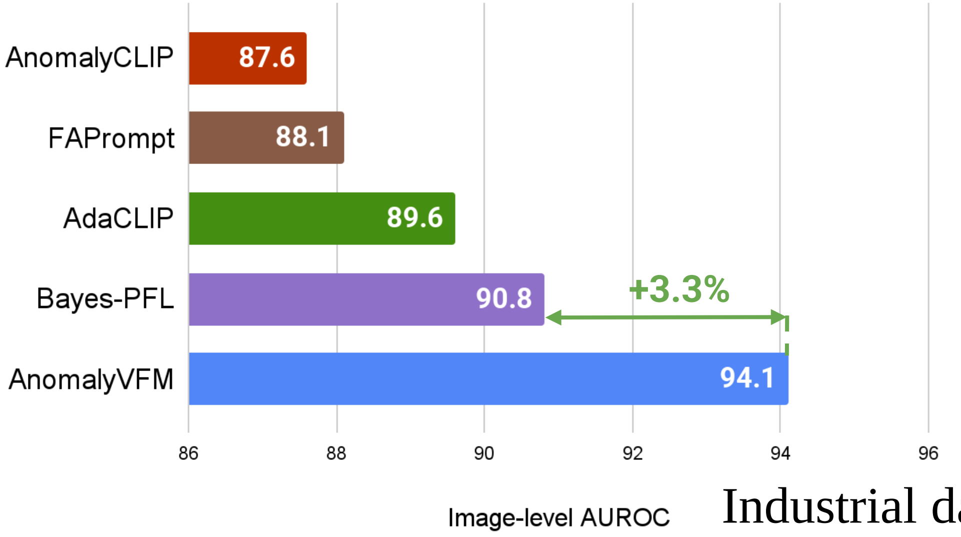

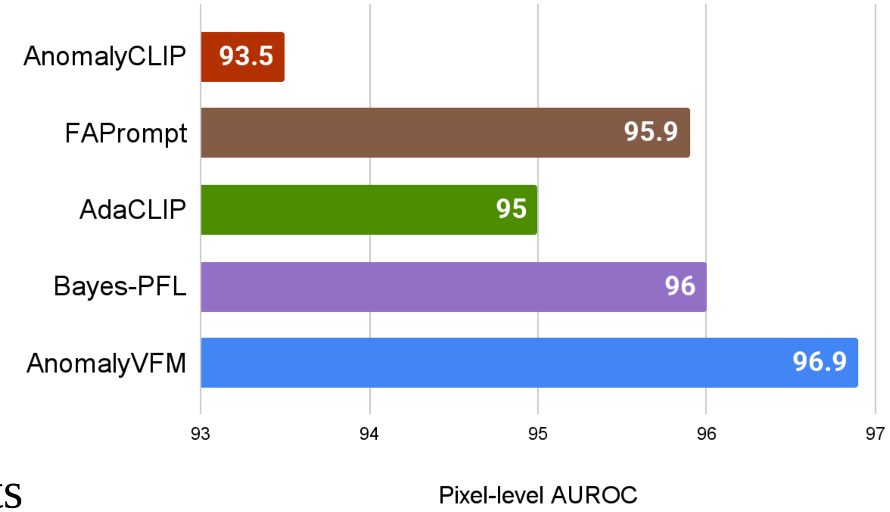

Industrial datasets

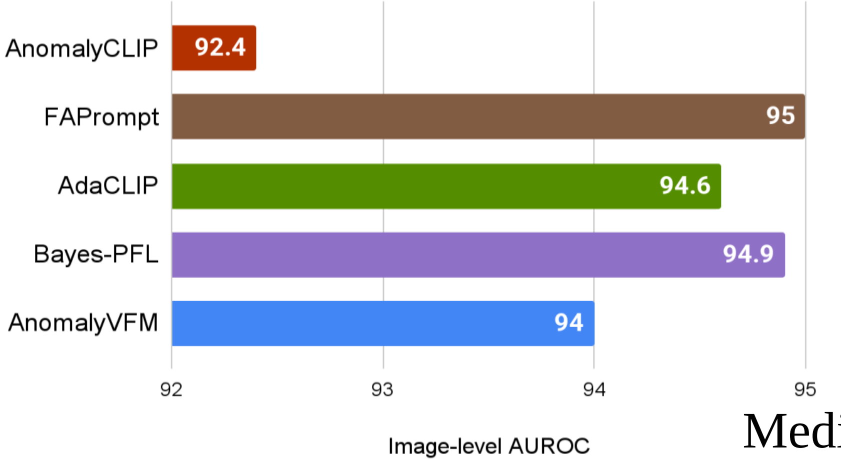

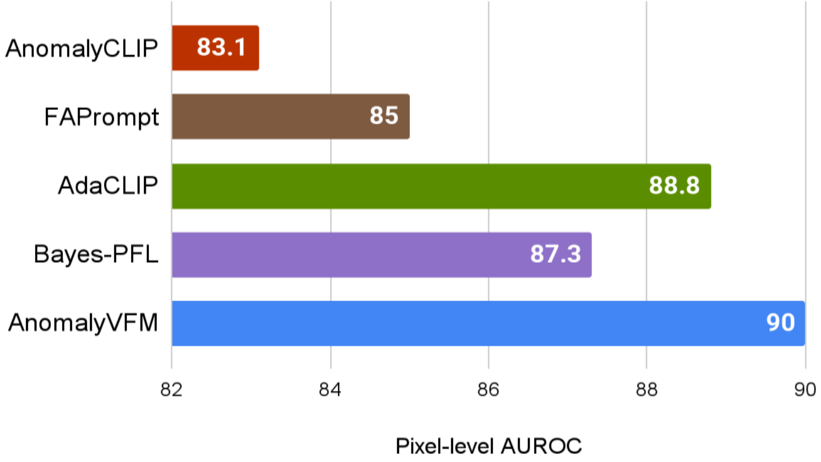

Medical datasets

Zero-shot AD – qualitative results

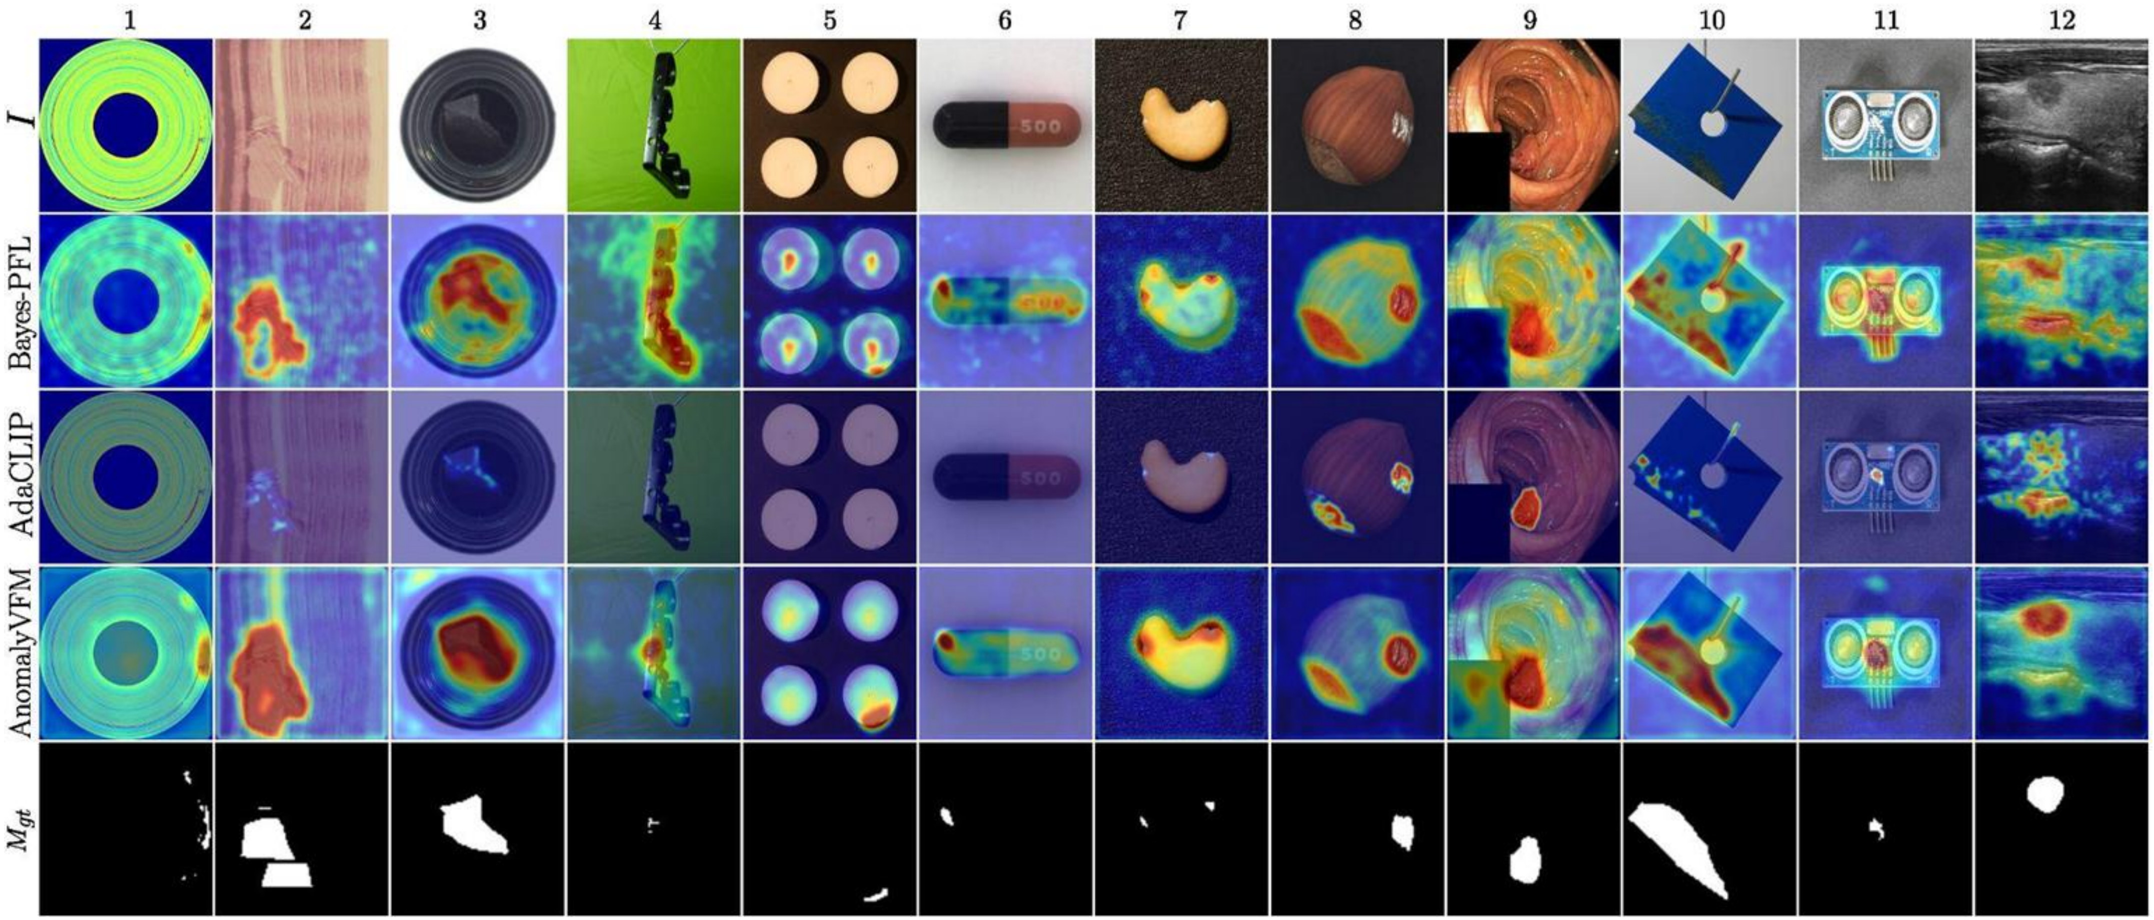

Few-shot anomaly detection

- Simply fine-tune zero shot model using few normal samples
- Results on MVTec AD:
- Better than few-shot SOTA!
- AnomalyVFM **strong backbone** for anomaly detection

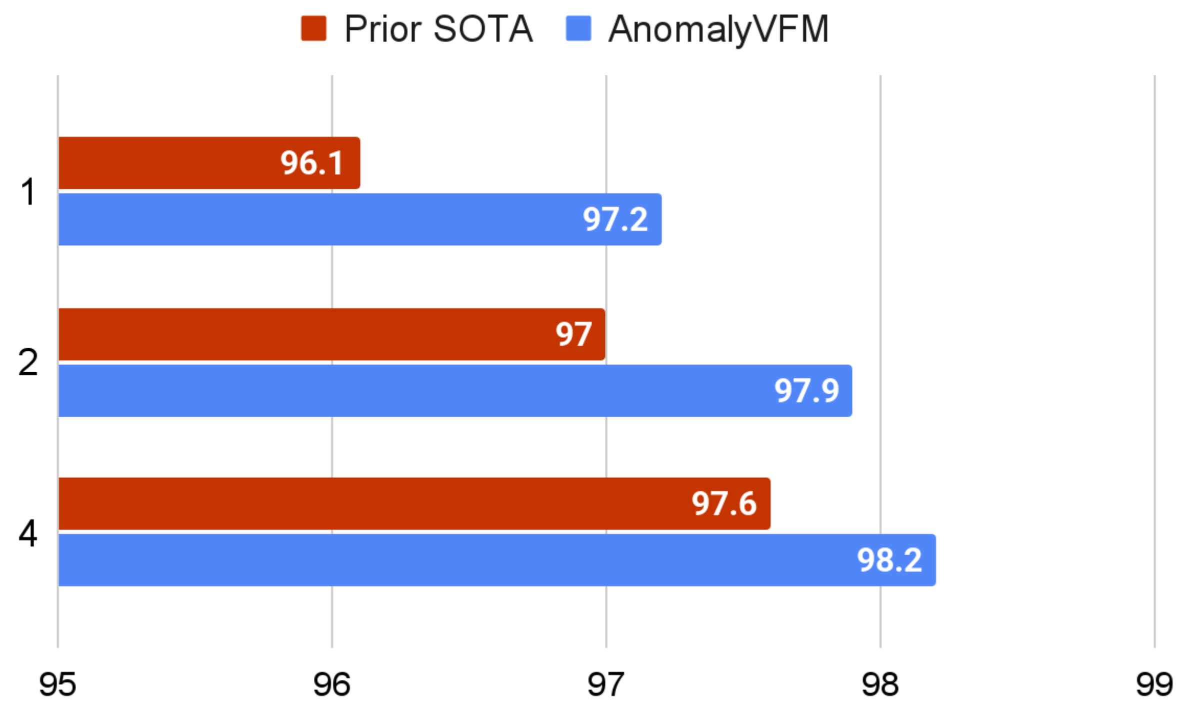

VFM for AD – using it as a backbone

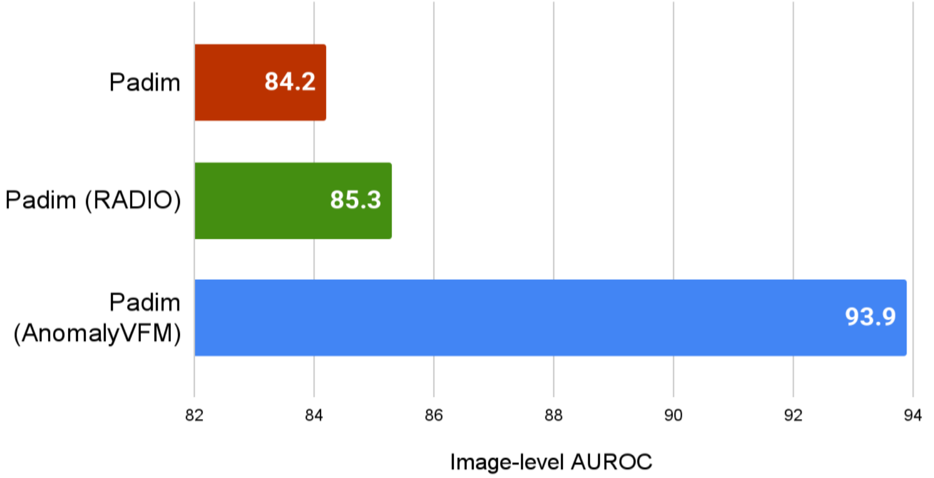

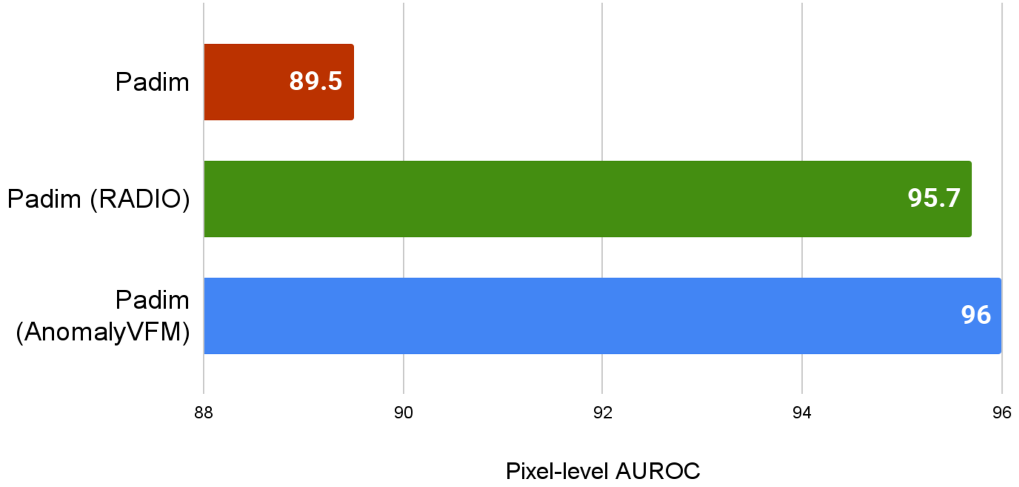

- Preliminary results
- Results show great potential
- MVTec AD under the multi-class setting (one model for all categories)

Practical considerations

Industrial appications

MV vs. DL

Development and mainainance

Academia vs. industry

Industrial applications

- Data-driven deep-learning approach to surface-defect detection

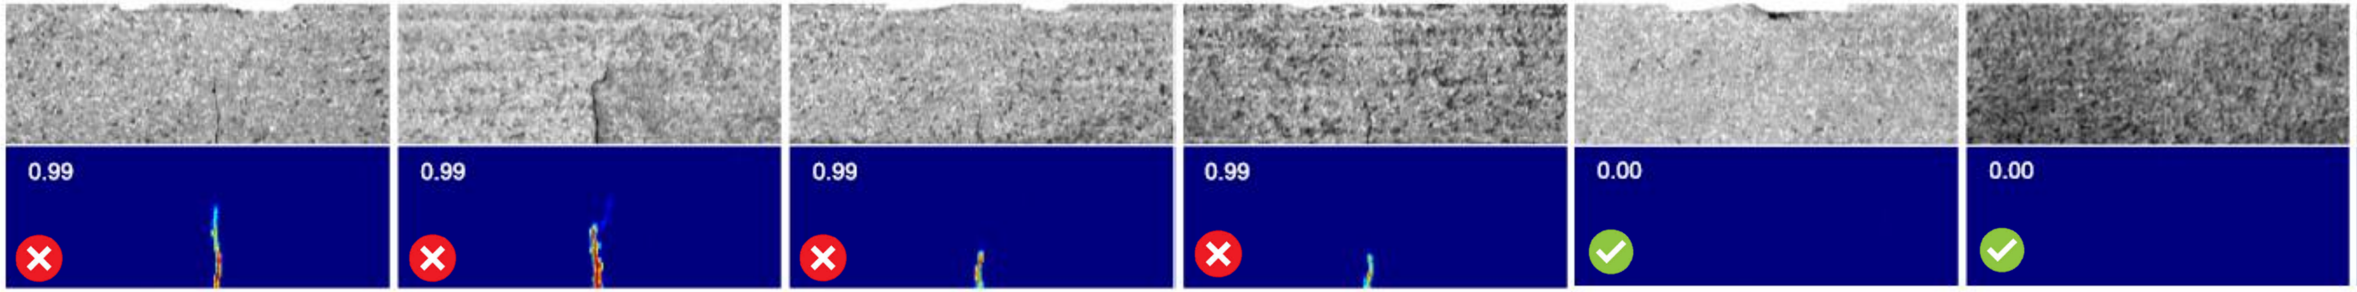

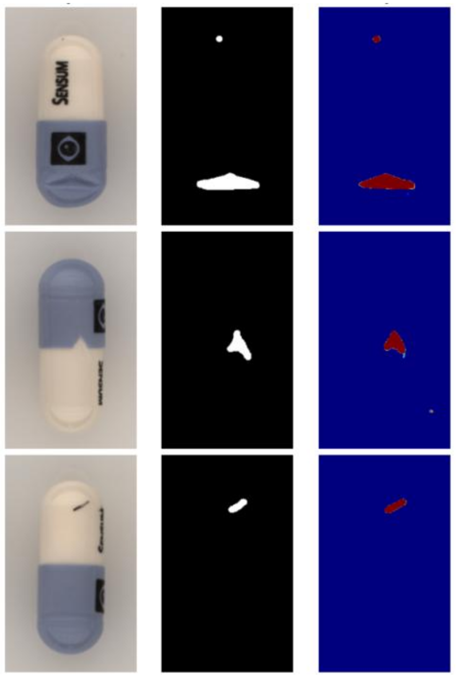

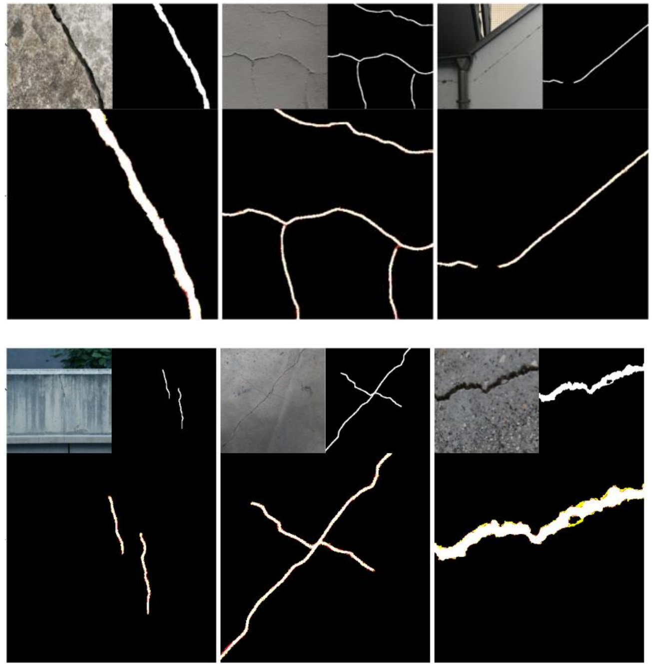

Conventional MV vs. DL

- O: „Deep learning is just a hype and a non-understandable and nonreliable black box“
- Y: „Deep learning is all you need“

- Use adequate HW (camera, lenses, illumination, background) to constrain the problem
  - garbage in garbage out
- Use good old MV techniques when they suffice
  - for less challenging or well-defined problems
  - in controlled environments
- Use MV techniques to constrain the problem
  - and make DL learning easier
  - requiring less training images
- Use DL where the problem is data-driven or hard
  - in less-controlled environments for more general tasks
  - or to speed up the development cycle

Adequate tools

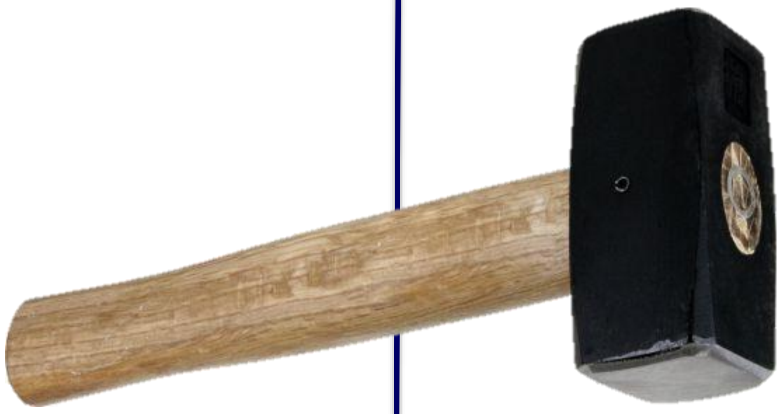

Deep learning

Routine solutions	Rule-based solutions	Data-driven solutions	General intelligence

Function approximator

- Deep model as a function approximator
- Different training possibilities:

function	known	unknown
$f(x) \doteq y$	x_{tr}, y_{tr}	f
$f(x)$	x_{tr}	f
$f(x) \doteq \hat{f}(x)$	$x_{tr}, \hat{f}$	f
$f(f^{-1}(y)) \doteq y$	y_{tr}, f^{-1}	f
$f(g(x)) \doteq y$	g, x_{tr}, y_{tr}	f
$g(f(x)) \doteq y$	g, x_{tr}, y_{tr}	f

Development and maintainance

- Data, data, data!
  - Sufficient amount of representative data
  - Correctly labelled data
- Adequate design of deep architecture
  - Adequate backbone, architecture, loss function,…
  - Learning, parameter optimisation
- Efficient implementation
  - Exectution speed
  - Integration
- Development and maintainance
  - Incremental improvement of the learned model
  - Adaptation to the changes in the environment

Real world considerations

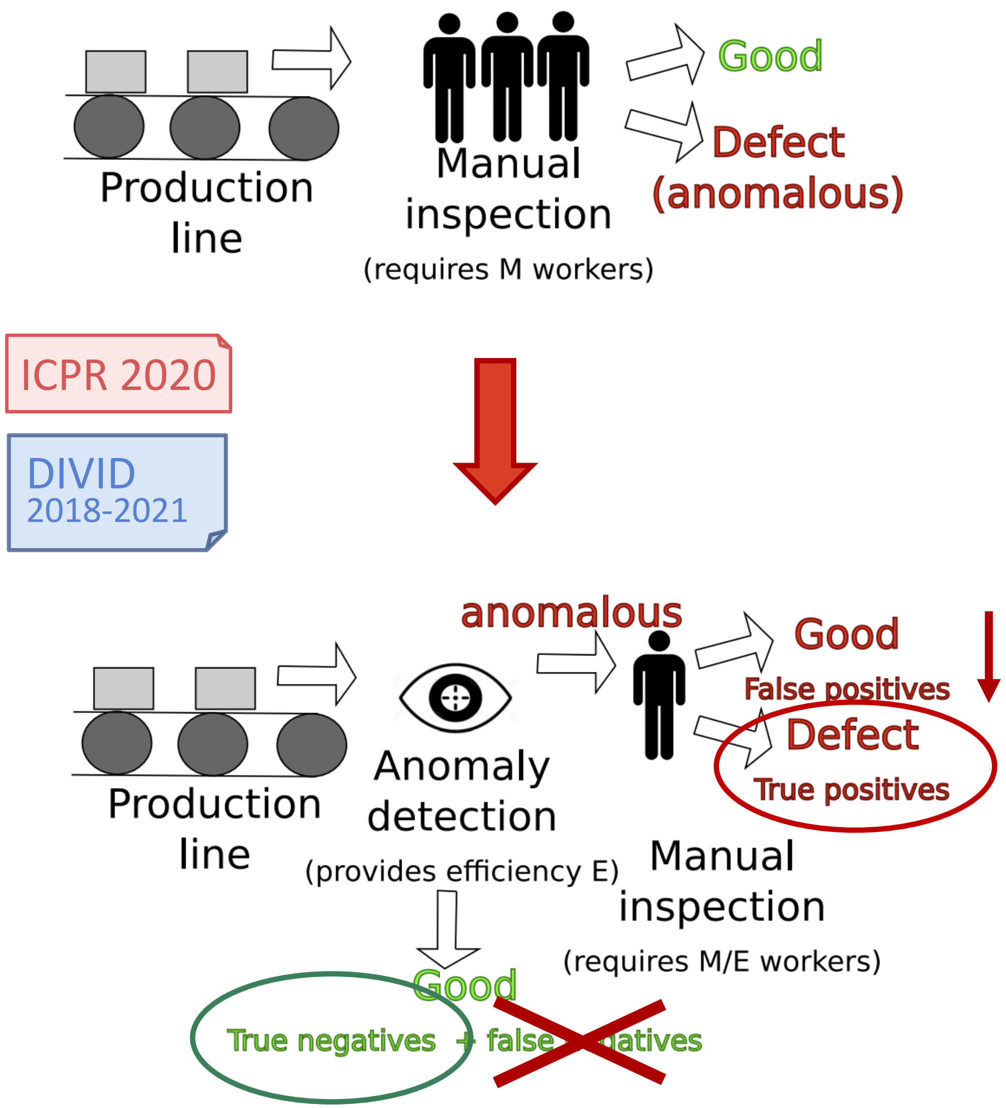

- Human in the loop
  - at least for some time
- Challenges
  - Robustness
  - Dependency on the training set
  - Domain shift
  - Non-adaptability
  - Non-interpretepability
- Opportunities
  - Learning under mixed supervision
  - Explainabilty
  - Tunable parameters
  - Compability with conventional MV
  - Agility, quick adaptability

Development and maintainance

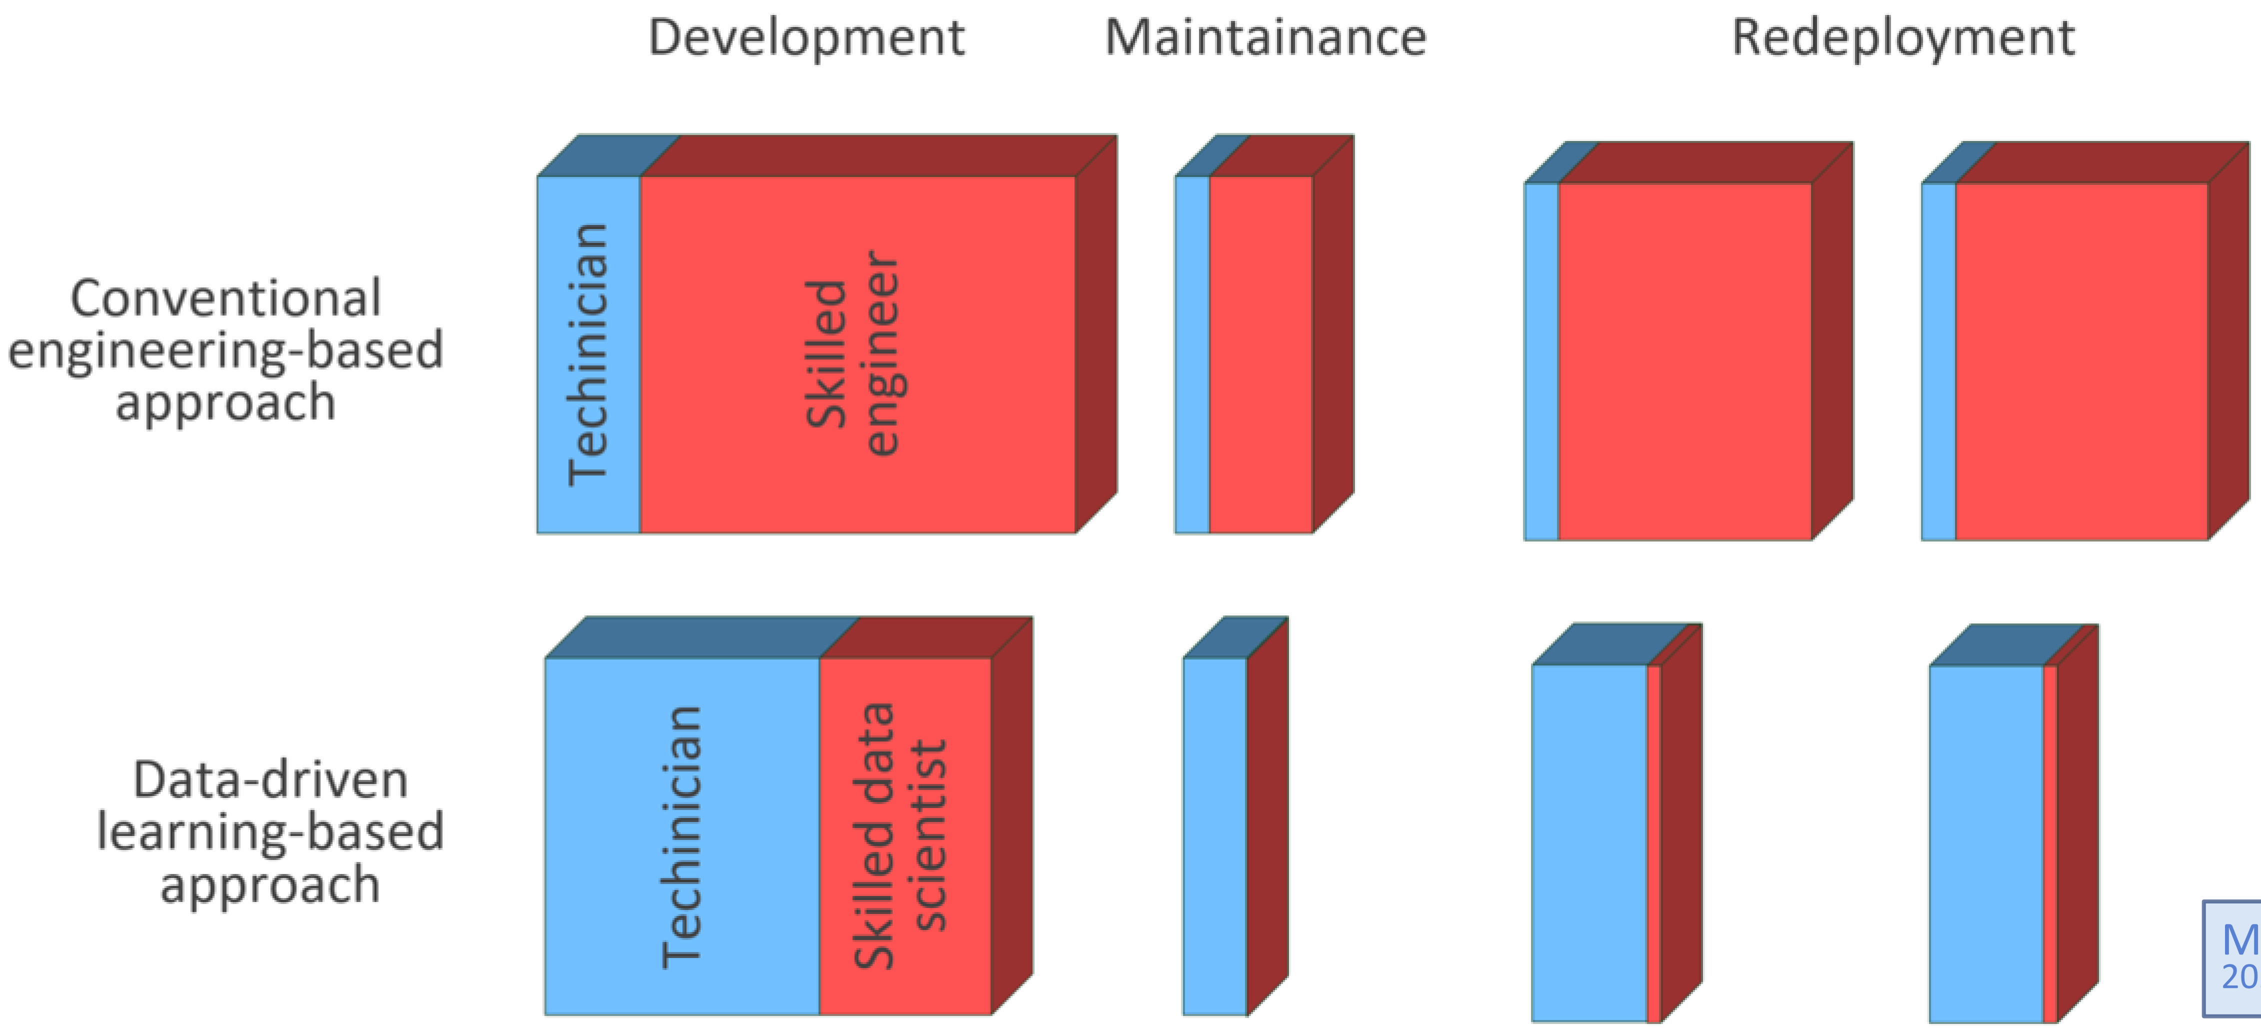

MV4.0
2021-2024

Code

- Supervised methods
  - SedDecNet
  - SuperSimpleNet
- Unsupervised methods
  - PatchCore
  - EfficientAD
  - DRAEM
  - INP-Former
- Mixed supervision
  - Mixed SegDecNet
  - SuperSimpleNet
- Few-Shot methods
  - PromptAD
  - AnomalyDINO
- Zero-Shot methods
  - AnomalyVFM
  - AdaClip
- Various methods and resources
  - Anomalib
  - Awesome Industrial Anomaly Detection
- And many many more…

Conclusion

- Data-driven learning-based solutions
- AI/DL/CV/MV/AD – key enabling technologies
- Wide applicability, interdisciplinarity
- Robustness
- Industry 4.0 and beyond
- New challenges, new opportunities
- Collaboration between academia and industrial partners