SIPTA SEMINAR

Robustness Quantification

using imprecise probabilities to assess the reliability of probabilistic classifiers

GHENT UNIVERSITY

Jasper De Bock, 18 December 2025

Want more than just these slides?

You can watch the entire talk on YouTube!




FLip

Foundations Lab for
imprecise probabilities

GHENT
UNIVERSITY



Adrián Detavernier

Rodrigo Lassance

MACHINE LEARNING

IMPRECISE PROBABILITIES

Adrián Detavernier

Rodrigo Lassance

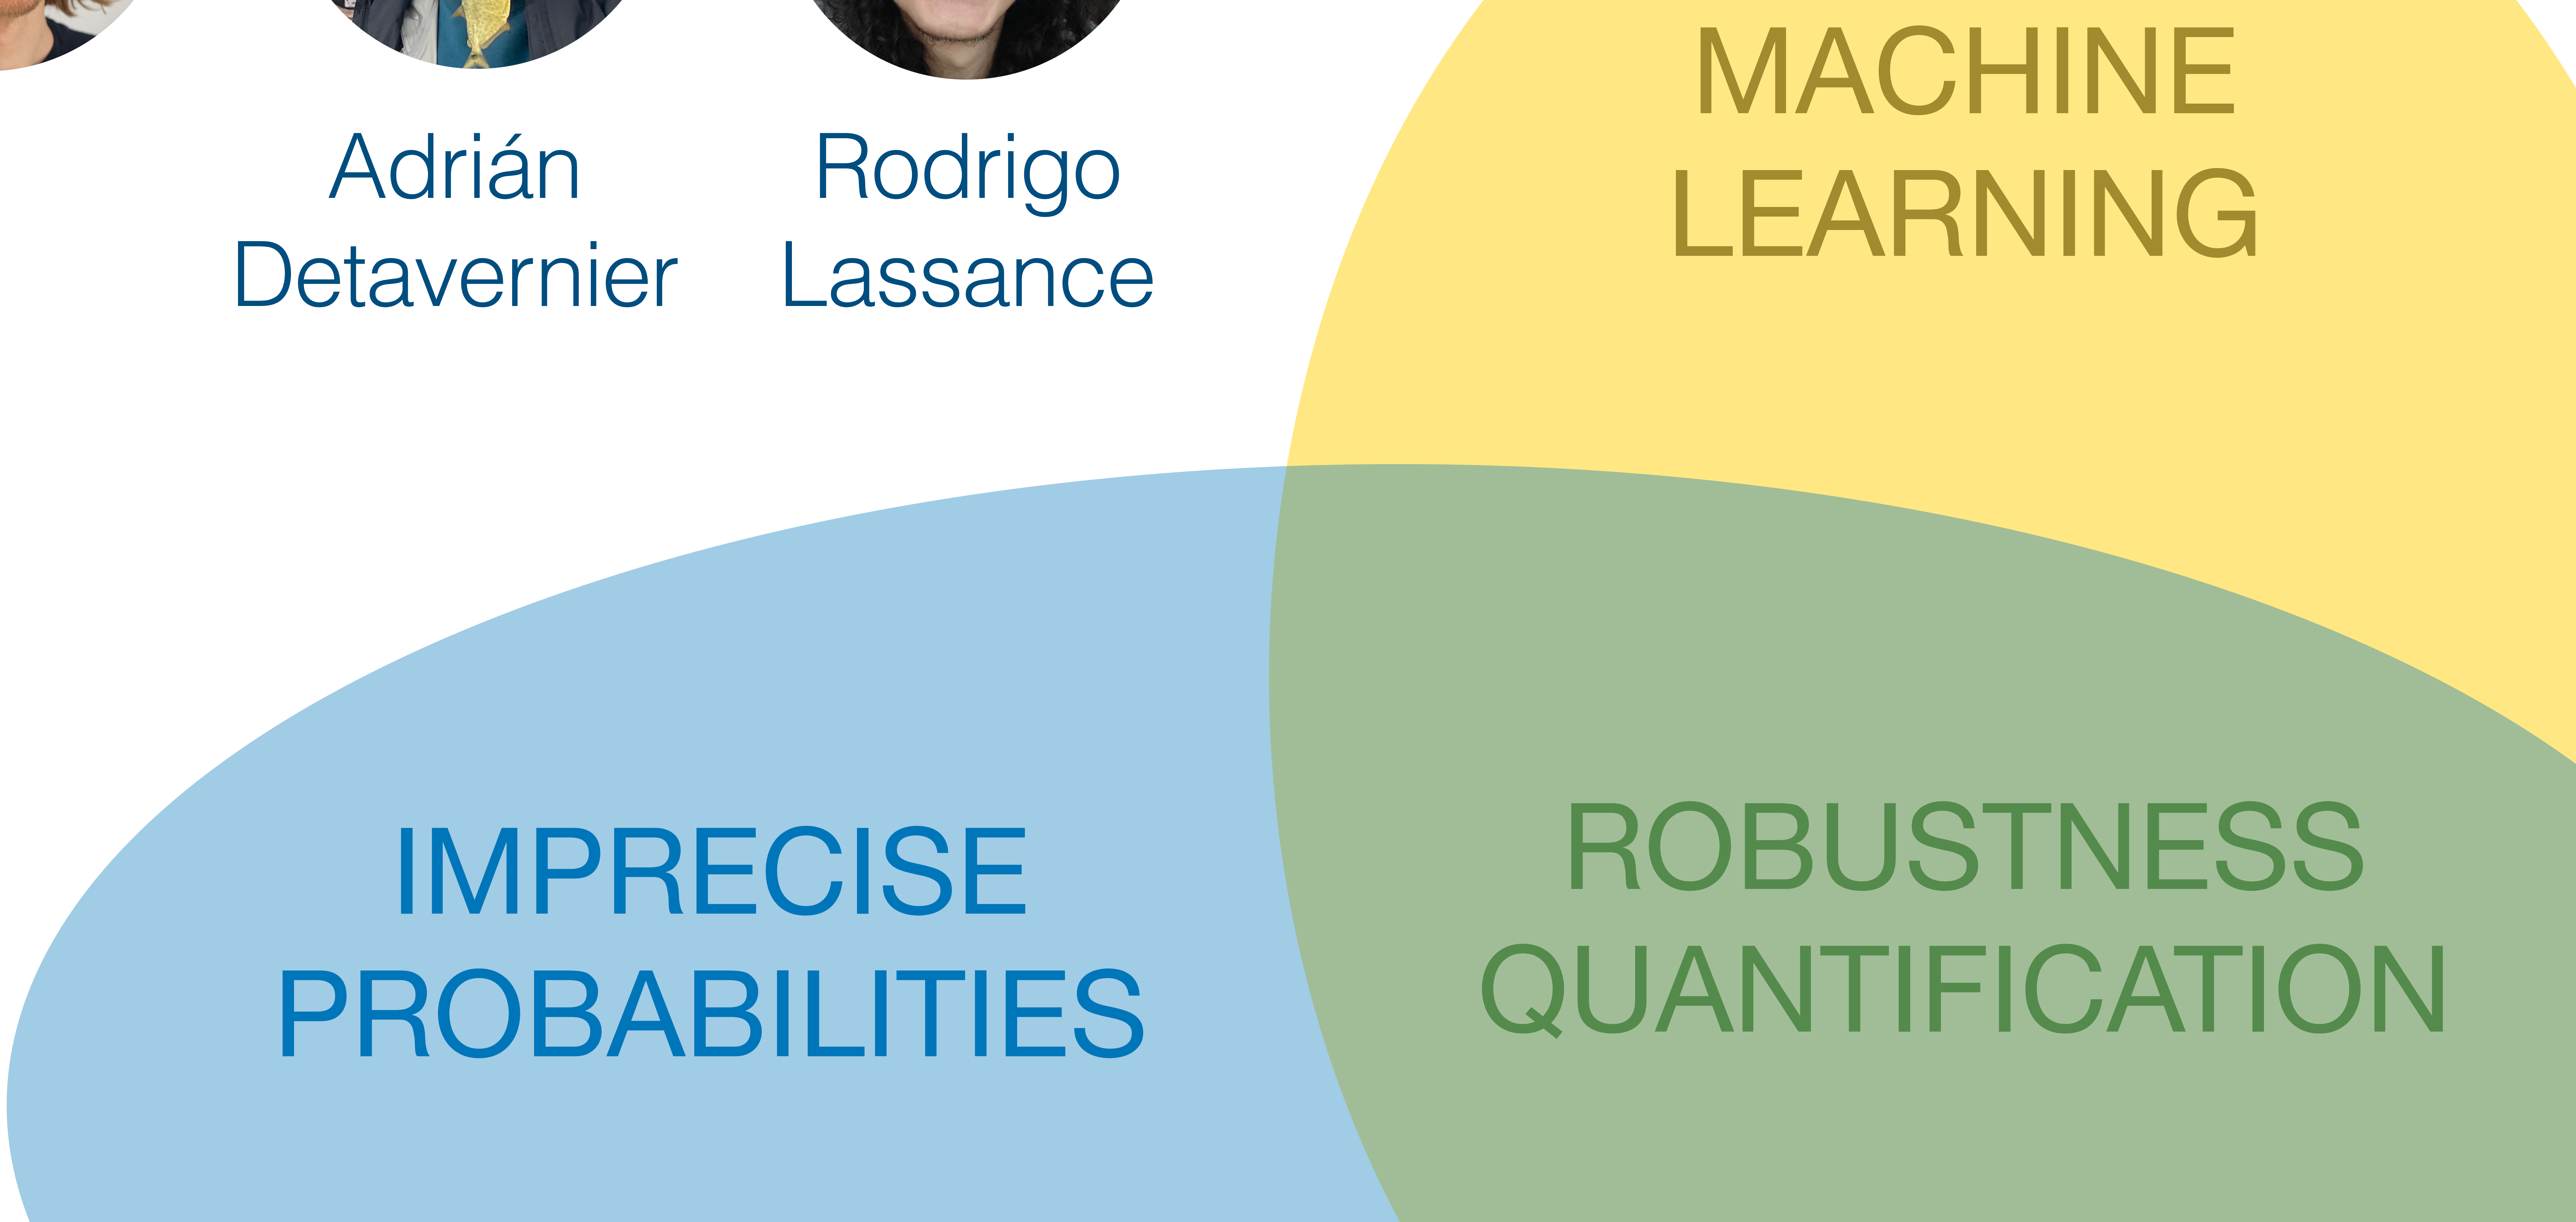

Adrián
Detavernier

Rodrigo
Lassance

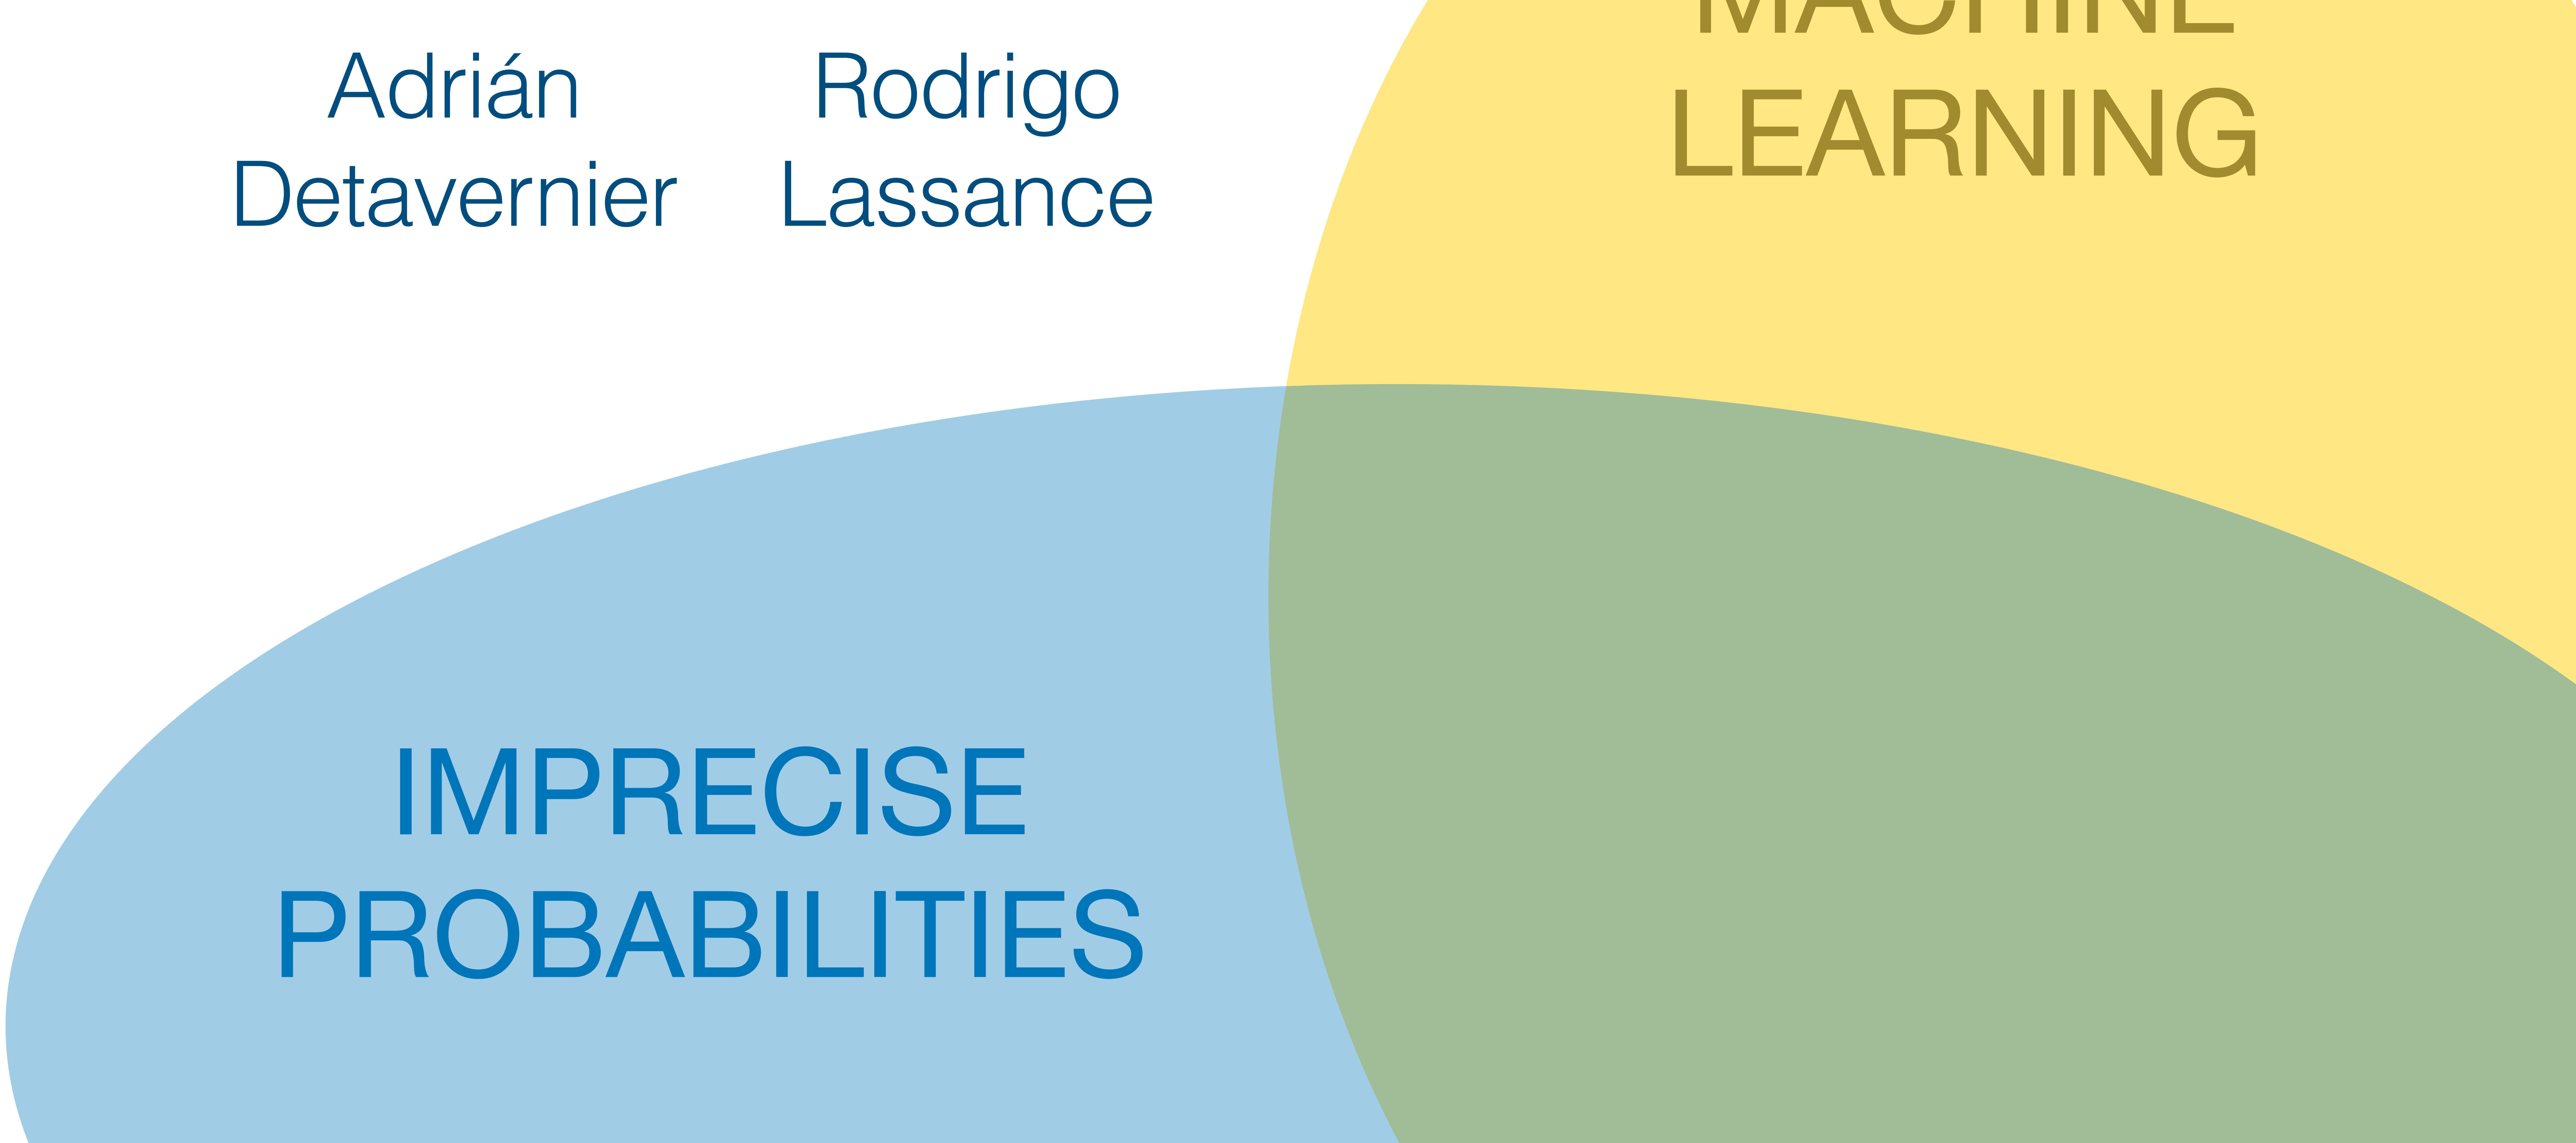

PROBLEM 1:

MACHINE LEARNING

… is unreliable

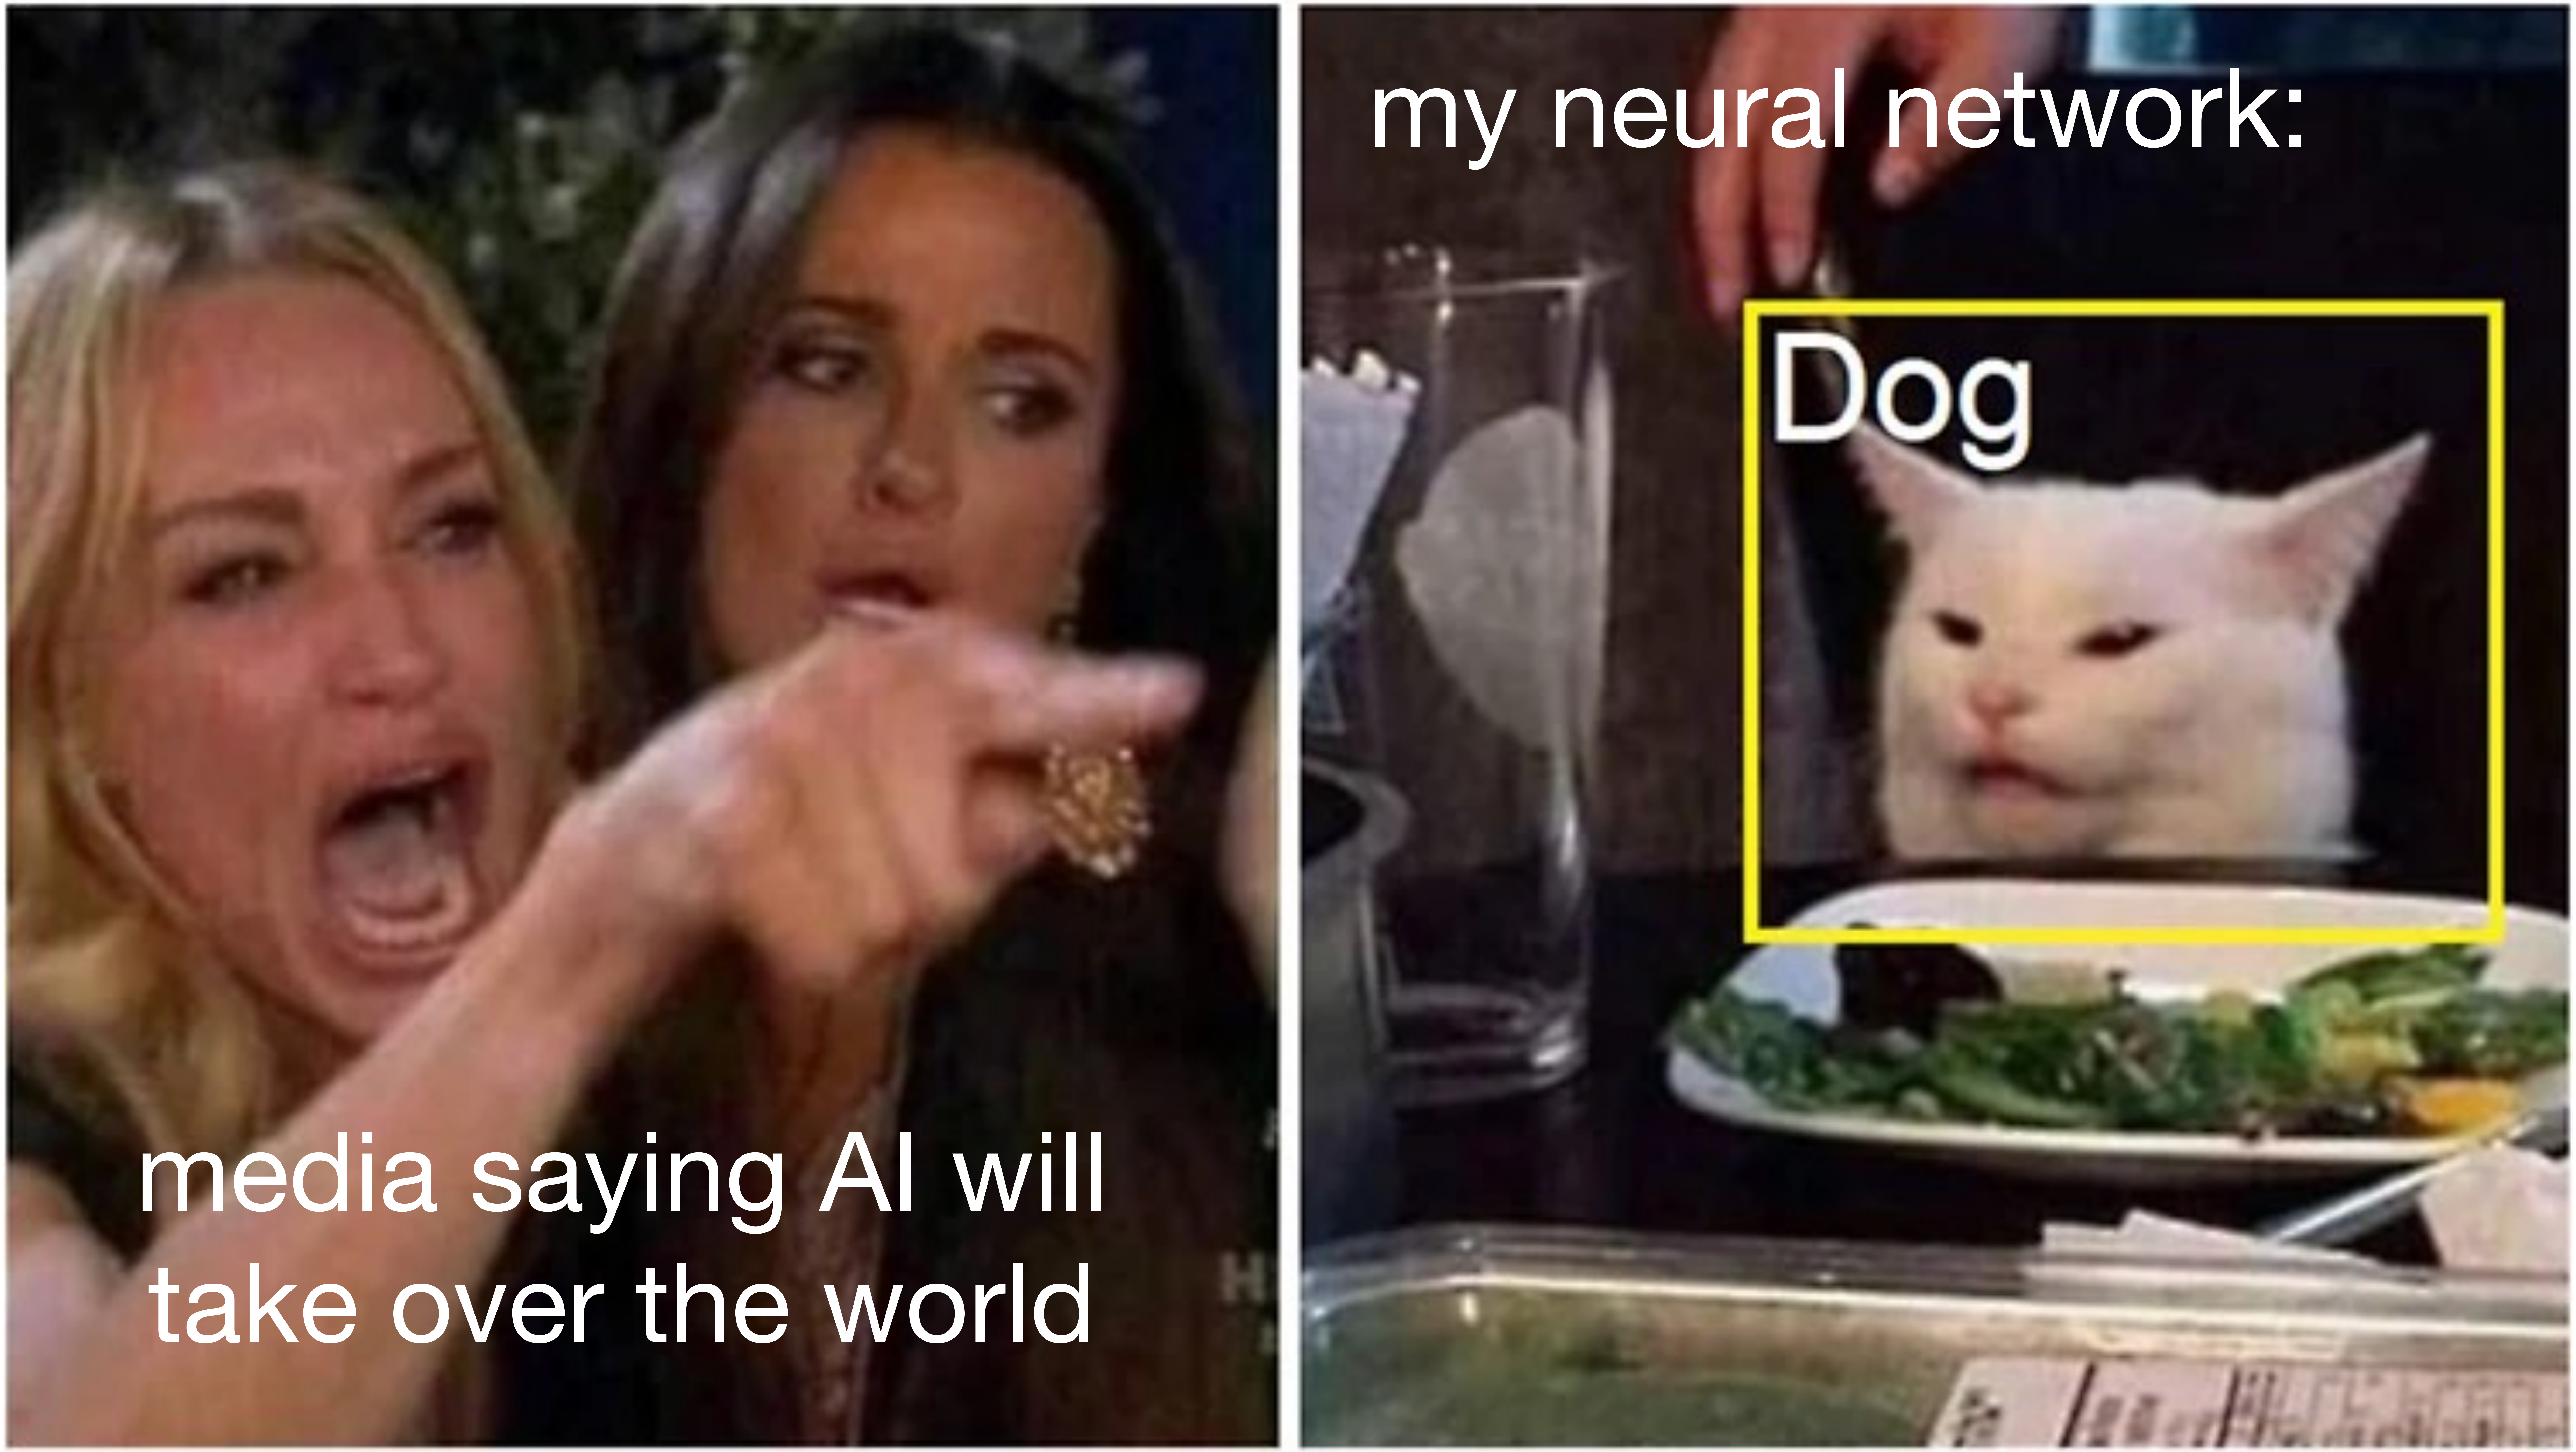
media saying AI will take over the world
my neural network:
Dog

CLASSIFICATION

features x

FEATURES

photo

medical data

...

CLASSIFICATION

features x

FEATURES

photo

medical data

...

try to predict

$$\text{FEATURES} \longrightarrow \text{CLASS}$$

CLASSIFIER h

predicted class $\hat{y} = h(x)$

CLASS

label

best treatment

...

CLASSIFICATION

features x

predicted class $\hat{y} = h(x)$

FEATURES $\xrightarrow{\text{try to predict}}$ **CLASS**

photo

medical data

...

label

best treatment

...

PROBABILISTIC CLASSIFIER h

predict the class y that maximizes $P_{\text{classif}}(y \mid x)$

CLASSIFICATION

features x

predicted class $\hat{y} = h(x)$

FEATURES $\xrightarrow{\text{try to predict}}$ CLASS

photo

medical data

...

label

best treatment

...

DISCRIMINATIVE PROBABILISTIC CLASSIFIER h

learn $P_{\text{classif}}(Y|x)$

predict the class y that maximizes $P_{\text{classif}}(y|x)$

CLASSIFICATION

features x

FEATURES

try to predict

predicted class $\hat{y} = h(x)$

CLASS

photo

medical data

…

label

best treatment

…

GENERATIVE PROBABILISTIC CLASSIFIER h

learn $P_{\text{classif}}(X, Y)$

predict the class y that maximizes $P_{\text{classif}}(y \mid x)$

CLASSIFICATION

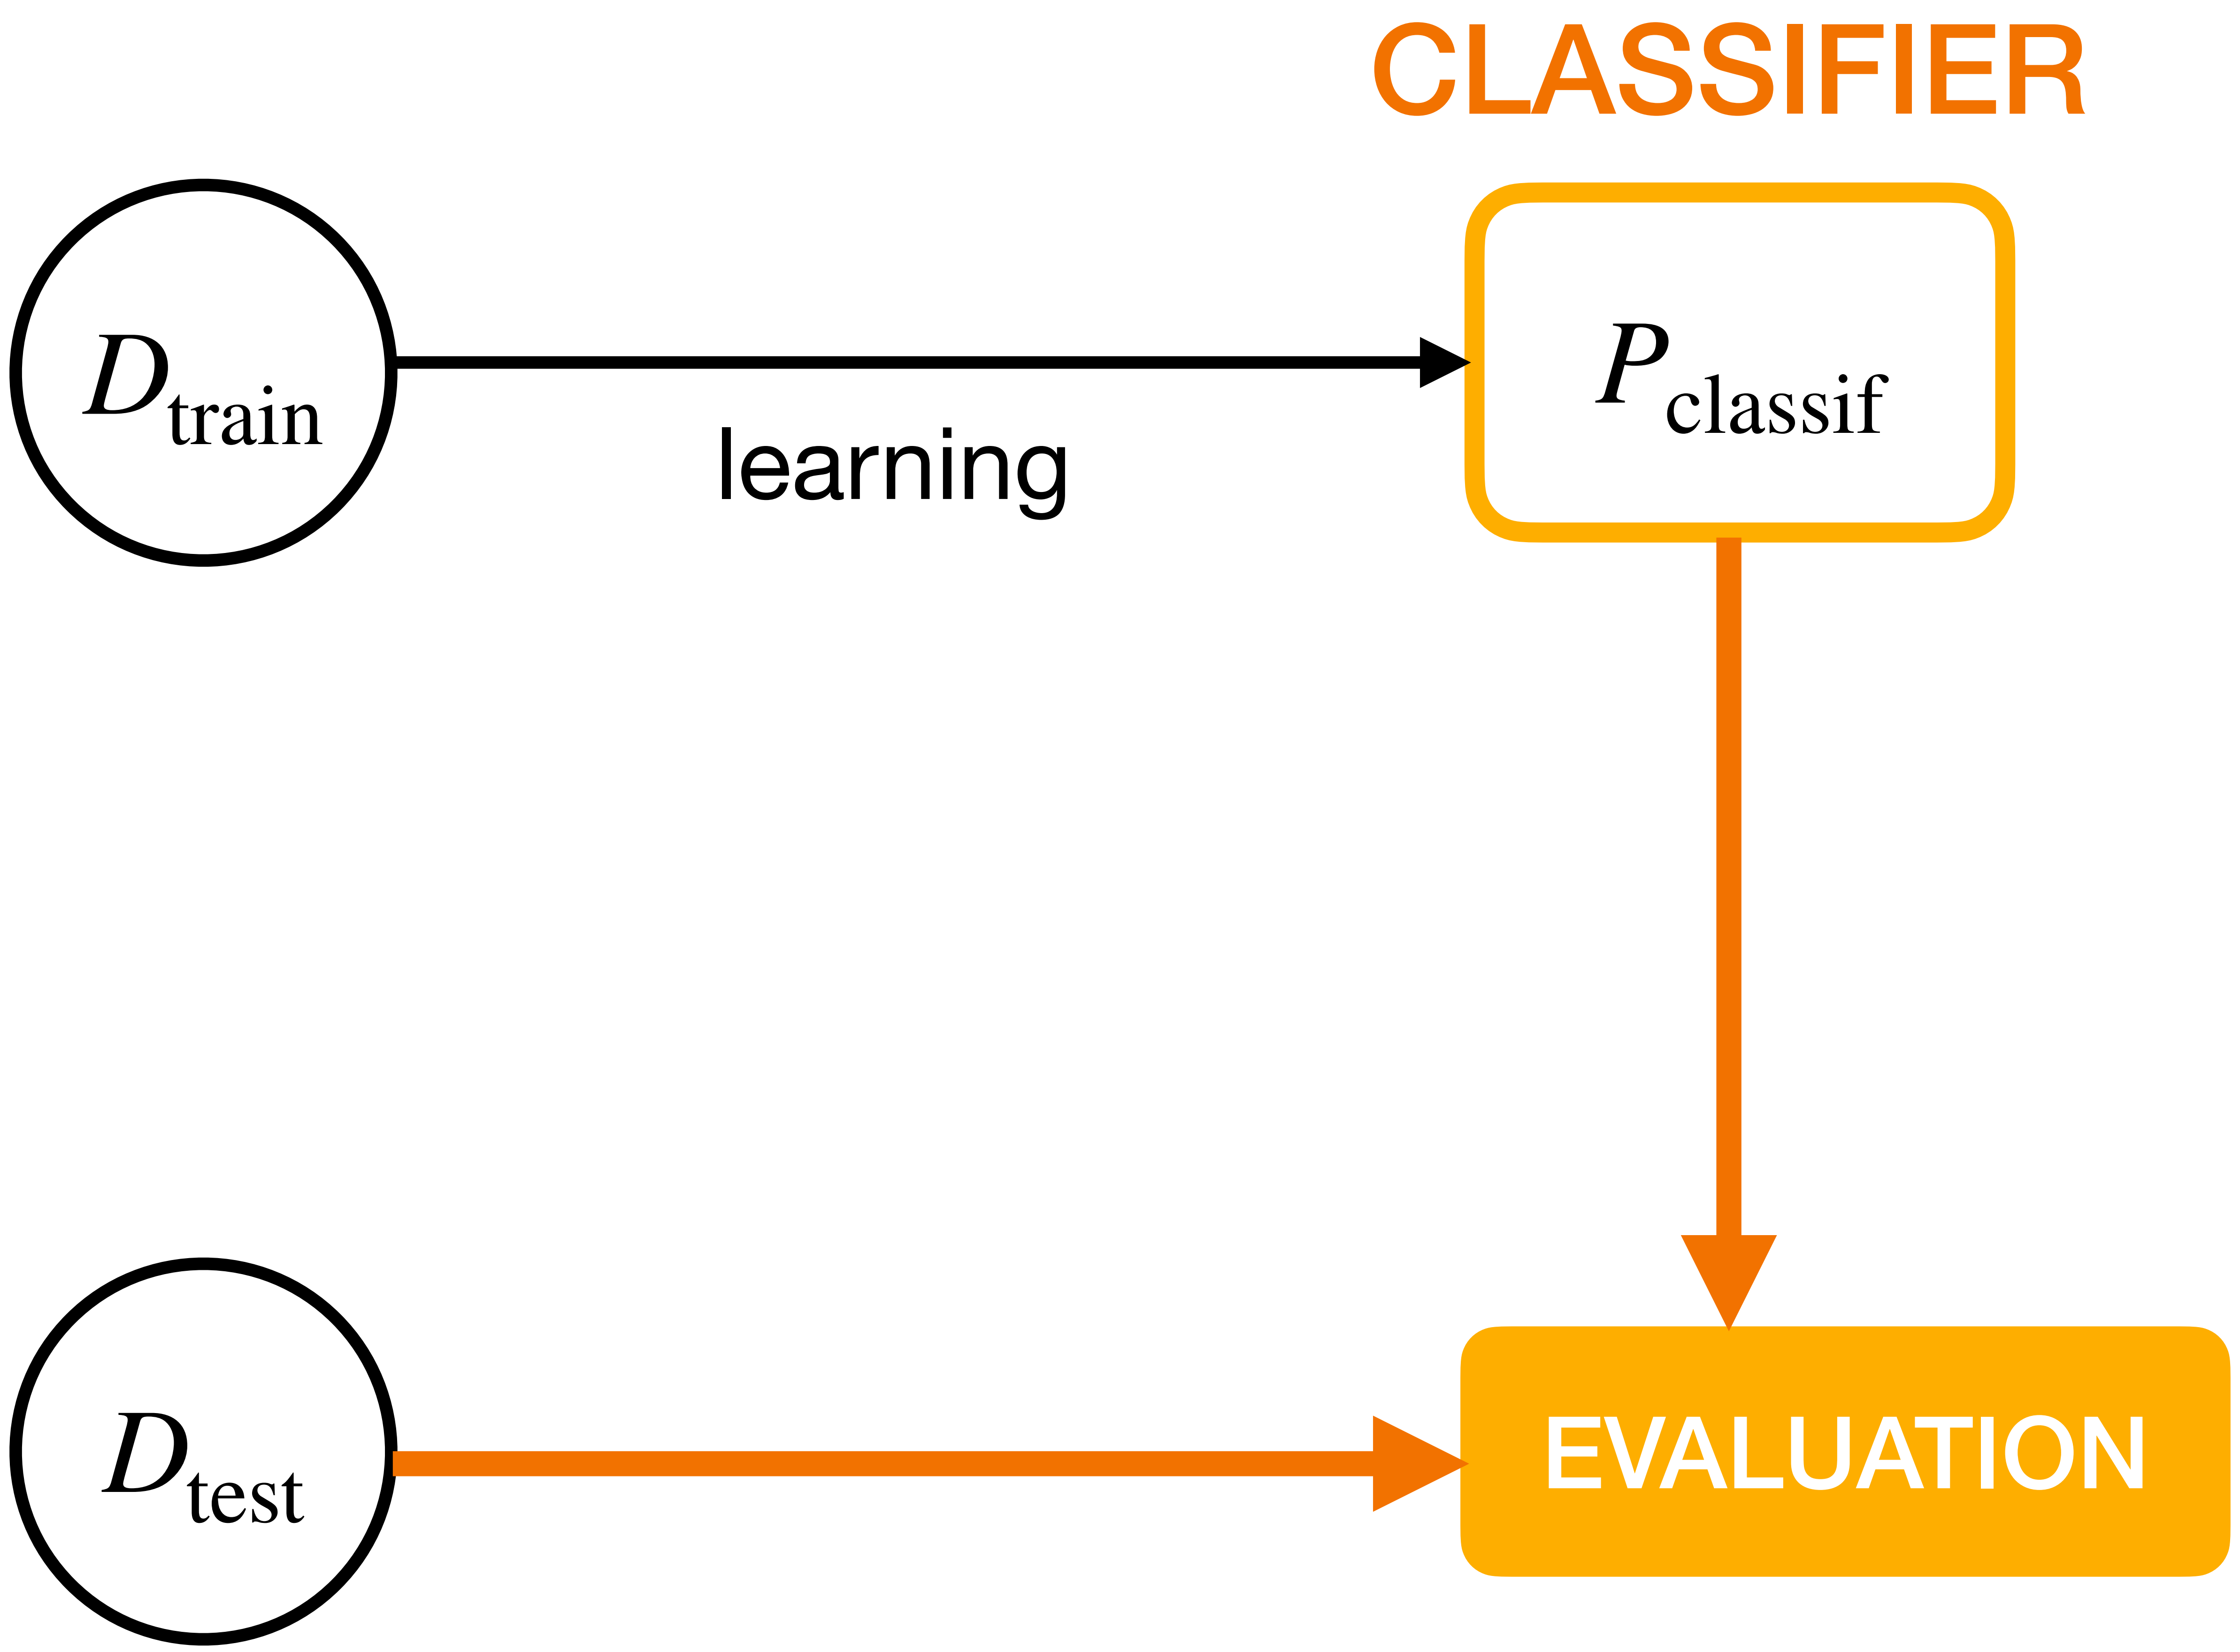

CLASSIFICATION

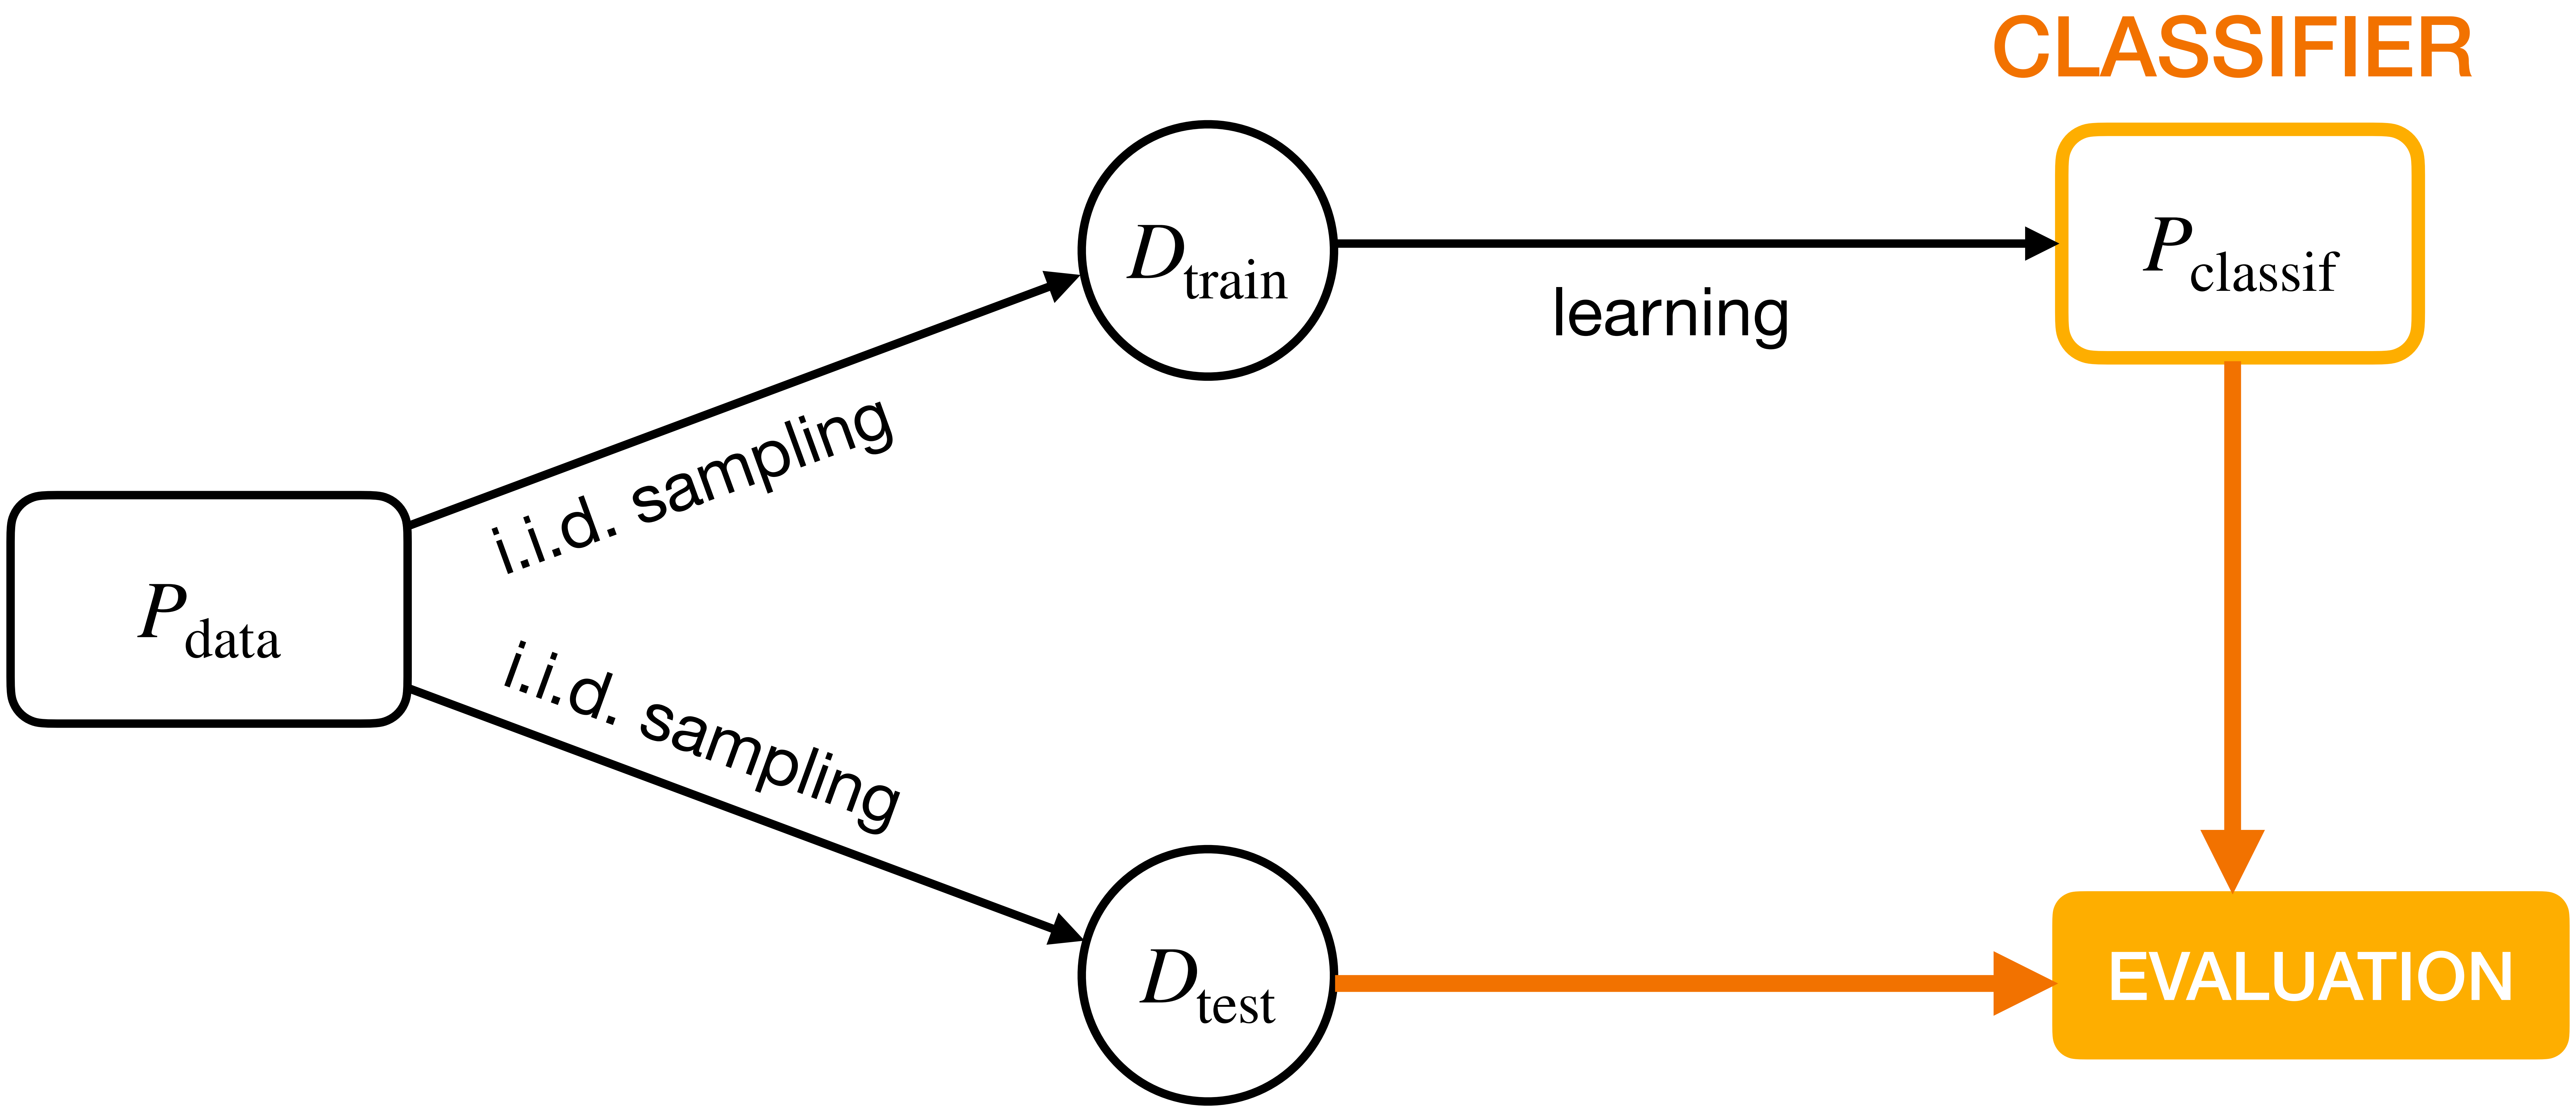

CLASSIFICATION

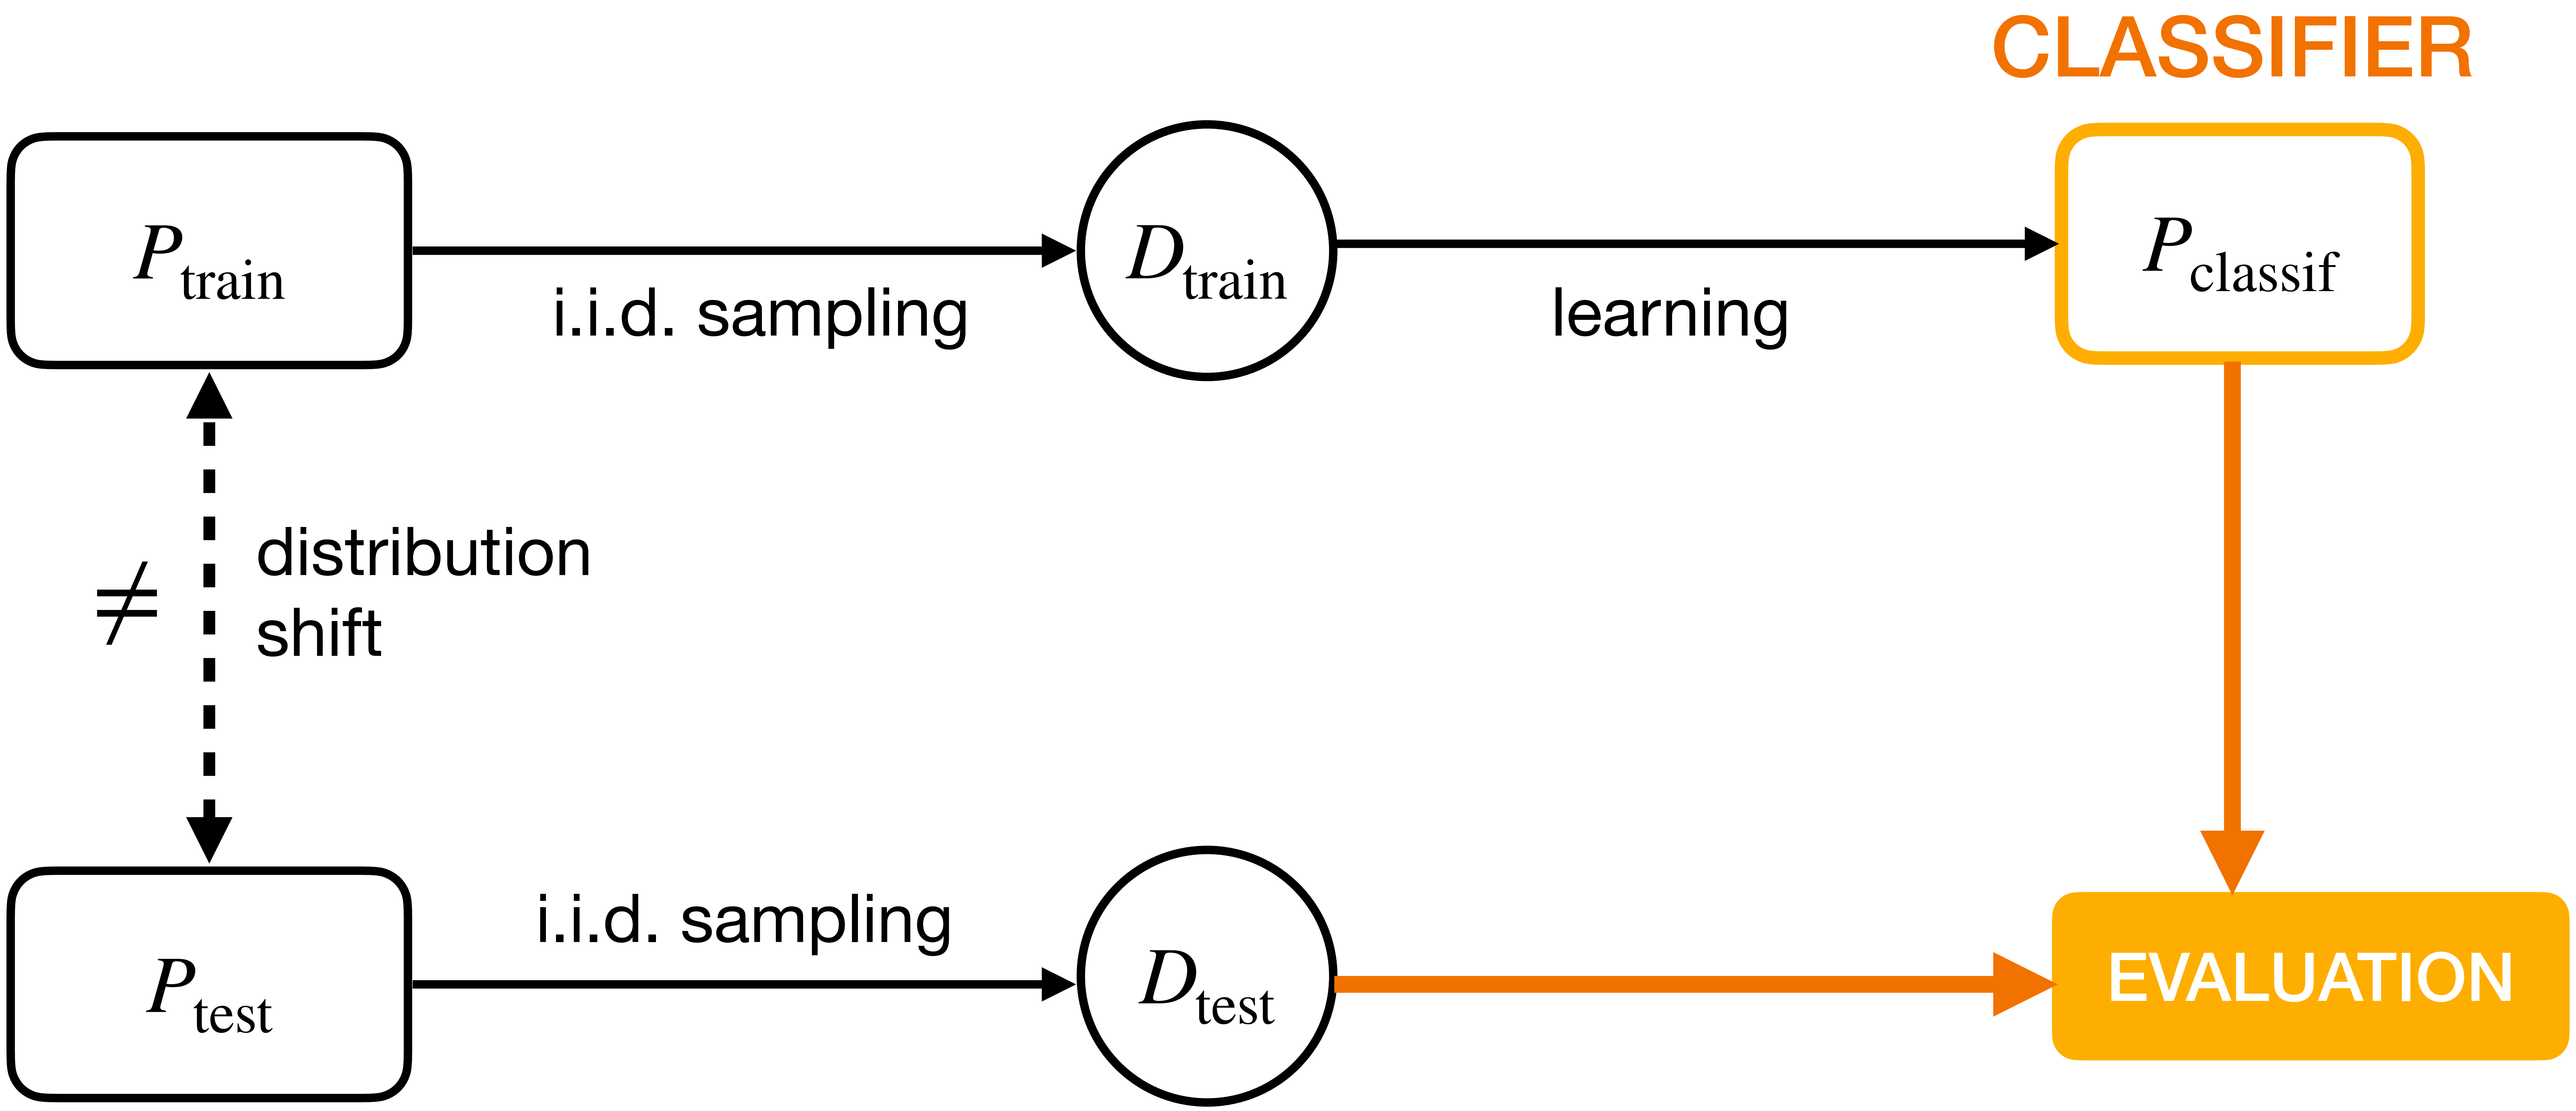

CLASSIFICATION ... is unreliable

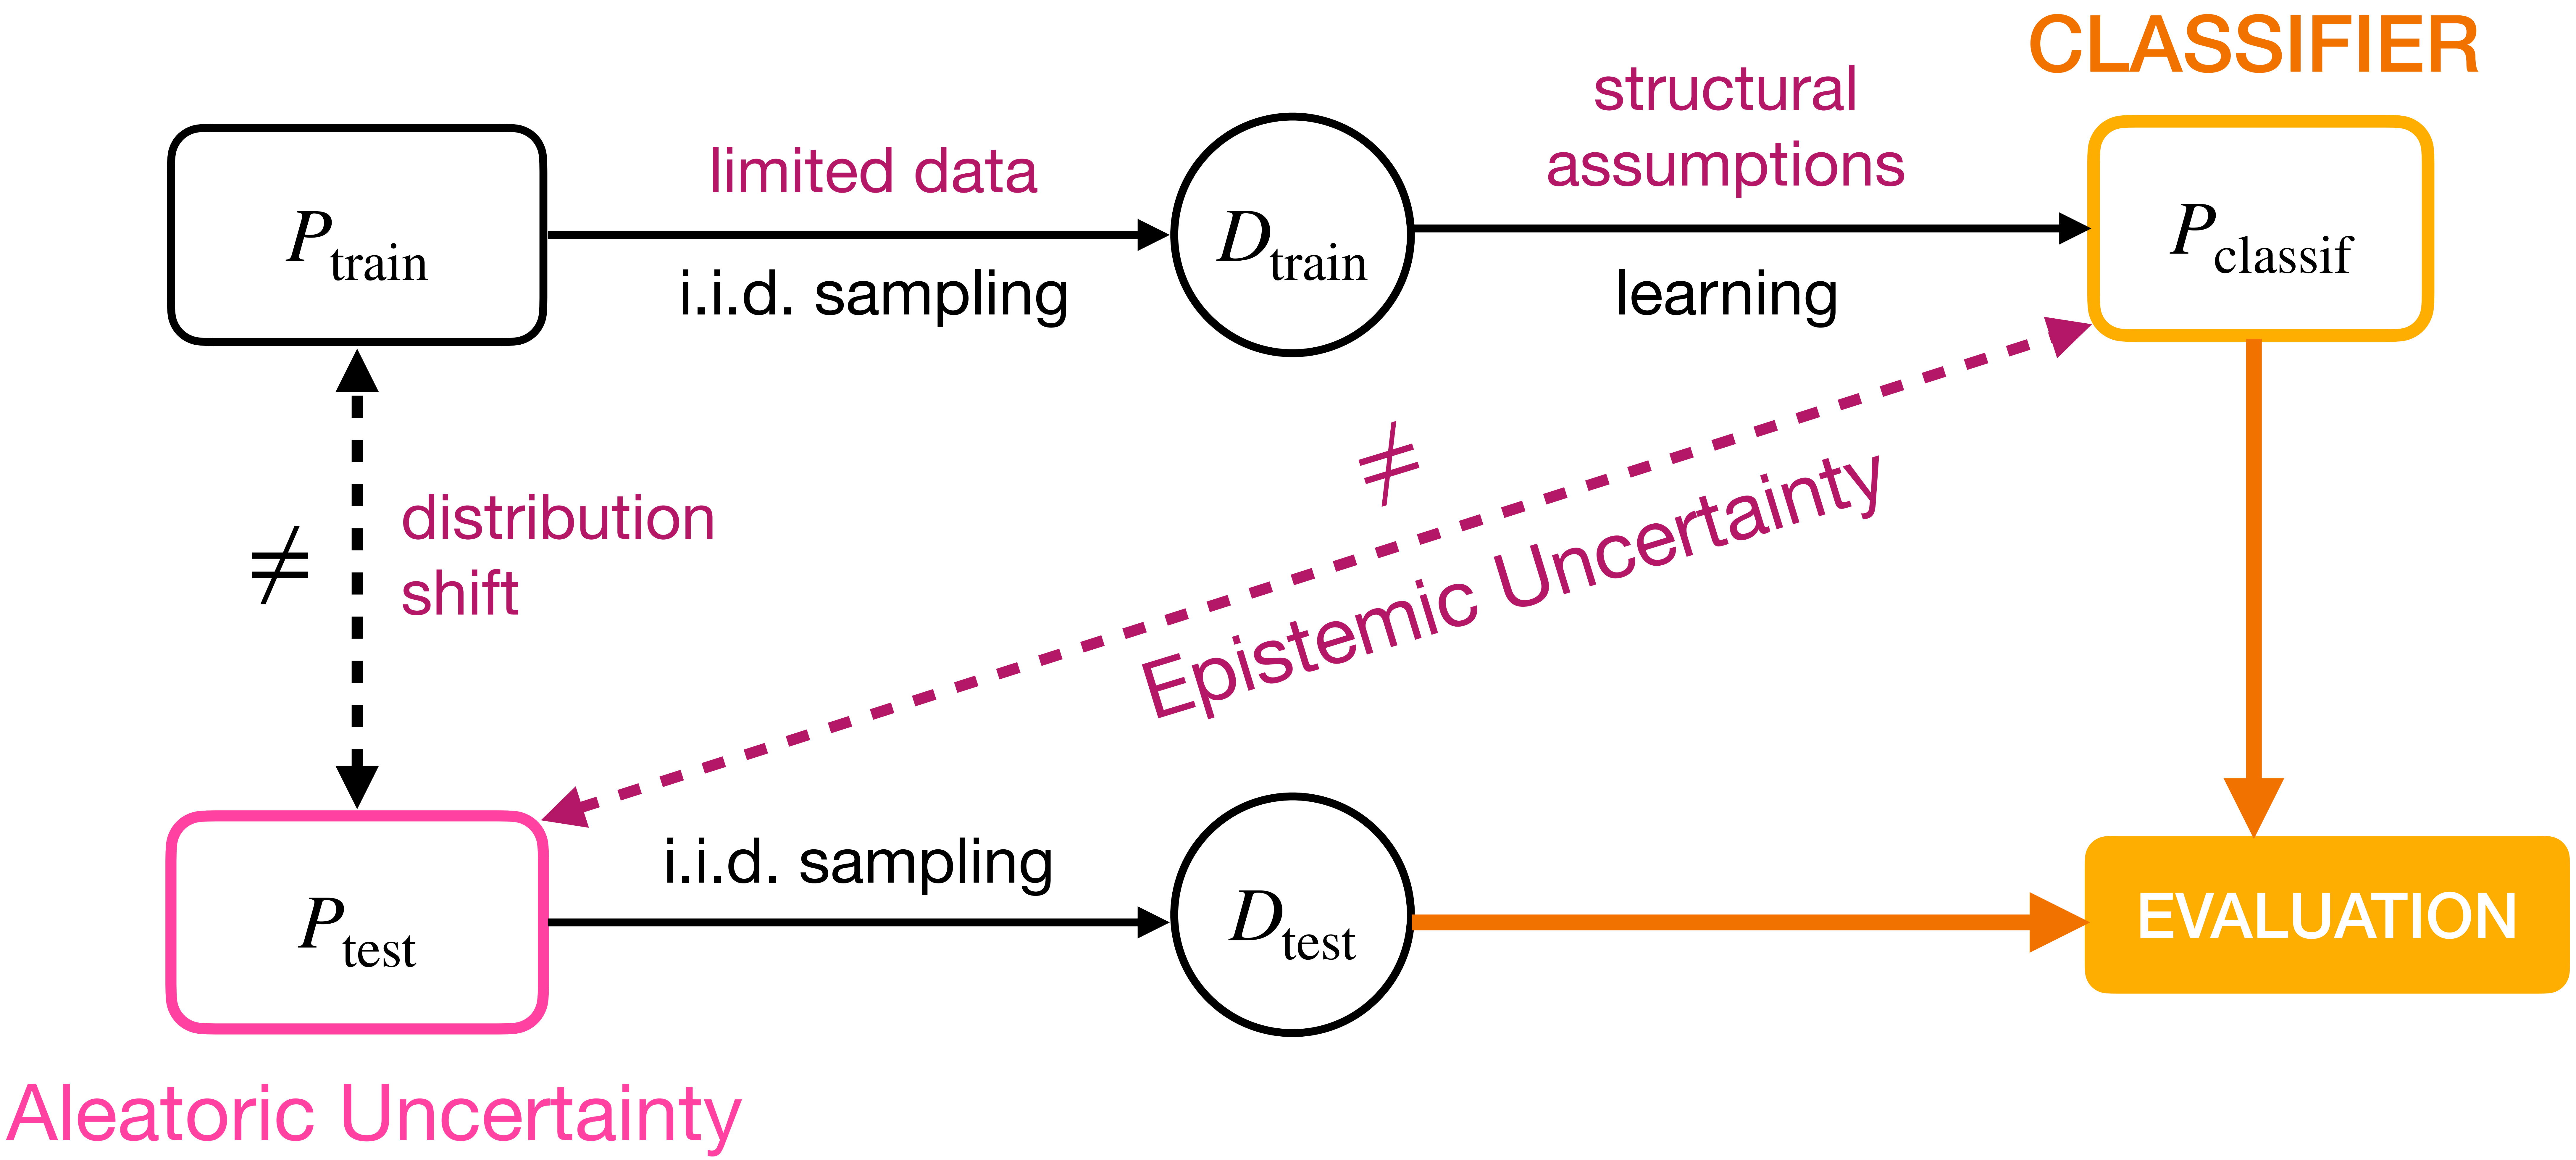

PROBLEM 2:

IMPRECISE
PROBABILITIES

… are arbitrary

IMPRECISE PROBABILITIES

PROBABILISTIC CLASSIFIER

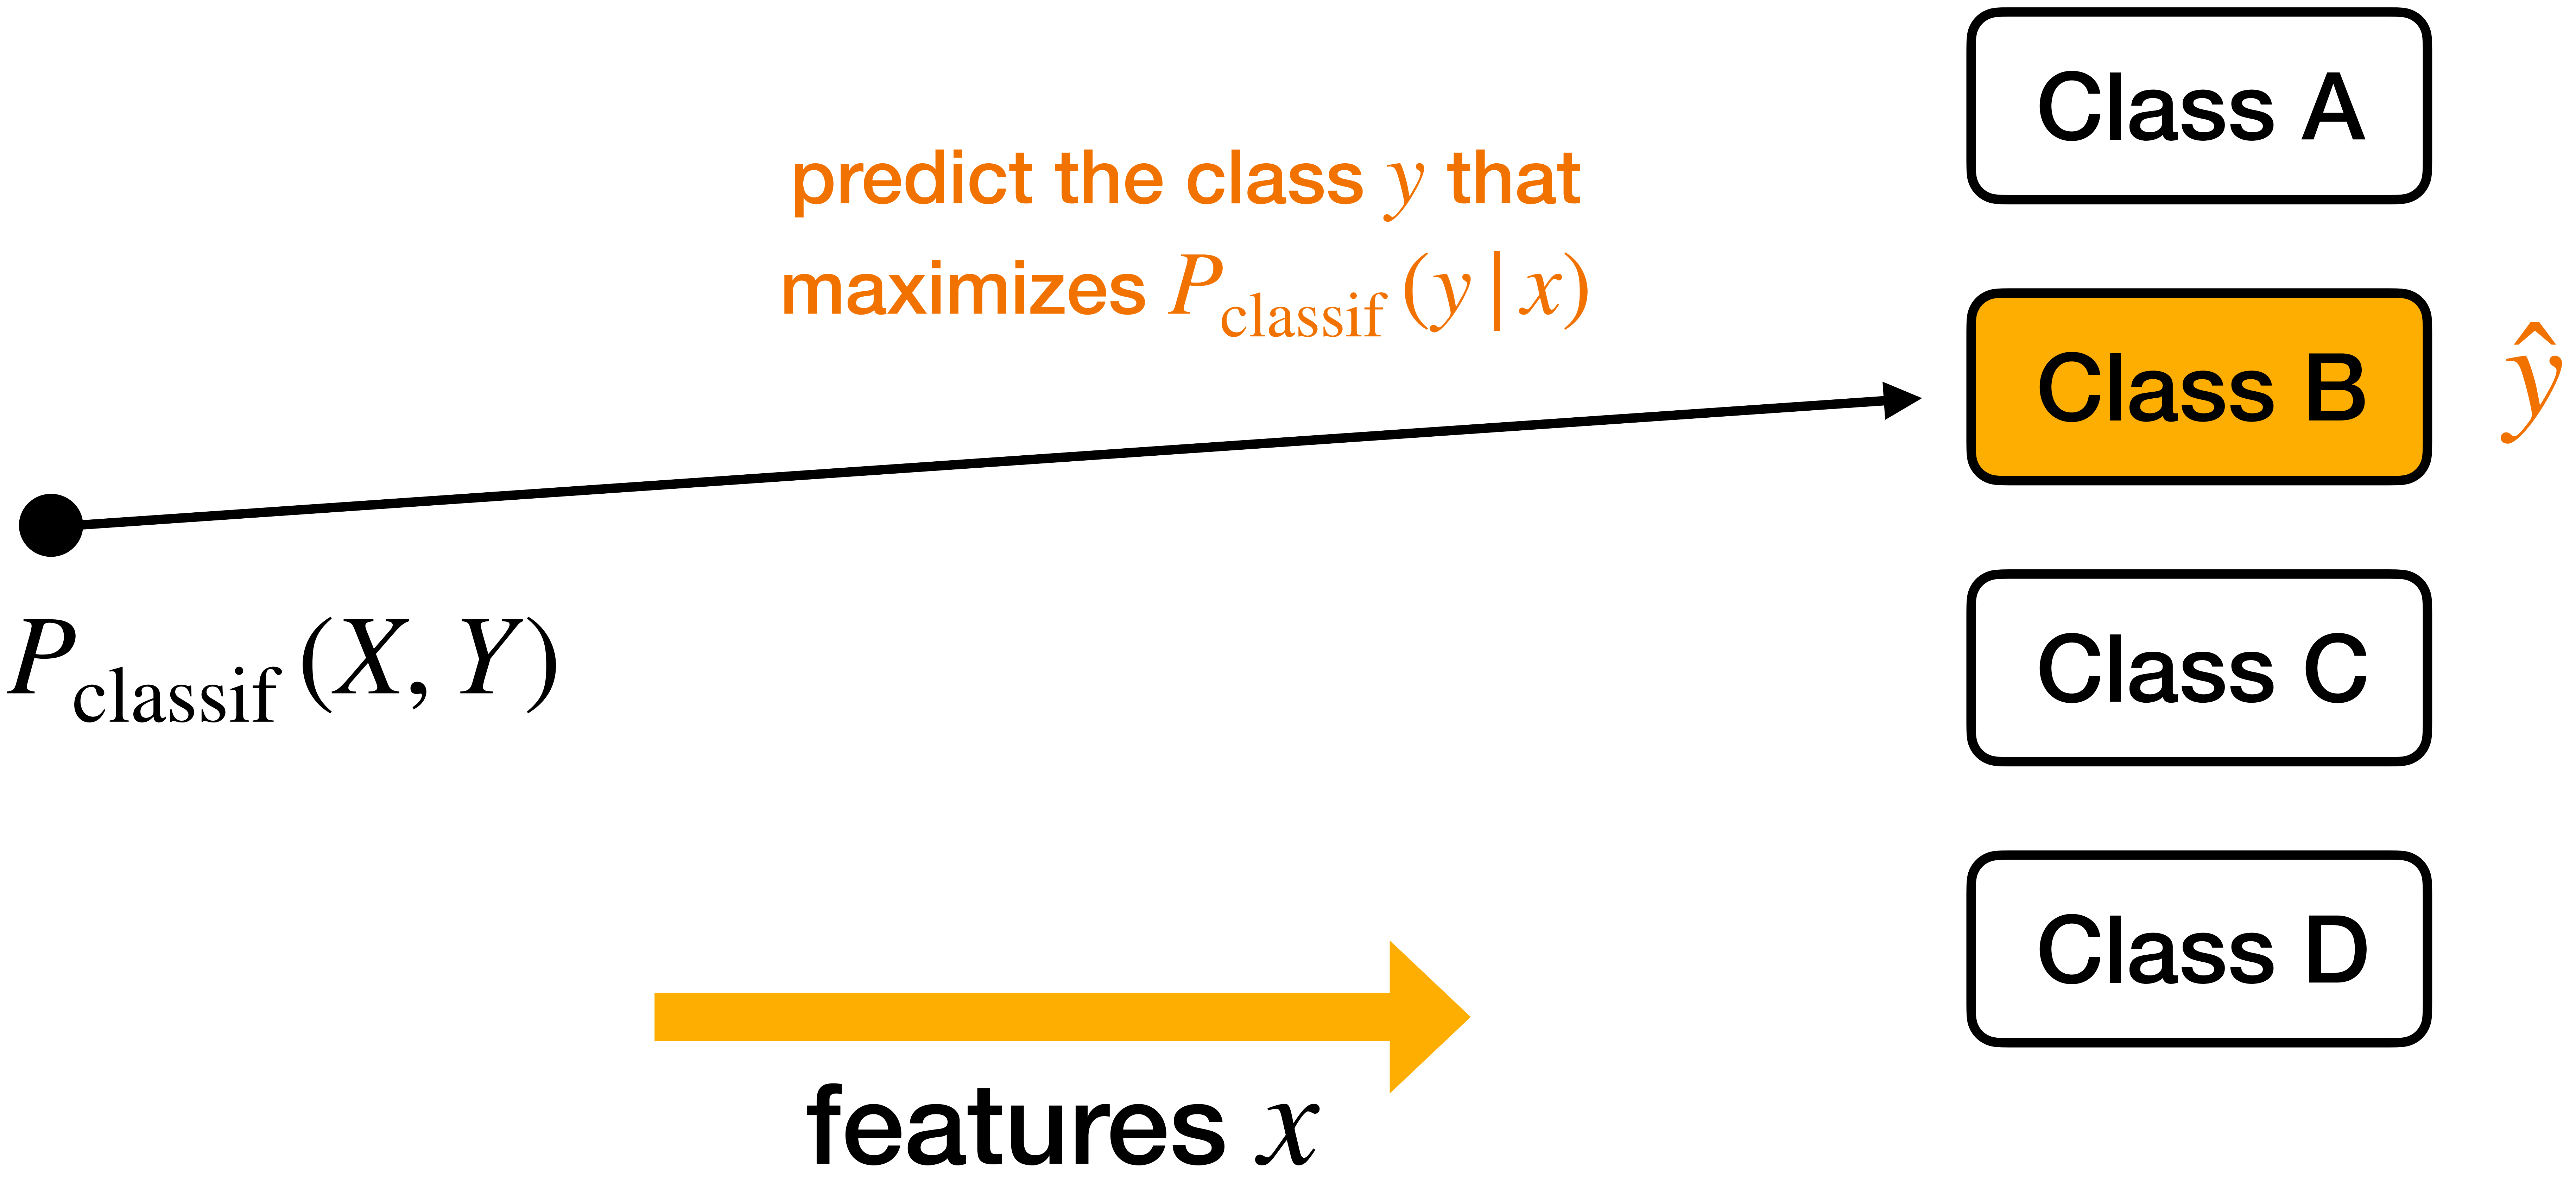

IMPRECISE PROBABILITIES

CREDAL CLASSIFIER

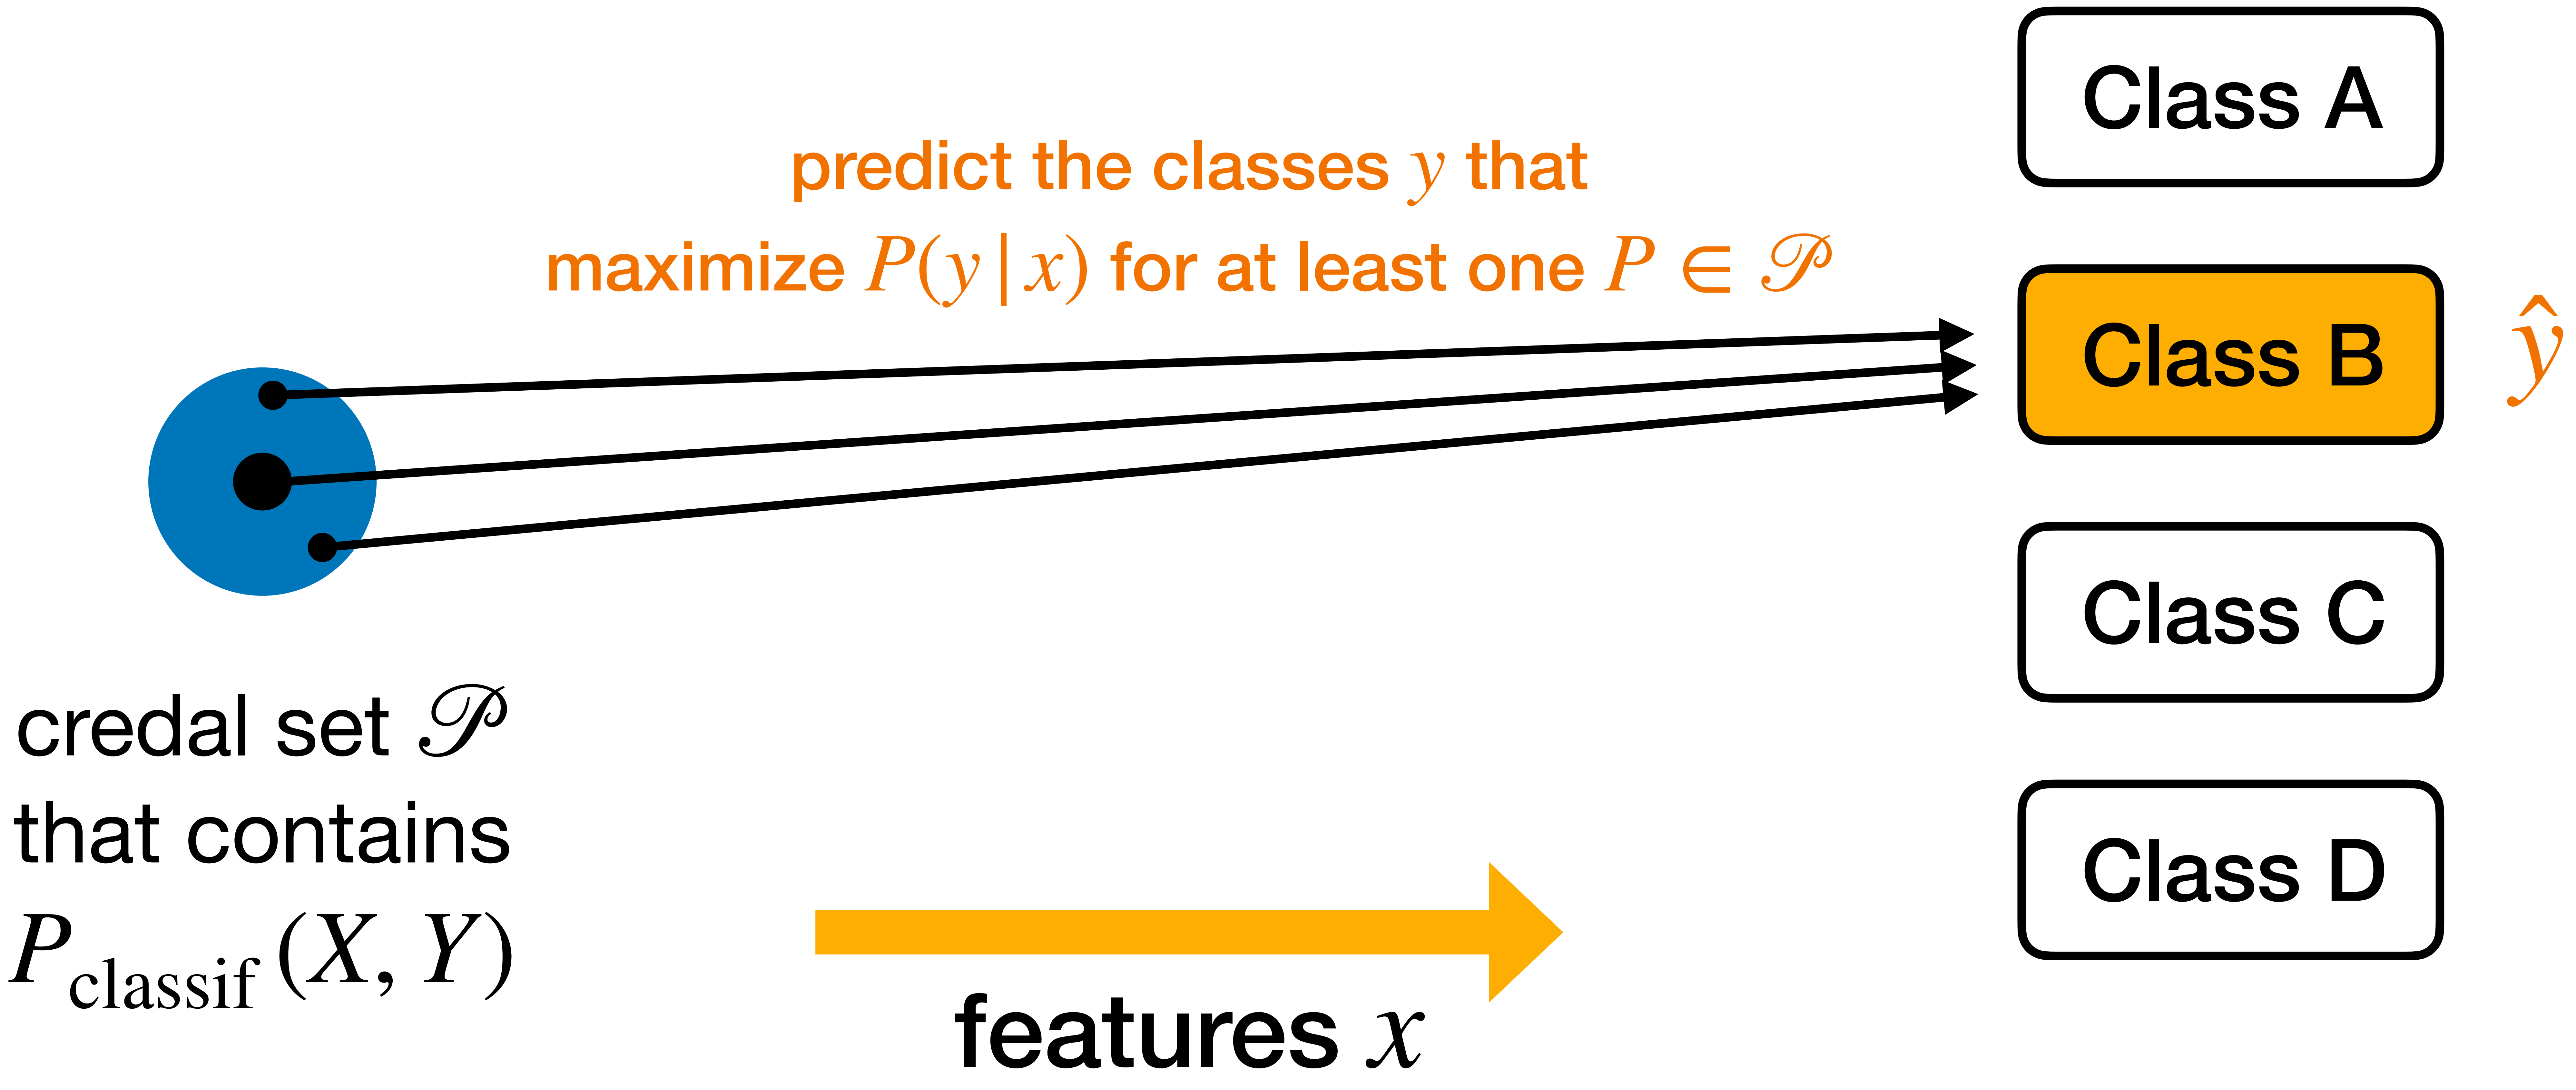

IMPRECISE PROBABILITIES

CREDAL CLASSIFIER

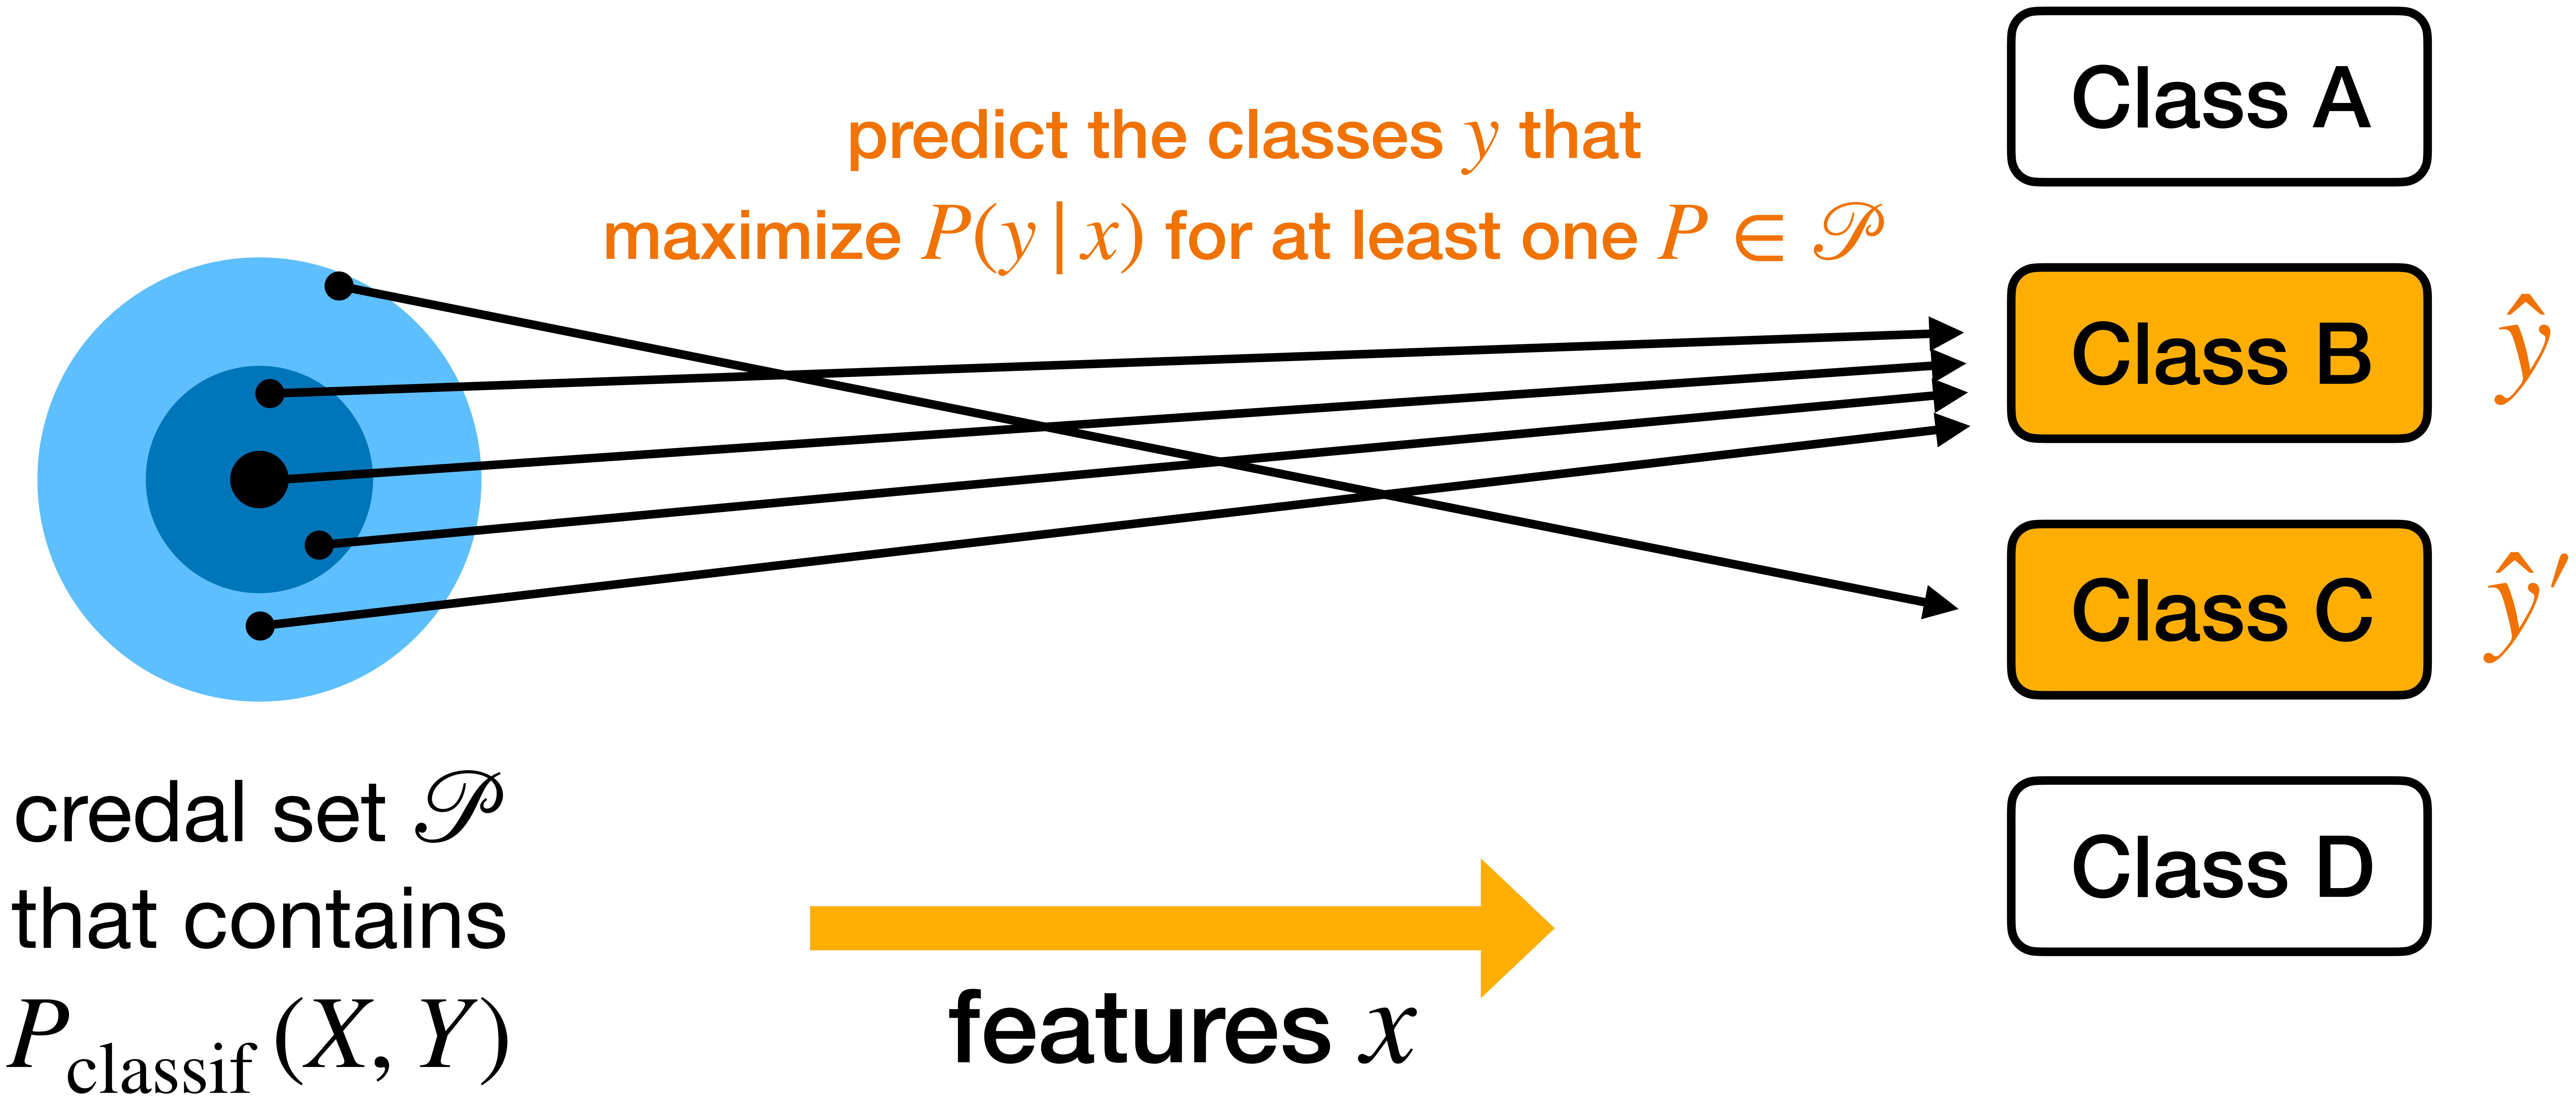

IMPRECISE PROBABILITIES

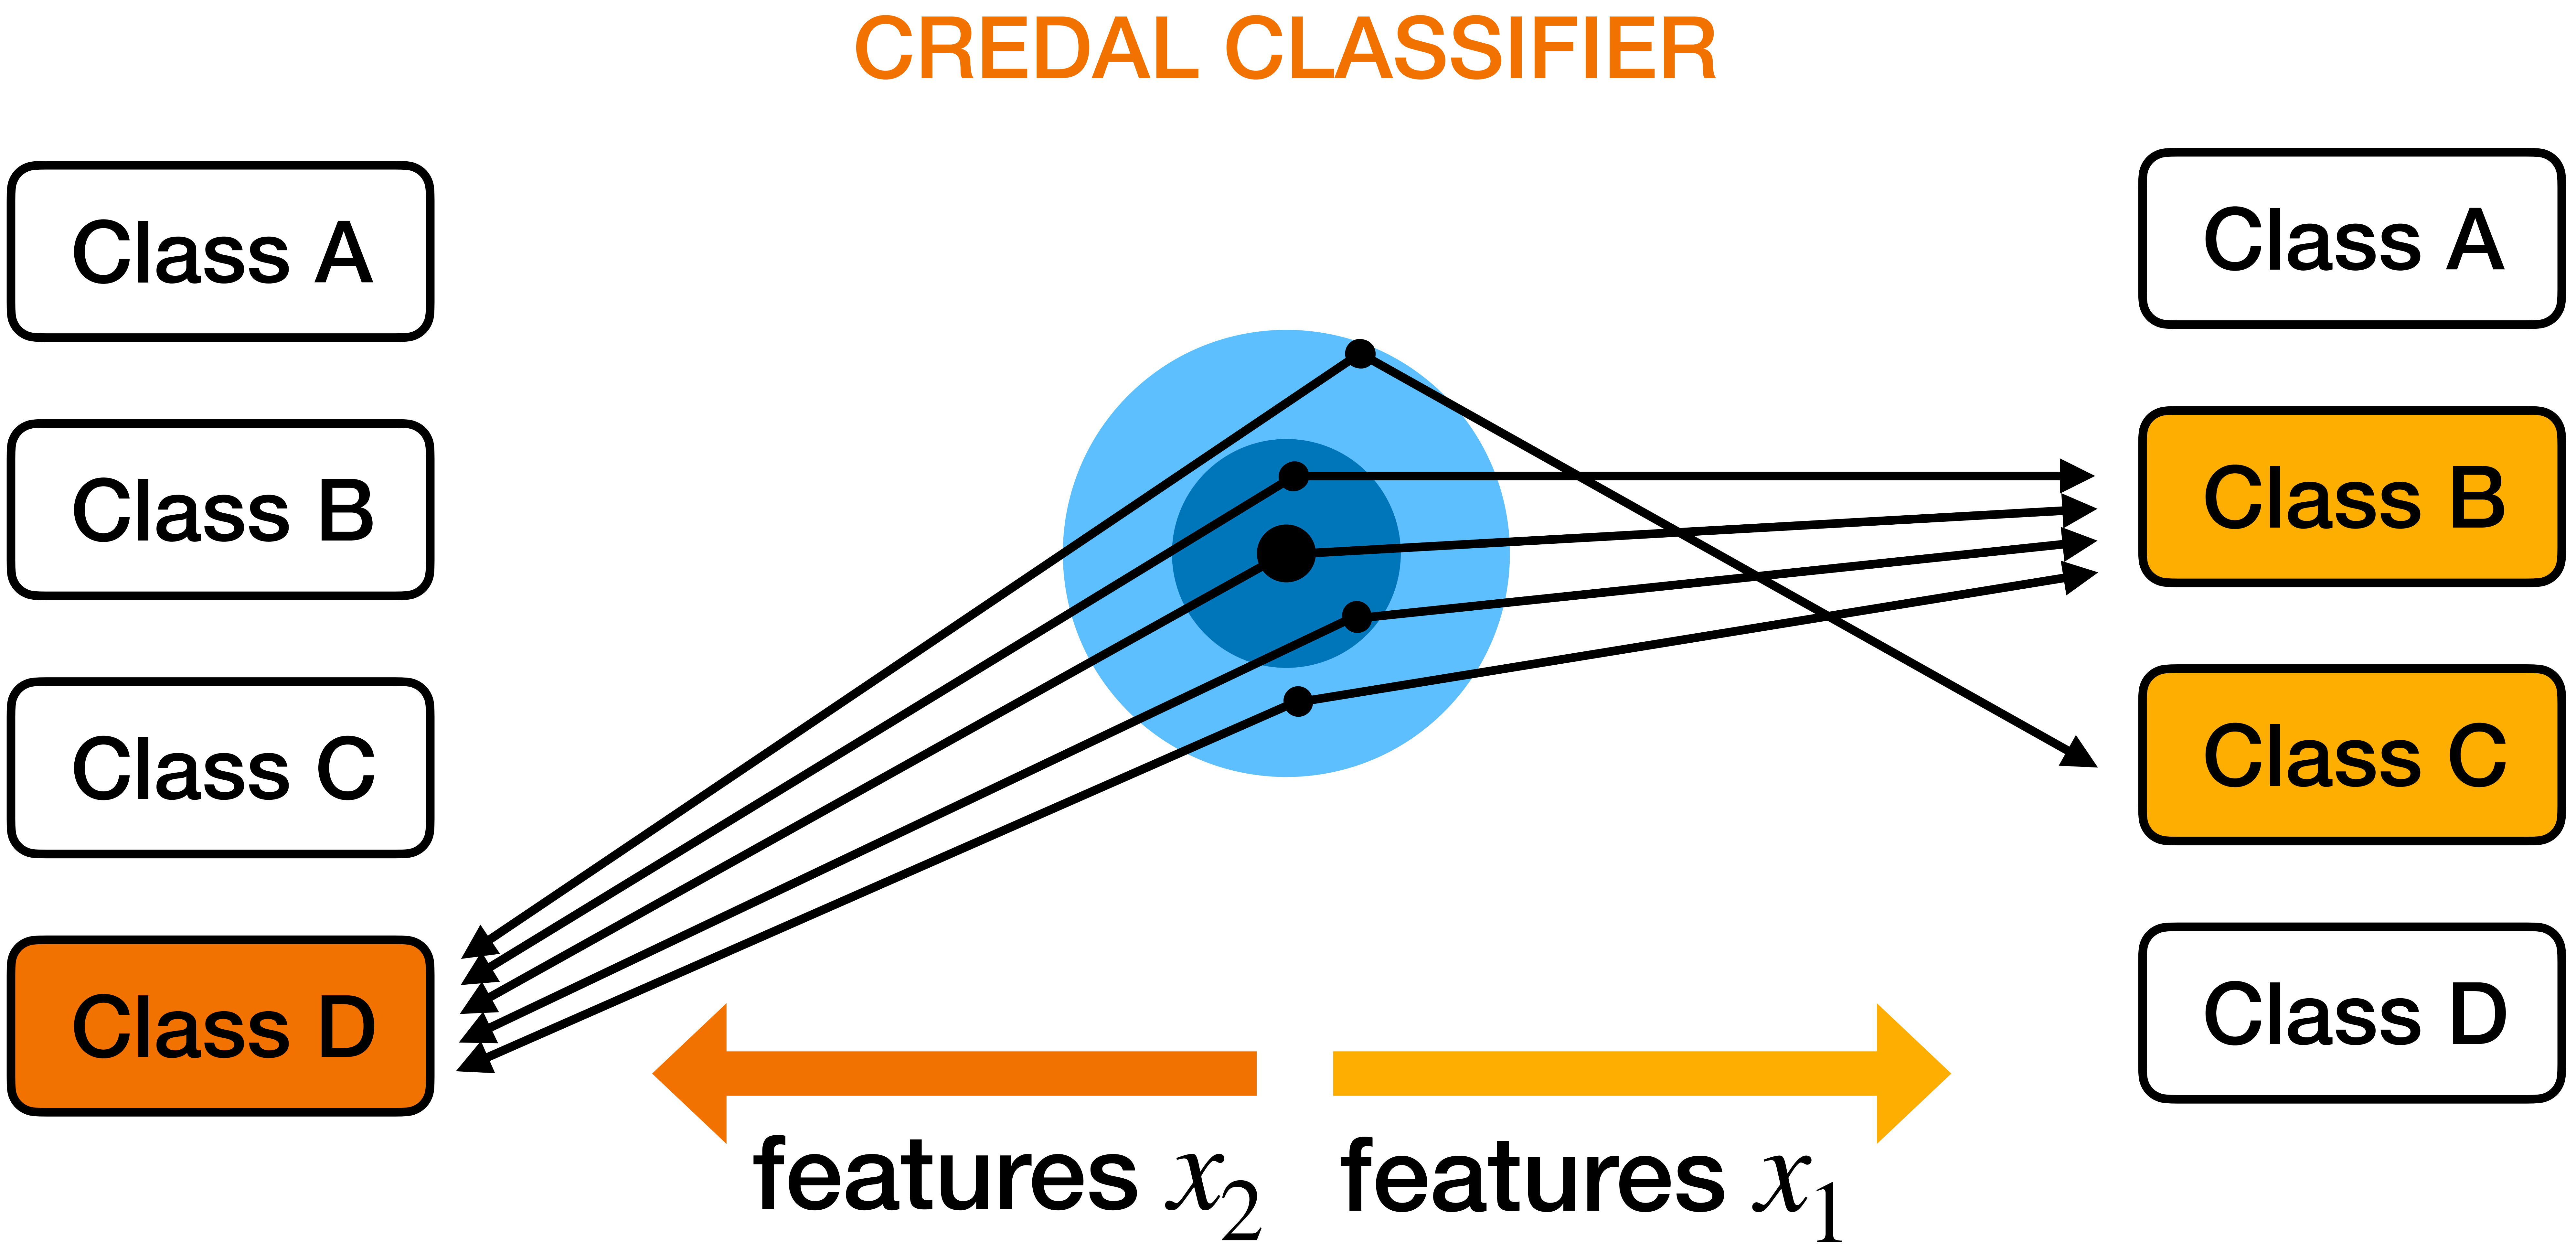

IMPRECISE PROBABILITIES

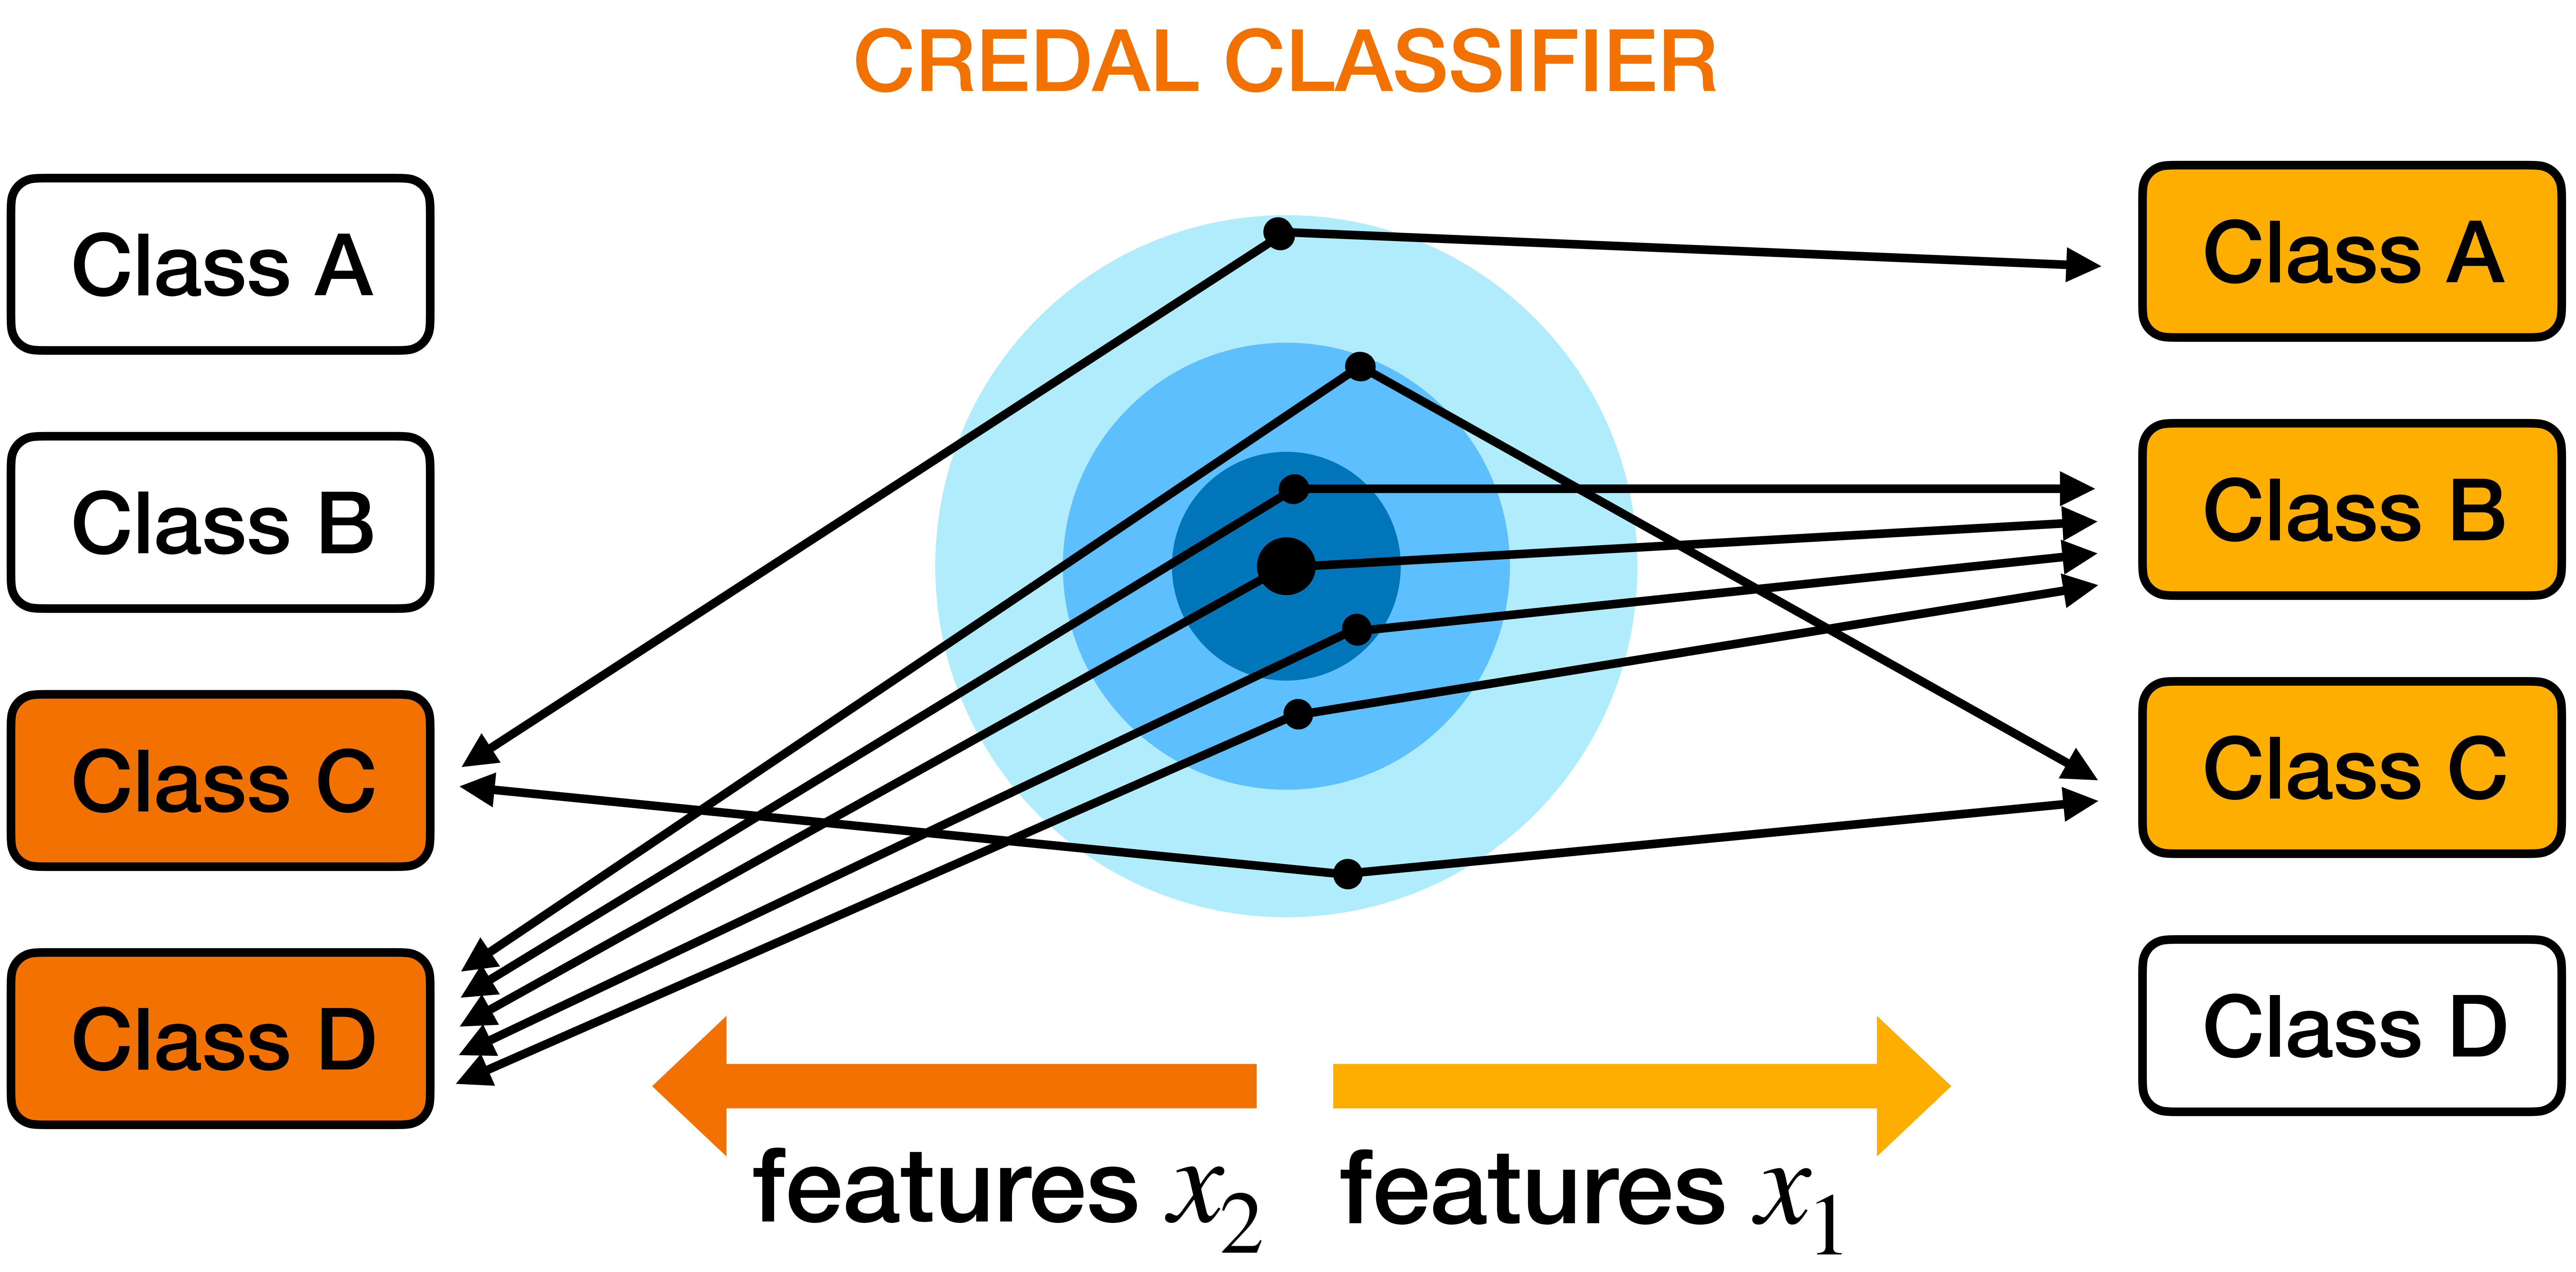

IMPRECISE PROBABILITIES ... are arbitrary

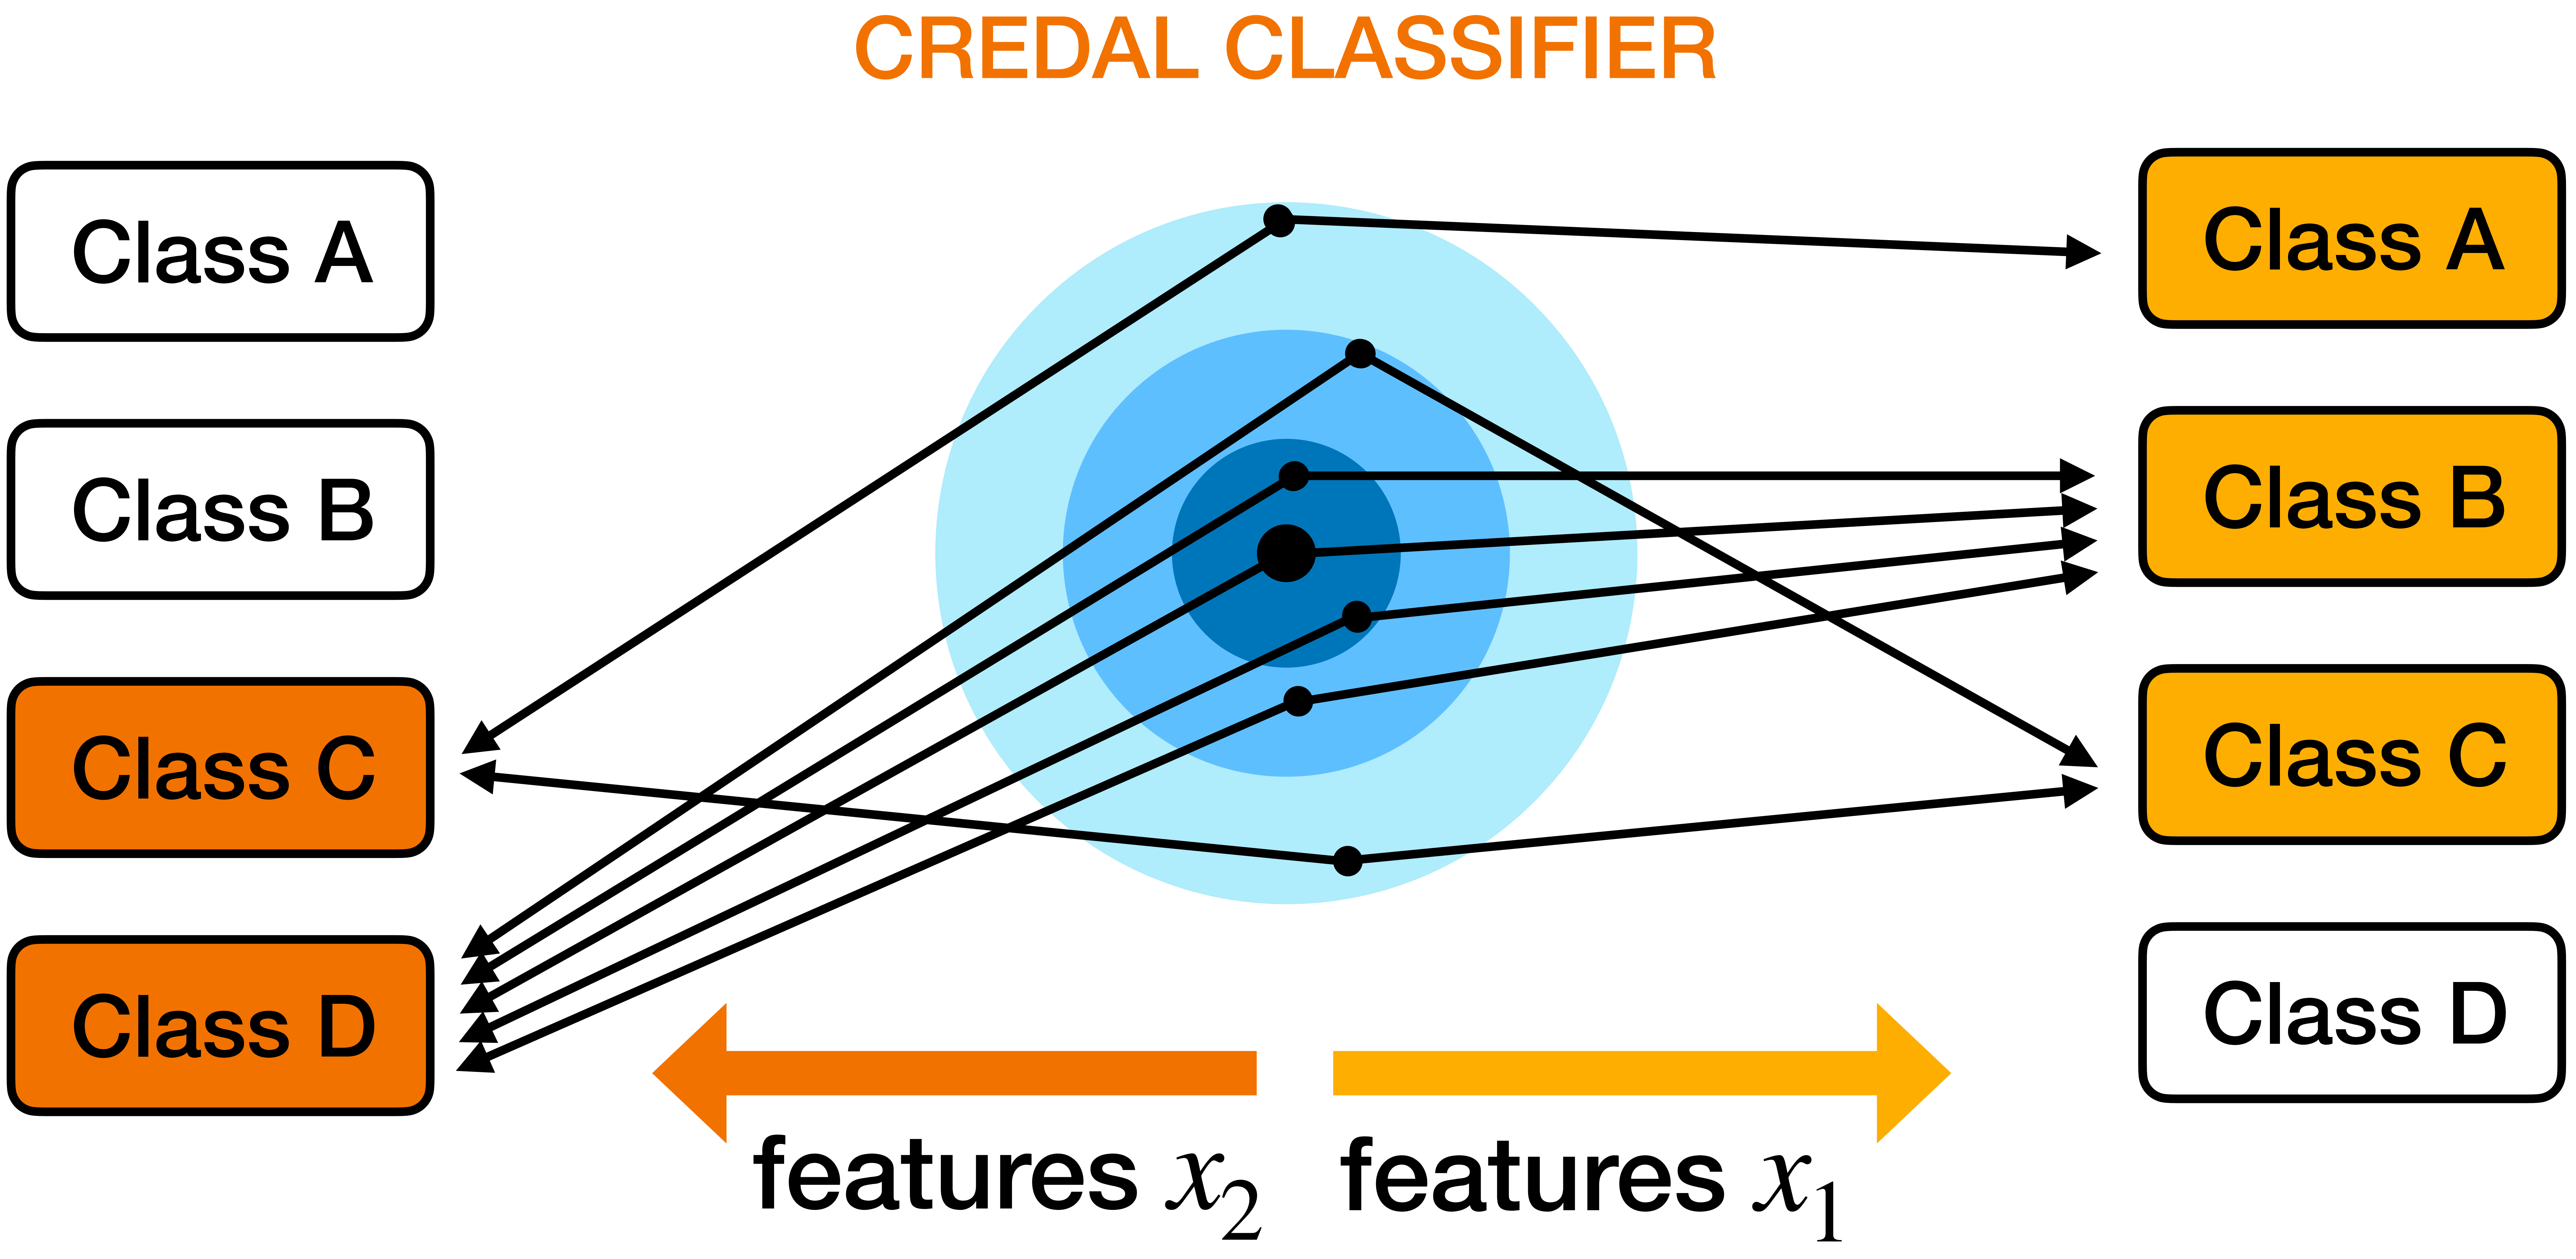

Adrián Detavernier

Rodrigo Lassance

MACHINE LEARNING

IMPRECISE PROBABILITIES

ROBUSTNESS QUANTIFICATION

ROBUSTNESS QUANTIFICATION

PROBABILISTIC CLASSIFIER

predict the class y that maximizes $P_{\text{classif}}(y \mid x)$

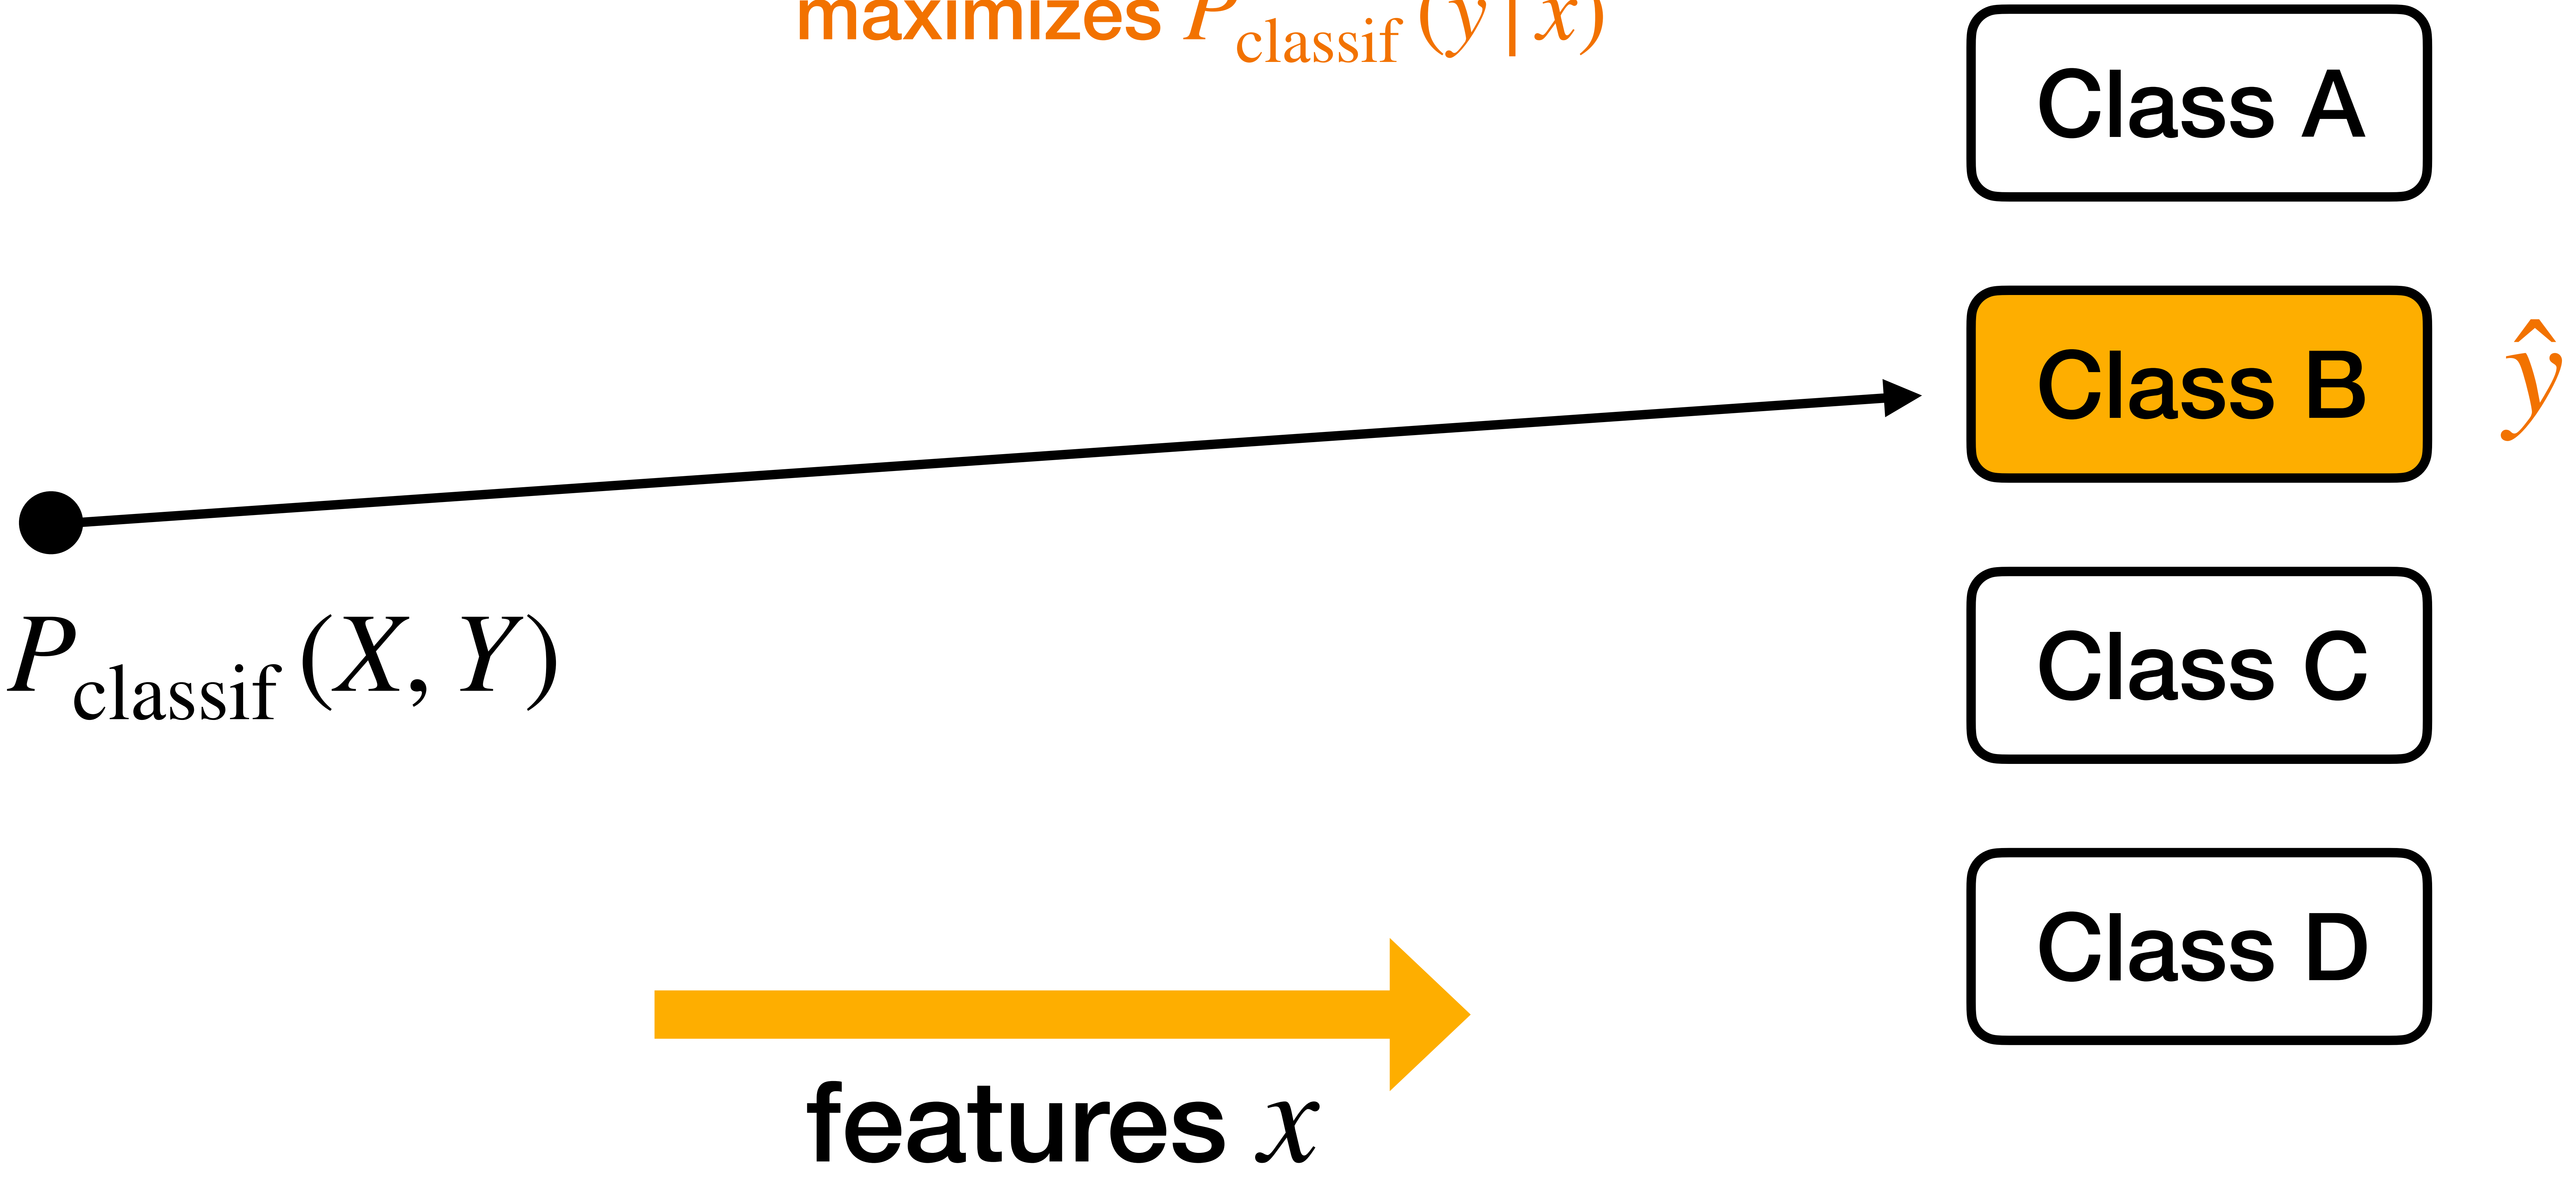

ROBUSTNESS QUANTIFICATION

ROBUST prediction

PROBABILISTIC CLASSIFIER

predict the class y that maximizes $P_{\text{classif}}(y \mid x)$

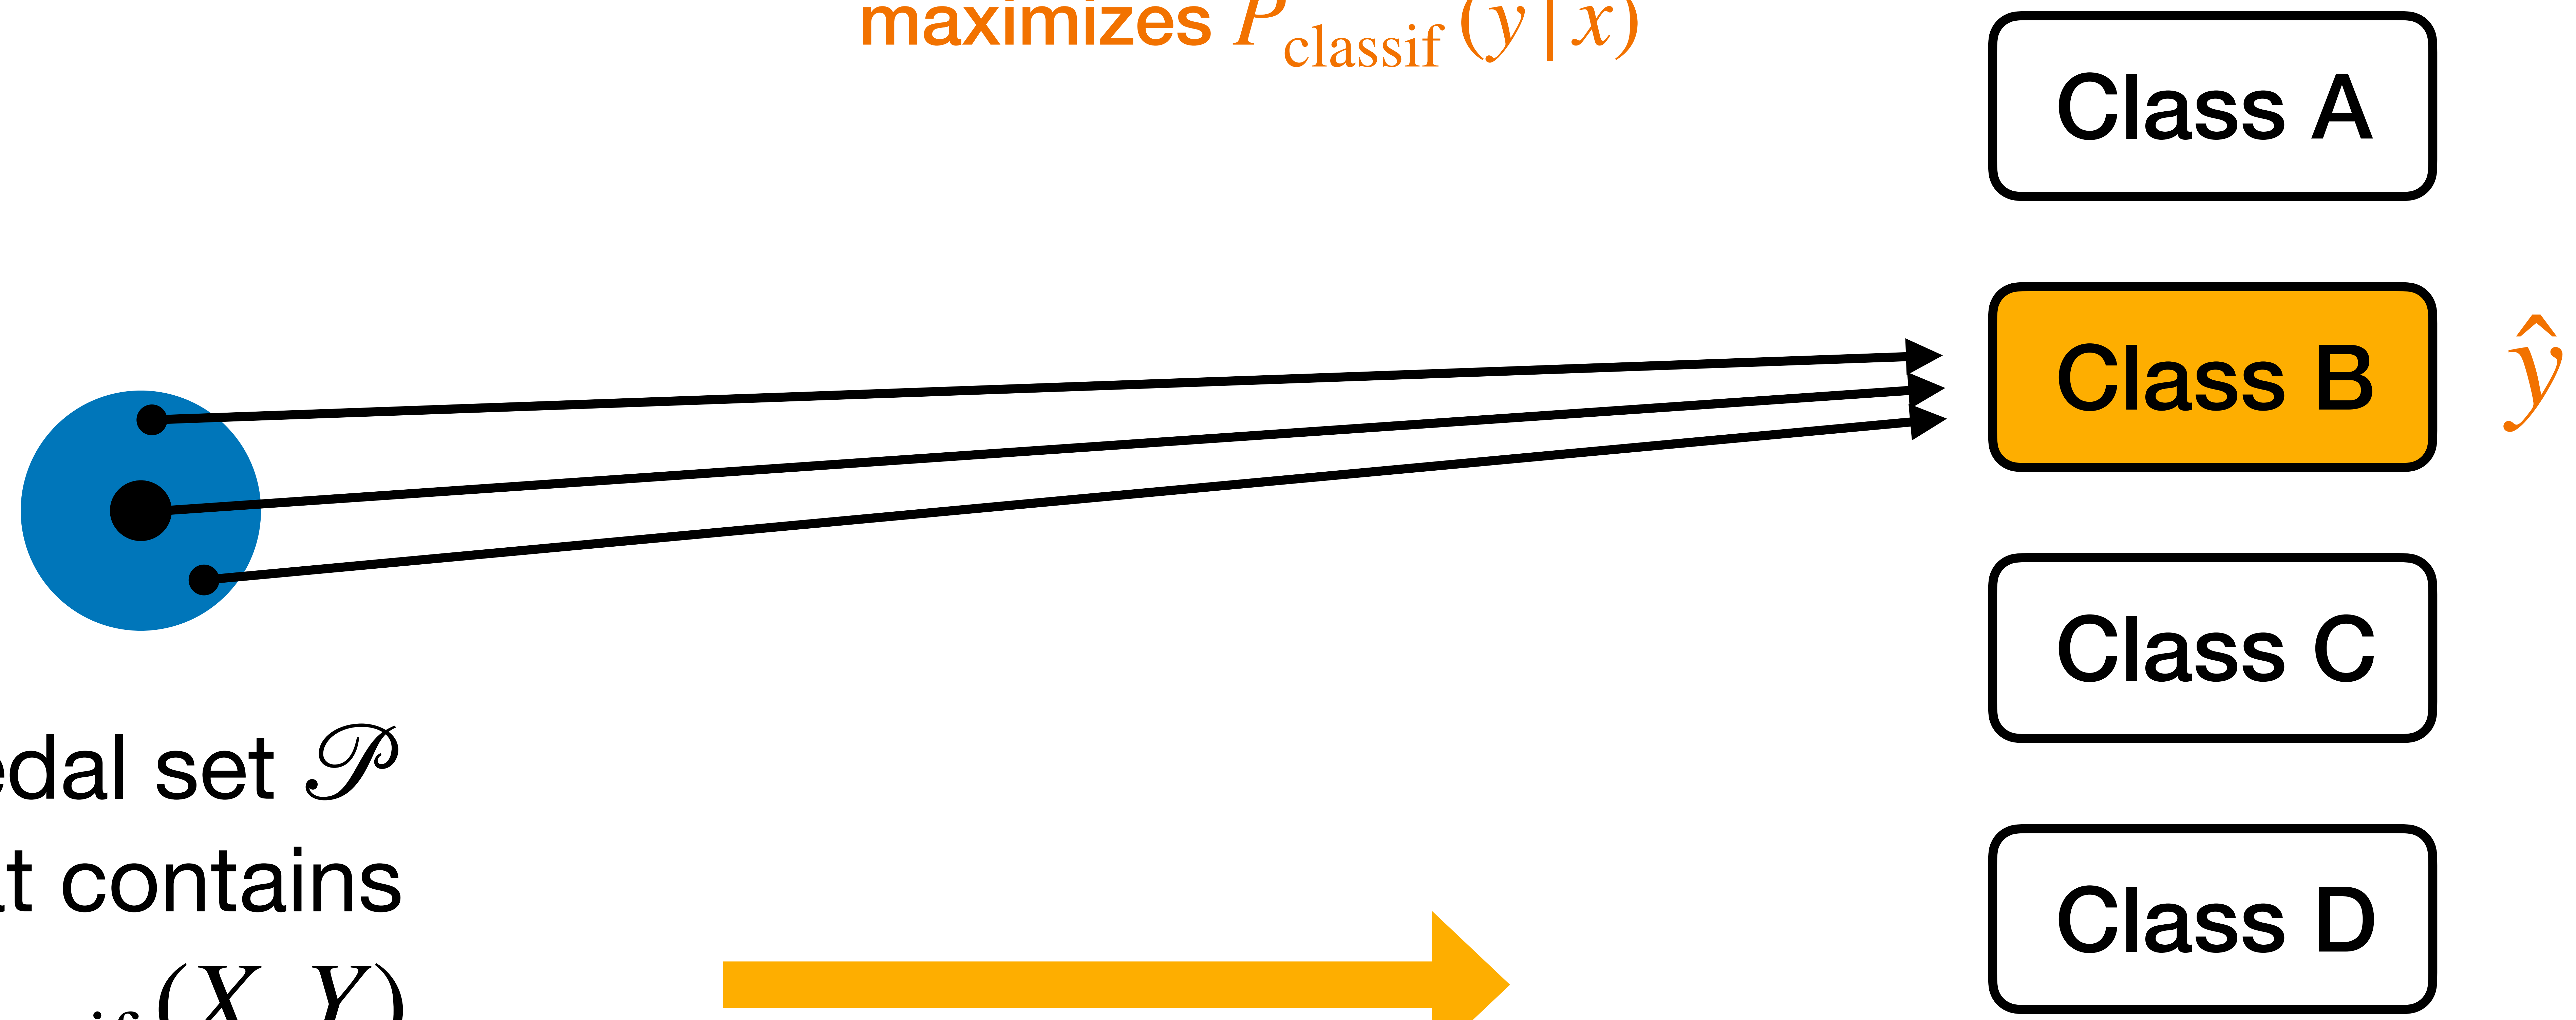

credal set $\mathscr{P}$ that contains $P_{\text{classif}}(X, Y)$

features x

ROBUSTNESS QUANTIFICATION

~~ROBUST prediction~~

PROBABILISTIC CLASSIFIER

predict the class y that maximizes $P_{\text{classif}}(y \mid x)$

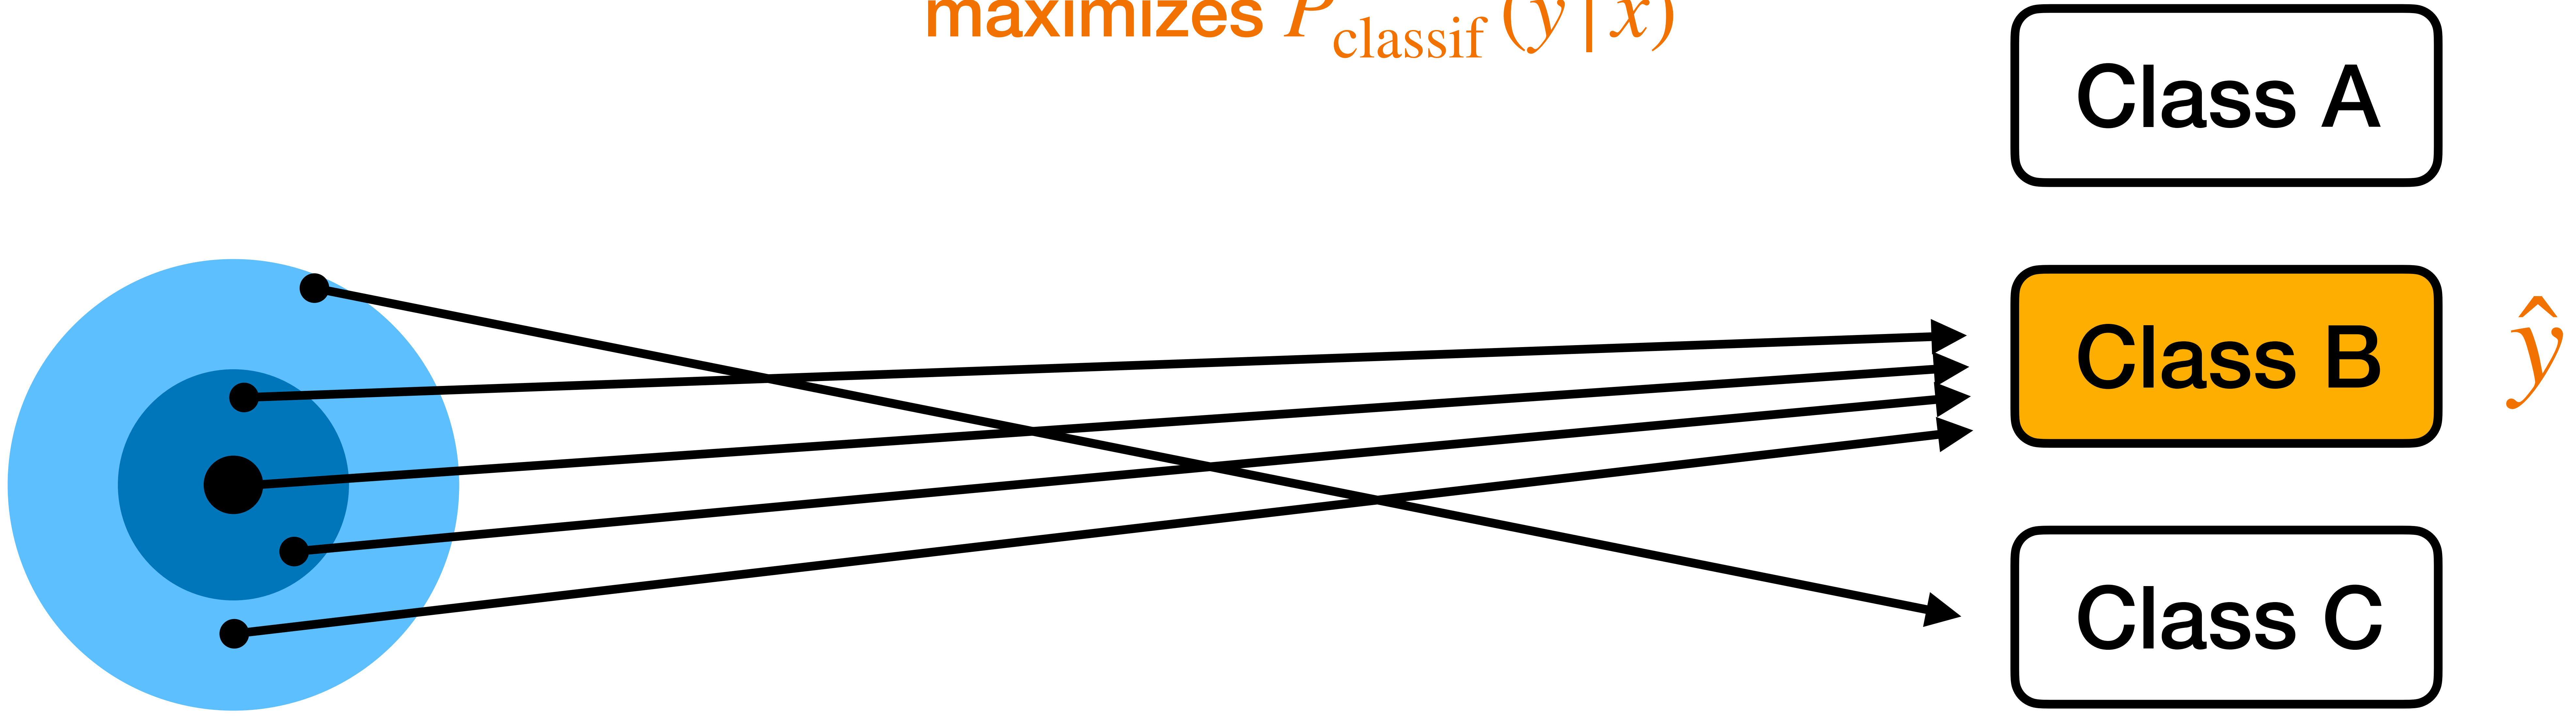

credal set $\mathscr{P}$ that contains $P_{\text{classif}}(X, Y)$

features x

ROBUSTNESS QUANTIFICATION

~~ROBUST prediction~~

PROBABILISTIC CLASSIFIER

predict the class y that maximizes $P_{\text{classif}}(y \mid x)$

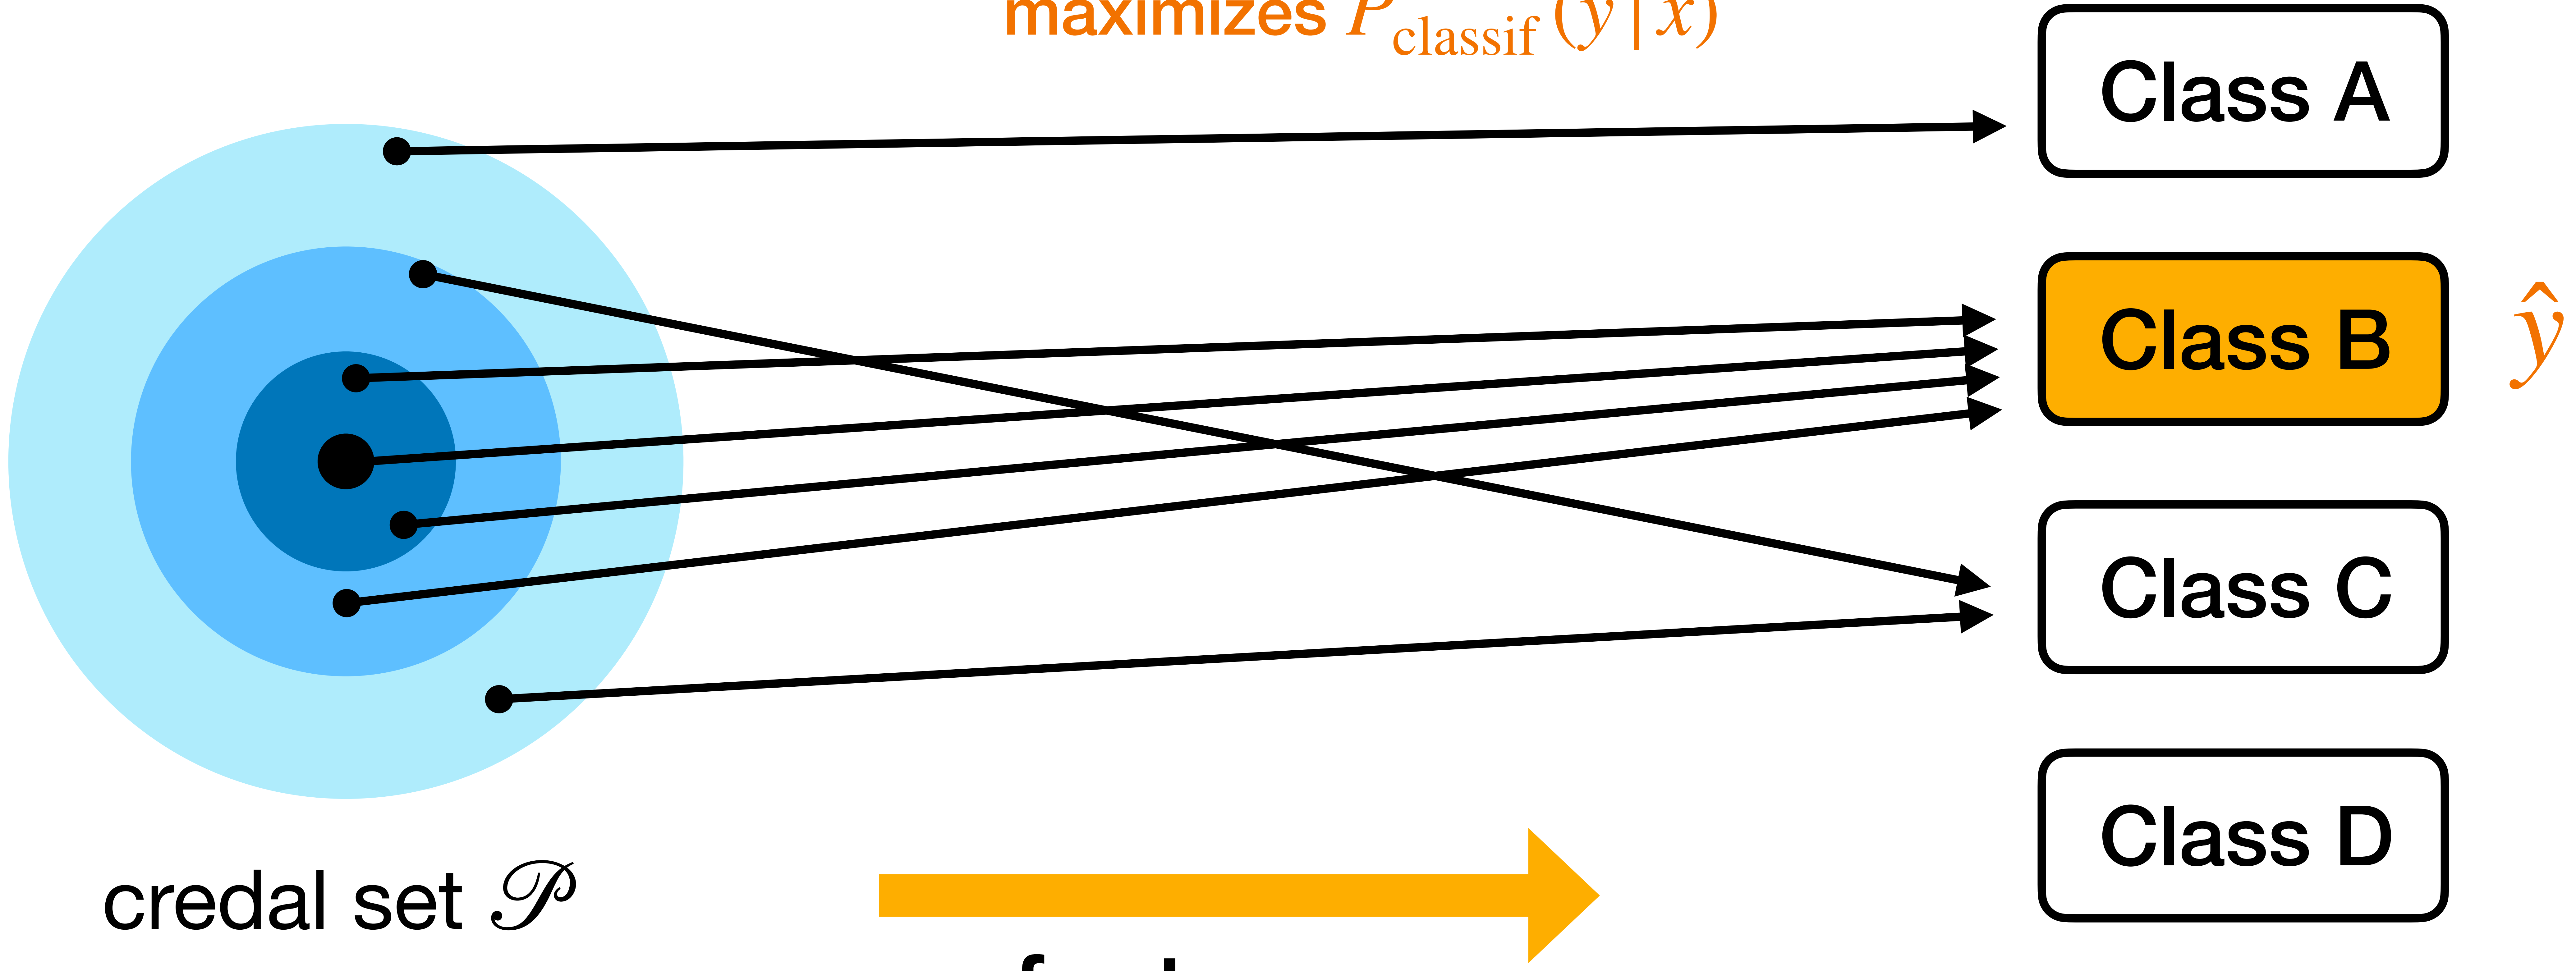

credal set $\mathscr{P}$ that contains $P_{\text{classif}}(X, Y)$

features x

ROBUSTNESS QUANTIFICATION

PROBABILISTIC CLASSIFIER

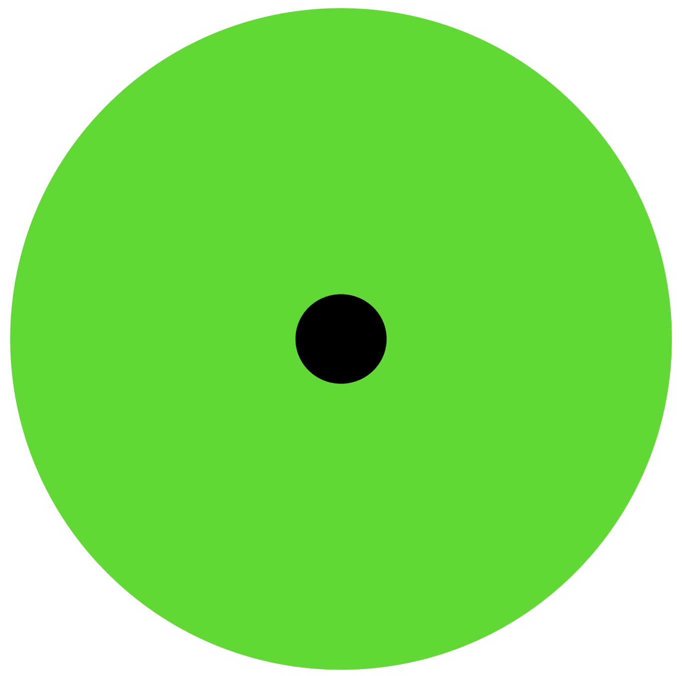

credal set $\mathscr{P}$
that contains
$P_{\text{classif}}(X, Y)$

ROBUSTNESS:
“size” of largest $\mathscr{P}$ for which the precise prediction is robust

Adrián Detavernier

Rodrigo Lassance

MACHINE LEARNING

IMPRECISE PROBABILITIES

ROBUSTNESS QUANTIFICATION

PROBLEM 2:

IMPRECISE PROBABILITIES

… are ~~arbitrary~~

PROBLEM 1:

MACHINE LEARNING

?

... is unreliable

ROBUSTNESS QUANTIFICATION

PROBABILISTIC CLASSIFIER

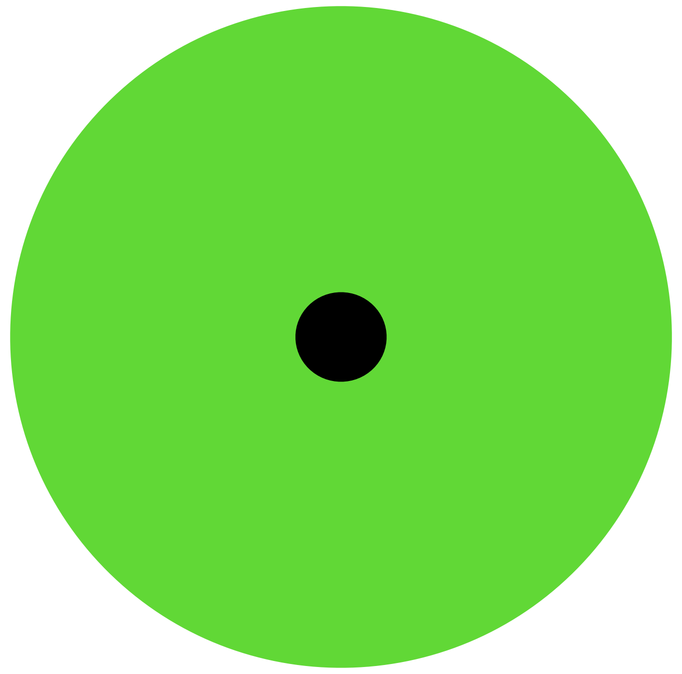

credal set $\mathscr{P}$
that contains
$P_{\text{classif}}(X, Y)$

ROBUSTNESS:
"size" of largest $\mathscr{P}$ for which the precise prediction is robust

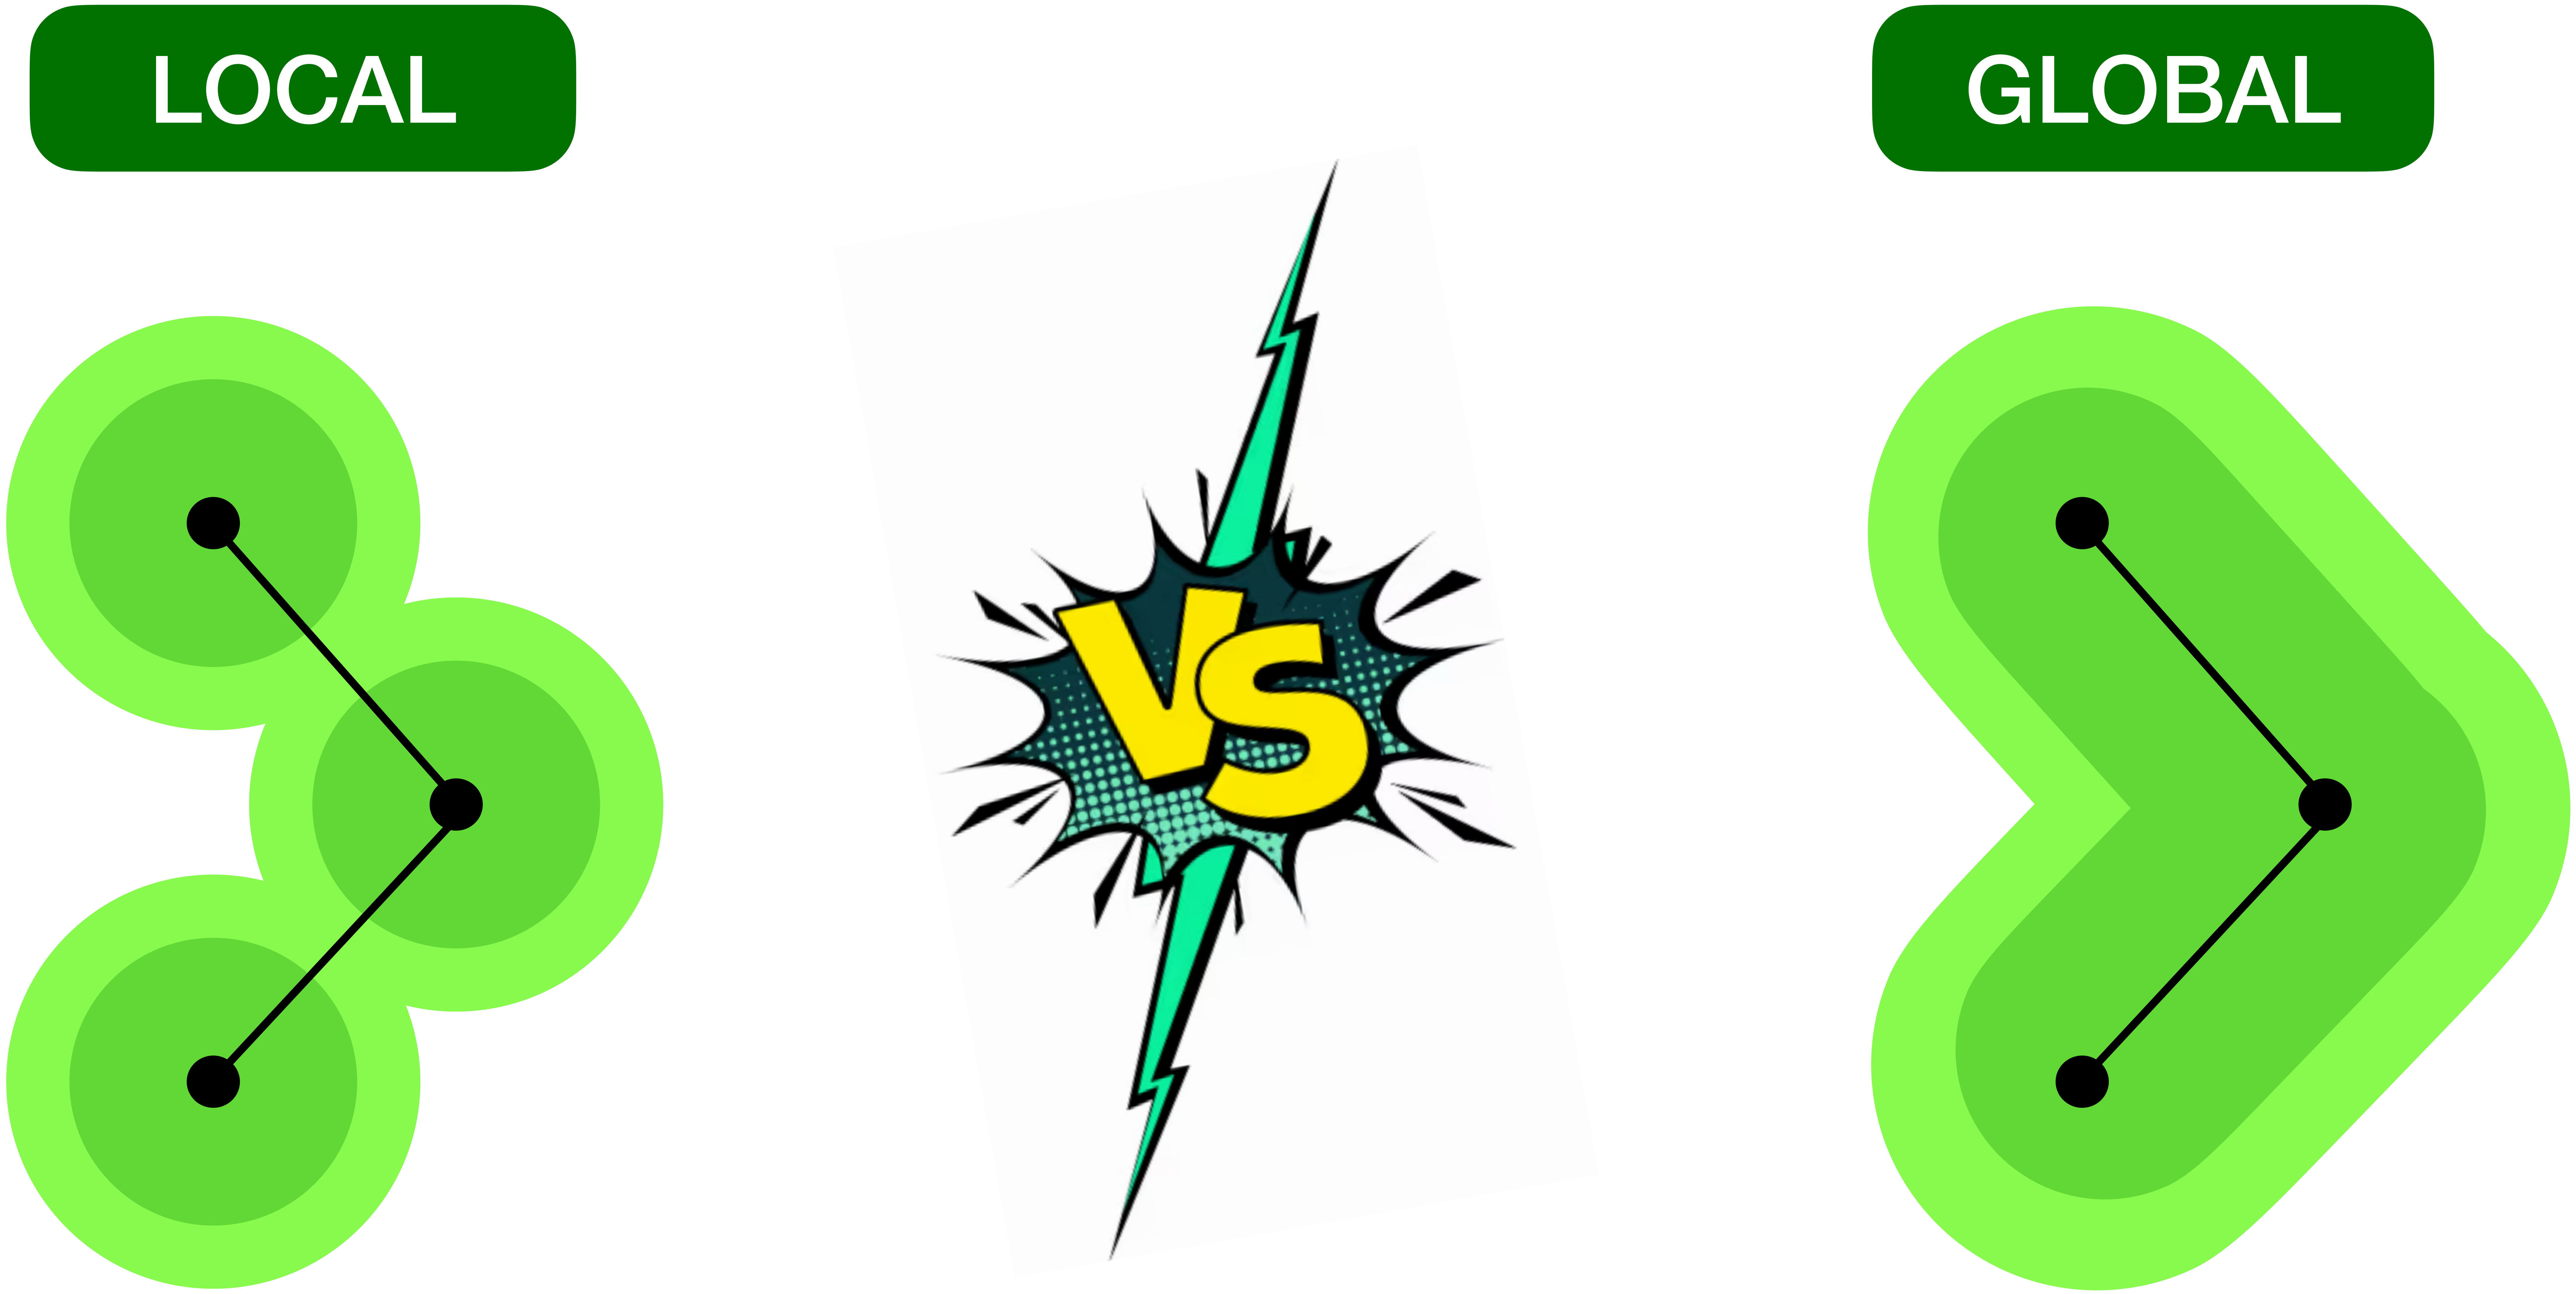
LOCAL
VS
GLOBAL

ϵ-CONTAMINATION

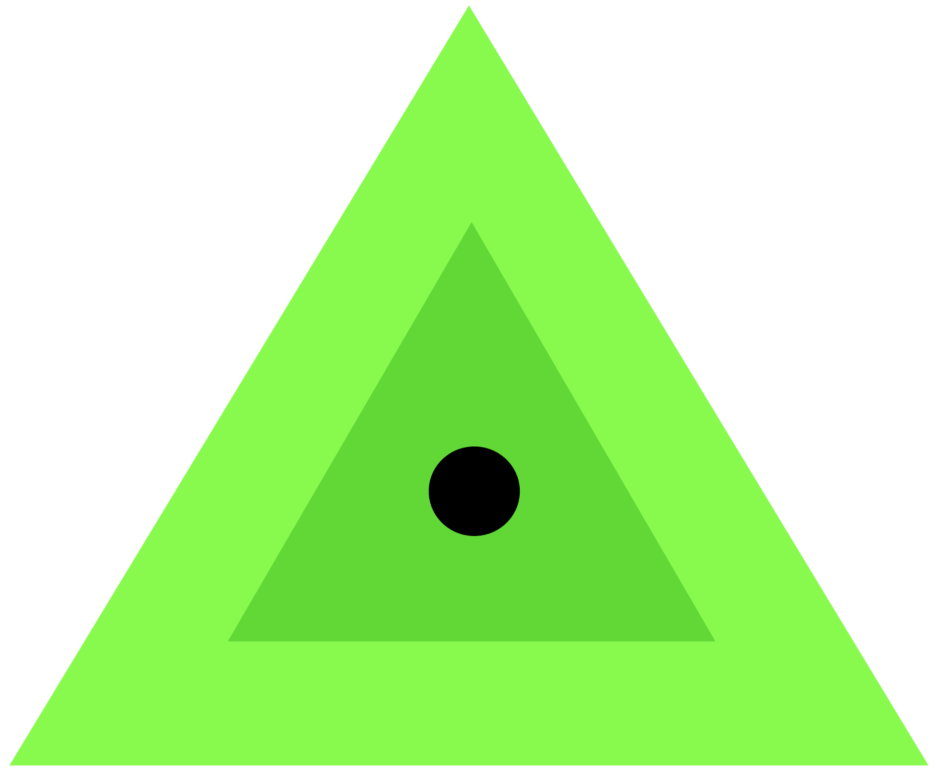

$$\mathscr{P}_\epsilon$$
$$\parallel$$
$$\left\{(1-\epsilon)P_{\text{classif}} + \epsilon P \colon P \in \Delta\right\}$$

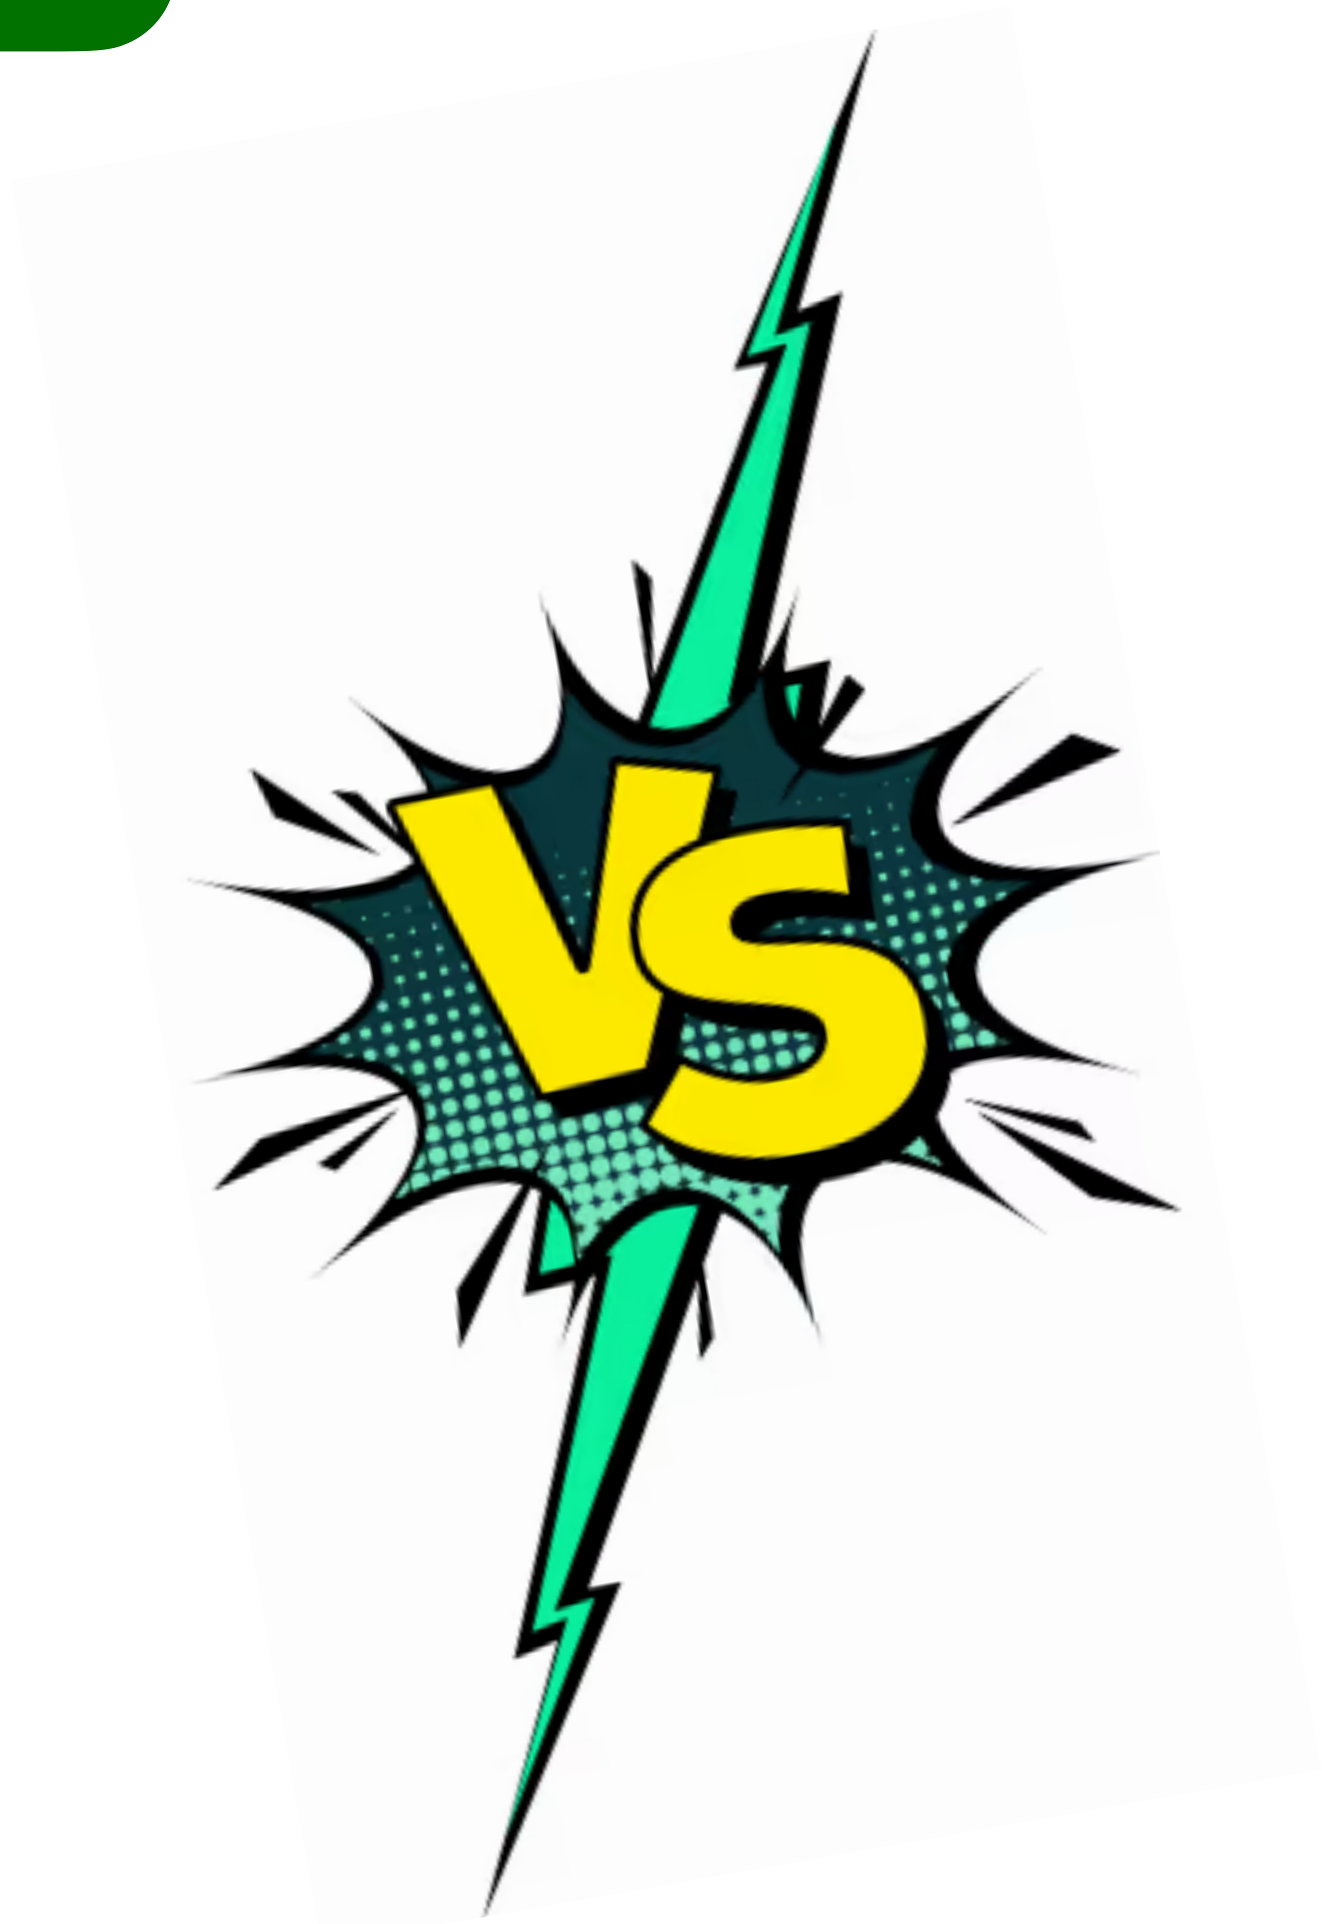

OTHER STUFF

distance-based, ...

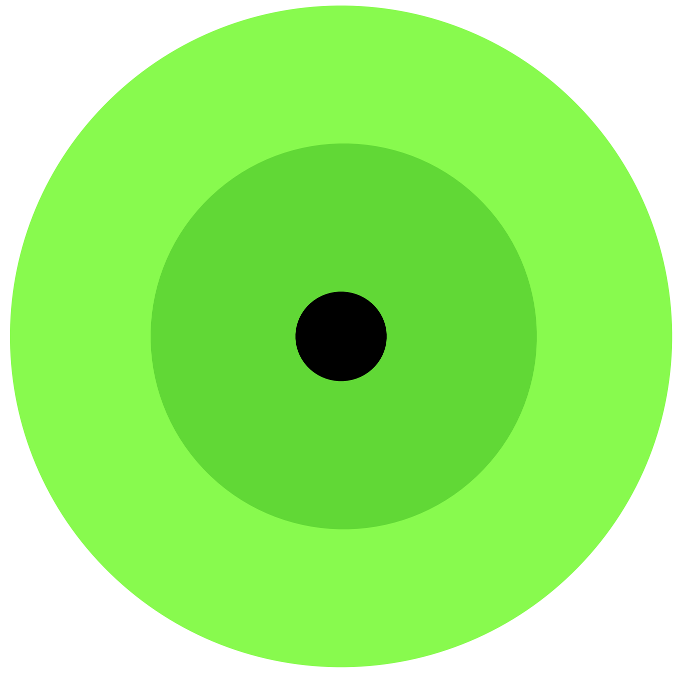

$$\mathscr{P}_\epsilon$$
$$\parallel$$
$$\left\{P \in \Delta \colon \mathrm{d}(P_{\text{classif}}, P) < \epsilon\right\}$$

ROBUSTNESS QUANTIFICATION

PROBABILISTIC CLASSIFIER

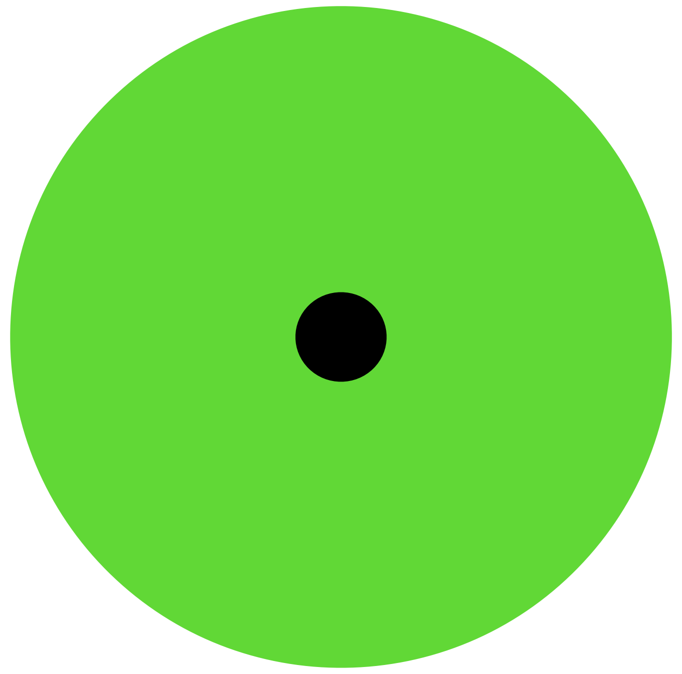

credal set $\mathscr{P}$
that contains
$P_{\text{classif}}(X, Y)$

ROBUSTNESS:
"size" of largest $\mathscr{P}$ for which the precise prediction is robust

2014

MRF

BN

Cassio de Campos

Alessandro Antonucci

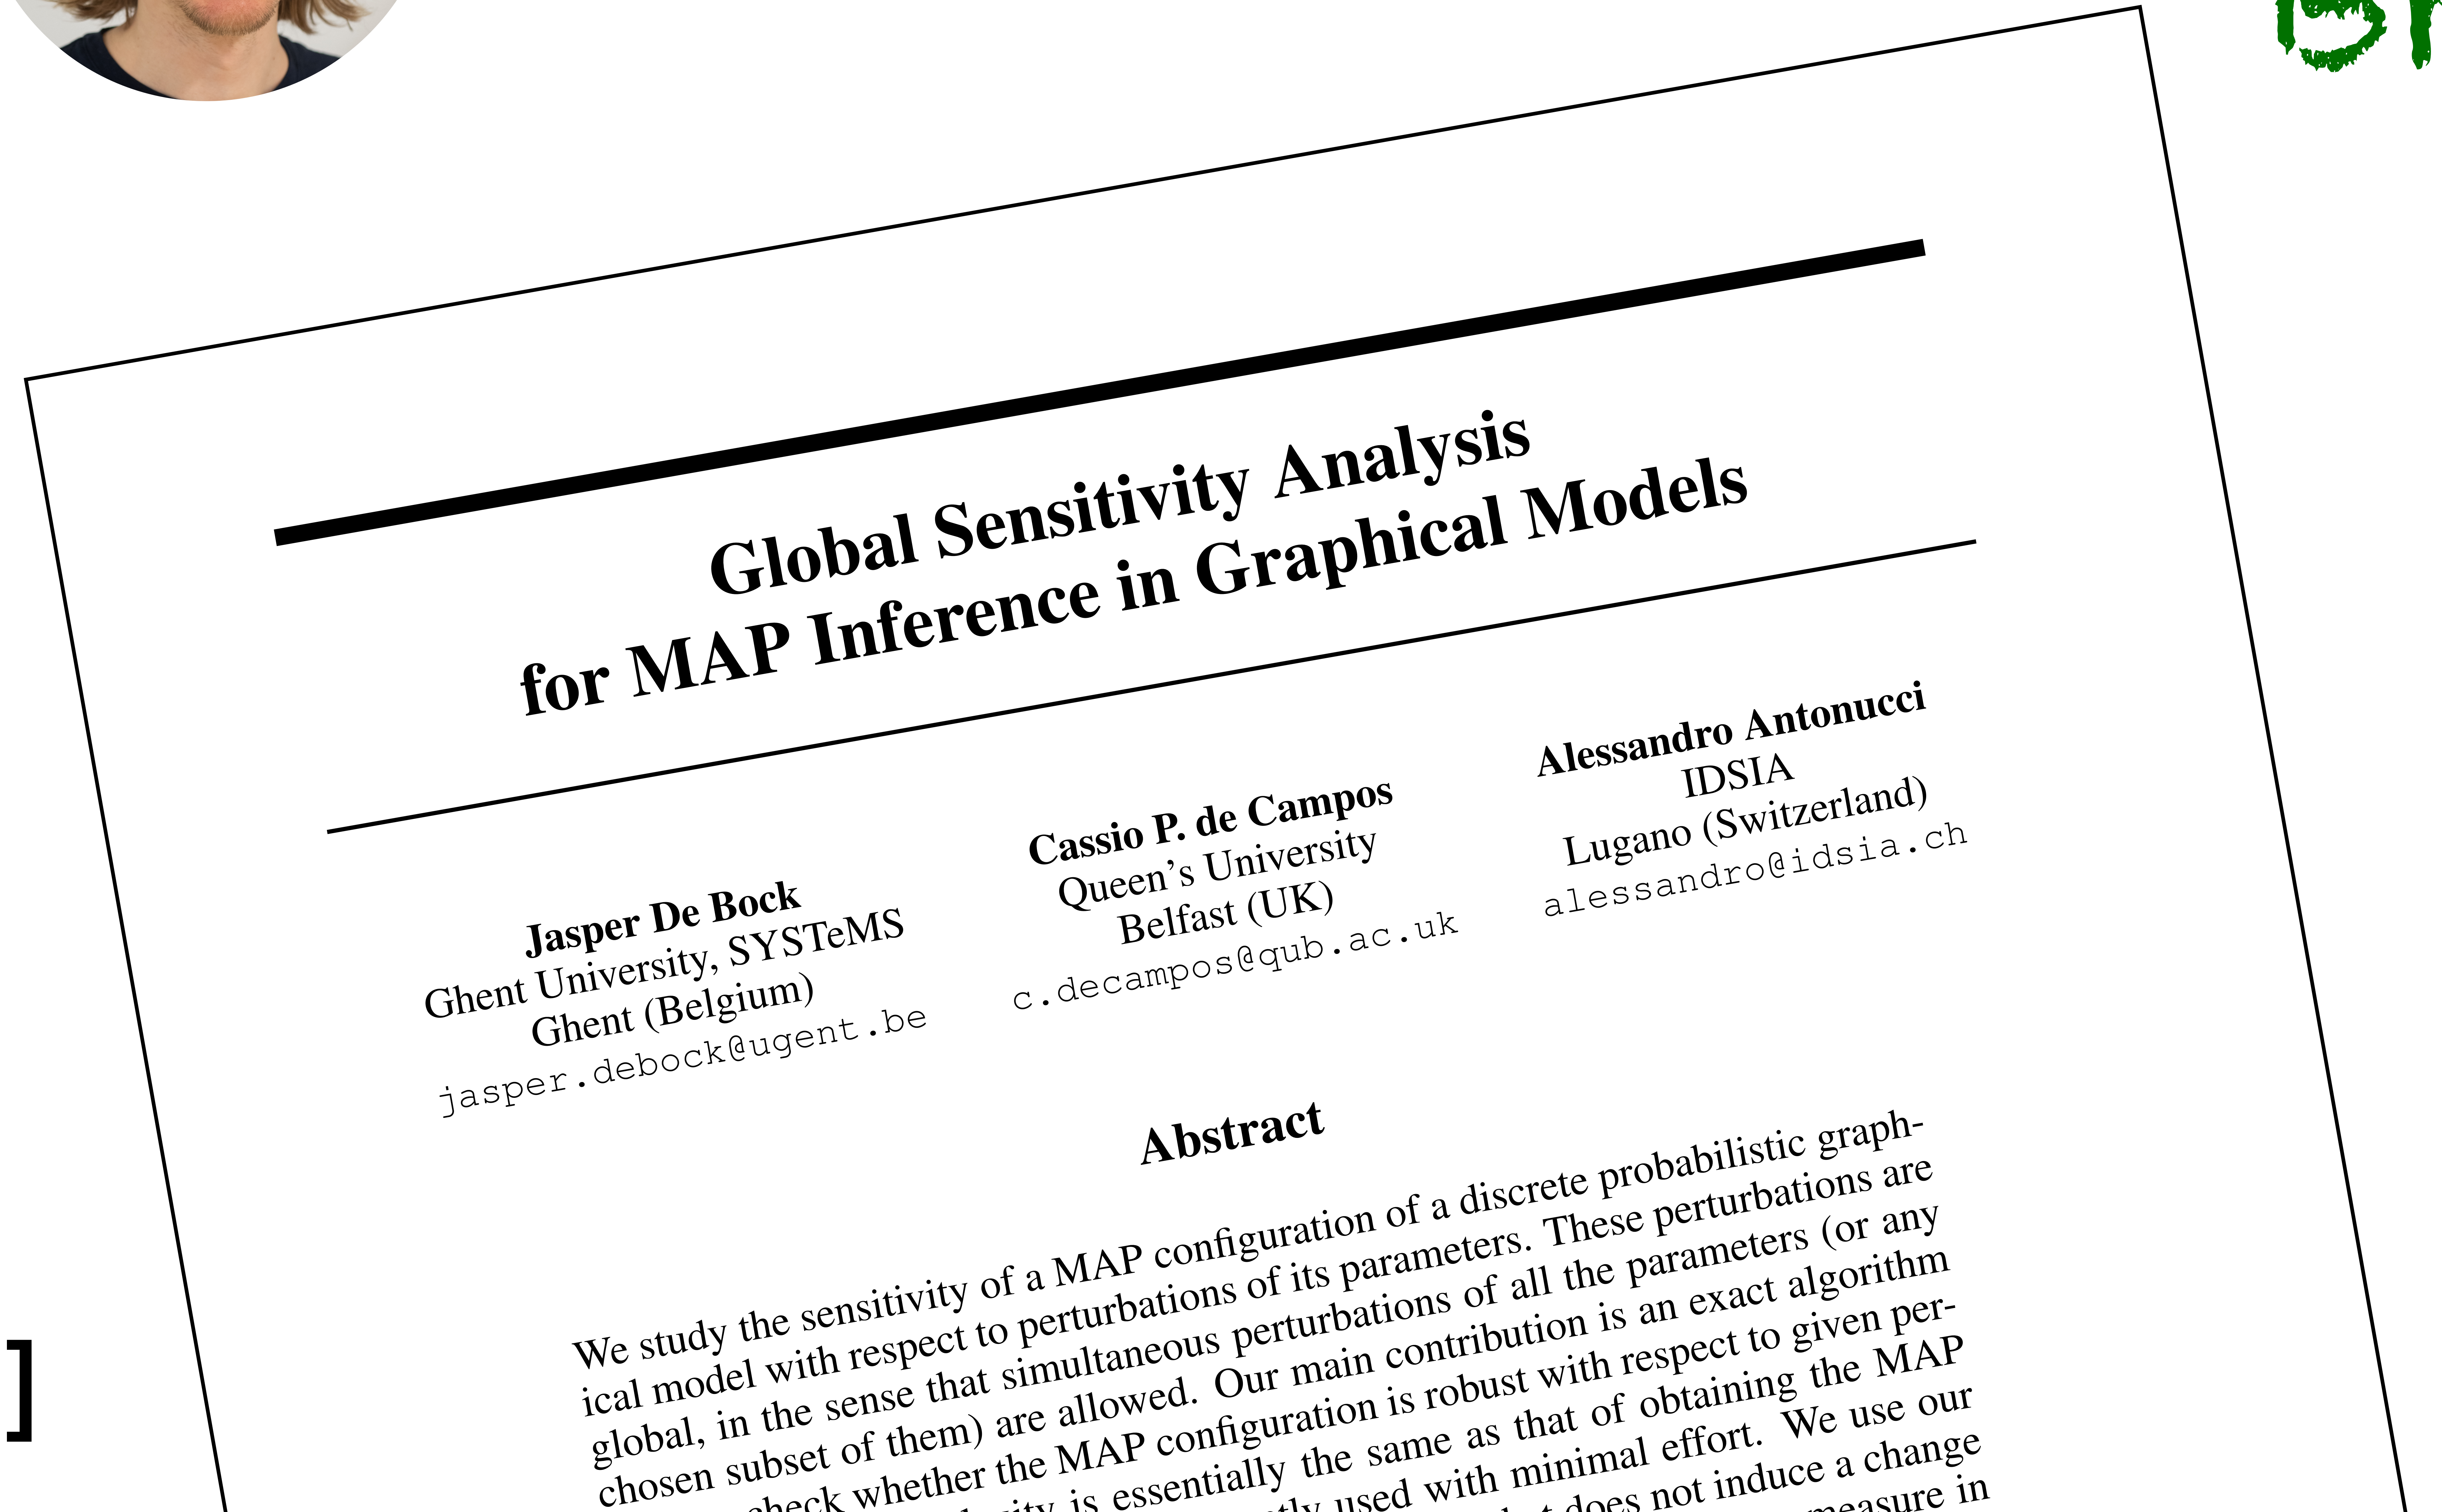

Global Sensitivity Analysis for MAP Inference in Graphical Models

Jasper De Bock
Ghent University, SYSTeMS
Ghent (Belgium)
jasper.debock@ugent.be

Cassio P. de Campos
Queen's University
Belfast (UK)
c.decampos@qub.ac.uk

Alessandro Antonucci
IDSIA
Lugano (Switzerland)
alessandro@idsia.ch

Abstract

We study the sensitivity of a MAP configuration of a discrete probabilistic graphical model with respect to perturbations of its parameters. These perturbations are global, in the sense that simultaneous perturbations of all the parameters (or any chosen subset of them) are allowed. Our main contribution is an exact algorithm ... check whether the MAP configuration is robust with respect to given per- ... ity is essentially the same as that of obtaining the MAP ... tly used with minimal effort. We use our ... t does not induce a change ... measure in

[1]

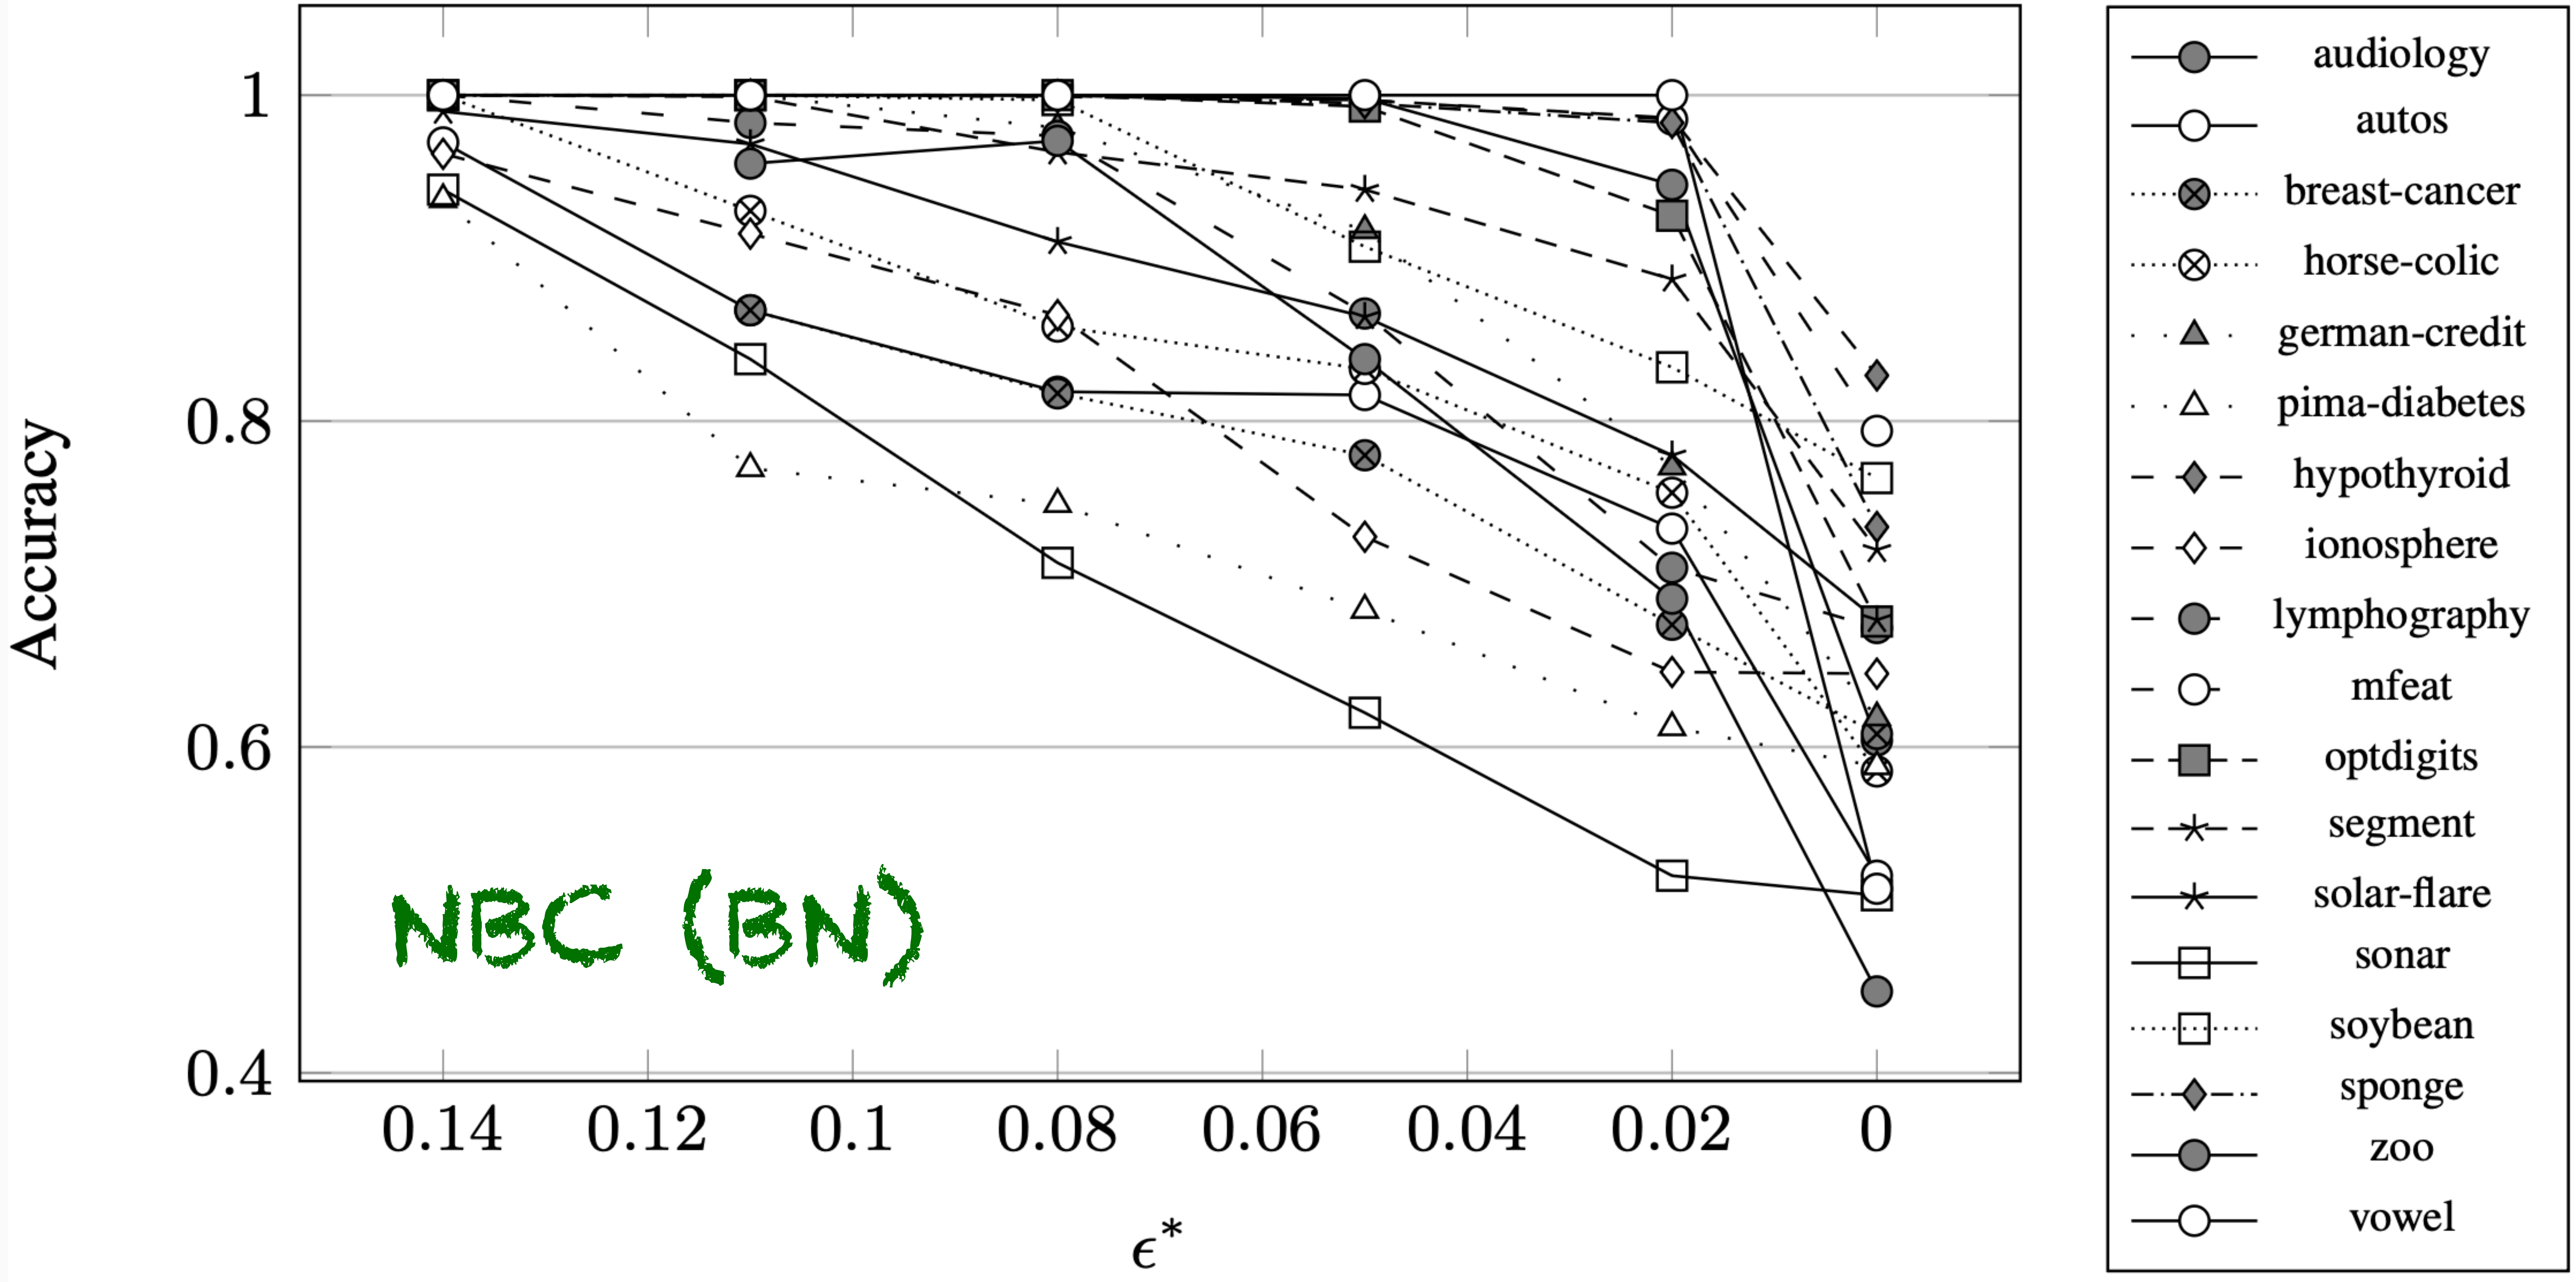
Accuracy
1
0.8
0.6
0.4
0.14
0.12
0.1
0.08
0.06
0.04
0.02
0
ϵ^*
NBC (BN)
audiology
autos
breast-cancer
horse-colic
german-credit
pima-diabetes
hypothyroid
ionosphere
lymphography
mfeat
optdigits
segment
solar-flare
sonar
soybean
sponge
zoo
vowel

[1]

MRF

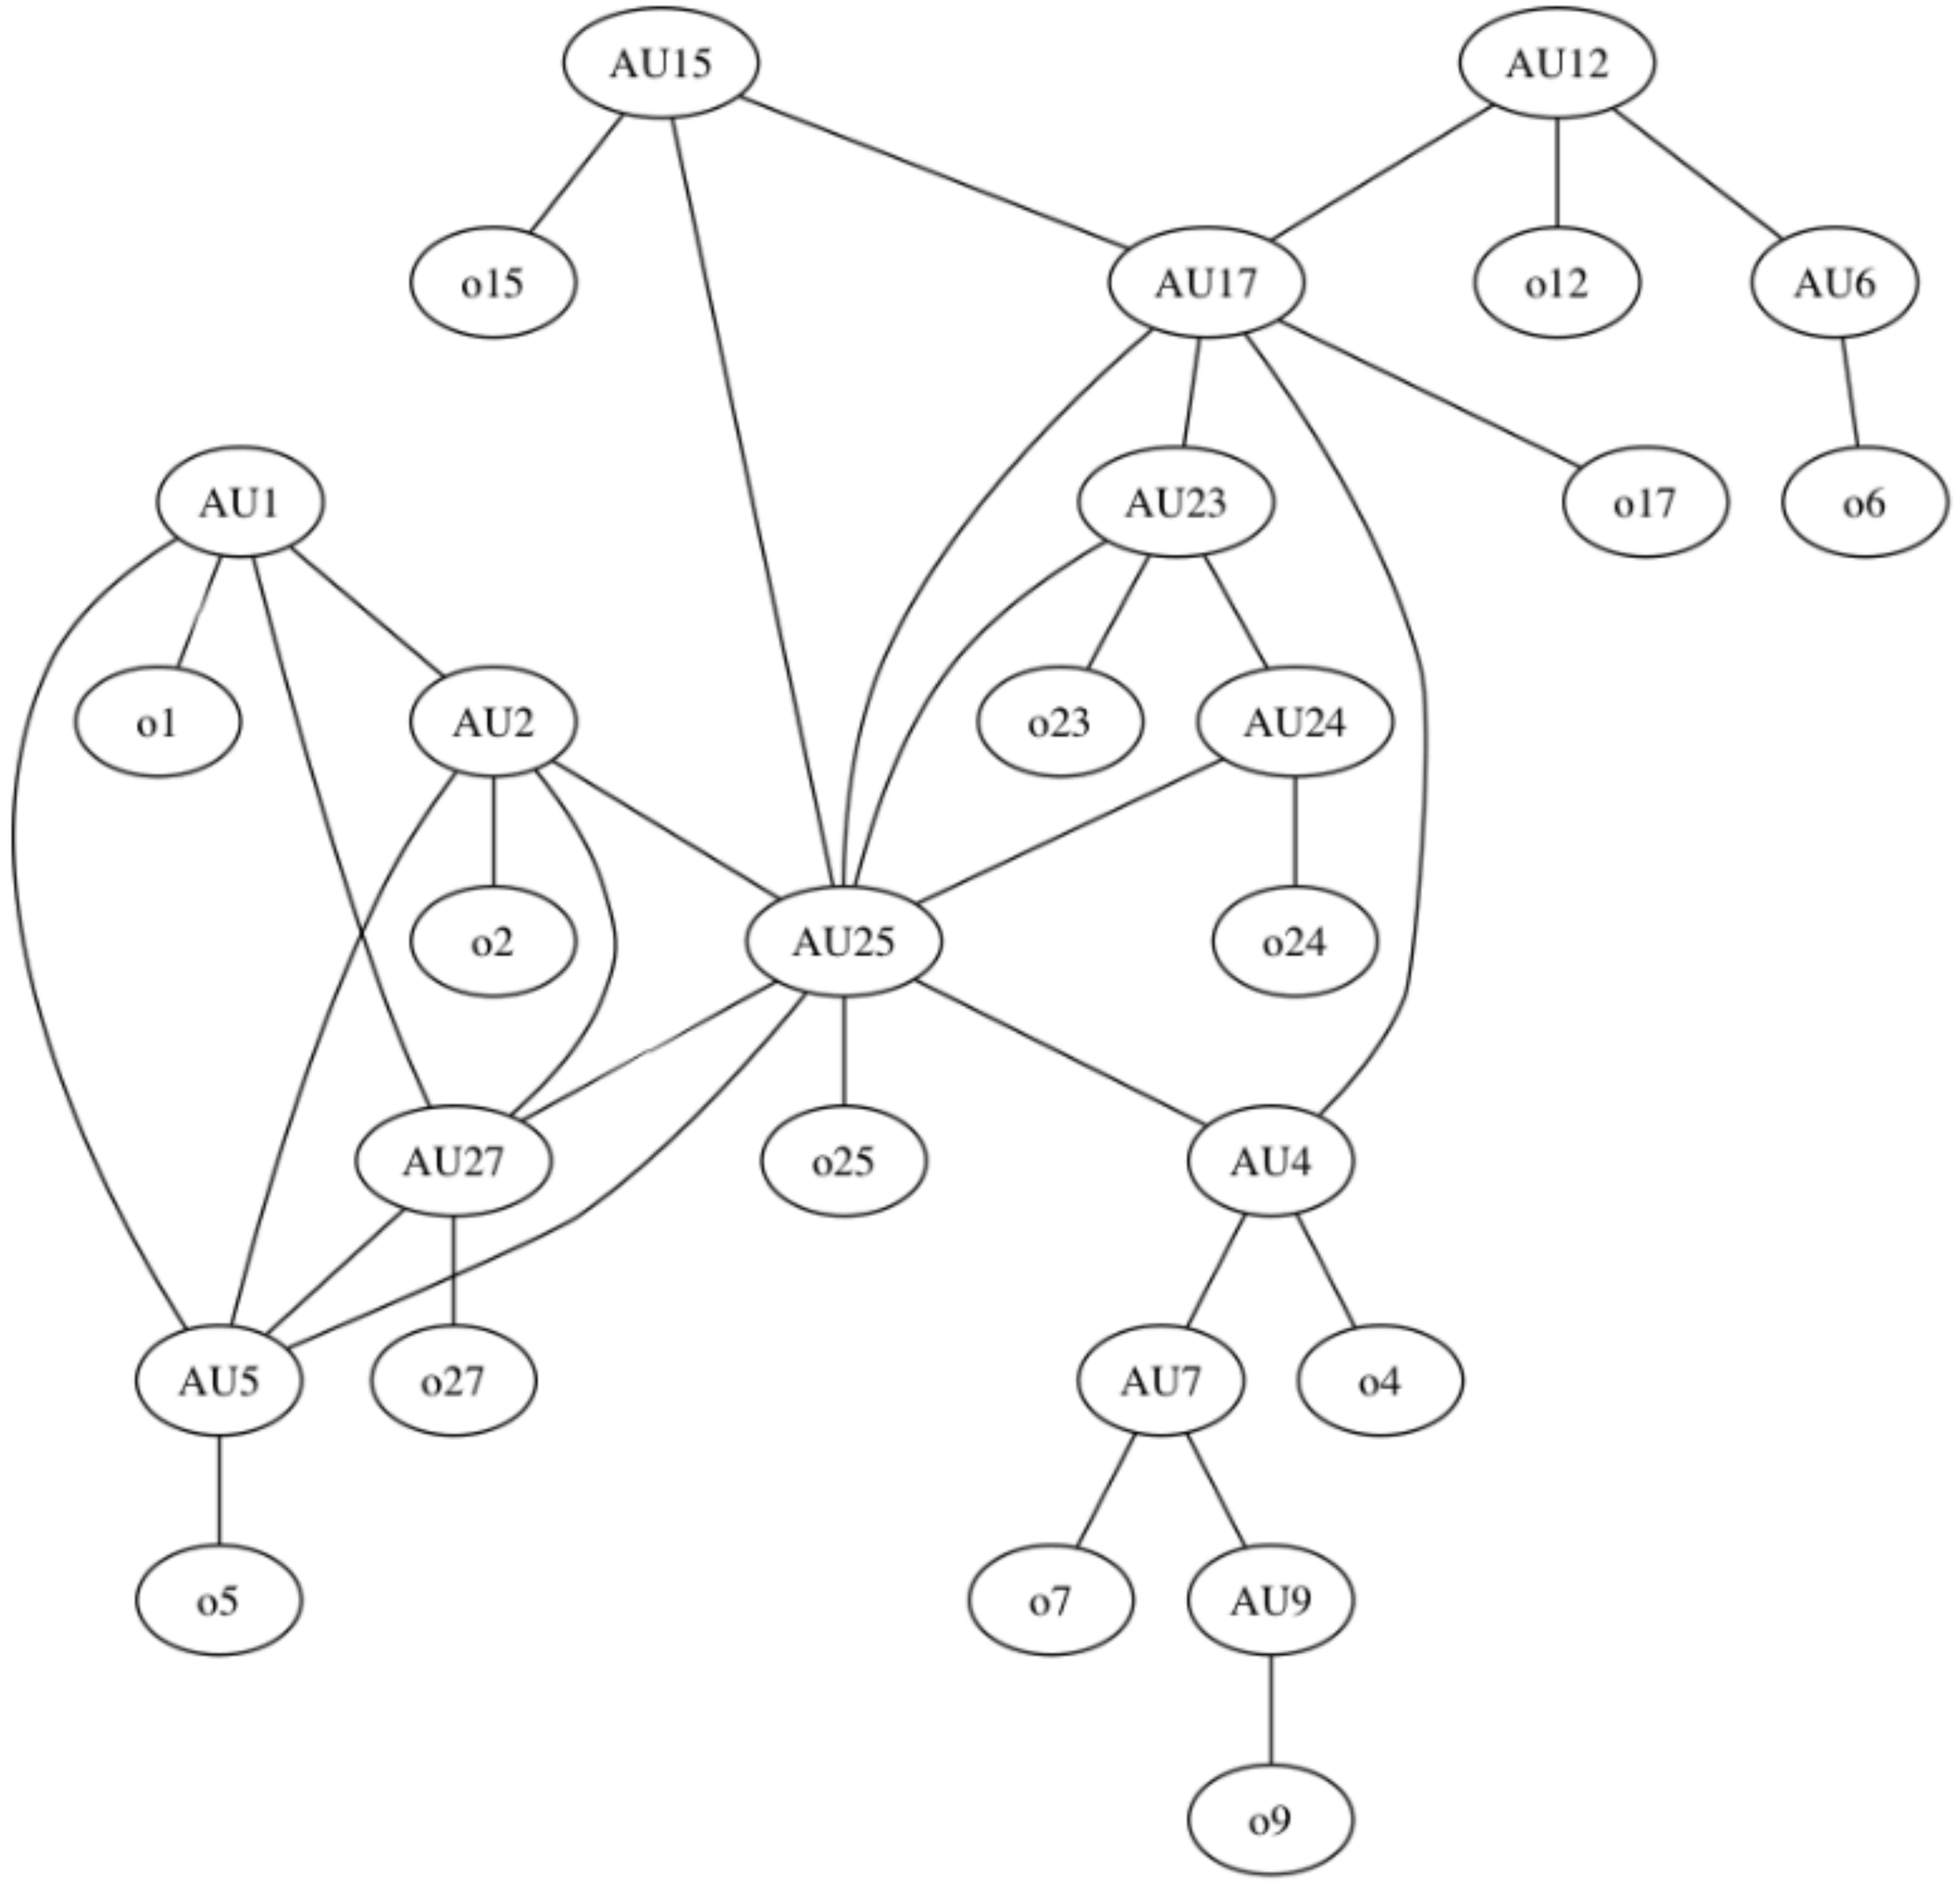

(a) MRF used in the computations.

[1]

MRF

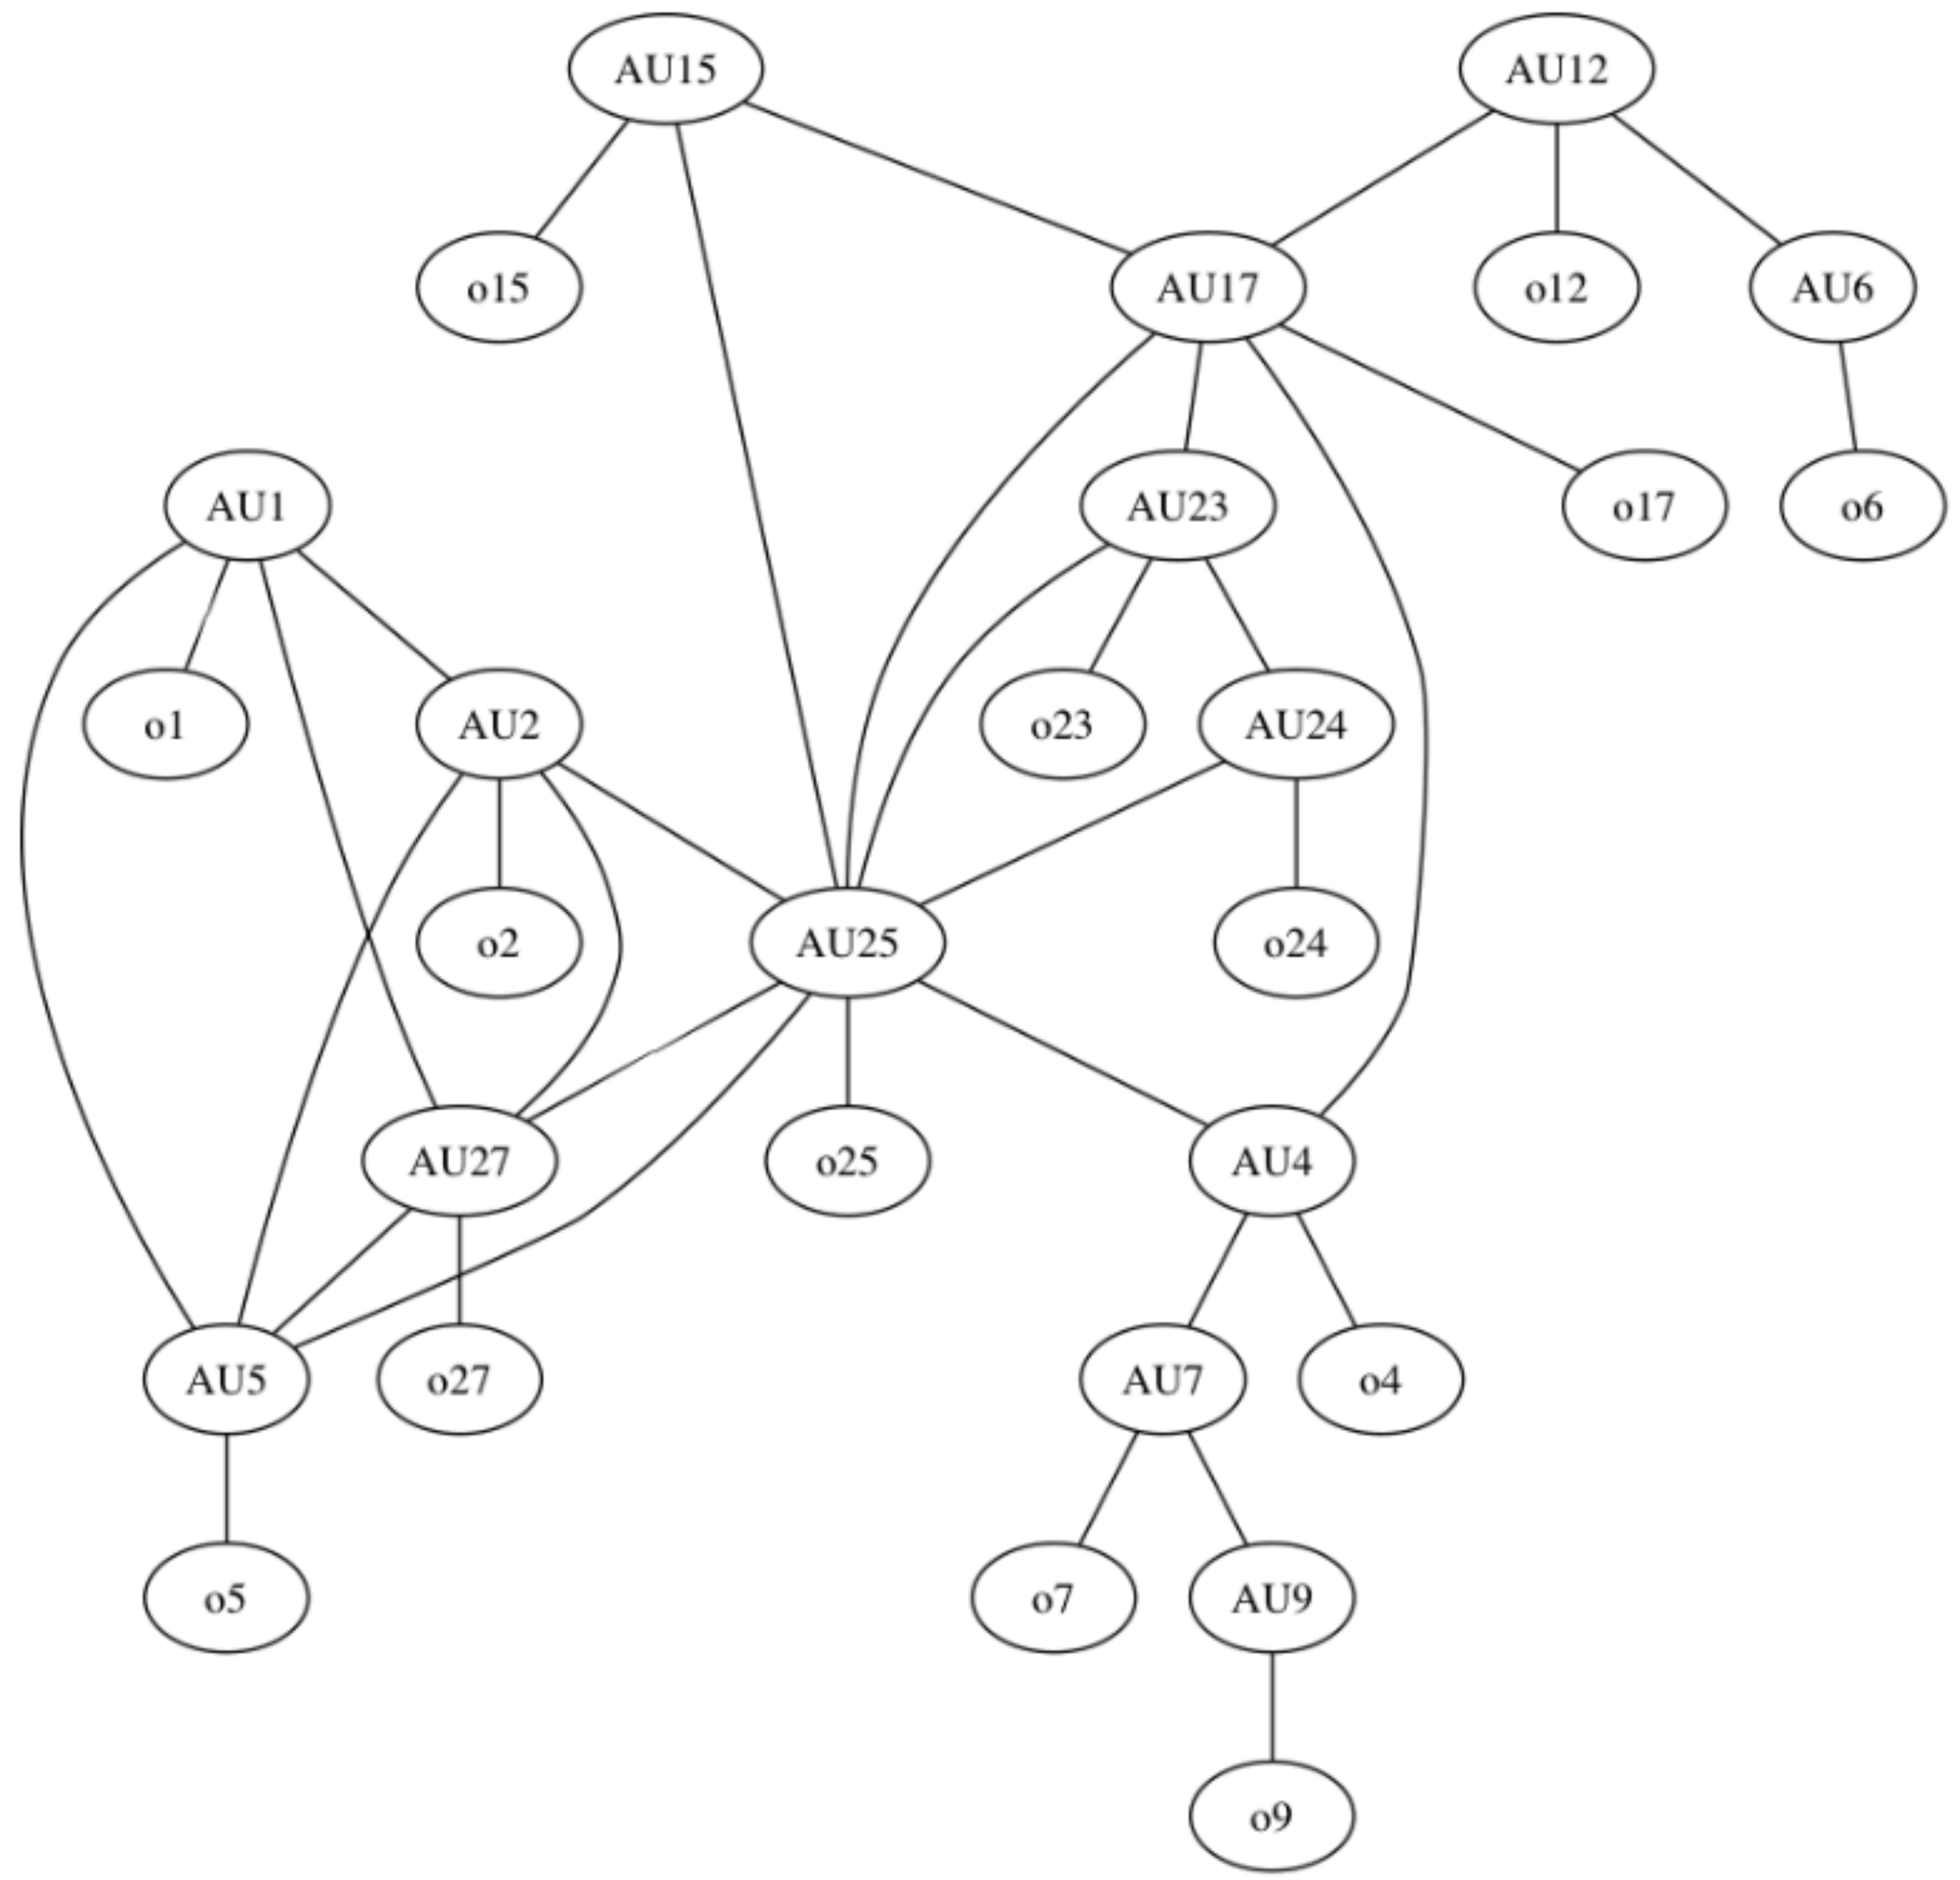

(a) MRF used in the computations.

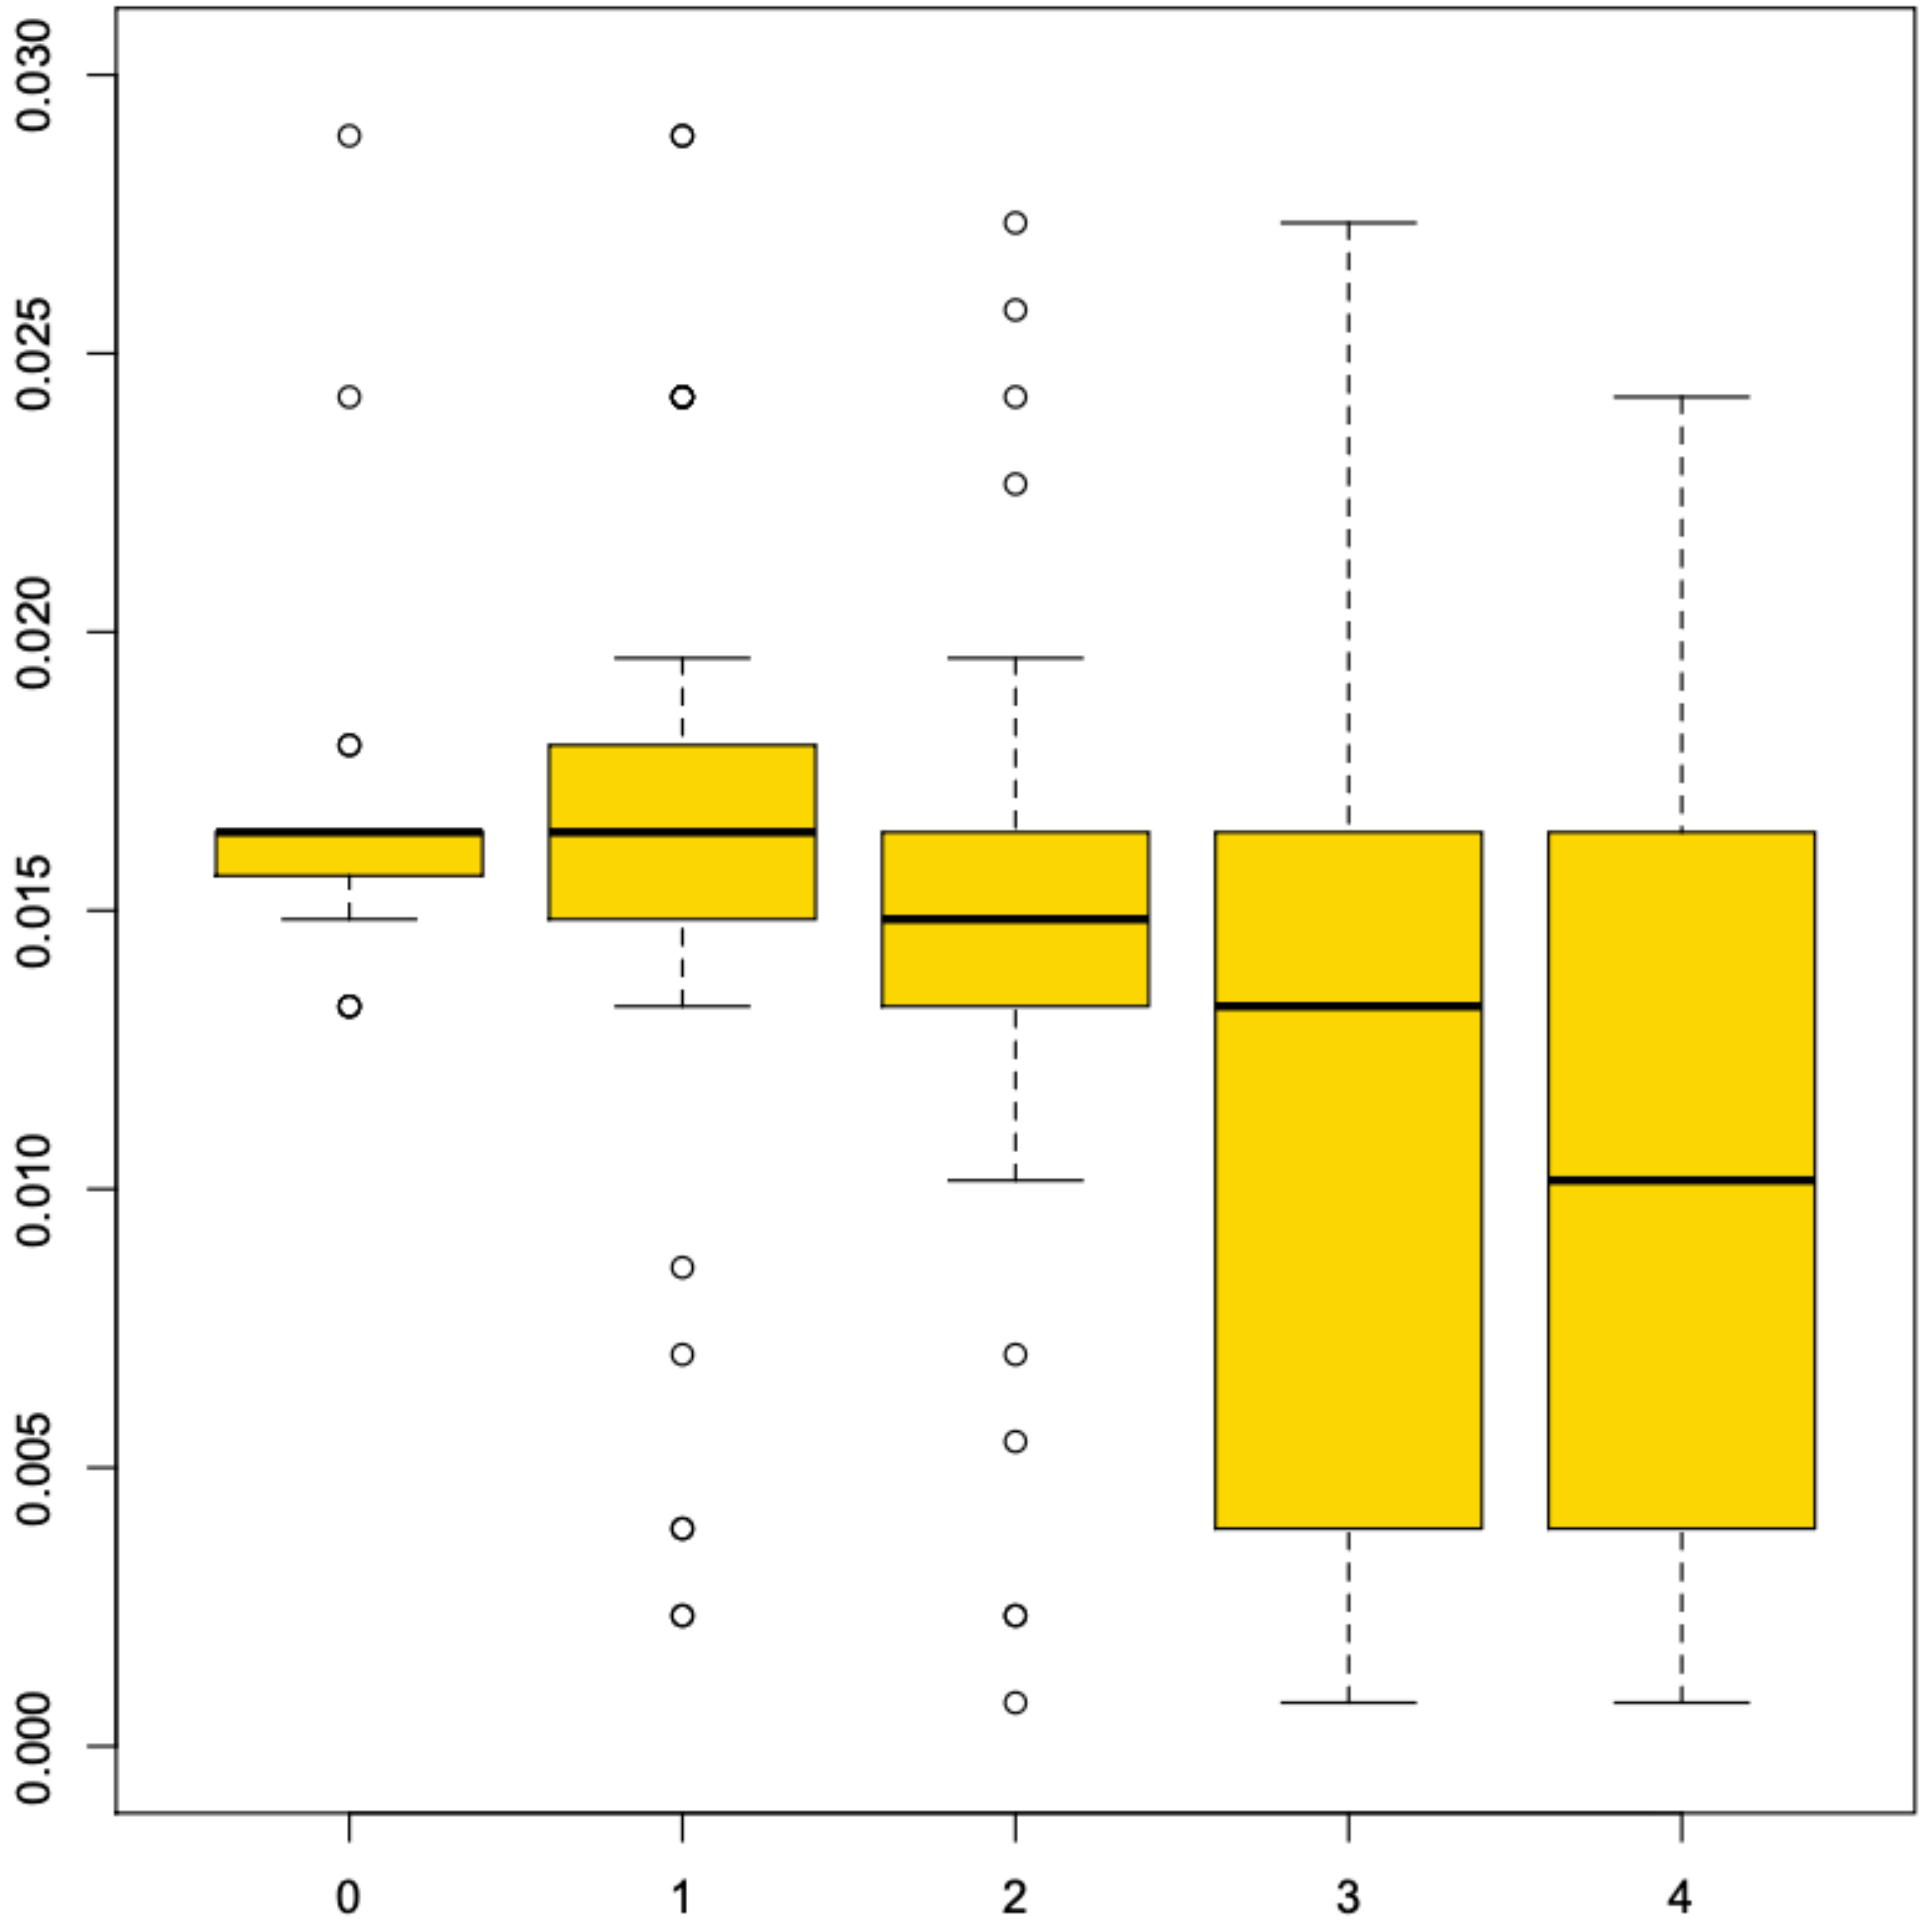

(b) Robustness split by Hamming distances.

[1]

2017 - 2020

SPN

[2-4, ...]

Cassio
de Campos
& various
co-authors

[5]

GeF+

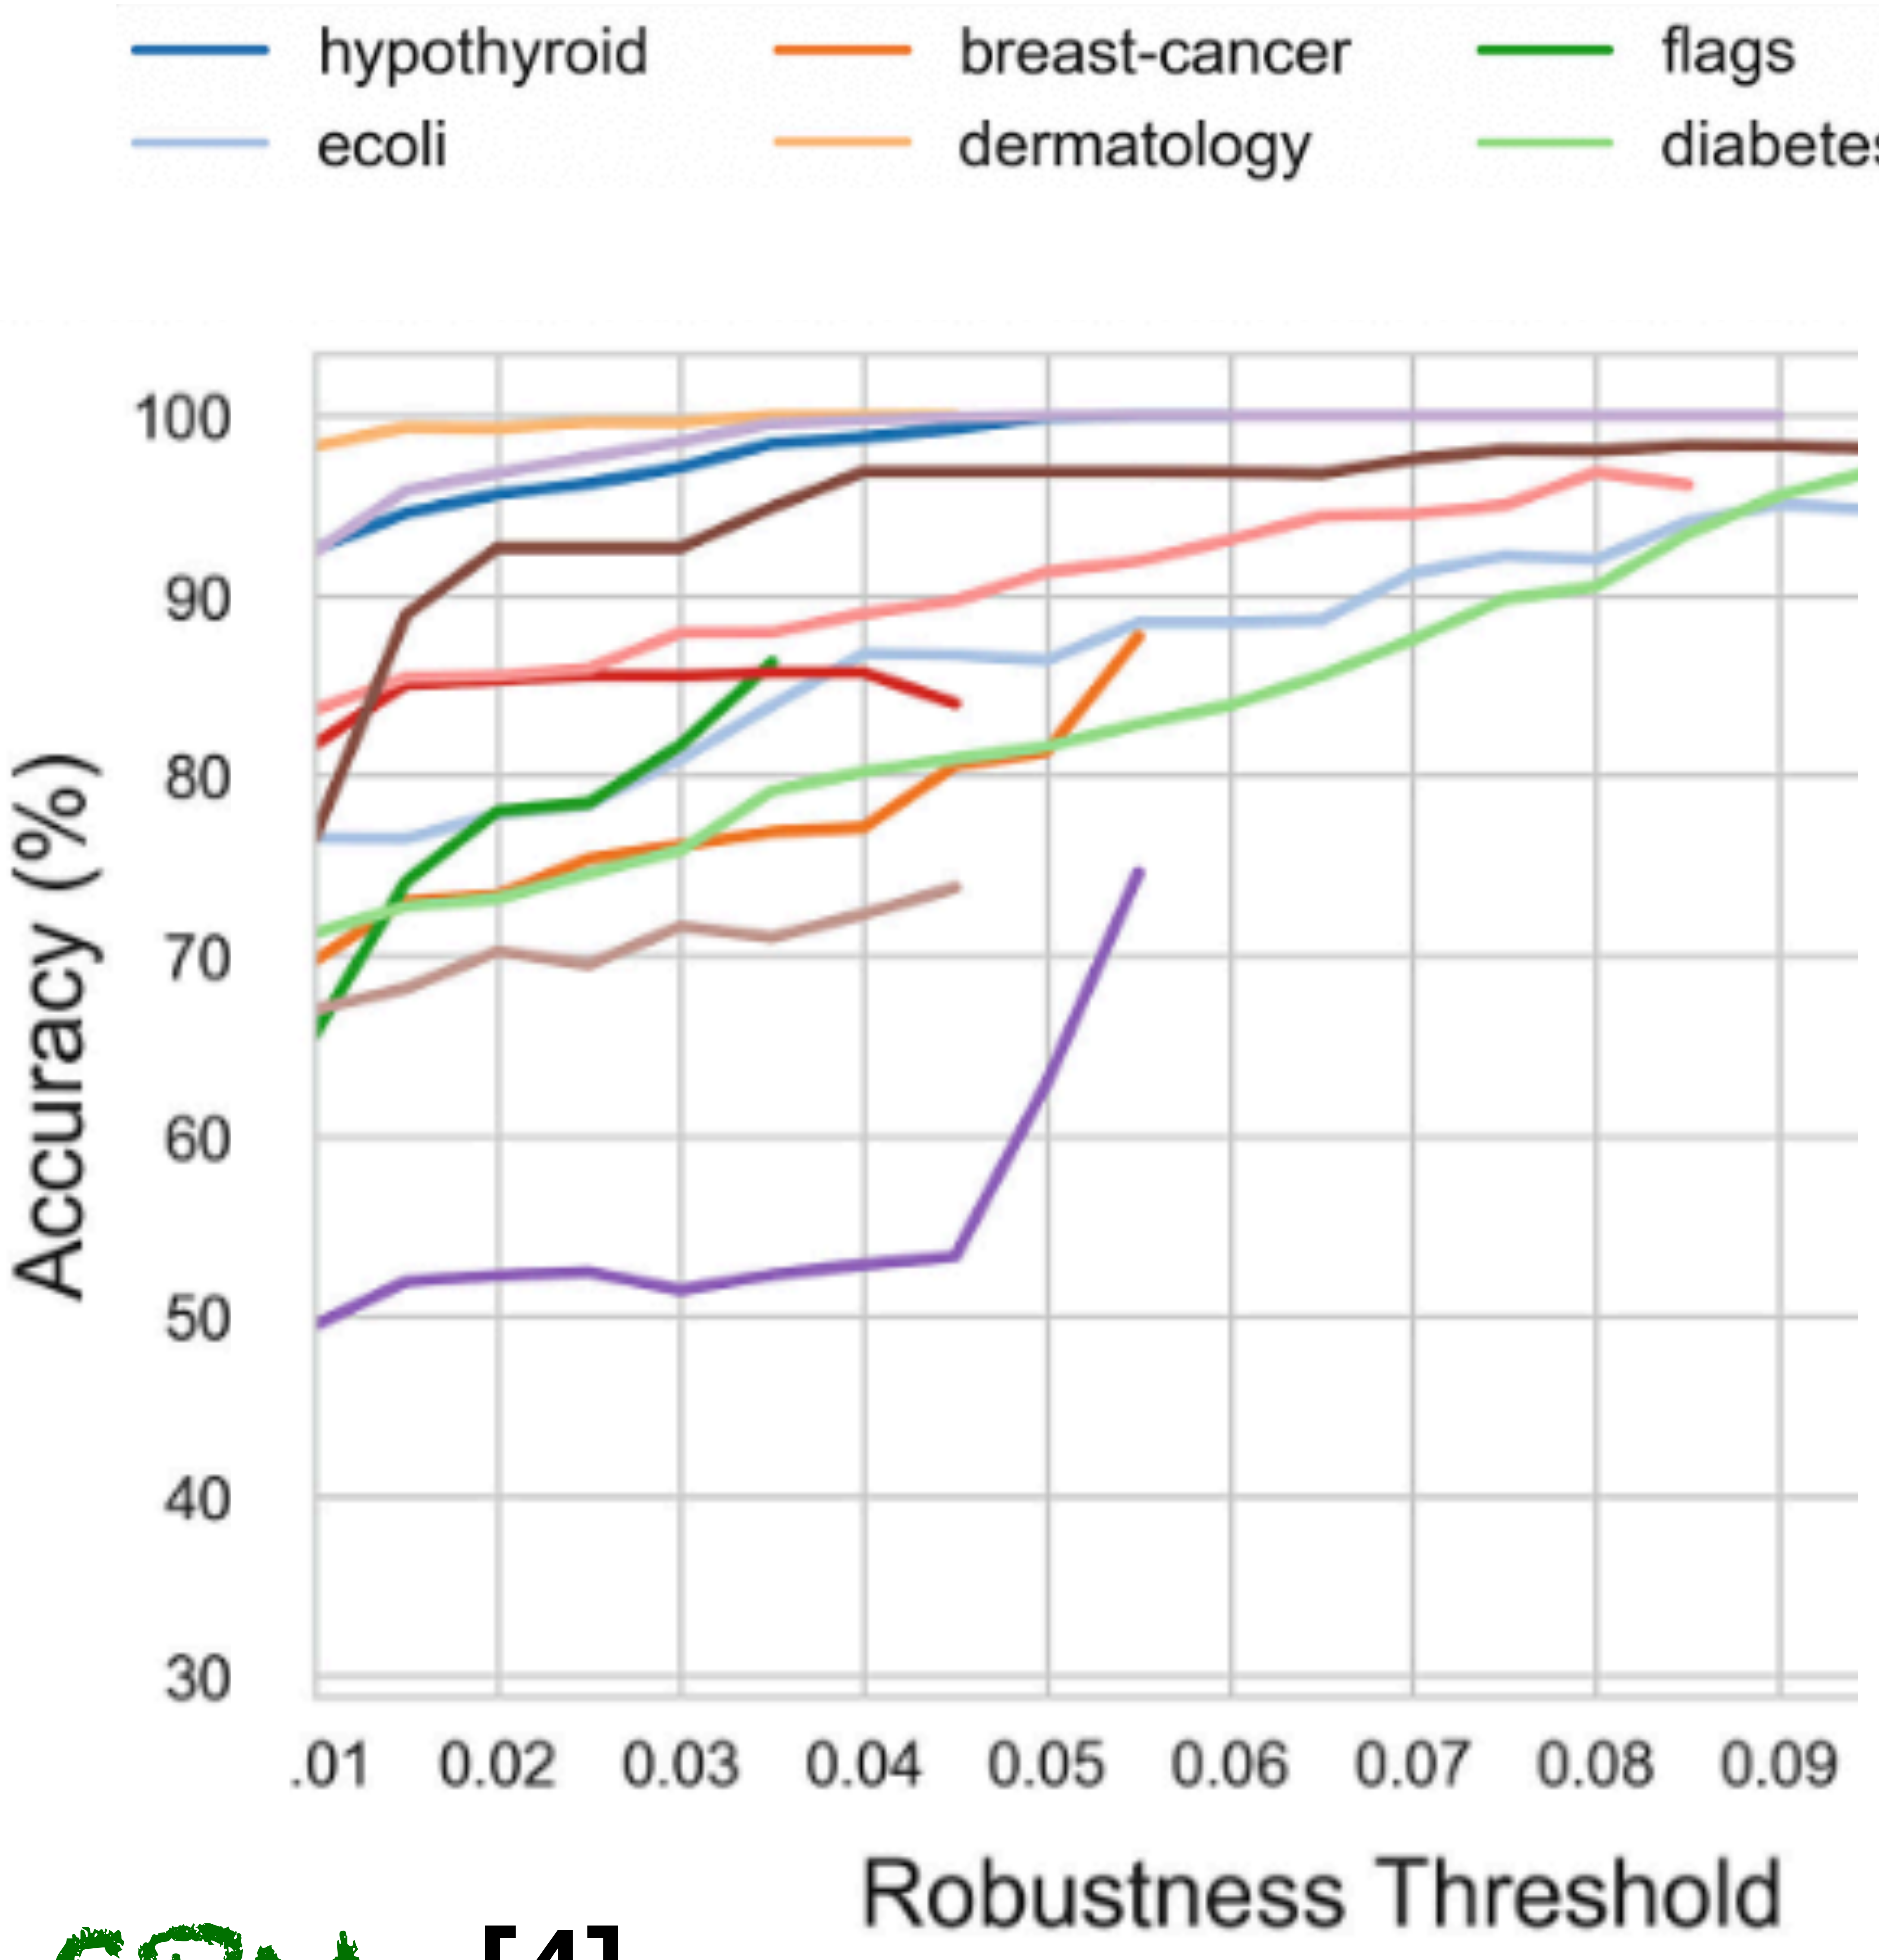
hypothyroid
breast-cancer
flags
colic
cmc
balance-scale
ecoli
dermatology
diabetes
heart-h
car
bridges
100
90
80
70
60
50
40
30
Accuracy (%)
.01
0.02
0.03
0.04
0.05
0.06
0.07
0.08
0.09
Robustness Threshold

SPN [4]

GeF+ [5]

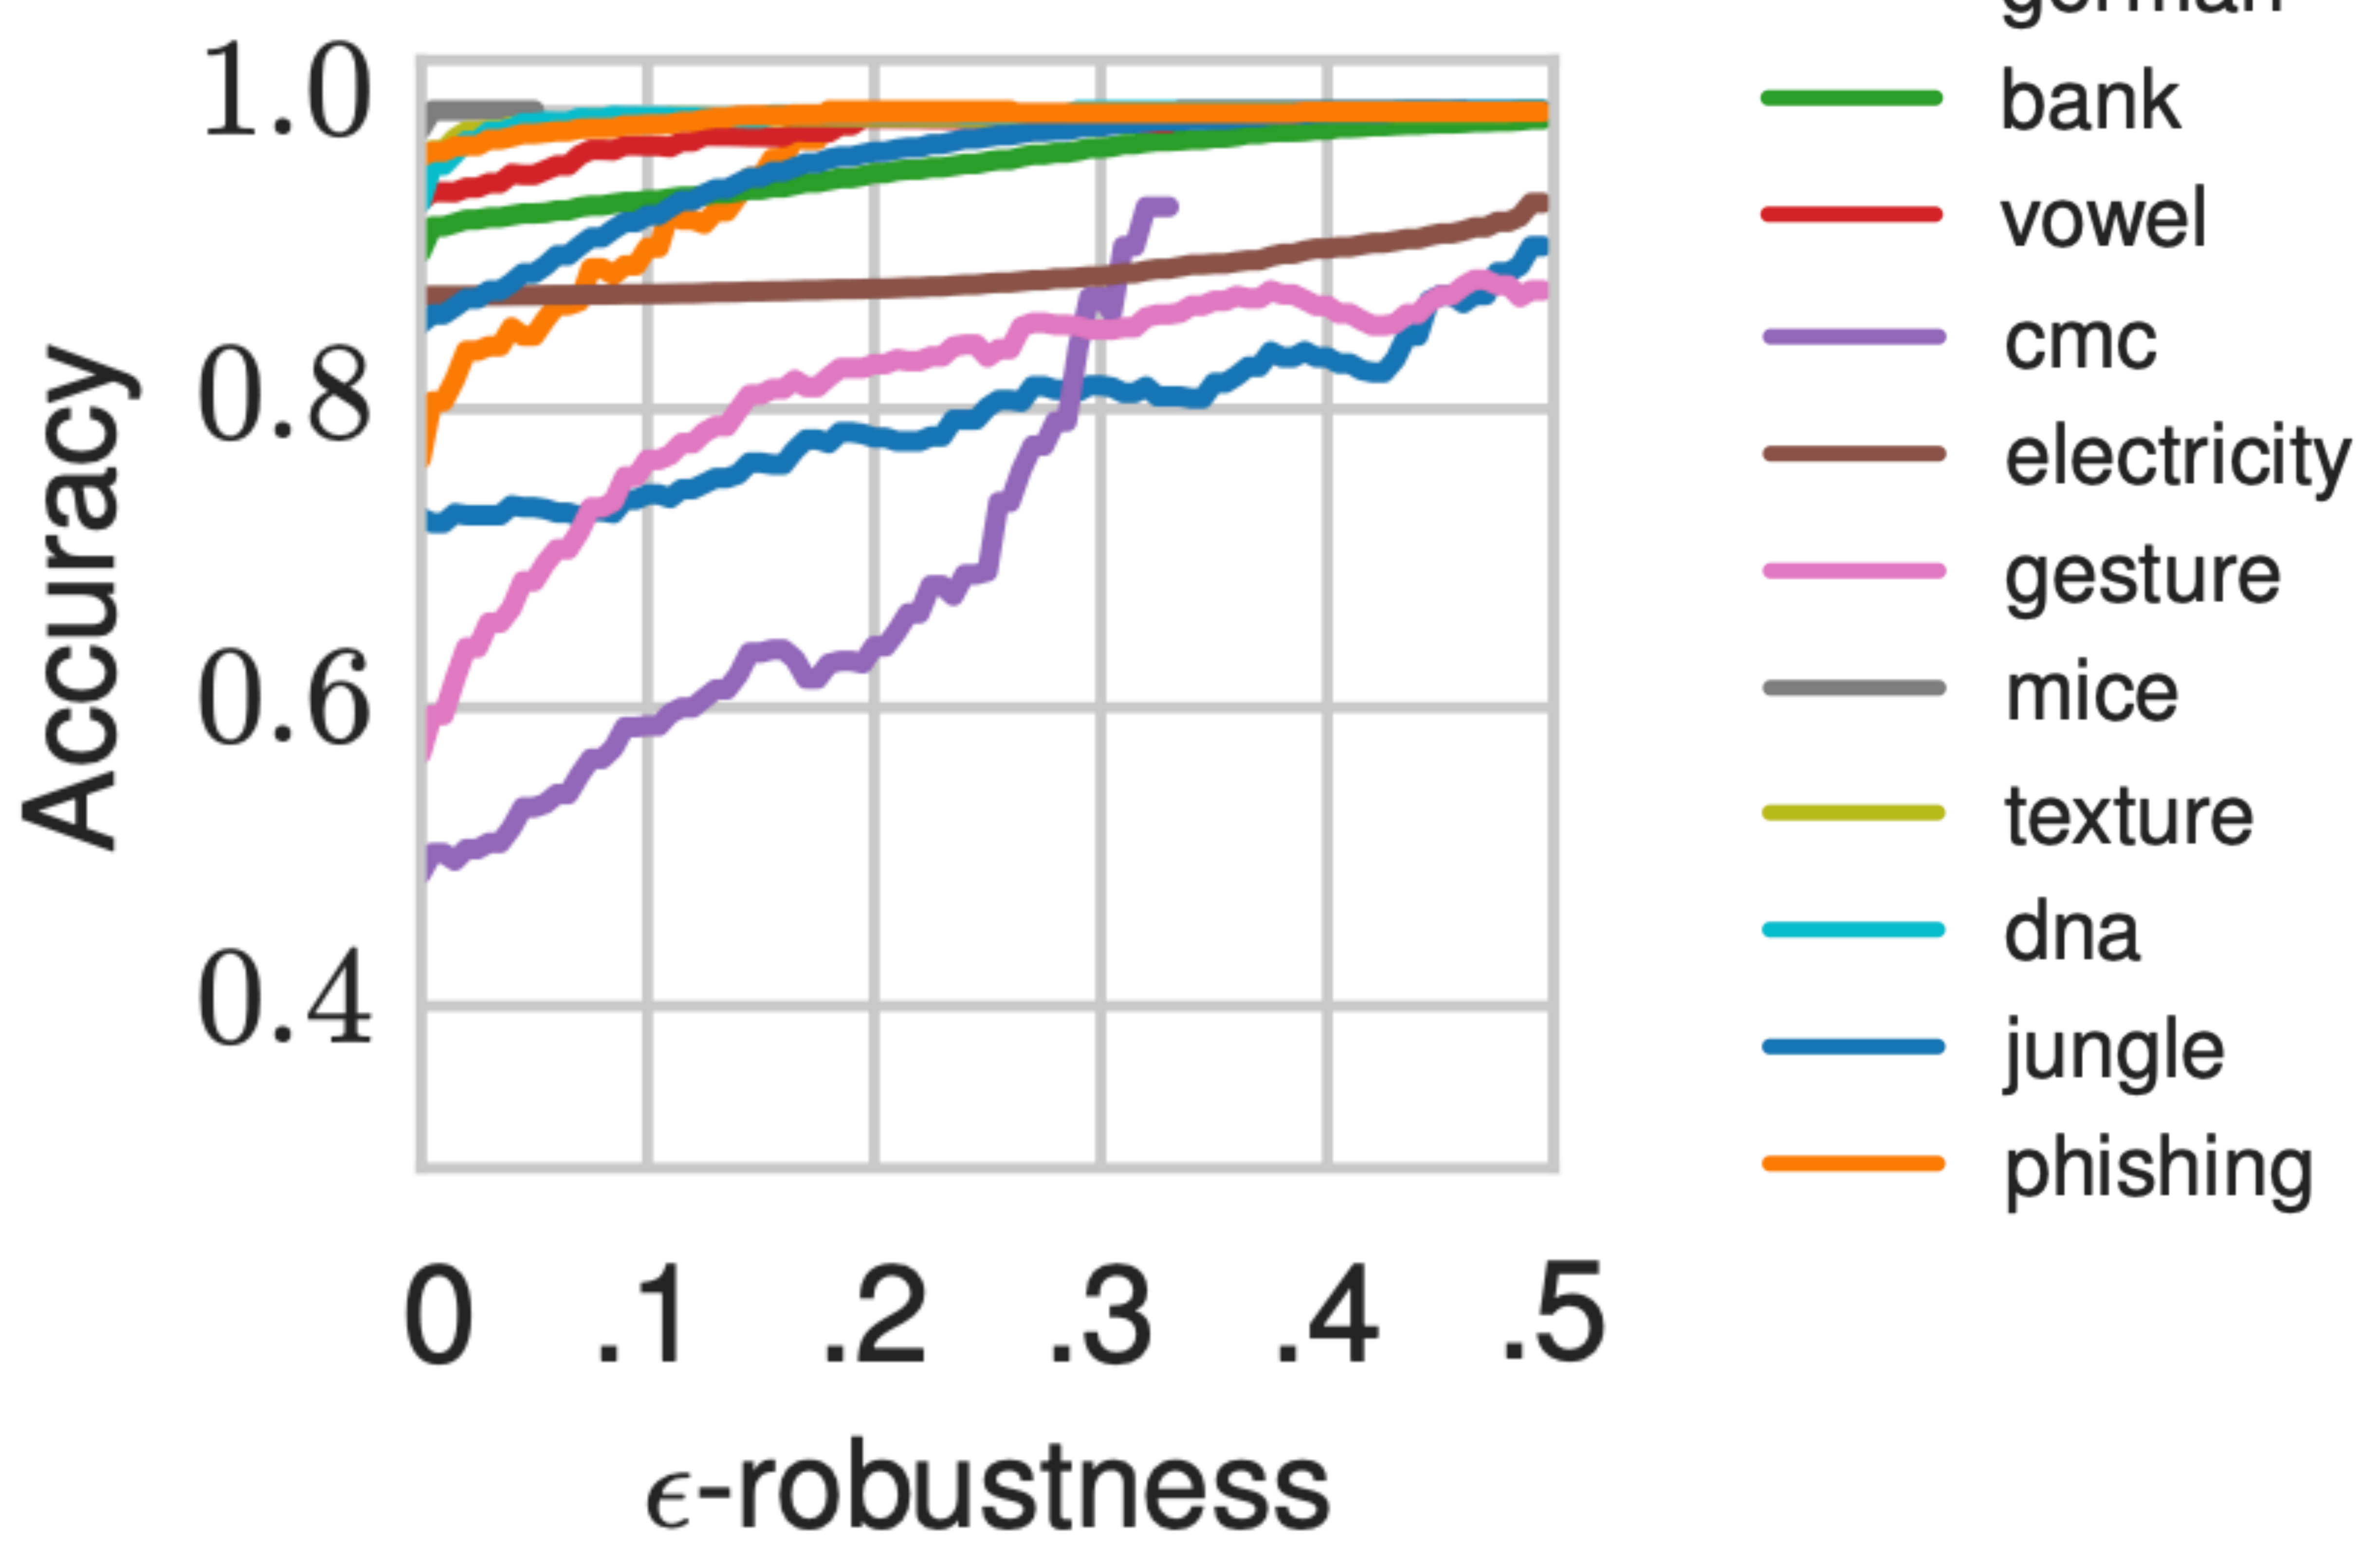
1.0
0.8
0.6
0.4
Accuracy
0
.1
.2
.3
.4
.5
ϵ-robustness
diabetes
german
bank
vowel
cmc
electricity
gesture
mice
texture
dna
jungle
phishing

ROBUSTNESS QUANTIFICATION

- correlates nicely with accuracy ✓
- works for different types of model architectures ✓

Adrián
Detavernier

Rodrigo
Lassance

2025 – ...

[6, 7, ...]

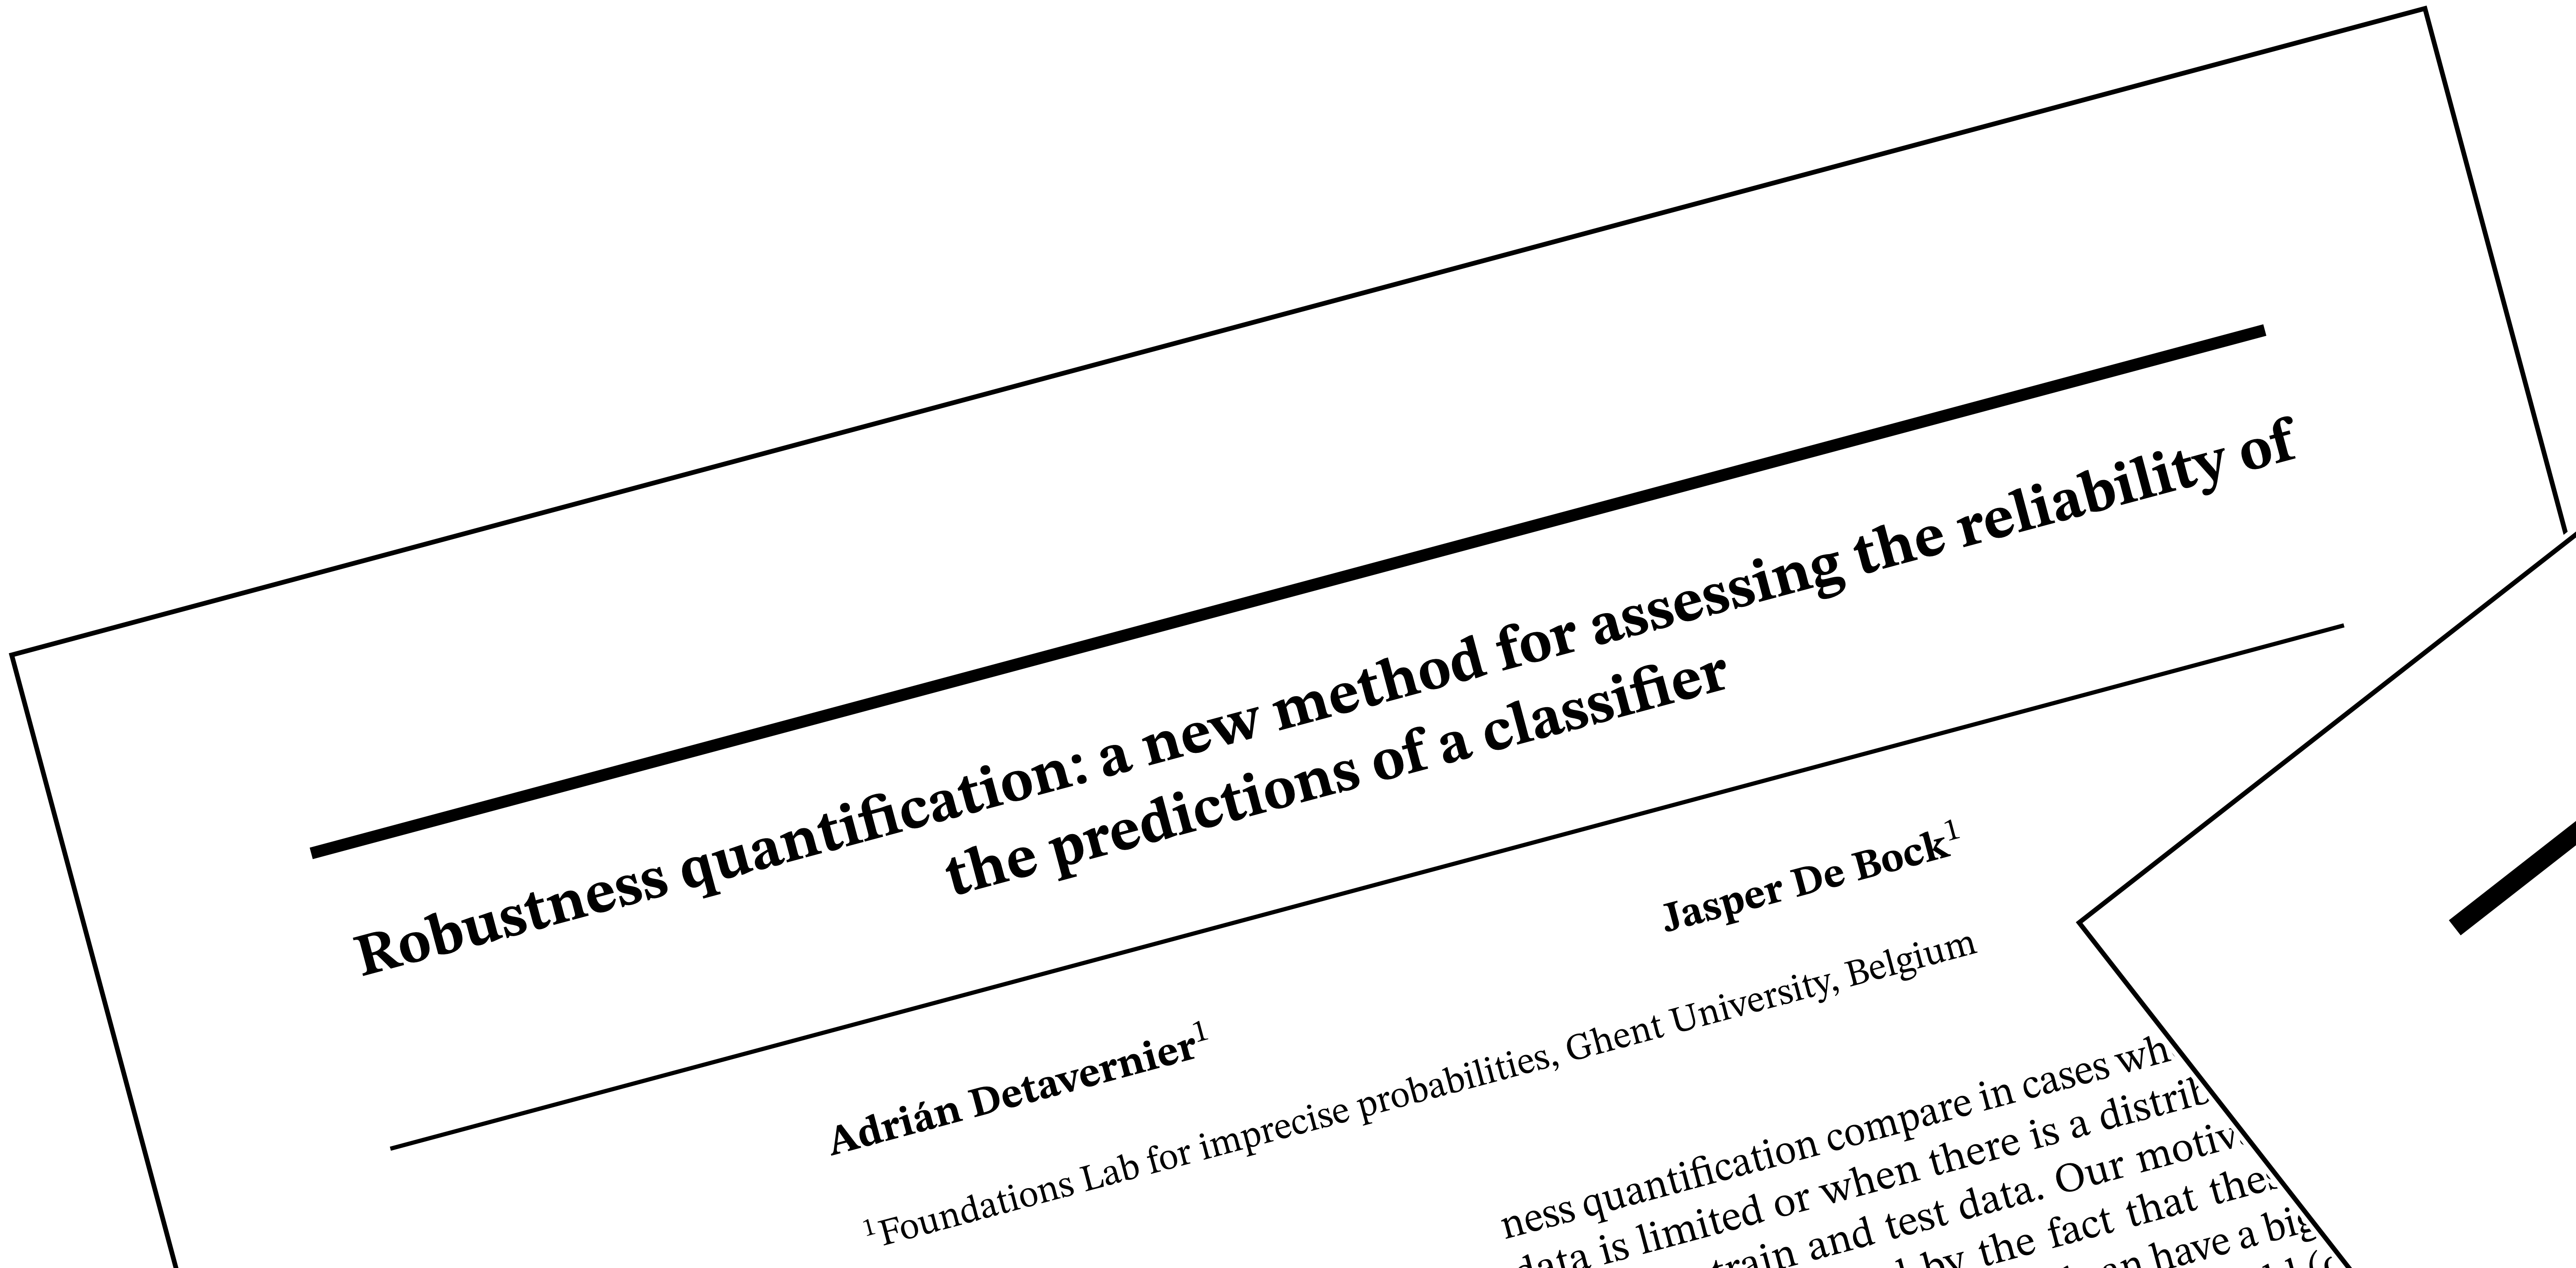

Robustness quantification: a new method for assessing the reliability of the predictions of a classifier

Adrián Detavernier[1] **Jasper De Bock**[1]

[1]Foundations Lab for imprecise probabilities, Ghent University, Belgium

Robustness and uncertainty: two complementary aspects of the reliability of the predictions of a classifier

Adrián Detavernier
Ghent University
Belgium

Jasper De Bock
Foundations Lab for imprecise probabilities
Ghent University
Belgium

Abstract

CLASSIFICATION … is unreliable

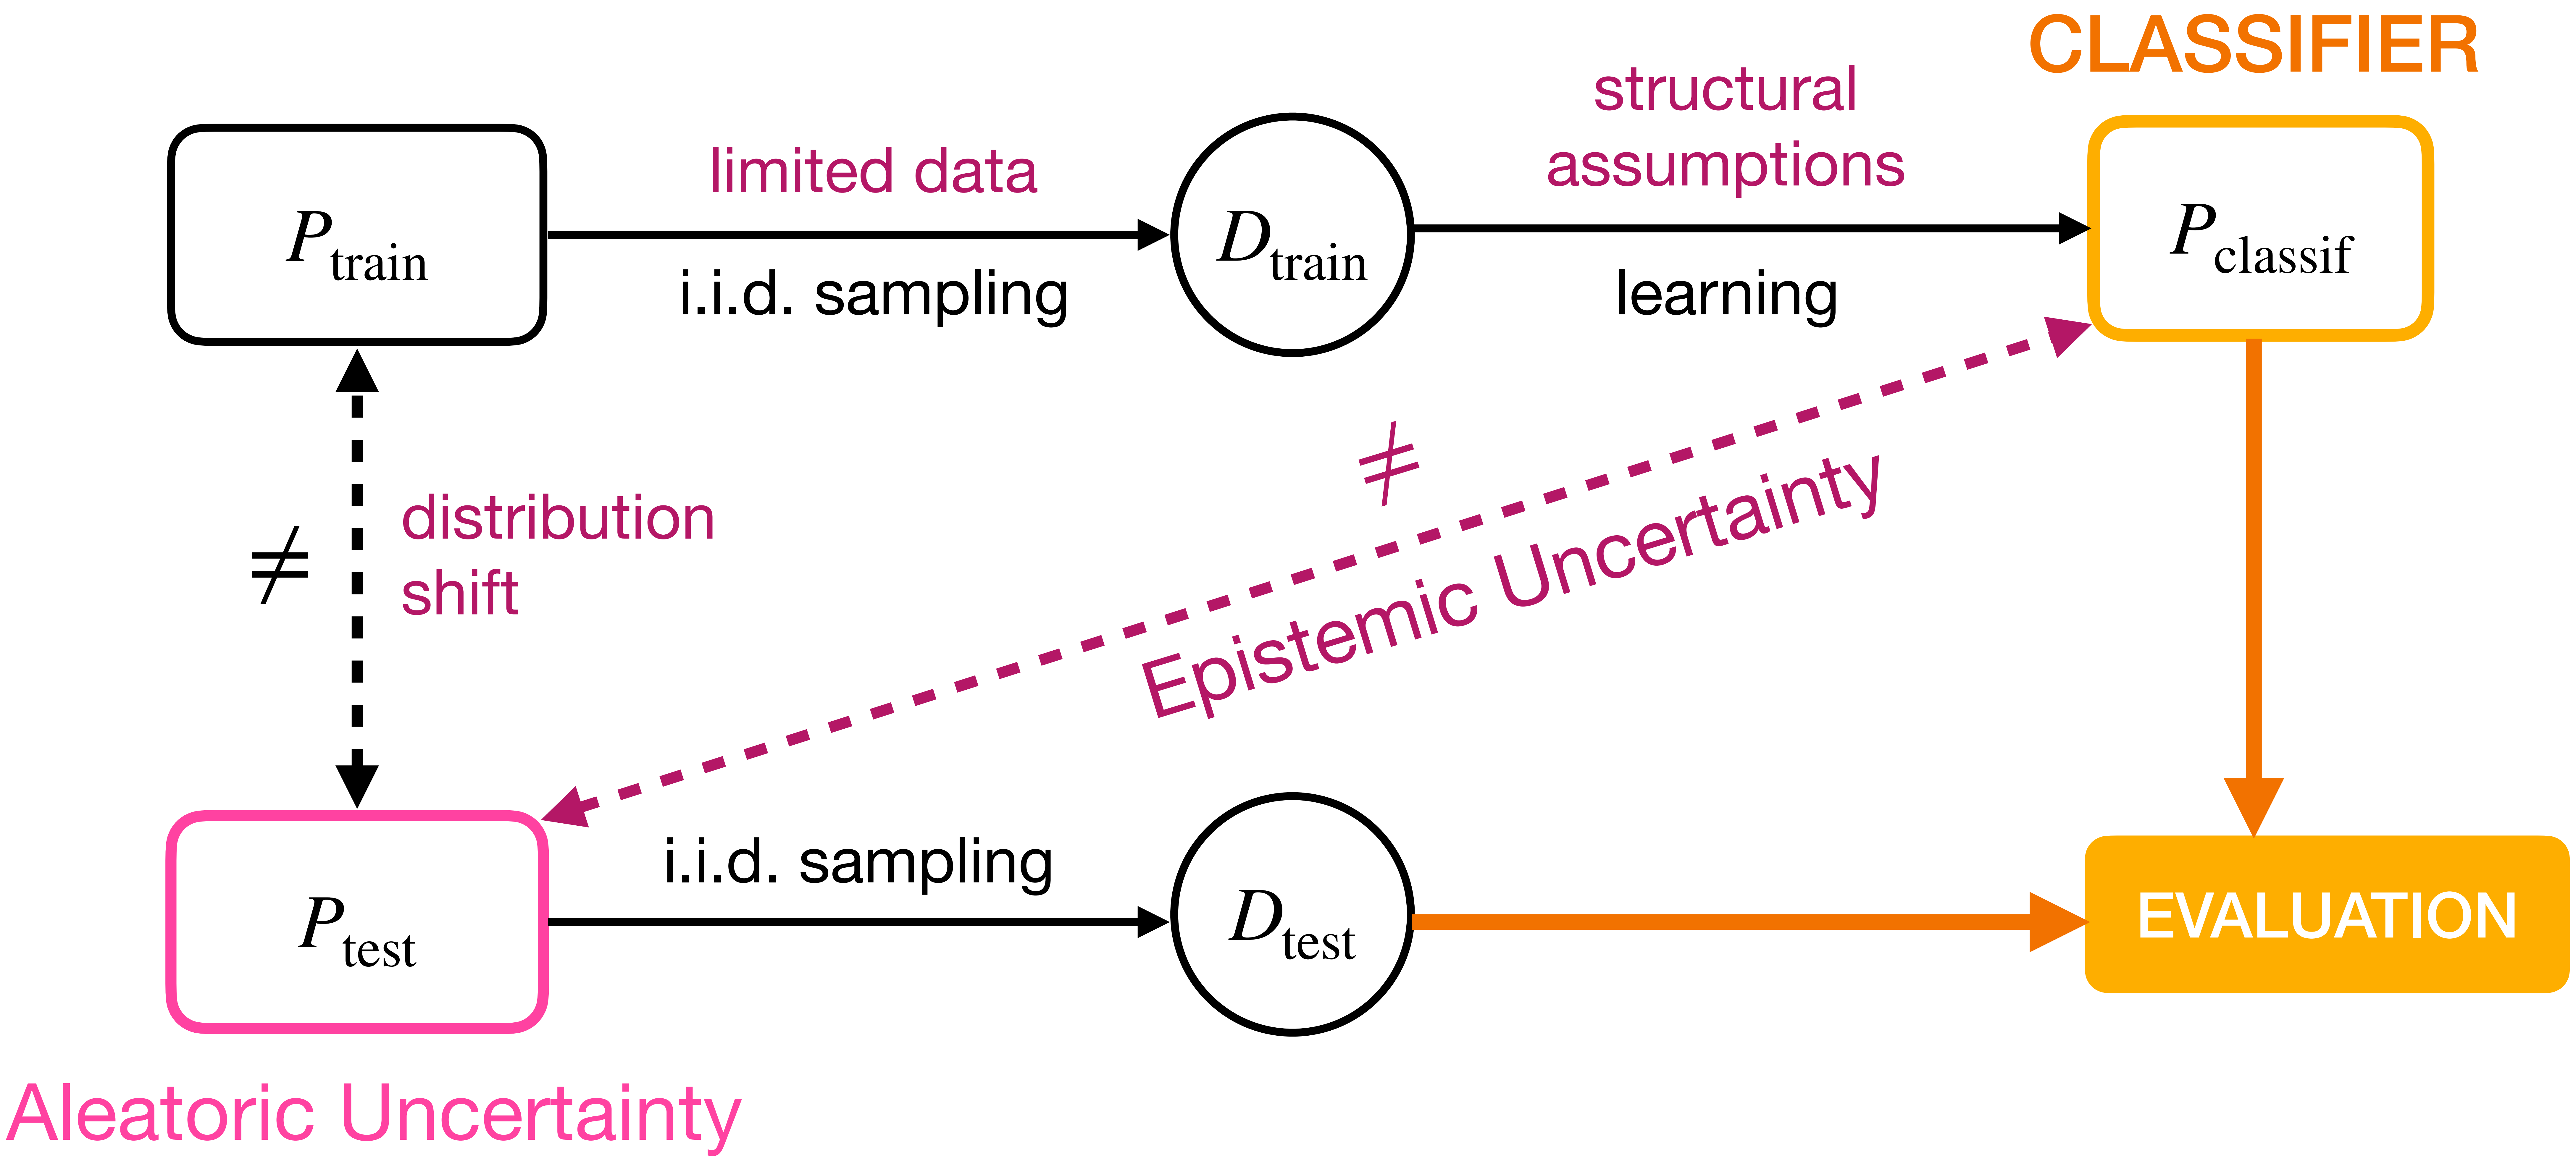

UNCERTAINTY QUANTIFICATION

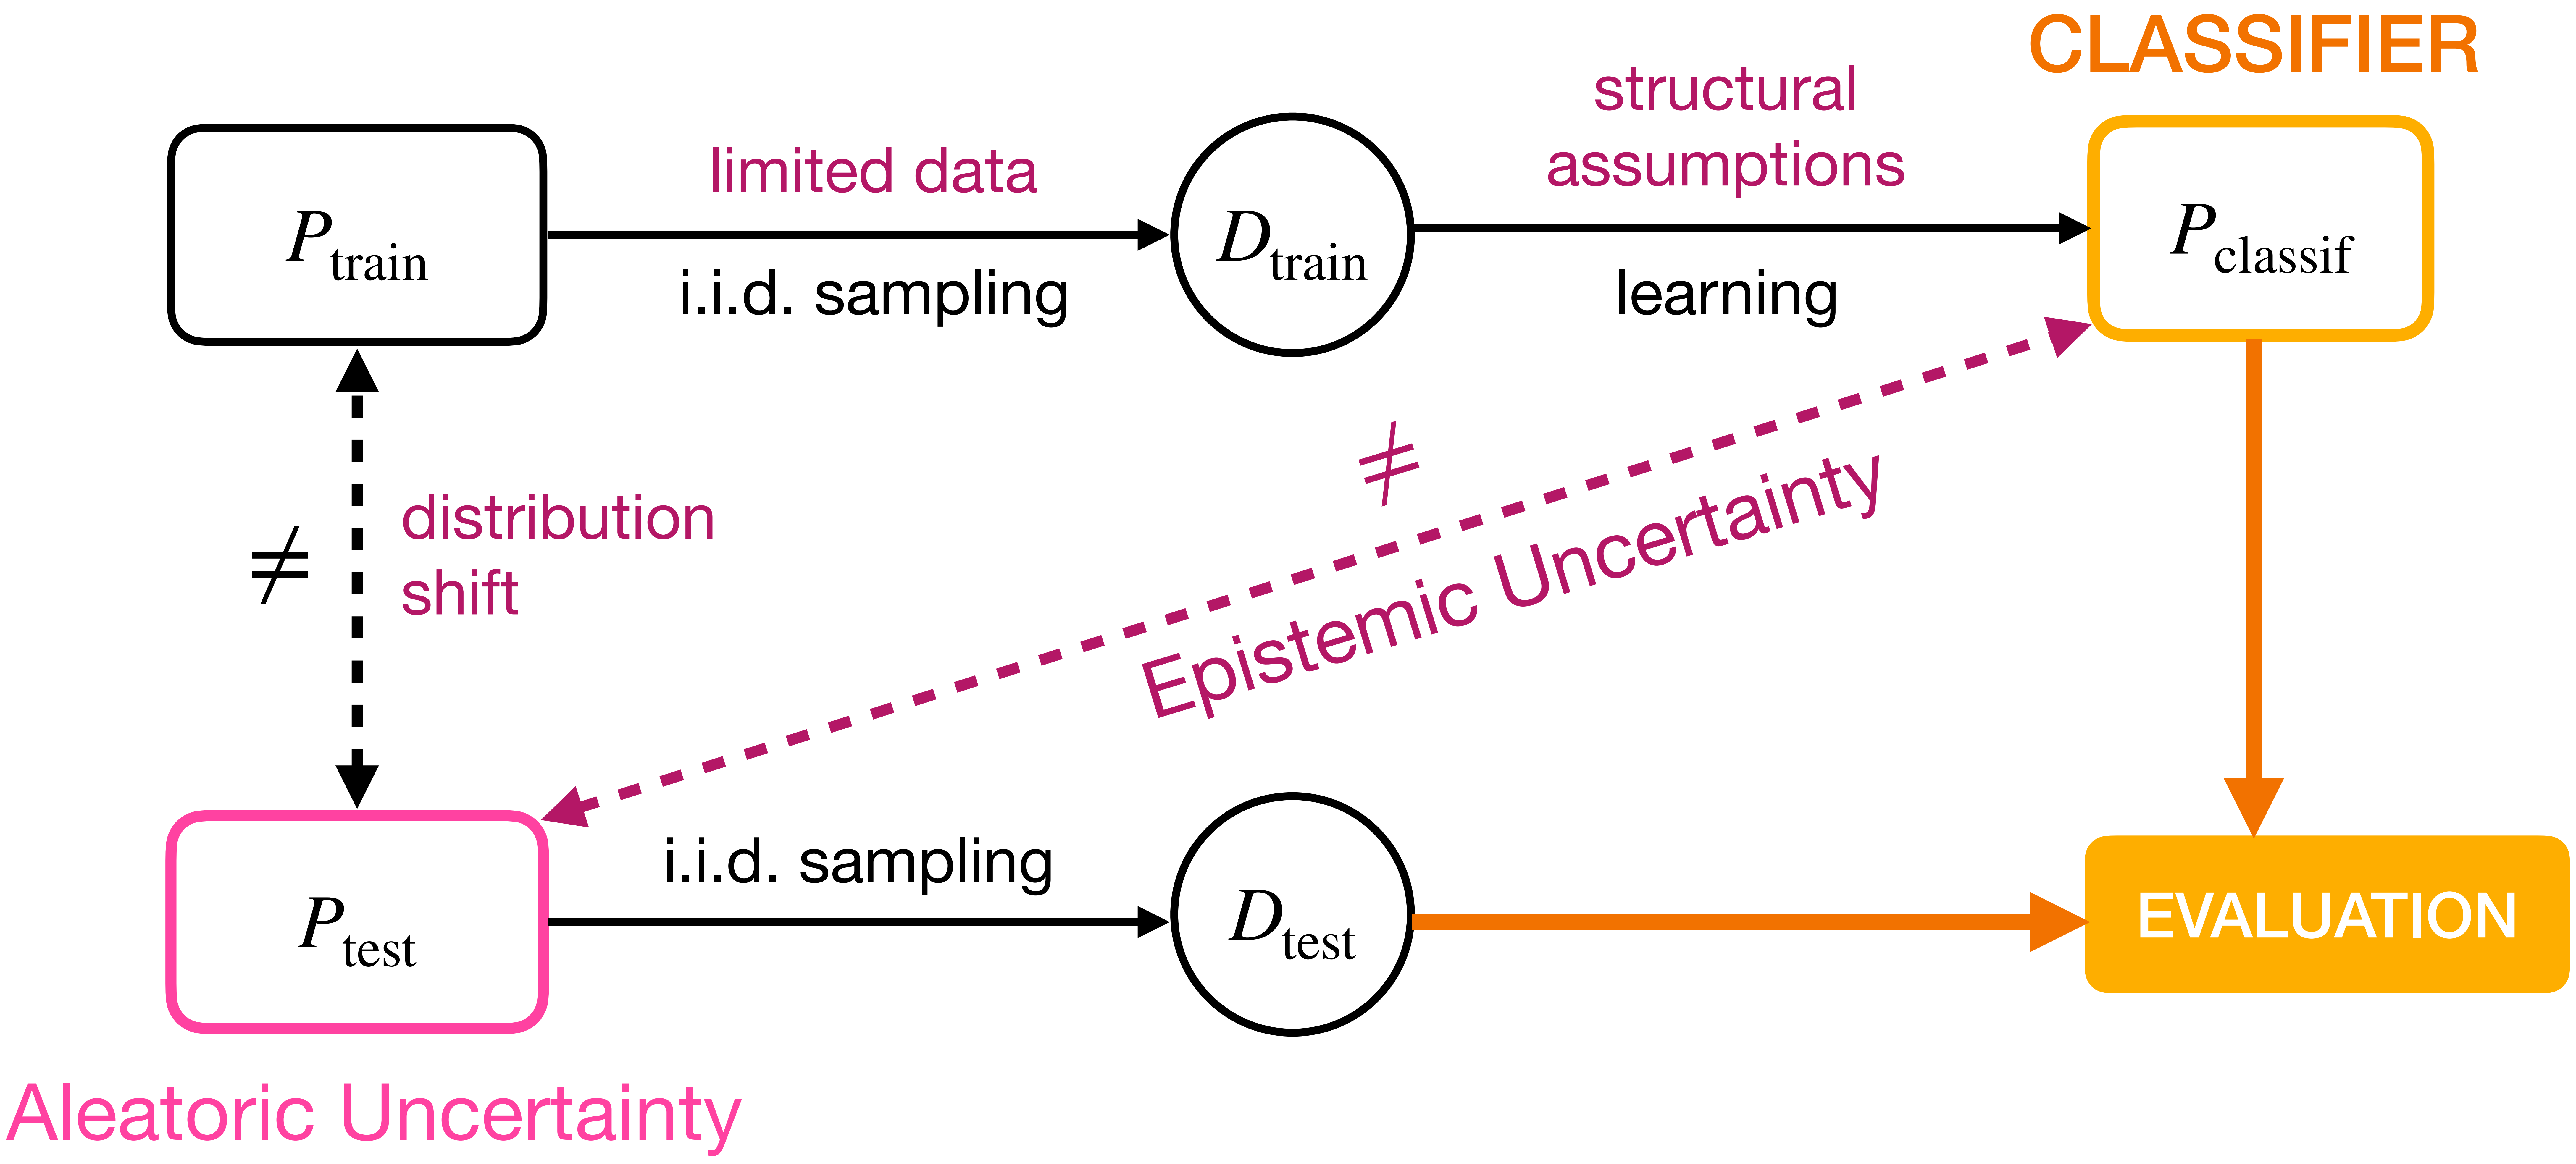

ROBUSTNESS QUANTIFICATION

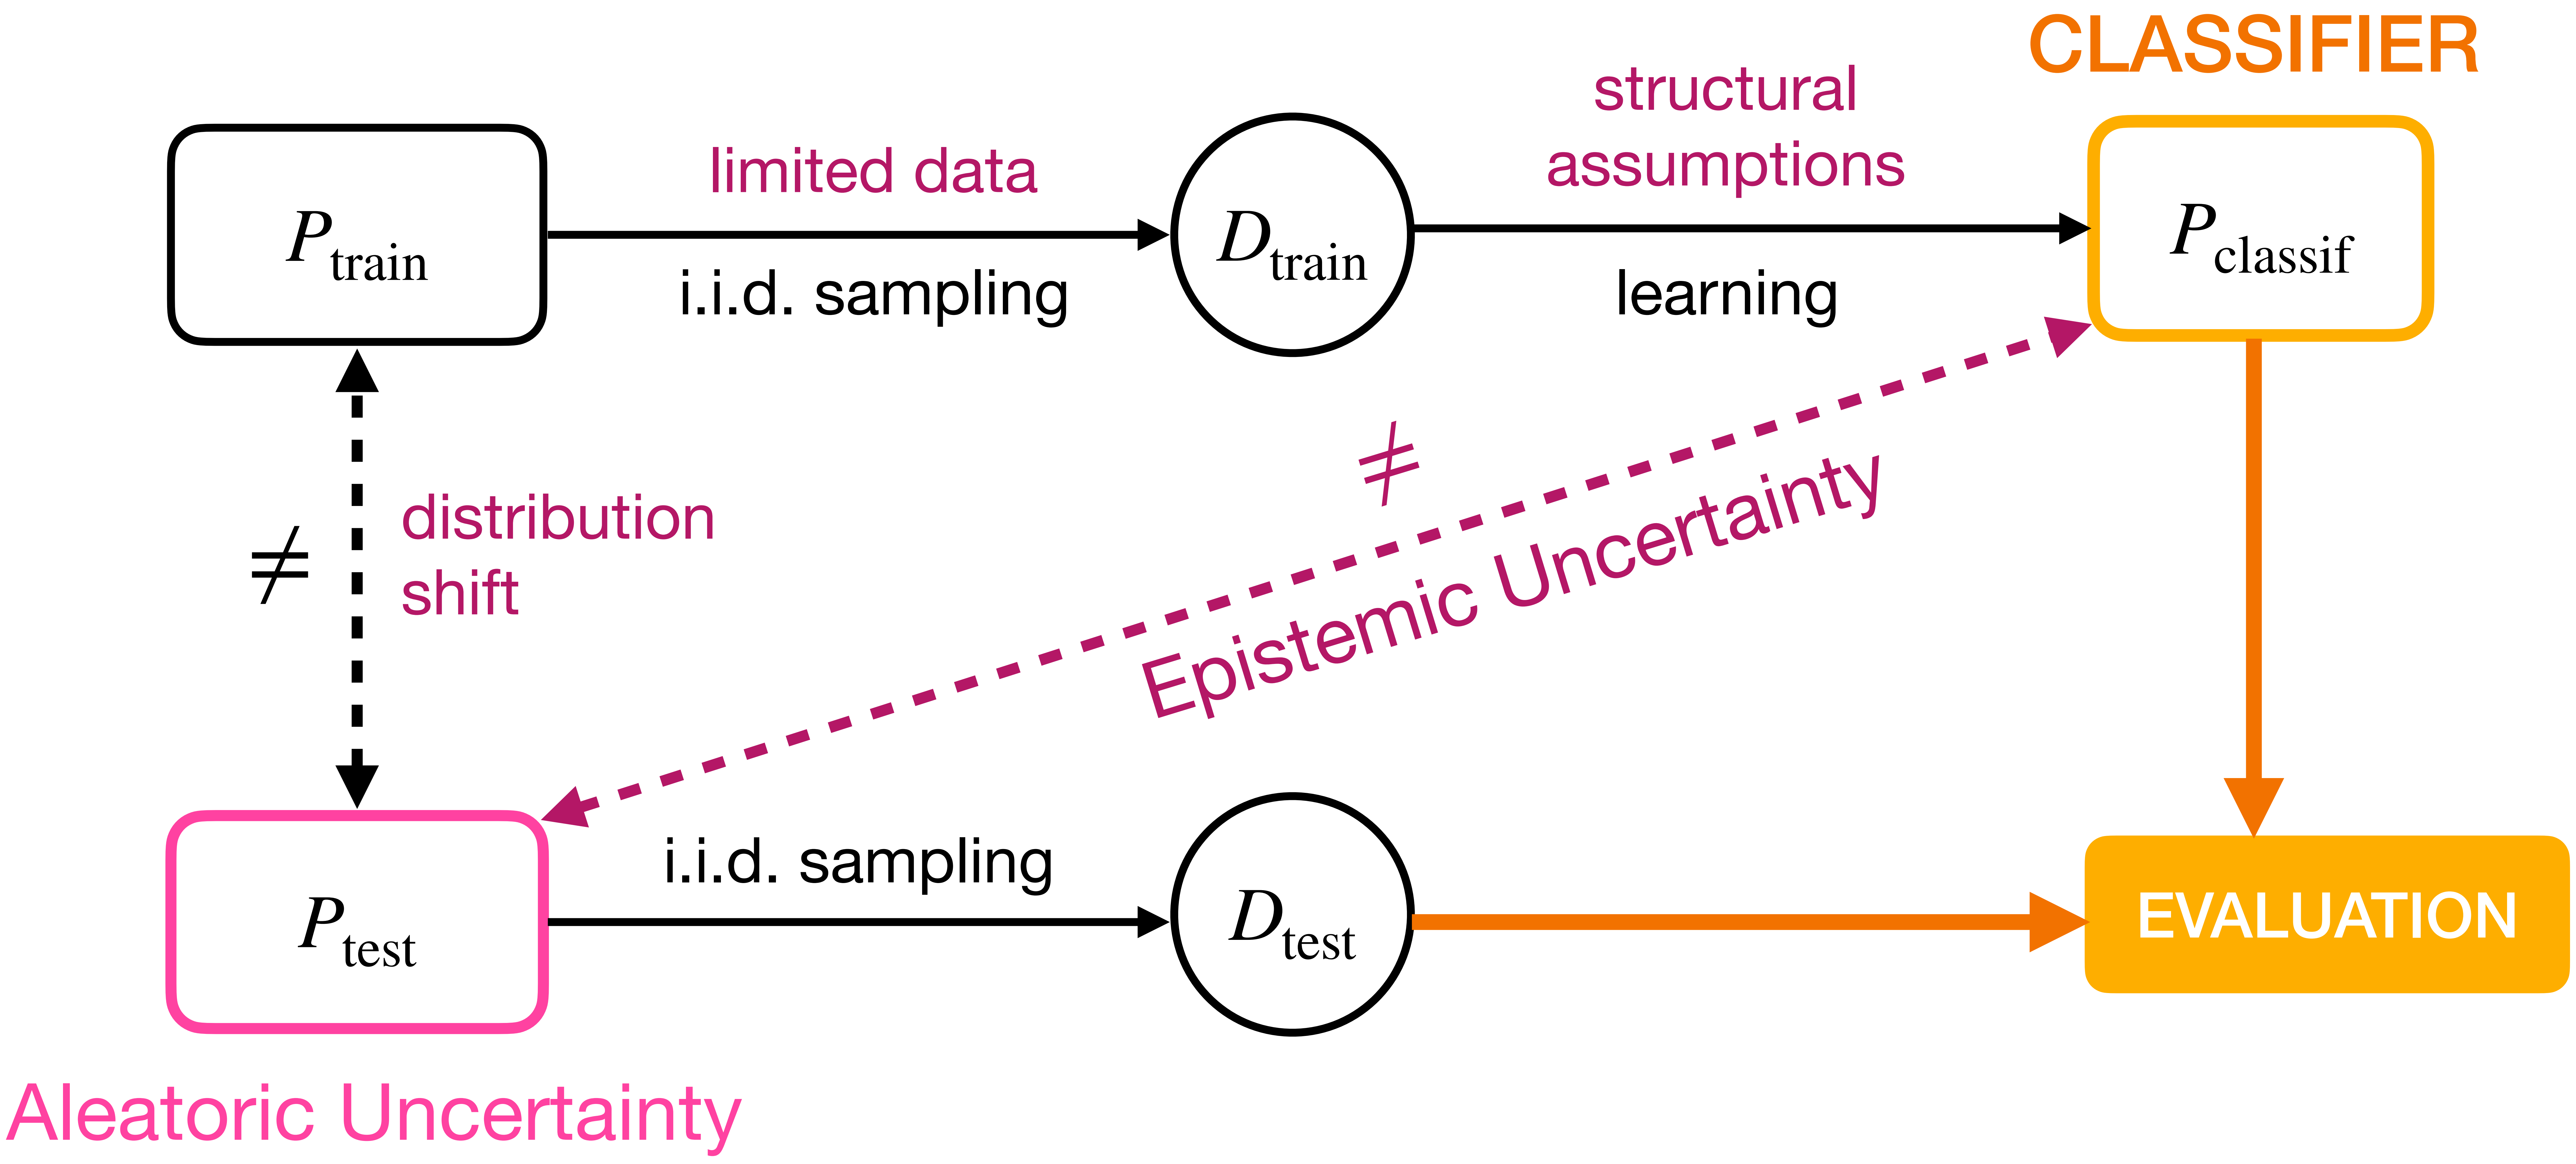

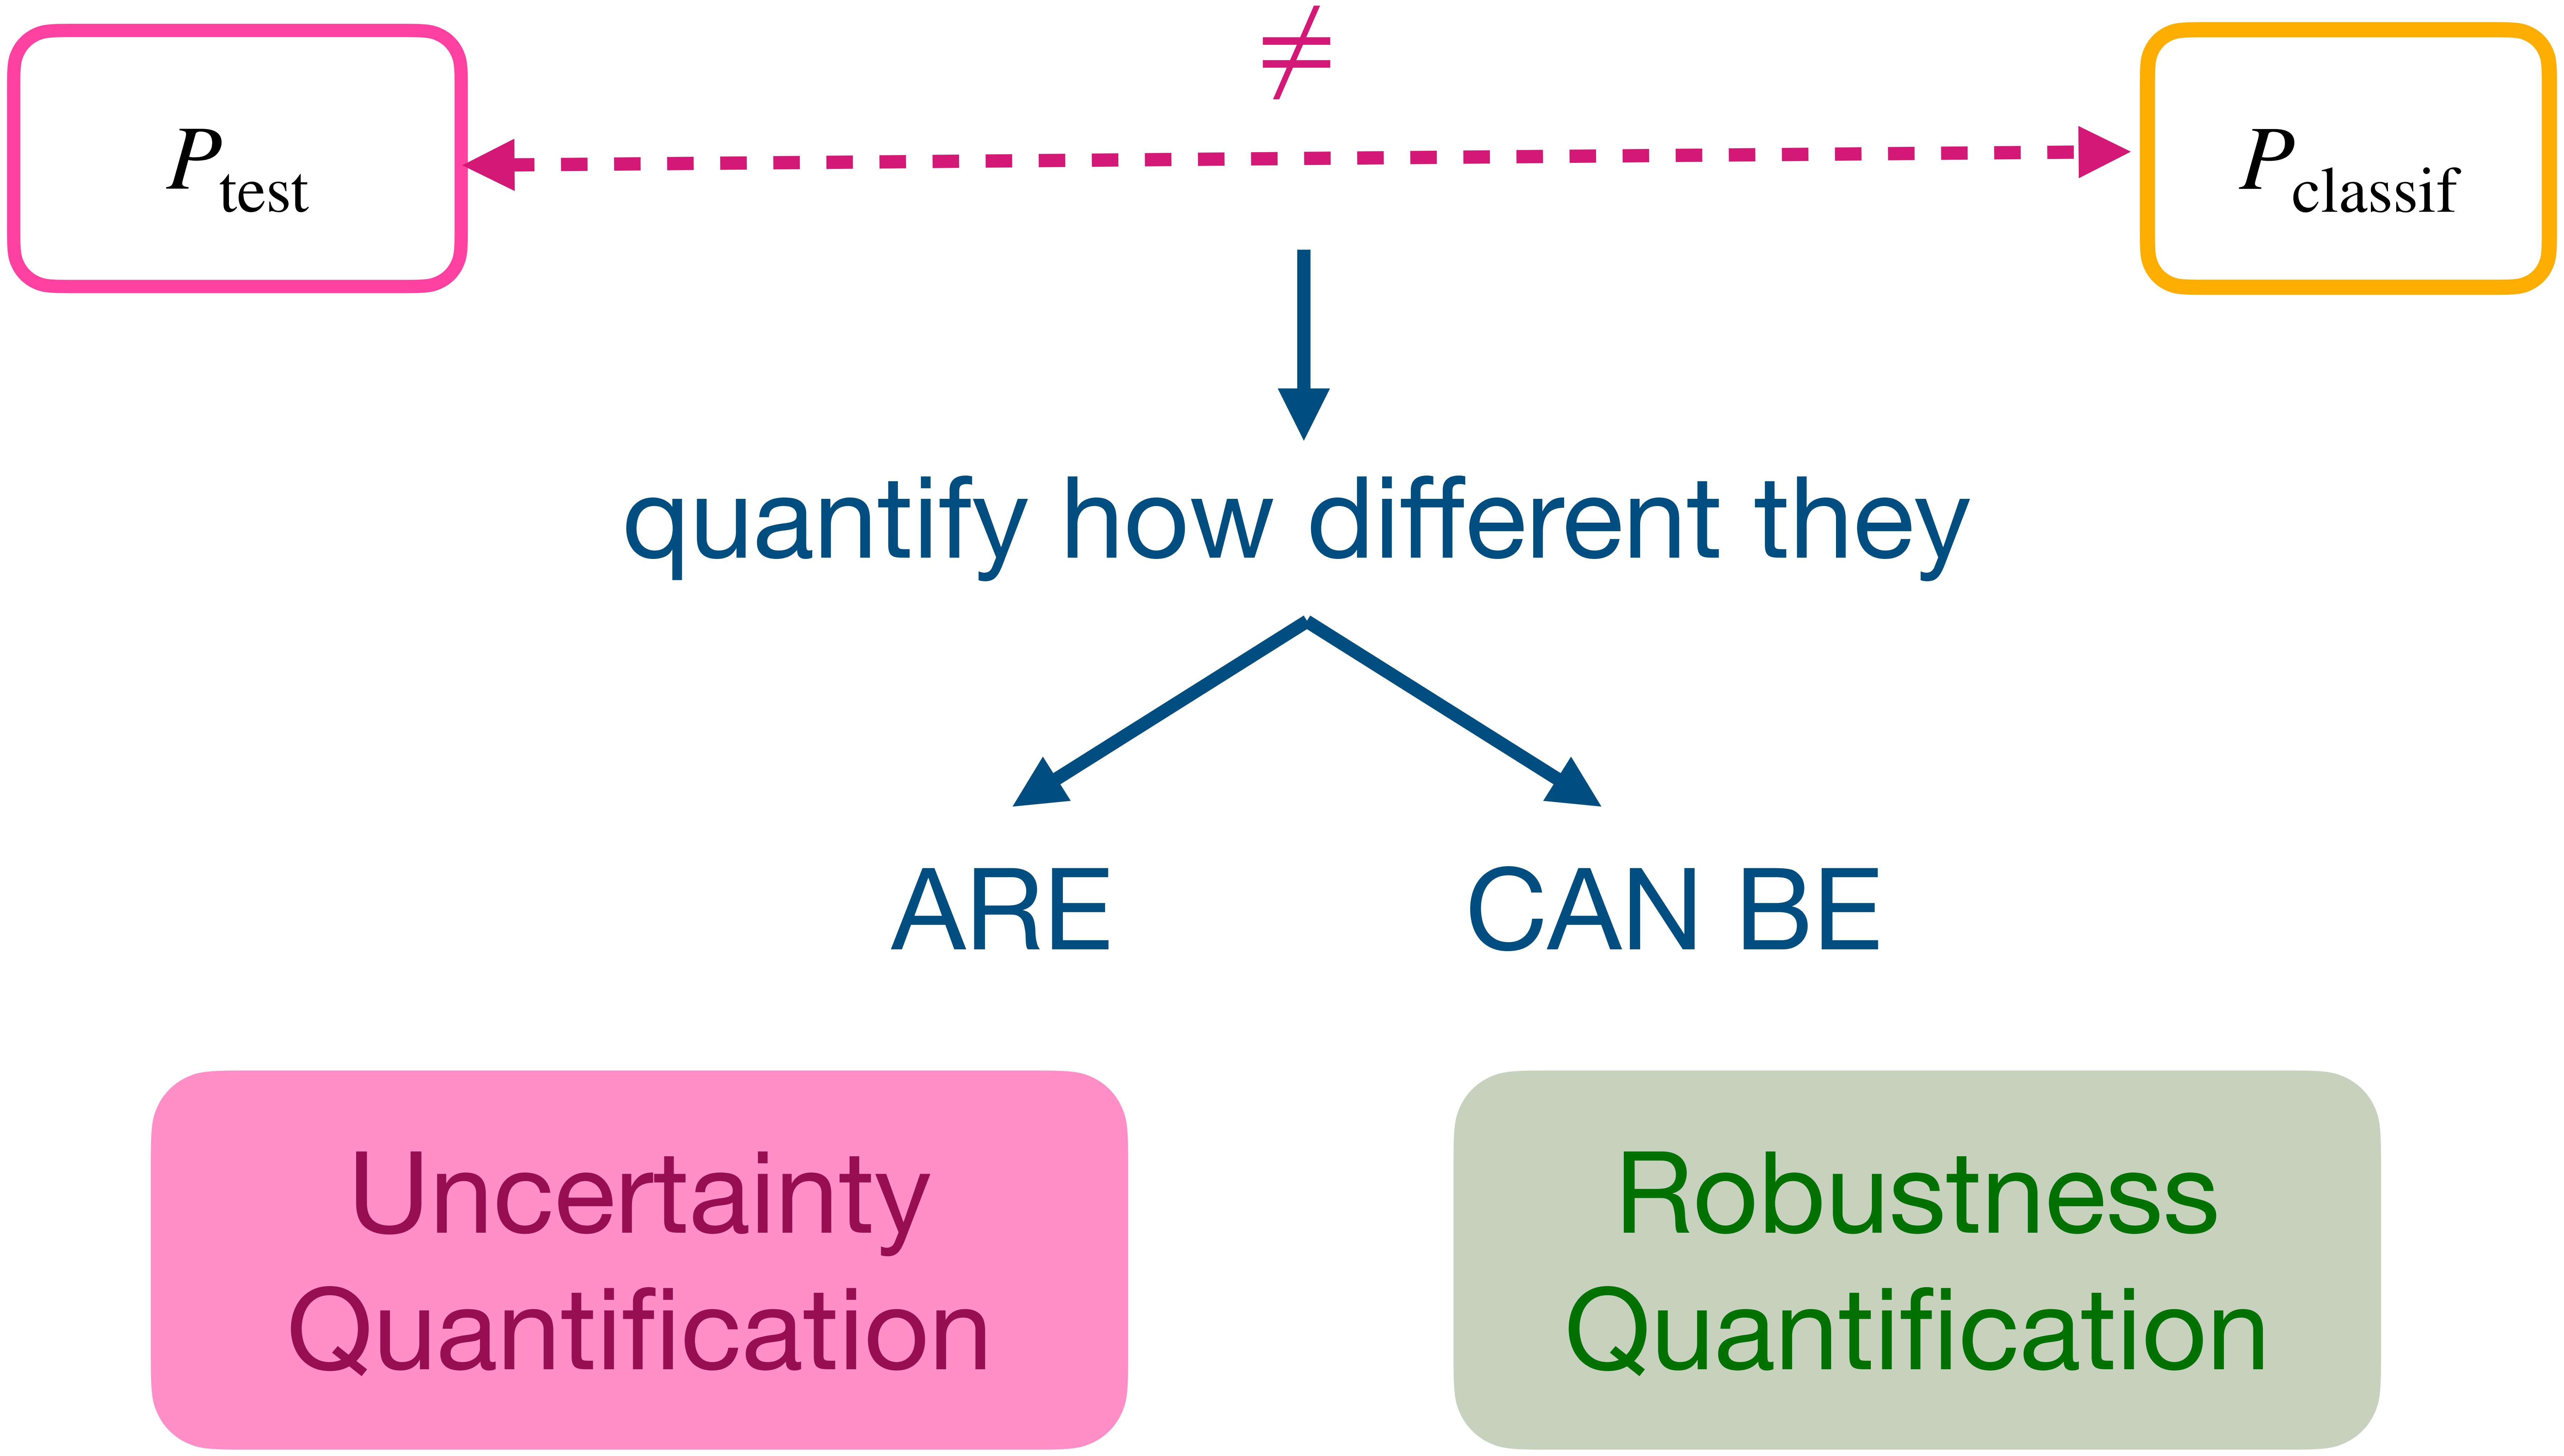
P_{test}
$\neq$
$P_{classif}$
quantify how different they
ARE
CAN BE
Uncertainty Quantification
Robustness Quantification

ROBUSTNESS QUANTIFICATION

- correlates nicely with accuracy
- works for different types of model architectures
- is conceptually different from UQ ✓

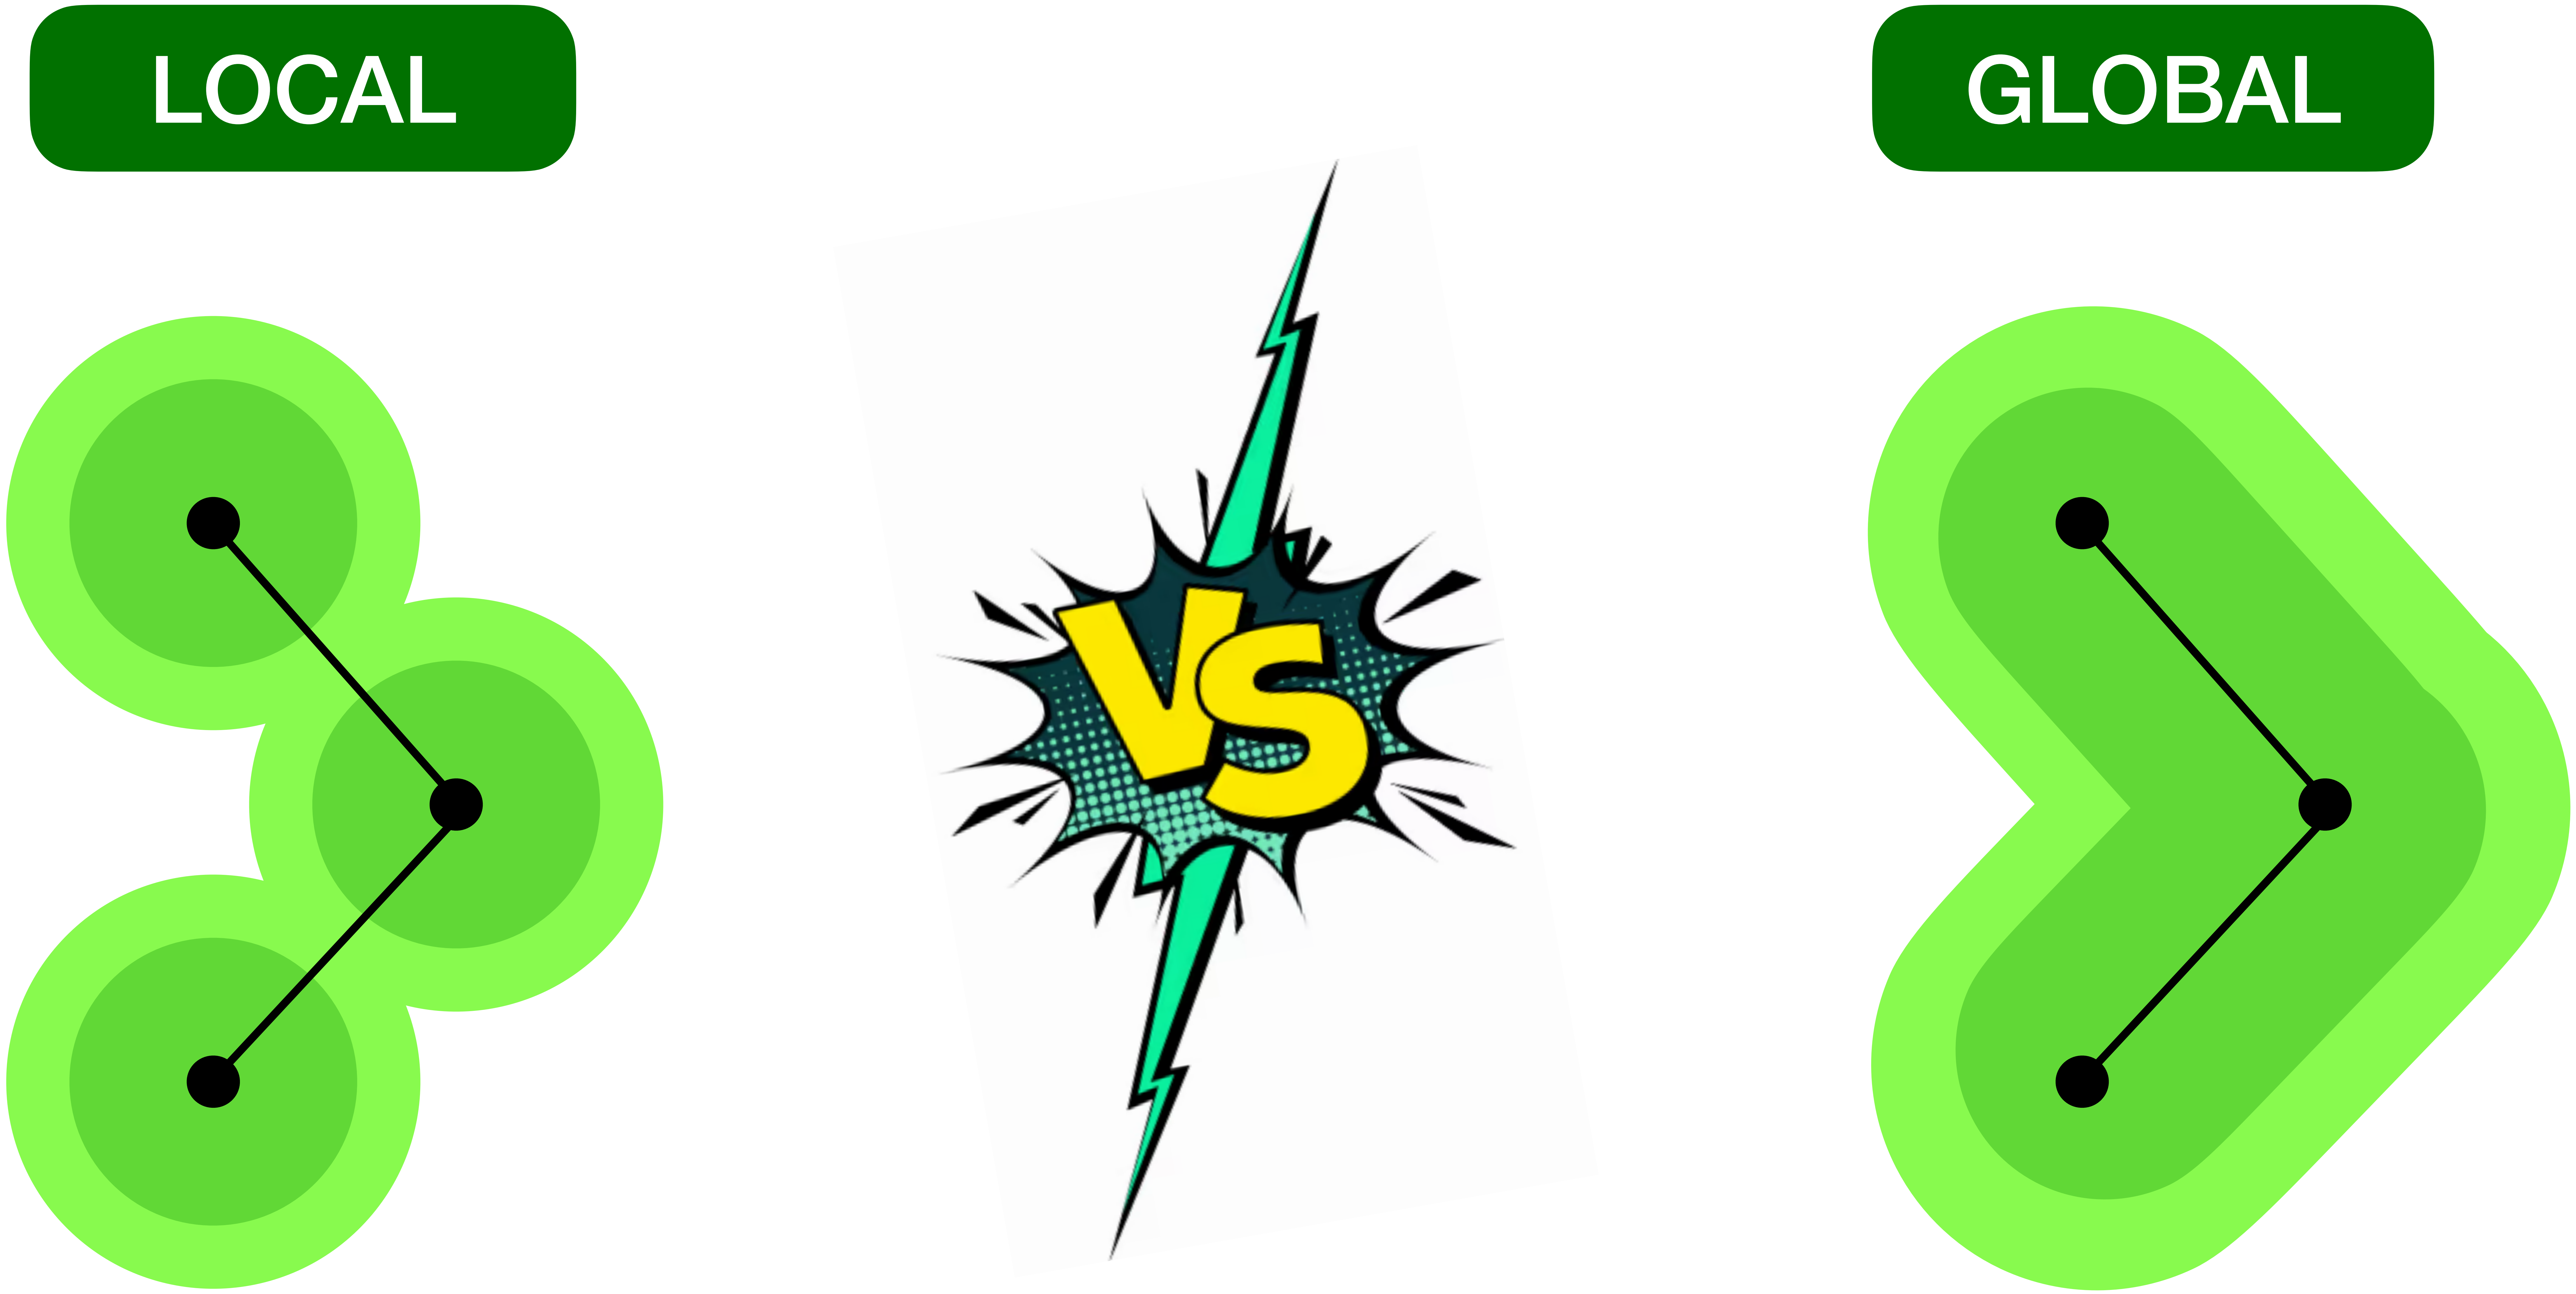
LOCAL
VS
GLOBAL

ϵ-CONTAMINATION

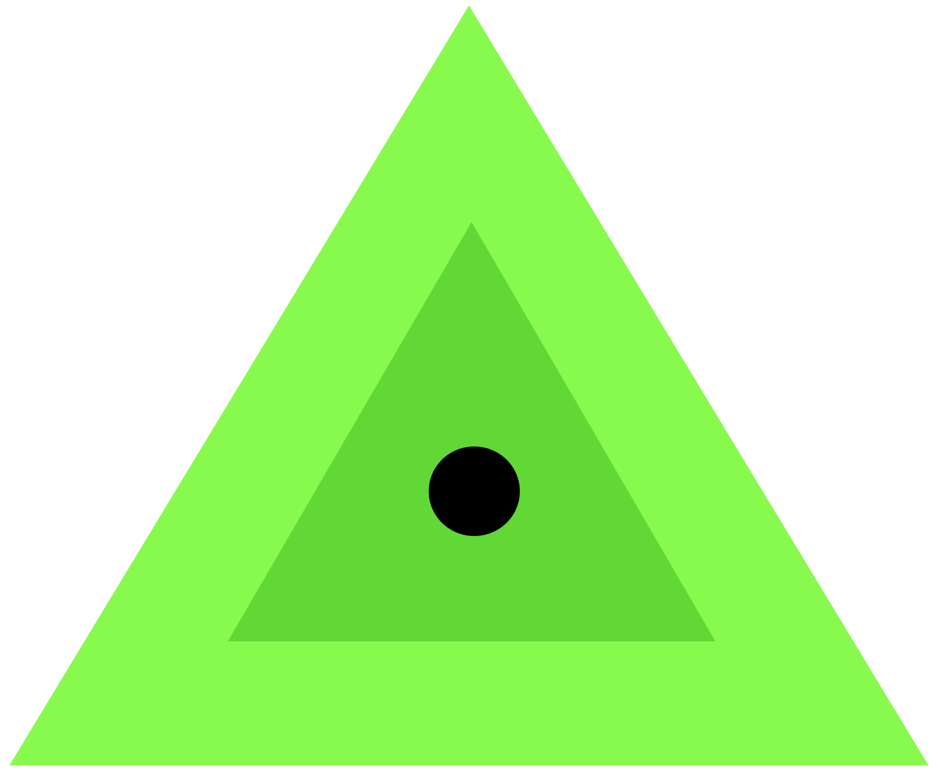

$$\mathscr{P}_\epsilon$$

$$\parallel$$

$$\left\{(1-\epsilon)P_{\text{classif}} + \epsilon P : P \in \Delta\right\}$$

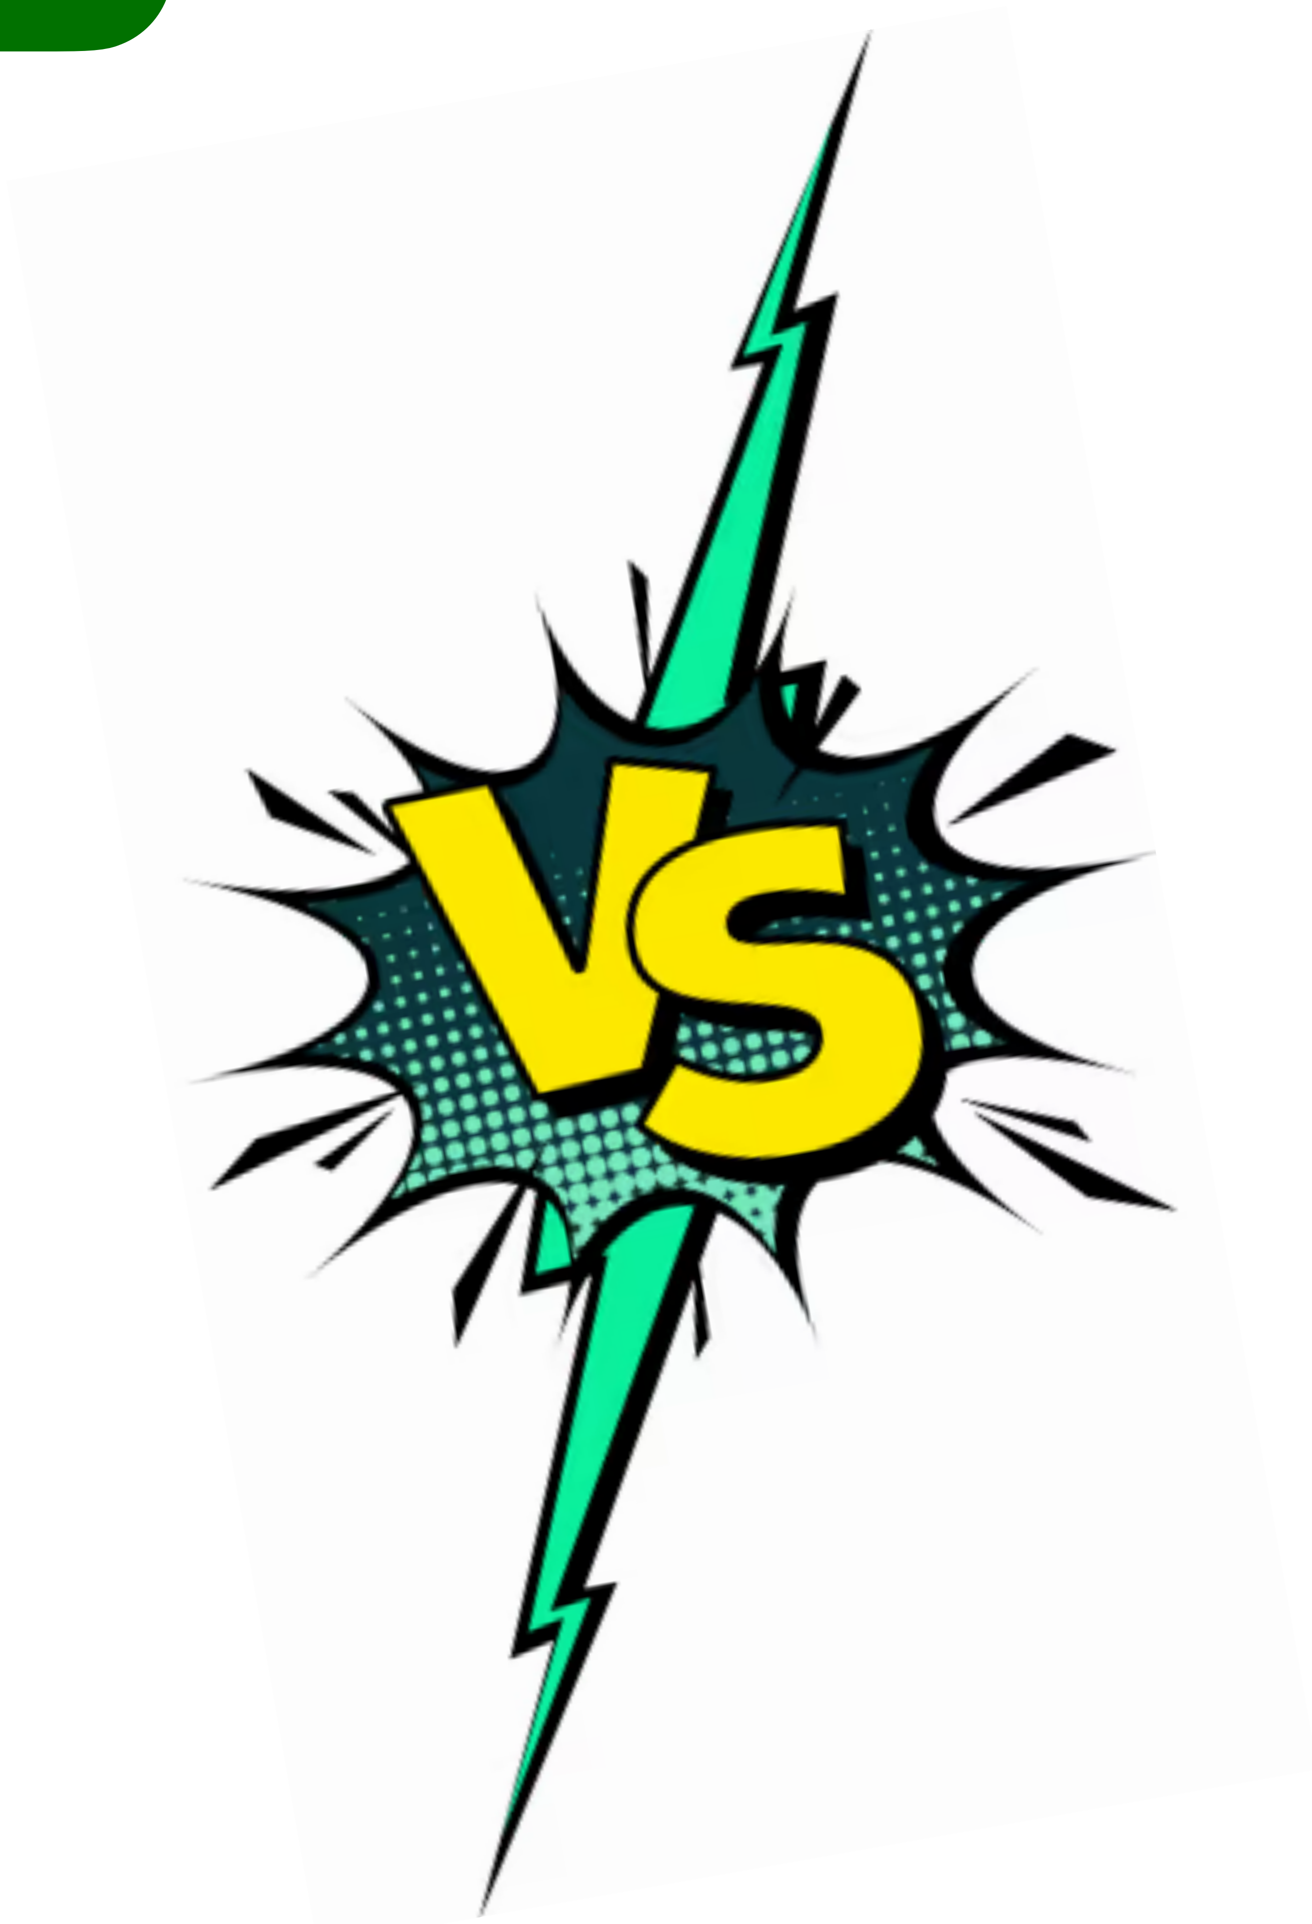

OTHER STUFF

distance-based, ...

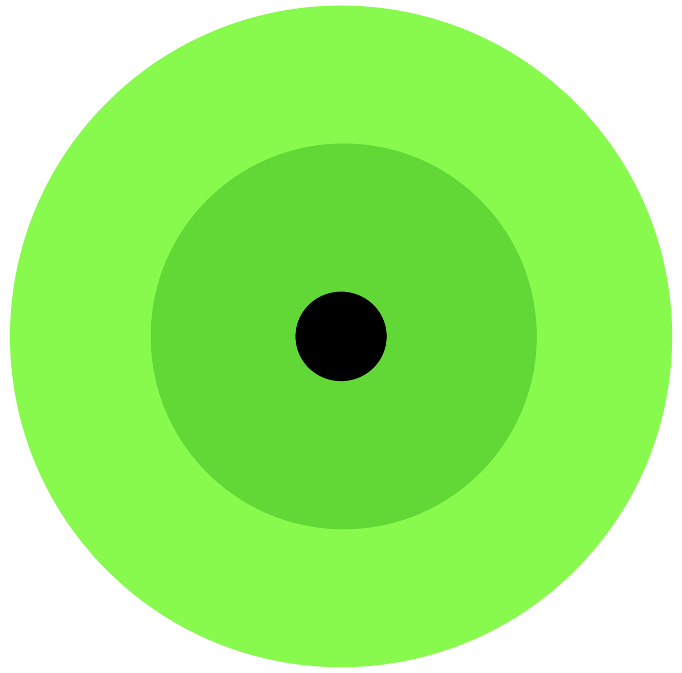

$$\mathscr{P}_\epsilon$$

$$\parallel$$

$$\left\{P \in \Delta : \mathrm{d}(P_{\text{classif}}, P) < \epsilon\right\}$$

ϵ-CONTAMINATION

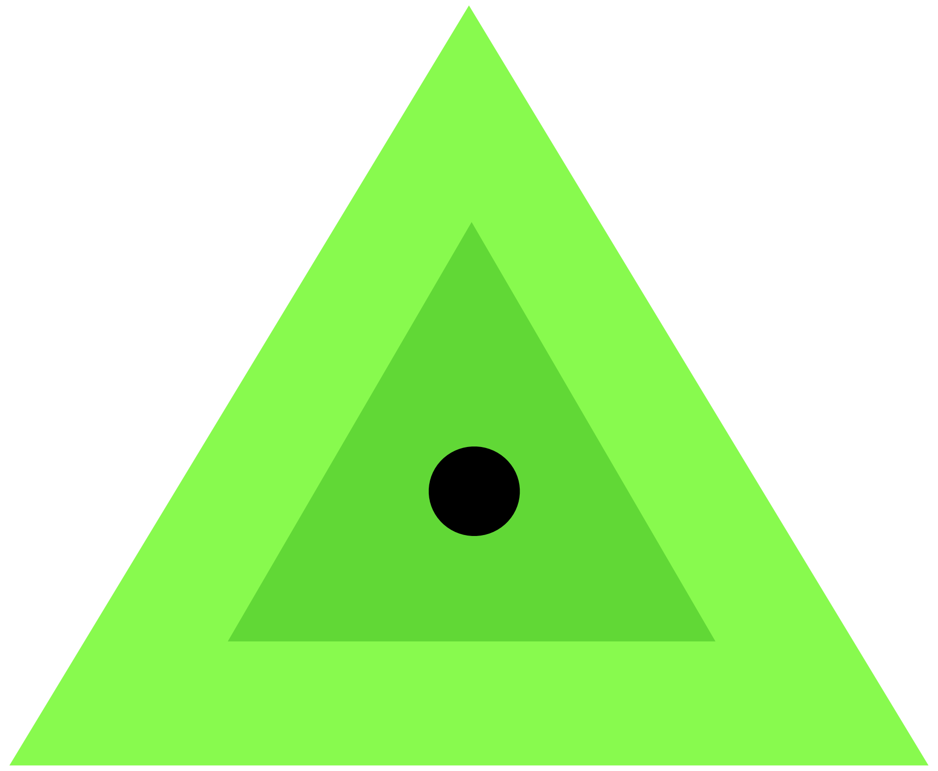

$$\mathscr{P}_\epsilon$$
$$\|$$
$$\left\{(1-\epsilon)P_{\text{classif}} + \epsilon P : P \in \Delta\right\}$$

+

GLOBAL

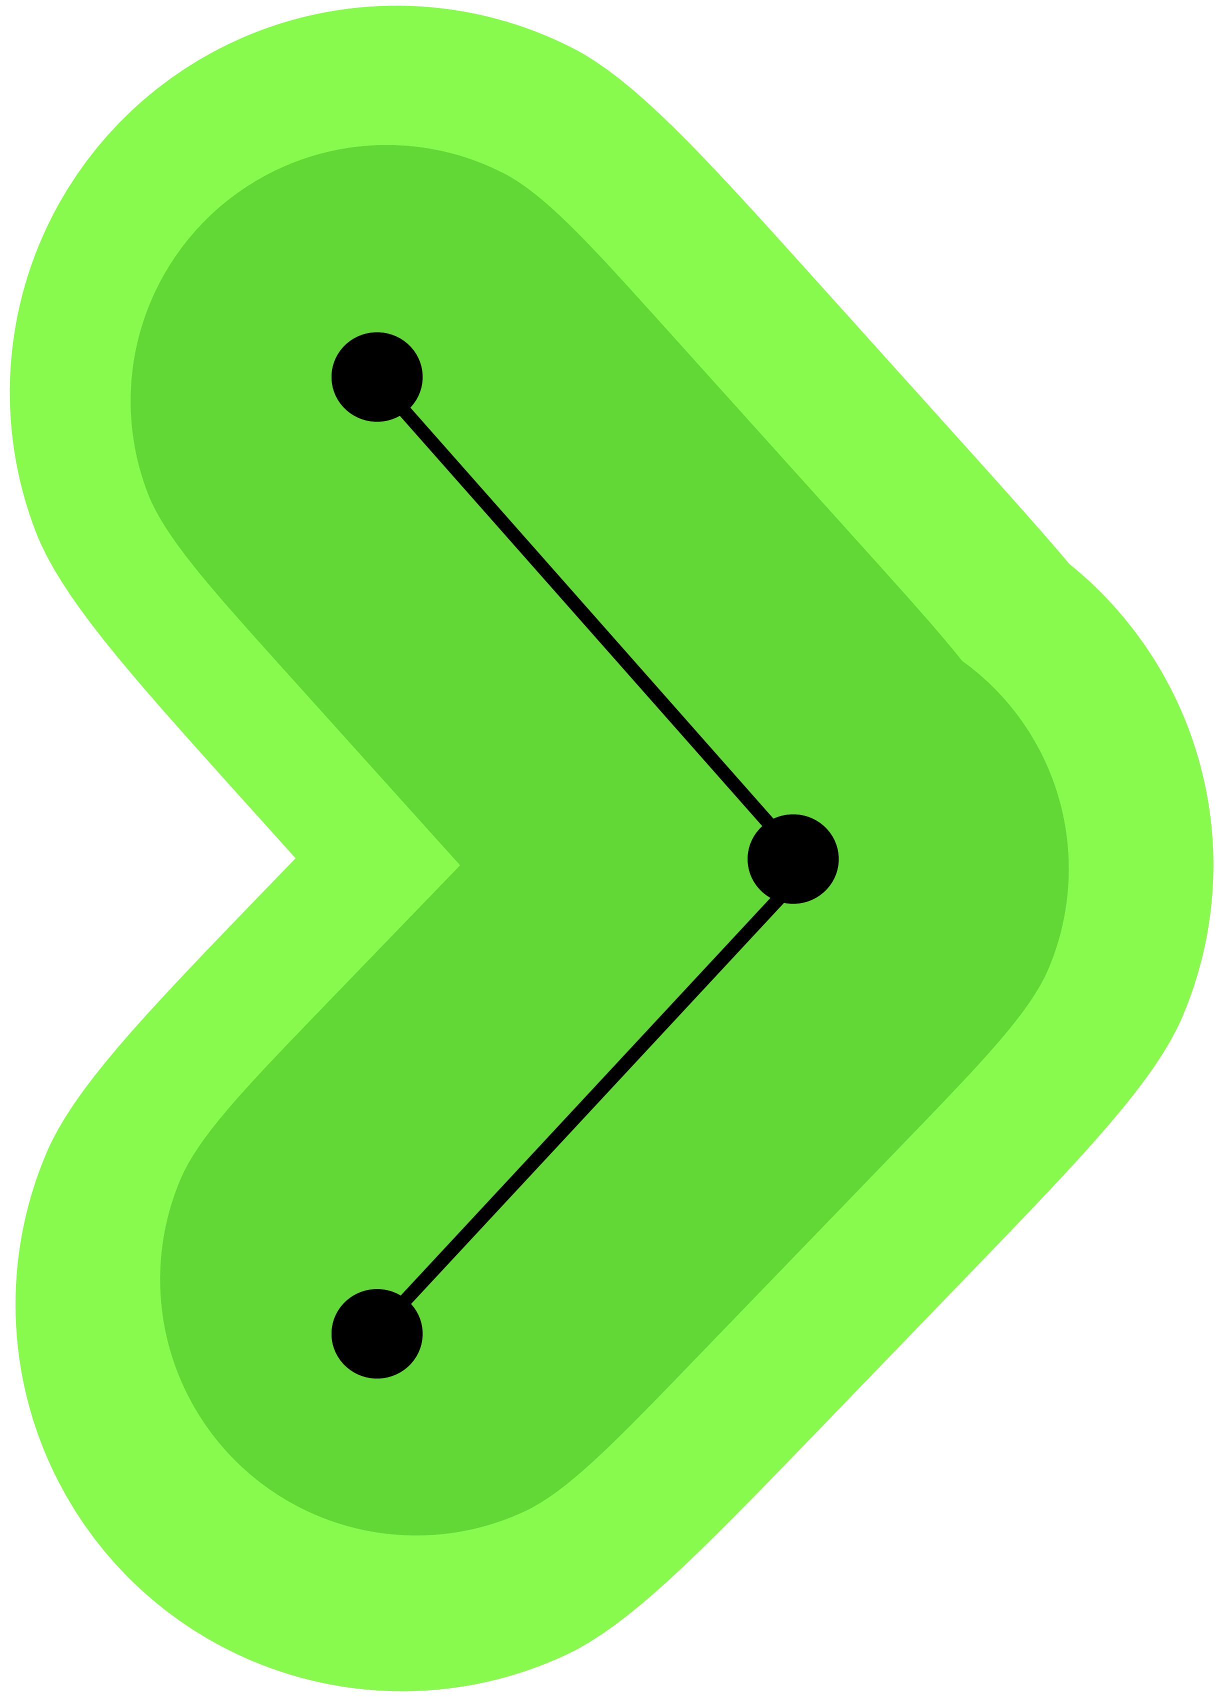

ϵ-CONTAMINATION

GLOBAL

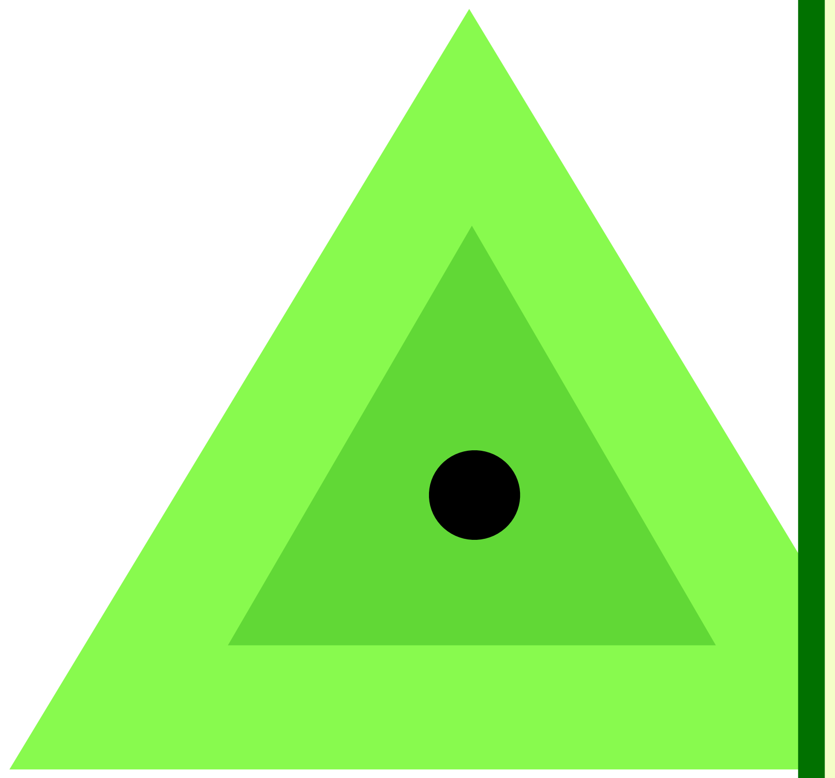

$$\epsilon_{\text{glob}} = \frac{\Delta}{1+\Delta}$$

$$\Delta = P_{\text{classif}}(x, \hat{y}) - P_{\text{classif}}(x, \hat{y}_2)$$

$$\hat{y}_2 = \arg \max_{y \in \mathscr{Y} \setminus \{\hat{y}\}} P_{\text{classif}}(y \mid x)$$

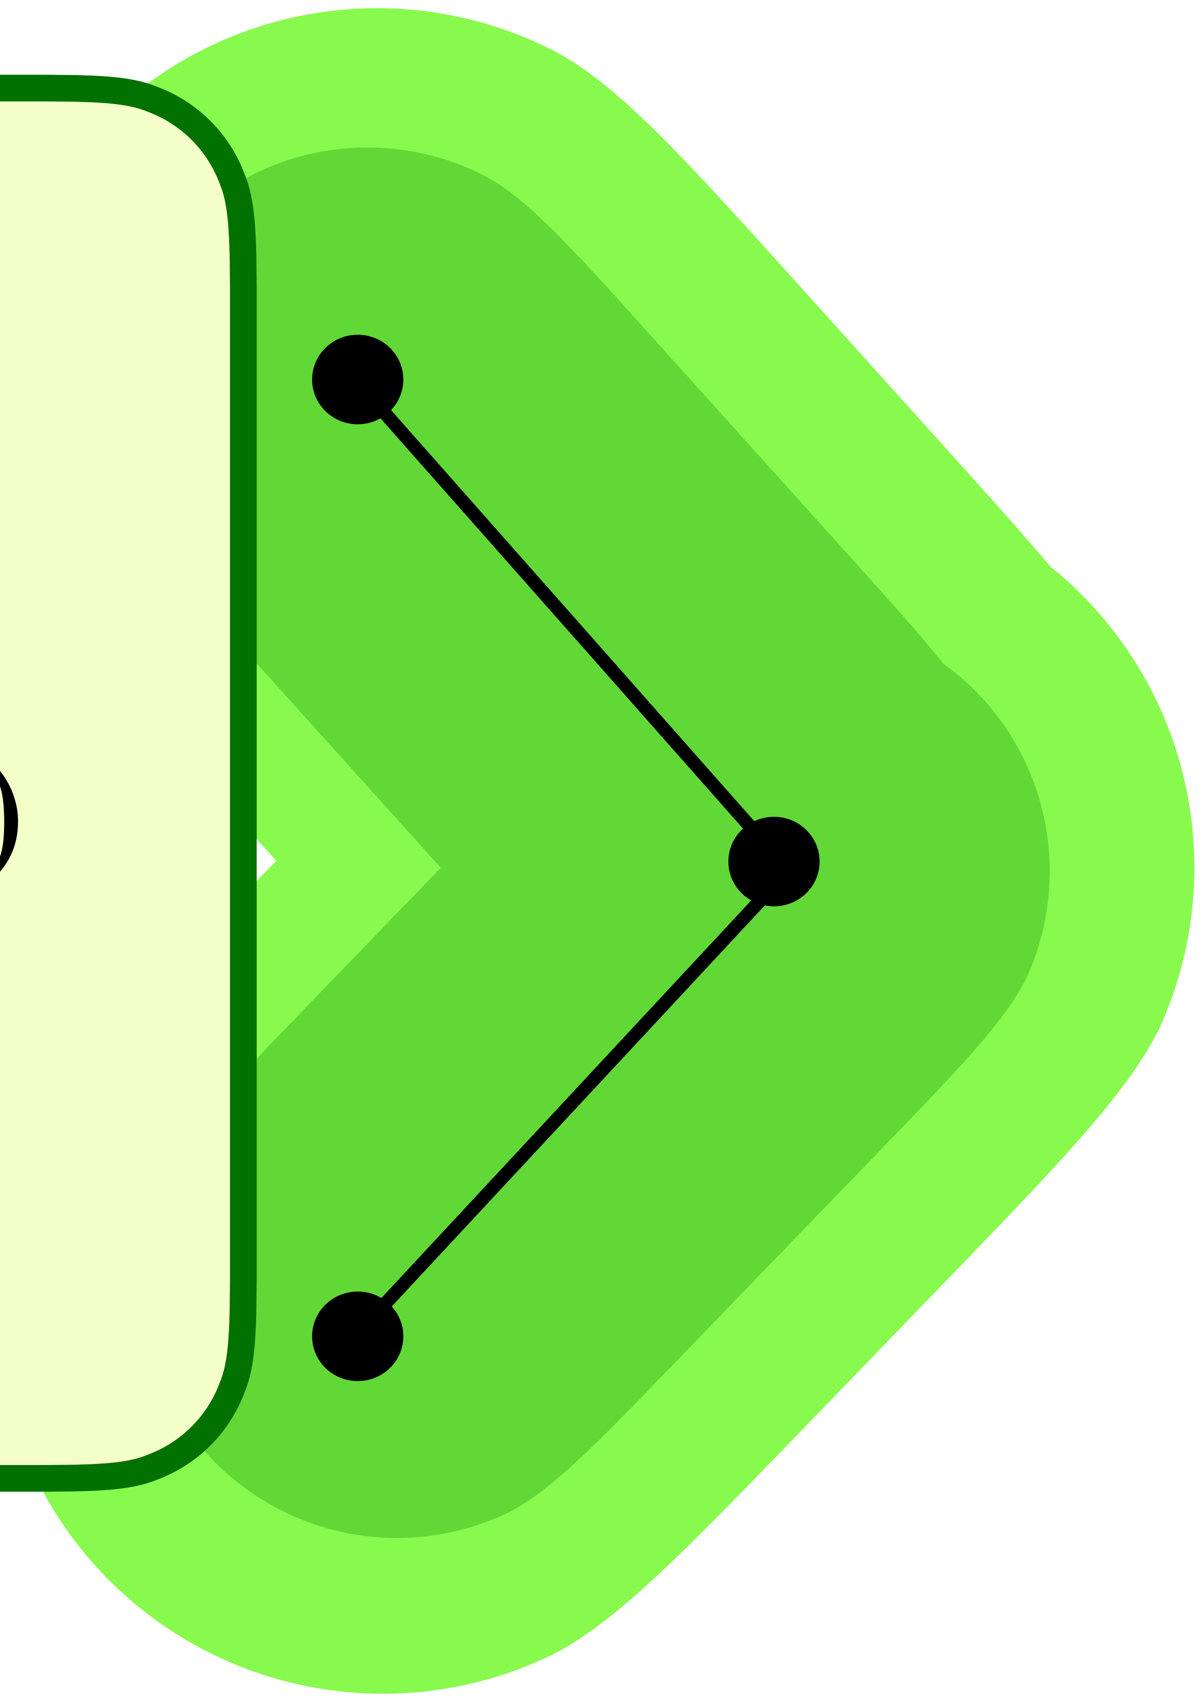

$$\mathscr{P}_\epsilon$$

$$\|$$

$$\left\{(1-\epsilon)P_{\text{classif}} + \epsilon P : P \in \Delta\right\}$$

ROBUSTNESS QUANTIFICATION

- correlates nicely with accuracy
- works for different types of model architectures
- is conceptually different from UQ
- also works with global perturbations ✓

NBC

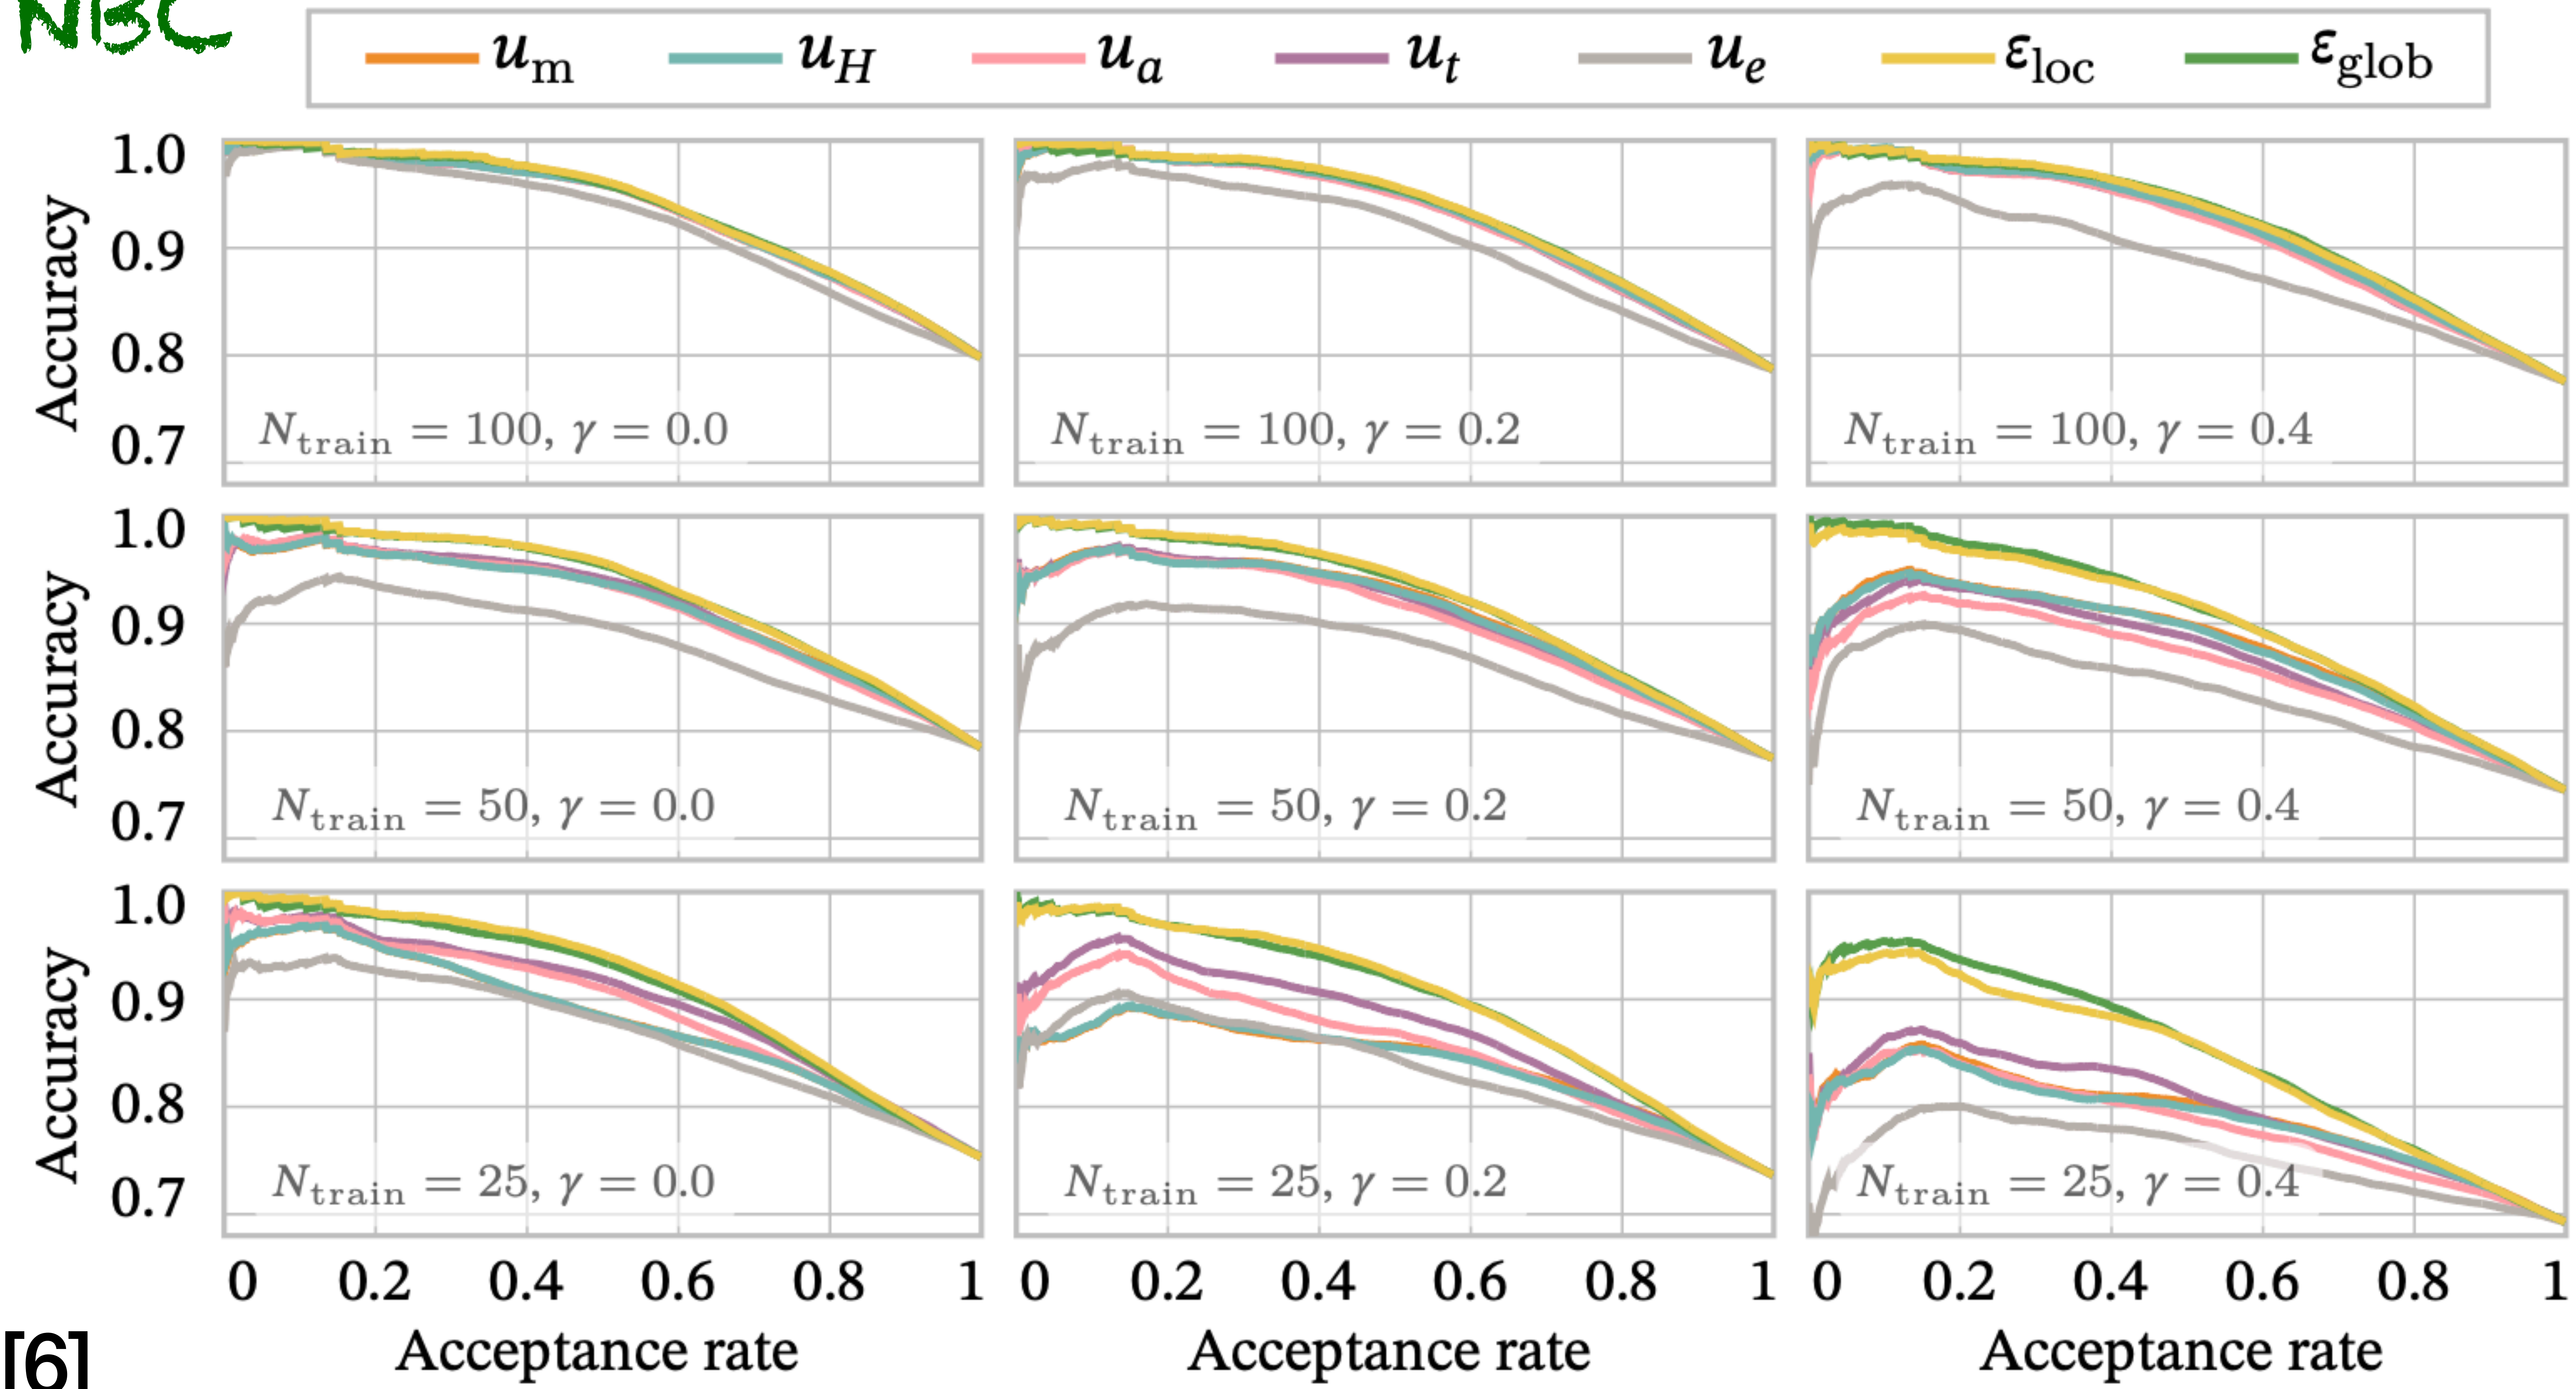

u_m
u_H
u_a
u_t
u_e
ε_{loc}
ε_{glob}
Accuracy
1.0
0.9
0.8
0.7
$N_{train} = 100, \gamma = 0.0$
$N_{train} = 100, \gamma = 0.2$
$N_{train} = 100, \gamma = 0.4$
$N_{train} = 50, \gamma = 0.0$
$N_{train} = 50, \gamma = 0.2$
$N_{train} = 50, \gamma = 0.4$
$N_{train} = 25, \gamma = 0.0$
$N_{train} = 25, \gamma = 0.2$
$N_{train} = 25, \gamma = 0.4$
0
0.2
0.4
0.6
0.8
1
Acceptance rate

[6]

ROBUSTNESS QUANTIFICATION

- correlates nicely with accuracy ✓
- works for different types of model architectures
- is conceptually different from UQ
- also works with global perturbations ✓
- is competitive with UQ ✓
- is good with distribution shift and small data sets ✓

NBC

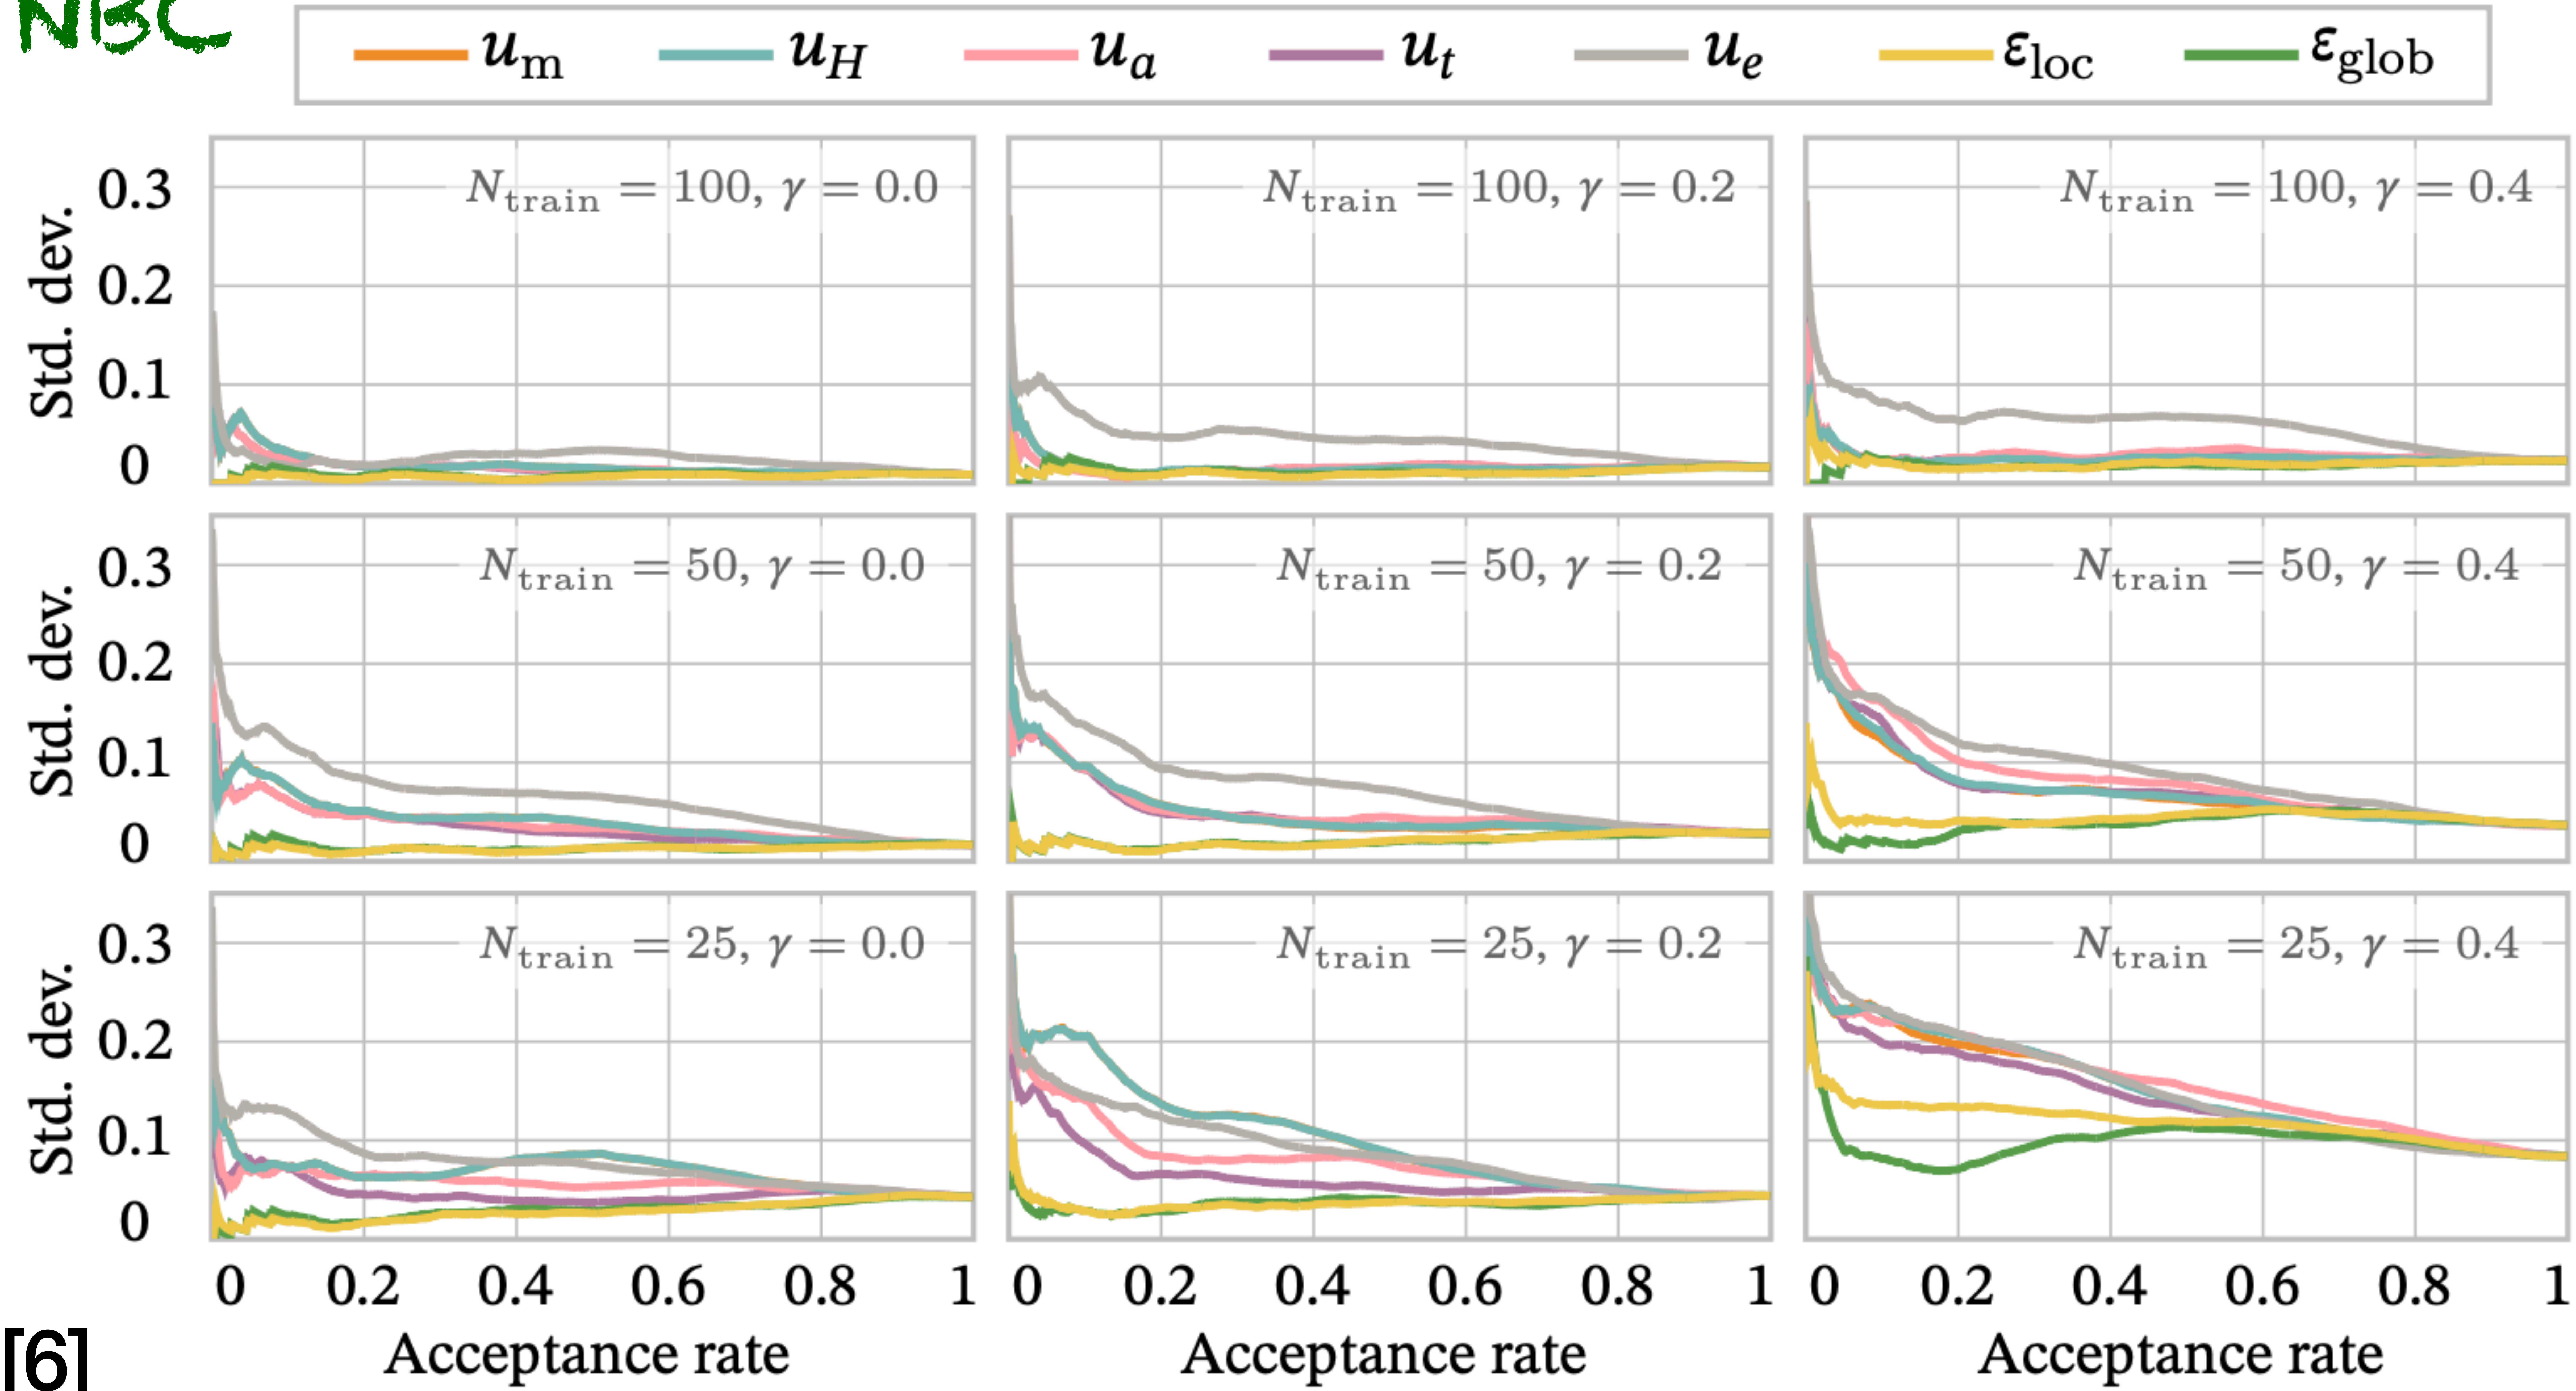

u_m
u_H
u_a
u_t
u_e
ε_{loc}
ε_{glob}
$N_{train} = 100,\ \gamma = 0.0$
$N_{train} = 100,\ \gamma = 0.2$
$N_{train} = 100,\ \gamma = 0.4$
$N_{train} = 50,\ \gamma = 0.0$
$N_{train} = 50,\ \gamma = 0.2$
$N_{train} = 50,\ \gamma = 0.4$
$N_{train} = 25,\ \gamma = 0.0$
$N_{train} = 25,\ \gamma = 0.2$
$N_{train} = 25,\ \gamma = 0.4$
Std. dev.
0.3
0.2
0.1
0
0
0.2
0.4
0.6
0.8
1
Acceptance rate

[6]

ROBUSTNESS QUANTIFICATION

- correlates nicely with accuracy
- works for different types of model architectures
- is conceptually different from UQ
- also works with global perturbations
- is competitive with UQ
- is good with distribution shift and small data sets
- is more stable than UQ ✓

NBC

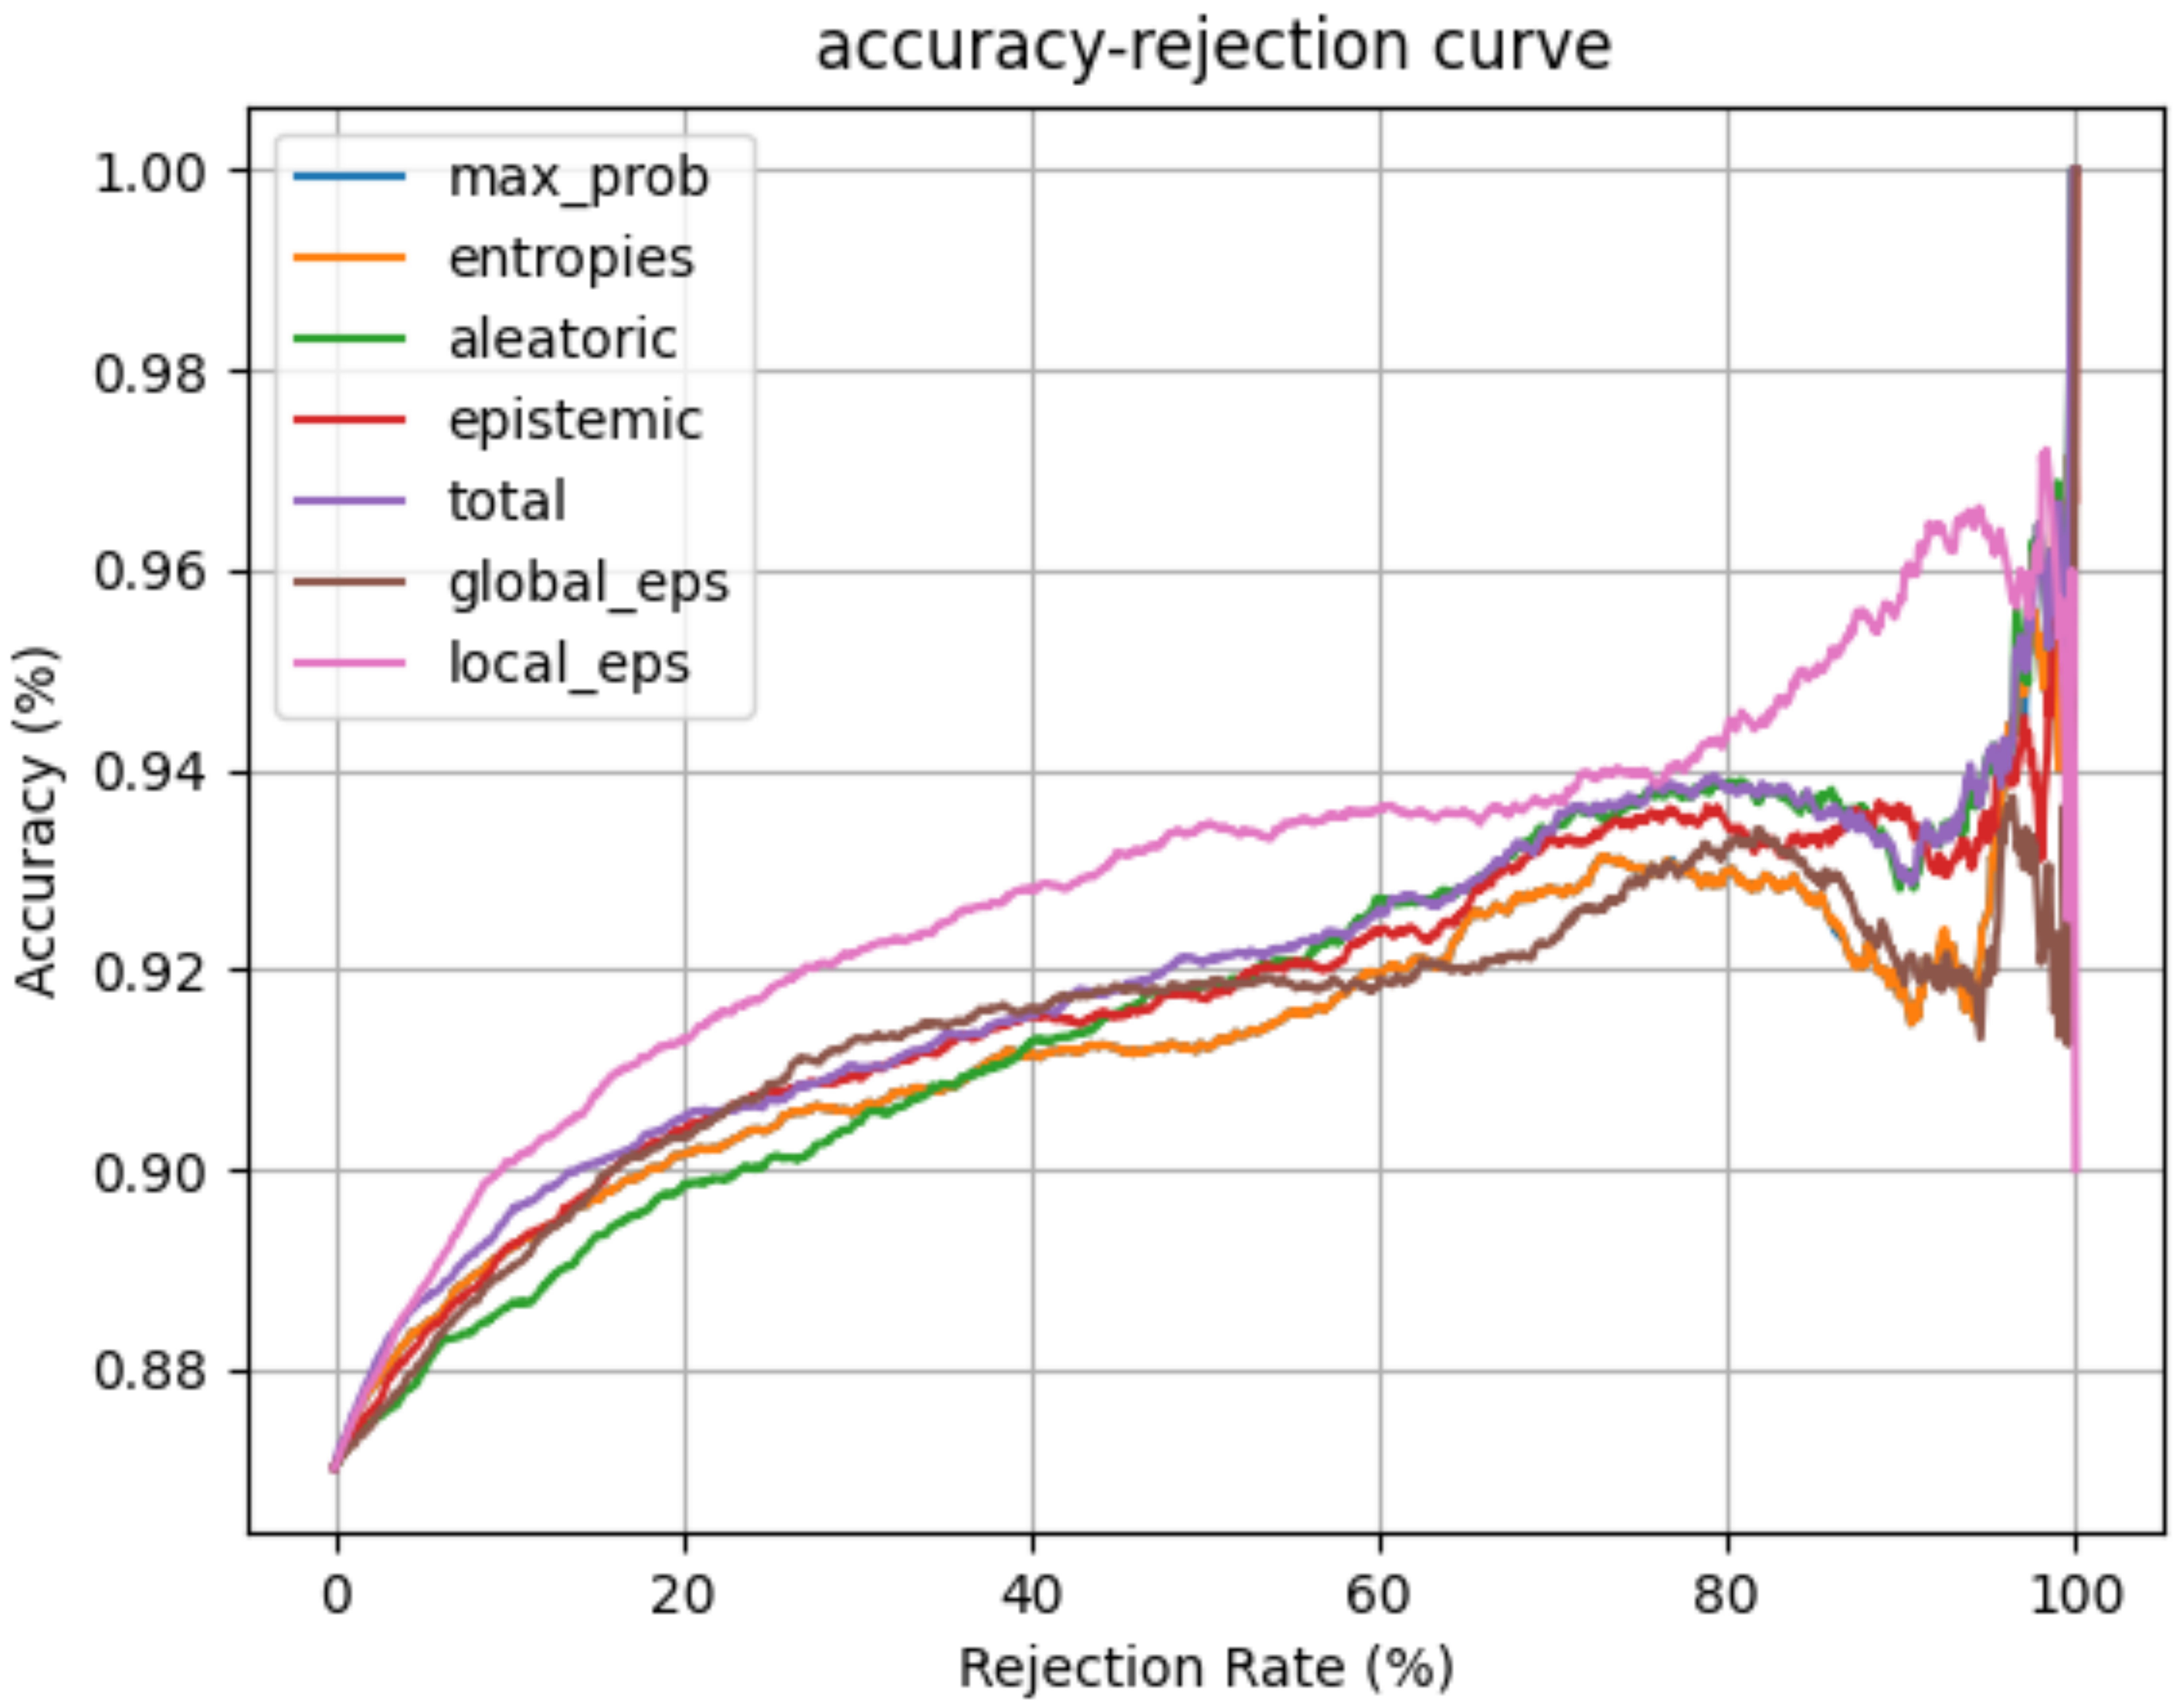

accuracy-rejection curve
max_prob
entropies
aleatoric
epistemic
total
global_eps
local_eps
1.00
0.98
0.96
0.94
0.92
0.90
0.88
Accuracy (%)
0
20
40
60
80
100
Rejection Rate (%)

NBC

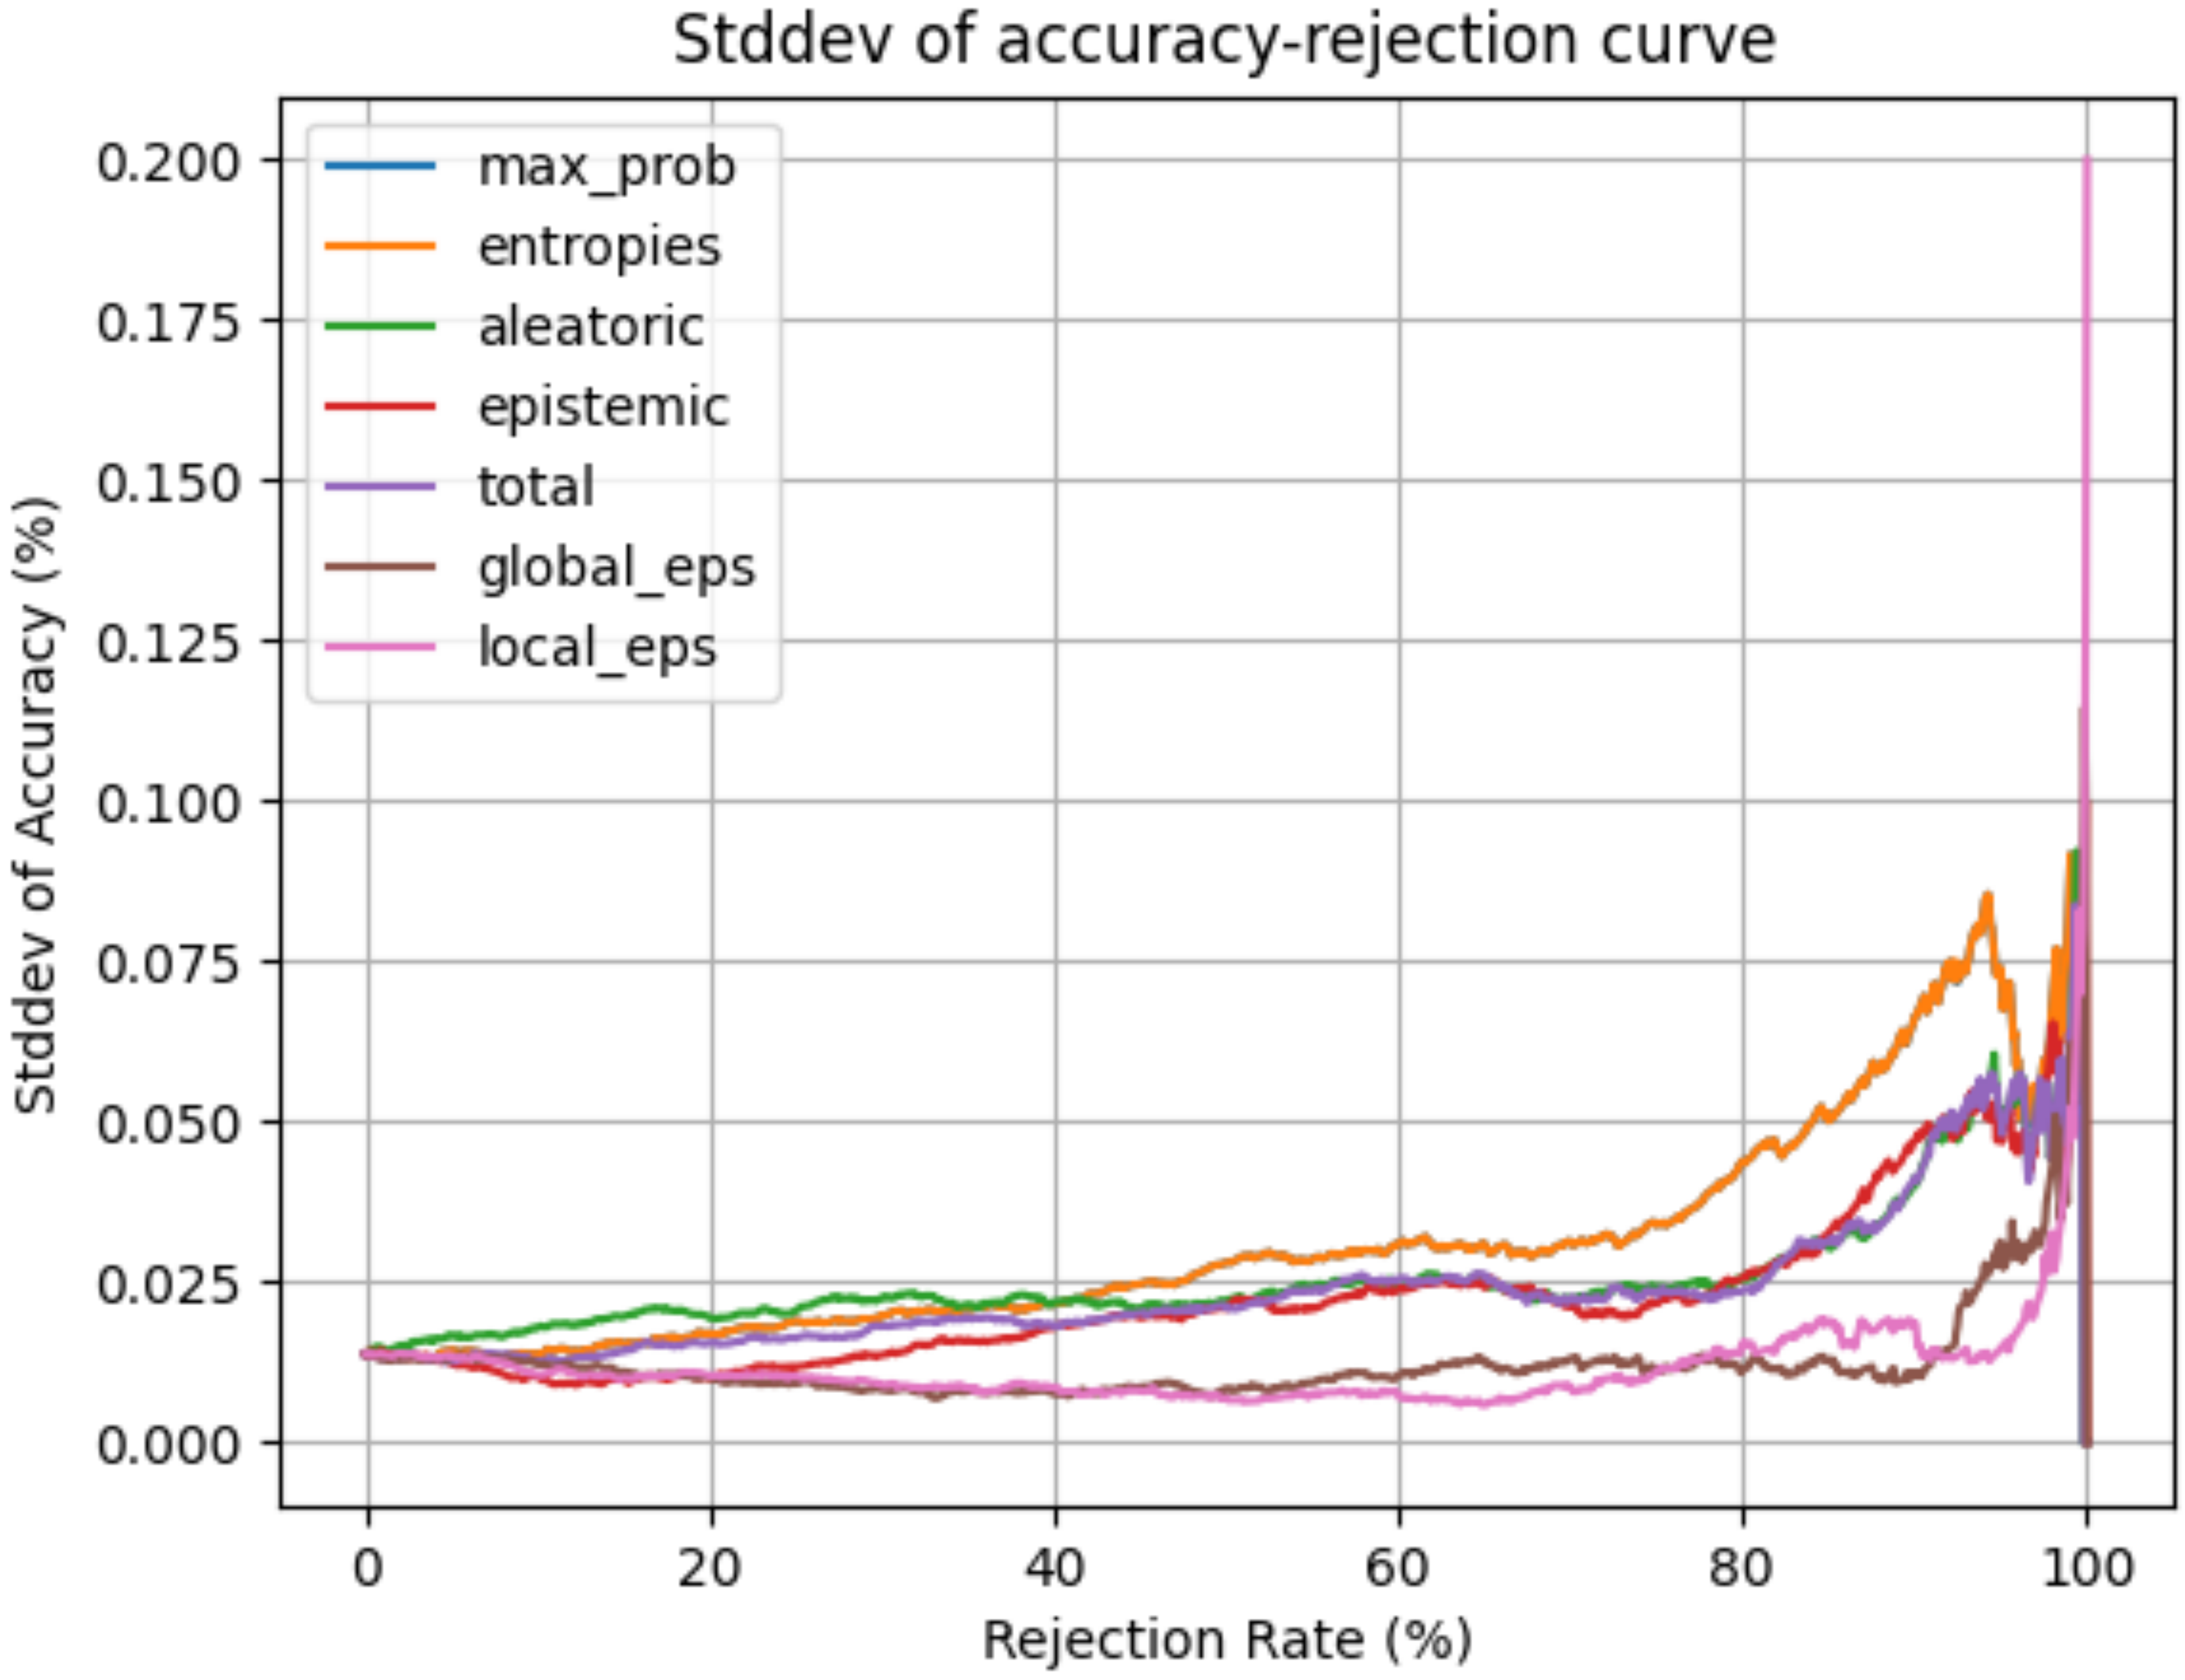

Stddev of accuracy-rejection curve
max_prob
entropies
aleatoric
epistemic
total
global_eps
local_eps
0.200
0.175
0.150
0.125
0.100
0.075
0.050
0.025
0.000
Stddev of Accuracy (%)
0
20
40
60
80
100
Rejection Rate (%)

ROBUSTNESS QUANTIFICATION

- correlates nicely with accuracy ✓
- works for different types of model architectures
- is conceptually different from UQ
- also works with global perturbations ✓
- is competitive with UQ ✓
- is good with distribution shift and small data sets
- is more stable than UQ ✓

NBC

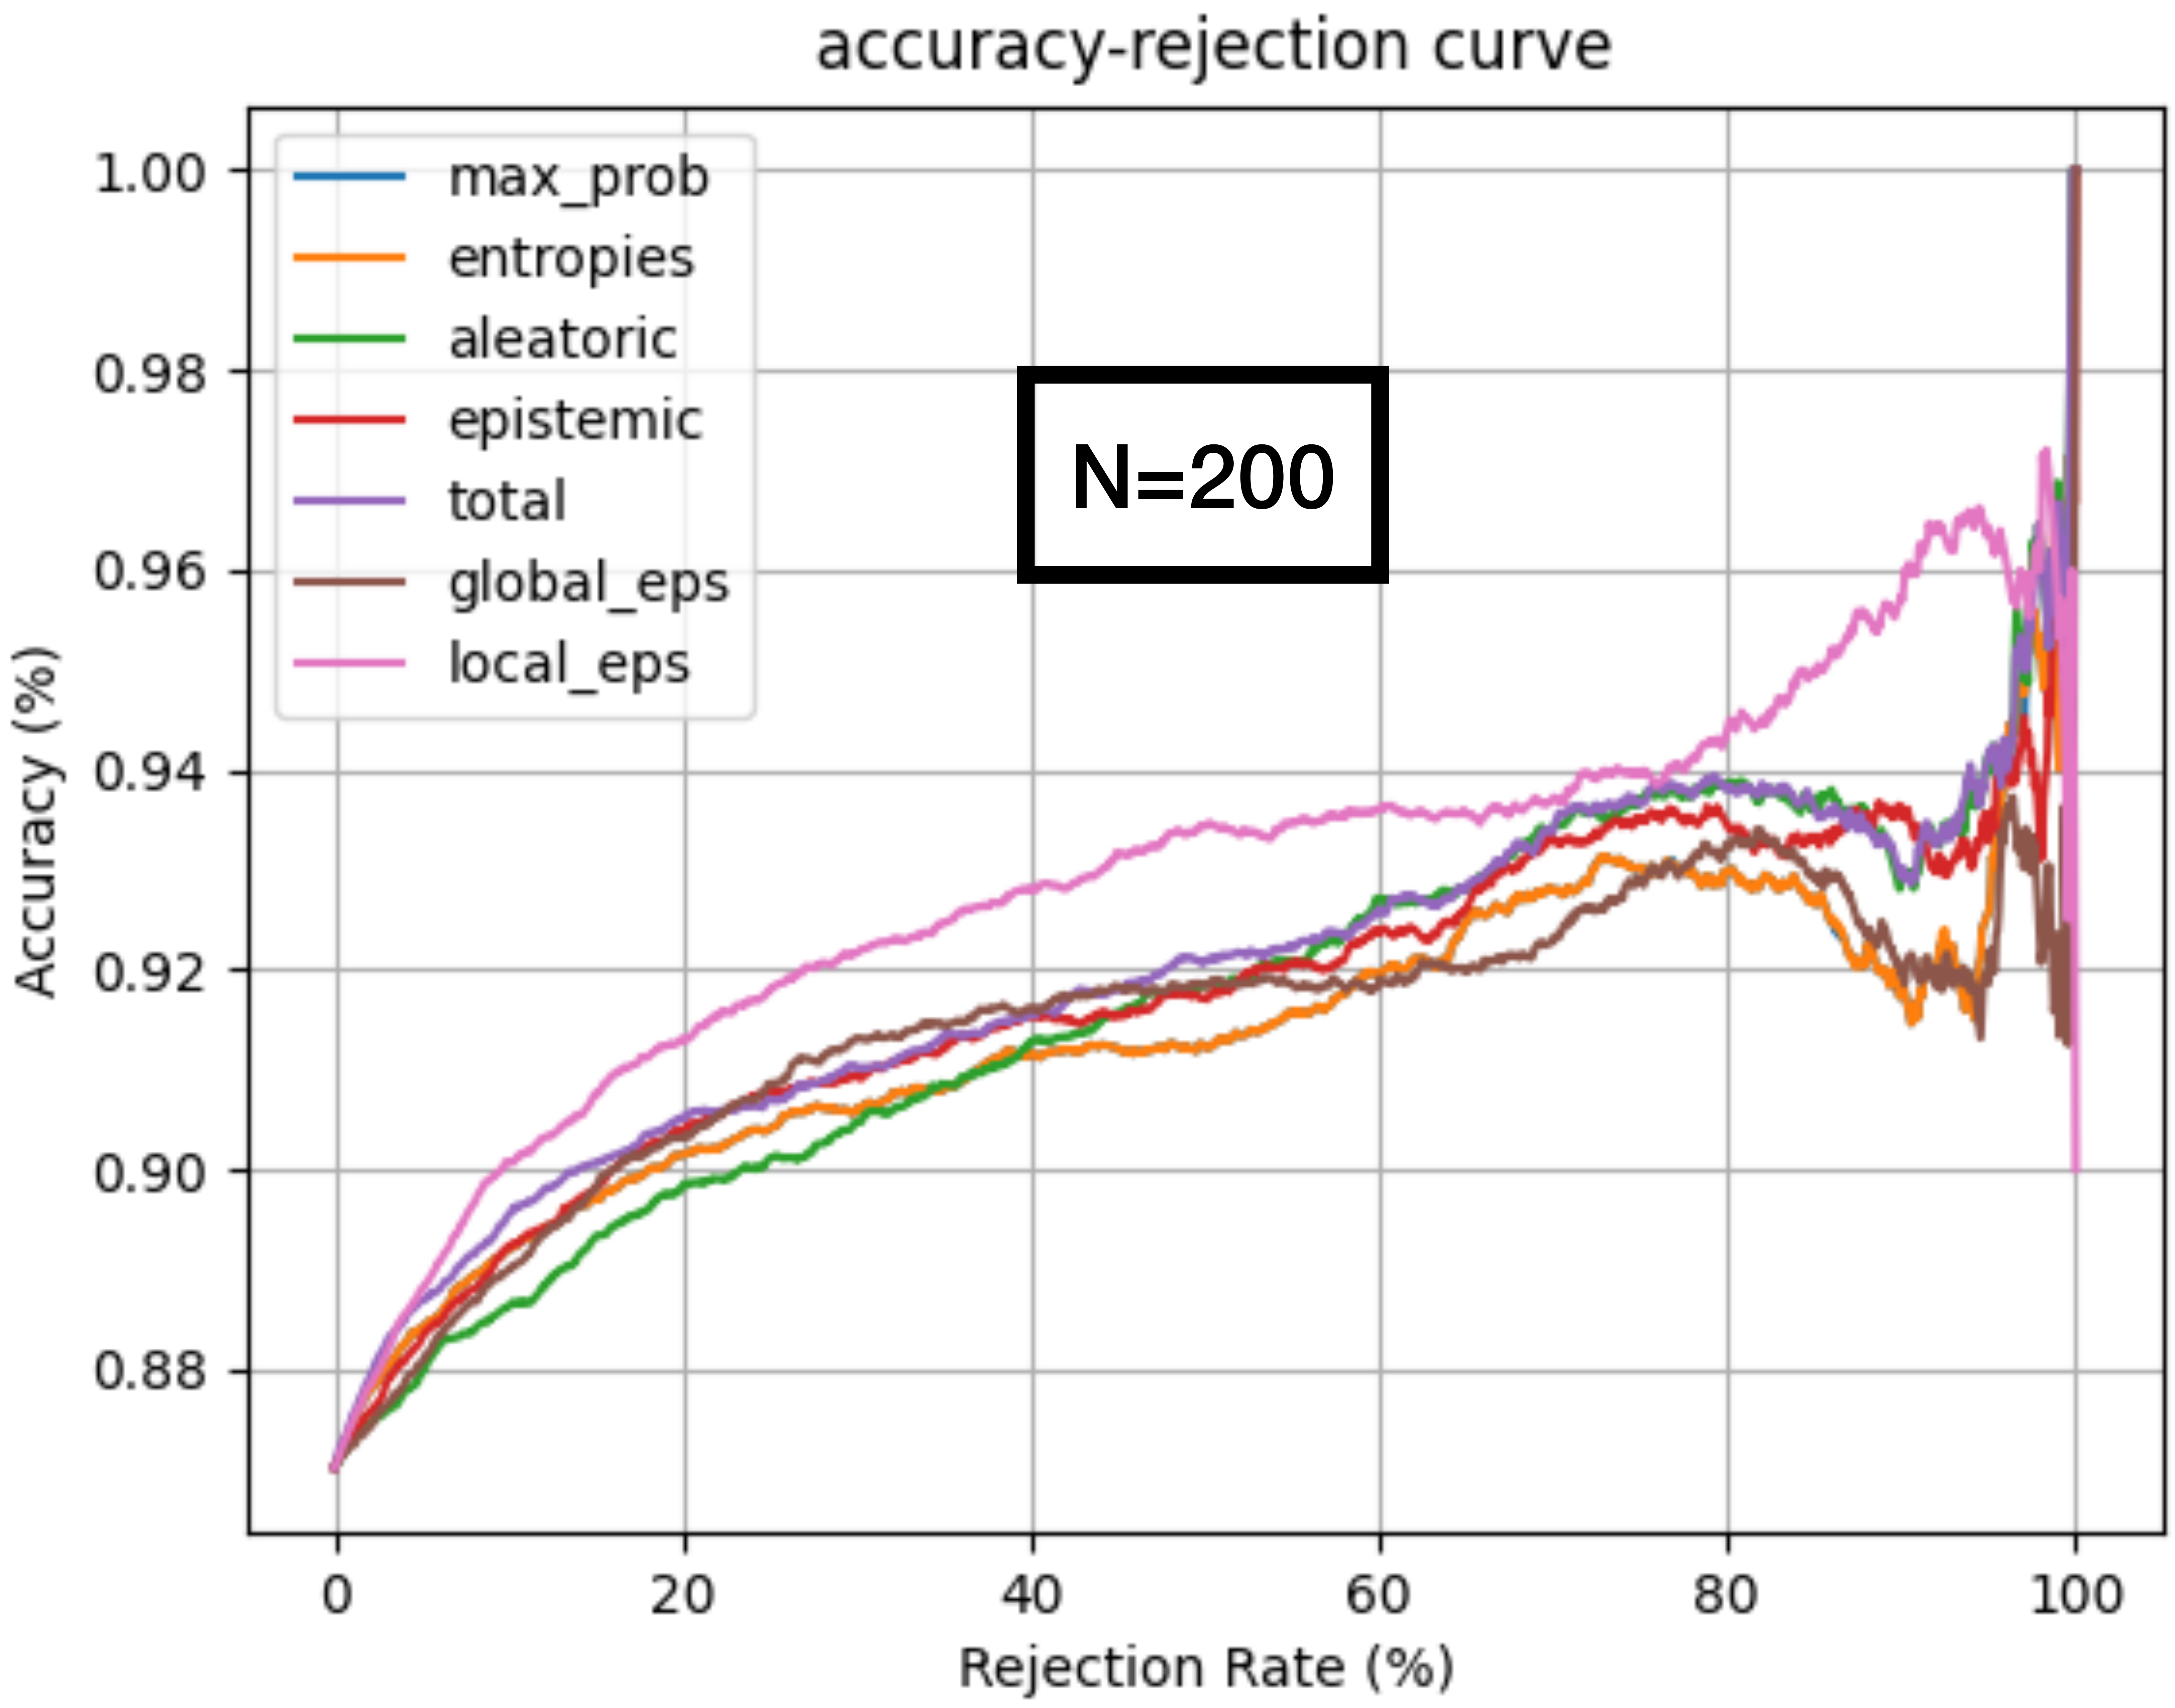
accuracy-rejection curve
max_prob
entropies
aleatoric
epistemic
total
global_eps
local_eps
N=200
Accuracy (%)
Rejection Rate (%)
1.00
0.98
0.96
0.94
0.92
0.90
0.88
0
20
40
60
80
100

NBC

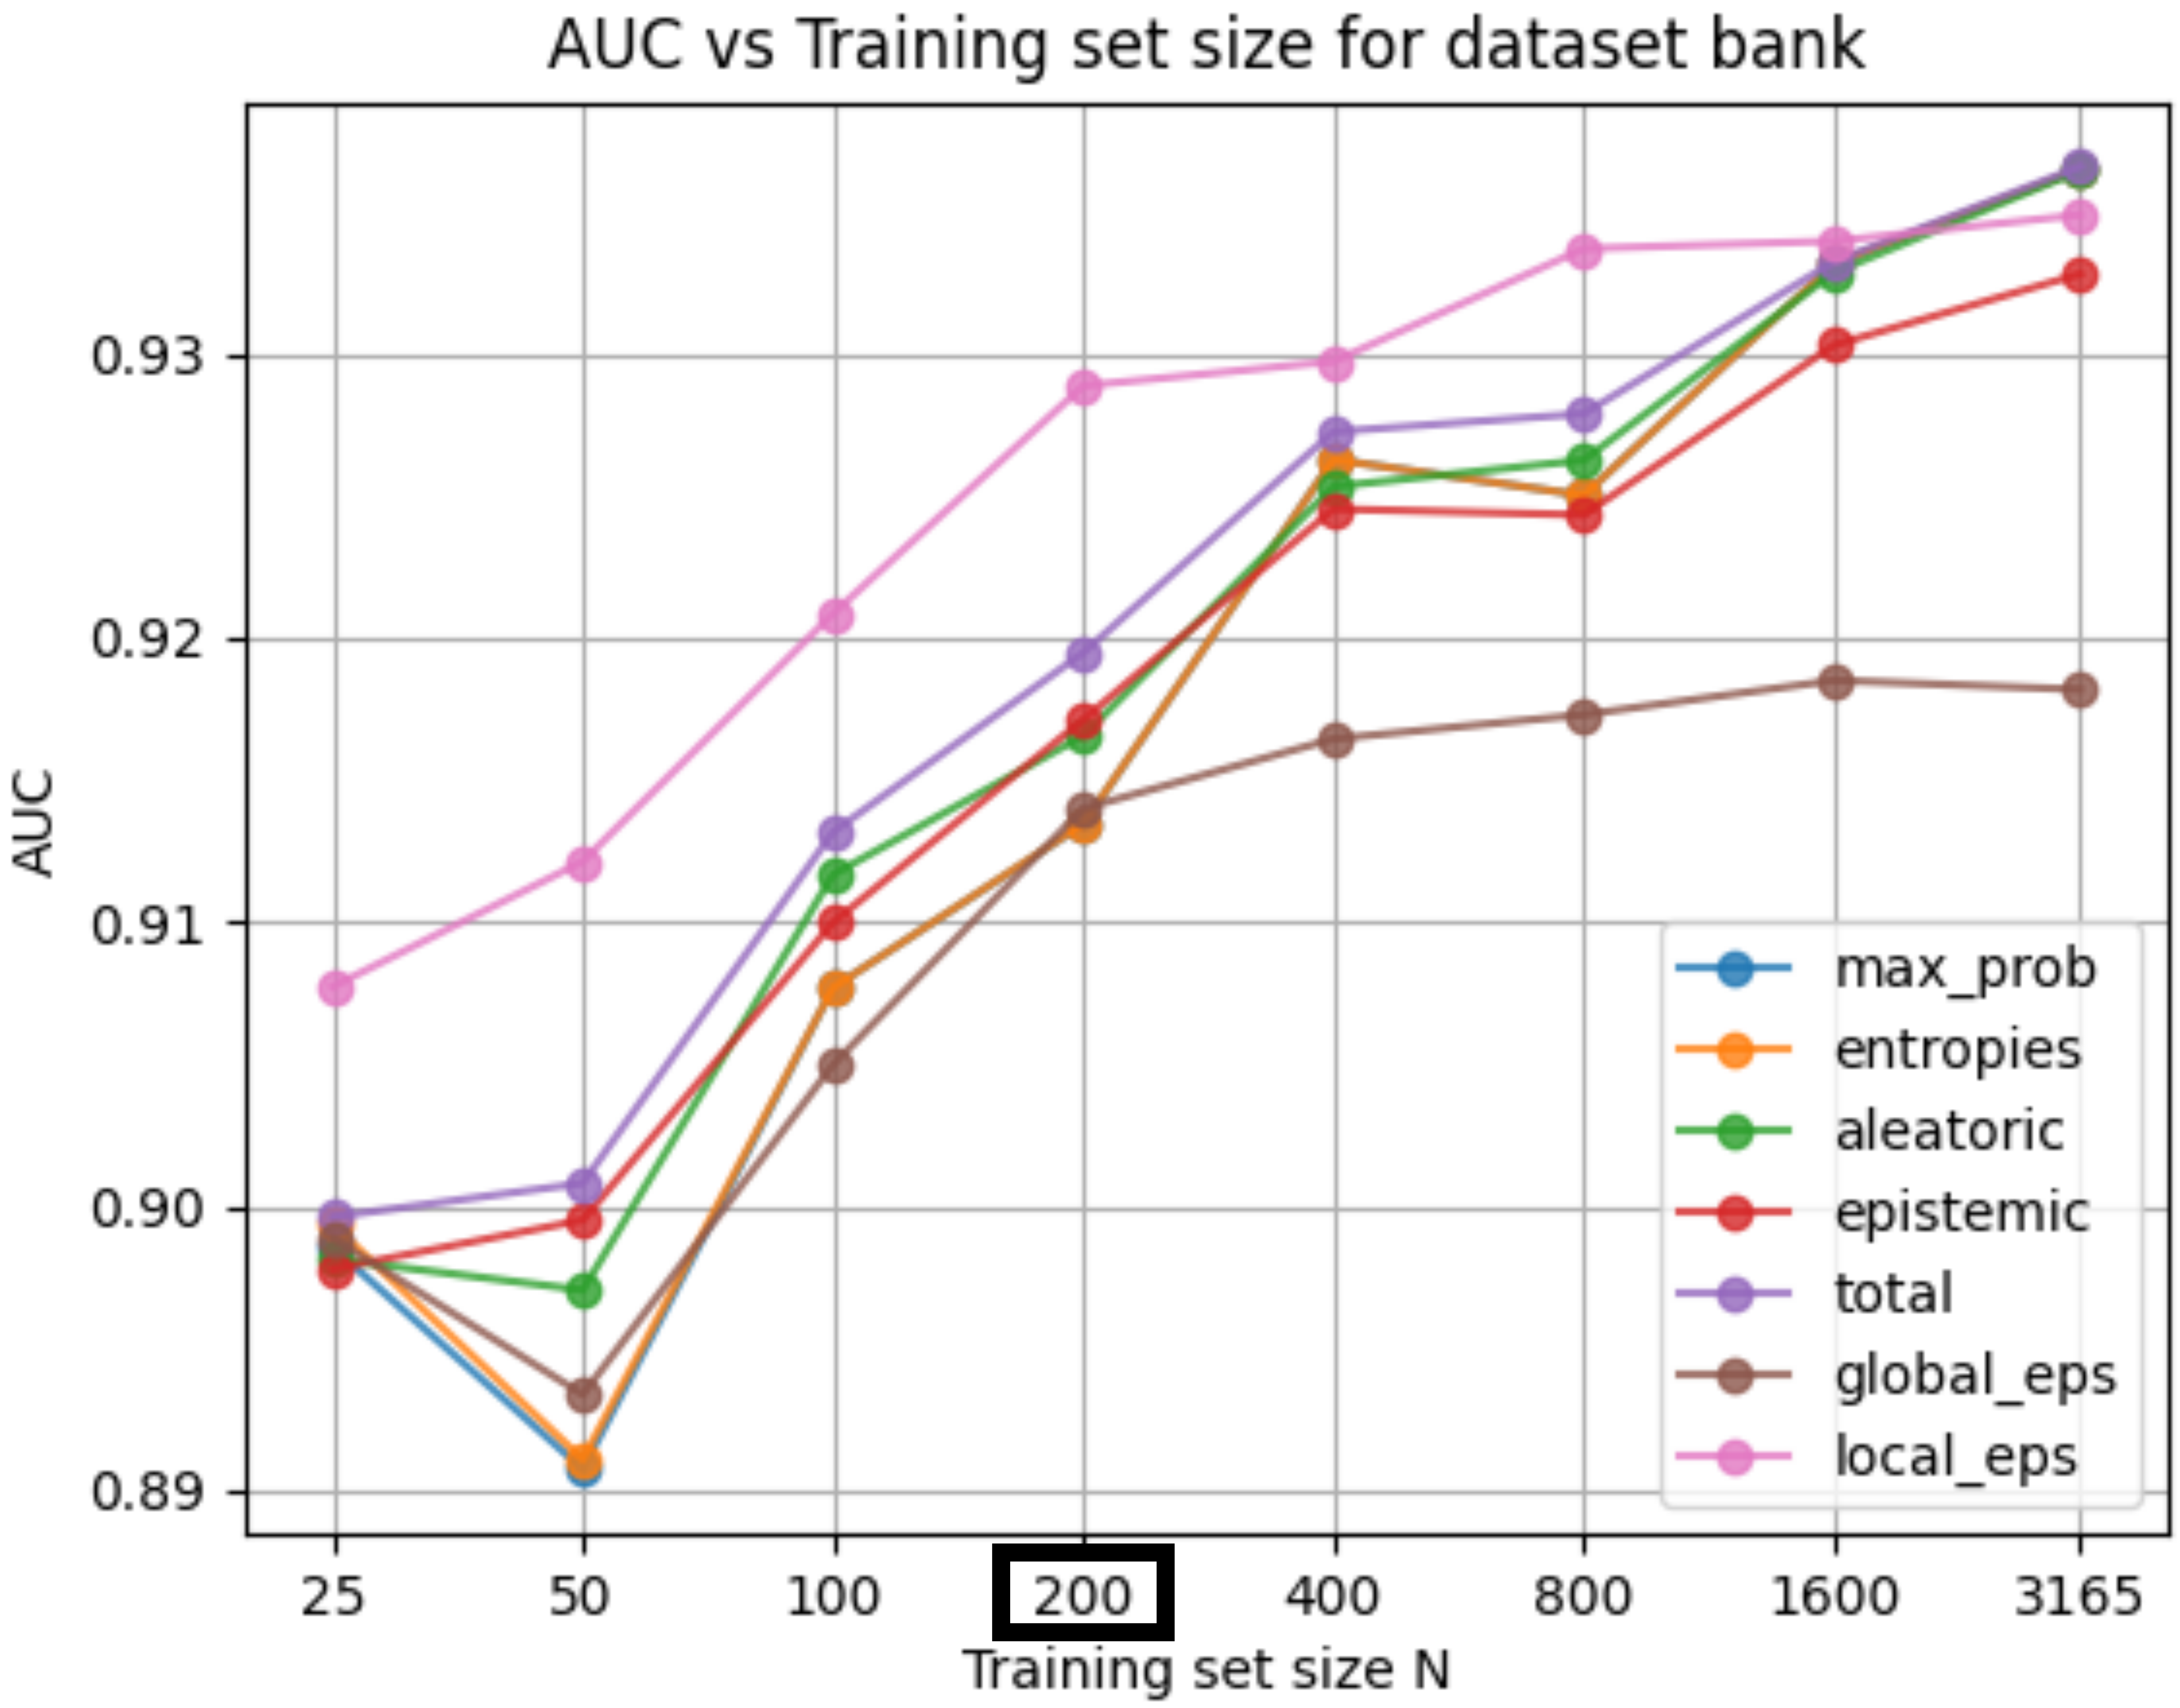

AUC vs Training set size for dataset bank
0.93
0.92
0.91
0.90
0.89
AUC
25
50
100
200
400
800
1600
3165
Training set size N
max_prob
entropies
aleatoric
epistemic
total
global_eps
local_eps

NBC

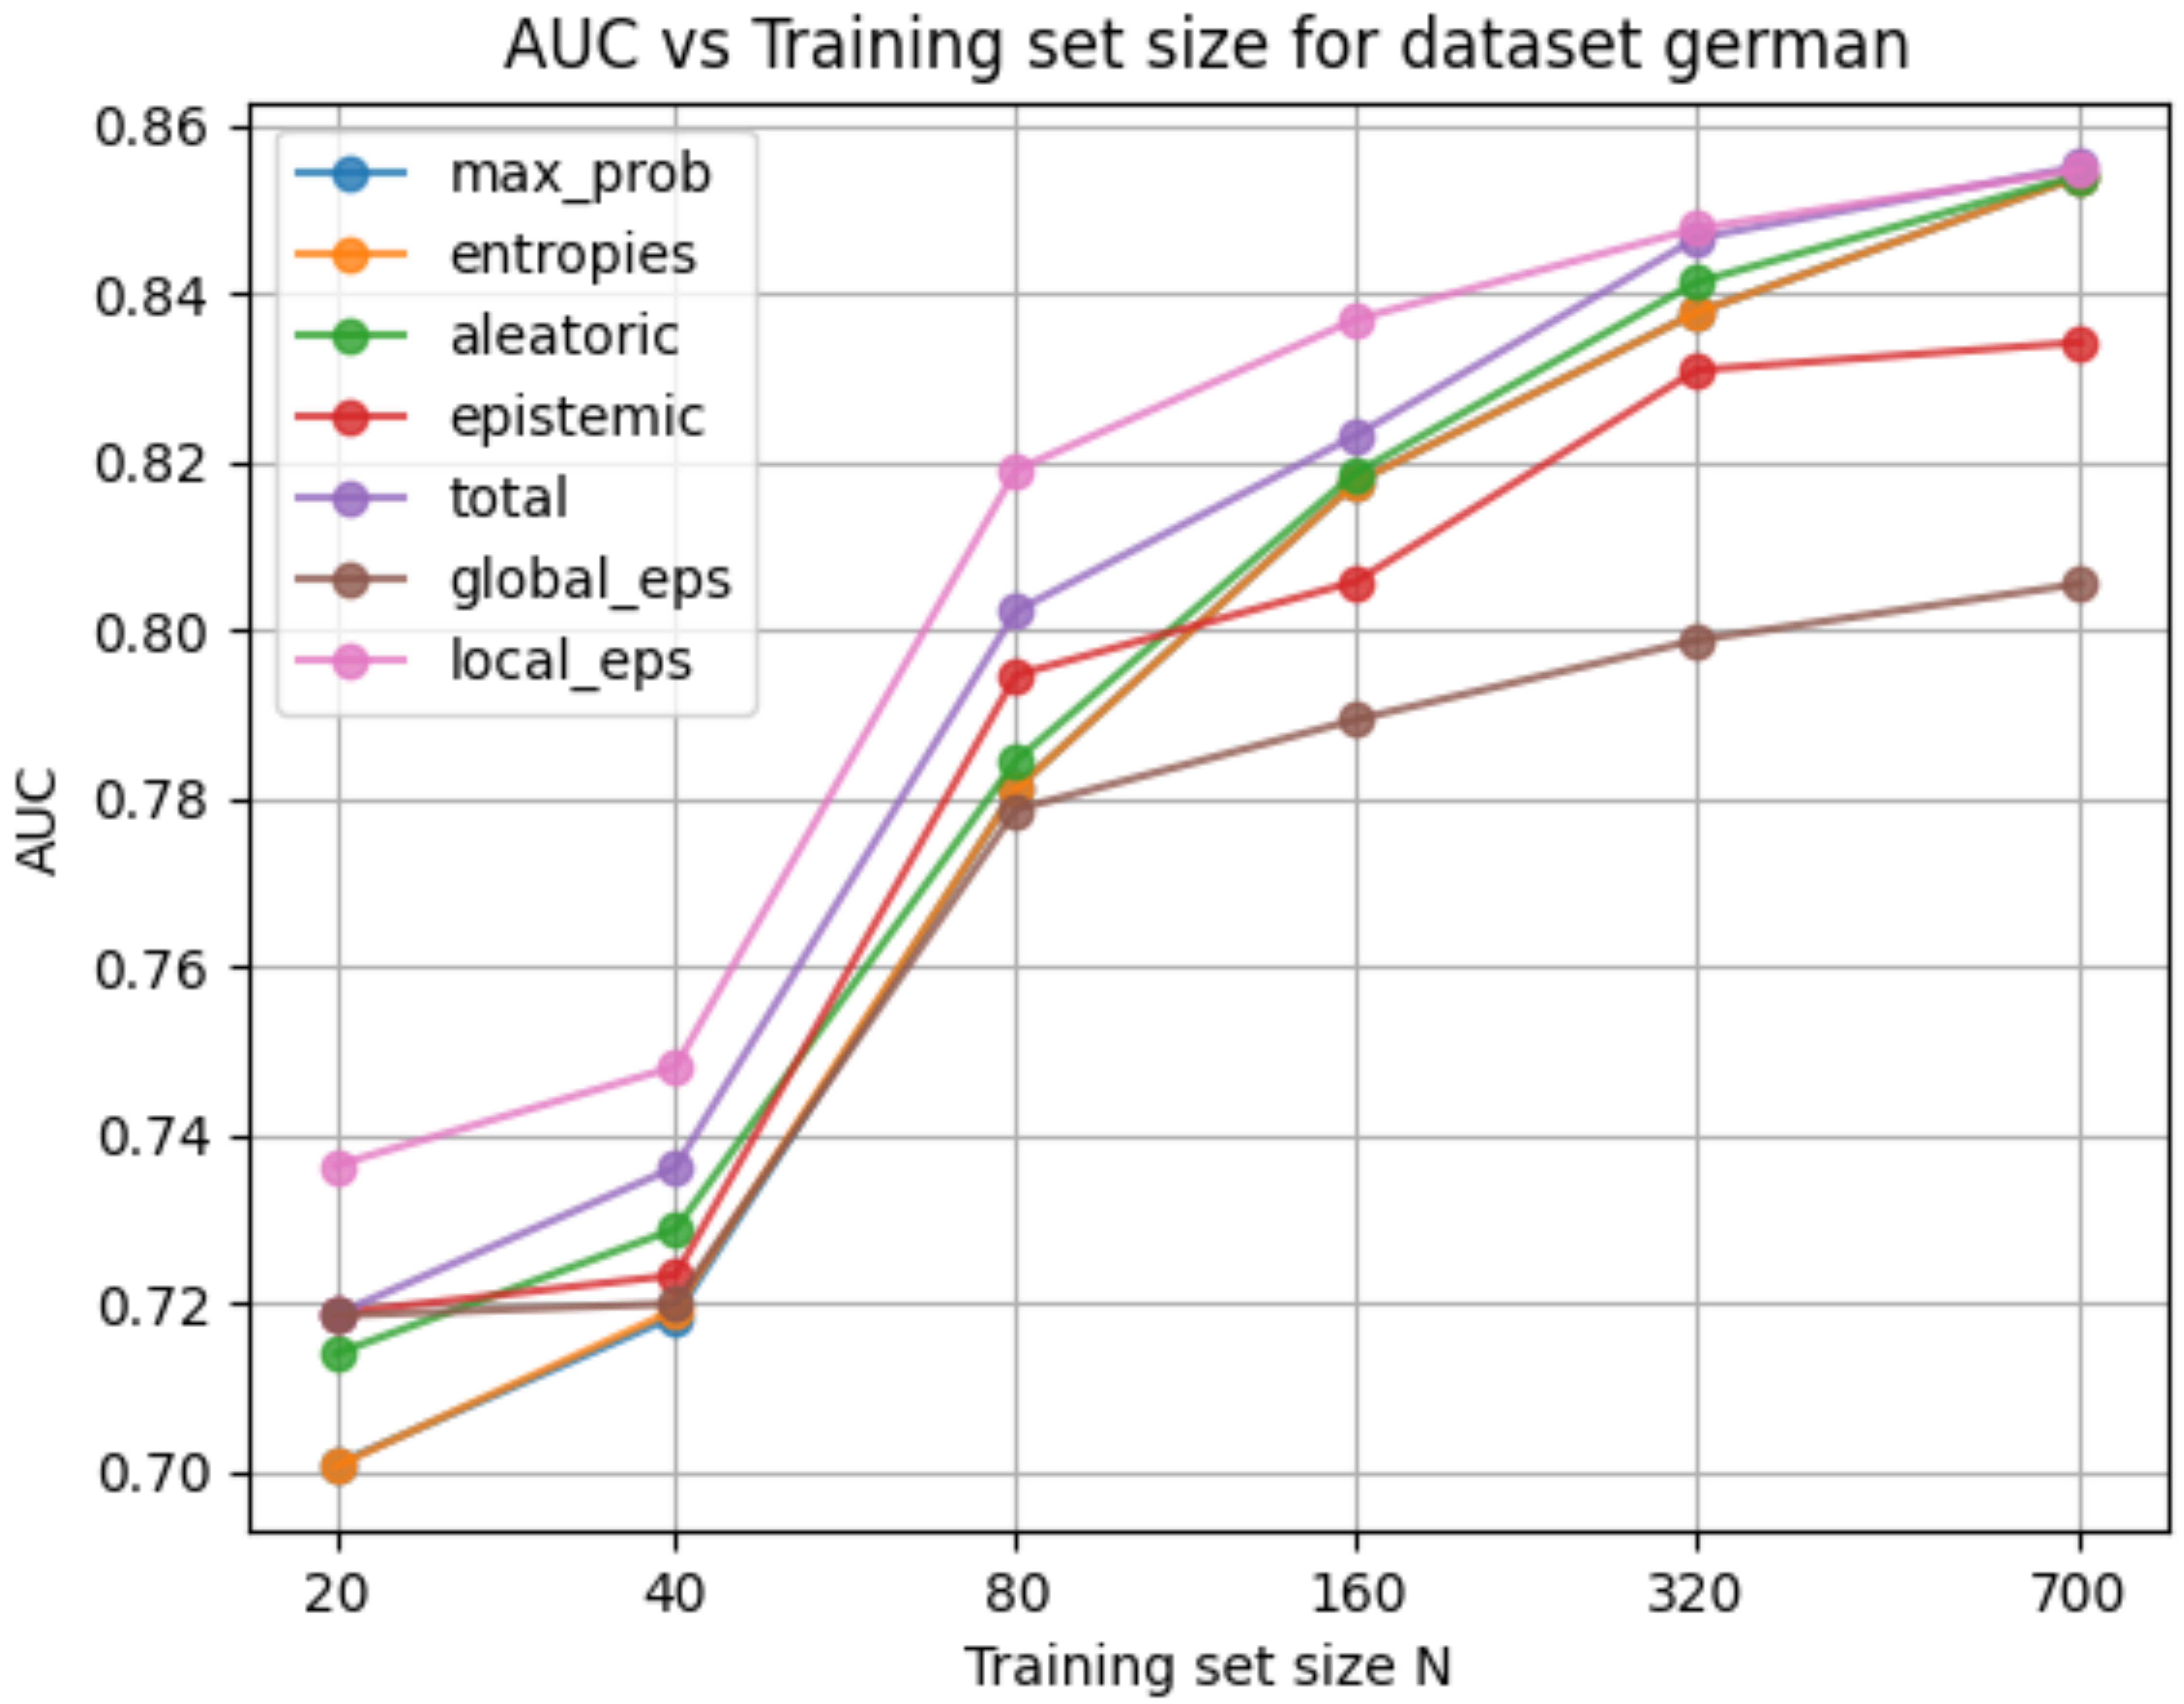

AUC vs Training set size for dataset german
max_prob
entropies
aleatoric
epistemic
total
global_eps
local_eps
AUC
Training set size N
0.70
0.72
0.74
0.76
0.78
0.80
0.82
0.84
0.86
20
40
80
160
320
700

NBC

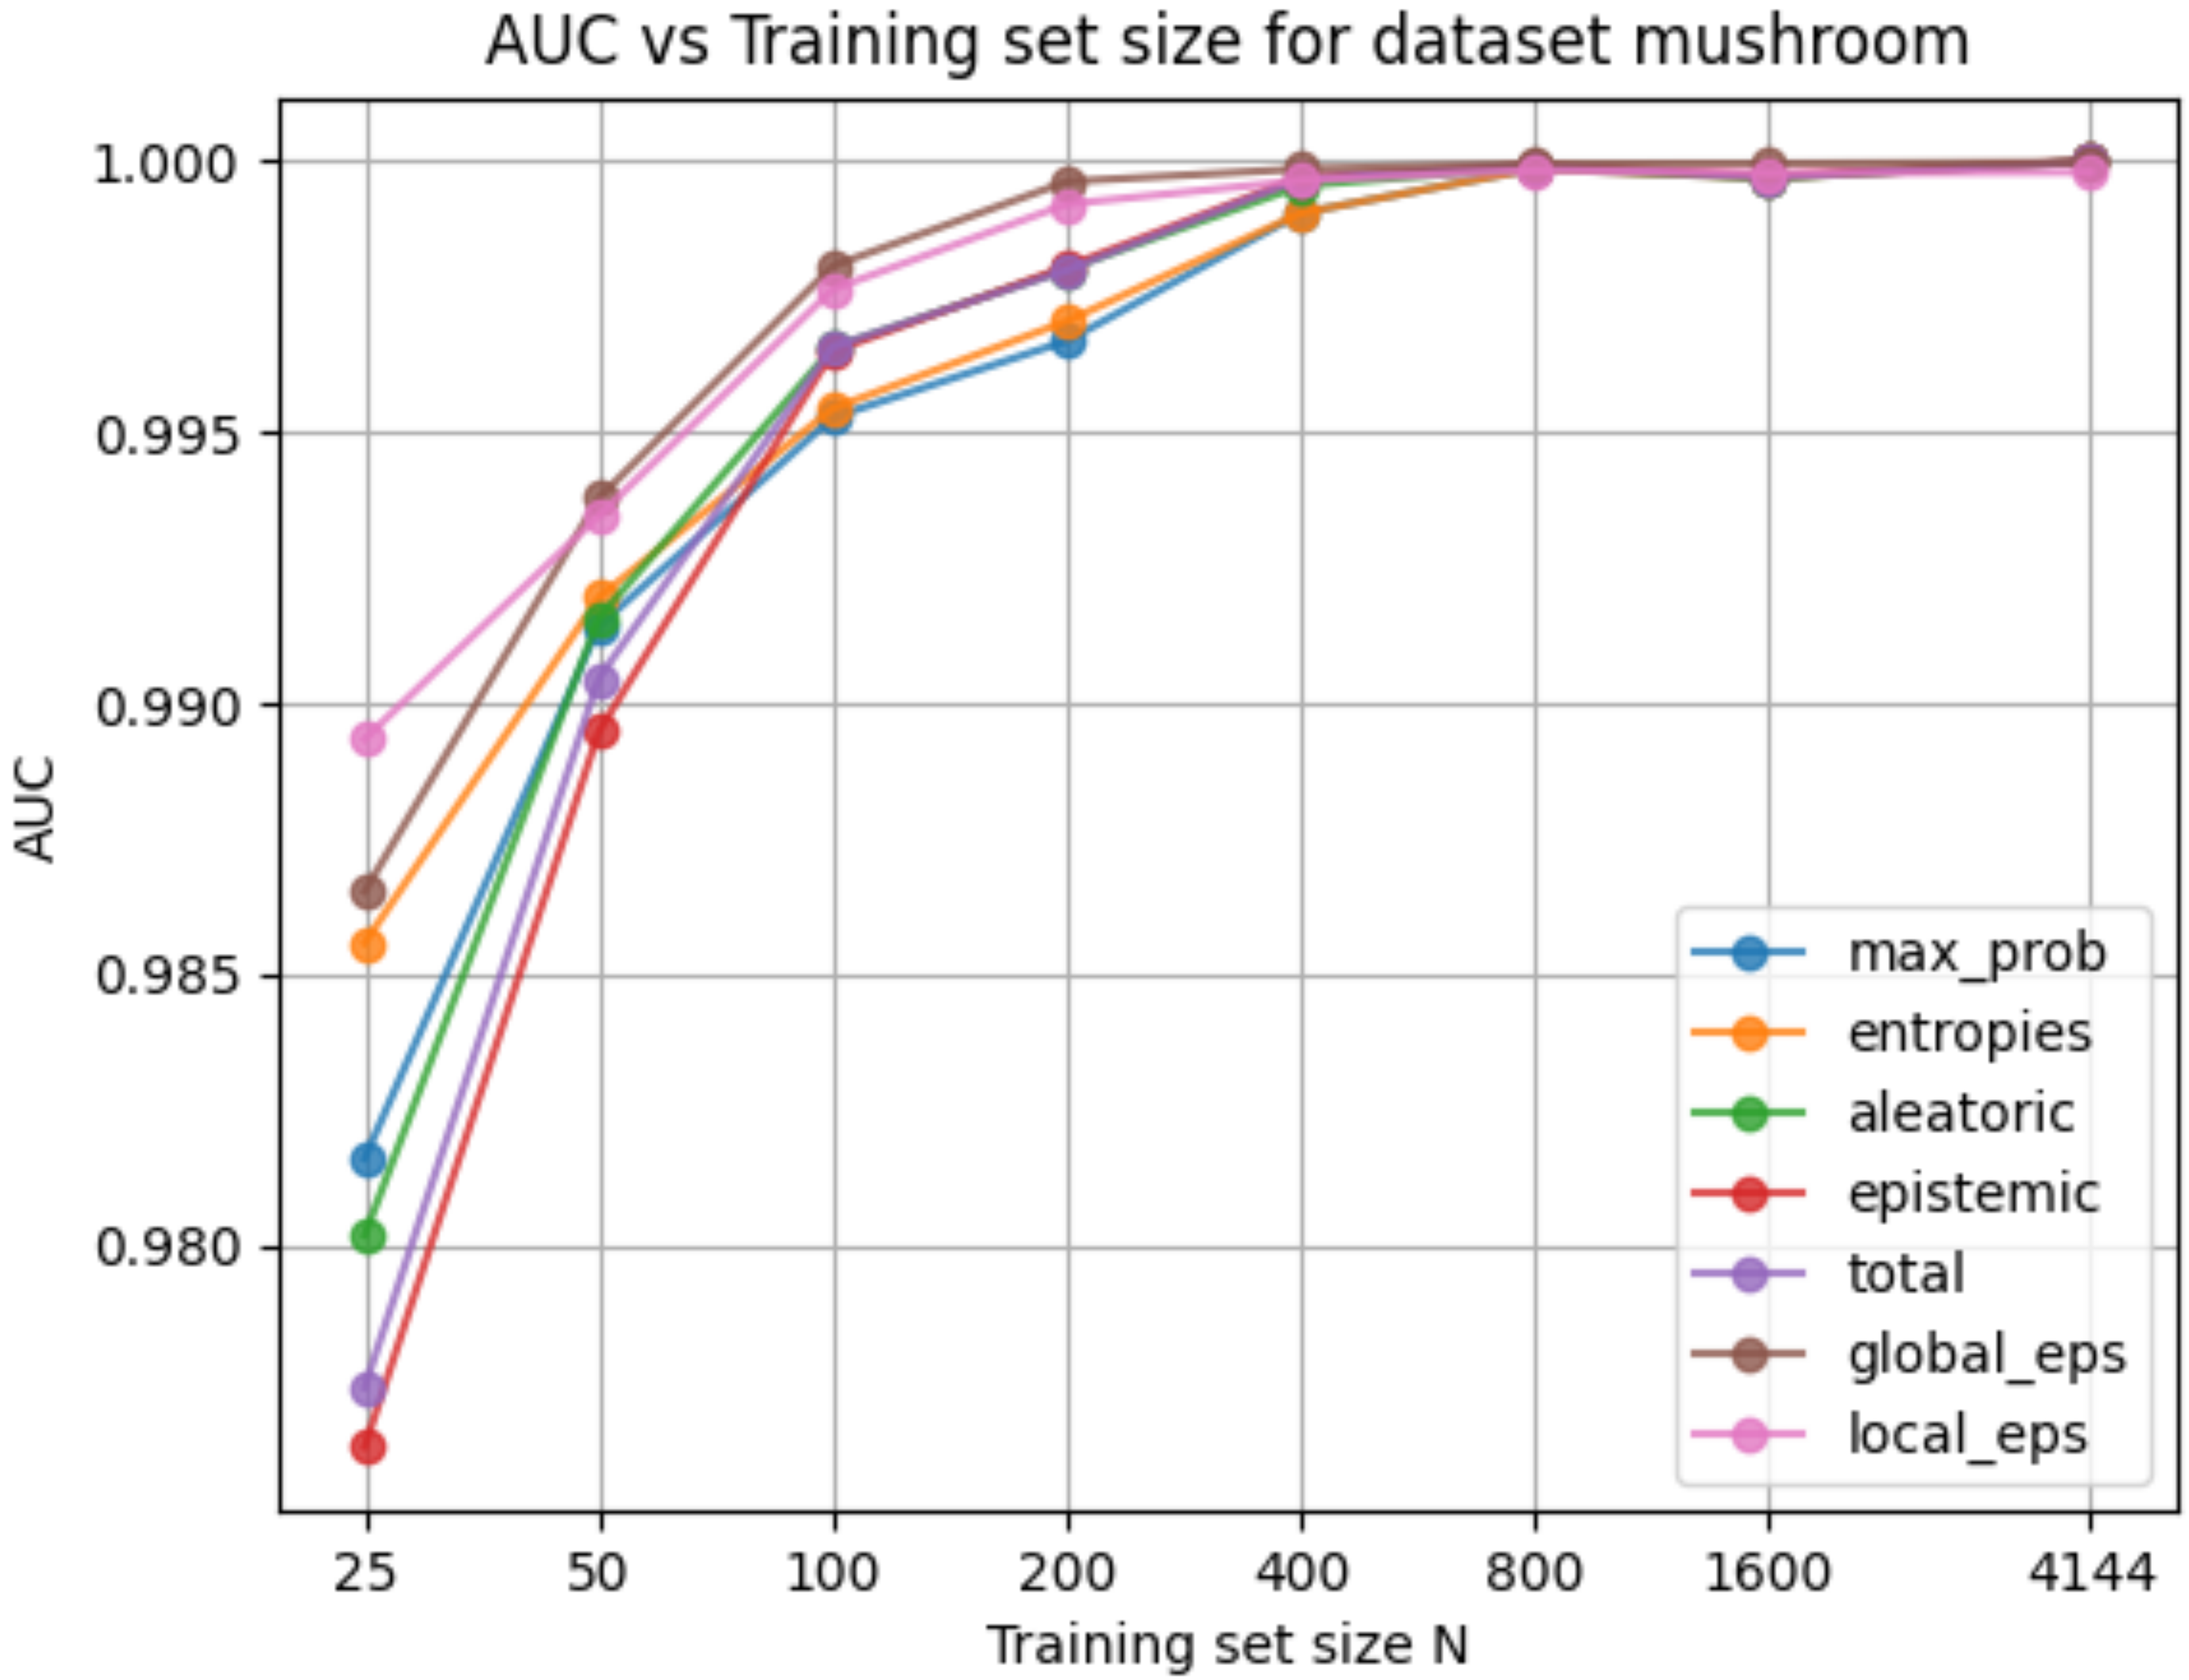

AUC vs Training set size for dataset mushroom
1.000
0.995
0.990
0.985
0.980
AUC
25
50
100
200
400
800
1600
4144
Training set size N
max_prob
entropies
aleatoric
epistemic
total
global_eps
local_eps

ROBUSTNESS QUANTIFICATION

- correlates nicely with accuracy
- works for different types of model architectures
- is conceptually different from UQ
- also works with global perturbations ✓
- is competitive with UQ ✓
- is good with distribution shift and small data sets ✓
- is more stable than UQ ✓

NBC

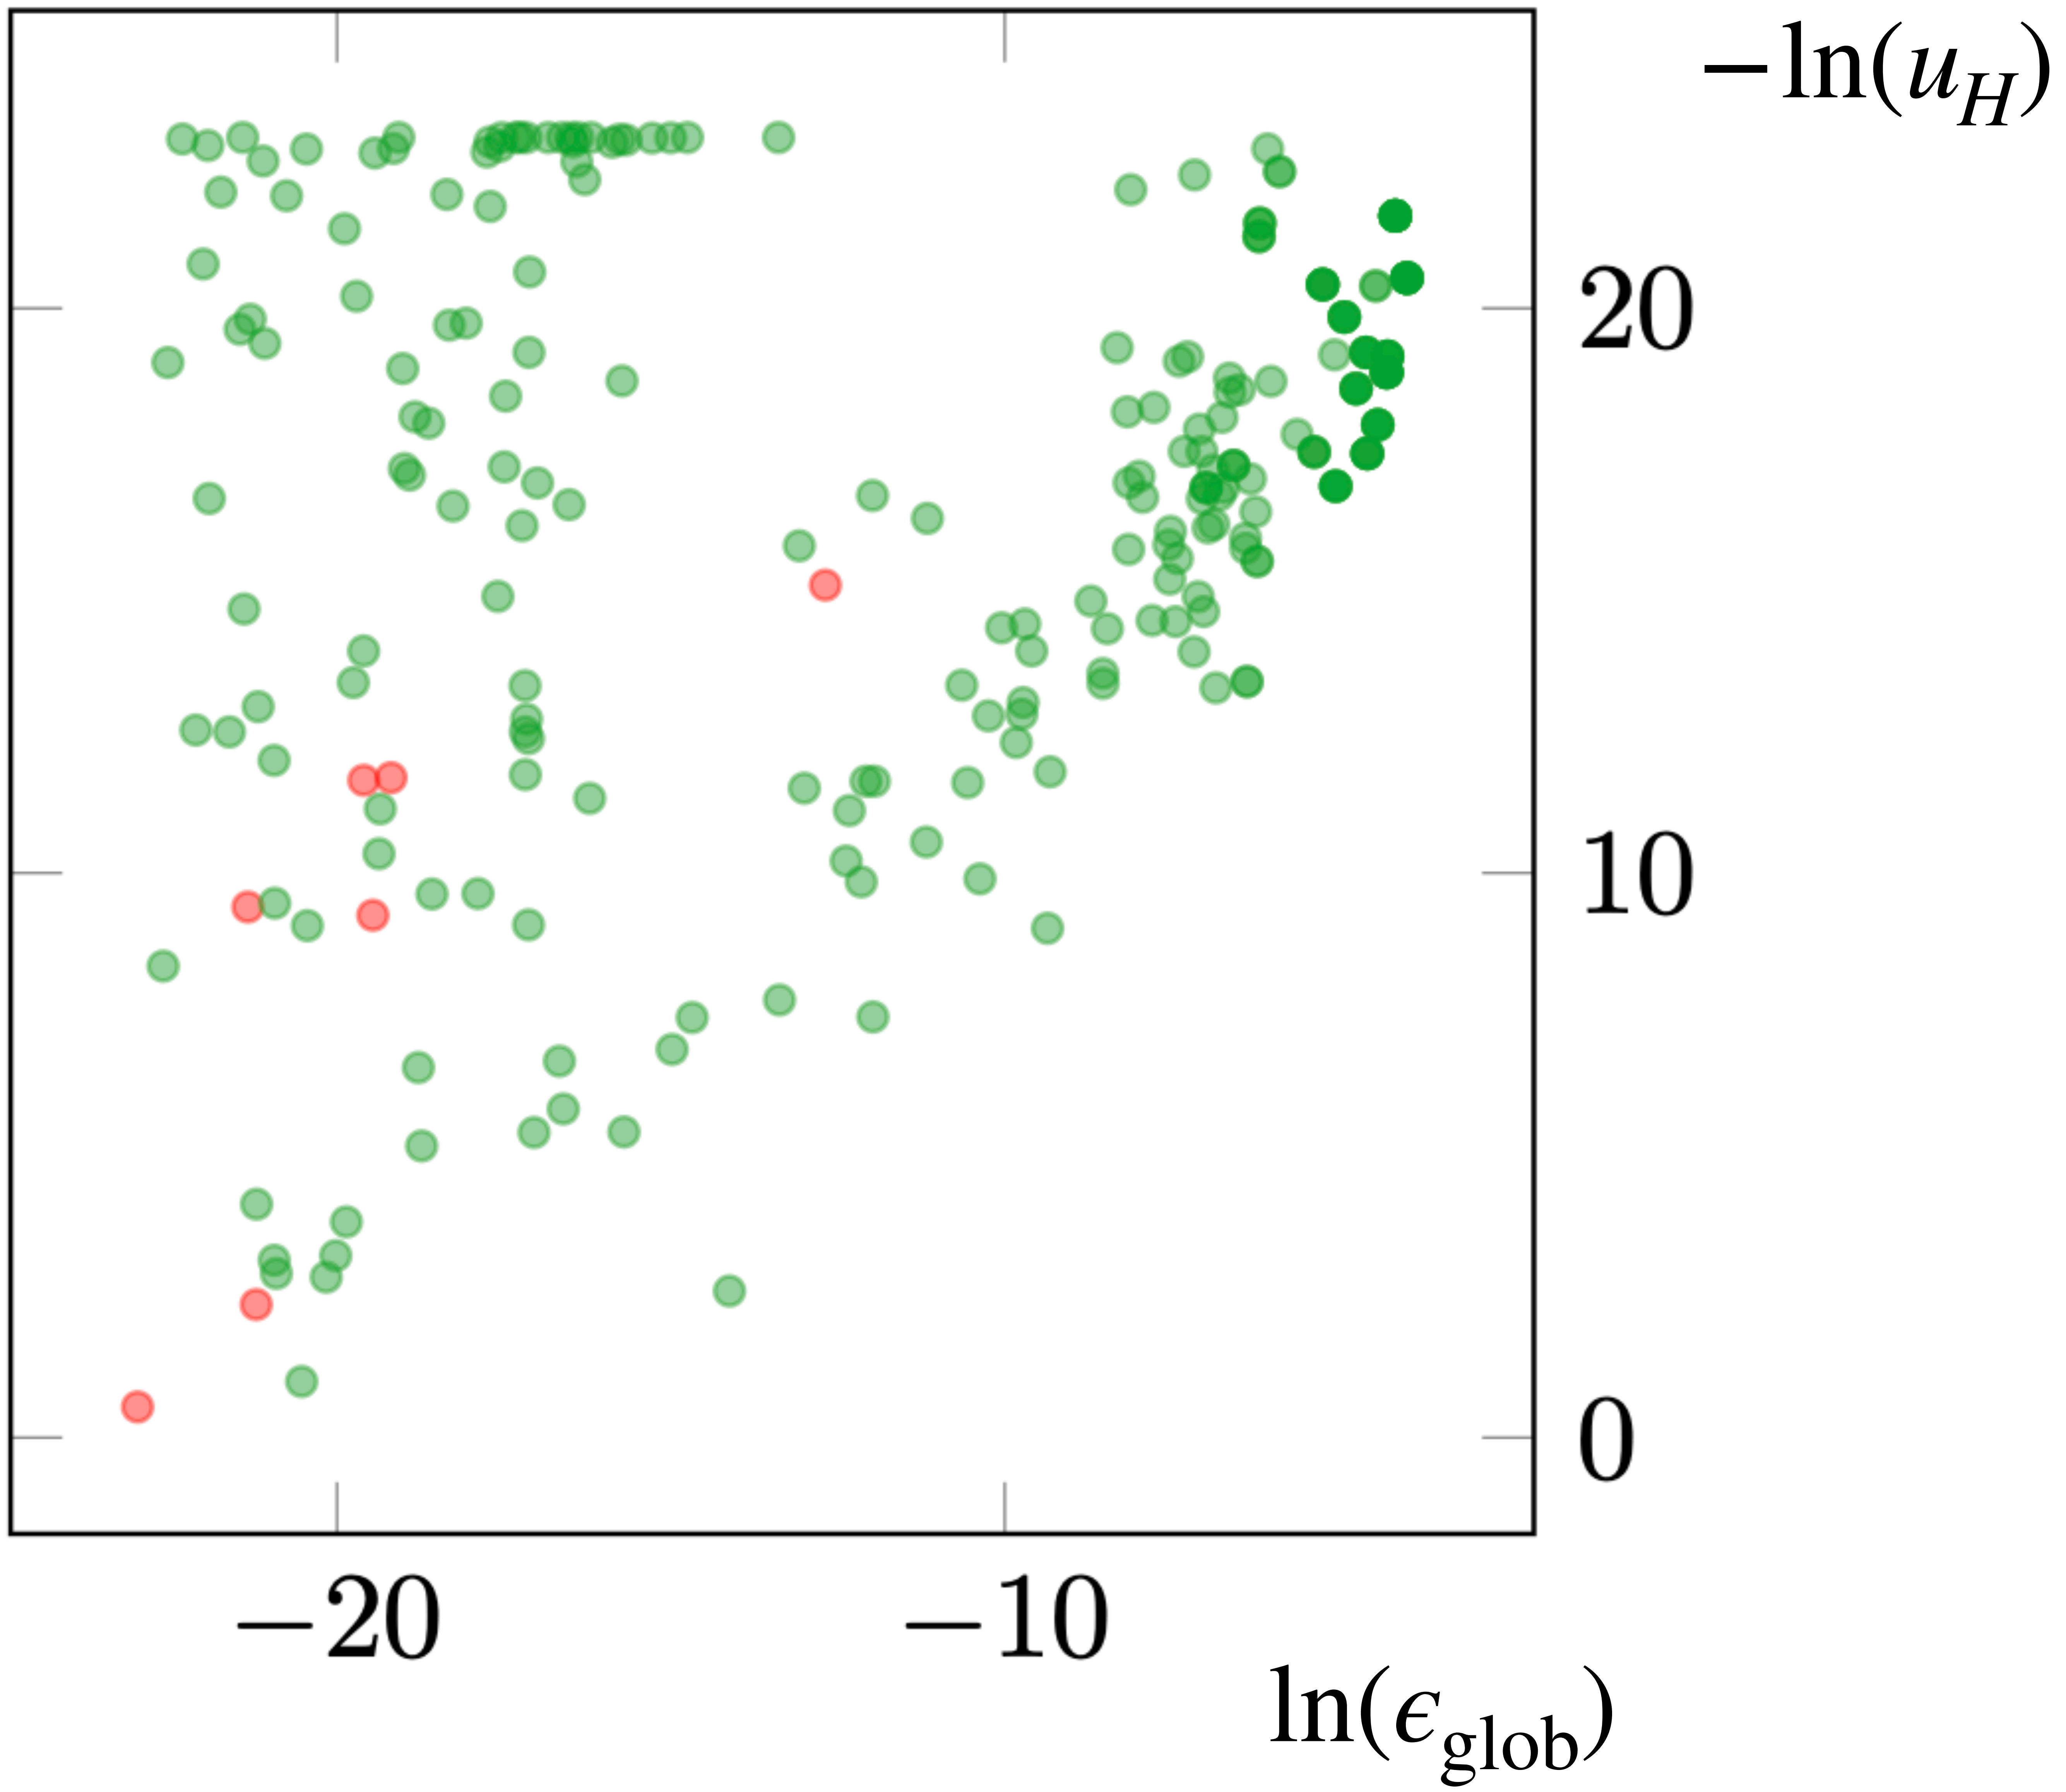

[7]

ROBUSTNESS QUANTIFICATION

- correlates nicely with accuracy ✓
- works for different types of model architectures
- is conceptually different from UQ
- also works with global perturbations ✓
- is competitive with and complementary to UQ ✓
- is good with distribution shift and small data sets
- is more stable than UQ

NBC

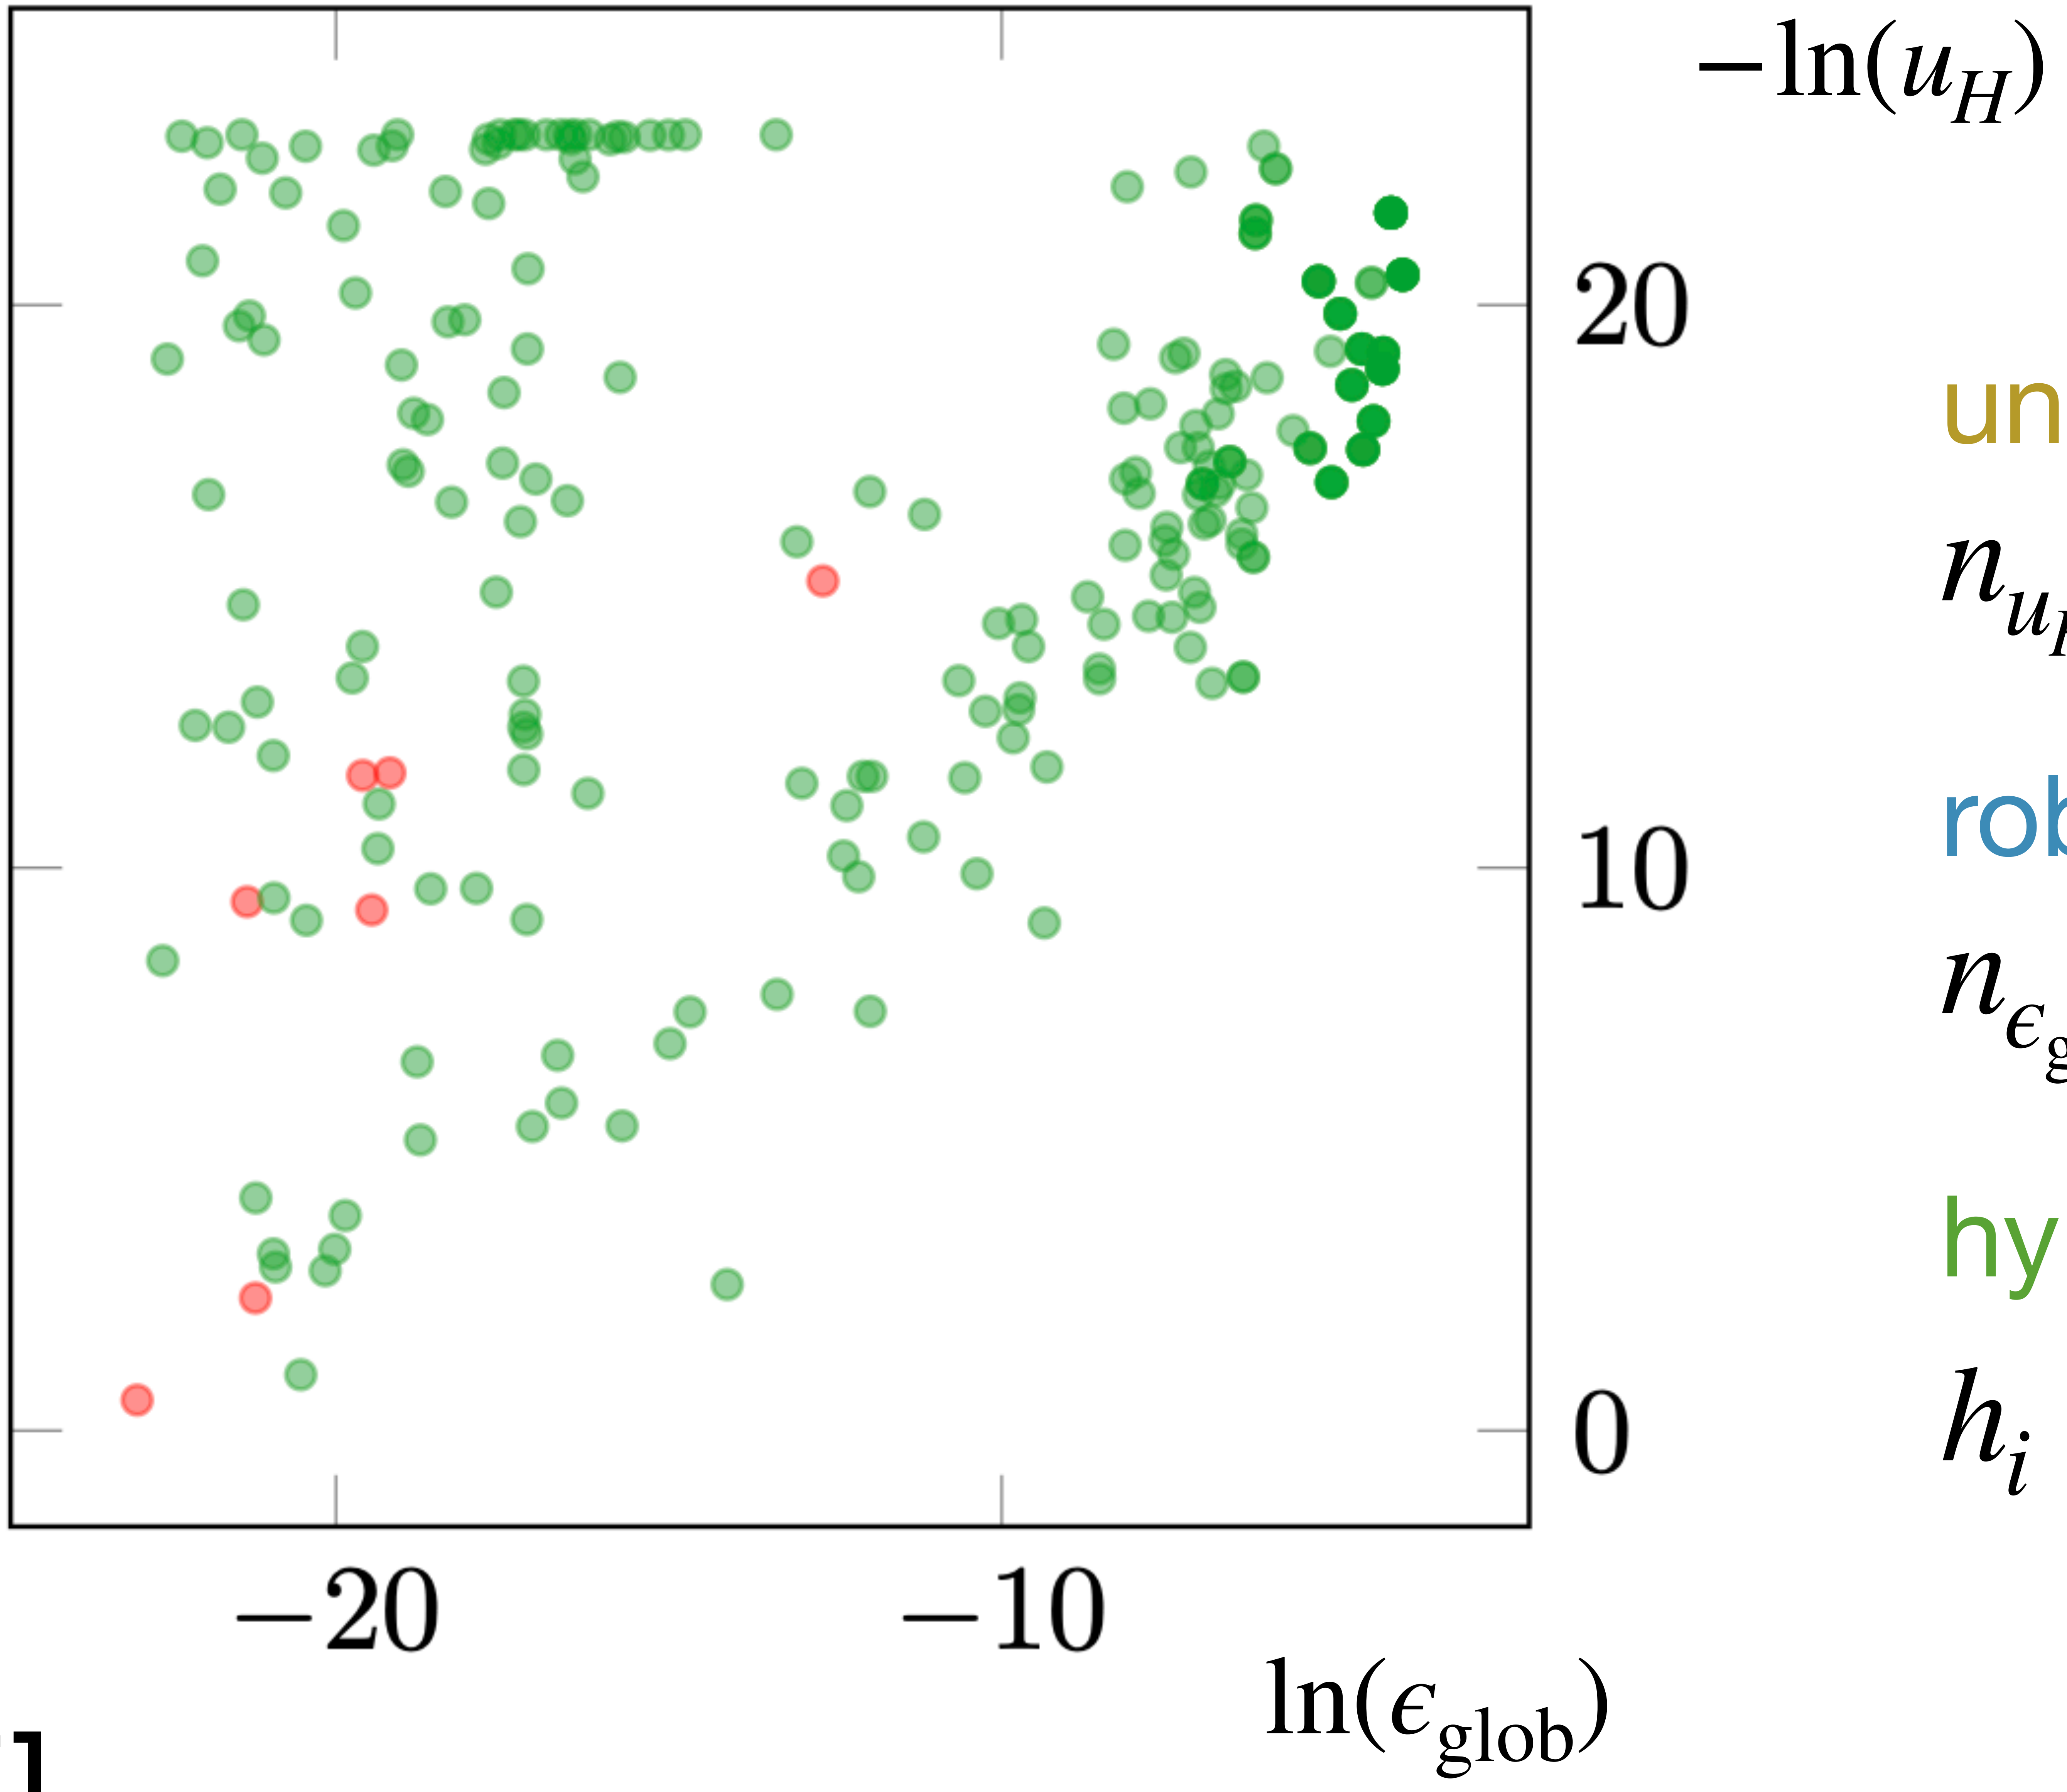

uncertainty ordering

$n_{u_H,i}$: order of i according to u_H

robustness ordering

$n_{\epsilon_{\text{glob}},i}$: order of i according to u_H

hybrid ordering

$$h_i = \gamma n_{u_H,i} + (1-\gamma) n_{\epsilon_{\text{glob}},i}$$

[7]

NBC

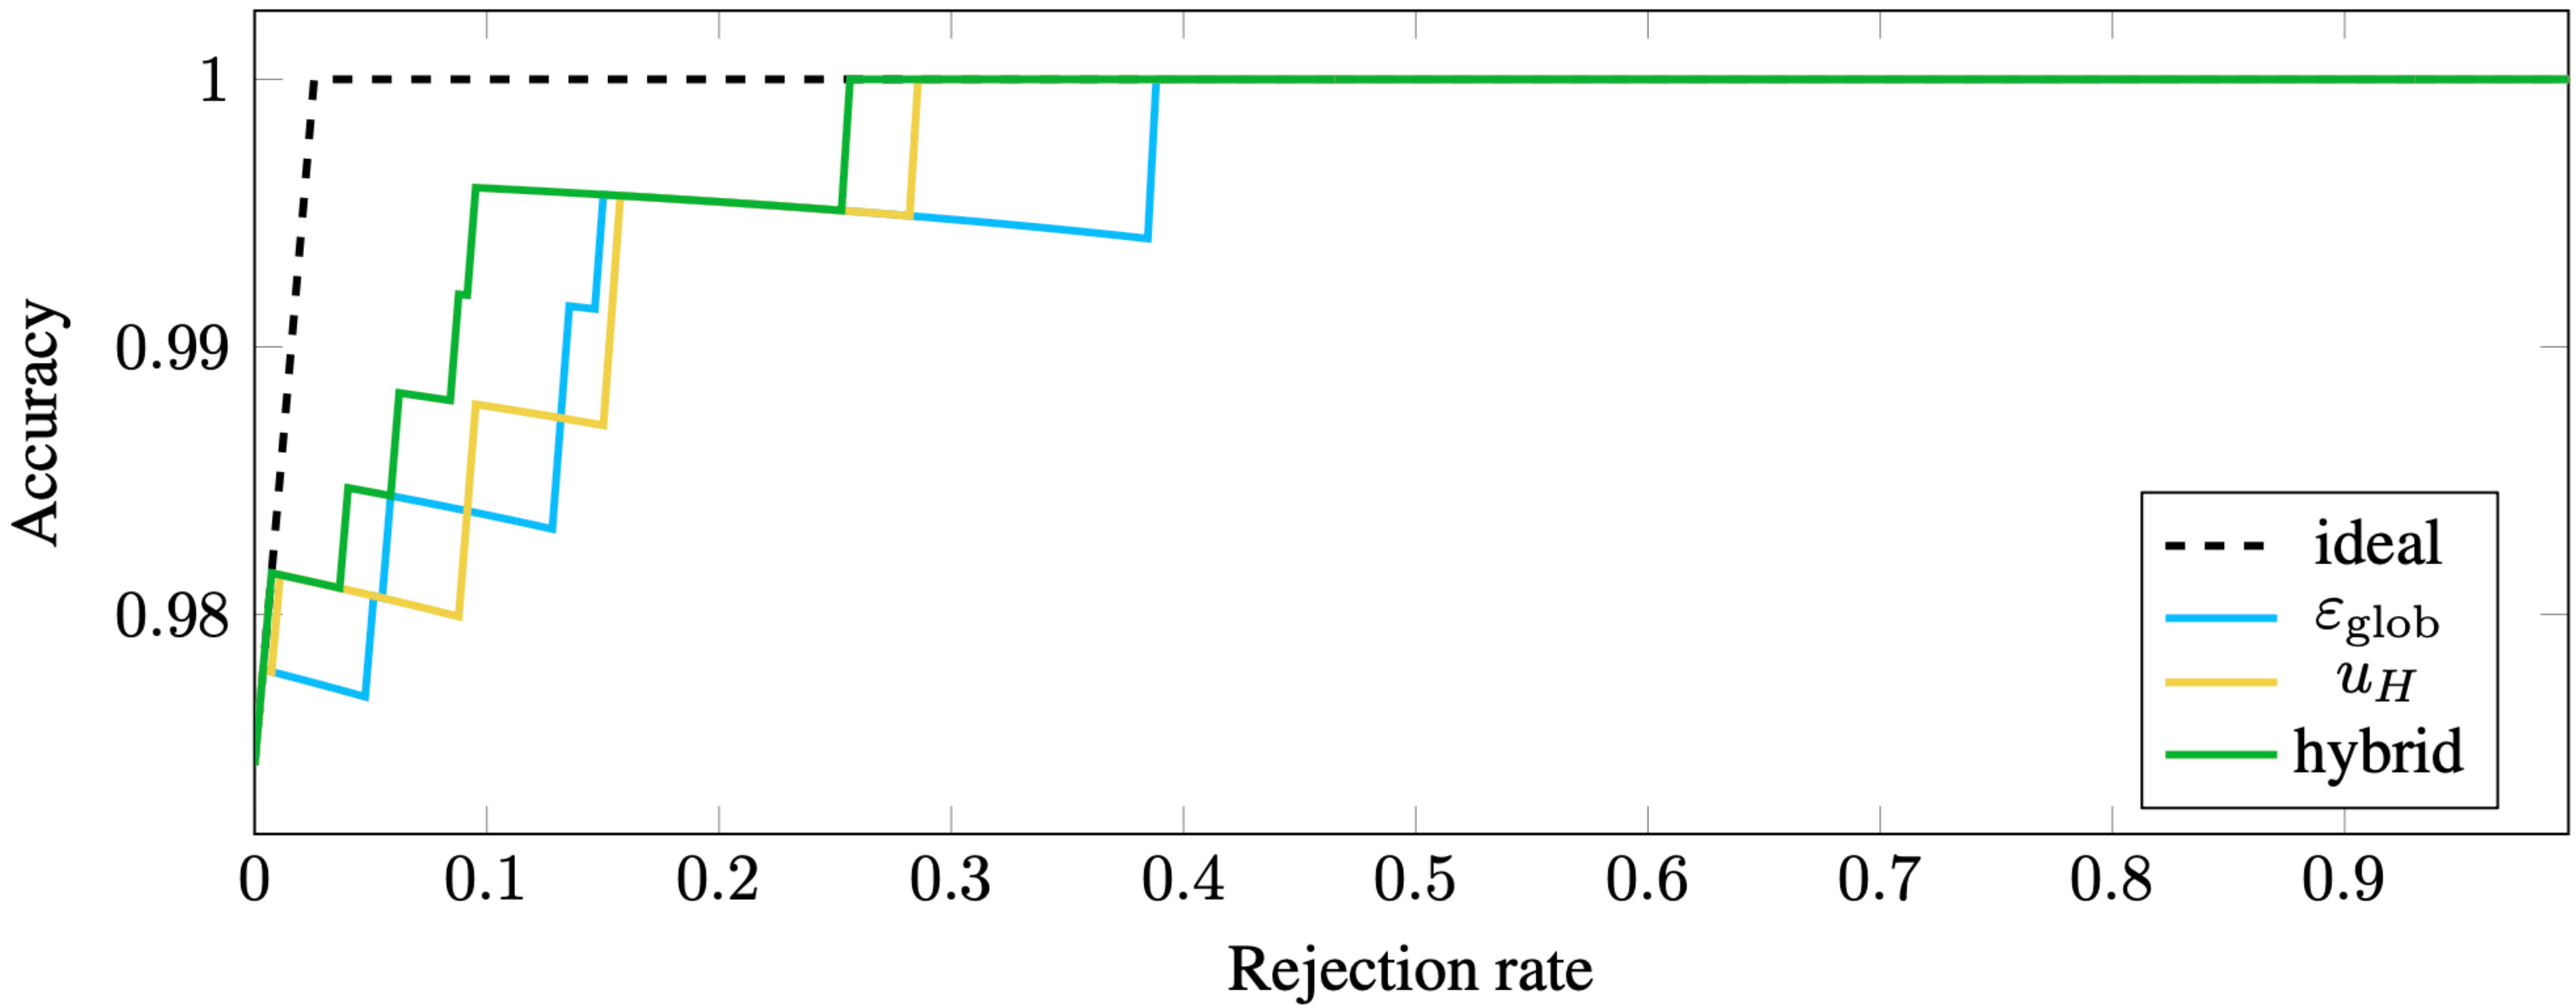

1
0.99
0.98
Accuracy
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
Rejection rate
ideal
$\varepsilon_{\text{glob}}$
u_H
hybrid

[7]

NBC

Dataset	u_{H}	$\varepsilon_{\mathrm{loc}}$	hybrid	γ	$\varepsilon_{\mathrm{glob}}$	hybrid	γ
Adult	0.9295	0.9066	**0.9295**	1.00	0.7690	**0.9295**	1.00
Austr. Cr.	0.9236	0.9139	**0.9265**	0.75	0.8872	**0.9246**	0.86
Bank M.	0.9485	0.9452	**0.9485**	0.55	0.9299	0.9481	0.88
BCW	0.9968	0.9962	**0.9974**	0.52	0.9961	**0.9978**	0.53
German Cr.	0.8338	**0.8380**	0.8378	0.53	0.7972	**0.8376**	0.85
Heart dis.	0.7602	0.7540	**0.7602**	0.95	0.6761	0.7600	0.95
Lymphogr.	0.9440	0.9419	0.9428	0.77	0.8981	0.9425	0.88
NPHA	0.4962	**0.5021**	0.4917	0.77	**0.5159**	0.4913	0.96
Nursery	0.9813	0.9822	**0.9824**	0.28	0.9730	**0.9814**	0.91
Solar (big)	0.8603	**0.8926**	0.8874	0.23	0.8693	**0.8836**	0.71
Solar (small)	0.8709	0.8597	0.8666	0.19	0.7990	**0.8797**	0.78
SPECT	0.9458	0.8915	**0.9458**	0.99	0.5738	0.9457	0.99
Stud. Math	0.9434	0.9465	**0.9468**	0.31	0.9205	**0.9445**	0.60
Stud. Port	0.8898	**0.9276**	0.9093	0.77	0.8952	**0.9067**	0.79

[7]

ROBUSTNESS QUANTIFICATION

- correlates nicely with accuracy ✓
- works for different types of model architectures
- is conceptually different from UQ
- also works with global perturbations ✓
- is competitive with and complementary to UQ ✓
- is good with distribution shift and small data sets
- is more stable than UQ

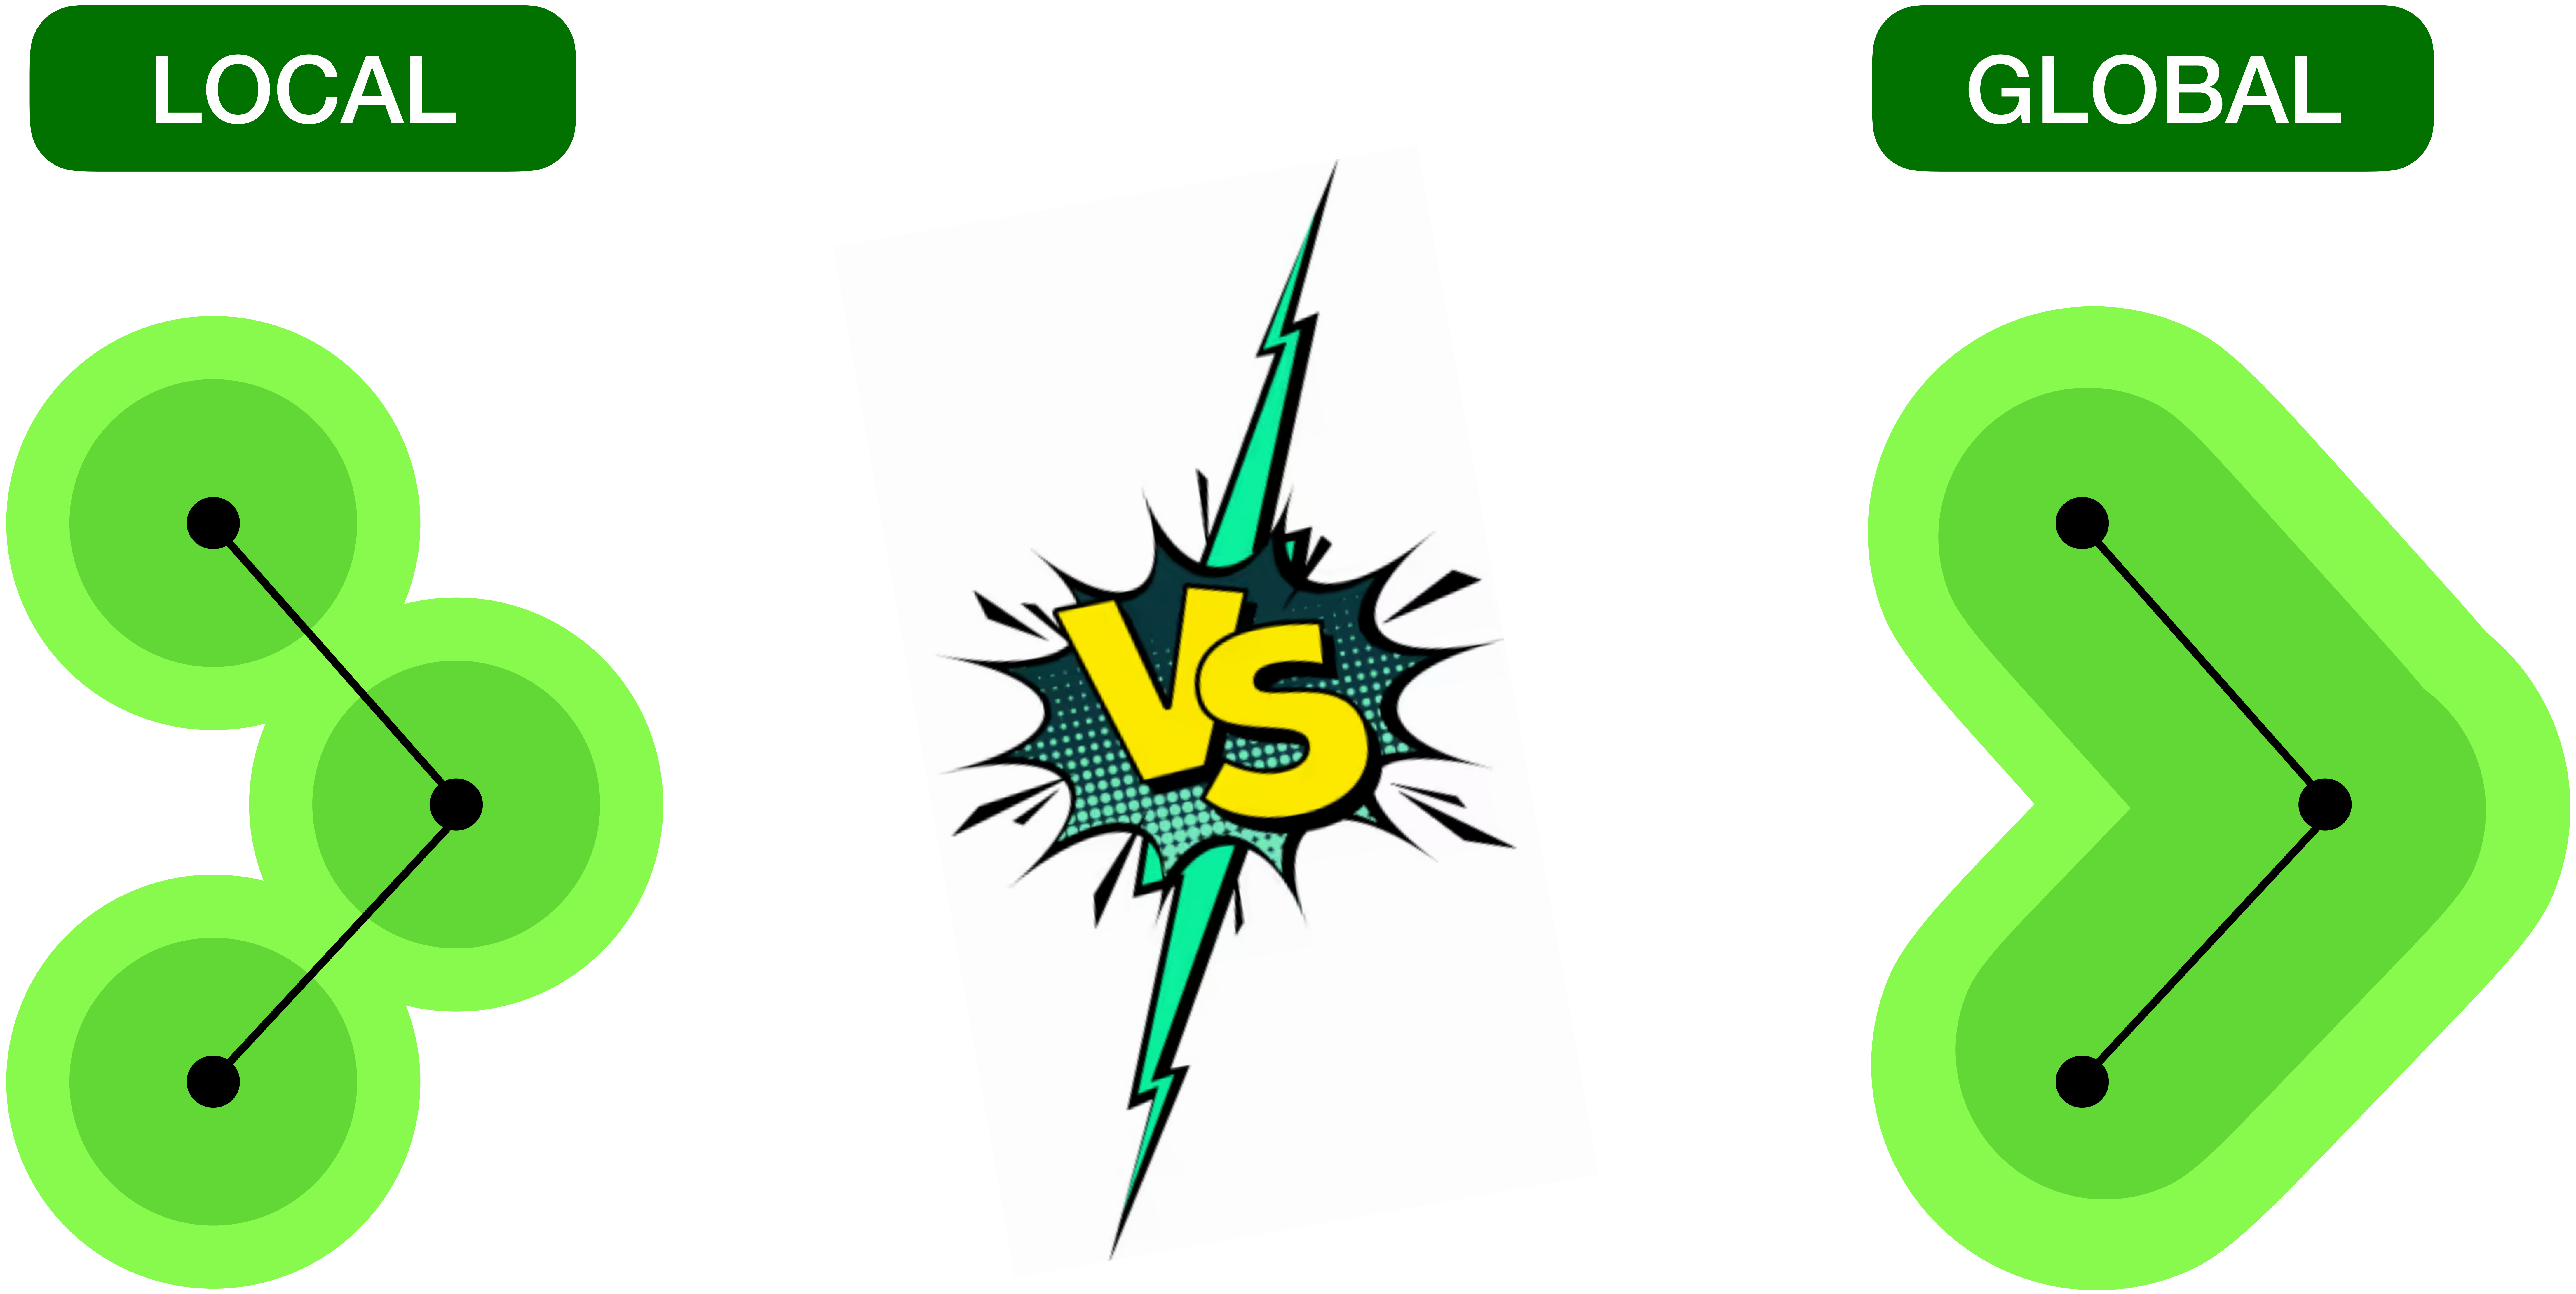
LOCAL
VS
GLOBAL

ϵ-CONTAMINATION

OTHER STUFF

distance-based, ...

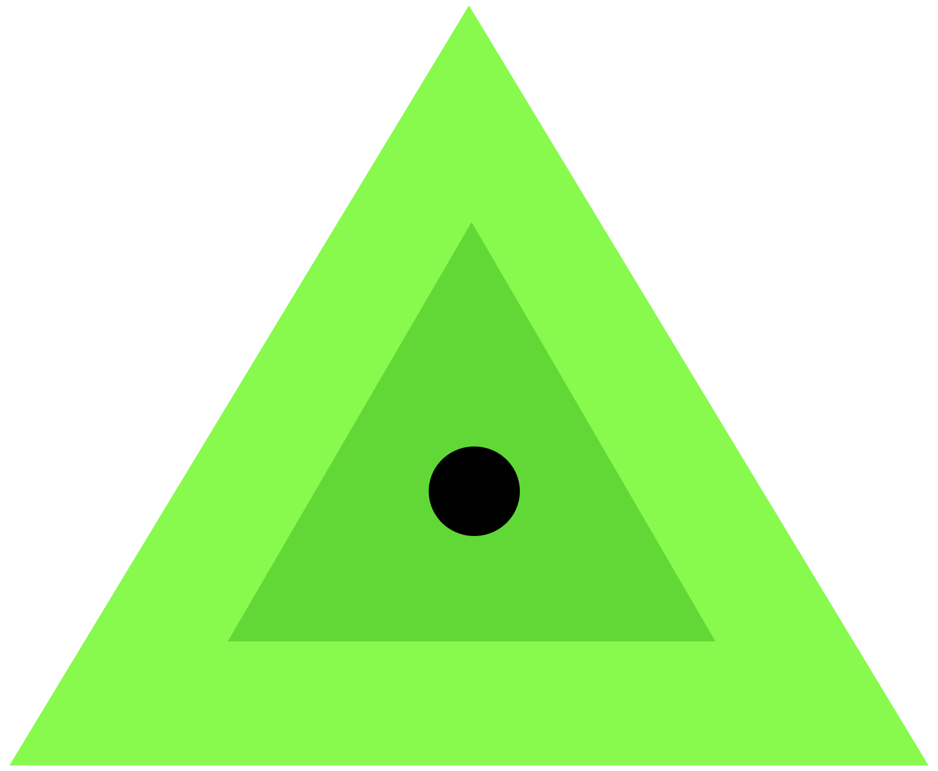

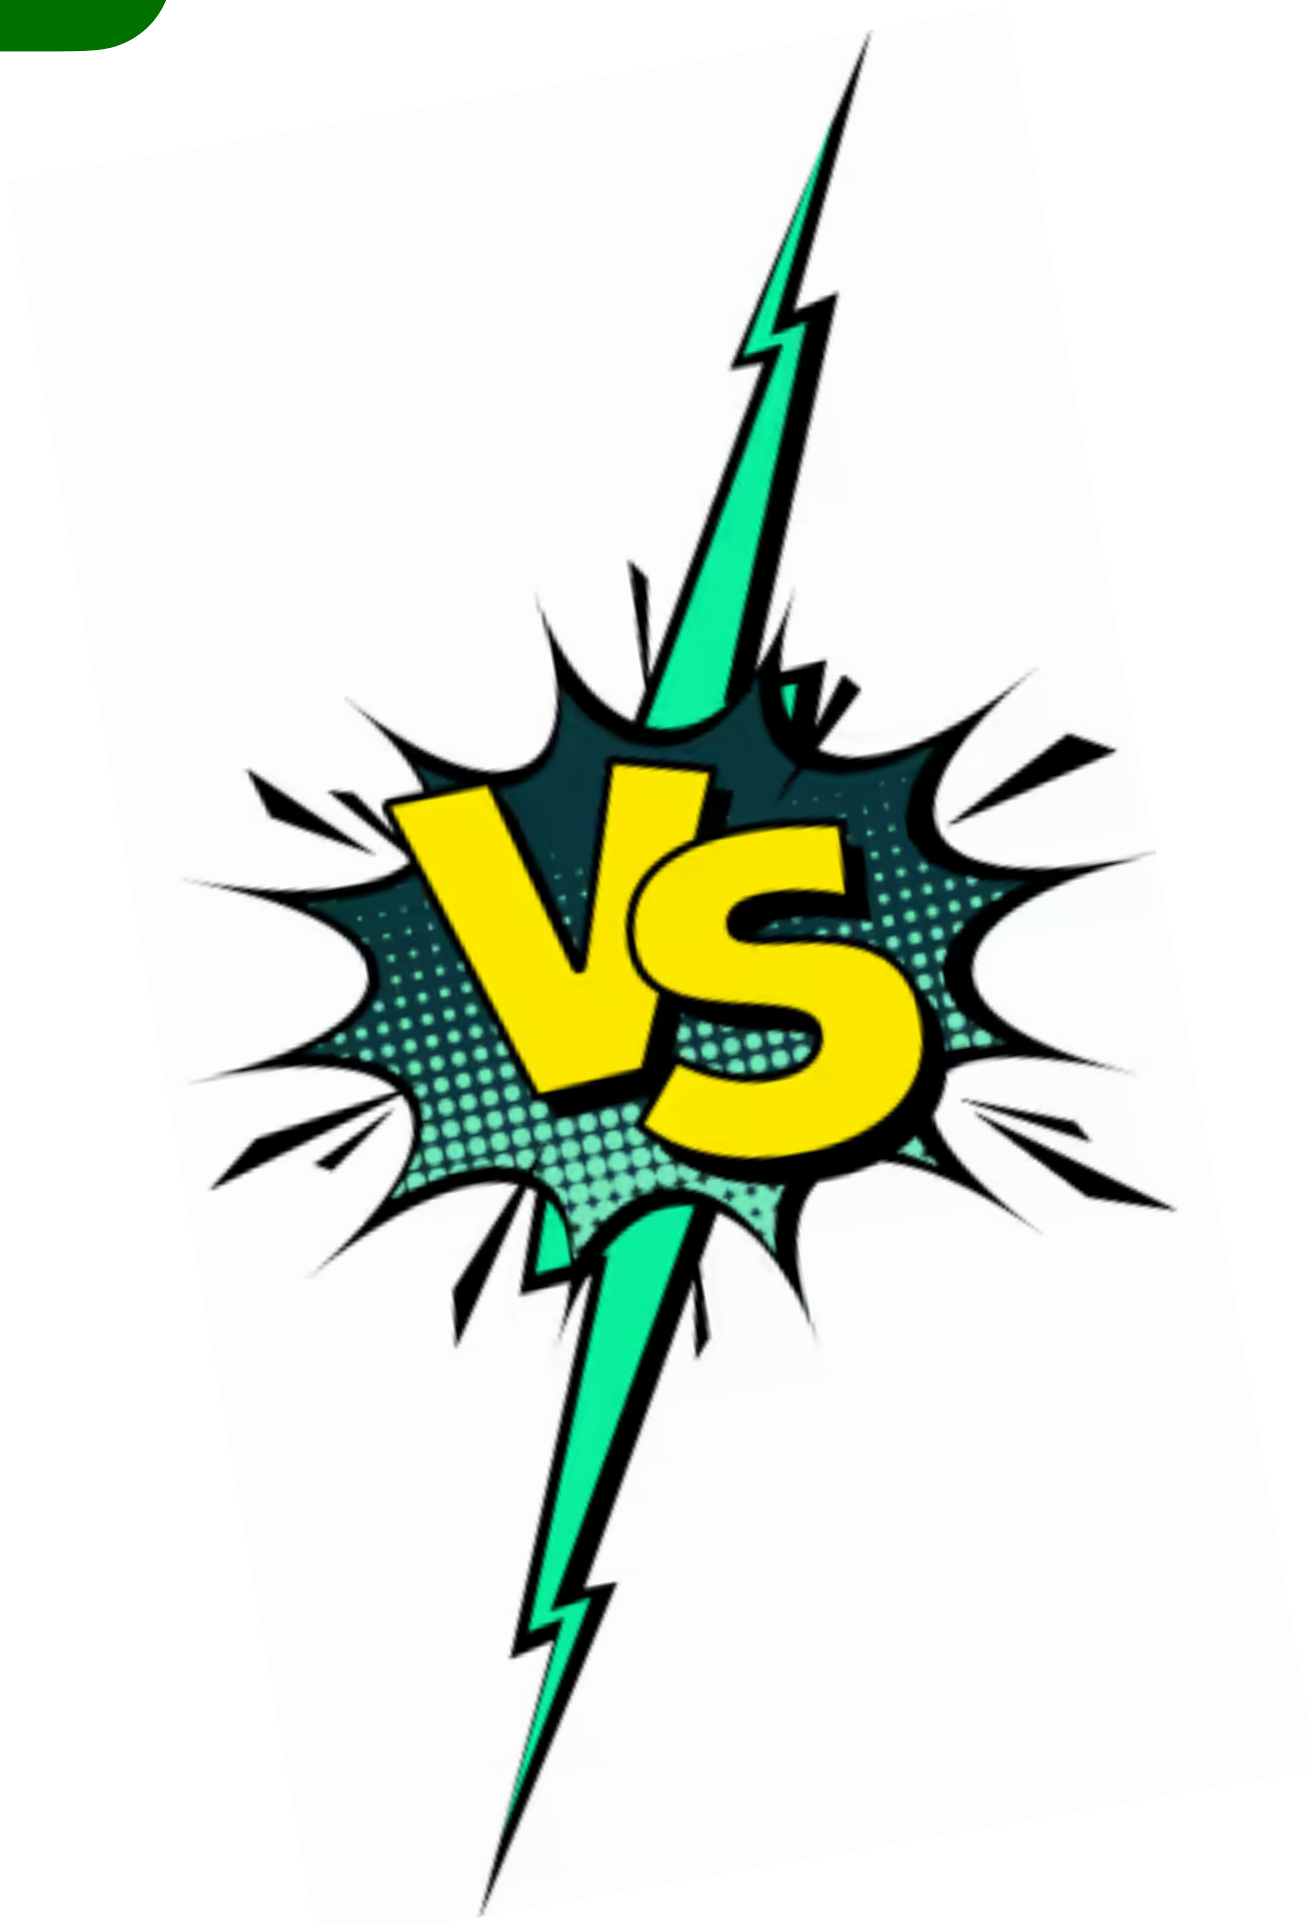

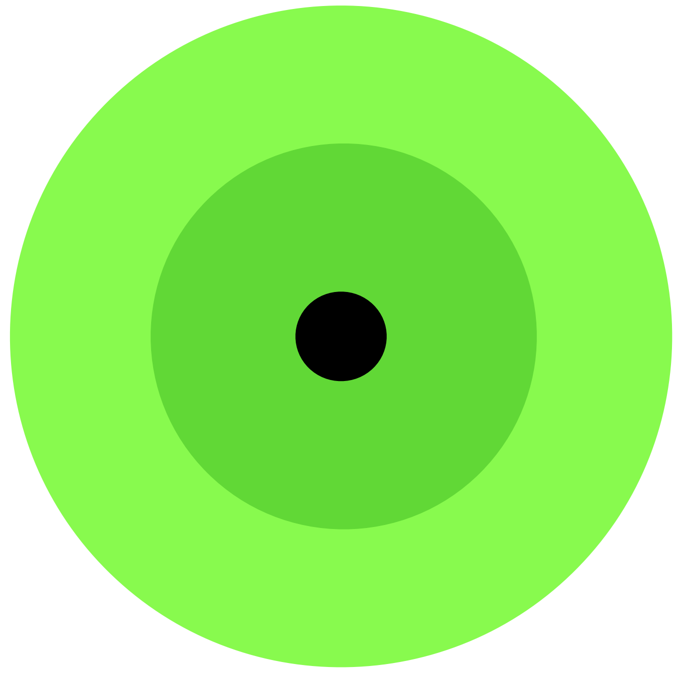

$$\mathscr{P}_\epsilon = \left\{(1-\epsilon)P_{\text{classif}} + \epsilon P : P \in \Delta\right\}$$

$$\mathscr{P}_\epsilon = \left\{P \in \Delta : \mathrm{d}(P_{\text{classif}}, P) < \epsilon\right\}$$

OTHER STUFF

distance-based, ...

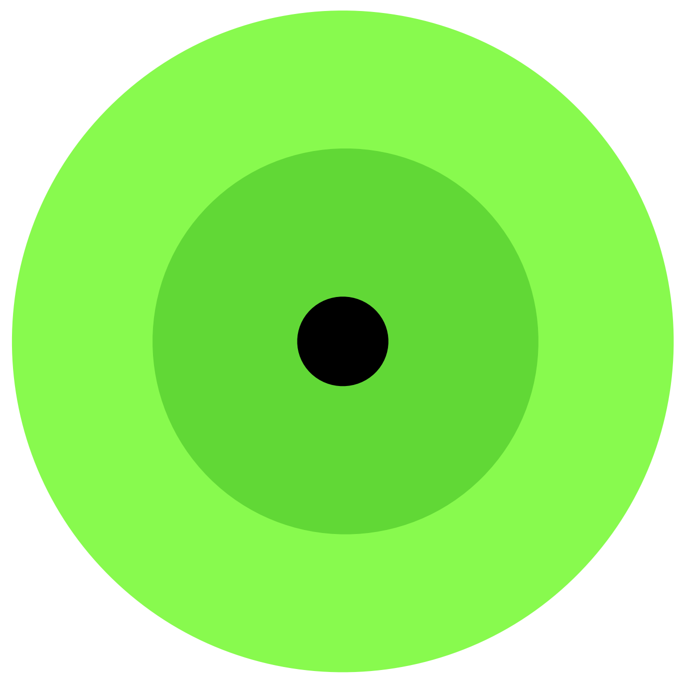

$$\mathscr{P}_\epsilon$$

$$\|$$

$$\left\{P \in \Delta : \mathrm{d}(P_{\text{classif}}, P) < \epsilon\right\}$$

$$\mathrm{d}(P_1, P_2) = \max\left\{\sup_A \frac{P_1(A)}{P_2(A)}, \sup_A \frac{P_2(A)}{P_1(A)}\right\}$$

GLOBAL

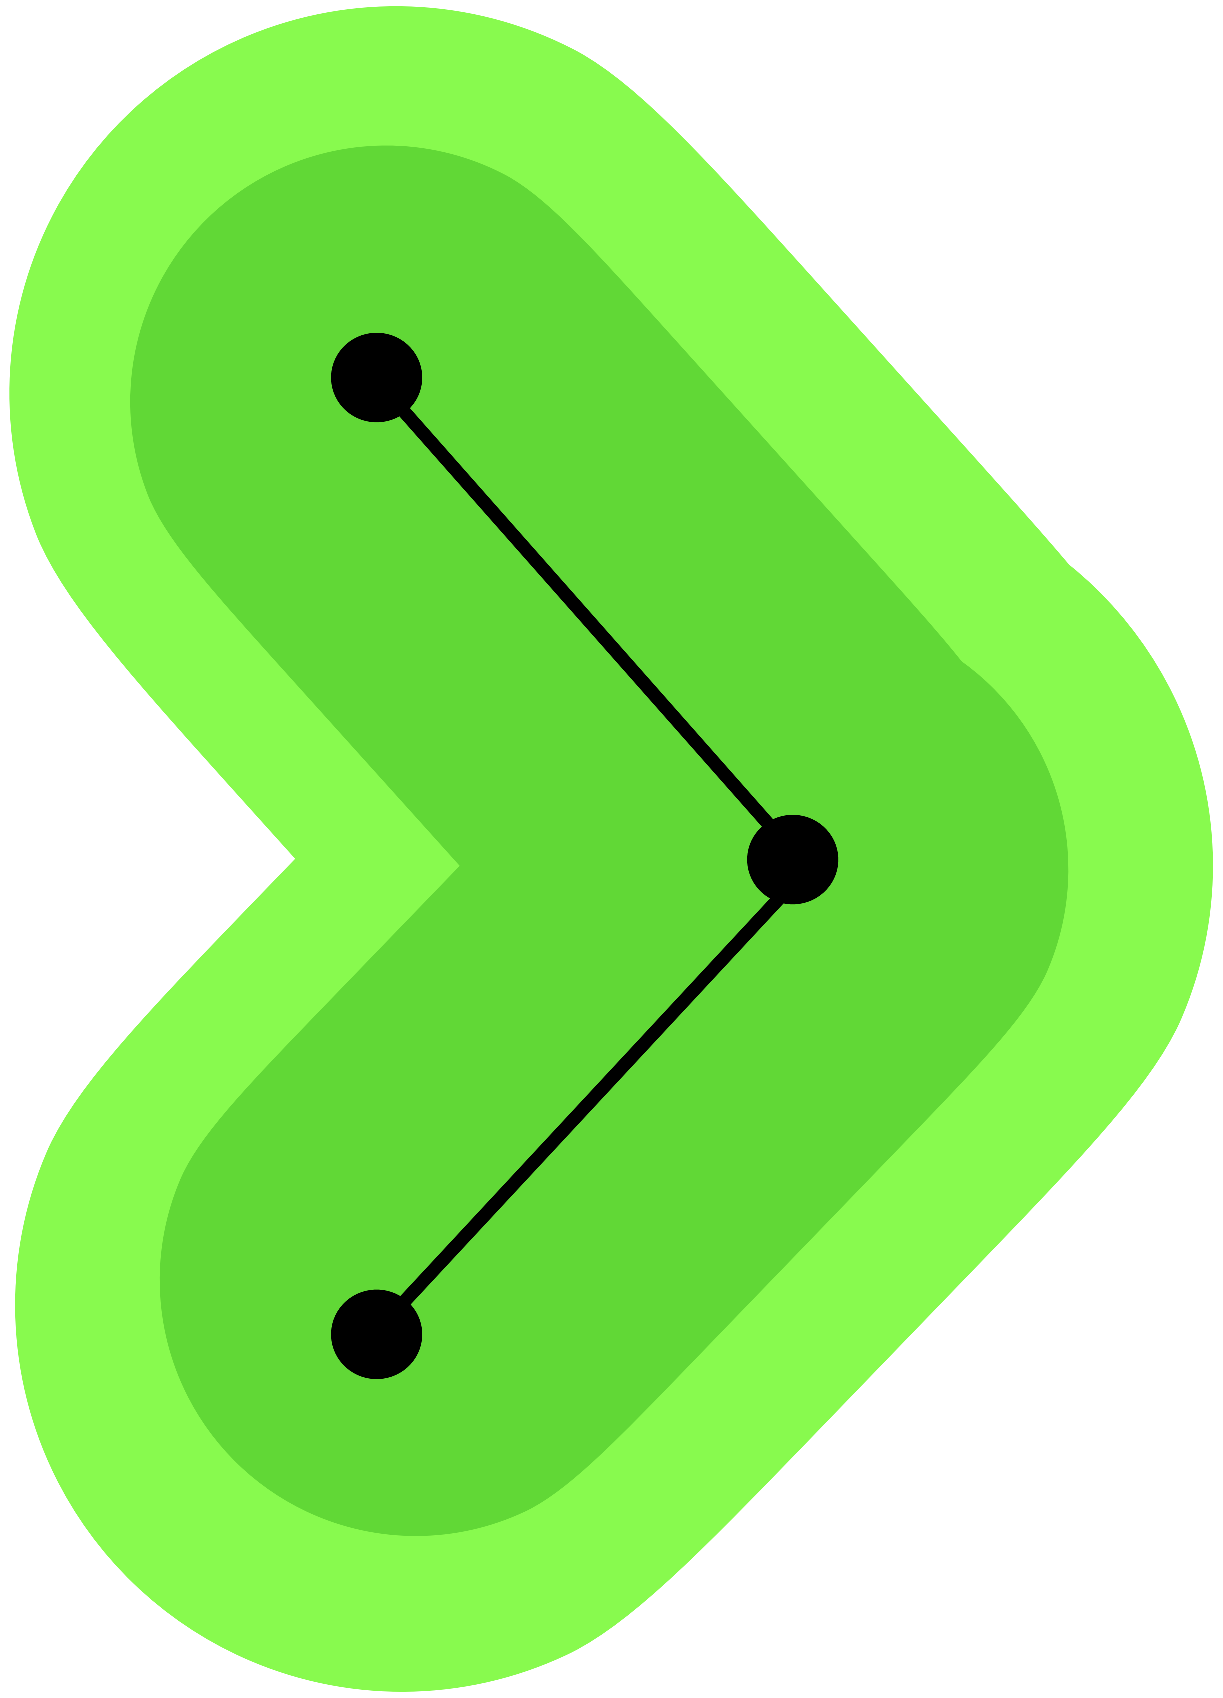

+

OTHER STUFF

distance-based, ...

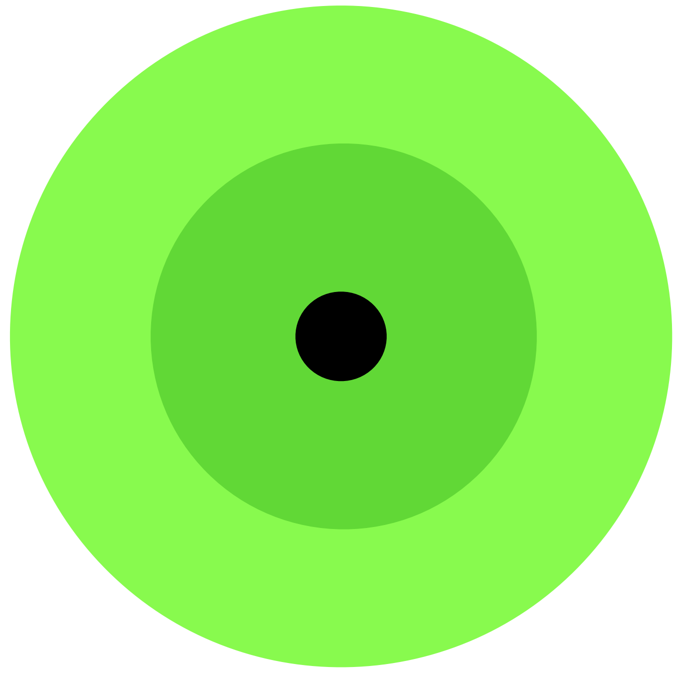

$$\mathscr{P}_\epsilon$$

$$\parallel$$

$$\left\{P \in \Delta : \mathrm{d}(P_{\text{classif}}, P) < \epsilon\right\}$$

GLOBAL

discrete features:

$$\epsilon_{\text{glob}} = \min \left\{ \frac{1}{1-\Delta}, \sqrt{\frac{P_{\text{classif}}(x, \hat{y})}{P_{\text{classif}}(x, \hat{y}_2)}} \right\}$$

$$\Delta = P_{\text{classif}}(x, \hat{y}) - P_{\text{classif}}(x, \hat{y}_2)$$

$$\hat{y}_2 = \arg \max_{y \in \mathcal{Y} \setminus \{\hat{y}\}} P_{\text{classif}}(y \mid x)$$

OTHER STUFF

distance-based, ...

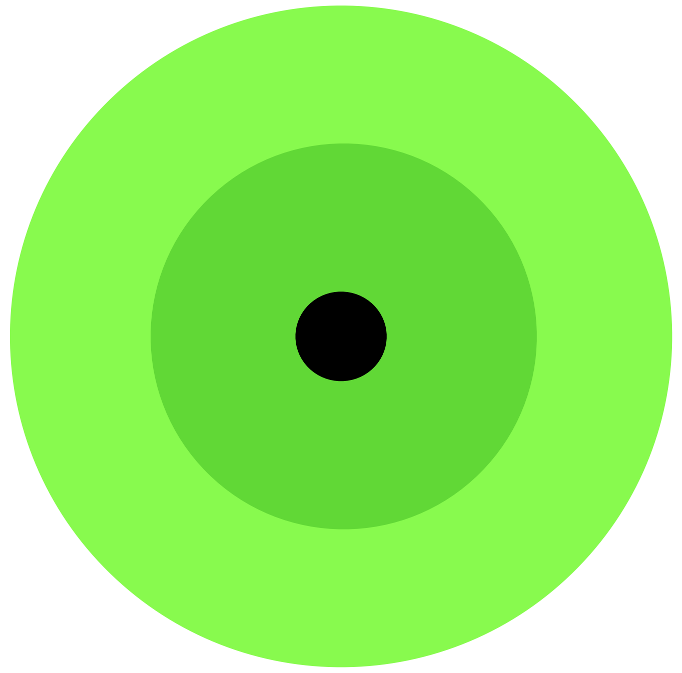

$$\mathscr{P}_\epsilon$$

$$\|$$

$$\left\{ P \in \Delta : \mathrm{d}(P_{\text{classif}}, P) < \epsilon \right\}$$

GLOBAL

OTHER STUFF

distance-based, ...

at least one continuous feature:

$$\epsilon_{\text{glob}} = \sqrt{\frac{P_{\text{classif}}(x, \hat{y})}{P_{\text{classif}}(x, \hat{y}_2)}}$$

$$\hat{y}_2 = \arg \max_{y \in \mathcal{Y} \setminus \{\hat{y}\}} P_{\text{classif}}(y \mid x)$$

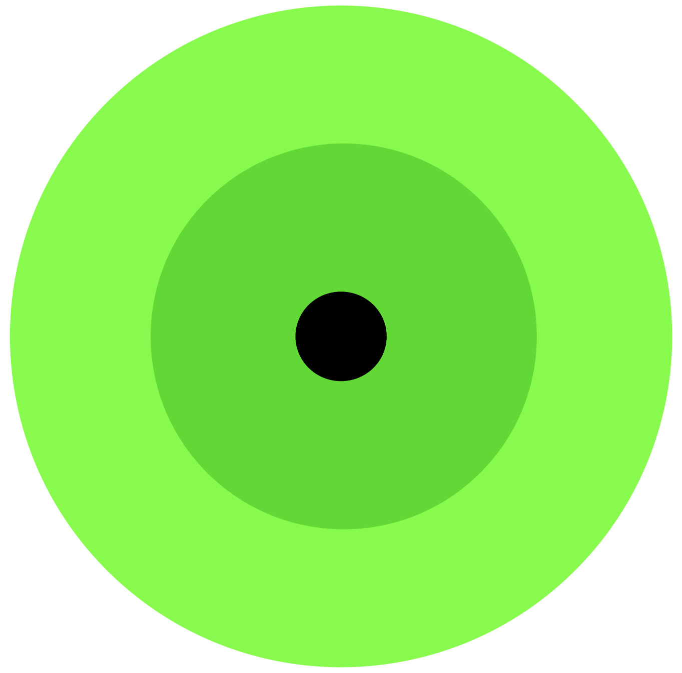

$$\mathcal{P}_\epsilon$$

$$\|$$

$$\left\{ P \in \Delta : \mathrm{d}(P_{\text{classif}}, P) < \epsilon \right\}$$

ROBUSTNESS QUANTIFICATION

- correlates nicely with accuracy
- works for different types of model architectures
- is conceptually different from UQ
- also works with global perturbations
- is competitive with and complementary to UQ
- is good with distribution shift and small data sets
- is more stable than UQ
- is not limited to epsilon-contamination ✓

CLASSIFICATION

features x

FEATURES

try to predict

predicted class $\hat{y} = h(x)$

CLASS

photo

medical data

…

GENERATIVE PROBABILISTIC CLASSIFIER h

label

best treatment

…

learn $P_{\text{classif}}(X, Y)$

predict the class y that maximizes $P_{\text{classif}}(y \mid x)$

CLASSIFICATION

features x

predicted class $\hat{y} = h(x)$

FEATURES $\xrightarrow{\text{try to predict}}$ CLASS

photo

medical data

...

label

best treatment

...

DISCRIMINATIVE PROBABILISTIC CLASSIFIER h

learn $P_{\text{classif}}(Y\,|\,x)$

predict the class y that maximizes $P_{\text{classif}}(y\,|\,x)$

GLOBAL

OTHER STUFF

distance-based, ...

at least one continuous feature:

$$\epsilon_{\text{glob}} = \sqrt{\frac{P_{\text{classif}}(x, \hat{y})}{P_{\text{classif}}(x, \hat{y}_2)}}$$

$$\hat{y}_2 = \arg \max_{y \in \mathscr{Y} \setminus \{\hat{y}\}} P_{\text{classif}}(y \,|\, x)$$

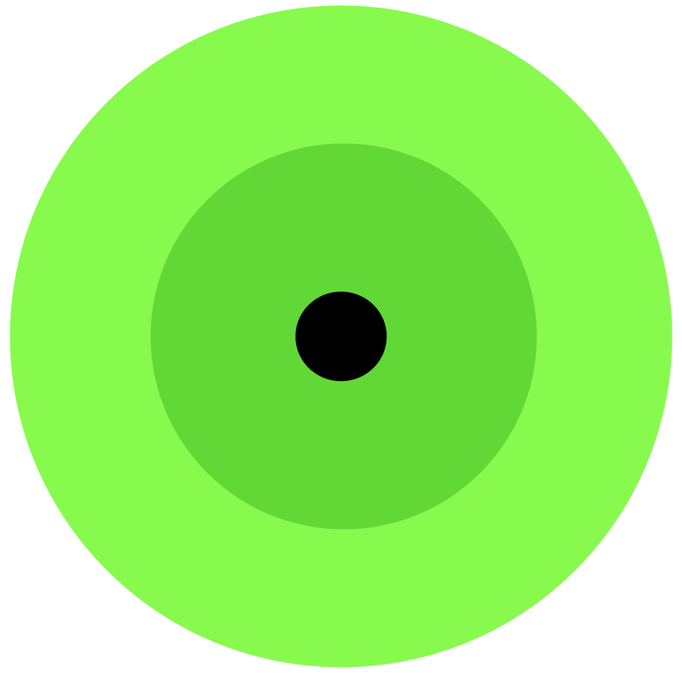

$$\mathscr{P}_\epsilon$$

$$\parallel$$

$$\left\{ P \in \Delta : \mathrm{d}(P_{\text{classif}}, P) < \epsilon \right\}$$

GLOBAL

OTHER STUFF

distance-based, ...

at least one continuous feature:

$$\epsilon_{\mathrm{glob}} = \sqrt{\frac{P_{\mathrm{classif}}(x,\hat{y})}{P_{\mathrm{classif}}(x,\hat{y}_2)}} = \sqrt{\frac{P_{\mathrm{classif}}(\hat{y}\,|\,x)}{P_{\mathrm{classif}}(\hat{y}_2\,|\,x)}}$$

$$\hat{y}_2 = \arg\max_{y\in\mathscr{Y}\setminus\{\hat{y}\}} P_{\mathrm{classif}}(y\,|\,x)$$

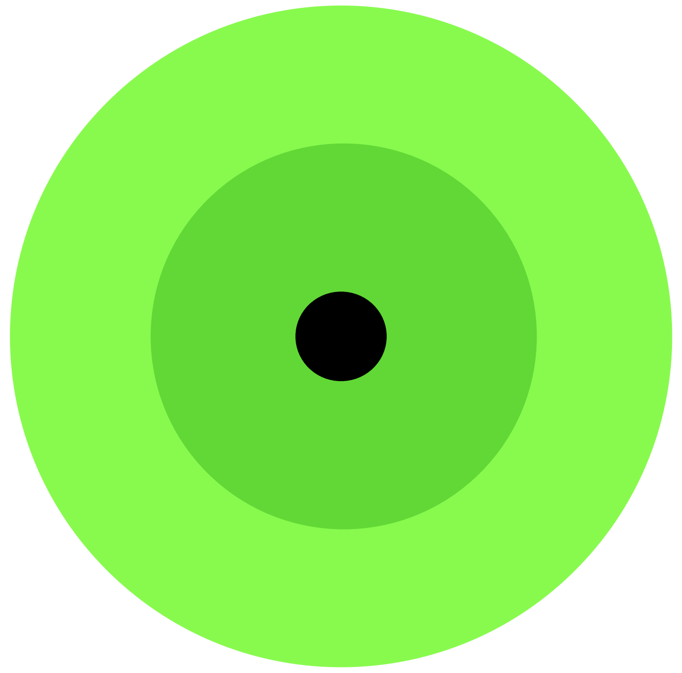

$$\mathscr{P}_\epsilon$$

$$\|$$

$$\left\{P\in\Delta : \mathrm{d}(P_{\mathrm{classif}},P) < \epsilon\right\}$$

ROBUSTNESS QUANTIFICATION

- correlates nicely with accuracy
- works for different types of model architectures
- is conceptually different from UQ
- also works with global perturbations
- is competitive with and complementary to UQ
- is good with distribution shift and small data sets
- is more stable than UQ
- is not limited to epsilon-contamination
- also works for discriminative classifiers ✓

RF

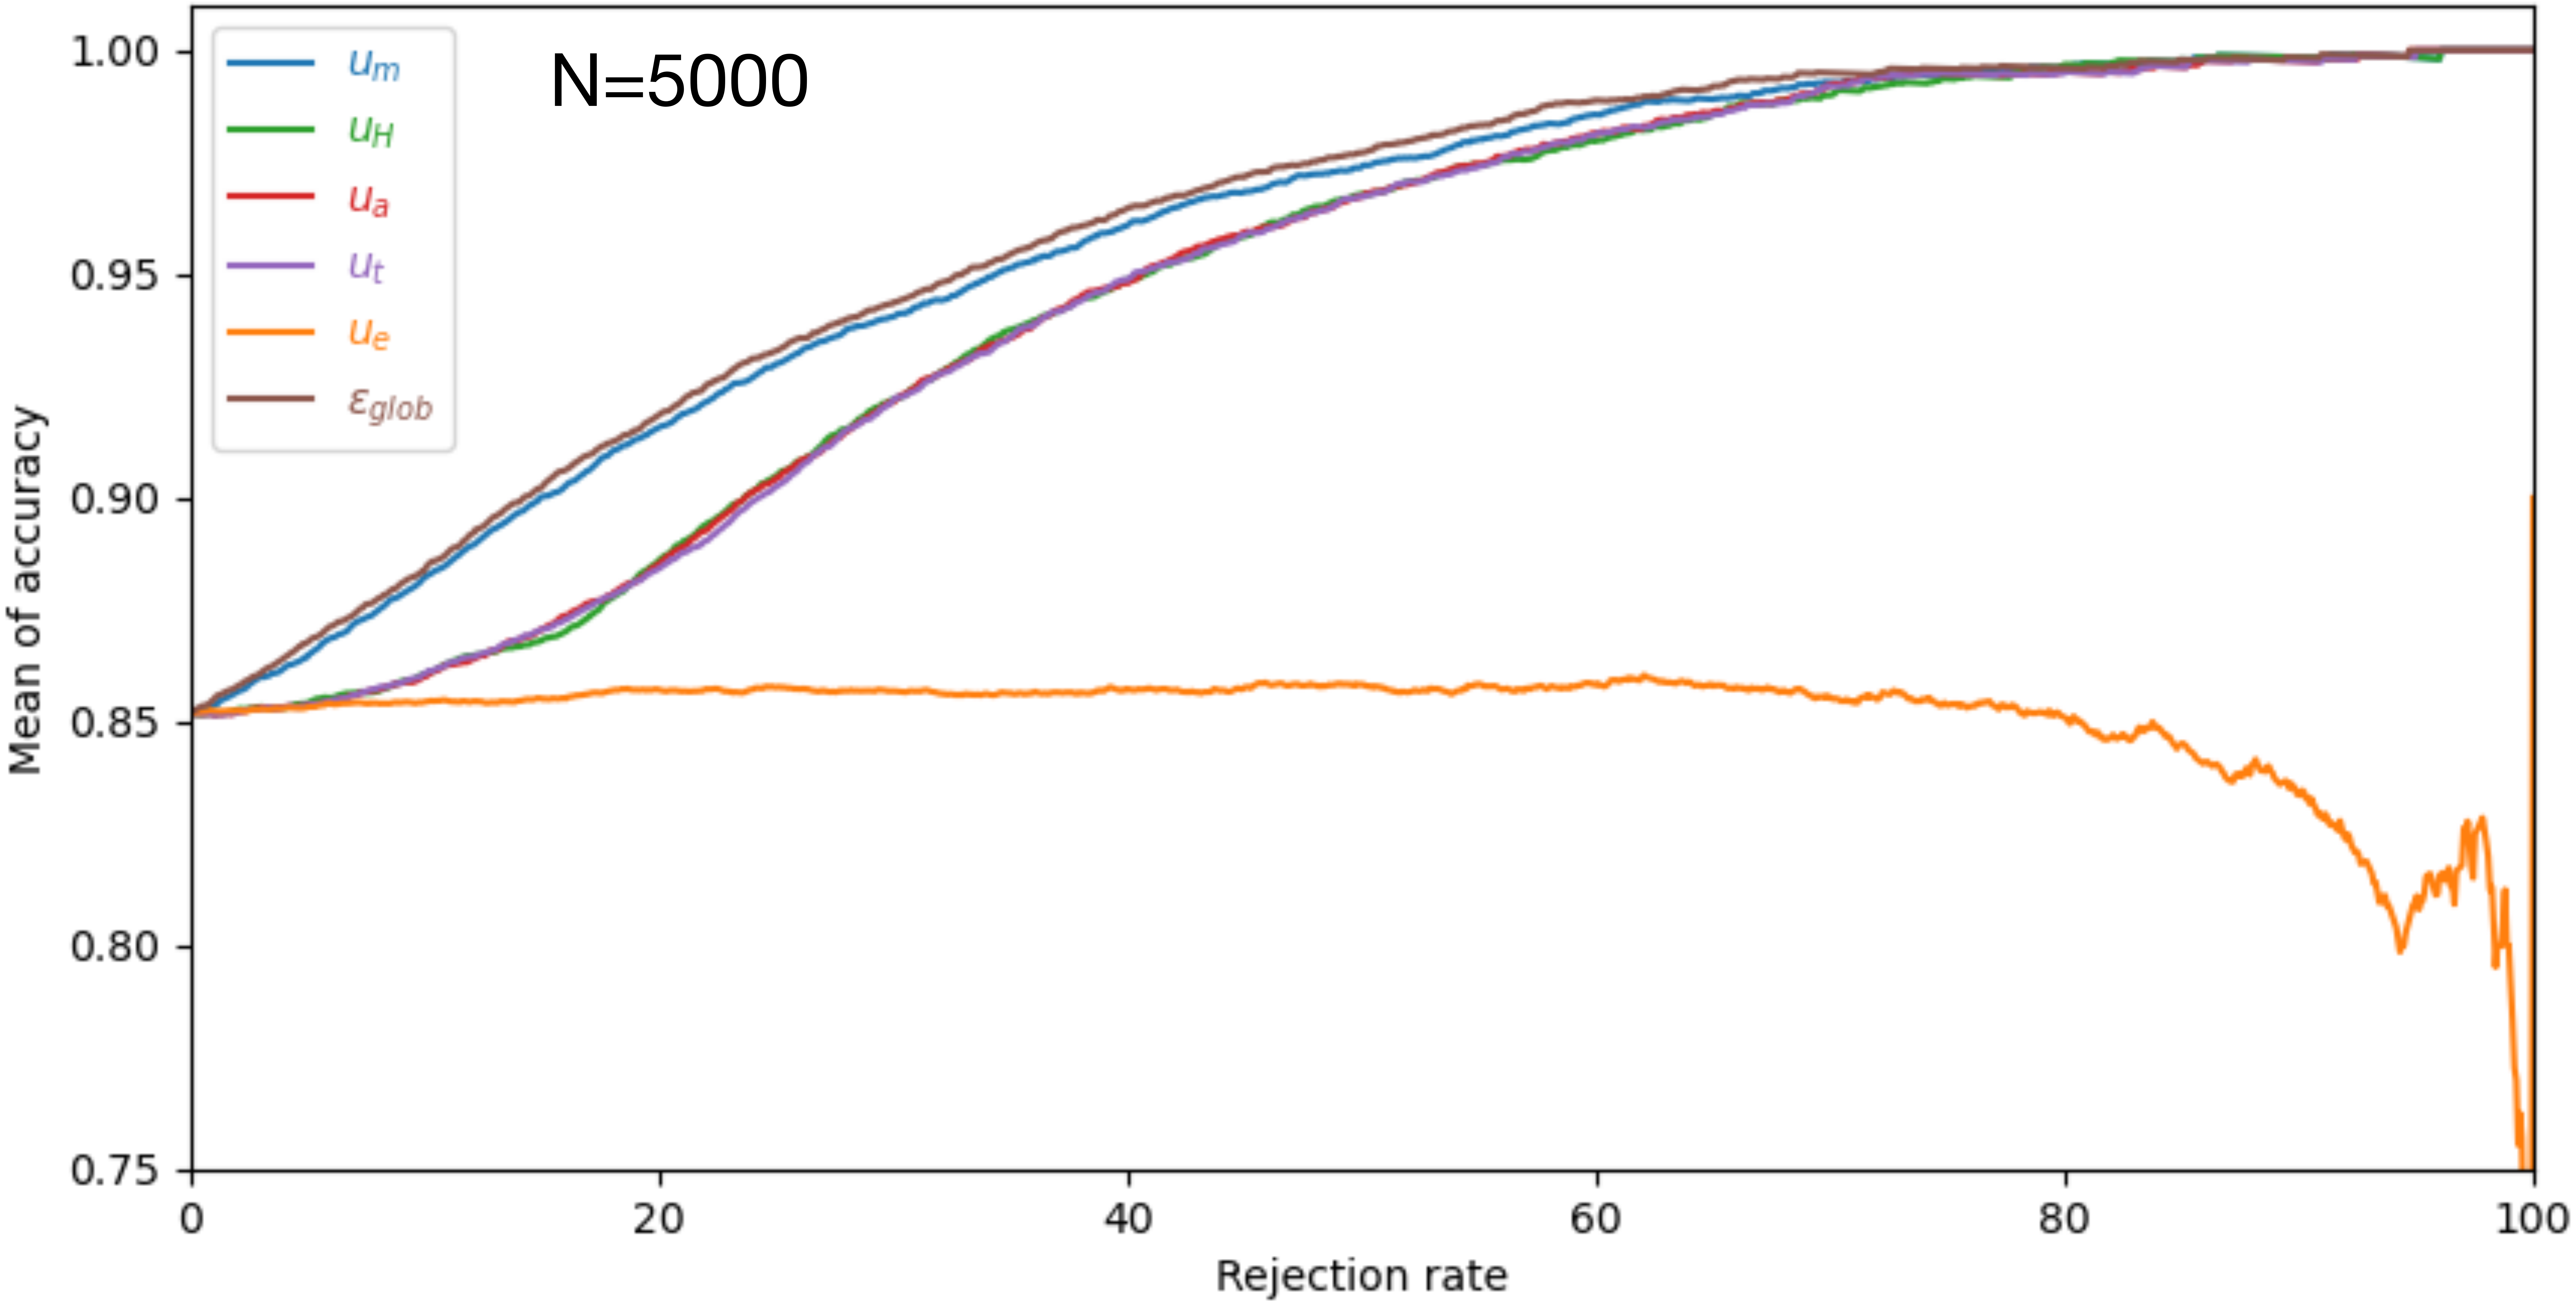
N=5000
u_m
u_H
u_a
u_t
u_e
ε_{glob}
Mean of accuracy
1.00
0.95
0.90
0.85
0.80
0.75
Rejection rate
0
20
40
60
80
100

RF

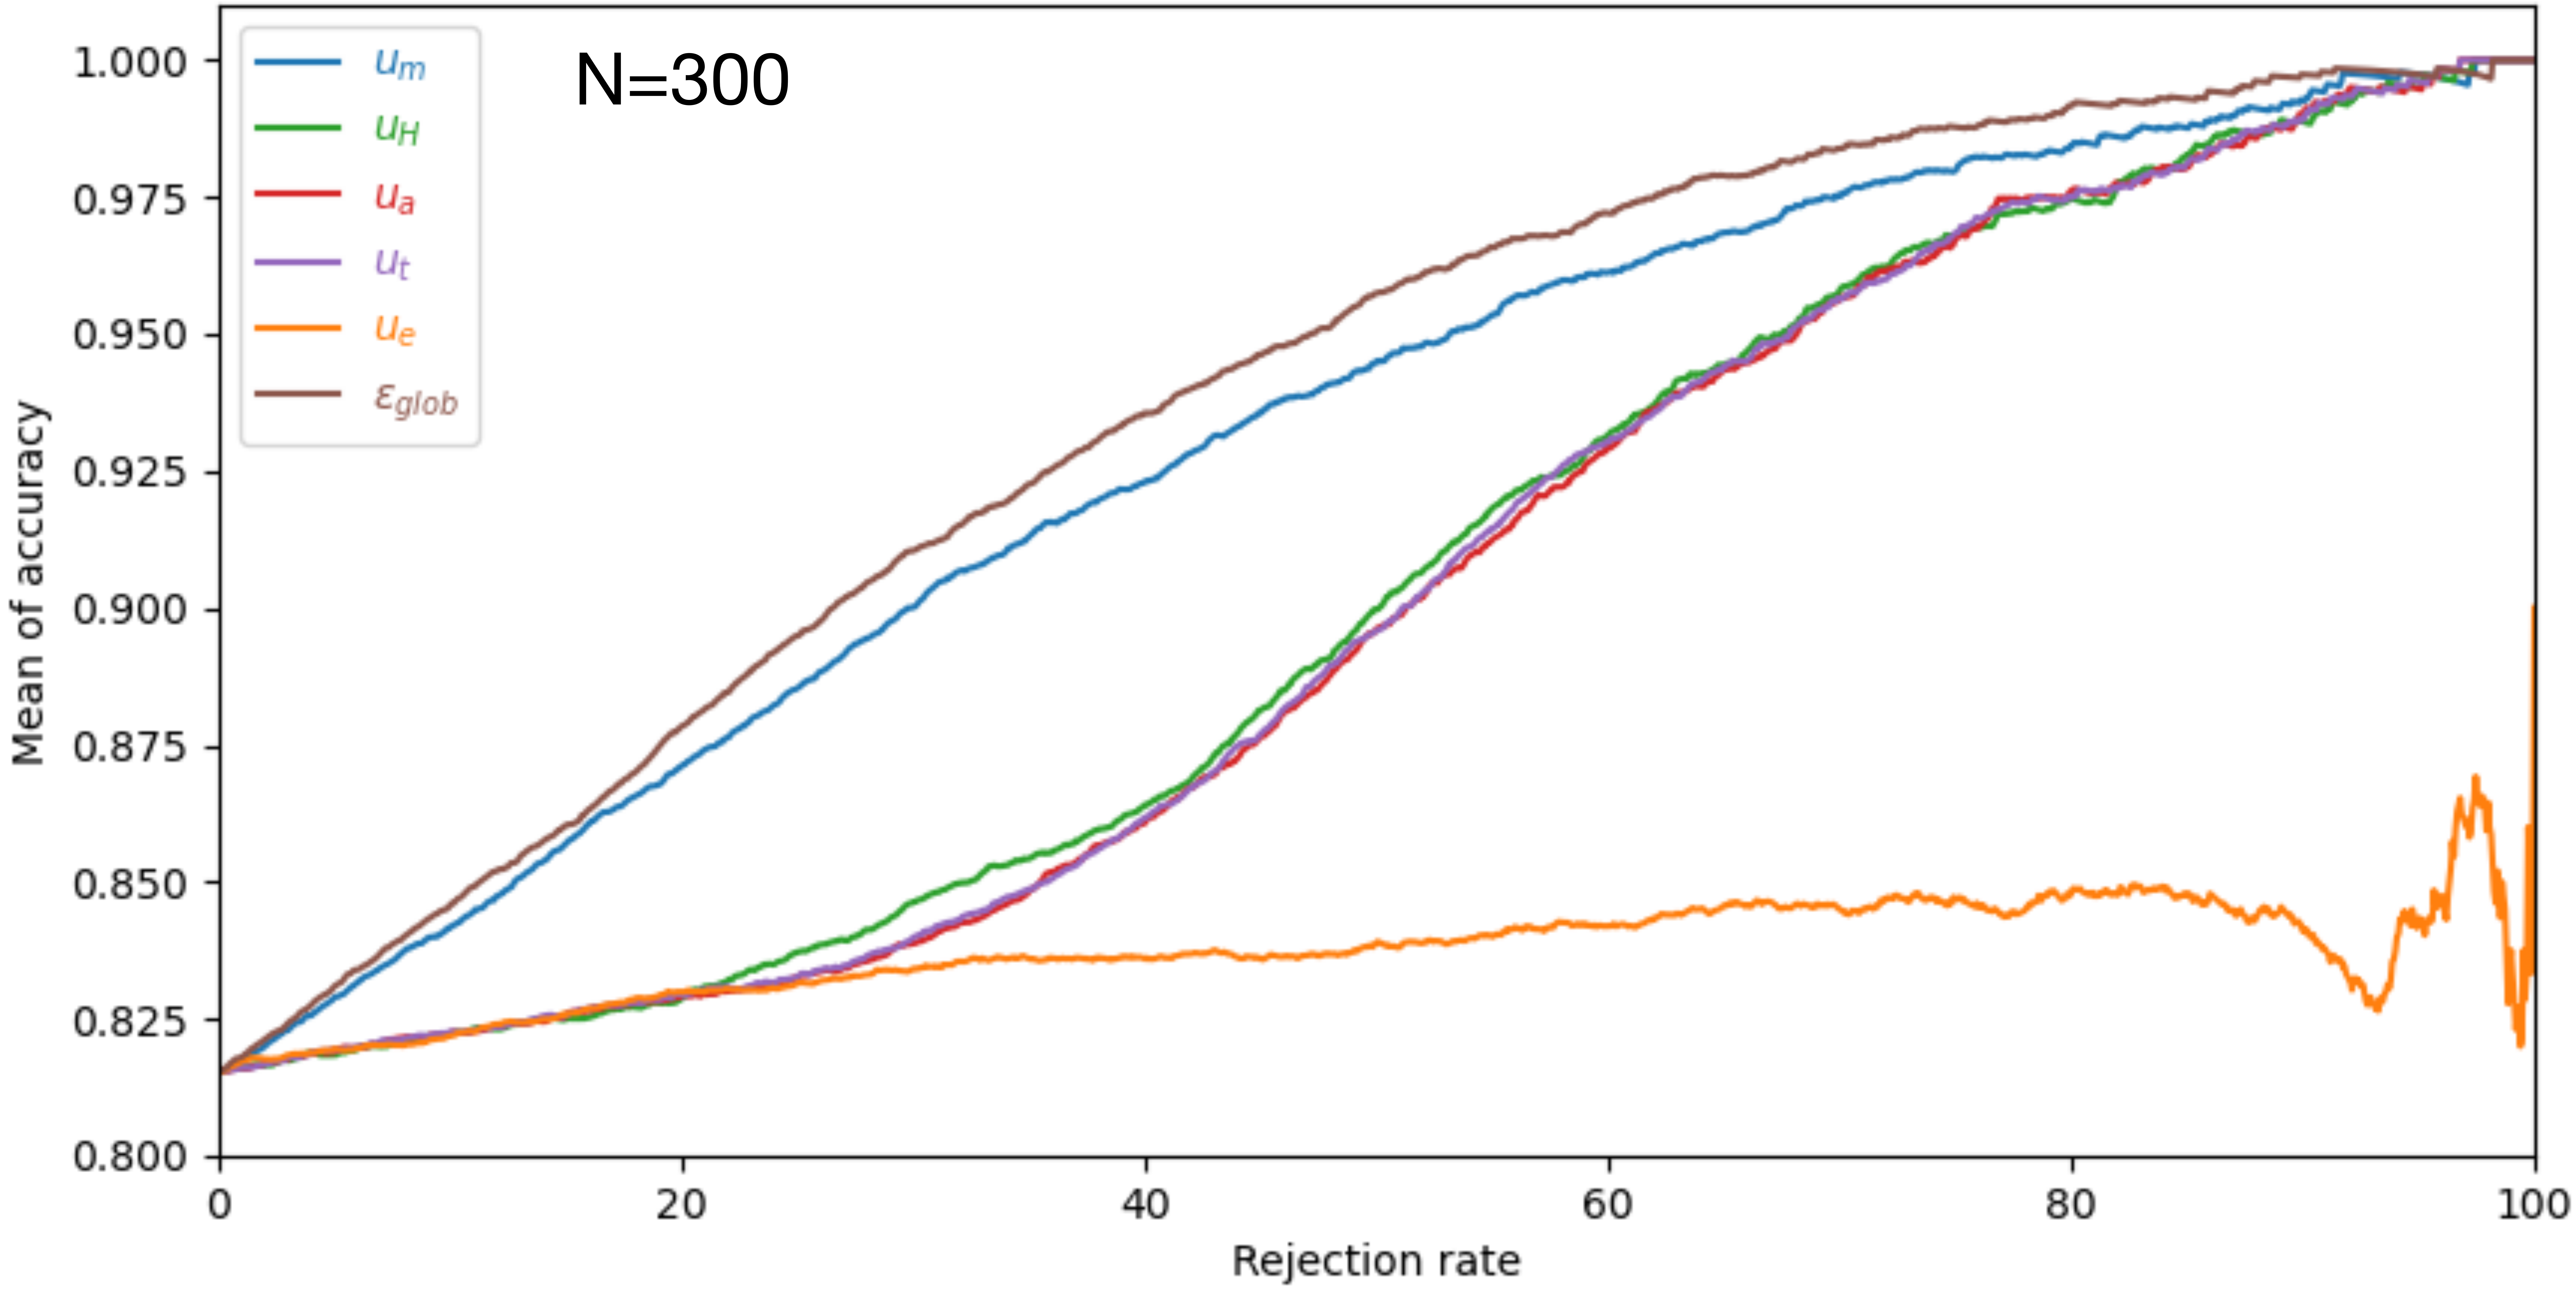
N=300
u_m
u_H
u_a
u_t
u_e
ε_{glob}
Mean of accuracy
1.000
0.975
0.950
0.925
0.900
0.875
0.850
0.825
0.800
0
20
40
60
80
100
Rejection rate

ROBUSTNESS QUANTIFICATION

- correlates nicely with accuracy ✓
- works for different types of model architectures ✓
- is conceptually different from UQ
- also works with global perturbations ✓
- is competitive with and complementary to UQ ✓
- is good with distribution shift and small data sets ✓
- is more stable than UQ
- is not limited to epsilon-contamination ✓
- also works for discriminative classifiers ✓

ROBUSTNESS QUANTIFICATION

- correlates nicely with accuracy ✓
- works for different types of model architectures ✓
- is conceptually different from UQ ✓
- also works with global perturbations ✓
- is competitive with and complementary to UQ ✓
- is good with distribution shift and small data sets ✓
- is more stable than UQ ✓
- is not limited to epsilon-contamination ✓
- also works for discriminative classifiers ✓


BUT....WAIT! THERES MORE

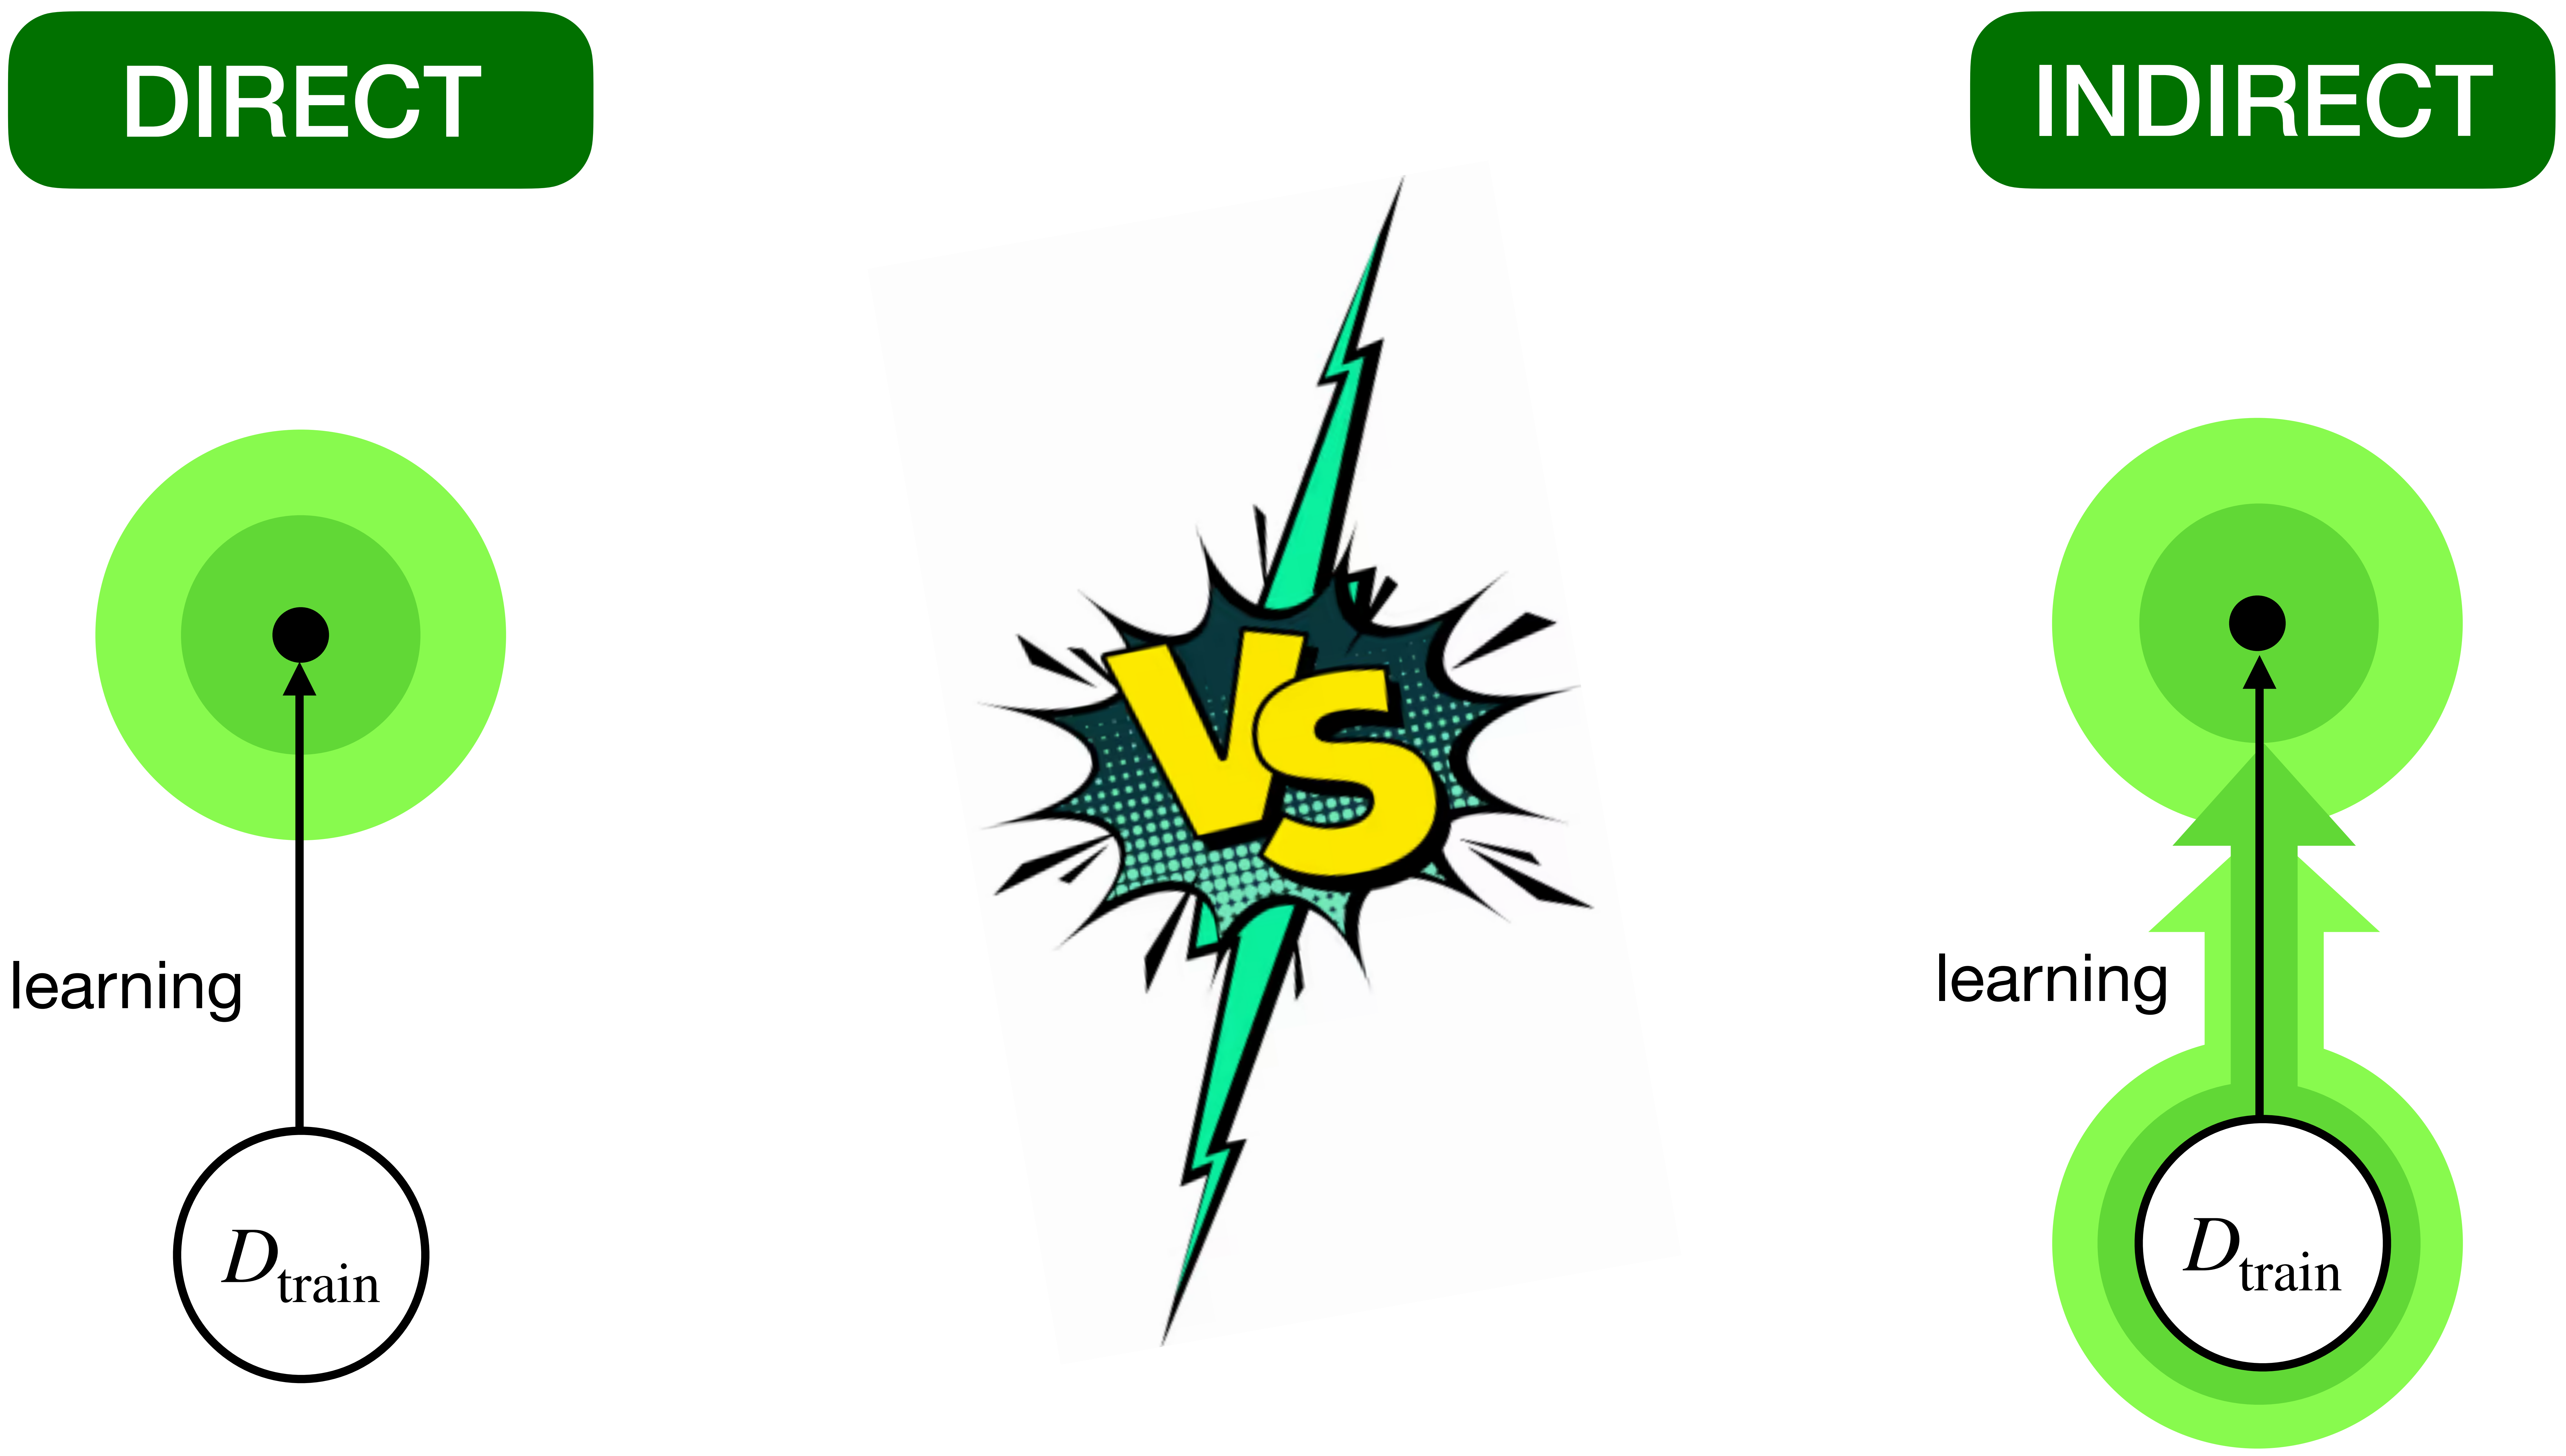
DIRECT
INDIRECT
VS
learning
D_{train}
learning
D_{train}

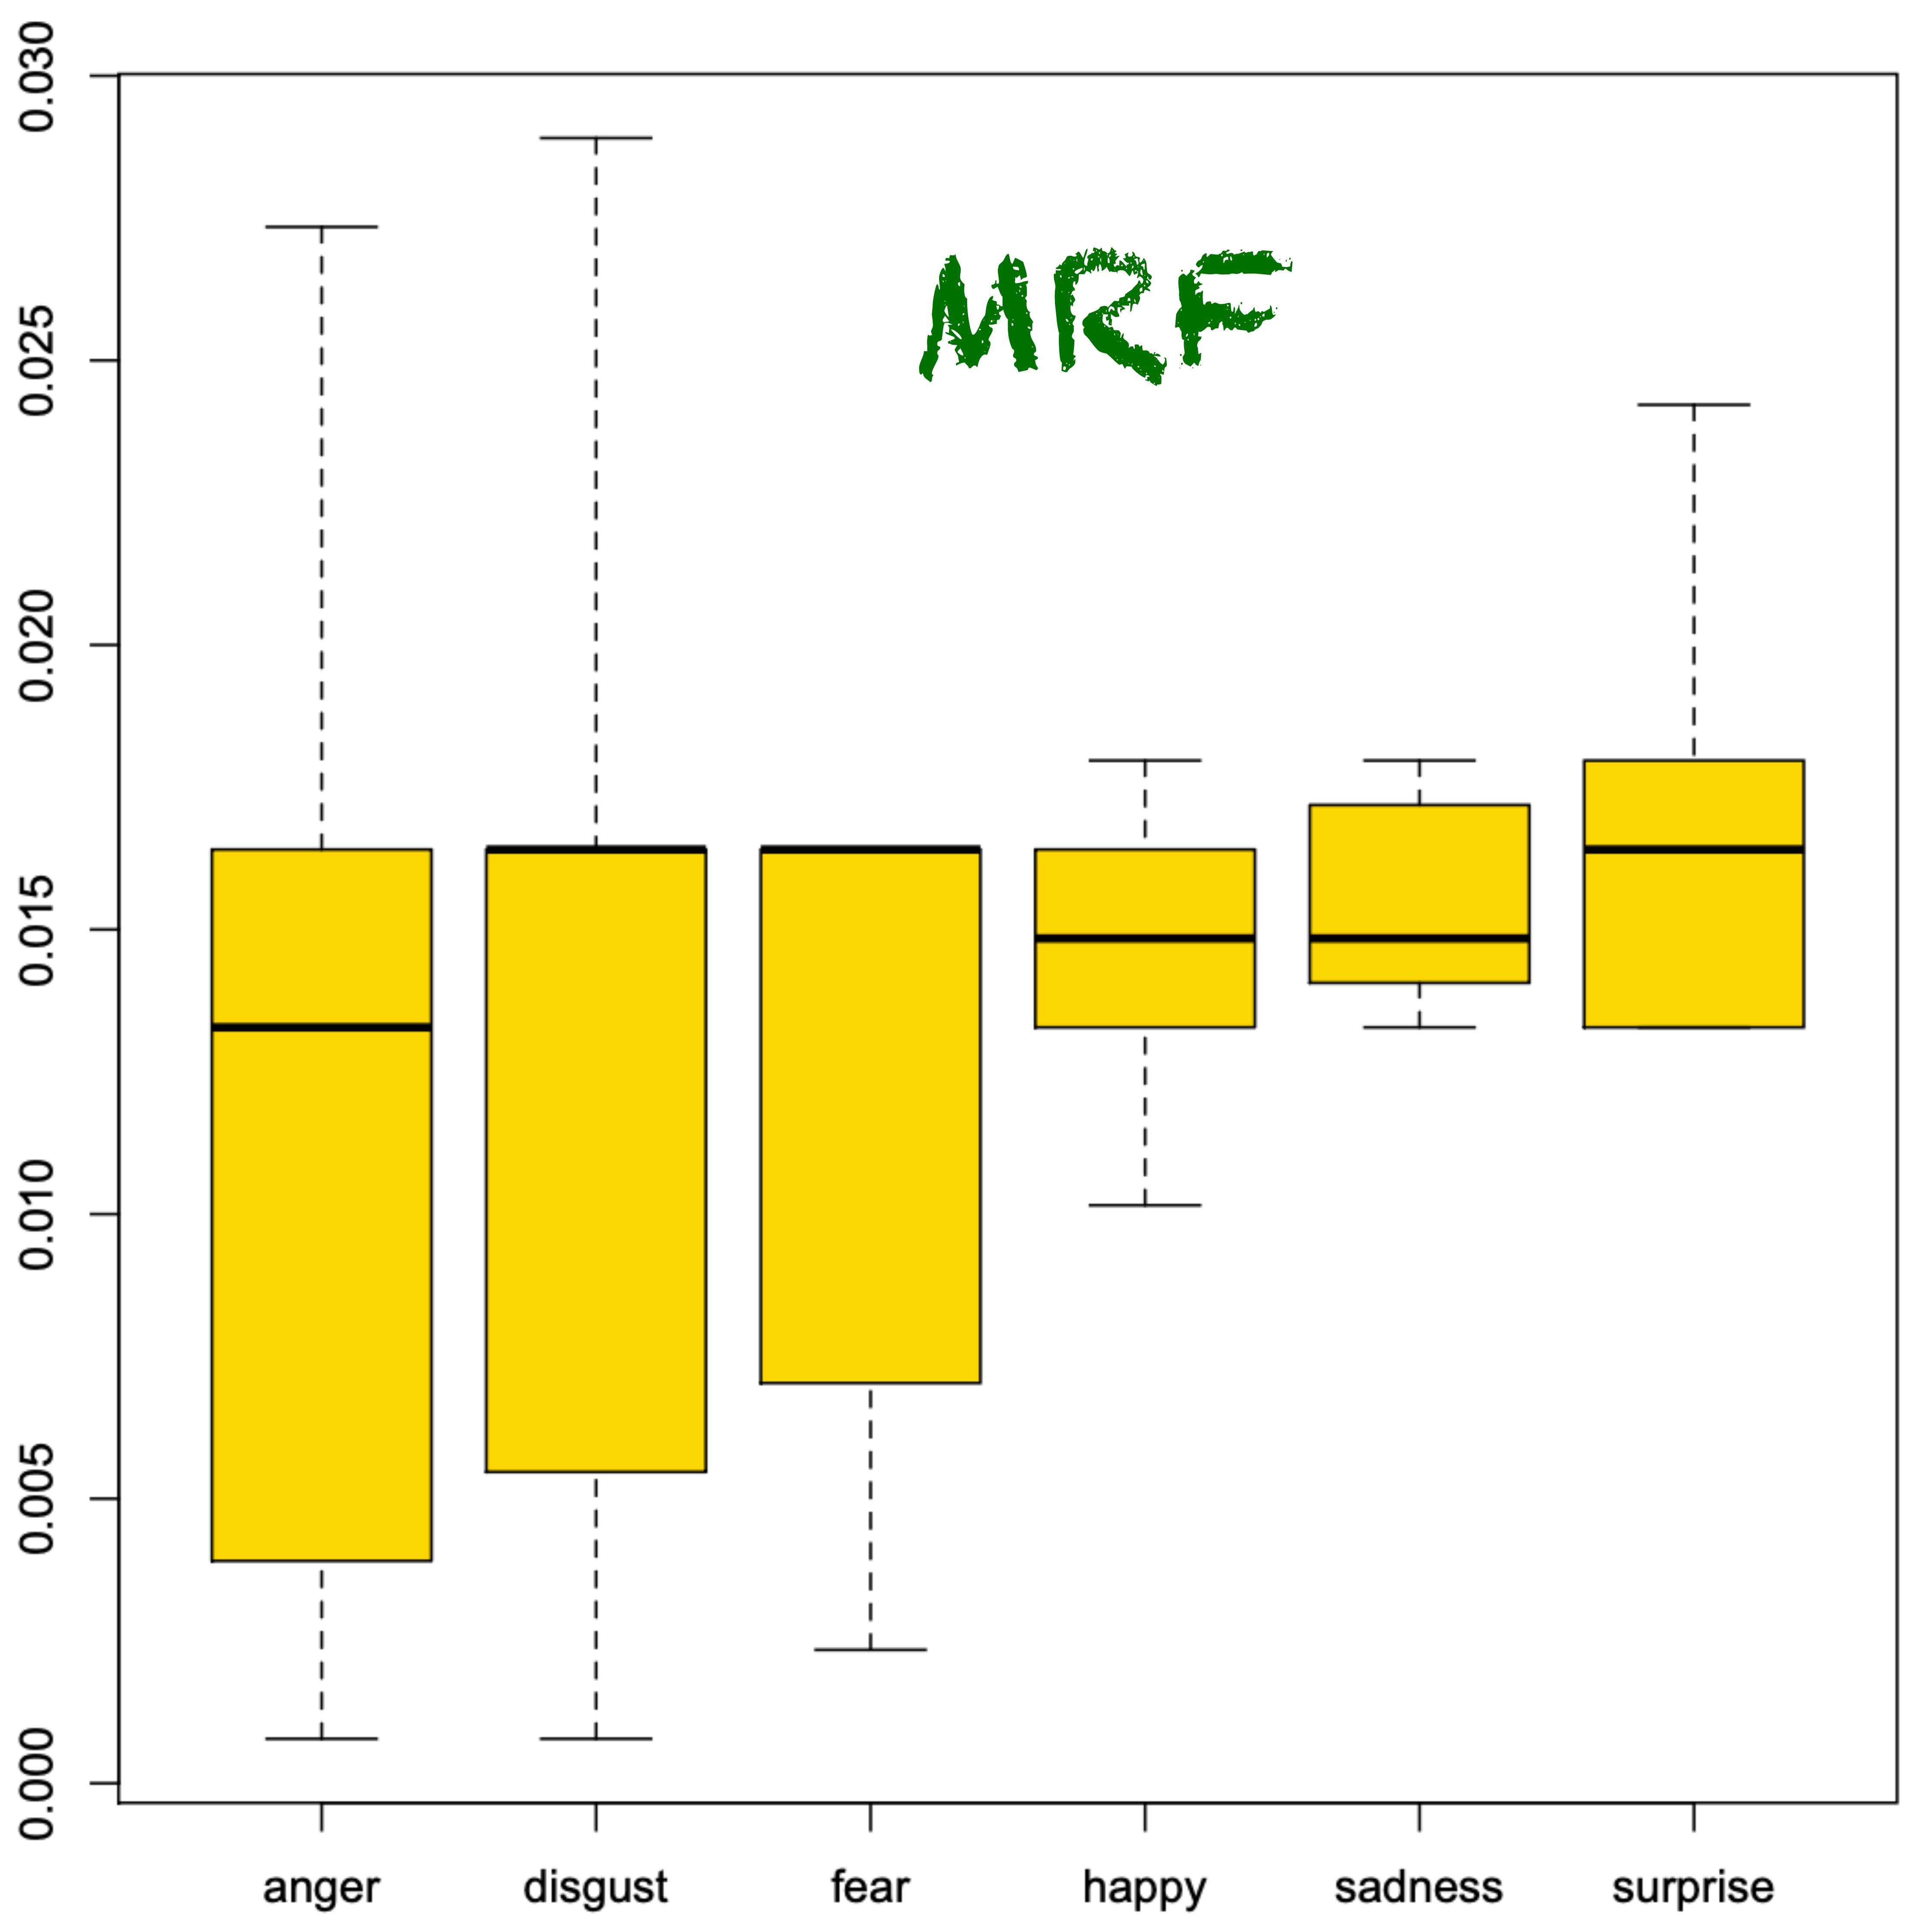

(a) Robustness split by emotions.

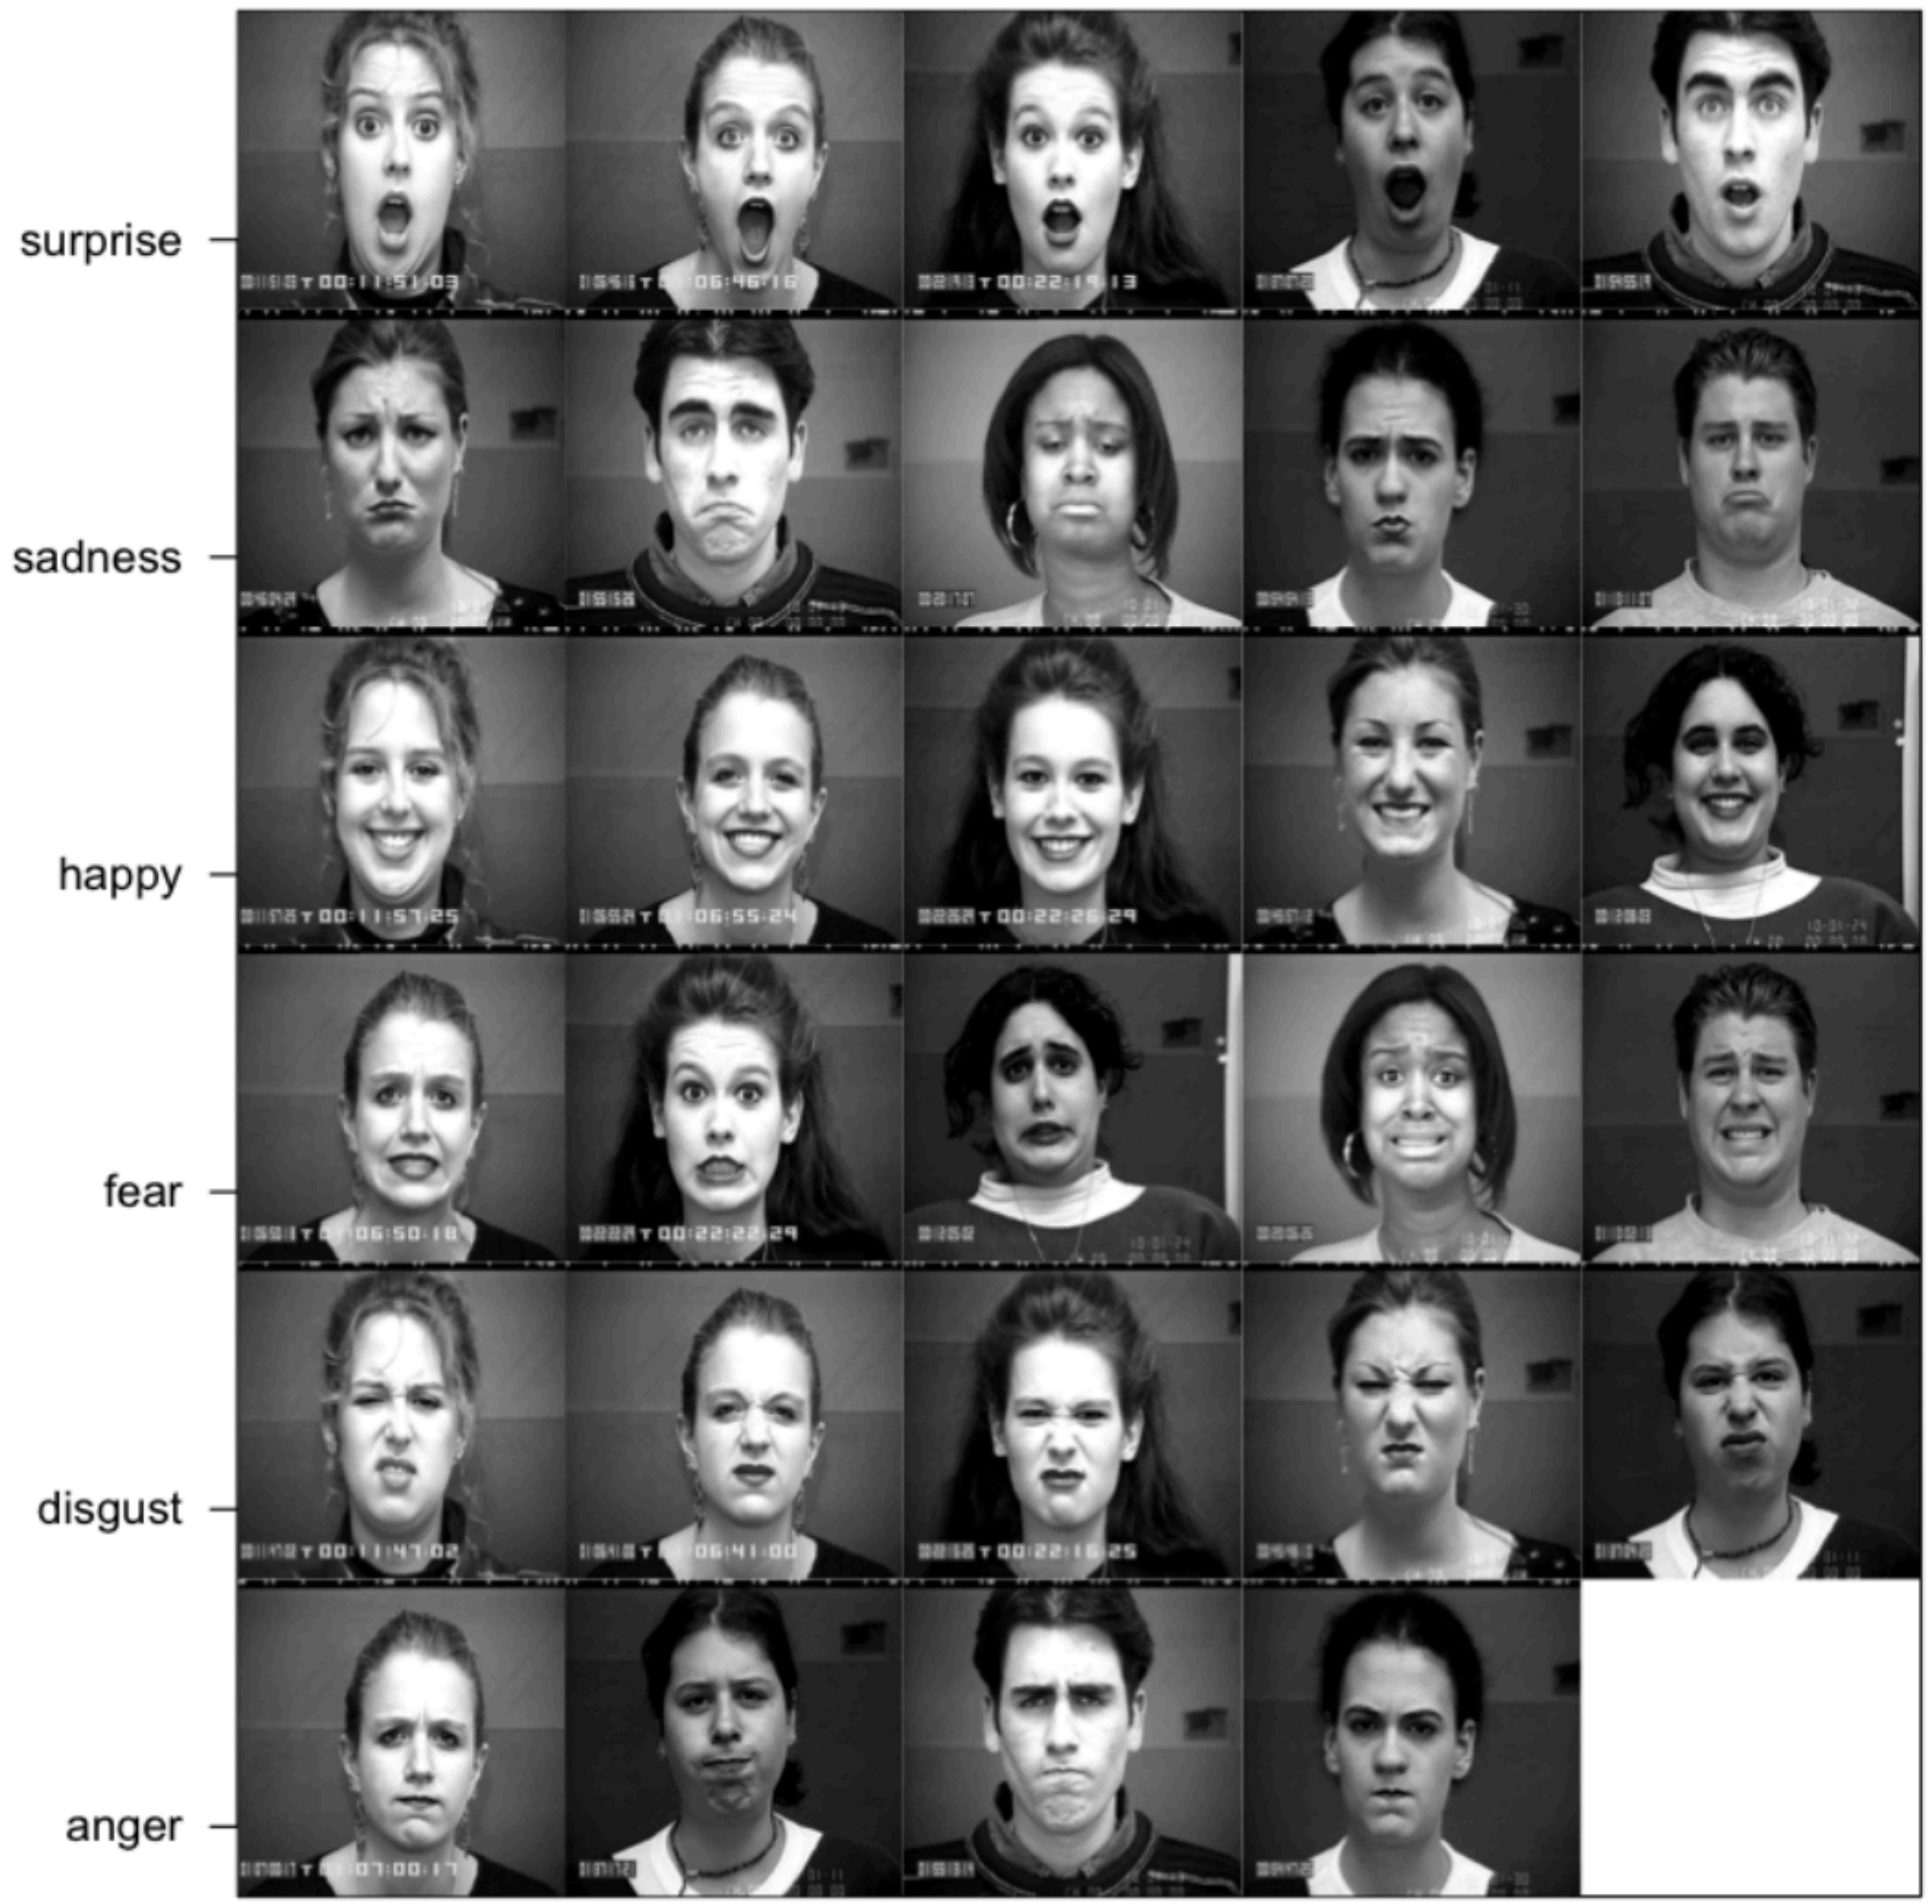

(b) Examples of emotions.

[1]

MODEL SELECTION

[8, 9, …]

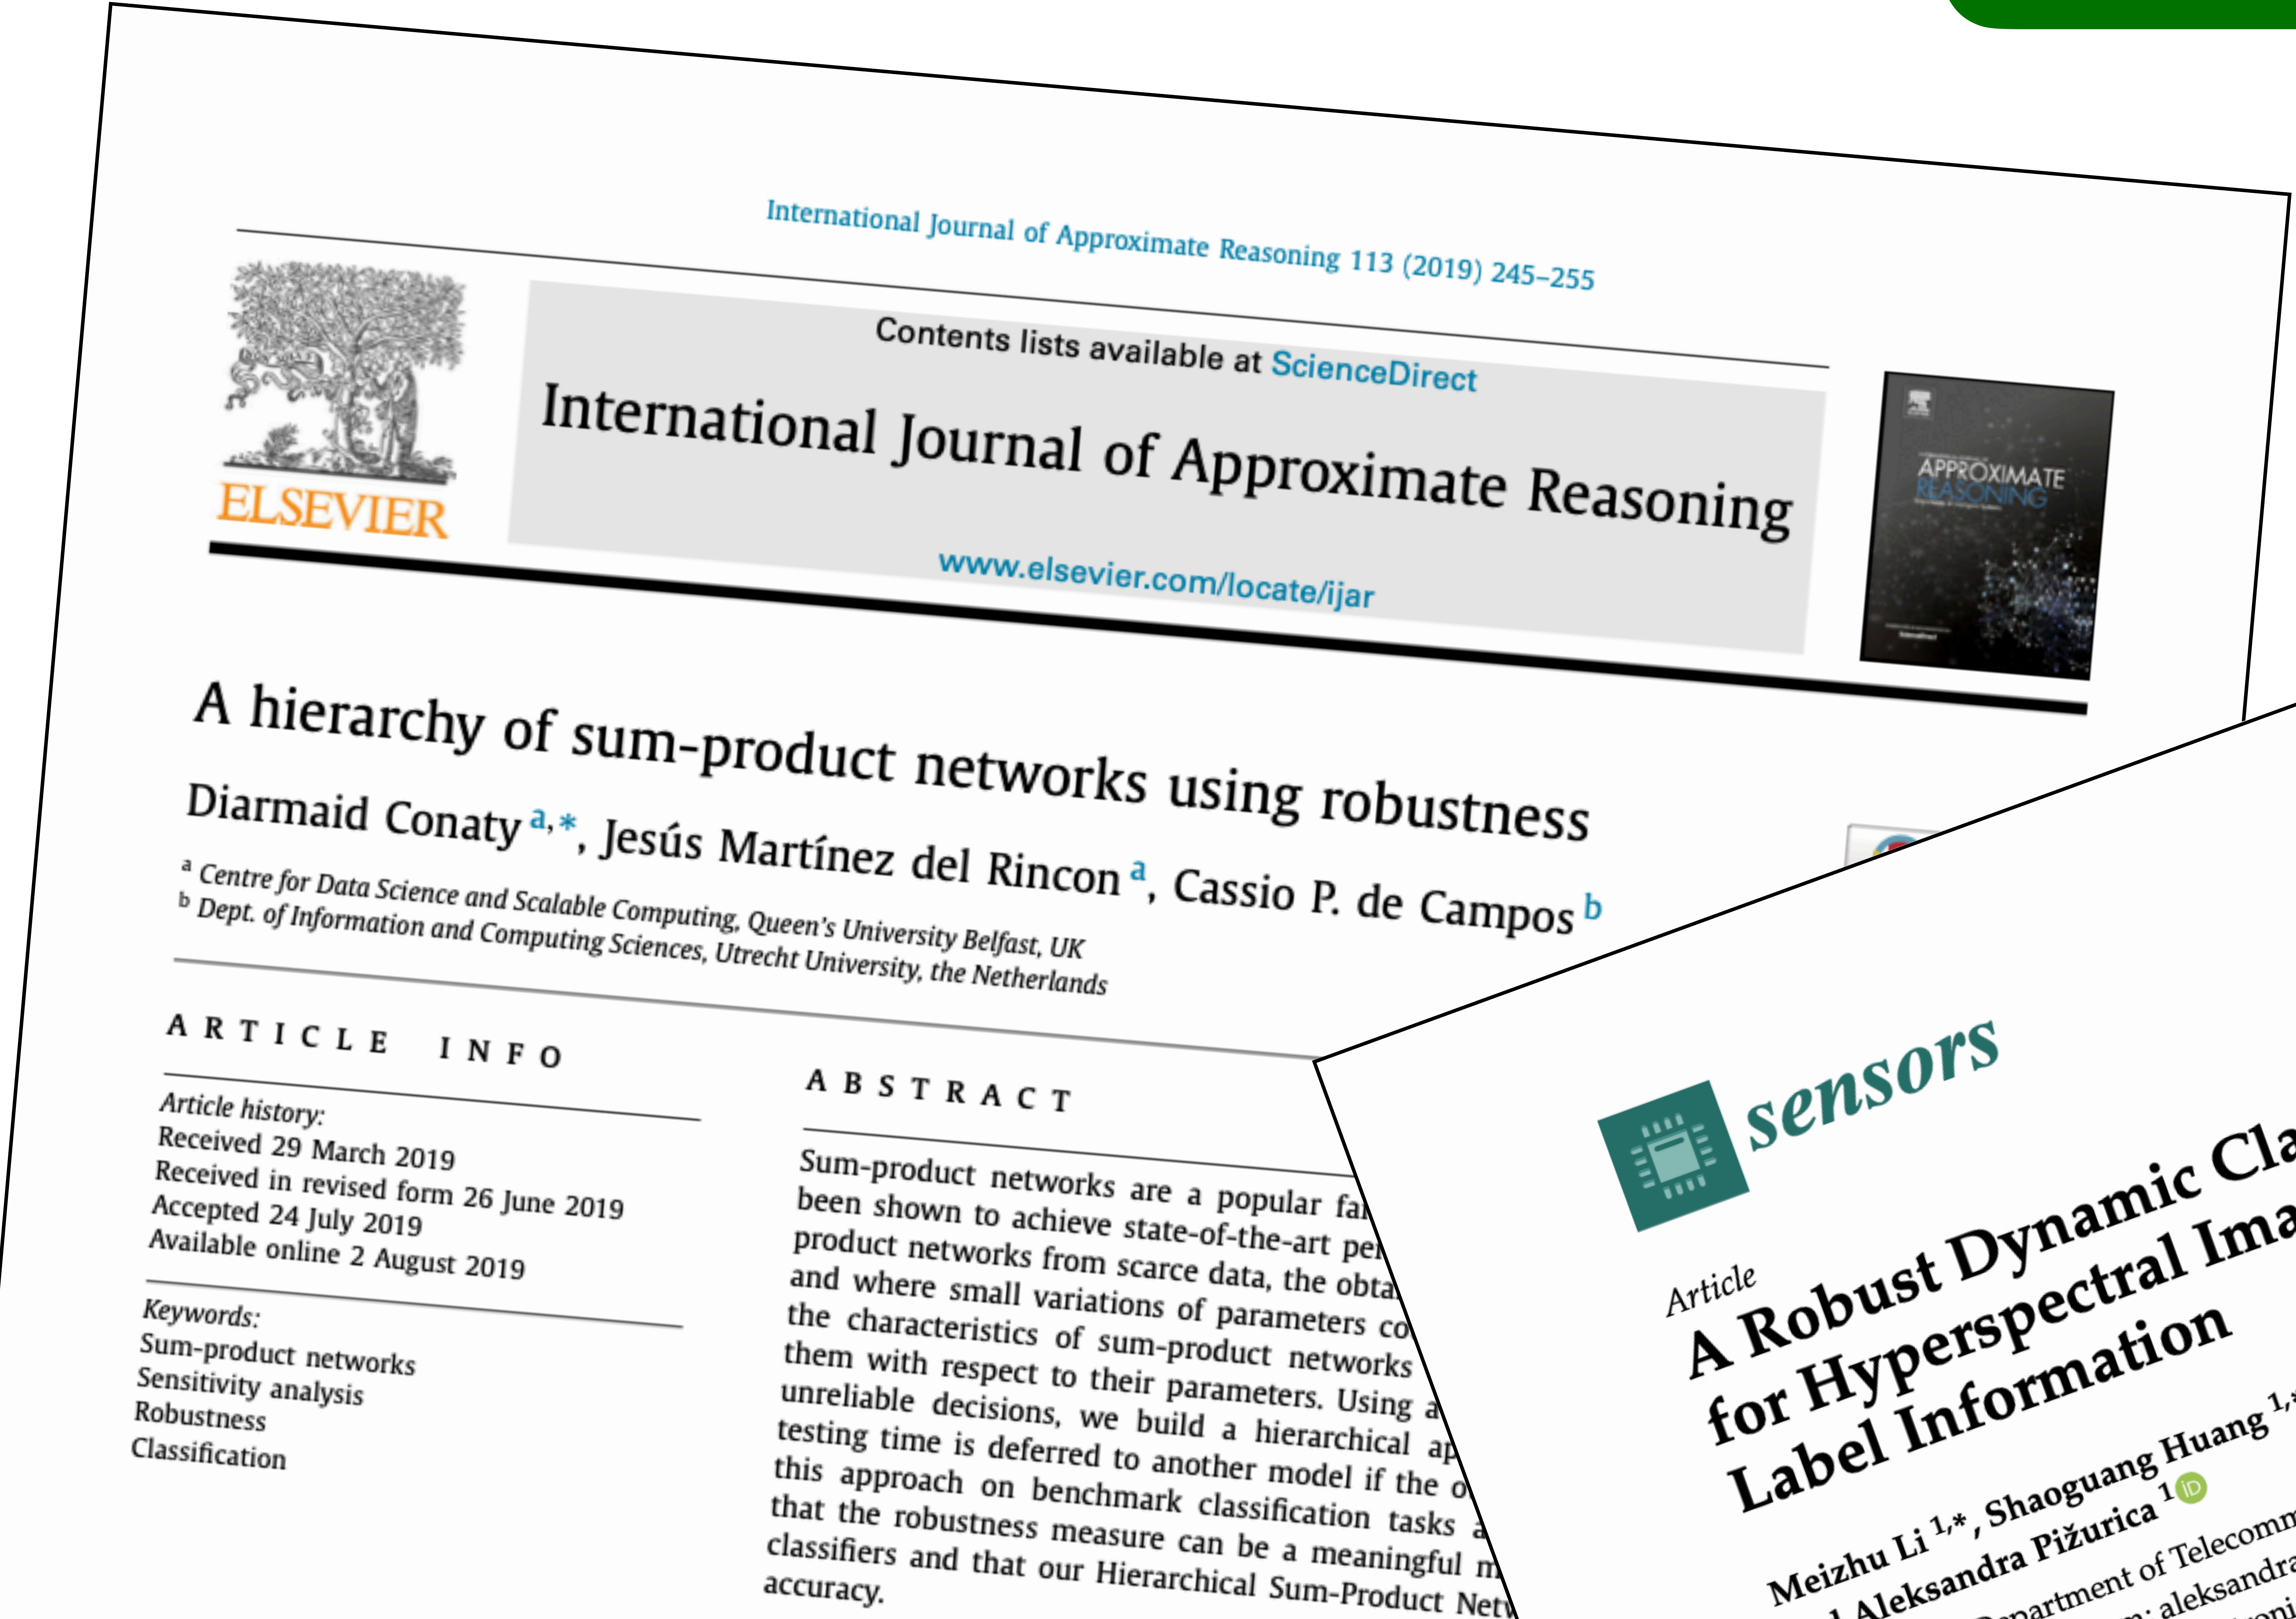

International Journal of Approximate Reasoning 113 (2019) 245–255

Contents lists available at ScienceDirect

International Journal of Approximate Reasoning

www.elsevier.com/locate/ijar

ELSEVIER

A hierarchy of sum-product networks using robustness

Diarmaid Conaty [a,*], Jesús Martínez del Rincon [a], Cassio P. de Campos [b]

[a] Centre for Data Science and Scalable Computing, Queen's University Belfast, UK
[b] Dept. of Information and Computing Sciences, Utrecht University, the Netherlands

ARTICLE INFO

Article history:
Received 29 March 2019
Received in revised form 26 June 2019
Accepted 24 July 2019
Available online 2 August 2019

Keywords:
Sum-product networks
Sensitivity analysis
Robustness
Classification

ABSTRACT

Sum-product networks are a popular fa
been shown to achieve state-of-the-art pe
product networks from scarce data, the obta
and where small variations of parameters co
the characteristics of sum-product networks
them with respect to their parameters. Using a
unreliable decisions, we build a hierarchical ap
testing time is deferred to another model if the o
this approach on benchmark classification tasks a
that the robustness measure can be a meaningful m
classifiers and that our Hierarchical Sum-Product Netw
accuracy.

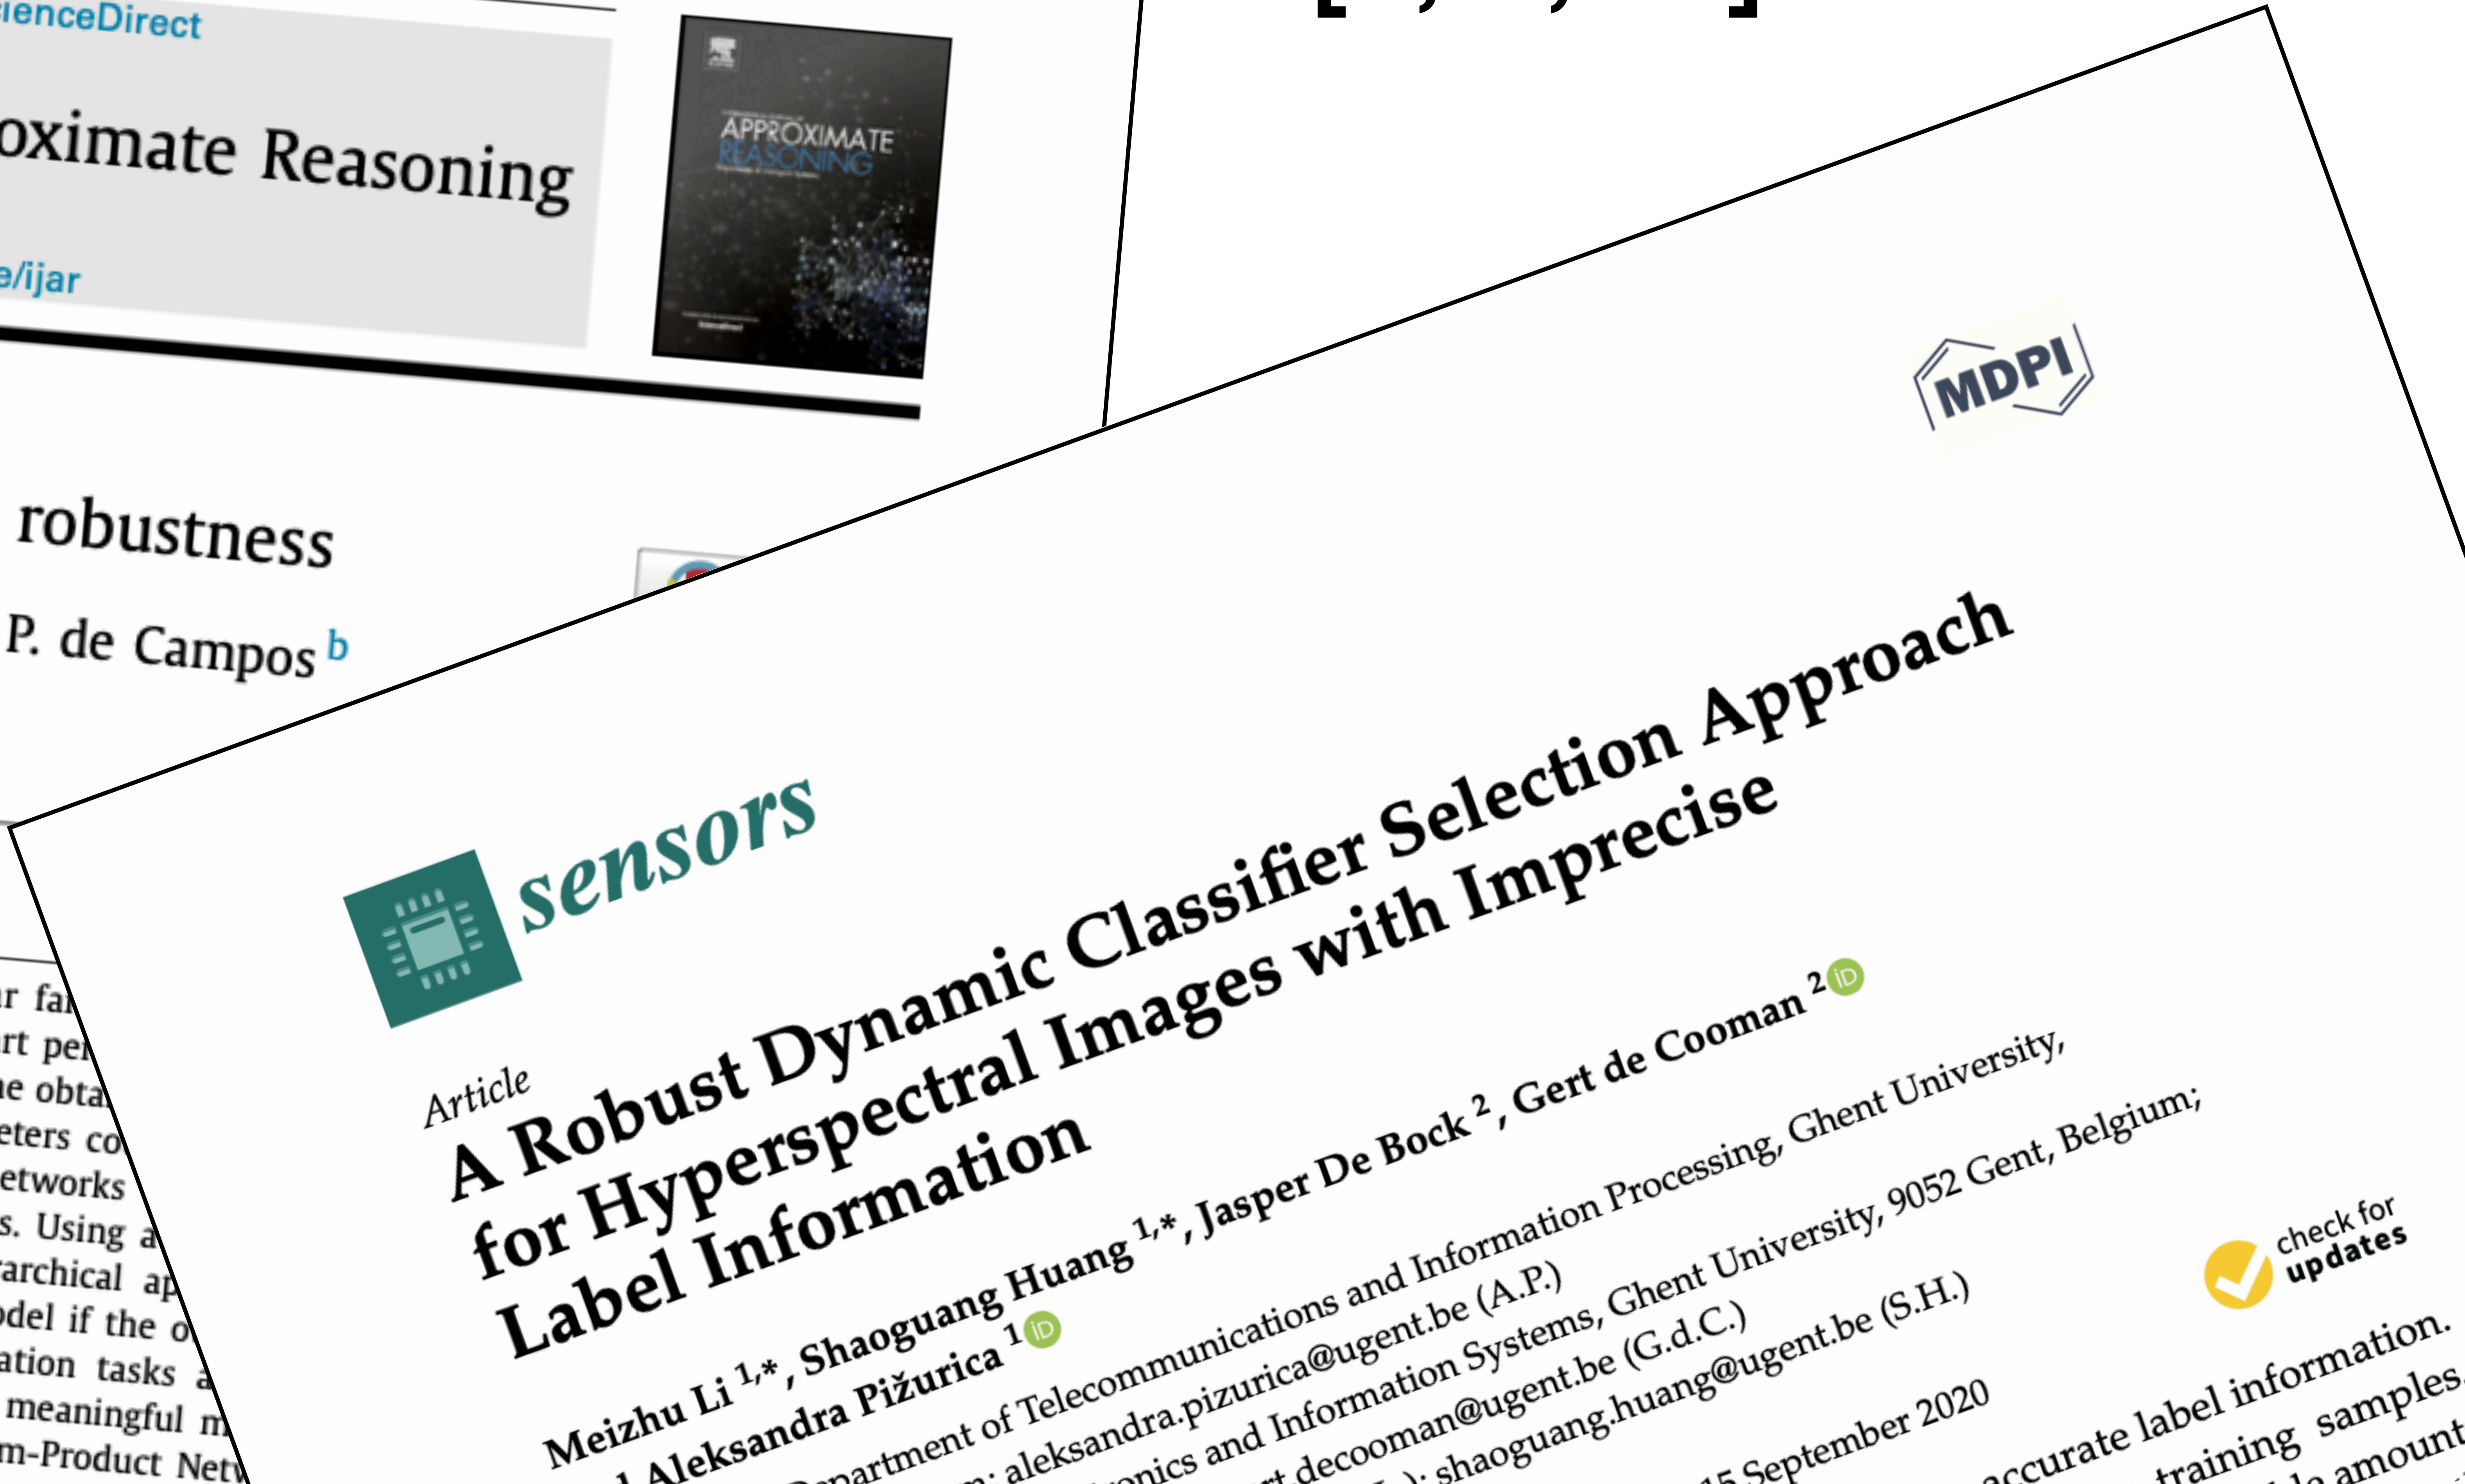

sensors

MDPI

Article

A Robust Dynamic Classifier Selection Approach for Hyperspectral Images with Imprecise Label Information

Meizhu Li [1,*], Shaoguang Huang [1,*], Jasper De Bock [2], Gert de Cooman [2]
Aleksandra Pižurica [1]

epartment of Telecommunications and Information Processing, Ghent University,
: aleksandra.pizurica@ugent.be (A.P.)
onics and Information Systems, Ghent University, 9052 Gent, Belgium;
rt.decooman@ugent.be (G.d.C.)
); shaoguang.huang@ugent.be (S.H.)

15 September 2020

check for updates

accurate label information.
training samples
amount


FUTURE WORK MEETINGS


CALIBRATION
FUTURE WORK MEETINGS


REGRESSION
FUTURE WORK MEETINGS


CONFORMAL PREDICTION
FUTURE WORK MEETINGS

jasper.debock@ugent.be

[1] Global Sensitivity Analysis for MAP Inference in Graphical Models. De Bock, de Campos & Antonucci. 2014.

[2] Credal sum-product networks.
Mauá, Cozman, Conaty & de Campos. 2017.

[3] Robustifying sum-product networks.
Mauá, Conaty, Cozman, Poppenhaeger & de Campos. 2018.

[4] Towards Scalable and Robust Sum-Product Networks.
Correia & de Campos. 2019.

[5] Towards Robust Classification with Deep Generative Forests.
Correia, Peharz & de Campos. 2020

[6] Robustness quantification: a new method for assessing the reliability of the predictions of a classifier.
De Tavernier, De Bock. 2025.

[7] Robustness and uncertainty: two complementary aspects of the reliability of the predictions of a classifier.
De Tavernier, De Bock. 2025.

[8] A hierarchy of sum-product networks using robustness.
Conaty, Martínez del Rincon & de Campos. 2019.

[9] A Robust Dynamic Classifier Selection Approach for Hyperspectral Images with Imprecise Label Information.
Li, Huang, De Bock & Pižurica. 2020.