# Background Check: A general technique to build more reliable and versatile classifiers

Miquel Perello-Nieto*[1], Telmo M. Silva Filho*[2], Meelis Kull[1] and Peter Flach[1]

* Equal contribution
[1]Department of Computer Science, University of Bristol, UK
[2]Centro de Informatica, Universidade Federal de Pernambuco, Brazil

Email: [1]{Miquel.PerelloNieto, Meelis.Kull, Peter.Flach}@bristol.ac.uk, [2]tmsf@cin.ufpe.br

December 13, 2016

# Representing uncertainty

1. Cautious classification
2. Outlier detection
3. Classification with confidence

# Representing uncertainty

1. Cautious classification
2. Outlier detection
3. Classification with confidence

| | $p(C_1\|x)$ | $p(C_2\|x)$ |
|---|---|---|
| A | | |
| B | | |
| C | | |
| D | | |
| E | | |

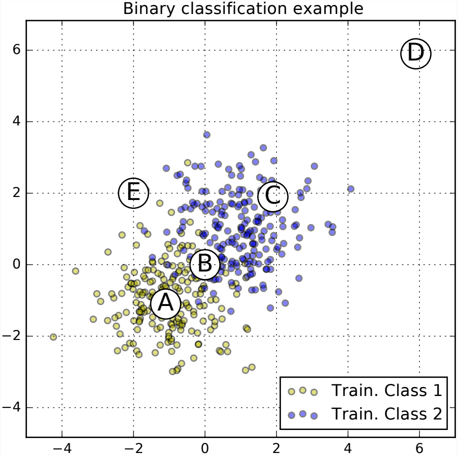

# Representing uncertainty

1. Cautious classification
2. Outlier detection
3. Classification with confidence

| | $p(C_1\|x)$ | $p(C_2\|x)$ |
|---|---|---|
| A | 1 | .0 |
| B | *.5* | *.5* |
| C | .0 | 1 |
| D | *.5* | *.5* |
| E | *.5* | *.5* |

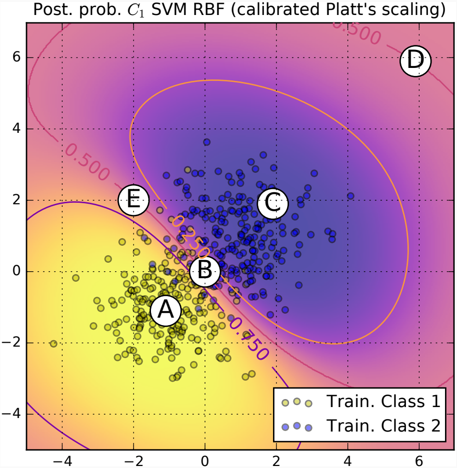

# Representing uncertainty

1. Cautious classification
2. Outlier detection
3. Classification with confidence

| | $p(C_1\|x)$ | $p(C_2\|x)$ | $p(b\|x)$ |
|---|---|---|---|
| A | 1 | .0 | |
| B | .5 | .5 | |
| C | .0 | 1 | |
| D | .5 | .5 | |
| E | .5 | .5 | |

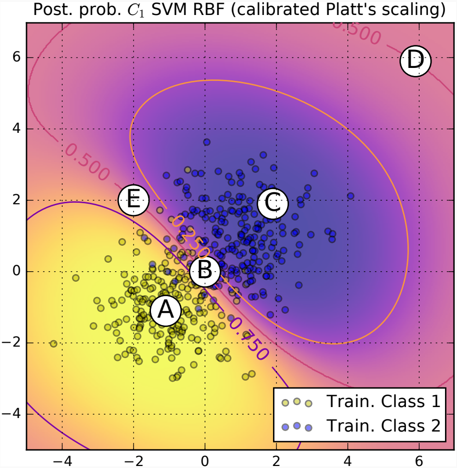

# Representing uncertainty

1. Cautious classification
2. Outlier detection
3. Classification with confidence

| | $p(C_1\|x)$ | $p(C_2\|x)$ | $p(b\|x)$ |
|---|---|---|---|
| A | $1 \rightarrow .9$ | $.0 \rightarrow .0$ | .1 |
| B | $.5 \rightarrow .5$ | $.5 \rightarrow .5$ | .0 |
| C | $.0 \rightarrow .0$ | $1 \rightarrow .5$ | .5 |
| D | $.5 \rightarrow .0$ | $.5 \rightarrow .0$ | 1 |
| E | $.5 \rightarrow .1$ | $.5 \rightarrow .1$ | .8 |

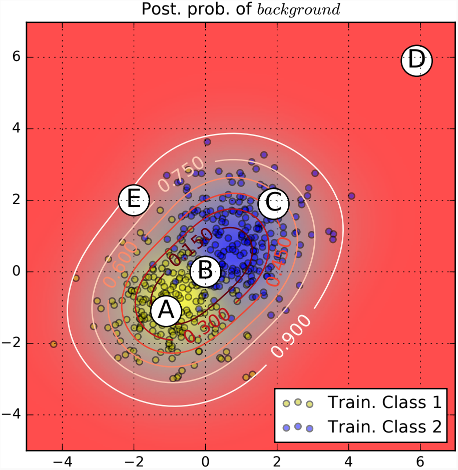

# Performing Background Check

- *Discriminative* approach
  - Pre-trained classifier
  - Generate background
  - Train binary classifier

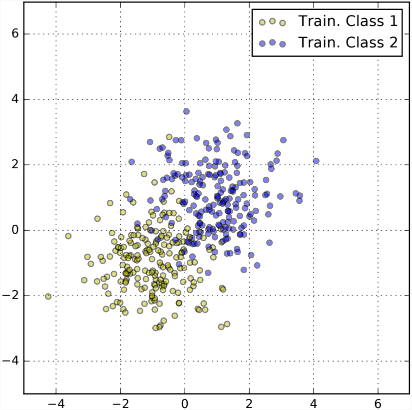

# Performing Background Check

- *Discriminative* approach
  - Pre-trained classifier
  - Generate background
  - Train binary classifier

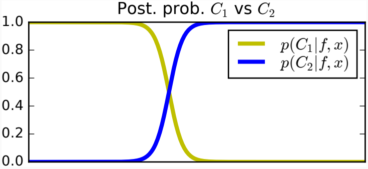


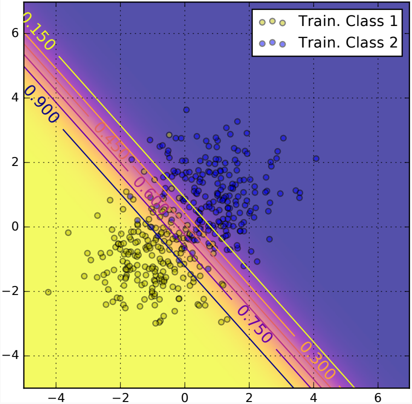

# Performing Background Check

- *Discriminative* approach
  - Pre-trained classifier
  - Generate background
  - Train binary classifier

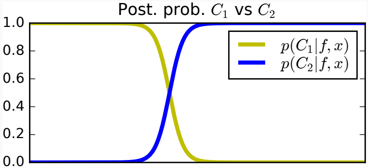


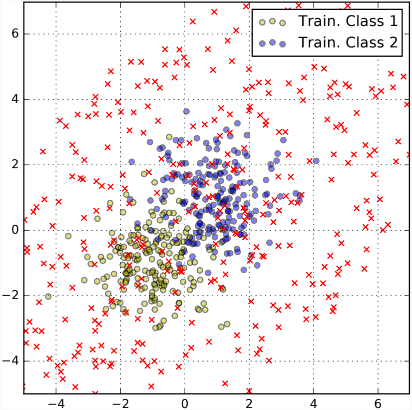

# Performing Background Check

- *Discriminative* approach
  - Pre-trained classifier
  - Generate background
  - Train binary classifier

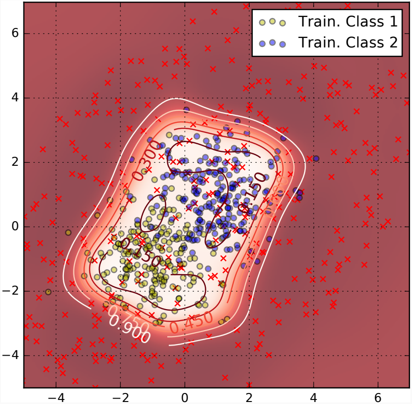


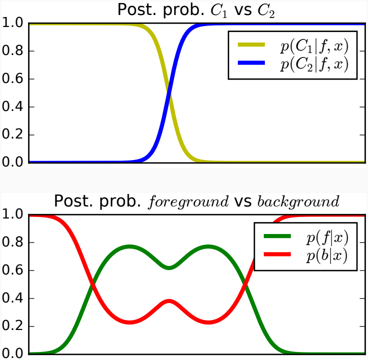

# Performing Background Check

- *Discriminative* approach
  - Pre-trained classifier
  - Generate background
  - Train binary classifier

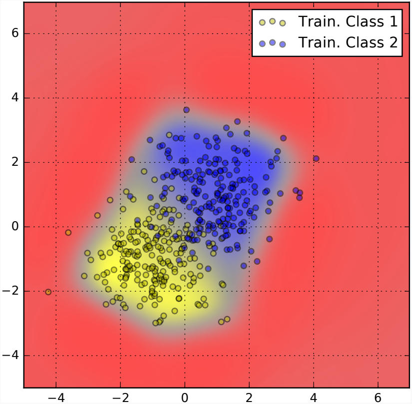


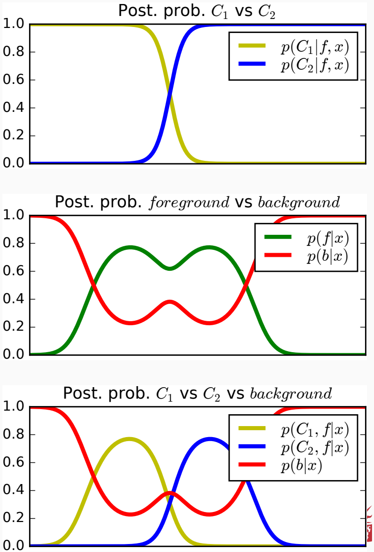


University of BRISTOL

# Performing Background Check

- *Familiarity* approach
  - Pre-trained classifier
  - Learn $q_f(x) \in [0, 1]$
  - Use inductive bias

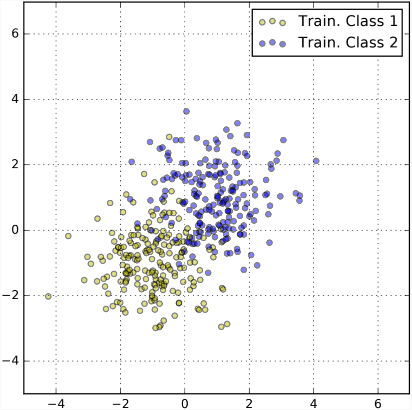

# Performing Background Check

- *Familiarity* approach
  - Pre-trained classifier
  - Learn $q_f(x) \in [0, 1]$
  - Use inductive bias

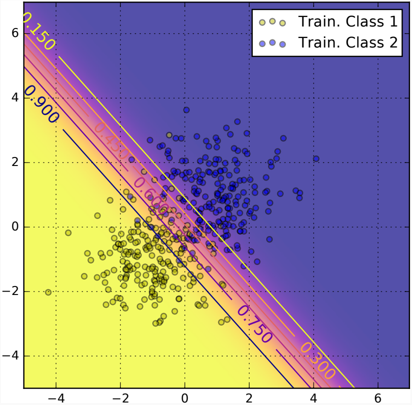


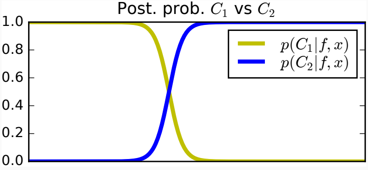

# Performing Background Check

- *Familiarity* approach
  - Pre-trained classifier
  - Learn $q_f(x) \in [0, 1]$
  - Use inductive bias

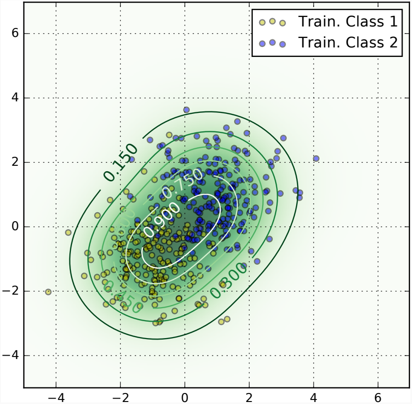


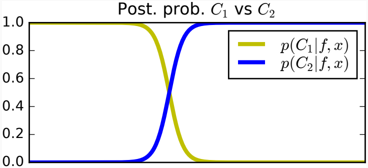


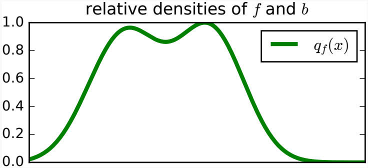

# Performing Background Check

- *Familiarity* approach
  - Pre-trained classifier
  - Learn $q_f(x) \in [0, 1]$
  - Use inductive bias

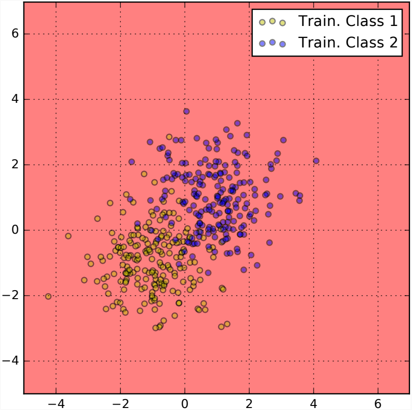


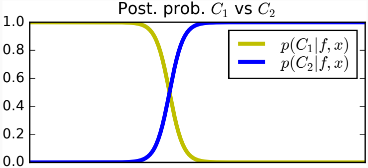


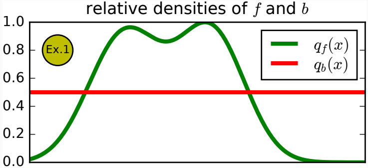

# Performing Background Check

- *Familiarity* approach
  - Pre-trained classifier
  - Learn $q_f(x) \in [0, 1]$
  - Use inductive bias

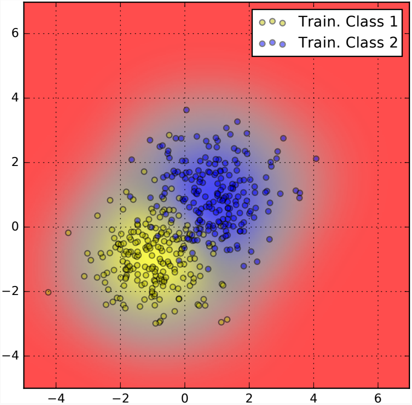


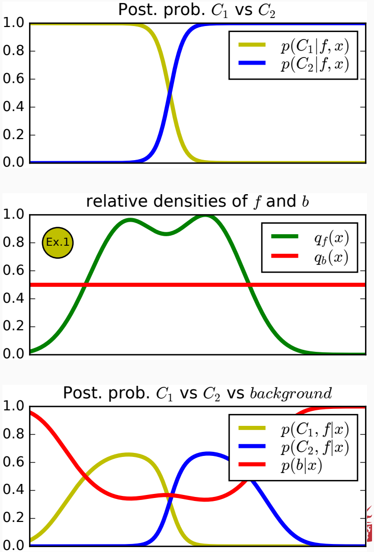

# Performing Background Check

- *Familiarity* approach
  - Pre-trained classifier
  - Learn $q_f(x) \in [0, 1]$
  - Use inductive bias

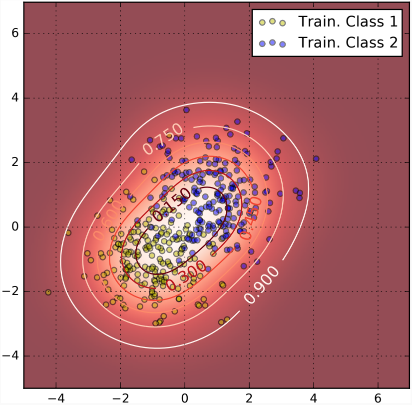


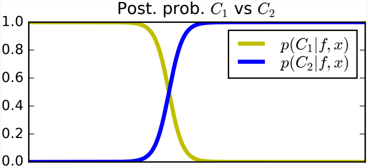


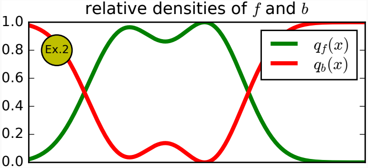

# Performing Background Check

- *Familiarity* approach
  - Pre-trained classifier
  - Learn $q_f(x) \in [0, 1]$
  - Use inductive bias

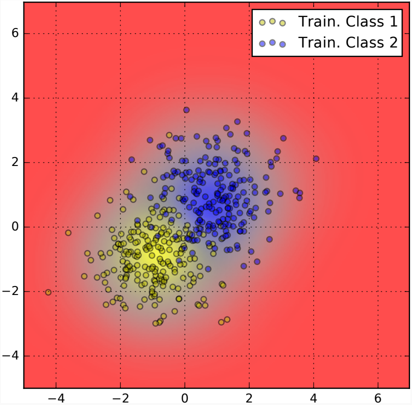


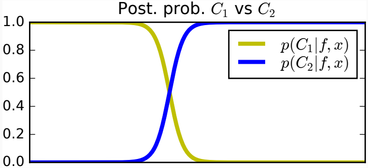


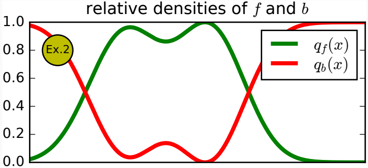


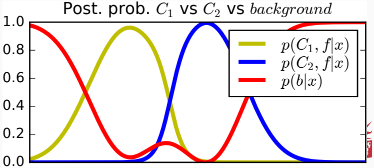

# Performing Background Check

- *Familiarity* approach
  - Pre-trained classifier
  - Learn $q_f(x) \in [0, 1]$
  - Use inductive bias

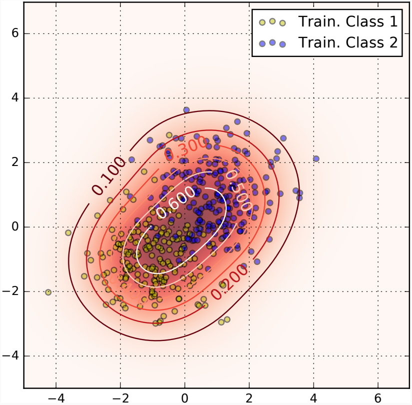


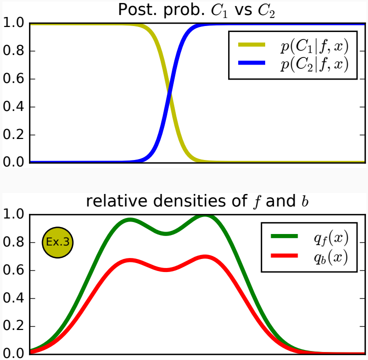

# Performing Background Check

- *Familiarity* approach
  - Pre-trained classifier
  - Learn $q_f(x) \in [0, 1]$
  - Use inductive bias

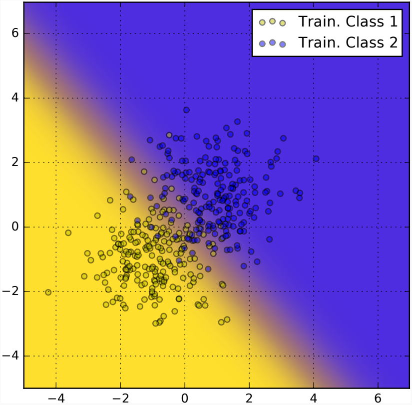


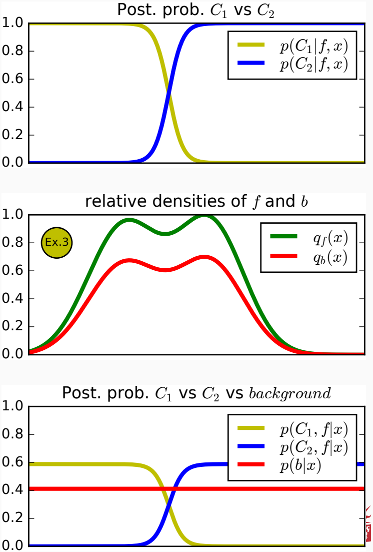


University of BRISTOL

# Performing Background Check

- *Familiarity* approach
  - Pre-trained classifier
  - Learn $q_f(x) \in [0, 1]$
  - Use inductive bias

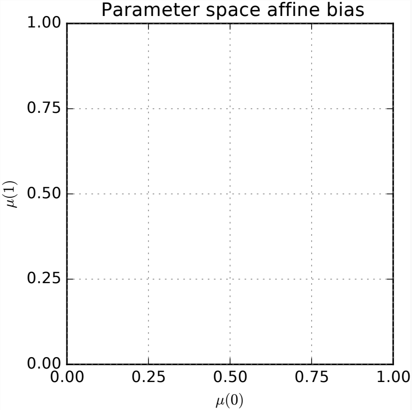


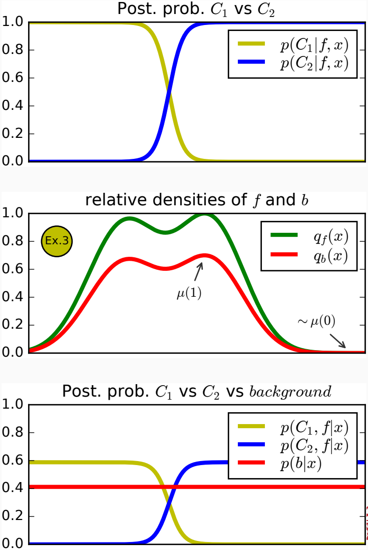


University of BRISTOL

# Performing Background Check

- *Familiarity* approach
  - Pre-trained classifier
  - Learn $q_f(x) \in [0, 1]$
  - Use inductive bias

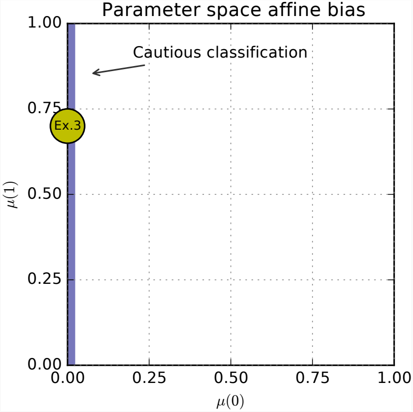


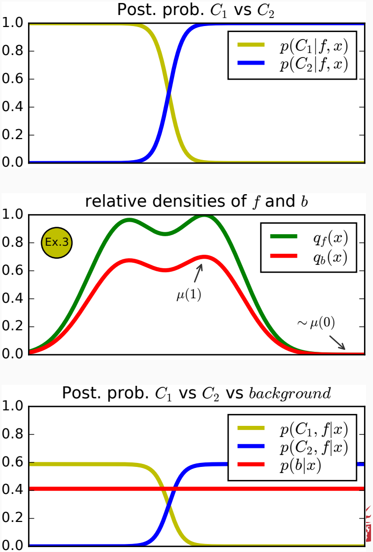


University of BRISTOL

# Performing Background Check

- *Familiarity* approach
  - Pre-trained classifier
  - Learn $q_f(x) \in [0, 1]$
  - Use inductive bias

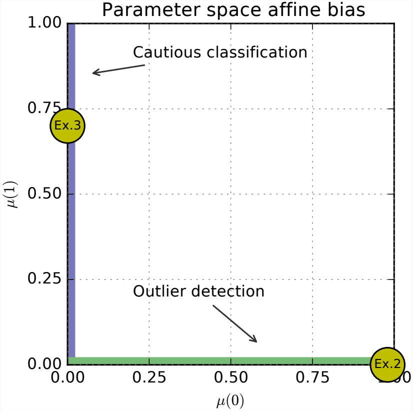


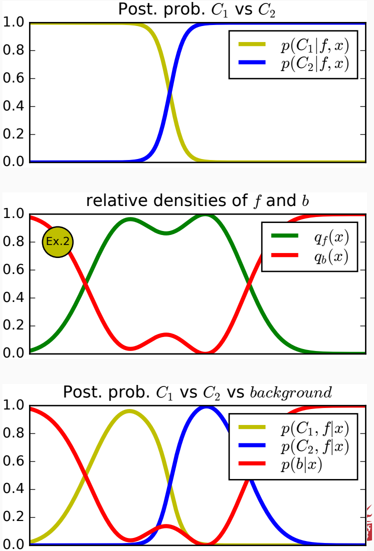


University of BRISTOL

# Performing Background Check

- *Familiarity* approach
  - Pre-trained classifier
  - Learn $q_f(x) \in [0, 1]$
  - Use inductive bias

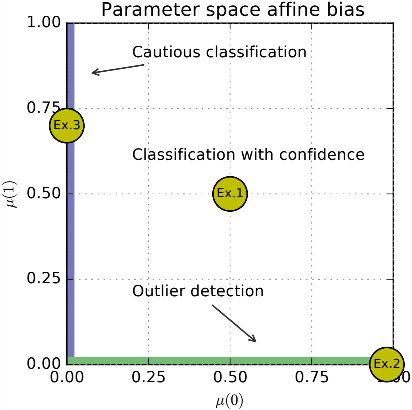


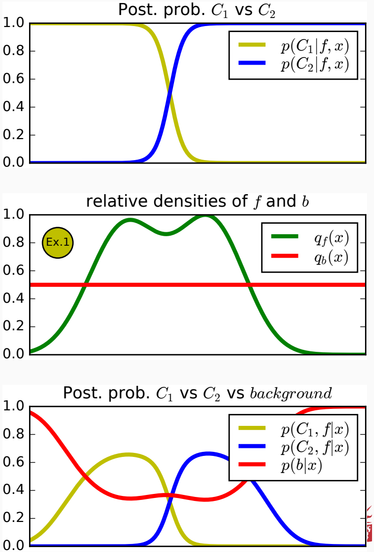

# Results

- Empirical evaluation
  - 41 multiclass datasets
  - 20 times 5-fold cross-validation
  - Classification with confidence
    - *Significantly better* than [Li et al., 2014] (Wilcoxon test $p < 0.005$)
  - Outlier detection
    - *Comparable results* with two specialized methods [Tax and Duin, 2008]
- Cautious Classification is *equivalent* to Chow's rule [Chow, 1970]

# Results

- Empirical evaluation
  - 41 multiclass datasets
  - 20 times 5-fold cross-validation
  - Classification with confidence
    - *Significantly better* than [Li et al., 2014] (Wilcoxon test $p < 0.001$)
  - Outlier detection
    - *Competitive results* with two specialized methods [Tax and Duin, 2008]
- Cautious Classification is *equivalent* to Chow's rule [Chow, 1970]

# Results

- Empirical evaluation
  - 41 multiclass datasets
  - 20 times 5-fold cross-validation
  - Classification with confidence
    - *Significantly better* than [Li et al., 2014] (Wilcoxon test $p < 0.001$)
  - Outlier detection
    - *Competitive results* with two specialized methods [Tax and Duin, 2008]
- Cautious Classification is *equivalent* to Chow's rule [Chow, 1970]

# Results

- Empirical evaluation
  - 41 multiclass datasets
  - 20 times 5-fold cross-validation
  - Classification with confidence
    - *Significantly better* than [Li et al., 2014] (Wilcoxon test $p < 0.001$)
  - Outlier detection
    - *Competitive results* with two specialized methods [Tax and Duin, 2008]
- Cautious Classification is *equivalent* to Chow's rule [Chow, 1970]

# Conclusion

- *General technique* to perform:
  - Cautious classification
  - Outlier detection
  - Classification with confidence
- Comparable and better results than *special purpose approaches*
- *Model agnostic*

# Conclusion

- *General technique* to perform:
  - Cautious classification
  - Outlier detection
  - Classification with confidence
- Comparable and better results than *special purpose approaches*
- *Model agnostic*

# Conclusion

- *General technique* to perform:
  - Cautious classification
  - Outlier detection
  - Classification with confidence
- Comparable and better results than *special purpose approaches*
- *Model agnostic*

Chow, C. (1970).
On optimum recognition error and reject tradeoff.
*IEEE Transactions on Information Theory*, 16(1):41–46.

Li, L., Hu, Q., Wu, X., and Yu, D. (2014).
Exploration of classification confidence in ensemble learning.
*Pattern Recognition*, 47(9):3120 – 3131.

Tax, D. and Duin, R. (2008).
Growing a multi-class classifier with a reject option.
*Pattern Recognition Letters*, 29(10):1565–1570.

# Background Check: A general technique to build more reliable and versatile classifiers

Miquel Perello-Nieto*[1], Telmo M. Silva Filho*[2], Meelis Kull[1] and Peter Flach[1]

* Equal contribution
[1]Department of Computer Science, University of Bristol, UK
[2]Centro de Informatica, Universidade Federal de Pernambuco, Brazil

Email: [1]{Miquel.PerelloNieto, Meelis.Kull, Peter.Flach}@bristol.ac.uk, [2]tmsf@cin.ufpe.br

December 13, 2016