Anomaly Ensembles

Nikitha Rao

Table of contents

? Problem statement

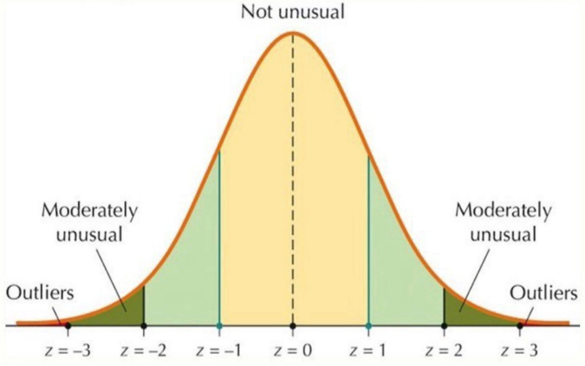

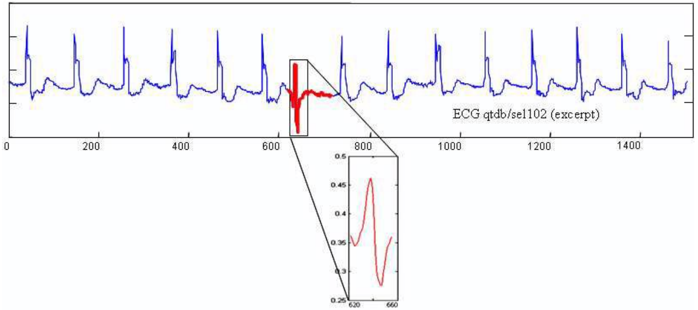

- **Anomaly detection - Identifying outliers**
- **Example: Left and right tails in a Normal distribution**
- **Several everyday application: Fraud detection, fault detection, defect detection etc.**
- **Main focus: Identify deviations from a “normal” data set.**

Image reference: https://medium.com/analytics-vidhya/how-anomaly-detection-works-4651f1786431 , Google slides: image search feature

Solution: Ensembles

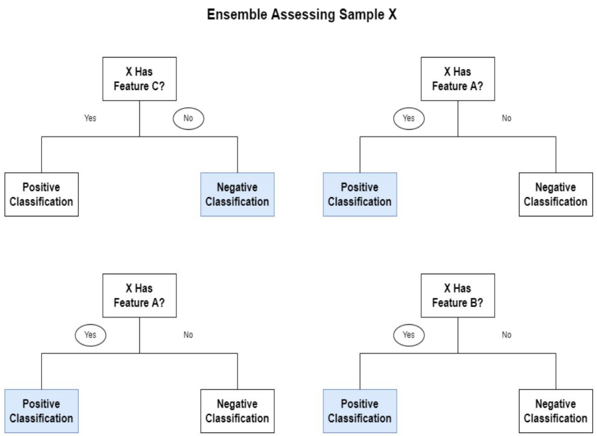

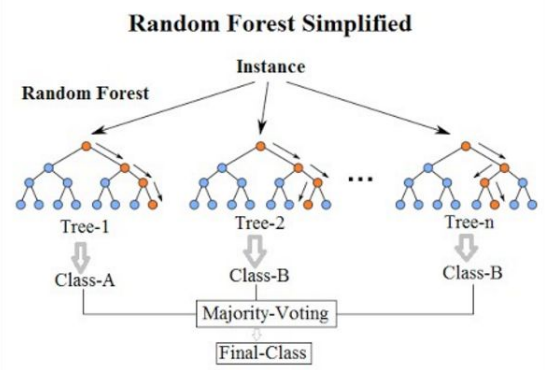

- **Combine several base models to produce one optimal model**
- **Bagging, stacking, boosting**
- **Example**
    - Decision trees, Random forests etc.

?: Challenge: Anomalies i.e. the class truth OR ground label - unknown!

Need: Class labels

Image reference: https://commons.wikimedia.org/wiki/File:Random_forest_diagram_complete.png, Wikipedia: Ensembles

Item response theory

- Models the ground truth as a latent trait (unobservable characteristic or attribute).
- Explains the relationship between
  - 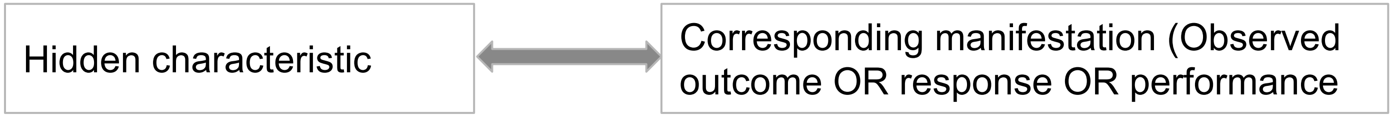

- Establishes a link between
  - 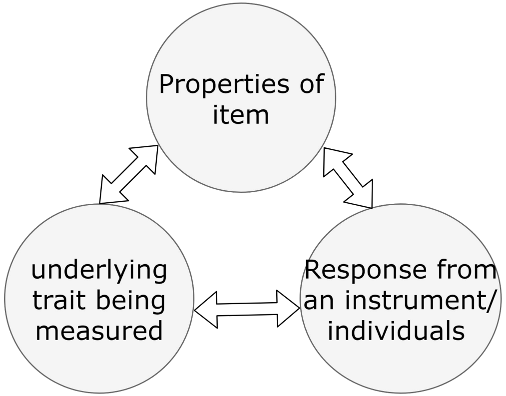

Item response theory

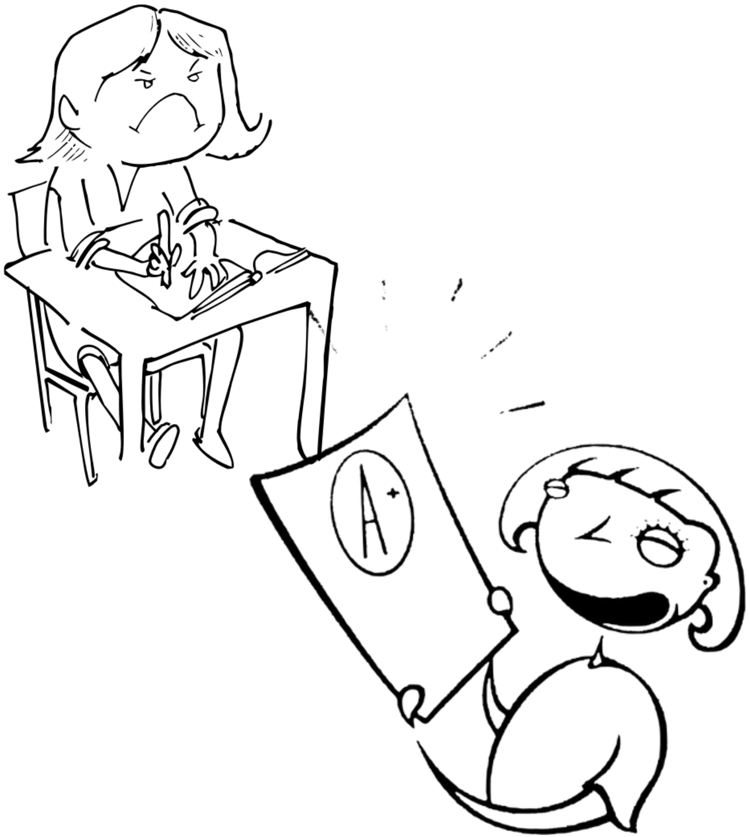

Image reference: Google slides: image search feature

- Psychometrics - Root of IRT
- Classic example: Examination
- 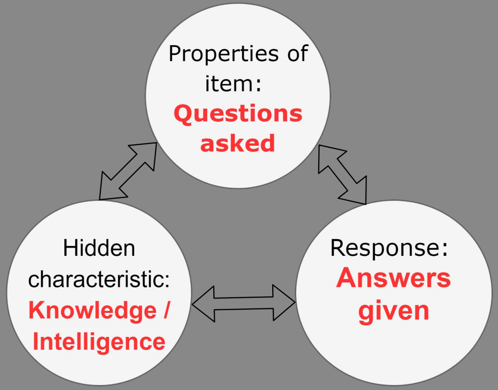

- Extend the concept to anomaly detection

Item response theory

- Earlier researches: AD as Supervised machine learning models that used pseudo-ground-truth labels
  - Limitations: Circular argument
- Current research: IRT ensembles for unsupervised anomaly detection.
  - Latent trait: Uncover the ground truth
- IRT framework in the paper also uses - different combination functions like average, maximum and other correlation based functions on the set of the heterogenous anomaly methods.
- This paper implements unsupervised anomaly detection using combination function on ensemble

Summarising AD & IRT

Several keywords in previous slide!!

- **Anomaly detection?**
- **Unsupervised anomaly detection?**
- **Unsupervised anomaly detection using ensemble!?**
- **Combination function?**
- **Unsupervised anomaly detection using combination function on ensemble!?**

Image reference: Google slides: image search feature

Summarising AD & IRT

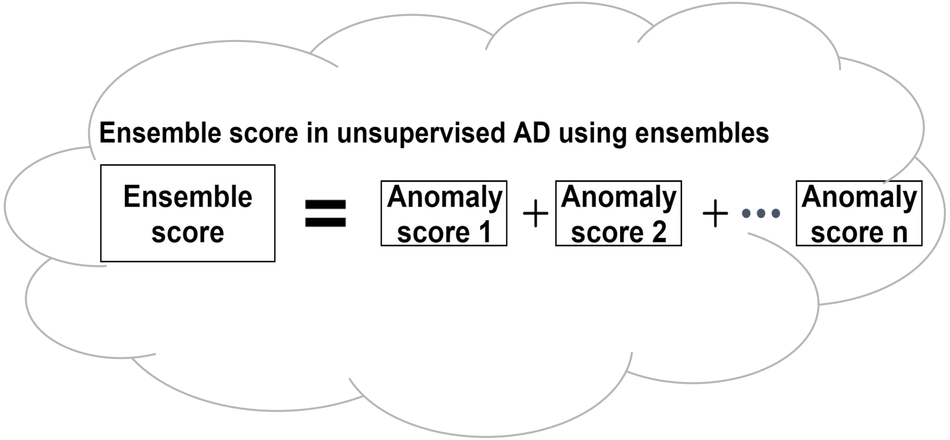

- **Anomaly detection:** *Identifying the outliers*
- **Unsupervised Anomaly detection:** *Identifying the outliers when the ground truth is unknown*
- **Unsupervised anomaly detection using ensembles**
  - *Unsupervised Anomaly Detection (AD) methods - gives - "anomaly scores" for each observation in the dataset*
  - *Larger scores - anomalous observations.*
  - *Ensemble scores = combination of normalised anomaly scores from several methods*

Summarising AD & IRT

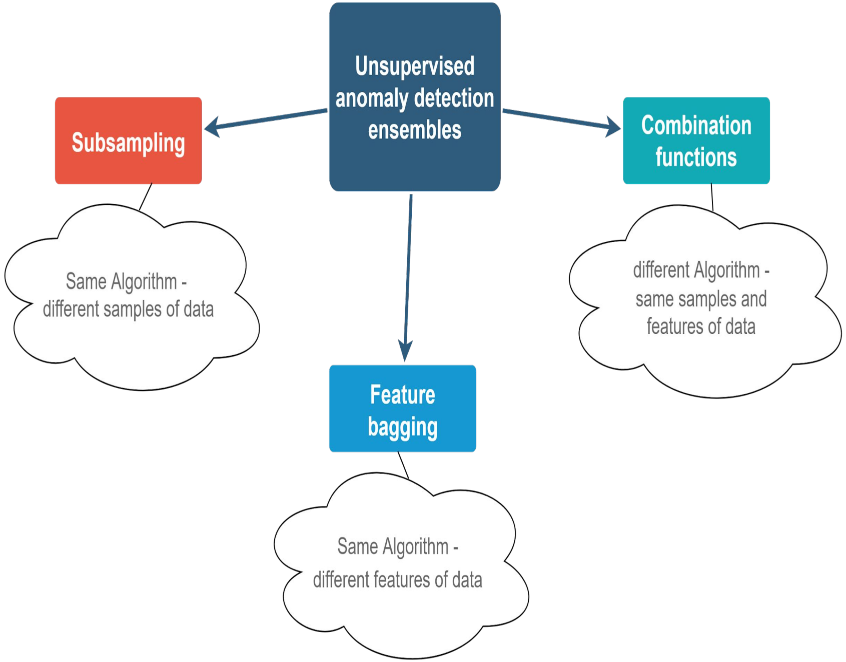

- **Combination function:** *One from the flow chart*
- **Unsupervised anomaly detection using combination function on ensemble:**
  - *Used in this research.*
  - Anomaly scores - by several different methods. Ex: LOF - density based outlier with decisions based on LOF score
  - Ensemble score by - *combination functions like average, greedy, averaged greedy, ICWA, Max, thresh along with IRT.*
  - *Best performing combination function as per ROC is analysed.*

Quick recap on ROC

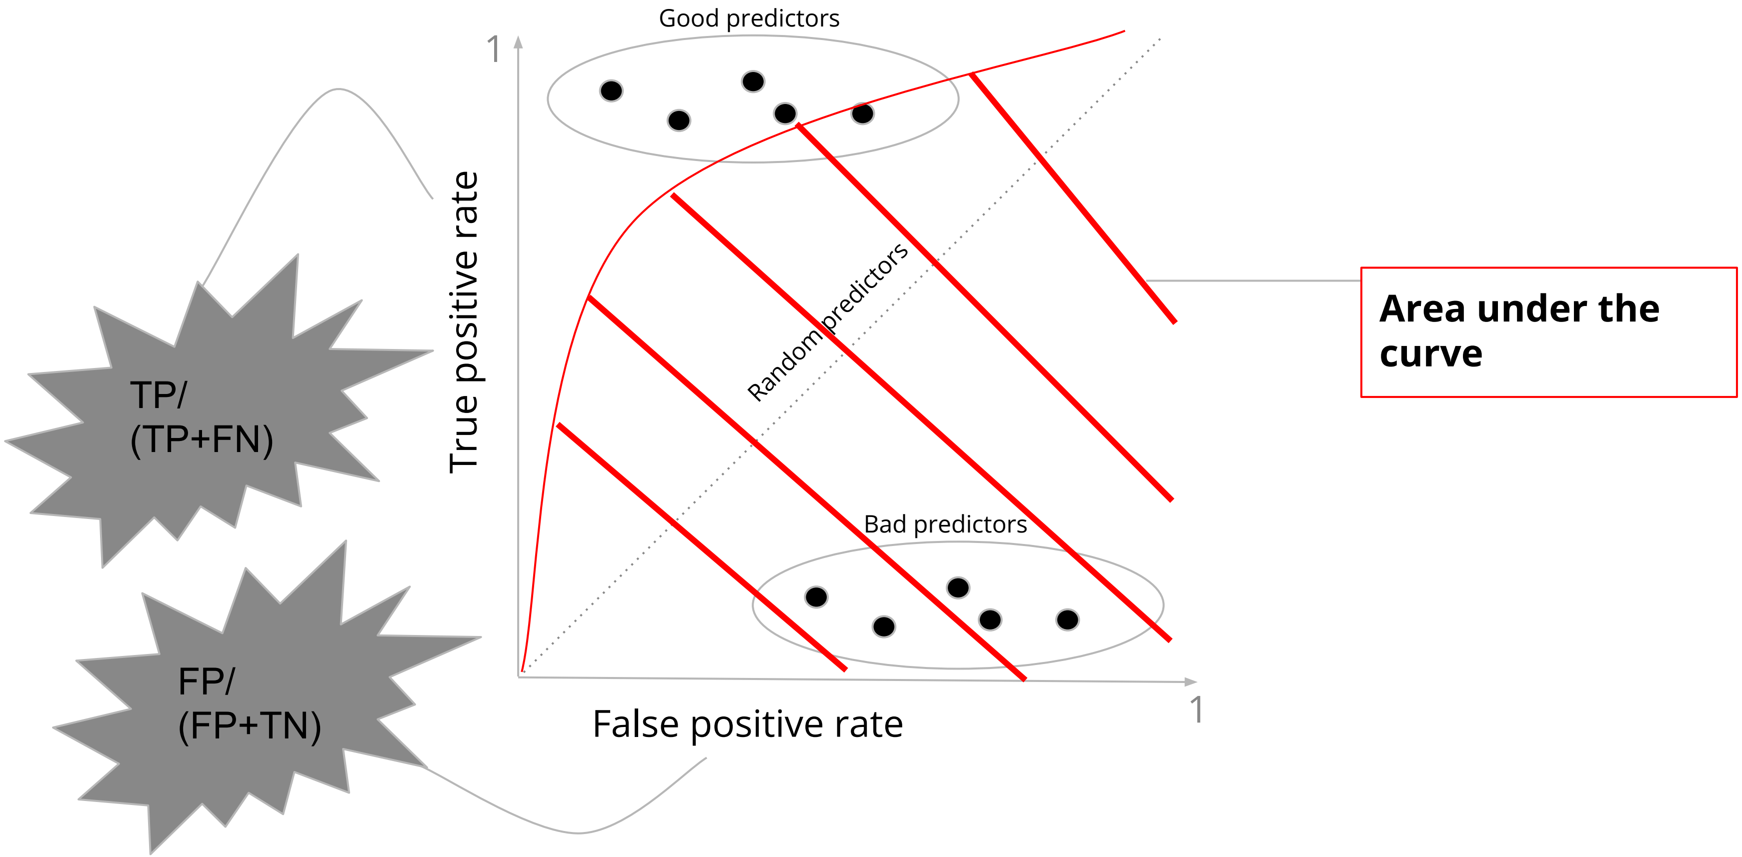

More on algorithms & implementation

Average

- Computes average anomaly score
- Pro: Benchmark, performs well in homogeneous data distribution
- Cons: Subpar performance in heterogeneous data

Greedy

- If something works, we favor it more.
- Selects best methods - using correlation and K number of anomalies
- Pro: able to handle correlation
- Cons: slower than average, might not arrive at optimal solution

Many more combination functions

● ● • •

More on algorithms & implementation

IRT

- From earlier slides: latent trait model, fitted using responses. There are several types of IRT based on the responses used in fitting

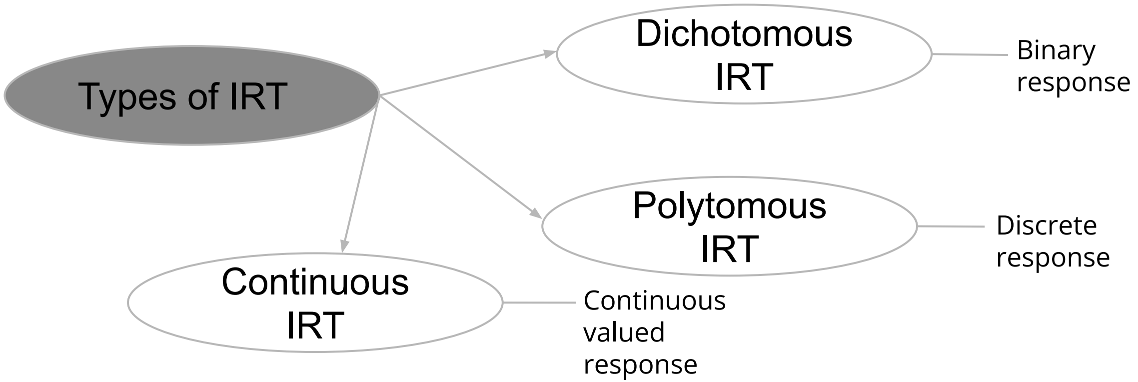

- Look deep into - Dichotomous IRT

More on algorithms & implementation

$$\Phi\left(y_{ij} = 1 | \theta_i, \alpha_j, \beta_j\right) = \frac{1}{1 + \exp\left(-\alpha_j\left(\theta_i - \beta_j\right)\right)}.$$

- **Assume:**
  - i = 1, 2, . . . , N participants
  - j = 1, 2, . . . , n test items
  - Y_{ij} ∈ {0, 1} - score or response of the i th participant to the j th test item
- **Parameters:**
  - Discrimination parameter α_{j}
  - difficulty parameter by β_{j}
- **Equation and 2-Parameter Logistic (2PL) mode as shown. θ_{i} is ability.**
- **Note: When α_{j} = β_{j}, , Φ = 0.5**

Image reference: Kandanaarachchi, Sevvandi. "Unsupervised anomaly detection ensembles using Item Response Theory." *arXiv preprint arXiv:2106.06243* (2021).

Mapping IRT to ensemble

- Two possibilities of mapping the response in psychometrics

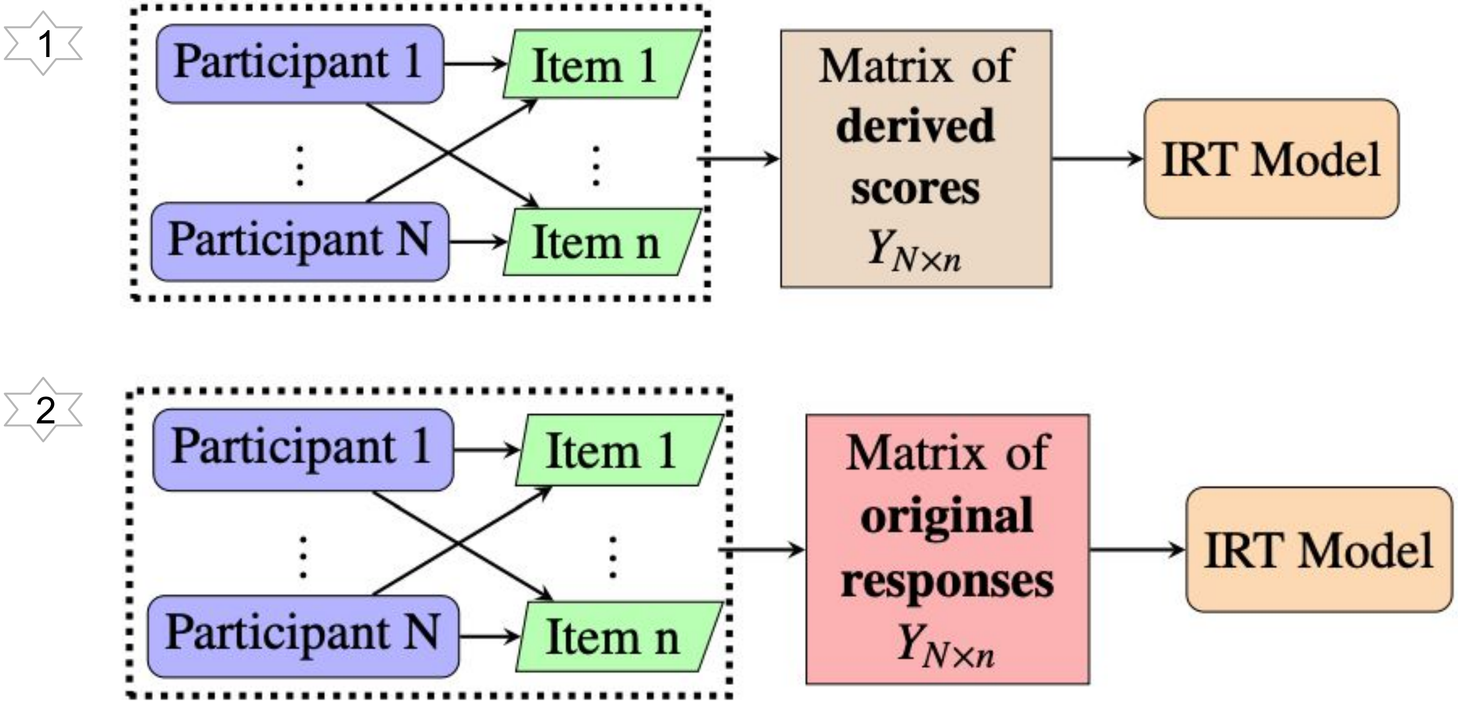

Image reference: Kandanaarachchi, Sevvandi. "Unsupervised anomaly detection ensembles using Item Response Theory." *arXiv preprint arXiv:2106.06243* (2021).

Mapping IRT to ensemble

- Two possibilities of algorithm evaluation

1

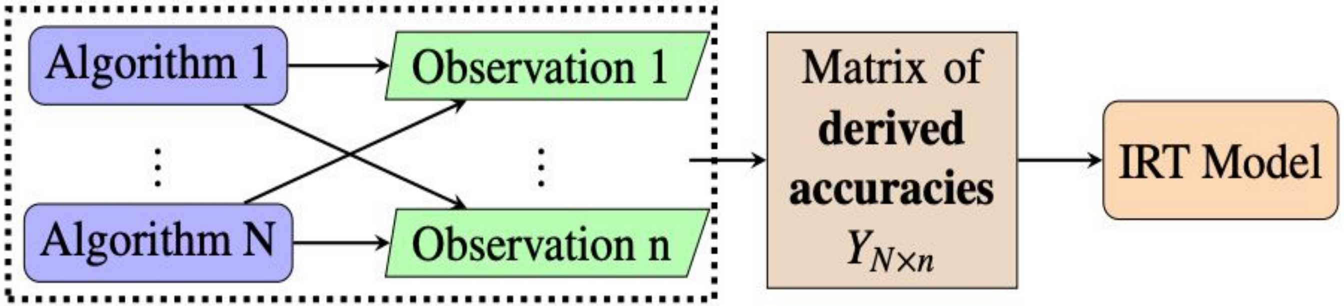

2

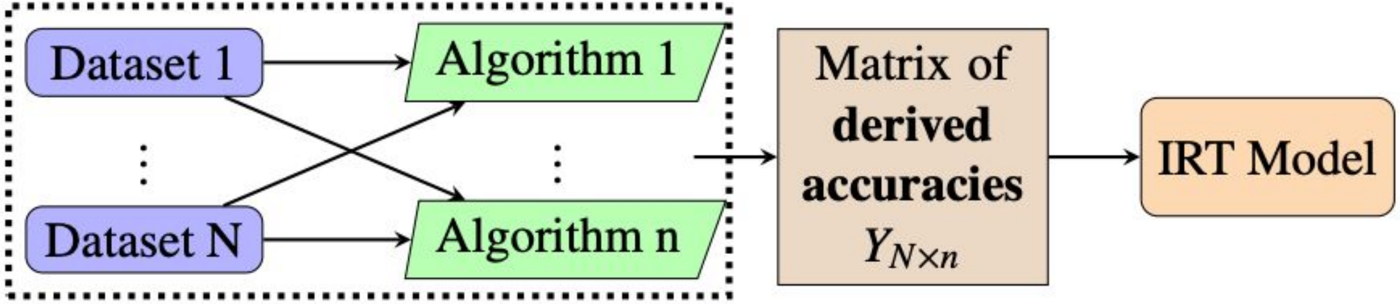

Image reference: Kandanaarachchi, Sevvandi. "Unsupervised anomaly detection ensembles using Item Response Theory." *arXiv preprint arXiv:2106.06243* (2021).

Mapping IRT to ensemble

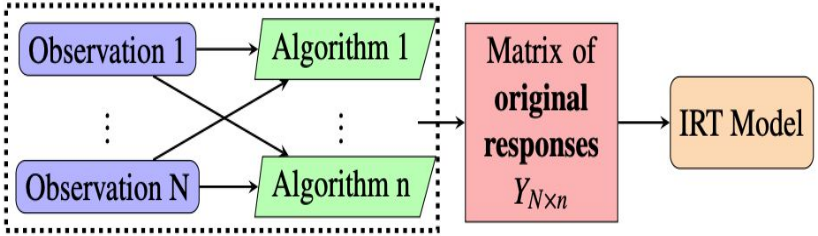

- Current research - Mixture of both
- Uses Standardised Original responses instead of an accuracy measure.
- Maps participants to dataset observations and test items to algorithms.

Image reference: Kandanaarachchi, Sevvandi. "Unsupervised anomaly detection ensembles using Item Response Theory." *arXiv preprint arXiv:2106.06243* (2021).

Experiments in Ensemble

Results:

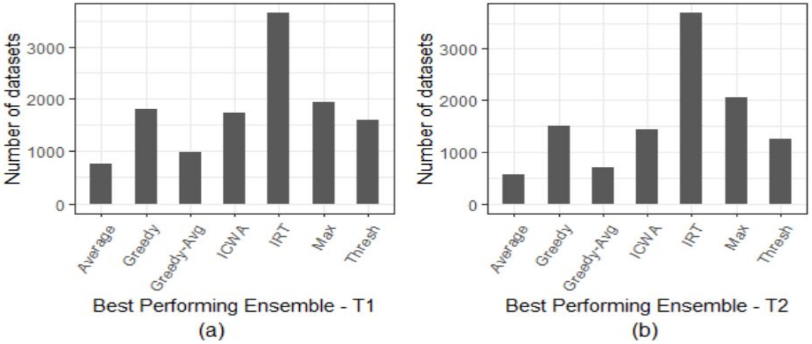

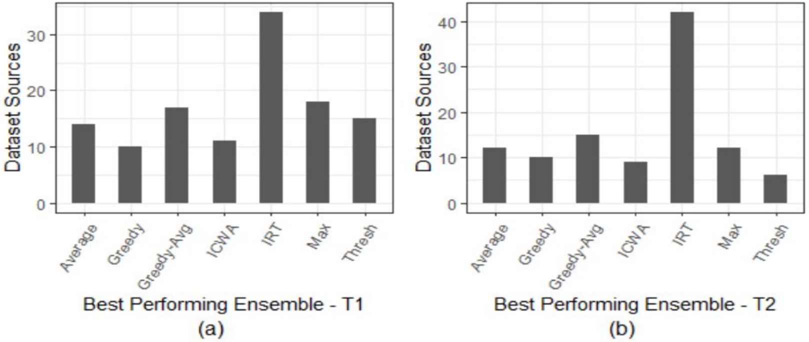

- **Data set:**
  - 12433 publicly available anomaly detection datasets
  - prepared from 119 source datasets
  - minority class is considered anomalous
- **Methods:** *KNN-AGG, LOF, COF, INFLO, KDEOS, LDF and LDOF*
- **Combination functions:** *average, greedy, averaged greedy, ICWA, Max, thresh, IRT.*
- **Parameter Setting:**
  - *T1: default values for k = k_min = 5 and k_max = 10*
  - *T2: k = k_min = max(N/10, 50) and k_max = k + 10*
- **Performance evaluation:** *AUC*

Image reference: Kandanaarachchi, Sevvandi. "Unsupervised anomaly detection ensembles using Item Response Theory." *arXiv preprint arXiv:2106.06243* (2021).

Conclusion and improvements

Conclusion

- In unsupervised learning - constructing ensemble using heterogeneous AD methods - challenge
- Introduced - IRT - uses latent trait to compute the ground truth for the first time.
- Evaluated the IRT ensemble w.r.t. 6 other AD ensemble techniques. Result - IRT performs the best

Improvements

- When and why ensembles work - is not explained
- Can justify more on the base models used
- Harder to justify Ensemble results
- AUC is a basic performance evaluator. Does not compare anomaly scores.
- IRT & Ensembles in general is computationally expensive.

References

- Kandanaarachchi, Sevvandi. "Unsupervised anomaly detection ensembles using Item Response Theory." *arXiv preprint arXiv:2106.06243* (2021).
- Zimek, Arthur & Campello, Ricardo & Sander, Joerg. (2014). Ensembles for unsupervised outlier detection: challenges and research questions. A position paper. ACM SIGKDD Explorations Newsletter. 15. 11-22. 10.1145/2594473.2594476.
- Embretson, Susan, and Steven Reise. (2013) 2013. Item Response Theory. 1st ed. Taylor and Francis.
- Martínez-Plumed, Fernando et al. “Making Sense of Item Response Theory in Machine Learning.” ECAI (2016).
- A. Chiang and Y. Yeh, "Anomaly Detection Ensembles: In Defense of the Average," 2015 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT), 2015, pp. 207-210, doi: 10.1109/WI-IAT.2015.260.
- Schubert, Erich & Wojdanowski, Remigius & Zimek, Arthur & Kriegel, Hans-Peter. (2012). On Evaluation of Outlier Rankings and Outlier Scores. 10.1137/1.9781611972825.90.

Thank you !