Anomaly Detection and Categorization Using Unsupervised Deep Learning

S6340

Thursday 7th April 2016
GPU Technology Conference

A. Stephen McGough, Noura Al Moubayed, Jonathan Cumming, Eduardo Cabrera, Peter Matthews, Toby P. Breckon, Ed Ruck-Keene, Georgios Theodoropoulos
Durham University, UK

intel

Intel Parallel Computing Centre

TOP 500

The List.

NVIDIA

CUDA RESEARCH CENTER

Why I'm here?

- UK has a major focus on Academic Impact
- Researchers collaborating with Industry
- Durham University has an Impact agenda
  - Which paid for this trip
- I'm actively seeking collaborations with Companies / Organizations

The Problem

- “90% of all the data in the world has been generated over the last two years”… IBM

- “85% of worldwide data is held in un-structured formats”… Berry and Kogan

- How can we understand it? ….or better still make use of it?
- How can we determine the most pertinent information? …and then act on it?
- How can we find the needle if we are not sure what it looks like or what hay looks like?

Anomaly Detection Framework

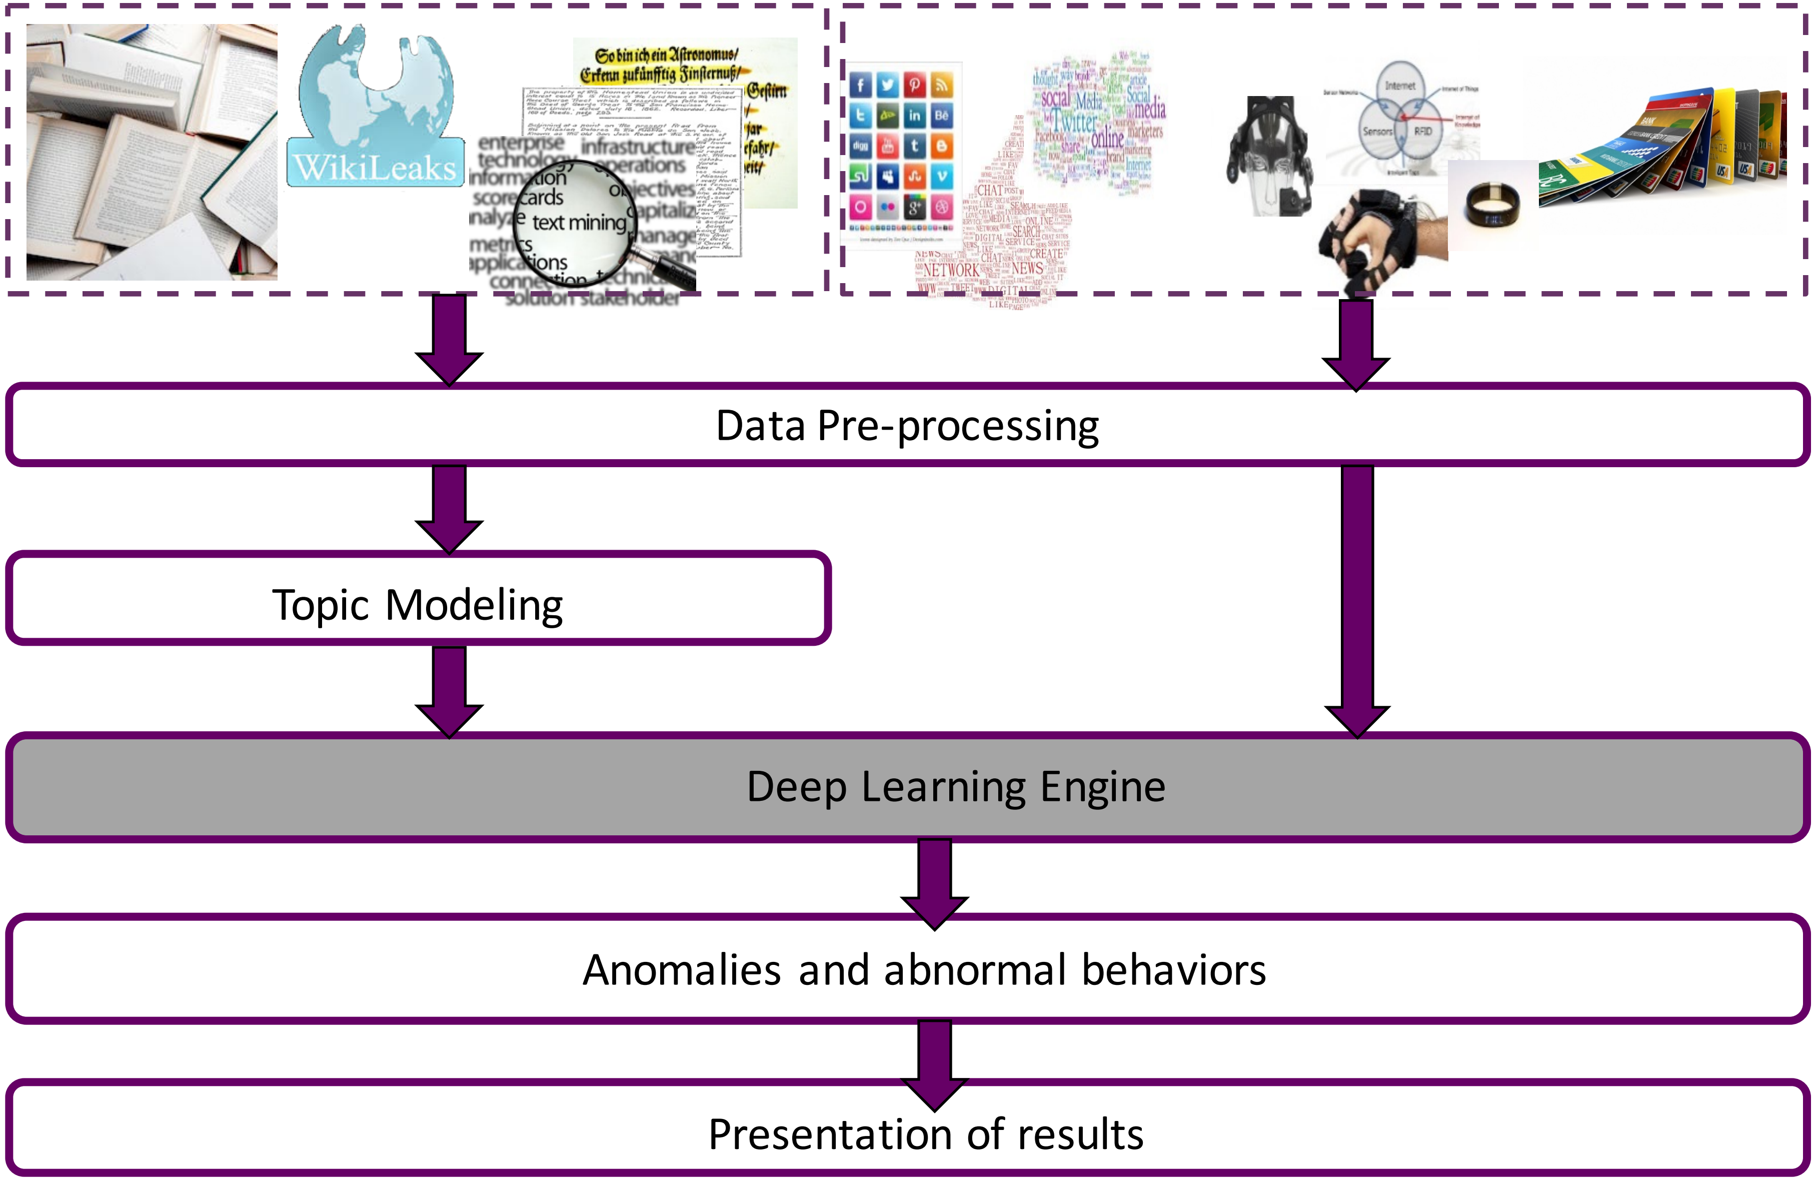

Topic Modelling

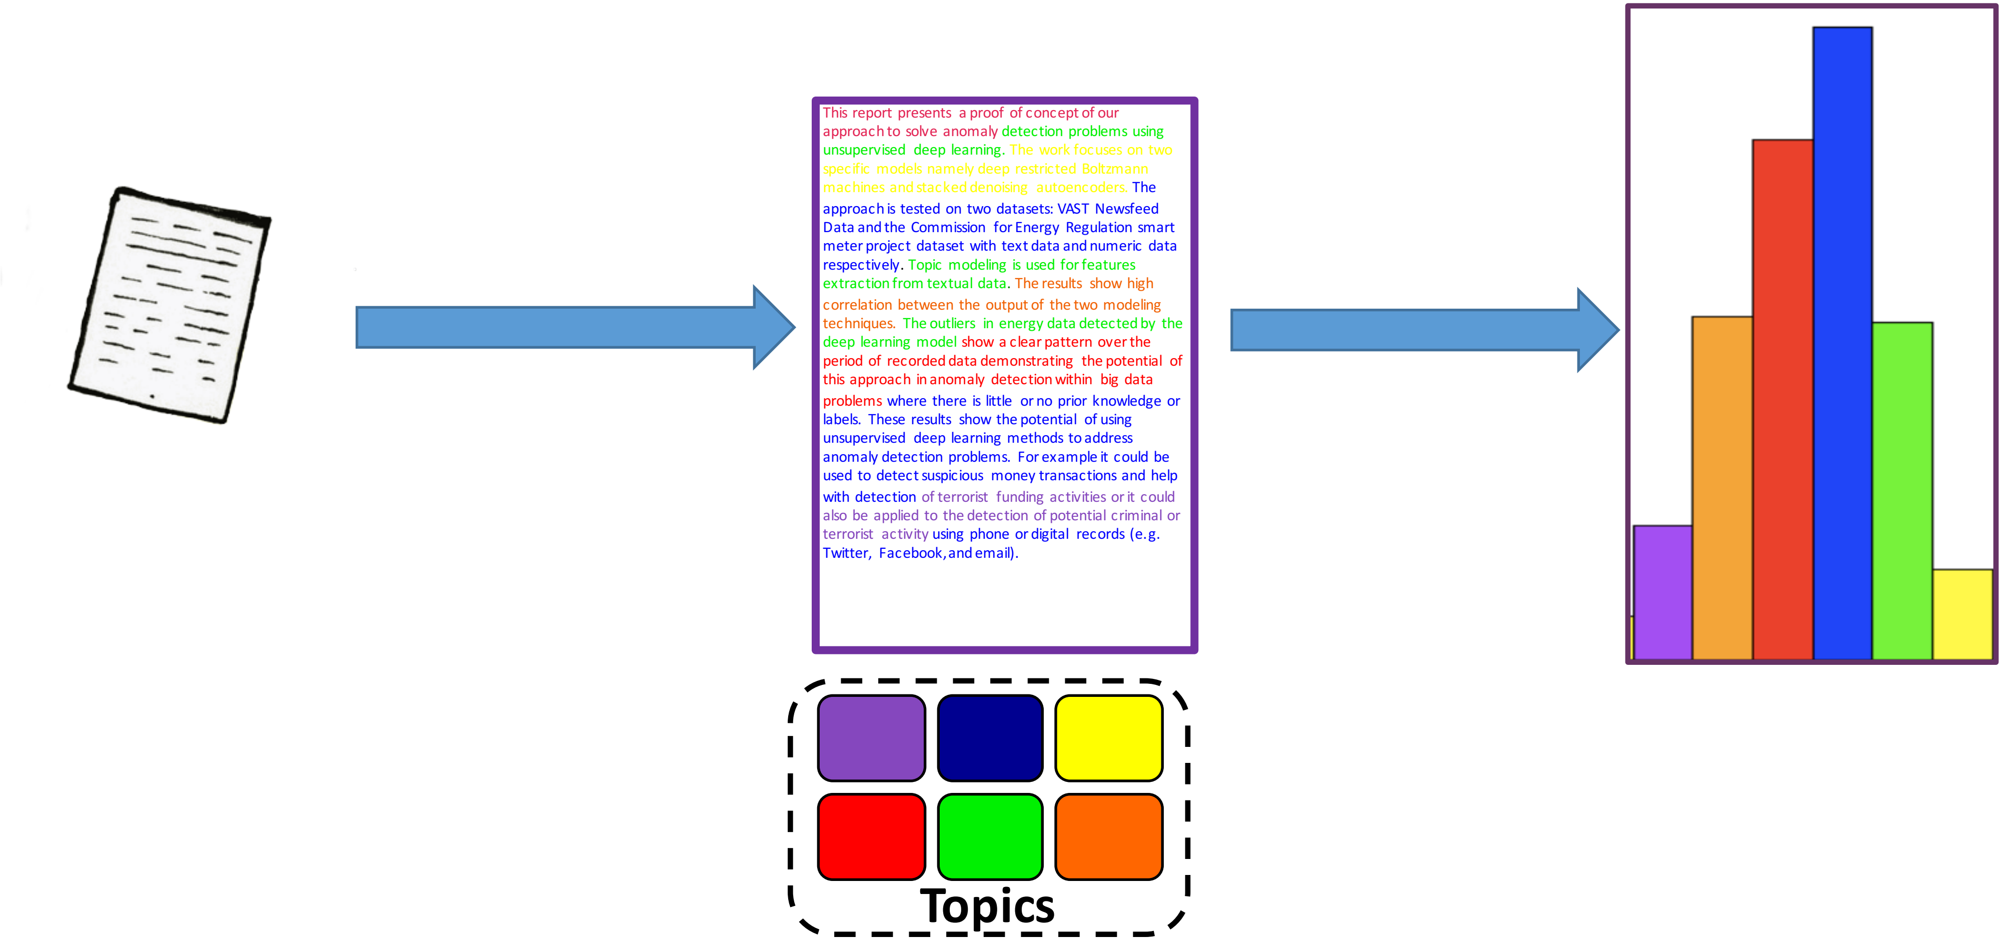

Topic Modelling

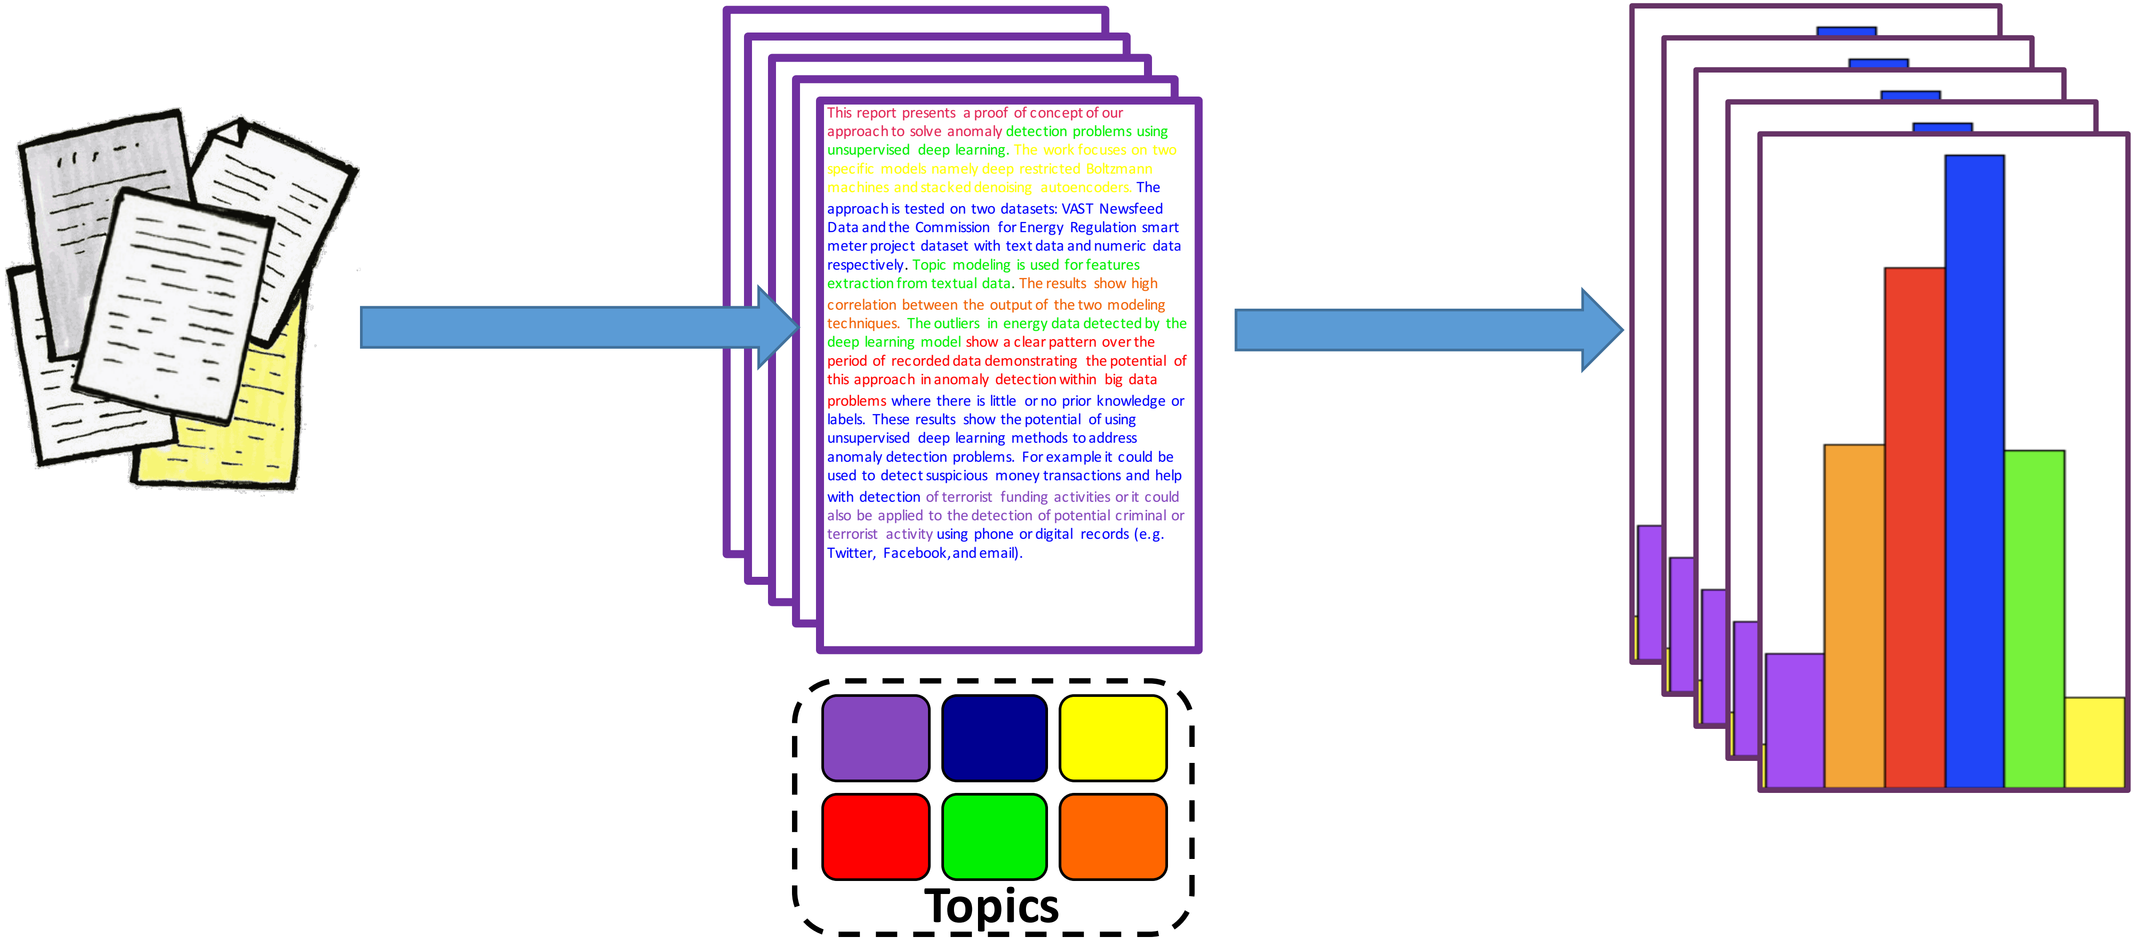

Topic Modelling

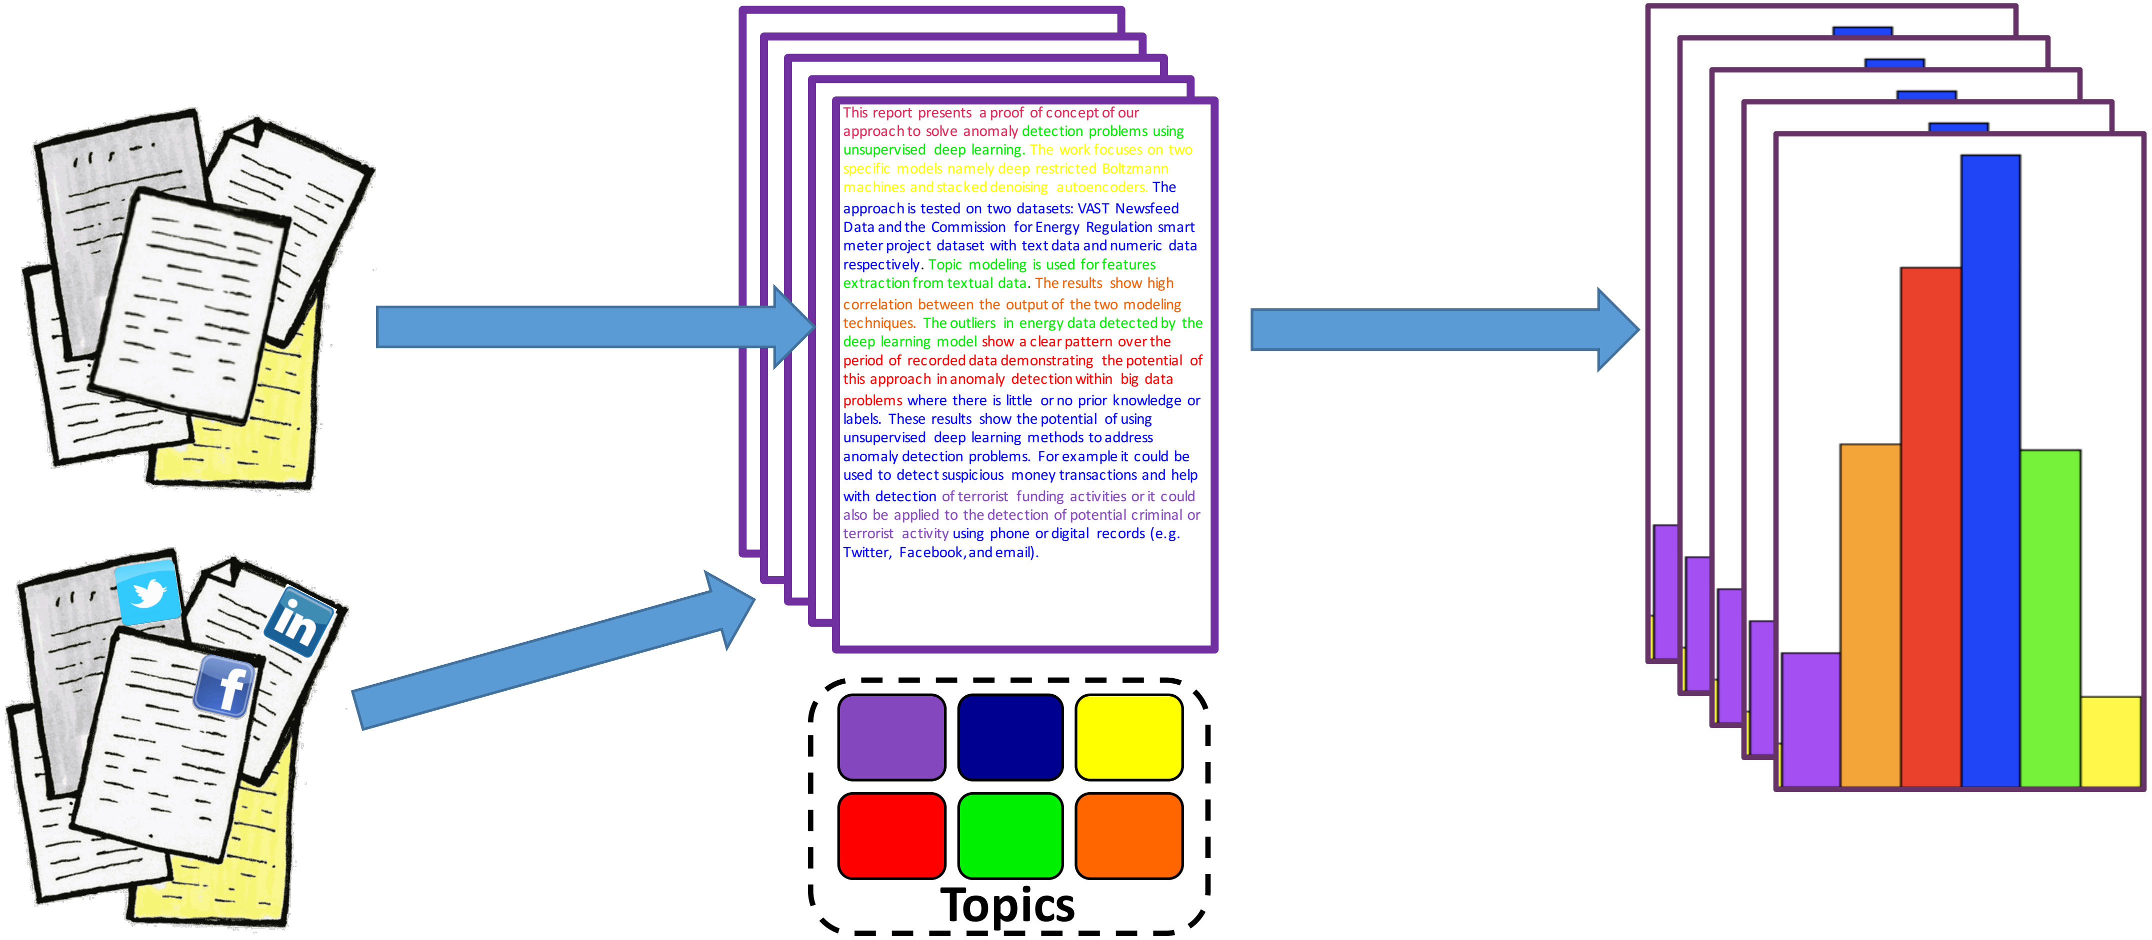

Probabilistic Topic Modelling

- Unsupervised analysis of text
  - Too many documents to label manually
- Allows us to uncover automatically themes that are latent in a collection of documents
- Same words may have different meanings depending on their co-occurrence with other words in a document
- Statistically identify the topics from a set of documents
  - Which words often found in the same document
- Statistically classify which topics appear in each document
  - Which topics appear in each document

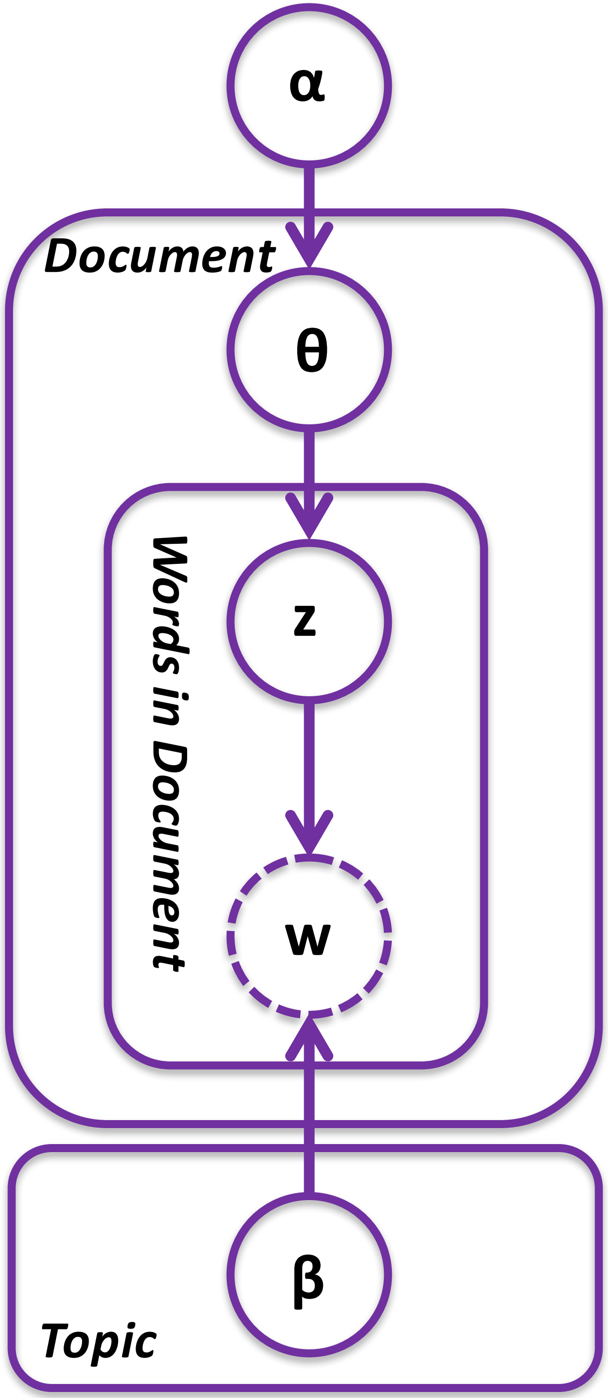

Anomaly Detection: Unsupervised Deep Learning

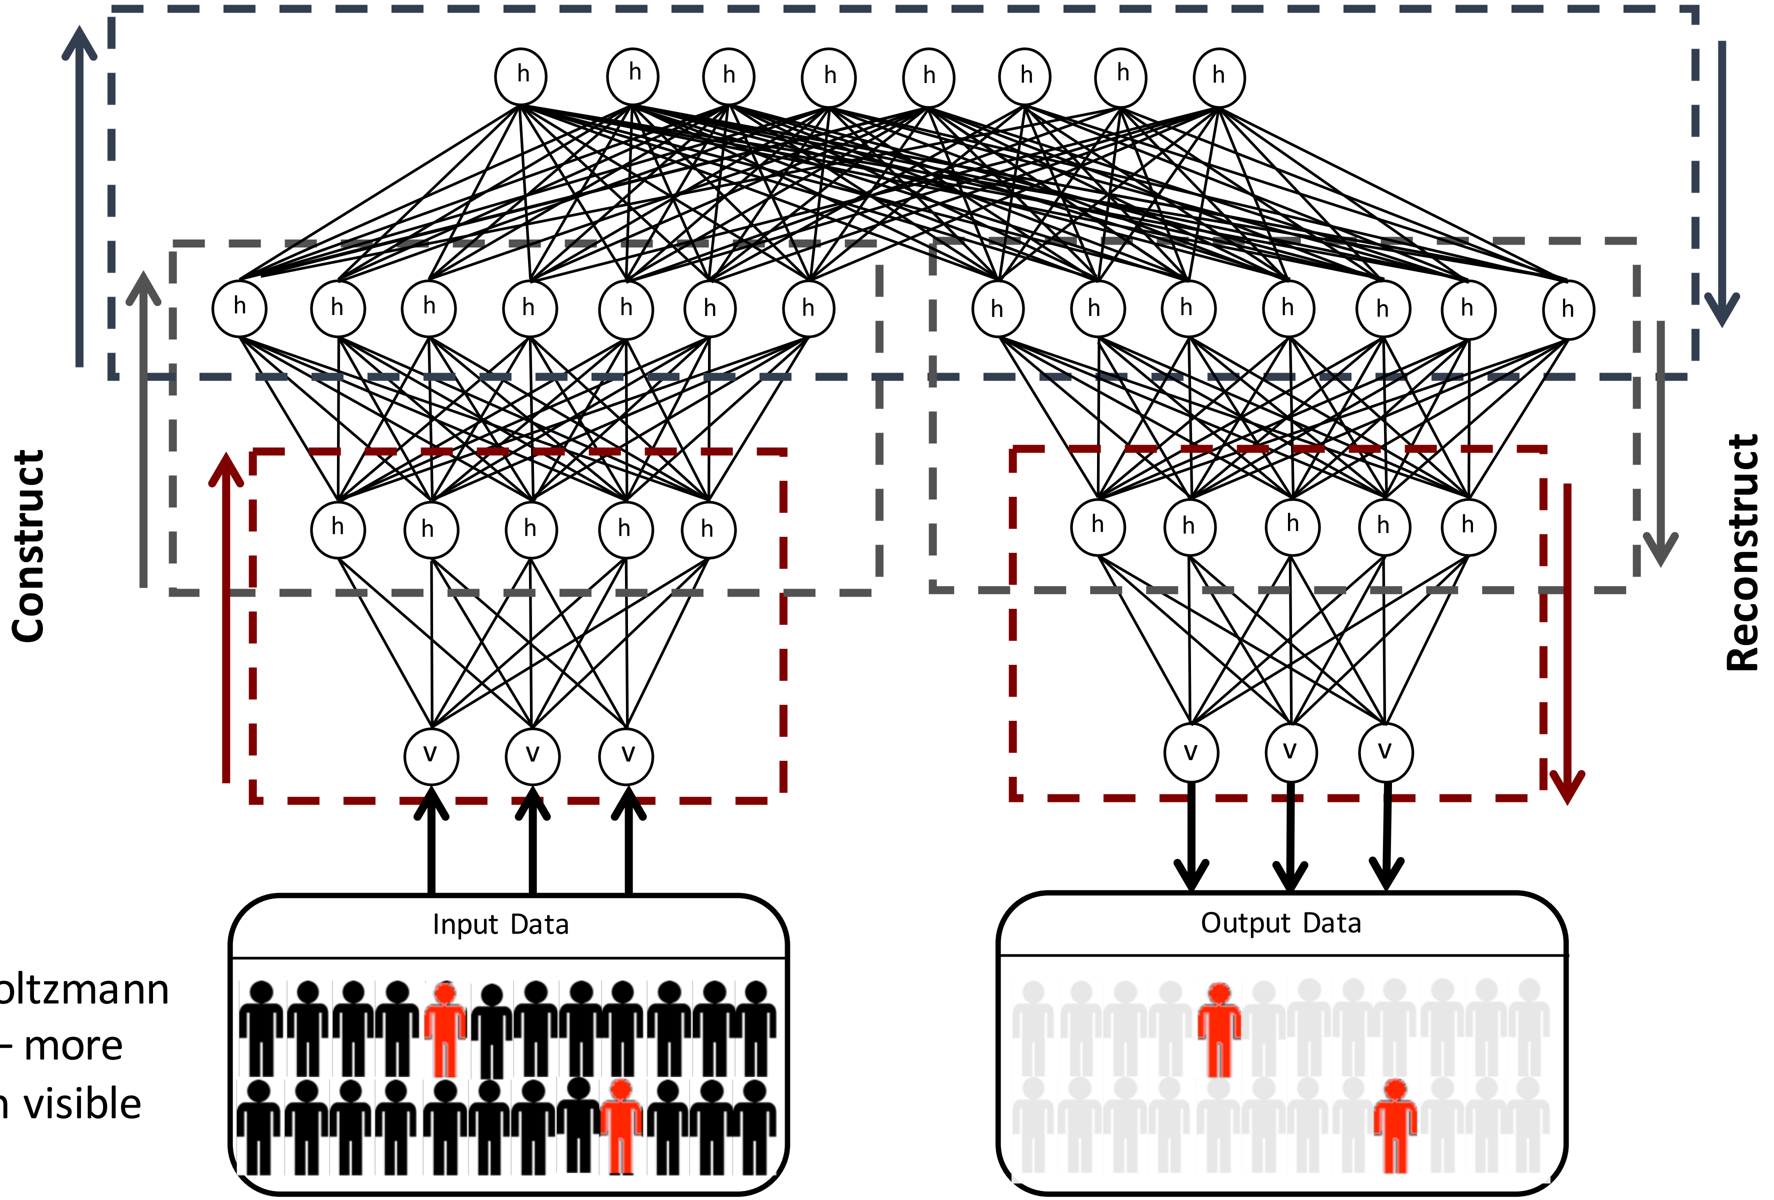

Deep Restricted Boltzmann Machine (DRBM) – more hidden nodes than visible nodes

Anomaly Detection: Unsupervised Deep Learning

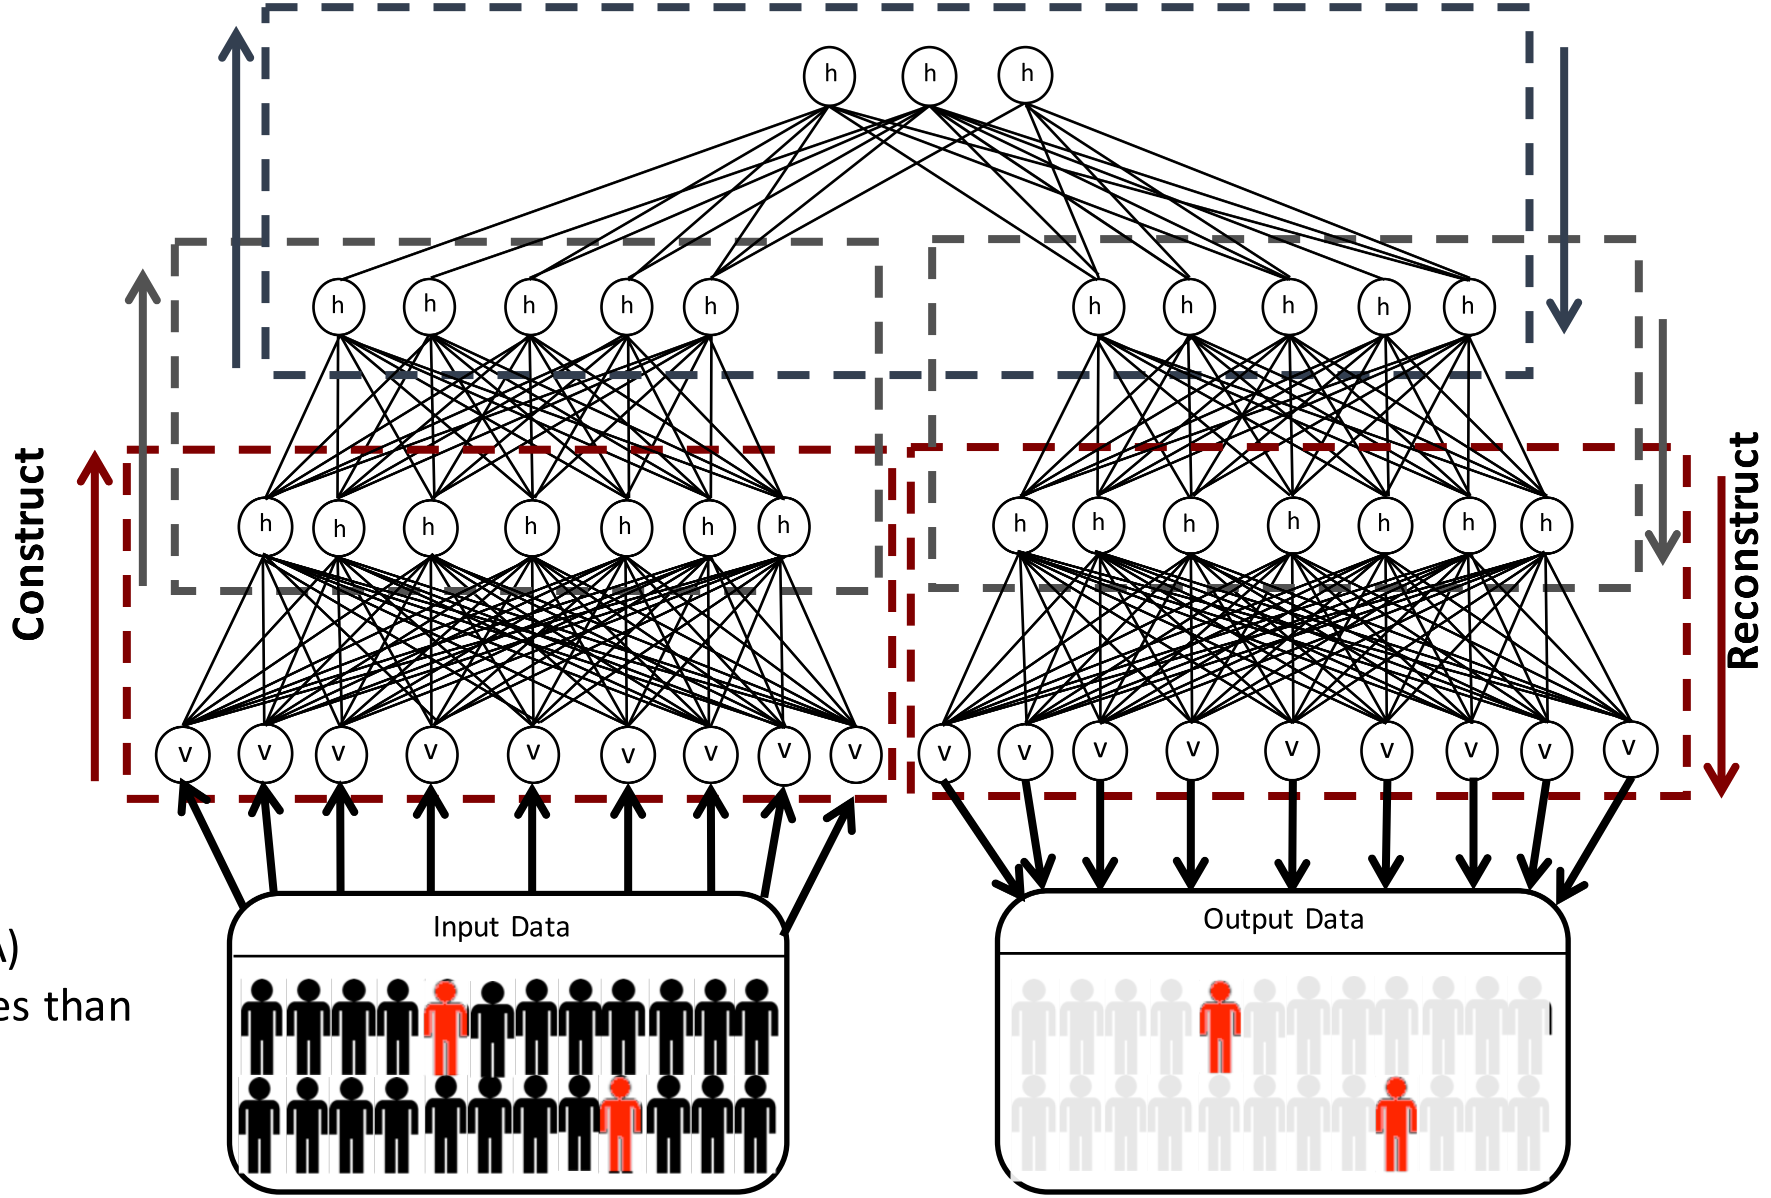

Stacked Denoising Autoencoder (SDA)

- Less hidden nodes than visible nodes

Overall Methodology

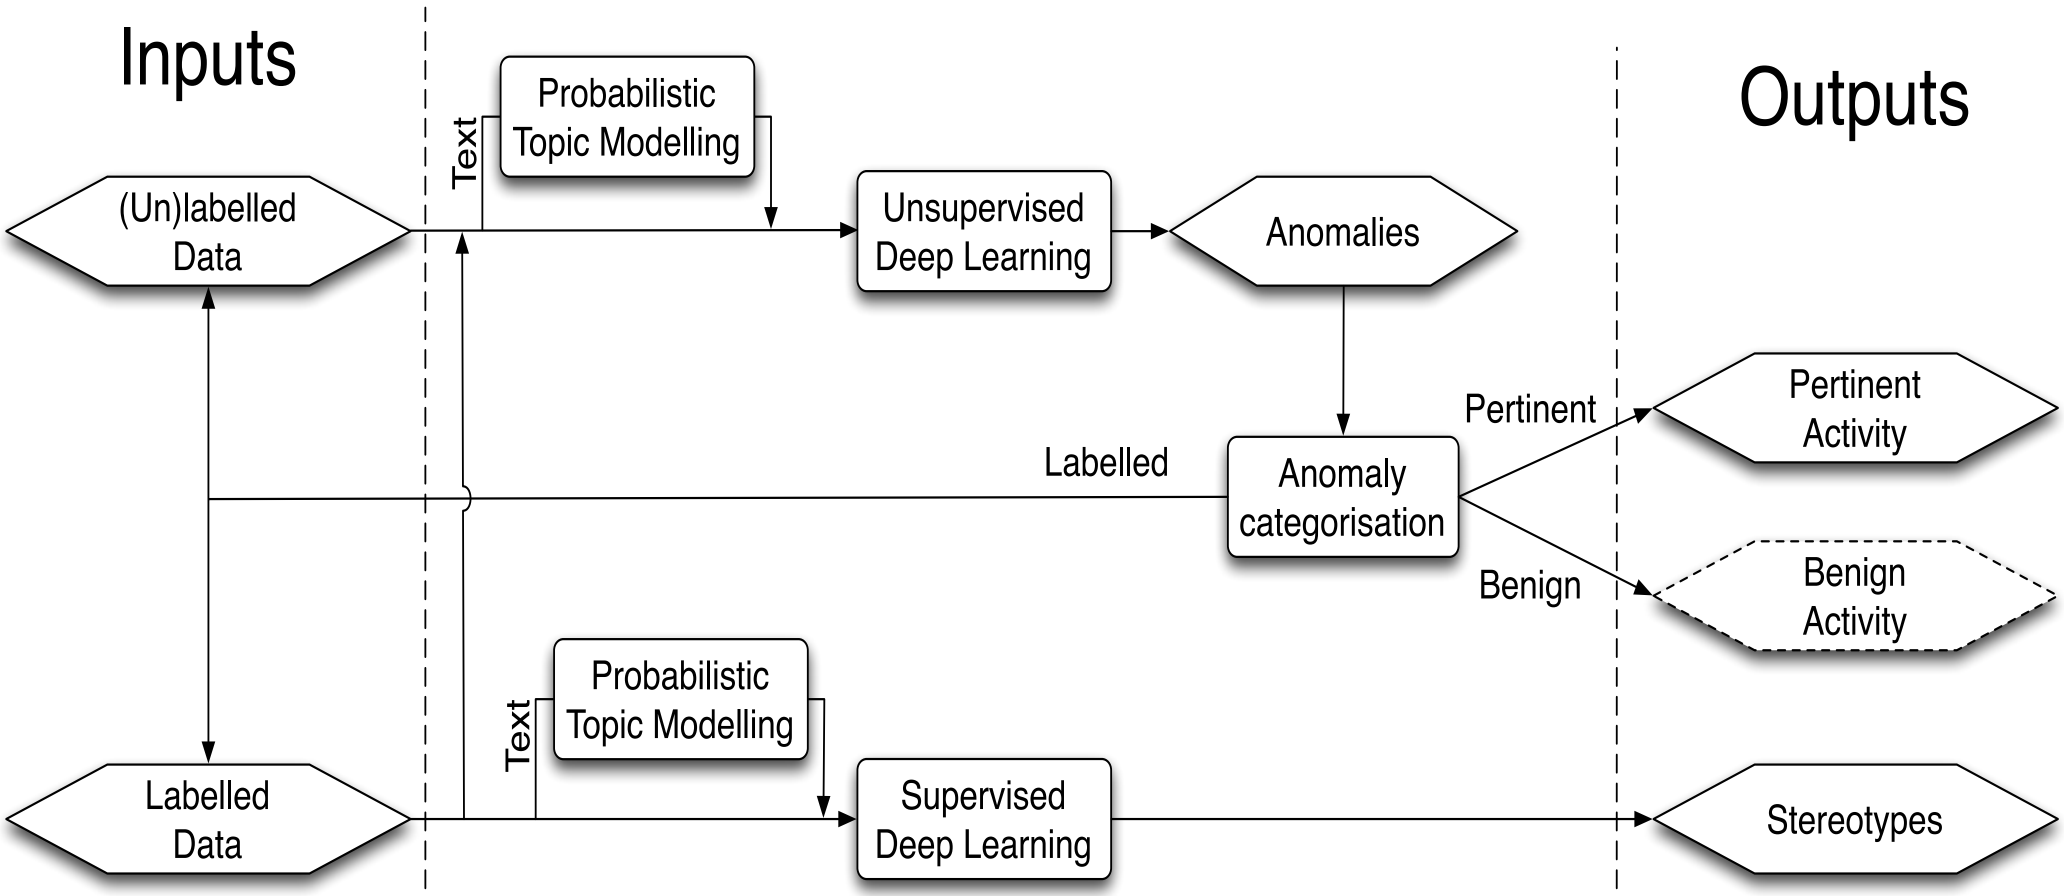

Overall Methodology

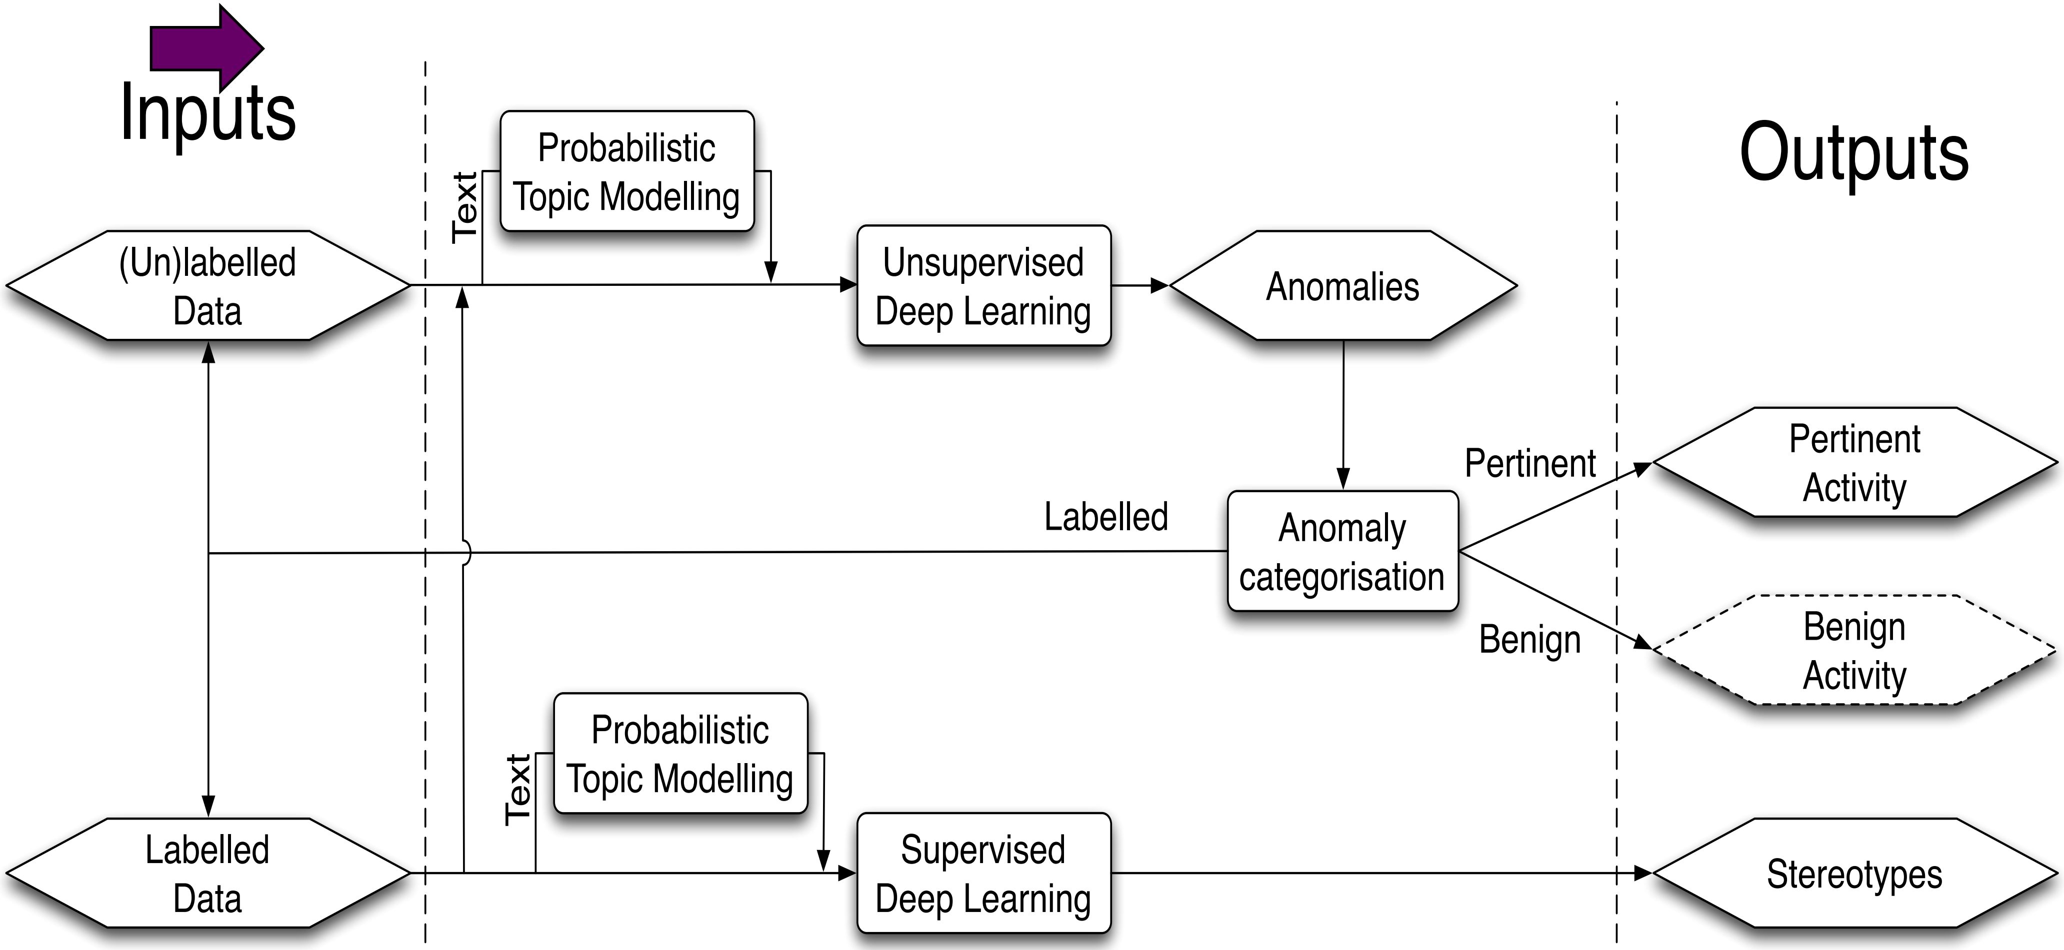

Overall Methodology

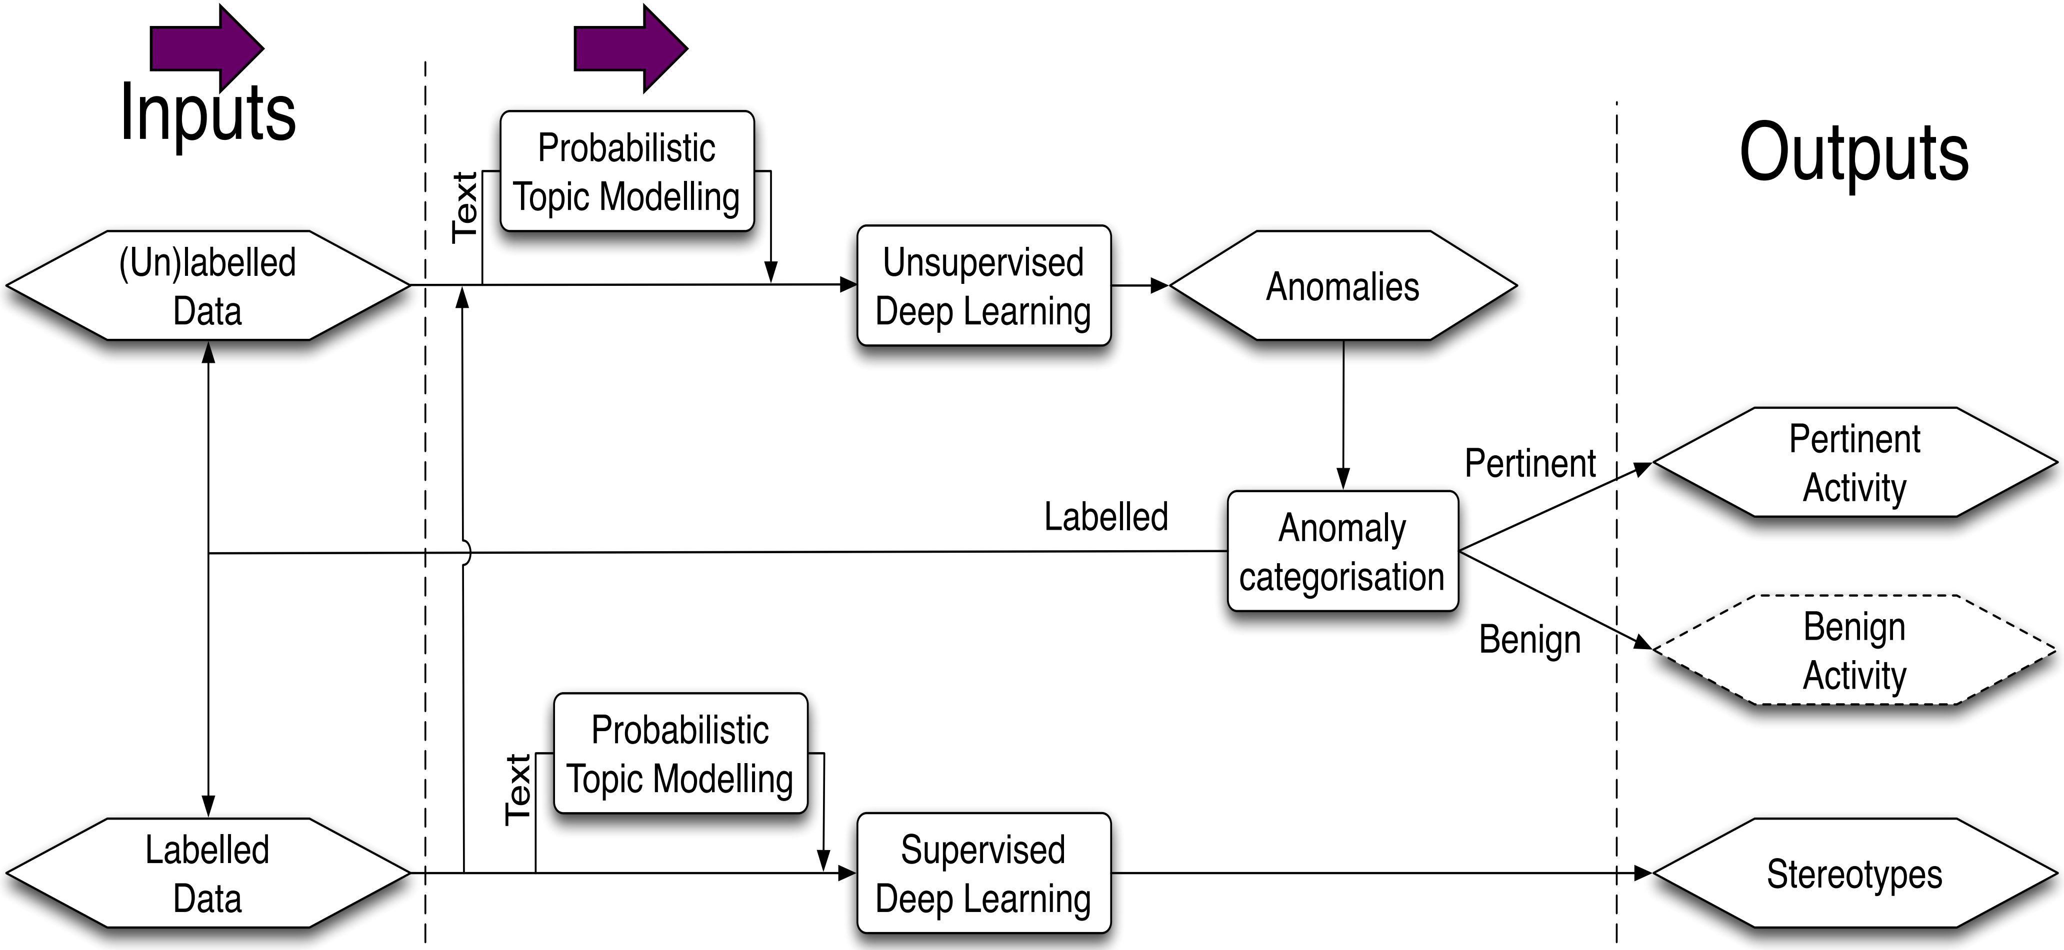

Overall Methodology

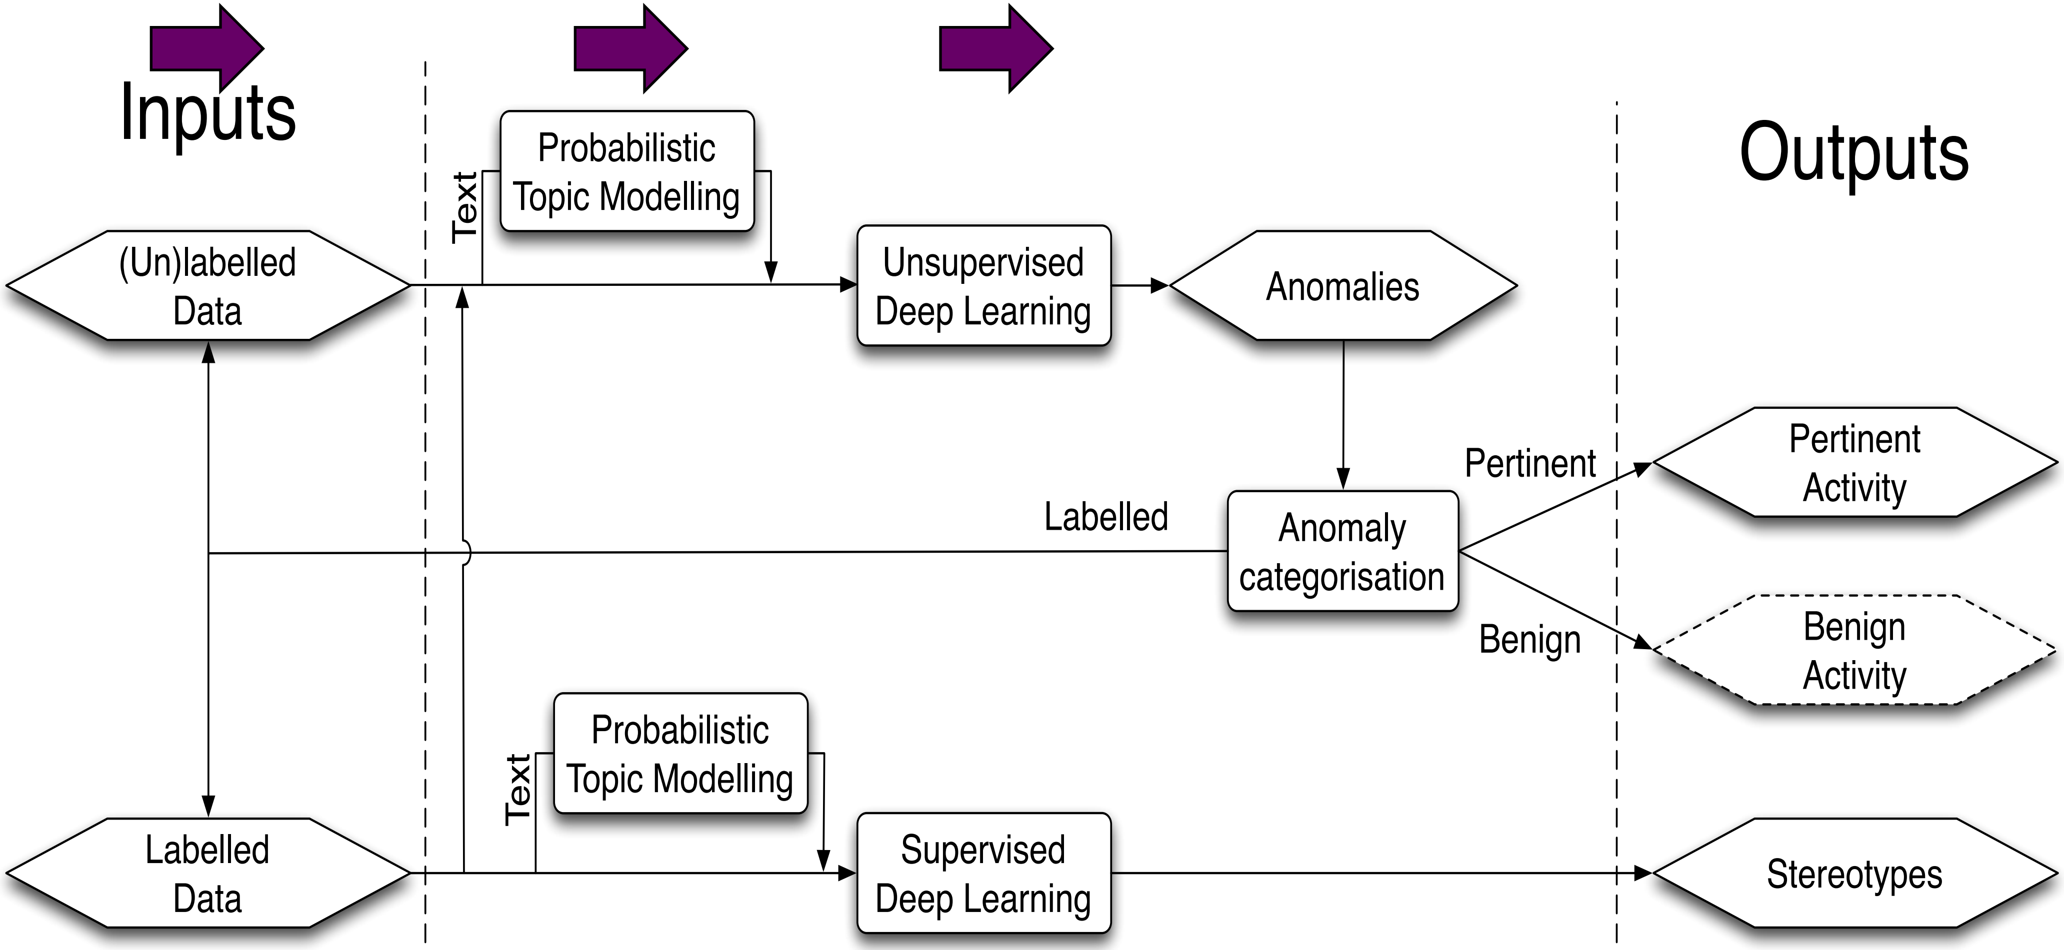

Overall Methodology

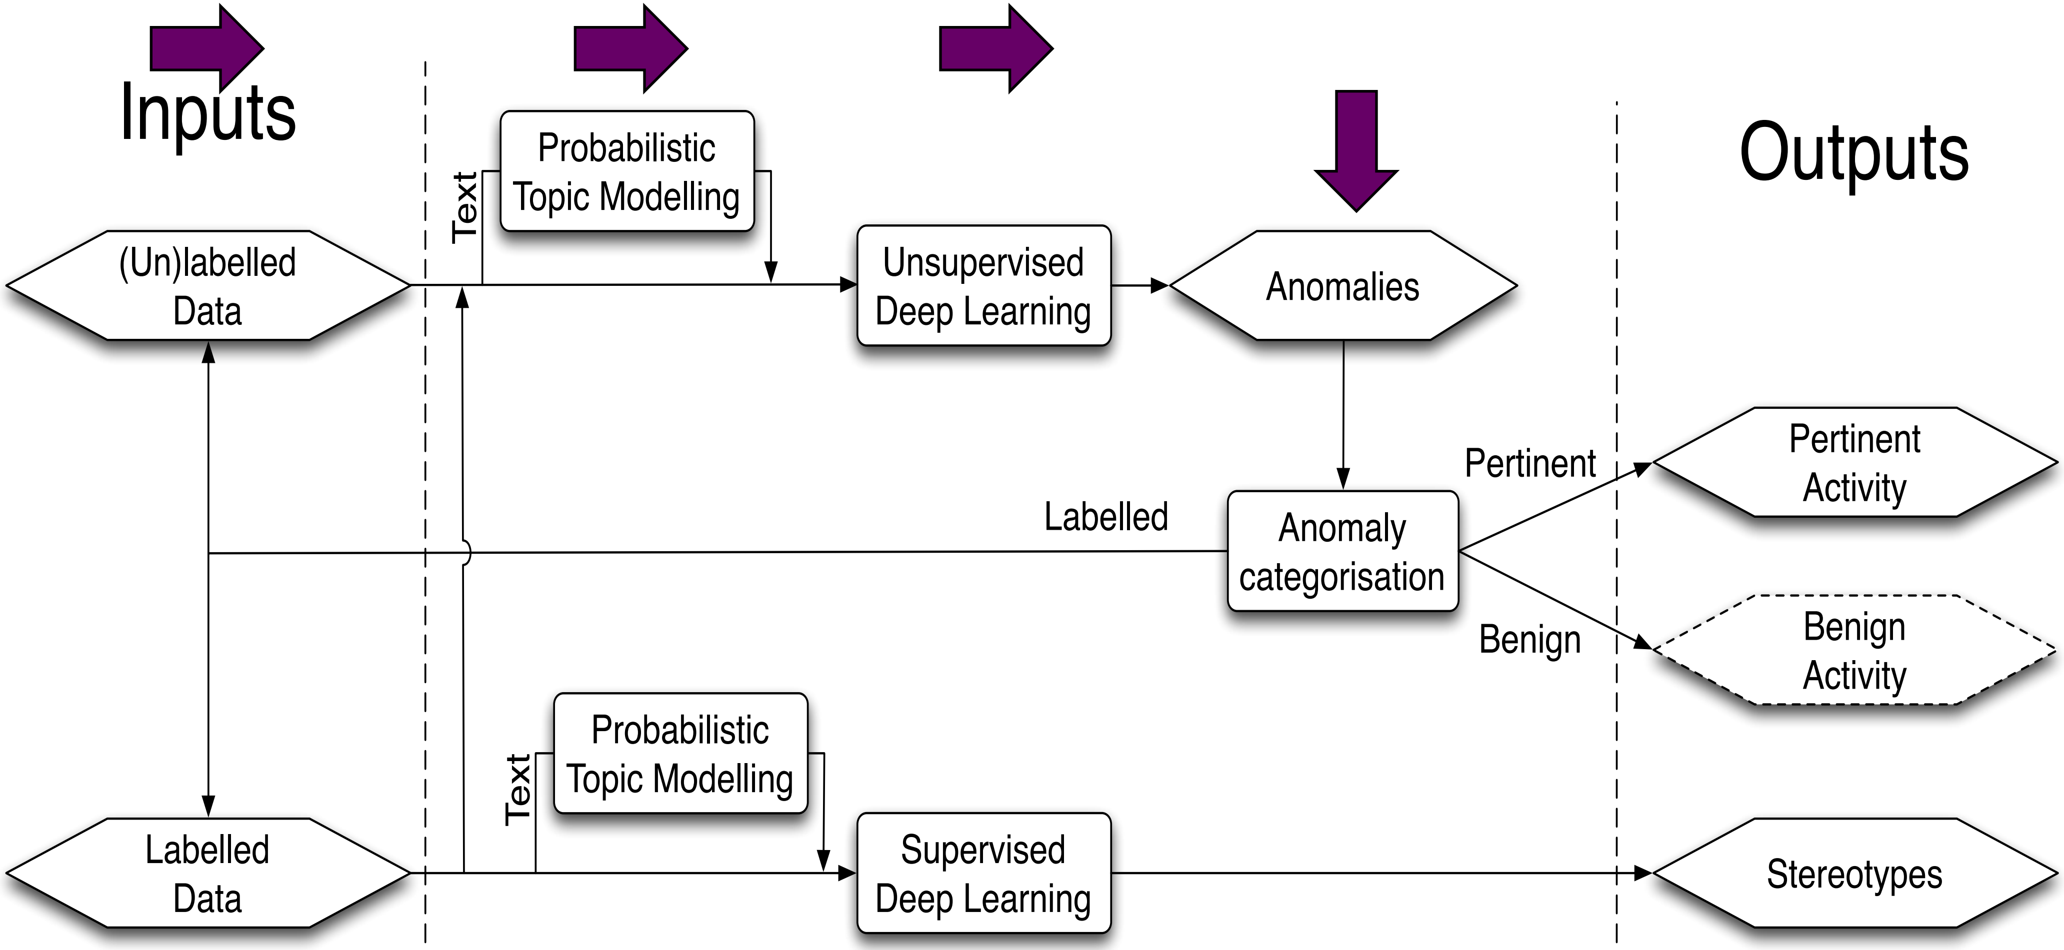

Overall Methodology

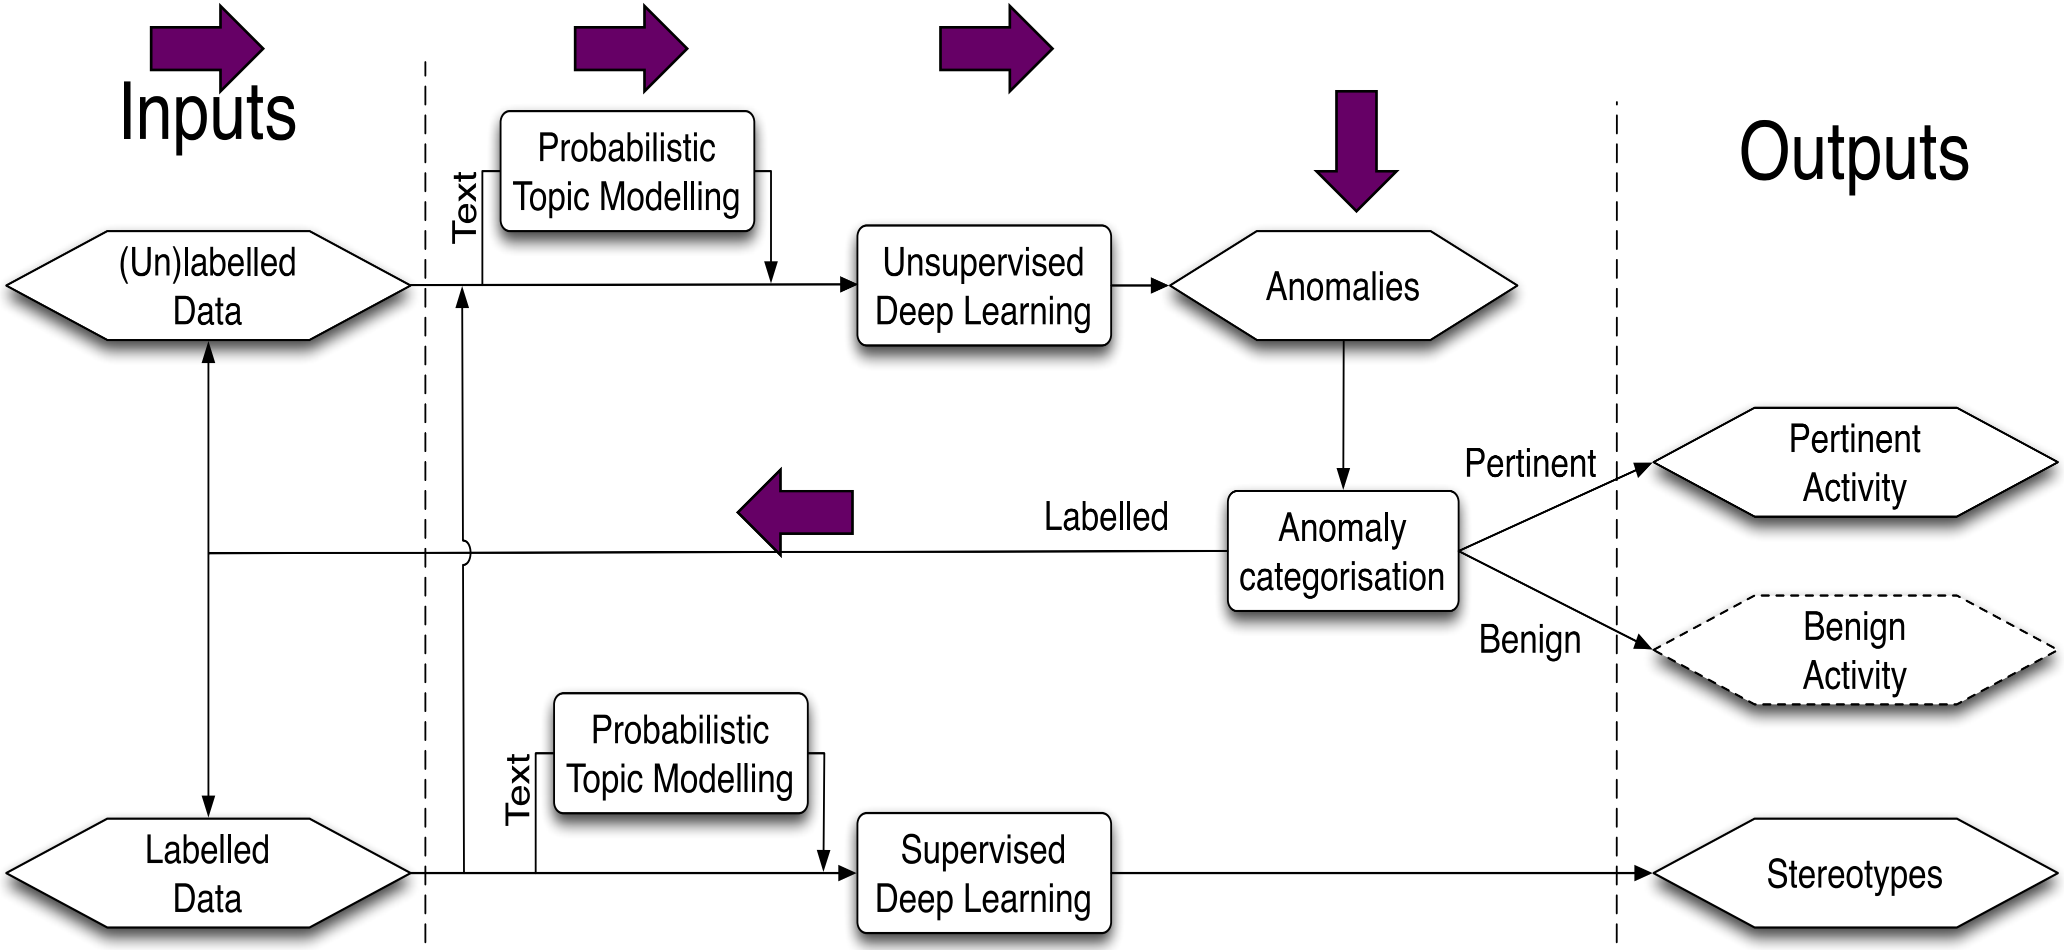

Overall Methodology

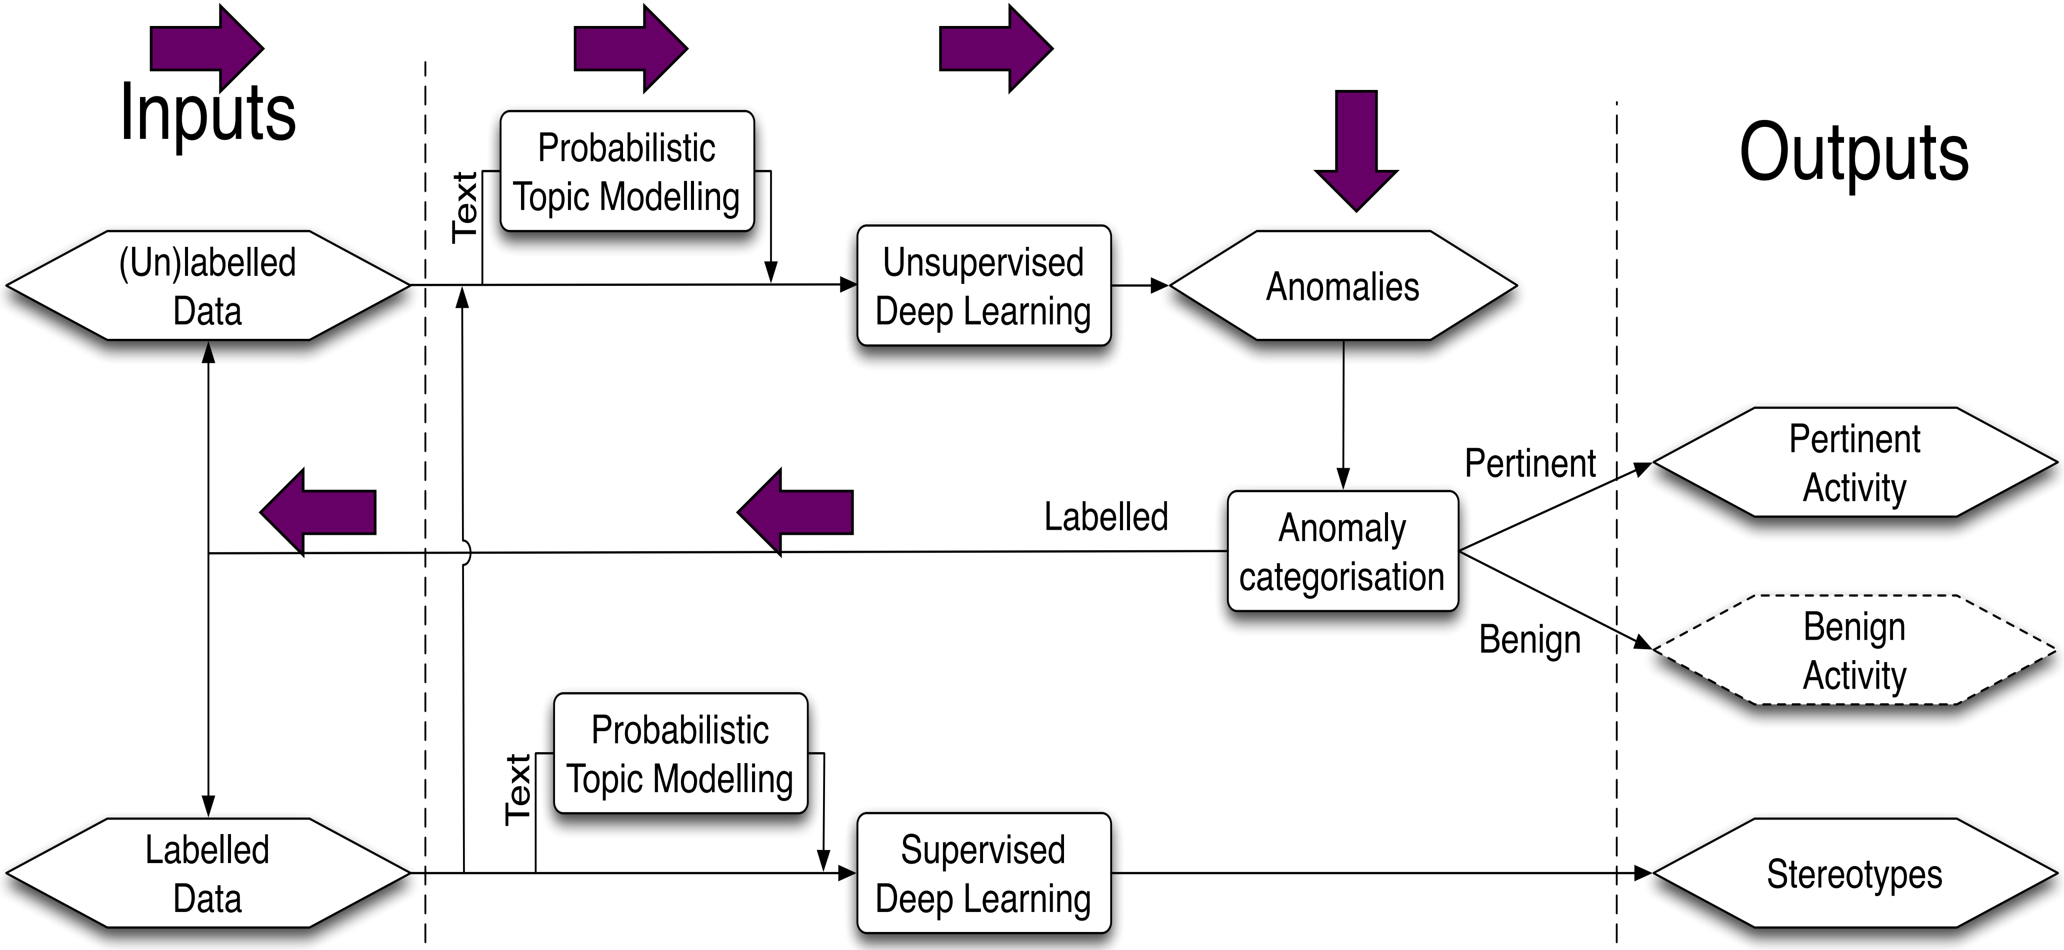

Overall Methodology

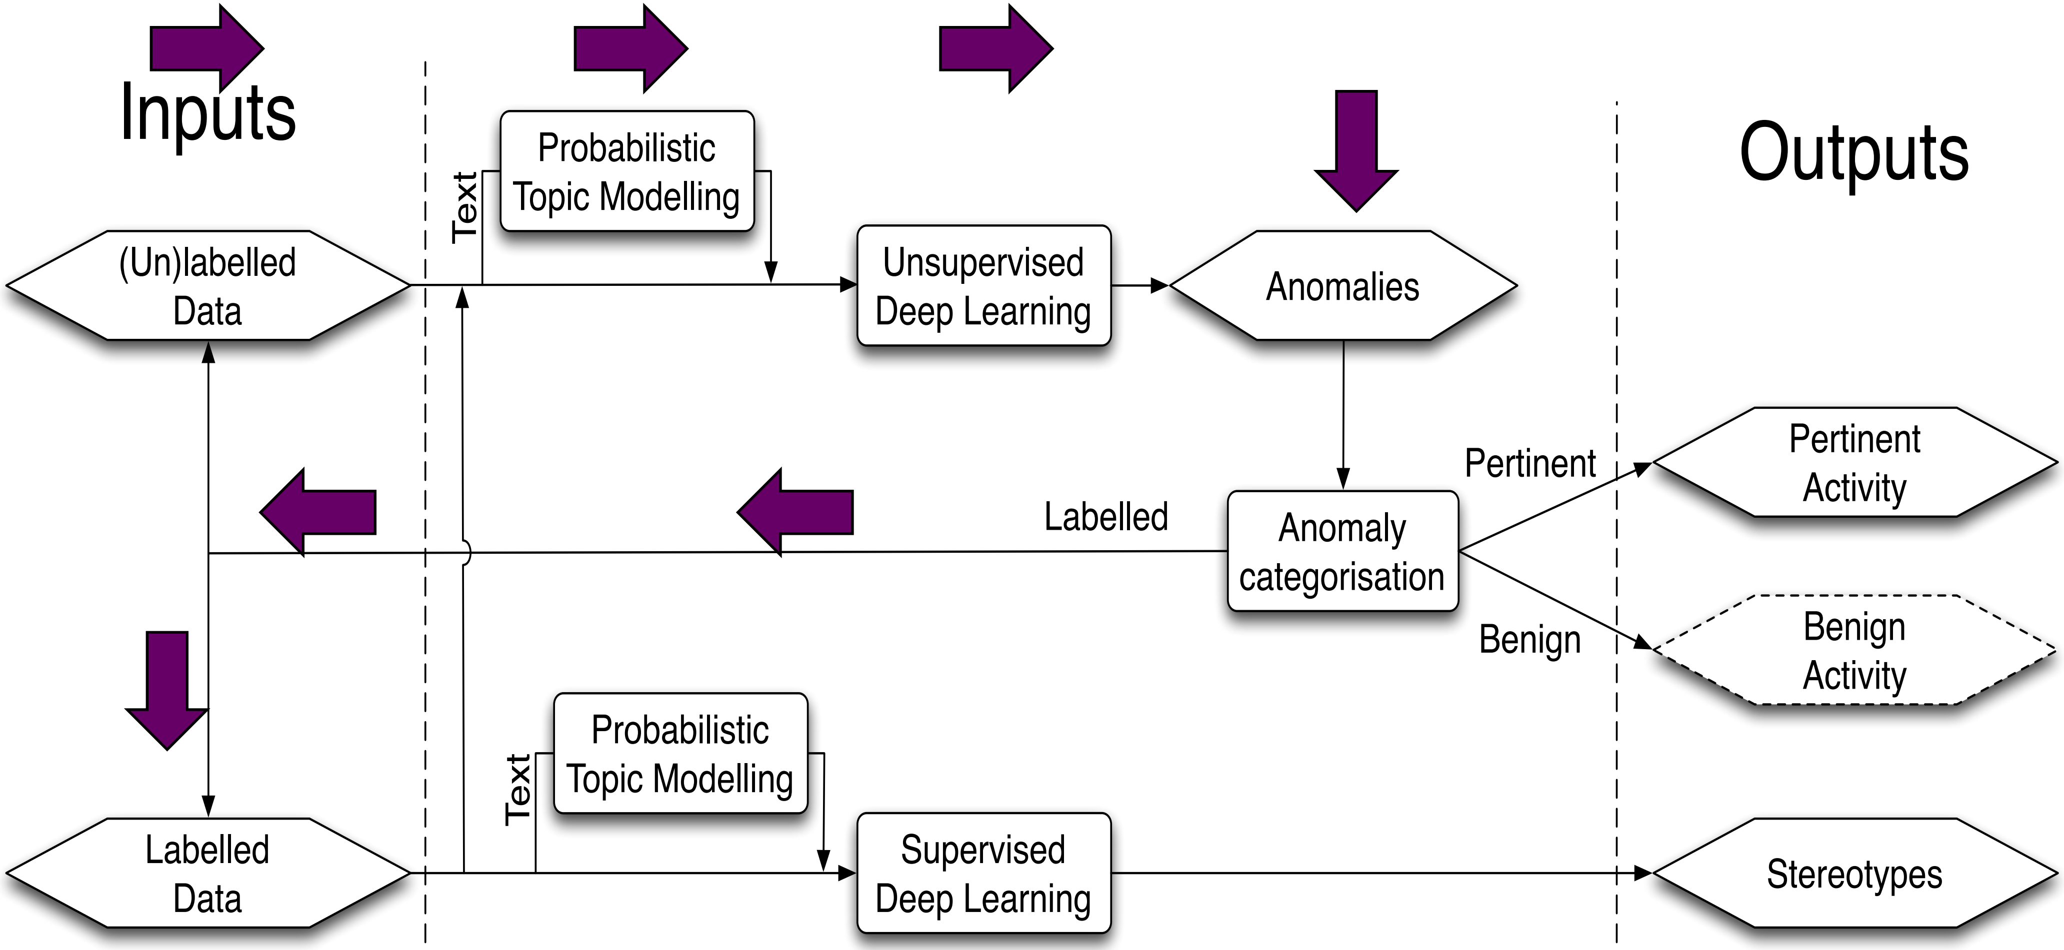

Overall Methodology

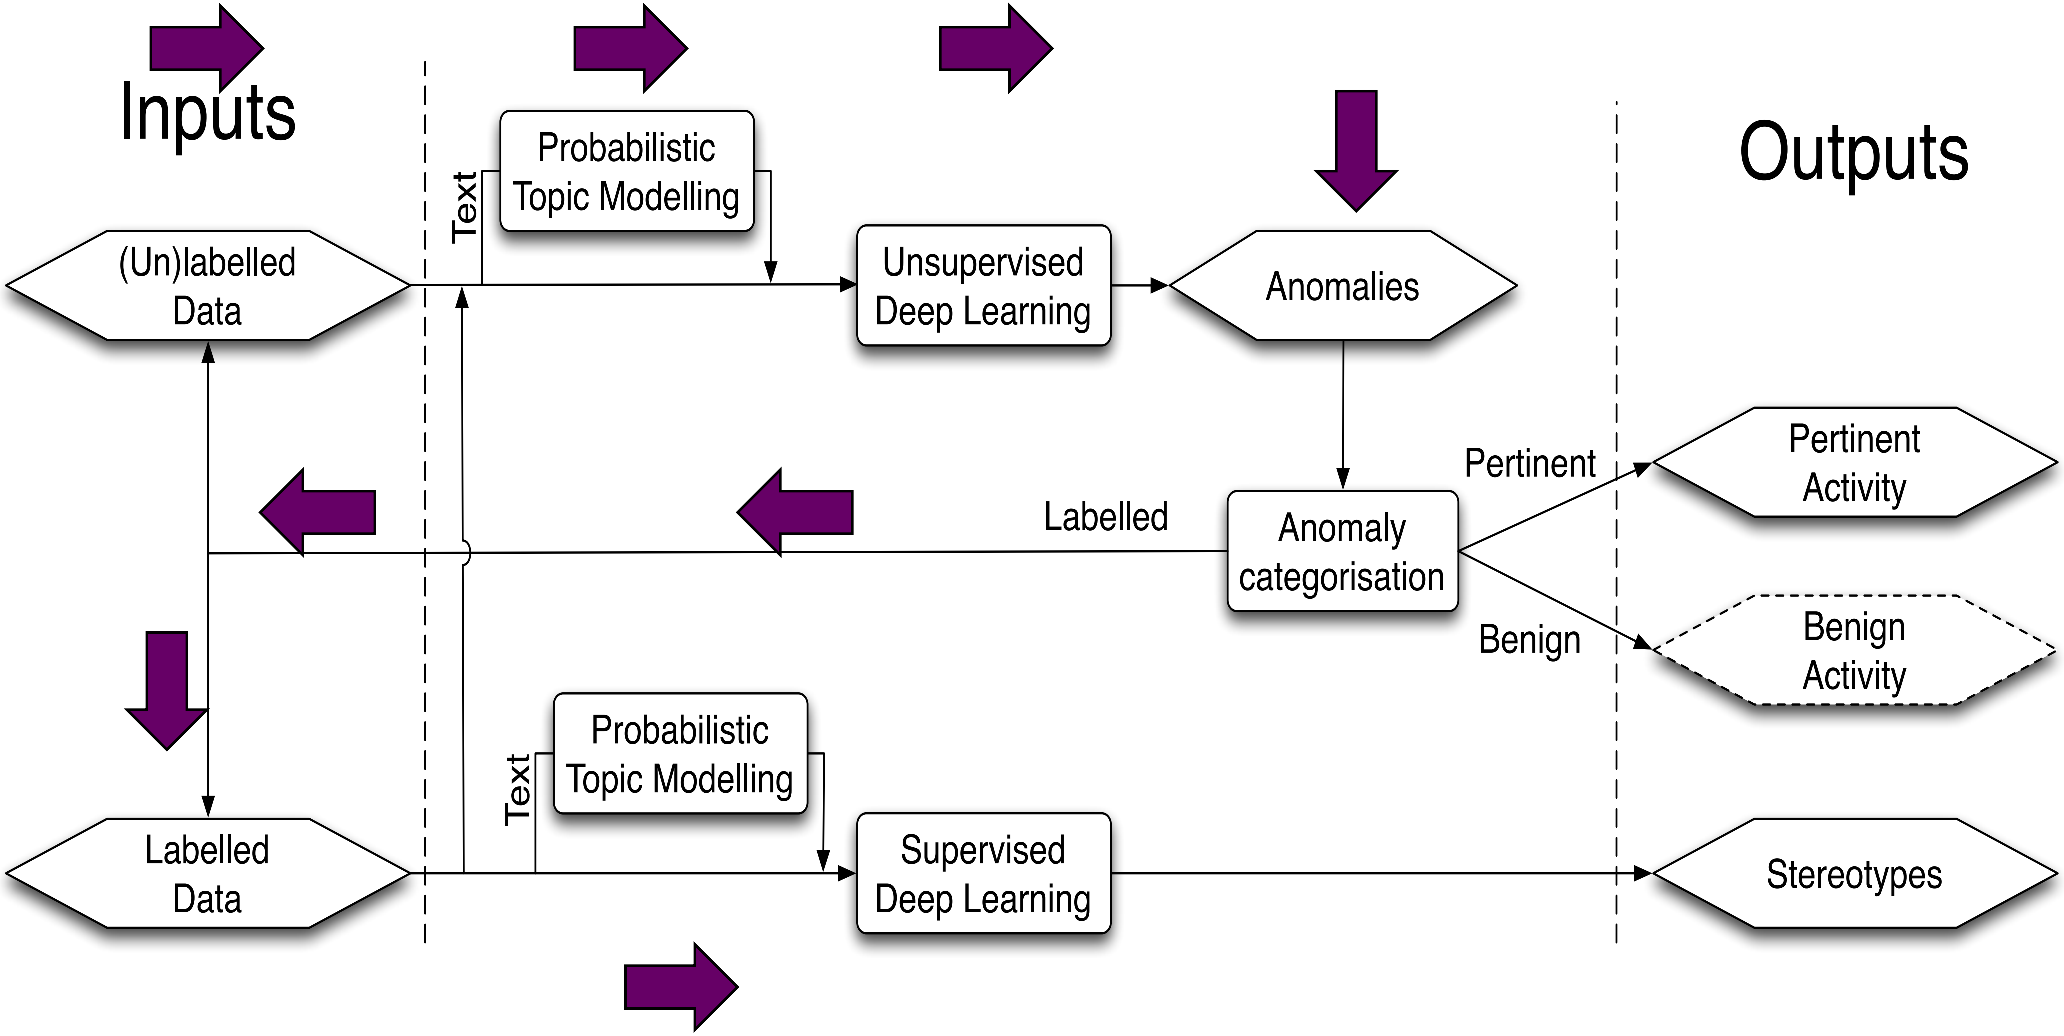

Overall Methodology

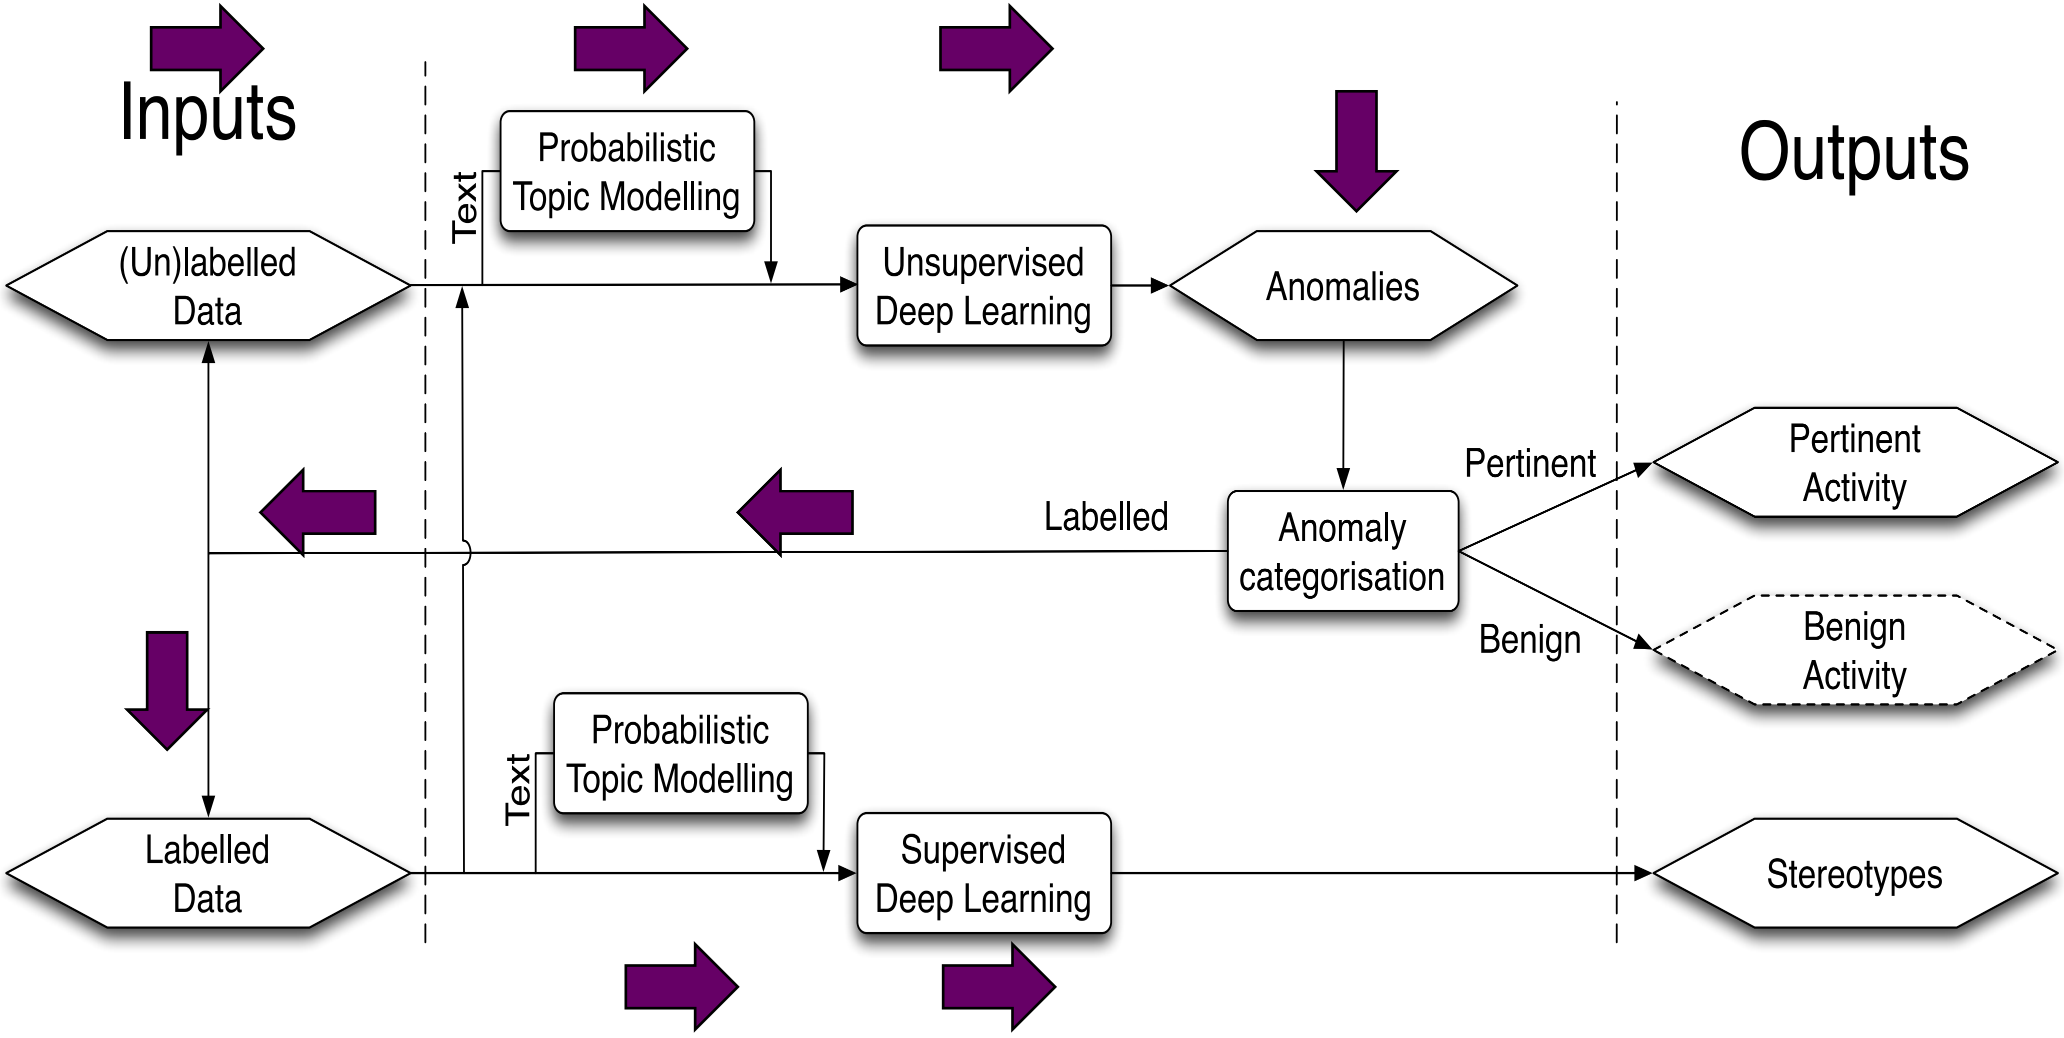

Overall Methodology

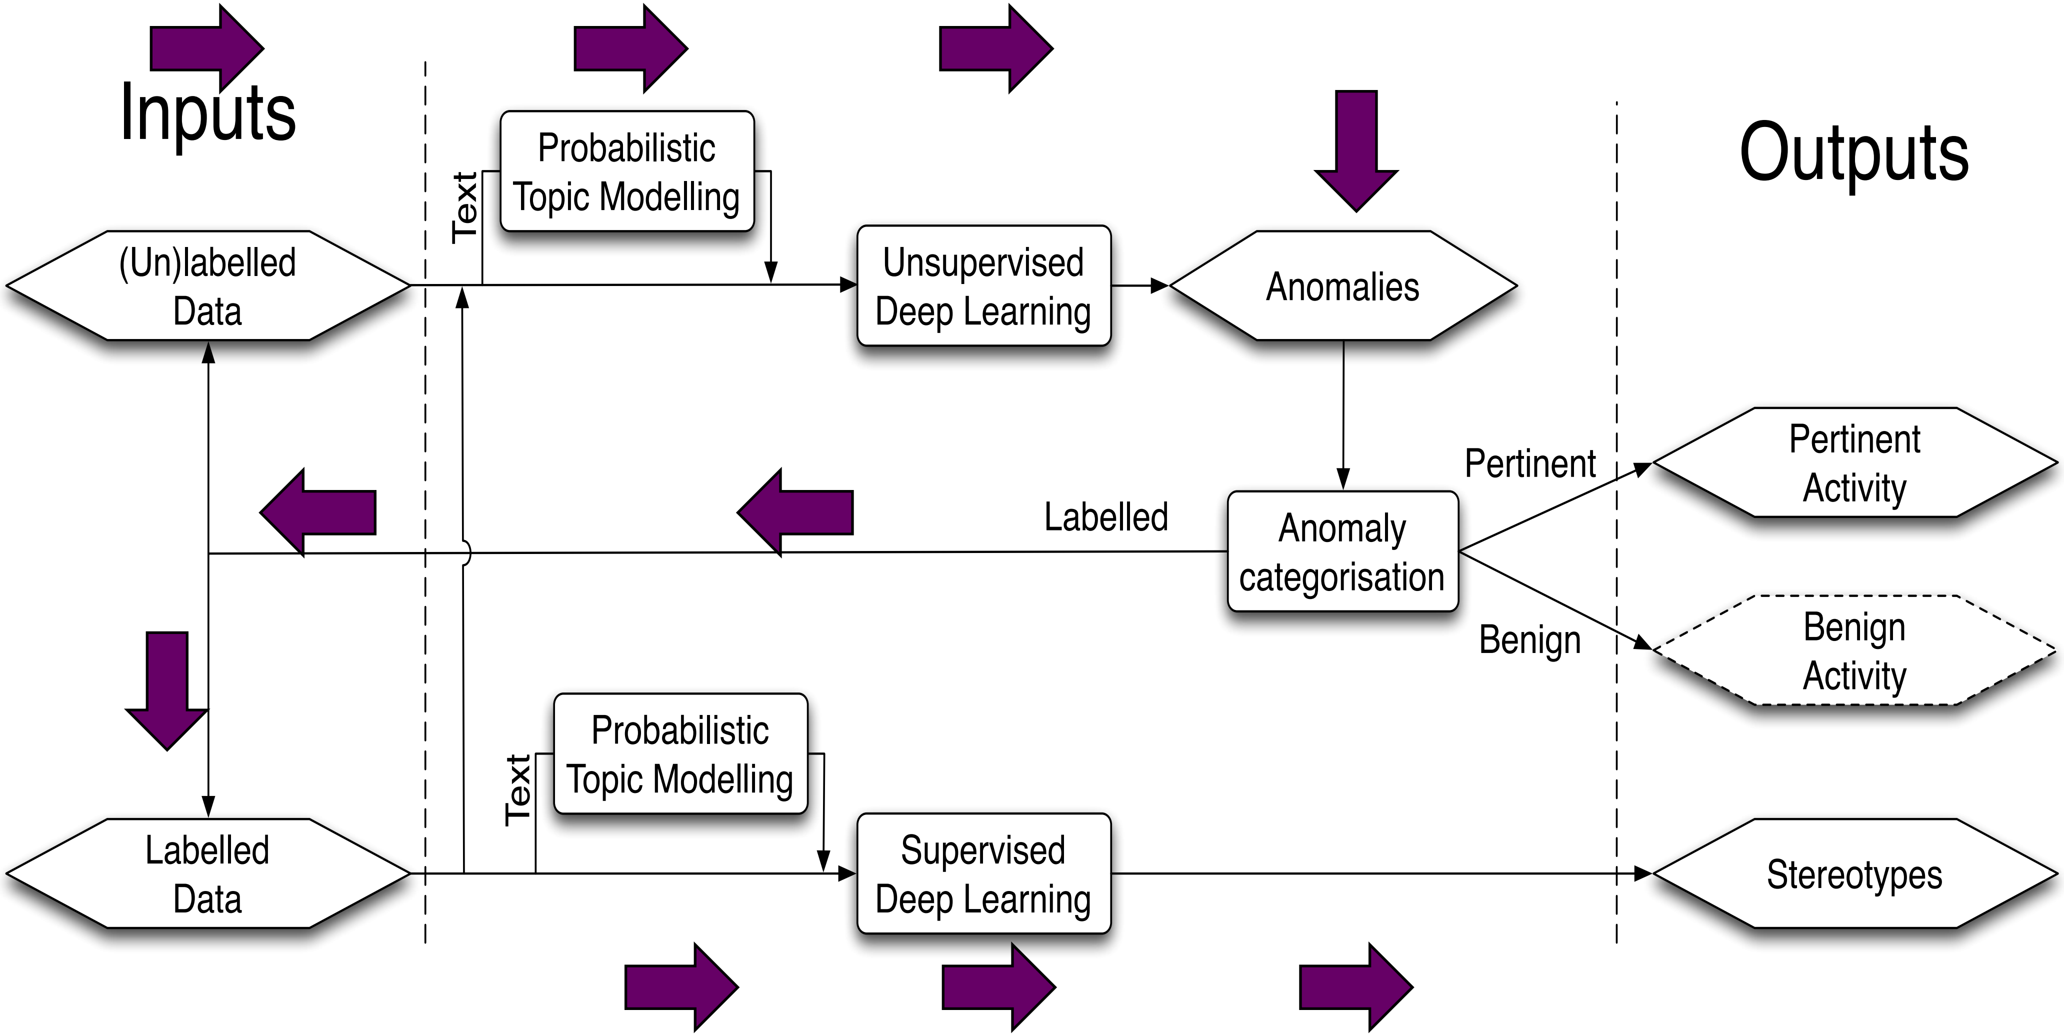

Example: Anomaly Identification

SPAM and HAM in SMS

- Auto-identification of SPAM from HAM in SMS messages

- 5574 SMS messages processed
- 4827 HAM messages
- 747 SPAM messages

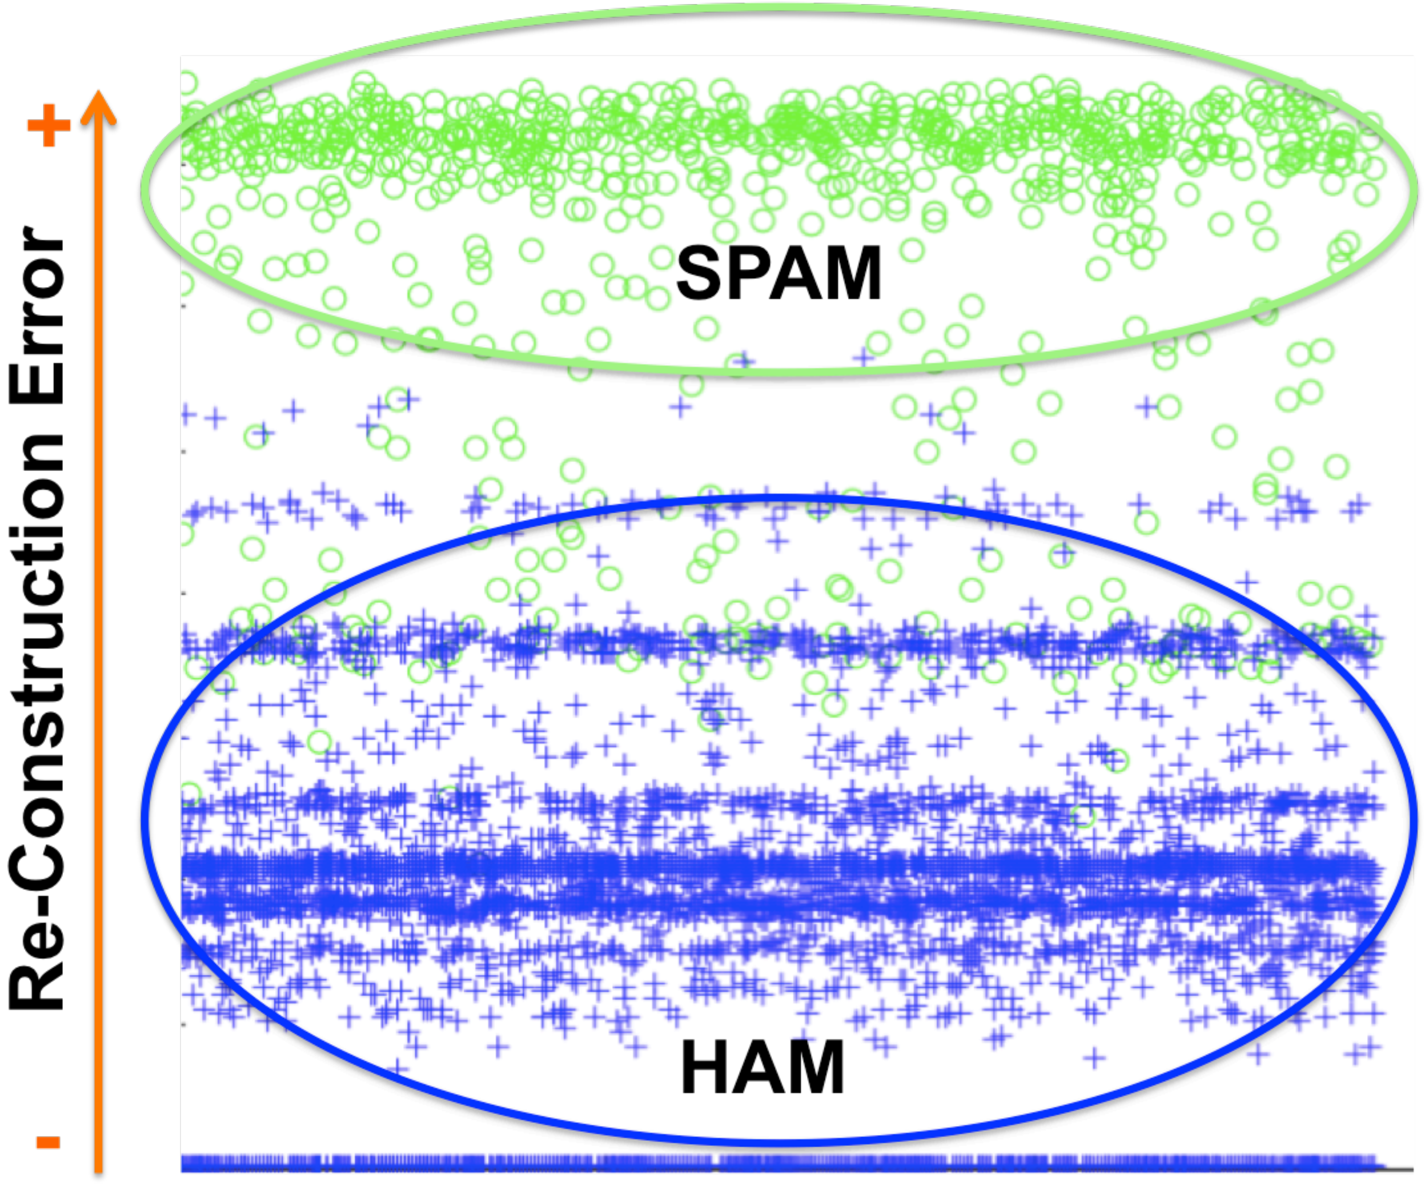

Anomaly Identification

SPAM and HAM in SMS

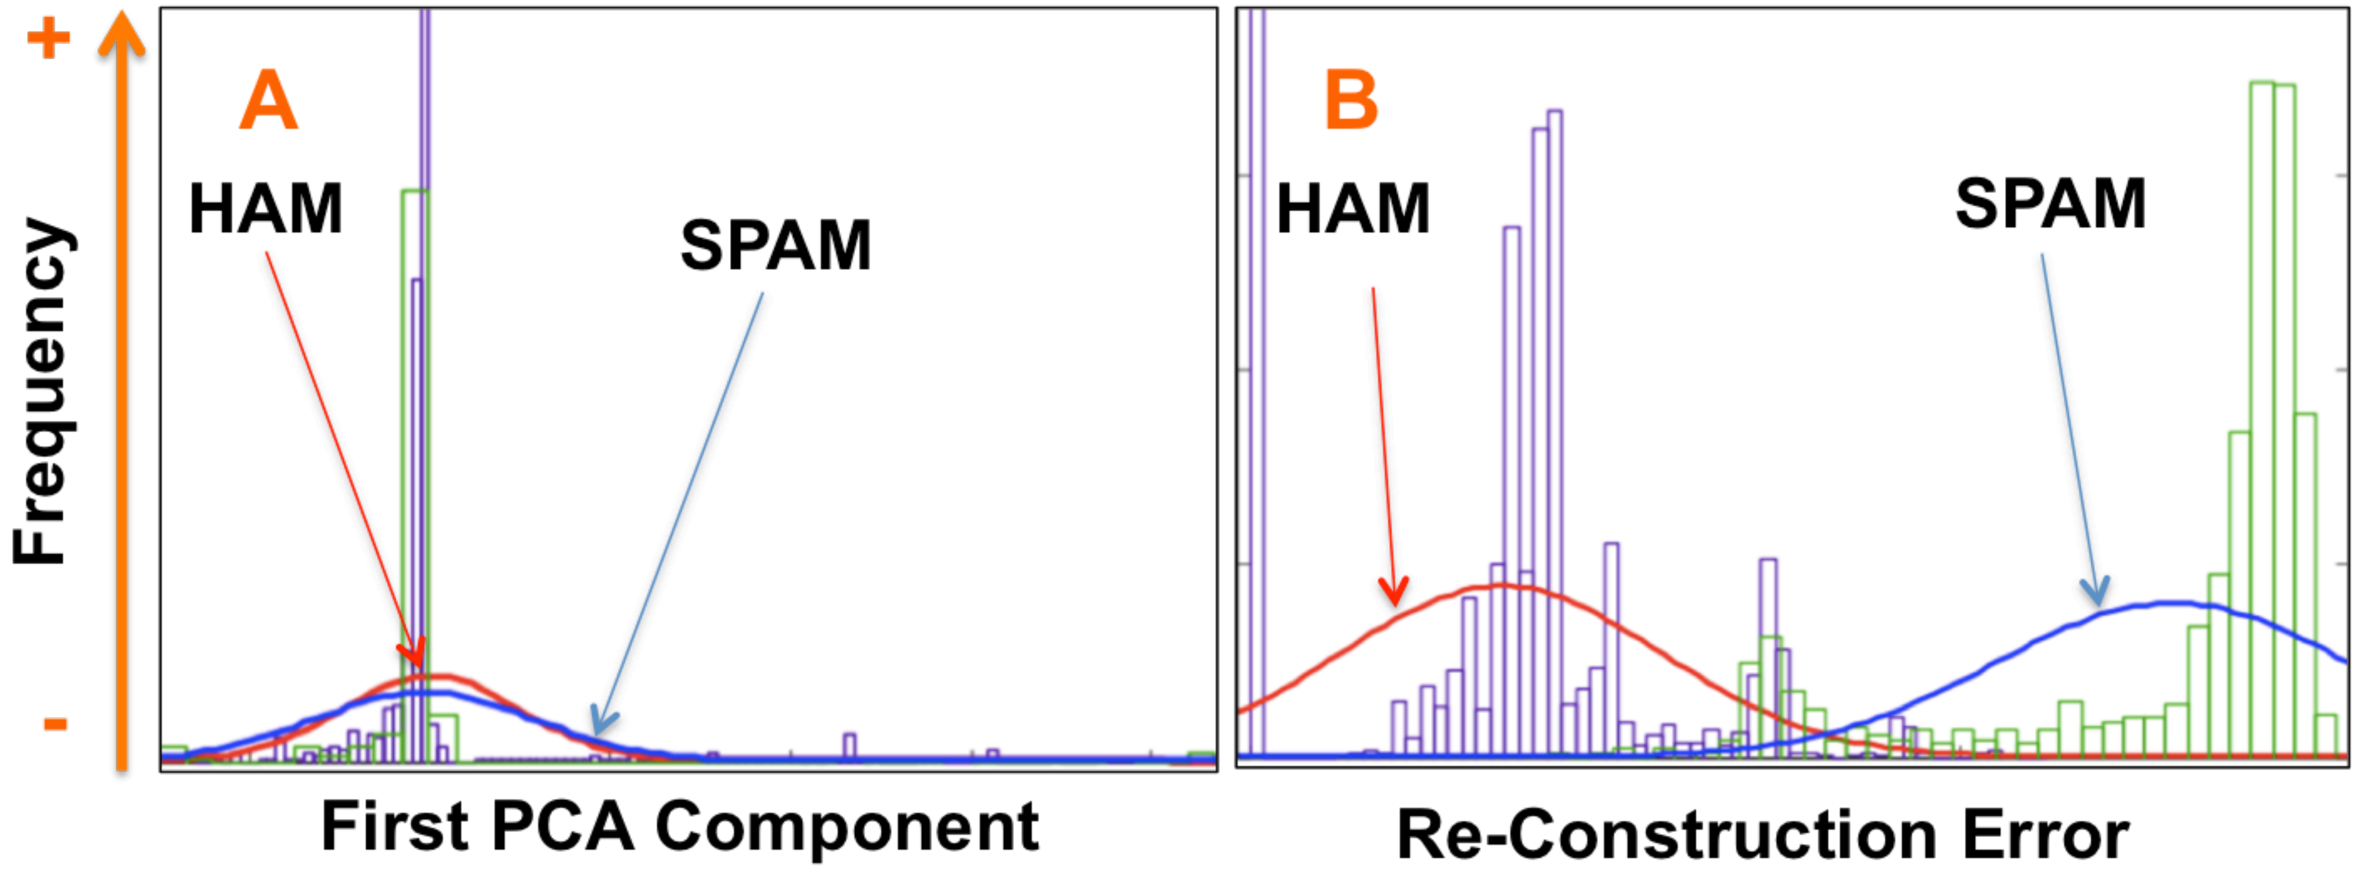

Comparison

SC% - SPAM Caught
BH% - Blocked HAM
Acc% - Accuracy
MCC% - Mathews Correlation Coefficient

Classifier	SC%	BH%	Acc%	MCC%
TM+SDA	85.59	0.62	97.51	0.899
Logistic Reg. + tok2	95.48	2.09	97.59	0.899
SVM + tok1	83.10	0.18	97.64	0.893
Boosted NB + tok2	84.48	0.53	97.50	0.887
SMO + tok2	82.91	0.29	97.50	0.887
Boosted C4.5 + tok2	81.53	0.62	97.05	0.865
MDL + tok1	75.44	0.35	96.26	0.826
PART + tok2	78.00	1.45	95.87	0.810
Random Forest + tok2	65.23	0.12	95.36	0.782
C4.5 + tok2	75.25	2.08	95.00	0.770
Bern NB + tok1	54.03	0.00	94.00	0.711
MN TF NB + tok1	52.06	0.00	93.74	0.697
MN Bool NB + tok1	51.87	0.00	93.72	0.695
1NN + tok2	43.81	0.00	92.70	0.636
Basic NB + tok1	48.53	1.42	92.05	0.600
Gauss NB + tok1	47.54	1.39	91.95	0.594
1Flex NB + tok1	47.35	2.77	90.72	0.536
Boolean NB + tok1	98.04	26.01	77.13	0.507
3NN + tok2	23.77	0.00	90.10	0.462
EM + tok2	17.09	4.18	85.54	0.185
TR	0.00	0.00	86.95	-

Performance

- Approach is computationally intensive
- Need to reduce execution time to tractable level
- Use of GPGPUs to improve the performance of the framework
- Have been used previously with Deep Learning showing significant benefits
- But focused on Dense Data (images / sound)
- This is a sparse data problem

Execution Time

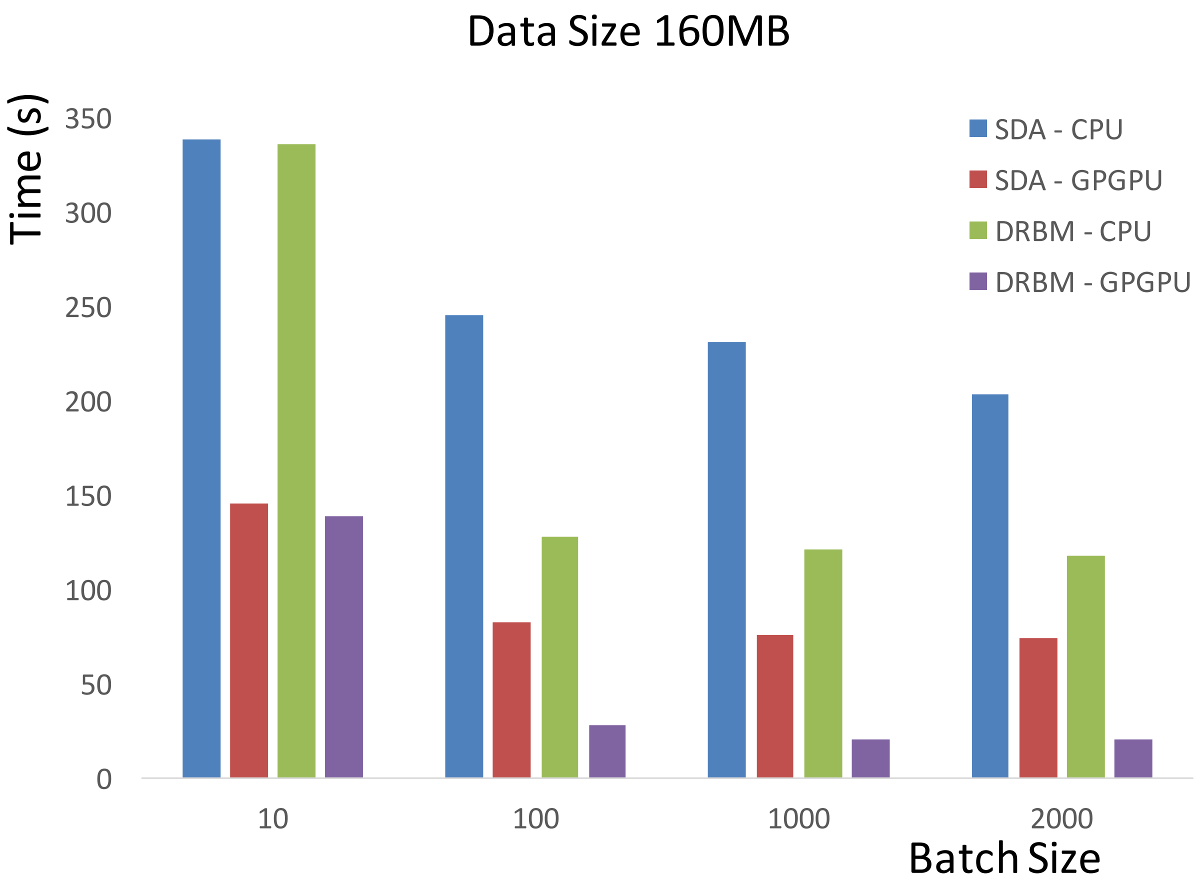

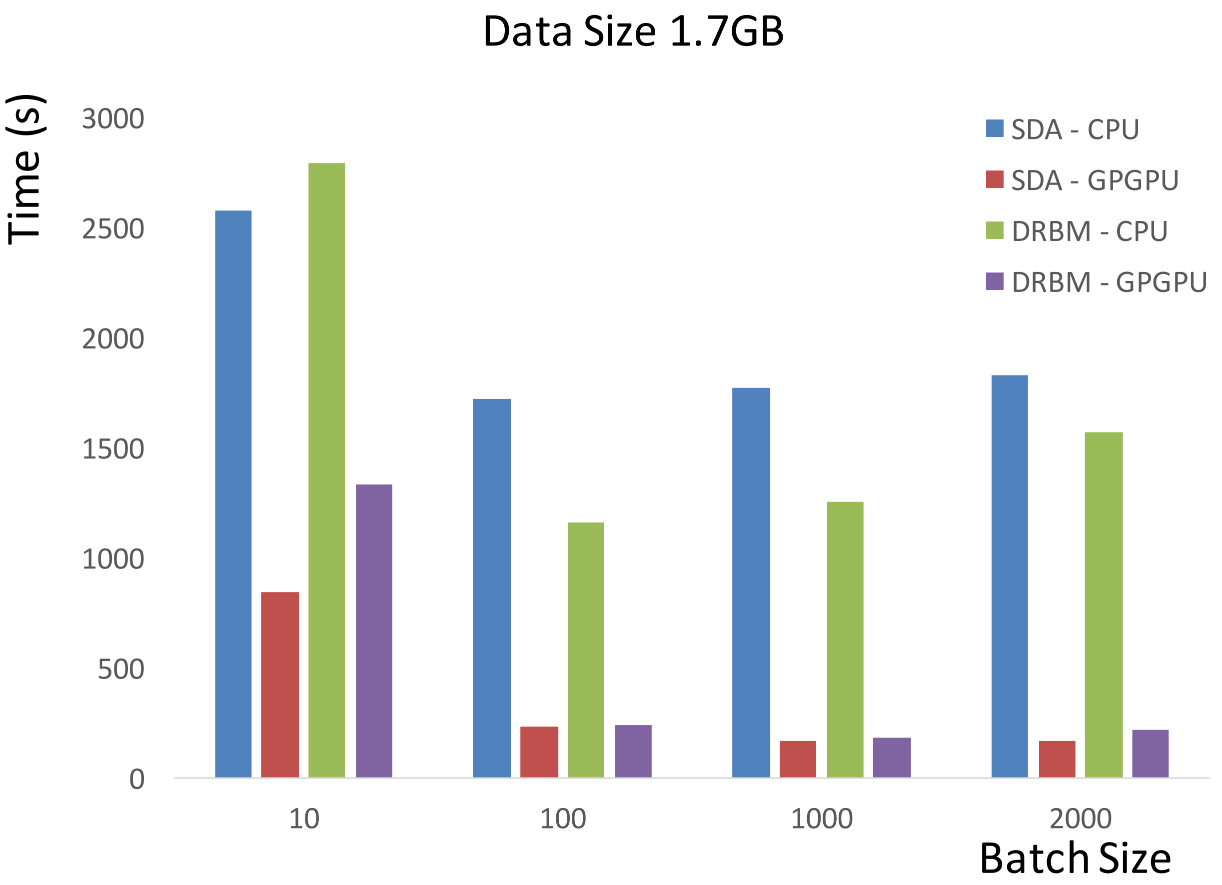

System: Intel Xeon E5-2650 v3 2.3GHz, 64GB RAM, 2 x 300GB 15k RPM SAS
GPU: NVIDIA K40
Dataset: Electric meter readings from Ireland

Speedup

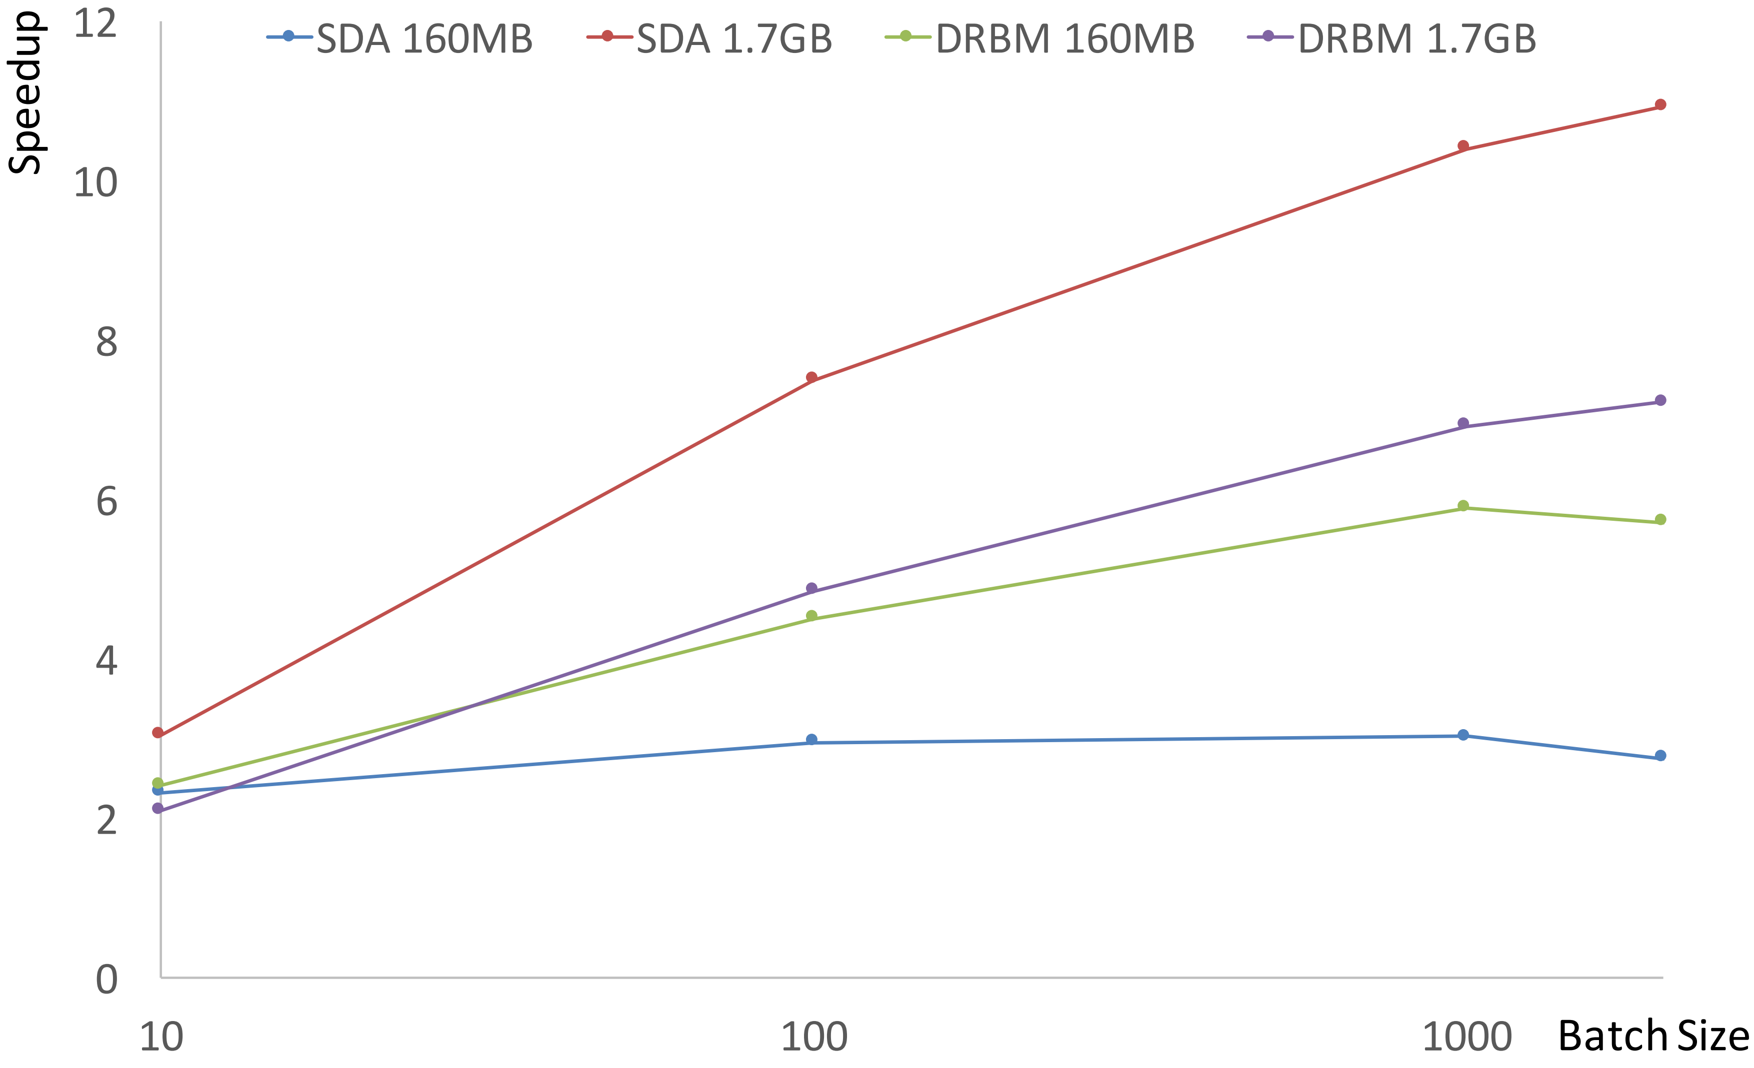

Speedup
12
10
8
6
4
2
0
SDA 160MB
SDA 1.7GB
DRBM 160MB
DRBM 1.7GB
10
100
1000
Batch Size

Possible Applications:

- Terrorist activity tracking
  - Acting out of character, predicting activity
- Ship/Flight tracking data
  - Hijacking, Flight deviations
- Police crime database
  - Criminal profiling, acting out of character
- Unwanted information release
  - Topic changes, specific damaging subjects,(e.g Wiki Leaks)
- Student applications
  - Identifying bogus attempts for visa
- Social media Tracking
  - Social grooming, political persuasion
- Safety camera tracks
  - Normal movements of people in area
- Illegal financial transactions
  - Fraud, laundering

A. Stephen McGough, Noura Al Moubayed, Jonathan Cumming, Eduardo Cabrera, Peter Matthews,
Toby P. Breckon, Ed Ruck-Keene, Georgios Theodoropoulos
Durham University, UK
stephen.mcgough@durham.ac.uk