# BayesWatch: An Anomaly Detection System

Kevin Korb, Steven Mascaro and Ann Nicholson

Email: {kbkorb@gmail.com, steven.mascara@bayesian-intelligence.com, ann.nicholson@monash.edu}

BAYESIAN INTELLIGENCE

# Anomaly Detection

Outlier:

A data point lying outside the norm

- Either unlucky
- or as a part of some *change*
- or it's a true anomaly

Anomaly detection aims at finding the last, or the last two.

# Applications

- Fraud detection
- Intrusion, security
- Fault detection
- Terrorism detection
- Crime identification
- Actually, it's an intrinsic aspect of data analysis in all applications, whether ignored or not.

# Clustering Approach

- Cluster the data into groups of different density
- Choose points in small cluster as candidate outliers
- Compute the distance between candidate points and non-candidate clusters.
  - If candidate points are far from all other non-candidate points, they are considered outliers

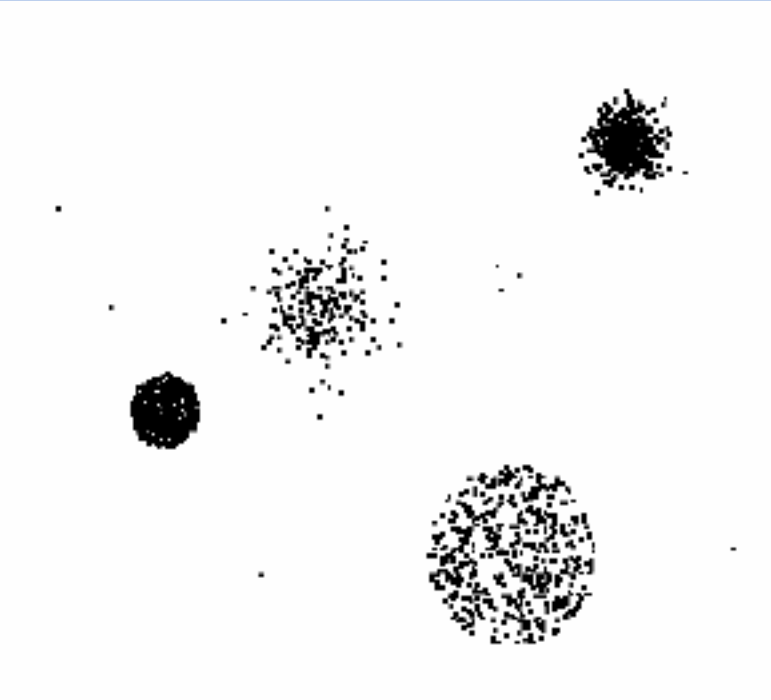

# Bayesian Net Approach

- Learn, or develop, a Bayesian net model of a system operating normally
- Compute the probability of a data point according to that model: P(x|BN)
- Either
  - If $P(x|BN) < \theta$, for some threshold, treat x as an anomaly
  - If $P(x|BN) < P(x|AN)$, for some anomaly model AN, treat x as an anomaly

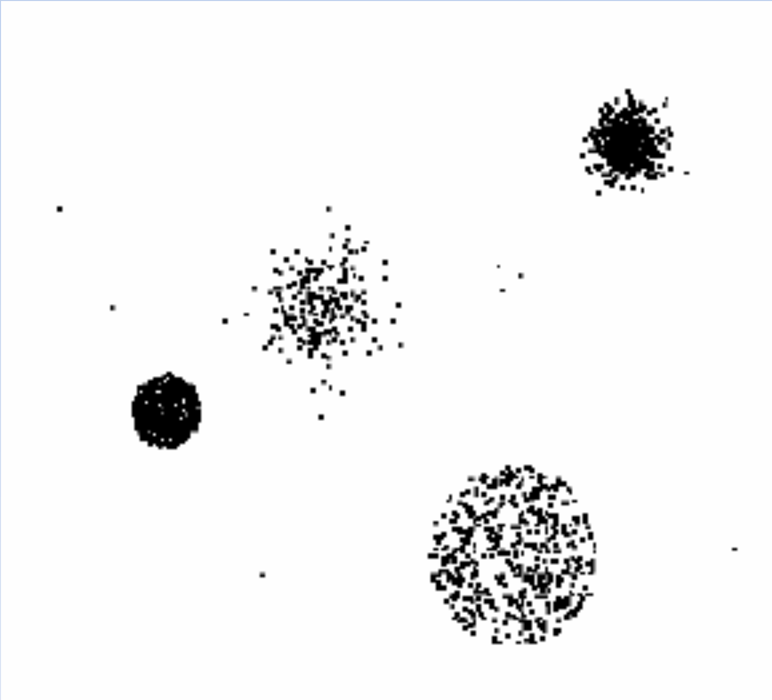

# Bayesian Net Approach

- Learn, or develop, a Bayesian net model of a system operating normally
- Compute the probability of a data point according to that model: P(x|BN)
- Either
  - If P(x|BN) < θ, for some threshold, treat x as an anomaly
  - If P(x|BN) < P(x|AN), for some anomaly model AN, treat x as an anomaly – This should be the goal!

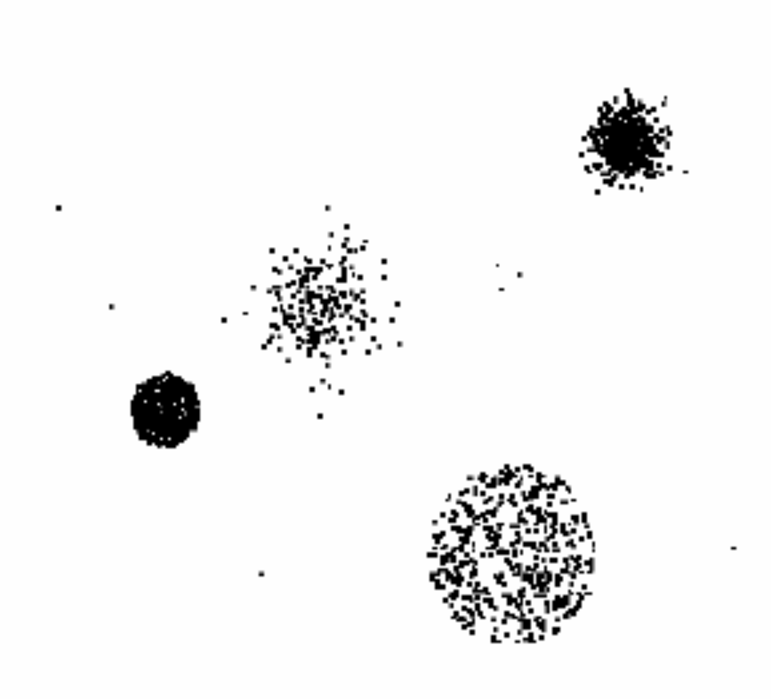

# BayesAnomalous: Anomaly detection in vessel tracks

- Mascaro, S., Korb, K.B., and Nicholson, A. E., Anomaly Detection in Vessel Tracks using Bayesian networks, To appear in Proc. of the 8th Bayesian Modeling Applications Workshop, Barcelona, Spain, July 14, 2011.

- Research contract with Australian "Defense Science and Technology Organisation" (DSTO)

# Anomaly detection in vessel tracks

- Use AIS data from vessels around Sydney Harbour
- Apply causal discovery (CaMML) to learn models of behavior
- Use likelihoods to identify anomalous tracks

# The AIS Data

- AIS data from May 1st to July 31st, 2009
- For a section of the NSW coast framing Sydney harbour

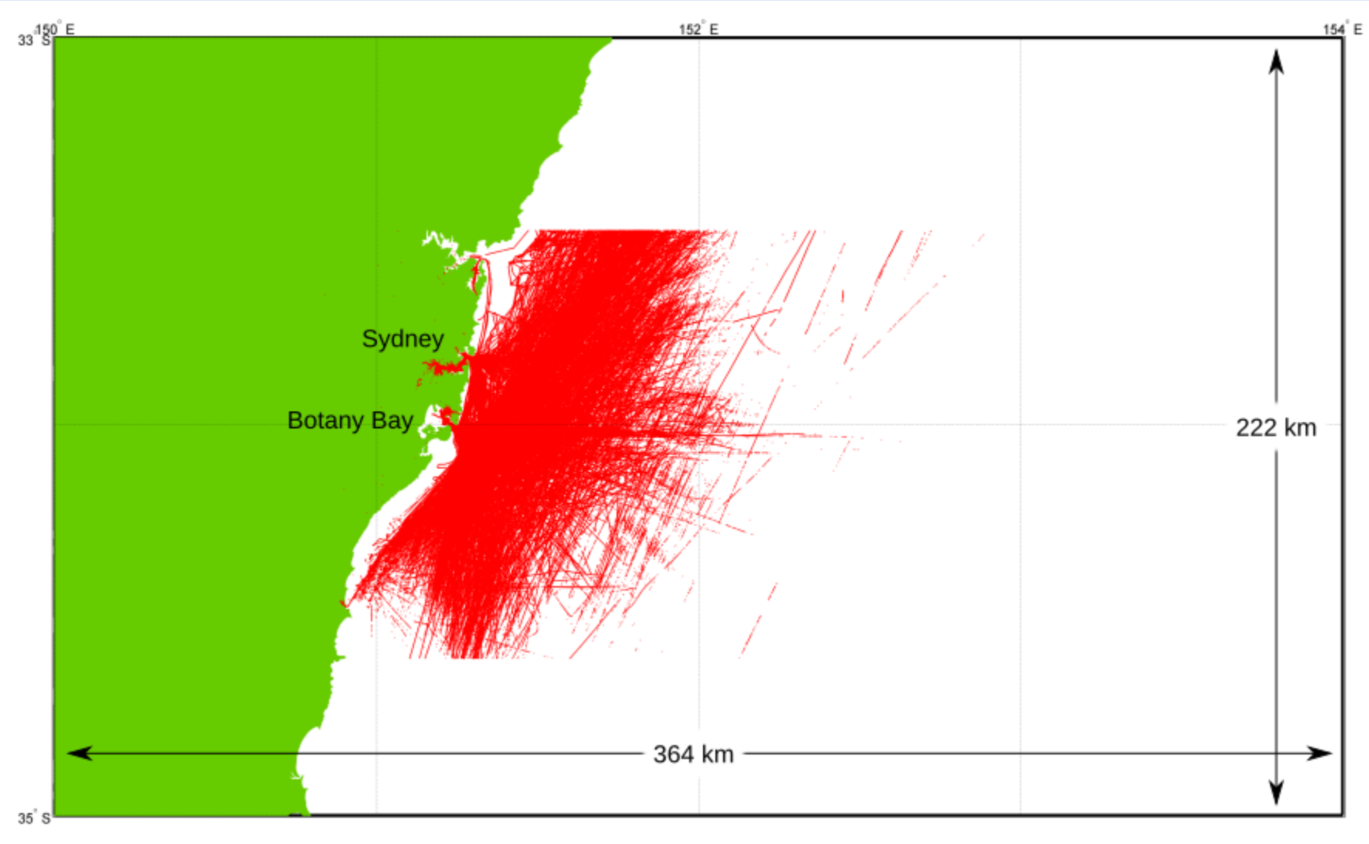

# The AIS data

- the vessel's MMSI (nine digit vessel ID)
- a timestamp
- the latitude and longitude of the vessel
- its reported speed, course and heading

| MMSI | Timestamp | Lat | Lon | Speed | Course | Hdng |
|---|---|---|---|---|---|---|
| X | 200905X | -33.X | 151.X | 18.7 | 49.9 | 46 |
| X | 200905X | -34.X | 151.X | 2.1 | 218 | 80 |
| X | 200905X | -33.X | 151.X | 0 | 0 | 511 |
| X | 200905X | -34.X | 151.X | 17.5 | 183 | 179 |
| X | 200905X | -33.X | 151.X | 1.2 | 28 | 64 |

- 9.2 million rows → 2,473 tracks across 544 unique MMSIs averaging 1,995 rows each

# Additional data

- Ship information (e.g. type, dimensions, weight) from marinetraffic.com, digital-seas.com & DSTO
- Weather (temperature, cloud cover, wind speed) from Bureau of Meteorology
- Natural temporal factors (hour of day, time since dawn/dusk
- Kinematic DBN nodes (e.g. speed)
- Info on vessel interactions

# The workflow

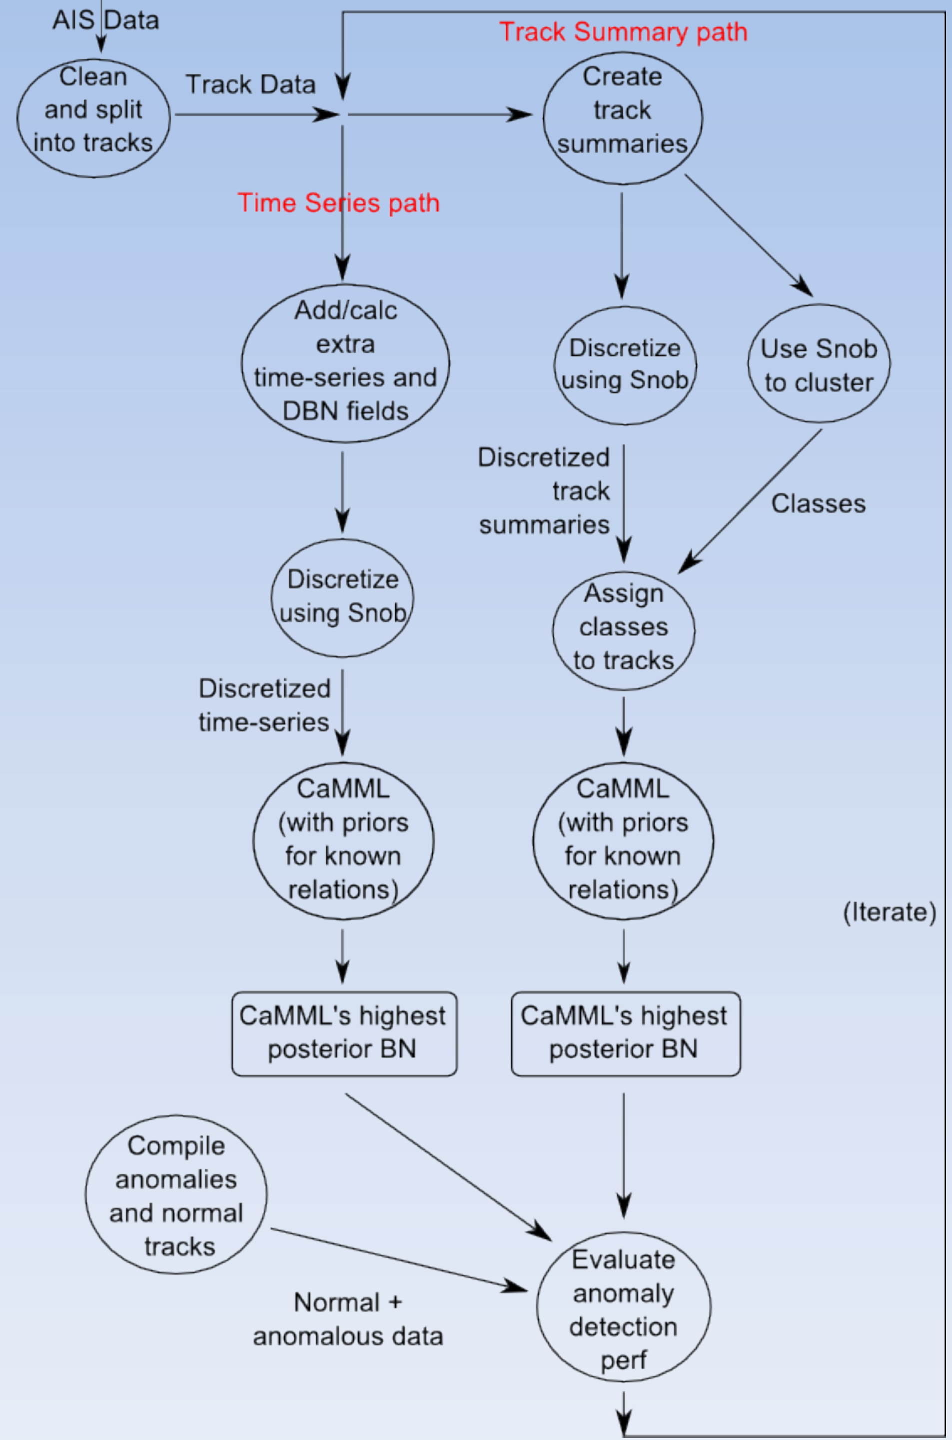

# Two main approaches

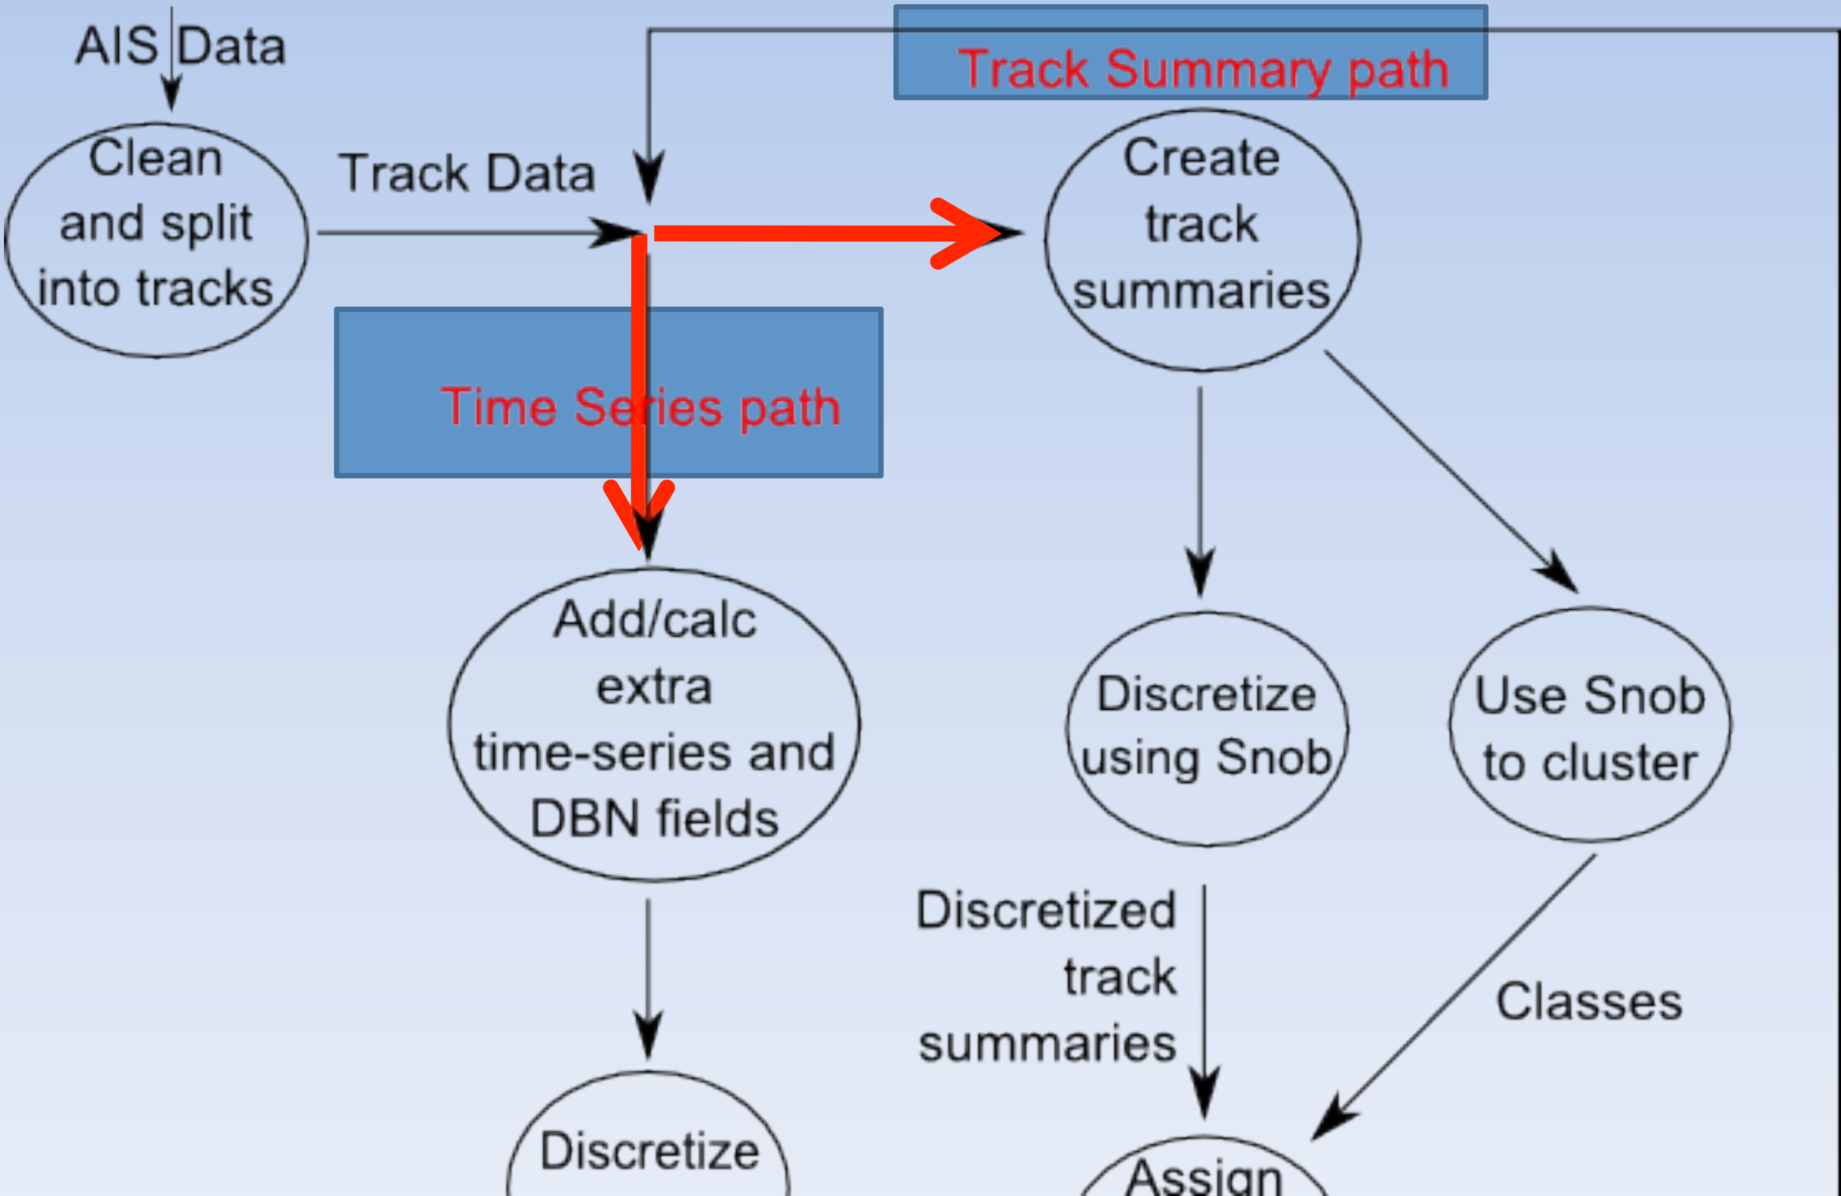

# Two approaches

1. Use original time series data, producing a DBN

2. Produce single summary records of each track, producing static BN
    - E.g. average speed and course, number of stops, major stopping points, % of time travelling straight

# Discretization

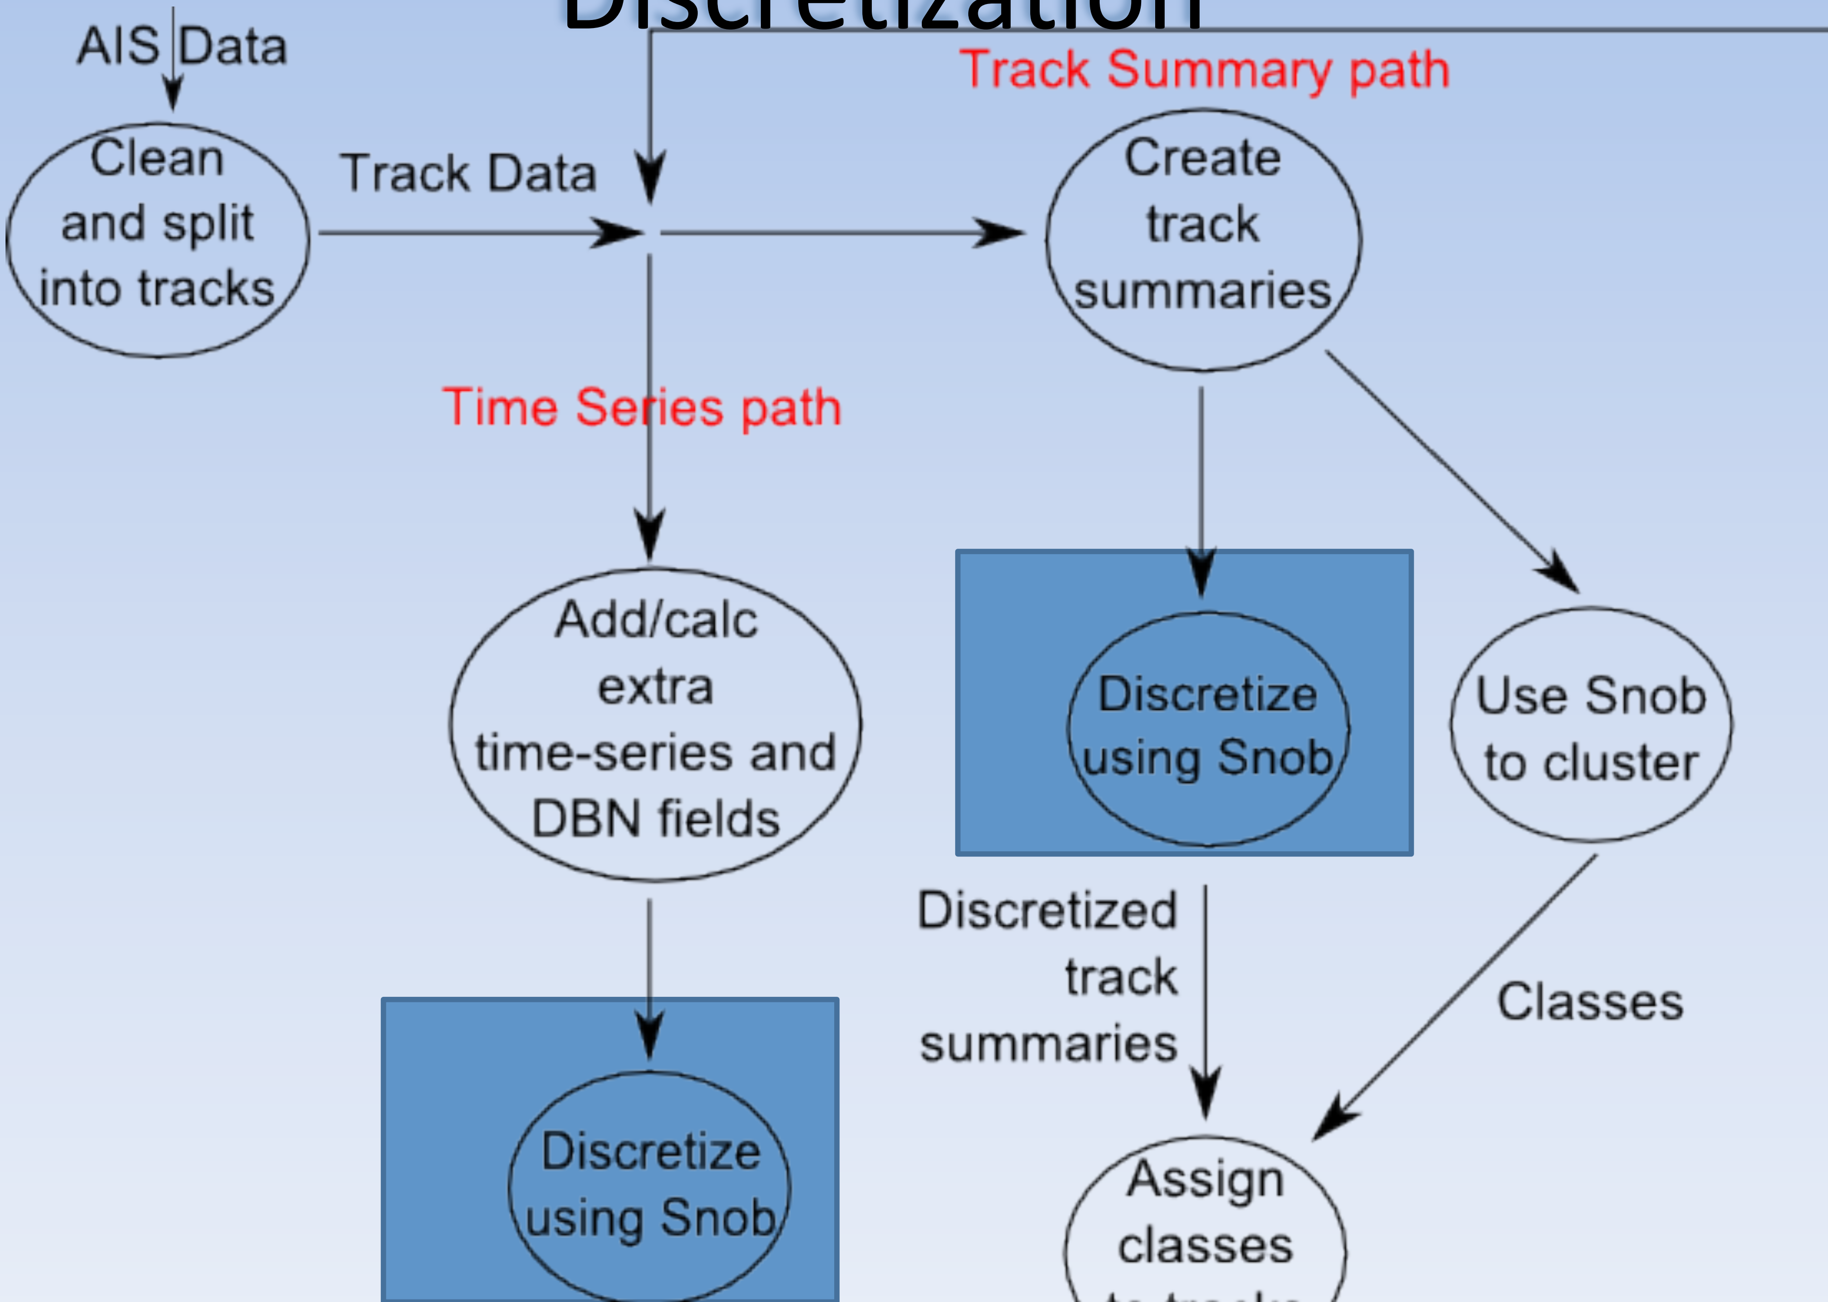

# Classification

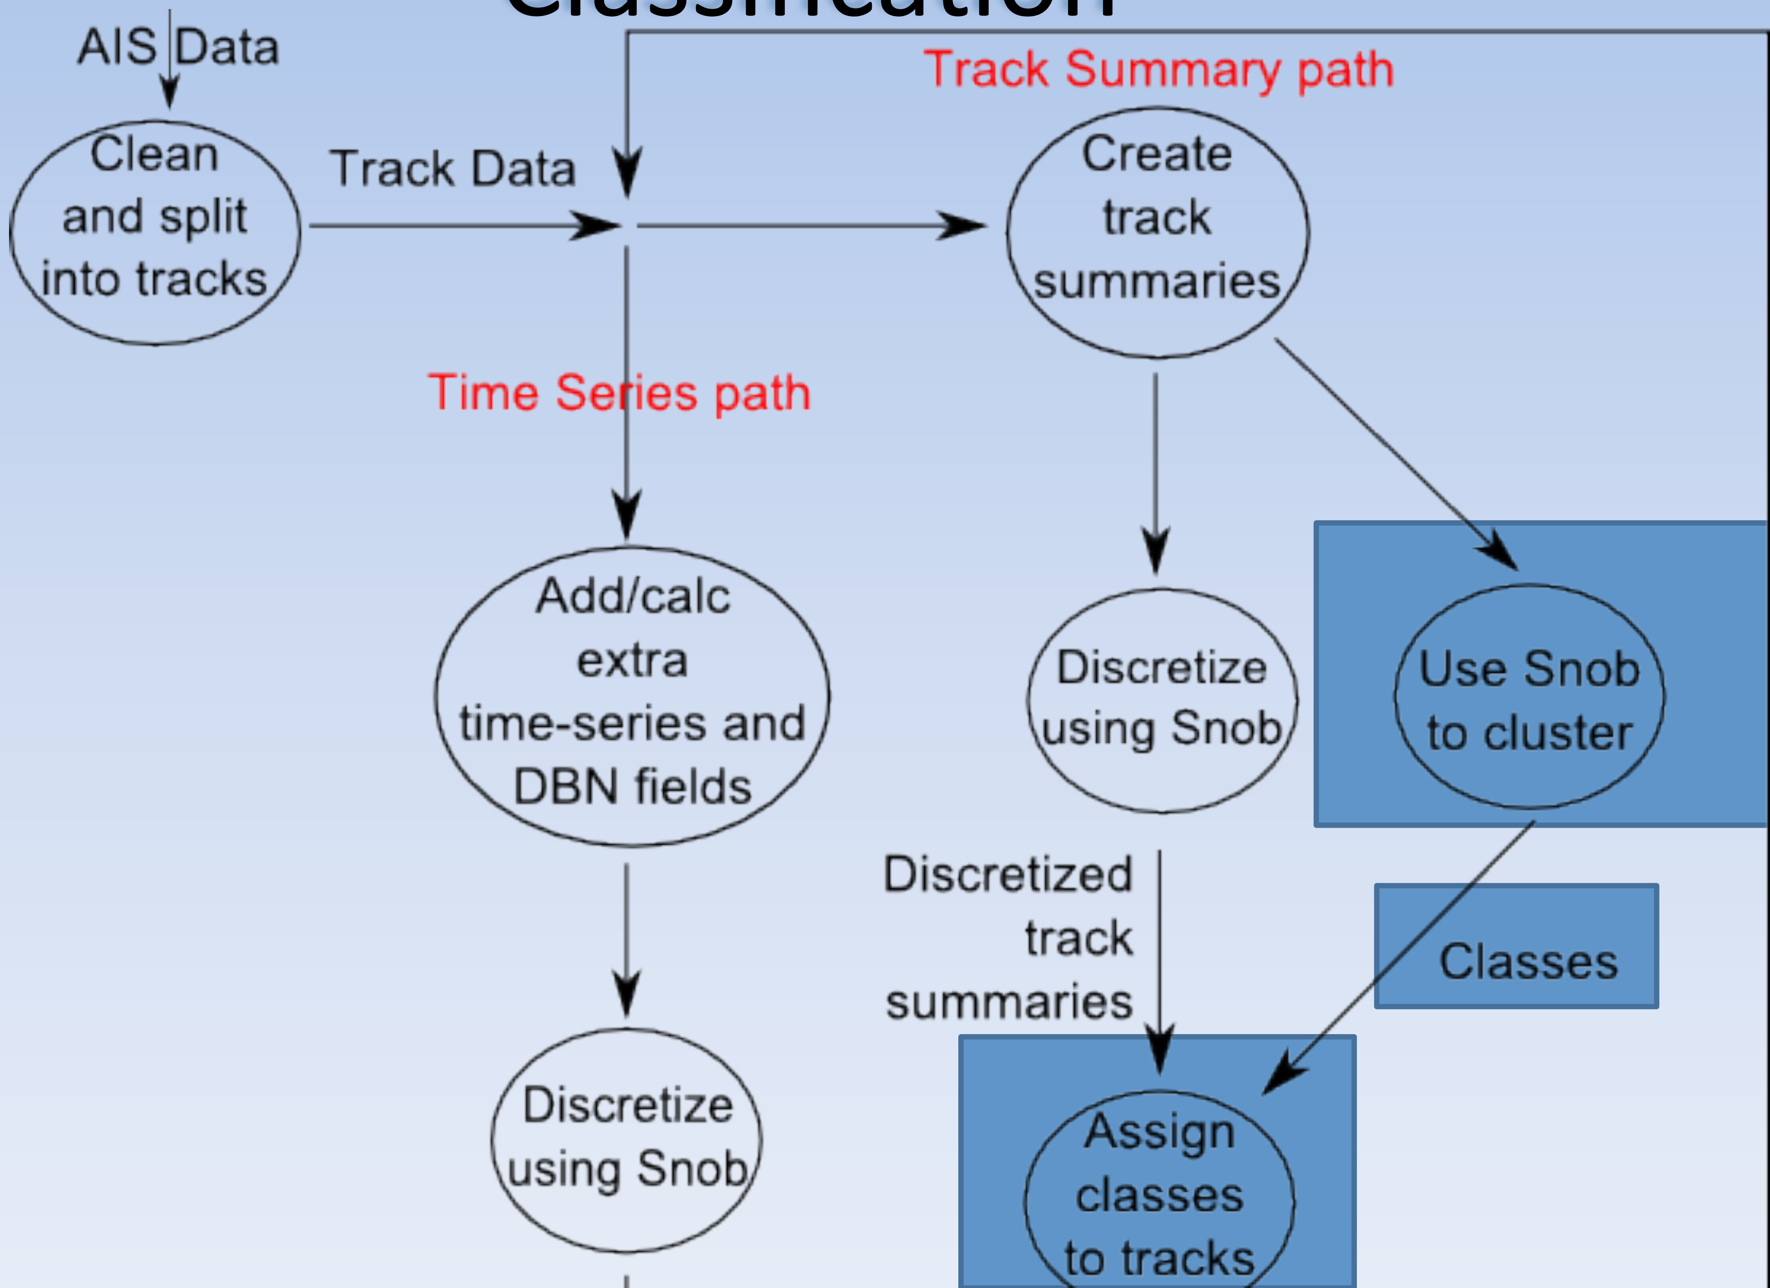

# BN Learning

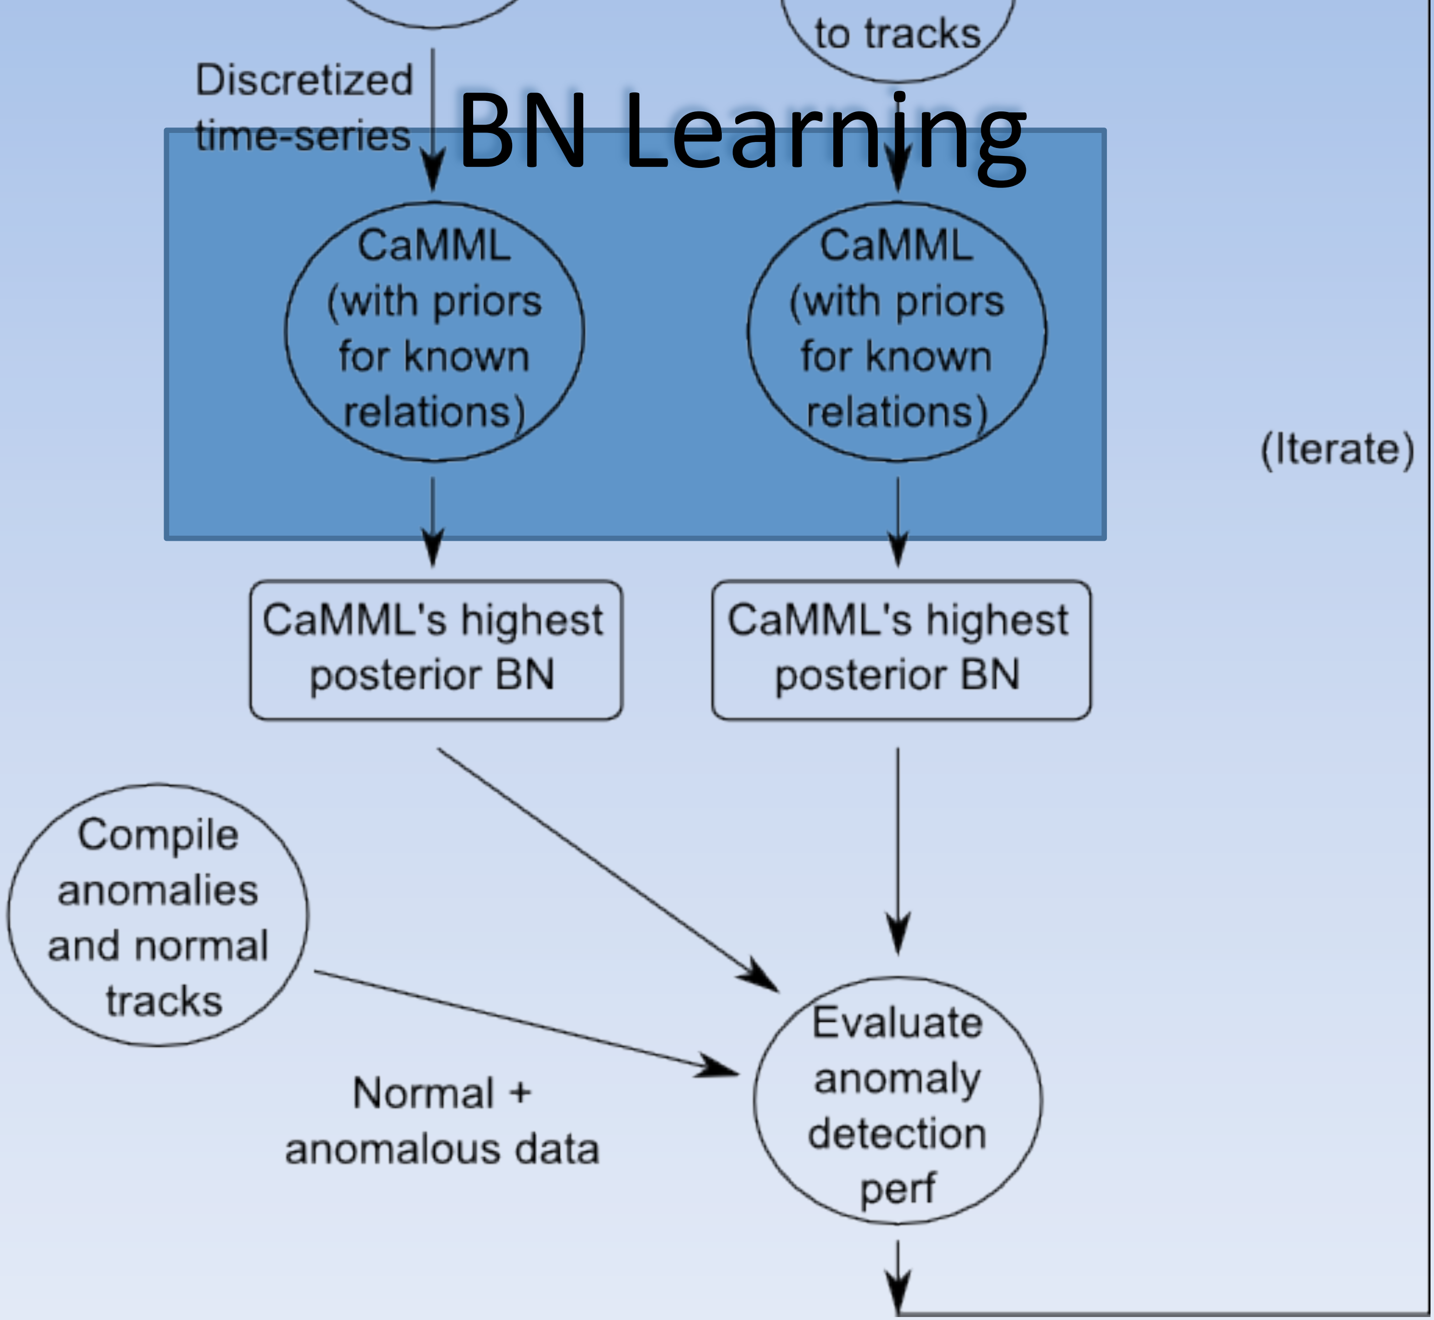
to tracks
Discretized time-series
CaMML (with priors for known relations)
CaMML (with priors for known relations)
(Iterate)
CaMML's highest posterior BN
CaMML's highest posterior BN
Compile anomalies and normal tracks
Evaluate anomaly detection perf
Normal + anomalous data

# BN Learner

- CaMML (Causal discovery via MML)
- Uses stochastic search (MCMC) and score approach (using minimum-message length)
- Parameterized model with standard counting-based procedure
- Allows user to specify a wide variety of expert priors on structure

# Structural priors

- Hard priors to enforce DBN relationships
- Temporal tiers

| | |
|---|---|
| **1st Tier** | ShipType, ShipSize, Rainfall, MaxTemp, EstWindSpeed, EstOktas |
| **2nd Tier** | Lat, Lon, Speed, Course, Heading, Acceleration, DayOfWeek, HourOfDay, CourseChangeRate, HeadingChangeRate, NumCloseInteractions, NumLocalInteractions, ClosestType, ClosestSpeed, ClosestCourse, ClosestDistance, SinceDawn, SinceDusk |
| **3rd Tier** | Lat-t2, Lon-t2, Course-t2, Heading-t2, Speed-t2, Acceleration-t2 |

# Experimental methodology

- Divided data 80% training, 20% testing
- Sets of 10 runs of CaMML , taking CaMML's reported "best" (highest posterior) BN each time

# Resultant BN: Time Series DBN

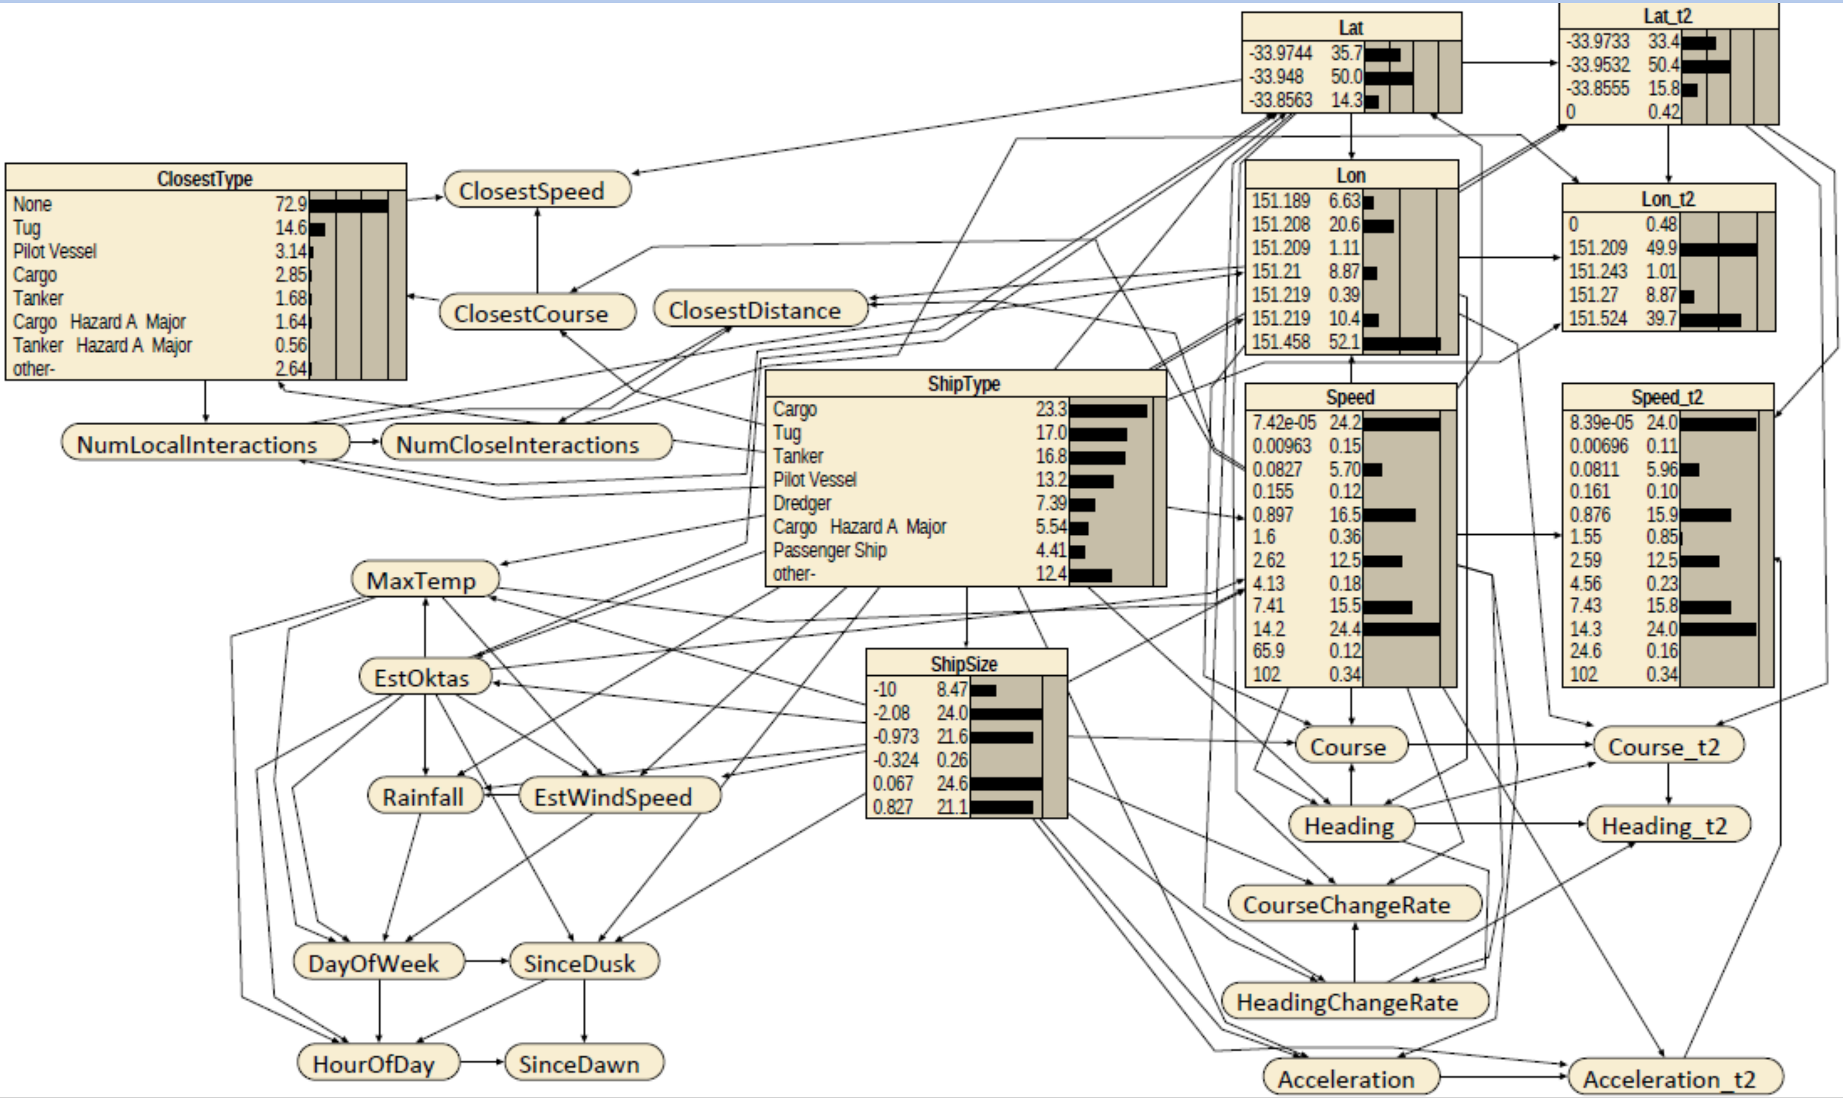

# Resultant BN: static track summary BN

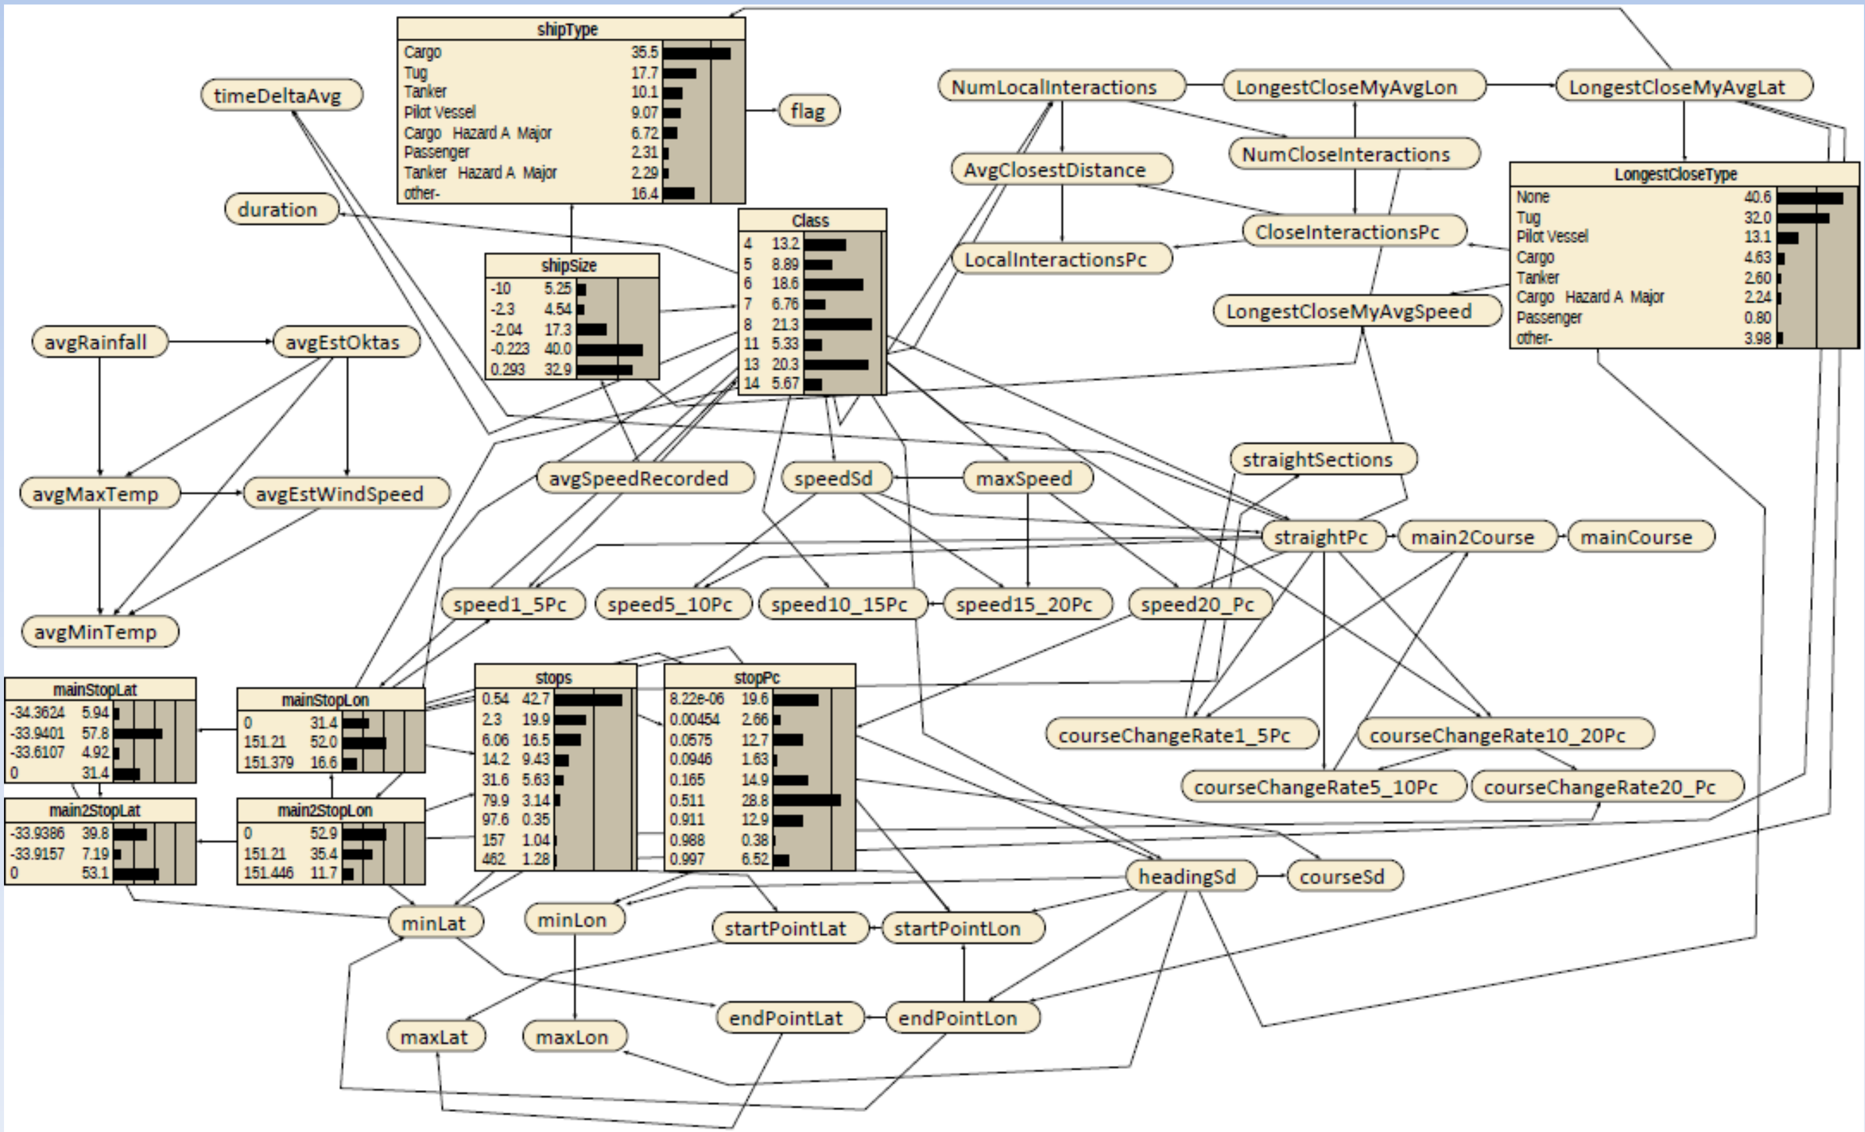

# Notes on models

- An isolated subnetwork on the left in the track summary model, which indicates that information about the weather (cloud cover, temperature and so forth) has no effect on the remainder of the network
- Learned time series models *included* weather variables, although their influence on kinematics variables was relatively weak

# Notes on models

- Few arcs in the learned networks represent intuitive direct causal relations, other than the DBN arcs (given as hard priors) and the weather variables.
- Many of the other variables are simultaneous properties of the vessel, which will be correlated by hidden common ancestors.
  - E.g. ship's peed, size and course related by? Vessel owner, purpose of trip, nature of crew & contents

# Notes on Models

- Hidden causes partly captured by the ShipType (see time series model)
  - e.g., the purpose of a trip employing a cargo ship is almost always transport.
- In the track summary network this common cause role is assumed by the ‘Class' variable instead

# Notes on Models

- Entering`Tug' or `Pilot Vessel' into the `ShipType' variable significantly increases the chance of another vessel being nearby.
- Cargo ships, on the other hand, travel mostly solo and tankers almost exclusively so.
- Ship sizes are also highly correlated with position , with larger vessels tending to appear in a restricted set of locations.

# Notes on Models

- The track summary model shows that cargo ships and tankers spend most of their time travelling straight, while tug directions are much more variable.
- Tugs also tend to stop in different locations from cargo ships, and they tend to be stopped for longer periods than cargo ships.

# Detecting anomalies using BNs

- Jensen and Nielsen (2007) "conflict measure"
- Loy et a. (2010) use learned DBNs to calculate log-likelihoods and compare against thresholds to maximize accuracy
- Cansado & Sato (2008): cases with low probability given the learned BN are considered anomalies

# Our anomaly score

- For track summary data:
- Compute each track's prior prob given normality model
- Information theoretic approach: take negative log to produce an "anomaly score"
  - The number of bits required to describe the data, given the model
  - The higher the anomaly score, the less probable the track
- For time series data – similar (average prob over all time steps)

# Anomaly scores for all data

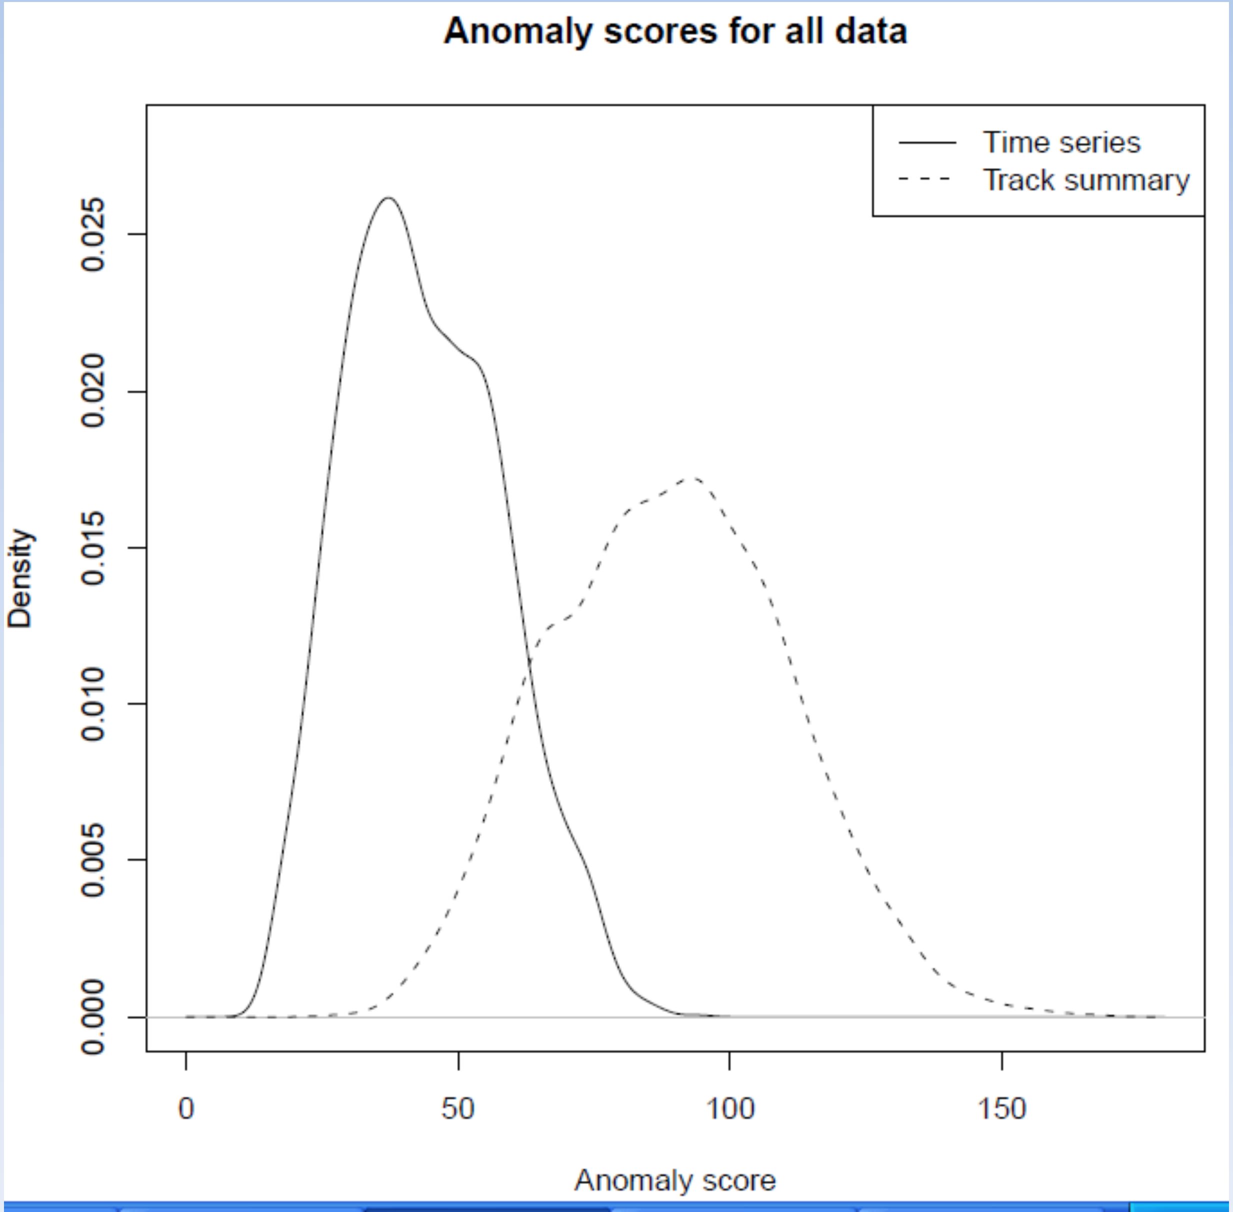

# Notes on anomaly scores

- These show a fair amount of diversity among anomaly scores, i.e. they do not simply clump around the lowest possible score.
- The scores produced by the two models are quite distinct
  - One likely reason is that the track summary scores are simply based on more variables, making each instance more specific and less probable.

# Notes on anomaly scores

- Surprisingly small correlation between the two sets of scores (*r*=0.159; $p<0.001$)
- The two models look at different aspects of each track, and, as we see below, reinforce each other when performing anomaly detection.

# Anomalous data

- Data did not include any known anomalous tracks.

- So we created our own by:

1. modifying instances by swapping incorrect ship type information (the false ship effect)
2. splicing tracks together
3. Drawing anomalous tracks

# Results - The False Ship Effect

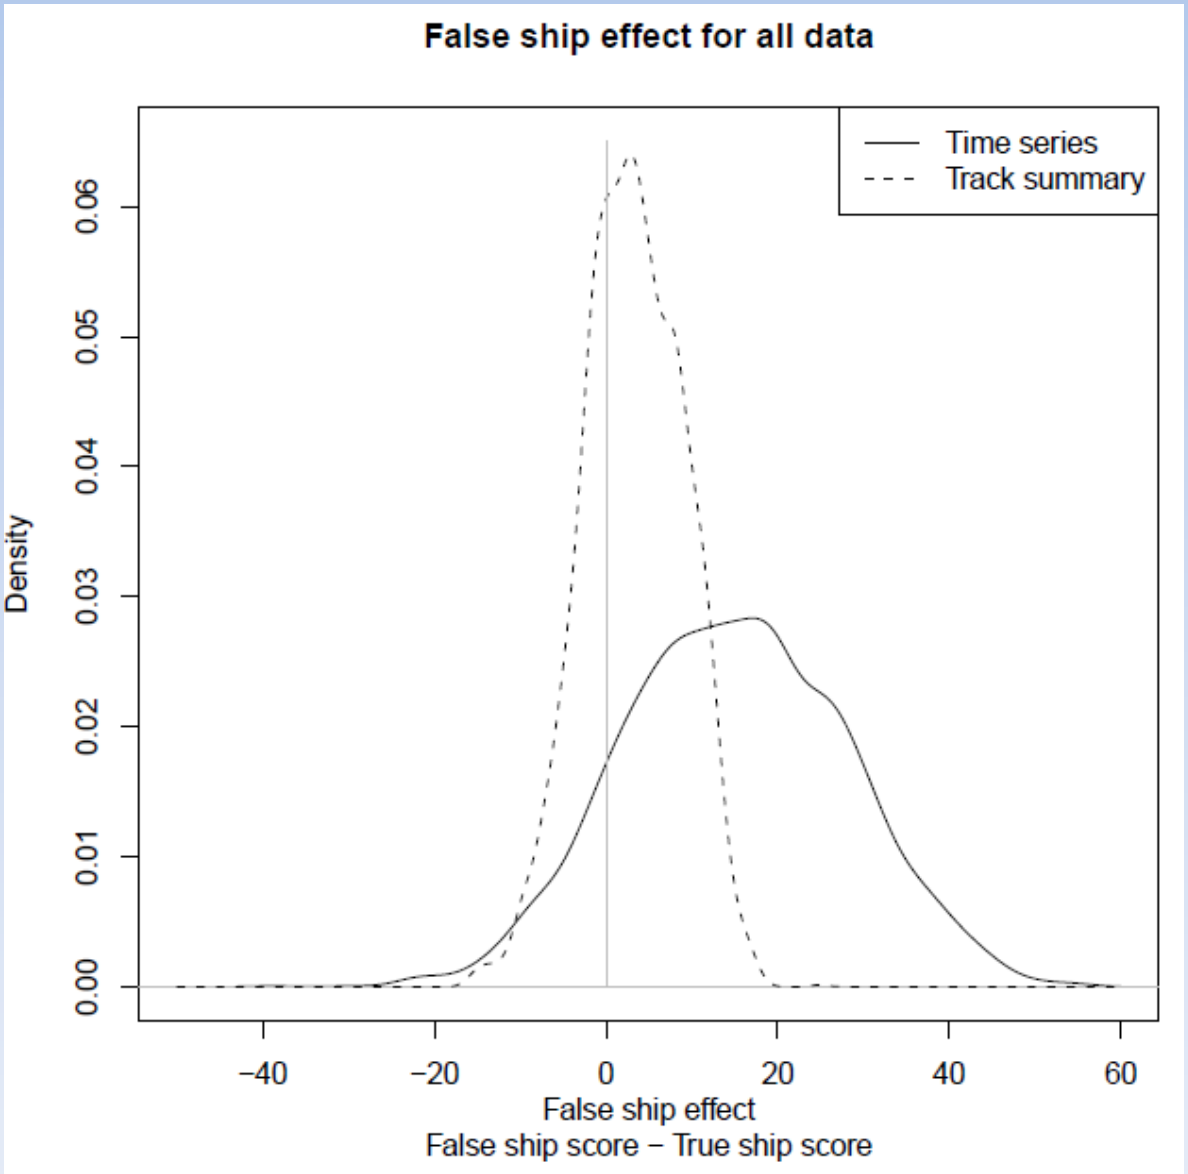

# Notes on false ship effect results

- In most cases this false ship effect is positive, increasing the anomaly score
- The false ship effect for the time series model is positive in around 87.2% of the cases as opposed to 69.4 of cases for the track summary model
- Sometimes tracks became more probable!
  - Some ship types are similar (e.g. sub-categories of cargo and tanker ships)
  - Other cases: the original track was anomalous (mis-labelled? Anomalous?)

# Results – The False Ship Effect

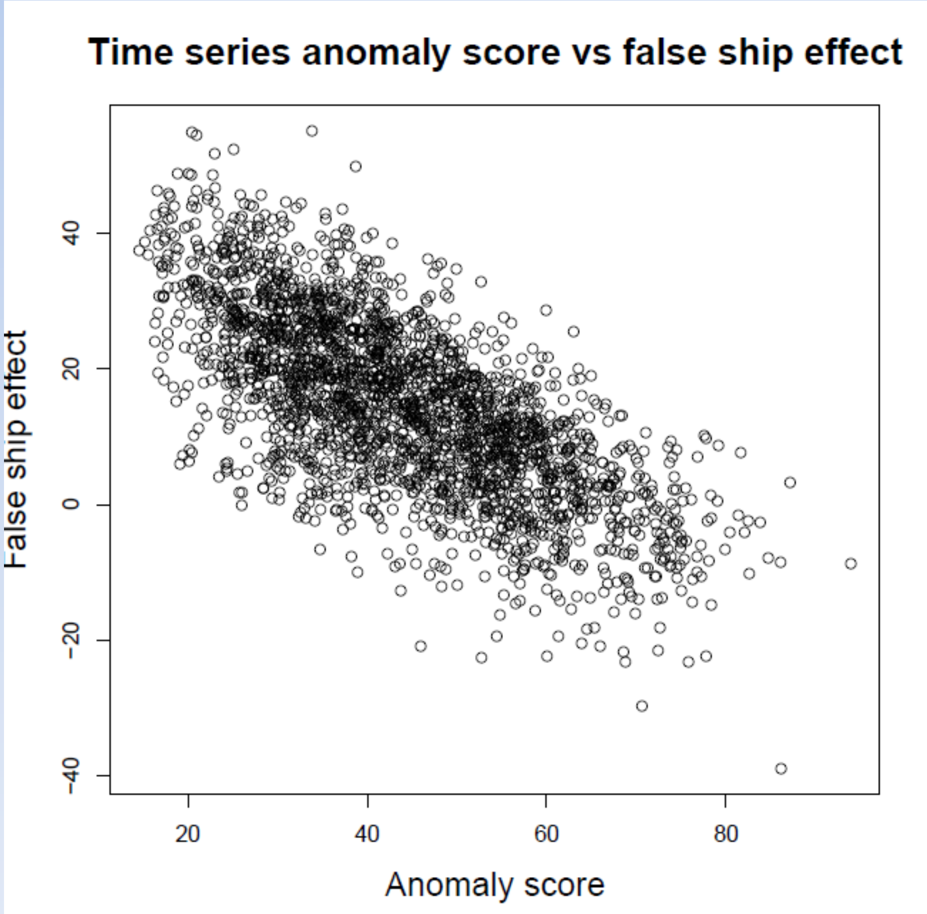


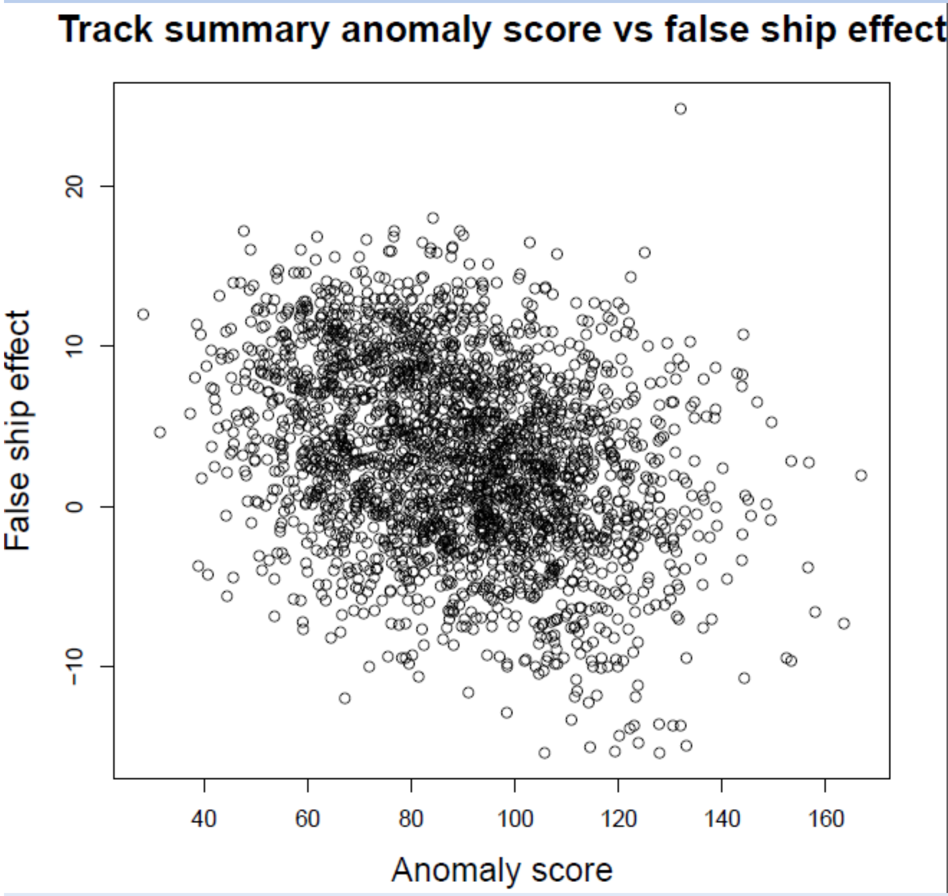

# Track Splices

| Tracks | Track Summary | Timeseries |
|---|---|---|
| **Same type** | 115.4 | 45.6 |
| **Different types** | 121.3 | 48.9 |
| **All data** | 89.0 | 43.8 |

- Shows advantage of higher level view of the track summaries

# Results – Manually drawn anomalies

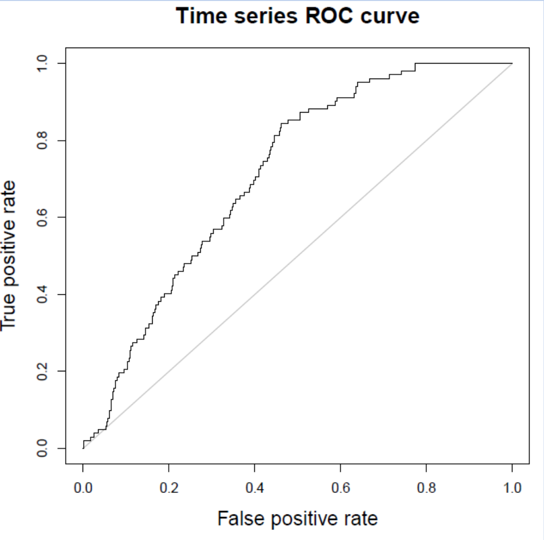


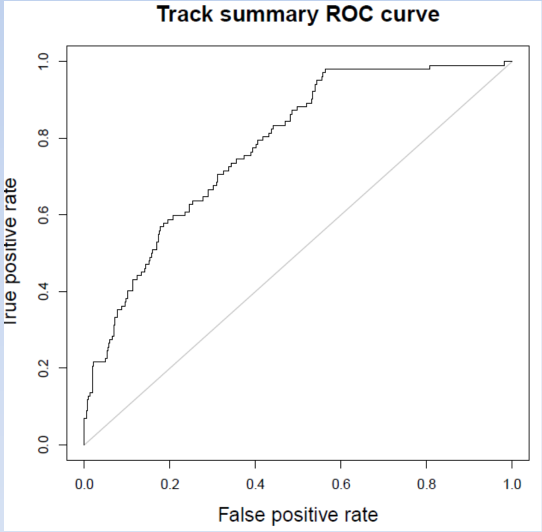


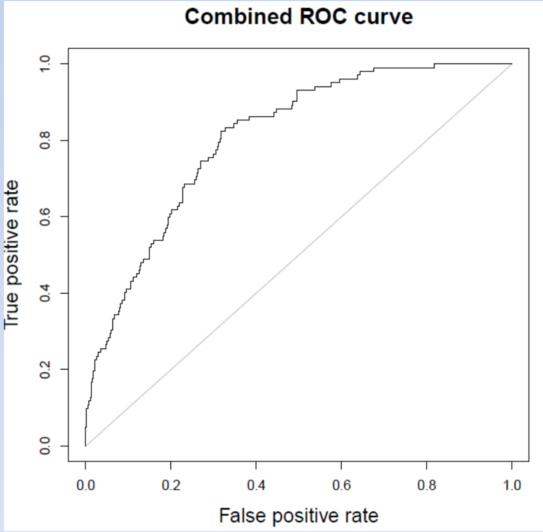


AUC= 0.712

AUC= 0.780

AUC= 0.809

# Results – Manually drawn anomalies

| Type | Track Summary Score | Delta | Time Series Score | Delta |
|---|---|---|---|---|
| Normal test tracks | 90.8 | (0) | 45.7 | (0) |
| Random movement in the middle of water | 102.4 | +11.7 | 50.8 | +5.1 |
| Closed tracks in the middle of water | 101.7 | +10.9 | 53.7 | +8.0 |
| Very short tracks | 95.5 | +4.7 | 62.7 | +17.0 |
| Unusual stops | 119.1 | +28.3 | 48.6 | +2.9 |
| Tracks with many interactions | 139.9 | +49.1 | 75.8 | +30.1 |
| Tracks with many loops | 126.2 | +35.4 | 52.7 | +7.0 |
| Travel over land | 122.2 | +31.4 | 60.2 | +14.5 |
| Appearing at edges of observable area only | 103.5 | +12.7 | 54.2 | +8.6 |
| Very noisy observations | 135.2 | +44.4 | 54.6 | +8.9 |
| Tracks behaving against type | 113.7 | +22.9 | 57.8 | +12.0 |
| Multiple anomalies | 126.9 | +36.1 | 53.9 | +8.2 |

# Notes on results for manually drawn anomalies

- Both models easily detect tracks containing too many close interactions
- Time series model detected overly short tracks best
- Track summary model outperformed timeseries for tracks containing unusual stops (as would be expected)
- In most cases, the track summary model outperformed the time series model

# BayesWatch: An Online Anomaly Detection System

- Again with DSTO
- Allows real-time monitoring or batch identification
- Track types: car, human, vessel, or new types
- Learning DBNs, static BNs or using uploaded BNs
- GUI interface for admin & user

FIN