Dan Gutfreund – manager, Machine Learning Technologies Group, IBM Research - Haifa
30 May 2013

Integrated Fraud Management

IBM Research

The world is our lab!

IBM Research – global labs

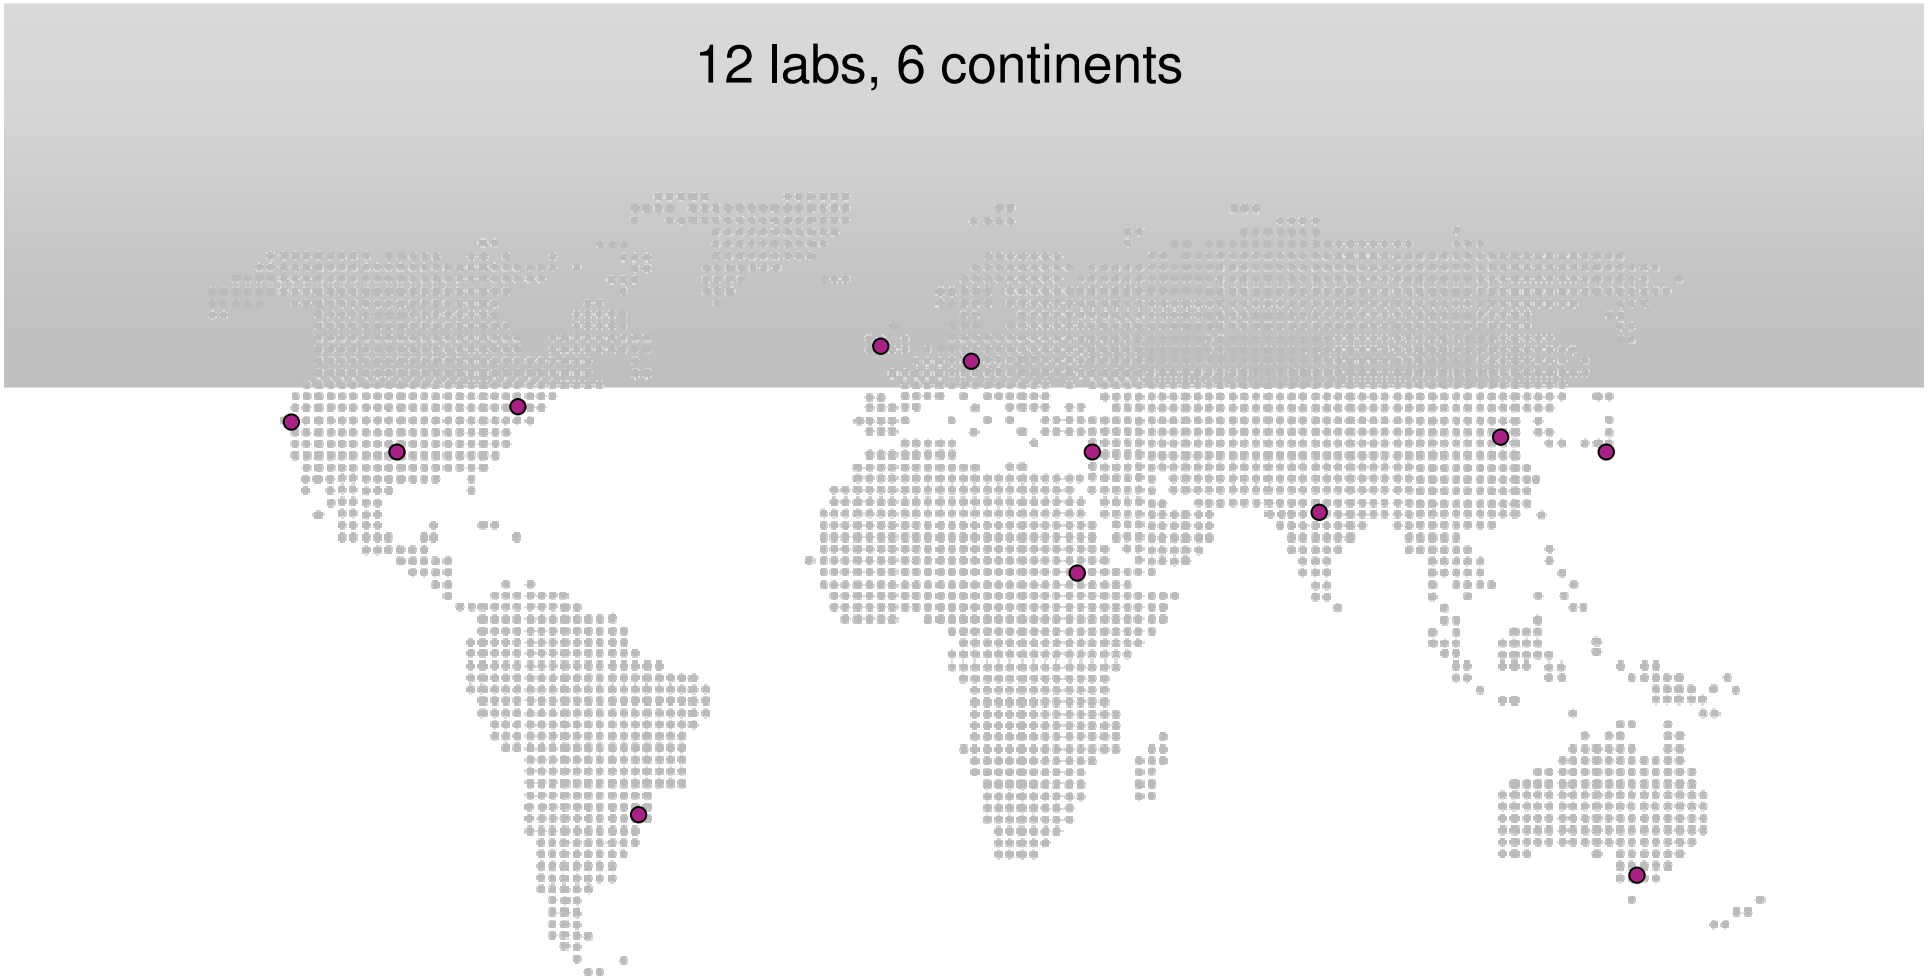

Plan for this talk

1. Banking fraud
2. DRIFT - a machine learning approach
3. Case studies and (advanced) capabilities

Different faces of fraud

Banking

- Money laundering
- Online fraud
- Transactional fraud
- Loan fraud

Insurance

- Claims fraud
- Policy fraud
- Internal/ employee fraud

Public Sector

- Tax/revenue
- Healthcare
- Food and nutrition programs
- Compliance and audit
- Social services

Health Care

- Provider/vendor fraud
- Pharmaceutical fraud
- Physician/ employee Fraud

5

Transactional fraud

Credit cards, debit cards, ATM, e-payments

Total amount of credit card fraud world-wide: $3.5B-5.5B (most in the US)

Fraud-related losses on UK-issued plastic cards in 2011: £341M [FFA UK]

Total amount of online frauds in 2011: $3.4B, 1% of revenue [JPMorgan]

Card not present becoming the most prevailing type of fraud

Active black market

DRIFT –

A machine learning approach to fraud detection

DRIFT (Detecting fRaud In Financial Transactions) - highlights

A real-time fraud detection system

Fraud detection portfolio: credit card, ATM, ACH, e-payments

Solution based on machine learning techniques

Outperforms offerings by leading vendors (proven in several PoCs)

Advanced capabilities (pre-fraud and costly fraud detection)

Machine learning background

Definition (Arthur Samuel): A field of study aiming to provide computers the ability to learn without being explicitly programmed

Teaches computers how to learn from a given data (training set) and generalize to new data (test set)

Data-driven approach (vs. human expert driven)

Plethora of applications

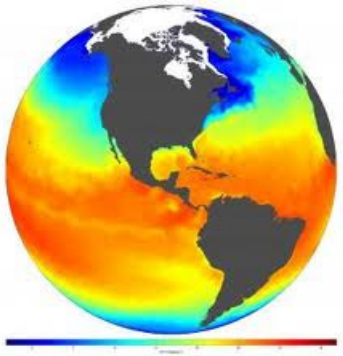

The machine learning paradigm

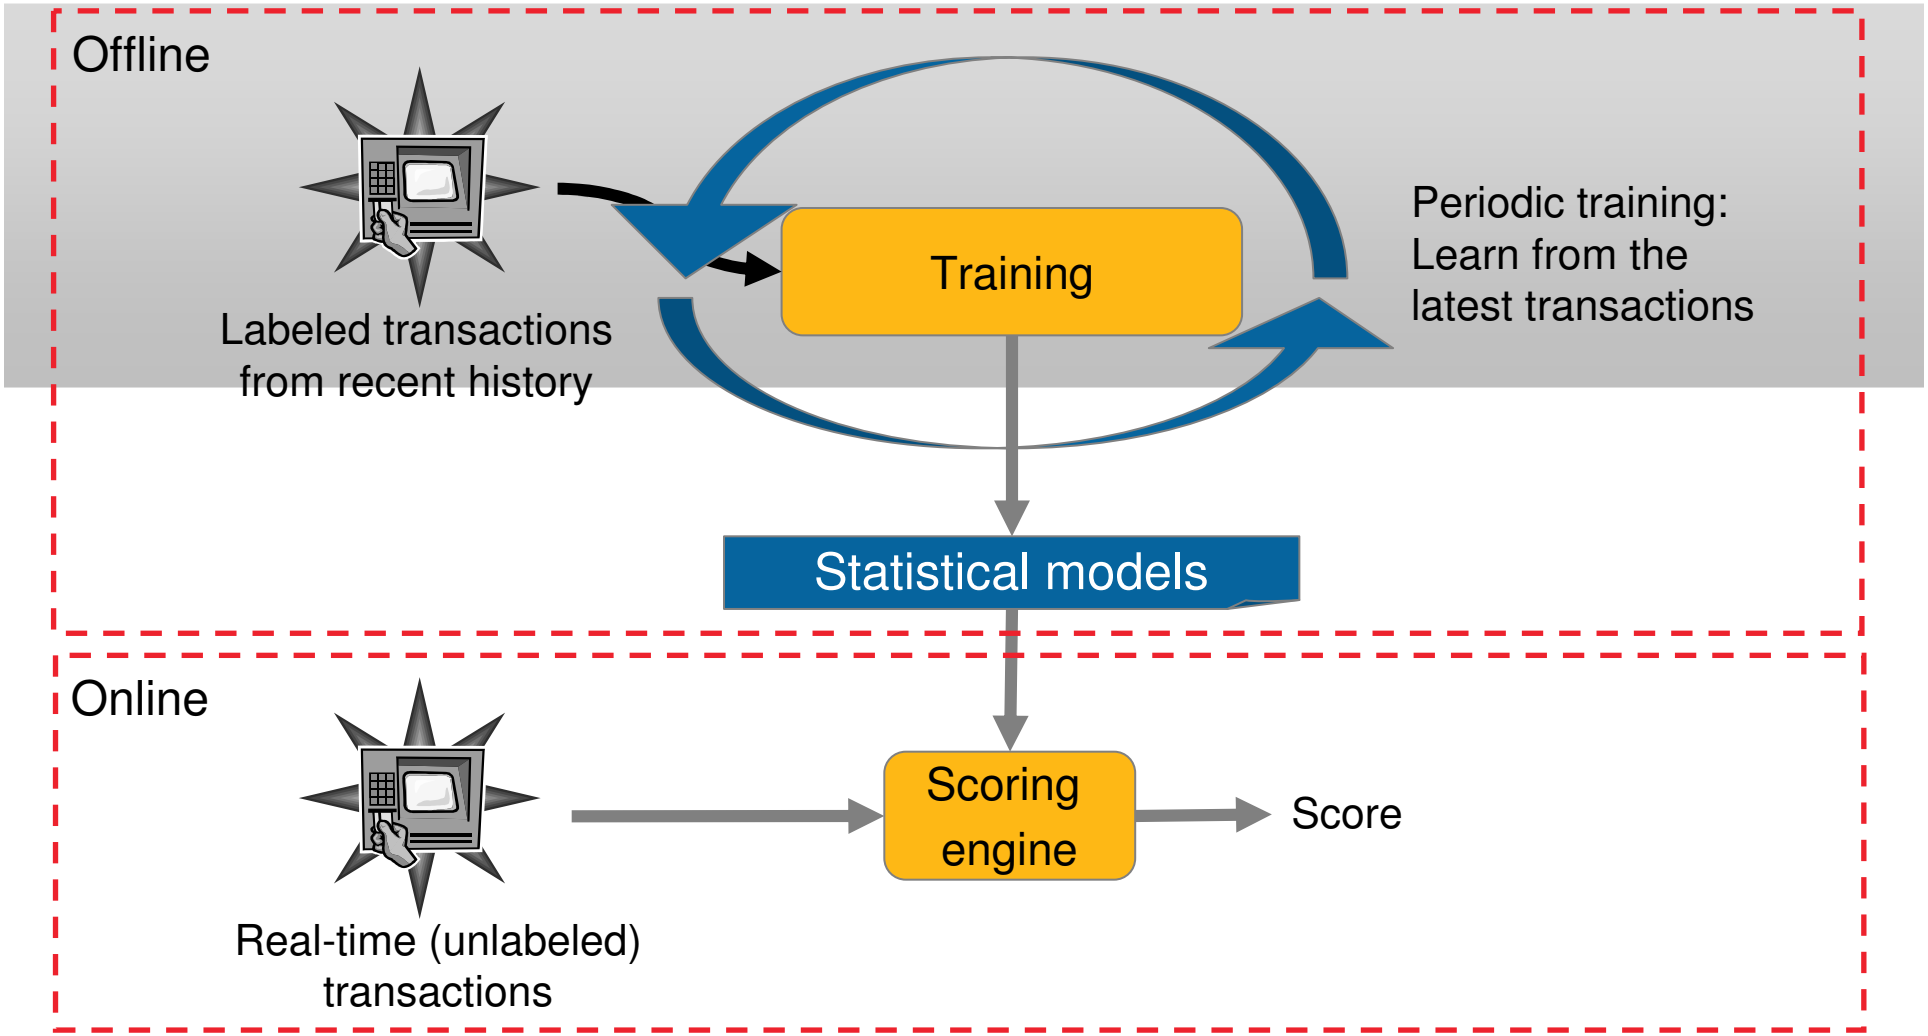

Architecture – high level

Online running: Detect fraud on streaming transactions, possible integration with existing system

Offline (re)training: Ensure solution continuously adjusts to meet new fraud patterns.

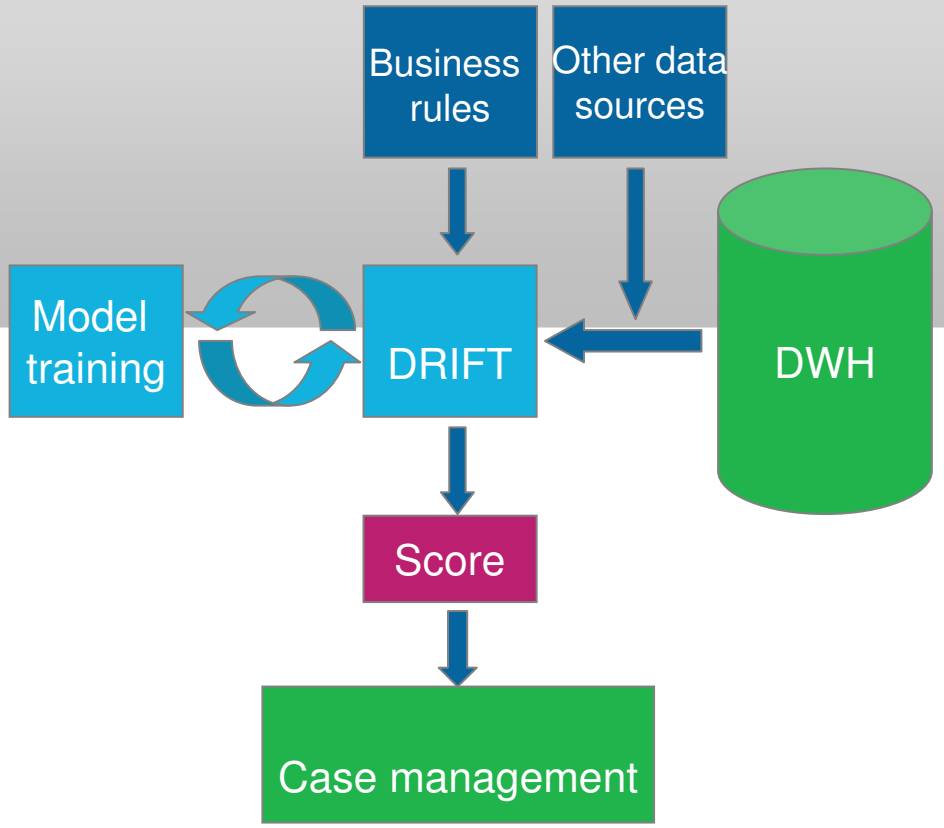

Runtime environment

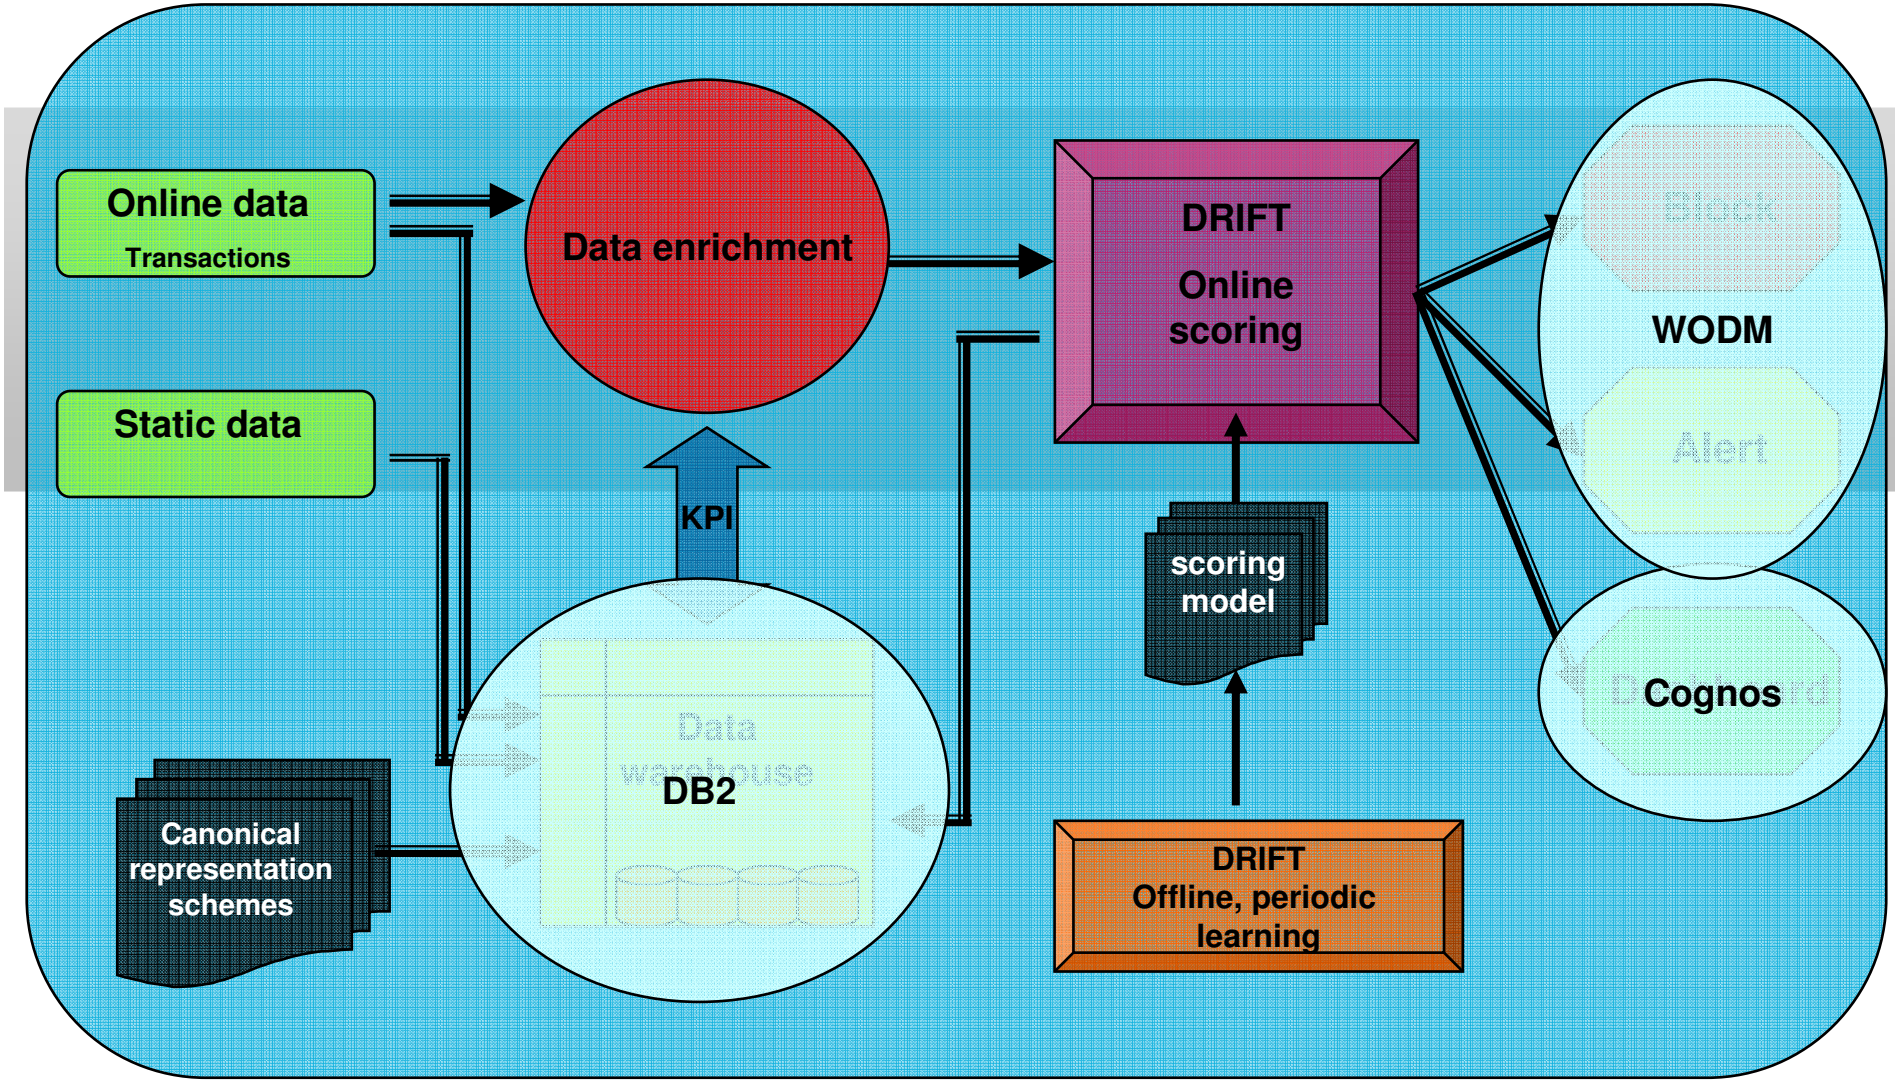

Classification problems

Input: Labeled data
Output: Classification rule

Error 1 = 20

Error 2 = 8

Error 3 = 6

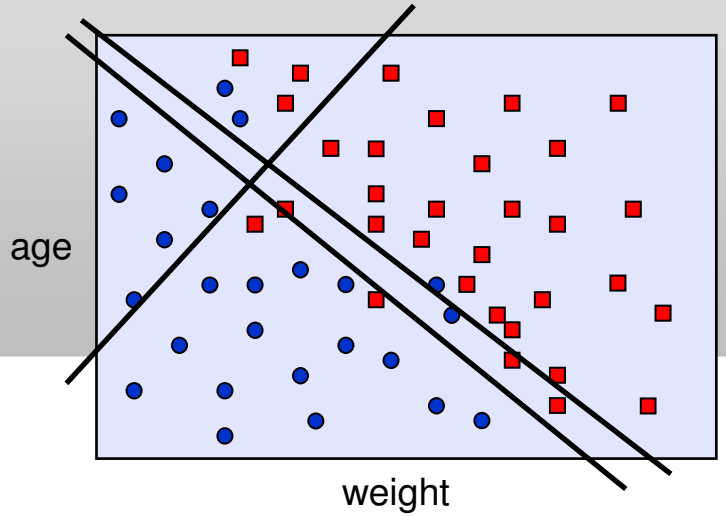

Learning algorithm: Finds a “good” classification rule from a predefined family

“Good” = Minimizes classification error on the training set

Example – Family of classifiers: Straight lines (linear separators)

Classifiers should be simple yet expressive

Decision trees

Classification error = 3

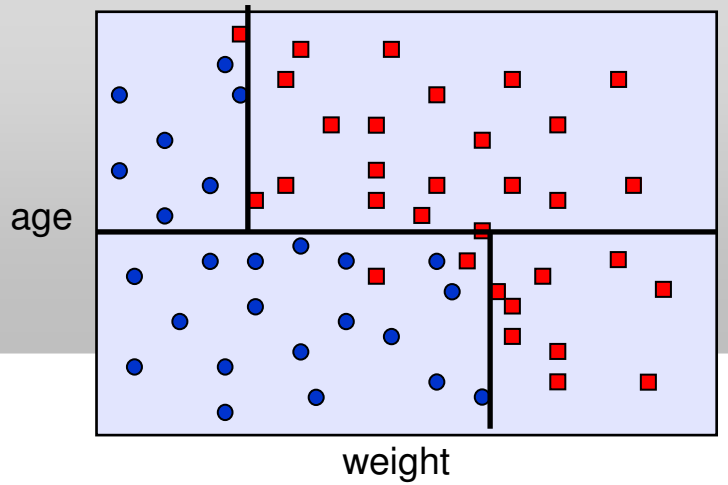

Linear separator

Classification error = 6

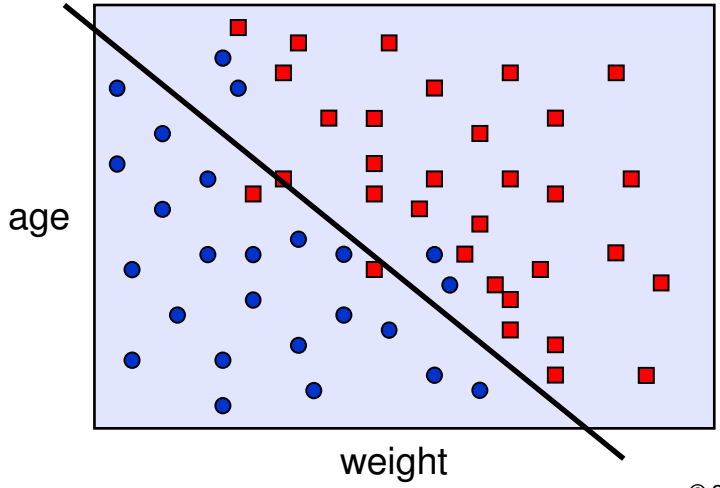

The role of data representation

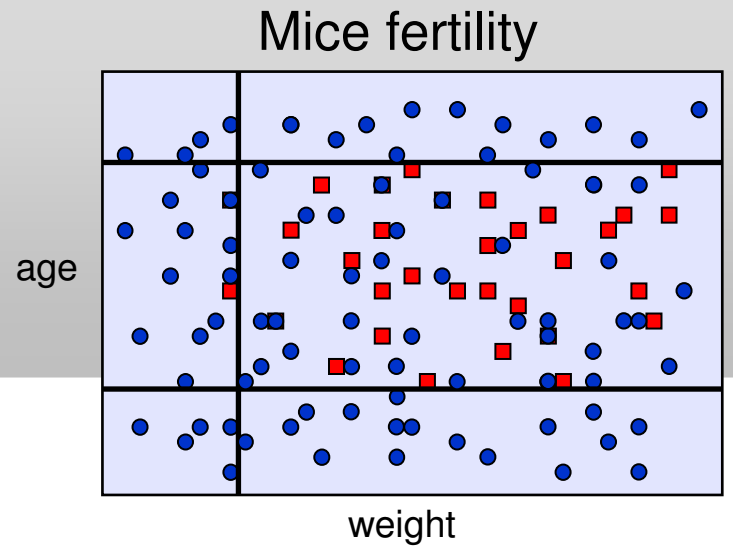

Data representation (which features to use) require domain knowledge
In DRIFT, we compute ~100 different features
Examples:

- Probability for a certain hour in the day
- Number of transactions in the last hour
- New country
- …

Simple illustration

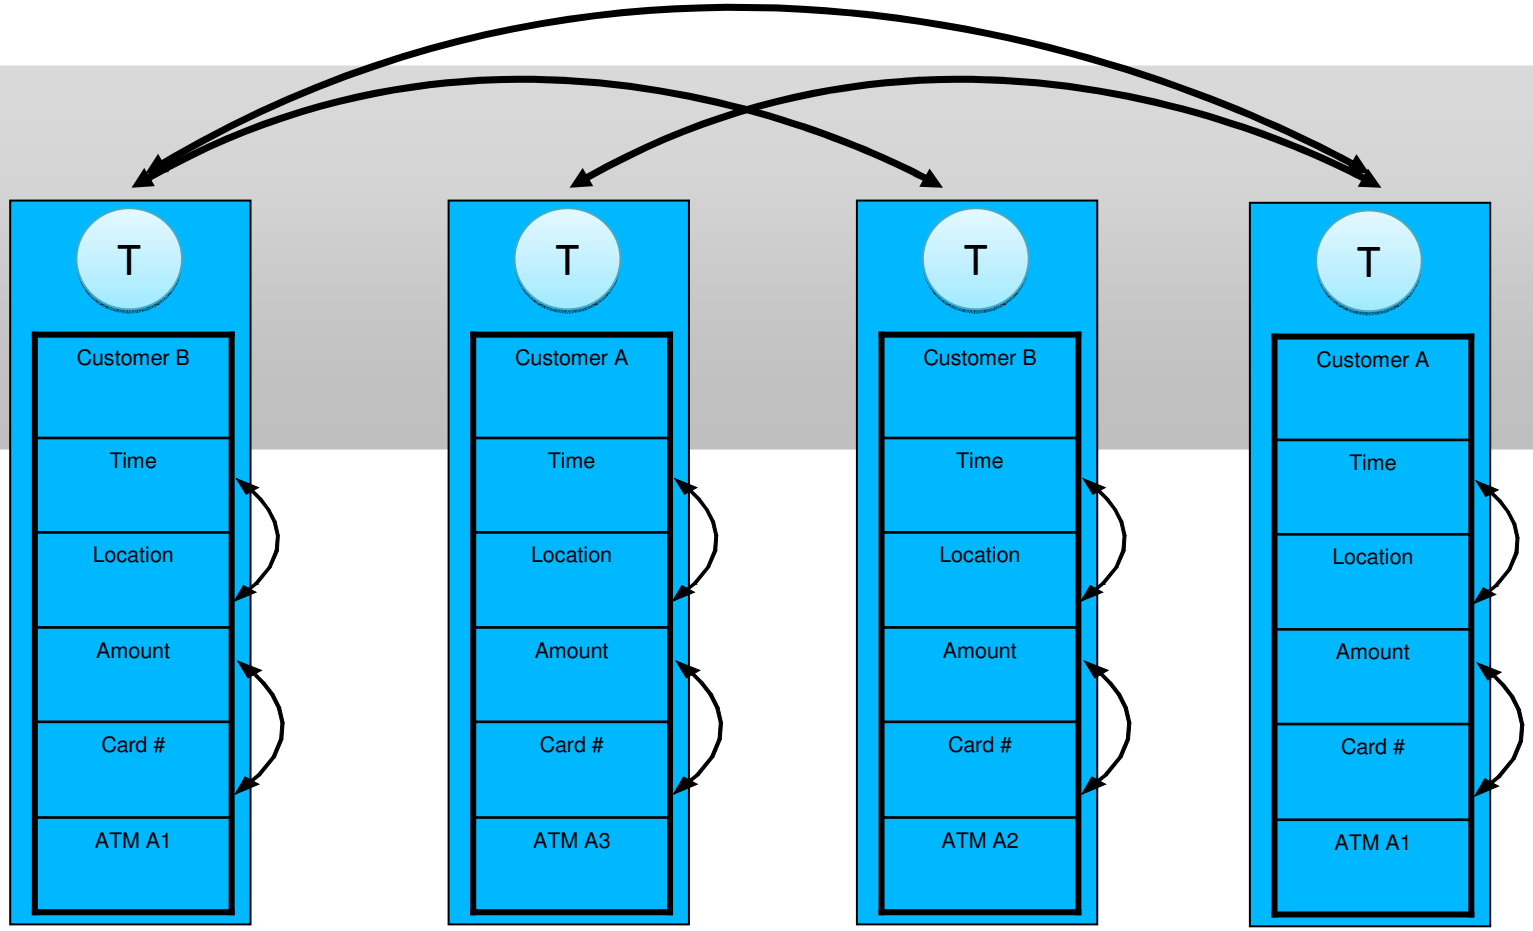

Data enrichment

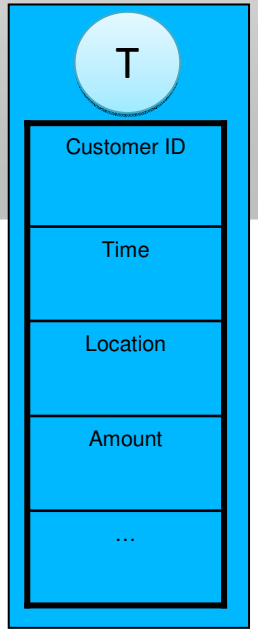

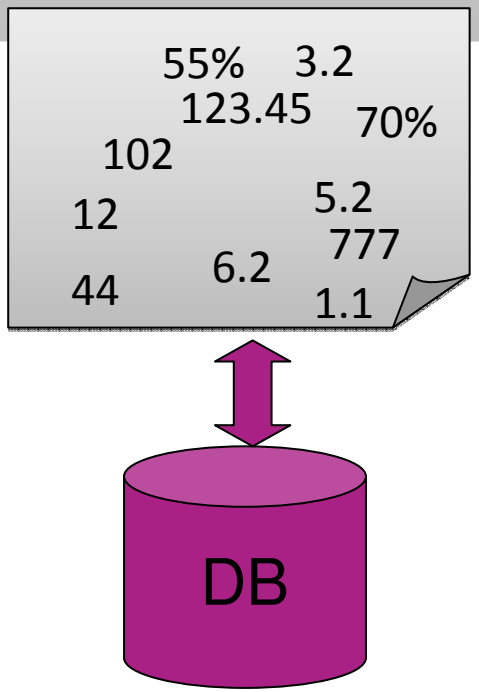

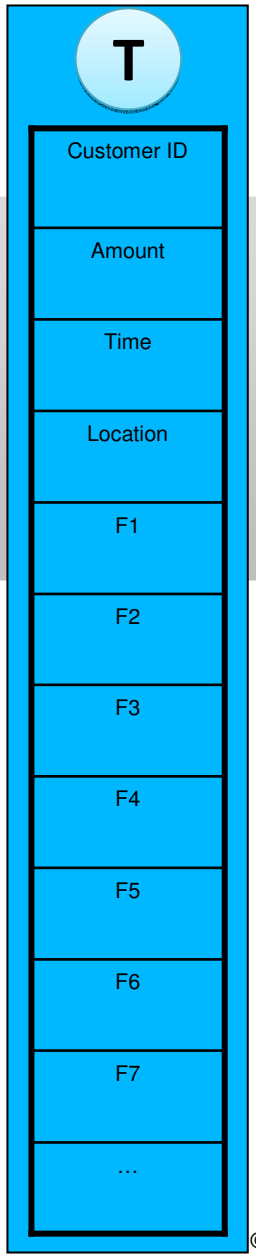

High dimensional data

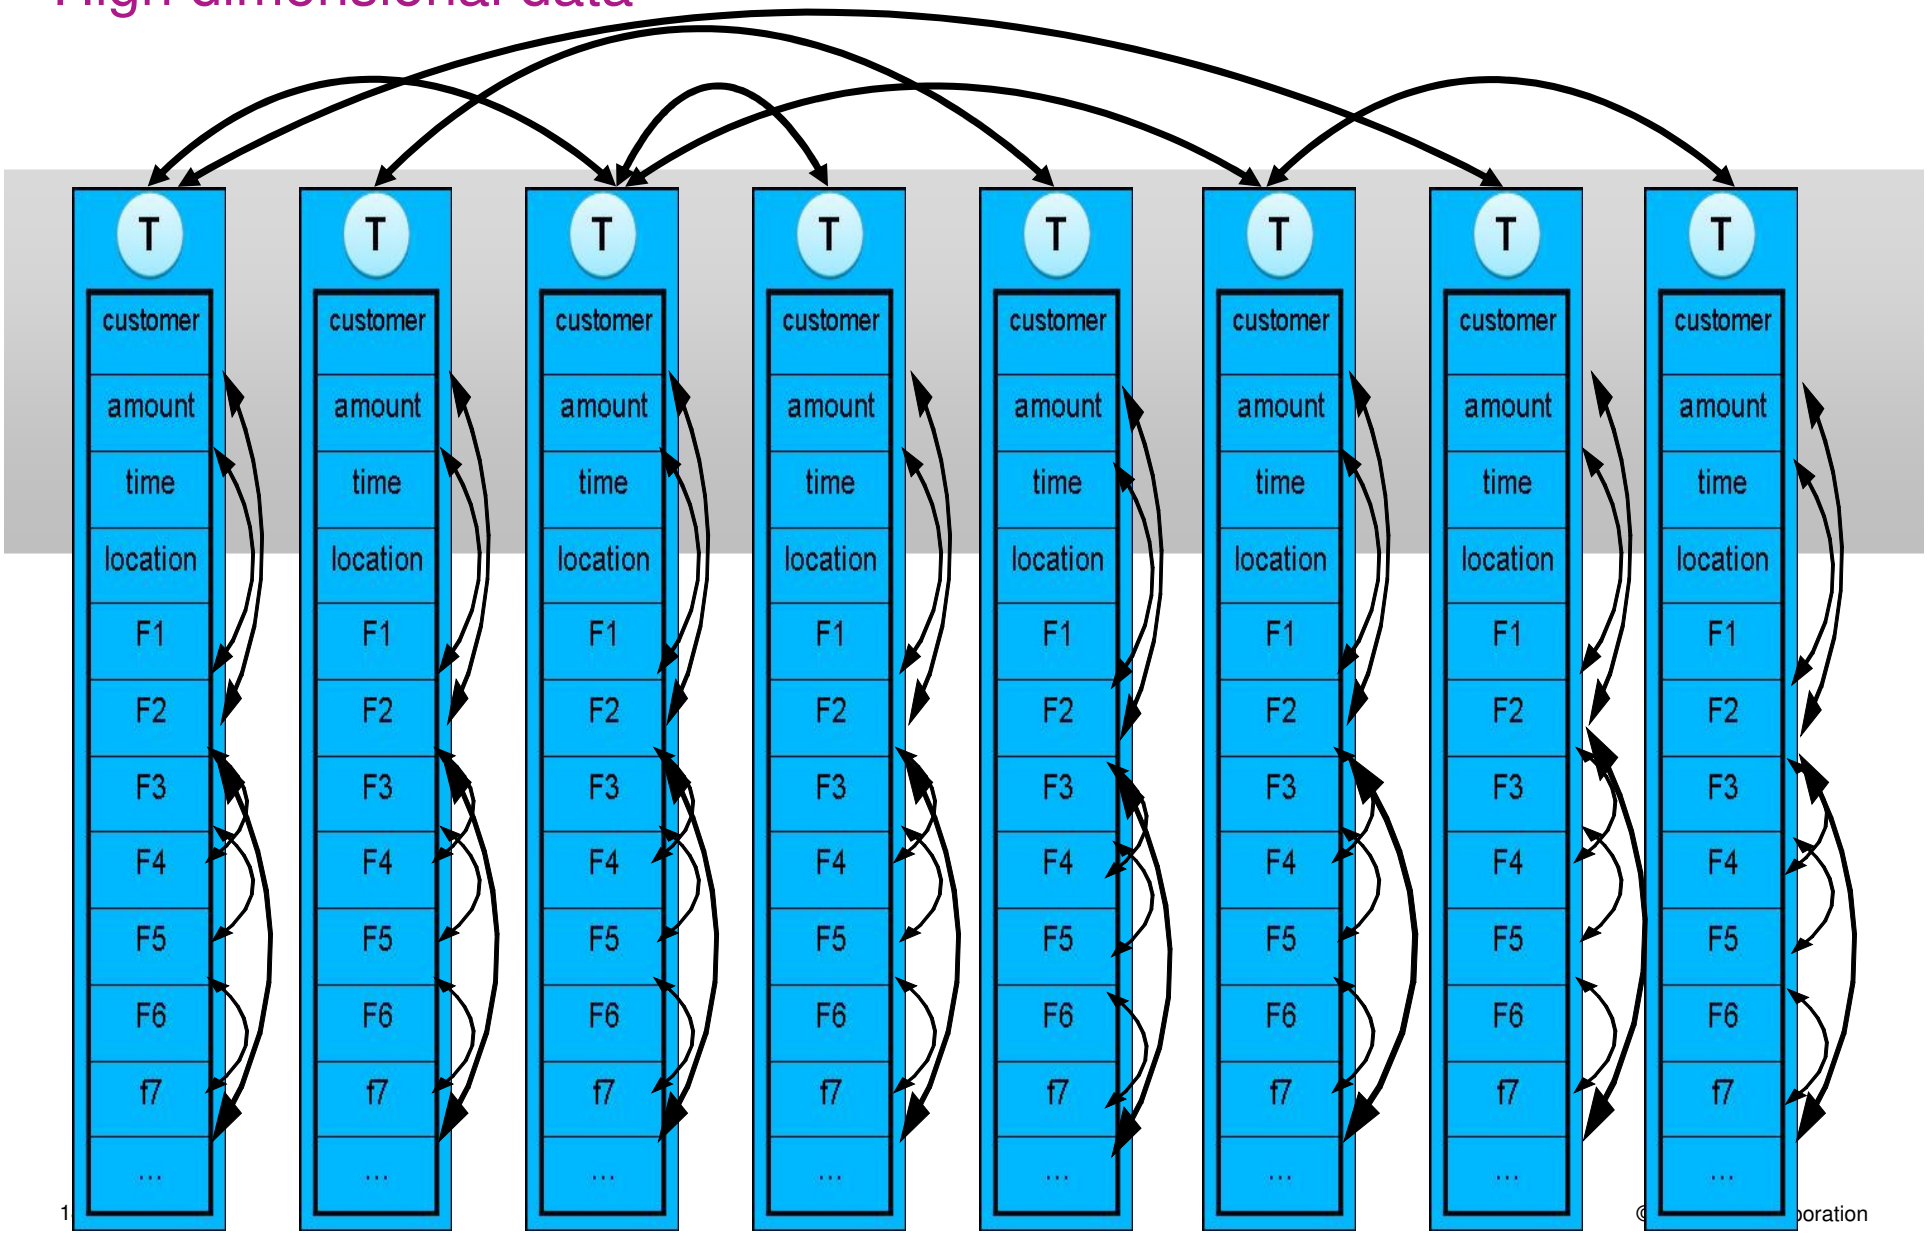

Case studies and advanced capabilities

Performance evaluation through ROC curves

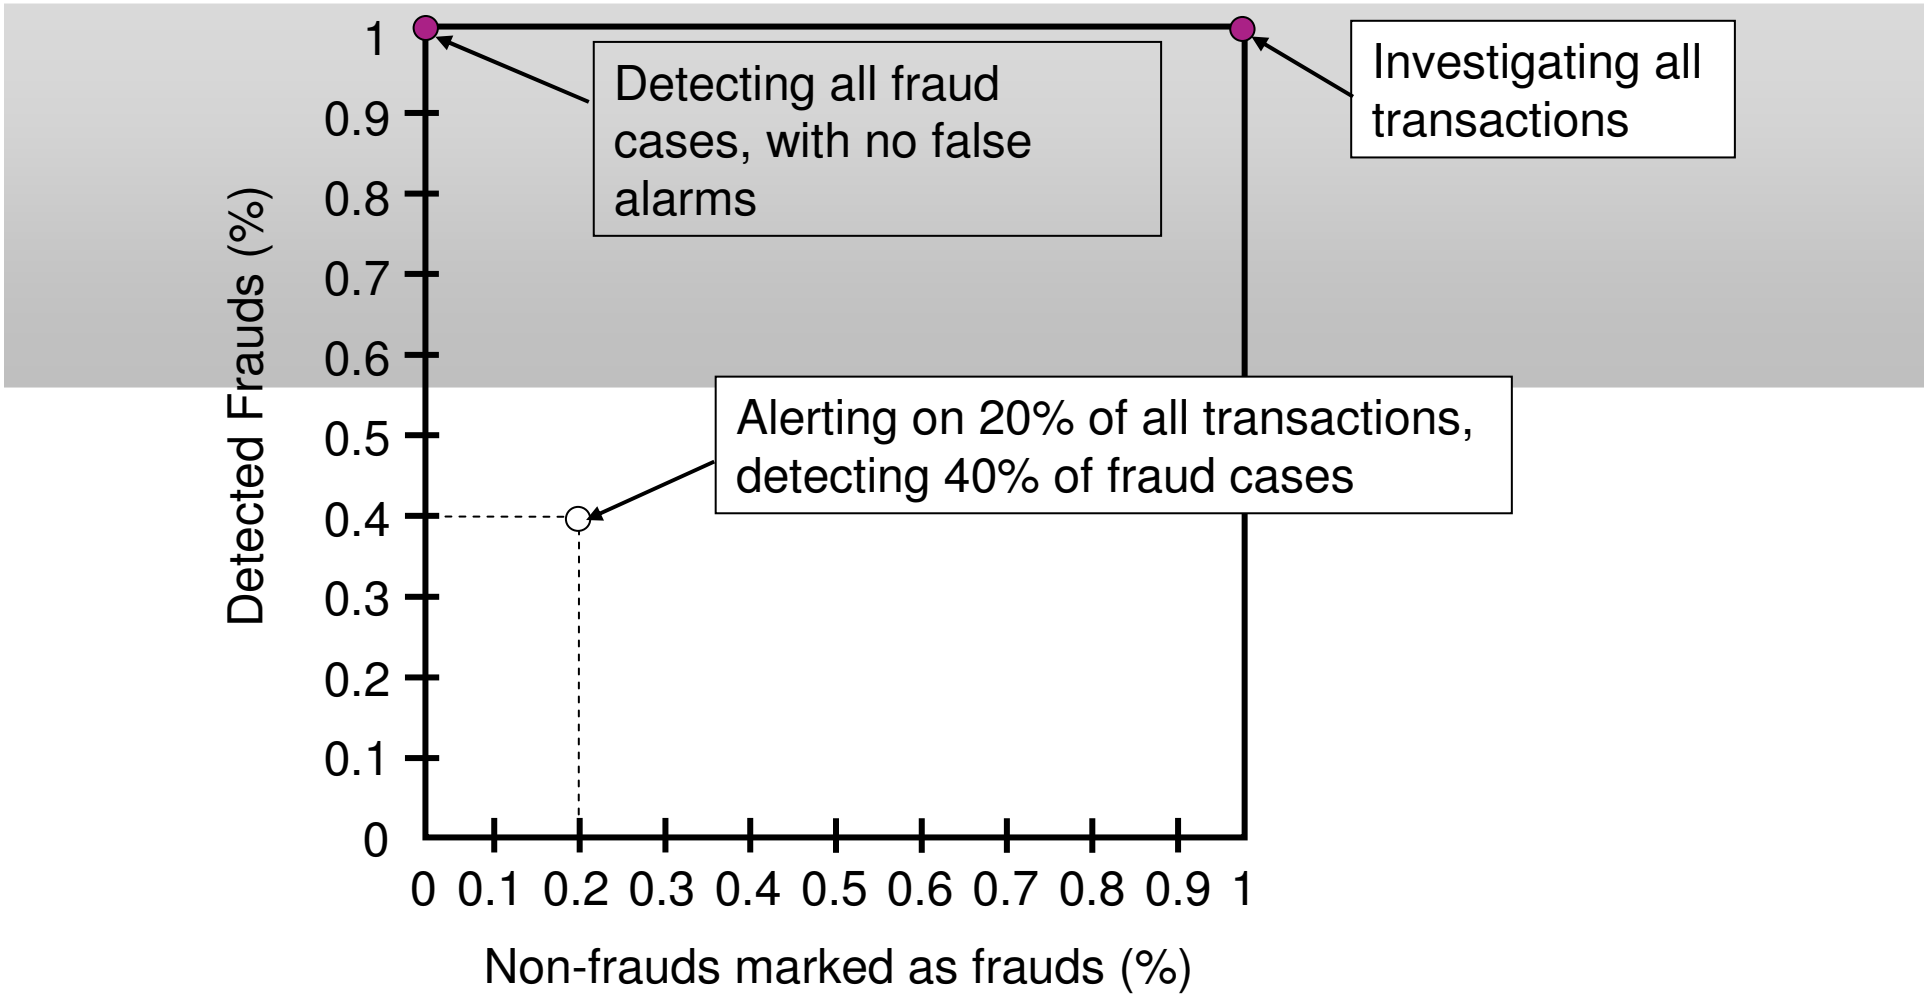

Performance evaluation through ROC curves

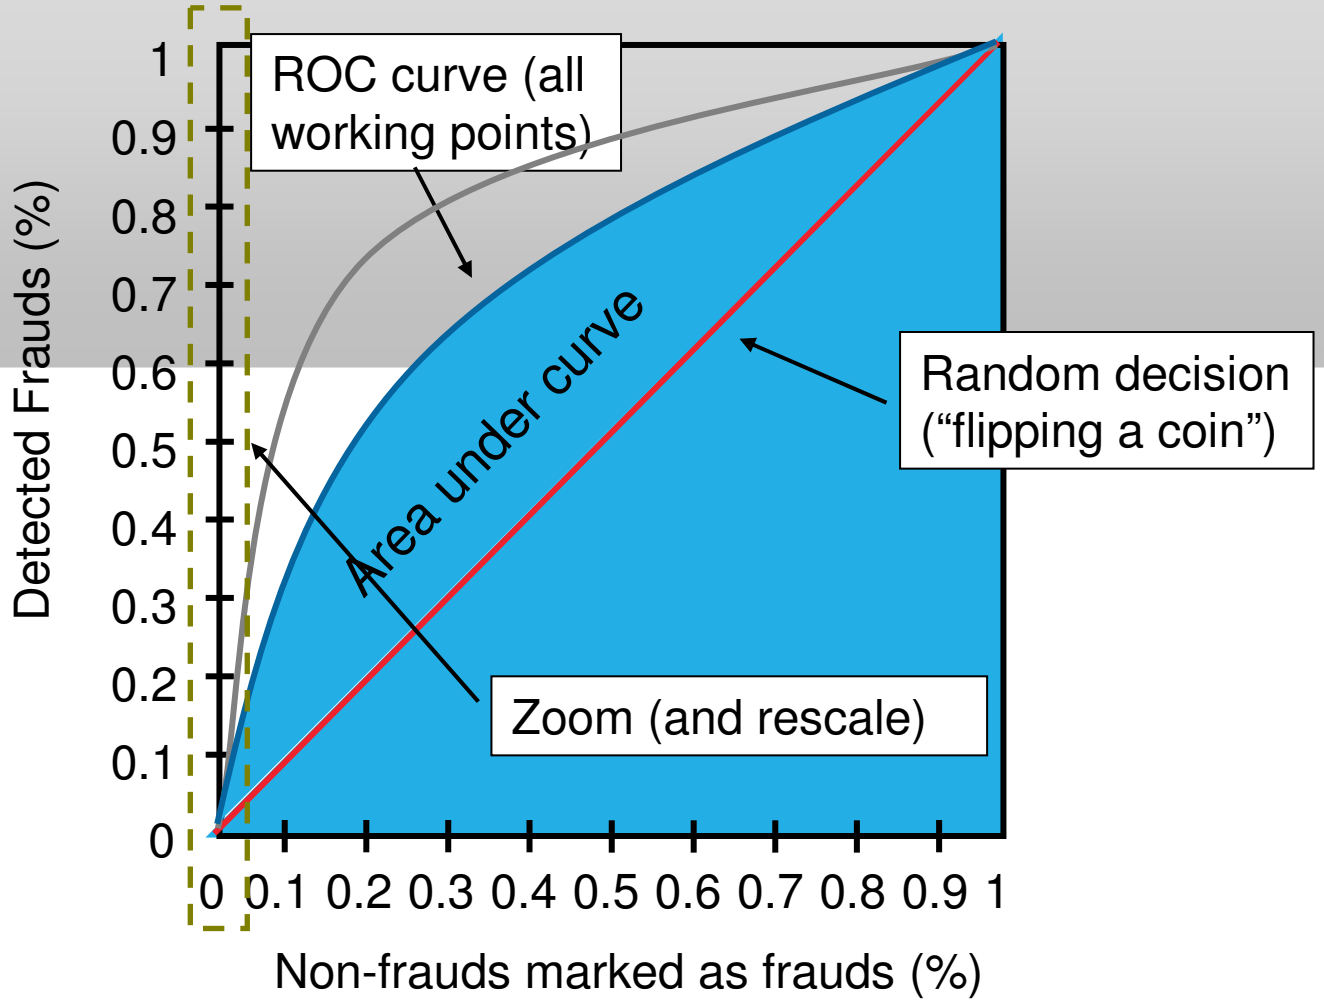

Case study 1: Credit cards – Unique capabilities

Outperforms a solution by a leading vendor

Differentiating capabilities:

1. Maximize value recovered
2. Pre-fraud detection identifies which activities precede fraud

Fraud/non-fraud prediction

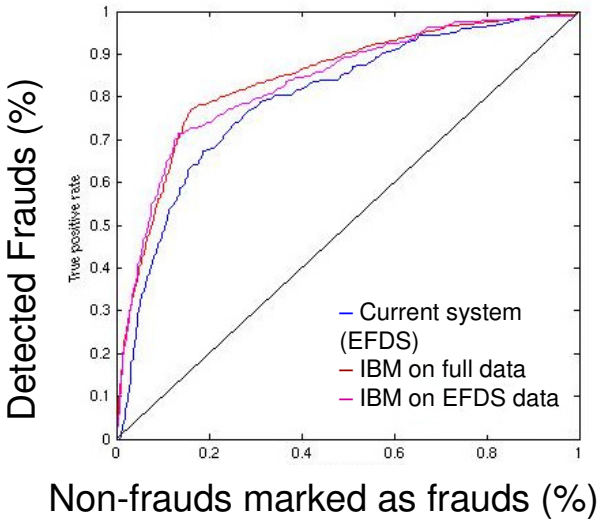

Hunt down costly frauds

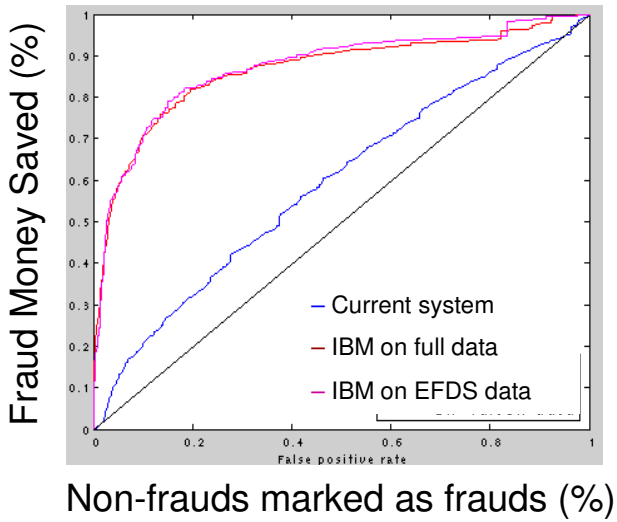

Pre-fraud detection

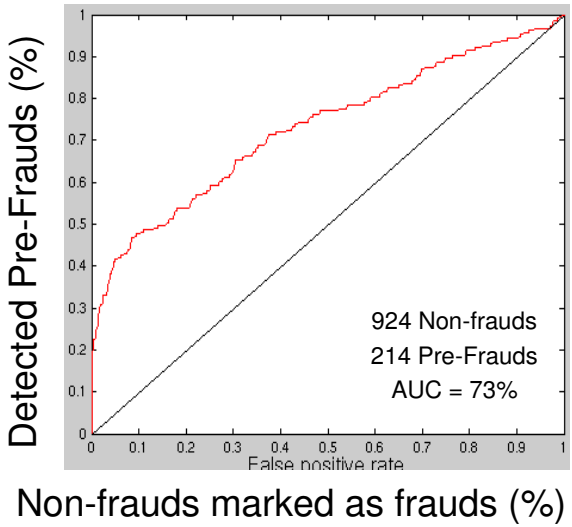

The value of integrating with an existing system

Improve results of existing systems by working in tandem to detect different types of fraud patterns

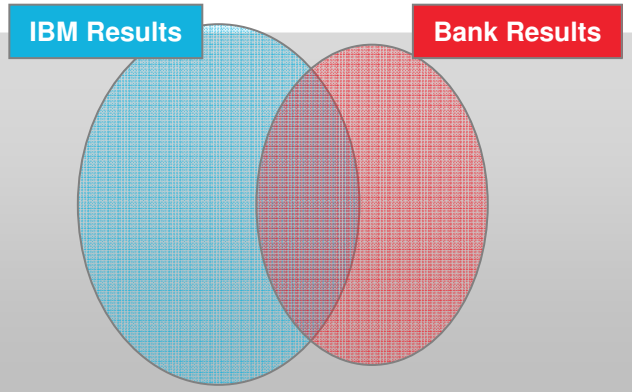

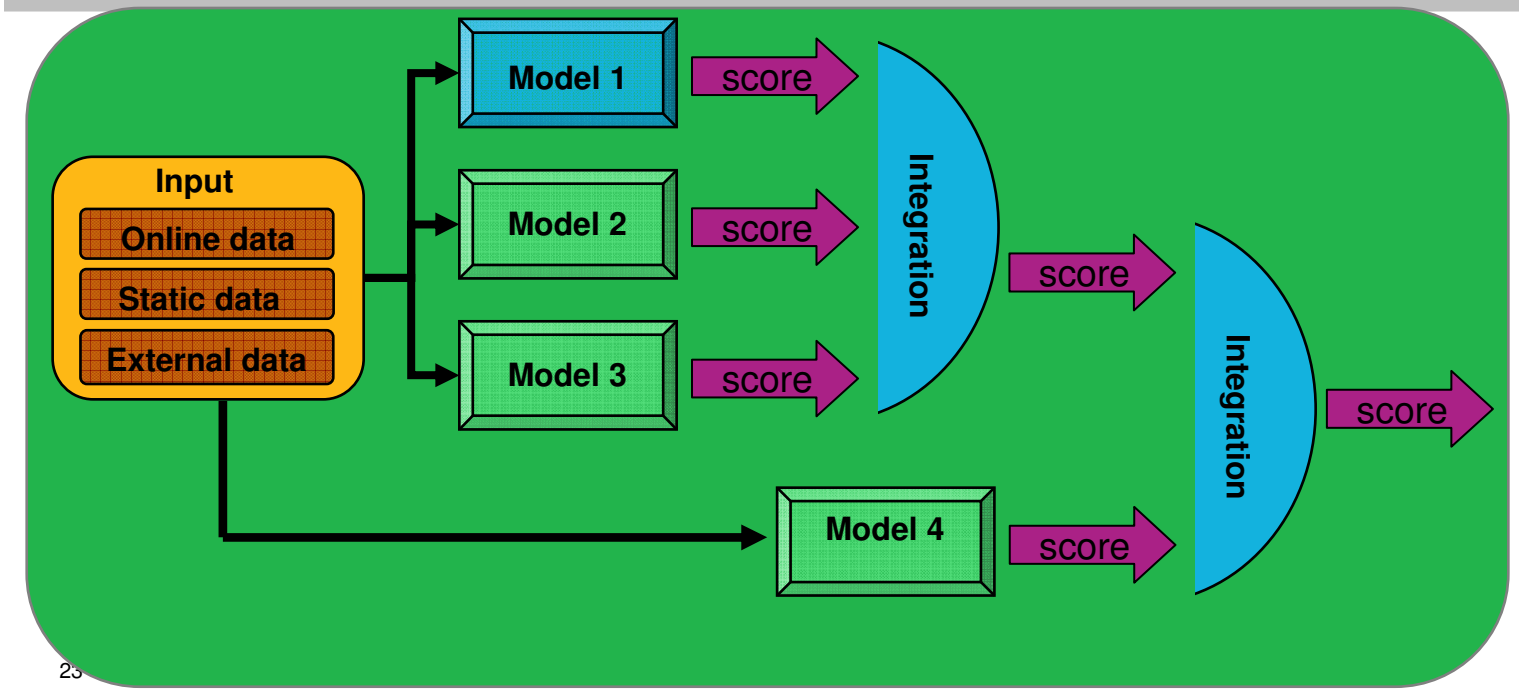

Case study 2: ATM – An integrated system

20% improvement in detected fraud (true positives) for the combined model at the current false positive rate of the existing system

-- OR --

50% reduction in false positives at the current fraud detection rate

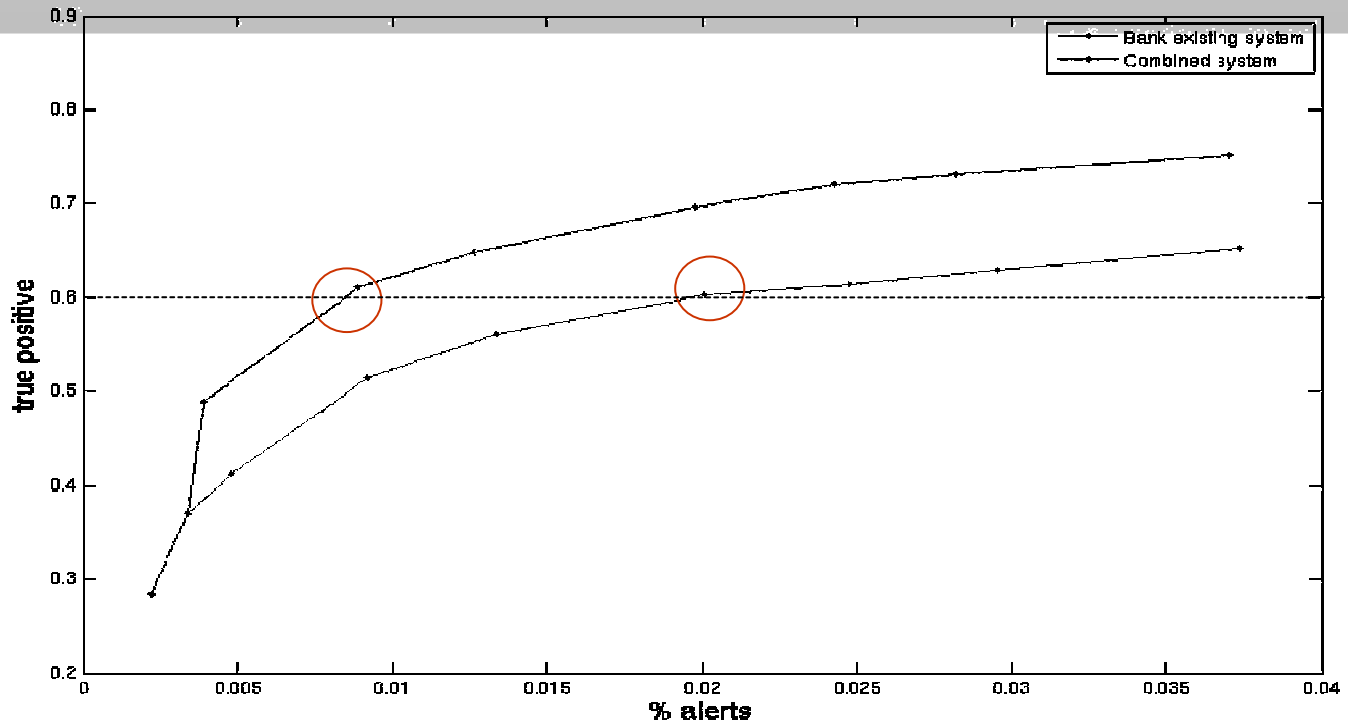

Case Study 3: Credit cards – An integrated system

15% improvement in detected fraud (true positives) for the combined model at the current false positive rate of the existing system

-- OR --

50% reduction in false positives at the current fraud detection rate

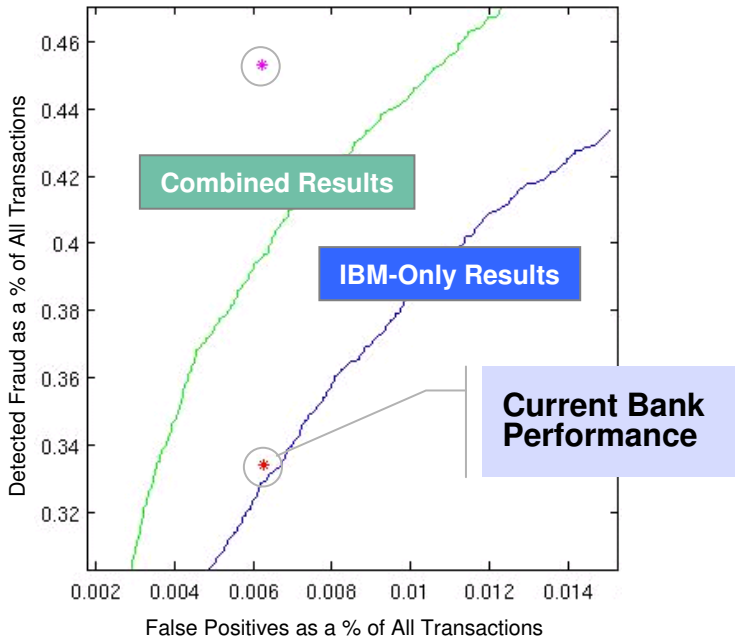

Case study 4: Credit cards – Chasing the money

The IBM model is **tuned to "chase the money"**

- Targets fraudulent transactions with high amounts
- Objective: maximize the value recovered

This strategy may cause lower TP and higher FP rates

Relative increase in amount saved

	Aug.	Sep.	Oct.
IBM "tuned"	57.2%	57.4%	68.6%
Integrated System	30.9%	43.9%	41%

Estimated annual increase in amount saved

- **IBM solution – $15M**
- **Integrated system – $10M**

Summary of capabilities

Credit cards (tier-1 bank in North America)

- 50% reduction in false positives
- 15% increase in detection level

ATM (one of the largest banks in Israel)

- 50% reduction in false positives
- 20% increase in detection level

E-payments (large European bank)

- 40% reduction in false positives

Advanced capabilities

- Chasing the money: estimated $15M savings
- Pre-fraud detection
- Integrating with existing systems

DRIFT summary

DRIFT is a real-time fraud detection system

Solution based on machine learning techniques

Outperforms offerings by leading vendors (proven in several PoCs)

Advanced capabilities (pre-fraud and costly fraud detection)

IBM Smarter Analytics Signature Solution – Anti-Fraud, Waste & Abuse