POLITECNICO DI MILANO

Association Rules

Information Retrieval and Data Mining

Prof. Matteo Matteucci

Learning Unsupervised Rules !?!

Market-Basket Transactions

Bread
Peanuts
Milk
Fruit
Jam

Bread
Jam
Soda
Chips
Milk
Fruit

Steak
Jam
Soda
Chips
Bread

Jam
Soda
Peanuts
Milk
Fruit

Jam
Soda
Chips
Milk
Bread

Fruit
Soda
Chips
Milk

Fruit
Soda
Peanuts
Milk

Fruit
Peanuts
Cheese
Yogurt

What is Frequent Pattern Mining?

- Frequent pattern: a pattern (a set of items, subsequences, substructures, etc.) that occurs frequently in a data set
- Motivation: Finding inherent regularities in data
  - What products were often purchased together? Beer and diapers?!
  - What are the subsequent purchases after a PC?
  - What kinds of DNA are sensitive to this new drug?
  - Can we automatically classify web documents?
- Applications
  - Basket data analysis, cross-marketing, catalog design, sale campaign analysis, Web log (click stream) analysis, DNA sequence analysis, etc.

What is Association Rule Mining?

TID	Items
1	Bread, Peanuts, Milk, Fruit, Jam
2	Bread, Jam, Soda, Chips, Milk, Fruit
3	Steak, Jam, Soda, Chips, Bread
4	Jam, Soda, Peanuts, Milk, Fruit
5	Jam, Soda, Chips, Milk, Bread
6	Fruit, Soda, Chips, Milk
7	Fruit, Soda, Peanuts, Milk
8	Fruit, Peanuts, Cheese, Yogurt

Examples

$\{bread\} \Rightarrow \{milk\}$
$\{soda\} \Rightarrow \{chips\}$
$\{bread\} \Rightarrow \{jam\}$

- Given a set of transactions, find rules that will predict the occurrence of an item based on the occurrences of other items in the transaction

Frequent Itemset

- **Itemset**
  - A collection of one or more items, e.g., {milk, bread, jam}
  - k-itemset, an itemset that contains k items
- **Support count (σ)**
  - Frequency of occurrence of an itemset
  - σ({Milk, Bread}) = 3
    σ({Soda, Chips}) = 4
- **Support**
  - Fraction of transactions that contain an itemset
  - s({Milk, Bread}) = 3/8
    s({Soda, Chips}) = 4/8
- **Frequent Itemset**
  - An itemset whose support is greater than or equal to a minsup threshold

TID	Items
1	Bread, Peanuts, Milk, Fruit, Jam
2	Bread, Jam, Soda, Chips, Milk, Fruit
3	Steak, Jam, Soda, Chips, Bread
4	Jam, Soda, Peanuts, Milk, Fruit
5	Jam, Soda, Chips, Milk, Bread
6	Fruit, Soda, Chips, Milk
7	Fruit, Soda, Peanuts, Milk
8	Fruit, Peanuts, Cheese, Yogurt

What is An Association Rule?

- Implication of the form X $\Rightarrow$ Y, where X and Y are itemsets
  - Example: {bread} $\Rightarrow$ {milk}
- Rule Evaluation Metrics, Suppor & Confidence
  - Support (s)
    - Fraction of transactions that contain both X and Y

$$s = \frac{\sigma(\{\text{Bread}, \text{Milk}\})}{\#\text{ of transactions}} = 0.38$$

  - Confidence (c)
    - Measures how often items in Y appear in transactions that contain X

$$c = \frac{\sigma(\{\text{Bread}, \text{Milk}\})}{\sigma(\{\text{Bread}\})} = 0.75$$

Support and Confidence Meaning

TID	Items
1	Bread, Peanuts, Milk, Fruit, Jam
2	Bread, Jam, Soda, Chips, Milk, Fruit
3	Steak, Jam, Soda, Chips, Bread
4	Jam, Soda, Peanuts, Milk, Fruit
5	Jam, Soda, Chips, Milk, Bread
6	Fruit, Soda, Chips, Milk
7	Fruit, Soda, Peanuts, Milk
8	Fruit, Peanuts, Cheese, Yogurt

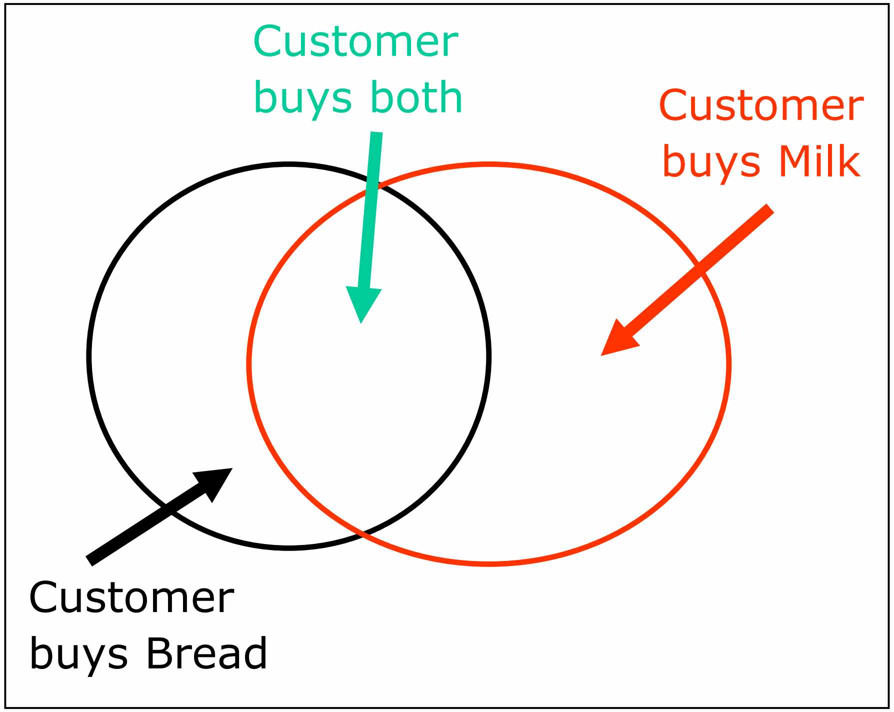

Support (s) = P(X,Y)

$$s = \frac{\sigma(\{\text{Bread}, \text{Milk}\})}{\#\text{ of transactions}} = 0.38$$

Confidence (c) = P(X,Y)/P(X)
= P(Y|X)

$$c = \frac{\sigma(\{\text{Bread}, \text{Milk}\})}{\sigma(\{\text{Bread}\})} = 0.75$$

What is the Goal?

- Given a set of transactions T, the goal of association rule mining is to find all rules having
  - support ≥ minsup threshold
  - confidence ≥ minconf threshold

- Brute-force approach
  - List all possible association rules
  - Compute the support and confidence for each rule
  - Prune rules that fail the minsup and minconf thresholds

- Brute-force approach is computationally prohibitive!

Mining Association Rules

{Bread, Jam} ⇒ {Milk}: s=3/8 c=3/4

{Bread, Milk} ⇒ {Jam}: s=3/8 c=3/3

{Milk, Jam} ⇒ {Bread}: s=3/8 c=3/3

{Bread} ⇒ {Milk, Jam}: s=3/8 c=3/4

{Jam} ⇒ {Bread, Milk}: s=3/8 c=3/5

{Milk} ⇒ {Bread, Jam}: s=3/8 c=3/6

TID	Items
1	Bread, Peanuts, Milk, Fruit, Jam
2	Bread, Jam, Soda, Chips, Milk, Fruit
3	Steak, Jam, Soda, Chips, Bread
4	Jam, Soda, Peanuts, Milk, Fruit
5	Jam, Soda, Chips, Milk, Bread
6	Fruit, Soda, Chips, Milk
7	Fruit, Soda, Peanuts, Milk
8	Fruit, Peanuts, Cheese, Yogurt

- All bove rules are binary partitions of the same itemset:
  {Milk, Bread, Jam}
- Rules originating from the same itemset have identical support but can have different confidence
- Decouple the support and confidence requirements!

Mining Association Rules in Two Steps

1. Frequent Itemset Generation
   - Generate all itemsets whose support $\geq$ minsup
2. Rule Generation
   - Generate high confidence rules from frequent itemset
   - Each rule is a binary partitioning of a frequent itemset

However frequent itemset generation is computationally expensive!

Frequent Itemset Generation

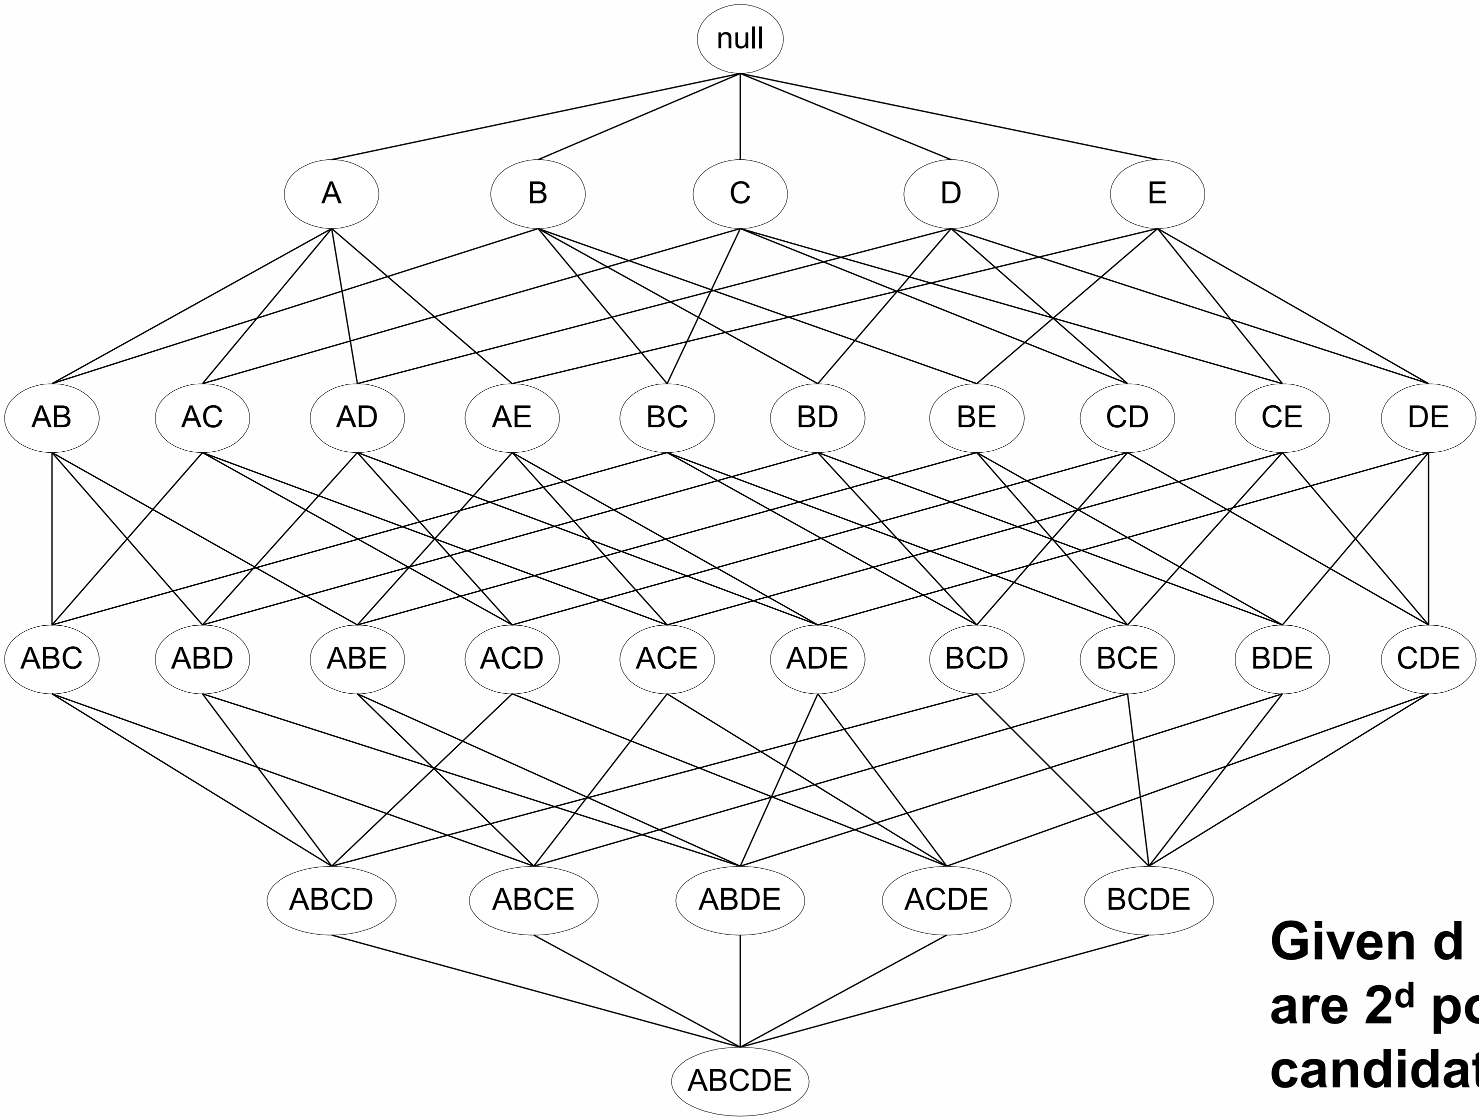

Given d items, there are 2^d possible candidate itemsets

Frequent Itemset Generation

- Brute-force approach:
  - Each itemset in the lattice is a candidate frequent itemset
  - Count the support of each candidate by scanning the database

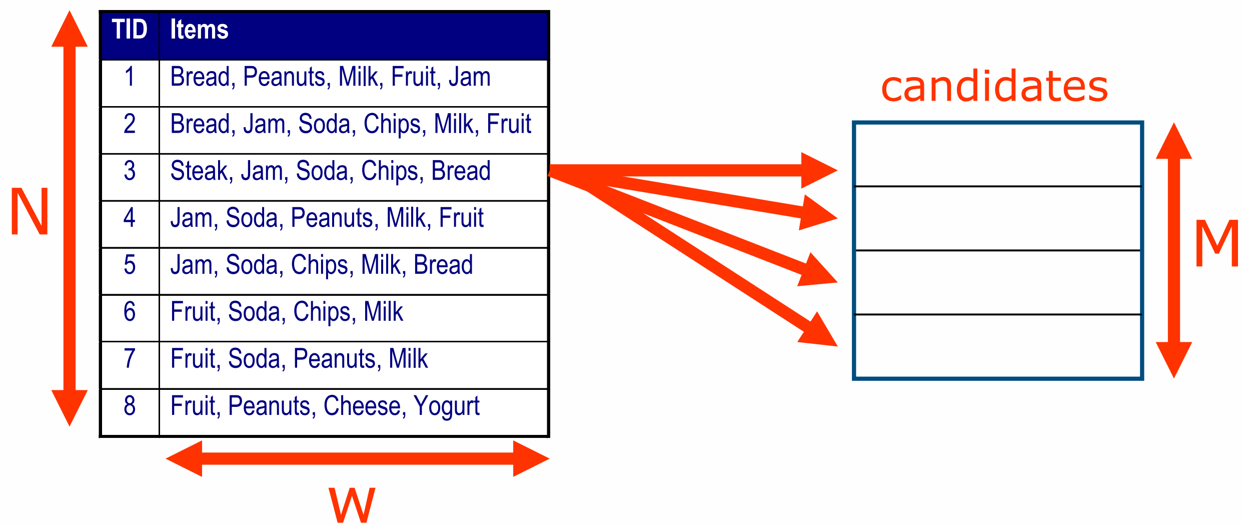

TID	Items
1	Bread, Peanuts, Milk, Fruit, Jam
2	Bread, Jam, Soda, Chips, Milk, Fruit
3	Steak, Jam, Soda, Chips, Bread
4	Jam, Soda, Peanuts, Milk, Fruit
5	Jam, Soda, Chips, Milk, Bread
6	Fruit, Soda, Chips, Milk
7	Fruit, Soda, Peanuts, Milk
8	Fruit, Peanuts, Cheese, Yogurt

  - Match each transaction against every candidate
  - Complexity ~ O(NMw) => Expensive since $M = 2^d$

Computational Complexity

- Given d unique items:
  - Number of itemsets: 2^d
  - Number of possible association rules: $$\sum_{k=1}^{d-1}\left[\binom{d}{k} \times \sum_{j=1}^{d-k}\binom{d-k}{j}\right] = 3^d - 2^{d+1} + 1$$

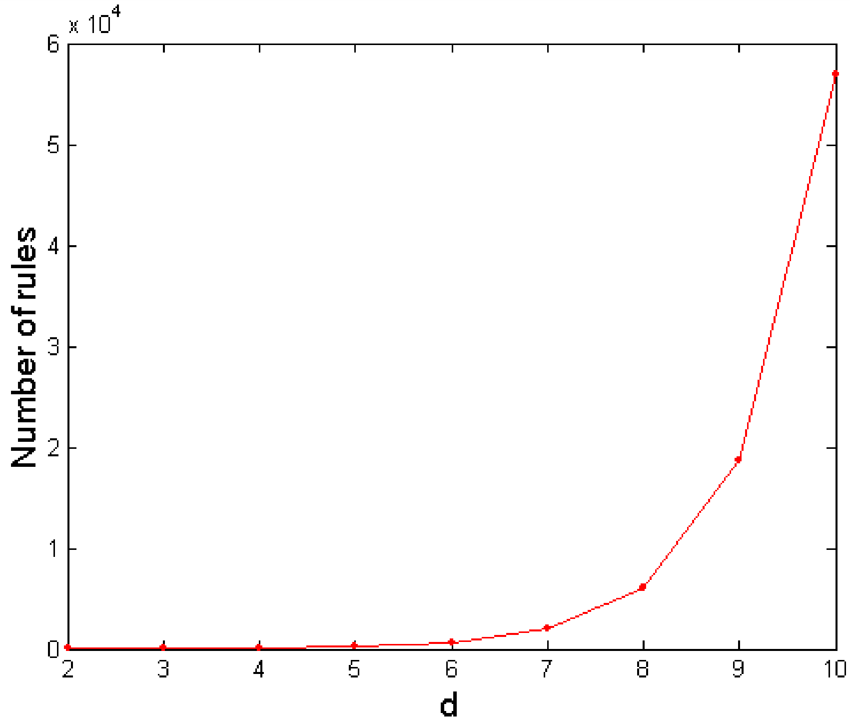

- For d=6, there are 602 rules

Frequent Itemset Generation Strategies

- Reduce the number of candidates (M)
  - Complete search: $M=2^d$
  - Use pruning techniques to reduce M
- Reduce the number of transactions (N)
  - Reduce size of N as the size of itemset increases
- Reduce the number of comparisons (NM)
  - Use efficient data structures to store the candidates or transactions
  - No need to match every candidate against every transaction

Reducing the Number of Candidates

- Apriori principle
  - If an itemset is frequent, then all of its subsets must also be frequent
- Apriori principle holds due to the following property of the support measure:

$$\forall X, Y : (X \subseteq Y) \Rightarrow s(X) \geq s(Y)$$

  - Support of an itemset never exceeds the support of its subsets, this is known as the anti-monotone property of support

Illustrating Apriori Principle

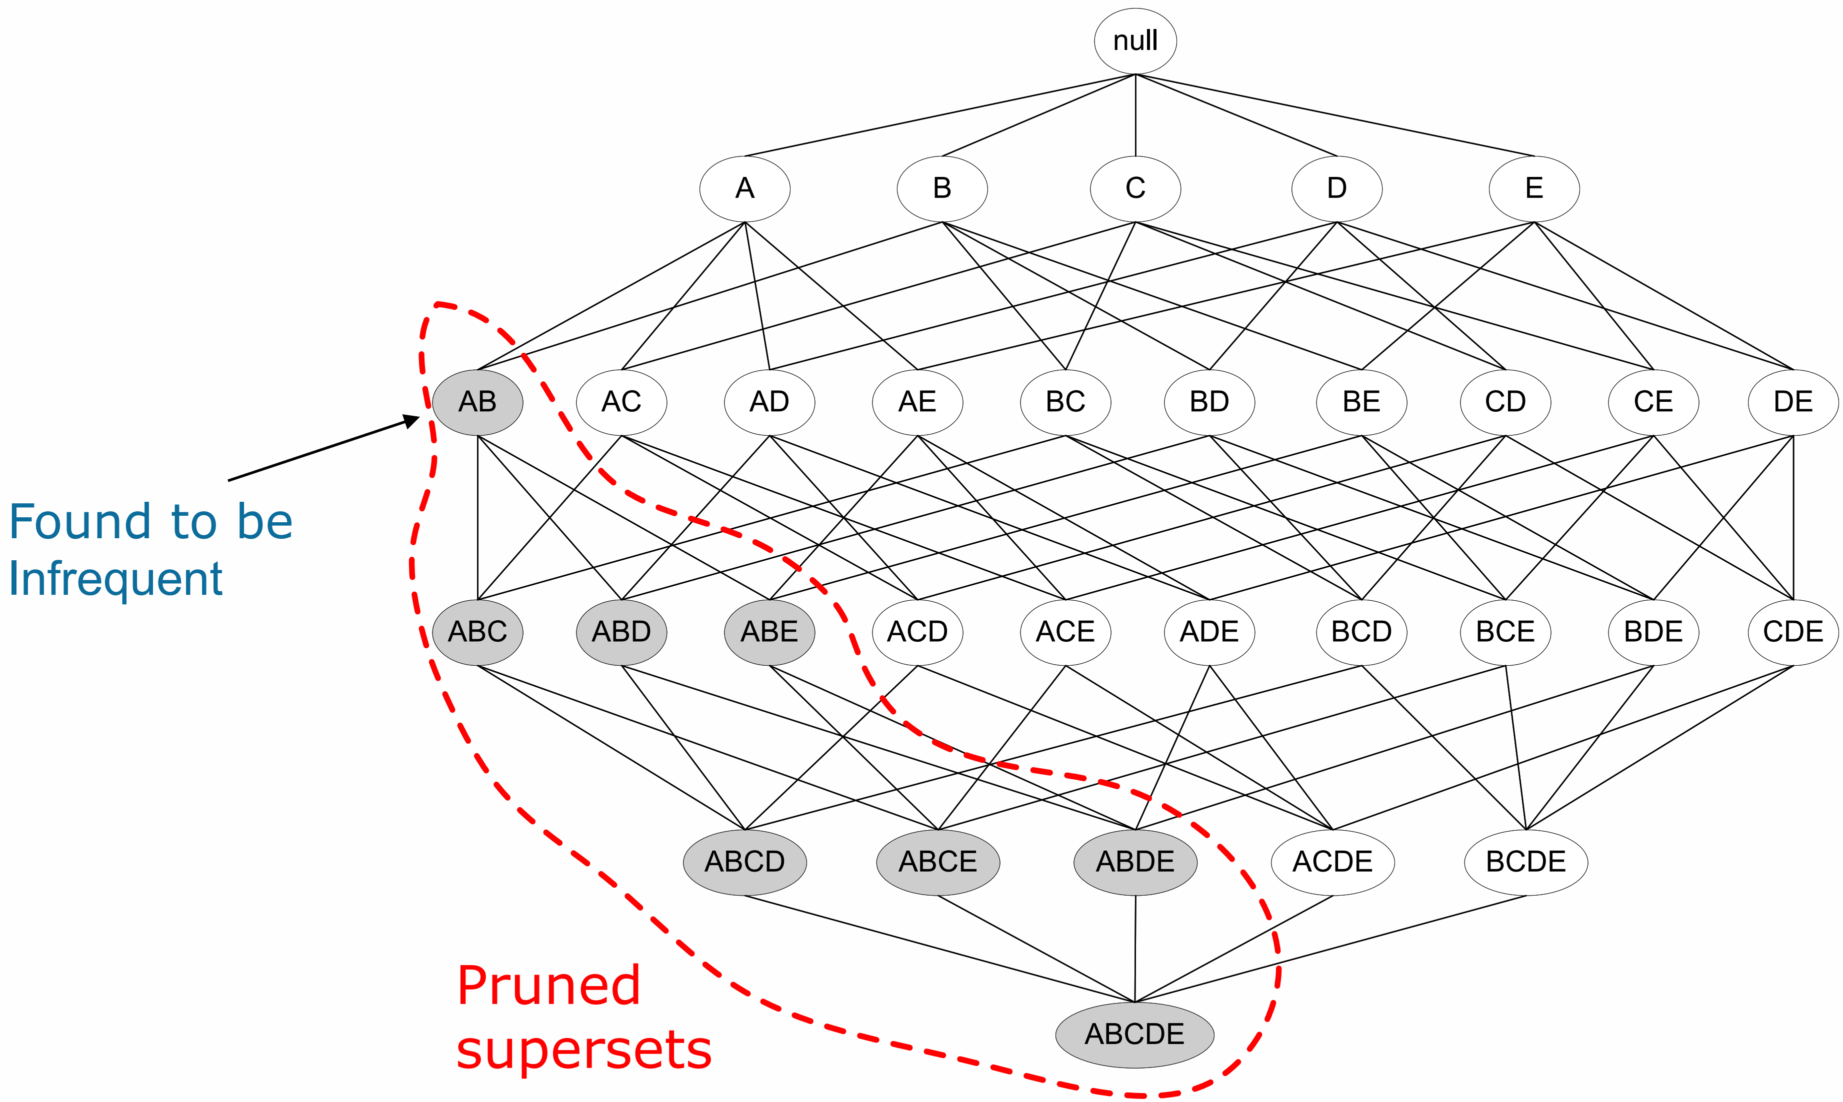

How does the Apriori principle work?

Items (1-itemsets)

Item	Count
Bread	4
Peanuts	4
Milk	6
Fruit	6
Jam	5
Soda	6
Chips	4
~~Steak~~	~~1~~
~~Cheese~~	~~1~~
~~Yogurt~~	~~1~~

Minimum Support = 4

2-itemsets

2-Itemset	Count
Bread, Jam	4
Peanuts, Fruit	4
Milk, Fruit	5
Milk, Jam	4
Milk, Soda	5
Fruit, Soda	4
Jam, Soda	4
Soda, Chips	4

3-itemsets

3-Itemset	Count
Milk, Fruit, Soda	4

The Apriori Algorithm Idea

- Let k=1
- Generate frequent itemsets of length 1
- Repeat until no new frequent itemsets are identified
  - Generate length (k+1) candidate itemsets from length k frequent itemsets
  - Prune candidate itemsets containing subsets of length k that are infrequent
  - Count each candidate support by scanning the DB
  - Eliminate candidates that are infrequent, leaving only those that are frequent

Important Details on Apriori Algorithm

- How to generate candidates?
  - Step 1: self-joining L_k
  - Step 2: pruning

- Example of Candidate-generation
  - L_3={*abc, abd, acd, ace, bcd*}
  - Self-joining: L_3*L_3
    - *abcd* from *abc* and *abd*
    - *acde* from *acd* and *ace*
  - Pruning:
    - *acde* is removed because *ade* is not in L_3
  - C_4={*abcd*}

The Apriori Algorithm

C_k: Candidate itemset of size k
L_k : frequent itemset of size k

L_1 = {frequent items};
for (k = 1; L_k !=$\emptyset$; k++) **do begin**
  C_{k+1} = candidates generated from L_k;
  for each transaction t in database do
      increment the count of all candidates in C_{k+1}
      that are contained in t
  L_{k+1} = candidates in C_{k+1} with min_support
  end
return $\cup_k L_k$;

An Example of Frequent Itemset

$Sup_{min} = 2$

Database TDB

Tid	Items
10	A, C, D
20	B, C, E
30	A, B, C, E
40	B, E

1st scan →

C_1

Itemset	sup
{A}	2
{B}	3
{C}	3
{D}	1
{E}	3

→ L_1

Itemset	sup
{A}	2
{B}	3
{C}	3
{E}	3

C_2

Itemset
{A, B}
{A, C}
{A, E}
{B, C}
{B, E}
{C, E}

← 2nd scan

C_2

Itemset	sup
{A, B}	1
{A, C}	2
{A, E}	1
{B, C}	2
{B, E}	3
{C, E}	2

← L_2

Itemset	sup
{A, C}	2
{B, C}	2
{B, E}	3
{C, E}	2

C_3

Itemset
{B, C, E}

3rd scan → L_3

Itemset	sup
{B, C, E}	2

Strategies to Improve Apriori Efficiency

- Hash-based itemset counting:
  - A k-itemset whose corresponding hashing bucket count is below the threshold cannot be frequent
- Transaction reduction:
  - A transaction that does not contain any frequent k-itemset is useless in subsequent scans
- Partitioning:
  - Any itemset that is potentially frequent in DB must be frequent in at least one of the partitions of DB
- Sampling:
  - mining on a subset of given data, lower support threshold + a method to determine the completeness
- Dynamic itemset counting:
  - add new candidate itemsets only when all of their subsets are estimated to be frequent

Rule Generation

- Given a frequent itemset L, find all non-empty subsets $f \subset L$ such that $f \rightarrow L - f$ satisfies the minimum confidence requirement

- If {A,B,C,D} is a frequent itemset, candidate rules are:

  ABC →D, ABD →C, ACD →B, BCD →A, A →BCD, B →ACD, C →ABD, D →ABC, AB →CD, AC → BD, AD → BC, BC →AD, BD →AC, CD →AB

- If $|L| = k$, then there are $2^k - 2$ candidate association rules (ignoring $L \rightarrow \emptyset$ and $\emptyset \rightarrow L$)

How to efficiently generate rules from frequent itemsets?

- Confidence does not have an anti-monotone property

$$c(ABC \rightarrow D) \text{ can be larger or smaller than } c(AB \rightarrow D)$$

- But confidence of rules generated from the same itemset has an anti-monotone property

$$L = \{A,B,C,D\}: c(ABC \rightarrow D) \geq c(AB \rightarrow CD) \geq c(A \rightarrow BCD)$$

- Confidence is anti-monotone with respect to the number of items on the right hand side of the rule

Rule Generation for Apriori Algorithm

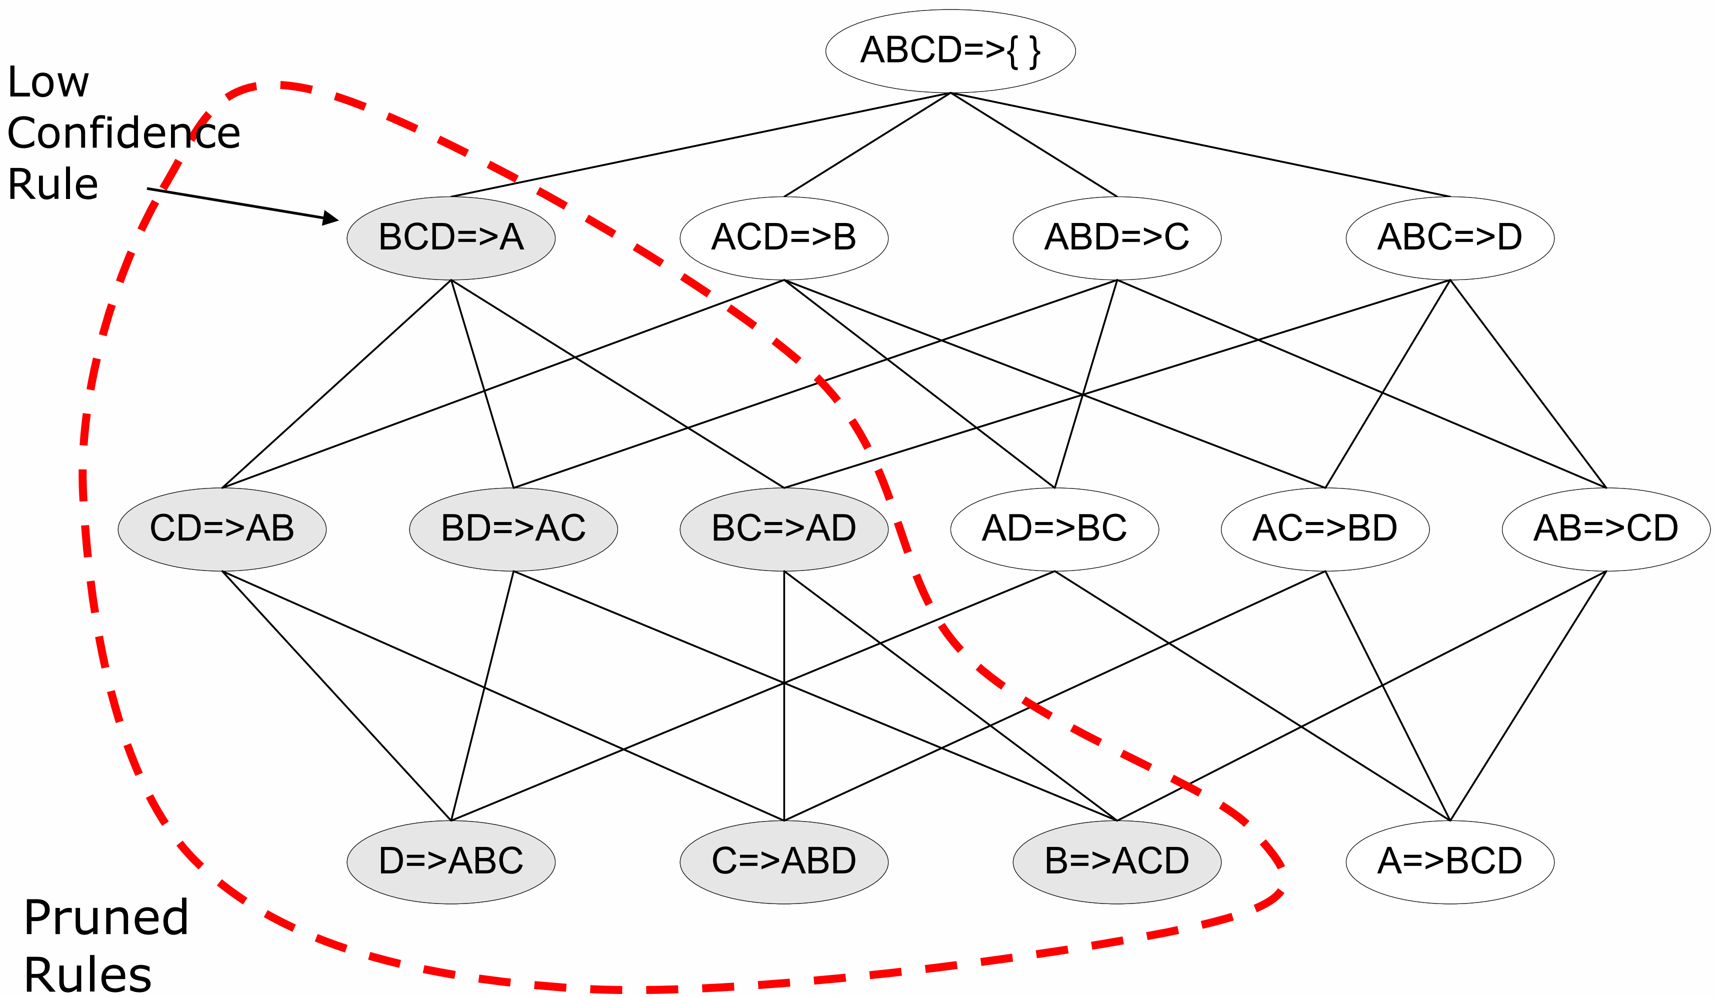

Rule Generation for Apriori Algorithm

Candidate rule is generated by merging two rules that share the same prefix in the rule consequent

- join(CD=>AB,BD=>AC) would produce the candidate rule D => ABC
- Prune rule D=>ABC if its subset AD=>BC does not have high confidence

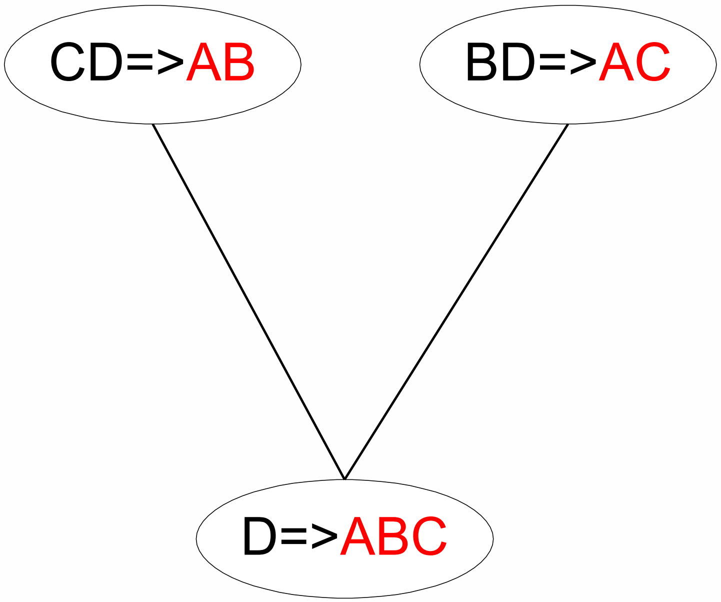

Example with Weka: Formatting the data

```
%
% Example of market basket data
%
@relation 'basket'
@attribute Bread {1}
@attribute Peanuts {1}
@attribute Milk {1}
@attribute Fruit {1}
@attribute Jam  {1}
@attribute Soda  {1}
@attribute Chips {1}
@attribute Steak {1}
@attribute Cheese {1}
@attribute Yogurt {1}
```

```
@data
1,1,1,1,1,?,?,?,?,?
1,?,1,1,1,1,1,?,?,?
1,?,?,?,1,1,1,1,?,?
?,1,1,1,1,1,?,?,?,?
1,?,1,?,1,1,1,?,?,?
?,?,1,1,?,1,1,?,?,?
?,1,1,1,?,1,?,?,?,?
?,1,?,1,?,?,?,?,1,1
```

Effect of Support Distribution

- Many real data sets have skewed support distribution

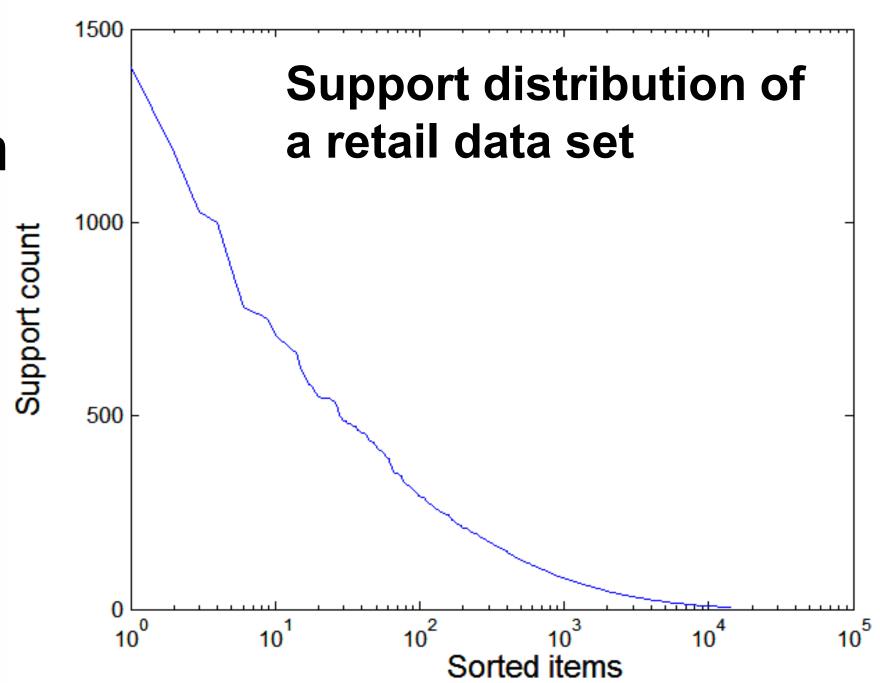

- If minsup is set too high, we could miss itemsets involving interesting rare (expensive)items
- If minsup is set too low, apriori becomes computationally expensive and the number of itemsets very large
- A single minimum support threshold may not be effective

Are there Interesting/Useful Rules?

Anything that is <u>interesting</u> happens significantly more than you would expect by chance.

<u>Example:</u> basic statistical analysis of basket data may show that 10% of baskets contain bread, and 4% of baskets contain washing-up powder.

What is the probability of a basket containing both bread *and* washing-up powder? The laws of probability say:

- if you choose a basket at random:
  - There is a probability 0.1 that it contains bread.
  - There is a probability 0.04 that it contains washing-up powder.
- If these two are independent:
  - There is a probability 0.1 * 0.04 = 0.004 it contains both

Are there Interesting/Useful Rules?

Anything that is interesting happens significantly more than you would expect by chance.

Example: basic statistical analysis of basket data may show that 10% of baskets contain bread, and 4% of baskets contain washing-up powder.

We have a prior expectation that just 4 baskets in 1000 should contain **both** bread and washing up powder:

- If we discover that really it is 20 in 1000 baskets, then we will be very surprised.
- Something is going on in shoppers' minds: bread and washing-up powder are connected in some way.
- There may be ways to exploit this discovery …

Another Example

ID	apples	beer	cheese	dates	eggs	fish	glue	honey	cream
1	1	1		1			1	1	
2			1	1	1				
3		1	1			1			
4		1				1			1
5					1		1		
6						1			1
7	1			1				1	
8						1			1
9			1		1				
10		1					1		
11					1		1		
12	1								
13			1			1			
14			1			1			
15								1	1
16				1					
17	1					1			
18	1	1	1	1				1	
19	1	1		1			1	1	
20					1				