

CMU SCS

Carnegie Mellon Univ. Dept. of Computer Science 15-415 - Database Applications

Data Warehousing / Data Mining
(R&G, ch 25 and 26)

CMU SCS

Data mining - detailed outline

- ➔ Problem
- Getting the data: Data Warehouses, DataCubes, OLAP
- Supervised learning: decision trees
- Unsupervised learning
 - association rules
 - (clustering)

Faloutsos CMU SCS 15-415 2

CMU SCS

Problem

Given: multiple data sources
Find: patterns (classifiers, rules, clusters, outliers...)

NY

sales(p-id, c-id, date, \$price)

SF

customers(c-id, age, income, ...)

???

PGH

Faloutsos CMU SCS 15-415 3

CMU SCS

Data Ware-housing

First step: collect the data, in a single place (= Data Warehouse)

How?
How often?
How about discrepancies / non-homegeneities?

Faloutsos CMU SCS 15-415 4

CMU SCS

Data Ware-housing

First step: collect the data, in a single place (= Data Warehouse)

How? A: Triggers/Materialized views
How often? A: [Art!]
How about discrepancies / non-homegeneities? A: Wrappers/Mediators

Faloutsos CMU SCS 15-415 5

CMU SCS

Data Ware-housing

Step 2: collect counts. (DataCubes/OLAP)
Eg.:

Faloutsos CMU SCS 15-415 6

OLAP

Problem: "is it true that shirts in large sizes sell better in dark colors?"

sales	ci-d	p-id	Size	Color	\$	C / S	S	M	L	TOT
				Blue	30	Red	20	3	5	28
	C10	Shirt	L	Blue	30	Blue	3	3	8	14
	C10	Pants	XL	Red	50	Gray	0	0	5	5
	C20	Shirt	XL	White	20	TOT	23	6	18	47
...										

Faloutsos CMU SCS 15-415 7

DataCubes

'color', 'size': DIMENSIONS
'count': MEASURE

C / S	S	M	L	TOT
Red	20	3	5	28
Blue	3	3	8	14
Gray	0	0	5	5
TOT	23	6	18	47

Faloutsos CMU SCS 15-415 8

DataCubes

'color', 'size': DIMENSIONS
'count': MEASURE

C / S	S	M	L	TOT
Red	20	3	5	28
Blue	3	3	8	14
Gray	0	0	5	5
TOT	23	6	18	47

Faloutsos CMU SCS 15-415 9

DataCubes

'color', 'size': DIMENSIONS
'count': MEASURE

C / S	S	M	L	TOT
Red	20	3	5	28
Blue	3	3	8	14
Gray	0	0	5	5
TOT	23	6	18	47

Faloutsos CMU SCS 15-415 10

DataCubes

'color', 'size': DIMENSIONS
'count': MEASURE

C / S	S	M	L	TOT
Red	20	3	5	28
Blue	3	3	8	14
Gray	0	0	5	5
TOT	23	6	18	47

Faloutsos CMU SCS 15-415 11

DataCubes

'color', 'size': DIMENSIONS
'count': MEASURE

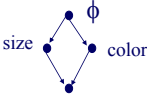
C / S	S	M	L	TOT
Red	20	3	5	28
Blue	3	3	8	14
Gray	0	0	5	5
TOT	23	6	18	47

Faloutsos CMU SCS 15-415 12

CMU SCS

DataCubes

'color', 'size': DIMENSIONS
'count': MEASURE



C / S	S	M	L	TOT
Red	20	3	5	28
Blue	3	3	8	14
Gray	0	0	5	5
TOT	23	6	18	47

DataCube

Faloutsos CMU SCS 15-415 13

CMU SCS

DataCubes

SQL query to generate DataCube:

- Naively (and painfully):


```
select size, color, count(*)
from sales where p-id = 'shirt'
group by size, color
```

```
select size, count(*)
from sales where p-id = 'shirt'
group by size
```

...

Faloutsos CMU SCS 15-415 14

CMU SCS

DataCubes

SQL query to generate DataCube:

- with 'cube by' keyword:


```
select size, color, count(*)
from sales
where p-id = 'shirt'
cube by size, color
```

Faloutsos CMU SCS 15-415 15

CMU SCS

DataCubes

DataCube issues:

Q1: How to store them (and/or materialize portions on demand)

Q2: Which operations to allow

Faloutsos CMU SCS 15-415 16

CMU SCS

DataCubes

DataCube issues:

Q1: How to store them (and/or materialize portions on demand) A: ROLAP/MOLAP

Q2: Which operations to allow A: roll-up, drill down, slice, dice

[More details: book by Han+Kamber]

Faloutsos CMU SCS 15-415 17

CMU SCS

DataCubes

Q1: How to store a dataCube?

C / S	S	M	L	TOT
Red	20	3	5	28
Blue	3	3	8	14
Gray	0	0	5	5
TOT	23	6	18	47

Faloutsos CMU SCS 15-415 18

DataCubes

Q1: How to store a dataCube?
 A1: Relational (R-OLAP)

Color	Size	count	C/S	S	M	L	TOT
Red	'all'	47	Red	20	3	5	28
Blue	'all'	14	Blue	3	3	8	14
Blue	M	3	Gray	0	0	5	5
Blue	M	3	TOT	23	6	18	47

...

Faloutsos CMU SCS 15-415 19

DataCubes

Q1: How to store a dataCube?
 A2: Multi-dimensional (M-OLAP)
 A3: Hybrid (H-OLAP)

C/S	S	M	L	TOT
Red	20	3	5	28
Blue	3	3	8	14
Gray	0	0	5	5
TOT	23	6	18	47

Faloutsos CMU SCS 15-415 20

DataCubes

Pros/Cons:
 ROLAP strong points: (DSS, Metacube)

Faloutsos CMU SCS 15-415 21

DataCubes

Pros/Cons:
 ROLAP strong points: (DSS, Metacube)

- use existing RDBMS technology
- scale up better with dimensionality

Faloutsos CMU SCS 15-415 22

DataCubes

Pros/Cons:
 MOLAP strong points: (EssBase/hyperion.com)

- faster indexing (careful with: high-dimensionality; sparseness)

HOLAP: (MS SQL server OLAP services)

- detail data in ROLAP; summaries in MOLAP

Faloutsos CMU SCS 15-415 23

DataCubes

Q1: How to store a dataCube
 Q2: What operations should we support?

Faloutsos CMU SCS 15-415 24

CMU SCS

DataCubes

Q2: What operations should we support?

C / S	S	M	L
Red	20	3	5
Blue	3	3	8
Gray	0	0	5

Faloutsos CMU SCS 15-415 25

CMU SCS

DataCubes

Q2: What operations should we support?

Roll-up

C / S	S	M	L
TOT	23	6	18

Faloutsos CMU SCS 15-415 26

CMU SCS

DataCubes

Q2: What operations should we support?

Drill-down

C / S	S	M	L
TOT	23	6	18

Faloutsos CMU SCS 15-415 27

CMU SCS

DataCubes

Q2: What operations should we support?

Slice

C / S	S	M	L
Red	20	3	5
Blue	3	3	8
Gray	0	0	5

Faloutsos CMU SCS 15-415 28

CMU SCS

DataCubes

Q2: What operations should we support?

Dice

C / S	S	M	L
Red	20	3	5
Blue	3	3	8
Gray	0	0	5

Faloutsos CMU SCS 15-415 29

CMU SCS

DataCubes

Q2: What operations should we support?

- Roll-up
- Drill-down
- Slice
- Dice
- (Pivot/rotate; drill-across; drill-through
- top N
- moving averages, etc)

Faloutsos CMU SCS 15-415 30

CMU SCS

D/W - OLAP - Conclusions

- D/W: copy (summarized) data + analyze
- OLAP - concepts:
 - DataCube
 - R/M/H-OLAP servers
 - 'dimensions'; 'measures'

Faloutsos CMU SCS 15-415 31

CMU SCS

Outline

- Problem
- Getting the data: Data Warehouses, DataCubes, OLAP
- ➔ Supervised learning: decision trees
- Unsupervised learning
 - association rules
 - (clustering)

Faloutsos CMU SCS 15-415 32

CMU SCS

Decision trees - Problem

Age	Chol-level	Gender	...	CLASS-ID
30	150	M		+
				...
				-
				??

Faloutsos CMU SCS 15-415 33

CMU SCS

Decision trees

- Pictorially, we have

Faloutsos CMU SCS 15-415 34

CMU SCS

Decision trees

- and we want to label '?'

Faloutsos CMU SCS 15-415 35

CMU SCS

Decision trees

- so we build a decision tree:

Faloutsos CMU SCS 15-415 36

CMU SCS

Decision trees

- so we build a decision tree:

Faloutsos CMU SCS 15-415 37

CMU SCS

skip

Outline

- Problem
- Getting the data: Data Warehouses, DataCubes, OLAP
- Supervised learning: decision trees
 - problem
 - approach
 - scalability enhancements
- Unsupervised learning
 - association rules
 - (clustering)

Faloutsos CMU SCS 15-415 38

CMU SCS

skip

Decision trees

- Typically, two steps:
 - tree building
 - tree pruning (for over-training/over-fitting)

Faloutsos CMU SCS 15-415 39

CMU SCS

skip

Tree building

- How?

Faloutsos CMU SCS 15-415 40

CMU SCS

skip

Tree building

- How?
- A: Partition, recursively - pseudocode:
 - Partition (Dataset S)
 - if all points in S have same label
 - then return
 - evaluate splits along each attribute A
 - pick best split, to divide S into S1 and S2
 - Partition(S1); Partition(S2)

Faloutsos CMU SCS 15-415 41

CMU SCS

skip

Tree building

- Q1: how to introduce splits along attribute A_i
- Q2: how to evaluate a split?

Faloutsos CMU SCS 15-415 42

CMU SCS skip

Tree building

- Q1: how to introduce splits along attribute A_i
- A1:
 - for num. attributes:
 - binary split, or
 - multiple split
 - for categorical attributes:
 - compute all subsets (expensive!), or
 - use a greedy algo

Faloutsos CMU SCS 15-415 43

CMU SCS skip

Tree building

- Q1: how to introduce splits along attribute A_i
- ➔ • Q2: how to evaluate a split?

Faloutsos CMU SCS 15-415 44

CMU SCS skip

Tree building

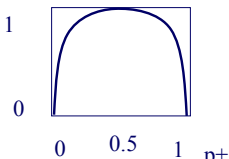
- Q1: how to introduce splits along attribute A_i
- ➔ • Q2: how to evaluate a split?
- A: by how close to uniform each subset is - ie., we need a measure of uniformity:

Faloutsos CMU SCS 15-415 45

CMU SCS skip

Tree building

entropy: $H(p+, p-)$ Any other measure?

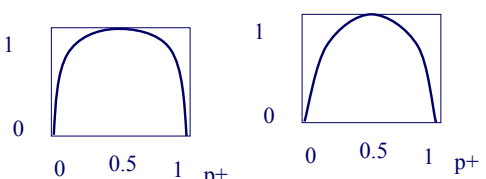


Faloutsos CMU SCS 15-415 46

CMU SCS skip

Tree building

entropy: $H(p+, p)$ ‘gini’ index: $1 - p_+^2 - p^2$



Faloutsos CMU SCS 15-415 47

CMU SCS skip

Tree building

entropy: $H(p+, p)$ ‘gini’ index: $1 - p_+^2 - p^2$

(How about multiple labels?)

Faloutsos CMU SCS 15-415 48

CMU SCS skip

Tree building

Intuition:

- entropy: #bits to encode the class label
- gini: classification error, if we randomly guess '+' with prob. p_+

Faloutsos CMU SCS 15-415 49

CMU SCS skip

Tree building

Thus, we choose the split that reduces entropy/classification-error the most: Eg.:

num. attr#2
(eg., chol-level)

num. attr#1 (eg., 'age')

Faloutsos CMU SCS 15-415 50

CMU SCS skip

Tree building

- Before split: we need $(n_+ + n_-) * H(p_+, p_-) = (7+6) * H(7/13, 6/13)$ bits total, to encode all the class labels
- After the split we need:
 - 0 bits for the first half and
 - $(2+6) * H(2/8, 6/8)$ bits for the second half

Faloutsos CMU SCS 15-415 51

CMU SCS skip

Tree pruning

- What for?

num. attr#2
(eg., chol-level)

num. attr#1 (eg., 'age')

Faloutsos CMU SCS 15-415 52

CMU SCS skip

Tree pruning

Shortcut for scalability: DYNAMIC pruning:

- stop expanding the tree, if a node is 'reasonably' homogeneous
 - ad hoc threshold [Agrawal+, vldb92]
 - (Minimum Description Language (MDL) criterion (SLIQ) [Mehta+, edbt96])

Faloutsos CMU SCS 15-415 53

CMU SCS skip

Tree pruning

- Q: How to do it?
- A1: use a 'training' and a 'testing' set - prune nodes that improve classification in the 'testing' set. (Drawbacks?)
- (A2: or, rely on MDL (= Minimum Description Language))

Faloutsos CMU SCS 15-415 54

CMU SCS skip

Outline

- Problem
- Getting the data: Data Warehouses, DataCubes, OLAP
- Supervised learning: decision trees
 - problem
 - approach
 - scalability enhancements
- Unsupervised learning
 - association rules
 - (clustering)

Faloutsos CMU SCS 15-415 55

CMU SCS skip

Scalability enhancements

- Interval Classifier [Agrawal+, vldb92]: dynamic pruning
- SLIQ: dynamic pruning with MDL; vertical partitioning of the file (but label column has to fit in core)
- SPRINT: even more clever partitioning

Faloutsos CMU SCS 15-415 56

CMU SCS

Conclusions for classifiers

- Classification through trees
- Building phase - splitting policies
- Pruning phase (to avoid over-fitting)
- For scalability:
 - dynamic pruning
 - clever data partitioning

Faloutsos CMU SCS 15-415 57

CMU SCS

Outline

- Problem
- Getting the data: Data Warehouses, DataCubes, OLAP
- Supervised learning: decision trees
 - problem
 - approach
 - scalability enhancements
- Unsupervised learning
 - association rules
 - (clustering)

Faloutsos CMU SCS 15-415 58

CMU SCS

Association rules - idea

[Agrawal+SIGMOD93]

- Consider 'market basket' case:
 - (milk, bread)
 - (milk)
 - (milk, chocolate)
 - (milk, bread)
- Find 'interesting things', eg., rules of the form:
 - milk, bread -> chocolate | 90%

Faloutsos CMU SCS 15-415 59

CMU SCS

Association rules - idea

In general, for a given rule

$$I_j, I_k, \dots I_m \rightarrow I_x | c$$

'c' = 'confidence' (how often people buy I_x , given that they have bought $I_j, \dots I_m$)

's' = support: how often people buy $I_j, \dots I_m, I_x$

Faloutsos CMU SCS 15-415 60

CMU SCS

Association rules - idea

Problem definition:

- given
 - a set of 'market baskets' (=binary matrix, of N rows/baskets and M columns/products)
 - min-support 's' and
 - min-confidence 'c'
- find
 - all the rules with higher support and confidence

Faloutsos CMU SCS 15-415 61

CMU SCS

Association rules - idea

Closely related concept: "large itemset"
 $I_j, I_k, \dots, I_m, I_x$
 is a 'large itemset', if it appears more than 'min-support' times

Observation: once we have a 'large itemset', we can find out the qualifying rules easily (how?)
 Thus, let's focus on how to find 'large itemsets'

Faloutsos CMU SCS 15-415 62

CMU SCS

Association rules - idea

Naive solution: scan database once; keep $2^{*|I|}$ counters

Drawback?
 Improvement?

Faloutsos CMU SCS 15-415 63

CMU SCS

Association rules - idea

Naive solution: scan database once; keep $2^{*|I|}$ counters

Drawback? 2^{*1000} is prohibitive...

Improvement? scan the db $|I|$ times, looking for 1-, 2-, etc itemsets

Eg., for $|I|=3$ items only (A, B, C), we have

Faloutsos CMU SCS 15-415 64

CMU SCS

Association rules - idea

(A)

100

(B)

200

(C)

2

first pass

min-sup: 10

Faloutsos CMU SCS 15-415 65

CMU SCS

Association rules - idea

(A,B)

~~(A,C)~~

~~(B,C)~~

(A)

100

(B)

200

~~(C)~~

2

first pass

min-sup: 10

Faloutsos CMU SCS 15-415 66

CMU SCS

Association rules - idea

Anti-monotonicity property:
if an itemset fails to be 'large', so will every superset of it (hence all supersets can be pruned)

Sketch of the (famous!) 'a-priori' algorithm
Let $L(i-1)$ be the set of large itemsets with $i-1$ elements
Let $C(i)$ be the set of candidate itemsets (of size i)

Faloutsos CMU SCS 15-415 67

CMU SCS

Association rules - idea

Compute $L(1)$, by scanning the database.
repeat, for $i=2,3,\dots$,
 'join' $L(i-1)$ with itself, to generate $C(i)$
 two itemset can be joined, if they agree on their first $i-2$ elements
 prune the itemsets of $C(i)$ (how?)
 scan the db, finding the counts of the $C(i)$ itemsets - set this to be $L(i)$
 unless $L(i)$ is empty, repeat the loop
(see example 6.1 in [Han+Kamber])

Faloutsos CMU SCS 15-415 68

CMU SCS

Association rules - Conclusions

Association rules: a new tool to find patterns

- easy to understand its output
- fine-tuned algorithms exist
- still an active area of research

Faloutsos CMU SCS 15-415 69

CMU SCS

Overall Conclusions

- Data Mining: of high commercial interest
- DM = DB+ ML+ Stat
- Data warehousing / OLAP: to get the data
- Tree classifiers (SLIQ, SPRINT)
- Association Rules - 'a-priori' algorithm
- (clustering: BIRCH, CURE, OPTICS)

Faloutsos CMU SCS 15-415 70

CMU SCS

Reading material

- Agrawal, R., T. Imielinski, A. Swami, 'Mining Association Rules between Sets of Items in Large Databases', SIGMOD 1993.
- M. Mehta, R. Agrawal and J. Rissanen, 'SLIQ: A Fast Scalable Classifier for Data Mining', Proc. of the Fifth Int'l Conference on Extending Database Technology (EDBT), Avignon, France, March 1996

Faloutsos CMU SCS 15-415 71

CMU SCS

Additional references

- Agrawal, R., S. Ghosh, et al. (Aug. 23-27, 1992). *An Interval Classifier for Database Mining Applications*. VLDB Conf. Proc., Vancouver, BC, Canada.
- Jiawei Han and Micheline Kamber, *Data Mining*, Morgan Kaufman, 2001, chapters 2.2-2.3, 6.1-6.2, 7.3.5

Faloutsos CMU SCS 15-415 72