


ORACLE®



**ENGINEERED
FOR INNOVATION**

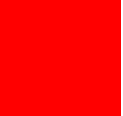
**ORACLE
OPEN
WORLD**

ORACLE[®]

In-Database Analytics: Predictive Analytics, Data Mining, Exadata & Business Intelligence

Charlie Berger
Sr. Director Product Management, Data Mining and Advanced Analytics
Oracle Corporation
charlie.berger@oracle.com
www.twitter.com/CharlieDataMine





The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions.


The development, release, and timing of any features or functionality described for Oracle's products remain at the sole discretion of Oracle.

Agenda


- Big Data Analytics
- BI, Data Mining and Predictive Analytics
- Oracle Data Mining
- Example use cases
- Applications Powered by PA and ODM
- P&G ODM presentation
- Q & A




What Makes it Big Data?




VOLUME



VELOCITY




VARIETY



VALUE

Big Data in Action



Make
Better
Decisions
Using
Big Data

BI & Analytics Spectrum

Queries & Reports

Extraction of detailed and roll up data

“Information”

Who purchased mutual funds in the last 3 years?

OLAP

Summaries, trends and forecasts

“Analysis”

What is the average income of mutual fund buyers, by region, by year?

Data Mining


Knowledge discovery of hidden patterns

“Insight & Prediction”

Who is likely to mutual fund in the next 6 months and why?

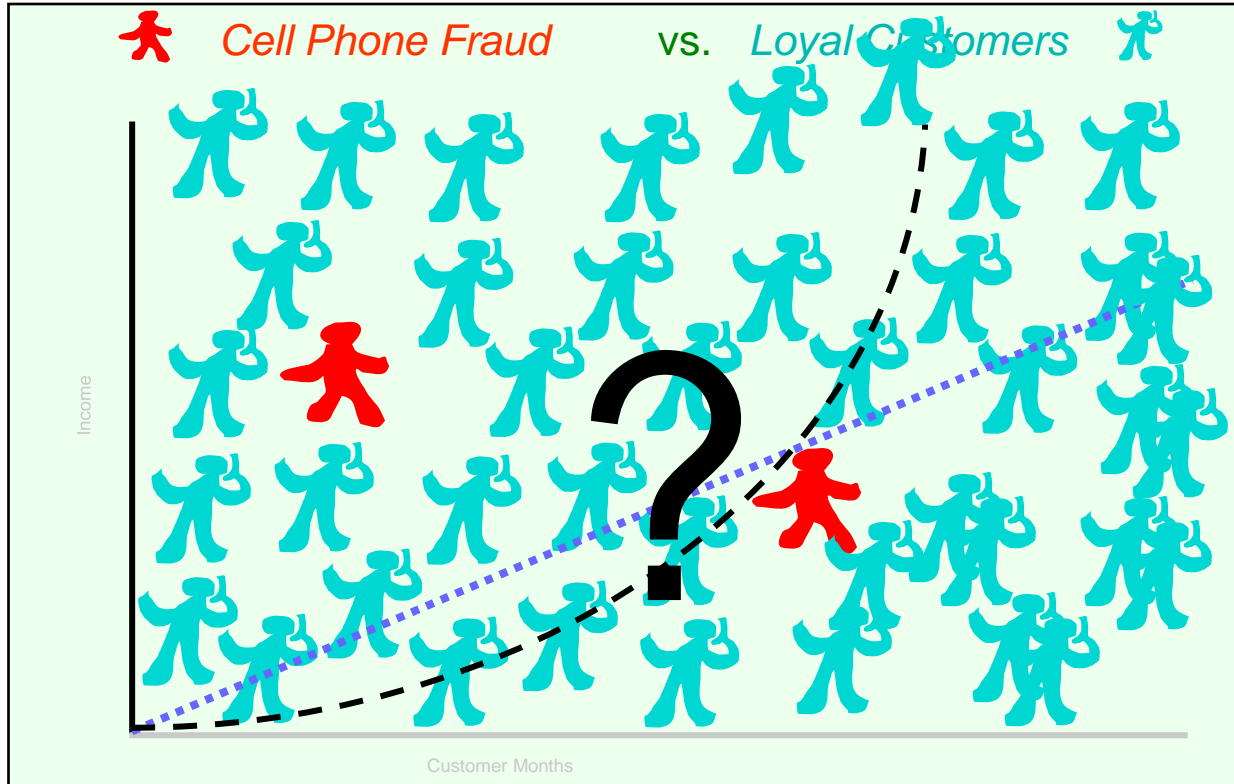
Data Mining Provides

Better Information, Valuable Insights and Predictions



Data Mining Provides

Better Information, Valuable Insights and Predictions



My Personal Experience



Purchases were made in pairs of \$75.00 purchases

Total purchases exceeds
time period average

May 22	1:14 PM	FOOD	Monaco Café	\$127.00
May 22	7:32 PM	WINE	Wine Bistro	\$28.00
...				
<u>June 14</u>	2:05 PM	MISC	Mobil Mart	<u>\$75.00</u>
<u>June 14</u>	2:06 PM	MISC	Mobil Mart	<u>\$75.00</u>
<u>June 15</u>	11:48 AM	MISC	Mobil Mart	<u>\$75.00</u>
<u>June 15</u>	11:49 AM	MISC	Mobil Mart	<u>\$75.00</u>
May 22	7:32	WINE	Wine Bistro	\$28.00
May 22	7:32	WINE	Wine Bistro	\$28.00
<u>June 16</u>	11:48 AM	MISC	Mobil Mart	<u>\$75.00</u>
<u>June 16</u>	11:49 AM	MISC	Mobil Mart	<u>\$75.00</u>

Gas Station?

France?

Pairs
of \$75?

All same \$75 amount?

Finding Needles in Haystacks

- Haystacks are usually **BIG**
- Needles are typically **small** and **rare**



Look for What is “*Different*”




Oracle Data Mining

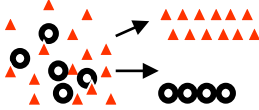

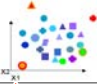
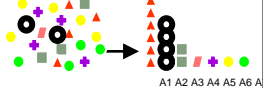
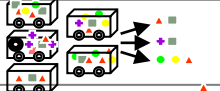
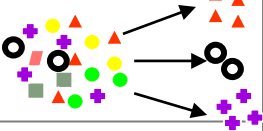
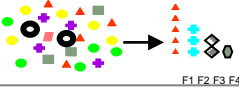
Anomaly Detection

- “One-Class” SVM Models
 - Fraud, noncompliance
 - Outlier detection
 - Network intrusion detection
 - Disease outbreaks
 - Rare events, true novelty

Problem: Detect rare cases



Oracle Data Mining Algorithms

Problem	Algorithm	Applicability
Classification 	Logistic Regression (GLM) Decision Trees Naïve Bayes Support Vector Machine	Classical statistical technique Popular / Rules / transparency Embedded app Wide / narrow data / text
Regression 	Multiple Regression (GLM) Support Vector Machine	Classical statistical technique Wide / narrow data / text
Anomaly Detection 	One Class SVM	Lack examples of target field
Attribute Importance 	Minimum Description Length (MDL)	Attribute reduction Identify useful data Reduce data noise
Association Rules 	Apriori	Market basket analysis Link analysis
Clustering 	Hierarchical K-Means Hierarchical O-Cluster	Product grouping Text mining Gene and protein analysis
Feature Extraction 	Nonnegative Matrix Factorization	Text analysis Feature reduction

Typical Data Mining Use Cases

- **Retail**

- Customer segmentation
- Response modeling
- Recommend next likely product
- Profile high value customers

- **Banking**

- Credit scoring
- Probability of default
- Customer profitability
- Customer targeting

- **Insurance**

- Risk factor identification
- Claims fraud
- Policy bundling
- Employee retention

- **Higher Education**

- Alumni donations
- Student acquisition
- Student retention
- At-risk student identification

- **Healthcare**

- Patient procedure recommendation
- Patient outcome prediction
- Fraud detection
- Doctor & nurse note analysis

- **Life Sciences**

- Drug discovery & interaction
- Common factors in (un)healthy patients
- Cancer cell classification
- Drug safety surveillance

- **Telecommunications**

- Customer churn
- Identify cross-sell opportunities
- Network intrusion detection

- **Public Sector**

- Taxation fraud & anomalies
- Crime analysis
- Pattern recognition in military surveillance

- **Manufacturing**

- Root cause analysis of defects
- Warranty analysis
- Reliability analysis
- Yield analysis

- **Automotive**

- Feature bundling for customer segments
- Supplier quality analysis
- Problem diagnosis

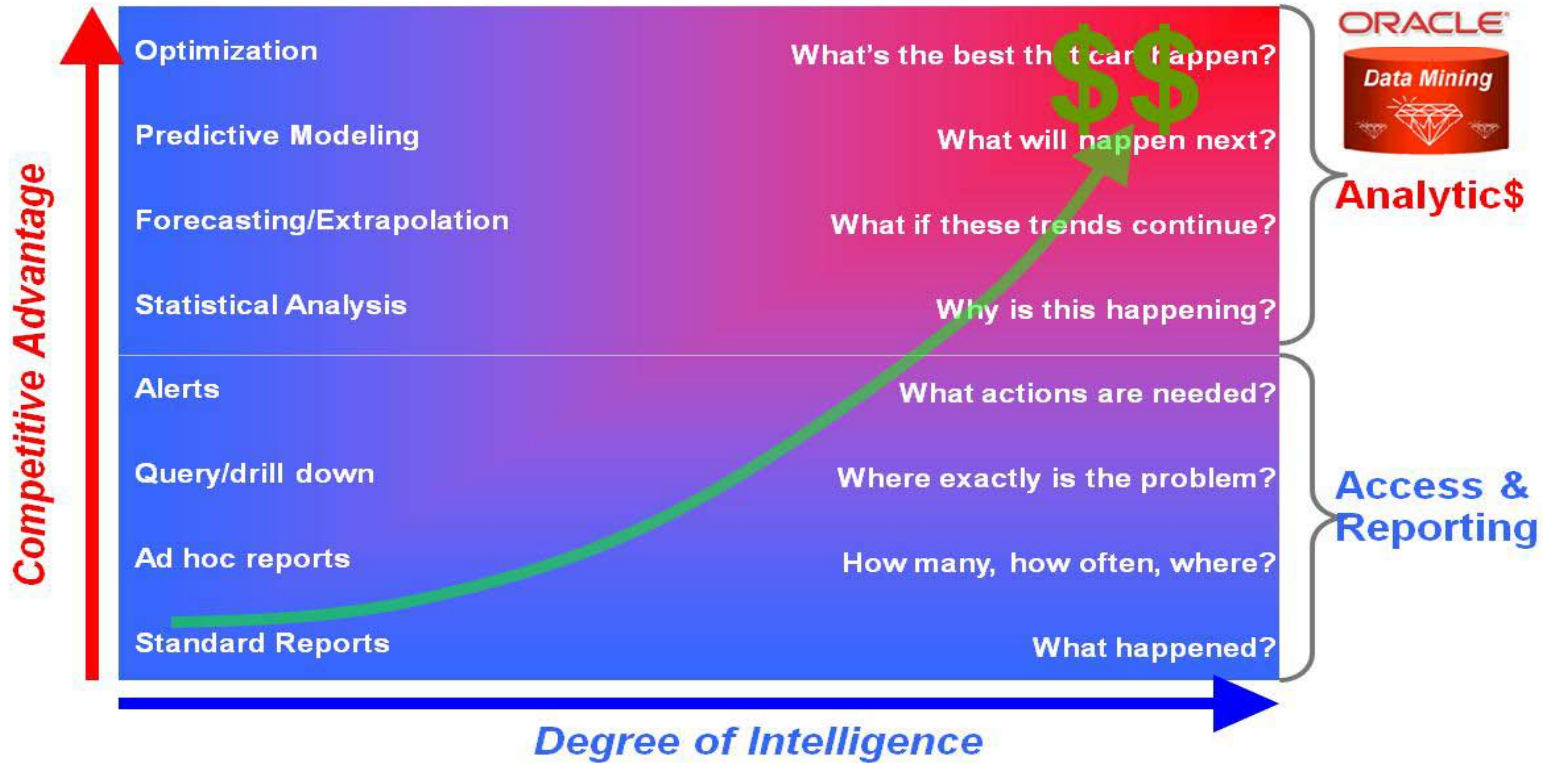
- **Chemical**

- New compound discovery
- Molecule clustering
- Product yield analysis

- **Utilities**

- Predict power line / equipment failure
- Product bundling
- Consumer fraud detection

Competitive Advantage

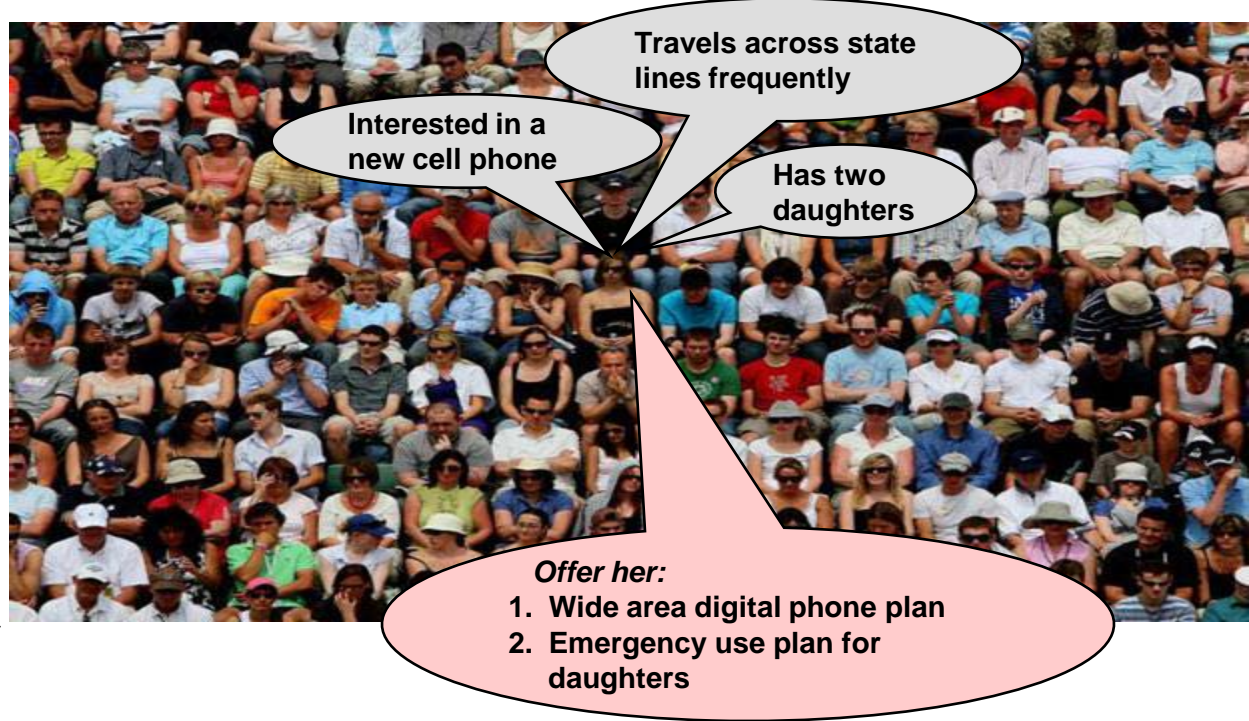


Source: Competing on Analytics, by T. Davenport & J. Harris

ORACLE

Targeting the Right Customers

- 1:1 Relationships
- Understand and predict individual customer behavior
- Offer products and services that anticipate customer needs
- Build loyalty and increase profitability



Oracle—Hardware and Software Engineered to Work Together



**Hardware and Software
Engineered to Work Together**

- Oracle is the world's most complete, open, and integrated business software and hardware systems company

- **Data Warehousing, VLDB and ILM**

- **Oracle Data Mining Option** 

- 12- in-DB data mining algorithms
- In-DB model build
- In-DB model apply
- In-DB text mining
- 50+ in-DB statistical functions

- **Oracle R Enterprise** 

- R for the Enterprise


*Oracle has taught the Database how to do
Advanced Math/Statistics/Data Mining, and more...*

ORACLE

What is Data Mining?

- Automatically sifts through data to find hidden patterns, discover new insights, and make predictions
- Data Mining can provide valuable results:
 - Predict customer behavior (*Classification*)
 - Predict or estimate a value (*Regression*)
 - Segment a population (*Clustering*)
 - Identify factors more associated with a business problem (*Attribute Importance*)
 - Find profiles of targeted people or items (*Decision Trees*)
 - Determine important relationships and “market baskets” within the population (*Associations*)
 - Find fraudulent or “rare events” (*Anomaly Detection*)

ORACLE®



ORACLE®


SQL Developer 3.0/Oracle Data Miner 11g Release 2 GUI

- Graphical User Interface for data analyst
- SQL Developer Extension (OTN download)
- Explore data—discover new insights
- Build and evaluate data mining models
- Apply predictive models
- Share analytical workflows
- Deploy SQL Apply code/scripts

The screenshot displays the Oracle SQL Developer 3.0 GUI with the Oracle Data Miner extension. The main workspace shows a workflow diagram for a customer analytics project. The workflow starts with a data source 'CUST_INSUR_LTV', which is processed through several steps: 'Explore Data 2', 'Column filter and AI', and 'Clust Build 21'. The results are then used to build '5 Response Mpols' and 'Model Details 17', which are applied to 'NEW CUST_INSUR_LTV1' data. A red starburst graphic with the text 'New GUI' is overlaid on the workflow diagram.

The 'Models' pane at the bottom right shows a table of model settings:

Name	Build	Test	Tune	Algorithm	Comment
CLAS_GLM_3_3	7/13/10 6:07...	7/13/10 6:07...	Automatic	Generalized Line...	
CLAS_SVM_3_3	7/13/10 6:06...	7/13/10 6:06...	Automatic	Support Vector ...	
CLAS_SVM_3_3	7/13/10 6:06...	7/13/10 6:07...	Automatic	Support Vector ...	
CLAS_DT_3_3	7/13/10 6:06...	7/13/10 6:06...	Automatic	Decision Tree	
CLAS_NB_3_3	7/13/10 6:06...	7/13/10 6:06...	Automatic	Naive Bayes	




- 12 years “stem celling analytics” into Oracle
 - Designed advanced analytics into database kernel to leverage relational database strengths
 - Naïve Bayes and Association Rules—1st algorithms added
 - Leverages counting, conditional probabilities, and much more
- Now, analytical database platform
 - 12 cutting edge machine learning algorithms and 50+ statistical functions
 - A data mining model is a schema object in the database, built via a PL/SQL API and scored via built-in SQL functions.
 - When building models, leverage existing scalable technology
 - (e.g., parallel execution, bitmap indexes, aggregation techniques) and add new core database technology (e.g., recursion within the parallel infrastructure, IEEE float, etc.)
 - True power of embedding within the database is evident when scoring models using built-in SQL functions (incl. Exadata)

```
select cust_id
from customers
where region = 'US'
and prediction_probability(churnmod, 'Y' using *) > 0.8;
```


In-Database Data Mining



Traditional Analytics



Hours, Days or Weeks



Oracle Data Mining



Secs. Mins or Hours



Savings

Results

- Faster time for "Data" to "Insights"
- Lower TCO—Eliminates
 - Data Movement
 - Data Duplication
- Maintains Security

- Model "Scoring"
Data remains in the Database
- Embedded data preparation
- Cutting edge machine learning algorithms inside the SQL kernel of Database
- SQL—Most powerful language for data preparation and transformation
- Data remains in the Database

You Can Think of It Like This...

Traditional SQL

- “Human-driven” queries
- Domain expertise
- Any “*rules*” must be defined and managed

- SQL Queries

- SELECT
- DISTINCT
- AGGREGATE
- WHERE
- AND OR
- GROUP BY
- ORDER BY
- RANK



Oracle Data Mining

- Automated knowledge discovery, model building and deployment
- Domain expertise to assemble the “right” data to mine


- ODM “Verbs”

- PREDICT
- DETECT
- CLUSTER
- CLASSIFY
- REGRESS
- PROFILE
- IDENTIFY FACTORS
- ASSOCIATE




Oracle Data Miner Nodes (Partial List)

Tables and Views




Transformations




Explore Data



Modeling

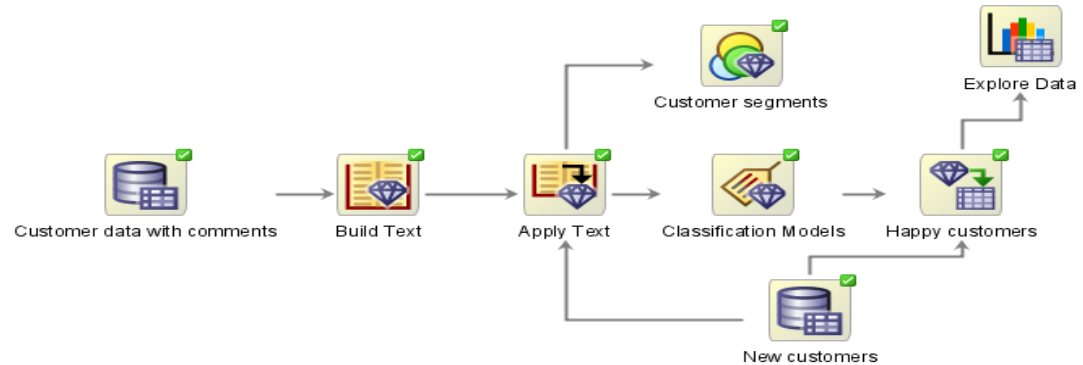


Text



Oracle Data Mining and Unstructured Data

- Oracle Data Mining mines unstructured i.e. “text” data
- Include free text and comments in ODM models
- Cluster and Classify documents
- Oracle Text used to preprocess unstructured text



Output Sample:

CUST_ID
101509
101510
101511
101512
101513
101514

COMMENTS

Shopping at your store is a hassle. I rarely shop there and usually forget to bring your new loyalty card and hence never get the items at the sale price. Can a store manager look up my account on-line?

COMMENTS_TOK


Name	Frequency
NEW	1
PRICE	1
RARELY	1
SALE	1
SHOP	1
SHOPPING	1
STORE	2
UP	1
USUALLY	1



Easier

Churn Demos

SOFTWARE.
HARDWARE.
COMPLETE.



Oracle Data Miner 11g Release 2 GUI

Churn Demo—Simple Conceptual Workflow

The screenshot displays the Oracle SQL Developer interface with the Oracle Data Miner component. The main window shows a conceptual workflow diagram. The workflow starts with a data source 'CHURNERS01'. From this source, three paths emerge: one to 'Explore Data', one to 'Filter Columns', and one to 'Clust Build'. 'Filter Columns' leads to '4 Classification models', which then leads to 'Customer Segments' and 'Likely Churners'. 'Clust Build' also leads to 'Customer Segments'. 'Explore Data' leads to 'Current Customers', which also leads to 'Likely Churners'.

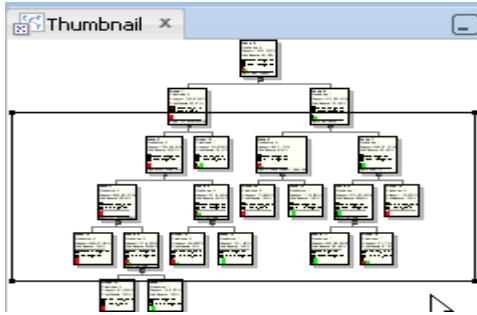
On the right side, the 'Component Palette' is visible, showing various models and transforms. The 'Models' section includes: Anomaly Detection, Association, Classification, Clustering, Feature Extraction, Evaluate and Apply, Apply, Test, Data, Create Table or View, Data Source, Explore Data, and Update Table. The 'Transforms' section includes: Aggregate, Filter Columns, Filter Columns Details, Filter Rows, Join, Sample, Text, Apply Text, Build Text, and Text Reference. The 'Linking Nodes' section includes: Link.

At the bottom, the 'Workflow Jobs' table is visible, showing the status of various workflows. Below that, the '4 Classification models - Property Inspector' window is open, displaying a table of model settings.

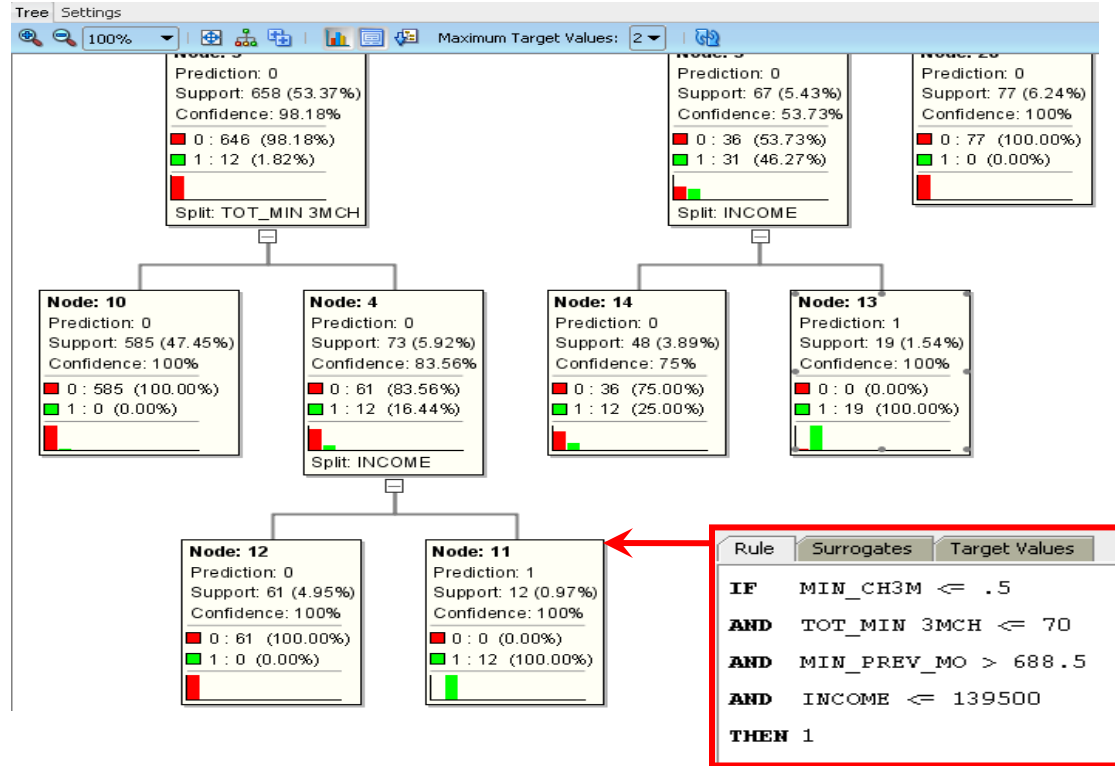
Name	Output	Build	Test	Tune	Algorithm	Comment
CLAS_GLM_1_47	...	✓ 4/26/11 11:30...	✓ 4/26/11 11:31...	Automatic	Generalized Linear Model	
CLAS_SVM_1_47	...	✓ 4/26/11 11:30...	✓ 4/26/11 11:31...	Automatic	Support Vector Machine	
CLAS_DT_1_47	...	✓ 4/26/11 11:31...	✓ 4/26/11 11:31...	Automatic	Decision Tree	
CLAS_NB_1_47	...	✓ 4/26/11 11:30...	✓ 4/26/11 11:31...	Automatic	Naive Bayes	

Oracle Data Miner 11g Release 2 GUI

Churn Demo—Simple Conceptual Workflow

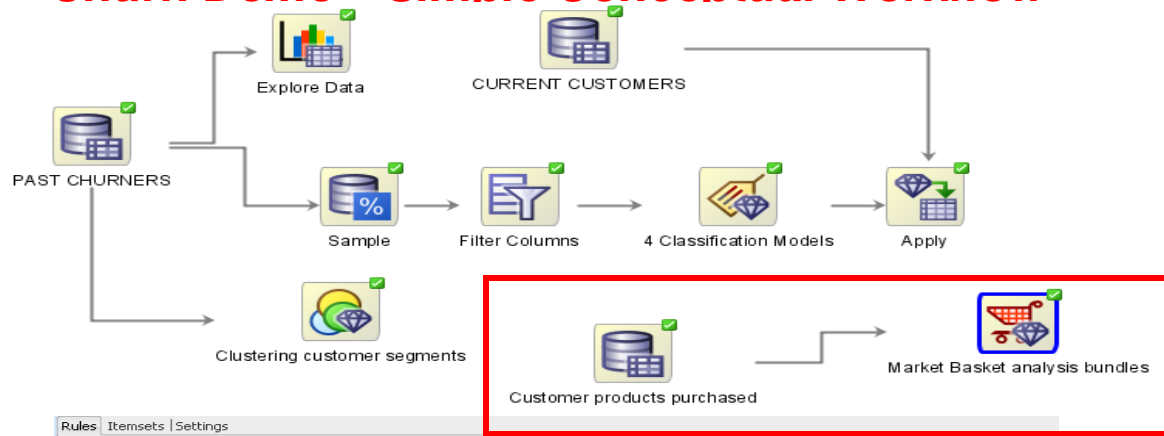


Churn models to product
and “profile” likely
churners



Oracle Data Miner 11g Release 2 GUI

Churn Demo—Simple Conceptual Workflow



Market Basket Analysis to identify potential product bundles

Rules | Itemsets | Settings

Sort by: Lift Descending

Fetch Size: 1,000

Rule Content: Name, Subname

Rules: 1,000 out of 20,988

ID	Antecedent	Consequent	Lift	Confidence(%)	Support(%)
17313	PROD_ID.137 AND PROD_ID.143 AND PROD_ID.138	PROD_ID.144	29.1235	70.8661	1.0689
17849	PROD_ID.138 AND PROD_ID.143 AND PROD_ID.139	PROD_ID.144	28.7389	69.9304	1.0528
17673	PROD_ID.137 AND PROD_ID.143 AND PROD_ID.142	PROD_ID.144	28.3021	68.8675	1.1003
20557	PROD_ID.139 AND PROD_ID.143 AND PROD_ID.142	PROD_ID.144	27.8527	67.7739	1.0975
17314	PROD_ID.137 AND PROD_ID.144 AND PROD_ID.138	PROD_ID.143	27.8265	77.5469	1.0689
18071	PROD_ID.138 AND PROD_ID.146 AND PROD_ID.144	PROD_ID.143	27.5925	76.8947	1.0207
17942	PROD_ID.138 AND PROD_ID.144 AND PROD_ID.140	PROD_ID.143	27.5082	76.6599	1.0486
18075	PROD_ID.138 AND PROD_ID.148 AND PROD_ID.144	PROD_ID.143	27.4991	76.6346	1.0563


Rule Details:

ID: 17313
IF
 PROD_ID.137 AND
 PROD_ID.143 AND
 PROD_ID.138
THEN
 PROD_ID.144

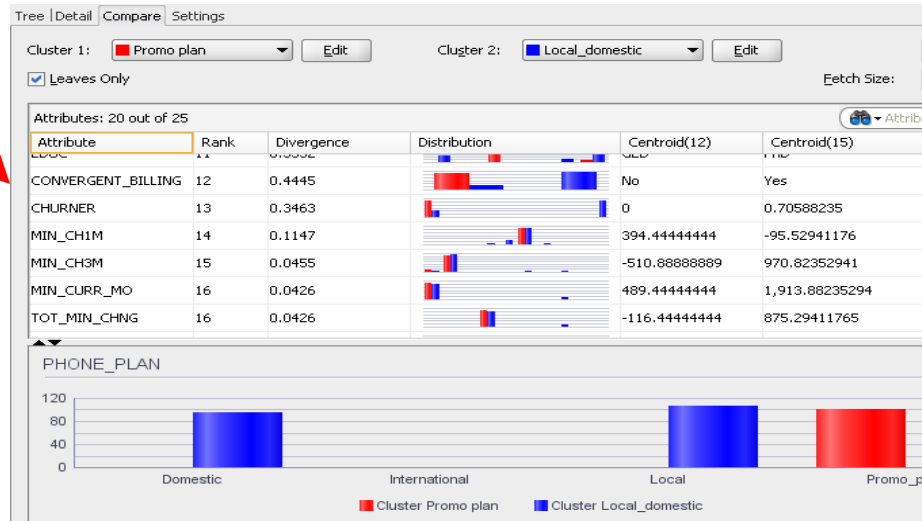
Lift	29.1235
Confidence(%)	70.8661
Support(%)	1.0689
Antecedent Support(%)	2.4333
Consequent Support(%)	1.5083
Item Count	3

Oracle Data Miner 11g Release 2 GUI

Churn Demo—Simple Conceptual Workflow



Clustering analysis to discover customer segments based on behavior, demographics, plans, equipment, etc.



Fraud Prediction Demo

```
drop table CLAIMS_SET;  
exec dbms_data_mining.drop_model('CLAIMSMODEL');  
create table CLAIMS_SET (setting_name varchar2(30), setting_value varchar2(4000));  
insert into CLAIMS_SET values  
    ('ALGO_NAME','ALGO_SUPPORT_VECTOR_MACHINES');  
insert into CLAIMS_SET values ('PREP_AUTO','ON');  
commit;
```

```
begin  
dbms_data_mining.create_model('CLAIMSMODEL', 'CLASSIFICATION',  
'CLAIMS2', 'POLICYNUMBER', null, 'CLAIMS_SET');  
end;  
/
```

```
-- Top 5 most suspicious fraud policy holder claims  
select * from  
(select POLICYNUMBER, round(prob_fraud*100,2) percent_fraud,  
    rank() over (order by prob_fraud desc) rnk from  
(select POLICYNUMBER, prediction_probability(CLAIMSMODEL, '0' using *) prob_fraud  
from CLAIMS2  
where PASTNUMBEROFCLAIMS in ('2 to 4', 'more than 4'))  
where rnk <= 5  
order by percent_fraud desc;
```

POLICYNUMBER	PERCENT_FRAUD	RNK
6532	64.78	1
2749	64.17	2
3440	63.22	3
654	63.1	4
12650	62.36	5

Automated Monthly "Application"! *Just add:*

```
Create  
View CLAIMS2_30  
As  
Select * from CLAIMS2  
Where mydate > SYSDATE - 30
```



ORACLE

Exadata + Data Mining 11g Release 2

“DM Scoring” Pushed to Storage!



- In 11g Release 2, SQL predicates and Oracle Data Mining models are pushed to storage level for execution

For example, find the US customers likely to churn:

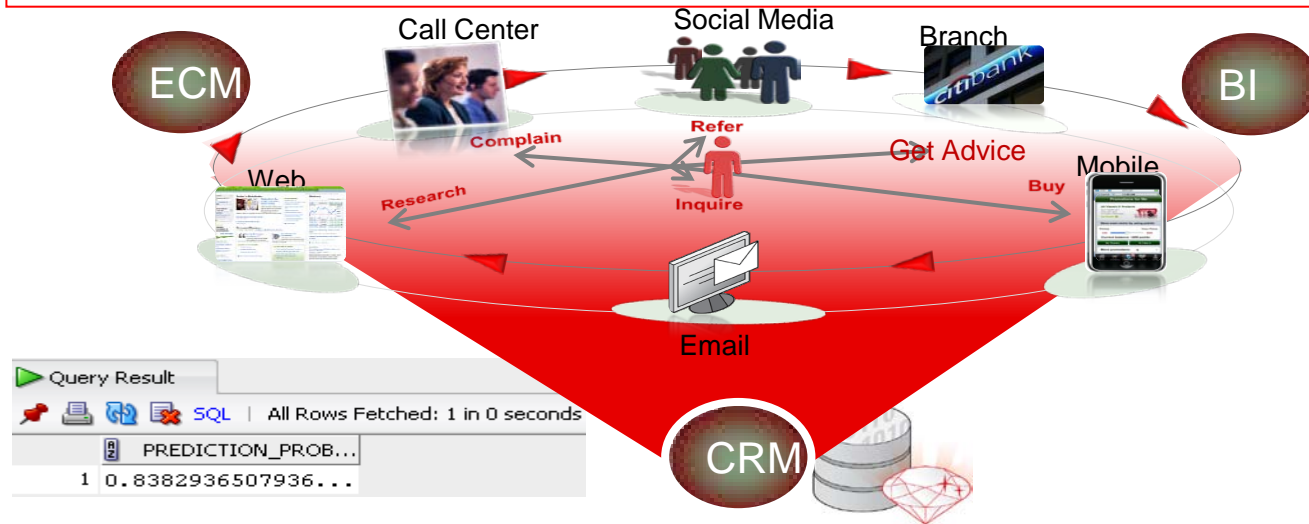
```
select cust_id
from customers
where region = 'US'
and prediction_probability(churnmod, 'Y' using *) > 0.8;
```

Scoring function executed in Exadata

Real-time Prediction for a Customer


- On-the-fly, single record apply with new data (e.g. from call center)

```
Select prediction_probability(CLAS_DT_5_2, 'Yes'  
  USING 7800 as bank_funds, 125 as checking_amount, 20 as  
  credit_balance, 55 as age, 'Married' as marital_status,  
  250 as MONEY_MONTHLY_OVERDRAWN, 1 as house_ownership)  
from dual;
```




Ability to Import/Export 3rd Party DM Models

- ODM 11g Release 2 adds ability to import 3rd party models (PMML), convert to native ODM models and score them in-DB
 - Supported models for ODM model export:
 - Decision Trees (PMML)
 - Supported algorithms for ODM model import:
 - Multiple regression models (PMML)
 - Logistic regression models (PMML)




• Benefits

- SAS, SPSS, R, etc. data mining models can scored on Exadata
 - Imported dm models become native ODM models and inherit all ODM benefits including scoring at Exadata storage layer, 1st class objects, security, etc.



11g Statistics & SQL Analytics (Free)



- Ranking functions
 - rank, dense_rank, cume_dist, percent_rank, ntile
- Window Aggregate functions (moving and cumulative)
 - Avg, sum, min, max, count, variance, stddev, first_value, last_value
- LAG/LEAD functions
 - Direct inter-row reference using offsets
- Reporting Aggregate functions
 - Sum, avg, min, max, variance, stddev, count, ratio_to_report
- Statistical Aggregates
 - Correlation, linear regression family, covariance
- Linear regression
 - Fitting of an ordinary-least-squares regression line to a set of number pairs.
 - Frequently combined with the COVAR_POP, COVAR_SAMP, and CORR functions

Descriptive Statistics

- DBMS_STAT_FUNCS: summarizes numerical columns of a table and returns count, min, max, range, mean, median, stats_mode, variance, standard deviation, quantile values, +/- n sigma values, top/bottom 5 values
- Correlations
 - Pearson's correlation coefficients, Spearman's and Kendall's (both nonparametric).
- Cross Tabs
 - Enhanced with % statistics: chi squared, phi coefficient, Cramer's V, contingency coefficient, Cohen's kappa
- Hypothesis Testing
 - Student t-test, F-test, Binomial test, Wilcoxon Signed Ranks test, Chi-square, Mann Whitney test, Kolmogorov-Smirnov test, One-way ANOVA
- Distribution Fitting
 - Kolmogorov-Smirnov Test, Anderson-Darling Test, Chi-Squared Test, Normal, Uniform, Weibull, Exponential

Note: Statistics and SQL Analytics are included in Oracle Database Standard Edition

ORACLE

Oracle Communications Industry Data Model Example

Better Information for OBIEE Dashboards

ODM's predictions & probabilities are available in the Database for reporting using Oracle BI EE and other tools

ORACLE Business Intelligence

Churn Report By customer Segment

Customer Segments Customer Segmentation Details

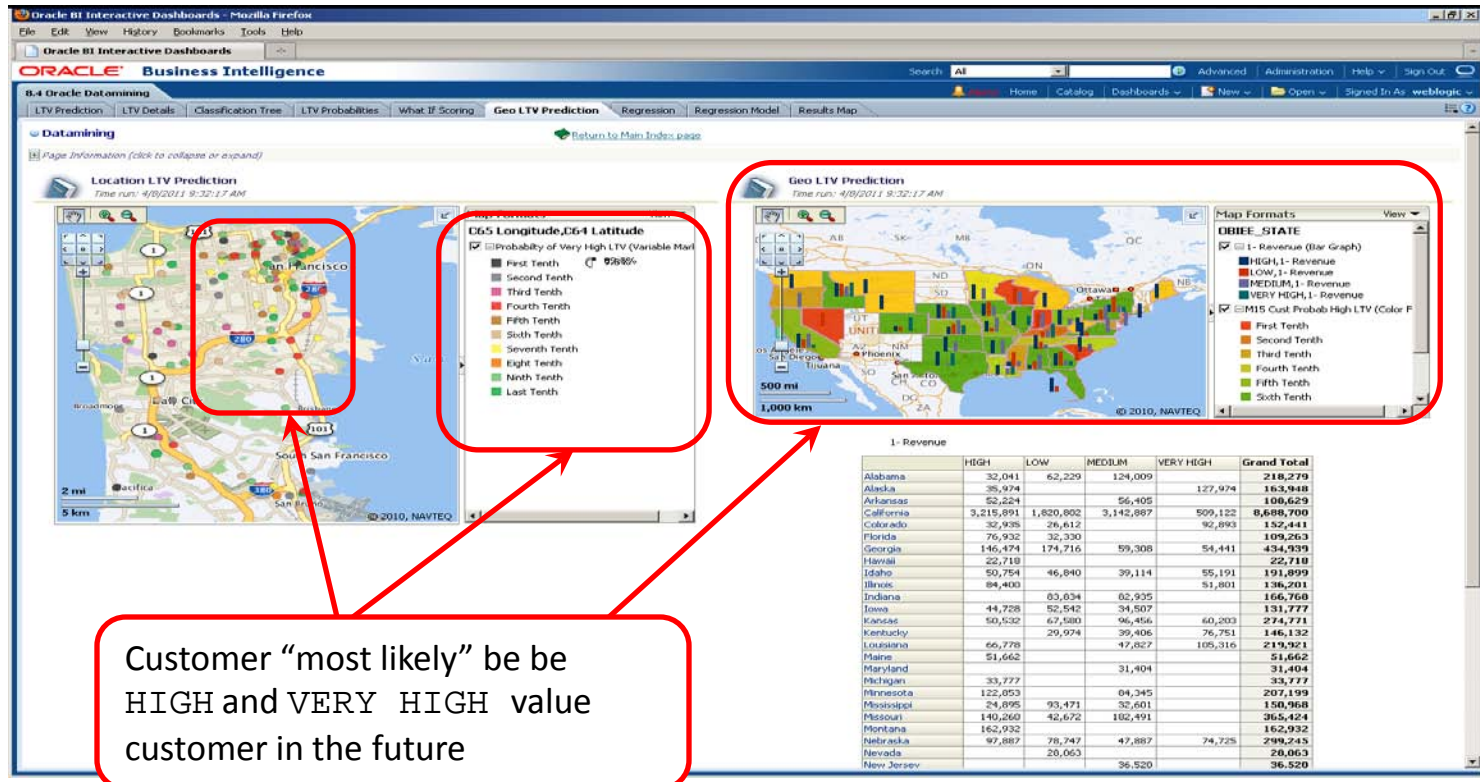
Customer Segment Name is equal to Age Young and PAY TV

Customer Segment	Customer Name	Cell Phone No	Contract Value	Month Revenue	Debt Value	LTV Band	LTV Value	LTV Months	ARPU Band	Churn Indicator	Sentiment	Churn Probability	Customer Segment Key	
Age Young and PAY TV user	Beverly Wan	9985007046	\$9,000.00	\$7,800.00	\$70.00	LTV_1	\$41,000.00	44	ARPU7500+		▲ +	59	104	
	Bradley Johnson	9985007589	\$9,000.00	\$8,100.00	\$222.00		\$49,000.00	32			▲ +	45	104	
	Ethan Nielley	9985006289	\$9,000.00	\$8,400.00	\$70.00		\$34,000.00	43		● Probability of Churning is very high	▼ -	71	104	
	Tobias Hamrick	9985008239	\$0.00	\$7,800.00	\$130.00		\$69,000.00	9			▲ +	43	104	
	Gale Lazar	9985003794	\$9,000.00	\$7,000.00	\$70.00		\$82,000.00	37	ARPU7500		▲ +	16	104	
	Mallory Lawson	9985008346	\$0.00	\$5,400.00	\$130.00		\$63,000.00	30			▲ +	57	104	
	Abbie Anderson	9985010557		\$0.00	\$2,769.23	\$222.00		\$90,000.00	14	ARPU5000	● Probability of Churning is very high	▼ -	79	104
				\$0.00	\$2,769.23	\$222.00		\$99,000.00	18		● Probability of Churning is very high	▼ -	85	104

Exadata with Analytics and Business Intelligence

Better Together

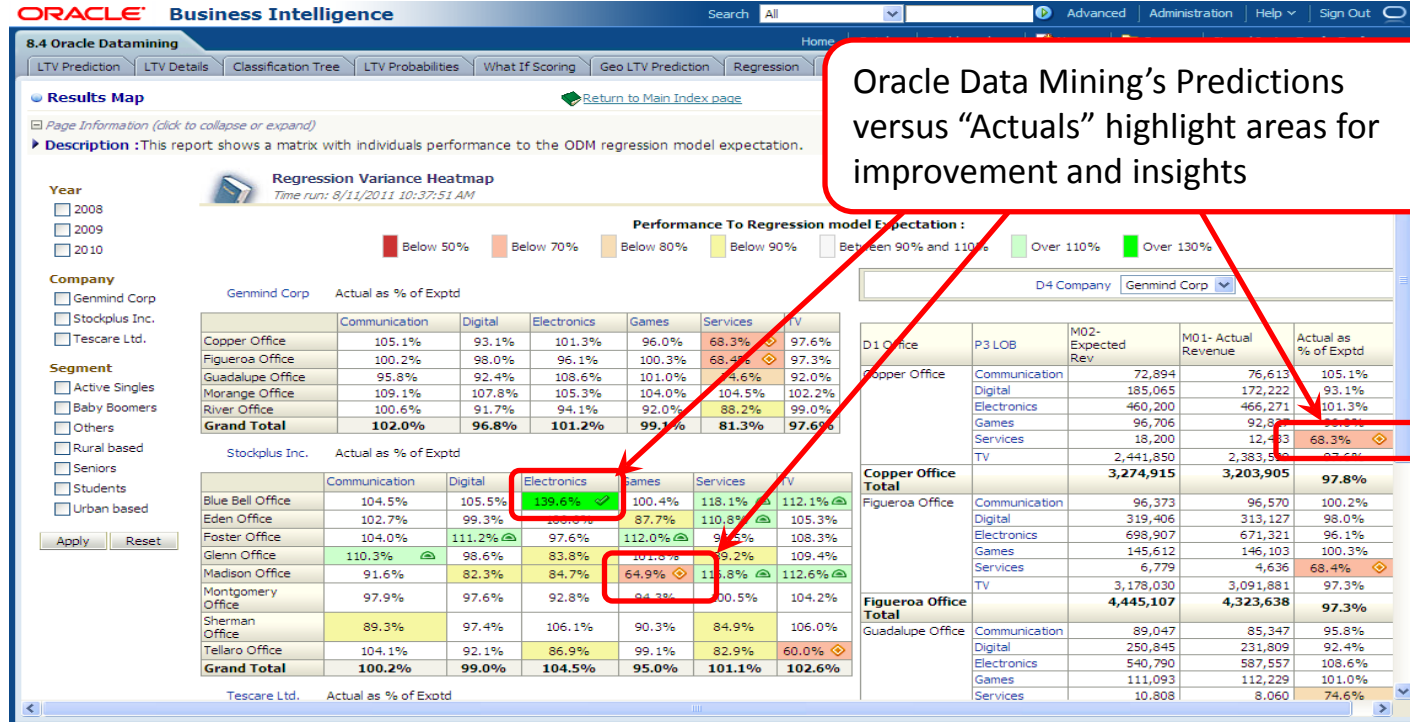
- In-database data mining builds predictive models that predict customer behavior
- OBIEE's integrated spatial mapping shows where



Exadata with Analytics and Business Intelligence

Better Together

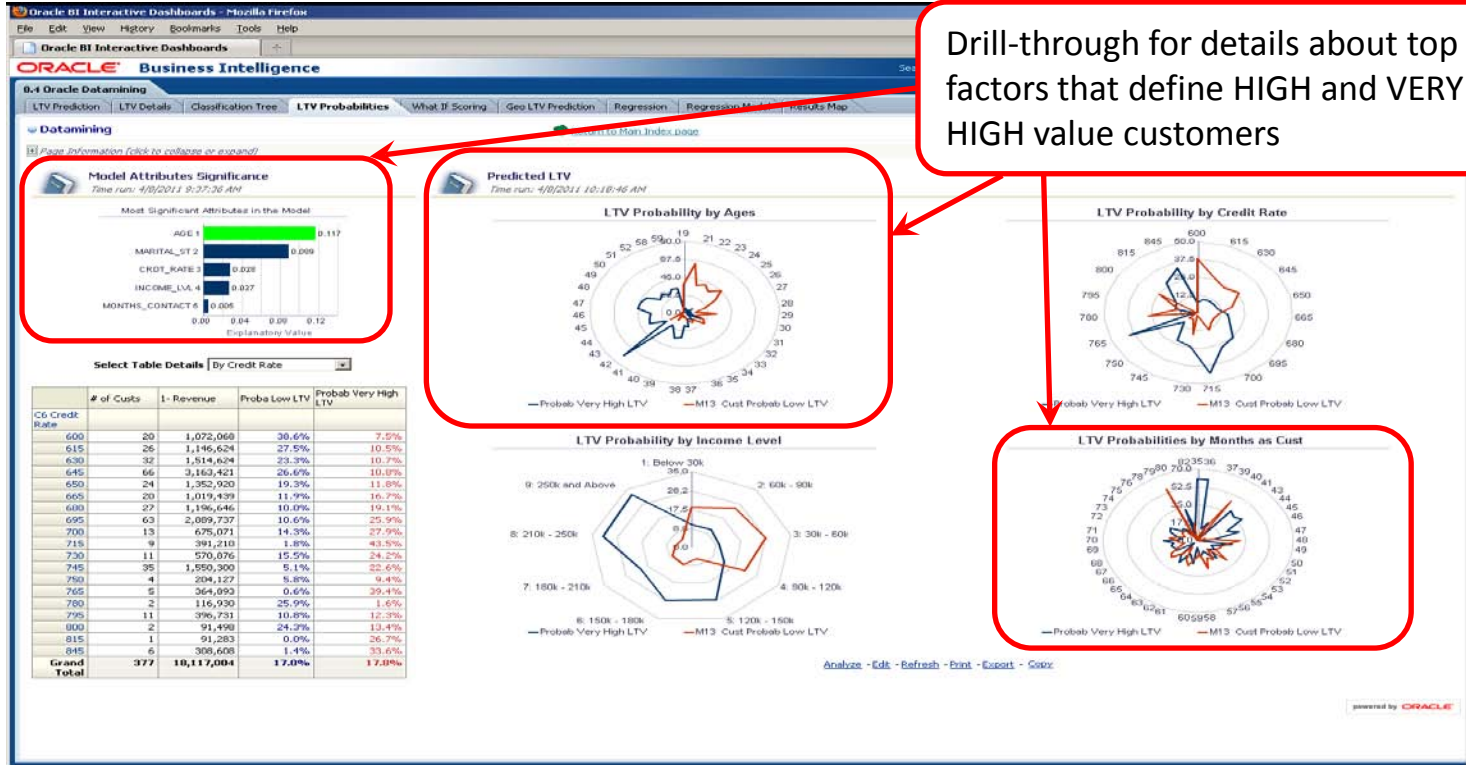
- Deliver advanced in-database analytics through OBIEE
- Ability to drill-through for detail
- Harness the power of Exadata for “Better BI & analytics”



Exadata with Analytics and Business Intelligence

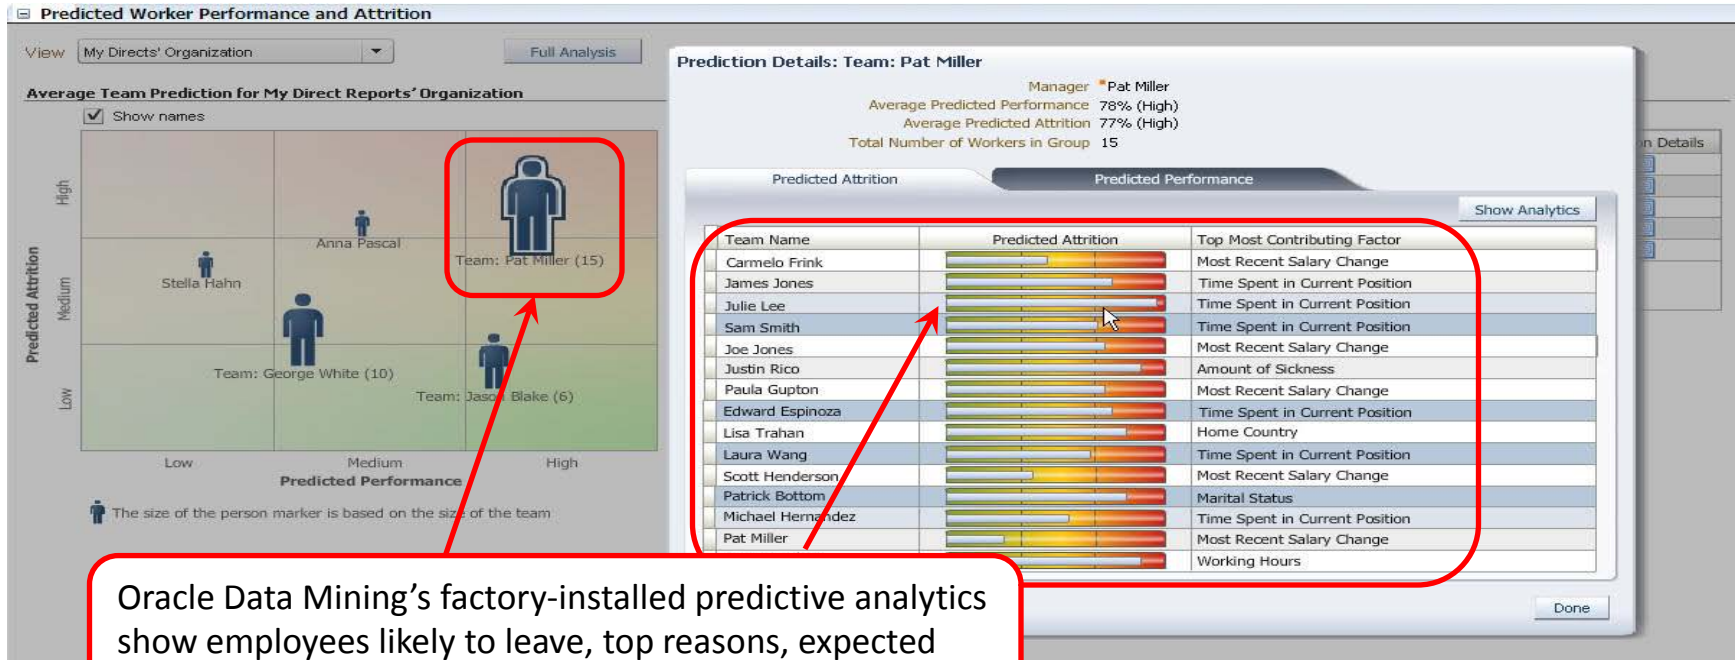
Better Together

- Exadata power
- OBIEE ease-of-use



Fusion HCM Predictive Analytics

Factory Installed PA/ODM Methodologies



Oracle Data Mining's factory-installed predictive analytics show employees likely to leave, top reasons, expected performance and real-time "What if?" analysis

Learn More



Oracle Data Mining 11g Release 2 OBE Series

The OBEs in this series provide you with instructions on how to use Oracle Data Mining 11g Release 2

[Provide Feedback](#)

[Bookmark this OBE](#)

Tags

[Application Development](#) | [Data Mining](#) | [SQLDEV](#)



Setting Up Oracle Data Miner 11g Release 2

This tutorial covers the process of setting up Oracle Data Miner 11g Release 2 for use within Oracle SQL Developer 3.0.

[OBE Details](#)



Release Date 11-MAR-2011 15 mins ★★★★★

Duration

Average Rating



Using Oracle Data Miner 11g Release 2

This tutorial covers the use of Oracle Data Miner to perform data mining against Oracle Database 11g Release 2. In this lesson, you examine and solve a data mining business problem by using the Oracle Data Miner graphical user interface (GUI). The Oracle Data Miner GUI is included as an extension of Oracle SQL Developer, version 3.0.

[OBE Details](#)



Release Date 11-MAR-2011 30 mins ★★★★★

Duration

Average Rating



Star Schema Mining Using Oracle Data Miner

This tutorial covers the use of Oracle Data Miner to perform star schema mining against Oracle Database 11g Release 2.

[OBE Details](#)



Release Date 11-MAR-2011 30 mins ★★★★★

Duration

Average Rating



Text Mining Using Oracle Data Miner

This tutorial covers the use of Oracle Data Miner to perform text mining against Oracle Database 11g Release 2.

[OBE Details](#)



Release Date 11-MAR-2011 30 mins ★★★★★

Duration

Average Rating


Oracle Data Mining PL/SQL Sample Programs

- The PL/SQL Sample Programs provide examples of mini-solutions and use cases for Oracle Data Mining
 - Excellent starting point when developing an ODM Application

<u>Mining Function</u>	<u>Algorithm</u>	<u>Sample Program</u>
Anomaly Detection	One-Class Support Vector Machine	dmsvodem.sql
Association Rules	Apriori	dmardemo.sql
Attribute Importance	Minimum Descriptor Length	dmaidemo.sql
Classification	Decision Tree	dmdtdemo.sql
Classification	Decision Tree (cross validation)	dmdtxvlddemo.sql
Classification	Logistic Regression	dmqlcdem.sql
Classification	Naive Bayes	dmnbdemo.sql
Classification	Support Vector Machine	dmsvcdem.sql
Clustering	k-Means	dmkmdemo.sql
Clustering	O-Cluster	dmocdemo.sql
Feature Extraction	Non-Negative Matrix Factorization	dmnmdemo.sql
Regression	Linear Regression	dmqlrdem.sql
Regression	Support Vector Machine	dmsvrдем.sql
Text Mining	Text transformation using Oracle Text	dmtxtfe.sql
Text Mining	Non-Negative Matrix Factorization	dmtxtnmf.sql
Text Mining	Support Vector Machine (Classification)	dmtxtsvm.sql

What is **ORACLE** **R** Enterprise?


Open Source



- Oracle R Enterprise brings R's statistical functionality closer to the Oracle Database
 1. Eliminate R's memory constraint by enabling R to work directly & transparently on database objects
 - Allows R to run on very large data sets
 2. Architected for Enterprise production infrastructure
 - Automatically exploits database parallelism without require parallel R programming
 - Build and immediately deploy
 3. Oracle R leverages the latest R algorithms and packages
 - R is an embedded component of the DBMS server

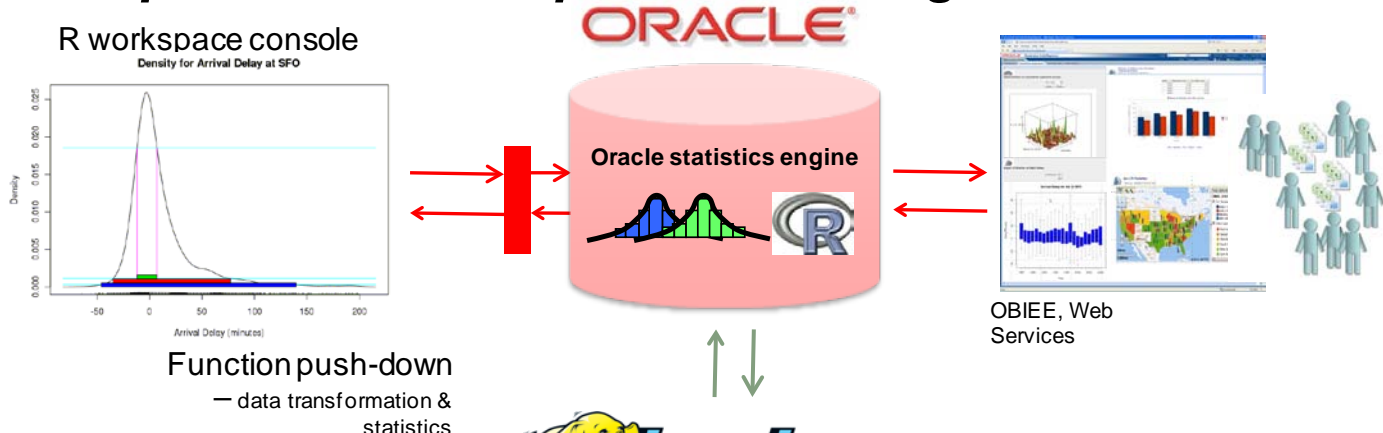
Architecture and Performance

- Transparently function-ships R constructs to database via R → SQL translation
 - Data structures
 - Functions
 - Data manipulation functions (select, project, join)
 - Basic statistical functions (avg, sum, summary)
 - Advanced statistical functions(gamma, beta)
- Performs data-heavy computations in database
 - R for summary analysis and graphics
- Transparent implementation enables using wide range of R “packages” from open source community




Big Data Appliance + R

For Compute Intensive Operations Using R





Function push-down
— data transformation & statistics

```
logreg <- function(input, iterations, dims, alpha){
  plane = rep(0, dims)
  g = function(z) 1/(1 + exp(-z))
  for (i in 1:iterations) {
    z = hdfs.get(hadoop.run( input,
      export = c(plane, g),
      map = logisticRegressionMapper,
      reduce = logisticRegressionReducer))
    gradient = c(z$val[1], z$val[2])
    plane = plane + alpha * gradient
  }
  plane
}
x = hdfs.push(WEBSESSIONS)
logreg(x, 10, 2, 0.05)
```




Oracle In-Database Advanced Analytics

Comprehensive Advanced Analytics Platform



Oracle R Enterprise

- Popular open source statistical programming language & environment
- Integrated with database for scalability
- Wide range of statistical and advanced analytical functions
- R embedded in enterprise apps & OBIEE
- Exploratory data analysis
- Extensive graphics
- Open source R (CRAN) packages
- Integrated with Hadoop for HPC



Oracle Data Mining

- Automated knowledge discovery inside the Database
- 12 in-database data mining algorithms
- Text mining
- Predictive analytics applications development environment
- Star schema and transactional data mining
- Exadata "scoring" of ODM models
- SQL Developer/Oracle Data Miner GUI

Statistics

Advanced Analytics

Data & Text Mining

Predictive Analytics

ORACLE

Together... You Can Think of It Like This...

Traditional SQL

- “Human-driven” queries
- Domain expertise
- Any “*rules*” must be defined and managed
- SQL Queries
 - SELECT
 - DISTINCT
 - AGGREGATE
 - WHERE
 - AND OR
 - GROUP BY
 - ORDER BY
 - RANK



In-Database Analytics

- Wide range of Oracle R Enterprise statistical functions
- Automated ODM knowledge discovery, model building and deployment
- ORE Statistics/Adv. Analytics
 - CORR
 - SUMMARY
 - ARIMA Time Series
- ODM “Verbs”
 - PREDICT
 - DETECT
 - CLUSTER





- **MARK YOUR CALENDARS!**

- **BIWA Summit @**

- **COLLABORATE 12**
April 22-26, 2012

Mandalay Bay Convention Center
Las Vegas, Nevada

<http://events.ioug.org/p/cm/ld/fid=15>



Q&A

Hardware and Software

ORACLE®

Engineered to Work Together

ORACLE®

ORACLE®