

With data mining, the best way to accomplish this is by setting aside some of your data in a vault to isolate it from the mining process; once the mining is complete, the results can be tested against the isolated data to confirm the model's \_\_\_\_\_.

**Validity**

Security

Integrity

None of above

The automated, prospective analyses offered by data mining move beyond the analyses of past events provided by \_\_\_\_\_ tools typical of decision support systems.

Introspective

Intuitive

Reminiscent

**Retrospective**

The technique that is used to perform these feats in data mining is called modeling, and this act of model building is something that people have been doing for a long time, certainly before the \_\_\_\_\_ of computers or data mining technology.

Access

**Advent**

Ascent

Avowal

Classification consists of examining the properties of a newly presented observation and assigning it to a predefined \_\_\_\_\_.

Object

Container

Subject

**Class**

During business hours, most \_\_\_\_\_ systems should probably not use parallel execution.

OLAP

DSS

Data Mining

**OLTP**

In contrast to statistics, data mining is \_\_\_\_\_ driven.

Assumption

**Knowledge**

Human

Database

Data mining derives its name from the similarities between searching for valuable business information in a large database, for example, finding linked products in gigabytes of store scanner data, and mining a mountain for a \_\_\_\_\_ of valuable ore.

Furrow

Streak

Trough

**Vein**

As opposed to the outcome of classification, estimation deal with \_\_\_\_\_ valued outcome.

Discrete

Isolated

**Continuous**

Distinct

The goal of ideal parallel execution is to completely parallelize those parts of a computation that are not constrained by data dependencies. The smaller the portion of the program that must be executed \_\_\_\_\_, the greater the scalability of the computation.

In Parallel

Distributed

**Sequentially**

None of above

Data mining evolve as a mechanism to cater the limitations of \_\_\_\_\_ systems to deal massive data sets with high dimensionality, new data types, multiple heterogeneous data resources etc.

**OLTP**

OLAP

DSS

DWH

The goal of ideal parallel execution is to completely parallelize those parts of a computation that are not constrained by data dependencies. The \_\_\_\_\_ the portion of the program that must be executed sequentially, the greater the scalability of the computation.

Larger

**Smaller**

Unambiguous

Superior

The goal of \_\_\_\_\_ is to look at as few blocks as possible to find the matching records(s).

**Indexing**

Partitioning

Joining

None of above

In nested-loop join case, if there are 'M' rows in outer table and 'N' rows in inner table, time complexity is

$O(M \log N)$

$O(\log MN)$

**$O(MN)$**

$O(M + N)$

Data mining derives its name from the similarities between searching for valuable business information in a large database, for example, finding linked products in gigabytes of store scanner data, and mining a mountain for a \_\_\_\_\_ of valuable ore.

Furrow

Streak

Trough

**Vein**

The technique that is used to perform these feats in data mining is called modeling, and this act of model building is something that people have been doing for a long time, certainly before the \_\_\_\_\_ of computers or data mining technology.

Access

**Advent**

Ascent

Avowal

There are many variants of the traditional nested-loop join. If there is an index and that index is exploited, then it is called

Naive nested-loop join

**Index nested-loop join**

Temporary index nested-loop join

None of these

Data mining is a/an \_\_\_\_\_ approach, where browsing through data using data mining techniques may reveal something that might be of interest to the user as information that was unknown previously

Non-Exploratory

**Exploratory**

Computer Science

None of these

With data mining, the best way to accomplish this is by setting aside some of your data in a \_\_\_\_\_ to isolate it from the mining process; once the mining is complete, the results can be tested against the isolated data to confirm the model's validity.

Cell

Disk

Folder

**Vault**

The automated, prospective analyses offered by data mining move beyond the analyses of past events provided by \_\_\_\_\_ tools typical of decision support systems.

Introspective

Intuitive

Reminiscent

**Retrospective**

There are many variants of the traditional nested-loop join. When the entire table is scanned it is called

Index nested-loop join

**Naive nested-loop join**

Temporary index nested-loop join

None of these

As opposed to the outcome of classification, estimation deal with \_\_\_\_\_ valued outcome.

**Discrete**

Isolated

Continuous

Distinct

Data mining evolve as a mechanism to cater the limitations of \_\_\_\_\_ systems to deal massive data sets with high dimensionality, new data types, multiple heterogeneous data resources etc.

**OLTP**

OLAP

DSS

DWH

With data mining, the best way to accomplish this is by setting aside some of your data in a vault to isolate it from the mining process; once the mining is complete, the results can be tested against the isolated data to confirm the model's \_\_\_\_\_.

**Validity**

Security

Integrity

None of these

In contrast to statistics, data mining is \_\_\_\_\_ driven.

Assumption

**Knowledge**

Human

Database

The goal of \_\_\_\_\_ is to look at as few blocks as possible to find the matching records(s).

**Indexing**

Partitioning

Joining

None

If every key in the data file is represented in the index file then index is

**Dense Index**

Sparse Index

Inverted Index

None

An optimized structure which is built primarily for retrieval, with update being only a secondary consideration is

OLTP

OLAP

DSS

**inverted Index**

Which of the following is not an “Orr’s Law of Data Quality”?

Data that is not used cannot be correct!”

Data quality is a function of its use, not its co

Data will be no better than its most stringent use!”

Data duplication can be harmful for the organization!”

During business hours, most \_\_\_\_\_ systems should probably not use parallel execution.

OLAP

DSS

Data Mining

**OLTP**

\_\_\_\_\_, if fits into memory, costs only one disk I/O access to locate a record by given key.

An Inverted Index

A Sparse Index

**A Dense Index**

None of these

\_\_\_\_\_, if too big and does not fit into memory, will be expensive when used to find a record by given key.

An Inverted Index

A Sparse Index

**A Dense Index**

None of these

The goal of ideal parallel execution is to completely parallelize those parts of a computation that are not constrained by data dependencies. The smaller the portion of the program that must be executed

\_\_\_\_\_, the greater the scalability of the computation.

- In Parallel
- Distributed
- Sequentially**
- None of these

A dense index, if fits into memory, costs only \_\_\_\_\_ disk I/O access to locate a record by given key.

- One**
- Two
- Linear
- Quadratic

Which of the following is not an activity of Data Quality Analysis Project?

- “Define”
- Measure”
- Analyze”
- “Compression”**

The goal of ideal parallel execution is to completely parallelize those parts of a computation that are not constrained by data dependencies. The \_\_\_\_\_ the portion of the program that must be executed in parallel, the greater the scalability of the computation.

- Smaller**
- Larger
- Shorter
- None

The automated, prospective analyses offered by data mining move beyond the analyses of past events provided by retrospective tools typical of \_\_\_\_\_ .

- OLTP
- OLAP
- Decision Support Systems**
- None of these

Data mining, the \_\_\_\_\_ of hidden predictive information from large databases, is a powerful new technology with great potential to help companies focus on the most important information in their data warehouses.

- Extraction**
- Insertion
- Inclusion
- Enclosure

An optimized structure which is built primarily for retrieval, with update being only a secondary consideration is

Select correct option:

- OLTP
- OLAP
- DSS

**Inverted Index**

Question # 4 of 10 ( Start time: 10:33:23 PM ) Total Marks: 1

If every key in the data file is represented in the index file then index is

Select correct option:

- Dense Index**
- Sparse Index

Inverted Index  
None of these

Question # 5 of 10 ( Start time: 10:34:47 PM ) Total Marks: 1

There are many variants of the traditional nested-loop join. If the index is built as part of the query plan and subsequently dropped, it is called

Select correct option:

Naive nested-loop join

Index nested-loop join

**Temporary index nested-loop join**

None of these

Question # 6 of 10 ( Start time: 10:36:08 PM ) Total Marks: 1

Data mining evolve as a mechanism to cater the limitations of \_\_\_\_\_ systems to deal massive data sets with high dimensionality, new data types, multiple heterogeneous data resources etc.

Select correct option:

**OLTP**

OLAP

DSS

DWH

Question # 7 of 10 ( Start time: 10:37:30 PM ) Total Marks: 1

A dense index, if fits into memory, costs only \_\_\_\_\_ disk I/O access to locate a record by given key.

Select correct option:

**One**

Two

Linear

Quadratic

Question # 8 of 10 ( Start time: 10:38:29 PM ) Total Marks: 1

Data mining derives its name from the similarities between searching for valuable business information in a large database, for example, finding linked products in gigabytes of store scanner data, and mining a mountain for a \_\_\_\_\_ of valuable ore.

Select correct option:

Furrow

Streak

Trough

**Vein**

The goal of ideal parallel execution is to completely parallelize those parts of a computation that are not constrained by data dependencies. The \_\_\_\_\_ the portion of the program that must be executed sequentially, the greater the scalability of the computation

Larger

**Smaller**

Unambiguous

Superior

\_\_\_\_\_, if fits into memory, costs only one disk I/O access to locate a record by given key.

An Inverted Index

A Sparse Index

**A Dense Index**

None of these

If someone told you that he had a good model to predict customer usage, the first thing you might try would be to ask him to apply his model to your customer \_\_\_\_\_, where you already knew the answer.

Base

Drive

**File**

Log

The automated, prospective analyses offered by data mining move beyond the analyses of past events provided by \_\_\_\_\_ tools typical of decision support systems.

Introspective

Intuitive

Reminiscent

**Retrospective**

If every key in the data file is represented in the index file then index is

**Dense Index**

Sparse Index

Inverted Index

None of these

A dense index, if fits into memory, costs only \_\_\_\_\_ disk I/O access to locate a record by given key.

**One**

Two

Linear

Quadratic

With data mining, the best way to accomplish this is by setting aside some of your data in a vault to isolate it from the mining process; once the mining is complete, the results can be tested against the isolated data to confirm the model's \_\_\_\_\_.

**Validity**

Security

Integrity

None of these

Data mining uses \_\_\_\_\_ algorithms to discover patterns and regularities in data.

Mathematical

Computational

**Statistical**

None of these

The goal of \_\_\_\_\_ is to look at as few blocks as possible to find the matching records(s).

**Indexing**

Partitioning

Joining

None of these

\_\_\_\_\_, if too big and does not fit into memory, will be expensive when used to find a record by given key.

An Inverted Index

A Sparse Index

**A Dense Index**

None of these

There are many variants of the traditional nested-loop join. If the index is built as part of the query plan and subsequently dropped, it is called

Naive nested-loop join

Index nested-loop join

**Temporary index nested-loop join**

None of these

\_\_\_\_\_, if fits into memory, costs only one disk I/O access to locate a record by given key.

An Inverted Index

A Sparse Index

**A Dense Index**

None of these

The goal of ideal parallel execution is to completely parallelize those parts of a computation that are not constrained by data dependencies. The smaller the portion of the program that must be executed \_\_\_\_\_, the greater the scalability of the computation.

In Parallel

Distributed

**Sequentially**

None of these

Data mining is a/an \_\_\_\_\_ approach, where browsing through data using data mining techniques may reveal something that might be of interest to the user as information that was unknown previously.

Non-Exploratory

**Exploratory**

Compute Science

none of these

Data mining evolve as mechanism to cater the limitations of \_\_\_\_\_ systems to deal massive data sets with high dimensionality , new data types, multiple heterogeneous data resources etc..

**OLTP**

OLAP

DSS

DWH

The goal of ideal parallel execution is to completely parallelize those parts of a computation that are not constrained by data dependencies. The smaller the portion of the program that must be executed \_\_\_\_\_, the greater the scalability of the computation.

None of these

**Sequentially**

In Parallel

Distributed

Data mining is a/an \_\_\_\_\_ approach, where browsing through data using data mining techniques may reveal something that might be of interest to the user as information that was unknown previously.

**Exploratory**

Non-Exploratory

Computer Science

Data mining evolve as a mechanism to cater the limitations of \_\_\_\_\_ systems to deal massive data sets with high dimensionality, new data types, multiple heterogeneous data resources etc.

**OLTP**

OLAP

DSS

DWH

There are many variants of the traditional nested-loop join. If the index is built as part of the query plan and subsequently dropped, it is called

Naive nested-loop join

Index nested-loop join

**Temporary index nested-loop join**

None of these

Data mining derives its name from the similarities between searching for valuable business information in a large database, for example, finding linked products in gigabytes of store scanner data, and mining a mountain for a \_\_\_\_\_ of valuable ore.



Furrow  
Streak  
Trough  
**Vein**

With data mining, the best way to accomplish this is by setting aside some of your data in a \_\_\_\_\_ to isolate it from the mining process; once the mining is complete, the results can be tested against the isolated data to confirm the model's validity.

Cell  
**Disk**  
Folder  
Vault

The automated, prospective analyses offered by data mining move beyond the analysis of past events provided by respective tools typical of \_\_\_\_\_.

OLTP  
OLAP

**Decision Support systems**

None of these

There are many variants of the traditional nested-loop join, if there is an index is exploited, then it is called.....

Naïve nested loop join index  
**Nested loop join temporary index**  
Index nested-loop joins