
**Information technology — Guidance
for the use of database language
SQL —**

**Part 5:
Row pattern recognition**

*Technologies de l'information — Recommandations pour l'utilisation
du langage de base de données SQL —*

Partie 5: Reconnaissance de formes de lignes



STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 19075-5:2021



COPYRIGHT PROTECTED DOCUMENT

© ISO/IEC 2021

All rights reserved. Unless otherwise specified, or required in the context of its implementation, no part of this publication may be reproduced or utilized otherwise in any form or by any means, electronic or mechanical, including photocopying, or posting on the internet or an intranet, without prior written permission. Permission can be requested from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office
CP 401 • Ch. de Blandonnet 8
CH-1214 Vernier, Geneva
Phone: +41 22 749 01 11
Email: copyright@iso.org
Website: www.iso.org

Published in Switzerland

Contents

Page

Foreword.....	vii
Introduction.....	ix
1 Scope.....	1
2 Normative references.....	2
3 Terms and definitions.....	3
4 Row pattern recognition: FROM clause.....	4
4.1 Context of row pattern recognition.....	4
4.2 Introduction to the FROM clause in row pattern recognition.....	4
4.3 Example of ONE ROW PER MATCH.....	4
4.4 Example of ALL ROWS PER MATCH.....	7
4.5 Summary of the syntax.....	9
4.6 The row pattern input table.....	10
4.6.1 Introduction to the row pattern input table.....	10
4.6.2 The row pattern input name.....	11
4.6.3 The row pattern input declared column list.....	12
4.7 MATCH_RECOGNIZE.....	13
4.8 PARTITION BY.....	13
4.9 ORDER BY.....	13
4.10 Row pattern variables.....	13
4.11 MEASURES.....	14
4.12 ONE ROW PER MATCH vs ALL ROWS PER MATCH.....	15
4.12.1 Introduction to use of ROWS PER MATCH.....	15
4.12.2 Handling empty matches.....	15
4.12.3 Handling unmatched rows.....	19
4.13 AFTER MATCH SKIP.....	21
4.14 PATTERN.....	22
4.14.1 Introduction to the PATTERN syntax.....	22
4.14.2 PERMUTE.....	23
4.14.3 Excluding portions of the pattern.....	24
4.15 SUBSET.....	25
4.16 DEFINE.....	26
4.17 The row pattern output table.....	27
4.17.1 Introduction to the row pattern output table.....	27
4.17.2 Row pattern output name.....	28
4.17.3 Row pattern output declared column list.....	28
4.18 Prohibited nesting.....	29
4.18.1 Introduction to prohibited nesting.....	29
4.18.2 Row pattern recognition nested within another row pattern recognition.....	30

4.18.3	Outer references within a row pattern recognition query.	30
4.18.4	Conventional query nested within row pattern recognition query.	31
4.18.5	Recursion.	32
4.18.6	Concatenated row pattern recognition.	32
5	Expressions in MEASURES and DEFINE.	33
5.1	Introduction to the use of expressions in MEASURES and DEFINE.	33
5.2	Row pattern column references.	33
5.3	Running vs. final semantics.	34
5.4	RUNNING vs.FINAL keywords.	38
5.5	Aggregates.	39
5.6	Row pattern navigation operations.	39
5.6.1	The four operations.	39
5.6.2	PREV and NEXT.	39
5.6.3	FIRST and LAST.	41
5.6.4	Nesting FIRST and LAST within PREV or NEXT.	42
5.7	Ordinary row pattern column references reconsidered.	43
5.8	MATCH_NUMBER function.	44
5.9	CLASSIFIER function.	44
6	Row pattern recognition: WINDOW clause.	48
6.1	Introduction to the WINDOW clause.	48
6.2	Example of row pattern recognition in a window.	48
6.3	Summary of the syntax.	50
6.3.1	Syntax components.	50
6.3.2	Syntactic comparison to windows without row pattern recognition.	51
6.3.3	Syntactic comparison to MATCH_RECOGNIZE.	52
6.4	Row pattern input table.	52
6.5	Row pattern variables and other range variables.	52
6.6	Windows defined on windows.	54
6.7	PARTITION BY.	55
6.8	ORDER BY.	55
6.9	MEASURES.	55
6.10	Full window frame and reduced window frame.	55
6.10.1	Introduction to window framing.	55
6.10.2	ROWS BETWEEN CURRENT ROW AND.	56
6.10.3	EXCLUDE NO OTHERS.	56
6.11	AFTER MATCH SKIP.	56
6.12	INITIAL vs SEEK.	57
6.13	PATTERN.	57
6.14	SUBSET.	57
6.15	DEFINE.	57
6.16	Empty matches and empty reduced window frames.	57
6.17	Prohibited nesting.	59
6.17.1	Restrictions on nesting.	59
6.17.2	Row pattern recognition nested within another row pattern recognition.	60
6.17.3	Outer references within a row pattern recognition query.	60
6.17.4	Conventional query nested within row pattern recognition query.	61
6.17.5	Recursion.	61

6.17.6	Concatenated row pattern recognition.	61
7	Pattern matching rules.	63
7.1	Regular expression engines.	63
7.2	Parenthesized language and preferment.	64
7.2.1	Introduction to parenthesized language and preferment.	64
7.2.2	Alternation.	65
7.2.3	Concatenation.	65
7.2.4	Quantification.	66
7.2.5	Exclusion.	67
7.2.6	Anchors.	68
7.2.7	The empty pattern.	68
7.2.8	Infinite repetitions of empty matches.	68
7.3	Pattern matching in theory and practice.	70
Index.		73

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 19075-5:2021

Tables

Table	Page
1 Sample data.	7
2 Results of ONE ROW PER MATCH.	7
3 Results of ALL ROWS PER MATCH.	8
4 Row pattern recognition syntax summary.	9
5 Analysis of sample data permitting empty matches.	16
6 Result of query permitting empty matches.	16
7 Results of query using SHOW EMPTY ROWS.	18
8 Results of query using OMIT EMPTY ROWS.	18
9 Results of ALL ROWS PER MATCH.	20
10 Original and renamed column names.	29
11 Ordered row pattern partition of data.	35
12 RUNNING and FINAL in MEASURES.	36
13 Ordered row pattern partition of data.	37
14 Ordered row pattern partition of data.	37
15 Example data set and mappings for FIRST and LAST.	41
16 Data set and mappings for nesting example.	43
17 Window example query results.	50
18 Row pattern recognition in windows — syntax summary.	51
19 Results for empty match and no match.	58
20 Computation of matches and window function results.	59
21 Input data.	71
22 Mapping of first element.	71
23 Mapping of second element.	72
24 Mapping of third element.	72

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular, the different approval criteria needed for the different types of document should be noted. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (see www.iso.org/directives or www.iec.ch/members_experts/refdocs).

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights. Details of any patent rights identified during the development of the document will be in the Introduction and/or on the ISO list of patent declarations received (see www.iso.org/patents), or the IEC list of patent declarations received (see patents.iec.ch).

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation of the voluntary nature of standards, the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the World Trade Organization (WTO) principles in the Technical Barriers to Trade (TBT) see www.iso.org/iso/foreword.html. In the IEC, see www.iec.ch/understanding-standards.

This document was prepared by Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 32, *Data management and interchange*.

This first edition of ISO/IEC 19075-5 cancels and replaces ISO/IEC TR 19075-5:2016.

This document is intended to be used in conjunction with the following editions of the parts of the ISO/IEC 9075 series:

- ISO/IEC 9075-1, sixth edition or later;
- ISO/IEC 9075-2, sixth edition or later;
- ISO/IEC 9075-3, sixth edition or later;
- ISO/IEC 9075-4, seventh edition or later;
- ISO/IEC 9075-9, fifth edition or later;
- ISO/IEC 9075-10, fifth edition or later;
- ISO/IEC 9075-11, fifth edition or later;
- ISO/IEC 9075-13, fifth edition or later;
- ISO/IEC 9075-14, sixth edition or later;
- ISO/IEC 9075-15, second edition or later;
- ISO/IEC 9075-16, first edition or later.

A list of all parts in the ISO/IEC 19075 series can be found on the ISO and IEC websites.

Any feedback or questions on this document should be directed to the user's national standards body. A complete listing of these bodies can be found at www.iso.org/members.html and www.iec.ch/national-committees.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 19075-5:2021

Introduction

This document discusses the syntax and semantics for recognizing patterns in rows of a table, as defined in ISO/IEC 9075-2.

The organization of this document is as follows:

- 1) **Clause 1, “Scope”**, specifies the scope of this document.
- 2) **Clause 2, “Normative references”**, identifies additional standards that, through reference in this document, constitute provisions of this document.
- 3) **Clause 3, “Terms and definitions”**, defines the terms and definitions used in this document.
- 4) **Clause 4, “Row pattern recognition: FROM clause”**, discusses Feature R010, “Row pattern recognition: FROM clause”.
- 5) **Clause 5, “Expressions in MEASURES and DEFINE”**, discusses scalar expression syntax in row pattern matching.
- 6) **Clause 6, “Row pattern recognition: WINDOW clause”**, discusses Feature R020, “Row pattern recognition: WINDOW clause”. Clause 6, “Row pattern recognition: WINDOW clause”, does not duplicate material already presented in Clause 4, “Row pattern recognition: FROM clause” and Clause 5, “Expressions in MEASURES and DEFINE”, which should be read even if the reader is only interested in Feature R020, “Row pattern recognition: WINDOW clause”.
- 7) **Clause 7, “Pattern matching rules”**, discusses the formal rules of pattern matching.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 19075-5:2021

Information technology — Guidance for the use of database language SQL —**Part 5:****Row pattern recognition****1 Scope**

This document discusses the syntax and semantics for recognizing patterns in rows of a table, as defined in [ISO/IEC 9075-2](#), commonly called “SQL/RPR”.

SQL/RPR defines two features regarding row pattern recognition:

- Feature R010, “Row pattern recognition: FROM clause”
- Feature R020, “Row pattern recognition: WINDOW clause”

These two features have considerable syntax and semantics in common, the principle difference being whether the syntax is placed in the FROM clause or in the WINDOW clause.

2 Normative references

The following documents are referred to in the text in such a way that some or all of their content constitutes requirements of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

ISO/IEC 9075-1, *Information technology — Database languages — SQL — Part 1: Framework (SQL/Framework)*

ISO/IEC 9075-2, *Information technology — Database languages — SQL — Part 2: Foundation (SQL/Foundation)*

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 19075-5:2021

3 Terms and definitions

For the purposes of this document, the terms and definitions given in ISO/IEC 9075-1 apply.

ISO and IEC maintain terminological databases for use in standardization at the following addresses:

- IEC Electropedia: available at <http://www.electropedia.org/>
- ISO Online browsing platform: available at <http://www.iso.org/obp>

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 19075-5:2021

4 Row pattern recognition: FROM clause

4.1 Context of row pattern recognition

The requirements for the material discussed in this document shall be as specified in ISO/IEC 9075-1 and ISO/IEC 9075-2.

4.2 Introduction to the FROM clause in row pattern recognition

Feature R010, “Row pattern recognition: FROM clause” of SQL/RPR enhances the capability of the FROM clause with a MATCH_RECOGNIZE clause to specify a row pattern. The syntax and semantics of a row pattern is discussed through examples presented throughout this Clause of this document.

There are two principal variants of the MATCH_RECOGNIZE clause:

- 1) ONE ROW PER MATCH, which returns a single summary row for each match of the pattern (the default).
- 2) ALL ROWS PER MATCH, which returns one row for each row of each match. There are three suboptions, to control whether to also return empty matches or unmatched rows.

4.3 Example of ONE ROW PER MATCH

The following example illustrates MATCH_RECOGNIZE with the ONE ROW PER MATCH option. Let Ticker (Symbol, Tradeday, Price) be a table with three columns representing historical stock prices. Symbol is a character column, Tradeday is a date column, and Price is a numeric column.

NOTE 1 — All examples in this document use mixed-case identifiers for the names of tables, columns, etc., whereas SQL key words are shown in uppercase. Unquoted identifiers are actually equivalent to uppercase, so the column headings of sample results will be shown with the identifiers converted to uppercase.

It is desired to partition the data by Symbol, sort it into increasing Tradeday order, and then detect maximal “V” patterns in Price: a strictly falling price, followed by a strictly increasing price. For each match to a V pattern, it is desired to report the starting price, the price at the bottom of the V, the ending price, and the average price across the entire pattern.

The following query may be used to solve this pattern matching problem:

```
SELECT M.Symbol, /* ticker symbol */
       M.Matchno, /* sequential match number */
       M.Startp, /* starting price */
       M.Bottomp, /* bottom price */
       M.Endp, /* ending price */
       M.Avgp /* average price */
FROM Ticker
     MATCH_RECOGNIZE (
       PARTITION BY Symbol
       ORDER BY Tradeday
       MEASURES MATCH_NUMBER() AS Matchno,
```

```

        A.Price AS Startp,
        LAST (B.Price) AS Bottomp,
        LAST (C.Price) AS Endp,
        AVG (U.Price) AS Avgp
ONE ROW PER MATCH
AFTER MATCH SKIP PAST LAST ROW
PATTERN (A B+ C+)
SUBSET U = (A, B, C)
DEFINE /* A defaults to True, matches any row */
        B AS B.Price < PREV (B.Price),
        C AS C.Price > PREV (C.Price)
) AS M

```

In the example above, the principal syntactic elements of MATCH_RECOGNIZE are presented on separate lines. In this example:

- Ticker is the name of the row pattern input table. In this example, the row pattern input table is a table or view. The row pattern input table may also be a derived table (in-line view).
- MATCH_RECOGNIZE introduces the syntax for row pattern recognition.
- PARTITION BY specifies how to partition the row pattern input table. The PARTITION BY clause is a list of columns of the row pattern input table. This clause is optional; if omitted, there are no row pattern partitioning columns, and the entire row pattern input table constitutes a single row pattern partition.
- ORDER BY specifies how to order the rows within row pattern partitions. The ORDER BY clause is a list of columns of the row pattern input table. This clause is optional; if omitted, the order of rows in row pattern partitions is completely non-deterministic. However, since non-deterministic ordering will defeat the purpose of most row pattern recognition, the ORDER BY clause will usually be specified.
- MEASURES specifies row pattern measure columns, whose values are calculated by evaluating expressions related to the match. The first row pattern measure column in this example uses the special nullary function MATCH_NUMBER(), whose value is the sequential number of a match within a row pattern partition. The third and fourth row pattern measure columns in this example use the LAST operation, which obtains the value of an expression in the last row that is mapped by a row pattern match to a row pattern variable. LAST is one of the row pattern navigation operations introduced by SQL/RPR, discussed in [Subclause 5.6, "Row pattern navigation operations"](#).

The result of the MATCH_RECOGNIZE clause is called the row pattern output table. When ONE ROW PER MATCH is specified, as in this example, the row pattern output table has one column for each row pattern partitioning column and one column for each row pattern measure column.

- ONE ROW PER MATCH specifies that the row pattern output table will have a single row for each match that is found in the row pattern input table.
- AFTER MATCH SKIP clause specifies where to resume looking for the next row pattern match after successfully finding a match. In this example, AFTER MATCH SKIP PAST LAST ROW specifies that pattern matching will resume after the last row of a successful match.
- PATTERN specifies the row pattern that is sought in the row pattern input table. A row pattern is a regular expression using primary row pattern variables. In this example, the row pattern has three primary row pattern variables (A, B, and C).
- SUBSET defines the union row pattern variable U as the union of the primary row pattern variables A, B, and C.
- DEFINE specifies the Boolean condition that defines a primary row pattern variable; a row shall satisfy the Boolean condition in order to be mapped to a particular primary row pattern variable. This example uses PREV, a row pattern navigation operation that evaluates an expression in the previous row. If a primary row pattern variable is not defined in the DEFINE clause, then the definition

4.3 Example of ONE ROW PER MATCH

defaults to a condition that is always true, meaning that any row can be mapped to the primary row pattern variable.

- AS M defines the range variable M to associate with the row pattern output table. This clause is optional; if omitted, then an implementation-dependent range variable is used. Since an implementation-dependent range variable is unknowable to the query writer, the AS clause should not be omitted if there are any other tables in the FROM clause aside from the MATCH_RECOGNIZE.

The processing of MATCH_RECOGNIZE is as follows:

- 1) The row pattern input table is partitioned according to the PARTITION BY clause. Each row pattern partition consists of the set of rows of the row pattern input table that are equal (more precisely, not distinct) on the row pattern partitioning columns.
- 2) Each row pattern partition is ordered according to the ORDER BY clause.
- 3) Each ordered row pattern partition is searched for matches to the PATTERN.
- 4) Pattern matching operates by seeking the match at the earliest row, considering the rows in a row pattern partition in the order specified by the ORDER BY. When there is more than one match at a row, then the most preferred match is taken. The precise rules of pattern matching are discussed in [Clause 7, “Pattern matching rules”](#).
- 5) After a match is found, row pattern matching calculates the row pattern measure columns, which are expressions defined by the MEASURES clause.
- 6) Using ONE ROW PER MATCH, as shown in the example, row pattern recognition generates one row for each match that is found.
- 7) The AFTER MATCH SKIP clause determines where row pattern matching resumes within a row pattern partition after a non-empty match has been found. In the example above, row pattern matching resumes at the next row after the rows mapped by a match (AFTER MATCH SKIP PAST LAST ROW).

[Table 1, “Sample data”](#), illustrates sample data for one row pattern partition of Ticker, shown sorted according to the ORDER BY clause. The sample data contains two matches to the pattern, indicated by arrows showing the mapping to primary row pattern variables in each match.

Table 1 — Sample data

SYMBOL	TRADEDAY	PRICE		
XYZ	2009-06-08	50		
XYZ	2009-06-09	60	→ A	}
XYZ	2009-06-10	49	→ B	
XYZ	2009-06-11	40	→ B	} first match
XYZ	2009-06-12	35	→ B	
XYZ	2009-06-15	45	→ C	}
XYZ	2009-06-16	45		
XYZ	2009-06-17	45	→ A	}
XYZ	2009-06-18	43	→ B	
XYZ	2009-06-19	47	→ C	} second match
XYZ	2009-06-22	52	→ C	
XYZ	2009-06-23	70	→ C	}
XYZ	2009-06-24	60		

The result of the example for this row pattern partition is shown in Table 2, “Results of ONE ROW PER MATCH”.

Table 2 — Results of ONE ROW PER MATCH

SYMBOL	MATCHNO	STARTP	BOTTOMP	ENDP	AVGP
XYZ	1	60	35	45	45.8
XYZ	2	45	43	70	51.4

4.4 Example of ALL ROWS PER MATCH

The previous example can be modified slightly to illustrate ALL ROWS PER MATCH, as follows:

```
SELECT M.Symbol, /* ticker symbol */
       M.Matchno, /* sequential match number */
       M.Tradeday, /* day of trading */
       M.Price, /* price on day of trading */
       M.Classy, /* classifier */
       M.Startp, /* starting price */
       M.Bottomp, /* bottom price */
```

4.4 Example of ALL ROWS PER MATCH

```

M.Endp, /* ending price */
M.Avgp /* average price */
FROM Ticker
MATCH_RECOGNIZE (
  PARTITION BY Symbol
  ORDER BY Tradeday
  MEASURES MATCH_NUMBER() AS Matchno,
             CLASSIFIER() AS Classy,
             A.Price AS Startp,
             FINAL LAST (B.Price) AS Bottomp,
             FINAL LAST (C.Price) AS Endp,
             FINAL AVG (U.Price) AS Avgp
  ALL ROWS PER MATCH
  AFTER MATCH SKIP PAST LAST ROW
  PATTERN (A B+ C+)
  SUBSET U = (A, B, C)
  DEFINE /* A defaults to True, matches any row */
         B AS B.Price < PREV (B.Price),
         C AS C.Price > PREV (C.Price)
) AS M

```

Note that the second row pattern measure column in this example shows the use of the special function CLASSIFIER(), which returns the name of the row pattern variable to which a row is mapped. CLASSIFIER is discussed in [Subclause 5.9, "CLASSIFIER function"](#).

The result of this query on the sample data is shown in [Table 3, "Results of ALL ROWS PER MATCH"](#).

Table 3 — Results of ALL ROWS PER MATCH

SYM BOL	MA TC HNO	TRADEDAY	PRICE	CLAS SY	STAR TP	BOT TOMP	ENDP	AVGP
XYZ	1	2009-06-09	60	A	60	35	45	45.8
XYZ	1	2009-06-10	49	B	60	35	45	45.8
XYZ	1	2009-06-11	40	B	60	35	45	45.8
XYZ	1	2009-06-12	35	B	60	35	45	45.8
XYZ	1	2009-06-15	45	C	60	35	45	45.8
XYZ	2	2009-06-17	45	A	45	43	70	51.4
XYZ	2	2009-06-18	43	B	45	43	70	51.4
XYZ	2	2009-06-19	47	C	45	43	70	51.4
XYZ	2	2009-06-22	52	C	45	43	70	51.4
XYZ	2	2009-06-23	70	C	45	43	70	51.4

ALL ROWS PER MATCH differs from ONE ROW PER MATCH in the following respects:

- 1) ALL ROWS PER MATCH returns one row for each row of each match of the pattern.

4.4 Example of ALL ROWS PER MATCH

- 2) The row pattern output table has a column corresponding to every column of the row pattern input table, not just the row pattern partitioning columns. (Note the column M.Price in the SELECT list. This is a column of the row pattern input table, not a row pattern measure column.)
- 3) The MEASURES clause supports two semantics for expression evaluation, running semantics and final semantics, indicated by the keywords RUNNING and FINAL.
- 4) ALL ROWS PER MATCH provides three suboptions for handling empty matches and unmatched rows. These options are not illustrated in this example; see Subclause 4.12.2, “Handling empty matches”, and Subclause 4.12.3, “Handling unmatched rows”, for examples of these options.

4.5 Summary of the syntax

The complete syntax for row pattern recognition in the FROM clause involves the components shown in Table 4, “Row pattern recognition syntax summary”.

Table 4 — Row pattern recognition syntax summary

Syntactic component	Optional?	Default	Cross reference
row pattern input table	no	—	Subclause 4.6, “The row pattern input table”
row pattern input name	yes	implementation-dependent	Subclause 4.6.2, “The row pattern input name”
row pattern input declared column list	yes	none	Subclause 4.6.3, “The row pattern input declared column list”
MATCH_RECOGNIZE	no	—	Subclause 4.7, “MATCH_RECOGNIZE”
PARTITION BY	yes	row pattern input table constitutes one row pattern partition	Subclause 4.8, “PARTITION BY”
ORDER BY	yes	non-deterministic ordering in each row pattern partition	Subclause 4.9, “ORDER BY”
MEASURES	yes	none	Subclause 4.11, “MEASURES”
ONE ROW PER MATCH or ALL ROWS PER MATCH	yes	ONE ROW PER MATCH	Subclause 4.12, “ONE ROW PER MATCH vs ALL ROWS PER MATCH ”
AFTER MATCH SKIP	yes	AFTER MATCH SKIP PAST LAST ROW	Subclause 4.13, “AFTER MATCH SKIP”
PATTERN	no	—	Subclause 4.14, “PATTERN”
SUBSET	yes	no explicit union row pattern variables	Subclause 4.15, “SUBSET”

Syntactic component	Optional?	Default	Cross reference
DEFINE	no	—	Subclause 4.16, “DEFINE”
row pattern output name	yes	implementation-dependent	Subclause 4.17.2, “Row pattern output name”
row pattern output declared column list	yes	none	Subclause 4.17.3, “Row pattern output declared column list”

4.6 The row pattern input table

4.6.1 Introduction to the row pattern input table

The row pattern input table is the input argument to MATCH_RECOGNIZE. In the examples above, the row pattern input table was Ticker, which is a table or view, or perhaps a named query (defined in a WITH clause). The row pattern input table can also be a derived table (also known as in-line view). For example:

```
FROM ( SELECT S.Name, T.Tradeday, T.Price
        FROM Ticker T, SymbolNames S
        WHERE T.Symbol = S.Symbol )
MATCH_RECOGNIZE ( ... ) AS M
```

The row pattern input table shall not be a <joined table>. The work-around is to use a derived table, such as:

```
FROM ( SELECT * FROM A LEFT OUTER JOIN B ON (A.X = B.Y) )
MATCH_RECOGNIZE ( ... ) AS M
```

Note that column names in the row pattern input table are unambiguous, since it is impossible to use range variables within the MATCH_RECOGNIZE clause to disambiguate. If the row pattern input table is a base table or a view, this is not a problem, since SQL does not allow ambiguous column names in a base table or view. This is only an issue when the row pattern input table is a derived table.

For example, consider a join of two tables, Emp and Dept, each of which has a column called Name. The following is a syntax error:

```
FROM ( SELECT D.Name, E.Name, E.Empno, E.Salary
        FROM Dept D, Emp E
        WHERE D.Deptno = E.Deptno )
MATCH_RECOGNIZE (
    PARTITION BY D.Name
    ... )
```

The preceding example is an error because the range variable D is not visible within the MATCH_RECOGNIZE (the scope of D is just the derived table). Rewriting like this is no help:

```
FROM ( SELECT D.Name, E.Name, E.Empno, E.Salary
        FROM Dept D, Emp E
        WHERE D.Deptno = E.Deptno )
MATCH_RECOGNIZE (
```

```
PARTITION BY Name
... )
```

This rewrite eliminates the use of the range variable D within the MATCH_RECOGNIZE. However, now the error is that Name is ambiguous, because there are two columns of the derived table called Name. The way to handle this is to disambiguate the column names within the derived table itself, like this:

```
FROM ( SELECT D.Name AS DName, E.Name AS EName,
           E.Empno, E.Salary
       FROM Dept D, Emp E
       WHERE D.Deptno = E.Deptno )
MATCH_RECOGNIZE (
  PARTITION BY DName
  ... )
```

4.6.2 The row pattern input name

Optionally, a correlation name for the row pattern input table may be declared, as in this example (equivalent to the example in [Subclause 4.3, “Example of ONE ROW PER MATCH”](#)):

```
SELECT M.Symbol,      /* ticker symbol */
       M.Matchno,     /* sequential match number */
       M.Startp,      /* starting price */
       M.Bottomp,     /* bottom price */
       M.Endp,        /* ending price */
       M.Avgp         /* average price */
FROM Ticker AS T
  MATCH_RECOGNIZE (
    PARTITION BY T.Symbol
    ORDER BY T.Tradeday
    MEASURES MATCH_NUMBER() AS Matchno,
              A.Price AS Startp,
              LAST (B.Price) AS Bottomp,
              LAST (C.Price) AS Endp,
              AVG (U.Price) AS Avgp
    ONE ROW PER MATCH
    AFTER MATCH SKIP PAST LAST ROW
    PATTERN (A B+ C+)
    SUBSET U = (A, B, C)
    DEFINE /* A defaults to True, matches any row */
           B AS B.Price < PREV (B.Price),
           C AS C.Price > PREV (C.Price)
  ) AS M
```

The row pattern input name in this example is T, as defined by the syntax “Ticker AS T”. It is also possible to omit the noise word AS, like this: “Ticker T”.

Specifying the row pattern input name is optional. The examples in [Subclause 4.3, “Example of ONE ROW PER MATCH”](#), and [Subclause 4.4, “Example of ALL ROWS PER MATCH”](#), do not show an explicit row pattern input name.

When the row pattern input name is not specified, the following defaults apply:

- 1) If the row pattern input table is a base table, view, or query name (the name of a query defined in a WITH clause), then the table name, view name or query name is the default row pattern input name.
- 2) Otherwise, an implementation-dependent row pattern input name, different from any other range variable in the query, is implicit. In practice, this means that the row pattern input name is unknowable and cannot be referenced elsewhere in the query.

4.6 The row pattern input table

The scope of the row pattern input name is the PARTITION BY and ORDER BY clauses of the MATCH_RECOGNIZE clause. This means that the row pattern input name can be used in the following contexts:

- 1) To qualify column names in the PARTITION BY clause.
- 2) To qualify column names in the ORDER BY clause.

The example above illustrates both of these uses.

The row pattern input name cannot be referenced in the MEASURES or DEFINE clauses, nor elsewhere in the query, such as the WHERE clause or the SELECT list.

4.6.3 The row pattern input declared column list

If an explicit row pattern input name is specified, it may be followed by a parenthesized list of column names, as in this example:

```
SELECT M.Sym,          /* ticker symbol */
       M.Matchno,      /* sequential match number */
       M.Startp,       /* starting price */
       M.Bottomp,      /* bottom price */
       M.Endp,         /* ending price */
       M.Avgp,         /* average price */
FROM Ticker AS T (Sym, Td, Pr)
  MATCH_RECOGNIZE (
    PARTITION BY T.Sym
    ORDER BY T.Td
    MEASURES MATCH_NUMBER() AS Matchno,
              A.Pr AS Startp,
              LAST (B.Pr) AS Bottomp,
              LAST (C.Pr) AS Endp,
              AVG (U.Pr) AS Avgp
    ONE ROW PER MATCH
    AFTER MATCH SKIP PAST LAST ROW
    PATTERN (A B+ C+)
    SUBSET U = (A, B, C)
    DEFINE /* A defaults to True, matches any row */
           B AS B.Pr < PREV (B.Pr),
           C AS C.Pr > PREV (C.Pr)
  ) AS M
```

The parenthesized list of column names (Sym, Td, Pr) is called the row pattern input declared column list. The row pattern input declared column list may be used to change the names of the columns of the row pattern input table. There shall be exactly the same number of column names in the list as there are columns in the row pattern input table. In this example, Symbol has been renamed to Sym, Tradeday has been renamed to Td, and Price has been renamed to Pr. Consequently, the columns cannot be referenced as Symbol, Tradeday, or Price within the MATCH_RECOGNIZE; instead, they are referenced by their new names, Sym, Td, and Pr. Note that this also changes the default names of the columns in the row pattern output table. Thus, in the SELECT list, the first item is M.Sym, because the input column names Symbol was renamed to Sym, which becomes the name of the corresponding output column.

4.7 MATCH_RECOGNIZE

MATCH_RECOGNIZE is the keyword that introduces the syntax for row pattern recognition in the FROM clause. Syntactically, MATCH_RECOGNIZE is a postfix operator following the row pattern input table. The MATCH_RECOGNIZE keyword is followed by a parenthesized list of syntactic components that collectively describe the row pattern recognition operation.

4.8 PARTITION BY

PARTITION BY is used to specify that the rows of the row pattern input table are to be partitioned by one or more columns. Note that the column names in the PARTITION BY may be unqualified, or they may be qualified by the row pattern input name. See the examples in Subclause 4.6.2, “The row pattern input name”.

If there is no PARTITION BY, then all rows of the row pattern input table constitute a single row pattern partition.

4.9 ORDER BY

ORDER BY is used to specify the order of rows within a row pattern partition. The ORDER BY clause of a MATCH_RECOGNIZE is similar to the ORDER BY clause of a cursor. As with the PARTITION BY clause, column names may be unqualified, or they may be qualified by the row pattern input name. See the examples in Subclause 4.6.2, “The row pattern input name”.

If the order of two rows in a row pattern partition is not determined by the ORDER BY, then the result of MATCH_RECOGNIZE is non-deterministic.

NOTE 2 — Syntactically, the row pattern output table is always regarded as non-deterministic because there is no way for the query engine to deduce at compile time whether the ordering is total. This means that MATCH_RECOGNIZE cannot be used in contexts that are required to be deterministic, such as check constraints and assertions. However, the query author can use ORDER BY to insure that the query is sufficiently deterministic for the author’s intended purpose.

4.10 Row pattern variables

Row pattern variables are range variables whose scope is limited to a MATCH_RECOGNIZE clause. As range variables, row pattern variables are used to qualify column references, in either the scalar expression of a row pattern measure column, or the Boolean condition of a DEFINE.

There are two kinds of row pattern variables:

- 1) Primary row pattern variables, which are declared in the PATTERN and defined by an associated Boolean condition specified in the DEFINE clause.
- 2) Union row pattern variables, which are declared in the SUBSET clause as a union of a list of primary row pattern variables. The primary row pattern variables are called components of the union row pattern variable.

A row pattern variable shall not be both a primary row pattern variable and a union row pattern variable. This means that a row pattern variable that is declared in PATTERN shall not also be declared on the left hand side of a SUBSET.

ISO/IEC 19075-5:2021(E)
4.10 Row pattern variables

Informally, a match consists of a set of contiguous rows in a row pattern partition of the row pattern input table. (For a more formal treatment, see [Clause 7, “Pattern matching rules”](#).) Each row of the match is mapped to a primary row pattern variable. The mapping of rows to primary row pattern variables conforms to the regular expression in the PATTERN clause, and is further constrained to insure that all Boolean conditions in the DEFINE clause are true.

Thus rows are mapped to row pattern variables. Conversely, each row pattern variable *RPV* has a set of rows that are mapped to *RPV*. For example, given:

```
PATTERN (A+ (B+ | C+) D)
SUBSET S = (B, D)
```

Suppose that consecutive rows R_3 , R_4 , R_5 , R_6 , and R_7 are mapped as follows:

$R_3 \rightarrow A$

$R_4 \rightarrow A$

$R_5 \rightarrow B$

$R_6 \rightarrow B$

$R_7 \rightarrow D$

Then:

- the set of rows mapped to A is $\{ R_3, R_4 \}$,
- the set of rows mapped to B is $\{ R_5, R_6 \}$,
- the set of rows mapped to C is empty, and
- the set of rows mapped to D is $\{ R_7 \}$.

The set of rows mapped to a union row pattern variable *URPV* can be obtained as the set union of rows mapped to each component of *URPV*. In this example:

- the set of rows mapped to S is $\{ R_5, R_6 \} \cup \{ R_7 \} = \{ R_5, R_6, R_7 \}$.

There is always one implicit union row pattern variable, called the universal row pattern variable, defined as the union of all primary row pattern variables. Thus, every row of a match is mapped to the universal row pattern variable. The universal row pattern variable is used to implicitly qualify unqualified column reference within the MEASURES or DEFINE clauses. There is no syntax available to the user to denote the universal row pattern variable. The query writer may, of course, define an explicit union row pattern variable that is the union of all primary row pattern variables. (The example in [Subclause 4.3, “Example of ONE ROW PER MATCH”](#), illustrates this technique.)

4.11 MEASURES

The MEASURES clause defines row pattern measure columns, which are columns of the row pattern output table whose value is computed by evaluating an expression related to a particular match. Note that SQL/RPR extends the scalar expression syntax of [ISO/IEC 9075-2](#), and provides special semantics for evaluating scalar expressions in the context of a row pattern match. This is discussed in [Clause 5, “Expressions in MEASURES and DEFINE”](#).

NOTE 3 — The MEASURES clause in a window definition does not define columns; instead, it defines named expressions which are accessed using a variant of the window function syntax, called row pattern measure functions. “Row pattern

measure” is the generic term for row pattern measure columns and row pattern measure functions, whose values are computed using the same rules.

4.12 ONE ROW PER MATCH vs ALL ROWS PER MATCH

4.12.1 Introduction to use of ROWS PER MATCH

ONE ROW PER MATCH indicates that the result has one row for each match. Columns of this row are defined by the PARTITION and MEASURES clauses. This is the default.

ALL ROWS PER MATCH indicates that the result has one row for each row of each match. (It is possible to exclude some rows using the exclusion syntax { - - } in the PATTERN; see [Subclause 4.14.3, “Excluding portions of the pattern”](#).)

ALL ROWS PER MATCH has three suboptions:

- ALL ROWS PER MATCH SHOW EMPTY MATCHES
- ALL ROWS PER MATCH OMIT EMPTY MATCHES
- ALL ROWS PER MATCH WITH UNMATCHED ROWS

These options are explained in the following subsections.

4.12.2 Handling empty matches

Some patterns permit empty matches. For example:

PATTERN (A*)

can be matched by zero or more rows that are mapped to A.

An empty match does not map any rows to primary row pattern variables; nevertheless, an empty match has a starting row. For example, there can be an empty match at the first row of a row pattern partition, an empty match at the second row of a row pattern partition, etc. An empty match is assigned a sequential match number, based on the ordinal position of its starting row, the same as any other match.

When using ONE ROW PER MATCH, an empty match results in one row of the row pattern output table. The row pattern measures for an empty match are computed as follows:

- The value of MATCH_NUMBER() is the sequential match number of the empty match.
- Any COUNT is 0.
- Any other aggregate, row pattern navigation operation, or ordinary row pattern column reference is null.

For example, the example in [Subclause 4.3, “Example of ONE ROW PER MATCH”](#), can be modified to permit empty matches, as follows:

```
SELECT M.Symbol, /* ticker symbol */
       M.Matchno, /* sequential match number */
       M.Firstp, /* starting price */
       M.Lastp /* ending price */
FROM Ticker
```

4.12 ONE ROW PER MATCH vs ALL ROWS PER MATCH

```

MATCH_RECOGNIZE (
  PARTITION BY Symbol
  ORDER BY Tradeday
  MEASURES MATCH_NUMBER() AS Matchno,
             FIRST A.Price AS Firstp,
             LAST (A.Price) AS Lastp
  ONE ROW PER MATCH
  AFTER MATCH SKIP PAST LAST ROW
  PATTERN (A*)
  DEFINE
    A AS A.Price > PREV (A.Price)
) AS M

```

Here the pattern has been changed to A*, and is used to detect runs of increasing prices. The sample data is now analyzed as shown in Table 5, “Analysis of sample data permitting empty matches”.

Table 5 — Analysis of sample data permitting empty matches

SYMBOL	TRADEDAY	PRICE		
XYZ	2009-06-08	50		match #1 (empty)
XYZ	2009-06-09	60	→ A	match #2
XYZ	2009-06-10	49		match #3 (empty)
XYZ	2009-06-11	40		match #4 (empty)
XYZ	2009-06-12	35		match #5 (empty)
XYZ	2009-06-15	45	→ A	match #6
XYZ	2009-06-16	45		match #7 (empty)
XYZ	2009-06-17	45		match #8 (empty)
XYZ	2009-06-18	43		match #9 (empty)
XYZ	2009-06-19	47	→ A	}
XYZ	2009-06-22	52	→ A	{ match #10
XYZ	2009-06-23	70	→ A	}
XYZ	2009-06-24	60		match #11 (empty)

The result of the preceding query on the sample row pattern partition is illustrated in Table 6, “Result of query permitting empty matches”.

Table 6 — Result of query permitting empty matches

SYMBOL	MATCHNO	FIRSTP	LASTP
XYZ	1		
XYZ	2	60	60

SYMBOL	MATCHNO	FIRSTP	LASTP
XYZ	3		
XYZ	4		
XYZ	5		
XYZ	6	45	45
XYZ	7		
XYZ	8		
XYZ	9		
XYZ	10	47	70
XYZ	11		

In the preceding result, note how the row pattern measures other than the match number are null for empty matches.

As for ALL ROWS PER MATCH, the question arises of whether to generate a row of output for an empty match, seeing that there are no rows in the empty match. To govern this, there are two options:

- 1) ALL ROWS PER MATCH SHOW EMPTY MATCHES: with this option, any empty match generates a single row in the row pattern output table.
- 2) ALL ROWS PER MATCH OMIT EMPTY MATCHES: with this option, an empty match is omitted from the row pattern output table. (This may cause gaps in the sequential match numbering.)

ALL ROWS PER MATCH defaults to SHOW EMPTY MATCHES. Using this option, an empty match generates one row in the row pattern output table. In this row:

- The value of a classifier function is null.
- The value of MATCH_NUMBER() is the sequential match number of the empty match.
- The value of any ordinary row pattern column reference is null.
- The value of any aggregate or row pattern navigation operation is computed using an empty set of rows (so any COUNT is 0, and all other aggregates and row pattern navigation operations are null).
- The value of any column corresponding to a column of the row pattern input table is the same as the corresponding column in the starting row of the empty match.

The following example alters the preceding example slightly, to use ALL ROWS PER MATCH SHOW EMPTY MATCHES:

```
SELECT M.Symbol,    /* ticker symbol */
       M.Matchno,   /* sequential match number */
       M.Tradeday,  /* day of trading */
       M.Price,     /* price on day of trading */
       M.Classy,    /* classifier */
       M.Firstp,    /* starting price */
       M.Lastp,     /* ending price */
FROM Ticker
     MATCH_RECOGNIZE (
```

4.12 ONE ROW PER MATCH vs ALL ROWS PER MATCH

```

PARTITION BY Symbol
ORDER BY Tradeday
MEASURES MATCH_NUMBER() AS Matchno,
          CLASSIFIER AS Classy,
          FINAL FIRST (A.Price) AS Firstp,
          FINAL LAST (A.Price) AS Lastp
ALL ROWS PER MATCH SHOW EMPTY MATCHES
AFTER MATCH SKIP PAST LAST ROW
PATTERN (A*)
DEFINE A AS A.Price > PREV (A.Price)
) AS M

```

The result of the preceding query on the sample row pattern partition is shown in Table 7, “Results of query using SHOW EMPTY ROWS”.

Table 7 — Results of query using SHOW EMPTY ROWS

SYMBOL	MATCH NO	TRADEDAY	PRICE	CLASSY	FIRSTP	LASTP
XYZ	1	2009-06-08	50			
XYZ	2	2009-06-09	60	A	60	60
XYZ	3	2009-06-10	49			
XYZ	4	2009-06-11	40			
XYZ	5	2009-06-12	35			
XYZ	6	2009-06-15	45	A	45	45
XYZ	7	2009-06-16	45			
XYZ	8	2009-06-17	45			
XYZ	9	2009-06-18	43			
XYZ	10	2009-06-19	47	A	47	70
XYZ	10	2009-06-22	52	A	47	70
XYZ	10	2009-06-23	70	A	47	70
XYZ	11	2009-06-24	60			

If, instead, ALL ROWS PER MATCH OMIT EMPTY MATCHES were used, the result would lack the rows for the empty matches, like illustrated in Table 8, “Results of query using OMIT EMPTY ROWS”.

Table 8 — Results of query using OMIT EMPTY ROWS

SYMBOL	MATCH NO	TRADEDAY	PRICE	CLASSY	FIRSTP	LASTP
XYZ	2	2009-06-09	60	A	60	60
XYZ	6	2009-06-15	45	A	45	45

SYMBOL	MATCH NO	TRADEDAY	PRICE	CLASSY	FIRSTP	LASTP
XYZ	10	2009-06-19	47	A	47	70
XYZ	10	2009-06-22	52	A	47	70
XYZ	10	2009-06-23	70	A	47	70

Note the gaps in the match numbering; also, the final empty match (number 11) is undetectable because there are no non-empty matches following it.

4.12.3 Handling unmatched rows

Some rows of the row pattern input table may be neither the starting row of an empty match, nor mapped by a non-empty match. Such rows are called unmatched rows.

The option ALL ROWS PER MATCH WITH UNMATCHED ROWS shows both empty matches and unmatched rows. Empty matches are handled the same as with SHOW EMPTY MATCHES. When displaying an unmatched row, all row pattern measures are null, somewhat analogous to the null-extended side of an outer join. Thus COUNT and MATCH_NUMBER may be used to distinguish an unmatched row from the starting row of an empty match. The exclusion syntax { - - } is prohibited as contrary to the spirit of WITH UNMATCHED ROWS.

The example in Subclause 4.4, “Example of ALL ROWS PER MATCH”, can be used to illustrate WITH UNMATCHED ROWS. The change in the query syntax is:

```

SELECT M.Symbol,      /* ticker symbol */
       M.Matchno,     /* sequential match number */
       M.Tradeday,    /* day of trading */
       M.Price,        /* price on day of trading */
       M.Classy,       /* classifier */
       M.Startp,       /* starting price */
       M.Bottomp,      /* bottom price */
       M.Endp,         /* ending price */
       M.Avgp,         /* average price */
FROM Ticker
      MATCH_RECOGNIZE (
  PARTITION BY Symbol
  ORDER BY Tradeday
  MEASURES MATCH_NUMBER() AS Matchno,
            CLASSIFIER AS Classy,
            A.Price AS Startp,
            FINAL LAST (B.Price) AS Bottomp,
            FINAL LAST (C.Price) AS Endp,
            FINAL AVG (U.Price) AS Avgp
  ALL ROWS PER MATCH WITH UNMATCHED ROWS
  AFTER MATCH SKIP PAST LAST ROW
  PATTERN (A B+ C+)
  SUBSET U = (A, B, C)
  DEFINE /* A defaults to True, matches any row */
         B AS B.Price < PREV (B.Price),
         C AS C.Price > PREV (C.Price)
) AS M

```

4.12 ONE ROW PER MATCH vs ALL ROWS PER MATCH

and the result on the data in the sample row pattern partition is shown in Table 9, “Results of ALL ROWS PER MATCH”.

Table 9 — Results of ALL ROWS PER MATCH

SYM BOL	MA TC HN O	TRADE DAY	PRICE	CLAS SY	STAR TP	BOT TOMP	ENDP	AVGP
XYZ		2009-06-08	50					
XYZ	1	2009-06-09	60	A	60	35	45	45.8
XYZ	1	2009-06-10	49	B	60	35	45	45.8
XYZ	1	2009-06-11	40	B	60	35	45	45.8
XYZ	1	2009-06-12	35	B	60	35	45	45.8
XYZ	1	2009-06-15	45	C	60	35	45	45.8
XYZ		2009-06-16	45					
XYZ	2	2009-06-17	45	A	45	43	70	51.4
XYZ	2	2009-06-18	43	B	45	43	70	51.4
XYZ	2	2009-06-19	47	C	45	43	70	51.4
XYZ	2	2009-06-22	52	C	45	43	70	51.4
XYZ	2	2009-06-23	70	C	45	43	70	51.4
XYZ		2009-06-24	60					

In the sample output, note the rows in which the row pattern measures are null. These rows correspond to unmatched rows in the row pattern input table.

It is not possible for a pattern to permit empty matches and also have unmatched rows. The reason is that if a row of the row pattern input table cannot be mapped to a primary row pattern variable, then that row can still be the starting row of an empty match, and will not be regarded as unmatched, assuming that the pattern permits empty matches. Thus, if a pattern permits empty matches, then the output using ALL ROWS PER MATCH SHOW EMPTY MATCHES is the same as the output using ALL ROWS PER MATCH WITH UNMATCHED ROWS. Thus WITH UNMATCHED ROWS is primarily intended for use with patterns that do not permit empty matches. However, the user may prefer to specify WITH UNMATCHED ROWS if the user is uncertain whether a pattern may have empty matches or unmatched rows.

Note that if ALL ROWS PER MATCH WITH UNMATCHED ROWS is used with the default skipping behavior (AFTER MATCH SKIP PAST LAST ROW), then every row of the row pattern input table will appear exactly once in the output (as the location of an empty match, as a row that is mapped by a non-empty match, or as an unmatched row).

Other skipping behaviors are permitted using WITH UNMATCHED ROWS, in which case it becomes possible for a row to be mapped by more than one match and appear in the row pattern output table multiple times. Unmatched rows will appear in the output only once.

4.13 AFTER MATCH SKIP

The AFTER MATCH SKIP clause determines the point to resume pattern matching after a non-empty match has been found. The default for the clause is AFTER MATCH SKIP PAST LAST ROW. The options are as follows (*RPV* denotes a row pattern variable):

- AFTER MATCH SKIP TO NEXT ROW: resume pattern matching at the row after the first row of the current match.
- AFTER MATCH SKIP PAST LAST ROW: resume pattern matching at the next row after the last row of the current match.
- AFTER MATCH SKIP TO FIRST *RPV*: resume pattern matching at the first row that is mapped to the row pattern variable *RPV*.
- AFTER MATCH SKIP TO LAST *RPV*: resume pattern matching at the last row that is mapped to the row pattern variable *RPV*.
- AFTER MATCH SKIP TO *RPV*: same as AFTER MATCH SKIP TO LAST *RPV*.

When using AFTER MATCH SKIP TO FIRST or AFTER MATCH SKIP TO [LAST], it is possible that no row is mapped to the <row pattern variable name>. For example, the row pattern variable A in

```
AFTER MATCH SKIP TO A
PATTERN (X A* X),
```

might have no rows mapped to A. If there is no row mapped to A, then there is no row to skip to, so a run-time exception is generated.

Another aberrant condition is that AFTER MATCH SKIP may try to resume pattern matching at the same row that the last match started. For example,

```
AFTER MATCH SKIP TO X
PATTERN (X Y+ Z),
```

In this example, AFTER MATCH SKIP TO X tries to resume pattern matching at the same row where the previous match was found. This would result in an infinite loop; consequently a run-time exception is generated for this scenario.

Note that the AFTER MATCH SKIP syntax only determines the point to resume scanning for a match after a non-empty match. When an empty match is found, one row is skipped (as if SKIP TO NEXT ROW had been specified). Thus an empty match never causes one of these exceptions.

A query that gets one of these exceptions should be rewritten. For example,

```
AFTER MATCH SKIP TO A
PATTERN (X (A | B) Y)
```

will cause a run-time error if alternative A does not match. Instead of this example, perhaps the following will serve the user's needs:

```
AFTER MATCH SKIP TO C
PATTERN (X (A | B) Y)
SUBSET C = (A, B)
```

In the revised example, no run-time error is possible, whether A or B is matched.

As another example:

```
AFTER MATCH SKIP TO FIRST A
PATTERN (A* X)
```

This example will always get an exception after the first match, either for skipping to the first row of the match (if A^* matches) or for skipping to a non-existent row (if A^* does not match). In this example, SKIP TO NEXT ROW might be a better choice.

When using ALL ROWS PER MATCH together with skip options other than AFTER MATCH SKIP PAST LAST ROW, it is possible for consecutive matches to overlap, in which case a row R of the row pattern input table might occur in more than one match. In that case, the row pattern output table will have one row for each match in which R participates. The MATCH_NUMBER function may be used to distinguish between the multiple matches in which a row of the row pattern input table participates. When a row participates in more than one match, its classifier may be different in each match as well.

4.14 PATTERN

4.14.1 Introduction to the PATTERN syntax

The PATTERN clause is used to specify a regular expression. The regular expression is enclosed in parentheses. It is built from primary row pattern variables, and may use the following operators:

- concatenation: indicated by the absence of any operator sign between two successive items in a pattern. Note that whitespace is required to delimit two successive primary row pattern variables.
- quantifiers: quantifiers are postfix operators with the following choices:
 - * — 0 or more iterations
 - + — 1 or more iterations
 - ? — 0 or 1 iterations
 - { n } — exactly n iterations ($n > 0$)
 - { n , } — n or more iterations ($n \geq 0$)
 - { n , m } — between n and m (inclusive) iterations ($0 \leq n \leq m$, $0 < m$)
 - { , m } — between 0 and m (inclusive) iterations ($m > 0$)reluctant quantifiers, indicated by an additional question mark (*?, +?, ??, { n }?, { n ,}?, { n , m }?, { m }?). See below for the difference between reluctant and non-reluctant quantifiers.
- alternation: indicated by a vertical bar (|). Alternatives are preferred in the order in which they are specified.
- grouping: indicated by parentheses.
- PERMUTE: see Subclause 4.14.2, “PERMUTE”.
- exclusion: parts of the pattern to be excluded from the output of ALL ROWS PER MATCH are enclosed between { - and - }. See Subclause 4.14.3, “Excluding portions of the pattern”.
- anchors (not permitted with row pattern matching in windows):
 - ^: matches the beginning of a row pattern partition
 - \$: matches the end of a row pattern partition
- (): empty pattern, matches an empty set of rows

The difference between non-reluctant (or “greedy”) and reluctant quantifiers appended to a single row pattern variable is illustrated as follows: A^* tries to map as many rows as possible to A (consistent with mapping the entire pattern), whereas $A^*?$ tries to map as few rows as possible to A (consistent with mapping the entire pattern). The semantics of quantifiers on complex regular expressions, such as $(A | B)^*$, cannot be expressed succinctly; see [Subclause 7.2.4, “Quantification”](#),

The precedence of the operators in a regular expression, in decreasing order, is as follows:

- primaries: primary row pattern variables, anchors, PERMUTE, parenthetic expressions, exclusion syntax, empty pattern
- quantifier; a primary may have zero or one quantifier
- concatenation
- alternation

Precedence of alternation is illustrated by this example:

PATTERN (A B | C D)

which is equivalent to

PATTERN ((A B) | (C D))

and is not equivalent to

PATTERN (A (B | C) D)

Precedence of quantifiers is illustrated by this example:

PATTERN (A B *)

which is equivalent to

PATTERN (A (B*))

and is not equivalent to

PATTERN ((A B)*)

A quantifier shall not immediately follow another quantifier. For example

PATTERN (A**)

is prohibited, whereas

PATTERN ((A*)*)

is permitted (though the latter pattern is no more powerful than just A^*).

It is permitted for a primary row pattern variable to occur more than once in a pattern. For example

PATTERN (X Y X)

4.14.2 PERMUTE

The PERMUTE syntax may be used to express a pattern that is a permutation of simpler patterns. For example,

PATTERN (PERMUTE (A, B, C))

is equivalent to an alternation of all permutations of three row pattern variables A, B and C, like this:

PATTERN (A B C | A C B | B A C | B C A | C A B | C B A)

Note that PERMUTE is expanded lexicographically. (In this example, since the three row pattern variables A, B, and C are listed in alphabetic order, it follows from lexicographic expansion that the expanded possibilities are also listed in alphabetic order.) This is significant because alternatives are attempted in the order written in the expansion. Thus a match to (A B C) will be attempted before a match to (A C B), etc.; the first attempt that succeeds is the “winner”.

As another example:

PATTERN (PERMUTE (X{3}, B C?, D))

is equivalent to

```
PATTERN (
  ( X{3} B C? D )
| ( X{3} D B C? )
| ( B C? X{3} D )
| ( B C? D X{3} )
| ( D X{3} B C? )
| ( D B C? X{3} ) )
```

4.14.3 Excluding portions of the pattern

When using ALL ROWS PER MATCH with either the OMIT EMPTY MATCHES or SHOW EMPTY MATCHES suboptions, rows matching a portion of the PATTERN may be excluded from the row pattern output table. The excluded portion is bracketed between {- and -} in the PATTERN clause.

For example, the following example finds the longest periods of increasing prices that start with a price no less than 10.

```
SELECT M.Symbol,      /* row's symbol */
       M.Tradeday,    /* row's trade day */
       M.Price,        /* row's price */
       M.Avgp,         /* average price */
       M.Matchno       /* row's match number */
FROM Ticker
     MATCH_RECOGNIZE (
       PARTITION BY Symbol
       ORDER BY Tradeday
       MEASURES FINAL AVG (S.Price) AS Avgp,
       MATCH_NUMBER() AS Matchno
       ALL ROWS PER MATCH
       AFTER MATCH SKIP TO LAST B
       PATTERN ( {- A -} B+ {- C -} )
       SUBSET S = (A, B)
       DEFINE A AS A.Price >= 10
              B AS B.Price > PREV (B.Price),
              C AS C.Price <= PREV (C.Price)
     ) AS M;
```

The row pattern output table will only have rows that are mapped to B; the rows mapped to A and C will be excluded from the output.

Although the excluded rows do not appear in the row pattern output table, they are not excluded from the definitions of union row pattern variables, nor from the calculation of scalar expressions in the DEFINE or MEASURES. For example, see the definitions of the primary row pattern variables A and C, the definition of union row pattern variable S, or the Avgp row pattern measure in the example above.

Also, excluded rows do not alter the behavior of AFTER MATCH SKIP. That is, excluded rows are still used in deciding where to resume looking for the next match. For example, in the example above, suppose the AFTER MATCH SKIP clause were changed to

```
AFTER MATCH SKIP PAST LAST ROW
```

while leaving the pattern the same:

```
PATTERN ( {- A -} B+ {- C -} )
```

In that case, a match to the pattern maps a row to the row pattern variable C, and the skip will be to the next row after the last row of the match; that is, after the row that is mapped to C, even though the row that is mapped to C is excluded from the output.

The exclusion syntax is not permitted with ALL ROWS PER MATCH WITH UNMATCHED ROWS.

The exclusion syntax is permitted with ONE ROW PER MATCH, though it has no effect since in this case there is only a single summary row per match.

4.15 SUBSET

The SUBSET clause is optional. It is used to declare union row pattern variables. For example:

```
FROM Ticker
MATCH_RECOGNIZE
( ORDER BY Tradeday
  MEASURES FIRST (X.time) AS x_firsttime,
    LAST (Y.time) AS y_lasttime,
    AVG (S.Price) AS xy_avgprice
  PATTERN (X+ Y+)
  SUBSET S = (X, Y)
  DEFINE X AS X.Price > PREV (X.Price),
    Y AS Y.Price < PREV (Y.Price)
)
```

This example declares a union row pattern variable, S, and defines it as the union of the rows mapped to X and the rows mapped to Y. See [Subclause 4.10, "Row pattern variables"](#), for an example of how such unions are formed.

There can be multiple union row pattern variables. For example:

```
PATTERN (W+ X+ Y+ Z+ )
SUBSET A = (X, Y),
  B = (W, Z)
```

The right hand side of a SUBSET item is a parenthesized, comma-separated list of distinct primary row pattern variables. This defines the union row pattern variable (on the left hand side) as the union of the primary row pattern variables (on the right hand side).

Note that the list of row pattern variables on the right hand side cannot include any union row pattern variables (there are no unions of unions).

4.16 DEFINE

DEFINE is a mandatory clause, used to specify the Boolean condition that defines a primary row pattern variable. In the example,

```
DEFINE X AS X.Price > PREV (X.Price),
      Y AS Y.Price < PREV (Y.Price)
```

X is defined by the condition $X.Price > PREV(X.Price)$, and Y is defined by the condition $Y.Price < PREV(Y.Price)$. (PREV is a row pattern navigation operation which evaluates an expression in the previous row; see [Subclause 5.6, “Row pattern navigation operations”](#), regarding the complete set of row pattern navigation operations.)

A primary row pattern variable does not require a definition; if there is no definition, the default is a predicate that is always true. Any row can be mapped to such a primary row pattern variable.

A union row pattern variable cannot be defined by DEFINE, but may appear in the Boolean condition of a primary row pattern variable.

The Boolean condition of a primary row pattern variable *RPV* may reference *RPV*, or other primary or union row pattern variables. For example:

```
FROM Ticker
MATCH_RECOGNIZE
( PARTITION BY Symbol
  ORDER BY Tradeday
  MEASURES FIRST (A.Tradeday) AS A_Firstday,
  LAST (D.Tradeday) AS D_Lastday,
  AVG (B.Price) AS B_Avgprice,
  AVG (D.Price) AS D_Avgprice
  PATTERN ( A B+ C+ D )
  SUBSET BC = (B, C)
  DEFINE A AS Price > 100,
         B AS B.Price > A.Price,
         C AS C.Price < AVG (B.Price),
         D AS D.Price > MAX (BC.Price)
) AS M
```

In this example:

- The definition of A implicitly references the universal row pattern variable (because of the unqualified column reference Price).
- The definition of B references the primary row pattern variable A.
- The definition of C references the primary row pattern variable B.
- The definition of D references the union row pattern variable BC.

The Boolean conditions are evaluated on successive rows of a row pattern partition in a trial match, with the current row being tentatively mapped to a primary row pattern variable *PRPV* as permitted by the pattern. To be successfully mapped to *PRPV*, the Boolean condition that defines *PRPV* shall evaluate to True.

In the preceding example:

```
A AS Price > 100
```

Here Price is an unqualified column reference, so it is implicitly qualified by the universal row pattern variable. All rows that are already mapped, including the current row, are mapped to the universal

row pattern variable. Also, Price is an ordinary row pattern column reference, so it is evaluated in the last row mapped to the universal row pattern variable, i.e., the current row. Thus this condition is equivalent to `A AS A.Price > 100`.

`B AS B.Price > A.Price`

Here B.Price and A.Price are ordinary row pattern column references. B.Price refers to the current row (since B is being defined), whereas A.Price refers to the last row mapped to A. In view of the pattern, the only row mapped to A is the first row to be mapped.

`C AS C.Price < AVG (B.Price)`

Here C.Price refers to the Price in the current row, since C is being defined. The aggregate `AVG (B.Price)` is computed as the average of all rows that are *already* mapped to B (but not to any rows that might be mapped to B later).

`D AS D.Price > MAX (BC.Price)`

This example is similar to the preceding, though it illustrates the use of a union row pattern variable in the Boolean condition.

The semantics of Boolean conditions are discussed in more detail in [Clause 5, “Expressions in MEASURES and DEFINE”](#).

4.17 The row pattern output table

4.17.1 Introduction to the row pattern output table

The result of `MATCH_RECOGNIZE` is called the row pattern output table. The shape (row type) of the row pattern output table depends on the choice of `ONE ROW PER MATCH` or `ALL ROWS PER MATCH`:

- If `ONE ROW PER MATCH` is specified or implied, then the columns of the row pattern output table are the row pattern partitioning columns in their order of declaration, followed by the row pattern measure columns in their order of declaration. Since a table is required to have at least one column, this implies that there shall be at least one row pattern partitioning column or one row pattern measure column.
- If `ALL ROWS PER MATCH` is specified, then the columns of the row pattern output table are the row pattern partitioning columns in their order of declaration, the ordering columns in their order of declaration, the row pattern measure columns in their order of declaration, and finally any remaining columns of the row pattern input table, in the order they occur in the row pattern input table.

The order of columns in the row pattern output table is only significant when using `SELECT *`. The order of columns is designed to facilitate comparing the output when the query is toggled between `ONE ROW PER MATCH` and `ALL ROWS PER MATCH`.

The names and declared types of the row pattern measure columns are determined by the `MEASURES` clause. The names and declared types of the non-measure columns are inherited from the corresponding columns of the row pattern input table.

4.17 The row pattern output table

4.17.2 Row pattern output name

Optionally, a correlation name may be assigned to the row pattern output table, like this:

```
SELECT T.Matchno
FROM Ticker
    MATCH_RECOGNIZE ( ...
        MEASURES MATCH_NUMBER () AS Matchno
        ...
    ) AS T
```

In the preceding example, M is the correlation name assigned to the row pattern output table. The noise word AS is optional.

The benefit to assigning a correlation name is that the correlation name may be used to qualify the column names of the row pattern output table, as in M.Matchno in the preceding example. This is especially important to resolve ambiguous column names if there are other tables in the FROM clause. For example, suppose Matchmaker is a table with a column named Matchno, to be joined with the row pattern recognition shown above. In that case the query might be written:

```
SELECT T.Matchno, M.Matchno
FROM Ticker
    MATCH_RECOGNIZE ( ...
        MEASURES MATCH_NUMBER () AS Matchno
        ...
    ) AS T, Matchmaker AS M
WHERE ...
```

4.17.3 Row pattern output declared column list

Optionally, the row pattern output name may be followed by a parenthesized list of column names, called the row pattern output declared column list. The row pattern output declared column list may be used to rename the columns of the row pattern output table. There shall be the same number of column names in the list as there are columns in the row pattern output table. The column names in the list are in one-to-one correspondence with the columns of the row pattern output table.

For example, the following is a modification of the example in Subclause 4.6.3, “The row pattern input declared column list”:

```
SELECT M.Cym,          /* ticker symbol */
       M.Mno,          /* sequential match number */
       M.Startprice,   /* starting price */
       M.Bottomprice,  /* bottom price */
       M.Endprice,     /* ending price */
       M.Avgprice      /* average price */
FROM Ticker AS T (Sym, Td, Pr)
    MATCH_RECOGNIZE (
        PARTITION BY T.Sym
        ORDER BY T.Td
        MEASURES MATCH_NUMBER() AS Matchno,
        A.Pr AS Startp,
        LAST (B.Pr) AS Bottomp,
        LAST (C.Pr) AS Endp,
        AVG (U.Pr) AS Avgp
        ONE ROW PER MATCH
        AFTER MATCH SKIP PAST LAST ROW
        PATTERN (A B+ C+)
```

```
SUBSET U = (A, B, C)
DEFINE /* A defaults to True, matches any row */
      B AS B.Pr < PREV (B.Pr),
      C AS C.Pr > PREV (C.Pr)
) AS M (Cym, Mno, Startprice, Bottomprice, Endprice, Avgprice)
```

The preceding example uses ONE ROW PER MATCH, so the columns of the row pattern output table are the row pattern partitioning column Sym, followed by the row pattern measure columns Matchno, Startp, Bottomp, Endp, and Avgp, for a total of six columns. The row pattern output declared column list renames these columns to Cym, Mno, Startprice, Bottomprice, Endprice, and Avgprice, respectively. Note that the SELECT list uses the column names of the row pattern output declared column list, since those are the final names of the columns.

In all, the preceding example has the originally defined column names and their renames as shown in Table 10, “Original and renamed column names”.

Table 10 — Original and renamed column names

row pattern input table		row pattern output table	
original column name	renamed column name	original column name	renamed column name
Symbol	Sym	Sym	Cym
Tradeday	Td		
Price	Pr		
		Matchno	Mno
		Startp	Startprice
		Bottomp	Bottomprice
		Endp	Endprice
		Avgp	Avgprice

4.18 Prohibited nesting

4.18.1 Introduction to prohibited nesting

The following kinds of nesting are prohibited by SQL/RPR:

- 1) Nesting one row pattern recognition within another is prohibited.
- 2) Outer references in MEASURES or DEFINE are prohibited. This means that a row pattern recognition cannot reference any table in an outer query block except the row pattern input table. (The row pattern input table is referenced using row pattern variables, not range variables defined outside the MATCH_RECOGNIZE.)
- 3) Subqueries in MEASURES or DEFINE cannot reference row pattern variables.

4.18 Prohibited nesting

4) Row pattern recognition cannot be used in recursive queries.

These restrictions are illustrated in the following Subclauses.

4.18.2 Row pattern recognition nested within another row pattern recognition

Nesting one row pattern recognition within another is prohibited. For example, the following is a syntax error:

```
SELECT ...
FROM Ticker
    MATCH_RECOGNIZE (
        ...
        DEFINE A AS EXISTS ( SELECT *
                             FROM Stock2
                             MATCH_RECOGNIZE (...) )
    )
```

A possible workaround is to relegate the nested row pattern recognition to a view or SQL-invoked function.

4.18.3 Outer references within a row pattern recognition query

Here is an example of row pattern recognition nested within an outer query. Note the underlined range variables T:

```
SELECT ( SELECT M.Avg_Price
          FROM Ticker
              MATCH_RECOGNIZE (
                  ORDER BY Tradeday
                  MEASURES AVG (T.Price) AS Avg_Price
                  PATTERN (T+)
                  DEFINE T AS T.Price >= AVG (T.Price)
              ) AS M
        )
FROM Toast AS T
```

In this example, T is both the range variable for Toast in the outer query, and also a row pattern variable in the scalar subquery.

SQL uses static scoping rules. This means that a range variable declared in an inner scope occludes a range variable of the same name declared in an outer scope. In the preceding example, there are two range variables named T. The row pattern variable T is declared in the PATTERN clause and visible in the DEFINE and MEASURES clauses, occluding the range variable T of the outer query block. Therefore, the scalar subquery (the row pattern recognition query) is not correlated with the outer query, the overall result will have one row for each row of Toast, and all rows will be identical.

This example is permitted because there are no outer references in the MATCH_RECOGNIZE. However, while legal, having multiple range variables with the same name can be confusing, so this example might be better written by changing one of the range variables. For example, changing the row pattern variable from T to X gives the equivalent query:

```
SELECT ( SELECT M.Avg_Price
          FROM Ticker
              MATCH_RECOGNIZE (
                  ORDER BY Tradeday
```

```

        MEASURES AVG (X.Price) AS Avg_Price
        PATTERN (X+)
        DEFINE X AS X.Price >= AVG (X.Price)
    ) AS M
)
FROM Toast AS I

```

On the other hand, the following is a syntax error:

```

SELECT ( SELECT M.Avg_Price
        FROM Ticker
        MATCH_RECOGNIZE (
            ORDER BY Tradeday
            MEASURES AVG (X.Price) AS Avg_Price
            PATTERN (X+)
            DEFINE X AS I.Price >= AVG (X.Price)
        ) AS M
    )
FROM Toast AS I

```

In the preceding example, the column reference T.Price in the DEFINE clause is an outer reference to the range variable T defined in the outer block; therefore, this example is a syntax error.

It may be possible to work around this limitation by placing the row pattern recognition in an SQL-invoked routine, passing as arguments the values that are prohibited as outer references.

4.18.4 Conventional query nested within row pattern recognition query

A subquery can be nested in an expression in MEASURES or DEFINE. Subqueries are permitted if they do not perform row pattern recognition themselves, and if they do not reference the row pattern variables of the outer query. Here is an example of the latter (note underlined A):

```

SELECT Firstday
FROM Ticker
    MATCH_RECOGNIZE (
        ORDER BY Tradeday
        MEASURES A.Tradeday AS Firstday
        PATTERN (A B+)
        DEFINE A AS A.Price > 100,
                B AS B.Price <
                (SELECT AVG (S.Price)
                 FROM Ticker S
                 WHERE S.Tradeday BETWEEN
                     A.Tradeday - INTERVAL '1' YEAR
                     AND A.Tradeday )
    )

```

In this example, the definition of B involves a subquery that is correlated with the row pattern variable A (note underlining). This is a syntax error, since subqueries of row pattern matching cannot reference row pattern variables.

It may be possible to work around this limitation by placing the correlated subquery in an SQL-invoked routine, passing as arguments the values that are prohibited as outer references.

4.18.5 Recursion

Row pattern matching is prohibited in recursive queries. For example, the following is a syntax error:

```
CREATE RECURSIVE VIEW Problem (Kolo, Xoro) AS
  SELECT Kolo, Xoro
  FROM T
  UNION
  SELECT Kolo + 1, Xoro
  FROM Problem
    MATCH_RECOGNIZE (
      ORDER BY Kolo
      MEASURES MATCH_NUMBER ( ) AS Xoro
      ALL ROWS PER MATCH
      PATTERN (A+)
      DEFINE A AS A.Xoro > PREV (A.Xoro)
    )
```

4.18.6 Concatenated row pattern recognition

Note that it is not prohibited to feed the output of one row pattern recognition into the input of another, as in this example:

```
SELECT ...
FROM ( SELECT *
      FROM Ticker
        MATCH_RECOGNIZE (...) )
      MATCH_RECOGNIZE (...)
```

In this example, the first MATCH_RECOGNIZE is in a derived table, which then provides the input to the second MATCH_RECOGNIZE.

5 Expressions in MEASURES and DEFINE

5.1 Introduction to the use of expressions in MEASURES and DEFINE

Scalar expression syntax as defined in [ISO/IEC 9075-2](#) is available in row pattern matching. This provides such familiar capabilities as arithmetic and aggregates. Note though that scalar expressions have special semantics in MEASURES and DEFINE; this is the subject of this Subclause.

In addition, SQL/RPR provides the following scalar expressions that are unique to row pattern matching:

- The COUNT aggregate has special syntax and semantics to count rows that are mapped to a row pattern variable by the current row pattern match.
- Row pattern navigation operations, using the functions PREV, NEXT, FIRST, and LAST. Row pattern navigation operations are discussed in [Subclause 5.6, “Row pattern navigation operations”](#).
- The MATCH_NUMBER function, which returns the sequential number of a row pattern match within its row pattern partition, discussed in [Subclause 5.8, “MATCH_NUMBER function”](#).
- The CLASSIFIER function, which returns the name of the primary row pattern variable to which a row is mapped, discussed in [Subclause 5.9, “CLASSIFIER function”](#).

Expressions in MEASURES and DEFINE clauses have the same syntax and semantics, with the following exceptions:

- 1) DEFINE clause only supports running semantics; MEASURES defaults to running semantics, but also supports final semantics. This distinction is discussed in [Subclause 5.3, “Running vs. final semantics”](#).
- 2) In DEFINE, the CLASSIFIER function does not return the classifier of rows after the current row, whereas in MEASURES, the CLASSIFIER function does return the classifier of rows after the current row. (This is only an issue when CLASSIFIER function is nested within the NEXT row pattern navigation operation; see [Subclause 5.9, “CLASSIFIER function”](#).)

5.2 Row pattern column references

A column reference is a column name qualified by an explicit or implicit range variable, such as

A.Price

A column name with no qualifier, such as Price, is implicitly qualified by the universal row pattern variable, which references the set of all rows in a match.

Column references may in general be nested within other syntactic elements, notably aggregates and subqueries. (However, nesting in row pattern matching is subject to limitations described in [Subclause 4.18, “Prohibited nesting”](#), for the FROM clause and [Subclause 6.17, “Prohibited nesting”](#), for the WINDOW clause.)

A column reference that is qualified by an explicit or implicit row pattern variable is called a row pattern column reference. Row pattern column references are classified as follows:

- Nested within an aggregate, such as SUM: aggregated row pattern column reference.

5.2 Row pattern column references

- Nested within a row pattern navigation operation (PREV, NEXT, FIRST, and LAST): a navigated row pattern column reference.
- Otherwise: an ordinary row pattern column reference.

All row pattern column references in an aggregate or row pattern navigation operation are qualified by the same row pattern variable. For example

```
PATTERN (A+ B+)
DEFINE B AS AVG (A.Price + B.Tax) > 1000
```

The preceding example is a syntax error, because A and B are two different row pattern variables. Aggregate semantics require a single set of rows; there is no way to form a single set of rows on which to evaluate A.Price + B.Tax. On the other hand, this is acceptable:

```
DEFINE B AS AVG (B.Price + B.Tax) > 1000
```

In the preceding example, all row pattern column references in the aggregate are qualified by B.

An unqualified column reference is implicitly qualified by the universal row pattern variable, which references the set of all rows in a match. For example

```
DEFINE B AS AVG (Price + B.Tax) > 1000
```

The preceding example is a syntax error, because the unqualified column reference Price is implicitly qualified by the universal row pattern variable, whereas B.Tax is explicitly qualified by B. On the other hand, this is acceptable:

```
DEFINE B AS AVG (Price + Tax) > 1000
```

In the preceding example, both Price and Tax are implicitly qualified by the universal row pattern variable.

5.3 Running vs. final semantics

Pattern matching in a sequence of rows is usually envisioned as an incremental process, with one row after another examined to see if it fits the pattern. With this incremental processing model, at any step until the complete pattern has been recognized, there is only a partial match and it is not known what rows might be added in the future, nor what variables those future rows might be mapped to. Therefore, in SQL/RPR, a row pattern column reference in the Boolean condition of a DEFINE clause has “running” semantics. This means that a row pattern variable represents the set of rows that have already been mapped to the row pattern variable, up to and including the current row, but not any future rows.

After the complete match has been established, it is possible to talk about “final” semantics. Final semantics is the same as running semantics on the last row of a successful match. Final semantics is only available in MEASURES, since in DEFINE there is uncertainty about whether a complete match has been achieved.

The keywords RUNNING and FINAL are used to indicate running or final semantics, respectively; the rules for these keywords are discussed in [Subclause 5.4, “RUNNING vs.FINAL keywords”](#).

The fundamental rule for expression evaluation in MEASURES and DEFINE is as follows:

- 1) When an expression involving a row pattern variable *RPV* is computed on a group of rows, then the set of rows *SR* that is mapped to *RPV* is used. If *SR* is empty, then COUNT is 0 and any other expression involving *RPV* is null.
- 2) When an expression requires evaluation in a single row, then the latest row of *SR* is used. If *SR* is empty, then the expression is null.

For example:

```
SELECT M.Symbol, M.Tradeday, M.Price, M.RunningAvg, M.FinalAvg
FROM TICKER
  MATCH_RECOGNIZE (
    PARTITION BY Symbol
    ORDER BY Tradeday
    MEASURES RUNNING AVG (A.Price) AS RunningAvg,
    FINAL AVG (A.Price) AS FinalAvg
    ALL ROWS PER MATCH
    PATTERN (A+)
    DEFINE A AS A.Price >= AVG (A.Price)
  ) AS M
```

Consider the ordered row pattern partition of data shown in Table 11, “Ordered row pattern partition of data”.

Table 11 — Ordered row pattern partition of data

Row	SYMBOL	TRADEDAY	PRICE
R_1	XYZ	2009-06-09	10
R_2	XYZ	2009-06-10	16
R_3	XYZ	2009-06-11	13
R_4	XYZ	2009-06-12	9

The following logic can be used to find a match.

- 1) On the first row of the row pattern partition, tentatively map row $R = R_1$ to row pattern variable A. At this point $SR = \{ R_1 \}$. To confirm whether this mapping is successful, evaluate the predicate

$A.Price \geq AVG (A.Price)$

On the left hand side, A.Price shall be evaluated in a single row, which is the last row of SR using running semantics. The last row of SR is R_1 ; therefore A.Price is 10.

On the right hand side, AVG (A.Price) is an aggregate, which is computed using the rows of SR. This average is $10/1 = 10$.

Thus the predicate asks if $10 \geq 10$. The answer is yes, so the mapping is successful. However, the pattern A+ is “greedy”, so the engine tries to match more rows if possible.

- 2) On the second row of the row pattern partition, tentatively map $R = R_2$ to row pattern variable A. At this point there are two rows mapped to A, so $SR = \{ R_1, R_2 \}$. Confirm whether the mapping is successful by evaluating the predicate

$A.Price \geq AVG (A.Price)$

On the left hand side, A.Price is evaluated in a single row, which is the last row of SR using running semantics. The last row of SR is R_2 ; therefore A.Price is 16.

On the right hand side, AVG (A.Price) is an aggregate, which is computed using the rows of SR. This average is $(10+16)/2 = 13$.

5.3 Running vs. final semantics

Thus the predicate asks if $16 \geq 13$. The answer is yes, so the mapping is successful.

- 3) On the third row of the row pattern partition, tentatively map $R = R_3$ to row pattern variable A. Now there are three rows mapped to A, so $SR = \{R_1, R_2, R_3\}$. Confirm whether the mapping is successful by evaluating the predicate

$A.Price \geq AVG(A.Price)$

On the left hand side, A.Price is evaluated in R_3 ; therefore A.Price is 13.

On the right hand side, $AVG(A.Price)$ is an aggregate, which is computed using the rows of SR. This average is $(10+16+13)/3 = 13$.

Thus the predicate asks if $13 \geq 13$. The answer is yes, so the mapping is successful.

- 4) On the fourth row of the row pattern partition, tentatively map $R = R_4$ to row pattern variable A. At this point $SR = \{R_1, R_2, R_3, R_4\}$. Confirm whether the mapping is successful by evaluating the predicate

$A.Price \geq AVG(A.Price)$

On the left hand side, A.Price is evaluated in R_4 ; therefore A.Price is 9.

On the right hand side, $AVG(A.Price)$ is an aggregate, which is computed using the rows of SR. This average is $(10+16+13+9)/4 = 12$.

Thus the predicate asks if $9 \geq 12$. The answer is no, so the mapping is not successful.

R_4 did not satisfy the definition of A, so the longest match to A+ is $\{R_1, R_2, R_3\}$. Since A+ has a greedy quantifier, this is the preferred match.

The averages computed in the DEFINE are always running averages. In MEASURES, especially with ALL ROWS PER MATCH, it is possible to distinguish final and running aggregates. Notice the use of the keywords RUNNING and FINAL in the MEASURES clause. The distinction can be observed in the result of the example, as shown in Table 12, "RUNNING and FINAL in MEASURES".

Table 12. — RUNNING and FINAL in MEASURES

SYMBOL	TRADEDAY	PRICE	RUNNINGAVG	FINALAVG
XYZ	2009-06-09	10	10	13
XYZ	2009-06-10	16	13	13
XYZ	2009-06-11	13	13	13

It is possible that the set of rows SR mapped to a row pattern variable RPV is empty. When evaluating over an empty set:

- 1) COUNT is 0.
- 2) Any other aggregate, row pattern navigation operation, or ordinary row pattern column reference is null.

For example:

```
PATTERN ( A? B+ )
DEFINE A AS A.Price > 100,
      B AS B.Price > COUNT (A.*) * 50
```

With the preceding example, consider the ordered row pattern partition of data illustrated in Table 13, “Ordered row pattern partition of data”.

Table 13 — Ordered row pattern partition of data

Row	PRICE
R_1	60
R_2	70
R_3	40

A match can be found in this data as follows:

- 1) Tentatively map $R = R_1$ to row pattern variable A. (The quantifier ? means to try first for a single match to A; if that fails, then an empty match is taken as matching A?). To see if the mapping is successful, the predicate

$A.Price > 100$

is evaluated. A.Price is 60; therefore the predicate is false and the mapping to A does not succeed.

- 2) Since the mapping to A failed, the empty match is taken as matching A?

- 3) Tentatively map $R = R_1$ to B. The predicate to check for this mapping is

$B.Price > COUNT(A.*) * 50$

No rows are mapped to A, therefore $COUNT(A.*)$ is 0. Since $B.Price = 60$ is greater than 0, the mapping is successful.

- 4) Similarly, rows R_2 and R_3 can be successfully mapped to B. Since there are no more rows, this is the complete match: no rows mapped A, $\{R_1, R_2, R_3\}$ mapped to B.

A row pattern variable can make a forward reference; that is, a reference to a row pattern variable that has not been matched yet. For example,

```
PATTERN (X+ Y+)
DEFINE X AS COUNT (Y.*) > 3,
      Y AS Y.Price > 10
```

is legal syntax. However, this example will never be matched since, at the time that a row is mapped to X, no row has yet been mapped to Y. Thus, $COUNT(Y.*)$ is 0 and can never be greater than 3. This is true even if there are four future rows that might be successfully mapped to Y. Consider the data set provided in Table 14, “Ordered row pattern partition of data”.

Table 14 — Ordered row pattern partition of data

Row	PRICE
R_1	2
R_2	11
R_3	12

Row	PRICE
R_4	13
R_5	14

Mapping $\{R_2, R_3, R_4, R_5\}$ to Y would be successful, since all four of these rows satisfy the Boolean condition defined for Y. In that case, one might think that one could map R_1 to X and have a complete successful match. However, the rules of SQL/RPR will not find this match, because, according to the pattern $X+Y+$, at least one row shall be mapped to X before any rows are mapped to Y.

5.4 RUNNING vs.FINAL keywords

RUNNING and FINAL are keywords used to indicate whether running or final semantics are desired. RUNNING and FINAL are available for use with aggregates and the row pattern navigation operations FIRST and LAST.

Aggregates, FIRST, and LAST can occur in the following places in a row pattern matching query:

- 1) In the DEFINE clause. When processing the DEFINE clause, the engine is still in the midst of recognizing a match; therefore, the only supported semantics is running.
- 2) In the MEASURES clause. When processing the MEASURES clause, the engine has finished recognizing a match; therefore, it becomes possible to consider final semantics. There are two subcases:
 - a) If ONE ROW PER MATCH is specified, or if row pattern matching is done in a window, then the engine is conceptually positioned on the last row of the match, and there is no real difference between running vs. final semantics.
 - b) If ALL ROWS PER MATCH is specified, then the row pattern output table will have one row for each row of the match. In this circumstance, the user may wish to see both running and final values, so SQL/RPR provides the RUNNING and FINAL keywords to support that distinction.

Based on this analysis, SQL/RPR specifies the following:

- 1) In MEASURES, the keywords RUNNING and FINAL may be used to indicate the desired semantics for an aggregate, FIRST, or LAST. The keyword is written before the operator, for example, RUNNING COUNT (A.*) or FINAL SUM (B.Price).
- 2) In both MEASURES and DEFINE, the default is RUNNING.
- 3) In DEFINE, FINAL is not permitted; RUNNING may be used for added clarity if desired.
- 4) In MEASURES with ONE ROW PER MATCH or in windows, all aggregates, FIRST, and LAST are computed after the last row of the match has been recognized, so that the default RUNNING semantics is actually no different from FINAL semantics. The user may prefer to think of expressions defaulting to FINAL in these cases. Alternatively, the user may choose to write FINAL for added clarity.
- 5) Ordinary column references always have running semantics. (To get final semantics in MEASURES, use the FINAL LAST row pattern navigation operation instead of an ordinary column reference.)

5.5 Aggregates

Aggregates (COUNT, SUM, AVG, etc.) may be used in both the MEASURES and DEFINE clauses. When used in row pattern matching, aggregates operate on a set of rows that are mapped to a particular row pattern variable, using either running or final semantics. For example:

```
MEASURES SUM (A.Price)          AS RunningSumOverA,  
          FINAL SUM(A.Price) AS FinalSumOverA  
ALL ROWS PER MATCH
```

In this example, A is a row pattern variable. The first row pattern measure, RunningSumOverA, does not specify either RUNNING or FINAL, so it defaults to RUNNING. This means that it is computed as the sum of Price in those rows that are mapped to A by the current match, up to and including the current row. The second row pattern measure, FinalSumOverA, computes the sum of Price over all rows that are mapped to A by the current match, including rows that may be later than the current row. Final aggregates are only available in MEASURES, not in DEFINE.

An unqualified column reference contained in an aggregate is implicitly qualified by the universal row pattern variable, which references all rows of the current row pattern match. For example:

```
SUM (Price)
```

computes the running sum of Price over all rows of the current row pattern match.

All column references contained in an aggregate shall be qualified by the same row pattern variable. For example:

```
SUM (Price + A.Tax)
```

is a syntax error, because Price is implicitly qualified by the universal row pattern variable, whereas A.Tax is explicitly qualified by A.

The COUNT aggregate has special syntax for row pattern matching, so that COUNT(A.*) may be specified. COUNT(A.*) is the number of rows that are mapped to the row pattern variable A by the current row pattern match, using running or final semantics as appropriate. As for COUNT(*), the * is implicitly qualified by the universal row pattern variable, so that COUNT(*) is the number of rows in the current row pattern match, with running or final semantics as appropriate.

5.6 Row pattern navigation operations

5.6.1 The four operations

There are four functions — PREV, NEXT, FIRST, and LAST — that enable navigation within the row pattern by either physical or logical offsets.

5.6.2 PREV and NEXT

The PREV function may be used to access columns of the previous row of a row pattern variable. If there is no previous row, the null value is returned. For example:

```
DEFINE A AS PREV (A.Price) > 100
```

5.6 Row pattern navigation operations

The preceding example says that a row R_n can be mapped to A if the preceding row R_{n-1} has a price greater than 100. If the preceding row does not exist (i.e., R_n is the first row of a row pattern partition), then $PREV(A.Price)$ is null, so the Boolean condition is not True, and therefore the first row cannot be mapped to A.

The $PREV$ function can accept an optional non-negative integer argument indicating the physical offset to the previous rows. Thus:

- $PREV (A.Price, 0)$ is equivalent to $A.Price$
- $PREV (A.Price, 1)$ is equivalent to $PREV (A.Price)$. (Note: 1 is the default offset.)
- $PREV (A.Price, 2)$ is the value of Price in the row two rows prior to the row denoted by A with running semantics. (If no row is mapped to A, or if there is no row two rows prior, then $PREV (A.Price, 2)$ is null.)

The offset is required to be a run-time constant (literal, embedded variable, and the like, but not a column or a subquery). There is an exception if the value of the offset is negative or null.

The $NEXT$ function may be used to reference rows in the forward direction in the row pattern partition using a physical offset. The syntax is the same as for $PREV$, except for the name of the function. For example,

```
DEFINE A AS NEXT (A.Price) > 100
```

The preceding example looks forward one row in the row pattern partition. Note that SQL/RPR does not support aggregates that look past the current row during $DEFINE$, because of the difficulty of predicting what row will be mapped to what row pattern variable in the future. The $NEXT$ function does not violate this principle, since it navigates to “future” rows on the basis of a physical offset, which does not require knowing the future mapping of rows.

For example, to find an isolated row that is more than twice the average of the two rows before and two rows after it: using $NEXT$, this can be expressed:

```
PATTERN ( X )
DEFINE X AS X.Price > 2 * ( PREV (X.Price, 2)
                        + PREV (X.Price, 1)
                        + NEXT (X.Price, 1)
                        + NEXT (X.Price, 2) ) / 4
```

This query can also be expressed:

```
PATTERN ( {- Y Y -} X {- Y Z -} )
SUBSET W = (Y, Z)
DEFINE Z AS X.Price > 2 * AVG (W.Price)
```

The second formulation (without $NEXT$) requires the use of exclusion syntax using $\{- \}$ and the non-intuitive definition of Z in terms of row pattern variables X and W. The row X is never really defined at all, though that is the only row of interest. The first formulation (using $NEXT$) avoids these issues.

Note that the row in which $PREV$ or $NEXT$ is evaluated is not necessarily mapped to the row pattern variable in the argument. For example, in the first formulation of the example, $PREV (X.Price, 2)$ is evaluated in a row that is not even part of the match. The purpose of the row pattern variable is to identify the row from which to offset, not the row that is ultimately reached. This point is discussed further in [Subclause 5.6.4, “Nesting FIRST and LAST within PREV or NEXT”](#).

$PREV$ and $NEXT$ may be used with more than one column reference; for example:

```
DEFINE A AS PREV (A.Price + A.Tax) < 100
```

When using a complex expression as the first argument of PREV or NEXT, all qualifiers shall be the same row pattern variable (in this example, A).

The first argument of PREV and NEXT shall have at least one column reference or CLASSIFIER function. For example, this is a syntax error:

PREV (1)

The preceding example is a syntax error because there is no row pattern column reference or CLASSIFIER function. Without a column reference or CLASSIFIER function, there is no way to determine the row that is the starting point for offsetting. (The use of CLASSIFIER function within PREV or NEXT is discussed in Subclause 5.9, “CLASSIFIER function”.)

PREV and NEXT always have running semantics; the keywords RUNNING and FINAL cannot be used with PREV or NEXT. To obtain final semantics, use, e.g., PREV (FINAL LAST (A.Price)) as explained in Subclause 5.6.4, “Nesting FIRST and LAST within PREV or NEXT”.

5.6.3 FIRST and LAST

FIRST returns the value of an expression evaluated in the first row of the group of rows mapped to a row pattern variable. For example:

FIRST (A.Price)

The preceding example evaluates A.Price in the first row that is mapped to A. If there is no row mapped to A, then the value is null.

Similarly, LAST returns the value of an expression evaluated in the last row of the group of rows mapped to a row pattern variable. For example:

LAST (A.Price)

The preceding example evaluates A.Price in the last row that is mapped to A (null if there is no such row).

The FIRST and LAST operators can accept an optional non-negative integer argument indicating a logical offset within the set of rows mapped to the row pattern variable. For example:

FIRST (A.Price, 1)

evaluates Price in the second row that is mapped to A. Consider the data set and mappings shown in Table 15, “Example data set and mappings for FIRST and LAST”.

Table 15 — Example data set and mappings for FIRST and LAST

Row	PRICE	mapping
R ₁	10	→ A
R ₂	20	→ B
R ₃	30	→ A
R ₄	40	→ C
R ₅	50	→ A

5.6 Row pattern navigation operations

Thus:

- FIRST (A.Price) = FIRST (A.Price, 0) = LAST (A.Price, 2) = 10
- FIRST (A.Price, 1) = LAST (A.Price, 1) = 30
- FIRST (A.Price, 2) = LAST (A.Price, 0) = LAST (A.Price) = 50
- FIRST (A.Price, 3) is null, as is LAST (A.Price, 3)

Note that the offset is a logical offset, moving within the set of rows $\{R_1, R_3, R_5\}$ that are mapped to the row pattern variable A. It is not a physical offset, as with PREV or NEXT.

The optional integer argument is required to be a run-time constant (literal, embedded variable, and the like, but not a column or subquery). There is an exception if the value of the offset is negative or null.

The first argument of FIRST or LAST shall have at least one row pattern column reference or CLASSIFIER function. (The use of CLASSIFIER function within FIRST or LAST is discussed in [Subclause 5.9, "CLASSIFIER function"](#).) Thus FIRST(1) is a syntax error.

The first argument of FIRST or LAST is permitted to have more than one row pattern column reference, in which case all qualifiers shall be the same row pattern variable. For example:

```
FIRST (A.Price + B.Tax)
```

is a syntax error, but

```
FIRST (A.Price + A.Tax)
```

is acceptable.

FIRST and LAST support both running and final semantics. RUNNING is the default, and the only supported option in DEFINE. Final semantics may be accessed in the MEASURES by using the keyword FINAL (particularly with ALL ROWS PER MATCH), as in:

```
MEASURES FINAL LAST (A.Price) AS FinalPrice
ALL ROWS PER MATCH
```

5.6.4 Nesting FIRST and LAST within PREV or NEXT

FIRST and LAST provide navigation within the set of rows already mapped to a particular row pattern variable; PREV and NEXT provide navigation using a physical offset from a particular row. These kinds of navigation may be combined by nesting FIRST or LAST within PREV or NEXT. This permits expressions such as the following:

```
PREV (LAST (A.Price + A.Tax, 1), 3)
```

In this example, A is required to be a row pattern variable. It is required to have a row pattern column reference or CLASSIFIER function, and all row pattern variables in the compound operator shall be equivalent (A, in this example). The use of CLASSIFIER function nested within row pattern navigation operations is discussed in [Subclause 5.9, "CLASSIFIER function"](#).

This compound operator is evaluated as follows:

- 1) The inner operator, LAST, operates on the set of rows that are mapped to the row pattern variable A. In this set, find the row that is "the last minus 1". (If there is no such row, the result is null.)
- 2) The outer operator, PREV, starts from the row found in step 1 and backs up 3 rows in the row pattern partition. (If there is no such row, the result is null.)

- 3) Let R be an implementation-dependent range variable that references the row found by step 2. In the expression $A.Price + A.Tax$, replace every occurrence of the row pattern variable A with R . The resulting expression $R.Price + R.Tax$ is evaluated and determines the value of the compound navigation operation.

For example, consider the data set and mappings illustrated in Table 16, “Data set and mappings for nesting example”.

Table 16 — Data set and mappings for nesting example

Row	PRICE	TAX	mapping
R_1	10	1	
R_2	20	2	→ A
R_3	30	3	→ B
R_4	40	4	→ A
R_5	50	5	→ C
R_6	60	6	→ A

To evaluate

$PREV (LAST (A.Price + A.Tax, 1), 3)$

the following steps may be used:

- 1) The set of rows mapped to A is $\{R_2, R_4, R_6\}$. LAST operates on this set, offsetting from the end to arrive at row R_4 .
- 2) PREV performs a physical offset, 3 rows prior to R_4 , arriving at R_1 .
- 3) Let R be a range variable pointing at R_1 . $R.Price + R.Tax$ is evaluated, giving $10+1 = 11$.

Note that this nesting is not defined as a typical evaluation of nested functions. The inner operator LAST does not actually evaluate the expression $A.Price + A.Tax$; it merely uses this expression to designate a row pattern variable (A) and then navigate within the rows mapped to that variable. The outer operator PREV performs a further physical navigation on rows. The expression $A.Price + A.Tax$ is not actually evaluated as such, since the row that is eventually reached is not necessarily mapped to the row pattern variable A. In this example, R_1 is not mapped to any row pattern variable.

5.7 Ordinary row pattern column references reconsidered

An ordinary row pattern column reference is one that is neither aggregated nor navigated. For example:

$A.Price$

Subclause 5.4, “RUNNING vs.FINAL keywords”, stated that ordinary row pattern column references always have running semantics. This means:

5.7 Ordinary row pattern column references reconsidered

- 1) In DEFINE, an ordinary column reference references the last row that is mapped to the row pattern variable, up to and including the current row. If there is no such row, then the value is null.
- 2) In MEASURES, there are two subcases:
 - a) If ALL ROWS PER MATCH is specified, then there is also a notion of current row, and the semantics are the same as in DEFINE.
 - b) If ONE ROW PER MATCH is specified, then conceptually the engine is positioned on the last row of the match. An ordinary column reference references the last row that is mapped to the row pattern variable, anywhere in the pattern. If the variable is not mapped to any row, then the value is null.

These semantics are the same as the LAST operator, with the implicit RUNNING default. Consequently, an ordinary column reference such as:

X.Price

is equivalent to:

RUNNING LAST (X.Price)

5.8 MATCH_NUMBER function

Matches within a row pattern partition are numbered sequentially starting with 1 in the order they are found. Note that match numbering starts over again at 1 in each row pattern partition, since there is no inherent ordering between row pattern partitions. MATCH_NUMBER () is a nullary function that returns an exact numeric value with scale 0 (zero) whose value is the sequential number of the match within the row pattern partition.

All previous examples of MATCH_NUMBER () have shown it used in MEASURES. It is also possible to use MATCH_NUMBER () in DEFINE, where it can be used to define conditions that depend upon the match number. For example:

```
PATTERN ( (A+ | B+) )
DEFINE A AS ( MOD (MATCH_NUMBER (), 2) = 1)
              AND A.Price > PREV (A.Price) ),
       B AS ( MOD (MATCH_NUMBER (), 2) = 0)
              AND B.Price < PREV (B.Price) )
```

The condition for A can only be true on odd-numbered matches, and the condition for B can only be true on even-numbered matches. Thus, the matches will alternate between A+ and B+.

MATCH_NUMBER () is not permitted except in MEASURES and DEFINE. For example, the following is a syntax error:

```
SELECT MATCH_NUMBER ()
FROM Ticker
```

5.9 CLASSIFIER function

The classifier of a row is the primary row pattern variable to which the row is mapped by a row pattern match. The classifier of a row that is not mapped by a row pattern match is null.

The CLASSIFIER function returns a character string whose value is the classifier of a row. The CLASSIFIER function has one optional argument, a row pattern variable, defaulting to the universal row pattern variable.

The simplest usage of CLASSIFIER is with no argument, as seen in the example in Subclause 4.4, “Example of ALL ROWS PER MATCH”:

```
MEASURES ...
    CLASSIFIER () AS Classy, ...
    ALL ROWS PER MATCH
```

In this example, the CLASSIFIER () function returns the classifier of the current row, which is assigned as the value of the row pattern measure Classy. Subclause 4.4, “Example of ALL ROWS PER MATCH”, shows the result of the example query on the sample data, illustrating how the CLASSIFIER function returns the classifier of each row of a row pattern match.

The classifier of the starting row of an empty match is the null value. This can be seen in the example in Subclause 4.12.2, “Handling empty matches”, of a query using ALL ROWS PER MATCH SHOW EMPTY MATCHES.

Optionally, the CLASSIFIER function may take a single argument, a row pattern variable *RPV*; the default is the universal row pattern variable. The argument is used to specify a set of rows using running semantics, namely, the set of rows up to and including the current row that are mapped to *RPV*.

The argument will typically be a union row pattern variable; the value returned tells which primary row pattern variable among the components of the union row pattern variable to which a row was mapped. For example:

```
MEASURES CLASSIFIER (AB) AS AorB
...
PATTERN (A | B | C)+
SUBSET AB = (A, B)
```

In this example, AB is a union row pattern variable. The value of the row pattern measure AorB is either A or B, whichever is the classifier of the last row that mapped to A or B. If no row mapped to A or B, the value is null.

The CLASSIFIER function may be used in an aggregate. For example:

```
ORDER BY Tradeday
MEASURES ARRAY_AGG (CLASSIFIER () ORDER BY Tradeday)
    AS ClassifierArray
ONE ROW PER MATCH
```

In the preceding example, one row is created for each row pattern match *RPM*, with a single row pattern measure, which is an array of character strings. The elements of the array are the classifiers of the rows in *RPM*, with one array element for each row of the row pattern input table that is mapped by *RPM*. Note that the array is ordered using the same ordering as the row pattern partition. This technique can be used to obtain a value reflecting the precise pattern that was detected, while using ONE ROW PER MATCH instead of ALL ROWS PER MATCH.

The CLASSIFIER function may be nested within a row pattern navigation operation. For example:

```
PREV (CLASSIFIER ())
```

The preceding example returns the classifier of the preceding row. This might be used in DEFINE so that the definition of one row pattern variable might depend on the classifier of a previously matched row. For example:

```
PATTERN ( (A | B) C )
```

ISO/IEC 19075-5:2021(E)

5.9 CLASSIFIER function

```
DEFINE A AS ...,
      B AS ...,
      C AS CASE
        WHEN PREV (CLASSIFIER ()) = 'A' AND Price > 100
        THEN 1
        WHEN PREV (CLASSIFIER ()) = 'B' AND Price < 100
        THEN 1
        ELSE 0
      END = 1
```

In this example, the first row might be mapped to either A or B. The definition of C can test PREV (CLASSIFIER ()) to learn to which row pattern variable the first row mapped. If the first row mapped to A, then C will be true if Price > 100; if the first row mapped to B, then C will be true if Price < 100.

This particular example would be more easily written as:

```
PATTERN ( (A AC | B BC) )
DEFINE A AS ...,
      B AS ...,
      AC AS Price > 100,
      BC AS Price < 100
```

However, the example does illustrate a general technique that might be useful for more complex patterns that want to inquire about the mapping chosen in earlier rows.

When nesting CLASSIFIER within NEXT, there is an important distinction between MEASURES and DEFINE. Consider, for example:

```
NEXT (CLASSIFIER ())
```

which asks for the classifier of the next row. This expression, used in DEFINE, will return null, because the next row has not been mapped yet when considering how to map the current row. Used in MEASURES, the entire pattern has been mapped, and the preceding example is able to return the classifier of the row after the current row. (If the current row is the last row that is mapped, the result is null.)

Here is an example of the CLASSIFIER function nested in a compound row pattern navigation operation:

```
NEXT (RUNNING LAST (CLASSIFIER (U), 2) 3)
```

This example would be evaluated as follows:

- 1) Find the set of rows mapped to U. In DEFINE, this can only be rows up to and including the current row; in MEASURES, this can be any rows mapped to U in the completed match.
- 2) RUNNING LAST restricts to the set of rows that map to U up to and including the current row. In DEFINE, this is no change from step 1. In MEASURES with ONE ROW PER MATCH, the position is on the last row, so this also is no change. In MEASURES with ALL ROWS PER MATCH, this may result in discarding some of the rows mapped to U.
- 3) In the set of rows remaining after step 2, find the row that is offset 2 from the end. This requires at least three rows remaining after step 2; if there are not that many, then the result is null. (This is the functionality of LAST.)
- 4) Now move forward in the row pattern partition three rows. If there are not enough rows in the row pattern partition, the result is null. (This is the functionality of NEXT.)
- 5) Finally, find the primary row pattern variable to which the row is mapped; this is the result. If the row is not mapped, the result is null. (This is the functionality of CLASSIFIER.)

The explicit or implicit argument of CLASSIFIER is a row pattern variable. This row pattern variable is used in the same fashion as the qualifier of a column reference in the argument of a row pattern navigation operation. For example:

```
NEXT (CLASSIFIER () || Name)
```

is a permissible expression, since CLASSIFIER () and Name both reference the universal row pattern variable. On the other hand,

```
NEXT (CLASSIFIER (A) || Name)
```

is not permissible, because the CLASSIFIER function references the row pattern variable A, whereas the column reference Name references the universal row pattern variable. Similarly:

```
NEXT (CLASSIFIER () || A.Name)
```

is not permissible.

The same rule applies to the argument of an aggregate: all row pattern variables referenced explicitly or implicitly shall be the same. Thus the following is a syntax error:

```
ARRAY_AGG (CLASSIFIER () || A.Name)
```

The CLASSIFIER function is not permitted except in MEASURES and DEFINE. For example, the following is a syntax error:

```
SELECT CLASSIFIER ()  
FROM Ticker
```

6 Row pattern recognition: WINDOW clause

6.1 Introduction to the WINDOW clause

Feature R020, “Row pattern recognition: WINDOW clause” of ISO/IEC 9075-2 enhances the WINDOW clause to include row pattern matching. In ISO/IEC 9075-2, a window is a data structure defined on the result of a <table expression> (the FROM...WHERE...GROUP BY...HAVING... clauses), producing a derived table. This data structure does three things:

- 1) Partitions the rows of the derived table according to zero or more columns.
- 2) Within each window partition, orders the rows of the derived table according to zero or more expressions.
- 3) For each row *R* in a window partition, defines a window frame, which is a subset of the ordered window partition. The endpoints of the window frame may be the beginning or end of the window partition, or may be defined relative to the current row using either a physical offset (row count), a logical offset (a value added to or subtracted from the only sort column), or a group count (number of groups, defined as peers under the ordering).

Using R020, row pattern recognition may be used to further reduce the window frame. Row pattern recognition is applicable only to window frames that start at the current row *R*. The window frame resulting from step 3 will be called the “full window frame” of *R*, and the window frame after reduction by pattern matching will be called the “reduced window frame” of *R*. When performing row pattern recognition in a window, the window partition serves as the row pattern partition and the window ordering serves as the row pattern ordering.

6.2 Example of row pattern recognition in a window

The example from Subclause 4.3, “Example of ONE ROW PER MATCH”, is adapted to use row pattern matching in the WINDOW clause below:

```
SELECT T.Symbol,      /* ticker symbol */
       T.Tradeday,    /* trade day */
       T.Price,       /* price on day of trading */
       Classy OVER W, /* classifier */
       Startp OVER W, /* starting price */
       Bottomp OVER W, /* bottom price */
       Endp OVER W,   /* ending price */
       Avgp OVER W    /* average price */
FROM Ticker AS T
WINDOW W AS
( PARTITION BY Symbol
  ORDER BY Tradeday
  MEASURES A.Price AS Startp,
           LAST (B.Price) AS Bottomp,
           LAST (C.Price) AS Endp,
           AVG (U.Price) AS Avgp
  ROWS
    BETWEEN CURRENT ROW
```

```

        AND UNBOUNDED FOLLOWING
    AFTER MATCH SKIP PAST LAST ROW
    INITIAL
    PATTERN (A B+ C+)
    SUBSET U = (A, B, C)
    DEFINE /* A defaults to True, matches any row */
        B AS B.Price < PREV (B.Price),
        C AS C.Price > PREV (C.Price)
    )

```

In the preceding example, the principle syntactic elements are presented on separate lines. In this example:

- The SELECT list contains window functions over the window W defined in the WINDOW clause. Note the window functions Classy OVER W, Startp OVER W, Bottomp OVER W, Endp OVER W, and Avgp OVER W. These are examples of row pattern measure functions, which are window functions used to access the row pattern measures defined in the MEASURES clause of the window definition.
- FROM introduces a conventional FROM clause. This example has a single table, Ticker, which is assigned the range variable T. Since there is no WHERE, GROUP BY, or HAVING clause, the result of the FROM clause is the row pattern input table in this example.
- WINDOW W declares the window name W.
- PARTITION BY specifies how to partition the row pattern input table. The PARTITION BY clause is a list of columns of the row pattern input table. This clause is optional; if omitted, there are no row pattern partitioning columns, and the entire row pattern input table constitutes a single row pattern partition.
- ORDER BY specifies how to order the rows within row pattern partitions of the row pattern input table. The ORDER BY clause is a list of columns of the row pattern input table. This clause is optional; if omitted, the order of rows is completely non-deterministic. However, since non-deterministic ordering will defeat the purpose of most row pattern recognition, the ORDER BY clause will usually be specified.
- MEASURES specifies row pattern measures, whose values are calculated by evaluating expressions related to the match. The values of row pattern measures are accessed using row pattern measure functions, as illustrated in the SELECT list.
- ROWS specifies the unit to use in defining the full window frame. The other choices, RANGE and GROUPS, are not permitted with row pattern matching in windows.
- BETWEEN CURRENT ROW AND UNBOUNDED FOLLOWING is one way to specify the full window frame. In this example, for any row R in a row pattern partition P, the full window frame consists of R and all rows that follow R in the row pattern partition P. The full window frame is subsequently reduced to just the rows constituting a pattern match.
- AFTER MATCH SKIP clause specifies where to resume looking for the next row pattern match after successfully finding a match. In this example, AFTER MATCH SKIP PAST LAST ROW specifies that pattern matching will resume after the last row of a successful match. When a row is skipped, its reduced window frame is empty.
- INITIAL specifies that the pattern always matches starting at the first row of the full window frame. If there is no such pattern match, the reduced window frame is empty. The alternative to INITIAL is SEEK, which specifies to seek the first row pattern match in the full window frame; if there is none, the reduced window frame is empty.
- PATTERN specifies the row pattern that is sought in the row pattern input table. A row pattern is a regular expression using primary row pattern variables. In this example, the row pattern has three primary row pattern variables (A, B, and C).

6.2 Example of row pattern recognition in a window

- SUBSET defines the union row pattern variables. In this example, U is defined as the union of the primary row pattern variables A, B, and C. The SUBSET clause is optional.
- DEFINE specifies the Boolean condition that defines a primary row pattern variable; a row shall satisfy the condition in order to be mapped to a particular primary row pattern variable. If a primary row pattern variable is not defined in the DEFINE clause, then its definition defaults to a condition that is always true, meaning that the primary row pattern variable can match any row.

The result of the preceding query on the sample row pattern partition is given in Table 17, “Window example query results”.

Table 17 — Window example query results

SYM BOL	TRADEDAY	PRICE	CLASSY	STARTP	BOTTOMP	ENDP	AVGP
XYZ	2009-06-08	50					
XYZ	2009-06-09	60	A	60	35	45	45.8
XYZ	2009-06-10	49					
XYZ	2009-06-11	40					
XYZ	2009-06-12	35					
XYZ	2009-06-15	45					
XYZ	2009-06-16	45					
XYZ	2009-06-17	45	A	45	43	70	51.4
XYZ	2009-06-18	43					
XYZ	2009-06-19	47					
XYZ	2009-06-22	52					
XYZ	2009-06-23	70					
XYZ	2009-06-24	60					

6.3 Summary of the syntax

6.3.1 Syntax components

The syntax of row pattern recognition in windows is summarized in Table 18, “Row pattern recognition in windows — syntax summary”.

Table 18 — Row pattern recognition in windows — syntax summary

Clause of window definition	Optional?	Notes	Cross-reference
Existing window name	yes	no default	Subclause 6.6, “Windows defined on windows”
PARTITION BY	yes	if omitted, the row pattern input table constitutes one row pattern partition	Subclause 6.7, “PARTITION BY”
ORDER BY	yes	if omitted, there is a non-deterministic ordering in each row pattern partition	Subclause 6.8, “ORDER BY”
MEASURES	yes	no default	Subclause 6.9, “MEASURES”
ROWS, RANGE, GROUPS	no	only ROWS is permitted with row pattern recognition	Subclause 6.10.2, “ROWS BETWEEN CURRENT ROW AND ”
BETWEEN CURRENT ROW AND ...	no	BETWEEN CURRENT ROW AND is required with row pattern recognition	Subclause 6.10.2, “ROWS BETWEEN CURRENT ROW AND ”
EXCLUDE NO OTHERS	yes	EXCLUDE NO OTHERS is the default; other EXCLUDE options are forbidden with row pattern matching	Subclause 6.10.3, “EXCLUDE NO OTHERS”
AFTER MATCH SKIP	yes	default is AFTER MATCH SKIP PAST LAST ROW	Subclause 6.11, “AFTER MATCH SKIP”
INITIAL, SEEK	yes	default is INITIAL	Subclause 6.12, “INITIAL vs SEEK”
PATTERN	no	same as MATCH_RECOGNIZE	Subclause 6.13, “PATTERN”
SUBSET	yes	same as MATCH_RECOGNIZE	Subclause 6.14, “SUBSET”
DEFINE	no	same as MATCH_RECOGNIZE	Subclause 6.15, “DEFINE”

6.3.2 Syntactic comparison to windows without row pattern recognition

Note the following differences between windows with and without row pattern matching:

- 1) Windows with row pattern matching always have the PATTERN and DEFINE clauses, and optionally may also have the MEASURES, AFTER MATCH SKIP, INITIAL, SEEK, and SUBSET clauses. Windows without row pattern matching have none of these clauses.

6.3 Summary of the syntax

- 2) Windows with row pattern matching always start with the current row, and shall specify ROWS (not RANGE or GROUPS).
- 3) The only permitted EXCLUDE clause with row pattern matching is EXCLUDE NO OTHERS, which is the default.

6.3.3 Syntactic comparison to MATCH_RECOGNIZE

The syntax for row pattern recognition in a window differs from MATCH_RECOGNIZE in the following respects:

- 1) Row pattern recognition in windows includes the window syntax of conventional windows, with some constraints described in Subclause 6.3.2, “Syntactic comparison to windows without row pattern recognition”.
- 2) Range variables declared in the FROM clause are visible in the PARTITION BY and ORDER BY clause of a window, unlike MATCH_RECOGNIZE. See Subclause 6.7, “PARTITION BY”, and Subclause 6.8, “ORDER BY”.
- 3) The ORDER BY clause may use scalar value expressions, not just columns. See Subclause 6.8, “ORDER BY”.
- 4) The options ONE ROW PER MATCH and ALL ROWS PER MATCH are not applicable to windows, and cannot be specified. (Row pattern recognition in windows is closer in spirit to ONE ROW PER MATCH, though it also has some similarity to ALL ROWS PER MATCH WITH UNMATCHED ROWS.)
- 5) Row pattern recognition in a window has a choice between INITIAL and SEEK.
- 6) The MATCH_NUMBER function is not supported.
- 7) Row pattern measures are not columns in the result of a window; instead, row pattern measures are referenced using OVER, like a window function.

6.4 Row pattern input table

The row pattern input table for row pattern recognition in a WINDOW clause is the result of the FROM, WHERE, GROUP BY, and HAVING clauses that precede the WINDOW clause.

The example in Subclause 6.2, “Example of row pattern recognition in a window”, does not have WHERE, GROUP BY, or HAVING clauses, so the row pattern input table in that example is the result of the FROM clause, that is, the table Ticker.

6.5 Row pattern variables and other range variables

There are two sets of range variables in a window with row pattern recognition:

- 1) The range variables declared in the FROM clause. (In the example in Subclause 6.2, “Example of row pattern recognition in a window”, T is such a range variable.)

These range variables may be used in the PARTITION BY and ORDER BY clauses, and of course in the SELECT list, WHERE clause, etc.

- 2) The row pattern variables declared in the PATTERN and SUBSET clauses. (In the example in Subclause 6.2, “Example of row pattern recognition in a window”, A, B, C, and U are row pattern variables.)

Row pattern variables may be referenced in the MEASURES, DEFINE, and SUBSET clauses. They cannot be used in the SELECT list.

Note that the two sets of range variables are declared in different clauses and have mutually exclusive scope. (Since they are walled off in mutually exclusive scopes, it is permitted to use the same range variables in each scope, though that is a confusing possibility that it is probably best to avoid.)

For example, the following is a syntax error:

```
SELECT Runlength OVER W
FROM Ticker T
WINDOW W AS ( PARTITION BY Symbol
               ORDER BY Tradeday
               MEASURES COUNT (T.*) AS Runlength
               ROWS BETWEEN CURRENT ROW
                     AND UNBOUNDED FOLLOWING
               PATTERN (A*)
               DEFINE A AS T.Price > PREV (T.Price) )
```

There are three syntax errors in this example. The first is COUNT (T.*). T is a range variable defined in the FROM clause and cannot be referenced in MEASURES. Instead of T, the variable to use here is A, since A is declared in the PATTERN. Similarly, in the DEFINE, the two instances of T are errors.

Any column names to be referenced in either the MEASURES or DEFINE are unique across the entire FROM clause, because the range variables in the FROM clause are not available to disambiguate in MEASURES or DEFINE. The workaround is to rename column names in the FROM clause as necessary to remove ambiguities.

For example, suppose both Emp and Dept have a column called Name. In the following query:

```
SELECT Aname OVER W, Bname OVER W
FROM Emp AS E, Dept AS D
WHERE E.Deptno = D.Deptno
WINDOW W AS ( PARTITION BY E.Deptno
               ORDER BY E.Empno
               MEASURES A.Name AS Aname,
                     B.Name AS Bname
               ROWS BETWEEN CURRENT ROW
                     AND UNBOUNDED FOLLOWING
               PATTERN (A B)
               DEFINE B AS B.Salary > A.Salary )
```

The expressions A.Name and B.Name in the MEASURES clause are ambiguous, because they could refer to either Emp.Name or Dept.Name. The solution is to rename at least one of them in the FROM clause, like this:

```
SELECT Aname OVER W, Bname OVER W
FROM ( SELECT Name AS Ename, Deptno, Salary
      FROM Emp ) AS E, Dept AS D
WHERE E.Deptno = D.Deptno
WINDOW W AS ( PARTITION BY E.Deptno
               ORDER BY E.Empno
               MEASURES A.Name AS Aname,
                     B.Name AS Bname
               ROWS BETWEEN CURRENT ROW
                     AND UNBOUNDED FOLLOWING
```

6.5 Row pattern variables and other range variables

```
PATTERN (A B)
DEFINE B AS B.Salary > A.Salary )
```

Note that row pattern variables are not available in the SELECT list. The following example is a syntax error:

```
SELECT A.Price
FROM Ticker AS T
WINDOW W AS ( PARTITION BY Symbol
               ORDER BY Tradeday
               MEASURES COUNT (A.*) AS Runlength
               ROWS BETWEEN CURRENT ROW
                   AND UNBOUNDED FOLLOWING
               PATTERN (A*)
               DEFINE A AS A.Price > PREV (A.Price) )
```

In this example, A is a row pattern variable, which makes it visible in MEASURES and DEFINE. A is not visible in the SELECT list. There is no loss of expressive power; any expression of row pattern variables can be placed in MEASURES and then referenced by its measure name, like this:

```
SELECT Lasta OVER W
FROM Ticker AS T
WINDOW W AS ( PARTITION BY Symbol
               ORDER BY Tradeday
               MEASURES LAST (A.Price) AS Lasta
               ROWS BETWEEN CURRENT ROW
                   AND UNBOUNDED FOLLOWING
               PATTERN (A*)
               DEFINE A AS A.Price > PREV (A.Price) )
```

6.6 Windows defined on windows

ISO/IEC 9075-2 allows one window to be defined on another window by referencing an existing window name. For example:

```
FROM Ticker AS T
WINDOW W1 AS ( PARTITION BY Symbol ),
              W2 AS ( W1 ORDER BY Tradeday ),
              W3 AS ( W2 ROWS BETWEEN CURRENT ROW
                   AND UNBOUNDED FOLLOWING )
```

Here window W2 inherits its partitioning from W1, and W3 inherits its partitioning and ordering from W2.

As an example of this capability using row pattern recognition:

```
SELECT Lastprice OVER W3
FROM Ticker AS T
WINDOW W1 AS ( PARTITION BY Symbol ),
              W2 AS ( W1 ORDER BY Tradeday ),
              W3 AS ( W2 MEASURES LAST(A.Price) AS Lastprice
                   ROWS BETWEEN CURRENT ROW
                       AND UNBOUNDED FOLLOWING
                   PATTERN (A+)
                   DEFINE A AS A.Price > PREV (A.Price) )
```

It is not possible to further subdivide the window definitions. For example, it is not permitted to put ROWS BETWEEN CURRENT ROW AND UNBOUNDED FOLLOWING in one window definition, and then inherit that in another window definition that adds the row pattern recognition features.

6.7 PARTITION BY

PARTITION BY is almost the same in windows and in MATCH_RECOGNIZE. The one difference is that range variables declared in the FROM clause are available in the PARTITION BY of a window, but not in MATCH_RECOGNIZE. Note that a row pattern partition is the same thing as a window partition when performing row pattern recognition in a window.

6.8 ORDER BY

ORDER BY is almost the same in windows and in MATCH_RECOGNIZE. The differences are:

- 1) Range variables declared in the FROM clause are available in the ORDER BY of a window, but not in MATCH_RECOGNIZE.
- 2) General scalar value expressions may be used in the ORDER BY of a window, but only column references may be used in the ORDER BY in MATCH_RECOGNIZE.

6.9 MEASURES

Row pattern measures in a window definition differ from row pattern measures in MATCH_RECOGNIZE as follows:

- 1) The MATCH_NUMBER function is not supported in windows.
- 2) Row pattern measures are referenced as window functions in the SELECT list using OVER, not as column references.
- 3) There is no real distinction between running and final semantics. The RUNNING and FINAL keywords may be used with aggregates, FIRST, and LAST, but the semantics is the same no matter which keyword is used. Row pattern measures are computed positioned on the last row of the match, where running and final semantics are identical.

6.10 Full window frame and reduced window frame

6.10.1 Introduction to window framing

A window associates with each row R a set of rows, called the window frame of R . The definition of the window frame is essentially a subtractive process:

- 1) At the outset, there is the entire window partition that contains R .

6.10 Full window frame and reduced window frame

- 2) Next, zero or more rows are removed from the window partition, based on their position relative to R in the ordering of rows of the window partition. The criterion at this stage is called the “window frame extent”.
- 3) Next, zero or more rows are removed, based on peer relationships to R , using the EXCLUDE clause.

The three steps above are used for all windows. The result is called the “full window frame”. When row pattern recognition is used, the window partition is also the row pattern partition, and there is one more step:

- 4) A match to the row pattern is sought within the full window frame; the rows that are mapped by this match (if any) constitute the “reduced window frame”. If there is no match, the reduced window frame is empty. (Skipped rows can also cause an empty reduced window frame; see [Subclause 6.16, “Empty matches and empty reduced window frames”](#).)

6.10.2 ROWS BETWEEN CURRENT ROW AND

When performing row pattern recognition in a window, only two options are allowed for specifying the window frame extent:

- ROWS BETWEEN CURRENT ROW AND UNBOUNDED FOLLOWING: this option specifies that the full window frame consists of the set of rows from the current row through the end of the row pattern partition.
- ROWS BETWEEN CURRENT ROW AND *offset* FOLLOWING: this option specifies that the full window frame extends from the current row through some positive offset, which shall be a positive integer, and specifies the number of rows after the current row. For example, ROWS BETWEEN CURRENT ROW AND 1 FOLLOWING specifies a full window frame with 2 rows, the current row and the one after it.

6.10.3 EXCLUDE NO OTHERS

The window EXCLUDE clause has four possibilities:

- 1) EXCLUDE CURRENT ROW: this is not permitted with row pattern recognition, since the design is that the full window frame always begins with the current row.
- 2) EXCLUDE GROUP: also not permitted with row pattern recognition, because this would exclude the current row, plus any ties under the window ordering.
- 3) EXCLUDE TIES: not permitted with row pattern recognition, because this could create a hole in the full window frame, which is contrary to the spirit of row pattern recognition.
- 4) EXCLUDE NO OTHERS: permitted with row pattern recognition. This is the default.

Thus the only permitted option with row pattern matching is the default, EXCLUDE NO OTHERS.

6.11 AFTER MATCH SKIP

The options for AFTER MATCH SKIP are the same as in MATCH_RECOGNIZE; see [Subclause 4.13, “AFTER MATCH SKIP”](#), for details.

As in MATCH_RECOGNIZE, it is a run-time error to skip to a non-existent row, or to skip to the first row of a match.

Since only one row pattern match per full window frame is sought, the semantics of AFTER MATCH SKIP in a window are as follows. Windows are processed in the sort order within a row pattern partition. If a row *R* is skipped as a consequence of a row pattern match in a full window frame prior to *R*, then the reduced window frame of *R* is set to empty, without attempting any row pattern match for *R*. This is illustrated in an example in [Subclause 6.16, “Empty matches and empty reduced window frames”](#).

6.12 INITIAL vs SEEK

If a row *R* has been skipped by a prior row *PR*, then the reduced window frame of *R* is empty.

If *R* has not been skipped, then a row pattern match is attempted in the full window frame of *R*. INITIAL and SEEK are two options that determine where to look for a match within the full window frame:

- 1) INITIAL is used to look for a match whose first row is *R*.
- 2) SEEK is used to permit a search for the first match anywhere from *R* through the end of the full window frame.

In either case, the reduced window frame comprises the rows that are mapped by the match; if there is no match, then the reduced window frame is empty. For a worked example, see [Subclause 6.16, “Empty matches and empty reduced window frames”](#).

The keyword INITIAL or SEEK is placed as a modifier before the PATTERN. The default is INITIAL.

6.13 PATTERN

This clause is precisely the same as in MATCH_RECOGNIZE, except that the anchors (^ and \$) are not permitted with row pattern matching in windows. See [Subclause 4.14, “PATTERN”](#), for details.

6.14 SUBSET

This clause is precisely the same as in MATCH_RECOGNIZE. See [Subclause 4.15, “SUBSET”](#), for details.

6.15 DEFINE

This clause is precisely the same as in MATCH_RECOGNIZE. See [Subclause 4.16, “DEFINE”](#), for details.

6.16 Empty matches and empty reduced window frames

An empty match will cause the reduced window frame to be empty. Empty reduced window frames can also arise if there is no match at all, as in these circumstances:

6.16 Empty matches and empty reduced window frames

- 1) AFTER MATCH SKIP on a prior row has caused the current row to be skipped, so no match is attempted.
- 2) The query specifies or implies INITIAL and there is no match starting at the current row.
- 3) The query specifies SEEK but there is no match anywhere between the current row and the end of the full window frame.

So there are two ways to get an empty reduced window frame: by finding an empty match, or by not finding a match at all.

The semantics for row pattern measures of empty reduced window frames are shown in [Table 19, “Results for empty match and no match”](#).

Table 19 — Results for empty match and no match

Measure	empty match	no match
CLASSIFIER ()	null	null
COUNT	0	null
other aggregates (e.g., SUM, AVG, etc.	null	null
row pattern navigation operations (e.g., PREV, NEXT, FIRST, LAST	null	null
ordinary column references	null	null

Thus COUNT (*) may be used to distinguish an empty match from no match at all. If an empty match is found, then COUNT (*) as a row pattern measure will be 0; if there is no match at all, then COUNT (*) as a row pattern measure will be null.

Note the following subtlety: If the query specifies COUNT (*) as a non-measure window function, then the count over an empty window frame is 0 in any case. It is only when COUNT (*) is used as a row pattern measure that it can be used to distinguish an empty match from no match at all. For example:

```
SELECT S, D,
       Kount OVER W AS "Measure",
       COUNT (*) OVER W AS "Window Function"
FROM T
WINDOW W AS ( ORDER BY S
               MEASURES COUNT (*) AS Kount
               ROWS BETWEEN CURRENT ROW
                     AND UNBOUNDED FOLLOWING
               AFTER MATCH SKIP PAST LAST ROW
               INITIAL PATTERN (A*)
               DEFINE A AS A.D = 'yes' )
```

Consider the following data, shown in the first two columns, with the other two columns of output shown in the next two columns, and the internal information (skip indicator, whether a match was found, and the reduced window frame) in the right three columns, as shown in [Table 20, “Computation of matches and window function results”](#).