OPENRULES®

Open Source Business Decision Management System

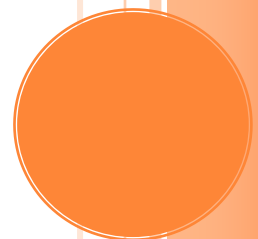## Release 6.2.6

User Manual

Table of Contents

INTRODUCTION

OpenRules® was developed in 2003 by OpenRules, Inc. as an open source Business Rules Management System (BRMS) and since then has become one of the most popular BRMS on the market. Over these years OpenRules® has been naturally transformed in a Business Decision Management System (BDMS) with proven records of delivering and maintaining reliable decision support software. OpenRules® is a winner of several software awards for innovation and is used worldwide by multi-billion dollar corporations, major banks, insurers, health care providers, government agencies, online stores, universities, and many other institutions.

Brief History

From the very beginning, OpenRules® was oriented to subject matter experts (business analysts) allowing them to work in concert with software developers to create, maintain, and efficiently execute business rules housed in enterprise-class rules repositories. OpenRules® avoided the introduction of yet another "rule language" as well as another proprietary rules management GUI. Instead, OpenRules® relied on commonly used tools such as MS Excel, Google Docs and Eclipse integrated with the standard Java. This approach enabled OpenRules® users to create and maintain inter-related decision tables directly in Excel. Initially each rules table included several additional rows, in which a software developer could place Java snippets to specify the exact semantics of rule conditions and actions.

In March of 2008, OpenRules® Release 5 introduced Rule Templates. Templates allowed a business analyst to create hundreds and thousands of business rules based on a small number of templates supported by software developers. Rule templates minimized the use of Java snippets and hid them from business users. Rule templates were a significant step in minimizing rule repositories and clearly separating the roles of business analysts and software specialists in maintaining the rules.

In March of 2011 OpenRules® introduced Release 6, which finally moved control over business logic to business users. OpenRules® 6 effectively removed any Java coding from rules representation allowing business analysts themselves to specify their decisions and supporting decision tables directly and completely in Excel. Business users can also create business glossaries and test cases in Excel tables. They may then test the accuracy of execute their decisions without the need for any coding at all.

Once a decision has been tested it can be easily incorporated into any Java or .NET environment. This process may involve IT specialists but only to integrate the business glossary with a specific business object model. The business logic remains the complete prerogative of subject matter experts.

OpenRules® Components

OpenRules® offers the following decision management components:

- Rule Repository for management of enterprise-level decision rules
- Rule Engine for execution of decisions and different business rules
- Rule Dialog for building rules-based Web questionnaires
- Rule Learner for rules discovery and predictive analytics
- Rule Solver for solving constraint satisfaction and optimization problems
- Finite State Machines for event processing and "connecting the dots".

Integration of these components with executable decisions has effectively converted OpenRules® from a B**R**MS to a B**D**MS, Business Decision Management System, oriented to "decision-centric" application development.

OpenRules, Inc. is a professional open source company that provides software, product documentation and technical support and other services that are highly praised by our customers. You may start learning about product with the document "Getting Started" which describes how to install OpenRules® and includes simple examples. Then you may look at a more complex example in the tutorial "Calculating Tax Return". This user manual covers the core OpenRules® concepts in greater depth. Additional OpenRules® components are described in separate user manuals: see Rule Learner, Rule Solver, and Rule Dialog.

Document Conventions

-   The regular Century Schoolbook font is used for descriptive information.
-   *The italic Century Schoolbook font is used for notes and fragments clarifying the text.*
-   `The Courier New font is used for code examples.`

CORE CONCEPTS

OpenRules® is a BDMS, Business Decision Management System, oriented to "decision-centric" application development. OpenRules® utilizes the well-established spreadsheet concepts of workbooks, worksheets, and tables to build enterprise-level rule repositories. Each OpenRules® workbook is comprised of one or more worksheets that can be used to separate information by types or categories.

To create and edit rules and other tables presented in Excel-files you can use any standard spreadsheet editor such as:

- MS Excel™
- OpenOffice™
- Google Docs™

Google Docs™ is especially useful for collaborative rules management.

OpenRules® supports different types of spreadsheets that are defined by their keywords. Here is the list of OpenRules® tables along with brief description of each:

Table Type (Keyword)	Comment
Decision	Defines a decision that may consist of multiple sub-decisions associated with different decision tables
DecisionTable or **DT** or **DecisionTableSingleHit** or **RuleFamily**	This is a single-hit decision table that uses multiple conditions on different defined on variables to reach conclusions about the decision variables
Glossary	For each decision variable used in the decision tables, the glossary defines related business concepts, as well as related implementation attributes and their possible domain
DecisionObject	Associates business concepts specified in the glossary with concrete objects defined outside the decision (i.e. as Java objects or Excel Data

	tables)
Rules	Defines a decision table that includes Java snippets that specify custom logic for conditions and actions. Read <u>more</u>. Some Rules tables may refer to <u>templates</u> that hide those Java snippets.
Datatype	Defines a new data type directly in Excel that can be used for testing
Data	Creates an array of test objects
Variable	Creates one test object
Environment	This table defines the structure of a rules repository by listing all included workbooks, XML files, and Java packages
Method	Defines expressions using snippets of Java code and known decision variables and objects
DecisionTable1 or **DT1** or **DecisionTableMultiHit**	A multi-hit decision table that allows rule overrides
DecisionTable2 or **DT2**	A multi-hit decision table that like DecisionTable2 executes all rules in top-down order but results of the execution of previous rules may affect the conditions of rules that follow
Layout	A special table type used by OpenRules® <u>Forms</u> and OpenRules® <u>Dialog</u>

The following section will provide a detailed description of these concepts.

SPREADSHEET ORGANIZATION AND MANAGEMENT

OpenRules® uses Excel spreadsheets to represent and maintain business rules, web forms, and other information that can be organized using a tabular format. Excel is the best tool to handle different tables and is a popular and widely used tool among business analysts.

Workbooks, Worksheets, and Tables

OpenRules® utilizes commonly used concepts of workbooks and worksheets. These can be represented and maintained in multiple Excel files. Each OpenRules® workbook is comprised of one or more worksheets that can be used to separate information by categories. Each worksheet, in turn, is comprised of one or more tables. Decision tables are the most typical OpenRules® tables and

are used to represent business rules. Workbooks can include tables of different types, each of which can support a different underlying logic.

How OpenRules® Tables Are Recognized

OpenRules® recognizes the tables inside Excel files using the following parsing algorithm.

1. The OpenRules® parser splits spreadsheets into "parsed tables". Each logical table should be separated by at least one empty row or column at the start of the table. Table parsing is performed from left to right and from top to bottom. The first non-empty cell (i.e. cell with some text in it) that does not belong to a previously parsed table becomes the top-left corner of a new parsed table.

2. The parser determines the width/height of the table using non-empty cells as its clues. Merged cells are important and are considered as one cell. If the top-left cell of a table starts with a predefined keyword (see the table below), then such a table is parsed into an OpenRules® table.

3. All other "tables," i.e. those that do not begin with a keyword are ignored and may contain any information.

The list of all keywords was described <u>above</u>. OpenRules® can be extended with more table types, each with their own keyword.

While not reflected in the table recognition algorithm, it is good practice to use a black background with a white foreground for the very first row. All cells in this row should be *merged*, so that the first row explicitly specifies the table width. We call this row the "*table signature*". The text inside this row (consisting of one or more merged cells) is the table signature that starts with a keyword. The information after the keyword usually contains a unique table name and additional information that depends on the table type.

If you want to put a table title before the signature row, use an empty row between the title and the first row of the actual table. Do not forget to put an

empty row after the last table row. Here are examples of some typical tables recognized by OpenRules®.

OpenRules® table with 3 columns and 2 rows:

Keyword <some text>		
Something	Something	Something
Something	Something	Something

OpenRules® table with 3 columns and still 2 rows:

Keyword	Something	Something
Something	Something	Something
Something	Something	Something

OpenRules® table with 3 columns and 3 rows (empty initial cells are acceptable):

Keyword <some text>		
Something	Something	
	Something	Something
		Something

OpenRules® table with 3 columns and 2 rows (the empty 3rd row ends the table):

Keyword <some text>		
Something	Something	Something
Something	Something	Something
Something	Something	Something

OpenRules® table with 2 columns and 2 rows (the empty cell in the 3rd column of the title row results in the 4th columns being ignored. This also points out the importance of merging cells in the title row):

Keyword	Something		Something
Something	Something	Something	Something
Something	Something	Something	Something

OpenRules® will not recognize this table (there is no empty row before the signature row):

Table Title		
Keyword <some text>		
Something	Something	
	Something	Something
		Something

Fonts and coloring schema are a matter of the table designer's taste. The designer has at his/her disposal the entire presentation power of Excel (including comments) to make the OpenRules® tables more self-explanatory.

OpenRules® Rule Table Example

Here is an example of a worksheet with two rules tables:

This workbook is comprised of three worksheets:

1. Worksheet "Decision Tables" - includes rule tables
2. Worksheet "Launcher" - includes a method that defines an order and conditions under which rules will be executed

3. Worksheet "Environment" - defines the structure of a rules repository by listing all included workbooks, XML files, and Java packages (if any).

The worksheet "Decision Tables" is comprised of two rule tables "defineGreeting" and "defineSalutation". Rule tables are a traditional way to represent business decision tables. Rule tables are decision tables that usually describe combinations of conditions and actions that should be taken when all of the conditions have been satisfied. In the table "defineGreeting", the action "Set Greeting" will be executed when an "hour," passed to this table as a parameter, is between "Hour From" and "Hour To". In the table "defineSalutation", an action "Set Salutation" will be executed when a customer's Gender *and* Marital Status correspond to the proper row.

These tables start with signature rows that are determined by a keyword in the first cell of the table. A table signature in general has the following format:

```
Keyword return-type table-name(type1 par1, type2 par2,..)
```

where table-name is a one-word function name and return-type, type1, and type 2 are types defined in the current OpenRules® configuration. For example, type may be any basic Java type such as int, double, Date, or Strin*g*.

The rule tables above are recognized by the keyword "Rules". All of the columns have been merged into a single cell in the signature rows. Merging cells B3, C3, and D3 specifies that table "defineGreeting" has 3 columns. A table includes all those rows under its signature that contain non empty cells: in the example above, an empty row 12 indicates the end of the table "defineGreeting".

Limitation. Avoid merging rule rows in the very first column (or in the very first row for horizontal tables) - it may lead to invalid logic.

Business and Technical Views

OpenRules® tables such as "Rules" and "Data" may have two views:

[1] Business View

[2] Technical View

These two views are implemented using Excel's outline buttons [1] and [2] at the top left corner of every worksheet - see the figure below. This figure represents a business view - no technical details about the implementation are provided. For example, from this view it is hard to tell for sure what greeting will be generated at 11 o'clock: "Good Morning" or "Good Afternoon"? If you push the Technical View button [2] (or the button "+" on the left), you will see the hidden rows with the technical details of this rules table:

	A	B	C	D
3		Rules void **defineGreeting**(int hour, Response response)		
4		C1	C2	A1
5		min <= hour	hour <= max	response.map.put("greeting", greeting);
6		int min	int max	String greeting
7		**Hour From**	**Hour To**	**Set Greeting**
8		0	11	Good Morning
9		12	17	Good Afternoon
10		18	22	Good Evening
11		23	24	Good Night

The technical view opens hidden rows 4-6 that contain the implementation details. In particular, you can see that both "Hour From" and "Hour To" are included in the definition of the time intervals. Different types of tables have different technical views.

Note. Using Rules Templates you may completely split business and technical information between different Excel tables. Decisions do not use technical views at all because they do not require any coding and rely on predefined templates.

DECISION MODELING AND EXECUTION

OpenRules® methodological approach allows business analysts to develop their executable decisions with underlying decision tables without (or only with a limited) help from software developers. You may become familiar with the major

decision modeling concepts from simple examples provided in the document "Getting Started" and several associated tutorials. First we will consider the simple implementation options for decision modeling, and later on we will describe more advanced OpenRules® concepts.

Starting with Decision

From the OpenRules® perspective a decision contains:

- a set of decision variables that can take specific values from domains of values
- a set of decision rules (frequently expressed as decision tables) that specify relationships between decision variables.

Some decision variables are known (decision input) and some of them are unknown (decision output). A decision may consist of other decisions (sub-decisions). To execute a decision means to assign values to unknown decision variables in such a way that satisfies the decision rules. This approach corresponds to the OMG standard known as "DMN".

OpenRules® applies a top-down approach to decision modeling. This means that we usually start with the definition of a Decision and not with rules or data. Only then we will define decision tables, a glossary, and then data. Here is an example of a Decision:

Decision DeterminePatientTherapy	
Decisions	**Execute Decision Tables**
Define Medication	DefineMedication
Define Creatinine Clearance	CalculateCreatinineClearance
Define Dosing	DefineDosing
Check Drug Interaction	WarnAboutDrugInteraction

Here the decision "DeterminePatientTherapy" consists of four sub-decisions:

- "Define Medication" that is implemented using a decision table "DefineMedication"
- "Define Creatinine Clearance" that is implemented using a decision table "DefineCreatinineClearance"
- "Define Dosing" that is implemented using a decision table "DefineDosing"
- "Check Drug Interaction" that is implemented using a decision table "WarnAboutDrugInteraction".

The table "Decision" has two columns "Decisions" and "Execute Decision Tables" (those are not keywords and you can use any other titles for these columns). The first column contains the names of all our sub-decisions - here we can use any combinations of words as decision names. The second column contains exact names of decision tables that implement these sub-decisions. The decision table names cannot contain spaces or special characters (except for "underscore").

OpenRules® allows you to use multiple (embedded) tables of the type "Decision" to define more complex decisions. For example, a top-level decision, that defines the main decision variable, may be defined through several sub-decisions about related variables:

Decision DecisionMain	
Decisions	**Execute Rules / Sub-Decisions**
Define Variable 1	DecisionTableVariable1
Define Variable 2	DecisionTableVariable21
Define Variable 2	DecisionTableVariable22
Define Variable 3	**DecisionVariable3**
Define Variable 4	DecisionTableVariable4

In order to Define Variable 2 it is necessary to execute two decision tables. Some decisions, like "Define Variable 3", may require their own separate sub-decisions such as described in the following table:

Decision DecisionVariable3	
Decisions	**Execute Rules**
Define Variable 3.1	DecisionTableVariable31
Define Variable 3.2	DecisionTableVariable32
Define Variable 3.3	DecisionTableVariable33

These tables can be kept in different files and can be considered as building blocks for your decisions. This top-down approach with Decision Tables and dependencies between them allows you to represent even quite complex decision logic in an intuitive, easy to understand way.

Some decisions may have a more complex structure than the just described sequence of sub-decisions. You can even use conditions inside decision tables. For example, consider a situation when the first sub-decision validates your data and a second sub-decision executes complex calculations but only if the preceding validation was successful. Here is an example of such a decision from the tax calculation tutorial:

Decision Apply1040EZ			
Condition		ActionPrint	ActionExecute
1040EZ Eligible		**Decisions**	**Execute**
		Validate	ValidateTaxReturn
Is	TRUE	Calculate	DetermineTaxReturn
Is	FALSE	Do Not Calculate	

Since this table "Decision Apply1040EZ" uses an optional column "Condition", we have to add a second row. The keywords "Condition", "ActionPrint", and "ActionExecute" are defined in the standard OpenRules® template "DecisionTemplate" – see the configuration file "DecisionTemplates.xls" in the folder "openrules.config". This table uses a decision variable "1040EZ Eligible" that is defined by the first (unconditional) sub-decision "Validate". We assume that the decision "ValidateTaxReturn" should set this decision variable to TRUE or FALSE. Then the second sub-decision "Calculate" will be executed only when

"1040EZ Eligible" is TRUE. When it is FALSE, this decision, "Apply1040EZ", will simply print "Do Not Calculate". In our example the reason will be printed by the decision table "ValidateTaxReturn".

Note. You may use many conditions of the type "Condition" defined on different decision variables. Similarly, you may use an optional condition "ConditionAny" which instead of decision variables can use any formulas defined on any known objects.   It is also possible to add custom actions using an optional action "ActionAny" – see "DecisionTemplates.xls" in the folder "openrules.config".

When you have completed defining all decision and sub-decisions, you may define decision tables.

Defining Decision Tables

OpenRules® decision modeling approach utilizes the classical decision tables that were in the heart of OpenRules® BDMS from its introduction in 2003. OpenRules® uses the keyword "**Rules**" to represent different types of classical decision tables. Rules tables rely on Java snippets to specify execution logic of multiple conditions and actions. In 2011 OpenRules® version 6 introduced a special type of decision tables with the keyword "**DecisionTable**" (or "**DT**") that do not need Java snippets and rely on the predefined business logic for its conditions and conclusions defined on already known decision variables. For example, let's consider a very simple decision "DetermineCustomerGreeting":

Decision DetermineCustomerGreeting	
Decisions	**Execute Rules**
Define Greeting Word	DefineGreeting
Define Salutation Word	DefineSalutation

It refers to two decision tables. Here is an example of the first decision table:

DecisionTable DefineGreeting					
Condition		Condition		Conclusion	
Current Hour		**Current Hour**		**Greeting**	
>=	0	<=	11	Is	Good Morning
>=	11	<=	17	Is	Good Afternoon
>=	17	<=	22	Is	Good Evening
>=	22	<=	24	Is	Good Night

Its first row contains a keyword "DecisionTable" and a unique name (no spaces allowed). The second row uses keywords "Condition" and "Conclusion" to specify the types of the decision table columns. The third row contains decision variables expressed in plain English (spaces are allowed but the variable names should be unique).

The columns of a decision table define conditions and conclusions using different operators and operands appropriate to the decision variable specified in the column headings. The rows of a decision table specify multiple rules. For instance, in the above decision table "DefineGreeting" the second rule can be read as:

"IF Current Hour is more than or equal to 11 AND Current Hour is less than or equal to 17 THEN Greeting is Good Afternoon".

Similarly, we may define the second decision table "DefineSalutation" that determines a salutation word (it uses the keyword "DT" that is a synonym for "DecisionTable"):

DT DefineSalutation					
Condition		Condition		Conclusion	
Gender		**Marital Status**		**Salutation**	
Is	Male			Is	Mr.
Is	Female	Is	Married	Is	Mrs.
Is	Female	Is	Single	Is	Ms.

If some cells in the rule conditions are empty, it is assumed that this condition is satisfied. A decision table may have no conditions but it always should contain at least one conclusion.

Decision Table Execution Logic

OpenRules® executes all rules within DecisionTable in a *top-down* order. When all conditions inside one rule (row) are satisfied the proper conclusion(s) from the same row will be executed, and all other rules will be ignored.

Note. OpenRules® decision tables can be used to implement a methodological approach described in the book "The Decision Model". It relies on a special type of decision tables called "Rule Families" that require that the order of rules inside a decision table should not matter. It means that to comply with the Decision Model principles, you should not rely on the default top-down rules execution order of OpenRules® decision tables. Instead, you should design your decision table (you even may use the keyword "RuleFamily" instead of "DT") in such a way that all rules are mutually exclusive and cover all possible combinations of conditions. The advantage of this approach is that when you decide to add new rules to your rule family you may place them in any rows without jeopardizing the execution logic. However, in some cases, this approach may lead to much more complex organization of rule families to compare with the standard decision tables.

AND/OR Conditions

The conditions in a decision table are always connected by a logical operator "AND". When you need to use "OR", you may add another rule (row) that is an alternative to the previous rule(s). However, some conditions may have a decision variable defined as an array, and within such array-conditions "ORs" are allowed. Consider for example the following, more complex decision table:

Here the decision variables "Customer Profile", "Customer Product", and "Offered Products" are arrays of strings. In this case, the second rule can be read as:

```
IF Customer Profile Is One Of New or Bronze or Silver
AND Customer Products Include Checking Account and
Overdraft Protection
AND Customer Products Do Not Include CD with 25 basis point
increase, Money Market Mutual Fund, and Credit Card
THEN Offered Products ARE CD with 25 basis point increase,
Money Market Mutual Fund, and Credit Card
```

Decision Table Operators

OpenRules® supports multiple ways to define operators within decision table conditions and conclusions. When you use a text form of operators you can freely use upper and lower cases and spaces. The following operators can be used inside decision table conditions:

Operator	Synonyms	Comment
Is	=, ==	When you use "=" or "==" inside Excel write'''="'' or'''=="'' to avoid confusion with Excel's own formulas
Is Not	!=, isnot, Is Not Equal To, Not, Not Equal., Not Equal To	Defines an inequality operator
>	Is More, More, Is More Than, Is Greater, Greater, Is Greater Than	For integers and real numbers, and Dates

>=	Is More Or Equal. Is More Or Equal To, Is More Than Or Equal To, Is Greater Or Equal To, Is Greater Than Or Equal To	For integers and real numbers, and Dates
<=	Is Less Or Equal, Is Less Than Or Equal To, Is Less Than Or Equal To, Is Smaller Or Equal To, Is Smaller Than Or Equal To, Is Smaller Than Or Equal To,	For integers and real numbers, and Dates
<	Is Less, Less, Is Less Than, Is Smaller, Smaller, Is Smaller Than	For integers and real numbers, and Dates
Is True		For booleans
Is False		For booleans
Is Empty		A string is considered "empty" if it is either "null" or contains only zero or more whitespaces
Contains	Contain	For strings only, e.g. "House" contains "use". The comparison is not case-sensitive
Starts With	Start with, Start	For strings only, e.g. "House" starts with "ho". The comparison is not case-sensitive
Match	Matches, Is Like, Like	Compares if the string matches a regular expression
No Match	NotMatch, Does Not Match, Not Like, Is Not Like, Different, Different From	Compares if a string does not match a regular expression
Within	Inside, Inside Interval, Interval	For integers and real numbers. The interval can be defined as: [0;9], (1;20], 5–10, between 5 and 10, more than 5 and less or equals 10 – see more
Is One Of	Is One, Is One of Many, Is Among, Among	For integer and real numbers, and for strings. Checks if a value is among elements of the domain of values listed through comma
Is Not One Of	Is not among, Not among	For integer and real numbers, and for strings. Checks if a value is NOT among elements of the domain of values listed through comma
Include	Include All	To compare two arrays. Returns true when the first array (decision variable) include all elements of the second array (value within decision table cell)

Exclude	Do Not Include, Exclude One Of	To compare an array or value with an array
Does Not Include	Include Not All	To compare two arrays. Returns true when the first array (decision variable) does not include all elements of the second array (value within decision table cell)
Intersect	Intersect With, Intersects	To compare an array with an array

If the decision variables do not have an expected type for a particular operator, the proper syntax error will be diagnosed.

The following operators can be used inside decision table conclusions:

Operator	Synonyms	Comment
Is	=, ==	Assigns one value to the conclusion decision variable. When you use "=" or "==" inside Excel write"'="" or"'=="" to avoid confusion with Excel's own formulas.
Are		Assigns one or more values listed through commas to the conclusion variable that is expected to be an array
Add		Adds one or more values listed through commas to the conclusion variable that is expected to be an array
Assign Plus	+=	Takes the conclusion decision variable, adds to it a value from the rule cell, and saves the result in the same decision variable.
Assign Minus	-=	Takes the conclusion decision variable, subtracts from it a value from the rule cell, and saves the result in the same decision variable.
Assign Multiply	*=	Takes the conclusion decision variable, multiplies it by a value from the rule cell, and saves the result in the same decision variable.
Assign Divide	/=	Takes the conclusion decision variable, divides it by a value from the rule cell, and saves the result in the same decision variable.

Using Regular Expressions in Decision Table Conditions

OpenRules® allows you to use standard <u>regular expressions</u>. Operators "Match" and "No Match" (and their synonyms from the above table) allow you to match the content of your text decision variables with custom patterns such as phone number or Social Security Number (SSN). Here is an example of a decision table that validates SSN:

DecisionTable testSSN		
Condition		Message
SSN		**Message**
No Match	\d{3}-\d{2}-\d{4}	Invalid SSN
Match	\d{3}-\d{2}-\d{4}	Valid SSN

The use of this decision table is described in the sample project "DecisionHelloJava".

Conditions and Conclusions without Operators

Sometimes the creation of special columns for operators seems unnecessary, especially for the operators "Is" and "Within". OpenRules® allows you to use a simpler format as in this decision table:

DT DefineGreeting	
If	Then
Current Hour	**Greeting**
0-11	Good Morning
11-17	Good Afternoon
17-22	Good Evening
22-24	Good Night

As you can see, instead of keywords "Condition" and "Conclusion" we use the keywords "If" and "Then" respectively. While this decision table looks much simpler in comparison with the functionally identical decision table defined

<u>above</u>, we need to make an implicit assumption that the lower and upper bounds for the intervals "0-11", "11-17", etc. are included.

Using Decision Variable Names inside Decision Table Cells

When your decision table contains too many columns it may become too wide and unmanageable. In practice large decision tables have many empty cells because not all decision variables participate in all rule conditions even if the proper columns are reserved or all rules. To make your decision table more compact, OpenRules® allows you to move a variable name from the column title to the rule cells. To do that, instead of the standard column's structure with two sub-columns

Condition	
Variable Name	
Oper	Value

you may use another column representation with 3 sub-columns:

ConditionVarOperValue		
Attribute		
Variable Name	Oper	Value

This way you may replace a wide table with many condition columns like the one below:

DecisionTable classifcationRules											
Condition		Condition			Condition		Conclusion		Conclusion	
C_OTH_EXPNS_AMT		A_ESTATE_TAX_AMT		...	Attribute		ClassifiedAs		HitRate	
>=	398	>=	10054		Oper	Value	Is	High	=	66
							Is	Low	=	63
>=	53						Is	Low	=	49
							Is	Low	=	86
							Is	Low	=	78
							Is	Other	=	56

to a much more compact table that may look as follows:

DecisionTable classifcationRules											
ConditionVarOperValue			ConditionVarOperValue			...	ConditionVarOperValue			Conclusion	Conclusion
Attribute			Attribute			...	Attribute			ClassifiedAs	HitRate
C_OTH_EXPNS_AMT	>=	398	A_ESTATE_TAX_A	>=	10054		Attribute	Oper	Value	Is	High	=	66
E_PRTSCRP_TOT_LC	<=	-6955	AGI_TPI_RATIO	<=	0.95993					Is	Low	=	63
C_OTH_EXPNS_AMT	>=	53	ORD_DIVIDENDS	<=	6617					Is	Low	=	49
TAXABLE_INC_TPI_R	<=	0.810447	TENT_TAX_AMT	>=	301630					Is	Low	=	86
DIVIDENDS_AND_INT	<=	12348	EXTNSN_PYMNT_	<=	30000					Is	Low	=	78
										Is	Other	=	56

You simply replace a column of the type "Condition" to the one that has the standard type "ConditionVarOperValue". Similarly, instead of a column of the type "Conclusion" you may use a column of the type "ConclusionVarOperValue" with 3 sub-columns that represent:
- Decision variable name
- Operator
- Value.

Using Formulas inside Decision Tables

OpenRules® allows you to use formulas in the rule cells instead of constants. The formulas usually have the following format:

$$::= (\text{expression})$$

where an "expression" can be written using standard Java expressions. Here is an example:

DecisionTable CalculateAdjustedGrossIncome	
Conclusion	
Adjusted Gross Income	
Is	::= (getReal("Wages") + getReal("Taxable Interest") + getReal("Unemployment Compensation"))

This decision table has only one conclusion that simply calculates a value for the decision variable "Adjusted Gross Income" as a sum of values for the decision variables "Wages", "Taxable Interest", and "Unemployment Compensation". This example also demonstrates how to gain access to different decision variables – you may write getReal("VARIABLE_NAME") for real decision variables. Similarly, you may use methods getInt(…), getBool(…), getDate(…), and getString(…).

You may also put your formula in a specially defined Method and then refer to this method from the decision table – observe how it is done in the following example:

Method double taxableIncome()				
return getReal("Adjusted Gross Income") - getReal("Dependent Amount");				

Decision Table CalculateTaxableIncome				
ConditionAny			Conclusion	
Condition			Taxable Income	
Is True	:= (taxableIncome() > 0)	Is	::= taxableIncome()	
Is False	:= (taxableIncome() > 0)	Is	0	

Here we defined a new method "*taxableIncome()*" that returns a real value using the standard Java type "double". Then we used this method inside both conditions and one conclusion of this decision table.

Note. Actually the formula format ::= (expression) is a shortcut for a more standard OpenRules® formula format := "" +(expression) that also can be used inside decision tables.

Direct References to Decision Variables

You may want to refer to values of some decision variables inside cells for different tables. To do that, you may simply put a dollar sign ("$") in front of the variable name. For example, in the following table

DT DefineWhomToCharge					
Condition		Condition		Conclusion	
Vendor		Provider		Charged Entity	
Is Empty	FALSE			Is	$Vendor
		Is Empty	TRUE	Is	UNKNOWN
		Is Not One Of	ABC, KLM, XYZ	Is	$Provider

the conclusion-column contains references $Vendor and $Provider to the values of decision variables Vendor and Provider. The reference $Vendor is similar to the formula ::= getString("Vendor"). You may also use similar references inside arrays. For example, to express a

condition that a Vendor should not be among providers, you may use the operator "Is Not One Of" with an array "ABC, $Vendor, XYZ".

Defining Business Glossary

While defining decision tables, we freely introduced different decision variables assuming that they are somehow defined. The business glossary is a special OpenRules® table that actually defines all decision variables. The Glossary table has the following structure:

Glossary glossary			
Variable	**Business Concept**	**Attribute**	**Domain**

The first column will simply list all of the decision variables using exactly the same names that were used inside the decision tables. The second column associates different decision variables with the business concepts to which they belong. Usually you want to keep decision variables that belong to the same business concept together and merge all rows in the column "Business Concept" that share the same concept. Here is an example of a glossary from the standard OpenRules® example "DecisionLoan":

Glossary glossary			
Variable	**Object**	**Attribute**	**Domain**
Monthly Income	Customer	monthlyIncome	0-5000000
Mortgage Holder		mortgageHolder	Yes,No
Outside Credit Score		outsideCreditScore	0-999
Loan Holder		loanHolder	Yes,No
Credit Card Balance		creditCardBalance	-1000000 - 100000000
Education Loan Balance		educationLoanBalance	-1000000 - 100000000
Internal Credit Rating		internalCreditRating	A,B,C,D,F
Internal Analyst Opinion		internalAnalystOpinion	High,Mid,Low
Income Validation Result	Request	incomeValidationResult	SUFFICIENT,UNSUFFICIENT,?
Debt Research Result		debtResearchResult	High,Mid,Low,?
Loan Qualification Result		loanQualificationResult	QUALIFIED, NOT QUALIFIED, ?
Total Income	Internal	totalIncome	0-500000
Total Debt		totalDebt	0-500000

All rows for the concepts such as "Customer" and "Request" are merged.

The third column "Attribute" contains "technical" names of the decision variables – these names will be used to connect our decision variables with attributes of objects used by the actual applications, for which a decision has been defined. The application objects could be defined in Java, in Excel tables, in XML, etc. The decision does not have to know about it: the only requirement is that the attribute names should follow the usual naming convention for identifiers in languages like Java: it basically means no spaces allowed. The last column, "Domain", is optional, but it can be useful to specify which values are allowed to be used for different decision variables. Decision variable domains can be specified using the naming convention for the intervals and domains described below. The above glossary provides a few intuitive examples of such domains. These domains can be used during the validation of a decision.

Defining Test Data

OpenRules® provides a convenient way to define test data for decisions directly in Excel without the necessity of writing any Java code. A non-technical user can define all business concepts in the Glossary table using Datatype tables. For example, here is a Datatype table for the business concept "Customer" defined above:

Datatype Customer	
String	fullName
String	SSN
int	monthlyIncome
int	monthlyDebt
String	mortgageHolder
int	outsideCreditScore
String	loanHolder
int	creditCardBalance
int	educationLoanBalance
String	internalCreditRating
String	internalAnalystOpinion

The first column defines the type of the attribute using standard Java types such as "int", "double", "Boolean", "String", or "Date". The second column contains the same attribute names that were defined in the Glossary. To create an array of objects of the type "Customer" we may use a special "Data" table like the one below:

Data Customer customers										
fullName	SSN	monthlyIncome	monthlyDebt	mortgageHol der	outsideCreditSc ore	loanHolder	creditCar dBalance	educationL oanBalanc e	internalCredi tRating	internalAnal ystOpinion
Borrower Full Name	Borrower SSN	Monthly Income	Monthly Debt	Mortgage Holder	Outside Credit Score	Loan Holder	Credit Card Balance	Education Loan Balance	Internal Credit Rating	Internal Analyst Opinion
Peter N. Johnson	157-82-5344	5000	2300	Yes	720	No	2500	0	A	Low
Mary K. Brown	056-45-8233	4300	2800	No	620	No	5654	23800	B	Low
Robert Cooper Jr.	241-56-9082	6400	2800	Yes	735	Yes	1200	0	C	Mid

This table is too wide (and difficult to read), so we could actually transpose it to a more convenient but equivalent format:

Data Customer customers				
fullName	Borrower Full Name	Peter N. Johnson	Mary K. Brown	Robert Cooper Jr.
SSN	Borrower SSN	157-82-5344	056-45-8233	241-56-9082
monthlyIncome	Monthly Income	5000	4300	6400
monthlyDebt	Monthly Debt	2300	2800	2800
mortgageHolder	Mortgage Holder	Yes	No	Yes
outsideCreditScore	Outside Credit Score	720	620	735
loanHolder	Loan Holder	No	No	Yes
creditCardBalance	Credit Card Balance	2500	5654	1200
educationLoanBalance	Education Loan Balance	0	23800	0
internalCreditRating	Internal Credit Rating	A	B	C
internalAnalystOpinion	Internal Analyst Opinion	Low	Low	Mid

Now, whenever we need to reference the first customer we can refer to him as customers[0]. Similarly, if you want to define a doubled monthly income for the second custromer, "Mary K. Brown", you may simply write

```
::= (customers[1].monthlyIncome * 2)
```

You can find many additional details about data modeling in this section.

Connecting the Decisions with Business Objects

To tell OpenRules® that we want to associate the object customers[0] with our business concept "Customer" defined in the Glossary, we need to use a special table "DecisionObject" that may look as follows:

DecisionObject decisionObjects	
Business Concept	Business Object
Customer	:= customers[0]
Request	:= loanRequests[0]
Internal	:= internal

Here we also associate other business concepts namely Request and Internal with the proper business objects – see how they are defined in the standard example "DecisionLoan".

The above table connects a decision with test data defined by business users directly in Excel. This allows the decision to be tested. However, after the decision is tested, it will be integrated into a real application that may use objects defined in Java, in XML, or in a database, etc. For example, if there are instances of Java classes Customer and LoanRequest, they may be put in the object "decision" that is used to execute the decision. In this case, the proper table "decisionObjects" may look like:

DecisionObject decisionObjects	
Business Concept	**Business Object**
Customer	:= decision.get("customer")
Request	:= decision.get("loanRequests")
Internal	:= internal

It is important that Decision does not "know" about a particular object implementation: the only requirement is that the attribute inside these objects should have the same names as in the glossary.

Note. You cannot use the predefined function "decision()" within the table "decisionObjects" because its content is be not defined yet. You need to use the internal variable "decision" directly.

Decision Execution

OpenRules® provides a template for Java launchers that may be used to execute different decisions. There are OpenRules® API classes OpenRulesEngine and Decision. Here is an example of a decision launcher for the sample project "DecisionLoan":

```
import com.openrules.ruleengine.Decision;

public class Main {
    public static void main(String[] args) {
        String fileName = "file:rules/main/Decision.xls";
        Decision decision = new Decision("DetermineLoanPreQualificationResults",fileName);
        decision.execute();
    }
}
```

Actually, it just creates an instance of the class Decision. It has only two parameters:

1) a path to the main Excel file "Decision.xls"
2) a name of the main Decision inside this Excel file.

When you execute this Java launcher using the provided batch file "run.bat" or execute it from your Eclipse IDE, it will produce output that may look like the following:

```
*** Decision DetermineLoanPreQualificationResults ***
Decision has been initialized
Decision DetermineLoanPreQualificationResults: Calculate Internal
Variables
Conclusion: Total Debt Is 165600.0
Conclusion: Total Income Is 360000.0
Decision DetermineLoanPreQualificationResults: Validate Income
Conclusion: Income Validation Result Is SUFFICIENT
Decision DetermineLoanPreQualificationResults: Debt Research
Conclusion: Debt Research Result Is Low
Decision DetermineLoanPreQualificationResults: Summarize
Conclusion: Loan Qualification Result Is NOT QUALIFIED
ADDITIONAL DEBT RESEARCH IS NEEDED from DetermineLoanQualificationResult
*** OpenRules made a decision ***
```

This output shows all sub-decisions and conclusion results for the corresponding decision tables.

Decision Analysis

Decision Testing

OpenRules® provides an ability to create a test harness comprised of an executable set of different test cases. It is important that the same people who design rules (usually business analysts) are able to design tests for them.

Usually they create test cases directly in Excel by specifying their own data/object types and creating instances of test objects of these types. Read more at the section Defining Test Data.

Decision Syntax Validation

OpenRules® allows you to validate your decision by checking that:
- there are no syntax error in the organization of all decision tables
- values inside decision variable cells correspond to the associated domains defined in the glossary.

The validation template is described in the standard file "`DecisionTableValidateTemplates.xls`".

If you use the Eclipse Plugin, it will display...

OpenRules® also provides a special plugin for Eclipse IDE, a de-facto standard project management tools for software developers within a Java-based development environment. Eclipse is used for code editing, debugging, and testing of rule projects within a single commonly known integrated development environment. OpenRules® has been designed to catch as many errors as possible in design-rime vs. run-time when it is too late. OpenRules® Plugin automatically diagnoses errors in the Excel-files and displays the proper error messages inside Eclipse views like at the picture below:
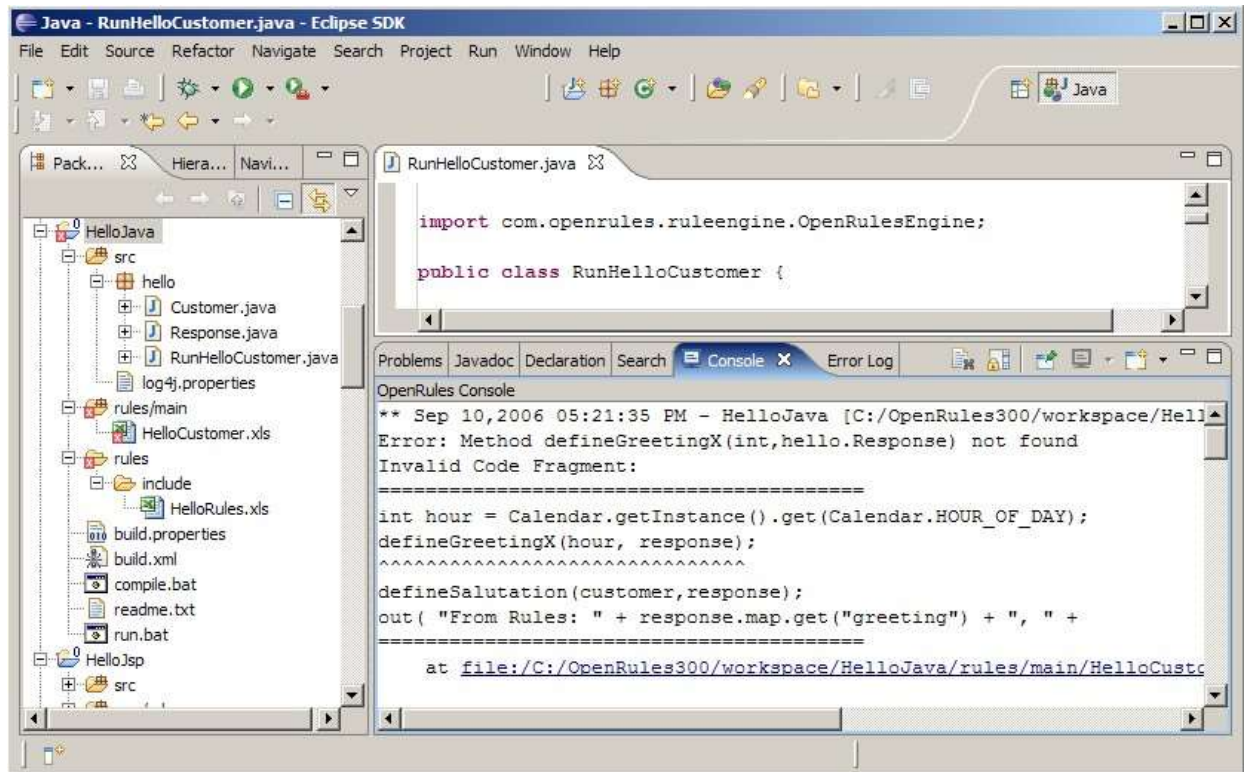
Eclipse Plugin diagnoses any errors in Excel-files before you even deploy or run your OpenRules-based application. To make sure that Eclipse controls your OpenRules® project, you have first to right-click to your project folder and "Add OpenRules Nature". You always can similarly "Remove OpenRules Nature". To be validated, your main xls-files should be placed into an Eclipse source folder while all included files should be kept in regular (non-source) folders. OpenRules® Plugin displays a diagnostic log with possible errors inside the Eclipse Console view. The error messages include hyperlinks that will open the proper Excel file with a cursor located in a cell where the error occurred.

Decision Execution Reports

OpenRules® provides an ability to generate decision execution reports in the HTML-format. To generate an execution report, you should add the following setting to the decision's Java launcher:

```
decision.put("report", "On");
```

before calling `decision.execute()`. By default, execution reports are not generated as they are needed mainly for decision analysis. Reports are regenerated for every decision run.

During decision execution, OpenRules® automatically creates a sub-directory "report" in your main project directory and generates a report inside this sub-directory. For every decision table, including single-hit, multi-hit, and rule sequencing tables, OpenRules® generates a separate html-file with the name `Report<n>.<DecisionTableName>.html`, where `n` is an execution order number for this particular decision table. For example, for the sample project "DecisionLoan" OpenRules® will generate the following files:

The first file contains a list of links to all executed decision tables:

OpenRules Execution Report of Sat Jan 05 11:19:49 EST 2013

Decision "DetermineLoanPreQualificationResults"

#	Decision Table
1	DetermineLoanPreQualificationResults
2	CalculateInternalVariables
3	DetermineIncomeValidationResult
4	DetermineDebtResearchResult
5	DetermineLoanQualificationResult

Below are other generated files (one per decision table) with lists of rules (rows) that were actually executed:

Decision Table "DetermineLoanPreQualificationResults" (Rule Sequence)

Executed Rule	ActionPrint	ActionExecute
#	**Decisions**	**Execute Rules**
1	Calculate Internal Variables	:= CalculateInternalVariables()
2	Validate Income	:= DetermineIncomeValidationResult()
3	Debt Research	:= DetermineDebtResearchResult()
4	Summarize	:= DetermineLoanQualificationResult()

Decision Table "CalculateInternalVariables" (Single-Hit)

Executed Rule	Conclusion	Conclusion
#	**Total Debt**	**Total Income**
1	Is := getInt("Monthly Debt") * getInt("Loan Term")	Is := getInt("Monthly Income") * getInt("Loan Term")

Decision Table "DetermineIncomeValidationResult" (Single-Hit)

Executed Rule	Condition	Conclusion
#	**Total Income**	**Income Validation Result**
1	Is More Than := getInt("Total Debt") * 2	Is SUFFICIENT

Decision Table "DetermineDebtResearchResult" (Single-Hit)

Executed Rule	Condition	Condition	Condition	Condition	Condition	Condition	Condition	Condition	Conclusion
#	Mortgage Holder	Outside Credit Score	Outside Credit Score	Loan Holder	Credit Card Balance	Education Loan Balance	Internal Credit Rating	Internal Analyst Opinion	Debt Research Result
1	Is Yes	-	-	-	-	-	-	-	Is High

Decision Table "DetermineLoanQualificationResult" (Single-Hit)

Executed Rule	Condition	Condition	Conclusion
#	**Income Validation Result**	**Debt Research Result**	**Loan Qualification Result**
3	Is SUFFICIENT	Is One Of Mid, High	Is QUALIFIED

These reports help a rule designer to analyze which rules were actually executed and in which order. The "Executed Rule #" corresponds to the sequential number of a rule inside its decision table.

Note. Execution reports are intended to explain the behavior of certain decision tables and are used mainly for analysis and not for production. If you turn on report generation mode in a multi-threaded environment that shares the same instance of OpenRulesEngine, the reports will be produced only for the first thread.

Decision Tracing

OpenRules® relies on the standard Java logging facilities for the decision output. They can be controlled by the standard file "log4j.properties" that by default looks like below:

```
log4j.rootLogger=INFO, stdout
log4j.appender.stdout=org.apache.log4j.ConsoleAppender
log4j.appender.stdout.layout=org.apache.log4j.PatternLayout
log4j.appender.stdout.layout.ConversionPattern=%m%n
#log4j.logger.org.openl=DEBUG
```

You may replace `INFO` to `DEBUG` and uncomment the last line to see OpenRules debugging information. To redirect all logs into a file "results.txt" you may change the file "log4j.properties" as follows:

```
log4j.rootLogger=INFO, stdout
#log4j.appender.stdout=org.apache.log4j.ConsoleAppender
log4j.appender.stdout=org.apache.log4j.FileAppender
log4j.appender.stdout.File=results.txt
log4j.appender.stdout.layout=org.apache.log4j.PatternLayout
log4j.appender.stdout.layout.ConversionPattern=%m%n
# uncomment the next line to see OpenRules debug messages
#log4j.logger.org.openl=DEBUG
```

You may control how "talkative" your decision is by setting decision's parameter "Trace". For example, if you add the following setting to the above Java launcher

```
decision.put("trace", "Off");
```

just before calling decision.execute(), then your output will be much more compact:

```
*** Decision DetermineLoanPreQualificationResults ***
Decision DetermineLoanPreQualificationResults: Calculate Internal
Variables
Decision DetermineLoanPreQualificationResults: Validate Income
Decision DetermineLoanPreQualificationResults: Summarize
ADDITIONAL DEBT RESEARCH IS NEEDED from DetermineLoanQualificationResult
*** OpenRules made a decision ***
```

You may also change output by modifying the tracing details inside the proper decision templates in the configuration files "DecisionTemplates.xls" and "DecisionTableExecuteTemplates.xls".

Rules Repository Search

To analyze rules within one Excel files you may effectively use familiar Search and Replace features provided by Excel or Google Docs.

When you want to search across multiple Excel files and folders, you may use a free while powerful tool called "IceTeaReplacer" that can be very useful for doing search & replace in OpenRules repositories. The following options are available:

- Perform search before replacing
- Match whole word only
- Ignore word case
- Do backup before replace
- Deselect files on which you don't want to perform replace.

Here is an example of its graphical interface:
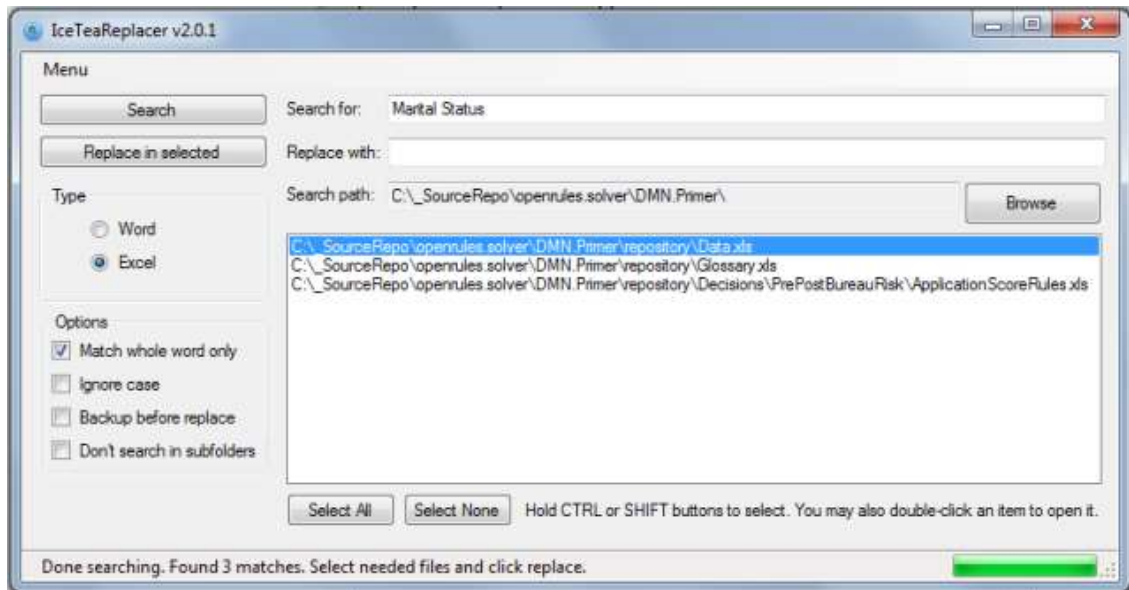
Consistency Checking

OpenRules® provides a special component Rule Solver™ that along with powerful optimization features allow a user to check consistency of the decision models and find possible conflicts within decision tables and across multiple decision tables. The detail description of the product can be found at http://openrules.com/pdf/RulesSolver.UserManual.pdf.

ADVANCED DECISION TABLES

In real-world project you may need more complex representations of rule sets and the relationships between them than those allowed by the default decision tables. OpenRules® allows you to use advanced decision tables and to define your own rule sets with your own logic.

Specialized Conditions and Conclusions

The standard columns of the types "**Condition**" and "**Conclusion**" always have two sub-columns: one for operators and another for values. OpenRules® allows you to specify columns of the types "**If**" and "**Then**" that do not require sub-columns. Instead, they allow you to use operators or even natural language expressions together with values to represent different intervals and domains of

values. Read about different ways to represent intervals and domains in this section below.

Sometimes your conditions or actions are not related to a particular decision variable and can be calculated using formulas. For example, a condition can be defined based on combination of several decision variables, and you would not want to artificially add an intermediate decision variable to your glossary in order to accommodate each needed combination of existing decision variables. In such a case, you may use a special type "**ConditionAny**" like in the example below:

DecisionTable CalculateTaxableIncome			
ConditionAny		Conclusion	
Condition		Taxable Income	
Is True	:= (taxableIncome() > 0)	Is	::= taxableIncome()
Is False	:= (taxableIncome() > 0)	Is	0

Here the word "Condition" does not represent any decision variable and instead you may insert any text, i.e. "Compare Adjusted Gross Income with Dependent Amount". When your conclusion, does not set a value for a single decision variable but rather does something that is expressed in the formulas within the cells of this column, you should use a column of type "**ActionAny**". It does not have sub-columns because there is no need for an operator.

Note. There is also a column of type "**Action**" that is equivalent to type "Then".

Specialized Decision Tables

Sometimes the default behavior of a DecisionTable (as single-hit rules tables) is not sufficient. OpenRules® provide two additional types of decision tables DecisionTable1 (or DT1) and DecisionTable2 (or DT2). While we recommend avoiding these types of decision tables, in certain situations they provide a convenient way around the limitations imposed by the standard DecisionTable.

DecisionTable1

Contrary to the standard DecisionTable that is implemented as a <u>single-hit</u> rules table, decision tables of type "DecisionTable1" or "DecisionTableMultiHit" are implemented as <u>multi-hit</u> decision tables. "DecisionTable1" supports the following rules execution logic:

1. All rules are evaluated and if their conditions are satisfied, they will be marked as "to be executed"
2. All actions columns (of the types "Conclusion", "Then", "Action", "ActionAny", or "Message") for the "to be executed" rules will be executed in top-down order.

Thus, we can make two important observations about the behavior of the "DecisionTable1":

- Rule actions cannot affect the conditions of any other rules in the decision table – there will be no re-evaluation of any conditions
- Rule overrides are permitted. The action of any executed rule may override the action of any previously executed rule.

Let's consider an example of a rule that states: "A person of age 17 or older is eligible to drive. However, in Florida 16 year olds can also drive". If we try to present this rule using the standard DecisionTable, it may look as follows:

DecisionTable ValidateDrivingEligibility					
Condition		Condition		Conclusion	
Driver's Age		US State		Driving Eligibility	
>=	17			Is	Eligible
Is	16	Is Not	Florida	Is	Not Eligible
Is	16	Is	Florida	Is	Eligible
<	16			Is	Not Eligible

Using a non-standard DecisionTable1 we may present the same rule as:

DecisionTable1 ValidateDrivingEligibility					
Condition		Condition		Conclusion	
Driver's Age		US State		Driving Eligibility	
				Is	Eligible
<	17			Is	Not Eligible
>=	16	Is	Florida	Is	Eligible

In the DecisionTable1 the first unconditional rule will set "Driving Eligibility" to "Eligible". The second rule will reset it to "Not Eligible" for all people younger than 17. But for 16 year olds living in Florida, the third rule will override the variable again to "Eligible".

DecisionTable2

There is one more type of decision table, "DecisionTable2," that is similar to "DecisionTable1" but allows the actions of already executed rules to affect the conditions of rules specified below them. "DecisionTable2" supports the following rules execution logic:

1. Rules are evaluated in top-down order and if a rule condition is satisfied, then the rule actions are immediately executed.
2. Rule overrides are permitted. The action of any executed rule may override the action of any previously executed rule.

Thus, we can make two important observations about the behavior of the "DecisionTable2":

- Rule actions can affect the conditions of other rules
- There could be rule overrides when rules defined below already executed rules could override already executed actions.

Let's consider the following example:

DecisionTable2 CalculateTaxableIncome				
Condition			Conclusion	
Taxable Income			Taxable Income	
		Is	::= (getReal("Adjusted Gross Income") - getReal("Dependent Amount"))	
Is Less	0	Is	0	

Here the first (unconditional) rule will calculate and set the value of the decision variable "Taxable Income". The second rule will check if the calculated value is less than 0. If it is true, this rule will reset this decision variable to 0.

Business Rules Defined on Collections of Objects

Previously, when OpenRules® users needed to run decisions against collections (arrays) of business objects, they needed to use Java loops. This release adds an ability to execute one decision against a collection of objects and to calculate values defined on the entire collection *within single decision run*.

Let's consider an example "**DecisionManyCustomers**" added to the standard OpenRules® installation. There is a standard Java bean Customer with different customer attributes such as name, age, gender, salary, etc. There is also a Java class "CollectionOfCustomers":

```
public class CollectionOfCustomers {
    Customer[] customers;
    int minSalary;
    int maxSalary;
    int numberOfRichCustomers;
    int totalSalary;
    ...
}
```

We want to pass this collection to a decision that will process all customers from this collection in one run and will calculate such attributes as

"minSalary", "totalSalary", "numberOfRichCustomers", and similar attributes, which are specified for the entire collection. Each customer within this collection can be processed by the following rules:

DecisionTable1 EvaluateOneCustomer

Condition	Conclusio		Conclusion	Conclusion	Conclusion	Conclusion
Salary	Wealth		Total Salary	Number of Rich Customers	Maximal Salary	Minimal Salary
			+= ::= getInt("Salary")		Max ::= getInt("Salary")	Min ::= getInt("Salary")
> 100000	Is	Rich		+= 1		

Pay attention that we use here a multi-hit table (DecsionTable1), so both rules will be executed. The first one unconditionally calculate the total. max, and min salaries. The second rule defines a number of "rich" customers inside the collection. To accumulate the proper values, we use the existing operator "+=" and newly introduced operators "Min" and "Max".

To execute the above decision table for all customers, we will utilize a new action "**ActionRulesOnArray**" within the following decision table:

DecisionTable CalculateCustomerTotals

Conclusion		Conclusion		Conclusion		ActionRulesOnArray		
Total Salary		Maximal Salary		Minimal Salary		Array of Objects	Object Type	Rules
Is	0	Is	0	Is	1000000	Customers	Customer	EvaluateOneCustomer

Here the first 3 actions (conclusions) simply define initial values of collection attributes. The last action has 3 sub-columns:

- The name of the array of objects as it is defined in the glossary ("Customers")
- The type of those objects ("Customer")
- The name of the decision table ("EvaluateOneCustomer") that will be used to processes each objects form this collection.

Thus, a combination of the two decisions tables (similar to the above ones) provides business users with a quite intuitive way to apply rules over collections of business objects without necessity to deal with programming constructions.

Decision Tables for Comparing Ranking Lists

In many real-world situations decisions are made based on comparison of attributes that belong to different predefined lists of values while the values inside these lists are ordered (ranked). For example, a business rule may sound as follows:

"If Diagnostic Need is Stronger Than Sensitivity Level

Then Document Access should be Allowed"

Here the Diagnostic Need could belong to the ranking list:

1. Immediately Life-Threatening
2. Life-Threatening
3. Acute
4. Chronic.

Similarly the Sensitivity Level could belong to this ranking list:

1. High
2. Mid
3. Low.

Newly defined custom templates allow us to present the relations between these two ranking lists in the following decision table of the new type "**DecisionTableCompareRanks**":

DecisionTableCompareRanks CompareDiagnosticNeedWithSensitivityLevel			
Condition	Action	Action	Action
Diagnostic Need Sensitivity Level	High	Mid	Low
Immediately Life-Threatening	Stronger	Stronger	Stronger
Life-Threatening	Weaker	Stronger	Stronger
Acute	Weaker	Weaker	Weaker
Chronic	Weaker	Weaker	Weaker

Then the above rule may be expressed using the following decision table of the new type "**DecisionTableRanking**":

DecisionTableRanking DefineDocumentAccess				
ConditionCompareRanks			Conclusion	
If <rank1> <stronger/weaker> <rank2>			Document Access	
Diagnostic Need	Stronger	Sensitivity Level	Is	Allow
Diagnostic Need	Weaker	Sensitivity Level	Is	Decline

To define "Stronger/Weaker" relations between these ranks, this decision table will automatically invoke the decision table with the dynamically defined name "*Compare<rank1>With<rank2>*" (after removing all spaces).

The benefits of these new types of decision tables become clear when you think about supporting hundreds of similar ranking lists. These tables may cover complex relationships between multiple ranking lists and at the same time they remain easy to understand and to be maintained by business users.

The complete working example "**DecisionRankingLists**" with the proper custom templates (see file "RankTemplates.xls") is included into the standard OpenRules® installation.

RULE TABLES

OpenRules® supports several ways to represent business rules inside Excel tables. Default decision table is the most popular way to present sets of related business rules because they do not require any coding. However, there classical decision tables can represent more complex execution logic that is frequently custom for different conditions and actions.

Actually, standard DecisionTable is a special case of an OpenRules® single-hit decision table that is based on a predefined template (see below). Since 2003, OpenRules® allows its users to configure different types of custom decision tables directly in Excel. In spite of the necessity to use Java snippets to specify custom logic, these tables are successfully used by major corporations in real-world decision support applications. This chapter describes different decision tables that go beyond the default decision tables. It will also describe how to use simple IF-THEN-ELSE statements within Excel-based tables of type "Method".

Simple Rule Table

Let's consider a simple set of HelloWorld rules that can be used to generate a string like "Good Morning, World!" based on the actual time of the day. How one understands such concepts as "morning", "afternoon", "evening", and "night" is defined in this simple rules table:

Rules void helloWorld(int hour)		
Hour From	**Hour To**	**Greeting**
0	11	Good Morning
12	17	Good Afternoon
18	22	Good Evening
23	24	Good Night

Hopefully, this rule table is not much more difficult to compare with the default DecisonTable. It states that if the current hour is between 0 and 11, the greeting should be "Good Morning", etc. You may change Hour From or Hour To if you want to customize the definition of "morning" or "evening". This table is also oriented to a business user. However, its first row already includes some technical information (a table signature):

```
Rules void helloWorld(int hour)
```

Here "**Rules**" is an OpenRules® keyword for this type of tables. "helloWorld" is the name of this particular rules table. It tells to an external program or to other rules how to launch this rules table. Actually, this is a typical description of a programming method (its signature) that has one integer parameter and returns nothing (the type "void"). The integer parameter "hour" is expected to contain the current time of the day. While you can always hide this information from a business user, it is an important specification of this rule table.

You may ask: where is the implementation logic for this rule table? All rule tables include additional hidden rows (frequently password protected) that you can see if you click on the buttons "+" to open the Technical View below:

1 2	A	B	C	D	
1					
2					
3		Rules void helloWorld(int hour)			
4		C1	C2	A1	
5		min <= hour	hour <= max	System.out.println(greeting + ", World!")	
6		int min	int max	String greeting	
7		**Hour From**	**Hour To**	**Greeting**	
8		0	11	Good Morning	
9		12	17	Good Afternoon	
10		18	22	Good Evening	
11		23	24	Good Night	
12					

This part of the rule table is oriented to a technical user, who is not expected to be a programming guru but rather a person with a basic knowledge of the "C" family of languages which includes Java. Let's walk through these rows step by step:

- Row "Condition and Action Headers" (see row 4 in the table above). The initial columns with conditions should **start with the letter "C"**, for example "C1", "Condition 1". The columns with actions should **start with the letter "A"**, for example "A1", "Action 1".
- Row "Code" (see row 5 in the table above). The cells in this row specify the semantics of the condition or action associated with the corresponding columns. For example, the cell B5 contains the code min <= hour. This means that condition C1 will be true whenever the value for *min* in any cell in the column below in this row is less than or equals to the parameter hour. If *hour* is 15, then the C1-conditions from rows 8 and 9 will be satisfied. The code in the Action-columns defines what should be done when all conditions are satisfied. For example, cell D5 contains the code:

```
System.out.println(greeting + ", World!")
```

This code will print a string composed of the variable greeting and ", World!", where greeting will be chosen from a row where all of the conditions are

49 ◎

satisfied. Again, if hour is 15, then both conditions C1 and C2 will be satisfied only for row 9 (because 9 <= 15 <= 17). As a result, the words "Good Afternoon, World!" will be printed. If the rule table does not contain a row where all conditions have been satisfied, then no actions will be executed. Such a situation can be diagnosed automatically.

- Row "Parameters" (see row 6 in the table above). The cells in this row specify the types and names of the parameters used in the previous row.
- Row "Display Values" (see row 7 in the table above). The cells in this row contain a natural language description of the column content.

The same table can be defined a little bit differently using one condition code for both columns "min" and "max":

Rules void **defineGreeting**(App app, int hour)		
C1		A1
min <= hour && hour <= max		app.greeting = greeting;
int min	int max	String greeting
Hour From	**Hour To**	**Set Greeting**
0	11	Good Morning
12	17	Good Afternoon
18	22	Good Evening
23	24	Good Night

How Rule Tables Are Organized

As you have seen in the previous section, rule tables have the following structure:

Row 1: Signature

Rules void tableName(Type1 par1, Type2 par2, ..) - Multi-Hit Rule Table

Rules <JavaClass> tableName(Type1 par1, Type2 par2, ..) - Single-Hit Rule Table

Row 2: Condition/Action Indicators

The condition column indicator is a word starting with "C".

The action column indicator is a word starting with "A".

All other starting characters are ignored and the whole column is considered as a comment

Row 3: Code

The cells in each column (or merged cells for several columns) contain Java Snippets.

Condition codes should contain expressions that return *Boolean* values.

If an action code contains any correct Java snippet, the return type is irrelevant.

Row 4: Parameters

Each condition/action may have from 0 to N parameters. Usually there is only one parameter description and it consists of two words:

> *parameterType parameterName*

Example: *int min*

parameterName is a standard one word name that corresponds to Java identification rules.

parameterType can be represented using the following Java types:

- Basic Java types: `boolean, char, int, long, double, String, Date`
- Standard Java classes: `java.lang.Boolean, java.lang.Integer, java.lang.Long, java.lang.Double, java.lang.Character, java.lang.String, java.util.Date`
- Any custom Java class with a public constructor that has a String parameter
- One-dimensional arrays of the above types.

Multiple parameters can be used in the situations when one code is used for several columns. See the standard example "Loan1.xls".

Row 5: Columns Display Values

Text is used to give the column a definition that would be meaningful to another reader (there are no restrictions on what text may be used).

Row 6 and below: Rules with concrete values in cells

Text cells in these rows usually contain literals that correspond to the parameter types.

For Boolean parameters you may enter the values "TRUE" or "FALSE" (or equally "Yes" or "No") without quotations.

Cells with Dates can be specified using `java.util.Date`. OpenRules® uses `java.text.DateFormat.SHORT` to convert a text defined inside a cell into `java.util.Date`. Before OpenRules® 4.1 we recommended our customers not to use Excel's Date format and define Date fields in Excel as Text fields. The reason was the notorious Excel <u>problem</u> inherited from a wrong assumption that 1900 was a leap year. As a result, a date entered in Excel as 02/15/2004 could be interpreted by OpenRules® as 02/16/2004. Starting with release 4.1 OpenRules® correctly interprets both Date and Text Excel Date formats.

Valid Java expression (Java snippets) may be put inside table cells by one of two ways:

- by surrounding the expression in curly brackets, for example: { driver.age+1; }
- by putting ":=" in front of your Java expression, for example: :=driver.age+1

    Make sure that the expression's type corresponds to the parameter type.

Empty cells inside rules means "whatever" and the proper condition is automatically considered satisfied. An action with an empty value will be

ignored. If the parameter has type String and you want to enter a space character, you must explicitly enter one of the following expressions:

```
:= " "
```

```
'= " "
```

```
{ " "; }
```

Note. Excel is always trying to "guess" the type of text inside its cells and automatically converts the internal representation to something that may not be exactly what you see. For example, Excel may use a scientific format for certain numbers. To avoid a "strange" behavior try to explicitly define the format "text" for the proper Excel cells.

Separating Business and Technical Information

During rules harvesting, business specialists initially create rule tables using regular Excel tables. They put a table name in the first row and column names in the second row. They start with Conditions columns and end with Action columns. For example, they can create a table with 5 columns [C1,C2,C3,A1,A2] assuming the following logic:

```
IF conditions C1 and C2 and C3 are satisfied

THEN execute actions A1 and A2
```

Then, a business specialist provides content for concrete rules in the rows below the title rows.

As an example, let's consider the rule table "defineSalutation" with the rules that define how to greet a customer (Mr., Ms, or Mrs.) based on his/her gender and marital status. Here is the initial business view (it is not yet syntactically correct):

Rules defineSalutation		
Gender	**Marital Status**	**Set Salutation**
Male		Mr.
Female	Married	Mrs.
Female	Single	Ms.

A business analyst has initially created only five rows:

- A signature "Rules defineSalutation" (it is not a real signature yet)
- A row with column titles: two conditions "Gender", "Marital Status" and one action "Set Salutation"
- Rows with three rules that can be read as:

1) IF Gender is "Male" THEN Set Salutation "Mr."
2) IF Gender is "Female" and Marital Status is "Married" THEN Set Salutation "Mrs."
3) IF Gender is "Female" and Marital Status is "Single" THEN Set Salutation "Ms."

While business specialists continue to define such rule tables, at some point a technical specialist should take over and add to these tables the actual implementation. The technical specialist (familiar with the software environment into which these rules are going to be embedded) talks to the business specialist (author of the rule table) about how the rules should be used. In the case of the "defineSalutation" rule table, they agree that the table will be used to generate a salutation to a customer. So, the technical specialist decides that the table will have two parameters:

1) a customer of the type Customer
2) a response of the type Response

The technical specialist will modify the signature row of the table to look like this:

```
Rules void defineSalutation(Customer customer, Response response)
```

Then she/he inserts three more rows just after the first (signature) row:

- Row 2 with Condition/Action indicators

- Row 3 with Condition/Action implementation

- Row 4 with the type and name of the parameters entered in the proper column.

Here is a complete implementation of this rule table:

Rules void define Salutation(Customer customer, Response response)		
C1	C2	A1
customer.gender.equals (gender)	customer.maritalStatus.equal s(status)	response.map.put("salutation", salutation);
String gender	String status	String salutation
Gender	**Marital Status**	**Set Salutation**
Male		Mr.
Female	Married	Mrs.
Female	Single	Ms.

The rules implementer will decide that to support this rule table, type Customer should have at least two attributes, "gender" and "maritalStatus", and the type Response should be able somehow to save different pairs (names,value) like*("salutation","Mr.")*. Knowing the development environment, s/he will decide on the types of attributes. Let's assume that both types Customer and Response correspond to Java classes, and the attributes have the basic Java type of String. In this case, the column "Gender" will be marked with a parameter "String gender" and the condition will be implemented as a simple boolean expression:

```
customer.gender.equals(gender)
```

The second column "C2" is implemented similarly with a String attribute and a parameter maritalStatus. Finally (to make it a little bit more complicated), we will assume that the class Response contains an attribute map of the predefined Java type HashMap, in which we can put/get pairs of Strings. So, the implementation of the action "Set Salutation" will look like:

```
response.map.put("salutation",salutation)
```

How Rule Tables Are Executed

The rules inside rule tables are executed one-by-one in the order they are placed in the table. The execution logic of one rule (row in the vertical table) is the following:

IF ALL conditions are satisfied THEN execute ALL actions.

If at least one condition is violated (evaluation of the code produces **false)**, all other conditions in the same rule (row) are ignored and **are not evaluated.** The absence of a parameter in a condition cell means the condition is always **true.** Actions are evaluated only if all conditions in the same row are evaluated to be **true** and the action has non-empty parameters. Action columns with no parameters are ignored.

For the default <u>vertical</u> rule tables, all rules are executed in top-down order. There could be situations when all conditions in two or more rules (rows) are satisfied. In that case, the actions of all rules (rows) will be executed, and the actions in the rows below can **override** the actions of the rows above.

For <u>horizontal</u> rule tables, all rules (columns) are executed in left-to-right order.

Relationships between Rules inside Rule Tables

OpenRules® does not assume any implicit ("magic") execution logic, and executes rules in the order specified by the rule designer. All rules are executed one-by-one in the order they are placed in the rule table. There is a simple rule that governs rules execution inside a rules table:

The preceding rules are evaluated and executed first!

OpenRules® supports the following types of rule tables that offer different execution logic to satisfy different practical needs:

- Multi-hit rule tables
- Single-hit rule tables
- Rule Sequences.

Multi-Hit Rule Tables

A multi-hit rule table evaluates conditions in ALL rows before any action is executed. Thus, actions are executed only AFTER all conditions for all rules have already been evaluated. From this point of view, the execution logic is different from traditional programming if-then logic. Let us consider a simple example. We want to write a program "swap" that will do the following:

```
If x is equal to 1 then make x to be equal to 2.

If x is equal to 2 then make x to be equal to 1.
```

Suppose you decided to write a Java method assuming that there is a class App with an integer variable x. The code may (but should not) look like this:

```
void swapX(App app) {
      if (app.x == 1) app.x = 2;
      if (app.x == 2) app.x = 1;
}
```

Obviously, this method will produce an incorrect result because of the missing "else". This is "obvious" to a software developer, but may not be at all obvious to a business analyst. However, in a properly formatted rule table the following representation would be a completely legitimate:

If x equals to	Then make x to be equal to
1	2
2	1

It will also match our plain English description above. Here is the same table with an extended technical view:

Rules void swapX(App app)	
C	A
app.x == oldValue	app.x = newValue
int oldValue	int newValue
If x equals to	Then make x to be equal to
1	2
2	1

Rules Overrides in Multi-Hit Rule Tables

There could be situations when all conditions in two or more rules (rows) are satisfied at the same time (multiple hits). In that case, the actions of all rules (rows) will be executed, but the actions in the rows below can **override** the actions of the rows above. This approach also allows a designer to specify a very natural requirement:

More specific rules should override more generic rules!

The only thing a designer needs to guarantee is that "more specific" rules are placed in the same rule table after "more generic" rules. For example, you may want to execute Action-1 every time that Condition-1 and Condition-2 are satisfied. However, if additionally, Condition-3 is also satisfied, you want to execute Action-2. To do this, you could arrange your rule table in the following way:

Condition-1	Condition-2	Condition-3	Action-1	Action-2
X	X		X	
X	X	X		X

In this table the second rule may override the first one (as you might naturally expect).

Let's consider the execution logic of the following multi-hit rule table that defines a salutation "Mr.", "Mrs.", or "Ms." based on a customer's gender and marital status:

Rules void **defineSalutation**(Customer customer, Response response)		
Gender	**Marital Status**	**Set Salutation**
Male		Mr.
Female	Married	Mrs.
Female	Single	Ms.

If a customer is a married female, the conditions of the second rules are satisfied and the salutation "Mrs." will be selected. This is only a business view of the rules table. The complete view including the hidden implementation details ("Java snippets") is presented below:

Rules void **defineSalutation**(Customer customer, Response response)		
C1	C2	A1
customer.gender. equals(gender)	customer.maritalStatus. equals(status)	response.map.put("salutation",s alutation);
String gender	String status	String salutation
Gender	**Marital Status**	**Set Salutation**
Male		Mr.
Female	Married	Mrs.
Female	Single	Ms.

The OpenRulesEngine will execute rules (all 3 "white" rows) one after another. For each row if conditions C1 and C2 are satisfied then the action A1 will be executed with the selected "salutation". We may add one more rule at the very end of this table:

Rules void **defineSalutation**(Customer customer, Response response)		
Gender	**Marital Status**	**Set Salutation**
Male		Mr.
Female	Married	Mrs.
Female	Single	Ms.
		???

In this case, after executing the second rule OpenRules® will also execute the new, 4th rule and will override a salutation "Mrs." with "???". Obviously this is not a desirable result. However, sometimes it may have a positive effect by avoiding undefined values in cases when the previous rules did not cover all possible situations. What if our customer is a Divorced Female?! How can this multi-hit effect be avoided? What if we want to produce "???" only when no other rules have been satisfied?

Single-Hit Rule Tables

To achieve this you may use a so-called "**single-hit**" rule table, which is specified by putting any return type **except** "**void**" after the keyword "**Rules**". The following is an example of a single-hit rule table that will do exactly what we need:

Rules String **defineSalutation**(Customer customer, Response response)		
Gender	**Marital Status**	**Set Salutation**
Male		Mr.
Female	Married	Mrs.
Female	Single	Ms.
		???

Another positive effect of such "single-hitness" may be observed in connection with large tables with say 1000 rows. If OpenRules® obtains a hit on rule #10 it would not bother to check the validity of the remaining 990 rules.

Having rule tables with a return value may also simplify your interface. For example, we do not really need the special object Response which we used to write our defined salutation. Our simplified rule table produces a salutation without an additional special object:

Rules String **defineSalutation**(Customer customer)		
C1	C2	A1
customer.gender.equals(gender)	customer.maritalStatus.equals(status)	return salutation;
String gender	String status	String salutation
Gender	**Marital Status**	**Set Salutation**
Male		Mr.
Female	Married	Mrs.
Female	Single	Ms.
		???

Please note that the last action in this table should return a value that has the same type as the entire single-hit table. The single-hit table may return any standard or custom Java class such as String or Customer. Instead of basic Java types such as "int" you should use the proper Java classes such as Integer in the table signature.

Here is an example of Java code that creates an OpenRulesEngine and executes the latest rules table "defineSalutation":

```
public static void main(String[] args) {
    String fileName = "file:rules/main/HelloCustomer.xls";
    OpenRulesEngine engine =
                    new OpenRulesEngine(fileName);
    Customer customer = new Customer();
    customer.setName("Robinson");
    customer.setGender("Female");
    customer.setMaritalStatus("Married");
    String salutation =
            (String)engine.run("defineSalutation", customer);
    System.out.println(salutation);
}
```

Rule Sequences

There is one more type of rule tables called "Rule Sequence" that is used mainly internally within templates. Rule Sequence can be considered as a multi-hit rule table with only one difference in the execution logic, conditions are not evaluated before execution of the actions. So, all rules will be executed in top-down order with possible rules overrides. Rule actions are permitted to affect the conditions of any rules that follow the action. The keyword "Rules" should be replaced with

another keyword "**RuleSequence**". Let's get back to our "swapX" example. The following multi-hit table will correctly solve this problem:

Rules void swapX(App app)	
C	A
app.x == oldValue	app.x = newValue; app.x;
int oldValue	int newValue
If x equals to	**Then make x to be equal to**
1	2
2	1

However, a similar rule sequence

RuleSequence void swapX(App app)	
C	A
app.x == oldValue	app.x = newValue; app.x;
int oldValue	int newValue
If x equals to	**Then make x to be equal to**
1	2
2	1

will fail because when x is equal to 1, the first rule will make it 2, and then the second rules will make it 1 again.

Relationships among Rule Tables

In most practical cases, business rules are not located in one file or in a single rule set, but rather are represented as a hierarchy of **inter-related rule tables** located in different files and directories - see Business Rules Repository. Frequently, the main Excel-file contains a main method that specifies the execution logic of multiple decision tables. You may use the table "Decision" for the same purposes. In many cases, the rule engine can execute decision tables directly from a Java program – see API.

Because OpenRules® interprets rule tables as regular methods, designers of rules frequently create special "processing flow" decision tables to specify the conditions under which different rules should be executed. See examples of processing flow rules in such sample projects as Loan2 and LoanDynamics.

Simple AND / OR Conditions in Rule Tables

All conditions inside the same row (rule) are considered from left to right using the AND logic. For example, to express

```
if (A>5 && B >10) {do something}
```

you may use the rule table:

Rules void testAND(int a, int b)		
C1	C2	A1
a > 5	b>10	System.out.println(text)
String x	String x	String text
A > 5	**B > 10**	**Do**
X	X	Something

To express the OR logic

```
if (A>5 || B >10) {do something}
```

you may use the rules table:

Rules void testOR(int a, int b)		
C1	C2	A1
a > 5	b>10	System.out.println(text)
String x	String x	String text
A > 5	**B > 10**	**Do**
X		Something
	X	

Sometimes instead of creating a decision table it is more convenient to represent rules using simple Java expressions inside Method tables. For example, the above rules table may be easily represented as the following Method table:

```
Method void testOR(int a, int b)

if (a > 5 || b>10)  System.out.println("Something");
```

Horizontal and Vertical Rule Tables

Rule tables can be created in one of two possible formats:

- Vertical Format (default)
- Horizontal Format.

Based on the nature of the rule table, a rules creator can decide to use a vertical format (as in the examples above where concrete rules go vertically one after another) or a horizontal format where Condition and Action are located in the rows and the rules themselves go into columns. Here is an example of the proper horizontal format for the same rule table "helloWorld":

Rules void helloWorld(int hour) //horizontal			
Hour From	0	12	18	23
Hour To	11	17	22	24
Greeting	Good Morning	Good Afternoon	Good Evening	Good Night

OpenRules® automatically recognizes that a table has a vertical or a horizontal format. You can use Excel's Copy and Paste Special feature to transpose a rule table from one format to another.

Note. When a rule table has too many rules (more than you can see on one page) it is better to use the vertical format to avoid Excel's limitations: a worksheet has a maximum of 65,536 rows but it is limited to 256 columns.

Merging Cells

OpenRules® recognizes the powerful Cell Merging mechanism supported by Excel and other standard table editing tools. Here is an example of a rule table with merged cells:

Rules void testMerge(String value1, String value2)				
Rule	C1	C2	A1	A2
	value1.equals(val)	value2.equals(val)	out("A1: " + text);	out("A2: " + text);
	String val	String val	String text	String text
#	**Name**	**Value**	**Text 1**	**Text 2**
1	B	One	11+21	12
2		Two		22
3		Three	31	32
4	D		41	42

The semantics of this table is intuitive and described in the following table:

Value 1	Value 2	Applied Rules	Printed Results
B	One	1	A1: 11+21 A2: 12
B	Two	2	A1: 11+21 A2: 22
B	Three	3	A1: 31 A2: 32
D	Three	4	A1: 41 A2: 42
A	Two	none	
D	Two	none	

Restriction. We added the first column with rules numbers to avoid the known implementation restriction that the very first column (the first row for horizontal rule tables) cannot contain merged rows. More examples can be found in the standard rule project "Merge" - click <u>here</u> to analyze more rules. When you use the standard decision tables, you may put the standard condition "C#" in the very first column and use numbers to mark each table's row.

Sub-Columns and Sub-Rows for Dynamic Arrays

One table column can consist of several sub-columns (see sub-columns "Min" and "Max" in the example <u>above</u>). You may efficiently use the Excel merge

mechanism to combine code cells and to present them in the most intuitive way. Here is an example with an unlimited number of sub-columns:

C6			
contains(rates,customer.rate)			
String[] rates			
AND **Internal Credit Rating**			
A	B	C	
D	F		
B	C	D	F
A	C	D	F

As you can see, condition C6 contains 4 sub-columns for different combinations of rates. The cells in the Condition, code, parameters and display values, rows are merged. You can insert more sub-columns (use Excel's menu "Insert") to handle more rate combinations if necessary without any changes in the code. The parameter row is defined as a String array, `String[] rates`. The actual values of the parameters should go from left to right and the first empty value in a sub-column should indicate the end of the array "*rates*". You can see the complete example in the rule table "Rule Family 212" in the file Loan1.xls.

If your rule table has a horizontal format, you may use multiple sub-rows in a similar way (see the example in file UpSell.xls).

Using Expressions inside Rule Tables

OpenRules® allows a rules designer to use "almost" natural language expressions inside rule tables to represent intervals of numbers, strings, dates, etc. You also may use Java expressions whenever necessary.

Integer and Real Intervals

You may use plain English expressions to define different intervals for integer and real decision variables inside rule tables. Instead of creating multiple columns for defining different ranges for integer and real values, a business user

may define from-to intervals in practically unlimited English using such phrases as: "500-1000", "between 500 and 1000", "Less than 16", "More or equals to 17", "17 and older", "< 50", ">= 10,000", "70+", "from 9 to 17", "[12;14)", etc.

You also may use many other ways to represent an interval of integers by specifying their two bounds or sometimes only one bound. Here are some examples of valid integer intervals:

Cell Expression	Comment
5	equals to 5
[5,10]	contains 5, 6, 7, 8, 9, and 10
5;10	contains 5, 6, 7, 8, 9, and 10
[5,10)	contains 5 but not 10
5 - 10	contains 5 and 10
5-10	contains 5 and 10
5- 10	contains 5 and 10
-5 - 20	contains -5 and 20
-5 - -20	error: left bound is greater than the right one
-5 - -2	contains -5 , -4, -3, -2
from 5 to 20	contains 5 and 20
less 5	does not contain 5
less than 5	does not contain 5
less or equals 5	contains 5
less or equal 5	contains 5
less or equals to 5	contains 5
smaller than 5	does not contain 5
more 10	does not contain 10
more than 10	does not contain 10
10+	more than 10
>10	does not contain 10
>=10	contains 10
between 5 and 10	contains 5 and 10
no less than 10	contains 10
no more than 5	contains 5

equals to 5	equals to 5
greater or equal than 5 and less than 10	contains 5 but not 10
more than 5 less or equal than 10	does not contain 5 and contains 10
more than 5,111,111 and less or equal than 10,222,222	does not contain 5,111,111 and contains 10,222,222
[5'000;10'000'000)	contains 5,000 but not 10,000,000
[5,000;10,000,000)	contains 5,000 but not 10,000,000
(5;100,000,000]	contains 5,000 and 10,000,000

You may use many other ways to represent integer intervals as you usually do in plain English. The only limitation is the following: *min* should always go before *max*!

Similarly to integer intervals, one may use the predefined type **FromToDouble** to represent intervals of real numbers. The bounds of double intervals could be integer or real numbers such as [2.7; 3.14).

Comparing Integer and Real Numbers

You may use the predefined type **CompareToInt** to compare a decision variable with an integer number that is preceded by a comparison operator. Examples of acceptable operators:

Cell Expression	Comment
<= 5	less or equals to 5
< 5	strictly less than 5
> 5	strictly more than 5
>= 5	more or equals to 5
!=	not equal to 5
5	equals to 5. Note that absence of a comparison operator means equality. You cannot use an explicit operator "=" (not to be confused with Excel's formulas).

Similarly to **CompareToInt** one may use the predefined type **CompareToDouble** to represent comparisons with real numbers. The comparison values may be presented as integer or real numbers, e.g. "<= 25.4" and "> 0.5".

Using Comparison Operators inside Rule Tables

A user can employ a comparison operators such as "<" for "less" or ">" for "more" directly inside the rules. There are several ways to accomplish this. Here is an example from the rule table "Rule Family 212" (Loan1.xls):

C4	
op.compare(c.creditCardBalance, value)	
Operator op	int value
Credit Card Balance	
Oper	Value
<=	0
>	0
>	0
<=	0
<=	0
	0

```
>
<
<=
>=
==
!=
```

You may use the Excel Data Validation menu to limit the choice of the operators:
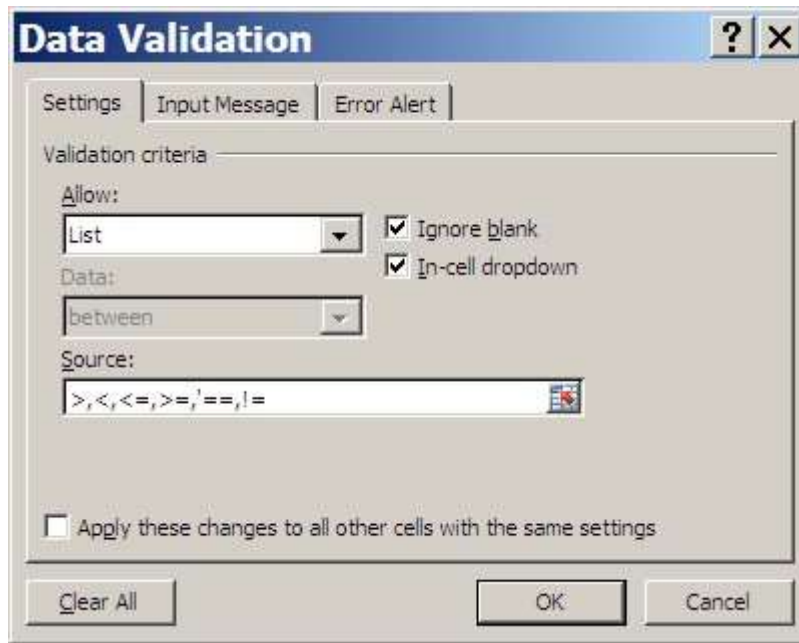
Here the sign "==" has an apostrophe in front to differentiate it from an Excel formula. The actual implementation of possible comparison operators is provided as an example in the project "com.openrules.tools" (see com.openrules.tools.Operator.java). You may change them or add other operators. In addition to values of the type "int" you may also use Operator to compare long, double, and String types.

Comparing Dates

You may use the standard java.util.Date or any other Java Comparable type. Here is an example of comparing Dates:

C1	
op.compare(visit.date,date)	
Operator op	Date date
Operator	**Date**
==	2/15/2007
!=	1/1/2007
<=	2/15/2007
>	2/15/2007
<	2/15/2007

Please note that the current implementation compares dates without time.

Another way to use operators directly inside a table is to use <u>expressions</u>. In the example above, instead of having two sub-columns "Operator" and "Value" we could use one column and put an expression inside the cell:

```
{ c.creditCardBalance <= 0; }
```

The use of expressions is very convenient when you do not know ahead of time which operator will be required for which columns.

Comparing Boolean Values

If a parameter type is defined as "boolean", you are allowed to use the following values inside rule cells:

- True, TRUE, Yes, YES
- False, FALSE, No, NO

You also may use formulas that produces a Boolean, .e.g.

```
{ loan.additionalIncomeValidationNeeded; }
```

Sometimes, you want to indicate that a condition is satisfied or an action should be executed. You may use any character like X or * without checking its actual value – the fact that the cell is not empty indicates that the condition is true. For example, in the following table (from the standard project VacationDays)

Rules void DecisionTable(Test t)											
C1	t.age >= max	int max	Age >=		18	18	18	45	45	45	60
C2	t.age < min	int min	Age <	18	45	45	45	60	60	60	
C3	t.service >= max	int max	Service >=			25	40		25	40	
C4	t.service < min	int min	Service <		25	40		25	40		
A1	t.days = 22	String X	Assign 22 days	X	X	X	X	X	X	X	X
A2	t.days += 5	String X	5 extra days	X							
A3	t.days += 2	String X	2 extra days			X	X	X	X	X	X
A4	t.days += 3	String X	3 extra days							X	X

only actions marked with "X" will be executed. You can use any other character instead of "X".

Representing String Domains

Let's express a condition that validates if a customer's internal credit score is one of several acceptable rates such as "A B C" and "D F". To avoid the necessity to create multiple sub-columns for similar conditions, we may put all possible string values inside the same cell and separate them by spaces or commas.   Here is an example of such a condition:

Condition
domain.contains(customer.internalCreditRating)
DomainString domain
Internal Credit Rating
A B C
D F
D F
A B C

Here we use the predefined type **DomainString** that defines a domain of strings (words) separated by whitespaces. The method "contains(String string)" of the class DomainString checks if the parameter "string" is found among all strings listed in the current "domain". You also may use the method

72 ◎

`"containsIgnoreCase(String  string)"` that allows you to ignore case during the comparison.

If possible values may contain several words, one may use the predefined type **DomainStringC** where "**C**" indicates that commas will be used as a string separator. For example, we may use **DomainStringC** to specify a domain such as "Very Hot, Hot, Warm, Cold, Very Cold".

Representing Domains of Numbers

If you need to represent domains of integer or double values, there are several predefined types similar to DomainString:

- `DomainInt`
- `DomainIntC`
- `DomainDouble`
- `DomainDoubleC`

For example, here is a condition column with eligible loan terms:

Condition
domain.contains(c.loanTerm)
DomainIntC domain
Eligible Loan Terms
24,36,72
36,72
72

Using Java Expressions

The use of Java expressions provides the powerful ability to perform calculations and test for complex logical conditions. While the writing of expressions requires some technical knowledge, it does not require the skills of a programmer. Real-world experience shows that business analysts frequently have a need to write these expressions themselves. It is up to the rule table designer to decide

whether to show the expressions to business people or to hide them from view. Let's consider a decision table for "Income Validation" from the provided standard example "Loan1":

Rules void **ValidateIncomeRules**(LoanRequest loan, Customer customer)	
C1	A1
customer.monthlyIncome * 0.8 - customer.monthlyDebt > loan.amount/loan.term	loan.incomeValidationResult = result;
boolean condition	String result
IF **Income is Sufficient for the Loan**	THEN **Set Income Vaidation Result**
No	UNSUFFICIENT
Yes	SUFFICIENT

Here the actual income validation expression is hidden from business people inside "gray" technical rows, and a business person would only be able to choose between "Yes" or "No". However, the same table could be presented in this way:

Rules void **ValidateIncomeRules**(LoanRequest loan, Customer customer)	
C1	A1
condition == true	loan.incomeValidationResult = result;
boolean condition	String result
IF **Condition is True**	THEN **Set Income Validation Result**
	UNSUFFICIENT
:= customer.monthlyIncome * 0.8 - customer.monthlyDebt > loan.amount/loan.term	SUFFICIENT

Now, a user can both see and change the actual income validation condition.

Notes: Do not use Excel's formulas if you want the content to be recognized by the OpenRules® engine: use OpenRules® expressions instead.

If you want to start your cell with "=" you have to put an apostrophe in front of it i.e. '= to direct Excel not to attempt to interpret it as a formula.

Expanding and Customizing Predefined Types

All the predefined types mentioned above are implemented in the Java package `com.openrules.types`. You may get the source code of this package and expand and/or customize the proper classes. In particular, for internationalization purposes you may translate the English key words into your preferred language. You may change the default assumptions about inclusion/exclusion of bounds inside integer and real intervals. You may add new types of intervals and domains.

Performance Considerations

The use of expressions inside OpenRules® tables comes with some price - mainly in performance, for large rule tables. This is understandable because for every cell with an expression OpenRules® will create a separate instance of the proper Java class during rules execution. However, having multiple representation options allows a rule designer to find a reasonable compromise between performance and expressiveness.

RULE TEMPLATES

OpenRules® provides a powerful yet intuitive mechanism for compactly organizing enterprise-level business rules repositories. Rule templates allow rule designers to write the rules logic once and use it many times. With rule templates you may completely hide rules implementation details from business users. OpenRules® supports several rule templatization mechanisms from simple rule tables that inherit the exact structure of templates to partial template implementations.

Simple Rules Templates

Rule templates are regular rule tables that serve as structural prototypes for many other rule tables with the same structure but different content (rules). A simple rule template usually does not have rules at all but only specifies the table structure and implementation details for conditions and actions. Thus, a simple rule template contains the first 5 rows of a regular decision table as in the following example:

Rules void **defineGreeting**(App app, int hour)			Signature with parameters
C1		A1	Conditions and Actions identifiers
min <= hour && hour <= max		app.greeting = greeting;	Java snippets describe condition/action semantics
int min	int max	String greeting	Parameter types and names
Hour From	**Hour To**	**Set Greeting**	Business names for conditions and actions

We may use this rule table as a template to define different greeting rules for summer and winter time. The actual decision tables will **implement** (or **extend**) the template table with particular rules:

Rules summerGreeting template defineGreeting		
Hour From	**Hour To**	**Set Greeting**
0	10	Good Morning
11	18	Good Afternoon
19	22	Good Evening
23	24	Good Night

and

Rules winterGreeting **template** defineGreeting		
Hour From	**Hour To**	**Set Greeting**
0	11	Good Morning
12	17	Good Afternoon

18	22	Good Evening
23	24	Good Night

Note that rule tables "summerGreeting" and "winterGreeting" do not have technical information at all - Java snippets and a signature are defined only once and reside in the template-table "defineGreeting".

Along with the keyword "**template**" you may use other keywords:

- **implements**
- **implement**
- **extends**
- **extend**

We will refer to these rule tables created based on a template as "*template implementations*".

Simple templates require that the extended tables should have exactly the same condition and action columns.

Defining Rules based on Templates

When many rule tables are created based on the same rule template, it could be inconvenient to keep the same default rules in all extended tables. As an alternative you may add the rules tables based on the same template. The location of the default rules depends on the types of your rules tables.

Templates for Single-Hit Rule Tables

Single-hit rule tables usually end their execution when at least one rules is satisfied. However, when conditions in all specified rules are not satisfied then a single-hit table usually uses the **last** rule(s) to specify the default action(s). The rules from the template will be executed **after** the actual rules defined inside the template implementation.

Let's consider an example. We have shown that without modification, the rule tables above would not produce any greeting if the parameter "hour" is outside of the interval [0;24]. Instead of adding the same error message in both "summer" and "winter" rules, we could do the following:

-   make our "defineGreeting" template a single-hit table by changing a return type from "void" to "String"
-   - add the default reaction to the error in "hour" directly to the template:

Rules String **defineGreeting**(App app, int hour)		
C1		A1
min <= hour && hour <= max		app.greeting = greeting; return greeting;
int min	int max	String greeting
Hour From	**Hour To**	**Set Greeting**
		ERROR: Invalid Hour

Signature now returns String
Conditions and Actions identifiers
"return greeting;" has been added
Parameter types and names
Business names for conditions and actions
This rule will be added at the end of all template implementations tables. The error message will be return instead of a greeting when all other rules fail.

A template for single-hit tables could include more than one rule with different conditions - they all will be added at the end of the template implementation tables to execute different default actions.

Templates for Multi-Hit Rule Tables

Multi-hit rule tables execute all their rules that are satisfied, allowing rules overrides. However, when conditions in all specified rules are not satisfied then a multi-hit table usually uses the first (!) rules to specify the default action. The rules from the template will be executed **before** the actual rules defined inside the extended tables.

Let's consider an example. You may notice that the rules tables above would not produce any greeting if the parameter "hour" is outside of the interval [0;24].

Let's assume that in this case we want to always produce the default greeting "How are you". To do this, simply add one default rule directly to the template:

Rules void **defineGreeting**(App app, int hour)		
C1		A1
min <= hour && hour <= max		app.greeting = greeting;
int min	int max	String greeting
		How are you

This rule will be added at the beginning of all template implementations. This greeting will be produced if all other rules in the rule tables fail

A template for multi-hit tables could include more than one default rule each with different conditions - they all will be added to the beginning of the template implementation tables and will execute different default actions.

Partial Template Implementation

Usually template implementation tables have exactly the same structure as the rule templates they extend. However, sometimes it is more convenient to build your own rule table that contains only some conditions and actions from already predefined rule templates. This is especially important when a library of rule templates for a certain type of business is used to create a concrete rules-based application. How can this be achieved?

The template implementation table uses its second row to specify the names of the used conditions and actions from the template. Let's consider an example. The DebtResearchRules from the standard OpenRules® example "Loan Origination" may be used as the following template:

Rules void **DebtResearchRules**(LoanRequest loan, Customer c)									
C1	C2		C3	C4		C5		C6	C7	A1
c.mortgageHolder.equals(YN)	c.outsideCreditScore>min && c.outsideCreditScore<=max		c.loanHolder.equals(YN)	op.compare(c.creditCardBalance,value)		op.compare(c.educationLoanBalance,value)		contains(rates,c.internalCreditRating)	c.internalAnalystOpinion.equals(level)	loan.debtResearchResult = level;
String YN	int min	int max	String YN	Operator op	int value	Operator op	int value	String[] rates	String level	String level
IF Mortgage Holder	AND Outside Credit Score		AND Loan Holder	AND Credit Card Balance		AND Education Loan Balance		AND Internal Credit Rating	AND Internal Analyst Opinion	THEN Debt Research Recommendations
	Min	Max		Oper	Value	Oper	Value			

We may create a rule table that implements this template using only conditions C1, C2, C5, C6 and the action A1:

Rules **MyDebtResearchRules** template **DebtResearchRules**									
C1	C2		C5		C6				A1
IF Mortgage Holder	AND Outside Credit Score		AND Education Loan Balance		AND Internal Credit Rating				THEN Debt Research Recommendations
	Min	Max	Oper	Value					
Yes									High
No	100	550							High
No	550	900							Mid
No	550	900	>	0					High
No	550	900	<=	0	A	B	C		High
No	550	900	<=	0	D	F			Mid
No	550	900							Low
No	550	900	<=	0					Low
No	550	900	>	0	D	F			High
No	550	900	>	0	A	B	C		Low

The additional second row specifies which conditions and actions from the original template are selected by this rule table. The order of conditions and actions may be different from the one defined in the template. Only names like "C2", "C6", and "A1" should be the same in the template and in its implementation. It is preferable to use unique names for conditions and actions inside templates. If there are duplicate names inside templates the first one

(from left to right) will be selected. You may create several columns using the same condition and/or action names.

Templates with Optional Conditions and Actions

There is another way to use optional conditions and actions from the templates. If the majority of the template implementations do not use a certain condition from the template, then this condition may be explicitly marked as optional by putting the condition name in brackets, e.g. "[C3]" or "[Conditon-5]". In this case *it is not necessary to use the second row* to specify the selected conditions in the majority of the extended tables. For example, let's modify the DebtResearchRules template making the conditions C3, C4, and C7 optional:

Rules void **DebtResearchRules**(LoanRequest loan, Customer c)							
C1	C2	**[C3]**	**[C4]**	C5	C6	**[C7]**	A1

Now we can implement this template as the following rule table without the necessity to name all of the conditions and actions in the second row:

Rules MyDebtResearchRules template DebtResearchRules										
IF Mortgage Holder	**AND Outside Credit Score**		**AND Education Loan Balance**		**AND Internal Credit Rating**					**THEN Debt Research Recommendations**
	Min	**Max**	**Oper**	**Value**						
Yes										**High**
No	100	550								**High**
No	550	900								**Mid**
No	550	900	>	0						**High**
No	550	900	<=	0	A	B	C			**High**
No	550	900	<=	0	D	F				**Mid**
No	550	900								**Low**
No	550	900	<=	0						**Low**
No	550	900	>	0	D	F				**High**
No	550	900	>	0	A	B	C			**Low**

However, a template implementation that does want to use optional conditions will have to specify them explicitly using the second row:

Rules MyDebtResearchRules template DebtResearchRules											
C1	C2		C3	C4		C5		C6			A1
IF Mortgage Holder	AND Outside Credit Score		AND Loan Holder	AND Credit Card Balance		AND Education Loan Balance		AND Internal Credit Rating			THEN Debt Research Recommendations
	Min	Max		Oper	Value	Oper	Value				
Yes											High
No	100	550									High
No	550	900	Yes	<=	0						Mid
No	550	900	Yes	>	0	>	0				High
No	550	900	Yes	>	0	<=	0	A	B	C	High
No	550	900	Yes	>	0	<=	0	D	F		Mid
No	550	900	No	>	0						Low

Similarly, optional actions may be marked as [A1]" or "[Action3]".

Implementation Notes:

o   Rule templates are supported for both <u>vertical and horizontal</u> rule tables.
o   The keywords "**extends**" or "**implements**" may be used instead of the keyword "**template**"
o   Template implementations cannot be used as templates themselves.

Templates for the Default Decision Tables

The rule tables of the type "DecisionTable" are implemented using several templates located in the following files inside the configuration project "openrules.config":

-   **DecisionTemlates.xls**: contains the following rule templates and methods for the decision tables:
    o   `DecisionTemplate(Decision decision)`: a template for the tables of type "`Decision`"
    o   `initializeDecision()`: the method that initializes the current decision
    o   `decision()`: the method that returns the current decision

- o `getGlossary()`: the method that returns the glossary
- o `getDecisionObject(String nameofBusinessConcept)`: the method that returns a business object associated with the `BusinessConcept`
- o `isTraceOn()`: returns `true` if the tracing of the decision is on
- o `DecisionObjectTemplate(Decision decision)`: a template for the table of the type "`DecisionObject`"
- o `GlossaryTemplate(Decision decision)`: a template for the table of type "`Glossary`"
- o Methods that return values of decision variables based on their names:
  - `int getInt(String name)`
  - `double getReal(String name)`
  - `String getString(String name)`
  - `Date getDate(String name)`
  - `boolean getBool(String name)`
- o Methods that set values of decision variables based on their names:
  - `void getInt(String name, int value)`
  - `void getReal(String name, double value)`
  - `void getString(String name, String value)`
  - `void getDate(String name, Date value)`
  - `void getBool(String name, Boolean value)`
- o Comparison methods that compare a decision variable with a given "name", against a given "value", or another decision variable using a given operator, "op":
  - `boolean compareInt(String name, String op, int value)`
  - `boolean compareInt(String name1, String op, String name2)`
  - `boolean compareReal(String name, String op, double value)`
  - `boolean compareReal(String name1, String op, String name2)`

- boolean compareBool(String name, String op, boolean value)
- boolean compareBool(String name1, String op, String name2)
- boolean compareDate(String name, String op, Date date)
- boolean compareDate(String name1, String op, String name2)
- boolean compareString(String name, String op, String value)
- boolean compareDomain(String name, String op, String domain)

o the Environment table that includes the following references:

- DecisionTable${OPENRULES_MODE}Templates.xls: where ${OPENRULES_MODE} is an environment variable that has one of the following values:

  - **Execute** – the default value for Decision Table execution templates
  - **Validate** –for Decision Table validation templates
  - **Solve** – for execution of decision models using Rule Solver.

- DecisionTableExecuteTemplates.xls: templates for execution
- DecisionTableValidateTemplates.xls: templates for validation

- **DecisionTableExecuteTemplates.xls**: contains the following rule templates:

  o DecisionTableTemplate(): a template for execution of the single-hit tables of the type "DecisionTable"

  o DecisionTable1Template(): a template for execution of the multi-hit tables of the type "DecisionTable1"

  o DecisionTable2Template(): a template for execution of the rule sequence tables of the type "DecisionTable2"

- o customInitializeDecision(): the method that can be used for initialization of custom objects
  - **DecisionTableValidateTemplates.xls**: contains the following rule templates:
    - o DecisionTableTemplate(): a template for validation of the tables of type "DecisionTable" against the domains defined in the glossary
    - o customInitializeDecision(): the method that can be used for the initialization of custom objects.

Decision Templates

The template "DecisionTemplate" contains two mandatory action columns with names "ActionPrint" and "ActionExecute" and three optional columns with the names "Condition", "ConditionAny", and "ActionAny". Here is an example of this template:

RuleSequence void DecisionTemplate()					
[Condition]	[ConditionAny]	ActionPrint	ActionExecute	[ActionAny]	[Message]
getGlossary().compare($COLUMN_TITLE, op, value);	op.compare(value);	String msg = "Decision " + $TABLE_TITLE + ": " + name; if (isTraceOn()) decision().log(msg);	if (method != null) decision().execute(method);		decision().log($TABLE_TITLE + ": " + message);
Oper op String value	Oper op boolean value	String name	String method	Object value	String message
Decision Variable	**Dynamic Condition**	**Decisions**	**Execute Decision or Table**	**Title for Action Any**	**Message**

Because you can use the same column "Condition" or "ConditionAny" many times in your own decision and sub-decision tables, you may create tables of type "Decision" that are based on this template with virtually unlimited complexity.

Three Major Decision Table Templates

The template "DecisionTableTemplate" serves as a template for all standard decision tables. All columns in this template are conditional meaning their names are always required. Here is an example of the first two rows of this template:

Rules String **DecisionTableTemplate**()							
[Condition]	[ConditionAny]	[If]	[Conclusion]	[Action]	[ActionAny]	[Then]	[Message]

The actual DecisionTable template is being upgraded with new OpenRules release and is much larger. Please look at the latest this and other decision table templates in the file "openrules.config/ DecisionTableExecuteTemplates.xls".

The template "DecisionTable1Template" serves as a template for all decision tables of type "DecisionTable1". Here is an example the first two rows of this template:

Rules void **DecisionTable1Template**()							
[Condition]	[ConditionAny]	[If]	[Conclusion]	[Action]	[ActionAny]	[Then]	[Message]

The template "DecisionTable2Template" serves as a template for all decision tables of type "DecisionTable2". Here is an example the first two rows of this template:

RuleSequence void **DecisionTable2Template**()							
[Condition]	[ConditionAny]	[If]	[Conclusion]	[Action]	[ActionAny]	[Then]	[Message]

You can use all these columns as many times as you wish when you may create concrete decision tables based on these templates. Please check the file "DecisionTableExecuteTemplates.xls" in your standard configuration project "openrules.config" to see the latest version of the decision table templates.

Customization

Customizing Default Decision Tables

A user may move the above files from "openrules.config" to different locations and modify the decision table templates (and possible other templates). For example, to have different types of messaging inside a custom decision, a user may add two more columns to the template "DecisionTableTemplate":

- **Warning**: similar to Message but can use a different log for warning only
- **Errors**: similar to Message but can use a different log for errors only.

Adding Custom Decision Tables

Users may add their own decision tables with conditions and actions specific to their applications by defining their own keywords by simply extending the keyword "DecisionTable" with they own identifier. For example, a user may add a new decision table type called "DecisionTableMy" by defining the proper custom conditions and actions inside a template with the name "DecisionTableMyTemplate". The standard installation includes a project "DecisionCustom" that demonstrates a custom decision table called "DecisionTableCustom" created based on a project-specific template "DecisionTableCustomTemplate". This template is placed in the project file "DecisionTableCustomTemplates.xls".

Adding Custom Methods to Decision and Decision Runs

The file "DecisionTemplates.xls" contains the default methods:

- `customInitializeDecision`
- `customInitializeDecisionRun`

that may be replaced by your own methods. For example, rewriting the method "`customInitializeDecision`" allows a user to initialize custom objects. These and other methods are described below. For a good example of customization look at the file "DecisionTableSolveTemplates.xls" that is used by Rule Solver instead of the file "DecisionTableExecuteTemplates.xls". Contact support@openrules.com if you need help with more complex customization of the decision templates.

OPENRULES® API

OpenRules® provides an Application Programming Interface (API) that defines a set of commonly-used functions:

- Creating a rule engine associated with a set of Excel-based rules
- Creating a decision associated with a set of Excel-based rules
- Executing different rule sets using application specific business objects
- Creating a web session and controlling client-server interaction.

OpenRulesEngine API

OpenRulesEngine is a Java class provide by OpenRule® to execute different rule sets and methods specified in Excel files using application-specific business objects. OpenRulesEngine can be invoked from any Java application using a simple Java API or a standard JSR-94 interface.

Engine Constructors

OpenRulesEngine provides an interface to execute rules and methods defined in Excel tables. You can see examples of how OpenRulesEngine is used in basic rule projects such as HelloJava, DecisionHellJava, HelloJsr94 and web applications such as HelloJsp, HelloForms, and HelloWS. To use OpenRulesEngine inside your Java code you need to add an import statement for `com.openrules.ruleengine.OpenRulesEngine` and make sure that `openrules.all.jar` is in the classpath of your application. This jar and all 3rd party jar-files needed for OpenRules® execution can be found in the subdirectory `openrules.config/lib` of the standard OpenRules® installation. You may create an instance of OpenRulesEngine inside of your Java program using the following constructor:

```
    public OpenRulesEngine(String xlsMainFileName)
```

where xlsMainFileName parameter defines the location for the main xls-file. To specify a file location, OpenRules® uses an **URL pseudo-protocol notation** with prefixes such as *"file:", "classpath:", "http://", "ftp://", "db:"*, etc. Typically, your main xls-file Main.xls is located in the subdirectory "rules/main" of your Java project. In this case, its location may be defined as "file:rules/main/Main.xls". If your main xls-file is located directly in the project classpath, you may define its location as "classpath:Main.xls". Use a URL like

```
http://www.example.com/rules/Main.xls
```

when `Main.xls` is located at a website. All other xls-files that can be invoked from this main file are described in the table "Environment" using include-statements.

You may also use other forms of the OpenRulesEngine constructor. For example, the constructor

```
OpenRulesEngine(String xlsMainFileName, String methodName)
```

allows you to also define the main method from the file xlsMainFileName that will be executed during the consecutive runs of this engine.

Here is a complete example of a Java module that creates and executes a rule engine (see HelloJava project):

```
package hello;
import com.openrules.ruleengine.OpenRulesEngine;
public class RunHelloCustomer {
  public static void main(String[] args) {
    String fileName = "file:rules/main/HelloCustomer.xls";
    String methodName = "helloCustomer";
    OpenRulesEngine engine = new OpenRulesEngine(fileName);
    Customer customer = new Customer();
    customer.setName("Robinson");
    customer.setGender("Female");
```

```java
    customer.setMaritalStatus("Married");


    Response response = new Response();
    Object[] objects = new Object[] { customer, response };
    engine.run(methodName,objects);
    System.out.println("Response: " +
        response.getMap().get("greeting") + ", " +
        response.getMap().get("salutation") +
        customer.getName() + "!" );
  }
}
```

As you can see, when an instance "engine" of OpenRulesEngine is created, you can create an array of Java objects and pass it as a parameter of the method "run".

Engine Runs

The same engine can run different rules and methods defined in its Excel-files. You may also specify the running method using

```java
    setMethod(String methodName);
```

or use it directly in the engine run:

```java
    engine.run(methodName,businessObjects);
```

If you want to pass to OpenRulesEngine only one object such as "customer", you may write something like this:

```java
    engine.run("helloCustomer",customer);
```

If you do not want to pass any object to OpenRulesEngine but expect to receive some results from the engine's run, you may use this version of the method "run":

```java
    String[] reasons = (String[]) engine.run("getReasons");
```

Undefined Methods

OpenRulesEngine checks to validate if all Excel-based tables and methods are actually defined. It produces a syntax error if a method is missing. Sometimes, you want to execute a rule method/table from an Excel file but only if this method is actually present. To do this, you may use this version of the method "run":

```
boolean mayNotDefined = true;
engine.run(methodName, businessObjects, mayNotDefined);
```

In this case, if the method "methodName" is not defined, the engine would not throw a usual runtime exception "*The method <name> is not defined*" but rather will produce a warning and will continue to work. The parameter "mayNotDefined" may be used similarly with the method "run" with one parameter or with no parameters, e.g.

```
engine.run("validateCustomer", customer, true);
```

How to invoke rules from other rules if you do not know if these rules are defined? It may be especially important when you use some predefined rule names in templates. Instead of creating an empty rules table with the needed name, you want to use the above parameter "mayNotDefined" directly in Excel. Let's say you need to execute rules tables with names such as "NJ_Rules" or "NY_Rules" from another Excel rules table but only if the proper state rules are actually defined. You may do it by calling the following method from your rules:

Method void runStateRules(OpenRulesEngine engine, Customer customer, Response response)
String methodName = customer.state + "_Rules";
Object[] params = new Object[2];
params[0] = customer;
params[1] = response;
engine.run(methodName, params, true);

We assume here that all state-specific rules ("NJ_Rules", "NY_Rules", etc.) have two parameters, "customer" and "response". To use this method you need to pass the current instance of OpenRulesEngine from your Java code to your main Excel file as a parameter "engine". If you write an OpenRules Forms application, this instance of the OpenRulesEngine is always available as `dialog().getEngine()`, otherwise you have to provide access to it, e.g. by attaching it to one of your own business objects such as Customer.

By default OpenRules will produce a warning when the required Excel rules table or method is not available. You may suppress such warnings by calling:

```
engine.turnOffNotDefinedWarning();
```

Accessing Password Protected Excel Files

Some Excel workbooks might be encrypted (protected by a password) to prevent other people from opening or modifying these workbooks. Usually it's done using Excel Button and then **Prepare** plus **Encrypt Document**. OpenRules Engine may access password-protected workbooks by calling the following method just before creating an engine instance:

```
OpenRulesEngine.setCurrentUserPassword("password");
```

Instead of "password" you should use the actual password that protects your main and/or other Excel files. Only one password may be used by all protected Excel files that will be processed by one instance of the OpenRulesEngine created after this call. This call does not affect access to unprotected files. The standard project "HelloJavaProtected" provides an example of the protected Excel file - use the word "password" to access the file "HelloCustomer.xls".

Note. The static method "*setCurrentUserPassword*" of the class OpenRulesEngine actually sets the BIFF8 encryption/decryption password for the current thread. The use of a "null" string will clear the password.

Engine Attachments

You may attach any Java object to the OpenRulesEngine using methods `setAttachment(Object attachment)` and *getAttachment()*.

Engine Version

You may receive a string with the current version number of the OpenRulesEngine using the method *getVersion()*.

Dynamic Rules Updates

If a business rule is changed, OpenRulesEngine automatically reloads the rule when necessary. Before any engine's run, OpenRulesEngine checks to determine if the main Excel file associated with this instance of the engine has been changed. Actually, OpenRulesEngine looks at the latest modification dates of the file xlsMainFileName. If it has been modified, OpenRulesEngine re-initializes itself and reloads all related Excel files. You can shut down this feature by executing the following method:

```
engine.setCheckRuleUpdates(false);
```

Decision API

Decision Example

OpenRules® provides a special API for decision execution using the Java class "Decision". The following example from the standard project "Decision1040EZ" demonstrates the use of this API.

```java
public class Main {

    public static void main(String[] args) {
        String fileName = "file:rules/main/Decision.xls";
        OpenRulesEngine engine =
            new OpenRulesEngine(fileName);
        Decision decision =
            new Decision("Apply1040EZ",engine);
        DynamicObject taxReturn =
                (DynamicObject) engine.run("getTaxReturn");
```

```
        engine.log("=== INPUT:\n" + taxReturn);
        decision.put("taxReturn",taxReturn);
        decision.execute();
        engine.log("=== OUTPUT:\n" + taxReturn);
    }
}
```

Here we first created an instance engine of the class OpenRulesEngine and used it to create an instance decision of the class Decision. We used the engine to get an example of the object taxReturn that was described in Excel data tables:

```
DynamicObject taxReturn =
            (DynamicObject) engine.run("getTaxReturn");
```

Then we added this object to the decision:

```
decision.put("taxReturn",taxReturn);
```

and simply executed decision:

```
decision.execute();
```

The Decision described in "Decision.xls" is supposed to modify certain attributes inside the object decision and objects which were put inside the decision after its execution.

Decision Constructors

The class Decision provides the following constructor:

```
 public Decision(String decisionName, String xlsMainFileName)
```

where "decisionName" is the name of the main table of the type "Decision" and "xlsMainFileName" is the same parameter as in the OpenRulesEngine's constructor that defines a location for the main xls-file.

There is also another constructor:

```
 public Decision(String decisionName, OpenRulesEngine engine)
```

where the parameter OpenRulesEngine engine refers to an already created instance of the OpenRulesEngine as in the above example.

Each decision has an associated object of type `Glossary`. When a decision is created, it first executes the table "glossary" that must be defined in our rules repository. It fills out the glossary, a step that applies to all consecutive decision executions. You may always access the glossary by using the method

```
Glossary glossary = decision.getGlossary();
```

Decision Parameters

The class Decision is implemented as a subclass of the standard Java class HashMap. Thus, you can put any object into the decision similarly as we did above:

```
decision.put("taxReturn",taxReturn);
```

You may access any object previously put into the decision by calling the method get(name) as in the following example:

```
TaxReturn taxReturn = (TaxReturn)decision.get("taxReturn");
```

You may set a special parameter

```
decision.put("trace","Off");
```

to tell your decision to turn off the tracing . You may use "On" to turn it on again.

Decision Runs

After defining decision parameters, you may execute the decision as follows:

```
decision.execute();
```

This method will execute your decision starting from the table of type "Decision" whose name was specified as the first parameter of the decision's constructor.

You may reset the parameters of your decision and execute it again without the necessity of constructing a new decision. This is very convenient for multi-transactional systems where you create a decision once by instantiating its glossary, and then you execute the same decision multiple times but with different parameters. To make sure that it is possible, the Decision's method execute() calls Excel's method "decisionObjects" each time before actually executing the decision.

If you know that the types of decision parameters are changed between different decision runs you may use the following variation of the method "execute":

```
decision.execute(true);
```

The actual execution of "**this**" decision involves engine runs for the following Excel methods (in this order):

```
-   engine.run("decisionObjects",this);
-   engine.run("initializeDecision",this);
-   engine.run("initializeDecisionRun",this);
-   engine.run(this); // run the main decision
-   engine.run("finalizeDecision",this);
```

All these methods are described in the standard file "DecisionTemplates.xls".
The method "initializeDecision" is executed only during the first decision run. It calls the method "customInitializeDecision" that may include an application specific decision initialization.
The method "initializeDecisionRun" is executed during every decision run. It calls the method "customInitializeDecisionRun" that may include a code that is specific for every decision run, e.g. it may analyze the parameters of this run and redefine some decision variables.
The method "finalizeDecision" is executed after the main Excel table of the type "Decision" that was specified in the decision's constructor.

Executing Decision Methods From Excel

There is one more form of this method:

```
decision.execute(String methodName);
```

It is used within Excel when you want to execute another Excel method. It is implemented as follows:

```java
public Object execute(String methodName) {
    return getEngine().run(methodName);
}
```

Decision Glossary

Every decision has an associated business glossary – see <u>above</u>. Glossaries are usually presented in Excel tables that may look like this table "glossary":

Glossary glossary		
Variable	**Business Concept**	**Attribute**
Gender		gender
Date of Birth	Customer	dob
Marital Status		maritalStatus
Greeting		greeting
Salutation	Response	salutation
Current Hour		hour

In large, real-world projects the actual content of business concepts such as the above "Customer" can be defined in external applications using Java-based Business Object Models or they may come from XML files, a database table, etc. The list of attributes inside business objects can be very large and/or to be defined dynamically. In such cases, you do not want to repeat all attributes in your Excel-based glossary and then worry about keeping the glossary synchronized with an IT implementation.

It is possible to programmatically define/extend the definition of the Glossary. For example, we may leave in the Excel's glossary only statically defined business concepts and their variables, e.g. in the above table we may keep only the variables of the concept "Response" and remove all rows related to the concept "Customer". Then in the Java module that creates an object "decision" of the predefined type Decision we may add the following code:

```
Decision decision = new Decision(fileName);
String[] attributes = getCustomerAttributes();
String businessConcept = "Customer";
for (int i = 0; i < attributes.length; i++) {
      String varName = attributes[i].getName();
      decision.getGlossary().put(varName,businessConcept,varName);
}
...
decision.put("customer", customer);
decision.execute();
```

Here we assume that the method getCustomerAttributes() returns the names of attributes defined in the class Customer. The variable name and the attribute name are the same for simplicity - of course you may define them differently.

You may add multiple concepts to the Glossary in a similar way. In all cases keep in mind that the table "Glossary glossary" always has to be present in your Excel repository even when it contain no rows. You also may find that the same method put(variableName, businessConcept, attributeName) of the class Glossary is used in the Glossary Template definition in the standard file "DecisionTemplates.xls".

Business Concepts and Decision Objects

OpenRules® Glossary specifies names of business concepts that contain decision variables. The connection (mapping) between business concepts and actual

objects that implement these concepts (decision objects) is usually specified in the Excel table "decisionObjects" that may look like:

DecisionObject decisionObjects	
Business Concept	**Business Object**
Customer	:= decision.get("customer")
Request	:= decision.get("loanRequests")
Internal	:= internal

The standard mapping is implemented in the DecisionObjectTemplate using the following Glossary's method:

```
void useBusinessObject(String businessConcept, Object object)
```

What if you want to change actual business objects on the fly during the decision execution? You can do it by using the same method inside your Excel rules. For example, you may want to apply the following decision table "EvaluateAssets" for all elements of an array "assets" of a given customer:

DecisionTable EvaluateAsset					
Condition		Condition		Conclusion	
Asset Name		**Asset Status**		**Customer's Assets Status**	
Is One Of	Asset12, Asset21, Asset23	Is	Active	Is	Sufficient

In this case you still may specify the business concept "Asset" in your glossary only once, but you may associate different elements of an array "assets" with the concept Asset multiple times in the loop similar to the one below:

```
Method void evaluateCustomerAssets(Customer customer)
Asset[] assets = customer.getAssets();
customer.customerAssetsStatus = "Insufficient";
for(int i=0; i<assets.length; i++) {
  getGlossary().useBusinessObject("Asset",customer.assets[i]);
  EvaluateAsset();
  if ("Sufficient".equals(customer.customerAssetsStatus))
    return;
}
```

Changing Decision Variables Types between Decision Runs

OpenRules® Glossary does not require a user to specify actual types of the variables - they are automatically defined from the actual types of decision parameters. It allows you to change types of decision parameters between decision runs without necessity to download all rules again. If you know that some attributes corresponding to your decision variables may change their types between different runs of the same decision, you may use the following Decision's method:

```
execute(boolean objectTypesVary)
```

If the parameter "objectTypesVary" is true then before executing the decision, the OpenRulesEngine will re-evaluate the decision's glossary and will reset types of all object attributes based on the actual type of objects passed to the decision as parameters. By default, the parameter "objectTypesVary" is false.

Decision Execution Modes

Before executing a decision you may validate it by setting a special "validation" mode. Here is a code example:

```
String fileName = "file:rules/main/Decision.xls";
System.setProperty("OPENRULES_MODE", "Validate");
Decision decision = new
Decision("DetermineDecisionVariable",fileName);
```

During the validation along with regular syntax check OpenRules® will validate if the values for conditions and actions inside all decision tables correspond to their glossary domains (if they are defined).

As you can see, the system property "OPENRULES_MODE" defines which mode to use. By default this property is set to "Execute". If you create an OpenRulesEngine before creation a Decision, you need to set this property first.

Generating Excel Files with Decision Tables

OpenRules® allows you to generate xls-files with multiple decision tables programmatically by providing the proper Java API. The Java class **DecisionBook** that corresponds to one Excel workbook (or an xls-file) allows you to add OpenRules® decision tables defined in Java. Multiple decision tables can be added to a preliminary created instance of the DecisionBook class. Each new decision table will be placed in a new worksheet of the same workbook. Then you may simply save this decision book as an Excel file.

Example with Explanations

Let's first consider an example provided in the standard OpenRules® installation as the "DecisionWithGeneratedRules" project. In this project we want to run a Java application (GenerateRules.java) to generate the following decision tables in Excel:

DecisionTable DefineGreeting		
If	If	Then
Current Hour	Current Hour	Result
>=0	<=11	Good Morning
>=12	<=17	Good Afternoon
>=18	<=21	Good Evening
>=22	<=24	Good Night

DecisionTable CheckGreeting			
ConditionVarOperValue			Message
<Var> <Oper> <Value>			Message
Result	Is Not	Good Afternoon	Error: Expected Good Afternoon
Result	Is	Good Afternoon	Good Result

Here is the proper Java class GenerateRules.java:

```java
import com.openrules.table.external.DecisionBook;

public class GenerateRules {

  public static void main(String[] args) {

    DecisionBook decisionBook = new DecisionBook();

    decisionBook.addDecisionTable(
        "DefineGreeting",                                       //table
        "DecisionTableTemplate",                                //template
        new String[] { "If", "If", "Then" },                   // labels
        new String[] { "Current Hour","Current Hour","Result" }, //variables
        new String[][] {                                       //rules
            new String[] {">=0","<=11","Good Morning"},
            new String[] {">=12","<=17","Good Afternoon"},
            new String[] {">=18","<=21","Good Evening"},
            new String[] {">=22","<=24","Good Night"}
        }
    );

    decisionBook.addDecisionTable(
        "CheckGreeting",                                        //table name
        "DecisionTableTemplate",                                //template name
        new String[] { "ConditionVarOperValue", "Message" },// labels
        new String[] { "<Var> <Oper> <Value>", "Message" }, //titles
        new String[][] {                                      //rules
            new String[] {"Result","Is Not","Good Afternoon",
                          "Error: Expected Good Afternoon"},
            new String[] {"Result","Is","Good Afternoon", "Good Result"}
        }
    );

    decisionBook.saveToFile("./rules/include/Rules.xls");

  }
}
```

The first statement `DecisionBook decisionBook = new DecisionBook();` simply creates an instance of the class DecisionBook. Then we add two rule tables to this decision book by using `decisionBook.addDecisionTable(…);`

Then you may easily map this Java structure to the above decision table "DefineGreeting". It is created based on the standard template `"DecisionTableTemplate".` Then the strings { `"If"`, `"If"`, `"Then"` } define the selected table columns from this template. The next array of strings { `"Current Hour"`, `"Current Hour"`, `"Result"` } defines the names of decision variables used

in these columns. Then we have a two-dimensional array of strings where each sub-array represents one rule (or the table row) such as

`new String[] {">=0","<=11","Good Morning"}`.

Depending on the column type, instead of the names of the decision variables the column titles may contain any text in plain English. For example, the first column in the second decision table "CheckGreeting" is defined as "`ConditionVarOperValue`", that according to the standard template has 3 sub-columns. The title of this column is defined as "`<Var> <Oper> <Value>`". Note that this title is "merged" while the content of the proper 3 sub-columns is defined using 3 strings such as `"Result","Is Not","Good Afternoon" i`n the unmerged format.

Finally, this decision book is saved to the file "`./rules/include/Rules.xls`" using the method `decisionBook.saveToFile("./rules/include/Rules.xls");`

Formal API

The Java class DecisionBook has a public constructor without parameters and the following public methods:

```
public void addDecisionTable(
   String tableName,       // table name
   String templateName,    // template name
   String[] labels,        // template column labels
   String[] descriptions,  // descriptions or variables
   String[][] rules        //  rules
);
```

This method adds a new decision table to the rule book. The first parameter is the name of the generated decision table (no spaces allowed). The second parameter is the name of the standard OpenRules template that has one of the following values:

- `DecisionTableTemplate` – for regular single-hit decision tables
- `DecisionTable1Template` – for multi-hit decision tables
- `DecisionTable2Template` – for rule sequences (see more)

The third parameter is an array of column labels selected from the proper template. The fourth parameter is an array of names that corresponds to the column type – it could be either a name of the decision variable or a title of the proper column. The fifth parameter is a two-dimensional array of strings where each sub-array represents one rule (or the decision table row).

The method

```
public void saveToFile(String xlsFile);
```

saves this decision book in the Excel file whose name is provided as a parameter.

The method

```
public int getNumberOfRuleTables();
```

returns a number of decision tables currently added to the decision book. Please note that the proper Excel file will contain a separate worksheet for each decision table.

Logging API

OpenRules® provides an API for decision logging. Assuming that "decision" is an instance of the class Decision, you may use the following logging methods:

- To log (print) any text string you may write

  decision.log(text);

- To memorize the execution log you may write

  decision.saveRunLog(true);
  decision.execute();

  Then all log-statements produced during this decision run will be saved internally.

- You may get this saved log as follows:
  Vector<String> log = decision.getRunLog();

You may print the saved log by the method
decision.printSavedRunLog()

or you may save it into a file by the method

decision.printSavedRunLog(filename).

This feature is very useful when your application wants to show the good results of the decision execution but also need to show the errors in the user-defined decision model - see, for example, how it is done in the latest <u>remote evaluation version</u>.

JSR-94 Implementation

OpenRules® provides a reference implementation of the JSR94 standard known as Java Rule Engine API (see <u>http://www.jcp.org/en/jsr/detail?id=94</u>). The complete OpenRules® installation includes the following projects:

JSR-94 Project	Description
lib.jsr94	This project contains the standard jsr94-1.0 library
com.openrules.jsr94	This is an OpenRules®'s reference implementation for the JSR94 standard and includes the source code. It uses OpenRulesEngine to implement RuleExecutionSet
HelloJsr94	This is an example of using JSR94 for simple rules that generate customized greetings
HelloJspJsr94	HelloJspJsr94 is similar to HelloJsp but uses the OpenRules® JSR-94 Interface to create and run OpenRulesEngine for a web application.

Multi-Threading

OpenRulesEngine is thread-safe and works very efficiently in multi-threaded environments supporting real parallelism. OpenRulesEngine is stateless, which allows a user to create *only one* instance of the class OpenRulesEngine, and then

share this instance between different threads. There are no needs to create a pool of rule engines. A user may also initialize the engine with application data common for all threads, and attach this data directly to the engine using the methods `setAttachment(Object attachment)`. Different threads will receive this instance of the rule engine as a parameter, and will safely run various rules in parallel using the same engine.

The complete OpenRules® installation includes an example "HelloFromThreads" that demonstrates how to organize a parallel execution of the same OpenRulesEngine's instance in different threads and how to measure their performance.

INTEGRATION WITH JAVA AND XML

Java Classes

OpenRules® allows you to externalize business logic into xls-files. However, these files can still use objects and methods defined in your Java environment. For example, in the standard example "`RulesRepository`" all rule tables deal with the Java object Appl defined in the Java package myjava.package1. Therefore, the proper Environment table inside file Main.xls (see above) contains a property "`import.java`" with the value "`myjava.package1.*`":

Environment	
import.java	myjava.packA1.*
include	SubCategoryA1/RulesA11.xls
	SubCategoryA1/RulesA12.xls

```
The property "import.java" allows you to define all classes
from the package following the standard Java notation, for
example "hello.*".  You may also import only the specific
class your rules may need, as in the example above.  You can
define a separate property "import.java" for every Java
package used or merge the property "import.java" into one cell
with many rows for different Java packages.  Here is a more
complex example:
```

Environment	
import.static	com.openrules.tools.Methods
import.java	my.bom.*
	my.impl.*
	my.inventory.*
	com.openrules.ml.*
	my.package.MyClass
	com.3rdparty.*
include	../include/Rules1.xls
	../include/Rules2.xls

Naturally the proper jar-files or Java classes should be in the classpath of the Java application that uses these rules.

If you want to use static Java methods defined in some standard Java libraries and you do not want to specify their full path, you can use the property "import.static". The static import declaration imports static members from Java classes, allowing them to be used in Excel tables without class qualification. For example, many OpenRules® sample projects use static methods from the standard Java library *com.openrules.tools* that includes class Methods. So, many Environment tables have property "import.static" defined as "com.openrules.tools.Methods". This allows you to write

```
out("Rules 1")
```

instead of

```
Methods.out("Rules 1")
```

XML Files

Along with Java classes, OpenRules® tables can use objects defined in XML files. For example, the standard sample project HelloXMLCustomer uses an object of type Customer defined in the file Customer.xml located in the project classpath:

```
<Customer
    name="Robinson"
```

```
        gender="Female"

        maritalStatus="Married"

        age="55"

/>
```

The xls-file, <u>HelloXmlCustomer.xls</u>, that deals with this object includes the following Environment table:

Environment	
import.static	com.openrules.tools.Methods
import.schema	**classpath:/Customer.xml**
import.java	hello.Response
include	include/HelloRules.xls

The property, "import.schema", specifies the location of the proper xml-file, in this case "classpath:/Customer.xml". Of course, you can use any other location in your local file system that starts with the prefix "file:". This example also tells you that this Excel file uses:

1. static Java methods defined in the standard OpenRules® package "com.openrules.tools.Methods"
2. xml-file "classpath:/Customer.xml"
3. Java class "Response" from a package "hello"
4. include-file "HelloRules.xls" which is located in the subdirectory "include" of the directory where the main xls file is located.

The object of the type "Customer" can be created using the following API:

```
Customer customer = Customer.load("classpath:/Customer.xml");
```

You may use more complex structures defined in xml-files. For example, the project HelloXMLPeople uses the following xml-file:

```
<?xml version="1.0" encoding="UTF-8"?>
<People type="Array of Person(s)">
  <Person name="Robinson" gender="Female" maritalStatus="Married"
age="55" />
          <Person name="Robinson" gender="Female"
maritalStatus="Single" age="23" />
```

```
            <Person name="Robinson" gender="Male"
maritalStatus="Single" age="17" />
            <Person name="Robinson" gender="Male"
maritalStatus="Single" age="3" />
</People>
```

The method that launches greeting rules for every `Person` from an array `People` is defined as:

```
Method void helloPeople()

int hour = Calendar.getInstance().get(Calendar.HOUR_OF_DAY);
App app = new App();
defineGreeting(hour, app);
// define and greet People from the XML file People.xml
People people = People.load("classpath:/People.xml");
for(int i = 0; i < people.Person.length; ++i)
{
    People.Person person = people.Person[i];
    defineSalutation(person,app);
    //greet Person
    System.out.println(app.greeting+", "+app.salutation+" "+person.name+"!");
}
```

DATA MODELING

OpenRules® includes an ability to define new data/object types and creates the objects of these types directly in Excel. It allows business analysts to do Rule Harvesting by defining business terms and facts without worrying about their implementation in Java, C#, or XML. It also provides the ability to **test** the business rules in a pre-integrated mode. To do standalone rule testing, a designer of rules and forms specifies his/her own data/object types as Excel tables and creates instances of objects of these types passing them to the rule tables. We describe how to do it in the sections below.

There is one more important reason why a business or even a technical specialist may need data modeling abilities without knowing complex software development techniques. In accordance with the SOA principle of loosely coupled

services, rule services have to specify what they actually need from the objects defined in an external environment. For example, if an object "Insured" includes attributes related to a person's military services, it does not require that all business rules that deal with the insured be interested in those attributes. Such encapsulation of only the essential information in the Excel-based data types, together with live process modeling, allows OpenRules® to complete the rule modeling cycle without leaving Excel.

OpenRules® provides the means to make business rules and forms independent of a concrete implementation of such concepts. The business logic expressed in the decision tables should not depend on the implementation of the objects these rules are dealing with. For example, if a rule says: "If driver's age is less than 17 then reject the application" the only thing this business rule should "know" about the object "driver" is the fact that it has a property "age" and this property has a type that support a comparison operator "<" with an integer. It is a question of configuration whether the Driver is a Java class or an XML file or a DB table from a legacy system. Similarly, if a form has an input field "Driver's Age", the form should be able to accept a user's input into this field and automatically convert it into the proper object associated with this field independently of how this object was implemented.

Thus, OpenRules® supports data source independent business rules (decision tables) and web forms. Your business rules can work with an object of type Customer independently of the fact that this type is defined as a Java class, as an XML file or as an Excel table. You can see how it can be done using examples HelloJava, HelloXML, and HelloRules from the OpenRules®'s standard installation. It is a good practice to start with Excel-based data types. Even if you later on switch to Java classes of other data types, you would always be able to reuse Excel-based types for standalone testing of your rules-based applications.

Datatype and Data Tables

OpenRules® allows a non-technical user to represent different data types directly in Excel and to define objects of these types to be used as test data. Actually, it provides the ability to

create Excel-based Data Models, which, in turn, define problem specific business terms and facts. At the same time, a data model can include data types specified outside Excel, for example in Java classes or in XML files. Here is an example of a simple data type "PersonalInfo":

Datatype PersonalInfo	
String	id
String	firstName
String	middleInitial
String	lastName
String	address
String	appartment
String	city
String	state
String	zipCode

Now we can create several objects of this type "`PersonalInfo`" using the following data table:

Data PersonalInfo personalInformation			
id	ID	He	She
firstName	First Name	John	Mary
middleInitial	Middle Initial	N.	A.
lastName	Last Name	Smith	Smith
address	Address	25 Maple Street	
appartment	appartment	Apt. 3C	
city	City	Edison	
state	State	NJ	
zipCode	ZipCode	08840	

We can reference to these objects inside rules or forms as in the following snippets:

```
out(personalInformation["He"].lastName);
if (personalInformation["She"].state.equals("NJ")) ...
```

You may use one datatype (such as `PersonalInfo`) to define a more complex *aggregate datatype,* like TaxReturn in this example:

Datatype TaxReturn	
PersonalInfo	taxPayer
PersonalInfo	spouse
boolean	marriedFillingJointly
boolean	claimedAsDependent
boolean	spouseClaimedAsDependent
double	wages
double	taxableInterest
double	unemploymentCompensation
double	adjustedGrossIncome
double	dependentAmount
double	taxableIncome
double	taxWithheld
double	earnedIncomeCredit
double	totalPayments
double	tax
double	refund

You may even create an object of the new composite type "TaxReturn" using references to the objects "He" and "She" as in this example:

Data TaxReturn **taxReturns**					
taxPayer	spouse	wages	taxableInterest	taxWithheld	earnedIncomeCredit
>personalInformation	>personalInformation				
TaxPayer	Spouse	Wages	Taxable Interest	Tax Withheld	Earned Income Credit
He	She	32026	1450	4530	230

Now we can reference these objects from inside rules or forms as in the following snippet:

```
out(taxReturn[0].taxPayer.lastName);
```

The above tables may remind you of traditional database tables simply presented in Excel. While these examples give you an intuitive understanding of OpenRules® Datatype and Data tables, the next sections will provide their formal descriptions.

You may use a type of table "**Variable**". These tables are similar to the Data tables but instead of arrays of variables they allow you to create separate instances of objects directly in Excel files. Here is a simple example:

Variable Customer mary			
name	age	gender	maritalStatus
Name	Age	Gender	Marital Status
Mary Brown	5	Female	Single

The variable "mary" has type Customer and can be used inside rules or passed back from an OpenRulesEngine to a Java program as a regular Java object. As usual, the object type Customer can be defined as a Java class, an Excel Datatype, or an xml structure.

How Datatype Tables Are Organized

Every Datatype table has the following structure:

Datatype tableName	
AttributeType1	AttrubuteName1
AttributeType2	AttrubuteName2
..	..
..	..

The first "signature" row consists of two merged cells and starts with the keyword "Datatype". The "tableName" could be any valid one word identifier of the table (a combination of letters and numbers). The rows below consist of two cells with an attribute type and an attribute name. Attribute types can be the basic Java types:

- `boolean`
- `char`
- `int`
- `double`
- `long`
- `String (java.lang.String)`
- `Date  (java.util.Date)`

You may also use data types defined:

-   in other Excel Datatype tables
-   in any Java class with a public constructor with a single parameter of the type String
-   as one-dimensional arrays of the above types.

The datatype "PersonalInfo" gives an example of a very simple datatype. We can define another datatype for a social security number (SSN):

Datatype SSN	
String	ssn1
String	ssn2
String	ssn3

and add a new attribute of this type to the datatype "PersonalInfo":

Datatype PersonalInfo	
String	id
String	firstName
String	middleInitial
String	lastName
String	address
String	appartment
String	city
String	state
String	zipCode
SSN	ssn

It is interesting that these changes do not affect the already existing data objects defined above (like `personalInformation["He"]`) - their SSNs just will not be defined.

Implementation Restriction. Make sure that the very first attribute in a Datatype table has type String or your own type but not a basic Java type like int.

The following example demonstrates how to create a Data table for a Datatype that includes one-dimensional arrays:

Datatype Order	
String	**number**
String[]	**selectedItems**
String[]	**offeredItems**
double	**totalAmount**
String	**status**

Here is an example of the proper Data table:

Data Order orders			
number	selectedItems	totalAmount	status
Number	**Selected Items**	**Total Amount**	**Status**
6P-U01	INTRS-PGS394 INTRS-PGS456 Paste-ARMC-2150	3700	In Progress

You may also present the same data in the following way:

Data Order orders				
number	selectedItems			totalAmount
Number	**Selected Items**			**Total Amount**
	Item 1	**Item 2**	**Item 3**	
6P-U01	INTRS-PGS394	INTRS-PGS456	Paste-ARMC-2150	3700

How Data Tables Are Organized

Every Datatype table has a vertical or horizontal format. A typical vertical Data table has the following structure:

Data datatypeName tableName			
AttributeName1 from "datatypeName"	AttributeName2 from "datatypeName"	AttributeName3 from "datatypeName"	...
Display value of the	**Display value of the**	**Display value of the**	**...**

AttributeName1	AttributeName2	AttributeName3	
data	data	data	...
data	data	data	...
...	...	...	...

The first "signature" row consists of two merged cells and starts with the keyword "Data". The next word should correspond to a known datatype: it can be an already defined Excel Datatype table or a known Java class or an XML file. The "tableName" is any one word valid identifier of the table (a combination of letters and numbers).

The second row can consists of cells that correspond to attribute names in the data type "datatypeName". It is not necessary to define all attributes, but at least one should be defined. The order of the columns is not important.

The third row contains the display name of each attribute (you may use unlimited natural language).

All following rows contain data values with types that correspond to the types of the column attributes.

Here is an example of the Data table for the datatype "PersonalInfo" defined in the previous section (with added SSN):

Data PersonalInfo personalInformation						
id	firstName	middleInitial	lastName	ssn.ssn1	ssn.ssn2	ssn.ssn3
ID	First Name	Middle Initial	Last Name	SSN1	SSN2	SSN3
He	John	N.	Smith	164	86	2298
She	Mary	A.	Smith	627	35	1293

The table name is "personalInformation" and it defines an array of objects of the type PersonalInfo. The array shown consists only of two elements personalInformation[0] for John and personalInformation[1] for Mary. You may add as many data rows as necessary.

The attributes after the SSN attribute have not been defined. Please, note that the references to the <u>aggregated</u> data types are defined in a natural way (`ssn.ssn1, ssn.ssn2, ssn.ssn3`) using the dot-convention.

As you can see from this example, the vertical format may not be very convenient when there are many attributes and not so many data rows. In this case, it could be preferable to use a horizontal format for the data tables:

Data datatypeName tableName					
AttributeName1 from "datatypeName"	Display value of the AttributeName1	data	data	data	...
AttributeName2 from "datatypeName"	Display value of the AttributeName2	data	data	data	...
AttributeName3 from "datatypeName"	Display value of the AttributeName3	data	data	data	...
...	..	...	...	...	...

Here is how our data table will look when presented in the horizontal format:

Data PersonalInfo personalInformation			
id	ID	He	She
firstName	First Name	John	Mary
middleInitial	Middle Initial	N.	A.
lastName	Last Name	Smith	Smith
ssn.ssn1	SSN1	164	627
ssn.ssn2	SSN2	86	35
ssn.ssn3	SSN3	2298	1293
address	Address	25 Maple Street	
appartment	appartment	Apt. 3C	
city	City	Edison	
state	State	NJ	
zipCode	ZipCode	08840	

Predefined Datatypes

OpenRules® provides predefined Java classes to create data tables for arrays of integers, doubles, and strings. The list of predefined arrays includes:

1. `ArrayInt` - for arrays of integer numbers, e.g.:

Method int[] getTerms()
return ArrayInt.getValues(terms);

Data ArrayInt terms
value
Term
36
72
108
144

2. `ArrayDouble` - for arrays of real numbers, e.g.:

Method double[] getCosts()
return ArrayDouble.getValues(costs);

Data ArrayDouble costs
value
Costs
$295.50
$550.00
$1,000.00
$2,000.00
$3,295.00
$5,595.00
$8,895.00

3. `ArrayString` - for arrays of strings, e.g.:

Method String[] getRegions()
return ArrayString.getValues(regions);

Data ArrayString regions
value
Region
NORTHEAST
MID-ATLANTIC
SOUTHERN
MIDWEST
MOUNTAIN
PACIFIC-COAST

These arrays are available from inside an OpenRules® table by just calling their names: `getTerms()`, `getCosts()`, `getRegions()`. You may also access these arrays from a Java program, using this code:

```
OpenRulesEngine engine =

        new OpenRulesEngine("file:rules/Data.xls");

int[] terms = (int[])engine.run("getTerms");
```

The standard installation includes a sample project "DataArrays", that shows how to deal with predefined arrays.

Accessing Excel Data from Java - Dynamic Objects

You can access objects created in Excel data tables from your Java program. These objects have a predefined type `DynamicObject`. Let's assume that you defined your own Datatype, Customer, and created an array of customers in Excel:

Data Customer customers			
name	maritalStatus	gender	age
Customer Name	**Marital Status**	**Gender**	**Age**
Robinson	Married	Female	24
Smith	Single	Male	19

Method Customer[] getCustomers()
return customers;

In you Java program you may access these objects as follows:

```
OpenRulesEngine engine =

        new OpenRulesEngine("file:rules/Data.xls");
DynamicObject[] customers =
 (DynamicObject[])engine.run("getCustomers");
System.out.println("\nCustomers:");
for(int i=0; i<customers.length; i++)
        System.out.println("\t"+customers[i]);
```

This code will print:

```
Customer(id=0){
        name=Robinson
        age=24
        gender=Female
        maritalStatus=Married
}

Customer(id=1){
        name=Smith
        age=19
        gender=Male
        maritalStatus=Single
}
```

You may use the following methods of the class DynamicObject:

```
public Object getFieldValue(String name);

public void setFieldValue(String name, Object value);
```

For example,

```
String gender = (String) customers[0].getFieldValue("gender");
```

will return "Female", and the code

```
customer.setFieldValue("gender", "Male");
customer.setFieldValue("age", 40);
```

will change the gender of the object customers[0] to "Male" and his age to 40.

How to Define Data for Aggregated Datatypes

When one Datatype includes attributes of another Datatype, such datatypes are usually known as *aggregated datatypes*. You have already seen an example of an aggregated type, `PersonalInfo,` with the subtype SSN. `Similarly,` you may have two datatypes, `Person` and `Address`, where type Person has an attribute "address" of the type `Address`. You may create a data table with type Person using aggregated field names such as "address.street", "address.city", "address.state", etc. The subtype chain may have any length, for example "address.zip.first5" or "address.zip.last4". This feature very

conveniently allows a compact definition of test data for complex interrelated structures.

Finding Data Elements Using Primary Keys

You may think about a data table as a database table. There are a few things that make them different from traditional relational tables, but they are friendlier and easier to use in an object-oriented environment. The very first attribute in a data table is considered to be its *primary key*. For example, the attribute "id" is a primary key in the data table "personalInformation" above. You may use values like "He" or "She" to refer to the proper elements of this table/array. For example, to print the full name of the person found in the array "personalInformation", you may write the following snippet:

```
 PersonalInfo pi = personalInformation["He"];

 out(pi.fisrtName + " " + pi.middeInitial + ". "

    + pi.lastName);
```

Cross-References Between Data Tables

The primary key of one data table could serve as a foreign key in another table thus providing a cross-reference mechanism between the data tables. There is a special format for data tables to support cross-references:

Data datatypeName tableName			
AttributeName1 from "datatypeName"	AttributeName2 from "datatypeName"	AttributeName3 from "datatypeName"	...
>referencedDataTable1		>referencedDataTable2	
Display value of the AttributeName1	Display value of the AttributeName2	Display value of the AttributeName3	...
data	data	data	...
data	data	data	...
...	...	...	...

This format adds one more row, in which you may add references to the other data tables, where the data entered into these columns should reside. The sign ">" is a special character that defines the reference, and "referencedDataTable" is the name of another known data table. Here is an example:

Data TaxReturn taxReturns					
taxPayer	spouse	wages	taxableInterest	taxWithheld	earnedIncomeCredit
>personalInformation	>personalInformation				
TaxPayer	Spouse	Wages	Taxable Interest	Tax Withheld	Earned Income Credit
He	She	32026	1450	4530	230

Both columns "TaxPayer" and "Spouse" use the reference ">personalInformation". It means that these columns may include only primary keys from the table, "personalInformation". In our example there are only two valid keys, He or She. If you enter something else, for example "John" instead of "He" and save your Excel file, you will receive a compile time (!) error "Index Key John not found" (it will be displayed in your Eclipse Problems windows). It is extremely important that the **cross-references are automatically validated at compile time** in order to prevent much more serious problems at run-time.

Multiple examples of complex inter-table relationships are provided in the sample rule project AutoInsurance. Here is an intuitive example of three related data tables:

Data Driver drivers				
name	age	gender	maritalStatus	dmvPoints
Name	Age	Gender	Marital Status	DMV Points
John Smith	24	Male	Single	2
Mary Smith	19	Female	Single	0

Data Vehicle vehicles				
id	make	model	year	hasAbs
ID	**Make**	**Model**	**Year**	**Has ABS**
Veh 1	Nissan	Maxima	2000	TRUE
Veh 2	Toyota	Corrola	1999	FALSE

Data Usage usages		
driver	vehicle	usage
> drivers	>vehicles	
Driver	**Vehicle**	**Usage(%)**
John Smith	Veh 1	100
Mary Smith	Veh 2	100

See more complex examples in the standard project "AutoInsurance".

OPENRULES® REPOSITORY

To represent business rules OpenRules® utilizes a popular spreadsheet mechanism and places rules in regular Excel files. OpenRules® allows users to build enterprise-level rules repositories as hierarchies of inter-related xls-files. The OpenRules® Engine may access these rules files directly whether they are located in the local file system, on a remote server, in a standard version control system or in a relational database.

Logical and Physical Repositories

The following picture shows the logical organization of an OpenRules® repository and its possible physical implementations:
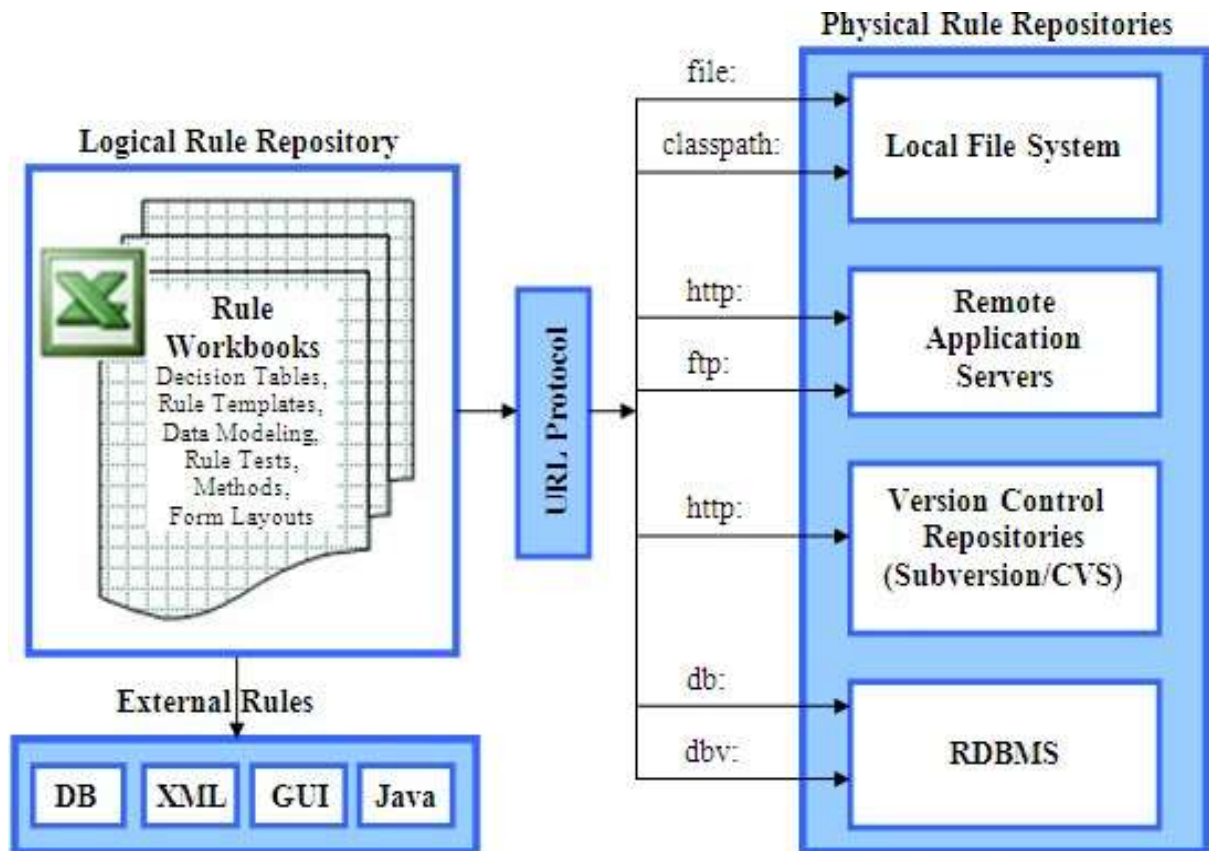
Logically, OpenRules® Repository may be considered as a hierarchy of rule workbooks. Each rule workbook is comprised of one or more worksheets that can be used to separate information by types or categories. Decision tables are the most typical OpenRules® tables and are used to represent business rules. Along with rule tables, OpenRules® supports tables of other types such as: Form Layouts, Data and Datatypes, Methods, and Environment tables. A detailed description of OpenRules® tables can be found here.

Physically, all workbooks are saved in well-established formats, namely as standard xls- or xml-files. The proper Excel files may reside in the local file system, on remote application servers, in a version control system such as Subversion, or inside a standard database management system.

OpenRules® uses an URL pseudo-protocol notation with prefixes such as `"file:"`, `"classpath:"`, `"http://"`, `"ftp://"`, `"db:"`, etc.

Hierarchies of Rule Workbooks

An OpenRules® repository usually consists of multiple Excel workbooks distributed between different subdirectories. Each rule workbook may include references to other workbooks thus comprising complex hierarchies of inter-related workbooks and rule tables.

Included Workbooks

Rules workbooks refer to other workbooks using so called "includes" inside the OpenRules® "Environment" tables. To let OpenRules® know about such include-relationships, you have to place references to all included xls-files into the table "Environment". Here is an example of an OpenRules® repository that comes with the standard sample project "RuleRepository":
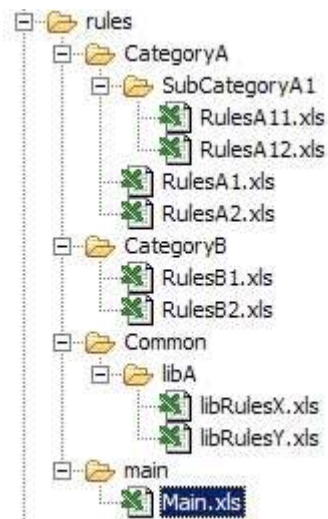
The main xls-file "Main.xls" is located in the local directory "rules/main". To invoke any rules associated with this file, the proper Java program creates an OpenRulesEngine using a string "file:rules/main/Main.xls" as a parameter.

There are many other xls-files related to the Main.xls and located in different subdirectories of "rules". Here is a fragment of the Main.xls "Environment" table:

include	../CategoryA/RulesA1.xls
	../CategoryA/RulesA2.xls
	../CategoryB/RulesB1.xls
	../CategoryB/RulesB2.xls
	../Common/libA/libRulesX.xls
	../Common/libA/libRulesY.xls

As you can guess, in this instance all included files are defined relative to the directory "rules/main" in which "Main.xls" resides. You may notice that files "RulesA11.xls" and "RulesA12.xls" are not included. The reason for this is that only "RulesA1.xls" really "cares" about these files. Naturally its own table "Environment" contains the proper "include":

Environment	
import.java	myjava.packA1.*
include	SubCategoryA1/RulesA11.xls
	SubCategoryA1/RulesA12.xls

125 ◎

Here, both "includes" are defined relative to the directory "CategoryA" of their "parent" file "RulesA1.xls". As an alternative, you may define your included files relative to a so called "include.path" - see sample in the next section.

Include Path and Common Libraries of Rule Workbooks

Includes provide a convenient mechanism to create libraries of frequently used xls-files and refer to them from different rule repositories. You can keep these libraries in a file system with a fixed "include.path". You may even decide to move such libraries with common xls-files from your local file system to a remote server. For instance, in our example above you could move a subdirectory "libA" with all xls-files to a new location with an http address http://localhost:8080/my.common.lib. In this case, you should first define a so-called "include.path" and then refer to the xls-files relative to this include.path using angle brackets as shown below:

include.path	http://localhost:8080/my.common.lib/
include	<libA/libRulesX.xls>
	<libA/libRulesX.xls>

Here we want to summarize the following important points:

- The structure of your rule repository can be presented naturally inside xls-files themselves using "includes"
- The rule repository can include files from different physical locations
- Complex branches on the rules tree can encapsulate knowledge about their own organization.

Using Regular Expressions in the Names of Included Files

Large rule repositories may contain many files (workbooks) and it is not convenient to list all of them by name. In this case you may use regular expression inside included file names within the Environment table. For example, consider in the following Environment table:

Environment	
include	../category1/*.xls
include	../category2/XYZ*.xls
include	../category3/A?.xls

The first line will include all files with an extension "xls" from the folder "category1". The second line will include all files with an extension "xls" and which names start with "XYZ" from the folder "category2". The third line will include all files with an extension "xls" that start with a letter "A" following exactly one character from the folder "category1".

Actually along with wildcard characters "*" or "?" you may use any standard <u>regular expressions</u> to define the entire path to different workbooks.

Imports from Java

OpenRules® allows you to externalize business logic into xls-files. However, these files still can use objects and methods defined in your Java environment. For example, in the standard example "RulesRepository" all rule tables deal with Java objects defined in the Java package myjava.package1. Therefore, the proper Environment table inside file Main.xls (see above) contains a property "import.java" with value "myjava.package1.*".

Usually, you only place common Java imports inside the main xls-file. If some included xls-files use special Java classes you can reference them directly from inside their own Environment tables.

Imports from XML

Along with Java, OpenRules® allows you to use objects defined in XML files. For example, the standard sample project "HelloXMLCustomer" uses an object of the type, Customer, defined in the file Customer.xml located in the project classpath:

```
<Customer
    name="Robinson"
    gender="Female"
    maritalStatus="Married"
```

```
        age="55"
    />
```

The xls-file "HelloCustomer.xls" that deals with this object includes the following Environment table:

Environment	
import.static	com.openrules.tools.Methods
import.schema	**classpath:/Customer.xml**
import.java	hello.Response
include	include/HelloRules.xls

The property `"import.schema"` specifies the location of the proper xml-file, in this case `"classpath:/Customer.xml"`. Of course, it could be any other location in the file system that starts with the prefix "**file:**". This example also tells you that this Excel file uses:

1. static Java methods defined in the standard OpenRules® package "com.openrules.tools.Methods"
2. xml-file "classpath:/Customer.xml"
3. Java class "Response" from a package "hello"
4. include-file "HelloRules.xls" that is located in the subdirectory "include" of the directory where the main xls file is located.

Parameterized Rule Repositories

An OpenRules® repository may be parameterized in such a way that different rule workbooks may be invoked from the same repository under different circumstances. For example, let's assume that we want to define rules that offer different travel packages for different years and seasons. We may specify a concrete year and a season by using environment variables YEAR and SEASON. Our rules repository may have the following structure:

```
rules/main/Main.xls

rules/common/CommonRules.xls
```

```
rules/2007/SummerRules.xls

rules/2007/WinterRules.xls

rules/2008/SummerRules.xls

rules/2008/WinterRules.xls
```

To make the OpenRulesEngine automatically select the correct rules from such a repository, we may use the following parameterized include-statements inside the Environment table of the main xls-file rules/main/Main.xls:

Environment	
import.java	season.offers.*
include	../common/SalutationRules.xls
include	../${YEAR}/${SEASON}Rules.xls

Thus, the same rules repository will handle both WinterRules and SummerRules for different years. A detailed example is provided in the standard project SeasonRules.

Rules Version Control

For rules version control you can choose any standard version control system that works within your traditional software development environment. We would recommend using an open source product "Subversion" that is a compelling replacement for CVS in the open source community. For business users, a friendly web interface is provided by a popular open source product TortoiseSVN. For technical users, it may be preferable to use a Subversion incorporated into Eclipse IDE. One obvious advantage of the suggested approach is the fact that both business rules and related Java/XML files will be handled by the same version control system.

You may even keep your Excel files with rules, data and other OpenRules® tables directly in Subversion. If your include-statements use http-addresses that point to a concrete Subversion repository then the OpenRulesEngine will dynamically

access SVN repositories without the necessity to move Excel files back into a file system.

Another way to use version control is to place your rule workbooks in a database and use DBV-protocol to access different versions of the rules in run-time - read <u>more</u>.

Rules Authoring and Maintenance Tools

OpenRules® relies on standard commonly used tools (mainly from Open Source) to organize and manage a Business Rules Repository:
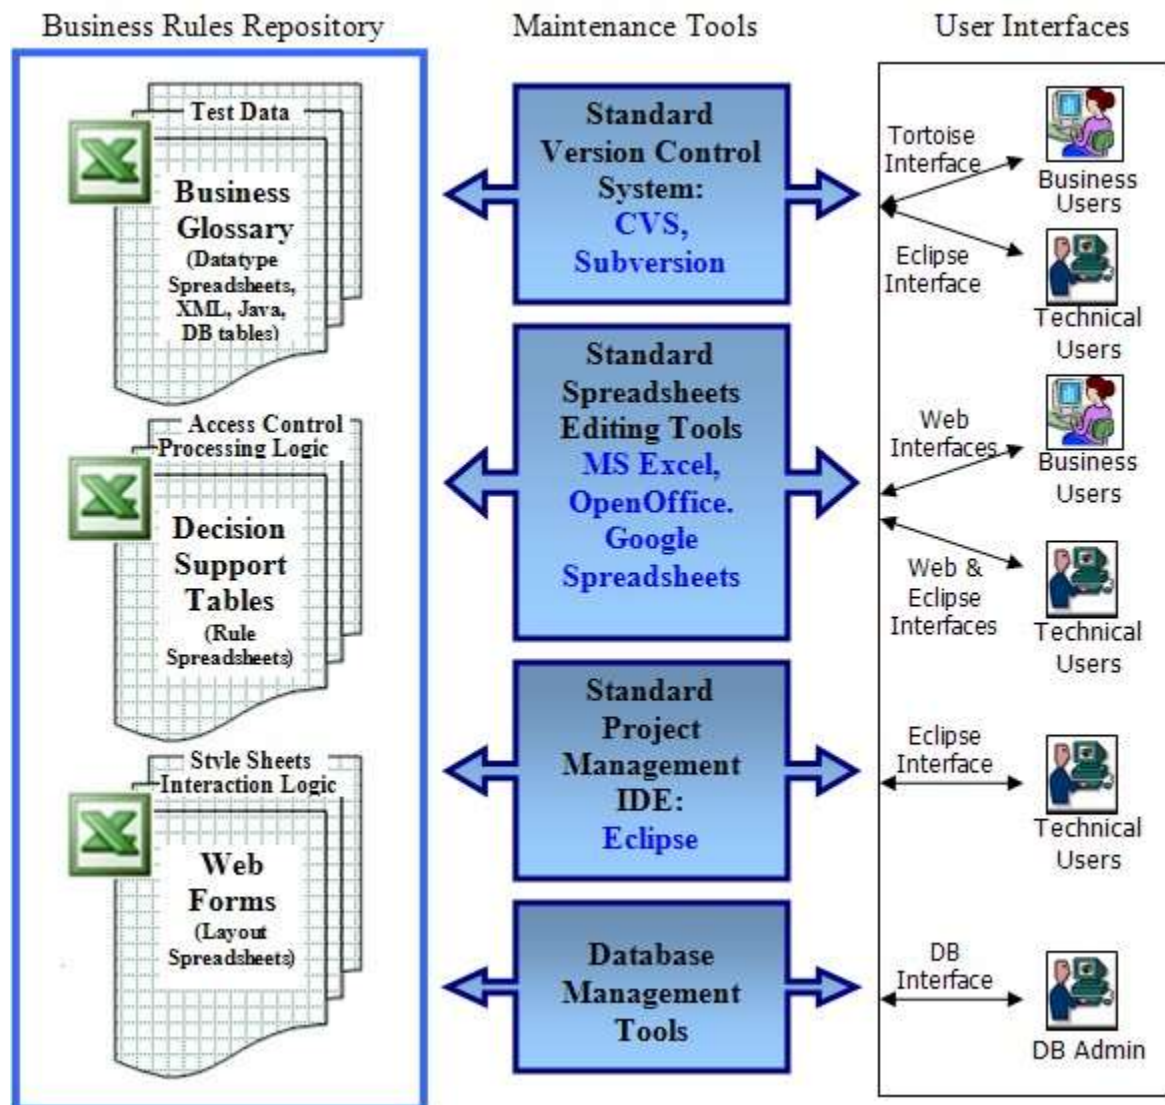
To create and edit rules and other tables presented in Excel-files you may use any standard <u>spreadsheet editors</u> such as:

- <u>MS Excel™</u>
- <u>OpenOffice™</u>
- <u>Google Spreadsheets™</u>

Google Spreadsheets are especially useful for sharing spreadsheet editing - see section <u>Collaborative Rules Management with Google Spreadsheets</u>.

For technical people responsible for rules project management OpenRules provides an <u>Eclipse Plug-in</u> that allows them to treat business rules as a natural part of complex Java projects.

DATABASE INTEGRATION

OpenRules® provides a user with ability to access data and rules defined in relational databases. There are two aspects of OpenRules® and database integration:

1. Accessing data located in a database
2. Saving and maintaining rules in a database as Blob objects.

The detailed description of database integration in provided at <u>http://openrules.com/pdf/OpenRulesUserManual.DB.pdf</u>.

EXTERNAL RULES

OpenRules® allows a user to create and maintain their rules outside of Excel-based rule tables. It provides a generic Java API for adding business rules from different external sources such as:

1. Database tables created and modified by the standard DB management tools
2. Live rule tables in memory dynamically modified by an external GUI

3. Java objects of the predefined type "RuleTable"

4. Problem-specific rule sources that implement a newly offered rules provider interface.

With external rules you may keep the business parts of your rules in any external source while the technical part (Java snippets) will remain in an Excel-based template, based on which actual rules will be created by the OpenRulesEngine. For example, you may keep your rules in a regular database table as long as its structure corresponds to the columns (conditions and actions) of the proper Excel template. Thus, the standard DB management tools, or your own GUI that maintains these DB-based rule tables, de-facto become your own rules management environment.

The external rules may also support a preferred distribution of responsibilities between technical and business people. The business rules can be kept and maintained in a database or other external source by business analysts while developers can continue to use Excel and Eclipse to maintain rule templates and related software interfaces.

The detailed description of external rules in provided at http://openrules.com/pdf/OpenRulesUserManual.ExternalRules.pdf.

OPENRULES® PROJECTS

Pre-Requisites

OpenRules® requires the following software:

- Java SE JDK 1.5 or higher
- Apache Ant 1.6 or higher
- MS Excel or OpenOffice or Google Docs (for rules and forms editing only)
- Eclipse SDK (optional, for complex project management only)

Sample Projects

The complete OpenRules® installation includes the following workspaces:

openrules.decisions - decision projects

openrules.rules - various rules projects

openrules.dialog – rules-based web questionnaires

openrules.web - rules-based web applications & web services

openrules.solver - constraint-based applications.

Each project has its own subdirectory, e.g. "DecisionHello". OpenRules® libraries and related templates are located in the main configuration project, "openrules.config", included in each workspace. A detailed description of the sample projects is provided in the Installation Guide.

Main Configuration Project

OpenRules® provides a set of libraries (jar-files) and Excel-based templates in the folder "`openrules.config`" to support different projects.

Supporting Libraries

All OpenRules® jar-files are included in the folder, "`openrules.config/lib`". For the decision management projects you need at least the following jars:

- `openrules.all.jar`
- `poi-3.6-20091214.jar`
- `commons-logging-1.1.jar (or higher)`
- `commons-logging-api-1.1.jar (or higher)`
- `commons-lang-2.3.jar (or higher)`
- `log4j-1.2.15.jar (or higher)`
- `commons-beanutils.jar (or higher)`

There is a supporting library

- `com.openrules.tools.jar`

contains the following optional facilities:

- operators described in the Java class `Operator` that can be used inside your own Rules tables and templates

- convenience methods like "`out(String text)`" described in the Java class `Methods`

- a simple JDBC interface `DbUtil`

- text validation methods like "`isCreditCardValid(String  text)`" described in the Java class `Validator`.

If you use the JSR-94 interface you will also need

- `com.openrules.jsr94.jar`

If you use external rules from a database you will also need

- `openrules.db.jar`
- `openrules.dbv.jar`
- `derby.jar`
- `commons-cli-1.1.jar`.

Different workspaces like "`openrules.decisions`", "`openrules.rules`", etc. include the proper versions of the folder "`openrules.config`".

Predefined Types and Templates

The Excel-based templates that support Decisions and Decision Tables included in the folder, "openrules.config":

- `DecisionTemplates.xls`
- `DecisionTableExecuteTemplates.xls`
- `DecisionTableValidateTemplates.xls`

Sample decision projects include Excel tables of the type "Environment" that usually refer to "`../../../openrules.config/DecisionTemplates.xls`". You may move all templates to another location and simply modify this reference making it relative to your main xls-file.

TECHNICAL SUPPORT

Direct all your technical questions to support@openrules.com or to this Discussion Group. Read more at http://openrules.com/services.htm.