



RippleDown Knowledge Builder Guide

This Knowledge Builder Guide provides detailed information including:

- RippleDown Rules
- Attributes
- Features and Derived Features
- Rule Wizard
- Reports
- Searching
- Performance Reports
- Help

More information regarding RippleDown, including training modules and product help, can be found at:

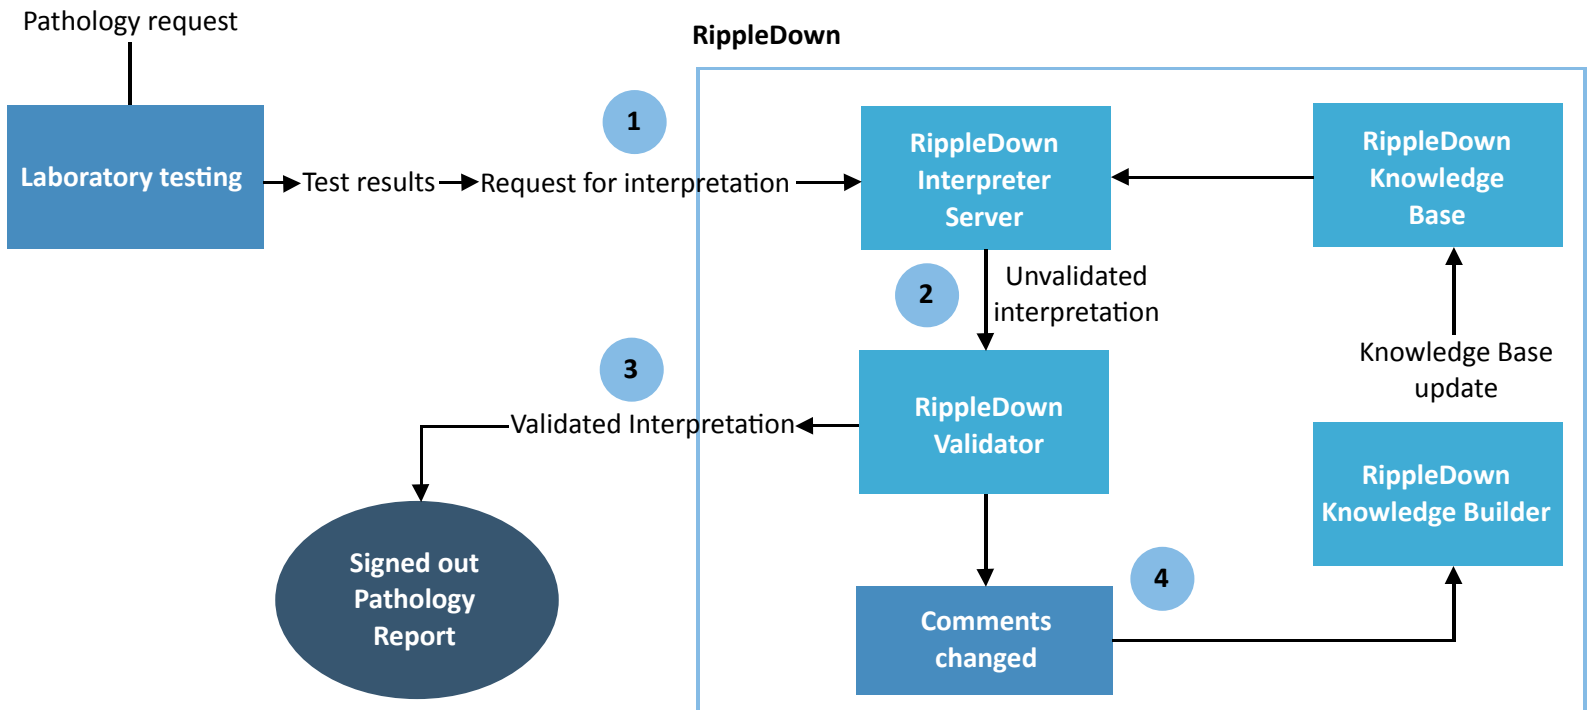
www.pks.healthcare/support

Overview

The Knowledge Builder allows users to create a Knowledge Base that captures their own clinical expertise and decision making process so it can be applied automatically. The Knowledge Base can be easily updated by users with new rules to ensure that interpretations and decisions are accurate and reflect the most recent practices.

Workflow

Laboratory Information System

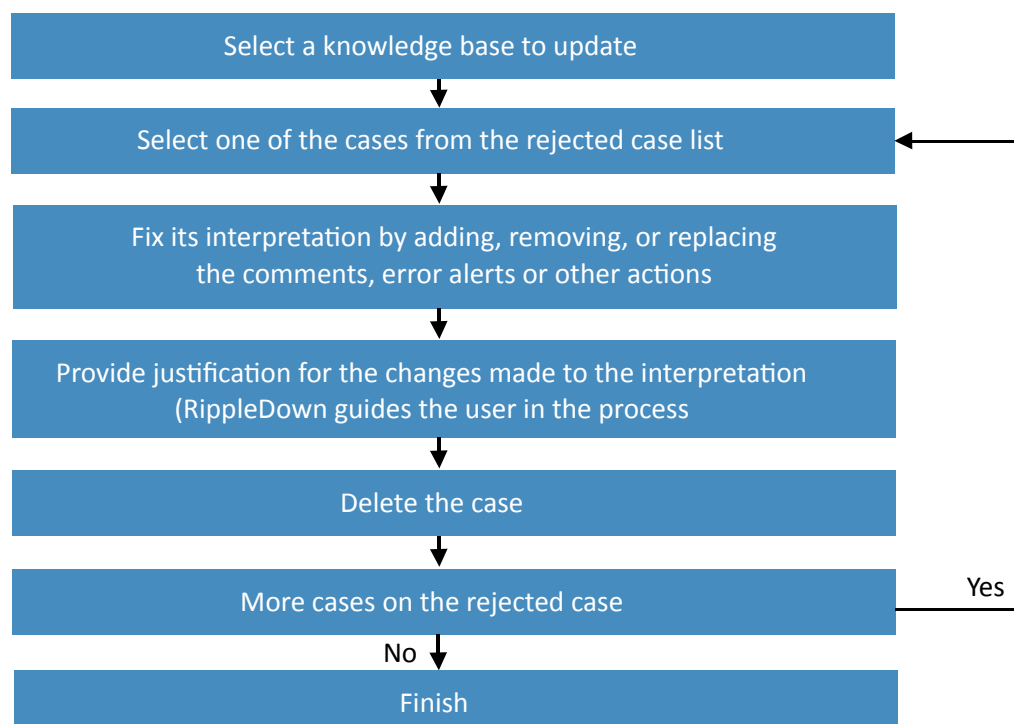


Knowledge Base Organisation

RippleDown automatically structures rules within the Knowledge Base to maintain a level of consistency and transparency over all rules that are built. This helps in knowledge retention and auditing within an organisation, as well as refining the Knowledge Base.

Knowledge Builder Tasks

Once the Knowledge Base is online the basic task for the Knowledge Builder is shown below:



Projects



Planning

Before building a project it is important to first plan the main objectives, the attributes required from the Online Information System, and the outputs expected from the Knowledge Base. You will also need to create an appropriate number of report sections.

Starting a new project

To start a new project, you will need input from:

1. the **RippleDown Administrator** to configure the Online Information System to:
 - send cases to RippleDown consisting of current and historical patient data
 - receive reports, alerts etc. from RippleDown
 - create project/s in RippleDown
 - ensure all relevant users have access to the project/s
2. the **Project Expert** to determine which data elements are needed in the project. These should include:
 - patient demographics
 - episode details
 - test codes and results
 - history files
3. the **Rule Builder** to create the Knowledge Base project in RippleDown.

Using the Knowledge Builder

To login in to Knowledge Builder, double click on the RippleDown icon and select 'Knowledge Builder' from the list.

Enter your username and password.

From the File menu, select 'Open Project' to open an existing project, or 'Create new project' to start a new one.



Contact your
RippleDown
Administrator if you
have any trouble logging
on, or require a
username and password

Case View

The Case View captures all information in an open case including:

- demographic data
- current and previous episodes
- current interpretations of the case
- rules associated with each of the comments in the interpretation
- any notes added to the case by the person who reviewed it in Validator
- any derived attributes or features that have been defined for the case

The screenshot displays the 'Case View' window in the Knowledge Builder application. It shows a patient's clinical data, including demographics (Age 71, Sex F), clinical notes (e.g., 'NormoSevAnaemi... WBC... false'), and a list of rules (e.g., 'NormoSevAnaemi... WBC... false'). The interface includes a menu bar (File, Edit, Case, Lab, Interpretation, Tools, Autovalidation, Options, Windows, Help) and a toolbar with buttons for 'Comments', 'Preview', 'Approval', 'Autovalidation', 'Queues', and 'Notes'. The status bar at the bottom indicates '1080 rules' and 'Autovalidation level: 100% By report'.

The first column displays the attribute / derived attribute / derived name.

The second column displays the reference ranges (if any exist) for the analyte attribute.

Columns from there onwards are for the episodes. Each episode has its own column, with results in the appropriate rows. The current episode column is the one farthest from the reference column.



These ranges are not set-up in RippleDown, but are sent from the Online Information System with the case, and therefore only relate to the current episode.


Case View Toolbar

The Toolbar at the top of the Case View screen includes a range of menus you can use to navigate and action items, such as:

File	<ul style="list-style-type: none">Used to open and close projects
Edit	<ul style="list-style-type: none">Provides access to Editors for constants, attributes, objectives, comments, queues and report sections
Case	<ul style="list-style-type: none">Used to add/delete/copy and search for cases
Interpretation	<ul style="list-style-type: none">Used for opening the Rule Wizard to add comments, features, derived features, autovalidation and queuesCan also be used to undo the last rule if required
Tools	<ul style="list-style-type: none">Provides access to the Project Explorer, Scratch pad and Project reports
Autovalidation	<ul style="list-style-type: none">Used for managing autovalidation and viewing comment statistics
Options	<ul style="list-style-type: none">Used to customise font type, size and format, set project settings such as case list properties, and show/hide removed comments
Windows	<ul style="list-style-type: none">Describes keyboard shortcuts for navigating through the Case View screen. Can also be used to hide panes on the main screen
Help	<ul style="list-style-type: none">Used to access the RippleDown online help facility

Action Panel

The panel below the results has 6 tabs referring to screen panes:



Comments	<ul style="list-style-type: none">The pane used to add comments
Preview	<ul style="list-style-type: none">Displays the report as presented in the Validator and the Online Information System
Approval	<ul style="list-style-type: none">Displays the report as approved by the RippleDown Validator and any differences between the original and approved reports.
Autovalidation	<ul style="list-style-type: none">Pane used to define auto-validation settings
Queues	<ul style="list-style-type: none">Pane used to define a special purpose queue
Notes	<ul style="list-style-type: none">Displays any notes added to the case by either the Validator or a Knowledge Builder

Case Lists

Several case lists can be found in the knowledge base, including:

Case list	Functionality
Rejected	<ul style="list-style-type: none">These are the cases in which a change has been made to the RippleDown interpretations during validation in either RippleDown Validator or directly in the Online Information System.Rejected cases are the ones that you are likely to use when building rules for these interpretations.Only the most recent 1,000 rejected cases will be kept on this list. You can delete a case from the 'Rejected' list once you have actioned it, for example by building rules to fix its report.You cannot edit a case on this list. If you wish to edit a case on this list, copy it to the 'Favourites' folder and then edit it in the Favourites list.
New reports	<ul style="list-style-type: none">Cases which have reports that you have not yet allowed to be auto-validated. You can allow a specific report for a case on this list using the Allow button on the Preview window. Only the most recent 1,000 cases will be kept on the New Reports list.
Favourites	<ul style="list-style-type: none">This list is manually created by the Knowledge Builder by copying cases from one of the other lists to the Favourites list.Once a case is on the Favourites list it may be kept for demonstration or training purposes and it can be edited for rule building and testing purposes. There is no limit to the number of cases on the Favourites list.
Cornerstones	<ul style="list-style-type: none">These are the cases that have been used to build rules, and so will normally provide the best examples of each type of interpretation given by the Knowledge Base.You cannot edit or delete a case from this list. You can however copy a case to the Favourites. There is no limit to the number of cases on the Cornerstone list.
Archive	<ul style="list-style-type: none">Stores 10,000 cases following interpretation by the RippleDown server.If a case is deleted from the Archive, it is permanently deleted from the project but may be retrieved from a backed up zip file stored by the server.
Search	<ul style="list-style-type: none">When a search by condition is performed on one of the other lists, the cases found are copied into the search list.

RippleDown Rules

A rule consists of one or more logical expressions or conditions, which, if true for the case means that an output such as a report comment, an alert, or a workflow instruction is generated from the project.

Rule Wizard

The Rule Wizard is the application used to build rules within RippleDown. The application will launch each time a rule is started, and will provide a step-by-step guide throughout the rule building process.

Any case which a rule is built around is identified as a 'Cornerstone' and is stored within the Knowledge Base. All subsequent rules are checked against these cornerstones to maintain the integrity of the Knowledge Base.

A user can undo the last rule that was added if required, and this action can be repeated for all rules in the Knowledge Base. Any changes made to the Knowledge Base, will take effect immediately and for subsequent interpretations made by the RippleDown Server.

Users can create conditions in rules, using near-natural language syntax. There are many Conditions, however some may refer to:

- the current or previous episodes
- trends in a test value
- previous or current reference ranges
- a specific time period or date
- comparisons between values
- textual values in clinical notes or other fields

ALL conditions in an expression must be met before a rule is 'true' or has passed. The full condition syntax list can be found in the RippleDown Help utility.

Attributes

Attributes are specific data items for patients, episodes or tests, for example Given Name, DOB, Sodium, Chloride.

Derived attributes are data items which have been defined and created by the user, based on the raw data (primary attribute results).

Primary Attributes

Primary attributes represent the raw data items that are sent to RippleDown from the Online Information System. They appear in the Attribute Editor section of RippleDown with both an external and internal name.

The external name refers to the unique codes supplied by the Online Information System. The internal name initially appears as an identical match to the external name, however it is able to be customised in the Attribute Editor in RippleDown to meet specific user requirements. The internal name is displayed in RippleDown Knowledge Builder and Validator case views.

Attribute Editor

Managing and customising attributes is done through the Attribute Editor, which can be accessed by selecting '**Attributes**' from the '**Edit**' menu.

You can search for attributes within the Attribute Editor using the search bar at the top of the window.

To move attributes to a new position in the list, simply click on the specific attribute you want to move and drag it to the preferred location. Multiple attributes can be moved by using CTRL-Click or SHIFT-Click to select multiple attributes.

Other case view properties of existing attributes can also be customised by right clicking on the desired attribute and select Attribute Display. The following functions will then become available:

Function	Description	
Name	● Enter the name to be displayed in Case View, with no spaces. Words can be separated by an underscore.	
Description	● Enter a description, if required, for the attribute.	
Position	● Enter the position number in which the attribute is to be displayed in Case View. The "in header" option will override the position number.	
Attribute display	● Select one of the options:	
	in header	To position the attribute and current value in the Case Header.
	highlighted	To add a highlight colour to the row containing the attribute. The colour can be customised in the Attribute Editor under ' Options ' and then ' Highlight colour '.
	always visible	The attribute is always shown.
	normal	The attribute and corresponding value at each episode are shown unless the sample value is empty. This the default.
	hidden	The attribute and sample values are not shown but may be used in creating rules.
Format	Default	The attribute will not show in the body of a case if all of the samples for the attribute for that case are 'uninteresting'. For a primary attribute, a calculated value, delta, a double list or a merged attribute, 'uninteresting' means blank. For ranges, lists, text attributes and episode attributes, 'uninteresting' means having the value 'false'.
	No formatting	The attribute value will appear as presented by the Online Information System.
	1,2,3 or 4 significant figures	Attribute sample can be formatted to either 1,2,3,4 significant figures. This affects the display. Calculations using the attribute value will be performed on the data as it received from the Online Information System.

Derived Attributes

Derived attributes are user-defined attributes developed for a particular project. Users can create derived attributes using the Knowledge Builder.

These derived attributes and their values can be used for rule building in exactly the same way as primary attributes. RippleDown will evaluate each derived attribute for each episode in the case.

The derived attributes available are:



Numeric Range

A numeric range gives a 'true' or 'false' value relative to a sub-range of values for a selected numerical attribute - such as a high, low or normal range. An example may be: The attribute 'Baby' is defined as a patient with an age in days between 7 and 365 days.



Derived Attributes cannot be deleted. To manage unwanted attributes simply right click on the attribute in the Attribute Editor and rename it - for example to "DoNotUse" - and then move it to the bottom of the list.



List Range

A list range is used to detect episodes in which the value for a primary attribute matches one of a list of values. A list attribute will return 'true' if the primary attribute on which it is based contains one of the entries from a discrete list.

An example might be a list of known Endocrinologists. The Endocrinologist list attribute will be 'true' if the referring doctor for the episode appears in the list (see image).

From the Case View it is possible to add elements to the list. To do this simply right click on the sample, select 'Add value to...' and select the appropriate list.



Delta Check

A Delta check refers to a detailed set of numeric ranges for a particular test that is defined in a tabular manner. These ranges may be separately defined according patient age, sex or other demographic information, and for sub-ranges of the test values.

An alert is generated if the current value is outside some defined range, or alternatively, if the value has increased or decreased from its previous value by an amount greater than what is allowed.

A Delta will give a value of "true" if the alert has been triggered for the current case. If the Delta is not triggered for the current case, it simply will not show.

Constraint	Constraint	Reference Range	Panic/Action Range	Variation Range	Allowable Increase	Allowable Decrease	Days
30 < DaysAge	30 < DaysAge	8.4-15.4	7.0-22.4	< 8.0	50.0 %	10.0 %	30
30 < DaysAge	30 < DaysAge	8.4-15.4	7.0-22.4	8-15	50.0 %	40.0 %	30
30 < DaysAge	30 < DaysAge	8.4-15.4	7.0-22.4	> 15	50.0 %	40.0 %	30
DaysAge < 31	DaysAge < 31	8.4-15.4	7.0-16.4	> 0	30.0 %	30.0 %	30



Delta Check

Each row of the Delta contains the following:

Constraints

- These are criteria that must be satisfied for the Delta processing to continue. Examples are, age > 50, sex is "F".
- Up to two constraints can be defined per row.
- To edit or add a constraint, right-click in a selected constraint cell and select "Edit Constraint" from the menu.
- The constraints can refer to primary or derived attributes of the case, and often refer to demographics such as age or sex, rather than the test attribute itself.
- All other ranges in the Delta refer to the test attribute.
- If you leave a constraint empty then it is assumed to be satisfied.
- If both constraints are satisfied by the current case, then the Delta processing will continue and look at the next Reference Range.

Reference range

- The Delta reference range is something you define for the test attribute - it may or may not correspond to the reference range supplied by the Online Information System as part of the case.
- If there is only 1 episode in the case, then the Delta will trigger if the attribute is outside the reference range.
- You should not leave the reference range blank, otherwise all single episode cases will trigger the Delta.
- To edit or add a reference range, right-click on the selected reference range cell and select "Edit reference range" from the menu.
- If there are two or more episodes in the case, and they are sufficiently close in time, then the reference range is ignored and the action range is used.

Action range

- Used if there are two or more episodes in the case, and the time between the current test value and the previous test value is less than the specified number of days. The Delta will trigger if the attribute is outside the action range.
- This enables you to define a wider action range than reference range, for those situations where you are closely monitoring the trend of a test value.
- To edit or add a action range, select the action cell and use the right-click menu "Edit action range"
- If the test value is within the action range (and we have the situation of the two episodes close in time), then the variation range and allowable Delta processing is done.

Delta Check (variation range, allowable increase, allowable decrease, units, days)

- A variation range defines the range of test values for which the permitted increase or decrease is checked. This allows you to define different permitted increases or decreases depending on the absolute value of the test attribute.
- To edit any of the Delta check fields, right-click on a selected Delta check cell and select "Edit delta check fields" from the menu.
- By default, each allowable increase or decrease will be interpreted as being in absolute units. However, you can change this to percentage.
- To add several variation ranges, each with a different allowable increase or decrease, use the right-click menu and select "Copy this row", then change just the Delta check fields for the new row.

Multiple rows can be defined if you require different reference ranges, action ranges, or Delta checks for different constraints, e.g. a new set of ranges if the patient is female, or over a certain age. To do this, use the right-click menu "Add a row", then edit the constraints and ranges for the new row.

Rows in the Delta can be re-ordered using the right-click menu and selecting "Move up" or "Move down". Rows can also be removed if required.

When you have finished defining the Delta, click the [Close] button. The new Delta will appear as an attribute in all cases where it is triggered.

Once the Delta ranges have been defined, the logic to trigger the check is as follows:

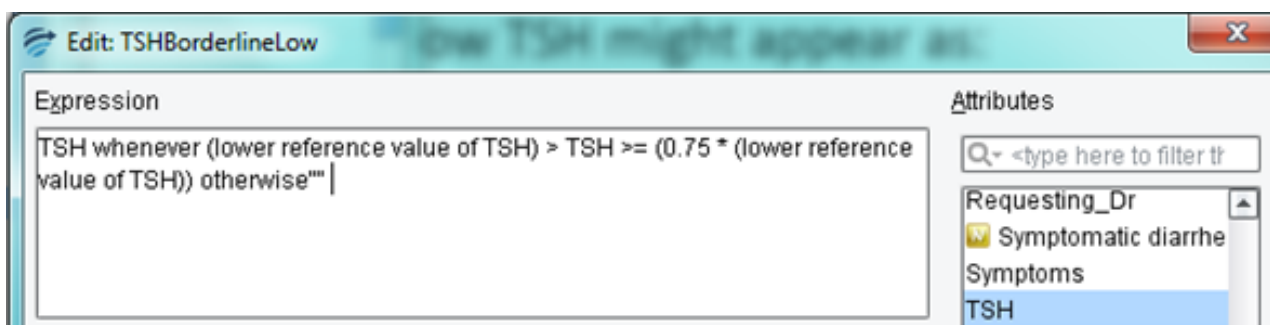
1. If there is only 1 episode for the case (i.e. the first presentation of this patient), the check is triggered if the test value is outside the reference range.
2. If there is more than 1 episode, and the previous episode is within the specified time period, the check is triggered if the test value is outside the panic range.
3. If the test value is inside the panic range, then the check is triggered if the change in the test value from the previous to the current episode is outside the allowable increase or decrease (specified either in absolute units or percentages).

Rules may be developed to take the appropriate action according to the value of a Delta check. For example, rules may be added to ensure that if a check is triggered, the case is never autovalidated. When the case is shown to the user in Validation, the reason why the check was triggered will be shown.



Calculated value

A calculated value refers to a formula involving one or more attributes. For example a determination of borderline low TSH might appear as the image below.

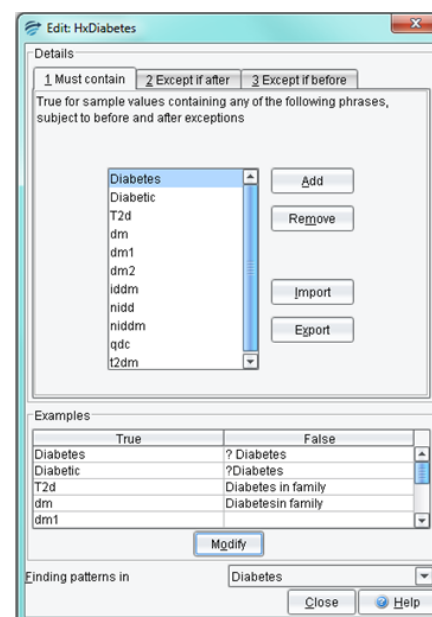


The formula can involve other calculated values as long as circular references are avoided i.e. A calculated attribute A can refer to calculated value attribute B as long as B does not already refer to A.



Text Attribute

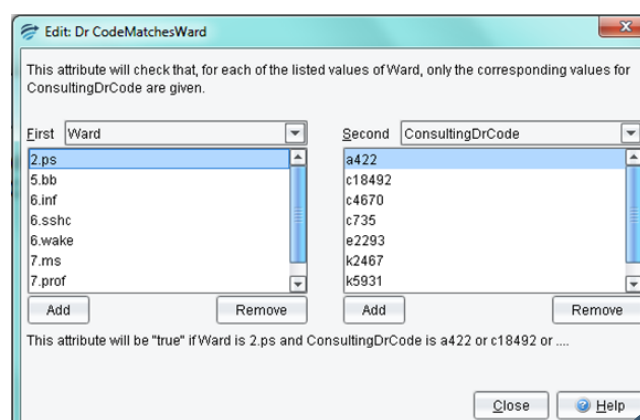
A text attribute is used to refer to the presence of a word or phrase in an attribute that has textual values, e.g. looking for “diabetic” in a patient’s clinical notes. Facilities are provided to refine the search so as to exclude similar but unwanted words or phrases, e.g. diabetic.



Double Lists

Double Lists are used to compare the values of two primary attributes, only when the value in the first list has been satisfied. So if the value is present in the first list, it will then compare that value in the case against the second list. If the values are satisfied in both lists, then the attribute will evaluate as ‘true’.

To demonstrate this, in the example pictured, if the Ward is ‘2ps’, it is then checked against the ConsultingDrCode in the second list which is ‘a422’, so given both values are present in the first and second list, the attribute will evaluate as ‘true’.





Merged Lists

A merged attribute is one that takes its value from two or more primary attributes. For example if a test is superseded in the Online Information System, the old test name may appear in an historical episode and the new name in the current episode. A merged attribute combines the two tests into a single attribute for use in the project. When defining a merged attribute, each test may be provided with a scaling factor to take into account changes in units between superseded and new test.

Define merged attribute - FT4_Merge

FT4_Merge will be created by merging values from the following primary attributes, taking into account the specified scale factors. See the Help for more details.

Primary Attribute	Scale factor
FT4	10
FT4new	1.0

Add Remove



Group Attribute

A group attribute refers to a set of related primary attributes. They can be useful in the situation where the project deals with many primary attributes, but where the comments and rules can be more simply expressed in terms of groups rather than the individual attributes.

Edit: LFTs

Q- <type here to filter the list>

Primary Attribute	Scale factor
FT4	10
FT4new	1.0

Add >> << Remove

For example, we can define the following groups of allergen attributes:

Group	Attributes in Group	
pollen	<input type="radio"/> birch	<input type="radio"/> mugwort
food	<input type="radio"/> milk	<input type="radio"/> peanut
animal	<input type="radio"/> dog	<input type="radio"/> cat



Episode Attribute

A merged attribute is one that takes its value from two or more primary attributes. For example if a test is superseded in the Online Information System, the old test name may appear in historical episode and the new name in the current episode. A merged attribute combines the two tests into a single attribute for use in the project. When defining a merged attribute, each test may be provided with a scaling factor to take into account changes in units between superseded and new test.

- the 'Pregnant' attribute has value "true"
- the 'CollectionType' attribute has value "Fasting"
- the 'Glucose' value is elevated.

The episode attribute, "GestDiab" will be true for when all three conditions are met, for example:

	05 Feb 12	06 Oct 12	07 Nov 12
Pregnant	false	true	false
Glucose	4.1	7.8*	4.5
CollectionType	R	F	F
GestDiab	False	True	False

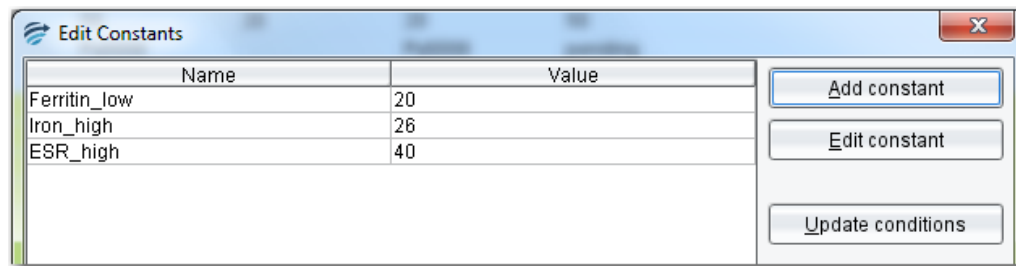
Constants

Constants are numerical or textual values used in conditions. For example, in a glucose Knowledge Base you might define a constant `GLUCOSE_HIGH = 5.5`. Then, instead of "hard-coding" numbers into your rules like `"glucose > 5.5"` you can use the constants, e.g. `"glucose > GLUCOSE_HIGH"`.

Advantages of using constants rather than hard-coded numbers:

- Is more easily understandable for users. For example a constant such as `'GLUCOSE_HIGH'` is more understandable than the number `"5.5"`, especially for someone who is exploring the Knowledge Base
- Saves time if changes need to be made. Once you have defined a constant, you can change its value at a later stage if required. This means that any conditions that use that constant will not have to be changed, saving you a lot of time.

Constants are defined and modified using the Constants Editor which can be found by selecting 'Constants' from the 'Edit' menu. Each constant consists of a name and a value. Multiple constants can be defined and will appear listed in a table:



The table can be sorted either by name or value simply by clicking on the appropriate heading.

To use constants:

Constants can be used by typing the constant name into expressions, for example: `Ferritin < Ferritin_Low`


Note: The constant name must be typed and it cannot be selected from the tabs in the Rule Wizard.

Updating an existing Knowledge Base to use constants

Once you have defined a set of constants, you can update any existing conditions in your Knowledge Base to use these constants. To do this, simply click on the '**Update conditions**' button. The system will then find all the conditions that use one of the constant values you have defined. You will then be asked to confirm each condition before the replacement is made.

Features and Derived Features

Features

Features are represented by the grey light globe icon: 


Each feature represents a high-level concept that is defined by rules which can also be used in other rules.

An example of a feature could be `'Is_Diabetic'` which can be added to a case for multiple purposes. For example, the `'Is_Diabetic'` may be applied where there are more than three instances of Glucose being higher than 11.0, or where clinical notes contains the term "diabetic". Once this feature has been defined, the interpretive comment "History of diabetes noted" can be added with condition `'Is_Diabetic' is true`.

There are several advantages to using features:

- Conditions using features may be more intelligible than using more complex expressions involving attributes
- A feature is a shorthand way of expressing a pair (or more) of conditions that often are used together
- A feature is easy to find and re-use in different rules
- The same feature may be defined in more than one way, for example `"IsDiabetic"` could also be given if the GTT 2 Hr result was at a certain level

Derived Features

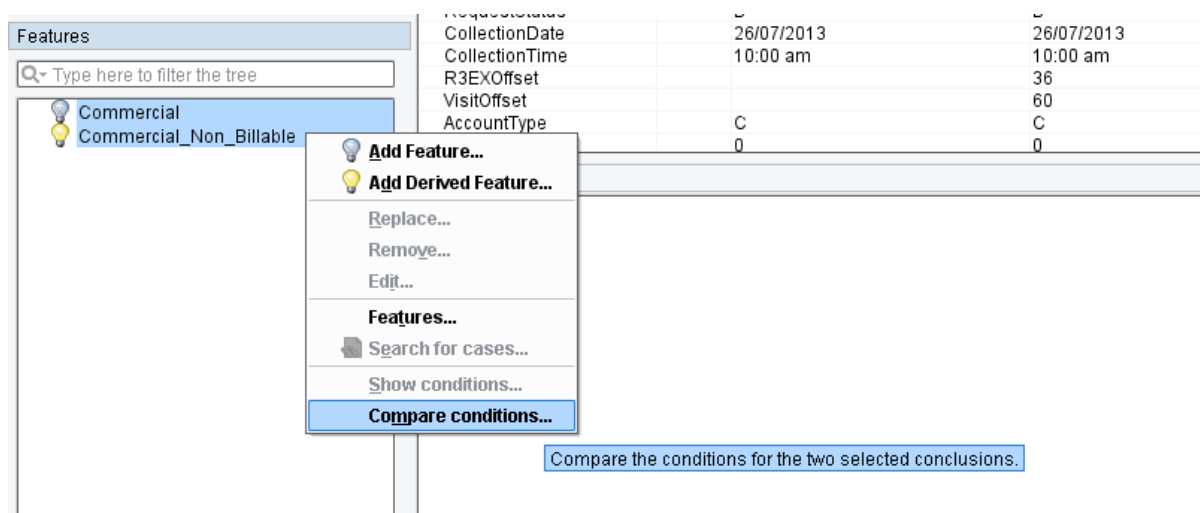
Derived features are represented by a yellow light globe icon in the application: 

Derived features are simply higher level features, with the only difference being that derived features can be defined using features.

For example, you might have already defined three features - for example XNormal, YNormal and ZNormal. There may be rules where you want to use these three features together. Rather than repeating these three features in conditions each time, you can then define a derived feature XYZNormal which is given if the other three are given, thus simplifying these rules.

Compare conditions for Features and Derived Features

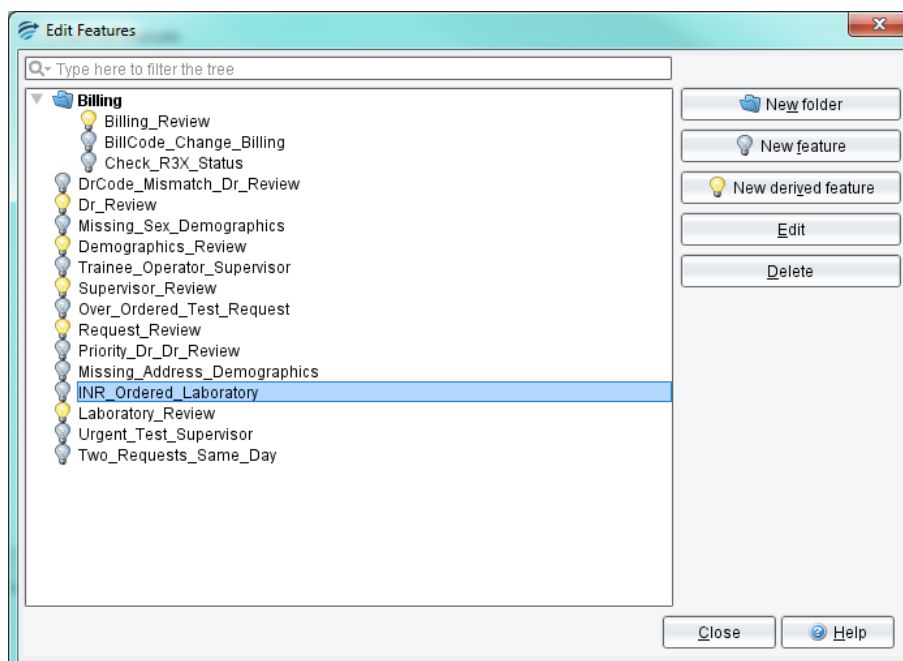
You can compare the conditions set for features and derived features by simply highlighting the features or derived features to be compared, right clicking your mouse, and selecting **"Compare conditions..."** from the menu.



Feature and Derived Feature Organisation

Features and derived features can be edited and organised into folders. This can be done through the 'Edit Features' window.

Features and derived features can be reorganised into folders by simply dragging and dropping them in the desired location.



Rules and Rule Wizard

A rule is a set of conditions used to generate a specific action - for example adding or deleting a comment. With RippleDown, rules can be built incrementally, and new rules do not break existing ones. RippleDown also has a safe rule building process with conflict checking.

RippleDown uses case-based reasoning in the rule building process, so a rule can be only be built against the conditions present in cases within the Case View.

To assist in the rule building process, RippleDown leverages the Rule Wizard application to guide users through the process of building rules to update the knowledge base for the case that is shown.

The Rule Wizard follows four main stages in the rule building process:

1. Define the rule action
2. Set an objective
3. Add conditions
4. Review cornerstone cases.

A detailed overview of each stage are outlined below.

Rule Wizard - Define the Rule Action

A variety of actions can be controlled by rules. These include:

Function	Action
Comments	<ul style="list-style-type: none">• Adding a comment• Replacing the selected comment by another one.• Removing the selected comment• Re-adding a removed comment• Enabling or disabling a report section
Features and Derived Features	<ul style="list-style-type: none">• Adding a feature• Replacing the selected feature by another one• Removing the selected feature
Autovalidation settings	<ul style="list-style-type: none">• Adding an auto-validation setting• Replacing the selected auto-validation setting by another one• Removing the selected auto-validation setting
Review queues	<ul style="list-style-type: none">• Adding a review queue• Replacing the selected review queue by another one• Removing the selected review queue
Report items	<ul style="list-style-type: none">• Enabling the selected report item• Disabling the selected report item

To initiate an action use the right-click function in the required tab.

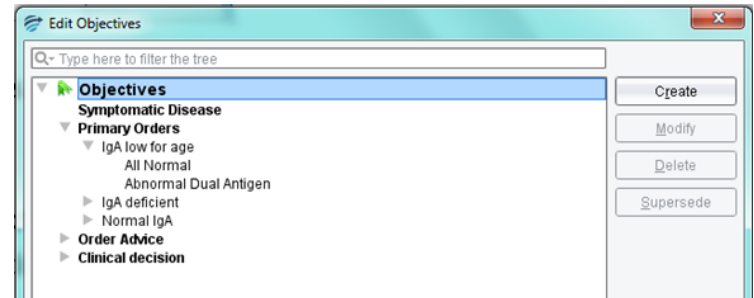
Rule Wizard - Set an Objective

Objectives are used to document a Knowledge Base by adding to each rule an explanation of why the rule is being built. An Objective can be further refined into a set of sub-Objectives. For example, an Objective could consist of a local commenting policy within the laboratory, and a sub-Objective could be a specific aspect of that policy. Therefore, it is possible to have several levels of Objectives.

The Objective Editor is used to view and manage objectives, and can be accessed in the 'Edit' menu under 'Objectives'. The Objective Editor enables users to:

- create new Objectives
- modify existing Objectives
- delete unused Objective
- supersede one Objective with another
- re-order or move an Objective under a new parent Objective using drag-and-drop

An extract of the Objective Editor is shown in the following image:

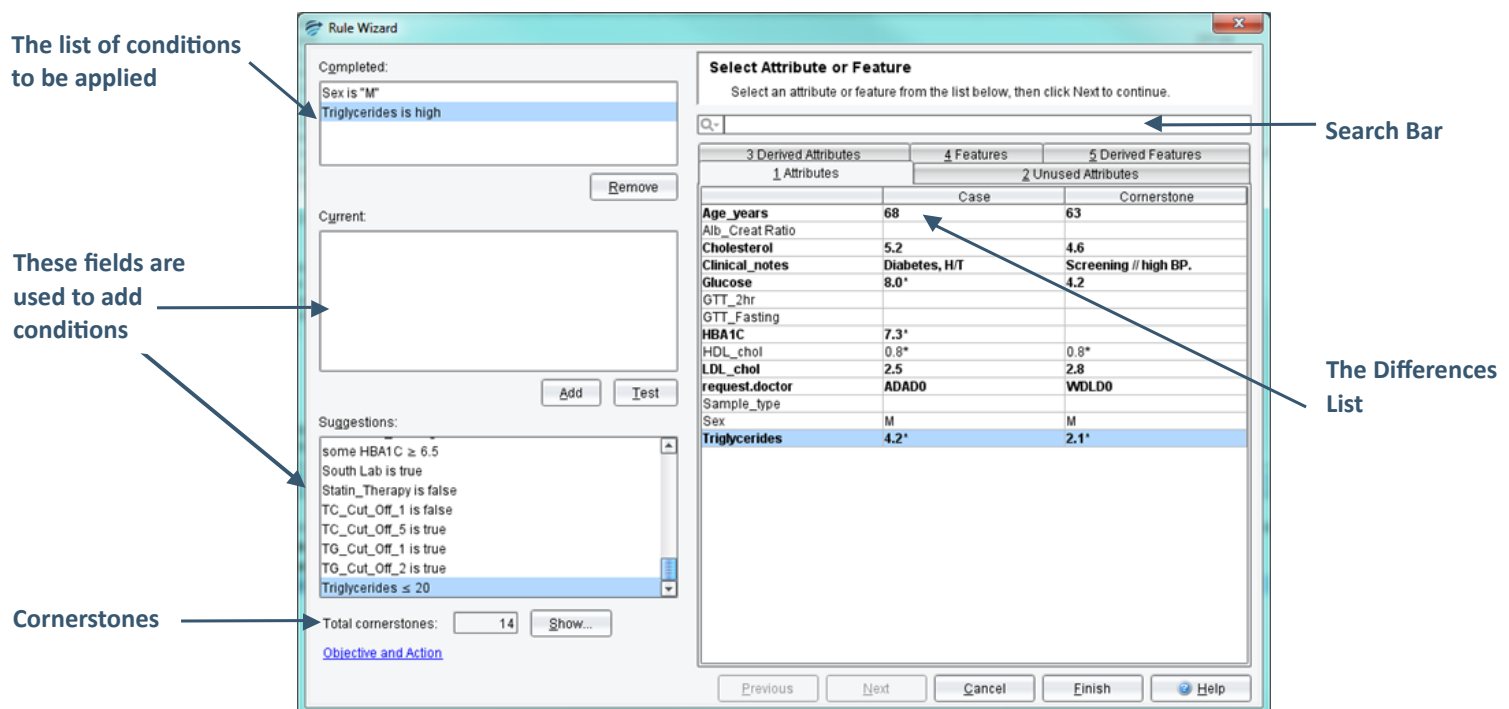


The very first item displayed in the Objectives Editor is the heading "Objectives". Underneath this heading are the top level Objectives shown in bold font.

Objectives can be created before rules are built using the Objective Editor from the main screen. Alternatively, an Objective can be created or selected at the time of building a rule by clicking on 'Select an Objective' in the relevant Rule Wizard step.

Rule Wizard - Add Conditions

Conditions are specifications that are used to justify the comments added in a rule. Each condition therefore evaluates to true or false for a case. Conditions are written in a near natural language syntax built up from various pre-defined functions. The following screen is used to add conditions:



The conditions that are to be applied will be listed in the 'Completed' pane. Further information on condition syntax is located in the next section.

There are three methods of adding conditions:

1. Using the current tab
2. Using the difference list
3. Using suggestions

Using the current tab:

1. Type the condition directly into the Current field (information on conditions can be found in the next chapter)
2. Use the Test button to determine:
 - a. Whether the condition is true for the case;
 - b. Whether the condition is true for the cornerstone case, if there is one;
 - c. How each part of the condition has evaluated.
3. If the condition is true, click the Add button.

Using the difference list:

The difference list includes the following tabs:

- Attributes (Only those that have been used in conditions)
- Unused Attributes (Only those not yet used in conditions)
- Derived Attributes (numeric range, list, text, or calculated value attributes)
- Features
- Derived Features
- Report Sections

Each tab contains a table that displays the names of each item (attributes, features etc.), their values for the current case, and their values for the current cornerstone case (if there is one). If an item's value is different for the case and the cornerstone case, it is shown in **bold type**. If an item has already been used in the current rule session, it is shown in colour in the difference list.

When an item is selected, it is automatically inserted into the 'Current' field. Pressing the 'Next' button will then show the templates that can assist the condition. Only items in the difference list that are relevant to the current rule session will appear.

With some items (derived attributes, features and derived features) you can simply double-click your mouse on the relevant item/s to quickly add them as conditions. For example, double-clicking a feature called 'Diabetic' will automatically add it as a condition if it is 'true' for the case. Otherwise the condition 'not Diabetic' will be added.

When selecting one of the tabs on the difference list, the suggested conditions will automatically be filtered. For example, when selecting the 'derived attributes' tab, only conditions that use a derived attribute will be shown as suggestions.

Using suggestions:

Suggestions are conditions that have previously been used in the Knowledge Base, and might be relevant. To add one of these suggestions, double click on the relevant item and it will be placed into the 'Completed' field.

Note: a single click will move the suggestion to the 'Current' field. Click the 'Add' button to move the suggestion to the 'Completed' field.

Rule Wizard - Check Cornerstone Cases

As conditions are added to the 'Completed' field, the conditions are compared to cornerstone cases. A cornerstone case is a case which has been used to build a rule. RippleDown keeps cornerstone cases in the 'Cornerstone case list' as a reference point for knowledge builders.

Cornerstone cases that are affected by new conditions are presented as 'Conflicting cornerstones'. The total number of conflicting cornerstones will be shown in the 'Total cornerstones' dialog box. Each of these cases must be reviewed to determine whether the rule is appropriate. To do this, simply click on the '**Show**' button to reveal the 'Check Cornerstone' screen shown below:

If the rule should be applied to the cornerstone case, then you should click the '**Apply change to cornerstone**' button.

If not, click '**Do Not Apply**' and add a condition that differentiates the rule case and the cornerstone case. This will exclude the cornerstone case and prevent the rule from affecting it.

An explanation as to how the cornerstone case and the report will be effected if the rule is applied can also be shown.

Once all cornerstone cases have been reviewed and the cornerstone count is shown as '0', you can then close the dialog box using the red X in the top right corner.

The Rule Wizard screen will then be shown again, and you will then need to click '**Finish**' to apply the rule.

The screenshot shows the 'Cornerstones' dialog box with a tabbed interface. The '1 Case and Cornerstone' tab is selected. The table displays the following data:

Case	Sex	Age	Range	01 Aug 01 08:00	22 Aug 02 00:00
1017989	M	63			
Attributes:					
Patient Details					
Age_inferences					
Clinical Notes					
Request Details					
Cholesterol					
Episodetest					
Critical_Cholesterol					
HDL_cholesterol					
LDL_cholesterol					

The 'Comments' section contains the following text:

Comments:
Dyslipidaemia is a risk factor for CVD. This pattern may be seen in the metabolic syndrome, type 2 diabetes, obesity, intercurrent if

Recommendations:
Therapeutic Guidelines
Internal

The bottom of the dialog box has a navigation bar with the following buttons: Features, Comments, Preview, Data Items, Forms, Sets, Autovalidation, Queue, Notes, Result Overwrite, Approved, Apply change to cornerstone, Do not apply, Explain cornerstone.

Case Free Rules

In addition to the traditional case-based rule building process, RippleDown also enables you to build rules without an example case, which is called “Case Free Rules”. These rules can be used in situations where one or more elements in the case data are in conflict with each other. An examples might be where a white cell profile is ‘normal’. whilst the case shows an elevated white blood cell count; or a financial class that indicates ‘non-billable’ where there is a billing code for some tests.

There are three types of Case Free Rules:

1. Mutual Exclusion Rules – when defined, consists of a list of primary attributes that should not be found together in the case.

This rule will alert the user when two or more of the defined attributes are present in a case.

The Mutual Exclusion Editor dialog box is used to define rules where certain attributes should not co-occur. It features a 'Name' field with the value 'Mutual Exclusion'. Below it is a search bar labeled '<type here to filter the list>'. A list of attributes (Pe, PH, pl, pN, Po, pP, PS) is shown on the left, with 'Add >>' and '<< Remove' buttons. On the right, a list contains 'Glucose' and 'GU glucose'. At the bottom, there are checkboxes for 'Add rule for current conflict' and 'Add rule for conflict within previous' (set to 0 days), along with 'Report Section' and 'Comments' dropdowns, and 'OK', 'Cancel', and 'Help' buttons.

2. Double Set Rules – when defined, consists of two lists of primary attributes - List 1 and List 2. The definition also allows the knowledge builder to set a complex relationship between the members of the two lists. This is based on selecting one of all/some/none from each list. In some configurations, this rule will alert the user that certain exclusive attributes are present together in the case OR, conversely attributes which should be present based on co-dependent attributes missing.

The Double Set Editor dialog box allows defining rules based on the relationship between two sets of attributes. The 'Name' field is 'Double Set'. It has two main sections: 'If' and 'Then'. The 'If' section has a dropdown set to 'some of the following are present' and a list containing 'GTT_1hr', 'GTT_2hr', and 'GTT_Fasting'. The 'Then' section has a dropdown set to 'none of the following is present' and a list containing 'Glucose dose'. Similar to the first dialog, it includes checkboxes for conflict types, report sections, and comments, with 'OK', 'Cancel', and 'Help' buttons at the bottom.

3. Qualified Set Rules – This type of rule is similar in functionality to the Double Set Rule however allows the insertion of expressions into List 1.

The Qualified Set Editor dialog box is used to define rules where a set of attributes must all be present. The 'Name' field is 'Qualified Set'. The 'If' section has a dropdown set to 'all of the following are true' and a list containing 'Lipid_Lowering_Therapy'. The 'Then' section has a dropdown set to 'all of the following are present' and a list containing 'Cholesterol', 'HDL_chol', 'LDL_chol', and 'Total Chol_HDL Ratio'. It also includes checkboxes for conflict types, report sections, and comments, with 'OK', 'Cancel', and 'Help' buttons at the bottom.

All three types of Case Free Rules can be defined for current and historical episodes. The generated comment can be allocated to any report section.

Conditions

RippleDown has a rich language for writing rule conditions for analysing numeric and text results. The functions range from simple comparisons to complex pattern matching.

When building rules in RippleDown you often want to restrict the conditions to some of the episodes of a case. For example, you might only be interested in glucose values where the patient was fasting.

The following table provides sample syntax that can be used in conditions:

Type	Condition	Example
Comparison - numeric	<, >, ≤, ≥, =	Glucose < 5.3 Glucose = 5.2
Comparison - text	is, is not, contains	Status is "Fasting"; ClinicalNotes contains "diabetic"
Existence	is available	Glucose is available
	is not available	Cholesterol is not available
Normal Range comparison	is normal is low/is high	Glucose is normal
The most recent available sample (could be current)	Most recent	Most recent Cholesterol is high Most recent Glucose is normal
Previous episode/s (i.e. before the current)	previous	Previous HbA1c is not available Previous Cholesterol is 5.3 Some previous Cholesterol is normal
All episodes	Some	Some Glucose ≥ 6.2
	None/no	No HbA1c is high
	At most, At least	At most Glucose are high At least 12 months since HbA1c is available
	Exactly	Exactly 3 Glucose are high
All episodes (including current) within a time period	Within last	Some Cholesterol within last year is normal No HbA1c within last year is available
All previous episodes within a time period (i.e. not the current)	Within last	Some previous (Cholesterol within last 6 months) is high
Change in result	After/where	Some Therapy is 'N', after Therapy is 'Y'

This table then shows how these conditions can be applied in patient test results:

		10 Jan 11	16 Jul 12	5 Dec 12	3 Mar 13
Status			Non Fasting	Fasting	Fasting
Therapy		Y		N	
ClinicalNotes		?DM			Is diabetic
Glucose	3.45 - 5.45	6.3*	5.9*	5.8*	5.2
HbA1c	3.3 - 5.7	4.1			
Cholesterol	<5.2	6.8*	5.0	5.3*	

Reports

A "Report" is the ordered set of comments provided by the Knowledge Base for a case and may be divided into report sections containing one or more comments.

The Knowledge Builder allows users to create a Knowledge Base that captures their own clinical expertise and decision making process so it can be applied automatically. The Knowledge Base can be easily updated by users with new rules to ensure that interpretations and decisions are accurate and reflect the most recent practices.

When a new project is created, there are two pre-defined Report Sections available:

1. A Report Section called "Comments" is enabled by default.
2. A Report Section called "Internal", which has the pre-defined internal format properties:
 - a. Neither the heading nor the comments in this section are included in the report.
 - b. Any comments for this section will appear in the Validator in a separate window as an "internal" comment for the reviewer.
 - c. Comments will not be released to the Online Information System if you are using the Validator to review cases.
 - d. The Internal Report Section can never be enabled by rules.

Report sections may be created, modified and ordered using the Report Section Editor which can be accessed by selecting 'Report Sections....' in the main 'Edit' menu. You can then select 'add' or 'modify' from the 'Edit' menu to show the properties of a specific Report Section. Within this component it is possible to name the report section and to select one of the pre-defined format options.

By default, report sections will appear in the order specified in the Report Section Editor. It is possible to specify that report sections should be ordered on a case-by-case basis. See the Help topic on "Report Section Ordering" for more information.

Interpretive Report

An interpretive report is created by adding comments using the rule syntax. The 'comments' section contains more detail on creating and managing comments.

The Knowledge Builder application enables users to view the current interpretation either in a "preview" mode (to view the interpretation as it is presented to the Validator or Online Information System), or in a more detailed "comments" view.

Any notes associated with the case can also be seen in the 'Notes' tab. Notes are typically memos from a reviewer in RippleDown to specify what sort of changes need to be made to the interpretation of the case before it is suitable for release.

Interpretive Actions

As well as producing interpretive reports, the Knowledge Builder can be configured to provide a defined set of actions. The same knowledge building methodology used for comments also applies to actions. The types of actions that can be generated are:

- Instructing the RippleDown Server to autovalidate cases with a particular report (refer to the Autovalidation section for more information)
- Instructing the RippleDown Server to refer cases to project or specific user queues of RippleDown Validator
- Generating alerts within the Online Information System
- Sending emails
- Producing PDF versions of the report
- Instructing the Online Information System to add a new test or test panel for a particular patient
- Instructing the Online Information System to set some data value in the patient record. For example a diabetic status flag, or billing flag

Disabling a report section

To disable a report section in the interpretation for a particular case (to stop any comments in this section from being reported), select the relevant report section heading in the "Comments" window and use the right-click menu and select 'Disable Report Section'. This will immediately open the Rule Wizard where you can follow the steps to explain why you want the report section to be disabled for that case.

Comments

The Knowledge Builder allows users to modify the current interpretation so that subsequent interpretations will give the appropriate interpretation. This includes:

- Creating and adding a new comment to the interpretation
- Adding an existing comment
- Replacing one comment with another one
- Removing a comment
- Changing the position of a comment in the report
- Showing why a comment is part of the interpretation
- Showing why a comment was not removed from the interpretation
- Showing other comments that have been removed from the interpretation
- Adding or removing a report section that groups related comments

Functions	Steps
Add	<ol style="list-style-type: none">1. Select the "Comments" window and use the right-click menu and select 'Add comment'2. Select the comment you want to add, or create a new one3. Use the 'Objective' screen to explain the reason for adding the comment
Replace	<ol style="list-style-type: none">1. Select the comment you want to replace, then use the right-click menu and select 'Replace comment'2. Choose the comment you want to be given as the replacement. You can also create a new comment as the replacement3. Use the 'Objective' to explain the reason for replacing the comment
Remove	<ol style="list-style-type: none">1. Select the comment you wish to remove and use the right-click menu and select 'Remove comment'2. The Rule Wizard will jump immediately to the steps that allow you to give the reasons why you wanted the comment removed.
Re-add a removed comment	<ol style="list-style-type: none">1. Use the 'Show Removed Comments' option in the Tools menu. Any comments that have been removed from this interpretation will be shown in blue in the "Comments" window.2. Select the comment you want to re-add, then use the right-click menu Re-add comment. The Wizard will jump immediately to the steps that allow you to give the reasons why you wanted the comment re-added.

Comment Editor

The Knowledge Builder also enables users to manage the list of all comments in the Knowledge Base. Comments are grouped by "report section".

Within the 'Comments Editor' users can:

- Add a new comment
- Insert a variable expression into a comment, e.g. referring to a test value
- Edit a comment
- Supersede one comment by another
- Delete a comment, provided it has not been used in rules
- Move a comment from one report section to another
- Reorder comments within a report section
- Filter facility to search for existing comments

It is also possible to import or export comments to and from external text files into RippleDown. To do this use the 'Edit menu' in the 'Comment Editor' and select 'Import values...' The comments must be in a text (.txt) format.

Variable Comments

The Knowledge Builder allows users to create comments that have embedded variables. When a case is interpreted, the specific value for these variables can be inserted into the report.

This feature enhances the ability of the Knowledge Builder to generate reports that are highly specific to the patient, referring doctor, and test results. For example: *"Last ferritin level 118 ug/L noted over 3 months ago on 29/May/13."*

These variables can be embedded into the comment as follows:

"Last ferritin level {last available Ferritin} ug/L noted over 3 months ago on {date at last available Ferritin}."

Internal Comments

Internal comments are intended to be seen in the process of reviewing reports in RippleDown Validator, and are **not** intended to be released. Typically these comments indicate a workflow instruction, a problem in the case, or an alert situation.

To create an internal comment for a case, use the Rule Wizard to add a comment in the relevant Report Section called "Internal".

If reports are being reviewed in RippleDown Validator, any internal comments will be shown in a separate window and will not be sent to the External Information System once the report is approved. These comments do not need to be deleted.

If reports are being reviewed using the facilities of the External Information System, any internal comments will be appended to the report and so will have to be manually removed by the person reviewing the case.

If an internal comment is added to a case it will never be autovalidated.

Tables

If you have specified that the report should be produced as a PDF document, you are able to insert a table into the comment section.


The table will be formatted as four columns containing attributes from the case and their corresponding values. The first and third columns will show the text you have entered to describe the attribute, and the second and fourth columns will show the values for those attributes for the case.

The table is added using the Rule Wizard in the same way as you would insert a text comment. Any attributes, either primary or derived, can be used in the table.

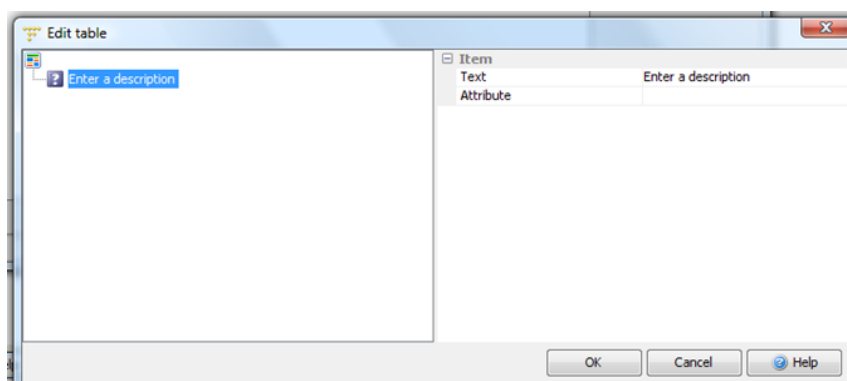
If the table is not in the correct place in the report section, you can re-order it using the 'Comment Editor' in the same way as you would re-order a text comment. That is, using the 'Move | Move up' or 'Move | Move down' menus.

To add a table:

1. Right click in the Report Section where you want the table to appear.
2. Select 'Add comment'
3. Click the 'Create new table' button and a blank table will appear as follows.



4. Enter a table name and description in the right-hand panel.
5. Right-click in the left hand pane and select the menu item "New item".



6. Enter a description in the 'Text' field. This will be the text that appears in the table representing the attribute.
7. Select the attribute by clicking the 'Attribute' field. This will bring up a list of all attributes in the project in alphabetical order. You can either scroll through this list to find the attribute you need, or else type the first few characters of the attribute name which will automatically search for that attribute.
8. Add further attributes as required by right-clicking on the left pane.
9. Attributes in the table can be ordered by dragging and dropping them with the mouse in the left hand pane.
10. Click 'OK' to finish creating the table. It will now appear with the list of comments for that report section.
11. The table can be ordered in the report section with respect to other tables or comments using the 'Move up' or 'Move down' menu items.
12. Click 'OK' and continue adding conditions. The most common scenario is to add the table unconditionally i.e. to all reports.
13. To view the table format use the 'Preview Tab' of 'Case View'.

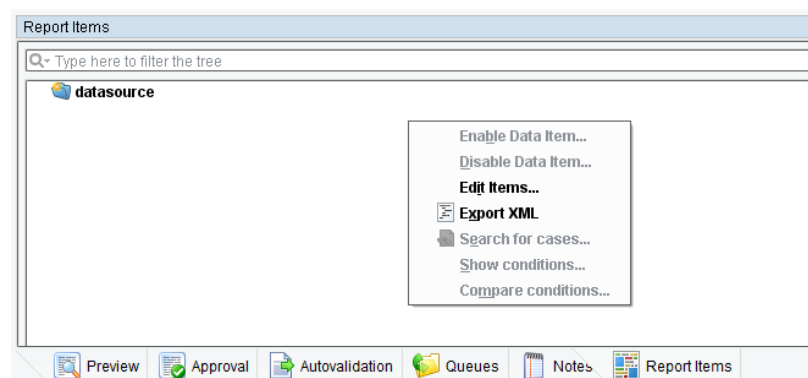
To edit an existing table

1. Open the Comment Editor by selecting 'Comments' in the 'Edit' menu.
2. Choose the report section containing the table.
3. Select the table in the list. You can scroll the list or select 'Edit' in the 'Find' menu.
4. To edit the table, double click on it or select 'Modify' from the 'Edit' menu.
5. To delete an attribute, select the relevant one from the left hand pane and press the DELETE key. Alternatively use the right-click menu item 'Delete'.
6. To add a new attribute, use the right click menu item "New item". Fill in the 'Text' and 'Attribute' fields as you would if you were creating a new table.
7. To edit the text of an existing attribute in the table, simply select the relevant attribute and edit the 'Text' field.
8. To re-order an attribute, drag and drop it to its required position in the left hand pane using the mouse.
9. Click 'OK' to finish editing the table.
10. The changes made to the table will be immediately applied to all subsequent reports where this table is included.

Report Items

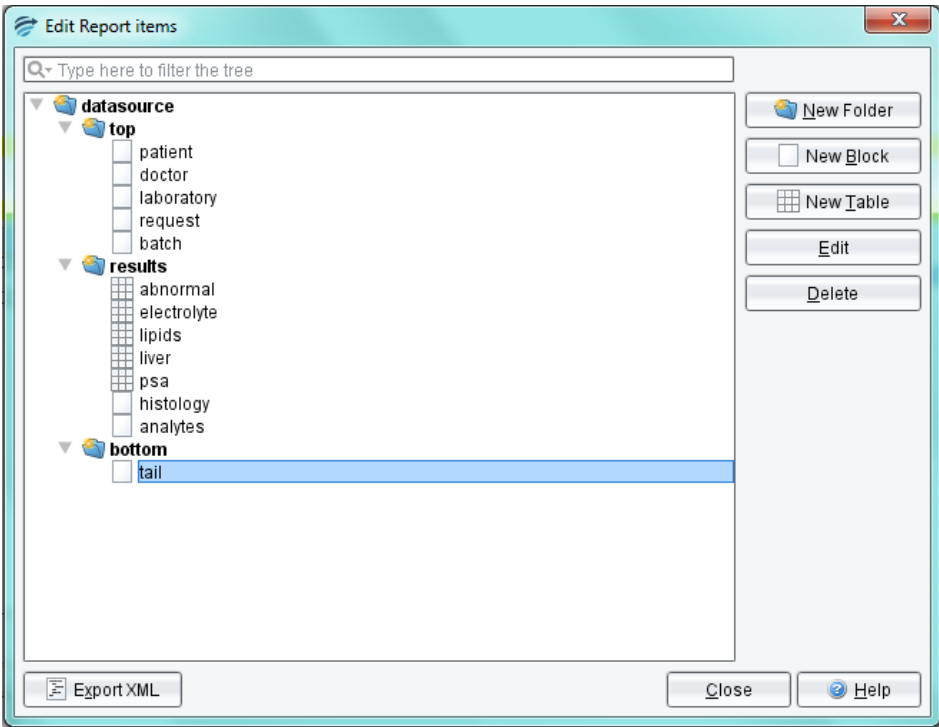
The 'Report Items' tab is used to generate an XML data source or schema for the reporting engine and report design tool.

This 'Report Items' docking pane is not shown by default. To display this pane select 'Report Items' from the 'Windows' menu. Then right-click in the tab and select 'Edit items...' as shown:



The 'Edit Report Items' dialog box will appear to display and/or edit the schema that consists of block and table items. Folders may be used to organise the report items.

The example shows three folders: top, results and bottom. The folders contain block items such as patient, and/or table items such as electrolyte.



Several functions are available for use:

Functions	Descriptions
New Folder	Allows report items to be grouped. The top level folder is always "datasource". Folders create tags in the exported XML schema. Note: Do not use spaces when naming folders or report items.
New Block	To define a new block item (a set of attributes for which the current value is to be displayed in the report)
Table item	Define a new table item (a set of attributes for which the current and historical values are to be displayed in the report).
Edit	Used to modify existing items
Delete	To permanently delete a report item
Export XML	Once all items are defined click on Export XML and select an appropriate location for the exported files. This file will be used as a data source by the reporting engine

Block Items

A block item is used when the current values of the attributes in the block item are to be printed in the report. The 'Edit Block Item' dialog box is used to name the block item and to add appropriate attributes. The example shows a block item for patient:

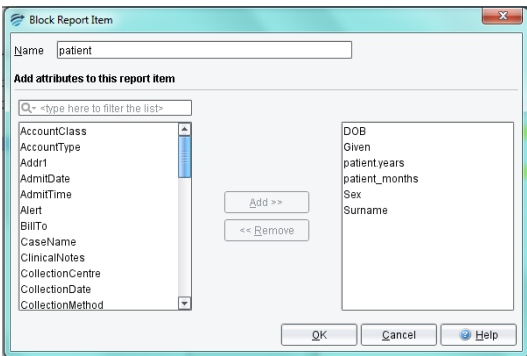
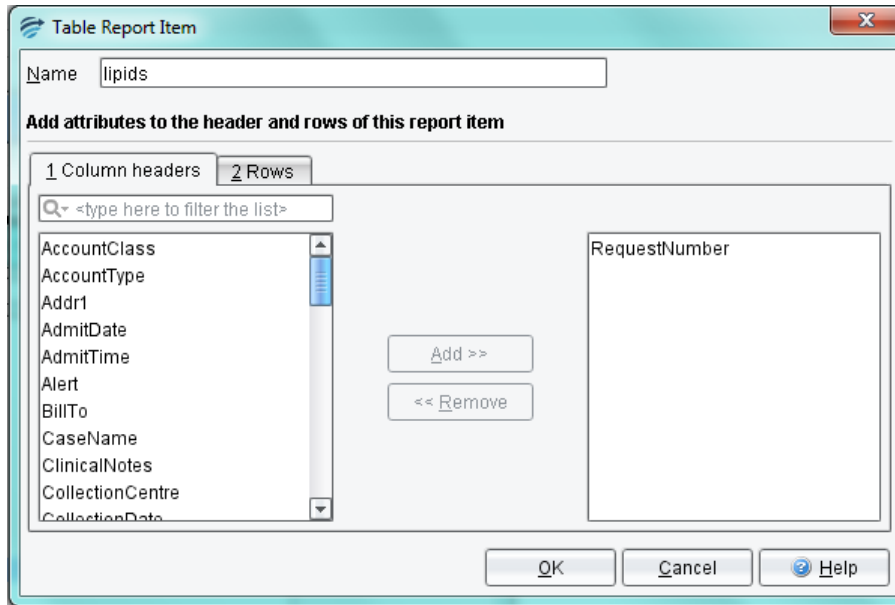


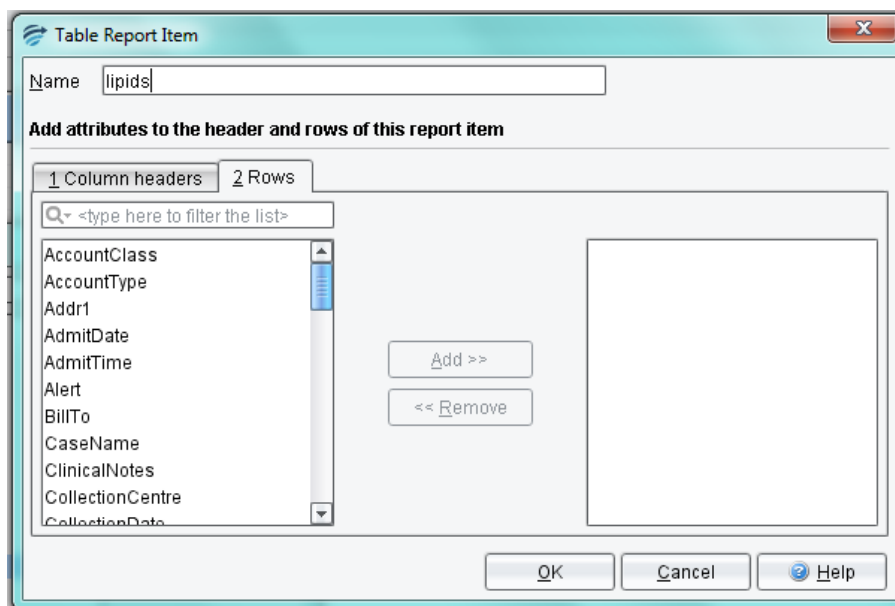
Table Items

Table items are used to display both current and historical attribute values in the report. The 'edit table items' dialog box is used to name the table item, as well as allocate table and row data. The example shows a table item for Lipids with the first tab listing 'Column headers' and the second listing 'Rows'.



By default, RippleDown inserts the episode date and time into the table.

The 'RequestNumber' is added as an additional item.



The rows will be populated by the primary attributes included in the group attribute called Lipids.

Report Items and Rules

By default, a new report item is enabled and appears as a report data source. However any report item may be disabled using a rule. An example might be in a situation where certain report items are not required by a particular GP practice. To disable a report item, right-click on the item and select 'Disable Data Item'. The Rule Wizard will then guide you through the process. Once complete, the disabled data item will appear in red.

Any disabled item can be re-enabled using the Rule Wizard.

XML EXPORT

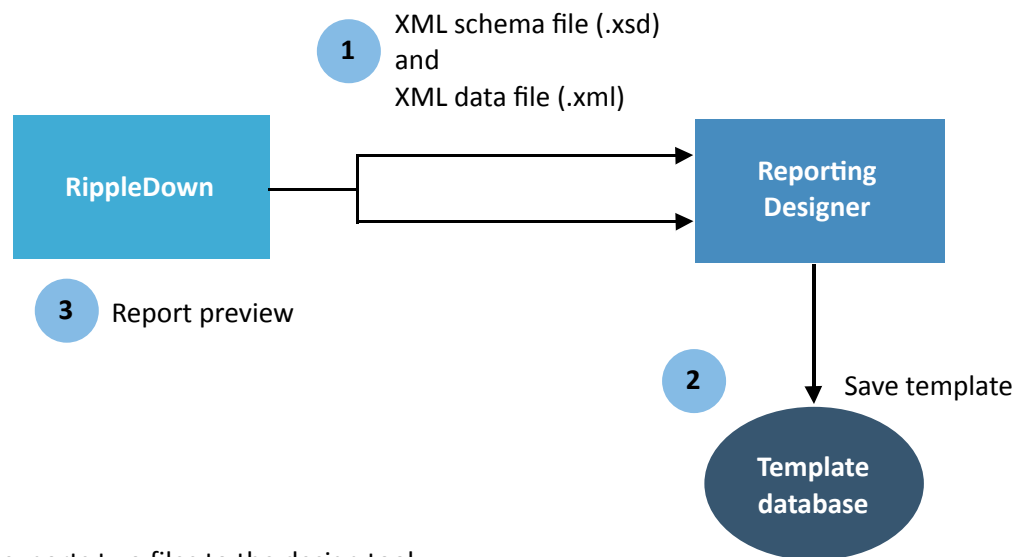
Once the report items are defined they can be made available for both a report designer tool and a report engine.

You can use the 'Export XML' button to export both the XML schema (and XSD file) and an example XML file for use by a report designer tool. The values for the report items will be taken from the case currently on display in the Knowledge Builder.

Template Design

Templates for pdf reports are created in a Report Design tool such as iReport and Adobe → LiveCycle → Workbench Designer using data defined in RippleDown.

Design Workflow



1. RippleDown exports two files to the design tool:

- the .xsd - a schema file
- and the .xml - contains data from the current case.

2. Template is created and saved in the template database.

3. To preview the pdf report in RippleDown open the 'Preview' tab and select the 'Pdf Preview' button. This action will export an XML file of the current case to the report engine which will then select the specific template, and return the pdf file to RippleDown.

Report Preview

In RippleDown it is possible to preview the PDF output generated by the reporting engine. This function requires configuration of the RippleDown server.

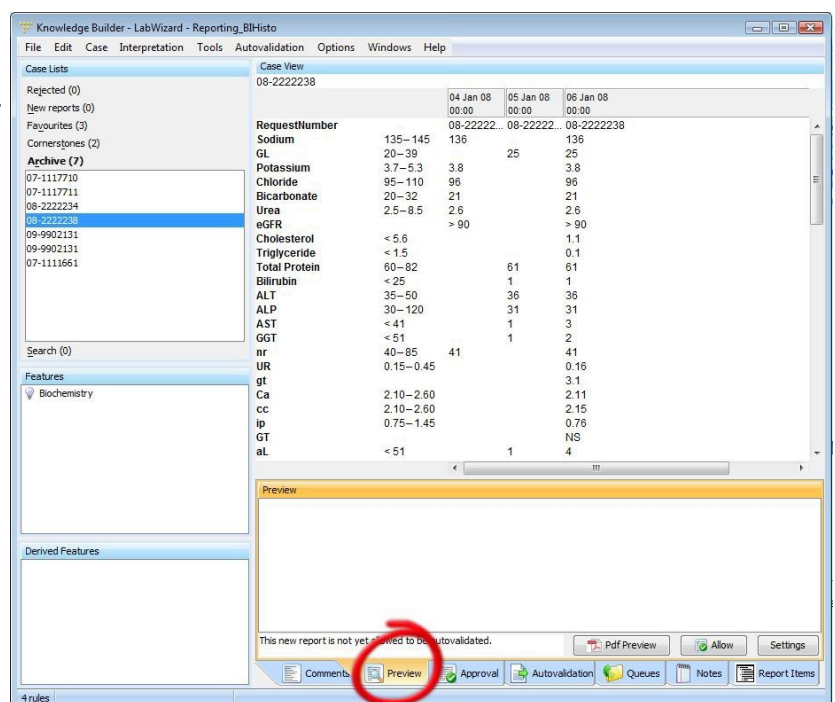
To do this:

1. Open the RippleDown Administrator client.
2. Use the 'Tools' menu and select 'Edit server settings'.
3. Enter the URL of the print service and click 'Apply'.
4. Enter the Username and Password for this service.
5. Close the Administrator.

In the Knowledge Builder, select the '**Preview**' tab (adjacent to the 'Comments' tab as shown):

Click the '**Pdf Preview**' button. This action will cause RippleDown to send an XML export to the reporting engine.

A PDF report will then be created and returned to RippleDown.



Project Explorer

Project Explorer is an invaluable tool that shows the history, and all major components of the project with a tracking user ID, date, and time stamp.

To start Project Explorer, select 'Project Explorer' from the 'Tools' menu. This may take a minute or two depending on the size of the project - a progress bar will be displayed at the bottom left hand corner of the screen.

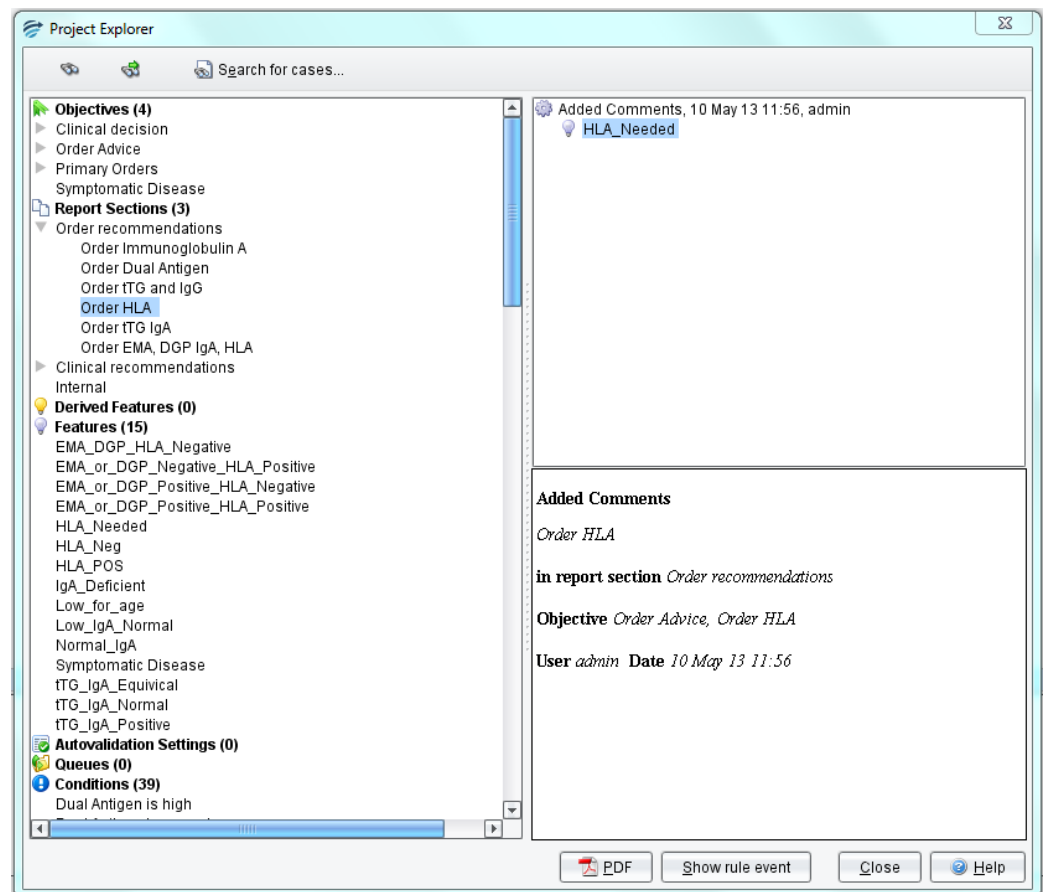
Each of the sections are listed in the index, along with the number of each in parenthesis.

To see the details of a section double click on the name.

The example shows the Report Section "Order recommendations" and the comment: "Order HLA"

In the upper right pane it shows the history of the comment and all the reasons for it being added or removed. The lower right pane displays the location and tracking details of the comment.

To view the specific details associated with the rule event select the 'Show Rule Event' button. A window will be displayed, showing the case on which the rule was made (i.e. the cornerstone case). The window will have two tabs, the first showing the case details before the rule was created, and the second showing the case details after the rule was created.



Expert Project Explorer

The Project Explorer may be exported in a PDF format:

1. Click the 'Pdf' button
2. Enter a name making sure it is suffixed by '.pdf'
3. Select a location and click 'Pdf'

Autovalidation

Autovalidation describes the process for when cases receiving particular types of interpretations are automatically released to the Online Information System without the need for human intervention.

Autovalidation may be appropriate when:

- You have reviewed a large number of cases giving a particular interpretation, and that interpretation is now always correct whenever it is given;
- The rules giving a certain type of interpretation are so general and simple that only minimal review is needed;
- The project is not being used to provide reports, but to automate processes in the laboratory such as reflexive testing or billing actions.

Managing Autovalidation

There are three ways of managing autovalidation: by reports; report sections; or by rules.

1. For "per report", a percentage autovalidation level is set for each unique report, regardless of the number of report sections that is generated by the knowledge base. This is the most restrictive policy, but the safest.
2. For "per report section", a percentage autovalidation level is set for each unique report section that is generated by the knowledge base. This is less restrictive than the "per report" policy, and is used if the report sections are relatively independent.
3. For rule-based autovalidation, use the Rule Wizard to add, modify or remove autovalidation settings that are given by the project, in exactly the same way as you would if you were adding, modifying or removing comments using rules. This is the most flexible approach.


For any particular project, the choices of "per report" or "per report section" are mutually exclusive. However, autovalidation by rules will take precedence over any "per report" or "per report section" based levels.

Type	Description
Report	A Report is the ordered set of comments provided by the Knowledge Base for a case.
Report Section	A Report Section is a component of a report. It consists of a series of comments and may or may not contain a heading (depending on settings defined in the report Section Editor).
New Report	A new report is automatically created when a comment is first given for some case, or when a new combination of comments is first given.

Allow Autovalidation

Autovalidation cannot occur unless the report or report section has been permitted for autovalidation. This can occur in one of three ways:

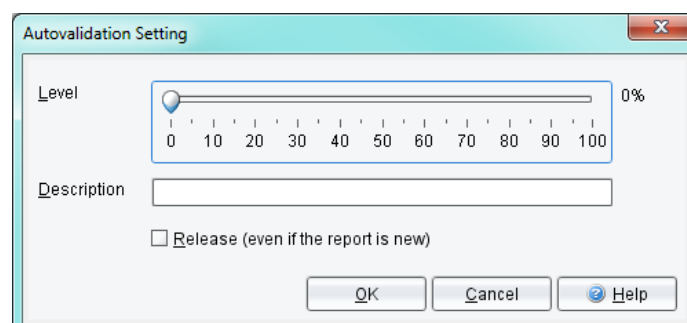
1. By selecting 'Allow' in the 'Preview' pane in the Knowledge Builder
2. Accessing the autovalidation controls by opening the 'Report Settings' and double clicking the specific report which will open the screen shown below. Then click the checkbox to allow autovalidation:



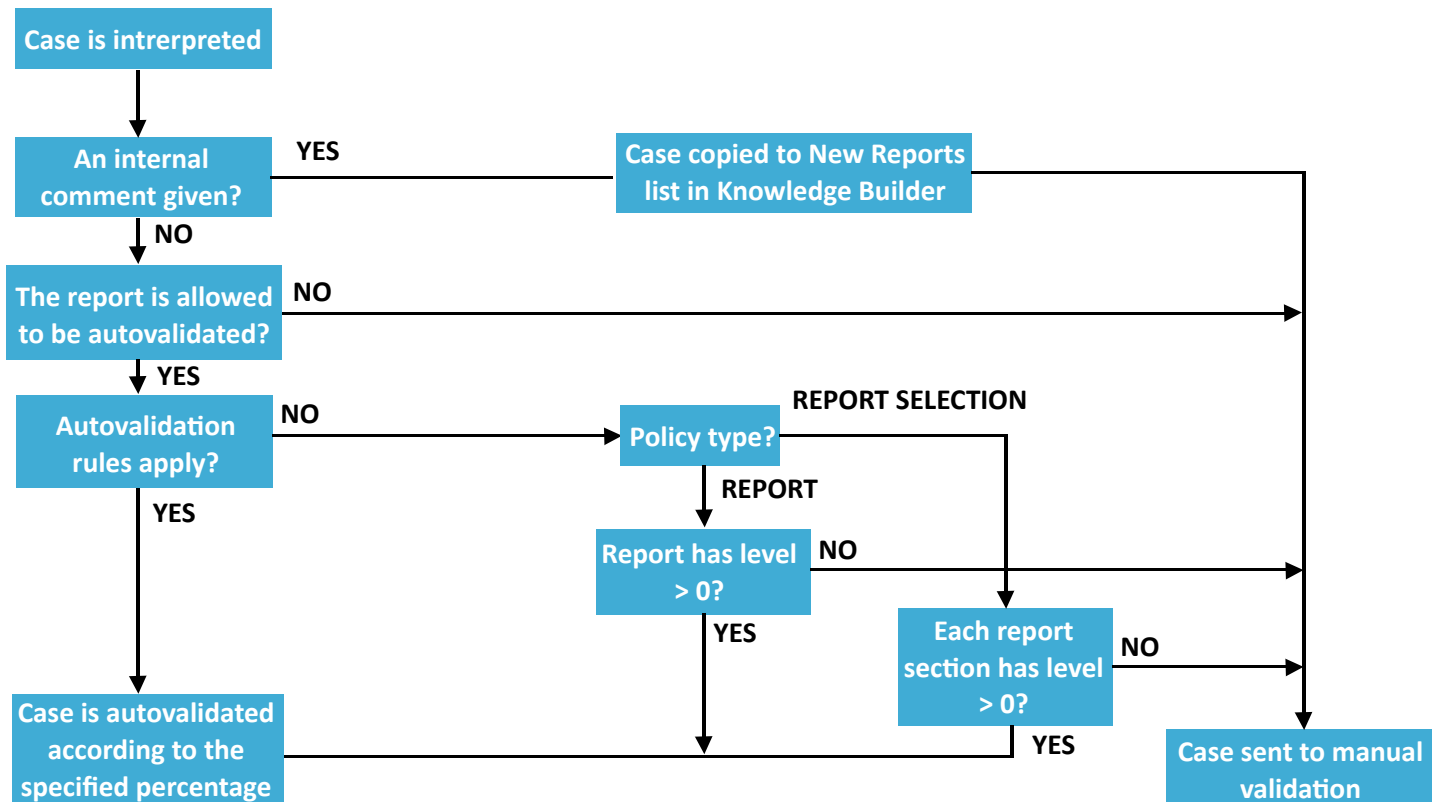
3. Using the RippleDown Validator to approve the case through normal lab workflow (this will only apply if the user has Knowledge Builder privilege)

Autovalidation setting to release a new report

There is a check box in the autovalidation setup screen that will release a report and apply the set autovalidation rate even if it has not been specifically allowed. This will occur once the report has been reviewed by a Knowledge Builder user or a RippleDown Validator user with the corresponding permission.



Autovalidation Options



Autovalidation Levels

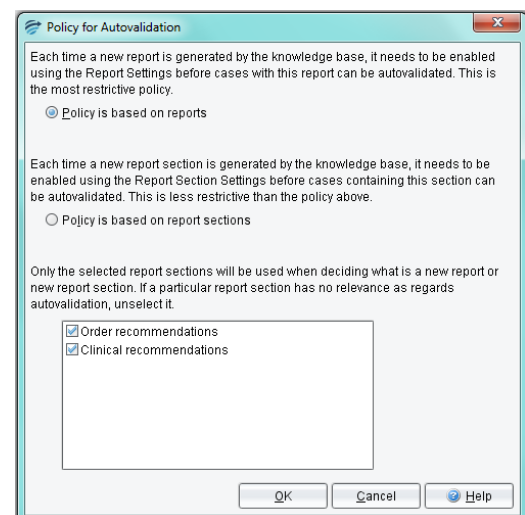
Users are able to define an autovalidation levels between 0% and 100% for each report, report section, or autovalidation rule setting.

Note: In most situations, it is advisable not set an autovalidation level above about 95%. This ensures that a small number of cases continue to appear for review in validation to enable an continual quality audit of the knowledge base.

Policy

The choice between the "per report" and "per report section" policy is done using the 'Policy' menu in the 'Autovalidation' menu.

The default setting for a project is "Policy based on reports" with all report sections being selected.



Policy based on reports

This option allows each unique report to be considered for autovalidation once the report has been allowed for autovalidation consideration.

To set autovalidation levels on a report-by-report basis:

1. Set the 'Policy to Policy' based on reports in the 'Autovalidation' section of the 'Policy' menu.
2. Use the 'Autovalidation' menu in 'Report Settings' to show a table of all reports that have been recently given by the project - i.e. within the last month.

3. Review each report in the table. Reports that already have some level of autovalidation set will have a tick in the 'Allowed' column, and possibly a percentage in the 'Level' column. Note that clicking on a column heading will sort the table by the values in that column. For example, clicking on the 'Given' column will order the reports according to how often each report has been given by the project in the last month.

4. The 'Rejected' column shows how many times the corresponding report was corrected during validation. Reports that are never, or rarely, rejected are possible candidates for autovalidation.

5. If you decide to set a non-zero autovalidation level for a report, click the 'Edit' button and set the desired autovalidation level for that report. The 'Allow autovalidation' check box will be checked automatically if you set a non-zero autovalidation level.

Allowed	Level	Description	Case	Given	Manual	Rejected
✓	10%	likely false-positive aTG		0	0	0
✓	15%	Celiac Disease is Unlikely		0	0	0
		Order Immunoglobulin A		0	0	0
		Order Dual Antigen		0	0	0
		Order EMA, DGP IgA, HLA Likely false-ne...		0	0	0
		Order HLA		0	0	0
		Celiac Confirmed		0	0	0
		Order HLA Celiac Disease is likely		0	0	0
		Order EMA, DGP IgA, HLA		0	0	0
				0	0	0
		Perform Biopsy		0	0	0
		Order tTG IgA		0	0	0
		Order tTG and IgG		0	0	0
		Celiac Disease is UnlikelyExclude histor...		0	0	0
		Likely false-negative HLA testConsider b...		0	0	0

Full description
Celiac Disease is Unlikely

In the last 31 days
Total: 0 Manual: 0 Auto: 0 Rejected: 0 None: 0

Changes made to autovalidation settings are effective immediately.

Policy based on report sections

This option allows each unique report section to be autovalidated. This is appropriate when comments given on one report section are independent of comments given in other report sections.

To set autovalidation levels on a report section basis:

1. Set the 'Policy to Policy' based on report sections in the 'Autovalidation' section of the 'Policy' menu.

2. Use the 'Autovalidation' menu in the 'Report Section' settings to show a table of all report sections that have been recently given by the project - i.e. within the last month.

3. Review each report section in the table. Report sections that already have some level of autovalidation set will have a tick in the 'Allowed' column, and a percentage in the 'Level' column. Note that clicking on a column heading will sort the table by the values in that column. For example, clicking on the 'Given' column will order the report sections according to how often each report section has been given by the project.

Description	Case	Given	Manual	Rejected	Section
likely false-positive aTG		0	0	0	Clinical rec...
Perform Biopsy		0	0	0	Clinical rec...
Order Immunoglobulin A		0	0	0	Order reco...
Celiac Disease is Unlikely		0	0	0	Clinical rec...
Order EMA, DGP IgA, HLA		0	0	0	Order reco...
Order HLA		0	0	0	Order reco...
Celiac Confirmed		0	0	0	Clinical rec...
Celiac Disease is likely		0	0	0	Clinical rec...
		0	0	0	
Order Dual Antigen		0	0	0	Order reco...
Order tTG IgA		0	0	0	Order reco...
Order tTG and IgG		0	0	0	Order reco...
Celiac Disease is UnlikelyExclude historyCons...		0	0	0	Clinical rec...
Likely false-negative HLA testConsider biopsy ...		0	0	0	Clinical rec...

Full description
Celiac Disease is Unlikely
Exclude history
Consider HLA kth4lh

In the last 31 days
Total: 0 Manual: 0 Auto: 0 Rejected: 0 None: 0

4. The 'Rejected' column shows how many times the corresponding sub-report was corrected during validation, that is, was inappropriately given by the project. Report sections that are never, or rarely, rejected are possible candidates for autovalidation.

5. If you decide to set a non-zero auto-validation level for a report section, click the 'Edit' button and set the desired autovalidation level for that report section. The 'Allow autovalidation' check box will be checked automatically if you set a non-zero autovalidation level.

Changes made to auto-validation settings will be effective immediately.

Autovalidation by Rules

To use rules to manage autovalidation, you add an autovalidation setting to a case using the Rule Wizard, in exactly the same way as you would add a comment to a case.

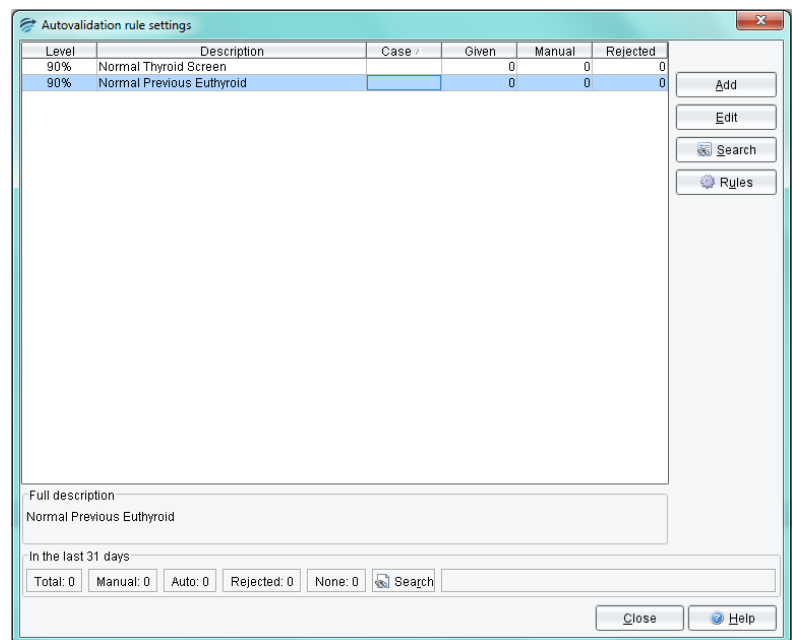
1. Find a case (perhaps in the 'Archive' or the 'Rejected' case list) which is representative of the type of case that could safely be autovalidated at some level.
2. Click on the 'Autovalidation' pane (adjacent to the 'Preview' pane)
3. Right click to start the Rule Wizard to add an "autovalidation setting", consisting of a percentage level and a textual description. You can select an existing setting, or create a new one.
4. Create or select an objective for this rule.
5. Add conditions for the rule, check cornerstone cases, and finish the rule session.

To replace one autovalidation setting for a case with another using a rule, you select the current setting, then use the Rule Wizard to replace it with another one. This is the same process you would use to replace a comment given for a case, giving only the conditions relevant to the change.

Finally, to remove the autovalidation setting for a case, you select the current setting and use the Rule Wizard to remove it.

You can see the list of all possible autovalidation settings using the menu from the main screen select 'Autovalidation' and then 'Autovalidation rule settings...'

This facility also allows you to edit the level of an autovalidation. The new level will apply to all cases with that autovalidation setting.



Notes:

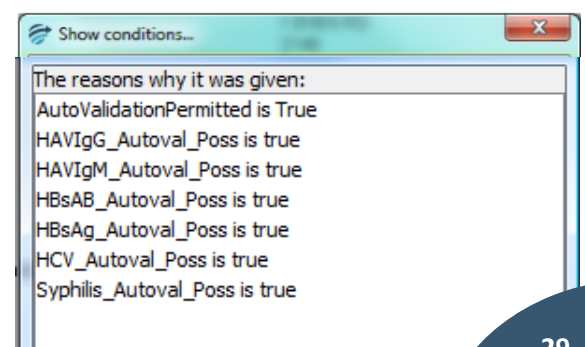
- If more than one setting is given for a case, the setting with the lowest autovalidation level will apply.
- Autovalidation settings using rules will take priority over any levels that have been set using reports or report sections.
- Even if an autovalidation setting is given for a case, it will **not** be autovalidated if the report for the case has yet to be allowed (or report section depending on the policy).
- Autovalidation will not take effect if an earlier occurrence of the case exists on a RippleDown Validation queue awaiting manual validation. NB: Interpretive comments **will** be updated.

Suggestion for setting Auto-validation by rules: Decide the characteristics of a report that may be suitable for Autovalidation e.g. all normal results or all negative results and/or relevant locations

1. Build a series of features that apply for the normal/negative situations
2. Build a derived feature that will apply when all the features apply
3. Build a rule to allow autovalidation based on the presence of the derived feature.

In the example above the Derived Feature allowing autovalidation is applied under the following conditions:

These are all features describing autovalidation criteria for the separate tests. All of the criteria must be satisfied for the autovalidation setting to be given.



Manual Validation and Queues

Whenever a case is processed by the RippleDown server it is assigned one of two possible status flags:

- Autovalidated (i.e. ready for release)
- Manual validation required (i.e. must be approved by a user before release)

If the case has been autovalidated, it is sent to the Online Information System with a status flag indicating that the report has been approved and is therefore ready for release.

Conversely, if manual validation is required, the following options are available for reviewing for the case, depending on how the Knowledge Base has been setup:

1. If a review queue has been defined by the Knowledge Base for this case, then the case will be queued to that special-purpose queue on the Validator
2. Otherwise, if the Knowledge Base has been setup by the RippleDown Administrator to use the RippleDown Validator by default the case will be sent to the default project queue on the RippleDown Validator,
3. Otherwise, the case will be sent to the Online Information System with the appropriate status flag indicating that it still needs to be approved. The facilities of the Online Information System can be used to review and possibly correct the report prior to release. This is known as external validation.

Searching

Find a Case in a Case List

To search a case by name, click on the 'Case List' title and start typing the case name. The "*" may be used in a random search.

Find Cases by Condition

You can search for cases that satisfy some condition, even if that condition has not been used in a rule.

To do this:

1. Select the case list you want to search (you can select any list, including the Search list),
2. Use the 'Case' menu or right-click on the Case List and select 'Search by condition'
3. Enter the condition you want to use (it does not have to be true for the currently selected case)
4. Press 'OK'

By default, each new search will clear any cases that were previously in the 'Search' list

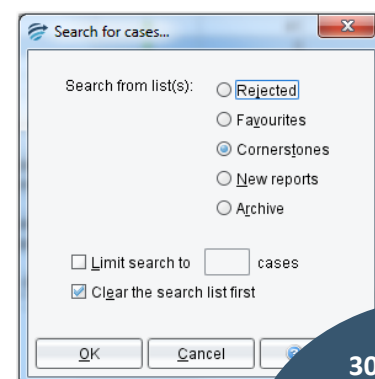
If you do not want to clear cases on the Search list before you do the search, then uncheck the option to "Clear the search list first". Any cases that are found will then be copied to the end of the Search list. In this way you can refine your search by using another search condition

Find Cases by Comment

You can search for cases that give a specific comment, feature or derived feature. If the comment (or feature) is given by the current case, right click on the comment and select the pop-up menu 'Search for cases'. The following dialog box will then appear allowing you to refine the search:

Any cases that are found will be copied to the 'Search' case list.

If the comment (or feature) that you are interested in is not given by the current case, open the project explorer and select the comment or feature, and click the 'Search' tool on the tool bar.

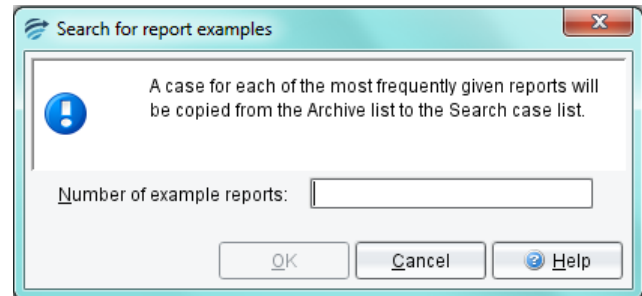


Find Cases for the most frequently given reports

A good way of identifying what comments a project is generating, is to search for example cases for the most frequently given reports.

To start a search, select the 'Tools' menu and click 'Search for report examples'. A dialog box will appear prompting you for how many reports you want examples of, as shown below:

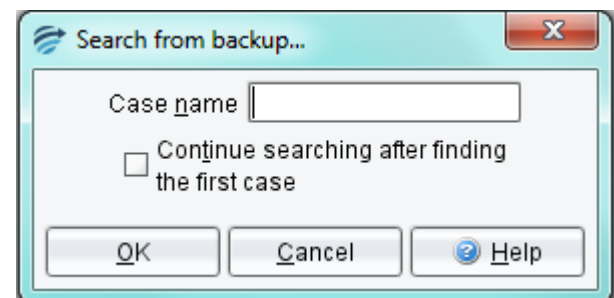
Once you enter this number, the system will start analysing cases in the Archive case list and will put copies of them in the 'Search' case list. Cases with the most frequently given reports will be first on the search list.



Find Cases from Backup

It is possible to search for a backed up copy of a case that has been interpreted by a project. To do this, activate the 'Case' menu and select 'Search from backup...'

A dialog box will appear, prompting you for the name (lab number) of the case. When you click 'OK', the server begins the search as a background process. A progress bar will be shown in the bottom left hand corner of the screen to indicate how many of the backup files have been searched. As the cases are found, they are put into the 'Search' case list.



Sometimes a case is sent several times by the Online Information System, over the course of a few days. If you want to see all of these versions of the case, choose the option 'Continue the search after the first case is found'. This will search all backups, and thus take longer, but can be cancelled at any stage.

Performance Reports

Comment Statistics

To see statistics on a particular comment given by the project, select 'Autovalidation' in the 'Comment statistics' menu. This will show a table of all comments that have been recently given by the project - i.e. within the last month.

Description	Case ▾	Given	Manual	Rejected
Celiac Disease is Unlikely	●	0	0	0
Order tTG IgA		0	0	0
Order Dual Antigen		0	0	0
likely false-positive aTG		0	0	0
Celiac Disease is likely		0	0	0
Order EMA, DGP IgA, HLA		0	0	0
Perform Biopsy		0	0	0
Likely false-negative HLA testConsider biopsy to ...		0	0	0
Order tTG and IgG		0	0	0
Celiac Disease is UnlikelyExclude historyConside...		0	0	0
Celiac Confirmed		0	0	0
Order HLA		0	0	0
Order Immunoglobulin A		0	0	0

Full description
Celiac Disease is Unlikely

In the last 31 days
Total: 0 Manual: 0 Auto: 0 Rejected: 0 None: 0 Search

Close Help

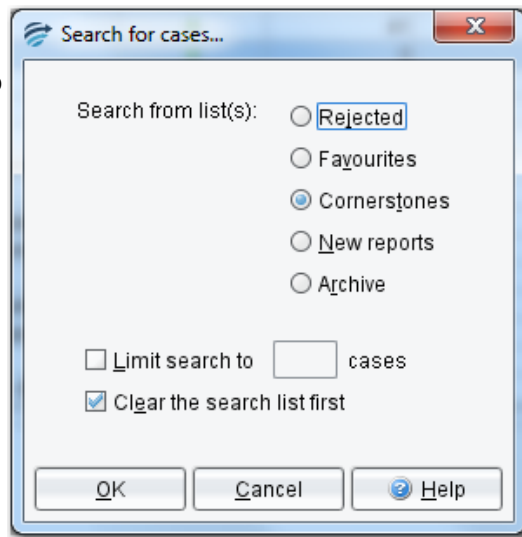
The statistics for each comment are:

- the number of cases with this comment that have been given
- the number of cases with this comment that have been manually validated
- the number of cases with this comment that have been rejected

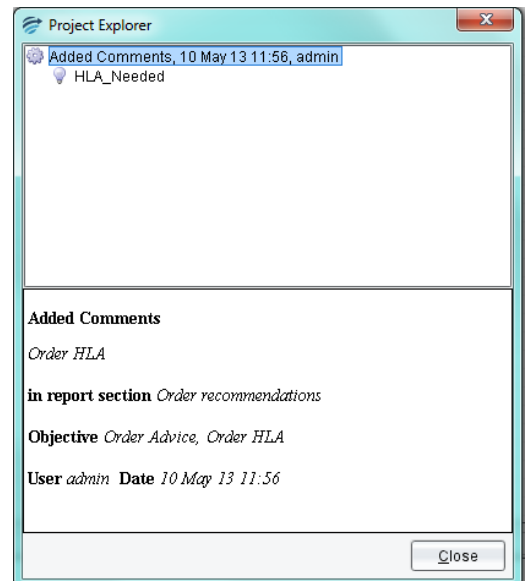
The summary panel also shows the number of cases that did not receive any comment within the last month.

The 'Search' button will activate a search facility for the comment that is highlighted. Simply select the case list for the search and enter a limit, if required. Click 'OK' to commence the search.

The cases that are found will be added to the Search Case List.



The 'Rules' button shows all the rules in the project that add or remove the comments:



Simulation Reports

RippleDown has a Simulation tool that enables you to analyse the interpretations of all cases in any one of the case lists, such as the Archive or Cornerstone list.

Simply select the list you want to analyse, then use the 'Tools' menu to select 'Simulation Report'. A dialog box will then appear that contains the following statistics:

- A summary statistic showing the number of cases that have an interpretation
- A set of detailed statistics for each:
 - Feature
 - Derived Feature
 - Comment
 - Report Section
 - Auto-validation setting
 - Review Queue
 - Report

The detailed statistics are arranged in a set of tables, with each table containing a particular item (e.g. Feature) and the associated count (i.e. the number of cases giving that Feature). You can sort the information in each table by clicking on the column heading.

If you select any item from one of the tables (e.g. a Feature), the Search button becomes enabled, allowing you to search for cases that contain this particular item. You can search any case list for the specified item, not just the list used for the Simulation Report.

The results of the search will go into the Search list on the main screen. You are able to access the main screen even while the Simulation Report dialog box is showing.

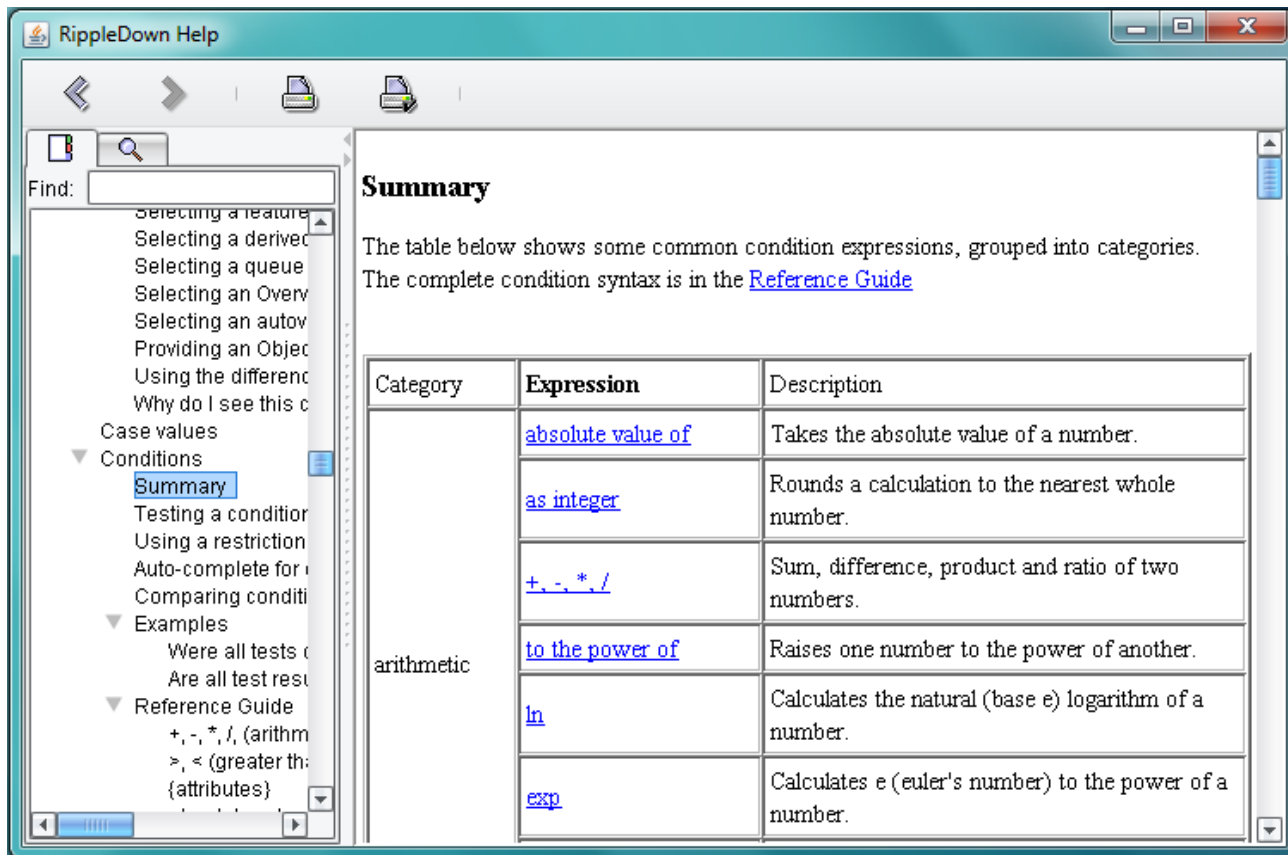
The 'Export' button enables you to export the Simulation Report to an Excel spreadsheet for further analysis.

NB: Ensure the name given to the file is suffixed with .xls

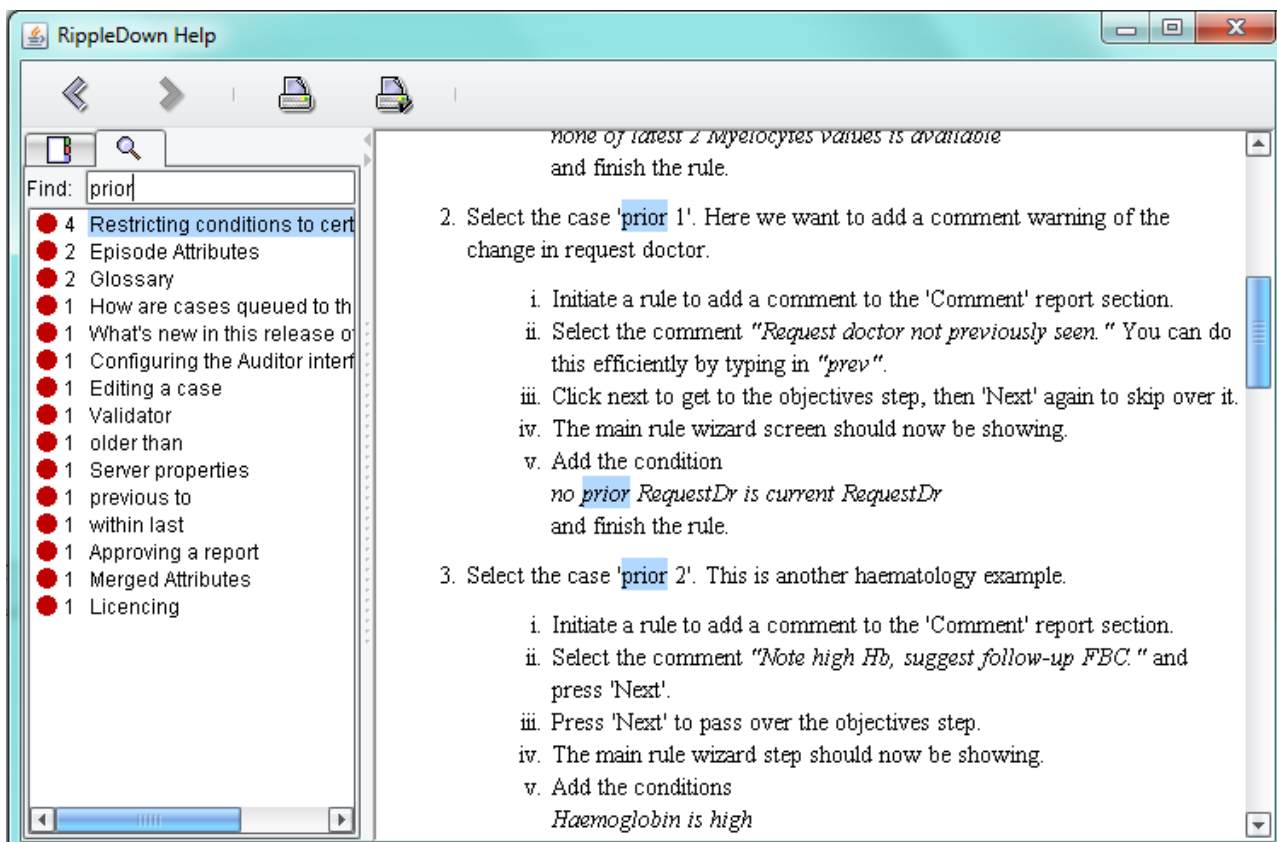
Help

The help facility can be launched from the main screen and from within the various application components by selecting the 'Help' button.

You are able to use the search facility using the index:



Or the search facility as shown:



For additional support, visit:
www.pks.healthcare/support

www.pks.healthcare/support

About PKS

Pacific Knowledge Systems (PKS) works with healthcare organizations around the world to better capture, manage and leverage their human expertise to improve the performance of their business and deliver better patient outcomes.

PKS owns a patented, new generation clinical expert intelligence system – RippleDown – which integrates patient data with a knowledge base that is managed by clinical domain experts, to deliver patient-specific reports, recommendation and alerts.