Decision Support Tool Suite Version 2

Deliverable D7.3.2

Version Final

Authors: W.R. Van Hage[1], T. Ploeger[1], T. Bogaard[1], J.E. Hoeksema[1]
Affiliation: (1) SynerScope B.V.

Building structured event indexes of large volumes of financial and economic data for decision making
ICT 316404

Grant Agreement No.	316404
Project Acronym	NEWSREADER
Project Full Title	Building structured event indexes of large volumes of financial and economic data for decision making.
Funding Scheme	FP7-ICT-2011-8
Project Website	http://www.newsreader-project.eu/
Project Coordinator	Prof. dr. Piek T.J.M. Vossen VU University Amsterdam Tel. + 31 (0) 20 5986466 Fax. + 31 (0) 20 5986500 Email: piek.vossen@vu.nl
Document Number	Deliverable D7.3.2
Status & Version	Final
Contractual Date of Delivery	January 2015
Actual Date of Delivery	February 2, 2015
Type	Report
Security (distribution level)	Public
Number of Pages	28
WP Contributing to the Deliverable	WP7
WP Responsible	WP7
EC Project Officer	Susan Fraser
Authors: W.R. Van Hage[1], T. Ploeger[1], T. Bogaard[1], J.E. Hoeksema[1]	
Keywords: decision support, second version, knowledgestore exporter	

Abstract: In this deliverable we describe the second version of the Decision Support Tool Suite (DSTS), the main point of end-user interaction with the event indexes extracted from news in the NewsReader project. We shortly describe SynerScope, a visual analytics application that delivers real time interaction with dynamic network-centric data. We focus mainly on changes made to SynerScope since the delivery of D7.3.1.

We also describe the second version of the KnowledgeStore connector which receives and converts the NewsReader data for loading into SynerScope. We also explain how to interact hands-on with the current (at the time of writing) NewsReader data using SynerScope.

Table of Revisions

Version	Date	Description and reason	By	Affected sections
0.1	30 Sep 2014	Deliverable skeleton	T. Ploeger	All
0.2	24 Nov 2014	Outline of content for all sections	T. Bogaard	All
0.3	19 Jan 2015	Initial content for 2 sections	T. Ploeger	1, 2
0.4	22 Jan 2015	Initial content for 3 sections	T. Ploeger	3, 4, 5
0.5	28 Jan 2015	Expanded content for all sections, first complete draft	T. Ploeger, W.R. van Hage	All
0.5	28 Jan 2015	Revision, final version	T. Bogaard	All
0.5	30 Jan 2015	Check by coordinator	VUA	-

Executive Summary

The NewsReader project aims to support decision making by building structured event indexes of large volumes of news articles and financial data. The main point of interaction with the NewsReader event indexes will be the Decision Support Tool Suite (DSTS). The DSTS is a graphical user interface that supports users making strategic decisions. The DSTS does this by providing insight into the sequences of events leading up to a current situation such that a user can extrapolate to what might happen in the future.

The technical implementation of the DSTS largely consists of the SynerScope tool, a visual analytics application that delivers real time interaction with dynamic network-centric data. SynerScope interfaces with the NewsReader KnowledgeStore through a connector tool.

SynerScope allows the user to display several visualization methods ('views') simultaneously. Interactions with one view (e.g. highlighting, selecting) are replicated in all other open views. The Hierarchical Edge Bundling View (HEB) is the primary network visualization in SynerScope. Each node is visualized as a point on a circle, and each link is visualized as a curved line between its source and target node. The node are grouped hierarchically, based on a user-defined ordering.

The NewsReader event data can be imported into SynerScope through a connector tool. The exported data can then be projected onto nodes and links in at least three different ways. Once imported and projected, the data can be navigated through search, network expansion, and slicing and dicing to look for interesting patterns or irregularities.

This deliverable emphasizes changes made to the SynerScope tool since the release of D7.3.1.

Contents

List of Tables

1 Introduction

The NewsReader project aims to support decision making by building structured event indexes of large volumes of news articles and financial data. The purpose of this deliverable is to describe *the second version* of the Decision Support Tool Suite (DSTS), the main point of interaction with the NewsReader event indexes for the end user.

The *design* of the DSTS is described in detail in D 7.1 and the first version of the DSTS in D 7.3.1. In this deliverable we focus on the second version of the DSTS. We start by giving a short recap of the functionality of the DSTS (Section 2). Next, we dive into the changes and new features since D 7.3.1 in Section 3 and demonstrate the different ways the NewsReader data can be visualized and interacted with once loaded into the DSTS (Section 3.1). This is followed by an overview of the different project setups that can be used for analyzing the data in SynerScope, illustrated by short screen recordings showing the use of SynerScope. Subsequently we conclude the deliverable (Section 5) with an overview of future work.

The authors have tried to make this deliverable as self-contained as possible, but will at times refer the reader to other project deliverables for more detailed information. Even so, this deliverable should provide the reader with enough information to be able to understand the second version of the DSTS with the NewsReader event data.

2 The Decision Support Tool Suite

The Decision Support Tool Suite (DSTS) is a visual analytics application that provides insight into the sequences of events leading up to a current situation, allowing a user to extrapolate what might happen in the future. The system does not directly predict future events, but rather aids users in recognizing predictive patterns.

Within the NewsReader project the DSTS will interface with the KnowledgeStore (see D 6.2.1 for more details on the KnowledgeStore) and with external Web sources. The technical implementation of the DSTS largely consists of the SynerScope application. This section contains a quick recap of SynerScope's features, as described in D 7.3.1. For full details on the features of the first version of the DSTS we refer the reader to D 7.3.1.

As stated in Section 1, the *design* of the DSTS - and thus also the design of SynerScope - was described in Deliverable 7.1. For more details on the design principles that lie at the foundation of SynerScope, its visualizations, and its interaction methods, the authors refer the reader to that deliverable.

2.1 Information Schema

SynerScope is designed to work with a basic information schema consisting of the two core elements of networks: **nodes** and **links**. How individual parts of a data source are mapped onto these two concepts is up to the user. This allows a user to create several 'projections' of the same data, each with a different perspective fine-tuned to the task at hand.

Nodes are typically static entities such as persons, bank accounts, or IP addresses. Nodes are connected to each other via links, which represent dynamic elements (in that they often have a temporal attribute), such as e-mails, money transfers, or network traffic. Both nodes and links can have additional structured and unstructured data attributes.

2.2 Visualizations

SynerScope visualizes nodes and links in multiple dimensions. Each visualization is called a 'view'. We will give a short description of each individual view below.

2.2.1 Network View

In the Network View (also called the Hierarchical Edge Bundling View), each node is visualized as a point on a circle, and each link as a curved line between its source and target node. Nodes can be ordered hierarchically[1]. The links between nodes of the same hierarchical category are bundled together (as if tied together with a cable tie). The thickness of the bundle is the overlapping of the lines (each line is still shown individually) and shows the amount of interaction between categories.

[1] In the case of geographical data where nodes represent cities a valid hierarchy would be "continent - country - city".

2.2.2 Timeline View

The Timeline View (also called the Massive Sequence View) is essentially a re-oriented and temporally ordered variant of the Network View. Each node gets a fixed position on a vertical axis, grouped in the same hierarchy as in the network view. Links between nodes are represented by a horizontal line between the respective positions of the nodes. When certain categories of nodes interact with each other in a temporal pattern, this becomes instantly recognizable in the timeline.

2.2.3 Map View

The Map View provides a geographical view in SynerScope. If a node or link has latitude and longitude attributes, these can be used to plot the nodes (not the links) on a map as dots to show point density. An additional numerical attribute can be selected to assign point size.

2.2.4 Scatter Plot

The Scatter Plot View relates the values of two numerical attributes of either nodes or links, drawn as dots on a two-dimensional chart. As with the map view, a third attribute can be selected to assign the size of the dots.

2.2.5 Table View

The Table View shows the complete data as a table of values. There are separate sheets for nodes and links. Rows in these sheets correspond to individual links or nodes, the columns represent the attributes. A selection can be exported from this view to a spreadsheet file for later processing with other tools.

2.2.6 Other Views

In addition to the visualization views described above, SynerScope also has several 'supporting' views:

- **Search and Filter** The Search and Filter View is an interactive view that allows the user to select nodes or links by (a range of) values. The user can type in a search query that is matched against the value of a certain attribute for a given node or link type.
- **Hierarchy Editor** In the Hierarchy Editor View it is possible to define a hierarchical ordering (of arbitrary depth) on nodes. Each layer of the hierarchical ordering groups nodes together based on some common attribute value.
- **Web View** The Web View is used to display user-developed, specialized views created with HTML and JavaScript.

2.3 Interacting

SynerScope is able to display all of the above views simultaneously and coordinates user interaction between views. This allows users to interact with data in one dimension and see the results in other dimensions. This principle enables the user to explore correlations between different facets of the data. For instance, it can be used to explore whether actors that interact are geographically co-located by selecting part of a social network in a network view, then checking which corresponding locations are selected in a geographical view and whether these are near to each other.

The user can interact with SynerScope views in several ways: By **selecting** and **highlighting**, **drilling down** or **up**, and **expanding** selections.

When a user hovers the cursor over a certain node, link, or category/bundle of nodes or links, SynerScope highlights these objects with a blue color. Clicking turns that highlight into an orange colored selection.

Drilling down re-scopes the views to the current selection. Essentially, it allows the user to zoom in on an interesting part of the data set to get a more detailed view. Multiple drill downs can be executed sequentially to keep zooming in.

Drilling up performs the reverse operation. It re-scopes the domain of investigation to the previous drill level, preserving the current Selection. This allows the user to select small parts of the data that are of interest in a drilled down view, after which the user can drill up to observe it in a larger context.

Network expansion allows the user to extend the currently selected set of nodes or links with respectively links or nodes that are directly connected. This way it is possible to let the entire Selection grow by 'hopping' over the link network. Network connectivity, or the lack thereof, can be investigated in this way.

A number of fine-grained expansion tools are also available, such as an 'internal' expansion within the domain of the current selection. This expands the currently selected set of nodes and links with all the links between currently selected nodes that are not already selected.

3 New Features

Several new features have been added to the DSTS since the first release (D 7.3.1) at the end of Y1. These include three new ways to project the NewsReader data onto nodes and links, a new 'word square' visualization for unstructured data, new developments in cloud deployment, and other functionalities.

3.1 Data Projections

As described in Section 2, SynerScope visualizes the two core elements of networks: nodes and links. How individual parts of a data source are mapped onto these two concepts is up to the user. This allows a user to create several 'projections' of the same data, each with a different perspective adapted to the task at hand.

The previous version of the DSTS (D7.3.1) used a very basic projection of all RDF triples (subject, predicate, object) stored in the KnowledgeStore onto nodes and links. Each triple was a node - link - node combination, where the left node is the object, the link is the predicate, and the right node is the subject of the triple. Although this is functional, it is hard to understand and use without knowledge of the underlying RDF subject-predicate-object model.

For the second version of the DSTS, three new projections were developed. One designed for end users, showing only pure event data. The second is developed to bridge the gap between mention and event data, showing both mentions and events side-by-side. The third looks at the high-level ESO relationships between entities derived from event data. Figure 1 shows the three projections side-by-side in terms of which data layers they use and where they overlap. We will describe each projection in more detail below. We also explain how to work with the different projections in Section 4.

3.1.1 Pure Event Projection

The pure event projection uses concepts that are very close to the real world. On a very high level, the NewsReader data model describes Events that have participating instances of Actors, occur at a certain instance of a Place, and take place at a particular Time. Two actor instances that participate in the same event are essentially connected to each other by that event. Similarly, an actor instance that participated in an event that took place at at certain place instance is connected to that particular place instance by the event.

This high level description of the data model maps quite well onto nodes and links in a network, the conceptual data model in SynerScope. This means that in the pure event projection actors and places are nodes while the events associating these actors and places are links.

In this setup an event that involves only one actor will simply connect back (i.e. a reflexive link) to the same actor. An event that involves more than two actors and/or places will be represented by multiple links connecting all 'pairs' of actors/places. An event that involves the actors John, Jack, and Harry playing soccer will be represented by

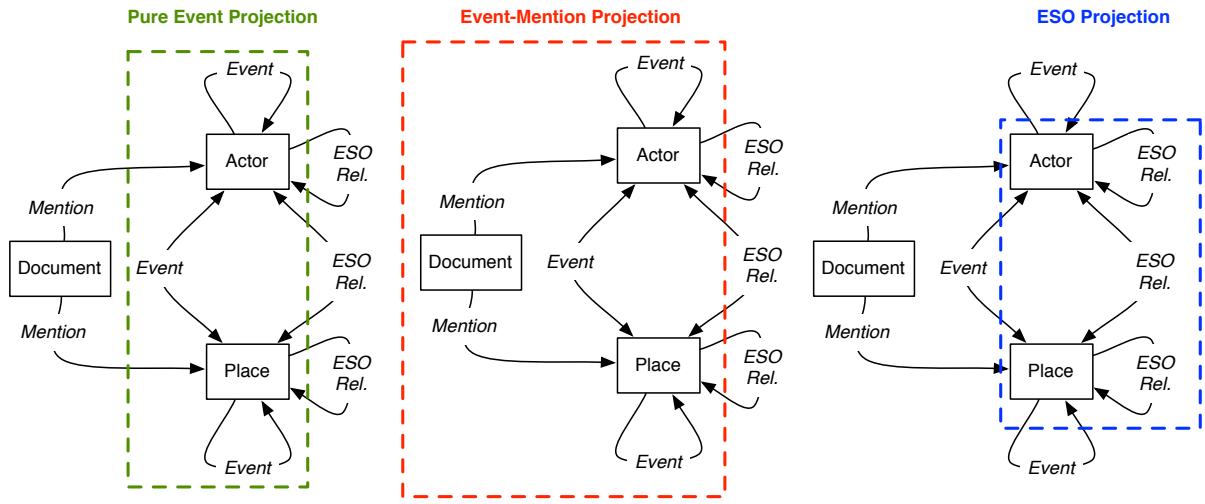

Figure 1: Figure showing the different data projections. Notice that they are all based on the same 'master schema' of elements, consisting of documents, actors, and places.

three links between the 3 actors: *John - Soccer - Jack, John - Soccer - Harry, and Jack - Soccer - Harry*. Because the links are not directed we do not need to include their inverse variants.

This projection allows users to explore who did what with whom, where, and when. Users can explore timelines of interactions between multiple actors, see the number of events that occurred in a certain place, find events of a specific type, and so forth.

3.1.2 Event-Mention Projection

The event-mention projection includes the pure event projection as described above, but also includes mention data. A mention (link) associates an actor or place (node) to a document (node) if they are mentioned in said document. With both the event and the mention data, it becomes possible to see the added value of the pure event data over the more traditional document-mention data. It also allows for quality inspection. It becomes possible to answer questions such as: Which events were derived from which mentions? Do they make sense given the content of the document? Which mentions do not have any events derived from them?

3.1.3 ESO Projection

The ESO projection does not show event or mention data itself, but rather shows information derived from these two data layers. It shows high-level relationships between actors and places, such as 'employed at' or 'at place'. These relationships were derived from the event data layer.

If there exists a 'hiring' event involving the actors Volkswagen and Jack on 2010-12-31, and a 'firing' event involving those same actors 5 years later, the linguistic pipeline will

derive an 'employed at' relation between Volkswagen and Jack with a begin time of 2010-12-31 and an end time of 2015-12-31. This relationship is projected onto a link, with the actors (or places in case of 'at place') as nodes.

This projection has primarily been developed to inspect the quality of the experimental ESO relationship data. We will use the DSTS to explore the derived relationships and see if they make sense. The hierarchical grouping of actors in the DSTS will allow us to immediately spot non-sensical relationships, such as inanimate objects being part of an employment relationship.

3.2 Word Square

The NewsReader event data has several attributes that are represented as unstructured text (e.g. event labels or parts of the original news articles) or categorical attributes (e.g. FactBank values, semantic classes, or syntactic classes). The word square visualization is a completely new type of interactive visualization, invented and developed by SynerScope within the NewsReader project, that aims at showing which textual values are the most (or the *least*, which is sometimes more informative than the *most*) representative of a set of selected data. The word square is an adaptation of the well known word cloud. As opposed to one set of terms, the word square can show two sets of terms that it contrasts to each other. One set, the selection, is depicted in a striking color (orange by default), while the rest of the data is shown in gray. Data values that occur relatively more often in the selection than the rest of the data are depicted larger, while relatively less frequent values are depicted smaller. The word square is designed to: 1) clearly show which terms are most or least representative of the selection, and how the importance of terms ranks relative to each other, 2) optimally use the available screen space, 3) allow the easiest possible interaction (clear labels that work like big buttons). An example is shown in Figure 2. Clicking on terms in the word square will select all data points that contain that term in the displayed attribute. This sets a new selection, and thus changes the word square. By clicking in sequence, the user can navigate the entire data set and so explore correlations between terms.

3.3 Cloud Deployment

SynerScope can run in a "private" cloud environment, as a graphically accelerated virtualized application, in the shape of a Dell rack mountable server with preinstalled software, including a Citrix virtualization layer. This allows users to user the Decision Support Tool Suite from their own computer regardless of the technical specification, even a tablet computer is sufficient, as long as the network access to the cloud server is sufficiently stable.

A setup in a public cloud environment is also possible. Specifically, we have had some success in running SynerScope on graphically accelerated virtual servers in a number of settings:

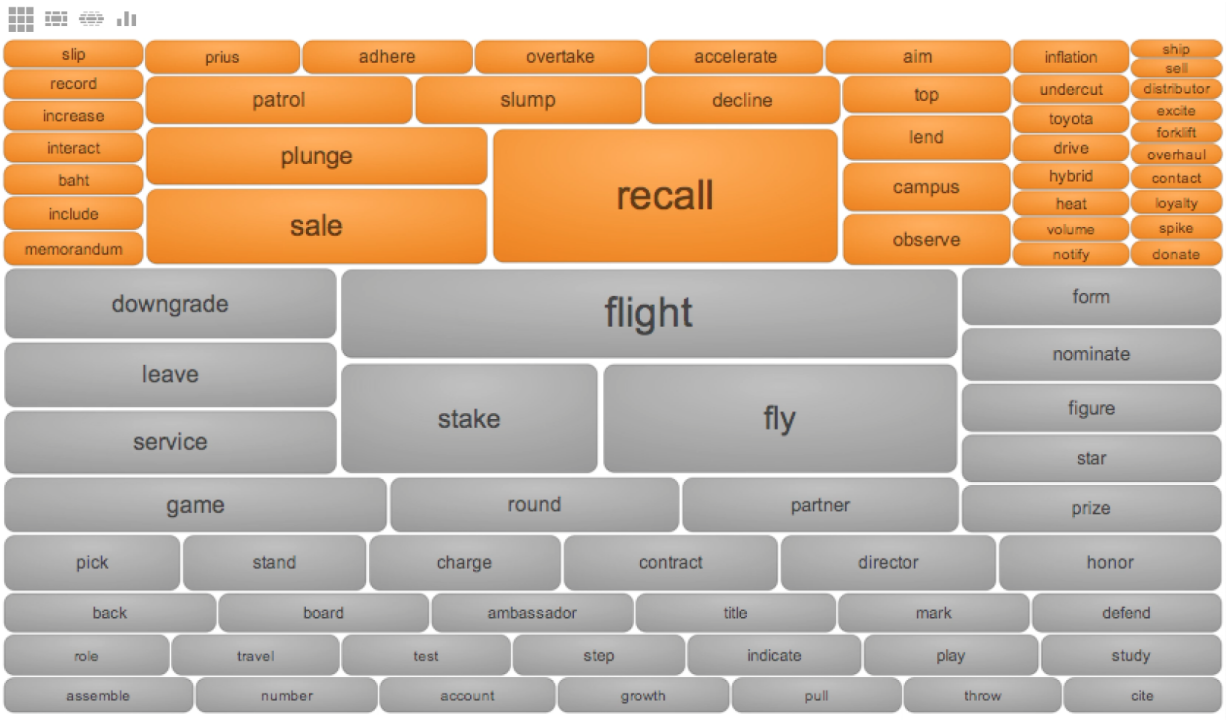

Figure 2: This word square visualization of event types shows which terms are most representative of the types of events that are currently selected (in orange) and most representative of the ones that are currently not selected (in gray). For instance, this selection is of events including Toyota and the word square shows that as compared to the rest of the data, Toyota is discussed more in the context of a "plunge" in "sales", and "recall" of products. The gray terms show that events having to do with Toyota are relatively unlikely to have anything to do with "flight" or "fly"[ing] or "downgrade" of something.

- With Amazon Web Services[2] (AWS),
  - In collaboration with OTOY[3]
  - In collaboration with Scalable Graphics[4]
- With Cloudalize[5]
- With MainFrame2[6]

Currently, each of these settings are sufficient for non-mission critical simple proof of concept cases. Depending on the location of the server parks and the latency and consistency of internet connectivity, the settings are sufficiently fast and interactive. However, data transfer rates and costs vary, and some settings require the installation of custom clients on the user's computer. In many organizations, specific users will not be able or allowed to install such a client. We are currently investigating a new setup together with Amazon Web Services that should solve network consistency and data transfer costs issues and should be able to run from a browser without the need to install a custom client application.

3.4 Navigating DSTS Visualizations

Based on the recommendations from the Y1 review, we have added more options to navigate through the visualizations in the DSTS. In SynerScope, it is possible to drill-down into a subset of the visualized data to view it more clearly, isolated from the rest of the data. These drill-down actions can be repeated as often as necessary, for example first drilling down to a particular time range and then to a certain geographical area. In this way you can 'zoom in' to an interesting subset of the visualized data step-by-step.

Zooming out of a subset of data ('drilling up') incrementally was not possible, it was only possible to go back up in one fell swoop to all data. As recommended by the Y1 reviewers, it is now possible drill up incrementally. This means it is now possible to go back step-by-step and retrace your drill-down actions, instead of losing them by having to drill back up to all data at once.

3.5 Saving Visualization Templates

Based on the recommendations from the Y1 review, we have added an option to save the visualization setup of the DSTS while refreshing the data that is shown. It is now possible to save the layout and settings of all visualizations in a SynerScope project to a separate file. This file can then be set as the default settings for a certain project. This settings file can also be shared with other users of the DSTS to exchange useful layouts and settings.

[2] AWS, http://aws.amazon.com
[3] OTOY Octane Cloud Workstation, https://aws.amazon.com/marketplace/pp/B00FGB1VTC
[4] Scalable Graphics, http://www.scalablegraphics.com
[5] Cloudalize, http://www.cloudalize.com
[6] MainFrame2, https://www.mainframe2.com

When using the saved settings, data can be refreshed in the background without breaking the layout and settings of the visualizations.

3.6 Other Features

Developing Training Material Based on the recommendations from the Y1 review, we have started developing training material for end users of the DSTS. As the DSTS itself is not designed to predict the future, but rather help users make their own predictions, users may need to be trained in operating and interpreting the visualizations by the DSTS to help them make their own predictions.

The training material starts with an overview of why visualizing data is important, and what SynerScope brings to the table in this area. Next, the material covers basic and advanced interaction patterns for use in the DSTS. The material concludes by describing common visual patterns that can occur for specific types of data.

Range Select In addition to the existing selection options (simple click, slice, dragging a bounding box), users can now use the Search & Filter view to search for a specific range of time or range of a certain numerical attribute (such as stock price or IP amount.

Collaboration Server Besides the new visualizations, we have developed an experimental collaboration method for the DSTS. Users of the DSTS can now save selections to a server using a combination of MongoDB[7] and Meteor.js[8]. This functionality allows users to save selections across investigations, as well as send your selection to another user somewhere else, for example for verification or further study.

[7] `http://www.mongodb.org/`
[8] `https://www.meteor.com/`

4 Working with the DSTS

We have designed ready to use SynerScope project files for each of the data projections described in Section 3.1. In this section we will step through each of these projects from the perspective of an end user, on the basis of example tasks that a user would perform.

Each task is accompanied by a short screen recording introducing the visualization setup and showing exactly how to perform the task step-by-step. We believe this is a more efficient and effective way of showing how to operate the DSTS than with a large number of static screenshots. The videos can be found at:

http://www.newsreader-project.eu/results/demos/synerscope/

Note that we do not attempt to give an exhaustive overview, but rather try to give a good impression of the possibilities.

4.1 Pure Event Projection

Events for a particular actor

1. Locate and select the actor you are interested in on the innermost circle in the network view. The type of the actors (according to the KnowledgeStore) is located on the outer circles of the network view. For example, all car models are within the 'MeansOfTransporation' segment. Alternatively, search for a specific actor by name using the search and filter view.

2. Once the actor is selected, press the 'expand nodes to links' button to select all events attached to the currently selected actor.

3. Observe that the events are now selected in orange in all open visualizations. See which other actors they connect to in the network view, when they occurred in the timeline view, where they occurred in the map view, and what type of event it was in the word square.

4. Optionally, you can now drill down into the selected set of events and repeat the process for another actor to refine your selection.

Events involving 2 (or more) specific actors

1. Locate and select the actors you are interested in on the innermost circle in the network view. Alternatively, search for specific actors by name using the search and filter view.

2. It is possible to select multiple actors by holding the CTRL key on you keyboard. This enable the additive selection mode.

3. Once all actors are selected, press the 'expand nodes to internal links' button to select all events attached to the currently selected actor. This is located in the submenu of the 'expand nodes to links button'.

4. Observe that the events involving all previously selected actors are now selected in orange in all open visualizations. See when they occurred in the timeline view, where they occurred in the map view, and what type of event it was in the word square.

5. Optionally, you can now drill down into the selected set of events and repeat the process for another actor to refine your selection.

Actors involved in events of a particular type

1. Locate and select the events you are interested in on the word square view. Alternatively, search for specific events by label using the search and filter view.

2. It is possible to select multiple labels by holding the CTRL key on you keyboard. This enable the additive selection mode.

3. Once all event labels are selected, press the 'expand links to nodes' button to select all actors and places attached to the currently selected events.

4. Observe that the events and actors are now selected in orange in all open visualizations. See which actors they connect to in the network view, when they occurred in the timeline view, and where they occurred in the map view.

5. Optionally, you can now drill down into the selected set of events and actors and repeat the process for another event label to refine your selection.

Events taking place on a particular location

1. Locate and select the place you are interested in on the innermost circle in the network view, or on the map view.

2. Once the place is selected, press the 'expand nodes to links' button to select all events attached to the currently selected place.

3. Observe that the events are now selected in orange in all open visualizations. See which actors and other places they connect to in the network view, when they occurred in the timeline view, where they occurred in the map view, and what type of event it was in the word square.

4. Optionally, you can now drill down into the selected set of events and repeat the process for another place and/or actor to refine your selection.

Events taking place at a certain point in time

1. Locate and select the events in the time period you are interested in by dragging a selection window from top to bottom in the timeline view. Tooltips on the borders of the selection window will indicate the length of the selected time period.

2. Once the events in the target period are selected, press the 'expand links to nodes' button to select all actors and places attached to the currently selected events.

3. Observe that the events and actors are now selected in orange in all open visualizations. See which actors they connect to in the network view and where they occurred in the map view.

4. Optionally, you can now drill down into the selected set of events and repeat the process for a more specific time period to refine your selection.

The interaction examples described above can be "mixed and matched" as desired, for example first filtering by location, then by time, and finally by event type. Drilling down after each step continuously narrows the scope of the visualizations to only show the parts that are of interest to the user.

4.2 Event-Mention Projection

As described in Section 3.1, the Event-Mention Projection is a proper superset of the Pure Event projection, meaning that all of the example tasks given for the Pure Event projection can also be performed in the Event-Mention projection. What follows are example tasks that can only be executed in the Event-Mention projection.

Comparing mentions and events for a certain actor In this task we will compare all mentions for a certain actor and see what events were derived from those mentions. Do they make sense? Are we missing events for certain mentions?

1. Locate and select the actor you are interested in on the innermost circle in the mention network view. Alternatively, search for a specific actor by name using the search and filter view.

2. Once the actor is selected, press the 'expand nodes to links' button to select all events *and* mentions attached to the currently selected actor.

3. Observe that the mentions and events are now selected in orange in all open visualizations. See which other actors and documents they connect to in the network views, when they occurred in the timeline view, where they occurred in the map view, and what type of event it was in the word square.

4. Pay specific attention to the types of events that were derived and which actors and/or places they link together. Do they make sense when you read the original mentions of the actors?

Comparing co-mentions and events In this task we will look at actors and/or places that were mentioned together in one document and see if the event layer contains events that make sense given the content of those documents.

1. Locate and select the actors you are interested in on the innermost circle in the mention network view. Alternatively, search for specific actors by name using the search and filter view. Select one of the actors and keep the other in mind.

2. Use the 'expand by name' button to select all mentions attached to the currently selected actor. Then use the 'expand by name' again, this time to select all documents attached to the mentions you just selected.

3. Now use the additive selection mode to add the second actor (in the mention network view) to your selection. Then use the 'expand to internal links' button to select all of the documents also mentioning the actor you just selected.

4. Drill down to the current selection. Notice that the event network view has been reduced to just one line (as we only have two actors left).

5. Browse through the mentions and documents and see if the events that were derived for these documents make sense. Pay specific attention if the document actually describes what the two actors were doing together, if anything.

4.3 ESO Projection

Evaluating entities related by ESO triples In this task we will take a look at a specific ESO relationship and see if the entities it relates make sense in the context of that relationship. For example, does it make sense that a specific model of car is employed by a person?

1. Locate and select the actor and/or place you are interested in on the innermost circle in the network view. The type of the actors (according to the KnowledgeStore) is located on the outer circles of the network view. For example, all car models are within the 'MeansOfTransporation' segment. Alternatively, search for a specific actor by name using the search and filter view.

2. Once the actor/place is selected, press the 'expand nodes to links' button to select all ESO relationships attached to the currently selected actor/place.

3. Observe that the ESO relationships are now selected in orange in all open visualizations. See which other actors/places they connect to in the network view, when these relationships began and when they ended in the timeline views, where they occurred in the map view, and what type of relationship it was in the word square.

4. Optionally, you can now drill down into the selected set of relations and repeat the process for another actor to fine tune your selection.

5. Pay specific attention to the types of relationships that were derived and which actors and/or places they link together. Do they make sense when you read look at the type of actor/place they relate? Are the begin and end timestamps reasonable? Do they overlap with another relationship in the same timeframe?

5 Conclusion

In this deliverable we documented the design of the Decision Support Tool Suite, a graphical user interface that is meant to support users when making strategic decisions. The DSTS provides insight into the sequences of events that led up to a current situation so that a user can extrapolate to what might happen in the future.

The DSTS will be the main point of end user interaction with the NewsReader event indexes. The technical implementation of the DSTS largely consists of the SynerScope tool, a visual analytics application that delivers real time interaction with dynamic network-centric data.

SynerScope allows the user to display several visualization methods ('views') simultaneously. Interactions with one view (e.g. highlighting, selecting) are replicated in all other open views. The Hierarchical Edge Bundling View (HEB) is the primary network visualization in SynerScope. Each Node is visualized as a point on a circle, and each Link is visualized as a curved line between its source and target Node. The Nodes are grouped hierarchically, based on an ordering defined by the user.

The NewsReader event data can be imported into SynerScope when downloaded from the KnowledgeStore. Once imported, the data can be navigated through search, network expansion, and slicing and dicing to look for interesting patterns or irregularities.

In this second revision of the DSTS, the following new features were implemented:

- Three new data projections that are more intuitive for end users and can help with evaluating the quality of the NewsReader data;
- Word square visualization for unstructured data;
- Effort has been made towards hosting the DSTS in the cloud;
- Navigating DSTS Visualizations has been made easier;
- Saving Visualization Templates while refreshing data;
- Range select feature for searching and filtering;
- Collaboration options, users can share results with each other and work together.

5.1 Future Work

We have drafted the following plan for the next year, which is subject to change pending the recommendations of this year's review.

- Completion of Deliverable D7.3.3, the third version of the Decision Support Tool Suite.
- Continue with developing support and training material for end users of the DSTS.

- Development of a graphical interface situated between WP06's KS and WP07's DSTS. This interface should facilitate intuitive importing and refreshing of subsets of the knowledge from the KnowledgeStore into the DSTS. We expect to be able to start full development of this interface now that the schema for the NewsReader data is stabilizing and the usage scenario's for end users are becoming more clear. As per the recommendation giving during the Y1 review, the interface will output pre-canned SynerScope projects that do not need to go through SynerScope's importer but can be opened directly. The interface will allow users to specify what they are interested in, in terms of event types, specific actors, specific places, certain timeframes, and so forth. The interface will also contain a mechanism to keep projects up-to-date with new information that becomes available as new news articles are processed.

References