Ref. Ares(2017)3839883 - 31/07/2017

D5.7 – Robot Semantic Sentiment Analysis

Project Acronym:	**MARIO**
Project Title:	**Managing active and healthy aging with use of caring service robots**
Project Number:	**643808**
Call:	**H2020-PHC-2014-single-stage**
Topic:	**PHC-19-2014**
Type of Action:	**RIA**

D5.7 – Robot Semantic Sentiment Analysis_Final version	
Work Package	WP5
Due Date	M30
Submission Date	31/07/2017
Start Date of the Project	01/02/2015
Duration of the Project	36 Months
Organisation Responsible of Deliverable	CNR
Version	2.1
Status	Final
Author name(s)	Valentina Presutti (CNR) Aldo Gangemi (CNR) Diego Reforgiato Recupero (R2M) Domenico Pisanelli (CNR) Paolo Ciancarini (CNR) Andrea Giovanni Nuzzolese (CNR) Luigi Asprino (CNR) Alessandro Russo (CNR)
Reviewer(s)	Christos Kouroupetroglou (CNET), André Freitas (R2M)
Nature	☐ R – Report ☐ P – Prototype ☐ D – Demonstrator ☒ O – Other
Dissemination Level	☒ PU – Public ☐ CO – Confidential, only for members of the Consortium (including the Commission) ☐ RE – Restricted to a group specified by the Consortium (including the Commission Services)

Revision History			
Version	**Date**	**Modified by**	**Comments**
0.1	16/06/2017	Valentina Presutti (CNR)	First draft starting from D5.3
0.2	19/06/2017	Alessandro Russo (CNR)	Revision and changes on the document structure
0.3	22/06/2017	Alessandro Russo (CNR)	Added content on polarity detection
0.4	26/06/2017	Andrea Giovanni Nuzzolese (CNR)	Revision and extension of semantic sentiment analysis content
0.4.1	27/06/2017	Andrea Giovanni Nuzzolese (CNR)	Unified notation for figures with knowledge graphs
0.5	05/07/2017	Alessandro Russo (CNR)	Initial content on use case scenario
0.5.1	06/07/2017	Alessandro Russo (CNR)	Content and figures for My Memories scenario
0.5.2	10/07/2017	Andrea Nuzzolese (CNR)	Added frame-based example in use case scenario
0.6	13/07/2017	Alessandro Russo (CNR)	Revision of document outline and overview sections
1.0	16/07/2017	Alessandro Russo (CNR)	Overall revision and minor edits to improve readability; draft version for internal review

continues

Revision History (cont.)			
Version	**Date**	**Modified by**	**Comments**
1.0	19/07/2017	Adam Santorelli (NUIG)	Comments and feedback on v1.0
1.0	21/07/2017	Christos Kouroupetroglou (CNET) André Freitas (PASSAU)	Reviewer's feedback on v1.0
1.1	25/07/2017	Alessandro Russo (CNR)	Revision and updates according to reviewer's feedback
2.0	27/07/2017	Alessandro Russo, Luigi Asprino, Andrea Giovanni Nuzzolese (CNR)	Revision of the overall document wrt reviewers' comments; added abbreviations list; minor formatting edits; version for coordinator
2.0	28/07/2017	Aisling Dolan (NUIG) Alexandros Gkiokas (ORTELIO)	Comments and feedback on v2.0
2.1	28/07/2017	Alessandro Russo (CNR)	Revision with additional edits; final version

Executive Summary

Deliverable 5.7 describes the software components that were designed and developed in the context of Task 5.3 (WP5) to provide the MARIO system architecture and software framework with Sentiment Analysis capabilities. This version represents the final, consolidated version of Deliverable 5.3.

Specifically, the components presented here constitute MARIO's Sentiment Analysis subsystem. The main component aims at providing **semantic sentiment analysis** capabilities, through a formal representation and evaluation of the sentiment expressed in text sentences. It is built as an extension of FRED (a machine reading component introduced in Deliverable 5.6) and relies on novel resources developed in the context of this task.

Table of Contents

Acronyms and Abbreviations

API	Application Programming Interface
CGA	Comprehensive Geriatric Assessment
JSON	JavaScript Object Notation
KB	Knowledge Base
MPI	Multidimentional Prognostic Index
MON	MARIO Ontology Network
NLP	Natural Language Processing
NLU	Natural Language Understanding
OWL	Web Ontology Language
PWD	Person/People with Dementia
RDF	Resource Description Framework
REST	Representational state transfer
S2T	Speech to Text

Introduction

Natural language is the primary means through which people with dementia (PWD) can interact with MARIO robots. Task 5.3 deals with MARIO's capability to extract sentiment information from natural language sentences expressed by a person with dementia (PWD). The ability to automatically extract and categorise sentiment information from the textual representation of natural language sentences enables MARIO to adjust and adapt its behaviour according to the sentiment or opinion potentially expressed by PWD.

This deliverable presents the approaches, algorithms and software components that enable MARIO to perform sentiment analysis over the textual representation of spoken natural language input. In particular, MARIO is equipped with semantic sentiment analysis capabilities, which exploit a frame-base formal representation of the textual input by augmenting the Natural Language Processing (NLP) capabilities discussed in Deliverable 5.6.

Work Package 5 Objectives

WP5 aims at developing the framework and tools that allow MARIO robots to interact with humans and understand their needs expressed through spoken natural language. As fundamental building blocks, understanding capabilities exploit machine reading/listening components and RDF/OWL ontologies to first produce and then process a formal encoding of the textual representation of natural language.

As such, the main objectives of WP5 are:

- to design and develop the Mario Ontology Network (MON) and Knowledge Base;
- to provide MARIO with the ability of transforming natural language into a formal representation, to enable reading and listening capabilities on the basis of FRED;
- to provide MARIO with the capability of recognising, storing and reusing sentiment information, on the basis of semantic sentiment analysis techniques.

Purpose and Target Group of the Deliverable

This deliverable aims at describing the software components designed and developed in Task 5.3.

These components provide the robot with sentiment analysis capabilities complementing the ability to process and understand spoken natural language. The role of sentiment analysis in human-robot interaction, with a focus on PWD, is considered.

Due to its technical nature, the deliverable is mainly targeting researches, practitioners and developers interested in sentiment analysis techniques and algorithms, and in particular semantic frame-based techniques and their application for service companion robots. In addition to the technical aspects, concrete use cases are considered, with the aim of providing health experts with an understanding on how these techniques can improve the interaction between PWD and companion robots.

Relations to other Activities in the Project

This deliverable directly relies on the results of Task 5.1 (see D5.1) as far as the background knowledge and knowledge models that it uses are concerned (as part of the MON – MARIO Ontology Network), and it is strongly related to the activities carried out in Task 5.2 concerning Natural Language Understanding (see D5.6).

The overall WP5 receives as input the user and functional requirements and the system architecture from WP1, while WP2 provides the Kompai robot and platform where the software components are deployed. These components provide sentiment analysis services to the applications and modules developed in WPs 3, 4 and 6 with specific focus on capturing and representing sentiment knowledge, in line with the integration procedures defined in WP7.

Validation activities in WP8 provide feedback to the iterative design and development process of the software components, and contribute to their evolution and refinement.

Document Outline

The rest of this deliverable is organised in three main sections. Specifically:

- the first section provides an overview of sentiment analysis and its role in supporting the interaction with PWD;
- the second section provides a description of MARIO's Sentiment Analysis subsystem, with a focus on the models, resources and algorithms that support and enable semantic sentiment analysis capabilities;
- the third section presents a representative use case scenarios for the Sentiment Analysis subsystem, by outlining how its capabilities are used in the context of the My Memories application for reminiscence.

About MARIO

MARIO addresses the difficult challenges of loneliness, isolation and dementia in older persons through innovative and multi-faceted inventions delivered by service robots. The effects of these conditions are severe and life-limiting. They burden individuals and societal support systems. Human intervention is costly but the severity can be prevented and/or mitigated by simple changes in self-perception and brain stimulation mediated by robots.

From this unique combination, clear advances are made in the use of semantic data analytics, personal interaction, and unique applications tailored to better connect older persons to their care providers, community, own social circle and also to their personal interests. Each objective is developed with a focus on loneliness, isolation and dementia. The impact centres on deep progress toward EU scientific and market leadership in service robots and a user driven solution for this major societal challenge. The competitive advantage is the ability to treat tough challenges appropriately. In addition, a clear path has been developed on how to bring MARIO solutions to the end users through market deployment.

Sentiment Analysis

Sentiment and Emotions when Interacting with PWD

The role and importance of recognising the *sentiment* or *emotional state* of PWD when interacting with them through a conversational approach is largely recognised.

> *Identify the emotional state of the response*
>
> *How is this person feeling? If they have been able to speak, what do the words convey?*
>
> http://www.scie.org.uk/dementia/after-diagnosis/communication/conversation.asp

> *Sometimes the emotions being expressed are more important than what is being said.*
>
> http://www.alz.org/care/dementia-communication-tips.asp

Sentiment Analysis – Overview (1/2)

Sentiment analysis concerns the study of intelligent algorithms capable of automatically mining (i.e., identifying and categorising) opinions from natural language content.

In the context of MARIO, we are interested in applying sentiment analysis to sentences in natural language spoken by PWD, especially when they are expected to express feedback about some recent activity or to converse about events, people and places that characterise their life history and past experiences.

This information can then be stored and used in order to contribute to a personalised user experience and influence MARIO's decision making when interacting with the users.

Sentiment Analysis – Overview (2/3)

The initial working and research hypothesis was based on the assumption that sentiment analysis capabilities could be mainly used to support the Comprehensive Geriatric Assessment (CGA) and Multidimensional Prognostic Index (MPI) robotic modules (investigated in WP4).

As detailed in Deliverables 4.3 and 5.6, the human-robot interaction process required to support CGA has to be based on a conversational approach driven by standardised clinical questionnaires.

Since the initial validation activities and on-the-field observations related to the development of the CGA module, it emerged that PWD's replies to the questions defined in the assessment questionnaires *do not carry sentiment information* (most of the questions assume a *yes/no* answer or are based on a multiple-choice structure), preventing sentiment analysis techniques to be effectively used and contribute to the MPI.

Sentiment Analysis – Overview (3/3)

As detailed later in this deliverable, an immediate application of sentiment analysis capabilities is to the *My Memories* app for reminiscence (presented in D3.3), which enables MARIO to interact with PWD to stimulate conversation and elicit memories, by prompting them and asking them to express their mood/feelings/opinion (e.g., about personal events or people) after looking at a picture with few words or a sentence, etc.

This information expressed through speech is captured by the *Understanding Subsystem* and processed by the *Sentiment Analysis* module, so as to produce data expressing sentiment and emotional information associated to specific entities or events.

At a later stage MARIO can use this information in order to decide whether to, e.g., show a picture again or talk about a specific person or past event, according to specific goals and conversational strategies it is implementing.

MARIO Sentiment Analysis Subsystem

Components, services and capabilities

Architectural Reference

Referring to MARIO architecture, this task contributes to the *Natural Language Understanding* subsystem (see D5.6), by introducing **sentiment analysis** modules and services.

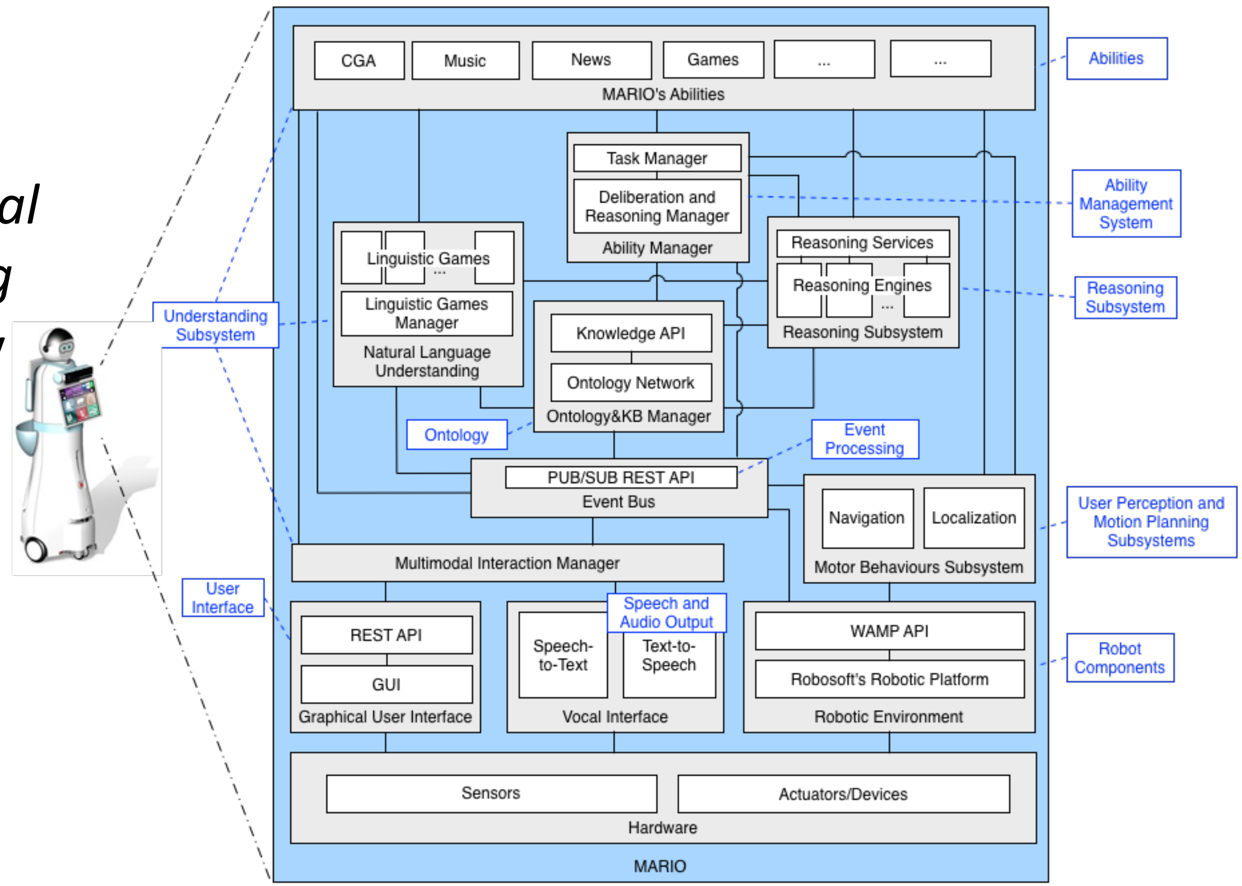

Approach and Contributions

In order to achieve the goal of this task and provide MARIO with multilingual sentiment analysis capabilities, we:

- adopted a **modular** and **service-oriented design** approach, to support multiple approaches having different capabilities;
- reused an **ontology for representing opinions**;
- integrated **sentiment lexical resources** into MARIO's background knowledge (relying on the Framester resource [13, 14], described in Deliverable 5.1);
- developed a **sentence-polarity** evaluation module;
- extended and implemented a **semantic frame-based sentiment analysis** algorithm.

Sentiment Analysis Components

The Sentiment Analysis Subsystem complements (and relies on) the capabilities provided by the Natural Language Understanding (NLU) Subsystem (described in D5.6).

It implements two main strategies, detailed in the next slides, for computing *sentiment* and *emotional scores*:

1. simple **sentiment polarity analysis**
   – it operates at a *sentence level*, to identify the overall *polarity* (Negative-Neutral-Positive) of a given sentence;
2. advanced **semantic sentiment analysis**
   – it operates on a *semantic frame-based representation* of a sentence, to identify the sentiment expressed by an *opinion holder* on a certain *entity* or *topic*;
   – it relies on the MARIO Knowledge Base (accessed via the REST API provided by Lizard; see Deliverable 5.1) and other resources already exploited by the NLU subsystem, with the addition of sentiment lexical resources;
   – this capability is provided by a software component accessible through a REST service[1].

[1] http://wit.istc.cnr.it/stlab-tools/sentilo/service

Components Overview

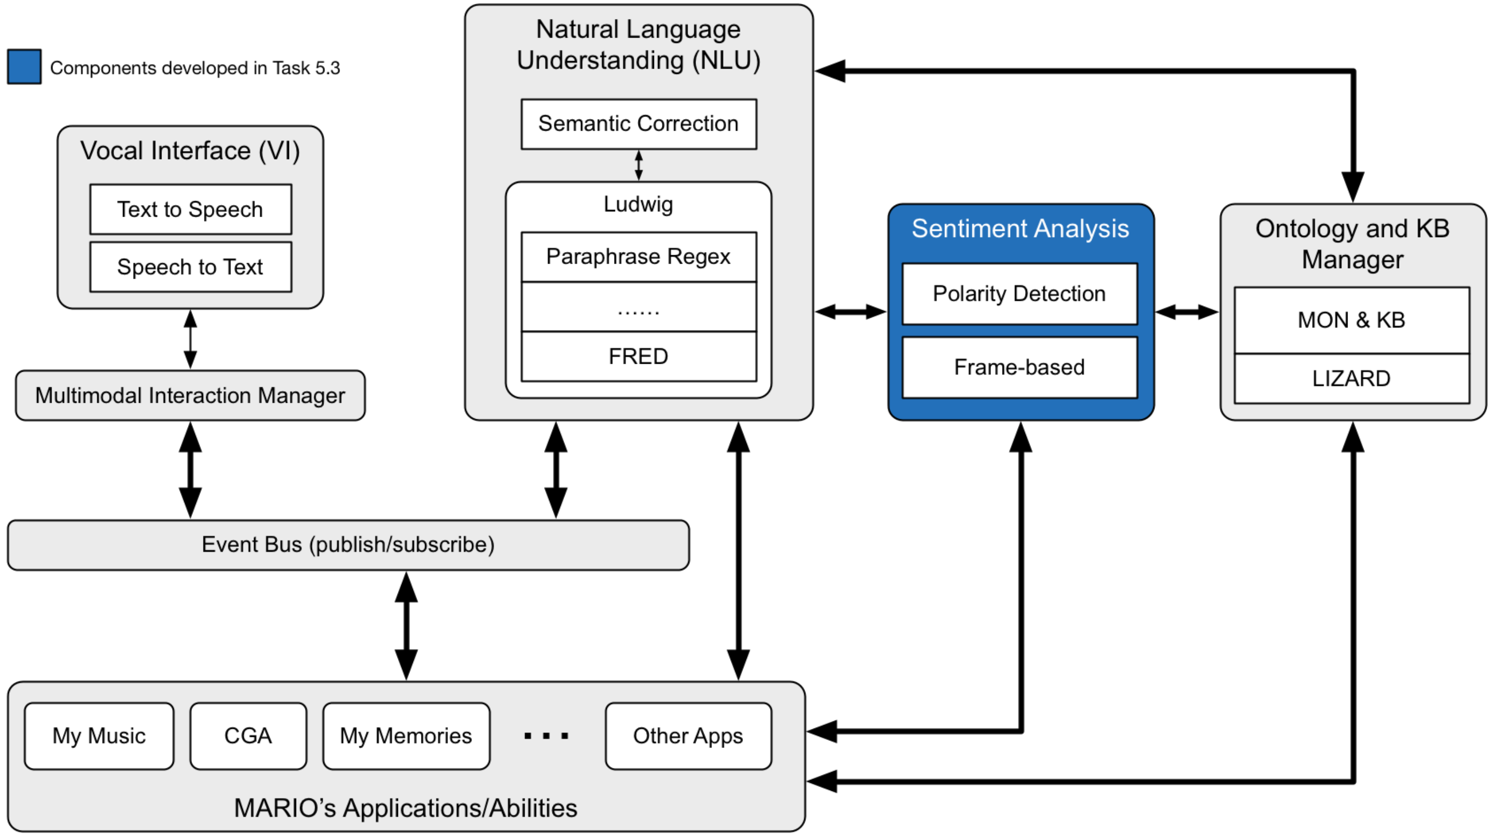

- This figure highlights the main dependencies and the flow among the components developed in this task and other components in the MARIO architecture.
- The Sentiment Analysis subsystem directly interacts with MARIO Knowledge Base via the REST API provided by *Lizard* (see Deliverable 5.1) for retrieving and storing data.
- The other components are developed in other project's tasks and WPs (T5.1, T5.2, WP3, WP4 and WP6).

Polarity Detection

Sentence-based Polarity Detection

Aims at identifying the *overall tonality/sentiment* (reflecting speaker's attitude) expressed in a *sentence*.

- It classifies a given sentence/text according to a *polarity scale* (typically, Negative-Neutral-Positive), potentially with a *ranking score* or *confidence value.*
- It maps to a classification problem, addressed with Machine Learning techniques;
  - e.g., Naïve Bayes, Maximum Entropy (MaxEnt), Support Vector Machines (SVM) or Deep Neural Network classifiers, exploiting syntactic features extracted from the sentences.
- The polarity classification step is preceded by Natural Language Processing (NLP) tasks:
  - basic NLP pipeline includes tokenization, sentence splitting and syntactic parsing;
  - part-of-speech (POS) tagging and lemmatisation steps can be added to the pipeline.

Polarity Detection Module

MARIO's **polarity detection module** relies on the Sentiment Analysis capabilities of the Stanford CoreNLP framework [1], which:

- uses a deep learning approach with Recursive Neural Networks, trained on the Stanford Sentiment Treebank dataset [2];
- is based on an underlying model that aims at capturing the effects of contrastive conjunctions as well as negation;
- classifies and scores input sentences according to *5 classes of sentiment*: very negative (-2), negative (-1), neutral (0), positive (1), and very positive (2);
- relies on a NLP pre-processing pipeline with text tokenization, sentence splitting and syntactic parsing.

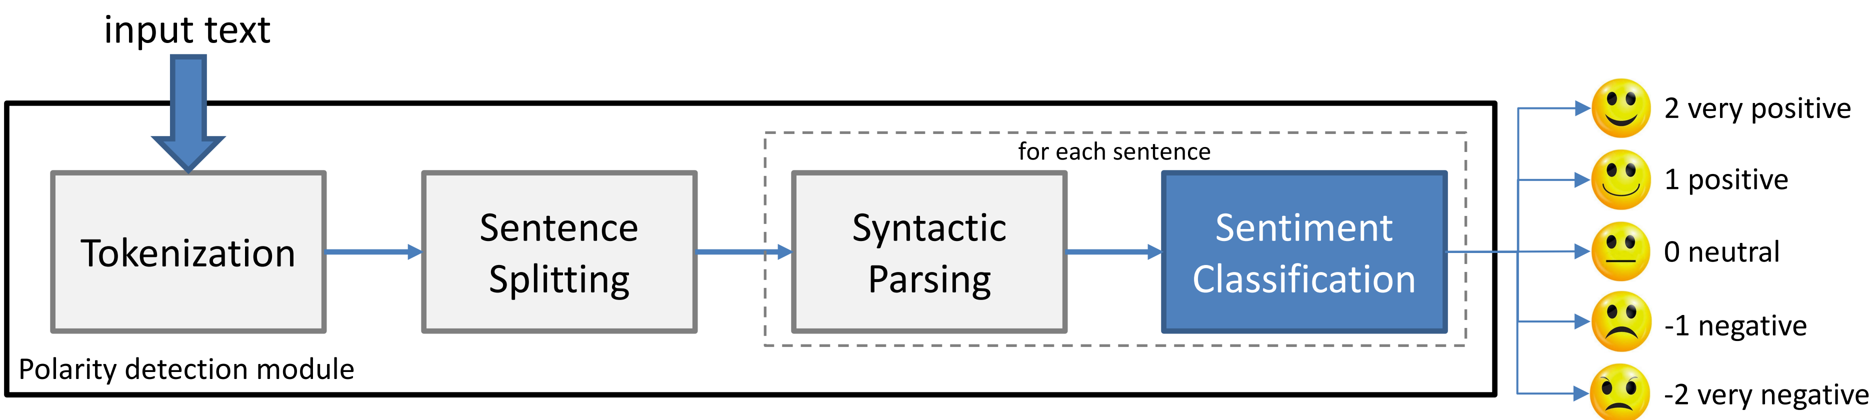

Polarity Detection

Pros:

+ pre-trained algorithms and models are available for the different steps of the processing pipeline;
+ good for sentences expressing a single, general opinion or sentiment;
+ simple to use and interpret the analysis results.

Limitations:

- unable to clearly distinguish multiple opinions expressed in a single sentence over different entities/topics
  - e.g., the sentence *"The food was great, but the weather was so bad!"* is classified as a whole as Negative, with no distinction about the different opinions expressed about the food and the weather
- does not exploit semantic features of the input.

Beyond polarity detection: ability to identify the *sentiment* expressed by an *opinion holder* on a certain *entity* or *topic*.

Semantic Sentiment Analysis

What's an opinion

An **opinion** can be defined as an intentional statement by somebody (*holder*) on some fact (*topic*) that is expressed with a possible *sentiment*.

A Sentiment Analysis system should be able to extract and characterise *opinions* by recognising the *attitude* (positive, negative or objective) of an *opinion holder* on a certain *topic*, or by evaluating the overall *tonality* of a sentence.

Formal definition

The goal of Sentiment Analysis is to detect quintuples *(e_j, a_{jk}, so_{ijkl}, h_i, t_l)* from unstructured text, where an opinion is a quintuple [3,4]:

$$(e_j,\ a_{jk},\ so_{ijkl},\ h_i,\ t_l)$$

where:

- e_j is a *target entity*;
- a_{jk} is an *aspect/feature* of the entity e_j;
- so_{ijkl} is the *sentiment value* of the opinion from opinion holder h_i on aspect a_{jk} of entity e_j at time t_l. so_{ijkl} is positive, negative or neutral, or a rating;
- h_i is an *opinion holder*;
- t_l is the *time* when the opinion is expressed.

Semantics into Sentiment Analysis

Traditional approaches hardly cope with subtle linguistic forms, combined and concurrent positive/negative opinions, and implicit judgements.

The literature shows evidence that the **inclusion of semantic features** in sentiment analysis algorithms **improves their overall performance**, e.g. [5].

Linked data, ontologies, controlled vocabularies, and lexical resources help aggregating the conceptual and affective information associated with natural language opinions.

Opinion Model Ontology

Opinion model ontology (1/5)

Our **opinion model** defines the main concepts characterising an opinion, and is used for **annotating** the semantic representation of a sentence, in order to identify its *opinion holder* and *topic*.

Opinion model ontology (2/5)

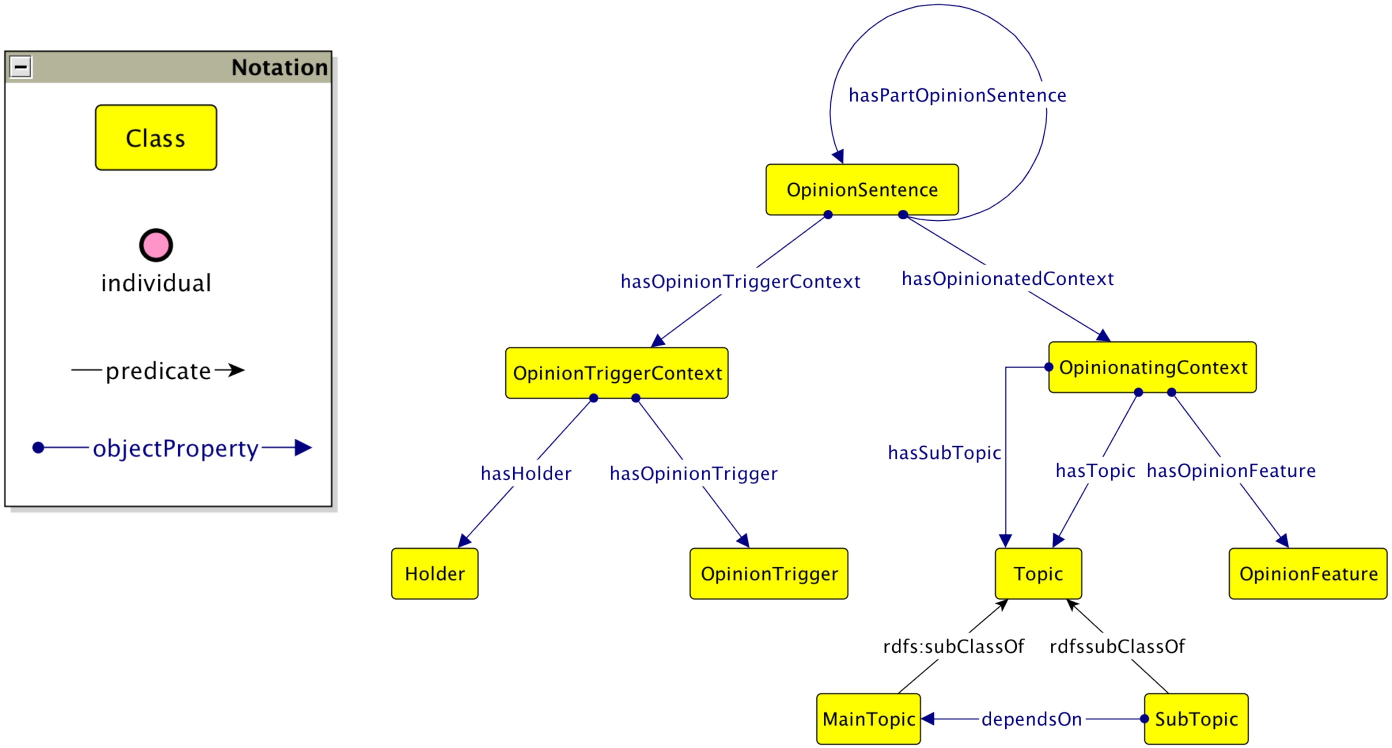

As shown in this Figure, we define the semantic representation of an opinion sentence as having two parts: *(i)* **opinion trigger context** and *(ii)* **opinionated context**.

The **opinion trigger context** is optional and identifies concepts that indicate the presence of an opinion expressed in the sentence, and its holder. It is composed of two parts: the entities that allow the identification of the opinion **holders** (i.e. holders), and the events e.g., think, say, support, etc., that act as **triggers of an opinion** (i.e. opinion triggers).

The **opinionated context** identifies concepts that express an opinion (possibly including sentiments). It is composed of two parts: the entities that identify the **opinion topics** (i.e. topics), and those expressing the opinion and its possible associated sentiments (i.e., the **opinion features**).

In some cases, terms that activate an opinion can also convey the opinion itself or a possible sentiment, e.g. “approve” or “deny”. When this happens, such terms play the role of **opinion triggers** as well as that of **opinion features**.

Opinion model ontology (3/5)

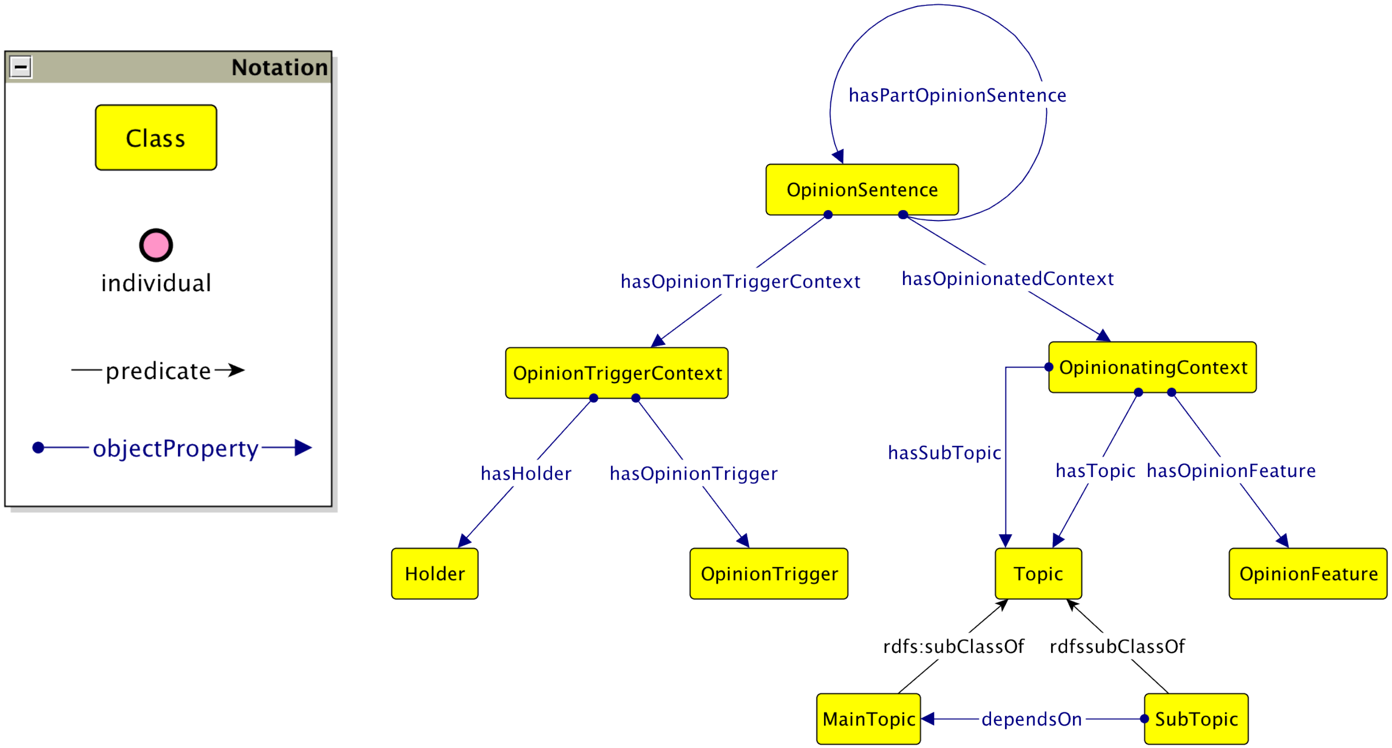

As for **topics**, we distinguish between **main topics** and **sub-topics**.

In fact, a topic can be a complex structure. e.g. an event or a situation, including other complex structures. For the purpose of Sentiment Analysis, it is important to distinguish all main topics, that are the direct targets of an opinion, from sub-topics, which could be indirect targets of an opinion.

For example, the main topic of the opinion sentence: *"Anna says the weather will become beautiful"* is the event *become*, while *weather* is a sub-topic.

Opinion model ontology (4/5)

The following axioms formalise the concepts of main topic and sub-topic by using a standard description logic syntax, directly translatable into OWL (the language used for representing all MARIO ontologies; see Deliverable 5.1).

$$(\texttt{MainTopic} \sqcup \texttt{SubTopic}) \sqsubseteq \texttt{Topic}$$

Topics (of opinion sentences) can be either main topics (direct targets of the opinion) or sub-topics (indirect targets of an opinion).

$$(\texttt{Topic} \sqcap (\exists \texttt{involvedIn}(\texttt{dul:Situation} \sqcap \texttt{MainTopic}))) \sqsubseteq \texttt{SubTopic}$$

When a main topic is a situation, its involved entities are subtopics.

$$(\texttt{Topic} \sqcap (\exists \texttt{dependsOn}(\texttt{dul:Event} \sqcap \texttt{MainTopic}))) \sqsubseteq \texttt{SubTopic}$$

When a main topic is an event, entities that have a dependency relation with it are sub-topics. Examples of dependency relations are: participation in the event (e.g., having a role in it), causality relations, temporal relations e.g., follows, precedes), etc.

Opinion model ontology (5/5)

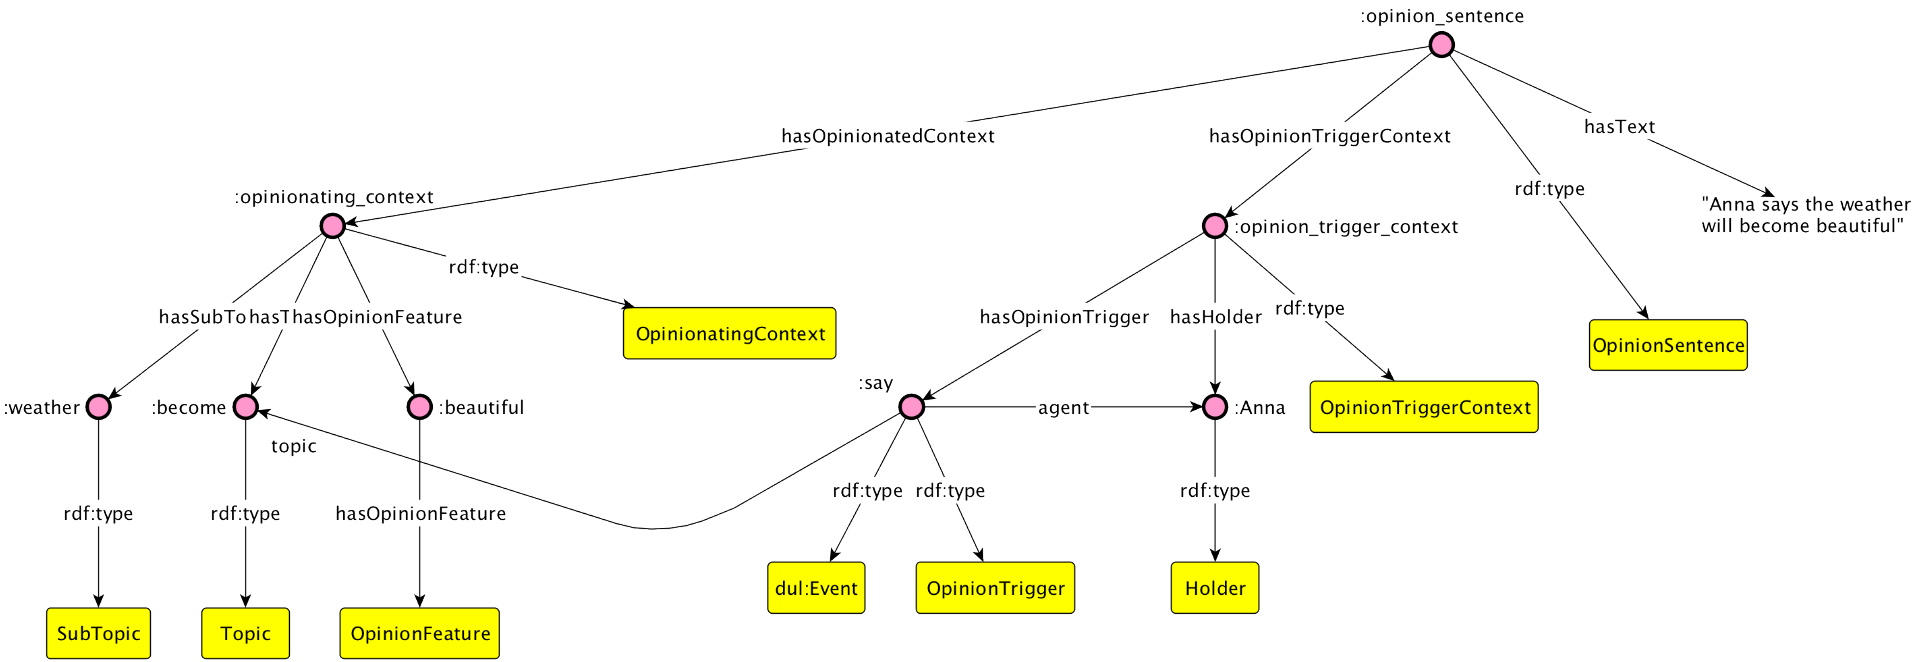

Following the opinion model, the semantic representation of the sentence *"Anna says the weather will become beautiful"* is depicted in the Figure above.

Sentiment Lexical resources

Sentiment lexical resources (1/2)

Lexical resources are key to *annotate* natural language sentence at the “lexical layer” with *sentiment information*.

They provide the *terms* (and possibly a *score* for them) that act as *triggers* of an opinion or as opinion features.

MARIO Semantic Sentiment Analysis module relies on different lexical resources, listed in the next slide.

Sentiment lexical resources (2/2)

- **Sentic-Net** [6]: a publicly available semantic and a affective resource for concept-level opinion and sentiment analysis. SenticNet is built by means of sentic computing, a paradigm that exploits both AI and Semantic Web techniques to better recognise, interpret, and process natural language opinions over the Web. We have transformed Sentic-Net to the RDF format and aligned it with Framester (see D5.1) so as to exploit it in the Sentiment Analysis module.
- **SentiWord-Net** [7]: a lexical resource for opinion mining. SentiWordNet assigns to each synset of WordNet three sentiment scores: positivity, negativity, objectivity. This resource is integrated within Framester with the same approach and aim as the integration of SenticNet.
- **DepecheMood** [8]: an emotion lexicon built by harvesting crowdsourced affective annotation from a social news network.
- **(a revision of) the Levin's classification of verbs** [9]: in this revision we have classified four classes of opinion verbs that imply the presence of an holder; we call them opinion trigger verbs. They are also included in MARIO background knowledge through alignment with Framester.
- **Sentilo-Net** [10]: a resource that annotates semantic roles of frames so as to enable the selection of subtopics that are indirectly affected by opinions expressed in a sentence, as well as the evaluation of their polarity. This resource is key to evaluate sentiment expressed in complex structures such as events and distinguish the different sentiments associated with their participants.

Frame-based sentiment analysis

Frame-based sentiment analysis

This module is implemented as an **extension of FRED**, hence it exploits a *deep parsing analysis* of a sentence and its *frame-based representation*.

FRED as well as the use of BabelNet[1] as part of Framester (see D5.1) guarantee the applicability of our method to both Italian and English, which are the target languages for MARIO.

The **frame-based representation of the sentence** is further annotated with the opinion model ontology (see slides 32-37)

The core of the module is a **sentiment propagation algorithm** that relies on the Sentilo-Net resource (see slide 40). The algorithm computes the **sentiment score** associated with each specific identified *topic* in the sentence representation, according to the role that they play in their participating frame.

In the following slides we show through an example how the algorithm works and how the different lexical sentiment resources are involved.

[1] http://babelnet.org/

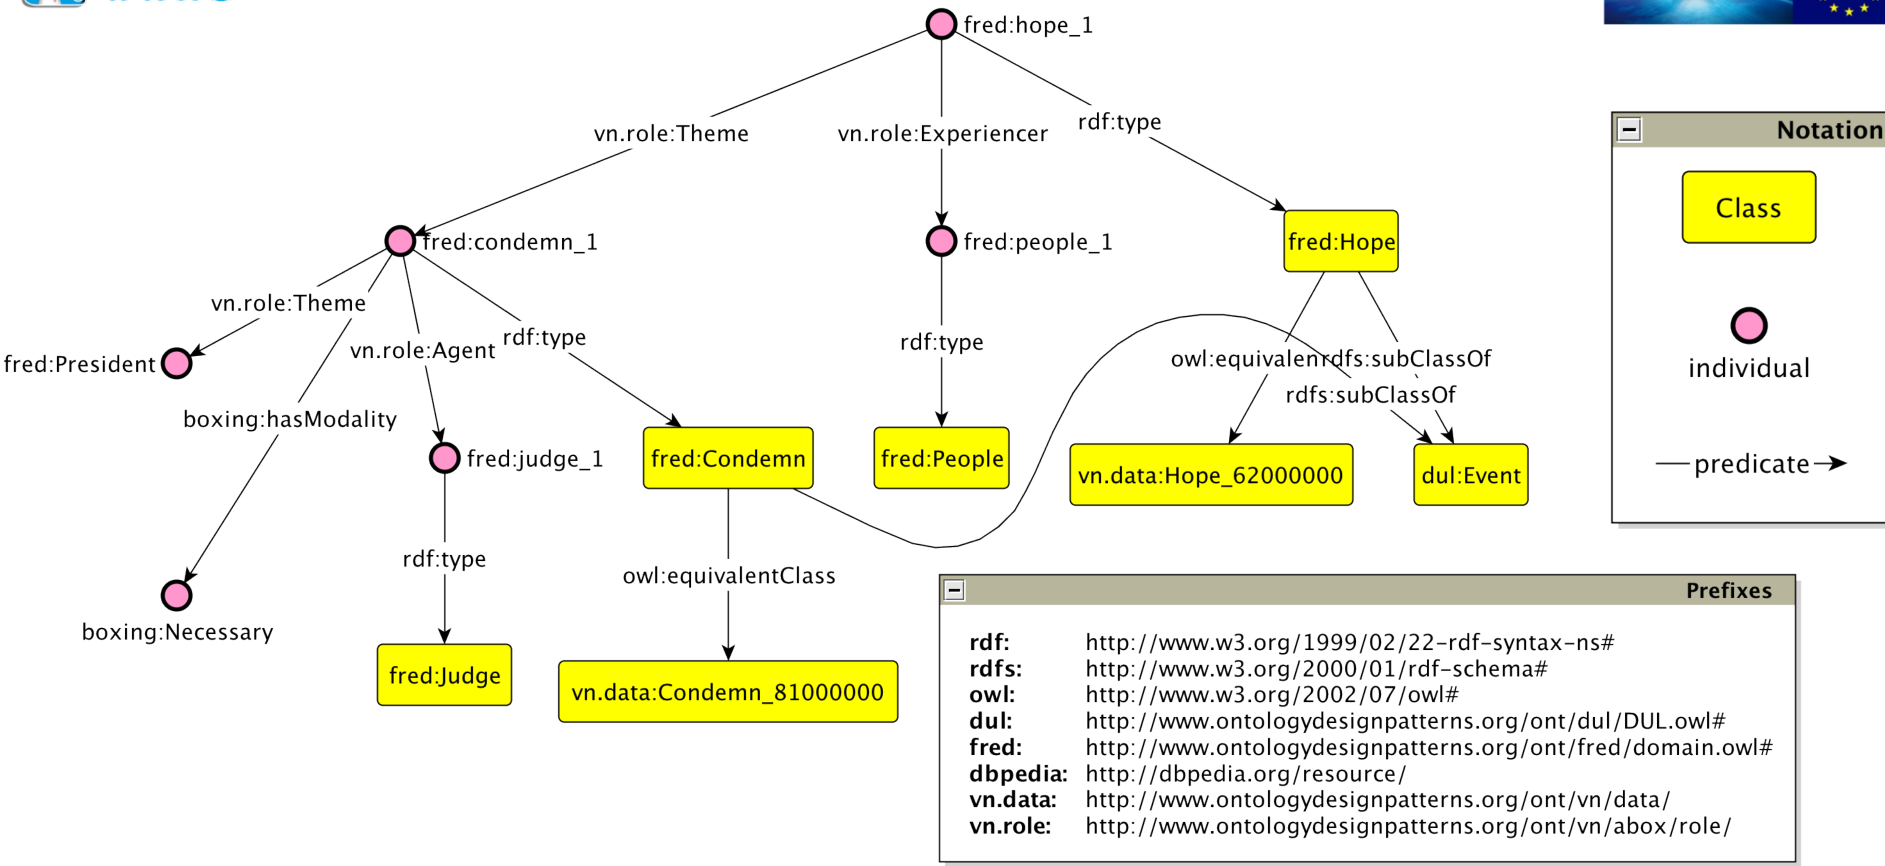

The graph above represents the sentence: *"People hope that the President will be condemned by the judges"*

The sentence is first **processed by FRED** (see D5.6), which provides as output its **frame-based graph representation**.

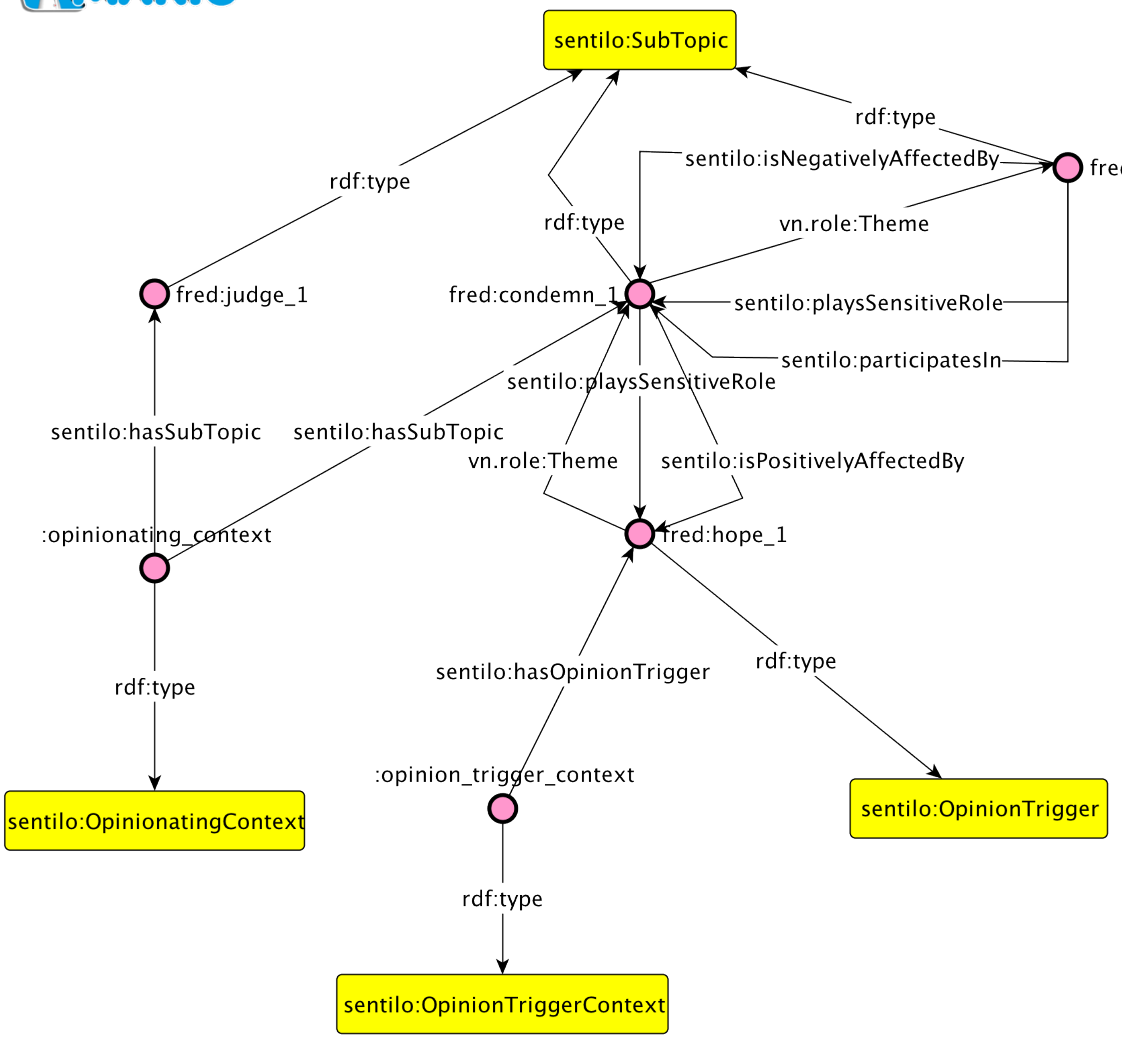

The graph is then passed to the frame-based sentiment analysis module, that **annotates** the resulting graph with the **opinion model ontology**.

This annotation is performed by identifying *triggering verbs*, *opinion holders* and *opinion topics* according to rules associated with the revised Levin's verb classification [9].

Furthermore, the semantic roles identified in the sentence are **annotated** according to the **Sentilo-Net resource**, in order to indicate for each sub-topic whether and how it is indirectly affected by the event (e.g., *sentilo:playSentisitveRole*).

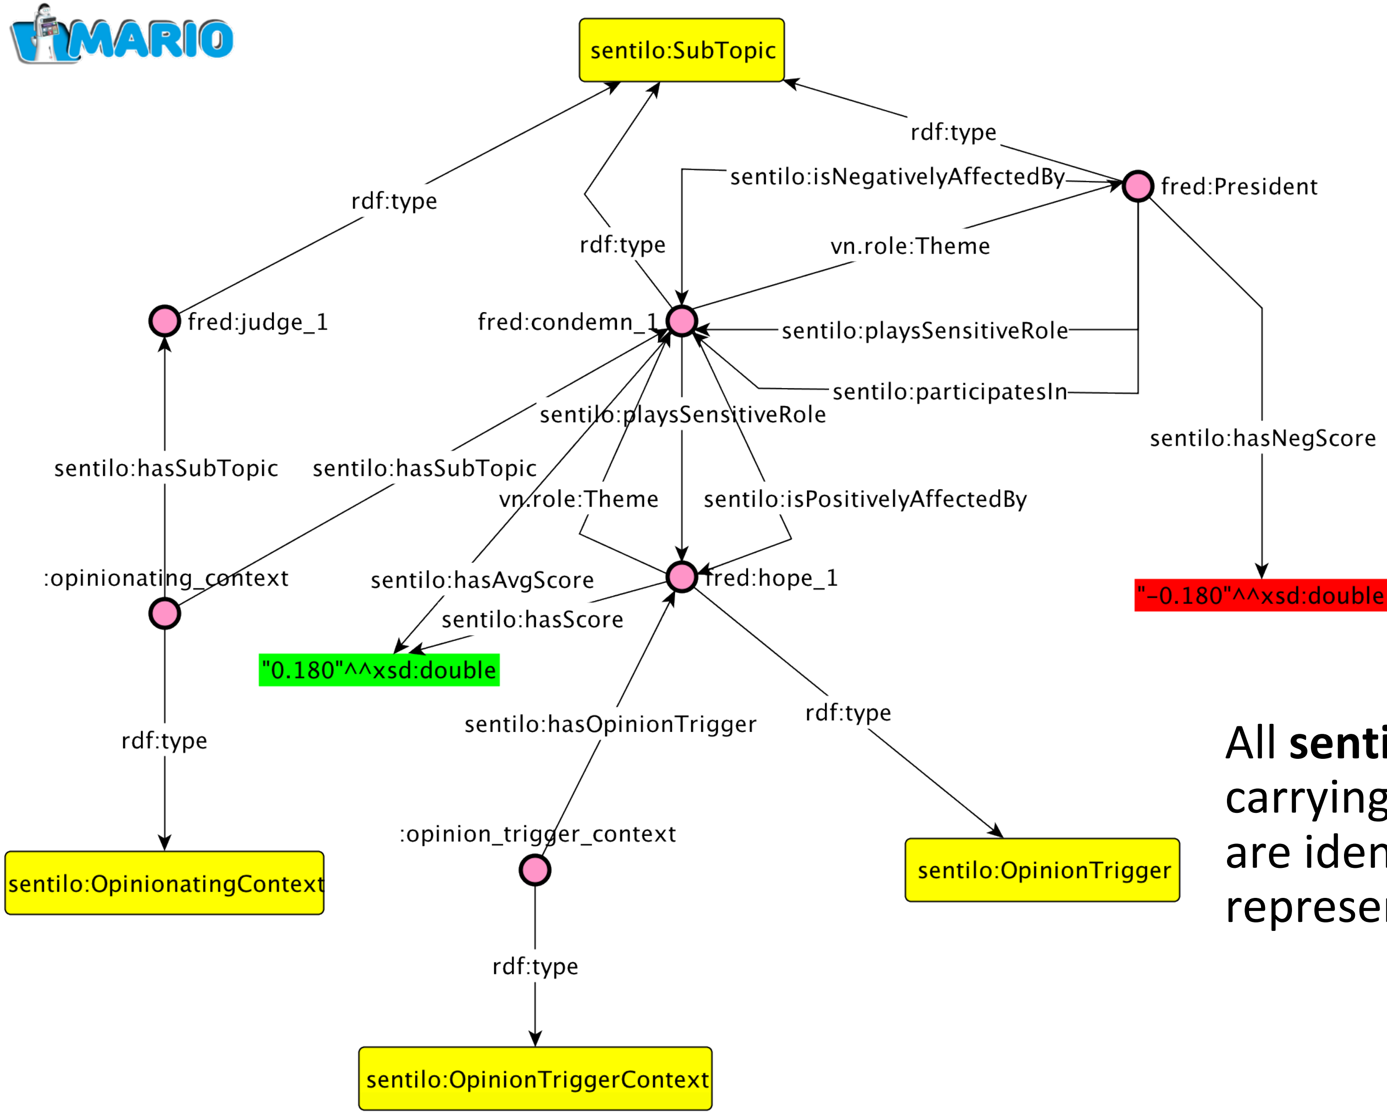

All **sentiment features** (terms carrying a sentiment polarity) are identified and their **scores** represented in the model.

Finally, based on the Sentilo-Net-related annotations, the **scores are propagated through the graph** in order to assign them to the acutal entities they are referred to and with the correct sign.

Sentiment propagation algorithm (1/4)

It relies on a sub-algorithm named *CombinedScore* which assigns an individual sentiment score (if applicable) to each element in an opinion sentence graph. To this aim, it relies on SentiWordNet and SenticNet.

The algorithm assigns a score to adjectives and adverbs that are identified by *dul:hasQuality* relation values, and to instances of *dul:Event* that are recognised as trigger events, i.e., identified by *sentilo:hasOpinionTrigger* relation values.

We have investigated and implemented two alternative approaches for score selection:

- the first approach assigns a score retrieved by querying the polarity attribute of a concept in SenticNet;
- the second one combines SenticNet and SentiWordNet scores.

Sentiment propagation algorithm (2/4)

From a set of empirical observations on using the Propagation algorithm with different approaches, we noticed that **the method that combines the scores from the two resources shows to be more reliable**.

This confirmed our expectations based on the following rationale: sometimes, SenticNet misses a score value for a required concept. Moreover, it provides one score per concept without distinguishing its possible different nuances. Hence, SenticNet score approximates an average value for the scores of all possible senses, or possibly indicates the most probable one.

For this reason, combining the SentiWordNet scores of most frequent senses and the SenticNet score can provide an appropriate balanced value. We have devised a *simple heuristics* for computing this combined score.

In the next slide the CombinedScore algorithm is sketched.

Sentiment propagation algorithm (3/4)

CombinedScore Algorithm

score = **CombinedScore**(w){
sNet = SenticNet score for w;
n = number of w senses;
T = { };
for $i \leftarrow 1$ ***to*** n **do**
    s_i = extract next sense of w from WordNet in decreasing order of tag_count;
    if *tag_count[s_{i-1}] > 10 × tag_count[s_i]* **then**
        break;
    T = T ∪ {s};
end
T' = SentiWordNet score values for each element in T;
sWN = AVERAGE(T'{..});
return AVERAGE(sWN,sNet);
}

Sentiment propagation algorithm (4/4)

Given an **entity**, identified as a *topic of an opinion* (either a main or subtopic), we compute its **sentiment score** by *combining the scores* of all its associated *opinion features* (i.e., values of *dul:hasQuality* relations), which are extracted from the RDF graph representing the opinionated sentence.

If a topic participates in an *event* or a *situation occurrence*, we **propagate** their sentiment scores to it, according to the semantics expressed by the frame-based thematic role (e.g., *vn.role:Agent*) that it plays, its sensitiveness and factual impact attribute values.

Sentiment score

The sentiment score $SC_{sentiment}$ of a topic *t* can be defined as a function *f* defined as:

$$SC_{sentiment} = f\left(\sum_{i=0}^{n} sc(q_i(t), \sum_{j=0}^{m} type_j(t), \sum_{i=0}^{n} cxt_k(t), truth(t), mod(t), sc(trig(sent))\right)$$

Where:

- *sc(x)* is the score of an entity *x* as provided by the CombinedScore algorithm;
- $q_i(t)$ is an object value of a triple *t dul:hasQuality* q_i. Such triples represent direct opinion features, i.e. adjectives and adverbs, associated with entities composing the opinion sentence;
- $type_j(t)$ is a type of t expressed in the RDF graph by means of *rdf:type* triples;
- $cxt_k(t)$ is a context of *t*, if any. It can be either a situation or an event, which *t* participates in;
- *truth(t)* is a truth value associated with *t*, where *t* is typically an event or situation occurrence, or a quality. If its value is false, it means that the entity is negated
  - for example, in a sentence such as ''John is not a good guy'', a RDF triple *situation_1 boxing:hasTruthValue boxing:False* would be included in the graph, and its effect would be to change the sign of the sentiment score assigned to the feature good;
- *mod(t)* is a marked modality of a topic *t*, if any
  - for example, in a sentence such as I would like a dog, an RDF relationship *fred:like_1 boxing:hasModality boxing:Necessary* would be included. At this time, the propagation algorithm does not yet use this information, but its abstract model, the *f* function, includes it;
- *trig(sent)* is an opinion trigger expression in the sentence containing *t*.

Sentiment propagation algorithm flowchart

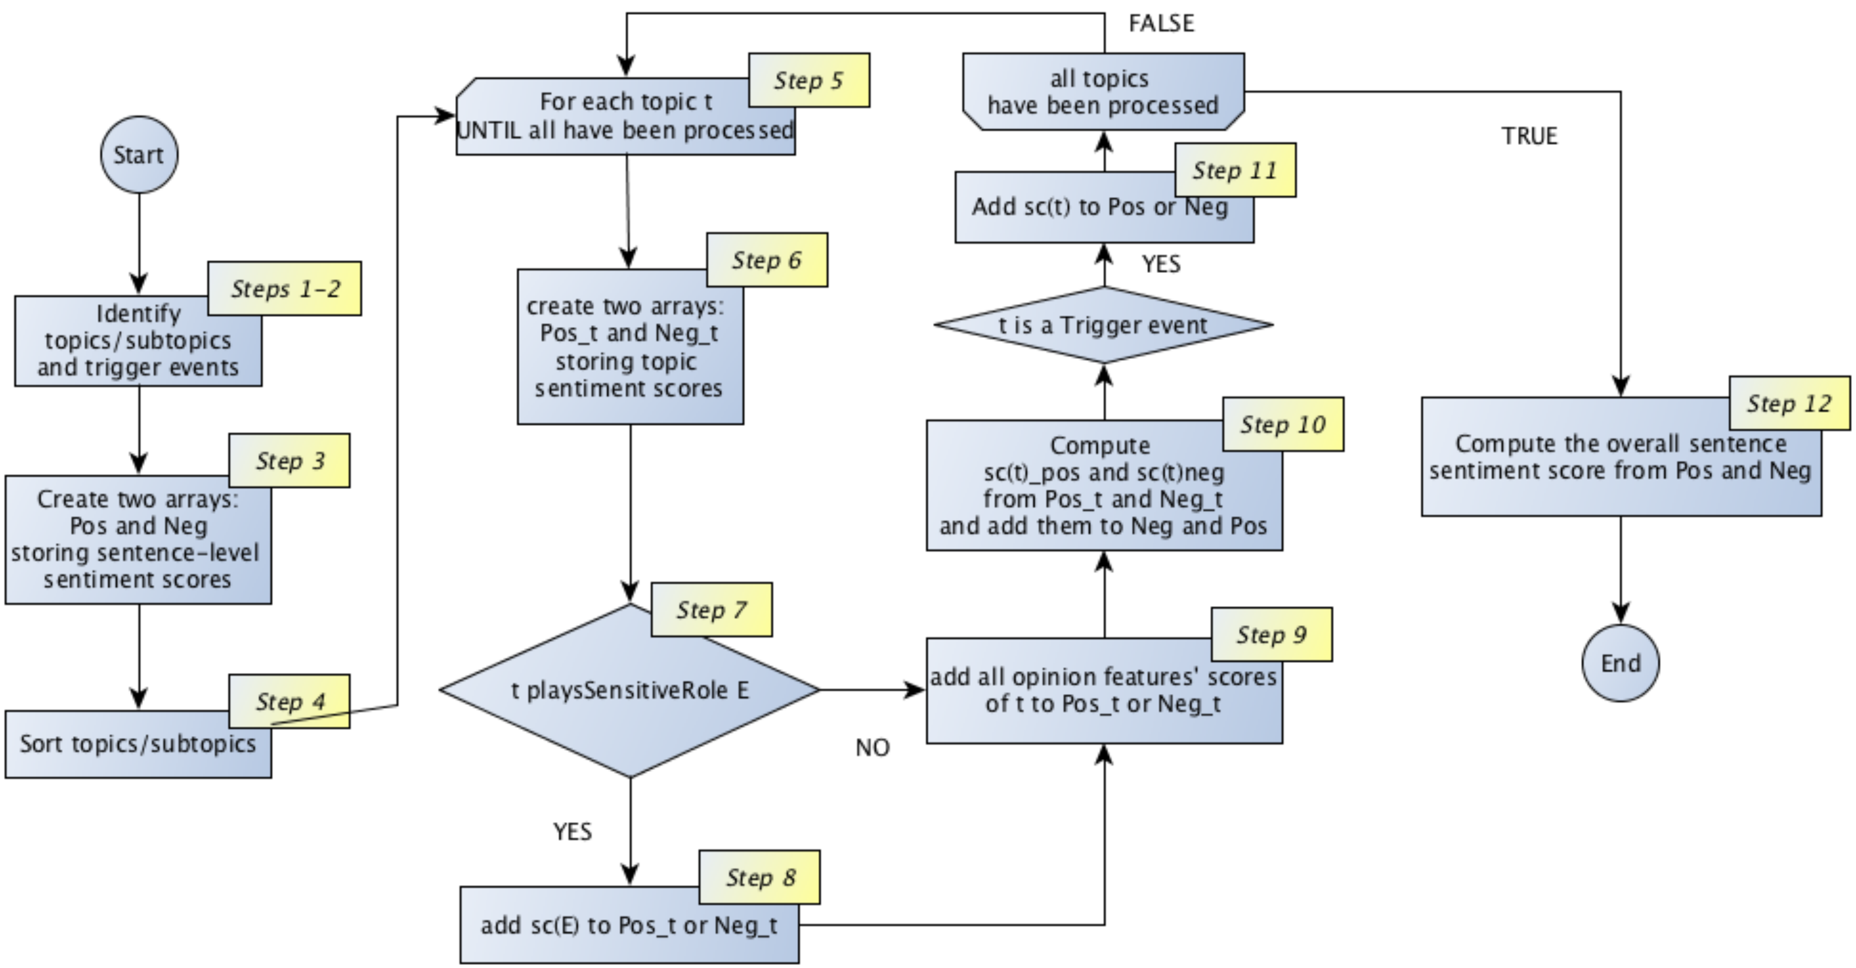

Representative Use Case Scenario

Sentiment Analysis in the My Memories Application

Sentiment analysis in the My Memories App

The **My Memories app** is a representative example of a MARIO application that combines the capabilities of **NLP modules and services** with the capabilities of **Sentiment Analysis modules and services**, as part of its logic and interaction/dialogue management strategy.

- App requirements, characteristics and design are detailed in Deliverable 3.3.
- App interaction with MARIO's NLP modules and dialogue management are presented in Deliverable 5.6.

In the following, we describe the peculiarities of this app in terms of Sentiment Analysis and their impact on the app's dialogue management approach

- for the sake of self-containedness and readability of this Deliverable, the main characteristics of the app reported in D5.6 are summarised again.

Research outcomes related to this apps have been presented at the *1st International Workshop on Application of Semantic Web technologies in Robotics (AnSWeR)*, co-located with the 14th Extended Semantic Web Conference (ESWC 2017) in Portoroz, Slovenia [11].

My Memories Application

Enables the robot to undertake interactive and personalised **reminiscence** sessions through a *conversational approach* based on user-specific knowledge and materials.

- **Knowledge Base (KB) support:** user profiles, family/social relationships, life events, tagged media objects (e.g., photographs).
- **interaction patterns:** prompt the user with focused memory triggers (verbal prompt + photograph).

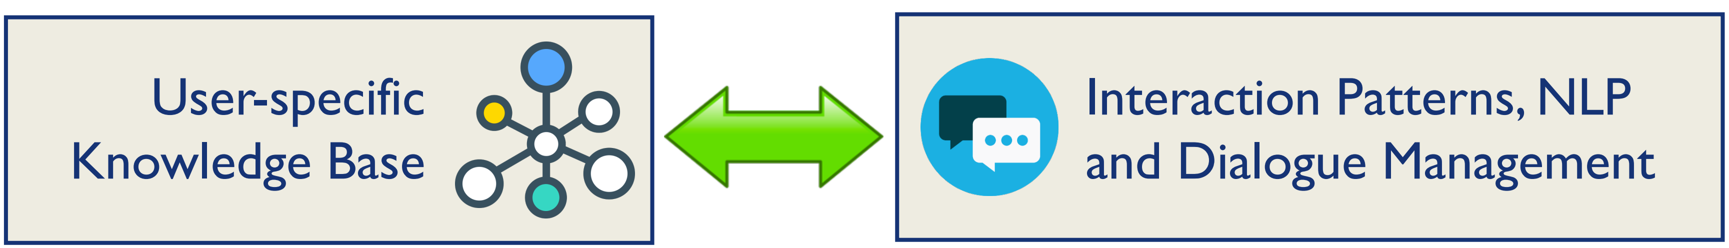

Architectural Reference

Manages reminiscence sessions and dialogue state and flow
- interaction pattern and prompt selection
- NLP services invocations and interpretation

Parametric interaction patterns provide dialogue blueprint
- question/prompt formulation
- NLP interpretation strategy

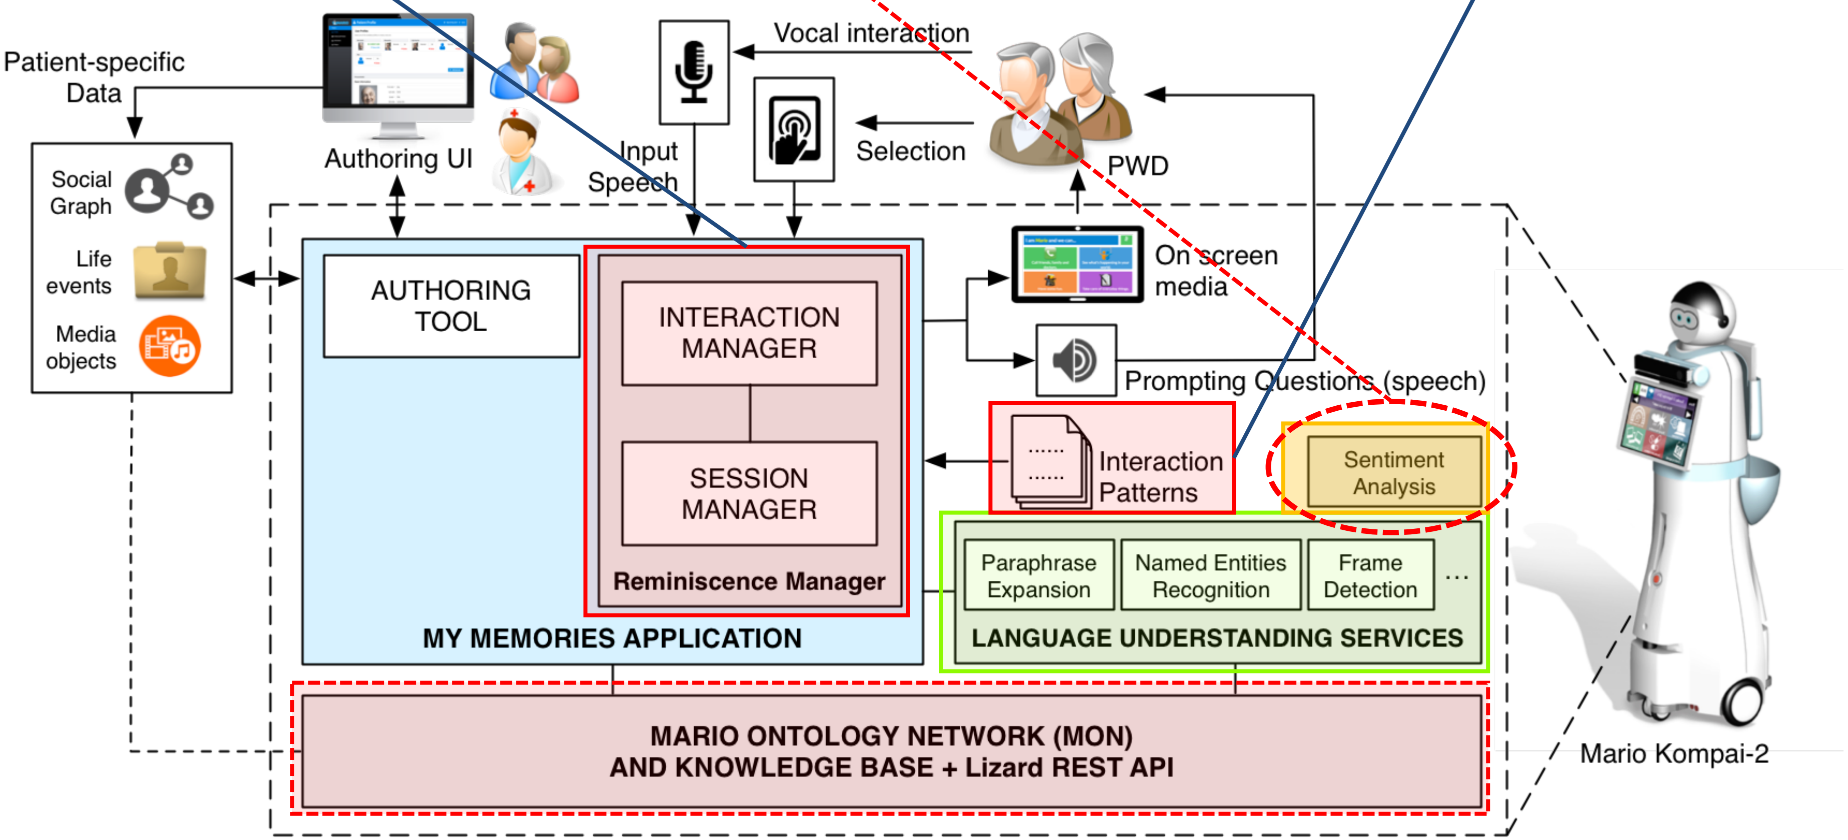

Interaction Patterns

The interaction with the user during a reminiscence session can follow two different conversational approaches, both based on *system-initiated* dialogue fragments defined in the form of interaction patterns.

1. **Question-based:** MARIO asks the user focused *closed-ended* questions related to the image contextually shown as memory trigger
   - detailed in D5.6
2. **Prompt-based:** MARIO prompts the user with *open-ended* prompts or questions related to the image:
   - prompts aim at stimulating reminiscence about people (e.g., *"What was you sister like as a child? Tell me more about her!"*), places (e.g., *"What was it like to grew up in London?"*) and life events (e.g., *"Tell me more about your wedding day!"*)
   - *sentiment analysis* capabilities enable MARIO to identify the polarity or sentiment expressed by the PwD and to reply and act accordingly

Prompt-based Interaction Patterns

Pattern structure and elements (JSON-encoded)

`precondition`	Constraint mapping to queries over the KB, defining under which conditions the prompting question can be used.
`prompt`	Multilingual parametric prompt templates to be instantiated with KB data (entities and their property values).
`ifPositiveSay`	Multilingual parametric system utterance template to be instantiated if user’s answer has **positive** sentiment.
`ifNeutralSay`	Multilingual parametric system utterance template to be instantiated if user’s answer has **neutral** sentiment.
`ifNegativeSay`	Multilingual parametric system utterance template to be instantiated if user’s answer has **negative** sentiment.

Prompt-based Interaction Process (1/2)

For an *interaction pattern* whose *applicability preconditions* are satisfied (wrt the Knowledge Base, and in particular for a specific photograph), the dialogue is managed according to the following main steps:

- the corresponding *prompt template* is instantiated and the prompt is issued to the PWD;
- the textual representation of PWD's vocal input is processed relying on the capabilities of the Sentiment Analysis subsystem;
- depending on the sentiment and score identified for PWD's answer in the sentiment analysis step, the corresponding utterance is issued by Mario, as defined in the interaction pattern.

The overall process is graphically summarised in the following slide and then illustrated with a concrete example.

Prompt-based Interaction Process (2/2)

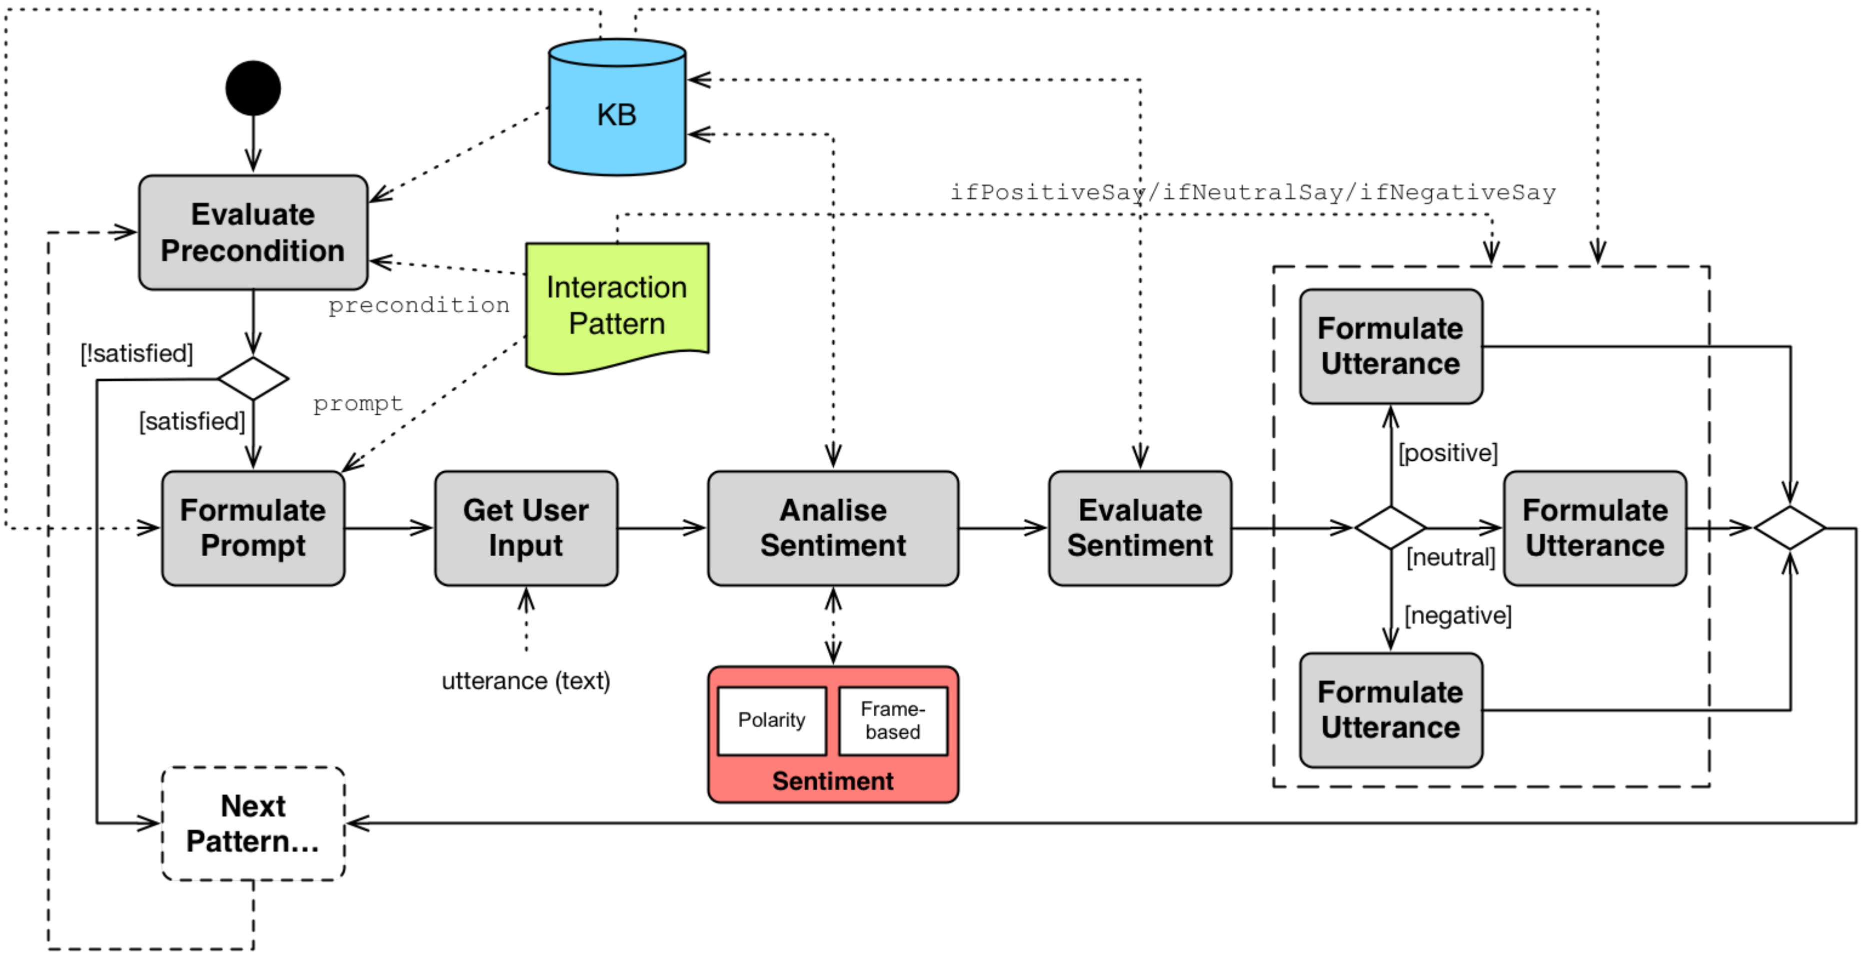

Example

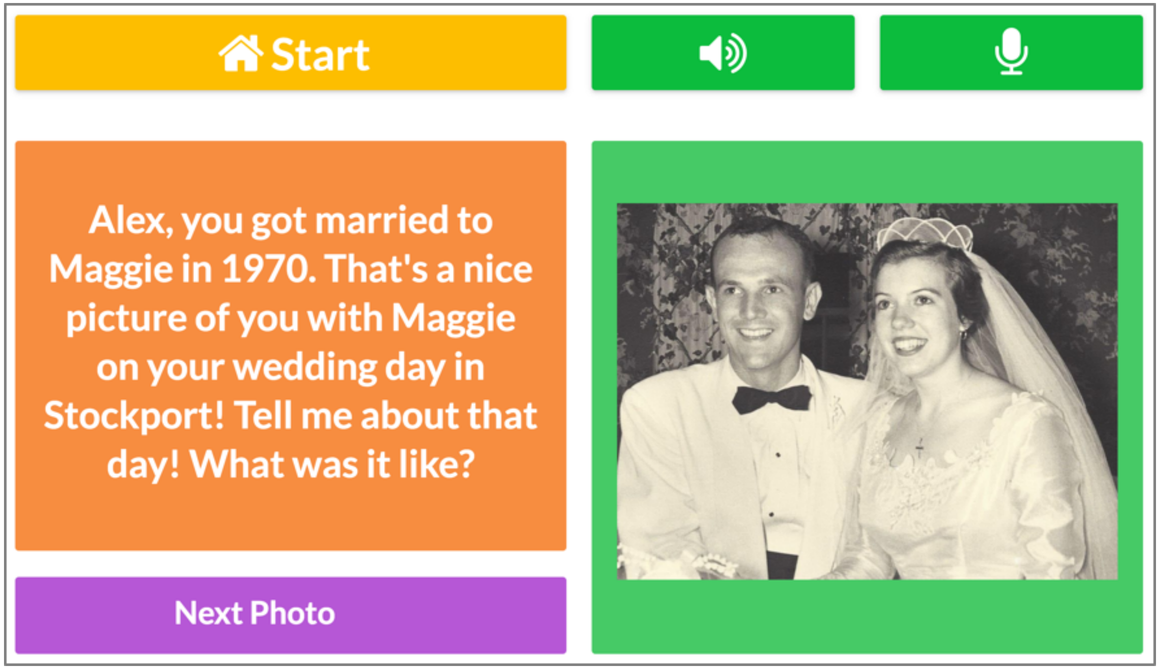

USER ANSWER PROCESSING

Task: **sentiment analysis** (polarity detection) over user's utterance

Possible outcomes

- ***Positive*** ⟶ encourage to tell more ("That sounds nice! Tell me more about your marriage!"), show other pics of same event
- ***Neutral*** ⟶ proposal ("Tell me more if you like, or ask me to show you another photo!")
- ***Negative*** ⟶ move to next photo

Example – Polarity detection (1/2)

“I remember it was a great day!”

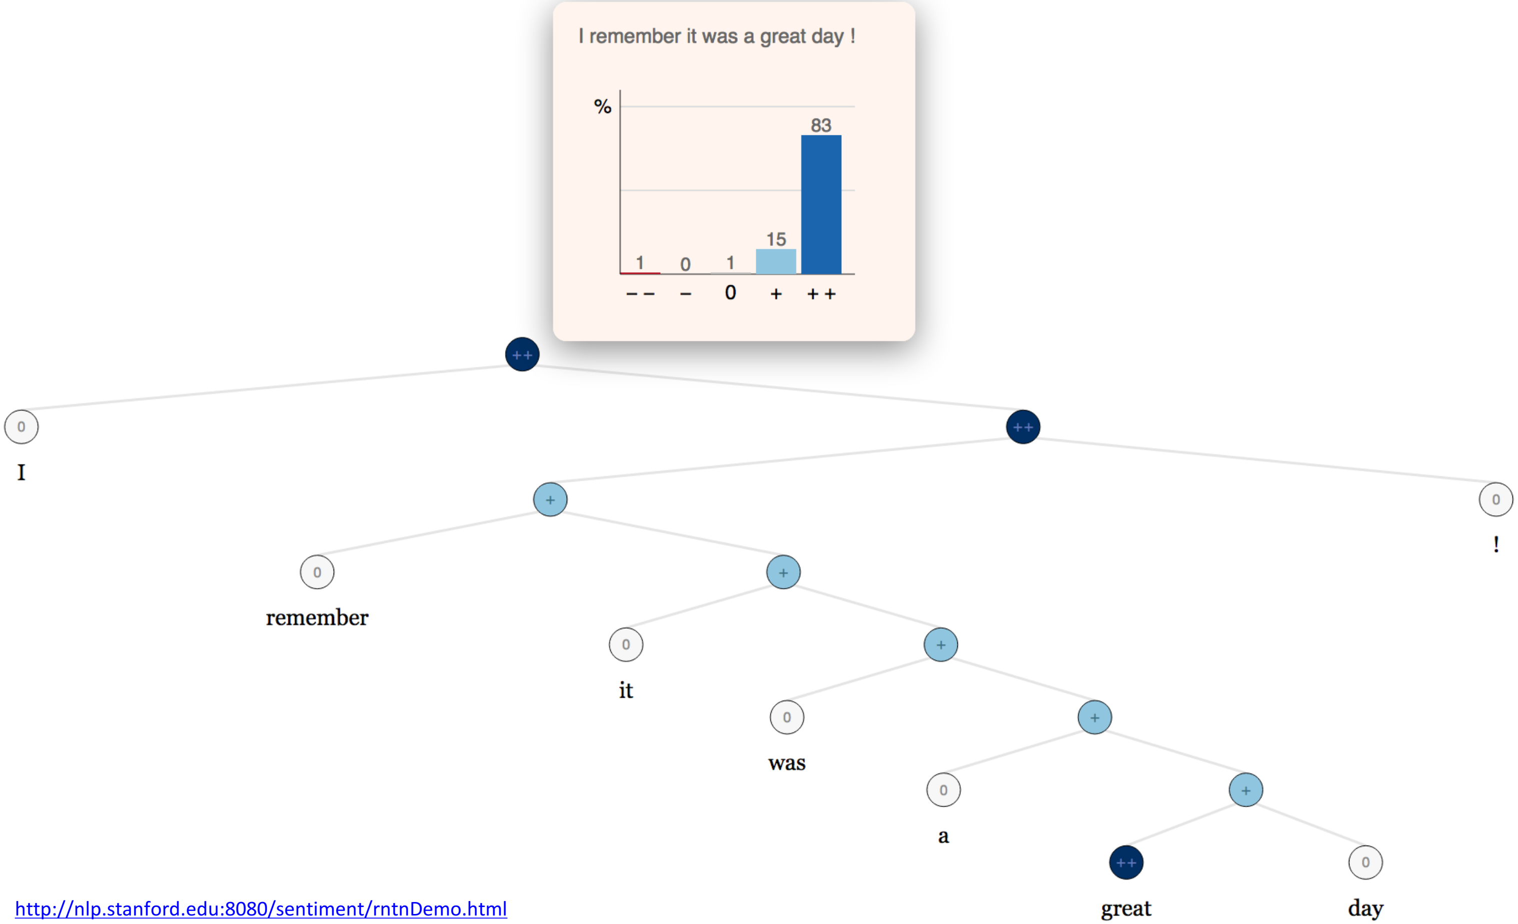

http://nlp.stanford.edu:8080/sentiment/rntnDemo.html

Polarity detection (1/2)

Simple polarity detection enables MARIO to identify user's reaction to a specific prompt and image.

Polarity info provided by the sentiment analysis module can be

- used to drive how MARIO proceeds in the reminiscence session
  - use different utterances and strategies for different polarities (as outlined in the previous slides: show empathy and ask to tell more for a positive reaction, propose to move to next photo for a negative reaction, etc.);
  - select next interaction pattern depending on polarity, for example by choosing or avoiding photos/prompts related to the current subject (life event, person, etc.);
- stored in relation to the image and subject of the prompt
  - reused in the image/pattern selection process in subsequent reminiscence sessions (e.g., favouring memory triggers generating positive feelings).

However, no info can be extracted on specific topic(s) or entities about which the sentiment is expressed (e.g., the wedding day itself, the wife, the picture, etc.)

- finer-grained sentiment polarity scores are provided by frame-based sentiment analysis.

Polarity detection (2/2)

The polarity detection process is also influenced by the characteristics of PWD’s utterances and their processing by the Speech to Text (S2T) component (described in D5.6).

- If an utterance coming from the S2T is recognised as a single sentence by the sentence split step in the NLP pipeline, the polarity of the sentence is analysed and evaluated, and MARIO reacts as outlined before (see slides 58-59).
  - For a given prompt, the PWD may express different opinions is different utterances, each corresponding to a single sentence and followed by MARIO’s reaction.
  - Each sentence is processed and analysed independently (e.g., a positive reaction will first lead MARIO to encourage the PWD to tell more about the current person/event/place; a subsequent sentence with negative polarity will then lead MARIO to suggest to move to another item).
- If an utterance coming from the S2T is recognised as multiple sentences by the sentence split step in the NLP pipeline, the polarity of each sentence is analysed and evaluated.
  - The overall polarity of the utterance can be assessed by combining the polarity scores of the sentences (e.g., by computing an average sentiment score over the number of sentences).
  - Again, no info can be extracted on specific topic(s) or entities about which (potentially different) sentiments are expressed.
  - For example, an utterance with a positive sentence and a negative one with similar scores would be considered as neutral, regardless of the mentioned topics or entities; no info is captured on the positive and negative attitudes and the corresponding topics or entities that characterise each sentence.

Example – Frame-based analysis (1/2)

MARIO can model the question by using the frame-based approach enabled by FRED.

What was your wedding day like?

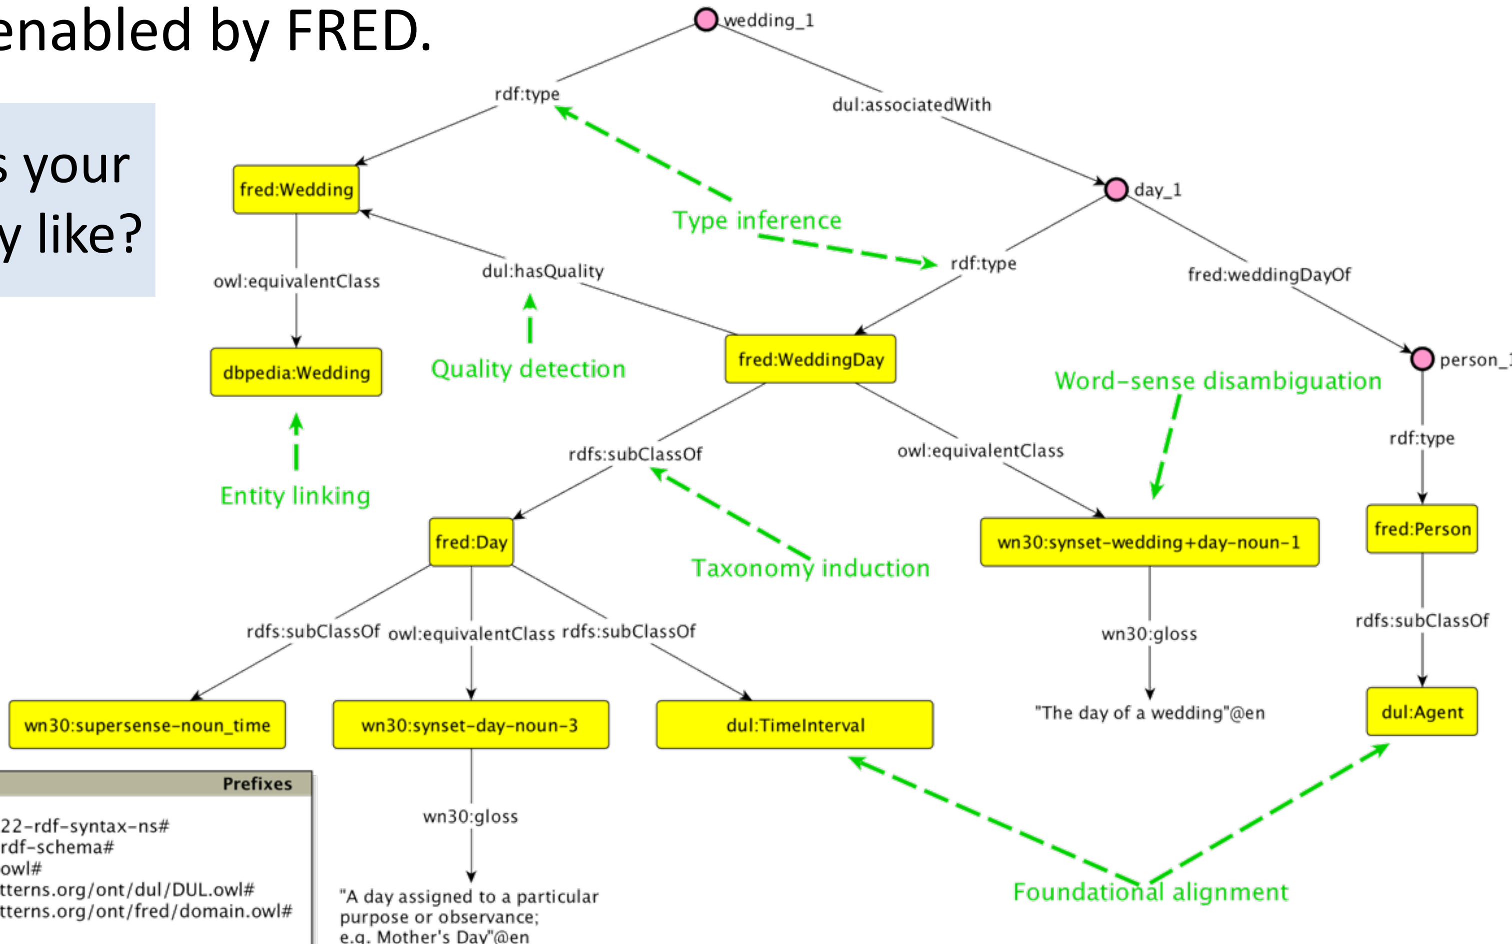

Notation

Class

individual

predicate

Prefixes

rdf:	http://www.w3.org/1999/02/22-rdf-syntax-ns#
rdfs:	http://www.w3.org/2000/01/rdf-schema#
owl:	http://www.w3.org/2002/07/owl#
dul:	http://www.ontologydesignpatterns.org/ont/dul/DUL.owl#
fred:	http://www.ontologydesignpatterns.org/ont/fred/domain.owl#
dbpedia:	http://dbpedia.org/resource/
wn30:	http://www.w3.org/2006/03/wn/wn30/instances/

Example – Frame-based analysis (2/2)

Frame-based representation of the user's opinion sentence

"I remember it was a great day!"

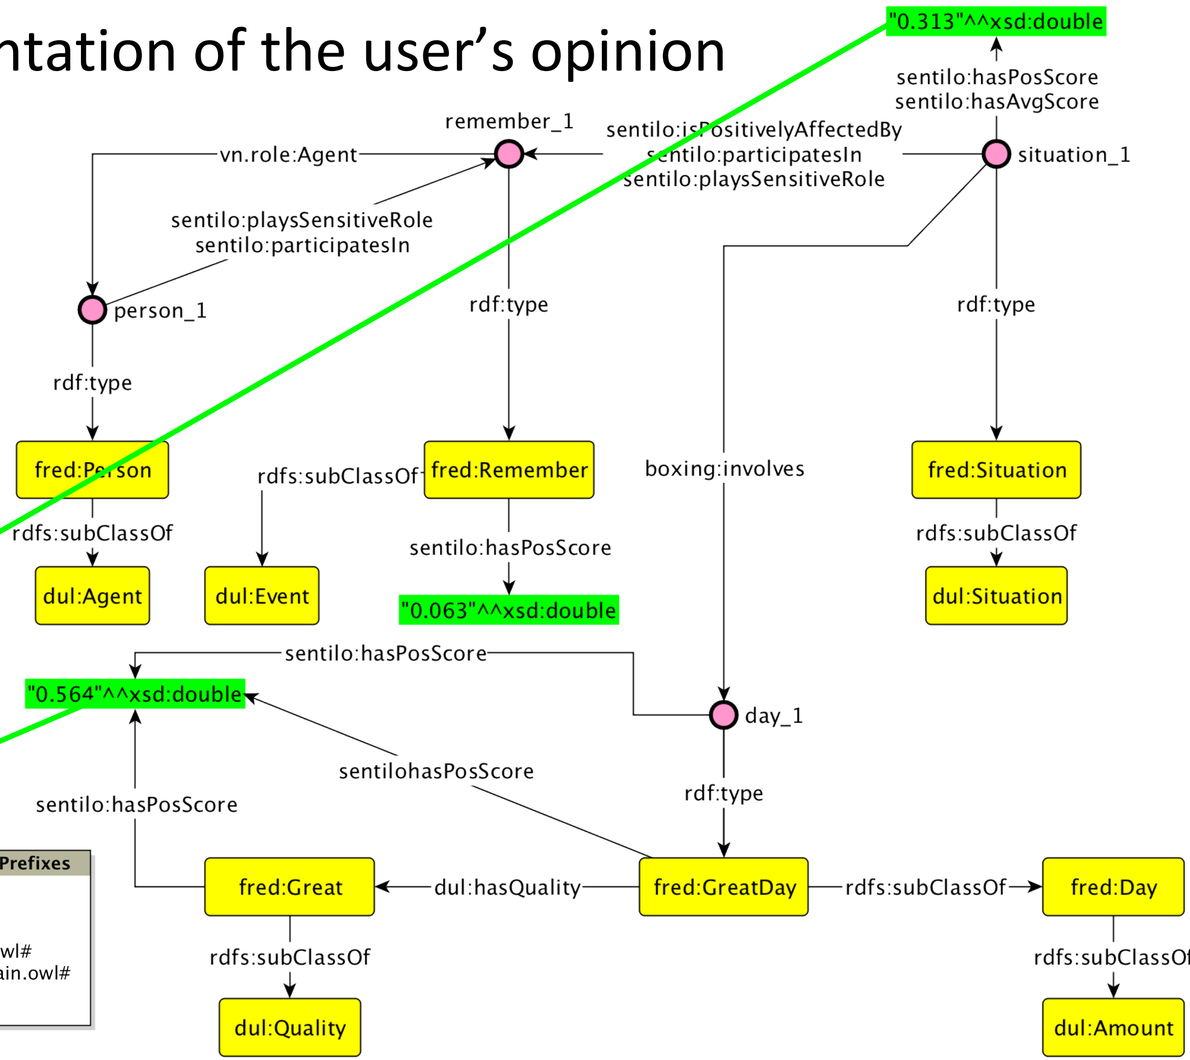

The reminiscence associated with the Patient's wedding day is recognised as a positive situation (0.313 polarity value)

A positive score (0.564) is also specifically associated with the entity representing the day

Prefixes

rdf:	http://www.w3.org/1999/02/22-rdf-syntax-ns#
rdfs:	http://www.w3.org/2000/01/rdf-schema#
dul:	http://www.ontologydesignpatterns.org/ont/dul/DUL.owl#
fred:	http://www.ontologydesignpatterns.org/ont/fred/domain.owl#
sentilo:	http://ontologydesignpatterns.org/ont/sentilo.owl#

Concluding Remarks (1/2)

- The software components presented in this document are part of the overall MARIO framework **deployed on 6 Kompai-2 robots** used for evaluation and validation activities in the 3 project pilot sites:
  - Geriatric Unit at the hospital Casa Sollievo della Sofferenza, Italy;
  - Nursing Centres at the National University of Ireland, Galway;
  - Communities in Stockport.
- At the time of writing, evaluation and validation activities are undergoing (according to the time-line defined in WP8).
- The data and feedback that will be collected during the final validation period (Aug-Nov 2017) will serve as a basis for assessing, in different settings, the capabilities of the Sentiment Analysis subsystem

Concluding Remarks (2/2)

- While simple polarity-based sentiment analysis techniques are currently used in multiple domains, the frame-based semantic sentiment analysis approach investigated in Task 5.3 represents a major contribution, in particular for the targeted domain
- Ongoing research activities, whose timeframe goes beyond WP5, aim to investigate how the approach can be extended for:
  - *aspect-based* sentiment analysis, to recognise the sentiment expressed on specific aspects/features related to a topic or entity
  - *emotion analysis*, to detect specific emotions or emotional states beyond polarity

References (1/2)

[1] Stanford CoreNLP Sentiment. Available at: http://nlp.stanford.edu/sentiment/code.html

[2] R. Socher, A. Perelygin, J. Wu, J. Chuang, C. Manning, A. Ng, and C. Potts. *"Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank"*. In: Proc. of the Conference on Empirical Methods in Natural Language Processing (EMNLP 2013), 2013

[3] Bing Liu. “Sentiment Analysis and Subjectivity”. Handbook of Natural Language Processing, 2010

[4] Bing Liu. “Sentiment Analysis and Opinion Mining”. Morgan & Claypool Publishers. May 2012

[5] H. Saif, Y. He, and H. Alani. “Semantic Sentiment Analysis of Twitter”. Boston, UA, pp. 508–524, Springer, 2012

[6] E. Cambria, C. Havasi, and A. Hussain, *“SenticNet 2: A semantic and affective resource for opinion mining and sentiment analysis”.* In: Proc. FLAIRS Conf., 2012.

[7] A. Esuli, S. Baccianella, and F. Sebastiani, *“SentiWordNet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining”.* In: Proc. 7th conf. Int. Language Resources Evaluation, 2010.

[8] DepecheMood. Available at: https://github.com/marcoguerini/DepecheMood/releases

[9] Levin opinion. Available at: http://www.stlab.istc.cnr.it/documents/sentilo/levin-opinion.zip

[10] Sentilo-Net. Available at: http://www.stlab.istc.cnr.it/documents/sentilo/sentilonet.zip

References (2/2)

[11] L. Asprino, A. Gangemi, A. G. Nuzzolese, V. Presutti, A. Russo: *“Knowledge-driven Support for Reminiscence on Companion Robots”*. In: Proceedings of 1st International Workshop on Application of Semantic Web technologies in Robotics (AnSWeR 2017), Portoroz, Slovenia, 2017

[12] A. Gangemi, V. Presutti, D. Reforgiato Recupero, A. G. Nuzzolese, F. Draicchio, M. Mongiovì: *“Semantic Web Machine Reading with FRED”*. Semantic Web Journal, 8(6), 2017

[13] A. Gangemi, M. Alam, L. Asprino, V. Presutti, D. Reforgiato Recupero. *“Framester: A Wide Coverage Linguistic Linked Data Hub”*. In: Proceedings of the 20th International Conference on Knowledge Engineering and Knowledge Management (EKAW 2016), pp. 19-23, 2016

[14] Framester - A giant hub of linguistic linked open resources. https://lipn.univ-paris13.fr/framester/