Knowledge-based Intelligence and Strategy Learning for Personalised Virtual Assistance in the Healthcare Domain

Eleni Kamateri*, Georgios Meditskos*, Spyridon Symeonidis*, Stefanos Vrochidis*,
Ioannis Kompatsiaris* and Wolfgang Minker†
*Information Technologies Institute
Centre for Research and Technology Hellas, 6th Km Charilaou-Thermi Road, Thessaloniki, Greece
Email: {ekamater, gmeditsk, spyridons, stefanos, ikom}@iti.gr
†Institute of Communications Engineering
Ulm University, 89081 Ulm, Germany
Email: wolfgang.minker@uni-ulm.de

***Abstract*—This paper introduces a virtual assistant framework that combines knowledge-based and statistical techniques to produce meaningful task-oriented conversations that are enhanced by "chatty" style dialogues in order to increase system's naturalness and user engagement. The paper describes how appropriate ontologies, semantic reasoning, dialogue management and policy learning techniques can be linked together and integrated through the dialogue process to enable a) the internal representation of the conversational state, b) the conversational awareness that drives the retrieval of appropriate information from the Knowledge Base (KB) and the inference of unrelated system actions with the current conversational state, and c) the dynamic selection of the most appropriate strategy at each dialogue turn, tackling both informational and social-related needs of individuals. The framework is exemplified by a use case from the healthcare domain where companionship and supportive care-related services are prerequisites for an efficient human-system interaction through a conversational agent.**

***Keywords*–Dialogue management; Knowledge representations; Reasoning; Strategy learning; Virtual assistance.**

I. Introduction

Nowadays, there is an increasing demand for intelligent agents. A challenging domain includes personalised virtual assistants that carry out human-like conversations taking into account the latest user's utterance, the dialogue history, as well as the background knowledge about the user. The development of such personalised systems requires a knowledge representation model for describing the semantics of various contexts and structuring the background knowledge about individuals.

Current task-oriented dialogue systems focus on one task at a time using frame-based [1] or agenda-based [2] mechanisms, while it was only recently, when some ontology-based dialogue systems (such as [3] and [4]) have been proposed using semantic models for the representation of user's utterance and the generation of the system's response. Access to a rich domain model and the conversation memory can deal with complex task-oriented dialogues. However, the typical problem of task-oriented dialogue solutions remains that is the difficulty of tackling user utterances that go beyond the agent's representational model and the smooth transition between task-oriented and "chatty" style dialogues.

To succeed this, we propose a hybrid dialogue framework that can be placed at the heart of any personalised virtual assistant to enhance its model-driven operation by "chatty" style responses. The proposed approach, which is an on-going work, combines knowledge representation and reasoning with statistical learning for the smooth transition between strategies, discussion topics and available knowledge with the aim to impose social skills in the personalised virtual assistants in order to efficiently realise meaningful task-oriented conversations, recover breakdowns in a natural way, and increase user engagement.

Our major contributions are summarised as follows:

1) **a domain and a dialogue representation model** are proposed and populated with local semantics coming from the language analysis of the user's utterance by means of semantic similarity and disambiguation techniques,
2) **a dialogue history representation model** is proposed and populated with global semantics of the entire dialogue session at each dialogue turn,
3) **semantic reasoning techniques** are applied on top of the semantically structured data with the aim to generate dynamically-inferred insights and actions,
4) **a dialogue management technique** analyses the system's confidence regarding the task-oriented response and produces a set of social-oriented action candidates, and
5) **a strategy selection technique** is used to select the appropriate strategy, i.e., action.

Such personalised virtual assistants can have many applications in the healthcare domain and provide a mixture of companionship and supportive care-related services, improving the quality of life of individuals. We selected to apply our framework in a rehabilitation setting, which involves people with motor, cognitive and behavioural disorders being in a clinical environment or after returning home.

The rest of the paper is structured as follows: Section II presents related work on dialogue systems. Section III describes the specifics of the proposed framework, elaborating on the representation, reasoning and dialogue management capabilities. Section IV presents an example use case in the rehabilitation domain, where the framework is currently being used. Finally, Section V concludes our work, mentioning future research directions.

II. Related Work

First conversational systems were mainly task-oriented (e.g., [5] realises restaurant reservations) lacking social competences. More recent personal assistants, such as the commercial platforms of Alexa, Siri, Google Assistant and Contana, have started to incorporate social features and support non-task-oriented dialogues as well, where users do not have a clear goal or intention. However, these systems are usually model-less, constrained to accessing the parameters of the last users utterance and thus, they are acceptable only for simple tasks that do not need to sustain the whole conversation memory.

On the other hand, non-task-oriented dialogue systems do not have a specific goal and are capable of addressing a wide range of topics. To succeed this, they are based on data-driven methods, such as the retrieval-based response selection [6] and the sequence-to-sequence recurrent neural networks [7]. Like most data-driven systems, they produce utterances that are incoherent or inappropriate from time to time and they require a big volume of data that may not be always available.

The combination of the two types of dialogue systems has only recently studied. Zu et al. [8] address the problems of task-oriented dialogue systems when the user's intention is not clear with a framework that incorporates non-task-oriented strategies to keep users interest in the conversation. Similarly, Papaioannou et al. [9] propose a system that combines task-oriented and chat-style dialogues. Both systems apply a reinforcement learning mechanism for selecting the appropriate strategy at each dialogue turn. Coronado et al. [10] propose a hybrid dialogue system that combines a Question Answering system with a conversational agent dealing with rest (small talk) phrases giving a social aspect to the system.

Although current works introduce social aspects through non-task-oriented strategies, we noticed that they mainly use retrieval-based methods with only exception the [11], which incorporates an extension of OwlSpeak dialogue manager [12] and decides whether to consult a knowledge-based module or react on its own. To the best of our knowledge, this is the first approach to combine knowledge-based and statistical techniques to produce task-oriented dialogues that will be used interchangeably with chatty style dialogues exploiting a rich domain model and sustaining the whole conversation memory.

III. Framework Overview

Our framework has four major components: (a) a Contextual Modelling and Representation (CMR) module, (b) a Semantic Intelligence (SI) module, (c) a Dialogue Management (DM) module and (d) a Strategy Selection (SS) module. Figure 1 shows the information flow among these components.

A user utterance is sent to the language understanding module that extracts useful information to help the CMR represent the parsed key entities and identify the discussion topic. Based on the CMR outcome, the SI updates the system's conversational picture, correlates it with background knowledge (e.g., the dialogue history) and infer unrelated insights and actions. Simultaneously, the DM accesses the discussion topic and produces topic-oriented action(s) along with a set of social-oriented actions. Finally, the SS selects among all the actions the most appropriate one and forwards it (along with relevant information from KB, if needed) to the language generation module to produce a system response.

A. Contextual Modelling and Representation

The module semantically represents and interlinks the user utterance against the system's cognitive models considering the information passed from the language understanding module.

To achieve this, the module employs existing ontologies and vocabularies. Existing ontologies form the basis of our domain model extended with application-specific aspects. Although there is a significant number of ontologies representing the domain knowledge, we found only few examples of respective ontologies for capturing the different features of the dialogue process. From these, we selected to reuse the well-established OwlSpeak ontology [12] extending it with domain-retrieved knowledge communicated within the user's utterance, exploiting the framework proposed in [4]. The dialogue turn, which is modelled by the *Move* concept, was extended with two new subclasses, the *UserMove* and the *SystemMove*, and each of them is broken down into a set of "generic" actions, which are common for both edges. For these actions, we used the list of typical actions for multi-agent dialogues presented in [13], including: Open/Greeting, Close/Goodbye, Pause, Resume, Ask, Inform, Affirm, Assert, Remind, and Alert, and extended them with "Repeat" and the "Recommend" action.

Each action is further specialised by a set of topic-oriented actions, which constitute the "discussion topics" that can be covered by the agent. Each topic might be associated with domain knowledge by means of a dialogue entity (*dialogueEntity*) which consist the target entity of each discussion topic. Additional entities extracted from the user's utterance might be associated with the dialogueEntity to further specify the requested entity.

The module semantically represents a user utterance using state-of-the-art disambiguation tools (e.g., UKB [14] or Babelfy [15]) that assign key entities extracted from the language understanding module to resource categories (i.e., synsets). These resource categories are then used to identify entities (synonyms) and topics against the domain and the dialogue ontology, respectively.

With respect to domain-driven mapping, we assume that $label(r)$, is the label of resource $r \epsilon KB$, $syn(k)$ is the synset of key entity $k \epsilon K$ and σ is a similarity function, the set $S(k)$ of all the relevant resources to k is defined as:

$$S(k) = argmax_{k \epsilon K} \sigma(k, label(r)) \qquad (1)$$

The UMBC Semantic Similarity Service [16] is used to calculate the semantic similarity σ between k and $label(r)$ combining Latent Semantic Analysis (LSA) word similarity and WordNet knowledge.

With respect to dialogue-driven mapping, a simple classification algorithm calculates the conditional probability of each discussion topic for all parsed resources, given that each discussion topic is described by means of a set of similar resources:

$$P(Topic_i \mid t_x) = \frac{P(Topic_i \cap t_x)}{P(t_x)} \qquad (2)$$

where $Topic_i$ is a topic defined by a set of resources $t_1, t_2, ...t_k$, while t_x is assumed to be a parsed resource from user utterance. This probability is then multiplied with respective probabilities for all parsed resources.

When a discussion topic is identified, the dialogue session is informed with the dialogue details including the dialogue topic, the dialogueEntity and associated entities populated with knowledge coming from the analysis of user utterance.

 ISBN: 978-1-61208-738-2

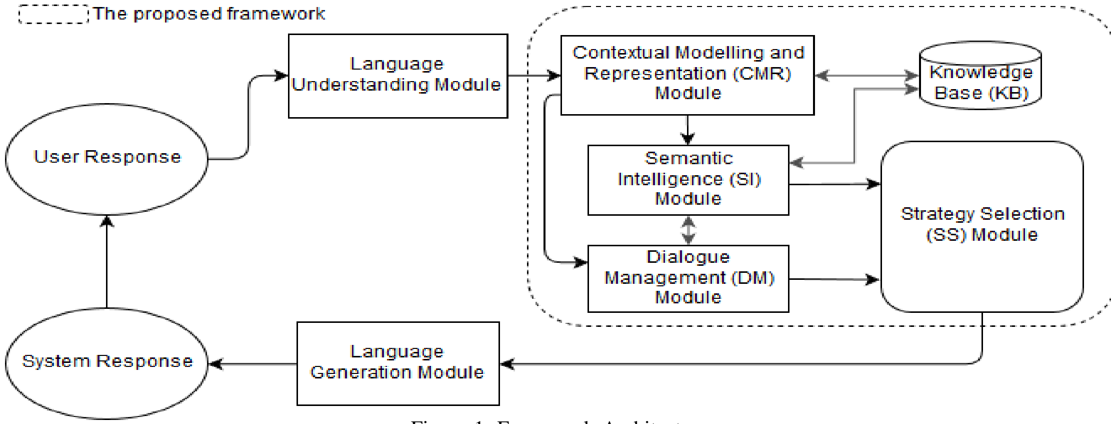

Figure 1. Framework Architecture.

B. Semantic Intelligence

The module utilises pattern-based models [17] to update domain models with new information communicated through the human-system interaction and inform the dialogue history with identified entities and topics at each dialogue turn. Moreover, it translates the system actions into actionable rules (SPARQL queries), which are then used to retrieve pertinent information from the underlying KB.

SPARQL Inferencing Notation (SPIN) are also applied to generate alerts, reminders and recommendations, which are triggered by the knowledge of the preceding discourse and the specific user profile. By this way, motivational or interventional actions are forwarded to the SS module, which might interrupt the usual flow and impose situation-oriented system responses. These actions consists of: (1) alert, (2) remind, (3) recommend and (4) repeat action.

C. Dialogue Management

The module processes the outcome of topic identification and decides the topic-oriented action to follow, selecting among: (5) predefined topic-based (re-)action, when the matching score of a topic exceeds a specific threshold, (6) clarification action, in case of partial topic identification with more than one topics receiving a significant matching score, and (7) say-again action, in case of incomplete topic identification.

Simultaneously, the module formulates a set of social-oriented action candidates considering the information received from the CMR and supportive information extracted from the KB. The social-oriented actions include (8) switch topic (a new topic is suggested based on user's preferences), (9) initiate a relevant topic, (10) end current topic and make an open question, (11) suggest to provide more info about the current topic, and (12) elicit more information.

D. Strategy Selection

This module chooses among all action candidates the most appropriate one with the aim to optimise the conversational flow towards natural and meaningful interaction. Different learning algorithms can be applied to train the strategy selection, such as Q-learning [8] and policy gradient [18]. Our strategy selection was implemented based on a simplified version of the reinforcement learning algorithm presented in [8]. The algorithm has a function that calculates the quantity of a state-action combination $Q : SxA- > R$, called Q table.

$$Q_{t+1}(s_t, a_t) \leftarrow Q_t(s_t, a_t) + a_t(s_t, a_t) \cdot \\ (R_{t+1} + \gamma max Q_t(s_{t+1}, a) Q_t(s_t, a_t))$$

For the reward function, we used domain experts' knowledge provided in [19] and [8]. According to them, the reward is calculated based on: turn index, number of times each strategy executed, sentiment polarity of previous utterances, most recently used strategy and coherence confidence of the response.

IV. A Use Case Example in Rehabilitation

As depicted in Figure 2, the system starts a conversation saying "Hello, what can I do for you?". Let us assume that the user replying "Can you tell me my workout exercises for today?".

For domain modelling, we reused COPDology [20], an ontology which was designed to facilitate the systematic monitoring of Chronic Obstructive Pulmonary Disease (COPD) patients, containing concepts pertinent to an individual's profile, the conditions they suffer from, and the medications/workout exercises they receive. We extended it with new properties, such as the *hasExecutionDay*, *hasExecutionSets* and *hasExecutionRepetitions*, to describe the execution guidelines for the scheduled workout exercises. Moreover, we assume that there is a *AskActivityForSpecificDay* topic, with the *Activity* being the target entity and the *Day* specifying the topic receiving a specific value, e.g., Monday.

The CMR annotates the key entities parsed from the language understanding module and identifies the "discussion topic". The incoming information "workout exercises" and "today" are associated with the *Activity* concept and the *Monday* instance of *Day* concept, while the *AskActivityForSpecificDay* topic is identified with a matching score of 0.8.

The SI module updates the dialogue history and enforces predefined rules. Emergency situations can be detected, for example, if the user asks more than a couple of times about the same topic, the system initially conceives it as repetition but if it happens more than a predefined amount of times (e.g., three times) the system enforces an emergency situation.

The DM evaluates the matching score of identified discussion topic and decides that a "predefined topic-based (re-)action" will be followed. This means that the *InformActivityForSpecificDay* system action, which is one-by-one associated

System response: "Hello, what can I do for you?".
User response: "Can you tell me my workout exercises for today?"
System response: "Each Monday, you have Straight Leg Raises and Glute Bridges".

System Analysis
CMR: Entities identification
"workout exercise"-> "Activity" concept
"today" -> "Monday" instance of "Day" concept
Topic identification
"AskActivityForSpecicDay", 0.8
SI: Social-oriented action
"alert", >3 repetitions
DM: Knowledge-oriented action
"predefined topic-based reaction": InformActivityForSpecicDay
(SI outcome) "dialogueEntity": "Straight Leg Raises", and "Glute Bridges".
DM: Social-oriented action
"switch topic": "Healthy diet", "initiate a relevant topic": "Diet" for "Monday"
"end current topic", "elicit more information",
"more info about the current topic": "execution sets" and "execution repetitions" (sibling properties to execution day)

Figure 2. Use case example.

with the user's action, will be enforced. In the meantime, the SI (upon DM's request) translates the system action and dialogue entities into SPARQL queries to retrieve instances of the "Activity" concept for Monday. Simultaneously, the module formulates a set of social-oriented action candidates.

Based on the learned Q table, the SS selects the most appropriate action and forwards it to language generation to produce the system response content.

V. Conclusion

The proposed framework combines dynamic knowledge-based features with social competences which are orchestrated by the means of a statistical policy learning that selects among action candidates the most appropriate one to optimise conversational effectiveness. The framework is currently validated in a running project involving clinicians and staff of a rehabilitation clinic. Our next steps is to establish an experimental set-up and evaluate it with real data. In addition, we plan to enrich the context understanding capabilities of the agent by integrating and fusing multimodal information, such as home activities and gestures, increasing the situational awareness of the agent.

Acknowledgment

This research has been cofinanced by the European Regional Development Fund of the European Union and Greek national funds through the Operational Program Competitiveness, Entrepreneurship and Innovation, under the call RESEARCH-CREATE-INNOVATE (project code: T1EDK-00686)

References

[1] D. G. Bobrow, R. M. Kaplan, M. Kay, D. A. Norman, H. Thompson, and T. Winograd, "Gus, a frame-driven dialog system," Artificial intelligence, vol. 8, no. 2, 1977, pp. 155–173.

[2] A. Rudnicky and W. Xu, "An agenda-based dialog management architecture for spoken language systems," in IEEE Automatic Speech Recognition and Understanding Workshop, vol. 13, no. 4, 1999.

[3] D. Altinok, "An ontology-based dialogue management system for banking and finance dialogue systems," CoRR, vol. abs/1804.04838, 2018. [Online]. Available: http://arxiv.org/abs/1804.04838

[4] M. Wessel, G. Acharya, J. Carpenter, and M. Yin, OntoVPA—An Ontology-Based Dialogue Management System for Virtual Personal Assistants. Cham: Springer International Publishing, 2019, pp. 219–233.

[5] F. Jurcícek, S. Keizer, M. Gasic, F. Mairesse, B. Thomson, K. Yu, and S. J. Young, "Real user evaluation of spoken dialogue systems using amazon mechanical turk," in INTERSPEECH 2011, 12th Annual Conference of the International Speech Communication Association, Florence, Italy, August 27-31, 2011, 2011, pp. 3061–3064.

[6] R. E. Banchs and H. Li, "Iris: a chat-oriented dialogue system based on the vector space model," in Proc. of the ACL 2012 System Demonstrations, 2012, pp. 37–42.

[7] O. Vinyals and Q. V. Le, "A neural conversational model," CoRR, vol. abs/1506.05869, 2015. [Online]. Available: http://arxiv.org/abs/1506.05869

[8] Z. Yu, Z. Xu, A. W. Black, and A. Rudnicky, "Strategy and policy learning for non-task-oriented conversational systems," in Proc. of the 17th annual meeting of the special interest group on discourse and dialogue, 2016, pp. 404–412.

[9] I. Papaioannou, C. Dondrup, J. Novikova, and O. Lemon, "Hybrid chat and task dialogue for more engaging hri using reinforcement learning," in 26th IEEE Int. Symposium on Robot and Human Interactive Communication (RO-MAN), 2017, pp. 593–598.

[10] M. Coronado, C. A. Iglesias, and A. Mardomingo, "A personal agents hybrid architecture for question answering featuring social dialog," in Int. Symposium on Innovations in Intelligent SysTems and Applications (INISTA), 2015, pp. 1–8.

[11] L. Pragst, J. Miehle, W. Minker, and S. Ultes, "Challenges for adaptive dialogue management in the kristina project," in Proceedings of the 1st ACM SIGCHI International Workshop on Investigating Social Interactions with Artificial Agents, ser. ISIAA 2017. New York, NY, USA: ACM, 2017, pp. 11–14. [Online]. Available: http://doi.acm.org/10.1145/3139491.3139508

[12] S. Ultes and W. Minker, "Managing adaptive spoken dialogue for intelligent environments," Journal of Ambient Intelligence and Smart Environments, vol. 6, no. 5, 2014, pp. 523–539.

[13] J. Baskar and H. Lindgren, "Human-agent dialogues and their purposes," in Proceedings of the European Conference on Cognitive Ergonomics 2017, ser. ECCE 2017. New York, NY, USA: ACM, 2017, pp. 101–104.

[14] [Online]. Available: http://ixa2.si.ehu.es/ukb/

[15] [Online]. Available: http://babelfy.org

[16] L. Han, A. L. Kashyap, T. Finin, J. Mayfield, and J. Weese, "Umbc_ebiquity-core: Semantic textual similarity systems," in 2nd Joint Conf. on Lexical and Computational Semantics (* SEM), Volume 1: Proc. of the Main Conf. and the Shared Task: Semantic Textual Similarity, 2013, pp. 44–52.

[17] G. Meditskos, S. Dasiopoulou, S. Vrochidis, L. Wanner, and I. Kompatsiaris, "Question answering over pattern-based user models," in Proc. of the 12th Int Conf. on Semantic Systems, 2016, pp. 153–160.

[18] J. Li, W. Monroe, A. Ritter, M. Galley, J. Gao, and D. Jurafsky, "Deep reinforcement learning for dialogue generation," arXiv preprint arXiv:1606.01541, 2016.

[19] F. Sukno and other, "A multimodal annotation schema for non-verbal affective analysis in the health-care domain," in Proc. of the 1st Int. Workshop on Multimedia Analysis and Retrieval for Multimodal Interaction, 2016, pp. 9–14.

[20] H. Ajami and H. Mcheick, "Ontology-based model to support ubiquitous healthcare systems for copd patients," Electronics, vol. 7, no. 12, 2018, p. 371.