



# Knowledge-Grounded Dialogue Flow Management for Social Robots and Conversational Agents

Lucrezia Grassi<sup>1</sup> · Carmine Tommaso Recchiuto<sup>1</sup> · Antonio Sgorbissa<sup>1</sup>

Accepted: 15 February 2022 / Published online: 18 March 2022  
© The Author(s) 2022

## Abstract

The article proposes a system for knowledge-based conversation designed for Social Robots and other conversational agents. The proposed system relies on an Ontology for the description of all concepts that may be relevant conversation topics, as well as their mutual relationships. The article focuses on the algorithm for Dialogue Management that selects the most appropriate conversation topic depending on the user input. Moreover, it discusses strategies to ensure a conversation flow that captures, as more coherently as possible, the user intention to drive the conversation in specific directions while avoiding purely reactive responses to what the user says. To measure the quality of the conversation, the article reports the tests performed with 100 recruited participants, comparing five conversational agents: (i) an agent addressing dialogue flow management based only on the detection of keywords in the speech, (ii) an agent based both on the detection of keywords and the Content Classification feature of Google Cloud Natural Language, (iii) an agent that picks conversation topics randomly, (iv) a human pretending to be a chatbot, and (v) one of the most famous chatbots worldwide: Replika. The subjective perception of the participants is measured both with the SASSI (Subjective Assessment of Speech System Interfaces) tool, as well as with a custom survey for measuring the subjective perception of coherence.

**Keywords** Social robotics · Conversational agents · Knowledge-grounded conversation

## 1 Introduction

Social Robotics is a research field aimed at providing robots with a brand new set of skills, specifically related to social behaviour and natural interaction with humans.

Social robots can be used in many contexts such as education [1], welcoming guests in hotels [2], cruise ships [3], malls [4], and elderly care [5]. A noteworthy application of Socially Assistive Robots (SARs) is in the healthcare field: it has been argued that robots can be used to make people feel less lonely and help human caregivers taking care of elders in care homes [6,7].

In particular, robots may help to cope with caregiver burden, i.e., the stress perceived by formal and informal caregivers, a relevant problem both in care homes and domestic environments. This subjective burden is one of the most important predictors for adverse outcomes of the care situa-

tion for the caregivers themselves, as well as for the one who requires care. Clinicians frequently overlook the caregiver burden problem [8].

Recently, social robots have been recognized as a very important resource by one of the most recent editorials of Science Robotics [9], which analyzes the potential of robots during the COVID-19 pandemic and underlines how Social Robotics is a very challenging area: social interactions require the capability of autonomously handling people's knowledge, beliefs, and emotions. In the last year, COVID-19 threatened the life of people with a higher risk for severe illness, i.e., older adults or people with certain underlying medical conditions. To slow down the spread of the virus [10], there is the need to reduce social contacts and, as a consequence, many older adults have been left even more socially isolated than before. However, in many cases, loneliness can spring adverse psychological effects such as anxiety, psychiatric disorders, depression, and decline of cognitive functions [11]. In this scenario, Social Robotics and Artificial Intelligence, in general, may have a crucial role: conversational robots and virtual agents can provide social interactions, without spreading the virus.

✉ Lucrezia Grassi  
lucrezia.grassi@edu.unige.it

<sup>1</sup> DIBRIS, University of Genoa, via all'Opera Pia 13, Genova, Italy

In many Social Robotics applications, the main focus is on the conversation. Based on its purpose, the conversation can be subdivided into:

- *task-oriented*: it is used to give commands to perform a task or retrieve information. Examples of task-oriented conversations can be found both during the interaction with a robot, i.e., “Go to the kitchen”, and during the interaction with a smart assistant, i.e., “Turn off the light”. In this case, the problem is to understand the semantics of what the user is saying, to provide the required service.
- *chit-chatting*: it has the purpose of gratifying the user in the medium to long term, keeping him/her engaged and interested. For this kind of conversation, the most relevant problem, rather than perfectly grasping the semantics of what the person says, is managing the conversation, showing the knowledge and competence on a huge number of different topics coherently and entertainingly [12].

In this scenario, the CARESSES project<sup>1</sup> is the first project having the goal of designing SARs that are culturally competent, i.e., able to understand the culture, customs, and etiquette of the person they are assisting, while autonomously reconfiguring their way of acting and speaking [13]. CARESSES exploited the humanoid robot Pepper<sup>2</sup> as the main robotic platform for interacting with people [6], in particular with older people in care homes, to make them feel less isolated and reduce the caregiver burden. Under these conditions, the capability of engaging the user to chit-chat with the system becomes of the utmost importance, as it has been shown [7] that this may have a positive impact on quality of life [16], negative attitude towards robots [17], and loneliness [18]. However, for a proper conversation to be possible, the system shall be able to talk about a huge number of topics that may be more or less relevant for different cultures, by properly managing the flow of the conversation to provide coherent replies, even when it is not able to fully understand what the user is talking about (which is likely to happen very often if the user can start discussing about any topic).

To increase the impact of the conversational systems of Social Robots, [19] argues that they shall be designed to overcome a number of limitations. Some desirable features that are not present in most systems are: (1) Breaking the “simple commands only” barrier; (2) Multiple speech acts; (3)

Mixed initiative dialogue; (4) Situated language and the symbol grounding problem; (5) Affective interaction; (6) Motor correlates and Non-Verbal Communication; (7) Purposeful speech and planning; (8) Multi-level learning; (9) Utilization of online resources and services.

The main contribution of the article is to investigate possible solutions to issue 7, which we aim to achieve by properly managing the conversation flow towards the execution of tasks or the exploration of relevant topics, thus ultimately leading to a more engaging interaction with the user. During the conversation, in many systems, it typically happens that a robot may provide answers that have nothing to do with what the user says: this has a very negative impact on the “suspension of disbelief” that is required to give the user the impression of genuine intelligence, and ultimately it generates frustration. A key element to avoid this is the capability to know when to further explore the current topic or choose the next conversation topic, coherently with the user sentence: if the algorithm does not pick a topic coherently, the agent’s reply will not be appropriate, even if the system had the knowledge to interact consistently.

Specifically, we proceed as follows.

First, we propose a novel system for knowledge-based conversation and Dialogue Management that relies on an Ontology for the description of all relevant concepts that may play a key role in the conversation: the Ontology is designed to take into account the possible cultural differences between different users in a no-stereotyped way, and it stores chunks of sentences that can be composed in run-time, therefore, enabling the system to talk about the aforementioned concepts in a culture-aware and engaging way [14,20].

Second, we test the developed solutions by comparing five Artificial Conversational Agents during a conversation with 100 recruited participants: (i) an agent addressing dialogue flow management based only on the detection of keywords, (ii) an agent addressing dialogue flow management based both on the detection of keywords and the Content Classification feature of Google Cloud Natural Language, (iii) an agent that picks random topics among those present in the Ontology, (iv) a human pretending to be a chatbot, and (v) one of the most famous chatbots worldwide: Replika. The subjective perception of the participants is measured both with the SASSI (Subjective Assessment of Speech System Interfaces) tool [21,22] and a custom survey for measuring the subjective perception of coherence in the conversation.

The remainder of this article is organized as follows. Section 2 presents an overview of previous works related to Knowledge-Grounded Conversation, the most popular validated tools to evaluate the User Experience, and introduces the concept of “coherence” in Dialogue Management. Eventually, it presents a typical classification for Artificial Conversational Agents and describes the up-to-date most famous chatbots. An overview of the CARESSES

<sup>1</sup> <http://caressrobot.org/>.

<sup>2</sup> CARESSES architecture has been recently revised to allow for a Cloud-based implementation [14], thus enabling virtually any device equipped with a network interface to perform long-term, culture-aware conversation with the user. Currently, the robots Pepper and NAO, the pill-dispenser Pillo, Prof. Einstein, and a virtual character implemented as an Android application have been provided with the onboard functionalities to implement the instructions received by the CARESSES Cloud [15].

knowledge-based conversational system is given in Sect. 3. Section 4 describes the experiment carried on to evaluate the user satisfaction when interacting with the agents, along with the statistical tools used to analyse the collected data. Section 5 presents the results obtained, discussed in detail in Sect. 6. Eventually, Section 7 draws the conclusions.

## 2 State of the Art

### 2.1 Knowledge-Grounded Conversation

A knowledge-grounded conversational system is a dialogue system that can communicate by recalling internal and external knowledge, similarly to how humans do, typically to increase the engagement of the user during chit-chatting. The internal knowledge is composed of things that the system already knows, while external knowledge refers to the knowledge acquired in run-time. To this end, knowledge-grounded systems must not only understand what the user says but also produce their response based on the available internal and external knowledge. There are many differences between the replies of a normal conversational system and those of a knowledge-grounded system. For example, if the user says “I love pizza”, the former might provide a general answer such as “I see, very interesting” as it has no previous knowledge about `Pizza`. On the other side, the latter may be more specific, i.e., “Oh, pizza is a very delicious Italian food”, if it had the chance to recall internal or external knowledge about pizza, e.g., using resources on the web.

Building a knowledge-grounded conversational system raises many challenges. For instance, the internal knowledge may be static or “expandable”, i.e., updated in run-time with new external information retrieved from the user utterances, which will be considered as internal knowledge in the next interactions. The external knowledge in most cases will come from websites: however, even if the system can find online resources associated with the current conversation topic by using a state-of-the-art information retrieval system, it may be difficult to extract knowledge from the search results because this would typically require complex Natural Language Processing (NLP) techniques to select the appropriate knowledge. This is done in [23] using a Generative Transformer Memory Network and in [24] using an SVM classifier, to name a few.

Another complex task is to generate appropriate responses that reflect the acquired knowledge as a consequence of the conversation history. The problems of knowledge extraction and response generation, among the others, are addressed in [25], where a knowledge-grounded multi-turn chatbot model is proposed (see also 2.3 for popular chatbots based on this principle).

A data-driven and knowledge-grounded conversation model, where both conversation history and relevant facts are fed into a neural architecture that features distinct encoders for the two entries, is proposed in [26]. Such model aims to produce more appropriate responses: the article shows that the outputs generated by a model trained with sentences and facts related to the conversation history were evaluated by human judges as remarkably more informative (i.e., knowledgeable, helpful, specific) with respect to those of a model trained with no facts.

In [27], a large corpus of status-response pairs found on Twitter has been employed to develop a system based on phrase-based Statistical Machine Translation, able to respond to Twitter status posts. As stated in the article, data-driven response generation will provide an important breakthrough in the conversational capabilities of a system. Authors claim that when this kind of approach is used inside a broad dialogue system and is combined with the dialogue state, it can generate locally coherent, purposeful, and more natural dialogue.

Finally, the problem of knowledge representation and knowledge-based chit-chatting, keeping into account the cultural identity of the person, is addressed in [20] and [14] using an Ontology designed by cultural experts coupled with a Bayesian network, to avoid rigid representations of cultures that may lead to stereotypes. Such an approach allows the system to have more control over what the robot says during the conversation: the idea of having a knowledge base designed by experts is particularly suitable for sensitive situations, i.e. when dealing with the more fragile population such as older adults or children. Approaches that automatically acquire knowledge from the Internet may not be optimal in these scenarios.

### 2.2 User Experience

In the Literature, tools exist that can be used to evaluate the User Experience (UX) and the overall quality of the conversation. The work carried on in [28] reviews the six main questionnaires for evaluating conversational systems: AttrakDiff, SASSI, SUISSQ, MOS-X, PARADISE, and SUS. Moreover, it assesses the potential suitability of these questionnaires to measure various UX dimensions.

As a measure of the quality of the flow of conversation, we aim to evaluate how “coherent” the system is: this shall be done not only taking into account the last user utterance but also the current conversation topic (i.e., what is referred to as context in [29,30]). To clarify what this means, suppose a conversational system that replies very accurately to everything the user says: if the user talks about football, the system will reply by talking about football. If the user talks about apples, the system will reply by talking about apples. But what if the person and the system are talking about

football, and the person says something that is not recognized as related to anything specific, such as “I will think about that”, “I think so”, “It’s so nice to talk with you”, and so on? A coherent Dialogue Management system shall be able to understand when it is more appropriate to further explore the current conversation topic (in this example, by taking the initiative to ask the person a question about his/her preferred team or players) or lead the dialogue to another topic, coherently with what the user said.

Unfortunately, the aforementioned UX tools are not meant to evaluate topic coherence in a mixed-initiative dialogue, as in our case. The most similar measure to what we want to evaluate is the *System Response Accuracy* scale of the SASSI questionnaire: such a questionnaire, described more in detail in Sect. 4.2.2, has been used during our experiments. The *System Response Accuracy* is defined as the system’s ability to correctly recognise the speech input, correctly interpret the meaning of the utterance, and then act appropriately. However, according to the aforementioned definition, what the System Response Accuracy scale aims to measure is not equivalent to the broader concept of coherence that has been given: it measures the quality of the system’s reply to the user input in a purely reactive fashion, without relying on the concept of “current topic of conversation”.

For this reason, a coherence measure in the spirit of [31] and [32] has been used to supplement the SASSI questionnaire, which requires users to rate individual sentences pronounced by the robot in light of the topic/context (details are given in Sect. 4.2.1). Notice that [32] and [33] also propose methods to automatically evaluate the coherence, as this is considered an important metric to evaluate multi-turn conversation in open domains: however, in this work we are only interested in evaluating the subjective perception of users, which makes a simple rating mechanism perfectly fitting our purposes.

### 2.3 Artificial Conversational Agents

Depending on the emphasis they put on a task-oriented conversation or chit-chatting, a typical classification for artificial conversational agents is based on their scope:

- *Question answering bots*: knowledge-based conversational systems that answer to users queries by analysing the underlying information collected from various sources like Wikipedia, DailyMail, Allen AI science and Quiz Bowl [34];
- *Task-oriented bots*: conversational systems that assist in achieving a particular task or attempt to solve a specific problem such as a flight booking or hotel reservation [35];
- *Social bots*: conversational systems that communicate with users as companions, and possibly entertain or give

recommendations to them [36]: recent notable examples are Microsoft Xiaoice [37] and Replika<sup>3</sup>.

In the following, we discuss more in detail the agents belonging to the third class pointing out differences and similarities with our solution, without making a distinction between robots and chatbots, unless strictly required.

For many decades, the development of social bots, or intelligent dialogue systems that can engage in empathetic conversations with humans, has been one of the main goals of Artificial Intelligence. As stated in [38], early conversational systems such as Eliza [39], Parry [40], and Alice [41] have been designed to mimic human behaviour in a text-based conversation, hence to pass the Turing Test [42] within a controlled scope. Despite their impressive successes, these systems, which were precursors to today’s social chatbots, worked well only in constrained environments. In more recent times, among the most successful conversational agents for general use that can act as digital friends and entertainers, Xiaoice, Replika, Mitsuku, and Insomnobot-3000 are the most frequently mentioned [43].

*XiaoIce* (“Little Ice” in Chinese) is one of the most popular chatbots in the world. It is available in 5 countries (i.e., China, Japan, US, India, and Indonesia) under different names (e.g., Rinna in Japan) on more than 40 platforms, including WeChat, QQ, Weibo, and Meipai in China, Facebook Messenger in the United States and India, and LINE in Japan and Indonesia. Its primary goal is to be an AI companion with which users form long-term emotional connections: this distinguishes XiaoIce not only from early chatbots but also from other recently developed conversational AI personal assistants such as Apple Siri, Amazon Alexa, Google Assistant, and Microsoft Cortana.

As stated in [37], the topic database of XiaoIce is periodically updated by collecting popular topics and related comments and discussions from high-quality Internet forums, such as Instagram in the US and the website douban.com in China. To generate responses, XiaoIce has a paired database that consists of query-response pairs collected from two data sources: human conversational data from the Internet, (e.g., social networks, public forums, news comments, etc.), and human-machine conversations generated by XiaoIce and her users. Even if the data collected from the Internet are subjected to quality control to remove personally identifiable information (PII), messy code, inappropriate content, spelling mistakes, etc., the knowledge acquisition process is automated and it is not supervised by humans: differently from our system, this solution may not be suitable in sensitive situations, e.g., when dealing with the more fragile population such as older people or children. Regarding issue

<sup>3</sup> <https://help.replika.ai/hc/en-us/articles/115001070951-What-is-Replika->

7 in Sect. 1 (i.e., the lack of flow in the conversation), the implementation of XiaoIce addressed this problem by using a Topic Manager that mimics human behaviour of changing topics during a conversation. It consists of a classifier for deciding at each dialogue turn whether or not to switch topics and a topic recommendation engine for suggesting a new topic. Topic switching is triggered if XiaoIce does not have sufficient knowledge about the topic to engage in a meaningful conversation, an approach that our system adopts as well by relying on an Ontology for representing relationships among different topics of conversation. Unfortunately, we could not try XiaoIce as it is not available in Italy.

*Replika* is presented as a messaging app where users answer questions to build a digital library of information about themselves. Its creator, a San Francisco-based startup called Luka, sees a whole bunch of possible uses for it: a digital twin to serve as a companion for the lonely, a living memorial of the dead, created for those left behind, or even, one day, a version of ourselves that can carry out all the mundane tasks that we humans have to do, but never want to.

To the best of the authors' knowledge, there is no detailed explanation of the implementation of *Replika*, hence we can only make assumptions on how issues related to Dialogue Management are faced. After having intensively tried *Replika*, it seems that the conversation is partitioned into "sessions" in which it has specific competencies, somehow playing a similar role as topics in XiaoIce and our system. Concerning issue 7 in Sect. 1, the chatbot uses a Neural Network to hold an ongoing, one-on-one conversation with its users, and over time, learn how to talk back [44]. The agent is trained on texts from more than 8 million web pages, from Twitter posts to Reddit forums, and it can respond in a thoughtful and human-like way. From time to time, it will push the user to have a "session" together: in this session, it will ask questions regarding what the user did during the day, what was the best part of the day, what is the person looking forward to tomorrow, and eventually, the user has to rate his/her mood on a scale of 1 to 10. During this session, *Replika* takes control of the conversation and insists that the user answers its questions about a specific matter, something that may be annoying and frustrating, and we avoid by letting the user free to easily switch to another topic at any time. *Replika*'s responses are backed by Open AI's GPT-2<sup>4</sup> text-generating AI system [45].

*Mitsuku*<sup>5</sup>, or Kuki as her close friends call her, is a chatbot created by Pandorabots<sup>6</sup>: an open-source chatbot framework that allows people to build and publish AI-powered chatbots

on the web, mobile applications, and messaging apps like LINE, Slack, WhatsApp, and Telegram.

The Pandorabots chatbot framework is based on the Artificial Intelligence Markup Language (AIML) scripting language, which developers can use to create conversational bots. The downside is that Pandorabots do not include machine learning tools that are common on other chatbot building platforms. A specific AIML file allows users to teach Mitsuku new facts: the user should say "Learn" followed by the fact (i.e., "Learn the sun is hot"). The taught information is emailed to its creator Steve Worswick, which will personally supervise the learning process. To address issue 7 in Sect. 1, specific areas of competence of the chatbot are managed by different AIML files, which are responsible for maintaining coherence in the responses: AIML files, therefore, play a similar role as topics/concepts in the Ontology that are the core elements of our system, but without the benefit of having a hierarchical structure as a key element for Dialogue Management. Mitsuku is a five-time Loebner Prize winner (in 2013, 2016, 2017, 2018, 2019), and it is available in a web version<sup>7</sup>, Facebook Messenger, Kik Messenger, and Telegram.

*Insomnobot-3000*<sup>8</sup> is the world's first bot that is only available to chat, exclusively via SMS, between 11 pm and 5 am regardless of the time zone. The bot was built by the mattress company Casper and, according to its creators, it was programmed to sound like a real person and talk about almost anything. *Insomnobot-3000* cannot learn new things and expand its knowledge base: however, it can generate over 2,000 different responses, depending on which category and emotion the keyword falls under. Concerning issue 7 in Sect. 1, when the bot receives a text, it chooses an appropriate response by identifying keywords: however, differently from the aforementioned systems and our solution, it lacks a more sophisticated mechanism for Dialogue Management to switch in-between different areas of competence or topics to provide contextual replies.

In addition to the aforementioned chatbot available to the large public, recent research on data-driven knowledge-grounded conversation based on sophisticated generative models is worth being mentioned.

*Meena* [29] is a generative chatbot model trained end-to-end on 40B words mined from public domain social media conversations, that addresses the problem of multi-turn conversation in open domains. According to its authors, this chatbot can conduct conversations that are more sensible and specific than existing state-of-the-art chatbots. Such improvements are measured by a new human evaluation metric, called Sensibleness and Specificity Average (SSA), which captures important attributes for human conversation.

<sup>4</sup> <https://en.wikipedia.org/wiki/GPT-2>.

<sup>5</sup> <https://www.pandorabots.com/mitsuku/>.

<sup>6</sup> <https://home.pandorabots.com/home.html>.

<sup>7</sup> <https://chat.kuki.ai>.

<sup>8</sup> <https://insomnobot3000.com>.

Meena is based on the concept of multi-turn contexts to evaluate the quality of a response: not only the last sentences but the recent conversation history in terms of (*context, response*) pairs is used to train and evaluate the network. Experiments show significant improvements in SSA with respect to competitors, providing replies that are coherent with what the user said and, at the same time, more specific to the context. Authors also report that issues related to inappropriate and biased language still need to be addressed: to the best of our knowledge, a publicly available version of this chatbot has not been released yet.

*BlenderBot* [30] is a very recent open-domain chatbot developed at Facebook AI Research. Like what we do, the authors put a stronger emphasis on aspects that are not only related to sentence generation: they argue that good conversations require a number of skills to be blended seamlessly, including providing engaging topics of conversation as well as listening and showing interest to what the user says, among the others. Current generative models, they argue, tend to produce dull and repetitive responses (and, we would add, do not take cultural appropriateness into account), whereas retrieval models may produce human written utterances that tend to include more vibrant (and culturally appropriate) language. To overcome the limitations of purely generative models, BlenderBot includes a retrieval step before generation according to a retrieve and refine model [46], which retrieves initial dialogue utterances and/or knowledge from a large knowledge base. As like Meena, the system is trained to choose the next dialogue utterance given the recent dialogue history, referred to as context: training is performed using a number of datasets that are key to exhibit different skills (i.e., engaging personality, emotional talking, knowledge-grounded conversation), as well as a Blended Skill Talk dataset that merges utterances taken from the three. Authors claim that their model outperforms Meena in human evaluations, however, they acknowledge that the system still suffers from a lack of in-depth knowledge if sufficiently interrogated, a tendency to stick to simpler language, and a tendency to repeat often used phrases.

DialoGPT [47] is a neural conversational response generation model trained on 147M conversation-like exchanges extracted from Reddit comment chains over a period spanning from 2005 through 2017 and it extends GPT-2. Authors claim that conversational systems that leverage DialoGPT generate more relevant, contentful, and context-consistent responses than strong baseline systems. As it is fully open-sourced and easy to deploy, users can extend the pre-trained conversational system to bootstrap training using various datasets. Authors claim that the detection and control of toxic output will be a major focus of a future investigation.

All generative models based on human-human data offer many advantages in producing human-like utterances fitting the context but have the major drawback that they can learn

undesirable features leading to toxic or biased language. Classifiers to filter out inappropriate language exist [48], but they still have limitations: this issue is particularly important when dealing with the more frail populations, especially by considering that classifying language as inappropriate or offensive typically depends on cultural factors, and the development of classifiers working properly in multiple cultures may be very challenging. Since we claim that a proper flow of conversation deserves higher attention than the automatic generation of sentences, we address these problems through an Ontology of conversation topics and topic-related chunks of sentences generated by cultural experts, that are then composed in run-time using the hierarchical structure of the knowledge-base to produce a huge variety of sentences. Finally, even if this is not the main objective of this article, it is worth reminding that our system may take cultural aspects into account in Dialogue Management, an element that is completely ignored by all state-of-the-art systems.

### 3 System Architecture

Even if the focus of this article is on knowledge-based chit-chatting and dialogue flow management, and not on cultural adaptation, it is necessary to briefly introduce the structure of the cultural knowledge base used by the CARESSES conversational system since it strongly influences the algorithms we implemented.

#### 3.1 Knowledge Representation

In CARESSES, the ability of the companion robot to naturally converse with the user has been achieved by creating a framework for cultural knowledge representation that relies on an Ontology [49] implemented in OWL2 [50]. According to the Description Logics formalism, concepts (i.e., topics of conversation the system is capable of talking about) and their mutual relations are stored in the terminological box (TBox) of the Ontology. Instead, instances of concepts and their associated data (e.g., chunks of sentences automatically composed to enable the system to talk about the corresponding topics) are stored in the assertional box (ABox).

To deal with representations of the world that may vary across different cultures [51], the Ontology is organized into three layers, as shown in Fig. 1. The TBox (layer I) encodes concepts at a generic, culture-agnostic level, which can be inherited from existing upper and domain-specific Ontologies (grey boxes) or explicitly defined to enable culture-aware conversation with the user (white boxes). An important point to be highlighted is that the TBox should include concepts that are typical of the union all cultures considered, whatever the cultural identity of the user is, to avoid stereotypes (an example related to different kinds of beverages is shown in

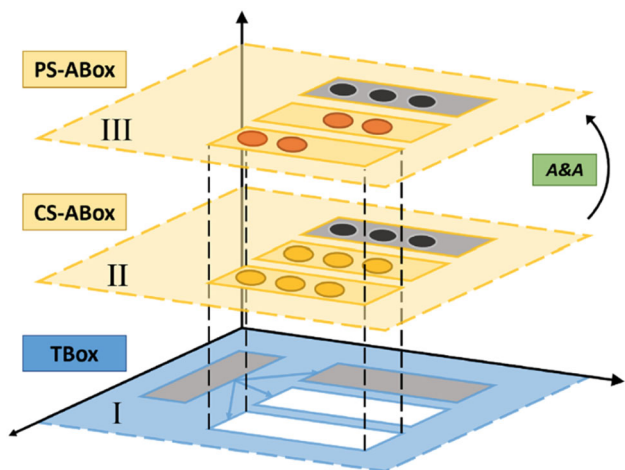


Fig. 1 Knowledge representation architecture for a culturally competent robot

Fig. 2). The system will initially guess the user beliefs, values, habits, customs, preferences based on the culture they declared to self-identify with, but it will then be open to considering choices that may be less likely in a given culture, as the user explicitly declares their attitude towards them. According to this principle, the system may initially infer that an English person may be more interested to talk about Tea rather than Coffee, and the opposite may be initially inferred for an Italian user. However, during the conversation, initial assumptions may be revised, thus finally leading to a fully personalized representation of the user attitude towards all concepts in the TBox, to be used for conversation.

To implement this mechanism, the Culture-Specific ABox layer comprises instances of concepts (with prefix EN\_ for “English” in Fig. 2) encoding culturally appropriate chunks of sentences to be automatically composed (Data Property hasSentence) and the probability that the user would have a positive attitude toward that concept, given that they belong to that cultural group (Data Property hasLikeliness).

Eventually, the Person-Specific ABox comprises instances of concepts (with prefix DS\_ for a user called “Dorothy Smith” in Fig. 2) encoding the actual user attitude towards a concept (Mrs Dorothy Smith may be more familiar with having tea than the average English person, hasLikeliness =Very High), sentences to talk about a topic explicitly taught by the user to the system (hasSentence=“You can never get a cup of tea large enough or a book long enough to suit me”) or other knowledge explicitly added during setup (e.g., the user name and the town of residence). At the first encounter between the robot and a user, many instances of the Ontology will not contain Person-Specific knowledge: the robot will acquire this awareness at run-time either from its perceptual system or during the interaction with the user, e.g., asking questions. Figure 1 also shows that some instances of existing Ontologies that are not culture- or person-dependent

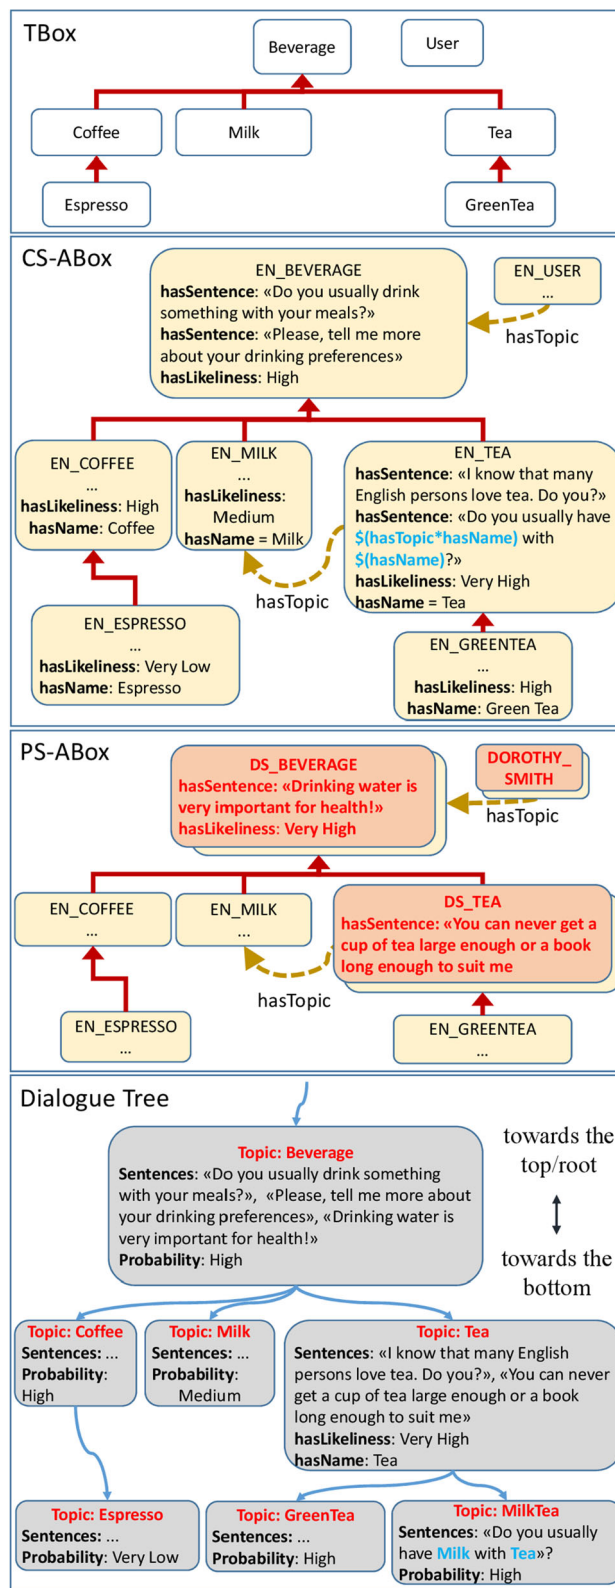


Fig. 2 The three layers of the Ontology: TBox, CS-ABox (for the English culture), PS-ABox (for the user Dorothy Smith), and the Dialogue Tree generated from the Ontology structure

(dark circles) may not change between the two ABox layers: this may refer, for instance, to technical data such as the serial number of the robot or physical quantities, if they need to be encoded in the Ontology. For a detailed description of the terminological box (TBox) and assertional box (ABox) of the Ontology, as well as the algorithms for cultural adaptation, see [14,20]

The Dialogue Tree (DT) (Fig. 2), used by the chit-chatting system (Section 3.2), is built starting from the Ontology structure: each concept of the TBox and the corresponding instances of the ABox are mapped into a conversation topic.

From here on, we define a topic as a node of the DT. The relation between topics is borrowed from the structure of the Ontology: specifically, the Object Property `hasTopic`, its sub-properties, and the hierarchical relationships among concepts and instances are analyzed to define the nodes and the branches of the DT. In the example of Figure 2, the instance of `Tea` for the English culture is connected in the DT to its child node `GreenTea` (which is a subclass of `Tea` in the Ontology), and its sibling `MilkTea` (since `EN_MILK` is a filler of `EN_TEA` for the Object Property `hasTopic`). In the DT the nodes corresponding to `Tea`, `GreenTea`, and `MilkTea` are all referred to as topics.

As mentioned, each conversation topic has chunks of culturally appropriate sentences associated with it that are automatically composed and used during the conversation. Such sentences can be of different types (i.e., positive assertions, negative assertions, different kinds of questions, or proposals for activities). They are selected in subsequent iterations, depending on internal rules and the input received from the user (until the topic is changed, see Sect. 3.2). Typically, when exploring the same topic in the DT, the system may first ask a question to understand if the user is familiar with that topic: if so, this may be followed by general considerations, proposals for activities, or open questions, until the system moves to the next topic if the conditions hold. Multiple sentences of the same type are present in the DT, randomly chosen to reduce the chance of repetitions. At any time, the user is free to take the initiative to express their considerations or lead the conversation to another topic.

Sentences may contain variables that are instantiated when creating the DT. For instance, a hypothetical sentence “Do you like \$hasName?”, encoded in the concept `Coffee`, might be used to automatically produce both “Do you like Coffee?” and “Do you like Espresso?” in the DT. That is, the variable `$hasName` automatically takes the values “Coffee” and “Espresso” in sentences associated with different topics, being `Espresso` a subclass of `Coffee`. Similarly, the sentence “I love \$hasName with \$hasActor\*hasName”, encoded in the concept `Movie`, can be used to produce several sentences in the DT, depending on `Movie`’s subconcepts and the relations with other

concepts. In this second case, when creating the DT, the variable `$hasName` is automatically instantiated depending on the subconcepts of `Movie`, whereas the variable `$hasActor*hasName` automatically takes the value from the concepts related to `Movie` along its `hasActor` property (where `hasActor` is a sub-property of `hasTopic` in the Ontology, relating movies to actors). Then, in the higher-level topics of the DT, the sentence would sound like “I love movies with great actors”. In the lower-level topics, the variables `$hasName` and `$hasTopic*hasName` will be instantiated with more specific types of movies and actors, producing sentences such as “You know... I love Bollywood movies with Amitabh Bachchan”. In the current version used for testing, exploiting variables and the taxonomy of the Ontology allowed us to easily produce a DT with 2,470 conversation topics and 24,033 sentences, with random variations made in run-time.

Since this concept is key for automatic sentence generation through variable instantiation and composition, a few additional words are worth spending. Currently, as in the aforementioned example, one sentence may contain variables whose value depends on (i) the specific concept that inherited the sentence from superconcepts in the Ontology, and (ii) related concepts through an Object Property. Then, when manually adding a sentence as a Data Property of a given concept, a number of sentences will be inherited and added automatically depending on the number of (i) its subconcepts and (ii) related concepts through an Object Property. By considering a DT generated from the Ontology and having a maximum branching factor  $B = B_s + B_p$ , where  $B_s$  is the branching factor due to subconcepts and  $B_p$  is the branching factor due to related concepts, adding a sentence in a concept at a height  $H$  from DT leaves produces a number of sentences equal to  $\sum_{i=0}^H B^i$ . Considering that all sentences are randomly associated in run-time with a prefix (e.g., “You know...”, “I heard that...”), and may be appended to each other, the number of possible variations increases even more. For instance, in the current configuration, the DT has  $H = 5$  and  $B = 34$ , which would ideally correspond to 46 million nodes and sentences in case of a complete tree. However, since the tree is not complete (the average branching factor is much lower than 34), in the current DT there are only 2,470 nodes, and adding a sentence at the root of the DT will produce a maximum of 2,470 inherited sentences (i.e., one for each node). For each topic, variables in sentences are properly instantiated as already explained. Such sentences may however lack specificity, as it is unlikely that the same pattern may be easily adapted to talk about everything in the Ontology: i.e., “I love talking about \$hasName with you” may work for all concepts/topics, but “I like eating \$hasName” will not.



### 3.2 Chit-Chatting

When the user pronounces a sentence that triggers the start of a conversation, the system switches to the *chit-chatting* state. In this case, the Dialogue Management algorithm operates to keep the conversation as engaging and natural as possible by implementing two different, intertwined mechanisms to navigate the DT based on what the user says:

1. *Navigating the DT according to the topic hierarchy*: the first mechanism implemented by the Dialogue Management algorithm simply chooses the topic based on the DT, following the branches from more general conversation topics close to the root (e.g., the topic `BreakfastHabits`) to more specific (e.g., the topic `HavingTeaForBreakfast`) or related ones (e.g., the topic `TakingCareOfOneself`). This is achieved by selecting the DT branches to maximise the probability that the user will be interested in the next topic (also based on cultural factors). When in a topic, the system asks questions, replies with positive or negative assertions, proposes activities or simply encourages the user to freely express his/her view, until the topic has been sufficiently explored. Probabilities associated with topics (`hasLikeliness` property, see Sect. 3.1) have been initially assigned with the help of Transcultural Nursing experts as well as looking for information on the Internet about habits, foods, sports, religions, etc. in different cultures. Please, notice however that these are only initial guesses, as the system will update probabilities during the conversation depending on the user preferences and attitudes to avoid stereotyped representations.
2. *Navigating the DT jumping between topics*: the second mechanism enables the Dialogue Management algorithm to jump to a different topic of the DT depending on what the user says (e.g., when talking about `BreakfastHabits` the person may start talking about a `Restaurant` they love). After doing so, the system reverts to the first mechanism but starting from the new topic: it asks questions, makes positive and negative comments, proposes activities, or allows the person to freely express their view about the new topic, and then re-start following branches as previously.

With reference to Figure 2, a possible dialogue pattern between the robot and Mrs Dorothy Smith, using the aforementioned mechanisms, would be the following (labels 1 and 2 identify the mechanism currently adopted by the system to navigate the DT):

*R. (1): Do you usually drink something with your meals?*  
*Mrs S.: Always! At least I try...*

*R. (1): Drinking water is very important for health! Please, tell me more about your drinking preferences.*

*Mrs S.: Well, I mostly drink water during lunch.*

*R. (1): I know that many English persons love tea. Do you?*

*Mrs S.: Yes, of course! But mostly in the afternoon.*

*R. (1): You can never get a cup of tea large enough or a book long enough to suit me. Do you usually have milk with your tea?*

*Mrs S.: Yes, and scones, too! But they are not very healthy.*

*R. (2): I love scones with jam and whipped cream!*

Figure 3 depicts a detailed flow-chart that clarifies what happens during the interaction with the user.

The diagram shows that, as soon as the person says something that is not a request to stop the execution (box `Wait For user input`), the Dialogue Management algorithm checks if the user sentence contains some command. We define a command as a sentence pronounced by the user aimed to trigger an activity that is not exclusively related to the conversation. Following the same rationale, we define an activity as any task that the robot might be requested to perform, not only related to the conversation, possibly using third-party services. For instance, the user might issue

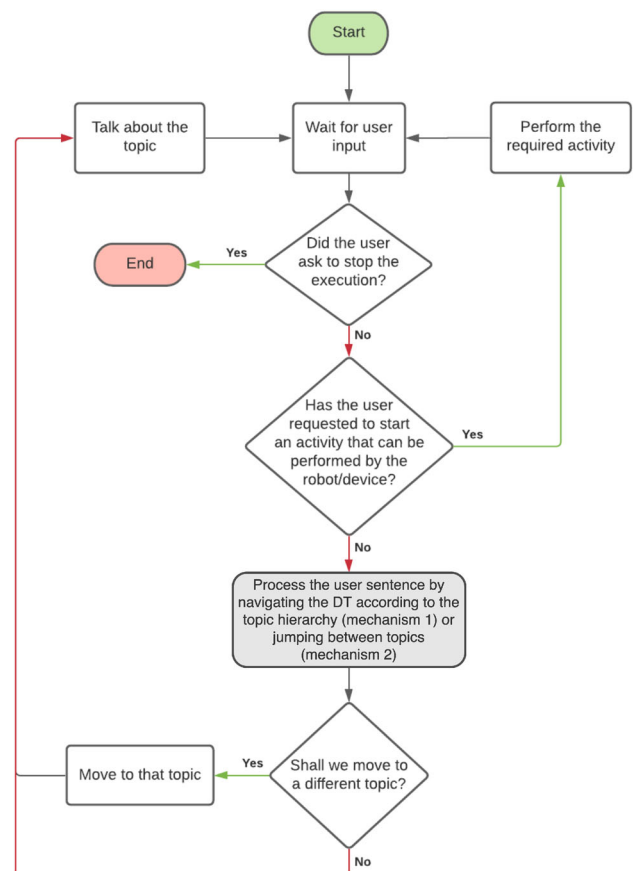


Fig. 3 Flow chart of the basic chit-chatting framework

the command “Show me the weather” to start the corresponding activity that acquires information about the weather (box `Perform the required activity`). Even if described for completeness, this part of the interaction is not the focus of this work. It has been implemented in our experiments with caring robots during the CARESSES project, but not in the tests described in this article, and it may rely on a third-party system for NLP<sup>9</sup> or a pre-existing command interface that is already implemented onboard the robot (e.g., verbal commands to clean the room in case of a robotic vacuum cleaner or to deliver medicines in case of a robotic pill dispenser).

If a request to start an activity is not found, the user input is processed (box `Process the user sentence by navigating the DT...`) by using the two mechanisms described so far: as a first step the user utterance is analysed to check if something relevant is detected, and possibly used to jump to another topic using mechanism 2; if no information is found to jump to another topic, mechanism 1 checks if the current topic shall be further explored or it is time to follow the branches of the DT from a parent node to a child node that is semantically related in the Ontology. Eventually, it should be mentioned that activities (e.g., vacuum cleaning or pill delivery) may even be proposed by the Dialogue Management System itself when related to the current topic of discussion (i.e., if the person is talking about cleaning and the robot has this capability). The interaction continues in this way until the user explicitly asks to stop it.

This approach to Dialogue Management, despite its simplicity, reveals to be very effective for chit-chatting. It should be reminded that here Dialogue Management is not meant to understand all the nuances of human language when giving a specific command to be executed (i.e., “Put the red ball on the green table”), but rather it aims at engaging the user in rewarding conversations: an objective that we try to achieve by enabling the system to coherently talk about a huge number of different topics while properly managing the conversation flow. Please notice that the approach has been already tested in the CARESSES project with more than 20 older people of different nationalities, for about 18 hours of conversation split into 6 sessions. Both the quantitative and qualitative analysis performed revealed very positive feedback from the participants [6,7]. However, the quantitative analysis performed with care home residents did not explicitly evaluate the quality of the conversation: the analysis performed was rather aimed to measure improvement in health-related quality of life (SF-36 [16]), negative attitude towards robots (NARS [17]), loneliness (ULS-8 [18]). Then, additional analysis with a wider population is needed to eval-

uate the quality of the conversation produced by our system in terms of issue 7 in Sect. 1 and motivates the present study.

Please notice that, when navigating in the knowledge base following the branches of the DT (aforementioned mechanism 1), coherence in the flow of conversation is straightforwardly preserved by the fact that two nodes of the DT are semantically close to each other by construction. However, to preserve coherence when deciding if it is needed to jump from one node to another (mechanism 2), proper strategies should be implemented to avoid the feeling that the system is “going off-topic” with respect to what the person says: on the one side, the system needs to be enough responsive to promptly move to another topic if the person wants to; on the other side, the system needs to avoid jumping from one topic to another, overestimating the desire of the person to talk about something else when they would be happier to further explore the current topic of conversation.

### 3.3 Jumping to a Different Discussion Topic

When using mechanism 2, described in the previous section, to navigate the DT, to preserve coherence in the conversation flow the system uses one of the two following jumping methods: keyword-based topic matching and keyword- and category-based topic matching. These methods differ in terms of complexity and in terms of the need to rely on third-party services to analyze the semantic content of sentences.

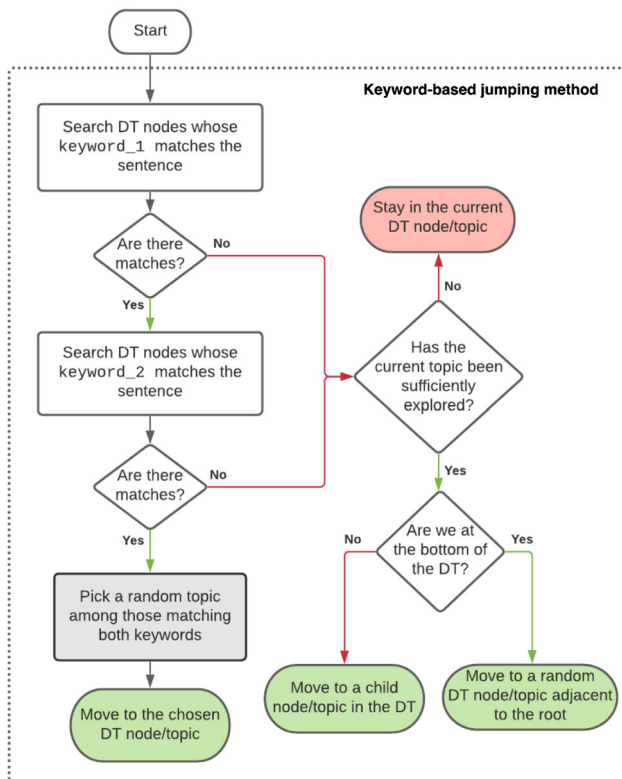
To avoid confusion, please notice that we use the word “mechanism” to refer to the two ways to navigate the DT (Mechanism 1: Navigating the DT according to the topic hierarchy; Mechanism 2: Navigating the DT jumping between topics), whereas we use the term “jumping methods” to refer to the two different ways to jump between topics when using mechanism 2.

#### 3.3.1 Keyword-Based Topic Matching

The first, and the simpler, jumping method is exclusively based on the detection of keywords in the user sentence (keywords are manually encoded in the Ontology in a corresponding Data Property). To match a topic, at least two keywords associated with that topic should be detected in the sentence pronounced by the user, using wildcards to enable more versatility in keyword matching. The use of multiple keywords allows the system to differentiate between semantically close topics (i.e., `Green Tea` rather than the more general concept of `Tea`).

Figure 4 shows a possible implementation of the box `Process user sentence` with Dialogue Management mechanism 1 or 2 in Fig. 3 when using the keyword-based topic matching jumping method:

<sup>9</sup> For instance, <https://dialogflow.cloud.google.com/>, that in the current implementation also manages small-talk requests triggered by sentences such as “Hello”, “How are you”, or similar.



**Fig. 4** Flow chart of the keyword-based topic matching and jumping method

- If the user sentence contains the two keywords corresponding to a topic in the Ontology (left part of the diagram), the algorithm returns a topic of the DT that matches those keywords. In case more topic matches are found, the algorithm randomly chooses one of the matching topics.
- Otherwise:
  - If there are still relevant questions to be asked or assertions to be made about that topic, the system stays on the same topic.
  - Otherwise:
    - If not at the bottom of the DT, the Dialogue Management System keeps on exploring the DT along its branches;
    - Otherwise the algorithm returns a random topic immediately below the root.

This basic approach, despite its simplicity, has been successfully exploited during the experimental trial in care homes, and its capability to provide (almost always) coherent, knowledge-grounded replies, was confirmed in public exhibitions. However, when focusing on coherence, the approach has obvious limitations. Let’s suppose that the user is having a conversation with the system and at some point, for some reason, the user says “My bank account has

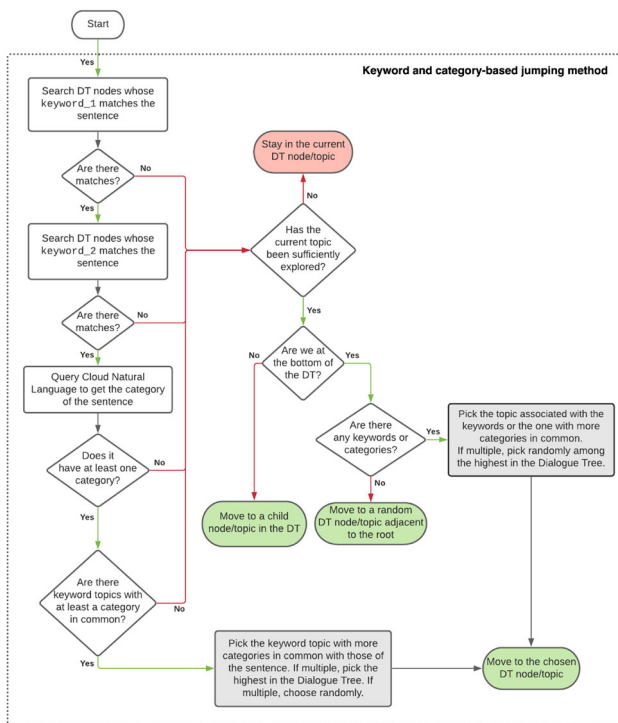
a high interest”. This sentence, like any other sentence, is provided as input to the keyword-based jumping method, to find the most appropriate topic to continue the conversation. The algorithm finds the word “interest” matching with the keyword\_1 of the topic HOBBY. Then, it checks if the sentence contains a word matching with keyword\_2 associated with the same topic: however, in this case, a wildcard is used in the topic HOBBY for keyword\_2 to guarantee that the first keyword alone is sufficient to determine the next topic. As a result, the algorithm will return the only matching node of the DT: the system will start asking the user about their hobbies (which does not look very appropriate, since the user was talking about bank interests). The second limitation related to the behaviour of this algorithm arises when no keywords are found and we are at the bottom of the DT: in this case, as already mentioned, the algorithm has no information to continue the conversation and *jumps* to a random topic close to the root.

### 3.3.2 Keyword- and Category-Based Topic Matching

The second jumping method is meant to solve the issues of the previous method, or at least reduce their frequency, by enabling a more complex understanding of the semantic content of the sentence. For this purpose, this jumping method not only exploits the keywords, but it also takes into account the “category” of the user sentence: in the current implementation, it uses the third-party services provided by Cloud Natural Language (CNL) API<sup>10</sup>, an advanced tool for Sentiment Analysis, Entity Recognition, and Content Classification, among the others.

To exploit the information related to the category of the sentence, the Content Classification of the topics contained in the Ontology is a fundamental operation and needs to be performed in a setup phase to create a mapping between the topics contained in the Ontology and the CNL category hierarchy. Off-line, before the system starts, the classification procedure is performed for all the topics. An algorithm puts together all the sentences associated with each topic (inserted into the Ontology by experts and then automatically composed, or added by the users during the conversation through a mechanism not described in this article [52]), sends them to CNL, and associates the returned categories to the corresponding topic in the Ontology and then in the DT. On-line, during the conversation, this mapping will allow the jumping method to find which topics of the Ontology match best (i.e., have more categories in common) with the category (if any) of the sentence pronounced by the user according to CNL. In the current Ontology, 122 topics have not been associated with a category, as none was found to be appropriate. Hence, the topics mapped to at least a CNL

<sup>10</sup> <https://cloud.google.com/natural-language?hl=en>.



**Fig. 5** Flow chart of the keyword- and category-based topic matching and jumping method

Category<sup>11</sup> are 2,348. Figure 5 shows a possible implementation of the box Process the user sentence by navigating the DT... in Fig. 3 when using both keywords and categories for topic matching. As a first step, the jumping method checks whether the sentence contains at least 20 tokens: the Content Classification feature of CNL does not work if the input does not satisfy this requirement. If the sentence is not long enough, it is replicated until it contains at least 20 words. Then it proceeds as follows:

- If the user sentence contains the two keywords corresponding to a topic in the Ontology *and* CNL returns at least a category associated with the sentence (left part of the diagram), the jumping method selects the topic(s) in the Ontology with the greatest number of categories in common: if there is more than one topic, it picks the closest to the DT root or, if there are multiple topics at the same level, it chooses randomly;
- Otherwise, if keywords do not match (i.e., there is no keyword<sub>1</sub> and/or keyword<sub>2</sub>) *or* the sentence does not have any category:
  - If there are still relevant questions to be asked or assertions to be made about that topic, the system stays on the same topic.

<sup>11</sup> <https://cloud.google.com/natural-language/docs/categories>.

– Otherwise:

- If not at the bottom of the DT, the Dialogue Management System keeps on exploring the DT along its branches;
- Otherwise:
  - If there are keywords *or* CNL categories associated with the sentence (we already know there are not both), the jumping method selects the topic(s) with matching keywords *or* the greatest number of categories in common: if multiple, it picks the one closer to the root in the DT or randomly if they are at the same level.
  - Otherwise, the jumping method returns a random topic immediately below the root.

This second jumping method is more complex and computationally expensive than the previous one (and, in the current implementation, it requires third-party NLP services), but it provides more stability and consistency when selecting the next conversation topic. Also, it diminishes the chance of performing random changes of conversation topic when at the bottom of the DT and no keywords are found: if the sentence has at least one category associated, recognized by CNL, there is still a chance of finding a coherent topic to jump to in order to continue the conversation.

## 4 Materials and Methods

Experimental tests with recruited participants have been performed for multiple purposes:

- Evaluate the impact of the different solutions for Dialogue Management we developed, either keyword-based or both keyword- and category-based, to improve the subjective perception of the system in terms of coherence and user satisfaction;
- Compare our solutions with Replika (one of the most advanced commercial social chatbots) as well as with baselines consisting of a system choosing replies randomly, and with a human pretending to be a chatbot.

Among the most famous chatbots mentioned in Sect. 2.3 we chose to use Replika for our experiment as it is available as Android<sup>12</sup> and IOS application, and in a web version. Moreover, as already mentioned, it is backed by Open AI's sophisticated GPT-2<sup>13</sup>: a transformer machine learning model that uses deep learning to generate text output that has

<sup>12</sup> Replika's Android application has been downloaded more than 5,000,000 times from the Google Play Store.

<sup>13</sup> <https://en.wikipedia.org/wiki/GPT-2>.

been claimed to be indistinguishable from that of humans [45].

#### 4.1 Participants and Test Groups

The participants for this test have been recruited through posts on Social Networks (Facebook and Twitter) as well as in Robotics and Computer Science classes at the University of Genoa, for a total of 100 volunteer participants aged between 25 and 65. The only inclusion criteria are that participants shall be able to read and write in English and to use a PC with a network connection from home. Since the test is anonymous and no personal data is collected, the whole procedure is fully compliant with the EU General Data Protection Regulation. Each time a new participant is recruited, they are randomly assigned to a group from 1 to 5 (where each group corresponds to a different conversation system), and given the instructions to perform the test. Based on the group they were assigned, participants had to interact with:

1. The system exploiting the keyword-based Dialogue Management algorithm;
2. The system exploiting the keyword and category-based Dialogue Management algorithm;
3. A system that chooses the next topic randomly (i.e., regardless of what the user says);
4. A human pretending to be a chatbot (i.e., in a Turing-test fashion);
5. Replika.

It shall be reminded that the purpose of the experiments is to test the impact of the Dialogue Management solutions we proposed, which emphasize the problem of switching between different topics and not sentence generation. Since Replika uses GPT-2, a language generation model, it is expected to perform better in generating individual sentences than systems 1, 2, and 3, and a human pretending to be a chatbot is expected to perform even better.

Our system is general purpose in the sense that, whatever the user says, it always tries to find the most suitable topic that matches the user expectations. However, in its current version, it composes pre-existing phrases to produce answers and does not exploit generative models. In principle, we conjecture that a language generation model could be included in the future to generate new sentences starting from the pre-stored phrases in the Ontology and then in the DT. However, this is out of the scope of this work: the purpose of the comparison is to show that, even in presence of a more advanced approach for producing individual sentences, our solution for Dialogue Management may still have a significantly positive impact on user experience.

For reasons due to the Covid-19 pandemic, participants could not take the tests with a robot, but they had to inter-

act with their assigned conversational system (1, 2, 3, 4, or 5) using their computer, by establishing a connection to a remote server in our laboratory. It is crucial to mention that the text-based user interface was identical for all the systems: therefore, participants assigned group 4 were not aware that they were interacting with a human, writing at a terminal, and participants assigned to group 5 were not aware that they were interacting with Replika (with a human manually copying and pasting sentences from Replika to the terminal and vice-versa).

In total, each session included 20 exchanges (i.e., a pair of utterances) between the participant and the system.

#### 4.2 Questionnaire

Immediately after the conversation session, participants were asked to fill in a questionnaire, divided into two parts.

The first section of the questionnaire required to evaluate what we defined as the Coherence of system's replies (Sect. 4.2.1 below), for a total of 20 replies, individually scored, in a 7-point Likert scale, where 1 means *not Coherent* and 7 means *Coherent*. The second part consisted of the SASSI questionnaire (Sect. 4.2.2).

##### 4.2.1 Coherence Measure

The instructions provided to each participant clarify what is meant by “evaluating the Coherence”. The aim is to have a way to determine whether the replies of the chatbot are *semantically consistent* with what the user says and/or with what the chatbot itself has previously said. Hence, if the user believes that the reply of the chatbot is perfectly consistent with what has just been said, then the user is instructed to assign 7 to that reply. Otherwise, if the reply of the chatbot has nothing to share with what has just been said in the previous exchange, such a reply should be evaluated with 1. Values in between shall be assigned to replies that are loosely consistent with the current conversation topic or the new topic raised by the user.

The average Coherence score assigned by a participant is computed over all replies; the average Coherence score of a system (i.e., 1, 2, 3, 4, or 5) is computed over all participants that interacted with it.

##### 4.2.2 SASSI Questionnaire

We decided to use this widely known tool to measure the user experience during the conversation, after performing some research and comparing it with other tools commonly used for this purpose [28].

Figure 6 reports the Subjective Assessment of Speech System Interfaces (SASSI) questionnaire.

System Response Accuracy	1. The system is accurate.
	2. The system is unreliable.
	3. The interaction with the system is unpredictable.
	4. The system didn't always do what I wanted.
	5. The system didn't always do what I expected.
	6. The system is dependable.
	7. The system makes few errors.
	8. The interaction with the system is consistent.
	9. The interaction with the system is efficient.
Likeability	10. The system is useful.
	11. The system is pleasant.
	12. The system is friendly.
	13. I was able to recover easily from errors.
	14. I enjoyed using the system.
	15. It is clear how to speak to the system.
	16. It is easy to learn to use the system.
	17. I would use this system.
	18. I felt in control of the interaction with the system.
Cognitive Demand	19. I felt confident using the system.
	20. I felt tense using the system.
	21. I felt calm using the system.
	22. A high level of concentration is required when using the system.
	23. The system is easy to use.
Annoyance	24. The interaction with the system is repetitive.
	25. The interaction with the system is boring.
	26. The interaction with the system is irritating.
	27. The interaction with the system is frustrating.
	28. The system is too inflexible.
Habitability	29. I sometimes wondered if I was using the right word.
	30. I always knew what to say to the system.
	31. I was not always sure what the system was doing.
	32. It is easy to lose track of where you are in an interaction with the system.
Speed	33. The interaction with the system is fast.
	34. The system responds too slowly.

**Fig. 6** The Subjective Assessment of Speech System Interfaces (SASSI) questionnaire

The published version has 34 items distributed across six scales, each item being scored with a 7-point Likert scale from Strongly Disagree to Strongly Agree:

1. *System Response Accuracy* (9 items): it refers to the user perceptions of the system as accurate and therefore doing what they expect. This will relate to the system's ability to correctly recognise the speech input, correctly interpret the meaning of the utterance, and then act appropriately;
2. *Likeability* (9 items): it refers to the user ratings of the system as useful, pleasant, and friendly;
3. *Cognitive Demand* (5 items): it refers to the perceived amount of effort needed to interact with the system and the feelings resulting from this effort;
4. *Annoyance* (5 items): it refers to the extent to which users rate the system as repetitive, boring, irritating, and frustrating;

5. *Habitability* (4 items): it refers to the extent to which the user knows what to do and knows what the system is doing;
6. *Speed* (2 items): it refers to how quickly the system responds to user inputs.

Controversial opinions about the usage of the SASSI can be found in [22]. It has been pointed out that, before being considered as a standard for measuring the quality of speech system interfaces, the SASSI questionnaire should be piloted and validated at a broader scale, and thoroughly revised in the process [53] to more firmly establish its psychometric properties. However, to the best of our knowledge, no other validated tools exist to capture the constructs that we want to measure: then, to strengthen the results provided by the SASSI, we introduced the aforementioned Coherence measure, whose correlation with the SASSI scales will be presented later in this Section, by computing Pearson's correlation index based on the collected data.

### 4.3 Data Collection

In total,  $100 \times 20 = 2000$  exchanges between the five systems and the participants (400 per group), individually scored for their Coherence, have been collected, as well as 100 completed SASSI questionnaires (20 per group). To analyze the results of this experiment, an Excel file containing three sheets has been created.

- The first sheet contains all the answers to the first part of the questionnaire, concerning the Coherence of the replies, divided per group (i.e., 1, 2, 3, 4, and 5);
- The second sheet contains all the answers to the second part of the questionnaire, divided per group (i.e., 1, 2, 3, 4, and 5) and per SASSI scale. For the Accuracy, Habitability and Speed scales, the scores assigned to negative statements (2, 3, 4, 5, 29, 31, 32, and 34) have been inverted; while for the Cognitive Demand, as a higher score has a negative meaning, we inverted the scores assigned to positive statements (19, 21, 23);
- The third sheet contains a statistical analysis of all data. All datasets have been tested for normality both with the Shapiro-Wilk Test and with criteria based on the descriptive statistics (data are normally distributed if the absolute values of both Skewness and Kurtosis are  $\leq 1$ ).

The internal consistency of SASSI data has been checked by computing the Cronbach's alpha for all scales: for Accuracy,  $\alpha=0.91$ ; for Likeability,  $\alpha=0.91$ ; for Cognitive Demand,  $\alpha=0.60$ , for Annoyance,  $\alpha=0.81$ ; for Habitability,  $\alpha=0.64$ ; for Speed,  $\alpha=0.92$ . Most scales/groups present a reliability of more than 0.80, ranging from good to excellent. However, Cognitive Demand and

Habitability present a reliability in the range of 0.6-0.7, which is considered questionable. We hypothesize that low alpha values can be due to the fact that we did not recruit only native English speakers for our experiments, since our inclusion criteria only required that participants could read and write in English: the fact that Cognitive Demand and Habitability have some negative and/or complex questions may have had an impact on the internal consistency of the collected scores. Results related to such scales are reported in Sect. 5, even if they should be taken *cum grano salis*.

The average scores of every group/scale are computed and pairwise compared. To assess whether there is a significant difference between the datasets, the *Mann-Whitney U test* has been used: this test is an alternative to the t-test when data are not normally distributed. In case the comparison is performed among two normally distributed samples, the *Welch's t-test* has been additionally performed.

The Ontology used by systems 1, 2, and 3, including all concepts as well as related sentences and keywords, the resulting DT with topics of conversation, the mapping between CNL categories and topics in the DT, and finally Excel files with individual replies to questionnaires as well as data analysis are openly available<sup>14</sup>

### 5 Results

This section presents the results obtained for the Coherence and the six scales of the SASSI. First, we computed the mean value and the standard deviation for each group, with the aim of comparing them two by two, making the null hypothesis that the compared groups are equal. That is, for Coherence and each SASSI scale, we test the null hypothesis that the scores corresponding to the compared groups, when averaged over participants to each group, are the same. To verify if such hypothesis shall be rejected, we performed the Shapiro-Wilk test for normality on each of the five datasets: if the distribution was not normal, we performed only the Mann-Whitney U test, while if the distribution was normal, we performed both the Mann-Whitney U test and the Welch's t-test. In case both tests were performed, we compared the p-values to verify if the outcomes were consistent.

For the Coherence and each SASSI scale, we report a histogram with the mean values and the standard deviations, and a table with the U-value and U-critical (corresponding to  $p = 0.05$  for the Mann-Whitney U test), and the p-values (computed with both tests, when appropriate). The green cells highlight the cases when the null hypothesis is rejected, with  $p < 0.05$ , i.e., there is a significant difference between the groups (in many cases, we found  $p < 0.01$ ).

Figures 7, 9, 11, 13, 15, 17, and 19, report the histograms representing the average and the standard deviation for each group respectively for the Coherence, Accuracy, Likeability, Cognitive Demand, Annoyance and Speed. Please notice that, in contrast with the other scales of the SASSI, a higher value of Cognitive Demand represents a negative aspect.

Figures 8, 10, 12, 14, 16, 18, and 20, present the tables containing the values computed through the Mann-Whitney U test and Welch's t-test. Note that, when both statistical tests are performed, the p-values are always consistent.

#### 5.1 Correlation between Coherence and SASSI scales

The correlation between Coherence measures and each SASSI scale has been computed by considering the average score totalled by each participant (yielding 100 values for Coherence and 100 for each SASSI scale) and computing the Pearson's correlation index by pairing Coherence measures with the corresponding scores of each SASSI scale. The value of the correlation found when comparing Coherence with Accuracy turns out to be  $r=0.78$ ; with Likeability  $r=0.69$ ; with Cognitive Demand  $r=-0.45$ ; with Annoyance  $r=-0.51$ ; with Habitability  $r=0.56$ ; with Speed  $r=-0.24$ . As expected, a high positive correlation is reported between Coherence and Accuracy and, to a minor extent, with Likeability and Habitability. A weaker, negative correlation is found with Cognitive Demand and Annoyance. An easily explainable small negative correlation with Speed is reported: system 4 (i.e., a human pretending to be a chatbot) obtains the highest

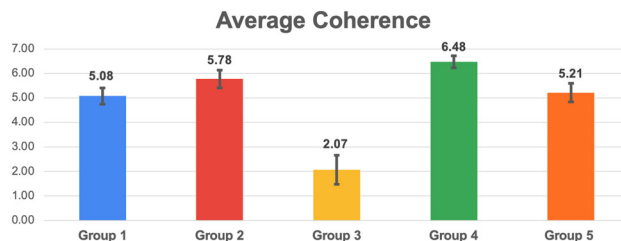


Fig. 7 Histogram reporting the average Coherence and the standard deviation of each group

Coherence	Mann-Whitney U Test			Welch's t-test
	U-value	U-critical	p-value	p-value
1 vs 2	106.5	127	0.01	/
1 vs 3	~0	127	< 0.01	/
1 vs 4	~0	127	< 0.01	/
1 vs 5	193	127	0.86	/
2 vs 3	~0	127	< 0.01	/
2 vs 4	76	127	< 0.01	/
2 vs 5	120.5	127	0.03	0.02
3 vs 4	~0	127	< 0.01	/
3 vs 5	~0	127	< 0.01	/
4 vs 5	~0	127	< 0.01	/

Fig. 8 Pairwise statistical comparison of Coherence for different groups

<sup>14</sup> <http://caressessrobot.org/IJSORO2021/>.

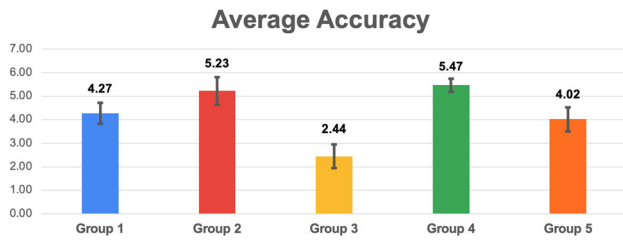


Fig. 9 Histogram reporting the average Accuracy and the standard deviation of each group

Accuracy	Mann-Whitney U Test			Welch's t-test
	U-value	U-critical	p-value	p-value
1 vs 2	106	127	0.01	0.01
1 vs 3	~0	127	< 0.01	< 0.01
1 vs 4	~0	127	< 0.01	< 0.01
1 vs 5	171.5	127	0.45	0.41
2 vs 3	~0	127	< 0.01	< 0.01
2 vs 4	188.5	127	0.76	0.41
2 vs 5	88.5	127	< 0.01	< 0.01
3 vs 4	~0	127	< 0.01	< 0.01
3 vs 5	58	127	< 0.01	< 0.01
4 vs 5	~0	127	< 0.01	< 0.01

Fig. 10 Pairwise statistical comparison of Accuracy for different groups

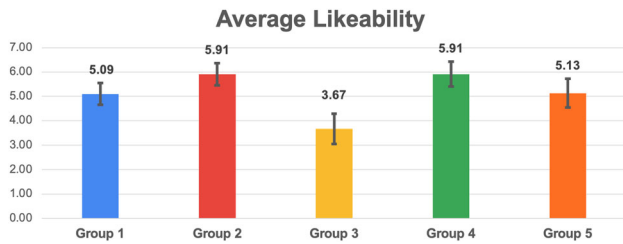


Fig. 11 Histogram reporting the average Likeability and the standard deviation of each group

Likeability	Mann-Whitney U Test			Welch's t-test
	U-value	U-critical	p-value	p-value
1 vs 2	104.5	107	0.01	/
1 vs 3	66	127	< 0.01	< 0.01
1 vs 4	96.5	127	0.01	/
1 vs 5	179.5	127	0.59	0.92
2 vs 3	~0	127	< 0.01	/
2 vs 4	194	127	0.88	/
2 vs 5	125	127	0.04	/
3 vs 4	~0	127	< 0.01	/
3 vs 5	71.5	127	< 0.01	< 0.01
4 vs 5	119	127	0.03	/

Fig. 12 Pairwise statistical comparison of Likeability for different groups

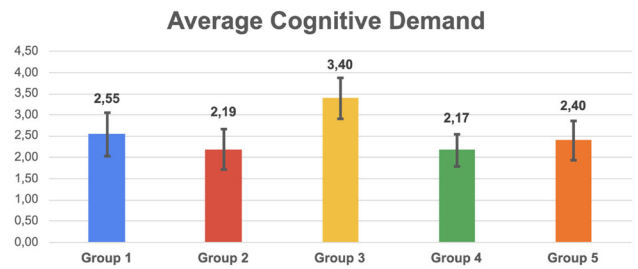


Fig. 13 Histogram reporting the average Cognitive Demand and the standard deviation of each group

Cognitive Demand	Mann-Whitney U Test			Welch's t-test
	U-value	U-critical	p-value	p-value
1 vs 2	156.5	127	0.25	0.26
1 vs 3	105	127	0.01	0.01
1 vs 4	156	127	0.24	0.19
1 vs 5	185.5	127	0.70	0.63
2 vs 3	68	127	< 0.01	< 0.01
2 vs 4	196	127	0.93	0.94
2 vs 5	169	127	0.41	0.48
3 vs 4	61	127	< 0.01	< 0.01
3 vs 5	90.5	127	< 0.01	< 0.01
4 vs 5	172.5	127	0.47	0.40

Fig. 14 Pairwise statistical comparison of Cognitive Demand for different groups

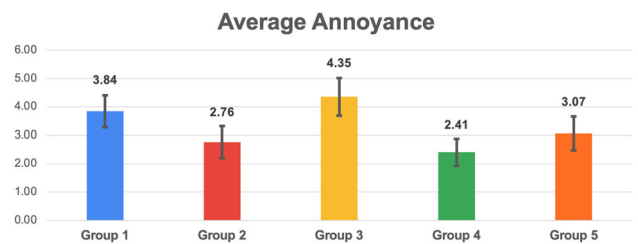


Fig. 15 Histogram reporting the average Annoyance and the standard deviation of each group

Annoyance	Mann-Whitney			Welch's t-test
	U-value	U-critical	p-value	p-value
1 vs 2	101.5	127	0.01	/
1 vs 3	142.5	127	0.12	0.19
1 vs 4	78	127	< 0.01	< 0.01
1 vs 5	126.5	127	0.05	0.04
2 vs 3	66	127	< 0.01	/
2 vs 4	168.5	127	0.40	/
2 vs 5	172.5	127	0.47	/
3 vs 4	~0	127	< 0.01	< 0.01
3 vs 5	94.5	127	< 0.01	< 0.01
4 vs 5	140	127	0.11	0.06

Fig. 16 Pairwise statistical comparison of Annoyance for different groups



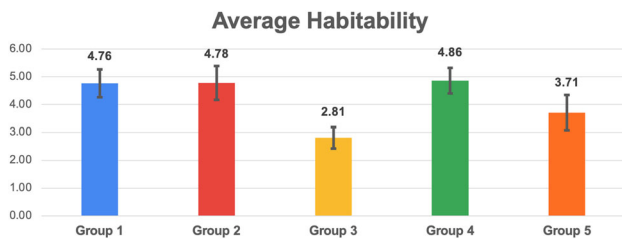


Fig. 17 Histogram reporting the average Habitability and the standard deviation of each group

Habitability	Mann-Whitney			Welch's t-test
	U-value	U-critical	p-value	p-value
1 vs 2	193	127	0.86	0.97
1 vs 3	~0	127	< 0.01	< 0.01
1 vs 4	200	127	0.99	0.75
1 vs 5	104.5	127	0.01	0.01
2 vs 3	~0	127	< 0.01	< 0.01
2 vs 4	180.5	127	0.61	0.80
2 vs 5	107.5	127	0.01	0.01
3 vs 4	~0	127	< 0.01	< 0.01
3 vs 5	121.5	127	0.03	0.01
4 vs 5	99	127	0.01	< 0.01

Fig. 18 Pairwise statistical comparison of Habitability for different groups

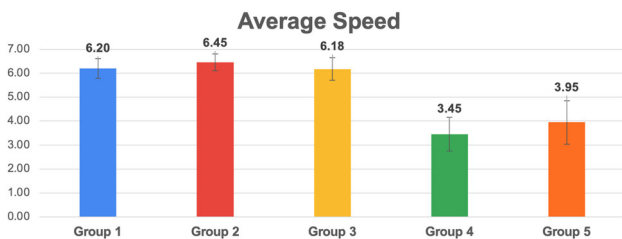


Fig. 19 Histogram reporting the average Speed and the standard deviation of each group

Speed	Mann-Whitney			Welch's t-test
	U-value	U-critical	p-value	p-value
1 vs 2	163.5	127	0.33	/
1 vs 3	199	127	0.99	/
1 vs 4	~0	127	< 0.01	/
1 vs 5	53	127	< 0.01	/
2 vs 3	166	127	0.36	/
2 vs 4	~0	127	< 0.01	/
2 vs 5	~0	127	< 0.01	/
3 vs 4	~0	127	< 0.01	/
3 vs 5	56	127	< 0.01	/
4 vs 5	167	127	0.38	0.33

Fig. 20 Pairwise statistical comparison of Speed for different groups

Coherence scores but needs a longer time to reply, whereas system 3 (i.e., the one that chooses topics randomly) obtains the lowest Coherence scores but has no delays when answering since no reasoning algorithms are involved.

## 6 Discussion

This section highlights and examines the major findings.

### 6.1 Coherence

Figure 7 shows that the standard deviations of groups 1, 2, and 5 are very similar and not too big: this means that, in general, the evaluations regarding the Coherence were quite homogeneous. Looking at the standard deviation of group 3, it is clear that the scores assigned by participants to its Coherence were more dissimilar. This may be because the random answers provided by the system during this test are more coherent during some conversations rather than others, or because some participants may be biased to interpret sentences to re-establish coherence with the context. Eventually, the standard deviation of group 4 is the smallest one: participants, without knowing that they were interacting with a real human, gave a homogeneously positive evaluation of the Coherence of the system's replies.

From Figure 8, and by looking at the averages of the Coherence reported in the histogram in Fig. 7, we can state that, as expected, group 3 has the lowest Coherence (with a significant difference with all the others), and group 4 has the highest Coherence (again with a significant difference from the others). The statistical analysis revealed a significant difference between groups 1 and 2 and between groups 2 and 5.

Results confirm that the additional cost for category extraction and matching (2) positively impacts the perceived Coherence with respect to using keywords only (1), and it is sufficient for beating Replika (5).

### 6.2 Accuracy

Figure 9 shows that the standard deviations of the Accuracy corresponding to different groups are very similar, except for the one of group 4 which is lower than the others: this indicates that the participants evaluated more homogeneously the Accuracy of this "system" (the human).

As for Coherence, the random system (3) is the worst one. An interesting result is that the system exploiting the keyword and category-based Dialogue Management algorithm (2) is the only one as accurate as a human (4) (no statistically significant difference found), other than being remarkably more accurate than all the other systems.

### 6.3 Likeability

The standard deviations shown in Fig. 11 are very similar; group 3 reports a slightly higher standard deviation, which indicates that participants had different opinions regarding the utility and the pleasantness of the systems. As happened

for the Coherence (Sect. 4.2.1), the higher standard deviation corresponding to group 3 may be related to chance, as well as to biases of participants to favourably interpret random replies to re-establish missing coherence.

As it could be expected, from the histograms it is immediate to notice that group 3 has the lowest Likeability. As regards the other groups, the average Likeability scores present less prominent differences.

Considering both the table in Fig. 12 and the histogram in Fig. 11, the most interesting result is that the system that exploits the keyword and category-based Dialogue Management algorithm (2), is significantly more likeable than the one exploiting keywords only (1) and Replika (5). Moreover, the Likeability of (2) is similar to that of a human (4), while this is not true when comparing the human with the system using keywords only (1), fully justifying the additional cost for category extraction and matching.

## 6.4 Cognitive Demand

The standard deviations in Fig. 13 are again very similar; group 4 reports a slightly lower standard deviation which, as always, indicates that the participants' evaluation regarding the Cognitive Demand of the system is homogeneous.

As regards the averages, group 3 deviates from the others. This means that a higher cognitive effort is needed to interact with the system: this is since the replies of the system are random, hence the participants found it more difficult to easily interact with it.

By looking at the p-values in Fig. 14 and the averages in Fig. 13, we can conclude that all the systems require the same amount of effort during the interaction, except the one providing random replies (3). A valuable result is that both systems we developed (1 and 2) present no significant difference, in terms of Cognitive Demand, with respect to a human (4) and Replika (5).

## 6.5 Annoyance

As shown in Fig. 15, the standard deviations are quite high and similar to one another. Again, the lowest standard deviation corresponds to group 4: when participants unknowingly interacted with a human, their ratings were more homogeneous. Considering the averages, as we could expect, the system used with group 3 appears to be the most annoying, while the lowest Annoyance score is associated with group 4.

Examining the table in Fig. 16 and the histogram in Fig. 15, we can conclude that system 2 (keywords plus categories) is significantly less annoying than 1 (keywords only), and it is as annoying as 4 (human) and 5 (Replika). Notice also that 1 (keywords only) is significantly more annoying than both 4 (human) and 5 (Replika).

## 6.6 Habitability

From the histograms in Fig. 17 it can be seen that the standard deviations regarding the Habitability are very similar for all groups, with the smallest one being that of group 3.

Regarding the averages, the smallest one is that corresponding to group 3, which presents also the lowest average Habitability. This result was expected, as the random system had also the lowest Accuracy, Likeability, and Annoyance, and the highest Cognitive Demand, whose scores are likely not completely independent of one another.

Considering the p-values in Fig. 18, and the averages in Fig. 17, it can be observed that the system used with group 2 (keywords plus categories) turned out to be more habitable than 5 (Replika), while 1 (keywords only) is less habitable than 5. Very interestingly, both our systems 1 and 2 are not distinguishable from 4 (human) in Habitability.

## 6.7 Speed

As it can be seen in Fig. 19, the system tested with group 3 (random replies) is faster than 1 and 2, since they are all connected to the CARESSES Cloud, but the former does not call the Dialogue Management algorithm. However, the standard deviation of group 3 is slightly higher than those of groups 1 and 2, even if the average is lower: we conjecture that the extremely low pleasantness of the overall interaction with this system, confirmed by all the results of the previous sections, negatively influences the perception of Speed.

Groups 4 and 5 report a higher standard deviation and a lower average. However, these results cannot be compared with those of the other tests. Since group 4 involves the unaware interaction with a human, the final score depends on the typing speed of the human. The same reasoning applies to group 5, which involved the unaware interaction with Replika: a human acted as an intermediary between the participant and Replika, by manually typing Replika's replies which required some additional time.

Considering both the table in Fig. 20, and the histogram in Fig. 19, the only relevant result is that there are no significant Speed differences between 1, 2, and 3. No conclusions shall be drawn when comparing groups 4 and 5.

## 7 Conclusion

The article proposes a novel system for knowledge-based conversation designed for Social Robots and other conversational agents. The proposed system relies on an Ontology for the description of all concepts that may be relevant conversation topics, as well as their mutual relationships. We compare two algorithms, based on the Ontology, for Dialogue Management that select the most appropriate conversation topics

depending on the user input: the two versions differ in their computational cost and/or the need for third-party NLP services. Moreover, they implement slightly different strategies to ensure a conversation flow that captures the user intention to drive the conversation in specific directions, while avoiding purely reactive responses.

Experiments performed with 100 volunteer participants, interacting with five different conversational systems (one of which is a human pretending to be a chatbot in a Turing-test fashion), support our intuitions about the importance of dialogue flow management, the improvements brought by the proposed solution based on the semantic category of the user sentence, as well as the positive correlations between Coherence in flow management and some scales of the SASSI (Subjective Assessment of Speech System Interfaces) questionnaire. Specifically, system (2), exploiting both keywords and categories for Dialogue Management, is statistically superior in Coherence and most SASSI scales (i.e., Accuracy, Likeability, Habitability and Speed) to Replika, one of the most popular chatbots worldwide. Since the SASSI is a widely known tool to measure the user experience during the conversation (and it positively correlates with our Coherence measure), we conclude that our approach may be a step forward towards properly managing the conversation flow to improve the user experience, thus supporting the aim of the paper. System (1), the one that exploits keywords only, is less computationally demanding as it does not require third-party services for category extraction and matching, and less performing than system (2): however, it can be a good “cheaper” alternative to produce an engaging conversation with the user. Indeed, it exhibits results not statistically discernible from those of Replika in Coherence, Accuracy, Likeability, Cognitive Demand, superior in Habitability and inferior in Annoyance.

The approach presented in this article has obvious limitations since it relies on an Ontology of concepts and the related sentences to talk about such concepts, which needs to be manually encoded by experts. This issue has been addressed in subsequent work, which explored strategies for expanding the knowledge base at run-time during the interaction with the user [52].

**Funding** Open access funding provided by Università degli Studi di Genova within the CRUI-CARE Agreement.

**Data Availability** The datasets generated during and/or analysed during the current study are available from the corresponding author on reasonable request.

## Declarations

**Conflict of interest** The authors declare that they have no conflict of interest.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Belpaeme T, Kennedy J, Ramachandran A, Scassellati B, Tanaka F (2018) Social robots for education: a review. *Sci Robot*, 3(21)
- Pan Y, Okada H, Uchiyama T, Suzuki K (2015) On the reaction to robot's speech in a hotel public space. *Int J Soc Robot* 7:911–920
- Pandey AK, Gelin R (2018) A mass-produced sociable humanoid robot: Pepper: the first machine of its kind. *IEEE Robot Autom Mag* 25(3):40–48
- Niemelä M, Heikkilä P, Lammi H, Oksman V (2019) A social robot in a shopping mall: Studies on acceptance and stakeholder expectations. *Soc Robot: Technol, Soc Ethical Aspects of Human-Robot Interact*. Springer
- Broekens J, Heerink M, Rosendal H (2009) Assistive social robots in elderly care: a review. *Gerontechnology* 8(2):94–103
- Papadopoulos C, Hill T, Battistuzzi L, Castro N, Nigath A, Randhawa G, Merton L, Kanoria S, Kamide H, Chong N-Y, Hewson D, Davidson R, Sgorbissa A (2020) The CARESSES study protocol: testing and evaluating culturally competent socially assistive robots among older adults residing in long term care homes through a controlled experimental trial. *Arch. Public Health* 78(1):1–20
- Papadopoulos C, Castro N, Nigath A, Davidson R, Faulkes N, Menicatti R, Khaliq AA, Recchiuto CT, Battistuzzi L, Randhawa G, Merton L, Kanoria S, Chong NY, Kamide K, Hewson D, Sgorbissa A (2021) The CARESSES randomised controlled trial: exploring the health-related impact of culturally competent artificial intelligence embedded into socially assistive robots and tested in older adult care homes. *Int J Soc Robot*
- Adelman RD, Tmanova LL, Delgado D, Dion S, Lachs MS (2014) Caregiver burden: a clinical review. *JAMA* 311(10):1052–1060
- Yang G-Z, Nelson BJ, Murphy RR, Choset H, Christensen H, Collins SH, Dario P, Goldberg K, Ikuta K, Jacobstein N, Kragic D, Taylor RH, McNutt M (2020) Combating COVID-19 - the role of robotics in managing public health and infectious diseases. *Sci Robot*, 5(40)
- Zhu N, Zhang D, Wang W, i X, Yang B, Song J, Zhao X, Huang B, Shi W, Lu R, Niu P, Zhan F, Ma X, Wang D, Xu W, Wu G, Gao GF, Tan W, (2020) China novel coronavirus investigating and research team. A novel coronavirus from patients with pneumonia in china. *New Eng J Med* 328(8):727–733
- Miller G (2020) Social distancing prevents infections, but it can have unintended consequences. *Science*
- Schuetzler RM, Grimes GM, Giboney JS (2020) The impact of chatbot conversational skill on engagement and perceived humanness. *JIMS* 37(3):875–900
- Bruno B, Chong NY, Kamide H, Kanoria S, Lee J, Lim Y, Pandey AK, Papadopoulos C, Papadopoulos I, Pecora F, Saffiotti A, Sgorbissa A (2017) Paving the way for culturally competent robots: a position paper. *IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 553–560

14. Recchiuto CT, Sgorbissa A (2020) A feasibility study of culture-aware cloud services for conversational robots. *IEEE Robot Autom Lett* 5(4):6559–6566
15. Recchiuto CT, Gava L, Grassi L, Grillo A, Lagomarsino M, Lanza D, Liu Z, Papadopoulos C, Papadopoulos I, Scalmato A, Sgorbissa (2020) A Cloud services for culture aware conversation: Socially assistive robots and virtual assistants. UR, p. 270–277
16. Ware J, Gandek B (1998) Overview of the sf-36 health survey and the international quality of life assessment (iqola) project. *J Clin Epidemiol* 51(11):903–12
17. Nomura T, Kanda T, Suzuki T (2006) Experimental investigation into influence of negative attitudes toward robots on human-robot interaction. *AI Soc* 20:138–150
18. Hays R, DiMatteo MR (1987) A short-form measure of loneliness. *J Pers Assess* 51(1):69–81
19. Mavridis N (2015) A review of verbal and non-verbal human-robot interactive communication. *Rob Auton Syst* 63:2–35
20. Bruno B, Recchiuto CT, Papadopoulos I, Saffiotti A, Koulouglioti C, Menicatti R, Mastrogiovanni F, Zaccaria R, Sgorbissa A (2019) Knowledge representation for culturally competent personal robots: requirements, design principles, implementation, and assessment. *Int J Soc Robot* 11:515–538
21. Lewis JR, Hardzinski ML (2015) Investigating the psychometric properties of the speech user interface service quality questionnaire. *Int J Speech Tech* 18:479–487
22. Lewis JR (2016) Standardized questionnaires for voice interaction design. *AVIXD*, 1(1)
23. Dinan E, Rollerand S, Shuster K, Fan A, Auli M, Weston J (2018) Wizard of Wikipedia: Knowledge-powered conversational agents. *Int Conf on Learning Representations (ICLR)*
24. Huang J, Zhou M, Yang D (2007) Extracting chatbot knowledge from online discussion forums. *Paper presented at the 20th Int J Conf on AI (IJCAI'07)*, p. 423–428
25. Kim S, Kwon O-W, Kim H (2020) Knowledge-grounded chatbot based on dual wasserstein generative adversarial networks with effective attention mechanisms. *Appl Sci*, 10(9)
26. Ghazvininejad M, Brockett C, Chang M-W, Dolan B, Gao J, Yih W-T, Galley M (2018) A knowledge-grounded neural conversation model. Paper presented at the Thirty-Second AAAI Conf. on AI
27. Ritter A, Cherry C, Dolan WB (2011) Data-driven response generation in social media. *EMNLP*, p. 583–593
28. Kocaballi AB, Laranjo L, Coiera E (2018) Measuring user experience in conversational interfaces: a comparison of six questionnaires. *Proceedings of the 32nd International BCS Human Computer Interaction Conference (HCI)*
29. Adiwardana D, Luong M-T, So DR, Hall J, Fiedel N, Thopilan R, Yang Z, Kulshreshtha A, Nemade G, Lu Y, Le QV (2020) Towards a human-like open-domain chatbot. arXiv preprint [arXiv:2001.09977](https://arxiv.org/abs/2001.09977)
30. Roller S, Dinan E, Goyal N, Ju D, Williamson M, Liu Y, Xu J, Ott M, Shuster K, Boureau Y-L, Smith EM, Weston J (2020) Recipes for building an open-domain chatbot. arXiv preprint [arXiv:2004.13637](https://arxiv.org/abs/2004.13637)
31. Venkatesh A, Khatri C, Ram A, Guo F, Gabriel R, Nagar A, Prasad R, heng M, Hedayatnia B, Metallinou A, Goel R, Yang S, Raju A (2018) On evaluating and comparing conversational agents. arXiv preprint [arXiv:1801.03625](https://arxiv.org/abs/1801.03625)
32. Cervone A, Riccardi G (2020) Is this dialogue coherent? learning from dialogue acts and entities. *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue*
33. Higashinaka R, Meguro T, Imamura K, Sugiyama H, Makino T, Matsuo Y (2014) Evaluating coherence in open domain conversational systems. Fifteenth Annual Conference of the International Speech Communication Association
34. Nimavat K, Champaneria T (2017) Chatbots: an overview types, architecture, tools and future possibilities. *Int J Sci Res Develop* 5(7):1019–1024
35. D'silva GS, Thakare S, More S, Kuriakose J (2017) Real world smart chatbot for customer care using a software as a service (saas) architecture. *IEEE Int Conf on IoT in Social, Mobile, Analytics and Cloud (I-SMAC)*, pp 658–664
36. Nuruzzaman M, Hussain OK (2018) A survey on chatbot implementation in customer service industry through deep neural networks. Paper presented at the 15th Int Conf on e-Business Engineering (ICEBE), p. 54–61
37. Zhou L, Gao J, Li D, Shum H-Y (2020) The design and implementation of xiaoice, an empathetic social chatbot. *Comput Linguist* 46(1):53–93
38. Shum H-Y, He X, Li D (2018) From eliza to xiaoice: Challenges and opportunities with social chatbots. *Front Inf Technol Electron Eng* 19:10–26
39. Weizenbaum J (1966) Eliza - a computer program for the study of natural language communication between man and machine. *Commun ACM*, 9(1)
40. Colby K (1975) Artificial paranoia: a computer simulation of paranoid processes. Elsevier
41. Wallace RS (2009) The anatomy of a.l.i.c.e. parsing the turing test. Springer, p. 181–210
42. Turing A (1950) Computing machinery and intelligence. *Mind* 59(236):433–460
43. Dilmegani C Top 30 successful chatbots of 2021 & reasons for their success. <https://research.aimultiple.com/top-chatbot-success/#bots-for-general-use>. Accessed: 2021-04-24
44. Pardes A The emotional chatbots are here to probe our feelings. <https://www.wired.com/story/replika-open-source/>. Accessed: 2021-04-24
45. Alleyne A Chat bots are becoming uncannily human. can they be our friends? <https://edition.cnn.com/style/article/tech-loneliness-replika-wellness/index.html>. Accessed: 2021-04-24
46. Weston J, Dinan E, Miller AH (2018) Retrieve and refine: Improved sequence generation models for dialogue. *Proceedings of the 2018 EMNLP Workshop SCAI: The 2nd International Workshop on Search-Oriented Conversational AI*
47. Zhang Y, Sun S, Galley M, Chen Y-C, Brockett C, Gao X, Gao J, Liu J, Dolan B (2020) DIALOGPT: Large-scale generative pre-training for conversational response generation. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*
48. Dinan E, Humeau S, Weston J, Chintagunta B (2020) Build it break it fix it for dialogue safety: Robustness from adversarial human attack. *Proceedings of EMNLP-IJCNLP 2019 - 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing*, p 4537–4546
49. Guarino N (1998) Formal ontology and information systems. *FOIS'98 Conf*, p. 81–97
50. Motik B, Patel-Schneider PF, Parsia B (2008) Owl 2 web ontology language structural specification and functional-style syntax. WWW Consortium
51. Carrithers M, Candea M, Sykes M, Holbraad M, Venkatesan S (2010) Ontology is just another word for culture: motion tabled at the 2008 meeting of the group for debates in anthropological theory. *Crit Anthropol*,
52. Grassi L, Recchiuto CT, Sgorbissa A (2021) Knowledge triggering, extraction and storage via human-robot verbal interaction. arXiv preprint [arXiv:2104.11170](https://arxiv.org/abs/2104.11170)
53. Weiss B, Wechsung I, Naumann A, Möller S (2008) Subjective evaluation method for speech-based uni- and multimodal applications. *Perception in Multimodal Dialogue Systems*, Springer, pp 285–288

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Lucrezia Grassi** is a Ph.D. student in Bioengineering and Robotics at the University of Genova. She got her master's degree in Robotics Engineering in 2020 with the thesis "A Knowledge-Based Conversation System for Robots and Smart Assistants", and she is currently pursuing her Ph.D. on multiparty interaction between humans and artificial agents. Her interests include Social Robotics, autonomous conversation systems, mixed and virtual reality.

**Carmine Tommaso Recchiuto** is Assistant Professor at the University of Genova, where he teaches Experimental Robotics, ROS programming, and Computer Science. His research interests include Humanoid and Social Robotics (with a specific focus on knowledge representation and human-robot interaction), wearable sensors, and Aerial Robotics. He has been the Coordinator of software integration and Head of Software Development for the CARESSES project, aimed at endowing social robots for older adults with cultural competence. He has also been the local Coordinator for the BrainHuRo project, developing Brain-Computer Interfaces for humanoid robots' remote control. He is the author of more than 40 scientific papers published in International Journals and conference proceedings.

**Antonio Sgorbissa** is Associate Professor at the University of Genova, where he teaches Real-Time Operating Systems, Social Robotics, and Ambient Intelligence in EMARO+, the European Master in Advanced Robotics. He is the Coordinator of National and EU research projects, among which the H2020 project CARESSES (<http://caressessrobot.org/>). Also, he is the local Coordinator of the ongoing IENE-10 project, aimed at preparing health and social care workers to work with intelligent robots in health and social care environments. He received several acknowledgements for his recent work: among the others, CARESSES has been acknowledged "Project of the month" by the EU, its technologies have been acknowledged by the EU Innovation Radar, and the project has been included among the "100 Italian Robotics & Automation Stories" in a report presented by Enel S.p.a. in February 2020. His research focuses on Autonomous Robotic Behaviour, with a focus on Culture-Awareness, Knowledge representation, Motion Planning, Wearable, and Ubiquitous Robotics. He is the author of about 150 articles indexed in international databases and has been a member of the Organizing Committee in the top-ranked robotic conference as well as Associate Editor for the International Journal of Advanced Robotic Systems edited by SAGE and Intelligent Service Robotics by Springer.