

Social Network Chatbots for Smoking Cessation: Agent and Multi-Agent Frameworks

Davide Calvaresi
HES-SO Valais-Wallis
Sierre, Switzerland
davide.calvaresi@hevs.ch

Jean-Paul Calbimonte
HES-SO Valais-Wallis
Sierre, Switzerland
jean-paul.calbimonte@hevs.ch

Fabien Dubosson
HES-SO Valais-Wallis
Sierre, Switzerland
fabien.dubosson@hevs.ch

Amro Najjar
Luxembourg University
Luxembourg
amro.najjar@uni.lu

Michael Schumacher
HES-SO Valais-Wallis
Sierre, Switzerland
michael.schumacher@hevs.ch

ABSTRACT

Asynchronous messaging is leading human-machine interaction due to the boom of mobile devices and social networks. The recent release of dedicated APIs from messaging platforms boosted the development of computer programs able to conduct conversations, (i.e., *chatbots*), which have been adopted in several domain-specific contexts. This paper proposes SMAG: a chatbot framework supporting a smoking cessation program (JDF) deployed on a social network. In particular, it details the single-agent implementation, the campaign results, a multi-agent design for SMAG enabling the modelization of personalized behavior and user profiling, and highlighting of coupling chatbot technology with and multi-agent systems.

CCS CONCEPTS

• **Applied computing** → **Health informatics**; • **Information systems** → *Personalization*; *Social networks*; • **Human-centered computing** → *Web-based interaction*; *Natural language interfaces*.

KEYWORDS

MAS, agent-based chatbots, smoking cessation

ACM Reference Format:

Davide Calvaresi, Jean-Paul Calbimonte, Fabien Dubosson, Amro Najjar, and Michael Schumacher. 2019. Social Network Chatbots for Smoking Cessation: Agent and Multi-Agent Frameworks. In *IEEE/WIC/ACM International Conference on Web Intelligence (WI '19)*, October 14–17, 2019, Thessaloniki, Greece. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3350546.3352532>

1 INTRODUCTION

In the past decades, a fast-paced technological evolution has deeply changed human habits which are increasingly tied to intelligent

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WI '19, October 14–17, 2019, Thessaloniki, Greece

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6934-3/19/10...\$15.00

<https://doi.org/10.1145/3350546.3352532>

systems (IS) (e.g., virtual assistants with multi-modal interactions). Human-machine interaction (HMI) is crucial for such a systems [25], which enabling vocal, video, and gesture-based interactions are disrupting the technological market exacerbating the competition (e.g., Siri/Apple, Cortana/Microsoft, Alexa/Amazon, and Google Now/Google). Besides these eye-catching systems, more discrete and asynchronous chat-like communications are still often preferred. Chatbots (CBTs) are programs able to conduct conversations currently booming [4, 36]. Perceivable as *conversational agents*, CBTs can be engaged in a wide range of domains. High availability & response immediateness, ubiquity, impersonal & confidential nature of communications, and social acceptance make CBTs outstand in motivational (e.g., social network campaigns [7]) and support (e.g., customer management [40]) applications (recently approaching eHealth [9] and assisted-living scenarios [17]). Personalized support and social aspects are crucial for health programs fostering behavioral change (e.g., smoking cessation) [9, 20, 38]. Technical performance (e.g., response and page-load time) has to be coupled with *non-technical* indicators and metrics (e.g., user's context, expertise, mood, personality traits, and expectations). Hence, the metrics have to be both Quality of Service (QoS) and Quality of Experience (QoE) [28, 29]. Recently, IS has enhanced personalization (while emulating human-like interactions). In particular, Multi-Agent Systems (MAS) gained a crucial role in this evolution [10, 11, 30]. Therefore, we envision MAS addressing behavior and coordination tasks in CBTs systems.

Contributions. This paper studies the implications of coupling MAS and CBT in the context of health-support programs. It presents SMAG (*i*) a framework based on a single-agent chatbot supporting a smoking cessation program deployed on Facebook and (*ii*) its evolution: a multi-agent design. The rest of the paper is organized as follows: Section 2 explores the state of the art, Section 3 outlines the case study, Section 4 introduce SMAG. Section 5 describes the implementation of SMAG and Section 6 explains the results of the program. Finally, Section 7 concludes the paper.

2 STATE OF THE ART

2.1 HMI & Chatbots

Nowadays platforms such as Telegram and Facebook started to release APIs to develop CBT, finding niche-application domains such as e-commerce and customer care support. Recently, several CBT

cloud services and frameworks emerged, fostering further developments in the area. For example, Amazon Lex [2] provides a service for CBT development, including natural language understanding, and automatic speech recognition. Powered by Google's machine learning techniques, Dialogflow [15] provides a framework aiming at understanding human conversation. Finally, Microsoft Bot Framework [1] is a tool-set supporting CBT development, including APIs for text and speech analysis. Most of these services rely on cloud management of user interaction data, often colliding with confidentiality and privacy restrictions enforced by ethical committees (especially for health-related programs [24]). Moreover, these frameworks are mostly dedicated to natural language and speech processing, with little support for conversation flow management, coordination, and user profiling.

2.2 MAS & Chatbots

To address these limitations, CBT have explored different approaches, including the adoption of MAS as underlying technology to support modeling human-like behaviors and dynamics. In [6], the combination of CBT and MAS is embodied as a stimulus-reply state automaton and a goal-driven probabilistic agent (defined as a Partially Observable Markov Decision Process). The CBT can react to the stimuli received from the user in either a predefined manner or through a conversation planning process. An agent can manage a set of actions linkable to two main goals over the whole conversation: (i) an *immediate goal* - achievable within a single step in the dialogue, (ii) a *global goal* - achievable by the end of the conversation. By doing so, elements of pragmatics have been added in the description of dialogues [22]. Thus, the CBT behavior can adhere to its mood, undertaking different interaction modes. Moreover, the agent-based CBT can store and elaborate on the stimuli received from the environment coupling them with the state of the user-system conversation.

Several initiatives, reporting mixed results, concerned the utilization of social networks for supporting health and well-being programs, including smoking cessation. For example, Tweet2Quit [33] showed promising potential for the development of online support systems. However, these results do not show yet conclusive evidence on the effect of social interactions or information messages. Other approaches have included WhatsApp and Facebook as social platforms, for instance targeting smoking relapse prevention [13], or smoking cessation in young adults [34]. Another study [14] analyzed social interactions among participants of such social networks for smoking cessation, identifying the importance of moderators, the different roles taken by users, and the limitations of non-personalized support.

2.3 Quality of Experience (QoE)

Quality of Experience emerged in the 2000's to bridge the gap between technical quality metrics (such as QoS) and the user subjective evaluation of the service quality [28]. QoE is typically useful for applications in which technical parameters are not enough to ensure end-user satisfaction.

Since it takes personal and technical factors into account, QoE, offers a comprehensive measure allowing to assess the end-user satisfaction. By definition, agents are self-interested and bound to

an individualistic view of the environment. For these reasons, as has been discussed by recent works on user modeling, intelligent agents have the potential to model user satisfaction, expectations, and seek to maximize the user QoE [31]. In these works, each user is represented by a personal agent that takes into account her current context and preferences and acts on her behalf [32].

Summarizing, bots are seen as potential tools to (i) disseminate information, (ii) profile users to provide personalized information, and (iii) motivate behavioral change. Merging CBT and MAS is still in an early stage. Although promising, the integration of the capabilities of conversational agents within the MAS dynamics has yet to be exploited.

3 SMOKING CESSATION: A CASE STUDY

Smoking cigarettes provokes three main types of addictions: chemical, routinary, and social [26]. Over the years, different measures such as media campaigns, increasing the prices of tobacco products, and smoke-free policies, have been implemented to support smoking cessation [19]. However, the hardest addictions to fight are related to routine and social behaviors. Hence, a personalized support system can achieve comparable benefits to those provided by an individual clinic-based intervention, but with lower costs. The main features a smoking cessation system must have are: (i) proactive interaction with the smokers, (ii) possible participants integration/grouping, (iii) support any individual in any stage of their challenge, (iv), guaranteeing easy access to the support application (e.g., over the internet), and (v) provide tailored interventions considering possible behavioral variations.

3.1 The JDF Program

The first edition of the smoking cessation program named «J'arrête de fumer» (JDF) was launched in 2015 in the canton of Valais and quickly expanded to all the French-speaking Swiss cantons. The program wanted to group Facebook communities of individuals willing to quit smoking on the same day. It allowed participants to share their approaches and techniques, and to support each other. After six months, out of 7000 participants, 13.5% definitely stopped smoking [18]. The main difficulty emerged from this first campaign was answering *timely* the huge number of messages received (e.g., more than 13K the first week). Timing is crucial to support the participant in fighting the "high-intensity" craving phases that last only for a few minutes.

A CBT-based approach was the envisioned solution. Indeed, CBTs are prompt-to-answer and scalable (i.e., able to handle thousands of messages simultaneously). The CBT-based JDF automated: profiling the participants, guiding them during the program, and providing immediate support to overcome the craving phases.

3.2 CBT design for JDF

Figure 1 shows the duration and sequence of the JDF-phases in which the assisting CBT has to operate: *P1 - pre-cessation*: tracking the smoking habits (cigarettes/time), duration; *P2 - ongoing-cessation*: timely assisting participants during the craving periods. Tracking the cigarettes in *P1* allows profiling the participants' consumption habits and identifying/clustering the distribution of the cigarettes (e.g., the morning or aperitif time). During the second

phase, the CBT assists the participants during the craving periods by operating accordingly to what has been defined in the previous phases.




Figure 1: Smoking cessation program: 2-phase timeline.

P1: The interaction begins with the bot asking demographic questions and assessing the type of smoking dependence. Then, the cigarette-tracking period starts (requiring the participants to communicate the bot whenever they are smoking). For every cigarette notified, the bot asks a few questions to enrich its classification (e.g., location –at home/office, social scenario– alone/accompanied, ongoing activity, and mood). Although several variables can impact patients habits, two weeks of monitoring are sufficient to sketch patient smoking patterns. Based on the smoking profile, participants can define distracting actions that they might be willing to follow during given cravings.

P2: Elaborating on the knowledge base built in *P1*, the CBT can provide any-time support by associating the request of support to corresponding actions. Considering the duration of *P2* (16 weeks), the suggested activities might become repetitive, thus losing effectiveness. Therefore, by reasoning on patient-similarity and clusters of cigarettes, new plans can be generated and suggested.

4 AGENT-BASED CBT FRAMEWORK

This section introduces SMAG, its design principles, and discusses both the architecture of the single (version employed in the JDF program) and multi-agent version of SMAG.

4.1 Design principles

SMAG is designed according to the following principles:

- DP1** User-centric data acquisition: the CBT profiles the user (enabling personalized recommendations, motivation messages, etc.).
- DP2** User-state tracking: A state-machine-based CBT design enable each user to progress independently, tracking the user behavior (i.e., sequence of states).
- DP3** Dynamic workflow: the engine allows to define arbitrarily complex interactions (adaptable to various use-cases – e.g., health questionnaires).
- DP4** Programmability: to allow integrations (e.g., with NLP and ML), the engine allows developers to programmatically define its behavior/integration with other software tools.

4.2 SMAG: Single-agent approach

Figure 2 schematize the single-agent SMAG. The CBT agent is modeled as a state-machine-powered component (Figure 2c). It interacts with users (Figure 2a) via the Facebook messenger platform (Figure 2b). Mapping the design principles on the agent components (Figure 3), we have: DP1 and DP4 serving the communicative capabilities, and DP2 and DP3 affecting the agent mind (autonomy, knowledge, and conversational behaviors).




Figure 2: Architecture design: single-agent SMAG.

SMAG manages four main aspects: (i) directly acquired knowledge through question-answer interactions (e.g., monitoring survey, & demographics), or indirectly through analysis of past interactions; (ii) a conversational model, which drives the types of actions dictated by the CBT logic; (iii) single-agent behavior & interactions. (iv) data collection and logging (e.g., activities, answers, tracking and storage of analytics and user behavioral change – persistence).




Figure 3: Single-agent schematization.

In terms of knowledge acquisition, the preparation phase is essential for the CBT to understand the cigarettes consumption pattern, context, and habits. Thus, the CBT would firstly ask the cigarette context (e.g., professional, personal alone, or driving), desire intensity, and mood. In turn, the bot can send each participant an overview of the smoking behavior, (Figure 4), including statistics, and the top 3 distracting alternatives to smoking (e.g., drinking water or chewing). This information is useful for the bot (profiling), and the participant (raising awareness and motivation).

The single-agent bot instantiates multiple state machines (one per user) managing the transitions' logic (e.g., after the name, ask the age). Users can repeat states depending on the type of situation (e.g., multiple cigarette-reporting states), or skip given states if necessary (e.g., avoiding context-irrelevant questions). The advantage of the state-machine model is its extensibility and simplicity (scalable implementation). CBT behavior depends on the conversational model. The behaviors implemented for JDF can be summarized as: (i) data acquisition interactions: i.e. demographic questions, dependency/habits questions, cigarette tracking, smoking context, midterm surveys; (ii) feedback interactions: i.e. sharing with the user his smoking habits statistics; (iii) motivation interactions, including motivational images, famous quotes, encouragement messages; and (iv) information interactions, i.e. impact on health, diseases related to smoking, etc.

The data layer in SMAG contains both (i) data representation of the possible states and conversations engaged by the agent (multiplexed according to how many users are participating), and (ii) the log of every single interaction. Thus, serving both business logic storage and tracking/auditing purposes.

4.2.1 Single-Agent SMAG: discussion. The single-agent approach brings several advantages: a shorter (almost zero) response time w.r.t. the manual approach adopted in traditional cessation campaigns; improved knowledge representation and understanding of the user context, and finally portability and lightweight execution. However, due to its simplicity, the single-agent approach implicates

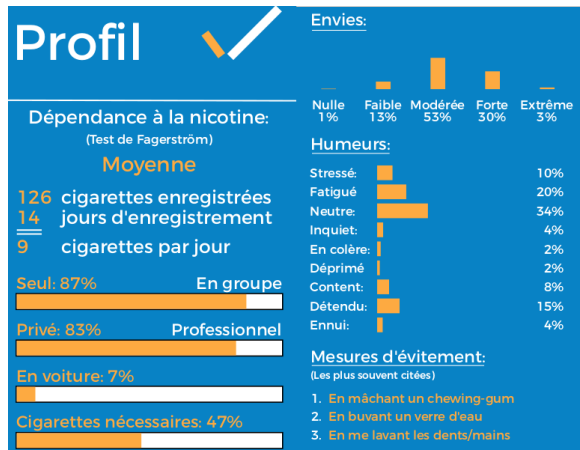


Figure 4: Smoking profile elaborated by the CBT.

several disadvantages. Delegating the management of all the participants to a single agent implies the sequential execution of the *support behaviors* or a virtual parallelization (Figure 5a) exploiting the multi-processing capability of the developed system. Despite the possibility of applying a multi-processing approach, there is no simple way to make processes interacting together, thus hampering information sharing between similar profiles. Hence, the need for distributed interactions among user-facing software entities is one of the main motivations for a multi-agent-based architecture. Moreover, in case new dynamics or more complex personalization would arise, to tailor ad-participant behaviors might be necessary, thus compromising the scalability. Finally, there is a lack of mechanisms to manage 3rd-party data analysis and aggregation (e.g., advanced NLP analytics or machine learning), which can provide additional value to the bot interactions. Summarizing, to overcome the limitations and enhance the benefits of this approach, the upgrade from a single to a multi-agent CBT platform is necessary.




Figure 5: Single-Agent vs. Multi-Agent platform workflow.

4.3 SMAG: a Multi-Agent Design

In the Multi-Agent SMAG, modules and structures are common to all the agents. Nevertheless, behaviors and knowledge modules can assume different identities after the bot start-up. This guarantees a better and smoother adaptation to the needs of the related participant (see Figure 5b). It is worth to notice, that w.r.t the Single-Agent, in the Multi-Agent approach (Figure 6) a frontier agent named “gateway-agent” has been introduced to dispatch incoming chat messages to the agent related to a given user (Figure 6). While rules/steps are fixed, personalizing the interaction can be encoded

in the behaviors module, and updated and tracked in the knowledge module. Moreover, thanks to the modularity of the Multi-Agent approach, this latter module can store the historical data based on the information collected from the participant and similar profiles. By doing so, similar agents can interact with each other at any time, supporting the understanding of the related participant and improving the assistance provided. Moving to the MAS approach allows




Figure 6: Multi-Agent architecture of the SMAG platform

personalization and improves users integration. Hence, to increase the chances of success of the smoking cessation, SMAG should have an estimate of each user satisfaction. To have a quantified estimate of end-user subjective satisfaction, the user agent can rely on Quality of Experience (QoE) [8, 28]. As mentioned above, QoE is influenced by several user-related (e.g., previous user experience), context-related, and technical (e.g., response time) influence factors [35]. Such an influence makes QoE subjective and varies from one user to another. In an application like JFD, user satisfaction is a key factor ensuring the user willingness to complete the treatment.

A personal agent representing the user, his/her preferences, and personality traits allows delivering a user-tailored personalized experience. However, in order to ensure a high QoE, the user agent should build an adequate model of user preferences, personality, and expertise related to his/her context. Thus, the agent is facilitated in estimating when is the best time to initiate *feedback* or *motivation interactions*. For instance, the user agent could maintain an estimate of the *user-bother cost* [37] and rely on this indicator to determine, given the context and the level of user attention, whether sharing habit and statistics with other users is recommended. The multi-agent architecture shown in Figure 6 allows such personalization given the direct binding user-personalized agent. Moreover, this agent can adapt its behavior to adhere the user evolution. Since the participants of a given campaign start the cessation on the same day, the multi-agent SMAG platform can run several trials simultaneously, ensuring information atomicity and isolation. Moreover, exploiting the dynamics encoded in the behavior module, new-released agents can learn from agents in a more advanced stage.

4.3.1 Multi-Agent SMAG discussion. Moving from single- to multi-agent SMAG can: (i) enable to reuse the agent knowledge in further campaign/studies, (ii) migrating capability, (iii) enhance the CBT knowledge about the smoking cessation and the participants’ paths, and (iii) facilitate a more efficient tailored support.

Unfortunately, some limitations have still to be acknowledged: (i) missing a standardized agent communication, (ii) to offer high QoE, an agent should construct an adaptive model of user behavior, and (iii) despite the aforementioned improvements, the design and development of the multi-agent SMAG are more complicated, demanding to change the underlying technological architecture.

5 IMPLEMENTATION

This section presents the internal architecture of the single-agent CBT engine, highlighting the technological stack used to develop and run the bot, and the core elements of the engine. The implementation of SMAG (Figure 7), has been developed in *Python*. It consists of a web server, written with the *Flask* framework served through *Nginx*, and using *uWSGI* as interpreter between *Nginx* requests and the *Python* code. *uWSGI* is also the layer allowing the parallel processing of requests, by spawning several processes running each one a different *Python* interpreter, and balancing requests over these processes. The whole system is deployed in production with the *Docker* [16] container system, and the data are stored in a *PostgreSQL* database.




Figure 7: SMAG single-agent deployment.

Messaging. The messages exchanged within the bot engine, and the underlying platform (Facebook FMP) are JSON objects dispatched through GET/POST HTTP requests. These JSON messages forwarded by the FMP, sent over HTTP requests, follow a structure that is also kept inside SMAG, not following multi-agent standard protocols yet.

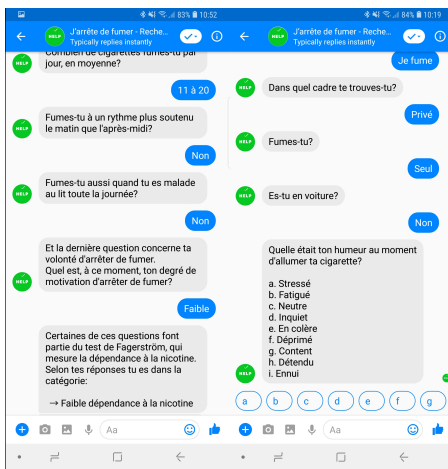


Figure 8: (L): dependence test; (R): cigarette tracking.

Technology stack. The CBT engine core has been developed using a *Python* virtual environment (enabling to replicate the setup and facilitating to populate the virtual environment). Moreover, the *Python*-based engine can facilitate (i) the future integration with NLP or ML processes; and (ii) the integration with web applications (e.g., Django and Flask) which may allow offering a web interface for audit and management.

The developed CBT could be integrated into several messaging and social networks. In particular, the current version runs on Facebook – platform used by the JDF program – (Figure 8). Facebook does not force the use of any particular programming language or framework (all interactions are through raw HTTP requests). The developed CBT uses *fbmq* [3] which deals with the needed HTTP requests behind a clean *Python* API.

The *Flask* framework has been used for setting up the REST endpoints (or URLs) that our project is listening to, for receiving events from the Facebook servers.

Engine core. Interactions with CBTs can be arbitrarily complex. For instance, questions may be skipped depending on previous answers to avoid repetitions (some questions may be asked in a random order, so to avoid structural biases). In this context, a purely declarative model, such as using JSON or XML files to represent a CBT discussion model, is not powerful enough.

While using *Python* makes it easy to build complex models, it also leaves room to undesired, and potentially wrong behaviors. To compromise between the ease of use and behavior complexity, our engine allows defining the interactions with the CBT as State Machines. With this architecture, developers only need to focus on the definition of states that represent different interactions with the users, while the engine itself ensures the rest of the machinery (i.e. performing and storing transitions between states, handling user events, and managing data storage)

The State Machine status is kept synchronized with the *User* table in the database. A *Python* class represents a state, and the fully qualified name of the python class is kept stored in this table. When a new event is triggered, the engine takes care of loading the current state from the class name in the database, and then it forwards the user’s event to this corresponding state. This state is then responsible for returning what should be the next state for a given user. Then, the engine stores this information back in the database. This way of abstracting the workflow as state transitions allows to model complex interactions.

6 CAMPAIGN RESULTS

The first CBT-based campaign JDF showed promising results, both in terms of usage of the platform, and the successful cessations. Out of 270 participants in the program (73% women), more than 200 actively used the CBT on the preparation phase, especially for tracking cigarette consumption. This task, allowed to compose smoking profiles, contributed to increasing the awareness on the consumption patterns, the necessity, circumstances, and possible substitution/avoidance strategies. Having a clear picture of addiction and behavioral patterns (cigarette consumption) increased the participants’ motivation. In this edition of the JDF program, 78 of the 270 participants (28.9%) succeeded in the smoking cessation goal, three months after the last cigarette. The cessations obtained

with the CBT-based campaign is 10% higher than the previous edition of JDF (no SMAG CBT). While there is still substantial room for improvement (e.g., during the cessation phase) the CBT has already shown to be crucial for the success of the program.




Figure 9: Aggregated cigarettes by time of day and mood.

We plan to improve the SMAG interactions by further exploiting the data collected in the preparation phase within multi-agent dynamics (enabling personalized motivation – e.g., time, context, and mood dependent (Figure 9).

7 CONCLUSIONS

This paper presented an agent-based CBT framework supporting smoking cessation over the Facebook social network. In particular, the implementation of a single-agent framework, and design of multi-agent approaches have been detailed, highlighting effectiveness and simplicity in the first case and the completeness, enhanced performance, and extended capabilities in the second. A real-life smoking cessation campaign was carried out with the single-agent CBT, showing promising results, and the potential of using this technology. The employment of the multi-agent approach is expected to provide additional improvements w.r.t. the current single-agent framework. One key improvement is to make user agents *social* [39], i.e. these user agents can communicate with others to share useful experiences or strategies.

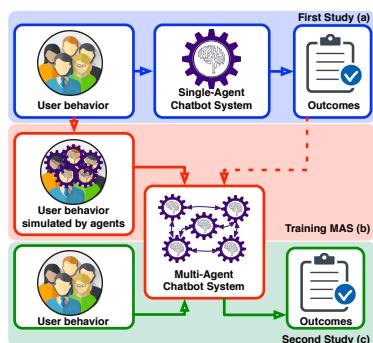


Figure 10: SMAG: behavior and simulation interactions.

The future works are illustrated in Figure 10. In particular, SMAG will rely on (i) the anonymized outcomes and user behaviors of

smoking cessation programs (either successful or unsuccessful outcomes – Figure 10a), supplemented with (ii) the behaviors of simulated users (Figure 10.b). By analyzing this data, SMAG will adapt its decision to provide a better service for users still under treatment. Moreover, we plan to develop mechanisms to implement and distribute the SMAG among several servers, thus better distributing the load according to the geographical requests, and finally, to implement mechanisms to handle the migration of an agent from one server to another according to the possible change of location of its related users. Another cardinal future research direction is to endow SMAG with explanatory behavior, allowing it to explain its actions and point out the motivation behind them [5, 12] (aligned with the Research in the domain of eXplainable Artificial Intelligence – XAI) [21]. The need of explainability is more accentuated in applications like SMAG where lay users would need to understand the rationale behind the actions of their personal agents in order to trust the system and use it in a continuous manner [23, 27].

REFERENCES

- [1] [n. d.]. <https://dev.botframework.com>. Accessed: 2018-09-29.
- [2] [n. d.]. AALex. <https://aws.amazon.com/lex/>. Accessed: 2018-09-29.
- [3] [n. d.]. fbmq. <https://github.com/conbus/fbmq>. Accessed: 2018-09-29.
- [4] Sumit Amar, Cheryl Lynn Wu, Abdulwafa Choudhary, James Cao, Dennis Franklin Olson, Srinivas Bhagavathula, Sunit Lohtia, Anjana Sarkar, and Mayur Vijay Jain. 2019. System and Method for Automated Chatbots. US Patent App. 16/219,124.
- [5] Sule Anjomshoae, Amro Najjar, Davide Calvaresi, and Kary Främling. 2019. Explainable agents and robots: Results from a systematic literature review. In *Proceedings of the 18th International Conference on Autonomous Agents and Multi-Agent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 1078–1088.
- [6] Carlo Bentivoglio, Diego Bonura, Vincenzo Cannella, Simone Carletti, Arianna Pipitone, Roberto Pirrone, Pier Rossi, and Giuseppe Russo. 2010. Intelligent Agents supporting user interactions within self regulated learning processes. *Journal of E-learning and Knowledge Society* 6, 2 (2010), 27–36.
- [7] Jacqueline Brixey, Rens Hoegen, Wei Lan, Joshua Rusow, Karan Singla, Xusen Yin, Ron Artstein, and Anton Leuski. 2017. SHIHbot: A Facebook chatbot for Sexual Health Information on HIV/AIDS. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*.
- [8] Kjell Brunnström, Sergio Ariel Beker, Katrien De Moor, Ann Dooms, Sebastian Egger, Marie-Neige Garcia, Tobias Hossfeld, Satu Jumisko-Pyykkö, Christian Keimel, Mohamed-Chaker Larabi, et al. 2013. Qualinet white paper on definitions of quality of experience. (2013).
- [9] Jean-Paul Calbimonte, Davide Calvaresi, Fabien Dubosson, and Michael Schumacher. 2019. Towards Profile and Domain Modelling in Agent-Based Applications for Behavior Change. In *International Conference on Practical Applications of Agents and Multi-Agent Systems*. Springer, 16–28.
- [10] Davide Calvaresi, Daniel Cesarini, Paolo Sernani, Mauro Marinoni, Aldo Franco Dragoni, and Arnon Sturm. 2017. Exploring the ambient assisted living domain: a systematic review. *Journal of Ambient Intelligence and Humanized Computing* 8, 2 (2017), 239–257.
- [11] Davide Calvaresi, Mauro Marinoni, Arnon Sturm, Michael Schumacher, and Giorgio Buttazzo. 2017. The challenge of real-time multi-agent systems for enabling IoT and CPS. In *Proceedings of the international conference on web intelligence*. ACM, 356–364.
- [12] Davide Calvaresi, Yazan Mualla, Amro Najjar, Stéphane Galland, and Michael Schumacher. 2019. Explainable Multi-Agent Systems through Blockchain Technology.
- [13] Yee Tak Derek Cheung, Ching Han Helen Chan, Chi-Keung Jonah Lai, Wai Fung Vivian Chan, Man Ping Wang, Ho Cheung William Li, Sophia Siu Chee Chan, and Tai-Hing Lam. 2015. Using WhatsApp and Facebook online social groups for smoking relapse prevention for recent quitters: A pilot pragmatic cluster randomized controlled trial. *J medical Internet research* 17, 10 (2015).
- [14] Heather Cole-Lewis, Adler Perotte, Kasia Galica, Lindy Dreyer, Christopher Griffith, Mary Schwarz, Christopher Yun, Heather Patrick, Kisha Coa, and Erik Augustson. 2016. Social network behavior and engagement within a smoking cessation Facebook page. *J of medical Internet research* 18, 8 (2016).
- [15] DialogFlow. 2018. <https://dialogflow.com>. Accessed: 2018-09-29.
- [16] Docker. [n. d.]. <https://docker.com>. Accessed: 2018-09-29.
- [17] Ahmed Fadhil and Silvia Gabrielli. 2017. Addressing challenges in promoting healthy lifestyles: the al-chatbot approach. In *Proceedings of the 11th EAI International Conference on Pervasive Computing Technologies for Healthcare*. ACM,

- 261–265.
- [18] Laura Folly, Gisana Riedo, Max Felder, Juan M. Falomir-Pichastor, and Olivier Desrichard. 2016. Rapport de l'évaluation externe du programme « J'arrête de fumer » sur Facebook. *Sur mandat du Fonds de Prévention du Tabagisme* (2016).
- [19] Centers for Disease Control, Prevention (CDC, et al. 2011. Quitting smoking among adults—United States, 2001–2010. *MMWR. Morbidity and mortality weekly report* 60, 44 (2011), 1513.
- [20] Amanda L Graham, George D Papandonatos, Bahar Erar, and Cassandra A Stanton. 2015. Use of an online smoking cessation community promotes abstinence: Results of propensity score weighting. *Health Psychology* 34, S (2015), 1286.
- [21] David Gunning. 2017. Explainable artificial intelligence (xai). *Defense Advanced Research Projects Agency (DARPA), nd Web 2* (2017).
- [22] Afsaneh Haddadi. 1996. *Communication and cooperation in agent systems: a pragmatic theory*. Vol. 1056.
- [23] Andreas Holzinger, Chris Biemann, Constantinos S Pattichis, and Douglas B Kell. 2017. What do we need to build explainable AI systems for the medical domain? *arXiv preprint arXiv:1712.09923* (2017).
- [24] Kit Huckvale, José Tomás Prieto, Myra Tilney, Pierre-Jean Benghozi, and Josip Car. 2015. Unaddressed privacy risks in accredited health and wellness apps: a cross-sectional systematic assessment. *BMC medicine* (2015).
- [25] Gunnar Johannsen. 2009. Human-machine interaction. *Control Systems, Robotics and Automation* 21 (2009), 132–62.
- [26] Andrew M Jones. 1994. Health, addiction, social interaction and the decision to quit smoking. *Journal of health economics* (1994), 93–110.
- [27] Pat Langley, Ben Meadows, Mohan Sridharan, and Dongkyu Choi. 2017. Explainable Agency for Intelligent Autonomous Systems. In *AAAI*.
- [28] Sebastian Möller and Alexander Raake. 2014. *Quality of experience: advanced concepts, applications and methods*.
- [29] A Najjar. 2015. *Multi-agent negotiation for qoe-aware cloud elasticity management*. Ph.D. Dissertation. PhD thesis, École nationale supérieure des mines de Saint-Étienne.
- [30] Amro Najjar, Olivier Boissier, and Gauthier Picard. 2017. Aquaman: an adaptive qoe-aware negotiation mechanism for saas elasticity management. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 1655–1657.
- [31] Amro Najjar, Christophe Gravier, Xavier Serpaggi, and Olivier Boissier. 2016. Modeling User Expectations & Satisfaction for SaaS Applications Using Multi-agent Negotiation. In *Web Intelligence (WI), 2016 IEEE/WIC/ACM International Conference on*. 399–406.
- [32] Amro Najjar, Xavier Serpaggi, Christophe Gravier, and Olivier Boissier. 2016. Multi-agent systems for personalized qoe-management. In *Teletraffic Congress (ITC 28), 2016 28th International*, Vol. 3. 1–6.
- [33] Cornelia Pechmann, Li Pan, Kevin Delucchi, Cynthia M Lakon, and Judith J Prochaska. 2015. Development of a Twitter-based intervention for smoking cessation that encourages high-quality social media interactions via automessages. *J medical Internet research* 17, 2 (2015).
- [34] Danielle E Ramo, Johannes Thrul, Kathryn Chavez, Kevin L Delucchi, and Judith J Prochaska. 2015. Feasibility and quit rates of the Tobacco Status Project: A Facebook smoking cessation intervention for young adults. *J of medical Internet research* 17, 12 (2015).
- [35] Ulrich Reiter, Kjell Brunnström, Katrien De Moor, Mohamed-Chaker Larabi, Manuela Pereira, Antonio Pinheiro, Junyong You, and Andrej Zgank. 2014. Factors influencing quality of experience. In *Quality of experience*. 55–72.
- [36] Bayan Abu Shawar and Eric Atwell. 2007. Chatbots: are they really useful?. In *LDV Forum*, Vol. 22. 29–49.
- [37] Ngoc Cuong Truong, Tim Baarslag, Sarvapali D Ramchurn, and Long Tran-Thanh. 2016. Interactive scheduling of appliance usage in the home. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*. AAAI Press, 869–875.
- [38] Janice Y Tsoh, Nancy J Burke, Ginny Gildengorin, Ching Wong, Khanh Le, Anthony Nguyen, Joanne L Chan, Angela Sun, Stephen J McPhee, and Tung T Nguyen. 2015. A social network family-focused intervention to promote smoking cessation in Chinese and Vietnamese American male smokers: A feasibility study. *Nicotine & Tobacco Res* 17, 8 (2015), 1029–1038.
- [39] Michael Wooldridge and Nicholas R Jennings. 1995. *Intelligent agents: Theory and practice*. *The knowledge engineering review* (1995).
- [40] Anbang Xu, Zhe Liu, Yufan Guo, Vibha Sinha, and Rama Akkiraju. 2017. A new chatbot for customer service on social media. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 3506–3510.