A Virtual Teaching Assistant for Personalized Learning

Extended Abstract

Luca Benedetto
Politecnico di Milano
Milan, Italy
luca.benedetto@polimi.it

Paolo Cremonesi
Politecnico di Milano
Milan, Italy
paolo.cremonesi@polimi.it

Manuel Parenti
Politecnico di Milano
Milan, Italy
manuel.parenti@mail.polimi.it

ABSTRACT

In this extended abstract, we propose an intelligent system that can be used as a Personalized Virtual Teaching Assistant (PVTA) to improve the students' learning experience both for online and on-site courses. We show the architecture of such system, which is composed of an instance of IBM's Watson Assistant and a server, and present an initial implementation, consisting in a chatbot that can be questioned about the content and the organization of the *RecSys* course, an introductory course on recommender systems.

KEYWORDS

intelligent agents, chatbots, conversation technology, NLP, education, virtual teaching assistants, personalized learning

CEUR-WS.org/Vol-2482/paper51.pdf

1 INTRODUCTION

Intelligent systems are extensively used in many domains, and they can bring some relevant advantages in education as well: indeed, they offer the opportunity to improve the learning experience and the quality of teaching, both in the case of online and on-site courses. So far, diverse applications have been explored: for instance, some research focused on the usage of recommender systems for suggesting new learning content [5] and to perform students' performance prediction [6]. Predicting students' performance is particularly important for e-learning, in order to improve retention and completion rates, which are one of the biggest limitations of online learning [4, 7]. Several works, such as [1, 3], discussed the possibility of using virtual teaching assistants (VTA) in order to reduce professors' workloads: indeed, VTAs can make education much more scalable since students can solve most of their problems without asking the professors for help. This extended abstract lies in this last branch of research: we propose a *Personalized Virtual Teaching Assistant (PVTA)* for "assisted learning", which consists in helping students with a series of services such as personalization of content, recommendation of learning material and student engagement, as well as other services. Also, we introduce an initial version of the PVTA, consisting in a chatbot - built leveraging the IBM's Watson Assistant - which is capable of answering students' questions about the content, the structure and the organization of the *RecSys* course, an introductory course on recommender systems. The main differences from previous research consist in the personalization offered by the system and our focus on its architecture, showing how it is built, how it works and how it will be expanded.

Copyright © CIKM 2018 for the individual papers by the papers' authors. Copyright © CIKM 2018 for the volume as a collection by its editors. This volume and its papers are published under the Creative Commons License Attribution 4.0 International (CC BY 4.0).

1.1 IBM's Watson Assistant

The IBM's Watson Assistant[1] (called "Assistant" from now on) is offered by IBM as part of the Watson suite, an AI engine that provides several NLP services. It is able to "understand natural-language input and use machine learning to respond to customers in a way that simulates a conversation between humans" and can be used to build virtual assistants. IBM does not share with customers the details of the core NLP model, thus Assistant can be used only as a black-box; however, we can adapt it to any desired application domain by feeding it with additional data in order to enlarge the training set and fine-tune the model. The training data we can feed the model with is made of objects belonging to two classes: *intents* and *entities*: *intents* identify the goals that we expect a user to have while interacting with the system, while *entities* affect the way in which Assistant reacts to a specific intent by giving it a context. Once defined intents and entities, it is necessary to create the dialog flow in order to teach Assistant how it should answer different requests.

2 RELATED WORK

In relation to this project, the most important work is the introduction of Jill Watson (JW) by A. Goel et al. [3], a VTA somewhat similar to the PVTA proposed in this document. However, although the authors showed the possible applications of JW, they never presented the details of the implementation and always considered it as a black-box; we aim at filling that gap, focusing on the architecture of our PVTA and explaining the role of the different components it is made of. Also, there are some differences between the two systems: JW aimed at completely replacing human TAs, thus it had to deal with situations outside of the educational domain, which caused some problems that still have to be addressed [2]. Our PVTA, instead, focuses on helping students in relation to the contents and the structure of the course, therefore we will not have

[1] https://console.bluemix.net/docs/services/conversation

to deal with that kind of issues. Lastly, JW did not provide any kind of personalization, while we are building a personalized system.

3 SYSTEM ARCHITECTURE

As shown in Figure 1, the PVTA is made of three main components: the front-end, a server and an instance of Assistant. The server contains the data about the course and the students, while the instance of Assistant contains the intents and entities we defined.

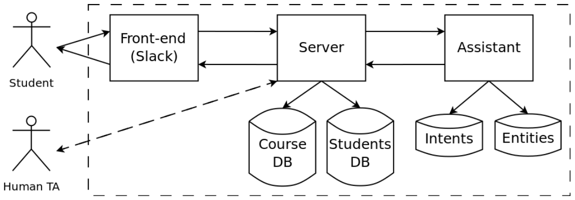

Figure 1: Architecture of the PVTA.

3.1 Front-end

The front-end of the PVTA is a simple application that enables the user to interact with Assistant and does not perform any operations on the input data; for the first implementation we used Slack[2].

3.2 Server

It is involved in different phases of the process: at the moment it performs up to three tasks for each question but this number will increase in the future, as we add new services to the PVTA.

3.2.1 Preprocessing. It is necessary if the server knows some information that Assistant is unaware of. Watson has no long-term memory to store the context of the interaction with a student and we must leverage the server in order to do so. As an example, if a student asks a question using the pronoun "it" referring to something he said in a previous question, the server has to modify the sentence in order for Assistant to understand such reference.

3.2.2 Post-processing. This is the analogous of preprocessing; for instance, when a student asks for the date of an exam, Assistant is able to understand what the user is interested in but it does not have access to the schedule (which is stored in the server). Thus, it sends an incomplete answer to the server, which fills the gap by adding the date and time of the exam.

3.2.3 Students' modeling. The server contains some intelligence as well: it collects data about students' behavior and clusters them in different groups using as similarity the intents and the entities they searched for.

3.2.4 Interaction with a human TA. Figure 1 shows that the server can interact with a human TA: Assistant assigns a confidence level to each answer and, if that is too low, the server forwards the question and the proposed answer to the human TA. He checks the proposed answer, possibly corrects it and sends it back to the server, which forwards it to the student and inserts the correct question-response pair in the training set (stored in Assistant).

3.3 Assistant

It is responsible for the NLP-related tasks of the VTA: it receives the question (possibly preprocessed) and provides an answer (which might require post-processing). In order to build a system capable of working in the educational domain, we had to define the *entities* and the *intents* related to such domain, as well as all the technical terms specific to the *RecSys* course. Each entity does not represent a unique concept, but a group of concepts; also, each concept might be referred to with different synonyms. So far, we defined 50 intents and 20 entities (more than 170 different concepts) but this number is likely to increase in the future, as we add new functionalities.

4 CONCLUSION AND FUTURE WORK

This document proposed an architecture for a PVTA capable of providing several services to students of online and on-site courses, moving towards the goal of "assisted learning". Also, we introduced an initial implementation of such system, consisting in a chatbot capable of answering the questions of students enrolled in a course about recommender systems. This chatbot is only a small part of the PVTA we propose, and we are working on new services to be implemented in the system. We are working on ways to continuously enrich the set of *intents*, *entities* and rules by monitoring the interactions between students and the PVTA: by means of knowledge extraction algorithms we aim to keep fine-tuning and improving the model even after deployment. We also plan to implement student engagement: looking at students' behavior, the PVTA can understand which are the students at risk of dropping-out and the ones not satisfied with the course; then, it can proactively intervene or send a warning to human TAs. Another aspect we are focusing on is analyzing whether it is possible to reduce the usage of the IBM's Watson Assistant and perform a bigger part of the NLP-related tasks in the server. Indeed, we have two objectives in mind: implementing some NLG (natural language generation) in the server in order to overcome the biggest limitation of Assistant, which is the impossibility to generate answers, and - as a second step - exploring the generation of personalized answers. Lastly, further work is focusing on the possibility of personalizing learning material, recommending different contents depending on the student and its interactions with the PVTA.

REFERENCES

[1] B. du Boulay. 2016. Artificial Intelligence as an Effective Classroom Assistant. *IEEE Intelligent Systems* 31, 6 (Nov.-Dec. 2016), 76-81.
[2] B. Eicher, L. Polepeddi, and A. K. Goel. 2018. Jill Watson Doesn't Care if You're Pregnant: Grounding AI Ethics in Empirical Studies. In *AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society (AIES'18).*
[3] A. K. Goel and L. Polepeddi. 2016. Jill Watson: A Virtual Teaching Assistant for Online Education. Retrieved July 16, 2018 from https://smartech.gatech.edu/handle/1853/59104
[4] M. Hlosta, Z. Zdrahal, and J. Zendulka. 2017. Ouroboros: Early identification of at-risk students without models based on legacy data. In *Proceedings of the Seventh International Learning Analytics and Knowledge Conference (LAK'17).* ACM, 6–15.
[5] A. Klasnja-Milicevic, M. Ivanovic, and A. Nanopoulos. 2015. Recommender systems in e-learning environments: a survey of the state-of-the-art and possible extensions. *Artificial Intelligence Review* 44, 4 (December 2015), 571-604.
[6] N. Thai-Nghe, L. Drumond, A. Krohn-Grimberghe, and L. Schmidt-Thieme. 2010. Recommender Systems for Predicting Students Performance. *Procedia Computer Science* 1, 2 (January 2010), 2811-2819.
[7] A. Wolff, Z. Zdrahal, A. Nikolov, and M. Pantucek. 2013. Improving retention: predicting at-risk students by analysing clicking behaviour in a virtual learning environment. In *Proceedings of the Third International Learning Analytics and Knowledge Conference (LAK'13).* ACM, 145–149.

[2] https://slack.com/