Intelligent Assistant Agent

Required: 2 diploma students   *Coordinator: Andrei Olaru* | cs@andreiolaru.ro

The goal of this project create a conversational agent that can assist the user in various tasks related to the user's digital life. The user will communicate with the agent in a limited subset of the natural language and the agent will use a visual representation (context graphs and patterns) to retain its knowledge; this knowledge can be shown to the user on demand, in a graphical form or in a textual form.

Using the agent, the user will be able to control the device and will be able to configure the assistant such that it will give reminders to the user or perform automatic actions when certain pre-described situations are detected.

The big IT companies each currently offer smart assistants that can help the users in their daily lives, by playing music, telling the time, the weather, or assisting the user in buying products. The "intelligence" of these assistants relies on processing in the cloud, using vast amounts of data obtained through machine learning.

We are exploring the possibility of creating a smart assistant that works locally on the user's device(s), relying exclusively on local data.

Good progress in this direction has already been made[1], obtaining an assistant that communicates with the user naturally (through speech recognition and voice output) and works internally by representing the user's statements and queries as graphs. Its main drawback is that (a) it is very slow and (b) the types of queries supported is limited.

Objectives

The goal of this project is to continue the work done (even if not using the same source code) improving the performance and quality of the assistance and also implement an Android version of the assistant.

An additional objective is to create an assistant which spans several of the user's devices (e.g. at least a PC and some Android devices) in order to harness computing power from more powerful machines.

Required skills

- Good knowledge of the Java programming language;
- Android development skills are a plus.

Skills learned

- Knowledge representation and reasoning;
- Experience with speech recognition and natural language processing;
- Experience in the problem of intelligent assistants.

This project is part of the AmIciTy initiative [http://aimas.cs.pub.ro/amicity/#Task]

---

[1]See Merlin Intelligent Assistant.

Lightweight Middleware for IoT

Required: 2 diploma students   *Coordinator: Andrei Olaru* | cs@andreiolaru.ro

The Internet of Things is composed of a large number of very different (heterogeneous) devices that communicate through different protocols and using different data formats. Most commercial IoT solutions are "walled gardens", allowing little interoperability with solutions from other providers.

The goal of this project is to create a software agent-based middleware that works on a variety of devices, from Raspberry Pi nodes to Android smartphones to full-size workstations. The project implies building a uniform infrastructure for all types of devices, allowing agents to easily work with any of them.

Using software agents and a distributed infrastructure, the project will help the implementation of *fog computing* in IoT.

This projects will use the vast experience of the tATAmI-2 platform for software agents, but will need a more lightweight and flexible structure.

Objectives

The main objective is to develop an abstraction for IoT devices and protocols so that the middleware can interoperate with several IoT protocols and can run on several computing platforms, such as Android and Raspberry Pi devices.

Required skills

- Good knowledge of the Java programming language;
- Android development skills are a plus;
- Raspberry Pi experience is a plus.

Skills learned

- IoT standards and protocols;
- Distributed software and communication protocols;
- Experience with software agents;
- Experience with complex software projects.

References

Olaru, Andrei. "tATAmI-2–A Flexible Framework for Modular Agents." 2015 20th International Conference on Control Systems and Computer Science. IEEE, 2015. download

This project is part of the AmIciTy initiative [http://aimas.cs.pub.ro/amicity/#tatami]

Activity Detection on Android Devices

Required: 1 diploma student　　　　*Coordinator: Andrei Olaru* | cs@andreiolaru.ro

While machine learning works great when computing resources abound, a privacy-protecting distributed implementation of context-aware, intelligent behavior requires learning on resource-constrained devices.

This means that we have to revert to simpler learning methods that prove effective in understanding the user's context.

Objectives

The goal of this research is to develop means to learn from the user's activity on smart devices (Android smartphones) only with the help of local resources.

The expected result is to be able to predict the activities of the user, using sensors on the device and also information available from local apps, in order to better assist the user in those activities and anticipate potential unwanted situations.

An additional goal is to build a dataset of sensory data using smartphones and using this dataset for training and testing.

Required skills

- Good knowledge of the Java programming language;
- Some experience with Android development is recommended.

Skills learned

- Machine learning algorithms;
- Data mining elements;
- Human activity recognition.

References

Krishnan, Narayanan C., and Diane J. Cook. "Activity recognition on streaming sensor data." Pervasive and mobile computing 10 (2014): 138-154.
Cook, Diane J. "Learning setting-generalized activity models for smart spaces." IEEE intelligent systems 2010.99 (2010): 1.
Olaru, Andrei, and Adina Magda Florea. "A Platform for Matching Context in Real Time." International Conference on Hybrid Artificial Intelligence Systems. Springer International Publishing, 2015.

This project is part of the AmIciTy initiative [http://aimas.cs.pub.ro/amicity/#incubator]