

e-ViTA workshop on Knowledge Graphs and Large Language Models.

In Evry on March 8th, 2024

<https://www.e-vita.coach/homepage/finalconference/>

14h : Gérard Chollet

Title : Introduction to the workshop on KGs and LLMs

Abstract : This introduction intends to motivate the workshop. It reports on recent publications demonstrating the complementarity of symbolic and stochastic approaches to language and world modelling. The number of publications is increasing exponentially ! 18 papers have been published on this topic since the 1st of January 2024. More than 160 papers have been published last year. Particular attention is focussed on the automatic construction of Knowledge Graphs using Large Language Models.

Biography : Gérard Chollet studied Linguistics, Electrical Engineering and Computer Science at the University of California, Santa Barbara where he was granted a PhD in Computer Science and Linguistics. In 1983, he joined a newly created CNRS research unit at ENST (now Institut Polytechnique de Paris). He supervised more than forty doctoral theses. He is currently Emeritus researcher within SAMOVAR (TSP) and consulting for Intelligent Voice Ltd, Speech Morphing Inc and Zaion.ai. Publications : <https://scholar.google.com/citations?user=ttVRli8AAAAJ&hl=fr&oi=ao>

14h20 : Haoyi Xiong

Title: Natural Language based Context Modeling and Reasoning for Ubiquitous Computing with Large Language Models

Abstract: This talk covers the increasing prevalence of Large Language Models (LLMs) in context-aware computing and their ability to model contexts and make decisions based on natural language understanding and reasoning capabilities for ubiquitous computing. This talk introduces the concept of LLM-driven Context-aware Computing (LCaC) as a paradigm for using LLMs to model context and perform reasoning without fine-tuning the model. The tutorial also demonstrates the use of texts, prompts, and autonomous agents to enable LLMs to perform context modeling and reasoning. Two showcases are presented to illustrate the practical applications of LCaC: operating a mobile z-arm in an apartment for assisted living, and planning a trip and scheduling the itinerary in a context-aware and personalized manner. The talk aims to provide an overview of the advancements in context-aware computing enabled by LLMs and the potential impact on various applications such as assisted living and personalized trip planning.

Biography: Dr. Haoyi Xiong is now a principal architect in Big Data Laboratory at Baidu Inc. Previously, he was a tenure track assistant professor at Missouri University of Science and Technology and a research associate at the University of Virginia. He earned his Ph.D. in computer science from Télécom SudParis jointly with Université Pierre-et-Marie-Curie, Paris 6, under supervision of Prof Daqing Zhang, Prof Monique Becker, and Dr. Vincent Gauthier. He is also affiliated to the University of Central Florida as a Graduate Faculty Scholar.

Publications : https://scholar.google.com/citations?user=f_Kcie0AAAAJ&hl=fr&oi=ao

14h40 : Graham Wilcock

Title : New technologies for spoken dialogue systems: LLMs, RAG and the GenAI Stack

Abstract : Spoken dialogue systems are making a paradigm shift from Conversational AI to Generative AI. New technologies support retrieval augmented generation (RAG) using open source large language models (LLMs). From these, I focus on three specific developments:

(1) The addition of vector indexing to Neo4j graph databases,
(2) The GenAI Stack (a software stack from a collaboration of Docker, Neo4j, LangChain and Ollama),
(3) The use of LLMs to generate database queries that access knowledge graphs in Neo4j graph databases. Example dialogue applications include RAG with PDFs, RAG with Wikipedia, and RAG with knowledge graphs.

Biography : Graham Wilcock worked in industry with ICL at EU HQ in Luxembourg and with Sharp Corporation in Japan. He has a PhD from University of Manchester (UMIST) and is Adjunct Professor at University of Helsinki. He co-organized the first Linguistic Annotation Workshop in 2007 and received an IBM UIMA Innovation Award in 2008. Since 2016 he has worked mainly on conversational AI for robots. He developed the CityTalk dialogue system in which robots answer queries by searching knowledge graphs in graph databases. In 2018-19 he was Visiting Professor at Kyoto University, where he worked with the ERICA female humanoid robot. Publications :

<https://scholar.google.com/citations?user=AvHe1S4AAAAJ&hl=fr&oi=ao>

15h : Eric de la Cergerie

Title : Coupling KG and LLM: a few directions

Abstract : Large Language Models are becoming more and more ubiquitous but suffer from a lack of reliability (hallucinations), adaptability to specialized domains, and accessibility to up-to-date information. Several approaches have been proposed to combine LLM with reliable, well-structured, and up-to-date knowledge bases, in particular represented by knowledge graphs. We briefly survey some of these approaches, with a clear distinction between integration and interaction.

Biography : <http://alpage.inria.fr/~clerger/>

15h20 : Kristiina Jokinen

Title : Conversational Grounding, Trustworthy AI and Generative AI - Exploring LLMs for Active Healthy Aging

Abstract : Generative AI has radically changed dialogue modelling research and made fluent chatting a common model of interaction. In the context of practical applications, e.g. to support active healthy aging, ability to provide relevant and accurate information is important, and solutions to overcome GenAI's tendency to fabricate facts and produce false information are an active research area.

In this talk I will discuss social robot interaction which needs to balance trustworthiness and fluency, and present explorations of using Retrieval Augmented Generation (RAG) in ongoing work dealing with knowledge-graphs and dialogue systems. In particular, I will also emphasise the aspects of

Conversational Grounding, i.e. the need to ground utterances in the given physical context, so they do not only sound like correct fluent contributions, but convey meanings connected to the real-world context.

Biography : Kristin Jokinen is Senior Researcher at AI Research Center at National Institute of Advanced Industrial Science and Technology (AIST) in Tokyo Waterfront, and Adjunct Professor at University of Helsinki. She is Life Member of Clare Hall at University of Cambridge, and Member of the European ELLIS network. Since her PhD (UMIST, Manchester), her research has focussed on natural multimodal interaction and AI-based dialogue modelling, including social robotics, within the general framework of Constructive Dialogue Model. Together with Graham Wilcock she developed WikiTalk, which won the award for Best Robot Design at ICSR 2017.

Publications : <https://scholar.google.com/citations?user=IBn2-9UAAAAJ&hl=fr&oi=ao>

15h40 : Anthony Alcaraz

Title : **Towards Hybrid Reasoning: Assimilating Structure into Subsymbolic Systems**

Abstract : Recent advances in large language models (LLMs) show impressive fluency and adaptability, but they struggle with deeper reasoning requiring compositional generalization, sustained causal chains, and creative hypothesis generation.

Knowledge graphs provide structured representations that can address the reasoning gaps of LLMs, but have challenges with scale, noise, incompleteness, sparsity, and difficulty of query formulation.

The talk proposes a coordinated approach that leverages the strengths of both symbolic knowledge graphs and subsymbolic language models. This involves iterative analysis, modularization, parallel evidence retrieval, assimilating results, evaluation/explanation, and leveraging asynchrony and concurrency.

Strategies are proposed for blending symbolic and subsymbolic representations, including joint vector embeddings, injecting symbolic graph schemas, and differentiable graph programming. The overall goal is to achieve "structured reasoning with LLMs in chaotic contexts" by directing and constraining the language models with precise, modular knowledge components retrieved from graphs. This orchestration aims to ground the reasoning process for improved consistency, completeness and explainability.

Biography : Chief Product Officer at Fribl, an AI-powered recruitment platform committed to pioneering fair and ethical hiring practices.

With recruiting topping the list of concerns for many CEOs and companies spending upwards of \$4,000 on average per hire, the need for innovation is clear.

All too often, the arduous recruiting process leaves both employers and applicants frustrated after 42 days of effort with uncertainty if the right match was made...

At Fribl , we are leading the charge to transform this status quo, starting with reinventing the screening process using our proprietary GenAI technology enhanced by symbolic AI.

16h : Christian Dugast and Hermann Ney

Title: AppTek's experience in building ClimateGPT, a factual domain specific LLM

Abstract: We will share our experience in building a factual domain specific LLM addressing Climate Change topics. This system answers each question along three different perspectives (e.g. economical, societal and scientific). We have built 5 different models, 2 of them being foundational models, the 3 others fine-tuned based on Llama2. On climate relevant benchmarks, our fine-tuned 7B parameter model is on par with the original Llama2-70B-Chat model, needing 12 times less energy to answer a question than the reference 70B model, for the same result.

We will present our data selection process, as much for Continuous Pre-Training (CPT) as for Instruction Fine Tuning (IFT), focusing on the value of automatically generated data. On the inference side, we will discuss the limits of Retrieval Augmented Generation (RAG) to address the hallucination problem. Finally, we will review our understanding of the value of automatic testing on tasks (knowledge) versus human evaluation on prompt/completion pairs (arguments).

Biography: Dr Christian Dugast, lead architect Natural Language Understanding (NLU) at AppTek, works on entity recognition, information extraction, language understanding, reasoning, and large language models. He started his scientific career on acoustic modelling while at Philips Research after a PhD in Artificial Intelligence from the Université Paul Sabatier in Toulouse. After a decade in research, he stepped over to business, building from scratch Nuance Communications in Europe. Another decade later he came back to research enjoying cracking language-processing nuts again. Publications : <https://scholar.google.com/citations?user=5O52kS0AAAAJ&hl=fr&oi=ao>

16h20 : Hugues Sansen

Title : LifeLine

Abstract : Extension of the e-ViTA project, me-ViTA started with the idea of using LLM on a mobile, totally out of the cloud services to preserve the privacy of the user. Unfortunately, the early tests of LLMs on an Android mobile made them incompatible with the response time, sometimes in excess of 1mn, needed for a fluid dialogue. less than 300ms. To cope with the issue, we adopted a pragmatic approach and revived the 2016/2017 project LifeLine, a dialogue based graph construction of the life of the user. For that purpose, a KG dbms and a rule based questionnaire have been developed on Android. The graph mixes temporal milestones with other knowledge vertices. The questionnaire constructs vertices and edges according to the answers. A short term context groups all the knowledge vertices and edges needed for or constructed during a specific part of the questionnaire while the session context groups all those used during a session. When Google's Gemini is available in Europe, it will be used to construct more flexible questionnaires, from the knowledge collected previously in the graph. We expect to use the LLM as the ghost writer of the user's biography, making me-ViTA a time capsule to be given to the heirs. In addition to the biography which can be considered as a teaser, non judging questionnaires are introduced to help the person think about her diet and everything related to the health and morbid habits. We plan inject all the data delivered by the sensors deployed in the e-ViTA project. If our KG respects the graph theory, the questionnaire is closer to object oriented programming. Both techniques may seem close, it has been discovered that it is necessary to unlearn object oriented techniques for the building of the KG in order to make it generic, flexible and easier to interrogate

by browsing the graph. Roberta will be connected to the smartphone to offer the embodied interface by reusing the developments of the Empathic project.

Biography : Hugues Sansen learned AI in Marseille in 1986. Instead of adopting Prolog, he switched to Smalltalk with which he developed Ship-Planner, the first expert system for the conception of the loading plans for container ships, operational in 1989. He coupled an ANN with the inference engine by 1989/1990 through the fuzzy logic capability of the latter. The second AI winter forced him to reorient to nuclear safety in the military domain. He worked for Gemstone System, the very first Object Oriented DB, in Smalltalk and Java, the inventor of the smart card, Tibco Systems and Soprano Design, an Australian telco oriented software house. He co-founded Peer2Phone in early 2003 for which he wrote the first code of the P2P phone "GoSip". Unfortunately the company was not able to raise funds and closed in 2007. He started Shankaa the same year and participated to the V-Assist in the mid 2010s, an AAL project that precluded Empathic and e-VITA. He created Roberta Ironside, an embodied avatar, with the idea to offer an affordable physical companion. He joined the e-VITA Telecom Sud Paris team in December 2023 to develop a mobile embedded version of e-Vita, me-VITA.

16:35 : Discussions starting with comments from Michael McTear and Maria Inès Torres

Michael McTear is an Emeritus Professor at Ulster University with a special interest in spoken language technologies. He has been researching in the field of spoken dialogue systems for more than 20 years and is the author of several books, including *Spoken Dialogue Technology: Toward The Conversational User Interface* (Springer, 2004), *Spoken Dialogue Systems* (Morgan and Claypool, 2010), with Kristiina Jokinen, *The Conversational Interface: Talking to Smart Devices* (Springer, 2016), with Zoraida Callejas and David Griol, and *Conversational AI* (Morgan & Claypool 2020). His latest book *Transforming Conversational AI*, co-authored with Marina Ashurkina, is due to be released in March 2024. Michael has delivered keynote addresses and tutorials at many academic conferences and workshops, including SpeechTEK, Conversational Interaction, ProjectVoice, REWORK AI Assistant Summit, the European Chatbot Conferences, and the Altrusia Conversational AI and Customer Experience Summits. Currently Michael is involved in several research and development projects investigating the use of conversational agents in socially relevant projects such as mental health monitoring and home monitoring of older adults.

17h : Drinks and further informal discussions