

The Seminars on “Information Technology Outlook” – PhD Program in Computer Science and Mathematics



Valentina Tamma

Senior Lecturer

University of Liverpool, UK

Thursday July 16, 2026

10:00



Aula Hume

Dept. of Computer Science

LLMs Can Do Some but Cannot Do It All: Revisiting Feigenbaum's Bottleneck in the Gen AI Hype

Knowledge acquisition is the central bottleneck in building intelligent systems that make use of expert knowledge. Yet the problem, articulated by Edward Feigenbaum in 1977, is far from solved. Decades later, ontology engineering still faces the same bottleneck, but in a new form: requirements elicitation via Competency Questions (natural language questions that an ontology must be able to answer, used to specify and bound the knowledge it needs to model) remains manual, inconsistent, and difficult to audit, and most deployed ontologies lack documented requirements altogether. This talk examines what Large Language Models contribute to this problem, and what they do not. Drawing on empirical studies, from retrofitting CQs onto existing ontologies to cross-domain characterisation of LLM-generated requirements, I show that LLMs draft requirements at scale but exhibit distinct, domain-dependent generation profiles, and consistently fail to identify the implicit, inferential requirements that domain experts routinely provide.

I then present IDEA², an expert-in-the-loop workflow in which LLMs act as drafting partners rather than autonomous knowledge engineers, collaborating with domain experts who judge their output: requirements obtained with LLM assistance become a form of incremental agreement formation on ontological commitments. The workflow is illustrated through an industrial use case with Unilever, where an ontology is generated from the elicited requirements and validated both against the agreed competency questions and by ensuring that it does not admit unintended models. This second, adversarial form of validation has limits of its own: it can only catch what the underlying foundational ontology can express, causing silent semantic mismatches when independently engineered ontologies are aligned. Even if validation is imperfect, autonomy is not on the table: LLMs' role in ontology engineering is assistive, not autonomous.

***SHORT BIO:** Valentina Tamma is a Senior Lecturer in the School of Computer Science and Informatics at the University of Liverpool. A graduate of the Department di Informatica at the University of Bari, she received her PhD from the University of Liverpool. Her research spans knowledge representation, ontology engineering, and multi-agent systems, with a long-standing focus on ontology engineering and ontology alignment and negotiation. She is currently focussing on translational research aimed at lowering the barriers to ontology engineering through LLM-assisted, expert-in-the-loop methods. Her collaborators include industrial partners such as Unilever, and museums such as the British Music Experience Museum. Her work appears in venues including IJCAI, KR, ISWC, AAMAS, and EKAW. She has served on numerous Programme Committees and on the editorial boards of the Semantic Web Journal and Transactions on Graph Data and Knowledge. She was Area Chair for Representation, Perception, and Reasoning at AAMAS 2026 and is General Chair of ISWC 2026, to be held in Bari in October 2026*