

Tutorial proposal

The Evolution of Natural Language Understanding and Prediction Technologies: from Formal Grammars to Large Scale Machine Learning

- *A two-sentence description of the tutorial, suitable for inclusion in the conference registration brochure.*

This tutorial is aimed at providing the IJCAI community an overview of the deployed natural language technologies and their historical evolution. We review two fundamental problems involving natural language: the language prediction problem and the language understanding problem. The presentation focuses on the theory and algorithms used to build voiced/text-based human-computer interaction systems from the early automated directory assistance to today's smart-phone virtual assistants and semantic web search.

- *A two-paragraph description of the tutorial, suitable for a web page overview.*

Scientists have long dreamed of creating machines humans could interact with by voice. Although one no longer believes Turing's prophecy that machines will be able to converse like humans in the near future, real progress has been made in the voice and text-based human-machine interaction. After five decades of research, natural language understanding and prediction technology has become an essential part of many human-machine interaction systems (and even human-to-human: automated translation and speech-to-speech systems). There are now voice-based personal assistants, search and transactional systems for most smart phone platforms. The technology is pushed even further by the search engines which have evolved from simple keyword search to semantic search (they can now provide direct answers to a wide range of questions).

This tutorial is aimed at providing the IJCAI community an overview of the deployed natural language technologies and their historical evolution. We review two fundamental problems involving natural language: the language prediction problem and the language understanding problem. The presentation focuses on the theory and algorithms used to build voiced/text-based human-computer interaction systems from the early automated directory assistance to today's smart-phone virtual assistants and semantic web search.

- *A detailed outline of the tutorial.*
 - Introduction and history of language understanding systems
 - Text-based: Knowledge/Grammar vs Data/Statistical systems

- Voice-based: Automated Directory Assistance, Call routing, Directed Dialog, Personal Assistants
- DARPA programs and evaluations
- Brief intro to Automated Speech Recognition
 - Acoustic modelling
 - Speech recognizer architecture
 - Current ASR performance and analytics
- Language prediction/modelling
 - Introduction & applications
 - Statistical formulation & naïve estimation
 - LM interpolation/smoothing techniques
 - Pruning vs full LM estimation; engineering challenges
 - LM coverage assessment: hit rates, perplexity, word error rates.
 - Modern approaches: (deep/recurrent) neural networks, exponential models
- Language Understanding
 - Problem characterization wrt input query space and output semantic space
 - Client only vs. client-server architectures; system components
 - Statistical/machine learning language understanding
 - Semantic data annotation & classification
 - Parsing based semantics & hybrid NLU systems
 - Integration of NLU modules with Dialog, Information retrieval and Action managers
 - Deployment issues and evaluation
- Intro to question/query answering, pro-active and deep learning technologies
 - Semantic search engines
- Future developments
- *A characterization of the potential target audience for the tutorial, including prerequisite knowledge.*

Potential target audience of the tutorial: AI/Machine learning researchers who would like an introduction to natural language technologies as well as practical considerations on building and deploying such systems

Prerequisite knowledge: Elementary probability/statistics. Some pattern recognition & machine learning knowledge may also be useful.

The following objectives are best served by the proposed tutorial:

Introduce expert non specialists to an AI subarea (natural language technologies).

Motivate and explain a topic of emerging importance for AI.

- *A description of why the tutorial topic would be of interest to a substantial part of the IJCAI audience.*

The last decade brought a proliferation of smart phones and wearable devices along with applications that can be driven by voice. They not only recognize the words being said but also their meaning and are many times able to give answers back or ask for additional information. Such systems are based on technologies from many AI domains: machine learning, knowledge representation, reasoning and retrieval, human-machine interfaces, etc. The proposed tutorial describes how theories and methods from these AI domains can be used to build intelligent systems a human can converse with.

- *Brief resume of the presenter:*

Nicolae Duta

Microsoft New England Research and Development Center

One Memorial Drive, Cambridge, MA, 02142, USA

Tel: (857) 4536243

niduta@microsoft.com

Nicolae Duta received the B.S. degree in applied mathematics from the University of Bucharest (Romania) in 1991, the D.E.A. degree in statistics from the University of Paris-Sud (France) in 1992, the M.S. degree in computer science from the University of Iowa in 1996 and the Ph.D. degree in computer science and engineering from Michigan State University in 2000. He is currently an applied scientist in the Applications & Services Group at Microsoft in Cambridge, MA. From 2006 to 2013 he was a member of the Natural Language Understanding and Language Modeling groups at Nuance Communications, Burlington, MA where he developed the first generation of voice-based personal assistants for smart phones. From 2000 to 2005 he was a scientist in the Speech and Language Processing department at BBN Technologies, Cambridge, MA. He also held temporary research positions at INRIA-Rocquencourt (France) in 1993 and Siemens Corporate Research (Princeton, NJ) from 1997 to 1999. He is a member of IEEE and his current research interests include computer vision, pattern recognition, language understanding, automatic translation, machine and biological learning.

Background in the tutorial area: 14 years of research in Language Modelling & Understanding at BBN Technologies, Nuance Communications and Microsoft. Authored 2 US/international patents and 3 pending patent applications in language technologies.

Citation to an available example of work in the area-- ideally, a published tutorial-level article on the subject
Nicolae Duta: Natural Language Understanding and Prediction: from Formal Grammars to Large Scale Machine Learning, Fundamenta Informaticae, Volume 131(3), 2014, pp. 425-440

Evidence of teaching experience: I previously presented machine learning tutorials at

[European Conference on Computer Vision 2002](http://www.it-c.dk/events/eccv02/tutorials.html) (<http://www.it-c.dk/events/eccv02/tutorials.html>)

International Conference on Advances in Pattern Recognition 2001 (<http://www.informatik.uni-trier.de/~ley/db/conf/icapr/icapr2001.html>)

Evidence of scholarship in AI/computer science: See my [Google Scholar profile](#)