

Education

2018–2022 PhD in Robotics and Autonomous Systems, Centre for Doctoral Training in the Edinburgh Centre for Robotics, Supervisors: Prof. Oliver Lemon, Dr.

Ioannis Konstas

Thesis: Visually Grounded Representation Learning using Language Games for Embodied AI

2017–2018 MSc by Research in Robotics and Autonomous Systems, Centre for Doctoral Training at the Edinburgh Centre for Robotics, Full marks with distinction

Thesis: “Dialogue Attention Buffer: A Neural Memory Module for Task-based Conversational Systems”

2014–2017 Master Degree in Computer Science, Knowledge Engineering and Machine Intelligence, University of Bari, Italy, Full marks with distinction

Thesis: “Converse-Et-Impera: a Conversational Content-based Recommender System based on Hierarchical Deep Reinforcement Learning”

2011–2014 Bachelor Degree in Computer Science, University of Bari, Italy, Full marks

with distinction

Thesis: “Content-based Recommender Systems based on Linked Open Data”

Work experience

February

2022 – Now

Tenure-track Research Fellow, Heriot-Watt University

I am one of the first hires of the National Robotarium (22.4M research centre), and

one of the Heriot-Watt Bicentennial Research Leaders. My research agenda focuses on developing novel machine learning models for artificial agents that can learn by interacting with the world and with other agents. I am currently the faculty advisor for EMMA, the Heriot-Watt team competing in the Amazon Simbot Challenge.

July 2021 –

October 2021

Applied Scientist Intern, Meta AI

Design and implementation of novel multi-modal techniques for video understanding and reasoning. The project was completed under the supervision of Dr. Seungwhan Moon.

March 2021 –

June 2021

Applied Scientist Intern, Amazon Alexa AI

Design and implementation of a state-of-the-art Deep Learning architecture for Embodied AI tasks. The project was completed under the supervision of Prof. Gaurav

Sukhatme and Dr. Jesse Thomason.

June 2020 –

Feb. 2021

NLP Software Engineer, AlanaAI Limited

Design and implementation of knowledge-based and data-driven components for the Alana Conversational AI. I was responsible for the Natural Language Understanding pipeline of the dialogue system as well as semantic-aware components based on ontologies and linked open data.

Oct. 2019 –

Aug. 2020

NLP Consultant, The AvantGarde Group Srl

Design and implementation of a hierarchical text-classification model based on BERT for financial documents written in Italian. The solution, evaluated on a real-world dataset, guaranteed a boost in terms of classification performance (measured as F1-measure) of 4.9% compared to feature-based machine learning models.

March 2018 –

Nov. 2018

Amazon Alexa Prize 2018, Heriot-Watt Universityl

Design and implementation of several components for the open-domain conversational agent “Alana” ranked 3rd in the Alexa Prize Challenge 2018. I was responsible for the main architecture of the socialbot as well as the Natural Language Understanding pipeline and several retrieval-based bots.

Apr. 2017 –

Sept. 2017

NLP Consultant, Plusimple Srl

Design and implementation of a context-aware search engine able to understand the user query intent and to reformulate the query according to intent and user-specific information collected during the platform usage time.

Research fundings

Feb. 2022 Equally Safe Online, EPSRC Grant (EP/W025493/1)

Nominated as collaborator for the EPSRC grant awarded to Prof. Verena Rieser about online gender-based violence. As part of this collaboration, I will supervise

the development of multi-modal deep learning models for the proposal.

Feb. 2022 –

Feb. 2023

Multi-modal Curriculum Learning for Embodied AI, Cirrus EPSRC

Project proposal to have access to the Cirrus EPSRC Tier-2 National HPC facility required to work on the project “Multi-modal Curriculum Learning for large-scale Embodied AI”. The proposal will cover GPU usage costs for an estimate of £113K (based on March 2022 prices).

Nov. 2020 –

Nov. 2021

Massively Multimodal Dialogue, Cirrus EPSRC

Project proposal to have access to the Cirrus EPSRC Tier-2 National HPC facility required to work on the project “Massively Multimodal Dialogue (MOD): The

Football Commentaries Challenge”. The proposal will cover GPU usage costs for an estimate of £320K (based on November 2020 prices).

Academic contributions

Jan. 2018 –

Now

Paper reviewing

Contributed in the paper reviewing process (as first or second reviewer) of several national and international conferences such as: ACL, EACL, IJCAI, AAAI, AI*IA, iCities, WebConf, LOD, WGNT, ALVR, and LANTERN.

Talks

August 2022 LingTech Summer School

Delivered a course on “Deep Learning for Conversational AI” at the LingTech Summer School in Lecce, Italy. The lesson presented a survey of the literature on

dialogue systems describing both traditional as well as cutting-edge deep learning methods for implementing state-of-the-art dialogue systems.

March 2021 -

November

2021

Samsung AI Seminar Series

Invited by Samsung AI (Toronto & Cambridge) to give a presentation at the Samsung AI “Multi-modal Machine Learning” seminar series about my work on interactive grounded language learning and Embodied AI.

July 2019 Glasgow Alexa Meet-up

Invited by Screenmedia to give a presentation at the Glasgow Alexa Meet-up about my experience with the Amazon Alexa Prize challenge and additional information about the design of Alana.

December

2018

Heriot-Watt Semantic Web Lab, Exploiting Semantic Web Technologies

in Open-domain Conversational Agents

I gave a lecture at the Heriot-Watt Semantic Web Lab seminar session where I presented how the Alana Conversational AI uses Linked Open Data to understand and reason over the entities mentioned during the conversation with its users.

May 2018 Amazon Alexa Prize Summit Science Fair

I was the Heriot-Watt University’s representative at the Amazon Alexa Prize Summit

Science Fair panel discussion where I talked about my experience as a PhD student, what motivated me to start a career in STEM, and I answered questions about Computer Science from local Seattle students.

Honours

2022 Shortlisted for Best PhD Thesis Award 2022, The Scottish Informatics and Computer Science Alliance (SICSA)

2021 Member of Nova

Member of the Nova excellence network which connects the top 3% of talent from 72 countries around the world.

2013–2014 Best graduated student award, University of Bari, Italy

Personal skills and competences

Languages

Italian (Native) English (Fluent)

Technical skills

Programming

languages

(preferred)

Python, Java Programming

languages

(known)

C, C++, Go, Haskell, Lua

Machine

Learning

libraries

scikit-learn, TensorFlow, Torch,

PyTorch, Weka

Semantic

Web

technologies

OWL, RDF, SPARQL

Database

Management

Systems

MySQL, MongoDB,

PostgreSQL, DynamoDB

Publications (see Google Scholar for the full list)

Francisco Javier Chiyah-Garcia et al. "Exploring Multi-Modal Representations for Ambiguity Detection & Coreference Resolution in the SIMMC 2.0 Challenge". In: arXiv preprint

arXiv:2202.12645 (2022).

George Pantazopoulos, Alessandro Suglia, and Arash Eshghi. "Combine to Describe: Evaluating Compositional Generalization in Image Captioning". In: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop. 2022, pp. 115–131.

Alessandro Suglia et al. "An empirical study on the generalization power of neural representations learned via visual guessing games". In: arXiv preprint arXiv:2102.00424 (2021).

Alessandro Suglia et al. "Embodied bert: A transformer model for embodied, language-guided visual task completion". In: arXiv preprint arXiv:2108.04927 (2021).

Olivia Foulds et al. "Predicting perceptual speed from search behaviour". In: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in

Information Retrieval. 2020, pp. 1989–1992.

Alessandro Suglia et al. “CompGuessWhat?!: A Multi-task Evaluation Framework for

Grounded Language Learning”. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics. 2020.

Alessandro Suglia et al. “Imagining Grounded Conceptual Representations from Perceptual Information in Situated Guessing Games”. In: Proceedings of the 28th International

Conference on Computational Linguistics. International Committee on Computational Linguistics. 2020.

Pierpaolo Basile et al. “Bridging the gap between linked open data-based recommender systems and distributed representations”. In: Information Systems 86 (2019), pp. 1–8.

Pierpaolo Basile et al. “Deep learning and hierarchical reinforcement learning for modeling a conversational recommender system”. In: Intelligenza Artificiale 12.2 (2018), pp. 125–141.

Amanda Cercas Curry et al. “Alana v2: Entertaining and informative open-domain social dialogue using ontologies and entity linking”. In: Alexa Prize Proceedings (2018).

Alessandro Suglia et al. “A deep architecture for content-based recommendations exploiting recurrent neural networks”. In: Proceedings of the 25th conference on user modeling,

adaptation and personalization. 2017, pp. 202–211.