



universität
wien

Faculty of Business,
Economics and Statistics

Workshop on the occasion of the 60th birthday of Immanuel Bomze



Organized by:

Irene Klein, Vera Lehmwald, Markus Leitner, Ivana Ljubic, Mario Ruthmair, Werner Schachinger



universität
wien

Workshop on the occasion of the 60th birthday of Immanuel Bomze

SHORT-LIVED HYPES AND FAILURES IN OR Panel Discussion

- Florian Jarre
- Joaquim Júdice
- Abdel Lisser
- Ivana Ljubic
- Marco Locatelli
- Fabio Schoen
- Fabio Tardella
- Tamás Terlaky

Moderator: Walter Gutjahr



SHORT-LIVED HYPES AND FAILURES IN OR

by

Ivana LJUBIC

What is AI?

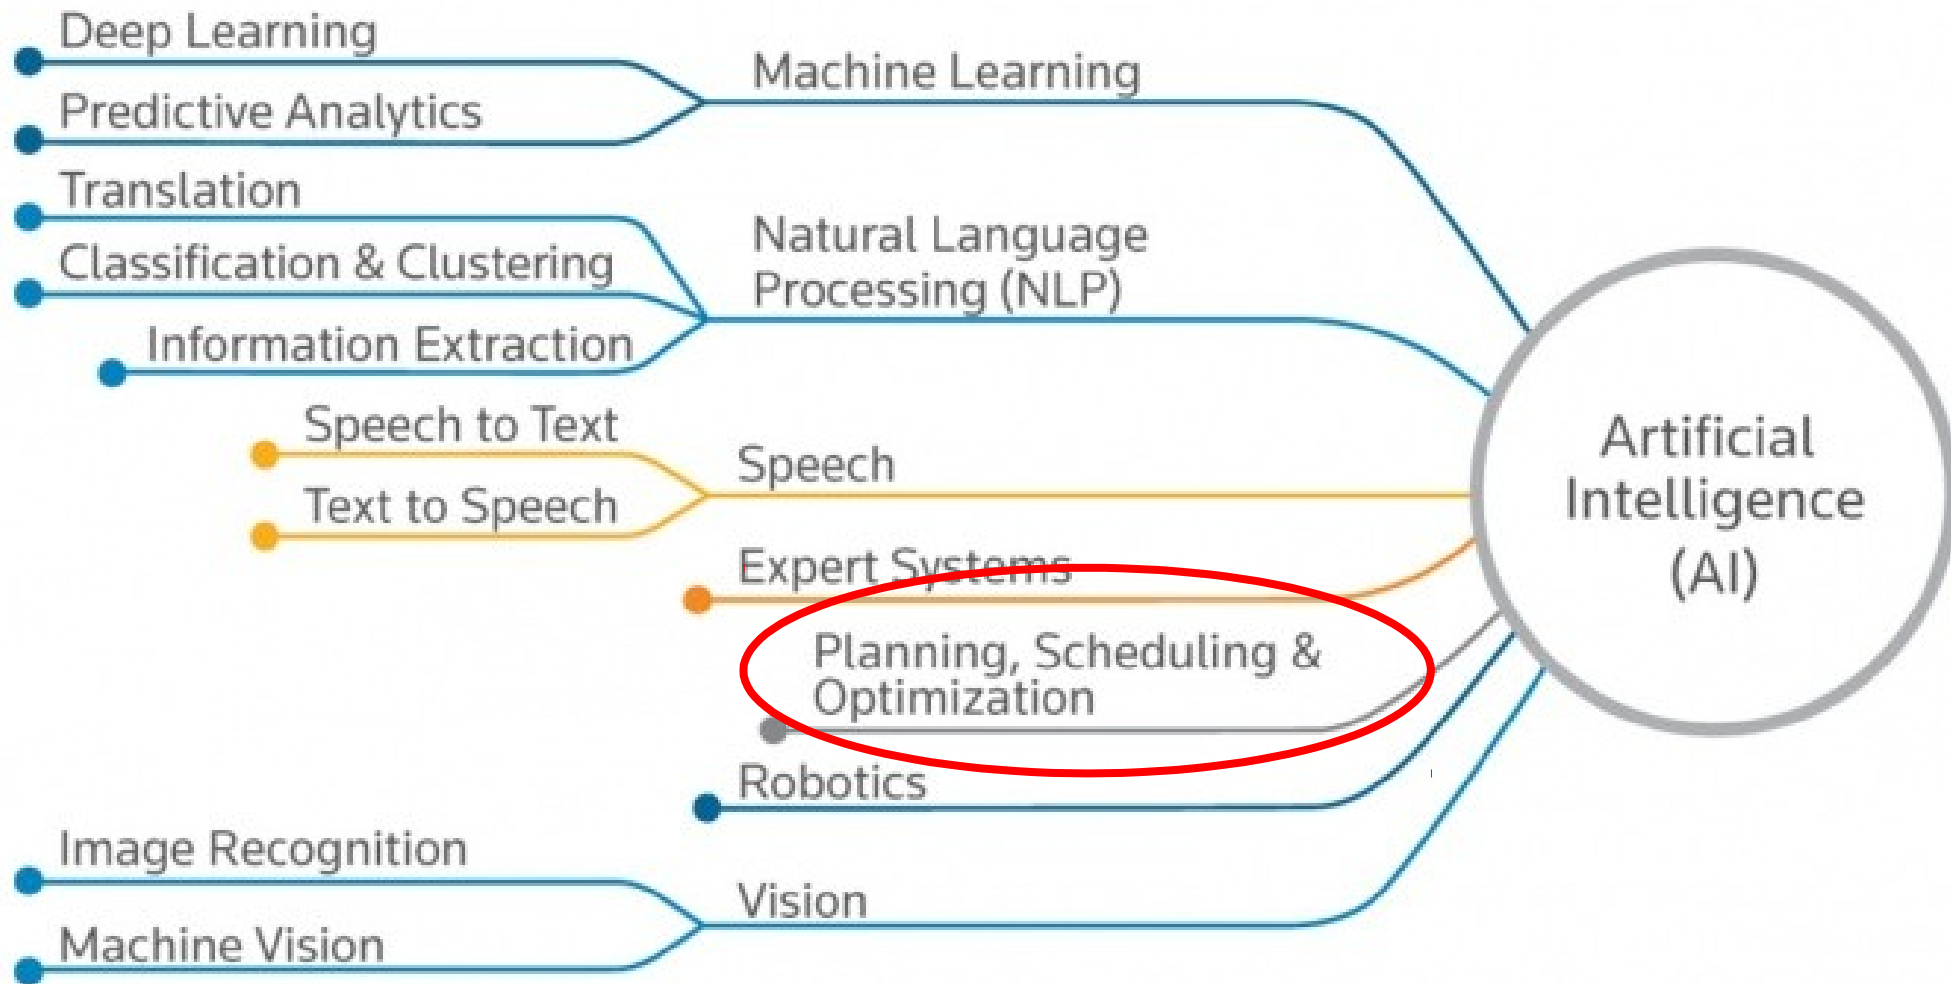
AI is whatever hasn't been done yet



Larry Tesler (invented copy-paste at Xerox), a computer scientist working on human interaction


The definition of AI changes as we get used to previous advances!

AI Taxonomy? Really?



Yes, AI was a short lived hype... Several times...

1951: Minsky and Edmonds have built first electronic learning system, a device **Snarc** (simulated the functioning of a simple neural network).



The book **Perceptrons** (Minsky&Papert, 1969): limitations of NNs at the time. NNs fall out of favor. **First AI winter (in the 70's).**

Expert systems: **second AI winter (end of '80s)**

Alpha Go (**DeepMind**) beats Lee Sedol in **March 2016**:
Deep Learning starts a **new AI hype cycle**

„We know more than we can tell“
Michael Polanyi





Edmond de Belamy, sold at **\$432,500** at Christie's (**Oct 2018**)

A.I. Researchers Are Making More Than \$1 Million, Even at a Nonprofit



Artificial intelligence experts are commanding eye-popping salaries. Including a signing bonus, OpenAI paid its top researcher, Ilya Sutskever, more than \$1.9 million in 2016.

About OpenAI

OpenAI is a non-profit AI research company, discovering and enacting the path to safe artificial general intelligence.




OpenAI (endowment: \$1bn) paid its top researcher, Ilya Sutskever, more than **\$1.9 million** in 2016.

DeepMind costs for 400 employees totaled \$138 million in 2016: **\$345,000** per employee, including researchers and other staff.

Neural Combinatorial Optimization with Reinforcement Learning

Irwan Bello, Hieu Pham*, Quoc V. Le, Mohammad Norouzi, Samy Bengio*

04 Nov 2016 (modified: 25 Jan 2017) ICLR 2017 conference submission Readers:  Everyone [Show Revisions](#)

Abstract: This paper presents a framework to tackle combinatorial optimization problems using neural networks and reinforcement learning. We focus on the traveling salesman problem (TSP) and train a recurrent neural network that, given a set of city coordinates, predicts a distribution over different city permutations. Using negative tour length as the reward signal, we optimize the parameters of the recurrent neural network using a policy gradient method. We compare learning the network parameters on a set of training graphs against learning them on individual test graphs. Without much engineering and heuristic designing, Neural Combinatorial Optimization achieves close to optimal results on 2D Euclidean graphs with up to 100 nodes. Applied to the KnapSack, another NP-hard problem, the same method obtains optimal solutions for instances with up to 200 items. These results, albeit still far from state-of-the-art, give insights into how neural networks can be used as a general tool for tackling combinatorial optimization problems.

TL;DR: This paper presents a framework to tackle combinatorial optimization problems using neural networks and reinforcement learning.

Conflicts: google.com

Keywords: Reinforcement Learning, Deep learning