Converse-Et-Impera: Exploiting Deep Learning and Hierarchical Reinforcement Learning for Conversational Recommender Systems

Claudio Greco, Alessandro Suglia, Pierpaolo Basile, and Giovanni Semeraro
claudiogaetanogreco@gmail.com, alessandro.suglia@gmail.com, {firstname.lastname}@uniba.it

Background

Conversational recommender systems assist online users in *decision making* tasks by supporting a *goal-oriented interactive process* with the aim of narrowing down the user interests until the desired item is obtained

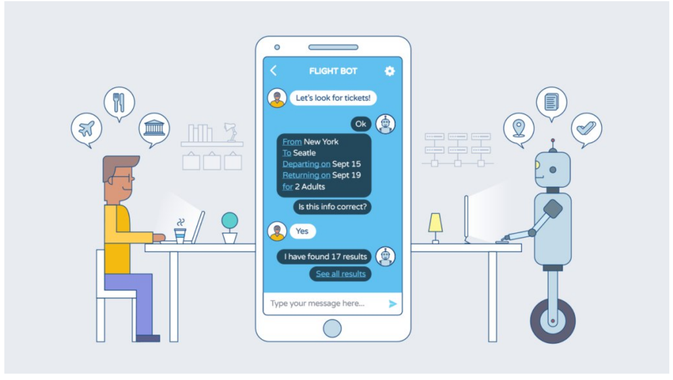

Research questions

1. Are goal-driven dialogue systems able to effectively support a conversation with the user by dividing it in simpler goals?
2. Are neural network models able to extract relevant evidences from textual documents?
3. Are sequence-to-sequence models able to generate personalized responses according to the user preferences?

Methodology

Converse-Et-Impera is equipped with:

- the meta-controller , which is responsible for selecting a goal in a given turn
- the controller , which selects an action given the state and the current goal

Applies a dedicated representation module according to the specific goal:

- chitchat, used to manage general conversations with the user
- qarecs, used to provide personalized suggestions according to the extracted evidences

Trained by using a two-stage training strategy:

1. supervised learning to learn a preliminary policy for the agent
2. reinforcement learning to allow the agent to learn from its experiences with a simulated user

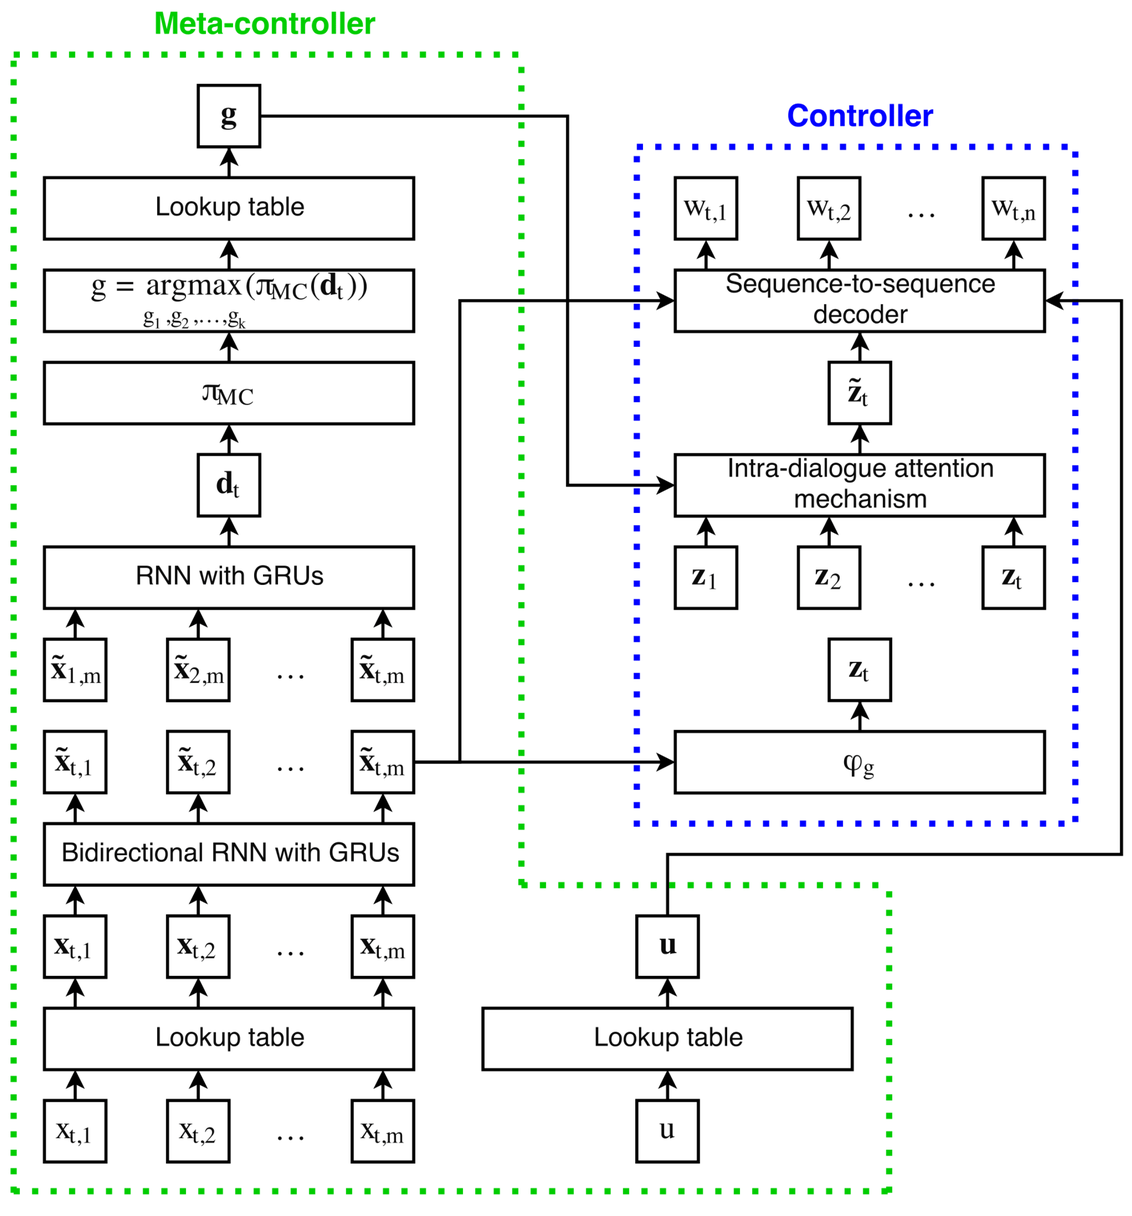

Results

Experimental evaluation conducted on synthetic dialogues generated from Movielens 1M (ML1M) and MovieTweetings (MT)

Supervised Learning

Configuration	MC precision		BLEU		F_1-measure		Per-user F_1-measure	
	MT	ML1M	MT	ML1M	MT	ML1M	MT	ML1M
Random	0.332	0.333	0	0.001	0	0	0	0.002
Seq2seq	NA	NA	**0.791**	0.839	0.022	0.01	0.044	0.029
CEI	**1.0**	**1.0**	0.784	**0.851**	**0.110**	**0.07**	**0.170**	**0.108**

Reinforcement Learning

ML1M

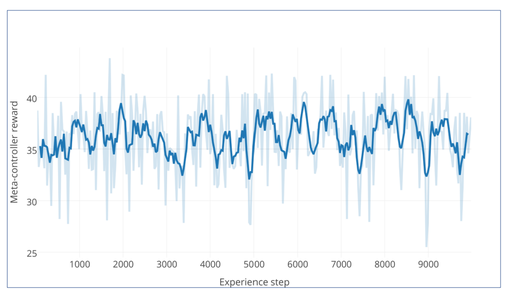

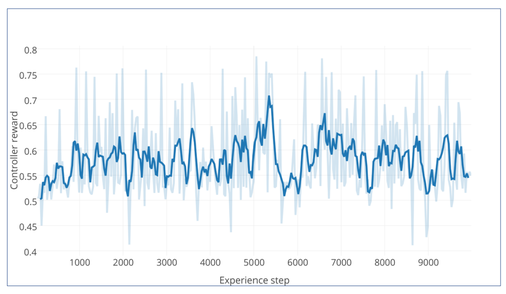

MT

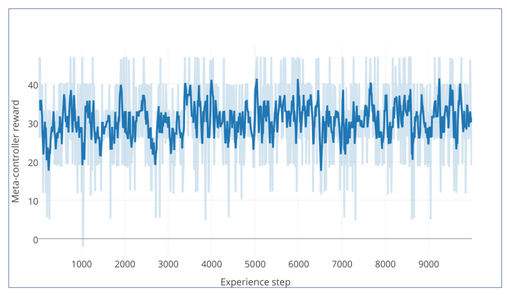

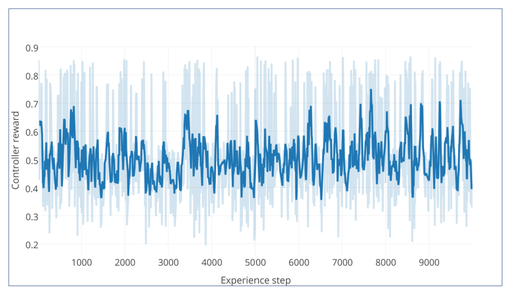

16th AI*IA, Bari, November 14-17, 2017