Available online at www.sciencedirect.com

ScienceDirect

Computer Speech & Language 47 (2018) 93–111

www.elsevier.com/locate/csl

A methodology for turn-taking capabilities enhancement in Spoken Dialogue Systems using Reinforcement Learning☆

Hatim Khouzaimi[a,b,1,*], Romain Laroche[a,2], Fabrice Lefèvre[b]

[a] *Orange Labs, Châtillon, France*
[b] *CERI-LIA, Avignon, France*

Received 23 August 2016; received in revised form 30 June 2017; accepted 7 July 2017
Available online 20 July 2017

Abstract

This article introduces a new methodology to enhance an existing traditional Spoken Dialogue System (SDS) with optimal turn-taking capabilities in order to increase dialogue efficiency. A new approach for transforming the traditional dialogue architecture into an incremental one at a low cost is presented: a new turn-taking decision module called the *Scheduler* is inserted between the Client and the Service. It is responsible for handling turn-taking decisions. Then, a User Simulator which is able to interact with the system using this new architecture has been implemented and used to train a new Reinforcement Learning turn-taking strategy. Compared to a non-incremental and a handcrafted incremental baselines, it is shown to perform better in simulation and in a real live experiment.

Keywords: Spoken Dialogue Systems; Turn-taking; Incremental dialogue; Reinforcement Learning

1. Introduction

Floor management in currently deployed dialogue systems is generally pretty basic: the user and the system must wait for each other to finish their respective utterances before taking the floor at each turn. This way of handling turns has the advantage of simplicity and maintainability. However, during the last 15 years, an important research thread has shown that this is not an optimal manner of managing turn-taking in human/computer dialogue (Aist et al., 2007; Skantze and Schlangen, 2009; El Asri et al., 2014) and that incremental[3] dialogue systems (Schlangen and Skantze, 2011; El Asri et al., 2014) offer a better user experience. A dialogue system is incremental when it is able to process the user's utterance as it is spoken which makes it able to take the floor anytime during the dialogue. Also, the user

☆ This paper has been recommended for acceptance by Prof. R. K. Moore.

* Corresponding author:
E-mail address: hatim.khouzaimi@askhub.io (H. Khouzaimi), romain.laroche@maluuba.com (R. Laroche).
[1] AskHub, Paris, France.
[2] Maluuba, Montreal, Canada.

[3] From an architectural point of view, incrementality refers to the ability to process increments separately. On the other hand, from the functional point of view, it designates the ability to process and act on speech as it is spoken. In this paper, we use the second frame since it is a feature that systems must have in order to improve their turn-taking capabilities.

http://dx.doi.org/10.1016/j.csl.2017.07.006

is allowed to interrupt the system. The concept of incremental processing has been proposed for the first time to build incremental compilers (Lock, 1965). According to Kilger and Finkler (1995), it has been introduced in the field of natural language processing in Wirén (1992). The idea of building incremental dialogue systems is directly inspired by human conversation since when people talk to each other, the listener most often tends to understand the speaker as she speaks even guessing the rest of her sentence before its end (Levelt, 1989; Tanenhaus et al., 1995).

In order to meet industrial constraints (Pieraccini and Huerta, 2005; Paek and Pieraccini, 2008; Laroche, 2010; Asri, 2016), in this paper, a new methodology for turn-taking capabilities enhancement at a low cost is proposed.

During the last decade, dialogue systems have been used to solve various tasks already, both in research and industry. As far as incremental dialogue systems are concerned, most of them are built from scratch with an initial intention of providing them with incremental capacities (Dohsaka and Shimazu, 1997; Allen et al., 2001; Schlangen and Skantze, 2011). In this paper, a new approach for transforming a traditional dialogue system into an incremental one at a low cost is introduced. A new turn-taking decision module is added to the traditional architecture (without modifying the dialogue manager): the *Scheduler*, firstly introduced in Khouzaimi et al. (2014a; 2014b).

In the field of dialogue systems, collecting data corpora is very costly. Therefore, user simulation techniques are widely used (Eckert et al., 1997; Schatzmann et al., 2006; Pietquin and Hastie, 2013; Laroche and Genevay, 2016). In this work, an incremental User Simulator is described (Khouzaimi et al., 2016). It is based on a new approach that is aimed to generate ASR instability (Selfridge et al., 2011). The implemented task is a slot-filling personal agenda management.

Since its first application to dialogue management (Levin and Pieraccini, 1997; Singh et al., 1999), RL (Sutton and Barto, 1998) has become one of the leading frameworks in the field. Here, it is applied to learn turn-taking decisions by taking dialogue duration and task completion as the only components of the reward function (Khouzaimi et al., 2015a), resulting in significantly more efficient and more robust dialogue systems. Compared to other techniques using supervised learning (Meena et al., 2013), no labelling effort is required here. Also, there is no need to make assumptions like the fact that the system should minimise gaps and overlaps (Sacks et al., 1974); the global dialogue quality is the only function to maximise.

Finally, a live experiment where users interact with a smart home (application called the Majordomo) has been run in order to validate this approach. It is shown that the RL strategy significantly improves the task completion ratio. Also, the subjective evaluation is slightly in favour of this new strategy.

Section 2 describes the related work and the positioning of this paper, then Section 3 presents the new methodology for transforming a traditional dialogue system into an incremental one. Based on that, Section 4 describes the simulated environment and the implementation of the rule-based turn-taking strategy. Finally, Section 5 introduces the RL model, Section 6 describes the live experiment as well as the associated results and Section 7 concludes and sheds light on planned future work.

2. Related work

Improving dialogue systems' turn-taking capacities has been an active research thread during the last two decades. Existing contributions can be classified in three design categories (handcrafted, supervised learning and RL) as well as two evaluation categories (indirect and direct) which results in six system categories. They are depicted in the following, given the turn-taking model they use and the way they evaluate it. This idea is synthesised in Table 1.

Some papers use rule-based models or supervised learning that are not directly evaluated with users through real interactions. Aist et al. (2007) introduce a handcrafted incremental setup that improves user satisfaction while reducing dialogue duration. Nevertheless, pre-recorded utterances that are perfectly understandable by the system were

Table 1
Related work classification according to the model and the evaluation method.

	Handcrafted	Supervised learning	RL
Indirect evaluation	Aist et al. (2007)	Meena et al. (2013); Zhao et al. (2015)	Jonsdottir et al. (2008); Selfridge and Heeman (2010); Lu et al. (2011); Dethlefs et al. (2012); Kim and Banchs (2014)
Direct evaluation	Raux and Eskenazi (2009); Skantze and Schlangen (2009); Ghigi et al. (2014); Zhao et al. (2015)	Paetzel et al. (2015)	This paper