



HAL
open science

Talking to Machines: do you read me?

Lina M. Rojas-Barahona

► **To cite this version:**

Lina M. Rojas-Barahona. Talking to Machines: do you read me?. Computation and Language [cs.CL].
Universite de Lorraine, 2024. tel-04620199

HAL Id: tel-04620199

<https://hal.science/tel-04620199v1>

Submitted on 21 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License



HDR: Talking to Machines: do you read me?

Part I: Scientific contributions

Lina María ROJAS BARAHONA

MEMOIRE

pour l'obtention de l'

Habilitation de l'Université de Lorraine

(Spécialité Informatique)

Ecole doctorale IAEM Lorraine

Rapporteurs

Prof. Olivier Pietquin. Université de Lille/Cohere.

DR. Chloé Clavel. INRIA-Paris.

Prof. Giuseppe Riccardi. University of Trento

Examineurs

Prof. Yannick Toussaint. LORIA/Université de Lorraine.

Prof. Slim Ouni. LORIA/Université de Lorraine.

MC. Jean-Charles Lamirel. LORIA/Université de Lorraine.

Invités

DR. Claire Gardent. LORIA/CNRS.

CR. Christophe Cerisara. LORIA/CNRS.

To my daughter Sophía, the bright light of my life. To my husband Hervé Vultrin, who has been always embracing me with his love, support and patience. To my brother Oscar and my sister Beatriz who have been there for me and gave me my beloved nephews and niece: Alejandro, José Francisco and Maríapaz. To my aunt Beatriz Barahona, who has always stood by my side being my second mother. I also dedicate this work as many others in the past to my parents, Aquiles Rojas and Ruth Barahona, who are always beside me, supporting me and giving me their strength even though they are enjoying the splendid experience of eternity. Thank you all for being the source of my inspiration.

To my beautiful country, my beloved Colombia.

Contents

1	Introduction	1
2	A Glance to the Research on Dialogue	3
2.1	Why is human conversation difficult?	4
2.2	Preliminary Approaches	5
2.3	Task Oriented Dialogue Systems	5
2.3.1	Definitions	6
2.3.2	Statistical Dialogue Systems	6
2.3.3	End-to-End Task-Oriented Systems	9
2.4	End-to-End Dialogues	10
2.5	Pre-trained Language Models	11
2.6	Conversational Question Answering	11
2.7	Positioning My Contributions in the State-of-the-Art	12
3	Contributions to Task-Oriented Dialogues	15
3.1	Language Understanding	16
3.1.1	NLU in a Serious Game	16
3.1.2	Spoken Language Understanding and Few-Shot Learning	23
3.2	Dialogue Manager	32
3.2.1	Bayesian Inverse Reinforcement Learning	33
3.2.2	Is the User Enjoying the Conversation?	37
3.2.3	Imitation Learning	43
3.2.4	Hierarchical Imitation Learning	47
4	Contributions to Conversational Question Answering (QA) and Other Contributions	53
4.1	Detection of Ellipsis and Co-reference in conversational corpora	54
4.1.1	Active learning (AL)	55
4.1.2	Experiments	57
4.1.3	Results	58
4.2	Question Rewriting	59
4.2.1	Annotations	60
4.2.2	Evaluation on Question Rewriting (QR)	60
4.2.3	Evaluation on Conversational Question Answering	62
4.3	A conversational QA corpus grounded in Wikidata	64
4.4	Other Contributions: Graph Embeddings	64

5	Data Collection, Annotation and Frameworks	73
5.1	Corpora	73
5.1.1	The French Portmedia Corpus	73
5.1.2	The French Emospeech Corpus	74
5.1.3	Annotating Posts Following Cognitive Behavioural Therapy Principles	76
5.1.4	Conversational Question Answering with Rewriting	77
5.1.5	Conversational Question Answering in French	78
5.1.6	Conversational QA grounded in Wikidata	78
5.2	Dialogue Frameworks	78
5.2.1	PyDial	78
5.2.2	Dialport	78
5.2.3	Conversational Search for General Knowledge	79
6	Scientific Project	81
6.1	LLMs for Task-Oriented Dialogue	82
6.1.1	Benchmarking and Evaluation	82
6.1.2	Interpreting LLMs	82
6.1.3	Retrieval Augmented Generation	83
6.1.4	Controlling Decoding	83
6.2	Multimodal task-oriented dialogues	84
6.2.1	Exploiting weak speech signals in the reward function	84
7	Conclusion	87

List of Acronyms

TOD	Task-Oriented Dialogues
DSs	Dialogue Systems
MDP	Markov Decision Process
POMDP	Partially Observable Markov decision process
NLP	Natural Language processing
SLU	Spoken Language Understanding
NLU	Natural Language Understanding
DST	Dialogue State Tracking
NLG	Natural Language Generation
ASR	Automatic Speech Recognition
DM	Dialog Manager
ML	Machine Learning
QA	Question Answering
IL	Imitation Learning
LR	Logistic Regression
SVM	Support Vector Machines
CNN	Convolutional Neural Network
LSTM	Long-Short Term Memory
GNN	Graph Neural Network
DIP	Domain Independent Parametrisation
IL	Imitation Learning
IRL	Inverse Reinforcement Learning
GNN	Graph Neural Network
QA	Question Answering
CoQAR	Conversational Question Answering with Rewriting
KGConv	Knowledge-base Conversations
KG	Knowledge Graph

CQA	Conversational Question Answering
CS4GK	Converstional Search for General Knowledge
DIANA	DIalogue in NATural Language
CBT	Cognitive Behavioural Therapy
CIFRE	Conventions industrielles de formation par la recherche
NL4XAI	Interactive Natural Language Technology for Explainable Artificial Intelligence
COBRA	Conversational Brains
ITN	Innovative Training Networks
RL	Reinforcement Learning
LLMs	Large Language Models
VQA	Visual Question Answering
RAG	Retrieval Augmented Generation
AI	Artificial Intelligence
IQ	Interaction Quality

Chapter 1

Introduction

‘Desde entonces no me gané un centavo que no fuera con la máquina de escribir, y esto me parece más meritorio de lo que podría pensarse, pues los primeros derechos de autor que me permitieron vivir de mis cuentos y novelas me los pagaron a los cuarenta y tantos años, después de haber publicado cuatro libros con beneficios ínfimos. Antes de eso mi vida estuvo siempre perturbada por una maraña de trampas, gambetas e ilusiones para burlar los incontables señuelos que trataban de convertirme en cualquier cosa que no fuera escritor.’
— *Gabriel García Marquez, Vivir para Contarla.*

‘Since then I haven’t earned a penny other than with the typewriter, and this seems to me more meritorious than one might think, since the first author rights that allowed me to live from my tales and novels were paid to me when I was in my forties, so many years after having published four books with just tiny benefits. Before that my life was always disturbed by a tangle of traps, tricks and illusions to circumvent the countless lures that tried to turn me into anything other than a writer.’ — *Gabriel García Marquez, Living to tell the tale.*

In this dissertation I would like to guide the reader to the research on dialogue but more precisely the research I have conducted during my career since my PhD thesis. Starting from modular architectures with machine learning/deep learning and reinforcement learning to end-to-end deep neural networks. Besides my work as research associate, I also present the work I have supervised in the last years. I proposed four PhD thesis Conventions industrielles de formation par la recherche (CIFRE) that Orange accepted to fund. Therefore, I could co-supervise four PhD candidates: Timothy Garwood supervised by Claire Gardent at CNRS, Thibault Cordier supervised by Fabrice Lefevre at the University of Avignon, Sebastien Montella supervised by Alexis Nasr at the university of Aix-Marseille, Léo Jacqmin supervised by Benoit Favre at the University of Aix-Marseille. During 5 years I was head of the industrial research project on dialogue, DIalogue in NATural Language (DIANA), which gave me the opportunity of supervising the work of the young researcher Quentin Brabant, other experimented researchers, a developer as well as students in internship and apprenticeship. The deliverables of DIANA project gather open-sourced datasets and neural models as well as scientific publications.

I review briefly the state of the art and highlight the open research problems on conversational agents in Chapter 2. Afterwards, I present my contribution to Task-Oriented Dialogues (TOD) in

Chapter 3, both as research associate and as the industrial supervisor of CIFRE theses. I discuss conversational QA in Chapter 4. Particularly, I present the work of two PhD candidates Thibault Cordier and Sebastien Montella; as well as the work of the young researcher Quentin Brabant. I present the scientific project in Chapter 6. Finally I present the conclusions in Chapter 7.

Chapter 2

A Glance to the Research on Dialogue

‘Please, then,’ said Alice, ‘how am I to get in?’
‘There might be some sense in your knocking,’ the Footman went on without attending to her [...]’ He was looking up into the sky all the time he was speaking, and this Alice thought decidedly uncivil. ‘But perhaps he can’t help it,’ she said to herself; [...] –How am I to get in?’ she repeated, aloud.
–[...] the Footman continued in the same tone, exactly as if nothing had happened. ‘How am I to get in?’ asked Alice again, in a louder tone.
‘Are you to get in at all?’ said the Footman. ‘That’s the first question, you know.’
It was, no doubt: only Alice did not like to be told so. ‘It’s really dreadful,’ she muttered to herself, ‘the way all the creatures argue. It’s enough to drive one crazy!’ — *Lewis Carroll*

The origins of dialogue systems date from 1966 when the Eliza chatterbox was presented (Weizenbaum, 1966). Eliza was the automated psychoanalyst that let us dream about intelligent systems able to converse as humans. However, Eliza was a simple *template-based* approach, limited by its poor understanding as well as the lack of expressivity and adaptability. The research on dialogue has come a long way since then. Several solutions have been proposed from *regular expressions* and *symbolic approaches* (formal grammars and formal logics) (Larsson and Traum, 2000; McTear, 2002) to *statistical approaches* (Pieraccini et al., 1992; Williams and Young, 2007; Young et al., 2010; Zhang et al., 2020b), which are data-driven techniques that use either machine learning or deep learning.

The availability of big data as well as the advances in processing units made deep learning approaches feasible and promising (Cuayáhuitl et al., 2015; Daubigney et al., 2013; Vinyals and Le, 2015; Sordani et al., 2015; Wen et al., 2017b; Serban et al., 2017), reviving the dream of creating artificial agents that can easily converse to people. We have already obtained promising results. For instance, we know that machines can learn optimal strategies for simple tasks in small domains (TOD)(Weisz et al., 2018a; Zhu et al., 2020). Moreover, we are already treating open domain dialogues by asking questions to online encyclopaedias such as Wikipedia (conversational reading comprehension) (Choi et al., 2018b; Reddy et al., 2018) and we are able to predict the best answer in chitchats (end to end neural approaches). The very recent breakthrough ChatGPT (Ouyang et al., 2022) confidently generates apparently coherent responses for a great

amount of domains and tasks. A variety of methods has been studied by the research community on dialogue during the last decade. From modular architectures to end-to-end neural networks. In this chapter I will describe some these approaches.

2.1 Why is human conversation difficult?

Who has not experienced the frustration of talking to an automatic system in a call centre. Typically, these systems struggle to understand. They are repetitive because they are unable to rectify misunderstandings. Users then must start all over from scratch. In the worst-case users need to call again, then they try hard to fool the system until the call is finally answered by a human. Conversations with automatic systems are unnatural because they do not deal correctly with misunderstandings, they do not adapt to novel situations and they constraint humans' great communication skills.

Dialogue-Acts and discourse obligations: A dialogue can be seen as a sequence of turns, in which every speaker takes a turn to speak and to contribute to the conversation. The philosopher (Austin, 1975a) stated that speakers perform actions while conversing and named these actions speech acts. Examples of these actions are : informing, requesting, offering, promising, answering, persuading, convincing, etc. The research community nowadays call these actions Dialogue Acts or Communicative Acts (Bunt et al., 2010). Adjacency-pairs are pairs of dialogue-acts in conversation. For example, after a question in a conversation, the speaker is waiting for an answer. After an offer the speaker is waiting for an acceptance or a rejection. These represent discourse obligations in human conversation.

Coreferences and Ambiguity: Moreover, humans can refer to concepts that were mentioned previously in the conversation. Indeed, a fluent conversation avoid repetitions. For instance, if you are talking about the president of France, you can say "Emmanuel Macron" the first time you mentioned him, later you can choose to say "*the president*". Moreover, if you want to further give your opinion about a recent political proposition he has made, you can say "I disagree with *his* retirement policy". These are well studied linguistic phenomena that made conversation difficult.

Another aspect is the ambiguity, the same sentence can mean different things in different contexts. "It's cold in here" can be understood as a request to close the window inside a closed room or it may mean "It's cool in here" in summer. In winter instead it could mean "I can't stand the weather; it is too cold".

Grounding: speakers are always checking that they are following each other. For instance, let us suppose you are receiving instructions about where to place a box in a room. If there are many similar boxes, and the instruction is "move the box to the right of the desk", you will probably ask "which one?". Then the instructor will provide more precise information such as "the yellow rectangular box". This coordination or mutual agreement is known as grounding, mutual knowledge or shared knowledge(Clark and Brennan, 1991).

Planning: The model Beliefs, Desires and Intentions (BDI) as the primary mental attitudes of an agent was first introduced by Bratman (Bratman, 1987). The beliefs are the agent's model

2.2 Preliminary Approaches

of the world. Desires, in turn, represent how the agent would like the world to be in the future; while intentions are the structured plan the agent has decided to perform. The agent interacts with the world by performing actions and by perceiving aspects of it, including changes which result from its own actions. Perceptions will influence the beliefs of the agent, while actions may change aspects of the world. This model was at the origin of modern Natural Language Understanding (NLU), in which the aim is to detect user's intentions or *intents*. However, the term intent usually means a dialogue-act with a set of concepts or a combination of them in semantic labels.

All these inherent characteristics in natural dialogue make implementing automated systems a very difficult task.

2.2 Preliminary Approaches

A range of approaches emerged in the history of Dialogue Systems (DSs), they were classified in conformity with their Dialog Manager (DM) (Allen et al., 2001; Churcher, 1997). According to this classification, ordered by increasing complexity, the simplest of these is the *finite-state scripts*, also called dialogue grammars, followed by *slot-filling*, *plan-based* and *agent-based models*. In a finite-state script the dialogue is represented as a script of prompts for the user. In slot-filling, questions are asked in order to enable the system to fill the necessary slots to perform a task. Conversely, plan-based theories claim that utterances infer acts that are part of a plan, thus, the system tries to identify users' underlying plan, collaborates in accomplishing that plan and responds appropriately. Agent-based models are at the highest level of complexity. They consider planning, executing, and monitoring operations in a dynamically changing world, possibly involving multi-modality. Examples of agent-based models are: the logic-based approaches, which uses inference engines of a higher complexity that in some cases are semi-decidable (in some cases the system will never halt), as well as reinforcement learning approaches, which need a large number of interactions to converge.

2.3 Task Oriented Dialogue Systems

Conversational agents have gained great interest in both academy and industry in the last decades. Typically, available conversational agents have been designed for the task of information-seeking. These agents act as a natural language interface to a database. First, the system tries to fill the constraints to query a database by inquiring the user. Then it retrieves the items that fulfill users' constraints and finally it communicates the results to the user in natural language. For instance, a person could call the system to check train timetables, she would provide the departure and arrival city as well as the departure date and time. Then, the system would inform her about the available trains.

TOD main goal is to complete a task in collaboration with the user (Pieraccini et al., 1992; Young, 2002; Rieser and Lemon, 2011; Young et al., 2013a). Examples of tasks are to search information about a restaurant, to reserve a hotel or to buy train tickets.

2.3.1 Definitions

As introduced in Section 2.1, *dialogue-acts* or communicative acts are the actions performed by the speakers when uttering sentences (e.g. Informing, Asking, Confirming, Greeting, etc.) (Austin, 1975b). A *domain* is formally defined in an *ontology* as a list of *slots* with their valid *values*. The most common task, the information seeking task, is usually modelled as a *slot-filling* data-query problem in which the system requests constraints to the user and proposes items that fulfil those constraints in a database. *The action* or *intention* is composed by a predicate: the dialogue-act, and a set of arguments: the slot-value pairs. For instance, let us suppose the user has uttered “I would like a restaurant in the center of town please”, this will be translated in the semantic form: $\text{inform}(\text{type} = \text{restaurant}, \text{area} = \text{center})$. This semantic representation is usually called *flat-semantics* because there is not hierarchy in the concepts of the ontology.

2.3.2 Statistical Dialogue Systems

One approach to automatic dialogue is to use Reinforcement Learning (RL) to select the system’s action (Levin et al., 2000; Litman et al., 2000). As in a Chess game, a dialogue involves two players, in which each of them takes turns to play. The system should decide its move by considering the environment (the other player’s moves) and the rewards is either win or lose the game. Dialogue is then formulated as an optimisation problem, in which the environment is the user action and the user’s feedback is the reward. The final goal of the system is then to maximise the accumulated reward at long run (Rieser and Lemon, 2011). The optimal *policy*, π , is a function that takes as argument the current state s and returns the optimal action a .

Markov decision process

Dialogue can be formalised as a Markov Decision Process (MDP), which is a tuple $M = (S, A, T, \gamma, R)$ where:

- S : A set of possible states that represent the dynamic environment.
- A : A set of possible actions.
- $T : S \times A \times S \rightarrow [0, 1]$ is a transition probability function. For any action $a \in A(s)$ taken in a state $s \in S$, the probability of transiting to the next state s' is given by $T(s, a, s')$.
- γ : A discounting factor in the range of $[0, 1]$, which controls the prediction horizon of the algorithm.
- R : The reward function that specifies the reward gained at every state. It contains the information that guides the agent towards the goal. R is a function of the state that is bounded in absolute value by R_{max} .

A stationary *policy* is a map $\pi : S \rightarrow A$ and the discounted infinite-horizon expected reward for starting in state s and following policy π thereafter is given by the value function $V^\pi(s)$ that satisfies the following *Bellman Equation*:

$$V^\pi(s) = R(s) + \gamma \sum_{s'} T(s, \pi(s), s') V^\pi(s') \quad (2.1)$$

2.3 Task Oriented Dialogue Systems

The discounted infinite-horizon expected reward for starting in state s , taking action a and following policy π thereafter is given by the *Q-function* $Q^\pi(s, a)$ that satisfies the following equation:

$$Q^\pi(s, a) = R(s) + \gamma \sum_{s'} T(s, a, s') V^\pi(s') \quad (2.2)$$

A policy π is optimal in M if, for all $s \in S$:

$$\pi(s) = \arg \max_{a \in A} Q^\pi(s, a) \quad (2.3)$$

Likewise $Q^*(s, a, \mathbf{R})$ is the optimal Q-function of the optimal policy π^* for a known reward function \mathbf{R} .

Partially Observable Markov decision process

Dialogues can be modelled as an optimisation problem with Partially Observable Markov decision process (POMDP)s. It simulates the inherent dynamic behaviour of human conversations while deals with the uncertainty of spoken language (Roy et al., 2000; Williams and Young, 2007; Young et al., 2013b).

A POMDP can be seen as a continuous-space Markov decision process (MDP) in terms of policy optimisation where the states are the belief states, which is partially observable. POMDPs have been proposed for spoken dialogue systems because the system is never sure about the user beliefs because of speech recognition errors due to noisy or spoken language disfluencies and hesitations (Roy et al., 2000; Young et al., 2013a). Since the state is uncertain, it is called the belief state $b(s)$. An example of a POMDP dialogue system is presented in Section 2.3.2. The task of predicting the $b(s)$ at a given time t is known as the task of Dialogue State Tracking (DST). The policy learning algorithm receives as input the $b(s)$ and returns the optimal policy π^* and a given time.

The belief state b_t is a vector encoding a probability distribution over the different goals, dialogue acts and concepts that are discussed in the dialogue. In the same way, the dialogue action a_t is a vector encoding a probability distribution over the possible agent dialogue actions.

Hierarchical Reinforcement Learning

MDP models have been proven to be inefficient for solving complex tasks. These models have trouble overcoming the cold start problem and/or suffer from the curse of dimensionality (Barto and Mahadevan, 2003). This pattern was also observed with models proposed recently (Mnih et al., 2013; Duan et al., 2016). To overcome this issue, (Parr and Russell, 1998) proposed to specify a hierarchy of tasks and to reuse parts of the state space across many sub-tasks, which can greatly improve both learning speed and agent performance.

The notion of temporal abstraction, in which a policy can be decomposed into sub-policies by calling temporally extended sub-tasks was first proposed by (Sutton et al., 1999). In order to consider hierarchical architectures with temporally extension, we have to generalise the MDP to the semi-Markov Decision Process (SMDP) (Parr and Russell, 1998) where actions can take a variable amount of time to complete. This creates a division between primitive actions that span over only one action and composite actions that involve an execution of a sequence of primitive

actions. This introduces a policy μ over options that selects option o in state s with probability $\mu(s, o)$, o 's policy might in turn select other options until o terminates and so on. The value function for option policies can be defined in terms of the value functions of the flat Markov policies (Sutton et al., 1999).

(Cuayáhuitl, 2009) was the first to propose Hierarchical Reinforcement Learning (HRL) based on the MAXQ algorithm for dialogue decomposition, making use of hierarchical abstract machines (Parr and Russell, 1998). However, the tabular approach of this algorithm prevents the efficient approximation of the state space and the objective function. To overcome this limitation (Budzianowski et al., 2017), uses Gaussian process, which provides uncertainty estimates which can be used to speed up learning and achieve more robust performance. Some recent work such as (Tang et al., 2018) aims to discover automatically sub-goals hierarchy in dialog.

Reward Functions for Dialogue Systems

Previous works on RL for learning dialogue strategies typically use reward functions that penalise long dialogues, returning a final positive reward for task completion or user satisfaction (Levin et al., 2000; Litman et al., 2000; Roy et al., 2000; Young et al., 2010; Rieser and Lemon, 2011). This might be an intuitive reward function for slot-filling applications, such as train ticket or restaurant reservation, in which usually customers know exactly what they want, and they expect to be accurately informed by the system as fast as possible.

However, this reward function might be inappropriate in other situations or for distinct users. For instance, a user might want more advice without caring about the duration of the call. This is especially true in tutorial dialogues, where learners usually have to complete a task and may not know exactly how to do it.

Architecture

The basic elements of a RL based spoken statistical dialogue system are shown in Figure 2.1. The words recognised by the speech recognition are converted to an abstract representation (the user dialogue acts) by the semantic parser, also known as semantic decoder or NLU. These user dialogue acts are then processed by a belief-state tracker which maintains a *dialogue state* s . This is typically a set of variables denoting the *slots* that the system must fill-in to complete the user's goal. For example, in a restaurant information system the slots might be `food` for the type of food offered and `area` for the location, and the state s might record the current value and confidence level of each slot. From the state, a *belief state* b (usually just a sub-set of the state vector) is extracted and an action a is decided based on a dialogue *policy*. The set of possible actions will include requesting new slot values, confirming already filled slot values and accessing the application for information. Once the appropriate action is determined, it is converted to a textual message m and then rendered by a speech synthesiser.

Deep Learning (DL) has been used to implement dialogue components such as semantic decoder or NLU (Rojas Barahona et al., 2016), belief tracker (Daubigney et al., 2013; Mrkšić et al., 2017) and generator (Wen et al., 2015). Deep Reinforcement Learning is used for policy learning (or DM) (Cuayáhuitl et al., 2015), such as deep q-network (DQN) (Casanueva et al., 2017a), the actor-critic algorithm (Su et al., 2017) and actor-critic with experience replay (ACER) (Weisz et al., 2018a). Benchmarks comparing these algorithms across different domains and different

2.3 Task Oriented Dialogue Systems

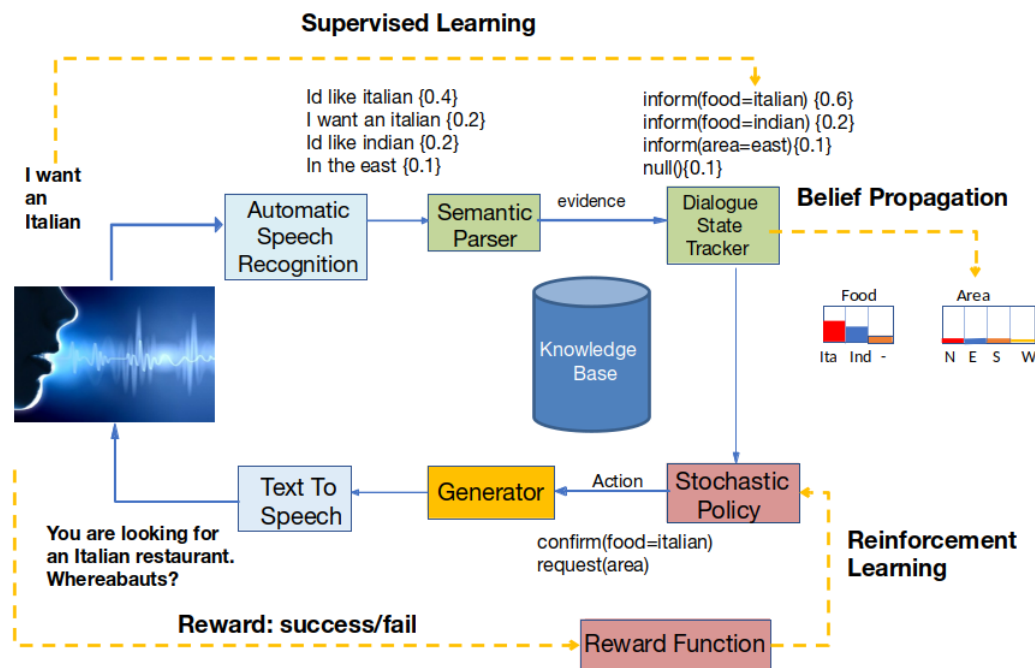


Figure 2.1: Basic elements of a statistical spoken dialogue system

environments have been published in (Casanueva et al., 2017a). It is worth noting that ACER has been trained on around 1000 action spaces and not on summary actions as the others (Weisz et al., 2018a). These algorithms are available in the open-source dialogue framework PyDial (Ultes et al., 2017b).

2.3.3 End-to-End Task-Oriented Systems

One preliminary approach to end-to-end dialogue is pipeline-based. It replicates the classical dialogue architecture (Figure 2.1), with the main difference that each module is a deep neural network. One example of this approach is the seminal work of (Wen et al., 2017b). This neural pipeline was later improved in (Wen et al., 2017a) through a latent variable model for learning the distribution of system actions. Although, authors claimed this approach to be end-to-end, modules are not trained jointly into a single learning unit that can be optimised by gradient-based methods such as back-propagation. Conversely, each module is trained separately in a cascade fashion where the outputs of one model are the inputs to the next one.

2.4 End-to-End Dialogues

Three approaches to end-to-end dialogue system are identify: retrieval-based, generative-based and the combination of both. These approaches are truly end to end, which means the model learns through gradient-based optimisation. *Retrieval-based* approach treats dialogue as an information retrieval problem (Lowe et al., 2015; Wu et al., 2018), in which there is a set of candidate responses from which one is selected as system response, here dialogue is evaluated as an information retrieval problem (in terms of precision/recall). *Generative-based* models (Vinyals and Le, 2015; Sordoni et al., 2015; Serban et al., 2016; Goo and Chen, 2018) on the contrary use natural language generation to generate the system response and dialogue is evaluated as as a generation problem (in terms of comparison with multi-references for instance, with BLEU score). Both approaches have been used for chit-chats or casual conversations. Retrieval-based models however have been also used for task-oriented solutions(Lowe et al., 2015). The combination of both, can use generation to paraphrase the retrieved answer. Another way is to compare generative and retrieved responses. The interested reader can find more details about recent end-to-end approaches in (Ni et al., 2022).

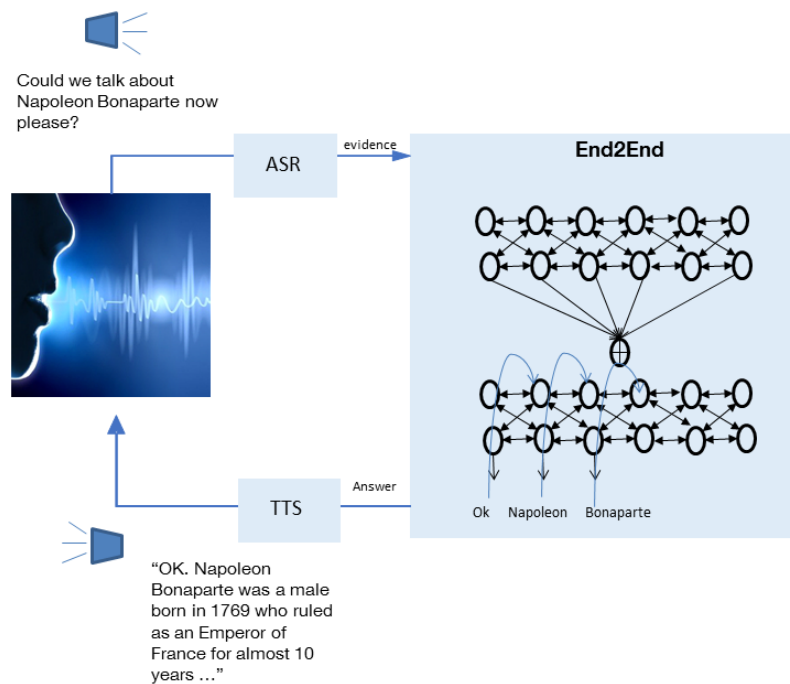


Figure 2.2: E2E neural architecture

Yet these approaches neglect the fact that dialogue is dynamic and highly dependent on the environment. Moreover, they do not consider the metrics usually used for evaluating dialogues such as task completion and user satisfaction. Generative approaches usually generate fluent answers, which are sometimes incoherent with the dialogue context. They can generate hallucinations, distortions or repetitions. Retrieval-based approaches are limited to a list of candidate responses, which needs to be created in advance, yielding answers without any context agreement. To overcome these limitations, (Sankar and Ravi, 2019) proposes a combination of reinforcement-learning and generative models. The successful ChatGPT is an example of generative model, carefully trained to follow instructions that also learns to rank its responses.

2.5 Pre-trained Language Models

2.5 Pre-trained Language Models

With the success of Language Models, Deep Learning and Transformers (Vaswani et al., 2017), recent solutions proposed pre-trained models with self-supervision. These models outperformed the state-of-the-art in different Natural Language processing (NLP) tasks. They have been initially trained on a large quantity of texts. For instance, BERT has been trained on 800M words and 2,500M words of the BooksCorpus and Wikipedia respectively. GPT-3 is a huge pre-trained model with 175B of parameters, trained on large quantity of data.

Recently, these pre-trained models have been also trained on conversations. For example, BlenderBot (Shuster et al., 2022) is an encoder decoder which has been trained on 1.5B of Reddit comments. DialoGPT is just a decoder that has trained on 147M of dialogues extracted from Reddit. (Santra et al., 2021). ConveRT has been trained on 727M of dialogs. ChatGPT is presumably the result of finetuning GPT3 with carefully curated instructions with reinforcement learning for correctly rank the generated responses (i.e. learning to rank) (Ouyang et al., 2022).

The trend nowadays is to initialise deep learning models with a pre-trained one and then fine-tune them for a specialised task. However, fine-tuning with little data will often degrade the initialised weights. Therefore, recent optimisation for fine-tuning are : prompting (Lester et al., 2021), prefixes (Li and Liang, 2021) and adapters (Hu et al., 2021).

Despite very promising, the great limitation is that pre-trained models are still static. If there are changes in the language or in the World, these changes will not be reflected in the model. Think about the BERT models pre-trained before the Covid pandemic, they do not contain any representation for the bunch of new vocabulary that emerged during the pandemics: Covid-19, Moderna, Pfizer, sanitary vaccination pass, test PCR, etc. It is worth noting that training these large models is computationally costly and they require Hyper-performance Computing (HPC). A research path is how to keep these models up to date avoiding catastrophic forgetting and reducing the carbon impact necessary during training.

2.6 Conversational Question Answering

Research on conversational question answering has gained increasing interest (Saha et al., 2018; Reddy et al., 2018; Choi et al., 2018b). It consists in sequences of question-answer pairs related to a document or to a knowledge graph. Complex sequential question answering (CSQA) contains open-domain conversations that treat complex linguistic phenomena such as co-references, ellipses, incompleteness (or under specification) as well as logical, comparative and quantitative reasoning (Saha et al., 2018). Two corpora containing discussions about a paragraph of a Wikipedia document have been made public, namely question answering in context (QuAC) (Choi et al., 2018b) and a conversational question answering challenge (CoQA) (Reddy et al., 2018). In addition, a workshop devoted to this topic has first created in 2017, search conversational artificial intelligence (SCAI), with the participation of academics and industrials¹.

Paragraph-based Question-Answering can be seen as a problem of reading comprehension, in which given a candidate document and a question, it finds the correct answer in the document. These systems use attention mechanisms (Seo et al., 2017) and memory networks (Weston et al., 2014).

¹<https://scai.info/>

2.7 Positioning My Contributions in the State-of-the-Art

Most of my work concerns task-oriented dialogue, which involves the different components presented in Section 2.3. I worked on NLU, DM and more recently on Natural Language Generation (NLG). I have adapted neural models for distinct dialogue components that are presented in Chapter 3. Since the past five years, I have also explored open-domain Conversational Question Answering (CQA) (Section 2.6) which I present in Chapter 4, with predictive and generative (encoder-decoder) models for question rewriting, reading-comprehension and knowledge-graph question generation. As an extension of the last approach, I have also explored graph-embeddings (Section 4.4) and graph verbalisation with language models (Montella et al., 2023). The contributions published the second half of this year or the work in progress are not included in this document.

My work has always followed the state-of-the-art at the time of publication. For a detail comparison of each contribution summarised in this manuscript with the related work of its time, we invite the curious reader to check the corresponding publications. To provide a brief positioning of some contributions, the work on data collection, NLU, dialogue management and human evaluation I made for the EmoSpeech corpus was the first of its kind: a set of dialogues (12 distinct types of dialogue) in a Serious Game and in French. Moreover, we were among the first to propose in 2013 data-augmentation with back-translation and distributional representations for balancing biased models (Gardent and Rojas-Barahona, 2013).

The work on inverse reinforcement learning followed on the seminal work of Ng et al. (2000). It differed from previous work Paek and Pieraccini (2008); Chandramohan et al. (2011); El Asri et al. (2012); Boularias et al. (2010) because of the Bayesian Inverse Reinforcement Learning (IRL) algorithm that was applied for learning the tutor (i.e. system) reward function from experts. Furthermore, experts were taken from twelve distinct types of conversations in a serious game that were Human-Human (i.e. The EmoSpeech corpus) and not Machine-Machine (not from simulated conversations).

I proposed to enrich a subset of the corpus CoQA (Reddy et al., 2019) releasing the corpus CoQAR² with up to three out-of-context question paraphrases per question in conversations (Chapter 4, Section 4.2). Unlike previous work, these paraphrases were made by professional English native annotators instead of using crowd-sourcing, guaranteeing fair earning and working conditions. We compared CoQAR to CANARD (Elgohary et al., 2019) that provided only one question rewriting per question in QuAC (Choi et al., 2018b). I experimented with RoBERTa for answer extraction in both datasets CoQA and CoQAR with a without rewriting. Surprisingly solving the context through rewritten questions confuses RoBERTa, which is already good to solve co-references by itself, specially in short dialogue contexts as in these datasets.

I also utilised contextual embeddings (DistilBERT (Sanh et al., 2020), TransformersXL (Dai et al., 2019)) for estimating the reward function in long dialogues. I pointed out at that time the limitation of BERT-like models to deal with long contexts, which, besides the notable improvements of recent years, is still an open research problem. We also explored neural generation models such as BART (Lewis et al., 2019) and T5 (Raffel et al., 2020) for contextual question generation, question rewriting (Section 4.2) and graph verbalisation (Montella et al., 2022, 2023). I explored together with Sebastien Montella and Johannes Heinecke structural adapters for graph verbalisation (Montella et al., 2023) just after low rank emerged (Hu et al., 2021) as a recom-

²<https://github.com/Orange-OpenSource/COQAR>

2.7 Positioning My Contributions in the State-of-the-Art

mended way to optimally fine-tune LLMs.

After the advent of Large Language Models (LLMs), I am now questioning the performance of these models in complex tasks that required planning, such as dialogue. First, we need an evaluation methodology to assess the performance of LLMs in these tasks. Then we need to compare LLMs based reasoning (Wei et al., 2022; Yao et al., 2023, 2022) with Reinforcement Learning (RL) approaches, and explore recent trends for learning complex strategies such as algorithm distillation (Laskin et al., 2022). I talk about these research paths in Chapter 6.

Chapter 3

Contributions to Task-Oriented Dialogues

‘José Arcadio Buendía pasó los largos meses de lluvia encerrado en un cuartito que construyó en el fondo de la casa para que nadie perturbara sus experimentos. Habiendo abandonado por completo las obligaciones domésticas, permaneció noches enteras en el patio vigilando el curso de los astros, y estuvo a punto de contraer una insolación por tratar de establecer un método exacto para encontrar el mediodía.’ — *Gabriel García Márquez, Cien años de soledad.*

‘José Arcadio Buendía spent the long months of the rainy season shut up in a small room that he had built in the rear of the house so that no one would disturb his experiments. Having completely abandoned his domestic obligations, he spent entire nights in the courtyard watching the course of the stars and he almost contracted sunstroke from trying to establish an exact method to ascertain noon.’ — *Gabriel García Márquez, One Hundred Years of Solitude.*

As introduced in Section 2.3, task-oriented dialogue systems search to accomplish a task. This task can be for instance, *information seeking*, in which the system search for items in a database according to the constraints obtained through natural language interaction. Thus, these constrains can be given by the user (*‘I am looking for a restaurant’*) or can be enquired by the system (*‘which price range?’*, *‘where about?’*). Once the desired items are retrieved, the system informs the results back to the user. This kind of dialogue has rich interactions with dialogue acts such as: *informing*, *requesting*, *clarifying*, *rectifying*. State-of-the-art systems are centred on the information seeking task for a variety of domains: hotels, restaurants, touristic attractions, trains, flights, taxis, etc. Unfortunately, these dialogues are very specific and have difficulties to generalise to new domains and to more complex tasks (i.e. beyond information seeking).

I present in this Chapter my contributions to TOD dialogues (Figure 2.1). I start by presenting my work on NLU in Section 3.1 and on Dialogue Management in Section 3.2. My contributions to NLU concerns the definition of an annotation schema for the *French corpus EmoSpeech*, in which 12 distinct dialogues were integrated in a *serious game*. The dialogues involved various characters representing *the system* and *the player* and they were triggered at different levels of the game quest. Dialogues were modelled as an information seeking task, in which the system is always providing information to the player. The set of dialogue acts and goals are introduced in Section 3.1, as well as the models that were trained. Data augmentation through paraphrases via back translation, dictionaries, lexical resources and distributional semantics is also presented

in Section 3.1.1. Moreover, the task of spoken language understanding is studied by using deep neural models and few-shot learning through risk minimisation in Section 3.1.2. Finally, I present the work of the PhD candidate Sebastien Montella on graph embeddings in Section 4.4.

The contributions on Dialogue management are presented in Section 3.2. I first present my own work on inverse reinforcement learning to find the implicit reward followed by humans in the EmoSpeech dialogue corpus. I also present a reward estimation by using deep learning. Finally, I present the work of the PhD candidate Thibault Cordier on hierarchical imitation learning.

3.1 Language Understanding

Natural language understanding (NLU) is the task of mapping natural language sentences to semantic concepts. As a component of a spoken system, it would map utterances to a semantic representation that describes user intentions. This representation can be a combination of dialogue acts¹ (e.g., greeting, request, inform, acknowledgement, confirm, etc.) and concept-value pairs, which are usually defined in a knowledge-base that describes the domain (e.g, Depart_City="New York"). This section presents the semantic annotation of the French EmoSpeech corpus, the proposed Machine learning models, and the ways to improve the performance of these models via data augmentation. This section also includes the contributions to spoken language understanding, in which the input corresponds to the output of the Automatic Speech Recognition (ASR). I show how the N-Best hypothesis are included as inputs to a Convolutional Neural Network (CNN) that generates the sentence representation and how the context is handle by a Long-Short Term Memory (LSTM). In addition, I introduce risk minimisation for zero-shot learning on this task. Finally, I present the work of the PhD candidate Sebastien Montella on graph embedding. In order to user richer representations as input to neural models, the hyperbolic space was explored for treating temporal relations.

3.1.1 NLU in a Serious Game

Machine Learning (ML) approaches such as logistic regression classifiers and conditional random fields were the state-of-the art back in 2011. I could collect a corpus through Wizard-of-Oz experiments, define an annotation scheme and train Logistic Regression (LR) and Support Vector Machines (SVM) multi-class classifiers for the task of NLU. These models were integrated within distinct dialogues in a serious game. I will start by describing briefly this work that gave origin to four publications (Rojas-Barahona et al., 2012b,c; Rojas-Barahona and Gardent, 2012; Gardent and Rojas-Barahona, 2013).

The serious game is a multiplayer quest where the players (3 teenagers) seek to build a video game joystick in order to free their uncle trapped in the game. To build this joystick, the players must explore a factory and achieve 17 mandatory goals (find the plans, get the appropriate mould, retrieve some raw material from the storing shed, etc). In addition, they can increase their score by achieving optional goals which, when reached, provide them with extra information about the industry (therefore increasing their knowledge). In total, the players can achieve up to 28 goals by conducting 12 separate subdialogs in various parts of the virtual world. That is, dialogs in the game are long dialogs involving multiple players in various settings.

¹dialogue acts are the performative action underlying a utterance(Austin, 1975b)

3.1 Language Understanding

Id	VC	Player	Mandatory Goals	Location
1	Lucas	Ben	Find the address of the enterprise.	Unlce's place.
2	M.Jasper	Lucas	The manufacturing first step	Enterprise reception
3	Samir	Julie	Find the plans of the joystick	Designing Office
4	Samir	Julie	Find out what to do next	Designing Office
5	Melissa	Lucas	Manufacturing process ...	Plant
6	Melissa	Lucas	Find the right machine	Plant
7	Melissa	Lucas	Find out what to do next	Plant
8	Operator	Julie	Knowing about the material space ...	Material Space
9	Serge	Ben	Perform quality tests	Laboratory Tests
10	Serge	Ben	Find out what to do next	Laboratory Tests
11	Sophia	Julie	Find the electronic components.	Finishing
12	Sophia	Julie	Finishing process	Finishing

Table 3.1: Description of the subdialogs in the MP Game.

Table 3.1 summarises the characteristics of the subdialogs conducted within the game highlighting three distinguishing features of game dialogs. First, the dialog participants vary whereby both the game agent and the player can change. Thus in the game, the player alternatively plays any of the three children involved in the quest while the game agent is successively, Lucas, M. Jasper, Samir, Melissa, an operator, Serge and Sophia. Second, game dialogs are task-driven whereby each subdialog is related to a step in the game and each dialog turn aims to achieve a game goal and improve the player score. Third, the context in which each subdialog takes place varies as the player moves around the world.

NLU annotation schema for dialogue is not necessarily dictated by speech act theory alone but might also consider more practical issues namely, how well it will support interpretation and/or dialogue. To enhance learning, the annotation schema designed for the game combines core communicative acts (Bunt et al., 2010) with domain specific information. The domain specific information specifies the goals being pursued/discussed/achieved etc. while the communicative act can be viewed as specifying how the current information state is updated by the speaker's utterance.

dialog: 01_dialogDirecteur-Tue Jun 14 11:04:23 2011

M.Jasper: Bonjour, je suis M.Jasper le directeur.
Hello, I am the director, Mr. Jasper.

→ *greet*

M.Jasper: Qu'est-ce que je peux faire pour vous?
What can I do for you?

→ *ask(task(X))*

Lucas: je dois sauver mon oncle
I must rescue my uncle

→ *first_step*

M.Jasper: Pour faire votre manette, il vous faut des plans. Allez voir dans le bureau d'études, ils devraient y être.
To build the joystick you will need the plans. You will find them in the Designing Office.

→ *inform(do(first_step))*

M.Jasper: J'aurais aimé continuer à répondre à vos questions mais je dois reprendre mon travail! Bonne Chance!
I have to go back to work! Good Luck!

→ *quit*

Figure 3.1: Excerpt from a dialogue in the EmoSpeech corpus. The corresponding user semantics is shown highlighted on the right.

Table 3.2 gives the full list of Dialogue Acts used for annotation together with the correspond-

Dialogue Act	Label	Gloss	Speaker
Welcome greeting	greet	Welcome greeting	P,S
Farewell greeting	quit	Farewell geeting	P, S
Adress Request	ask(Goal)	Request to pursue Goal	S
Adress Request	help	Request for help	P
Confirm	yes	Confirms previous query	P
Disconfirm	no	Disconfirms previous query	P
Provide Information	inform(do(Goal))	Provides information about how to achieve Goal	S
Provide Information	Goal	Provides information about the goal Goal	P
Positive Feedback	ack	Acknowledges understanding of preceding turn	S
Propositional Question	ask(do(more(X)))	Asks whether other topics should be discussed	S
Set Question	ask(topic(X))	Asks which other topics should be discussed	S
Out of Context	other	Out of context turn	P,S
Misunderstanding	reqRep	Request for rephrasing	S

Table 3.2: The Annotation Scheme. P and S stands for Player and System respectively.

ing dialog acts and a gloss of their meaning. As can be seen the labels used are very specific to the game to facilitate the integration within the game (same goals as defined in the serious game) and to bypass much of the pragmatic reasoning necessary to associate a dialog turn with a communicative function. For instance, in the dialog above, the turn *je dois sauver mon oncle* (*I must rescue my uncle*) does not explicitly state that the player (i) is seeking to achieve the game goal “rescueing one’s uncle” and (ii) is asking the game agent for the first step towards achieving that goal.

Experimental setup

We experimented with both an SVM and an LR² classifier using different sets of features on different data sets with and without TF*IDF (term frequency*Inverse Document Frequency) filtering.

	Whole Dialog		Subdialogs	
	w/o Tf*Idf	w/ Tf*Idf	w/o Tf*Idf	w/ Tf*Idf
LR	79.74	90.26	86.41	88.22
SVM	78.79	88.55	76.45	83.99
SVM (P)			78	83.55

Table 3.3: Global Results for the Logistic Regression (LR), the SVM (SVM) and the SVM Classifier with Penalisation (SVM(P))

We compared a single classifier on the whole dataset (the whole game) against 12 distinct classifiers, one for each subdialog. In both cases the categories to be learned are restricted to the speaker’s intent (greet,quit,inform(Goal), ack, ask(do(more(X))), ask(topic(X)), other in Table

²We used MALLET (McCallum, 2002) for the LR classifier with L1 Regularisation.

3.1 Language Understanding

3.2). Taking into account the game goals, the total number of categories to be learned is 27. When learning on subdialogs, the number of categories to be learned is smaller but so is the size of the training set. The features for the machine learning models were bag of words, in which stop words were filtered out, utterances were deaccented and converted to lower-case. In addition, we experimented with various context length using as features the 0 to 4 previous dialogue acts. Subdialog identifiers were also used when training the classifier on the whole dialogue. More details are given in (Rojas-Barahona et al., 2012a).

We also experimented using tf*idf filtering to limit the impact of frequent uninformative words. Moreover, we experimented penalising those categories with more training instances, since the data was highly skewed. Dialogue acts that relate to optional goals were often not followed up by the players resulting in data sparseness.

Results

Table 3.3 shows the results for the 6 main configurations: training on the whole dialog or on subdialogs, with and without tf*idf filtering and using LR, SVM or SVM with penalisation. The best results are obtained using the LR classifier on the whole dataset with tf*idf filtering. Penalising improved slightly the accuracy of the SVM when classifying without tf*idf filtering or when having a reduced context (0 or 2 previous acts in Table 3.4).

*Impact of the tf*idf filtering.* Globally, the tf*idf filtering has a positive impact leading to an increase in accuracy ranging from 2.81 to 11.52 points. For the SVM classifier, the tf*idf filtering consistently lead to better results. However, for the LR classifier the filtering adversely impacts performance on short subdialogs (6 and 7), where one unique goal is being discussed. We conjecture that for these cases, the tf*idf filtering removed words which helped the classifier distinguish between turns about the unique goal from other turns. SVM with penalisation yields worse results with the tf*idf filtering than without, thus suggesting overfitting. In the next section we present how can we exploiting synonyms to improve generalisation.

Impact of contextual features. Having a notion of context is crucial for correctly interpreting dialog acts. As mentioned above, we use the dialog acts of the previous turns to model context. However the further back we look into the previous turns, the more features there will be to train on. In other words, depending on the number of previous turns considered, the data to learn from will be more or less sparse. We experimented with 3 setups: a null context, the dialog acts of the two previous turns and the dialog acts of the four previous acts. Table 3.4 shows the results.

	Whole Dialog			Subdialogs		
	0	2	4	0	2	4
LR	88.43	90.26	90.26	84.43	87.59	88.22
SVM	84.36	86.76	88.55	78.04	82.06	83.99
SVM(P)				79.32	83.12	83.55

Table 3.4: The impact of context on accuracy. 0,2 and 4 indicates that the context is captured by having as features the dialog acts of 0, 2 and 4 previous turns respectively

Impact of dialog acts. The accuracy varies per dialog acts from 48% to 99%. with most of the acts having an accuracy above 80%. Unsurprisingly, the acts with lowest accuracy are also the

acts with fewest training data. The data is split randomly for the 30-fold evaluation with the risk of having insufficient data for optional goals.

Estimating the User Satisfaction : We applied the evaluation framework Paradise (Walker et al., 1997a) to assess these dialogues in (Rojas-Barahona and Gardent, 2012). We compared a rule-based dialogue manager and a dialogue manager that picks the answer randomly from a set of candidate responses. We found out that users prefer to talk to the second system because conversations are more fluid and there are less repetitions and misunderstandings. However, with the second dialogue manager some times the player got stuck in the game and needed to repeat the dialogue with the virtual agent, impacting negatively the user satisfaction.

Data Augmentation

This work on data augmentation was published in (Gardent and Rojas-Barahona, 2013). We explored four ways of modifying the content features used for classification: lemmatising the training and the test data; augmenting the training data with automatically acquired paraphrases; and substituting unknown words with synonyms or its distributional neighbours at run-time.

For Lemmatisation, we used the French version of TreeTagger³ to lemmatise both the training and the test data. Lemmas without any filtering were used to train classifiers. We then compare performance with and without lemmatisation. As we shall see, the lemma and the POS tag provided by TreeTagger are also used to lookup synonym dictionaries and EuroWordNet when using synonym handling at run-time.

We were among the first to exploit automatically acquired paraphrases and to use these not only to increase the size of the training corpus but also to better balance it⁴. We proceed as follows.

First, we generated paraphrases using a pivot machine translation approach where each user utterance in the training corpus (around 3610 utterances) was translated into some target language and back into French. Using six different languages (English, Spanish, Italian, German, Chinese and Arabian), we generated around 38000 paraphrases. We used Google Translate API for translating.

Second, we eliminate from these paraphrases, words that are likely to be incorrect lexical translations by removing words with low normalised term frequency (< 0.001) across translations i.e., lexical translations given by few translations and/or translation systems. We then preprocessed the paraphrases in the same way the utterances of the initial training corpus were preprocessed i.e., utterances were unaccented, converted to lower-case and stop words were removed, the remaining words were filtered with TF*IDF. After preprocessing, duplicates were removed.

Third, we added the paraphrases to the training data seeking to improve the balance between dialog acts per dialog. The process to balance data was guided by the deviation of the category with lowest examples compared to the standard deviation. If the deviation is lower than the standard deviation then we add paraphrases by keeping as much as possible the data balanced after replacement. We invite the interested reader to find more details about the algorithm proposed for balancing data in the paper (Gardent and Rojas-Barahona, 2013).

³<http://www.ims.uni-stuttgart.de/projekte/corplex/TreeTagger/>

⁴The Emospeech data is highly skewed with some classes being populated with many utterances and others with few.

3.1 Language Understanding

Substituting Synonyms for Unknown Words A word is unknown, if it is a well-formed French word⁵ and if it does not appear in the training corpus. When an unknown word w is detected in a player utterance at run-time, we search for a word w' which occurs in the training data and is either a synonym of w or a distributional neighbour. After disambiguation, we substitute the unknown word for the synonym.

H		Lemmatisation			
H-H	Orig.	Lemmas	+EWN	+DIC	+RI
Orig.	65.70% ± 5.62	66.04% ± 6.49	68.17% ± 6.98	67.92% ± 4.51	66.83% ± 5.92
Parap.	70.89% ± 6.45	74.31% ± 4.78*	74.60% ± 5.99*	73.07% ± 7.71*	72.63% ± 5.82*
H-C	Orig.	Lemmas	+EWN	+DIC	+RI
Orig.	59.71% ± 16.42	59.88% ± 7.19	61.14% ± 16.65	61.41% ± 16.59	60.75% ± 17.39
Parap.	59.82% ± 15.53	59.48% ± 14.02	61.70% ± 14.09*	62.01% ± 14.37*	61.16% ± 14.41*

Table 3.5: Accuracy on the H-H and on the H-C corpus. The star denotes statistical significance with the Wilcoxon test ($p < 0.005$) used for the HH corpus and the McNemar test ($p < 0.005$) for the HC corpus.

To identify synonyms, we make use of two lexical resources namely, the French version of EuroWordNet (EWN) (Vossen, 1998), which includes 92833 synonyms, hyperonyms and hyponyms pairs, and a synonym lexicon for French (DIC)⁶ which contains 38505 lemmas and 254149 synonym pairs. While words are categorised into Noun, Verbs and Adjectives in EWN, DIC contains no POS tag information.

To identify distributional neighbours, we constructed semantic word spaces for each subdialog in the EmoSpeech corpus⁷ using random indexing (RI) on the training corpus expanded with paraphrases. Using the cosine measure as similarity metrics, we then retrieve for any unknown word w , the word w' which is most similar to w and which appear in the training corpus.

For lexical disambiguation, two methods are compared. We use the POS tag provided by TreeTagger. In this case, disambiguation is syntactic only. Or we pick the synonym with highest probability based on a trigram language model trained on the H-H corpus.

Results and Discussion

Table 3.5 summarises the results obtained in four main configurations: (i) with and without paraphrases; (ii) with and without synonym handling; (iii) with and without lemmatisation; and (iv) when combining lemmatisation with synonym handling. We also compare the results obtained when evaluating using 10-fold cross validation on the training data (H-H dialogs) vs. evaluating the performance of the system on H-C interactions.

Overall Impact The largest performance gain is obtained by a combination of the three techniques namely, data expansion, synonym handling and lemmatisation (+8.9 points for the cross-

⁵A word is determined to be a well-formed French word if it occurs in the LEFFF dictionary, a large-scale morphological and syntactic lexicon for French (Sagot, 2010)

⁶DICOSYN (<http://elsap1.unicaen.fr/dicosyn.html>).

⁷We also used *distributional semantics* from the Gigaword corpus but the results were poor probably because of the very different text genre and domains between the the Gigaword and the game.

validation experiment and +2.3 for the H-C evaluation).

Impact of Lexical Substitution at Run Time We found that lexical resources are only useful when combined with lemmatisation. This is unsurprising since synonym dictionaries and EuroWordNet only contain lemmas. Indeed when distributional neighbours are used, lemmatisation has little impact (e.g., 65.11% using distributional neighbours without lemmatisation on the H-H corpus without paraphrases vs. 66.41% when using lemmatisation).

Another important issue when searching for a word synonym concerns lexical disambiguation: the synonym used to replace an unknown word should capture the meaning of that word in its given context. We tried using a language model trained on the training corpus to choose between synonym candidates (i.e., selecting the synonym yielding the highest sentence probability when substituting that synonym for the unknown word) but did not obtain a significant improvement. In contrast, it is noticeable that synonym handling has a higher impact when using EuroWordNet as a lexical resource. Since EuroWordNet contain categorial information while the synonym dictionaries we used do not, this suggests that the categorial disambiguation provided by TreeTagger helps identifying an appropriate synonym in EuroWordNet.

Finally, it is clear that the lexical resources used for this experiment are limited in coverage and quality. We observed in particular that some words which are very frequent in the training data (and thus which could be used to replace unknown words) do not occur in the synonym dictionaries. For instance when using paraphrases and dictionaries (fourth row and fourth column in Table 3.5) 50% of the unknown words were solved, 17% were illformed and 33% remained unsolved. To compensate this deficiency, we tried combining the three lexical resources in various ways (taking the union or combining them in a pipeline using the first resource that would yield a synonym). However the results did not improve and even in some cases worsened due probably to the insufficient lexical disambiguation. Interestingly, the results show that paraphrases always improves synonym handling presumably because it increases the size of the known vocabulary thereby increasing the possibility of finding a known synonym.

In sum, synonym handling helps most when (i) words are lemmatised and (ii) unknown words can be at least partially (i.e., using POS tag information) disambiguated. Moreover since data expansion increases the set of known words available as potential synonyms for unknown words, combining synonym handling with data expansion further improves accuracy.

Impact of Lemmatisation When evaluating using cross validation on the training corpus, lemmatisation increases accuracy by up to 3.42 points indicating that unseen word forms negatively impact accuracy. Noticeably however, lemmatisation has no significant impact when evaluating on the H-C corpus. This in turn suggests that the lower accuracy obtained on the H-C corpus results not from unseen word forms but from unseen lemmas.

Impact of Paraphrases On the H-H corpus, data expansion has no significant impact when used alone. However it yields an increase of up to 8.27 points and in fact, has a statistically significant impact, for all configurations involving lemmatisation. Thus, data expansion is best used in combination with lemmatisation and their combination permits creating better, more balanced and more general training data. On the H-C corpus however, the impact is negative or insignificant suggesting that the decrease in performance on the H-C corpus is due to content words that are

3.1 Language Understanding

new with respect to the training data i.e., content words for which neither a synonym nor a lemma can be found in the expanded training data.

While classifiers are routinely trained on dialog data to model the dialog management process, the impact of such basic factors as lemmatisation, automatic data expansion and synonym handling has remained largely unexplored. The empirical evaluation described here suggests that each of these factors can help improve performance but that the impact will vary depending on their combination and on the evaluation mode. Combining all three techniques yields the best results. We conjecture that there are two main reasons for this. First, synonym handling is best used in combination with POS tagging and lemmatisation because these supports partial lexical semantic disambiguation. Second, data expansion permits expanding the set of known words thereby increasing the possibility of finding a known synonym to replace an unknown word with.

3.1.2 Spoken Language Understanding and Few-Shot Learning

The following work was published in Coling 2016 (Rojas Barahona et al., 2016). At that time the task of Spoken Language Understanding (SLU), namely semantic decoding, was seen as a sequence tagging problem with models trained and tested on datasets with word-level annotations (Tür et al., 2013; Mesnil et al., 2015; Yao et al., 2013; Sarikaya et al., 2011; Deoras and Sarikaya, 2013; Sarikaya et al., 2014). Nevertheless, spoken language understanding from *unaligned data*, in which utterances are annotated with an abstract semantics, faces the additional challenge of not knowing which specific words are relevant for extracting the semantics. This problem was tackled in (Zhou and He, 2011), by using conditional random fields (CRFs) driven by finely-tuned hand-crafted features. Other discriminative approaches that deal with unaligned data use some form of *delexicalisation* or mapping of the input to known ontological concepts (Henderson et al., 2012, 2014a). The main disadvantage of delexicalisation is the difficulty in scaling it, not only to larger and more complex dialogue domains but also to handle the many forms of language variation.

We proposed a semantic decoder that learns from unaligned data (Figure 3.2) and that exploits rich semantic distributed word representations instead of delexicalisation. The semantic decoder predicts the dialogue act and the set of slot-value pairs from a set of n-best hypotheses returned by an ASR. The prediction is made in two steps. First, a deep learning architecture is used for the joint prediction of dialogue acts and the presence or absence of slots. Second, the same architecture is reused for predicting the values of the slots that were detected by the first joint-classifier. The deep architecture combines sentence and context representations. A CNN (Collobert et al., 2011) is used to generate the sentence representation, while a LSTM network (Hochreiter and Schmidhuber, 1997) is used to generate the context representation. A non-linear function then combines the top layers of these neural networks and distinct Softmax layers are used to predict the dialogue act and slots in the first joint model. In the second model, a single Softmax predicts the possible values for each slot.

We evaluated our models on two datasets DSTC2 (Henderson et al., 2014b) and In-car (Tsiakoulis et al., 2012) using accuracy, f-measure and the Item Cross Entropy (ICE) score (Thomson et al., 2008). Our models outperform previous proposed models, without using manually designed features and without any pre-processing of the input (e.g., stop words filtering, delexicalisation). They do this by exploiting distributed word representations and we claim that this allows semantic decoders to be built that can easily scale to larger and more complex dialogue domains.

voip-922209b777-20130325_155209

System: Hello , welcome to the Cambridge restaurant system?
You can ask for restaurants by area, price range or food type.
How may I help you?

→ *welcomemsg*

User: i am looking for a moderately priced restaurant
in the north part

→ *inform(area=north,pricerange=moderate)*

System: meghna is a nice restaurant
in the north of town in the moderate price range

→ *offer(name=meghna),
inform(pricerange=moderate,area=north)*

User: is everything else

→ *reqalts()*

System: golden wok is a nice restaurant
in the north of town in the moderate price range

→ *offer(name=golden...),
inform(pricerange=moderate, area=north)*

User: can i get the phone number

→ *request(slot=phone)*

System: The phone number of golden wok is ...

→ *offer(name=golden wok),inform(phone=01..)*

User: type of food do they serve

→ *request(slot=food)*

Figure 3.2: Excerpt from a dialogue in the DSTC2 corpus. The top-best ASR hypothesis is shown highlighted on the left, and the corresponding user semantics is shown highlighted on the right.

Deep Learning Semantic Decoder

We split the task of semantic decoding into two steps: (i) training a joint model for predicting the dialogue act and presence or absence of slots and (ii) predicting the values for the most probable slots detected in (i). As shown in Figure 3.3, we use the same deep learning architecture in both steps for combining sentence and context representations to generate the final hidden unit that feeds one or many softmax layers. In the first step, as shown in the Figure, there are distinct softmax layers for the joint optimisation of the dialogue act and each possible slot. In the second step there is a single softmax layer that predicts the value of each specific slot. In the following we explain this architecture in more detail.

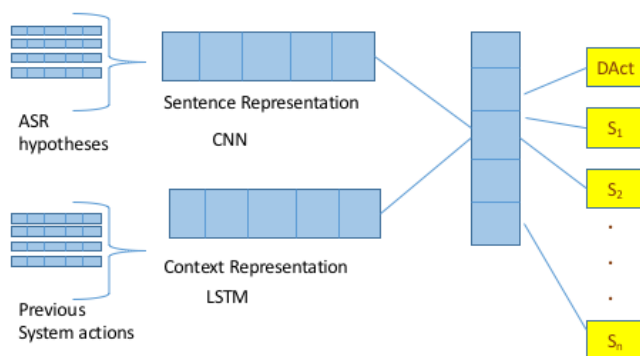


Figure 3.3: Combination of sentence and context representations for the joint prediction of dialogue acts and slots.

3.1 Language Understanding

Sentence Representation

A CNN is used for generating the hypothesis representation, then these representations are weighted by their confidence scores and then summed up to obtain the sentence representation (Figure 3.4).

The CNN is a variant of (Kim, 2014), in which the inputs are the word vectors in each ASR hypothesis. Let x_i be a k -dimensional word embedding for the i -th word in a hypothesis. A hypothesis of length m is represented as: $x_{1:m} = \mathbf{x}_1 \oplus \mathbf{x}_2 \oplus \dots \oplus \mathbf{x}_m$ where \oplus is the concatenation operator. A convolutional operation is applied to a window of l words to produce a new feature.

$$c_i = f(\mathbf{w} \cdot \mathbf{x}_{i:i+l-1} + b) \quad (3.1)$$

where f is the hyperbolic tangent function; $w \in \mathbb{R}^k$ is a filter applied to a window of l words and $b \in \mathbb{R}$ is a bias term. The filter is applied to every window of words in the sentence to produce a feature map.

$$\mathbf{c} = [c_1, c_2, \dots, c_{n-l+1}] \quad (3.2)$$

with $\mathbf{c} \in \mathbb{R}^{n-l+1}$. A max pooling operation is then applied to give the maximum value $c = \max\{\mathbf{c}\}$ as the representative feature for that filter. Multiple filters can be applied by varying the window size to obtain several adjacent features for a given hypothesis. These features \hat{f}_j for the hypothesis $j \in H$ are then multiplied by the ASR confidence score p_j ⁸ and summed over all ASR hypotheses to generate a representation for the sentence s_t (Equation 3.3), as shown in Figure 3.4.

$$s_t = \sum_{j \in H} \hat{f}_j * p_j \quad (3.3)$$

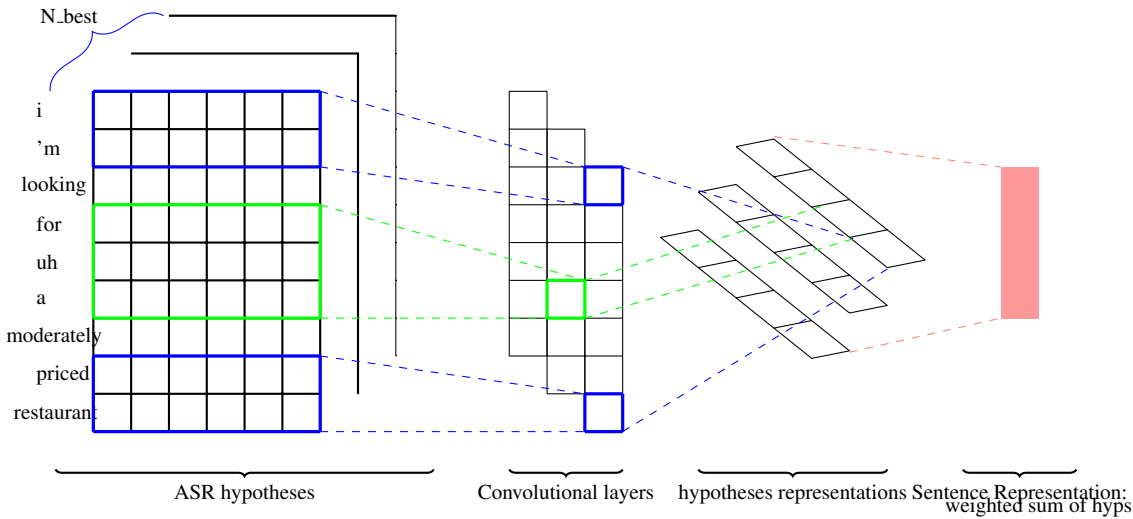


Figure 3.4: Sentence Representation: after applying convolution operations on the N-best list of ASR hypotheses, the resulting hidden layers are weighted by the ASR confidence scores and summed.

⁸The posterior probability of hypothesis j in the N-best list.

Context Representation

An LSTM (Hochreiter and Schmidhuber, 1997) is used for tracking the context implied by previous dialogue system actions. The top layer of this LSTM network then provides the context representation for decoding the current input utterance.

An LSTM is a sequence model that utilises a memory cell capable of preserving states over long periods of time. This cell is recurrently connected to itself and it has three multiplication units, an input gate, a forget gate and an output gate. These gating vectors are in $[0,1]$. The cell makes selective decisions about what information is preserved, and when to allow access to units, via gates that open and close.

As shown in Figure 3.2, system actions are encoded in the form of a system dialogue act plus one or more slot-value pairs. To track the history of system actions, slots and values are treated as words and the input x_t is formed from its corresponding word vectors. The length of the context can vary. We consider all the system actions previous to the current user utterance, or a window l of the previous system actions. For instance, if we are currently processing the last user input in Figure 3.2, in which L is the total number of system actions, we can consider all previous system actions ($L=4$), or the last l system actions, where $l < L$.

Combining Sentence and Context

We study in this paper two ways of combining the sentence s_t and the context h_t representations. The first straightforward way is to apply a non linear function to their weighted sum:

$$\hat{h}_t = \tanh(\mathbf{W}_s \cdot s_t + \mathbf{W}_c \cdot h_t) \quad (3.4)$$

The second way is to let the sentence representation be the last input to the LSTM network, then $\hat{h}_t = h_t$. For classification a softmax layer is used for each prediction. The result of the prediction is the most probable class. The back-propagation optimisation is done by minimising the negative log-likelihood loss function through stochastic gradient descent.

Experimental Evaluation

In this section we introduce the corpora, and describe the experiments performed and the evaluation metrics used.

Corpora

Experimental evaluation used two similar datasets: DSTC2 (Henderson et al., 2014b) and In-car (Tsiakoulis et al., 2012). Both corpora were collected using a spoken dialogue system which provides restaurant information system for the city of Cambridge. Users can specify restaurant suggestions by area, price-range and food type and can then query the system for additional restaurant specific information such as phone number, post code and address. The first dialogue corpus was released for the dialogue state tracking challenge and we use here the semantic annotations that were also provided⁹. The trainset has 2118 dialogues and 15611 turns in total while the testset has 1117 dialogues and 9890 turns in total.

⁹The DSTC2 corpus is publicly available in: <http://camdial.org/~mh521/dstc/>

3.1 Language Understanding

The second corpus contains dialogues collected under various noisy in-car conditions. In a stationary car with the air conditioning fan on and off, in a moving car and in a car simulator (Tsiakoulis et al., 2012)¹⁰. The trainset has 1508 dialogues and 10532 turns in total and the testset has 641 dialogues and 4861 turns in total. Because of the noise, the average word error rate (WER = 37%) is significantly higher than for DSTC2 (around 29%).

Experiments

Step I: Joint classification of dialogue-acts and slots: We evaluated five different model configurations for the joint classification of dialogue-acts and presence or absence of slots.

- **CNN:** the softmax layers for the joint classification of dialogue acts and slots are connected directly to the CNN sentence representation with no context.
- **CNN+LSTM:** we study the influence of context by considering the previous system actions (Section 3.1.2, Eq. 3.4), here we study the different context length, by using a context window of 1, 4, and all the previous system actions, namely **CNN+LSTM_w1**, **CNN+LSTM_w4** and **CNN+LSTM_w** respectively.
- **LSTM_all:** Finally, we study the impact of long distance dependencies, by using mainly the LSTM model, with the previous system actions as input, but we inject the sentence representation as the last LSTM input.

Step II: Classification of slot value pairs: We select the best model in step I for predicting the presence of slots, then for each slot present we predict the value, by using again the best architecture from the previous step.

Evaluation Metrics

We evaluate the performance of our models by using the conventional metrics for classification, namely accuracy, precision, recall and F-measure (F1-score). In addition, we used the ICE score to measure the overall quality of the distribution returned by the models taken into account the hypotheses and the reference semantics (ie. ground-truth)(Thomson et al., 2008).

Results and Discussion

In this section we report the results on DSTC2 and In-car dialogue corpora.

Step I: Joint classification of dialogue-acts and slots: For this step, the classifiers must predict jointly 14 dialogue acts and 5 slots for the DSTC2 dataset as well as 14 dialogue acts and 7 slots for the In-car dataset. We evaluate both (i) using 10 fold cross-validation on the trainsets and (ii) on the corpora' testsets.

Our results on 10 fold cross-validation results on both corpora suggest that for DTSC2, the context representation is not significantly impacting the prediction. Although, the model with a window of 4 ,**CNN+LSTM_w4**, improves slightly the accuracy and f1-score. On the In-car

¹⁰This corpus has been obtained in an industry funded project and therefore it is not available for public use.

dataset, however, including the context does help to disambiguate the semantic predictions from ill-formed hypotheses. This is expected, since this data set has a much higher error rate and hence higher levels of confusion in the ASR output. Although there is no significant difference on the f1-score when using the immediate previous system act ($w1$) or a longer context, **CNN+LSTM_w** gives a better accuracy and a lower ICE score on this dataset.

Table 3.6 shows the results on the test sets. Consequently, when evaluating on the DSTC2 test set, a window of 4 ($w4$), performs slightly better than other window sizes and better than the simple CNN model. On the In-car testset, a context window of 4 outperforms all the other settings: **CNN+LSTM**. However, on this test set using the sentence representation as the last input to the LSTM context neural network (section 3.1.2) improves the f1-score and reduces the ICE error.

Corpus	Metric	CNN	CNN+LSTM			LSTM_all
			w1.	w4	w	
-	-	-	w1.	w4	w	-
DSTC2	acc.	96.03%	95.79%	95.79%	95.69%	95.59%
	P.	89.73%	88.69%	88.95%	88.38%	88.15%
	R.	84.74%	85.09%	86.02%	85.96%	84.76%
	F1	87.14%	86.83%	87.43%	87.12%	86.42%
	ICE	0.268	0.278	0.292	0.297	0.308
In-car	acc.	87.60%	82.19%	82.25%	82.14%	82.3%
	P.	69.96%	79.52%	79.29%	80.25%	78.12%
	R.	62.14%	71.09%	71.59%	70.9%	74.04%
	F1	65.53%	74.89%	75.15%	75.02%	75.9%
	ICE	1.332	1.344	1.333	1.421	1.106

Table 3.6: Evaluation of the Step I on DSTC2 and In-car testsets. We also compare two ways of combining sentence and context representation: **CNN+LSTM** models (combining sentence and context representation through a non linear function) and **LSTM_all** model (embedding the sentence representation into the context model).

Step II: Prediction of slot value pairs For evaluating Step II, we selected the best model obtained during the 10-fold cross-validation experiments in terms of F1 score. For both corpora, this was the **CNN+LSTM_w4** configuration. For DSTC2, it was the 4th-fold crossvalidation with $Acc = 90.42\%$, $F1 = 88.69\%$ and $ICE = 0.251$. For In-car, it was the 5th-fold crossvalidation with $Acc = 93.13\%$, $F1 = 81.49\%$ and $ICE = 0.393$. We used these models to classify whether a given slot appears in a given hypothesis or not. Then for that slot, we train another **CNN+LSTM_w4** classifier for predicting its values. In the In-car corpus the slot "type" has only one possible value "restaurant". Similarly, the slot "task" can only be the value "find". For these slots with only one value, we report values using the model of Step I, since it is enough to detect the slot in the utterance.

Given that there is no domain specific delexicalisation, the models achieve a good level of performance overall (Table 3.7). Note that the slot "food" has 74 possible values in DSTC2 and 25 in In-car. Hence, this slot has much higher cardinality than all the other slots.

Overall performance A baseline for assessing overall performance is provided by the model presented in (Henderson et al., 2012), in which the vector representation is obtained by summing

3.1 Language Understanding

Slot	DSTC2					In-car				
	Acc.	P.	R.	F1	ICE	Acc.	P.	R.	F1	ICE
Slot ¹¹	95.29%	90.89%	95.72%	93.24%	0.478	89.92%	74.73%	61.56%	67.51%	0.743
Area	91.77%	92.66%	92.83%	92.74%	0.563	72.03%	72.56%	74.28%	73.41%	1.676
Food	71.37%	73.19%	76.02%	74.58%	1.989	66.46%	64.27%	68.70%	66.41%	2.309
Price	94.62%	91.33%	94.49%	92.89%	0.729	93.96%	88.77%	92.03%	90.37%	0.632
This ¹²	98.70%	96.79%	93.92%	95.33%	0.113	97.16%	96.14%	84.72%	90.07%	0.214
Type	-	-	-	-	-	95.56%	95.09%	86.69%	90.69%	0.290
Task	-	-	-	-	-	97.12%	83.24%	64.93%	72.95%	0.175
Mean	90.35%	88.97%	90.60%	89.76%	0.774	87.47%	82.11%	76.13%	78.77%	0.863
St.Dev.	0.109	0.091	0.082	0.085	0.715	0.128	0.121	0.118	0.112	0.821

Table 3.7: Evaluation of the step II: the slot-value pairs classification on DSTC2 and In-car.

up the frequency of n-grams extracted from the 10-best hypotheses, weighted by their confidence scores. Here we compare our performance against Henderson’s model with and without context features, namely WNGRAMS+Ctxt and WNGRAMS respectively. Henderson reported his results on the In-car dataset. A similar model, namely SLU1, was evaluated on DSTC2 in (Williams, 2014). Both implementations consist of many binary classifiers for dialogue act and slot-value pairs.

Corpus	Model	F1	ICE
DSTC2	SLU1 (Williams, 2014)	80.2%	1.943
	CNN+LSTM_w4	83.59%	0.758
In-car	WNGRAMS (Henderson et al., 2012)	70.8%	1.76
	WNGRAMS+Ctxt (Henderson et al., 2012)	74.2%	1.497
	CNN+LSTM_w4	73.06%	1.106

Table 3.8: Overall performance of the setting CNN+LST_w4 semantic decoder.

In terms of the ICE score, the model CNN+LSTM W4 outperforms all the baselines (Table 3.8). In terms of the F1 score, the model significantly outperforms the SLU1 and WNGRAMS baselines. However it is slightly worse than WNGRAMS+Ctxt, which has been enhanced with context features on In-car. Remember however, that our model uses only word-embeddings for automatically generating sentence and context representations without having any manually designed features or using explicit application specific semantic dictionaries.

Few-shot Learning through Risk Minimisation (RM)

We treat rarely seen slots by following two steps. (i) We optimise jointly in a deep neural network the weights that feed multiple binary Softmax units. (ii) We further tune the weights learned in the previous step by minimising the theoretical risk of the binary classifiers as proposed in (Balasubramanian et al., 2011). In order to apply the second step, we rely on two assumptions: the rank of the class marginal is assumed to be known and the class-conditional linear scores are assumed to follow a Gaussian distribution. In (Balasubramanian et al., 2011), this approach has been proven to converge towards the true optimal classifier risk. We conducted experiments on the dialogue

corpus released for the third dialogue state tracking challenge, namely DSTC3 (Henderson et al., 2014c) and we show positive results for detecting rare slots as well as zero-shot slot-value pairs.

We use the unsupervised approach proposed in (Balasubramanian et al., 2011) for risk minimisation (RM). We assume a binary classifier that associates a score $f_{\mathbf{W}_0}(\mathbf{h})$ to the first class 0 for the hidden unit $\mathbf{h} = (h_1, \dots, h_n)$ of dimension n :

$$f_{\mathbf{W}_0}(\mathbf{h}) = \sum_i^n w_i h_i$$

where the parameter $w_i \in \mathbb{R}$ represents the weight of the feature indexed by i for class 0.

The objective of training is to minimize the classifier risk:

$$R(\mathbf{W}) = E_{p(\mathbf{h}, Y)}[\mathcal{L}(Y, f_{\mathbf{W}}(\mathbf{h}))] \quad (3.5)$$

where Y is the true label and $\mathcal{L}(Y, f_{\mathbf{W}}(\mathbf{h}))$ is the loss function. The risk is derived as follows:

$$R(\mathbf{W}) = \sum_{y \in \{0,1\}} P(y) \int_{-\infty}^{+\infty} P(f_{\mathbf{W}}(\mathbf{h}) = \alpha|y) \mathcal{L}(y, \alpha) d\alpha \quad (3.6)$$

We use the following hinge loss:

$$\mathcal{L}(y, \alpha) = (1 + \alpha_{1-y} - \alpha_y)_+ \quad (3.7)$$

where $(z)_+ = \max(0, z)$, and $\alpha_y = f_{\mathbf{W}_y}(\mathbf{h})$ is the linear score for the correct class y . Similarly, $\alpha_{1-y} = f_{\mathbf{W}_{1-y}}(\mathbf{h})$ is the linear score for the wrong class.

Given y and α , the loss value in the integral (Equation 3.6) can be computed easily. Two terms remain: $P(y)$ and $P(f_{\mathbf{W}}(\mathbf{h}) = \alpha|y)$. The former is the class marginal and is assumed to be known. The latter is the class-conditional distribution of the linear scores, which is assumed to be normally distributed. This implies that $P(f_{\mathbf{W}}(\mathbf{h}))$ is distributed as a mixture of two Gaussians (GMM):

$$P(f_{\mathbf{W}}(\mathbf{h})) = \sum_{y \in \{0,1\}} P(y) \mathcal{N}(f_{\mathbf{W}}(\mathbf{h}); \mu_y, \sigma_y)$$

where $\mathcal{N}(z; \mu, \sigma)$ is the normal probability density function. The parameters $(\mu_0, \sigma_0, \mu_1, \sigma_1)$ can be estimated from an **unlabeled corpus** \mathcal{U} using a standard Expectation-Maximization (EM) algorithm for GMM training. Once these parameters are known, it is possible to compute the integral in Eq. 3.6 and thus an estimate $\hat{R}(\mathbf{W})$ of the risk without relying on any labeled corpus. In (Balasubramanian et al., 2011), it has been proven that: (i) the Gaussian parameters estimated with EM converge towards their true values, (ii) $\hat{R}(\mathbf{W})$ converges towards the true risk $R(\mathbf{W})$ and (iii) the estimated optimum converges towards the true optimal parameters, when the size of the unlabeled corpus increases infinitely. This is still true even when the class priors $P(y)$ are unknown.

The unsupervised algorithm is as follows:

Unsupervised tuning for the binary classifier c , where $c = 1, \dots, C$

- 1: **input:** \mathbf{h} the top hidden layer and the weights \mathbf{W}^c , as trained by the deep learning decoder (Section 3.1.2).
- 2: **output:** The tuned weights $\hat{\mathbf{W}}^c$
- 3: **repeat**
- 4: **for** every index i in \mathbf{h} **do**,
- 5: Change the weights $\mathbf{W}_i^c = \mathbf{W}_i^c + \delta$,

3.1 Language Understanding

```
6:   Estimate the Gaussian parameters using EM
7:   Compute the risk (Eq. 3.6)13 on the unlabeled corpus  $\mathcal{U}$  (i.e. the evaluation set).
8:   Compute the gradient using finite differences
9:   Update the weights accordingly  $\mathbf{W}_i^c \hat{=} \mathbf{W}_i^c$ 
10:  end for
11:  until convergence
```

Experiments

The supervised and unsupervised models are evaluated on DSTC3 (Henderson et al., 2014c) using the macro F-Measure¹⁴. We compare then three distinct models, (i) independent neural models for every binary classifier; (ii) neural models optimised jointly and (iii) further tuning of the weights through RM.

Dataset As displayed in Table 3.9a in DSTC3 new slots were introduced relative to DSTC2. The training set contains only a few examples of these slots while the test set contains a large number of them. Interestingly, frequent values per slots in the trainset such as *area=north*, are absolutely absent in the testset. In DSTC3 the dialogues are related to restaurants, pubs and coffee shops. The new slots are: *childrenallowed*, *hastv*, *hasinternet* and *near*. Known slots, such as *food*, can have zero-shot values as shown in Table 3.9b. The corpus contains 3246 dialogues, 25610 turns in the trainset and 2264 dialogues, 18715 turns in the testset.

(a) Frequency of slots in DSTC3.

Slot	#Train	#Test
hastv	1	239
childrenallowed	2	119
near	3	74
hasinternet	4	215
area	3149	5384
food	5744	7809

(b) Some zero-shot values per slots in DSTC3.

Slot	Value	#Train	#Test
near	trinity college	0	5
food	american	0	90
food	chinese takeaway	0	87
area	romsey	0	127
area	girton	0	118

The Gaussianity Assumption As explained in Section 3.1.2, the risk minimisation tuning assumes the class-conditional linear scores are distributed normally. We verified this assumption empirically on our unlabeled corpus \mathcal{U} (i.e. DSTC3 testset) and we found that for the slots: *childrenallowed*, *hastv* and *hasinternet* this assumption holds. However, the distribution for *near* has a negative skew. When verifying the values per slot, this assumption does not hold for *area*. Therefore, we can not guarantee this method will work correctly for *area* values on this evaluation set.

¹³A closed-form is used to compute the risk for binary classifiers. (Rojas Barahona and Cerisara, 2015)

¹⁴The macro F-score was chosen because we are evaluating the capacity of the classifiers to predict the correct class and both classes positive and negative are equally important for our task. Moreover, being nearly zero-shot classifiers, it would be unfair to evaluate only the capacity of predicting the positive category.

(a) Results for learning rare slots on DSTC3 evaluation set. (b) Results for learning zero shot slot-value pairs on DSTC3 evaluation set.

Deep Learning Independent Models	
Slot	F-Measure
childrenallowed	49.84%
hastv	49.68%
hasinternet	49.72%
near	49.90%
Deep Learning Joint Optimisation	
childrenallowed	58.76%
hastv	59.16%
hasinternet	58.77%
near	56.65%
Risk Minimisation Tuning	
childrenallowed	61.64%
hastv	61.35%
hasinternet	60.87%
near	58.60%

Deep Learning Independent Models		
Slot	Value	F-Measure
near	trinity college	49.99%
food	american	49.88%
	chinese take away	49.88%
area	romsey	49.83%
	girton	49.84%
Deep Learning Joint Optimisation		
near	trinity college	61.25%
food	american	59.93%
	chinese take away	61.02%
area	romsey	51.30%
	girton	55.19%
Risk Minimisation Tuning		
near	trinity college	62.08%
food	american	62.52%
	chinese take away	63.79%
area	romsey	48.76%
	girton	51.45%

Results

Tables 3.10a and 3.10b display the performance of the models that predict slots and values respectively. The low F-Measure in the independent models shown their inability to predict positive examples. The models improve significantly the precision and F-Measure after the joint-optimisation. Applying RM tuning results in the best F-Measure for all the rare slots (Table 3.10a) and for the values of the slots *food* and *near* (Table 3.10b). For *area*, the joint optimisation improves the F-Measure but the improvement is lower than for other slots. The performance is being affected by its low cardinality (i.e. 20), the high variability of new places and the fact that frequent values such as *north* and *east*, are completely absent in the test set. As suspected, the RM tuning degraded the precision and F-Measure because the Gaussianity assumption does not hold for *area*. However, RM will work well in larger evaluation sets because the Gaussian assumption will hold when the unlabelled corpus tends to infinite (please refer to (Balasubramanian et al., 2011) for the theoretical proofs).

3.2 Dialogue Manager

My work on dialogue management regards learning the reward function. First, I explored inverse reinforcement learning to infer the reward function from human conversations on the EmoSpeech dataset (Rojas-Barahona and Cerisara, 2014). Second, I trained a predictor of the interaction quality to infer the reward function in the PyDial dialogue framework (Rojas-Barahona, 2020). Last but not least, I co-supervised a PhD thesis on imitation learning to solve the problem of

3.2 Dialogue Manager

the scarce reward signal in dialogue systems. Besides imitation learning (Cordier et al., 2020), we also explored graph neural networks for handling policies in multi-domain and multi-task environments (Cordier et al., 2022). Furthermore, we use both imitation and graph neural policies for few-shot learning (Cordier et al., 2023).

3.2.1 Bayesian Inverse Reinforcement Learning

Inverse reinforcement learning (IRL) was defined in (Ng et al., 2000) as the problem of recovering the reward function from experts’ demonstrations. It learns an optimal reward, which leads to a decision policy that follows as closely as possible the examples provided by experts, while maximising the expected accumulated reward in the long run.

In (Ramachandran and Amir, 2007) we used Bayesian Inverse Reinforcement Learning (BIRL) to infer human behaviour in the context of the Emospeech serious game (Section 3.1.1), given evidence in the form of stored dialogues provided by experts, who played the role of several conversational agents in the game. We also reduce the computational complexity in large state spaces by using the approach proposed by (Michini and How, 2012). Instead of designing in advance the reward function to “properly instruct players”, which is a difficult and subjective task, we rather propose to learn it from humans.

We evaluated BIRL in terms of policy loss (Michini and How, 2012) and is compared against two baselines. The first one uses random rewards, while the second one exploits corpus-estimated locally optimal rewards (i.e., supervised learning). The results show that the proposed approach converges relatively quickly and consistently outperforms both baselines. This suggests that taking into account the dynamic properties of the environment leads to virtual characters that better reproduce the behaviour of experts. Qualitatively, our models have thus learned to adequately inform users and provide help when needed.

States, Actions and Transitions

As shown in Table 3.1, there are 12 distinct conversations in the game between 7 virtual characters (VC) and 3 players. Each of these dialogues talks about mandatory and optional goals. The player either asks for information about these goals or asks for help. Accordingly, the virtual character either informs about the goals or provides help. It can also handle out of domain topics, misunderstandings or request information (see example of dialogue in Figure 3.1).

We designed coarse-grained states containing user and system contributions to the dialogue; either by explicitly asking about the domain specific tasks (i.e. the dialogue goals) or by producing general dialogue acts (e.g., greeting, asking for help, acknowledgments, etc). A binary variable that indicates whether the dialogue has finished is also included. With this state representation we have 32 states for the shortest dialogue (the first dialogue in Table 3.1), and 432 states for the longest dialogue (i.e., the third dialogue in Table 3.1 with 5 goals).

State variables

1. Has any of the characters ended the dialogue with a farewell action? : 1 for setting a terminal state, 0 otherwise.

2. The last goal either informed or requested by the system: 0 when the system has not informed/requested about any goal, otherwise the goal id (e.g., from 1 to up to 5 for the longest dialogue).
3. The last goal either asked or confirmed by the player: 0 when the user has not yet asked/confirmed about any goal, otherwise the id of the goal (e.g., from 1 to up to 5 for the longest dialogue).
4. The last general dialogue act produced by the system: 0 for absence of general dialog act, 1 when providing help, and 2 when asking the player about the task to be solved (e.g., "How may I help you").
5. The user has asked for help: 0 if the user has not asked for help, 1 otherwise.

Actions We are considering only the following actions in our experiments.

- *quit*: farewell greeting.
- *inform(do(g_i))*: informing about how to achieve goal g_i .
- *inform(help)*: providing help
- *ask(task(X))*: Asking the player about the task, it corresponds to a general welcome sentence (e.g., "How may I help you"). Note that this action neither occurs in dialogue 1 nor in dialogue 7.
- *WAIT*: the system gives the turn back to the user.
- *ack*: the system acknowledges understanding.
- *other*: the system answers to out of context turns.

Virtual characters always greet the player at the beginning; thus we do not need to learn this behaviour.

Transition Function The transition function is not deterministic when the next state reflects an (unpredictable) user action. This is typically the case after the WAIT system action. However, BIRL requires this transition function to be given, and we have thus estimated such non-deterministic transition probabilities using smoothed counts from the observed corpus as follows:

$$P(s'|s, a) = \frac{N(s, a, s') + \alpha}{N(s, a) + N_\chi \alpha}$$

Where $N(s, a, s')$ and $N(s, a)$ are respectively the number of times the transition (s, a, s') and the state-action pair (s, a) have been observed in the corpus, and N_χ is the number of observed state-action pairs. α is a smoothing constant arbitrarily set to 0.1.

The other transitions that reflect a system action are deterministic and have been defined as:

$$P(s'|s, a) = \begin{cases} 1, & \text{if } s' = \text{next}_s(s, a) \\ 0, & \text{otherwise} \end{cases}$$

3.2 Dialogue Manager

Where $next_s(s, a)$ is a function that computes the next state given a system action a . For instance, when the system informs about the first goal, g_1 , the action $a_t = inform(do(g_1))$ yields the next state s' to have the state variable 2 set to 1.

Bayesian Inverse Reinforcement Learning

The IRL problem as defined in (Ng et al., 2000) is described as follows: given a finite state space S , a set of actions $A = \{a_1, a_2, \dots, a_k\}$, a transition probability $P_{ss'}^a$, a discount factor γ , and a policy π , determine a set of possible reward functions R such that π is the optimal policy for the given MDP. The IRL problem is an ill-posed problem (Abbeel and Ng, 2004), because potentially an infinite number of rewards may be optimal. Bayesian IRL approaches model this uncertainty by inferring the posterior distribution of the reward vector \mathbf{R} , treating the demonstration sequences as the evidence and relying on a prior on the reward function (Ramachandran and Amir, 2007).

The IRL agent receives a sequence of observations of the expert's behaviour:

$O_\chi = \{(s_1, a_1), (s_2, a_2), \dots, (s_k, a_k)\}$, which means that at time step i , the virtual character χ that mimics the expert is in state s_i and takes the action a_i . After applying Bayes Theorem, the posterior can be written as:

$$Pr(\mathbf{R}|O_\chi) = \frac{Pr(O_\chi|\mathbf{R})Pr(\mathbf{R})}{Pr(O_\chi)} \quad (3.8)$$

We model next the reward function by a simple n -dimensional real vector, where n is the number of different states. Then, $Pr(\mathbf{R}|O_\chi)$ is the posterior distribution of the reward vector given the observed state-action pairs of the expert. $Pr(O_\chi|\mathbf{R})$ is the likelihood of the observed expert state-action pairs given the reward vector \mathbf{R} . This likelihood is modelled in (Ramachandran and Amir, 2007) with a parameter α representing the degree of confidence we have in the expert's ability to choose a good action as follows:

$$Pr(O_\chi|\mathbf{R}) = \frac{1}{Z} e^{\alpha \sum_i Q^*(s_i, a_i, \mathbf{R})} \quad (3.9)$$

$Pr(\mathbf{R})$ is the prior distribution and $Pr(O_\chi)$ is the probability of the evidence over the entire space of reward vectors \mathbf{R} , which is not needed in the BIRL algorithm. The original BIRL algorithm, namely PolicyWalk, follows a Markov Chain Monte Carlo (MCMC) technique iterating as follows: Given a reward vector \mathbf{R} , it performs random walks over the neighbours of \mathbf{R} on a grid of length δ , finding a new proposal $\bar{\mathbf{R}}$, such that: $\bar{\mathbf{R}}(s) = \mathbf{R}(s) \pm \delta$. The proposal is accepted with probability $\min\{1, \frac{Pr(\bar{\mathbf{R}}|O)}{Pr(\mathbf{R}|O)}\}$, where the posterior is given by Eq (3.8).

The expected value of the reward given this posterior is then computed over all these samples. Note that the normalising constants cancel out in the ratio used to accept the proposed $\bar{\mathbf{R}}(s)$ and that finding Q^* in Eq (3.9) requires to solve the MDP at every MCMC iteration. This can be done for example with the policy iteration (PI) algorithm (Sutton and Barto, 2018).

BIRL converges slowly when applied to large state spaces. One reason for this is that it infers the reward of every state, although many states have little expert evidence. Second, searching over a reward function space easily increases the number of MCMC iterations needed to approximate the mean of the posterior. To solve these limitations, (Michini and How, 2012) proposed a modified BIRL (MBIRL) that:

- infers only those states that are similar to the observed ones according to a *kernel-based relevance function*.

- uses simulated annealing to focus the sampled distribution around its maximum, hence reducing the number of samples needed to converge. Therefore, they use a modified acceptance probability of $\left(\frac{Pr(\bar{\mathbf{R}}|O)}{Pr(\mathbf{R}|O)}\right)^{\frac{1}{T_i}}$ where T_i is a decreasing *cooling schedule*.

Experiments

We introduce the baselines, the evaluation metrics, and the experiment setup for 12 dialogues in the game.

Baselines We evaluate the performances of the proposed system by comparing it with two baselines:

- Using random rewards (RR);
- Exploiting "locally-estimated" rewards (LR), i.e., rewards that are trained on the corpus with the additional assumptions that the reward prior $Pr(\mathbf{R})$ is uniform, that the states are conditionally independent given the reward $P(O_\chi|\mathbf{R}) = \prod_i P(s_i|\mathbf{R})$ and that the state likelihood is multinomial with parameters representing the reward $P(s = k|\mathbf{R}) = R_k$, so that the path that maximizes the cumulated reward also maximizes the likelihood. Then:

$$\begin{aligned} \arg \max_R Pr(\mathbf{R}|O_\chi) &= \arg \max_R Pr(O_\chi|\mathbf{R}) \\ &= \arg \max_R \prod_i P(s_i|\mathbf{R}) \end{aligned} \quad (3.10)$$

Let n_k be the number of times the k^{th} state occurs in the expert observations: $n_k = |\{(s_i = k, a_i)\}_{i \in O_\chi}|$

Then we want to maximize the likelihood $\prod_k P(s = k|\mathbf{R})^{n_k}$ under the constraint $\sum_k R_k = 1$, which gives the locally optimum reward:

$$\hat{R}_k = \frac{n_k}{N_\chi}$$

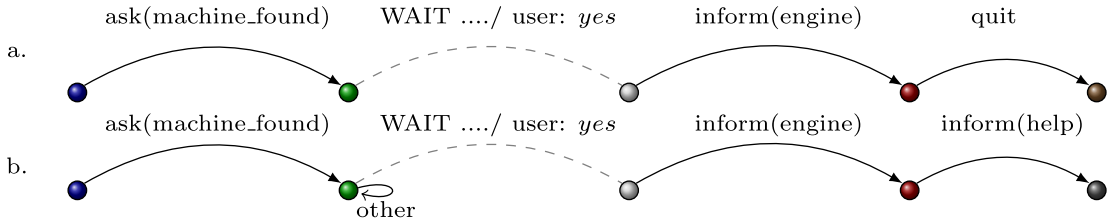
with $N_\chi = |\{(s_i, a_i)\}_{i \in O_\chi}|$ the number of observed state-action pairs.

Evaluation Metrics We consider two evaluation metrics: the policy loss (Michini and How, 2012) and the system training time.

- *Policy loss*: The policy loss is the ratio $\frac{n_\neq}{N_\chi}$, where $n_\neq = |\{(s_i, a_i \neq \pi(s_i))\}_{i \in O_\chi}|$ is the number of expert state-action pairs that disagree with the learned policy π and $N_\chi = |\{(s_i, a_i)\}_{i \in O_\chi}|$ is the number of observed state-action pairs.
- *Elapsed time*: The time in milliseconds it takes to MBIRL and to the policy iteration algorithms to finish.

3.2 Dialogue Manager

D. 7



D. 2

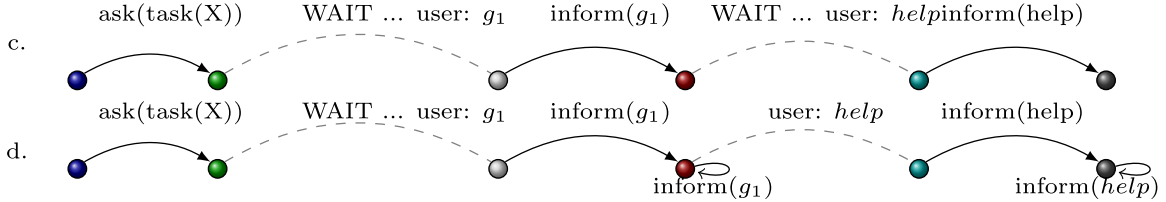


Figure 3.5: Comparison of expert vs. MBIRL trajectories for dialogues 7 (top) and 2 (bottom). (a) and (c) depict expert trajectories, while (b) and (d) show the trajectories of MBIRL optimal policy π .

Performance: Two important issues affect performance: the size of the state-space and the limited number of expert observations. In general MBIRL outperforms both locally-optimal and random rewards. However, with a larger state-action space such as in dialogue 3, 4 and 8, the models do not improve over the locally-optimal reward, which suggests that the number of samples that are generated is not large enough. Moreover, for state spaces greater than 300 states, MBIRL takes a prohibitively long running time to finish. The huge computational expense for large state spaces is an important limitation of MBIRL since it needs to solve one RL problem per iteration. A potential solution to this issue might be to use appropriate function approximation both for the Q function and for modelling the reward function \mathbf{R} , but this is left for future work.

Unsurprisingly, the MBIRL policy loss is higher for optional dialogues (such as dialogue 10 and 12) showing that the scarce number of observations significantly affects performance.

Trajectories: Figure 3.5 shows two dialogue trajectory excerpts with both the gold (or expert) trajectory and the trajectory inferred by MBIRL. Interestingly, in most of the dialogues, both trajectories coincide in the first state and in those states where the system has to inform about mandatory goals just after explicitly requested by the user. This is also the case of the states where the system properly provides help as requested by the user. On the other hand, the learned policy usually fails to close the dialogue and it sometimes contains repetitions e.g., once it has informed about a goal, it may inform again later on.

For more details please refer to (Rojas-Barahona and Cerisara, 2014).

3.2.2 Is the User Enjoying the Conversation?

The impact of user satisfaction in policy learning for task-oriented dialogue systems has long been a subject of research interest (Walker et al., 1997b; Schmitt et al., 2011; Ultes et al., 2015; Ultes, 2019). Similarly, sentiment analysis has been widely adopted to analyse massive blogs, recommendations tweets and reviews (Rojas-Barahona, 2016; Do et al., 2019). Although senti-

ment analysis can be used to infer user satisfaction, most of the work that incorporates sentiment analysis in dialogue focused on the generation of empathetic responses in end-to-end chitchat dialogues (Lee et al., 2018; Ma et al., 2020). Moreover, most sentiment analysis solutions focus on the analysis of out-of-context short texts (Rojas-Barahona, 2016). In this work we are interested in the study of user satisfaction for measuring the quality of the interaction in task-oriented dialogues, in which dialogue is modelled as a POMDP (Young et al., 2013a).

First we study distinct neural networks that use distributed representations for predicting satisfaction scores. At this stage we would like to answer the following question: does relying only on distributed representations improve the performance of neural models? Therefore, we evaluate the performance of hierarchical networks and state-of-the-art Transformers. Second, we evaluate the impact of using the best trained network for computing the reward function within a POMDP dialogue framework (Ultes et al., 2017b). We would like to determine how realistic it is to incorporate user satisfaction estimators that rely solely on distributional semantics in reinforcement learning (RL) dialogue systems. This approach can be used for instance to train satisfaction predictors from large human-human chats, in which satisfaction has been self-scored by users.

The case study is the English LEGO corpus of human-machine spoken conversations (Schmitt et al., 2012), which has been annotated at each system turn with the Interaction Quality (IQ), ranging from 1 (poor quality) to 5 (good quality). Our results suggest that distributed representations do outperform state-of-the-art models trained on fine-tuned features. We also show that using IQ estimators in the reward function greatly improves the task success rate for dialogues in the same domain the networks were trained on, which in this case is the Let’s Go domain (Raux et al., 2005).

Networks for Estimating the User Satisfaction

We study three distinct neural networks: hierarchical bi-directional Gated Recurrent Units (GRUs) (Cho et al., 2014) with attention, Transformers for generating contextual embeddings that feed a GRU layer and solely Transformers (Vaswani et al., 2017). We are interested in studying the impact of context-length in transformers, because real dialogues can easily attain a context of thousands of tokens (Table 3.11). Therefore, we explore BERT (Devlin et al., 2018), DistilBERT (Sanh et al., 2019) and Transformers eXtra-Large (Transformers-XL) (Dai et al., 2019).

(i) Hierarchical GRUs: Figure 3.6(a) shows the network. It has a Bidirectional GRU layer (Bi-GRU) at the lower level that returns the turn representation h_{tk} . Attention is used to weight relevant units in the turn hidden representation. Then, a GRU layer is then used to process dialogues as a sequence of turns. The last layer is a Softmax that predicts the most probable IQ class from 0 up to 5.

(ii) Contextual embeddings + GRU: We investigate the use of transformers as turn representations thus we propose a BERT-like transformer, which inputs are the tokens of the turn. The turn representations then feed a GRU layer. The output of the GRU then feeds a Softmax layer for predicting the score.

(iii) Transformer: This network is depicted in Figure 3.6(c). It consists in a transformer that takes as input the tokens tok_1, \dots, tok_D of the previous and current utterances. Then the output

3.2 Dialogue Manager

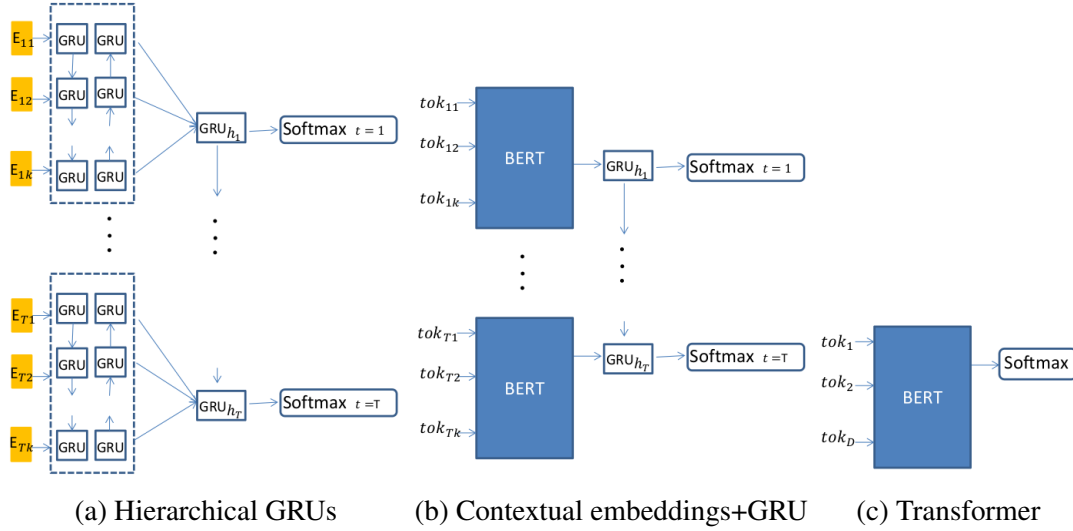


Figure 3.6: The proposed neural architectures for predicting the user satisfaction. E_{Tk} in (a) is the embedding for the k th token of the last turn ($t = T$). In (b) $tok_{T1} = [\text{CLS}]$ and $tok_{Tk} = [\text{SEP}]$ after applying the WordPiece tokenization to the last turn ($t = T$). In (c) tok_1, \dots, tok_D are the tokens of the dialogue after applying the WordPiece tokenization, in which $tok_1 = [\text{CLS}]$ and the token $[\text{SEP}]$ marks turn separation. tok_{D-1} is the last token of the last utterance in the dialogue and $tok_D = [\text{SEP}]$.

$[\text{CLS}]$ of the transformer feeds a Softmax layer for predicting the score of the current utterance.

We evaluate the prediction y_t at each system turn t . The back-propagation optimisation is done by minimising the cross-entropy loss function Tieleman and Hinton (2012) through stochastic gradient descent.

The Reward Function in a POMDP Dialogue System

Figure 3.7, shows the architecture connected to a reward estimator. In the Let’s Go bus-scheduled information system, the slots might be `origin` for the bus departure place and `time` for the departure time. The state s might record the current value and confidence level of each slot. From the state, a *belief state* b is extracted and an action a is decided based on a dialogue *policy*. Once the appropriate action is determined, it is converted to a textual message and then rendered by a speech synthesiser.

The reward function most commonly adopted for task-oriented dialogues penalises every dialogue turn with -1 and sums a reward of $+20$ at the end of the dialogue whenever the system provided the right information to the user or 0 otherwise (Eq. 3.11) (Gašić and Young, 2013).

$$R_{TS} = T \cdot (-1) + \mathbb{1}_{TS} \cdot 20 \quad (3.11)$$

In this work the reward estimator is based on the IQ as defined in (Ultes, 2019).

$$R_{IQ} = T \cdot (-1) + (iq - 1) \cdot 5 \quad (3.12)$$

Where R_{IQ} describes the final reward, iq is the IQ value predicted by the classifier (Section 3.2.2), which is a number from 1 to 5, where 1 represents poor quality and 5 good quality.

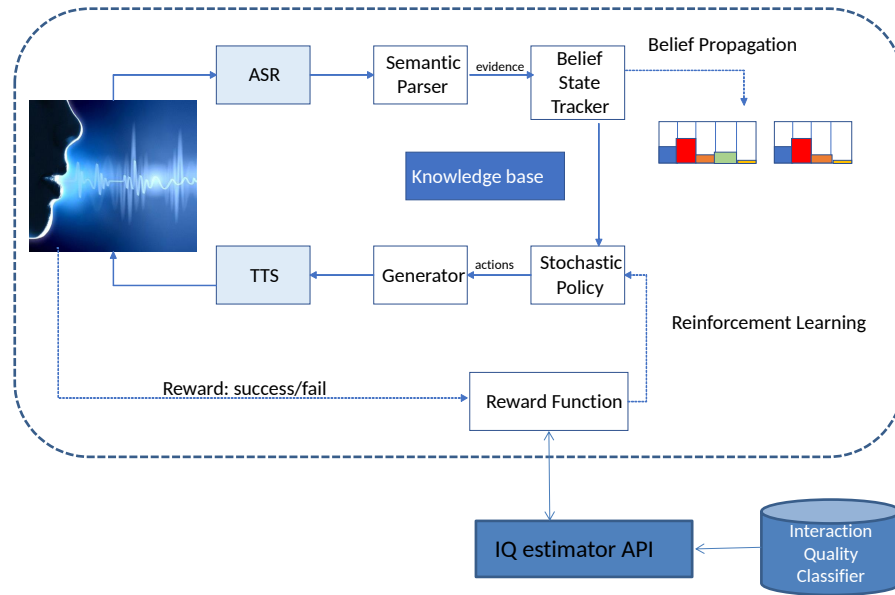


Figure 3.7: A POMDP spoken dialogue system with the interaction quality estimator.

We used PyDial (Ultes et al., 2017b), the publicly available POMDP dialogue framework and we implemented an application programming interface (API) that returns the IQ estimation predicted by the neural models presented above. Usually, RL systems first learn the policy on a simulated user until an optimal performance is reached, then they are ready to be tested by humans.

Experiments

In this section we introduce the corpus as well as describe the experiments and the evaluation metrics.

The Dataset The LEGO corpus collects *spoken* dialogues between users and the Let’s Go dialogue system Raux et al. (2005), which provides bus schedule information to the Pittsburgh population during off-peak times. 400 dialogues in the corpus have been manually annotated with the IQ score Schmitt et al. (2012). Since conversations in LEGO are system-initiative they have a lot of system interactions such as misunderstandings, confirmations and repetitions, producing quite long dialogues (i.e. hundreds of turns).

The complexity of the corpus is presented in Table 3.11, containing long dialogues of up to 200 turns with up to 76 tokens per turn.

User Satisfaction Estimators We compared our networks with the networks presented in Ultes (2019) for predicting the IQ with the following evaluation metrics: the unweighted average recall (UAR), which is the arithmetic average of all class-wise recalls, as well as a linearly weighted

3.2 Dialogue Manager

N. Dialogues	Dialogue length (max/mean/median)	Max.turn length (max/mean/median)	Max. toks p/dial. (max/mean/median)
400	200/65/53	76/26/24	6590/1444/1089

Table 3.11: The LEGO corpus with IQ annotations. Its complexity is measured by the maximum dialogue length (number of turns per dialogue), the maximum turn length (number of tokens per turn) and the maximum number of tokens per dialogue.

version of Cohen’s κ and Spearman’s ρ . The experiments were conducted in a 10-fold cross-validation, assuring that the same dialogue did not slip into different folds (i.e. dialogue-wise cross validation). We studied the length of the **dialogue context**: the turn for which we are predicting the score and the previous turns. We vary the context length and find an optimal context length of up to 100 turns per dialogue for BiGRUs.

Table 3.12 shows that the BiGRUs network trained on word embeddings outperformed the state-of-the-art (BiLSTM+att) networks in all performance measures, obtaining an absolute improvement of +1 for UAR, +9 for κ and +2 for ρ . It is worth noting that the state-of-the-art (BiLSTM+att) networks were trained on fine-tuned *turn features*.

These results are encouraging and suggest that distributed representations impact positively the performance of satisfaction estimators in hierarchical networks. We would like to study in the next section whether these models can be used to predict task success in dialogue systems.

Predicting IQ			
Model	UAR	κ	ρ
SVM_feats Ultes (2019)	44%	53%	69%
BiLSTM+att_feats Ultes (2019)	54%	65%	81%
BiGRUs	55%	74%	83%
DBert+GRU	35%	57%	42%
Trans-XL(ctxt \approx 1K)	47%	59%	67%
Trans-XL(ctxt \approx 2K)	44%	58%	67%

Table 3.12: Performance of the proposed models. The BiGRUs with FastText embeddings outperform all the networks trained on fine-tuned *turn features*

BERT-based Transformers do not perform well for this task on this dataset when using them to get the turn representations. This can be explained by the large number of turns dialogues have (i.e., up to 200 turns), the large number of parameters a transformer needs and the quite short annotated dataset (\approx 400 dialogues). We first tried (BERT+GRU) and due to its large memory requirement, we could only learn weights of up to 7 dialogue turns per dialogue in a cluster of 32GB-GPU machines (all the other turns representations were frozen). However, with DistilBERT we could treat up to 15 turns per dialogue while maintaining the same performance. Fortunately, we could process larger contexts with Transformers-XL, reaching an optimal performance with \approx 1K dialogue tokens. The results of Transformers-XL are comparable with the SVM baseline trained on fine-tuned features, yielding a better UAR (+3) and κ (+6) as well as a slightly lower ρ (-2). Having larger contexts \approx 2K do not seem to impact significantly their performance. These results suggest that in 32GB-GPU nodes the large context length (i.e. up to 6.5K tokens per

dialogue) is affecting transformers performance as they will require a prohibited usage of GPU memory to process the whole context.

The Impact on the Reward Function We evaluated the impact of the IQ estimators presented in Section 3.2.2 on the reward function for the Let’s Go (LetsGo) domain by using PyDial (Ultes et al., 2017b).

The Let’s Go dialogue system provides information about bus time-schedule according to the constraints: `origin`, `destination`, `time` and `route`, which corresponds to LetsGo(4). LestGo(6) also considers `origin neighbourhood` and `destination neighbourhood`. It is important to note that the Let’s Go domain is far more complex in terms of the number of database items than other domains available in PyDial.

The experiments run on simulated dialogues as in (Ultes, 2019; Casanueva et al., 2017b). We implemented a template-based generator for the user and the system utterances for the Let’s Go domain because our models rely on textual inputs (i.e. distributed representations). We compared our models in an environment without noise because unlike (Casanueva et al., 2017b) and (Ultes, 2019) the simulator in this work runs at the surface-level and not at the semantic-level and the noise used for User-Simulation in PyDial alters the semantic-dialogue acts regardless the surface form. In addition, we would like to apply these methods to chatbots, thus simulating ASR noise would not be appropriate and studying a more appropriated noise is out of the scope of this paper.

Domain	Reward	Task Success Rate(\uparrow)	Average Turns(\downarrow)
LetsGo(4)	R_{TS}	99% \pm 1.4	6.5 ± 0.96
	R_{IQ}	99% \pm 2.1	6.3 \pm 0.8
LetsGo(6)	R_{TS}	$81\% \pm 7.78$	11 ± 1.09
	R_{IQ}	97% \pm 3.98	8.9 \pm 1.26

Table 3.13: Task success rate of the simulated experiments for the Let’s Go domain over three runs with distinct seeds. Each value is computed after 1000 training dialogues/100 evaluation.

We used a policy model based on the GP-SARSA algorithm (Gašić and Young, 2013), which is a sample efficient Gaussian process approximation to the value function. We used the focus tracker (Henderson et al., 2014b) for belief tracking. The policy decides on summary actions of the dialogue state tracker which are based on dialogue acts (e.g., *request*, *inform* or *confirm*). The *task success rate* was the metric used to measure the dialogue performance (Casanueva et al., 2017b).

We observe in Table 3.13 that the reward computed with R_{IQ} outperforms the classical R_{TS} reward when having more constraints, namely LetsGO (6). Moreover, dialogues rewarded by R_{IQ} tend to be significantly shorter. Although there is not significant distinction between R_{TS} and R_{IQ} in terms of the task success for LetsGO (4), dialogues are slightly short with R_{IQ} . We also conducted preliminary experiments on domain transfer by evaluating R_{IQ} on the Cambridge Restaurants domain, obtaining a success rate of 44 ± 25.7 , compared to 99 ± 1.83 for R_{TS} . Unsurprisingly, R_{TS} and feature-based R_{IQ} (Ultes, 2019) are more robust to unknown domains than embedding-based R_{IQ} because both task-success and quality features are domain-agnostic.

For more details, please refer to (Rojas-Barahona, 2020).

3.2 Dialogue Manager

3.2.3 Imitation Learning

This section introduces the work of the PhD candidate Thibault Cordier, who I co-supervised together with Dr.Tanguy Urvoy and Professor Fabrice Lefevre. This work explores imitation learning for learning the policy on single domains. It has been published in the NeurIPS workshop Human in the loop dialogue systems (Cordier et al., 2020).

Deep RL (DRL) (Li, 2018) has achieved significant success on many complex decision-making problems and in particular in conversational AI (Gao et al., 2018). The ability to learn from few interactions is essential in dialogue applications because human interactions are scarce and costly. Unfortunately, standard RL algorithms usually require a large amount of interactions with the environment to reach good performances. One solution to speedup the learning process is to guide the agent’s exploration like stochastic learning policies such as soft-kind (Harnoja et al., 2017, 2018; Gao et al., 2019). The first question we address here is: (i) can dialogue policy learning be improved with stochastic off-policy learning methods?

Several methods search a way to exploit demonstrations to accelerate the learning with the conviction that demonstrations are the solution to *the sparse reward* (Hester et al., 2017). They can be based on Imitation Learning (IL) to learn an ”optimal” policy function or on Inverse Reinforcement Learning (IRL) to learn an ”optimal” reward function. Demonstrations may be useful to guide efficiently the exploration. We consider that human expertise can be used in different ways. The most classical one is to use demonstrations that humans have already produced. Another way is to use a rule-based agent, namely handcrafted, that has been designed, evaluated and fine-tuned by humans. Indeed, (Casanueva et al., 2017c) have shown that handcrafted approaches still perform better than policy learning approaches.

Moreover it is well known that human interactions and manually crafted rules are not only costly but also time consuming; while simulated interactions are cheaper and easier to collect (Su et al., 2016; Schatzmann et al., 2006). Therefore, the second issue that raised in our work is: (ii) can we use demonstrations without supervision and only as a way to guide exploration in an on-line learning?

Proposed Approach

The proposed approach to policy learning is based on Boltzmann sampling, and to which we integrate demonstrations directly into the RL process in order to better guide exploration and make relevant exploitation. RL is usually implemented as either value-based method as Q-learning or policy-based method as actor-critic. Both are our baselines with which we will make our contributions.

Baselines Q-learning is the combination of Double and Duelling Deep-Q-Network with Experience Replay (DQN/D3QN) (van Hasselt et al., 2015; Wang et al., 2016b). In short, DQN is an approximation function that searches to estimate the optimal state-action value function (or Q-value). The Double DQN architecture is an alternative method that mitigates the problem of overoptimistic value estimation and the Duelling DQN architecture is for learning more efficiently by decoupling value and advantage functions. On top of that, experience replay can be added for reducing sample correlation and for improving sample efficiency.

The actor-critic baseline is the Trust Region Policy Optimisation for Actor-Critic with Experience Replay (A2C/ACER/TRACER) (Wang et al., 2017; Weisz et al., 2018b). In brief, an

actor π tries to maximise the expected reward when in the same time a critic Q learned separately evaluates the actor decisions.

The importance sampling truncation with bias correction is used to correct the perceived sampling distribution induced by experience replay in order to reduce variance. We apply the Retrace algorithm (Munos et al., 2016) to recursively estimate the advantage function in safe and efficient way with small bias and variance. Finally, we use the trust region policy optimisation (TRPO) (Schulman et al., 2015) for adjusting the policy gradient in order to learn in a safe parameters region limiting the deterioration of the policy performance.

Exploration Strategy

The exploration strategy learns and plays a stochastic policy related to energy-based model. We propose to learn an energy-based policy in the dialogue environment where the agent samples his actions according to Boltzmann’s stochastic sampling (Haarnoja et al., 2017). We opt for using energy-based policies of the following form:

$$\pi(a_t|b_t) \propto \exp(-\mathcal{E}(b_t, a_t)) \quad (3.13)$$

When using value-based method, the energy function can be represented by the Q-function with parameter τ , the temperature, where we set $\mathcal{E}(b_t, a_t) = -\frac{1}{\tau}Q^\pi(b_t, a_t)$. When using policy-based method, the energy function is directly represented by the policy network and so is learned implicitly.

A stochastic sampling can be achieved by the Boltzmann sampling. Contrary to the commonly used strategy as ϵ -greedy where actions a_t are sampled from $\arg\max_a (1 - \epsilon) \arg\max_a \pi(a|b_t) + \epsilon U(\mathcal{A})$ where $U(\mathcal{A})$ is an uniform distribution over action space, the Boltzmann sampling strategy samples actions from $\exp(-\mathcal{E}(b_t, a_t))$, hence $a_t \sim \pi(a_t|b_t)$.

One of its advantages is that policy learning is less influenced by the policy function changes. Conversely, the ϵ -greedy sampling faces sudden jumps in action choices due to the argmax operator. So in theory, the Boltzmann strategy can make learning more stable than the ϵ -greedy.

Another advantage is that the temperature parameter can control the exploration-exploitation balance. So, to counter the weakness of its exploitation, it can be interesting to define correctly the temperature.

In practice, we decide to add random exploration in such a way that the actions are sampled according to $a_t \sim (1 - \epsilon)\pi(a_t|b_t) + \epsilon U(\mathcal{A})$, namely ϵ -Boltzmann sampling, with decreasing ϵ parameter in order to explore enough before following the stochastic policy with a fixed τ temperature parameter.

Imitation Learning Strategies

During the reinforcement learning process, demonstrations can serve as an efficient way to explore the environment. Indeed, they can lead the agent to receive rewards promptly and so can lead it to exploit confident winning trajectories quickly. A handcrafted agent is used to simulate near-optimal demonstrations and feed-backs. It offers good performance compared with the other deep learning methods. Also, it has been designed, evaluated and fine-tuned by humans. Thus it is a way to mimic human expertise and can be served as a near-optimal expert in our experimentation.

3.2 Dialogue Manager

We propose two demonstration sampling strategies corresponding to two different ways of using the knowledge of an expert.

A) Learning with demonstrations: Let us assume that we learn with an offline expert i.e. demonstrations are given before learning. We propose that in β (in percent) of dialogues, the agent plays for itself. Otherwise in $1 - \beta$ of dialogues, the expert gives to the agent one of its expert trajectories as demonstration and the agent replays the dialogue as if it was the one who played it. Therefore, the agent learns as if they are two datasets. The first one contains its trajectories and the second the expert demonstrations.

B) Learning with feed-backs: Let us assume that we learn with an online expert i.e. demonstrations are given during learning. We propose that in β (in percent) of dialogue actions, the agent plays for itself. Otherwise in $1 - \beta$ of dialogue actions, the expert gives to the agent its expert action as feed-back and the agent plays the dialogue action as if it was its choice.

This technique let the agent explore the environment by playing relevant actions in a given state. In other words, it learns about its trajectories in which the expert can redirect it at any moment to more relevant action, as DAgger does (Ross et al., 2011).

Experiments

In our experiments the Pydial framework (Ultes et al., 2017c) is used, which implements an agenda-based user simulation (Schatzmann et al., 2007). As in (Casanueva et al., 2017c) we tested our algorithms for policy learning on different domains and in different environments by increasing the inputs' noise. The domains in PyDial differ from each other by the ontology size, impacting the state and action space dimensions.

We evaluate the learned policies according to three levels of noise with respect to the semantic error rate (SER). This corresponds to the noise that comes from the ASR and the NLU channels. In Pydial, this is modelled at the semantic level whereby the true user action is corrupted by noise to generate an N-best-list with associated confidence scores.

Table 3.14 shows the compared policy models. HDC corresponds to the handcrafted policy learning, which is a rule-based approach written by experts. DQN and ACER are the baselines enhanced with stochastic (stoc) exploration and either behaviour cloning (BC) or feedbacks (FB).

Method name	Abbrev.
Handcrafted Policy	HDC
Stochastic Q-learning	Stoc-DQN
Stochastic Q-learning with Demonstrations	Stoc-DQN-BC
Stochastic Q-learning with Feed-backs	Stoc-DQN-FB
Stochastic Actor-Critic	Stoc-ACER
Stochastic Actor-Critic with Demonstrations	Stoc-ACER-BC
Stochastic Actor-Critic with Feed-backs	Stoc-ACER-FB

Table 3.14: Overview of proposed methods

We performed a long training stage over 10 000 dialogues. This will evaluate the contribution of stochastic sampling strategy during training and testing stages. Here we search to answer the question: can we improve dialogue policy learning with stochastic off-policy learning methods in order to compete the handcrafted agent? All methods are evaluated after training over 1 000 dialogues during which the learned policy is fixed. For the second experiments, we decide to compute the average performance over the last five checkpoints from dialogue indices 6 000 to 10 000 with a step of 1 000 dialogues. This calculation is done in order to reduce variance induced by RL when we estimate the performance of the models.

Results

The results are presented in Table 3.15. In most environments, methods learn very well compared to the benchmarks (Casanueva et al., 2017c). Furthermore, some of them can compete with the handcrafted agent whether it is a stochastic Q-learning approach or a stochastic actor-critic approach. For instance, Stoc-DQN-BC outperforms handcrafted expert with 30% SER for laptops and SFR. Stochastic ACER was more robust to the different environments, showing better performance, particularly for Stoc-ACER-BC with 30% SER for laptops.

<i>Task</i>		Stoc-DQN		Stoc-DQN-BC		Stoc-DQN-FB		HDC	
		Suc.	Rew.	Suc.	Rew.	Suc.	Rew.	Suc.	Rew.
0% SER	CR	97.48%	12.67	98.34%	13.30	98.88%	13.46	100.0%	14.00
	SFR	94.18%	10.99	87.40%	9.58	95.62%	10.94	98.2%	12.40
	LAP	95.08%	11.10	98.10%	11.76	98.40%	11.77	97.0%	11.70
15% SER	CR	92.22%	10.29	94.16%	11.26	95.76%	11.64	96.7%	11.00
	SFR	90.64%	8.56	88.70%	8.06	89.22%	8.21	90.9%	9.00
	LAP	90.34%	8.59	92.56%	9.40	91.86%	9.19	89.6%	8.70
30% SER	CR	84.32%	7.73	85.36%	8.64	85.46%	8.65	89.6%	9.30
	SFR	82.48%	5.11	80.26%	5.05	80.34%	4.63	79.0%	6.00
	LAP	82.24%	5.89	84.62%	6.25	83.40%	6.00	76.1%	5.30
<i>Task</i>		Stoc-ACER		Stoc-ACER-BC		Stoc-ACER-FB		HDC	
		Suc.	Rew.	Suc.	Rew.	Suc.	Rew.	Suc.	Rew.
0% SER	CR	99.60%	14.02	99.64%	14.03	99.30%	13.87	100.0%	14.00
	SFR	97.66%	12.36	96.48%	11.98	96.34%	11.98	98.2%	12.40
	LAP	95.24%	11.23	95.34%	11.28	95.40%	11.24	97.0%	11.70
15% SER	CR	97.56%	12.69	95.80%	12.32	96.78%	12.59	96.7%	11.00
	SFR	88.62%	9.26	87.64%	8.91	86.98%	8.67	90.9%	9.00
	LAP	88.74%	8.72	88.00%	8.55	86.26%	8.24	89.6%	8.70
30% SER	CR	89.82%	10.19	89.38%	9.98	89.32%	10.04	89.6%	9.30
	SFR	72.74%	4.38	78.22%	5.45	71.02%	4.37	79.0%	6.00
	LAP	75.80%	4.86	78.66%	5.40	77.82%	5.23	76.1%	5.30

Table 3.15: Results of Experiment 2. Long term learning, average from 6 000 to 10 000 training dialogues, for 1 000 testing dialogue. Each bold result represent better models than the handcrafted agent.

These results are encouraging and suggest that stochastic sampling makes it possible to learn a policy that performs as well as the handcrafted agent, which has been designed and fine-tuned by

3.2 Dialogue Manager

humans, even in hard environments. Also, these results show that demonstrations can significantly contribute to improve the performance at early learning stages.

3.2.4 Hierarchical Imitation Learning

This section also presents the work of the PhD candidate Thibault Cordier, who I co-supervised together with Dr.Tanguy Urvoy and Professor Fabrice Lefevre. This work has been published in Cordier et al. (2022).

We explore Graph Neural Network (GNN) for learning the policy in multi-domain and multi-task environments, in which several domains and tasks can be evoked in the same conversation.

In practice, real applications like personal assistants or chatbots must deal with multiple tasks: the user may first want to **find** a hotel (first task), then **book** it (second task). Moreover, the tasks may cover several domains: the user may want to find a hotel (first task, first domain), book it (second task, first domain), and then find a restaurant nearby (first task, second domain).

One way of handling this complexity is to rely on a *domain hierarchy* which decomposes the decision-making process; another way is to switch easily from one domain to another by scaling up the policy.

Although *structured dialogue policies* can adapt quickly from a domain to another Chen et al. (2020), covering multiple domains remains a hard task because it increases the dimensions of the state and action spaces while the reward signal remains sparse. A common technique to circumvent this reward scarcity is to guide the learning by injecting some knowledge through a teacher policy ¹⁵.

We study how structured policies like *graph neural networks* (GNN) combined with some degree of *imitation learning* (Imitation Learning (IL)) can be effective to handle multi-domain scenarios.

We provide large scale experiments in a dedicated framework (Zhu et al., 2020) in which we analyse the performance of different types of policies, from multi-domain policy to generic policy, with different levels of imitation learning.

Dialogue State / Action Representations

One way of standardising the slot representation into a common feature space is to use Domain Independent Parametrisation (DIP) (Wang et al., 2015) parametrisation. We adopt DIP as state and action representations, which are not reduced to a flat vector but to a set of sub-vectors: one corresponding to the domain parametrisation (or *domain representation*), the others to the slots parametrisation (or *slot representations*). For any active domain, the input to the *domain representation* is the concatenation of the previous *domain* user and system actions (see examples of the output below, and a formal definition in Section 3.2.4), the number of entities fulfilling the user’s constraints in the database, the booleans indicating if the dialogue is terminated and whether an offer has been found / booked. The output corresponds to action scores such as REQMORE, OFFER, BOOK, GREAT, etc. Regarding the *slot representation*, its input is composed of the previous *slot-dependent* user and system actions (see output below), the booleans indicating if a value is known and whether the slot is needed for the *find / book* tasks. Its output are actions scores such

¹⁵ For deployment the teacher is expected to be a human expert, however, for experimentation purposes we used the handcrafted policy as a proxy (Casanueva et al., 2017c).

as INFORM, REQUEST and SELECT. The parameterisation used depends on the representation of the deterministic states of CONVLAB which does not consider the uncertainty in the predictions made by the *natural language understanding* (NLU) module.

Graph Neural Network

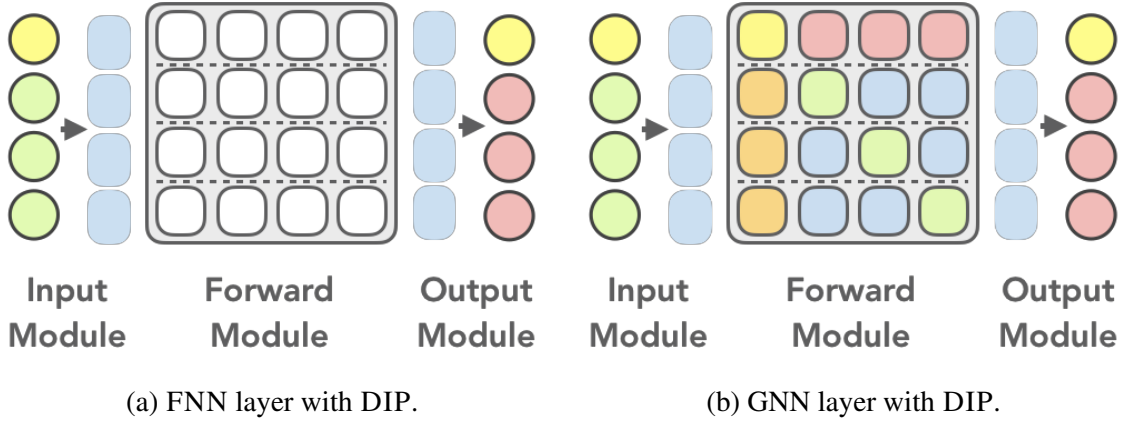


Figure 3.8: Structure of the layers with DIP. The central box represents the weight matrix of a layer W^l . It can be decomposed into sub-matrices $W_{i,j}^l$. The white sub-matrices represent any sub-weight and the coloured sub-matrices represent shared sub-weights. The circles represent input and output graph nodes. The *domain representations* are depicted in yellow; the *slot representations* in green and red.

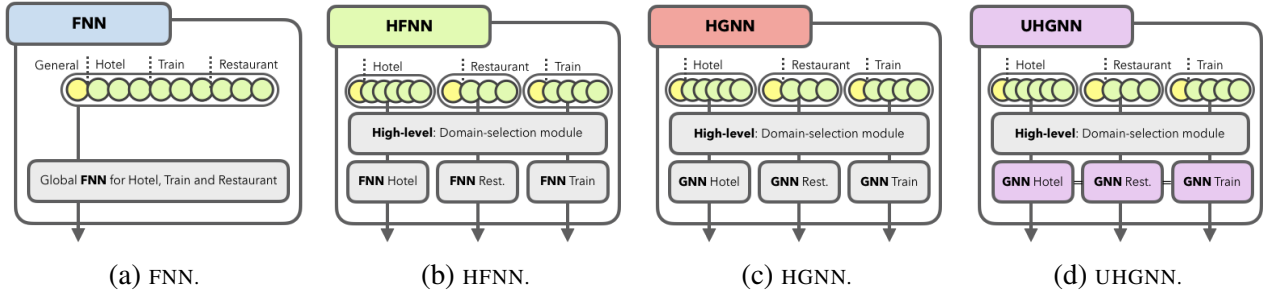


Figure 3.9: Policy and input data structures. Different levels of structure are presented from classical *feed-forward neural network* (FNN) to *graph neural network* (GNN). The prefix H- corresponds to a hierarchical policy and UH- corresponds to a unique sub-policy for all domains. For a FNN layer, the input data is the concatenation of all DIP slot representations. For a GNN layer, the input keeps its structure.

Prior knowledge can be integrated in our models by constraining the layer structure imposing symmetries in the neural policies. Without prior knowledge, the standard structure used is the *feed-forward neural network* layer (FNN) as represented in Figure 3.8a. This unconstrained structure does not assume any symmetry in the network.

3.2 Dialogue Manager

Assuming that sub-policies associated with the slots are the same, a better alternative is to use the *graph neural network* layer (GNN) presented in Figure 3.8b. This structure assumes that the state and action representations have a graph structure that are identically parameterised by DIP. The GNN structure is a fully connected and directed graph, in which each *node* represents a sub-policy associated with a slot and a directed *edge* between two sub-policies represents a message passing. We identify two roles for sub-policies: the general node as I-NODE associated to the *domain representation* and the slot nodes denoted as S-NODE associated to the *slot representations*. Both representations were introduced in Section 3.2.4. We also identify the relations: I2S for I-NODE to S-NODE, S2I and S2S respectively.

We formally define the GNN structure as follows. Let n be the number of slots and L the number of layers. Let x be the dialogue state, $\mathbf{x}_0 = \phi_0(x)$, $\mathbf{h}_0^l \forall l \in [0, L-1]$ and \mathbf{y}_0 be respectively the input, hidden and output I-NODE representations. Let the input, hidden and output S-NODES representations be respectively $\forall i \in [1, n]$, $\mathbf{x}_i = \phi_i(x)$, $\mathbf{h}_i^l \forall l \in [0, L-1]$ and \mathbf{y}_i . First, the GNN transforms inputs:

$$\begin{aligned} \forall i \in [0, n], \quad \mathbf{h}_i^0 &= F_i^0(\phi_i(\mathbf{x})) \\ \text{with } F_i^0(\mathbf{h}) &= \sigma^0(\mathbf{W}_i^0 \mathbf{h} + \mathbf{b}_i^0) \end{aligned} \quad (3.14)$$

Then, at the l -th layer, it computes the hidden nodes representations (Eq. 3.15a) by following message sending¹⁶ (Eq. 3.15b), message aggregation (Eq. 3.15c) and representation update (Eq. 3.15d):

$$\forall i \in [0, n], \quad \mathbf{h}_i^l = F_i^l(\mathbf{h}^{l-1}) \quad (3.15a)$$

$$\mathbf{m}_{i \leftarrow j}^l = M_{i \leftarrow j}^l(\mathbf{h}_j^{l-1}) \quad (3.15b)$$

$$\mathbf{m}_i^l = A_i^l(\mathbf{m}_{i \leftarrow *}) \quad (3.15c)$$

$$\mathbf{h}_i^l = U_i^l(\mathbf{m}_i^l) = \sigma^l(\mathbf{m}_i^l) \quad (3.15d)$$

The message sending function $M_{i \leftarrow j}^l$ is a linear transformation with bias. The message aggregation function A_i^l is the average pooling function. The representation update function U_i^l compute the new hidden representation with RELU activation function and dropout technique during learning stage. Finally, the GNN concatenates (\oplus symbol) all final nodes representations and computes the policy function with the Softmax activation function.

$$\mathbf{y} = \sigma^L\left(\bigoplus_{i=0}^n \mathbf{W}_i^L \mathbf{h}_i^{L-1} + \mathbf{b}_i^L\right) \quad (3.16)$$

Imitation Learning

In addition to the structured architecture, we use some level of IL to guide the agent’s exploration. In our experiments, we used CONVLAB’s handcrafted policy as a *teacher*¹⁵, but other policies could be used as well. *Behaviour cloning* (BC) is a pure supervised learning method that tries to mimic the teacher policy. Its loss function is the cross-entropy loss as in a classification problem. *Imitation Learning From Oracle Demonstrations* (ILFOD) is a RL method which allows the agent

¹⁶The notation $i \leftarrow j$ denotes a message sending from slot j to slot i . It also corresponds to the directed relation between the slots j and i . The notation $i \leftarrow *$ denotes all messages sending to slot i .

to play oracle actions as demonstrations and to inject them in its *replay buffer*. The same presented in Section 3.2.3. In our experiments, we kept half of the agent’s own actions in the buffer along with those generated by the oracle. *Imitation Learning From Oracle Supervision* (ILFOS) is the combination of supervised and reinforcement learning when the agent learns with a supervised loss, namely the margin loss Hester et al. (2018).

Experiments on GNN and Imitation Learning

We performed an ablation study: (i) by progressively extending the baseline to our proposed GNNs and (ii) by guiding the exploration with IL. All the experiments were restarted 10 times with random initialisations and the results evaluated on 500 dialogues were averaged. Each learning trajectory was kept up to 10,000 dialogues with a step of 1,000 dialogues in order to analyse the variability and stability of the methods.

Models The baseline is **ACER** which is a sophisticated actor-critic method (Wang et al., 2016a). After an ablation study, we progressively added some notion of hierarchy to FNNs to approximate the structure of GNNs. **FNN** is a feed-forward neural network with DIP parametrisation. Thus, the agent actions are single-actions. **FNN-REF** is a FNN with the native parametrisation (no DIP) with multiple-actions of CONVLAB¹⁷. **HFNN** is a hierarchical policy with domain-selection module and based on FNNs for each domain. **HGNN** is a hierarchical policy with domain-selection module and based on GNNs. **UHGNN** is a HGNN with a unique GNN for all domains.

Metrics: We evaluate the performance of the policies for all tasks. For the find task, we use the precision, the recall and the F-score metrics: the **inform rates**. For the book task, we use the accuracy metric namely the **book rate**. The dialogue is marked as **successful** if and only if both inform’s recall and book rate are 1. The dialogue is considered **completed** if it is successful from the user’s point of view (*i.e* a dialogue can be completed without being successful if the information provided is not the one objectively expected by the simulator).

Evaluation of the Dialogue Manager We performed an ablation study based on ACER as reported in Figure 3.10. First, all RL variants of ACER (Figure 3.10a) have difficulties to learn without supervision in contrast to BC variants (Figure 3.10b). In particular, we see that hierarchical decision making networks (HFNN in green), graph neural network (HGNN in red) and generic policy (UHGNN in purple) drastically improve the performance compared to FNNs. Similarly, using IL like ILFOD (Figure 3.10c) and ILFOS (Figure 3.10d) notably improves the performance. Therefore, learning generic GNNs allows collaborative gradient update and efficient learning on multi-domain dialogues.

Conversely, we observe that hierarchical decision making with HFNNs does not systematically guarantee any improvement. These results suggest that GNNs are useful for learning dialogue policies on multi-domain which can be transferred during learning across domains on-the-fly to improve performance. Finally, regarding ILFOD variants (Figure 3.10c), we can observe that all architectures are affected by a large variability. This shows that multi-domain dialogue management is difficult despite the use of demonstrations and that learning with reward is not sufficient to robustly succeed.

¹⁷The native parametrisation manually groups multi-actions based on MULTIWOZ Budzianowski et al. (2018).

3.2 Dialogue Manager

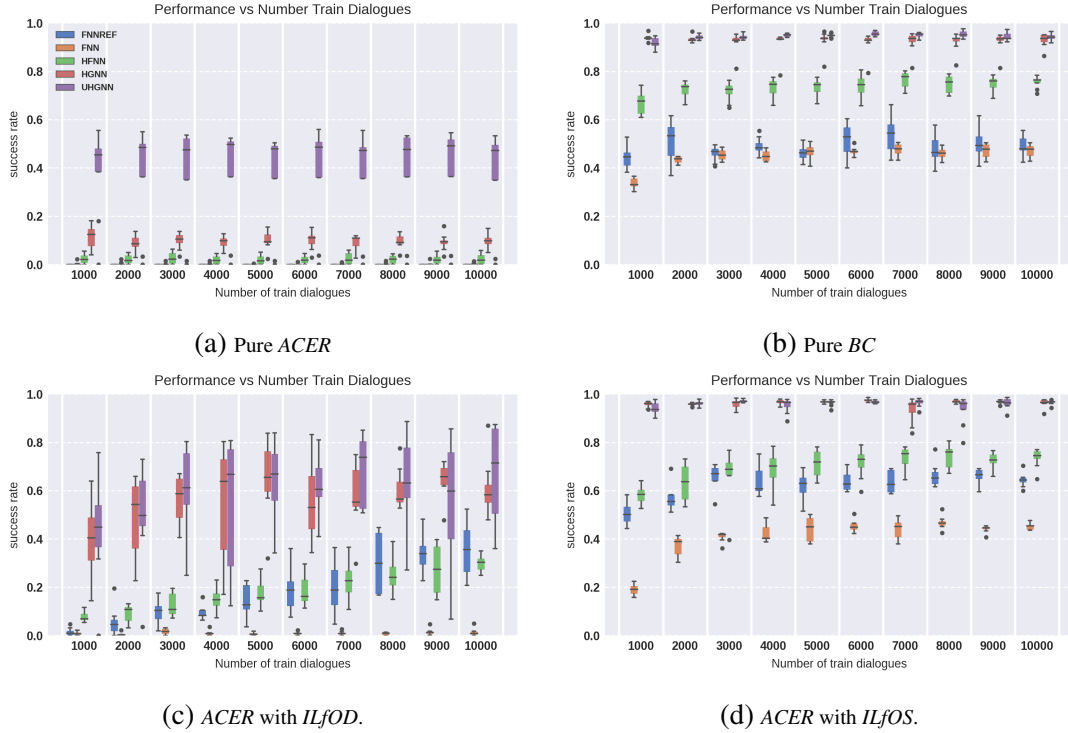


Figure 3.10: Distribution via boxplot of the performance of the proposed approaches on CONVLAB, with 10 different initializations and without pre-training. The coloured area represents the interquartile Q1-Q3 of the distribution, the middle line represents its median (Q2) and the points are outliers.

Evaluation of the Dialogue System We evaluate the policy learning algorithms in the entire dialogue pipeline, in particular our best DM policy ACER-ILFOS-UHGNN under a shorter name **ACGOS**.

The results of our experimentation are presented in the paper (Cordier et al., 2022). We observe that the performance of our approach is closed to the handcrafted policy (the teacher) when directly passing the dialogue acts, when using BERT NLU Devlin et al. (2018) and template-based NLG. Moreover, the performance of our approach is better than baselines with a significant difference. These results highlight the benefit of structured policies against standard policies.

A limitation of current policies in CONVLAB, including ours, is that the robustness to noisy inputs is not specifically addressed as it had been done in PyDial Ultes et al. (2017b). It could be also interesting to study the impact of incorporating real human feed-backs and demonstrations instead of a handcrafted teacher.

The GNN structured policies combined with imitation learning avoid sparsity, while being data efficient, stable and adaptable. They are relevant for covering multi-domain task dialogue problems.

A continuation of this work for few-shot learning will soon be published in the findings EACL2023, for more details please refer to (Cordier et al., 2023).

Chapter 4

Contributions to Conversational QA and Other Contributions

‘Estaba perdiendo la vista y el oído, parecía confundir a los interlocutores con personas que conoció en épocas remotas de la humanidad, y contestaba a las preguntas con un intrincado batiburrillo de idiomas.’ — *Gabriel García Márquez, Cien años de soledad.*

‘He was losing his sight and his hearing, he seemed to confuse the people he was speaking to with others he had known in remote epochs of mankind, and he would answer questions with a complex hodgepodge of languages.’ — *Gabriel García Márquez, One Hundred Years of Solitude.*

Conversational QA is a relatively recent area of research that groups reading comprehension, QA and dialogue. Typically, it consists in a sequence of questions and answers related to a paragraph or to a knowledge graph. My contribution in this field was to enrich existing datasets with (i) information about the ellipsis and coreferences (Section 4.1) and (ii) question rewriting (Section 4.2). I also contribute to the creation of a new dataset Knowledge-base Conversations (KGConv) (Section 4.3), that we will soon make public. This work was part of the industrial research project DIANA, of which I was the head. The work described in here was made mainly by the young researcher Quentin Brabant under my supervision. For the work presented in (Section 4.2), I proposed the idea of enriching the corpus CoQA with question rewriting in 2020. I wrote with Timothy Garwood a document describing the annotations and we manually annotated 10 dialogues. Then, I was in charge of the administrative process to formalise the collaboration with ELRA to produce the annotations, this took about a year. I also checked the annotations and conducted experiments on conversational question answering by using RoBERTa. Regarding this work, Gwenole Lecorve worked on question generation, while Quentin Brabant on question rewriting. The work presented in Sections 4.1 and 4.3 were joint work with Claire Gardent as part of the European Innovative Training Networks (ITN) project Interactive Natural Language Technology for Explainable Artificial Intelligence (NL4XAI), that involves industrial partners as Orange and academic institutions as the CNRS. I was representing Orange in this project.

4.1 Detection of Ellipsis and Co-reference in conversational corpora

This work is joint work with Quentin Brabant and Claire Gardent. It was published in (Brabant et al., 2021). We made several contributions to the task of ellipsis and coreference detection in conversational corpora. We created labelled data by enriching three existing datasets with annotations indicating whether a turn contains an ellipsis and/or a coreference. As these annotations were incomplete, we drew on inferential relations between incompleteness, pronominalisation, ellipsis and coreference to both extend and complement these annotations. We then use these annotated data to train a classifier based on DistilBERT (Sanh et al., 2020), which assigns to each question in a conversation two labels indicating whether it contains an ellipsis and/or a coreference. We also explore how active learning, multilabel approaches and fine-tuning can be used to train this model.

A *coreference* occurs when an entity is referred via two or more expressions in the same conversation. However, we are only interested in detecting a particular kind of coreference. We say that a coreference happens in a turn if and only if (1) it contains an expression referring to an entity already mentioned in a previous turn and (2) this entity cannot be identified outside of the conversational context. The *resolution* of a coreference consists in replacing the referring expression by an unambiguous reference to the entity.

In linguistics, an *ellipsis* is the omission of one or several words from a clause that preserves the meaning in the context. When a turn is not understandable without its context (i.e. without the conversation history), we call it *incomplete*. In this paper, we assume that any conversation turn contains an ellipsis if and only if it is still incomplete after coreferences have been resolved. It follows from this definition that an incomplete sentence contains either a coreference, an ellipsis, or both.

A conversation is a sequence of alternating questions and answers that starts with a question and ends with an answer: $(q_1, a_1, q_2, a_2 \dots, q_n, a_n)$. In many available conversational question answering datasets questions are sentences produced by humans (e.g. (Choi et al., 2018a; Christmann et al., 2019; Elgohary et al., 2019; Quan et al., 2019; Reddy et al., 2019)), while answers are often given by an automated system, and often not in the form of a sentence. For this reason, we focus on ellipsis and coreference detection in questions. Moreover, we will sometimes use the term question to refer to turns that are not question per-se, but that are produced by a user and not by an automated system (see Section ??, GECOR dataset).

We propose a model to predict whether any given question q_i of a conversation $(q_1, a_1, \dots, q_n, a_n)$ contains an ellipsis and/or a coreference; since any turn can normally be understood based on the context of previous turns, our task can be seen as the classification of q_i with the given context $c = (q_1, a_1, \dots, q_{i-1}, a_{i-1})$. We thus formulate our task as a 2-labels classification: for a given input question q_i and an input context c , output two values $(coref, ellipsis) \in \{0, 1\}^2$ where 1 denotes the presence of the phenomenon and 0 denotes its absence. We call *instance* of our task the couple formed by a question, and its context. An instance is *annotated* when it is associated with an annotation of the form $(coref, ellipsis)$.

4.1 Detection of Ellipsis and Co-reference in conversational corpora

4.1.1 Active learning (AL)

We use the following values for annotating the datasets: 1 for the presence of a phenomenon (positive class), 0 for its absence (negative class). Cases where no label is assigned are denoted by the value -1. Note that -1 does not denote a class, but only the absence of information about the actual class. We describe how we process each dataset in order to obtain train instances for our task.

ConvQuestions Christmann et al. (2019). Many conversations of ConvQuestions are centered on the same entity; those conversations tend to be similar to each other, as they often have questions in common. In order to maximise the benefits of manual annotations, we created subsets of the original data containing exactly one conversation per topic entity. This resulted in train/dev/test sets containing respectively 905/330/335 questions in total. Based on these new sets, we created an instance of our task for each question (except the first one) of each conversation. Some of these conversations were manually annotated with (*coref*, *ellipsis*) values. We obtained train/dev/test of 247/329/331 annotated instances.

GECOR Quan et al. (2019). We create instances as follows. For each conversation $(q_1, a_1, \dots, q_n, a_n)$ in the GECOR dataset, each $i \in \{2, \dots, n\}$, and each variant $q'_i \in \{q_i(e), q_i(r), q_i(c)\}$ of the question q_i : if q'_i is not empty, then we create the instance $((q_1, a_1, \dots, a_{i-1}), q'_i)$ and annotate it with (*coref*, *ellipsis*) values. Those values can sometimes be deduced by using the following rules:

- $q_i(e)$ contains an ellipsis;
- $q_i(r)$ contains a coreference;
- $q_i(c)$ contains no ellipsis nor coreference;
- if $q_i(e) = q_i(r)$ we infer that both $q_i(e)$ and $q_i(r)$ contain an ellipsis and a coreference;
- if $q_i(e)$ is empty, we infer that q_i contains no ellipsis and thus $q_i(r)$ neither;
- if $q_i(r)$ is empty, we infer that q_i contains no coreference and thus $q_i(e)$ neither.

These rules are not sufficient to deduce ellipsis and coreference label values in all cases. By default, the value -1 is assigned.

CANARD Elghary et al. (2019). Instances were extracted similarly as from the GECOR dataset. The two main differences are: for each created conversation, two variants (original and complete) of the last question are used. When the complete variant is used, we assign 0 to both *coref* and *ellipsis*; otherwise, we assign -1. An example is given in Table 4.1.

At this point many labels are missing in the instances of the task. In particular, instances from CANARD do not contain any positive label. We addressed this issue via two approaches: multilabel learning and label filling.

Multilabel classification can be seen as a particular case of multitask learning, since a single model is trained on several binary classification tasks. One justification for using this approach (instead of one model per classification) is that the parameters are shared during training which has been shown in the literature to be beneficial to all the classifiers.

The 4-labels classification task considers the following labels: coreference, ellipsis, incompleteness, and pronoun detection. Formally, it means that annotations of the form (*coref*, *ellipsis*)

Piece of conversation from CANARD:	
q_1	What is On the Sunday of Life?
$q_1(c)$	What is On the Sunday of Life?
a_1	In 1992, Delerium released On the Sunday of Life as an edition of 1,000 copies, complete with a deluxe gatefold sleeve.
q_2	Did it do well?
$q_2(c)$	Did Porcupine Tree, On the Sunday of Life do well?
a_2	On the Sunday of Life... had accumulated sales of more than 20,000 copies.
q_3	Was it rereleased?
$q_3(c)$	Was Porcupine Tree, On the Sunday of Life rereleased?

Corresponding instances of the task:							
Context	Question	Coref	Ellipsis	Coref	Ellipsis	Incomp.	Pronoun
(q_1, a_1)	q_2	-1	-1	1	-1	1	1
(q_1, a_1)	$q_2(c)$	0	0	0	0	0	0
(q_1, a_1, q_2, a_2)	q_3	-1	-1	1	-1	1	1
(q_1, a_1, q_2, a_2)	$q_3(c)$	0	0	0	0	0	0

Table 4.1: Example of conversation from CANARD and the corresponding instances of the task. Columns with gray headers show the result of label filling.

4.1 Detection of Ellipsis and Co-reference in conversational corpora

are replaced by annotations of the form (*coref*, *ellipsis*, *inc*, *pronoun*). We used automatic pronoun detection to provide a 0 or 1 value to *pronoun* in all questions. By default, the value of *inc* is set to -1, except for instances from CANARD where the value is known.

We then replace some of the -1 values by taking advantage of the logical dependencies between labels: a pronoun always indicates a coreference; incompleteness is either due to a coreference or an ellipsis; coreferences and ellipses always cause incompleteness. We therefore applied the following rules to each instance, in order:

1. if *pronoun* = 1 then *coref* \leftarrow 1,
2. if *coref* = 1 or *ellipsis* = 1 then *inc* \leftarrow 1,
3. if *coref* = 0 and *ellipsis* = 0 then *inc* \leftarrow 0,
4. if *inc* = 0 then *coref* \leftarrow 0 and *ellipsis* \leftarrow 0.

Remark that in some cases these rules are not sufficient to get rid of all unknown values. Such cases can be found in the examples of Table 4.1.

Active learning (AL) is a human-in-the-loop method that aims at maximizing the performance gains relatively to the number of manual annotations. It is especially interesting when few labeled data are available and only a small fraction of unlabeled data can be manually annotated in reasonable time. We apply several rounds of AL for labeling (separately) ellipses and coreferences. Each round consists in the following steps:

1. *Train and evaluate a model.* We use CANARD/GECOR as a training set. All CANARD instances that have already been manually annotated during previous rounds are included. The evaluation is done on ConvQuestions test set.
2. *Run the model on unlabeled data.* The model trained in step 1 associates a prediction (*coref**, *ellipsis**) to each instance.
3. *Select a subset of unlabeled data.* We select the 50 CANARD conversations on which the model displays the least certainty. Since one conversation is the source of several instances, we define the certainty of a conversation as the average certainty of the corresponding instances. The certainty of the model (for a given label, on a given instance) is defined as the distance from 0.5 of the output corresponding to the predicted label value, i.e.: $|coref^* - 0.5|$ for coreference and $|ellipsis^* - 0.5|$ for ellipsis.
4. *Manually label the selected subset.* We label the selected conversations (either for ellipsis or coreference). Labeled conversations are used during training in the next loop.

We stop repeating these steps when the evaluation score stops increasing.

4.1.2 Experiments

We evaluate the following model variants.

- *Baseline.* The baseline is a DistilBert(Sanh et al., 2019) model trained on the 4-label classification task on CANARD/ GECOR.

- *Fine tuning only.* The model is fine-tuned on the 4-label classification task on the training set of ConvQuestions.
- *Baseline + AL.* The model is fine-tuned on the 4-label classification task on CANARD/GECOR, but labelled instances of CANARD are added via AL. Each round of AL adds 50 instances that are labelled for either coreference or ellipsis. We evaluate several versions of this variant: three versions use instances that were annotated for coreference via, respectively, 1, 2, and 3 rounds of AL. Three others versions use instances that were annotated for ellipsis via 1, 2, and 3 rounds.
- *Baseline + all AL.* Identical to baseline + AL, but using all annotations produced for coreference and ellipsis (3 rounds for each).
- *Baseline + all AL + fine tuning.* Identical to *Baseline + all AL.*, but training on CANARD/GECOR is followed by a fine-tuning step on the training set of ConvQuestions.
- *2-label variants.* We evaluate three of them. They are respectively identical to *baseline*, to *baseline + all AL*, and to *baseline + all AL + fine tuning*, with the difference that the model is trained on the 2-labels classification task.

We use GECOR and CANARD for training our models, while ConvQuestions is used for evaluation and fine tuning. In this way we can better assess how well the classifier behaves on unseen data, data that is different from the data the model was trained on. During training, labels with -1 value are simply ignored (no error is retro-propagated). During evaluation, we measure the recall, precision, and F-measure on ellipsis and coreference detection.

4.1.3 Results

The results are displayed in Table 4.2. Each line corresponds to a variant of the model.

Generally, the results show that coreference detection performs better than ellipsis detection. Moreover, by looking at lines 2 to 9 in the table, we see that AL is clearly beneficial; the *all AL labels* variant improves F1 scores for coreference and ellipsis detection by 10 and 13 points compared to the baseline. The same conclusion is drawn when comparing lines 11 and 12. The effects of training on 4 labels versus 2 are less clear: by comparing lines 2, 9, 10 to lines 11, 12, 13, we see that 4-labels variants perform roughly as well as their 2-labels counterparts on coreference detection. For ellipsis detection, they score significantly higher on F1 score when no fine tuning is applied, but the scores are too low to propose a meaningful interpretation. Fine tuning increases scores for both ellipsis and coreference detection; however the increase is way larger in the case of ellipsis. In fact, coreference detection arguably performs reasonably well without fine-tuning, contrary to ellipsis detection. A possible explanation is that the kinds of ellipses occurring in one dataset can be different from those occurring in another. In contrast, coreferences cover a narrower set of phenomena.

In addition to measuring performances, we looked at the output of the model on the test set: we noticed that coreferences due to pronouns use are well recognized, while many false negatives correspond to cases where an entity is referred to via its type or function, as in: “To which continent does Germany belong? What size is the country?”.

4.2 Question Rewriting

		Coreference			Ellipsis		
		P	R	F1	P	R	F1
1	fine tuning only	81	65	72	51	67	57
2	baseline	97	64	77	64	36	46
3	+ AL for ellipsis (1 round)	92	63	75	71	48	56
4	+ AL for ellipsis (2 rounds)	89	72	80	83	41	55
5	+ AL for ellipsis (3 rounds)	85	79	82	74	48	57
6	+ AL for coref. (1 round)	87	84	85	72	46	56
7	+ AL for coref. (2 rounds)	92	81	86	71	39	50
8	+ AL for coref. (3 rounds)	95	79	86	67	31	43
9	+ all AL labels	94	81	87	84	46	59
10	+ fine tuning	94	93	94	83	71	77
11	baseline, 2-labels variant	89	68	77	100	10	19
12	+ all AL labels	91	86	89	88	35	50
13	+ all AL labels + fine-tuning	94	93	93	84	70	76

Table 4.2: Results of the experiments. Scores are given as percentages.

4.2 Question Rewriting

As mentioned before, CQA (Reddy et al., 2019; Choi et al., 2018a; Saha et al., 2018) is a task in which a system interacts with a user. The interaction takes the form of a conversation, where the user always asks questions that the system answers. In this work, we focus on the case where the system searches for answers in a passage, although settings relying on structured data (e.g. knowledge bases) also exist (Saha et al., 2018), as the one presented in the following section (Section 4.3). Compared to QA, the system faces an additional difficulty: each question is asked in a *conversational context* that consists in previous turns. Therefore, implicit references to the context may happen in the form of ellipses and coreferences, making the understanding of questions more difficult for the system. One way to overcome this difficulty is Question Rewriting (QR), which consists in rewriting each original (*in-context*) question into an *out-of-context* question that is understandable by itself, i.e., that can be answered without knowing the conversational context.

We present the corpus Conversational Question Answering with Rewriting (CoQAR), which is an annotated subset of the CQA corpus CoQA (Reddy et al., 2019). CoQAR was obtained by asking specialised native speakers to annotate original questions with at least two and at most three distinct *out-of-context* rewritings. This work was published in (Brabant et al., 2022). Our contribution is two-fold.

Firstly, we provide CoQAR, which contains high-quality questions rewritings. The corpus is publicly available¹; moreover, its annotations were conducted in accordance to ethical concerns: every annotator involved was properly hired.

Secondly, we assess the quality of the annotations of CoQAR through several experiments. We train Question Rewriting (QR) models. We then rate these models’ outputs via human evaluation. We also evaluate these models as preprocessing steps of (conversational and non-conversational)

¹The COQAR dataset is publicly available at <https://github.com/Orange-OpenSource/COQAR>

	Number of rewritings				total
	0	1	2	3	
train	365	108	31,378	13,210	45,061
dev	9	0	37	7,937	7,983

Table 4.3: Number of questions depending on the number of rewritings.

QA models. To this end, we compare the performance of a state-of-the-art QA model with and without QR.

Our results support the claim of (Vakulenko et al., 2021) that QR models can be successfully used in combination with existing QA models. Indeed, we found that adding QR as a preprocessing step boosts the performances of QA models and allows reusing non-conversational state-of-the-art QA systems while reducing performance degradation on CQA.

4.2.1 Annotations

We decided to hire two specialised native-speakers’ annotators. Their task was to annotate original (*in-context*) questions from CoQA with at least two and at most three distinct *out-of-context* rewritings. To make sure that they understand what was expected, we ourselves annotated a conversation and provided it as an example. An example of conversation annotated by the annotators is provided in Table ??.

While annotators were told to preserve the meaning of the original sentences, they were also asked to paraphrase in their rewritings. As a result, these annotations contrast with those of CANARD, where the structure of the original question is usually preserved in the rewriting. In total, 4.1k conversations of CoQA train set were annotated as well as all 500 conversations of the dev set. Since the test set of CoQA is not available, no conversation was annotated from it. The train and dev sets of CoQAR respectively contain 45k and 8k questions. Table 4.3 summarises the number of questions that have 0,1,2 or 3 rewritings.

Overall, passages contain from 75 to 1079 words, with an average of 275. Conversation length distribution is displayed in Figure 4.1.

On average, out-of-context rewritings are longer (8.8 words) than the original questions (5.5 words); Figure 4.2 shows the question length distribution.

Most conversations were annotated by only one annotator, but 50 conversations were annotated by both. We relied on these conversations to analyse the annotations. We extracted two rewritings per question and per annotator and, using a pair of rewritings as references and the other as hypothesis, we computed the SacreBLEU score (Post, 2018) and the BERT-score (Zhang et al., 2020a). SacreBLEU gives us an insight on the similarity of the surface form of rewritings, while BERT-score gives us an insight on the semantic similarity. We obtained a SacreBLEU score of 32.67 and a BERT-score of 90.22: this suggests that the rewritings have diverse surface form while being close in terms of meaning.

4.2.2 Evaluation on Question Rewriting (QR)

In QR, the model receives as input an in-context question, its conversational context, and the associated passage. Its task is to generate an out-of-context rewriting of the question. We conducted

4.2 Question Rewriting

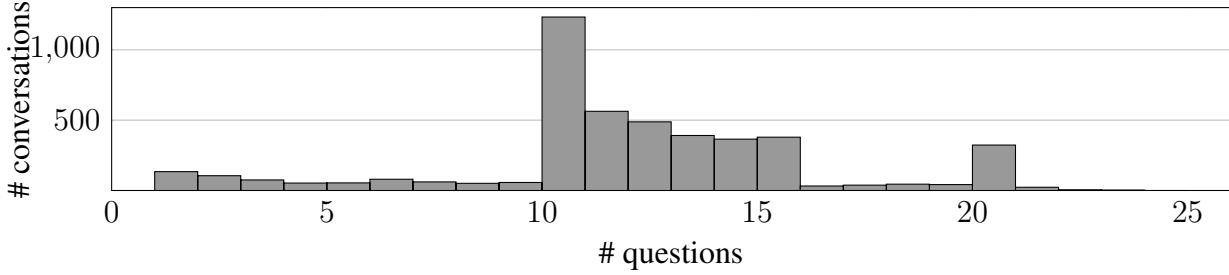


Figure 4.1: Distribution of conversations' length.

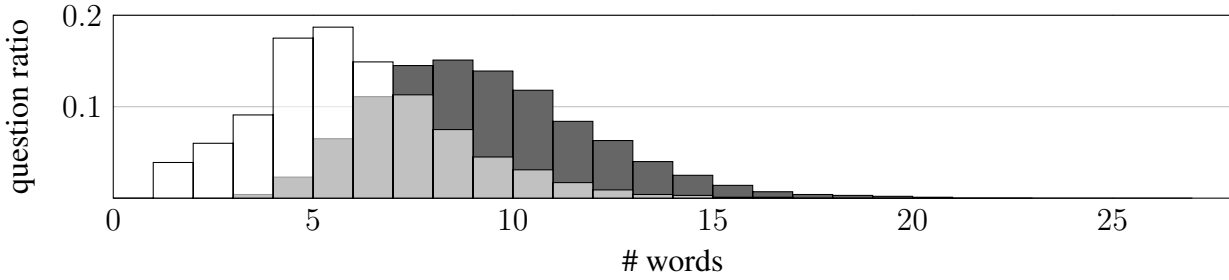


Figure 4.2: Distribution of length for original questions (white) and out-of-context rewritings (dark grey). Overlap of the distribution is light grey.

the following experiment: (1) training QR models on CoQAR and CANARD; (2) evaluating these models, via standard metrics and human evaluation. Furthermore, we evaluate these QR models on downstream conversational question answering as presented in the next section.

Datasets. For training and evaluation, we rely on CANARD and CoQAR. For CANARD, we use the original train/dev/test splits. For CoQAR, we use the original dev set as test set, and split the original train set into a train set and dev set, in such manner that CANARD and CoQAR dev sets have the same size. For training, we also make use of a mixture of CANARD and CoQAR, that we refer to as CoQAR+CANARD, whose train and dev sets are, respectively, the union of both corpora's train and dev sets. We train three variants of the QR model: one variant is trained on CANARD, one is trained on CoQAR, and the third one is trained on a mixture of both datasets.

Test set	Model	Meaning preservation		Linguistic correctness	
		MOS	(Std dev.)	MOS	(Std dev.)
CoQAR	Human rewriting	4.5	(0.86)	4.86	(0.45)
	T5(CoQAR)	3.82	(1.42)	4.66	(0.82)
CANARD	Human rewriting	4.60	(0.96)	4.7	(0.89)
	T5(CANARD)	3.92	(1.34)	4.43	(1.08)
	T5(CoQAR+CANARD)	3.96	(1.47)	4.76	(0.77)

Table 4.4: Results of the human evaluation of QR.

Model: We train a QR model based on T5 on three datasets: CoQAR, CANARD, and CoQAR+CANARD. For each dataset, we fine-tune the small 1.1 version of T5². The model is evaluated on the dev set using METEOR.

Two Mean Opinion Score (MOS) evaluations were carried out on 8 human testers who were asked to judge the quality of rewritten questions. We sampled 50 original questions from CoQAR and 50 original questions from CANARD. Each original question was then paired with several rewritings:

- one rewriting from the corpus, to which we refer as the *reference*;
- one or several rewritings generated by different T5 models: each source question from CoQAR is paired with a rewriting generated by T5(CoQAR), while each source question from CANARD is paired with one rewriting generated by T5(CANARD) and one rewriting generated by T5(CoQAR+CANARD).

The pairs were then used in two evaluations.

In the first evaluation, rewritten questions were presented to human testers, together with the original question and its context (preceding turns and the corresponding text passage). Testers assessed the semantic similarity of the rewritten and original questions. In the second evaluation, rewritten questions were presented alone to the testers for them to assess linguistic correctness. Both semantic similarity and linguistic correctness were evaluated on the 5-points scale. In the end, each rewritten question received one rating for semantic similarity and one for linguistic correctness. The results are reported in Table 4.4.

We see that QR models obtain scores that are clearly below human performance in terms of meaning preservation. We also observe that the T5 model that was trained on CoQAR and CANARD obtains higher linguistic correctness scores than the model that was only trained on CANARD, and this result does not seem due to chance (a Mann-Whitney U test gives a p -value of 0.026). It is plausible that, although adding data from CoQAR to the training set does not improve meaning preservation, it improves linguistic correctness because of its greater diversity in term of rewritings’ surface forms.

4.2.3 Evaluation on Conversational Question Answering

Typically, the inputs to a CQA neural model are: a question, its conversational context (i.e. the sequence of previous questions and answers), and the associated passage. The output is the answer, either in the form of a span from the passage or in the form of valid tokens such as “yes”, “no” or “unknown”.

A challenge for conversational question answering was also released with CoQA³. The models are evaluated with the F1 score (Reddy et al., 2019). Transformers have been successfully used in this task: to the time this paper was written, the best model (a RoBERTa-based model (Ju et al., 2019)) got 90.7 of overall F1 measure, overcoming human performance 88.8.

Our goal is to indirectly assess the quality of QR by comparing the performance of a model taking original questions and their context as inputs with a model using out-of-context rewritings instead. In other words, we would like to know whether replacing the original question with its

²https://huggingface.co/google/t5-v1_1-small

³<https://stanfordnlp.github.io/coqa/>

4.2 Question Rewriting

QR mechanism	F1	EM
None (question+context)	68.13	49.63
Human rewriting	63.26	45.10
T5(CoQAR+CANARD)	63.30	44.97

Table 4.5: Results of the CQA evaluation.

conversational context by the out-of-context rewriting has a positive impact on answer extraction. First, we evaluate the impact of rewritten questions in the performance of a RoBERTa baseline (Liu et al., 2019). Second, in order to assess the reusability of QR models trained on CoQAR, we further evaluate a state-of-the-art non-conversational QA model trained on SQuAD (Rajpurkar et al., 2018) by testing it with the rewritten questions.

We would like to assess the impact of QR on state-of-the-art models for CQA by answering the following question: would the models be able to extract the correct answer from the passage without dealing with the conversational context? To this aim we propose three experiments in which we train and evaluate a transformer on several variations of QR: no rewriting, human rewriting, and model rewriting.

Datasets and Variants. We use CoQAR, with distinct rewriting.

- i No rewriting: the original dataset, taking into account the conversational context.
- ii Human rewriting: the dataset containing only the question rewritten by human annotators, ignoring completely the conversational context.
- iii QR model: instead of using human annotations we use questions that were generated automatically by the T5(CoQAR+CANARD) model presented in Section 4.2.2.

Model. For the CQA experiments, we train and evaluate a RoBERTa⁴ transformer on CoQAR with the distinct rewriting mechanisms described above.

Evaluation. Results are presented in Table 4.5. Surprisingly, resolving the context with human question rewriting does not seem to help RoBERTa to better identify the answer in terms of F1 and exact match (EM) as defined in Rajpurkar et al. (2016). We obtained an F1 and EM gain of 4.87 and 4.53 respectively of the original in-context questions over the out-of-context human rewritings.

Unlike (Vakulenko et al., 2021), where results of the same task are reported on CANARD, the setting relying on original questions (referred to as CANARD_O) and the one relying on human-written questions (CANARD_H) respectively obtain 53.65 and 57.12 F1 scores, which correspond to a gain of 3.47 points for human rewriting. We suspect that the self-attention mechanism of RoBERTa solves the coreferences and ellipsis present in short in-context questions limited by the separation token from the context and the passage. While processing a long self-contained rewriting might be more difficult. These results confirm the good performance of RoBERTa on the original task of CQA (Ju et al., 2019).

⁴<https://huggingface.co/>

Interestingly, automatically rewritten questions trained on both CoQAR and CANARD obtained similar performance than human rewritings, although human rewriting, got a slightly better EM. These results are comparable with the ones reported on CANARD in (Vakulenko et al., 2021).

We also conducted experiments on the re-usability of QA systems by solving the context through question rewriting. The results are promising, rewriting out-of-context question will let us reuse existing QA systems. We invite the reader to look for details in the paper (Brabant et al., 2022).

4.3 A conversational QA corpus grounded in Wikidata

After the great success of ChatGPT that spread out nonfactual generative neural models to the big audience, guiding semantically these models to enable explainability and to reduce their typical errors: hallucinations, distortions, omissions and repetitions (Faille et al., 2021; Narayan et al., 2022; Nie et al., 2019) is an urgent need. We propose KGConv⁵, a corpus of Conversational Question Answering (CQA) grounded on Wikidata⁶ to constraint the generation with a Knowledge Graph (KG).

KGConv is composed of conversations between two participants, one that always asks questions and another one that always answers based on facts. Thus, it contains sequences of question-answer pairs. The grounded sequences are composed of Wikidata triples of the form: (s, p, o) , in which s is the subject, p is the property and o corresponds to the object of a fact belonging to a KG.

In total KGConv gathers 71K conversations (604K question-answer pairs in total), where each pair relates to an underlying fact from the public KG Wikidata⁷. Each conversation is focused on a given *root* entity. As illustrated by Table 4.6, the first question is directly about this root entity, while the next ones explore new facts about any entity discovered during the conversation (including the root entity itself). This corpus can be used for distinct tasks such as factual question generation, question rewriting as well as generation of sequence of questions and answers from a given Knowledge-graph or vice-versa.

This is ongoing work. The corpus has been released publicly and a paper presenting the corpus will be submitted for publication soon. This dataset was also used by the PhD student Juliette Faille supervised by Claire Gardent as part of the collaboration with Orange in the ITN European Project NL4XAI. Particularly, it was her subject of study during her secondment at Orange. Her work focused on studying explainability and factual question generation.

4.4 Other Contributions: Graph Embeddings

This work was made by the PhD candidate Sebastien Montella in co-supervision with Dr. Johannes Heinecke at Orange and it was published in (Montella et al., 2021). While most of dialogue systems store their knowledge using simple structures, namely a set of slot-value pairs, world knowledge is usually stored in Knowledge Graphs. A KG is a collection of triples $\langle s, p, o \rangle$; where s, p and o stand for the *subject*, *predicate* and *object* respectively. An *entity* denotes whether

⁵<https://github.com/Orange-OpenSource/KGConv>

⁶<https://www.wikidata.org/>

⁷<https://www.wikidata.org/>

4.4 Other Contributions: Graph Embeddings

#1	Question variants	Triple	<i>(NGC 4833, part of, Milky Way)</i>
		original	NGC 4833 is part of what astronomical object?
		subject	NGC 4833
		rewritten	NGC 4833 is part of what astronomical object?
		original	Where is NGC 4833 located?
		subject	NGC 4833
		rewritten	Where is NGC 4833 located?
	Answer	Milky Way	
#2	Question variants	Triple	<i>(NGC 4833, discoverer or inventor, Nicolas Louis de Lacaille)</i>
		original	Who was behind the discovery of NGC 4833?
		subject	NGC 4833
		rewritten	Who was behind the discovery?
		original	What was the name of the discoverer of NGC 4833?
		subject	NGC 4833
		rewritten	Who discovered this object?
		original	Who found NGC 4833?
		subject	NGC 4833
		rewritten	Who found this object?
	Answer	Nicolas Louis de Lacaille	
#3	Question variants	Triple	<i>(Nicolas Louis de Lacaille, religion or worldview, Catholic Church)</i>
		original	What was his religion?
		subject	his
		rewritten	What was his religion?
		original	What faith did he follow?
		subject	he
		rewritten	What faith did he follow?
	Answer	Catholic Church	

Table 4.6: Excerpt of a question-answer conversation along with the related triples. The root entity is NGC 4833, from the theme “space object”. The rewritten corresponds to the in-context question that has been automatically generated by a T5 model.

a subject or an object and a *relation* denotes a predicate that links two entities. The main contributions of this work are: (i) it explores graph embeddings in the hyperbolic space instead of the Euclidian space; (ii) it considers the time parameter; (iii) it propose a hyperbolic model aware of time and (iv) it compares the performance of state-of-the-art models that takes into consideration time versus models that use negative sampling.

Since KGs are sometimes incomplete, one important task of NLP is Link Prediction (LP), which consists in predicting the missing connections between entities. This task can help for instance to build knowledge graphs on the fly. Each entity and relation are map into a vector space to learn low-dimensional embeddings such that, valid triples maximise a defined scoring function and that fallacious triples minimise it. An approach is efficient if it can model multiple relational patterns. Some predicates are symmetric (*e.g. marriedTo*), asymmetric (*e.g. fatherOf*), an inversion of another relation (*e.g. fatherOf* and *childOf*) or a composition (*e.g. grandfatherOf*). Hierarchical relations have remained challenging to model in Euclidean space, while hyperbolic

	entities	properties	triples	conv.	questions			
					train	dev	test	total
person	31671	327	71915	25918	184939	29352	11386	225677
country	2171	171	3475	703	5085	817	214	6116
ideology	1220	169	1677	450	3112	581	228	3921
space object	2586	116	6360	5961	0	0	50158	50158
molecular entity	17798	151	38314	23033	154511	24587	9531	188629
historical event	4695	189	7770	4972	35270	5684	2247	43201
food	2532	166	4012	2099	15050	2230	1011	18291
taxon	3190	215	5408	1902	0	0	16099	16099
with_unseen_properties	13651	404	24123	5558	0	0	51813	51813
whole dataset	63345	458	142691	70596	397967	63251	142687	603905

Table 4.7: For each theme, the table gives: the number of different entities and properties appearing in conversations, the number of conversations, and the number of questions for each split. Note that in the entities and properties columns, the “total” values are not the sum of the cells above; this is because some entities and properties appear in several themes.

geometry reveals to be a strong asset to capture hierarchical patterns. Nevertheless, the aforementioned approaches represent embeddings as invariant to time. For example, the triple $\langle Donald Trump, presidentOf, U.S. \rangle$ is not longer correct in 2022.

This work shows that an optimised number of negative samples enables the state-of-the-art model ATTH (Chami et al., 2020) to reach competitive or even better performance on temporal link prediction while being unaware of the temporal aspect. It also introduces an extension of ATTH, namely HERCULES⁸.

This was the first attempt to leverage the curvature of a manifold to coerce time-aware representation. An ablation study of distinct curvature definitions has been done to investigate the compelling results of ATTH over time-aware models.

Problem Definition

Lets consider a valid quadruplet $\langle s, p, o, t \rangle \in \mathcal{S} \subset \mathcal{E} \times \mathcal{R} \times \mathcal{E} \times \mathcal{T}$, with \mathcal{E} , \mathcal{R} and \mathcal{T} the sets of entities, relations and timestamps respectively and \mathcal{S} the set of correct facts. A scoring function $f : \mathcal{E} \times \mathcal{R} \times \mathcal{E} \times \mathcal{T} \rightarrow \mathbb{R}$ is defined such that $f(s, p, o, t)$ is maximised for any quadruplet $\in \mathcal{S}$, and minimised for corrupted quadruplet ($\notin \mathcal{S}$). Throughout the optimisation of the foregoing constraint, representations of entities, relations and times are learned accordingly. The resulting embeddings should then capture the multi-relational graph structure. Thus, f is measuring the probability that an entity s is connected to an entity o by the relation p at time t .

Hyperbolic Geometry

Hyperbolic geometry belongs to non-Euclidean geometry. In contrast to Euclidean geometry relying on Euclid’s axioms (Heath and Euclid, 1956), non-Euclidean geometry rejects the fifth

⁸Hyperbolic Representation with Time and Relational Curvatures for Temporal Knowledge Graphs

4.4 Other Contributions: Graph Embeddings

axiom known as the parallel postulate. It states that given a point x and a line l_1 , there exists a unique line l_2 parallel to l_1 passing through x . This is only possible due to a (constant) zero curvature of the space. The curvature defines how much the geometry differs from being flat. The higher the absolute curvature, the curvier. Euclidean space has a zero curvature hence called *flat space*. When represented in an Euclidean space, straight lines become curved, termed as *geodesics* (Fig. 4.3).

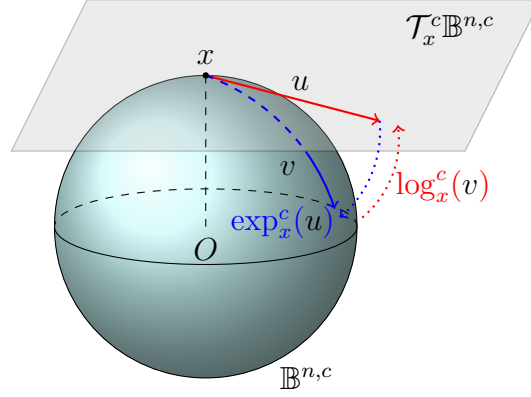


Figure 4.3: Illustration of the exponential and logarithmic maps between the Poincaré ball $\mathbb{B}^{n,c}$ and the tangent space $\mathcal{T}_x^c \mathbb{B}^{n,c}$.

Hyperbolic geometry comes with a constant negative curvature. The interested reader is referred to formal definitions in (Montella et al., 2021).

From ATTH to HERCULES

Given a quadruplet, $\langle s, p, o, t \rangle$, we note e_s^H , r_p^H and e_o^H the hyperbolic embeddings of the subject, predicate and object respectively.⁹ ATTH uses relation-specific embeddings, rotations, reflections and curvatures. The curvature is defined as depending on the corresponding relation p involved. Precisely, a relation p is attributed with an individual parametric curvature c_p . The curvature c_p is defined in Eq. 4.1 as:

$$c_p = \sigma(\mu_p) \quad (4.1)$$

where μ_p is a trainable parameter $\in \mathbb{R}$ and σ is a smooth approximation of the ReLU activation function defined in $[0, +\infty]$. With such approach, the geometry of the manifold is learned, thus modified for a particular predicate. The curvature dictates how the manifold is shaped. Changing the curvature of the manifold implies changing the positions of projected points. This means that for distinct relations, the same entity will have different positions because of the different resulting geometries for each relation. For example, lets consider the triples $t_1 := \langle Barack Obama, visit, France \rangle$ and $t_2 := \langle Barack Obama, cooperate, France \rangle$. The Euclidean representations of entities *Barack Obama* and *France* from both facts will be projected onto the riemannian manifold. However, the structure (i.e. curvature) of the manifold changes as a function of the relation of each fact (i.e. 'visit' and 'cooperate'). Therefore, the resulting hyperbolic embedding of *Barack Obama* of t_1 will not be the same resulting hyperbolic embedding of *Barack Obama* in t_2 . By analogy, the same holds for entity *France*.

⁹Since ATTH is not considering time, the parameter t is not used.

Datasets	$ \mathcal{E} $	$ \mathcal{R} $	$ \mathcal{T} $	Training	Validation	Test
ICEWS14	7,128	230	365	72,128	8,941	8,963
ICEWS05-15	10,488	251	4017	368,962	46,275	46,092

Table 4.8: ICEWS14 and ICEWS05-15 Datasets Statistics

In order to learn rotations and reflections, ATTH uses 2×2 Givens transformations matrices (Chami et al., 2020). Those transformations conserve relative distances in hyperbolic space and can therefore directly be applied to hyperbolic embeddings (isometries). Furthermore, ATTH utilizes an hyperbolic attention mechanism to represent complex relations that can be a mixture of rotation and reflection. The attention scores are computed in the tangent space by projecting the hyperbolic rotation embedding and hyperbolic reflection embedding with the logarithmic map into the euclidian space, as shown in Figure 4.3. Then, the attention vector is mapped back to manifold using the exponential map. We propose HERCULES, a time-aware extension of ATTH. HERCULES redefines the curvature of the manifold as being the product of both relation and time. The main intuition of HERCULES is that both relation and time directly adjust the geometry of the manifold such that the positions of projected entities are relation-and-time-dependent. This is advantageous in that no additional temporal parameters per entity are needed. Since the whole geometry has changed for specific relation and time, all future projections onto that manifold will be aligned to the corresponding relation and timestamp. We investigate different curvature definitions and time translation in our experiments (see the next Section). The scoring function of HERCULES remains same as ATTH.

When learning hyperbolic parameters, the optimisation requires to utilise a Riemannian gradient (Bonnabel, 2013). However, proven to be challenging, we instead learn all embeddings in the Euclidean space. The embeddings can then be mapped to the manifold using the exponential map. This allows the use of standard Euclidean optimisation strategies.

Experiments

Datasets For fair comparisons, we test our model on same benchmark datasets used in previous works, *i.e.* ICEWS14 and ICEWS05-15. Both datasets were constructed by (García-Durán et al., 2018) using the Integrated Crisis Early Warning System (ICEWS) dataset (Boschee et al., 2018). ICEWS provides geopolitical information with their corresponding (event) date, *e.g.* $\langle \text{Barack Obama, visits, France, 2009-03-11} \rangle$. More specifically, ICEWS14 includes events that happened in 2014 whereas ICEWS05-15 encompasses facts that appeared between 2005 and 2015. We give the original datasets statistics in Table 4.8. To increase the number of samples, for each quadruplet $\langle s, p, o, t \rangle$ we add $\langle s, p^{-1}, o, t \rangle$, where p^{-1} is the inverse relation of p . This is a standard data augmentation technique usually used in LP (Balažević et al., 2019; Goel et al., 2020; Han et al., 2020).

Evaluation Protocol & Metrics Given a (golden) test triple $\langle s, p, o, t \rangle$, for each entity $s' \in \mathcal{E}$, we interchange the subject s with s' and apply the scoring function f on the resulting query $\langle s', p, o, t \rangle$. Since replacing s by all possible entity s' may end up with a correct facts, we filter out those valid quadruplets and give them extremely low scores to avoid correct quadruplets to be scored higher than the tested quadruplet in final ranking (Bordes et al., 2013). We then rank the entities

4.4 Other Contributions: Graph Embeddings

based on their scores in descending order. We store the rank of the correct entity s noted z_s . Thus, the model should maximize the returned score for the entity s such that $z_s = 1$. The same process is done using the object o .

To evaluate our models, we make use of the Mean Reciprocal Rank (MRR). We also provide the Hits@1 (H@1), Hits@3 (H@3) and Hits@10 (H@10) which assess on the frequency that the valid entity is in the top-1, top-3 and top-10 position, respectively.

Results We provide link prediction results on ICEWS14 and ICEWS05-15 for ATTH, HERCULES and different models from the literature. As (Han et al., 2020), we adopted a dimension analysis to investigate behaviors and robustness of approaches. When possible, we re-run official implementation of models. Otherwise, official or best results in literature are reported. Results are shown in Table 4.9.

As expected, hyperbolic-based strategies (*i.e.* DYERNIE, ATTH and HERCULES) perform much better at lower dimensions, outperforming most of other approaches with ten times less dimensions. We report an average absolute gain of 11.6% points in MRR with only 10 dimensions over the median performance of other approaches with 100 dimensions. This strengthens the effectiveness of hyperbolic geometry to induce high-quality embeddings with few parameters.

Astonishingly, we notice that ATTH model is highly competitive despite the absence of time parameter. ATTH exhibits new state-of-the-art or statistically equivalent performances compared to DYERNIE and HERCULES. We remark no statistically significant differences in performances between hyperbolic models.¹⁰ Importantly, unlike other research carried out in this area, time information here does not lead to any notable gain. This seems to indicate that other parameters should be considered. We examine this phenomenon in section 4.4.

On ICEWS14, for $dim \in \{20, 40, 100\}$, both ATTH and HERCULES outperform DYERNIE by a large margin. We witness an improvement of 2.5% and 5% points in MRR and Hits@1 with 100-dimensional embeddings. On ICEWS05-15, ATTH and HERCULES yield comparable achievements with the state-of-the-art. In contrast to DYERNIE, it is noteworthy that ATTH and HERCULES utilize a single manifold while reaching top performances.

We also distinguish tempered results on Hits@10 metric for ATTH and HERCULES models. This suggests that during optimization, ATTH and HERCULES favor ranking some entities on top while harming the representation of others.

Time Awareness vs Negative Sampling.

First, besides time translation, we probe different curvature definitions to identify fluctuation in performances. We analyse how time information alters the LP results by adding time as part of the curvature (*i.e.* HERCULES) and as a translation. We also explore if incorporating the Euclidean dot product of the subject and object embeddings (noted $\langle e_s^E, e_o^E \rangle$) into the curvature helps to learn a better geometry. An ablation study is given in Table 4.10.

Albeit counter-intuitive, we observe that our results corroborate with our initial finding: time information is not the culprit of our high performances. More strikingly, a simple relational curvature (*i.e.* ATTH) is sufficient to perform best on ICEWS14 ($dim = 40$). Neither the inclusion

¹⁰We performed the Mixed-Factorial Analysis of Variance (ANOVA), in which the independent variables are *the dimension* and *the model* and the dependent variable is *the metric*. We consider two groups one for each dataset. We report p -values of 0.842, 0.872, 0.926 and 0.229 for MRR, H@1, H@3 and H@10 respectively.

Datasets		ICEWS14 (filtered)				ICEWS05-15 (filtered)			
<i>dim</i>	Model	MRR	H@1	H@3	H@10	MRR	H@1	H@3	H@10
10	ATISE [†]	18.0	3.03	23.9	48.7	15.9	4.35	19.22	41.0
	TERO [†]	7.25	2.39	6.40	16.6	10.3	3.54	10.1	23.2
	DYERNIE*	46.2	36.0	51.1	<u>66.3</u>	58.9	50.5	63.2	75.1
	HERCULES	<u>46.0</u>	<u>34.9</u>	52.4	66.0	<u>54.7</u>	<u>43.8</u>	<u>61.8</u>	73.2
	ATTH	45.6	34.2	<u>52.0</u>	66.4	49.9	34.4	61.6	<u>73.6</u>
20	ATISE [†]	19.1	1.28	28.2	54.7	24.5	7.67	32.3	59.2
	TERO [†]	24.5	13.8	28.01	46.3	27.1	13.5	33.3	54.1
	DYERNIE*	53.9	44.2	58.9	72.7	64.2	56.5	68.2	79.0
	HERCULES	55.5	47.2	<u>59.4</u>	<u>71.4</u>	63.2	55.2	<u>67.7</u>	<u>77.6</u>
	ATTH	<u>55.2</u>	<u>46.7</u>	59.7	<u>71.4</u>	<u>63.5</u>	<u>55.8</u>	<u>67.7</u>	77.5
40	ATISE [†]	38.4	23.3	47.6	67.3	35.7	19.2	44.3	69.1
	TERO [†]	35.1	22.7	40.5	60.8	28.3	12.7	35.3	60.5
	DYERNIE*	58.8	49.8	63.8	76.1	68.9	61.8	72.8	82.5
	HERCULES	<u>61.2</u>	<u>54.3</u>	<u>64.7</u>	74.1	<u>68.5</u>	62.1	<u>72.0</u>	<u>80.9</u>
	ATTH	61.7	54.5	65.4	<u>75.4</u>	<u>68.5</u>	<u>62.0</u>	71.9	80.6
100	TRANSE*	30.0	14.8	42.7	60.1	30.4	13.3	42.4	61.1
	DISTMULT*	57.5	46.9	64.2	77.9	47.1	33.6	55.1	72.5
	COMPLEX*	49.3	36.6	56.2	74.2	39.0	22.9	49.2	68.4
	TTRANSE*	34.4	25.7	38.3	51.3	35.6	15.4	51.1	67.6
	TCOMPLEX*	31.8	12.9	45.7	63.0	45.1	36.3	49.2	62.0
	HYTE*	33.1	6.8	54.5	73.6	38.1	7.6	65.0	80.4
	ATISE [†]	52.2	41.0	60.0	72.7	47.0	32.4	55.5	76.4
	TERO [†]	45.4	34.0	52.2	67.0	41.1	26.3	48.9	71.7
	DYERNIE*	66.9	<u>59.9</u>	<u>71.4</u>	79.7	73.9	<u>67.9</u>	77.3	85.5
	HERCULES	<u>69.4</u>	65.0	<u>71.4</u>	77.9	73.5	68.6	<u>76.1</u>	<u>82.9</u>
ATTH	69.5	65.0	71.5	<u>78.2</u>	<u>73.6</u>	68.6	76.0	<u>82.9</u>	

Table 4.9: Link prediction results on ICEWS14 and ICEWS05-15 datasets: (†) results are obtained using the official implementation of (Xu et al., 2020), (*) results are taken from (Han et al., 2020). For each dimension (*i.e.* *dim*), best results are in bold and second-to-best underlined. No statistically significant differences in performance are observed between DYERNIE, HERCULES and ATTH.

Relation	Time Curvature	Time Translation	$\langle e_s^E, e_o^E \rangle$ Curvature	MRR	H@1	H@3	H@10
✓	✗	✗	✗	61.7	54.5	65.4	75.4
✓	✓	✗	✗	61.2	54.3	64.7	74.1
✓	✓	✓	✗	60.1	52.1	64.5	75.0
✓	✓	✓	✓	49.5	38.9	55.4	69.2

Table 4.10: Ablation study: Link prediction results on ICEWS14 using ATTH (*dim* = 40) with different curvature definitions and time translation applied.

4.4 Other Contributions: Graph Embeddings

of a time translation, similarly to TTRANSE, nor the Euclidean dot product provide interesting outcomes.

We then probe the sensitivity of HERCULES towards temporal feature by performing LP with incorrect timestamps. Our intuition is to inspect whether feeding invalid timestamps during evaluation exhibits significant variation or not compared to the reference performances, *i.e.* LP results with initial (non-corrupted) testing samples. To do so, for each testing quadruplet, we replace the (correct) time parameter with each possible timestamp from \mathcal{T} . We therefore collect multiple LP performances of HERCULES corresponding to each distinct timestamp. Our finding is that despite erroneous timestamps, LP results show insignificant discrepancies with the initial HERCULES performance (dashed red line). This indicates that HERCULES gives little importance to the time parameter and thus only relies on the entity and the predicate to perform knowledge graph completion. This further highlights our finding that timestamp is not responsible for our attracting performances.

We therefore assume that the optimisation procedure may be involved. We consequently question the effect of negative sampling. Precisely, we train HERCULES with $dim = 40$ by tuning the number of negative samples between 50 to 500. For both, ICEWS14 and ICEWS05-15, negative sampling shows considerable gain as the number of samples increases. We record an absolute gain of 5% points in MRR from 50 to 500 samples. We can see a rapid growth in MRR when the number of samples is inferior to 200. Adding 50 samples is equivalent to about 2% points gain in MRR. Then, performances reach a plateau around 300 negative samples. We conjecture that a diversity in negative samples is enough to learn good representations. Notwithstanding that a large number of negative samples heavily constraints the location of entities in space, the resulting embeddings might benefit from it to be better positioned relatively to others.

We conclude that despite the present time parameter, an optimal negative sampling enables to reach new state-of-the-art outcome. Therefore, we argue that time is not the only parameter that should be considered when performing LP. We highlight that one should be raising awareness when training TKG representations to identify if time is truly helping to boost performances.

Chapter 5

Data Collection, Annotation and Frameworks

This chapter summarises the corpora, annotations, and frameworks in which I worked on. It also describes briefly the national and international research projects in which I have been involved.

5.1 Corpora

This section describes briefly the work I have done in data collection and annotation.

5.1.1 The French Portmedia Corpus

The French MEDIA corpus collects about 70 hours of spontaneous speech (1258 dialogues, 46k utterances, 494.048 words and 4068 distinct words) for the task of hotel reservation and tourist information (Bonneau-Maynard et al., 2009). Calls from 250 speakers to a simulated reservation system (i.e. the Wizard-of-Oz) were recorded and transcribed. Dialogues are full of disfluencies, hesitations, false starts, truncations or fillers words (e.g., euh or ben). I worked on the semantic annotations of this corpus as I was involved in the French ANR project PORTMEDIA (Rojas-Barahona et al., 2011; Rojas-Barahona and Quignard, 2011). 330 utterances were manually annotated with semantic relations (i.e. High-Level Semantics). This gold corpus gathers 653 head segments and 1555 argument segments, from which around 20 are both arguments and heads, such as *une chambre* in Figure 4. This work focuses on annotating the semantic structure, by segmentating users' utterances into concepts.

This ontology identifies the concepts that can have arguments, and we thus use this information to further distinguish between *head* segments that can have arguments.

Besides the gold annotation, a silver annotation of the whole MEDIA dataset was also provided. It was generated automatically after a pipeline of syntactic analysis, semantic role labelling and extraction of semantic frames. I implemented the whole pipeline as well as the annotation tool used by the annotators. I also collaborate to the creation of a Bayesian model for inferring these annotations (Lorenzo et al., 2013).

5.1.2 The French Emospeech Corpus

The French Emospeech corpus was already introduced in Section 3.1.1. As described in (Rojas-Barahona et al., 2012b), to collect Human-Game dialog data, we developed a Wizard-of-OZ (WOZ) interface using the MITRE Dialog Toolkit Midiki (Burke et al., 2003). Midiki, is a portable toolkit for building dialogue managers in Java. It implements the information-state model of dialogue (Traum and Larsson, 2003) where in essence, the information state models the progression of dialog while update rules formalise the way that information state is changed as the dialog progresses.

We first extended Midiki to support a multi-agent architecture and the configuration from a relational database. We then used this extended Midiki (i) to develop a rule-based dialog system for the MP game and (ii) to implement two Wizard-of-OZ interfaces for data collection: the free- and the semi-automatic WOZ interface.

The *free WOZ interface* aims to simulate mixed-initiative dialogs by allowing the wizard to chat with the player as she moves around the game while simultaneously storing all interactions in a database. A virtual dialog manager ensures that the wizard respects the game logic, starting the appropriate subdialogs at the appropriate place in the virtual world. In this setup, the interactions between the wizard and the player simulate a direct Human-Human dialog in the context of the MP game.

In contrast, *the semi-automatic wizard* favours system-driven dialogs by connecting the Wizard not only with the player and the game but also with the rule-based dialog manager (Figure ??). This dialog manager supports the Wizard by automatically interpreting the player's input and selecting a possible response. As the Wizard interacts with a player, she can then either accept the response suggested by the rule-based dialog manager (if this response is appropriate) or enter a different response (whenever the response suggested is incorrect or inappropriate).

Figure 5.3 shows the architecture of the WOZ interface: the dialog manager (either Midiki or a virtual DM), the MP game, the Wizard of Oz interface and an automatic speech recognition module (ASR)¹ communicate together within the Open Agent Architecture (OAA) (Cheyer and Martin, 2001). The WOZ interface is implemented as a web service and all interactions are logged into a relational database.

To support data collection for different game scenarios, we also developed a Dialogue Con-

¹Although the Wizard Framework supported both speech and written input, we did not record speech in our first experiments. All data is therefore written data.

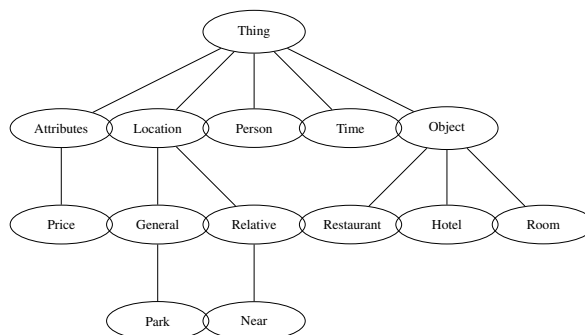


Figure 5.1: Excerpt of MEDIA ontology

	Subjects	Dialogs	Utctes	Tokens	Player U.	Player Tokens	Player Token Types
Semi-Aut.	40	591	4874	77854	1321	12901	1427
Free	50	658	5580	90655	2288	18712	1542
Total	90	1249	10454	168509	3609	31613	2969

Figure 5.4: Data collected

between 1 and 1.5 hour.

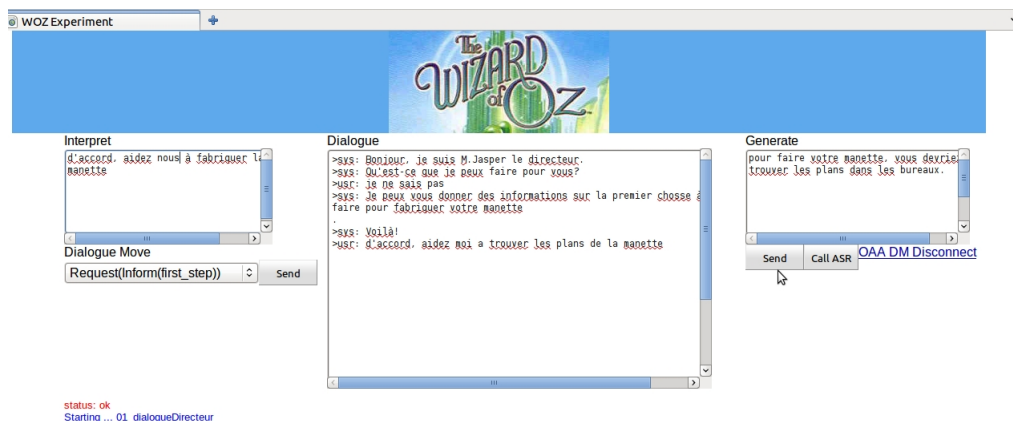


Figure 5.5: Wizard of OZ Graphical User Interface (GUI).It is split in three parts: the interpretation, the dialog and the generation. The wizard receives the input sentence at the left-side, she can see the whole dialog in the centre, and she edits the generated utterance at the right side. She can introduce also the dialog move associated to the input and output sentences.

After configuration of the WOZ tool using the Dialog Configuration Tool mentioned in the preceding section, the free-WOZ was also used by Master students from Rennes University for collecting dialogs in a game simulating the visit of an exhibition on Alice in Worderland. In this way, they collected 25 dialogs using Lewis Carroll’s subrealist.

5.1.3 Annotating Posts Following Cognitive Behavioural Therapy Principles

This work was published in (Rojas-Barahona et al., 2018). The main goal of annotations based on Cognitive Behavioural Therapy (CBT) was to develop the understanding component of a health assistant for preventive intervention of mental health. The corpus consists of 500K written posts that users anonymously posted on the Koko platform². This platform was based on the peer-to-peer therapy proposed by (Morris et al., 2015). In this set-up, a user anonymously posts their problem (referred to this as the *problem*) and is prompted to consider their most negative take on the problem (referred to this as the *negative take*). Subsequently, peers post responses that attempt

²<https://itskoko.com/>

5.1 Corpora

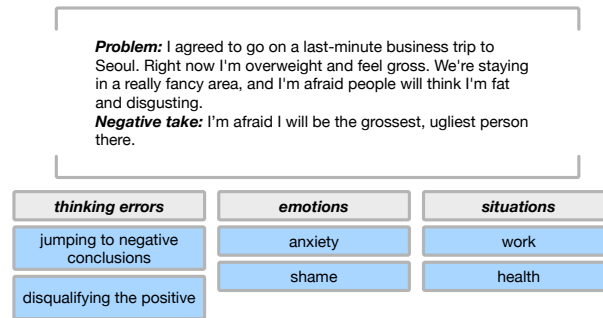


Figure 5.6: An example of an annotated Koko post.

to offer a re-think and give a more positive angle on the problem. Initially, any first-time Koko user would be given a short introductory tutorial in the art of 're-thinking'/'re-framing' problems (based on CBT principles), before being able to use the platform; this however changed over time, as the age group of the users decreased, and a different introduction, emphasizing empathy and optimism, was used in relation to suggesting helpful responses (less CBT-based than the 're-thinking'). Some of the data annotated in this study was drawn from this later phase. When first developed, this framework was shown to be more efficacious than expressive writing, an intervention that has been shown to improve physical and emotional well-being (Morris et al., 2015). Since then, the company has developed an app that has collected a very large number of posts and associated responses. In this work we only focus on analysing the posts. Figure 5.6 gives an example of an annotated post.

We draw from principles of Cognitive Behavioural Therapy (CBT) to define the ontology. CBT is derived originally from Beck's Cognitive Therapy model theory (Beck, 1976; Beck et al., 1979) which says that our emotions and behaviour are influenced by the way we think and by how we make sense of the world. Thus, if the patient changes the way he or she thinks about their problem that will in turn change the way he or she feels and behaves. A major underlying principle of CBT is the idea of cognitive distortion, and the value in challenging this. In CBT, patients are helped to test their assumptions and views of the world in order to check if they fit with reality. When patients learn that their perceptions and interpretations are distorted or unhelpful, they then work at correcting them. Within the realm of cognitive distortion, CBT identifies a number of specific self-defeating thought processes, or thinking errors. There is a core of around 10 to 15 thinking errors, with their exact titles having some fluidity. A strong component of CBT is teaching the client to be able to recognise and identify the thinking errors themselves, and ultimately discard the negative thought process, and 're-think' their problem.

We consider the first step that a machine should be able to perform is to adequately decode these 'thinking error' concepts, along with identifying the key emotion(s) expressed, and situational context, within a particular presented problem. Therefore, our ontology consists of *thinking errors*, *emotions*, and *situations*.

5.1.4 Conversational Question Answering with Rewriting

I lead the extension of the corpus CoQA with question rewriting as explained in Section 4.2.

5.1.5 Conversational Question Answering in French

I proposed to extend the French corpus Calor with sequence of questions and answers. The annotations were done by the University of Aix-Marseille. This is a joint work with Geraldine Damnati and Frederic Bechet. This work has been published in (Béchet et al., 2022).

Calor-Dial is an enriched version of the Calor corpus (Marzinotto et al., 2018), collected from French encyclopedic data in order to study Information Extraction on domain specific data. The corpus was initially annotated in semantic Frames (Calor-Frame (Béchet et al., 2017)) and enriched with a first set of questions for Machine Reading Question Answering (Calor-Quest (Béchet et al., 2019)). Calor-Dial addresses the scope of conversational Question Answering. The main originality is that different types of questions are annotated, including more challenging configurations than in classical QA corpora.

5.1.6 Conversational QA grounded in Wikidata

I contribute to the creation of the corpus KGConv and I lead the work from the Orange side, as part of the research project DIANA. This work is in collaboration with Claire Gardent (CNRS) as part of the partnership of the ITN-European project NL4XAI. This corpus is presented in Section 4.3.

5.2 Dialogue Frameworks

During the course of my research career, I was directly involved in the development of dialogue frameworks. Starting from AdaRTE (Rojas-Barahona and Giorgino, 2009) during my PhD studies to the most widespread one, PyDial (Ultes et al., 2017a). In this Section I will briefly present PyDial.

5.2.1 PyDial

PyDial is an open-source end-to-end statistical spoken dialogue toolkit developed by the University of Cambridge (Ultes et al., 2017b). It provides implementations of statistical approaches for all dialogue modules: NLU, DST, Policy, NLG. The term statistical means that: (i) the framework preserves the confidence probability of each module, thus each module outputs are the N-Best list of hypotheses; (ii) Deep learning models can be easily integrated in it; (iii) The framework implements POMDP dialogue systems. Thus, reinforcement learning dialogue management is fully supported, with the user simulator and reward estimators. Moreover, it has been extended to support multiple domains. It offers easily extensible to other domains and or specialised module implementations. It also offers domain-independent implementations of the dialogue modules (see Figure 5.7). The toolkit is available for download under the Apache 2.0 license.

5.2.2 Dialport

Dialport (Lee et al., 2017) is an academic dialogue portal to collect large amounts of real user data for spoken dialog systems (SDS). Sophisticated statistical representations in state-of-the-art SDS, require large amounts of data, which is difficult to obtain by academic teams. With one central portal, connected to many different systems, the task of advertising and affording user

5.2 Dialogue Frameworks

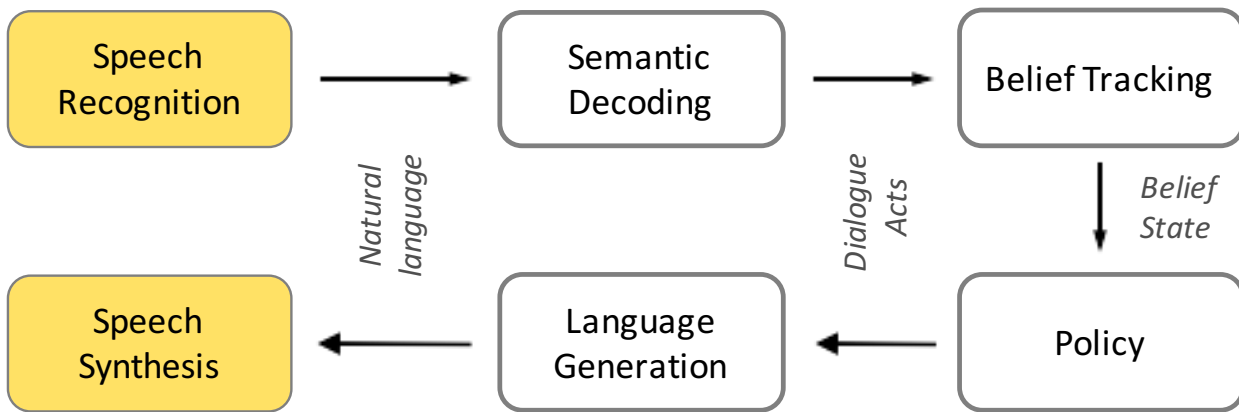


Figure 5.7: Architecture of a modular Spoken Dialogue System.

access can be done in one centralised place that all systems can connect to. DialPort provides a steady stream of data, allowing system creators to focus on developing their systems. The portal decides what service the user wants and connects them to the appropriate system which carries on a dialog with the user, returning control to the portal at the end. Dialport was connected to the Cambridge restaurant information system, which helps users find a restaurant in Cambridge, UK based on the area, the price range or the food type. The connection is done through an **API!** (API!) that connects Dialport to PyDial. An assesment of DialPort is presented in (Lee et al., 2019), to summarise, 28.8% of chatbot utterances were non-understanding recovery turns, such as “can you please rephrase that?”. 62.85% of the times DialPort successfully recommended users and 78.40% it correctly directed users to the appropriate system.

5.2.3 Conversational Search for General Knowledge

Converstional Search for General Knowledge (CS4GK) is a spoken conversational question answering proof of concept that is able to answer questions about general knowledge from Wikidata³ (Rojas-Barahona et al., 2019). The dialogue component does not only orchestrate various components but also solve coreferences and ellipsis.

The architecture of the proposed system consists of a speech-processing front-end, an understanding component, a context manager, a generation component, and a synthesis component. The context manager provides contextualised mediation between the dialogue components and several question answering back-ends, which rely on data provided by Wikidata. Interaction with a human user is achieved through a graphical user interface (GUI). Figure 5.8 depicts the components together with their interactions.

³<https://www.wikidata.org>

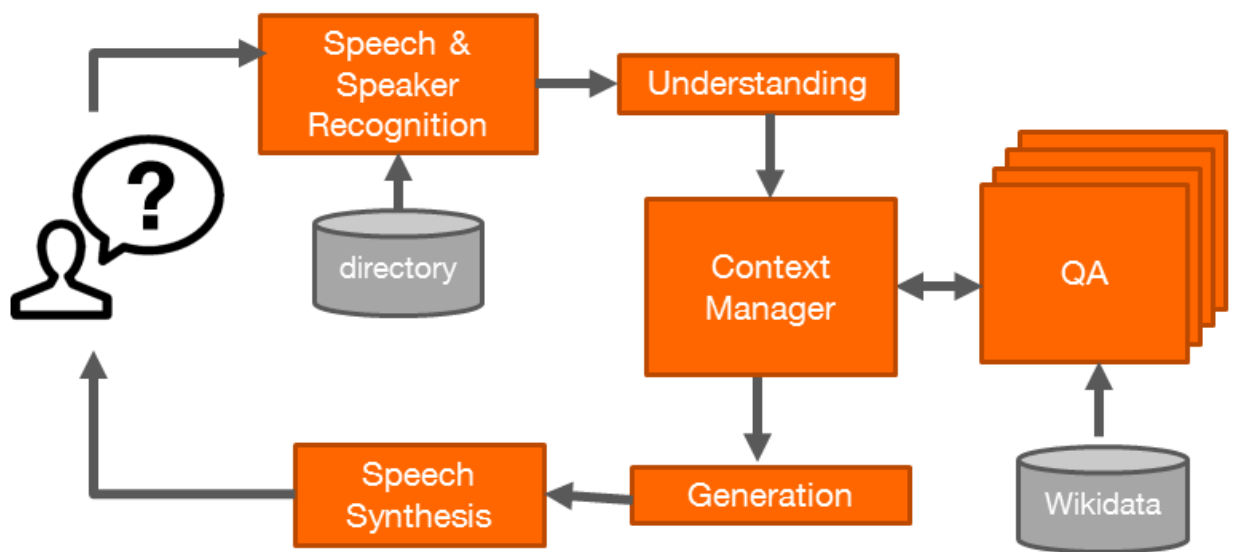


Figure 5.8: High-level depiction of the proposed spoken conversation question answering system. Arrows indicate data flow and direction.

Chapter 6

Scientific Project

After the breakthrough of ChatGPT (Ouyang et al., 2022), Large Language Models have been widely used for distinct daily activities such as summarisation, translation, sentiment analysis, question answering, redaction, code-generation, etc. Nevertheless, it is not clear how these models can be used to solve complex decision-making tasks, such as task-oriented dialogue. Although promising approaches have recently emerged (Wei et al., 2022; Yao et al., 2023, 2022), utilizing LLMs to solve complex tasks is not straightforward. In the case of task-oriented dialogue, one important point concerns the lack of control (Section 6.1.4). Although the conversation is fluid and pleasant is it following the necessary steps to solve the task? Are these steps validated by experts? Are these steps correctly grounded in the World knowledge? These questions bring us to an important issue the evaluation (Section 6.1.1). Can we estimate whether the task was accomplished, and whether the sequence of selected actions was indeed optimal? Retrieval Augmented Generation (RAG) (Mao et al., 2021; Asai et al., 2022), can be used to retrieve useful information that can be injected to the LLM as a prompt. This is a way of grounding the LLM in "factual" information. However, is there a way to be sure the retrieved information is indeed factual?(Section 6.1.3) Moreover, the effort of prompt-engineering can not be neglected. Are we coming back to hand-crafted solutions by hand-crafting prompts? Would it be difficult to maintain and to keep these prompts up-to-date?

This scientific project proposes first to study recent LLM-based reasoning approaches for task-oriented dialogue, providing a rigorous evaluation in terms of task-completion and success rate. We can also, inspired by the seminal evaluation framework Paradise (Walker et al., 1997a), think in ways to find the correlation between task-completion and user satisfaction. For instance, unlike the study presented in Paradise, in which at that time users were not enjoying long conversations with repetitive systems, maybe users now really enjoy talking to LLMs. However, a strong indicator of poor performance might be to confirm that users usually need to call again because despite having a fluid and natural interaction, their problems were not solved at all.

Multimodality is another interesting research topic we will discuss in this proposal. Concerning dialogue, emergent approaches are Visual Question Answering (VQA) and speech-aware emotion detection. The first one can be used in TOD wherein heterogeneous sources of knowledge are identified. The second, can be used to personalise dialogue according to users' mood detected by speech analysis.

6.1 LLMs for Task-Oriented Dialogue

Recently synergising reasoning and acting in large language models (ReAct) (Yao et al., 2022) has shown promising results employing few-shot prompting in a LLM with a sequence of thoughts, actions, and observations. LLMs are indeed capable of performing complex tasks. Thoughts refer to the internal reasoning that decompose a problem into sub-problems. For example, if the model is asked the age of Barack Obama’s wife power of 3, the thoughts might be as follows. First, I need to find out who Barack Obama’s wife is (by calling a question-answering API such as Google Search). Second, I need to find the age of Barack Obama’s wife (by calling a question-answering API). Finally, I will calculate the age power of 3 (by calling a calculator API). Examples have shown promising results for tasks such as multi-hop question-answering, WebShop, and a textual version of a butler in a virtual environment, ALFWorld: the butler can report on actions he has taken to solve a given task in a kitchen (Yao et al., 2022).

We need to study the recent state-of-the-art in task-decomposition utilizing LLMs(Wei et al., 2022; Yao et al., 2023, 2022). We can also explore Algorithm Distillation (Laskin et al., 2022), wherein the logs of cross-episodic events generated during the learning process of a RL algorithm, are used to feed in a LLM, which as consequence can learn the optimal strategy. Since the state-of-the-art is moving impressively fast, this proposal is open to upcoming approaches.

6.1.1 Benchmarcks and Evaluation

We will propose an evaluation framework that takes into account long-term memory, beyond one dialogue session (Xu et al., 2022). For evaluating LLMs an evaluation framework that is both model-agnostic and domain-independent is suitable. Therefore, we should think in ways to detect the conversation goal together with performance indices, with or without humans in the loop. In an initial state we can compare our previous results presented in Section 3.2.4 on hierarchical reinforcement learning for dialogue (Cordier et al., 2022, 2023) with LLMs that follows ReAct and Algorithm Distillation by using the metrics introduced in ConvLab (Zhu et al., 2020). Therefore, we can start this study by using the dataset MultiWoz, in which the user goal is formally defined and provided. This corpus also includes the instructions given to annotators during data collection. Later, we can move to complex tasks in which the user goal is not provided. We can also study more realistic cases such as commercial or technical-support systems. We will perform human evaluation to perceive their satisfaction with the system. As in Paradise (Walker et al., 1997a) and in (Rojas-Barahona and Gardent, 2012), we would like to make a correlation study between objective metrics (indices of performance that can be computed automatically) and subjective metrics (e.g. user satisfaction). This framework can be extended to even more complex interactions such as in multi-modal dialogue systems (Section 6.2).

6.1.2 Interpreting LLMs

After evaluating LLMs performance for task-oriented dialogue, we can use model agnostic black-box interpretability methods (Cafagna et al., 2023) to understand how these models are able to solve complex tasks. Where in the LLM architecture this behaviour is being produced? It is worth noting that these methods can be applied only to Open-Source LLMs, in which we can have access to the model (e.g., Llama (Touvron et al., 2023), Falcon (Penedo et al., 2023)). Interpretability

6.1 LLMs for Task-Oriented Dialogue

might give us an insight of how we can correct LLMs to avoid undesirable behaviour (e.g. forgetting an important instruction, biases). Methods to correct LLMs are introduced in the following Sections.

6.1.3 Retrieval Augmented Generation

One way to find grounded knowledge to feed into LLMs is through information retrieval (Mao et al., 2021; Asai et al., 2022). Despite these techniques rank documents according to their relevance, an important aspect is factuality (Thorne et al., 2018). Another aspect is that the sources of knowledge are heterogeneous: they concern not only documents, but also images, recorded interactions, knowledge graphs, tabular data, among others. Once we have retrieved crucial information there are several ways to inject this information into LLMs. The most adopted one is prompt engineering, but controlling decoding might be more interesting in terms of scientific research.

6.1.4 Controlling Decoding

There are distinct sampling mechanisms used for decoding such as greedy ($\mathcal{T} = 0$ in Equation 6.2), beam search, top-k, nucleus or penalised sampling (Fan et al., 2018; Holtzman et al., 2019; Keskar et al., 2019). An emergent method to control decoding is by adjusting sampling weights in beam search (Ghazvininejad et al., 2017).

Let p_θ be a pre-trained generative language model which has learned the distribution over token sequences by optimising:

$$\mathcal{L} = - \sum_t \log p_\theta(x_t | x_{<t}) \quad (6.1)$$

The next token can be sampled by applying Softmax with temperature \mathcal{T} because the final decoder layer predicts logits o , over the vocabulary space:

$$p_i \propto \frac{\exp(o_i/\mathcal{T})}{\sum_j \exp(o_j/\mathcal{T})} \quad (6.2)$$

Beam search is a breadth-first search algorithm which explores the β tokens which best scores at each level. Then, the likelihood of sampling for the next token x_{t+1} at step t can be augmented by a scoring function:

$$score(x_{t+1}, b_t) = score(b_t) + \log p(x_{t+1}) + \sum_i \alpha_i f_i(x_{t+1}) \quad (6.3)$$

Where the log-likelihood predicted by the pre-trained language model is defined by $\log p(x_{t+1})$. The accumulated score of the already-generated words in the current beam state b_t is $score(b_t)$. A set of feature functions that define the preferences is $f(\cdot)$, which can be a binary classifier that predicts whether a sample is from the true data distribution (Grover et al., 2019). Finally, α_i are the associated weights that work like "control knobs". Therefore, the classifier can be used to constrain factuality by predicting whether the token is grounded in knowledge or not. This is an interesting path that can be further explored.

Another interesting approach is Chain of hindsight (Liu et al., 2023), in which pairs of (answer, feedback) are provided as input to the model during fine-tuning. Thus, the model learns to condition the generation on feedbacks during training.

6.2 Multimodal task-oriented dialogues

Generative models were also trained for process multiple modalities (Antol et al., 2015; Zhu et al., 2016; Srivastava et al., 2021), supporting image captioning and VQA. After the release of multi-modal LLMs, the borders between domains, tasks and modalities have been vanished. One can think in VQA applications that interact with elders, visual impaired (Chen et al., 2022) or more broadly patients or practitioners in the medical domain. Efforts for ethical Artificial Intelligence (AI) can not be diminished before putting VQA to interact with sensitive population. This is particularly true in Europe, after the AI regulation (Hacker et al., 2023). Therefore, important research areas concern interpretability and evaluation of multi-modal LLMs to correct model biases that might produce any harm. Generally, the datasets used for training these large models were built without including minorities or people at risk. Therefore, these models can not respond to their needs. A rigorous study to evaluate LLMs, involving these individuals must be carried out. Interpretability methods (Section 6.1.2) can help us to understand biases deep inside the model and to provide insights to correct them. This can be done in public models for image captioning or VQA such as OFA (Wang et al., 2022).

Multimodality might also be highly related to RAG in both text-to-image and image-to-text models (Yasunaga et al., 2023), wherein the information is available as images, audio or video (Section 6.1.3). Finally, multi-modality also involves speech, in this proposal I talk about ways to adapt task-oriented dialogues that use either RL or LLMs for policy learning and that take into account the speech signal for emotion detection .

6.2.1 Exploiting weak speech signals in the reward function

Treating speech and text together is a research topic widely studied in Spoken Language Understanding (SLU). An interesting research path would be to go beyond intention recognition up to response generation. We worked with the PhD candidates Leo Jacqmin, Lucas Druart and other researchers in a cascade approach that integrated the ASR (i.e. Whisper)(Radford et al., 2022) with a generative model for DST and we participate to the challenge DSTC-11 (Jacqmin et al., 2023)¹. RL can be used for learning the strategy (Section 2.3.2). However, it assumes there is a reward function. In dialogue systems the reward signal is scarce, because it is unbearable for a user to send a satisfaction signal at each dialogue turn. For fluidity, the reward function is asked at the end of the conversation, which produces a scarce signal, complicating the task of policy learning. Thus, the policy spectrum will have a large variance from very poor to very good.

We explored already two ways of solving this problem: (i) predicting the user satisfaction at each dialogue turn by using the interaction quality as reward signal (Section 3.2.2) and (ii) using imitation learning to guide the policy learning (Section 3.2).

One can think in improving the reward scarcity by exploiting the information in the speech signal. Studying weak signals in the speech, such as the emotion, could be a way to improve the reward signal. Emotion detection is important for developing conversational systems that could adapt better to the users' needs, improving as consequence the user satisfaction. For instance, an early detection of distress would entail changing the dialogue strategy to quickly solve the misunderstanding. Emotion detection involves an active research community producing datasets

¹https://storage.googleapis.com/gresearch/dstc11/dstc11_20221102a.html

6.2 Multimodal task-oriented dialogues

build from actors playing the emotion (Burkhardt et al., 2005; Schröder et al., 2007), from wizard of oz or more recently from more natural content (Zadeh et al., 2018, 2020; Dhall et al., 2012; Scheidwasser-Clow et al., 2022). Neural methods for emotion detection have been proposed in (Trigeorgis et al., 2016), since there are not many dataset for training deep models from scratch transfer learning or distillation from other modalities have been also proposed (Pepino et al., 2021; Albanie et al., 2018).

The main research questions behind this study will be:

- Concerning the emotion detection: we need to study domain adaptation for emotion recognition to better exploit the available cross-domain datasets to dialogue.
- Combine multi-modality to produce the reward, one can think in using speech signal and other domain agnostic metrics, such as counting the number of repetitions, counting the number of dialogue turns, measuring the misunderstandings, detecting the sentiment from text.
- Evaluate whether the emotion-based reward is suitable for policy learning
- Compare RL with LLM-based policies (e.g. algorithm distillation (Laskin et al., 2022), Chain of Hindsight (Liu et al., 2023)).

Chapter 7

Conclusion

I presented in this dissertation a selected number of contributions I made to the areas of task-oriented dialogue systems, conversational question answering and graph embeddings.

The contributions to task-oriented dialogue were in the fields of NLU, SLU and DM. Particularly, I explored classical machine learning techniques for NLU, convolutional and recurrent neural networks to deal with noisy inputs for SLU, as well as data-augmentation techniques. Although not mentioned in this work, I am currently supervising with Benoit Favre from the University of Aix-Marseille a PhD thesis on DST, recently we competed in the challenge DSTC-11 and we were awarded the first and second place (Jacqmin et al., 2023)¹. Regarding the DM, I explored Inverse Reinforcement Learning, Deep Reinforcement Learning, Imitation Learning and Structured Policy Learning.

The contributions to conversational QA regard the annotation of existing datasets with information about ellipsis and coreferences, and with question rewriting to transform in-context questions into out-of-context questions. The generative models released with these annotations for the tasks of answer extraction, question generation and question rewriting; as well as the models implemented for predicting ellipsis and coreferences were also presented. I also briefly introduced the corpus KGConv grounded in Wikidata. Moreover, I presented our work on graph embeddings in the hyperbolic space. Finally, I summarised the released resources such as datasets and frameworks, in which I contributed to their creation, annotation and development.

This document consolidates my own contributions as young researcher, the work of two PhD candidates supervised jointly with academics under the CIFRE convention. I could also collaborate with academics in two ITN projects: Conversational Brains (COBRA) and NL4XAI. As head of the project DIANA during 5 years, I could also define the main workpackages of the project and supervise a dynamic team of researchers, developers, students in internship and apprenticeship.

I would like to focus my future research in proposing a framework to evaluate LLMs performance in complex task-oriented dialogue. Moreover, I would like to explore interpretability, retrieval augmentation and semantically control to understand LLMs internally, to support grounding and to generate factual information. I would like to explore recent reasoning LLMs approach (e.g., ReAct, algorithm distillation, etc.) to study LLMs capabilities to make long-term decisions. It is worth noting that these research paths also cover multi-modal interactions.

¹https://storage.googleapis.com/gresearch/dstc11/dstc11_20221102a.html

Bibliography

- Pieter Abbeel and Andrew Y Ng. 2004. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, page 1. ACM.
- Samuel Albanie, Arsha Nagrani, Andrea Vedaldi, and Andrew Zisserman. 2018. Emotion recognition in speech using cross-modal transfer in the wild. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 292–301.
- J. Allen, D. Byron, M. Dzikovska, G. Ferguson, and L. Galescu. 2001. Towards conversational human-computer interaction.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- Akari Asai, Matt Gardner, and Hannaneh Hajishirzi. 2022. Evidentiality-guided generation for knowledge-intensive NLP tasks. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2226–2243, Seattle, United States. Association for Computational Linguistics.
- J. Austin. 1975a. *How to Do Things With Words (William James Lectures)*. Harvard University Press.
- John Langshaw Austin. 1975b. How to do things with words: (william james lectures). *Harvard University Press*.
- Krishnakumar Balasubramanian, Pinar Donmez, and Guy Lebanon. 2011. Unsupervised supervised learning II: Margin-based classification without labels. *Journal of Machine Learning Research*, 12:3119–3145.
- Ivana Balažević, Carl Allen, and Timothy Hospedales. 2019. Multi-relational poincaré graph embeddings. In *Advances in Neural Information Processing Systems*.
- Andrew G Barto and Sridhar Mahadevan. 2003. Recent advances in hierarchical reinforcement learning. *Discrete event dynamic systems*, 13(1-2):41–77.
- Frédéric Béchet, Cindy Aloui, Delphine Charlet, Geraldine Damnati, Johannes Heinecke, Alexis Nasr, and Frédéric Herledan. 2019. Calor-quest: generating a training corpus for machine reading comprehension models from shallow semantic annotations. In *MRQA: Machine Reading*

for Question Answering-Workshop at EMNLP-IJCNLP 2019-2019 Conference on Empirical Methods in Natural Language Processing.

- Frédéric Béchet, Géraldine Damnati, Johannes Heinecke, Gabriel Marzintotto, and Alexis Nasr. 2017. CALOR-Frame : un corpus de textes encyclopédiques annoté en cadres sémantiques. In *ACor4French – Les corpus annotés du français - Atelier TALN*, Orléans, France.
- Frédéric Béchet, Ludivine Robert, Lina Rojas-Barahona, and Géraldine Damnati. 2022. Calordial: a corpus for conversational question answering on french encyclopedic documents. In *CIRCLE (Joint Conference of the Information Retrieval Communities in Europe)*.
- A.T. Beck. 1976. *Cognitive Therapy and the Emotional Disorders*. New York, International Universities Press.
- A.T. Beck, J. Rush, B. Shaw, and G Emery. 1979. *Cognitive Therapy of Depression*. New York, Guilford Press.
- Silvère Bonnabel. 2013. Stochastic gradient descent on riemannian manifolds. *IEEE Transactions on Automatic Control*, 58:2217–2229.
- Hélène Bonneau-Maynard, Matthieu Quignard, and Alexandre Denis. 2009. Media: a semantically annotated corpus of task oriented dialogs in french: Results of the french media evaluation campaign. *Language Resources and Evaluation*, 43:329–354.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Durán, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2, NIPS’13*, page 2787–2795, Red Hook, NY, USA. Curran Associates Inc.
- Elizabeth Boschee, Jennifer Lautenschlager, Sean O’Brien, Steve Shellman, and James Starz. 2018. *ICEWS Weekly Event Data*. Harvard Dataverse.
- Abdeslam Boularias, Hamid R Chinaei, and Brahim Chaib-draa. 2010. Learning the reward model of dialogue pomdps from data. In *NIPS Workshop on Machine Learning for Assistive Techniques*. Citeseer.
- Quentin Brabant, Gwénolé Lecorvé, and Lina M. Rojas Barahona. 2022. CoQAR: Question rewriting on CoQA. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 119–126, Marseille, France. European Language Resources Association.
- Quentin Brabant, Lina Maria Rojas-Barahona, and Claire Gardent. 2021. Active Learning and Multi-label Classification for Ellipsis and Coreference Detection in Conversational Question-Answering. In *12th International Workshop on Spoken Dialog System Technology (IWSDS 2021)*, Singapour/Virtual, Singapore.
- M. E. Bratman. 1987. *Intention, Plans, and Practical Reason*. Harvard University Press, Cambridge, MA.

BIBLIOGRAPHY

- Paweł Budzianowski, Stefan Ultes, Pei-Hao Su, Nikola Mrkšić, Tsung-Hsien Wen, Iñigo Casanueva, Lina M. Rojas-Barahona, and Milica Gašić. 2017. Sub-domain modelling for dialogue management with hierarchical reinforcement learning. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 86–92, Saarbrücken, Germany. Association for Computational Linguistics.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gašić. 2018. MultiWOZ - a large-scale multi-domain Wizard-of-Oz dataset for task-oriented dialogue modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5016–5026, Brussels, Belgium. Association for Computational Linguistics.
- Harry Bunt, Jan Alexandersson, Jean Carletta, Jae-Woong Choe, Alex Chengyu Fang, Koiti Hasida, Kyong Lee, Volha Petukhova, Andrei Popescu-Belis, Laurent Romary, Claudia Soria, and David Traum. 2010. Towards an iso standard for dialogue act annotation. In *Proceedings of LREC 2010, the Seventh International Conference on Language Resources and Evaluation*, Malta.
- Carl Burke, Christy Doran, Abigail S Gertner, Andy Gregorowicz, Lisa Harper, Joel Korb, and Dan Loehr. 2003. Dialogue complexity with portability? research directions for the information state approach. In *Proceedings of the HLT-NAACL 2003 workshop on Research directions in dialogue processing*, pages 13–15.
- Felix Burkhardt, Astrid Paeschke, Miriam Rolfes, Walter F Sendlmeier, Benjamin Weiss, et al. 2005. A database of german emotional speech. In *Interspeech*, volume 5, pages 1517–1520.
- Michele Cafagna, Lina M. Rojas-Barahona, Kees van Deemter, and Albert Gatt. 2023. Interpreting vision and language generative models with semantic visual priors. *Frontiers in Artificial Intelligence*, 6.
- Iñigo Casanueva, Paweł Budzianowski, Pei-Hao Su, Nikola Mrkšić, Tsung-Hsien Wen, Stefan Ultes, Lina Rojas-Barahona, Steve Young, and Milica Gašić. 2017a. A benchmarking environment for reinforcement learning based task oriented dialogue management. *arXiv preprint arXiv:1711.11023*.
- Iñigo Casanueva, Paweł Budzianowski, Pei-Hao Su, Nikola Mrkšić, Tsung Hsien Wen, Stefan Ultes, Lina Rojas-Barahona, Steve Young, and Milica Gašić. 2017b. A benchmarking environment for reinforcement learning based task oriented dialogue management. *Deep Reinforcement Learning Symposium, NIPS*.
- Iñigo Casanueva, Paweł Budzianowski, Pei-Hao Su, Nikola Mrkšić, Tsung-Hsien Wen, Stefan Ultes, Lina Rojas-Barahona, Steve Young, and Milica Gašić. 2017c. A Benchmarking Environment for Reinforcement Learning Based Task Oriented Dialogue Management. *arXiv:1711.11023 [cs, stat]*. ArXiv: 1711.11023.
- Ines Chami, Adva Wolf, Da-Cheng Juan, Frederic Sala, Sujith Ravi, and Christopher Ré. 2020. Low-dimensional hyperbolic knowledge graph embeddings. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6901–6914, Online. Association for Computational Linguistics.

- Senthilkumar Chandramohan, Matthieu Geist, Fabrice Lefevre, and Olivier Pietquin. 2011. User simulation in dialogue systems using inverse reinforcement learning. In *Interspeech 2011*, pages 1025–1028.
- Chongyan Chen, Samreen Anjum, and Danna Gurari. 2022. Grounding answers for visual questions asked by visually impaired people. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19098–19107.
- Zhi Chen, Xiaoyuan Liu, Lu Chen, and Kai Yu. 2020. Structured hierarchical dialogue policy with graph neural networks. *arXiv preprint arXiv:2009.10355*.
- Adam Cheyer and David Martin. 2001. The open agent architecture. *Autonomous Agents and Multi-Agent Systems*, 4(1-2):143–148.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar. Association for Computational Linguistics.
- E. Choi, H. He, M. Iyyer, M. Yatskar, W-T. Yih, Y. Choi, P. Liang, and L. Zettlemoyer. 2018a. QuAC: Question Answering in Context. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2174–2184, Brussels, Belgium. Association for Computational Linguistics.
- Eunsol Choi, He He, Mohit Iyyer, Mark Yatskar, Wen-tau Yih, Yejin Choi, Percy Liang, and Luke Zettlemoyer. 2018b. QuAC: Question answering in context. In *Proceedings of EMNLP 2018*, pages 2174–2184, Brussels, Belgium.
- P. Christmann, R. Saha Roy, A. Abujabal, J. Singh, and G. Weikum. 2019. Look before you Hop: Conversational Question Answering over Knowledge Graphs Using Judicious Context Expansion. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM '19*, pages 729–738, New York, NY, USA. Association for Computing Machinery.
- G. Churcher. 1997. Dialogue management systems: a survey and overview.
- Herbert H Clark and Susan E Brennan. 1991. Grounding in communication. *American Psychological Association*.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *Journal of Machine Learning Research*, 12(Aug):2493–2537.
- Thibault Cordier, Tanguy Urvoy, Fabrice Lefèvre, and Lina M. Rojas Barahona. 2022. Graph neural network policies and imitation learning for multi-domain task-oriented dialogues. In *Proceedings of the 23rd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 91–100, Edinburgh, UK. Association for Computational Linguistics.

BIBLIOGRAPHY

- Thibault Cordier, Tanguy Urvoy, Fabrice Lefevre, and Lina M Rojas-Barahona. 2023. Few-shot structured policy learning for multi-domain and multi-task dialogues. *The 17th Conference of the European Chapter of the Association for Computational Linguistics. Findings EACL2023*.
- Thibault Cordier, Tanguy Urvoy, Lina M Rojas-Barahona, and Fabrice Lefèvre. 2020. Diluted near-optimal expert demonstrations for guiding dialogue stochastic policy optimisation. In *Human in the loop dialogue systems Workshop at 34th Conference on Neural Information Processing Systems*.
- Heriberto Cuayáhuitl. 2009. Hierarchical reinforcement learning for spoken dialogue systems.
- Heriberto Cuayáhuitl, Simon Keizer, and Oliver Lemon. 2015. Strategic dialogue management via deep reinforcement learning. *arXiv preprint arXiv:1511.08099*.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc V Le, and Ruslan Salakhutdinov. 2019. Transformer-xl: Attentive language models beyond a fixed-length context. *arXiv preprint arXiv:1901.02860*.
- Lucie Daubigney, Matthieu Geist, and Olivier Pietquin. 2013. Model-free pomdp optimisation of tutoring systems with echo-state networks. In *Proceedings of the SIGDIAL 2013 Conference*, pages 102–106.
- Anoop Deoras and Ruhi Sarikaya. 2013. Deep belief network based semantic taggers for spoken language understanding. In *INTERSPEECH*, pages 2713–2717.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Abhinav Dhall, Roland Goecke, Simon Lucey, Tom Gedeon, et al. 2012. Collecting large, richly annotated facial-expression databases from movies. *IEEE multimedia*, 19(3):34.
- Hai Ha Do, PWC Prasad, Angelika Maag, and Abeer Alsadoon. 2019. Deep learning for aspect-based sentiment analysis: a comparative review. *Expert Systems with Applications*, 118:272–299.
- Yan Duan, Xi Chen, Rein Houthoofd, John Schulman, and Pieter Abbeel. 2016. Benchmarking deep reinforcement learning for continuous control.
- Layla El Asri, Romain Laroche, and Olivier Pietquin. 2012. Reward function learning for dialogue management. In *STAIRS 2012*, pages 95–106. IOS Press.
- A. Elgohary, D. Peskov, and J. Boyd-Graber. 2019. Can You Unpack That? Learning to Rewrite Questions-in-Context. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5918–5924, Hong Kong, China. Association for Computational Linguistics.
- Juliette Faille, Albert Gatt, and Claire Gardent. 2021. Entity-based semantic adequacy for data-to-text generation. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1530–1540, Punta Cana, Dominican Republic. Association for Computational Linguistics.

- Angela Fan, Mike Lewis, and Yann Dauphin. 2018. Hierarchical neural story generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898, Melbourne, Australia. Association for Computational Linguistics.
- Jianfeng Gao, Michel Galley, and Lihong Li. 2018. Neural Approaches to Conversational AI. *arXiv:1809.08267 [cs]*. ArXiv: 1809.08267.
- Yang Gao, Huazhe Xu, Ji Lin, Fisher Yu, Sergey Levine, and Trevor Darrell. 2019. Reinforcement Learning from Imperfect Demonstrations. *arXiv:1802.05313 [cs, stat]*. ArXiv: 1802.05313.
- Alberto García-Durán, Sebastijan Dumančić, and Mathias Niepert. 2018. Learning sequence encoders for temporal knowledge graph completion. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4816–4821, Brussels, Belgium. Association for Computational Linguistics.
- Claire Gardent and Lina Maria Rojas-Barahona. 2013. Using Paraphrases and Lexical Semantics to Improve the Accuracy and the Robustness of Supervised Models in Situated Dialogue Systems. In *Conference on Empirical Methods in Natural Language Processing*, pages 808–813, Seattle, United States. SIGDAT, the Association for Computational Linguistics special interest group on linguistic data and corpus-based approaches to NLP.
- Milica Gašić and Steve Young. 2013. Gaussian processes for pomdp-based dialogue manager optimization. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 22(1):28–40.
- Marjan Ghazvininejad, Xing Shi, Jay Priyadarshi, and Kevin Knight. 2017. Hafez: an interactive poetry generation system. In *Proceedings of ACL 2017, System Demonstrations*, pages 43–48, Vancouver, Canada. Association for Computational Linguistics.
- Rishab Goel, Seyed Mehran Kazemi, Marcus Brubaker, and Pascal Poupart. 2020. Diachronic embedding for temporal knowledge graph completion. In *AAAI*.
- Chih-Wen Goo and Yun-Nung Chen. 2018. Abstractive dialogue summarization with sentence-gated modeling optimized by dialogue acts.
- Aditya Grover, Jiaming Song, Ashish Kapoor, Kenneth Tran, Alekh Agarwal, Eric J Horvitz, and Stefano Ermon. 2019. Bias correction of learned generative models using likelihood-free importance weighting. *Advances in neural information processing systems*, 32.
- Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, and Sergey Levine. 2017. Reinforcement Learning with Deep Energy-Based Policies. *arXiv:1702.08165 [cs]*. ArXiv: 1702.08165.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. 2018. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. *arXiv:1801.01290 [cs, stat]*. ArXiv: 1801.01290.
- Philipp Hacker, Andreas Engel, and Marco Mauer. 2023. Regulating chatgpt and other large generative ai models. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, pages 1112–1123.

BIBLIOGRAPHY

- Zhen Han, Peng Chen, Yunpu Ma, and Volker Tresp. 2020. DyERNIE: Dynamic Evolution of Riemannian Manifold Embeddings for Temporal Knowledge Graph Completion. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7301–7316, Online. Association for Computational Linguistics.
- Thomas L. Heath and Euclid. 1956. *The Thirteen Books of Euclid’s Elements, Books 1 and 2*. Dover Publications, Inc., USA.
- M. Henderson, B. Thomson, and S. J. Young. 2014a. Word-based Dialog State Tracking with Recurrent Neural Networks. In *Proceedings of SIGdial*.
- Matthew Henderson, Milica Gašić, Blaise Thomson, Pirros Tsiakoulis, Kai Yu, and Steve Young. 2012. Discriminative Spoken Language Understanding Using Word Confusion Networks. In *Spoken Language Technology Workshop, 2012. IEEE*.
- Matthew Henderson, Blaise Thomson, and Jason Williams. 2014b. The second dialog state tracking challenge. In *15th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, volume 263.
- Matthew Henderson, Blaise Thomson, and Jason Williams. 2014c. The third dialog state tracking challenge. In *Proceedings IEEE Spoken Language Technology Workshop (SLT)*. IEEE Institute of Electrical and Electronics Engineers.
- Todd Hester, Matej Vecerik, Olivier Pietquin, Marc Lanctot, Tom Schaul, Bilal Piot, Dan Horgan, John Quan, Andrew Sendonaris, Gabriel Dulac-Arnold, Ian Osband, John Agapiou, Joel Z. Leibo, and Audrunas Gruslys. 2017. Deep Q-learning from Demonstrations. *arXiv:1704.03732 [cs]*. ArXiv: 1704.03732.
- Todd Hester, Matej Vecerik, Olivier Pietquin, Marc Lanctot, Tom Schaul, Bilal Piot, Dan Horgan, John Quan, Andrew Sendonaris, Ian Osband, et al. 2018. Deep q-learning from demonstrations. In *Thirty-second AAAI conference on artificial intelligence*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2019. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Léo Jacqmin, Lucas Druart, Valentin Vielzeuf, Lina Maria Rojas-Barahona, Yannick Estève, and Benoît Favre. 2023. Olisia: a cascade system for spoken dialogue state tracking.
- Y. Ju, F. Zhao, S. Chen, B. Zheng, X. Yang, and Y. Liu. 2019. Technical report on conversational question answering. *arXiv preprint arXiv:1909.10772*.

- Nitish Shirish Keskar, Bryan McCann, Lav R Varshney, Caiming Xiong, and Richard Socher. 2019. Ctrl: A conditional transformer language model for controllable generation. *arXiv preprint arXiv:1909.05858*.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882*.
- S. Larsson and D. Traum. 2000. Information state and dialogue management in the trindi dialogue move engine toolkit. *Natural Language Engineering*, 6:323–340.
- Michael Laskin, Luyu Wang, Junhyuk Oh, Emilio Parisotto, Stephen Spencer, Richie Steigerwald, DJ Strouse, Steven Hansen, Angelos Filos, Ethan Brooks, et al. 2022. In-context reinforcement learning with algorithm distillation. *arXiv preprint arXiv:2210.14215*.
- Chih-Wei Lee, Yau-Shian Wang, Tsung-Yuan Hsu, Kuan-Yu Chen, Hung-yi Lee, and Lin-shan Lee. 2018. Scalable sentiment for sequence-to-sequence chatbot response with performance analysis. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6164–6168. IEEE.
- Kyusong Lee, Tiancheng Zhao, Yulun Du, Edward Cai, Allen Lu, Eli Pincus, David Traum, Stefan Ultes, Lina M. Rojas-Barahona, Milica Gasic, Steve Young, and Maxine Eskenazi. 2017. DialPort, gone live: An update after a year of development. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 170–173, Saarbrücken, Germany. Association for Computational Linguistics.
- Kyusong Lee, Tiancheng Zhao, Stefan Ultes, Lina Rojas-Barahona, Eli Pincus, David Traum, and Maxine Eskenazi. 2019. An assessment framework for dialport. In *Advanced Social Interaction with Agents: 8th International Workshop on Spoken Dialog Systems*, pages 79–85. Springer.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The power of scale for parameter-efficient prompt tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Esther Levin, Roberto Pieraccini, and Wieland Eckert. 2000. A stochastic model of human-machine interaction for learning dialog strategies. *IEEE Transactions on speech and audio processing*, 8(1):11–23.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4582–4597, Online. Association for Computational Linguistics.
- Yuxi Li. 2018. Deep Reinforcement Learning. *arXiv:1810.06339 [cs, stat]*. ArXiv: 1810.06339.

BIBLIOGRAPHY

- Diane Litman, Satinder Singh, Michael Kearns, and Marilyn Walker. 2000. NJFun- a reinforcement learning spoken dialogue system. In *ANLP-NAACL 2000 Workshop: Conversational Systems*.
- Hao Liu, Carmelo Sferrazza, and Pieter Abbeel. 2023. Chain of hindsight aligns language models with feedback.
- Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Alejandra Lorenzo, Lina Rojas-Barahona, and Christophe Cerisara. 2013. Unsupervised structured semantic inference for spoken dialog reservation tasks. In *Proceedings of the SIGDIAL 2013 Conference*, pages 12–20, Metz, France. Association for Computational Linguistics.
- Ryan Lowe, Nissan Pow, Iulian Serban, and Joelle Pineau. 2015. The ubuntu dialogue corpus: A large dataset for research in unstructured multi-turn dialogue systems. *arXiv preprint arXiv:1506.08909*.
- Yukun Ma, Khanh Linh Nguyen, Frank Z Xing, and Erik Cambria. 2020. A survey on empathetic dialogue systems. *Information Fusion*.
- Yuning Mao, Pengcheng He, Xiaodong Liu, Yelong Shen, Jianfeng Gao, Jiawei Han, and Weizhu Chen. 2021. Reader-guided passage reranking for open-domain question answering. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 344–350, Online. Association for Computational Linguistics.
- Gabriel Marzinotto, Jeremy Auguste, Frederic Bechet, Géraldine Damnati, and Alexis Nasr. 2018. Semantic frame parsing for information extraction : the calor corpus. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC-2018)*. European Language Resource Association.
- Andrew Kachites McCallum. 2002. Mallet: A machine learning for language toolkit. [Http://mallet.cs.umass.edu](http://mallet.cs.umass.edu).
- Michael F. McTear. 2002. Spoken dialogue technology: enabling the conversational user interface. *ACM Comput. Surv.*, 34(1):90–169.
- Grégoire Mesnil, Yann Dauphin, Kaisheng Yao, Yoshua Bengio, Li Deng, Dilek Hakkani-Tur, Xiaodong He, Larry Heck, Gokhan Tur, Dong Yu, et al. 2015. Using recurrent neural networks for slot filling in spoken language understanding. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 23(3):530–539.
- Bernard Michini and Jonathan P How. 2012. Improving the efficiency of bayesian inverse reinforcement learning. In *2012 IEEE International Conference on Robotics and Automation*, pages 3651–3656. IEEE.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*.

- Sebastien Montella, Lina M Rojas Barahona, Frederic Bechet, Johannes Heinecke, and Alexis Nasr. 2022. Transfer learning and masked generation for answer verbalization. In *Proceedings of the Workshop on Structured and Unstructured Knowledge Integration (SUKI)*, pages 47–54.
- Sebastien Montella, Alexis Nasr, Johannes Heinecke, Frederic Bechet, and Lina M Rojas-Barahona. 2023. Investigating the effect of relative positional embeddings on amr-to-text generation with structural adapters. *arXiv preprint arXiv:2302.05900*.
- Sebastien Montella, Lina M. Rojas Barahona, and Johannes Heinecke. 2021. Hyperbolic temporal knowledge graph embeddings with relational and time curvatures. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3296–3308, Online. Association for Computational Linguistics.
- RR Morris, Schueller SM, and Picard RW. 2015. Efficacy of a Web-Based, Crowdsourced Peer-To-Peer Cognitive Reappraisal Platform for Depression: Randomized Controlled Trial. *J Med Internet Res*, 17(3).
- Nikola Mrkšić, Diarmuid Ó Séaghdha, Tsung-Hsien Wen, Blaise Thomson, and Steve Young. 2017. Neural belief tracker: Data-driven dialogue state tracking. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1777–1788, Vancouver, Canada. Association for Computational Linguistics.
- Rémi Munos, Tom Stepleton, Anna Harutyunyan, and Marc G. Bellemare. 2016. Safe and Efficient Off-Policy Reinforcement Learning. *arXiv:1606.02647 [cs, stat]*. ArXiv: 1606.02647.
- Shashi Narayan, Gonçalo Simões, Yao Zhao, Joshua Maynez, Dipanjan Das, Michael Collins, and Mirella Lapata. 2022. A well-composed text is half done! composition sampling for diverse conditional generation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1319–1339, Dublin, Ireland. Association for Computational Linguistics.
- Andrew Y Ng, Stuart J Russell, et al. 2000. Algorithms for inverse reinforcement learning. In *ICML*, volume 1, page 2.
- Jinjie Ni, Tom Young, Vlad Pandelea, Fuzhao Xue, and Erik Cambria. 2022. Recent advances in deep learning based dialogue systems: A systematic survey. *Artificial Intelligence Review*, pages 1–101.
- Feng Nie, Jin-Ge Yao, Jinpeng Wang, Rong Pan, and Chin-Yew Lin. 2019. A simple recipe towards reducing hallucination in neural surface realisation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2673–2679.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *arXiv preprint arXiv:2203.02155*.
- Tim Paek and Roberto Pieraccini. 2008. Automating spoken dialogue management design using machine learning: An industry perspective. *Speech communication*, 50(8-9):716–729.

BIBLIOGRAPHY

- Ronald Parr and Stuart J Russell. 1998. Reinforcement learning with hierarchies of machines. In *Advances in neural information processing systems*, pages 1043–1049.
- Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli, Hamza Alobeidli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. 2023. The refinedweb dataset for falcon llm: Outperforming curated corpora with web data, and web data only.
- Leonardo Pepino, Pablo Riera, and Luciana Ferrer. 2021. Emotion Recognition from Speech Using wav2vec 2.0 Embeddings. In *Proc. Interspeech 2021*, pages 3400–3404.
- Roberto Pieraccini, Evelyne Tzoukermann, Zakhar Gorelov, J-L Gauvain, Esther Levin, C-H Lee, and Jay G Wilpon. 1992. A speech understanding system based on statistical representation of semantics. In *Acoustics, Speech, and Signal Processing, 1992. ICASSP-92., 1992 IEEE International Conference on*, volume 1, pages 193–196. IEEE.
- Matt Post. 2018. A call for clarity in reporting BLEU scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- J. Quan, D. Xiong, B. Webber, and C. Hu. 2019. GECOR: An End-to-End Generative Ellipsis and Co-reference Resolution Model for Task-Oriented Dialogue. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4547–4557, Hong Kong, China. Association for Computational Linguistics.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2022. Robust speech recognition via large-scale weak supervision. *arXiv preprint arXiv:2212.04356*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know What You Don’t Know: Unanswerable Questions for SQuAD. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, Melbourne, Australia. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392.
- Deepak Ramachandran and Eyal Amir. 2007. Bayesian inverse reinforcement learning. In *IJCAI*, volume 7, pages 2586–2591.
- Antoine Raux, Brian Langner, Dan Bohus, Alan W Black, and Maxine Eskenazi. 2005. Let’s go public! taking a spoken dialog system to the real world. In *Ninth European conference on speech communication and technology*.

- S. Reddy, D. Chen, and C. D. Manning. 2019. CoQA: A Conversational Question Answering Challenge. *Transactions of the Association for Computational Linguistics*, 7:249–266.
- Siva Reddy, Danqi Chen, and Christopher D Manning. 2018. Coqa: A conversational question answering challenge. *arXiv preprint arXiv:1808.07042*.
- Verena Rieser and Oliver Lemon. 2011. *Reinforcement learning for adaptive dialogue systems: a data-driven methodology for dialogue management and natural language generation*. Springer Science & Business Media.
- Lina M Rojas-Barahona. 2020. Is the user enjoying the conversation? a case study on the impact on the reward function. *34st Conference on Neural Information Processing Systems (NeurIPS 2020)*.
- Lina M Rojas-Barahona, Pascal Bellec, Benoit Besset, Martinho Dossantos, Johannes Heinecke, Munshi Asadullah, Olivier Leblouch, Jeanyves Lancien, Géraldine Damnati, Emmanuel Mory, et al. 2019. Spoken conversational search for general knowledge. In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 110–113.
- Lina M Rojas-Barahona and Christophe Cerisara. 2014. Bayesian inverse reinforcement learning for modeling conversational agents in a virtual environment. In *International Conference on Intelligent Text Processing and Computational Linguistics*, pages 503–514. Springer.
- Lina M. Rojas Barahona, Milica Gašić, Nikola Mrkšić, Pei-Hao Su, Stefan Ultes, Tsung Hsien Wen, and Steve Young. 2016. Exploiting sentence and context representations in deep neural models for spoken language understanding. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 258–267, Osaka, Japan.
- Lina M. Rojas-Barahona, Alejandra Lorenzo, and Claire Gardent. 2012a. Building and Exploiting a Corpus of Dialog Interactions between French Speaking Virtual and Human Agents. In *Proceedings of the 8th International Conference on Language Resources and Evaluation*.
- Lina M Rojas-Barahona and Matthieu Quignard. 2011. An incremental architecture for the semantic annotation of dialogue corpora with high-level structures. a case of study for the media corpus. In *Proceedings of the SIGDIAL 2011 Conference*, pages 332–334.
- Lina M. Rojas-Barahona, Bo-Hsiang Tseng, Yinpei Dai, Clare Mansfield, Osman Ramadan, Stefan Ultes, Michael Crawford, and Milica Gašić. 2018. Deep learning for language understanding of mental health concepts derived from cognitive behavioural therapy. In *Proceedings of the Ninth International Workshop on Health Text Mining and Information Analysis*, pages 44–54, Brussels, Belgium. Association for Computational Linguistics.
- Lina Maria Rojas-Barahona. 2016. Deep learning for sentiment analysis. *Language and Linguistics Compass*, 10(12):701–719.
- Lina Maria Rojas-Barahona, Thierry Bazillon, Matthieu Quignard, and Fabrice Lefevre. 2011. Using mmil for the high level semantic annotation of the french media dialogue corpus. In *Ninth International Conference on Computational Semantics-IWCS 2011*.

BIBLIOGRAPHY

- Lina Maria Rojas Barahona and Christophe Cerisara. 2015. Weakly supervised discriminative training of linear models for Natural Language Processing. In *3rd International Conference on Statistical Language and Speech Processing (SLSP)*, Budapest, Hungary.
- Lina Maria Rojas-Barahona and Claire Gardent. 2012. What should i do now? supporting conversations in a serious game. In *SeineDial 2012-16th Workshop on the Semantics and Pragmatics of Dialogue*.
- Lina Maria Rojas-Barahona and Toni Giorgino. 2009. Adaptable dialog architecture and runtime engine (adarte): a framework for rapid prototyping of health dialog systems. *international journal of medical informatics*, 78:S56–S68.
- Lina Maria Rojas-Barahona, Alejandra Lorenzo, and Claire Gardent. 2012b. Building and exploiting a corpus of dialog interactions between french speaking virtual and human agents. In *The eighth international conference on Language Resources and Evaluation (LREC)*, pages 1428–1435.
- Lina Maria Rojas-Barahona, Alejandra Lorenzo, and Claire Gardent. 2012c. An end-to-end evaluation of two situated dialog systems. In *Proceedings of the 13th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 10–19.
- Stéphane Ross, Geoffrey J Gordon, and J Andrew Bagnell. 2011. A Reduction of Imitation Learning and Structured Prediction to No-Regret Online Learning. *Proceedings of the Workshop on Artificial Intelligence and Statistics, AISTATS*, page 9.
- Nicholas Roy, Joelle Pineau, and Sebastian Thrun. 2000. Spoken dialogue management using probabilistic reasoning. In *Proceedings of the 38th annual meeting of the association for computational linguistics*, pages 93–100.
- Benoît Sagot. 2010. The Lefff, a freely available and large-coverage morphological and syntactic lexicon for French. In *7th international conference on Language Resources and Evaluation (LREC 2010)*, Valletta, Malta.
- Amrita Saha, Vardaan Pahuja, Mitesh M Khapra, Karthik Sankaranarayanan, and Sarath Chandar. 2018. Complex sequential question answering: Towards learning to converse over linked question answer pairs with a knowledge graph. *arXiv preprint arXiv:1801.10314*.
- V. Sanh, L. Debut, J. Chaumond, and T. Wolf. 2020. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *arXiv:1910.01108 [cs]*. ArXiv: 1910.01108.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Chinnadhurai Sankar and Sujith Ravi. 2019. Deep reinforcement learning for modeling chit-chat dialog with discrete attributes. *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*.
- Bishal Santra, Sumegh Roychowdhury, Aishik Mandal, Vasu Gurram, Atharva Naik, Manish Gupta, and Pawan Goyal. 2021. Representation learning for conversational data using discourse mutual information maximization. *arXiv preprint arXiv:2112.05787*.

- R. Sarikaya, G. E. Hinton, and B. Ramabhadran. 2011. Deep belief nets for natural language call-routing. In *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5680–5683.
- Ruhi Sarikaya, Geoffrey E Hinton, and Anoop Deoras. 2014. Application of deep belief networks for natural language understanding. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 22(4):778–784.
- Jost Schatzmann, Blaise Thomson, Karl Weilhammer, Hui Ye, and Steve Young. 2007. Agenda-based user simulation for bootstrapping a POMDP dialogue system. In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Companion Volume, Short Papers on XX - NAACL '07*, pages 149–152, Rochester, New York. Association for Computational Linguistics.
- Jost Schatzmann, Karl Weilhammer, Matt Stuttle, and Steve Young. 2006. A survey of statistical user simulation techniques for reinforcement-learning of dialogue management strategies. *Knowledge Engineering Review*, 21(2):97–126.
- Neil Scheidwasser-Clow, Mikolaj Kegler, Pierre Beckmann, and Milos Cernak. 2022. Serab: A multi-lingual benchmark for speech emotion recognition. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7697–7701. IEEE.
- Alexander Schmitt, Benjamin Schatz, and Wolfgang Minker. 2011. Modeling and predicting quality in spoken human-computer interaction. In *Proceedings of the SIGDIAL 2011 Conference*, pages 173–184. Association for Computational Linguistics.
- Alexander Schmitt, Stefan Ultes, and Wolfgang Minker. 2012. A parameterized and annotated spoken dialog corpus of the CMU let’s go bus information system. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC’12)*, pages 3369–3373, Istanbul, Turkey. European Language Resources Association (ELRA).
- Marc Schröder, Laurence Devillers, Kostas Karpouzis, Jean-Claude Martin, Catherine Pelachaud, Christian Peter, Hannes Pirker, Björn Schuller, Jianhua Tao, and Ian Wilson. 2007. What should a generic emotion markup language be able to represent? In *Affective Computing and Intelligent Interaction: Second International Conference, AII 2007 Lisbon, Portugal, September 12-14, 2007 Proceedings 2*, pages 440–451. Springer.
- John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. 2015. Trust region policy optimization. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 1889–1897, Lille, France. PMLR.
- Min Joon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Bidirectional attention flow for machine comprehension. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.

BIBLIOGRAPHY

- Iulian Serban, Alessandro Sordoni, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2017. A hierarchical latent variable encoder-decoder model for generating dialogues. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31.
- Iulian V Serban, Alessandro Sordoni, Yoshua Bengio, Aaron Courville, and Joelle Pineau. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *Thirtieth AAAI Conference on Artificial Intelligence*.
- Kurt Shuster, Jing Xu, Mojtaba Komeili, Da Ju, Eric Michael Smith, Stephen Roller, Megan Ung, Moya Chen, Kushal Arora, Joshua Lane, et al. 2022. Blenderbot 3: a deployed conversational agent that continually learns to responsibly engage. *arXiv preprint arXiv:2208.03188*.
- Alessandro Sordoni, Yoshua Bengio, Hossein Vahabi, Christina Lioma, Jakob Grue Simonsen, and Jian-Yun Nie. 2015. A hierarchical recurrent encoder-decoder for generative context-aware query suggestion. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*, pages 553–562. ACM.
- Yash Srivastava, Vaishnav Murali, Shiv Ram Dubey, and Snehasis Mukherjee. 2021. Visual question answering using deep learning: A survey and performance analysis. In *Computer Vision and Image Processing: 5th International Conference, CVIP 2020, Prayagraj, India, December 4-6, 2020, Revised Selected Papers, Part II 5*, pages 75–86. Springer.
- Pei-Hao Su, Paweł Budzianowski, Stefan Ultes, Milica Gašić, and Steve Young. 2017. Sample-efficient actor-critic reinforcement learning with supervised data for dialogue management. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 147–157, Saarbrücken, Germany. Association for Computational Linguistics.
- Pei-Hao Su, Milica Gašić, Nikola Mrkšić, Lina M. Rojas-Barahona, Stefan Ultes, David Vandyke, Tsung-Hsien Wen, and Steve Young. 2016. On-line active reward learning for policy optimisation in spoken dialogue systems. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2431–2441, Berlin, Germany. Association for Computational Linguistics.
- Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.
- Richard S Sutton, Doina Precup, and Satinder Singh. 1999. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*, 112(1-2):181–211.
- Da Tang, Xiujun Li, Jianfeng Gao, Chong Wang, Lihong Li, and Tony Jebara. 2018. Subgoal discovery for hierarchical dialogue policy learning. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2298–2309, Brussels, Belgium. Association for Computational Linguistics.
- Blaise Thomson, Kai Yu, Milica Gasic, Simon Keizer, Francois Mairesse, Jost Schatzmann, and Steve J Young. 2008. Evaluating semantic-level confidence scores with multiple hypotheses. In *INTERSPEECH*, pages 1153–1156.

- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. FEVER: a large-scale dataset for fact extraction and VERification. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 809–819, New Orleans, Louisiana. Association for Computational Linguistics.
- Tijmen Tieleman and Geoffrey Hinton. 2012. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural networks for machine learning*, 4(2):26–31.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- David R. Traum and Staffan Larsson. 2003. The information state approach to dialogue management. In Jan Kuppevelt, Ronnie W. Smith, and Nancy Ide, editors, *Current and New Directions in Discourse and Dialogue*, volume 22 of *Text, Speech and Language Technology*, pages 325–353. Springer Netherlands. 10.1007/978-94-010-0019-2_15.
- George Trigeorgis, Fabien Ringeval, Raymond Brueckner, Erik Marchi, Mihalis A Nicolaou, Björn Schuller, and Stefanos Zafeiriou. 2016. Adieu features? end-to-end speech emotion recognition using a deep convolutional recurrent network. In *2016 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 5200–5204. IEEE.
- Pirros Tsiakoulis, Milica Gašić, Matthew Henderson, Joaquin Planells-Lerma, Jorge Prombonas, Blaise Thomson, Kai Yu, Steve Young, and Eli Tzirkel. 2012. Statistical methods for building robust spoken dialogue systems in an automobile. *Proceedings of the 4th applied human factors and ergonomics*.
- Gökhan Tür, Anoop Deoras, and Dilek Hakkani-Tür. 2013. Semantic parsing using word confusion networks with conditional random fields. In *INTERSPEECH*, pages 2579–2583.
- S. Ultes, L. M. Rojas-Barahona, P-H. Su, D. Vandyke, D. Kim, I. Casanueva, P. Budzianowski, N. Mrkšić, T-H. Wen, M. Gašić, and S. Young. 2017a. PyDial: A Multi-domain Statistical Dialogue System Toolkit. In *Proceedings of ACL 2017, System Demonstrations*, pages 73–78, Vancouver, Canada. Association for Computational Linguistics.
- Stefan Ultes. 2019. Improving interaction quality estimation with BiLSTMs and the impact on dialogue policy learning. In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 11–20, Stockholm, Sweden. Association for Computational Linguistics.
- Stefan Ultes, Matthias Kraus, Alexander Schmitt, and Wolfgang Minker. 2015. Quality-adaptive spoken dialogue initiative selection and implications on reward modelling. In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 374–383, Prague, Czech Republic. Association for Computational Linguistics.
- Stefan Ultes, Lina M. Rojas Barahona, Pei-Hao Su, David Vandyke, Dongho Kim, Iñigo Casanueva, Paweł Budzianowski, Nikola Mrkšić, Tsung-Hsien Wen, Milica Gasic, and Steve Young. 2017b. PyDial: A Multi-domain Statistical Dialogue System Toolkit. In *Proceedings*

BIBLIOGRAPHY

- of *ACL 2017, System Demonstrations*, pages 73–78, Vancouver, Canada. Association for Computational Linguistics.
- Stefan Ultes, Lina M. Rojas Barahona, Pei-Hao Su, David Vandyke, Dongho Kim, Iñigo Casanueva, Paweł Budzianowski, Nikola Mrkšić, Tsung-Hsien Wen, Milica Gasic, and Steve Young. 2017c. PyDial: A Multi-domain Statistical Dialogue System Toolkit. In *Proceedings of ACL 2017, System Demonstrations*, pages 73–78, Vancouver, Canada. Association for Computational Linguistics.
- S. Vakulenko, S. Longpre, Z. Tu, and R. Anantha. 2021. Question rewriting for conversational question answering. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, pages 355–363.
- Hado van Hasselt, Arthur Guez, and David Silver. 2015. Deep Reinforcement Learning with Double Q-learning. *arXiv:1509.06461 [cs]*. ArXiv: 1509.06461.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Oriol Vinyals and Quoc Le. 2015. A neural conversational model. *arXiv preprint arXiv:1506.05869*.
- Piek Vossen, editor. 1998. *EuroWordNet: a multilingual database with lexical semantic networks*. Kluwer Academic Publishers, Norwell, MA, USA.
- Marilyn A. Walker, Diane J. Litman, Candace A. Kamm, and Alicia Abella. 1997a. PARADISE: A framework for evaluating spoken dialogue agents. In *35th Annual Meeting of the Association for Computational Linguistics and 8th Conference of the European Chapter of the Association for Computational Linguistics*, pages 271–280, Madrid, Spain. Association for Computational Linguistics.
- Marilyn A Walker, Diane J Litman, Candace A Kamm, and Alicia Abella. 1997b. Paradise: A framework for evaluating spoken dialogue agents. *arXiv preprint cmp-lg/9704004*.
- Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. 2022. Ofa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. In *International Conference on Machine Learning*, pages 23318–23340. PMLR.
- Zhuoran Wang, Tsung-Hsien Wen, Pei-Hao Su, and Yannis Stylianou. 2015. Learning domain-independent dialogue policies via ontology parameterisation. In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 412–416.
- Ziyu Wang, Victor Bapst, Nicolas Heess, Volodymyr Mnih, Remi Munos, Koray Kavukcuoglu, and Nando de Freitas. 2016a. Sample efficient actor-critic with experience replay. *arXiv preprint arXiv:1611.01224*.

- Ziyu Wang, Victor Bapst, Nicolas Heess, Volodymyr Mnih, Remi Munos, Koray Kavukcuoglu, and Nando de Freitas. 2017. Sample Efficient Actor-Critic with Experience Replay. *arXiv:1611.01224 [cs]*. ArXiv: 1611.01224.
- Ziyu Wang, Tom Schaul, Matteo Hessel, Hado van Hasselt, Marc Lanctot, and Nando de Freitas. 2016b. Dueling Network Architectures for Deep Reinforcement Learning. *International conference on machine learning*, pages 1995–2003.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.
- Gellért Weisz, Paweł Budzianowski, Pei-Hao Su, and Milica Gašić. 2018a. Sample efficient deep reinforcement learning for dialogue systems with large action spaces. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 26(11):2083–2097.
- Gellért Weisz, Paweł Budzianowski, Pei-Hao Su, and Milica Gašić. 2018b. Sample Efficient Deep Reinforcement Learning for Dialogue Systems with Large Action Spaces. *arXiv:1802.03753 [cs, stat]*. ArXiv: 1802.03753.
- Joseph Weizenbaum. 1966. Eliza: A computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9(1).
- Tsung-Hsien Wen, Milica Gasic, Nikola Mrksic, Pei-Hao Su, David Vandyke, and Steve Young. 2015. Semantically conditioned lstm-based natural language generation for spoken dialogue systems. *arXiv preprint arXiv:1508.01745*.
- Tsung-Hsien Wen, Yishu Miao, Phil Blunsom, and Steve Young. 2017a. Latent intention dialogue models. In *International Conference on Machine Learning*, pages 3732–3741. PMLR.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrkšić, Milica Gašić, Lina M. Rojas-Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2017b. A network-based end-to-end trainable task-oriented dialogue system. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 438–449, Valencia, Spain. Association for Computational Linguistics.
- Jason Weston, Sumit Chopra, and Antoine Bordes. 2014. Memory networks. *CoRR*, abs/1410.3916.
- Jason D Williams. 2014. Web-style ranking and slu combination for dialog state tracking. In *Proceedings of the 15th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, pages 282–291.
- Jason D Williams and Steve Young. 2007. Partially observable markov decision processes for spoken dialog systems. *Computer Speech & Language*, 21(2):393–422.
- Yu Wu, Zhoujun Li, Wei Wu, and Ming Zhou. 2018. Response selection with topic clues for retrieval-based chatbots. *Neurocomputing*, 316:251–261.

BIBLIOGRAPHY

- Chengjin Xu, Mojtaba Nayeri, Fouad Alkhoury, Hamed Shariat Yazdi, and Jens Lehmann. 2020. TeRo: A time-aware knowledge graph embedding via temporal rotation. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1583–1593, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Jing Xu, Arthur Szlam, and Jason Weston. 2022. Beyond goldfish memory: Long-term open-domain conversation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5180–5197, Dublin, Ireland. Association for Computational Linguistics.
- Kaisheng Yao, Geoffrey Zweig, Mei-Yuh Hwang, Yangyang Shi, and Dong Yu. 2013. Recurrent neural networks for language understanding. In *INTERSPEECH*, pages 2524–2528.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of thoughts: Deliberate problem solving with large language models. *arXiv preprint arXiv:2305.10601*.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2022. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*.
- Michihiro Yasunaga, Armen Aghajanyan, Weijia Shi, Richard James, Jure Leskovec, Percy Liang, Mike Lewis, Luke Zettlemoyer, and Wen-Tau Yih. 2023. Retrieval-augmented multimodal language modeling. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 39755–39769. PMLR.
- Steve Young, Milica Gašić, Simon Keizer, François Mairesse, Jost Schatzmann, Blaise Thomson, and Kai Yu. 2010. The hidden information state model: a practical framework for pomdp-based spoken dialogue management. *Computer Speech & Language*, 24(2):150–174.
- Steve Young, Milica Gašić, Blaise Thomson, and Jason D Williams. 2013a. Pomdp-based statistical spoken dialog systems: A review. *Proceedings of the IEEE*, 101(5):1160–1179.
- Steve Young, Milica Gasic, Blaise Thomson, and Jason D. Williams. 2013b. POMDP-Based Statistical Spoken Dialog Systems: A Review. *Proceedings of the IEEE*, 101(5):1160–1179.
- Steve J Young. 2002. Talking to machines (statistically speaking). In *INTERSPEECH*. Citeseer.
- Amir Zadeh, Yan Sheng Cao, Simon Hessner, Paul Pu Liang, Soujanya Poria, and Louis-Philippe Morency. 2020. Cmu-moseas: A multimodal language dataset for spanish, portuguese, german and french. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing. Conference on Empirical Methods in Natural Language Processing*, volume 2020, page 1801. NIH Public Access.
- AmirAli Bagher Zadeh, Paul Pu Liang, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. 2018. Multimodal language analysis in the wild: Cmu-mosei dataset and interpretable dynamic fusion graph. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2236–2246.

- T. Zhang, V. Kishore, F. Wu, K.Q. Weinberger, and Y. Artzi. 2020a. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020b. DIALOGPT : Large-scale generative pre-training for conversational response generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 270–278, Online. Association for Computational Linguistics.
- Deyu Zhou and Yulan He. 2011. Learning conditional random fields from unaligned data for natural language understanding. In *European Conference on Information Retrieval*, pages 283–288. Springer.
- Qi Zhu, Zheng Zhang, Yan Fang, Xiang Li, Ryuichi Takanobu, Jinchao Li, Baolin Peng, Jianfeng Gao, Xiaoyan Zhu, and Minlie Huang. 2020. Convlab-2: An open-source toolkit for building, evaluating, and diagnosing dialogue systems. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 142–149.
- Yuke Zhu, Oliver Groth, Michael Bernstein, and Li Fei-Fei. 2016. Visual7w: Grounded question answering in images.