Complementing RRL for Dialogue Summarisation*

Norton Trevisan Roman[1] and Ariadne Maria Brito Rizzoni Carvalho[2]

[1] School of Arts, Sciences and Humanities – University of São Paulo, São Paulo, Brazil
norton@usp.br
[2] Institute of Computing – University of Campinas, Campinas, Brazil
ariadne@ic.unicamp.br

Abstract. This paper describes the use of NECA's RRL – Rich Representation Language – in the aid of automatic dialogue summarisation. We start by defining a set of summary acts, and then present a mapping between these acts, taken from human generated summaries, and their corresponding speech acts from the source dialogues. This mapping can then be used by a summary planner to identify portions in the dialogue with higher probability to support some argument it is about to introduce in the summary. Depending on the application, it could also be used to gather, in the source dialogue, specific information to fulfill summary templates. In following this procedure, we hope researchers will be able to build more human-like summaries, by linking corpora summaries to their source dialogues.

1 Introduction

RRL – *Rich Representation Language* [11, 12] – is an XML-based coding language used to represent, amongst other things, the gestures, facial expressions, speech and semantic information exchanged by conversational agents. Originally developed for the NECA platform[1] [4], this language is also designed for representing each speech act in a dialogue (called Dialogue Act in NECA), as well as its semantic content, along the lines of Kamp and Reyle's Discourse Representation Theory (DRT) [6].

The fact that, within NECA, the user can build some Embodied Conversational Agents (ECAs), setting up their personality traits (such as their politeness degree, for instance) and then automatically generating dialogues between these agents, makes it a suitable platform for other dialogue related tasks. One such task is the automatic generation of summaries which, given an RRL, is to produce a summary for the dialogue the RRL represents (in the same representation language). Also, although NECA has been developed to generate dialogues for

* This research was sponsored by CNPq – Conselho Nacional de Desenvolvimento Científico e Tecnológico – and CAPES – Coordenação de Aperfeiçoamento de Pessoal de Nível Superior.

[1] Net Environment for Embodied Emotional Conversational Agents.

A. Kuri-Morales and G. Simari (Eds.): IBERAMIA 2010, LNAI 6433, pp. 376–385, 2010.

two distinct domains, in this paper we will only focus on one of them – *eShowRoom* – which illustrates the interaction between a car seller and a prospective buyer.

However useful, the greatest problem with NECA's RRL lies in the fact that, even though each and every utterance is assigned a dialogue act (*e.g. inform*, *greeting* etc), some of them have no explicitly defined semantic content. As a matter of fact, in the *eShowRoom* scenario there seems to be only two instances where such information is found[2] (see [8]), to wit, (a) when some of the car characteristics are informed or asked for; or (b) when some specific topic is brought up for discussion, as in "I'd like to know more about this car". This limitation turns out to be rather damaging for any automatic system that must rely on a deeper semantic analysis of the dialogue to build its outputs.

In this paper, we describe our attempt to address the above limitation, by mapping the semantic content of human generated summaries (all abstracts), taken from previous research [14–16], to their counterparts in the RRL of the source dialogues. To do so, we first define a set of summary acts – those acts human summarisers seem to perform when building dialogue summaries – and then annotate summaries accordingly. Next, we relate each of these acts to their corresponding dialogue acts in the source dialogue, creating a mapping. Differently from current literature (*e.g.* [2, 3]), however, which deals almost exclusively with lexical and syntactical information, our approach works at the semantic level, thereby making it possible to produce summaries in a language different from that of the summarised dialogue.

With this mapping at hand, automatic summarisers, during the document planning stage, could follow the links between the information they intend to add to the summary, and those spots in the source dialogue which are more likely to support that information. In doing so, they would be able, for example, to backtrack and build a new plan for the document, should the dialogue fail to back it up, thereby overcoming the difficulty the lack of a deeper semantic representation for the dialogue poses to the summarisation process[3]. Depending on the summarisation method, the same mapping could also be used to mine, in the source dialogue, specific information needed to fulfill some summary template, for instance (*cf.* [5, 7, 13]).

The rest of this paper is organised as follows. Section 2 describes a set of summary acts we have defined to represent part of the summary semantics. It also describes the corpus annotated with this information, in order to build the mapping between the semantics of human-produced summaries and their source dialogues. Next, in Section 3, we determine those summary acts that were more used by human summarisers, providing an overview of their importance. Section 4, in turn, presents the resulting mapping, identifying which summary acts were linked to just a few dialogue acts, and which might be supported by a good deal of them. Finally, Section 5 presents a conclusion to this work.

[2] The reason why this is so were left unspecified in NECA's documentation.

[3] These were the actual steps followed when building our own automatic dialogue summariser, described in [16].

2 Corpus Annotation: Extending the RRL

Our first step towards building a mapping that links the content of summaries to their source in the summarised dialogues was to start determining its semantic content, by defining a set of summary acts. Based on Searle's definition of speech acts [17], a summary act represents the illocutionary act (*cf.* [1]) executed by the summariser when writing some clause in the summary, *i.e.* the act underlying that clause. In this research, clauses, *i.e.* a unit consisting, as a minimum, of a verb and its complements [9], were taken as our basic unit due to the relative objectivity of this concept.

To get to this set of summary acts, we analysed 240 summaries (1,773 clauses) from the corpus described in [14, 15], leading to a total of nine summary acts, shown on Table 1. Each clause from the corpus was then annotated with a summary act and a predicate-arguments pair (in first-order logic), representing the clause's semantic content. The construction of each pair builds on the theory presented by Parsons [10], according to which the information presented in a simple sentence – called eventuality – can be classified either as (i) a state, describing some characteristic of an object or person; something that, at a specific moment, is simply true or false (*E.g.* "He is a good seller"); (ii) an event, presenting some fact, action or even state change (*E.g.* "He grew nervous" and "The buyer walked into the shop"; or (iii) a process, *i.e.*, an event with no apparent finishing point, composed of various sub-events identical to the main one, as in "He kept on babbling about it".

Based on the above distinctions, each clause in the corpus was assigned a set of predicates, joined by the symbol '&', so that each predicate represents a single eventuality in that clause. To do so, predicates are named after the verb (either single or complex), in the case of events and processes. Along with this information, they are also assigned a list of complementary attributes, built from the remaining elements in the clause. As for states, we kept NECA's representation, naming the predicate as "attribute" and adding to its complementary list the attribute of the object or participant it deals with. This list, whenever possible, should contain information about:

- Who executed the action (for events or processes), or bears a determined attribute (for states), as in "*The vendor* politely showed the car to the customer", or "It was a very good *car*". This information is compulsory.
- To whom the action was directed, as in "The vendor politely showed the car to *the client*". This information does not apply to states, being optional for the rest (*i.e.*, it is included only if found in the clause).
- What object was involved, as in "The vendor politely showed *the car* to the client". This information does not apply to states and, depending on the situation, can be optional for the rest.
- What modifier was used, as in "The vendor *politely* showed the car to the customer". This information is optional.

Under this classification, clauses like "she (buyer) asked me (vendor) about the luggage compartment" become *ask(buyer, vendor, luggage_compartment)*,

Table 1. Summary acts defined for this research

Summary Act	*Description*
Advice	When the summariser, by writing the clause, advises the reader to do something. *E.g.* "If you ever need a new car, *you should see him*".
Evaluation	When the clause, either directly or indirectly, evaluates some product or dialogue participant; or else when it describes an evaluation by some participant. *E.g.* "Ritchie was ambiguous", "Tina did not like the car", or "It seemed to be a good car".
Auxiliary	Either used with interjections, or with clauses whose semantic content acts solely as a complement to some other clause's semantics, as in "*It was precisely* what I was looking for" (Auxiliary + Evaluation).
Closure	When the clause describes the way the dialogue finished. *E.g.* "It was a smooth sell", "Still I bought the car".
DescrSituation	When the clause summarises the entire dialogue. *E.g.* "This is a dialogue between a vendor and a buyer".
Inform	When the clause only informs about some object's characteristic. *E.g.* "The car's got ABS".
InformAction	When the clause reports on some action executed by some participant. *E.g.* "I informed the client (...)" (incomplete) and "Ritchie described some positive points of that car to Tina".
Opening	The clause describes the way the dialogue began, or the situation that happened at its beginning. *E.g.* "The buyer asked about a car" and "Tina walked into the shop".
Opinion	When the clause explicitly presents the summariser's personal opinion. *E.g.* "Although I find it strange only 100 mph for a sport car". Through this act, the summariser explicitly shows him/herself to the reader.

whereas "An argumentative client came up to the shop" becomes *come_up(buyer, shop)& argumentative(buyer)*. Even though this codification can represent with reasonable precision the meaning of a clause, it fails in capturing its emotional load (a necessary feature to deal with the data from [15]). This limitation was nevertheless addressed by adding an extra predicate to the clause's meaning, responsible for the identification of its polarity (positive or negative), along with the loaded element (either attribute or predicate). As such, clauses like "At least she (buyer) was patient" can be codified as *patient(buyer)&positive(patient)*.

3 Corpus Analysis

By analysing the density of each summary act within summaries, *i.e.*, the mean number of clauses per summary labelled with that act (Figure 1), along with the overall proportion of clauses labelled with them (*i.e.* their coverage, shown on Figure 2), we can see that some summary acts were considered more relevant than others. One such case relates to reporting on someone's behaviour, *i.e.* the action s/he performed (represented by *informAction*). This act's density was

almost twice as high as the second-ranked one, being found in almost 34% of all 1,773 clauses in our corpus, which is a good indication of the importance of such information to the summary [16]. At the other end of the scale, the least used *Acts* where those related either to overviews of the entire situation (*DescrSituation*), or to the personal opinions summarisers add to the summary, both explicitly (*opinion*) and indirectly, through some advice given to the reader (*advice*).

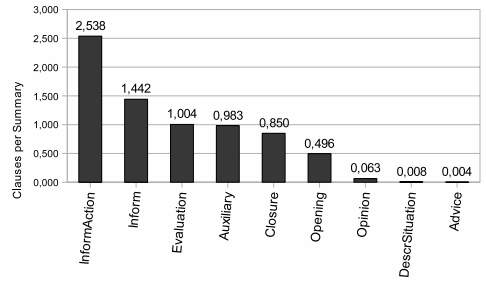

Fig. 1. Summary Acts density: Number of clauses per summary

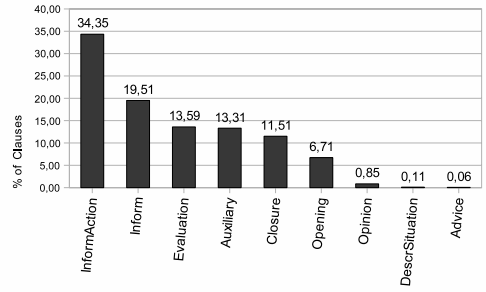

Fig. 2. Clause coverage for the Summary Acts

Interestingly, the summary act *inform*, which conveys purely technical information about the negotiated object, comes in second place in both analyses, indicating that, although not the most widely used, such information still plays, as we might expect, a prominent role in the summaries. If, however, we move to the analysis of each act's coverage amongst summaries, *i.e.* the proportion of summaries with some clause labelled with it, distribution changes considerably, specially regarding to the purely technical information (Figure 3). When comparing both clause and summary coverage, we see this information moving from second down to sixth place, being found in 32.5% of all summaries. This result indicates that, although some people might have considered it paramount, as reflected by the amount of clauses labelled so, the majority of human summarisers decided to let it out of their summaries.

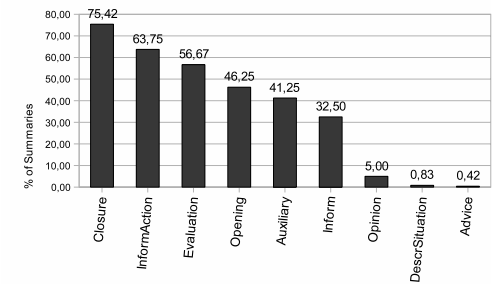

Fig. 3. Summary coverage for the Summary Acts

On the other hand, some *Acts*, such as *closure* for instance, moved from the fifth position (11.5%), in the clause coverage rank, to first amongst dialogues (75.4%). This is somewhat expected, since people are supposed to mention the way the dialogue was closed by its participants, while not bothering to talk too much about it. Similar results can be observed for *opening*, which moved from sixth (6.7%) to fourth position (46.3%). Finally, *Acts* responsible for subjective or behaviour-related information, such as some *evaluation* and *informAction*, were not so much affected by moving from clause to summary distribution (although they could still be found in over 50% of both sets).

4 Mapping Summaries to Dialogues

Once determined the semantics of all clauses in the summaries, our next step was to link those clauses to their counterparts in the source dialogue. Ideally, the automatic summariser should look directly in the source dialogue for the desired propositional content. However, as illustrated on Table 2 (taken from [8]), since over 50% of NECA's utterances present no other semantics apart from their dialogue acts, there usually is no straightforward way of automatically determining the propositional content of all utterances in the source dialogue. Hence, the only alternative we are left with is to handcraft a mapping between each summary act (from the summary clauses) and all dialogue acts (in the source dialogue) that might support that clause. Such a mapping could then be used by the automatic summariser to choose, amongst all competing propositional contents taken from the human produced summaries, the most probable for the dialogue at hand, given the summary act elected by the planner to take part in the final summary.

As expected, the resulting mapping has, as one of its characteristics, the existence of many-to-many relations, *i.e.*, situations in which each summary act can relate to many different dialogue acts (and vice-versa), as shown on Figure 4[4]. Another interesting point about this mapping is the considerable

[4] *Auxiliary* was left out of this figure because it is not directly related to the summary, but the summariser instead.

Table 2. Dialogue acts in NECA

Dialogue Act	*Context*	*Contains Semantics*
agree	The vendor agrees with some negative comment from the buyer	sometimes
initiateClosingNegative	The buyer rejects the car, closing the conversation	no
completeClosingNegative	Vendor's answer to initiateClosingNegative	no
initiateClosingPositive	The buyer takes the car, closing the dialogue	no
completeClosingPositive	Vendor's answer to initiateClosingPositive	no
confirm	The vendor confirms some yes/no question	no
disagree	The vendor disagrees on the buyer's negative remark	no
disconfirm	The vendor denies some yes/no question	no
feedback	Client's reaction, in the absence of a positiveResponse or negativeResponse	no
greeting	The vendor's greetings to the customer at the beginning of the dialogue	no
inform	The vendor informs the buyer about some attribute	yes
negativeEvaluation	The vendor informs the buyer about some negative feature	yes
negativeResponse	The buyer makes a negative remark on some attribute	yes
openingComplaint	The buyer starts the dialogue with a complaint	no
openingComplaintResponse	Vendor's answer to an openingComplaint	no
openingQuestion	The vendor poses a question to start the conversation	no
openingResponse	Buyer's answer to openingQuestion	yes
positiveEvaluation	The vendor informs the client about some positive attribute	yes
positiveResponse	The buyer makes a positive remark on some attribute	sometimes
refuseAnswer	The vendor refuses to answer some question	no
refuseAnswerResponse	Buyer's answer to refuseAnswer	no
requestIf	The buyer poses a yes/no question	yes

Table 2. *(continued)*

Dialogue Act	*Context*	*Contains Semantics*
requestInfo	The buyer's first question	yes
requestValue	The buyer asks for some specific attribute	yes

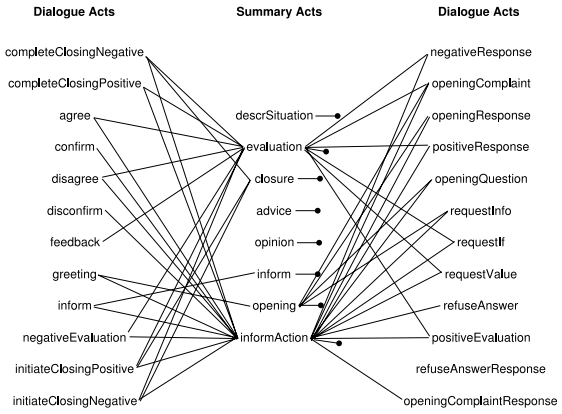

Fig. 4. Mapping between summary acts and dialogue acts

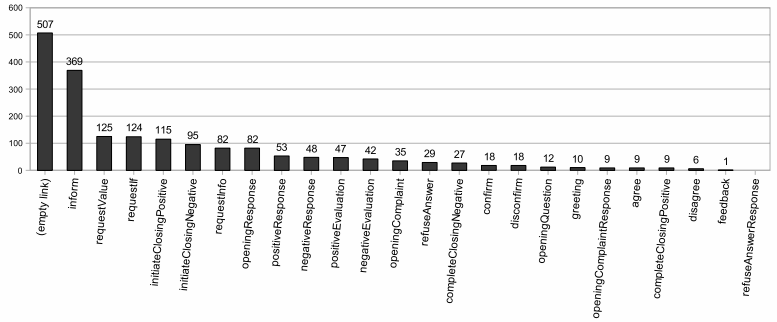

Fig. 5. Distribution of links amongst targeted Dialogue Acts

amount of summary acts with no direct link to some dialogue act (around 27% of all links), represented by the symbol —• in the figure. In this case, the absence of such a link can be interpreted not as the inclusion of information from outside the dialogue, but instead, as the inclusion of information that, although present in the dialogue, has no explicitly defined semantic content within its RRL representation, *i.e.*, it neither has a predicate-arguments pair nor

a dialogue act assigned to. At the other end, we have refuseAnswerResponse with no link whatsoever, meaning that no human summary reported on its content. Figure 5 summarises these results.

5 Conclusion

In this paper we described the measures taken to use RRL in the process of dialogue summarisation. Such measures included (i) the identification and definition, from experimental data, of a set of summary acts used to annotate these data with their semantic meaning; and (ii) the mapping between these summary acts and the *Dialogue Acts* produced by the NECA platform. Differently from current literature (*e.g.* [2, 3]), however, our approach works at the semantic level only, instead of relying on syntactical features.

With this mapping at hand, automatic summarisers could identify, at the document planning stage, those portions in the dialogue with higher probability to support some argument, or even to gather, in the source dialogue, specific information to fulfill some predefined templates (as it was already done in [16], with very encouraging results). As an additional benefit, the method can also be used to draw some conclusions, from experimental data, about the importance of certain summary acts within specific domains, thereby influencing the way automatic summarisers choose (or learn to choose) their output summary content.

Avenues for future work include determining to what extent is our inventory of summary acts specific to the car sales domain. It would be also interesting testing this approach on different domains, such as NECA's *Socialite* scenario, for example, extending the already defined set of *Summary Acts* if necessary. Finally, since the summary acts were determined exclusively by one of the researchers, it is necessary to test them with other people, comparing the obtained results.

References

1. Austin, J.L.: How to do Things with Words, 2nd edn. Oxford University Press, Oxford (1975)
2. Banko, M., Mittal, V., Kantrowitz, M., Goldstein, J.: Generating extraction-based summaries from hand-written summaries by aligning text spans (pacling 1999). In: Proceedings of the Pacific Association for Computational Linguistics, Waterloo, Canada (1999)
3. Daum-III, H., Marcu, D.: Induction of word and phrase alignments for automatic document summarization. Computational Linguistics 31(4), 505–530 (2005)
4. van Deemter, K., Krenn, B., Piwek, P., Klesen, M., Schröder, M., Baumann, S.: Fully generated scripted dialogue for embodied agents. Artificial Intelligence 172(10) (June 2008)
5. Kameyama, M., Kawai, G., Arima, I.: A real-time system for summarizing human-human spontaneous spoken dialogues. In: Proceedings of the 4th International Conference on Spoken Language (ICSLP 1996), Philadelphia, USA, vol. 2, pp. 681–684 (1996)

6. Kamp, H., Reyle, U.: From Discourse to Logic: Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory. Springer, Heidelberg (1993)
7. Kipp, M., Alexandersson, J., Reithinger, N.: Understanding spontaneous negotiation dialogue. In: Proceedings of the IJCAI 1999 Workshop on Knowledge and Reasoning in Practical Dialogue Systems, Stockholm, Sweden, pp. 57–64 (1999)
8. Klesen, M.: Dialogue act specification for the neca domains. Manuscript, DFKI (June 2002)
9. Miller, J.: An Introduction to English Syntax. Edinburgh University Press Ltd., Edinburgh (2002) ISBN 0 7486 1254 8
10. Parsons, T.: Events in the Semantics of English: A Study in Subatomic Semantics. MIT Press, Cambridge (January 1994) ISBN 0-262-66093-8
11. Piwek, P.: NECA deliverable D3a: Specification of scene descriptions for the neca domains. Tech. rep., ITRI – University of Brighton, Brighton, UK (2002); nECA IST-2000-28580 Deliverable D3a
12. Piwek, P., Krenn, B., Schröder, M., Grice, M., Baumann, S., Pirker, H.: RRL: A Rich Representation Language for the description of agent behaviour in NECA. In: Proceedings of the AAMAS Workshop on Embodied Conversational Agents - Let's Specify and Evaluate them!, Bologna, Italy (2002)
13. Reithinger, N., Kipp, M., Engel, R., Alexandersson, J.: Summarizing multilingual spoken negotiation dialogues. In: Proceedings of the 38th Annual Meeting on Association for Computational Linguistics (ACL 2000), Hong Kong, China, pp. 310–317 (2000)
14. Roman, N.T., Piwek, P., Carvalho, A.M.B.R.: Politeness and summarization: an exploratory study. In: Qu, Y., Shanahan, J., Wiebe, J. (eds.) Papers from the 2004 AAAI Spring Symposium: Exploring Attitude and Affect in Text: Theories and Applications (AAAI-EAAT 2004), March 22-24, pp. 112–117. The AAAI Press, Palo Alto (2004)
15. Roman, N.T., Piwek, P., Carvalho, A.M.B.R.: Politeness and Bias in Dialogue Summarization: Two Exploratory Studies. In: Computing Attitude and Affect in Text: Theory and Applications. The Information Retrieval Series, January 9, vol. 20, pp. 171–185. Springer, Netherlands (2006)
16. Roman, N.T., Piwek, P., Carvalho, A.M.B.R.: Emotion and behaviour in automatic dialogue summarisation. In: Proceedings of the Sixth Workshop in Information and Human Language Technology (TIL 2008), Vila Velha, Brazil, October 27-28, vol. 2, pp. 304–309 (2008)
17. Searle, J.: Speech Acts: An Essay in the Philosophy of Language. Cambridge University Press, Cambridge (1969)