

On Three Notions of Grounding of Artificial Dialog Companions

Andy Lücking and Alexander Mehler
Goethe University Frankfurt

Appears in:

Science, Technology & Innovation Studies

exp. Oct. 2013

Draft - do not copy or distribute without written permission of the author.

Abstract

We provide a new, theoretically motivated evaluation grid for assessing the conversational achievements of Artificial Dialog Companions (ADCs). The grid is spanned along three grounding problems. Firstly, it is argued that *symbol grounding* in general has to be intrinsic. Current approaches in this context, however, are limited to a certain kind of expression that can be grounded in this way. Secondly, we identify three requirements for *conversational grounding*, the process leading to mutual understanding. Finally, we sketch a test case for symbol grounding in the form of the *philosophical grounding problem* that involves the use of modal language. Together, the three grounding problems provide a grid that allows us to assess ADCs' dialogical performances and to pinpoint future developments on these grounds.

1 Object, Aim and Research Questions

This paper deals with embodied conversational agents (Cavazza et al., 2010) as potential interlocutors of human users (Wachsmuth, 2008; Wilks, 2005, 2007, 2009; Wilks et al., 2010). In the literature, there are a lot of names and acronyms for these kinds of systems. Candidate designations include *Artificial Companions* (Wilks, 2005),¹ *Artificial Conversational Companions* (Danilava, Busemann, and Schommer, 2012), *Embodied Conversational Agents* (Cassell, 2001), *Dialogue Agents* (Wilks, 2009), *Conversational Agents* (Kopp and Wachsmuth, 2004), and *Dialog Companions* (Wilks, 2005). We focus on those systems that are able to communicate with human users by means of a natural language. We concentrate on the linguistic facilities of those systems and abstract over issues of anthropomorphic design or ethics of behavior - that is, we stress their dialogue aspect over their companions aspect. Throughout this paper, we call such agents *Artificial Dialog Companions* (or simply ADCs).

The aim of ADCs is to provide long-term companions that accompany their human users in a way that they learn the habits, interests and cognitive states of their users in order to better meet, for example, their conversational needs. The operational scenarios of ADCs range from task-oriented dialogs to free conversation (Cavazza et al., 2010; Wachsmuth, 2008; Wilks, 2005). Building on some adaptable knowledge resource (based, for example, on Wikipedia (Gabrilovich and Markovitch, 2009; Waltinger, Breuing, and Wachsmuth, 2011)), some inference mechanism (build-

ing, for example, on semantic-web technologies (Wilks et al., 2010)) and some dialog management system (Traum and Larson, 2003), ADCs process and generate data to keep track of the conversation with their human interlocutors (Gilroy, Porteous, Charles, and Cavazza, 2012; Salem, Kopp, Wachsmuth, Rohlfing, and Joublin, 2012; Wachsmuth and Knoblich, 2005). The data processed by ADCs comprise a wide range of data that includes verbal, linguistic data as well as multimodal sensory input. Currently, models of ADCs are under research that are said to allow even for the emotional control and reflection of their conversations (Rehm, André, and Nakano, 2009).

In this paper, we discuss possible limits of the conversational behaviour of ADCs partly in an abstract, partly in an exemplary manner. We deal with scenarios under which the conversation of an ADC with a human user can be said to be unnatural, dysfluent or even unsuccessful. From the point of view of cognitive science, limits of this sort are affected by what an ADC can *intrinsically* learn without being *extrinsically* pre-programmed by its human designer (Ziemke, 1999). In this line of reasoning, we view language learning as being critical for the acceptability of an ADC as it affects the flexibility of its conversational behavior. In order to analyze the conversational flexibility of ADCs with regard to the dynamics of natural language conversations, we consider three notions of grounding that relate to different conversational abilities of ADCs:

1. Starting with the notion of *grounding in AI* (Harnad, 1990), we consider the possibilities of an intrinsic semantics that goes beyond intersective predicates, which are anchored in perceptual experience. From this point of view, we discuss the requirement that ADCs should be able to answer questions about factual states of the world as, for example, *“What is the temperature outside?”*
2. Utilizing the notion of grounding in dialog theory (Clark, 1996), we discuss the flexibility of the conversational behavior of ADCs beyond managing typical speech acts and adjacency pairs (Sacks, Schegloff, and Jefferson, 1974; Searle, 1971). From this point of view, we ask for the ability of ADCs to manage states of informational uncertainty of dialog acts, for example, by means of clarification requests of the sort *“Whom do you mean by Hans?”*
3. Finally, referring to the notion of grounding in philosophy, we discuss the need of an *intensional semantics* (Montague, 1974) to be intrinsically learnt by an ADC. From this point of view, we ask for ADCs that can answer questions about *possible* states of the world as exemplified by the question *“What would you recommend: what I shall do if two of my friends would have the same birthday?”*

Based on these three notions of grounding, we argue that ADCs are limited with regard to their categorical (1), conversational (2) and intensional (3) grounding. As a result of these constraints, we state that, currently, ADCs cannot converse with human interlocutors to a degree that is natural for a conversation with a human being. In a nutshell: we argue that ADCs do not yet function as interlocutors – currently, they are not sufficiently equipped to be called *dialog* companions.

Irrespective of this assessment, we are very sympathetic with the highly ambitious approach that underlies ADCs. There are many possible application areas in which ADCs can help (e.g., in supporting caregiving or everyday tasks). Smart HCI systems of this sort are partly an object of our own research (Mehler and Lücking, 2012). However, we are also convinced that ADCs cannot be applied usefully unless they are able to communicate on a near-human level. This is not only due to security reasons (which are of highest importance, e.g., in the context of caregiving), but also to possible frustration as a result of insufficient interaction and understanding. In order to get a better estimation of the achievements and potentials of ADCs, we describe some “milestones” in terms of the grounding problem that full-blown ADCs should have mastered. These grounding steps make an (incomplete) grid that may accompany or even replace costly user evaluation studies.

The paper is organized as follows: Section 2 sketches three notions of grounding according to the accent of their academic provenance: grounding in terms of AI, dialog theory and philosophy. Sections 3, 4 and 5 utilize these notions to successively specify requirements with regard to the conversational capabilities of ADCs. In this context, Section 3 analyzes the limits of categorization games as a model of learning an intrinsic semantics on the part of ADCs. Section 6 sums up our findings in assessing the conversational interactivity of up-to-date technologies of ADCs.

2 Three Notions of Grounding

Dialogical communication on the side of ADCs involves at least two dimensions of meaning:

- The symbols used in conversations have a meaning that is known to the ADC. We call this the *symbol* dimension. The key problem here is how agents acquire an *intrinsic semantics* (Harnad, 1990). Generally speaking, the semantics of an artificial agent is said to be *extrinsic* if the meanings of the signs that it uses are externally determined by its designer. In contrast to this, the semantics is said to be internal to the agent, that is, *intrinsic* if it generates the mapping of sign vehicles and meanings independently of its designer.ⁱⁱ

- Within a dialogical exchange, symbols are used and acknowledged according to certain exchange rules. This pertains to the *interaction* dimension of dialog. Key issues here are turn-taking and ensuring mutual understanding.

In order for a system to be a *dialog* companion, it has to master both the symbol and the interaction dimension. We identify three grounding problems that allow us to assess an ADC's achievements on these dimensions. Each grounding problem is exemplified by a paradigmatic question.

GP_{symb} Grounding Problem_(symbols). The grounding problem in AI, robotics and technical systems dealing with language in general has been defined by Harnad (1990, p. 335) as follows: "How can the semantic interpretation of a formal symbol system be made *intrinsic* to the system, rather than just parasitic on the meanings in our heads?" (emphasis in original). ADCs that have mastered GP_{symb} can answer a question like "What are you seeing (right now)?"

GP_{conv} Grounding Problem_(conversation). Every act of speaking presupposes information - background knowledge shared by conversational participants (Stalnaker, 1978, 2002; Lewis, 1969; Schiffer, 1972). This background knowledge is often termed *common ground* and is a core component of any theory of language use (Clark, 1992). The linguistic grounding problem consists in spelling out what information is part of common ground, how it is represented, and how it is maintained and updated in the course of conversation. Conversational grounding enables ADCs to talk about mutually known persons, amongst others, for example answering a question like "Have you seen Maynard recently?"

GP_{mod} Grounding Problem_(modality). In philosophy, the grounding problem originates from material coincidence, for instance, a statue of Goliath and the lump of clay it is made of sharing a spatio-temporal portion of the world (Gibbard, 1975). Now we can ask: "If the statue gets destroyed, will the lump of clay still exist?" If the answer is *yes*, then both the statue and the lump of clay differ in at least one modal property, from which follows, that the statue and the lump of clay are not identical. The philosophical puzzle now is how it can be that two different objects can occupy the same spatial region at the same time. However that may be, the question exemplifies that people do not only talk about factual events or currently perceived scenes, but also about possible or future events. How would an ADC answer such a question? The key problem here is that an ADC has to be able to process counterfactuals and modality in order to understand or formulate the question. Dealing with counterfactual conditionals and grammatical mood is part and parcel of the GP_{mod}. These topics are bound up with philosophical work on, amongst others, modal logic, temporality, necessity, and causation and situational

regularities (Reichenbach, 1947; Lewis, 1973b,a; Kripke, 1980; Prior, 1967; Montague, 1974; Vendler, 1957; Barwise, 1989, Chap. 5), which in turn make up the backbone of respective linguistic modeling (e.g., Dowty, 1979; Parsons, 1994; Kamp and Reyle, 1993; Krifka, 1992). Thus, the philosophical grounding problem of the statue and the lump of clay is used as an example case for modal speech, which for this reason is referred to as the grounding problem of modality in this paper.

GP_{symb} and GP_{mod} pertain to the symbol dimension of dialogs. They both focus on intrinsic meaning constitution of ADCs. In this context, GP_{symb} denotes a minimal requirement of symbolic grounding, whereas GP_{mod} highlights an advanced level. GP_{conv} , on the other hand, focuses on the interaction dimension. Conversational grounding is a complex process that, if successful, leads to dialogic understanding.

GP_{conv} and GP_{symb} affect the speech handling of ADCs directly: the former, for it relates to the dialog management of the ADC, the latter, for it concerns how agents are able to share intrinsic semantics in the first place. ADCs cannot ponder the philosophical grounding problem before they have mastered the other two. Agents, however, that have acquired synonyms within their lexicon in the course of a language game (cf. e.g. Baronchelli, Loreto, and Steels, 2008) should be able to question whether there holds indeed an identity relation between the referents of the synonymous expression by reflecting, *inter alia*, the spatial, temporal, and modal properties of these referents.

We want to emphasize that we do not claim that the three grounding aspects or the two meaning dimensions distinguished above are independent from another. The opposite is true: grounding modal speech is a special case of the general symbol grounding problem (cf. Lücking and Mehler, 2011, p. 30), and symbol grounding depends on conversationally interacting agents (Lewis, 1969; Puglisi, Baronchelli, and Loreto, 2008). However, notwithstanding the interrelationships that may hold between GP_{symb} , GP_{conv} and GP_{mod} , they have different foci that should not be confused in discussing achievements and requirements of ADCs.

Note further, that we do not take the three grounding aspects to be an exhaustive list of grounding phenomena in the context of dialog companions. The grounding problems identified above are confined to verbal speech, ignoring, for instance any nonverbal or social properties of ADCs.ⁱⁱⁱ A common feature of our grounding problems is, however, that they are standardly labeled as “grounding” and therefore can potentially give rise to confusion, if not properly kept apart.

3 ADCs and GP_{symb}

Starting from the notion of grounding in terms of GP_{symb} , our basic argument with regard to the limits of the conversational flexibility of ADCs can be summarized as follows:

1. **Limited interactivity as a result of insufficient grounding:** At present, ADCs implement an extrinsic semantics (see above at beginning of Section 2). This means that the semantics of their conversational items is mainly predefined and prescribed by the system designer. As a result, ADCs have a limited learning capacity. Because of this limitation, ADCs are not sufficiently interactive in terms of a natural conversational interaction among human interlocutors (Brennan, 1998). ADCs with such a limited capacity of *artificial interactivity*^{iv} may have problems with regard to their acceptability as interlocutors of human users.
2. **Grounding ADCs with the help of evolutionary Models of Language Evolution (MoLE):** A possible way out of this problem starts with the notion of grounding in AI (Cangelosi, Greco, and Harnad, 2002; Steels, 2008; Ziemke, 1999). In line with this, we think of ADCs that interact with their environment in an intrinsic manner such that their behavior-generating patterns are not prescribed by the system designer.^v Such systems may be flexible enough so that they successfully “hide” their artificiality from the point of view of their human users. To achieve this goal we need an approach that endows ADCs with a learning capacity that enables them to intrinsically acquire a semantic to a degree that they solve the GP_{symb} . Since Ziemke (1999) has already shown the limits of the cognitivist approach (Fodor, 1997; Fodor and McLaughlin, 1995) and of the enactive approach (Varela, Thompson, and Rosch, 1991) to grounding in AI, an alternative approach is needed. Such an approach exists in terms of the paradigm of language evolution (cf. Steels, 2008, 2011): “[...] the most promising path toward successful synthesis/modeling of fully grounded and truly intelligent agents, will probably be what might be called ‘evolutionary and developmental situated robotics’, i.e. the study of embodied agents/species developing robotic intelligence bottom-up in interaction with their environment, and possibly on top of that a ‘mind’ and ‘higher-level’ cognitive capacities.” (Ziemke, 1999, p. 187). In line with this approach, we may think of ADCs that *intrinsically learn* the semantics of conversational items by interacting with human users or some other artificial interlocutors in order to evolve a common language that is not prescribed to them.

3. **Limits of MoLE as a means of grounding ADCs:** Notwithstanding the attractiveness of MoLE, this approach has limits with regard to the task under consideration. To simplify our argument, we focus on learning a semantics beyond the level of intersective predicates (see below) in the framework of the predominant model of evolutionary semantics, that is, the *Categorization Game* (CG) (Baronchelli, Gong, Puglisi, and Loreto, 2010; Puglisi, Baronchelli, and Loreto, 2008; Vogt, 2005).^{vi} Starting from Lücking and Mehler (2012), we briefly recapitulate that the CG is limited in that it does not go beyond learning the semantics of intersective predicates. As a result of this recapitulation, we state that the CG needs to be extended before it can be considered an alternative to solving the GP_{symb} . In any event, our diagnosis is that, presently, the CG is not expressive enough to provide an intrinsic semantics for ADCs and, therefore, limits their conversational competence.

In what follows, we substantiate this argumentation scheme. The GP_{symb} , that has been formulated in terms of the *Symbol Grounding Problem* (SGP) by (Harnad, 1990), tackles the possibility of an *intrinsic* semantics (see above) for AI applications. Solving the SGP or, equivalently, the GP_{symb} , means meeting the requirement of autonomy of interpretation on the part of the artificial agent. Any model that claims to solve the SGP has to explain at least three phenomena (Harnad, 1990):

1. Firstly, it has to explain how sensory input is projected onto corresponding *iconic representations*.
2. Secondly, it has to explain how *categorical representations* are learnt from iconic representations, for example, by means of identifying invariant features in the sensory projections.
3. Finally, it has to explain how atomic *symbolic representations* are learnt as names for categorical representations (i.e., statements of class membership) according to the detection of invariant features. This includes an account of the organization of atomic symbols into taxonomies and their combination into complex symbolic representations, for example, by means of logical connectives (“and”, “or”, “not”, “all”, and so on).

In a nutshell: symbols are said to be groundable if they can be traced back to something perceptible in the sense of this enumeration.

Since the time of the formulation of the SGP, much successful and seminal work has been done on letting agents learn an intrinsic semantics, most prominently within the *Naming Game* paradigm and its extension in terms of the *Categorization Game* (Baronchelli, Loreto, and Steels, 2008; Steels, 1996). This work

has been convincing to such an extent that Steels (2008) stated that “[t]he Symbol Grounding Problem has been solved” for “groundable symbols” (Steels, 2008, p. 223) in the sense that “[t]here is no human prior design to supply the symbols or their semantics, neither by direct programming nor by supervised learning.” (Steels, 2008, p. 239). Steels (2008, p. 239) clarifies this notion of an intrinsic semantics by claiming that “[e]ach agent builds up a semiotic network relating sensations and sensory experiences to perceptually grounded categories and symbols for these categories.”

In order to provide a pretest of this statement, consider an attribute-noun construction like “slow slug”. A term like “slug” is certainly groundable in the sense of the GP_{symb} (cf. work on pattern matching and classification as reviewed in Tenenbaum *et al.*, 2011). But what is with “slow”? One reading of this adjective refers to a perceptual magnitude, namely distance per time unit. Obviously, there is no fixed magnitude that makes up the perceptual counterpart of “slow”. Rather, the semantics of “slow” is contextsensitive in the sense that it is calibrated (Kamp and Partee, 1995) in the context of its argument, that is, the head noun that it modifies: the speed of a slow slug differs, for example from the speed of a slow hunting-leopard such that both cannot be said to belong to the same class of slow animals (for related examples see Lahav 1989). Obviously, the meaning of an adjective like “slow” is open in the sense that it is non-trivially affected by its usage context (Hörmann, 1983). In terms of the SGP, there is neither a simple perceptually grounded representation of “slow” nor a compositional representation on the symbolic level.

This example recapitulates the data basis that has been used by Lücking and Mehler (2012) to show that the semantic expressivity of the current version of the CG is limited by an intersective semantics.^{vii} According to such a semantics, the meaning of an attribute-noun construction is the intersection of the meanings of its constituents - disregarding any kind of context-sensitive calibration. In other words, we state that the CG does not yet implement more complex cases of context-sensitive meaning calibration as described, for example, by Kamp and Partee (1995). Thus, the CG as the predominant model of the evolution of natural language semantics is restricted with regard to the semantic complexity of the predicates it can deal with - below the level of the semantics of a natural language. As a corollary, we state that this restriction is *extrinsic* in the sense that it is *prescribed by the designer* of the CG. This prescription is a consequence of the way the designer defines single rounds of a CG, the underlying meaning space and the way artificial agents can generate new signs. In a nutshell: CGs extrinsically restrict the semantics that artificial agents can learn as part of a CG. Thus, CGs do not yet provide grounding in the desired way, that is, in terms of the GP_{symb} . Note that this assessment does not imply that CGs im-

plement a sort of supervised learning. Rather, we say that the current implementation of CGs is supervised on a higher level on which it prescribes semantic expressivity.

At this point, one may object that the naming and the categorization game have been said to solve the grounding problem for *groundable* predicates whose semantics is anchored in perceivable objects or processes (Steels, 2008). However, as our example of “slow” shows: even predicates that are assumed to be groundable in this sense can be affected by a context-sensitive semantics. Suppose in contrast to this assessment that “slow” has an intersective semantics so that “slow slug” denotes the intersection of all perceivable objects that are said to be slow and all perceivable objects that are categorized as slugs. In order to learn such a semantics, an ADC would need to learn the meaning m of “slow”, subject to its different usage contexts so that m turns out to be the union of all result sets of all these context-sensitive meaning constitutions. It is this that we do not see in current implementations of the CG and what is more intuitively represented in terms of a substantive semantics where the meaning of “slow slug” is learnt, resulting in a subset of the meaning of “slug”. Under this regime, an ADC never needs to represent the meaning of “slow” as something that is the union of all things that are said to be slow – there is no need for such a representation. Rather, the ADC just needs to learn how to apply the attribute “slow” as an operator to the meanings of its arguments (that operates in a certain quality dimension in the sense of Gärdenfors 2000).

In line with this argument, we also question the status of semantic networks in the CG (see above): CGs implement many-to-many relations between sign vehicles and their denotations where syntagmatic and paradigmatic relations of signs are mapped insofar as they provide a compositional semantics (Vogt, 2005). The meaning relation between sign vehicles and their denotations can be seen to span a bipartite graph (Newman, 2010). Any such graph induces a neighborhood graph, for example, on the side of the sign vehicles such that vehicles that are related to the same or similar denotations, are interlinked. This allows us to account for, for example, relations of (partial) synonymy. It is obvious how to derive more complex semantic relations (e.g., hyperonymy or co-hyponymy) based on this representation format – see Loreto, Mukherjee, and Tria (2012) for an example of this research branch. However, in many implementations of the CG, this relational network of signs does not play a role as a dependent variable, that is, as a possible outcome of the CG. In this sense, we do not see how the present version of the CG generally provides a model that allows for learning both a sign-meaning relation on the one hand and a semantic network (Mehler, 2008; Steyvers and Tenenbaum, 2005) on the other.

Based on this argument we conclude that the GP_{symb} has not been completely solved.^{viii} As we are convinced that CGs provide a partial solution to the GP_{symb} , we need to specify this part in more detail. This can be done with the help of Coradeschi and Saffiotti (2003, p. 85), who introduce the *anchoring problem* as the “problem of connecting, inside an artificial system, symbols and sensor data that refer to the same physical objects in the external world.” From our point of view, this part of the GP_{symb} has been solved by the CG and related approaches. However, “[s]ymbol grounding” as Coradeschi and Saffiotti (2003, p. 93) continue, “is a more general problem than anchoring. It concerns the philosophical issues related to the meaning of symbols in general.”

We do not claim that the the CG fails to offer a solution for GP_{symb} *in principle*. Rather, we tried to show that currently the CG does not account for the full range of semantic classes of natural language predicates as systematized, for example, by Kamp and Partee (1995). Respective enhancements are necessary in order to endow ADCs with the desired learning capacity.

4 ADCs and GP_{conv}

Communication between two (or more) interlocutors is a coordinated activity and a joint achievement (Clark, 1992).^{ix} For instance, even an apparently simple question like “Have you seen Maynard recently?” can only be answered by the addressee if he knows who Maynard is. In other words, both dialog partners are required to have mutual knowledge of a certain person named Maynard. Furthermore, as communication proceeds, the dialog contributions cannot simply be taken for granted – contributions may fail at various levels, as pointed out by Clark and Schaefer (1987, 1989). Given the example question from above (“Have you seen Maynard recently?”), possible reactions include:

“Huh?” (*I didn’t hear what you said.* – form aspect),

“Maynard?” (*Who are you talking about?* – meaning aspect), or

“recently?” (*‘Recently’ is the wrong word, I haven’t seen him for years.* – meta-communicative aspect)

Note that (failed) grounding may concern the whole utterance or any part of it (Ginzburg, 2012; Poesio and Rieser, 2010). Thus, in communication an utterance – as locution as well as illocution or perlocution (Austin, 1962) – cannot simply be added to the dialog fact sheet; rather, it has to be *acknowledged* first, or exposed to *clarification* or even to *repair*, whenever this is necessary. This mutual process of dialog management that is performed by interlocutors by alternatingly contributing communication events and giving feedback is known as *grounding*. The conversational events that have been acknowledged or presupposed make up the so-called *common ground* (Stalnaker, 2002).

Conversational grounding has to be seen as a *sine qua non* for the dialog management module of ADCs, since “[m]any of the errors that occur in human-computer interaction can be explained as failures of *grounding*, in which users and systems lack enough evidence to coordinate their distinct knowledge states.” (Brennan, 1998, p. 201) Accordingly, the GP_{conv} can be formulated as follows: *How can ADCs keep track of grounding in user interactions with their human interlocutors?* If an ADC is not able to master the linguistic grounding problem, successful conversation with this ADC will not be possible, because grounding errors *block* mutual understanding. From the viewpoint of a requirement analysis for ADCs Danilava, Busemann, and Schommer (2012) conclude: “The interaction with an ACC [Artificial Conversational Companion] cannot be modelled as just a simple stimulus-response based exchange of utterances.”

In order to evaluate ADCs in terms of GP_{conv} , we can give the following requirements specification:

- Processing of contributions has to be incremental (Schlangen and Skantze, 2011), since elements from single words to whole sentences can be subject to acknowledgment, clarification or repair.
- ADCs have to deal with contributions that do not project onto full sentences – so called *non-sentential utterances* (Fernández and Ginzburg, 2002).
- ADCs have to keep track of the form, the meaning and the meta-communicative function of contributions, since interlocutors can inquire about these features for any conversational element (cf. the Maynard example above).

How do ADCs perform in comparison to these requirements of GP_{conv} ? The first thing to note is that the dialog systems used in constructing an ADC have turn management and dialog act tagging at their disposal (see the overview given in Wilks, Catizone, Worgan, and Turunen, 2011). Since dialog acts are related to the conversational and pragmatic role of turns and, furthermore, ADCs are equipped with models for the meaning of those turns (see e.g. Catizone, Dingli, Pinto, and Wilks, 2008), ADCs can be said to fulfil a great deal of the last-mentioned criterion.^x We haven’t found explicit, written evidence, however, whether the ADCs’ dialog modules provide a retrievable representation of the *form* of an utterance. Such locutionary information is needed, for example, to handle form-related clarifications like “Did you say ‘Maynard’? Did I hear it correctly?”.

As regards non-sentential utterances, ADCs seem to be able to handle at least short answers (cf. the example SC: “*When was this photo taken?*”, R: “*last year*” of Wilks et al., 2011, p. 142). However, there are various kinds of non-sentential utterances (Ginzburg, 2012, pp. 219 sq., distinguishes 15 classes of non-sentential utterances). To our knowledge, ADCs are not able, for

instance, to process a meta-communicatively used reprise fragment like “10 euros?” as a response to “This costs 10 euros.” or perspective takeovers (for example, personal pronoun adjustments like A: “You should do this”, B: “Me?”.). As far as one can get from the literature, ADCs probably can handle only such non-sentential utterances whose “missing parts”^{xi} can be filled with recourse to dialog act structure (such as *Question-Response* adjacency pairs (Sacks, Schegloff, and Jefferson, 1974)). In sum, the processing of non-sentential utterances seems to fall behind their elaborate manners of use in human-human conversation.

The “normal scenario” of HCI is as follows: “ECA talks, then there is a pause, then user talks” (Crook, Smith, Cavazza, Pulman, Moore, and Boye, 2010, p. 30). Additionally, backchannel signals are allowed during speech. Under certain conditions (e.g., talking duration and loudness of interjection), overlapping speech is treated as an interruption (Crook, Smith, Cavazza, Pulman, Moore, and Boye, 2010). Interruptions, however, are treated on the level of whole turns: after an interruption of a turn has been identified and processed, the system has to decide whether to “continue, replan [or] abort” the turn (Crook, Smith, Cavazza, Pulman, Moore, and Boye, 2010, p. 30). This decision is “very challenging” (Crook, Smith, Cavazza, Pulman, Moore, and Boye, 2010, p. 31), partly due to the not yet achieved processing need that “the interrupting utterance must to be considered in the context of the ECA utterance that provoked the interruption” (Crook, Smith, Cavazza, Pulman, Moore, and Boye, 2010, p. 32). Since interruptions can occur *at any given point in dialog*, an incrementally growing semantic representation is needed. Any increment reached at some point *t* in a conversation can be acknowledged or put to clarification or repair, and that in fact on the form, the meaning, or the meta-communicative level (cf. above).^{xii}

In formal dialog theory, incrementality and the semantics of discourse is a chief issue in the framework of Poesio, Traum and Rieser (PTT, Poesio and Traum, 1997; Poesio and Rieser, 2010). To our knowledge, there is no PTT implementation yet. Actually, incremental construction of dialog representations appears to be a very recent topic; we know of three approaches (namely Peldszus and Schlangen, 2012; Purver, Eshghi, and Hough, 2011; Visser, Traum, DeVault, and Akker, 2012). Since none of these approaches seems to be employed within an ADC as discussed here, the first-given requirement, incrementality, is probably not yet fulfilled. Our diagnosis is supported by work on grounding in human-computer interaction: Peltason, Rieser, Wachsmuth and Wrede (2013, p. 116) report that “[t]he robot does not KNOW turn taking rules, so it cannot project (anticipate) sequences in the CA [Conversation Analysis] sense.” (emphasis in original).^{xiii}

5 ADCs and GP_{mod}

We think that GP_{mod} and the philosophical grounding problem provide a neat test case for language grounding in AI systems. The reason is the following: agents eventually learn synonyms, that is, two different names that refer to the same thing (say, “statue” and “lump of clay”). Synonymy relations can change in the course of language learning. However, such changes are due to broadening or narrowing the perceptual categories associated with these names – supposing they are groundable in terms of Steels (2008). Consequently, agents can learn that two terms are synonymous (or not) by experience, which is perfectly in line with the notion of symbol grounding. The intrinsic semantics of ADCs at present is factual: meaning is triggered by perception (as in the Naming Game paradigm Steels, 1996) or by information retrieval (as in the Companions project Catizone, Dingli, Pinto, and Wilks, 2008). The content of conversations is always tied to sensoric representations (anchoring, cf. above) or to the facts in a knowledge base. Such systems are able to draw inferences (again, see Catizone, Dingli, Pinto, and Wilks, 2008) of the form “If X is the case and Y is the case, then Z holds.”, where X, Y and Z denote content available through the resource (i.e., perception or knowledge base).

Part of mastering language, however, is to be able to talk not only about factual events, but also about events from the past or the future, or events that might be the case. Once a semantics has been acquired for a given symbol s , then s can also be used independently of its external source (be it perception or knowledge base), that is, without immediate factual underpinning. In addition to factual speech, *modal* speech also becomes possible. This kind of language ability is asked for when one wants to discuss modal properties of things, as is done in the context of the philosophical grounding puzzle. ADCs that are said to have acquired an intrinsic semantics should be able to perform counterfactual speech of the following form: “If X would be the case, then Y ”.

The interesting observation of the philosophical grounding problem and GP_{mod} is that modal speech requires a use of symbols that is detached from its factual anchors and grounding sources. For instance, the use of “destroy” in a question like “If I would destroy the statue, would the lump of clay still exist?” does not refer to a factual event; rather, the event talked about is shifted into the realms of possibility by the conjunctive mood of “would”. Symbol use that is independent from external triggers in this sense can be called “intrinsic” properly. On this account, GP_{mod} provides a test case for assessing whether an ADC has acquired an intrinsic semantics even in the strong, modal sense.

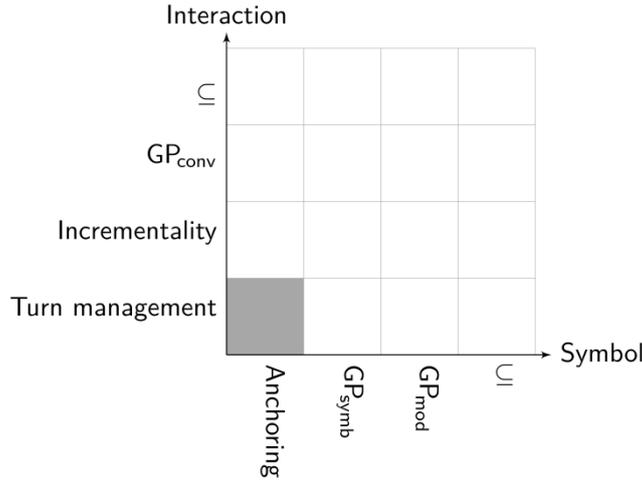


Figure 1: Grounding Steps for ADCs.

6 Conclusion

We have identified three grounding problems for the semantics of symbols used by artificial dialog companions. Firstly, we argued that the intrinsic semantics of symbols acquired according to the basic symbol grounding problem (GP_{symb}) is limited and that therefore GP_{symb} has not been solved in general yet. Nevertheless, current approaches have implemented ways to master the anchoring problem (connecting sensory and symbolic information), which is a subset of the GP_{symb} . Secondly, the dialog aspect of ADCs requires a model of linguistic grounding as a centerpiece. We identified the principal items of creating and managing common ground. We noted that full conversational grounding rests on turn management (contributing, acknowledging, repairing, clarifying) and incrementality. Thirdly, we posed the philosophical grounding problem as a test case for the intrinsic meaning of the symbols in ADCs' lexicons. If an artificial dialog agent is able to talk about possible states of affairs that question the co-referentiality of synonymous terms, then this agent has acquired an intrinsic concept of meaningfulness. Such a test is, to our knowledge, still missing in discussions of ADCs but is needed in order to assess their symbol grounding achievements.

If we map the grounding problem onto the two dialog dimensions (*Interaction vs. Symbol* - cf. Section 2 above), we receive the two-dimensional grid from Figure 1. The grid stakes out the space of grounding as delimited here into nine fields (we added an additional row and column for further grounding steps). The grid can be used to assess in more detail the dialogical effectiveness of ADCs. Figure 1 accordingly indicated the current achievements of conversational agents by gray highlighting of fields. As argued in the main text above, ADCs have solved the anchoring problem on the symbol axis and have been equipped with turn-taking modules. The visual representation allows us to spot quickly that there are still some steps to go until an ADC can become a cooperative conversational partner.

7 References

- Austin, John L. (1962). *How To Do Things With Words*. 2nd ed. Cambridge, MA: Harvard University Press.
- Baronchelli, A., V. Loreto, and L. Steels (2008). "In-Depth Analysis of the Naming Game Dynamics: The Homogeneous Mixing Case." In: *International Journal of Modern Physics C* 19.5, pp. 785–812.
- Baronchelli, Andrea, Tao Gong, Andrea Puglisi, and Vittorio Loreto (2010). "Modeling the emergence of universality in color naming patterns." In: *PNAS* 107.6, pp. 2403–2407. doi: 10.1073/pnas.0908533107.
- Barwise, Jon (1989). *The Situation in Logic*. CSLI lecture notes 17. Menlo Park: CSLI.
- Brennan, Susan E. (1998). "The Grounding Problem in Conversations With and Through Computers." In: *Social and cognitive psychological approaches to interpersonal communication*. Ed. by S. R. Fussell and R. J. Kreuz. Hillsdale, NJ: Lawrence Erlbaum, pp. 201–225.
- Cangelosi, Angelo, Alberto Greco, and Stevan Harnad (2002). "Symbol Grounding and the Symbolic Theft Hypothesis." In: *Simulating the Evolution of Language*. Ed. by Angelo Cangelosi and Domenico Parisi. London: Springer. Chap. 9, pp. 191–210.
- Cassell, Justine (2001). "Embodied conversational agents: representation and intelligence in user interfaces." In: *AI magazine* 22.4, p. 67.
- Catizone, Roberta, Alexiei Dingli, Hugo Pinto, and Yorick Wilks (2008). "Information Extraction tools and methods for Understanding Dialogue in a Companion." In: *Proceedings of the Sixth International Conference on Language Resources and Evaluation*. LREC'08. Marrakech, Morocco, pp. 3250–3254.
- Cavazza, Marc et al. (2010). "Persuasive Dialogue Based on a Narrative Theory: An ECA Implementation." In: *PERSUASIVE*. Ed. by Thomas Ploug, Per F. V. Hasle, and Harri Oinas-Kukkonen. Vol. 6137. Lecture Notes in Computer Science. Springer, pp. 250–261.
- Clark, Herbert H., ed. (1992). *Arenas of Language Use*. Chicago ; Stanford: University of Chicago Press & CSLI.
- (1996). *Using Language*. Cambridge: Cambridge University Press.
- Clark, Herbert H. and Edward F. Schaefer (1987). "Collaborating on contributions to conversations." In: *Language and Cognitive Processes* 2.1, pp. 19–41. doi: 10.1080/01690968708406350.
- (1989). "Contributing to Discourse." In: *Cognitive Science* 13.2, pp. 259–294. doi: 10.1016/0364-0213(89) 90008-6.

- Coradeschi, Silvia and Alessandro Saffiotti (2003). "An introduction to the anchoring problem." In: *Robotics and Autonomous Systems* 43, pp. 85-96.
- Crook, Nigel, Cameron Smith, Marc Cavazza, Stephen Pulman, Roger Moore, and Johan Boye (2010). "Handling User Interruptions in an Embodied Conversational Agent." In: *Proceedings of the AAMAS International Workshop on Interacting with ECAs as Virtual Characters*. Toronto, Canada, pp. 27-33.
- Danilava, Sviatlana, Stephan Busemann, and Christoph Schommer (2012). "Artificial Conversational Companions. A Requirement Analysis." In: *4th International Conference on Agents and Artificial Intelligence*. ICAART 2012.
- Dowty, David R. (1979). *Word Meaning and Montague Grammar*. Dordrecht: Reidel.
- Fernández, Raquel and Jonathan Ginzburg (2002). "Non-Sentential Utterances: A Corpus Study." In: *Traitement Automatique de Langues* 43.2, pp. 13-42.
- Fodor, Jerry Alan (1997). "Connectionism and the Problem of Systematicity (Continued): why Smolensky's Solution still doesn't Work." In: *Cognition* 62.1, pp. 109-119.
- Fodor, Jerry Alan and Brian P. McLaughlin (1995). "Connectionism and the Problem of Systematicity: Smolensky's Solution Doesn't Work." In: *Connectionism: Debates on Psychological Explanation*. Ed. by C. MacDonald and G. MacDonald. Blackwell Publishers, pp. 199-222.
- Fricke, Ellen (2012). "Grammatik multimodal. Wie Wörter und Gesten zusammenwirken". De Gruyter: Berlin ; Boston.
- Gabrilovich, Evgeniy and Shaul Markovitch (2009). "Wikipedia-based Semantic Interpretation for Natural Language Processing." In: *Journal of Artificial Intelligence Research*, pp. 443-498.
- Gärdenfors, Peter (2000). *Conceptual Spaces*. Cambridge, MA: MIT Press.
- Gibbard, Alan (1975). "Contingent Identity." In: *Journal of Philosophical Logic* 4, pp. 187-221.
- Gilroy, Stephen, Julie Porteous, Fred Charles, and Marc Cavazza (2012). "PINTER: interactive storytelling with physiological input." In: *Proceedings of the 2012 ACM international conference on Intelligent User Interfaces*. IUI '12. Lisbon, Portugal: ACM, pp. 333-334.
- Ginzburg, Jonathan (2012). *The Interactive Stance: Meaning for Conversation*. Oxford, UK: Oxford University Press.
- Grice, Herbert Paul (1975). "Logic and Conversation." In: *Syntax and Semantics*. Ed. by Peter Cole and Jerry L. Morgan. Vol. 3, Speech Acts. New York: Academic Press, pp. 41-58.

- Harnad, Stevan (1990). "The symbol grounding problem." In: *Physica D: Nonlinear Phenomena* 42.1-3, pp. 335-346. doi: 10.1016/0167-2789(90)90087-6.
- Hauser, Marc D., Noam Chomsky and Tecumseh W. Fitch (2002). "The Faculty of Language: What Is It, Who Has It, How Did It Evolve?" In: *Science* 298.5598, pp. 1569-1579.
- Hörmann, Hans (1983). *Was tun die Wörter miteinander im Satz oder Wieviele sind einige, mehrere und ein paar?* Göttingen: Hogrefe.
- Kamp, Hans and Barbara Partee (1995). "Prototype Theory and Compositionality." In: *Cognition* 57.2, pp. 129-191.
- Kamp, Hans and Uwe Reyle (1993). *From Discourse to Logic. Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Dordrecht: Kluwer.
- Kopp, Stefan and Ipke Wachsmuth (2004). "Synthesizing multimodal utterances for conversational agents." In: *Journal of Computer Animation and Virtual Worlds* 15, pp. 39-52.
- (2012). "Artificial interactivity." In: *Handbook of Technical Communication*. Ed. by Alexander Mehler, Laurent Romary, and Dafydd Gibbon. Berlin/Boston: De Gruyter, pp.
- Krifka, Manfred (1992). "Thematic Relations as Links Between Nominal Reference and Temporal Constitution." In: *Lexical Matters*. Ed. by Ivan A. Sag and Anna Szabolcsi. CSLI Lecture Notes. Stanford: CSLI. Chap. 2, pp. 29-53.
- Kripke, Saul A. (1980). *Naming and Necessity*. Oxford: Blackwell.
- Lahav, Ran. (1989). "Against Compositionality: The Case of Adjectives." In: *Philosophical Studies* 57.3, pp. 261-279.
- Lewis, David (1969). *Conventions. A Philosophical Study*. Cambridge: Harvard U.P.
- (1973a). "Causation." In: *Journal of Philosophy* 70, pp. 556-567.
- (1973b). *Counterfactuals*. Oxford: Blackwell.
- Loreto, Vittorio, Animesh Mukherjee, and Francesca Tria (2012). "On the origin of the hierarchy of color names." In: *PNAS* 109.18, pp. 6819-6824.
- Lücking, Andy and Alexander Mehler (2011). "A Model of Complexity Levels of Meaning Constitution in Simulation Models of Language Evolution." In: *International Journal of Signs and Semiotic Systems* 1.1, pp. 18-38.
- Lücking, Andy and Alexander Mehler (2012). "What's the Score of the Naming Game? Constraints on Semantic Categorization." In: *Proceedings of the 9th International Conference on the Evolution of Language*. Evolang IX. Kyoto, Japan, pp. 196-203.

- Mehler, Alexander (2008). "On the Impact of Community Structure on Self-Organizing Lexical Networks." In: *Proceedings of the 7th Evolution of Language Conference*. Evolang 7, Barcelona, pp. 227-234.
- (2009). "Artifizielle Interaktivität. Eine semiotische Betrachtung." In: *Medienwandel als Wandel von Interaktionsformen – von frühen Medienkulturen zum Web 2.0*. Ed. by Tilmann Sutter and Alexander Mehler. Wiesbaden: VS.
- Mehler, Alexander and Andy Lücking (2012). "WikiNect: Towards a Gestural Writing System for Kinetic Museum Wikis." In: *Proceedings of the International Workshop on User Experience in e-Learning and Augmented Technologies in Education (UXeLATE 2012)* in Conjunction with ACM Multimedia 2012, Nara, Japan.
- Montague, Richard (1974). *Formal Philosophy: Selected Papers*. Ed. by Richmond H. Thomason. New Haven: Yale University Press.
- Newman, Mark E. J. (2010). *Networks: An Introduction*. Oxford: Oxford University Press.
- Parsons, Terence (1994). *Events in the Semantics of English: A Study in Subatomic Semantics*. Vol. 19. Current Studies in Linguistics Series. MIT Press.
- Peldszus, Andreas and David Schlangen (2012). "Incremental Construction of Robust but Deep Semantic Representations for Use in Responsive Dialogue Systems." English. In: *Proceedings of the Coling Workshop on Advances in Discourse Analysis and its Computational Aspects*. ADACA 2012, Mumbai, India, pp. 59-76.
- Peltason, Julia, Hannes Rieser, Sven Wachsmuth and Britta Wrede (2013). "On Grounding Natural Kind Terms in Human-Robot Communication". In: *Künstliche Intelligenz 27.2*, pp. 107-118.
- Poesio, Massimo and Hannes Rieser (2010). "Completions, Coordination, and Alignment in Dialogue." In: *Dialogue and Discourse 1*, pp. 1-89. doi: 10.5087/dad.2010.001.
- Poesio, Massimo and David R. Traum (1997). "Conversational Actions and Discourse Situations." In: *Computational Intelligence 13.3*, pp. 309-347.
- Prior, Arthur N. (1967). *Past, Present and Future*. Oxford: Clarendon Press.
- Puglisi, Andrea, Andrea Baronchelli, and Vittorio Loreto (2008). "Cultural route to the emergence of linguistic categories." In: *PNAS 105.23*, pp. 7936-7940. doi: 10.1073/pnas.0802485105.
- Purver, Matthew, Arash Eshghi, and Julian Hough (2011). "Incremental semantic construction in a dialogue system." In: *Proceedings of the 9th International Conference on Computational Semantics*. IWCS 2011. Oxford, UK, pp. 365-369.

- Rehm, Matthias, Elisabeth André, and Yukiko Nakano (2009). "Some Pitfalls for Developing Enculturated Conversational Agents." In: *Human-Computer Interaction. Ambient, Ubiquitous and Intelligent Interaction*. Ed. by Julie Jacko. Vol. 5612. Lecture Notes in Computer Science. Springer Berlin / Heidelberg, pp. 340–348. doi: 10.1007/978-3-642-02580-8_37.
- Reichenbach, Hans (1947). *Elements of Symbolic Logic*. New York: The Macmillan Company.
- Sacks, Harvey, Emanuel A. Schegloff, and Gail Jefferson (1974). "A Simplest Systematics for the Organization of Turn Taking for Conversation." In: *Language* 50, pp. 696–735.
- Salem, Maha, Stefan Kopp, Ipke Wachsmuth, Katharina Rohlfing, and Frank Joublin (2012). "Generation and Evaluation of Communicative Robot Gesture." In: *International Journal of Social Robotics* 4 (2), pp. 201–217. doi: 10.1007/s12369-011-0124-9.
- Schiffer, Stephen R. (1972). *Meaning*. Oxford, NY: Oxford University Press.
- Schlangen, David and Gabriel Skantze (2011). "A General, Abstract Model of Incremental Dialogue Processing." In: *Dialogue and Discourse* 2.1, pp. 83–111.
- Schulz, Ruth, Arren Glover, Gordon Wyeth, and Janet Wiles (2010). "Robots, communication, and language: An overview of the Lingodroid project." In: *Proceedings of the Australasian Conference on Robotics and Automation*. ACRA 2010. Brisbane, Australia.
- Searle, John (1971). *Sprechakte*. Frankfurt/Main: Suhrkamp.
- Searle, John R. (1978). "Literal meaning." In: *Erkenntnis* 13 (1), pp. 207–224. doi: 10.1007/BF00160894.
- (1980). "Minds, Brains, and Programs." In: *The Behavioral and Brain Sciences* 3, pp. 417–457.
- Stalnaker, Robert (1978). "Assertion." In: *Syntax and Semantics* 9, pp. 315–332.
- (2002). "Common Ground." In: *Linguistics and Philosophy* 25, pp. 701–721.
- Steels, Luc (1996). "Self-organising vocabularies." In: *Proceedings of Artificial Life V*. Nara, Japan, pp. 179–184.
- (2008). "The Symbol Grounding Problem Has Been Solved. So What's Next?" In: *Symbols and Embodiment: Debates on Meaning and Cognition*. Ed. by M. de Vega. Oxford: Oxford University Press. Chap. 12.
- ed. (2011). *Design Patterns in Fluid Construction Grammar*. Amsterdam: John Benjamins.

- Steyvers, Mark and Josh Tenenbaum (2005). "The large-scale structure of semantic networks: Statistical analyses and a model of semantic growth." In: *Cognitive Science* 29.1, pp. 41-78.
- Taddeo, Mariarosaria and Luciano Floridi (2005). "Solving the symbol grounding problem: a critical review of fifteen years of research." In: *Journal of Experimental & Theoretical Artificial Intelligence* 17.4, pp. 419-445. doi: 10.1080/09528130500284053.
- Tenenbaum, Joshua B., Charles Kemp, Thomas L. Griffiths and Noah D. Goodman (2011). "How to Grow a Mind?" In: *Science* 331.6022, pp. 1279-1285.
- Traum, David R. and Staffan Larsson (2003). "The Information State Approach to Dialogue Management." In: *Current and New Directions in Discourse and Dialogue*. Ed. by Jan Kuppevelt, Ronnie W. Smith, and Nancy Ide. Vol. 22. Text, Speech and Language Technology. Amsterdam, The Netherlands: Springer, pp. 325-353.
- Varela, Francisco J., E. Thompson, and Eleanor Rosch (1991). *The Embodied Mind. Cognitive Science and Human Experience*. Cambridge: MIT Press.
- Vendler, Zeno (1957). "Verbs and Times." In: *Philosophical Review* 56, pp. 143-160.
- Visser, Thomas, David Traum, David DeVault, and Rieks op den Akker (2012). "Toward a Model for Incremental Grounding in Spoken Dialogue Systems." In: *Proceedings of the 12th International Conference on Intelligent Virtual Agents*. IVA 2012. Santa Cruz, CA.
- Vogt, Paul (2005). "The emergence of compositional structures in perceptually grounded language games." In: *Artificial Intelligence* 167.1-2, pp. 206-242.
- Wachsmuth, Ipke (2008). "'I, Max' - Communicating with an artificial agent." In: *Modeling Communication with Robots and Virtual Humans*. Ed. by I. Wachsmuth and G. Knoblich. LNAI 4930. Berlin: Springer, pp. 279-295.
- Wachsmuth, Ipke and Günther Knoblich (2005). "Embodied Communication in Humans and Machines - A Research Agenda." In: *Artificial Intelligence Review* 24 (3), pp. 517-522.
- Waltinger, Ulli, Alexa Breuing, and Ipke Wachsmuth (2011). "Interfacing Virtual Agents With Collaborative Knowledge: Open Domain Question Answering Using Wikipedia-based Topic Models." In: *Proceedings of the International Joint Conference on Artificial Intelligence*. IJCAI 2011, Barcelona, Spain.
- Wilks, Yorick (2005). "Artificial companions." In: *Interdisciplinary Science Reviews* 30.2, pp. 145-152.

— (2007). “Is There Progress on Talking Sensibly to Machines?” In: *Science* 318, pp. 927–928.

— (2009). “Artificial Companions as Dialogue Agents.” In: *Proceedings of SIGDIAL 2009: the 10th Annual Meeting of the Special Interest Group in Discourse and Dialogue*, Queen Mary University of London, p. 216.

— ed. (2010). *Close Engagements with Artificial Companions: Key social, psychological, ethical and design issues*. Vol. 8. Natural Language Processing. Amsterdam and Philadelphia: John Benjamins.

Wilks, Yorick, Roberta Catizone, Simon Worgan, and Markku Turunen (2011). “Some background on dialogue management and conversational speech for dialogue systems.” In: *Computer Speech and Language* 25.2, pp. 128–139. doi: 10.1016/j.csl.2010.03.001.

Wilks, Yorick et al. (2010). “A prototype for a conversational companion for reminiscing about images.” In: *Computer Speech and Language* 25, pp. 140–157.

— (2011). “A prototype for a conversational companion for reminiscing about images.” In: *Computer Speech and Language* 25.2, pp. 140–157. doi: 10.1016/j.csl.2010.04.002.

Wittgenstein, Ludwig. (1953). *Philosophical Investigations*. Oxford, UK: Blackwell Publishing.

Ziemke, Tom (1999). “Rethinking Grounding.” In: *Understanding Representation in the Cognitive Sciences. Does Representation Need Reality?* Ed. by Alexander Riegler, Markus Peschl, and Astrid von Stein. New York/Boston/Dordrecht: Kluwer/Plenum, pp. 177–190.

ⁱ Strictly speaking, Artificial Companion is a hypernym of the kind of conversational systems that we focus on here, since it additionally encompasses, for example, companions like artificial pets, which we exclude from our discussion.

ⁱⁱ To keep a short argumentation, we circumvent any discussion of the notion of independence in terms of algorithmic determinism etc. The interested reader should refer to Ziemke (1999) and related references.

ⁱⁱⁱ Note that a notion of *language* may include social communities (Wittgenstein, 1953), nonverbal communication means (Fricke, 2012) and brain structures (Hauser, Chomsky and Fitch, 2002).

^{iv} For this notion see, for example, Kopp and Wachsmuth (2012) and Mehler (2009).

^v As we do not require ADCs to be intelligent, we want to circumvent any discussion of hard versus soft AI (Searle, 1980).

^{vi} For an overview of these approaches see Steels (2011). A very advanced project in this area is probably the *Lingodroids* project (Schulz, Glover, Wyeth, and Wiles, 2010).

^{vii} The interested reader may consult Lücking and Mehler (2012) for the details of this argumentation.

^{viii} See also Taddeo and Floridi (2005), who argue that so far no approach to the symbol grounding problem accomplished full intrinsicity of meaning (what the authors refer to as *zero semantical commitment condition*).

^{ix} There is a bunch of work that corroborates the cooperative nature of dialog, but Herbert Clark probably sketched this issue most explicitly and extensively.

^x Since a great variety of different and differently scaled phenomena are subsumed under the heading of pragmatics – for instance, conversational implicatures (Grice, 1975) or wide background knowledge (Searle, 1978) – we deem it unfair to construct pragmatic counterexamples in this context.

^{xi} We use quotation marks here, since we do not assume that such non-sentential utterances are somehow deficient – quite the contrary (see also the analysis of Ginzburg, 2012, Chap. 7).

^{xii} In a recent anthology of artificial companions (Wilks, 2010), the term “grounding” is used only once, namely in a footnote where a dialogical repair situation is distinguished from decreasing engagement in conversation.

^{xiii} The authors also argue that grounding of natural kind terms in human-computer interactions does not climb the complete Clarkian action ladder (Clark, 1996), but remains on a level that in the context of the present paper can be described as “public anchoring”.