
Toward Givenness Hierarchy Theoretic Natural Language Generation

Poulomi Pal

POULOMIPAL@MINES.EDU

MIRRORLab Colorado School of Mines, 1600 Illinois Street, Golden, CO 80401 USA

Tom Williams

TWILLIAMS@MINES.EDU

MIRRORLab Colorado School of Mines, 1600 Illinois Street, Golden, CO 80401 USA

Abstract

Language-capable interactive robots participating in dialogues with human interlocutors must be able to naturally and efficiently communicate about the entities in their environment. A key aspect of such communication is the use of anaphoric language. The linguistic theory of the *Givenness Hierarchy* (GH) suggests that humans use anaphora based on the *cognitive statuses* their referents have in the minds of their interlocutors. In previous work, researchers presented GH-theoretic approaches to robot anaphora understanding. In this paper we describe how the GH might need to be used quite differently to facilitate robot anaphora generation.

1. Introduction

For language-enabled collaborative robots to be effectively introduced into human society, they must be able to communicate about the objects, locations, and people in human environments and must be able to understand and use pronominal forms like *it*, *this*, and *that* just like humans. According to the theory of the *Givenness Hierarchy* (GH) (Gundel et al., 1993), when humans use anaphora rather than definite description, they signal their subconscious assumptions about the *cognitive status* their target referent has in the mind of their interlocutor. Thus, for robots to understand and generate human-like language they must be able to model the cognitive status of their referents.

Previously, Williams & Scheutz (2018) (see also Williams et al. (2016)) presented the first full computational implementation of the GH for the purpose of robotic natural language understanding (NLU), using a set of hand-crafted rules informed by the GH literature. Specifically, while Williams & Scheutz (2018) demonstrated that for anaphora understanding what is needed is only a set of searchable GH-theoretic data structures, as we will show in this paper, efficient anaphora generation requires these data structures to be augmented with a model that directly maps (some limited set of) entities to their presumed cognitive status.

In previous work (Pal et al., 2020), we proposed a statistical per-entity model of cognitive status in which a *Cognitive Status Engine* comprised of per-entity *Cognitive Status Filters* is used to maintain a distribution over cognitive statuses for each entity the robot believes to be at least familiar to all parties within the conversation. In this paper, we describe how this model can be used in conjunction with data structure oriented methods to facilitate efficient and effective algorithms for deciding when and how to use anaphora.

2. Related Work

2.1 Cognitive Status

The Givenness Hierarchy, originally presented by Gundel et al. (1993), is comprised of six hierarchically nested tiers of cognitive statuses: *in focus*, *activated*, *familiar*, *uniquely identifiable*, *referential*, and *type identifiable*, associated with each of which is a set of referring (or pronominal) forms that can be used when referring to an entity with that status (Gundel et al., 2006; Hedberg, 2013). The GH coding protocol, presented by Gundel et al. (2006), provides guidelines as to what features of linguistic and environmental context can be used to determine the cognitive status of a given entity. Since the GH has been shown to be well validated across a wide variety of languages beyond English (Gundel et al., 2010), many researchers such as Kehler (2000); Chai et al. (2004, 2006); Williams et al. (2016) and Williams & Scheutz (2018) have sought to computationally implement it for the purposes of reference resolution. These approaches use the GH to (1) justify sets of data structures that hold representations of candidate referents with different presumed cognitive statuses, and (2) determine how to use those data structures when resolving referring expressions.

While this data structure-oriented approach facilitates the efficiency of NLU by providing restricted subsets of candidates to be considered, during referring expression *generation*, the speaker already knows what entity they wish to refer to, and needs to determine what cognitive status that entity likely has in order to determine how to refer to it. Thus, what is needed, during natural language generation (NLG) (in *addition* to these data structures, as we will later describe) is a means of quickly determining not what entities have a given cognitive status, but what cognitive status is most likely for a given entity. Accordingly, the first contribution of this paper is to fill this gap in the literature by proposing a technique for jointly maintaining GH-theoretic cognitive status buffers needed for efficient language understanding and per-entity cognitive status models needed for NLG.

2.2 Referring Expression Generation

A key task in NLG is referring expression generation (REG): deciding how to refer to an entity so as to disambiguate it from a contrast set. Recent work in robot dialogue systems has sought to extend classic REG algorithms to handle the knowledge constraints imposed in realistic robotic domains and by realistic robotic architectures (Williams & Scheutz, 2017). Central to these algorithms is identification of the set of distractors that must be eliminated by a generated referring expression. Critically, while these algorithms may naively use as a set of distractors all objects that the speaker knows of, this is typically unnecessary, due to the way that attentional focus, and discourse context often work together to constrain the set of reasonable referents (Grosz & Sidner, 1986).

We argue that the GH stands to play a critical role in mediating the interaction between low-level dialogue features and the set of distractors that must be ruled out by REs. For example, when a speaker uses “this *NP*”, it indicates that the speaker assumes their target referent to be “Activated” to the listener; a GH-theoretic inspired REG algorithm could thus determine that only entities assumed to be at least Activated must be ruled out in the generated *NP*. However, to the best of our knowledge, the GH has not previously been used to inform REG algorithms in any way. Accordingly, the second contribution of this paper is to fill this gap in the literature by proposing a technique for using the GH to inform the set of distractors to be eliminated by REG algorithms.

2.3 Anaphora Generation

While much research has been performed on referring expression generation, relatively little attention has been paid to anaphora generation Gatt & Krahmer (2018). Previous work on anaphora generation has explored factors such as discourse structure, coherence, and salience (Poesio et al., 2004; McCoy & Strube, 1999; Kibble & Power, 2004; Ge et al., 1998; Callaway & Lester, 2002; McCoy & Strube, 1999), but these works directly associate these low-level features with pronoun use, rather than using cognitive status as a mediating factor as suggested by the linguistics literature.

The GH’s claim that choice of referring form depends on presumed cognitive status Rosa & Arnold (2011) suggests that the GH should be a powerful tool for deciding when and how to use anaphoric expressions. Naively, one might assume the straightforward use the fact that each tier of the GH is associated with a different set of referring forms that can be used to refer to objects assumed to have that status (or higher). However, because progressively lower tiers become associated with progressively greater numbers of possible referring forms, and because forms like "it", "this", and "that" can be ambiguous, these tier-form associations are not sufficient on their own for language generation. Indeed, although the knowledge of cognitive status affects language understanding and generation, no significant linguistic work exists that identifies the mechanisms by which it is actually *used* for selection of referring forms Arnold (2016). Accordingly, the third contribution of this paper is to fill this gap in the literature by proposing a GH-theoretic model of NLG that leverages the GH-theoretic cognitive status modeling and REG capabilities described above.

3. Cognitive Status Estimation and Memory Modeling

As described in Pal et al. (2020), we formulate cognitive status modeling as a Bayesian filtering problem. Let a dialogue consist of a set of utterances U_0, \dots, U_n . For object o , let $S_o^t \in \{I, A, F\}$ denote the cognitive status of o at a particular timestep t after utterance U_t (either In Focus, Activated or Familiar), and let $L_o^t \in \{N, M, T\}$ denote the linguistic status of o in utterance U_t (either not mentioned in the utterance, mentioned in the utterance in a non-topic role, or mentioned in the utterance in a topic role). Using this formalism, our goal is to recursively estimate, for a given object, the probability distribution over cognitive statuses for object o at time t :

$$p(S_o^t) = p(S_o^{t-1})p(L_o^t)p(S_o^t | S_o^{t-1}, L_o^t) \quad (1)$$

Algorithm 1 GetStatus(O)

- 1: CSE: Cognitive Status Engine comprised of Cognitive Status Filters $\{CSF_0 \dots CSF_n\}$
 - 2: **if** $\exists(CSF \in CSE \mid CSF.entity = O)$ **return** $\underset{S}{argmax}P(S|O)$
 - 3: **else return** *UID*
 - 4: **endif**
-

We define a Bayesian filter of this form as a *Cognitive Status Filter* (CSF) for a given object o . Given a set of known objects, $O = \{o_1, \dots, o_n\}$, our goal is then to estimate this distribution for each $o \in O$ at each time step (Pal et al., 2020). We maintain distributions over cognitive statuses for each

object presumed to be Familiar or higher through a Cognitive Status Engine (CSE) comprised of CSFs for each such object (if an entity is not tracked by a CSF, it is assumed to be at most Uniquely Identifiable (UID)). We also propose to maintain a set of data structures associated with each tier of cognitive status, each of which contains pointers to the objects assumed to have (at most) that status. Whenever the most probable cognitive status for an object changes, this reference is updated. The $\text{GetStatus}(O)$ function (Algorithm 1) takes an object O and returns its most likely cognitive status. If no CSF exists for O within the CSE, “UID” is returned; otherwise the most probable cognitive status for O (as determined by the distribution maintained by O ’s CSF) is returned.

Algorithm 2 Describe (O)

```

1: S = GetStatus(O)
2: D = GetDistractors(O,S)
3: if S == “UID” then
4:   return “the” + REG(O,D)
5: else if S == “FAM” then
6:   return “that” + REG(O,D)
7: else if S == “ACT” then
8:   if D ==  $\emptyset$  then
9:     if dist(O) == close return “this” else return “that”
10:  else
11:    return DREG(O,D)
12:  end if
13: else if S == “IF” then
14:   if D ==  $\emptyset$  return “it” else return DREG(O,D)
15: end if

```

4. Algorithmic Approach

We now propose the *Describe* algorithm (Alg. 2). Our ultimate goal is to decide how to refer to a target O based on its cognitive status (Line 1). We do so using the concept of a *Cognitive Status Engine* (Sec. 3, Alg. 1). Once we have determined the cognitive status of O , we use it to decide how to refer to O . Each tier of cognitive status is associated with a set of referring forms. One might thus naively assume that to decide how to refer to O , one could simply use a referring form associated with O ’s presumed cognitive status. While we may do so when O is at most Familiar or Uniquely Identifiable (Lines 3-6), when O is In Focus or Activated, this is infeasible for two reasons. First, there may be multiple candidate referents that could plausibly be referred to using that referring form (e.g., when a form like “this” is used to cue an activated object, and more than one activated object exists). Accordingly, when this type of ambiguity is identified (i.e., when the set of status-sensitive distractors for an object are non-empty)¹ (Lines 8 and 14), it may be avoided

1. To construct a status-sensitive distractor set, we must identify all objects with at least the same status as our target. The cognitive status of the target can be quickly retrieved from the CSE (Section 3). The distractors to be ruled out are then the union of the data structures associated with the target’s current most probable cognitive status and higher.

by selecting a referring form *that involves a definite description* at either the same level or a lower level of cognitive status, e.g. selecting “this *NP*” rather than “it”. Second, some referring forms (e.g., “this” and “that”) are only appropriate for objects that are physically (or cognitively²) close. When this type of conflict is identified³, it may be avoided by selecting a referring form that *does not violate such constraints* at either the same level or a lower level of cognitive status, e.g., selecting “that *N*” rather than “this *N*”, as shown by the distance-sensitivity displayed both in Alg. 2 and Alg. 3.

Once we have determined what referring form to use, if that form includes a constituent *NP*, we must generate that referring expression using an REG algorithm such as the Incremental Algorithm (Dale & Reiter, 1995) or DIST-PIA (Williams & Scheutz, 2017), which require (1) a target referent, and (2) a set of distractors that must be ruled out. Critically, because we are using a cognitive status theoretic approach, once we decide what referring form to use for an object, we can intelligently identify a reduced set of distractors to rule out, as we can expect that the listener will not even consider objects of lower cognitive status than that cued by the chosen referring form (Alg. 3).

Algorithm 3 DREG(O,D)

```

1: if dist(O)== close return “this” + REG(O, Dthis)
2: else if dist(O) == far return “that” + REG(O, Dthat)
3: else return “the” + REG(O, D)
4: endif

```

5. Conclusion

We have presented techniques for: (1) jointly maintaining GH-theoretic inspired data structures required for effective NLU and the statistical per-entity cognitive status models for NLG, (2) using the GH to identify distractors to be ruled out during REG, and (3) GH-theoretic NLG. In future work we plan to integrate and evaluate all described elements that are not yet implemented within our robotic architecture, to enable effective and efficient anaphora generation in realistic HRI scenarios. We will evaluate our algorithm using a human subject experiment where we will judge human preference for referring forms generated by the algorithm (and presented through a robot) and assess the accuracy and efficiency of those referring forms.

References

- Arnold, J. E. (2016). Explicit and emergent mechanisms of information status. *Topics in Cog. Sci.*
 Callaway, C. B., & Lester, J. C. (2002). Pronominalization in generated discourse and dialogue. *Proc. ACL* (pp. 88–95).

2. An entity may be “cognitively close” if it is reachable within a certain number of edges from topics currently in focus or activated within a semantic memory network, or if it is sufficiently recent within episodic memory.
 3. I.e., when the weighted sum ($dist(O)$) over physical and episodic distance measures to object *O* (either globally or with respect to some contrast set) crosses some threshold

- Chai, J. Y., Hong, P., & Zhou, M. X. (2004). A probabilistic approach to reference resolution in multimodal user interfaces. *Intelligent User Interfaces (IUI)*.
- Chai, J. Y., Prasov, Z., & Qu, S. (2006). Cognitive principles in robust multimodal interpretation. *Journal of Artificial Intelligence Research*, 27, 55–83.
- Dale, R., & Reiter, E. (1995). Computational interpretations of the gricean maxims in the generation of referring expressions. *Cognitive science*, 19, 233–263.
- Gatt, A., & Krahmer, E. (2018). Survey of the state of the art in natural language generation: Core tasks, applications and evaluation. *Journal of Artificial Intelligence Research*, 61, 65–170.
- Ge, N., Hale, J., & Charniak, E. (1998). A statistical approach to anaphora resolution. *Sixth Workshop on Very Large Corpora*.
- Grosz, B. J., & Sidner, C. L. (1986). Attention, intentions, and the structure of discourse. *Computational linguistics*, 12, 175–204.
- Gundel, J. K., Bassene, M., Gordon, B., Humnick, L., & Khalfaoui, A. (2010). Testing predictions of the givenness hierarchy framework: A crosslinguistic investigation. *Journal of Pragmatics*.
- Gundel, J. K., Hedberg, N., & Zacharski, R. (1993). Cognitive status and the form of referring expressions in discourse. *Language*, (pp. 274–307).
- Gundel, J. K., et al. (2006). Coding protocol for statuses on the givenness hierarchy. *Unpublished manuscript (1993/2006)*. http://www.sfu.ca/hedberg/Coding_for_Cognitive_Status.pdf.
- Hedberg, N. (2013). Applying the givenness hierarchy framework: Methodological issues. *International Workshop on Information Structure of Austronesian Languages*.
- Kehler, A. (2000). Cognitive status and form of reference in multimodal human-computer interaction. *AAAI/IAAI*.
- Kibble, R., & Power, R. (2004). Optimizing referential coherence in text generation. *Comp. Ling.*
- McCoy, K. F., & Strube, M. (1999). Generating anaphoric expressions: pronoun or definite description? *The Relation of Discourse/Dialogue Structure and Reference*.
- Pal, P., Zhu, L., Golden-Lasher, A., Swaminathan, A., & Williams, T. (2020). Givenness hierarchy theoretic cognitive status filtering. *Proceedings of the Annual Conference on Cognitive Science*.
- Poesio, M., Stevenson, R., Eugenio, B. D., & Hitzeman, J. (2004). Centering: A parametric theory and its instantiations. *Computational linguistics*, 30, 309–363.
- Rosa, E. C., & Arnold, J. E. (2011). The role of attention in choice of referring expressions. *Proceedings of PRE-Cogsci: Bridging the gap b/w computational, emp. and theoretical app. to reference*.
- Williams, T., Acharya, S., Schreitter, S., & Scheutz, M. (2016). Situated open world reference resolution for human-robot dialogue. *Proc. Int'l Conf. on Human-Robot Interaction (HRI)*.
- Williams, T., & Scheutz, M. (2017). Referring expression generation under uncertainty: Algorithm and evaluation framework. *Proc. Int'l Conf. on Natural Language Generation*.
- Williams, T., & Scheutz, M. (2018). Reference in robotics: A givenness hierarchy theoretic approach. *The Oxford handbook of reference*.