Does informativity modulate linearization preferences in reference production?

Muqing Li (muqingli@coli.uni-saarland.de)[1]
Noortje J. Venhuizen (N.J.Venhuizen@tilburguniversity.edu)[2]
Torsten Kai Jachmann (jachmann@coli.uni-saarland.de)[1]
Heiner Drenhaus (drenhaus@coli.uni-saarland.de)[1]
Matthew W Crocker (crocker@coli.uni-sb.de)[1]

[1] Department of Language Science and Technology, Saarland University, 66123 Saarbrücken, Germany
[2] Department of Cognitive Science and Artificial Intelligence, Tilburg University, 5037 AB Tilburg, Netherlands

Abstract

During referential communication, speaker choices regarding the syntactic encoding of their expressions can modulate the linear ordering of the properties necessary to identify the referent. We investigated whether such syntactic choices are influenced by the *informativity* of these properties in a given visual context, as quantified by Referential Entropy Reduction (RER). In two experiments, a maze-based sentence completion task was used to examine whether informativity of a particular property (animal or action) influenced the decision to produce pre- versus post-nominal modifications when describing animal-performing-action referents in a visual scene. While many participants used a fixed strategy, informativity did significantly influence linearization for the remaining participants, consistent with a maximal informativity strategy in which the high RER property is be encoded first. This suggests that speakers who vary their encodings are indeed sensitive to the informativity of properties in a visual scene, preferring syntactic linearization in which informative properties appear early.

Keywords: visually situated language production; reference; informativity; referential entropy reduction; linearization

Introduction

In visually situated communication tasks, speakers will often need to formulate referring expressions that identify a specific, co-present referent to a listener. To accomplish this, the speaker must both 1) determine the relevant properties necessary for the listener to identify the target, and 2) select a syntactic encoding that conveys these properties, resulting in a particular linear ordering. The current study focuses on the latter type of encoding choice, namely encoding decisions regarding linearization. Specifically, we investigate whether linearization choices between pre-nominal versus post-nominal modification – e.g., "the crying rabbit" vs. "the rabbit that's crying" – are influenced by the *informativity* of a particular property/word ("rabbit" vs. "crying"), in identifying the intended referent.

Informativity and Referential Entropy Reduction

Following an information theoretic approach (Shannon, 1948), informativity is quantified by Referential Entropy Reduction (RER, Tourtouri, Delogu, Sikos & Crocker, 2019). Referential Entropy measures the uncertainty of identifying the intended target among all possible referents x in a visual scene (Equation 1). Referential Entropy Reduction (RER) quantifies the degree to which the uncertainty about the intended referent is reduced by each property word w (which applies to a subset of referents) in the unfolding of an utterance (Equation 2).

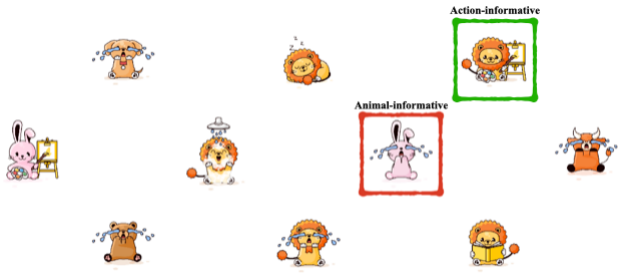

Figure 1. Example stimulus with targets in the Action-informative and Animal-informative Conditions. The informative property narrows down the selection scope from 10 to 2 figures, while the uninformative one narrows from 10 to 5. For each subject in experiments, only one target was highlighted in a red square in each trial.

$$H(X) = -\Sigma P(x)\log_2 P(x) \quad (1)$$

$$\Delta H_w = H_{w-1} - H_w \quad (2)$$

A word is more informative if it has higher RER, as it reduces uncertainty about the intended referent to a greater extent: For example, in the visual scene depicted in Figure 1, the initial Referential Entropy is 3.3 bits ($\log_2 10$), indicating referential uncertainty among the 10 possible referents. If a referential expression is uttered that starts with the word "rabbit", its RER ΔH_{rabbit} is 2.3 bits ($\log_2 10 - \log_2 2$), as the word "rabbit" narrows down the selection scope from ten referents to two referents. By comparison, if the expression starts with the word "crying", RER is lower ($\Delta H_{crying} = \log_2 10 - \log_2 5 = 1$ bit), because it only reduces the selection scope from ten to five referents, and hence it is less informative. In other words, RER is inversely correlated to the selection scope of a property given a visual context, i.e., the number of referents that can fit such property: The smaller the selection scope, the higher the RER of a property/word, and thus, the higher the informativity of the word.

In M. Goldwater, F. K. Anggoro, B. K. Hayes, & D. C. Ong (Eds.), *Proceedings of the 45th Annual Conference of the Cognitive Science Society.*

It has been shown that informativity plays a crucial role in property selection. Not only do speakers tend to *overspecify*, i.e., to include more property information than necessary for picking out the target (e.g., Koolen, Gatt, Goudbeek & Krahmer, 2011), but some speakers are more likely to overspecify for a property when it locally reduces RER to a greater extent, possibly to facilitate the incremental identification of the target by the listener (Tourtouri et al., 2019). However, while this provides some support for the claim that informativity influences production choices in over- vs. minimal- specifications, we are not aware of any studies that have directly examined whether informativity influences the production of meaning-equivalent minimal specifications, purely with regard to the order in which properties are encoded.

Therefore, in the current study, we first set out to determine whether the relative informativity of the two properties that are required to identify a referent influences the preference for the speaker to use a particular linearization of those properties, as reflected by their choice to use either a pre-nominal or post-nominal encoding. Although evidence has suggested that the planning of producing a pre-nominal or a post-nominal phrase may be qualitatively different (e.g., Brown-Schmidt & Konopka, 2008) and pre-nominal modifiers are generally preferred by English, German and Dutch speakers in previous referential communication studies (Fukumura, 2018; Tourtouri et al., 2019; van Gompel, van Deemter, Gatt, Snoeren & Krahmer, 2019), we are primarily interested in whether the effect of informativity can nonetheless modulate such encoding preferences.

If we find that situationally determined informativity does play a role, even over and above the overarching syntactic preference of pre-nominal modifications, we then further examine whether the preference favors placing more RER properties earlier or later in the encoding, and what the consequences are for extant theories about linearization strategies.

Linearization strategies

One potential linearization strategy for encoding referential expressions is to place the most informative expression as early as possible in the utterance, which we refer to as the *maximal informativity* hypothesis. One motivation for such orderings is that, in a visual context, the most informative word, i.e., with high RER, can help the listener to identify the target more rapidly. Fukumura (2018) reported cases where speakers can place the most discriminative property, i.e., the more informative adjective, as the first pre-nominal modifier to be uttered. By comparing between color overspecifications as the pre-nominal modifier in English and the post-nominal modifier in Spanish, Rubio-Fernández (2016) showed that the early ordered, pre-nominal color adjective is uttered more frequently, which contributes towards more efficient target identification. A similar notion of maximal informativity is also assumed by Rational Speech Act theory (RSA), which asserts that the most informative word is the most probable one to be uttered based on rational, Gricean-based assumptions (e.g., Frank & Goodman, 2012; Goodman & Frank, 2016; Grice, 1975; Horn, 1984). However, there is to date little empirical evidence whether maximal informativity can influence orderings when they entail syntactically different, but meaning equivalent, encodings. (An exception is Cohn-Gorden, Goodman & Potts, 2019, who touch upon the ordering issue in an incremental RSA computation model.)

Another possible strategy is that speakers linearize their utterances in line with the Uniform Information Density Hypothesis (UID, Jaeger, 2010; Frank & Jaeger, 2008). UID predicts that among alternative encodings, speakers prefer to choose the one that can distribute information more uniformly across the utterance, to avoid peaks of surprisal, the amount of information transmitted by a linguistic unit in a given context (Shannon, 1948), which can be viewed as a different, but related, notion of informativity to RER. A more uniform distribution will avoid surprisal peaks (and troughs), thus avoiding excessive comprehension effort for readers (e.g., Sikos, Greenberg, Drenhaus, & Crocker, 2017). Evidence for UID at the syntactic level has focused on syntactic reduction behavior as evidenced by corpus data (e.g., the likelihood of reduced relative clauses, in Levy & Jaeger, 2006; and *that*-complementizer omission, in Jaeger, 2010). Less is known as to whether a preference for a more uniform distribution of information density can be extended to the choice of different word orders. In contrast to the maximal informativity hypothesis, UID – if extended to the notion of RER in reference production settings – would predict a tendency for speakers to avoid the informativity peak, by preferring encodings that distribute informativity uniformly by ordering low RER properties earlier in the utterance.

The current study

The current experiment investigates whether informativity, quantified by RER, affects linearization choices, and if so, whether these choices better align with predictions of maximal informativity or other patterns, such as UID. We compare the use of pre-nominal versus post-nominal referring expressions in visually situated referential communication, by presenting participants with visual scenes in which the individual properties differ in their selection scopes such that the expressions describing these properties differ in informativity (RER).

We employ a set of visual stimuli depicting animals performing different actions (see Figure 1). Unlike simple object properties, such as color or pattern, which are highly constrained by grammatical convention to be encoded as the pre-nominal modifiers in English or German (e.g., Fukumura, 2018; Tourtouri et al., 2019), the *action* property in German is not obligatorily encoded as a pre-nominal structure (e.g., *der weinende Hase*, "the crying rabbit"), but can also be naturally placed in a post-nominal relative clause (e.g*., der Hase, der weint*, "the rabbit that cries"). Further, the conceptual and visual prominence of the *animal* property, encoded as a noun, may increase the overall likelihood of

eliciting a post-nominal expression, which are generally less preferred.

The flexibility of the two encoding orders with equivalent meanings (pre-nominal: action-animal or post-nominal: animal-action), offers an ideal basis for testing speakers' linear encoding decisions when the informativity of the two properties varies: In two critical conditions, the Animal-informative Condition and Action-informative Condition (see Figure 1), the description of the informative property renders a larger RER (narrowing the selection scope from 10 to 2 figures), compared with the uninformative property (narrowing the selection scope from 10 to 5 figures).

The current study adopts an innovative maze-based sentence completion task to investigate linearization choices in the description of an intended target. Participants are tasked with offering a description of a given target by consecutively selecting two parts of their expressions from a list of options (see Figure 2). In the first completion step (Step1), the referential expression is initiated by selecting an expression with either an adjective/action (*der weinende*, "the crying") or a noun/animal (*der Hase*, "the rabbit") from a list of four expressions (including two distractors). Depending on the choice of the expression in Step1, the referring expression is completed in Step2 by either selecting a matching noun (for pre-nominal modification) or a matching relative clause (for post-nominal modification). Critically, the list of expressions presented in Step2 will only contain expressions that fit the selection scope of the expression that are chosen in Step1 within the current visual scene. This means that the informativity of expressions in Step1 affects the number of options to be chosen Step2: if the choice at Step1 was informative, Step2 will only present two button options, because the informative property narrowed down the selection scope from 10 to 2 figures; Instead, if the choice at Step1 was uninformative, Step2 will list out five button options corresponding the five visual figures that can be modified by the expression at Step1. Such trade-off between informativity in Step1 and the number of button options at Step2 was designed to implicitly draw participants' attention to the RER contribution of the two properties in the visual scenes and in the individual expressions.

Our aim in using the maze task is to detect whether RER influences participants' choice regarding the initial word of the referring expression, which effectively determines the entire encoding. By presenting all encoding continuations to the participants, this paradigm may increase the opportunity for more varied linearizations, and mitigate experimentally-induced behaviour that may have inflated consistent responses in previous studies (e.g., Tourtouri et al., 2019; Tarenskeen, Broersma & Geurts, 2015).

In two experiments, participants perform the maze task in the context of an online communication game. If speakers do not utilize informativity for linearization at all, they would be insensitive to RER, and thus would not vary their encodings across the conditions. If informativity does play a role, speakers would adhere to either a maximal informativity strategy, by placing the most informative property first, or a UID strategy in which the lower RER property are encoded first.

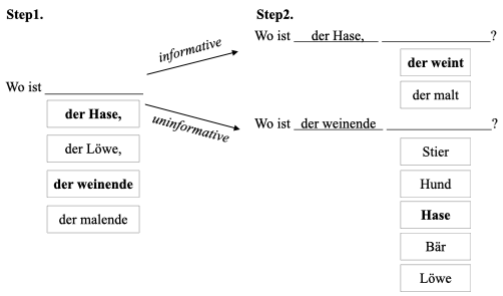

Figure 2. Example of the maze-based completion task. The target was "the crying rabbit" in Figure 1. The two steps were presented sequentially. Only one stem was presented for the subject, based on their decisions at Step1.

Experiment 1

Methods

Participants 80 participants (33 female and 47 male) were recruited online via the subject recruiting platform Prolific (www.prolific.co). Participants were native German speakers with an age range from 18 to 55 (mean = 28.73). Participants were paid according to German minimum wage (13.94€/hour).

Materials and Conditions The stimuli were designed based on 25 pictures depicting an animal performing an action (e.g., a crying rabbit). These figures were combined into 48 ten-figure displays (Figure 1 as an example). Each display consisted of figures with overlapped animal properties (e.g., 2 rabbits and 5 lions in Figure 1) and figures with overlapped action properties (2 painting and 5 crying animals). The locations of each figure, and of the target were pseudo-randomized.

Stimuli were distributed across 4 pseudo-randomized lists, consisting of 48 trials: 12 Animal-informative trials (e.g., crying rabbit in Figure 1), and 12 Action-informative trials (painting lion), as well as 12 Animal-singleton fillers (crying dog, as there is only one dog among the ten figures), and 12 and Action-singleton fillers (reading lion, as the only reading figure). Each filler trial was placed before a critical trial.

Procedure The experiment was implemented online via the platform LabVanced (Finger, Goeke, Diekamp, Standvoß & König, 2017). The experiment was narrated as a communication game, where the participants should collaborate with a partner online to identify whether the displays of both sides were identical, which was explained to be directly determined by the row location of the target. The experiment first presented a practice session where participants were guided to perform the tasks of both roles.

In the main experiment, each trial started with a ten-figure display without the indication of the target. After 2000ms, the

target figure was highlighted by a red square. Participants were guided to complete a question sentence about the target: *Wo ist __ __?* ("Where is__ __?"). The two blanks to be filled were presented consecutively (see Figure 2). A choice of an adjective in Step1 (e.g., "*der weinende*", "the crying") would lead to a pre-nominal description, while starting with the noun (e.g., "*der Hase,*", "the rabbit,") resulted in a post-nominal structure. The linguistic candidates were presented as buttons to be clicked, stacked vertically with a pseudo-randomized order. Step2 was conditioned by the choice at Step1, with two (if the more informative option was selected first) or five (if the less informative option was selected first) options to choose from. In the fillers, Step1 also provided one-word options that can minimally specify the target (e.g., "der Hase?", "der Weinende?"). The two distractors presented in Step1 did not semantically match with the target properties but could still be selected to proceed into Step2, fitting the selection scope of the distractors.

Once Step2 was finished, the sentence was replaced by a text reply from the partner (in reality, a simulated listener), depicting the row in which the target was situated from the listener's perspective. Participants then verified if the displays were identical with their partners based on the reply. Replies from the simulated listener in one third of the trials suggested a mismatched display pair, in which the target was found in a different row.

Results and discussion

Analysis One participant was excluded due to low accuracy: An accurate trial was one that was correct both in the target description and in the verification of the display. Participants were excluded if the proportion of inaccurate trials was more than 10%. Another 40 individual inaccurate trials (1.05%) from other participants were further excluded.

A generalized mixed modeling analysis was conducted (Bates, Mächler, Bolker, & Walker, 2015). The dependent variable was participants' choice at Step1 in the critical trials: whether the description started with the action- or the animal-property, leading to one of two syntactic structures for a description: pre-nominal or post-nominal modifications, respectively. The predictor of the model was the condition of the critical trials (Action-informative or Animal-informative Condition). The random intercepts for both items and subjects, as well as the random slopes for subjects, were included. Predictor Condition was dummy-coded, with the Action-informative Condition as the reference level.

Results Overall, there was a general preference for pre-nominal modification (i.e., starting with the action property), regardless of conditions (78.47% pre-nominal vs. 21.53% post-nominal modifications). Among the post-nominal modifications across conditions (i.e., starting with the animal property), they were more frequently used in the Animal-informative Condition (the proportion of post-nominal modifications within the Animal-informative Condition was 25.21%, 95% Clopper–Pearson binomial CI of the proportion was 22.46% - 28.12%), compared with the Action-informative Condition (17.83%, CI: 15.42% - 20.44%). This trend was, however, not significant ($\beta = 1.00, SE = .81, z = 1.24, p = .21$).

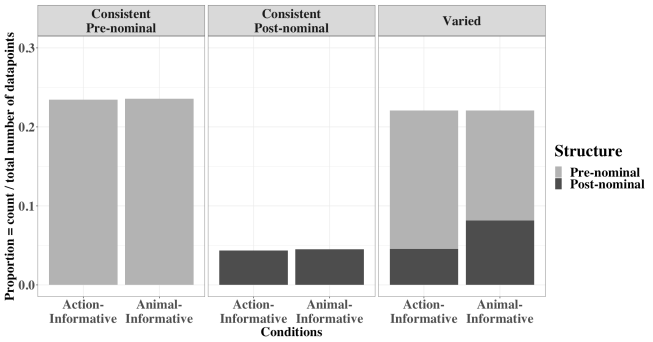

Figure 3. Overall proportion of modification structures by conditions and by subject groups in Experiment 1. A pre-nominal expression starts with the action property, while a post-nominal modification starts with the animal property. See text for Confidence Intervals (CI).

It was observed that subjects can be categorized based on whether they exhibit any variability in their encoding structures. A post hoc analysis divided the 79 participants into three groups based on their variability throughout the experiment: 37 participants (46.84%) always began with the action property, i.e., they consistently produced pre-nominal modifications for target descriptions (Group *Consistent Pre-nominal*); 7 participants (8.86%) consistently began with the animal property (Group *Consistent Post-nominal*); The remaining 35 participants (44.30%) showed variations in the choices at Step1 (Group *Varied*). Figure 3 presents the distribution of encodings for the three groups.

A generalized mixed modeling analysis was conducted for the Group *Varied* (N=35), in which the frequency of the use of post-nominal modifications was significantly higher in the Animal-informative Condition (36.89% within the group and the condition) than in the Action-informative Condition (20.63%) ($\beta = 1.08, SE = 0.18, z = 5.76, p < .001$).

Discussion In Experiment 1, pre-nominal modification was overwhelmingly preferred for describing our stimuli depicting animals performing actions using the maze task. This was contributed by two subject groups: Group *Consistent Pre-nominal*, who only adopted such structures throughout the critical trials of the experiment, as well as Group *Varied*, who showed variation in their linearization.

Participants in group *Varied* were sensitive to our manipulation of informativity: The more informative property, with higher RER, was placed early in the description more frequently when comparing the two critical conditions. This was especially evidenced by the higher frequent use of post-nominal modifications in the Animal-informative Condition compared with the Action-informative Condition.

Experiment 2

In Experiment 2, participants first perform a Listener Task by listening to auditory target descriptions that are encoded in a maximally informative manner, followed by the same Speaker Task as in Experiment 1. We hypothesized that sensitivity to the informativity of expressions may be enhanced if participants are first exposed to the more informative encodings in the role of Listener. This hypothesis is based on previous work on the influence of role-switching during communication: for instance, Sikos, Venhuizen, Drenhaus & Crocker (2021) found that more pragmatic reasoning was engaged if the listener performed the role of a speaker first. Therefore, in Experiment 2, we ask if the benefit of role-switching obtains: namely, whether having participants perform a Listener Task first would amplify the influence of informativity on the ordering of expressions in the Speaker Task.

Methods

Participants 159 native German speakers (74 female and 85 male) were recruited via Prolific (mean age = 27.77).

Materials and Conditions Both the Listener Task and the Speaker Task consisted of 24 trials per list (12 critical trials and 12 fillers). The critical trials utilized the same displays as Experiment 1 but were reduced to 6 items per condition.

Procedure After practice, participants first performed the Listener Task. Participants were required to listen to auditory descriptions of a target, and to identify the row in which it was situated. The auditory descriptions were synthesized speech generated via The CereVoice TTS system's Alex voice (Version 3.2.0).

Each listener trial first presented the ten-figure display. The auditory description started to play after 4000ms. Participants were then required to choose the corresponding row in which the target described in the speech was located. The procedure of the Speaker Task was identical to Experiment 1.

Results and discussion

Analysis For the Listener Task, accuracy was measured by the correct identification of the row in which the target was located. Accuracy in the Speaker Task was measured in the same way as in Experiment 1. Subjects with more than 10% inaccurate trials in either task were excluded. 10 participants were excluded in this step. Another 85 individual inaccurate trials (4.83%) from other participants were also excluded. Generalized mixed modeling analysis was conducted as in Experiment 1, except that both random intercepts and random slopes of items and subjects were included.

Results Overall, a main effect of Condition was observed ($\beta = 1.44, SE = .52, z = 2.65, p < .01$): The frequency of the use of post-nominal modifications was significantly higher in the Animal-informative Condition (21.59%, CI: 18.91% - 24.46%) compared with the Action-informative Condition (17.54%, CI: 15.08% - 20.22%). Subjects were further categorized by the same criterion as in Experiment 1: Group *Consistent Pre-nominal* obtained 74 subjects (49.66%); Group *Consistent Post-nominal* had 5 subjects (3.35%); 70 subjects (46.98%) belonged to Group *Varied*. Within Group *Varied*, the use of post-nominal modifications was significantly more frequent in the Animal-informative Condition (38.93%) than in the Action-informative Condition (30.27%) ($\beta = .68, SE = .29, z = 2.39, p < .05$). Figure 4 presents the distribution of encodings for the three groups visually.

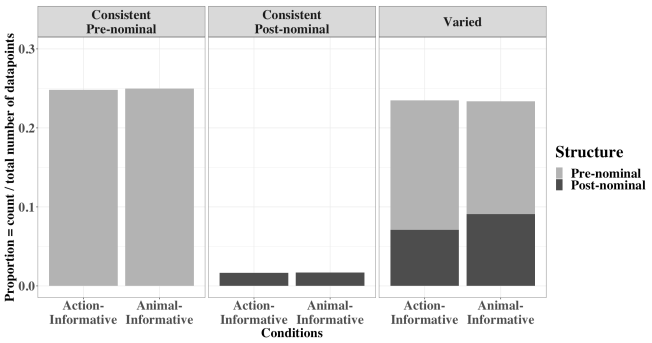

Figure 4. Overall proportion of modification structures by conditions and by subject groups in Experiment 2. See text for CIs.

Discussion Experiment 2 replicated the main result of Experiment 1, with a strong preference for pre-nominal encodings overall (80.43%). Three groups of subjects who differed in the consistency of description structures were again identified. Group *Varied* performed similarly as in Experiment 1: post-nominal modifications were used more frequently in the Animal-informative Condition compared with the Action-informative Condition, supporting a maximal informativity strategy. However, we did not observe a significant increase in the proportion of subjects for Group *Varied* with respect to Experiment 1 (44.30% in Experiment 1 vs. 46.84% in Experiment 2), indicating that the additional Listener Task did not increase participants' sensitivity to the informativity profile of the different linearization choices.

General discussion

The current study examined whether the informativity of a particular property affects speakers' linearization choices during referential production. Specifically, we investigated participants' choices in using pre-nominal or post-nominal modifications when describing visual scenes containing animals performing actions, where the two properties (animal and action) varied in their selection scopes, rendering different informativity profiles, quantified by RER.

In two experiments, we observed an overall strong preference for pre-nominal encodings. Participants in Group *Consistent* was completely invariant in the choice of syntactic encoding for every critical trial, with the majority choosing to begin their utterance with the action property, encoded as

the pre-nominal modifier. This group of speakers was insensitive to our manipulation of RER in the visual scenes, and at the same time, ignored the benefit of being informative at Step1 in the maze task, which led to fewer selection options required in Step2.

One explanation for this invariant behavior could be the reduced engagement in online experiments, where there is no co-present listener. Van Gompel (2021), for instance, observed differences between in-lab and online experiments with respect to informativity: without the presence of a listener, speakers not only tended to produce more overspecifications in online experiments, but also became insensitive to the visual manipulation of "discriminative power" (a similar notion to informativity). In future studies, we plan to conduct in-lab experiments, setting up with a co-present listener partner, to investigate whether a more engaging experimental setting can reduce the proportion of speakers in the *Consistent* Groups.

It is important to note, however, that the effort of maintaining such consistency in the current experimental setup was not trivial. In the maze task, not only are all possible continuations explicitly presented, but the positions of the response buttons were fully randomized such that participants were required to actively seek and click on their preferred expression, among four buttons in Step1, and among either two or five buttons in Step2, across multiple trials. Nevertheless, participants in the *Consistent* Groups managed to adhere to a single syntactic encoding structure for referential descriptions. This consistent behavior is at odds with the availability account (e.g., Bock, 1982), as well as with efficiency-based predictions (e.g., Pechmann, 1989; Rubio-Fernández, 2016): Although the animal property is more visually salient and presumably more available, the animal-first post-nominal expressions were not preferred among the participants, despite the fact that the maze task provided these as explicit choices and they are more informative in the Animal-informative Condition. Further research is required to obtain more insight into the mechanisms underlying the consistent preference for pre-nominal modifications in referential communication tasks.

Importantly, we also observed that speakers in the Group *Varied*, who showed variations in the use of pre-nominal and post-nominal encodings, were sensitive to informativity, as reflected in their linearization choices: when comparing the use of post-nominal modifications between the two conditions, speakers in Group *Varied* were more likely to place the animal property first in the Animal-informative Condition, leading to increased post-nominal modifications.

The observed variation in linearization choices purely at the syntactic level is a novel result. As far as we are aware, research concerning syntactic differences in meaning-equivalent referential encodings are not well-studied among studies adopting referential communication tasks. Current debates still mainly focus on over- or minimal- specifications of properties that are encoded as a string of pre-nominal modifiers. A notable exception is the work by van Gompel et al. (2019), where post-nominal overspecifications were elicited in describing the *border* property, e.g., "a plane in a diamond/square". However, mixed results were observed across conditions concerning the effect of informativity on post-nominal modifications in this study. Our results suggest that further research is required for a more thorough understanding of the effect of informativity on production and linearization strategies, not only with respect to pre- versus post-nominal modifications but also with respect to other syntactic encoding choices (e.g., double objects, coordination structures).

By prioritizing the most informative property in the utterance, the linearization choices of participants in the *Varied* Group are in line with the *Maximal informativity* hypothesis. It is worth noting that the current result may have been facilitated by the manipulation of the number of options in Step2 conditioned by Step1, which was intended to increase participants' awareness of the RER profile. Thus, in future studies we plan to examine the validity of this manipulation for a more solid effect of RER and informativity, by investigating if the informative-first preference remains without explicit task benefits.

Nevertheless, our results are consistent with the hypothesis that speakers prefer the maximally informative ordering, motivated by the facilitation of target identification for the listeners (e.g., Tourtouri et al., 2019; Fukumura, 2018). From the comprehension perspective, such facilitation effects have been previously observed: e.g., reduced retrieval effort in comprehension was detected in Staudte, Ankener, Drenhaus, & Crocker (2021), where meaning restrictions of a verb led to high RER for the subsequent object selection in a visual scene (e.g., the verb "spill" narrowed down RER to a smaller scope of liquid objects in the visual stimuli). Rubio-Fernández & Jara-Ettinger (2020) also found that target identification processing is incrementally aligned with the word order: Native language users can efficiently resolve temporary ambiguity by processing the noun contrast first if exposed to post-nominal expressions. Our results are thus in line with the idea that at least some speakers take the listener into account when encoding their utterance – i.e., audience design – possibly at the cost of increased production effort (in terms of a "trade-off between production ease and communication goals"; Kurumada & Jaeger, 2015). Future work is needed to investigate in which settings and to what extent the additional effort would be preferred to be expended by speakers.

Conclusion

Does informativity modulate linearization preferences in reference production? While some speakers consistently prefer pre-nominal modification, speakers in the *Varied* group in our study show that informativity does modulate linearization choices in terms of pre- versus post- nominal modification, where the most informative property is more likely to be encoded earlier in the utterance.

Acknowledgments

Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) — Project-ID 232722074 — SFB/CRC 1102.

References

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, *67*(1), 1–48.

Bock, J. K. (1982). Toward a cognitive psychology of syntax: Information processing contributions to sentence formulation. *Psychological Review*, *89*, 1–47.

Brown-Schmidt, S., & Konopka, A. E. (2008). Little houses and casas pequeñas: Message formulation and syntactic form in unscripted speech with speakers of English and Spanish. *Cognition*, *109*(2), 274–280.

Cohn-Gordon, R., Goodman, N., & Potts, C. (2019). An incremental iterated response model of pragmatics. *Proceedings of the Society for Computation in Linguistics (SCiL) 2019*, 81–90.

Finger, H., Goeke, C., Diekamp, D., Standvoß, K., & König, P. (2017). LabVanced: A unified JavaScript framework for online studies. *International Conference on Computational Social Science*.

Frank, A. F., & Jaeger, F. T. (2008). Speaking rationally: Uniform information density as an optimal strategy for language production. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 939–944.

Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, *336*(6084), 998–998.

Fukumura, K. (2018). Ordering adjectives in referential communication. *Journal of Memory and Language*, *101*, 37–50.

Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, *20*(11), 818–829.

Grice, H. P. (1975). Logic and conversation. In *Syntax and semantics 3: Speech acts*. Brill.

Horn. R. L. (1984). Toward a new taxonomy for pragmatic inference: Q-based and R-based implicature. In D. Schiffrin (Eds.), *Meaning, form, and use in context: Linguistic applications.* Georgetown University Press.

Jaeger, T. F. (2010). Redundancy and reduction: Speakers manage syntactic information density. *Cognitive Psychology*, *61*(1), 23–62.

Koolen, R., Gatt, A., Goudbeek, M., & Krahmer, E. (2011). Factors causing overspecification in definite descriptions. *Journal of Pragmatics*, *43*(13), 3231–3250.

Kurumada, C., & Jaeger, T. F. (2015). Communicative efficiency in language production: Optional case-marking in Japanese. *Journal of Memory and Language*, *83*, 152–178.

Levy, R., & Jaeger, T. F. (2006). Speakers optimize information density through syntactic reduction. *Proceedings of the 19th International Conference on Neural Information Processing Systems*, 849–856.

Pechmann, T. (1989). Incremental speech production and referential overspecification. *Linguistics*, *27*(1), 89–110.

Prolific. (2022). Available from: www.prolific.co

Rubio-Fernández, P. (2016). How redundant are redundant color adjectives? An efficiency-based analysis of color overspecification. *Frontiers in Psychology*, *7*, 153.

Rubio-Fernández, P., & Jara-Ettinger, J. (2020). Incrementality and efficiency shape pragmatics across languages. *Proceedings of the National Academy of Sciences*, *117*(24), 13399–13404.

Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, *27*, 623–656.

Sikos, L., Greenberg, C., Drenhaus, H., & Crocker, M. W. (2017). Information density of encodings: The role of syntactic variation in comprehension. *Proceedings of the 39th Annual Conference of the Cognitive Science Society (CogSci 2017)*, 3168–3173.

Sikos, L., Venhuizen, N., Drenhaus, H., & Crocker, M. W. (2021). Speak before you listen: Pragmatic reasoning in multi-trial language games. *Proceedings of the 43rd Annual Meeting of the Cognitive Science Society*, 1465–1471.

Staudte, M., Ankener, C., Drenhaus, H., & Crocker, M. W. (2021). Graded expectations in visually situated comprehension: Costs and benefits as indexed by the N400. *Psychonomic Bulletin & Review*, *28*, 624–631.

Tarenskeen, S., Broersma, M., & Geurts, B. (2015). Overspecification of color, pattern, and size: salience, absoluteness, and consistency. *Frontiers in psychology*, *6*, 1703.

Tourtouri, E. N., Delogu, F., Sikos, L., & Crocker, M. W. (2019). Rational over-specification in visually-situated comprehension and production. *Journal of Cultural Cognitive Science*, *3*, 175–202.

van Gompel, R. P. G., van Deemter, K., Gatt, A., Snoeren, R., & Krahmer, E. J. (2019). Conceptualization in reference production: Probabilistic modeling and experimental testing. *Psychological Review*, *126*, 345–373.

van Gompel. R. P. G. (2021, September 2-4). *The role of discriminatory power in reference production* [Conference presentation abstract]. Architectures and Mechanisms for Language Processing, Université de Paris, France. https://amlap2021.github.io/program/91.pdf.