

Speakers Use More Informative Referring Expressions to Describe Surprising Events

Christian Stegemann-Philipps (stegemann.christian@gmail.com)^{1,3}
Martin V. Butz^{1,3}, Susanne Winkler^{2,3}, Asya Achimova^{1,3}
{martin.butz, susanne.winkler, asya.achimova}@uni-tuebingen.de

¹ Department of Psychology and Department of Computer Science, Sand 14, 72076 Tübingen, Germany

² Department of English, Wilhelmstr. 50, 72074 Tübingen, Germany

³ Research Training Group “Ambiguity”, Keplerstr. 2, 72074 Tübingen, Germany

Abstract

Production of referring expressions (*the dog, it, Snoopy*) is a complex process regulated by a combination of linguistic and cognitive constraints. In this paper, we explore the impact of world knowledge on the types of references speakers produce, focusing on predictability of event progressions. We argue that speakers are more likely to use a full noun phrase rather than drop the subject or use a pronoun when they describe an event progression they find surprising. In order to avoid the influence of distributional properties of event descriptions, we created an artificial world and trained subjects to recognize typical collision-event progressions within. Speakers then described novel scenes, which either conformed to their expectations or violated them, in a free production experiment. The results reveal that unpredictable event progressions lead to a more frequent production of full noun phrases, in contrast to reduced linguistic expressions (pronouns and null subjects). We conclude that speakers choose more informative descriptions to talk about surprising events.

Keywords: Referring expressions; speech production experiment; event-predictive cognition; ambiguity avoidance; world knowledge

Introduction

Consider the following pair of sentences (1-2)¹:

- (1) I put the butterfly wing_{*i*} on the table_{*j*} and it_{*i*/**j*} broke.
- (2) I put the heavy book_{*i*} on the table_{*j*} and it_{*i*/*j*} broke.
(Davis, 2019)

Each of these sentences contains two potential antecedents for the pronoun *it*: the object of the verb *put* (*the butterfly wing* in (1) or *the heavy book* in (2)) and the noun phrase *the table* inside the prepositional phrase *on the table*. While syntax allows both binding options *it_{*i*/*j*}*, each of these sentences has only one reading that is coherent with our understanding of how the world normally works: in (1) *it* refers to the butterfly wing and not the table, while in (2) the pronoun *it* refers to the table rather than the heavy book. Such structurally parallel sentence pairs, known as *Winograd Schemas*², show an effect of semantic content on pronominal disambiguation.

¹We use the linguistic notation **_{*j*}* to denote unacceptable co-indexing. When the noun phrase and the pronoun bear the same index, it means that they refer to the same object/individual.

²Winograd schemas were first proposed by Terry Winograd as examples of particularly challenging machine translation problems (Winograd, 1972) and later turned into a competition by Hector Levesque (Levesque, 2014). These sentence pairs were designed to test the power of computational models of language and of world knowledge.

In this paper, we focus not on the disambiguation process but rather on the production side, hence, the speaker. We ask whether a speaker would refrain from producing an ambiguous sentence, such as (1), and use a full noun phrase *the butterfly wing* instead to avoid confusion. We further probe whether the hypothesized ambiguity avoidance is mediated by world knowledge—the general patterns of behavior of the world that are familiar to a person. Such knowledge allows the speaker to form expectations and leads to differing levels of surprise when she is confronted with unexpected or new events (Baldwin & Kosie, 2020; Butz, Achimova, Bilkey, & Knott, 2021; Kuperberg, 2020). We test whether the predictability of the referent affects how likely the speaker is to avoid ambiguity. Assuming that speaker and listener share background knowledge (such as that a butterfly wing is more fragile than a table), we ask how likely the speaker is to produce an ambiguous anaphoric reference as in (1). We furthermore probe whether the speaker factors in how easily a reduced anaphoric reference can be disambiguated by the listener. Our results show that world knowledge, operationalized as knowledge of common event patterns, affects the production of anaphoric references.

Tailoring A Message To The Listener

When speakers need to refer to an object or a person they have several options to choose from: a definite noun phrase (3), a pronoun (4) or a null subject (5). While the form in (5) may not be available in standard English, it is common in telegraphic speech (Barton, 1998; Haegeman, 2013).

- (3) The letter finally arrived today.
- (4) It finally arrived today.
- (5) Finally arrived today.

A number of linguistic factors potentially affect the choice of a referential expression. Speakers are less likely to produce a pronoun if the potential referents for it share relevant characteristics (Fukumura, van Gompel, Harley, & Pickering, 2011). The authors demonstrated that in a test scenario, if both male characters in a scene were shown on a horse, participants were less likely to describe the event where the first character got off the horse using a pronoun (“He got off the

horse”), and instead preferred to use a noun phrase. Topicality is another factor that affects the use of reduced referential expressions: Rohde and Kehler (2014) demonstrated that in their experiments speakers were more likely to produce pronouns if the referent was the topic. The effect of thematic roles on reference is less clear: Rosa and Arnold (2017) argued that thematic roles influence the frequency of producing a pronoun, while Fukumura and van Gompel (2010) found no such link.

From a production-centered perspective (for a review see Jaeger & Buz, 2017), a speaker should choose the referential expression that is easiest to produce: either a null subject (5) or a pronoun (4). This effort-saving behavior on the part of the speaker is mediated by the ability of the listener to infer the actual antecedent from context. If speakers focus on ease of production, anaphoric reduction should be dominant, independent of whether it is possible to infer the antecedent pragmatically.

On the other hand, if speakers focus on comprehension, they will not use reduced reference in cases where this may cause the listener problems to infer the correct reference, for example, in the light of ambiguity that may be introduced by the reduction. However, Fukumura and van Gompel (2012) argue that in producing pronouns instead of full noun phrases, speakers do not account for the knowledge of the listener but simply rely on their own privileged knowledge. The ability of the speaker to take into account the interpretability of the chosen referential expression for the listener has been called into question in a large body of literature on audience design and ambiguity avoidance (see Ferreira, 2019, for a review). Speakers have been shown to rarely avoid either lexical (Ferreira, Slevc, & Rogers, 2005) or structural ambiguity (Haywood, Pickering, & Branigan, 2005).

At the same time, speakers have been shown to adjust their description when giving directions based on whether the listener seems to know the local area or not (Kingsbury, 1968). They are also able to attenuate the characteristics of a story depending on whether the listener has heard the story before or not, including events mentioned, number of words, amount of detail, and even the clarity of articulation (Galati & Brennan, 2010).

While studies show that ambiguity avoidance was not a sufficient motivation to insert a disambiguating complementizer *that* to mark a new clause (Haywood et al., 2005), experimental work on the role of predictability in similar sentences paints a different picture. Predictability has been shown to affect the acoustic intensity of a word with more acoustically prominent words marking surprising information (Lam & Watson, 2010). In a task testing the effect of predictability on sentence structure (Jaeger, 2010), speakers were more likely to use the *that* in a phrase like “[m]y boss confirmed (that) we were absolutely crazy” if they judged the contents of the sentence to be less predictable. Jaeger (2010) therefore argues that speakers actively smooth the information density of an utterance, i.e. the rate of surprising information, call-

ing this the uniform information density hypothesis. Since it is difficult to judge how predictable or surprising words in a sentence are, Jaeger (2010) approximated this with next-word statistics derived from a corpus. Corpus studies further reveal that referential expressions in written text alternate between full descriptive reference, proper name, and pronoun depending on predictability (Tily & Piantadosi, 2009).

Predictability has also been shown to affect the grammatical encoding of an utterance. In an experimental work on Japanese, Kurumada and Jaeger (2015) have shown that speakers tend to insert case markings strategically when sentence properties, such as the animacy of the object or plausibility in general, suggest a different structure. Note that Kurumada and Jaeger (2015) quantified the plausibility of events by means of norming studies with native speakers rather than merely using word statistics. While the participating native speakers were most likely guided by their world knowledge when assigning their ratings, it is virtually impossible to exclude the possibility that they were also influenced by the distributional properties of target word combinations in spoken and written language, since such properties naturally reflect the structure of events they describe. Knowledge of event structures, sometimes also referred to as semantic knowledge, may be difficult to separate from the knowledge of word co-occurrences. Willits, Amato, and MacDonald (2015) show that the relative weight of these factors is also modulated by the nature of the task.

Bunger, Papafragou, and Trueswell (2013) have explored the effects of world knowledge more directly by a prime/scene description paradigm. They showed that conceptual overlap between the prime sentence and the target scene (such as motion through space) led to a tendency to assimilate structural components of the prime sentence into utterance production. Thus, abstract conceptual world knowledge, such as the type of an event, influenced production behavior.

In this paper, we present a more immediate way to explore the influence of world knowledge in the form of event predictability on language production. We ask whether world knowledge, and event prediction in particular, influences scene descriptions. To avoid confounding event knowledge with distributional regularities of its descriptions, we trained participants to learn about typical interactive event progressions in an artificial world. We show that speakers use full noun phrases more often when they describe the events that they find unusual.

Experiment

In the course of the experiment, participants watched animated scenes that pictured characters interacting in an artificial world. The participants were then asked to describe similar event progressions verbally. We used a free production set up in order to elicit maximally natural descriptions of the observed event progressions. This type of narration differs from scripted story continuation adopted in other psycholinguistic studies (e.g. Rosa & Arnold, 2017).

We manipulated the predictability of events that speakers had to describe. In half of the trials, the characters interacted according to the event progression patterns that the participants previously learned, while in the other half an unexpected event progression was shown. We predict that participants will use more pronominal or null references when the involved entities interact as expected. In contrast, when an unusual interaction is observed, we expect the speakers to produce more full noun phrases.

We use scenes from an artificial mini world called *Brain-Control* (Schrodt, Röhm, & Butz, 2017; Schrodt, Kneissler, Ehrenfeld, & Butz, 2017)³ to confront participants with an environment they are unfamiliar with.

Design

The scenes featured three types of moving entities: long blue cells, spiky red cells⁴, and round green viruses⁵. These entities can move left, move right, or rest. A collision between two entities forces one of them to flee. The blue cell flees from the green virus and the green virus flees from the red cell. Cells never interact. All possible interactions are shown in (6). Figures 1 and 2 provide examples of collision events shown to the participants.

- (6)
- Green virus attacks blue cell, blue cell runs away.
 - Green virus attacks red cell, green virus runs away.
 - Red cell attacks green virus, green virus runs away.
 - Blue cell attacks green virus, blue cell runs away.




Figure 1: The blue cell attacks the virus and flees




Figure 2: The green virus attacks the blue cell, the cell flees

The actual experiment started with an introduction of the entity categories and the participants were asked to predict the outcome of collisions. In the learning phase, participants were then shown the four interactions given in (6a-d): first

³This work is based on code available at github.com/CognitiveModeling/BrainControl

⁴We tested variations in colors in a pilot-study but found no impact of the entities' color.

⁵We collected the data for this project before the Covid-19 pandemic unfolded. It is an unfortunate coincidence that our stories feature viruses as characters.

the interactions by themselves in a randomized order, then a larger scene with several interactions occurring within the scene (2x(6b), 1x(6a,d)).

We then tested the participants' knowledge of the basic event patterns in (6). Participants were shown an initial situation (e.g. a green virus facing a red spiky cell) and two possible outcomes (e.g. only the green virus and only the red cell remaining) and were asked to choose which of the outcomes was more likely. This test was done twice. If a participant chose the correct outcome both times, the experiment proceeded to the production phase. Otherwise, the participant repeated the learning phase with a new random ordering and the test phase with another pair of situation-outcome questions. The experiment then continued independent of the participant's answers in the second testing phase, however if participants did not answer at least 2 test cases correctly across both testing phases, their data was excluded.

The production phase included four possible interactions in (6), as well as their surprising counterparts, where the other character fled, i.e. a spiky cell fled from a green virus and a green virus fled from a blue cell. Figure 3 provides an example of such a surprising version, directly reversing the pattern shown in Figure 1.




Figure 3: The blue cell attacks the green virus, the virus flees

The production phase started with pretrial fillers and then presented blocks of two trials (i.e. events), with the blocks separated by fillers. As fillers, we used video clips where either one or two entities are present and moving, but no interactions happen. The pretrial fillers featured a longer sequence of such events to let the participants get used to the task. The participants were asked to follow the prompt: *Simply describe what is going on!* The participants' linguistic descriptions involved unconstrained spontaneous speech production.

There were no restrictions or priming to influence descriptions apart from the instructions shown in (7), which may have primed the usage of the verb *attack*. We explicitly avoided priming the use of NPs or reduced utterances as well as any specific grammatical constructions.

- (7) Let us first look at interactions between a virus and a cell. Notice how it doesn't matter who attacks whom.

All events consisted of two sub-events: one character attacking another character, and as a result of this collision one of the characters flees the scene. We thus anticipated that the structure of the target event will make the participants produce one of the descriptions in (8) for the event shown in Figure 1.

- (8)
- The blue cell attacks the virus and the blue cell flees.

- b. The blue cell attacks the virus and it flees.
- c. The blue cell attacks the virus and flees.

Table 1: The 2 x 2 design of the trial stimuli

	agent fleeing	patient fleeing
expected	B attacks A, B flees <i>Condition 1</i>	A attacks B, B flees <i>Condition 3</i>
surprising	A attacks B, A flees <i>Condition 2</i>	B attacks A, A flees <i>Condition 4</i>

The 8 trials adhere to a 2 x 2 design as shown in Table 1, varying the thematic role of the fleeing character in the first sub-event (agent or patient) and whether the event followed an expected pattern (Figures 1 and 2) or surprising pattern, as in Figure 3. In Figure 1, the attacker of the first sub-event flees (agent). In Figure 2 on the other hand, the one that was attacked in the first sub-event flees (patient). Participants viewed and described two scenes for each of the four conditions.


We conducted the experiment online and recruited 300 participants with US IP addresses through the online platform *Amazon Mechanical Turk*. Participants received compensation for their participation with a rate of \$10/hour.

Annotation And Analysis

Technical assistants blind to the purpose of the experiment transcribed the audio recordings that we obtained from the study participants. Participants produced a range of utterances that conformed to our expectations (e.g. 9 - 12).

- (9) The red cell just attacked the virus and it ran away.
- (10) The virus attacks the blue cell, the blue cell runs away to the left.
- (11) The green virus attacks and hits the blue cell, but then runs away.
- (12) Virus eats the red blood cell and then the virus runs.

We annotated the utterances as to whether the subject of the second verb in the event description was a full noun phrase (NP) or a pronoun. As an alternative, sometimes the speakers produced a conjoined verb phrase (8c), a possible simplified structure shown in (13). We assigned such responses to category “null” for null separate subject of the second verb (the whole verb phrase has only one subject)⁶.



⁶An alternative analysis would be to postulate coordination at the level of IP but this structure appears unlikely (Fasold & Connor-Linton, 2014)

Responses that did not fit any of the categories were assigned to a separate category X. Such descriptions either did not conform to subject verb structure, which allows the use of a full NP or a pronoun (14), or the speakers provided a summary of a whole scene rather than describing the sub-events (15).

- (14) Virus and cell attack each other.
- (15) The cell defeats the virus.

For the analysis, the utterances were grouped according to the type of referring expression that was used in the second conjunct. Utterances with pronouns and null descriptions formed the reduced category, utterances with full NPs formed their own category. We hypothesize that a full NP is justified in cases when the speakers find the event surprising. We expect this effect to be similar to an m-implicature (manner implicature): speakers choose a marked form of a linguistic expression to describe an unusual event (Levinson, 2000).

Out of 300 participants, we excluded 31 participants because of bad audio quality and 25 participants because they did not pass the learning test block. We further excluded 11 participants because their descriptions did not conform to our schema at all, mostly because they produced predictions instead of descriptions. Data from 233 participants (1864 descriptions) were included in the analysis.

Results

Participants produced both full NPs and reduced forms when both sub-events contained the same character as the agent (Conditions 1 and 2), using reduced expressions in 34.4% and full NPs in 65.6% of the cases. In the conditions where the patient of the first sub-event became the agent of the second sub-event (Conditions 3 and 4), speakers overwhelmingly produced a full NP as the subject of the second clause (97% of the time), independent of whether the event structure violated the patterns they learned or not. Since that we observed a ceiling effect in Conditions 3 and 4, we concentrate on Conditions 1 and 2, where the same character acted as the agent in both sub-events.

Prior to the main analysis, we plot the production rates for different types of responses as a function of trial order. Figure 4 indicates that trial 1 shows an unusual pattern of responses compared to other trials. Here participants were much more likely to produce reduced phrases compared to other trials, we therefore excluded trial 1 from the main analysis and proceed further with data from trials 2-8.⁷

Figure 5 shows a distribution of different referring expressions depending on condition. We fitted a generalized linear mixed-effects model using *lme4* package (Bates, Mächler, Bolker, & Walker, 2015). The type of referring expression served as a binomial dependent variable, surprise was treated

⁷Despite the different overall base rate of pronouns and NPs in the first trials, the qualitative effect of surprise remains the same there.

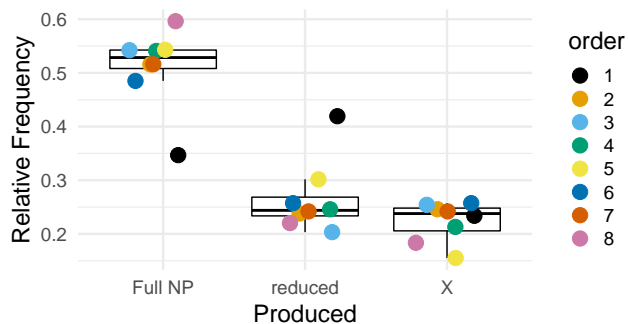


Figure 4: Production of referring expressions as a function of trial order. The first trial elicits more reduced forms overall. X denotes production choices that do not fall into either category.

as the independent variable. The random effect structure included items and subjects as random intercepts, as well as random slopes for surprise per subject. As we predicted, speakers produced more full NPs when they encountered a surprising event ($\beta = 1.417$, $SE = 0.678$, $z = 2.091$, $p < 0.05$). There was no effect of trial order in this case. Figure 6 shows the odds ratios for the two factors and their interaction, revealing the relative magnitude of the effects.⁸

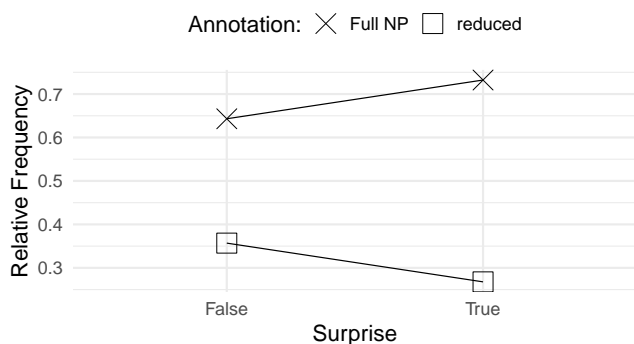


Figure 5: Full NP vs reduced form frequencies for surprising and unsurprising events, same agent in both sub-events, first trial excluded.

Even though it seems plausible that longer NPs with more discriminating adjectives could be used to mark atypical referents (Degen, Hawkins, Graf, Kreiss, & Goodman, 2020), there was no such effect in our study, possibly because the number of objects was small and color was a sufficient discriminating property ($\beta = -0.045$, $SE = 0.609$, $z = 0.074$, $p = 0.94$). There was also no effect of learning speed

⁸The odds ratio for the effect of surprise reduces to 2.09 if we include first trials into the analysis, turning the overall effect non-significant, since first trials show a qualitatively different pattern of responses.

on produced phrases ($\beta = 0.158$, $SE = 0.263$, $z = 0.601$, $p = 0.55$). For this we measured learning speed by whether participants needed to repeat training a second time.

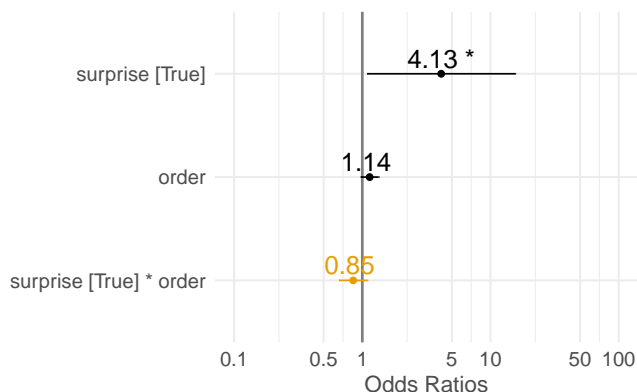


Figure 6: Effects as determined by the model

Discussion

By creating an artificial world set up, we were able to isolate the effect of knowledge of event structures from the distributional regularities of the language stream. Therefore, our work directly targets the effect of world knowledge on the production of referential expressions, and language production more broadly.

In line with the Gricean maxim of manner, to be clear but brief (Grice, 1989), speakers typically refrain from using a full noun phrase when referring to a previously mentioned entity. In fact, repeating a name causes a disruption in discourse coherence and slows down reading times. This effect is known as “Repeated-name Penalty” (Gordon, Grosz, & Gilliom, 1993). Interestingly, the penalty applies if the antecedent of a repeated reference occurs in subject position but not in other positions, since the subject position marks entities salient in discourse. Gor (2020) summarizes evidence from a number of studies showing that more informative nominals, such as full NPs in our case, often refer to less salient entities in discourse, while less informative references (pronouns and possibly null subjects) refer to more salient entities. It is therefore not surprising that in our data, speakers were more likely to use full NPs, or in other words, they repeated a name, when the antecedent occurred in the object position in the previous clause. Moreover, in this condition, the character switched its thematic role from being a patient in the first sub-event, to becoming an agent in the second sub-event. The need to alert the listener to the change of role justifies using a full NP in the second clause and results in the ceiling effect in the use of full NPs—when the object in the first sentence became the agentive subject in the second.

The standard Repeated-name Penalty effect has been described for sentences with a single potential antecedent for a repeated name. However, it appears that the effect holds to a certain extent even when there are multiple competing

antecedents in the previous clause. Our data reveals that the effect is stronger when speakers describe an event where the agent engages in familiar interactions with the patient. In such cases, a less informative nominal (a pronoun or a null subject) should normally be preferable for a predictable agent in the second event. However an unexpected turn in the unfolding of an event progression calls for a referring expression that could boost the salience of the agent to the listener. These findings are in line with the uniform information density hypothesis (Jaeger, 2010; Levy & Jaeger, 2007), that predicts that full NPs should be preferred over pronouns if the NP referent is less predictable. In that situation, the NP carries essential information about the referent and smoothes out information density. We can then reformulate the Gricean requirement for ambiguity avoidance in information theoretic terms: ambiguity arises when the meaning of an expression cannot be straightforwardly predicted from prior context. Therefore, if a pronoun could pick out multiple potential referents, and the context (or world knowledge) does not provide sufficient cues to disambiguate, speakers should opt for a more informative referring expression.

Jaeger (2010) viewed ambiguity avoidance and smoothing information density as distinct hypotheses and showed that in his study participants were *not* more likely to use a disambiguating complementizer “that” when a sentence was more likely to lead to a garden-path. Rather, what predicted the use of a complementizer was a desire to make information density uniform: adding an extra word allowed to spread out a more informative (less predictable) content across more words, therefore equalizing information density per word. The case of referential ambiguity we are considering in this paper offers another angle to explore the interaction of ambiguity avoidance and information density. Unlike in the case of structural ambiguity caused by the absence of “that”, the absence of an overt subject (when grammatically allowed) is not what causes the ambiguity. Rather it is the informativity of the referring expression itself that is crucial: choosing a pronoun that can be co-referenced with multiple potential antecedents induces referential ambiguity. Our data reveals that if the identity of the referent is predictable and can be easily recovered by relying on knowledge of event structures, speakers are more likely to use a pronoun or drop the subject and opt for a conjoined verb phrase. Speakers do not need to avoid potential referential ambiguity because the reference can be easily and correctly disambiguated. What makes speakers utter a full NP in the case of a surprising agent of the second sub-event is not ambiguity avoidance as a strategy, but rather avoiding misleading reference, since it can compromise sentence comprehension. Note that this information-theoretic approach does not take the potential referential ambiguity of a pronoun per se into account, but rather how likely the listener is to recover the intended referent given the actual knowledge about the world.

It remains an open question whether speakers intentionally avoid misunderstandings. We view the preference to use a full

NP when talking about surprising situations and event progressions as a type of automatic behavior (Lieberman, 2007; Dale et al., 2018), which developed through monitoring the effect of previous interactions. While it is aimed at simplifying the comprehension process, it is not necessarily driven by a consideration of particular beliefs of the listener, but may rely on the speaker’s own belief of what event structures she finds likely.

When generating referring expressions speakers perform intricate mental calculations balancing the effect of multiple linguistic and extra-linguistic factors. Our goal in this paper was to investigate the nature of these interactions, and spell out precisely which components of the world knowledge affect the choice of utterances. Alongside with the theoretical contribution, this paper documents an experimental setup that may serve as a general paradigm to examine the influence of world knowledge on language further. Participants adapted to the task surprisingly well. The majority of utterances produced were in line with our expectations, without priming utterance structure or prescribing parts of the utterances. Thus, we believe that the paradigm may be useful to further study intricate interactions between world knowledge, event prediction, and language. Here, we focused on the role of predictability of an event progression and demonstrated that discriminating referring expressions (full NPs) are preferred over reduced forms to mark the agent of a surprising event transition. This effect illustrates the interaction of linguistic constraints on reference with world knowledge.

Acknowledgements

We would like to thank Tizian Thieringer, Jannis Strecker and Loredana Columbo for their help in setting up the experiment and annotating data. We would also like to thank the anonymous CogSci 2021 reviewers for their insightful comments and literature pointers. This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)—Project number 198647426. Martin Butz is a member of the Machine Learning Cluster of Excellence, EXC number 2064/1 – Project number 390727645.

References

- Baldwin, D. A., & Kosie, J. E. (2020). How does the mind render streaming experience as events? *Topics in Cognitive Science*. doi: 10.1111/tops.12502
- Barton, E. L. (1998). The grammar of telegraphic structures: sentential and nonsentential derivation. *Journal of English Linguistics*, 26(1), 37–67.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48.
- Bunger, A., Papafragou, A., & Trueswell, J. C. (2013). Event structure influences language production: Evidence from structural priming in motion event description. *Journal of Memory and Language*, 69(3), 299–323.
- Butz, M. V., Achimova, A., Bilkey, D., & Knott, A. (2021). Event-predictive cognition: A root for conceptual human

- thought. *Topics in Cognitive Science*, 13, 10-24. doi: 10.1111/tops.12522
- Dale, R., Galati, A., Alviar, C., Contreras Kallens, P., Ramirez-Aristizabal, A. G., Tabatabaeian, M., & Vinson, D. W. (2018). Interacting timescales in perspective-taking. *Frontiers in psychology*, 9, 1278.
- Davis, E. (2019). *Collection of Winograd schemas*. www.cs.nyu.edu/faculty/davise/papers/WinogradSchemas/WSCollection.html.
- Degen, J., Hawkins, R. D., Graf, C., Kreiss, E., & Goodman, N. D. (2020). When redundancy is useful: A Bayesian approach to “overinformative” referring expressions. *Psychological Review*, 127(4), 591–621.
- Fasold, R. W., & Connor-Linton, J. (2014). *An introduction to language and linguistics*. Cambridge university press.
- Ferreira, V. S. (2019). A mechanistic framework for explaining audience design in language production. *Annual Review of Psychology*, 70(1).
- Ferreira, V. S., Slevc, L. R., & Rogers, E. S. (2005). How do speakers avoid ambiguous linguistic expressions? *Cognition*, 96(3), 263–284.
- Fukumura, K., & van Gompel, R. P. G. (2010). Choosing anaphoric expressions: Do people take into account likelihood of reference? *Journal of Memory and Language*, 62(1), 52–66.
- Fukumura, K., & van Gompel, R. P. G. (2012). Producing pronouns and definite noun phrases: Do speakers use the addressee’s discourse model? *Cognitive Science*, 36(7), 1289–1311.
- Fukumura, K., van Gompel, R. P. G., Harley, T., & Pickering, M. J. (2011). How does similarity-based interference affect the choice of referring expression? *Journal of Memory and Language*, 65(3), 331–344.
- Galati, A., & Brennan, S. E. (2010). Attenuating information in spoken communication: For the speaker, or for the addressee? *Journal of Memory and Language*, 62(1), 35–51.
- Gor, V. (2020). *Experimental investigations of principle c at the syntax-pragmatics interface*. Unpublished doctoral dissertation, Rutgers University-School of Graduate Studies.
- Gordon, P. C., Grosz, B. J., & Gilliom, L. A. (1993). Pronouns, names, and the centering of attention in discourse. *Cognitive science*, 17(3), 311–347.
- Grice, P. (1989). *Studying in the way of words*. Cambridge, MA: Harvard University Press.
- Haegeman, L. (2013). The syntax of registers: Diary subject omission and the privilege of the root. *Lingua*, 130, 88–110.
- Haywood, S. L., Pickering, M. J., & Branigan, H. P. (2005). Do speakers avoid ambiguities during dialogue? *Psychological Science*, 16(5), 362–366.
- Jaeger, T. F. (2010). Redundancy and reduction: Speakers manage syntactic information density. *Cognitive Psychology*, 61(1), 23–62.
- Jaeger, T. F., & Buz, E. (2017). Signal reduction and linguistic encoding. In E. M. Fernández & H. S. Cairns (Eds.), *The handbook of psycholinguistics* (1st ed., pp. 38–81). Wiley.
- Kingsbury, D. (1968). *Manipulating the amount of information obtained from a person giving directions*. Unpublished doctoral dissertation, Harvard University.
- Kuperberg, G. R. (2020). Tea with milk? A hierarchical generative framework of sequential event comprehension. *Topics in Cognitive Science*. doi: 10.1111/tops.12518
- Kurumada, C., & Jaeger, T. F. (2015). Communicative efficiency in language production: Optional case-marking in Japanese. *Journal of Memory and Language*, 83, 152–178.
- Lam, T. Q., & Watson, D. G. (2010). Repetition is easy: Why repeated referents have reduced prominence. *Memory & cognition*, 38(8), 1137–1146.
- Levesque, H. J. (2014). On our best behaviour. *Artificial Intelligence*, 212, 27–35.
- Levinson, S. C. (2000). *Presumptive meanings: The theory of generalized conversational implicature*. MIT press.
- Levy, R., & Jaeger, T. F. (2007). Speakers optimize information density through syntactic reduction. *Advances in neural information processing systems*, 19, 849.
- Lieberman, M. D. (2007). Social cognitive neuroscience: a review of core processes. *Annu. Rev. Psychol.*, 58, 259–289.
- Rohde, H., & Kehler, A. (2014). Grammatical and information-structural influences on pronoun production. *Language, Cognition and Neuroscience*, 29(8), 912–927.
- Rosa, E. C., & Arnold, J. E. (2017). Predictability affects production: Thematic roles can affect reference form selection. *Journal of Memory and Language*, 94, 43–60.
- Schrod, F., Kneissler, J., Ehrenfeld, S., & Butz, M. V. (2017). Mario becomes cognitive. *Topics in Cognitive Science*, 9(2), 343–373.
- Schrod, F., Röhm, Y., & Butz, M. V. (2017). An event-schematic, cooperative, cognitive architecture plays super mario. In R. Chrisley, V. C. Müller, Y. Sandamirskaya, & M. Vincze (Eds.), *Proceedings of the EUCognition meeting (european society for cognitive systems) “cognitive robot architectures”* (pp. 10 – 15).
- Tily, H., & Piantadosi, S. (2009). Refer efficiently: Use less informative expressions for more predictable meanings. In *Proceedings of the workshop on the production of referring expressions: Bridging the gap between computational and empirical approaches to reference*.
- Willits, J. A., Amato, M. S., & MacDonald, M. C. (2015). Language knowledge and event knowledge in language use. *Cognitive psychology*, 78, 1–27.
- Winograd, T. (1972). Understanding natural language. *Cognitive Psychology*, 3(1), 1–191.