

Conceptualisation in reference production: Probabilistic modelling and experimental testing

Roger P.G. van Gompel

Kees van Deemter

Albert Gatt

Rick Snoeren

Emiel J. Krahmer

Psychological Review. © American Psychological Association, 2018. This paper is not the copy of record and may not exactly replicate the authoritative document published in the APA journal.

Author Note

Roger P.G. van Gompel, School of Social Sciences, University of Dundee; Kees van Deemter, Department of Information and Computing Sciences, Utrecht University; Albert Gatt, Institute of Linguistics and Language Technology, University of Malta and Department of Communication and Cognition, Tilburg University; Rick Snoeren, Department of Communication and Cognition, Tilburg University; Emiel Krahmer, Department of Communication and Cognition, Tilburg University.

Parts of this work were previously reported at Architectures and Mechanisms of Language Processing (AMLaP) in 2012 (Riva del Garda) and 2016 (Bilbao), at the PRE-Cogsci meeting in 2013 (Berlin) and at REFNET Round Table Event 2016 (Aberdeen).

Address all correspondence to Roger van Gompel, Psychology, School of Social Sciences, University of Dundee, Dundee, DD1 4HN, Scotland, United Kingdom. Contact: r.p.g.vangompel@dundee.ac.uk

Abstract

In psycholinguistics, there has been relatively little work investigating conceptualisation – how speakers decide which concepts to express. This contrasts with work in natural language generation (NLG), a subfield of AI, where much research has explored content determination during the generation of referring expressions. Existing NLG algorithms for conceptualisation during reference production do not fully explain previous psycholinguistic results, so we developed new models that we tested in three language production experiments.

In our experiments, participants described target objects to another participant. In Experiment 1, either its size, its colour, or both its size and colour distinguished the target from all distractor objects; in Experiment 2, either colour, type or both colour and type distinguished it from all distractors; In Experiment 3, either colour, size or the border around the object distinguished the target. We tested how well the different models fit the distribution of description types (e.g., “small candle”, “grey candle”, “small grey candle”) that participants produced.

Across these experiments, the PRO model provided the best fit. In this model, speakers first choose a property that rules out all distractors. If there is more than one such property, then they probabilistically choose one based on a preference for that property. Next, they sometimes add another property, with the probability again determined by its preference and speakers’ eagerness to overspecify.

Keywords: Reference production; referring expressions; conceptualisation; overspecification; computational models.

Disclosures and acknowledgements

RvG acknowledges support from ESRC award ES/P001866/1. EK and AG received financial support from The Netherlands Organization for Scientific Research, via a Vici Grant (NWO Grant 27770007), and EK additionally via NWO Grant PR-14-87. KvD was supported by the RefNet project (EPSRC award EP/J019615/1).

All authors contributed significantly to the design of the experiments, the modelling, analyses, interpretation of the data and the writing.

None of the authors have any conflicts of interest that that may have inappropriately impacted or influenced the research and interpretation of the findings.

We would like to thank Ben Vincent for his suggestions on Bayesian model testing, as well as Jenna Lyons and Manon Yassa for their help with the experiments.

Conceptualisation in reference production:

Probabilistic modelling and experimental testing

Human speech production — going from “intention to articulation” (Levelt, 1989), or from “mind to mouth” (Bock, 1995) — is a complex process, yet our understanding of it is increasing (Goldrick, Ferreira, & Miozzo, 2014). Many researchers agree that speech production involves different stages (e.g., Dell, 1986; Garrett, 1984; Levelt, 1989; Levelt, Roelofs, Meyer, 1999; Vigliocco & Hartsuiker, 2002). A speaker first has to decide what she wants to say¹, a decision referred to as *conceptual preparation* or *conceptualisation*, which results in a *preverbal message*. The second stage, often referred to as *formulation*, involves lexical access and planning of the structure of the utterance. Finally, during *phonological encoding*, the utterance plan is phonologically encoded and articulated.

Most experimental studies and psycholinguistic models have concentrated on the last two stages of speech production, and as a result these are better understood than the first, conceptualisation stage. It is interesting, however, that conceptualisation is an important research topic in a different field, known as *natural language generation* (NLG, see e.g., Gatt & Krahmer, 2018; Mellish et al., 2006; Reiter & Dale, 2000). NLG is a subdomain of Artificial Intelligence dedicated to the automatic conversion of data into text. The goal of the current paper is to see whether NLG conceptualisation algorithms can further our understanding of how human speakers conceptualise their utterances.

One NLG area in which conceptualisation has been intensively studied is the generation of referring expressions, especially noun phrase descriptions of objects, such as “the grey candle”. This line of work dates back at least to Winograd (1972), and more systematically, from Appelt (1985) onwards (see Krahmer & Van Deemter, 2012, and Van Deemter 2016 for recent discussions). The focus on such descriptions is understandable, given that reference

has been argued to be central to communication (Johnson-Laird, 1983). NLG reference generation systems usually distinguish between *content determination* and *linguistic realisation*, similar to the distinction between conceptualisation and formulation in psycholinguistic models. During the content determination phase, the NLG system decides which attributes of a referent (like its colour, size or type)² to include in the description. Next, during realisation, it decides how to express selected attributes in words and phrases. NLG models have often focused on the selection of attributes in full noun phrases that aim to single out the target from its distractors in one go (*one-shot descriptions*), and this is what we shall do in this paper as well.

Here, we ask to what extent these NLG models can serve as a stepping-stone towards computational models of *human* conceptualisation during the production of referring expressions. As we will demonstrate, it is possible to derive from NLG models precise, quantitative predictions about the referring expressions that humans produce. We contrast the predictions of different NLG models in a set of language production experiments, to further our understanding of human conceptualisation of reference production.

Computational Models of Conceptualisation in Reference

Computational models of content selection address language production from an applied perspective, in the context of larger NLG systems that automatically convert data into text (Mellish et al., 2006; Reiter & Dale, 2000). Such systems are useful in the automatic generation of, for example, textual weather reports (Goldberg, Driedger, & Kittredge, 1994), and summaries of patient information in intensive care (Portet et al., 2009).

Typically, content selection models determine which set of properties distinguishes a target object from the distractors in a given domain, assuming that we know the relevant properties of the domain objects. For example, a content selection model has to determine

which properties to select to refer to a target such as the circled object in Fig. 1 in order to distinguish it from its distractors (i.e., the two other objects); it could either select the property “red” or “small” (or both) for inclusion in the referring expression. In many cases, NLG algorithms start from handcrafted domain models, but in recent years they increasingly start from raw images (e.g., Kazemzadeh, Ordonez, Matten, & Berg, 2014; Mao et al., 2016).

Insert Fig. 1 about here

Early NLG models of content selection. Algorithms differ primarily in the way they decide on the inclusion of attributes. In early NLG models, the focus was on computing the shortest possible distinguishing description that distinguishes the target from all its distractors. This is in line with Grice’s (1975) influential *maxim of quantity*, which states that speakers should make their contribution as informative as required but not more informative. The full brevity algorithm (Dale, 1989), for example, first checks whether there is a single property of the target that rules out all distractors. If this fails, it considers all possible combinations of two properties, and so on. Unfortunately, this approach suffers from two problems, making it less suitable as a psycholinguistic model of conceptualization. First, it is computationally expensive, which implies that producing a description could take a long time in larger domains and for longer descriptions (Appelt, 1985; Dale, 1989; Dale & Reiter, 1995). Although this would not be an issue with a small number of attributes and distractors (as in many psycholinguistic experiments), it would face difficulties with more complex scenes. Second, the algorithm never produces overspecified descriptions, that is, descriptions that include more properties than necessary for unique identification of the target. An example of this is “small red broom” in Fig. 1.

There is much evidence from psycholinguistic experiments that human speakers often produce overspecified descriptions. For example, Pechmann (1989) asked speakers to refer to a target object (a car, ball or chair, for example) among three, five or seven distractors, differing in their type, colour and/or size. Pechmann found that while only 4% of the descriptions speakers produced were underspecified, 21% of the descriptions included at least one property that was not required for identification by the addressee. Others report comparable low levels of underspecification (e.g., Koolen, Gatt, Goudbeek & Kraemer, 2011). Percentages of overspecified descriptions vary from one study to the next, but are always substantial. Pechmann (1989), as noted above, reported 21% overspecified descriptions, but other studies found percentages between 40% and 80% (e.g., Koolen, Goudbeek, & Kraemer, 2013; Engelhardt, Bailey, & Ferreira, 2006; Engelhardt, Demiral, & Ferreira, 2011; Maes, Arts, & Noordman, 2004, Tarenskeen, Broersma, & Geurts, 2015). Several factors may influence the rate of overspecification, for instance, speakers are more likely to overspecify in situations where misunderstandings would be costly (Arts, 2004; Arts, Maes, Noordman, & Jansen, 2011), when instructions make participants aware of potential communication breakdown (Rubio-Fernandez, 2016), when domains are spatially complex (Paraboni & Van Deemter, 2014), and when the distractors in the scene are more varied (e.g., they differ in type or colour, Koolen et al., 2013; Rubio-Fernandez, 2016).

Another early content determination model is the greedy heuristic algorithm (Dale, 1989, 1992). This algorithm selects attributes based on how many distractors they rule out, which is usually referred to as the *discriminatory power* or *utility* of an attribute. The algorithm iteratively selects that property which rules out most of the distractors not previously ruled out, incrementally extending the description based on what property has most discriminatory power at each stage. It has a number of interesting properties from a psycholinguistic perspective. First, it may give rise to overspecified descriptions because it is

incremental: When a property (e.g., colour) becomes redundant following the selection of a later property (e.g., size), the initial property cannot be “unselected”. Note that this notion of incrementality is different from that at the surface, word level: Although colour may be selected before size, this does not mean that colour information is necessarily realised before size at the word level. In addition, since the algorithm does not try to compute the shortest possible description, it is computationally considerably more efficient than the full brevity algorithm (Dale & Reiter, 1995). A limitation of the model is that ties are not resolved: applying the standard greedy algorithm to the example in Fig. 1 results in two possible descriptions —“the red one” and “the small one”— since both rule out the same distractors.

The rational speech act model. A more recent model in which discriminatory power plays a critical role is Frank and Goodman’s (2012) rational speech act model, which aims to account for both the production and comprehension of reference (see also Goodman & Frank, 2016; Goodman & Stuhlmüller, 2013). It assumes that language users are rational actors: Speakers try to produce referring expressions that are optimally useful for listeners, and listeners, in turn, assume the speaker is maximally helpful.

During reference production, the probability that a speaker chooses an attribute is directly proportional to the informativeness of this property, formalised as a function of surprisal, which is an information-theoretic measure expressing (in this case) how much a given property reduces the listener’s uncertainty about which object is the target. For example, if one attribute rules out twice as many distractors as another attribute, then the most discriminatory attribute is chosen twice as often as the less discriminatory one.

Frank and Goodman (2012) showed that their model made very accurate predictions for an experiment in which participants were asked to bet which single property (e.g., “blue” or “circle”) would be most likely to be selected by a speaker referring to a target (e.g., a blue circle) that was presented together with a distractor set (e.g., a blue square and a green

square). However, as it stands, the model does not select multiple properties and therefore cannot account for psycholinguistic evidence from reference production tasks that speakers often use multiple properties in overspecified descriptions. In addition, the use of a single property is often not sufficient to distinguish the target from all distractors; yet, as we mentioned above, speakers rarely use underspecified descriptions. The model needs further development to account for these experimental findings. Nevertheless, an attractive feature of the model is that it is *non-deterministic*: It assumes that, if there is more than one distinguishing property (a property that holds true of the referent but is false of at least one of the distractors), each of these properties is selected some proportion of the time. This contrasts with many other computational algorithms that are deterministic, that is, they always generate the same referring expression in a particular situation (e.g., in an experimental condition). As we have argued in Van Deemter, Gatt, Van Gompel, and Krahmer (2012), reference production is normally non-deterministic; different speakers may produce different descriptions in a given context. In fact, speakers themselves are not deterministic either, but may produce different descriptions in the same context on different occasions as well (Viethen & Dale, 2010).

The incremental algorithm. Arguably the most influential computational model of reference generation has been the incremental algorithm proposed by Dale and Reiter (1995). It works from the assumption that some attributes are preferred over others (cf. Pechmann, 1989). The algorithm presupposes the existence of a fixed preference order for a given domain, which is a complete ranking of all relevant attributes. Such a preference order can be used to model the relative importance of different attributes to the communicative task, their salience, or their (a)typicality for the objects in question. The algorithm iterates through this preference order, adding an attribute if it helps ruling out a distractor not previously ruled out and terminating when a set of attributes has been selected that collectively rules out all

distractors. Like the greedy heuristic algorithm, it can produce overspecified descriptions due to its incremental nature.

To see how the incremental algorithm works, let us consider the examples in Fig. 2, which will be tested as three different conditions in the current study. Fig. 2a is similar to Fig. 1: Both colour and size are fully discriminatory in that each attribute alone distinguishes the circled target from its distractors. In Fig. 2b, size is fully discriminatory, whereas colour is only partially discriminating because it rules out only distractor. Finally, in Fig. 2c, colour is fully discriminatory whereas size is only partially discriminatory.

Insert Fig. 2 about here

Let us say that the preference order for our examples is $\langle \text{type, colour, size} \rangle$. As a result, the incremental algorithm first considers type in all examples; our target is a candle, but so are both distractors, hence the algorithm ignores this attribute and considers colour next. In Fig. 2a, the target is grey, while both distractors have different colours. Hence, the algorithm selects colour, and since this rules out all distractors, the search stops and the algorithm does not consider size. There is a final check whether the type attribute was selected, and if not, the algorithm includes it anyway, since (so the reasoning goes) an object description without type information (usually expressed by the head noun) is undesirable. The algorithm then terminates with a set of properties that could be realised linguistically as “the grey candle”. Thus, here, the algorithm generates a minimally specified description.

This is different in Fig. 2b. Following the same reasoning as before, type is not selected, but colour is (because it rules out the red candle). Since this does not rule out *all* distractors, the model now considers size, which is also selected (since it rules out the large grey candle). All distractors are ruled out, and the algorithm terminates with a set of

properties that could be realised as “the small grey candle”. This is an overspecified description, since “the small candle” would have been sufficient to rule out both distractors. In other words, the selection of size makes the previous inclusion of colour redundant, but due to the incremental nature of the model, colour is realised nonetheless.

Finally, in Fig. 2c, colour is also first selected and because it rules out all distractors, size is not considered. As a result, the algorithm ends up with properties that could be realised as “the grey candle”, which is a minimally specified expression.

The incremental algorithm was strongly influenced by psycholinguistic work that had shown that speakers frequently overspecify and have a preference for certain attributes such as colour over others, such as size. Subsequent experimental evaluations also reveal that the incremental algorithm produces descriptions that are more like those produced by human speakers than the full brevity and greedy heuristic algorithms, if an appropriate preference order is assumed (Van Deemter, Gatt, Van der Sluis, & Power, 2012). However, one issue with the incremental algorithm is that it is deterministic. Given a preference order and a domain, the incremental algorithm always predicts the same output in a specific condition. The candles in Figs. 2a and 2c are always referred to using the minimal description “the grey candle”, while the one in Fig. 2b is always referred to with the overspecified description “the large grey candle”. As we have discussed above, this prediction is not in accordance with the psycholinguistic data described above, which reveal that speakers produce descriptions that sometimes are and sometimes are not overspecified.

Developing more psychologically plausible conceptualisation accounts

The discussion above indicates that current computational models of reference generation are not yet completely psychologically realistic, as they cannot account for the full range of experimental evidence from human speakers. We therefore set out to improve these models,

taking the incremental algorithm as the starting point, and then to test their predictions in a series of reference production experiments.

Non-deterministic incremental algorithms. As we have seen, a major issue with the incremental algorithm is its deterministic nature, which is due to the complete, fixed ordering of attributes in the preference list. In Van Deemter et al. (2012), we proposed (but did not test) a variant of the incremental algorithm (which we call the “non-deterministic incremental algorithm”, or non-deterministic IA for brevity), which non-deterministically varies the order in which attributes are checked.

To see how this works, reconsider Figs. 2a-c. The algorithm checks colour before size with a probability c and it checks size before colour the rest of the time, with a probability $1-c$. Thus, c represents the colour-size preference. In Fig. 2a, in those cases where colour is checked first, the algorithm includes this in the description. Because it rules out all distractors, size is not subsequently checked. Assuming (as before) that type is always included, this results in a description such as “the grey candle” with a probability of c . The rest of the time (with a probability of $1-c$), size is checked first. Because this also rules out all distractors, colour will not be added, resulting in “the small candle” with a probability of $1-c$. The algorithm does not generate overspecified descriptions (“the small grey candle”) in Fig. 2a, because once a fully distinguishing description has been found, it terminates. In Fig. 2b, speakers also first select colour with a probability of c , but because this does not result in a fully discriminatory expression, they add size, resulting in the overspecified “small grey candle” with a probability of c . The rest of the time, they first choose size, and because this *is* fully discriminatory, this results in “small candle” with a probability of $1-c$. Finally, in Fig. 2c, initial selection of colour results in a fully discriminatory expression, so it produces “grey candle” with a probability of c , but when they first choose size, they need to add colour,

resulting in “small grey candle” with a probability of $1-c$. Thus, unlike the original IA, the non-deterministic IA generates overspecified descriptions in Fig. 2c some proportion of time.

Insert Table 1 about here

Table 1 shows the formulas for generating the different non-deterministic IA descriptions. Assuming we can determine the colour-size preference c , the non-deterministic IA makes quantitative predictions about cases such as those in Fig. 2. For example, if the colour-size preference is .8, then “grey candle” in Figs. 2a and 2c and “small grey candle” in Fig. 3b should be produced in 80% of cases, while “small candle” in Figs. 2a and 2b and “small grey candle” in Fig. 2c should be generated in 20% of cases. Thus, if we can determine the colour-size preference c from the descriptions that speakers produce in one condition, we can predict the descriptions they should produce in another condition.

In sum, the non-deterministic IA appears to be a promising candidate for modelling conceptualisation during the human production of descriptions: It retains all positive aspects of the original incremental algorithm (incrementality, overspecification), and also captures the inherent non-determinism of human speakers. Until now, the quantitative predictions of the non-deterministic IA have not been tested, so we did this in the current study.

However, as noted above, the non-deterministic IA predicts no overspecification in Fig. 2a, when both attributes are fully discriminatory. However, Goudbeek and Kraemer (2012) and Viethen, Van Vessel, Goudbeek and Kraemer (2017) showed that overspecification occurred between 11 and 40% of the time in such situations. These results suggest that the choice of an overspecified rather than a minimally specified description in Fig. 2a is also non-deterministic; speakers do not always produce minimal descriptions. We therefore also tested a modified version of the non-deterministic IA.

In this modified algorithm, speakers sometimes add a second attribute even if the first attribute they selected is fully discriminatory. Fig. 3 shows the decision tree for the three conditions in Fig. 2 and Table 1 shows the formulas. The likelihood with which a particular description is chosen can be worked out by starting at “select first attribute” and following the arrows to the particular description. In the colour-or-size condition (Fig. 2a), speakers first choose colour or size depending on the preference (respectively c and $1-c$). Either attribute is fully discriminating, but to account for previous evidence that speakers sometimes overspecify in this condition, the modified non-deterministic IA predicts that, in some proportion of trials, speakers add a second attribute. We assume that the probability that this happens is determined by the preference for this attribute: c for colour and $1-c$ for size. The probability of first selecting colour and then adding size in the colour-or-size condition is the product of these two selections: $c*(1-c)$. But because speakers can also first select size and then add colour $(1-c)*c$, the total probability of colour-and-size descriptions is the sum of the two routes through the decision tree: $c*(1-c)+(1-c)*c$. In contrast, there is only one route to a colour-only description ($c*c$) and to a size-only $(1-c)*(1-c)$ description. In a similar way, Fig. 3 can be used to derive formulas for the colour-only and size-only conditions. Critically, the algorithm predicts that after selecting colour in the colour-only condition, speakers sometimes add size and after selecting size in the size-only condition, they sometimes add colour.

Insert Fig. 3 about here

PRO: A model combining discriminatory power with preferences. So far, we have discussed two classes of potential conceptualisation models for human reference production: models relying either on preferences, such as the incremental algorithm, or on discriminatory

power, such as the greedy algorithm and the rational speech act theory. Recently, we have proposed a probabilistic conceptualisation model that does both (Gatt, Van Gompel, Van Deemter, & Kraemer, 2013; Van Gompel, Gatt, Kraemer, & Van Deemter, 2012). We termed this model the Probabilistic Referential Overspecification model (PRO), because the aims of this model are to account (1) for the fact that human reference production is probabilistic and (2) for the finding that humans frequently overspecify. In PRO, discriminatory power is assumed to play a role if an attribute rules out all distractors in one fell swoop. If there is a single fully discriminatory attribute, such as size in Fig. 2b or colour in Fig. 2c, then this attribute is always chosen first. However, when there are two fully discriminating attributes (as in Fig. 2a) preference comes into play: The model non-deterministically selects one of the fully-discriminating attributes, with the chance that a particular attribute is chosen depending on its preference. For example, the chance that colour is selected can be c , whereas the chance that size is chosen is $1-c$, as in the non-deterministic IA. If a single attribute is fully discriminating, the model could terminate after selecting either colour or size. But PRO assumes that the decision to stop or to overspecify is made non-deterministically as well. This is the second place where preference parameter c has an effect: The more preferred an attribute is, the more likely that it is added. In addition, we assume a second, “overspecification eagerness” parameter e that affects the likelihood that speakers overspecify. This parameter is motivated by evidence that the likelihood of overspecification is affected by the speaker’s task and by the type of speaker. For example, as we have seen, Arts (2004) and Arts et al. (2011) found that overspecification was more common in fault-critical situations, Rubio-Fernandez (2016) found that the frequency of overspecification depended on the instructions to the participants, and Deutsch and Pechmann (1982) found that adults overspecify more often than children. Note that because the overspecification eagerness parameter e is not a probability, but a parameter that

influences the probability with which speakers overspecify, it can be either positive or negative: Positive values make overspecification more likely than it would be if this parameter were not included, whereas negative values make it less likely.

PRO can be formalised as an algorithm, shown in Appendix 1. This algorithm allows us to make predictions for a wide range of situations with different targets and distractors. The present paper will focus on decision trees derived from the algorithm, which exemplify the workings of the PRO algorithm. Fig. 4 shows the decision tree for the conditions in Fig. 2. To work out the probabilities for referring expressions we start at “select first attribute”. Next, we follow the arrows to the appropriate condition and then the particular expression. For example, the probability that a speaker produces “small grey candle” in the colour-only fully distinguishing condition (Fig. 2c) is the product of the chance that she first selects the fully distinguishing attribute colour (1: the speaker always selects it) and the probability that size is added, with the latter probability being determined by the preference for size ($1-c$) plus the eagerness to overspecify (e). The probability that a speaker produces “grey candle” in the same condition is the product of first selecting colour (1) and subsequently not adding size, with the latter being the preference for *not* size (c) minus the overspecification preference e . Note that because the probabilities need to add up to one, the value for e is either added to or subtracted from the colour or size preference (rather than e.g., multiplied or divided by it). Also note that in the colour-or-size condition (Fig. 2a), colour-and-size descriptions can arise because speakers can first select colour and then add size $c(1-c+e)$ or they can first select size and then colour $(1-c)(c+e)$. As a result, the likelihood that a speaker chooses a colour and size description in this condition is $c(1-c+e)+(1-c)(c+e)$. Table 1 shows a summary of the formulas.

Insert Fig. 4 about here

PRO relies on insights from other models, but crucially combines them in a novel way. In common with the greedy heuristic algorithm and the rational speech act theory, it assumes that discriminatory power plays a role in concept selection, although in contrast to these models, it only plays a role when an attribute rules out *all* distractors and unlike in the greedy heuristic, ties of such fully discriminatory attributes are resolved based on preferences. It builds on the incremental algorithm because it assumes that attributes' preferences play an important role and that concept selection is incremental in nature. Like the non-deterministic IA and rational speech act theory, PRO assumes that concept selection is probabilistic, and finally, the assumption that attribute selection does not always stop once a fully discriminating description has been found is shared with the modified non-deterministic IA.

Overview of the Experiments

We report three experiments that contrasted the predictions of PRO and the different variants of the non-deterministic incremental algorithm. The conditions in the experiments were specifically designed so we could accurately determine the parameter values in the models and derive quantitative predictions from them.

Experiment 1 tested scenarios such as in Fig. 2, where colour and size discriminated the target from the distractors. To test how well the predictions of the models generalised to a different combination of attributes, Experiment 2 investigated scenarios where the target object's type and colour were distinguishing features. We conducted the experiments in English and Dutch, to determine whether the models can account for the choice of referring expressions in both languages. Finally, in Experiment 3, three attributes discriminated the

target from the distractors, making the production of reference potentially more complex because speakers have a larger choice of attributes and can produce longer expressions.

The models under consideration all deal with reference in its simplest form, in that they do not account for the production of, say, complex descriptions including negation (“not the broom”) or reference to sets (“two candles”), so we used simple scenes such as in Fig. 2, where complex descriptions would be unlikely. In order to minimise the role of attention, we also used a small number of distractors: In domains with larger distractor set sizes, speakers may not see or may not pay attention to all distractors, so they would not be able to take them into account during reference production. Note that it is relatively common to have a limited number of objects in visual attention, for example the objects on a table in front of you.

Speakers described the target pictures to an addressee, who had to select the picture that the speaker described. Both participants saw the same pictures, but in different positions so that the speaker could not use the picture location (“middle candle”). In this way, we could investigate how speakers produced reference for an addressee while controlling for other factors that affect interactive dialogue such as speaker adjustment due to feedback from the addressee, alignment with the addressee and the influence of the preceding discourse (priming from a previous description or anaphoric reference). In future research, it will be helpful to test more complex situations with a more interactive task, but we believe it is important to examine simple situations first, because the models should make particularly precise predictions in such cases.

Experiment 1

Method

Participants. Thirty pairs of undergraduate students, all native speakers from the University of Dundee, took part in the English experiment. In the Dutch experiment, we also

tested 30 participant pairs; all participants were native speakers of Dutch from Tilburg University. All speakers reported normal colour vision.

Materials. The materials in both the English and Dutch experiment consisted of 36 experimental item sets, each consisting of three conditions. Fig. 2 gives an example. In each condition, three pictures of objects (e.g., three candles) that only differed from each other in size or colour (or both size and colour) were shown. We used blue, grey, red and green as colours and the size of the smaller objects was two thirds of the larger objects. The target object was circled; its position was counterbalanced across items. The pictures were constructed using a version of the Snodgrass and Vanderwart (1980) line drawings with colour and texture (Rossion & Pourtois, 2004).

In the colour-or-size condition, mention of either the target's size or colour was sufficient for the addressee to identify the target, because the target differed from both distractors in both size and colour. In the size-only condition, the use of size was sufficient to identify the target. Colour distinguished the target from only one of the distractors. Finally, in the colour-only condition, the use of colour ruled out all distractors, whereas the size of the target distinguished it from only one of the distractors.

The experiment included 108 filler items. It was run together with Experiment 2, so thirty-six of the fillers consisted of the experimental items from Experiment 2. In the other 72 fillers, the target differed from both distractors in only its type (36 fillers), its orientation (44), part of its colour (6) or the number of objects (22).

Design. The 36 experimental items were tested in three conditions. For each language, we constructed three experimental lists consisting of 12 items in each condition. The conditions were rotated over the lists according to a Latin square design so that all items were presented in all conditions, but each list contained only one condition of an item. Ten

participants were randomly assigned to each list. The experimental items and 108 fillers were presented in a random order that was the same across lists.

Procedure. Participants were tested in pairs. The experimenter randomly assigned the roles of speaker and addressee to the two participants. Participants were sat behind computer monitors, facing each other. The participants were informed that they would see a series of scenarios containing objects. They were told that they would see the same objects, but that these appeared in different locations. This was done so that speakers would avoid producing referring expressions containing spatial information. The speaker was instructed to describe the circled objects so that the addressee could identify them and the addressee was asked to use the computer mouse to tick the described object on their screen.

Coding. Following transcription of the recordings, we scored whether speakers used colour-and-size descriptions (e.g., “the large blue helicopter”), colour-only descriptions (e.g., “the red kite”), and size-only descriptions (e.g., “small racket”). All trials on which speakers made a speech error and then repaired their utterance were excluded. There were 51 repairs (4.7%) in the English experiment and 94 (9.8%) in the Dutch experiment. We also excluded 8 cases (0.7%) from the English experiment and 7 (0.7%) from the Dutch experiment where speakers produced a modifier after the noun (usually in a relative clause, e.g., “big light bulb that’s green”), because they may be similar to a repair. Finally, 4 trials (0.4%) were removed from the Dutch experiment because the speaker mentioned neither colour nor size.

Results and Discussion

Fig. 5 shows the proportions of colour-and-size, colour-only and size-only descriptions in the English and Dutch experiments along with the predictions of the various models (see below). Because we were mainly interested in testing how well the models predicted the data rather

than in analysing differences between conditions, we focus on comparisons between the observed data and the predictions of the models below.

As is clear from Fig. 5, the pattern of data was very similar in English and Dutch. Speakers overspecified in all conditions. Overspecification occurred most frequently in the size-only condition, but also occurred, to a lesser extent, in the colour-or-size and size-only conditions. In contrast, speakers very rarely underspecified: We found only 4 such cases. The experiment also showed a clear preference for colour over size: In the colour-or-size condition, speakers used colour-only descriptions considerably more often than size-only descriptions and colour-only descriptions were more frequent in the colour-only condition than size-only descriptions in the size-only condition. The critical question is: Can the models we discussed in the Introduction account for the numerical patterns we observed?

Model Testing

Parameter setting in the non-deterministic algorithm and PRO. The goodness-of-fit of the non-deterministic IA and PRO is dependent on their parameter settings. To get an impression of the fit of these models, and to compare them with each other, we used the *generalisation criterion methodology* (Busemeyer & Wang, 2000). This method is particularly suitable to studies with experimentally controlled conditions, as it splits the data into two subdesigns based on the experimental conditions in the experiment. It uses one subdesign consisting of one or more conditions to determine the best-fitting parameter values for this subdesign (the *calibration stage*) and then uses these values to make predictions for a different *extrapolation condition* (the *generalisation stage*). As pointed out by Busemeyer and Wang (2000), this cross-validation method compares *a priori* predictions of models, because it tests to what extent the parameter settings in the training condition(s) make accurate predictions for a new set of condition(s).

In our experiments, we used two conditions as calibration conditions to determine the best-fitting parameter values and then used the remaining condition as the extrapolation condition. Thus, we found the best-fitting parameter values for the size-only and colour-only conditions to make numerical predictions in the colour-or-size condition. Similarly, the colour-or-size and colour-only conditions were used to make predictions in the size-only condition and the colour-or-size and size-only conditions for predictions in the colour-only condition. In this way, we obtained *a priori* predictions for all conditions. Because at least in theory, the findings from English and Dutch could be different, we did this separately for each language.

We used the binomial and multinomial probability functions *binopdf* and *mnpdf* in MatLab version 8.0.0.783 to calculate the parameter values that resulted in the highest $p(\text{data}/\text{model})$ in each of the pairs of calibration conditions. The obtained maximum likelihood parameter values from the pairs of calibration conditions were then used in the formulas for each of the models to calculate predictions in each of the conditions.

Using the binomial and multinomial probability functions, we then calculated $p(\text{data}/\text{model})$ across all conditions. We also calculated a value for the Bayesian Information Criterion (*BIC*), which provides an index of the goodness of fit of a model (the lower, the better the fit) and takes into account the number of free parameters (Schwarz, 1978). Finally, to compare two models directly, we calculated the Bayes factor (*B*): $B =$

$e^{-\frac{1}{2}(BIC_{\text{model1}} - BIC_{\text{model2}})}$. A Bayes factor larger than 10 is generally considered strong evidence for the better fitting model (Lewandowky & Farrell, 2010; Wasserman, 2000).

Evaluation of the non-deterministic IA. As explained in the introduction, the original, deterministic IA can be made non-deterministic by assuming that the colour-over-size preference is probabilistic. Table 1 shows a summary of the formulas used to calculate the predictions. We determined the maximum likelihood value for *c* in each pair of conditions

in each language by finding the value for c that resulted in the best fit, the highest $p(\text{data}/\text{model})$, for each pair of conditions. These values were then used to make predictions in the remaining, third condition.

Table 2 shows all values of c that we used to derive predictions from the non-deterministic IA in each language. In all cases, c is larger than .5, indicating that the non-deterministic IA assumes a preference for colour over size.

Insert Table 2 about here

Fig. 5 shows the predicted probabilities in English and Dutch and Table 3 shows the goodness-of-fit and BIC of the non-deterministic IA together with those of the other models discussed below. The modelling results show that the non-deterministic IA, which has one free parameter (c), provides a reasonable fit to the data in both English and Dutch. However, it predicts no overspecifications in the colour-or-size condition, whereas the data showed that overspecifications occurred fairly often (17% in English, 23% in Dutch).

Insert Table 3 about here

Evaluation of the modified non-deterministic IA. To account for the observed overspecifications in the colour-or-size condition, we assume that the incremental algorithm does not always stop after it has found a fully-discriminating expression, but sometimes adds a further attribute. See Fig. 3 for the decision tree and Table 1 for the formulas.

We again determined the maximum likelihood value for c in each pair of conditions, and then used this value for predictions in the remaining condition. Table 2 shows the c -values that we used and Fig. 5 shows the predicted proportions. We see that the algorithm

now correctly predicts overspecifications in the colour-or-size condition and Table 3 shows that overall, the fit of the model is fairly good. The fit is better than that of the original non-deterministic incremental algorithm: $B = 3.04 \cdot 10^{75}$ in English and $B = 8.57 \cdot 10^{103}$ in Dutch. However, the modified version predicts more frequent overspecifications in the colour-or-size condition than we observed and the predictions in the size-only and colour-only conditions are less good than in the original non-deterministic IA.

The modified algorithm may not be completely successful in accounting for our data because the probability of adding a further attribute to an already fully-distinguishing expression may not just be determined by the preference for that attribute. As mentioned in the Introduction, there is evidence that the likelihood of overspecification is affected by both task and type of speaker (Arts et al, 2011, Rubio-Fernandez, 2016; Deutsch & Pechmann, 1982). We therefore added an overspecification eagerness parameter e and tested whether including it would improve the predictions of the modified incremental model. In this model, the probability of adding size after selecting fully discriminatory colour is the preference for size ($1-c$) plus the overspecification eagerness (e) and the probability of not adding size in this case is the remaining probability ($c-e$). Similarly, the probability of adding colour after choosing fully discriminatory size is the preference for colour (c) plus the overspecification eagerness value (e), while the probability of not adding colour is the remainder ($1-c-e$).

For brevity, we do not present the full model testing here, because including the overspecification parameter did not help. The maximum likelihood value of e for predicting expressions in the colour-or-size condition was very negative ($-.46$ in English, $-.57$ in Dutch). A negative e -value is theoretically possible and means that overspecification is predicted to be less likely than it would have been without this parameter (see modelling of PRO below), but in this case it resulted in a negative value for the proportion of colour-and-size descriptions in the colour-or-size condition in both English and Dutch. Thus, including

an overspecification parameter resulted in impossible values, indicating that it is inappropriate to include it in this model.

Evaluation of the PRO model. The modelling above shows that neither the original non-deterministic IA nor the modified version provides a very good fit of all data. We therefore tested PRO, which assumes that speakers first choose an attribute that rules out all distractors; if there is more than one, then the probability of choosing an attribute is being determined by its preference c . Next, speakers may add a second attribute, with the chance of adding it depending on the preference for this attribute (c) and the overspecification eagerness of the speaker (e). See Fig. 5 for the decision tree and Table 1 for the formulas.

We determined the maximum likelihood values and calculated the predicted proportions of expressions in the same way as before. Table 2 shows the parameter values and Fig. 5 shows the predictions. Note that the overspecification eagerness parameter e is negative, as it was in the version of incremental algorithm that included this parameter, but in this case, it did not result in negative predicted proportions. Inclusion of this parameter merely makes overspecification less likely than if it had not been included. (We also tested PRO without this parameter and although the fit was better than any of the versions of the incremental algorithm, it was less good than PRO with this parameter.) The model testing results in Table 3 show that the fit of PRO is excellent. It is considerably better than that of the next-best model, the modified non-deterministic IA: English $B = 6.55 \cdot 10^{89}$, Dutch $B = 2.11 \cdot 10^{125}$. PRO predicts all observed proportions within .07.

Discussion

Comparison of the models and the observed data showed that PRO provided the best fit to the data from both the English and Dutch experiments. The non-deterministic incremental algorithm failed to account for overspecifications in the colour-or-size condition, because the algorithm terminates when either colour or size is selected, as either attribute is fully

distinguishing. The modified version of the incremental algorithm was designed to deal with this issue, because it sometimes adds a further attribute even if a fully distinguishing attribute has already been selected. This algorithm did indeed generate overspecifications in the colour-or-size condition, but provided a less good fit in the size-only and colour-only conditions than the original non-deterministic algorithm. Additional modelling showed that the fit did not improve when an overspecification eagerness parameter was added, suggesting that the better fit of PRO is not merely due to the extra parameter it has.

Experiment 2

Experiment 1 investigated whether models of reference generation can successfully account for the use of two frequently used attributes, colour and size. Experiment 2 investigated the use of a different commonly used attribute, the object's type or category, which is usually encoded by the head noun, such as "bed" or "ashtray". In this respect, it is different from other attributes, which are generally encoded by a modifier. Thus, its selection may involve different mechanisms from those of other attributes. Specifically, the original incremental algorithm as proposed by Dale and Reiter (1995) assumes that type is included in the description regardless of whether it rules out any distractors, that is, speakers would say "the grey bed" rather than "the grey one" even if all objects in the context are beds. In our modelling of the non-deterministic IA, we made the same assumption.

In contrast, the PRO model does not make any additional assumptions to account for the use of the type attribute: The decision tree in Fig. 4 also applies to cases where type is a distinguishing attribute. Similarly, we do not need to make any special assumptions for the use of type in the modified non-deterministic IA, because it assumes that, after selecting a fully distinguishing attribute, speakers sometimes add a further attribute. Hence, after

selecting an attribute such as colour, they may add type even though colour is fully distinguishing.

Thus, Experiment 2 tested whether models need to incorporate a special mechanism for dealing with type. Fig. 6 shows an example of the conditions in Experiment 3. The conditions were the same as in Experiment 1, but rather than manipulating colour and size, we manipulated type and colour. To test the various models, we examined whether speakers mentioned the type alone, the colour alone or both the type and colour in the different conditions. As before, we tested both English and Dutch.

Insert Fig. 6 about here

Method

Participants. Experiment 2 was run together with Experiment 1, so the participants were the same.

Materials, design, and procedure. These were the same as in Experiment 1, except that the experimental materials now had conditions where either only colour ruled out all distractors, only type ruled out all distractors or both type and colour ruled out all distractors.

Coding. The coding was done in the same way as in Experiment 1. In the English experiment, 25 trials (2.3%) were excluded because speakers repaired their utterance, 2 (0.2%) because they used a postnominal modifier and 4 (0.4%) because they used a negation (e.g., “not the scissors”). In the Dutch experiment, 38 trials (3.5%) were repairs, 3 (0.3%) contained a postnominal modifier, 5 (0.5%) were excluded because the response could not be understood, and 18 (1.7%) because of recording failure.

Results and Discussion

Fig. 7 shows the observed proportions of type-and-colour, type-only and colour-only descriptions in English and Dutch along with the predictions of the models. As in Experiment 1, the results from English and Dutch were similar: Speakers overspecified in all conditions and they rarely underspecified (only 4 cases). The experiment showed a clear preference for type over colour: In the type-or-colour condition, speakers produced type-only descriptions much more frequently than colour-only descriptions. In addition, they produced type-only descriptions in the type-only condition more often than colour-only descriptions in the colour-only condition.

Model Testing

Evaluation of the non-deterministic IA. Similar to Experiment 1, we turned the deterministic IA into a non-deterministic variant by assuming that the type-colour preference is probabilistic, that is, speakers first check colour with a probability of c and first check type with a probability of $1-c$. Type is fully discriminatory in the type-or-colour condition, so when speakers first check type, they should not add colour, resulting in the expression *bed* with a proportion of $1-c$. Colour is also fully discriminatory in this condition, so when speakers check colour first, they do not need to add type to rule out all distractors. However, to account for our finding that speakers rarely produced colour-only descriptions, we follow Dale and Reiter (1995) in assuming that the head noun is always produced, that is, type is always added. Thus, speakers should produce *black bed* with a proportion of c in the type-or-colour condition. In the colour-only condition, speakers should again first choose type with a probability of $1-c$, but because it is not fully discriminatory in this case, they should subsequently add colour. The rest of the time (c), they should first choose colour. However, although colour is fully discriminatory, they should again add type because it is encoded on the head noun. Thus, regardless of whether speakers first choose type or colour in the colour-

only condition, they always end up including both attributes. Finally, in the type-only condition, type is fully discriminatory, so when they check this attribute first, they do not need to add colour, resulting in *bed* with a proportion of $1-c$. The rest of the time, they first check colour, but because this is not fully discriminatory, they need to add type, resulting in *black bed* with a proportion of c . Table 4 shows the formulas for the non-deterministic IA.

Insert Table 4 about here

As before, we determined the value for c in a particular condition and language by calculating the maximum likelihood value for c in the two other conditions of that language. See Table 5 for the values of c and Fig. 7 for the predictions.

Insert Table 5 about here

Table 6 shows that the model makes fairly accurate predictions. In particular, it is notable that the non-deterministic IA correctly predicts overspecification in the type-or-colour condition. The reason is that in cases where colour is selected first, the head noun, encoding type, is always added, resulting in an overspecified expression. Other predictions for Experiment 1 are also in line with the observed data. In particular, the non-deterministic IA predicts that colour-only descriptions are never produced (because it assumes that type is always included), and the observed data show that such descriptions are indeed very rare. It also correctly predicts that the majority of descriptions in the colour-only condition are type-and-colour descriptions and that in the type-or-colour and type-only conditions, both type-and-colour and type-only descriptions are produced a fair amount of time. However, the

model predicts more type-only descriptions in the type-or-colour than type-only condition, whereas the data show the opposite pattern.

Insert Table 6 about here

Evaluation of the modified non-deterministic IA. The modified version of the IA is the same as in Experiment 1 (Fig. 3), except that size is replaced by type. Table 4 shows the formulas. As mentioned above, it does not require additional assumptions about the inclusion of the head noun. This is because the model assumes that a second attribute may be added even if the initially selected attribute is fully distinguishing. Assuming that type is strongly preferred, speakers should usually add it after selecting colour even if colour is fully distinguishing. Therefore, few colour-only descriptions should be produced.

We determined the type-colour preference value c as before, shown in Table 5 and resulting in the predictions in Fig. 7. The predictions are quite similar to those of the original non-deterministic algorithm, except that the modified model correctly predicts some colour-only descriptions in the colour-only and type-and-colour conditions. Table 6 shows that the fit of the model is good, considerably better than the original non-deterministic IA: $B = 2.80 \cdot 10^{47}$ for English and $B = 4.43 \cdot 10^{44}$ for Dutch. However, like the original model, it incorrectly predicts more type-only descriptions in the type-or-colour than type-only condition. We also tested whether adding an overspecification eagerness parameter similar to that used in PRO improved the model. Unlike in Experiment 1, this did not result in negative predicted proportions, but the model with the additional parameter provided a less good fit.

Evaluation of the PRO model. To account for the data in Experiment 2, PRO makes the same assumptions and uses the same formulas as for the choice between colour and size in Experiment 1. The decision tree in Fig. 4 can be used by replacing size with type. Table 4

shows the formulas. Similar to the modified non-deterministic IA, PRO does not make any additional assumptions for the inclusion of the head noun. Thus, it predicts that type is omitted in a proportion of cases in the type-or-colour and colour-only conditions.

We determined parameter values as before, shown in Table 5. In contrast to Experiment 1, the overspecification eagerness values e are positive, indicating that the model assumes that speakers are more eager to overspecify when the distinguishing attributes are type and colour than when they are colour and size. This may seem surprising given that Experiments 1 and 2 were run together with the same participants. One possible explanation is that speakers do not always consider type-and-colour descriptions as true overspecifications when only type or only colour is sufficient, because omitting the head noun is uncommon in English and Dutch (as our experiments show). Therefore, they may be more likely to overspecify in such cases. This suggests that the overspecification parameter e does not just capture how eagerly speakers overspecify, but also whether they consider type-and-colour descriptions to be true overspecifications in cases where mentioning one of the attributes is sufficient. This is in line with Dale and Reiter's (1995) proposal, which uses a special mechanism to ensure that a type is included in every description; however, our model arrives at this prediction automatically, without requiring any special mechanism.

Fig. 7 shows the predictions of PRO. All predicted data points are within .07 of the observed data, resulting in a very good fit of the model, as shown in Table 6. The fit is better than the modified non-deterministic IA in English ($B = 1.86 \cdot 10^{13}$) and somewhat better in Dutch ($B = 10.96$).

Discussion

As in Experiment 1, PRO provided an excellent fit to the data from both the English and Dutch experiments. Although the two versions of the non-deterministic IA that we tested

provided a decent fit to the data, they performed less well. The original modified incremental IA did not predict any colour-only descriptions in the colour-only and type-and-colour conditions. The modified version did, but it incorrectly predicted that type-only descriptions should be more frequent in the type-or-colour than in the type-only condition.

An interesting conclusion from the modelling of the experimental results is that we do not need to assign a special status to the type attribute, that is, we do not need to assume that speakers add it because they want to avoid a referring expression without a head noun, unlike the incremental algorithm. Where the type is distinguishing, both PRO and the modified non-deterministic IA can account for the paucity of descriptions without this attribute by assuming that it is strongly preferred.

Experiment 3

The previous experiments showed that PRO accounts well for the production of referring expressions when colour and size discriminated the target from the distractors (Experiment 1) and when type and colour did (Experiment 2). However, accounting for reference production in these experiments was arguably relatively simple. To test the scope of PRO and to compare it more thoroughly with the next-best model, the modified non-deterministic IA, we included a third discriminating attribute in Experiment 3. This makes the reference generation process in the models more complex because (1) speakers need to choose between three rather than two different attributes that each have their own preference, (2) they can produce longer expressions by choosing all three attributes and (3) because the attributes can be combined in more different ways, they have a larger number of expressions to choose from.

We used a border (e.g., a square or a triangle) around the objects as the third attribute because speakers normally realise the border attribute post-nominally in a prepositional

phrase (e.g., *the airplane in the square*). This allowed us to explore the possibility that concepts expressed as postnominal modifiers are selected differently from prenominal modifiers. Brown-Schmidt and Tanenhaus (2006), for example, showed that fixations to a size-contrast distractor were later when a size adjective was produced post-nominally (e.g., *the triangle with small diamonds*) than when it was produced pre-nominally (e.g., *the small triangle*), suggesting that pre- and postnominal modifiers are encoded at different points in time (see also Brown-Schmidt and Konopka, 2008). However, other research suggests that at least in some instances, speakers encode postnominal modifiers simultaneously with the earlier part of the noun phrase. For example, Garrett's (1975) analysis of speech errors showed that word exchange errors occur between different phrases in a clause (e.g., *I broke a dinghy in the stay*, where *I broke a stay in the dinghy* is intended), suggesting that lexical planning has a clausal scope (see also Meyer, 1996 for evidence from the picture-word interference method).

Experiment 3 did not aim to investigate the scope of conceptual planning directly, but because there is some evidence that pre- and postnominal modifiers are planned differently, we wanted to see whether PRO could account for the selection of concepts in postnominal modifiers. If the selection of concepts in postnominal modifiers is similar to that in prenominal modifiers, PRO would not need to make any additional assumptions to account for the use of the border attribute. But if concept selection is different in postnominal modifiers, then the PRO algorithm that we tested in Experiments 1 and 2 would not make accurate predictions for this case and additional mechanisms would need to be considered. To test the predictions of PRO and the modified non-deterministic IA, we used the conditions in Fig. 8, in which we manipulated the colour, size and border of the objects.

Insert Fig. 8 about here

Method

Participants. Thirty-five new participants from the same population as in the English Experiments 1 and 2 took part.

Materials. The materials consisted of 84 experimental items in 7 conditions. Fig. 8 shows an example. Each condition showed three objects of the same type. The objects were selected from Rossion and Pourtois (2004). The target object was indicated by a red arrow underneath it and its position was counterbalanced across items. In the seven conditions, we manipulated whether the colour, size or border of the target was different from that of the distractors. The border attribute was manipulated by putting either a square, circle, diamond or triangle around the objects. As in the previous experiments, we used four colours: blue, grey, red and green, but the size of the smaller objects was only half the size of the larger objects (it was two-thirds in Experiments 1 and 2).

In the colour-only condition, colour distinguished the target from both distractors, whereas size and border distinguished the target from only one of the distractors. Similarly, in the size-only and border-only conditions, respectively size and border distinguished the target from both distractors, whereas the other two attributes only distinguished it from one distractor. In the colour-and-size, colour-and-border and size-and-border conditions, there were always two fully discriminating attributes as indicated by the labels for the conditions, whereas the remaining attribute ruled out only one distractor. Finally, in the colour-size-or-border condition, all three attributes each ruled out all distractors.

We also included 96 fillers. In these fillers, the target differed from both distractors in both its colour and type (12 fillers), both its colour and border (12), its orientation (20), the number of objects (20), type only (14), size only (6), border only (6) or part of its colour (6).

Design. All 84 experimental items had seven conditions. We constructed seven experimental lists consisting of 12 items per condition, using a Latin square design as before. Five participants were assigned to each list. The experimental items and fillers were presented in a random order that was the same across lists.

Procedure. The procedure was the same as in the previous experiments except as follows. Instead of an addressee who was a real participant in the experiment, we used a confederate. The confederate was a male native English speaking student from the University of Dundee. None of the participants showed awareness that their partner was a confederate. To ensure that attributes were used as pre-modifiers or post-modifiers (rather than say, e.g., “blue in the circle”), the speaker was asked to name the object in each expression.

Coding. The coding was the same as in the previous experiments except that we now also coded whether speakers used the border attribute. We excluded trials on which participants produced a speech error ($N = 208$, 7.1%) or a postnominal modifier that did not include border ($N = 42$, 1.4%). We excluded 8 trials (0.3%) on which speakers did not include a head noun (e.g., “green”) because we wanted to test whether the models could account for the use of attributes that were expressed both before and after the noun. Finally, 3 trials (0.1%) were excluded because speakers did not use either colour, size or border.

Results and Discussion

Insert Fig. 9 about here

Fig. 9 shows the observed proportions of referring expressions in each condition along with the predictions of the modified non-deterministic IA and PRO. As in Experiments 1 and 2, the results showed that overspecification was common in the conditions where the most preferred attribute (here: colour) did not rule out all distractors. Even in the conditions where colour was fully discriminating, participants used overspecified descriptions in over 35% of trials. Across conditions, participants used all combinations of attributes that led to unambiguous descriptions, including descriptions that contained all three attributes colour, size and border. Finally, participants again rarely produced underspecified expressions.

Model testing

Evaluation of the modified non-deterministic IA. Because participants in Experiment 3 could choose between three attributes (colour, size and border), the modified non-deterministic IA had two parameters: one for the colour preference (c) and one for the size preference (s). The border preference was the remaining proportion ($1-c-s$). We only calculated the predictions of the model without the overspecification eagerness parameter because in both Experiments 1 and 2, the fit was better without it.

The algorithm cannot be presented in a single decision tree for all conditions, so instead we exemplify its workings with a decision tree for the colour-only condition (Fig. 10). The algorithm assumes that speakers may either select colour (c), size (s) or border ($1-c-s$) first in this condition. If they select colour, they may not add a further attribute because colour is fully discriminatory. The probability that no attribute is added is the colour preference (c). However, the modified version of the IA does not always terminate when a fully discriminating description is found, so either size or border may be added. The probability of adding either size or border is determined by the preference for each (respectively s and $1-c-$

s). After the second attribute is added, speakers may add a third attribute and the probability of this is again determined by the preference for that attribute.

Insert Fig. 10 about here

The algorithm proceeds slightly differently if size is first selected in the colour-only condition, because this attribute is not fully distinguishing. Thus, speakers must add a further attribute. The probability that they add colour is the preference for colour normalised over the preferences of the remaining attributes colour and border ($c/(1-s)$), and similarly, the probability of adding border is the normalised preference for border. If speakers first select size and then colour, they may either add border (determined by its preference) or not, but if they first select size and then border, they must add colour, because neither size nor border is fully distinguishing. Finally, if speakers first select border, then the algorithm proceeds along similar lines as when they first select size.

As before, the likelihood that a particular expression is produced is calculated by multiplying the formulas at each stage in the decision tree and summing the different routes that can be taken. Appendix 2 shows all formulas derived from the algorithm.

We determined the parameter values of c and s in each condition by finding their maximum likelihood values in the other six conditions. Table 7 shows the values we used in each condition and Fig. 9 shows the resulting predictions. The fit of the model is quite poor: $p(\text{data}/\text{model}) < 1.00 \cdot 10^{-320}$, $BIC > 1490$. Especially, in the colour-only, size-only and border-only conditions, it overpredicts the probability that speakers use all three attributes and underpredicts the probability that they use a single attribute. In the colour-or-size, colour-or-border and colour-size-or-border conditions, it underpredicts the proportion of colour-only

descriptions and in the size-or-border condition, it overpredicts colour-and-border descriptions.

Insert Table 7 about here

Evaluation of PRO. Like the modified non-deterministic IA, PRO has parameters for the colour (c) and the size (s) preference; the border preference is the remainder ($1-c-s$). In addition, it includes the overspecification parameter (e), but in contrast to the previous experiments, we did not calculate its value using the other conditions in the current experiment. Given that our participants came from the same population and the method was similar to that in the previous experiments, the overspecification eagerness value in the current experiment should be similar. In particular, it should be similar to the value in Experiment 1, because neither in the current experiment nor in Experiment 1 was the head noun a distinguishing attribute. Experiment 2 showed that the overspecification value is different when the head noun *is* a distinguishing attribute. (We argued this is because speakers do not consider overspecification by the head noun as true overspecification, because omitting the head noun is uncommon.) Thus, we used the overspecification value from Experiment 1 ($-.0531$).

Fig. 11 shows the PRO decision tree for all conditions. We exemplify it using the colour-or-size condition. Both colour and size are fully distinguishing attributes, so speakers should always first choose either colour or size. The probability of choosing each is determined by the preference for the attribute normalised over the two fully distinguishing attributes. The starting box on the left in Fig. 11 shows this. For example, the probability of first selecting colour in the colour-or-size (CS) condition is $c/(c+s)$. Next, speakers may add no further attribute, add size or add border. The probability that speakers do not add a further

attribute is the colour preference (i.e. not the size plus border preference) minus the overspecification eagerness value ($c-e$). The probability that speakers *do* add another attribute is the remainder ($1-c+e$) multiplied by the preference for the attribute (e.g., size) normalised over the two remaining attributes (e.g., size and border, $s/(1-c)$). Finally, speakers may add a third attribute, with the chance being determined by the preference for that attribute plus the overspecification eagerness value (e.g., $1-c-s+e$ for adding border after colour and size). Again, the probability that a particular description is used is calculated by multiplying the formulas at each stage and summing the different routes that lead to the expression. Appendix 2 shows all formulas.

Insert Fig. 11 about here

We determined the c and s parameter values as before. Table 7 shows these values and Fig. 9 shows the predictions of PRO. In all conditions, PRO correctly predicts what the most frequently used description is and in all conditions except the size-or-border condition, it also predicts what the second most frequent description is. All proportions are predicted within .15. The model appears to have no particular difficulty with descriptions that used the postnominal border modifier and accurately predicts the proportions of descriptions with all three attributes. Overall, the fit of the model is good: $p(\text{data}/\text{model}) = 8.91 * 10^{-108}$, $BIC = 508.77$. Note that the BIC value cannot directly be compared to that in Experiments 1 and 2, because it is based on more conditions. The fit of PRO is considerably better than that of the modified non-deterministic IA: $B > 6.11 * 10^{234}$.

Discussion

The main aim of the experiment was to test whether PRO could account for reference production in more complex scenarios with more distinguishing attributes, resulting in a wider choice of attributes and in potentially longer expressions. The modelling results suggested that it can: The predicted proportions were generally very close to the observed results for all descriptions.

Most critically, PRO made better predictions than the modified non-deterministic IA, the model closest to PRO in Experiments 1 and 2. Although the modified non-deterministic IA made fairly accurate predictions in the previous experiments, it did not fare well in Experiment 3. In particular, in all conditions except the size-or-border condition, it predicted too many descriptions with all three attributes and in all conditions, it predicted too few with a single attribute. Unlike PRO, the modified non-deterministic IA does not appear to generalise to the more complex scenarios in Experiment 3, suggesting that its reference production mechanisms are different from those of human speakers.

Experiment 3 showed that the production mechanisms underlying PRO do generalise to situations that are more complex than in the previous experiments. In addition, PRO made accurate predictions for postnominal modifiers that include the border attribute; it did not require additional assumptions to deal with these modifiers. PRO assumes that properties are not necessarily selected in the order in which they are realized, that is, whether they occur pre- or postnominally. Instead, an attribute's discriminatory power and preference affect the order in which concepts are selected. PRO is currently a concept selection model only; it does not account for the linearization of the selected attributes, but assumes that this is a separate task. The modelling results from Experiment 3 showed that this is a feasible account, as PRO made accurate predictions about both pre- and postnominal attributes. Although Experiment 3 did not directly test whether human speakers select postnominal concepts after prenominal

concepts, the fact that PRO can account for both without postulating additional mechanisms suggests that human speakers do not distinguish between pre- and postnominal modifiers during concept selection.

To derive the PRO predictions, we used the same overspecification eagerness value e as in Experiment 1, so the number of free parameters of PRO in Experiment 3 was the same as that of the modified non-deterministic IA. Thus, the number of free parameters does not explain the better fit of PRO compared to the modified non-deterministic IA. The finding that the overspecification eagerness value from Experiment 1 resulted in accurate predictions for Experiment 3 is striking for two reasons. First, if postnominal modifiers were added as an “afterthought” when the prenominal modifiers did not rule out all distractors (similar to a speech repair), then the overspecification value for postnominal attributes should have been different from that of prenominal attributes, because speakers would have been less eager to add an overspecifying postnominal attribute. The fact that we could use the same overspecification value for pre- and postnominal attributes again suggests that their concepts are selected in the same way. Second, one might have expected that the more attributes speakers have already selected, the less likely it is that they select further attributes. As a result, the overspecification value should have been lower during the selection of the third attribute than the second. This did not appear to be the case: Using the same value, PRO accounted for referring expressions that included two as well as three attributes. In other words, a third attribute is as likely to be added as a second. However, expressions with three attributes are less frequent than those with two because, following the selection of two attributes, the overall chance that no further attribute is added is higher than that an attribute *is* added.

General Discussion

Evaluation of the Models

As we mentioned in the Introduction, existing computational models are at odds with important aspects of the way in which people perform the conceptualization of referring expressions. In particular, evidence that speakers use a range of referring expressions in the same condition is inconsistent with deterministic computational models such as the full brevity, greedy and incremental algorithms, while the rational speech act theory currently does not deal with overspecification. Taking into account that human speakers are non-deterministic and frequently overspecify, we therefore developed new models and tested them in three reference production experiments.

To account for the inherently probabilistic nature of reference production, we modified the original incremental algorithm to make it non-deterministic. Unlike in the original algorithm, where the preference order of attributes is fixed, we postulated that speakers are most likely to first select the attribute that is most preferred, but sometimes instead select the less preferred attribute, with the likelihood being determined by the relative preference for the two attributes. We tested two non-deterministic versions of the incremental algorithm. In the first version, probabilistic selection of attributes continues until all distractors are ruled out. However, this version failed to account for overspecification when all distinguishing attributes were fully distinguishing, because the selection of any distinguishing attribute causes the algorithm to stop. We therefore tested a second version according to which speakers probabilistically add further attributes even after a fully distinguishing description has been found. This model indeed provided a better overall fit of the data in Experiments 1 and 2 than the first version, but in Experiment 3, it predicted too many descriptions with all three attributes and too few with a single attribute.

Our results were more consistent with the predictions of the PRO model, a probabilistic model that assumes that attributes' preferences as well as discriminatory power play a role in conceptualisation during reference production. It claims that speakers first choose the attribute that rules out all distractors. When more than one attribute is fully discriminating, they probabilistically choose one depending on its preference. Next, they add further attributes depending on the degree to which each of these is preferred and on speakers' eagerness to overspecify. Across experiments, PRO made accurate predictions about the proportions with which different referring expressions were used in the conditions.

The excellent fit of PRO cannot be explained by the number of free parameters in the model. First, the parameter values in each condition were fixed on the basis of the data from the other conditions so they were not truly free. Second, versions of the non-deterministic IA that had the same number of parameters as PRO made less good predictions. In Experiments 1 and 2, we tested the modified non-deterministic IA with an overspecification parameter, but its predictions were less good than PRO. In Experiment 3, PRO had the same number of parameters as the modified non-deterministic IA, because the overspecification parameter value was independently obtained from Experiment 1.

The main conclusion we can draw from our PRO modelling results is that it is possible to develop an algorithmic model of reference that mimics human reference production very closely, much more closely than classic reference generation algorithms of the 1990s and current versions of the rational speech act theory. The success of the model comes from making crucial aspects of the algorithm probabilistic and taking into account both the preference of attributes and a form of discriminatory power. PRO can be seen as a synthesis between the classic 1990s algorithms and the rational speech act theory. From earlier reference generation algorithms, it borrows the ideas that some attributes are intrinsically more preferred than others and that discriminatory power plays a role. From the rational

speech act theory, it borrows the ideas of probabilistic generation and discriminatory power. By combining these ideas in a unique way, PRO makes accurate predictions about human reference production.

The complexity of the formulas in the PRO algorithm (especially for longer expressions and conditions where several attributes are fully distinguishing, as in Experiment 3) might seem at odds with the fluency with which human speakers usually produce reference. The reason that the formulas are so complex is that PRO predicts averaged probabilities of referring expressions across multiple speakers and items using multiple routes through the decision trees. However, on each individual trial, the referring process is much simpler because averaged probabilities need not be calculated. Rather, the process of producing reference can be seen as throwing a biased dice at each point in the decision tree that determines whether they will follow one route or another. This can be a simple, fast process for human speakers.

Although PRO can boast a considerable degree of empirical adequacy, one might argue that this achievement is bought at the expense of the attractive simplicity possessed by such models as the greedy algorithm and the rational speech act theory. An interesting question from this perspective is what the origins of the attribute preference and overspecification parameters in PRO are. Why do speakers prefer certain attributes over others and why does the likelihood of overspecification depend on the type of speaker and her task? Regarding attribute preference, it seems plausible that this is affected by how easy it is to detect differences in this attribute. Belke and Meyer (2002) found that participants detected differences in colour between two objects faster than differences in size. They suggested that this may be because colour is an absolute attribute, that is, it can be determined independently of the colour of other objects, whereas size is a relative attribute that can generally only be determined by comparing it to other objects in the context. However, we have recently found

that colour is not always preferred to size in reference production (Van Gompel, Gatt, Krahmer, & Van Deemter, 2014): When the difference in size between the target and distractors is sufficiently large, speakers use size as often as colour. This is more consistent with the idea that attribute preferences are due to their perceptual saliency rather than the distinction between relative and absolute attributes.

Other factors may also affect attribute preference. For example, Goudbeek and Krahmer (2012) found that the use of specific attributes in object descriptions could be primed: Speakers used an attribute more often when they had just heard a referential description with this attribute than when they had not (e.g., they said “the fan facing left” more often after hearing “the front-facing chair” than “the red chair”). This finding suggests that in addition to perceptual saliency, the conceptual saliency of attributes (due to priming) also affects their preference. This may also explain why Westerbeek, Koolen, and Maes (2015) and Rubio-Fernandez (2016) found that speakers mentioned colour more often when it was an atypical attribute of the object (e.g., a pink banana): Atypical colours are conceptually incongruent and may therefore be more salient than typical colours.

The overspecification parameter is likely to be affected by speakers’ cooperativeness with the listener: As Arts et al. (2011) and Rubio-Fernandez (2016) observed, when it is particularly important that the listener correctly identifies the object because it is critical for his task, speakers tend to overspecify more. Time pressure (cf., Horton & Keysar, 1996) and the complexity of the domain (Koolen et al., 2013; Paraboni and Van Deemter, 2014) may also affect how often speakers overspecify. The amount of overspecification also appears to depend on the type of speaker. As we mentioned before, Deutsch and Pechmann (1982) found that children overspecified less often than adults. It is possible that this speaker effect reflects cooperativeness: Certain speakers such as children may be less cooperative with the

listener than other speakers (Fukumura, 2016; Matthews, Lieven, Theakston, & Tomasello, 2006).

Future directions

Because there is an almost unlimited number of target attributes and different distractors in the real world, speakers' ability to refer potentially faces an infinite number of possible situations. Since our decision trees do not cover all possible inputs (e.g., where there are more than three distinguishing attributes or where there is no single fully distinguishing attribute), they are not a complete algorithm. In Appendix 1, we show the "complete" PRO *algorithm*, which is defined under a much wider class of possible situations; since the algorithm is under-determined by the available experimental data, it contains an element of extrapolation. The algorithm is only complete under certain assumptions: for example, it only deals with one-shot descriptions to singular objects (as opposed to sets), so it does not yet account for repeated reference and the use of anaphora. In future, PRO may be extended to cover an even wider class of situations.

Many features of the PRO algorithm remain untested. For example, our experiments tested situations with only two distractors. When there are many distractors, speakers may not keep all information in their focus of attention. As a result, the first stage in PRO, where speakers choose a fully distinguishing property, may become probabilistic: Although a property may not rule out all distractors, speakers may sometimes first choose it because it rules out all distractors in their current attention. Speakers may also ignore a fully discriminating property during the first stage because they may sometimes fail to check whether it has discriminatory power. Additional mechanisms may be called for. For example, it might be that speakers only take into account distractors and their attributes if they have

fixated them; PRO could deal with this by postulating that only these distractors and attributes are in the attentional domain.

Another question is which attributes are considered for inclusion in the referring expression. We have implicitly assumed that PRO only selects attributes that distinguish the target from at least one of the distractors. For instance, it cannot generate expressions like “grey candle with the flame” for the situations in Fig. 2, because “with the flame” does not discriminate the target from any of the distractors. This prediction was generally borne out by the data from our experiments. The only exception was the use of the type attribute in Experiment 1, where it did not rule out any distractors: Participants used this in 88% of their English descriptions and 93% of their Dutch descriptions. (In Experiment 2, PRO can explain the occurrence of references with type because in that experiment, type ruled out at least one distractor. In Experiment 3, participants were instructed not to omit the noun.)

One possible solution for this problem would be to assume that, as it stands, PRO focuses on discriminating attributes, but that later processes can add attributes that do not rule out any distractors, subject to the preference for this attribute (so type would be added relatively often). A different solution would be to assume that the decision trees of PRO (Figs. 4 and 11) should include not just colour, size and border, but also type. We did not test whether this would account for the use of the type attribute, because it would have resulted in another parameter, type preference, and this would have resulted in too many parameters for the number of data points (i.e. conditions and types of expressions) we had. The inclusion of another attribute, such as type, to the model would not change the relative proportions with which the descriptions including colour, size and border are used; the only thing that changes is that the model would sometimes generate these descriptions with and sometimes without the additional attribute.

As described, PRO and the other computational models we have discussed are conceptualisation or content determination models. They were not designed to account for word order preferences, for example, whether speakers say “large grey candle” or “grey large candle”. In line with most language production models, we have assumed that speakers in our experiments first determined which concepts to express, and following this, they accessed the lexical properties of the words for these concepts and put the words in order during the realisation stage (e.g., Bock & Levelt, 1994; Dell, 1986; Levelt et al., 1999).

PRO’s strict separation between conceptualisation and realisation may suggest that speakers put all concepts that they select into a memory buffer before they lexicalise them and put them in order. In our experiments, this seems plausible, because there was little word order variation: For example, speakers virtually always produced size before colour in Experiments 1 and 3. Although in theory, it is possible that speakers always conceptualised size before colour and put the words in the same order as they accessed the concepts, this seems unlikely given that colour was preferred over size. Furthermore, in the colour-only conditions, colour was fully discriminatory, whereas size was not. Thus, it seems plausible that at least in some cases, speakers accessed the concept for colour first and kept it in a conceptual buffer until they had accessed the size concept. It is likely that this also happened in Experiment 2, where colour was produced before type even though type is more preferred. In sum, when properties have a rigid word order, speakers often appear to hold concepts in a buffer before they order them.

However, when word order is more flexible, the order in which concepts are selected may affect their order of mention. Fukumura (2018) investigated the ordering of colour and pattern (green spotted bow vs. spotted green bow), where both word orders are relatively common. She found that both discriminability and attribute preference affected word order, suggesting that speakers may immediately have lexicalised the property that was

conceptualised first, rather than holding it into a buffer. If the latter is true, PRO should also make predictions about word order in such cases; it would be interesting to explore this in future work.

Finally, PRO is currently neutral with regard to whether speakers select particular attributes because they are *allocentric*, in order to help the listener identify the target, or *egocentric*, selecting attributes that are easiest for themselves to produce (e.g., Brennan & Clark, 1996; Brown & Dell, 1987; Heller, Gorman, & Tanenhaus, 2012; Horton & Keysar, 1996; Fukumura & Van Gompel, 2012; Wardlow-Lane & Ferreira, 2008). For example, speakers may first choose a fully discriminating attribute because the fact that it rules out all distractors makes it most helpful for the listener. Alternatively, they may first choose a fully discriminating attribute because this target attribute is most contrastive and therefore perceptually most salient for themselves. Similarly, speakers may prefer one attribute over another because they assume it is perceptually or conceptually more salient for the listener (and therefore more helpful) or because it is more salient for themselves. PRO is compatible with either type of account.

In sum, we hope to have demonstrated that computational models of reference generation can be useful models of human reference production and conceptualisation, because they make explicit quantitative predictions that can be tested in psycholinguistic experiments. We see the PRO model, and the manner in which we have tested its quantitative predictions, as the start of a new avenue of research. We have shown that PRO provides an excellent fit to our current data, but more experimental work is needed to test more of PRO's assumptions and predictions. This should result in an even better understanding of human reference production, and in even more accurate and more general reference production algorithms.

References

- Appelt, D. (1985). Planning English referring expressions. *Artificial Intelligence*, 26, 1–33.
- Arts, A. (2004). *Overspecification in instructive texts*. Nijmegen: Wolf Publishers.
- Arts, A., Maes, A., Noordman, L., & Jansen, C. (2011). Overspecification in written instruction. *Linguistics*, 49, 555–574.
- Belke, E., & Meyer, A. (2002). Tracking the time course of multidimensional stimulus discrimination: Analysis of viewing patterns and processing times during same-different decisions. *European Journal of Cognitive Psychology*, 14, 237–266.
- Bock, J.K. (1995). Producing agreement. *Current Directions in Psychological Science*, 4, 56–61.
- Bock, K., & Levelt, W. J. M. (1994). Language production: Grammatical encoding. In M.A. Gernsbacher (Ed.), *Handbook of Psycholinguistics* (pp. 945-984). London: Academic Press.
- Brennan, S. E., & Clark, H. H. (1996). Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22, 1482–1493.
- Brown, P.M., & Dell, G.S. (1987). Adapting production to comprehension: The explicit mention of instruments. *Cognitive Psychology*, 441-472.
- 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, 274-280.
- Brown-Schmidt, S. & Tanenhaus, M.K. (2006). Watching the eyes when talking about size: An investigation of message formulation and utterance planning. *Journal of Memory and Language*, 54, 592-609.

- Bussemeyer, J.R. & Wang, Y-M. (2000). Model comparison and model selections based on generalization criterion methodology. *Journal of Mathematical Psychology*, 44, 171-189.
- Dale, R. (1989). Cooking up referring expressions. In *Proceedings of the 27th Annual Meeting of the Association for Computational Linguistics (ACL)*, pp. 68–75.
- Dale, R. (1992). *Generating Referring Expressions: Constructing Descriptions in a Domain of Objects and Processes*. Cambridge: The MIT Press.
- Dale, R., & Reiter, E. (1995). Computational interpretations of the Gricean maxims in the generation of referring expressions. *Cognitive Science*, 18, 233–263.
- Dell, G.S. (1986). A spreading-activation theory of retrieval in sentence production. *Psychological Review*, 93, 283–321.
- Deutsch, W., & Pechmann, T. (1982). Social interaction and the development of definite descriptions. *Cognition*, 11, 159–184.
- Engelhardt, P.E., Bailey, K. G., & Ferreira, F. (2006). Do speakers and listeners observe the Gricean Maxim of Quantity? *Journal of Memory and Language*, 54, 554–573.
- Engelhardt, P.E., Demiral, S. B., & Ferreira, F. (2011). Over-specified referring expressions impair comprehension: An ERP study. *Brain and Cognition*, 77, 304–314.
- Frank, M.C., & Goodman, N.D. (2012). Predicting pragmatic reasoning in language games. *Science*, 336, 998.
- Fukumura, K. (2016). Development of audience design in children with and without ASD. *Developmental Psychology*, 52, 71-87.
- Fukumura, K. (2018). Ordering adjectives in referential communication. *Journal of Memory and Language*, 101, 37-50.
- 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, 1289–1311.

- Garrett, M.F. (1975). The analysis of sentence production. In G. Bower (Ed.), *The psychology of learning and motivation* (pp. 133–177). New York: Academic Press.
- Garrett, M.F. (1984). The organization of processing structure for language production: Application to aphasic speech. In D. Caplan, A.R. Lecours, & A. Smith (Eds.), *Biological perspectives on language* (pp. 172–193). Cambridge, MA: MIT Press.
- Gatt, A. and 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
- Gatt, A., Van Gompel, R.P.G., Van Deemter, K., & Krahmer, E. (2013). Are we Bayesian referring expression generators? In *Proceedings of the 3rd PRE-Cogsci meeting*, Berlin, Germany.
- Goldberg, E., Driedger, N., & Kittredge, R. (1994). Using natural-language processing to produce weather forecasts. *IEEE Expert*, 9, 45–53.
- Goldrick, M., Ferreira, V., & Miozzo, M. (2014). *The Oxford Handbook of Language Production*. Oxford: Oxford University Press.
- Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, 20, 818–829.
- Goodman, N.D. & Stuhlmüller (2013). Knowledge and implicature: Modeling language understanding as social cognition. *Topics in Cognitive Science*, 5, 173-184.
- Goudbeek, M., & Krahmer, E. (2012). Alignment in interactive reference production: Content planning, modifier ordering and referential overspecification. *Topics in Cognitive Science*, 4, 269-289.
- Grice, H. (1975). Logic and conversation. In P. Cole & J. Morgan (Eds.), *Syntax and semantics: Speech acts* (Vol. III, pp. 41–58). New York: Academic Press.

- Heller, D., Gorman, K. S., & Tanenhaus, M. K. (2012). To name or to describe: Shared knowledge affects referential form. *Topics in Cognitive Science*, 44, 290–305.
- Horton, W.S., & Keysar, B. (1996). When do speakers take into account common ground?. *Cognition* 59, 91-117.
- Johnson-Laird, P.N. (1983). *Mental models*. Cambridge: Cambridge University Press.
- Kazemzadeh, S. Ordonez, V., Matten, B., Berg, T.L. (2014). Referring to Objects in Natural Scenes. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP-2014)*, Doha, Qatar, 787-798.
- Koolen, R., Gatt, A., Goudbeek, M., & Krahmer, E. (2011). Factors causing overspecification in definite descriptions. *Journal of Pragmatics*, 43, 3231-3250.
- Koolen, R., Goudbeek, M. & Krahmer, E. (2013). The effect of scene variation on the redundant use of color in definite reference. *Cognitive Science*, 37, 395-411.
- Krahmer, E., & Van Deemter, K. (2012). Computational Generation of Referring Expressions: A Survey. *Computational Linguistics*, 38, 173-218.
- Levelt, W.J.M. (2013). *A History of Psycholinguistics: The Pre-Chomskyan Era*. Oxford: Oxford University Press.
- Levelt, W.J.M., Roelofs, A., & Meyer, A.S. (1999). A theory of lexical access in speech production. *Behavioral and Brain Sciences*, 22, 1–75.
- Lewandowsky, S., & Farrell, S. (2010). *Computational Modeling in Cognition: Principles and Practice*. Thousand Oaks, CA: Sage.
- Maes, A., Arts, A., & Noordman, L. (2004). Reference management in instructive discourse. *Discourse Processes*, 37, 117-144.
- Mao, J., Huang, J., Toshev, A., Camburu, O., Yuille, A., & Murphy, K. (2016). Generation and Comprehension of Unambiguous Object Descriptions. In *Proceedings of the 2016 Conference on Computer Vision and Pattern Recognition (CVPR'16)*.

- Matthews D., Lieven E.V., Theakston A.L., & Tomasello, M. (2006). The effect of perceptual availability and prior discourse on young children's use or referring expressions. *Applied Psycholinguistics*, 27, 403-422.
- Mellish, C., Scott, D., Cahill, L., Paiva, D., Evans, R., & Reape, M. (2006). A reference architecture for natural language generation systems. *Natural Language Engineering*, 12, 1-34.
- Meyer, A.S. (1996). Lexical access in phrase and sentence production: Results from picture-word interference experiments. *Journal of Memory and Language*, 35, 477-496.
- Paraboni, I. & Van Deemter, K. (2014). Reference and the facilitation of search in spatial domains. *Language and Cognitive Processes*, 29, 1002-1007.
- Pechmann, T. (1989). Incremental speech production and referential overspecification. *Linguistics*, 27, 98-110.
- Portet, F., Reiter, E., Gatt, A., Hunter, J., Sripada, S., Freer, Y., & Sykes, C. (2009). Automatic generation of textual summaries from neonatal intensive care data. *Artificial Intelligence*, 173, 789-816.
- Reiter, E., & Dale, R. (2000). *Building natural language generation systems*. Cambridge, UK: Cambridge University Press.
- Rossion, B., & Pourtois, G. (2004). Revisiting Snodgrass and Vanderwart's object databank : The role of surface detail in basic level object recognition. *Perception* 33, 217-236.
- Rubio-Fernandez, P. (2016). How redundant are redundant color adjectives? An efficiency-based analysis of color overspecification. *Frontiers in Psychology*, 7 (153).
- Schwarz, G.E. (1978), Estimating the dimension of a model, *Annals of Statistics*, 6, 461-464.
- Snodgrass, J. G., & Vanderwart, M. (1980). A standardized set of 260 pictures: Norms for name agreement, image agreement, familiarity and visual complexity. *Journal of Experimental Psychology: Human Learning and Memory* 6, 174-215.

- Tarenskeen, S., Broersma, M., & Geurts, B. (2015). Overspecification of color, pattern, and size: salience, absoluteness, and consistency. *Frontiers in Psychology*, 6, 1703.
- Van Deemter, K. (2016). *Computational models of referring: a study in Cognitive Science*. Cambridge MA: MIT Press.
- Van Deemter, K., Gatt, A., Van der Sluis, I., & Power, R. (2012). Generation of referring expressions: Assessing the incremental algorithm. *Cognitive Science*, 36, 799-836.
- Van Deemter, K., Gatt, A., Van Gompel, R.P.G., & Krahmer, E. (2012). Towards a computational psycholinguistics of reference production. *Topics in Cognitive Science*, 4, 166-184.
- Van Gompel, R.P.G., Gatt, A., Krahmer, E., & Van Deemter, K. (2012). PRO: A computational model of referential overspecification. Poster presented at the 18th AMLaP conference, Riva del Garda, Italy.
- Van Gompel, R.P.G., Gatt, A., Krahmer, E., & Van Deemter, K. (2014). Overspecification in reference: Modelling size contrast effects. Poster presented at the 20th AMLaP conference, Edinburgh, United Kingdom.
- Viethen, J., & Dale, R. (2010). Speaker-dependent variation in content selection for referring expression generation. In *Proceedings of the 8th Australasian Language Technology Workshop (ALTA 2010)*, Melbourne, Australia.
- Viethen, J., Van Vesseem, T., Goudbeek, M., & Krahmer, E. (2017). Color in reference production: The role of color similarity and color codability. *Cognitive Science*, 41, 1493-1514.
- Vigliocco, G., & Hartsuiker, R.J. (2002). The interplay of meaning, sound, and syntax in sentence production. *Psychological Bulletin*, 128, 442-472.

- Wardlow-Lane, L., & Ferreira, V. S. (2008). Speaker external vs speaker-internal forces on utterance form. Do cognitive demands override threats to referential success? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34, 1466–1481.
- Wasserman, L. (2000). Bayesian model selection and model averaging. *Journal of Mathematical Psychology*, 44, 359-397.
- Westerbeek, H., Koolen, R., & Maes, A. (2015). Stored object knowledge and the production of referring expressions: The case of color typicality. *Frontiers in Psychology*, 6, 1–12.
- Winograd, T. (1972). *Understanding Natural Language*. Academic Press, New York.
- Xu, F., & Tenenbaum, J. (2007). Word learning as Bayesian inference. *Psychological Review*, 114, 245-272.
-

¹ Following a common convention, we will refer to speakers with “she” and listeners with “he”.

² Here and elsewhere, we refer to colour and size as attributes, while “red” and “small” (values of respectively the colour and size attributes) are properties of an object.

Table 1: Formulas of the models in Experiment 1.

	description		
	colour and size	colour-only	size-only
Non-deterministic incremental algorithm			
colour-or-size condition	0	c	$1-c$
size-only condition	c	0	$1-c$
colour-only condition	$1-c$	c	0
Modified non-deterministic incremental algorithm			
colour-or-size condition	$c*(1-c)+(1-c)*c$	$c*c$	$(1-c)*(1-c)$
size-only condition	$c+(1-c)*c$	0	$(1-c)*(1-c)$
colour-only condition	$c*(1-c)+1-c$	$c*c$	0
PRO			
colour-or-size condition	$c*(1-c+e)+(1-c)*(c+e)$	$c*(c-e)$	$(1-c)*(1-c-e)$
size-only condition	$1*(c+e)$	0	$1*(1-c-e)$
colour-only condition	$1*(1-c+e)$	$1*(c-e)$	0

Table 2: Parameter values for the predictions in each language (Experiment 1).

	English		Dutch	
	<i>c</i>	<i>e</i>	<i>c</i>	<i>e</i>
Non-deterministic incremental algorithm				
colour-or-size condition	.8799		.8389	
size-only condition	.9376		.9240	
colour-only condition	.8916		.8381	
Modified non-deterministic incremental algorithm				
colour-or-size condition	.8454		.7969	
size-only condition	.9180		.8944	
colour-only condition	.7867		.7164	
PRO				
colour-or-size condition	.8777	-.0449	.8360	-.0786
size-only condition	.8473	-.0789	.8122	-.0911
colour-only condition	.8664	-.0510	.8161	-.0631

Table 3: Model testing results for Experiment 1.

	<i>p(data/model)</i>	<i>BIC</i>
English		
Non-deterministic IA	$4.70 \cdot 10^{-179}$	826.16
Modified non-deterministic IA	$1.43 \cdot 10^{-103}$	480.55
PRO	$2.99 \cdot 10^{-12}$	66.93
Dutch		
Non-deterministic IA	$1.35 \cdot 10^{-239}$	1106.91
Modified non-deterministic IA	$1.16 \cdot 10^{-135}$	628.28
PRO	$7.56 \cdot 10^{-9}$	51.14

Table 4: Formulas of the models in Experiment 2.

	description		
	type and colour	type-only	colour-only
Non-deterministic incremental algorithm			
type-or-colour condition	c	$1-c$	0
type-only condition	c	$1-c$	0
colour-only condition	1	0	0
Modified non-deterministic incremental algorithm			
type-or-colour condition	$c*(1-c)+(1-c)*c$	$(1-c)*(1-c)$	$c*c$
type-only condition	$c+(1-c)*c$	$(1-c)*(1-c)$	0
colour-only condition	$c*(1-c)+1-c$	0	$c*c$
PRO			
type-or-colour condition	$c*(1-c+e)+(1-c)*(c+e)$	$(1-c)*(1-c-e)$	$c*(c-e)$
type-only condition	$1*(c+e)$	$1*(1-c-e)$	0
colour-only condition	$1*(1-c+e)$	0	$1*(c-e)$

Table 5: Parameter values for the predictions in each language (Experiment 2).

	English		Dutch	
	<i>c</i>	<i>e</i>	<i>c</i>	<i>e</i>
Non-deterministic IA				
type-or-colour condition	.3136		.2507	
type-only condition	.4859		.3601	
colour-only condition	.3997		.3045	
Modified non-deterministic IA				
type-or-colour condition	.1914		.1563	
type-only condition	.2477		.2051	
colour-only condition	.2125		.1679	
PRO				
type-or-colour condition	.1864	.1272	.1503	.1004
type-only condition	.2065	.1508	.1602	.0953
colour-only condition	.1932	.1451	.1834	.0604

Note. IA = incremental algorithm; PRO = probabilistic referential overspecification model.

Table 6: Model testing results for Experiment 2.


	$p(\text{data}/\text{model})$	BIC
English		
Non-deterministic IA	$4.89 \cdot 10^{-69}$	321.54
Modified non-deterministic IA	$1.37 \cdot 10^{-21}$	103.03
PRO	$8.25 \cdot 10^{-7}$	41.93
Dutch		
Non-deterministic IA	$1.57 \cdot 10^{-64}$	330.77
Modified non-deterministic IA	$6.95 \cdot 10^{-20}$	95.15
PRO	$2.42 \cdot 10^{-17}$	90.43

Table 7: Parameter values for the predictions in Experiment 3.

	parameter values used for predictions	
	<i>c</i>	<i>s</i>
Modified non-deterministic IA		
colour-only condition	.6145	.1283
size-only condition	.6557	.0860
border-only condition	.6814	.1316
colour-or-size condition	.6241	.1195
colour-or-border condition	.6209	.1246
size-or-border condition	.6817	.0907
colour-size-or-border condition	.6306	.1178
PRO		
colour-only condition	.6633	.1130
size-only condition	.6596	.1115
border-only condition	.6583	.1144
colour-or-size condition	.6592	.1160
colour-or-border condition	.6437	.1103
size-or-border condition	.6591	.0910
colour-size-or-border condition	.6450	.1131



Fig. 1: Three brooms; the target broom is circled. Note that both colour and size distinguish the target from its distractors. The colours of the objects are indicated in brackets.



a. Either colour or size fully distinguishes the target from the distractors.



b. Only size fully distinguishes the target from the distractors.



c. Only colour fully distinguishes the target from the distractors.

Fig. 2: Three examples of a target (circled) and distractors. The colours of the objects are indicated in brackets; participants saw the actual colours, not the words in brackets.

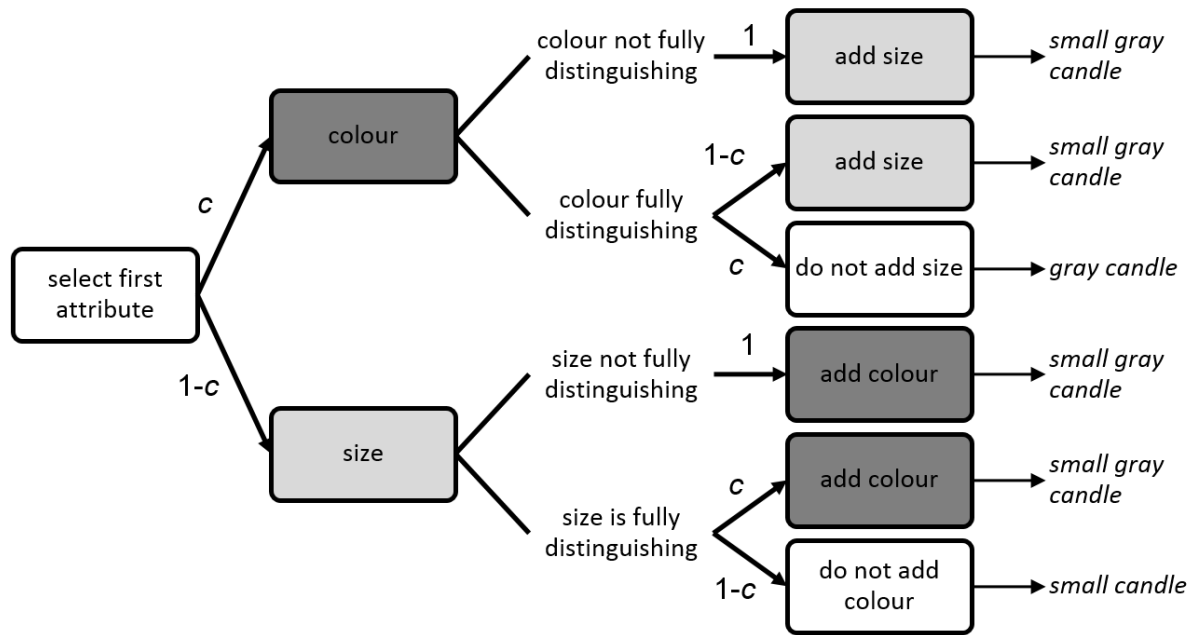


Fig. 3: Modified non-deterministic IA decision tree for the conditions in Experiment 1.

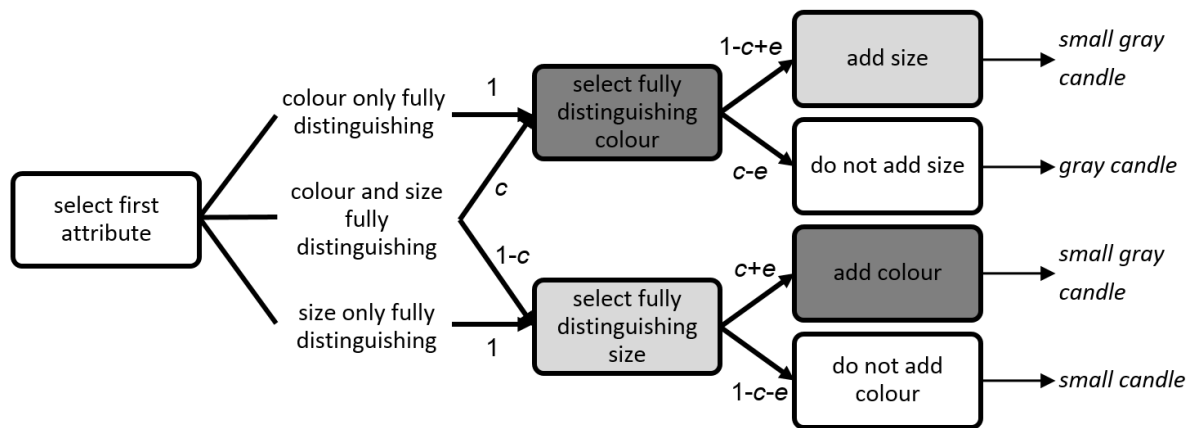


Fig. 4: PRO decision tree for the conditions in Experiment 1.

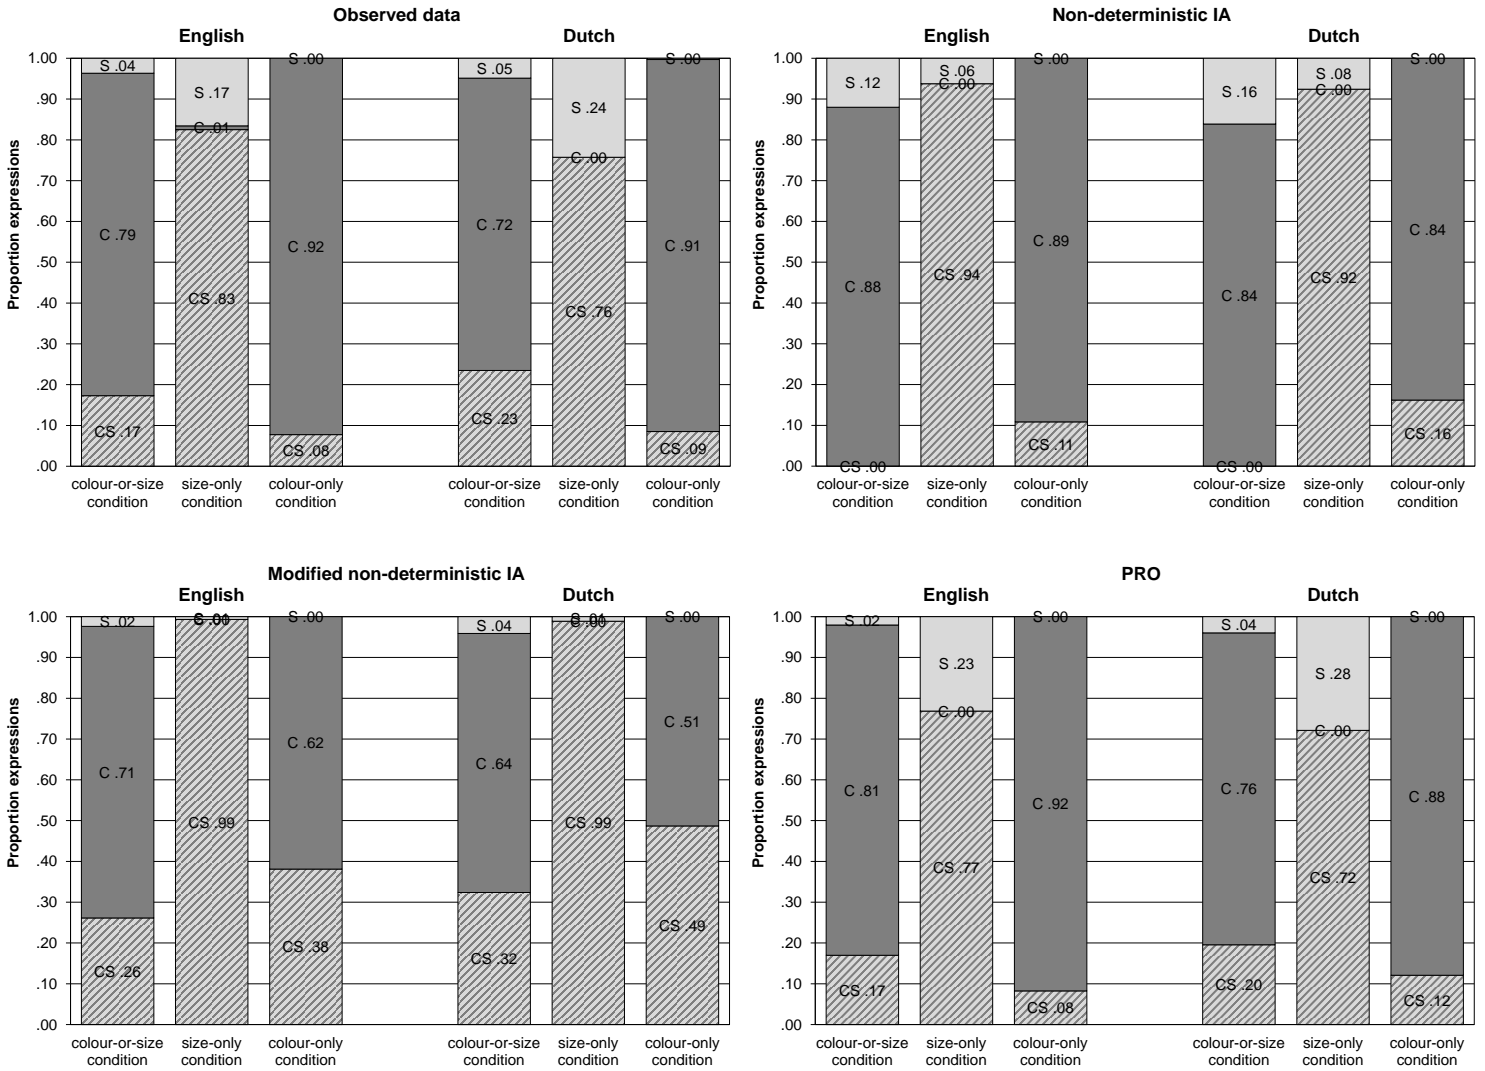




Fig. 5: Observed results and model predictions for Experiment 1. C = colour-only descriptions, S = size-only descriptions, CS = colour-and-size descriptions.



a. Either type or colour fully distinguishes the target from the distractors.



b. Only colour fully distinguishes the target from the distractors.



c. Only type fully distinguishes the target from the distractors.

Fig. 6: Three examples of a target (circled) and distractors in Experiment 2. The colours of the objects are indicated in brackets; participants saw the actual colours, not the words in brackets.

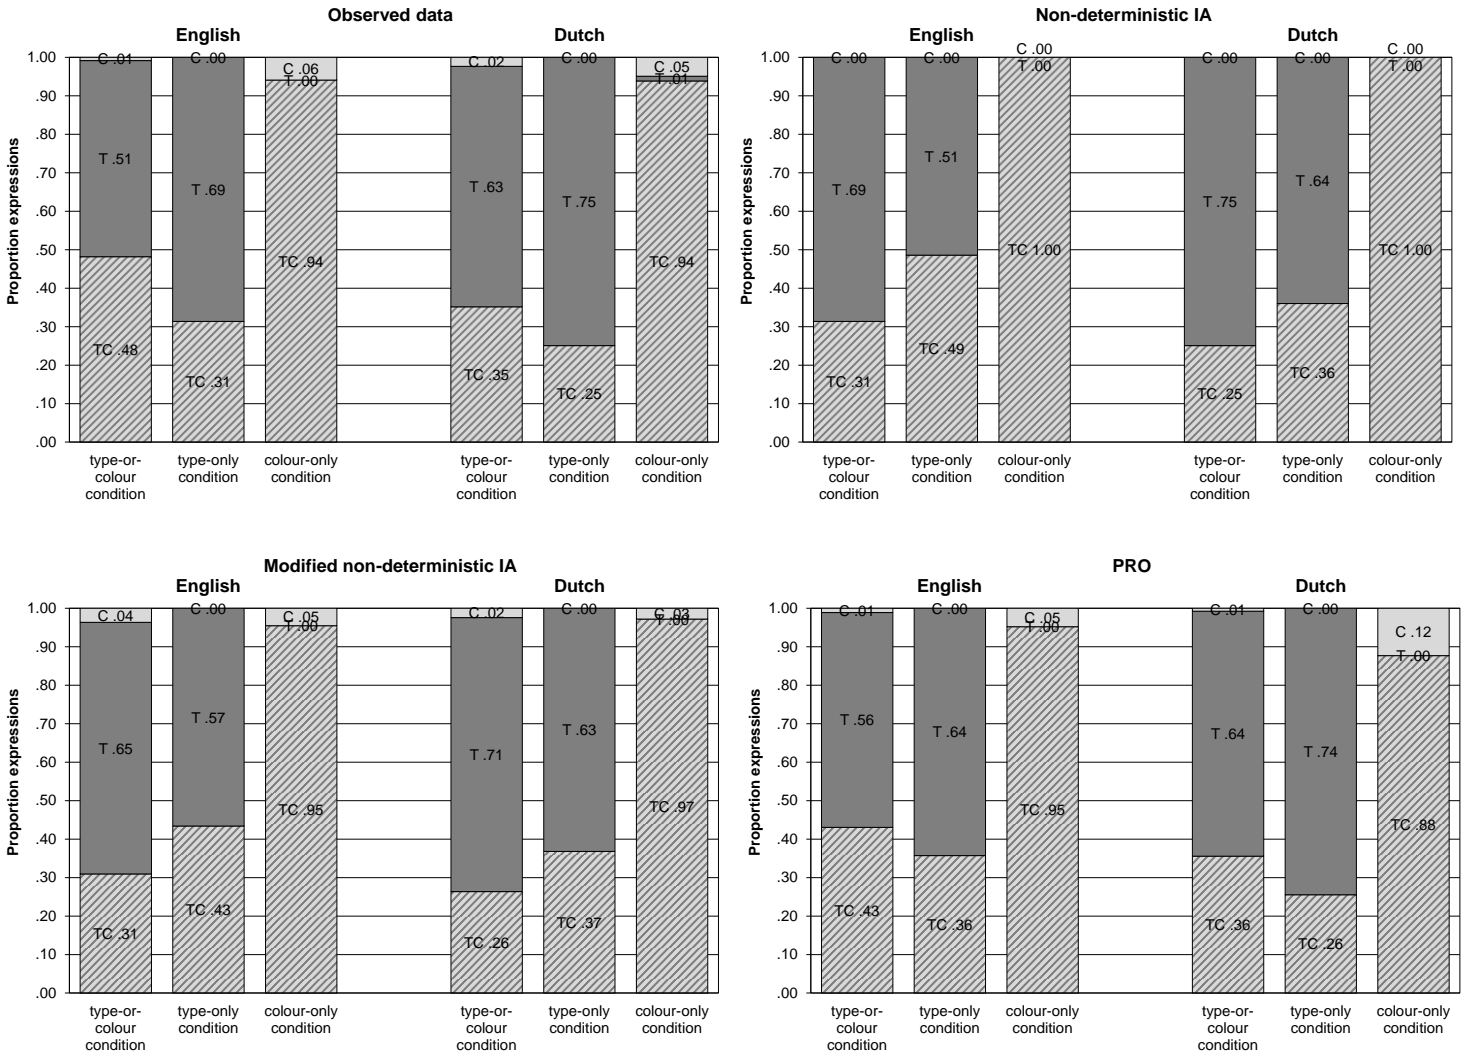





Fig. 7: Observed results and model predictions for Experiment 2. C = colour-only descriptions, T = type-only descriptions, TC = type-and-colour descriptions.




a. Only colour fully distinguishes the target from the distractors.




b. Only size fully distinguishes the target from the distractors.




c. Only border fully distinguishes the target from the distractors.




d. Both colour and size fully distinguish the target from the distractors.



e. Both colour and border fully distinguish the target from the distractors.



f. Both size and border fully distinguish the target from the distractors.



g. Colour, size or border all fully distinguish the target from the distractors.

Fig. 8. Seven examples of a target (indicated by an arrow) and distractors. The colours of the objects are indicated in brackets; participants saw the actual colours, not the words.




Fig. 9: Observed results and model predictions for Experiment 3. C = colour-only descriptions, S = size-only descriptions, B = border-only descriptions, CS = colour-and-size descriptions, CB = colour-and-border descriptions, SB = size-and-border descriptions, CSB = colour-size-and-border descriptions.

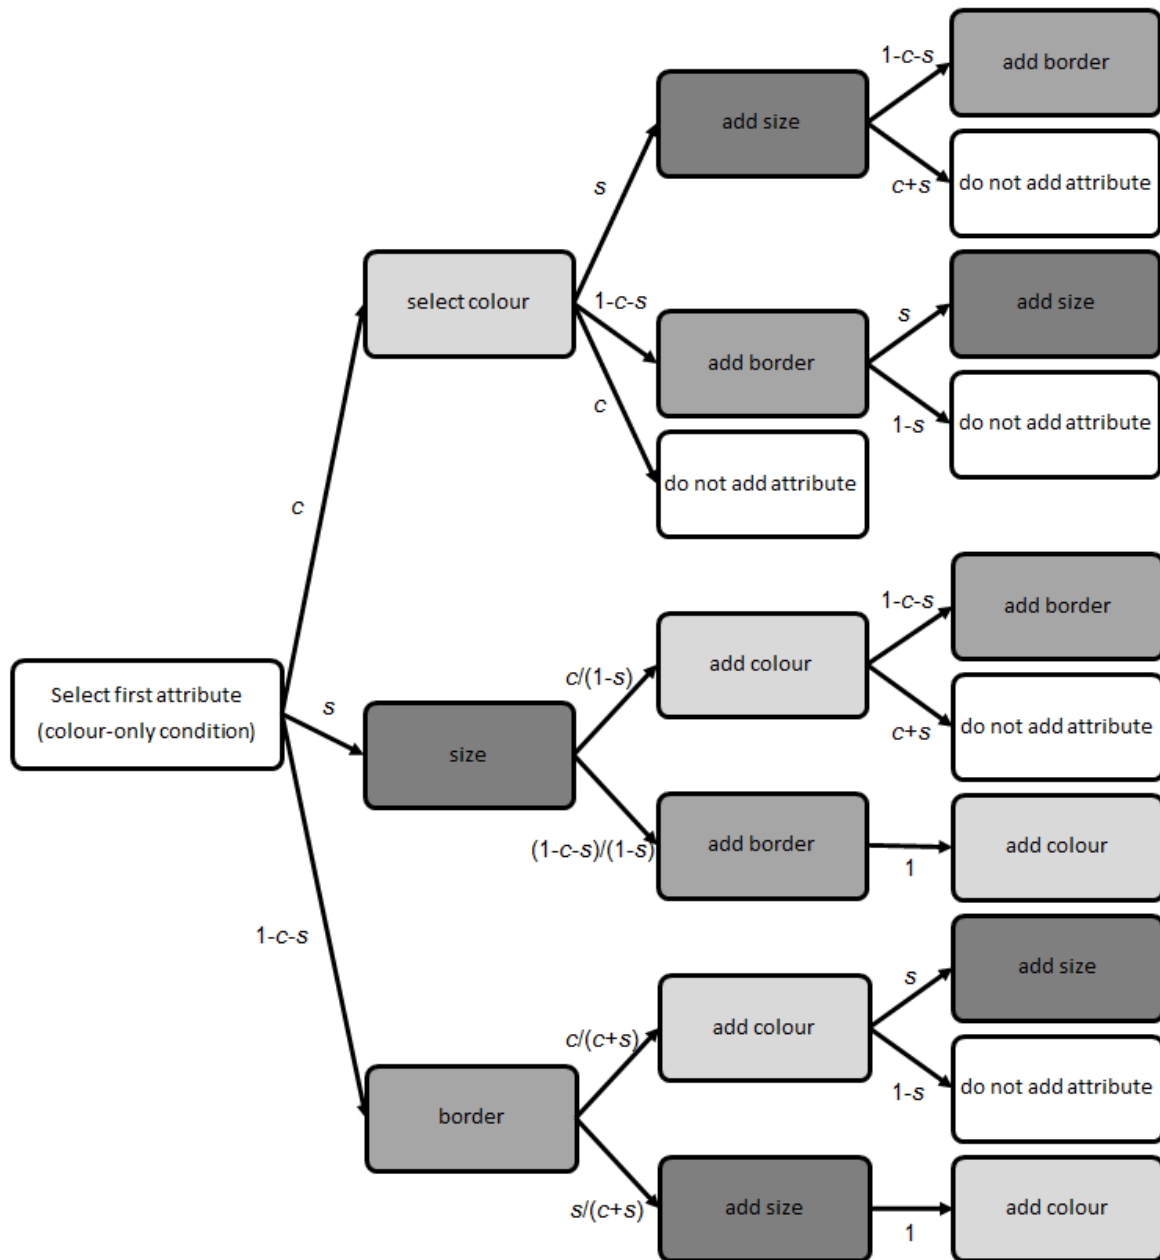


Fig. 10. Modified non-deterministic IA decision tree for the colour-only condition (Fig. 8a) in Experiment 3.

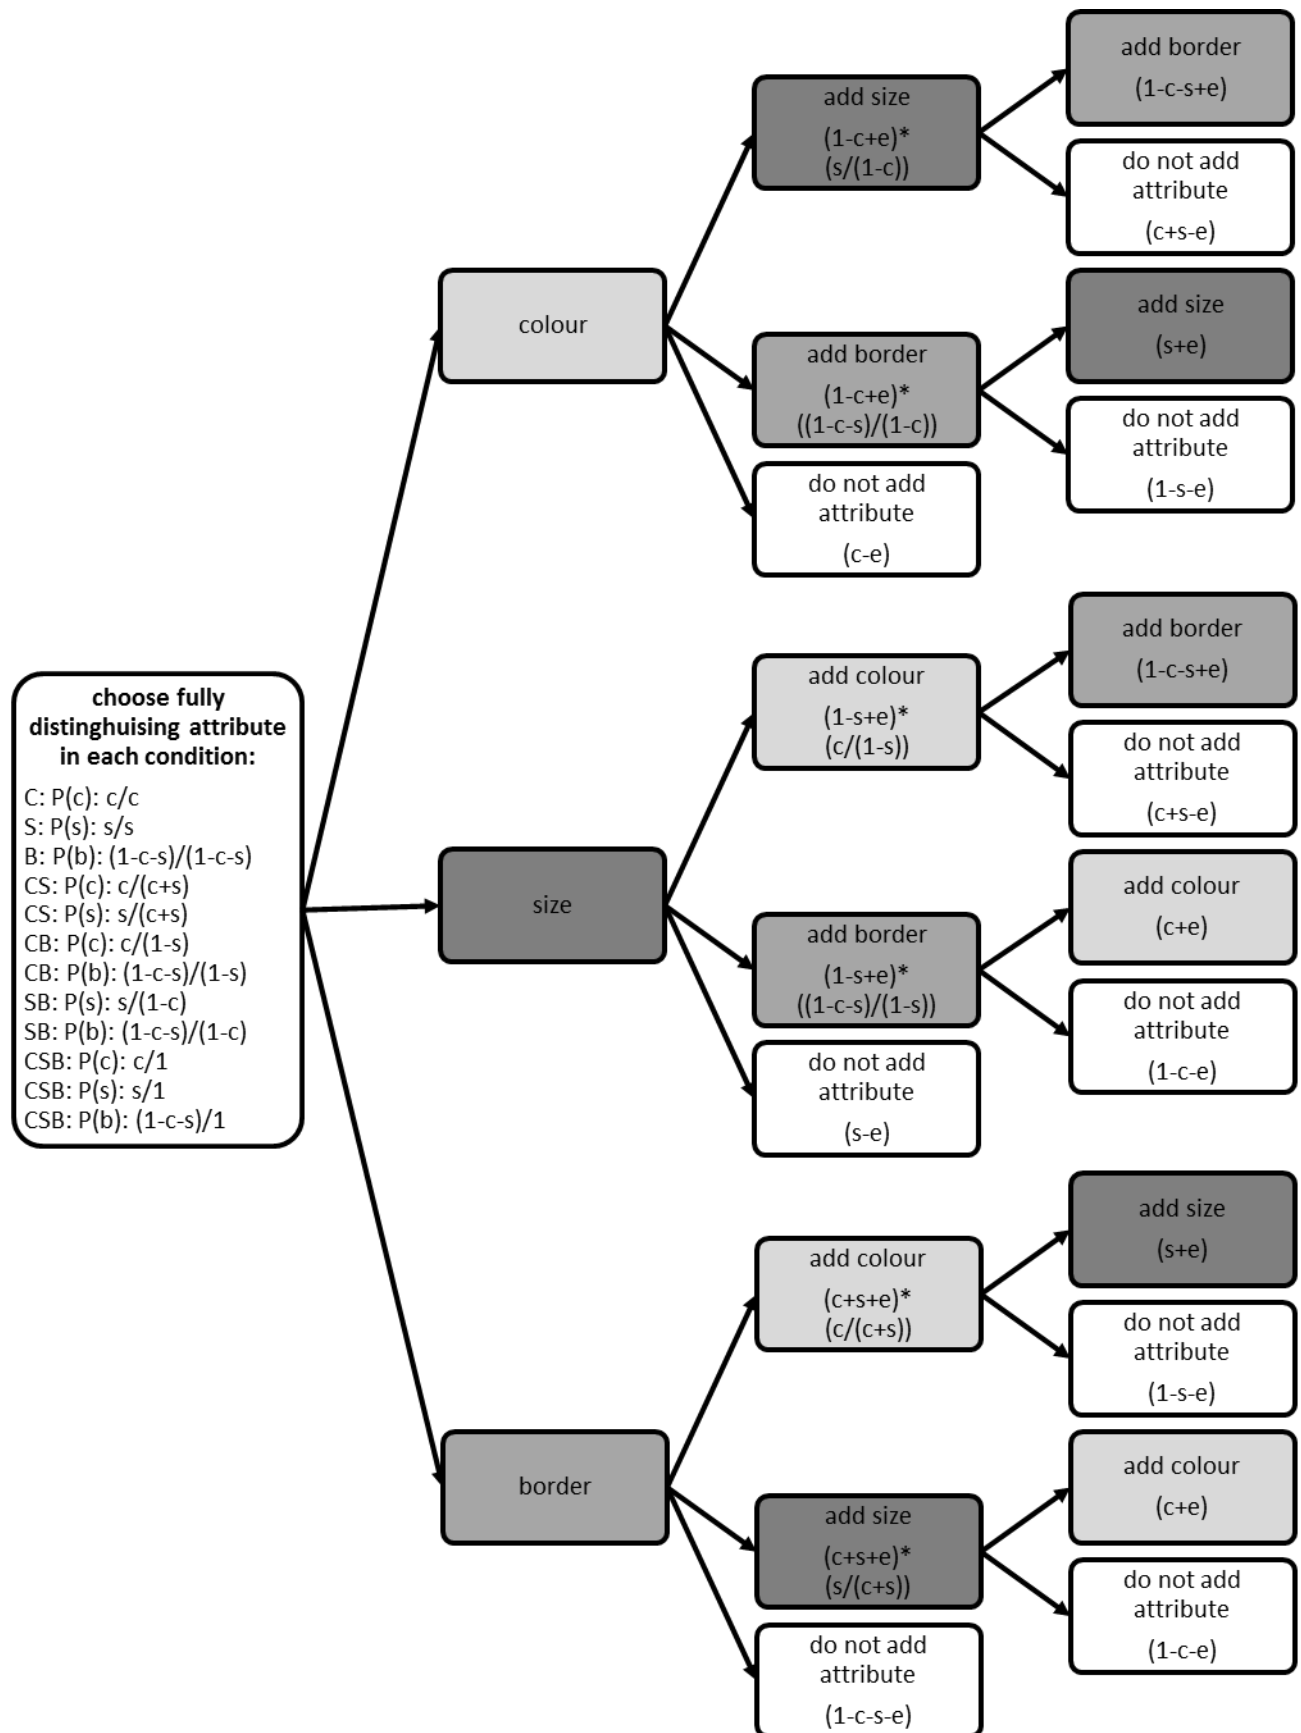


Fig. 11. PRO decision tree for the conditions in Fig. 8.

Appendix 1: PRO algorithm

Input: A domain of objects, containing a target referent r and a non-empty set M of distractors. A set \mathbf{A} of attributes true of r , each of which is false for at least one distractor in M . Furthermore, for each attribute, a probability that reflects the Preference degree associated with the attribute (this is to inform *P-choose*); and a probability that determines the likelihood of overspecification (this is to inform *R-choose*; details below).

Output: If a distinguishing description of r (given \mathbf{A} and M) exists then the algorithm returns one of these as the value of D .

```

1: if some property in  $\mathbf{A}$  individually suffices to remove all distractors then
2:   P-choose one such property, for example  $P$ 
3:    $D := \{P\}$ 
4: while True do
5:   if  $\mathbf{A} = \text{empty}$  then
6:     return  $D$ 
7:   else if  $D$  is not distinguishing yet then
8:     P-choose a property from  $\mathbf{A}$ 
9:     update  $D$ ,  $\mathbf{A}$ , and  $M$ 
10:  else
11:    R-choose between STOP and the  $n$  properties remaining in  $\mathbf{A}$ 
      (making  $n+1$  options in total)
12:    if a further property is chosen then
13:      Update  $D$ ,  $\mathbf{A}$ , and  $M$ 
14:    else
15:      return  $D$ 

```

The PRO algorithm, sketched in pseudo-code above, iterates through a number of steps, each of which either adds an attribute or none at all. The algorithm has two phases. The first phase (lines 1-3) checks whether there is one single attribute that is able to identify the referent; if there is, then such an attribute is included in the description. If there is more than one attribute that is fully discriminatory, then one is probabilistically included into the description (*P-choose*, line 2). *P-choose* (short for Property-choose) is the function that selects properties probabilistically, in proportion to their relative degree of preference. *P-choose* uses $p - 1$ preference values, where p is the number of fully distinguishing attributes (that have not yet been included in the description). For example, there is one preference value in the colour-or-

size condition of Experiment 1 (c) and two in the colour-size-or-border condition of Experiment 3 (c and s). The preference values are normalised over all fully distinguishing attributes, so in the colour-or-size condition of Experiment 3, the colour preference is $c/(c+s)$ and the size preference is $s/(c+s)$ (border is not fully distinguishing in this condition).

The second phase of the algorithm (lines 4-15) takes place whether or not the first phase has found an attribute to include: Either way, this second phase keeps looking for opportunities to include (further) attributes in the description (line 4) until a description is returned. \mathbf{A} is the set of all available properties that hold true of the referent r but are false for at least one distractor; \mathbf{A} shrinks as properties are added to the set of properties \mathbf{D} that is generated. \mathbf{M} is the set of distractors, which likewise gets smaller. If there are no remaining properties left in \mathbf{A} (a situation that did not occur in our experiments), then the description \mathbf{D} cannot be updated and the algorithm outputs \mathbf{D} (lines 5-6).

If \mathbf{D} does not yet distinguish the target from all distractors, then another attribute is added, with the probability determined by its preference value (lines 7-9). As before, P-choose uses $p - 1$ preference values, where p is the number of remaining distinguishing attributes and the preference values are normalised over these attributes. As a result of the “while **True**” loop (line 4), there may be several iterations of P-choose before the attributes fully distinguish the target from the distractors. If \mathbf{D} is distinguishing, then the algorithm probabilistically chooses either to stop and return \mathbf{D} , or to add a further attribute (lines 11-15). The decision to stop or continue is determined by R-Choose (short for Redundancy-choose), which reflects both an attribute’s preference and the speaker’s overspecification eagerness.

Updating \mathbf{D} , \mathbf{A} and \mathbf{M} : Whenever a property \mathbf{P} is chosen, then \mathbf{D} is updated by adding \mathbf{P} as an element of \mathbf{D} ; \mathbf{A} is updated by removing \mathbf{P} from \mathbf{A} (because it is no longer available);

and M is updated by intersecting M with the set of all domain objects for which the property is true (i.e., by removing distractors).

R-Choose uses $r - 1$ preference values, with r being the number of remaining attributes that rule out at least one distractor. Again, because of the “while **True**” loop, the algorithm may go through several iterations of R-Choose. As explained in Experiment 3, during each iteration, the probability of adding a particular attribute (e.g., colour, size or border) is the probability of adding *any* attribute multiplied by the preference value for that particular attribute normalised over all remaining attributes.

Appendix 2: Formulas used in Experiment 3**Modified non-deterministic IA**

Colour-only condition

$$C = c * c$$

$$CS = c * s * (c + s) + s * (c / (1 - s)) * (c + s)$$

$$CB = c * (1 - c - s) * (1 - s) + (1 - c - s) * (c / (c + s)) * (1 - s)$$

$$CSB = c * s * (1 - c - s) + c * (1 - c - s) * s + s * (c / (1 - s)) * (1 - c - s) + s * ((1 - c - s) / (1 - s)) * 1 + (1 - c - s) * (c / (c + s)) * s + (1 - c - s) * (s / (c + s)) * 1$$

Size-only condition

$$S = s * s$$

$$CS = c * (s / (1 - c)) * (c + s) + s * c * (c + s)$$

$$SB = s * (1 - c - s) * (1 - c) + (1 - c - s) * (s / (c + s)) * (1 - c)$$

$$CSB = c * (s / (1 - c)) * (1 - c - s) + c * ((1 - c - s) / (1 - c)) * 1 + s * c * (1 - c - s) + s * (1 - c - s) * c + (1 - c - s) * (c / (c + s)) * 1 + (1 - c - s) * (s / (c + s)) * c$$

Border-only condition

$$B = (1 - c - s) * (1 - c - s)$$

$$CB = (1 - c - s) * c * (1 - s) + c * ((1 - c - s) / (1 - c)) * (1 - s)$$

$$SB = (1 - c - s) * s * (1 - c) + s * ((1 - c - s) / (1 - s)) * (1 - c)$$

$$CSB = c * (s / (1 - c)) * 1 + c * ((1 - c - s) / (1 - c)) * s + s * (c / (1 - s)) * 1 + s * ((1 - c - s) / (1 - s)) * c + (1 - c - s) * s * c + (1 - c - s) * c * s$$

Colour-or-size condition

$$C = c * c$$

$$S = s * s$$

$$CS = c * s * (c + s) + s * c * (c + s)$$

$$CB = c * (1 - c - s) * (1 - s) + (1 - c - s) * (c / (c + s)) * (1 - s)$$

$$SB = s * (1 - c - s) * (1 - c) + (1 - c - s) * (s / (c + s)) * (1 - c)$$

$$CSB = c * s * (1 - c - s) + c * (1 - c - s) * s + s * c * (1 - c - s) + s * (1 - c - s) * c + (1 - c - s) * (c / (c + s)) * s + (1 - c - s) * (s / (c + s)) * c$$

Colour-or-border condition

$$C = c * c$$

$$B = (1 - c - s) * (1 - c - s)$$

$$CS = c * s * (c + s) + s * (c / (1 - s)) * (c + s)$$

$$CB = c * (1 - c - s) * (1 - s) + (1 - c - s) * c * (1 - s)$$

$$SB = s * ((1 - c - s) / (1 - s)) * (1 - c) + (1 - c - s) * s * (1 - c)$$

$$CSB = c * s * (1 - c - s) + c * (1 - c - s) * s + s * (c / (1 - s)) * (1 - c - s) + s * ((1 - c - s) / (1 - s)) * c + (1 - c - s) * s * c + (1 - c - s) * c * s$$

Size-or-border condition

$$S = s*s$$

$$B = (1-c-s)*(1-c-s)$$

$$CS = c*(s/(1-c))*(c+s)+s*c*(c+s)$$

$$CB = c*((1-c-s)/(1-c))*(1-s)+(1-c-s)*c*(1-s)$$

$$SB = s*(1-c-s)*(1-c)+(1-c-s)*s*(1-c)$$

$$CSB = c*(s/(1-c))*(1-c-s)+c*((1-c-s)/(1-c))*s+s*c*(1-c-s)+s*(1-c-s)*c+(1-c-s)*c*s+(1-c-s)*s*c$$

Colour-size-or-border condition

$$C = c*c$$

$$S = s*s$$

$$B = (1-c-s)*(1-c-s)$$

$$CS = c*s*(c+s)+s*c*(c+s)$$

$$CB = c*(1-c-s)*(1-s)+(1-c-s)*c*(1-s)$$

$$SB = s*(1-c-s)*(1-c)+(1-c-s)*s*(1-c)$$

$$CSB = c*s*(1-c-s)+c*(1-c-s)*s+s*c*(1-c-s)+s*(1-c-s)*c+(1-c-s)*c*s+(1-c-s)*s*c$$

PRO

Colour-only condition

$$C = (c-e)$$

$$CS = (1-c+e)*(s/(1-c))*(c+s-e)$$

$$CB = (1-c+e)*((1-c-s)/(1-c))*(1-s-e)$$

$$CSB = (1-c+e)*(s/(1-c))*(1-c-s+e)+(1-c+e)*((1-c-s)/(1-c))*(s+e)$$

Size-only condition

$$S = (s-e)$$

$$CS = (1-s+e)*(c/(1-s))*(c+s-e)$$

$$SB = (1-s+e)*((1-c-s)/(1-s))*(1-c-e)$$

$$CSB = (1-s+e)*(c/(1-s))*(1-c-s+e)+(1-s+e)*((1-c-s)/(1-s))*(c+e)$$

Border-only condition

$$B = (1-c-s-e)$$

$$CB = (c+s+e)*(c/(c+s))*(1-s-e)$$

$$SB = (c+s+e)*(s/(c+s))*(1-c-e)$$

$$CSB = (c+s+e)*(c/(c+s))*(s+e)+(c+s+e)*(s/(c+s))*(c+e)$$

Colour-or-size condition

$$C = (c/(c+s))*(c-e)$$

$$S = (s/(c+s))*(s-e)$$

$$CS = (c/(c+s))*(1-c+e)*(s/(1-c))*(c+s-e)+(s/(c+s))*(1-s+e)*(c/(1-s))*(c+s-e)$$

$$CB = (c/(c+s))*(1-c+e)*((1-c-s)/(1-c))*(1-s-e)$$

$$SB = (s/(c+s))*(1-s+e)*((1-c-s)/(1-s))*(1-c-e)$$

$$CSB = (c/(c+s))*(1-c+e)*(s/(1-c))*(1-c-s+e)+(c/(c+s))*(1-c+e)*((1-c-s)/(1-c))*(s+e)+(s/(c+s))*(1-s+e)*(c/(1-s))*(1-c-s+e)+(s/(c+s))*(1-s+e)*((1-c-s)/(1-s))*(c+e)$$

Colour-or-border condition

$$C = (c/(1-s))*(c-e)$$

$$B = ((1-c-s)/(1-s))*(1-c-s-e)$$

$$CS = (c/(1-s))*(1-c+e)*(s/(1-c))*(c+s-e)$$

$$CB = (c/(1-s))*(1-c+e)*((1-c-s)/(1-c))*(1-s-e)+((1-c-s)/(1-s))*(c+s+e)*(c/(c+s))*(1-s-e)$$

$$SB = ((1-c-s)/(1-s))*(c+s+e)*(s/(c+s))*(1-c-e)$$

$$CSB = (c/(1-s))*(1-c+e)*(s/(1-c))*(1-c-s+e)+(c/(1-s))*(1-c+e)*((1-c-s)/(1-c))*(s+e)+((1-c-s)/(1-s))*(c+s+e)*(c/(c+s))*(s+e)+((1-c-s)/(1-s))*(c+s+e)*(s/(c+s))*(c+e)$$

Size-or-border condition

$$S = (s/(1-c))*(s-e)$$

$$B = ((1-c-s)/(1-c))*(1-c-s-e)$$

$$CS = (s/(1-c))*(1-s+e)*(c/(1-s))*(c+s-e)$$

$$CB = ((1-c-s)/(1-c))*(c+s+e)*(c/(c+s))*(1-s-e)$$

$$SB = (s/(1-c))*(1-s+e)*((1-c-s)/(1-s))*(1-c-e)+((1-c-s)/(1-c))*(c+s+e)*(s/(c+s))*(1-c-e)$$

$$CSB = (s/(1-c))*(1-s+e)*(c/(1-s))*(1-c-s+e)+(s/(1-c))*(1-s+e)*((1-c-s)/(1-s))*(c+e)+((1-c-s)/(1-c))*(c+s+e)*(c/(c+s))*(s+e)+((1-c-s)/(1-c))*(c+s+e)*(s/(c+s))*(c+e)$$

Colour-size-or-border condition

$$C = c*(c-e)$$

$$S = s*(s-e)$$

$$B = (1-c-s)*(1-c-s-e)$$

$$CS = c*(1-c+e)*(s/(1-c))*(c+s-e)+s*(1-s+e)*(c/(1-s))*(c+s-e)$$

$$CB = c*(1-c+e)*((1-c-s)/(1-c))*(1-s-e)+(1-c-s)*(c+s+e)*(c/(c+s))*(1-s-e)$$

$$SB = s*(1-s+e)*((1-c-s)/(1-s))*(1-c-e)+(1-c-s)*(c+s+e)*(s/(c+s))*(1-c-e)$$

$$CSB = c*(1-c+e)*(s/(1-c))*(1-c-s+e)+c*(1-c+e)*((1-c-s)/(1-c))*(s+e)+s*(1-s+e)*(c/(1-s))*(1-c-s+e)+s*(1-s+e)*((1-c-s)/(1-s))*(c+e)+(1-c-s)*(c+s+e)*(c/(c+s))*(s+e)+(1-c-s)*(c+s+e)*(s/(c+s))*(c+e)$$