

Adapting the use of attributes to the task environment in joint action: results and a model

Markus Guhe

University of Edinburgh
Linguistics and English Language & HCRC
40 George Square
Edinburgh, EH8 9LL
m.guhe@ed.ac.uk

Ellen Gurman Bard

University of Edinburgh
Linguistics and English Language & HCRC
40 George Square
Edinburgh, EH8 9LL
ellen@ling.ed.ac.uk

Abstract

Speakers use referring expressions to identify an object in the environment. To generate a referring expression, features of the intended referent have to be selected that distinguish the object from the other potential referents. Current accounts of referring expressions consider a number of factors that influence the choice of features but ignore the influences of the task environment. In particular, they do not address how these influences change the generation of referring expressions over an extended period of time. We present results of how colour terms are used to describe landmarks in a task oriented dialogue (a route communication task) and describe a computational cognitive model of the observed adaptations over time.

1 Introduction

Much attention in recent computational as well as psychological research on language has been given to the linguistic problem of the use and generation of referring expressions. Referring expressions are linguistic expressions that identify either a referent entity in the real world or a discourse entity in the form of an antecedent. Referring expressions serve the purpose of distinguishing the target or referent from the set of other possible referents in the given context, called the distractor set. For example, in the set of objects in Figure 1, *the black cup* and *the small, black cup* would both succeed in distinguishing the cup at the lower left (the referent) from the other two objects (the distractor set).

A speaker wanting to pick out that small, black, cup at the lower left of the array could use

any of the attributes in the expressions just given. Computational approaches to generating referring expressions often produce expressions that, if possible, uniquely and minimally select the target object. But such algorithms are computationally costly and may not be helpful in modelling human behaviour: People (1) produce non-minimal expressions, which contain redundant information (e.g., Pechmann 1989) and (2) interpret such expressions more easily (e.g., Paraboni, van Deemter and Masthoff 2007).



Figure 1: A simple domain of reference: for each object, the other are distractors

A prominent account of how human-like, non-minimal referring expression can be generated is the algorithm by Dale and Reiter (1995), which by now has many extensions (see van der Sluis (2005) for a recent overview). This algorithm incrementally tests whether using an attribute in a referring expression will rule out distractor objects. The attributes are tested according to a preference list that is fixed beforehand. For the domain used in Figure 1, for example, this preference list could be $\langle \text{type, colour, size} \rangle$. Identifying the object to the right would then produce the non-minimal expression *large, white cup* by first adding the type attribute (which has a special status and is always added), then by adding *white* (because it removes the object in the lower left from the distractor set), finally by adding *large* (because it removes the object in the top

left). Non-minimal expressions arise simply because a selected attribute is never de-selected, even if a subsequently selected attribute makes it redundant.

While these approaches deal with which of the available possibilities to describe the target object is chosen, they do not account for the adaptations that a speaker makes over time to the demands of the current task environment. The computational as well as the psycholinguistic paradigms typically lack history: On each trial a participant (or algorithm) is presented with a picture like Figure 1 and instructed to produce a suitably distinguishing expression. The trial terminates without feedback and is followed by others, presenting different objects and distinguishing features. How the fourth target is distinguished from its distractors might actually owe something to the participant's experience with the first three, and our work attempts to discover and model such effects of experience.

We examine referring expressions in an unrestricted, task-oriented dialogue in which the interlocutors get natural feedback on failures of reference and refer to many different objects. We use a variant of the HCRC Map Task (Anderson et al. 1991) in which a player who can see the route on a schematic map describes it to a fellow player who must reproduce it. Each map is populated with cartoon landmarks, distinguished by several different features. We have shown that the use of features changes across first mentions as players pursue their task (Guhe and Bard 2008). In the present paper we ask how and why the changes take place. Colour is a perceptually salient property, usually one of the first tested in the incremental Dale and Reiter type algorithms. In our experiment, however, we set unreliability against salience: Colour is an unreliable distinguisher. In contrast, each map allows for use of a reliable attribute, too, (shape, number, kind or pattern). Thus, our participants need to use the adaptive attributes but waste time and can cause misunderstandings using the unreliable one.

In this paper, we report how the use of colour terms changes over the course of the experiment and present a simple computational cognitive model of this change. More precisely, we describe how the utility of the colour feature influences the Instruction Giver's choice of whether to use colour in introductory referring expressions. The model offers an explanation of this change in terms of Anderson's rational analysis (Anderson 1990; Anderson and Schooler 1991). Rational analysis is the core mechanism in ACT-

R's utility-based production selection (Anderson 2007) and is a variant of utility learning mechanisms found in reinforcement learning or the delta rule (Sutton and Barto 1998). In brief, rational analysis says that human memory reflects the frequency of events in the environment, making more frequent experiences easier to retrieve and corresponding behaviours more likely to be used. By using rational analysis our model goes beyond existing accounts of use and generation of referring expressions in that it reveals the environmental influences on these processes.

2 Comparison to existing research

The problem of whether the use of features changes with the demands of the task environment has scarcely been addressed in the literature. Although Brennan and Clark's (1996) conceptual pacts address changes in referring expressions, these changes are about how speakers refer to objects after they have been introduced. However, our questions here address the overall use of features in referring expressions over the course of many interactions. To exclude effects of conceptual pacts we are only analysing the use of introductory (first) mentions of landmarks.

Garrod and Doherty (1994) describe how a community of speakers establishes a sub-language in referring to entities. We are concerned with the internal structure of the referring expressions themselves and propose a utility-based explanation instead of one based on precedence and salience.

There is some evidence that extra-linguistic factors play a role in generating referring expressions. For example, Arnold and Griffin (2007) show that the presence of a second character influences the choice of whether to use a pronoun or the character's name for references following the introductory mention. This is true even if the characters differ in gender, so that the name does not disambiguate any more than the pronoun. Arnold and Griffin argue that the reasons for this behaviour lie in the speakers' cognitive load when they generate the referring expression.

This is part of another strand of findings in which the cooperative view on dialogue (e.g. Clark 1996) is changed towards a speaker-oriented view (e.g. Bard et al. 2000). In this view, the speaker makes the general assumption that what he/she knows is shared knowledge. Only if problems arise in the dialogue, e.g. by explicit feedback from the listener, might the speaker adapt to the listener's needs. In fact,

even if overspecified referring expressions (Dale and Reiter 1995; Paraboni, van Deemter and Masthoff 2007; Pechmann 1989) help the listener to identify the target object, the speaker also profits in terms of a generation process of greatly reduced complexity. Since both – speaker and listener – benefit from using such referring expressions, the communicative strategy cannot be attributed uniquely to concerns for the listener’s needs. In our task, however, the colour feature is counterproductive in the majority of cases, because it does not match between the two maps. So the speaker’s assumption about the usefulness of the salient feature colour are mistaken.

Another related line of research is the use of machine learning techniques to extract the way attributes are selected for modified versions of the Dale and Reiter algorithm (Jordan and Walker 2005). Although these algorithms already incorporate psychological findings, e.g., conceptual pacts, they only provide global adaptations to properties of linguistic corpora and do not account for changes over time and for adaptations to the properties of the task environment.

3 Experiment

3.1 Task

The experiment is a modified Map Task (Anderson et al. 1991). The Map Task is an unscripted route-communication task in which an Instruction Giver and an Instruction Follower each have a map of the same fictional location. The Giver’s map contains a route that is missing on the Follower’s map. The dyad’s goal is to recreate the Giver’s route on the Follower’s map.

The dialogue partners use the landmarks on the maps to navigate from START (shared) to FINISH (only on the Instruction Giver’s map).

3.2 Materials, procedure, data collection

Materials. Some landmarks differ between the two maps. In our experiment they can differ by:

1. Being absent on one of the maps or present on both;
2. Mismatching in a feature between the two maps (most notably colour);
3. Being affected by ‘ink damage’ that obscures the colour of some landmarks on the Instruction Follower’s map.

There are four attributes which also distinguish landmarks. Each serves for two different kinds of landmarks:

1. *Number* (bugs, trees),
2. *Pattern* (fish, cars),
3. *Kind* (birds, houses/buildings),
4. *Shape* (aliens, traffic signs).

Three crossed independent variables determine the nature of Giver–Follower map pairs:

1. *Homogeneity*: whether the landmarks on a map are of just one kind (single) or of different kinds (mixed).
2. *Orderliness*: whether the ink blot on the Instruction Follower’s map obscures a contiguous stretch of the route (orderly) or a non-contiguous stretch (disorderly). The number of obscured landmarks is constant.
3. *Animacy*: whether the landmarks on a map are animate or inanimate (thus, on the mixed maps there are only landmarks from the 4 inanimate or the 4 animate kinds of landmarks).

The maps in Figure 2 are a pair of Giver and Follower maps for the disorderly, mixed tree condition. Thus, the maps contain mainly trees but also other inanimate objects (mixed), and the Follower’s map shows multiple, non-contiguous ink blots (disorderly).

Procedure. Participants are told that the maps are ‘of the same location but drawn by different explorers’. They thus know that the maps can differ but not where or how. They are instructed to recreate the route on the Follower’s map as accurately as possible.

Each dyad did 2 simple training maps and then completed a set of 8 maps, one for each kind of landmark. The maps were counterbalanced with respect to the experimental conditions. After the fourth map, the role of Instruction Giver and Instruction Follower were exchanged.

To reduce the variability of words and concepts used in the unrestricted dialogues, each participant was prompted textually to provide standard type names for a few landmarks that would occur on the following map.

Setup and data collection. Participants sat in front of individual computers, facing each other, but separated by a visual barrier.

This research is part of a larger multimodal project. The communication was recorded using 5 camcorders. The Giver was eye tracked using a remote eye tracker. Speech was recorded using a

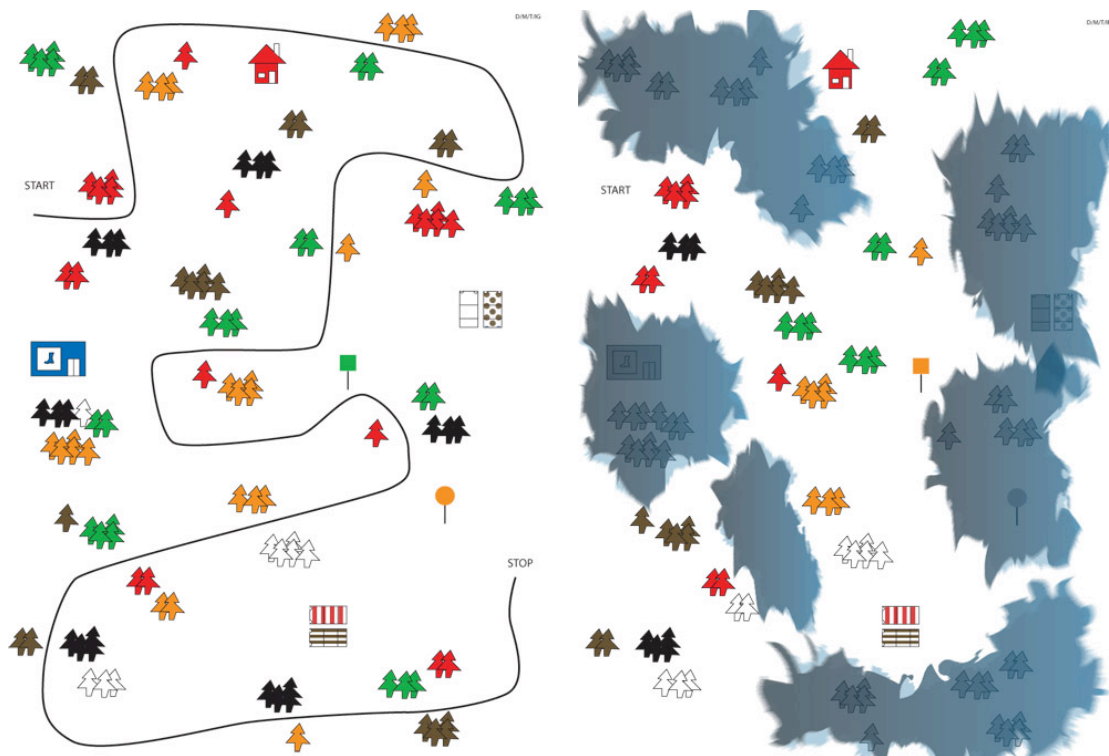


Figure 2: A pair of example maps; Instruction Giver left, Instruction Follower right

Marantz PMD670 recorder whereby Giver and Follower were recorded on two separate channels using two AKG C420 headset microphones. The speech was transcribed manually. The routes drawn by the Follower were recorded by the computer.

As the participants were in the same room, they could hear each other's speech. They could also see each other in the left half of their monitor, which showed the dialogue partner's upper torso video stream. The right half of the monitor showed the map.

Participants. In exchange for course credit, 64 undergraduates of the University of Memphis participated in pairs. In 4 dyads the participants knew each other previously.

3.3 Analysis and results

The recorded dialogues were coded for referring expressions. We present results for the first mentions of landmarks by the Instruction Giver. Introductory mentions should be both maximally independent of one another (as repeated mentions reflect precedence in naming a given object) and maximally detailed (as reductions in form characterise anaphora). Mentions of colour in landmark introductions were calculated as a proportion of opportunities

1. Over the course of single dialogues (by quartiles),
2. Across successive maps (1–8) and
3. Between those where the Instruction Giver lacked or already had experience as Instruction Follower.

The changes in the ratio of colour term use is depicted in Figure 3.

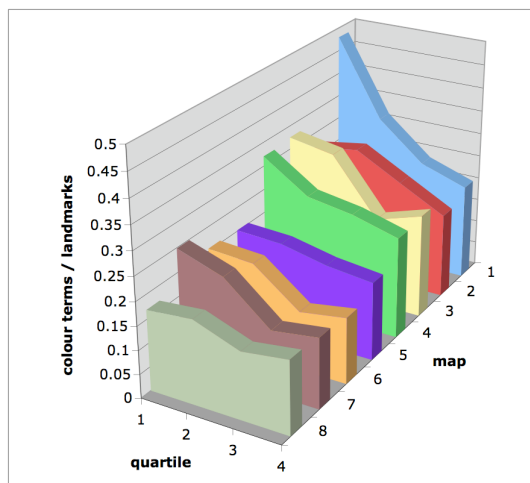


Figure 3: Change of the use of colour terms over quartiles of the eight maps

The use of colour terms significantly decreased over an average dialogue (effect of quartile within experience (2) x map encountered as Instruction Giver (4) x quartile (4) ANOVA on the arcsine transformed proportion of colour terms: $F_1(2, 54.8) = 15.57, p < 0.001$). Although there was no significant reduction across dialogues with the same Instruction Giver, the Givers used significantly fewer colour terms when they had served earlier as Follower (0.267 colour terms on average in the first four maps vs. 0.175 in the second four). This is a significant effect of experience ($F_1(1, 28) = 7.90, p < 0.01$).

Note that the orderliness of the ink blots on the Instruction Follower's maps did not have a significant effect. In contrast to colour, distinguishing features (number, kind, shape, pattern) are significantly more common in the maps where they are critical (used in more than 80%) and significantly increase within a dialogue. Thus, the decrease and low overall use of colour terms is not due to a general decrease in use of feature terms. There is also no effect of prior experience as Giver for useful features. The detailed results are presented in Guhe and Bard (2008).

3.4 Discussion

The participants adapted their use of colour to its low utility in the given task environment. The adaptation was distributed between speaker and listener. The use of colour terms does not fall significantly over the 4 dialogues a participant has the role of Instruction Giver, but there is a significant drop when the participants exchange roles: experience trying to match colour terms to grey-scale objects as Instruction Follower discourages to mention colour as Instruction Giver. Any listener-centric effect is outweighed or fuelled by a speaker-centric appreciation of utility.

4 Utility and task environment

4.1 Utility and selection probability

This is not the place to delve into the depths of the ACT-R theory, see Anderson (2007) for the most recent account. For the model described below it is only relevant that in ACT-R procedural knowledge (such as to decide whether to use colour or not) is encoded as production rules, or productions for short. A production is basically an if-then rule: *if* a certain set of conditions are given *then* execute a specified action.

In ACT-R, each production has a utility value. The utility is an estimate of how likely the use of

the production results in achieving the current goal (here: successfully describing the landmark to the interlocutor).

Productions' utilities are important in the cases in which more than one production is applicable for a given set of conditions. Then, the utilities serve to compute the probabilities with which a production is selected. This selection probability is computed as:

$$P_i = \frac{e^{U_i/s}}{\sum_j e^{U_j/s}}$$

with:

P_i : selection probability for production i

U_i : utility of production i

s : noise in the utilities (defaults: $s = 1$)

j : set of all applicable productions (including i)

Utility values are learnt over time. After a production has been used, its utility is updated depending on whether it was successful according to the following equation:

$$U_i(n) = U_i(n-1) + \alpha[R_i(n) - U_i(n-1)]$$

with:

U_i : utility of production i

n : number of applications of the production

α : learning rate

R : reward

If the production is applied successfully, the utility is updated with a positive reward, if it is unsuccessful, it receives a negative reward.

Anderson (2007, p 161) points out that this is basically the Rescorla-Wagner learning rule (Rescorla and Wagner 1972) or the delta rule by Widrow and Hoff (1960). So there is nothing special 'ACT-R-ish' about this rule; it is a general learning rule.

4.2 Structure of the task environment

In the maps about half of the landmarks on the Instruction Follower's map are obscured by ink blots, and, therefore, don't have colour. Additionally, some of the route critical landmarks mismatch in colour. Overall this means that using colour to describe a landmark is successful in only about 40% of cases. By comparison, using the distinguishing feature of a map is successful in about 92% of cases.

5 Model

5.1 Introduction

The following analyses compare the model's performance to the introduction of the first 33 landmarks of each map by the Instruction Giver. The 33rd landmark is still mentioned in 206 of the possible 256 cases (32 dyads with 8 maps each). The 34th landmark is introduced only 186 times.

There are three main patterns in the data. Firstly, map 1 behaves differently than the other maps in that the number of colour terms shows a pronounced drop from 0.6 to 0.25 (taken from the means of the first and last three values). Secondly, maps 2 to 4 each show a decrease of colour rate from 0.3 to 0.2. Thirdly, in maps 5 to 8 – after the role change – the colour rate drops in each map from 0.2 to 0.15. (This lower colour rate is the basis for the effect of role change.)

Thus, between maps the colour rate is going up again. Explanations may be that the longer-term utility of colour (learnt over a lifetime) or the textual prompting between dialogues exert some influence. The fact that the colour rate in maps 5 to 8 starts at the same rate as it ends in maps 2 to 4 may be due to the utility learning during the time as Instruction Follower. But a more detailed model is needed to explain this.

5.2 The model

The model is not a fully implemented ACT-R model, but just uses the two equations for updating production utility and probability of production selection introduced above. The model contains two competing 'productions' one for using colour, one for not using colour. Because the Instruction Giver always has colour available to describe a landmark, the model assumes that both productions are applicable for each landmark. Thus, the model is similar to the ACT-R model for an experiment by Friedman et al. (1964), described by Anderson (2007, p. 165–169; in this experiment participants have to predict which one of two lights will be lit.) Using the other features would be modelled as analogous sets of productions.

For each decision, the model selects one of the productions according to their utilities and corresponding selection probabilities at that time. After the decision has been made, the usefulness of colour is determined according to the structure of the task environment (thus, using colour is successful in 40% of cases) and the utility of the selected production is updated accordingly. For a

successful application the production receives a reward of $R = 14$; if it is unsuccessful it receives a reward of $R = 0$ (cf. Anderson 2007, p. 162).

The results reported in the remainder of this section were obtained by 500 runs of the model. However, just 32 runs – matching the number of dyads in the experiment – suffice to get significant results; more runs of the model just produce a smoother curve.

5.3 Map 1

For the first map the model starts with the following estimated utilities:

$$U_{\text{colour}}(1) = 5.5$$

$$U_{\text{no-colour}}(1) = 5$$

These values mean that the colour-production has a probability of being selected of 0.622, which is close enough to the mean of the first three values of 0.594. (Using $U_{\text{colour}}(1) = 5.4$ would give an initial probability of 0.599, but one can be too fussy.)

The final average utilities are:

$$U_{\text{colour}}(33) = 4.6$$

$$U_{\text{no-colour}}(33) = 7.7$$

Choosing these initial utilities gives an excellent fit to the data, see Figure 4. A regression using the model as predictor for the data shows a significant correlation ($\beta_1 = 0.90$, $p < 0.001$) that accounts for 72% of the variance ($R^2 = 72\%$, $F(1, 31) = 79.5$, $p < 0.001$).

However, the initial values are not that important, and the model matches the data significantly for a wide range of start values, as long as $U_{\text{colour}}(1) > U_{\text{no-colour}}(1)$ and the values are not close to the extremes of 0 and 14. The same holds for the following simulations.

5.4 Maps 2 to 4

For maps 2 to 4 (see Figure 5) the initial utilities were set to:

$$U_{\text{colour}}(1) = 5.5$$

$$U_{\text{no-colour}}(1) = 6.5$$

resulting in final average utilities of:

$$U_{\text{colour}}(33) = 4.5$$

$$U_{\text{no-colour}}(33) = 7.5$$

The regression shows that the model accounts for 66% of the variance ($R^2 = 66.3\%$, $F(1, 31) = 61.0$, $p < 0.001$) with $\beta_1 = 2.44$ ($p < 0.001$).

5.5 Maps 5 to 8

Finally, for maps 5 to 8 (see Figure 6) the initial utilities were set to:

$$U_{\text{colour}}(1) = 3$$

$$U_{\text{no-colour}}(1) = 4$$

resulting in the final average utilities

$$U_{\text{colour}}(33) = 3.3$$

$$U_{\text{no-colour}}(33) = 7.7$$

The model accounts for 52.7% of the variance ($R^2 = 52.7\%$, $F(1, 31) = 34.6$, $p < 0.001$) with $\beta_1 = 0.84$ ($p < 0.001$).

6 Conclusions

There are two main conclusions from the research presented here. Firstly, the dialogue partners indeed adapt their naming behaviour to the task environment. More specifically, they adapt to the fact that colour is an unreliable distinguisher for the landmarks on the maps. (This is amplified by the fact that the participants do not make a substantial effort to identify the parts of the maps that are obscured by ink, which shows in the absence of an orderliness effect.)

Secondly, the simple computational cognitive model accounts for this change. In particular, the model shows that the change in behaviour is indeed an adaptation to the structure of the task

environment, because the rate of the probabilities and the changes in the probabilities with which

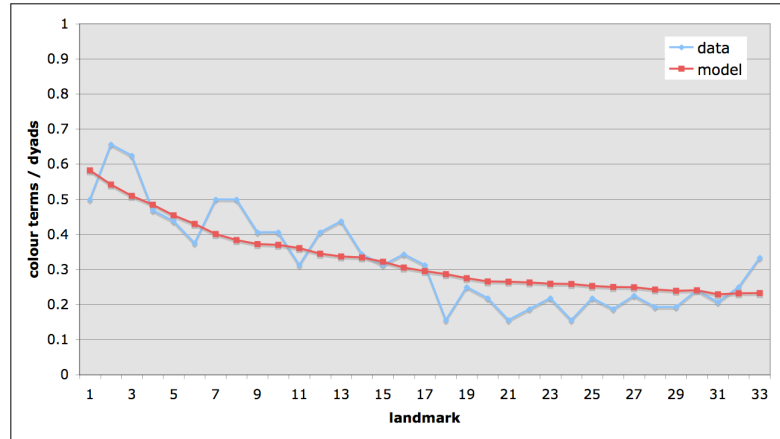


Figure 4: Comparison of data and model for the first 33 landmarks in map 1.

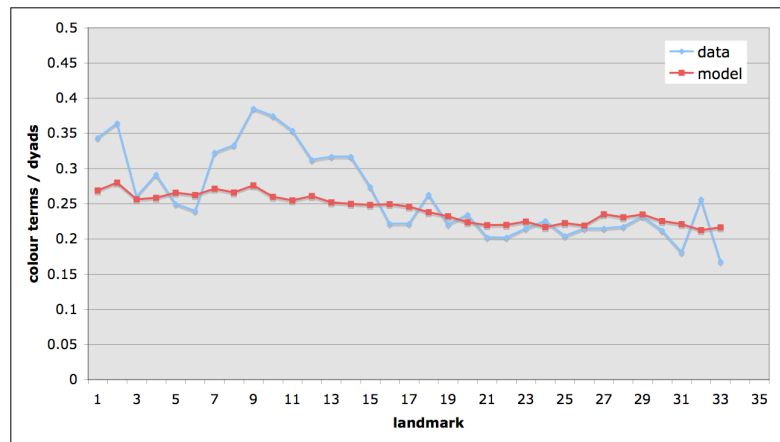


Figure 5: Data and model for maps 2 to 4.

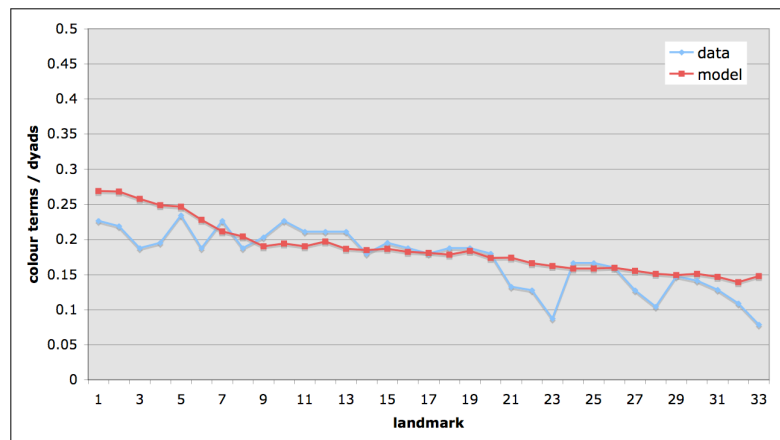


Figure 6: Data and model for maps 5 to 8.

colour is used as a descriptor is a direct result of the fact that colour can be successfully used for about 40% of the landmarks on the maps. Thus, rational analysis (the fact that memory reflects the probabilities encountered in the environment) explains the observed phenomenon.

Although – after the fact – it may not be too surprising that rational analysis explains the observed phenomenon, the result is more far-reaching, because the influences of the task environment on naming behaviour (the generation of referring expressions) has not yet been reported.

7 Future work

Our future research will address a number of direct follow-up issues. Firstly, the model will be extended to account for the changes in the mentions of the distinguishing features (number, pattern, kind, shape). Secondly, after a more detailed analysis of the data we will extend the model to account for individual adaptation patterns in the sense that the model can account for groups of dyads showing similar dialogue histories. For this, we will model the landmark introductions made by the Instruction Follower as well. This model serves as starting point for a comprehensive ACT-R model of how referring expressions (including repeated mentioned of landmarks) are generated in the given task.

Acknowledgements

This research was supported by grant NSF-IIS-0416128 to Max Louwerse, Art Graesser, Mark Steedman, and Ellen Gurman Bard. Thanks to Antje van Oosten and Jonathan Kilgour for help with the data coding and extraction and to Max Louwerse, Nick Benesh, Gwyneth Lewis and Megan Zirnstein for providing the transcriptions.

References

- Anderson, A. H., Bader, M., Bard, E. G., Boyle, E., Doherty, G., Garrod, S., Isard, S., Kowtko, J., McAllister, J., Miller, J., Sotillo, C., Thompson, H. S., and Weinert, R. (1991). The HCRC Map Task corpus. *Language and Speech*, 34(4), 351–366.
- Anderson, J. R. (1990). *The Adaptive Character of Thought*. Hillsdale, NJ: Erlbaum.
- Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* New York: Oxford University Press.
- Anderson, J. R. and Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science*, 2(6), 396–408.
- Arnold, J. E. and Griffin, Z. M. (2007). The effect of additional characters on choice of referring expression: Everyone counts. *Journal of Memory and Language*, 56(4), 521–536.
- Bard, E. G., Anderson, A. H., Sotillo, C., Aylett, M. P., Doherty-Sneddon, G., and Newlands, A. (2000). Controlling the intelligibility of referring expressions in dialogue. *Journal of Memory and Language*, 42(1), 1–22.
- Brennan, S. E. and Clark, H. H. (1996). Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(6), 1482–1493.
- Clark, H. H. (1996). *Using Language*. Cambridge, MA: Cambridge University Press.
- Dale, R. and Reiter, E. (1995). Computational interpretations of the Gricean maxims in the generation of referring expressions. *Cognitive Science*, 19(2), 233–263.
- Friedman, M. P., Burke, C. J., Cole, M., Keller, L., Millward, R. B., and Estes, W. K. (1964). Two-choice behavior under extended training with shifting probabilities of reinforcement. In R. C. Atkinson (ed.), *Human Brain Function* (2nd ed.) San Diego, CA: Academic Press.
- Garrod, S. and Doherty, G. (1994). Conversation, coordination and convention: An empirical investigation of how groups establish linguistic conventions. *Cognition*, 53, 181–215.
- Guhe, M. and Bard, E. G. (2008). Adapting referring expressions to the task environment. In: *Proceedings of CogSci 2008*.
- Jordan, P. W. and Walker, M. A. (2005). Learning content selection rules for generating object descriptions in dialogue. *Journal of Artificial Intelligence Research*, 24:157–194.
- Kronmüller, E. and Barr, D. J. (2007). Perspective-free pragmatics: Broken precedents and the recovery-from-preemption hypothesis. *Journal of Memory and Language*, 56(3), 436–455.
- Paraboni, I., van Deemter, K., and Masthoff, J. (2007). Generating referring expressions: Making referents easy to identify. *Computational Linguistics*, 33(2), 229–254.
- Pechmann, T. (1989). Incremental speech production and referential overspecification. *Linguistics* 27, 98–110.
- Sutton, R. S. and Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. Adaptive Computation and Machine Learning. Cambridge, MA: MIT Press.
- van der Sluis, I. (2005). *Multimodal Reference: Studies in Automatic Generation of Multimodal Referring Expressions*. PhD thesis, Tilburg University, The Netherlands.