Investigating the Effect of Random Noise on the Evolution of Colour Terms

Mike Dowman
Theoretical and Applied Linguistics
School of Philosophy, Psychology &
Language Sciences
The University of Edinburgh
mdowman@ling.ed.ac.uk

Abstract- The effect of adding noise to an expression-induction model of language evolution was investigated. The model consisted of a number of artificial people who were able to infer the denotation of basic colour terms from examples of colours which the words had been used to identify, using a Bayesian inference procedure. The artificial people would express colours to one-another, so producing data from which other people could learn. Occasionally they would be creative, which allowed new words to enter the language. When certain points in the colour space were made especially salient, so that the artificial people were more likely to remember colours at these points, the languages emerging over a number of generations in evolutionary simulations replicated the typological patterns seen in the 110 languages of the world colour survey. It was found that if random noise was added to the data from which the artificial people learned, this had no major effect on the emergent languages, demonstrating that the Bayesian inference procedure is able to learn effectively despite the presence of random noise, even when placed in an evolutionary context.

1 Introduction

This paper describes computational modelling experiments performed to explain empirical data concerning basic colour terms. Colour terms are simply words in natural languages which are used to denote the property of colour. In most, if not all, languages, a special subset of such words can be identified, which Berlin and Kay (1969) named *basic colour terms*. These are colour words which are known by all speakers of a language, which are highly salient for all speakers, and which don't just name a subset of the colours denoted by another colour word. The application of these criteria to English, results in the set of 11 basic colour words, *red, yellow, green, blue, orange, purple, pink, brown, grey, black* and *white*.

Berlin and Kay (1969) found that different languages had different numbers of basic colour words, varying from as few as two, up to the eleven terms of English. A few languages such as Russian or Hungarian might have a twelfth basic colour term (MacLaury, 1997a), although this is somewhat controversial, because not all researchers accept that the twelfth terms should be considered basic. Even in languages which have only a very small number of colour words, it is extremely unusual for languages to leave any area of the colour space without a corresponding colour word (Kay and Maffi, 1999).

Heider and Olivier (1972) found that if a language has only two colour words, these will divide the colour space up so that one refers to light colours, including white, red and yellow, while the other refers to dark colours. Languages with only two colour words are extremely rare, but research has shown that all languages with only two colour words either divide the colour space up in this way, or else simply into lighter and darker colours (MacLaury, 1997a).

Berlin and Kay (1969) found that in languages with more colour terms, there were also predictable patterns in the way in which the colour words divide up the colour space. For example, the third colour term to be added will always correspond to red. It will have it's focus in the same place as red terms in all other languages, although the boundaries of the range of colours which such words denote may vary greatly.

Berlin and Kay (1969) proposed that such typological patterns are due to a cultural evolutionary process, in which languages gradually add more colour terms over time. They proposed that the order in which colour terms were added was the same across languages, although they did note that a minority of languages appeared not to conform to their classification. This theory has undergone a number of revisions over time, essentially allowing for more variability between languages in the order in which they add colour terms, and allowing for some alternative, less common evolutionary sequences. Kay and Maffi (1999) propose that most languages (83% of those in the 110 language world colour survey (Kay, Berlin, Maffi and Merrifield, 1997)) lie on the evolutionary trajectory shown in Figure 1, although their classification does not attempt to predict the order of appearance of the terms *orange, pink, brown, purple* or *grey* so precisely.

white-red-yellow + black-green-blue

↓

white + red-yellow + black-green-blue

↓

white + red-yellow + black + green-blue

↓

white + red + yellow + black + green-blue

↓

white + red + yellow + black + green + blue

Figure 1 The Main Line of the Evolutionary Trajectory

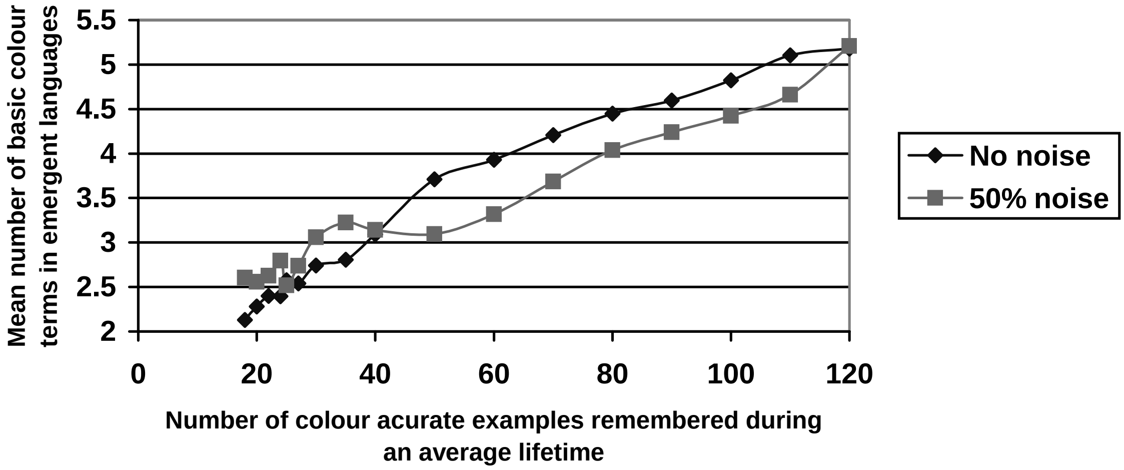

Figure 2 The Number of Basic Colour Terms in Emergent Languages

Kay and Maffi (1999) also acknowledge the presence of some alternative trajectories, in one of which red splits from yellow before black splits from composite green-blue. Following this stage, either black will split from black-green-blue to return to the main evolutionary trajectory, or green will split from black-green-blue to leave a black-blue term.

We also see systems with yellow-green-blue terms, the origin of which is somewhat mysterious, because it is not clear how such systems could develop from any attested type of colour term system. Kay and Maffi (1999) suggest that these systems may develop in languages which previously did not have colour words for some areas of the colour space, while MacLaury (1997a) suggests that these terms can be explained as developing in languages which previously divided the colour space simply on the dimension of lightness, ignoring distinctions of hue. Each of these hypotheses seem plausible, though I do not think that there is strong evidence to support either of them. Systems with yellow-green-blue terms can then return to the main line of the evolutionary trajectory if yellow splits from yellow-green blue, or blue can split, resulting in yellow-green terms.

Kay and Maffi (1999) also note that a very small number of languages appear to have colour term systems which do not correspond to any of these types. However, the colour term systems of the vast majority of languages can be classified as being of one of these types. The only types of colour term that meet Berlin and Kay's (1969) criteria for basic status, are those contained in the evolutionary trajectories together with *orange*, *purple*, *brown*, *pink*, *grey* and possibly another type of desaturated lavender colour (MacLaury, 1997a). We might expect to see basic colour terms between blue and green (*turquoise*) or between yellow and green (*lime* or *chartreuse*), but no such term has been reported in the literature.

While Kay and Maffi included neither *purple* nor *orange* terms in their evolutionary trajectories, MacLaury notes that *purple* tends to emerge before *orange*, and is hence more common cross-linguistically, but otherwise the order of emergence of the colour terms other than *red*, *yellow*, *green*, *blue*, *black* and *white* seems less predictable.

Kay and McDaniel (1978) attempted to link these typological patterns to the neurophysiology of the human colour vision system. We have four kinds of cell in the retina of our eyes that have a maximum firing rate in the presence of particular hues of either red, yellow, green or blue, and such cells were also hypothesized to exist for black and white (de Valois and Jacobs, 1968). In addition, Berlin and Kay (1969) showed that colour terms across languages tended to be focussed on one of these colours, and MacLaury (1997b) provides further corroboration of this finding using data from the world colour survey. Hence Kay and McDaniel proposed that these universal patterns might be the direct result of the physiology of the human visual system.

There is also psychological evidence suggesting that these colours are especially salient, and that they are better remembered than other colours (Heider, 1971, 1972; Rosch, 1973), although more recent work has cast doubt on some of the earlier evidence (Lucy, 1992; Roberson, Davies and Davidoff, 2000). Some researchers have suggested that an adequate explanation of colour term typology is only possible if cultural factors are considered, (Foley, 1997; Saunders, 1992) and that colour term universals may be more influenced by cultural practices than by neurophysiology.

However, the universal prototypes apparent in basic colour term systems appear to relate to the colours of light that produce maximal firing rates in cells in the eye, so most attempts to explain colour term typology have attempted to do so primarily in terms of the neurophysiology of colour vision. While Kay and McDaniel (1978) suggested a connection between neurophysiology and colour words, they did not provide an explanation as to why we see composite colours such as red-yellow, yellow-green, and black-green-blue, but not other possible composites such as blue-red or white-green-yellow. The computer model described in this paper aims to give a fuller explanation of colour term typology, suggesting that it is the product of evolutionary processes under the influence of physiological biases.

2 Expression-Induction Modelling

The methodology used in this paper is a version of expression induction modelling[1] (Hurford, 2002). This kind of methodology is a development of the computational evolutionary linguistic modelling, originating with Hurford (1987), who used it to account for aspects of number word typology. The methodology involves creating a number of artificial people (often referred to as *agents*), who are able to infer a simple language based on example utterances, and then to express themselves using this language.

The expressions of one or a number of agents will form the input from which the next generation of learners will induce their grammars (although agents within the same generation may learn from each other, so there need not be any clear cut generational gap). Generally, after a period of time, all speakers will come to share a common language, although the internalised languages of each individual person may in some cases be slightly different. Because of the bottleneck through which language passes, it is also possible that transmission between generations may be imperfect, resulting in language drift, which may be viewed as a kind of cultural evolutionary process.

The expression-induction modelling methodology has now been applied to a wide range of problems, including explaining syntactic compositionality (Kirby, 1999), and vowel system typology (de Boers, 1999). Most relevant to the work reported here, however, is the model of Belpaeme (2002), who also constructed an expression induction model of colour term evolution. Most of Belpaeme's simulations contained ten artificial people, each of which was able to represent colour categories using adaptive networks, a kind of neural network. Colour in the model was represented in terms of the CIE-LAB space, which represents colour in terms of three dimensions, one of which corresponds to its degree of redness or greenness, one to the degree of yellowness or blueness, and the third to the lightness or darkness of the colour[2]. The networks acted as fuzzy membership functions, allowing colour categories corresponding to a volume of the three dimensional CIE LAB space of almost any size or shape to be represented. Each artificial person could also remember a number of word forms, each of which could be paired with a colour category.

In the initial state of the simulation, the artificial people did not know any colour categories or colour words, so, the first time one of them spoke, they would have to create a new category and corresponding word. In general, communication proceeded by first choosing one colour to be a topic, and another to be a context, and then choosing one person to be a speaker, and another to be a hearer. The speaker would then try to communicate to the hearer which of the colours was the topic, and which was the context, by choosing a word which included the topic, but not the context, in its denotation. If the word that the speaker used was known by the hearer, and the colour category which the hearer had associated with that word included only one of the colours, then the hearer would understand that that colour was the topic. If this was correct, then communication would have been successful, and the association between the topic colour and the colour word would be strengthened. If communication was not successful, then the hearer would be shown the correct topic, and the word's colour category would be adapted, so that it would be a better match for the topic colour. Categories and words which were persistently not useful in communication would eventually be forgotten.

In some simulation runs, the same artificial people would exist for the whole of the simulation, though in others evolution over a number of generations would be simulated, by periodically replacing one of the people with a new one who had not learned any colour words. However, similar results were obtained in both these conditions. The most important result was that, over a period of time, coherent colour lexicons emerged which were shared by all the artificial people. The colour lexicons would divide the colour space into a number of colour regions, each of which would be associated with a particular colour word. The people never agreed completely about the exact meaning of each colour word, but their languages were consistent enough for them to achieve rates of communicative success in excess of 85%. However, the colour categories emerging in Belpaeme's model did not resemble the colour terms of real languages, as they did not conform to the typological restrictions observed in colour term systems cross-linguistically[3].

Belpaeme and Bleys (2005) present new results, obtained using a modified version of Belpaeme's original model. Colour categories were no longer represented using adaptive networks. Instead, they were represented by fixing their centre at a point in the colour space, and using the Euclidean distance to that point as a membership function for the colour category. This simplifies the representation of colour categories, but would seem to place some restrictions on the shape of categories that could be represented.

Example colours were then chosen either completely at random, as before, or were selected randomly from colours within digital photographs of natural scenes. Belpaeme and Bleys were able to show that there was a similarity, in terms of where category centres occurred, between the results of their simulations and those of the world colour survey. The simulation results were most similar to those of the world colour survey when the colours from digital photographs were used, and when the people tried to communicate with one another, rather than simply trying to maximise the discriminative ability of their own colour categories. However, Belpaeme and Bleys did not demon-

[1] This kind of model is also called an *iterated learning model*.

[2] This colour space was chosen because Lammens (1994) showed that his computer model of colour naming worked best in this space.

[3] Belpaeme (2002) did suggest that the split into light and dark colours seen in languages with only two colour terms might be explainable in terms of his model, because this might be the easiest way to divide up the colour space, but, in its present form, the model was not able to account for any other aspects of colour term typology.

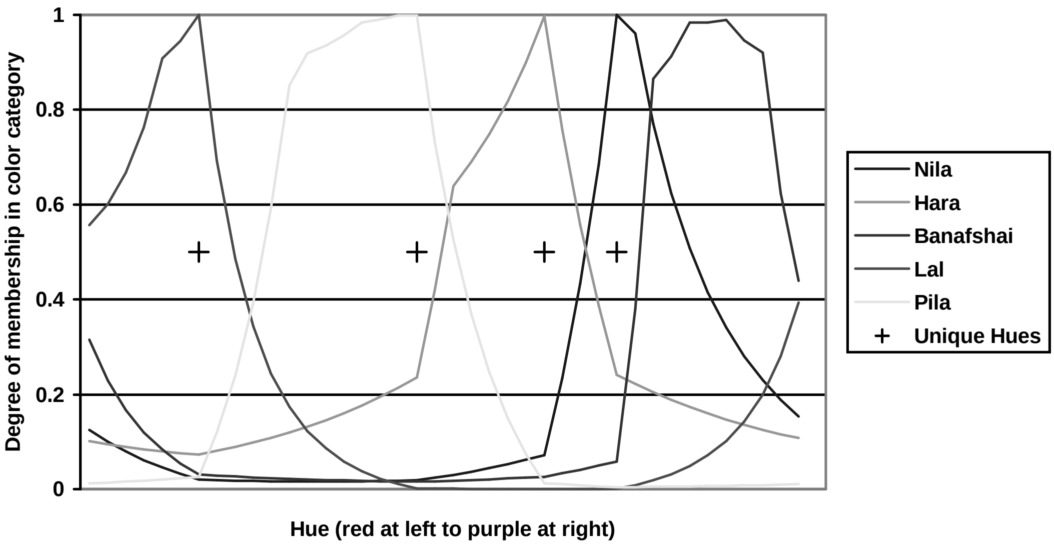

Figure 3 Denotations Learned for Urdu Colour Terms

strate a reproduction of the evolutionary trajectories reported by Kay and Maffi (1999), something which the expression-induction model presented here was able to achieve.

3 A Bayesian Evolutionary Model

The expression-induction model of colour term evolution uses a Bayesian acquisitional model based closely on that of Dowman (2001), which allows the denotation of colour words to be inferred based on examples of colours which they have been used to identify. However, Dowman (2003) modified the acquisitional model of Dowman (2001) by making the four neurophysiological foci, red, yellow, green and blue, especially salient[4]. Heider (1972) provided psycholinguistic evidence to suggest that people find these colours especially salient, and find them easier to remember. Hence the model was modified so that examples of possible denotata of colour words were more likely to be remembered by the artificial people when they corresponded to these colours.

It was also proposed that foci of green and blue were the closest together of any neighbouring neurophysiological foci in the conceptual colour space which people use to classify colours, whilst blue and red were hypothesized to be furthest apart. The red and yellow and yellow and green foci were placed at intermediate distances, with the yellow and green foci somewhat closer together than the red and yellow foci. This colour space is shown in Figure 4, together with the locations of each of the unique hues. (The units are arbitrary, but the total size of the colour space is 40 units.) Independent verification for these parameter settings is somewhat weak, although MacLaury (1997a) does suggest that there is some evidence to suggest that the green and blue foci are more similar than the other foci. Hence the primary justification for these parameters is provided by the results of simulations using the model.

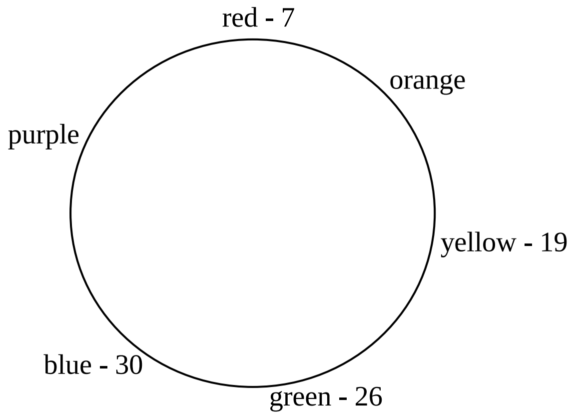

Figure 4. The Conceptual Colour Space

Figure 3 provides an example of the kind of denotations that can be learned by the acquisitional part of the expression-induction model. Example colours for each chromatic Urdu colour term were created, each randomly selected from the range of hues within the term's denotation. These examples were then shown to the model, until it had remembered 40 examples. The degree of membership of each hue in each colour category was then calculated and plotted in the graph. We can see that each term's denotation has prototype properties, with a single best example, and membership in the colour category gradually decreasing the more a colour is dissimilar to the prototype. Each term that contains a unique hue has its prototype at that colour, which is consistent with empirical findings.

[4] The Bayesian acquisition model simplifies the colour domain by ignoring the dimensions of lightness and saturation, so that it is only concerned with the dimension of hue. Hence it cannot learn denotations for black, white, grey, pink, or brown terms, as these terms are differentiated from others principally on the basis of lightness or saturation. However the model is able acquire the denotations of red, yellow, green, blue, purple and orange terms.

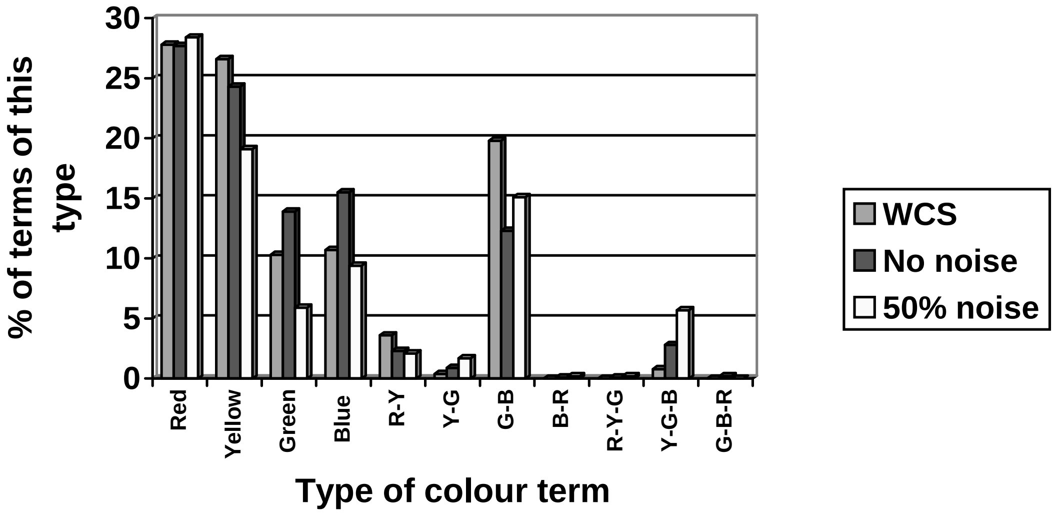

Figure 5 The Frequencies of Each Type of Colour Term

This shows that the acquisitional model explains some of the universal properties of basic colour term systems, but it cannot completely explain the restrictions on emergent language types, as it can also learn colour term systems of unattested types.

4 Evolutionary Simulations

Dowman (2003) performed simulations over several generations using groups of ten artificial people who learned using this model. They would successively name randomly chosen colours, and the chosen name and corresponding colour would be observed as data by another agent, and used by it to try to infer the colour term's full denotation, and hence which colour term to use when its turn to speak came. Occasionally agents would be creative and make up a completely new colour term, hence allowing new words to enter the language.

425 evolutionary simulations were performed, and it was found that colour term systems which evolved using this model tended to conform to the types contained in Kay and Maffi's evolutionary trajectories (Kay and Maffi, 1999), as shown in Figure 2. There were a small number of systems which did not fit these patterns, which is consistent with the empirical data, and a small proportion of colour words were of types not attested empirically as basic colour terms. It was clear that the majority of colour term systems and individual colour terms were clearly of types attested typologically.

However, an aspect of all these simulations appears to be unrealistic, as the data from which the artificial people learned was completely free from noise. The data from which the artificial people learned consisted of colour words paired with specific colours that those words could name. The justification for using this kind of data is that this is presumably the same type of data from which real people learn colour words, as children appear to learn the meanings of words largely by observing the speech of other people (Bloom, 2000). In order to learn meanings, children must infer not only the words which people use, but also what such words were used to mean. In the case of colour words this would correspond to particular colours.

However, inferring the intended referent of a word used by another person would seem to be a somewhat difficult task, and so it seems unlikely that this could be accomplished without ever pairing a colour with a word which cannot correctly denote that colour. Furthermore, there are added complications, because the speaker could also use an incorrect word, or other errors could occur, such as the learner mishearing a word. For all these reasons it would seem that not all the data from which children learn colour words is likely to be accurate.

Dowman (2001) designed the acquisitional model to be able to cope with erroneous data, and Dowman (2002) demonstrated that the model was able to learn even when as much as 80% of the data presented to it was random noise (although in such circumstances the model needs to observe a greater number of examples before inferring a word's denotation accurately). However, in the evolutionary models of Dowman (2003) no random noise or erroneous data was added to the data from which the artificial people learned.

The research described in this paper was conducted with the aim of investigating whether coherent colour term vocabularies would emerge in the presence of large quantities of random noise, and whether the colour term systems would reflect the typological patterns. The same model was used as in Dowman (2003), although 50% of the time, instead of the data from which an artificial person learned being produced by another artificial person, a random colour was paired with the colour word produced by the speaker. In other respects the model was identical to that of Dowman (2003).

Essentially the model consists of a Bayes' optimal classifier which infers colour term denotations based on examples of colours which the term has been used to denote. The colour space is represented as a one dimensional hue space, which can be indexed simply with a single number. However this conceptual space is circular, as red and purple, despite being at opposite ends of the spectrum are perceptually similar. There is a wealth of evidence to support the existence of such a conceptual colour space; see for example Thompson (1995) for a review of some of the evidence.

This acquisitional model was then included in a simple evolutionary simulation, in which there were ten artificial people, who were in turn made to choose the colour term which they thought most likely to be the best example of a colour, and pass that colour together with the chosen term to another speaker as an example. However, one time in a thousand a speaker will simply make up an entirely new colour word, as otherwise there would be no way for new colour words to enter the language, or for the number of colour words in the language to grow. Periodically one of the older speakers would die and be replaced by a new person who did not know any colour words at all.

The simulation was run 170 times, 10 times in each condition, where the number of accurate colour examples which each artificial person observed during their lifetime on average was varied between 18, 20, 22, 24, 25, 27 , 30, 35, 40, 50, 60, 70, 80, 90 , 100, 110 and 120. In each case, there would be one random example for each accurate one, constituting a level of random noise equal to 50%. Figure 2 shows the results of these simulations compared to those in which there was no random noise. The average number of basic colour terms emerging in each condition was measured and plotted on the graph. A term was considered basic if a person had seen at least 4 examples of it. Terms were included only if they were known by at least half the artificial people in the community whose age was over half the average lifespan. (This final restriction was added because younger people would be likely to know less terms, as they would not have had sufficient opportunity to learn all of the terms which were widely used in their community.)

We can see that the number of colour words emerging in the languages is, on average, roughly proportional to the average number of colour examples observed by the artificial people. This result might seem to be unsurprising, because if people use colour words more often, then they might find it useful to have more such words in their languages. However, in these simulations there are no truly functional pressures, because the artificial people receive no benefit or reward for achieving successful communication. Hence this model suggests that when we use terminology within a particular domain frequently, we might gain more words making more fine grained distinctions within that domain, regardless of whether such distinctions have any functional advantage.

Perhaps surprisingly, the number of words emerging seems to be dependent solely on the number of accurate examples of colour words which people observe during their lifetimes. Even though in the condition with 50% noise, twice as many examples were observed by each person as the people who observed the same number of accurate examples but no random noise, the number of colour terms emerging seems to be essentially the same in either condition. (The small differences between the no noise and 50% noise conditions can be attributed simply to random variation.)

It seems that in the condition with 50% percent random noise, the simulations have performed in almost exactly the same way as though that noise wasn’t present. This result seems somewhat counter intuitive, because no parameter was changed between the simulations which would have given the model any indication that there were varying amounts of random noise, and there was no indication given to any of the artificial people which would allow them to distinguish accurate from random examples.

The most important consideration, however, was whether the simulations would still mirror the typological patterns when there was so much random noise. Figure 5 compares the proportions of basic colour terms which were classified as red, yellow, green, or blue, or as composites of these terms in each condition of having no noise, or 50% noise, to the proportions of terms which were classed as each of these types in the languages on the evolutionary trajectories in Kay and Maffi (1999). (This data is derived from the world colour survey, and so is labelled WCS.)

We can see that the typological patterns in the relative frequencies of each type of colour term are roughly reproduced in each condition. The only major differences between the condition with no noise and that with noise is that there are fewer green and blue terms when there is a high level of noise, and a greater proportion of yellow-green-blue terms. These differences might be due simply to random variation, or it is possible that by altering parameters concerning the location of the neurophysiological foci the results with 50% noise might more closely reflect those with no noise. In any event, it is clear that in some ways adding noise has resulted in the simulations more closely reflecting the empirical data, especially in that the proportion of blue terms is now almost the same as that found in the world colour survey. However, in other ways adding noise has caused the simulations to diverge somewhat from the empirical data, most significantly because there are now considerably more yellow-green-blue and yellow-green terms.

Kay and Maffi (1999) did not provide data on the occurrence of purple and orange terms, but it is reported that purple terms occur more frequently than orange ones (MacLaury, 1997a). In the noiseless condition 76.9% of terms without a neurophysiological focus were purple, while 19.2% of such terms were orange, while with 50% noise these figures were 60.6% for purple and 26.8% for orange. Hence in both conditions the empirical finding that orange is less common than purple was supported by the simulations.

The corresponding figures for lime and turquoise terms were 3.8% and 0% with no noise, and 9.9% and 0.3% with 50% noise. These results are consistent with the empirical data, in that in general basic lime and turquoise terms are not found in natural languages. (Neither of these terms is generally considered basic in English.) Possibly the occurrence of basic lime terms is somewhat more frequent than should be expected, although it is not clear how often such terms are simply ignored in linguistic analyses as there is a theoretically motivated expectation that they will not be basic.

5 Conclusion

This paper has shown that the typological patterns observed in basic colour term systems cross-linguistically can be accounted for in terms of neurophysiological biases acting on an evolutionary process. Adding large quantities of random noise to the simulation, which ought to have made it more realistic, did not prevent it from accounting for the empirical data, and has not radically affected the results compared to the noiseless condition. This would seem a very desirable property for a model of language evolution, as it would be a very poor model which was unable to account for empirical data when attempts were made to reproduce conditions similar to those in which real language evolution takes place.

Acknowledgments

This work was supported by IPA and IPRS scholarships from the University of Sydney and the Australian Government respectively, and an Economic and Social Research Council post-doctoral fellowship. I would like to thank my Ph.D. supervisor Judy Kay, and everyone else who has assisted me with the research reported here.

References

Belpaeme, T. (2002) *Factors influencing the origins of colour categories*. Ph.D. Thesis, Artificial Intelligence Lab, Vrije Universiteit Brussel.

Belpaeme, T. & Bleys, J. (2005) Colourful Language and Colour Categories. Paper presented at *Second International Symposium on the Emergence and Evolution of Linguistic Communication*, Hatfield, UK.

Berlin, B. & Kay, P. (1969). *Basic Colour Terms.* Berkeley: University of California Press.

Bloom, P. (2000). *How Children Learn the Meanings of Words*. Cambridge, MA: MIT Press.

de Boer, B. (1999). Evolution and self-organisation in vowel systems. *Evolution of Communication*, 3 (1).

Dowman, M. (2001). A Bayesian Approach to Colour Term Semantics. *Lingu@scene*, 1.

Dowman, M. (2002). Modelling the Acquisition of Colour Words: A Bayesian Approach to Fuzzy Sets. In B. McKay and J. Slaney (eds.) *Advances in Artificial Intelligence*. Berlin: Springer-Verlag.

Dowman, M. (2003). Explaining Colour Term Typology as the Product of Cultural Evolution using a Multi-agent Model. In *Proceedings of the 25th Annual Conference of the Cognitive Science Society.*

Foley, W. A. (1997). *Anthropological linguistics: an introduction*. Malden, MA: Blackwell Publishers.

Heider, E. R. (1971). "Focal" Colour Areas and the Development of Colour Names. *Developmental Psychology,* 4(3):447-445.

Heider, E . R. (1972). Universals of Colour Naming and Memory. *Journal of Experimental Psychology,* 93:10-20.

Heider, E. R. & Olivier, D. C. (1972). The structure of the colour space in naming and memory for two languages. *Cognitive Psychology,* 3:337-354.

Hurford, J. R. (1987). *Language and Number The Emergence of a Cognitive System*. New York, NY: Basil Blackwell.

Hurford (2002). Expression/induction models of language evolution: dimensions and issues. In T. Briscoe (Ed.) *Linguistic Evolution through Language Acquisition: Formal and Computational Models*. Cambridge: Cambridge University Press

Kay, P., Berlin, B., Maffi, L. & Merrifield, W. (1997). Colour Naming Across Languages. In C. L. Hardin & L. Maffi (eds.) Colour Categories in Thought and Language. Cambridge: Cambridge University Press.

Kay, P. & McDaniel, K. (1978). The Linguistic Significance of the Meanings of Basic Colour Terms. *Language,* Volume 54, Number 3.

Kay, P. & Maffi, L. (1999). Colour Appearance and the Emergence and Evolution of Basic Colour Lexicons. *American Anthropologist,* Volume 101, pages 743-760.

Kirby, S. (1999) Learning, Bottlenecks and the Evolution of Recursive Syntax, in Briscoe, Edward, Eds. *Linguistic Evolution through Language Acquisition: Formal and Computational Models*. Cambridge: Cambridge University Press.

Lammens, J. M. G. (1994). *A Computational Model of Colour Perception and Colour Naming*. Doctor of Philosophy dissertation, State University of New York at Buffalo.

Lucy, J. A. (1992). *Language Diversity and Thought A Reformulation of the Linguistic Relativity Hypothesis.* Cambridge: Cambridge University Press.

MacLaury, R. E. (1997a). *Colour and Cognition in Mesoamerica: Constructing Categories as Vantages*. Austin: University of Texas Press.

MacLaury, R. E. (1997b) Ethnographic evidence of unique hues and elemental colours. *Behavioural and Brain Sciences*, 20:202-203.

Roberson, D., Davies, I. & Davidoff, J. (2000). Colour Categories are Not Universal: Replications and New Evidence from a Stone-Age Culture. *Journal of Experimental Psychology: General*, 129(3): 369-398.

Rosch, E. H. (1973). Natural Categories. *Cognitive Psychology*, 4:328-350.

Saunders, B. A. C. (1992). *The Invention of Basic Colour Terms*. Utrecht: ISOR.

Thompson, E. (1995). *Colour Vision: A Study in Cognitive Science and the Philosophy of Perception*. New York: Routledge.

de Valois, R. L. & Jacobs, G. H. (1968). Primate colour vision. *Science,* 162, 533-540.