

MODELING INFERENTIAL MINDS IN CONCEPTUAL SPACE

Carlos Barra
Independent Researcher
Las Pimpinelas 880, Concon, Chile
E-mail: c.barra@vtr.net

Enrique Canessa
Facultad Ingeniería y Ciencias, CINCO
Universidad Adolfo Ibáñez
Av. P. Hurtado 750, Viña del Mar, Chile
E-mail: ecanessa@uai.cl

Sergio E. Chaigneau
Escuela Psicología, CINCO
Universidad Adolfo Ibáñez
Av. D. Las Torres 2640, Santiago, Chile
E-mail: sergio.chaigneau@uai.cl

KEYWORDS

Conceptual Agreement Theory, evolution of concepts, Agent-based modeling, abstract concepts.

ABSTRACT

We present an Agent Based Model (ABM) named MIMICS (Modeling Inferential Minds in Conceptual Space), which shows how a social group develops abstract concepts for achieving agreement in communication. Agents describe concepts by assigning properties to them based on learning and communication interactions, trying to develop a conceptual space that discriminates as much as possible between two concepts (i.e., they try to assign properties to concepts decreasing the overlap among the properties that describe them). Contrarily to concrete concepts, those properties come from the social group and not from objects' physical properties. The results show that agents in MIMICS develop abstract concepts that exhibit the same characteristics that are found in studies of real concepts: non-uniform frequency distributions of properties, inter-subjective variability and stable concepts that are useful for the simulated social group by providing agreement in communication.

INTRODUCTION

Concrete concepts are typically associated to physical properties. For example, a type of face could be characterized by properties such as a given nose length, eye separation, mouth width, etc. (Tversky 1977). The standard theory of concrete concepts holds that concepts are learned by observing category exemplars, extracting the relevant properties (Schyns et al. 1998), and estimating the frequency distribution of those relevant properties (Ashby and Alfonso-Reese 1995; Griffiths et al. 2011). This then allows organizing a semantic structure and making category judgments (e.g., How typical is a given exemplar of category X ? Does the exemplar belong to category X or category Y ? How central is property j for category X ?).

Perhaps one of the most important findings about concrete concepts, is that concepts relate only probabilistically to conceptual properties (Rosch 1973). This means that, among other things, two concepts that may be applied to a situation or object are not discriminable through a logical rule (i.e., by necessary and sufficient properties), but show a typicality structure instead (Rosch et al. 1976). Those exemplars that exhibit

frequent properties are more typical than those exhibiting less frequent properties (Rosch and Mervis 1975) (e.g., an ostrich would be a low typicality exemplar of the BIRD category, whereas a dove would be a typical exemplar). The typicality structure also means that an object can be a member of different categories, although to different degrees (e.g., a Chihuahua may be a low typicality exemplar of the DOG category, while simultaneously being a relatively more typical case of the PET category). Note here that the fact that an exemplar may belong to more than one category, implies that concepts must share properties to a certain extent (e.g., *being friendly to people* may be a property of the concept DOG, but also of the concept PET). Henceforth, we will refer to this as “conceptual overlap”.

In contrast to concrete concepts, relatively little is known about abstract concepts (e.g., freedom, democracy, personality). This is a problem, given that a large proportion of the concepts that we use are abstract concepts (estimated to be more frequent than concrete words, Rechia and Jones 2012). Though the standard concrete concept theory assumes that it is a valid description of all kinds of concepts, there is evidence that abstract concepts do not respond to the same characteristics.

When researchers study concrete or abstract concepts they frequently ask a sample of individuals to produce lists of conceptual properties (e.g., Wu and Barsalou 2009). However, for abstract concepts subjects do not produce physical properties. Rather, they produce verbal associations (e.g., for the concept EMERGENCY, we might obtain *danger* as one of its properties; Della Rosa et al. 2010). When these lists are coded and aggregated, non-uniform or non-homogeneous frequency distributions of conceptual properties are obtained (these are called norming studies).

Another difference between concrete and abstract concepts is the following. Though concrete concepts may be learned without supervision (e.g., Love 2002), it does not seem possible to learn an abstract concept without some kind of supervision. A concrete concept may be learned by perceiving a sequence of exemplars, while extracting common properties. It is dubious that the same could be achieved for an abstract concept. Though there is no empirical support for this claim, it is difficult to imagine a list of exemplars that would allow learning, e.g., the concept of SECURITY without some kind of feedback. Furthermore, many abstract concepts refer to internal states that are not directly perceptible (e.g.,

INTENTION, Wu and Barsalou 2009) and are therefore clearly different from concrete concepts.

A working hypothesis of the work we report here is that the two issues discussed above (i.e., supervision requirement and verbal associations) are related. If abstract concepts cannot be learned by observing exemplars, it is possible that learning them requires attending to the behavior of other members of the social group when they use the concept. For example, it is possible that some kind of explicit teaching of the concept is necessary for someone to learn which other concepts are associated with the focal concept. If so, the need for supervision may cause that the conceptual content of abstract concepts is basically verbal.

Another important feature of property frequency distributions (for abstract as well as concrete concepts) is that they show inter-subjective variability in conceptual content (Barsalou 1993). In fact, given that subjects produce different lists of conceptual properties, non-uniform distributions follow.

Given our discussion above, our general research question in the current work is why abstract concepts (just as concrete concepts) should produce a probabilistic structure of properties (i.e., verbal associates), given that abstract concepts seem not to be learned as concrete concepts are. In the next section, we develop a meta-theory that offers an explanation, which we later implement as an Agent Based Model (ABM). Note that although we don't strictly follow the ODD protocol (Grimm et al. 2006) to present the ABM, we comply with including most of the material suggested in it.

A META-THEORY ABOUT CONCEPTS

Our meta-theory implies several factors that operate simultaneously. First, we assume that people can interact with a concept in two different manners: either learning or using it. At any given moment, individuals should make a decision regarding how to interact with a concept (not necessarily a conscious one). We assume that this decision depends on how much an individual knows how to differentiate the concepts in question. Though there are several potential ways in which individuals could determine if they know a concept well enough to use it confidently (e.g., they could pay attention to feedback from others regarding whether they are using the concept correctly), in the current work we assume that individuals attempt to discriminate as much as possible the focal concept from other potentially applicable concepts. Thus, the lesser they are able to discriminate, the more they are prone to learn something new about the concept in question.

Given that it does not seem possible to learn an abstract concept merely by perceiving exemplars, it is likely that these concepts are learned through explicit information acquired from others. There are several ways in which this could happen (e.g., individuals could directly ask others about the associated properties and construct their own frequency distribution in a piecemeal fashion). In the current work, we assume that individuals can learn

the content of an abstract concept by asking explicitly if a property corresponds to a concept.

As a consequence of learning more about the focal concept, we assume that individuals increasingly tend to decide to use a concept rather than continue learning. Classically, it would be assumed that a concept coded in language would be used to make reference (e.g., the word "dog" could be used to refer to a specific dog or to the category DOG). In contrast, here we assume that when using abstract concepts, individuals are trying to understand the point of view of a conversational partner (i.e., if she conceptualizes a situation as a case of the focal concept or as a case of an alternative concept). Here, again, there are several ways in which this could happen (e.g., an individual could observe the conceptual content produced by someone and by an associative process could gain information about which concept is being used). In the current work, we assume that individuals first adopt a given point of view (i.e., they conceptualize the situation as a case of a given concept) and look for confirmation that their conversational partner has the same point of view (Chaigneau et al. 2012).

Though searching for confirmation is a strategy that will lead to errors (Nickerson 1998), in our work we assume that a social group could use it to keep useful concepts (i.e., those that allow inferring the likely mental state of others). Looking for confirmation is a very simple strategy, which is likely to be used more than sophisticated processes (e.g., disconfirmation), and that does not require assuming elaborate cognitive processing.

DESCRIPTION OF MIMICS ABM

We designed an ABM that implements a specific version of the meta-theory described above (MIMICS; Modeling Inferential Minds in Conceptual Space). This theory assumes specific solutions to the topics discussed above, though — as also discussed above — other solutions are possible. Just to refresh them, the topics are the distinction between learning and using a concept, how an abstract concept may be learned, and what does it mean to use an abstract concept. Thus, our specific goal is to test if the ABM formalization is able to produce the pattern of results exhibited by abstract concepts: probability distributions of properties, absence of an objective criterion to define concepts and inter-subjective variability, and, despite all that, stability and usefulness of concepts.

In MIMICS, agents play two types of roles: observers (*O*) and actors (*A*), and act as *O*s and *A*s depending on the type of interaction executed (see Table 1 and associated explanations). Regardless of the role, agents know there are two concepts that can apply to a situation (*C*₁, *C*₂), and that there are properties (*j*) that can describe them. They also have a finite universe of *P* potential properties $j \in \{0, 1, 2, \dots, P - 1\}$ that can describe any of the two concepts $c \in \{1, 2\}$. These are not properties in a traditional sense (i.e., they are not independently

discriminable perceptual features), but rather verbal tags associated to concepts.

Agents develop their concepts either communicating with other agents or learning from them. For each concept ($C1$, $C2$), agents keep track of the number of occasions (f_j^c) in which they have found property p_j when interacting with the given concept c , and of the number of times in which they have interacted successfully (d_j^c) with that property p_j relative to that given concept c (see below for an explanation of what constitutes a successful interaction). In general, the greater d_j^c is in relative terms, the greater the evidence is for that property p_j to belong to that concept c . Note that the potential property j becomes a known property p_j (we will explain this process later on).

MIMICS has 2 mechanisms for concept development based on social interactions:

One is an implicit mechanism in which O is not attempting to learn, but to decide if A is in the same mental state as he is (we call this process, communication). In this process, O believes it knows the concept sufficiently and that there is no need to continue learning it. Then, in the communication mode, O assumes that the situation can be described by $C1$ (or $C2$), and waits for evidence that A conceptualizes it similarly. Then, A selects a concept c and a property p_j that belongs to c (p_j^c), and offers that property to O . If that property is in O 's concept $C1$ (or $C2$), then O assumes that both agree about the situation's definition. Consequently, O increases d_j^c and f_j^c for that property in concept $C1$ (or $C2$), otherwise, O increases only f_j^c (not d_j^c) for that property in concept $C1$ (or $C2$). Note that agreement can be true (A is really also thinking of concept $C1$ (or $C2$)) or it can be illusory (A is not really thinking in $C1$ but in $C2$ (or not in $C2$ but in $C1$)). In other words, ABM agents cannot read other agents' minds, and can only infer their mind states based on the evidence.

The other mechanism is one of explicit learning. If O believes it needs to learn more about concept $C1$ (or $C2$), then it looks for more information. For that, O queries A with a c, j pair (i.e., asks whether j is a property of c in A 's mind). If the query receives a negative answer, then O increases f_j^c but does not increase d_j^c (i.e., signaling that j has been experienced, but that it is not part of the focal concept). If the query receives a positive answer, then O increases f_j^c and d_j^c .

For each property p_j^c (i.e., each j in each concept $C1, C2$), agent O computes a success probability ($SP_j^c = d_j^c / f_j^c$) for interacting with that property p_j^c in that concept c . SP_j^c is the probability, computed from an agent's own experience, that it can achieve agreement when using a given property p_j^c in a given concept c .

The information obtained in communication and learning is used by O for two things:

First, it uses it to decide to which concept to assign a property p_j^c . The probability of p_j^c being assigned, e.g., to concept $C1$, increases probabilistically as the normalized absolute difference between the SP s for property p_j^c also

increases (i.e., how much an individual knows how to differentiate the concepts). In general, as the number of successful interactions when using a property p_j^c increases (i.e., those interactions that produce agreement), the evidence for that property belonging to that concept, and not to an alternative concept, also increases. In other words, to assess the possibility of discriminating a property between both concepts, the agents use $|SP_j^1 - SP_j^2|$. A small absolute difference shows that property p_j^c is not very discriminable (i.e., it produces about the same success probability for both concepts). This value is normalized to obtain what we define as the discrimination probability:

$$DP = \frac{|SP_j^1 - SP_j^2|}{\forall p_j \max |SP_j^1 - SP_j^2|} \quad (1)$$

In eq. (1), the absolute difference in SP for property p_j^c for both concepts ($|SP_j^1 - SP_j^2|$) is divided by the maximum difference across all known properties in O 's mind (see below for an explanation of how an agent knows properties), so that DP will always fall in the $[0,1]$ interval. Thus, using DP , an O agent will probabilistically decide if it has enough information to discriminate. If that is the case, the discrimination process is accomplished by comparing SP_j^1 with SP_j^2 , so that if $SP_j^1 < SP_j^2$, p_j^c is assigned to $C2$ and it is withdrawn from $C1$ (and vice-versa). This built-in preference for clearly separable concepts has been posed as a basic tendency in human categorization. If possible, people prefer to form linearly separable categories (Blair and Homa 2001).

Second, O uses the information obtained in communication and learning to decide if its next interaction with an A should be in the learning or in the communication mode (as described earlier). To this end, O computes a measure of the "separation" that the properties p_j^c have achieved. In MIMICS, this measure is the average absolute difference of all the properties' SP . Based on this average, O probabilistically decides in which mode to interact. An increase in this average value, signals an increase in separation, and results in a decreased learning probability (LP) for O (i.e., the probability that O decides to continue learning). However, because an agent knows A properties for a given concept, LP is really computed as a representative average:

$$LP = \frac{1}{A} \cdot \sum_{\forall p_j} 1 - |SP_j^1 - SP_j^2| \quad (2)$$

Note that because agents discriminate and also decide to stop learning depending on their own experience with conceptual properties, inter-subjective variability follows naturally in MIMICS. Table 1 presents the pseudo-code of the learning and communication interactions. MIMICS randomly selects without replacement an agent from the list of all agents and that agent acts as O and O randomly selects another agent as an A , following the

actions defined in Table 1. This process is executed until all agents have been O_s , which constitutes a simulation step.

Table 1: Pseudo-code of Learning and Communication Interactions

OBSERVER O	ACTOR A
Preparation of the interaction	
1. Randomly selects an A actor from the rest of the agents 2. Randomly selects a concept c and property j from the P potential ones 3. Decides interaction mode: $i Rdm(1) \leq LP?$	
$i Rdm(1) \leq LP? = \text{FALSE (Learning mode)}$	
	4. Selects same c as O 5. If $c = \emptyset$ (<i>auto-learning</i>) randomly selects a property j from the P potential ones ($p_j^c = j$) and increments f_j^c and d_j^c . 6. Assigns $p_j^c = j$ of O 7. $\exists p_j^c \Rightarrow result = 1$ $\neg \exists p_j^c \Rightarrow result = 0$
8. Increments f_j^c 9. $result = 1 \Rightarrow d_j^c = d_j^c + 1$ 10. $SP_j^c = d_j^c / f_j^c$ 11. $d_j^c \geq 0 \Rightarrow p_j \in c$ 12. Discrimination Inference: $Rdm(1) \leq DP \Rightarrow$ a) $SP_j^1 < SP_j^2 \Rightarrow p_j \notin 1, p_j \in 2$ b) $SP_j^1 > SP_j^2 \Rightarrow p_j \in 1, p_j \notin 2$	
$i Rdm(1) \leq LP? = \text{TRUE (Communication mode)}$	
	4. Randomly selects concept c and property p_j^c ($j = p_j^c$). 5. If $c = \emptyset$ (<i>auto-learning</i>) randomly selects a property j from the P potential ones ($p_j^c = j$) and increments f_j^c and d_j^c .
6. Assigns $p_j^c = j$ of A 7. $\exists p_j^c \Rightarrow result = 1$ $\neg \exists p_j^c \Rightarrow result = 0$ 8. Increments f_j^c 9. $result = 1 \Rightarrow d_j^c = d_j^c + 1$ 10. $SP_j^c = d_j^c / f_j^c$ 11. $d_j^c \geq 0 \Rightarrow p_j \in c$ 12. Discrimination Inference: $Rdm(1) \leq DP \Rightarrow$ a) $SP_j^1 < SP_j^2 \Rightarrow p_j \notin 1, p_j \in 2$ b) $SP_j^1 > SP_j^2 \Rightarrow p_j \in 1, p_j \notin 2$	

Note: LP see eq. (2), DP see eq. (1)

The initial conditions of a run consist in instantiating f_j^c , d_j^c and SP_j^c and the lists of properties that belong to concept $C1$ or $C2$ in each agent to null. Notably, interactions among agents depend on the interaction rules described above, and the only exogenous parameters are the number of agents (N) in a simulation and the number of potential properties for describing concepts $C1$ and $C2$ (P). This means that all results presented here can be

attributed to the interaction and decision rules (i.e., the meta-theory and MIMICS' solutions to the topics discussed earlier), and not to the way specific parameters were set in the experiments, except for N and P .

To simplify Table 1, we did not include the process by which agents know the properties. Agents know the existence of potential properties when properties are used in any of the interaction modes shown in Table 1. An agent knows property j by initializing $p_j^c = j$ and $d_j^c = f_j^c = SP_j^c = 0$, $\forall c$. There are two exceptional cases: 1) in communication mode, given that A presents to O a known property and does not receive anything from O , new properties are not incorporated by A as known and 2) when $c = \emptyset$ (i.e., when there is no conceptual content in c at initial conditions) for A in learning mode, A does an auto-learning process; that is, A not only initializes the property j , but also assigns it to c ($d_j^c = f_j^c = 1$) (i.e., property j becomes a known and assigned property p_j^c), which always happens at the beginning of a simulation run, when agents don't have any information or structure in their particular conceptual space. Finally, each run is ended when, at the social group system's level, the conceptual space structure is stable. That is determined when no further change is observed for the properties incorporated into concepts at the group's level. This occurs when the standard deviation of the average SP of both concepts (across all agents), calculated in a sliding window of 3,000 simulation steps, does not show significant variations; i.e., the standard deviation of the standard deviation of the average SP of both concepts is equal to or less than 0.004.

EXPERIMENTS AND RESULTS

Our general hypothesis is that a process based on social interactions (communication and learning), where conceptual properties come from the social group and not from objects' physical properties, is able to produce concepts characterized by non-uniform probability distributions and inter-subjective variability in conceptual content, while making minimal assumptions about agents' cognitive machinery. Specifically, we expect that, for a wide range of experimental conditions (N and P values), MIMICS will produce stable concepts that are useful for the simulated social group, but not at the expense of homogeneity in conceptual content (i.e., MIMICS should exhibit inter-subjective variability in conceptual content). Also, as a direct consequence of this, MIMICS should produce non-uniform frequency distributions of properties similar to those found in norming studies. For the experiments we set up $N = \{14, 40, 60\}$ and $P = \{10, 50, 100\}$. We selected those values for representing small, medium and large groups of agents and number of potential properties. Each of the nine experimental conditions was run 20 times and in all the graphs that show averages, these were computed using the output values of the 20 replications. We don't present std. deviations, given that they are very small and only would have cluttered the graphs. We performed an ANOVA for all the presented results (where suitable),

which indicates that all of them are highly statistically significant (all p -values ≤ 0.005). In the following paragraphs we present the results for concept $C1$, given that the ones for concept $C2$ are similar. For those interested in replicating our experiments, the program is available at <http://ccl.northwestern.edu/netlogo/models/community/>. You will have to search for the file MIMICS v-CSI.netlogo, found under the March 2016 heading, and download it to your computer. Then you need to download and install the Netlogo platform, version 4.0.4 at <http://ccl.northwestern.edu/netlogo/oldversions.shtml>. To assess the usefulness of concepts, we use the probability of true ($p(a1)$) and illusory agreement ($p(a2)$) per Conceptual Agreement Theory (CAT, Chaigneau et al. 2012). According to CAT, when human beings talk about abstract concepts (e.g., democracy, political views, masculinity, personality traits), they try to infer agreement, i.e., to infer whether other people's mind-content is similar to their own content or not. To illustrate, imagine two individuals, O and A , that are having a conversation about a given topic, and that O has a hypothesis $C1$ about how entity x is being jointly conceptualized (i.e., that they are talking about x as an instance of $C1$). However, because concepts are events in individual minds, O can only infer whether $C1$ is the case for A or not. To make this inference, O observes A , and when A describes x as having a property p_i , O evaluates if p_i is consistent with $C1$ in her mind or not. If it is consistent, then O infers that A is also talking about x conceptualized as $C1$. If A is in fact talking about x conceptualized as $C1$, then this is true agreement (event $a1$ and its probability is $p(a1)$). If A is talking about x conceptualized as $C2$, then illusory agreement happens (event $a2$ and its probability is $p(a2)$). Note that this situation corresponds to the idealized communication interaction shown in Table 1. In MIMICS, to compute $p(a1)$, each time agents engage in a communication interaction and both are using concept $C1$, a counter f_{a1} is incremented. On the other hand, if O is using concept $C1$ and A is using $C2$, a counter f_{a2} is incremented. If agents infer agreement and that is actually true agreement (both agents are actually thinking of $C1$), then a counter $a1$ is incremented. Contrarily, if agent O is thinking of $C1$ and agent A is thinking of $C2$, then a counter $a2$ is incremented. Calculating $p(a1)$ and $p(a2)$ amounts to dividing $a1$ by f_{a1} and $a2$ by f_{a2} . Given that concepts should afford a $p(a1)$ larger than $p(a2)$ to be useful in communication among members of a group (i.e., more true than illusory agreement; Chaigneau et al. 2012), we should observe the same in MIMICS' outputs. As Figure 1 shows, that is the case. For all the nine experimental conditions, always $p(a1)$ is larger than $p(a2)$, which means that agents develop a conceptual space that promotes true agreement in communication.

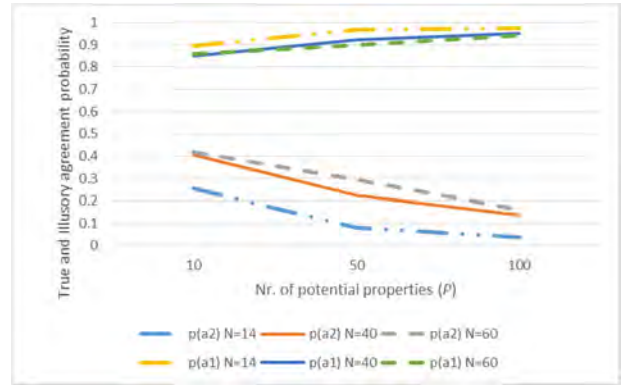


Figure 1: Avg. $p(a1)$ and $p(a2)$ for the 9 Experimental Conditions

On the other hand, although a high true agreement is reached, agents exhibit inter-subjective variability in conceptual content. To illustrate this, we can inspect the properties assigned to concept $C1$ by two agents in a given experimental condition ($N = 40$, $P = 100$). For example, agent 0's content for $C1$ is [2 4 5 8 10 16 17 43 63 67 68 71 74 77 78 90 97], whereas agent 10's content is [8 10 27 33 34 38 62 64 68 75 77 85 97]. To generalize this claim, Figure 2 shows the frequency distribution of the properties across agents for $C1$, for the same experimental condition. It can be seen that the distribution is non-uniform (which also supports our assertion that MIMICS would produce non-uniform frequency distributions of properties). Given that the distribution is non-uniform, the only way that may happen is if agents have diverse conceptual contents. To more generally back up our claim, Figure 3 shows MIMICS' outputs k_i and s_i . Variable k_i corresponds to the total number of properties for concept $C1$ in a population of individuals, and s_i to the average number of properties coherent with concept $C1$ in an individual's mind (Chaigneau et al. 2012).




Figure 2: Frequency Distribution of Properties across Agents for Concept $C1$ ($N = 40$, $P = 100$)

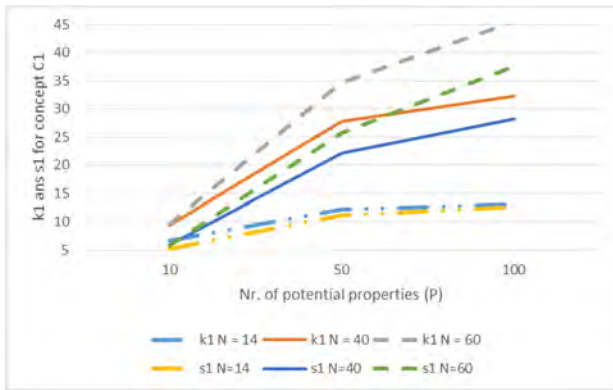


Figure 3: Avg. k_1 and s_1 for the 9 Experimental Conditions

From Figure 3, we can see that for all the experimental conditions, k_1 is always larger than s_1 . That may happen only if the number of properties assigned to $C1$ in agents' minds is smaller than the total number of properties assigned to $C1$ across all agents, which proves that inter-subjective variability in conceptual content must exist. As already discussed, Figure 2 supports our claim that MIMICS would produce non-uniform frequency distributions of properties similar to those found in norming studies. To generalize this finding to all the nine experimental conditions, Figure 4 presents the standard deviation for the properties for $C1$ (i.e., for the numbers that represent those properties).

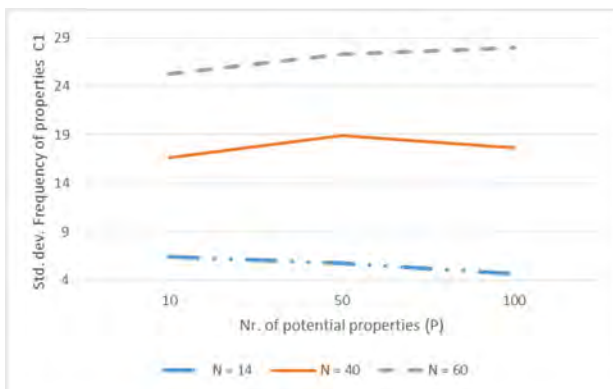


Figure 4: Avg. Std. Deviation of Properties of $C1$ for the 9 Experimental Conditions

The standard deviations different from zero confirm that for all experimental conditions, the frequency distributions of properties are non-uniform.

Finally, as pointed out earlier, another characteristic of these frequency distributions obtained in norming studies, is that the properties, which describe concepts, exhibit some conceptual overlap (i.e., some properties describe more than one concept, as one can see from Figure 2). To generalize that finding, Figure 5 shows a normalized RMSE of the frequencies of the properties that describe $C1$ and $C2$ in MIMICS. The normalized RMSE was calculated by first dividing the frequency of each property that describes $C1$ and $C2$ by the respective maximum frequency. Then, the sum of the squared difference between the normalized frequencies of $C1$ and

$C2$ was divided by the number of frequencies that are larger than zero in both concepts, and the square root of that mean is the RMSE.




Figure 5: Avg. Normalized RMSE of Properties of $C1$ and $C2$ for the 9 Experimental Conditions

This form of calculating the RMSE assures obtaining a superposition index between the $C1$'s and $C2$'s property frequency distributions that is comparable across different N and P . That will be an important issue when further analyzing MIMICS' results in future work. Note from Figure 5 (and also from Figure 2), that under all nine experimental conditions, a RMSE above zero, implies that indeed there exists conceptual overlap between concepts, just as has been found in empirical studies.

DISCUSSION/CONCLUSIONS

As we discuss in the introduction, when subjects are asked to produce conceptual content for abstract or concrete concepts, non-uniform frequency distributions of properties are obtained. Also, there is inter-subjective variability in conceptual content. For concrete concepts, these properties are descriptors of the concrete objects that belong to the category, and non-uniform distributions occur because some properties are more frequent than others in the exemplars that belong to the category (e.g., most dogs bark). In contrast, for abstract concepts properties are verbal or conceptual associations, and it is unclear why non-uniform property distributions and inter-subjective variability should be obtained.

In the current work, we present MIMICS, which is a theory of how a social group develops a system of abstract concepts. MIMICS makes three important assumptions about abstract concepts (widely supported by the literature that we have cited throughout this paper). First, it assumes that abstract concepts are states of mind or points of view about a situation. As such, they cannot be directly observed and need to be inferred. Second, it assumes that individuals are motivated to know if other individuals share their own particular point of view. Third, it assumes that—as is true of concrete concepts—people attempt to learn linearly separable concepts. Noteworthy about MIMICS, is that conceptual content develops from social interaction and not from the environment's structure. There are two kinds of interactions: learning from other group members, and

communicating (i.e., using conceptual properties to infer a conversational partner's state of mind).

In our computational experiments with MIMICS, we found that for a wide range of experimental conditions (i.e., combinations of N and P values), MIMICS reproduces the type of results that are obtained in conceptual norming studies. MIMICS' rules of interaction are successful in producing non-uniform property frequency distributions, concepts that are not neatly discriminated, and agents with non-homogeneous conceptual content, though agents do not extract this structure from an environment. Just as importantly, MIMICS produces concepts that, despite inter-subjective variability, allow communication. Concepts developed by MIMICS allow agents more often than not to correctly infer conceptual agreement with other agents.

There are several situations in which a researcher may need to explore the use of abstract concepts in a social group. An anthropologist may want to know whether a social group holds a shared view on a socially relevant topic (e.g., how political parties are characterized); a marketing expert may want to know whether a social group holds a shared view on a brand's image). We envision using MIMICS and the theoretical insights derived from our simulations as providing tools to analyze problems such as these.

There are many other conclusions that can be drawn from our results. However, they are beyond the scope of this paper and will be part of our future work with MIMICS. What we want to stress here is that MIMICS shows that abstract concepts may be advantageously viewed as devices developed by a social group to allow agreement and mind-reading.

ACKNOWLEDGMENT

This work was supported by FONDECYT (Fondo Nacional de Ciencia y Tecnología of the Chilean Government) grant Nr. 1150074 to the last two authors.

REFERENCES

- Ashby, F.G. and Alfonso-Reese L.A. 1995. "Categorization as Probability Density Estimation." *Journal of Mathematical Psychology* 39, No 2, 216-233.
- Barsalou, L. W. 1993. Flexibility, structure, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In A. C. Collins, S. E. Gathercole, & M. A. Conway (Eds.), *Theories of memories*. London: Erlbaum. pp. 29-101.
- Blair, M. and Homa, D. 2001. "Expanding the search for a linear separability constraint on category learning." *Memory and Cognition* 29, No8, 1153-1164.
- Chaigneau, S.E.; Canessa, E.; and Gaete, J. 2012. Conceptual agreement theory. *New Ideas in Psychology* 30, Nr 2, 179-189.
- Della Rosa, P.A.; Catricalà, E.; Vigliocco, G.; and Cappa, S. 2010. Beyond the abstract-concrete dichotomy: mode of acquisition, concreteness, imageability, familiarity, age of acquisition, context availability, and abstractness norms for a set of 417 Italian words. *Behavior Research Methods* 42, 1042-1048.
- Griffiths, T.L.; Sanborn, A.N.; Canini, K.R.; Navarro, D.J.; and Tenenbaum, J.B. 2011. "Nonparametric Bayesian models of categorization." In E.M. Pothos and A.J. Wills (Eds.) *Formal Approaches in Categorization* 2011, pp. 173-198. Cambridge: Cambridge University Press.
- Grimm, V. et al. 2006. "A standard protocol for describing individual-based and agent-based models". *Ecological Modelling* 198, 115-126.
- Love, B. C. 2002. "Comparing supervised and unsupervised category learning." *Psychonomic Bulletin & Review* 9, 829-835.
- Nickerson, R.S. 1998. "Confirmation bias: A ubiquitous phenomenon in many guises." *Review of General Psychology* 2, Nr 2, 175-220.
- Recchia, G. and Jones, M.N. 2012. "The semantic richness of abstract concepts." *Frontiers in Human Neuroscience* 6, Art. 315, 1-16.
- Rosch, E. 1973. "On the internal structure of perceptual and semantic categories." In T. E. Moore (Ed.), *Cognitive Development and the acquisition of Language* 1973. New York: Academic Press.
- Rosch, E. and Mervis, C.B. 1975. "Family resemblances: Studies in the internal structure of categories." *Cognitive Psychology* 7, 573-605.
- Rosch, E.; Simpson, C.; and Miller, R.S. 1976. "Structural bases of typicality effects." *Journal of Experimental Psychology: Human Perception and Performance* 2, Nr 4, 491-502.
- Schyns, P.G.; Goldstone, R.L.; and Thibault, J.P. 1998. "The development of features in object concepts." *Behavioral and Brain Sciences* 21, Nr 1, 40-41.
- Tversky, A. 1977. "Features of similarity." *Psychological Review* 84, No 4, 327-352.
- Wu, L.L. and Barsalou, L.W. 2009. "Perceptual simulation in conceptual combination: Evidence from property generation." *Acta Psychologica* 132, 173-189.

AUTHOR BIOGRAPHIES

CARLOS BARRA holds a postgraduate degree in Integrated Political Science (2005) from the Chilean Naval Academy, an MBA (1999) from the Institute for Executive Development (IEDE - Chile) and a Master of Science in Computer Engineering (1996) from F. Santa María University, Chile. His research interests include the study of complex systems.

ENRIQUE CANESSA is an associate professor at Universidad Adolfo Ibáñez, Chile. He holds a PhD in MIS/CIS (2002), a Certificate of Graduate Studies in Complex Systems (2001) and an MBA (1991) from the University of Michigan, USA. His research interests include the study of organizations, sociology and cognitive psychology using ABM.

SERGIO E. CHAIGNEAU is full professor at Universidad Adolfo Ibáñez, Chile. He holds a PhD in Cognitive and Developmental Psychology (2002) from Emory University, USA, and a Master of Arts (1995) in Experimental Psychology from the University of Northern Iowa, USA. His research interests include the study of causal categorization and inter-subjective agreement.