Ordered models for concept analysis

Michael Freund

Ordered models for concept analysis

Michael Freund
IHPST
University of Paris Sorbonne
13 rue du Four 75006 Paris, France
micfreund@gmail.com

Abstract

In this paper, I propose a simple ordered model for categorization theory. When a concept is grasped through an auxiliary set of features, stemming from a description of the concept or from a resemblance to some typical exemplars, categorial membership, typicality and resemblance can be accounted for by means of weak order relations. These orders render precise the notion of an object 'falling more' than another under this concept, being 'a more typical exemplar' or 'resembling more' this concept. The membership and typicality orders can be naturally extended to compound concepts, without the drawbacks that are classically encountered in conjunction theory. Moreover, these orders reveal themselves to be a particularly adequate tool for the resolution of problems linked with category-based induction.

Keywords categorization, prototype theory, categorial membership, typicality, deep learning, category-based induction, resemblance theory.

1 Introduction

How do we form concepts ? What makes us decide that different objects may be gathered under a common label ? And once this label has been set, how do we account for its adequacy for new items ? These questions stand at the root of all problems concerning the categorization process. Several models has been proposed since the primitive formalism of Frege (Frege,

1879) in which concepts were assimilated with simple one-place predicates. Apart from letting aside the problem of typicality, this view reduced categorization to an all-or-nothing process that only two-valued truth functions could account for. After the work of Eleanor Rosch however, it appeared that categorial membership and typicality should be evaluated by more sophisticated tools: degrees functions (Rosch, 1975), membership functions (Zadeh, 1965), geometrical measures (Gärdenfors, 2000), or quantum mechanics (Aerts, 2009). In this perspective several models were proposed, all aiming at a unified quantitative theory for categorization.

My approach differs from the preceding ones, in that I consider categorial membership and typicality as *comparative* notions. As I argued in (Freund, 2008) and (Freund, 2014), the relative strength with which a concept applies to an object is best accounted for by a (pre)order relation among the objects of the universe. To take an example, an agent may be unable to assign a precise *birdhood* degree to a bat or to a tortoise, while ready to attribute more *birdhood* to one of these items than to the other. Similarly, one may consider the duck as a *more typical* exemplar of a bird than the penguin, without having an idea of what their typical degree could be.

In this framework, every concept endows the agent's universe of discourse with a couple of *conceptual orders*. The first one, the membership order, compares the strength with which a concept applies to the objects of the universe. It may be seen as a way of comparing objects that do not necessarily fall under the concept. The second one, the typicality order, is meant to compare the relative typicality of the exemplars of the concept. The knowledge of these two orders is equivalent to that of the associated concept: two concepts agree if and only if they have same membership order and same typicality order. hotel plein ciel

The determination of the conceptual orders is not always possible. In (Freund, 2008) and (Freund, 2009), an explicit construction was proposed for concepts that belonged to a specific class - the family of *constructible* concepts. However, this construction suffered from two drawbacks: first, the set of constructible concepts is difficult to circumscribe, as it rests on the difference between *sharp* and *vague* concepts. Second, the construction of the membership and the typicality orders required the use of recursively defined degree functions, thus contradicting the philosophy of a supposedly purely comparative theory. The aim of the present paper is to remedy these failures and propose, for a sufficiently large family of concepts, a coherent and general theory based on natural conceptual orders, sufficiently robust to

address the principal problems encountered in categorization theory.

This paper is organized as follows: in Section 2, the difference between concepts and features is evoked and discussed. This leads to the notion of *featured concepts*, which is central in this paper. A construction of the membership order, based on the defining feature set for simple concepts, is then developed in Section 3. This order gives rise to an auxiliary definition of membership degree. Membership orders and degrees can be extended to the class of compound concepts obtained through the determination operator; this provides an interesting framework for the study of the conjunction effect. Section 4 deals with typicality for simple and compound concepts, while questions linked with resemblance theory are discussed in Section 5. All these notions being settled, it is then possible to tackle some classical problems linked with category-based induction: this is the object of Section 6.

.

2 Featured concepts

2.1 Back to the attributional theory

The classical theory of categorization viewed concepts as theoretical objects that could be defined by a certain number of properties: vertebrates that have beak and feathers were labelled *birds*, a *car* could be defined as a road vehicle powered by an engine able to carry a small number of persons, and *democracy* was a system of government by all the eligible members of a state. In this perspective, where each concept was endowed with a set of attributes, categorization relative to a concept boiled down to categorization relative to its features. This attributional view was advocated by some authors in the late seventies (Smith, Shoben, & Rips, 1974) and (Smith & Medin, 1981). It gave rise to the so-called binary model, in which two auxiliary sets are attached to a concept. On the one hand, the *defining feature set* provided the conditions that an item should satisfy in order to fall under a given concept; on the other hand, the *characteristic set* listed the features that an object should have to be qualified as a *typical* instance of this concept. Given for instance the concept *to-be-a-fruit*, one may take as defining feature set the set consisting of the two elements *to-be-a-vegetable* and *to-heave-seeds*, while the characteristic set would include features like *to-grow-on-trees*, *to-be-sweet*, *to-be-raw-edible*, *to-yield-juice*.

The attributional theory was thereafter rejected by most researchers, as it appeared that concepts defined by a conjunction of features formed an exceptional subclass. Fodor for instance (Fodor, 1998) argued that there exists practically no examples of successful definition around. Without being so radical, it is clear that a great deal of concepts are deprived of any set of defining features: what list of attributes could be attached for concepts like *to-be-a-citrus fruit*, *to-be-a-lie*, or *to-be-a-heap* ? However, in spite of this, it appears that the attributional view is justified for certain well-defined families of concepts: such is for instance the case for most *nominal* concepts, i.e. concepts that are conventionally defined, like *to-be-a-mammal*, *to-be-a-theft* or *to-be-a-refugee*. In particular, this remains true for scientific concepts, like *to-be-a-vortex* or *to-be-a-square*. Furthermore, it may happen that a concept, first grasped through its exemplars, is thereafter sharpened with the help of a set of defining features: thus, the concept *to-be-a-bird*, a natural kind concept, was revisited by naturalists and turned into the pseudo-nominal *to-have-feathers + to-have-a-beak + to-have-wings*. Note also that even for non-definable concepts, class membership may still depend on auxiliary features: this is the case when categorial membership is induced through resemblance to a prototypical exemplar. Then, it is the features of this prototype that play the role of defining features.

These considerations show that the family of attributional concepts is large enough to deserve a treatment of its own. Since attributional concepts can be studied with the help of quite elementary mathematical tools - contrary to the examples evoked in the preceding paragraph - there is no reason not to devote to them a special study. Such is the aim of the present paper.

It is necessary to emphasize at this point that this study does not aim at providing an *objective* model for categorization, even for the restricted area of attributional concepts. Indeed, the defining feature set and the characteristic set on which this construction is based are by no means well-defined. They vary from an agent to another and, for a given agent, they may change with time. It may even happen that a concept is seen as attributional by an agent, while deprived of this quality by another one. It should be therefore understood that the model proposed in this paper is meant to reflect, at a given precise time, the single point of view of a subjective particular agent.

2.2 Concepts and features

The terms *concept* and *feature* cover different notions. Formally, concepts are most often introduced through the auxiliary *to-be*, followed by a noun: *to-be-a-bird, to-be-a-vector-space, to-be-a-democracy.* Features may be presented through a verb (*to-fly*), the auxiliary *to-have*, followed by a noun (*to-have-a-beak)*, or the auxiliary *to-be*, followed by an adjective (*to-be-tall*). While concepts appear as unary predicates, this condition is no more necessary for features. Features, like concepts, apply to the objects at hand but, contrary to concepts, they are context-sensitive: they borrow part of their significance from the concept they are attached to. Properties like *to-be-tall*, *to-be-rich* or *to-be-red* take their full meaning only in a given context, that is when qualifying a well-defined entity. Even simple verbal forms like *to-fly*, *to-run*, *to-live-in-water*, *to-be-made-of-metal* need a principal referent concept to fully seize the strength with which they apply to different items. Thus, the concept a feature applies to may be seen itself as a contextual determination of this feature. To summarize, the meaning of a feature depends on the context in which this feature is used, contrary to the meaning of a concept, which exists by itself.

It does not seem at this stage that any formalism can fully account for the difference between features and concepts. It is true that in Description Logics, a different treatment is applied for one and two-places predicates: binary predicates characterize indeed the *roles* of the language, which are used to express relationship between the concepts (Nardi & Brachman, 2003). In this famework, *to-be-a-tree* will be a concept, expressible by a single symbol A, but *to-have-green-leaves* is a role, expressed by a formula of the type '$\forall$ hasLeaves.Green'. However, no difference is made in Description Logics between the unary predicates that translate a notion of concept and those that translate a notion of feature.

It should be noted that, when considering a feature f of a concept α, the way f applies to an object is generally itself related to the categorial membership of an auxiliary concept β. Consider for instance the feature *to-have-wings*. In order to evaluate to which degree this feature applies to an item x, we must be able to determine what exactly covers the concept *to-be-a-wing*, and which objects fall under it (for a discussion on the difference between *knowing what an X is* and *knowing what it is to have an X*, see (Fodor, 1998). This observation shows that circularity cannot be avoided when trying to build a general mathematical model for categorial member-

ship. However, such is not the purpose of this work, which aims at analyzing how the knowledge of a concept can be drawn from the knowledge of its associated set of features, without questioning the nature and the sources of the latter.

2.3 Applicability functions and applicability orders for concept features

The strength with which a feature f of a concept α applies to an object is usually measured in a given ontology through a percentage or an *applicability degree function* δ_f^α that takes its values in the unit interval. Concerning the range of this function, it is important to observe that it can be most often circumscribed to a *finite* subset of $[0, 1]$: this is clearly true for fuzzy features like *to-be-tall*, *to-be-rich* or *to-be-warm*, since the measure of their applicability is always approximative (to an inch, a cent or a degree). On the Brittany coast of France, for instance, the set of water temperatures t in July ranges for from 15 to to 25 degrees Celsius degrees, thus covering 11 possible (integral) values. In this context, the function associated with the concept *warm* may be given by $\delta_f^\alpha = t/10 - 1.5$. The finiteness of the range of δ_f^α is even more obvious in the general process of categorization: ranking the objects relatively to a feature of a given concept only yields a small number of equivalence classes. To determine, for instance, to which extent a flower may be considered as a *poppy*, one roughly evaluates its redness, its shape and the size of its petals. Concerning the redness, comparison with other objects shows that only a finite number of discernible reds separate the color of that particular flower from that of an ideal poppy. Thus, in the context of a *to-be-a-poppy*, there exists only a small number of possible degrees of redness. The same observation can be made concerning the other features that define or describe poppies, like the shape and the size of the petals.

For this reason, this paper will be devoted to the family of concepts whose features can be weighed on a finite scale. These concepts form the family of *featured concepts*. Given a featured concept α, the way any of its associated feature f applies to an item is accounted for by an applicability degree function δ_f^α that takes only a finite number of values. Equivalently, we may say that f, as an α-feature, generates an *applicability order* $\preceq_f^\alpha$ defined for any objects x, y of the universe of discourse by: $x \preceq_f^\alpha y$ if and if $\delta_f^\alpha(x) \leq \delta_f^\alpha(y)$. The relation thus defined is a total order that has only

finitely many equivalence classes. There exists therefore only a finite number of intermediate states between an object x totally deprived of f and an object y to which the feature f fully applies.

Apart from the finiteness condition, featured concepts will be required to satisfy an *agreement condition* to guarantee the existence of at least an object to which simultaneously apply all the defining and characteristic features associated with a concept α.

This leads to the following definition:

A featured concept α is a concept for which there exists a finite set of defining features $\Delta(\alpha)$ and a finite set of characteristic features $\Xi(\alpha)$ that satisfy the two properties :

1. *for every defining or characteristic feature f, the corresponding applicability function δ_f^α takes a finite number of values*

2. *There exists at least an item z such that $\delta_f^\alpha(z) = 1$ for all defining or characteristic features f.*

As was recalled in paragraph 2.1, in the agent's mind the set $\Delta(\alpha)$ is constituted by attributes that appear to be *essential* for the concept realization: the only objects to which α applies are those that possess these attributes. On the contrary, the characteristic features are only present in the *typical* exemplars of the concept.

The family of featured concepts will be denoted by $\mathcal{F}$. In order to lighten the notations, the superscript of the applicability functions and degrees will be omitted, so that $\preceq_f$ and δ_f stands for $\preceq_f^\alpha$, and δ_f^α. However, it is necessary to keep in mind that the applicability order or the applicability function of a feature is always defined relatively to the concept it qualifies.

3 Membership orders

In general, the human mind has no tool at its disposal to directly evaluate the categorial membership of an object: in the extremal cases, we may be able to decide that a given object is or is not a member of the concerned category, but we do not know in the intermediate states how to quantify its partial membership: a conventional bomb is definitely not a *weapon-of-mass-destruction* (WMD), and the same is true for a machine gun, yet, we

are unable to attribute a precise degree of membership to any of these items taken alone. We can only compare them. Similarly, we are unable to assign a precise membership degree to a sink as a *piece of furniture*, while being ready to admit that it is 'more' a piece of furniture than a heat-pipe, and 'less' a piece of furniture than a window.

As a matter of fact, concerning membership evaluation, and apart from the binary distinction between members and non-members of a category, the best thing the human mind seems to be capable of is to *compare* two objects and decide which one, if any, falls 'more' under the concerned concept. Thus, the concept *to-be-a-weapon-of-mass-destruction* will be generally considered as applying *more* to a machine-gun than to an arquebus, and *less* to a spear than to an arquebus. Clearly, this judgement shows the existence of a basic ordering induced by the concept *to-be-a-weapon-of-mass-destruction* in the universe of discourse. This ordering is by no means a consequence of a supposed degree assignment that the agent has set *a-priori* on the objects at her disposal: if directly questioned what membership degree should be attributed to a machine-gun considered as a WMD, an agent will be generally unable to provide a sensible answer. Of course, such an assignment may be established once comparison has been made between the items at hand. For instance, a non-decreasing ranking like *bludgeon* $\leq$ *sword* $\leq$ *crossbow* $\leq$ *arquebus* $\leq$ *gun* $\leq$ *machine-gun* $\leq$ *flamethrower* $\leq$ *conventional bomb* $\leq$ *scud* $\leq$ *atomic bomb* may yield *a posteriori* a membership degree of the concerned items, which can be readily visualized from their position on a [0,1] scale: thus, an *atomic bomb* will be considered as being 100% a WMD, a *scud* as 90%, a *conventional bomb* as 80% and so on. The point is that these numerical values will appear as a consequence of a pre-recognized order among the different weapons that are part of the agent's universe: they will not be at the origin of it. This construction of a membership degree as a secondary tool, stemming from a membership order, will be examined in details in section 3.3.

Order relations therefore appear to provide the most adequate model to account for categorial membership as perceived by a cognitive agent. Appealing systematically to relations of this type whenever it is possible avoids the drawbacks that may result from the application of more sophisticated theories. It is true that in some cases, order relations may be insufficient to fully treat categorial membership - this will be for instance the case for *fuzzy concepts*, or for vague concepts of a continuous type. However, concerning the specific class of featured concepts, the tool provided by weak order

relations is sufficiently powerful to fully address the categorization problem.

In this perspective, the fact that a concept α may apply *more* (or *better*) to an item x than to an item y will be translated by a *preorder* (reflexive and transitive) relation on the set $\mathcal{O}$ of objects, real or imaginary, that are part of the agent's universe of discourse. This relation will be denoted by $\preceq^{\mu}_{\alpha}$. The expression '$x \preceq^{\mu}_{\alpha} y$' therefore translates the sentence 'the concept α applies at least as much to the object y as to the object x'.

The corresponding strict partial order will be denoted by $\prec^{\mu}_{\alpha}$, that is $x \prec^{\mu}_{\alpha} y$ if and only if $x \preceq^{\mu}_{\alpha} y$ and not $y \preceq_{\alpha} x$.

In the general case, and contrary to most of the existing theories, the relation $\preceq^{\mu}_{\alpha}$ will not be supposed to be total. There is no reason indeed to *a priori* eliminate the case where the membership of two objects is incomparable. Taking again the *bird* example, this enables us for instance to deal with the case where an agent refuses to compare the birdhood of a bat with that of a tortoise.

3.1 The case of featured concepts

Let now α be a featured concept and suppose that its defining feature set $\Delta(\alpha)$ consists of the k features $f_1, f_2, \ldots, f_k$. In the perspective of the *attributional view*, the applicability orders $\preceq_{f_1}, \preceq_{f_2}, \ldots, \preceq_{f_k}$ induced on $\mathcal{O}$ by these features are part of the agent's knowledge. The assumption that the knowledge of the features is sufficient to acquire knowledge of the target concept requires that the membership order $\preceq^{\mu}_{\alpha}$ should naturally stem from the $\preceq_{f_1}, \preceq_{f_2}, \ldots, \preceq_{f_k}$. One could then simply consider that $\preceq^{\mu}_{\alpha}$ is the intersection of the $\preceq'_{f_i}$s, but this would ignore the internal structure of the set $\Delta(\alpha)$. Indeed, the defining features of α cannot be considered as equivalent. In the agent's mind they do not necessarily weigh the same weight. For instance, a particular agent may associate with the concept *to-be-a-bird* the defining set {*to-be-a-vertebrate, to-be-oviparous, to-have-feathers, to-have-a-beak, to-have-wings*}, and consider moreover that *having wing*s is a more important feature for birdhood than *having a beak*. For this agent consequently, a bat will be given more birdhood than a tortoise. To account for this phenomenon, it is necessary to equip the set $\Delta(\alpha)$ with a *salience* relation, which will be simply translated by a strict partial order $>_{\Delta(\alpha)}$.

The membership order induced by α can be now defined by:

$x \preceq^{\mu}_{\alpha} y$ iff for each defining feature f_i such that $y \prec_{f_i} x$, there exists a

defining feature f_j, $f_j >_{\Delta(\alpha)} f_i$, such that $x \prec_{f_j} y$.

Thus, the concept α applies at least as much to y as to x if each of its defining feature that applies more to x than to y is dominated by a defining feature that applies more to y than to x (see (Freund, 2014) for a justification of this construction).

The relation $\preceq^{\mu}_{\alpha}$ sets a partial (weak) order on the set $\mathcal{O}$. Its associated strict partial order reads

$x \prec^{\mu}_{\alpha} y$ if and only if $x \preceq^{\mu}_{\alpha} y$ and there exists some defining feature f_i for which $x \prec_{f_i} y$.

Example 1 *Let α be the concept* to-be-a-bird*, and suppose that, from the point of view of an agent, its defining feature set in the context of* living beings *is the set* {to-have-two-legs, to-lay-eggs, to-have-a-beak, to-have-wings} *with the following salience order:* to-have-a-beak $>_s$ to-lay-eggs $>_s$ to-have-two-legs*, and* to-have wings $>_s$ to-lay-eggs$>_s$ to-have-two-legs*. Suppose for the sake of simplicity that, in the agent's mind, membership to any of these features is a two-valued function. Let s, m, t, b and d respectively stand for a sparrow, a mouse, a tortoise, a bat and a dragonfly. Then the induced membership order is determined by the following arrow :*

	$two-legs$	$lay-eggs$	$beak$	$wings$
sparrow	$\star$	$\star$	$\star$	$\star$
mouse				
tortoise		$\star$	$\star$	
bat	$\star$			$\star$
dragonfly		$\star$		$\star$

One readily checks that $d \prec^{\mu}_{\alpha} s$, $m \prec^{\mu}_{\alpha} t$, and $m \prec^{\mu}_{\alpha} b$. Note that one has $b \preceq^{\mu}_{\alpha} d$, since the concept to-have-two-legs *under which the bat falls, contrary to the dragonfly, is dominated by the concept* to-lay-eggs *that applies to the dragonfly and not to the bat. On the other hand, one does not have $d \preceq^{\mu}_{\alpha} b$, as nothing compensates the fact that the dragonfly lays eggs and the bat does not. This yields $b \prec^{\mu}_{\alpha} d$. Note also that the tortoise and the bat are incomparable: one has neither $b \preceq^{\mu}_{\alpha} t$, nor $t \preceq^{\mu}_{\alpha} b$.*

The strict α-membership order therefore reads:

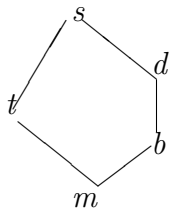

3.2 Essence and Extension

The membership order $\preceq^{\mu}_{\alpha}$ adequately translates the notion of membership to a category. Indeed it follows from the *agreement condition* that there exist $\preceq^{\mu}_{\alpha}$-maximal elements in $\mathcal{O}$, namely those that fall under all the defining features of α. These objects form the *extension*, or the *category*, of the concept α, denoted by $Ext\,\alpha$. Its elements constitute the *exemplars*, or the *instances*, of the concept.

This definition by means of maximal membership conforms with the intuition: an object x (fully) falls under a concept α if α cannot apply more to an object y than to x. Whatever salience order is set on $\Delta(\alpha)$, an object falls under the featured concept α if and only if it falls under each of its defining features. In other words one has $Ext\,\alpha = \bigcap(Ext\,k)_{k\in\Delta(\alpha)}$. One retrieves here the classical characterization of a defining feature set as a set of features that are 'individually necessary and jointly sufficient to ensure membership relative to α'.

As mentioned above, the set $\Delta(\alpha)$ has no objective character: it consists of the features that a given agent would choose as most suitable to define the extensional aspect of a concept. Another agent may choose a different defining feature set $\Delta'(\alpha)$, even if both agents agree on the definition of α. The set $Ext\,\alpha$ provides a more objective representation of α: if two agents agree on α, they will agree on $Ext\,\alpha$. More generally, the role played by the defining feature set may be extended to a larger set, $Ess\,\alpha$, called the *essence* of α. This set gathers the concepts and features that apply to every element of $Ext\,\alpha$, that is: $Ess\,\alpha = \{\beta \in \mathcal{F}; Ext\,\alpha \subseteq Ext\,\beta\}$. One may consider the elements of this set as the *essential attributes* of the concept. It readily follows from the definitions that the extension of α can be retrieved from its essence: one has indeed $Ext\,\alpha = \bigcap(Ext\,k)_{k\in Ess\,\alpha}$. In the language or Formal Concept Analysis, this shows that the pair $(Ext\,\alpha, Ess\,\alpha)$ can be considered as a formal concept in the context $(\mathcal{O}, \mathcal{F}, I)$, where $(x, \alpha) \in I$ iff

x falls under α.

The notion of *subconcept* can be now precisely defined : a concept β will be said to be a subconcept of α iff $\preceq^{\mu}_{\beta} \subseteq \preceq^{\mu}_{\alpha}$. This implies in particular $Ext\,\beta \subseteq Ext\,\alpha$, and consequently $Ess\,\alpha \subseteq Ess\,\beta$.

Two remarks have to be made concerning the above definitions. First, the *essence* of a concept as defined above corresponds to the classical one (see (Desclés & Pascu, 2011)), provided one remembers that the universe of discourse includes *imaginary*, as well as real, objects. Restricting the set $\mathcal{O}$ to real objects would indeed lead to counterintuitive results: for instance consider the concept *to-be-US-president*, and suppose that it is known that all past and present US presidents were golf players. Then the attribute *to-be-a-golf-player* applies to all instances of *to-be-US-president*; if only real existing objects were considered, the feature *to-be-a-golf-player* would become part of the *essence* of *to-be-US-president*...! Enlarging the set of real objects to imaginary ones avoids this drawback, because we can imagine an US president that does not play golf.

Another remark is that the essence of a concept, as defined above, gathers two different families of concepts. On the one hand, we find the *specific* attributes that are attached to this concept and help distinguishing it from neighboring concepts. In the example of *to-be-US-president*, such is for instance the case of the concepts *to-sleep-in-the White-House*, *to-convene-Congress*, *to-command-the-US-armed-forces* - in the current literature, this set is usually referred to as the 'Intension' of the concept. On the other hand, we have the *generic* attributes, which the concept has inherited from some super-concept: for instance, the generic attributes of the concept *to-be-US-president* include the features *to-be-mortal*, *to-be-a-vertebrate* or *to-have-a-heart*, which are part of the essence of *to-be-a-man*.

The distinction between specific and non specific features plays an important role in category-based induction. It will be more precisely studied in section 6.2.2.

3.3 Membership distance and membership degree

Since $\preceq^{\mu}_{\alpha}$ is generally not a total order, there exists no membership degree function δ_α that would satisfy $x \preceq^{\mu}_{\alpha} y$ iff $\delta_\alpha(x) \leq \delta_\alpha(y)$. It is however possible to approximate the notion of membership degree by considering increasing chains of objects, similarly to what was evoked in the WMD example. The following construction stems from the fact that any strictly increasing chain

of the form $x_1 \prec^\mu_\alpha x_2 \prec^\mu_\alpha \ldots \prec^\mu_\alpha x_k \prec^\mu_\alpha \ldots$ must be of finite bounded length (this because of the finite number of defining features, together with the finiteness condition of paragraph 2.3). Given an object x, it therefore makes sense to consider the *maximal length* of a chain $x \prec^\mu_\alpha x_1 \prec^\mu_\alpha \ldots \prec^\mu_\alpha x_n$ with last term $x_n \in Ext\,\alpha$. This length measures the distance that separates x from $Ext\,\alpha$. It will be referred to as the *membership distance* of x, and denoted by $\mu_\alpha(x)$.

Let now N_α be the length of a maximal $\prec^\mu_\alpha$-chain in $\mathcal{O}$. The membership degree δ^μ_α is defined, for all objects x, by $\delta^\mu_\alpha(x) = 1 - \frac{\mu_\alpha(x)}{N_\alpha}$. One has readily $\delta^\mu_\alpha(x) = 1$ if and only $x \in Ext\,\alpha$, and $\delta^\mu_\alpha(x) = 0$ if and only if x is maximally distant from $Ext\,\alpha$. Note that $\delta^\mu_\alpha(x) < \delta^\mu_\alpha(y)$ whenever $x \prec^\mu_\alpha y$.

Example 2 *In Example 1, the membership order yields* $\mu_\alpha(t) = 1$ *provided that, for the agent, there exists no animal* z *such that* $t \prec^\mu_\alpha z \prec^\mu_\alpha s$. *Similarly* $\mu_\alpha(d) = 1$ *and* $\mu_\alpha(b) = 2$. *Concerning the mouse, the agent may consider that the chain* $m \prec^\mu_\alpha k \prec^\mu_\alpha b \prec^\mu_\alpha d \prec^\mu_\alpha s$ *is a maximal one, where* k *stands for* a monkey, *so that* $\mu_\alpha(m) = 4$. *The maximal length of a chain is* 4, *and the membership degrees are* $\delta^\mu_\alpha(m) = 0$, $\delta^\mu_\alpha(b) = 1/2$, $\delta^\mu_\alpha(t) = \delta^\mu_\alpha(d) = 3/4$) *and* $\delta^\mu_\alpha(s) = 1$.

3.4 Categorial membership for compound concepts

3.4.1 The determination connective

It is sometimes possible to *determine* a concept α by another concept β. We obtain in this way a compound concept, which will be denoted by $\beta \star \alpha$. This determination is most often realized by the combination of an adjective or an adjectived verb with a noun, like in the compositions *to-be-a-carnivorous-animal, to-be-a-flying-bird, to-be-a-french-student, to-be-a-red-apple*. It can also take the form of a noun-noun combination, like in *to-be-a-pet-fish, to-be-a-barnyard-bird*, and, more generally, of a relative clause that will be globally encapsulated by the concept β (e.g. *to-be-an-American*-who-lives-in-Paris). Typically, the concept β becomes a simple feature of the compound concept $\beta \star \alpha$. Its role can be therefore considered as secondary, compared with that played by the principal concept α: *to-be-red* is a feature of the composed concept *to-be-a-red-car*, and *to-be-a-woman* becomes a feature of a *to-be-physician-that-is-a-woman*.

Not to mention the simple conjunction (*to-be-an-history-teacher* & *to-be-a-geography-teacher*), which provides equal importance to both constituents,

other types of connective exist in which the principal role is attributed to the modifier, like in *to-be-a-Picasso-painting*. These will not be studied in the present paper.

As a last remark, it is important to keep in mind that only the *intersective* conceptual combinations are accounted for: the objects that fall under the composed concept $\beta\star\alpha$ are exactly the ones that both fall under α and under β (see (Kamp & Partee, 1995) for the distinction between intersective and non-intersective modifiers). This shows that the determination connective $\star$ is only a *partial* operator: given arbitrary α and β, it may be meaningless to build the concept $\beta \star \alpha$. For instance, there is no sense in talking of a *sailing-number* or a *wooden-salience*. Such pseudo-concepts correspond to nothing, and no object, real or fictitious, can be thought of falling under them, contrary to imaginary concepts like a *pink-elephant*, a *striped-apple* or a *flying-cow*: these latter definitely have a non-empty extension, because we can imagine a pink elephant, a striped apple or a flying cow.

Similarly, concepts determined through qualitative or quantitative adjectives, like a *big-piano*, or a *nice-house* will not be taken in consideration, as one cannot consider that *to-be-big* or *to-be-nice* are well-defined concepts. When forming the composition $\beta \star \alpha$, it is always understood that the intersection $Ext\,\alpha \cap Ext\,\beta$ is a non-empty well-defined set.

3.4.2 Membership orders for compound concepts

Let α and β be two concepts for which corresponding membership orders have been defined, and suppose that the determination of α by β is meaningful, that is $Ext\,\alpha \cap Ext\,\beta \neq \emptyset$. We want to set a membership order on the concept $\beta\star\alpha$ that gives preeminence to α over β. This can be simply done by setting:

$$x \preceq^{\mu}_{\beta\star\alpha} y \text{ if } x \preceq^{\mu}_{\alpha} y \text{ and, either } x \prec^{\mu}_{\alpha} y, \text{ or } x \preceq^{\mu}_{\beta} y.$$

In this framework, the concept *to-be-a-flying-bird* will be considered as applying more to a penguin than to a bat: indeed, the principal concept is that of *being-a-bird*, while *to-fly* appears as a simple feature, less important than the concept it modifies.

The relation thus defined is reflexive and transitive. Since it is clearly a subrelation of $\preceq^{\mu}_{\alpha}$, it makes $\beta \star \alpha$ a subconcept of α.

The hypothesis that $Ext\,\alpha \cap Ext\,\beta \neq \emptyset$ implies that full categorial membership can be recovered through the $\preceq^{\mu}_{\beta\star\alpha}$-maximal elements. If we define indeed the *extension* $Ext\,(\beta\star\alpha)$ of $\beta\star\alpha$ as the set of $\preceq^{\mu}_{\beta\star\alpha}$-maximal elements,

it holds $Ext\,(\beta \star \alpha) = Ext\,\alpha \cap Ext\,\beta$. Categorial membership is therefore *compositional*: an object falls under a determined concept if and only if it falls both under the concept and under its determiner. The category of *red-cars* exactly covers all the items that are red and that are cars.

Compositionality for $\preceq^{\mu}_{\beta\star\alpha}$-maximal elements does not extend to arbitrary items: an object y may fall more than an object x under $\beta \star \alpha$, while falling less than x under β: *to-be-a-flying-bird* for instance applies more to a penguin than to a bat, although *to-fly* applies more to a bat than to a penguin.

3.4.3 Membership degree and the conjunction effect

As was done for simple concepts, a membership distance can be defined for compound concepts of the form $\beta \star \alpha$: for any object x, let $\mu_{\beta\star\alpha}(x)$ denote the length of a maximal $\prec^{\mu}_{\beta\star\alpha}$-chain starting from x and ending in $Ext\,(\beta \star \alpha)$. Note that $\mu_{\alpha}(x) \leq \mu_{\beta\star\alpha}(x)$. If $N_{\beta\star\alpha}$ is the length of a maximal $\prec^{\mu}_{\beta\star\alpha}$-chain in $\mathcal{O}$, the $(\beta\star\alpha)$-membership degree of x may be then defined by $\delta^{\mu}_{\beta\star\alpha}(x) = 1 - \frac{\mu_{\beta\star\alpha}(x)}{N_{\beta\star\alpha}}$.

An interesting side-effect of this construction is that it renders possible a modelling of the so called *conjunction effect*. This phenomenon was observed in 1981 by Osherson and Smith (Osherson & Smith, 1981). It has been since at the origin of numerous research and experiments (see in particular (Hampton, 1988), (Tversky, 1977), (Kamp & Partee, 1995), (Aerts, 2009), (Franco, 2009), or (Hampton, 2017) for a general review. The conjunction effect can be described by the fact that an item may be found to be more strongly a member of the conjunction of two concepts than a member of one of them. Thus, a guppy appears to be more a member of the concept *to-be-a-pet-fish* than a member of *to-be-a-pet* or *to-be-a-fish*. This problem is different from the one evoked at the end of the preceding section: here, one does not compare the membership of two objects relative to a single concept, but the membership of one object relative to two concepts. The use of membership degrees renders possible an explanation of the conjunction effect. It turns out indeed that nothing stands against the fact that the $(\beta\star\alpha)$-membership degree of an item might be greater than its α-membership degree, as can be seen in the following example:

Example 3 *Let us take again the degree computations of example 2. The chain* $m \prec^{\mu}_{\alpha} k \prec^{\mu}_{\alpha} b \prec^{\mu}_{\alpha} d \prec^{\mu}_{\alpha} s$ *being supposed to be maximal, one has*

$\delta^{\mu}_{\alpha}(k) = 1/4$. *If* β *is the feature* to-be-black, *equipped with a two-valued membership degree function, one may think of a maximal* $\beta \star \alpha$*-chain like*

$$m \prec^{\mu}_{\beta\star\alpha} m' \prec^{\mu}_{\beta\star\alpha} k' \prec^{\mu}_{\beta\star\alpha} k \prec^{\mu}_{\beta\star\alpha} b \prec^{\mu}_{\beta\star\alpha} b' \prec^{\mu}_{\beta\star\alpha} d \prec^{\mu}_{\beta\star\alpha} d' \prec^{\mu}_{\beta\star\alpha} s \prec^{\mu}_{\beta\star\alpha} r,$$

with m*: a white mouse;* m'*: a black cat;* k'*: a (red-brown) kangaroo;* k*: a black macaque;* b*: a brown bat;* b'*: a black bat;* d*: a blue dragonfly;* d'*: a black fly;* s*: a sparrow;* r*: a raven. This leads to* $N_{\beta\star\alpha} = 9$*, and* $\delta^{\mu}_{\beta\star\alpha}(k) = 1/3$*, showing that, in this model, the macaque is more a black bird than a bird.*

3.4.4 Objections to the compositional theory

Several arguments have been developed against a theory of compositionality in which membership to a compound concept simply boils down to membership to each of its constituents. It has been objected for instance that there is a difference between the concept $s \star g$ (*being-a-game-that-is-a-sport* and the concept $g \star s$ (*being-a-sport-that-is-a-game*), and that they do not have the same exemplars, although, in the proposed model, the concepts $s \star g$ and $g \star s$ would have the same extension. But even though they share the same exemplars, these concepts are not equal, as they have different membership orders. Moreover, although some experiments seem to contradict this assumption, an item which is considered as *a sport that is a game* must be both a sport and a game - even if it is before all considered as a sport. Confusion comes when one mixes membership and typicality, as is the case of the unary model, where a unique scale is used to measure these two notions. Things would be of course different if we considered a composed concept like 'to-be-a-sport-that-is-*secondarily*-a-game'. Nevertheless, such a complex concept cannot be accounted for through the determination connective $\star$ the way it has been circumscribed.

Another example that would tend to reject compositionality is provided by the composed concept *to-be-school-furniture*. Experiments show that people consider a blackboard as a clear exemplar of this concept, although not an exemplar of the unmodified *to-be-furniture*. The explanation is that here again the compound resulting concept is not really obtained through the determination connective $\star$, the same way *to-be-a-urban-furniture* is not a determination of *to-be-a-furniture*. Indeed, by default, *to-be-a-furniture* refers to *to-be-home-furniture*, and school-furniture is *not* home-furniture that can be found in school. This shows that rather than a determination, the operation that yields *school-furniture* from *furniture* is a *modification*. Clearly

there is no contradiction in the fact that an instance of a modified concept is not an instance of the unmodified concept.

It is interesting to note that the compound concept *to-be-school-furniture* cannot be translated in French through a determination or a modification of the words *meubles* or *ameublement*, which are used for usual (home) furniture. The French word that translates the furniture dedicated to some specific use is *mobilier*: one speaks then of *mobilier d'école* (school-funiture), *mobilier urbain* (urban-furniture), or *mobilier de bureau* (office furniture). The compound concept obtained in this way is fully compositional.

4 Typicality

The relative typicality of objects falling under a concept α is accounted for through an order relation that stems from the *characteristic* set associated with α. This set consists of features that, from the point of view of an agent, are sufficient to characterize the typical exemplars of α. For instance, if α is the concept *to-be-a-bird*, the associated characteristic set $\Xi(\alpha)$ may include the concepts *to-fly*, *to-sing*, *to-live-in-the-trees*. It may also include attributes that are related to the average size, the shape or the weight of the concept typical exemplars, like *to-be-small* or *to-be-light*. On the contrary, concepts like *to-have-feathers* and *to-have-wings*, which apply to all instances of α, will not be part of the characteristic feature set, but they may figure in the defining feature set $\Delta(\alpha)$.

As was the case for the defining feature set, the characteristic set is supposed to be equipped with a partial *salience order*, denoted $>_{\Xi(\alpha)}$. This order is meant to compare the relative importance of the different characteristic features.

4.1 The typicality order

Let α be a featured concept. For two elements x and y of $Ext\,\alpha$, we set

$x \preceq^{\tau}_{\alpha} y$ iff for each characteristic feature f of α such that $y \prec_f x$, there exists a characteristic feature g, $g >_{\Xi(\alpha)} f$, such that $x \prec_g y$.

It is easily seen that this relation is a (weak) partial order on $Ext\,\alpha$. The set of $\preceq^{\tau}_{\alpha}$-maximal elements of $Ext\,\alpha$ will be denoted by $Typ\,\alpha$. It gathers the *typical exemplars* of the concept. Clearly, an exemplar z of α is α-typical

if and only if it falls under all the characteristic features of α. It follows that $Typ\,\alpha = \bigcap(Ext\,k)_{k\in\Sigma(\alpha)}$.

It is possible to parallel the construction proposed in paragraph 3.2, and define the *intension* of the concept α as the set of concepts and features that apply to every element of $Typ\,\alpha$. One has therefore $Int\,\alpha = \{\beta \in \mathcal{F}; Typ\,\alpha \subseteq Ext\,\beta\}$. This set can be substituted to the (subjective) characteristic set used for the definition of $Typ\,\alpha$. Conversely, it holds $Typ\,\alpha = \bigcap(Ext\,k)_{k\in Int\,\alpha}$. This shows that $(Typ\,\alpha, Int\,\alpha)$ is a formal subconcept of $(Ext\,\alpha, Ess\,\alpha)$.

The intension of a concept includes its essence; the elements of $Int\,\alpha$ that are not in $Ess\,\alpha$ are the *typical attributes* of the concept: they consist of the attributes that are *generally* but not always true of the exemplars of the concept. They apply to every typical instance of the concept, without applying to all of its exemplars.

Contrary to categorial membership, typicality is not preserved by embedding: typical ostriches are not typical birds. A subconcept β of α will be said to be *smooth* if its typical instances are typical for α, that is if $Typ\,\beta \subseteq Typ\,\alpha$. The concept *to-be-a-robin* for instance may be seen as a smooth subconcept of *to-be-a-bird*.

Membership and typicality orders fully determine a concept as they describe how the universe of discourse is structured relative to this concept. Two concepts α and β will be considered are *equivalent*, written $\alpha \equiv \beta$, if $\preceq^{\mu}_{\alpha} = \preceq^{\mu}_{\beta}$ and $\preceq^{\tau}_{\alpha} = \preceq^{\tau}_{\beta}$.

Note that two concepts may have same extension and same typical elements without being equivalent. For instance, such will be the case if their respective defining set and characteristic sets are identical, but equipped with different corresponding salience orders. The concepts α and β will be said to be *similar*, written $\alpha \simeq \beta$, if they have same extension and same set of typical elements

4.2 The case of compound concepts

Let $\beta\star\alpha$ be the determination of α by a concept β, and define the typicality order $\preceq^{\tau}_{\beta\star\alpha}$ on the set $Ext\,\alpha \cap Ext\,\beta$ by:

$$x \preceq^{\tau}_{\beta\star\alpha} y \text{ if } x \preceq^{\tau}_{\alpha} y \text{ and either } x \prec^{\tau}_{\alpha} y, \text{ or } x \preceq^{\tau}_{\beta} y.$$

This relation is reflexive and transitive. Because of the finiteness of $\Xi(\alpha)$, the set of $\preceq^{\tau}_{\beta\star\alpha}$-maximal elements of $Ext\,(\beta\star\alpha)$ is not empty. This set will

be denoted by $Typ(\beta\star\alpha)$, and its elements referred to as the *typcal instances* of $\beta\star\alpha$.

As was the case for membership, the typicality order for composed concept gives preeminence to the initial concept: a typical exemplar of $\beta\star\alpha$ has maximal α-typicality among the elements of $Ext\,(\beta\star\alpha)$. Taking this side however leads to seemingly paradoxical results. It has been objected for instance that this construction makes a gull a more typical exemplar of *Antarctic-bird* than a penguin, while, for most people, the penguin appears to be *the* typical Antarctic bird, The explanation is that there is a mix-up between a *typical Antarctic bird* and a *bird that typically lives in Antarctic*, this latter being usually interpreted as a *bird that mainly lives in Antarctic*. Such is indeed the case for the penguin, which, contrary to the gull, is an Antarctic *endemic* bird. Moreover, the penguin's features are so different from those of a familiar European birds that they may appear as 'typical' of this atypical species. Typicality indeed is often understood as a differentiation tool: the most typical exemplars of an atypical subcategory tend to be chosen among the less typical exemplars of this category... This interpretation however does not correspond to the usual definition of typicality by means of characteristic features. It covers a different notion, that accounts for representativeness rather than typicality. To precisely define this notion, it would be necessary to introduce a *representativeness* order inside a subcategory: for instance, given two exemplars x, y of a subconcept β of α, one may say that, relatively to α, *y is a better representative of β than x* if y is both more β-typical and less α-typical than x. It is this order relation, different from typicality, that would render the penguin more representative as an Antarctic bird than the skua or the petrel.

Contrary to membership, typicality is *not* compositional. The typical instances of $\beta\star\alpha$ cannot be retrieved from the typical instances of α and β. This comes from the fact that the set $Typ\,\alpha\cap Typ\,\beta$ may be empty, like in the example (*to-be-an-ostrich*)$\star$(*to-be-a-bird*). However it is not difficult to see that if the set $Typ\,\beta\cap Typ\,\alpha$ is not empty, one has $Typ\,\beta\cap Typ\,\alpha = Typ\,(\beta\star\alpha)$: a typical black olive is a typical olive that is typically black.

Concerning the determination connective, two interesting properties deserve to be mentioned. The first one is *associativity*: for any concepts α,β,γ, one has

$$\gamma\star(\beta\star\alpha)\equiv(\gamma\star\beta)\star\alpha.$$

Indeed, a straightforward computation shows that the membership and

the typicality orders of $\gamma \star (\beta \star \alpha)$ and $(\gamma \star \beta) \star \alpha$ are identical.

The second property is that of *idempotence*: for any concept α, it holds:

$$\alpha \star \alpha \equiv \alpha.$$

There is no difference between *to-be-a-bird-that-is-a-bird* and *to-be-a-bird*.

5 On resemblance

The notion of resemblance is an important one in the categorization process. It has been at the center of numerous theoretical and experimental studies. Searchers tried to precisely circumscribe this notion, either by defining a resemblance degree between two items, or by determining the link between resemblance, membership and typicality (Rosch, 1975), (Tversky & Kahneman, 1983) and (Wittgenstein, 1953).

Resemblance may be considered as a binary relation between objects (*Henry resembles his brother*), between concepts (*the wolf resembles a dog*), or between an object and a concept (*this picture resembles a Picasso*). To study this notion first requires to determine which type of resemblance one is dealing with. In this paper, only resemblance between objects and concepts or between concepts will be studied.

5.1 Resemblance between objects and concepts

In a first approximation, resemblance between an object and a concept seems to be directly linked with categorial membership : saying that *this piece of music resembles Beethoven* means that what is heard is close of falling under the concept *to-be-Beethoven's work*. However, it is clear that resemblance relative to a concept is first perceived as resemblance with the *typical instances* of this concept. The *Beethoven* referred to in the preceding example is not the 'young Beethoven', whose compositions still reflect Haydn influence, but the later Beethoven of the second or third period. Similarly, saying that 'Peter's bedroom looks like a boat cabin' refers to a typical boat cabin, excluding for instance a destroyer's cabin...

It follows that, when an agent asserts that a particular item x resembles a concept α, one may infer, first, that x is not known by the agent to be an *instance* of α, and, secondly that, for this agent, x resembles a *typical* exemplar of this concept, sharing with it a certain amount of *typical attributes*.

Thus, looking at a bat, one may say it resembles a bird, just because it has wings, it flies, and it has the size or the shape of a bird. Conversely, an animal may be declared *not* to resemble a bird if it does not resemble a *typical* bird, even though this animal is known to be a bird. For instance, looking at a penguin, an assertion like 'this animal does not resemble a bird' is perfectly understandable. Resemblance first deals with the typical attributes of a concept.

In the framework of featured concepts, it seems that resemblance to a concept may be analyzed and treated through a preorder relation in a way similar to what was done for membership and typicality.

The simplest way to do so would be to define on the set $\Delta(\alpha) \cup \Xi(\alpha)$ a salience order stemming from the salience orders of $\Delta(\alpha)$ and $\Xi(\alpha)$, and then build on $\mathcal{O}$ a weak order relation, as was done for membership or typicality. However, in such a model, resemblance would be very much dependent on salience, leading to counterintuitive results: a single α-feature with maximal salience, that applied to x and not to y, would make x more α-resemblant than y, even though many other α-features may apply to y and not to x. To take an example, suppose that α is the concept *to-be-a-bird* and that *to-fly* has maximal salience in $\Delta(\alpha) \cup \Xi(\alpha)$. Then bird-resemblance will principally rest on the ability of an item to fly; consequently, bats will be more bird-resemblant than kiwis, although *having feathers*, *singing* and *building nests*, taken together, should, at least, compensate the fact that kiwis do not fly.

The following construction seems to remedy this drawback, as it takes into account the number of α-attributes that apply to an object x as well as the preeminence of the characteristic features over the defining ones.

Denote by $\Pi(\alpha)$ the set $\Delta(\alpha) \cup \Xi(\alpha)$ - this set is referred to as the *stereotypical* set of α (see (Connolly, Fodor, Gleitman, & Gleitman, 2007), (Fodor, 1994), or (Jönsson & Hampton, 2007)).

Recalling that $\Delta(\alpha)$ and $\Xi(\alpha)$ are supposed to be disjoint sets, let $>_{\Pi(\alpha)}$ be the salience order on $\Pi(\alpha)$ that extends the salience orders $>_{\Delta(\alpha)}$ and $>_{\Xi(\alpha)}$, and satisfies moreover $f >_{\Pi(\alpha)} g$ for all $f \in \Xi(\alpha)$ and $g \in \Delta(\alpha)$. The study of the ordered set $(\Pi(\alpha), >_{\Pi(\alpha)})$ may be seen as providing a bridge between the binary model and the unary model.

Given an element f of $\Pi(\alpha)$, define the *salience degree* of f as the number $s_f = 1 + |\{h \in \Pi(\alpha); f >_{\Pi(\alpha)} h\}|$.

Using the f-applicability function $\delta_f(x)$ defined in 2.3 renders possible the definition of the α-*resemblance degree* of an element x of $\mathcal{O}$ by setting:

$$\delta^{\rho}_{\alpha}(x) = \frac{\Sigma_{f\in\ \Pi(\alpha)} s_f \delta_f(x)}{\Sigma_{f\in\Pi(\alpha)} s_f}.$$

The elements of $\mathcal{O}$ that are maximally resemblant to a concept α are the typical instances of α. Their resemblance degree is equal to 1.

Example 4 *Taking again the bird example, suppose that the defining feature set is as in example 1, that is* $\Delta(\alpha)$ =$\{$to-have-two-legs, to-lay-eggs, to-have-a-beak, to-have-wings$\}$ *with salience order* to-have-a-beak $>_s$ to-lay-eggs $>_s$ to-have-two-legs, *and* to-have wings $>_s$ to-lay-eggs$>_s$ to-have-two-legs. *Suppose that its characteristic set consists of the three features* to-build-nests, to-sing, *and* to-fly, *equipped with the order* to-fly$>_s$to-build-nests *and* to-fly$>_s$to-sing. *Computing the bird-resemblance degrees of a bat b, a penguin p and a kiwi k yields* $\delta^{\rho}_{\alpha}(b) = 11/26$, $\delta^{\rho}_{\alpha}(p) = 9/26$ *and* $\delta^{\rho}_{\alpha}(k) = 19/26$. *The bat is less bird-resemblant than the kiwi, but more bird-resemblant than the penguin.*

Note that, in this model, categorial membership cannot be directly retrieved through resemblance: there exists no degree threshold above which an item can be declared to be an exemplar of the target concept. All what can be said, is that an item with α resemblance degree equal to 1 must be a typical instance of α. This comes from the fact that, in evaluating resemblance, one compares objects with the *typical instances* of a concept, rather than to arbitrary exemplars of this concept.

5.2 Resemblance between concepts

One may again interpret concept resemblance as resemblance between typical instances : judging that 'Braque resembles Picasso' amounts to say that a typical painting of Braque resembles a Picasso. In other words, one will say that *a concept* α *resembles a concept* β if *every typical instance* of α resembles to the concept β in the sense of the preceding paragraph.

This observation leads to a possible construction of the β-*resemblance degree* of a concept α, which can be defined as

$$\delta^{\rho}_{\beta}(\alpha) = Min_{x\in Typ\,\alpha}\delta^{\rho}_{\beta}(x).$$

The resemblance between α and β is therefore measured by taking among the typical elements of α those that are the least β-resemblant.

Following this definition, the maximally β-resemblant concepts are the concepts α such that $Typ\,\alpha \subseteq Typ\,\beta$. In particular, such will be the case

for the *smooth subconcepts* of β. In general, however, α may have maximal β-resemblance without being a subconcept of β. For instance, although not all *Caravage paintings* are *figure paintings*, we can consider that *Caravage paintings* maximally resemble *figure representations*, because typical paintings from Caravage are typical figure representations.

6 Category-based induction

Category-based induction is the process through which, from a given category, one infers information relative to another category, this inference being based on a link between the source and the target category. In particular, one may ask which attributes of a category remain true for a subcategory or, inversely, which attributes of a subcategory may be raised to the whole category. One may also be interested in resemblance between categories, and look for the attributes that transpose from one category to a resemblant one. These problems can be treated with the tools developed in the preceding sections.

6.1 Within-category induction

While the essential attributes of a concept apply to the exemplars of any of its subconcepts, this is no more true for its typical attributes. Typical birds fly, but typical ostriches don't. However, typical birds have feathers with remiges, and such is the case for ostriches. This poses the problem of determining which attributes of a concept remain valid for some subconcept. Knowing for instance that ducks have webbed feet, can we deduce that quacking ducks have webbed feet (Connolly et al., 2007)? The inheritance problem is of particular interest in the case of subconcepts obtained through concept determination. To quote (Jönsson & Hampton, 2012) 'Given that an attribute is not universally true of a concept (...), how should one determine whether the predicate should also be considered generically true of the complex concept formed when an adjectival or nominal modifier is applied to the noun'. As will be seen, the answer depends on the *smoothness* of the target subconcept (recall that a subconcept β of α is said to be *smooth* in α if $Typ\,\beta \subseteq Typ\,\alpha$).

6.1.1 Smooth subconcepts and non-exceptional modifiers

A determiner β will be said to be *non-exceptional* for α if $Ext\,\beta \cap Typ\,\alpha \neq \emptyset$, that is if there exists at least an exemplar of β that is a typical instance of α. *To-be-red* for instance is non-exceptional for the concept *to-be-an-apple*, while *to-be-blue* would be an exceptional determiner (in fact, a modifier). This notion of non-exceptional determiner may be seen as a formal definition of the *compatible modifiers* introduced by (Smith, Osherson, Rips, & Keane, 1988). It is tightly linked with that of smoothness: indeed, it can be shown (Freund, 2008) that $\beta \star \alpha$ is a smooth subconcept of α if and only if β is non exceptional for α. Conversely, it is not difficult to prove that, up to similarity, the concepts obtained from the determination of α by a non exceptional attribute β generate *all* the smooth subconcepts of α.

If β is a non-exceptional modifier of α, $\beta \star \alpha$ is smooth in α, and we have therefore $Typ\,(\beta \star \alpha) \subseteq Typ\,\alpha$, showing that every typical attribute of α applies to the typical instances of $\beta \star \alpha$. For instance, knowing that there exist typical birds that are white and that birds generally sing, we can conclude that white birds generally sing.

This provides a first answer to the within-category induction problem:

if β is non-exceptional for α, inheritance in $\beta \star \alpha$ is total.

Note as a particular case, that the above result holds if β is itself a typical attribute of α: *flying-birds* will inherit all the typical properties of birds.

6.1.2 Comparison with experimental results: the modifier effect

It has been observed that whatever type of modifier is used, likelihood for an attribute tends to be reduced after modification of a principal concept (Jönsson & Hampton, 2012)). Thus, the authors wrote that 'Sentences were assigned systematically decreasing likelihood ratings as the head noun was modified by a single typical modifier (eg., *flightless penguins*), by a single atypical modifier (*solitary penguins*, or by two atypical modifiers (*solitary migrant penguins*)'.

As far as non-exceptional modifiers are concerned, this definitely contradicts the property evoked above, and therefore questions the proposed model. However, the experiments did not concern simple truth of statements, like 'do quacking ducks have webbed feet': using numbers from 1 to 10, the participants had to indicate the likely truth of a sentence (Ducks have webbed feet, quacking ducks have webbed feet...). It would be interesting to see what

answers of the type Y/N would have been given successively to sentences first unmodified, then typically modified, then atypically modified. At any rate, the disparity of some answers with the above result would show that a weak but maybe significant percentage of the participants do not follow this type of model in their categorization process.

6.1.3 The case of exceptional modifiers

When β is exceptional for α, $\beta \star \alpha$ is no longer smooth in α and the full inheritance property does not apply anymore. Non-flying birds do not inherit the bird typical property of flying. The problem is to determine which attributes of α are preserved in $\beta \star \alpha$. For this it is useful to determine what makes β exceptional for α.

Suppose that α is a featured concept for which β is exceptional. Since $Typ\,\alpha = Ext\,\alpha \cap (\bigcap(Ext\,k)_{k\in\Xi(\alpha)})$, the set $Ext\,\beta \cap Ext\,\alpha \cap (\bigcap(Ext\,k)_{k\in\Xi(\alpha)})$ is empty.

The simplest case where this situation occurs happens when, for a single α-characteristic feature k_0, one has $Ext\,\beta \cap Ext\,\alpha \cap Ext\,k_0 = \emptyset$, while $Ext\,\beta \cap Ext\,\alpha \cap (\bigcap(Ext\,k)_{k\in\Xi(\alpha)-\{k_0\}}) \neq \emptyset$. No element of $Typ\,(\beta\star\alpha)$ (fully) falls under k_0. Suppose more precisely that the k_0-applicability degree of each element of $Typ\,(\beta \star \alpha)$ is equal to 0. Then one observes that **any characteristic feature of α different from k_0 remains a typical attribute of $\beta \star \alpha$.** Indeed, let h be such a feature and suppose by *reductio ad absurdum* that there exists an element z of $Typ\,(\beta \star \alpha)$ that does not fall under h. Choose an arbitrary element y of $Ext\,\beta \cap Ext\,\alpha \cap (\bigcap(Ext\,k)_{k\in\Xi(\alpha)-\{k_0\}})$. Then $z \prec_h y$ and $z \preceq_k y$ for any characteristic feature k of α, $k \neq k_0$. From this follows that $z \prec^{\tau}_{\alpha} y$. But this contradicts the fact that z is a $\preceq^{\tau}_{\beta\star\alpha}$-maximal element of $Ext\,(\beta \star \alpha)$. This shows that h applies to $Typ\,(\beta \star \alpha)$, and h is therefore a typical attribute of $\beta \star \alpha$.

By this result, we see that inheritance to an exceptional subcategory still holds for the characteristic features that do not directly contradict the determiner. If α is the concept *to-be-a-bird*, with characteristic feature set $\Xi(\alpha) = \{$to-fly, to-build-nests, to sing to-eat-seeds$\}$, we can conclude that, typically, walking birds sing, walking birds build nests and walking birds eat seeds.

6.2 Over-category induction

Over-category induction is the process through which a property that is known to hold for a subcategory can be raised to a whole category. For instance (Connolly et al., 2007), knowing that setters are susceptible to dysplasia, can one conclude that all dogs are susceptible to dysplasia ? This problem may be seen as the converse of that of within-category induction. However, it slightly differs in that one has to distinguish between *essential* and *characteristic* attributes. The essential attributes pose no problem in within-category induction, because any essential attribute of a concept remains an essential attribute of its subconcepts (recall (see 3.2) that the essential attributes of a concept are those that apply to every exemplar of the concept). On the contrary, an essential attribute of a subcategory need not hold for the whole category : ducks have webbed feet, but this is not generally true for birds. The problem of over-category induction is therefore the following one: what are the essential attributes of a subcategory that *generically* apply to the whole category ?

6.2.1 The case of compound concepts

Over-category induction has an evident solution when the considered subconcept is obtained by determining the basic one through one of its typical attributes. In this case indeed **if β is a typical attribute of α, every essential attribute of $\beta \star \alpha$ becomes a typical attribute of α**. This comes from the fact that one has $Typ\,\alpha \subseteq Ext\,\beta$, hence $Typ\,\alpha \subseteq (Ext\,\beta \cap Ext\,\alpha) = Ext\,(\beta \star \alpha)$. Any essential attribute of $\beta \star \alpha$ therefore applies to $Typ\,\alpha$.

From the fact that birds generally fly and that all flying-birds have feathers with remiges, we can thus conclude that birds generally have feathers with remiges.

6.2.2 Typical subconcepts

Let us consider again the concept *to-be-a-bird*. When an agent says that the robin is a *typical* exemplar of this concept, she rests her assertion on the fact that robins inherit all the attributes that a typical bird should have: to fly, to sing, to live in the trees, etc. This inheritance property then extends from individual items to the whole *category* of robins, thus becoming an essential property of the *concept* 'to-be-a-robin'. Each exemplar of *to-be-a-robin* is a typical instance of *to-be-a-bird*. This leads to the following notion: a concept

β is a *typical subconcept* of α if each exemplar of β is a typical instance of α, that is if $Ext\,\beta \subseteq Typ\,\alpha$. Note that such a concept is necessarily smooth in α.

Let now β and γ be two typical subconcepts of α, and h an essential feature of β. Suppose that over-category induction holds for h, so that h becomes a typical attribute of α. Then, since $Ext\,\gamma \subseteq Typ\,\alpha$, h must be also an essential feature of γ. This shows that *a feature of β can extend to a typical attribute of α only if it is an essential feature of* all *typical subconcepts of α*. In other words, this feature must not be *specific* to β.

The distinction between specific and non-specific features intervened during the experiments conducted on category-based induction (Connolly et al., 2007). The predicates that were the best candidates for induction to the mother category were the so-called *blank* predicates. For the participants, these predicates did not bear any special meaning, they were liable to apply to any category that was considered close enough to the tested one. On the contrary, features that were tightly related or specific to the source category had to be discarded.

To make this notion precise, let us say that an essential attribute is *non-specific* for a typical subconcept β of α if it is shared by all the typical subconcepts of α. For instance, *to-have-webbed-feet* is a non specific feature of ducks as a subcategory of aquatic-birds.

This distinction provides an answer to the over-category induction problem, at least in the case where α is a featured concept. Indeed, let β be a typical subconcept of α and k an essential feature of β. If k remains a typical attribute for α, we have seen that k is non-specific. Conversely, suppose that k is a non-specific feature of β. Denote by $f_1, f_2, \ldots f_n$ the characteristic feature features of α, and let γ be the 'pseudo-concept' $\gamma = f_1 \star f_2 \star \ldots \star f_n \star \alpha$. We may see γ as the concept *to-be-a-typical-α* that parallels the notion of *typical object* introduced by Desclés (Desclés & Pascu, 2011). Note that $Ext\,\gamma = Ext\,\alpha \cap Typ\,\alpha$, that is $Ext\,\gamma = Typ\,\alpha$, showing that γ is a typical subconcept of α. This implies now that k becomes an essential feature of γ, and therefore applies to $Ext\,\gamma$, hence to $Typ\,\alpha$. This shows that h is a typical attribute of α.

The above result shows that **an essential feature of β extends to a typical attribute of α if and only if it is non-specific to β.** For instance, *to-have-webbed feet* applies to the category of aquatic birds, but not to the whole category of birds.

6.3 Induction through resemblance

When a concept α is sufficiently close to a concept β, one expects that some attributes that hold for the typical instances of β will also hold for the typical attributes of α: for example some attributional properties should transpose from bats to birds, from horses to cows or from camels to dromedaries.

If α is maximally β-resemblant, it follows from section 5.1 that one has $Typ\,\alpha \subseteq Typ\,\beta$: any typical attribute of β therefore typically applies to α. Apart from this trivial case, one may wonder if a similar answer may hold in the case where the β-resemblance degree of α, although not maximal, is very close to 1.

Recall that the β resemblance degree of α is $\delta^{\rho}_{\beta}(\alpha) = Min_{x \in Typ\,\alpha}\, \delta^{\rho}_{\beta}(x)$, where $\delta^{\rho}_{\beta}(x) = \frac{\Sigma_{f \in\ \Pi(\beta)} s_f \delta_f(x)}{\Sigma_{f \in \Pi(\alpha)} s_f}$.

Suppose that $\delta^{\rho}_{\beta}(\alpha) \geq 1 - \frac{|\Delta(\beta)|}{\Sigma_{f \in \Pi(\beta)} s_f}$. Then, **given a typical instance x of α, one has $\delta_f(x) \neq 0$ for every characteristic feature of β**. Indeed, if $\delta_h(x) = 0$ for some element h of $\Xi(\beta)$, this would imply $\delta^{\rho}_{\beta}(x) < 1 - \frac{|\Delta(\beta)|}{\Sigma_{f \in \Pi(\beta)} s_f}$, because the salience degree of h, as a characteristic feature of β, is greater than $|\Delta(\beta)|$.

In the particular case where the applicability of the characteristic features of β are measured by a two-valued degree function, this result means that **if α sufficiently resembles to β, every characteristic feature of β becomes a typical attribute of α**.

This result shows that Induction through resemblance is accounted for by the proposed model, as was the case for within and over-category induction.

7 Conclusion

Order relations seem to provide an economical as well as a powerful tool for the study of categorization. For the class of featured concepts, they accurately model the way a given agent may understand a concept. They constitute an interesting alternative that explains and foresees classical problems linked with concept combination or category-based induction. It is to be hoped that experimental studies will confirm the adequacy between this model and the reality.

Acknowledgement

Thanks are due to Daniel Lehmann and James A. Hampton for their remarks and suggestions after reading a preliminary version of this paper.

References

Aerts, D. (2009). Quantum structure in cognition. *Journal of Mathematical Psychology*(53-5), 314-348.

Connolly, A., Fodor, J., Gleitman, L., & Gleitman, H. (2007). Why stereotypes don't even make good defaults. *Cognition*(103-1), 1-22.

Desclés, J.P., & Pascu, A. (2011). Logic of determination of objects: How to articulate 'extension' with 'intension' and 'objects' with 'concepts'. *Logica Universalis*(5).

Fodor, J.A. (1994). Concepts: a potboiler. *Cognition*(50), 95-113.

Fodor, J.A. (1998). *Concepts: where cognitive science went wrong*. Oxford University Press.

Franco, R. (2009). The conjunction fallacy and interference effects. *Journal of Mathematical Psychology*(53-5), 415-422.

Frege, G. (1879). *Begriffsschrift, eine der arithmetischen nachgebildete formelsprache des reinen denkens*. Halle.

Freund, M. (2008). On the notion of concept I. *Artificial Intelligence*(172), 570-590.

Freund, M. (2009). On the notion of concept II. *Artificial Intelligence*(173), 167-179.

Freund, M. (2014). On categorial membership. *Erkenntnis*, *79*(3), 1045-1068.

Gärdenfors, P. (2000). *Conceptual spaces: The geometry of thought*. MIT Press.

Hampton, J.A. (1988). Overextension of conjunctive effects: Evidence for a unitary model of concept typicality and class inclusion. *Journal for Experimental Psychology : Learning, Memory, and Cognition*(14), 12-32.

Hampton, J.A. (2017). *Compositionality and concepts.*

Jönsson, M.L., & Hampton, J.A. (2007). On prototypes as defaults (comment on connolly, fodor, gleitman and gleitman, 2007). *Cognition*(In press).

Jönsson, M.L., & Hampton, J.A. (2012). The modifier effect in within-category induction: Default inheritance in complex noun phrases. *Language and Cognitive Processes*(27-1), 90-116.

Kamp, H., & Partee, B. (1995). Prototype theory and compositionality. *Cognition*(57), 129-191.

Nardi, D., & Brachman, R. (2003). An introduction to description logics. In F. Baader (Ed.), *The description logic handbook* (p. 1-44). Cambridge University Press.

Osherson, D., & Smith, E. (1981). On the adequacy of prototype theory as a theory of concepts. *Cognition*(11), 237-262.

Rosch, E. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology*(104), 192-233.

Smith, E., & Medin, D. (1981). *Categories and concepts*. Cambridge: Harvard University Press.

Smith, E., Osherson, D., Rips, L., & Keane. (1988). Combining prototypes: a selective modification model. *Cognitive Science*(152), 485-527.

Smith, E., Shoben, E., & Rips, L. (1974). Structure and process in semantic memory: a featural model for semantic decisions. *Psychological Review*(81), 214-241.

Tversky, A. (1977). Features of similarity. *Psychological review*(84), 327-352.

Tversky, A., & Kahneman, D. (1983). Extension versus intuitive reasoning: The conjunction fallacy in probability judgement. *Psychological Review*(90), 141-168.

Wittgenstein, L. (1953). *Philosophical investigations*. New York, Macmillan.

Zadeh, L. (1965). Fuzzy sets. *Information and Control*(8), 338-353.