Lexical Pragmatics and Memory Traces*

Catalina Urquiza

Abstract

In recent relevance-theoretic work on lexical pragmatics, it is claimed that cases of lexical narrowing and broadening are the different outputs of a single inferential pragmatic process which results in an ad hoc concept. The input to the process of lexical adjustment is described as "encoded concepts, contextual information and pragmatic expectations or principles" (Wilson & Carston, 2007: 230). In this paper, I explore a fourth possible source of input to this inferential process: memory traces of previous uses of a word. Memory traces include the linguistic form, the contexts in which it has appeared before, and the inferences that pragmatic expectations on those occurrences legitimated. First, I argue that children's acquisition of word meanings benefits from repeated encounters and that this lends support to a multiple-trace memory model. Second, I contest the assumption that, once acquisition is complete, traces are disregarded and all that needs to be stored is a single, reliable, encoded meaning. Rather, utterance traces continue to accumulate in memory and to have an effect on adult interpretation processes.

Keywords: Relevance Theory, lexical pragmatics, memory traces, ad hoc concepts.

1 Introduction

The study of human communication has progressively moved away from what has often been called the code model. It seems widely accepted today that even if a speaker and a hearer share the same language (English, for instance), actually *interpreting* an utterance requires cognitive capacities that go well beyond linguistic decoding. Furthermore, since Grice's 'logic of conversation,' it has been compellingly argued that human communication is not even essentially the decoding of signals but rather the expression and recognition of intentions (see Grice 1967; Sperber & Wilson 1986/95; Wilson & Sperber 2004). Grice demonstrated that by adopting a rational stance to the behaviour of others, one individual can *infer* the intentions (specifically, the *m-intentions*) of another and that *this*, rather than the code, is what communication hinges on[1]. Relevance theory (henceforth RT) broadly adopts and elaborates on this *inferential* model of communication. The RT contribution, among other things, is to redefine the fundamental notions of the interpretive process by placing it within the framework of cognitive science. The cognitive perspective critically reveals the central role played by *relevance*, defined as a potential property of both any external stimulus and any internal representation that provides input to cognitive processes (Wilson and Sperber 2004: 608). This in turn allows RT to reveal an important characteristic of the interpretive process: it is triggered by the inherent presumption of relevance of communicative input. More recently, Wilson and Carston (2007) show that the approach to utterance interpretation developed by RT (Sperber & Wilson 1986/95) applies to words and phrases just as much as sentences. In their 'unitary account' of lexical pragmatics, adjustments to word meaning are presented as the outputs of an inferential interpretive

* Many thanks to Professor Robyn Carston for very helpful discussions and comments on earlier drafts of this paper.

[1] For a thorough presentation of Grice's pragmatics, see Levinson (1983), Sperber and Wilson (1986/95).

process that, as with communication in general, hinges on speakers' communicative intentions and hearers' expectations of relevance.

The proposal I present in this paper largely follows from this perspective. I address the issue of relevance as a potential property of a particular type of input: memory of previous experiences with language. 'Memory traces' of previous utterances include the linguistic form, the contexts in which they have appeared, and the inferences that pragmatic expectations of those occurrences legitimated. The interpretational process both creates new ad hoc concepts and benefits from previous occurrences. Consequently, hearers' expectations of relevance can be taken as orienting the inferential process both as episodic memory stores the construction of an ad hoc concept and as it activates memories of previous uses according to their relevance in the present context[2]. First, I briefly outline the key notions of the relevance-theoretic approach. Following this, I focus on how this inferential approach extends to lexical pragmatics. Then, I introduce the notion of memory traces through a presentation of Douglas Hintzman's multiple-trace memory model. A first compatibility between RT and a multiple-trace memory model is identified: both models activate input on the basis of its relevance to a current context and operate in loops until expectations are met or abandoned. Next, I illustrate how children might employ memory traces in order to learn the meanings of some, if not most, of their words. After that, I address some possible objections to a view of memory traces as affecting adult language use and suggest other sources of supporting evidence. In the last section, I illustrate how memory traces could accumulate and play an important role in later interpretational processes.

2 The pragmatics of word meaning

2.1 Relevance Theory

There are two basic principles of relevance. First, the ***cognitive*** *principle of relevance*: *cognition tends to be geared to the maximisation of relevance*. In other words, our minds permanently conduct a search for the most relevant input; furthermore, because this feature of the mind is taken as an outcome of evolution, the pursuit of relevance is not simply a maxim or norm that we could deliberately or mistakenly opt out of.[3] It is inherently a potential property of all external and internal stimuli.[4] Lastly, it can be defined as a positive function of its beneficial cognitive effects and a negative function of the processing effort required to interpret it: *other things being equal, the greater the positive cognitive effects achieved by processing an input, the greater the relevance of the input to the individual at that time*; and, *other things being equal, the greater the processing effort expended, the lower the relevance of the input to the individual at that time* (Wilson & Sperber 2004: 609).

A theory of communication naturally follows from the above first principle of relevance: assuming that the bid to maximise relevance is an unavoidable part of our information processing and that it is universal in our species, we can rely on it not only to arrive at the most relevant interpretation of some stimuli for ourselves, but *to read the minds*

[2] Episodic memory is responsible for personal, autobiographical memories of specific events, it is therefore this memory system that captures the details associated with a learning episode (see Neath & Surprenant, 2003: 153 – 154).

[3] The cognitive perspective implies that whatever mechanisms pragmatics proposes, they should clearly fit in with how human cognition has evolved (see Wilson & Sperber 2004, Scott-Phillips 2010).

[4] As a result, the objects of research in RT are no longer restricted to utterances, as they might be outside of RT, but could include, for instance, thoughts, and most importantly for my purposes, *memories*.

of others when they overtly direct our attention to a stimulus they are producing. RT characterises this as the interaction of *ostensive stimuli* and *optimal relevance*. The second, ***communicative*** *principle of relevance* states that *every ostensive stimulus conveys a presumption of its own optimal relevance*. An *ostensive stimulus* is defined as a behaviour which comes with two levels of intention: *the intention to inform the audience of something* and *the intention to inform the audience of one's informative intention*. These are called the *informative* and the *communicative* intention respectively. The *presumption of relevance* implies that *the ostensive stimulus is relevant enough to be worth the audience's processing effort* and *that it is the most relevant one compatible with the communicator's abilities and preferences* (Wilson & Sperber 2004: 613).

The above principles can be elaborated in the form of a comprehension procedure which licenses addressees to form interpretive hypotheses in order of accessibility and to accept the first one that satisfies their specific expectations of relevance (see Wilson & Sperber 2004 for detailed examples of this procedure in practice). Roughly, communicative behaviour triggers a comprehension process that constructs an appropriate hypothesis about the explicit and implicit content of an utterance. Relevance theory correctly points out that it is not only implicit content but also explicit content that requires pragmatic enrichment processes (see Carston 2002a, 2002b, and 2009). For this reason, constructing hypotheses about the explicit content of the utterance is not independent from constructing hypotheses about the intended contextual assumptions or the intended implications. The different inferential activities of the on-line comprehension procedure feed into each other and progress in parallel. In fact, the often considerable gap that the hearer must bridge between the 'literal' meaning coded linguistically by the utterance and the meaning that is actually intended by the speaker often requires loops of decoding, assumption building, pragmatic enrichment and deriving of possible implicatures (Wilson & Sperber 2004).

Embedded in this inferential account of the overall comprehension procedure is another "genuinely inferential" account concerning the pragmatic processes that go into *lexical* adjustment (Wilson & Carston 2007). According to RT, just as the linguistically encoded sentence meaning underdetermines the intended meaning of the utterance, the encoded linguistic meaning of individual words falls short of their intended meaning in context.

2.2 Lexical pragmatics

Deirdre Wilson and Robyn Carston (2007) present an account of lexical pragmatics that extends Relevance Theory's inferential account from sentence level to word (and phrase) level. Three points are central to this claim. First, that it is not only sentence meaning but also word meaning that must be pragmatically adjusted in context. Second, that the same cognitive and communicative principles that are at work in recovering the speaker's meaning at sentence level are used at word level to fine-tune individual words or expressions. The second point sets Wilson and Carston apart from most pragmatic and philosophical accounts of lexical adjustment. Where others conceive of lexical adjustment as a distinct process from general pragmatic processes, Wilson and Carston see a single process, they see lexical adjustment as an homogenous part of the overall interpretive process. Where others further subdivide the lexical pragmatic process into distinct processes of metaphoric extension, narrowing and approximation, Wilson and Carston see these as different *outcomes* of a single mechanism that, again, is a homogenous part of the overall interpretive process. Lastly, the relevance-theoretic approach to lexical pragmatics holds that whatever processes are at work, they contribute to the forming of an occasion-specific sense, an "ad hoc concept" which is a constituent of the (intuitive) truth-conditional content of the utterance (i.e. the explicature).

Lexical broadening is described as involving "the use of a word to convey a more general sense than the encoded one, with a consequent expansion of the linguistically-specified denotation" and *lexical narrowing* as "the use of a word to convey a more specific sense than the encoded one" (Wilson & Carston 2007: 234 & 232). As mentioned above, most of the work on this topic has assumed that these outcomes were the products of different pragmatic processes. In contrast, Wilson and Carston's radical unified account holds, in accordance with RT, that a single interpretive process can adequately account for all cases. In other words, narrowing and broadening, including approximation, hyperbole and metaphor, mobilise the same non-demonstrative inference processes as 'literal' interpretations.

One particularity of the interpretive process in the case of lexical adjustment is that it constructs *ad hoc concepts*. In accordance with earlier work on pragmatic adjustment, Wilson and Carston believe that "lexical interpretation typically involves the construction of an *ad hoc concept* or occasion-specific sense, based on interaction among encoded concepts, contextual information and pragmatic expectations or principles" (Wilson & Carston 2007: 230). Notice that the interaction producing the ad hoc concept basically takes three types of input, the first of which is linguistically encoded concepts. In fact, in the type of communicative situation of interest here – that is, one requiring *lexical* adjustment – the speaker provides evidence of his or her intention to convey a certain meaning by uttering a *linguistically* encoded piece of evidence: a particular word or expression. Another important point is that the *interaction* mentioned above is part of an *inferential* interpretive process: the *only kind* of associations that can affect lexical adjustment are *inferential ones*. This point is important both for the lexical pragmatic approach described here and for the view of memory traces that I present in this paper. Let's consider an example of an inferential interpretive process that creates an ad hoc concept. You receive the following message on your phone:

(1) Be an *angel* and pick up some bread on your way home[5].

According to the lexical pragmatic approach, the word 'angel' has a linguistically encoded meaning, ANGEL, that activates logical properties in the form of inferences (such as the property of being SUPERNATURAL BEINGS), and which enables some deductive inferences (I am being asked to be a SUPERNATURAL BEING). It also activates encyclopaedic properties of different kinds of angels (good angels, guardian angels, avenging angels, and so on) which also enable conclusions to be drawn. In the context of (1), a good angel is likely to be the most accessible and so the encyclopaedic properties that would receive the most activation would be kindness, helpfulness, and watchfulness. From there, I could tentatively consider that I am being asked to be the good type of angel. But am I being asked to be supernatural or simply kind and helpful? The account holds that I tentatively adopt the implications that fit with my expectation of relevance and look for an interpretation of ANGEL that would warrant them. To do this, I create an ad hoc concept ANGEL* that is narrower (includes only good angels) and broader (includes non-supernatural beings that share properties such as being kind and helpful with good angels) (Wilson & Carston 2007). In the account that I outline in sections 3 and 4, the creation of this ad hoc concept is an experience with language that memory records. This means that the next time I come across 'angel' in a similar situation, I can activate the linguistically encoded properties of ANGEL but also a memory trace containing my own previous interpretation, arguably minimising my cognitive effort (see Sperber & Wilson 2002). Wilson and Carston (2007: 238-239 and 241) also mention that what starts out as an ad hoc concept could, as a result of frequent use, result in a new encoded

[5] This example is adapted from *Sally is an angel* (Wilson & Carston 2007).

meaning and they add that this has perhaps already happened, at least for some people, with 'angel'.

Additionally, broadening and narrowing are best understood as a matter of degree. They are products of a single interpretive mechanism and each simply occupies different points on its own continuum. This can be illustrated with sentences/utterances that can be given different interpretations. For instance:

(2) I *made* you some coffee/tea/cookies.

Decoding the word 'made' in (2) would activate the linguistically encoded concept PREPARED and accompanying encyclopaedic properties. The second source of input to the process is context: the interpretation of "I made you some coffee" differs significantly depending on the situation. After all, it's unlikely that I would be equally impressed by the minimal act of pressing a button on a coffee machine as I would be by the grinding and percolating and/or by the foaming of milk and the stirring in of sugar. A context can even be imagined where visiting a friend on a coffee plantation, the phrase "I made you some coffee" implies much more. Pragmatic expectations of relevance, the third element above, drive my interpretations. In the first two, the sense seems narrowed, 'made' as in PREPARED is restricted to the different (long or short) series of actions required for 'make' to apply to the situation. The last interpretation could arguably be labelled a case of broadening. If my friend has planted, harvested and created a special blend of coffee especially just for me, then uttering "I made you some coffee" is an understatement. Similarly, the interpretation of "I made you some cookies" differs depending on the extent of your expectations. If you expect 'made' to mean that I started from scratch, then knowing that I have bought the dough in a tube at the grocery store might invalidate my claim to having 'made' the cookies, but if you are already surprised to see that I have used an oven, then where the dough came from might be unimportant. Furthermore, notice that it is not because 'make' is a polysemous word that broadening and narrowing are necessary. It was specified that 'made' should be understood as PREPARED (arguably one of its established senses) from the start.

What's more, the undeniable intuition that narrowing, approximation, understatement and other kinds of loose uses create strong effects in utterances is not to be left unexplained. In fact, the relevance-theoretic approach gives an explanation to the phenomena that underlines a fundamental aspect of the construction of ad hoc concepts: *the influence of encyclopaedic knowledge and background assumptions*. This can also be illustrated with an example:

(3) It's a marathon[6].

Depending on the context in which it is used, this sentence can be analysed as literal, approximate, hyperbolical or metaphorical. These differences depend not on which mechanism of interpretation is used - *under the unified account there is only one* - but rather on the encyclopaedic assumptions that are deployed in a particular context for the utterance to satisfy specific expectations of relevance. The encyclopaedic properties and background assumptions associated with a marathon include the following:

(i) A COMPETITION ORGANISED SINCE THE FIRST MODERN OLYMPIC GAMES THAT INVOLVES RUNNING A DISTANCE OF 42.195 KILOMETRES IN THE FASTEST POSSIBLE TIME, ETC.

(ii) INVOLVES ORGANISING COMMITTEES, ROAD BLOCKS, PARTICIPANTS MUST ENROL, ETC.

[6] This example is adapted from one used in Wilson & Carston (2007): the water is *boiling*.

(iii) SUCCESSFULLY COMPLETING A MARATHON REQUIRES HIGH LEVEL OF FITNESS, IS OFTEN PRECEDED BY A PERIOD OF TRAINING AND PREPARATION, FOLLOWED BY PHYSICAL EXHAUSTION, ETC.
(iv) REQUIRES PSYCHOLOGICAL STAMINA, SELF-DISCIPLINE, DETERMINATION, ENDURANCE, WILLINGNESS TO MAKE SACRIFICES, ETC.

The addressee's mental representations concerning marathons include the properties in (i)–(iv) which are not all activated to the same degree on different occasions of utterance. For instance, if I try to drive into London and upon encountering a road block ask a policeman if it's because of a parade, he might answer "It's a marathon". Processing this utterance is straightforward because the discourse context, that is, for instance, the road block and the preparations I can see around me, makes certain encyclopaedic properties of MARATHONS, such as those in (ii), particularly highly accessible. I therefore easily arrive at an interpretation, 'literal' in this case, that satisfies my expectations of relevance. This same utterance can receive a different interpretation on a different occasion. Suppose a friend of mine spends his weekends training for all sorts of competitions. I ask him why he is currently running *X* times around the park and expect an answer to do with distances. When he replies "It's a marathon," the properties in (i) associated with distance will be particularly accessible. Through a process of mutual parallel adjustment (see below), I tweak 42.195 KILOMETRES to *APPROXIMATELY* 42 KILOMETRES. This will enable sufficient contextual implications for my expectations of relevance to be met. We can call this use of the word 'marathon' an approximation because some of the most accessible encyclopaedic properties, such as those in (ii), are not those that yield the overall interpretation that satisfies my expectations. A metaphorical interpretation might especially depend on specific contextual elements. My friend is planning her wedding, she complains about not having any time for anything else and when I ask if it's really that much work, she answers "It's a marathon". I would not take her as committed to saying that she is actually running a race or that she is doing any kind of competitive running. Rather, the encyclopaedic properties in (iv) associated with psychological endurance are highly activated; perhaps those in (ii) associated with detailed and lengthy organisation are also activated, albeit to a lesser extent. As with the ANGEL example above, I create an ad hoc concept that would warrant the kind of implications that would best satisfy my expectations: It's a MARATHON*, paraphraseable as 'an exceedingly demanding task that will give her no rest until completed, etc.'

These examples illustrate the process of mutual parallel adjustment whereby "tentative hypotheses about contextual assumptions, explicatures and contextual implications are incrementally modified so as to yield an overall interpretation which satisfies the hearer's expectations of relevance" (Wilson & Carston 2007: 248). As always with the process of interpreting an utterance (or other ostensive stimulus), the goal is to achieve the best yield of cognitive effects for the least cognitive effort, modulo recognition of the speaker's competence and preferences. Here, the cognitive effort is expended on the activation of encyclopaedic properties and the process of mutual adjustment, and the cognitive effects gained are the contextual implications that are thereby inferentially confirmed. The subject's goal in this case is therefore to arrive at the optimally relevant interpretation of "It's a marathon," whether that be literal, approximate or metaphoric, by deriving implications (cognitive effects) from the most highly activated among the set of encyclopaedic properties.

3. Introduction to memory traces

The previous section summarises the existing relevance-theoretic account of lexical pragmatics. The remainder of this paper focuses on presenting and discussing a relatively

new notion, 'memory traces,' that would add an interesting new component to the picture sketched above. To be more precise, I explore the possibility that memory records what lexical pragmatics describes as the products of the interpretation process and that future interpretation processes benefit from what memory has recorded in the past. I argue for this claim by presenting original research ranging from theories of memory to accounts of language change. Research on memory traces is only in its initial stages, and applications of memory traces to language studies are scarce. As a result, much of the material presented here supports my claim only indirectly and there is little experimental evidence to report. As I hope to show, however, this does not mean that the case for memory traces as a potentially fruitful input to lexical pragmatics is weak, only that, thus far, it has not been explored and reported thoroughly.

In this section, I first introduce the notion of memory traces through a brief outline of the research of Douglas Hintzman on multiple-trace memory models (1986, 1988). Next, I discuss a possible fit between multiple-trace memory models and the current first language acquisition frameworks of cognitive and developmental psychologists Paul Bloom and Michael Tomasello. The final part of this section addresses possible objections to my interpretation of this fit and provides further evidence that memory traces play a role in how memory records our experiences with language.

3.1 Hintzman's multiple-trace memory model

Many researchers have adopted multiple-trace type models since they began to appear in the 1980's (e.g., Hintzman, Raaijmakers, Shiffrin, Steyvers, Dennis). A review of this literature is beyond the scope of this paper (a partial review for the non-specialist is available in Neath and Surprenant 2003), but in-depth treatment of the advantages of this type of model was offered in two major articles by Hintzman (1986, 1988) and with some precautions, his presentation can serve to explore the topic. The articles discuss the basic design features of *MINERVA 2*, a computer model designed to simulate human memory. Contrary to previous models, Hintzman's model does not assume that information coming into memory is filtered by some form of executive control. Nor does it assume that the information coming into the system is fundamentally transformed before it can be captured in memory.

As their name suggests, memory *traces* are evidence left behind by the passage of an experience through awareness. Traces preserve, perhaps imperfectly, the configuration of properties making up the experience. According to Hintzman, previous models that assume memory records *abstractions*, or, in other words, schematic representations of decontextualised content, are faced with the difficulty of explaining *how* these abstractions are formed at the time of input. This is especially so now that the cognitive sciences have recognised the fallacy of positing centralised control of modules that are best described as automatic. In Hintzman's model, information flow is largely unfiltered, and there is no homunculus 'inside the machine' performing abstractions and discarding (allegedly) irrelevant information. Rather, memory preserves in a memory trace the episode *as it was experienced*. Additionally, following Relevance Theory, we can assume that the input captured in the trace is maximally relevant because memory records what we attend to and the cognitive principle of relevance would predict that our awareness is attracted to optimally relevant aspects of our environment.

Another important feature of Hintzman's model, and one that sometimes differentiates it from other trace models, is that *each* experience attended to creates a trace in memory and if an experience is repeated, a pre-existing trace is not strengthened, but rather a *new* memory trace is created. This means that the 'weight' of traces does not come from the degree of saturation of a particular trace as in previous models that would add up instances and

represent them as 'thicker' or 'thinner' single traces. According to Hintzman, only *multiple*-traces adequately capture all the detail necessary to model something as complex as human memory. In single trace accounts, information from different episodes is combined and, therefore, arguably, the structure of the individual episodes would be lost, and along with it, information such as recency. In contrast to this, the design features of MINERVA 2 allow it to combine the information contained in memory traces *only* when that information is demanded. Hintzman's major insight is that the subtlety needed to explain how memory works requires the system to be able to combine information stored in memory and information in the current situation in *the most flexible way*, which calls for multiple-trace models.

The importance of the current context to Hintzman's model also facilitates the link I establish between this perspective of memory and Relevance Theory's cognitive principle of relevance. According to Hintzman, when the system needs to access stored information, the information is there in its most complete form, and memory can *differentially activate particular elements in existing memory traces according to their usefulness to the present situation*. Therefore, in both Hintzman's model and in RT accessing information in memory means that only what is optimally relevant gets activated. Notice that following the general cognitive principle of relevance, this input is twice optimally relevant. It was optimally relevant when it was first captured in a memory trace and, at a later time, a pool of such relevant traces is differentially activated according to potentially different expectations of relevance. Hintzman describes this retrieval process as involving a *probe* and an *echo*. Roughly, the probe contains the elements of the current situation and traces are activated according to their degree of similarity with this probe. The echo is the summed contributions from activated traces. Successive instances of probing and echoes can operate in loops, the echo of a first activation process can become itself a probe in order to narrow or broaden the spread of activation *until the subject is satisfied or the search abandoned*.

Once again, Hintzman's perspective seems particularly well adapted to an application in RT as RT already posits interpretation processes that operate in loops (mutual adjustment) until the subject is satisfied or abandons his search. Bringing the two perspectives together could reveal similarities between memory and interpretation concerning how relevance orients cognitive tasks. This paper is only a preliminary exploration of this possibility.

I now turn to a discussion of the role of memory traces in children's acquisition of word meaning.

3.2 Memory traces and children's acquisition of word meaning

As suggested above, memory traces can be understood as the *imprint* that the passage of an experience through awareness leaves behind. A *trace* is also a minute amount of something, almost undetectable, yet present. An application of this metaphor to the process of children learning language suggests that they may build their word meanings up out of small amounts of evidence but also that the clues they need deposit themselves in memory effortlessly, unintentionally, almost as a side effect to awareness. To my knowledge, no one has yet put forward an account describing how such traces left in memory by early experience with language would participate in children's learning of word meaning. To explore this topic, I propose a brief discussion of approaches to first language learning illustrated with some examples. The objective of the theoretical discussion is not to offer a complete review, but rather to point to a possible compatibility between established frameworks and the view of memory traces defended in this paper.

The earliest formation of memory traces concerning language can be illustrated with two opposing examples. One of the examples is an extremely simple case of learning a word

in an unambiguous situation. For instance, the child's attention is attracted to a dog and noticing this, his caretaker utters 'That's a dog' to which the child responds /daw/. This simplified example is nonetheless a learning episode in which memory creates a novel trace; however, because in this situation the referent is immediately recognizable, the need for repeated occasions of the same use, and their resulting multiple-traces, is obscured. Another type of example is needed to illustrate how the child might need to accumulate traces before being able to correctly pick out the referent.

Suppose a child is learning his first words of English and during his daily bath, while sitting and splashing in a tub full of water, he successfully isolates the phonological word form for 'bath' from the speech flow of his adult caretaker. He has surely heard /bɑːθ/ before, but let's suppose that he has only marginally attended to it and, consequently, only a very faint trace in memory is activated by /bɑːθ/ now that he has begun focusing his attention on this word. In any event, he now creates a memory trace that tentatively associates the linguistic form /bɑːθ/ with his present situation, which includes other words (/beɪbi/, /toɪ/), and salient aspects of his interaction with the adult. Notice that the 'correct' association is probably not straightforward for the child. How is he to know which elements of the situation the conventional label /bɑːθ/ corresponds to? Is it BEING WET? Or SPLASHING AROUND? Or yet the GAME PLAYED WITH THE FLOATING TOYS? All of these elements are potentially salient aspects of the child's situation and of his interaction with the adult who is uttering the word. This issue is often raised by referring to Willard V. O. Quine's famous analogy involving a visiting linguist and a local informant who utters "gavagai" as a rabbit dashes by. Quine's point was that since there is an infinity of logically possible meanings, *or translations*, for 'gavagai,' the linguist could never be sure he got the right one.

Modern theorists point out that despite the difficulty of learning words and word meanings, normal children do it effortlessly and at an amazing speed. Paul Bloom (2000) carefully describes the underlying capacities children bring to the task of learning the meanings of words. Roughly, there are three: (i) 'theory of mind,' or, in other words, a child's ability to understand, *or read*, the thoughts of others; (ii) the ability to grasp the associated concept (the concept DOG or BATH); and (iii) the ability to segment speech in order to make the form (/dog/, /bɑːθ/) of the word available. These capacities are remarkable, yet they are still not enough to learn the meanings of all words. Despite children's mind reading capacities, they are not telepathic and so, in many situations, like our bath example above, the overt behaviour of their adult caretakers would not suffice. For these cases, Bloom suggests that children rely on their sensitivity to syntactic cues, their global understanding of naming and, importantly, on "information gained by hearing a word used on multiple occasions" (Bloom 2000: 86).

Furthermore, as the occasions in which the child hears a word add up, his understanding of the meaning of that word is more complete. The memory traces brought into play in any instance of word learning involve any previous episode that can be associated with the current situation. In the case of 'bath,' this includes any traces of having heard that word before or other words often appearing with it, plus other syntactic clues; additionally, any traces having to do with concepts such as TUB, BATHROOM, BEING WET, SPLASHING and so on; finally, any (even) tentative conclusions involving the child's previous trials of naming or general intention reading in similar situations.

In the future, encountering the sounds /dog/ or /bɑːθ/ would potentially simultaneously activate *all* of their associated traces, not only the one created in the instance of naming, when the child first associated /dog/ with the animal before him, but also other traces that are more or less closely associated with this one. The more the meaning of a word is unequivocal, the more the traces related to the meaning of the word 'overlap'. In the case of /bɑːθ/, these might initially be quite 'scattered,' that is, they might contain both truly

relevant aspects of the bath situation for the *meaning* of the word 'bath' and some indirectly relevant, or wholly irrelevant, aspects. For instance, it is relevant to the child, and to his interaction with the adult, that this situation is an opportunity to play, but this is not the most relevant aspect of the situation for the meaning of the word 'bath'. A bath is rather an instance of washing. For reasons that will become clear in subsequent sections, this is not a disadvantage of the multiple-trace view. For now, we can see this apparently disorienting situation for the child as only a phase, since as the child's experience with language grows, the traces involving 'bath' accumulate, and progressively, those elements that the child can *reliably* associate to the use of the word 'bath' emerge from a background of more loosely associated elements.

Another interesting perspective on this issue is provided by Michael Tomasello's (2003, 2008, and 2009) in-depth account of how children acquire language. His *socio-pragmatic* approach highlights three aspects of the interactions between children and caretakers that are important for our purposes. First, that prior to their first use of words, children already participate in tasks requiring 'shared intentionality'.[7] In other words, they attend to whether the adult is attending to them as well as to the task and share some understanding with the adult as to their common objectives. Second, that language learning takes place within a structured series of routines and a rich social-interactive environment. In all cultures, even those where pre-linguistic children are not directly addressed, they participate in the everyday routines of households: they at least sit by while the adults gather supplies, prepare food, and wash up while engaging in conversation. Third, that shared intentionality in this social environment constitutes the social-cognitive equipment that children need in order to read the intentions of others, especially those of the adults with whom they interact (2003: 87). Tomasello believes that, thus equipped, infants can 'tune into' word use in their everyday routines. The diversity of strategies that they mobilise to do this is an object of growing research (see Bloom 2000, Tomasello 2003 *and the references therein*, Csibra 2010) but a promising lead is that "children use all kinds of interpretive strategies based on pragmatic assumptions that utterances are somehow *relevant* to the ongoing social interaction" (Tomasello 2003: 89).

Word learning is not a simple mapping between linguistic conventions and experiences. In our 'bath' example, the child is unlikely to form arbitrary associations between the form /bɑːθ/ and elements of the situations in which it was uttered because, even before his entry into language, an infrastructure of shared intentionality has been established, with, for instance, mother-child dyads and early pointing gestures. Memory traces record instances of joint attention in which children can determine to a large extent where the attention of the adult is focused when he utters a particular word or phrase. This limits, although it does not eliminate, potentially misleading associations the child might record. The fact that mapping mistakes are so rare points not to telepathic abilities in children but rather to the robustness of the learning process described above. In our example, the child could still mistakenly believe 'bath' refers to PLAYING WITH WATER. After all, he has often experienced this type of interaction while taking his bath and, on the contrary, might be oblivious to the adult's intention of washing him (a concept he might not yet fully grasp). However, if we consider that the child has a good stock of examples of adult uses of the word 'bath' stored as memory traces, we can begin to understand how he can manage to be so cautious. Suppose that he has traces associating a certain relevant aspect of the bath situation (i.e., PLAYING WITH WATER) with the word 'bath,' he could still refrain from using the word with this sense if he does not

[7] While Michael Tomasello uses the term *joint attention* in his earlier work, more recently, he has adopted *shared intentionality*, which is more inclusive.

have corroborating evidence, in the form of memory traces, that adults actually use this word in this way.

Lastly, it should be noted that there is no incompatibility between the fact that, in particular circumstances, children can *use* a new word after only one exposure to it and that they *learn* most words after repeated exposures. The same theorists who design experiments that support the first point stress that word learning greatly benefits from multiple occurrences. Paul Bloom, for instance, uses a different-bits-on-different-occasions conception of word learning to explain how it can be so fast. A good conservative number for the average vocabulary of an 18-year-old is 60,000 words (Pinker 1994, cited by Bloom 2000). If word learning starts at 12 months, then that means that a child learns a word every waking 90 minutes, or 3,750 words per year. However, "Learning the precise meaning of certain words, especially verbs, might be a long process requiring many trials (…). On any given day, then, it might not be that children are learning 10 words; they might be instead learning one-hundredth of each of a thousand different words" (Bloom 2000: 25). Memory traces are these fragments deposited in the child's memory as he accumulates more experience with language.

3.3 Memory traces beyond first language acquisition

A first foreseeable objection to the picture presented above is that memory traces are only necessary while the child is coming to grips with word meaning. According to this position, the child, or anyone learning a new word, would accumulate traces only until he could extract the central gist for each word and could thereafter do away with his overwhelming collection of memory traces. He would no longer need them to reliably understand and use the words of his language community. Notice that this explanation does not challenge the claim that memory traces are at some stage created and used. It simply assumes that the only reason they are accumulated is *in order to* form a single representation sufficiently similar to the 'linguistic meaning' of the language community and consequently that confirmed language speakers would have *no other need for them*. Yet, if we accept that traces provide a record of our experiences quite generally, there is no a priori reason why memory traces would not continue to capture our experiences with language.

Indirect evidence that previous traces continue to be useful even after a great number of traces have been accumulated in memory is supported by Hintzman's success at simulating human memory. His model crucially depends on the availability of a large number of traces that represent a subject's experiences. These traces are not discarded once they have been used, quite the contrary, they are reactivated each time they are relevant and if they continue to be relevant, there is no reason to discard them. New traces can be added on, they don't necessarily replace the old ones. When old traces *are* discarded, it is because they have not been activated in a long period of time. Recency of use and frequency of use are important characteristics of memory traces, accounting to a large extent for their degree of accessibility - how this works is an issue I plan to address in future research.

The view that new traces are continually added on (rather than only in a first learning phase) and that old traces are discarded only when they are no longer useful is also supported by evidence from language change. Above I argued that the outcome of an interpretive process that adjusts a core meaning is just the kind of experience that memory traces could capture. It stands to reason that the trace thus created might be useful in the future when a similar situation calls for a similar adjustment. Furthermore, if the contexts consistently *pull* the sense in the same direction, the traces of a whole language community could be the object of adjustments that would themselves be the object of other adjustments in a continuous chain of meaning change. Such a reading can be given to the famous example of 'going to' given in

the grammaticalization literature. Roughly, the lexical item ‘go’ came to serve as an auxiliary in certain linguistic contexts. For instance, when a locative adverb is not present and ‘go’ can be interpreted as intentional movement, an inference of futurity became available. An interpretation that beforehand would have been bracketed [I am going [to see the king]] became [I am going to [see the king]]. With this change, ‘going to’ - which till then was only the progressive aspect of ‘go’ with the directional ‘to’- acquired the possibility of expressing immediate future (Hopper & Traugott 2003: 3, see also Traugott 2004, Bybee 2010). A detailed account of cases involving language change would be of interest to the topic at hand, yet limits of time and space do not allow it. Clearly however, the role of memory traces in this perspective would be to record the variability and frequency that makes language change possible.

Further supporting evidence for linguistic memory traces comes from quantitative studies on language variation and change. Pierrehumbert (2001) reports on Bybee’s (2000) study where a statistically significant relationship between the deletion of /t/ and /d/ and word frequency is revealed. Pierrehumbert, following Bybee, reasons that if deletions are significantly more common in frequent words than in infrequent words, then it is possible that specific phonetic details are stored with the word itself. This in turn challenges the received view that all language categories are abstract, or, in the terms that we have been using, that memory records only context-independent content that has been abstracted away from the context of use. The model that grew out of these first observations is called the Exemplar Model. Detailed phonetic memories associated with individual words, or *exemplars*, are recorded in memory. Frequency effects are then logically explained by the corresponding larger populations of exemplars in memory. Exemplar models quickly spread to other areas of language studies where ‘categories,’ whether they be phonemes, words or constructions, show variability and frequency effects (see Wedel 2006, Croft 2008, Bybee 2010). Once again, I plan a more thorough treatment of this topic in future research. For now, it seems reasonable to suggest, with some confidence, that multiple traces and exemplars capture the same phenomena. Just as phonetic exemplars have been demonstrated to capture variability in word pronunciation, we can maintain that *pragmatic* exemplars would be the recorded outcomes of previous pragmatic processes of adjustment (see Ariel, 2008; Bybee, 2010) and would impact on future uses.

To summarise this section, I imagined a child at the beginning of his language learning process having too few memory traces associated with the form /bɑːθ/ to reliably use this word. I further imagined that his memory traces, or *exemplars*, would add up progressively as his experience with language grew and that he would thus come to grips with the meaning of the word ‘bath’. I also considered independent evidence suggesting that their role in the acquisition of the meaning of words is not the only way in which memory traces affect our use of language. For my purposes, the most important hypothesis arising from the exploration above is that the meaning of a word is not represented in memory by a single trace or abstracted sense. Rather, it seems likely that memory stores meaning in multiple traces.

Thinking back to the cognitive principle of relevance, reasoning that relevance is a potential property of memories seems largely validated by the evidence presented so far. In the next section, I present an interpretation process illustrating the interest of memory traces to a relevance-oriented lexical pragmatics.

3. Memory traces and lexical pragmatics

The main objective of this paper is to suggest another source of input for the non-demonstrative inference process described above in section 2.2 on lexical pragmatics. Three basic types of input have already been mentioned: (i) the concept encoded by the word at

issue (and the concepts encoded by the other words in the utterance), (ii) contextual information and, (iii) the addressee's pragmatic expectations of relevance. The fourth source of input would be memory traces of previous encounters with the word. The claim is that memory records each previous utterance together with the context in which it appeared and the relevance-driven inferences that the pragmatic expectations of that occurrence legitimated. Furthermore, the claim is that this information, gathered by memory through language use, is routinely activated by relevance-oriented cognitive processes.

As mentioned in previous sections, one of the difficulties of studying memory traces is finding a way to observe them as directly as possible. My aim in this section is to present a series of observations that highlight their role in everyday language use. To do this, a multitude of memory traces that can be associated with a single word is presented, each of which, in accordance with the claims made in this paper, is the stored trace of a previous experience with the word. Lexical pragmatics tells us that such an experience can be the theatre of pragmatic processes of interpretation that combine, on-line, linguistically encoded concepts and contextual information (including relevant encyclopaedic information associated with the decoded concept) guided and constrained by the pragmatic expectations of relevance of that particular moment. The amount and variability of experience with a certain word determines what the 'cloud' of memory traces looks like. If previous experience with the word has not created any ad hoc concepts, then the available traces can be imagined as concentrated around a 'core'. This core is the linguistically encoded sense. In such a case, the original three elements combine and form an *initial* ad hoc concept, or occasion-specific sense, as described above in the section on lexical pragmatics. This ad hoc concept is recorded in a trace and becomes potentially available for future experiences with this word. Of course, if previous experiences with the word *are* available, that is, if memory traces of previous occurrences were recorded, then they are activated according to their relevance to the current situation, thus becoming the fourth element in the combination that creates a new ad hoc concept. As a result, an important feedback loop is created. Each time that an ad hoc concept is created, this interpretive experience is stored in memory and can be called upon again.

The word I have chosen to illustrate this process is 'gene'. It is a relatively new word in order to keep the length of the comments manageable, but its history is also complex enough to demonstrate how memory traces work. The term serves our purposes particularly well because, despite coming from the very technical domain of *genetics*, it has already been adopted into everyday language and is often subject to quite loose uses. Much could be said about the scientific developments that led to our present understanding of genes, but rather than trying to summarise them here, I focus on the way the word has been used. It was coined in the early 1900's to refer to the material entities of heredity discovered by Mendel. Before 'gene' was coined, the word 'genetic' already existed. It was used to speak of origin, development and common evolutionary origin *but not heredity*. The scientific discovery of genes forever changed this and, consequently, today the notion of heredity is also attached to 'genetic'[8].

Following RT and lexical pragmatics, let's first consider what properties are activated when I encounter the word 'gene'. I do not mean these to be exhaustive nor scientifically precise. The waves of enthusiasm with which the succeeding scientific discoveries have been greeted by the general public are much more important to the everyday use of the word than

[8] For instance, before this change in the meaning of 'genetic', today's 'developmental psychology' was termed 'genetic psychology' (Oxford English Dictionary –online version).

the actual content of these discoveries. Our general background knowledge, or in other words, what we use to understand the concept GENE, could thus be listed as follows:

(i) A GENE IS (according to one of the most important discoveries of the 20th century) *THE* UNIT OF HEREDITY OF THE LIVING WORLD.

(ii) A GENE IS THE *MATERIAL* SUPPORT FOR INHERITABLE TRAITS.

(iii) A GENE IS A SPECIFIC SEQUENCE OF DNA LOCATED AT A SPECIFIC SITE ON A CHROMOSOME.

(iv) GENES ARE THE BASIS FOR FILIATION, GENES ARE THE MATERIAL THAT IS PASSED ON FROM GENERATION TO GENERATION, SHARED BY THE MEMBERS OF THE SAME FAMILY, etc.

According to the lexical pragmatics account outlined in section 2.2, these encyclopaedic properties, or at least those of them that a particular addressee has stored, are potentially available to be activated when 'gene' is encountered in an utterance and decoded. Let's look at an example. After spending Christmas with Peter's extended family:

(4) Mary: Why does everyone in your family know how to cook so well?
Peter: It's in our genes.

The encoded concept GENE interacts with the conversational context of utterance (4) and, as a result, encyclopaedic properties such as those in (i)–(iv) are activated to varying degrees. Since Mary is expecting an answer to her question concerning why a particular skill is so well represented among the members of Peter's family, the encyclopaedic properties in (iv) are likely to receive more activation than those in (i)–(iii). On activation of (iv), Mary can tentatively adopt some contextual implications:

(5) The members of Peter's family all know how to cook so well
Because they are all members of the same family (triggered in part by the linguistic clue 'our').
Because they share X (call it *genes*).
Because they have a common origin.

The fact that Mary is not going to take Peter to be communicating that there is a specific genetic code that predisposes to good cooking means that her interpretation cannot simply rely on the encoded meaning of 'gene'. She must create an ad hoc concept better suited to the present situation. She narrows the concept GENE to that which is shared by the members of a family and broadens it to include non-material elements that can be shared. The result is GENE*. Now Mary can derive an implicature that will satisfy her expectations of relevance: Peter has uttered 'It's in our genes,' and communicated the explicature GOOD COOKING IS IN OUR FAMILY'S GENES*. The property denoted by GENES* is CONCERNS ABILITIES THAT ARE SHARED WITHIN A FAMILY AS A RESULT OF PROLONGED INTERACTION, NOT NECESSARILY AS A RESULT OF DNA OR HEREDITY. The lexical pragmatic account of this use of 'gene' would stop here. This use of 'gene' is not a case where an ad hoc concept has become conventionalised and so the lexical pragmatic approach would hold that, if I encountered 'gene' again, I would again decode GENE and derive implicatures as above until my expectations of relevance were satisfied or abandoned.

Alternatively, the claim defended in this paper is that the interpretation of (4) would involve not only the three elements mentioned above but also the memory traces of previous uses. Let's suppose that our imaginary conversation took place in the year 2000 and let's look at some possible occurrences of the word 'gene' leading up to that moment.

Thanks to Google's Ngram project, a massive online database containing approximately 4% of all books ever printed, the appearance of the word 'gene' can be visualised and its frequency monitored up to the year 2008[9]. The occurrences of the word grow evenly from the 1920's to the 1970's and show exponential growth from there until the year 2000. Between 2000 and 2008 the trend is downward (see Annexe, Figure 1). Occurrences of the plural form 'genes' paired with 'good' and 'bad' also appear around the 1920's and their frequency is relatively steady until the end of the 1980's. From 1990 to just after the year 2000, 'bad genes' doubles in frequency, 'good genes' shows exponential growth from 1990 to just after the year 2000 (see Figure 2). This pattern of explosive frequency is quite common among expressions containing 'gene'. It unquestionably reflects growing interest in genes, and, at least initially, the word is likely to have been used predominantly in its 'literal', scientific sense. However, as in (4), at some point, we know that it began to be used in a loose way, in a *figurative* sense.

Can we use these numbers to talk about memory traces? The raw data that the Ngram viewer uses to illustrate the frequency of words over time is the relative frequency of these words in printed books. To arrive at their numbers, the team at Google counted how often a word appeared in the publications of a particular year and then divided it by the total number of words also appearing in the works of that year (Michel *et al*. 2011). How well does this reveal the underlying frequency of use by people in conversations? And even more difficult, how well does this reveal the quantity of memory traces they store in memory with regards to a particular word? Perhaps not directly, but the sheer size of the corpus employed (5 million books and a corpus of 361 billion words) means that the trends are extremely robust and that trajectories as clear as the ones pictured in the graphs at the end of this paper deserve an explanation. Having taken these precautions, and assuming that we accept that memory traces record uses of the word, we can, at least tentatively, link frequent uses with a greater number of associated memory traces and thus higher accessibility. Let's look at some further examples. The interpretive processes are described very briefly and with very little context, as the objective is not to comment on how the occasion-specific sense is created but rather to note that occasion-specific senses are common occurrences rather than marginal phenomena.

(6) gene pool

Encountering this expression in a neutral context in the 1950's probably led to the activation, among other things, of encyclopaedic property (ii), and to an interpretation along the lines of a 'gene pool' is THE SUM OF THE MATERIAL ENTITIES *CALLED* 'GENES'. But the frequency of the expression 'gene pool' steadily grew from the 1950's to the mid 1970's, experienced a brief decline and grew again from the 1990's to the year 2000 (see figure 3). According to my account, as the decades passed, the individuals exposed to these uses stored occasion-specific senses as memory traces. These traces were thereafter activated if relevant or discarded if they remained unused. Reconstructing exactly what these memory traces contained is impossible, but some clues are available. For instance, the frequency of the expression 'gene pool' increased by approximately 15 percentage points (from 0.00003 to 0.000045%) in the

[9] Google Ngrams (http://books.google.com/ngrams) is a convenient way of visualising general trends in word use. It must be used with precaution because it contains two types of errors: those linked to the accuracy of character recognition (the data base is linked to Google's scanning project) and, more important to us, the year attributed to the source is sometimes mistaken. For the very general trends that interest us here, these errors are not an issue. For a detailed presentation of Google Ngrams, see Michel *et al.* (2011).

1990's. It is unlikely that this popularity is solely due to interest in scientific discoveries[10]. Rather, perhaps 'gene pool' captures technical aspects of genes in an easy to understand compound ('*X* pool') and thus is easily used loosely but keeps its evocative power. Today, frequent uses of 'gene pool' include "a shrinking gene pool" which, it should be noted, can be understood both literally and figuratively. 'Gene pool' can be used in clearly metaphorical expressions like 'put some chlorine in your gene pool'. Roughly, an inferential interpretive process like the one adopted in this paper takes all the relevant input from the linguistically encoded concepts of SHRINK, GENE and POOL, a context of utterance, *among other things*, and constructs as a tentative hypothesis that a "shrinking gene pool" is A DIMINISHING POTENTIAL TO REMAIN COMPETITIVE, TO ADAPT TO YOUR SURROUNDINGS. Putting chlorine in your gene pool is a way of saying that A GENE POOL, LIKE ANY OTHER POOL, CAN BECOME 'CONTAMINATED' BY BAD GENES, IN WHICH CASE, IT NEEDS TO BE 'CLEANED'. Storing ideas such as health, inbreeding, adaptability, genetic diversity, and so on together with the expression 'gene pool' is a good way of ensuring that they remain accessible for future occasions of use, perhaps reducing processing costs and increasing cognitive benefits. Also, positing this stored diversity might facilitate explaining how figurative senses arise.

Here is another frequent pairing with the word gene:

(7) faulty genes

Before the discovery of genes, it was already known that certain diseases ran in families. The discovery of the physical support for this heredity led to an explosion of highly publicised scientific research. The expression 'faulty genes' appeared around the 1960's, at the same time as expressions such as 'genetic disorder,' 'genetic mutation,' and 'cancer gene' (see figure 4). In a neutral context, these expressions would activate encyclopaedic properties such as (iii), *the DNA that codes the gene has mutated*, and (iv) *(this mutated gene causes a disorder that) is passed on from generation to generation*, etc. So, faulty genes (or bad genes) predispose their bearers to diseases. Notice that this very important implication is not *directly* entailed by any of our initial encyclopaedic properties of genes. It somehow underlies our encyclopaedic knowledge: *diseases run in families, genes are shared by members of the same family*. In any event, arguably, our experience with expressions such as those above creates memory traces containing this idea, among others, and making it available for future inferential interpretation processes.

(8) The gay gene

This expression appeared in the 1990's and in the span of 10 years went from nearly zero occurrences to a relative frequency of 0.000004%, then from the year 2000 to 2008 dropped again to less than half of its previous high (see Figure 5). As with other loose uses of expressions with 'gene' that we have encountered, 'gay gene' probably leads the hearer to adjust her interpretation rather than taking the speaker as asserting a scientific fact about genes[11]. 'Gay gene', suggests, in a markedly loose sense of the word 'gene', that the bearer of such a gene is inherently gay rather than becoming gay through life experiences. This interpretation probably itself relies on the loose use of the associated word 'innate'. From the

[10] An in-depth look at the thousands of instances of this expression in that decade would be necessary to establish precisely what senses the expression took in context and when more loose uses started to appear. I plan to look into how this could be done.

[11] This would also be the case for less frequent but quite evocative expressions such as 'longevity gene'.

perspective of memory traces, the frequency with which we have observed such loose uses within this small corpus of expressions with ‘gene’ is worth noticing.

Going back to this section’s original example, we can now mention some of the elements that memory traces could bring to the interpretation process. When Mary decodes the word ‘genes’, she accesses two types of information from memory, the encyclopaedic properties (i)–(iv) made available by the linguistically encoded concept GENE and a host of memory traces of previous experiences with this word. The information contained in memory is activated by an inferential interpretive process, in other words, context and expectations of relevance interact with memory so that relevant aspects receive more activation. In the memory model outlined in section II, this means that the utterance plus the context of the utterance act as a probe in an activation process that produces as a result an echo. For (4), we can imagine that the echo contains something like SHARED IN A FAMILY, PREDISPOSING TO *X*, HERITAGE FROM PREVIOUS GENERATIONS, MAKING SOMEONE INHERENTLY *X*, etc.

Let’s suppose that Mary knows Peter’s family fairly well. After all, the interpretation depends, among other things, on contextual assumptions that Peter can assume he and Mary share. Mary tentatively adopts the following implications:

(9) The members of Peter’s family know how to cook so well
Because something they share (as a family) predisposes them to good cooking.
Because something they have inherited from previous generations predisposes them to good cooking.
Because something acting as mysteriously as genes makes them inherently good cooks.

Compared to my original analysis, the contextual implications hypothesized here by Mary are more detailed. The construction of the ad hoc concept still requires her to narrow and broaden GENE as before, but now, the biggest difference is that by taking memory traces of previous utterances into consideration, Mary has additional relevant input on hand to exclude inappropriate properties associated with GENE from her ad hoc concept, such as THE UNIT OF HEREDITY, THE MATERIAL SUPPORT FOR INHERITABLE TRAITS, DNA, and so on. She also has a multitude of memory traces in which ‘gene’ and ‘genes’ have been used in loose and metaphorical ways so that this possibility is naturally more readily accessible to her.

An important part of the situation of utterance concerns how much Mary knows about Peter’s family, aspects of which may play the role of contextual assumptions. These elements will be key in further broadening and/or narrowing of the concept under construction. For instance, Mary may know that Peter’s grandmother is a great cook and that Peter and his family have learned from her, and furthermore, that they now feel they share this as a family. Mary could narrow GENE to that which Peter’s family has inherited from their grandmother and broaden the material aspect to something that can be learned. The result is GENE*. Now Mary can take Peter as meaning that cooking has been transmitted in his family, cooking is something they share as they share genetic material. Critically, whatever Mary concludes, including tentative aspects of her conclusions, will be stored as further memory traces. They only remain in memory if future activations confirm their relevance; if not, they are discarded.

In this section, I have illustrated how the accumulation of memory traces could be envisaged. An interpretation process was proposed and the memory traces potentially feeding into it were presented and discussed. Finally, the original interpretation process was revisited and a role for memory traces was made explicit. The account given here is a first exploration of the possible role of memory traces in inferential interpretive processes. This view would greatly benefit from a more careful consideration of how the different sources of input are activated and from analysis of more numerous examples.

4. Conclusion

This paper set out to introduce the notion of memory traces and presented arguments in favour of an approach to lexical pragmatics that includes a key role for memory traces of previous utterances. Arguing for such an approach faces various challenges. Among the most important is the fact that this notion of memory traces is relatively new and that research connecting memory traces to language processing is quite scarce. As a result, methodologies for studying memory traces in everyday use of language are still largely to be developed. Additionally, as already mentioned, the effects of memory traces are by definition subtle; consequently, observing them directly will critically depend on discovering sensitive paradigms of measurement.

Despite these challenges, however, the arguments provided so far seem to indicate certain directions for further investigation. For instance, each of the sources of evidence for the account mentioned in this paper would greatly benefit from closer inspection. Furthermore, many more examples of memory traces in everyday interpretation processes need to be analysed and discussed and their role in lexical pragmatics in particular tracked in more detail. A new objective would be to uncover the effects of memory traces *in action*, as they influence a particular interpretation. In this paper, to illustrate the accumulation of memory traces and link them with particular interpretations, I have resorted to the slight artifice of looking at how different uses of a single word have built up over the *word's* lifetime. I am assuming that this is analogous to how memory traces accumulate both in a linguistic community's collective memory and in an individual's lifelong experiences with a word or expression. What is common to the three is that different situations of utterance affect ensuing interpretations and these interpretations are, by hypothesis, recorded; moreover, certain psycholinguistic effects, such as frequency, recency and salience, can be linked to the number of traces recorded, their temporal occurrence and their degree of availability.

An experiment could perhaps be designed to isolate an effect by setting up a situation and an utterance containing an expression that could be interpreted in two or more ways and manipulating the previous experiences with the expression to test what exactly tips the interpretational balance one way or another. It would surely be necessary to use artificial, made-up stimuli and to expose an experimental subject to it on multiple occasions. A clue as to how this could be successfully accomplished might be found in recent research on 'emotional memory'. Experimental and imaging findings in the cognitive and neuro-sciences point to a robust relation between emotion and memory (Kensigner & Corkin 2003, LaBar 2009). It seems that emotionally charged words are remembered better than neutral words. This could mean that fewer memory traces would be necessary for a subject to recall an emotionally polarised made-up stimulus and simultaneously, that the context and emotional content of the stimuli could be controlled in order to tip the interpretational balance one way or the other. Since our interpretational processes are very well-tuned to emotional valence, it stands to reason that memory traces of inferences using emotional stimuli would record both the linguistic form and the emotional content of the situation of utterance (positive or negative), as well as the inferences that were warranted. On subsequent testing, the degree to which the context and the inferences were recalled could be measured.

Research into memory traces, and particularly into the effects of memory traces on language processing, is only in its infancy. It is my aim to develop these ideas, in particular with regard to the interpretation of words in context, and to chart more closely the role of memory traces of previous utterances in the pragmatic process of lexical adjustment as set out within relevance theory.

Annexe

Figure 1 Ngram for ‘gene’ and ‘genes’:

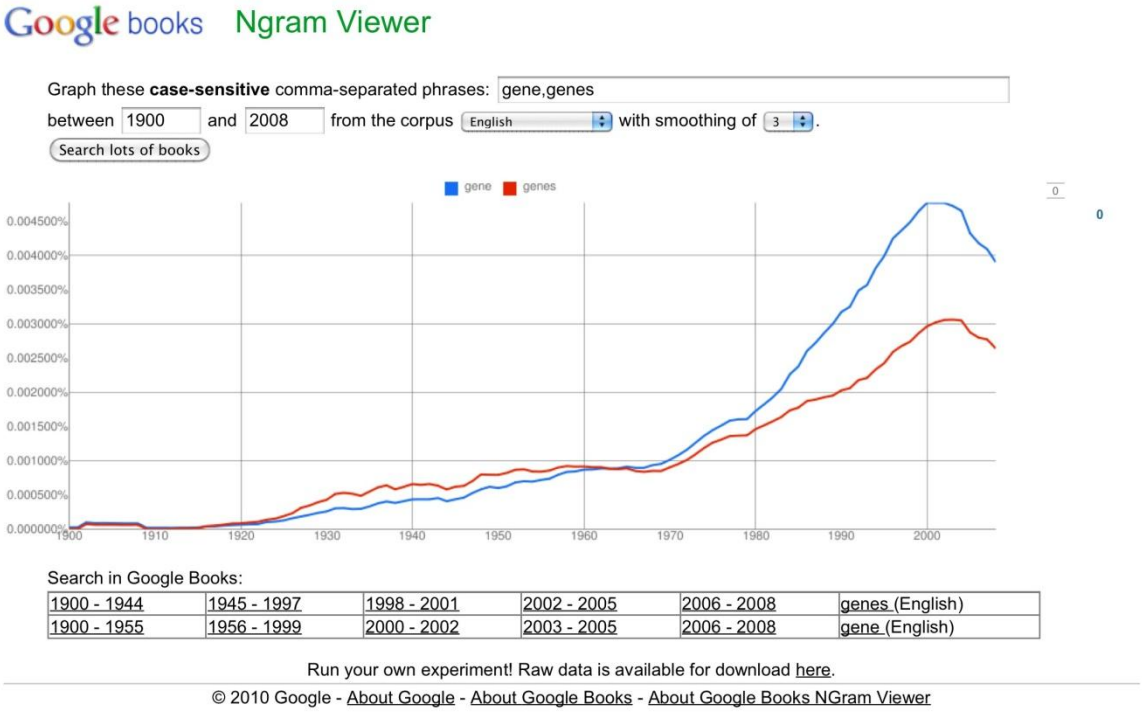

Figure 2 Ngram for ‘good genes’ and ‘bad genes’:

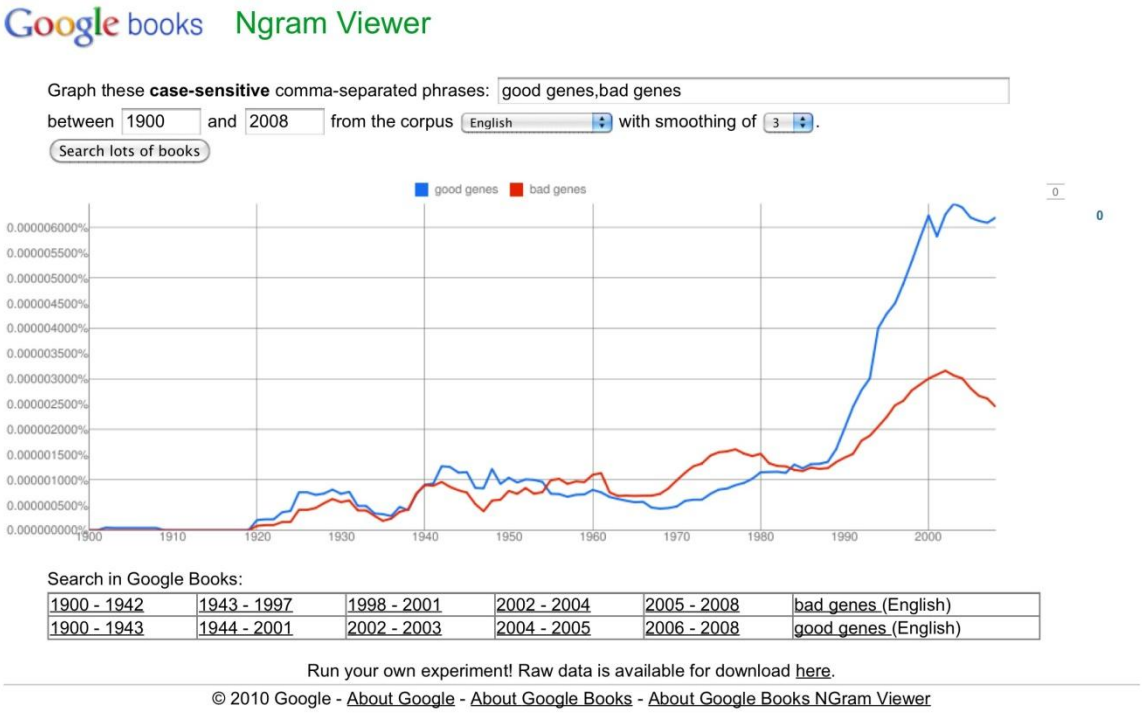

Figure 3 Ngram for ‘gene pool’:

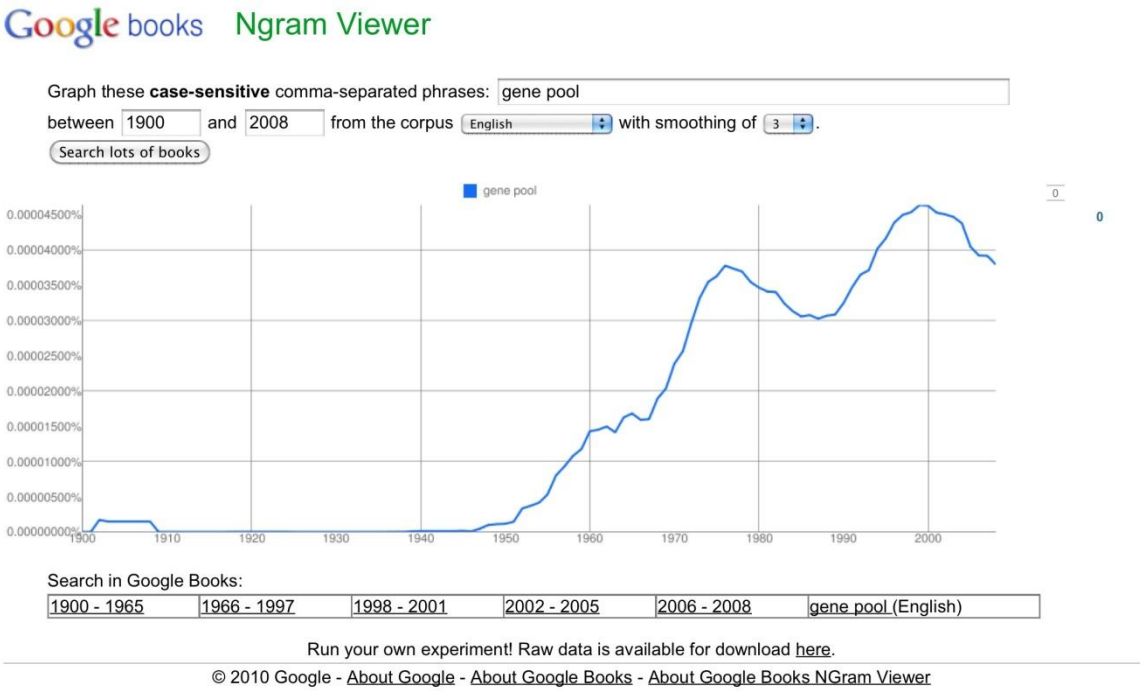

Figure 4 Ngram for ‘faulty genes’, ‘genetic disorder’, ‘genetic mutation’, ‘cancer gene’:

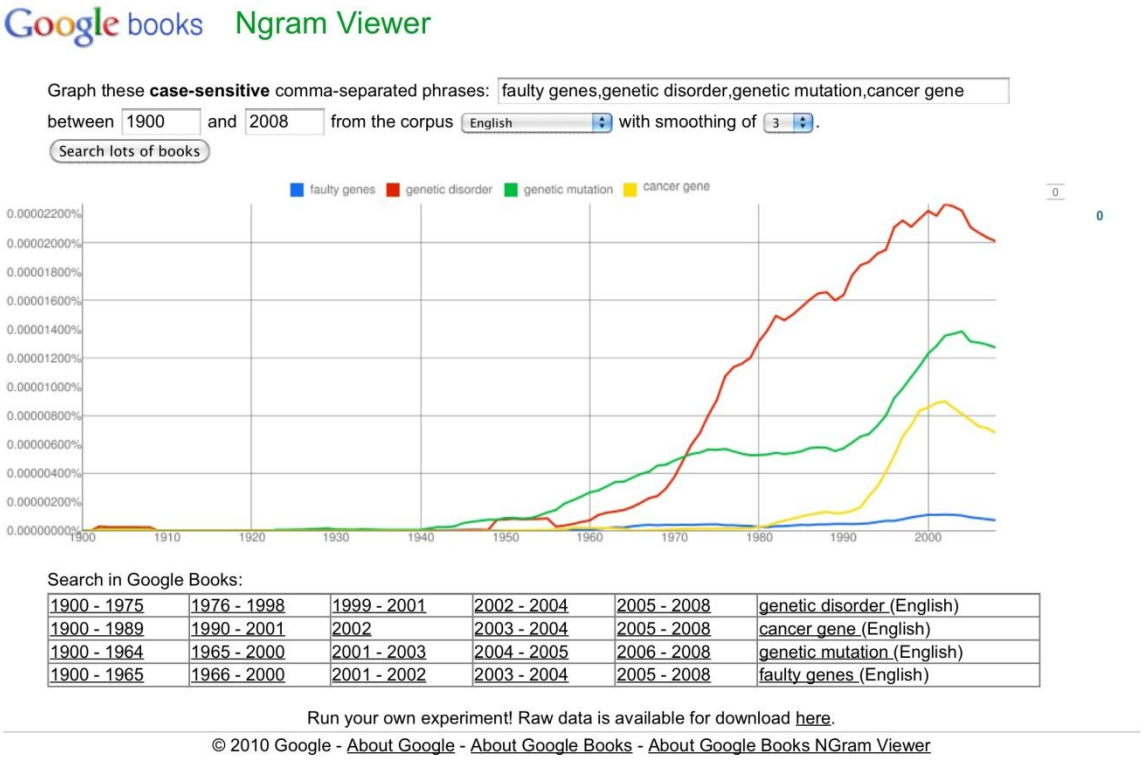

Figure 5 Ngram for ‘gay gene’:

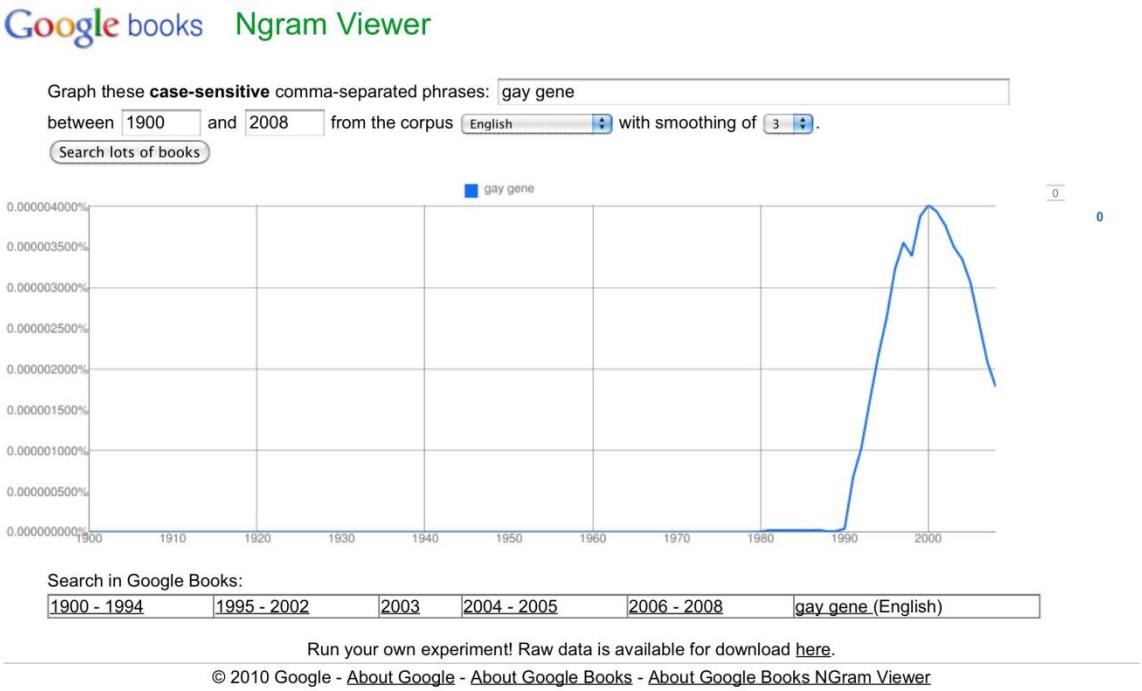

References

Ariel, Mira. (2008). *Pragmatics and grammar*. Cambridge: Cambridge University Press.

Bloom, Paul. (2000). *How children learn the meanings of words*. London: MIT Press.

Bybee, Joan L. (2000) The phonology of the lexicon; evidence from lexical diffusion. In Michael Barlow and Suzanne Kemmer (eds.). *Usage-based models of language*. Stanford : CSLI. 65-85.

Bybee, Joan L. (2010). *Language, usage and cognition*. Cambridge: Cambridge University Press.

Carston, Robyn. (1999). The semantics/pragmatics distinction: A view from relevance theory. In Ken Turner (ed.). *The semantics / pragmatics interface from different points of view*. Elsevier Science. 85-125.

Carston, Robyn. (2002a). *Thoughts and utterances : The pragmatics of explicit communication*. Oxford: Blackwell.

Carston, Robyn. (2002b). Linguistic meaning, communicated meaning and cognitive pragmatics. *Mind & Language, 17*(1-2), 127-148.

Carston, Robyn. (2002c). Metaphor, ad hoc concepts and word meaning –more questions than answers. *UCL Working Papers in Linguistics*. 14 : 83-105.

Carston, Robyn. (2009). The explicit/implicit distinction in pragmatis and the limits of explicit communication. *International Review of Pragmatics, 1*(1), 35-62.

Croft, William. Exemplar semantics. *unpublished manuscript*. Available at <http://www.unm.edu/~wcroft/Papers/CSDL8-paper.pdf>

Csibra, Gergely. (2010). Recognizing communicative intentions in infancy. *Mind & Language, 25*(2), 141-168.

Dennis, Simon. (2005). A memory-based theory of verbal cognition. *Cognitive Science, 29*(2), 145-193.

Grice, H. Paul. (1967). *Logic and conversation*. William James Lectures. Repr. in H. Paul Grice, 1989. *Studies in the way of words*. Cambridge, Mass.: Harvard University Press. 1-143.

Grice, H. Paul. (1989). *Studies in the way of words*. Cambridge, Mass.: Harvard University Press.

Hintzman, Douglas. L. (1986). Schema abstraction in a multiple-trace memory model. *Psychological Review, 93*(4), 411-428.

Hintzman, Douglas. L. (1988). Judgments of frequency and recognition memory in a multiple-trace memory model. *Psychological Review, 95*(4), 528-551.

Hopper, Paul J., and Traugott, Elizabeth Closs. (2003). *Grammaticalization* (2nd ed.). Cambridge: Cambridge University Press.

Horn, Laurence R., and Ward, Gregory L. (eds.). (2004). *The handbook of pragmatics*. Oxford: Blackwell.

Kensinger, Elizabeth A., and Corkin, Suzanne. (2003). Memory enhancement for emotional words: Are emotional words more vividly remembered than neutral words? *Memory & Cognition, 31*(8), 1169-1180.

Lakoff, George. (1987). *Women, fire, and dangerous things : What categories reveal about the mind*. Chicago: University of Chicago Press.

LaBar, Kevin S. (2009). Imaging emotional influences on learning and memory. In Frank Rösler, Charan Renganath, Brigitte Röder and Rainer H. Kluwe (eds.), *Neuroimaging of human memory. Linking cognitive processes to neural systems*. Oxford: Oxford University Press. 331-348

Levinson, Stephen C. (1983). *Pragmatics*. Cambridge: Cambridge University Press.

Michel, Jean-Baptiste, Shen, Yuan K., Aiden, Aviva P., Veres, Adrian, Gray, Matthew K., Pickett, Joseph P., The Google Books Team… (2011). Quantitative analysis of culture using millions of digitized books. *Science, 331*(6014), 176-182.

Murphy, G. (1997). Polysemy and the creation of novel word meanings. In Thomas Ward, Steven Smith and Jyotsna Vaid (eds.), *Creative thought: An investigation of conceptual structures and processes*. Washington D.C.: American Psychological Association. 235 -265

Nelson, Katherine. (1973). *Structure and strategy in learning to talk*. *Monographs of the Society for Research in Child Development 38*(Serial No. 149), 1-135.

Neath, Ian, & Surprenant, Aimée M. (2003). *Human memory : An introduction to research, data, and theory* (2nd ed.). Australia; London: Thomson/Wadsworth.

Nyckees, Vincent. (2003). La perspective médiationniste en linguistique. In Maryse Siksou (dir.). *Variation, construction et instrumentalisation du sens*, Paris: Lavoisier.

Nyckees, Vincent. (2006). Rien n'est sans raison : les bases d'une théorie continuiste de l'évolution sémantique. In *Aspects diachroniques du vocabulaire*, Danielle Candel and Françoise Gaudin (dir.), Publications des Universités de Rouen et du Havre : 15-88.

Pierrehumbert, Janet B. (2001). Exemplar dynamics: Word frequency, lenition and contrast. In Joan Bybee and Paul Hopper (eds.), *Frequency and the emergence of linguistic structure*. Amsterdam : John Benjamins. 137-157.

Pinker, Steven. (1994). *The language instinct : The new science of language and mind*. London: Allen Lane.

Raaijmakers, Jeroen G. W., and Shiffrin, Richard M. (1981). Search of associative memory. *Psychological Review, 88*(2), 93-134.

Recanati, François. (2003). *Literal meaning*. Cambridge: Cambridge University Press.

Recanati, François. (2010). *Truth-conditional pragmatics*. Oxford: Oxford University Press.

Scott-Phillips, T. C. (2010). The evolution of relevance. *Cognitive Science, 34*(4), 583-601.

Shiffrin, Richard M., & Steyvers, Mark. (1997). A model for recognition memory: Rem - retrieving effectively from memory. *Psychonomic Bulletin & Review, 4*(2), 145-166.

Sperber, Dan, and Wilson, Deirdre. (1986). *Relevance : Communication and cognition*. Oxford: Blackwell.

Sperber, Dan, and Wilson, Deirdre. (2002). Pragmatics, modularity and mind-reading. *Mind & Language, 17*(1-2), 3-23.

Tomasello, Michael. (2003). *Constructing a language : A usage-based theory of language acquisition*. London: Harvard University Press.

Tomasello, Michael. (2008). *Origins of human communication*. Cambridge, Mass.: MIT Press.

Tomasello, Michael. (2009). *Why we cooperate*. Cambridge, Mass.: MIT Press.

Traugott, Elizabeth Closs. (2004). Historical pragmatics. In Laurence R. Horn & Gregory Ward (eds.), *The handbook of pragmatics*. Oxford: Blackwell. 538-561.

Wedel, Andrew B. (2006). Exemplar models, evolution and language change. *The Linguistic Review*(23), 247-274.

Wilson, Deirdre, and Sperber, Dan (2004). Relevance theory. In Laurence R. Horn & Gregory Ward (eds.), *The handbook of pragmatics*. Oxford: Blackwell. 607-632.