

ОСНОВЫ ПСИХОЛИНГВИСТИКИ 1 курс САН, АН

SEMINAR IV

Tasks:

1. Read the text “Selective Processing in Text Understanding” and give definitions of the terms in bold.
2. Choose any piece of text (5–6 sentences), prove the coherence of this text. What types of inferences does the text include?
3. Prepare 3 more questions on the text to check the understanding of it by your group-mates.

SELECTIVE PROCESSING IN TEXT UNDERSTANDING*

It is commonly assumed that the goal of reading a piece of text is to end up with a mental representation of that text. Texts which are unclear may give rise to poorly structured mental representations, while well-written texts, if read properly, will give rise to well-structured representations which are in some sense the result of the comprehension process. Whatever the details of such a process, the text should induce a coherent representation in the mind of the willing reader. Of course, the mental representation will not be a verbatim copy of the text itself unless special rote-learning procedures are adopted. Rather, subjects typically retain the gist of a text. The point is easily appreciated that one might read and remember the gist of a novel, but not generally retain it verbatim. In short, most of what is remembered is the product of comprehending the text rather than learning it.

The problem is bound up intrinsically with which **inferences** are drawn on the basis of a text. Ideas which were not included in a message but which are captured by the internal representation of that message are called inferences. Since a text could support any number of inferences of varying degrees of plausibility, there must be some mechanism supporting the selection of some kinds of inference over others. The simplest idea is that only inferences necessary for the establishment of cohesion are drawn, but as we shall see, this criterion leaves room for interpretation, and there is evidence that elaborative inferences are made, and that inferences necessary for cohesion are not always drawn.

Coherence in text and mind

The term COHERENCE has been used in relation to texts and to the mental representation of texts. The idea of coherence in text is closely related to the notion of well-formedness. Thus Reinhart (1980) claims that it is coherence that distinguishes a text from a set of sentences which are unconnected. If a text is seen as an object, then it is clearly important that there are visible aspects of the text which represent the connections between sentences. These connectors are termed by Halliday and Hasan **cohesion devices**, because they hold the text together, giving it coherence. Halliday and Hasan distinguish five kinds of tie: **conjunction** (e.g., *and*, *but*, *because*), **coreference**, **substitutions**, **ellipsis**, and **lexical cohesion**.

The presence of surface markers of this type seems to be too weak a criterion for coherence in some instances: It does not rule out sets of sentences which we would not want to consider to be coherent in an everyday sense, such as (1).

(1) *I bought a Ford. The car in which President Wilson rode down the Champs-Elysee was black. Black English has been widely discussed. The discussions between the presidents ended last week. A week has seven days. etc.*

The presence of **coreference** links (Ford–The car, black–Black, etc.) is not enough to specify what could reasonably be called a coherent text. (Enkvist, 1978, calls such examples **pseudo-coherent**, because they do hold together in any fashion.) Reinhart (1980) suggested that the following conditions are a better specification of what might make a text coherent. The discourse should be:

1. **Connected.** The sentences (clauses) of a text will be formally connected, in that each adjacent pair is either referentially linked, or linked by a semantic sentence connector.

2. **Consistent.** Each sentence has to be logically consistent with the previous sentence.

3. **Relevant.** Each sentence must be relevant to an underlying discourse topic and to the context of the utterance.

The first criterion of formal connection means that cohesion markers should be present. Again, it is possible to look for these in a text itself. The second is more problematic, for while it is possible to test for inconsistency (e.g., contradictions) in a text, it is not possible to test for consistency. Presumably, the second criterion is that each successive sentence is not inconsistent. Moving to relevance, it is not clear how relevance can be found in a text. Rather, it must be the product of inference and so makes Reinhart's definition partly dependent on a reader.

We have seen that surface cohesion markers are not sufficient for a text to be called coherent. Are they necessary, as Reinhart's formulation requires?

The answer is *no*, because it is possible to have a text which we would want to call coherent, but which does not depend on the presence of any markers.

(2) *At dinner last night, John burnt his mouth. The soup was too hot.*

These utterances are connected only by way of an inference: The second sentence explains how John came to burn his mouth. Such constructions are commonplace in normal, easily comprehended writing.

In summary, coherence does not seem to be a property of text; rather it is a property of the mental representation (interpretation) of a text. A text can yield a coherent mental representation even when it does not contain appropriate cohesion markers. From now on, we shall suppose that coherence occurs in the mind of the reader and is the establishment of a mental representation which consists of a connected set of ideas based on appropriately interpreted discourse. Both connectivity and appropriate interpretation are matters of degree, and we shall examine some influences on selective processes which underlie coherence, starting with the issue of inference control.

Selective inference in the service of coherence

It is clear that inferences are necessary for coherence on any account of text processing. But a text can lead to an infinite number of inferences, and it is therefore important to determine just which inferences are made during reading. By far the

simplest approach to this question has been made within a framework that divides inferences into those which are necessary to support the construction of a coherent representation of the text in the mind of the reader (henceforth **necessary inferences**) and those which are not necessary, but which are mere elaborations (henceforth **elaborative inferences**). The essence of the distinction is illustrated by the materialism.

(3) *No longer able to control his anger, the husband threw the delicate porcelain vase against the wall. It cost him well over one hundred dollars to replace.*

In this case, the inference that the vase broke is a **plausible inference** and must be made in order to realize an important link between the first and the second sentence. The inference is therefore necessary. Note that the inference is not necessary in the logical sense; rather, it is necessary for coherence. This of course raises the question of what is coherence. For the moment, we note that in order to link the two sentences, an explanation for the second is required, and the most plausible explanation is that the vase broke. Also, note that this inference does not need to be drawn until the second sentence is encountered.

Consider a second example.

(4) *No longer able to control his anger, the husband threw the delicate porcelain vase against the wall. He had been feeling angry for weeks, but had refused to seek help.*

Here, any inference that the vase broke would not be obviously useful for cohesion. Rather, the second sentence links to the previous one by providing an explanation and elaboration of it. If the inference that the vase broke had indeed been made, it would be merely **elaborative**. The contrast pair illustrates the distinction between forward and backward inferences. A **forward inference** is one which is made before the text requires it in order to establish a cohesive link. It is therefore elaborative at the time it is made. A **backward inference** is one which is drawn to link a previous text fragment with a later one. This distinction has been used as the foundation of an account of which inferences will and will not be made during reading; Only inferences necessary for a coherent interpretation will be made, that is, necessary inferences.

SEMINAR V

Tasks:

1. Read the text “Second Language Acquisition” and give definitions of the terms in bold.
2. Compare the processes of first and second language acquisition.
3. Explain approaches to second language acquisition.
4. Prepare 2 more questions on the text to check the understanding of it by your group-mates.

SECOND LANGUAGE ACQUISITION*

Language teachers were once swayed by an argument that the most natural way of acquiring a second language was to emulate the process of first language acquisition. However, modern practice reflects a realisation that the two situations are very different. Compared with an infant acquiring its first language, an adolescent or adult acquiring a second:

- has less time for learning;
- is cognitively developed – possessing concepts such as causality or aspect;
- is primed by experience to seek for patterns in data and so responds to input analytically;
- already has a first language, which provides a lens through which the second is perceived;
- has access to a language of explanation, and is therefore capable of understanding (even if not applying) theoretical explanations;
- is accustomed to expressing their personality in L1, and may find their limited powers of expression in L2 a chastening experience;
- has pragmatic experience of a range of social circumstances in L1 and extensive world knowledge.

Second language acquisition: approaches

Linguistic. In the linguistic tradition, research and analysis are usually based on the assumption that the acquisition of our first language is supported by an innately acquired **Universal Grammar** (UG). In this context, six different positions can be adopted in relation to second language (L2) acquisition:

- The L2 learner retains access to the same UG as was available for L1.
- UG supports L1 acquisition only, and is then lost. The process of L2 acquisition is therefore very different.
- UG supports L1 acquisition only, but L2 acquisition is able to model itself upon residual traces of our experience of acquiring L1.
- UG survives until early adolescence and then decays. There is thus a **critical period** for second language acquisition.
- Universal principles are retained and continue to guide L2 acquisition. However, the user's parameters are adjusted to L1 values and therefore need to be re-set to L2 values.
- Universal linguistic criteria (perhaps based on **markedness**) determine which linguistic concepts are the easiest to acquire and which are the most difficult.

One research approach is **theory-driven**: with researchers applying L1 linguistic theory (usually Chomskyan) to second language learning and use. Researchers often ask subjects to make L2 **grammaticality judgements**, which are said to tap in to their competence. A second approach is **observational**, with researchers obtaining longitudinal evidence of the order in which particular areas of L2 syntax are acquired and the variants which the learner employs at different stages. The data is then compared with patterns of L1 acquisition and interpreted in a framework of grammatical theory and of concepts such as **parameter-switching**.

Cognitive. A theoretical assumption is adopted that language is part of general cognition. It is therefore valid to trace parallels between the techniques adopted by a

second language learner and those employed in acquiring other types of **expertise**. Cognitivist accounts of second language acquisition (SLA) cover both acquisition (how learners construct a representation of L2) and use (how they employ their knowledge of L2 in order to communicate).

There has been discussion of the relationship between **explicit knowledge** gained in the form of L2 instruction and **implicit knowledge** gained by acquisition in an L2 environment. The former is likely to provide linguistic information in an **analysed form**, while linguistic information that is acquired naturalistically is often in the form of unanalyzed **chunks**. A similar contrast exists between circumstances where accuracy is a requirement and the use of L2 may therefore be subject to careful **control** – and others where fluency is called for and it is desirable to aim for a high degree of **automaticity**.

If linguistic information is initially acquired in explicit/controlled form, then it has to be reshaped in order to support spontaneous spoken performance in the target language. A case has been put for treating second language acquisition as a form of skill acquisition not unlike learning to drive or becoming an expert chess player.

Much attention in SLA research has been given to *transfer*, the effect of the native language upon performance in L2. Early accounts of transfer drew upon behaviourist theory. Language use was depicted as habitual behaviour, with the habits of the first language having to be replaced by those of the second. Current models treat the issue in terms of the relative cognitive demands made by L2 as against those made by L1. These might reflect the extent to which a grammatical feature is *marked* in one language but not the other. Or it might reflect differences between languages in the importance attached to linguistic cues such as word order, inflection or animacy.

Another approach considers L2 acquisition in terms of the way in which the learner's language develops. At any given stage, a learner is said to possess an **interlanguage**, an interim form of self-expression which is more restricted than the native-speaker target but may be internally consistent. Longitudinal studies have examined changes in interlanguage: for example, the different forms used to express the interrogative or negative. Most learners appear to proceed through similar stages; an explanation is found in the relative difficulty of the cognitive operations involved rather than in constraints imposed by UG.

Some commentators have suggested that SLA involves continual **restructuring** in which knowledge structures are reorganised in order to accommodate new linguistic insights. The *Multi-Dimensional Model* sees restructuring as part of a developmental process in which two important cognitive factors determine a learner's performance. The first is the developmental stage that the learner has reached, development being represented as the gradual removal of limitations upon the linguistic structures that the learner is capable of forming. The second is the extent to which each individual engages in a process of simplification, reducing and over-generalising the L2 grammar so as to make it easy to handle.

Another line of research has concerned itself with the learner as an active participant in the learning process. A non-native speaker's chief goal in L2 communicative contexts is to extract meaning, but the question arises of whether they also have to specifically '*notice*' (direct attention to) the form of the words that is used in order to add to their own syntactic repertoire.

A further area of study that is relevant to psycholinguistics considers the way in which second-language learners handle communicative encounters, and the strategies they adopt in order to compensate for their incomplete knowledge of the lexis and grammar of the target language. There has been interest in the **communication strategies** adopted in spoken production, but rather less is known about the strategies employed in extracting meaning from written or spoken texts.

SEMINAR VI

Tasks:

1. Read the text “Bilingualism” and give definitions of the terms in bold.
2. Name types of bilingualism mentioned in the text.
3. Present the hypotheses of bilingual acquisition.
4. Speak about three aspects of bilingualism that Psycholinguistics takes into account.
5. Prepare 3 more questions on the text to check the understanding of it by your group-mates.

BILINGUALISM*

Bilingualism is the ‘habitual, fluent, correct and accent-free use of two languages’ (Paradis, 1986) – or of more than two languages. However, on this definition, few individuals qualify as complete bilinguals. It often happens that a bilingual is not equally competent in different aspects of the two languages: they might, for example, have a more restricted vocabulary in one than in the other or might exhibit different abilities in respect of speaking, listening, reading and writing. Furthermore, many bilinguals use their languages in ways that are domain-specific: one language might be used in the family and one reserved for educational contexts.

An early account of bilingualism (Weinreich, 1968) proposed three types. In **compound bilingualism**, conditions in infancy are equally favourable for both languages, and words in both are attached to one central set of real-world concepts. **Coordinate bilingualism** occurs when conditions in infancy favour one language over the other; the consequence is that the infant develops two independent lexical systems, though meanings overlap. **Subordinate bilingualism** occurs when the second language is acquired some time after the first, and so remains dependent upon it.

These categories have proved difficult to substantiate. However, the stage at which the two languages are acquired remains an important consideration in recent accounts, which often distinguish **simultaneous bilingualism** (both languages acquired concurrently), early **successive or sequential bilingualism** (both languages acquired in childhood but one preceding the other) and late bilingualism (the second language acquired after childhood).

Simultaneous bilingualism arises during ‘primary language development’, which commentators regard variously as occurring during the first three or the first five years of life. Exposed to two languages, infants initially mix vocabulary and syntax from both. In naming objects and actions, they often adopt the first word they encounter, regardless of which language it comes from; though in their morphology they may exhibit a preference for the less complex of their languages.

The **unitary language hypothesis** concludes that these infants start out with undifferentiated language systems. They begin to distinguish between the two sets of data by restricting each language to particular interlocutors, situations or pragmatic intentions. At the next stage of development, the infant distributes its vocabulary between two separate lexical systems, and becomes capable of translating words from one language to the other. However, the same syntactic rules are usually applied to both systems. In a final stage, the languages become differentiated syntactically, and mixing declines.

An alternative **separate development hypothesis** maintains that the two languages are distinguished from the start by the infant and that the phenomenon of mixing simply shows two incomplete systems operating in parallel.

Simultaneous bilingual acquisition appears to follow a very similar path to monolingual acquisition. There is no evidence that the acquisition process is delayed when more than one language is involved, though early vocabulary levels may be slightly lower in bilingual children. Nor do similarities between the two target languages appear to assist acquisition: an English-French bilingual does not develop language faster than an English-Chinese one.

In successive bilingualism, there is much greater variation between individuals. The time of acquisition of the second language (during the primary period/before puberty/in adulthood) may be a factor; while mastery of the later language may be limited to certain domains. In some cases, the acquisition of the later language is additive, resulting in the use of two systems in parallel. In others, the effect may be subtractive, with the later language replacing the first in some, many or all domains. The acquisition of a second language by an immigrant may even lead to the attrition of the original language if the speaker has to communicate mainly or exclusively with members of the host community.

A distinction is made between adult bilinguals who are **balanced** and those for whom one language is **dominant**. A balanced bilingual has been represented (Thiery, 1978) as somebody who is accepted as a native speaker in two linguistic communities at roughly the same social level, has learnt both languages before puberty and has made an active effort to maintain both of them. Fully balanced bilinguals are said to be rare.

Bilinguals may not always be aware of which language is their dominant one, and it has not proved easy to establish dominance. One approach has been to ask individuals which language they are conscious of having spoken first; though many recall acquiring both simultaneously. Another is to ask individuals to express a preference for one of their languages. There may be a relationship between dominance and anxiety, with the dominant language resorted to in times of stress or tiredness. Experimental methods to determine dominance have included rating bilinguals’ language skills across languages, self-rating questionnaires, fluency tests, tests of flexibility (checking the ability to

produce synonyms or draw upon a range of senses for a particular word), and dominance tests where bilinguals read aloud cognates which could be from either of their languages. Even where dominance is established, the situation may not remain constant: the relationship between languages may shift as the individual's linguistic needs and circumstances change.

Psycholinguistic research has especially considered three aspects of bilingualism:

- *Storage*. Are the two languages stored separately in the user's mind or together? Possible evidence for separate stores comes from the phenomenon of code-switching where, often prompted by a change of topic, bilinguals shift with ease between their languages. However, it has been noted that **code-switching** takes place almost exclusively at important syntactic boundaries (the ends of clauses, phrases, sentences) and that these boundaries are often common to both languages.

- *Cross-linguistic influence*. Is performance in one language affected by the user's knowledge of the other? Constituents from one language are sometimes introduced into an utterance involving the other in an effect called **code-mixing**. The transfer can occur at many different linguistic levels: phonological, orthographic, morphological, semantic and phrasal, and can involve structural features such as word order. Cross-linguistic lexical influence is seen in borrowing, where a word is transferred from one language to the other with its pronunciation and morphology adjusted accordingly.

- *Costs and benefits*. Does being bilingual have positive or negative consequences? The consequences might be linguistic, educational, cultural, affective or cognitive. In terms of linguistic development, a balance theory suggests that the possession of two languages makes increased demands on working memory, and thus leads to some decrement in proficiency in at least one of the languages. There has been little evidence to support this. An alternative view is that there is a language-independent 'common underlying proficiency' which controls operations in both languages. Early studies in bilingual contexts such as Wales led to the conclusion that bilingualism had an adverse effect on educational development; but these are now generally discredited. Recent research has tended to stress the positive outcomes of bilingualism: it appears that bilinguals may benefit from more flexible thought processes and from a heightened language awareness.

SEMINAR VII

Tasks:

1. Read the text "How good is 'Good Enough'?" and give definitions of the terms in bold.
2. What factors influence our language processing?
3. Give your own examples of ambiguous clauses.
4. Prepare 2 more questions on the text to check the understanding of it by your group-mates.

HOW GOOD IS “GOOD ENOUGH”?*

Previously linguists have assumed that **comprehenders** are trying to process the language that they hear or read as fully as possible – that they interpret each word completely and build complete syntactic structures (sometimes even more than one) when they encounter sentences. But, recent work has cast some doubt on just how complete these representations really are.

Suggestive evidence comes from the early 1980s and the Moses illusion (Erickson & Matteson, 1981), named after one of the most commonly cited sentences that exemplifies it. This illusion centers on how people initially respond when asked to say whether certain sentences are true or false, such as:

1. Moses put two of each sort of animal on the Ark

Most people, when they first encounter this sentence, answer that it is true. However, the problem is that it was not Moses, but rather Noah who put animals on the Ark. So, how do people who are fully familiar with the story of Noah and the Ark fail to notice this? Clearly, they are not paying attention to all the details of the question. Other similar effects are found when people are asked questions like, “After an air crash, where should the survivors be buried?” Barton and Sanford (1993) found that half of participants who were asked this question responded with an answer that indicated that they should be buried where their relatives wished them to be buried. These participants failed to notice that survivors are living people.

All of these cases involve some pretty complicated situations. For the Moses illusion and the survivor case, the problem hinges on not paying attention to the meaning of a particular word with respect to the bigger context. But, they also depend on trusting the assumptions of the question or statement, particularly in the survivor question: the question presupposes that survivors should be buried, and so people may sometimes be lulled into going along with this. However, the larger point that the responses to these types of sentences make is intriguing: if there are times when people are not fully processing language, when are these times? Essentially, when else are people not fully processing language?

This is not a trivial issue – until recently it was taken as established and assumed by most theories of **language processing** that people process language as fully as possible, in real time. The debates in language processing have really centered on what kinds of information or cues get used during processing (and when they get used), not whether people are not actually processing language fully. Thus, if it really is the case that people are not always processing language fully, but rather relying on shallow processing strategies to create “good enough” representations of what they are reading or hearing, then we will need to reconsider almost every theory of how language is processed in real time.

The researchers thus ran two further experiments using a special class of verbs that are well known to require a reflexive interpretation in that there is no object given. What does this mean? These verbs (called **reflexive absolute transitive**, or RAT, verbs) generally deal with actions related to personal hygiene, such as wash, bathe, shave, and

dress. If someone says “Mary bathed” then we interpret the unsaid object as a reflexive (herself). So, *Mary bathed* means *Mary bathed herself*. If there is an object given, then the subject does the action on that verb: *Mary bathed the baby* means that *Mary did not bath herself, but gave a bath to a baby*. These verbs are useful here because if there is no object provided in the sentence, people should give the reflexive interpretation. So, in “While Anna dressed the baby that was small and cute spit up on the bed,” Anna dressed herself and the baby spat up. They compared these types of verbs with verbs like before, in which the object is optional, but does not have a reflexive interpretation. However, even for these RAT verbs, participants gave incorrect “yes” answers to questions like “Did Anna dress the baby” 65.6% of the time. The optionally transitive verbs were even worse—with 75% incorrect answers.

These are intriguing results – could it be that people were really not completely correcting their initial mistakes while reading? It appears so. Of course, this opens up a whole new issue – are people completely interpreting what they read the first time?

It is plausible that people may process less fully when in a challenging linguistic environment. This possibility was addressed by Ferreira (2003). She tested people’s ability to assign roles to the referents in very simple sentences. For example, if we have a sentence like “The mouse ate the cheese” we know that the mouse did the eating and it was the cheese that was eaten. In linguistic terms, the verb eat assigns two roles – an agent role (to the subject, mouse) and a theme or patient role (to the object, cheese). Ferreira tested for whether people were able to accurately assign roles in these simple sentences by having the participants listen to the sentences and then respond verbally to a prompt word that asked to them to say who/what was the “acted-on” thing in the sentence or what was the “do-er” in the sentence. So, if you heard the sentence above and then saw “Do-er” the correct answer would be for you to say “mouse.” Ferreira coded accuracy and also looked at how long it took people to respond. She was particularly interested in whether people would be less accurate for relatively simple, unambiguous sentences that nonetheless had the order of the agent/patient reversed. In the first experiment, she compared active and passive sentences. The passive version of our earlier sentence would be “The cheese was eaten by the mouse.” She also looked at sentences like these that were nonreversible (e.g., cheese cannot eat mice), sentences that were reversible but implausible if reversed, and sentences in which either referent could be the agent or patient/theme. She gave these sentences in both active and passive versions, and both active and passive tested assigning the nouns to both roles. Examples of these sentences in active and passive are as follows:

2. *Nonreversible, Plausible: The mouse ate the cheese/The cheese was eaten by the mouse.*

3. *Nonreversible, Implausible: The cheese ate the mouse/The mouse was eaten by the cheese.*

4. *Reversible, Plausible: The dog bit the man/The man was bitten by the dog.*

5. *Reversible, Implausible: The man bit the dog/The dog was bitten by the man.*

6. *Symmetrical, Version 1: The woman visited the man/The man was visited by the woman.*

7. *Symmetrical, Version 2: The man visited the woman/The woman was visited by the man.*

Ferreira found that while overall accuracy was high, as one might expect, there were still some interesting patterns in the data, as well as some conditions in particular that were surprisingly low. For all three sentence types, passives took longer to answer and were answered incorrectly more than actives.

From a certain perspective, passives are more complicated than active sentences and so perhaps it is the case that passives are more difficult simply because they are more complicated.

Ferreira argues that comprehenders do not necessarily fully process language as they encounter it, but instead rely on heuristics to provide a “good enough” representation. In English, it is so often the case that the noun before the verb is the agent of the verb and that any noun following the verb is not the agent, that English speakers can simply assign agent status to the first noun.

The degree to which people bother to process language fully could be under strategic control. This does not mean that we consciously decide to process or not process fully (though it could). Instead, perhaps we only process language as fully as necessary for the needs of the current communicative situation (Ferreira, Bailey, & Ferraro, 2002).

Further evidence in support of good enough hypothesis comes from Swets, Desmet, Clifton, and Ferreira (2008), who tested specifically whether task demands could influence depth of processing. They focused on a previous finding concerning the attachment of **ambiguous** clauses, such as in (8):

8. The maid of the princess who scratched herself in public was terribly humiliated.

The sentence is fully ambiguous with respect to who did the scratching: it could be the maid or the princess.

Previous work (Traxler, Pickering, & Clifton, 1998; van Gompel, Pickering, & Traxler, 2001; van Gompel, Pickering, Pearson, & Liversedge, 2005) found that with sentences like these, readers actually took a shorter amount of time to read the ambiguous sentence compared with the unambiguous versions. The details of the various accounts of this effect are not important for our present purposes, but they assume, in line with accounts of other phenomena, that readers attach the ambiguous phrase to rest of the sentence when they encounter it. But, perhaps readers don't actually do this – perhaps they do not try to interpret who scratched themselves unless they have to. Swets et al. (2008) wanted to test this hypothesis: that task demands (e.g., knowing that attachment was required or not required to correctly answer questions about the sentences) would influence whether readers actually did bother to interpret (or attach) phrases like “who scratched herself” to the maid or princess. That is, would readers underspecify the representation of the sentence if there was no reason to fully interpret it? They had people read the same sets of sentences, but for some readers all of the questions required a full interpretation of the sentence (e.g., Did the princess scratch in public?) while for other readers the questions were more superficial (e.g., Was anyone humiliated?). Swets et al. (2008) found that the pattern of reading did in fact change depending on the type of questions asked. These results suggest that readers may only interpret sentences as fully as necessary for the task, and supports the idea that we may have underspecified or “good enough” representations of language unless it is necessary

to create a more detailed interpretation of a sentence. Why have “shallow” or underspecified processing? Sanford (2002) suggested that a system with finite resources should be able to allocate those resources flexibly – that is, just as in other systems (e.g., like vision) we can’t pay attention to everything all the time, and so we process fully only what we need to.

This is the information structure of a sentence and linguists have long been interested in both how languages encode information structure and how information structure interacts with meaning (Lambrecht, 1994). Recently, psycholinguists have been interested, too, in how information structure influences language processing.

Swets et al. show that task is an influence – knowing that you need to know a particular piece of information.

So we’ve seen that we process language less fully than previously thought, but that this processing depth appears to be dynamic – we can process language more completely when the task requires it, or if our attention is drawn to a particular part of the sentence by linguistic, pragmatic, or even typographical cues.