Thematics : Interdisciplinary Studies,
M. Louwerse & W. van Peer (Editors).
Amsterdam: Benjamins.
2002, pp. 157-170.

On the notions of theme and topic in psychological process models of text comprehension

Walter Kintsch
Department of Psychology
University of Colorado

Running Head: Process models of comprehension

Abstract:
Latent semantic analysis (LSA) is used to define the theme of a text and to generate summaries automatically. The theme information – the already known information - in a text can be represented as a vector in semantic space; the text provides new information about this theme, potentially modifying and expanding the semantic space itself. Vectors can similarly represent subsections of a text. LSA can be used to select from each subsection the most typical and most important sentence, thus generating a kind of summary automatically.

Keywords:
process models of comprehension; mental representation; macrostructure; summaries; latent semantic analysis.

Distinctions like given-new, topic-comment, or theme-rheme have played an important role in psycholinguistic investigations of sentence processing (e.g., Havilland & Clark, 1974; Gernsbacher, 1990). On the other hand, theme and topic are not basic concepts in psychological models of text comprehension, at least not explicitly. However, certain features in these models correspond, not always directly, to the notions of theme and topic. I shall describe these features here as a first step in exploring the analogies and correspondences between psychological process models of language understanding and the linguistic notions of theme and topic.

I shall focus here on a specific process model of text comprehension, the construction-integration (CI) theory of Kintsch (Kintsch, 1998; based on earlier work of Kintsch & van Dijk, 1978; van Dijk & Kintsch, 1983; Kintsch, 1988; for a broader discussion of the psychological literature in this area see Gernsbacher, 1994). Only those features of the construction-integration theory directly relevant to the notions of topic and theme will be discussed here.

The construction-integration model of text comprehension attempts to simulate the computations involved in the construction of a mental representation of a text in human comprehension. Fundamental to the theory is a distinction between different aspects of the mental representation of a text. The <u>textbase</u> comprises that part of the mental representation that is directly derived from the text. Naturally, a great deal of knowledge on the part of the comprehender is necessary to construct a textbase, both linguistic knowledge and general world knowledge, but this is knowledge employed in the service of understanding what is directly contained in the text. The mental representation of a text, however, comprises more than this textbase - the comprehender's prior knowledge is used to fill in missing links in the text and to elaborate what the text actually says. Thus, the final product of comprehension, the <u>situation model</u>, is a mixture of structures that are derived directly form the text and structures that have been added from prior knowledge. Just what this mixture will be depends on numerous factors, which are being explored in the psychological literature on text comprehension (e.g. Zwaan & Radvansky, 1998)

Textbases represent the meaning of a text both at a local and a global level. Specifically, in the CI model a textbase represents the meaning of the phrases and sentences of a text by means of an interlinked series of propositions. Propositions in this use of the term correspond roughly to (simple) sentences, but are meaning units, a level removed from the words used to express them linguistically. In addition to this local propositional representation, there exists also a global propositional representation in terms of <u>macropropositions</u> (van Dijk, 1972). Macropropositions are organized hierarchically, subsuming ever-larger units of texts. Thus, macropropositions represent the meaning of paragraphs, sections of a text, clusters of sections, and eventually the whole text itself. They are formal counterparts of the intuitive notions of gist and summary.

Macrorules are linguistic operators that permit to trace the construction of a macroproposition. Van Dijk & Kintsch (1983) discuss these rules in some detail. There are three rules: deletion, generalization, and construction. Deletion deletes propositions from the textbase that are not macrorelevant; generalization substitutes a general concept for a number of specific ones; construction substitutes a general event for a sequence of specific ones. While these rules are not powerful enough computationally to generate macrostructures because their conditions of application are not fully specified, they can be used after the fact to explain how a particular macrostructure was generated. For instance, one can trace out how an ideal summary was obtained for a text, or show why some element that someone included in a summary does not belong into a summary because it is a detail that should have been generalized. Macrorules, therefore, can provide a post-hoc account of a how a summary was generated, but they cannot generate a summary, because of their vagueness. Hence they are inadequate as computational mechanisms.

Macrostructures pertain to texts. The global organization of the mental representation of the text, however, may deviate in certain respects from the macrostructure of the text because situation models are jointly determined by the content and organization of the text and by the content and organization of the comprehender's background knowledge. A knowledgeable and goal-oriented reader may override the author's intended macrostructure and choose his or her own idiosyncratic organization for the mental representation of a text.

Texts are linked with relevant areas of the comprehender's background knowledge in such a way that information in the text provides direct access to prior knowledge. The theory of long-term working memory (Ericsson & Kintsch, 1995; Kintsch, 1998) describes how this linkage is achieved and under what conditions it will be successful or not. Basically, it is the case that comprehenders will automatically form links between new information in a text and relevant background knowledge in areas where they have a great deal of well-structured background knowledge and experience. The successful linkage between information in a text and background knowledge is critical for deep understanding of the text and for learning from text, that is, the ability to employ the text information in novel ways and integrate it with prior knowledge. Without such integration the mental representation of a text remains inert - sufficient to reproduce the text, but not to use the information for other purposes. A reader whose comprehension processes result merely in a textbase will be able to recall and recognize the text for a time, but his or her future thinking, problem solving and comprehension will be unaffected thereby. On the other hand, a reader who achieves a well-integrated situation model will be able to use the information later when it becomes relevant, even when the text itself may be forgotten.

The construction-integration theory of text comprehension thus focuses on two aspects of text comprehension: the text itself, which yield the textbase component of the mental representation, and the background knowledge of the comprehender, which when integrated with the textbase yields a situation model for the text in question. In this model, there are two kinds of global features that characterize the mental representation of a text: the macrostructure of the text and the situational and world knowledge the text is linked to. In order to describe how knowledge can be modeled in the framework of psychological theories of text comprehension, a brief discussion of Latent Semantic Analysis's necessary.

Latent Semantic Analysis: LSA.

LSA provides a way to simulate human verbal knowledge (e.g., Landauer; & Dumais, 1997; Landauer, Foltz, & Laham, 1998). It is a fully automatic procedure using standard mathematical techniques to analyze a large corpus of digitized text. The analysis starts with collecting data on word use: counts of which words are used in which contexts. Such word frequency data are affected by many chance factors, such as an author's decision to use this particular word rather some other alternative. LSA extracts word meaning from such data by disregarding the accidents of word use in specific contexts and focusing on what is common in all contexts. A matrix algebra technique called Singular Value Decomposition allows us to partition the information[1] about word use into two components: the semantic essence underlying word use and the information associated with specific contexts. The latter is discarded, and the former provides the semantic representation of LSA: a high-dimensional space (typically 300 dimensions) in which the semantic content of words, sentences, and whole texts can be represented as vectors. Such vectors are simply lists of 300 numbers and can only be interpreted by comparison with other vectors. That is, if I want to know whether the vector I have computed for the meaning of "mother" is meaningful or not, I need to compare it to related words like "father", "child", or "woman" to which it should be similar, and to unrelated words, to which it should not be similar. The angle between two vectors provides an index of their similarity. The angle (in 300-dimensional space) is measured by the cosine between two vectors: two vectors that are almost identical will have a cosine close to 1, whereas two vectors that are orthogonal to each other will have a cosine of 0. Thus, the cosine behaves somewhat like the more familiar correlation coefficient.

[1] The term information is used here in its non-technical sense.

LSA learns the meaning of words by noting in which contexts these words are uttered – a bit like children who acquire word meanings not through explicit definitions but by observing how the word is used. But while LSA starts with word-use statistics, it ends up with something quite different and much more powerful: a true semantic representation, a space that captures the essential semantic relationships. For instance, words that are similar in meaning are often sued in different contexts. Thus, the singular and plural forms of nouns are (which are different words as far as LSA is concerned) are usually not employed both at the same time. If someone talks about a "mountain" he doesn't also talk about "mountains." In the LSA space, however, the correlation between "mountain" and "mountains" is quite high (. 84 in this case), because although the two words don't often appear together, both of them tend to be used in similar contexts. Hence LSA infers that the two are similar in meaning.

We assume that people, as they interact with the world they live in, keep track of co-occurrences and perform an operation analogous to dimension reduction in LSA on these co-occurrence statistics. Of course, what matters for real people are not just words, but also images, noises, feelings, events, and actions. Thus the raw experiential base is far richer for people than for LSA, which only can deal with digitized texts, and LSA spaces are therefore not equivalent to human semantic spaces. Nevertheless, if LSA is trained on a large enough and sufficiently varied corpus of texts, the results should come reasonably close to human prototypes. People write about much of what they know. The space used in the example below is based on a corpus of about 11 million words in texts read by typical American school children in grades 3 to college. A real child, of course, not only reads as much as that in this period of life, but listens to a lot more, and has rich additional sources of information from his or her interaction with the world which the brain manages to integrate in ways that we don't understand with the verbal material into one knowledge base of which the purely text-based LSA space is merely a pale reflection. But it is a reflection nevertheless, and a lot better than we can do with other techniques.

If we represent a text (the words, sentences, paragraphs, sections, the whole text) as vectors in the LSA space we obtain a representation of what the reader already knows about that text. The textbase/situation-model is a representation of the text and the new information it conveys. The text uses known words in novel combinations (sentences, paragraphs) and therefore the textbase/situation-model represents what is new and distinct about this text. In contrast, the LSA representation of text represents what is already known about this text - how the text looks in the already existing semantic space. Of course we could modify that space to include the new information by recomputing the singular value decomposition based on the old texts as well as the one just read. In practice, adding a single text to a huge base will have little effect, however. In real life, such additions occur continuously and with a significant cumulative effect.

The LSA representation of a text, therefore, tells us what the text is about in terms of what we already know. It is. therefore, the old information in a discourse, its theme or topic. The textbase/situation-model representation of the text provides the new information. It tells us what is distinctive about this text, how our old knowledge is recombined in novel ways to generate new meaning. Both kinds of representations exist at the local level (sentences, propositions) as well as the global level (macrostructures, whole sections of the text in LSA). Thus, at the global level we can say that a text is about a certain facet of knowledge (theme, topic), and we can specify this facet precisely by means of the LSA vector representing the text and its subsections. At the same time we can specify globally the novel, distinct information the text provides in terms of its macrostructure.

The distinction made here is commonly made in everyday language use. Texts have a topic or theme, but then they tell us something new about that .We say, for instance, that a given newspaper article is about "the forest fires in Florida," which specifies a particular knowledge domain, as an LSA analysis of the article would. We can also summarize the article as "the people who have been evacuated are returning to their homes as rains have doused the fires," expressing in words the macrostructure of the article and updating our information about that topic.

<u>Selecting a macrostructure: A sample analysis</u>.

In order to illustrate how LSA can be used to generate the macrostructure of a text, and how the notion of theme relates to that macrostructure, I present a brief computational example. The text to be analyzed is a chapter from a junior-high science textbook (reprinted in the Appendix) entitled "Wind Energy." It is 960 words long and divided by its author into six sections, each with its own subtitles. Thus, the author indicates the intended macrostructure and even provides appropriate macropropositions, in the form of six subtitles. Macrorules can be used to explain where these subtitles come from. Consider the following paragraph (the second subsection of the chapter):

(1) *The history of windmills.*

Since ancient times, people have harnessed the wind's energy. Over 5,000 years ago, the ancient Egyptians used the wind to sail ships on the Nile River. Later, people built windmills to grind wheat and other grains. The early windmills looked like paddle wheels. Centuries later, the people in Holland improved the windmill. They gave it propeller-type blades. Holland is still famous for its windmills. In this *country, the colonists used windmills to grind wheat and corn, to pump water, and to cut wood at sawmills. Today people still sometimes use windmills to grind grain and pump water, but they also use new wind machines to make electricity*

The macrorule of Construction can be used to compress sentences 2-4 into

People used windmills in Egypt.

Similarly, the other sections of the paragraph can be reduced to

People used windmills in Holland.

People used windmills in the colonies.

People use windmills today.

which can be transformed by the macrorule of Generalization into

(2) *People used windmills throughout history*

or

(3) *The history of windmills.*

Thus, macrorules allow us to postdict or explain what the author did. But the application of these rules depends on our intuitions about the text and our knowledge about it. By themselves, these rules cannot compute anything.

LSA provides a computational mechanism that can compute macrostructures of a kind. For instance, we can compute a vector in LSA space that is the centroid of all the words in paragraph (1). Such vector may seem to be totally useless - it is, after al, a list of 300 uninterpretable numbers! - but that is not so. It can be quite useful, for instance to decide how good a proposed subtitle is. The cosine between the paragraph vector and the proposed subtitle is a measure of how close the subtitle is to the paragraph as a whole. For instance, (2) and (3) have rather similar cosines with the centroid of the paragraph, .39 and .48, respectively, indicating that they are both acceptable summary statements. But suppose we had chosen an ill-considered subtitle for the paragraph, say “Holland is still famous,” or something totally inappropriate such as “Rain douses forest fires.” The cosine measure would have allowed us to reject these choices (the cosine is .26 in the first case and only .05 in the second, both much lower than the cosines for (2) and (3)). Of course, that is exactly what intuition tells us - but what I am interested here is finding a computational procedure that gives the same results as our intuitions.

There are other uses for vector representation of a macrostructure, too. For instance, we can compute how closely related the sections of a text are to each other. This kind of information can be of interest in various ways. If two sections of a text are very closely related, one might consider combining them. Or if two similar sections are separated by a dissimilar one in the text, one might consider reordering the sections of the text.

We can also obtain a measure of how important a section is to the overall text. One way to do this is to compute the cosine between the whole text and each section. This is not the best index, however, because long sections necessarily are advantaged by this measure. In the case of the "Wind Energy" text, which has six sections, one of the sections is far longer than the others (more than a third of the total text) and it would be selected as the most important one if we use this measure. That does not seem right intuitively, because it is a somewhat specialized section (entitled by the author "Wind power plants"). A better measure might be how typical a section is, that is, how strongly it is related to all the other sections of the text. Thus, we can compute the average cosine between each section and all the other sections in the text. By this measure, Section 3, entitled "Today's wind machines," is selected as the most typical one - more in line with our intuitions about the text.

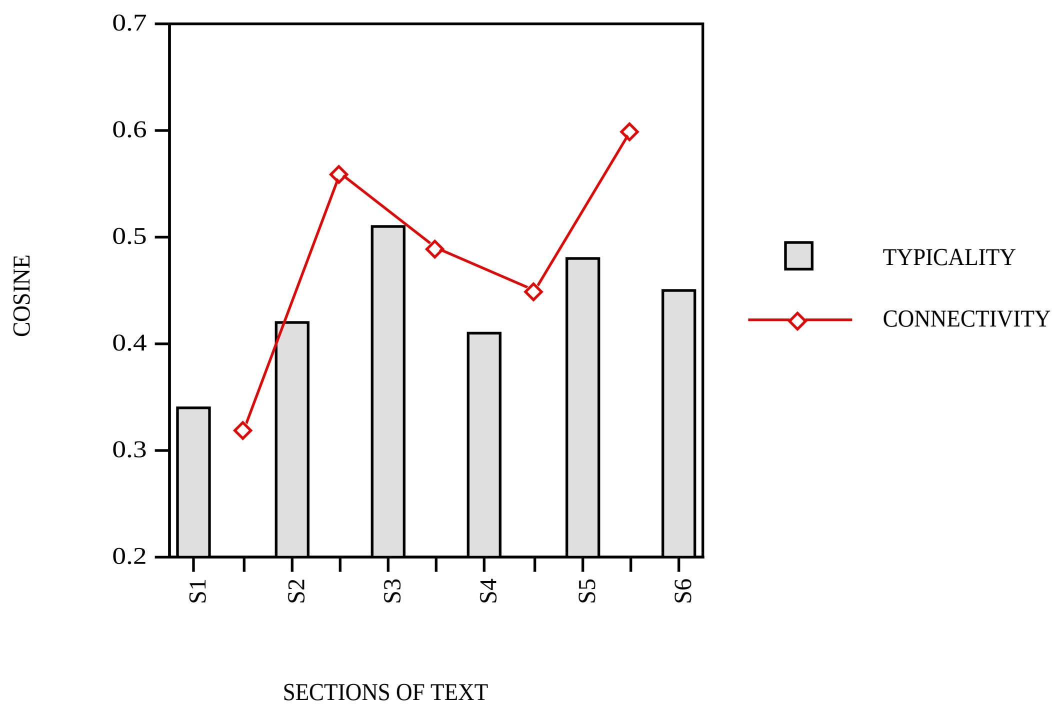

Figure 1: Typicality values of the six section of the "Wind Energy" text as well as measures of connectivity between sections.

Figure 1 shows the typicality values of the six section of the “Wind Energy” text. (the average of the cosines between a section and all other sections). Typicality correlates only moderately with importance (the cosine between a section and the whole text), r =. 56 in this case, because importance is highly correlated with the length of a section, r=. 87. Figure 1 also shows the connectivity of the text, how similar each section is to the next, as measured by the cosine between the two sections. Thus, Figure 1 provides a great deal of information about the macrostructure of this text: Section 1 is least typical of this section and least well connected (it is entitled “What is wind?” whereas all other sections deal with windmills in some way). Sections 2 and 3, and 5 and 6 are rather similar to each other in content. The text seems well-structured, although the structure is somewhat unusual: the first paragraph, which often states the main idea of a section, is here an aside, not closely related to the main body of the chapter.

While Figure 1 is informative about some aspects of the text’s macrostructure, it does not tell us in words what these sections are all about. That is, it does not provide the actual macropropositions. To generate the full range of macropropositions is beyond the scope of LSA; operations such as generalization and construction are not readily modeled within this framework. But we can generate a degenerate macrostructure using only the selection operation. For each section we can find the most typical sentence in the section (i.e. the one with the highest average cosine to all the other sentences in the section, as in Figure 1) and select that sentence as the macroproposition. This will not always yield the best result, because the ideal macroproposition may involve generalization or construction, but it will serve as a reasonable approximation. Table 1 shows the subtitles the author has selected for each section – presumably, the ideal macroproposition – as well as the sentence selected on the basis of typicality.

Section	Subtitle	Typical Sentence	Cosine(subtitle: sentence
S1	What Is Wind? (.33)	Wind is simply air in motion. (.48)	.55
S2	The History of Wind Machines	Today people still sometimes use windmills to grind grain and pump water, but they also use new wind machines	.55

	(.52)	to make electricity. (.73)	
S3	Today's Wind Machines (.63)	Like old-fashioned windmills, today's wind machines use blades to collect the wind's kinetic (motion) energy. (.82)	.62
S4	Types of wind machines (.77)	There are two types of wind machines commonly used today: horizontal-axis wind machines and vertical-axis wind machines. (.82)	.56
S5	Wind Power Plants (.59)	Wind power plants, or wind farms as they are sometimes called, are clusters of wind machines used to produce electricity. (.75)	.75
S6	How Much Energy Do We Get from the Wind? (.47)	One reason wind plants don't produce more electricity is that they only run about 25 percent of the time. (.71)	.26

Table 1. Subtitles for the six sections of the “Wind Energy” text and the typical sentence for each section selected by LSA. (The number in brackets behind each statement is the cosine between that statement and the corresponding sub-section of the text).

Table 1 lists the typical sentences selected in this manner for each section of the “Wind Energy” text. For sections 1,3,4 and 5, the subtitle is a subset of the words in the selected sentence. The title for section 2 could be obtained by generalizing the selected sentence. In all of these cases, the similarity between the author's subtitles and the selected statements is high, as shown in the last column of Table 1. For section 6, the subtitle and the selected statement are less similar. However, a generalization of the selected sentence, such as "Limitations of wind machines" would appear to be a viable alternative to the title the author has selected. The cosines between each subtitle and each statement and the corresponding section of the text are also shown in Table 1. As one would hope, they are all reasonably high, and, of course, they are higher for the longer statements than for the more concise subtitles.

Thus, intuitively, the sentences selected as typical appear to be appropriate. Furthermore, they correspond quite well with the summaries people have written for this text. Three expert summaries were obtained for this text. Each summary included a sentence or more for each section of the text. These expert summaries can be used as standards against which all sentences of the text can be evaluated: that sentence from the text should be selected that is most similar (by the cosine measure) to the experts' summary. Indeed, for four of the six sections that sentence is the shown in Table 1 as the most typical sentence. In general, sentences with higher typicality scores tended to be more similar to the expert summaries than sentences with lower scores. In addition to typicality (the average cosine between the sentence and all other sentences in that section), we also calculated for each sentence importance (cosine between the sentence and the whole section), and fit (cosine between the sentence and the expert section summary. Over the 62 sentence so the text, importance and typicality correlated r =. 56; typicality correlated significantly with fit r = .69, whereas importance correlated with fit r =. 57. Importance was highly correlated with word length, r = .87, typicality less so, r = .39.

As the example analyzed here shows, some progress can be made towards a computational model of macrostructure generation. LSA allows us to generate an abstract vector representation of the macrostructure of a text (at least in those cases where the subsection of the text are clearly indicated, as in the example above). Furthermore, procedures can be devised to select for each section of a text the most typical sentence. However, that doe not make a summary yet, and the operations for reducing the selected typical sentence to an essential phrase or fragment depend on more analytic procedures that go beyond LSA. For instance, if we had a parser that would segment the sentence into its syntactic constituents, we could then select the most typical phrase from the sentence in the same way that we have selected the most typical sentence from the section. Or alternatively, if we had an automatic way to generate propositional structures, we could use the construction-integration model to find the most highly activated propositional elements and base our summary on them. Indeed, with hand coding of sentences into phrases or propositions, both procedures yield promising results.

Conclusion.

Macrostructures are mental representations of text at a global level. They may simply mirror the structure of the text from which they were derived, or they may reflect, to varying degrees, the comprehender's own prior knowledge structure that has been imposed upon the text in the creation of a situation model.

Macrostructures as envisaged by van Dijk (1972) and discussed in van Dijk & Kintsch (1983) are hierarchies of propositions. Macropropositions put into words are summary statements, at different levels of generality. They subsume what the different sections of a text are about. They are derived from the text by the operations of selection, generalization, and construction, but propositional macrostructures cannot be computed automatically from a text. The macrorules merely help us to explain what can be done, but they are not algorithms or computational procedures that generate macropropositions from a text automatically.

A computationally more feasible but in other ways more limited alternative for the representation of macrostructures is provided by LSA. Instead of representing the meaning of a sentence by a proposition, its meaning can be represented as a vector in an existing high-dimensional semantic space. For some purposes, such a representation is all that is needed: e.g., one can compare new texts, such as summaries students write, with these macro-vectors, one can compute the importance or typicality of sentences from the text, and so on. For other purposes, verbal statements corresponding to macropropositions are needed. The best we can do with LSA, so far, however, is to select the most typical sentence for a section of a text. To transform that sentence into a true macro-statement involves operations such as generalization and construction and requires a detailed understanding of the syntactic and propositional structure of the sentence. LSA alone , in its current form, is not powerful enough to do so. Various alternatives to extend LSA, or to combine it with other psychological theories (specifically, the construction-integration theory of Kintsch,1998) will be explored in future work.

The suggestion made here is that LSA allows us to derive a precise mathematical representation of the topic or theme of a text. It is a non-verbal representation, but as was shown above, it nevertheless can be useful for certain purposes. The textbase and situation model, and at the global level the macrostructure of the text, represents both the given and new information. Learning, in the sense of knowledge acquisition, occurs when the LSA space is modified to incorporate the new information provided by the text.

Footnote

The preparation of this chapter was supported by a grant from the Army Research Institute. I thank Tom Landauer and Lise Menn for helpful discussions.

References

Ericsson, K. A., & Kintsch, W. (1995). Long-term working memory. Psychological Review, 102, 211-245.

Gernsbacher, M. A. (1990) Language comprehension as structure building. Hillsdale, NJ: Erlbaum.

Gernsbacher, M. A. (1994) Handbook of psycholinguistics. San Diego: Academic Press.

Havilland S. E. & Clark, H. H. (1974) What's new? Acquiring new information as a process in comprehension. Journal of Verbal Learning and Verbal Behavior, 13, 515-521.

Kintsch, W. (1988). The use of knowledge in discourse processing: A construction-integration model. Psychological Review, 95, 163-182.

Kintsch, W. (1998). Comprehension: A paradigm for cognition. New York: Cambridge University Press.

Kintsch, W., & van Dijk, T. A. (1978). Towards a model of text comprehension and production. Psychological Review, 85, 363-394.

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The Latent Semantic Analysis theory of acquisition, induction and representation of knowledge. Psychological Review, 104, 211-240.

Landauer, T. K., Foltz, P., & Laham, D. (1998). An introduction to Latent Semantic Analysis. Discourse Processes, 259-284.

van Dijk, T. A. (1980). Macrostructures. The Hague: Mouton.

van Dijk, T. A., & Kintsch, W. (1983). Strategies of discourse comprehension. New York: Academic Press.

Zwaan R. A. & Radvansky, G. A. (1998) Situation models in language comprehension and memory. Psychological Bulletin, 123, 162-185.

Appendix:

The “Wind” text.

What Is Wind?

Wind is simply air in motion. It is caused by the uneven heating of the earth by the sun. Since the earth's surface is made up of land, desert, water, and forest areas, the surface absorbs the sun's heat differently.

During the day, the air above the land heats up more quickly than air above water. The warm air over the land expands and rises, and the heavier, cooler air rushes in to take its place, creating winds. Likewise, the large atmospheric winds that circle the earth are created because the land near the earth's equator is heated more by the sun than land near the North and South Poles.

Today people can use wind energy to make electricity. Wind is called a renewable energy source because we will never run out of it.

The History of Wind Machines

Since ancient times, people have harnessed the wind's energy. Over 5,000 years ago, the ancient Egyptians used the wind to sail ships on the Nile River. Later, people built windmills to grind wheat and other grains. The early windmills looked like paddle wheels. Centuries later, the people in Holland improved the windmill. They gave it propeller-type blades. Holland is still famous for its windmills. In this country, the colonists used windmills to grind wheat and corn, to pump water, and to cut wood at sawmills. Today people still sometimes use windmills to grind grain and pump water, but they also use new wind machines to make electricity.

Today's Wind Machines

Like old-fashioned windmills, today's wind machines use blades to collect the wind's kinetic (motion) energy. Wind machines work because they slow down the speed of the wind. The wind flows over the blades causing lift, like the effect on airplane wings, causing them to turn. The blades are connected to a drive shaft that turns an electric generator to make electricity. The new wind machines are still wrestling with the problem of what to do when the wind isn't blowing. They usually have a battery to store the extra energy they collect when the wind is blowing hard.

Types of Wind Machines

There are two types of wind machines commonly used today: horizontal-axis wind machines and vertical-axis wind machines.

Horizontal-axis wind machines have blades that go crosswise and look like airplane propellers. A typical horizontal wind machine stands as tall as a 10-story building and has two or three blades that span 60 feet across. The largest wind machines in the world have blades longer than a football field! Wind machines stand tall and wide to capture more wind. Vertical-axis wind machines have blades that go from top to bottom and look like giant egg beaters. The typical vertical wind machine stands about 100 feet tall and is about 50 feet wide.

Wind Power Plants

Wind power plants, or wind farms as they are sometimes called, are clusters of wind machines used to produce electricity. A wind farm usually has hundreds of wind machines in all shapes and sizes. Unlike coal or nuclear electric power plants, most wind plants are not owned by public utility companies. Instead they are owned and operated by business people who sell the electricity produced on the wind farm to electric utility companies.

Operating a wind power plant is not as simple as plunking down machines on a grassy field. Wind plant owners must carefully plan where to locate their machines. One important thing to consider is how fast and how much the wind blows. Scientists use an instrument called an anemometer to measure how fast the wind is blowing. An anemometer looks like a modern-style weather vane. It has three spokes with cups that spin on a revolving wheel when the wind blows. It is hooked up to a meter that tells the wind speed. (By the way, a weather vane tells you the direction of the wind, not the speed.)

As a rule, wind speed increases with height and over open areas with no wind breaks. Good sites for wind plants are the tops of smooth, rounded hills, open plains or shorelines, and mountain gaps where the wind is funneled. The three biggest wind plants in California are located at mountain gaps.

Wind speed varies throughout the country. It also varies from season to season. In Tehachapi, California, the wind blows more during the summer than in the winter. This is because of the extreme heating of the Mojave desert during the summer months. The hot desert air rises, and the cooler, denser air from the Pacific Ocean rushes through the Tehachapi mountain pass to take its place. In Montana, on the other hand, the wind blows more in the winter. By happy coincidence, these seasonal variations perfectly match the electricity demands of the regions. In California, people use more electricity during the summer cooling months. In Montana, people use more electricity during the winter heating months.

How Much Energy Do We Get from the Wind?

Every year wind energy produces enough electricity to serve 300,000 households, as many as in a city the size of San Francisco or Washington, D.C. Still this is only a tiny amount of the electricity this country uses. One reason wind plants don't produce more electricity is that they only run about 25 percent of the time. (That means they would run 25 hours in a 100 hour period—or six hours a day.) Wind machines only run when the wind is blowing around 14 mph or more. Coal plants, on the other hand, run 75 percent of the time because they I can run day or night, during any season of the year.

Wind machines do have some advantages over coal-fired power plants though. First, wind machines are clean. They do not cause air or water pollution because no fuel is burned to generate electricity. Second, we may run out of coal some day, but we won't run out of wind.