

3

Writing about What We Know: Generating Ideas in Writing

David Galbraith

An old cliché about writing has it that writers should write about what they know. The intention being to exhort writers, young writers in particular, to write about things in terms of their own experience rather than, for example, to make up stories about spaceships and pop stars. There is some truth in this: writers do write more fluently and coherently about familiar topics than they do about unfamiliar topics. However, as we shall see, simply writing about what one knows has also been identified as one of the main weaknesses of novice writers.

There are two senses in which simply writing about what you know is problematic. The first is that writing does not just involve the unmediated transcription of knowledge. It also involves actively transforming knowledge in order to satisfy broader communicative goals. The second sense in which ‘writing what you know’ is a problematic piece of advice in that it implies that the writer already knows what they will be going to write about. However, as E.M. Forster famously asked, ‘How can I know what I think until I see what I say’. Exhortations to write about what you know imply that the writer should stick to thoughts that they

already have rather than that they should explore the possibilities of thoughts that they do not as yet fully understand.

These two senses suggest that over and above writing about what one knows, one needs also to (i) go beyond this in order to create a coherent text that achieves the writer’s goals with respect to the reader, and (ii) that ‘what we know’ is something that may only emerge in the course of writing rather than being something pre-existing about which we can write. In this chapter, I will outline the classic cognitive models of how the writer’s knowledge is realized as text and then describe how they account for the transformation and/or constitution of knowledge during writing. I will then consider an alternative account of these processes based on a different conception of how knowledge is mentally represented.

KNOWLEDGE-TELLING

Hayes and Flower’s original model of the writing process (Hayes and Flower, 1980; see Alamargot and Fayol, this volume, for an

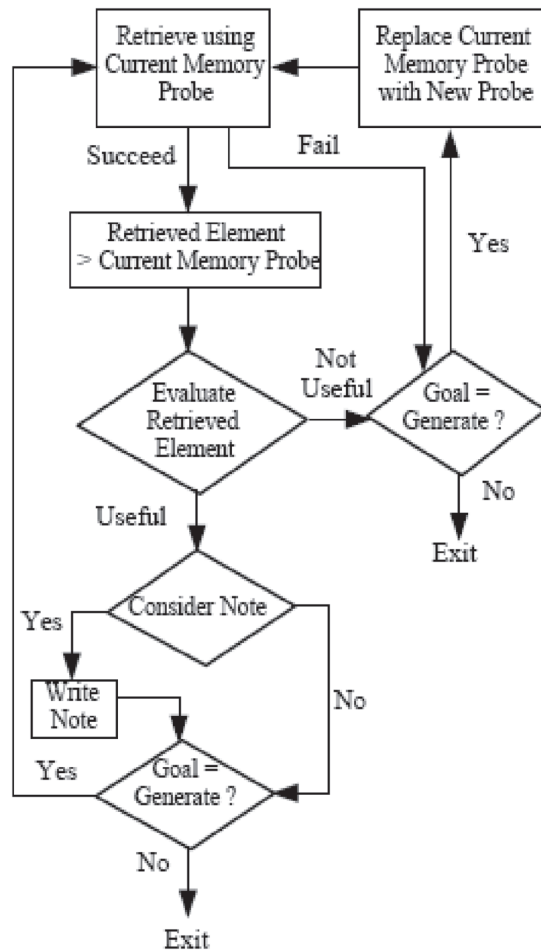
overview) was based on their analysis of protocols collected from writers thinking aloud as they wrote. Their aim was to identify the main components of the process and create a model of how these were combined in the course of writing. This model was deliberately constructed as a high-level model corresponding to writer's introspections about how writing was carried out. Fine-grained detail about how different components of the process were implemented by the cognitive system was left to be filled out by a combination of current research into basic features of the cognitive system and future research about how these were instantiated during writing. Their analysis of the protocols suggested that three main kinds of process were involved: *planning*, which involved the formulation of goals, and the generation and organization of ideas in order to satisfy those goals; *translation*, which, as the name implies, involved the translation of ideas into language; and *reviewing*, which involved reading over the text that the writer had already written, and editing it so that it satisfied their goals better. These processes used resources from the task environment (including the writing assignment and the text produced so far) and from the writer's long-term memory and were coordinated by a 'monitor' or central executive, which was responsible for deciding which processes should be carried out when.

For present purposes, the key feature of this model is the strong distinction between planning, where the creation of content takes place, and translation, where content is formulated as language. This means that what the writer says in their text is exclusively determined by the planning component of the writing process: translation is concerned with how best to express what the writer thinks, and reviewing is concerned with reading and editing the way this thought is expressed. To understand their claims about how writers generate the content of their text, therefore, we need to look in detail at their model of the planning process and how it interacts with the task environment and the

writer's long-term memory. Figure 3.1 shows their model of the generating subcomponent of the process.

According to this model, in order to work out what to say about the topic, the writer starts by using the specifications in the writing assignment to construct a set of cues with which to probe long-term memory. If content is successfully retrieved and then positively evaluated, it is either noted down on paper or stored mentally for later translation into text. This content then acts as a new probe for memory so that each retrieval episode consists of associative chains of content being retrieved from memory. If they cannot retrieve appropriate content, then they (or more strictly speaking the monitor) have to decide what to do next. This may be to pursue a different goal – for example, to go and read a book about the topic, or to read the assignment more closely – or it may be to carry on generating content by probing memory again with a different set of cues.

Note, also, that the monitor plays a role even if the writer successfully generates content in that the writer has to decide what to do next. This decision process is under the control of the monitor and depends in part on the writer's overall writing strategy. The writer might, for example, decide to translate the content into full text or to consider how to organize the idea in relation to other already generated content. Alternatively, they might decide to carry on generating further content in note-form. Hayes and Flower characterized differences in the overall strategy used by writers as corresponding to different configurations of the monitor. Some writers, for example, prefer to generate content freely before they turn to organizing: they would be characterized in terms of the model as prioritizing the goal of generating in the early stages of the writing process. Other writers may prefer to attempt a perfect first draft and engage in full cycles of generating, organizing, translating, and reviewing before moving on to the next episode of generating.



Query: Please check quality not good.

Figure 3.1 The generating component of the writing process as conceived in Hayes and Flower's model (1980:13).

This model of idea generation as a process of retrieving content from long term memory is a general assumption in current cognitive models of writing. It is virtually identical to the knowledge-telling model described by Bereiter and Scardamalia (1987) (see Figure 3.2). Although they use this model to characterize the writing processes of younger and less expert writers, and contrast it with the knowledge-transforming model employed by older and more expert writers, it is still embedded as the idea generation component of the more advanced model.

COGNITIVE OVERLOAD

Given this model of how ideas are generated during writing, the main finding of research on the factors affecting retrieval is that idea generation can be reduced when it has to be combined with other components of writing. Subsequent research has focused on strategies for reducing the effects of this conflict between different components of writing.

Caccamise (1987), for example, asked college students to carry out what was essentially a knowledge telling task. They were

Query: Please check quality not good.

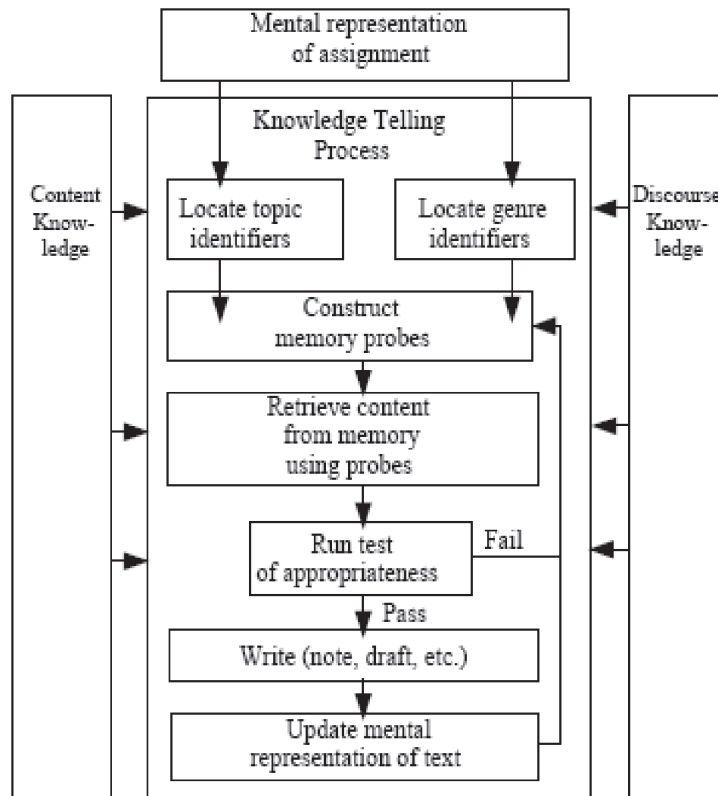


Figure 3.2 Bereiter and Scardamalia's (1987:8) model of the knowledge telling process.

asked to say everything that came to mind about a topic without worrying about repetitions or how well organized or expressed what they said was. She gave them four different topics varying in how familiar they were and asked them either to assume an audience of fifth grade school children or an audience of fellow college students. Her first, straightforward, finding was that the students produced many more ideas about familiar topics than unfamiliar ones, and that ideas produced about the familiar topic were much more organized than the ideas produced about unfamiliar ones. More interestingly, she also found that when children were given as the audience, although the students produced just as many ideas, these were much less organized and were produced much less fluently. She concluded that the audience

was not typically included as part of the memory probes used to search memory, but that the audience constraint was imposed after retrieval as part of an editing process. She also suggested that, because of the limited capacity of short-term memory this editing process could not be carried out very effectively, and recommended that under such circumstances writers should be encouraged to generate their ideas first without regard to the audience and only concern themselves with editing once ideas had been retrieved.

Caccamise's results suggest that idea generation can be impaired by conflicts with relatively high level processes like adapting to audience constraints. A series of studies by Bourdin and Fayol (1994) comparing written and spoken recall with varying age groups

suggest that low level processes involved in spelling and handwriting can also impair retrieval. In simple word-recall tasks, Bourdin and Fayol found that both second and fourth grade children recalled substantially fewer items when their responses were written (i.e., using relatively less practised handwriting and spelling skills) than when their responses were spoken (i.e., they could rely on more automatic speech production skills). There was no corresponding difference for adults. In fact, recall was slightly better with written responses for adults. Similar results were found for a more complex sentence production task (Bourdin and Fayol, 1996). However, when the composition task was substantially more complex (Bourdin and Fayol, 2002), they found that even adults perform worse in writing compared to speaking. This suggests that even when spelling and handwriting are very well practised, they can still have a residual effect on memory retrieval if working memory resources are overloaded by other resource demanding processes.

A study by Glynn et al. (1982) suggests that idea generation is reduced whenever it has to be combined with other components of the writing process. Following Hayes and Flower (1980), they identified four distinct operations in writing: (a) generating ideas, (b) sequencing or organizing ideas, (c) expressing ideas in sentences, and (d) complying with spelling and grammatical conventions. The extent to which these operations had to be carried out at the same time was manipulated by instructing writers to divide the writing of a brief letter into two separate ten-minute sessions and varying the number of operations, which had to be carried out during the preliminary draft. The resulting four preliminary draft formats successively removed one of the operations. In the polished sentences condition, writers had to write a complete, polished version of the letter on the first draft. In the mechanics-free sentences condition, they had to write the complete text, but without worrying about mechanics (spelling and punctuation). In the ordered notes condition, they were instructed to write their ideas down

in brief three or four word notes, and to ensure these were organized into a logical order. Finally, in the unordered notes condition, they were instructed to jot their ideas down in note-form as before, but not to worry about the order in which they were expressed. Glynn et al. found that the number of ideas generated in the preliminary drafts was progressively lower as the number of constraints present increased, with the fewest ideas being generated in the polished sentences conditions and the most ideas being generated in the unordered notes condition. They concluded that generating ideas was more productive when it was carried out in note-form prior to the production of text than when it was carried out at the same time as producing the text.

One implication of this general line of research is that it is important for other components of the writing process to be carried out as automatically as possible. Being able to write or type fluently and having well-developed language skills should reduce cognitive overload and facilitate more fluent retrieval of content from long-term memory. In addition, strategies for managing the writing process, which help reduce cognitive load should also enable more fluent idea generation. The most thorough investigation of the effectiveness of different drafting strategies was carried out by Kellogg in a series of experiments (Kellogg, 1988, 1990; see Kellogg, 1994, for a review). Kellogg (1988) distinguished two different ways in which writers might reduce cognitive load during writing. An outline strategy, in which writers generate and organize their ideas prior to writing, before focusing their attention on translation and revision processes, should minimize the attention required for translating ideas during planning. A rough-drafting strategy, which involves translating text without worrying about how well expressed it is, leaving monitoring of expression to revision of the draft after writing, should further reduce the burden on formulating content during text production. In combination, then, outlining followed by rough drafting should

in theory provide for the most efficient distribution of resources during writing since it will separate both planning from translation and translation from revision. Kellogg (1988) tested this by manipulating two variables. Writers were instructed either to make a hierarchical outline before writing, or to start writing immediately. Then, when they produced the text itself, they were instructed either to write the text freely, without worrying about how well it was expressed, returning later to revise, or to attempt to produce a polished text on the first draft. The effect on the distribution of processes during writing (as indicated through directed retrospection) and the quality of the final text were measured. The results were very clear. First, the manipulations did indeed lead to a redistribution of processing during writing. In the outline conditions, writers showed much less evidence of planning during text production, presumably because this had largely been completed prior to writing. In the rough draft conditions, revision was reduced during the initial draft and postponed until later. Second, outlining was associated with higher quality final drafts but rough drafting showed no effect, despite the fact that revision had been postponed and should, therefore, have been able to draw on more attentional resources.

A later study by Kellogg (1990) suggested that the construction of a hierarchically organized outline prior to writing is associated with a higher quality final product than is the construction of an ordered list of ideas, and that this in turn is associated with higher quality final text than a simple clustering strategy. Furthermore, this is true despite the fact that more ideas tend to be generated using a clustering strategy than when an outline is constructed. Kellogg's (1994) general conclusion is that the effectiveness of the outlining strategy is a consequence of the fact that it enables writers to organize their ideas better prior to writing, as well as that it then enables them to devote more resources to formulating these ideas effectively in text. (But see Galbraith and Torrance, 2004, for

evidence that these studies may underestimate the effectiveness of certain forms of rough-drafting strategy).

KNOWLEDGE-TRANSFORMING

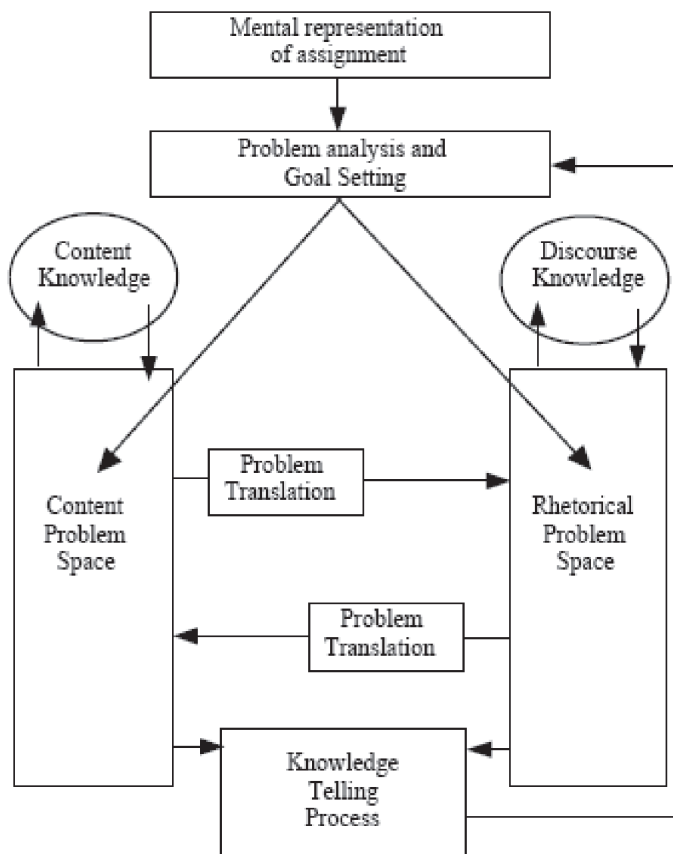
The research we have considered so far has concentrated on how retrieval of ideas can be improved by automating other components of the writing process – particularly the low-level skills involved in transcription and the language skills involved in formulating ideas in text – and by strategies for managing the writing process itself. Studies comparing the writing processes of experts and novices (see Hayes and Flower, 1986, for a review of research on adults, and Bereiter and Scardamalia, 1987 for a review focusing on the developmental literature) have, however, suggested that these are not the only ingredients involved. Experts not only just manage the cognitive conflict involved more effectively but also appear to direct their writing towards different goals. While novice writers appear to define writing as primarily a matter of expressing what they know about a topic, expert writers define it as a matter of achieving communicative goals. On the other hand, in Flower's (1979) words, whereas novices produce 'writer-based' prose; experts produce 'reader-based' prose.

This difference in the goals towards which writing is directed represents a fundamental shift in focus and has wide-ranging consequences for the way that writing is carried out. Experts typically develop much more elaborate and interconnected sets of goals for their writing, building these networks gradually and modifying them in the course of writing. As a result, they spend longer planning, during, as well as before, writing: Bereiter and Scardamalia (1987) report that adult protocols typically contain ratios of thought to text of around 4:1, whereas children of 10 or so years old show ratios of thought to text more like 1:1. Flower and Hayes (1980) report that, whereas novices

generate the majority (70%) of their ideas in response to the topic alone, experts generate the majority of their ideas (60%) in response to their rhetorical goals. Moreover, revision is no longer a matter of assessing how well the writer's ideas are expressed in language but becomes a matter of assessing how well the text satisfies the writer's goals. As a reflection of this, Hayes's (1996) revised model of writing no longer treats the revision process as a simple matter of reading and editing text or as an independent component of the writing process. Instead, it is treated as a combination of the more basic processes of text interpretation, reflection, and text production, and involves generating and organizing ideas, both during and after writing, in order to satisfy rhetorical goals.

AQ: "Hayes's 1996" is not listed in Reference list.

Bereiter and Scardamalia (1987) formalized these differences in their knowledge transforming model of writing (see Figure 3.3), stressing that this should not be seen simply as an evolution of the knowledge telling model but that it involved a radical change in the way that the writing task is defined by the writer and in the way that it is carried out. Thus, although it retains the knowledge telling model (and hence Hayes and Flower's model of idea generation) as a characterization of the process whereby content is retrieved from memory, this is embedded within a dialectic between content and rhetorical problem spaces. This is intended to capture two features of the writing process. First, it reflects the fact that ideas are represented, not just as a reflection of the writer's



Query: Please check quality not good.

Figure 3.3 Bereiter and Scardamalia's (1987) knowledge-transforming model of writ-

knowledge (content space), but also in terms of their rhetorical function within the text (rhetorical space). Second, writing is not simply a matter of adapting content to the rhetorical context, but is an emergent process in which content is formulated as the text develops. Thus, not only is content retrieved in response to a more elaborated representation of the assignment as a rhetorical problem, but it is also formulated in the context of, and as a contribution to, the series of rhetorical acts gradually emerging in the text. The contrast between the knowledge telling model and the knowledge transforming model is a contrast between writers who ask themselves questions like, ‘What do I know about this?’, ‘Does this sentence correspond to the idea I want to express?’, ‘What else do I know about this?’, and writers who, having thought about the goals they want to achieve, say things like, ‘If I want to achieve this, then the first step I need to take is ...’, ‘I can do this by saying ...’ ‘Having said that, what do I need to do next ...’. Content is retrieved, as and when it is needed, in order to develop the text further towards the writer’s goals. In consequence, writing well involves not just satisfying rhetorical goals, but in the course of doing so, developing one’s understanding of the topic.

A training study conducted by Scardamalia et al. (1984) with a group of sixth grade students provides a good illustration of how this translates into practice. Training took place over 19 weeks and included procedural facilitation, modelling of writing by both teachers and pupils, and explicit instruction in various problem-solving strategies. The key theoretical ingredient was procedural facilitation. This involved providing students with prompt cards to be used during planning. These prompt cards were grouped according to their rhetorical function (e.g., elaborate, improve, goals, and organization) and consisted of cues like ‘A better argument would be’, ‘An example of this would be’, ‘My main point is’, ‘But many readers won’t agree that’, and so on. These prompts were used extensively throughout the course of

instruction – in demonstrations by the teacher, in modelling by students for one another, and during practice writing exercises – and were designed to help the students set rhetorical goals and use these to guide the generation of content during writing. In post-tests, Bereiter et al. found that, in contrast to a control group, the training group increased the number of reflective thoughts included in think aloud protocols collected during planning, and produced what were judged to be more reflective essays. They noted, however, in a more detailed analysis, that students did less well on ratings concerned with ‘developing a coherent and well-thought-out position’ on a topic, and suggested that this could reflect a gap between attempt and execution.

One feature that distinguishes Bereiter and Scardamalia’s approach is that they place less emphasis on strategies like outlining. This is partly because they see the knowledge transforming model as applying across the writing process as a whole, and partly because they want to emphasize that, what matters about planning is not so much *when* it takes place as *how* it is carried out. In a study investigating how writers of different ages (students from Grades four, six, and eight, as well as a group of adults) set about constructing the main point of their texts, Bereiter, Burtis and Scardamalia (1988) found that, over and above the amount of time spent planning, the main difference between older and younger writers was in how the writers went about constructing their main points. As part of the study, the participants were asked to think aloud while planning and the resulting protocols were scored for evidence of six different kinds of constructive moves, which Bereiter et al. assumed to be characteristic of a knowledge transforming approach to writing. Although, they did find a strong relationship between the length of time spent planning and the quality of the main points identified in the plans, they also found that the better main points were associated with a greater number of constructive moves. Notably, this was not just a consequence of the difference in the ages of the groups: the

Query: Year needs to be inserted here?

Query: Year needs to be inserted here?

same relationships were apparent within the different age groups. A subsequent path analysis showed that grade level influenced the number of constructive moves the writer made, which then influenced the level of the main point in the plan, which in turn influenced the level of the main point in the text. The crucial feature being that grade level did not have a direct influence on level of main point (as would be expected if the results were just a consequence of differences between the age groups in existing knowledge) but only influenced it via the number of constructive moves made during planning. (See also, a more recent study by Galbraith et al. (2005) showing that when the amount of time spent preplanning is held constant the quality of the finished text correlates strongly with the extent to which extra content is added to satisfy rhetorical goals during preplanning).

In summary, the knowledge transforming model represents an important extension of Hayes and Flower's model of idea generation. Although it shares the same basic mechanism of generating ideas – the retrieval of content from long-term memory – it emphasizes the rhetorical nature of the goals towards which writing is directed, and claims that this involves a redefinition of writing, rather than simply an evolution of a knowledge telling approach. This means that, although Bereiter and Scardamalia might accept that using outlining to reduce cognitive load could facilitate planning, they would argue that it is not sufficient. Only once writing is redirected towards rhetorical goals does it become a means of developing the writer's understanding, rather than simply a means of making content generation more fluent.

Although this model of writing captures important features of the writing process, I want to argue in the remainder of this chapter that it is only a partial account and that it has two particularly problematic features. First, although one of its attractive features is its claim to account for the common experience of writing as a source of discovery, this is only an implication of the model, which was not

directly tested in the development of the model. Second, the knowledge-telling model, which is embedded within the knowledge-transforming model as its account of how content is generated to satisfy goals, does not on the face of it explain how novel content is formulated during writing.

DISCOVERY THROUGH WRITING

The research we have discussed so far has focused on the processes involved in writing rather than on the effects of these processes on the writer's thought. The claim that writing develops understanding is an implication of the model, rather than something that has been empirically tested. Subsequent research has investigated this more directly, and examined the conditions under which writers discover new ideas through writing (Galbraith, 1992, 1999; Galbraith et al., 2006).

These experiments used the same general procedure. Writers were asked to rate how much they felt they knew about a topic and then to list their ideas about the topic, both before and after writing. They were then asked to rate the similarity of the ideas contained in the two lists. In some experiments, they were also asked to rate the similarity of the ideas within lists, again, both before and after writing. This enabled Galbraith and his colleagues to identify whether the processes carried out during writing had led to a change in what the writer thought about the topic or in how organized their thoughts were, and whether any such changes were associated with subjective changes in knowledge.

This research has focussed on the conditions under which two different types of writer, selected using Snyder's self-monitoring scale (Snyder, 1986), generate new ideas during writing. High self-monitors are, in Snyder's words, 'particularly sensitive to the expression and self-presentation of relevant others in social situations and use these cues as guidelines for monitoring (that is regulating and controlling) their own verbal and

non-verbal self-presentation'. By contrast, low self-monitors' 'self-presentation and expressive behaviour ... seems, in a functional sense, to be controlled from within by their affective states (they express it as they feel it) rather than moulded and tailored to fit the situation'. Galbraith selected these two types of writers, because they seemed to embody the contrast between knowledge-telling and knowledge-transforming approaches to writing: low self-monitors would be expected to prioritize the direct expression of their beliefs about the topic, whereas high self-monitors would be expected to generate content to satisfy their communicative goals. (A recent study by Klein et al., 2004, showing that high self-monitors vary the thoughts they list before a discussion with different audiences, whereas low self-monitors do not, supports this assumption).

Galbraith (1992) asked low and high self-monitors either to write notes in preparation for an essay (planning) or to write an essay without preplanning (text production), and measured the extent which writers developed new ideas in these different conditions using the methods described in the preceding section. If, as Bereiter and Scardamalia suggest, discovery depends on the extent to which writers generate content in response to rhetorical goals, one would expect the high self-monitors to produce more new ideas after writing than the low self-monitors. If, furthermore, the process involves deliberate problem solving and this is impaired when the capacity of working memory is overloaded, then one would expect a greater number of new ideas to be produced after planning in note-form than when writers had to produce full text at the same time as planning. This was, partly, what Galbraith found. The high self-monitors discovered more new ideas after writing notes than the low self-monitors did, and this was reduced when the high self-monitors had to write full text. This suggests that discovery depends on the extent to which content generation is directed towards the satisfaction of rhetorical goals and that it is reduced when the writer has to

deal with the extra cognitive load of producing well-formed text. However, if this were the case, one would also expect the low self-monitors to generate even less new ideas when writing full text. In fact, Galbraith found that the low self-monitors generated a high number of new ideas after writing full text, just as many in fact as the high self-monitors did when they made notes. In other words, discovery also appeared to occur when writing was assumed not to be directed towards rhetorical goals, and when cognitive load should be at its highest. Furthermore, in direct contrast to what the knowledge transforming model would predict, the new ideas produced by the high self-monitors after writing notes were not associated with subjective increases in writers' knowledge. By contrast, there was a clear positive correlation between the number of new ideas produced by the low self-monitors after writing full text and increased knowledge of the topic.

Overall, Galbraith (1992) concluded that, although there was evidence that adapting thought to rhetorical goals does affect the generation of content, this was not associated with the development of the writer's understanding. In addition, there was also evidence that dispositionally guided text production, far from being a matter of retrieving existing ideas from memory, involved actively creating novel content, and that this led to the development of the writer's understanding.

In a later study, using a similar method but focusing on different forms of text production, Galbraith (1999) found that the number of new ideas produced by low self-monitors was at its greatest when writing was not planned beforehand, and that outlining before writing reduced the difference between the low and high self-monitors, with low self-monitors experiencing a reduction in new ideas and high self-monitors an increase in new ideas compared to nonplanned writing.

More recently, Galbraith et al. (2006) again replicated the finding that low-self-monitors discover more new ideas after

nonplanned text production than high self-monitors, but also found evidence of a difference in the effect of outline-planned writing on low and high-self-monitors' thought. Thus, although, as in previous experiments, the high self-monitors produced more new ideas when they were allowed to make an outline before writing, Galbraith et al. found that the high self-monitors' ideas were less coherently organized after writing. In other words, although high self-monitors did appear to produce more new ideas when they were asked to make an outline before writing than when they just wrote an unplanned draft, as one might expect if outlining helps them to develop a fuller representation of the rhetorical problem, this appeared to lead to less coherently organized thoughts about the topic. By contrast, low self-monitors produced more new ideas after writing a spontaneous draft, and in both planned and unplanned writing maintained the coherence of their thoughts before and after writing.

For present purposes, the key point about these results is that they appear to contradict important features of Bereiter and Scardamalia's knowledge transforming model. First, they consistently suggest that, although writers who are sensitive to rhetorical content, and readily adapt their thought to the communicative context, do generate new ideas during writing, these new ideas are not associated with the development of understanding, and in fact appear to reduce the coherence of the writer's thought. Second, they consistently suggest that writers whom one would expect to translate their existing ideas into text, and who typically do not adapt their thought to the demands of the communicative context (low self-monitors), also generate new ideas during writing, and that these new ideas are coherently organized and associated with the development of their understanding. In other words, precisely the wrong kinds of people, according to the knowledge-transforming model, appear to develop their understanding through writing.

KNOWLEDGE CONSTITUTING

By themselves, these empirical findings are suggestive, but not conclusive, evidence against the knowledge transforming model. Further research is needed directly examining the processes involved when low and high self-monitors develop their ideas before we can conclusively rule out a knowledge transforming account of the findings. The problem is that these empirical findings tie in with a more conceptual problem in the model. This stems from the assumption that generating ideas is a matter of retrieving pre-existing content from memory. While this can account for the fact that different content is retrieved when the rhetorical context drives memory search than when content is retrieved associatively, it does not explain how this can be new content that develops the writer's understanding. Bereiter and Scardamalia (1987: 349–351) acknowledge this problem when they discuss memory search procedures, pointing out that 'One of the most formidable challenges to theories of language use is to explain how it is that skillful speakers and writers are able so quickly to think of material fitting multiple constraints'. They consider the possibility that this could be a consequence of spreading activation within a network of fixed semantic units but conclude that it would be impossible for such a network to store explicitly all the different possibilities that might be required by a different rhetorical contexts. Instead, they suggest a process of heuristic search, in which the rhetorical problem is progressively redefined until it 'provides cues that activate appropriate nodes in memory'. However, it is not clear how, given their assumption that memory nodes consist of fixed units, this resolves the issue. Surely, all that heuristic search – or more refined rhetorical planning – achieves in this context is retrieval of more specific content. In other words, although they recognize the problem, it is not clear that heuristic search provides a solution to it.

Both Hayes and Flower and Bereiter and Scardamalia appeal to information processing

Query: Year
needs to be
inserted here?


models current at the time they were writing. These have come to be characterized as symbolic models of representation and typically assumed that content was explicitly represented in long-term memory. At just about this time, alternative, connectionist theories of representation began to appear (McClelland, Rumelhart and the PDP group, 1986). In such models, content is not stored as fixed, explicit representations in a separate long-term memory, waiting to be retrieved for manipulation in short-term memory. Instead, it is *synthesized* as and when it is needed in a contextually appropriate form.

To give you an idea of the principles of the processing involved, consider the very simple, 'toy' network shown in Figure 3.4.

This network consists of simple units, very loosely analogous to neurons in the brain, with three basic properties: (i) each unit sums up the activation it receives via connections (shown as arrows in the diagram) from other units in the network; (ii) if the total activation it receives is higher than its threshold for activation, then (iii) the unit will fire and pass

activation on to all the other units it is connected to in the network. In themselves, the units or nodes in such networks do not represent anything at all, they simply sum up activation from incoming connections and transmit it through outgoing connections to units further up the network. In the network shown in Figure 3.5, these units are organized in three layers. For the sake of this example, you can think of the input layer as a perceptual layer, receiving input from the external environment (each unit might be a photoreceptor in the retina for example), the hidden layer as the network's conceptual representation of this input, and the output layer as the network's response to the perceptual input (for example, the name of an object perceived in the environment).

The key feature of this network is that its knowledge is not represented by the hidden units within the network. These same set of units has different patterns of activation for different inputs and only develop these patterns of activation in the presence of a particular input. Rather it is the connections



Query: Please check quality not good.

Figure 3.4 A simple feed-forward network.

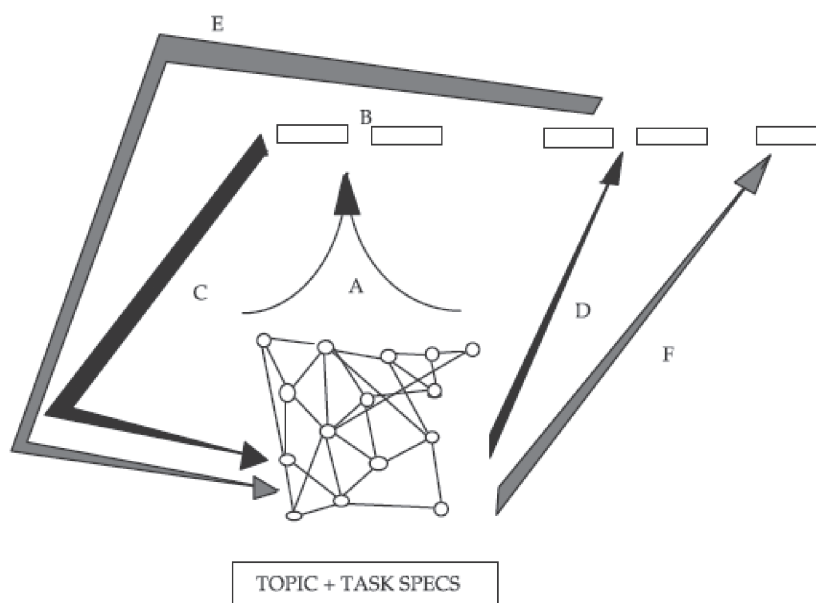


Figure 3.5 Writing as a knowledge-constituting process (Galbraith, 1999).

between the units, which are fixed and which guide processing, and which constitute the network's knowledge of the world. This knowledge is not directly accessible, but reveals itself, implicitly, in the response it causes the network to synthesize in a particular context. Thus, when the visual image of a cat for example appears on the input layer, the connections linking the input layer will pass activation on to the hidden layer, and cause a particular pattern of activation to appear across these hidden units, which will in turn pass activation on through the connections linking the hidden layer to the output layer, causing a particular pattern of activation at the output layer corresponding to the word 'cat'. Equally, if the input pattern were to be the visual image of a dog, for example, the *same* set of connections would result in a different pattern across the hidden layer, and the production of the response 'dog' at the output layer. In other words, the network's knowledge is represented by the strength of the connections between units, and these have their effects, implicitly, by guiding activation round the network, enabling it to

synthesize different responses to different inputs. In such networks then knowledge is synthesized anew in a contextually specific form each time it is used rather than being stored explicitly as a fixed piece of information in long-term memory.

Galbraith (1999) invoked these principles to provide an account of writing as a knowledge constituting process. The first important feature to note about this model is that it does not deny that content is represented sometimes as individual events which can be accessed via a process of retrieval. Writers do sometimes remember (or try to recall) what it is they want to say. Rather, it claims that such knowledge is stored in a separate – *episodic* – memory system, operating (perhaps) according to the principles assumed by Hayes and Flower and Bereiter and Scardamalia, whereas language production draws on a different – *semantic* – memory system, and operates according to connectionist principles. (See McClelland et al. (1995) for evidence of two such complementary memory systems, and Rogers and McClelland, 2004, for a connectionist model of semantic memory.) According to Galbraith,

the knowledge constituting process is prompted specifically by the requirement to formulate thought in explicit propositions.

In the knowledge-constituting model, shown in Figure 3.5, the writer's semantic memory is represented by the network of units at the centre of the diagram, with the writer's knowledge, or disposition towards the topic as Galbraith (1999) calls it, represented by the fixed weights connecting the units. Given an input (TOPIC + TASK SPECS), the units within the network are activated according to the strength of the connections between the input pattern and the units in the network. The units then pass activation between themselves until they settle into stable state. This represents the 'best fit' the network can find between the various pieces of information in the network, and corresponds to the message the writer wishes to convey (labelled A in the diagram). In terms of the 'toy' example I have just described, the TOPIC + TASK SPECS are analogous to the input layer, the writer's semantic memory to the hidden layer and A to the output layer. This message is then formulated as an utterance (labelled B in the diagram). (Note that the transformation of A (the message) into B (the utterance) could be considered to be a further network, with an input layer at A passing activation through a hidden layer (not shown) to an output layer at B).

By themselves, steps A and B could be seen as a replacement for the idea generation component of Bereiter and Scardamalia's knowledge-transforming model. The only difference being that, rather than simply involving the retrieval of existing content, this would involve synthesizing content, as if for the first time, in response to a contextually specific set of goals. A particular virtue being that this would provide a mechanism whereby novel content could be created in response to altered rhetorical goals. Furthermore, since knowledge is represented implicitly in such networks, the writer would, quite literally, be 'finding out what they think by seeing what they say'.

However, the second part of the model (as represented by the steps labelled C, D, E, and F in the diagram) makes further claims which mean that it cannot simply be added to the knowledge transforming model as a more up to date account of how content is synthesized in text. The first claim is that the initial proposition formulated by the network (at B in the diagram) is only a partial, 'best fitting', representation of the input. When the input initially occurs, it activates a wide range of units within semantic memory according to the strength of the links between the input and the nodes within the network. However, once these units have been activated, constraint satisfaction between the units will mean that some ultimately 'lose out' and hence are not incorporated in the message to be formulated in language. Furthermore, limitations in the expressive capabilities of the writer's linguistic resources mean that not all of the content will be formulated in language.

In order, then, to fully capture the writer's implicit knowledge, further cycles of text production are required. Galbraith (1999) suggests that this involves negative connections (labelled C in the diagram) from the output at B to the units in semantic memory. (Mannes and Kintsch, 1991, include a similar assumption in their model of planning by constraint satisfaction). This has the effect of reducing the activation of units corresponding to the initial proposition, and enables previously suppressed units to influence the next output synthesized by the network. This following synthesis of content is labelled D in the diagram. The network continues with succeeding cycles of negative feedback (labelled E) and utterance synthesis (labelled F) until all the content activated by the input has been formulated.

There are alternative ways in which feedback from the initial proposition could take place. For example, feedback may involve adding the previous output utterance to the input to the network, rather than a direct influence on the network itself. For present purposes, however, the key point is that the writer's implicit disposition will only be fully

realized in the text to the extent that it is allowed to continue to guide the production of successive propositions. This only occurs if the sequence of propositions is allowed to unfold without interruption by external, rhetorical constraints. The writer's response to the initial prompt is not represented by a single proposition but is instead spread, discursively, across the set of propositions. Thus, in this model, progressive refinement of thought is achieved in the text itself, and involves successive dispositional responses to emerging propositions rather than a progressive redefinition of rhetorical constraints.

This emphasis on the dispositionally-driven nature of text production brings the hypothesis into direct conflict with the knowledge transforming model's claim that discovery depends on the adaptation of thought to external rhetorical goals. Taken together, the different aspects of the model explain why writers in Galbraith's experiments develop their understanding when they write full text (express thought in explicit propositions) rather than when they plan in note-form and why it is the low self-monitors (dispositionally guided writers) rather than the high self-monitors who do so. It also suggests why, in Scardamalia, Bereiter and Steinbach's (1984) training study, in which students were asked to consult cue cards after each succeeding sentence, the children did not produce particularly coherent text. The knowledge constituting model assumes that this was because external rhetorical goals interrupted the text production process before the writer's implicit disposition towards the topic had been fully constituted in the text.

CONCLUSION

The knowledge constituting model suggests that, during text production, idea generation involves the synthesis rather than the retrieval of content. Nevertheless, it does not deny that retrieval of existing content or explicit

planning to satisfy rhetorical goals are an important aspect of writing. Reflectively surveying, memory plays a valuable role in identifying relevant content, and explicitly formulating rhetorical goals ensures that it is appropriate to the rhetorical context. However, in order for this potential content to be realized in the text, and in order for it to capture the writer's implicit understanding, it has to be dispositionally synthesized in the course of text production. This means that a dual process model is required to account for writing as a whole.

According to the dual process model, then, both rhetorical planning and dispositional text production are required for effective writing. Rhetorical planning is assumed to operate on an episodic memory of previously entertained propositions (ideas that the writer has read, heard, or formulated themselves in the past). Ideas are retrieved from this memory in the way described in Hayes and Flower's model and writers vary in the extent to which memory search is guided by their rhetorical goals. Because it involves explicit consideration of content and possible ways of organizing it, it is subject to working memory constraints, and operates best when thought is represented economically in note form. Dispositional text production is assumed to operate through parallel constraint satisfaction within the writer's semantic memory, and is responsible for constituting content suggested by planning in a series of explicit propositions in the text. In general, the two processes interact with each other. Planning delivers potential content for realization in the text, and unpredicted formulations in the text lead to revision of the writer global plan. Writers vary in the extent to which they prioritize the two processes. Low self-monitors appear to prioritize dispositional goals, with the result that they are more likely to constitute their thought in text, but perhaps at the expense of producing less clearly structured text. High self-monitors appear to prioritize rhetorical goals, with the result that they are more likely to consider a wider range of rhetorically appropriate

content, but at the expense of constituting their own understanding in the text.

The dual process model casts many of the phenomena I have reviewed here in a different light. First, it suggests that, over and above the cognitive overload engendered by having too many things to think of at once, another fundamental conflict in writing is between writers' dispositional and rhetorical goals. The writer's disposition, when conceived of as the fixed weights in a constraint satisfaction network, consists of the writer's distinctive point of view about the world. It is who they are. This means that writing arouses profoundly conflicting emotions: elation when one finds oneself discovering a new insight; fear when one feels oneself entering unexplored territory with unpredictable consequences; alienation when external constraints or our own preconceptions prevent us from constituting our thought; loneliness when what we say is misunderstood by others.

Second, the model implies a different view of the role of language in writing. The main focus in cognitive models of writing to date has been on the thinking behind the text and on the need to reduce the load on working memory to enable the writer to do it more effectively. Hence, the emphasis has been either on the need to plan separately from producing text, or on the need to automate language processes so that they do not consume the working memory resources required for thinking. If the knowledge constituting process is as is suggested here, then language skills have a much more essential role. They are the vehicle through which thought is constituted externally, and hence enable the writer to gain access to their implicit disposition about a topic. Text production in general, rather than being something that gets in the way of thinking, is in fact where thinking takes place.

Let me conclude with a final comment on 'writing about you know'. I suggested initially that this was bad advice, partly because it could encourage a knowledge-telling model of writing, and partly because it assumes that

the writer already knows what they want to say. The knowledge constituting model suggests that, insofar as they are taken to mean that one should let one's writing be dispositionally guided, there is also a sense in which this, and other romantic exhortations like 'be true to yourself', are good advice. According to the dual process model, however, the more analytic process of adapting thought to rhetorical goals is equally important, and the conflict between the two processes is a necessary one. The reason why there is a perennial conflict between romantic approaches to writing, which focus on personal expression, and more classical approaches which focus on rhetorical skill, is that they both capture an essential component of the writing process.

REFERENCES

- Bereiter, C. and Scardamalia, M. (1987) *The Psychology of Written Composition*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Bereiter, C., Burtis, P.J., and Scardamalia, M. (1988) 'Cognitive Operations in Constructing Main Points in Written Composition', *Journal of Memory and Language*, 27: 261–278.
- Bourdin, B. and Fayol, M. (1994) 'Is Written Language Production more Difficult than Oral Language Production – a Working-memory Approach', *International Journal of Psychology*, 29(5): 591–620.
- Bourdin, B. and Fayol, M. (1996) 'Mode Effects in a Sentence Production Span Task', *Cahiers De Psychologie Cognitive-Current Psychology of Cognition*, 15(3): 245–264.
- Bourdin, B. and Fayol, M. (2002) 'Even in Adults, Written Production is still more Costly than Oral Production', *International Journal of Psychology*, 37(4): 219–227.
- Caccamise, D.J. (1987) 'Idea Generation in Writing', in A. Matsuhashi (ed.), *Writing in Real Time: Modelling Production Processes*, New York, NY: Longman Inc. pp.224–253.
- Flower, L. (1979) 'Writer-based Prose: A Cognitive Basis for Problems in Writing', *College English*, 41(1): 19–37.
- Flower, L.S. and Hayes, J.R. (1980) 'The Cognition of Discovery: Defining a Rhetorical Problem', *College Composition and Communication*, 31: 21–32.

- Galbraith, D. (1992) 'Conditions for Discovery through Writing', *Instructional Science*, 21: 45–72.
- Galbraith, D. (1999) 'Writing as a Knowledge-constituting Process', in M. Torrance and D. Galbraith (eds), *Knowing What to Write*, Amsterdam, NL: Amsterdam University Press. pp.139–160.
- Galbraith, D. and Torrance, M. (2004) 'Revision in the Context of Different Drafting Strategies', in L. Chanquoy, L. Allal, and P. Largy (eds), *Revision in Writing*, Dordrecht, NL: Kluwer Academic Publishers. pp.63–86.
- Galbraith, D., Ford, S., Walker, G., and Ford, J. (2005) 'The Contribution of Different Components of Working Memory to Planning in Writing', *L1 – Educational Studies in Language and Literature*, 15: 113–145.
- Galbraith, D., Torrance, M., and Hallam, J. (2006) 'Effects of Writing on Conceptual Coherence', Proceedings of the 28th Annual Conference of the Cognitive Science Society. pp. 1340–1345.
- Glynn, S. M., Britton, B., Muth, D., and Dogan, N. (1982) 'Writing and Revising Persuasive Documents: Cognitive Demands', *Journal of Educational Psychology*, 74: 557–567.
- Hayes, J.R. and Flower, L.S. (1980) 'Identifying the Organization of Writing Processes', in L.W. Gregg and E.R. Steinberg (eds.), *Cognitive Processes in Writing*, Hillsdale, NJ: Lawrence Erlbaum Associates. pp.3–30.
- Hayes, J.R. and Flower, L.S. (1986) 'Writing Research and the Writer', *American Psychologist*, 41(10): 1106–1113.
- Hayes, J.R. and Nash, J.G. (1996) 'On the Nature of Planning in Writing', in C.M. Levy and S. Ransdell (eds), *The Science of Writing: Theories, Methods, Individual Differences and Applications*, Mahwah, NJ: Lawrence Erlbaum Associates. pp.29–56.
- Kellogg, R.T. (1988) 'Attentional Overload and Writing Performance: Effects of Rough Draft and Outline Strategies', *Journal of Experimental Psychology: Learning, Memory and Cognition*, 14: 355–365.
- Kellogg, R.T. (1990) 'Effectiveness of Prewriting Strategies as a Function of Task Demands', *American Journal of Psychology*, 103: 327–342.
- Kellogg, R.T. (1994) *The Psychology of Writing*, New York: Oxford University Press.
- Klein, O., Snyder, M., and Livingston, R.W. (2004) 'Prejudice on the Stage: Self-monitoring and the Expression of Group Attitudes', *British Journal of Social Psychology*, 43: 299–314.
- Mannes, S.M. and Kintsch, W. (1991) 'Routine Computing Tasks: Planning as Understanding', *Cognitive Science*, 15(3): 305–342.
- McClelland, J.L., McNaughton, B.L., and O'Reilly, R.C. (1995) 'Why there are Complementary Learning Systems in the Hippocampus and Neocortex: Insights from the Successes and Failures of Connectionist Models of Learning and Memory', *Psychological Review*, 102: 419–457.
- McClelland, J.L., Rumelhart, D.E., and the PDP research group (1986) *Parallel Distributed Processing: Volumes 1 and 2*, Cambridge, MA: MIT Press.
- Rogers, T. and McClelland, J. (2004) *Semantic Cognition*, Cambridge, MA: MIT Press.
- Scardamalia, M., Bereiter, C., and Steinbach, R. (1984) 'Teachability of Reflective Processes in Written Composition', *Cognitive Science*, 8: 173–190.
- Snyder, M. (1986) *Public Appearances, Private Realities: The Psychology of Self-Monitoring*, New York: W. H. Freeman and Company.