The Inner Voice in Writing

N. ANN CHENOWETH
The University of Texas–Pan American
JOHN R. HAYES
Carnegie Mellon University

This study explores the connection between writing and working memory, specifically the role of the subvocal articulatory rehearsal process (or inner voice). The authors asked the 18 participants to type sentences describing 24 multipanel cartoons. In some conditions, the participants were required to repeat a syllable continuously while writing. This activity, called articulatory suppression, interferes with the articulatory rehearsal process. Results indicated that interfering with the articulatory rehearsal process (or inner voice) interferes with writing by slowing the rate of writing, increasing mechanical errors, changing the temporal microstructure of text production, and increasing the perceived difficulty of the writing task. The authors applied their model of written text production to provide a theoretical account for these results.

Keywords: *writing models; working memory; text production; articulatory suppression; revision*

We have all had the experience when writing of thinking of something to write and then hurrying to write it down before it was forgotten. If we did not get to our computer in time or find our pen soon enough, the words may well have been lost. Our fleeting hold on the words we have intended to write reminds us that the path from intention to text involves temporary memory resources. Psychologists have called these resources working memory.

Another phenomenon that many writers experience is a perceived inner voice that seems to speak the words that they intend to write. Peter Elbow (1994) has described a similar phenomenon in which many writers hear the text that they are reading. The presence of this inner voice suggests another connection between writing and working memory. To explore this connection further, we need to review some aspects of working memory.

WRITTEN COMMUNICATION, Vol. 20 No. 1, January 2003 99-118
DOI: 10.1177/0741088303253572

Baddeley and his colleagues (e.g., Baddeley & Hitch, 1974; Gathercole & Baddeley, 1993) have proposed a model of working memory that consists of the following three parts: a phonological loop for storing verbal information, a visuospatial sketchpad for storing visual information, and a central executive that, among other functions, manages the other two parts.

The phonological loop consists of two parts. The first part, called the phonological short-term store, represents verbal material in a phonological code that decays rapidly with time. Attending to the phonological code would be experienced as hearing the verbal material. The second part is a subvocal, articulatory rehearsal process that refreshes the material in the short-term store. Engaging in articulatory rehearsal is experienced as speaking to oneself.

Kellogg (1996) has predicted that the phonological loop in Baddeley's model is involved in the process of formulating text. Levy and Marek (1999) have provided some empirical support for this prediction by showing decrements in a sentence-writing task when participants were required to listen to irrelevant speech while performing the task. If we were able to show that articulatory suppression interfered with writing, we would have stronger evidence for this prediction. If the inner voice is instrumental in the act of writing, then articulatory suppression should interfere with writing.

According to the Gathercole and Baddeley (1993) model, when a person hears a telephone number, that number is represented in the short-term store. If the number cannot be dialed immediately, it may be forgotten unless the person rehearses the number subvocally. Baddeley and Lewis (1981) has referred to the phonological short-term store as "the inner ear" and the articulatory rehearsal process as "the inner voice." It is as if the inner voice speaks to the inner ear to remind it of the material to be remembered.

It is possible to interfere with a person's articulatory rehearsal processes by having that person repeat a syllable over and over, for example, *the*, *the*, *the* or *tap*, *tap*, *tap*. This process, which psychologists have called articulatory suppression, impairs memory for verbal material by preventing rehearsal. For example, if a person is trying to remember a telephone number, repeating a syllable such as "the" or "tap" will prevent rehearsal of the number. Articulatory suppression may reduce memory for words by as much as 50% (Longoni, Richardson, & Aiello, 1993).

Chenoweth and Hayes (2001) examined the relation of linguistic experience in a second language (defined by the number of semesters

spent studying the language) to written fluency in that language. They found that increased linguistic experience was related to greater fluency, reflected in a higher rate of text production and a lower frequency of revisions. In addition, they found that writers with greater linguistic experience produced texts in longer bursts. That is, they wrote more words between breaks in text production. Chenoweth and Hayes (2001) suggested that burst length "is fundamentally related to fluency" in that it reflects "the capacity of the translator to handle complex language structures" (p. 94).

In this study, we explored the impact of articulatory suppression on writing both when the texts that the participants were writing were visible to them and when they were not. We included an invisible text condition because we believed that participants might use the visible text as a memory aid while composing, especially when articulatory suppression made articulatory rehearsal difficult. For this reason, we expected to find a larger effect of articulatory suppression on composing when the text was invisible than when it was visible.

METHOD

Participants

The 18 participants in this study were each paid $10 for their time. They are all adult native speakers of English.

Procedure

Each participant was asked to view a series of 24 cartoons and to write a one-sentence summary of each cartoon on a word processor. Participants were shown two examples of cartoons together with "good" and "less good" sentences describing the cartoons. The purpose of the examples was to emphasize the importance of completeness in describing the cartoons.

During all trials, the cartoon remained in view on the computer screen while the participants typed their sentences, and there was a metronome clicking 120 times a minute. In the voice-tap condition, participants said "tap" in time to the metronome; in the foot-tap condition, they tapped their foot in time to the metronome; and in the no-tap condition, they were not required to respond to the sound of the

metronome. The foot-tap and no-tap conditions, which do not involve articulatory suppression, may be thought of as control conditions for the voice-tap condition, which does involve articulatory suppression. For half of the trials, the sentences that the participants wrote were invisible to them. Thus, there were the following six experimental conditions: voice-tap + visible text; foot-tap + visible text; no-tap + visible text; voice-tap + invisible text; foot-tap + invisible text; and no-tap + invisible text. The 24 trials consisted of four blocks of six trials. All six conditions were presented in each block of six trials. The first block (trials 1 to 6) was a practice block, designed to provide the participants with experience in writing under each of the experimental conditions. The remaining three blocks (trials 7 to 24) were the experimental blocks. The order of conditions was randomized in each experimental block separately for each participant with the constraint that on each trial, three participants engaged in each of the six experimental conditions. Only data from the experimental blocks were used in the following analysis. The 24 cartoons were presented to the participants in a single fixed order, and the conditions for the practice sessions were all presented in the same order. The presentation of the cartoons and the collection of data was managed by a program written in the cT language[1] and run on a Macintosh G4.

At the end of the experiment, the researcher asked participants, "Did you always have a complete sentence in mind when you started to type?" "Which of the conditions was most difficult?" and "Which was the next most difficult?"

Design

The experimental design was a within-participants design with the following three repeated measures: tap condition (voice-tap, foot-tap, or no-tap), text visibility condition (visible or invisible), and block.

Analysis

The software started timing when the cartoon first appeared, recorded the identity of the key struck and the time of each keystroke (including deletes), and stored the final version of the sentence that the participant wrote. These data allowed us to compute the following measures of fluency[2]:

- Start time (the time between the appearance of the picture and the participant's first key stroke).

- Number of words (the number of words in the final version of the sentence).
- Writing time (the time between the participant's first and last keystroke).
- Rate (number of words divided by writing time).
- Revision time (the time between the beginning of a series of deletes and the resumption of typing new text).
- Number of revisions (the number of revisions in each sentence).
- Burst length (the number of newly proposed words in a burst).
- Number of bursts (the number of bursts in each sentence).

A burst is a group of words produced by a writer for inclusion in a text that is bounded by breaks in the production process. A break is defined either as a pause of at least 2 seconds or a grammatical discontinuity indicating that the prior language was being revised.

Bursts that were terminated by a pause we call P-bursts. Bursts that were terminated by a revision we call R-bursts. Data for P-bursts and for R-bursts were analyzed separately.

We also evaluated each sentence for the following measures of quality:

- Complexity (based on the number of clauses in the final sentence).
- Mechanical errors (based on the number of grammatical, typing, or spelling errors in the final sentence).
- Holistic quality (based on the quality ranking of the sentence compared to all of the sentences in the same text visibility condition describing the same cartoon).[3]

Prior to carrying out ANOVAs, we applied Hartley's test (Rosenthal & Rosnow, 1991, pp. 338-340) for homogeneity of variance. For four of our measures—revision time, number of revisions, P-burst length, and mechanical errors—it was necessary to apply a logarithmic transformation to achieve acceptable homogeneity.

RESULTS

Fluency Measures

Start time. Start time was not significantly influenced by any of the independent variables. The average start time over all conditions was 10.03 seconds.

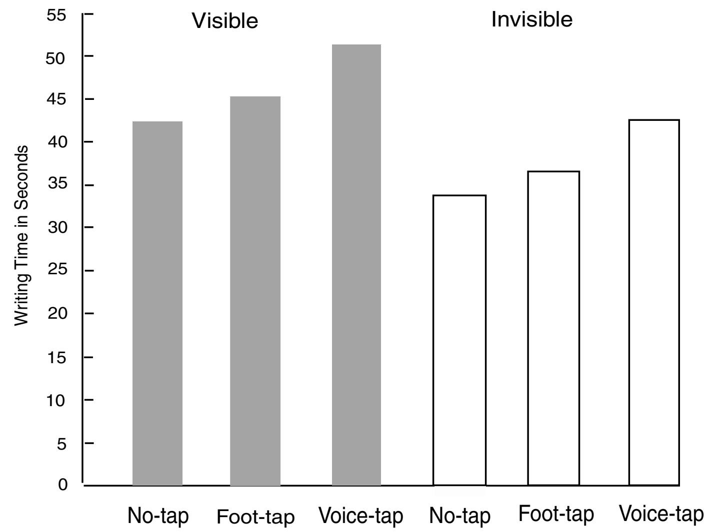

Figure 1. Average writing time for the six experimental conditions

Number of words. Number of words was not significantly influenced by any of the independent variables. The average number of words per sentence over all conditions was 21.65.

Writing time. Figure 1 shows the average writing time for the six experimental conditions. The difference between visible and invisible text was significant ($F = 8.608$, $df = 1, 15$, $p = .0103$) as were the differences among tap conditions ($F = 6.445$, $df = 2, 30$, $p = .0047$). The differences between voice-tap and no-tap and between voice-tap and foot-tap were significant, the respective p values being .0015 and .0192, but the difference between no-tap and foot-tap was not significant.[4] There were no significant differences due to block.[5]

Rate. Figure 2 shows the average rate in words written per second for each of the six experimental conditions. The difference between visible and invisible text was significant ($F = 14.772$, $df = 1, 14$, $p = .0018$) as were the differences among tap conditions ($F = 11.079$, $df = 2, 28$, $p = .0003$). The differences between voice-tap and no-tap and between voice-tap and foot-tap were significant, the respective p values being .0001 and .0050, but the difference between no-tap and foot-tap was not significant. There was a significant interaction between the text visibility and the tap conditions ($p = .0113$). Articulatory suppres-

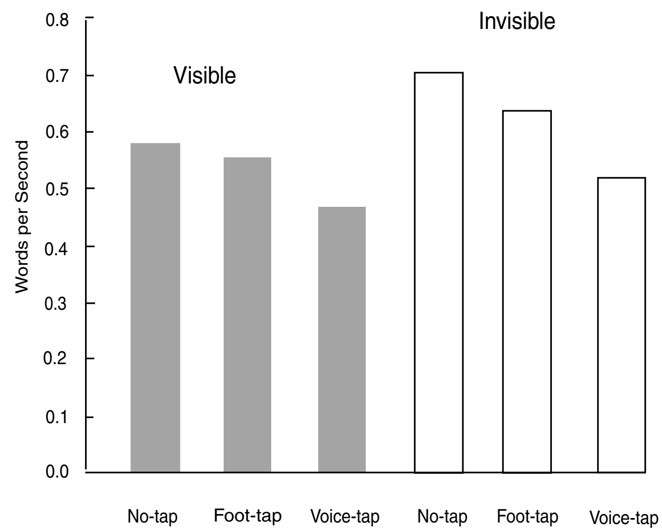

Figure 2. Writing rate for the six experimental conditions

sion had a stronger effect in the invisible condition than in the visible condition. There were no significant differences due to block.

Revision time. Text visibility had a significant main effect on log revision time ($F = 29.901$, $df = 1, 16$, $p < .0001$). There were no other main effects or interactions. The average percentage of writing time spent in revision was 17.3% in the visible condition and 8.2% in the invisible condition.

Number of revisions. Text visibility also had a significant main effect on the log of the number of revisions ($F = 36.035$, $df = 1, 15$, $p = .0001$). The average number of revisions in the visible condition was 3.01 per sentence and in the invisible condition, 1.34 per sentence. Thus, making the text invisible reduced the number of revisions by about 56%. There were no other main effects or interactions.

P-burst length. Figure 3 shows the average P-burst length for the six experimental conditions. The differences in log P-burst length among tap conditions were significant ($F = 18.73$, $df = 2, 32$, $p < .0001$). The differences between the voice-tap condition and the two other conditions were significant and large. The *p* values were both less than .0001. However, the difference between the no-tap condition and the foot-tap condition was not significant. Figure 4 shows the distribution

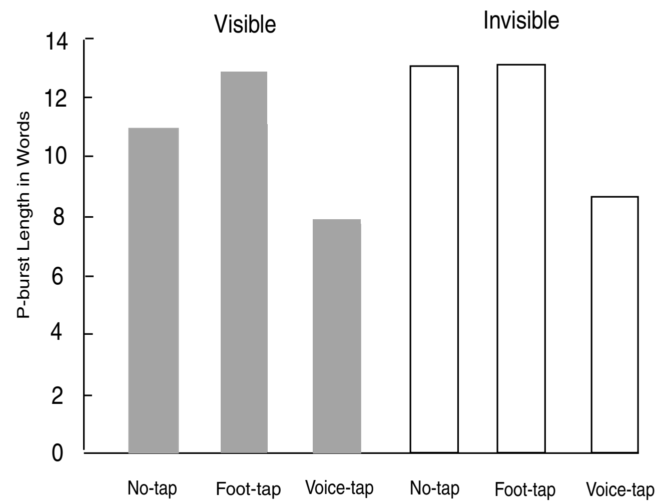

Figure 3. P-burst length for the six experimental conditions

of P-burst lengths for the voice-tap condition (that is, the condition with articulatory suppression) and the two conditions with no articulatory suppression. In the voice-tap condition, the whole distribution of P-burst lengths is shifted to the left, indicating shorter P-burst lengths when the articulatory rehearsal process is interfered with. Mean P-burst lengths for the no-tap, foot-tap, and voice-tap conditions were 12.07, 13.00, and 8.22 words, respectively.

There was a marginally significant main effect of text visibility (F = 3.453, df = 1, 16, p = .0816). P-bursts were 10.53 words in length in the visible condition and 11.66 words in the invisible condition. There were no other significant main effects or interactions.

Number of P-bursts. As would be expected, analysis of the number of bursts per sentence also showed a significant main effect for tap conditions (F = 16.092, df = 2, 32, p < .0001). The number of P-bursts in the voice-tap condition was significantly larger (p > .0001) than in either of the other response conditions. The foot-tap and no-tap conditions were not significantly different. The only other significant effect was an interaction between the tap conditions and block (F = 2.891, df = 4, 64, p = .029).

R-burst length. R-bursts were less frequent than P-bursts. Of the 917 bursts measured in this study, 87.1% were P-bursts and 12.9% were R-

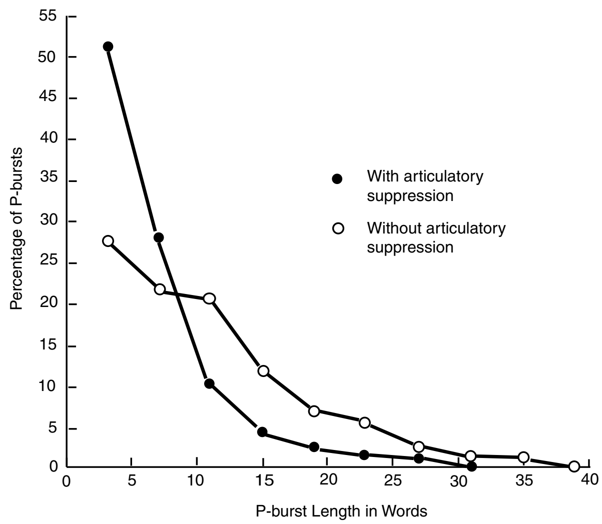

Figure 4. The distribution of P-burst lengths with and without articulatory suppression

bursts. Chenoweth and Hayes (2001, p. 91) found that for writers writing in their first language, 13% of bursts were R-bursts. Because the number of R-bursts is small, there were too many empty cells to allow for an ANOVA of the effects of the experimental variables on R-burst length. However, averaging the nonempty cell provided a crude estimate of R-burst length of 6.39 words. This suggests that in this study[6] R-bursts are substantially shorter than P-bursts, which average 11.10 words in length. This is the result that we would expect if R-bursts were P-bursts interrupted by revisions.

Number of R-bursts. An analysis of the numbers of R-bursts in each sentence revealed a significant main effect for text visibility (F = 10.139, df = 2, 32, p = .0058), a marginally significant main effect for block (F = 2.545, df = 2, 32, p = .0942) and no other main effects or interactions.

Quality Measures

Number of clauses. Block was the only independent variable that had a (marginally) significant main effect on the number of clauses (F = 3.333, df = 2, 28, p = .0503). Sentences in block 2 (trials 13 through

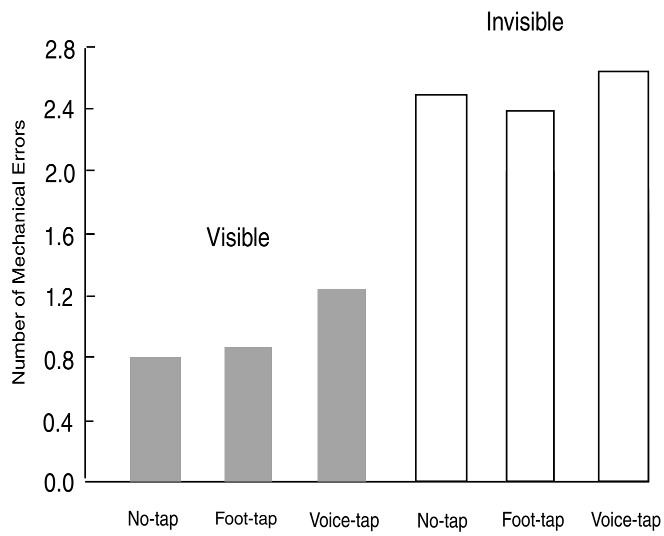

Figure 5. Mechanical errors for the six experimental conditions

18) had significantly more clauses than sentences in block 3 (trials 19 through 24). There was a significant interaction between block and tap condition ($F = 3.399$, $df = 4, 56$, $p = .0147$). On average, participants' sentences included 3.44 clauses.

Mechanical errors. Figure 5 shows the number of mechanical errors in each of the six experimental conditions. Text visibility had a significant main effect on the log of the number of mechanical errors ($F = 48.234$, $df = 1, 14$, $p < .0001$). The mean number of errors per sentence was 0.97 in the visible condition and 2.50 in the invisible condition. The tap conditions also had a significant main effect ($F = 4.182$, $df = 2, 28$, $p = .0258$). The mean numbers of errors in the voice-tap, foot-tap, and no-tap conditions were 2.0, 1.6, and 1.6, respectively. There were no other main effects or interactions.

Holistic quality. Results for the two judges who evaluated holistic quality were not consistent. For one judge, there was a significant main effect of tap conditions ($F = 3.398$, $df = 2, 28$, $p = .0477$). For this judge, quality in the voice-tap condition was significantly lower than in the two control conditions. However, for the second judge, there were no significant effects of the tap conditions. Furthermore, the Spearman correlation between the judges' rankings was only .50.

Given these inconsistent results, it is unclear whether articulatory suppression influenced text quality or not.

Answers to Postsession Questions

"Did you always have a complete sentence in mind when you started to type?" Four of the 18 participants answered "yes" or "I think so." The remaining 14 answered "no."

"Which of the conditions was most difficult?" "Which was next most difficult?" Of the 18 participants, 3 judged that the two conditions with voice tapping (one with visible and one with invisible text) were tied as most difficult; and 1 indicated that voice tapping with visible text was the most difficult. For these 4 participants, requiring voice tapping made the task more difficult but making the text invisible did not. The remaining 14 judged that the voice-tap + invisible text condition was the most difficult. For these participants, we can estimate the relative contributions to task difficulty of requiring voice tapping versus making the text invisible by noting which condition was judged the second most difficult. If voice tapping contributed the most to difficulty, then the second most difficult condition should be voice tapping with visible text. In contrast, if making the text invisible contributed most to difficulty, then one of the other invisible conditions should be judged most difficult. For 6 of these 14 participants, invisibility made the greatest contribution, for 7, voice tapping made the greatest contribution, and for 1, the contributions were about equal. Thus, voice tapping contributed most to problem difficulty for 11 of the participants, text invisibility for 6 of the participants, and the two variables contributed equally for the remaining participant. We can conclude, then, that for most participants, writing with articulatory suppression was as difficult or more difficult than writing without the visual feedback provided in the visible text condition.

DISCUSSION

Our participants were given the task of writing one-sentence descriptions of cartoons and, in all of the experimental conditions, they were able to do so. In all of the conditions, participants wrote sentences that were about equally long (21 to 22 words) and equally complex (three to four clauses). However, the writers' performance was

influenced by the experimental conditions in a number of interesting ways. We summarize these below.

First, articulatory suppression slowed the rate of text production by about 21%.

In contrast, making the text invisible had the unexpected consequence of increasing the rate of text production by about 16%. This increase can be attributed in large part to a substantial reduction in the amount of time participants spent revising during the invisible text condition.

Table 1 shows that much of the effect of text visibility on writing time and rate can be accounted for by the effect that this variable has on revision. If we subtract the amount of revision time from total writing time, the differences between the text visibility conditions both in writing time and rate are reduced by 55% to 60%. But what about the remaining 40% to 45%? We speculate that in the course of writing, participants may pause to consider revisions that they decide not to carry out. Such pauses would not be scored as revision time in this study because they do not accompany overt revisions. It is possible, then, that the remaining difference between writing times and rates in the visible and invisible conditions may result from revisions that were considered but not carried out.

We speculated that articulatory suppression would have a larger effect in the invisible condition where external memory resources were not available than in the visible condition where they were available. This speculation was borne out in our observations on the rate of text production. The significant interaction between the tap and text visibility conditions reflects a greater reduction in rate due to articulatory suppression in the invisible condition (23%) than in the visible condition (18%).

Second, neither the total revision time nor the number of revisions per sentence was affected by articulatory suppression.

Third, P-burst length was markedly decreased (34%) by articulatory suppression. There are at least two possible explanations for this decrease.

Explanation 1: Articulatory suppression might affect the output of the translation process by limiting the length of translated passages that can be stored in the phonological loop. If this were the case, we would expect that long bursts would be less frequent under articulatory suppression but that short bursts would be relatively unaffected.

Table 1
Writing Time and Rate Adjusted for Revision Time

Variable		Visible	Invisible	Difference
Writing time	Not adjusted for revision time	46.44	37.41	9.03
Writing time	Adjusted for revision time	38.36	34.34	4.02
Rate	Not adjusted for revision time	0.481	0.571	–0.090
Rate	Adjusted for revision time	0.587	0.624	–0.037

Explanation 2: Articulatory suppression might affect the translator directly by reducing its capacity to translate passages of any length. If this were the case, we would expect that P-bursts of all lengths would be shortened. Examination of Figure 4 suggests that the second explanation is more consistent with our results than the first.

We also found that P-burst length was slightly increased (11%) in the invisible text conditions. This effect may be tied to the reduction in the number of revisions that occurs in the invisible text condition. Chenoweth and Hayes (2001) suggested that revisions are more likely to occur in long bursts than in short ones. If so, frequent revisions would have the effect of reducing mean P-burst length. Thus, reducing the number of revisions should increase mean P-burst length.

Fourth, the frequency of R-bursts was strongly reduced (56%) in the invisible text condition. This is exactly the percentage by which the number of revisions was reduced in the invisible text condition. Notice that these two percentages need not have been the same. Although every R-burst is associated with a revision, most revisions are very brief and are not associated with R-bursts. In contrast, articulatory suppression had no effect on R-burst frequency.

Fifth, mechanical errors were increased dramatically (160%) in the invisible text condition. Articulatory suppression caused a smaller (21%) but statistically significant increase in mechanical errors.

Sixth, perceived task difficulty was increased both by articulatory suppression and by text invisibility.

And lastly, the time between the appearance of the cartoon and the initiation of writing (start time) was not influenced by any of the experimental variables.

These results indicate that articulatory rehearsal, a process that appears to correspond to the inner voice we experience when writing, is definitely involved in text production. Articulatory suppression

slowed the rate of writing, increased mechanical errors, sharply reduced P-burst length, and increased the perceived difficulty of the writing task. We can conclude that interfering with the inner voice interferes with writing.

To provide a theoretical interpretation of the results obtained in the current study as well as studies by Chenoweth and Hayes (2001) and Kaufer, Hayes, and Flower (1986), we propose the following narrative account of the text-production process. This account is represented graphically in Figure 6. It is based on the authors' model of text production (Chenoweth & Hayes, 2001, pp. 83-85). This model includes the following four processes:

- A proposer that proposes ideas for expression.
- A translator that converts the proposed ideas into linguistic strings.
- A transcriber that converts the linguistic strings into text.
- A reviser that evaluates and revises proposed and written language.

Moving From Intention to Text

At the beginning of a writing episode, the proposer has some general plan for writing. The plan may include information about communicative goals, content, order, audience, genre, and so forth. However, the plan may be a simple as "Here's what I want to write first, and after I write that, I'll figure out what to write next." Whatever the plan, the proposer selects an initial package of ideas and passes them on to the translator.

The proposer derives the ideas to be expressed through reflection on the plan and on information relevant to the plan. Relevant information may be in either linguistic or nonlinguistic form and may be accessed from memory or from external sources such as texts or pictures. In some cases, ideas accessed from text or speech may be ready to transcribe. For example, a student writing a research paper may include a quote stored in memory or in notes. In such cases, the materials can be passed directly to the transcriber. More generally, whether ideas are derived from linguistic or nonlinguistic sources, they will be translated into new language. For example, in writing a summary of a story, the writer formulates ideas by reading the story but expresses those ideas in language that may be entirely different from the language of the story. Of course, when ideas are derived from a linguistic source, the source language may facilitate the creation of new text by suggesting words or phrases or an organization of ideas that can be

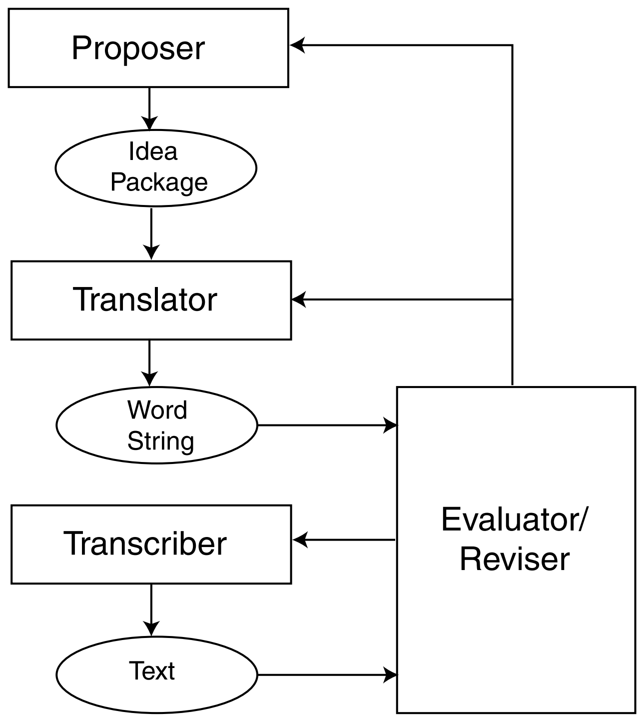

Figure 6. A Model of text production

incorporated in the new text. In the present study, ideas were presented in visual form to minimize such linguistic facilitation.

The proposer (because it does not deal with articulated language in this study) should not be influenced by articulatory suppression. If it were influenced by articulatory suppression, we would expect that start time, that is, the time required for the writer to begin processing the first package, would be increased in the articulatory suppression conditions. We found no evidence of such an effect. Similarly, the visibility of the text cannot initially influence start time because there is no text. Indeed, we found no effect of text visibility on start time.

When the translator receives the package, it translates the ideas into a linguistic string, selecting lexical items, ordering them, and supplying inflections consistent with the information in the package.

To translate a package, the translator draws on its long-term memory and working memory resources. In particular, it draws on linguistic knowledge stored in long-term memory, on the storage capabilities of the phonological short-term store, and on the articulatory rehearsal process (the inner voice). When the translation of a package is complete, there will be a pause in text production while a new package is readied for translation. These pauses mark the boundaries of P-bursts.

The package of ideas that the proposer sends to the translator may or may not be large enough to allow for the writing of a whole sentence. Kaufer et al. (1986) found that writers often stopped in mid-sentence and made comments to themselves such as, “What do I want to say?” (p. 126). In the present study, 14 of the 18 participants indicated that they did not always have a complete sentence in mind when they began typing their cartoon descriptions.

The number and complexity of the ideas in the package is determined by the translator’s capacity for finding language to encode complex ideas. (We assume that the proposer will not send, and the translator will not accept, larger packages than the translator can typically handle.) The translator’s capacity is determined, in part, by linguistic knowledge and by the available working memory resources. Thus, the size of these packages, and therefore the length of P-bursts, is larger the greater the writer’s linguistic knowledge and access to articulatory rehearsal. Chenoweth and Hayes (2001) showed that P-burst length increases with linguistic experience. The present study shows that articulatory suppression substantially reduces the length of P-bursts and increases the number of P-bursts per sentence. In contrast, articulatory suppression did not affect the number of R-bursts per sentence but text visibility did. R-bursts were significantly more frequent when the text was visible than when it was not. These results suggest that articulatory suppression interferes with the translator, which we hypothesize determines P-burst length, but not the reviser, which determines R-burst length. The absence of an effect of articulatory suppression on the reviser is consistent with results by Daneman and Stainton (1991), who found that articulatory suppression did not influence performance on a proofreading task.

The translated package may then be passed to the reviser, where it is examined for acceptability. The proportion of the proposed bursts that are accepted by the reviser depends on the linguistic competence of the writer. As Chenoweth and Hayes (2001) have shown, this proportion increases with linguistic experience. For native speakers, the

proportion ranges from .75 to about .90, depending on the complexity of the writing task. We do not have comparable data on acceptance rates in this study because we do not have protocol data that could identify language that the writer proposed but did not write.

If the reviser finds the package acceptable, the transcriber writes the text and the production process begins again with the selection of a new package of ideas. However, once text has been written, that text can interact with the text-production processes in at least two ways. First, the written text may trigger revision. The results of this study indicate that the revision process is influenced by text visibility. We have found that revision is most frequent when the text is visually present. However, some revision occurs even when it is not visible, presumably supported by a memory trace of the text or the intention that produced it.

Second, the text may be used to support the generation or the translation of the next text segment. Kaufer et al. (1986) found that writers frequently reread the text just written as they were preparing to write the next text segment.

We have found main effects on writing rate for both articulatory suppression and for text visibility. We have attributed the effect of articulatory suppression to the limitation of memory resources in the translator. Furthermore, we have attributed the effect of text visibility on rate to the removal of visual stimuli that trigger revision. So far, so good.

However, there is a problem. We have also identified a significant interaction between the effects of articulatory suppression and text visibility on rate. The effect of articulatory suppression was greater when the text was invisible than when it was visible. But, if the effects of these variables were mediated entirely by separate processes, it is not obvious how to account for the interaction.

We propose that the interaction actually occurs within the translation process when the translator is using the text just written to guide the translation of the next package. We assume that it makes use of information from the text just written to shape the next text segment. To maintain this information, the writer could rely on external memory provided by the visible text, on the articulatory rehearsal process, or both. We would expect that interference with articulatory rehearsal (articulatory suppression) would have a greater effect when the text is invisible than when it is visible. That is the result that we found. Notice that this interpretation implies that making the text invisible does have an effect on the translation process but that the effect occurs

only after some text has been written and that it is greatest during articulatory suppression.

In this study, articulatory suppression reduced writing time and rate but not sentence length or complexity. In the instructions, we emphasized the importance of writing complete descriptions of the cartoons and did not emphasize speed. If instead we had emphasized the importance of speed rather than completeness, articulatory suppression might have reduced sentence length and complexity.

SUMMARY

The results of this study show that articulatory rehearsal, which appears to correspond to the inner voice we experience when writing, plays an important role in the writing process. In particular, it plays a role in the translation process that converts ideas into language. This result has a practical consequence because it suggests that any factor that interferes with articulatory rehearsal, for example, being exposed to others speaking or to song lyrics, will make writing more difficult. Indeed, Madigan, Johnson, and Linton (1994) have found that exposing writers to irrelevant speech reduces the quality of the texts they write.

The results also show that being able to see the text as it is written facilitates writing. Making the text invisible seriously interfered with revision, increased mechanical errors, and increased the perceived difficulty of the writing task. This result also has a practical consequence because it suggests that anything that interferes with text visibility, for example, writing with a device that has a very small or hard-to-read screen, will make writing more difficult.

Finally, the results have allowed us to provide a more specific theoretical account of the processes by which writers move from intention to text.

NOTES

1. cT is a programming language developed at Carnegie Mellon University by David Andersen, Bruce Sherwood, Judith Sherwood, and Kevin Whitley.

2. All but the first three of these measures were the fluency measures used by Chenoweth and Hayes (2001).

3. We chose to rank sentences in the visible and invisible text conditions separately because we felt that the obvious difference in the number of mechanical errors between the two conditions might lead judges to focus only on that aspect of sentence quality. Because we already had a measure of mechanical errors, we thought it would be useful if judges were able to focus on other aspects of sentence quality.

4. In all cases, the significance of pair-wise contrasts was evaluated by Fisher's PLSD test.

5. The absence of differences due to block of presentation is important because it indicates that the content of the cartoons themselves, which were presented to all participants in a fixed order, did not interact with the experimental conditions. Put positively, the effects that we found were representative of both participants and cartoons (Clark, 1973).

6. Chenoweth and Hayes (2001) did not find a difference in the length of R-bursts and P-bursts.

REFERENCES

Baddeley, A. D., & Hitch, G. J. (1974). Working memory. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 8, pp. 47-89). New York: Academic Press.

Baddeley, A. D., & Lewis, V. J. (1981). Inner active processes in reading: The inner voice, the inner ear, and the inner eye. In A. M. Lesgold & C. A. Perfetti (Eds.), *Interactive processes and reading* (pp. 107-129). Hillsdale, NJ: Erlbaum.

Chenoweth, N. A., & Hayes, J. R. (2001). Fluency in writing: Generating text in L1 and L2. *Written Communication*, *18*, 80-98.

Clark, H. H. (1973). The language-as-fixed-effect fallacy: A critique of language statistics in psychological research. *Journal of Verbal Learning and Verbal Behavior*, pp. 335-359.

Daneman, M., & Stainton, M. (1991). Phonological recoding in silent reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*, 618-632.

Elbow, P. (1994). About voice and writing. In P. Elbow (Ed.), *Landmark essays on voice and writing* (pp. xi-xlvii). Davis, CA: Hermagoras Press.

Gathercole, S. E., & Baddeley, A. D. (1993). *Working memory and language*. Hillsdale, NJ: Lawrence Erlbaum.

Kaufer, D. S., Hayes, J. R., & Flower, L. (1986). Composing written sentences. *Research in the Teaching of English*, *20*, 121-140.

Kellogg, R. T. (1996). A model of working memory in writing. In C. M. Levy & S. E. Ransdell (Eds.), *The science of writing: Theories, methods, individual differences and applications* (pp. 57-71). Mahwah, NJ: Lawrence Erlbaum.

Levy, C. M., & Marek, P. (1999). Testing components of Kellogg's multicomponent model of working memory in writing: The role of the phonological loop. In M. Torrance & G. C. Jeffrey (Eds.), *The cognitive demands of writing: Processing capacity and working memory in text production* (pp. 25-42). Amsterdam: Amsterdam University Press.

Longoni, A. M., Richardson, A. T. E., & Aiello, A. (1993). Articulatory rehearsal and phonological storage in working memory. *Memory & Cognition*, *21*, 11-22.

Madigan, R. J., Johnson, S. E., & Linton, P. W. (1994, June). *Working memory capacity and the writing process*. Paper presented at the American Psychological Society, Washington, DC.
Rosenthal, R., & Rosnow, R. L. (1991). *Essentials of behavioral research: Methods and data analysis* (2nd ed.). New York: McGraw-Hill.

N. Ann Chenoweth is an assistant professor of English at The University of Texas–Pan American, where she coordinates the first-year writing program. Her primary interests are in the development of writing ability, writing processes, and course assessment.

John R. Hayes is a professor of psychology at Carnegie Mellon University, where he is director of the graduate program. His primary interests are in theories of writing and empirical studies of writing.