

Multitasking as Skill Acquisition

Frank J. Lee (fjl@rpi.edu)

Department of Cognitive Science, Rensselaer Polytechnic Institute
Troy, NY 12180 USA

Niels A. Taatgen (niels@ai.rug.nl)

Artificial Intelligence, University of Groningen
Groningen, The Netherlands

Abstract

Multitasking allows people to cope with the ever changing and complex world that we live in. However, as much as cognitive scientists have learned about the details of human cognition, the question of how people acquire multitasking ability remains a mystery. In this paper, we argue that multitasking can be best understood as a product of skill acquisition. In particular, we describe *production composition*, a computational theory of procedural skill acquisition, which can account for the acquisition of multitasking skill. We explore this idea in this paper as part of our effort to develop a cognitive model of a simulated air-traffic controller Task.

Introduction

Multitasking is a critical ability that allows people to cope with and flourish in the complex world that we live in. However, as much as cognitive scientists have learned about the inner workings of human cognition, our ability to multitask remains a mystery. In this paper, we argue that we can best understand multitasking as a product of *production composition* (Taatgen & Lee, submitted), a computational theory of procedural skill acquisition that has been implemented within the ACT-R framework (Anderson & Lebiere, 1998). Production composition has been used successfully to account for skill acquisition in a wide variety of domains including language learning (Taatgen & Anderson, submitted) and individual differences in complex skill acquisition (Taatgen, 2001). We believe that it can also be used to account for the acquisition of multitasking skill.

Multitasking

Multitasking is the ability to handle the demands of multiple tasks simultaneously. At the most basic level, this may involve executing multiple perceptual-motor actions at the same time, such as moving your attention to the next lane and turning the steering wheel. At a more complex level, this may involve interleaving the steps of many complex tasks, such as shifting down to a lower gear while navigating a curve and carrying on a conversation.

Important insights into people's ability to multitask come from the dual-task performance literature. One such insight is that while there is some interference between the two tasks that are being performed (with a caveat regarding the modality of stimuli and responses), people can consciously trade off performing one task for the other (Wickens & Gopher, 1977). Another is that people's performances in both tasks depend highly on their skill in the individual tasks (Allport, Antonis, & Reynolds, 1972). That is, being skilled in one task allows a person to perform it and other tasks with negligible impact on the overall performance of both tasks. For example, a skill driver might have little difficulty talking with a friend while driving, whereas a novice driver might find it difficult.

Skill Acquisition

Anderson (1982) proposed a theory of skill acquisition in terms of transitioning from declarative knowledge to procedural knowledge through a process called *knowledge compilation*. Initially, knowledge is in declarative form and is interpreted. Interpreting declarative knowledge is slow and may lead to errors, especially if the relevant knowledge cannot be retrieved when needed or erroneous knowledge is retrieved instead. With practice, declarative knowledge is compiled into procedural knowledge and is fast and free of errors. Newell and Rosenbloom (1981) proposed an alternate theory of skill acquisition called *chunking* that became an important component of the Soar cognitive architecture (Newell, 1990). Within Soar, skill acquisition is a function of combining multiple procedures into a single procedure and converting the current goal context into a more specialized procedure.

Production Composition

Production composition is a theory of skill acquisition that incorporates aspects of both Anderson's and Newell and Rosenbloom's account. It involves compiling declarative knowledge into procedural knowledge and combining multiple procedures into a new procedure. Consider the process of retrieving

information from declarative memory in ACT-R, which is usually done in two steps. In the first step, a production rule¹ issues a request to declarative memory for a certain piece of knowledge, while in the next step another production rule acts on the retrieved knowledge. Production composition eliminates the retrieval process and creates a single production rule out of the two original rules while substituting the retrieval into this new production rule. Through this process, general rules can be specialized into task-specific rules.

Figure 1 graphically illustrates this process. Before production composition takes place, a production rule requests an instruction from declarative memory on what to do next. Declarative memory returns with an instruction that the Enter key should be pressed in the current context of “Land plane 3”. In response a production rule issues a motor command to press the Enter key, which initiates the motor system to actually press the Enter key. Production composition eliminates the retrieval from declarative memory and combines both of these production rules into one single rule, producing a task-specific production rule that issues a motor command to push the Enter key once it is in the context of “Land plane 3”.

Production composition has been used successfully to model learning in a simulated air traffic controller task (Taatgen & Lee, submitted), inflection of the English past tense (Taatgen & Anderson, submitted) and the German plural (Taatgen, 2001), and strategy development in the balanced-beam task (van Rijn, van Someren & van der Maas, submitted). Perhaps reflecting the utility of production composition, it has been incorporated into the current version of the ACT-R cognitive architecture. In the next section, we discuss how production composition can be used to account for the acquisition of multitasking skill.

Multitasking and Production Composition

In the example given in Figure 1, all steps are executed serially. According to the ACT-R theory, however, each of the different subsystems, Hand, Declarative Memory and Production, as well as the Visual and other sensory-motor systems can work asynchronously and in parallel (Byrne & Anderson, 1998). This is not always possible: sometimes one subsystem must wait for information from another. The goal of multitasking in such cases is to exploit these gaps in processing by slipping in other useful processes. The production composition mechanism is capable of modeling this aspect of multitasking. Figure 2 gives a graphical example of this process, in which two tasks have to be carried out: Task

¹ Within ACT-R and other cognitive theories, procedural knowledge is often represented by *production rules* that have the form of “IF-THEN” condition-action rules.

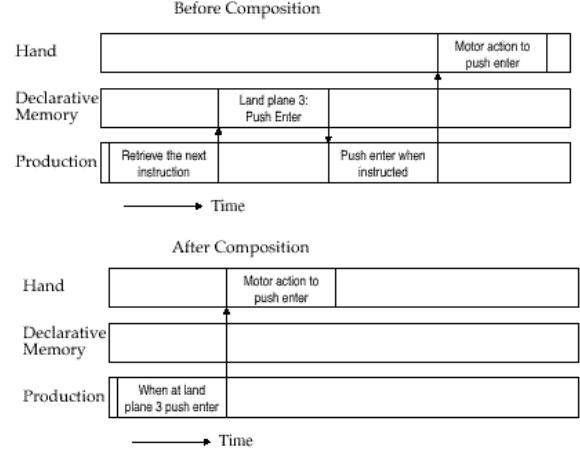


Figure 1: The production composition process.

A to push the Enter key and then check whether a light has gone on, and Task B to say “yes”. Obviously these tasks do not make any sense without context, but consider them as part of a larger task, for example in context of the simulated air-traffic controller task that we will discuss later on.

The top left panel of Figure 2 depicts Task A for the novice, comparable to Figure 1. First, a production rule issues a request for the next instruction. Declarative memory produces the instruction to push the Enter key. Next, a production rule issues a motor command to the Hand to do this. Although it takes the motor system some time to execute this command, a production rule immediately fires to retrieve the next instruction.

The retrieved instruction requests the visual system to check the light after the button has been pushed. The production rule that carries out this request has to wait for the instruction and the completion of the previous motor command. Only then can it issue a command to the visual system to check the light. Task B has a similar, although slightly simpler structure: an instruction is retrieved, after which the speech system is instructed to say, “Yes”.

If both task A and B rely on declarative instructions, it is impossible to carry them out concurrently because declarative memory is busy almost all of the time. Once production composition has taken care of some of the declarative retrievals, multitasking is possible. Suppose task A has been composed into task-specific production rules but not task B. Now task A is carried out as in the top right panel of Figure 2: a rule issues the motor command, the motor command is carried out, after which a production rule issues the perceptual command. But now there is time left after the first production rule to do something else, for example slip in task B. The bottom left panel shows how this is done. After a rule has fired to initiate the Hand command, a new rule fires

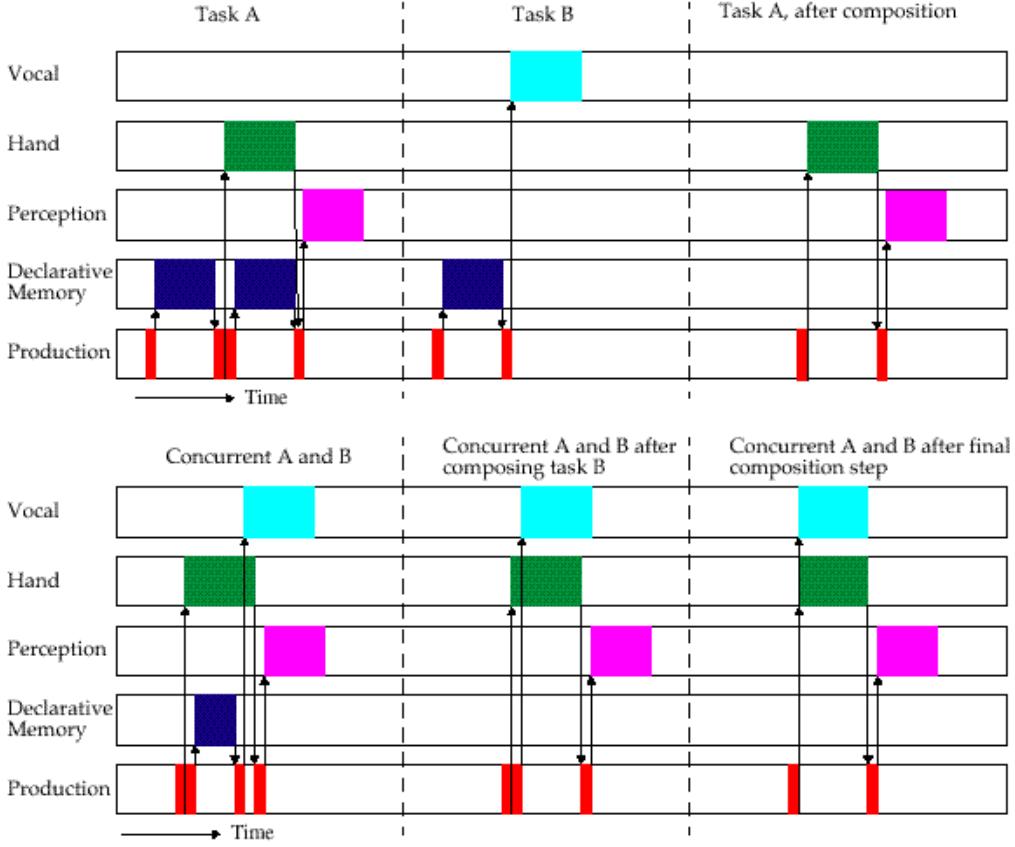


Figure 2: Development of multitasking production rules through production Composition.

to retrieve the instruction for task B. Depending on whether the Hand or retrieval from declarative memory is faster (we chose declarative memory in this case), the Vocal or the Perceptual command is issued. The other follows directly afterwards. The composition process does not stop here, because the retrieval in task B can also be eliminated, producing the situation in the bottom middle panel. Finally, the rule that initiates task A can be combined with the rule that initiates task B, producing the final state of the bottom right panel of Figure 2.

The Task

The task that we use in this paper to explore the concept of multitasking as skill acquisition is the Kanfer-Ackerman Air Traffic Controller (KA-ATC) Task (Ackerman, 1988; Ackerman & Kanfer, 1994). The KA-ATC task is composed of the following elements displayed on the screen: (a) 12 hold positions, (b) 4 runways, (c) information on current score, landing points, penalty points, conditions of the runways, and wind direction and speed, (e) a queue of planes waiting

to enter the hold, and (f) 3 message windows, 1 for notifying of weather changes, 1 for providing feedback on errors, and 1 for displaying of the rules of the task in response to information requests by the participants. The 12 hold positions are divided into 3 levels corresponding to altitude, with hold level 3 being the highest and hold level 1 being the lowest. A typical display of the KA-ATC task is presented in Figure 3.

Six rules govern participant's actions in this task: (1) Planes must land into the wind, (2) Planes can only land from hold level 1, (3) Planes can only move 1 hold level at a time, but to any open position in that level, (4) Ground conditions and wind speed determine the runway length required by different plane types. In particular, 747's always require long runways, DC10's can use short runways only when runways are DRY or WET (i.e. not ICY), and wind speed is less than 40 knots, 727's can use short runways only when the runways are dry or wind speed is 0-20 knots, and PROP's can always use short runways, (5) Planes with less than 3 minutes of fuel remaining must be landed immediately, and (6) Only one plane at a time can occupy a runway.

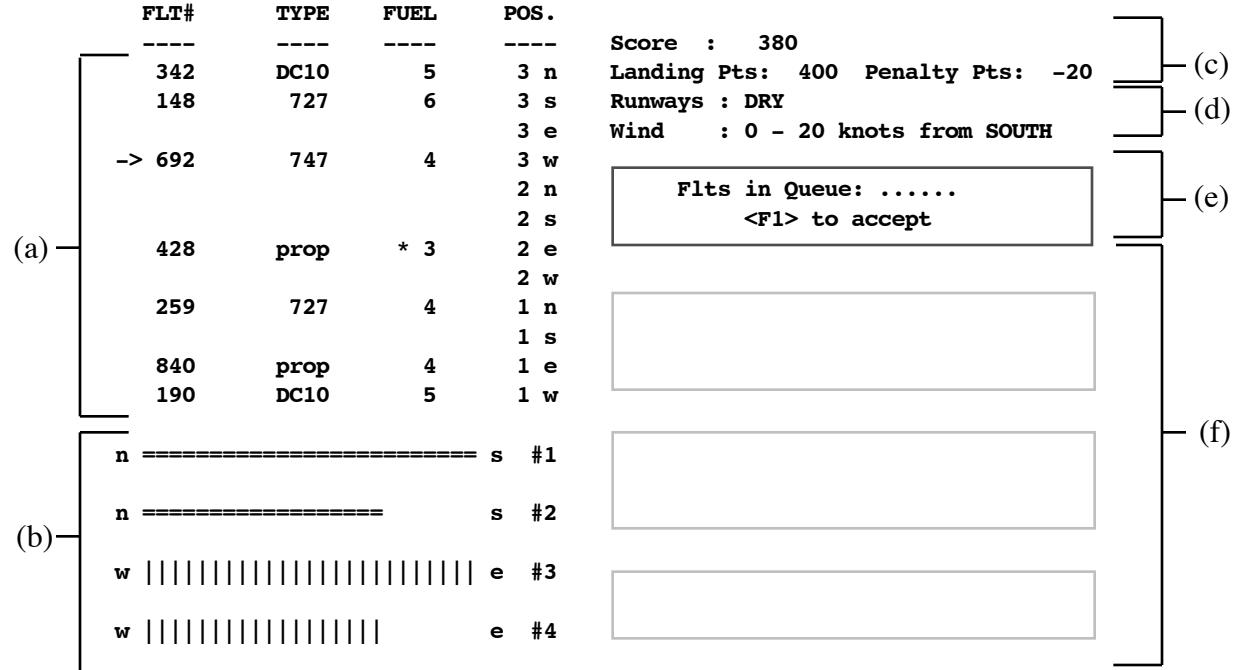


Figure 3: Kanfer-Ackerman ATC Task.

Participants can execute three actions in this task: (a) they can accept a plane from the queue into an open hold-position, (b) they can move a plane between the three hold-levels, and (c) they can land a plane on a runway. They can accomplish these actions by using four keys: the Up-arrow and the Down-arrow keys, \uparrow and \downarrow ; the F1 function key, **F1**; and the Enter key, \downarrow . They can move the cursor up and down the hold-positions and the runways using the \uparrow key and the \downarrow key. They can accept a plane from the queue into an open hold-position using the **F1** key. And, they can select a plane in the hold, place a selected plane in an open hold-position (either from the queue or from another hold-position), or land a plane on a runway using the \downarrow key. In addition, participants can press the number keys 1 - 6 to examine the rules 1 - 6 any time during the task.

Participants are given 50 points for landing a plane, penalized 100 points for crashing a plane, and penalized 10 points for violating one of the six rules. A plane crashes when the fuel-level of a plane falls to 0 minutes. Planes are added to the queue approximately every 7 seconds and it takes 15 seconds for a plane to clear a runway. Once planes enter the hold position from the queue, they have between 4 - 6 minutes of fuel and begin to lose fuel in real time.

In Ackerman (1988), participants performed in the fair-weather condition where the wind speed was fixed to 0 - 20 knots and the runway condition was fixed to

DRY. Under this condition, Rule 4 simplifies to the rule that all planes, except 747s, can land on a short runway.

The Model

Taatgen and Lee (submitted) have developed a model of the initial learning of the task. In this paper we describe a modification of this model to include some aspect of multitasking. The general idea of the model is that the participant in the experiment first encodes the instructions declaratively, forming a (often incomplete) plan on how to do the task. As interpretation of these instructions is slow, initial performance is also slow, resulting in poor performance. But due to production composition, a speed-up is realized that can account for the increase in performance.

The former model consisted of a fairly linear plan to land planes, to decide between the tasks to land planes, to move planes between hold levels and to get new planes from the queue. The model does not allow for much concurrency because of the linear structure of the plan that persists even after proceduralization. In order to test the new approach, we took out one aspect of the plan to land a plane, namely the checking of the wind direction. Wind direction has to be checked to see which of the runways can be used at the moment, and as it periodically changes, it has to be rechecked occasionally. In the original model, the wind was checked as one of the first steps in the landing procedure. We took this checking step out of the main plan, and made an alternative plan to check for the wind

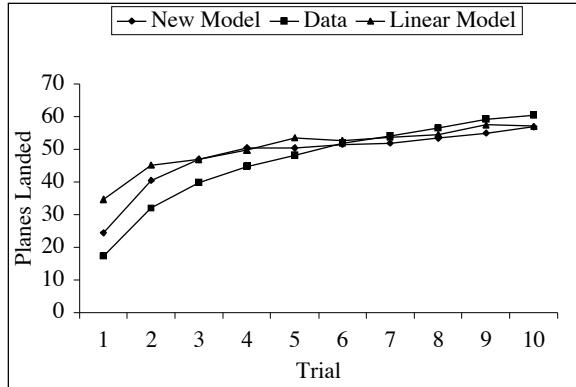


Figure 4: Number of planes landed.

at moments of “slack time”, for example when the arrow is moved to a plane or a runway. These arrow movements take multiple key-presses, allowing for some time for the model take a quick peek at the wind direction. This checking procedure will only succeed after the relevant steps have been proceduralized themselves, similar to the example in Figure 2.

The main question to be answered now is whether this change from a linear to a more parallel model improves the fit with the data. An interesting dependent measure in this respect is the time it takes for the participant (or model) to notice a change in wind direction. Although this cannot be measured directly (at least not in human participants), a measure (also used by Ackerman) is the elapsed time between a change in wind direction and the first landing of a plane on a runway in the new direction.

We compare the model predictions with data from Study 2 in the ONR data set (Ackerman & Kanfer, 1994), as reported in Ackerman (1988). The data from Study 2 were from 65 college undergraduates who completed 27 trials of the KA-ATC task with each trial lasting 10 minutes. For our model comparisons we only use trials 1 through 10, all in the fair-weather condition.

Figure 4 shows the overall score in terms of the number of planes landed in each 10-minute trial. Both the scores for the original “linear model” are shown (from Taatgen & Lee, submitted), and the predictions by the new model. Although the new model is more accurate than the old, linear model, the difference is slight. A larger difference can be seen in the time to notice a change in wind direction, the measure closely tied to the change in the model. Figure 5 shows the results. Although the linear model also predicts an improvement in this reaction time, because all processing is faster due to proceduralization, the new model matches the data much more closely especially in the first three trials where the model still has trouble interleaving checking the wind with other behavior.

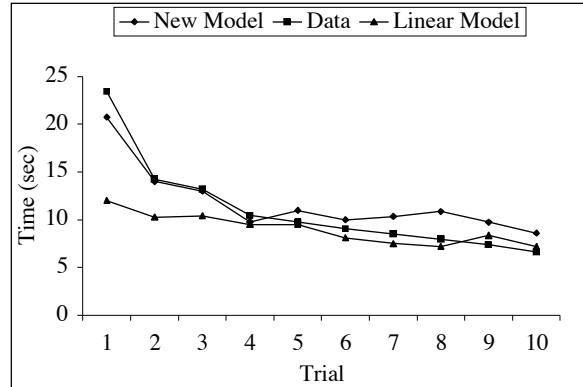


Figure 5: The elapsed time between a change in the wind direction and the first plane landed.

The current model is not a full implementation of the principle of multitasking within a complex task like the KA-ATC. It just demonstrates one aspect of in the act of checking the wind direction. Fortunately this aspect can be verified empirically. The current model is not yet capable of explaining improvements in performance after trial 10, where human participants still gradually improve but the model does not. This can only be explained within the ACT-R theory by a more efficient schedule of perceptual, cognitive and motor processes (Lee & Anderson, 2000).

Discussion

Multitasking and Planning

It is worthwhile in our discussion to see how multitasking and production composition might be related to other areas of human behavior. Especially relevant is the area of planning that researchers in Artificial Intelligence have looked at closely. The mechanism of production composition is an automated process that is below a person’s conscious control that automatically generates new procedural knowledge that are then tuned in their utility with use.

Planning, on the other hand, is largely seen as a deliberative and conscious process. In the context of multitasking, one can clearly imagine people reasoning about the structure of the multiple tasks that they must engage in, and explicitly devising a “plan” to interleave the tasks. This can happen at a larger time scale, such as when attempting to cook several dishes at the same time for a 7-course meal, or at a smaller time scale, such as trying to press the clutch and change the gear when learning to drive a car with a manual transmission.

From our perspective, planning and any other weak-method problem solving is completely consistent with production composition. Weak-methods, such as using instructions, examples, and planning, are all an aspect

of the declarative problem solving process that generates sequential actions that can be exploited by production composition to develop more efficient (i.e. multitasking) procedures.

We believe production composition can provide a resolution to the debate in the AI community between traditional planning versus reactive planning (c.f. Russell and Norvig, 1995). Traditional planning posits that agents reason over situations and actions in order to formulate a plan before taking the requisite actions. Reactive planning on the other hand posits that agents simply find the most applicable action in the current situation and executes them. Within ACT-R cognitive architecture, one can view weak-method problem solving as a mechanism for traditional planning and production composition as a mechanism for reactive planning. Weak methods generate sequential actions that are then used by production composition to generate reactive production rules. This makes perfect sense from the perspective of studying human behavior, since people display both types of planning.

Is Multitasking a Mystery?

Some may wonder, “Why is multitasking a mystery? Isn’t it a well understood phenomenon?” Our response is that while the fact that people are able to multitask is well understood, it is less clear how multitasking skill is learned. In particular, most of the current models of multitasking, with very few exceptions, include explicit control structures to switch between multiple tasks. However, the question still remains as to *how* these control structures came into being. We have argued in this paper that multitasking is a product of production composition and is able to learn new control knowledge for performing multiple tasks.

Conclusion

The learning in the KA-ATC task can be conceptualized as retrieving instructions from memory and executing them (Taatgen & Lee, submitted). The two steps are then compiled by production composition into a single production rule. One result of production composition in the KA-ATC task is the acquisition of production rules that encapsulates the execution of a keystroke-level task (c.f. Lee & Anderson, 2001). However, another result of production composition is the combining of two (or more) learned keystroke-level production rules into a single rule. This second result, we argue, reflects the acquisition of multitasking skill.

References

Ackerman, P.L. (1988). Determinants of individual differences during skill acquisition: Cognitive abilities and information processing. *Journal of Experimental Psychology: General*, 117, 288-318.

Ackerman, P.L., & Kanfer, R. (1994). *Kanfer-Ackerman air traffic controller task© CD-ROM database, data collection program, and playback program*. Office of Naval Research, Cognitive Science Program.

Allport, D.A., Antonis, B., & Reynolds, P. (1972). On the division of attention: A disproof of the single channel hypothesis. *Quarterly Journal of Experimental Psychology*, 24, 255-265.

Anderson, J.R. (1982). Acquisition of cognitive skill. *Psychological Review*, 89, 369-406.

Anderson, J.R. (1995). *Cognitive Psychology and Its Implications*. NY: Freeman.

Anderson, J.R., & Lebiere, C. (1998). *The Atomic Components of Thought*. Mahwah, NJ: Erlbaum.

Byrne, M. D., & Anderson, J. R. (1998). Perception and Action. In J. R. Anderson & C. Lebiere (Eds.), *The atomic components of thought* (pp. 167-200). Mahwah, NJ: Erlbaum.

Fitts, P.M. (1964). Perceptual-motor skill learning. In A.W. Melton (Ed.), *Categories of human learning*. New York, NY: Academic Press.

Lee, F.J. & Anderson, J.R. (2000). Modeling Eye-Movements of Skilled Performance in a Dynamic Task. In N.A. Taatgen, & J. Aasman (Eds.), *Proceedings of the Third International Conference on Cognitive Modeling*. Veenendaal, The Netherlands: Universal Press.

Lee, F.J. & Anderson, J.R. (2001). Does learning a complex task have to be complex? A Study in Learning Decomposition. *Cognitive Psychology*, 42, 267-316.

Newell, A. (1990). *Unified theories of cognition*. Cambridge, M.A.: Harvard University Press.

Newell, A., & Rosenbloom, P.S. (1981). Mechanisms of skill acquisition and the law of practice. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition* (pp. 1-55). Hillsdale, NJ: Erlbaum.

Taatgen, N.A. (2001). Extended the past tense debate: a model of the German plural. In K. Stenning & J. Moore (Eds.), *Proceedings of the Twenty-Third Annual Meeting of the Cognitive Science Society*. Mahwah, NJ: Erlbaum.

Taatgen, N.A., & Anderson, J.R. (submitted). Why do children learn to say “Broke”? A model of the past tense without feedback.

Taatgen, N.A., & Lee, F.J. (submitted). Production composition: A simple theory of complex skill acquisition.

van Rijn, H., van Someren, M, & van der Maas, H. (submitted). Modeling developmental transitions on the balance scale task.

Wickens, C.D., & Gopher, D. (1977). Control theory measures of tracking as indices of attention allocation strategies. *Human Factors*, 19, 249-366