The Effects of Performance Motivation: A Computational Exploration of a Dynamic Decision Making Task

James D. Brooks
Department of Industrial and Systems Engineering
Rensselaer Polytechnic Institute
Troy, NY 12180

Nicholas Wilson Ron Sun
Cognitive Science Department
Rensselaer Polytechnic Institute
Troy, NY 12180

***Abstract*—Effects of performance motivation on task performance have been observed empirically [1]. In this study, we explore one such result, the effects of assigned goals on performance of the well-studied Kanfer-Ackerman air traffic control task. We argue that the phenomenon observed in this task can be attributed to computationally expressible cognitive mechanisms defined by the CLARION cognitive architecture. In particular, we show that the performance variation between goal conditions can be explained by differences in explicit and implicit processing that are the result of external goal-setting.**

I. Introduction

Control of complex dynamic systems (dynamic decision making) is a task common in many practically and theoretically interesting scenarios [2]. One such dynamic control problem is that of scheduling. Scheduling problems are those where the goal is to temporally situate and manage discrete objects to achieve some objective [3]. A common paradigm for studying human scheduling performance is that of class scheduling [4]–[6]. Another is control of production or other dynamic systems [2], [7]–[9]. However, our research is focused on a specific scheduling paradigm that has received significant attention in applied psychology, the air traffic control task [10], [11].

The air traffic control task can be considered complex according to standard task complexity measures [6], [12]; it is also considered cognitively complex for a scheduling task [3], thus making it an interesting and substantial domain to study. Similar empirical findings, as obtained in [10], have been replicated using an human resources staffing task [13] and a class scheduling task [6]. This suggests that the underlying phenomena is significant and somewhat generalizable, making it a reasonable paradigm for simulation using a unified computational theory of cognition (such as CLARION).

In the next section of this paper, we describe the Kanfer Ackerman air traffic control (KA-ATC) task and the experimental result of interest. We will then discuss (in sections III and IV) the prior work in cognitive modeling and analyses of the task, the motivation for the current work, and the theoretical foundations. We then (in section V) present the details for how these theories can be computationally expressed using the CLARION cognitive architecture. The next section demonstrates how the model can be used to capture and provide explanations for the KA-ATC task performance by comparing the results of our simulations to the human data. Finally, we will conclude with a general discussion and suggest some implications for future work.

TABLE I
Rules for use of short runways.

Type	*Wind*		*Condition*
747	Never		Never
727	0-20 Knots	or	Dry
DC10	$<$ 40 Knots	and	Not Icy
Prop	Any		Any

Note: All aircraft can always use long runways.

II. Task Description

The task of interest is the well-used Kanfer-Ackerman air traffic control (KA-ATC) task [10], [14]. In this task, participants act as air traffic controllers, whose job is to land incoming aircraft safely and efficiently. These incoming aircraft arrive every seven seconds on average, and each aircraft must be maneuvered through a series of holding pattern levels (three levels, four positions in each level) before being assigned a runway. The runway that may be used is dependent on the aircraft type and the current weather and runway conditions as (shown in Table I). Furthermore, each aircraft is given only four to six minutes of fuel before it crashes. The following rules govern proper operation:

1) Aircraft must move down the holding pattern one level at a time
2) Aircraft must land into the wind
3) Aircraft must land from the bottom level of the holding pattern
4) Only one aircraft can be on a runway at any given time
5) The runway on which an aircraft can land is determined by the aircraft type and weather conditions (see Table I)
6) Aircraft with less than three minutes of fuel must land immediately

When an aircraft is called into the holding pattern from the incoming queue, it may be placed into any available holding

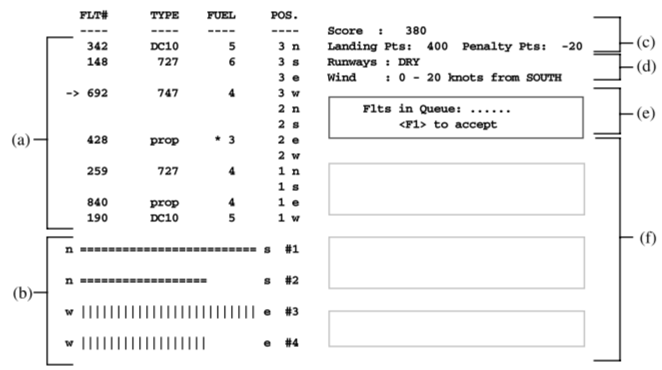

Fig. 1. KA-ATC task screen layout from [15]. The holding pattern is depicted in section (a), the runways in section (b), the current score in section (c), current conditions in section (d), incoming queue in section (e), and other messages and rules in section (f).

pattern positions. Once placed, the participant becomes aware of the aircraft type and the fuel remaining (4-6 minutes). No fuel is consumed while an aircraft is waiting in the incoming queue.

Once an aircraft is assigned a runway on which to land from the bottom level of the holding pattern (Level 1), it takes fifteen seconds for that aircraft to use and exit the runway. During this time, no other aircraft can be assigned to that runway (Rule #4).

Points are awarded to the participants based on the number of aircraft landed (+50 each). Points are deducted for each rule that is violated (-10 per violation or per minute of violation of Rule #6) and for each aircraft that crashes (-100 each).

Finally, we note that subjects performed the task at a computer terminal as depicted in Fig. 1.

Result

Of particular interest is the impact of performance motivation that was observed in Kanfer and Ackerman's experiment #3 [10]. In this experiment, subjects were divided into no-goal ("do your best") and difficult performance goal ("achieve a score of 2200") groups. The target achievement level for the performance goal group was chosen, based on earlier experiments, such that only 10% of the participants are expected to be capable of success. It was observed that both high- and low-ability participants (as determined by ASVAB scores) performed *worse* in the performance goal condition under procedural pre-training. The procedural pre-training enables the participants to learn the *mechanics* of the task (i.e., the required keystrokes), but not the rules that govern the system. This degradation in performance was initially thought to be a contradiction to goal-setting theory, and was thus an interesting result.

We restrict our work to this procedural pre-training case as we are more interested in the higher level cognitive dynamics rather than the motor control aspects. Lee and Anderson [16] showed that motor speedup accounts for a significant amount of performance improvement in an earlier study, providing further support for exploring only the procedural pre-training case (i.e., to avoid additional confounds).

Note that a second test condition, in which the participants were engaged in a pre-training session that focused on systematically learning the rules, did not exhibit the same performance detriment that was associated with the difficult goal condition. This test condition is not covered here (but see *Further Work* section).

Along with the theory presented in [10], both of these results can alternatively be explained by goal-setting theory [1], [5], [6]. The following section will briefly review the relevant aspects of this theory and will discuss how they can be computationally expressed using CLARION.

III. Prior Analysis and Modeling Work

Different aspects of the KA-ATC task have been well studied by other researchers. All of the work described below utilized data published by the Office of Naval Research which is no longer available. The study that was predominantly analyzed and modeled which used undergraduates as subjects and did not contain motivational considerations [14].

Lee *et al.* [17] identified two strategic indicators of performance evident in the data: initial hold level, and efficient use of runways on wind direction change (i.e., time from change to first aircraft landed). They found that half of the subjects learn to insert aircraft into the first level of the holding pattern only. This saves the participant six to twelve keystrokes for each landing. This "hold one" strategy was further explored by John and Lallement [18]. They consider the initial hold level for each aircraft over time as well as observed patterns of filling and emptying the hold levels. Three patterns are identified: stacked, sequential, and opportunistic. Stacked pattern is one where the participant fills the hold levels before landing them all. Sequential is the case where the participant attends to one aircraft at a time. Cases where the participant is attending to multiple aircraft at a time to take advantage of any slack time is called opportunistic behavior. This opportunistic behavior was found to significantly improve performance.

Shifts in strategy were also considered in [18]. Many participants were found to shift strategies throughout the trials, while some were not found to exhibit any consistent strategy. Both abrupt and gradual shifts were observed, however most of the shifts (78%) were gradual.

A detailed task analysis was performed in [16] which described functional, unit-level, and keystroke-level subtasks. The authors showed that learning was uniform across the identified subtasks. All of the subtasks identified (unit-task, functional, and keystroke) exhibited power law (i.e., consistent) learning. Lee and Anderson [16] performed another experiment using the KA-ATC task while capturing the output of eye-tracking software. Task-relevant regions of the screen were found to be aircraft type, fuel level, and runway. Though weather conditions are important, they are relatively static and therefore result in low gaze percentages.

An ACT-R model of skilled performance was created by Lee and Anderson [19]. This model captured only expert performance. A later ACT-R model has been created to account for learning of the task which implements a production composition theory of procedural skill learning [15], [20]. This theory, which is built on the idea of chunking [21] and knowledge compilation [22], posits that new productions are created from exiting procedures which result in a reduction in memory retrieval and rule specialization. In this way, the authors are able to model the speed-up learning which takes place in participants. The declarative memory and goal structures are built *a priori* using their earlier task analysis results [16]. The resulting model interacted directly with the task and maintained default ACT-R parameters. The authors note that only one possible strategy was implemented and that the declarative knowledge that was built into the model had to be learned by the participants during the early trials.

Implications

These prior results provide several important implications for the modeling work presented herein. Two forms of learning are found: strategy development, and speed up [16], [18]. Further, strategy shifts are predominantly gradual [18] suggesting an implicit learning process. Nonetheless, the existence of abrupt changes also suggests that learned strategies can be reasoned upon [18], and thus should be explicit.

The prior work which employed eye tracking during task performance suggest that the relevant inputs are aircraft type for each hold level, fuel level (only during early stages of learning), and runway status [16]. We will further include weather conditions as it is clear that though participants spend only a small percentage of the time looking at the weather area of the screen, knowledge of the current conditions is required. Note that in this work we are not concerned with the presentation of the information and requisite visual processing, but rather the dynamic decision making after these relevant inputs have been accurately perceived.

The work reviewed here did not consider the strategic learning process itself (i.e., how the strategies were developed) or the motivational impact that was observed in [10]. This is the focus of the study undertaken herein. We show how the theoretical work in goal-setting and anxiety effects map to motivational mechanisms and parameters in the CLARION cognitive architecture and explore the resulting learning and performance effects.

IV. Theoretical Background

In this section, we review the theoretical background underpinning the effects of motivation on learning. While Kanfer and Ackerman posit an attention theory [10], there has also been a motivation-based theory [6] which is more relevant to the modeling work undertaken in this study.

A. Goal-Setting Theory

Edward Lock and Gary Latham have been studying motivational effects of goal setting on performance for four decades [1], [23]. Notice that goal setting is a self-regulating mechanism [1]. This theory originally posited that performance is positively correlated with the difficulty of the specific goal. The theory has been expanded to suggest that in cases where the participant has yet to gain complete knowledge of how to perform the task a performance goal will decrease performance while a (difficult) learning goal will increase performance [5]. Goals have several mechanisms, one of which is to direct attention toward relevant cognitive activities and away from irrelevant activities. Several moderators of the goal-performance relation have been identified. Two relevant (and related) ones are commitment and self-efficacy. They stress that goals cannot only be "assigned", but must be committed to by the participant and that the participant's confidence in meeting the goal (self-efficacy) affects the personal goal and thus performance. In cases where the participant has low self-efficacy (i.e., the participant has low confidence in his or her ability to meet the goal), anxiety is developed which inhibits learning and subsequent performance [1].

This theory suggests that cases wherein a difficult performance goal is present and self-efficacy is high (i.e., participants have learned the task) that attention (i.e., explicit processing) will be allocated to the achievement of the goal and thus improve performance. However, when a task is not yet learned, this added attention inhibits learning and thus performance. The experiment of interest falls into this second category as the participants are relatively unfamiliar with the task.

B. Anxiety Effects

Several studies focusing on the effect of anxiety on task performance (e.g., see [24]–[28]) suggest that two (somewhat opposing) outcomes may occur as a result of elevated anxiety levels. The first possibility is that, in situations where anxiety levels are relatively low, individuals may choose to refocus their attention towards explicit control over a task. This increase in explicit monitoring [27], [29]–[31] has been shown to have two possible effects. In situations where tasks are naturally more explicit, performance may be improved. However, in situations where tasks are either well learned (i.e., proceduralized) or naturally more implicit, focusing additional attention on the explicit steps may have a deleterious affect on performance.

The second possible outcome of elevated anxiety levels (usually referred to as distraction theory [30]) is that, in situations where anxiety levels are especially heightened, cognitive resources must be reallocated in order to directly attend to the anxiety experience itself. Accordingly, task-related decisions must be made using more reactive (implicit) processes. The result is that task performance is typically hindered [26], [27]. It should be noted here that, while most studies on anxiety tend to confirm this outcome, other studies have also suggested that heightened anxiety levels may, in fact, have a somewhat different effect on learning tasks, especially when those tasks are naturally more implicit [28].

While these theories may seem opposing to one another, we have previously contended that they can actually be unified using the inverted U-curve theory ([24], [25], [32], [33]).The working hypothesis is that when anxiety increases, it leads the individual to become more controlled (more explicit) while making action decisions. However, when anxiety reaches a certain higher level, it can begin reducing control and the individual may revert to more automatic, implicit processes. Depending on the specific dynamics of a task, the effects of anxiety can either enhance or degrade learning and performance.

By combining relevant concepts from these theories, we can posit that, for the complex air-traffic control task, difficult performance goals should have an overall negative impact on learning and performance. However, for those individuals with high self-efficacy, anxiety should generally be lower than for other individuals (with lower self-efficacy). This allows them to apply the more instinctual (implicit) bottom-level processes, which ultimately improves performance (as implicit processes are more attuned to capturing the subtle complexities of the task). However, when heightened anxiety levels result from low self-efficacy, it forces individuals to shift their decision-making strategies towards more explicit processes. For the air-traffic control task, this increase in explicit processing has the effect of hindering the implicit learning mechanisms, thus degrading performance.

V. A CLARION-based Computational Model

The following subsections will outline the details of the proposed computational model and discuss how this model can be applied to the simulation and explanation of the participants' performance of the air-traffic control task using the theoretical framework presented previously.

A. The CLARION Cognitive Architecture

CLARION (the computational cognitive architecture) has previously been used as a means for computationally simulating and thus theoretically interpreting various existing findings related to explicit-implicit processing (i.e., cognitive control versus automaticity) in the context of motivational (i.e., drive level) changes [24], [25]. Therefore, it is our belief that CLARION can also be applied to performance-oriented anxiety-inducing situations.

CLARION is a well-established cognitive architecture (see, e.g., [34]–[38]) that is based on two basic assumptions: representational differences and learning differences of two different types of knowledge — implicit vs. explicit [36], [37] — among other essential assumptions and hypotheses [34].

CLARION takes note of the fact that the inaccessible nature of implicit knowledge is best captured by sub-symbolic, distributed representations (such as in a backpropagation network). It is widely agreed upon [39] that distributed representational units, for example in the hidden layers of a backpropagation network, are capable of accomplishing processing but are generally not individually meaningful. This characteristic of distributed representations, which renders the representational form less accessible, accords well with the relative inaccessibility of implicit knowledge [40], [41].

In contrast, explicit knowledge may be best captured in computational modeling by symbolic or localist representations [35], [37], in which each unit is more easily interpretable and has a clearer conceptual meaning. This characteristic of symbolic or localist representations captures the characteristic of explicit knowledge being more accessible and manipulable [35].

The dichotomous difference in the representations of the two different types of knowledge has led to a two-level architecture, whereby each level uses one kind of representation and captures one corresponding type of process (i.e., implicit or explicit). While this two-level structuring is the key foundation of CLARION, additional distinctions are also made.

The theoretical considerations from the previous section can be operationalized in CLARION using the motivational (MS), meta-cognitive (MCS), and action-centered (ACS) subsystems [34]. Note that, since the air-traffic control task is essentially novel to the human participants, we contend that mainly motivational, control, and procedural knowledge mechanisms are applied, and that declarative reasoning processes are not used in any significant way. Therefore, the current study focuses specifically on the interaction between motivation (in the MS) and decision-making (in the ACS) as well as the interaction between implicit and explicit processing within the action-centered subsystem (ACS).

B. Model Implementation

In the following section, we will discuss the implementation of the CLARION-based model. Each subsystem will be described in turn, followed by other considerations. Due to space limitations, the task environment itself will not be discussed.

1) Action Centered Subsystem (ACS): The action centered subsystem contains all action-related processing. The subsystem takes inputs from the environment and other subsystems to make decisions about appropriate courses of action. These actions can be either external (e.g., move object) or internal (e.g., memory storage/retrieval) [34].

The present simulation utilizes an implicit decision network (IDN) in the bottom-level of the ACS along with top-level explicit rules. The IDN uses simplified Q-learning and contains input, hidden, and output layers. The input layer is composed of dimension value (DV) pairs describing current situations, while the output layer is concerned with external actions. There are sixty-six input DV pairs and twenty seven output DV pairs. The input layer, for example, contains a dimension for each holding pattern position each of which contains two values: occupied or empty. See Table II for a complete listing of the input dimension value pairs. The output external actions are shown in Table III. Each output corresponds to an external action. The default number of hidden nodes was used (46 total).

The top-level of the ACS contains explicit action rules. These rules can be pre-loaded into an agent manually as well as learned bottom-up through experience using rule extraction [36]. The agents were given a limited set of action rules which correspond to the initial instruction given to the

TABLE II
INPUT DIMENSIONS AND THEIR ASSOCIATED VALUES (66 TOTAL).

Description	*Dimension*	*Values*
Aircraft(s) waiting in queue	queueTF	{true, false}
Holding Pattern positions	[hold{1-3}-{N,S,E,W}]	{occupied, empty}
Runway status	[{NS, EW}_{Short, Long}]	{occupied, empty}
Aircraft Type (level 1 only)	{Ntype,Stype,Etype,Wtype}	{prop, DC10, 727, 747}
Fuel level (level 1 only)	{Nfuel,Sfuel,Efuel,Wfuel}	{high, low}
Wind Speed	windSp	{0_20, 25_35, 40_50}
Wind Direction	windDir	{NS, EW}
Runway Conditions	Runways	{dry, wet, icy}

TABLE III
EXTERNAL ACTIONS IN OUTPUT LAYER OF ACS IDN (27 TOTAL).

Action	*Enumerations*
Bring in aircraft from queue (choose level)	{1, 2, 3}
Move aircraft down (levels 2, 3)	[dwn{2-3},{N,S,E,W}]
Land aircraft (from level 1 only)	[{N,S,E,W}_{NS, EW}_{Short, Long}]

TABLE IV
PRE-LOADED RULES CORRESPONDING TO INITIAL TRAINING.

Condition	*Action*
Aircraft waiting	Enter at {1,2,3}
Aircraft at x (levels 2,3)	Move down from x
Aircraft at x (level 1)	Land from x

TABLE V
COMBINED REWARD STRUCTURE.

Condition	*Reward*
Enqueue aircraft in level (3,2,1)	(0.6,0.75,0.9)
Move down from level (3,2)	(0.6, 0.8)
Landing	1*
Incorrect moves	0.2*
Crash	0*

*Note: Rewards with * are defined in the formal task structure. All remaining are distance based.*

participants. These rules describe available (i.e., potential) moves. For example, it only makes sense to attempt to move an aircraft (from a certain position) down in the hierarchy if that position is indeed occupied by an aircraft. The rules that were pre-loaded are detailed in Table IV.

Learning: The task has a series of built-in reinforcement mechanisms through the use of assigned scores. These scores are displayed directly to the participants in real time during the completion of the task. The reward structure in the model reflects this scoring system. In addition to these scores, it is likely that the participants imposed additional reward structure to the task. It was well known by the participants that the goal was to land aircraft quickly and thus it is reasonable to assume that having aircraft in the lowest level of the holding pattern would be perceived to be a better system state than having aircraft in the highest level in the pattern. Indeed the use of the lowest level (along with runway efficiency and reaction to weather changes) was positively correlated to high performance [17]. As a result, we utilize a distance reduction heuristic in addition to the explicit reward structure of the task [42]. This heuristic rewards successful moves in inverse proportion to the distance from the subject aircraft to the runways (the goal state for each aircraft). The details of this combined reward structure are shown in Table V. Notice that the reward is always in the interval [0,1] according to the default range of values in CLARION.

It is well known that learning (both implicit and explicit) continues to occur off-line after interaction with the particular task of interest (e.g., [43]). Robertson *et al.* [43] further show that explicit learning requires rapid eye movement (REM) sleep, while implicit learning occurs even without sleep time. Though the participants in [10] performed the trials on the same day, there was time between trials where surveys were completed which would have facilitated further implicit learning. We capture this effect in the model by training the bottom level network using random occurrences from the prior trial. Five times as many random training instances are used as was experienced during the trial. This implies that each iteration from the trial is expected to be seen during this offline training process five times on average. This offline learning was critical to achieving results comparable in magnitude to the human data. Evidence for offline learning in the earlier ATC study (see [14]) is also observed in [18].

Parameter Values: The majority of parameters within the ACS are left at their default values. Exceptions are: selection temperature, which was reduced from the default of 0.1 to 0.01 (more deterministic selection); enabling the agent to select top-level rules based on their past performance (as captured using match statistics); disabling explicit rule refinement; and the neural network learning rate was set to 0.5 (default: 0.1). Finally, the level probability parameters are varied by the motivation and meta-cognitive subsystems according to the test condition as described next.

2) Motivational, Meta-Cognitive Subsystems (MS & MCS): The motivation subsystem of the CLARION architecture implements an agent's internal drives and goals [44]. Within the CLARION architecture, each drive strength, ds_i, is updated according to the current stimulus, the perceived deficit, and the gain and bias parameters, α_i and β_i, according to the linear function $ds_i = \alpha_i(\text{stimulus}_i)(\text{deficit}_i) + \beta_i$.

Each drive in the MS is associated with goals. Goal strengths are determined by the drive strengths and the associated relevance weights. Using these goals, the MCS can direct attention through input filtering, modify learning method and reinforcement levels, change reasoning mechanisms, monitor other subsystems, and change parameters [44]. It is this last capability which we utilize for this model.

Drive strengths have also been used in previous CLARION models to modify parameters. For example, in Wilson *et al.* [24], [25], the strength of avoidance drives, which was hypothesized to capture the level of anxiety, is assumed to determine top-level selection probability (i.e., degree of cognitive control or amount of explicit processing). We propose a similar approach here.

This study focuses on a complex performance task which is likely to stimulate an approach drive (such as "recognition and achievement") when self-efficacy is high. Conversely, when self-efficacy is low, the same goal would be more likely to stimulate an avoidance drive. We can see that the parameter α_i for avoidance drives (which determine anxiety levels) can be inversely correlated to the theoretical concept of self-efficacy while the "stimulus" (for both approach and avoidance drives) represents the external goals provided to the participants (low "stimulus" for no goal condition and high "stimulus" for goal condition).

Because these effects are fixed for each participant for the scenario of interest, the end result of these combined drive and goal mechanisms (parameter changes) was enacted directly (through the use of hard-coded values for each goal condition). A more detailed and complete integration will be discussed in the *Further Work* section.

As described in the *Theoretical Background* section, goal-setting theory proposes that the end result of the different self-efficacy and personal (i.e., committed, internal) goal conditions is to affect the amount of attention (i.e., explicit processing) allocated to the task. Correspondingly, the parameter changes within our model are the top- and bottom-level probabilities. These parameters affect the probability of a given perception-action cycle choosing an action from each level. A broad range for these parameters was simulated to explore the performance effects which we will address in a more detailed technical report.

C. Timing Considerations

Each of the human trials in [10] was 10 minutes in duration. At the beginning of the trials, the average participant could only land 10 aircraft per trial. By the end of the sixth trial, an average of 35 aircraft could be landed by the high-ability participants. Each landing requires, at most, 4 correct actions from incoming queue to runway (accept into level 3, move down to 2, move down to 1, land). Participants made an average of 10 errors. This suggests at least 150 external actions per 10-minute trial for average performance. Another consideration is the time scale of the task itself. The incoming aircraft arrive at a rate of one every seven seconds on average. An aircraft remains on a runway for fifteen seconds after beginning to land. In order for the cognitive model to perceive the average aircraft arrival, seven second increments would be required. This suggests at least 90 perceptions per 10-minute trial. Thus, the actions required to land aircraft is the limiting factor. The simulations in this study use 175 perception-action cycles of the CLARION-based model to simulate each 10-minute trial.

VI. Simulation Results

First we briefly review the human data found in [10]. The left panel of Fig. 2 shows the mean number of landed aircraft per trial found in experiment three under the procedural pre-training condition. We observe that the mean performance changes from 8-18 aircraft landed in the first trial up to 28-35 in the final trial. The goal condition decreases performance of both high- and low-ability participants across trials.

We test the hypotheses regarding the effect of level probabilities on learning and performance described above through a series of simulations. These parameters are chosen to sum to one. For the difficult goal case, the top-level probability was 0.20 (bottom-level was 0.80). The no goal case was modeled with a top-level probability of 0.05 (bottom-level of 0.95). At first these may appear to be unusually skewed toward implicit processing. However, other researchers who have experienced the task first-hand note that at the beginning "it feels as though the system is driving you" [18]. This reaction suggests a very large amount of implicit processing.

For each condition, twenty agents were initialized with individual ability differences (modeled as random Boltzman distribution selection temperatures – randomly chosen uniformly between 0.01 and 0.03) and subjected to ten trials. As done in [10], the random agents were separated into two groups according to the median ability (selection temperature)

We regard the first four trials as a training set and the remaining six as true test trials. The right panel of Fig. 2 shows the mean number of aircraft successfully landed by the agents for each trial. The results show similar qualitative behavior between the model and the human data. In particular, the effect of a difficult performance goal (increasing explicit processing) leading to decreased performance in both high- and low-ability groups is observed. We are unable to provide detailed comparisons as the original data are unavailable.

One will further observe that the human subjects appear to learn more rapidly in the early trials than the computational agents. As one would expect, experimentation shows that the shape of these curves is affected by the number of implicit learning iterations performed between trials.

We finally note that the cognitive models crashed a comparable number of aircraft per trial (mean=7.9, std.dev=2.0) as the human participants.

VII. General Discussion

We have argued that the performance effects of Lock and Latham's Goal-setting theory can be explained by a level (i.e., implicit vs. explicit) probability cognitive mechanism. Further, this study has shown how this mechanism can be

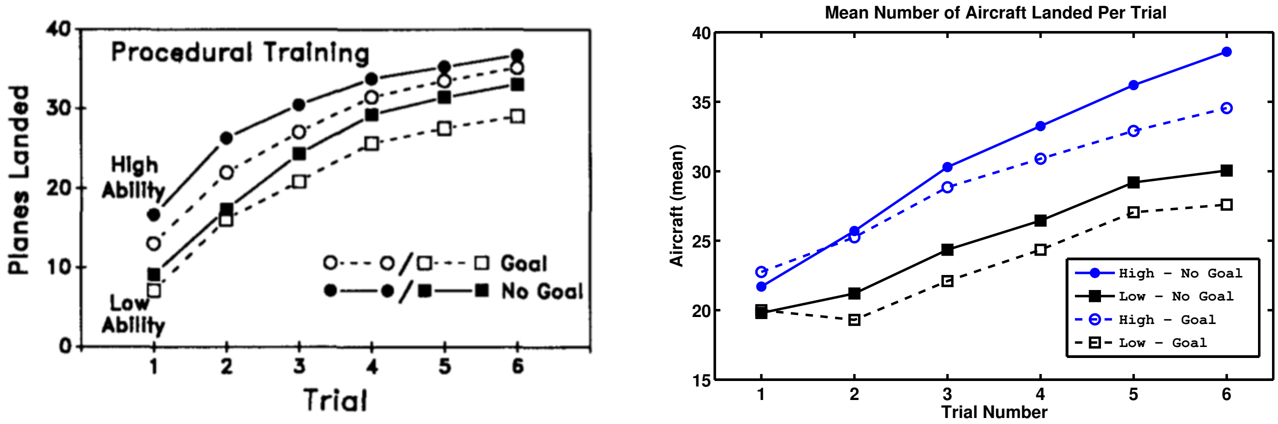

Fig. 2. Mean landing results from human trails from [10] (left) and of CLARION agents for each condition (right). The solid curves show the mean number of aircraft successfully landed in each trial for the *no goal* condition for high- (circle) and low-ability (square) groups. Similarly, the dashed curves show the data for the *goal* condition.

enacted through the use of drives. The CLARION cognitive architecture provides a framework for showing plausibility of these motivational mechanisms in explaining observed learning and performance behaviors.

The results provided suggest that additional explicit attention on a task before a significant number of rules that can enhance the task performance have been generated hinder learning and thus future performance. Further work is needed to understand this tradeoff between state of knowledge (i.e., number of rules) and the benefit of explicit processing for high performance.

A. Limitations

There are a couple of limitations of the current model which should be discussed. Though we have captured initial bottom-up learning, the full development of strategies as they are discussed in the literature (e.g., [18]) are not included. As strategies are learned, interaction with the non-action centered subsystem (NACS) would become necessary and these interactions would be included as outputs of the action centered subsystem (ACS). Several inputs have been omitted for simplicity (and are largely justified in [16]). These include fuel and aircraft type at levels two and three in the holding pattern. Further, the detailed point deductions for low fuel as well as the textual feedback mechanism present in the task for incorrect responses is not used here. The implementation has been simplified to avoid certain types of error that a participant may have been able to make (e.g., landing an aircraft from level 2, moving more than one level at a time) or to simplify an operation (e.g., specifying which position to move down to).

Despite these limitations, the model presented herein shows the plausibility of the level probability cognitive mechanism in producing the observed effects. Further work to include the items mentioned above as well as correspondingly more detailed analysis would improve the insight to be gained from this model.

B. Further Work

The model and results presented here show promising results that suggest that the differences in performance due to the presence of a goal can be explained by a cognitive mechanism which modulates the explicit processing devoted to the task. Further analysis could be done which examines the strategies exhibited by the agents to compare to the human strategies reported in [18] as well as an exploration of the extracted rules. Addressing the limitations discussed earlier would build on this foundation to more fully understand strategy formation and motivational impacts on performance.

Individual differences and the moderating effects of ability have been modeled in the present work using differences in the Boltzman distribution temperature setting. Additional individual differences should be explored with include varied learning rates, stress responses, probability of rule encoding, and so on.

A further integration of the model with the complete goal-setting theory would allow extensions to model the benefits of a difficult learning goal. We propose that this integration include a representation of the agent's self-efficacy and goal production. Self-efficacy can be modeled as an internal state of the meta-cognitive subsystem which can be derived from monitoring the state of the top-level of the ACS (e.g., number of existing rules) in addition to the assigned goal. A conceptual representation of this calculation is shown in Fig. 3.

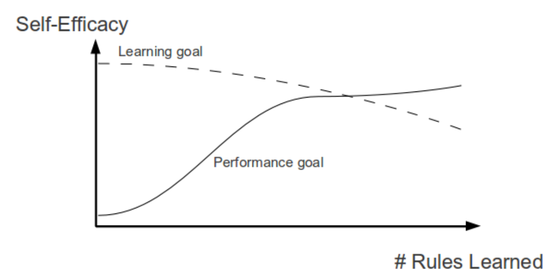

Fig. 3. Conceptual representation of self-efficacy calculation in the MCS.

Individual differences also occur in the perceived deficit levels for drives. Kanfer and Ackerman [10] suggest that high ability individuals are more likely to set difficult goals for themselves without an externally imposed goal, thus they have larger perceived deficits. It is important to note that this deficit is in relation to expected performance regardless of the stimulus (assigned goal).

References

[1] E. A. Locke and G. P. Latham, "Building a practically useful theory of goal setting and task motivation: A 35-year odyssey." *American Psychologist*, vol. 57, no. 9, pp. 705–717, 2002.

[2] M. Osman, "Controlling uncertainty: A review of human behavior in complex dynamic environments," *Psychological Bulletin*, vol. 136, no. 1, pp. 65–86, 2010.

[3] J. Cegarra, "A cognitive typology of scheduling situations: A contribution to laboratory and field studies," *Theoretical Issues in Ergonomics Science*, vol. 9, no. 3, pp. 201–222, May 2008.

[4] P. Earley, "Influence of information, choice and task complexity upon goal acceptance, performance, and personal goals." *Journal of Applied Psychology*, vol. 70, no. 3, pp. 481–491, 1985.

[5] D. Winters and G. P. Latham, "The effect of learning versus outcome goals on a simple versus a complex task," *Group & Organization Management*, vol. 21, no. 2, pp. 236–250, 1996.

[6] G. P. Latham, G. Seijts, and D. Crim, "The effects of learning goal difficulty level and cognitive ability on performance." *Canadian Journal of Behavioural Science*, vol. 40, no. 4, pp. 220–229, 2008.

[7] K. McKay, "An information systems based paradigm for decisions in rapidly changing industries," *Control Engineering Practice*, vol. 3, no. 1, pp. 77–88, 1995.

[8] B. MacCarthy, J. Wilson, and S. Crawford, "Human performance in industrial scheduling: A framework for understanding," *Human Factors and Ergonomics in Manufacturing*, vol. 11, no. 4, pp. 299–320, 2001.

[9] P. P. Stoop and V. C. Wiers, "The complexity of scheduling in practice," *International Journal of Operations & Production Management*, vol. 16, no. 10, pp. 37–53, 1996.

[10] R. Kanfer and P. L. Ackerman, "Motivation and cognitive abilities: An integrative/aptitude-treatment interaction approach to skill acquisition." *Journal of Applied Psychology*, vol. 74, no. 4, pp. 657–690, 1989.

[11] S. Kozlowski and B. Bell, "Disentangling achievement orientation and goal setting: Effects on self-regulatory processes." *Journal of Applied Psychology*, vol. 91, no. 4, pp. 900–916, 2006.

[12] R. E. Wood, "Task complexity: Definition of the construct," *Organizational Behavior and Human Decision Processes*, vol. 37, no. 1, pp. 60–82, 1986.

[13] M. Mone and C. Shalley, "Effects of task complexity and goal specificity on change in strategy and performance over time," *Human Performance*, vol. 8, no. 4, pp. 249–262, 1995.

[14] P. L. Ackerman, "Determinants of individual differences during skill acquisition: Cognitive abilities and information processing," *Journal of Experimental Psychology*, vol. 117, no. 3, pp. 288–318, 1988.

[15] F. J. Lee and N. A. Taatgen, "Multitasking as Skill Acquisition," in *Proceedings of the 24th Annual Conference of the Cognitive Science Society*, 2002, pp. 572–577.

[16] F. J. Lee and J. R. Anderson, "Does learning a complex task have to be complex? A study in learning decomposition." *Cognitive Psychology*, vol. 42, no. 3, pp. 267–316, May 2001.

[17] F. Lee, J. Anderson, and M. Matessa, "Components of dynamic skill acquisition," in *Proceedings of 17th Annual Conference of the Cognitive Science Society*, 1995, pp. 506–511.

[18] B. E. John and Y. Lallement, "Strategy use while learning to perform the Kanfer-Ackerman Air Traffic Controller c 1 task," in *Proceedings of the 19th Annual Conference of the Cognitive Science Society*, 1997, pp. 337–342.

[19] F. Lee and J. Anderson, "Modeling eye-movements of skilled performance in a dynamic task," in *Proceedings of the Third International Conference on Cognitive Modeling*, 2000, pp. 194–201.

[20] N. A. Taatgen and F. J. Lee, "Production compilation: A simple mechanism to model complex skill acquisition," *Human Factors: The Journal of the Human Factors and Ergonomics Society*, vol. 45, no. 1, pp. 61–76, 2003.

[21] A. Newell and P. Rosenbloom, "Mechanisms of skill acquisition and the law of practice," in *Cognitive skills and their acquisition*, J. R. Anderson, Ed. Hillsdale, NJ: Erlbaum, 1981, pp. 1–55.

[22] J. Anderson, "Acquisition of cognitive skill," *Psychological Review*, vol. 89, no. 4, pp. 369–406, 1982.

[23] E. A. Locke and G. P. Latham, *A theory of goal setting and task performance*. Englewood Cliffs, NJ: Prentice Hall., 1990.

[24] N. R. Wilson, R. Sun, and R. C. Mathews, "A motivationally-based simulation of performance degradation under pressure," *Neural Networks*, vol. 22, no. 5-6, pp. 502–8, 2009.

[25] N. Wilson, R. Sun, and R. Mathews, "A motivationally based computational interpretation of social anxiety induced stereotype bias," in *32nd Annual Conference of the Cognitive Science Society*, 2010, pp. 1750–1755.

[26] A. Lambert, B. Payne, L. Jacoby, L. Shaffer, A. Chasteen, and S. Khan, "Stereotypes as dominant responses: On the "social facilitation" of prejudice in anticipated public contexts." *Journal of Personality and Social Psychology*, vol. 84, no. 2, pp. 277–295, 2003.

[27] S. Beilock and T. Carr, "On the fragility of skilled performance: What governs choking under pressure?" *Journal of Experimental Psychology: General*, vol. 130, no. 4, pp. 701–725, 2001.

[28] S. Beilock, C. Kulp, L. Holt, and T. Carr, "More on the fragility of performance: Choking under pressure in mathematical problem solving." *Journal of Experimental Psychology: General*, vol. 133, no. 4, pp. 584–600, 2004.

[29] E. Langer and L. Imber, "When practice makes imperfect: Debilitating effects of overlearning." *Journal of Personality and Social Psychology*, vol. 37, no. 11, pp. 2014–2024, 1979.

[30] B. Lewis and D. Linder, "Thinking about choking? attentional processes and paradoxical performance," *Personality and Social Psychology Bulletin*, vol. 23, no. 9, pp. 937–944, 1997.

[31] R. S. Masters, "Knowledge, knerves and know-how: The role of explicit versus implicit knowledge in the breakdown of a complex motor skill under pressure," *British Journal of Psychology*, vol. 83, no. 3, pp. 343–358, 1992.

[32] R. Yerkes and J. Dodson, "The relation of strength of stimulus to rapidity of habit-formation," *Journal of Comparative Neurology and Psychology*, vol. 18, no. 5, pp. 459–482, 1908.

[33] L. Hardy and G. Parfitt, "A catastrophe model of anxiety and performance," *British Journal of Psychology*, vol. 82, no. 2, pp. 163–178, 1991.

[34] R. Sun, "A tutorial on CLARION 5.0," Rensselaer Polytechnic Institute, Tech. Rep., 2003.

[35] ——, *Duality of the mind: A bottom-up approach toward cognition*. Lawrence Erlbaum, 2002.

[36] R. Sun, E. Merrill, and T. Peterson, "From implicit skills to explicit knowledge : A bottom-up model of skill learning," *Cognitive Science*, vol. 25, pp. 203–244, 2001.

[37] R. Sun, P. Slusarz, and C. Terry, "The interaction of the explicit and the implicit in skill learning: A dual-process approach." *Psychological review*, vol. 112, no. 1, pp. 159–92, Jan. 2005.

[38] R. Sun and I. Naveh, "Social institution, cognition, and survival: A cognitive–social simulation," *Mind & Society*, vol. 6, no. 2, pp. 115–142, 2007.

[39] D. Rumelhart, J. McClelland, and The PDP Research Group, *Parallel distributed processing: Explorations in the microstructures of cognition*. MIT Press, 1986.

[40] A. Cleeremans, A. Destrebecqz, and M. Boyer, "Implicit learning: News from the front," *Trends in Cognitive Sciences*, vol. 2, no. 10, pp. 406–416, 1998.

[41] A. Reber, "Implicit learning and tacit knowledge." *Journal of Experimental Psychology: General*, vol. 118, no. 3, pp. 219–235, 1989.

[42] F. Dandurand, T. R. Shultz, and A. Rey, "Including cognitive biases and distance-based rewards in a connectionist model of complex problem solving," *Neural Networks*, vol. 25, pp. 41–56, 2012.

[43] E. M. Robertson, A. Pascual-Leone, and D. Z. Press, "Awareness modifies the skill-learning benefits of sleep," *Current Biology*, vol. 14, no. 3, pp. 208–212, 2004.

[44] R. Sun, "Motivational Representations within a Computational Cognitive Architecture," *Cognitive Computation*, vol. 1, no. 1, pp. 91–103, 2009.