ANNEX

WPP: A Dynamic decision making task

In WPP the decision maker plays the role of an operator of a Water Purification Plant, whose goal is to distribute water to different locations on time. To do the task, the decision maker activates/de-activates pumps to distribute water with maximum number of activated pumps restricted to 5. The simulation runs from 2:00 PM to 10:00 PM or 8 simulation hours labeled one trial, and the simulation time is shown in the top left corner of the screen. The pace of the simulation is defined before the trial starts and an eight-minute trial, for example, indicates that each simulated hour will run in one real minute. The slowest pace is thirty real minutes, indicating that each simulated hour will run in 3.75 minutes.

The description in this section makes reference to Figure 2 in the paper. Throughout one simulation trial, the decision maker needs to meet several deadlines. Meeting a deadline means that all of the water in the tanks leading to that deadline has been pumped out by the operator.

There are 23 tanks in the system interconnected in a tree structure. Consecutive tanks such as tank 3, tank 15, and tank 21 form a chain and thus have the same deadline, 8:00PM. For example, the deadline for tanks 0 and 2 is 5:00 PM and the deadline for tank 1 is 8:00 PM. Each tank has two pumps that may be activated/deactivated to transfer the water to the next tank. A pump may be activated by clicking on it once with the mouse while it is idle (red color), and de-activated by clicking on it while it is working (green color). The simulation provides an indicator to track the number of pumps in use. Activating both pumps on a side of a tank will double the pumping speed. Each of the pumps delivers water at a rate of one gallon every two minutes; however, when two pumps in one tank are active the delivery rate is one gallon per minute.

WPP runs according to a scenario that describes the pattern of water arrival to the system. The scenario defines the arrival time, the amount of water dispersed, and the destination tank. For example, an entry in the WPP scenario may indicate that at 2:02 PM 10 gallons of water will be assigned to tank 2. Using this scenario, any pattern of water arrival can be controlled, making it more or less difficult to perform the task. The scenario is unknown to the decision maker, therefore, the total number of gallons of water to deliver to each deadline is unknown. The decision maker s goal is to empty all the tanks, pumping the water detected in the system out towards the deadlines. The decision maker can keep track of his/her progress by regarding the indicator of the water gallons missed in the simulation positioned in the left upper corner of the screen. The decision maker s actions are interconnected and determine whether the water will be delivered on time. For example, once the decision maker decided to activate pumps corresponding to tank 4, water will accumulate in tank 16 and activation of its pumps will be necessary in the future. These sequential and interrelated decisions may be affected by water arriving to the system at unpredicted times and destinations according to the scenario. The decision maker needs to effectively use the pumps and be aware of incoming water in order to deliver the water on time.

Water flows from left to right (from the root towards the deadlines). For example, water in tank 1 is distributed to tanks 3 to 8. Each chain of tanks has a chain value. For example, tank 22 has a chain value of 0 (no tanks ahead), while tank 4 has a chain value of 2. Tank 0 has chain value of 4, tank 1- 3, and tank 3- 2. The higher the value of a chain, the longer it takes to deliver the water.

WPP and its successful solution

In WPP it is theoretically possible to reach the best performance of zero gallons missed under the employed scenario. The *standard* scenario is a scenario used as the baseline for this study with 1008-gallons of water distributed within the 23 tanks. Assuming optimum utilization of all 5 available pumps during the simulation, the system s capacity is 2400 minutes (5 pumps X 8 hours X 60 minutes/hour). 1008 gallons at a rate of 1 gallon every two minutes can be distributed in 2016 minutes. This calculation per deadline assures efficient use of the system s capacity to solve this task successfully. Although it is possible to solve this task successfully, in practice it has proven to be a challenge. Of approximately 500 WPP simulation users only two have reached the best performance score of 0 after 18 trials.[1]

Verbal protocols have also been conducted of participants in this task. The protocols indicate that time, water volume, and chain value are important variables for making accurate decisions. The awareness of these variables, however, depends on the amount of practice users have on the task. At the beginning some participants report making choices randomly or based exclusively on the deadlines. With practice, participants become more aware of the relationship of other variables, such as volume of water and the chain value.

[1] The authors of this paper have also reached best performance, but of course we are aware of the scenario and the external arrivals of water.

As a reasonable upper-limit yardstick for performance, a program was created to make random activations while never allowing idle pumps, thus making the best use of the system s capacity. This strategy is called the zero intelligence scheduler. The results for 30 replications of random assignments produced a mean of 182.9 missed gallons with a standard deviation of 28.4, therefore, reasonable performance is between 0 and 200. Other strategies include the time heuristic, which activates pumps of tanks with closer deadlines first. The use of the time heuristic over the standard scenario gives a score of 58 gallons missed with no idle pumping time. The time-volume considers two features: deadlines and water amounts. The water amounts are transformed into time units. The use of the time-volume heuristic over the standard scenario gives a score of 103 gallons. Finally, the time-volume-chain heuristic considers the chain in addition to the deadlines and water volumes. The water amounts are transformed into time units and multiplied by the chain value. The use of this heuristic gives a score of 88 gallons missed. The results of applying these heuristics are much better than the zero intelligence scheduler but far from the optimal solution. These numbers also indicate that the simplest heuristic, time, works better than the other two more complex heuristics. The analysis of applying heuristics to the task indicates that the optimal solution in WPP is more than a consistent adherence to a particular heuristic. The dynamic nature of the task, environmental uncertainty, and constrained resources indicate that a heuristic might be good in some situations and bad in others.

Algorithmic description of the production rules for CogIBLT in the WPP task

Let variable **current_state** define the current state of the simulation
variable **tanks** define the set of all the tanks in the simulation
variable **situation_decision_chunks** are produced according to the missed buckets (SDUs)

Let deadline = 5:00PM
REPEAT WHILE deadline < 10:00PM
[
Function *goal_meet* (deadline)
{
Production 1. **[If** *reached* (current_deadline) **= True**
Then *update* (situation_decision_chunks)
Let new_deadline = *what_is_new_deadline* (current_state),
goal_meet (new deadline) //goal is updated
Else *goal_make* (decision)}**]**

Function *goal_make* (decision)
{
Let (tank_to_be_activated, tank_to_be_deactivated) = *evaluate_urgent*(tanks)
If *more_urgent* (tank_to_be_activated, tank_to_be_deactivated) = tank_to_be_deactivated
Then
Production 5. **[** *deactivate_pump*(tank_to_be_deactivated),
goal_meet(deadline) **]**
Else If *number_of* (tank_to_be_activated) = 0
Then
Production 3. **[Let** tank_to_activate = *select_to_activate* (current_state, tanks),
goal_evaluate (tank_to_activate) **]**
Else If pumps_in_use = 5
Then
Production 2. **[** tank_to_deactivate = *select_to_deactivate* (current_state, tanks),
goal_evaluate (tank_to_deactivate) **]**
Else
Production 4. **[** *activate_pump* (tank_to_be_activated),
goal_meet (deadline) **]**
}

Function *goal_ evaluate* (tank)
{
If *similarity_of_evaluations*(current_state, past_decisions) = similar
Then
Production 6. **[Let** new_evaluation = *ACT_R_blending_mechanism* (tank, past_decisions, current_state),
add_situation_decision_chunk(new_situation_decision_chunk, past_decisions),
goal_make (decision) **]**
Else
Production 7. **[** *use_rule*(tank),
Let new_evaluation = *use_rule* (tank, current_state)
add_situation_decision_chunk (new_situation_decision_chunk, past_decisions),
goal_make (decision) **]**
}

Exploring individual data

Two humans with different learning curves were selected from the 14 participants by visually determining a flat and a steep curve (HumanP1 and HumanP2). To obtain the simulated subjects that best fit the humans, performance r^2 were calculated for these two participants compared to all the 14 simulated subjects over the 18 trials. The two simulated participants with the highest r^2 were selected (Model163 and Model172).

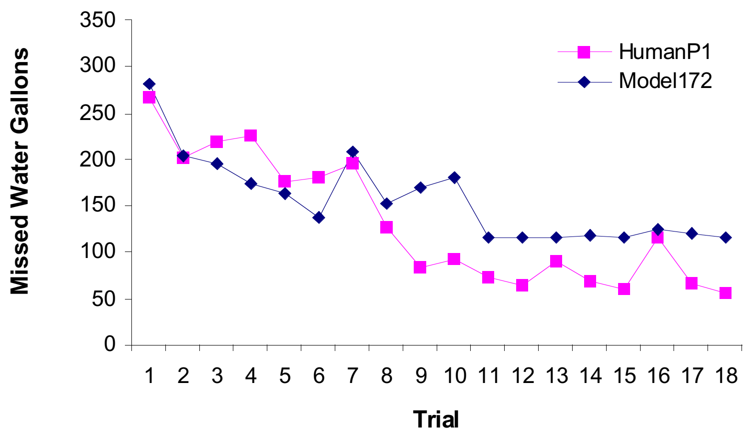

a.

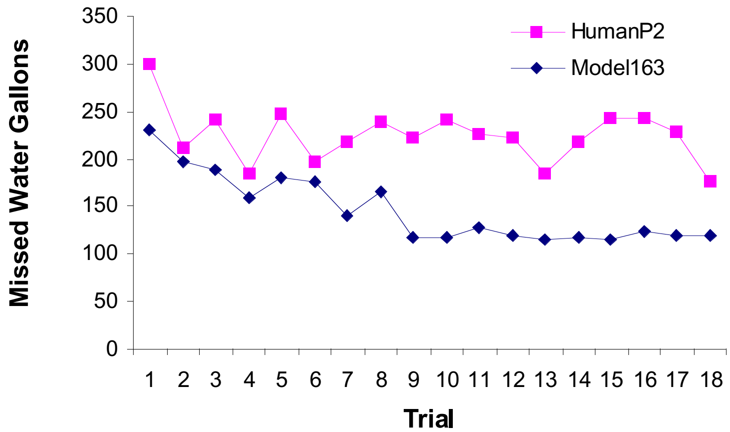

b.

Figure 20. The learning curves for two human participants compared to simulated participants. a. Shows the best fit to a fast learner, while b shows the best fit to a slow learner.

Fig. 20 shows human and model data for a fast and a slow learner. Figure 20a fits the model data with an r^2 value of .79 and Figure 20b fits the model with an r^2 of .27. In our previous research, we have found that performance differences are not caused by participants simply failing to monitor the environment, for example, by leaving the pumps unused, but rather by their working memory capacity (Gonzalez, Qudrat, & Lerch, manuscript submitted for publication). A measure of the user s activity during a trial is the number of made decisions. IBLT predicts that users reduce the number of made decisions overtime.

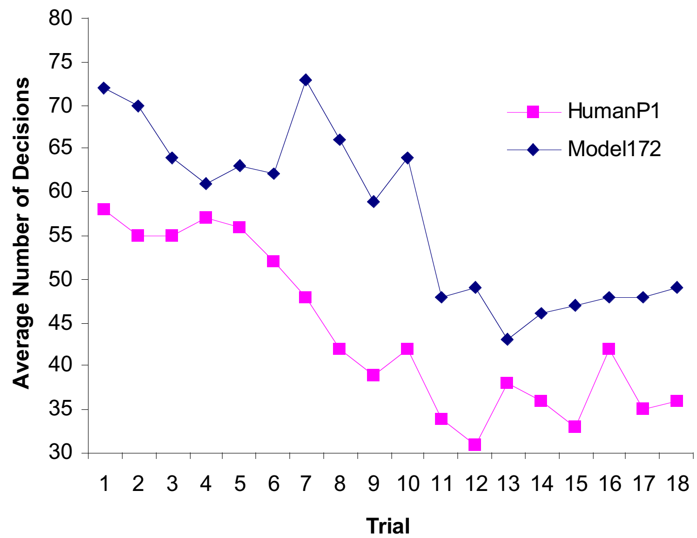

a.

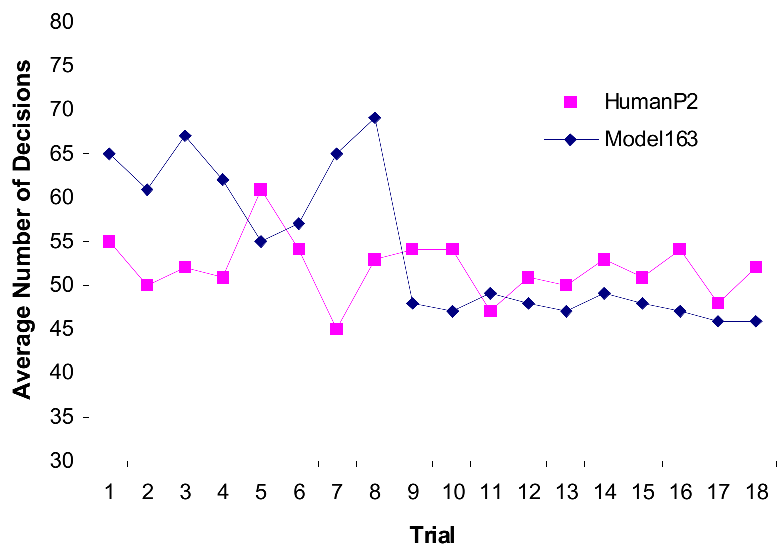

b.

Figure 21. The average number of decisions per trial for two human participants compared to simulated participants. a. Shows the best fit to a fast learner, while b shows the best fit to a slow learner.

Fig. 21 shows number of decisions per trial for these two individuals compared to the model s participants. Human data shows a bigger reduction in the number of decision for the fast learner, HumanP1, compared to the slow learner, HumanP2. This result suggests that the fast learner waited longer than the slow learner before taking action, confirming the IBLT predictions. CogIBLT captures this effect well for the fast learner with an $r^2 = .62$ and a similar total reduction in a number of decisions, but poorly for the slow learner with almost zero correlation, although the average number of decisions are comparable.

In sum, with standard ACT-R parameters and within the same experimental condition (Model 6-Low) CogIBLT can produce different learning curves. CogIBLT, however, fits the human fast learners better than slow learners in the Model 6-Low condition. This is not surprising, considering that Model 6-Low includes most of the learning propositions by IBLT. Thus the differences in learning rate between model runs may provide a better fit to individual subject differences.