Example-based models of control problems

Dieter Wallach
Department of Psychology
University of Basel
CH.4056 Basel
+41 61 267 3525
`wallachd@ubaclu.unibas.ch`

Christian Lebiere
Department of Psychology
Carnegie Mellon University
Pittsburgh, PA, 5000 Forbes Avenue
+1 412 268 2815
`cl+@cmu.edu`

ABSTRACT
The paper reports work in progress on the investigation of the new ACT-R *blending mechanism* in its scope to account for data from two experiments with Broadbent's TRANSPORTATION task. Two simple models are proposed that store and retrieve problem solving episodes to learn controlling the TRANSPORTATION scenario. Preliminary results with the models provide clear evidence for the empirical adequacy of the *blending mechanism*.

Keywords
Complex problem solving, implicit learning, blending mechanism, instance-based learning.

INTRODUCTION
In reviewing problem solving theories of system control, Funke (1995, p. 262) concluded that "the development of problem solving theories is in a rather desolate condition". Recently some theoretical progress has been made by using the ACT-R architecture as an integrative framework for modeling complex learning processes in dynamic system control (Lebiere, Wallach & Taatgen, 1998; Wallach & Lebiere, 1998). This paper extends our previous work by introducing example-based models for Broadbent's TRANSPORTATION scenario (Broadbent, 1997), a semi-classic computer-based task that has been applied in a large number of studies to investigate implicit learning processes in controlling dynamic systems. The proposed ACT-R model presents an application of the new *blending mechanism* that is used as the foundation of an example-based learning approach that seems to be especially powerful in continuous domains. The blending mechanism retrieves from declarative memory not a single chunk but the value(s) that best satisfy the constraints expressed by an entire set of chunks, with each chunk weighted by its probability of retrieval. This allows relations that are implicit in declarative memory to have a bearing upon the retrieval process without needing to be explicitly formulated.

THE TRANSPORTATION SCENARIO
In two experimental studies with 40 subjects each, participants were told to imagine that they are controlling the transport in a city. In the TRANSPORTATION scenario subjects can vary two input variables: *[t]* (time interval between buses entering the city) and *[f]* (amount charged for the use of the city's car parks). Altering these quantities affects two output variables: *[L]* (load on the buses) and *[S]* (number of empty spaces remaining in the car parks). Unbeknownst to the subjects the behavior of Transportation is governed by the following equations:

$$(1) \quad L = 220t + 80f$$

$$(2) \quad S = 4.5f - 2t$$

In the experiments subjects are asked to manipulate the input variables *t, f* to produce specific value pairs *[L, S]* of the output variables.

Previous research by Broadbent and his associates found no correlation between ability to control the system (as judged by the number of attempts to reach target values) and scores on a post-task questionnaire. While most subjects in Broadbent's experiments discovered the direct (or salient) influence of *t* on *L* and *f* on *S*, they were not able to verbalize knowledge on the less direct (or crossed) influence of *f* on *L* and *t* on *S*. However, Berry and Broadbent (1987, p. 9) assume that "Subjects must take the crossed relationships, as well as the direct ones, into account when controlling the system … successful performance cannot be based on the salient relationships alone". Since subjects seem to lack verbalizable knowledge of the non-salient relationships, Broadbent and Berry refer to an *implicit* learning mechanism to explain the performance of the subjects.

PROPOSED ACT-R MODELS
The proposed ACT-R approach challenges Broadbent's view by substantiating that example-based learning that only represents encountered input-output pairs (without explicitly encoding structural knowledge about causal relationships between variables) is clearly sufficient to successfully control the TRANSPORTATION task.

Two model variants were developed:

- *Model 1* encodes and retrieves separate pairs of input/output values of the saliently connected scenario variables. This representation implicitly assumes that subjects are aware of the *salient relationship* between $t \rightarrow L$ and $f \leftarrow S$. This separate representation can be justified by the results of a questionaire that was given to the subjects *prior* to the experiments. The data of experiment 1 shows that 37 out of 40 subjects (experiment 2: 35/40) assume a positive relationship between *f* and *S* and 33 out of 40 (experiment 2: 31/40) assume a positive relationship between *t* and *L* *before* they were exposed to the task. In post-task questionaires virtually all subjects were able to verbalize knowledge of the respective salient relationships.

Figures 1a and 1b show the central productions of the model.

```
(p fee-retrieval
   =goal>
      isa spaces-fee
      spaces =spaces
   =fact>
      isa spaces-fee
      spaces =spaces
      fee =fee
==>
   =goal>
      fee =fee
)
```

Figure 1a: Retrieval of *spaces*-instances

```
(p interval-retrieval
   =goal>
      isa load-interval
      load =load
   =fact>
      isa load-interval
      load =load
      interval =interval
==>
   =goal>
      load =interval
)
```

Figure 1b: Retrieval of *interval*-instances

- *Model 2* represents instance chunks as quadruples that store the settings of both input variables with their resulting output values. Figure 2 shows the central production of the model.

```
(p retrieve
   =goal>
      isa transportation
      spaces =spaces
      load =load
   =fact>
      isa transportation
      spaces =spaces
      load =load
      fee =fee
      interval =interval
==>
   =goal>
      load =load
      fee =fee
)
```

Figure 2: Retrieval of instances (*joint representation*)

For both models one architectural parameter (*activation noise*) was varied to fit the data.

EMPIRICAL EVALUATION

In an empirical evaluation the models were compared to the human data of the experiments that we conducted. In accounting for the data, we distinguish three classes of models:

- *Prototypical model:* A prototypical model is trained with a generic set of training instances (input-output pairs) for each model run without making use of actual subject data. 10 instances that exemplify a representative sample of the problem space of TRANSPORTATION are used to train model 1. Model 2 was not used as a prototypical model yet.
- *Individualized model:* An individualized model is trained with the input-output-pairs that the respective subject produced in the experiment when working on the first two of the eight problems[1] that were supposed to be solved in the experiments. The model is then run to predict the performance (in terms of trials until problem solution) for the final six problems. The noise parameter was *not* fit individually for each subject but was set globally for the whole population.
- *Individual models:* As with the individualized models, individual models are trained with the first two problem solving episodes of a respective subject. However, the noise parameter is fitted individually for each subject. The results of the individual models are not reported in this paper.

The evaluation of the models is thus not only carried out on an *aggregate* or *prototypical level*, but also involves detailed comparisons of learning trajectories on an *individualized level.* With respect to the levels of resolution for comparing model behavior to empirical data, Simon and Wallach (1999) have proposed criteria on different grain sizes. On a coarse level of comparison we use the *average total number* of trials for the solution of all 6 problems, on a finer level we compare the *number of trials on each problem* to empirically evaluate the model.

Average total number of trials

Figures 3 and 4 show the average total number of trials in experiment 1 and 2 in comparison to the performance of the models.

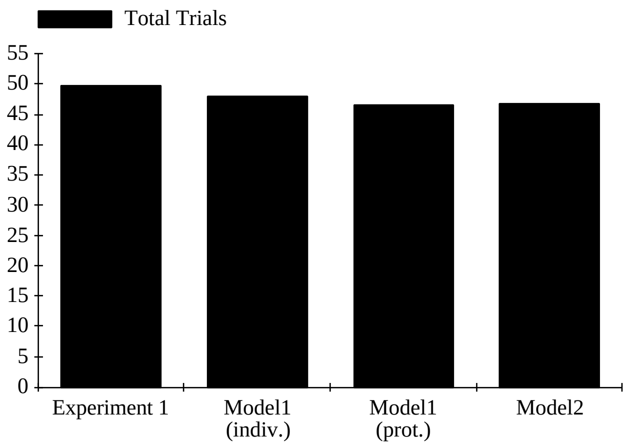

Figure 3: Average total number of trials (Experiment 1)

[1] The problems were the same that Broadbent used in his experiments.

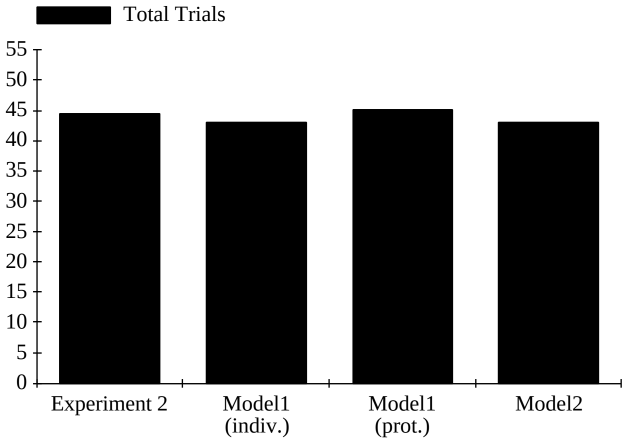

Figure 4: Average total number of trials (Experiment 2)

With regard to the *average total number of trials* for the solution of six problems, the proposed example-based ACT-R models all predict a performance that is clearly comparable to the data observed in experiment 1 and experiment 2. To evaluate the models on a finer grain-size, the next section reports the performance in terms of the number of trials on each problem.

Number of trials on each problem

Figure 5 shows the behavior of Model 1 (prototypical run, i.e. trained with 10 generic instances) on each of the six problems in the TRANSPORTATION task. The model was run with activation noise set at 0.25.

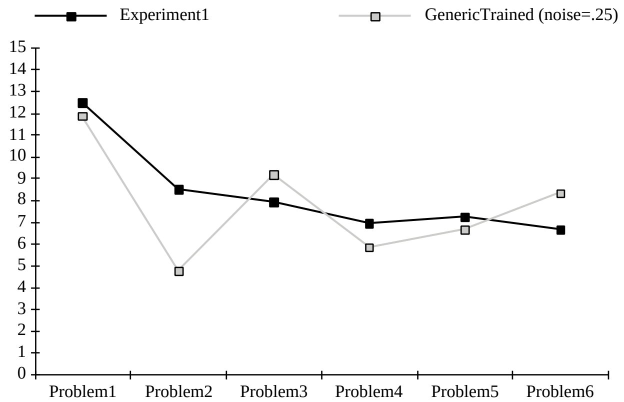

Figure 5: Model 1 (prototypical run), experiment 1

Figure 6 shows the performance of the individualized model[2] 1 on the six TRANSPORTATION problems. In comparison to the prototypical version, the individualized model clearly provides a better fit to the empirical data.

[2] The better performance of all models on problem 2 is a effect that is due to the specific output pair that was given as the target pair (see Lebiere & Wallach, in prep.).

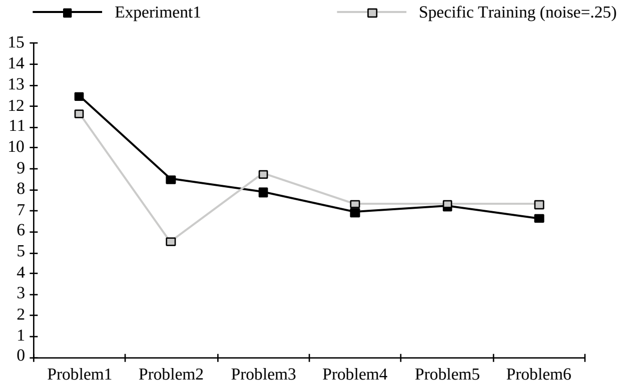

Figure 6: Model 1 (individualized run), experiment 1

As Figure 7 shows, the individualized variant of model 2 also matches the data more convincingly than the prototypical run of model 1.

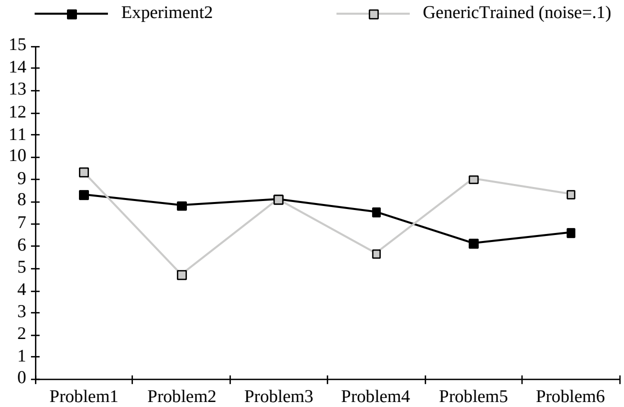

Figure 7: Model 2 (individualized run), experiment 1

When comparing the two model variants of model 1 on experiment 2 (Figure 8 and 9) we essentially observe the same picture: again, the individualized model results in a better fit to the empirical data. Note that the activation noise parameter was set to 0.1 for the results reported in Figures 8 and 9.

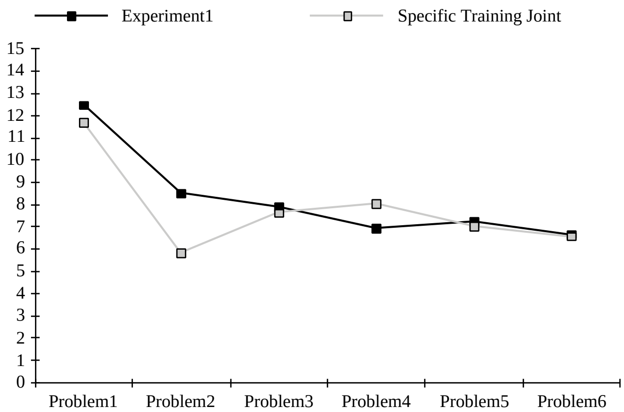

Figure 8: Model 1 (prototypical run, experiment 2)

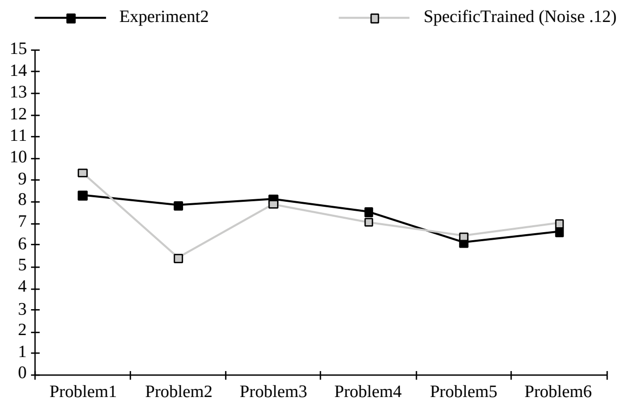

Figure 9: Model 1 (individualized run, experiment 2)

Individualized runs of model 2, however, do — in contrast to experiment 1 — not result in a better performance than the prototypical version of model 1. Further examination of model 2 revealed that it generally seems to need more training instances to produce higher performance levels (Lebiere & Wallach, in press).

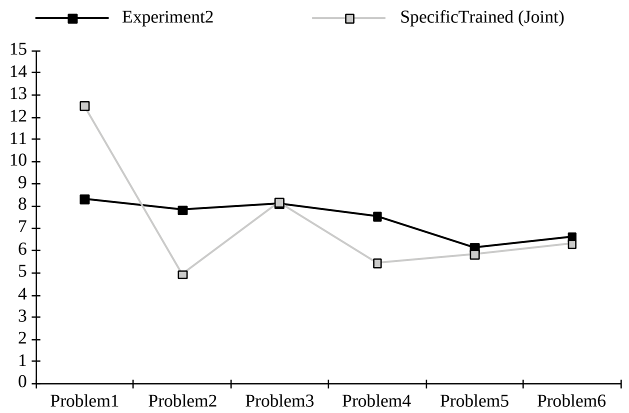

Figure 11: Separate retrieval of *interval*-instances

CONCLUSIONS

The paper reports work in progress on the scope of the ACT-R *blending mechanism* on the TRANSPORTATION control task. Despite the striking simplicity of the models proposed, both showed remarkably good fits to the empirical data that point to the empirical adequacy of *blending* as an architectural mechanism.

REFERENCES

Broadbent, D.E. (1977). Levels, hierarchies, and the locus of control. *Quarterly Journal of Experimental Psychology, 29, 181-201.*

Broadbent, D.E. & Berry, D. (1987). Interactive tasks and the implicit-distinction. *British Journal of Psychology, 79, 251-272.*

Funke, J. (1995). *Experimental research on complex problem solving.* In P. Frensch & J.Funke (Eds.). *Complex problem solving: The european perspective,* pp. 243-268. Hillsdale, NJ: Erlbaum.

Lebiere, C., Wallach, D. & Taatgen N. (1998). Implicit and explicit learning in ACT-R. In F. E. Ritter & R. Young (Eds.). *Proceedings of the 2nd European conference on cognitive modeling,* pp. 183-192, Nottingham: Nottingham University Press.

Lebiere, C. & Wallach, D. (in prep.). Example-based models of control problems. *Manuscript in preparation.*

Wallach, D. & Lebiere, C. (1998). Modellierung von Wissenserwerbsprozessen bei der Systemregelung. In U. Kottlamp & W. Krause (Eds.). *Intelligente Informationsverarbeitung,* pp. 93-100. Wiesbaden: DUV.

Wallach, D. & Simon, H.A. (1999). Cognitive modeling in perspective. *Kognitionswissenschaft, 8, 1-5.*