Title

Is Feedforward Learning more Efficient than Feedback Learning in Smart Meters of Electricity Consumption?

Authors

Guath, Mona
Juslin, Peter
Millroth, Philip

Is Feedforward Learning more Efficient than Feedback Learning in Smart Meters of Electricity Consumption?

Mona Guath (mona.guath@psyk.uu.se)
Department of Psychology, Uppsala University
SE-751 42, Uppsala, Sweden

Peter Juslin (peter.juslin@psyk.uu.se)
Department of Psychology, Uppsala University
SE-751 42, Uppsala, Sweden

Philip Millroth (philip.millroth@psyk.uu.se)
Department of Psychology, Uppsala University
SE-751 42, Uppsala, Sweden

Abstract

The most popular way to improve consumers' control over their electricity cost is by providing frequent and detailed feedback with "in-home displays" (IHD). In this study, we examined alternative ways to train experimental participants to control and optimize their use of electricity by "feedforward" training to map energy consuming behaviors to costs. The participants were trained in one of four experimental conditions, one feedback ("IHD") and three feedforward conditions before they had to control the electricity consumption in a simulated household. Results showed that one of the feedforward conditions produced somewhat higher utility and as good or better satisfaction of a monthly budget than the feedback training condition, despite never receiving *any* feedback about the monthly cost, but the generalization to a new budget constraint proved to be slightly poorer.

Introduction

The use of so-called "smart electricity meters" is rapidly becoming common. It has been estimated that within the European Union alone some 51 billion euro is being invested in smart meters (Faruqui, Harris, & Hledik, 2009). In many countries, household energy consumption is still billed once a month, but smart meters can offer feedback that is detailed and more frequent with so called In Home Displays (IHDs). Intuitively, the latter kind of feedback system seems more beneficial, and, indeed, many early studies suggested energy reductions up to 15%. However, more recent studies point at consumption reductions at 2-4%, few of them being significant (Klopfert & Wallenborn, 2011). In the present study, we focus at in-home displays (IHDs), which only display the electrical consumption at different time intervals, and, unlike smart meters, they do not have a two-way communication with the central system. In a previous laboratory experiment (Guath, Millroth, Elwin, & Juslin, 2012), we investigated how feedback about electricity consumption is best presented to electricity consumers in order to control and optimize their use of electricity. To measure a participant's energy efficiency in an experimentally controlled environment, the participants took on the role of an inhabitant in a simulated household, performing different types of energy consuming behaviors within a given budget (Figure 1). The goal of decreasing electricity consumption is often emphasized, but the participant's task is actually an optimization problem that requires an appropriate balance between the cost of the electricity consumed and the benefit or *utility* obtained. The problem is illustrated in Figure 2, where the utility of electricity consumption is plotted against cost. The maximum utility obtainable at a given cost, assumed to be a decelerating function of the cost, is illustrated by the curve in Figure 2. The hypothetical utility obtained at a cost by a consumer is illustrated with a dot. The task is to move closer to the line for "maximal utility", however, this is associated with two constraints: achieving sufficient utility to make life bearable and not surpassing a constrained budget.

Guath et al. (2012) showed that in a deterministic system, frequent and detailed feedback was advantageous, but in probabilistic system, feedback aggregated over time was better, because it filtered out random noise.

The Present Study

In the present study, we wanted to evaluate if the same improvement could be obtained by feedforward training, rather than feedback training (as in most IHDs), hence, minimizing the negative effects from feedback interventions as conceptualized in Kluger and DeNisi's (1996) study. Specifically, we wanted to avoid the decrease of effectiveness when attention is moved away from the task to the self, thus, making the effects of the training short-term. Another motive was to make the mapping task more flexible, not being dependent on the simulated household (Figure 1). Detailed and frequent feedback (an IHD) was compared to three feedforward conditions. Feedforward is defined as a process where knowledge is used to act directly to control the system, thus anticipating the changes that will occur (Basso & Olivetti Belardinelli, 2006). In the present task, partici-

1

pants had to control the monthly cost of electrical consumption. Feedback training involved feedback about this criterion variable of daily and monthly cost of electrical consumption from experience with the task (running the simulated household in Figure 1). Feedforward training involved no feedback about the criterion variable (monthly electricity cost), but three different training schemes in various ways teaching the participants to directly map energy consuming behaviors to their costs ("map" refers to the mathematical concept of associating each element in a set with an element of another set, here the electrical cost to a certain electrical post).
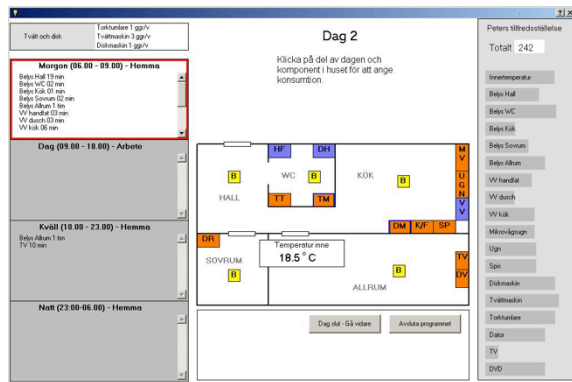

Figure 1. The computer display in the simulated household in the experiment.
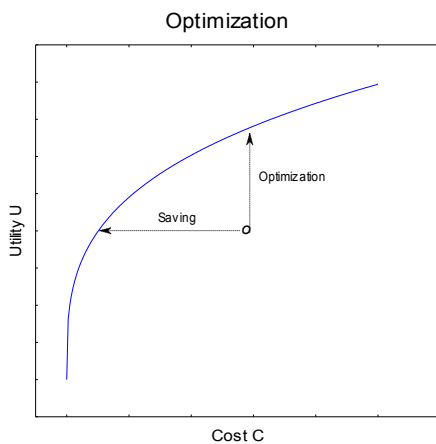

Figure 2. Illustration of the two ways of obtaining maximum utility at a specific cost (the solid nonlinear function), either by saving or by optimizing the use.

The choice of feedforward conditions was, in part, inspired by Pachur and Olsson's (2012) study of how learning tasks affect performance and strategy selection. They investigated two learning tasks, *direct criterion learning* and *learning by comparison*, and how these affected performance depending on the type of test (paired-comparison, classification, estimation) and decision environment (linear vs. non-linear). Pachur and Olsson (2012) concluded that direct criterion learning invites exemplar memory processes (Nosofsky, 1986), while learning by comparison invites processes of cue abstraction. Because our task is non-linear, if anything, exemplar memory should be a more efficient process than the abstract processes involved in cue abstraction,

which are constrained to mainly capture linear and additive tasks (Juslin, Karlsson, & Olsson, 2008).

Our first feedforward condition, *metric mapping* (corresponding to direct criterion learning), informed about the function that relates the consumption to its cost, as studied in research on function learning (e.g., Kalish, Lewandowsky & Kruschke, 2004). When training function concepts, a continuous stimulus variable is associated with a continuous response variable, in this case an electricity post (e.g., inner temperature) with its monthly cost. The metric mapping consisted of learning to map a certain electricity consuming activity (e.g., using hot water 15 min/day) to its cost (i.e., 262 SEK/month).

The second condition was *rank-order mapping* (corresponding to learning by comparison) as conceptualized in decision by sampling (DbS) (Stewart et al., 2006). In DbS, it is assumed that people do not store metric knowledge in memory but only perform ordinal comparisons. Instead frequency accumulation in pair-wise comparisons are used for evaluating a target attribute against a decision sample. Indeed, the results in Pachur and Olsson (2012) suggested that at least in linear tasks learning by pairwise comparisons was more efficient than training with metric mapping, despite that the pairwise comparisons provide no explicit metric information about the criterion. On the other hand, if people also need to store metric knowledge in our task, then metric mapping should be more efficient. The rank-order evaluations are elicited by questions concerning the relation between two electricity consuming device (e.g., Which of the following has the highest monthly cost: A: Having the lights on for 60 minutes per day or B: Having the computer on for 10 minutes per day?).

The third feedforward condition was *causal mapping training*, in which the participant is encouraged to experiment with the individual and total monthly cost of the electrical posts in a minimalistic computer program. The causal mapping condition is inspired by the theory of causal nets (Holyoak & Cheng, 2011) that accounts for how people learn about strength and structure as well as direction of causal relations. In view of this literature, we expected that invitation to manipulate the system in real time and experiment by changing individual variables during training should produce a more accurate (causal) model of relationships in the system.

Given that our decision task is non-linear, where the linear and additive integration afforded by cue abstraction is less appropriate, and the results suggesting that metric mapping invites exemplar memory (Pachur & Olsson, 2012), performance in the metric mapping condition is expected to be better than in the rank-order and the causal mapping conditions. We also looked at the ability to generalize knowledge to a new budget (from 2000 SEK á month to 1500 SEK á month).

Method

Participants

One-hundred-and-twenty-nine students at Uppsala University volunteered to participate and were compensated with a cinema ticket (worth approximately $10) or by

course credit. The sample consisted of 89 females and 40 males, with mean age 24.5 years (SD=4.66).

Material and Procedure

The experiment consisted of four parts presented in the following order: pre-test, systematic learning, post-test for effects of the systematic learning and post-test for ability to generalize to a new budget. Participants were given written and verbal instructions for each part.

The participant was presented with a sketch of a home on the computer screen indicating various energy consuming appliances, in all 18 posts (Figure 1). The task was to regulate the electricity consumption in the house so that its fictive inhabitant (an avatar called "Peter") received as much utility as possible from the consumption. The avatar was a way to incorporate the two parameters of utility and cost in the task in a comprehensible way. In Pretest and Posttest 1 of all conditions, the participant regulated the energy consumption each day for a period of 30 days with the aim to maximize the utility from the energy consumption within a budget of 2000 Swedish Crowns (SEK), approximately $300, per month (changed to 1500 SEK in Posttest 2). On each new day, the participant could adjust the indoor temperature, the number of times of use per week for the dishwasher, washing machine, and tumble drier etc. When the settings had been made, they could not be changed for that day. On each day, the previous day´s settings were presented as default, but they could be changed by the participant.

The utility of consumption for each appliance was presented by a bar on the right side of the screen. A separate bar for each post increased with the utility of consumption associated with this post, and a global sum in the upper right corner increased with the overall utility of the energy consumption. The utility $u_i(t_{ij})$ obtained by consumption t_{ij} of post i (i=1…18) at level j was,

$$u_i(t_{ij}) = \sum_{i=1}^{18} w_i \cdot (t_{ij}^{\alpha_i} / r_i^{\alpha_i}),$$ (1)

where w_i is the linear weight in the overall summed utility (Σw_i=1), r_i is a ceiling on the allowable consumption, and α_i is a parameter for the curvature of the utility function for post i. Eq. 1 defines utility functions with diminishing marginal return, where the posts differ both in the rate of the diminishing return (α_i) and in their weight in the total utility (w_i). The parameters were selected to approximate realistic utility functions.

The total utility U was the sum of the utility of each of the 18 posts,

$$U = \sum_{i=1}^{18} u_i(t_{ij}).$$ (2)

In pretests and posttests, feedback was presented *only* for utility. Participants were given no feedback regarding the cost of their settings. After the pretest, the participants were assigned to one of four learning conditions.

Detailed and Frequent Feedback ("IHD")

In the feedback condition, participants continued with 120 days in the simulated household. Furthermore, and most importantly, they also received feedback about the cost of their consumption. After each day the participants received a bill containing feedback on the cost of energy consumption where the feedback was presented in terms of used kWh and the cost in SEK, as based on a fixed price of 1.40 SEK per kWh. The bill showed the cost for each appliance as well as the total sum for all appliances. This detailed and immediate feedback resembles that of an IHD. If the budget was exceeded, the total cost was red-lighted; if not, it was green-lighted. A normally and independently distributed random error, with standard deviation equal to 5 % of the cost, was added to the cost of each post to simulate probabilism[1]. The total cost C was the sum of the consumption cost $c(t_{ij})$ of the individual posts:

$$C = \sum_{i=1}^{18} c(t_{ij}).$$ (3)

Metric mapping

For participants in the metric mapping condition, the task was to learn to map consumption of each electricity post directly to its cost. They were presented with questions such as: "What is the monthly cost for having the computer on for 15 minutes a day?" Intervals of 5, 10, 15, 30 and 60 minutes were used to give the participants a broad spectrum of the cost for each appliance. Participants reported their responses for one question at a time and were then given feedback on whether the response was correct or not and, if not, what the correct answer was. The program coded answers within ±20 percent of the correct answer as correct. A stop criterion was set for three correct responses in one block of five questions (one block involved the cost for 5, 10, 15, 30 and 60 minutes of use). When the participants achieved the stop criterion for one appliance they continued with the next appliance, until they had gone through all appliances in the house. Electric posts for lighting and hot water (shower and tap water) was lumped together, creating 12 different appliances from the 18 electricity posts in the simulated household. For appliances that were run on a weekly basis, such as the dish washing machine, participants were asked about the cost for number of runs per week.

Rank-order Mapping

For participants in the rank-order mapping condition the task was to learn which of pairs of electricity consuming activities that is most costly. They answered questions such as "What has the highest monthly cost? A: Having the lights on for 60 minutes per day or B: Having the computer on for 10 minutes per day?" After each guess they were provided with feedback on whether the response was correct or not. The items were sampled from a pool of questions created by crossing the 12 applianc-

---

[1] The probabilism is intended to capture all factors that contribute to the imperfect measurement value of a specific post at a randomly chosen time. This includes both limits in the precision of measurement as such and exogenous factors that affect the cost but are unknown to the consumers.

es (as in the direct-mapping condition) by the five time intervals (5, 10, 15, 30, 60 minutes, except for appliances run on a weekly basis). The items were sampled randomly with one constraint: each appliance had to appear at least once during the training session. Each participant received a unique sample. The participants started with 200 items; thereafter, they continued until they answered 19 of the latest 20 questions correctly.

Causal Mapping
For participants in the causal mapping condition the task was to learn the relationship between consumption and electricity cost by interacting with the appliances on a real time basis, in order to obtain a sense of the cause and effect relationships. Again, the participants trained on 12 appliances with the same time intervals as in the other mapping conditions (5, 10, 15, 30, 60 minutes, except for appliances run on a weekly basis). Participants were presented with a program where they could manipulate the usage of appliances on slide bars. Each slide bar had five levels for the time intervals of 5, 10, 15, 30 and 60 minutes. Next to each slide bar, the cost of the appliance was indicated. The cost changed simultaneously with the manipulation of the slide bar. Participants could observe the monthly cost for their current settings on the top of the screen, also changing simultaneously with the manipulation of a slide bar. All appliances were presented on the same screen. The order in which the appliances were presented on the screen was randomized for each participant. A time limit of 15 minutes was set, and the participants were told to experiment and learn as much as possible in this time, with the goal to optimize their behavior in the household.

Posttests
After participants had finished their respective training session, they all continued with another 30 days in the simulated household, similar to the 30 days in the pretest. After that first post-test, the participants continued with another round of 30 days in the simulated household, but this time the budget constraint was set to 1500 SEK instead of 2000 SEK, with the goal of investigating participants' ability to generalize the knowledge they had acquired in the systematic learning conditions.

Design
The experiment involved a 4x2 mixed factorial design, with learning condition (detailed and frequent feedback, metric mapping, rank-order mapping, and causal-model mapping) as between-subjects independent variable, and budget constraint (2000 and 1500 SEK) as the within-subjects independent variable. The participants were randomized to one of the between-subjects conditions, resulting in app. 30 participants in each condition. Dependent measures were the cost and utility of the use of electricity, with a particular eye to the maximization of utility within the indicated budget constraints of 2000 SEK and 1500 SEK per month.

Results
In pretest, there were no significant differences between the conditions and all conditions exceeded the budget of

67 SEK/day. All conditions reduced their median cost from the Pretest to Posttest 1 (Wilcoxon Test, T=1135, Z=7.187, $p < .001$ across all four conditions; the same holds separately within each condition, all ps < .005).

In Figure 3, the median utility is plotted against the median cost for Posttest 1. The rank order condition was unable to satisfy the budget. The three conditions satisfying budget performed similarly, although metric mapping produced somewhat more utility than the other two conditions, which both fell below budget, as observed previously with feedback training (Guath et al, 2012).
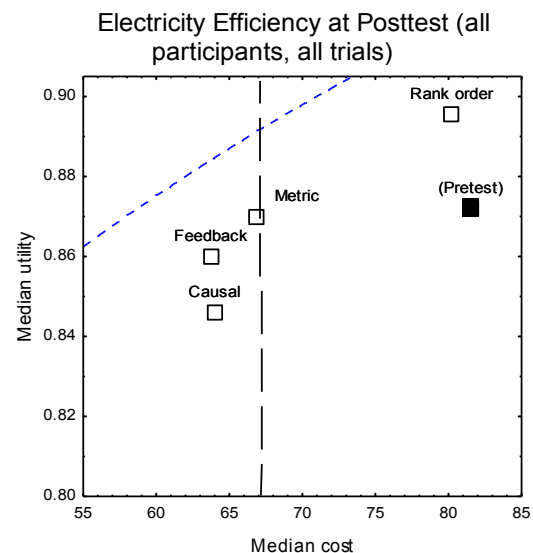

Electricity Efficiency at Posttest (all participants, all trials)

Figure 3. The median utility obtained across the participants plotted as a function of the median cost in Pretest and Posttest 1 in each of the four training conditions (N≈30). The vertical line represents the budget, while the curve is the maximum utility obtainable as a function of the cost.

There was a significant difference between the conditions in the median cost at Posttest 1 (Kruskal-Wallis test: H (3, N= 129) =18.607 p =.003), where pairwise multiple-comparisons indicate a significant difference between rank-order mapping and feedback training (p<.001) and between rank-order and causal mapping (p = .005)[2]. In both cases, rank order mapping has a significantly higher cost. There was also a significant difference between the conditions in median utility at Posttest 1 (Kruskal-Wallis test: H (3, N= 129) =12.557 p =.006), where pairwise multiple-comparisons indicate a significant difference between rank-order mapping and feedback training (p=.019) and between rank-order mapping and causal mapping (p = .009). In both cases, rank-order mapping has a significantly higher utility. The main differences is between the rank-order condition and the other conditions, with small differences between the other conditions, albeit with a slight hedge for metric mapping that comes closest to optimal performance (i.e., where the lines in Figure 3 intersect).

Figure 4 reports a more strict analysis only including those participants that roughly satisfied the budget (i.e.,

---

[2] In the multiple-comparisons, we report raw p-values assuming α=.05. The Bonferroni corrected α-level is app. .008. The same is true for the multiple-comparisons reported below.

fell within ± 5 units of 67 SEK/day). While app. 50 % of the participants were able to satisfy the budget with feedback training and metric mapping only a minority of participants were able to satisfy the budget with causal mapping and rank order mapping (27% and 12%, respectively). Among the only two conditions with many participants satisfying the budget, metric mapping produced significantly more utility than feedback training (Mann Whitney: U=75, Z=2.179, p=.029). Thus, if anything, metric mapping appears to produce somewhat better performance than feedback training in Posttest 1.
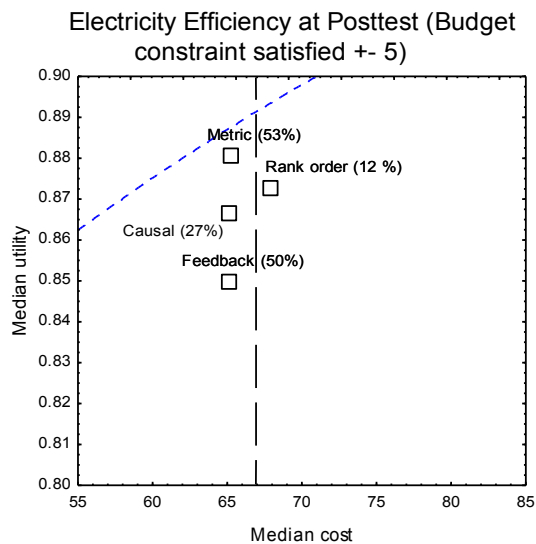

Figure 4. The median utility across the participants plotted as a function of the median cost in Posttest 1 in each of the four training conditions for the participants able to (app.) satisfy the budget (cost between 62 and 72 SEK/day). The vertical line represents the budget, while the curve is the maximum utility obtainable as a function of the cost. The percentages in parenthesis refer to the proportion of participants in each condition that satisfied the budget constraint.

All training conditions reduced their cost from Posttest 1 to Posttest 2 (new budget) (Wilcoxon test: all ps < .001), and reduced their utility (Wilcoxon test: all ps < .001). As shown in Figure 5, in all conditions the median cost exceeded the budget, especially in the rank order condition. There is again modest difference between the other three conditions, although metric mapping exceeds the budget more than the other conditions.

There was a significant difference between the conditions in median cost at Posttest 2 (Kruskal-Wallis test: H (3, N= 129) =17.606 p <.001), where pairwise multiple-comparisons indicate significant differences between rank-order mapping and feedback training (p< .001), between rank-order and causal mapping (p = .033), and between feedback training and metric mapping (p=.045). There was also a significant difference between the conditions in median utility at Posttest 2 (Kruskal-Wallis test: H (3, N= 129) =11.800 p =.008), where pairwise multiple-comparisons indicate a significant difference only between rank-order mapping and feedback training (p=.011). The differences between rank order and metric mapping (p=.061) and between rank order and causal mapping (p=.050), however, ap-

proach significance. Rank order mapping produced a higher utility (but exceeds the budget). The rank-order condition thus provides the poorest performance, while, among the other three conditions, metric mapping seems to suffer most going from Posttest 1 to Posttest 2.
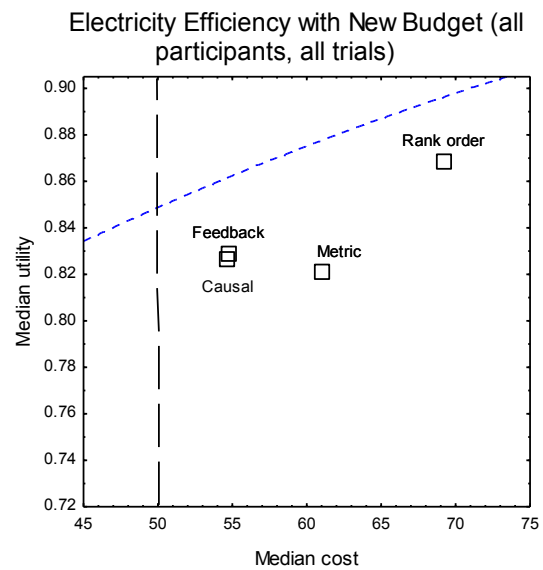

Figure 5. The median utility across participants plotted as a function of the median cost in Posttest 2 (new budget) in each of the four training conditions (N≈30). The vertical line represents the budget, while the curve is the maximum utility obtainable as a function of the cost.
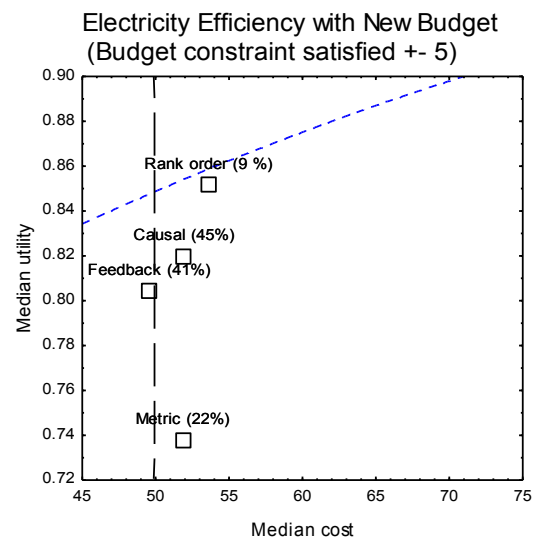

Figure 6. The median utility across the participants plotted as a function of the median cost in Posttest 2 (new budget), in each of the four training conditions for participants that were app. able to satisfy the budget (cost between 45 and 55 SEK a day). The vertical line represents the budget, while the curve is the maximum utility obtainable as a function of the cost. The percentages in parenthesis refer to the proportion of participants in each condition that satisfied the budget constraint.

Figure 6 shows the results of a stricter analysis only including participants with a cost falling within ±5 units of the budget (50 SEK/day). This figure also illustrates that metric mapping training suffered more in the generalization test, with only 22% of the participants satisfy-

ing the new budget, in contrast to 41% with feedback and 45% with causal mapping training. In sum: the rank order condition again produced poor performance, while feedback training and causal mapping appear to allow better generalization of the knowledge obtained to satisfy also a new budget, as compared to metric mapping.

Discussion

The results indicate that at an immediate test of performance direct mapping training is as good as or better than the detailed and frequent feedback in a typical "smart-meter". It should be noted that the performance of the participants in the direct mapping condition is quite impressive, with a cost virtually exactly on the budget and higher utility, despite *never receiving any feedback about the total monthly cost in the house*. An objection, of course, could be that these participants did not really learn anything from the metric mapping training, they just happened to be right because their prior conceptions about electricity consumption happened to be correct in regard to the simulated household (which is intended to be "realistic"). That all training groups changed their behavior significantly from the pretest to the posttest to accommodate the budget speaks against this explanation. Participants clearly learned to satisfy the monthly budget from metric mapping training, despite that they never received any feedback about it.

The good performance with metric mapping is in line with the results in Pachur and Olsson (2012), where the participants with direct criterion learning performed better than learning by comparison in a non-linear context due to the exemplar strategy. In the context of their interpretation, this suggests that our participants relied on exemplar memory rather than cue abstraction. An obvious question, in that case, is how people generate exemplar representations of the complex stimuli used in our experiment. On plausible possibility, perhaps, is that they rely on "exemplars" in the form of partial configurations of electricity consumption that were associated with very successful (or unsuccessful) performance.

When generalizing to another budget, the feedback group performs somewhat better than the metric mapping group. This result was unexpected and we can only speculate as to what explains this difference. One possibility is that participants with metric mapping relied on a more exemplar-based strategy, which is known to offer less flexible generalization than the more analytic knowledge of cue-criterion relations. Another possible explanation could be the testing effect (Roediger & Karpicke, 2006). Tests enhance retention more than additional study of the material, even when tests are given without feedback. In that perspective, the feedback training could be seen as a first test, and the post tests as yet another tests. Another interpretation could be that feedback training and causal mapping provided the participants with better knowledge of the underlying causes, that in turn, facilitated the generalization task.

In this task we found no evidence that people are especially apt at learning rank orders from pairwise comparison (Stewart et al., 2006), considering that the rank order condition allowed few participants to satisfy the budget. In Posttest 1, only 27% of the participants satisfied the budget criterion (and only 9 % in Posttest 2). It might be the case that metric information is more crucial in the cost-benefit optimization task in our experiment.

To further investigate feedforward training, future experiments will explore the flexibility with which metric mapping and feedback training can adapt to new budgets and compare mapping involving different metrics (e.g., metrics instead referring to negative environmental effects). Also, it would be interesting to investigate the testing effect, and how it pertains to this context.

The results reported in this study opens for the possibility that shorter and cheaper feedforward training, for example, involving a 15-minute session with a computer program, can be a cost effective alternative to the large scale implementation of complex information technology to monitor consumption and cost in real time.

References

Basso, D. & Olivetti Belardinelli, M. (2006). The role of the feedforward paradigm in cognitive psychology. *Cognitive Processes, 7*, 73-88.

Faruqui, A., Sergici, S., & Sharif, A. (2010). The impact of informational feedback on energy consumption - survey of the experimental evidence. *Energy, 35*, 1598-1608.

Guath; M., Millroth, P., Elwin, E., & Juslin, P. (2012). *Re-Evaluating Smart Meters with Detailed and Frequent Feedback of Electricity Consumption.* Manuscript submitted for publication. Department of Psychology, Uppsala University, Sweden

Holyoak, K. J., & Cheng, P. W. (2011). Causal learning and inference as a rational process: The new synthesis. *Annual Review of Psychology*, *62*, 135-163.

Juslin, P., Karlsson, L., & Olsson, H. (2008). Information integration in multiple-cue judgment: A division-of-labor hypothesis. *Cognition, 10*6, 259-298.

Kalish, M. L., Lewandowsky, S., & Kruschke, J. K. (2004). Population of linear experts: Knowledge partitioning and function learning. *Psychological Review, 111,* 1072-1099.

Klopfert, F. & Wallenborn, G. (2011). Empowering Consumers Through Smart Metering. Report for BEUC, The European Consumer Organization.

Kluger, A. & DeNisi, A. (1996). The Effects of Feedback Interventions on Performance: A Historical Review, a Meta-Analysis, and a Preliminary Feedback Intervention Theory. *Psychological Bulletin, 2,* 254-284.

Nosofsky, R. M. (1986). Attention, Similarity, and the Identi-fication-Categorization Relationship. *Journal of Experimental Psychology: General*, *115*, 39-57.

Pachur, T. & Olsson, H. (2012). Type of learning task impacts performance and strategy selection in decision making, *Cognitive Psychology, 65,* 207-240.

Roediger, H.L., III and J.D. Karpicke (2006). The power of testing memory: basic research and implications for educational practice. *Perspectives on Psychological Science* 1(3): 181-210.

Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology, 53,* 1-26.