

# Knowledge-based Modeling in Dynamic Decision Making

**Angel Iglesias (angel.iglesias@csic.es)**

Bioengineering Group, Spanish National Research Council (CSIC)  
Arganda del Rey, 28500-Madrid, Spain

**M. Dolores Del Castillo (lola@iai.csic.es)**

Bioengineering Group, Spanish National Research Council (CSIC)  
Arganda del Rey, 28500-Madrid, Spain

**J. Ignacio Serrano (jignacio.serrano@csic.es)**

Bioengineering Group, Spanish National Research Council (CSIC)  
Arganda del Rey, 28500-Madrid, Spain

**Jesus Oliva (jesus.oliva@csic.es)**

Bioengineering Group, Spanish National Research Council (CSIC)  
Arganda del Rey, 28500-Madrid, Spain

## Abstract

A knowledge-based model that emulates human behavior in a Dynamic Decision Making task is proposed. The model, MAIDEN-DSF, uses a connectionist representation of knowledge and a value function to compute the best alternative. In order to validate MAIDEN-DSF, two data sets have been used: a training set and a test set that contain the behavior of participants that performed the task with different conditions. The results suggest that MAIDEN-DSF is a considerable framework in order to model human behavior. The aim of this paper is to use MAIDEN-DSF to prove that participants do not perceive delay conditions when dealing with Dynamic Decision Making tasks.

**Keywords:** Decision Making; Computational Models; Connectionism.

## Introduction

Dynamic Decision Making lies in tasks that require a series of decisions where the state of the world changes, both autonomously and as a consequence of the decisions made (Brehmer, 1992). The Dynamic Stocks and Flows (DSF) task (Gonzalez & Dutt, 2011), emulates such situations with three elements: a single stock represented by a water tank; inflows, which increase the level of the stock; and outflows, which decrease the level of the stock. The goal of this task is to keep the stock at a certain level over 100 time periods. In every time period, a participant can control the stock adding or removing water via two inputs that represent the decision: the user inflow ( $UI$ ) and the user outflow ( $UO$ ) values. Besides, there are two external inputs that represent the environment inflow ( $EI$ ) and outflow ( $EO$ ) values, both of them are not directly controlled by the participant. The dynamics of the environment is unknown to the participant, so the  $EI$  and  $EO$  values have to be predicted using the experience acquired in previous time periods. When the participant makes a decision, the DSF determines the level of water in the tank by adding the  $UI$  and  $EI$  to the level in the tank and subtracting the  $UO$  and the  $EO$ . Then the DSF presents the resulting level, the goal level and the last  $EI$ ,  $EO$ ,  $UI$  and  $UO$  values within the following time period.

This seemingly simple stock problem is actually unintuitive and difficult. In fact, there is evidence that suggests that people do not perceive stocks and flows dynamics correctly. For example, in an experiment where graduate students were asked to sketch the evolution of the water level in a bathtub over time (given simple patterns for the environmental inflow and outflow), only 36% of the students answered correctly (Sterman, 2002).

The DSF has been studied under different conditions (Cronin & Gonzalez, 2007) and there is evidence that the hardest condition for controlling the water tank is when the user inputs are delayed for three time periods (Lebiere, Gonzalez, & Warwick, 2010). Participants have some understanding about the dynamics of the tank, for example, if 1 gallon of water is added, then the level is increased 1 gallon. However, with this delay, participants seem to have incorrect beliefs about the relationship between stocks and delayed flows or participants do not perceive the delay and, therefore, they behave as if there was no delay.

The aim of this paper is to use a knowledge-based model that emulates human behavior in the DSF in order to find cues that prove that participants do not actually perceive the delay. The proposed model is based on MAIDEN, a Model of Assessment and Inference of DEcisions based on a Net of concepts (Iglesias, Del Castillo, Serrano, & Oliva, 2010).

## A knowledge-based model

MAIDEN-DSF, the implementation of MAIDEN for dealing with DSF, is divided in two main phases. The first phase lies in the estimation of  $EI$  and  $EO$  values using a connectionist representation of knowledge called decision net. In the second phase, MAIDEN-DSF uses a value function to choose the best alternative, which will be the one that sets the amount in the tank at the goal level corresponding with the estimated  $EI$  and  $EO$  values of the first phase.

## Decision net

MAIDEN-DSF is a knowledge-based model that uses a weighted net of concepts called decision net to predict the values of  $EI$  and  $EO$ . The concepts of the decision net are arranged in five layers: perception, short-term memory, working memory, deliberative and output layer. If the concepts of the decision net change, then the behavior of MAIDEN-DSF also changes. Therefore, the selection of suitable concepts is a key point in the implementation of the decision net. The decision net must be as generic as possible and the concepts of the decision net must represent general knowledge about the decision making task and the past experience.

1. The perception layer contains concepts that are directly shown by the DSF in every time period: the current level in the tank ( $Level_t$ ), the goal level ( $Goal$ ), the environment outflow and inflow in the previous time period ( $EO_{t-1}$  and  $EI_{t-1}$ ), and the last decision of the participant ( $UO_{t-1}$  and  $UI_{t-1}$ ).
2. The short-term memory layer takes into account the recent past experience and contains the last three values of the tank level ( $Level_{t-1}$ ,  $Level_{t-2}$  and  $Level_{t-3}$ ), the environment inflow ( $EI_{t-2}$ ,  $EI_{t-3}$  and  $EI_{t-4}$ ) and the environment outflow ( $EO_{t-2}$ ,  $EO_{t-3}$  and  $EO_{t-4}$ ).
3. The working memory layer provides elaborated information via concepts that represent the increase in the last environment inflow and outflow, the decrease in the last environment inflow and outflow and the positive or negative difference between the current level and the goal level.
4. The deliberative layer aggregate other concepts contained in the previous three layers. This layer is composed of two deliberative concepts.
5. The output layer contains the concepts used by MAIDEN-DSF in the value function: the estimated environment inflow and outflow ( $EI$  and  $EO$ ).

Every node of the decision net has got an associated positive activation value. The decision net propagates activation to one node by performing the weighted sum of the incoming activations from the nodes connected to it. The following expression shows how a node computes the total weighted input  $net$ :

$$net_i = \sum_j w_{ij} \cdot a_j \quad (1)$$

In expression 1,  $net_i$  represents the total weighted input of the  $i^{th}$  node,  $a_j$  is the activation of the  $j^{th}$  node and  $w_{ij}$  is the weight of the connection between the  $i^{th}$  and the  $j^{th}$  node. If  $net_i \geq 0$ , then the activation of a node is equal to its  $net$  value. Otherwise activation is zero.

The connections among the different nodes of the decision net have three constraints. First, nodes of the same layer are not interconnected. Second, nodes of the deliberative layer

can only propagate activation to the output layer. Third, a concept belonging to the perception, the short-term memory or the working memory layer can propagate activation to the deliberative layer or the output layer, but it cannot propagate activation to both layers according to neurophysiological evidence (Damasio, 1994; Romanski & LeDoux, 1992).

The concepts of the perception layer are easily identifiable in the DSF. The concepts of the short-term memory and the working memory layer have been extracted from interviews of participants who performed the DSF. Notice that the activation of these concepts must be updated every time period depending on the decisions and consequent outcomes up to the current time period. Figure 1 shows a graphical representation of the decision net.

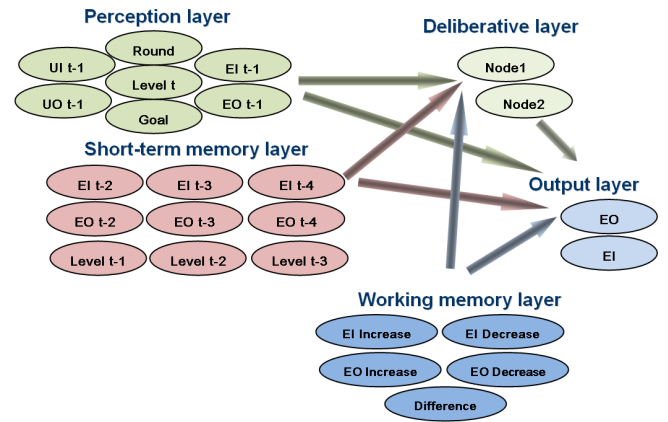


Figure 1: Representation of the decision net.

This design is coherent with Norman's definition of affordances (Norman, 1988): all action possibilities latent in the environment that are perceivable by a participant. These perceived affordances result from the mental interpretation of things based on past knowledge and experience. In MAIDEN-DSF, alternatives or actions are also evaluated based on past knowledge and experience.

## Value function

The goal of the task is to keep the water in the tank at a given level. The expression 2 represents the water level at time period  $t + 1$  ( $Level_{t+1}$ ), which depends on  $UI_t$ ,  $UO_t$ ,  $EI_t$ ,  $EO_t$  and  $Level_t$ .

$$Level_{t+1} = Level_t + UI_t - UO_t + EI_t - EO_t \quad (2)$$

The difference between the water level and the goal level at time period  $t + 1$  is calculated using the following expression:

$$Diff = Level_{t+1} - Goal \quad (3)$$

Since the best decision is the one that sets the stock at the goal level ( $Goal$ ), the value of  $Diff$  must be 0 and the best  $UI_t - UO_t$  value is given by the combination of equations 2 and 3:

$$UI_t - UO_t = Goal - Level_t - EI_t + EO_t \quad (4)$$

The expression 4 computes the best decision ( $UI_t - UO_t$ ) and requires the values of  $Goal$  and  $Level_t$ , which are known to the participant, and the values of  $EI_t$  and  $EO_t$ , which are unknown. Therefore, the value function used by MAIDEN-DSF replaces the unknown variables  $EI_t$  and  $EO_t$  with the activation of the concepts  $EI$  and  $EO$  of the decision net estimated with the information available in the current time period  $t$ , as shown in figure 2.




Figure 2: Scheme of MAIDEN's value function.

## Evaluation

The connection weights of the decision net of MAIDEN-DSF have been adjusted to model, as well as possible, the behavior of a set of participants who performed the DSF. The evaluation consists in the prediction of the behavior of another set of participants in different conditions of the DSF. The training and test data has been collected from the DSF challenge (Lebiere et al., 2010) whose homepage presents a wealth of information about the task (<http://www.hss.cmu.edu/departments/sds/ddmlab/modeldsf>). The measure used to evaluate the model fit was Pearson's linear correlation coefficient  $R$ , comparing the decisions of the model and the participants over all 100 time periods. In order to compute the decisions of MAIDEN-DSF, in each time period the activation of the concepts of the perception, short-term memory and working memory layers has been generated with the participant's actual sequence of decisions and consequent water levels up to the current time period. This method ensures that MAIDEN-DSF makes a decision with the same available knowledge as the participant whose behavior is wanted to be modeled.

The values of the connection weights that optimize the correlation of the model and the participants' decisions have been calculated by the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) (Hansen, 2006), which is a kind of Evolutionary Algorithm (Goldberg, 1989; De Jong, 2006) with considerable potential for searching in complex spaces. CMA-ES has been run ten times to find the best connection weights that optimize the correlation of the decisions of MAIDEN-DSF and the decisions of the participants who participated in the training conditions. Then, the correlation of the best connection weights is presented.

## Performance in the training conditions

The DSF with the training conditions used an environment inflow that was either an increasing function or a decreasing function over 100 time periods. The function could be linear or non linear and, therefore, there were four training conditions. The environment outflow function was constant and set to zero throughout the task. The following expression shows the linear increasing function:

$$EI_t = 2.0 + 0.08 \cdot t \quad (5)$$

The goal was to maintain the level of water in 4 gallons during all 100 time periods. The initial water level in the tank was fixed to 2 gallons. The training data corresponds with the behavior of 61 participants (linear increasing = 15 participants, linear decreasing = 11 participants, non linear increasing = 17 participants and non linear decreasing = 18).

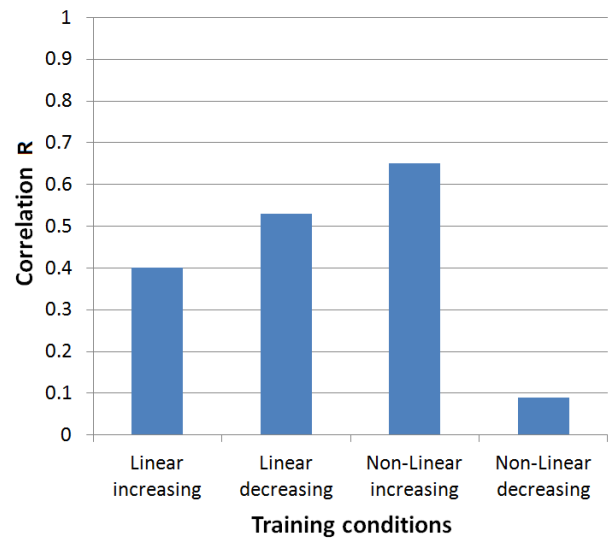


Figure 3: Correlation between MAIDEN-DSF and human decisions with the training conditions.

Figure 3 shows the correlation between the decisions of MAIDEN-DSF, whose connection weights have been optimized in order to fit the training data, and the decisions of the participants with the training conditions. All the  $R$  values are statistically significant using a Student's  $t$  distribution for a transformation of the correlation ( $p < 0.001$ ). These values correspond with the best set of connection weights found by CMA-ES after ten executions.

In the training set, the worst correlation is obtained in the non linear decreasing condition which appears to be the hardest task for the participants. The whole training set of participants has been modeled with the same connection weights. This low correlation value may be due to connection weights that model well the behavior of participants in different conditions but not in the non linear decreasing condition. Note

that CMA-ES finds the connection weights that maximize the overall correlation taking into account the four conditions.

### Performance in the test conditions

The values of the connection weights of MAIDEN-DSF were estimated to optimize the correlation of the model and the participants' decisions with the training conditions. This subsection shows the predictions made by this MAIDEN-DSF in the test conditions in order to validate the model.

The DSF with the test conditions used an environment inflow that was either a sequence-based function or a delayed function. There were three different sequence-based functions that generated a repeated sequence of length 2, a repeated sequence of length 4 and the same sequence of length 4 with noise. There were also two delayed functions where participant's decisions were delayed for either two or three time periods. Therefore, there were five test conditions. The environment outflow function was also constant and set to zero throughout the task. The goal was to maintain the level of water in 6 gallons during all 100 time periods and the initial water level in the tank was fixed to 4 gallons. The test data corresponds with 100 participants (delay 2 = 20 participants, delay 3 = 20 participants, sequence 2 = 20 participants, sequence 4 + noise = 20 participants and sequence 4 = 20 participants).




Figure 4: Correlation between MAIDEN-DSF and human decisions with the test conditions.

Figure 4 shows the correlation between the decisions of MAIDEN-DSF and the decisions of the participants with the test conditions. The  $R$  values obtained with the test conditions are all statistically significant ( $p < 0.001$ ).

It is noteworthy that the decision net of MAIDEN-DSF

does not contain any concept representing the delay; however, it models the human behavior with the delay conditions with a correlation higher than 0.4, even when the training conditions do not contain any delay. This result suggests that participants may not take into account the delay or may not understand it while performing the DSF.

In the delay 3 condition the correlation is reduced. This can be due to the high variability observed in the human decisions, for example, a participant decided to increase the level of water in  $1E + 09$  gallons.

With respect to the sequence-based conditions, the short-term memory layer of the decision net contains concepts from the previous four time periods up to the current one. A good performance in the sequence-based conditions implies the knowledge of the last two environmental inputs, if it is a sequence of length 2, or the last four environmental inputs, if it is a sequence of length 4. MAIDEN-DSF contains concepts representing the last four environmental inputs, so it may suitably model behaviors with these conditions.

### Delay conditions

The aim of this work is to find cues that prove whether participants do not perceive the delay during the DSF within the delay conditions. Although the results obtained in the previous experiment may already suggest that participants ignore delays, this section presents another cue to support the experiment.

A new version of MAIDEN-DSF with concepts that represent the delay has been also built. This new version will be denoted as MAIDEN-DSFd. The new decision net is composed of the concepts already explained in this paper with the addition of four new concepts:

- The short-term memory layer takes also into account the following user decisions:  $UI_{t-2}$ ,  $UI_{t-3}$ ,  $UO_{t-2}$  and  $UO_{t-3}$ .

Note that with the delay conditions the user inputs are delayed for two or three time periods, so the decisions relevant for the current tank level  $Level_t$  are taken in  $t - 2$  and  $t - 3$ .

The connection weights of MAIDEN-DSF and MAIDEN-DSFd have been adjusted to model the behavior of the 40 participants that participated in the delay conditions. CMA-ES has been applied ten times for each model version to estimate the connection weights that maximize the correlation between the model decisions and the decisions of the participants.

Table 1: Mean correlation (E) and standard deviation (SD) obtained by MAIDEN-DSF and MAIDEN-DSFd with the delay conditions.

Model	Delay 2		Delay 3	
	E	SD	E	SD
MAIDEN-DSF	0.54710	1.0E-06	0.43931	5.0E-09
MAIDEN-DSFd	0.54711	1.3E-06	0.43931	6.0E-09

Table 1 shows the mean correlation and the standard deviation of the ten solutions found by CMA-ES for each version

of MAIDEN. The  $R$  values obtained with both implementations are all statistically significant ( $p < 0.001$ ).

Table 2: Mean (E) and standard deviation (SD) of the root mean square error obtained by MAIDEN-DSF and MAIDEN-DSFd with the delay conditions.

Model	Delay 2		Delay 3	
	E	SD	E	SD
MAIDEN-DSF	24275.04	0.033	6227685.70	0.098
MAIDEN-DSFd	24274.98	0.036	6227685.59	0.059

Table 2 shows the mean and the standard deviation of the root mean square error obtained by the ten solutions found by CMA-ES for each version of MAIDEN. The error seems to be high due to the high variability of the human decisions within the delay conditions. In fact, in the DSF challenge (Lebiere et al., 2010) the best ranked model obtained a root mean square error of 25560.84 in the delay 2 condition and 5930065.84 in the delay 3 condition.

MAIDEN-DSF and MAIDEN-DSFd obtained similar correlations in both conditions. Within the delay 2 condition they only differ in 0.00001 and within the delay 3 condition the mean correlation is exactly the same. The standard deviation is also similar. These results point out that the behavior of the participants might be modeled without using the  $UI_{t-2}$ ,  $UI_{t-3}$ ,  $UO_{t-2}$  and  $UO_{t-3}$  decisions.

## Conclusion

A knowledge-based model that emulates human behavior in the Dynamic Stocks and Flows task has been proposed. MAIDEN-DSF operation is divided in two main phases. First, it estimates the environment inflow and outflow values using a decision net. Second, it uses a value function to choose the best alternative, which is the one that sets the water level of the tank at the goal level.

The model has been evaluated with two data sets from the DSF challenge (training set and test set). The evaluation lies in the prediction of human behavior in the test set given the training set. The results achieved support that MAIDEN-DSF is a considerable framework in order to model human behavior in the DSF. Besides, the proposed model participated in the DSF challenge and reached rank 2 regarding the ranking provided by the DSF challenge organizers based on the  $R^2$  correlation and the root mean square error.

MAIDEN-DSF has been used to find cues that point out that the behavior of the participants who performed the DSF task with the delay conditions can be modeled without taking into account any concept representing the delay. This evidence suggests that participants do not take into account or are not aware of the delay while making a decision during the delay conditions.

The results suggest that it is very difficult to configure a model in order to emulate the performance of several human beings with the same parameter values (connection weights in this case). This might lead to develop comparison procedures

focusing on individual modeling and adaptation. MAIDEN is mainly based on the knowledge acquired through past experience, the knowledge extracted from the environment and the relationships between the concepts that represent these two kinds of knowledge. A very interesting feature is that the connection weights of the decision net can show what concepts are the most important for a participant. In this experiment, the whole set of participants has been modeled with the same connection weights. A future work might consist in the individual modeling of each participant in order to study the connection weights that better fit each participant's behavior. An individual modeling may indicate which participants had realized that there was a delay in their decisions. In this experiment the connection weights of the concepts representing  $UI_{t-2}$ ,  $UI_{t-3}$ ,  $UO_{t-2}$  and  $UO_{t-3}$  in MAIDEN-DSFd affect the behavior of the model. An absence of connection between these concepts implies that the participants do not use this knowledge. However, CMA-ES seeks the best combination of connection weights that better fit the whole set of human behaviors, therefore, if there is one person that uses this knowledge, CMA-ES will find a connection weight for these concepts. This is the reason for the better suitability of individual modeling.

A global modeling is suitable for experiments where there are at least two groups: a control group and an experimental group. A future work may study the connection weights that better model the behavior of participants who performed well in the delay conditions (control group) and the participants who obtained poor results (experimental group). The connection weights can point out what concepts affect a decision for a certain participant. For instance, this information may be used to improve participants' performance by explicitly revealing relevant information that participants who obtained poor results tend to ignore but participants who performed well in the delay conditions take into account. This may lead to better understanding of dynamic decision making.

## Acknowledgments

This work has been supported by the JAE Program (Predoctoral research training leading to PhD Theses).

## References

- Brehmer, B. (1992). Dynamic decision making: Human control of complex systems. *Acta Psychologica*, 81(3), 211-241.
- Cronin, M. A., & Gonzalez, C. (2007). Understanding the building blocks of dynamic systems. *System Dynamics Review*, 23(1), 1-17.
- Damasio, A. R. (1994). *Descartes' error: emotion, reason, and the human brain*. New York: Putnam.
- De Jong, K. A. (2006). *Evolutionary computation: A unified approach*. Cambridge, Massachusetts: The MIT Press.
- Goldberg, D. E. (1989). *Genetic algorithms in search, optimization and machine learning*. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc.

- Gonzalez, C., & Dutt, V. (2011). A generic dynamic control task for behavioral research and education. *Computers in Human Behavior*, 27(5), 1904 - 1914.
- Hansen, N. (2006). The cma evolution strategy: A comparing review. In J. Lozano, P. Larranaga, I. Inza, & E. Bengoetxea (Eds.), *Towards a new evolutionary computation* (p. 75-102). Springer.
- Iglesias, A., Del Castillo, M. D., Serrano, J. I., & Oliva, J. (2010). A psychologically and neurophysiologically plausible model for emulating human behavior in decision making tasks. In *Proceedings of the brain inspired cognitive systems - bics 2010*. Madrid: Universidad Politecnica de Madrid.
- Lebiere, C., Gonzalez, C., & Warwick, W. (2010). Editorial: Cognitive architectures, model comparison and agi. *Journal of Artificial General Intelligence*, 2(2), 1-19.
- Norman, D. (1988). *The psychology of everyday things*. New York: Basic Books.
- Romanski, L. M., & LeDoux, J. E. (1992). Equipotentiality of thalamo-amygdala and thalamo-cortico-amygdala circuits in auditory fear conditioning. *Journal of Neuroscience*, 12, 4501-4509.
- Sterman, J. D. (2002). All models are wrong: Reflections on becoming a systems scientist. *System Dynamics Review*, 18(4), 501-531.