Paper presented at the 32nd International Conference of the System Dynamics Society:
Good Governance in a Complex World. Delft, Netherlands, July 20-24, 2014

# Decision Heuristics and Human Performance in a Policy Management Simulation

**Daniel Lafond[1], Jean-François Gagnon[2], Marie-Ève St-Louis[2], Sylvain Pronovost[2], Michel B. DuCharme[3], Sébastien Tremblay[2]**
[1]Thales Canada, [2]Université Laval, [3]Defence R&D Canada

Thales Research and Technology Canada
Thales Canada Inc., 1405 Parc-Technologique,
Quebec City, Canada, G1P 4P5
Tel: 418-651-0606 x4510652 Fax: 418-651-1953
e-mail: daniel.lafond@ca.thalesgroup.com

**Abstract**

*Psychometric research has delivered reliable means for assessing various forms of intelligence, yet there has been relatively little success in predicting the human ability to solve complex problems in dynamic environments. The present work aims to profile dynamic decision making strategies using dynamic simulations in order to predict individual complex problem solving performance. We report an experiment assessing decision heuristics with the goal to predict complex problem solving ability. We used the COmplex DEcision Making (CODEM) system dynamics testbed to assess information seeking behaviors and the similarity of decision patterns to different types of heuristics. The Democracy 2 serious game is then used as an objective measure of complex problem solving ability. Democracy 2 is a realistic government management simulation requiring strong planning and systems thinking skills. A set of three new metrics is proposed to quantify similarity to different heuristics. Three models are compared on the basis of their predictive accuracy: a linear regression model, an artificial neural network and a support vector machine. Results show that the support vector machine has the most potential due to its superior results in a cross-validation test. We conclude with a discussion on future model extensions and generalization tests.*

**Keywords:** Complex problem solving, policy making, decision making, heuristics, experimentation, simulation, modeling.

## 1. Introduction

Given the complexity of the interrelated network of factors in a society (e.g., economy, education, health, security, justice and environment), individuals in charge of its governance ought to possess strong complex problem solving (CPS) skills to increase the likelihood of implementing successful policies.

CPS ability is known to vary considerably across individuals (Fischer, Greiff, & Funke, 2012; Rouwette, Größler, & Vennix, 2004), yet there is a lack of psychological models able to account for this variability at a quantitative level. A number of studies have found a relationship between scores on intelligence tests and decision making performance (Elg, 2005; Gonzalez, Thomas & Vanyukov, 2005; Goode & Beckmann, 2010). However, these results are not consistent across studies (Kluwe, Shilde, Fisher, & Oellerer, 1991; Rigas & Brehmer, 1999). Based on a review of experimental research, Wenke, Frensch and Funke (2005) conclude that intelligence as a general ability may be an overly broad concept to allow reliable prediction of individual differences in complex problem solving. Interestingly, CPS ability, as measured using dynamic simulations, has been found to predict differences in academic achievement and supervisor ratings beyond those explained by intelligence tests (Danner et al., 2011; Wüstenberg, Greiff, & Funke, 2012).

Dörner (1986) introduced the notion of *operative intelligence* to describe the higher order cognitive abilities necessary for CPS. Operative intelligence is essentially about problem solving capabilities such as circumspection (e.g., anticipating of long-term and side-effects of interventions), strategy selection (e.g., trial-and-error, systematic analysis, adoption of heuristics), and the ability to set and manage subgoals.

One potential limitation of using a CPS simulation as a psychometric test is its volatility (Rigas, Carling, & Brehmer, 2002). Complex problems can be unforgiving, allowing little room for error despite having a generally good approach to problem solving. The duration of CPS simulations tends to make it impractical to perform multiple tests to increase measurement reliability. We therefore suggest that a more robust measure of CPS ability (i.e., operative intelligence) may be the overall decision making process as measured in a diagnostic scenario. Accordingly, the goal of the present work is to characterise such a pattern over the course of a scenario in a generic way and to determine if it can reliably predict CPS performance in a different scenario (i.e., better than mere performance on the first scenario). The proposed approach combines the use of information acquisition behavioural markers and of newly defined metrics that indicate the overall similarity of the decision pattern to three high-level heuristic signatures.

The present paper is organised as follows. Following this introduction, Section 2 presents the apparatus and methodology used for data collection. Section 3 presents the data analysis procedure and model selection results. Section 4 discusses the implications of our findings and directions for future work.

## 2. Method

*Participants.* 22 adults (12 women and 10 men; mean age: 23.0 y, SD: 7.73) participated in a 4-hour experiment split in two 2-hour sessions.

*Design and Procedure.* Participants were all assigned to the same experimental condition. This group constitutes a baseline condition for upcoming work on CPS training. Session 1 includes a tutorial, a familiarization scenario, a practice scenario, and the diagnostic "Stability Operations" scenario, all played within the COmplex DEcision Making experimental platform (CODEM) system dynamics testbed. Session 2 involves a tutorial and test scenario in the Democracy 2 simulation.

*Apparatus.* The experiment was run on a standard personal computer in a laboratory setting, using CODEM (Defence R&D Canada) and Democracy 2 (Positech Games).

## 2.1 CODEM

The CODEM system dynamics simulator is a "microworld" or "interactive learning environment" platform for the design and administration of complex problem solving tasks. The underlying problem structure is defined using stocks (situation variables) and flows (relations between variables). CODEM provides extensive experimental manipulation and data logging capabilities for research purposes. It controls dynamic decision making scenarios where players (individuals, teams, or adversaries) can allocate their resources amongst different possible interventions in order to influence the state of the system. The flexible scenario editor allowed creating a highly challenging fictional stability operations scenario (Lafond & DuCharme, 2011). In this scenario, participants are in charge of stabilizing a failing state in the midst of a rising insurgency. Participants can allocate resources called "action points" into seven different intervention types (Table 1, left column). Furthermore, the state of the situation in a given turn is described through nine variables ranging from 0 to 20 (Table 1, right column).

*Table 1. Possible interventions and situation variables in the stability operations scenario.*

| **Possible Interventions** | **Situation Variables** |
|---|---|
| Security operations | Host-nation governance |
| Influence operations | Population allegiance |
| Cultural training | Local media |
| Humanitarian aid | Criminality suppression |
| Training of local forces | Socio-economic welfare |
| Infrastructure development | Local forces |
| Governance capacity building | Infrastructures |
| | Cultural understanding |
| | Insurgency suppression |

In the simulation, the nine situation variables mutually influence each other so that each decision results in a chain of effects within the system. Depending on its current value, each variable can be in a desirable or undesirable state as described by a three-color scale that goes from green to orange to red (for simplicity, higher values correspond to more favorable states for all variables). Feedback on the changes occurring in the situation is provided during the transition from one turn to the next.

The goal of the participant is to bring all eight dimensions (cultural understanding is a mediating variable but not a sub-goal) outside of the "critical" (red) state in seven turns or less (this goal can be achieved in four turns). The mission has failed if the allegiance of the local population falls to zero. The underlying model captures several key characteristics of complex dynamic systems (reinforcing and balancing feedback loops, delayed effects, uncertainty, partial opacity, etc.). Figure 1 illustrates the different tabs of the CODEM interface.

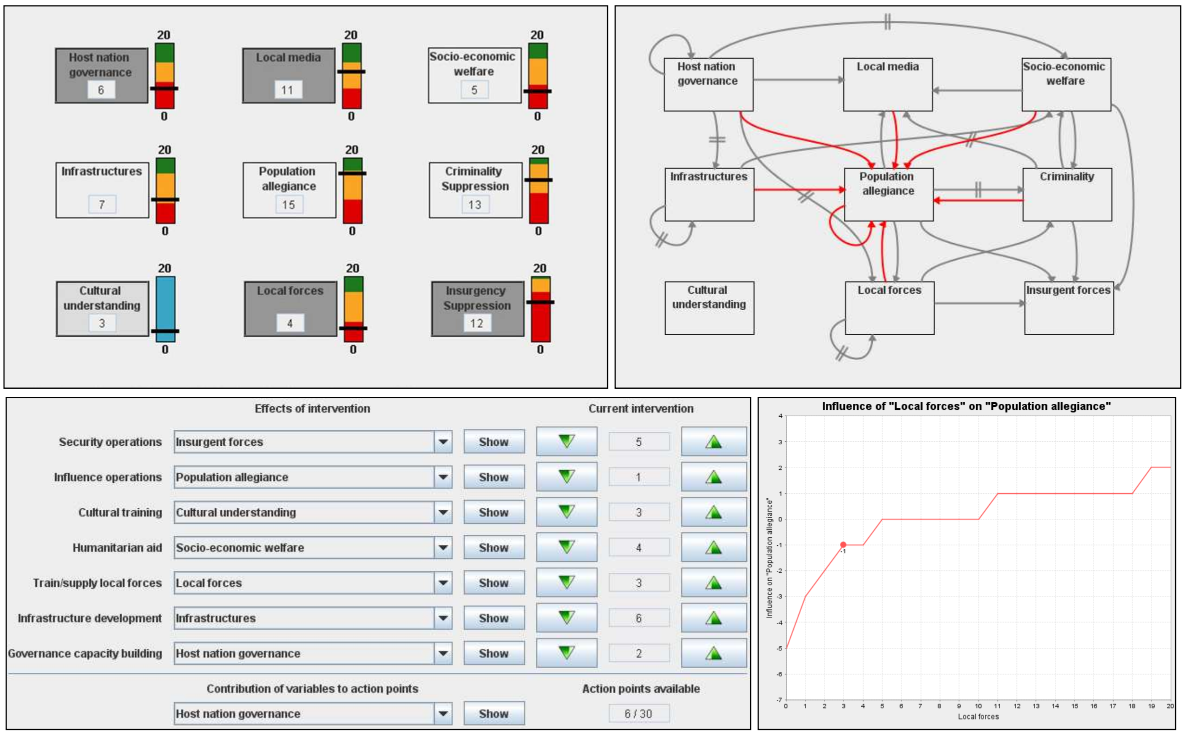

Figure 1. CODEM interface (Stability Operations Scenario)

The Situation tab (upper left of Fig. 1) shows the current value of system variables. White boxes are for standard variables, dark boxes refer to agents, and grey boxes indicate mediating variables. The Relations tab (upper right of Fig. 1) shows the interactions between variables. Double-bars indicate delayed effects. The Decision tab (lower left of Fig. 1) shows the different intervention options available, the current amount of resources (action points) available, and factors influencing action points that will be available in future turns.

Effects in the relations tab and the decision tab can be visualized in a graph (by clicking on a relation or on "show"). The bottom right of Figure 1 shows a graph illustrating the effects of a variable or intervention (*x*-axis) on another variable (*y*-axis). Effects can include delays and can be roughly linear or highly non-linear. Effects can also be conditional and vary according to the current situation. Feedback on the changes in the situation is shown after each turn. The scenario ends when the desired end-state is achieved or when the maximum number of turns is reached.

## 2.2 Democracy 2

Democracy 2 is an award-winning government simulation game that recreates the major systems thinking challenges in the areas of welfare, economy, taxes, public services, law and order, foreign affairs and transportation that policy makers and society as a whole actually face every day. It is a *serious game*, i.e., a simulation environment designed for training and education purposes rather than pure entertainment. Figure 2 shows the main game interface.

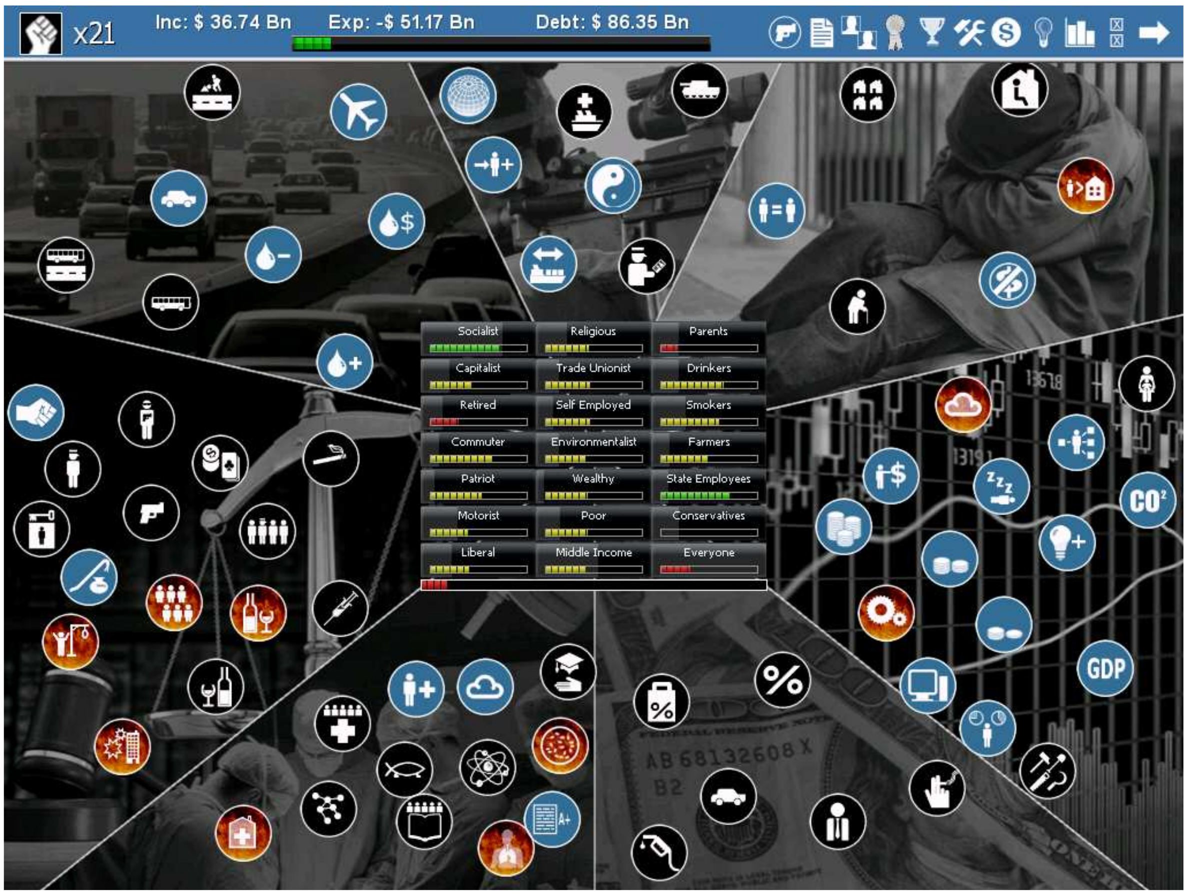

Figure 2. Main Democracy 2 interface. Blue icons are situation variables. Black icons are active policies. Red icons are critical problems. The table at the center of the screen shows the opinion of different population groups. The top icons, from left to right, refer to action points, income, expenses, debt, proportion of the mandate completed, intelligence reports, quarterly report, cabinet members, party membership, player achievements, game options, budget report, new policies, polls, review of promises, and start next turn.

This serious game models a government's mandate, from the moment of taking office, to the next election, at which time the simulated scenario is either successful or a failure, depending on the election results (i.e., the simulation is limited to a single mandate to constrain experiment duration). By making choices which influence voters, partisans, and cabinet members, the player must strive to balance the popularity of his or her government with the effectiveness of policies which have direct or indirect impacts on matters such as economy, social programs, education, and health care. Unexpected events such as natural disasters, financial crises, civil uprisings, military conflicts, or even assassination attempts can also occur.

Democracy 2 is a turn-based game where resources (the finances of the state of which the government is in charge) are used to implement a number of weighted choices (different policies organized in categories such as education, health, economy, environment, etc.). Decisions impact a number of variables such as pressure groups, voters' intents, and the situation variables targeted by the interventions (health, economy, education, etc.). There are also time delays in policy implementations and feedback. The intuitive interface facilitates exploration of causal relations to understand the system's dynamics. The system's causal structure is highly transparent, i.e., the quantitative relationships between policies and target variables are accessible to direct observation, and may be leveraged to enhance the comprehension of the system's inner workings. Figure 3 shows an example of the relations that appear when hovering the mouse over an icon.

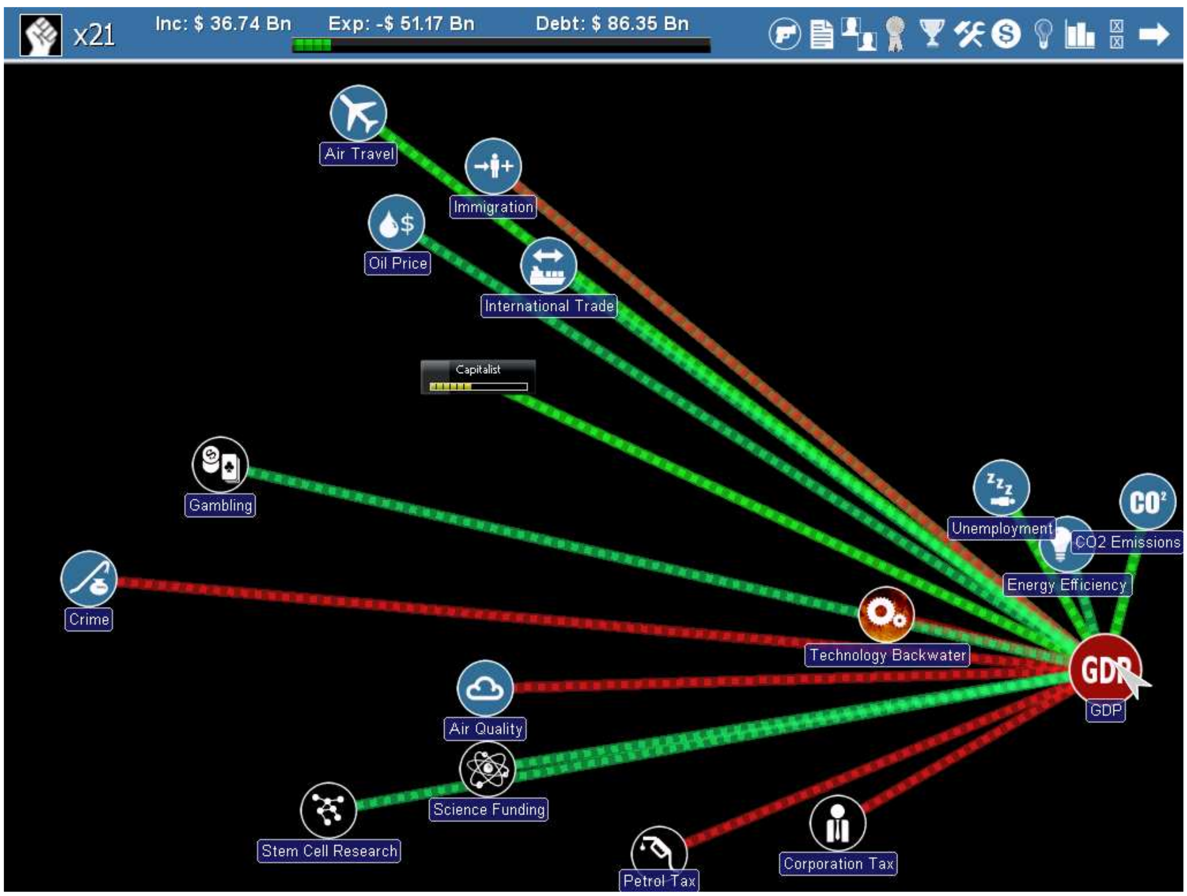

Figure 3. Interrelations between variables displayed in Democracy 2. Green/red lines indicate that the impact is to increase/decrease the value of the target variable. The

green/red lines are dynamic, creating a visual flow effect moving in the directionality of the effect at a speed that is proportional to the strength of the effect.

When introducing a new policy or modifying an existing one, the player sees a menu that indicates policy effects and allows setting the policy value using a slider bar, as seen in Figure 4.

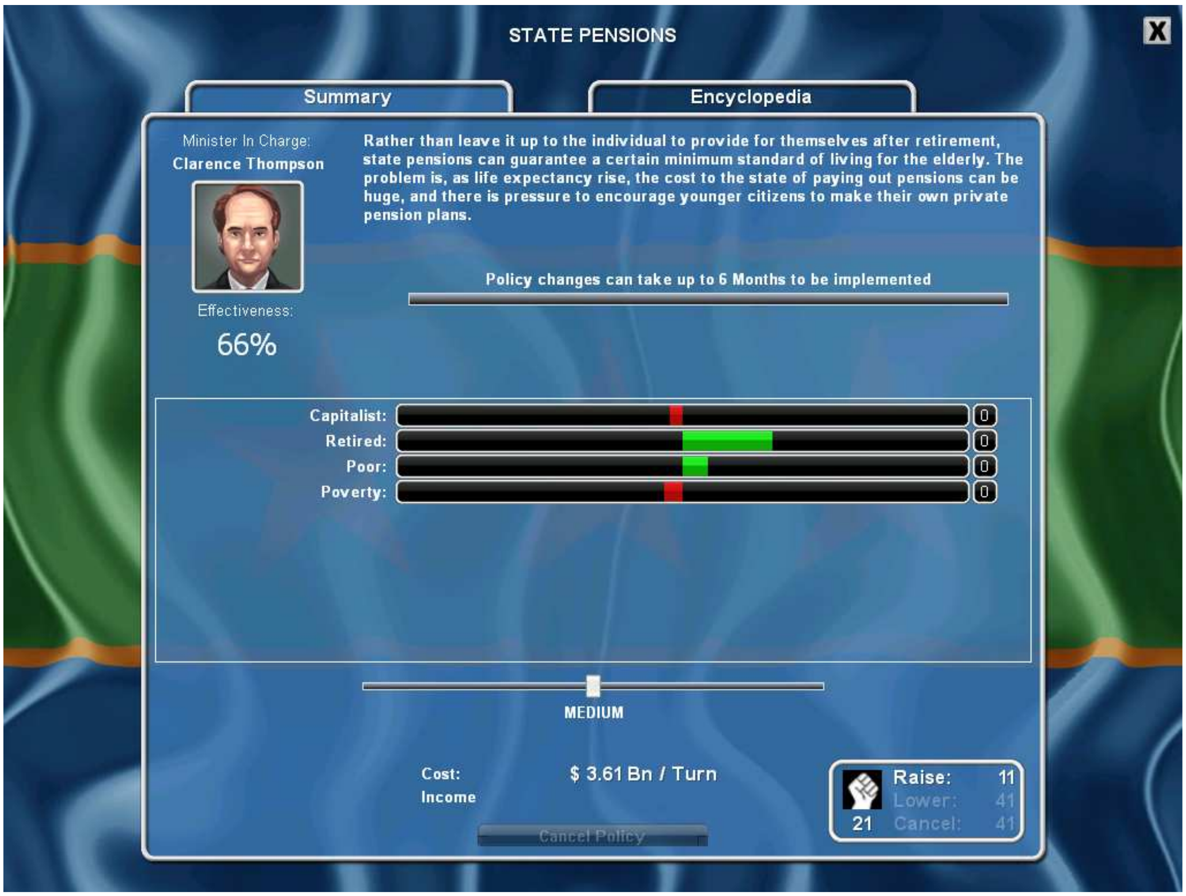

Figure 4. Policy introduction/modification menu.

The specific Democracy 2 scenario selected for the experiment is called Malaganga, a fictional debt-ridden state where voting is compulsory. Instructions given to the participants are to obtain the highest possible number of votes, while keeping debt as low as possible.

## 2.3 Measures

Several measures are collected in the Stability Operations scenario to be used for prediction purposes, while a single performance measure is collected in the Democracy 2 scenario.

*Score in the Stability Operations scenario.* Performance is measured by the relative distance from the eight sub-goals, and is based on the proportion of the seven-year mandate completed, resulting in a scale ranging from 0 to 100. Hence, reaching the mission goal at the end of the seventh turn (the last turn) does not yield a score of 100. A score of 100 is attributed to reaching the goal in as few turns as possible (i.e., on Turn 4 in this particular scenario).

*Information seeking frequency in the Stability Operations scenario.* CODEM logs the frequency of requests for graphs displaying the effects between variables (relations tab) and effects of interventions (decision tab).

*Decision Pattern in the Stability Operations scenario.* CODEM logs the situation and decisions made on each turn. This allows the systematic assessment of the similarity of each participant's decision pattern to three types of heuristics.

- *Flat heuristic*. One simple heuristic can be to assign an equal importance to the various intervention options. This assumes that a balanced approach is a good general strategy to solving complex problems.
- *Static heuristic*. A strategy to save time and effort is to assume that there is a single resource allocation with the best possible trade-offs, that is robust enough to be repeated throughout the scenario without the need to adjust to changes in the situation.
- *Goal-Distance heuristic*. Since there are several well-defined subgoals to achieve in the stability operations scenario, it is possible to operationalize a heuristic that calculates the goal-distance for each variable on each turn. The heuristic then assigns a resource allocation proportional to the expected benefits minus adverse effects of each intervention, weighted by goal-distance.

The heuristic similarity metrics are obtained by calculating the root of the mean squared error (RMSE), between the observed proportion of resources allocated to each intervention, and the proportion corresponding to each heuristic. Since RMSE varies from zero to one in the present case, the similarity metric corresponds to 1-RMSE. Note that the heuristic similarity metrics described above are seen as non-mutually exclusive dimensions that together provide a set of features to characterise the overall strategy used by a participant.

*Proportion of Action Points Used (Stability Operations Scenario).* Since participants had the option to use up all their action points in a given turn or not, this metric simply captures that decision. This measure may help differentiate those who understand the severity of the initial situation and those who do not see that the system is on the verge of a major crisis.

*Score in Democracy 2.* For the purposes on the present study, the score in Democracy 2 is derived by dividing the accumulated debt by the number of votes. Scores are then standardized and the sign is inverted so that positive Z-scores correspond to relatively good results and negative Z-scores correspond to relatively poor results, using the average as the main reference point. A Z-score of 0 means that the score is the same as the mean. A Z-score of 1 means that the participant's result is one standard deviation above the mean.

## 3. Data Analysis and Results

A first analysis involved looking at the correlation between the Democracy 2 score and the set of individual predictors considered in the present study. Table 2 shows the set of distinct predictors investigated. In the case of the information seeking metric, the first turn is considered separately from other turns because it was previously found to be critical in predicting success (within the same simulation; Gagnon, Lafond, DuCharme, St-Louis & Tremblay, 2012). Indeed, Turn 1 typically includes extensive initial planning on the part of the participants.

*Table 2. CODEM behavior/outcome markers extracted from the logs.*

| **Markers** |
|---|
| Score in the Stability Operations scenario |
| Information seeking (Relations) total frequency |
| Information seeking (Decision) total frequency |
| Information seeking (Relations) frequency Turn 1 |
| Information seeking (Decision) frequency Turn 2+ |
| Information seeking (Relations) frequency Turn 1 |
| Information seeking (Decision) frequency Turn 2+ |
| Proportion of action points used on Turn 1 |
| Closeness to Flat heuristic |
| Closeness to Static heuristic |
| Closeness to Goal-distance heuristic |

The correlation analysis showed that the score in the Stability Operations scenario was not significantly related to the score in Democracy 2, $r(20) = .027$, n.s. However, the analysis showed that closeness to the Goal-distance heuristic was significantly related to the score in Democracy 2, $r(20) = .487$, $p = .022$ (two-tailed).

### Model Comparison

Three candidate models were considered in order to attempt to predict CPS performance in Democracy 2: a linear regression model, an artificial neural network, and a support vector machine. These candidate models were implemented in the RapidMiner (RapidMiner Inc.), integrated environment for data mining, machine learning, and predictive analytics. The objective of this effort is to come up with a model with a superior predictive accuracy compared to a baseline single factor regression model based on the closeness metric to the goal-distance heuristic.

The first candidate model was W-Linear Regression. This model employs the M5 prime feature selection method which relies on the Akaike Information Criterion (AIC; Akaike, 1974). The AIC uses information entropy to allow a trade-off between the number of factors in the linear regression and the error (i.e., residual sum of squares). Through successive iterations, the method selects the attribute with the smallest standardized coefficient, removes it and performs another regression. An attribute is dropped if removing it results in an improved AIC. This procedure is repeated until no remaining attribute can be dropped. The second candidate model was the W-Multilayer Perceptron, a classic artificial neural network using the backpropagation supervised learning technique (Rumelhart, Hinton, & Williams, 1986). The third candidate model was a Support Vector Machine (SVM; also called support vector regression when used to predict a continuous variable) implemented by Rüping (2001). An SVM model represents examples as points in a multidimensional space (i.e., defined by the examples' features), using a mapping that separates as much as possible examples with different target values. New examples are then mapped into that space in order to predict their target value. The special property of SVMs is that by maximizing the geometric margin, they tend to minimize generalization error (Drucker, Burges, Kaufman, Smola, & Vapnik, 1997).

The baseline single-factor model was found to account for 23% of the variance in CPS ability (using the $R^2$ goodness of fit indicator). Yet more importantly, we performed a cross-validation test (using the leave-one-out resampling procedure; see Browne, 2000) in order to estimate the average prediction error. The result from the baseline single factor model was a RMSE of .707 (i.e., the unit being a standard deviation from the mean).

The resulting linear regression model retained four predictors (plus an intercept of -18.05):

- Information seeking (relations) total frequency (coefficient = .0084)
- Closeness to Flat heuristic (coefficient = -15.56)
- Closeness to Static heuristic (coefficient = 8.25)
- Closeness to Goal-distance heuristic (coefficient = 30.23)

The multilayer perceptron will not be described in detail herein, but suffice to say that the two hidden layer nodes in the model relied on exactly the same four predictor inputs as the linear regression model. The SVM did not perform any feature selection and therefore used all the available predictors to some extent. Table 3 shows results of the three candidate models in terms of goodness of fit and average prediction accuracy.

*Table 3. Goodness-of-fit and average prediction error of each model.*

| **Model** | **Statistic** | |
|---|---|---|
| | $R^2$ | RMSE |
| W-Linear Regression | .44 | .777 |
| W-Multilayer Perceptron | .94 | .764 |
| Support Vector Machine | .16 | .470 |

The key finding is that the SVM was the only model that achieved a lower prediction error compared to the baseline single factor model. The two other models were clearly overfitting the data at the detriment of predictive accuracy (Myung, Pitt, & Kim, 2005).

## 4. Discussion

The objective of the present paper was to use a human-in-the-loop system dynamics simulation to collect behavioral markers and investigate whether they could reliably predict individual differences in CPS performance using a different scenario. We proposed new heuristic similarity metrics to characterise key features of a participant's overall strategy. A correlation analysis showed that the score in the Stability Operations scenario was not significantly related to the score in Democracy 2. This result was not particularly surprising given the well-known variability in CPS simulations. Indeed, our hypothesis was that a greater sensitivity to individual differences can be achieved by taking into account the process rather than only the outcome of CPS. We compared three candidate models to capture the relationship between the predictors and the Democracy 2 scores: 1) a linear regression model; 2) an artificial neural network; and 3) a support vector machine. Results indicate that only the support vector machine was able to improve on a basic single-factor model due to its superior predictive accuracy ascertained using a cross-validation test.

Future research on the development of a cognitive model of CPS ability could benefit from an integration of behavioral markers and assessments of fundamental cognitive abilities such as working memory capacity (e.g., Gonzalez et al., 2005), fluid intelligence (Raven, Raven, & Court, 1998) and updating ability (Rondeel, 2013). Such a model could help identify talented individuals for leadership and advisor positions, and help design better education and training procedures through a greater understanding of the factors that explain individual differences in complex problem solving ability.

## Acknowledgements

We thank Jean-Denis Latulippe-Thériault and Philippe Carpentier for their valuable support in the preparation of the experimental apparatus and the data collection. Special thanks to Positech Games for the Democracy 2 serious game. This work was funded by the Natural Sciences and Engineering Research Council of Canada (NSERC), and Thales Research & Technology Canada, through a collaborative research and development grant to S. Tremblay. D. Lafond is supported by a postdoctoral industrial R&D fellowship from NSERC, and J-F. Gagnon received support in the form of a doctoral scholarship from the joint program of university and industry collaboration (bourse en milieu pratique) from the Fond de Recherche du Québec - Nature et Technologies and NSERC.

## References

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control 19,* 6, 716-723.

Browne, M. (2000). Cross-validation methods. *Journal of Mathematical Psychology, 44,* 108-132.

Danner, D., Hagemann, D., Holt, D. V., Hager, M., Schankin, A., Wüstenberg, S., et al. (2011). Measuring performance in a complex problem solving task: Reliability and validity of the Tailorshop simulation. *Journal of Individual Differences*, *32,* 225–233.

Dörner, D. (1986). Diagnostik der operativen Intelligenz [Assessment of operative intelligence]. *Diagnostica, 32,* 4, 290-308.

Drucker, H., Burges, C. J. C., Kaufman, L., Smola, A. J., & Vapnik, V. N. (1997). Support Vector Regression Machines. *Advances in Neural Information Processing Systems, 9,* 155-161.

Elg, F. (2005). Leveraging intelligence for high performance in complex dynamic systems requires balanced goal. Theoretical Issues in Ergonomics Science, 6, 63-72.

Fischer, A., Greiff, S., & Funke, J. (2012). The Process of Solving Complex Problems. *Journal of Problem Solving, 4*(1), 19-42.

Gagnon, J.-F., Lafond, D., DuCharme, M. B., St-Louis, M.-E., & Tremblay, S. (2012). Identification of adaptive behaviors and decision heuristics in a simulated strategic decision making task. P*roceedings of the IEEE Conference on Cognitive Methods in Situation Awareness and Decision Support.* New Orleans, LA, March 6-8.

Gonzalez, C., Thomas, R. P., & Vanyukov, P. (2005). The relationships between cognitive ability and dynamic decision making. *Intelligence*, *33*, 169-186.

Goode, N., & Beckmann, J.F. (2010). You need to know: There is a causal relationship between structural knowledge and control performance in complex problem solving tasks. *Intelligence, 38,* 345-352.

Greiff, S., Holt, D. V., Wüstenberg, S., Goldhammer, F., & Funke, J. (2013). Computer-based assessment of Complex Problem Solving: concept, implementation, and application. *Educational Technology Research & Development, 61*, 407-421.

Kluwe, R. H., Shilde, A., Fisher, C., & Oellerer, N. (1991). Problemlöseleistungenbeim umgang mit komplexen systemen und intelligenz (Problem solving performance when interacting with complex systems and intelligence). *Diagnostica, 37,* 291–313.

Lafond, D., & DuCharme, M.B. (2011). Complex decision making experimental platform (CODEM): A counter-insurgency scenario. *Proceedings of the IEEE Symposium on Computational Intelligence for Security and Defence Applications.* Paris, FR, April 11-15, pp. 72-79.

Myung, I. J., Pitt, M. A., & Kim, W. (2005). Model evaluation, testing and selection. In K. Lambert and R. Goldstone (Eds.), *The handbook of cognition* (pp. 422–436). Thousand Oaks, CA: Sage.

Raven, J., Raven, J. C., & Court, J. H. (1998). Manual for Raven's Advanced Progressive Matrices (1998 edition). Oxford, England: Oxford Psychologists Press.

Rigas, G., & Brehmer, B. (1999). Mental processes in intelligence tests and dynamic decision making tasks. London: Lawrence Erlbaum Associates.

Rigas, G., Carling, E., & Brehmer, B. (2002). Reliability and validity of performance measures in microworlds. *Intelligence, 30,* 463-480.

Rondeel, E. W. M, (2013). Cognitive Control in Goal-Directed Decision Making. Doctoral Thesis. Radboud University Nijmegen.

Rouwette, E.A.J.A., Größler, A., & Vennix, J.A.M. (2004). Exploring influencing factors on rationality: A literature review of dynamic decision-making studies in system dynamics. *Systems Research and Behavioral Science, 21,* 351-370.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, *323,* 533-536.

Rüping, S. (2001). Incremental learning with support vector machines. *Proceedings of the 2001 IEEE International Conference on Data Mining,* 641-642.

Wenke, D., Frensch, P. A., & Funke, J. (2005). Complex Problem Solving and intelligence: Empirical relation and causal direction. In R. J. Sternberg & J. E. Pretz (Eds.), *Cognition and intelligence: Identifying the mechanisms of the mind* (pp. 160-187). New York: Cambridge University Press.

Wüstenberg, S., Greiff, S., & Funke, J. (2012). Complex problem solving. More than reasoning? *Intelligence, 40,* 1-14.