

Running head: COMBINING PSYCHOLOGICAL MODELS WITH MACHINE LEARNING

Combining Psychological Models with Machine Learning to Better Predict People's Decisions

Avi Rosenfeld¹, Inon Zukerman², Amos Azaria³, Sarit Kraus^{3,4}

¹Department of Industrial Engineering

Jerusalem College of Technology, Jerusalem, Israel 91160

² Department of Industrial Engineering and Management,

Ariel University Center of Samaria, Ariel, Israel 40700

³Department of Computer Science

Bar-Ilan University, Ramat-Gan, Israel 92500

⁴ Institute for Advanced Computer Studies

University of Maryland, College Park, USA 20742

rosenfa@jct.ac.il, inonzu@ariel.ac.il, azariaa1@mail.biu.ac.il, sarit@cs.biu.ac.il

Abstract

Creating agents that proficiently interact with people is critical for many applications. Towards creating these agents, models are needed that effectively predict people's decisions in a variety of problems. To date, two approaches have been suggested to generally describe people's decision behavior. One approach creates a-priori predictions about people's behavior, either based on theoretical rational behavior or based on psychological models, including bounded rationality. A second type of approach focuses on creating models based exclusively on observations of people's behavior. At the forefront of these types of methods are various machine learning algorithms. This paper explores how these two approaches can be compared and combined in different types of domains. In relatively simple domains, both psychological models and machine learning yield clear prediction models with nearly identical results. In more complex domains, the exact action predicted by psychological models is not even clear, and machine learning models are even less accurate. Nonetheless, we present a novel approach of creating hybrid methods that incorporate features from psychological models in conjunction with machine learning in order to create significantly improved models for predicting people's decisions. To demonstrate these claims, we present an overview of previous and new results, taken from representative domains ranging from a relatively simple optimization problem and complex domains such as negotiation and coordination without communication.

Combining Psychological Models with Machine Learning to Better Predict People's Decisions

Introduction

The challenge of predicting people's decisions is of utmost importance for many economists, psychologists and artificial intelligence researchers (Chalamish, Sarne, & Kraus, 2008; Gigerenzer & Goldstein, 1996; Manisterski, Lin, & Kraus, 2008). Within the field of economics and psychology, validly encapsulating human decision-making is critical for predicting the short and long term effects of a given policy (Neumann & Morgenstern, 1944; Selten, 1998; Kahneman & Tversky, 1979). To computer scientists, accurately predicting people's actions is critical for mixed human-computer systems such as entertainment domains (Maes, 1995), Interactive Tutoring Systems (Murakami, Sugimoto, & Ishida, 2005), adversarial interactions (Cheng, Zuckerman, Nau, & Golbeck, 2011) and mixed human-agent trading environments (Manisterski et al., 2008). Within these and similar domains, creating agents that effectively understand and/or simulate people's logic is particularly important (Chalamish et al., 2008).

To date, two approaches have been proposed for predicting people's decisions by social and behavioral scientists. One classic approach, often advocated by economists, has modeled people's behavior based on classic decision theory. This direction, originally proposed by Von Neumann and Morgenstern (Neumann & Morgenstern, 1944), assumes that people's decisions can be modeled mathematically and rationally based on expected utility. Even when people are faced with uncertainty, these models assume that people will adhere to strict mathematical formulas based on the probability that each event will occur. Game theory follows this approach, and equilibrium strategies such as the Nash equilibrium (Nash, 1951) apply an expected utility to situations where two or more people interact in order to predict their decisions. These solution concepts have proven effective in some applications (Neumann & Morgenstern, 1944; Russell & Norvig, 2003). However, research into people's decisions has shown that people do not necessarily

always adhere to these rigid models (Gigerenzer & Goldstein, 1996; Selten, 1998; Kahneman & Tversky, 1979).

A second class of approaches, often advocated by psychologists and experimental economists, builds cognitive models based on people’s subjective perception of a problem. These approaches posit that theoretical outcomes are less important, and models must instead be constructed based on modeling people’s observed behavior. Examples of this direction include Kahneman and Tversky’s Prospect Theory (Kahneman & Tversky, 1979), which models how people deviate from expected utility when faced with risk, and Gigerenzer and Goldstein’s fast and frugal heuristics (Gigerenzer & Goldstein, 1996), which assume that people use simplistic heuristics to guide their decisions. Models of bounded rationality lie within this group, as they posit that people search for non-optimal alternatives to fulfill their goals. Simon coined the term “satisfies” to capture that bounded decision makers seek solutions that are “good enough” and not optimal ones (Simon, 1957). We considered one such theory, Selten’s Aspiration Adaptation Theory (Selten, 1998), whereby people make decisions by attempting to satisfy only one goal variable at a time, called their “aspiration”.

In contrast to these cognitive models, computer scientists often model peoples’ decisions through machine learning techniques (Mitchell, 1997). These models are based on statistical methods such as Bayes’ Rule, Neural Networks or Decision Tree algorithms, and are exclusively built on observed decisions instead of general predictions about peoples’ behavior. As a result, these models do not make any claims for their general applicability as they were created exclusively based on observations in a specific setting.

The key contribution of this paper is an exploration of how one can combine the **cognitive decision making approaches** proposed by social scientists with **machine learning approaches**. At times, and as we demonstrate in this paper, sufficient data exists to confirm various behavior models. In these types of cases, decision tree models can be constructed to demonstrate the typical decision model of people within a given domain. This result is particularly significant if several behavior models are theoretically possible—allowing us to

confirm without bias which behavior model is best.

Other times, multiple cognitive models are available, but insufficient data exists to allow machine learning to form a complete cognitive model on its own. However, in these types of domains, sufficient data does exist to help fine-tune parameters within a relevant cognitive model. Often this process can be completely automated with recognized mathematical packages or through accepted machine learning algorithms (Azaria, Rabinovich, Kraus, Goldman, & Tsimhoni, 2012; Ratcliff & Smith, 2004; Nilsson, Rieskamp, & Wagenmakers, 2011). For example, within computer science Azaria et. al used machine learning to fit parameters within a path selection model. Within psychology and behavioral economics, Ratcliff and Smith used the Simplex algorithm to fit parameters within a 2-choice decision model (Ratcliff & Smith, 2004) and Nilsson et. al contrasted how Bayesian and maximum likelihood approaches could be used to set parameters within prospect theory (Nilsson et al., 2011).

However, these approaches (Azaria et al., 2012; Ratcliff & Smith, 2004; Nilsson et al., 2011) are not appropriate for domains that are simply too complex to understand which behavior model should be applied or how parameters within a given model can be tuned. These domains suffer from a phenomenon often known as the “curse of dimensionality” (Evangelista, Embrechts, & Szymanski, 2006; Bengio, Delalleau, & Roux, 2005). The curse of dimensionality is, unfortunately, particularly evident in many real-world situations, as people can potentially act within a very large set of actions (a high dimension of possible actions), and they often do not consistently choose the same actions. Even with sophisticated mathematical algorithms or through using accepted machine learning procedures, people’s behavior within these domains cannot be easily quantified because the possible range of actions that a person might take is too large and the data given to learn a person’s behavior is insufficient given this level of complexity. Particularly in these types of domains, novel approaches are necessary to accurately predict people’s behavior.

To address this challenge, and a key contribution of this paper, is an exploration of the idea that one can **combine** the cognitive decision-making approaches proposed by social scientists with classic machine learning approaches, as a way of creating improved prediction models of

people's behavior. Specifically, we demonstrate that there are cases in which **hybrid models** that combine machine learning with a cognitive theory can be more successful than any of its individual parts. Moreover, we found that the advantage of using a hybrid model increases with the complexity of the underlying prediction task.

To demonstrate the results, we present an overview of previous and new results considering two different psychological models: Aspiration Adaptation Theory (Selten, 1998) and Focal Points Theory (Schelling, 1963). Specifically, in the second section we show that the Aspiration Adaptation Theory proved to be the best predictor in a relatively simple optimization problem, and that machine learning methods demonstrated without bias that the AAT model was most appropriate, even though alternate cognitive models were possible. We also demonstrate that adding AAT features into a machine learning model significantly increased the prediction accuracy of people's bids within a complex negotiation domain. In the third section we present the Focal Points Theory that describes a low-level cognitive ability to pick prominent solutions in the absence of communication. We found that adding information from this model significantly increased the accuracy of a machine learning based prediction model in problems where people had to coordinate without communication.

Aspiration Adaptation Theory

Aspiration Adaptation Theory (AAT) was proposed by Selten as a general economic model that describes how people make certain economic decisions without any need for expected utility (Selten, 1998). AAT was originally formulated to model how people make decisions where utility functions cannot be constructed. For example, assume you need to relocate and choose a new house in which to live. There are many factors that you need to consider, such as the price of each possible house, the distance from your work, the neighborhood and neighbors and the schools in the area. How do you decide which house to buy? While in theory utility based models could be used, many of us do not create rigid formulas involving numerical values to weigh trade-offs between each of these search parameters.

AAT provides an alternative to utility theory for how decisions can be made in this and

other problems. First, m goal variables are sorted in order of priority, or their *urgency*. Accordingly, the order of G_1, \dots, G_m refers to goal urgency, or the priority by which a solution for the goal variables is attempted. Each of the goal variables has a desired value, or its *aspiration level*, that the agent sets for the current period. This desired value is not necessarily the optimal one, and the agent may consider the variable “solved” even if it finds a sub-optimal yet sufficiently desired value. The agent’s search starts with an initial aspiration level and is governed by its *local procedural preferences*. The local procedural preferences prescribe which aspiration level is most urgently adapted upward if possible, the second most desired aspiration, etc. and which partial aspiration level is *retreated from* or adapted downward if the current aspiration level is not feasible. Here, all variables except for the goal variable being addressed are assigned values based on *ceteris paribus* (given all other goals being equal, a better value is preferred to a worse one).

We studied what decision models, AAT or others, were used to solve two types of problems – a relatively simple optimization problem and a complex negotiation problem. In the first optimization problem, we consider a problem where a person must minimize the price in buying a commodity (a television) given the following constraints: Assume a person must personally visit stores in order to observe the posted price of the commodity. However, some cost exists from visiting additional stores. For any given discrete time period, the person must decide if she wishes to terminate the search without incurring an additional cost. Full details of our implementation can be found in our previously published work (Rosenfeld & Kraus, 2009, 2011).

In addition to AAT, other strategies, bounded and strictly rational, were possible here. A clear optimal strategy existed within the implementation of the commodity search domain. In the settings with which we experimented, the specific strategy was to buy the commodity if the price in the current store is less than 789. Thus, classical expected utility theory would predict that people would similarly buy the commodity at this price. We also recognize that AAT is not the only possible option of a bounded model within this domain. Following Gigerenzer and Goldstein’s fast and frugal heuristics (Gigerenzer & Goldstein, 1996), we would expect people to formulate simple strategies involving only one variable (e.g. search until price $< X$, or visit Y

stores and buy in the cheapest store). However, using an AAT-based model for prediction would assume that some type of combination strategy exists where one variable is first searched for but then retreated from, assuming that value could not be satisfied. For example, a person might initially search for a price less than 650, but will settle on an even a higher price (e.g. the lowest found so far) after unsuccessfully finding this price after searching 5 stores. In fact, our previous work did find that people typically used these AAT strategies instead of optimal or fast and frugal heuristics (Rosenfeld & Kraus, 2009).

We also analyzed a previously presented negotiation domain (Lin, Kraus, Wilkenfeld, & Barry, 2008). We consider a negotiation session that takes place after a successful job interview between an employer and a job candidate. In this session both sides wish to formalize the hiring terms and conditions of the applicant: her **Salary**, **Job Description**, **Car Benefits**, **Pension benefits** and **Working hours**. In the problem setting considered, each side could pick from a list of possible values for each of the parameters. For example, the employee might ask for a salary of 20,000 per month, with the job title of Project Manager, with a car, pension benefits and an 8-hour work day, while the employer might counter with the same offer, but a salary of only 12,000 per month and without the pension benefits. The goal of this study is to accurately predict what each side would offer. Here again, equilibrium strategies were possible based on strictly rational behavior. Following Gigerenzer and Goldstein's model of fast and frugal heuristics, we would have expected that simple compromise heuristics could be used. Possibilities of such heuristics include always countering the middle position between the previous offer of both sides or offering the middle position between all previous offers of both sides. Nonetheless, our overall finding was that people create aspiration-based strategies where they negotiate for specific issues in a specific order. For example, we found that negotiations first focused on the salary parameter and only then moved on to other parameters such as pension or car benefits. We found that adding these aspirations explicitly as a parameter for the machine learning models to consider helped to significantly improve the accuracy in predicting people's offers.

Focal Points Theory

Focal Points were introduced by Schelling (Schelling, 1963) as a prominent subset of solutions for *tacit coordination games*, which are coordination games where communication is not possible. In such games (also known as *matching games* in game theory terminology) the players only have to **agree** on a possible solution, regardless of the solution itself. In other words, they receive a reward for selecting the same solution, regardless of the solution. When their solutions differ, both players lose and do not get any reward. A solution is said to be “focal” (also “salient”, or “prominent”) when, despite similarity among many solutions, the players somehow converge to this solution.

A classic example of Focal Points coordination is the solution most people choose when asked to divide \$100 into two piles, of any size; they should attempt only to match the unseen player’s choice. More than 75% of the subjects in Schelling’s experiments created two piles of \$50 each; that solution is what Schelling dubbed a Focal Point. Here again, other behavioral models are possible – using equilibrium strategies would result in a random selection among the 101 possible divisions, as the probability distribution is uniform.

Several attempts have been made to formalize Focal Points from a game theoretic, human interaction point of view ((Janssen, 1998) provides a good overview). However, that research does not provide the practical tools necessary for predicting people’s actions. In a meta-analysis of previous Focal Points experiments we developed some general properties that “focalize” an answer: (1) Centrality, (2) Extremeness, (3) Firstness and (4) Singularity. For further details and examples of these rules, we encourage the reader to refer to our previous work (Zuckerman, Kraus, & Rosenschein, 2011).

The task of learning which of these properties will be used by people is far from trivial due to the large number of possibilities, resulting in the “curse of dimensionality”. To overcome this difficulty, we present a *Focal Points Learning* approach which combines this psychological approach and machine learning. This approach is created by preprocessing raw domain data and then placing it into a new representation space based on the Focal Points properties. Specifically,

we take every game instance of the domain encoded in the most naive raw encoding (called the original feature space) and we apply a transformation to a new feature space that includes Focal Points information.

We designed a simple and intuitive tacit coordination game that represents a simplified version of a domain where an agent and a human partner need to agree on a possible meeting place. The game, coined “Pick the Pile” is played on a 5-by-5 square grid. Each square of the grid can either be empty or can contain either a pile of money or the game agents. Each square on the game board is white, yellow or red. The players were instructed to pick the pile of money from the three identical piles which most other players, playing exactly the same game, would pick.

Experimental Results

In this section we present an overview of previous and new results that demonstrate when and how machine learning techniques can benefit from behavioral theories. In general, we found that in the relatively simple optimization problem, strictly rational, AAT models and machine learning converged on nearly identical results. In the more complicated negotiation domain, adding information about people’s aspirations increased the predictive accuracy of models built based upon machine learning. Strictly rational models performed far worse. In an even more complex coordination without communication domain, Focal Points information again improved the accuracy of a model based upon machine learning models. Strictly rational models and models built upon Focal Points without machine learning performed far worse.

Results from an Optimization Problem

In the first task, a relatively simple optimization problem, we wished to predict if a person would stop their commodity search in any given store. In this domain an optimal search strategy exists, namely that in the specific settings that we considered, the person should stop the search in the first store with a price less than 789. Note that this solution can be mathematically calculated and does not require any input from observed behavior. At the other extreme, we can create a prediction model based exclusively on machine learning techniques. Previously we used

decision trees to create this model. The advantage to specifically using this type of model lies in the output – we can check if the decision tree’s decision model is consistent with the optimal solution or with other bounded models. We considered two such bounded models: simple heuristics and AAT. Based on the fast and frugal approach, we would expect people to use a simple decision-making process. Specifically, we assume that they would stop their search based on only one parameter, such as the number of stores visited to date, or the price of the commodity in any given store. This could be considered a classic example of the fast and frugal take-the-best heuristic (Gigerenzer & Goldstein, 1996). According to AAT we would expect to see more complicated strategies with multiple parameters and some type of ordering and retreat between them. Our previous work (Rosenfeld & Kraus, 2011) did in fact find that the decision trees’ output was consistent with AAT strategies as people would typically immediately buy the commodity if its price was below a certain threshold, but would settle on a higher price after visiting a certain number of stores.

In this paper, we focus on when and how we can combine various decision theories in order to better predict people’s decisions. In this domain, this included comparing the following models:

1. An optimal model based on expected utility – e.g. people buy only if the price is less than 789.
2. A machine learning model based on observed decisions.
3. A combination model. In this problem, the combination model involved adding information about the average price where people stopped their search, and the average number of stores after which they were willing to settle on a more expensive commodity. Note that here, as well as in all of the domains we consider, this hybrid approach assumes that we have some general information about a given population.

For this domain, we found that adding general information about people’s aspirations was useful, but only slightly. Table 1 presents the accuracy of different models in predicting when 41 people stopped their commodity search. Each of these people was presented with a simulation of the commodity search domain and ran at least 25 trials each ending in their buying the commodity, logging a total of nearly 5000 instances where these people either decided to buy the

commodity or to continue their search. The first column of Table 1 presents a baseline *Naive* model that classifies all decisions based on the majority class, here assuming that people will always continue the search. In the second column, we present the predictive ability of the optimal model. Column 3 presents the results from the machine learning method which performed similarly to the optimal model with 82.67% accuracy. Adding information from people's aspirations did help, but only slightly, with 83.45% accuracy achieved through knowing the average values of these people's aspirations. Note that this value serves as an upper baseline, as we collected this aspiration data from the same population being evaluated. A more realistic aspiration model is the Sparse AAT model which used only 50 randomly selected decisions to help model peoples' decisions (or less than 1% of the total logged data). Nonetheless, even this model did slightly outperform both the optimal and based machine learning methods with 83% accuracy. This result is even more striking when you consider that machine learning models were validated through cross-validation of 90% of the data used for training the model, while this sparse model used less than 1% of the data. Thus, we conclude that in this relatively basic domain, differences between the predictive abilities of the different models was not large. Nonetheless, a slight improvement in prediction accuracy was obtained through limited information about people's aspirations.

AAT in a Negotiation Domain

According to AAT, one would expect people to rank the importance of each of the negotiation parameters according to his or her individual aspiration scale. Assuming people often have the same aspiration scales, we would also see an order where issues are addressed, e.g. certain parameters are typically negotiated first, second, etc. Our premise is that as the negotiation domain is more complex than the optimization problem, one should add people's aspiration information into traditional models such as a decision tree model to more accurately predict what bids people will offer.

To test this hypothesis, we proceeded to study what gain, if any, did adding AAT information have in predicting how people will negotiate. In the problem we considered, the

parameters to be negotiated could have between 2 and 4 discrete values. In order to study this point we considered several models for the negotiation problem (see Table 3). The goal of all of these models was to predict the next value for each parameter. First, we considered the **Majority Rule** model. Given the full log file, this rule assumes that a person would offer the most popular value for any given parameter. For example, in the employer / employee domain, the most popular **title** was “Programmer”. Second, we implemented two models based on the **equilibrium strategy**. These strategies are based on previous work in these problems (Lin et al., 2008). However, as the equilibrium strategy depends on which person is allowed to offer the last bid, we checked both in order to know which equilibrium strategies would best predict all parameters. Next, we created a baseline strategy that uses a decision tree (**D.T.**) algorithm (C4.5) to predict the next offer for each parameter. This model used historical information about the previous offer and the current negotiation iteration. Next, we created a **D.T. with AAT statistical information** prediction model. As we previously demonstrated, each parameter had different urgencies. Thus, we attempted to create a more accurate model by adding information about which parameters were typically raised or lowered for any given iteration. Specifically, we added a field with a binary flag value in order to differentiate between the iterations for which people typically changed a given parameter’s value with a frequency of ≥ 0.5 , and those which were typically not changed and where additional information likely would not help. This was done to avoid over-fitting the AAT statistics for any training / testing pair, and thus to retain the generality of the results. Finally, we created a **D.T. + Complete Behavior Knowledge** model. This final baseline had knowledge about what the previous offer was, and also added perfect knowledge if the person would revise upwards, downwards, or leave unchanged their previous offer. In cases where only two options exist, one would expect this baseline to guarantee 100% accuracy. However, when more than 3 values exist for a given parameter, even this model cannot guarantee 100% accuracy. For example, if a previous salary offer was \$7,000 per month and we know that the next offer will be higher, we still do not know if it will be raised to \$12,000 or \$20,000. Nonetheless, the goal of this model was to provide an upper bound for how much

AAT-based information could theoretically help.

Table 3 demonstrates the effectiveness of adding AAT information to boost prediction accuracy. The first row of this table show the parameter to be negotiated and the number of possible values. The second row presents the majority rule baseline. The third and fourth rows present how effective the equilibrium policies were in predicting what people actually offered. Note that both of these policies fall well below the naive majority baseline. This again demonstrates the ineffectiveness of using equilibrium theoretical policies to predict how people actually behave. The fifth row presents the accuracy of the learned decision tree model. This model represents the effectiveness of this traditional learning method in predicting each of the parameters. We then added AAT information and reran the same decision tree algorithm, the results of which are in the sixth row. Note that the significant improvement gained from the AAT information is significant and only one parameter did not gain from the added aspiration information. In this parameter, few instances existed where people had clear general aspiration changes, preventing any accuracy boost from this approach. Finally, the last line in the table presents the accuracy of the decision tree algorithm with complete behavior knowledge, or perfect information about whether a person will retreat from (decrease) a given parameter value or upwardly revise its aspiration (increase). Note that, as expected, even complete AAT information could not yield 100% prediction accuracy for parameters with more than 2 values.

Experimental Results for Focal Points in the Pick the Pile domain

In order to evaluate the effectiveness in using Focal Points information to predict people’s actions, we collected data using an Internet website which logged approximately 3,000 game instances. We then compared the correct classification performance of both decision trees (again the C4.5 algorithm) and FFBP neural network classifiers. The comparison was between a **domain data agent** — an agent that was trained only on the raw domain encoding (original domain without any explicit Focal Points information), a **Focal Points agent** (FP) — an untrained agent that used only the Focal Points rules for prediction, weighted uniformly (without running any machine learning algorithm), and a **Focal Points learning agent** (FPL) — an agent that

was trained on the data following the preprocessing stage in which the Focal Points rules were injected explicitly. “Correct classification” means that the agent made the same choice as that of the particular human player who played the same game. Obviously the learning problem is extremely difficult as there is no simple function that can capture the notion that for some games, different human players can select different choices. We optimized our classifiers’ performance by varying the network architecture and learning parameters, until we attained the best results. We used a learning rate of 0.3, momentum rate of 0.2, 1 hidden layer, random initial weights and no biases of any sort. Before each training procedure, the data set was randomly divided into a test and a training set. Each instance of those sets contained the game description (either the binary or Focal Points encoding) and the human answer to it. The classification results using the neural network and the decision tree algorithms were very close (maximum difference of 3%).

Examining the results in Table 4, we see a significant improvement when using the Focal Points learning approach to train classifiers, rather than the domain data agent ($p < 0.01$ in two-proportion z-tests in all domains). In this domain, the domain data agent is not able to generalize sufficiently, thus achieving classification rates that are only about 5%–10% higher than a random guess (which is 33%). Using FPL, the classification rate improved to more than 65% correct classification. Since even humans do not have 100% success with one another in these games, FPL is correspondingly the more impressive. The results also show that even the classical FP agent, which does not employ any learning algorithm, performs better than the domain data agent, with 48% correct classification. In an additional analysis that was done on the FP agent, we saw a tendency in which, when facing coordination problems with low *focalilty difference*, the FP agent’s performance deteriorates to that of random guesses.

An additional advantage of using FPL is the reduction in training time (e.g., in the *Pick the Pile* domain we saw a reduction from 4 hours on the original data to 3 minutes), due to the reduction of input size. Moreover, the learning tree that was created using FPL was smaller and can be easily converted to a rule-based system as part of the agent’s design.

Discussion and Conclusion

Predicting people’s decisions is an important but complex task. To address this task, researchers often propose general behavior models such as rationality theory or purely statistical methods such as machine learning algorithms. However, there often exist specialized cognitive models or theories that describe various tendencies or biases that are commonly used by the majority of the people. Such theories include bounded rationality theories, various risk attitudes and use of heuristics.

This paper addresses how one can take a potentially relevant cognitive theory and use machine learning methods to help augment it in order to provide added value in predicting human behavior. We showed how two cognitive theories, Aspiration Adaptation theory and the Focal Points theory, could be used in conjunction with machine learning algorithms to create an improved classifier. Possibly equally significant is the result that strictly rational models, and even many specialized cognitive models, often do not accurately predict people’s decisions.

Our results also show some positive correlation between the complexity of the problem domain and the improvement in performance when augmenting the cognitive model. To demonstrate this phenomenon, we present our results, which are found in Table 4. In relatively simple domains, such as the optimization problem we considered, machine learning methods can accurately predict people’s behavior given enough training data. In these types of problems, cognitive models will not help improve the prediction accuracy but they can help predict people’s behavior with less training data. Additionally, we found that using machine learning models allowed us to confirm the accuracy of a given cognitive model without human bias, here AAT, even though other cognitive theories were potentially applicable. In complex problems, the number of parameters that need to be set within the cognitive model is large and cannot be readily identified even with large training data sets. Because of the range of possible actions, these problems suffer from the “curse of dimensionality” where even sophisticated machine learning approaches and parameter tuning alone will not suffice (Evangelista et al., 2006; Bengio et al., 2005). In these problems, the novel hybrid approach we present— using machine learning

algorithms as the base of the solution, but adding features from the cognitive models – creates significant improvements over both base approaches, often by large amounts.

We believe that the hybrid approaches presented in this paper are successful for the following reason. At the root of the curse of dimensionality is that the space of possible actions is too large for mathematical approaches to find an optimal value. Thus, any information that will help reduce the range of possible values will lead to improved models if this allows the curse of dimensionality to be removed, allowing machine learning models, or other mathematical approaches, to produce an accurate model from the data. We believe that this also explains why hybrid models become increasingly important as the complexity of the domain rises. Increased problem complexity contains higher dimensionality. As problem complexity rises, we found that traditional approaches to model people's behavior became increasingly less successful in creating an accurate model. However, once our hybrid models introduce information about how people behave, and just as importantly, how they won't behave, they reduce the space in which machine learning algorithms must search for an accurate behavior model. This in turn transforms the modeling problem into one the machine learning algorithms could solve, resulting in significantly more accurate predictions.

As we present a generalized approach for how to combine cognitive theories with machine learning algorithms, we expect this approach to be generally applicable to a variety of new domains as well. We have already discovered new examples from the representative categories within this paper. For example, we found that using a previously developed driving model (Fancher & Bareket, 1996) was useful in predicting drivers' preferred settings within their adaptive cruise control (Rosenfeld et al., 2012). In this domain, sufficient data existed to learn the drivers' cognitive model, similarly paralleling the simple optimization problem within this study. At the other extreme, we have also found additional examples in complex problems where hybrid models with machine learning and cognitive models have yielded more accurate prediction models than those based on cognitive models or machine learning alone. Specifically, we also found that a hybrid Focal Points / machine learning model was more successful in predicting people's

preferences in other coordination settings including shape selection and job candidate selection (Zuckerman et al., 2011). Other researchers have also used machine learning to infer how social preferences affect different types of people in human-computer negotiation (Gal, Grosz, Kraus, Pfeffer, & Shieber, 2010). Our hope is that additional researchers will also be successful in creating many additional applications using the hybrid cognitive model / machine learning approach described in this paper.

Acknowledgements

This work was supported in part by ERC grant #267523, MURI grant number W911NF-08-1-0144, ARO grants W911NF0910206 and W911NF1110344 and JCT grant 963.

References

- Azaria, A., Rabinovich, Z., Kraus, S., Goldman, C. V., & Tsimhoni, O. (2012). Giving advice to people in path selection problems. In *Aamas*.
- Bengio, Y., Delalleau, O., & Roux, N. L. (2005). *The curse of dimensionality for local kernel machines* (Tech. Rep.).
- Chalamish, M., Sarne, D., & Kraus, S. (2008). Programming agents as a means of capturing self-strategy. In *Aamas '08* (pp. 1161–1168).
- Cheng, K.-L., Zuckerman, I., Nau, D. S., & Golbeck, J. (2011). The life game: Cognitive strategies for repeated stochastic games. In *Socialcom/passat* (p. 95-102).
- Evangelista, P. F., Embrechts, M. J., & Szymanski, B. K. (2006). Taming the curse of dimensionality in kernels and novelty detection. In *In ajith abraham, bernard de baets, mario kppen, and bertam nickolay, editors, applied soft computing technologies: The challenge of complexity* (pp. 431–444). Springer Verlag.
- Fancher, P., & Bareket, Z. (1996). A comparison of manual versus automatic control of headway as a function of driver characteristics. *3rd Annual World Congress on Intelligent Transport Systems*.
- Gal, Y., Grosz, B., Kraus, S., Pfeffer, A., & Shieber, S. (2010). Agent decision-making in open mixed networks. *Artificial Intelligence*, 174(18), 1460–1480.
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: models of bounded rationality. *Psychology Review*, 103(4), 650–669.
- Janssen, M. C. W. (1998). Focal points. In P. Newman (Ed.), *The new palgrave of economics and the law* (pp. 150–155). MacMillan.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47, 263–291.
- Lin, R., Kraus, S., Wilkenfeld, J., & Barry, J. (2008). Negotiating with bounded rational agents in environments with incomplete information using an automated agent. *Artificial Intelligence*, 172(6-7), 823-851.

- Maes, P. (1995). Artificial life meets entertainment: lifelike autonomous agents. *Commun. ACM*, 38(11), 108–114.
- Manisterski, E., Lin, R., & Kraus, S. (2008). Understanding how people design trading agents over time. In *Aamas '08* (pp. 1593–1596).
- Mitchell, T. M. (1997). *Machine learning*. New York: McGraw-Hill.
- Murakami, Y., Sugimoto, Y., & Ishida, T. (2005). Modeling human behavior for virtual training systems. In *Aaai* (p. 127-132).
- Nash, J. (1951, September). Non-Cooperative Games. *The Annals of Mathematics*, 54(2), 286–295.
- Neumann, J. V., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton University Press.
- Nilsson, H., Rieskamp, J., & Wagenmakers, E.-J. (2011). Hierarchical bayesian parameter estimation for cumulative prospect theory. *Journal of mathematical psychology (Print)*, 55(1), 84–93.
- Ratcliff, R., & Smith, P. L. (2004, April). A comparison of sequential sampling models for two-choice reaction time. *Psychol Rev*, 111(2), 333–367.
- Rosenfeld, A., Bareket, Z., Goldman, C. V., Kraus, S., LeBlanc, D. J., & Tshimoni, O. (2012). Learning drivers behavior to improve the acceptance of adaptive cruise control. In *Innovative applications of artificial intelligence*. AAAI.
- Rosenfeld, A., & Kraus, S. (2009). Modeling agents through bounded rationality theories. In *Ijcai* (p. 264-271).
- Rosenfeld, A., & Kraus, S. (2011). Using aspiration adaptation theory to improve learning. In *Aamas* (p. 423-430).
- Russell, S. J., & Norvig, P. (2003). *Artificial intelligence: A modern approach*. Prentice Hall.
- Schelling, T. (1963). *The strategy of conflict*. New York: Oxford University Press.
- Selten, R. (1998). Aspiration adaptation theory. *Journal of Mathematical Psychology*, 42, 1910–214.

Simon, H. A. (1957). *Models of man*. New York: John Wiley & Sons.

Zuckerman, I., Kraus, S., & Rosenschein, J. S. (2011). Using focal point learning to improve human-machine tacit coordination. *Autonomous Agents and Multi-Agent Systems*, 22(2), 289-316.

Table 1

Comparing the Prediction Accuracy between Optimal, Machine Learning and AAT Based Models

Naive	Optimal	Learning	Learning + Complete AAT	Sparse AAT
78.56	82.8	82.67	83.45	83

Table 2

Comparing the Prediction Accuracy between AAT and non-AAT Based Models in the Employer / Employee Negotiation Domain

	Salary-3	Title-4	Car-2	Pension-3	Promotion-2	Hours-3	Average
Majority Rule	60.19	67.59	57.41	70.37	62.96	62.96	63.58
Equilibrium 1	44.44	67.59	69.44	66.67	41.67	67.59	59.57
Equilibrium 2	25.93	17.59	69.44	19.44	43.51	61.11	39.51
D.T. Without AAT	61.11	68.52	68.52	67.59	83.33	69.44	69.75
D.T. with AAT	62.96	68.52	75.93	71.30	91.67	76.85	74.54
D.T. + Complete	95.37	89.81	100	96.28	100	96.30	96.30

Table 3

Results from “Pick the Pile” domain

Random guess	Raw Encoding	Only Focal Points Rules	Focal Points Learning
33%	40%	48%	65%

Table 4

Results Summary

Problem Type	Cognitive Model	Hybrid Type	Prediction Improvement
Optimization	AAT	Parameter Learning	Slight / Sparse Improvement
Negotiation	AAT	Cognitive Features	5%
Pick the Pile	Focal Points	Cognitive Features	17%