Optimally Designing Games for Cognitive Science Research

Anna N. Rafferty (rafferty@cs.berkeley.edu)
Matei Zaharia (matei@cs.berkeley.edu)
Computer Science Division, University of California, Berkeley, CA 94720 USA

Thomas L. Griffiths (tom_griffiths@berkeley.edu)
Department of Psychology, University of California, Berkeley, CA 94720 USA

Abstract

Collecting cognitive science data using games has the potential to be a powerful tool for recruiting participants and increasing their motivation. However, designing games that provide useful data is a difficult task that often requires significant trial and error. In this work, we consider how to apply ideas from optimal experiment design to designing games for cognitive science experiments. We use Markov decision processes to model players' actions within a game, and then make inferences about the parameters of a cognitive model from these actions. We present a general framework for finding games with high expected information gain based on this approach. We apply this framework to Boolean concept learning, inferring the difficulty of Boolean concepts from participants' behavior. We show that using games with higher expected information gain allows us to make this inference more efficiently.

Keywords: optimal experiment design; Markov decision process; computer games; concept learning

Introduction

Computer games have become increasingly popular tools for gathering psychological data and for educational purposes (e.g., Michael & Chen, 2005; Von Ahn, 2006; Siorpaes & Hepp, 2008; Klopfer, 2008), providing a way to recruit large numbers of motivated participants. However, creating games that actually result in useful data requires significant engineering, normally based on trial and error. In this paper we propose a method for automating the process of designing games, using ideas from optimal experiment design.

The key problem in designing games for cognitive science research is finding a game that provides as much information as possible about the research question being addressed. For traditional experiments, the field of optimal experiment design seeks to choose the design that will give the most information about the dependent variable (Atkinson, Donev, & Tobias, 2007). We adapt this method to identify the game that will give the most information about the parameters of a cognitive model. By automating the process of game design, we limit the trial and error necessary to find a game that will provide useful data, while still reaping the benefits of using games rather than traditional experiments.

Adapting optimal experiment design methods to game design requires predicting people's behavior within games, which may differ from behavior in traditional experiments. For instance, in a categorization experiment a participant's response to a stimulus may roughly correspond to whether she believes the stimulus is in the category. However, in a game this relationship is complicated by competing incentives. For example, available actions in a game may be contingent on the success of previous actions, leading to complex strategies. We propose using Markov decision processes (MDPs) to predict people's actions in games. MDPs incorporate the current and future benefit of an action, and thus allow us to take into account the incentive structures and rules of a particular game. By combining MDPs with ideas from optimal experiment design, we create a framework for finding the game that will provide the highest expected information gain.

The plan of the paper is as follows. The next section provides background on optimal experiment design and MDPs. We then show how to combine these ideas in a framework for optimal game design. The remainder of the paper applies this general framework to the specific case of learning Boolean concepts, illustrating the benefits of optimal game design. We introduce a novel concept learning game, and use our approach to optimize the game parameters. Two behavioral experiments show that our optimized game results in more efficient estimation of the difficulty of learning different kinds of Boolean concepts, and that the actual amount of information obtained from players is positively correlated with the expected information gain of the game.

Background

The optimal game design framework we propose relies on ideas from Bayesian experiment design and Markov decision processes, which we will introduce in turn.

Bayesian Experiment Design

Bayesian experiment design, a subfield of optimal experiment design, seeks to choose the experiment that will maximize the expected information gain about a parameter θ (Atkinson et al., 2007; Chaloner & Verdinelli, 1995). In cognitive science, this procedure and its variations have been used to design more informative experiments that allow for clearer discrimination between alternative hypotheses (Myung & Pitt, 2009). Throughout this paper, let ξ be an experiment (or game) design and y be the data collected in the experiment. Then the Bayesian experimental design procedure is as follows:

$$\begin{aligned} \text{maximize} \quad U(\xi) &= \int p(y|\xi)U(y,\xi)\mathrm{d}y \\ \text{where} \quad p(y|\xi) &= \int p(y|\xi,\theta)p(\theta)\mathrm{d}\theta \\ \text{and} \quad U(y,\xi) &= \int \left(H(p(\theta|y,\xi)) - H(p(\theta))\right)\mathrm{d}\theta, \end{aligned} \tag{1}$$

where $H(p)$ is the Shannon entropy of a probability distribution p, defined as $H(p) = \int p(x)\log(p(x))\mathrm{d}x$. Thus, the

procedure maximizes the expected utility of an experiment ξ, defined as the information gain over all outcomes y weighted by their probabilities $p(y|\xi)$ under the current prior.

Markov Decision Processes

The Bayesian experiment design procedure uses $p(\theta|y,\xi)$ to calculate the information gain from an experiment. This quantity represents the impact that the data y collected from experiment ξ have on the parameter θ. In a game, the data y are a series of actions, and to calculate $p(\theta|y,\xi)$, we must interpret how θ affects those actions. Via Bayes' rule, we know $p(\theta|y,\xi) \propto p(y|\theta,\xi)p(\theta)$. We thus want to calculate $p(y|\theta,\xi)$, the probability of taking actions y given a particular value for θ and a game ξ. To do so, we turn to Markov decision processes (MDPs), which provide a natural way to model sequential actions. MDPs and reinforcement learning have been used previously in game design for predicting player actions and adapting game difficulty (Erev & Roth, 1998; Andrade, Ramalho, Santana, & Corruble, 2005; Tan & Cheng, 2009).

MDPs describe the relation between an agent's actions and the state of the world and provide a framework for defining the value of taking one action versus another (see Sutton & Barto, 1998). Formally, an MDP is a tuple $\langle S,A,T,R,\gamma\rangle$, where S is the set of possible states and A is the set of actions that the agent may take. The transition model T gives the probability $p(s'|s,a)$ that the state will change to s' given that the current state is s and the agent takes action a. The reward model $R(s,a,s')$ describes the probability of receiving a reward $r \in \mathbb{R}$ given that action a is taken in state s and the resulting state is s'. Finally, the discount factor γ represents the relative value of immediate versus future rewards. The value of taking action a in state s is defined as the expected sum of discounted rewards and is known as the Q-value:

$$Q(s,a) = \sum_{s'} p(s'|s,a)\left(R(s,a,s') + \gamma \sum_{a' \in A} p(a'|s')Q(s',a')\right), \tag{2}$$

where $p(a'|s')$ is the probability that an agent will take action a' in state s' and is defined by the agent's policy π. We assume that people's actions can be modeled as a Boltzmann policy, as in Baker, Saxe, and Tenenbaum (2009):

$$p(a|s) \propto \exp(\beta Q(s,a)), \tag{3}$$

where higher values of β mean the agent is more likely to choose the best action, while $\beta = 0$ results in random actions.

Optimal Game Design

We can now define a procedure for optimal game design, identifying the game with maximum expected information gain for θ. We assume there is an existing game design with parameters to adjust, corresponding to point values, locations of items, or any other factor that can be varied. To apply Bayesian experiment design to choosing a game, we define the utility of a game ξ as the expectation of information gain over the true value of θ and the actions chosen by the players:

$$U(\xi) = E_{p(\theta,\mathbf{a})}[H(p(\theta)) - H(p(\theta|\mathbf{a},\xi))], \tag{4}$$

where $\mathbf{a}$ is the set of action vectors for all players. The expectation is approximated by sampling θ from the prior $p(\theta)$, and then simulating players' actions given θ by calculating the Q-values for the MDP and sampling from Equation 3.

The remaining quantity in Equation 4 is $p(\theta|\mathbf{a},\xi)$. Intuitively, this quantity connects actions taken in the game with the parameter of the cognitive model that we seek to infer, θ. For a game to yield useful information, it must be the case that people will take different actions for different values of θ. Concretely, we expect that players' beliefs about the reward model and the transition model may differ based on θ. For instance, in a categorization task with two objects A and B, θ might determine the probability that A is a positive instance and B is a negative instance of the category. If taking a particular action leads to positive rewards only when a positive instance is observed, then we would expect that the value of θ is large if many players take that action when observing A.

The process of inferring θ from actions assumes that each θ corresponds to a particular MDP. If this is the case, we can calculate a distribution over values of θ based on the observed sequences of actions $\mathbf{a}$ of all players in the game ξ:

$$\begin{aligned} p(\theta|\mathbf{a},\xi) &\propto p(\theta)p(\mathbf{a}|\theta,\xi) && (5)\\ &= p(\theta)p(\mathbf{a}|\mathrm{MDP}_\theta,\xi) && (6)\\ &= p(\theta)\prod_i p(\mathbf{a}_i|\mathrm{MDP}_\theta,\xi), && (7) \end{aligned}$$

where $\mathbf{a}_i$ is the vector of actions taken by player i and MDP_θ is the MDP derived for the game based on the parameter θ. Calculating this distribution can be done exactly if there is a fixed set of possible θ or by using Markov chain Monte Carlo (MCMC) methods if the set of θ is large or infinite (see Gilks, Richardson, & Spiegelhalter, 1996).

Now that we have defined $p(\theta|\mathbf{a},\xi)$, we can use this to find the utility of a game. Equation 4 shows that this calculation follows simply if we can calculate the entropy of the inferred distribution. In the case of a fixed set of possible θ, $H(p(\theta|\mathbf{a},\xi))$ can be calculated directly. If MCMC is used, one must first infer a known distribution from the samples and then take the entropy of that distribution. For example, if θ is a multinomial and $p(\theta)$ is a Dirichlet distribution, one might infer the most likely Dirichlet distribution from the samples and find the entropy of that distribution.

We have now shown how to (approximately) calculate $U(\xi)$. To complete the procedure for optimal game design, any optimization algorithm that can search through the space of games is sufficient. Maximizing over possible games is unlikely to have a closed form solution, but stochastic search methods can be used to find an approximate solution to the maximum utility game. For example, one might use simulated annealing (Kirkpatrick, Gelatt, & Vecchi, 1983). This method allows optimization of both discrete and continuous parameters, where neighboring states of current game are formed by perturbations of the parameters to be optimized.

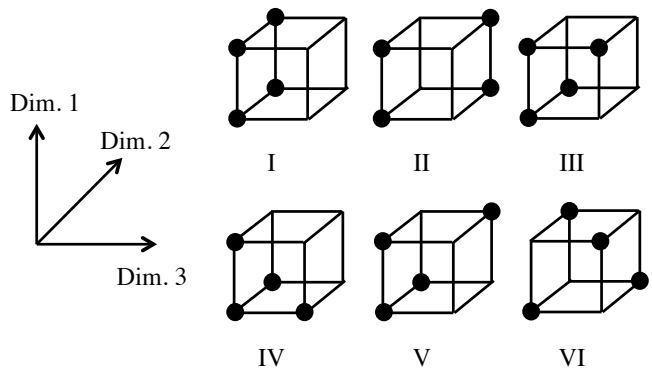

Figure 1: Boolean concept structures. In each structure, eight objects differing in three binary dimensions are grouped into two categories of four elements. Each object is represented as a corner of a cube based on its combination of features, and the objects chosen for one category in each problem type are represented by dots.

Optimal Games for Boolean Concept Learning

We have described a general framework for automatically finding games that are potentially highly informative about model parameters. To test this framework, we applied it to a particular question: What is the relative difficulty of various Boolean concept structures? This question has been studied in past work (e.g, Shepard, Hovland, & Jenkins, 1961; Griffiths, Christian, & Kalish, 2008), so we can compare our results to those produced using more traditional methods. We first describe Boolean concept learning, and then turn to the game we created and the application of optimal game design.

Boolean Concepts

In Boolean concept learning, one must learn how to categorize objects that differ along several binary dimensions. We focus on the Boolean concepts explored in Shepard et al. (1961). In these concepts, there are three feature dimensions, resulting in 2^3 possible objects, and each concept contains four objects. This results in a total of 70 concepts with six distinct structures, as shown in Figure 1. Shepard et al. (1961) found that the six concept structures differed in learning difficulty, with a partial ordering from easiest to most difficult of $\mathrm{I} > \mathrm{II} > \{\mathrm{III}, \mathrm{IV}, \mathrm{V}\} > \mathrm{VI}$. Similar results were observed in later work (Nosofsky, Gluck, Palmeri, McKinley, & Glauthier, 1994; Feldman, 2000) although the position of Type VI in the ordering can vary (Griffiths et al., 2008).

To model learning of Boolean concepts, we assume learners' beliefs about the correct concept h can be captured by Bayes' rule (Griffiths et al., 2008):

$$p(h|\mathbf{d}) \propto p(h)p(\mathbf{d}|h) \tag{8}$$

$$= p(h)\prod_{d\in\mathbf{d}} p(d|h), \tag{9}$$

where each $d \in \mathbf{d}$ is an observed stimulus and its classification, and observations are independent given the category. The likelihood $p(d|h)$ is then a simple indicator function:

$$p(d|h) \propto \begin{cases} 1 & \text{if } h \vdash d \\ 0 & \text{otherwise} \end{cases}, \tag{10}$$

where $h \vdash d$ if the stimulus classification represented by d matches the classification of that stimulus in hypothesis h. We seek to infer the prior $p(h)$, which represents the difficulty of learning different concepts and thus gives an implicit ordering on structure difficulty. In our earlier terminology, θ is a prior distribution on concepts $p(h)$. For simplicity, we assume all concepts with the same structure have the same prior probability, so θ is a 6-dimensional multinomial.

Corridor Challenge

To teach people Boolean concepts we created the game Corridor Challenge, which requires learning Boolean concepts to achieve a high score. Corridor Challenge places the player in a corridor of islands, some of which contain a treasure chest, joined by bridges (Figure 2).[1] The islands form a linear chain and the bridges can be crossed only once, so players cannot return to previous chests. Some chests contain treasure, while others contain traps; opening a chest with treasure increases the player's score and energy, while opening a chest with a trap decreases these values. Each chest has a symbol indicating whether it is a trap; symbols differ along three binary dimensions and are categorized as a trap based on one of the Boolean concepts. Players are shown a record of the symbols from opened chests and their meanings (see the right hand side of Figure 2). Players are told to earn the highest score possible without running out of energy, which is depleted by moving to a new island or opening a trapped chest. When a player runs out of energy, the level is lost and she cannot explore the rest of the level; surviving a level earns the player 250 points. Corridor Challenge games may consist of several levels. Each level is a new corridor with different chests, but the same symbols are used and they retain the same meaning as on the previous level. At the start of each level, the player's energy is restored, but points are retained from level to level.

Optimizing Corridor Challenge

Applying optimal game design to Corridor Challenge requires specifying the parameters to optimize in the search for the optimal game, formulating the game as an MDP, and specifying the model for how the player's prior on concepts (θ) relates to the MDP parameters. The structure of Corridor Challenge allows for many variants that may differ in the expected information gain. To maximize expected information gain while keeping playing time relatively constant we limited the game to two levels, with five islands per level. We then used optimal game design to select the number of points gained for opening a treasure chest, points lost for opening a trap chest, the energy lost when moving, the symbols that

[1] Corridor Challenge uses freely available graphics from http://www.lostgarden.com/2007/05/dancs-miraculously-flexible-game.html

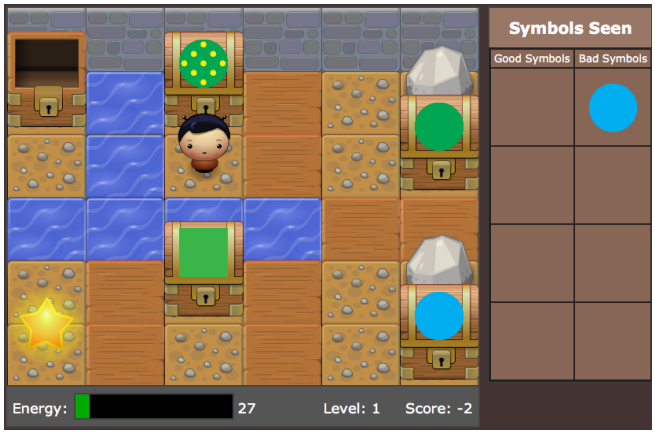

Figure 2: Corridor Challenge game UI (Level one of the random game in Experiment 1). In this screenshot, the player has opened the first chest and moved to the second island.

appeared on the chests, and the Boolean concept used to categorize the chests. For simplicity, we assumed that the number of points gained (or lost) for a particular action is equal to the amount of energy gained (or lost) for that particular action.

Given particular specifications for these variants of the game, we can define an MDP. Note that we define the MDP based on a player's beliefs, since these govern the player's actions, and these beliefs do not include knowledge of the true concept that governs the classification of the symbols:

States: The state is represented by the player's energy, her current level and position in the level, and the symbols on all chests in the current level.

Actions: The player can open the current chest (if there is one) or move to the next island.

Transition model: The player transitions to a new state based on opening a chest or moving to a new island. In both cases, the symbols on the chests stay the same, with the current symbol removed if the player opens the chest. If a player chooses to move, she knows what state will result: her position will move forward one space and her energy will be depleted by a known constant. If the result is negative energy, then the game transitions to a loss state. However, if a player opens a chest, her beliefs about what state will occur next is dependent on $p(h|\mathbf{d})$, her beliefs about the true concept given the data $\mathbf{d}$ she has observed so far. The player will gain energy if the symbol x on the current chest is in the concept. Taking an expectation over possible concepts h, this probability is

$$p(x \text{ in concept}) = \sum_h I(h \vdash x) p(h|\mathbf{d}), \tag{11}$$

where $I(h \vdash x) = 1$ if x is in the concept h and 0 otherwise. The probability of decreased energy is $(1 - p(x \text{ in concept}))$. Based on the Bayesian model presented above, the player's current beliefs $p(h|\mathbf{d})$ are dependent on the prior probability distribution over concepts. Thus, the transition model assumed by the player is dependent on the parameter θ that we would like to estimate, which is this prior distribution.

Reward model: When the player moves from one island to another, the reward model specifies $R(s,a,s') = 0$, and when the player opens a chest, $R(s,a,s')$ is a fixed positive number of points with probability $p(x \text{ in concept})$ and a fixed negative number of points with probability $(1 - p(x \text{ in concept}))$.

By using the MDP framework and assuming that the player updates her beliefs after seeing information, we ignore the value of information in understanding people's decisions; that is, we assume people make decisions based on their current information and do not consider the effect that information gained now will have on their future decisions.

Experiment 1: Inferring Difficulty

To test our framework, we first used the optimal game design procedure to find a version of Corridor Challenge with high expected information gain, and then ran an experiment in which players played the optimized game or a randomly chosen game with lower expected information gain.

Search Methods

We used simulated annealing to search for the best design for Corridor Challenge (Kirkpatrick et al., 1983). The expected information gain of a game was approximated by sampling 35 possible θ vectors uniformly at random (reflecting a uniform prior on θ), simulating the actions of $n = 25$ players in the game, and using the simulated data to infer $p(\theta|\xi,\mathbf{a})$. We approximated $p(\theta|\xi,\mathbf{a})$ using the Metropolis-Hastings MCMC algorithm (Gilks et al., 1996), with a Dirichlet proposal distribution centered at the current state. The parameter β for the Boltzmann policy was set to 1.

To execute the search, we parallelized simulated annealing by using several independent search threads. Every five iterations, the search threads pooled their current states, and each thread selected one of these states to continue searching from, with new starting states chosen probabilistically such that states with high information gain were more likely to be chosen. Each search state is a game, and the next state was found by selecting a parameter of the current game to perturb. If the selected parameter was real-valued, a new value was chosen by sampling from a Gaussian with small variance and mean equal to the current value; if the selected parameter was discrete, a new value was selected uniformly at random.

Search Results

The stochastic search found games with significantly higher information gain than the initial games, regardless of starting point. This demonstrates that the evaluation and search procedure may be able to eliminate some trial and error in designing games for experiments. Qualitatively, games with high information gain tended to have a low risk of running out of energy, at least within the first few moves, and a diverse set of stimuli on the chests. These games also generally had positive rewards with larger magnitudes than the negative rewards. The game with the highest information gain used a true concept of Type II, although several games with similarly

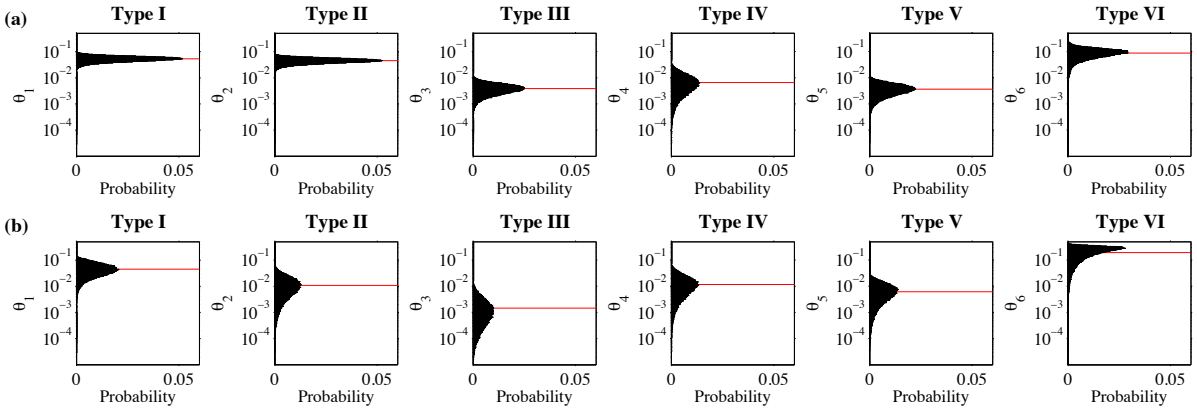

Figure 3: Results of Experiment 1, in the form of posterior distributions on concept difficulty from participants' responses in (a) the optimized game and (b) the random game; red lines indicate the mean of each distribution. Each panel shows the distribution over the inferred difficulty of a concept with the given structure (Types I-VI), as reflected by its prior probability in the Bayesian model. Concepts with higher prior probability are easier to learn. Note the logarithmic scale.

high information gain had true concepts with different structures. While the information gain found for any given game is approximate, since we estimated the expectation over possible priors with only 35 samples, this was sufficient to separate poor games from relatively good games; we explore this relationship further in Experiment 2.

Experiment Methods

Participants. A total of 50 participants were recruited online and received a small amount of money for their time.
Stimuli. Participants played Corridor Challenge with parameters set based either on an optimized game (expected information gain of 3.4 bits) or on a random game (expected information gain of 0.6 bits). The symbols differed along the dimensions of shape, color, and pattern.
Procedure. Half of the participants were randomly assigned to each game design, and played the game in a web browser. The participants were shown text describing the structure of the game, and then played several practice games to familiarize them with interface. The first practice game simply had chests labeled "Good" and "Bad"; the next three games used Boolean concepts of increasing difficulty based on previous work. All practice games used different symbols from one another and from the final game. Practice games used the point and energy values from the game chosen for their condition (i.e., the random game or the game found by the search) in order to make players aware of these values, but the symbols in the practice games were identical across conditions. The final game differed by condition. After completing the final game, participants were asked to rate how fun and how difficult the game was, both on 7-point Likert scales. Additionally, they were shown the stimuli and categorization information that they observed during the final game, and asked to classify the remaining stimuli from the game that were not observed.

Results

Figure 3 shows the inferred distribution over the prior probability of each concept (θ_i) based on participants' actions for the optimized game and the random game; if a concept has higher prior probability, it will be easier to learn. These distributions were obtained via MCMC using a Metropolis-Hastings algorithm on both the prior and the noise parameter β. Results show samples generated from five chains with 100,000 samples each; the first 10% of samples from each chain were removed for burn-in.

Qualitatively, the distributions inferred from the optimized game appear more tightly concentrated than those from the random game; this is confirmed by the actual information gain, which was 3.30 bits for the optimized game and 1.62 bits for the random game. This implies that we could halve the number of participants by running the optimized game.

For both games, the ordering of the mean prior probabilities of each type, shown by red lines in Figure 3, is the same as that found in previous work, except for Type VI. Our inferred distributions for Type VI placed significant probability on many values, suggesting that we simply did not gain much information about its actual difficulty. We do infer that Type VI is easier than Types III, IV, or V, which has a precedent in the results of Griffiths et al. (2008).

Experiment 2: Estimating Information Gain

To verify the relationship between actual and expected information gain, we conducted a second experiment in which players played games with a range of information gains. In order to isolate the impact of the symbols on the chests and the true concept we fixed the point structures to those found for the optimized game in Experiment 1 and conducted new searches over the remaining variables. We then selected games from the search paths that had varying expected information gains, demonstrating that even without changing the incentive structure a range of information gains was possible.

Methods

Participants. A total of 175 participants were recruited online and received the same payment as in Experiment 1.
Stimuli. Participants played one of seven new games.

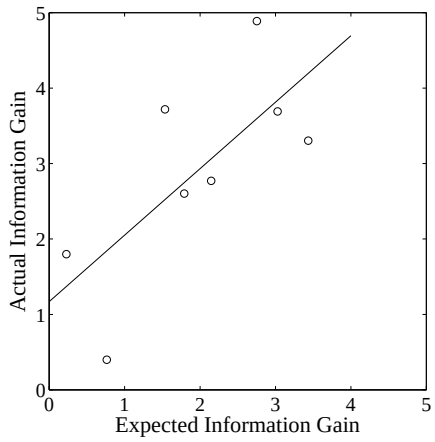

Figure 4: Results of Experiment 2, showing expected versus actual information gains ($r = 0.72$). Each circle represents a game, and the least-squares regression line is shown.

Procedure. Procedure matched Experiment 1.

Results

We compared the actual and expected information gains for the seven new games and the optimized game from Experiment 1, all of which used the same point structure. As shown in Figure 4, we found that expected and actual information gain were positively correlated ($r(6) = 0.72$, $p < 0.05$). This demonstrates that the design of the game does influence how much information we can infer from human players' actions, and that this gain is predicted by our estimates.

Conclusion

We have presented a general framework for the optimal design of games for cognitive science experiments that adapts ideas from Bayesian experiment design. Our methodology hinges on using Markov decision processes to infer cognitively relevant information from players' actions. By combining this model with a stochastic search method, we were able to find appreciably more informative games for Boolean concept learning. Our experimental results demonstrate that using this framework to infer concept difficulty gives similar results to prior work, and that the expected information gain of a game predicts the actual information gain when the game is played by real players. One of the chief advantages of this framework is its applicability to a wide variety of scenarios, from creating other types of games that examine different psychological questions to designing optimally informative assessments within educational software.

There are a number of ways in which this initial test of a framework for designing games could be expanded. In this work, we ignored the value of information for computational tractability. However, people may consider how their future knowledge will be affected by their current actions, leading to deviations from our model's predictions; comparing the fit of a partially observable Markov decision process model to the fit of an MDP model would help to determine whether this is an issue. In the optimal game design procedure, one might also want to explore other metrics than entropy for measuring a game's expected utility. For instance, one might use a loss metric that considers the distance of samples from the true value of θ. Finally, it would be interesting to explore whether there are advantages beyond motivation for using games rather than traditional experiments. While there remain many areas for future exploration, this work gives a starting point for designing highly informative games and gives experimental support that these games can provide meaningful data for cognitive science.

Acknowledgements. This work was supported by a DoD NDSEG Fellowship to ANR, a Google Ph.D. Fellowship to MZ, and NSF grant number IIS-0845410 to TLG.

References

Andrade, G., Ramalho, G., Santana, H., & Corruble, V. (2005). Extending reinforcement learning to provide dynamic game balancing. In *Proceedings of the Workshop on Reasoning, Representation, and Learning in Computer Games, 19th IJCAI* (pp. 7–12).

Atkinson, A., Donev, A., & Tobias, R. (2007). *Optimum experimental designs, with SAS* (Vol. 34). New York: Oxford University Press.

Baker, C. L., Saxe, R. R., & Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition*, *113*(3), 329–349.

Chaloner, K., & Verdinelli, I. (1995). Bayesian experimental design: A review. *Statistical Science*, *10*(3), 273–304.

Erev, I., & Roth, A. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review*, *88*(4), 848–881.

Feldman, J. (2000). Minimization of Boolean complexity in human concept learning. *Nature*, *407*, 630-633.

Gilks, W., Richardson, S., & Spiegelhalter, D. J. (Eds.). (1996). *Markov chain Monte Carlo in practice*. Suffolk, UK: Chapman and Hall.

Griffiths, T. L., Christian, B. R., & Kalish, M. L. (2008). Using category structures to test iterated learning as a method for identifying inductive biases. *Cognitive Science*, *32*, 68-107.

Kirkpatrick, S., Gelatt, C., & Vecchi, M. (1983). Optimization by simulated annealing. *Science*, *220*(4598), 671–680.

Klopfer, E. (2008). *Augmented learning: Research and design of mobile educational games*. Cambridge, MA: The MIT Press.

Michael, D., & Chen, S. (2005). *Serious games: Games that educate, train, and inform*. Boston, MA: Thomson Course Technology.

Myung, J., & Pitt, M. (2009). Optimal experimental design for model discrimination. *Psychological Review*, *116*(3), 499.

Nosofsky, R. M., Gluck, M., Palmeri, T. J., McKinley, S. C., & Glauthier, P. (1994). Comparing models of rule-based classification learning: A replication and extension of Shepard, Hovland, and Jenkins (1961). *Memory & Cognition*, *22*, 352-369.

Shepard, R. N., Hovland, C. I., & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological Monographs*, *75*. (13, Whole No. 517)

Siorpaes, K., & Hepp, M. (2008). Games with a purpose for the semantic web. *Intelligent Systems*, *23*(3), 50–60.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning*. Cambridge, MA: MIT Press.

Tan, C., & Cheng, H. (2009). IMPLANT: An integrated MDP and POMDP learning ageNT for adaptive games. In *Proceedings of The Artificial Intelligence and Interactive Digital Entertainment Conference* (pp. 94–99).

Von Ahn, L. (2006). Games with a purpose. *Computer*, *39*(6), 92–94.