

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Modeling individual differences in socioeconomic game playing

Permalink

<https://escholarship.org/uc/item/44g005bd>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 34(34)

ISSN

1069-7977

Authors

Asher, Derrick
Zhang, Shunan
Zaldivar, Andrew
et al.

Publication Date

2012

Peer reviewed

Modeling individual differences in socioeconomic game playing

Derrick E. Asher¹, Shunan Zhang¹, Andrew Zaldivar¹, Michael D. Lee¹ and Jeffrey L. Krichmar^{1,2}

¹ Department of Cognitive Sciences, University of California, Irvine

² Department of Computer Science, University of California, Irvine

Abstract

Game theory has been useful for understanding risk-taking and cooperative behavior. In the present study, subjects played the Hawk-Dove game with simulated and embodied (robotic) neural agents which used a neurobiologically plausible model of action selection and adaptive behaviors. Subjects had their serotonin levels temporarily altered through acute tryptophan depletion (ATD). The traditional assumption for subject data from Game-theory-ATD or human robot interaction (HRI) studies is that all participants come from the same underlying distribution or same group. We used probabilistic graphical models in order to determine potential sub-group affiliations based on the subjects' responses while playing the Hawk-Dove game. The results from the models indicate sub-groups within a subject population exist. We find that two-group, one that tends toward cooperation and the other that tends toward aggression, best describes the effect of subject behavior in response to ATD and embodiment.

Keywords: Adaptive systems; Human robot interaction; Neurotransmitters; Cognitive Robotics; Bayesian inference; Graphical models; Individual Differences.

Introduction

Economic game theory has had a long, productive history of predicting and describing human behavior in cooperative and competitive situations (Maynard Smith, 1982 ; Nowak, Page, & Sigmund, 2000 ; Skyrms, 2001). The theory of games has also been used to illuminate the neural basis of economic and social decision-making (Lee, 2008 ; Rilling & Sanfey, 2011). However, these studies typically have people play against opponents with set strategies and predictable behavior. Moreover, in most of these studies, subjects are making decisions while sitting in front of an antiseptic computer screen. The present study addresses these issues by having subjects play a socioeconomic game, known as Hawk-Dove, against an autonomous robot with the ability to adapt its behavior to the game situation.

Neuromodulatory systems, such as dopamine and serotonin, appear to be applicable to decision-making in social situations. The serotonergic (5-HT) and dopaminergic (DA) systems oppose each other with respect to predicting punishment (5-HT) versus predicting reward (DA) (Boureau & Dayan, 2011).

We developed a computational model of neuromodulation and action selection based on the assumptions, that dopamine levels are related to the expected reward of an action, and serotonin levels are related to the expected cost or risk of an action (Asher, Zaldivar, & Krichmar, 2010 ; Zaldivar, Asher, & Krichmar, 2010). The model of neuromodulation and action selection demonstrated the ability to adapt to the game situation and its opponent's strategy. The model was embedded in both simulated and embodied neural agents to investigate reciprocal social interactions in games of cooperation

and conflict with people (Asher, Zaldivar, Barton, Brewer, & Krichmar, submitted).

Subjects played a series of Hawk-Dove games against robotic and simulated agents. The effects of serotonergic levels on adaptive behavior in these games were tested by simulating serotonergic lesions in the neural agent, which results in a more aggressive agent, or lowering the CNS serotonin levels of people through a dietary manipulation called acute tryptophan depletion (ATD), which has been shown to decrease cooperation and lower harm-aversion (Crockett, Clark, Tabibnia, Lieberman, & Robbins, 2008 ; Wood, Rilling, Sanfey, Bhagwagar, & Rogers, 2006).

A major finding of the study was that people changed their overall strategies in response to changes in the neural agents state. Subjects tended to deploy either Tit-For-Tat (T4T) or Win-Stay, Lose-Shift (WSLS) strategies during game play. In a T4T strategy, a subject copies the most recent move of the opposing player. In a WSLS strategy, a subject selects the same action that led to a positive payoff in the previous game (Win-Stay), or a different action from the previous game if that action led to zero or negative payoff (Lose-Shift). When playing against a more aggressive neural agent, which had a lesion to its serotonergic system, subjects switched from a Win-Stay, Lose-Shift (WSLS) strategy to a Tit-For-Tat (T4T) strategy. This change in strategy was independent of whether the neural agent was a robot or a computer simulation, and independent of subject tryptophan levels.

In the present study, we test whether embodiment and lowering serotonin has an effect on individual subject behavior during Hawk-Dove game playing by using hierarchical latent mixture models with Bayesian inference. This framework for developing and evaluating structured cognition offers a principled and comprehensive approach for modeling individual differences and their use of cognitive strategies (Lee, 2008 ; Lee, Zhang, Munro, & Steyvers, 2011). The hierarchical nature of the models allows variation in the parameters controlling cognitive processes across individuals to be accommodated. We find that two categories of subjects, one that tends to be more aggressive and one that tends to be more cooperative, best describes subject behavior in response to ATD and embodiment.

Experiment

Subjects

Eight subjects (three female; mean age: 26.6 years; standard deviation of age: 3.8 years) participated in this study.

Hawk-Dove Game

The Hawk-Dove game consisted of a human and a neural agent choosing a single action in response to a territory of interest (TOI). The Hawk-Dove game, which is similar to Prisoner's Dilemma, was chosen because it is amenable to a physical instantiation with a robot. Moreover, it has an additional strategic element since choices are different depending on who arrives at the TOI first. At the start of the game, the TOI and the human subjects' location were randomly placed on a playing grid. The current location of the robot was used as a starting position. The player who arrived at the neutral TOI first had the opportunity to take one of two possible actions: Escalate (i.e. an aggressive, confrontational tactic) or Display (i.e. a nonviolent, cooperative tactic). The player who arrived second responded with one of the two aforementioned actions. After each game, a payoff was calculated. If both players chose Escalate, they received a penalty that is set before the game. If both players chose Display, they split the value of the TOI resource. If one player chose Escalate and the other chose Display, the player who chose Escalate received the entire value of the resource.

We also developed a simulated variant of the Hawk-Dove game, where subjects played against a robot icon on an interactive screen. The same neural model used for the real robot dictated the control of the robots icon and its decision-making. This simulated setup allowed us to judge whether physical embodiment had an effect on human behavior.

Acute Tryptophan Depletion (ATD)

ATD was used to temporarily alter the levels of serotonin in the brain via a decrease in blood plasma tryptophan, the amino acid precursor to serotonin. Because free blood plasma tryptophan levels, and the corresponding serotonin levels in the brain, vary with the amount of dietary tryptophan and the rate of protein synthesis, these levels can be altered by a low protein diet in combination with a specially prepared 'protein shake'. This protein shake contains an amino acid load (lacking tryptophan), which has two effects. First, it stimulates protein synthesis in the liver, which uses up blood plasma tryptophan. Second, the amino acids that are given in the protein shake compete with tryptophan for transport across the blood-brain barrier, which restricts entry of tryptophan into the brain and leads to lower levels of serotonin in the brain (e.g. Bell, Hood, & Nutt, 2005 ; Hood, Bell, & Nutt, 2005).

General Procedures

In a double-blind study, human subjects were randomly assigned on the first experimental day to receive either the control mixture (Tryp+) with tryptophan or the mixture without tryptophan (Tryp-). The mixtures were administered as a specially prepared protein shake. The Tryp+ and Tryp- shakes contained 16 and 15 amino acids respectively.

Each subject then returned to participate in the other condition at least seven days later to ensure the return to baseline blood plasma tryptophan levels between experimental days. On the morning of each experimental day, a blood sample

was drawn to determine baseline blood plasma tryptophan levels. Following the blood draw, subjects ingested one of the mixture drinks (either Tryp+ or Tryp-). A second blood sample was drawn approximately 5 hours after ingestion of the mixture to confirm reduction (Tryp- condition) or maintenance (Tryp+ condition) of blood plasma tryptophan levels. Roughly 5.5 hours after consumption of the mixture, human subjects then participated in a series of Hawk-Dove games against a neural agent.

Experimental Conditions and Data

We were interested in two experimental conditions per subject. In the Simulation vs. Robot condition, subjects would play games against a computer agent or against the robot; In the Tryp+ vs. Tryp- condition, subjects would play games against a neural agent with an intact simulated neuromodulatory system or a simulated lesion of its serotonergic system.

The Escalate or Display decisions for each game were collected for both the subject and neural agent for all games in each condition. The dependent variables of interest are the percentage of games per condition a subject chose the Escalate tactic and the use of either the Tit-For-Tat (T4T) or Win-Stay, Lose-Shift (WSLS) strategies. Human subjects played 20 games of Hawk-Dove per condition. For detailed experimental conditions, see Asher et al. (submitted).

Data Analysis

Bayesian Hierarchical Model Approach

To investigate the influences of lowering serotonin levels and of agent embodiment on individual decision-making, we used hierarchical latent mixture models with Bayesian inference. Hierarchical Bayesian inference has been demonstrated as a flexible and interpretable way of extending simple models of cognitive processes (e.g. Lee, 2008 ; Rouder, Lu, Speckman, Sun, & Jiang, 2005 ; Wetzels, Vandekerckhove, Tuerlinckx, & Wagenmakers, 2010). The hierarchical nature of the models allows variation in the parameters controlling cognitive processes across individuals to be accommodated. The latent mixture nature of the models allows the use of entirely different cognitive strategies across individuals to be modeled.

Formally, we recast the cognitive models as probabilistic graphical models and used Markov Chain Monte Carlo (MCMC) methods for computational Bayesian inference. This approach to Bayesian inference over richly-structured cognitive models has been applied to data covering a diverse set of cognitive skills. For example, Bayesian graphical models have been used to make inferences about the use of strategies, such as WSLS or T4T, from sequences of choice data in bandit problems and other sequential decision-making tasks (e.g. Lee et al., 2011 ; Newell & Lee, 2011).

Using hierarchical latent mixture models, we addressed the question of how ATD and embodiment can affect subjects' decisions to compete (i.e., escalate) or cooperate (i.e., display). We modeled the probability of escalating through a logistic model. Specifically, the logit of the probability of es-




Figure 1: Generic process of escalation or strategy usage under the effect of ATD or embodiment. ϕ : baseline tendency of escalation or strategy usage; ψ^u , ψ^d and ψ^s : additive effects of escalation or strategy usage associated with the experimental conditions; t : indicator of experimental condition; θ : probability of escalation or strategy usage; d : observed escalation decision or strategy usage; Z : latent group indicator; $\mu^\phi, \sigma^\phi, \mu^u, \mu^d, \sigma^c$: hyper-parameters of prior distributions. Experimental conditions are either Tryp- vs. Tryp+, or Robot vs. Sim.

calating for each subject in each condition is assumed to follow a Gaussian distribution defined by its mean and variance (hyperparameters in the hierarchical model), with the mean modeled as the sum of the baseline level of escalating for the subject, and an additive effect associated with ATD. To give a full account of the data, the hierarchical model is designed to address individual differences at two levels, the baseline level, which depends on the subjects, and the additive level, which depends on the interaction between subjects and experimental conditions. Our justification is that it is possible that the effect of ATD or embodiment can vary across different individuals, resulting in either an increase or decrease in the likelihood of escalating a fight.

The Graphical Model

We use graphical models to describe the relationship between subject decision-making (escalation, T4T and WSLS) and

predictors of interest (ATD and embodiment). In total we built six graphical models for every decision-making and predictor pair. In this section, we show a concrete example of how we model the relationship between escalation and ATD; all other models were built along a similar line.

As shown in Figure 1, nodes represent variables of interest, and the graph structure is used to indicate dependencies between the variables. Arrows run into nodes from their direct influences (parents). Formally, the model represents the assumption that, given its parent nodes, each node is independent of all other nodes in the graph except its descendants.

Each subject is assumed to produce data through the same generative model with different parameters. The plate with i subjects corresponds to independent replications for all subjects. Each subject is assumed to have their own baseline of escalation, ϕ_i , independently drawn from a Gaussian distribution with mean μ^ϕ and variance σ^ϕ (termed as hyper-parameters). Each subject i is also associated with a latent group identity, Z_i , with $Z_i = 1$ indicating that the subject comes from a ‘down group’ that shows less escalation decisions with ATD, and $Z_i = 2$ indicating that the subject comes from an ‘up group’ that shows more escalation decisions with ATD.

The plate with j conditions corresponds to independent replications for all experimental conditions. t_j is an observed variable, with $t_j = 1$ indicating the control condition (Tryp+), and $t_j = 2$ indicating the treatment condition where the subjects received ATD (Tryp-).

In the control condition, random fluctuations around the baseline are drawn from a Gaussian distribution with mean of 0 and a variance σ^c . In the ATD condition, the increment in escalation depends on which group the subject is from. For example, if subject i is from the ‘up group’, the increment of escalation, denoted by ψ_{ij}^u in the graphical model, will be drawn from a Gaussian with a positive mean μ^u and variance σ^c and added to the baseline ϕ_i . If subject i is from the ‘down group’, however, the increment ψ_{ij}^d will be drawn from a Gaussian with a different positive mean μ^d and variance σ^c and deducted from the baseline¹.

The probability of escalation for subject i on the k th game of day j , denoted by θ_{ijk} , is determined through the logistic function by the overall level of escalation. The observed escalation decision of each subject on each game, d_{ijk} , is a binary variable which is assumed to be independent and identically generated by θ_{ijk} through a Bernoulli distribution. Overall, one would determine the likelihood of the observed data for all subjects for each combination of the hyperparameters, and each choice of individual parameters.

We use the conventions of representing continuous variables with circular nodes and discrete variables with square nodes, and unobserved variables without shading and observed variables with shading, and stochastic variables with

¹Even though in theory, ψ_{ij}^u and ψ_{ij}^d could flip in sign following the way this model is specified, this almost never happened in the sampling process when the MC chains have converged.

single borders and deterministic variables with double borders. In addition, the plate encloses subsets of the graph that have independent replications in the model. For example, the probability of escalation for each subject in each game is a continuous variable that is not directly observable, but is determined by the overall level of escalation, therefore it is represented by a circular, unshaded, double-bordered node. On the other hand, the binary escalation decision is generated with probability and directly observable, therefore it is represented by a square, shaded node with a single border.

We built another model that captures the relationship between escalation and the effect of embodiment (Robot vs. Simulation). There are only two differences from the graphical model shown in Figure 1. First, t_j represents embodiment with t_1 indicating robot and t_2 indicating simulation; secondly, $Z_i = 1$ indicates the latent group that shows more escalation when playing against simulation, and $Z_i = 2$ indicate the group that shows more escalation when playing against robot.

In addition to escalation decisions, we are also interested in whether general strategies that subjects employed for the Hawk-Dove game were related to ATD and embodiment. Similar to the generative model shown in Figure 1, the observed T4T-type (WSLS-type) decisions are Bernoulli variables generated by probabilities of T4T (WSLS) usage, which are determined by an overall level of T4T (WSLS) usage that is the sum of a baseline level of usage and additive effect from ATD or embodiment.

Results

Inferring Strategies

All the graphical models were implemented using WinBUGS which uses MCMC methods. We evaluated all six graphical models by drawing 1000 posterior samples after a ‘burn-in’ period (early steps of MCMC where samples are not recorded so that the Markov chain is allowed time to converge) of 100 samples.

The results showed that both serotonin levels and the embodiment of a robot were influential factors in individual subject decision-making (See Figure 2). For example, panel (a) provides evidence that there are at least two sub-groups within a subject population with respect to how escalation may be altered by ATD. Subject 2, 6 and 8 fall in the ‘down’ group, and subject 1, 3, 4, 5 and 7 fall in the ‘up’ group. Panel (b) shows two sub-groups within a subject population with respect to how escalation may be altered by embodiment. Subject 3, 5 and 6 fall in the ‘down’ group (less escalation when playing with a robot), while all other subjects fall in the ‘up’ group (more escalation when playing with a robot). Individual differences in strategy usage were also affected by ATD and embodiment, as shown in Figure 2c through f. For example, Figure 2c shows individual differences in the effect of ATD on the proportion of T4T-type decisions. The red (dark) dots are subjects who had more T4T-type decisions when tryptophan depleted and the green (light) dots are subjects

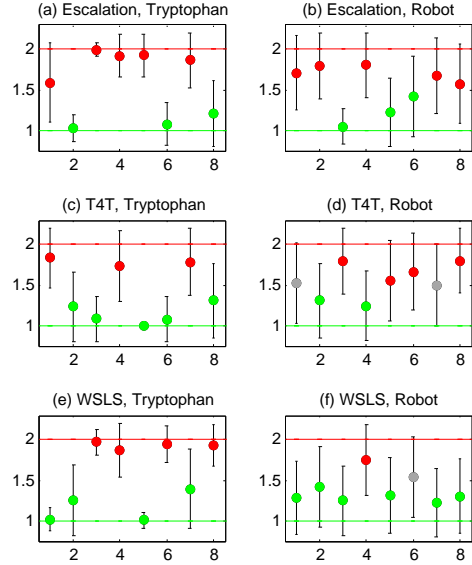


Figure 2: (a): Estimated group identities based on ATD’s effect on escalation. Horizontal axis shows subject indices. Vertical axis shows the posterior mean of the latent variable Z , with $Z = 1$ indicating ‘down’ group and $Z = 2$ indicating ‘up’ group based on the effect of tryptophan depletion. Red (dark) dots are subjects who were more likely to escalate when tryptophan depleted and green (light) dots are subjects who were less likely to escalate when tryptophan depleted. Error bars show the 95% Bayesian CI of the posterior mean. Gray dots imply ambiguous group identities. (b): Estimated group identities based on embodiment’s effect on escalation. Red dots are subjects who were more likely to escalate when playing against a robot. (c): Estimated group identities based on ATD’s effect on T4T usage. Red dots indicate more T4T-type decisions when tryptophan depleted. (d): Estimated group identities based on embodiment’s effect on T4T usage. Red dots indicate more T4T-type decisions when playing against a robot. (e)-(f): Estimated group identities on the effect of ATD and embodiment on WSLS usage.

who had less T4T-type decisions when tryptophan depleted. In Figure 2d, red (dark) dots represent subjects who had more T4T-type decisions when playing against robot. Similarly, Figure 2e and 2f represent individual differences in how ATD and embodiment affect their usage of WSLS. All conditions show a tendency that subjects can be categorized into sub-groups with regard to how ATD or embodiment may affect strategy usage.

In general, subgroups are more clear when the effect of tryptophan depletion rather than the embodiment type is considered. Potential connections of subgroup identities across conditions for each subject is an interesting point raised here and will be addressed in future research.

Predictions of Behavioral Patterns

To check the models' ability to describe the data accurately, we examined the posterior predictions of escalation decision and strategy usage for all subjects in all conditions. The posterior predictive is the prediction about observed data for each possible combination of parameter values under the model, where each combination is weighted according to its posterior probability. Our goal is not game-by-game prediction, rather, it is the prediction of overall rate of escalation decision and strategy usage in a specified condition, as captured in Figure 3. The x-axis shows experimental conditions, and the y-axis shows the proportion of escalation decisions or strategy usage. Each gray line is a subject. The colored (dark) lines are posterior predictives (summarized in the same way as the data) for each model, with filled vs. open circles representing whether subjects were inferred in the 'up' or 'down' groups corresponding to the condition (ATD or embodiment). It is clear that our models were able to capture individual differences and fit the data well, especially for the effect of ATD and embodiment on escalation and T4T usage. Predictions of proportions of WSLs-type decisions had larger fluctuations, but the general pattern of change in WSLs usage associated with ATD and embodiment were captured.

Conclusions

In contrast to the results observed in our population analysis of subject behavior during Hawk-Dove game play (Asher et al., submitted), we found strong influences of lowering serotonin levels and of agent embodiment on individual decision-making by using hierarchical latent mixture models with Bayesian inference. The results from Figure 2 and Figure 3 show that individual differences likely exist within a subject population. Specifically, it appears that there are two groups of subjects within a given population. Our hypothesis is that within Game-theory-ATD and HRI studies, there exist two opposing subgroups within any given population of typical human subjects. These subgroups may exist due to inherent genetic variation. It is interesting to note that the results from Game-theory-ATD and HRI studies are mild when considering a single subject group. However, if one considers that there exist two groups with opposing affects, the results for these kinds of studies may be significantly more robust.

ATD had a strong effect on subject behavior and this behavior could be categorized in two groups: subjects who were more aggressive when tryptophan depleted and subjects who were less aggressive when tryptophan depleted. Aggressive, uncooperative behavior has been reported in behavioral studies in which serotonin levels were lowered through ATD (Schweighofer et al., 2008 ; Tanaka et al., 2007). However, individual variation both due to experience and genetic background can affect behavior. For example, there is widespread variation in the serotonin transporter gene 5-HTTLPR (Homberg & Lesch, 2011). Subjects carrying the short allele variant of the 5-HTTLPR outperform subjects with the long allele in an array of cognitive tasks and show

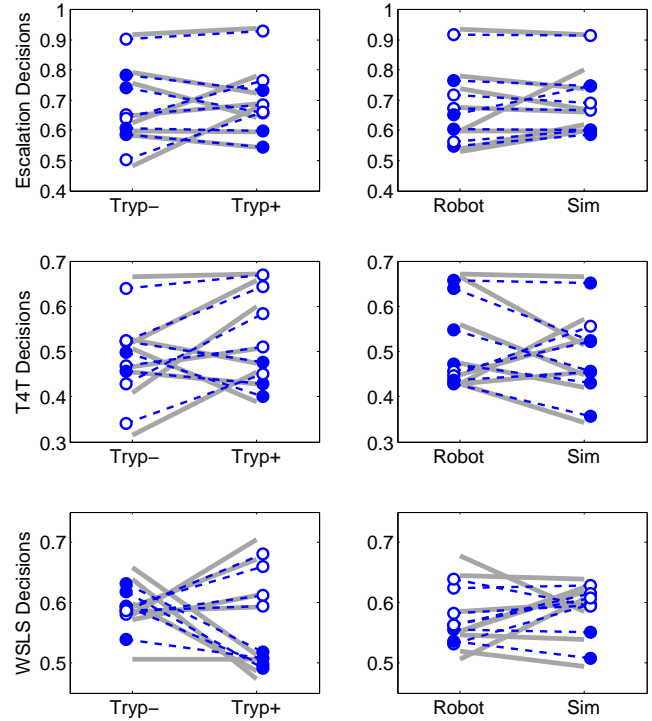


Figure 3: Predicted escalation rate by the model for all subjects with a comparison to the real data. Tryptophan depleted and non-depleted conditions labeled as 'Tryp-' and 'Tryp+'. The broken lines are predictions from the models, with filled vs. open circles representing subjects who show more escalation when tryptophan depleted vs. subjects who show less escalation when tryptophan depleted. The solid lines are the data. Labels of 'Robot' and 'Sim' indicate whether the game was played against the robot or the simulation.

increased social conformity under normal conditions. However, subjects carrying the long allele variant perform better under stressful conditions. The prevalence of these and other genetic polymorphisms in the human population suggests that there is an evolutionary advantage for this variability, such as optimizing competition or cooperation in different situations.

Embodiment had a strong effect on subject behavior and, similar to the ATD effect, this behavior could be categorized in two groups: subjects who were more aggressive when playing against a robot and subjects who were less aggressive when playing against a robot. Playing an opponent who is interactive and personified has previously been observed to evoke strong responses in subjects. For example, in the Ultimatum Game, subjects rejected more offers made by a human partner than those offers made by a computer, suggesting that participants had a stronger emotional reaction to unfair offers from humans than from a computer (Sanfey, Rilling, Aronson, Nystrom, & Cohen, 2003). Indeed, such embodied models have been shown to elicit strong reactions in humans (Breazeal & Scassellati, 2002 ; Kidd & Breazeal, 2004)

and to exhibit more natural and complex behavior than pure simulations (Krichmar & Edelman, 2002, 2005). However, it appears that individuals responded differently and idiosyncratically to the presence of a robot in the present study.

Our results highlight the following ideas: first, the hierarchical latent mixture model's ability to capture individual differences; secondly, serotonin levels have differing effects on subject decision-making, and lastly embodiment plays a role in how likely subjects are willing to cooperate with an agent. Our results suggest that there are at least two opposing subgroups with respect to Game-theory-ATD and HRI studies. We suggest the possibility that these subgroups may have emerged as a result of genetic variation. A next step towards investigating this hypothesis involves genetic testing of polymorphisms responsible for subject variance and their decision-making in competitive and cooperative games of Game theory along with human robot interactions.

Acknowledgment

This work was supported by the National Science Foundation (EMT/BSSE Award No.: 0829752) and the Office of Naval Research (Award No.: N000140910036). MDL and SZ acknowledge support from the Air Force Office of Scientific Research Award FA9550-11, and Australian Research Council Grant DP110100797.

Références

- Asher, D., Zaldivar, A., Barton, B., Brewer, A., & Krichmar, J. (submitted). Reciprocity and retaliation in social games with adaptive agents. *IEEE Transactions on Autonomous Mental Development*.
- Asher, D., Zaldivar, A., & Krichmar, J. (2010). Effect of neuromodulation on performance in game playing: A modeling study. In *IEEE 9th international conference on development and learning* (p. 155-160).
- Bell, C., Hood, S., & Nutt, D. (2005). Acute tryptophan depletion. part II: clinical effects and implications. *Aust N Z J Psychiatry*, 39, 565-574.
- Boureau, Y., & Dayan, P. (2011). Opponency revisited: competition and cooperation between dopamine and serotonin. *Neuropsychopharmacology*, 36, 74-97.
- Breazeal, C., & Scassellati, B. (2002). Robots that imitate humans. *Trends in Cognitive Sciences*, 6, 481-487.
- Crockett, M., Clark, L., Tabibnia, G., Lieberman, M., & Robbins, T. (2008). Serotonin modulates behavioral reactions to unfairness. *Science*, 320, 1739.
- Homberg, J., & Lesch, K. (2011). Looking on the bright side of serotonin transporter gene variation. *Biol Psychiatry*, 69, 513-519.
- Hood, S., Bell, C., & Nutt, D. (2005). Acute tryptophan depletion. part I: clinical effects and implications. *Aust N Z J Psychiatry*, 39, 558-564.
- Kidd, C., & Breazeal, C. (2004). Effect of a robot on user perceptions. In (p. 3559-3564).
- Krichmar, J., & Edelman, G. (2002). Machine psychology: autonomous behavior, perceptual categorization and conditioning in a brain-based device. *Cerebral Cortex*, 12, 818-830.
- Krichmar, J., & Edelman, G. (2005). Brain-based devices for the study of nervous systems and the development of intelligent machines. *Artificial Life*, 11, 63-77.
- Lee, D. (2008). Game theory and neural basis of social decision making. *Nature Neuroscience*, 11, 404-409.
- Lee, M. D. (2008). Three case studies in the Bayesian analysis of cognitive models. *Psychonomic Bulletin & Review*, 15, 1-15.
- Lee, M. D., Zhang, S., Munro, M., & Steyvers, M. (2011). Psychological models of human and optimal performance on bandit problems. *Cognitive Systems Research*, 12, 164-174.
- Maynard Smith, J. (1982). *Evolution and the theory of games*. Cambridge University Press.
- Newell, B., & Lee, M. (2011). The right tool for the job? comparing an evidence accumulation and a naive strategy selection model of decision making. *Journal of Behavioral Decision Making*, 24, 456-481.
- Nowak, M., Page, K., & Sigmund, K. (2000). Fairness versus reason in the ultimatum game. *Science*, 289, 1773-1775.
- Rilling, J., & Sanfey, A. (2011). The neuroscience of social decision making. *Annual Review Psychology*, 62, 23-48.
- Rouder, J., Lu, J., Speckman, P., Sun, D., & Jiang, Y. (2005). Unknown title. *Psychonomic Bulletin & Review*, 12, 195-223.
- Sanfey, A., Rilling, J., Aronson, J., Nystrom, L., & Cohen, J. (2003). The neural basis of economic decision-making in the ultimatum game. *Science*, 300, 1755-1758.
- Schweighofer, N., Bertin, M., Shishida, K., Okamoto, Y., Tanaka, S., Yamawaki, S., et al. (2008). Low-serotonin levels increase delayed reward discounting in humans. *J Neurosci*, 17, 4528-4532.
- Skyrms, B. (2001). The stag hunt. In *Presidential address pacific division of the american philosophical association*.
- Tanaka, S., Schweighofer, N., Asahi, S., Shishida, K., Okamoto, Y., Yamawaki, S., et al. (2007). Serotonin differentially regulates short- and long-term prediction of rewards in the ventral and dorsal striatum. *Public Library of Science ONE*, 12, 2.
- Wetzels, R., Vandekerckhove, J., Tuerlinckx, F., & Wagenmakers, E. (2010). Bayesian parameter estimation in the expectance valence model of the Iowa gambling task. *Journal of Mathematical Psychology*, 54, 14-27.
- Wood, R., Rilling, J., Sanfey, A., Bhagwagar, Z., & Rogers, R. (2006). Effects of tryptophan depletion on the performance of an iterated prisoner's dilemma game in healthy adults. *Neuropsychopharmacology*, 31, 1075-1084.
- Zaldivar, A., Asher, D., & Krichmar, J. (2010). *Simulation of how neuromodulation influences cooperative behavior*. Springer-Verlag Lecture Notes on Artificial Intelligence (LNAI 6226).