

DATA-DRIVEN DYNAMIC DECISION MODELS

John J. Nay

School of Engineering
Vanderbilt University
PMB 351826

2301 Vanderbilt Place
Nashville, TN 37235-1826, USA

Jonathan M. Gilligan

Department of Earth & Environmental Sciences
Vanderbilt University
PMB 351805

2301 Vanderbilt PL
Nashville, TN 37235-1805, USA

ABSTRACT

This article outlines a method for automatically generating models of dynamic decision-making that both have strong predictive power and are interpretable in human terms. This is useful for designing empirically grounded agent-based simulations and for gaining direct insight into observed dynamic processes. We use an efficient model representation and a genetic algorithm-based estimation process to generate simple approximations that explain most of the structure of complex stochastic processes. This method, implemented in C++ and R, scales well to large data sets. We apply our methods to empirical data from human subjects game experiments and international relations. We also demonstrate the method's ability to recover known data-generating processes by simulating data with agent-based models and correctly deriving the underlying decision models for multiple agent models and degrees of stochasticity.

1 MODEL REPRESENTATION

This article describes a modeling method designed to understand data on dynamic decision-making. We have created a practical, easy-to-use software package implementing the method. Although our method is more broadly applicable, the motivation for the model representation was prediction of individual behavior in strategic interactions, i.e. games. Most *behavioral* game-theoretic treatments of repeated games use action-learning models that specify the way in which attractions to actions are updated by an agent as play progresses (Camerer 2003). Action learning models can perform poorly at predicting behavior in games where cooperation (e.g. Prisoner's Dilemma) or coordination (e.g. Bach or Stravinsky) are key (Hanaki 2004). Also, they often fail to account for the effects of changes in information and player matching conditions (McKelvey and Palfrey 2001). In this paper, we model repeated game strategies as decision-making procedures that can explicitly consider the dynamic nature of the environment, e.g. if my opponent cooperated last period then I will cooperate this period. We represent decision-making with finite-state machines and use a genetic algorithm to estimate the values of the state transition tables. This combination of representation and optimization allows us to efficiently and effectively model dynamic decision-making.

Traditional game theories define strategies as complete contingent plans that specify how a player will act in every possible state; however, when the environment becomes even moderately complex the number of possible states of the world can grow beyond the limits of human cognition (Miller 1996, Fudenberg et al. 2012). One modeling response to cognitive limitations has been to exogenously restrict the complexity of repeated game strategies by representing them as Moore machines – finite state machines whose outputs depend only on their current state (Moore 1956) – with a small number of states (Rubinstein 1986, Miller 1996, Hanaki et al. 2005). Moore machines can model bounded rationality, explicitly treating procedures of decision-making (Osborne and Rubinstein 1994). A machine modeling agent i responding to the actions of agent j is a four-tuple $[Q_i, q_i^0, f_i, \tau_i]$, where Q_i is the set of states, $q_i^0 \in Q_i$ is the initial state, $f_i : Q_i \rightarrow A_i$ is the output function mapping a state to an action, and $\tau_i : Q_i \times A_j \rightarrow Q_i$ (where $j \neq i$) is the transition function

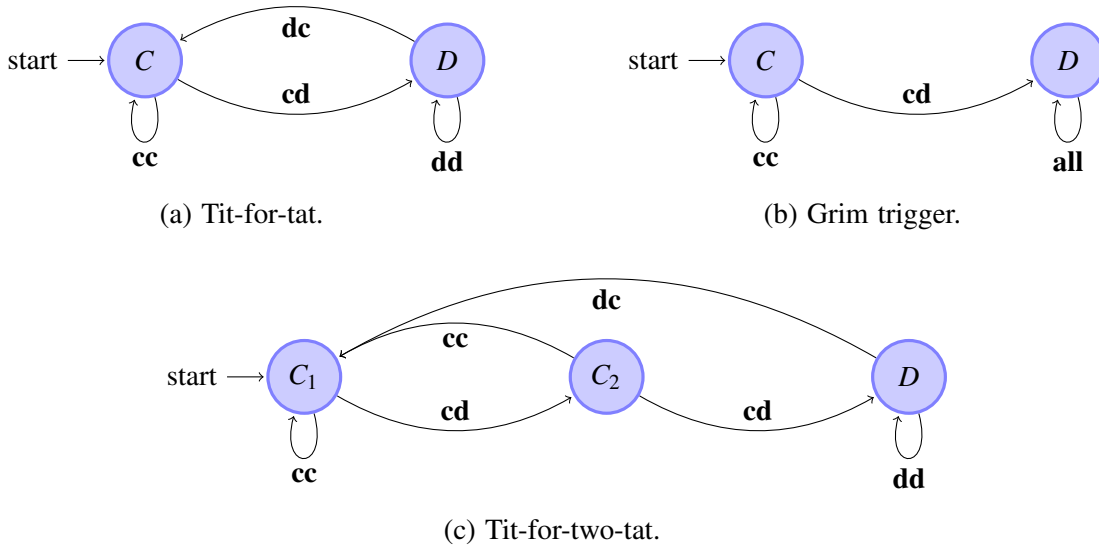


Figure 1: Game theoretic strategies represented as finite state machines (Rubinstein 1986, Miller 1996, Fudenberg et al. 2012). Lower-case letters denote the possible outcomes of the previous period of play, e.g. **cd** means the player cooperated and her opponent defected.

mapping a state and another agent’s action to a state (Osborne and Rubinstein 1994). We generalize this model beyond games by allowing for more inputs in τ_i than A_j , and by providing empirical rankings of these inputs that can be used to induce sparsity in more context-rich environments. The Moore machine can have many underlying states for a single observable action, allowing it to represent arbitrarily complex decision processes. The complexity is directly controlled by the number of states, which is a tuning parameter of our method that can be optimized by Algorithm 2 for predictive performance.

Fig. 1 shows examples of finite state machines (FSMs) representing strategies for the Iterated Prisoner’s Dilemma game (see Section 3 for game details): The possible states are cooperate (C) and defect (D), and after initialization the current state is determined by the history of the player and her opponent cooperating or defecting (**cc**, **cd**, **dc**, **dd**) in the previous period.

2 ESTIMATION

Genetic algorithms (GAs) have been used to model agents updating beliefs based on endogenously determined variables in a general equilibrium environment (Bullard and Duffy 1999), and agents learning to make economic decisions (Arifovic 1994, Arifovic and Eaton 1995, Marks et al. 1995, Midgley et al. 1997). In contrast to investigations of GAs as models of agent learning and behavior, we use GAs to automatically generate interpretable agent decision models from empirical data. This is similar to work by Fogel (1993), Miller (1996), and Miller and Page (2007), in which GAs evolved FSMs based on their interactions with one another in simulated games, but whereas these were theoretical exercises, we are estimating models to explain and predict observed interactions among real agents. We use GAs as optimization routines for estimation because they perform well in rugged search spaces to quickly solve discrete optimization problems, are a natural complement to our binary string representation of FSMs (Goldberg and Holland 1988), and are easily parallelized.

Duffy and Engle-Warnick (2002) combined empirical experimental data with genetic programming (GP) to model behavior. GP, with the same genetic operations as most GAs (Koza 1992), is a process that can evolve arbitrary computer programs (Duffy 2006). We apply genetic operations to FSM representations rather than to all predictor variables and functional primitives because we are interested in deriving decision

models with a particular structure: FSMs with latent states, rather than models conditioning on observable variables with any arbitrary functional form. With data-driven modeling, it is desirable to impose as many constraints as can be theoretically justified on the functional form of the model (see Miller and Page (2007) for interesting theoretical results related to FSM agents interacting in games). This avoids overfitting by constraining the model to a functional form that is likely generalizable across contexts, allows genetic selection to converge better, and reduces the computational effort required to explore parameter space. An additional challenge in implementing GP is specifying the genetic operations on the function primitives while ensuring that they will always produce syntactically valid programs that represent meaningful decision models. This requires fine-tuning to specific problems, which we avoid because we are designing a general method applicable across domains.

Our choice to use Moore machines as the building blocks of our decision modeling method ensures that estimation will produce easily interpretable models with latent states that can be represented graphically (see Fig. 1 for examples). Our process represents Moore machines as Gray-encoded binary strings consisting of an action vector followed by elements that form the state matrix (Savage 1997). For details, see Fig. 3a and our `build_bitstring`, `decode_action_vec`, and `decode_stat_mat` functions. This way, genetic operators can have free reign to search the global parameter space guided by the ability to predict provided data with the decoded binary strings.

Data: Time series of actions taken by agents and the relevant predictors of each action.
Result: Finite state machine (FSM) with highest predictive performance.
 Set convergence criteria (maximum number of generations or number of generations without improvement in performance of best FSM) based on the number of parameters to estimate;
 Create initial population at step $k = 0$ of p individuals $\{\theta_1^0, \theta_2^0, \dots, \theta_p^0\}$ (FSMs, encoded as binary strings);
while *convergence not satisfied* **do**
 Decode each individual's string into an FSM, evaluate its predictive performance on training data, and set this as the individual's fitness for step k , $f(\theta_i^k)$;
 Assign each individual a probability for selection proportional to its fitness, $p_i^k \propto f(\theta_i^k)$;
 Select individuals by sampling from the population with replacement;
 Create next generation, $\{\theta_1^{k+1}, \theta_2^{k+1}, \dots, \theta_p^{k+1}\}$, by applying random *crossover* and *mutation* to the selected sub-population;
end
 Return $\text{argmax}_{\theta_i^k} f(\theta_i^k)$: the individual with the greatest predictive power;
 Check each element of the solution's state transition matrix for contribution to predictive performance, and evaluate solution on test data, if supplied.

Algorithm 1: Evolving finite state machines with a genetic algorithm.

The vast majority of computation time for Algorithm 1 is the evaluation of the predictive accuracy of the FSMs (not the stochastic generation of candidate FSMs). To improve performance we implement this evaluation in C++ using the `Rcpp` package (Eddelbuettel 2013), and, because it is embarrassingly parallel, distribute it across processor cores. We have incorporated our code into an R package with an API of documented function calls and using the `GA` package (Scrucca 2013) to perform the GA evolution. A user can generate an optimized FSM by calling `evolve_model(data)`, where `data` is an R `data.frame` object with columns representing the time period of the decision, the decision taken at that period, and any predictor variables. There are many additional optional arguments to this `evolve_model` function, but they have sensible default values. Our package then generates C++ code for a fitness function and uses it to evaluate automatically generated candidate models. Once the convergence criteria of this iterative search process is satisfied, the best FSM is identified, and each predictor variable is assessed by checking its

	C	D
C	(3,3)	(0,4)
D	(4,0)	(1,1)

Figure 2: Payoff table for a Prisoner’s Dilemma game. C is cooperate and D is defect.

identifiability and computing its importance in that decision model. The return value contains a descriptive summary of all results, including those shown in Fig. 3.

The number of states in the FSM and the number of predictor variables to include are hyper-parameters that control the complexity of the model. Beginning with the simplest possible model and increasing complexity by adding states and variables, we often observe that at first, out-of-sample predictive accuracy grows because bias falls more quickly than variance rises; but eventually, adding further complexity reduces bias less than it increases variance so accuracy decreases (Hastie et al. 2009). We can use cross-validation on the training data to find values for the hyper-parameters that maximize predictive accuracy (Algorithm 2). We assess the out-of-sample predictive accuracy of the final model with a hold-out test set of data, distinct from the cross-validation test-sets in Algorithm 2. Increasing complexity to optimize predictive accuracy introduces a new trade-off because more complex decision models are harder to interpret in human terms, so the “best” solution will depend on the goals of the analysis.

Data: A dataset, a performance metric (e.g. accuracy or area under the ROC curve), and possible hyper-parameter values (e.g. number of states or predictor variables to include).

Result: An estimate of the hyper-parameters that lead to best predictive performance.

for each row in design matrix of hyper-parameter sets **do**

 Sample data into k (e.g., 10) groups;

for each k **do**

 Set group k as a testing set and everything else as a training set;

 Evolve model on training set with Algorithm 1;

 Predict testing set and compare to actual values based on performance metric;

end

 Calculate average performance across all predictions on all k testing sets;

end

Return hyper-parameter set with best average performance;

Algorithm 2: Use cross-validation to optimize FSM hyper-parameters for predictive performance.

3 EXPERIMENTAL GAME DATA

The Iterated Prisoner’s Dilemma (IPD) is often used as a model of cooperation (Axelrod 1984). A one-shot PD game has a unique equilibrium in which each player chooses to defect even though both players would be better off if they cooperated. Suppose two players play the simultaneous-move PD game in Fig. 2, observe the choice of the other person, and then play the same simultaneous-move game again. Even in the (finitely) repeated version, no cooperation can be achieved by rational income maximizers. This tension between maximizing collective and individual gain is representative of a broad class of social situations (e.g. the “tragedy of the commons” (Hardin 1968)).

We applied our procedure to data from laboratory experiments on human subjects playing IPD games for real financial incentives. Nay (2014) gathered and integrated data from many experiments, conducted

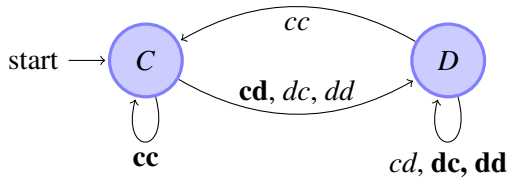
raw string: 0, 1, 0, 0, 1, 0, 0, 0, 0, 0
 encoded av, sm: 0, 1 0, 0, 1, 0, 0, 0, 0, 0
 decoded av, sm: 0, 1 0, 0, 1, 1, 1, 1, 1, 1

Action vector: 1, 2

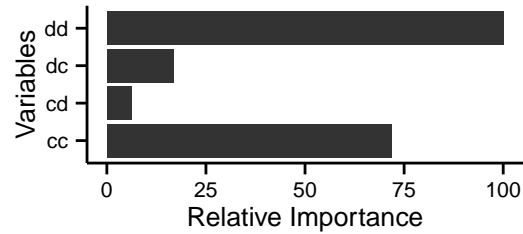
State	cc	dc	cd	dd
C (1)	1	2	2	2
D (2)	<i>1</i>	2	<i>2</i>	2

(a) Raw string; string split into Gray-encoded action vector (av) and state matrix (sm); and decoded binary representations. Decoded sm consists of column-wise elements that index into the action vector to determine the action for each state.

(b) Bitstring decoded into action vector and state matrix. For this action vector, 1 corresponds to cooperation (C) and 2 to defection (D). Columns of state matrix correspond to the observed behaviors at $t - 1$ and rows correspond to the state.



(c) Corresponding graph representation.



(d) Relative variable importance. Most important is 100.

Figure 3: A finite state machine estimated on 108,305 decisions with 82% accuracy on a hold-out test data set of 27,083 decisions. Transitions that would be accessible in strictly deterministic play are represented in boldface and inaccessible transitions in italic. Because the human players did not follow exact deterministic strategies, and the italicized transitions were taken in simulating this model with actual game play, the values of these transitions were identifiable.

and analyzed by Bereby-Meyer and Roth (2006), Duffy and Ochs (2009), Kunreuther et al. (2009), Dal Bo and Frechette (2011) and Fudenberg et al. (2012). All of the experiments share the same underlying repeated Prisoner’s Dilemma structure, although the details of the games differed. Nay’s data set comprises 135,388 cooperation decisions, which is much larger than previous studies of repeated game strategies.

Fudenberg et al. (2012) and Dal Bo and Frechette (2011) modeled their IPD experimental data with repeated game strategies; however, they applied a maximum likelihood estimation process to estimate the prevalence of a relatively small predefined set of strategies. In contrast, our estimation process automatically searches through a very large parameter space that includes all possible strategies up to a given number of states and does not require the analyst to predefine any strategies, or even understand the game.

We used 80% of our data for training and reserved the other 20% as a hold-out test set. Fig. 3 shows different representations of the the fittest two-state machine of a GA population evolved on the training data: The raw Gray-encoded and binary string (Fig. 3a), the bitstring decoded into state matrix and action vector form (Fig. 3b), and the corresponding graph representation (Fig. 3c). We measure variable importance (Fig. 3d) by switching each value of an estimated model’s state matrix to another value in its feasible range, measuring the decrease in goodness of fit to the training data, normalizing the values, then summing across each column to estimate the relative importance of each predictor variable (in this case, the moves each player made in the previous turn).

Fig. 4 illustrates the GA run that evolved the FSM of Fig. 3 by predicting cooperation decisions in IPD training data games. This GA run, which only took a few seconds on a modest laptop, used common algorithm settings: a population of 175 FSMs initialized with random bitstrings. If the analyst

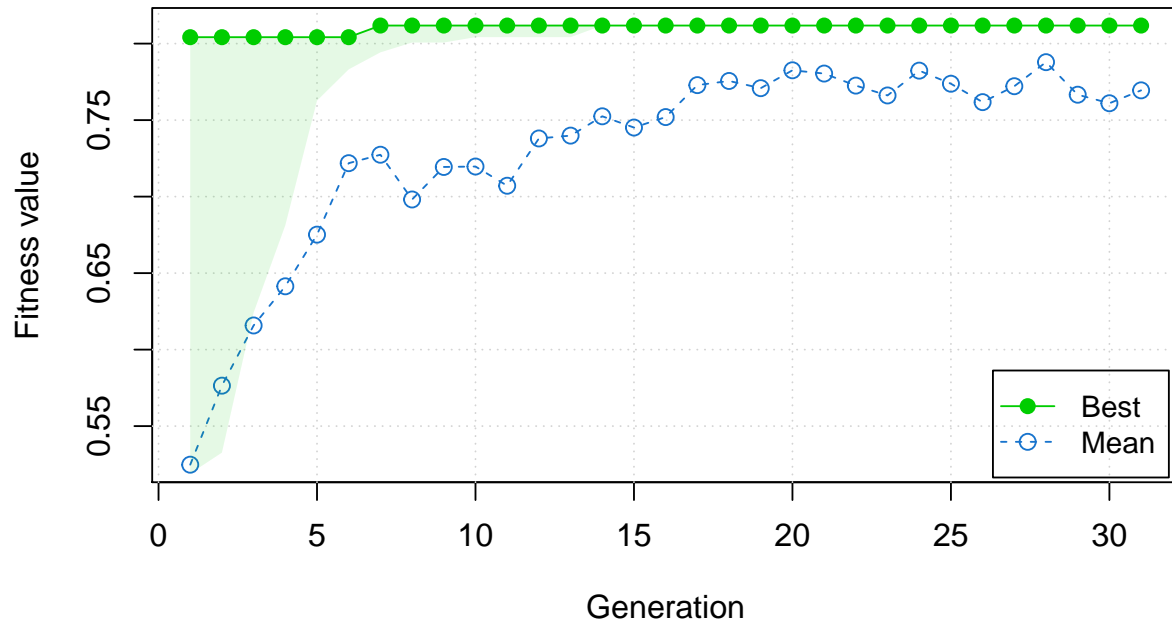


Figure 4: Proportion of correct predictions of human decisions in IPD games by an evolving population of decision models (strategies). On average, a random choice strategy will have 50% accuracy. Shaded area is the difference between the median of the population of strategies and the best individual strategy.

has an informed prior belief about the subjects' decision models, she can initialize the population with samples drawn from that prior distribution, but this paper focuses on deriving useful results from random initializations, corresponding to uniform priors, where the analyst only provides data. A linear-rank selection process used the predictive ability of individuals to select a subset of the population from which to create the next generation. A single-point crossover process was applied to the binary values of selected individuals with 0.8 probability, uniform random mutation was conducted with probability 0.1, and the top 5% fittest individuals survived each generation without crossover or mutation, ensuring that potentially very good solutions would not be lost (Scrucca 2013). These are standard GA parameter settings and can be adjusted if convergence is taking particularly long for a given dataset.

Using theoretical agent-based simulations and a fitness measure that is a function of simulated payoffs, Axelrod (1997) demonstrated the fitness of the tit-for-tat (TFT) strategy. Using a fitness measure that is a function of the ability to explain human behavior, we discovered a hybrid of TFT and grim trigger (GT), which we call noisy grim (NG). TFT's state is determined solely by the opponent's last play. GT will never exit a defecting state, no matter what the opponent does.

With traditional repeated game strategies such as TFT and GT, the player always takes the action corresponding to her current state (boldface transitions in Fig. 3c), but if we add noise to decisions so the player will sometimes choose the opposite action from her current state (italic transitions in Fig. 3c), then the possibility arises for both the player and opponent to cooperate when the player is in the defecting state (i.e. to reach the second row first column position of the state matrix in Fig. 3b). This would return the player to the cooperating state (see, e.g., Chong and Yao (2005)). Noisy grim's predictions on the hold-out test data are 82% accurate, GT's are 72% accurate, and TFT's are 77% accurate. We also tested

16 other repeated game strategies for the IPD from (Fudenberg et al. 2012). Their accuracy on the test set ranged from 46% to 77%. Our method uncovered a deterministic dynamic decision model that predicts IPD play better than all of the existing theoretical automata models of IPD play that we are aware of and has interesting relationships to the two most well-known models: TFT and GT.

This process has allowed us to estimate a highly interpretable decision model (fully represented by the small image of Fig. 3c) that predicts most of the behavior of hundreds of human participants, merely by plugging in the dataset as input. We address the potential concern that the process is too tuned to this specific case study by inputting a very different dataset from the field of international relations and obtaining useful results. However, before moving to more empirical data—where the data-generating process can never be fully known—to test how robustly we can estimate a known model, we repeatedly simulate a variety of known data-generating mechanisms and then apply the method to the resulting choice data.

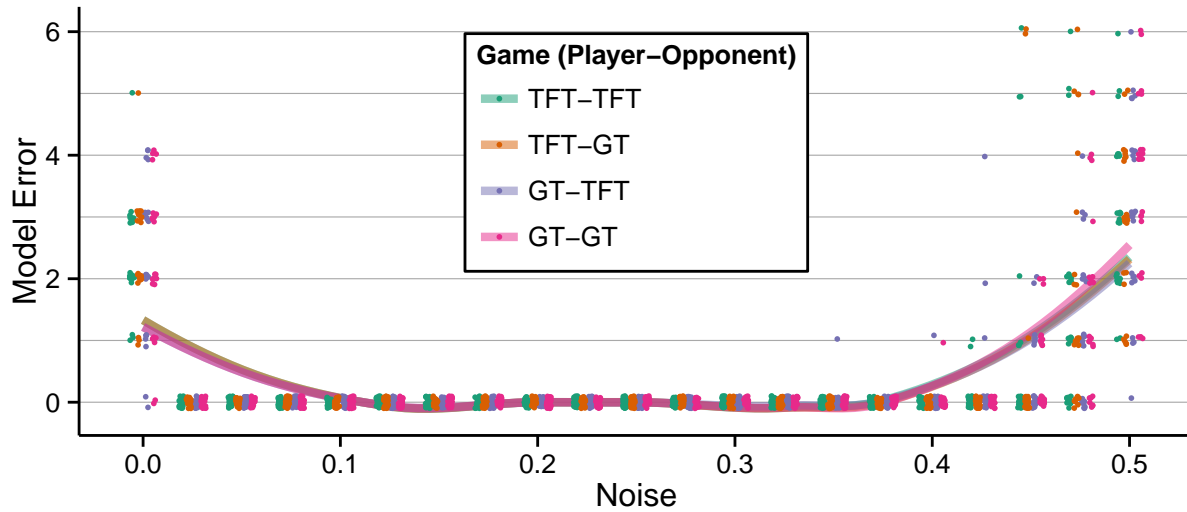
4 SIMULATED DATA

In the real world, people rarely strategically interact by strictly following a deterministic strategy (Chong and Yao 2005). Whimsy, strategic randomization, or error may induce a player to choose a different move from the one dictated by her strategy. To study whether our method could determine an underlying strategy that an agent would override from time to time, we followed the approach of Fudenberg et al. (2012) and created an agent-based model of the IPD in which agents followed deterministic strategies, such as TFT and GT, but made noisy decisions: At each time period, the deterministic strategy dictates each agent’s preferred action, but the agent will choose the opposite action with probability p , where p ranges from 0 (perfectly deterministic play) to 0.5 (completely random play). The noise parameter, p , is constant across all states of a strategy of a particular agent for any given simulation experiment we conducted.

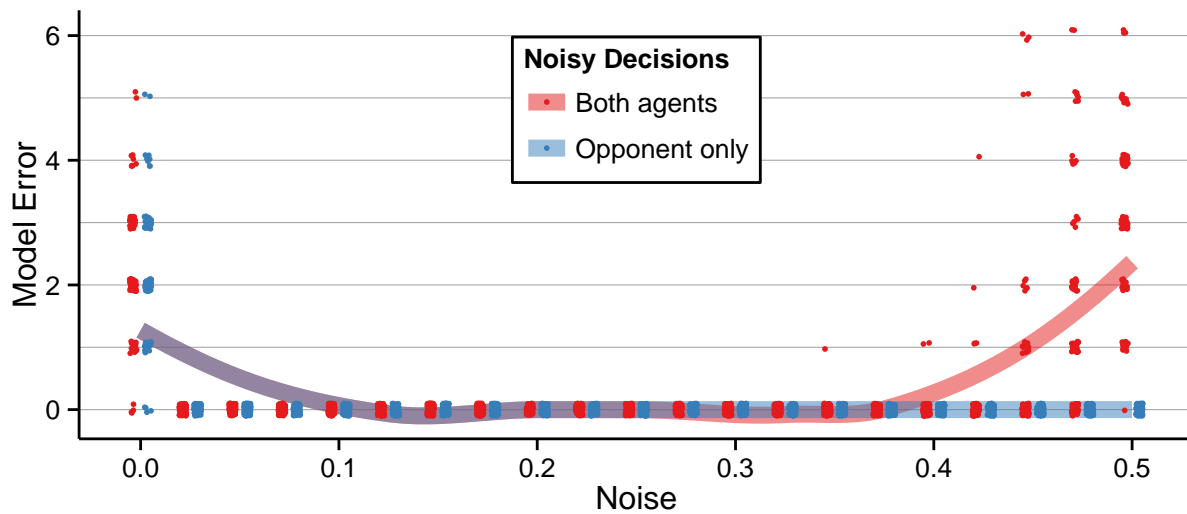
When a player follows an unknown strategy, characterized by latent states, discovering the strategy (the actions corresponding to each state and transitions between the states) requires observed data that explores as much as possible of the state transition matrix defined by all possible combinations of state and predictor values (for these strategies the predictors are the history of play). Many deterministic strategy pairings can quickly reach equilibria in which players repeat the same moves for the rest of the interaction. If the player and opponent both use TFT and both make the same first move, every subsequent move will repeat the first move. If the opponent plays GT, then after the first time the player defects the opponent will defect for the rest of the session and the data will provide little information on the player’s response to cooperation by the opponent. However, if the opponent plays with noise, the play will include many instances of cooperation and defection by the opponent, and will thus sample the accessible state space for the player’s strategy more thoroughly than if the opponent plays deterministically. Indeed, this is why Fudenberg et al. (2012) added noise to action choices in their human subjects experimental games.

We simulated approximately 17 million interactions, varying paired decision models of each agent [(TFT, TFT), (TFT, GT), (GT, TFT), (GT, GT)] and also varying the noise parameter p (0, 0.025, ..., 0.5) for each of two noise conditions: where both players made equally noisy decisions, and where only the opponent made noisy decisions while the player under study strictly followed a deterministic strategy. We ran 25 replicates of each of the 168 experimental conditions, with 4,000 iterations of game play for each replicate, and then applied the FSM estimation method to each replicate of the simulated choice data to estimate the strategy that the agent player under study was using.

Being in state/row k (e.g. 2) corresponds to the player taking action k (e.g. D) in the current turn. All entries in row k corresponding to the player taking action k in the current period (e.g. columns 2 and 4 for D) are identifiable. Entries in row k that correspond to not taking action k in the current period (e.g. columns 1 and 3 for row 2) represent transitions that cannot occur in strictly deterministic play, so their values cannot affect play and thus cannot be determined empirically. We take this into account when testing the method’s ability to estimate underlying deterministic models: this is why only 6 elements of a 10-element TFT or GT matrix can be identified (Fig. 5). We also take this into account when estimating models from empirical data, where the data-generating process is assumed to be stochastic: each element



(a) Effect of decision models and noise in choice on ability to estimate true model when a noisy player faces an equally noisy opponent. Color indicates the pair of decision models used by the player and opponent. For noise probabilities between 0.025 and 0.325 the estimated model always exactly matches the true model.



(b) Comparison of a purely deterministic player facing a noisy opponent versus a noisy player facing a noisy opponent. As noise approaches 0.5, the decision model's impact on actual choices vanishes, so the model error rises sharply when the player's decisions are noisy.

Figure 5: Role of noise in correctly identifying the player's strategy. Points have been jittered for visualization. Lines are LOESS-curves. Noise represents the probability, in any period, that an agent will take the opposite action from its decision model. Both the GT and TFT strategies can be represented as a 10 element matrix, of which 6 elements have identifiable values. The model error represents the number of mismatched elements between the estimated and true matrices, so it ranges from 0 (best) to 6 (worst).

of the matrix that would be inaccessible under deterministic play is identified, and the fitness is calculated with a strategy matrix in which that element is changed to its complement (“flipped”). If flipping the element does not change the fitness, then the two complementary strategies are indistinguishable and the element in question cannot be determined empirically. If each element decreases the fitness when it is flipped, then the strategy corresponds to a deterministic approximation of a stochastic process and all of the elements of the state matrix can be identified.

When the noise parameter was zero, most of the models estimated by the GA had at least two incorrect elements. However, for moderate amounts of noise ($p = 0.025$ – 0.325), all of the models estimated by the GA were correct (see Fig. 5a). For noise levels above $p = 0.325$ in the player, the amount of error rose rapidly with p , as expected because at $p = 0.5$ the action the player chooses moves completely at random so there is no strategy to discover. When a strictly deterministic player faced a noisy opponent, the GA correctly identified the player’s strategy for all noise levels above $p = 0.025$ (see Fig. 5b).

5 OBSERVATIONAL DATA

In order to extend this method to more complex situations the predictor variables (columns of the state matrices) can include any time-varying variable relevant to an agent’s decision. In context-free games such as the IPD, the only predictor variables are the moves the players made in the previous turn, but models of strategic interactions in context-rich environments may include other relevant variables.

We find it difficult to interpret graphical models with more than four predictors, but an analyst who had many potentially relevant predictor variables and was unable to use theory alone to reduce the number of predictors sufficiently to generate easily interpretable models with our method could take four courses of action (listed in order of increasing reliability and computation time):

1. Before FSM estimation, apply a (multivariate or univariate) predictor variable selection method.
2. Before FSM estimation, estimate an arbitrary predictive model that can produce variable importance rankings and then use the top $p < 4$ predictors for FSM estimation.
3. After FSM estimation with $p \geq 4$ predictors, inspect the returned predictor variable importance ranking, and remove all but the top $p < 4$ from her dataset and re-run estimation.
4. Conduct FSM estimation with all combinations of $p < 4$ predictors out of all relevant predictors and choose the estimated model with the best performance (usually highest out-of-sample accuracy).

We illustrate the use of extra predictor variables by applying our method to an example from international relations involving repeated water management-related interactions between countries that share rivers. We use data compiled by Brochmann (2012) on treaty signing and cooperation over water quality, water quantity, and flood control from 1948–1999 to generate a model for predicting whether two countries will cooperate. We used three lagged variables: whether there was water-related conflict between them in the previous year, whether they cooperated around water in the previous year, and whether they had signed a water-related treaty during any previous year. This data set was too small to divide into training and hold-out subsets for assessing predictive accuracy, so we report models’ accuracy in reproducing the training data (a random choice model is 50% accurate). A two-state decision model (Fig. 6a) is 73% accurate, a three-state model (Fig. 6c) is 78% accurate, and a four-state model is 82% accurate, but its complexity makes it difficult to interpret visually so it is not shown.

Accuracy can be a problematic measure when the classes are imbalanced, i.e. if a class the model is trying to predict is rare. Many alternatives to accuracy are available that illuminate different aspects of predictive power. For instance, precision is the proportion of (cooperation) event signals predicted by our models that are correct and recall is the proportion of events that are predicted by our models. For this subset of the dataset, cooperate and not cooperate were almost evenly distributed and to maintain a comparison to the experimental and simulated data we used accuracy as the fitness measure.

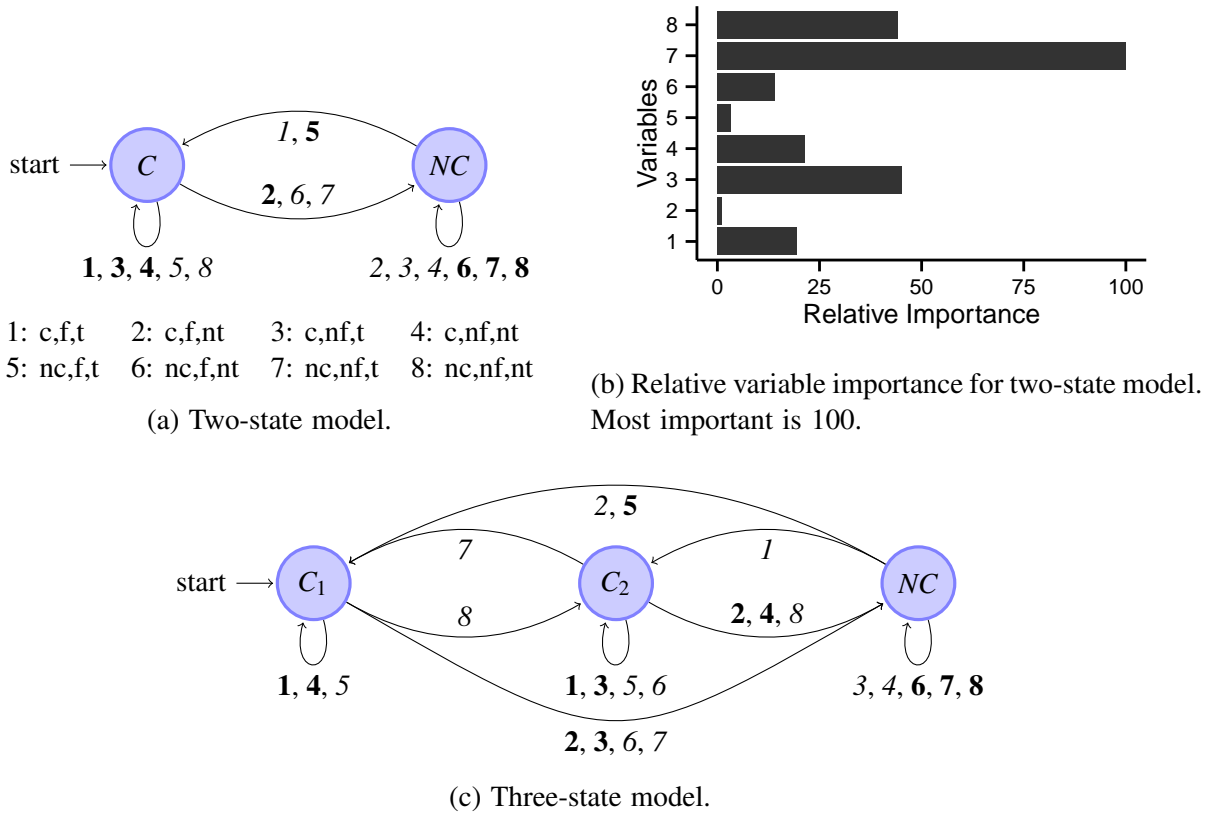


Figure 6: Models of cooperation over water issues for sets of countries. The italicized transitions would not be feasible outcomes if this strategy was played deterministically, but they are realized in the data. The inputs (c,nc) represent cooperation and non-cooperation, respectively; (f,nf) conflict and non-conflict; and (t,nt) signing and not signing a treaty.

In the two-state model, whether or not the countries cooperated in the previous year, the combination of conflict and treaty-signing in the previous year always produces cooperation, whereas conflict without treaty-signing in the previous year always produces non-cooperation. In the three-state model, three of the four outcomes that include conflict lead to a transition from non-cooperation to cooperation, and four of the six outcomes that cause transitions from cooperation (states C_1 and C_2) to non-cooperation are non-conflict outcomes. While this does not tell us something decisive about the role of conflict, it suggests that there may be a counter-intuitive role of conflict in promoting cooperation. Brochmann (2012, p. 159), using a bivariate probit simultaneous equation model, has a similar finding: “In the aftermath of conflict, states may be particularly eager to solve important issues that could cause future problems.”

6 DISCUSSION

This paper outlined a method for estimating interpretable models of dynamic decision-making. By estimating a global, deterministic, simple function for a given dataset, imposing constraints on the number of predictor variables, and providing options for reducing the number of predictor variables, our process facilitates capturing a significant amount of information in a compact and useful form. The method can be used for designing empirically grounded agent models in agent-based simulations and for gaining direct insight into observed behaviors of real agents in social and physical systems. Combining state matrices and a genetic algorithm has proven effective for simulated data, experimental game data, and observational

international relations data. With the simulated data, we successfully recovered the exact underlying models that generated the data. With the real data, we estimated simple deterministic approximations that explain most of the structure of the unknown underlying process. We discovered a theoretically interesting dynamic decision model that predicted IPD play better than all of the existing theoretical models of IPD play that we were aware of.

We have released [an open-source R package](#) that implements the methods described here to estimate any time series classification model that uses a small number of binary predictor variables and moves back and forth between the values of the outcome variable over time. Larger sets of predictor variables can be reduced to smaller sets by applying one of the four methods outlined in Section 5. Although the predictor variables must be binary, a quantitative variable can be converted into binary by division of the observed values into high/low classes. Future releases of the package may include additional estimation methods to complement GA optimization.

ACKNOWLEDGMENTS

We gratefully acknowledge the authors of [R](#) (R Core Team 2014). This manuscript was prepared using knitr (Xie 2014). We would like to thank Yevgeniy Vorobeychik for discussions on predicting games.

This work was supported by U.S. National Science Foundation grants EAR-1416964 and EAR-1204685.

REFERENCES

- Arifovic, J. 1994. “Genetic Algorithm Learning and the Cobweb Model”. *Journal of Economic Dynamics and Control* 18:3–28.
- Arifovic, J., and C. Eaton. 1995. “Coordination via Genetic Learning”. *Computational Economics* 8:181–203.
- Axelrod, R. M. 1984. *The Evolution of Cooperation*. New York: Basic Books.
- Axelrod, R. M. 1997. *The Complexity of Cooperation: Agent-based Models of Competition and Collaboration*. Princeton: Princeton University Press.
- Bereby-Meyer, Y., and A. E. Roth. 2006. “The Speed of Learning in Noisy Games: Partial Reinforcement and the Sustainability of Cooperation”. *The American Economic Review* 96:1029–1042.
- Brochmann, M. 2012. “Signing River Treaties: Does It Improve River Cooperation?”. *International Interactions* 38:141–163.
- Bullard, J., and J. Duffy. 1999. “Using Genetic Algorithms to Model the Evolution of Heterogeneous Beliefs”. *Computational Economics* 13:41–60.
- Camerer, C. F. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton: Princeton University Press.
- Chong, S., and X. Yao. 2005. “Behavioral Diversity, Choices and Noise in the Iterated Prisoner’s Dilemma”. *IEEE Transactions on Evolutionary Computation* 9:540–551.
- Dal Bo, P., and G. R. Frechette. 2011. “The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence”. *American Economic Review* 101:411–429.
- Duffy, J. 2006. “Agent-Based Models and Human Subject Experiments”. In *Handbook of Computational Economics*, Volume 2, 949–1011. Amsterdam: Elsevier.
- Duffy, J., and J. Engle-Warnick. 2002. “Using Symbolic Regression to Infer Strategies from Experimental Data”. In *Evolutionary Computation in Economics and Finance*, 61–82. New York: Springer.
- Duffy, J., and J. Ochs. 2009. “Cooperative Behavior and the Frequency of Social Interaction”. *Games and Economic Behavior* 66:785–812.
- Eddelbuettel, D. 2013. *Seamless R and C++ Integration with Rcpp*. New York: Springer.
- Fogel, D. B. 1993. “Evolving Behaviors in the Iterated Prisoner’s Dilemma”. *Evolutionary Computation* 1:77–97.
- Fudenberg, D., D. G. Rand, and A. Dreber. 2012. “Slow to Anger and Fast to Forgive: Cooperation in an Uncertain World”. *American Economic Review* 102:720–749.

- Goldberg, D. E., and J. H. Holland. 1988. "Genetic Algorithms and Machine Learning". *Machine Learning* 3:95–99.
- Hanaki, N. 2004. "Action Learning Versus Strategy Learning". *Complexity* 9:41–50.
- Hanaki, N., R. Sethi, I. Erev, and A. Peterhansl. 2005. "Learning Strategies". *Journal of Economic Behavior & Organization* 56:523–542.
- Hardin, G. 1968. "The Tragedy of the Commons". *Science* 162:1243–1248.
- Hastie, T., R. Tibshirani, and J. Friedman. 2009. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Second Edition*. 2nd ed. New York, NY: Springer.
- Koza, J. R. 1992. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. Cambridge, MA: Bradford.
- Kunreuther, H., G. Silvasi, E. T. Bradlow, and D. Small. 2009. "Bayesian Analysis of Deterministic and Stochastic Prisoner's Dilemma Games". *Judgment and Decision Making* 4:363–384.
- Marks, R. E., D. F. Midgley, and L. G. Cooper. 1995. "Adaptive Behaviour in an Oligopoly". In *Evolutionary Algorithms in Management Applications*, edited by J. Biethahn and V. Nissen, 225–239. New York: Springer.
- McKelvey, R. D., and T. R. Palfrey. 2001. "Playing in the Dark: Information, Learning, and Coordination in Repeated Games". Technical report, California Institute of Technology, Pasadena.
- Midgley, D. F., R. E. Marks, and L. C. Cooper. 1997. "Breeding Competitive Strategies". *Management Science* 43:257–275.
- Miller, J. H. 1996. "The Coevolution of Automata in the Repeated Prisoner's Dilemma". *Journal of Economic Behavior & Organization* 29:87–112.
- Miller, J. H., and S. E. Page. 2007. *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton: Princeton University Press.
- Moore, E. 1956. "Gedanken-experiments On Sequential Machines". *Automata Studies* 34:129–153.
- Nay, J. J. 2014, December. "Predicting Cooperation and Designing Institutions: An Integration of Behavioral Data, Machine Learning, and Simulation". In *Winter Simulation Conference Proceedings*. Savannah, GA.
- Osborne, M. J., and A. Rubinstein. 1994. *A Course in Game Theory*. Cambridge, MA: MIT Press.
- R Core Team 2014. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Rubinstein, A. 1986, June. "Finite automata play the repeated prisoner's dilemma". *Journal of Economic Theory* 39 (1): 83–96.
- Savage, C. 1997, January. "A Survey of Combinatorial Gray Codes". *SIAM Review* 39 (4): 605–629.
- Scrucca, L. 2013. "GA: A Package for Genetic Algorithms in R". *Journal of Statistical Software* 53:1–37.
- Xie, Y. 2014. *Dynamic Documents with R and knitr*. Boca Raton: Chapman & Hall/CRC.

AUTHOR BIOGRAPHIES

JOHN J. NAY is a Ph.D. Candidate at Vanderbilt University. He develops and applies machine learning and simulation methods to better predict behavior in complex systems. His research involves collaboration with social scientists, environmental scientists, computer scientists and engineers. His email address is john.j.nay@vanderbilt.edu, and website: <https://johnnay.github.io/>.

JONATHAN M. GILLIGAN is Associate Professor of Earth and Environmental Sciences at Vanderbilt University and Associate Director for Research at the Vanderbilt Climate Change Research Network. His research and teaching interests center on integrating natural science, social science, behavioral science, and engineering to study coupled human-natural systems. He holds a Ph.D. in Physics from Yale University. His email address is jonathan.gilligan@vanderbilt.edu, and website: <https://my.vanderbilt.edu/jonathangilligan/>.