

Grounding Social Interaction with Affective Intelligence

Joshua D. A. Jung¹, Jesse Hoey¹, Jonathan H. Morgan²,
Tobias Schröder³, and Ingo Wolf³

¹ University of Waterloo, Dept. of Computer Science, Waterloo, Ontario, Canada
{j35jung, jhoey}@cs.uwaterloo.ca,

² Duke University, Dept. of Sociology, Durham, NC, USA
jonathan.h.morgan@duke.edu

³ University of Applied Sciences, Institute for Urban Futures, Potsdam, Germany
{schroeder, wolf}@fh-potsdam.de

Abstract. Symbolic interactionist principles of sociology are based on the idea that human action is guided by culturally shared symbolic representations of identities, behaviours, situations and emotions. Shared linguistic, paralinguistic, or kinesic elements allow humans to coordinate action by enacting *identities* in social situations. Structures of identity-based interactions can lead to the enactment of social orders that solve social dilemmas (e.g., by promoting cooperation). Our goal is to build an artificial agent that mimics the identity-based interactions of humans. This paper describes a study in which humans played a repeated prisoner’s dilemma game against other humans or one of three artificial agents (bots). One of the bots has an explicit representation of identity and demonstrates more human-like behaviour than the other bots.

1 Introduction

The prisoner’s dilemma has long been studied, starting with the work of Axelrod [3]. Recent work has looked at modelling both rational choice and social imitation to simulate more human-like behaviour in networked PD games [16]. Others have looked at using emotional signals to influence play in PD games, for example by changing expectations of future games [4]. Emotions have also been linked with intrinsic reward and exploration bonuses [14]. It has become increasingly clear that human handling of an infinite action space (not limited to the realm of the prisoner’s dilemma) may be governed largely by affective processes [1,10]. Shared affective structures allow agents to focus on the subset of possibilities that provide interactions aligning with the shared structure. This subset of possibilities forms the set of “cultural expectations” for behaviours that are “rational relative to the social conventions and ethics” ([1], p200).

A recent product of these ideas is *BayesAct* [8], which models the emotional control of social interaction by humans and can explain the emergence of stable role relations and patterns of interaction [13]. Here, we empirically study the class of interactions in the iterated prisoner’s dilemma, a fundamental paradigm

in the social sciences aimed at understanding the dynamics of human cooperation vs. competition. Our results are encouraging in terms of supporting the validity of the *BayesAct* agent as a mechanistic model of human social interactions.

BayesAct [2,7,8,13] is a partially observable Markov decision process model of affective interactions between a human and an artificial agent. *BayesAct* arises from the sociological (symbolic interactionist) “Affect Control Theory” (ACT) [6]. *BayesAct* generalises this theory by modeling affective states as probability distributions, and allowing decision-theoretic reasoning about affect. *BayesAct* proposes that humans learn and maintain a set of *shared* cultural affective *sentiments* about people, objects, behaviours, and about the dynamics of interpersonal events. Humans use a simple affective mapping to appraise individuals, situations, and events as sentiments in a three dimensional vector space of evaluation (E: good vs. bad), potency (P: strong vs. weak) and activity (A: active vs. inactive). These “EPA” mappings can be measured, and the culturally shared consistency has repeatedly been demonstrated to be extremely robust in large cross-cultural studies [12]. Many believe this consistency “gestalt” is a keystone of human intelligence. Humans use it to make predictions about what others will do, and to guide their own behaviour. Further, it defines an affective heuristic (a *prescription*) for making decisions quickly in interactions. Humans strive to achieve consistency by choosing actions that maximally increase *alignment* (decrease *deflection* in ACT terms) in shared affective cultural sentiments. The shared sentiments and dynamics, the affective prescriptions, and the resulting *affective ecosystem* of vector mappings, result in an equilibrium or *social order* [5], which is optimal for the group as a whole, rather than for individual members. Humans living at the equilibrium “feel” good and want to stay there, with positive evolutionary consequences. However, agents with sufficient resources can plan beyond the prescription, allowing them to manipulate other agents to achieve individual profit in collaborative games [2].

For example, in the repeated prisoner’s dilemma, cooperation has a different emotional signature than defection: it is usually viewed as nicer (higher evaluation). Rationality predicts an agent will try to optimize over his expected total payout, perhaps modifying this payout by some additional intrinsic reward for altruism. The *BayesAct* view is quite different: it says that an agent will take the most aligned action given her estimates of her own and her partner’s affective *identity*. Thus, *friends* will do nice things to *friends* and cooperate, but will be more likely to defect against a *scrooge* or a *traitor*. *Scrooges* will defect, as this is consistent with a more negative identity, but may cooperate to manipulate.

As elucidated by Squazzoni [15], models of social networks must take into account the heterogeneity of individuals, behaviours, and dynamics in order to better account for the available evidence. In this paper we argue that the principles encoded in *BayesAct* can capture this heterogeneity. As evidence, here we present results from an experiment in which participants played a repeated prisoner’s dilemma (PD) game against each other and against a set of computer programs, one of them *BayesAct*.

2 Experiments and Results

The prisoner’s dilemma is a classic two-person game in which each person can either *defect* by taking \$1 from a (common) pile, or *cooperate* by giving \$2 from the same pile to the other person. There is one Nash equilibrium in which both players defect, but when humans play the game they often are able to achieve cooperation. A rational agent will optimise over his expected long-term payoffs, possibly by averaging over his expectations of his opponent’s type (or strategy).

A *BayesAct* agent computes what *affective* action (an EPA vector) is prescribed in the situation, given his estimates of his and the other’s (called the *client*) identities, and of the affective dynamics, and then seeks the propositional action ($\in \{cooperate, defect\}$) that, according to a stored cultural definition, is most consistent with the prescribed affective action. As the game is repeated, the *BayesAct* agent updates his estimates of identity (for self and other/*client*), and adjusts his play accordingly.

For example, if *agent* thought of himself as a *friend* (EPA: {2.75, 1.88, 1.38}) and knew the other agent to be a *friend*, the deflection minimizing action would likely be something good (high E). Indeed, a simulation shows that one would expect a behaviour with EPA={1.98, 1.09, 0.96}, with closest labels such as *treat* or *toast*. Intuitively, cooperate seems like a more aligned propositional action than defect. This intuition is confirmed by the distances from the predicted (affectively aligned) behaviour to *collaborate with* (EPA: {1.44, 1.11, 0.61}) and *abandon* (EPA: {-2.28, -0.48, -0.84})⁴ of 0.4 and 23.9, respectively, clearly showing the closer proximity of collaboration to this affectively aligned action.

The *agent* will predict the *client*’s behavior using the same principle: compute the deflection minimising affective action, then deduce the propositional action based on that. Thus, a *friend* would predict that a *scrooge* would defect, but would still want to cooperate in order to *reform* or *befriend* the other agent. If a *BayesAct* agent has sufficient resources, he could search for an affective action near to his optimal one, but that would still allow him to defect. Importantly, he is *not trading off costs in the game with costs of disobeying the social prescriptions*: his resource bounds and action search strategy are preventing him from finding the more optimal (individual) strategy, implicitly favouring those actions that benefit the group and solve the social dilemma.

In order to compare the predictions of *BayesAct* to human play, we recruited 70 students (55 male and 15 female) from a senior undergraduate class on artificial intelligence at the University of Waterloo⁵. The participants played a total of 360 games in a computer lab environment. The length of each game was randomly chosen between 12 – 18 rounds (plays of cooperation or defection). Each game a participant played was against either (1) another randomly chosen participant; (2) an automated *tit-for-tat* player; (3) a *BayesAct* agent as described above; or (4) a fixed strategy of cooperate three times followed by always defect,

⁴These are representative of the affective meaning of the actions in the game [2]

⁵The study was reviewed and approved by the UW Office of Research Ethics. For further discussion of experiment procedures, see [9].

hereafter referred to as *jerkbot*. The *BayesAct* agent reward is only over the game (e.g. 2, 1, or 0), and we use a two time-step game in which both *agent* and *client* choose their actions at the first time step, and then communicate this to each other on the second step.

Participants were assigned some order in which to play each opponent, but that order was randomized for each participant. Further, participants were told that all of their opponents were human. Upon sign-up, and after each game (of between 12–18 rounds), participants were asked the following by providing them with a slider for each E,P,A dimension, known as a semantic differential [6]:

- how they felt about the plays in the game (take 1 or give 2), out of context. *BayesAct* agents then interpret the affective signature of actions in the game by comparing the EPA vectors to these two vectors.
- how they felt about themselves (their self identity). This gives *BayesAct* its self-identity, as we want it to replicate a participant. We use the raw data from all student responses across all questions as this self-identity *BayesAct*.
- how they felt about their opponent in the game they just played. Before the first game we asked them how they felt about a generalised opponent in this game, giving the *BayesAct* client identity.

A total of 89 samples were used for identities (resampled to get N=2000 samples used in the *BayesAct* particle filter) and an average of 89 samples used for the SCB. From this sample, we measured for *Give 2* an EPA of {1.4, 0.10, 0.18}, and for *Take 1*, {-0.65, 0.85, 0.70}. *Take 1* is seen as more negative and more powerful and active. Additionally, the self is seen as more positive than the opponent or “other” (with average E value 1/0.25 for self/other), but about the same power (0.56/0.64) and activity (0.41/0.33).

Table 1. Summary statistics. coops: number of cooperations after 10th game.

opponent	num. avg. game		agent (human)		client (human or bot)	
	games	length	payoff	coops	payoff	coops
jerkbot	83	15.01	15.86 ± 3.00	0.09 ± 0.24	22.33 ± 6.00	0.00 ± 0.00
bayesact	73	14.85	27.05 ± 5.92	0.54 ± 0.40	22.19 ± 7.87	0.69 ± 0.32
human	35	15.43	24.11 ± 7.55	0.56 ± 0.45	26.00 ± 5.92	0.51 ± 0.47
titfortat	82	14.82	27.66 ± 5.39	0.81 ± 0.35	26.96 ± 6.00	0.83 ± 0.34

Table 1 (cols 2,3) shows the statistics of game numbers and lengths against the different opponents. Figure 1 shows the mean, standard deviation, and median reward gathered at each step of the game, for each of the opponents. The blue lines show the human play, while the red lines show the opponent (one of human, *BayesAct*, *tit-for-tat*, or *jerkbot*). We see that humans mostly manage to cooperate together until about 4-5 games before the end. The *tit-for-tat* strategy ensures more even cooperation, but is significantly different from humans. *Jerkbot* is obvious, as a few defections after three games convinces the human to




Fig. 1. Blue=human; Red=agent (human, bayesact, titfortat and *jerkbot*); dashed=std.dev.; solid (thin, with markers): mean; solid (thick): median.

defect thereafter. The *BayesAct* agent play is very similar to the human play, but the human participants take advantage of the *BayesAct* agents late in the game. This may be because the *BayesAct* agent is using a short (5 second) planning timeout, and we would need to compare to a zero timeout (so only using the ACT prescriptions) and to longer timeouts to see how this behaviour changes.

Table 2. Means of pre-game and post-game impressions for each opponent type.

opponent	give 2			take 1			self (human)			other (human/bot)		
	E	P	A	E	P	A	E	P	A	E	P	A
(initial)	1.4	0.1	0.2	-0.6	0.9	0.6	1.1	0.6	0.3	0.2	0.6	0.3
jerkbot	1.3	-0.3	-0.1	-1.3	0.8	0.7	1.3	-0.1	0.9	-1.9	0.4	0.5
bayesact	1.3	0.1	0.0	-0.9	1.1	1.0	0.7	1.4	1.2	0.4	-0.1	-0.3
human	1.7	0.7	0.3	-1.2	0.4	0.3	1.5	1.2	1.0	0.5	0.0	0.1
titfortat	2.3	1.2	1.1	-1.2	0.5	0.3	1.9	1.7	1.7	2.2	1.1	1.1

To further investigate the differences between the different opponents, we measure the mean fraction of cooperative actions on the part of the human after (and including) the 10th game (see Table 1). We find that, when playing against another human, humans cooperate in 0.56 ± 0.45 of these last games. This number was almost the same when playing *BayesAct* agent at 0.54 ± 0.40 . Against *tit-for-tat*, there was much more cooperation (0.81 ± 0.35). Finally, against *jerkbot*, it was very low 0.09 ± 0.24 . We also computed the mean EPA ratings of the self and other after each game, as shown in Table 2. We found that *jerkbot* (EPA:{-1.9, 0.4, 0.5}) is seen as much more negative, and *tit-for-tat* (EPA:{2.2, 1.1, 1.1}) much more positive, than human (EPA:{0.5, 0.0, 0.1}) or *BayesAct* (EPA:{0.4, -0.1, -0.3}), and that the human participants felt less powerful when playing *jerkbot* (EPA of self:{1.3, -0.1, 0.9}) than when playing *BayesAct* (EPA of self:{0.7, 1.4, 1.2}), or another human (EPA of self:{1.5, 1.2, 1.0}). Human participants felt more powerful, positive and active when playing *tit-for-tat* (EPA of self:{1.9, 1.7, 1.7}).

3 Conclusion

We have presented a model for affectively guided play in the prisoner’s dilemma. Our aim is to design agents that are human-like in their behaviours using symbolic interactionist principles, which prescribe socially expected actions given the identities of the actor and her opponent. In this paper, we have shown how these principles result in more human-like play in the iterated prisoner’s dilemma. We are currently running simulations of *BayesAct* agents (learned from human data) in a networked prisoner’s dilemma setting. Other research avenues include assistive technologies [11], intelligent tutoring [8] and other games [2].

References

1. Antonio, D.: Descartes’ error: Emotion, reason, and the human brain. GP PutnamSons, New York (1994)
2. Asghar, N., Hoey, J.: Monte-Carlo planning for socially aligned agents using Bayesian affect control theory. In: Proc. Uncertainty in Artificial Intelligence (UAI). pp. 72–81 (2015)
3. Axelrod, R., Hamilton, W.D.: The evolution of cooperation. *Science* 211(4489), 1390–1396 (1981)
4. De Melo, C.M., Carnevale, P., Read, S., Antos, D., Gratch, J.: Bayesian model of the social effects of emotion in decision-making in multiagent systems. Proc. AAMAS Volume 1. pp. 55–62.(2012)
5. Erving Goffman. *Behavior in Public Places*. The Free Press, New York, 1963.
6. Heise, D.R.: Expressive order: Confirming sentiments in social actions. Springer Science & Business Media (2007)
7. Hoey, J., Schröder, T.: Bayesian affect control theory of self. In: AAAI. pp. 529–536. Citeseer (2015)
8. Hoey, J., Schröder, T., Alhothali, A.: Affect control processes: Intelligent affective interaction using a partially observable Markov decision process. *Artificial Intelligence* 230, 134–172 (2016)
9. Jung, J.D.A., Hoey, J., Morgan, J.H., Schröder, T., Wolf, I.: Comparison of Affect-Control Theoretic Agents and Humans in the Prisoner’s Dilemma Tech. rep. CS-2015-18, University of Waterloo School of Computer Science (2015)
10. LeDoux, J.: The emotional brain: The mysterious underpinnings of emotional life. Simon and Schuster (1998)
11. Malhotra, A., Hoey, J., König, A., van Vuuren, S.: A study of elderly people’s emotional understanding of prompts given by virtual humans. In: Proc. 10th EAI Conference on Pervasive Computing Technologies for Healthcare (2015)
12. Osgood, C.E., May, W.H., Miron, M.S.: Cross-cultural universals of affective meaning. University of Illinois Press (1975)
13. Schröder, T., Hoey, J., Rogers, K.B.: Modeling dynamic identities and uncertainty in social interactions: Bayesian affect control theory. *In Press, Am. Soc. Rev.* (2016)
14. Sequeira, P., Melo, F.S., Paiva, A.: Learning by appraising: an emotion-based approach to intrinsic reward design. *Adaptive Behaviour* 22(5), 330–349 (2014)
15. Squazzoni, F., Jager, W., Edmonds, B.: Social simulation in the social sciences a brief overview. *Social Science Computer Review* 32(3), 279–294 (2014)
16. Vilone, D., Ramasco, J.J., Sánchez, A., San Miguel, M.: Social imitation versus strategic choice, or consensus versus cooperation, in the networked prisoner’s dilemma. *Physical Review E* 90(2), 022810 (2014)