

Can agents regulate fast emotions to be more sociable?

David C. Moffat
Digital Design Techniques
Glasgow Caledonian University, UK.
D.C.Moffat@gcu.ac.uk

Abstract. Agents that are required to interact with humans socially may need to simulate or emulate emotional behaviour, as well as understand the emotions behind human behaviour. However the way that human emotion operates is not well understood, and people often appear to act irrationally because of their emotions.

Research in psychology and economics laboratories has used formal games to investigate how people behave. One is the Ultimatum Game, in which players typically diverge from what a perfect "rational actor" model would predict, presumably because of their emotional reactions in the game.

This paper uses the Ultimatum Game as a test domain, and shows how a computational model of an agent may be developed that manifests social emotions in the game. Instead of a procedural approach, such as is used in many typical artificial agents, this model is built on a rule-based system architecture, in which emotional and rational actions are implemented similarly. As a result, the architecture promises to allow the agent to regulate its emotional reactions depending on circumstances.

When people play the UG, they often react angrily, in ways that are seen as irrational by the standard economics view of rationality as utility maximisation. This can be modeled as a case of vengeance, borne of moral outrage at being treated unfairly. The agent has rules which encapsulate this sense of fairness and which potentially cause a similarly "vengeful" reaction. Because the rules are similar to the other rules in the agent, and even the rules for vengeance are cast in the same formal scheme, the model shows how emotion and reason might be seen as much more compatible than they are usually thought to be.

The agent architecture thus throws some light on philosophical debates about the supposed irrationality of the passions, but it also offers possibilities for more elaborate modeling of emotional phenomena such as action readiness and regulation. In future if robots, and other agents that we wish to build, are going to be able to live amongst us as we hope, they will need emotions; but they will also have to regulate their emotions, much as we do.

1 Introduction

Are emotions rational? While there is much AI research into multi-agent systems, and robotics, the focus generally is on making them able to plan, act and communicate logically. There is research in the field of affective computing [6] that aims to make robots and other computer systems able to recognise human emotions, such as from facial expressions. It aims to make virtual agents express emotions too, with their artificial faces. However, there are fewer researchers who wish to make agents actually emotional in order to make them

more effective. The reason may be a general opinion that emotions are not wanted in an artificial, intelligent agent, because they are contrary to rationality and could make the agent less effective, instead of more so.

Against this view, it is obvious that natural agents, including humans and animals, are highly emotional beings. Far from handicapping us all, emotions have evolved in mammals principally, in order to increase their survival. As such, emotions can be seen generally as a rational design, even if not every occurrence of an emotion is beneficial.

Furthermore, emotions pervade our social life, colouring our speech and writing, except perhaps for the most formal prose; and they determine our personal relationships. In the case of formal or professional relationships too, or transient relationships with strangers, there is no call to exclude emotion. One could even say that there is no such thing as a proper human relationship without some emotion.

Therefore, we are called to devote more efforts to research into the nature of emotion, also in the areas of AI agents and robotics; and this, not only to make the artificial systems more effective if they would, but to make them much better able to navigate the unclear and stormy waters of social interaction. If agents remain ignorant and unfeeling of human emotion, then what chance have they of interacting with us as equals; or come to that, even as inferiors?

If we recognise that emotion permeates our social lives, and regulates our interactions with other people, then we could have little reason to leave artificial agents without the capacity. We should therefore study the ways in which human emotion regulates social interactions, and the fields of psychology and other behavioural sciences are a natural launch-pad for this.

In behavioural economics, for one, critical situations are formalised as artificial games, in laboratory experiments. The Ultimatum Game is used to examine some kinds of social behaviour, and will be used in this paper as the domain for a proposed agent player.

After describing the game, and its relevance for our purpose, we set out a design of a player of the game. The player agent has a rule-based system architecture, which is intended to allow emotions to integrate with and compete with more (conventionally) rational cognition. In consequence, the agent can potentially play the game in quite a human-like way, choosing similar moves in similar cases, and following both its emotional inclinations and its rational calculations of optimal play, alternately.

2 The Ultimatum Game

The Ultimatum Game (UG) is an artificial mathematical game that is used in laboratory experiments to probe participants' judgements of fairness in social interactions.

There are two players in the game: the *proposer* and the *responder*, and a sum of money that they have to split between them as follows. The proposer offers a split, which we may express as a percentage of the sum. The responder then chooses to accept the offer, or reject it. Accepting the offer means that both players get their part of the split; but rejecting it means that both get nothing.

For example, if the proposer offers 50% then the responder would surely accept it, and both players would get half the sum. But if the proposer offers much less, say only 4%, then any human responder is likely to reject it. It is easy to see why (if you are also human): the responder is angered by the tiny offer, in which proposer keeps nearly all the money for himself. However, that angry human has behaved irrationally, according to standard economic theory and mathematical game theory. The responder should accept any offer made to him, to maximise his gain in "utility", because even a tiny amount of money is better than nothing.

The fact that people are consistently and robustly "irrational" in this way is what makes the UG such an interesting game for researchers. Is it really true that humans are an inherently irrational species? Is it our emotions that make us irremediably irrational? Or is there something deeply wrong with standard economic theory?

There are some indications in the literature that it is indeed emotion to be blamed, and probably the emotions of anger or disgust. For example, dosing participants with the oxytocin before they play the UG makes them less likely to reject the offer [7]. As oxytocin is a hormone that fosters affiliative feelings in mammals, (and as we are mammals,) the suggestion is that responders feel more forgiving toward the proposers, and are thus less inclined to punish them.

It seems clear that the kind of economic and other transactions, that are intended to be modelled by the Ultimatum Game, are moderated by social emotions and attitudes on both sides.

3 Relevance of the UG for social emotions

The Ultimatum Game asks the players to consider what would be a fair offer, linking into their sense of morality and cultural norms. Both players are soon aware that an "unfair" (unequal) offer invites a rejection from the responder, and the more unfair (lower) it is, the more likely it is to be rejected [5].

Economic transactions like bartering might be seen as elaborate forms of the game, in which each side makes an offer or a counter offer, until a "fair" price is agreed. In heavily bartering cultures, the barter sets the tone of the transient relationship between buyer and seller, suggesting that the structure of the game would permit a kind of relationship between the players to form, at least if it were iterated with multiple rounds of play. Indeed, the UG is typically played in laboratory experiments as a single round, without further iterations, and the reason for this is specifically to avoid the possibility of reciprocation arising between the two players.

The sense of fairness, and its opposite, the sense of unfairness or exploitation, are part of the moral sense that all humans seem to have. Even in the animal kingdom, there are many species that live in social groups, in which relationships are maintained by exchange of resources or of symbolic tokens of respect. Examples would be the "grooming" behaviours of the great apes, and reciprocation of favours such as defending each other when a fight occurs.

Negative gestures and threat displays are also common in social animals, serving to warn off individuals in the competition for resources. Examples would include territorial marking, and aggression between males in the mating season, from the blackbird's song to the ram's head-butt.

In all such behavioural interactions, emotions are evident, or the animal form of them (perhaps without some of our more cognitively invested character). Not only in humans then, but in many other animals too, emotions serve to drive and regulate social interactions and build relationships. As these benefits are crucial to the survival of the social species, it is clear that the social emotions are vital to successful interaction, and to Homo Sapiens in particular.

According to this line of argument, we should aim to build a capacity for authentic social emotions into artificial agents if they are intended to interact with people. A basic test of such a social agent would be to see if it can be made to play the UG in a human-like manner.

4 The RBS architecture

The design we propose here is a rule-based-system (RBS architecture) [4]. In its simplest form, this kind of system has a working memory space, or database, and a set of rules that are applied to the data to add new inferences. Rules generally take the form of IF-THEN operations, so that if the antecedent IF part matches some content in the database, then the consequential changes that are given in the THEN part are applied, and the database is thus modified. Some of the facts in the database correspond to the agent's intentions to act. The way the rules apply is determined by the system's resolution policy. If there is only one rule whose IF part can match only one part of the database, there is no conflict and the rule fires, so that it adds to or changes those database contents. If there is no rule that matches the database, then no rule can fire and the execution terminates. But if there are two or more candidate rules that could fire, a choice must be made between them, and this is determined by the resolution policy.

Typical factors in a policy include the rule's *specificity*, so that rules with IF-parts that match more specific data are preferred. Another common factor is a rule's *priority*, so that more time-critical rules are preferred to fire first, for example. There may be various other factors included in the policy to help resolve conflicts between rules, and to prevent the system going into an infinite loop, and failing to terminate or produce any result. It is common for rules to be demoted after they fire, for example, by going to the end of the queue, so that they will be able to fire again until after the other rules have had their chance.

For a player who takes the role of *responder* in the Ultimatum Game, we may add a rule that angrily reacts to a low offer from the *proposer* by immediately rejecting it.

```
R-anger:
    IF      offer(X)   AND   X < low
    THEN   want( reject(X) )
```

The above rule (called R-anger) is one that makes the agent take revenge. The `offer` is the amount offered by the *proposer* (a percentage). The threshold value `low` represents the lowest offer that the agent would accept, and could be modified in some circumstances. The predicate `want` is to assert the agent's intention to reject the offer.

This is a rule that the players could easily create when the game is explained to them in the laboratory, before they sit down to play it.

They might modify the value of `low` as they digest the significance of accepting or rejecting potential offers of different sizes, but we assume here that the threshold is set by the time the game begins.

Because the game is supposed to engage rational calculation about possible payoffs, which depend on what the responder does, we add rules to predict the possible outcomes.

```
R-accept:
  IF      offer(X)
  THEN    choice([accept(X)])

R-accept-win:
  IF      choice([accept(X)])
  THEN    choice([accept(X), win(X)])

R-reject:
  IF      offer(X)
  THEN    choice([reject(X)])

R-reject-win:
  IF      choice([reject(X)])
  THEN    choice([reject(X), win(0)])

R-prefer:
  IF      choice([Xact, win(X)])
  AND     choice([Yact, win(Y)])
  AND     X > Y
  AND     NOT prefer(Xact, Yact)
  THEN    prefer(Xact, Yact)

R-want:
  IF      choice([Xact, _])
  AND     NOT prefer(_, Xact)
  THEN    want(Xact)
```

The rules are written in a notation similar to Prolog, so that variables begin with capital letters (like `Xact` which matches a possible action). The special variable `"_"` is a wild-card, which matches anything. The first rule says that if an offer has been made by the proposer (and has some value `X`), then it is possible to respond with the action `accept(X)`, to accept the offer in the game. The later rule `R-reject` declares the other possible move, which would be to reject the offer.

Then the agents needs to predict the consequences of the two possible actions, and these are explicated by the rules `R-accept-win` and `R-reject-win`, which extend the sequence of events with the resulting wins for the agent. Either the agent wins the offered amount `X`, or else nothing if the offer is rejected. More Prolog notation is used here, to represent the growing sequence of events as lists (like `[Xact, win(X)]` and so on).

The rule `R-prefer` tells the agent which of its possible actions are to be preferred. Given any pair of alternative actions (the choices `Xact` and `Yact`), the agent will prefer the one with the higher pay-out or win.

The order of execution of these rules follows the resolution policy, but we take the default basis to be in file order, as is usual. The rules thus apply in order from top to bottom, with the last rule applying when the other rules cannot match any longer or have already fired on the data they do match. Finally, then, the rule `R-want` will fire, and it matches any possible action that is a most preferable one, and then produces an impulse, to want to perform that act.

These rules collectively enable the agent to choose its best move according to what the predicted and preferred outcomes are. It is a form of planning, simplified enough to fit into the RBS scheme, but

still able to plan effectively in the chosen domain of the Ultimatum Game.

5 Example of the agent at play

In order to see how the agent performs, there is a summary of the key events or steps in its cognition shown in Table 1. Each step shows the application of a rule, and the change that it adds to the database. Later rules (in the file order) tend to use the data produced by earlier ones, and this dependency is one reason why they fire later.

The first step is shown as the agent's observation of the proposer's offer, which is an action of perception to record the fact into the database. When none of the rules can fire any longer, the system stops, and the next action phase can begin. In this case there is only one impulse to action, as the agent wants to accept the offer of 50%.

Table 1. Proposer offers 50% and the low threshold is 30%

step	rule	result
1	observe	offer(50)
2	R-accept	choice([accept(50)])
3	R-accept-win	choice([accept(50),win(50)])
4	R-reject	choice([reject(50)])
5	R-reject-win	choice([reject(50),win(0)])
6	R-prefer	prefer(accept(50),reject(50))
7	R-want	want(accept(50))
		terminates

This run of the agent shows that it can accept an offer if it is not too low (in this case anything above 30%). Human players of the Ultimatum Game invariably accept fair offers like 50%, or a little less than that.

The agent thus models that typical human response, but what about those cases where people often respond in an economically "irrational" manner, such as by rejecting a low offer? Table 2 shows what can happen in that case.

Table 2. Proposer offers 20% and the low threshold is 30%

step	rule	result
1	observe	offer(20)
2	R-anger	want(reject(20))
3	R-accept	choice([accept(20)])
... as before ...
8	R-want	want(accept(20))
		terminates

This time the rules are mostly the same as before, but a new rule (`R-anger`) has fired first. This rule is a short and simple one, which matches the offer and sees that it is too low to be acceptable. The unfair offer thus produces an immediate impulsive desire to act, namely to reject the offer.

The agent does not necessarily act impulsively on any action it wants to execute, however. In this case, the impulse is produced, but cognition continues to explore other options. When the execution terminates, if it gets that far before the agent reacts angrily, then there are two competing impulses in the database for the action system to execute.

In a more complete model, there would be further rules to decide which of the impulses should go forward to the action system for execution. They are not given here, to keep the simulation simple, but key ideas include:-

- impulses can have a strength parameter
- rules can be added to delay action

The strength of emotional reactions could then be set to be stronger than impulses that result from other planning. In both cases the strength of reaction could be determined by the other event qualities, like the unfairness of the offer, or the size of the win that would be sacrificed by a rejection. Rules would then say how strong an impulse needs to be before it is put forward to execution rather than delayed until further cognition can be performed (by allowing more rules to fire).

These new rules would therefore govern the operation of the system, and can be seen as metaheuristics. Exactly what their settings should be is open to experimentation, and indeed varying them would correspond to giving the agent different personalities, ranging from purely cognitive and rational (in the economists' sense), through emotional to generally impulsive, and extremely irascible.

6 Discussion

The artificial player of the Ultimatum Game has been designed so that it can react in ways that appear to be similar to those of human players of the game. Rules in the agent can be set to make the agent accept an offer of around 50%, which is how humans react too. These rules amount to an implementation of a decision process, or what would be called a planning system in traditional AI. They combine together to both predict the outcomes (payouts) of the responder's possible moves (accept or reject); and they notice that acceptance pays out more than rejection. As the agent prefers to win more money (as human players of the Ultimatum Game also understand the purpose to be), it then forms a desire to accept the offer. According to the rational actor view of economics, the agent should simply act on this desire, and in cases where the offer is around 50% it does so. Because human players do the same, we can conclude that they and the agent are all being rational in such cases.

Human players react differently when the offer is low, though. The agent model can be made to react similarly again, depending on how its parameters in the rules are set. The most important parameter is the threshold of what it considers to be a "low" offer, so that anything below that is seen as unfair or insulting, as a human might perceive it. In case of a low offer, the agent has another rule which also fires, and produces a desire to reject the offer. The existence of this rule is assumed to model the natural understanding that a human player would develop while hearing the game described by the experimenter. In that case the rule is like a pre-set intention to act in vengeful way if the proposer makes an "unfair" offer. However this is to simplify what would in reality be a more complicated process, in which the rule would only be half pre-formed, prior to playing the game. It is likely that human players reconsider the matter in the event that the proposal is low or nearly low. There is also the question of where such a vengeful rule could come from — in humans, but ultimately in artificial agents too. A natural way to view this is to hypothesise that such specific reaction rules come from more generic emotion rules, that become instantiated to fit into any particular context, such as in our case the Ultimatum Game. Similarly, one would aim to make artificial agents that can adapt their programmed generic rules (like "innate emotions") to any situation that is relevant to them.

However the rules are derived, their effect is to create in the agent an "emotional" reaction tendency, or impulse. In this case the agent can thus have two impulses at the same time, which are contradictory, and must somehow choose between them. This would require further

complications in the model, for which some possibilities have been described. They introduce more parameters that could be tuned to give the agent different response tendencies, or personalities. Some agents could then react with some types of emotional reaction more strongly than others, so they could be seen as relatively hostile or vengeful or passive and forgiving. They could also react at different speeds, and in this architecture it is an emergent phenomenon that depends on the complexity of the rules involved.

The main point of the architecture design given here is to show how the emotional reactions and the so-called "rational" cognition might not be such different phenomena as we suppose. In fact they are all implemented in the system as reaction rules, of the same type. What distinguishes them is only how "fast" they can react, and how "strong" are the impulses which they produce.

The model also manifests reactions that can be seen as emotional in a way consistent with human psychology. According to leading theory, including that of Frijda [1, 2], emotions are characterised, among other things, by *control precedence*, *action readiness* and *regulation*. These mean that emotion can dominate and occupy our attention if it is strong enough, and redirect our mental resources to whatever the emotion is about. It can produce a felt tendency to act in a certain way, but the tendency can be resisted depending on circumstances. In a social context, for example, it might be necessary to suppress an angry reaction, and so emotion generally is regulated. When human players are given more time to react in the Ultimatum Game, for example, their behaviour becomes less angry [3]

In the simulation of the model applied the the Ultimatum Game, such as in Table 2 and discussion, it is seen how the emotional response, depending on its strength, demands mental effort (so that the high priority rule can fire first), and could occupy attention for longer too, if there were more rules to elaborate on the consequences of the emotional response. The impulse to reject was not immediately enacted however, and in the example run the agent took time to think through its options, and come to an alternative plan of action. In this way it might manifest an emotional response but only mentally, without overtly expressing it in action. Thus the agent can be seen to partially model Frijda's theory of emotion.

To make a fuller model of an emotional agent needs further work of course, some of which has been outlined above in the paper, even if only to play the Ultimatum Game, which is in a way an extremely simple domain. It is an interesting domain for social cognition however, and would help us to examine a wide variety of phenomena that occur in social interactions. These include, or at least touch upon, issues of morality, reputation, attitudes to other agents, social status and culture.

In the future, many agents we develop will need to interact with people, for many purposes, including economic transactions. We should therefore build in a capacity for social emotion, so that they can both understand human behaviour, and predict it; and so that they can react in appropriate ways that will seem natural and intuitive, and not perplex people. Moreover, they will presumably be more effective agents too, if they follow rules of behaviour and interaction protocols that have been evolved over millions of years.

Finally, note that the emotions are necessary for social life in two opposing ways. While they help drive and guide our social interactions on the one hand, and can do so strongly, they also need to be partially controlled on the other hand. If emotions are not guided or muted, depending on the context, then they can lead to extreme and unregulated behaviour. Since agents will have to be responsible citizens in the future, they will have to regulate their emotions just as we do ourselves. The architecture presented here can form the basis

to allow this, by integrating emotion and cognition, to get the best of both. Then agents will not only be emotional, but *intelligently* so.

REFERENCES

- [1] Nico H. Frijda, *The emotions*, CUP Press, 1986.
- [2] Nico H. Frijda, *The laws of emotion*, Erlbaum, 2007.
- [3] Veronika Grimm and Friederike Mengel, 'Let me sleep on it: Delay reduces rejection rates in ultimatum games', *Economics Letters*, **111**(2), 113–115, (2011).
- [4] Crina Grosan and Ajith Abraham, 'Rule-based expert systems', in *Intelligent Systems*, volume 17 of *Intelligent Systems Reference Library*, 149–185, Springer Berlin Heidelberg, (2011).
- [5] Martin A. Nowak, Karen M. Page, and Karl Sigmund, 'Fairness versus reason in the ultimatum game', *Science*, **289**(5485), 1773–1775, (2000).
- [6] Rosalind W Picard, *Affective computing*, MIT press, 1997.
- [7] Paul J. Zak, Angela A. Stanton, and Sheila Ahmadi, 'Oxytocin increases generosity in humans', *PLoS ONE*, **2**(11), e1128, (11 2007).