

An Intelligent Virtual Agent to Increase Involvement in Financial Services

Tibor Bosse¹, Ghazanfar F. Siddiqui^{1,2}, and Jan Treur¹

¹ Vrije Universiteit Amsterdam, Department of Artificial Intelligence,
De Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands

² Quaid-i-Azam University Islamabad, Department of Computer Science, 45320, Pakistan
{tbosse, ghazanfa, treur}@few.vu.nl ghazanfar@qau.edu.pk
<http://www.few.vu.nl/~{tbosse, ghazanfa, treur}>

Abstract. In order to enhance user involvement in financial services, this paper proposes to combine the idea of adaptive personalisation with intelligent virtual agents. To this end, a computational model for human decision making in financial context is incorporated within an intelligent virtual agent. To test whether the agent enhances user involvement, a web application has been developed, in which users have to make a number of investment decisions. This application has been evaluated in an experiment for a number of participants interacting with the system and afterwards providing their judgement by means of a questionnaire. The preliminary results indicate that the virtual agent can show appropriate emotional expressions related to states like happiness, greed and fear, and has high potential to enhance user involvement.

Keywords: user involvement, finance, greed and risk, adaptive personalisation.

1 Introduction

In recent years, there has been a huge increase in the amount of services that are being offered via the Internet. These services include, among others, financial services such as Internet banking [17]. Despite the success of such services, an existing challenge in this area concerns the question how to make people more *involved* in such financial applications. According to [1], customer involvement in financial services can be defined as ‘an unobservable state of motivation, arousal or interest’ (taken from [14]). In order to increase this state of involvement in users of financial applications, some authors claim that personalisation is an important criterion (e.g., [2, 6]): by having the system learn certain characteristics of the customer, this person will feel more understood and will be more likely to accept the service that is offered. However, there is also research that suggests that personalisation alone is not sufficient for financial services to attract users for longer periods (e.g., [8]).

To deal with this last issue, the current paper proposes to enhance user involvement in financial applications by combining adaptive personalisation with the use of an intelligent virtual agent. As pointed out by various authors (e.g., [9, 12]), human-like

virtual agents have the ability to increase a user's presence in virtual environments. This finding was the inspiration to develop a personalised intelligent agent which supports persons that have to make financial (investment) decisions. As known from behavioural economics, humans do not behave completely rationally when they have to decide between alternatives that involve risk (as, for example, in financial situations). Since then, from time to time it has been argued that theories of economic decision making need to incorporate psychological factors such as greed and fear [5, 11, 13, 16]. Thus, the main goal of this paper is to develop a virtual agent that has insight in and adapts to the individual psychological characteristics and states over time of persons that are working with financial applications. The virtual agent should exploit this on the one hand by providing appropriate support, in following these (dynamical) states and characteristics in an adaptive personalised manner. On the other hand, by showing the appropriate emotions at the right moment the virtual agent encourages involvement and reflection by the person through mirroring his or her states and decision making processes; for example, the agent may show the person how greedy he or she behaves.

In order to develop such a supporting virtual agent, as a basis a solid computational model of human decision making in financial context is needed. To this end, the model presented in [3] is taken. This model takes some of the main principles underlying the Modern Portfolio Theory (MPT) [7, 15] as a point of departure, and extends these with mechanisms to incorporate psychological factors (inspired, among others, by [11, 13]). In the current paper, this model is incorporated within an intelligent virtual agent. To test whether the agent enhances user involvement, a simple web application has been developed, in which users make a number of investment decisions. This application has been evaluated by a number of participants in an experiment in which they interacted with the agent and afterwards provided their judgement by means of a questionnaire.

The remainder of this paper is structured as follows. In Section 2, the basic model for financial decision making (taken from [3]) is summarised, and an intelligent virtual agent application is introduced that incorporates the model. Section 3 introduces the experiment that was performed to evaluate the virtual agent, and Section 4 presents the results. Section 5 concludes the paper with a discussion.

2 The Virtual Agent Application

The model for financial decision making described in [3] is based on the assumption that a person's greed is determined by her (long-term) personality profile (e.g., some persons are more risk seeking than others), combined with observations about recent events (e.g., if many investments have provided high returns recently, persons are more likely to increase their greed, and as a consequence take more risk). These assumptions can also be found in existing literature such as [11, 13]. By incorporating this model within a virtual agent, the agent is able to analyse a human's decision making by observing her decisions and the received returns, while tuning the risk profile to the person. Within this analysis not only this personal risk profile is available, which is assumed static for the person, but also the more dynamic greed

level that actually determines the decisions. By having this, at each point in time the agent can predict what a reasonable decision would be for the human, given her personal background and history. In particular at all stages of the process it can estimate and show the type and level of emotions expected. These emotions can be shown to the human at runtime.

To design and implement the virtual agent, Haptek's Peopleputty software [19] has been used. Through this software the face of the virtual agent was created. More specifically, twelve different faces were designed using the built in sliders for happy, sad, anger, mellowness, suspicion, and curiosity (which are related to facial expressions), and ego, aggression and energy (which are related to head movement). Each of these twelve faces represented a particular emotional state; one for each possible combination of the three required levels of the emotions *happiness* (slightly_happy, happy and very_happy), *sadness* (slightly_sad, sad and very_sad), *fear* (slightly_scared, scared and very_scared) and *greed* (slightly_greedy, greedy and very_greedy). In addition, a face for the state neutral was used. A web-based application was implemented, within which the virtual agent was embedded as a Haptek player. For this the scripting language JavaScript [18] was used, in combination with scripting commands provided by the Haptek software [19], to control the Haptek player within a web browser.

Within the application, a human can make a number of consecutive investment decisions, while the virtual agent mimics the estimated emotional states related to happiness, sadness, greed and fear of the human (see the screenshot in Figure 1, where the agents shows a slightly_greedy face). In this application, in total 10 (represented by letters from A to J) products are given. The characteristics of these products are represented by the two variables X and Y, which are shown on the right hand part of the screen. Here, as in [3], X represents the expected risk of the product, and Y represents the expected return of the product. Note that in the model both X and Y have a value in the domain $[0, 1]$, but in the application the values of Y have been scaled to the domain $[0, 1000]$, to have them represent US dollars (see Section 3 and 6 of [3] for the exact formulae used within the model).

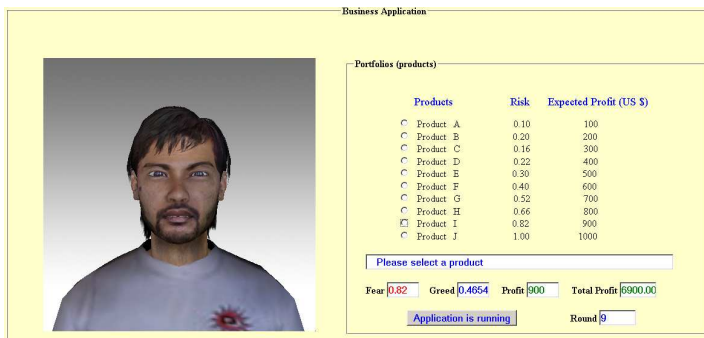


Fig. 1. Screenshot of the Application




Fig. 2. Fluctuation of W

During a number of rounds, the human is asked to select a product from the given products (from A to J). After she selected a product, some time will pass, until a message is shown on the screen that the “result of your investment will soon be

announced". Next, it again takes some seconds until the real result is shown on the screen. To determine what this result will be, the formulae introduced in [3] are used. Note that these formulae make use of an additional parameter W (in the domain $[0, 1]$) which represents the economic situation of the world (e.g., a high value indicates a strong economy, thus a higher probability to receive return on investment). The value of this parameter fluctuates over time, and is not known to the user.

In every round, the virtual agent shows emotional facial expressions at appropriate moments. The following fixed scheme determines when to show which type of emotional expression: 1) human is asked to select a product, 2) agent shows face related to greed, 3) human selects a product, 4) agent shows face related to greed, 5) message is shown stating that result will soon be announced, 6) agent shows face related to fear, 7) result of investment is announced, 8) agent shows face related to happiness/sadness, 9) go back to 1.

The criteria that determine the exact faces that are displayed are as shown in Table 1 (where risk equals the X value of the selected product, profit equals the result of the investment (i.e., either 0 or the Y value of the selected product), and greed equals the value of G as estimated by the model).

Table 1. Criteria for the displayed Face Expressions

	<i>Criterion</i>	<i>Displayed</i>
Fear	$\text{risk} \leq 0.5$	neutral face
	$\text{risk} > 0.5 \ \& \ \text{risk} \leq 0.7$	scared face
	$\text{risk} > 0.7 \ \& \ \text{risk} \leq 1$	very scared face
Happiness/Sadness	$\text{profit} = 0$	very sad face
	$\text{profit} > 0 \ \& \ \text{profit} \leq 300$	slightly happy face
	$\text{profit} > 300 \ \& \ \text{profit} \leq 600$	happy face
	$\text{profit} > 600 \ \& \ \text{profit} \leq 1000$	very happy face
Greed	$\text{greed} \leq 0.1$	neutral face
	$\text{greed} > 0.1 \ \& \ \text{greed} \leq 0.3$	slightly greedy face
	$\text{greed} > 0.3 \ \& \ \text{greed} \leq 1$	very greedy face

When the agent shifts from one facial expression to another, it would be undesirable if the emotions of the agent would shift too quickly. Therefore, these shifts are performed in a more fluent manner. For instance, if the agent shifts from very happy to very greedy, the following faces are shown consecutively:

very happy à happy à slightly happy à neutral à slightly greedy à greedy à very greedy

Such a scenario is used when the agent shifts from any emotional state to another emotional state.

While the application is running, some information about the user is displayed in the bottom right part of the screen (see Figure 1). This information concerns the user's estimated amount of fear and greed (in the domain $[0, 1]$), her current amount of profit received, and her total (cumulative) amount of profit.

3 Experiments

A number of experiments were performed to test to what extent users of the application feel involved with the agent. In total, 15 participants were recruited to

perform the experiment. The age of the participants ranged between 24 and 34, with a mean age of 29 and a standard deviation of 2.78. Among the participants, 11 were male and 4 were female. Two variants of the experiment were designed, one in which the virtual agent was showing the appropriate emotions and one in which it did not show any emotions. All of the participants were asked to perform both variants (where we used counterbalancing to determine the order in which they were performed).

Before they started the experiment, each participant was first asked to read the following instructions:

*Imagine that you are an investor in a stock market. During a number of subsequent rounds, you have to select a product from a given set of products. Each round, the same 10 products are available. The characteristics of these products are represented by two variables (called X and Y), which are shown on the screen. Here, X is a value in the domain [0, 1] which represents the risk of the product (i.e., a higher value for X means that it is more likely that you will not receive the corresponding return), and Y is a value in the domain [0, 1000] which represents the expected return of the product in US dollars (i.e., a higher value for Y means that you will earn more profit). The value of X is related to the probability p of not receiving the corresponding return Y according to the formula $p = X * (1 - W)$. Here, W is a number in the domain [0, 1] which represents the economic situation of the world (i.e., a higher value for W means that it is more likely that you will receive the corresponding return). However, the value of W fluctuates during the simulation, and is not shown to you. After you have selected a product, some time will pass, until a message is shown on the screen that the result of your investment will soon be announced. Next, it still takes some seconds until the real result is shown on the screen. As mentioned above, the probability of receiving the profit also depends on the economic status of the world. After the result of your investment has been announced, a new round starts, in which you are asked to make a new investment. In total, the experiment lasts 20 rounds.*

Next, a small training was given to each participant, and after that the participants performed the actual experiment. When the experiment was finished, the person was asked to fill in a questionnaire. In this questionnaire (cf. [4]), the participants were asked to evaluate, using a 7 point Likert scale [10] (with 1=strongly disagree, 7=strongly agree and 4=neutral), various properties of the agent related to involvement. In particular, they were asked whether they thought the virtual agent was friendly, trustworthy, showing emotions adequately, realistic, showing happiness, showing sadness, showing greed, showing fear, and behaving human-like.

In each experiment, the value of the economic state W fluctuated between 0 and 1, as shown in Figure 2. However, the participants were not aware of this.

4 Results

The answers provided by the participants to the questions about their involvement with the virtual agent were analysed by means of paired sample t-tests. The results are shown in Table 2. As shown in the table, for all properties except 'greed', the participants scored the virtual agent with emotions above moderate. Similarly, for all properties except 'greed', the participants appreciated this variant more than the virtual agent without emotions. The participants were also asked to give suggestions or comments about the application. Some participants said, for example, that the fear emotion should be more intense, while others said that the greed emotion should be improved, as they did not see this very well. Participants also indicated that they were more involved with the virtual agent with emotions. In addition, some participants were of the opinion that the agent should speak as well.

The fact that greed did not score very well in this first test may depend on the type of face expression chosen for greed. To explore how the perception of greed could be

enhanced, another small experiment was conducted. For this purpose 9 different alternative faces for greed were created using the Peopleputty software [19]. Six new participants were asked to rate each face on a 7 point scale, for its appropriateness to express greed. After all the participants gave their responses, for each face the average score over all participant responses was taken, and the face with the highest average value was selected for a next experiment (again with 6 new participants). This experiment was identical to the experiment of which the results are given in Table 2, only in this case the new face was used to display the greed. Table 3 shows (part of) the results of this experiment. As can be seen, this time the greed was evaluated much better, and resembles the evaluations of the other emotions.

Table 2. Results of the Questionnaire

Q #	Question	With Emotions		Without Emotions		Paired Sample Test	
		Mean	SD	Mean	SD	t	Sig(2tailed)
1	Friendly	4.47	1.506	3.53	1.992	2.168	.048
2	Trustworthy	4.13	1.598	2.47	1.246	4.183	.001
3	Adequate emotions	4.93	1.223	2.40	1.844	4.219	.001
4	Realistic	4.93	1.534	3.47	1.922	2.442	.028
5	Happiness	5.84	0.834	2.13	1.457	9.153	.000
6	Sadness	5.67	1.113	2.20	1.474	7.124	.000
7	Greed	2.67	1.877	2.00	1.195	1.323	.207
8	Fear	4.27	1.751	2.00	1.254	5.264	.000
9	Happiness at right time	5.60	0.986	2.20	1.656	7.462	.000
10	Sadness at right time	5.87	0.990	2.00	1.363	9.648	.000
11	Fear at right time	4.13	1.598	1.93	1.387	4.036	.001
12	Human-like	4.67	1.291	2.93	1.667	2.578	.022

Table 3. Results of the Questionnaire for the Additional Experiment related to Greed

Q #	Question	With Emotions		Without Emotions		Paired Sample Test	
		Mean	SD	Mean	SD	t	Sig(2-tailed)
1	Greed	5.167	0.983	2.33	1.366	3.782	.013

5 Discussion

In this paper an application was presented combining adaptive personalisation with intelligent virtual agents, in order to enhance user involvement in financial services. To this end a computational model for human decision making in financial context (taken from [3]) was incorporated within an intelligent virtual agent. This computational model enables the virtual agent to have a form of understanding of the person's (dynamical) states and decision making processes in an adaptive manner. Moreover, a second way in which the agent was made adaptive was by equipping it with a model to tune the risk profile parameter to the person.

A web application has been developed, in which users make a number of investment decisions. This application has been used to test whether the virtual agent enhances user involvement. This has been evaluated in an experiment for a number of participants interacting with the system and afterwards providing their judgement by

means of a questionnaire. The preliminary results indicate that the virtual agent can be given facial expressions showing emotional states like happiness, greed and fear that are evaluated as appropriate (regarding the type of facial expression as well as the moments on which the expressions are shown). In particular, getting an appropriate expression for greed was nontrivial.

For future work the virtual agent may be tested in a real environment to analyse whether it makes humans perform better in financial decision making, for example in the form of a smart phone application. One of the factors that may need some more attention is the level of awareness of the person of the state of the world.

References

1. Aldlaigan, A.H. and Buttle, F.A. (2001). Consumer involvement in financial services: An empirical test of two measures. *Int. Journal of Bank Marketing*, 19(6), pp. 232-246.
2. Blom, J. and Monk, A. (2001). One-to-one e-commerce: who's the one? In: *Proc. of the Conference on Human Factors in Computing Systems (CHI'01)*. ACM Press, pp. 341-342.
3. Bosse, T., Siddiqui, G.F., and Treur, J., Supporting Financial Decision Making by an Intelligent Agent Estimating Greed and Risk. In: *Proc. of the Fourth Int. Workshop on Human Aspects in Ambient Intelligence, HAI'10*. IEEE CS Press, 2010, to appear.
4. Bowling A. (1997). Questionnaire design. In: *Research Methods in Health*. Buckingham: Open University Press, 1997.
5. Brandstätter, H. (1993). Should *Economic Psychology Care about Personality Structure?* *Journal of Economic Psychology*, vol.14, pp. 473-494.
6. Chau, P.Y.K. and Lai, V.S.K. (2003). An empirical investigation of the determinants of user acceptance of internet banking. *Journal of Organizational Computing & Electronic Commerce*, 13(2), pp. 123-145.
7. Haugen, R.A. (1997). *Modern Investment Theory*. Prentice-Hall, Inc.
8. Ho, S.Y. (2006). The Attraction of Internet Personalization to Web Users. *Electronic Markets*, vol. 161, pp. 41-50.
9. Leite, I., Pereira, A., Martinho, C., and Paiva, A. (2008). Are Emotional Robots More Fun to Play With? *Proceedings of the IEEE RO-MAN 2008 Conference*, pp. 77-82.
10. Likert, R. (1932). A Technique for the Measurement of Attitudes. *Archives of Psychology*, vol. 140, pp. 1-55.
11. Lo, A., Repin, D.V., and Steenbarger, B.N. (2005). Fear and Greed in Financial Markets: A Clinical Study of Day-Traders. *American Economic Review* 95, pp. 352-359.
12. Nowak, K.L. and Biocca, F. (2003). The effect of the agency and anthropomorphism on users' sense of telepresence, copresence, and social presence in virtual environments. *Presence: Teleoperators and Virtual Environments* 12(5), pp. 481-494.
13. Rabin, M. (2002). A Perspective on Psychology and Economics. *European Economic Review*, vol. 46, pp. 657-685.
14. Rothschild, M.L. (1984). Perspective on Involvement: Current Problems and Future Directions. In: Kinnear, T.C. (ed.), *Advances in Consumer Research*, ACR, vol. 11.
15. Sabal, J. (2002). *Financial Decisions in Emerging Markets*. Oxford Univ. Press, Inc., NY.
16. Simon, H.A. (1987). Behavioural Economics. In: *The New Palgrave: A Dictionary of Economics*. London, MacMillan.
17. Tan, M. and Teo, T.S.H. (2000). Factors Influencing the Adoption of Internet Banking. *Journal of the AIS* 1(5), pp. 1-42.
18. http://developer.mozilla.org/en/docs/About_JavaScript
19. <http://www.haptek.com>