

Personalized Alert Agent for Optimal User Performance

Avraham Shvartzon¹, Amos Azaria², Sarit Kraus¹, Claudia V. Goldman³, Joachim Meyer⁴ and Omer Tsimhoni³

¹ Dept. of Computer Science, Bar-Ilan University, Ramat Gan 52900, Israel

² Dept. of Machine Learning, Carnegie Mellon University, Pittsburgh, PA 15213

³ General Motors Advanced Technical Center, Herzliya 46725, Israel

⁴ Dept. of Industrial Engineering, Tel Aviv University, Ramat Aviv, Tel Aviv 69978, Israel

{shvarta,sarit}@cs.biu.ac.il, azariaa@cs.cmu.edu, {claudia.goldman, omer.tsimhoni}@gm.com, jmeyer@tau.ac.il

Abstract

Preventive maintenance is essential for the smooth operation of any equipment. Still, people occasionally do not maintain their equipment adequately. Maintenance alert systems attempt to remind people to perform maintenance. However, most of these systems do not provide alerts at the optimal timing, and nor do they take into account the time required for maintenance or compute the optimal timing for a specific user. We model the problem of maintenance performance, assuming maintenance is time consuming. We solve the optimal policy for the user, i.e., the optimal timing for a user to perform maintenance. This optimal strategy depends on the value of user's time, and thus it may vary from user to user and may change over time. Based on the solved optimal strategy we present a personalized maintenance agent, which, depending on the value of user's time, provides alerts to the user when she should perform maintenance. In an experiment using a spaceship computer game, we show that receiving alerts from the personalized alert agent significantly improves user performance.

Introduction

In our daily life we rely on various types of mechanical equipment and electrical devices such as our computers, cars or bicycles, smartphones, washing machines, dryers, heating and air conditioning systems and so on. These devices are very important to us, but they tend to malfunction occasionally, which can greatly disrupt our lives. Most devices come with maintenance recommendations, which, if followed, aim to decrease the probability and frequency of such malfunctions. Unfortunately, performing maintenance actions is both costly and time consuming. Therefore many people fail to perform these maintenance actions in a timely manner, and consequently they suffer the consequences, when necessary devices suddenly malfunction (occasionally when they are most needed).

This non-optimal behavior with regards to performing maintenance may be attributed to the following reasons: 1. Procrastination and forgetfulness: People are known to forget and procrastinate, especially in regard to tasks for which the consequences are not immediately evident (Ariely and Wertenbroch 2002; Weber 2006). Therefore, when it is time

to perform a maintenance task (according to the user manual), people often either forget or procrastinate. 2. Non-optimal recommendations: Occasionally it is not optimal for the user to perform all maintenance tasks as required by the manual. Performing maintenance for all devices according to their manual may be tedious and not cost effective. Users may actually benefit, in terms of time and cost, from performing maintenance less often than requested by the manual. 3. Non-personalized recommendations: The instructions which appear in manuals are not personalized and thus do not account for people who have different values of time or different costs associated with performing maintenance and repairs. Therefore performing maintenance tasks according to the manual may not be the optimal behavior for a specific user.

In an attempt to increase people's awareness to perform maintenance, and solve the problem of procrastination and forgetfulness, many devices include alert systems or reminders which recommend performing maintenance activities. For example, an air conditioning system may turn on a warning light when it is recommended to change its filter. A car may have similar warning lights, which are set to go on when an oil change or other periodic treatments are recommended. There are also software programs which alert users when a computer system has not been backed-up for some time, recommending them to back-up their system. Reminders and alerts have been shown to have a positive impact on people's tendency to perform a task (Barreau and Nardi 1995).

We model the problem of providing maintenance alerts, taking into account the fact that performing maintenance and repairing malfunctions are both wealth and time consuming. Inspired by the Bellman Equation, we solve the optimal policy for the user, i.e., the optimal timing for a user to perform maintenance. It is important to note that in order to accurately capture the cost of time consuming actions, the model must take into account the value of time for each user (or user performance). Therefore, our solution depends on this value of time, and thus the optimal performance time varies from user to user and may change over time.

We present the *spaceship game*, which allows us to evaluate people's tendency to perform maintenance and repairs, with and without recommendations obtained from a maintenance alert system. In this game, a player controls a space-



Figure 1: A screen-shot of the spaceship game in progress.

ship, which shoots asteroids (see Figure 1 for a screen-shot). The player is required to perform maintenance actions on his or her spaceship. Occasionally, and depending on the frequency at which the player performs maintenance, the spaceship may suffer malfunctions which are repaired at a significant cost.

We present a personalized alert agent which provides maintenance recommendations. This agent tries to overcome the three causes mentioned above for non-optimal human behavior with respect to taking maintenance actions. Given the user's performance so far the agent predicts the user's expected future performance, and, using our general solution, the agent identifies the optimal timing for a player to perform maintenance and recommends the user to perform maintenance at this optimal timing. We show in an experiment, that when subjects are presented with personalized recommendations, they significantly perform better than when they are presented with non-personalized recommendations (which are optimal only for average performance) and when they receive no alerts at all. This result may encourage maintenance alert system designers to not only urge users to take maintenance actions, but to do so according to timing which is optimal for each specific user and his or her needs, which may change over time. Such personalized maintenance alert systems may significantly improve users' overall performance.

Related Work

An empirical study that examined users' tendency to perform preventive maintenance actions with or without indications from an alert system that indicated the need for intervention was conducted in an abstract laboratory setting (Bitan and Meyer 2007). The study showed that users do not optimally perform preventive actions, and that they can be aided by an alert system, especially when the system is reliable. The study of alert provision is part of the behavioral shaping field, in which an agent aims to change the behavior of a human, either for the human's self benefit, or for the benefit of a different party (represented by the agent). Much work in AI is dedicated to this field (Hajaj, Ha-

zon, and Sarne 2014; Azaria, Richardson, and Kraus 2014; Azaria et al. 2011; Hajaj et al. 2013; Azaria, Aumann, and Kraus 2014; Azaria et al. 2014; Rosenfeld et al. 2015b; Azaria et al. 2015).

Many applications exist to help equipment owners maintain equipment. In (Dekker 1996), the author describes several optimization models for maintenance, including stochastic and deterministic models, which are distinguished by simplicity (with a single component) and complexity. Such an optimization model is deployed by (Li, Mourelatos, and Singh 2012). Their optimal maintenance model is based on the expected lifecycle of a given equipment, taking into account the costs of production, inspection, etc. So far, to the best of our knowledge, no studies exist that discuss and evaluate different reminder strategies for maintenance of cars or other equipment. Sherif and Smith (1981) provide a review on optimal maintenance models. However, neither of the papers in the review account for the time required to perform maintenance, and thus there is no personalized model.

Many studies on reminder strategies exist in the field of medicine. For example, a survey by Vervloet et al. (2012) examines the effectiveness of interventions using electronic reminders in improving patients' adherence to chronic medication. This review provides evidence for the effectiveness of electronic reminders in improving adherence by patients taking chronic medication. However, these works did not consider even the trivial personalization, of not sending a reminder after the medication was taken. They emphasize that further research is needed to investigate the influence of the frequency with which reminders are sent on adherence. The effectiveness of this non-automated type of electronic reminder for adherence is currently being investigated (Vasbinder et al. 2013). Similarly, a survey by Tao et al. (2014) supports these findings, affirming that the use of electronic reminders seems to be an effective way to improve medication adherence of patients with chronic conditions. They recommend that future research should aim to identify optimal strategies for the design and implementation of electronic reminders, with which the effectiveness of the reminders is likely to be augmented.

Several works have considered interruption management, questioning when would be the best time to interrupt the user (Adamczyk and Bailey 2004; Bailey and Iqbal 2008; Sarne and Grosz 2007; Rosenfeld et al. 2015a). Shrot et al. (2014) clustered the users into groups and then used collaborative filtering in order to determine which interaction method would work best with each type of user. This approach yields different policies for different users. However, in contrast to this paper, they did not try to change the user's behavior, but convey information at the best timing.

Several additional approaches use a form of personalization for maintenance. A model based on equipment personalization is discussed in (Last 2011). This model predicts equipment failures, based on sensor measurements and warranty claims. Huang et al. (2009) suggest a maintenance model, which indeed uses a form of personalization, albeit the model does not attempt to optimize the maintenance, but rather utilizes collected data to inform the user about recommended tasks that needed to be performed, based on the

personalized usage history.

The Spaceship Game

In this work we use a variant of the spaceship game (Shvartzon et al. 2015) in order to model the problem of maintenance alert provision. In this game the player controls a spaceship (see Figure 1). Throughout the flight the spaceship should shoot down meteors which fly in space. Every time a meteor is shot down, the player gains money (points). The player must also avoid getting hit by the meteors, and he loses money if he is hit by them. As the players achieve points, they obtain additional cannons.

In order to reduce the probability of malfunctions, the spaceship needs to occasionally carry out maintenance actions. Each of these actions is both time consuming and incurs a monetary cost. While performing a maintenance action (which lasts several seconds), the spaceship ‘freezes’ and the player cannot shoot down any asteroids and thus is unable to gain any points (but at the same time cannot lose points either). If a malfunction occurs it is repaired. This repair is associated with both a score and time cost (i.e. the spaceship freezes during the repair time). In this paper we study how an agent providing maintenance alerts may improve a player’s performance in the spaceship game.

Formal Model

In this section we build a formal model of a maintenance game. This model has the following characteristics (which are motivated by real world maintenance settings): 1. A user must determine at which time (t) to perform maintenance. 2. Such maintenance actions impact the probability of faults (as will be explained below). 3. If a malfunction occurs, the user must fix it in order to continue. 4. Maintenance and repair actions are associated with a maintaining / repairing cost and maintaining / repairing time. While the equipment is being either maintained or repaired, it may not be used. We denote by c_m the monetary cost of maintenance and by c_r the monetary cost of repair (usually $c_m < c_r$). We use w_m to denote the waiting time required (cost of time in seconds) once the player performs maintenance and w_r to denote the waiting time of repairing a malfunction. 5. Future discounting. We assume that any future action is discounted, i.e., multiplied by an exponentially decreasing discounting factor (γ). Future discounting has justification from both psychology and economics (Broome 1994) (although the exact function that should be used is often in dispute (Rubinstein 2003)). In our domain, a discount factor may be justified by a probability for a sudden break-down (which is unrelated to the maintenance-repair actions), such as a total loss in a car accident, the loss of a phone, sudden-death of the player etc.

We use Π to denote a maintenance policy, which determines at any given time, whether the user should perform maintenance or not. We assume the following sampling method procedure for determining when and whether a malfunction will occur: 1. At the beginning of the game a time at which a malfunction may potentially occur is sampled using some probability density function (*pdf*) (or a cumulative density function, *CDF*). 2. If the player performed

maintenance before that time, this malfunction is removed. Once the maintenance ends, a new potential malfunction is sampled using the same *pdf*, but from the current time. The player is associated with both the monetary and time costs of maintaining. 3. If the player has reached the time at which a malfunction may potentially occur, a malfunction occurs and is fixed (and the player is associated with both the monetary and time costs of repairing). A new future potential malfunction is sampled similarly to the above.

The following two propositions relate to the properties of the optimal policy. We show that the optimal policy has the form of performing maintenance every X seconds since the last maintenance or repair action has ended (or X seconds since the beginning of the game, if no malfunction has occurred).

From the sampling method we observe that once the player performs maintenance (and the maintenance cost is applied), or a malfunction occurs (and is fixed, along with its cost), the player faces the exact situation as in $t = 0$.

Proposition 1. *The optimal policy may be determined by some value t . When following the optimal policy, the player performs maintenance once t time units have passed since the last time maintenance was performed or a malfunction has occurred.*

Proof. Assume an optimal policy Π^* . Since this policy may be general, it may take into account all previous actions and occurrences. Running Π^* on a sampled game, and assuming no malfunction occurred from the beginning until the first maintenance, denote by X the first time Π^* implied a maintenance action. Clearly (since Π^* must be deterministic), in any game, unless a malfunction occurs before t , t is fixed. Once a maintenance action is performed or a malfunction occurs (resulting in a repair action) at time t_1 , according to the above observation the game becomes identical to the starting position and thus the optimal policy is identical and, unless a malfunction occurs, the next maintenance should be performed at $t_1 + t$. \square

A corollary of Proposition 1 is that the domain of policies may be determined by a single parameter t . Thus, we use $\Pi(t)$ to denote a policy in which maintenance is performed when t seconds have elapsed since the last maintenance or repair actions have ended.

Due to the complexity of our solution, we begin by presenting a simpler model in which there is no time cost associated with performing maintenance or repairing a malfunction, i.e., there are only monetary costs associated with maintenance and repairing malfunctions.

Actions Have Monetary Cost Only

In this section we assume that there is no time cost associated with performing maintenance or repairing a malfunction. This simplifies the model and makes it easier to identify the optimal strategy for maintenance. In this section we will compute the policy which brings to minimum the expected cost of maintenance / repair actions, over time.

According to Proposition 1, finding the optimal policy is equivalent to determining t such that the overall expected

cost under a policy $\Pi(t)$ is minimized. We use $C(t)$ to denote the expected cost of maintenance and repairs under a policy in which maintenance is performed when t seconds have elapsed since the last maintenance or repair actions have ended. Thus the optimal policy, Π^* is $\Pi(t)$ such that: $t = \arg \min_t C(t)$.

In order to calculate $C(t)$ we will first calculate $C(t)$ given no malfunction occurs the first time, i.e., the first malfunction is sampled *after* t . Denote the first malfunction sampling time by t_{m1} . We are therefore interested in $C(t | t_{m1} > t)$. The next lemma will show that $C(t | t_{m1} > t)$ has a form similar to that of a Bellman Equation, i.e., the expected cost appears on the right side of the formula multiplied by the discount factor.

Lemma 1. $C(t | t_{m1} > t)$ has the following property:

$$C(t | t_{m1} > t) = \gamma^t (c_m + C(t))$$

Proof. Since no malfunction occurs, at time t , according to the policy, the user performs maintenance at a cost of c_m . This cost is discounted by γ^t (since it occurs at time t). Once a maintenance is performed, the player faces the same situation as in $t = 0$ and thus the expected cost is $C(t)$ (multiplied by the discount factor). \square

Denote $p_r(\cdot)$ as the probability for an event. We now calculate $C(t | t_{m1} < t) \cdot p_r(t_{m1} < t)$, the expected cost given that the malfunction actually occurs in the first time, i.e. the malfunction is sampled *before* t , multiplied by the probability that the malfunction actually occurs before time t .

Lemma 2. $C(t | t_{m1} < t) \cdot p_r(t_{m1} < t)$ has the following property:

$$C(t | t_{m1} < t) \cdot p_r(t_{m1} < t) = \int_{x=0}^t pdf(x) \gamma^x (c_r + C(t)) dx$$

Proof. Once a malfunction occurs, the player encountered cost c_r and returns to the starting position $C(t)$. Assuming this malfunction occurred at time x , this cost and the starting position are multiplied by γ^x . Given an infinitesimally small time x , the probability that the malfunction occurs at this time is $pdf(x)dx$. Therefore, the portion of the expected cost when the malfunction actually occurs *before* t is $\int_{x=0}^t pdf(x) \gamma^x (c_r + C(t)) dx$. \square

Finally we define an equation which allows us to calculate $C(t)$:

Theorem 1. The expected cost follows: $C(t) =$

$$\int_{x=0}^t pdf(x) \gamma^x (c_r + C(t)) dx + (1 - CDF(t)) \gamma^t (c_m + C(t))$$

Proof. According to the law of total probability: $C(t) = C(t | t_{m1} < t) \cdot p_r(t_{m1} < t) + C(t | t_{m1} > t) \cdot p_r(t_{m1} > t)$ By Lemmas 1 and 2: $C(t) = \int_{x=0}^t pdf(x) \gamma^x (c_r + C(t)) dx + \gamma^t (c_m + C(t)) \cdot p_r(t_{m1} > t)$. The probability for a single malfunction to be sampled after t , $p_r(t_{m1} > t)$, is $(1 - CDF(t))$, since by definition, $CDF(x)$ is the probability that $X \leq x$. \square

Time consuming actions

In this subsection we take into account the fact that performing maintenance or repair may be time consuming. As we will show, this solution implies that the player's optimal time to perform maintenance actually depends on the

player's performance (in the spaceship game this would be the average points gained per second).

In order to account for the value of time and thus model the loss which the player will encounter by waiting, we need to consider the performance of the player, i.e., the points or value the player obtains per unit of time (excluding any maintenance or repair costs). This value will be denoted p . Since the player is assumed to have some value of time, we now denote $U_p(t)$ as the expected *utility* for a player with performance p using a policy that performs maintenance every t seconds, and our problem becomes a *maximization* problem. The optimal policy Π^* is now $\Pi(t)$, such that: $t = \arg \max_t U_p(t)$.

Calculating the expected utility, is fairly similar to the calculation of the expected cost in the previous section, though must take into account the player's performance and the time cost of both maintenance actions and malfunctions (or repair). We first calculate the utility of $U_p(t)$ given no malfunction occurs the first time, i.e., the first malfunction is sampled *after* t ($t_{m1} > t$).

Lemma 3. $U_p(t | t_{m1} > t)$ has the following property:

$$U_p(t | t_{m1} > t) = \gamma^t (-c_m + \gamma^{w_m} U_p(t)) + \int_{y=0}^t p \gamma^y dy$$

Proof. Since no malfunction occurs, at time t , according to the policy, the user performs maintenance at a cost of c_m (discounted by γ^t). *After waiting the time cost of maintenance*, w_m , and thus applying an additional discounting factor of γ^{w_m} , the player faces the same situation as in $t = 0$. Since the player played until time t performing no maintenance, and no malfunction occurs, the utility from performance alone up to time t is given by: $\int_0^t p \gamma^y dy$. \square

We now calculate $U_p(t | t_{m1} < t) \cdot p_r(t_{m1} < t)$, the expected utility of the player given that the malfunction is sampled *before* t and multiplied by the probability that the malfunction actually occurs before time t .

Lemma 4. $U_p(t | t_{m1} < t) \cdot p_r(t_{m1} < t)$ equals:

$$\int_{x=0}^t pdf(x) \left(\gamma^x (-c_r + \gamma^{w_r} U_p(t)) + \int_{y=0}^x p \gamma^y dy \right) dx$$

Proof. Until the malfunction occurs (at time x), the player gains a utility of $\int_{y=0}^x p \gamma^y dy$. Once the malfunction occurs, the player encountered cost c_r and returns to the starting position $U_p(t)$ *after waiting the time cost of repairing*, w_r , thus an additional factor of γ^{w_r} . All other factors are identical to Lemma 2. \square

Finally we define an equation which calculates $U_p(t)$:

Theorem 2. The expected utility of a policy $\Pi(t)$ performing maintenance every t time units since the previous malfunction or maintenance, and assuming performance p follows the following equation:

$$U_p(t) = \int_{x=0}^t pdf(x) \left(\gamma^x (-c_r + \gamma^{w_r} U_p(t)) + \int_{y=0}^x p \gamma^y dy \right) dx + (1 - CDF(t)) \left(\gamma^t (-c_m + \gamma^{w_m} U_p(t)) + \int_{y=0}^t p \gamma^y dy \right) \quad (1)$$

Proof. According to the law of total probability: $U_p(t) = U_p(t | t_{m1} < t) \cdot p_r(t_{m1} < t) + U_p(t | t_{m1} > t) \cdot p_r(t_{m1} > t)$. The rest follows from Lemmas 3 and 4: \square

Equation 1 can be written as:

$$U_p(t) = \frac{(-c_r + \frac{p}{\ln(\gamma)}) \int_{x=0}^t pdf(x) \gamma^x dx - \frac{p}{\ln(\gamma)} \int_{x=0}^t pdf(x) dx}{1 - (1 - CDF(t)) \gamma^{(t+w_m)} - \int_{x=0}^t pdf(x) \gamma^{(x+w_r)} dx} + \frac{(1 - CDF(t)) \left(\gamma^t \left(-c_m + \frac{p}{\ln(\gamma)} \right) - \frac{p}{\ln(\gamma)} \right)}{1 - (1 - CDF(t)) \gamma^{(t+w_m)} - \int_{x=0}^t pdf(x) \gamma^{(x+w_r)} dx} \quad (2)$$

In order to find the optimal policy, we must differentiate Equation 2 with respect to t and solve $U'_p(t) = 0$ (and test the two extreme cases of $t = 0$ and $t = \infty$).

Maintenance Timing in the Spaceship Game

We use a simple form of the beta distribution with $\alpha = 2$ and $\beta = 1$, spread over T seconds. That is, we calibrate the beta distribution over a time-line of T seconds, i.e., $pdf(t) = \frac{2}{T^2} t$ and $CDF(t) = (\frac{t}{T})^2$. This function has the following crucial property related to repairing and maintenance: the more time that has passed since the latest maintenance, the more likely it is that a malfunction will occur at each moment.

We use our results from Theorem 2 to solve the optimal maintenance timing in the spaceship game and plug the given beta distribution into Equation 2 (the equation and its derivative are omitted due to lack of space). Given the total length of the game, t_{tot} , we use: $\gamma = \frac{1}{1-t_{tot}}$ to determine the discount factor. Assuming the discount factor is viewed as a way to model the probability that the game will suddenly end, the above formula ensures that the expected length of such a game when using the discount factor will be t_{tot} .

Agents for Maintenance Alerts

We considered two maintenance alert agents. The first, is a non-personalized agent (NPA) that only uses the *average* human performance when providing maintenance alerts. This agent, calculates the optimal time to perform maintenance, using the formulas developed above and assuming the performance of an average player.

The second is our Fully Personalized Agent (FPA). This agent predicts the expected performance of the current player, and, using this information, the FPA calculates the optimal time to perform maintenance (using the formulas developed above). We considered four different methods to compute the expected performance, based on the previous data of the current player. The first method is simply to use the performance in the previous second to predict the performance in the current second. The second method is to use the average performance of the current player from the beginning of the (current) game. The third method, the moving window method, requires a parameter x . This method uses the average performance of the past x seconds to predict the current performance. The fourth method, known as exponential smoothing ((Gans, Knox, and Croson 2007)) predicts current performance by multiplying the previous expected performance by some discount factor, δ , and the performance of the previous second by $1 - \delta$. Exponential smoothing was shown to work well when predicting human

Method	MSE
Previous second	1393
Full average	766.5
Moving window	767.1
Exponential smoothing	649.5

Table 1: Prediction methods and their mean squared error (MSE) on training data.

behavior ((Acuna and Schrater 2008)). Both the exponential smoothing and moving window methods require some initial value (a parameter that must be determined).

In order to determine the average performance for NPA and the method to be used to predict the user performance along with its parameters for FPA, we collected training data from 20 subjects. Table 1 presents the mean squared error (MSE) of the prediction methods employing a tenfold-cross-validation on the training data (lower values indicate a better fit). As shown, the exponential smoothing method exhibited a higher fit-to-data than all other performance prediction methods. Therefore the exponential smoothing method was selected to be implemented in the FPA.

Evaluation

Experimental Setup

In order to evaluate our agent, we recruited a total of 56 subjects via Amazon’s Mechanical Turk (AMT) to play the spaceship game. The set of subjects consisted of 57% males and 43% females. Subjects’ ages ranged from 18 to 51, with a mean of 32. All subjects were residents of the USA. The subjects were paid 20 cents to participate in the experiment, and, depending on their performance, could achieve an additional payment of up to \$2.00. Before playing, the subjects filled out a short demographic questionnaire. They were then presented a tutorial, played a training game to help them understand the game rules and played a two minute practice game. Each subject played 3 actual games which lasted 4 minutes each. In these three games the subjects either received alerts from FPA, NPA or received no alerts and were simply told that they should perform maintenance every 20 seconds (which is optimal for a player with average performance). The order of these different games may have a significant impact on the performance of the user, since the user may gain experience over time. Note, that there are 6 possible permutations on the order of the three conditions. Therefore, every time a new subject joined, she was assigned to the permutation with the least subjects so far.

In order to increase compliance, when presenting a maintenance alert, the screen flashed for a second and a tone was played. The alerts were presented one second in advance (i.e., one second before the agent wanted the user to actually perform maintenance). Table 2 presents the settings we used for the spaceship game. After playing all games, the subjects filled out a questionnaire.

Requiring full attention for all these games, is more than most AMT workers are used to. Therefore, we expected that some subjects will not play seriously for so long, and de-

Action	Score (\$)	Time cost (spaceship frozen)
Maintenance	-\$5	8 secs
Repair	-\$500	3 secs
Hit by meteor	-\$10	N/A
Hit a meteor	+\$30	N/A

Table 2: Settings used in the spaceship game




Figure 2: Average performance of subjects in FPA, NPA and no alert games.

cided on an exclusion criterion of removing all subjects who had a negative score game in any of the three conditions. This exclusion criterion was reinforced by the fact that several subjects stated (in the comments part of the ending questionnaire) that over time they have lost interest in the game. Unfortunately, 20 subjects have met this exclusion criterion.

Results

The fully personalized agent (FPA) accounted for an improvement of between 17% and 18% over the non-personalized agent (NPA) and the no alert condition (5637 vs. 4756 and 4833). We ran an ANOVA test with repeated measurements on the subject score and set the agent type as our dependent variable and the game order as a controlled variable. The impact that the agent type had on the subject’s score was statistically significant ($F(2, 30) = 3.467$, $p < 0.05$). In pairwise comparisons FPA significantly outperformed both NPA and the no alert condition (with Bonferroni adjustment for multiple comparisons, single tailed, $p < 0.05$). Any differences between NPA and no alert were minor and not significant. Figure 2 illustrates these results.

The average standard deviation of the scores within each of the groups was 3573, which is in fact 70% of the average score. This is clearly a very large standard deviation, emphasizing the need for a personalized agent which is most required when people differ from one another. Interestingly, the subjects performed more maintenance actions when presented with alerts from the NPA (7.8 per game), than with the FPA (6.2 per game) and when they received no alerts at all (6.2 per game). This indicates that the failure of the NPA was not due to the subjects not following its alerts but due to the fact that these alerts were not presented at the optimal timing according to the individuals’ performance.

When discarding the exclusion criterion and including all subjects the average performance with FPA was 3770, still outperforming NPA (with 3613) and No-alert (with only

2950). Though, clearly, subjects who lost interest in a game, added a tremendous amount of noise.

The subjects seemed to enjoy the game, giving it an average of 7.7 on a 1 to 10 scale. We found low correlation between the average score of a player and the answer to this question (0.29), and negligible correlation between the player’s age and the answer to this question (0.06).

Discussion

Any maintenance agent is limited in the ability to increase the overall performance of the user, due to the fact that the major contributor to the user’s score is actually the ability of the player to shoot down asteroids and avoid being hit by them, and we assume that the maintenance alert agent has no impact on such abilities. Therefore, we believe that our results, which show that FPA increases the average performance of the players by 17% is a very significant result.

The consequences of not performing maintenance may vary among different environments. The more immediate and obvious is the link between not performing maintenance and malfunctioning, the more likely the user is to follow the advice, and the more important it may be to provide personally fitted alerts. If the result of maintenance is more vague, the agent may need to provide additional information on top of merely alerting the user, so that if and when a malfunction occurs, the user will be able to link it with her own actions, and perhaps follow the advice next time.

In this paper we use past performance to predict future performance. While in real life, the value of time may be more challenging to obtain, some bodies (insurance, lawsuits etc.) try evaluating it. These bodies usually consider the earning wage, i.e. the amount the person earns a month or year divided by the number of working hours and sometimes ask the subject, either direct questions (e.g. How much would you be willing to spend to save an hour of your time?), or indirect questions (e.g. How much time would you be willing to spend to save \$50 on your groceries?).

The cost (both in time and in money) of maintenance and repairing is often available. The probability density function may be trickier to obtain. In order to learn a reasonable approximation for it, one needs to collect enough data. However, in the field of car maintenance, for instance, most new cars have logs stating when each treatment was performed and when each malfunction appeared, etc. This information, if gathered from enough cars, may be a valid source for computing the probability density function. Additional domains exist in which the performance may be easier to predict, e.g. a company using machines which produce a value to the company as long as they work, and have some cost maintaining them.

Conclusions

In this paper we study to what extent a maintenance recommendation system can be beneficial to a user. We present the spaceship game, which is a game wherein a user controls a spaceship that requires maintenance to avoid malfunctions. This game allows the quantification and scoring of the user’s performance and enables us to measure to what extent recommendation alerts affect her behavior and performance.

In this paper we show the effectiveness of a personalized maintenance recommendation agent. When interacting with a user, the agent predicts the expected performance of that user. Based on this prediction and our analytically derived solution, the agent alerts the user to perform maintenance at the optimal maintenance timing. To the best of our knowledge, this is the first paper that attempts to solve the optimal time for performing maintenance, when taking into account the value of time for the individual users and thus resulting in personalized recommendations.

References

- Acuna, D., and Schrater, P. 2008. Bayesian modeling of human sequential decision-making on the multi-armed bandit problem. In *Cognitive Science Society*, volume 100, 200–300. Washington, DC: Cognitive Science Society.
- Adamczyk, P. D., and Bailey, B. P. 2004. If not now, when?: the effects of interruption at different moments within task execution. In *CHI*, 271–278. ACM.
- Ariely, D., and Wertenbroch, K. 2002. Procrastination, deadlines, and performance: Self-control by precommitment. *Psychological Science* 13(3):219–224.
- Azaria, A.; Aumann, Y.; and Kraus, S. 2014. Automated agents for reward determination for human work in crowdsourcing applications. *JAAMAS* 28(6):934–955.
- Azaria, A.; Rabinovich, Z.; Kraus, S.; and Goldman, C. V. 2011. Strategic information disclosure to people with multiple alternatives. In *AAAI*.
- Azaria, A.; Kraus, S.; Goldman, C. V.; and Tsimhoni, O. 2014. Advice provision for energy saving in automobile climate control systems. In *IAAI*.
- Azaria, A.; Gal, Y.; Kraus, S.; and Goldman, C. V. 2015. Strategic advice provision in repeated human-agent interactions. *Autonomous Agents and Multi-Agent Systems* 1–26.
- Azaria, A.; Richardson, A.; and Kraus, S. 2014. An agent for the prospect presentation problem. In *AAMAS*, 989–996.
- Bailey, B. P., and Iqbal, S. T. 2008. Understanding changes in mental workload during execution of goal-directed tasks and its application for interruption management. *TOCHI* 14(4):21.
- Barreau, D., and Nardi, B. A. 1995. Finding and reminding: file organization from the desktop. *ACM SigChi Bulletin* 27(3):39–43.
- Bitan, Y., and Meyer, J. 2007. Self-initiated and respondent actions in a simulated control task. *Ergonomics* 50(5):763–788.
- Broome, J. 1994. Discounting the future. *Philosophy & Public Affairs* 23(2):128–156.
- Dekker, R. 1996. Applications of maintenance optimization models: a review and analysis. *Reliability Engineering & System Safety* 51(3):229–240.
- Gans, N.; Knox, G.; and Croson, R. 2007. Simple models of discrete choice and their performance in bandit experiments. *Manufacturing & Service Operations Management* 9(4):383–408.
- Hajaj, C.; Hazon, N.; Sarne, D.; and Elmalech, A. 2013. Search more, disclose less. In *AAAI*, 401–408.
- Hajaj, C.; Hazon, N.; and Sarne, D. 2014. Ordering effects and belief adjustment in the use of comparison shopping agents. In *AAAI*.
- Huang, Y.; Gardoni, M.; Coulibaly, A.; et al. 2009. A decision support system designed for personalized maintenance recommendation based on proactive maintenance. In *ICED*.
- Last, M. 2011. Vehicle failure prediction using warranty and telematics data. In *Learn*, volume 29, 245–260.
- Li, J.; Mourelatos, Z.; and Singh, A. 2012. Optimal preventive maintenance schedule based on lifecycle cost and time-dependent reliability. Technical report, SAE Technical Paper.
- Rosenfeld, A.; Agmon, N.; Maksimov, Oleg, A. A.; and Kraus, S. 2015a. Intelligent agent supporting human-multi-robot team collaboration. In *IJCAI*.
- Rosenfeld, A.; Azaria, A.; Kraus, S.; Goldman, C. V.; and Tsimhoni, O. 2015b. Adaptive advice in automobile climate control systems. In *AAMAS*, 543–551.
- Rubinstein, A. 2003. “Economics and psychology”? the case of hyperbolic discounting*. *International Economic Review* 44(4):1207–1216.
- Sarne, D., and Grosz, B. J. 2007. Sharing experiences to learn user characteristics in dynamic environments with sparse data. In *AAMAS*, 43–50.
- Sherif, Y., and Smith, M. 1981. Optimal maintenance models for systems subject to failure—a review. *Naval Research Logistics Quarterly* 28(1):47–74.
- Shrot, T.; Rosenfeld, A.; Golbeck, J.; and Kraus, S. 2014. Crisp: an interruption management algorithm based on collaborative filtering. In *CHI*, 3035–3044. ACM.
- Shvartzon, A.; Azaria, A.; Kraus, S.; Goldman, C. V.; Tsimhoni, O.; and Meyer, J. 2015. A game for studying maintenance alerts’ effectiveness. In *AAMAS*, 1655–1656.
- Tao, D.; Xie, L.; Wang, T.; and Wang, T. 2014. A meta-analysis of the use of electronic reminders for patient adherence to medication in chronic disease care. *Journal of telemedicine and telecare* 1357633X14541041.
- Vasbinder, E. C.; Janssens, H. M.; Rutten-van Mülken, M. P.; van Dijk, L.; de Winter, B. C.; de Groot, R. C.; Vulto, A. G.; van den Bemt, P. M.; et al. 2013. e-monitoring of asthma therapy to improve compliance in children using a real-time medication monitoring system (rtmm): the e-matic study protocol. *BMC medical informatics and decision making* 13(1):38.
- Vervloet, M.; Linn, A. J.; van Weert, J. C.; De Bakker, D. H.; Bouvy, M. L.; and Van Dijk, L. 2012. The effectiveness of interventions using electronic reminders to improve adherence to chronic medication: a systematic review of the literature. *Journal of the American Medical Informatics Association* amiajnl-2011.
- Weber, E. U. 2006. Experience-based and description-based perceptions of long-term risk: Why global warming does not scare us (yet). *Climatic Change* 77(1-2):103–120.