Hindawi Publishing Corporation
International Journal of Computer Games Technology
Volume 2012, Article ID 625476, 10 pages
doi:10.1155/2012/625476

Research Article

Dynamic Difficulty Balancing for Cautious Players and Risk Takers

Guy Hawkins,[1] Keith Nesbitt,[2] and Scott Brown[1]

[1] *School of Psychology, Newcastle Cognition Laboratory, University of Newcastle, Callaghan, NSW 2308, Australia*
[2] *School of Design Communication and IT, University of Newcastle, Callaghan, NSW 2308, Australia*

Correspondence should be addressed to Keith Nesbitt, keith.nesbitt@newcastle.edu.au

Received 13 March 2012; Accepted 27 April 2012

Academic Editor: Daniel Thalmann

Dynamic balancing of game difficulty can help cater for different levels of ability in players. However, performance in some game tasks depends on not only the player's ability but also their desire to take risk. Taking or avoiding risk can offer players its own reward in a game situation. Furthermore, a game designer may want to adjust the mechanics differently for a risky, high ability player, as opposed to a risky, low ability player. In this work, we describe a novel modelling technique known as particle filtering which can be used to model various levels of player ability while also considering the player's risk profile. We demonstrate this technique by developing a game challenge where players are required to make a decision between a number of possible alternatives where only a single alternative is correct. Risky players respond faster but with more likelihood of failure. Cautious players wait longer for more evidence, increasing their likelihood of success, but at the expense of game time. By gathering empirical data for the player's response time and accuracy, we develop particle filter models. These models can then be used in real-time to categorise players into different ability and risk-taking levels.

1. Introduction

In designing a popular game, it would be beneficial to have a model of the ideal player. The designer could use this player profile to design just the right amount of difficulty and emotional impact into their game. No player would become bored with easy challenges or overburdened by difficult ones. All players could be fairly compensated for taking risks by a well-calculated reward structure. Everyone who played the game would receive the same optimal experience and level of entertainment. Unfortunately for the designer, it is unlikely that such an ideal player model exists.

To aid the design process, it is not unusual to categorise players into different groups. A typical division of such player types is the casual versus hard-core player, or the experienced versus inexperienced player. Such categories provide designers with general levels of player abilities allowing them to design corresponding levels of difficulty into the game mechanics.

The motivation for the designer is to balance game difficulty with player ability in such a way that the game is sufficiently challenging that it maintains interest and entertains across the broadest possible range of player abilities. This assumes that, to reward a player, their overall satisfaction is solely dependent on game difficulty and their ability to succeed in each challenge. However, there are other types of reward that can also be important for players. For example, how well the game allows the players to exercise their desired level of risk taking.

In this paper, we develop a novel model of the player using a technique known as particle filtering. Such a model can incorporate various levels of player ability while also considering the player's risk profile. Once developed, such a particle filter model is well suited to making dynamic adjustments in game difficulty. To develop the model, however, first requires the gathering of empirical data and fitting this data into an appropriate particle filter for the game scenario. This paper focuses on demonstrating the development of

such a particle filter model and illustrating how it can be incorporated into the gameplay.

2. Utility and Risk

When considering player reward in terms of overall satisfaction, it is convenient to use terminology associated with economics, where overall satisfaction is described as "utility". The concept of utility can then be related to a person's preferred level of risk taking. For example, a typical division of management styles would be risk seeking, risk neutral, and risk averse [1]. Figure 1 relates these styles to utility, showing that a risk seeker is only satisfied when the payback is high, where conversely a risk averse individual is equally satisfied at low returns.

While risk taking profiles are more typically related to management styles or real-world activities such as gambling or stock market trading, they are also relevant to game designers. Risk taking profiles are particularly relevant for gameplay that involves making decisions based on incomplete information. In a game situation, where players must choose between alternatives, we could expect some players to take more risks than others. We might also expect that an individual's level of risk is associated with their level of enjoyment in the game. Risk seekers do not want to play it safe and cautious players do not want to risk it all.

How then can a designer deliver the best entertainment across the spectrum of risky and cautious players while also catering for different levels of ability in players? In this paper, we address this question by considering both the player's ability level and their risk-taking approach and describe how to dynamically recognise and adapt the gameplay based on these player attributes.

The dynamic modelling technique we use here is known as particle filtering [2]. Particle filters are simulation-based models that use sequential Monte Carlo methods. It is a novel approach in gaming terms and although it has been proposed as a way to incorporate intelligence into nonplayer characters [3] we are not aware of it being previously used for modelling players. Particle filters are promising models for describing cognition, and in particular decision making, as they involve updating beliefs about the state of the world as evidence accumulates over time.

Particle filters can readily model various levels of ability in human performance by just varying the number of particles in the filter. For example, a large number of particles can model statistically optimal behaviour, while a smaller number generates predictions similar to flawed, human-like behaviour.

Particle filters have been successfully used to account for behaviour in a number of psychological domains including language comprehension [4], categorisation [5], change detection [6], and determining reward rate payoffs [7].

Another useful feature of particle filters is that they can be updated in real time with a minimal computational overhead. This makes the algorithm ideal for dynamic balancing since the model can be included in the game mechanics without impacting on the speed of the game loop.

In the following section, we will briefly discuss dynamic balancing and then go on to describe a specific decision-making game challenge. We use this challenge to develop a casual game for gathering empirical data about player response time and accuracy in the challenge. The empirical data allows us, via the particle filter model, to generate estimates of both the player's ability and their risk profile. Finally, we illustrate how this model can be used to dynamically recognise and adapt the gameplay for a continuum of both novice to expert players, and risky to cautious players.

3. Balancing Difficulty

How should the designer adjust difficulty in a game challenge? Certainly, the difficulty of any game should normally increase as the game progresses, trying to match the player's increasing level of skill or in game powers. But where does the level of difficulty begin for each unique player? If the difficulty level starts too low, the player may become bored; if it starts too high, the player may become overwhelmed.

For the designer, the simplest approach is to allow players to choose their own difficulty level. Over time, many games in different genres have used this technique. The Atari 2600 console even provided a hardware switch to choose between two difficulty levels in games like Adventure [8] and Asteroids [9]. More recent games such as Quake [10], Halo [11], and Devil May Cry [12] provided more typical software-based difficulty levels that the player could select. Many other games provide adjustable difficulty settings allowing for easy, normal, hard, and extreme play. These difficultly levels are often given exotic names such as "piece of cake," "let's rock," "come get some," and "damn I'm good" as used in Duke Nukem [13].

The second approach to solving the difficulty level problem is to dynamically measure player performance during the game and adjust the difficulty based on how well the player is performing. This approach has also been utilised in a number of games. Some good examples are the third-person shooter, Max Payne [14], Far Cry [15], Left 4 Dead [16], and the Mario Kart [17] series of games. Indeed, the technique was commonly used in racing games and became known as "rubber-banding" as the mechanics of the game were adjusted to ensure the player was always held close to other cars, as if all the racers were held together by rubber bands [18].

A key feature of good dynamic balancing is transparency to the player. There is the danger that, when the mechanics are adjusted, the difficulty no longer matches the narrative. Left 4 Dead [16] attempts to overcome this problem by using their "AI Director" which dynamically adjusts the game's dramatics and pacing along with the difficulty. For example, spawning enemies are appropriately placed and numbered based on the player's current situation.

Dynamic balancing first requires identifying player ability followed by an appropriate adjustment of difficulty. Player ability can be measured by any number of parameters such as successful shots, life points, time to complete a task, or indeed any values used to calculate the game score. This

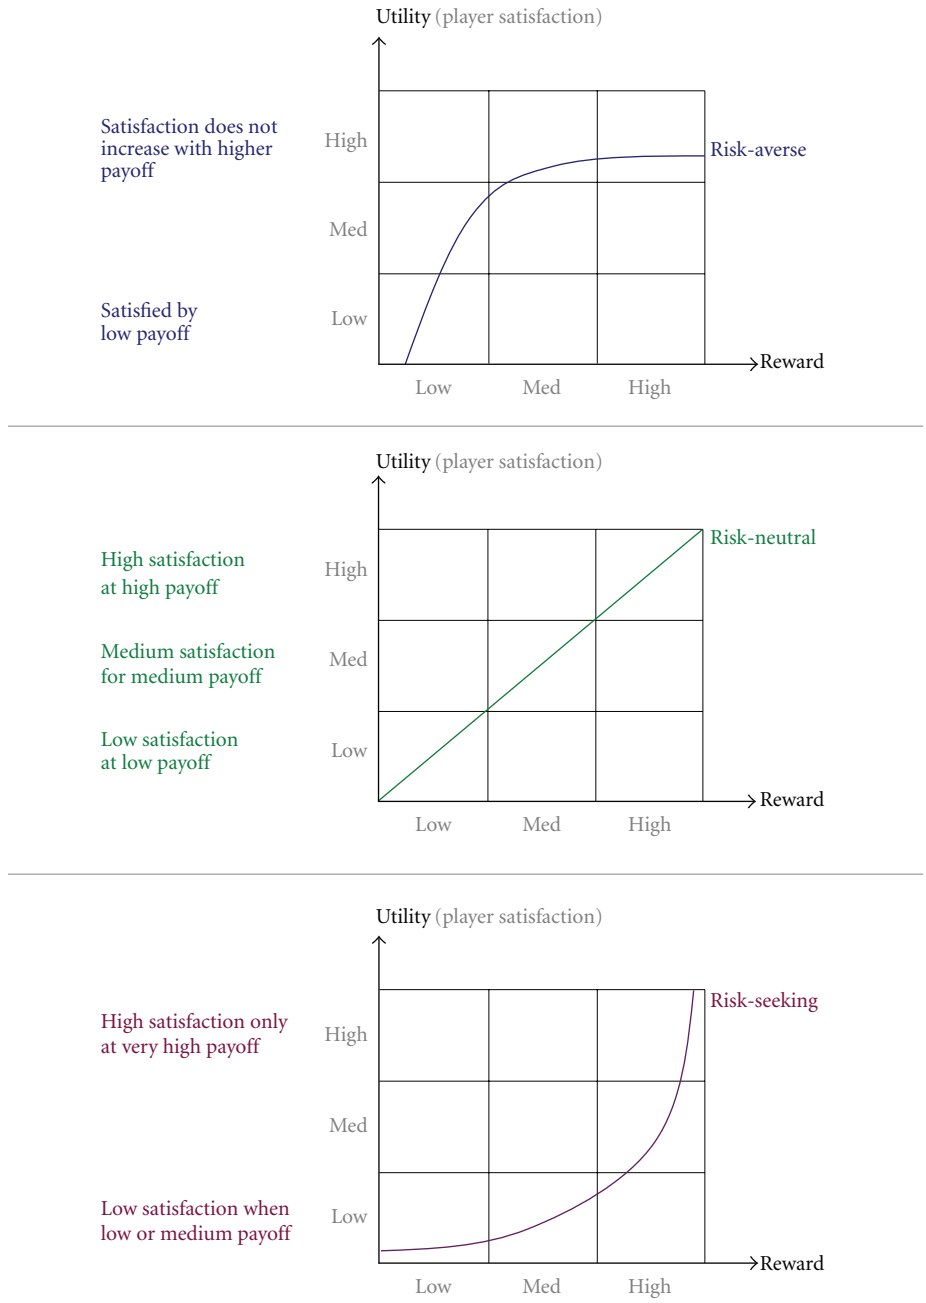

Figure 1: Three different risk profiles relating utility and reward. Note how risk seekers are only satisfied with large rewards accompanied by higher risks, while risk averse players prefer lower risk and reward decisions.

calculation has been called a "challenge function" [19] as it relates to how challenging the player is finding the game in its current state.

Depending on the game genre, typical adjustments to game mechanics either enhance player ability or adjust the ability of competing NPCs. Some examples include adjusting the speed, health, power, number, or spawn rate of enemies or the frequency or strength of player power ups. For example, in Max Payne the game dynamically adjusts the strength of enemies and can also provide different levels of aiming assistance for players. In the Mario Kart games, lower ranked players are more likely to receive items that improve their speed in future races.

Different approaches have been suggested for dynamic balancing. Hunicke and Chapman [20] and Hunicke [21] developed a first person shooter, called Hamlet, that automatically estimates the player's current requirement for core inventory items such as health, ammunition, shielding, and weapons.

Adjusting the intelligence of the NPCs is another approach that has been described. This can be achieved using sets of rules with a probability or weight attached to each

rule and then dynamically adjusting the weight [22]. So, for a novice player, the NPC might be more likely to behave based on rules that are less effective and thus give the player a better chance. Other techniques based on reinforcement learning [23], and evolutionary algorithms [19], have also been used to adapt the intelligence of NPCs for the player's skill level.

The balancing approach by Yannakakis and Hallam [24] is more closely aligned with our own technique, as it uses experimental data from gameplay to develop a player model. Their models use different types of artificial neural networks that are trained through evolutionary techniques based on game features and player entertainment. Their aim is to predict the level of player satisfaction and adjust the game accordingly.

Our study also uses experimental data to develop player models but in the context of a decision-making challenge that we describe in the next section. To create our player models, we employ a novel technique called particle filtering which allows us to model both player risk taking as well as ability.

The idea of dynamic adjustment, however one implements it, rests on first measuring a player's ability, and then knowing how large an adjustment to make. Our model provides help on both these problems. It aids in measuring ability because it will not be "fooled" about a player's ability just because that player adopts an unusual level of risk or caution. Also, it helps in knowing how large the adjustments should be because it provides a predictive model of player behaviour.

4. A Simple Decision Challenge

In our study, we first developed a simple decision-making challenge where the difficulty became easier over time. The challenge requires a player to choose between possible alternatives where only one is correct. We describe the single correct alternative as the "target" and the other, incorrect possibilities as the "distracters." The difficulty of the challenge can be adjusted by two means: increasing the total number of distracters; or by increasing the similarity of the distracters to the target.

Players perform best when they choose the correct alternative as quickly as possible. However, the nature of the challenge is that the target becomes more evident as time passes. Because both response time and accuracy are important measures of success in the challenge, the player can be risky by responding earlier rather than later, but in doing this they run a greater chance of choosing the wrong alternative.

In our challenge, the possible alternatives consisted of 20 empty squares on a screen (Figure 2). As time passed, some of these squares gradually filled with dark blue dots. This was likened to raindrops filling a bucket and the player's task was to predict which bucket was filling the fastest. The filling process was based on a probability distribution. For example, time passed in discrete steps, and at each discrete time step, the distracter squares had a 40 percent chance of gathering a new fill event (a blue dot), while the target square had a 50 percent chance.

The player must choose the target square as quickly and accurately as possible. As time passes, the target square is more evident as we expect the actual distribution of raindrops to approach the probability settings. The closeness of the probability distributions, between distracters and target, affects the difficulty of the challenge and this is one of the parameters under the designer's control. The other parameter that can be controlled is the number of alternatives the player must choose from. In our challenge, we allowed for up to 20 alternatives.

As the display evolves over time, we expect the decision to become easier as the filling approaches the probability distribution. As Figure 2 shows, even with a 10 percent difference between the distracters and the target square, the task is not trivial. We expect risky players to make a decision quickly, based on little accumulated evidence. Because they respond quickly with insufficient information, we also expect them to make a number of incorrect choices. Note that players with high ability in this task may also respond quickly but would be more accurate. This demonstrates why response time alone is not sufficient to distinguish the risk profile of players and why we must also consider it in relation to their accuracy in the task.

5. Collecting Player Data

To allow us to develop our player model for the task, we first prototyped this simple challenge using Flash and Actionscript in a nongame context. It was deployed online and subsequently played by 31 first year psychology students from the University of Newcastle. Each player completed a total of 140 decision challenges from which we recorded response time and accuracy.

The number of active squares (K) displayed on any challenge was randomly chosen from $K \in \{2, 4, 6, 8, 10\}$, subject to the condition that each K appeared equally often for every player. The target square was randomly allocated to one of the active squares.

During each challenge, the display evolved in discrete steps of 15 events per second. We monitored this frame rate during the game and only used data from players whose computers met this frame rate. On each time step, each active square either accumulated a new dot or not. The distracters always filled with a probability of 40 percent while the target filled with a probability of 50 percent. This means that, on average, the target square accumulated 7.5 dots per second while each distracter square accumulated approximately 6.0 dots every second.

At the start of a challenge all squares began with a completely empty white background. Each time a new dot was accumulated a 2×2 pixel area within the square changed to a dark blue colour. The position of the new dot was chosen randomly from the remaining unfilled area of the square.

Players were instructed to identify the target as quickly as possible, but if they responded too early, they may incorrectly select a square that had, by chance, collected the most dots

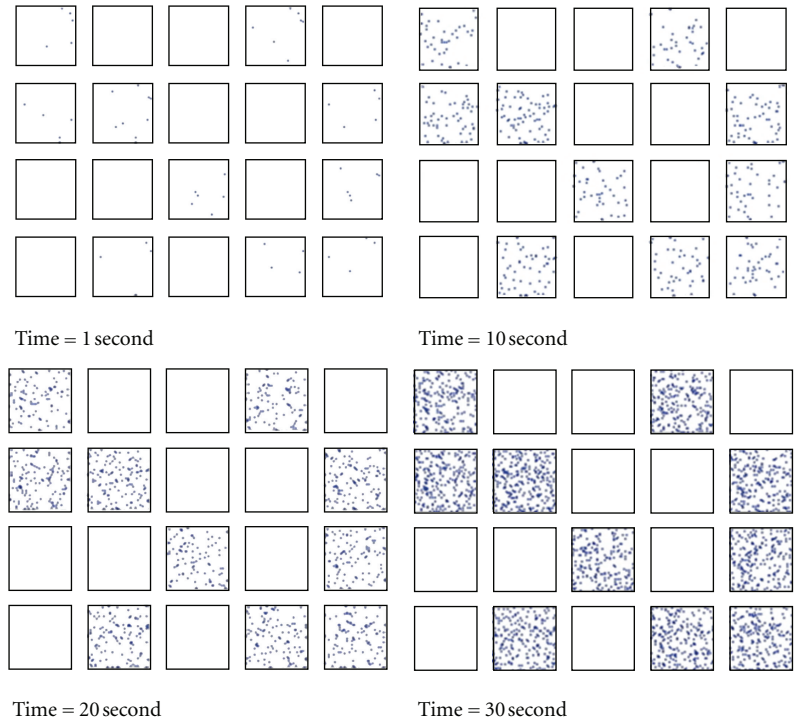

Figure 2: Four screen shots showing a decision challenge with ten alternatives at different time periods. As time progresses each of the active squares fill with dots based on independent probability distributions. In this case, the correct alternative is the square in the second column of the second row.

so far in the challenge. Participants were free to watch the display until they felt confident enough to make their decision. They recorded their choice by simply clicking on their chosen square.

After the participant chose a target square, a fast animation illustrated many more fill events very quickly. This rapid filling of the squares provided feedback to the player on whether they had selected the target square. If the player had correctly selected the target, a green outline was displayed around the chosen square. If the player's response was incorrect the chosen square's border was outlined in orange, and the true target square was outlined in green.

The mean performance of all players in terms of response time and accuracy is shown in Figure 3. Note how the average performance decreases with the number of alternatives but is higher than expected if players respond purely by guessing. This data is consistent with what we know about such challenges from Hick's Law [25, 26]. Hick's Law can be expressed in a number of ways, the most simple stating that mean response time (RT) and the logarithm of the number of choice alternatives (K) are linearly related: $RT = a + b\log(K)$. Hick's law generally provides good descriptions of data across many different types of decision-making tasks [27, 28].

6. The Game Scenario

We next transferred the simple decision challenge into a game scenario and made it available online. While the mechanics of the challenge remained the same, however, we provided a more elaborate backstory and integrated the challenge into a simple gameworld (Figure 4).

Players were introduced to a game titled "EMFants: Last Light." A mission brief informed participants they were commander of Dark-Stealth-6, a spaceship with time-hop propulsion, a "shadow-scope" to detect alien EMFants, and "blue-ray" armament. The goal of the game was identified as locating and destroying EMFants. Participants were provided with a backstory describing the electromagnetic-feeding (EMF) habits of the EMFant species. The EMFants escaped from a twin universe and have been detected in numerous galaxies. The player's goal was to destroy the EMFants before they rapidly spread to all known galaxies.

After the few introduction screens, players were informed of the layout of the game. The game mirrored a typical psychological experiment in structure, consisting of many trials within multiple blocks, although the trials in the game were described as "missions" that must be manually controlled (i.e., click a "next mission button"). At the start of each new block, players manually engaged Dark-Stealth-6's time-hop capabilities to navigate from one galaxy to another. This initiated a short animation representing the time hop.

Players were required to use their shadow-scope to detect the EMFant colonies. The shadow scope consisted of a number of squares that were being filled with dots (as in the original experiment). The EMFant colony growing at the fastest rate indicated the home of the EMFant queen (the target square). By clicking on the target, players

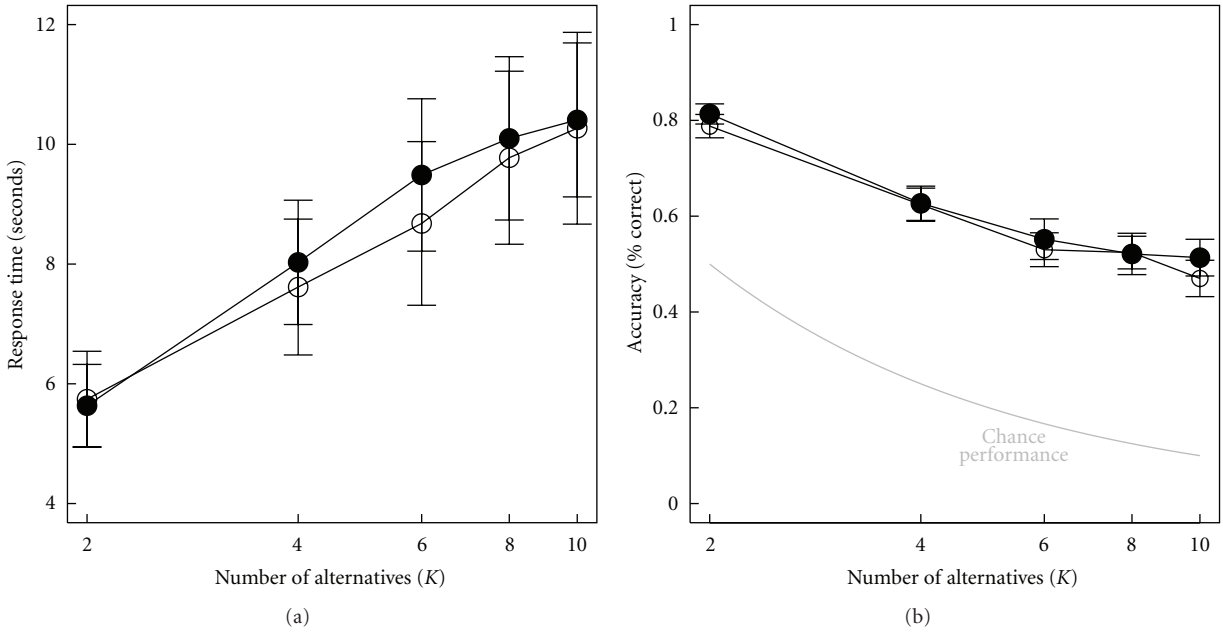

Figure 3: Mean response time (a) and decision accuracy (b) for the simple challenge and contextualised challenge (unfilled and filled circles, resp.).

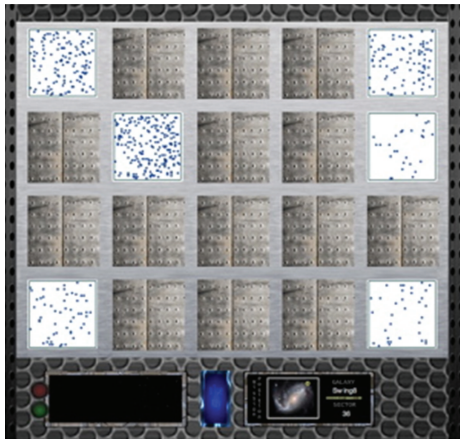

Figure 4: The decision challenge as it appeared in the EMFants game.

fired their blue-ray, described as an intense pulse of long-wavelength radiation, to destroy the EMFant colony. Players were informed that speed was essential to prevent EMFants spreading to other galaxies. Players were also instructed that accuracy was essential, since they only had one chance in each mission to fire the blue-ray, and if they did not destroy the colony of the queen, the EMFants would duplicate and spread to other galaxies.

Once each decision challenge began it proceeded in a statistically identical manner to the original simple experiment described previously. When a player selected a square the entire display quickly flashed blue as the blue-ray fired, followed by an outline of green (for a correct answer) or orange (for an incorrect answer) on the selected EMFant colony. A correct answer was accompanied by a sound of a cheering crowd. An incorrect answer produced a disappointed "urrrgghh."

The EMFants game was completed online by 28 first year psychology students from the University of Newcastle. Each player completed 140 decision challenges, from which we recorded response time and decision accuracy. The parameters for the decision challenges were the same as the original challenge. Figure 3 demonstrates that player performance from the EMFants game is almost identical to the data from the simple version of the decision challenge. We, therefore, used the combined data from 59 players and over 8000 decision challenges to build our player model.

7. Modelling Players

Having collected adequate performance data, the next step of the process was to design an adaptive model of players that could both recognise the player's ability level and their risk-taking.

Assessing a player's risk level is not quite as straightforward as it sounds because of the interplay between risk taking and underlying ability. For example, a player who responds faster than average might be relatively risky, or they might instead have a very high ability level and are thus able to make fast responses without taking undue risks.

To disentangle risk level and ability, we apply a particle filter model to our empirical data to represent the player. Particle filtering is a recent development in cognitive theory [29] and provides a novel way of measuring a player's risk-taking profile and their underlying ability, without changes in one of these two constructs contaminating measurement of the other construct.

Particle filters are sequential Monte Carlo methods that approximate Bayesian posterior distributions. Particle filters allow estimated posterior distributions to be updated as new data arrive. These update algorithms do not require integration over the entire history of observed data (as in other integration methods, such as Markov Chain Monte Carlo). The calculations, therefore, remain psychologically plausible since they do not become increasingly taxing each time new data are observed.

A particle filter begins with a set of particles, each of which is treated as a sample from the posterior distribution of interest. For example, in our game challenge each particle represents a "guess" about which of the K choice alternatives is the correct target. On each frame of the game challenge, the particles are "evolved" to incorporate the new data that arrive regarding the fill rates of the squares. This evolution step usually involves resampling the particles according to their likelihood. Particles consistent with the new datum have a higher probability of being resampled. In contrast, unlikely particles are inconsistent with the observed datum and hence become rare over time.

We used the particle filter developed by Hawkins et al. [29] to model data from the game challenges. This particle filter model is illustrated conceptually in Figure 5. The particle filter model includes a mechanism to track the probability that each response option is the true target. This mechanism corresponds to the player's ability to differentiate the fill rates of each alternative and so detect evidence about which square is filling the fastest. A higher level of ability is represented by more particles. The particle filter also contains a decision mechanism to trigger a response based on the evidence probabilities. In terms of game players, the player's risk profile is captured in this response triggering mechanism. Higher risk players require less evidence than cautious players.

In the model, each particle holds a number from $1 - K$ corresponding to a belief about which square is the target. At the beginning of a decision, particles are randomly sampled from a uniform prior distribution. An illustrative set of $P = 10$ particles for a decision between $K = 4$ alternatives is shown in the top row of the right-hand side of Figure 5. In this example, three particles hypothesize that square 1 is the target (which it actually is), two particles that the target is 4 and so on.

On each frame of the game challenge, a fill event either occurred or did not occur in each square, and these are represented by the "evidence increments" in the shaded rectangle on the left of Figure 5. The uppermost row illustrates that on the first time step of the decision challenge a dot appeared in both of squares 3 and 4, but not in squares 1 or 2. The probability of this sequence of dots across the squares can easily be calculated under the hypothesis of each particle (assuming the true target and distracter fill rates are perfectly known). These probabilities are used to resample a new set of P particles for the next time step, with replacement. The outcome of this resampling is shown by the second row of particles.

After each time step of the decision challenge, the particle filter estimates the posterior probability that each square is the target by calculating the proportion of particles representing that square, illustrated by the histograms on the far right side of Figure 5. These probability estimates represent the output of the evidence tracking mechanism. The number of particles in the filter controls the performance level, which is analogous to the player's ability. More particles make for better performance as this represents a larger sampling size and so a better approximation to the actual fill rates of the target and distracter squares.

The model predicts that a response is triggered whenever the largest posterior probability exceeds a criterion threshold (c). This criterion parameter determines the risk profile of the model, because a high probability threshold requires a lot of evidence to make a decision, so responses are slow but accurate (and vice versa for low probability thresholds). For example, in Figure 5, if the threshold was set at $c = 0.8$ the particle filter would have incorrectly responded (with square 4) after the fourth time step, since eight out of ten particles represented square 4 at that time.

For any particular ability level (i.e., number of particles, P) and risk profile (i.e., decision threshold, c), the particle filter model predicts a particular combination of accuracy and mean response time. By comparing these predictions to measurements from the game, we can abstract from raw data measurements (accuracy and response time) to the deeper psychological constructs of real interest: player ability and risk profile.

Figure 6 illustrates this process using data from our experiment. Each player's data are represented on the graph by a single plot point-determined by their mean accuracy (y-axis) and mean response time (x-axis). The grey lines on the graph show the particle filter's predictions for varying parameters. The close-to-straight lines show the predictions for a fixed level of risk (either a low, medium, or high value of the threshold parameter) and varying ability levels. The curved lines show the converse—different levels of risk for fixed ability. Comparing data against these predictions allows easy categorisation of player ability and risk. For example, data falling above the top curved line indicate very high ability, and data falling to the right of the right-most straight line indicate very cautious risk settings.

Figure 6 shows that the data from our experiments almost all fall nicely within the range of data patterns that the particle filter can predict, which suggests that the model provides a useful description of performance in this task. Note that the data from the two experiments (simple challenge and EMFants challenge) are represented separately in Figure 6 (unfilled and filled circles, resp.) but there appears to be little difference in terms of accuracy and mean response time.

8. Adjusting Game Difficulty

Having developed a particle filter model of the player, in this section, we demonstrate how to use the model to adjust the game mechanics appropriately.

Dynamic balancing first requires a player to be categorised in terms of their risk taking and ability. Our experimentally collected response data is used as the basis

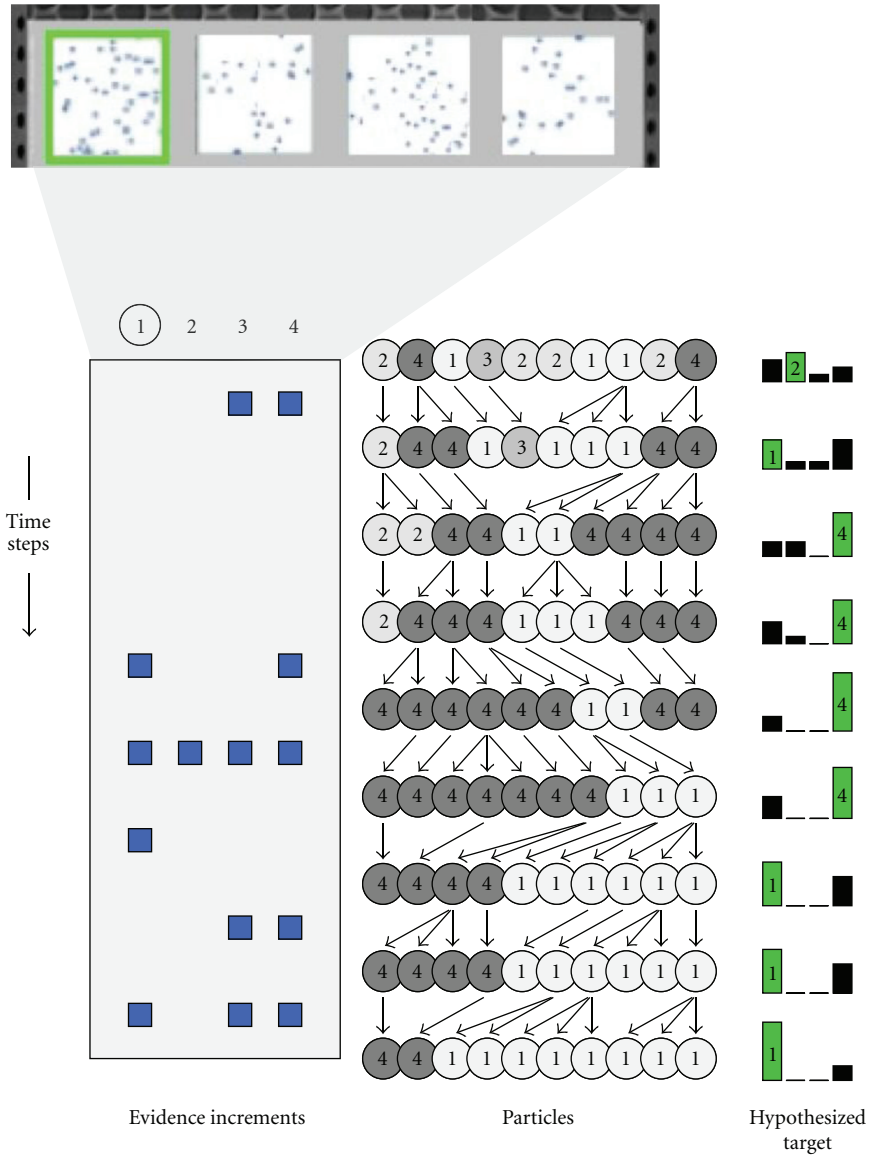

Figure 5: Example of how the particle filter incorporates new evidence on each frame (Figure adapted from Hawkins et al. [29]). Note how the player tracks evidence adjusting their belief at each time step. The number of particles in the model relate to the player's ability to do this tracking. Players will make a choice about the target when evidence exceeds some decision threshold. The amount of evidence required is related to the players risk profile.

of this grouping by considering the player's comparative response time and performance in the task. We can do this by calculating the player's mean accuracy and response time, either from an introductory level or continuously during the game by using a moving average.

The second aspect of dynamic balancing is the alteration of the mechanics to meet the designer's requirements for the game. We should note that having a mechanism for modelling the player in no way reduces the burden on the designer to develop appropriate adaptations to the gameplay. For means of example, we now provide description of one such adaptation that might be useful.

Consider a situation where the designer would like all players to meet an equivalent level of performance. For illustrative purposes, we defined the target performance level as the mean of response times and accuracy across all our experimental players, which was about 8 seconds and 60 percent accuracy. We then use the particle filter to make theoretical predictions of how we should change the probabilities for target and distracter buckets to try to achieve a set level of performance in two example players.

Previously, in Figure 1 we illustrated three basic types of risk taking and the associated player utility. In this work, we took two representative participants, one with high ability

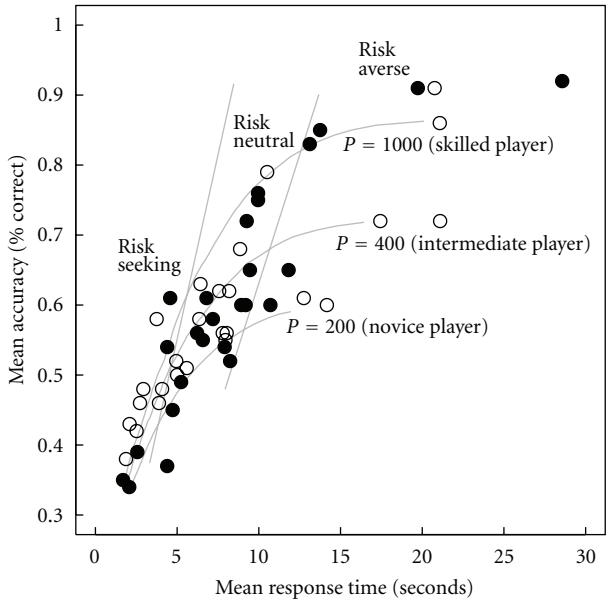

Figure 6: Data and particle filter model predictions from the experiment.

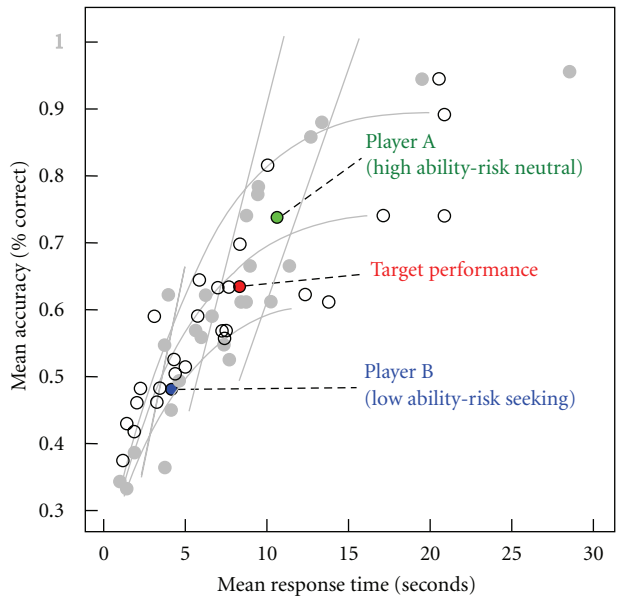

Figure 7: Player ability and risk profile can be measured dynamically to identify particle filter size (P) and decision threshold (c) parameters. These parameters can be used to adjust the degree of similarity between distracters and target and move different types of players to the same target performance.

who was risk neutral and one with low ability who was risk-seeking, marked in Figure 7 with the green and blue dots, respectively.

To approach the target performance level, we want the risk-seeking low ability player to increase mean task accuracy from their current score of about 0.48 to 0.6, and increase mean response time from about 4 seconds to 8 seconds. To increase response time, we need to reduce the fill probabilities (so squares fill more slowly), and to increase accuracy, we need to increase the difference between target and distracter fill probabilities (so the target "stands out" more than the distracters, and the participant makes fewer errors). For this participant's ability level and risk profile, the particle filter predicts that we should change the target fill rate from 50 percent to 18 percent, and distracter fill rates from 40 percent to 10 percent.

In contrast, for the risk neutral high-ability player to approach the target performance level, we want to decrease mean task accuracy from their current score of about 0.72 to 0.6, and decrease mean response time a little, towards 8 seconds. To maintain response time, we need to increase the overall fill probabilities a little (so squares fill a little more quickly), and to decrease accuracy, we need to decrease the difference between target and distracter fill probabilities (so the target is harder to differentiate from the distracters, and the participant makes more errors). For this participant, the particle filter predicts that we change the target fill rate from 50 to 80 percent, and distracter fill rates from 40 to 72 percent.

9. Discussion

Particle filters have considerable potential as models of cognition, and in particular decision making, as they involve updating beliefs about the state of the environment as evidence accumulates over time. By varying the number of particles in the filter the model can approximate a range in performance. For example, a large number of particles can model statistically optimal behaviour, while a smaller number generates predictions similar to flawed, human-like behaviour.

We investigated the use of particle filters for modelling players who undertake a decision-making challenge, where there are multiple alternatives and the player accumulates evidence over time. For such a task, it is desirable to understand players' risk versus caution profile, as well as their underlying ability level. Neither of these underlying psychological properties can be unambiguously inferred from raw observation of accuracy and response time.

Rather than developing ad hoc methods of combining the two measures into some composite, we have imported an appropriate decision theory from cognitive psychology that supports fast and efficient estimation of players' risk profile and ability. This approach does require some previous testing with players to gather empirical data. In this work, we illustrated this approach by gathering experimental data about player performance and using this to develop a model of player performance. Using this model, we demonstrated how difficulty level could be adapted during gameplay using the predictions of such a particle filter.

In this case, particle filters provided an efficient mechanism to develop such dynamic player models, providing parameters, P (the number of particles), and c (the decision threshold) that intuitively relate back to the underlying decision challenge. This intuitive relationship can be important for the designer, where a good understanding of

the underlying mechanics can assist with other adaptations required in the gameplay. Making such online gameplay adaptations is an important step towards designing and developing games that ensure maximal reward for the widest possible range of player profiles.

References

[1] K. Schwalbe, *Information Technology Project Management*, Thomson Course Technology, Boston, Mass, USA, 4th edition, 2006.

[2] M. K. Pitt and N. Shephard, "Filtering via simulation: auxiliary particle filters," *Journal of the American Statistical Association*, vol. 94, no. 446, pp. 590–599, 1999.

[3] C. Bererton, "State estimation for game AI using particle filters," in *Proceedings of the 19th National Conference on Artificial Intelligence*, pp. 36–40, San Jose, Calif, USA, July 2004.

[4] R. Levy, F. Reali, and T. L. Griffiths, "Modeling the effects of memory on human online sentence processing with particle filters," in *Proceedings of the 22nd Annual Conference on Neural Information Processing Systems*, 2009.

[5] A. N. Sanborn, T. L. Griffiths, and D. J. Navarro, "A more rational model of categorization," in *Proceedings of the 28th Annual Conference of the Cognitive Science Society*, 2006.

[6] S. D. Brown and M. Steyvers, "Detecting and predicting changes," *Cognitive Psychology*, vol. 58, no. 1, pp. 49–67, 2009.

[7] S. K. M. Yi, M. Steyvers, and M. Lee, "Modeling human performance in restless bandits with particle filters," *Journal of Problem Solving*, vol. 2, pp. 81–101, 2009.

[8] Inc. Atari, "Adventure," [Atari 2600], USA: Inc. Atari, 1979.

[9] Inc. Atari, "Asteroids," [Atari 2600], USA: Inc. Atari, 1981.

[10] Id Software, "Quake," [PC], USA: GT Interactive, 1996.

[11] Bungie, "Halo," [Xbox], USA: Microsoft Game Studios, 2001.

[12] Capcom, "Devil May Cry," [PlayStation 2], USA: Capcom, 2001.

[13] Apogee Software, "Duke Nukem," [PC], USA: Apogee Software, 1991.

[14] Remedy Entertainment, "Max Payne," [Windows], USA: Gathering Of Developers, 2001.

[15] Crytek, "Far Cry," [Windows], France: Ubisoft, 2005.

[16] Valve Corporation, "Left 4 Dead," [Windows], USA: Valve Corporation, 2008.

[17] Nintendo, "Mario Kart," [Super Nintendo], Japan: Nintendo, 1992.

[18] O. Missura and T. Gärtner, "Player modeling for intelligent difficulty adjustment," *Lecture Notes in Computer Science*, vol. 5808, pp. 197–211, 2009.

[19] P. Demasi and A. Cruz, "Online coevolution for action games," in *Proceedings of the 3rd International Conference on Intelligent Games And Simulation (GAME-ON'2002)*, pp. 113–120, London, UK, 2002.

[20] R. Hunicke and V. Chapman, "AI for dynamic difficulty adjustment in games," in *Proceedings of the 19th National Conference on Artificial Intelligence*, pp. 91–96, San Jose, Calif, USA, July 2004.

[21] R. Hunicke, "The case for dynamic difficulty adjustment in games," in *Proceedings of the ACM SIGCHI International Conference on Advances in Computer Entertainment Technology (ACE '05)*, pp. 429–433, June 2005.

[22] P. Spronck, I. Sprinkhuizen-Kuyper, and E. Postma, "Difficulty scaling of game AI," in *Proceedings of the 5th International Conference on Intelligent Games and Simulation (GAME-ON'2004)*, p. 33, London, UK, 2004.

[23] G. Andrade, G. Ramalho, H. Santana, and V. Corruble, "Challenge-sensitive action selection: an application to game balancing," in *Proceedings of the IEEE/WIC/ACM International Conference on Intelligent Agent Technology*, pp. 194–200, Compiègne, France, September 2005.

[24] G. N. Yannakakis and J. Hallam, "Towards capturing and enhancing entertainment in computer games," in *Proceedings of the 4th Hellenic Conference on Artificial Intelligence*, pp. 432–442, Springer, Heraklion, Greece, 2006.

[25] W. E. Hick, "On the rate of gain of information," *Quarterly Journal of Experimental Psychology*, vol. 4, pp. 11–26, 1952.

[26] R. Hyman, "Stimulus information as a determinant of reaction time," *Journal of Experimental Psychology*, vol. 45, no. 3, pp. 188–196, 1953.

[27] W. H. Teichner and M. J. Krebs, "Laws of visual choice reaction time," *Psychological Review*, vol. 81, no. 1, pp. 75–98, 1974.

[28] A. T. Welford, *Reaction Times*, Academic Press, London, UK, 1980.

[29] G. Hawkins, S. D. Brown, M. Steyvers, and E.-J. Wagenmakers, "A particle filter account of multi-alternative decisions," in *Proceedings of the 43rd Annual Meeting of the Society for Mathematical Psychology*, 2010.