

The Role of Sensorimotor Contingencies and Eye Scanpath Entropy in Presence in Virtual Reality: A Reinforcement Learning Paradigm

Esen Küçüktütüncü, Francisco Macia-Varela, Joan Llobera, and Mel Slater, *Member, IEEE*

Abstract - Sensorimotor contingencies (SC) refer to the rules by which we use our body to perceive. It has been argued that to the extent that a virtual reality (VR) application affords natural SC so the greater likelihood that participants will experience Place Illusion (PI), the illusion of 'being there' (a component of presence) in the virtual environment. However, notwithstanding numerous studies this only has anecdotal support. Here we used a reinforcement learning (RL) paradigm where 26 participants experienced a VR scenario where the RL agent could sequentially propose changes to 5 binary factors: mono or stereo vision, 3 or 6 degrees of freedom head tracking, mono or spatialised sound, low or high display resolution, or one of two color schemes. The first 4 are SC, whereas the last is not. Participants could reject or accept each change proposed by the RL, until convergence. Participants were more likely to accept changes from low to high SC than changes to the color. Additionally, theory suggests that increased PI should be associated with lower eye scanpath entropy. Our results show that mean entropy did decrease over time and the final level of entropy was negatively correlated with a post exposure questionnaire-based assessment of PI.

Index Terms— Presence, virtual reality, sensorimotor contingencies, reinforcement learning, entropy, eye tracking, scanpath

INTRODUCTION

Virtual Reality systems were first invented in the 1960s, and came to public prominence in the 1990s, and then again in the past decade when new hardware and software were developed – with high performance and at a fraction of the cost of the 1990s systems. However, throughout this time some of the scientific questions have remained the same. One of the principle issues is the illusion of 'presence' first discussed in the early 1990s [1]. Presence is the illusion of being in the virtual environment - 'being there' – the scientific puzzle being that in spite of obvious differences to reality (for example, the level of visual realism, the poor resolution of displays compared to normal vision in the physical world, the relatively imprecise representations of physics if any, the non-realistic animations applied to humanoid characters, amongst others) people nevertheless tend to experience a strong illusion of presence. They feel and act as if the virtual world that is displayed to their senses is the actual world in which they are participating. So, the first question concerns understanding the causes of presence. The second issue is how to measure it. Typically, this is carried out with a questionnaire or physiological measures, for example [2-4]. While there are other methods, each has its own problems. For example, measuring behavioral and physiological responses to events in the VR [5] requires specific events to be deliberately added that may cause these responses, which may not be appropriate to the application.

Slater [6] put forward a theory for the occurrence of presence as 'being there' – referred to as 'Place Illusion' (PI) the illusion of being in the virtual place. The idea is that to the extent that the VR system

affords perception through natural sensorimotor contingencies [7] so the probability is enhanced that participants will experience PI. Sensorimotor contingencies refer to how the body is used for perception. For example, moving the upper body and head to look behind objects, reaching out, bending down to look underneath something, looking around by turning the head and the eyes, amongst others. If the VR system allows perception through the use of the body similarly to how this is carried out in physical reality, then the simplest hypothesis for the brain to adopt is that what you perceive signifies where you are. This not only refers to the act of moving the body, but how perception changes must also match expectations from reality. For example, if you turn your head to the right then your view of the world (or sounds heard) must update accordingly. You would expect to see near objects in depth, but less depth perception for objects that are further away. If you move your head close to an object and it dissolves into pixels, then this is a failure of sensorimotor contingencies.

VR devices do, to varying extents, support these types of visual and auditory sensorimotor contingencies through the 6 degrees of freedom head tracking, and higher resolution displays with wide field-of-view, and spatialized sound. However, this is a difficult theory to test, because if sensorimotor contingencies are altered for an experiment, then there is a high probability that participants will experience simulator sickness. For example, if we introduced latencies into the head tracking, then such sickness would be likely to occur, although reduced latency has been associated with higher reported presence [8].

The second problem, how to approach the measurement of presence, without reliance solely on questionnaires, or artefacts introduced for physiological measures, is an important issue for the engineering of presence-inducing systems. A method based on the impact on presence of real-time changes in various factors in the scenario (such as the level of visual realism, the field-of-view of the display, and others) was introduced by Slater, et al. [9]. This method offers the participant a number of system configurations, for example: wide field-of-view or narrow field-of-view, realistic visual rendering or simple visual rendering, vision from first- or third-person perspective, having a virtual body or not having one. The original idea of the method is that at any moment the virtual environment is constituted by k such factors $F_{i1}, F_{i2}, \dots, F_{ik}$ forming a configuration C_i . If each factor is binary then $i = 0, 1, \dots, 2^k - 1$. For example, F_{i1} is either 0 (monovision) or 1 (stereovision). The participant first enters the environment with all factors at their

Esen Küçüktütüncü is with the Institute of Neurosciences of the University of Barcelona, Spain. E-mail: esenkucukututuncu@ub.edu

Francisco Macia-Varela is with the University of Barcelona, Spain. E-mail: frankmmacia@gmail.com. Current address: Full-Body Interaction Lab, Universitat Pompeu Fabra, Barcelona, Spain.

Joan Llobera is with Artanim Foundation, Geneva, Switzerland. E-mail: joan.llobera@artanim.ch

Mel Slater is with Institute of Neurosciences of the University of Barcelona, Spain. E-mail: melslater@ub.edu. Corresponding Author. Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org.

Digital Object Identifier: xx.xxxx/TVCG.201x.xxxxxx/.

highest level C_H (e.g., F_{H1} is stereovision) and pays attention to the corresponding sensation of PI. They are also trained on the meaning of the various factors and how to change their levels. Then they enter the environment now with all factors at their lowest or at randomly chosen levels C_0 . At various times during their exposure they are able to change the level of one of the factors (under a cost constraint) in order to attain the sensation of PI that they had during their exposure to C_H . Each such change leads to a transition $C_i \rightarrow C_j$. They continue to make transitions until they declare a *match* with their original sensation of PI while they experienced C_H . From this, over all participants, we obtain a large number of transitions from which we can compute a $2^k \times 2^k$ Markov probability matrix \mathbf{P} with elements p_{ij} , the probability that conditional on the configuration being C_i that the next configuration would be C_j (including the possibility that $i = j$). From these data various interesting probabilities can be computed, for example the equilibrium state of the Markov chain, and conditional probabilities such as $P(\text{match} | C_i)$, the probability of a match being declared given that the current configuration is C_i . We refer to this method as the Multi-Modal Matching (3M) method.

The method was extended in [10] where instead of participants directly choosing which element in a configuration to change, changes were offered to them by a Reinforcement Learning (RL) agent. The agent initially offers configuration changes at random, but then, depending on participant choices (to accept a proposed change or not) eventually converges, meaning that it has estimated probabilities of participants accepting proposed changes, and thus is unlikely to propose changes with low probability of acceptance. Across many participants a consistent pattern of selected configurations emerges. We refer to this as Adaptive 3M (A3M).

In the experiment reported here we used A3M with respect to the problem of assessing factors that contribute to PI. However, all but one of the factors included relate to sensorimotor contingencies. The objective was to study whether the RL would converge to a configuration that is likely to include factors that involve changing sensorimotor contingencies – which were selected to minimize the risk of simulator sickness – but not include the factor that is unrelated to sensorimotor contingencies.

During the exposure, since participants would be selecting or rejecting changes to the configuration in order to reach a higher level of PI, the likelihood of PI occurring on the average across participants, should be increasing. It has previously been proposed, and supported by experimental evidence, that the entropy of gaze direction is inversely associated with PI. This is based on the idea that in stable perceptions the eye scanpath involves repeatedly moving between a few salient points in the environment, with perception highly influenced by top down rather than bottom up processing [11]. Internal top down models, rather than what is in the environment actively drive vision. For example, when observing ambiguous figures, the eye scanpath is different depending on whether the viewer perceives one interpretation or another, even though the figure itself is fixed. This was suggested as one mechanism by which virtual reality ‘works’ even in very basic and simply rendered scenes – people perceive their internal models from a set of minimal cues provided by the environment [12]. Based on this, the theory was proposed that presence as PI is associated with eye scanpath entropy [13] – that the onset of PI is signified by a *reduction* in entropy.

Entropy measures the degree of disorder in a system. Note that entropy is not the same as variance – the variance of eye scanpaths might be high, but if the same set of visually salient areas are visited regularly and repeatedly then the entropy is low. Jordan and Slater [13] showed this by exposing participants to a gradually forming environment with its component triangles randomly becoming visible over time, until finally the environment was perceived by participants as their standing on top of a large column. Head-gaze movement scanpath entropy decreased at approximately the same time as skin conductance measuring arousal increased, suggesting that at the moment that the percept was formed, PI was established. Participants experienced arousal on realizing that they were standing

on top of the column (higher skin conductance and lower entropy). In one control group where participants finally found themselves to be standing at ground level, the entropy decreased at approximately the same time as in the experimental condition but the skin conductance was unchanged. In the experiment described here we used eye tracking to obtain the eye scanpaths and showed that on the average the entropy decreased with time, and that the final entropy level was negatively correlated with PI as assessed from a questionnaire administered after the VR exposure.

1 MATERIALS AND METHODS

1.1 Participants and Ethics

Participants were recruited from the University of Barcelona campus. Thirty were recruited in total but due to technical failures there are 26 full results for the Reinforcement Learning and 29 for the eye scanpath entropy analysis. The experiment was approved by the Comisió de Bioètica de la Universitat de Barcelona, and procedures were carried out in conformance with the approved procedures. Participants gave written and informed consent and were informed that they could leave the experiment at any time without giving reasons or losing benefits. They were paid 10 euros as compensation for participating to the experiment.




Figure 1 - The string quartet scenario. (A) The learning phase where participants were in a room showing some loudspeakers and practiced changing the settings, in the case shown the resolution. (B) An overview of the scenario – this image has been slightly vertically stretched for alignment purposes. (C) The participant chooses whether or not to make a change to the audio. (D) The scene is shown in the alternate color scheme and the participant chooses whether or not to accept this change.

1.2 The scenario

In the main phase of the experiment participants were in a VR that depicts a string quartet playing some classical music, reusing the same scenario as in [14]. At various times during the performance, they were offered the chance to change one factor in the scenario configuration or leave it as it is. This continued for a maximum of 20 minutes. The possible changes for the participant to make were chosen by a Reinforcement Learning algorithm. When the participant had rejected 8 successive proposed changes, and provided that at least 10 minutes had elapsed, this was taken as a sign of convergence, and the scenario ended. The scenario is illustrated in Figure 1. See also Supplementary Video S1 (a higher resolution version is available on https://youtu.be/AUJJqMUR5_E).

1.3 Factors Manipulated

The RL algorithm offered changes to the following factors:

1. Vision: (0) monovision (1) 3D stereovision
2. Parallax: (0) 3 degrees of freedom head-tracking (1) 6 degrees of freedom head-tracking
3. Audio: (0) Mono audio (1) Spatialized audio.
4. Resolution: (0) 1282 × 672 (1) 3664×1920, the normal resolution of the HMD
5. Color: (0) Original color (1) Alternative color.

Items 1-4 are different aspects of sensorimotor contingencies, whereas item 5 is not. It was expected that on items 1-4, changes from (0) to (1) would be far more likely to be accepted than (1) to (0), whereas in item 5 there will be an approximately equal chance of selecting (0) or (1), since these are not related to sensorimotor contingencies.

There are 32 configurations and after each suggestion by the RL there is a potential transition from one configuration to another. The configurations are shown in the Table 1. A transition occurs when there is a change from one configuration in the set (0,...,31) to another in the same set. This also includes null transitions where there is no change.

Table 1 – The 32 configurations

No.	Meaning	Binary representation
0	Monovision - No Parallax - Mono Audio - Low Resolution - Original Color	00000
1	Monovision - No Parallax - Mono Audio - Low Resolution - Alternative Color	00001
2	Monovision - No Parallax - Mono Audio - High Resolution - Original Color	00010
3	Monovision - No Parallax - Mono Audio - High Resolution - Alternative Color	00011
4	Monovision - No Parallax - Spatial Audio - Low Resolution - Original Color	00100
5	Monovision - No Parallax - Spatial Audio - Low Resolution - Alternative Color	00101
6	Monovision - No Parallax - Spatial Audio - High Resolution - Original Color	00110
7	Monovision - No Parallax - Spatial Audio - High Resolution - Alternative Color	00111
8	Monovision - Parallax - Mono Audio - Low Resolution - Original Color	01000
9	Monovision - Parallax - Mono Audio - Low Resolution - Alternative Color	01001
10	Monovision - Parallax - Mono Audio - High Resolution - Original Color	01010
11	Monovision - Parallax - Mono Audio - High Resolution - Alternative Color	01011
12	Monovision - Parallax - Spatial Audio - Low Resolution - Original Color	01100
13	Monovision - Parallax - Spatial Audio - Low Resolution - Alternative Color	01101
14	Monovision - Parallax - Spatial Audio - High Resolution - Original Color	01110
15	Monovision - Parallax - Spatial Audio - High Resolution - Alternative Color	01111
16	Stereovision - No Parallax - Mono Audio - Low Resolution - Original Color	10000
17	Stereovision - No Parallax - Mono Audio - Low Resolution - Alternative Color	10001
18	Stereovision - No Parallax - Mono Audio - High Resolution - Original Color	10010
19	Stereovision - No Parallax - Mono Audio High Resolution - Alternative Color	10011
20	Stereovision - No Parallax - Spatial Audio - Low Resolution Original Color	10100
21	Stereovision - No Parallax - Spatial Audio - Low Resolution - Alternative Color	10101
22	Stereovision - No Parallax - Spatial Audio - High Resolution - Original Color	10110
23	Stereovision - No Parallax - Spatial Audio - High Resolution - Alternative Color	10111
24	Stereovision - Parallax - Mono Audio - Low Resolution - Original Color	11000
25	Stereovision - Parallax - Mono Audio - Low Resolution - Alternative Color	11001
26	Stereovision - Parallax - Mono Audio - High Resolution - Original Color	11010
27	Stereovision - Parallax - Mono Audio - High Resolution - Alternative Color	11011
28	Stereovision - Parallax - Spatial Audio - Low Resolution - Original Color	11100
29	Stereovision - Parallax - Spatial Audio - Low Resolution - Alternative Color	11101
30	Stereovision - Parallax - Spatial Audio - High Resolution - Original Color	11110
31	Stereovision - Parallax - Spatial Audio - High Resolution - Alternative Color	11111

1.4 Equipment

For this experiment the Pico Neo 3 Eye head-mounted display (HMD) was used, with a resolution of 3664 × 1920, and 773 pixels per inch, and it has field of view is 98°. The headset has a refresh rate of 90Hz and is equipped with 2 eye tracking cameras, each being 400 × 400 running at 90 Hz. The controllers of the headset were used for interaction, which have 6DoF optical positioning and linear resonant actuators. For audio equipment Skullcandy Riff Wireless Headphones were used in order to ensure comfort and immersion.

1.5 Procedures

The experiment took place at the University of Barcelona (UB), Mundet Campus. It consisted of one session. Information about the experiment was told verbally to the participant by the experimenter and also explained in a Qualtrics Questionnaire (www.qualtrics.com). After reading the information sheet the participants were informed about their freedom to withdraw from the experiment at any time and were asked to sign a consent form.

1.6 Questionnaire

The questionnaire was divided into two parts, prior to entering the virtual environment (Pre Questionnaire in Table 2) which was concerned with demographic and background issues, and after completing the VR experience (Post Questionnaire in Table 2).

The questionnaire included demographic questions, level of programming and familiarity with VR experiences. After completion of the Pre Auestionnaire the experiment started. In the first phase of the experiment the participants were immersed in a reduced

environment that did not include the virtual musicians (Figure 1A). There, each of the configurations were explained with a written description, and they were able experience each of the 2 possible choices using a pop-up panel on their left arm. Once they had tried both of the options, they were asked to press the ‘I understand both options’ button and then proceeded to the next factor. Each factor that they could change was represented by an icon, with complete list and their meanings shown in Table 3.

After this learning period was completed, the scene faded into the same environment but now with 4 musicians playing in front of them, which started the main phase. The experimental factors were randomized at the beginning of this phase, so each participant started with a randomly chosen configuration from the list in Table 1. Their task was to listen to and watch the musical performance. At uniformly random intervals between 35 to 45 seconds, they would feel the left hand controller vibrate and a corresponding sound effect, where a virtual panel would pop up on their wrist and give them the option to either change the current configuration or leave it as it is. During this task, the interface presented choices to the participants using the method presented during the pre-immersion phase, except that the button ‘I understand both options’ was replaced by ‘I feel more present’.

As indicated above they could decide to continue with the current configuration or change it to the alternative offered. For example, if the scene was in low resolution, they could either keep it the that way by selecting the low-resolution option or change it to high resolution. The factor that was displayed was selected by the RL agent and it continued until convergence, or the 20 minute maximum time, with the application ending by fading out.

Table 2 – Response variables from the questionnaire. IQR is Interquartile Range.

Variable Name	Question	Result
Pre Questionnaire		
gender	What is your gender? Male (0) Female (1) Non-binary (2) Other (3) Prefer not to say (4)	12 17
age	18-24 (1) 25-34 (2) 35-44 (3) 45-54 (4) 55-64 (5) 65-74 (6) 75- (7) Prefer not to say (8)	12 12 4 1
		Median (IQR)
vr	What is your level of expertise with programming for VR? 1 not at all ... 7 expert	1 (1)
vr_design	What is your level of expertise when designing an experience for VR? 1 not at all ... 7 expert	1 (0)
vr_eval	What is your level of expertise on evaluating VR experiences? 1 not at all ... 7 expert	1 (4)
Post Questionnaire		
	During the virtual reality session, you were asked to make decisions that changed the environment or how you interacted with it. These questions are about those decisions.	
decisions	How confident were you in general making these decisions? 1 not at all ... 7 very confident	6 (2)
	Please rank them [the 5 factors] in order of how important they were to improving your feeling of being in the virtual environment.	
order_stereo	rank from 1 to 5	2 (2)
order_parallax	rank from 1 to 5	3 (1)
order_sound	rank from 1 to 5	3 (2)
order_res	rank from 1 to 5	1 (2)
order_color	rank from 1 to 5	5 (2)











PI Questions These refer to the last 5 minutes of the experience:		
there	Please rate your feeling of being in the virtual environment, from 1 to 7, where 7 represents your normal feeling of being in a physical place. 1 - Not like a physical place 7 - Like a physical place	5 (1)
reality	To what extent were there times when the virtual environment was reality for you? 1 - Never 7 - Almost all the time	5 (4)
place	When you think about the experience, does it feel more like images you saw, or a place you visited? 1 - Images I saw 7 - Place I visited	5 (2)
virtualplace	During the experience, which was stronger, your feeling of being in the virtual environment or in the real laboratory? 1 - Real Laboratory 7 - Virtual Environment	6 (1)

Then, they completed the post-experiment questionnaire and were asked to order the experimental factors from the most important to the least. They were also asked to leave comments regarding the experiment. Overall, the entire process lasted between 35 to 45 minutes with 20 minutes being the VR exposure.

1.7 The Reinforcement Learning Method

The Reinforcement Learning (RL) method worked in the same way as in [10]. Each configuration listed in Table 1 was associated with a state of the RL algorithm. We used Q learning from Sutton and Barto [15] (Chapter 6). For every configuration there were 5 possible actions associated, one for each configuration change. The algorithm attempted to find the configuration change that the participant was most likely to accept. Accepted changes constituted a positive reward for the algorithm.

Table 3 – Icons used to describe each possible change. Level 0 represents the lower SC and level 1 the higher except for Color.

Experimental factor	Level 0	Level 1
Vision		
Head-tracking		
Audio		
Resolution		
Color		

For the RL method, we chose the same configuration values as in [10]: $l = 0.2$, for the learning rate; $\gamma = 0.15$, as the discount factor. All values in the Q-table were initialized to 0. In addition, since it was possible that different participants had different preferences, and the numbers of episodes per participant were variable, we chose to run the RL algorithm separately for each participant. The only differences were that there were 5 binary factors instead of 4 and the number of configurations visited

changed between participants, since the experiment would stop when the participant rejected 8 successive proposed changes.

1.8 Eye Scanpath Entropy

Eye scanpaths were obtained by analyzing the raw data from the Tobii Eye-tracking system in the HMD. The system consists of illuminators placed in a ring-like structure around the HMD lenses, that create a pattern of light reflections on the participants' eyes. This reflection pattern is captured by high-resolution cameras and processed in real-time by machine learning algorithms. This process occurs at a rate of 90 Hz. This raw eye direction data is converted into a three-dimensional vector through the HMD's SDK, and can be used with the head position and rotation data from the HMD in order to calculate where a participant is looking in the virtual space at any given point.

There were two predictions with respect to eye scan paths:

1. Mean entropy over all participants would decrease over time (with increasing PI due to manipulation of the factors).
2. There would be a negative correlation between the final level of entropy per participant and their subjectively reported level of PI as measured by the questionnaire after the experience.

The entropy was calculated as:

$$e = - \sum_{i=1}^s p_i \log(p_i)$$

where there are s possible 'states' and p_i is the probability of being in the i th state. The states are defined quite straightforwardly as a partition of the space in front of the participant into equal angle intervals (Figure 2). Hence p_i is the proportion of times that the target of the eye tracking entered into the i th state during the time period concerned. For example, if the eyes hardly moved and stayed in one interval during the time period, then $e = 0$. If the eyes moved uniformly randomly across all segments, then the probabilities $p_i = 1/s$ and $e = \log(s)$.

The overall time is also partitioned into successive segments, and entropy calculated for each participant for each segment.

1.9 Statistical Methods

We first present results descriptively in order to provide an easy to follow overview of the outcomes. We use Bayesian methods for formal statistical analysis. Using classical null hypothesis significance tests would run into the problem that there would be many tests, which undermines the meaning of the significance level due to these multiple comparisons. Although there are *ad hoc* methods to overcome this problem (such as Bonferroni corrections) there is no such problem with Bayesian methods since there is one

overall model that includes all parameters, and as many probability statements as required can be drawn from their posterior joint distribution without diminishing their validity. In order to carry out the analysis we used the Stan statistical programming language [16] in RStudio (<https://posit.co/download/rstudio-desktop/>) with the rstan interface (<https://mc-stan.org/users/interfaces/rstan>). Some of the graphs were constructed with Stata (www.stata.com) (version 16.1).

2 RESULTS

2.1 Participant statistics

The characteristics of the participants are shown in the Pre Questionnaire section of Table 2. There were 5 more females than males, the majority were aged between 18 and 34 and they had very little prior experience of VR. They had little technical knowledge (a low level of computer programming).




Figure 2 – Partition of the space into equal angle sectors.

2.2 Ordering of the factors

Overall participants were confident in their decisions in responses to requests for change, where the median level of confidence was 6 (out of a maximum 7) with interquartile range 2 (the variable *decisions* in Table 2). Figure 3 shows box plots of the subjective ordering of the 5 factors with respect to their importance. It is clear that color is considered the least important, and resolution the most important, followed by vision.




Figure 3 – Box plots of the ordering of the importance of the factors from the questionnaire. The thick black horizontal lines are the medians, the boxes are the interquartile ranges (IQR), and the whiskers range from lower quartile – max(smallest value, 1.5*IQR), and upper quartile + min(max value, 1.5*IQR). Points outside these limits are shown individually.

2.3 The Place Illusion questions

The box plots for the PI questions (Table 2) are shown in Figure 4. It can be seen that the subjective level is high, with all medians at

least 5, with the highest being *virtualplace*. The variable *overall* refers to the medians across the 4 questionnaire variables. The entire IQR for this is greater than the mid-point of the scale.

2.4 Reinforcement Learning

The mean \pm SD of the number of proposals per individual for changes was 16.7 ± 4.4 . There were 433 transitions over the 26 participants. Table 4 shows the basic outcomes. For example, with respect to the Vision factor, when the current configuration included the Monovision state and a change was proposed there were 6 rejections, and 10 acceptances of the change. When it included the Stereovision state, there were 20 rejections and 2 acceptances.

Table 4 also shows the different responses to a proposal for a change by the different factors. For example, in the case of Vision, when the current configuration included Monovision 62% of the proposals for change were accepted, but when the configuration included Stereovision only 9% of proposals for a change were accepted. It is notable that a configuration that included Monosound always transitioned to one that included Stereosound. With respect to Color, proposals for change had the lowest proportions of acceptances.




Figure 4 – Box plots for the PI questions from Table 2. Overall refers to the median of the 4 questions.

2.5 Terminating Configurations

Table 5 shows the frequency distribution of the terminating configurations. The highest frequency configuration is the one that contains all of the sensorimotor factors but not the color. Note that Spatial Audio appears in all 8, High Resolution appears in 7/8 and all the rest appear in 4/8.

Table 4 – Frequency distribution of the response to a proposal where $a \rightarrow b$ means that the corresponding factor was in the a (0,1) state and the change was b (0,1) where 0 is rejected and 1 is accepted. The 'Off' and 'On' columns refer to the proportion of accepted changes when the factor is off (e.g., Monovision, No Parallax) or on (Stereovision, Parallax).

Factor	0→0	0→1	1→0	1→1	Total	Off	On
Vision	6	10	20	2	38	0.62	0.09
Parallax	15	16	45	8	84	0.52	0.15
Sound	0	15	69	2	86	1.00	0.03
Resolution	4	11	78	1	94	0.73	0.01
Color	49	24	30	28	131	0.33	0.48

2.6 Markov Chain Analysis

Every proposed change results in a transition, even if it is not to accept, and therefore the configuration is unchanged. From the

collection of all transitions across all participants we can build a Markov Transition matrix P , where the elements

$$P_{ij} = P(\text{transition to configuration } j | \text{configuration } i)$$

the probability of transitioning from configuration i to j . This is a 32×32 sparse matrix, given the small number of trials relative to the total number of transitions possible. Moreover, the matrix shows that some configurations were never reached, and there was never a transition from them. These eliminated configurations are shown in Table 6. This reduces the size of the matrix to 27×27 .

Table 5 – The frequency distribution of the terminating configurations in decreasing order of frequency

No.	Config.		Freq	Prop.
30	11110	Stereovision - Parallax - Spatial Audio - High Resolution - Original Color	9	0.35
31	11111	Stereovision - Parallax - Spatial Audio - High Resolution - Alternative Color	5	0.19
14	01110	Monovision - Parallax - Spatial Audio - High Resolution - Original Color	4	0.15
22	10110	Stereovision - No Parallax - Spatial Audio - High Resolution - Original Color	3	0.12
15	01111	Monovision - Parallax - Spatial Audio - High Resolution - Alternative Color	2	0.08
5	00101	Monovision - No Parallax - Spatial Audio - Low Resolution - Alternative Color	1	0.04
6	00110	Monovision - No Parallax - Spatial Audio - High Resolution - Original Color	1	0.04
23	10111	Stereovision - No Parallax - Spatial Audio - High Resolution - Alternative Color	1	0.04

Suppose π is a 1×27 row-vector of probabilities of being in the configurations. Then πP is the probability distribution over the configurations after a single transition. In general, P^k for any positive integer k is the k -step transition matrix, i.e., the (i, j) th element is the probability of being in configuration j , k transitions after starting in configuration i . Hence πP^k is the probability distribution over the configurations after k transitions. For increasing values of k we can consider whether πP^k reaches a stable state, in other words an equilibrium distribution has been reached so that further transitions make no difference to this probability distribution. The equilibrium distribution can be found by solving:

$$\pi P = \pi$$

reflecting the idea that if π is the equilibrium distribution then the application of a further transition matrix does not change it.

Table 7 shows the equilibrium probability distribution over the configurations. There are 6 configurations with probabilities of at least 0.05. Spatial Audio and High Resolution appear in all 6, Stereopsis and Parallax in 4, and Alternate Color in 3. It can be seen that the sensorimotor contingency variables dominate, although Alternative Color does appear in the second place. We cannot know whether this distribution is by chance, and whether the choice or not of the color change is different from the rest of the factors, though from this table and Table 5 above it would seem so. Next, we turn to a formal statistical analysis to resolve these issues.

Table 6 – Eliminated Configurations – from which there are no transitions and which were never reached from other states.

8	01000	Monovision - Parallax - Mono Audio - Low Resolution - Original Color
17	10001	Stereovision - No Parallax - Mono Audio - Low Resolution - Alternative Color
24	11000	Stereovision - Parallax - Mono Audio - Low Resolution - Original Color
25	11001	Stereovision - Parallax - Mono Audio - Low Resolution - Alternative Color
26	11010	Stereovision - Parallax - Mono Audio - High Resolution - Original Color

Table 7 –Probability distribution of the equilibrium sorted by decreasing probability. Configurations not shown have probability 0.

No.	Config	Meaning	Prob
30	11110	Stereopsis - Parallax - Spatial Audio - High Resolution - Original Color	0.476
31	11111	Stereopsis - Parallax - Spatial Audio - High Resolution - Alternative Color	0.210
22	10110	Stereopsis - No Parallax - Spatial Audio - High Resolution - Original Color	0.090
14	01110	Monocular - Parallax - Spatial Audio - High Resolution - Original Color	0.066
15	01111	Monocular - Parallax - Spatial Audio - High Resolution - Alternative Color	0.064
23	10111	Stereopsis - No Parallax - Spatial Audio - High Resolution - Alternative Color	0.053
6	00110	Monocular - No Parallax - Spatial Audio - High Resolution - Original Color	0.011
7	00111	Monocular - No Parallax - Spatial Audio - High Resolution - Alternative Color	0.008
19	10011	Stereopsis - No Parallax - Mono Audio High Resolution - Alternative Color	0.005
27	11011	Stereopsis - Parallax - Mono Audio - High Resolution - Alternative Color	0.004
21	10101	Stereopsis - No Parallax - Spatial Audio - Low Resolution - Alternative Color	0.003
29	11101	Stereopsis - Parallax - Spatial Audio - Low Resolution - Alternative Color	0.003
10	01010	Monocular - Parallax - Mono Audio - High Resolution - Original Color	0.002
18	10010	Stereopsis - No Parallax - Mono Audio - High Resolution - Original Color	0.002
20	10100	Stereopsis - No Parallax - Spatial Audio - Low Resolution Original Color	0.002
28	11100	Stereopsis - Parallax - Spatial Audio - Low Resolution - Original Color	0.001

Table 8– Predictor variables for accept

Long variable name	Short Name	Factor	Meaning
monovision	V_0	Vision	Mono
stereovision	V_1	Vision	Stereo
noparallax	P_0	Parallax	No Parallax
parallax	P_1	Parallax	Parallax
monoaudio	A_0	Audio	Mono audio
spatialaudio	A_1	Audio	Spatial audio
lowres	R_0	Resolution	Low resolution
highres	R_1	Resolution	High resolution
origcolor	C_0	Color	Original color
alternativecolor	C_1	Color	Alternative color

2.7 Statistical analysis

During the course of the VR experience participants made choices to *accept* or reject the offer of a change. We use as a response variable *accept*, which is 0 when the proposed change is rejected, and 1 when the proposed change is accepted. These offers were in the contexts shown in Table 8.

For example, the variable $V_0 = 1$ when a change to the Vision factor is offered and the current configuration includes Monovision, and $V_0 = 0$ otherwise. Or $R_1 = 1$ when a change to resolution is offered, and the current configuration includes High Resolution. We use these as predictor variables for *accept*. If the idea that sensorimotor contingencies are at the basis of PI is valid then we would expect that the X_0 variables would be associated with an increase in the probability of *accept* = 1 (e.g., participants would be more likely to accept a change from monoaudio to spatialaudio) and the X_1 variables would be associated with a decrease in the probability of *accept* = 1 (e.g., they would be less likely to accept a change from Stereovision to Monovision).

Note that we cannot include all 10 variables simultaneously because of aliasing between them (knowing the values of 9 of them can predict exactly the 10th) therefore we consider the ‘0’ and ‘1’ variables in two equations within the same overall model.

In order to assess the influence of these variables on *accept*, we use a standard Bernoulli logit model. Let $accept_i$ be a binary response variable that indicates that the proposal was accepted and 0 when it was rejected, over all proposals made ($i = 1, 2, \dots, n = 443$).

$$\begin{aligned} P(accept_i = 1) &= \theta_i \\ P(accept_i = 0) &= 1 - \theta_i \end{aligned} \quad (1)$$

where θ_i is the probability of the i th response being *accept*. To relate θ_i to the variables in Table 8 above we use the linear predictors:

$$\begin{aligned} \eta_{0i} &= \beta_{00} + \beta_{01}V_{0i} + \beta_{02}P_{0i} + \beta_{03}A_{0i} + \beta_{04}R_{0i} + \beta_{05}C_{0i} \\ \eta_{1i} &= \beta_{10} + \beta_{11}V_{1i} + \beta_{12}P_{1i} + \beta_{13}A_{1i} + \beta_{14}R_{1i} + \beta_{15}C_{1i} \end{aligned} \quad (2)$$

In the Bernoulli logit model, the relationship between (1) and (2) is through the logit link, i.e.,

$$\log\left(\frac{\theta_i}{1 - \theta_i}\right) = \eta_i \quad (3)$$

or equivalently,

$$\theta_i = \frac{1}{1 + \exp(-\eta_i)}$$

where η_i represents η_{0i} or η_{1i} .

This also guarantees that θ_i is in the range 0 to 1 for all possible values of η_i . The interpretation of β_j is that it is the change in log-odds of *accept* compared to reject for a unit change in the corresponding variable, for all else held constant.

Let the prior distributions for the β parameters be *normal(mean=0, SD=10)*, which results in prior 95% credible intervals ± 20 , i.e., prior to utilizing the observed data there is 0.95 probability of being within these limits. These are wide intervals corresponding to weakly informative priors [17]. The posterior distributions are summarized in Table 9. Notice that the posterior

95% credible intervals are considerably narrower than the prior intervals, reflecting the impact of the data.

It can be seen that there are strong positive associations between *accept* and the ‘0’ variables indicating that when the configuration included these, acceptance of a change was likely. This also holds for original color. However, for the ‘1’ variables there is a strong negative association between these variables and *accept* indicating that when a configuration included these settings a change back to the ‘0’ setting was unlikely. For example, a change in vision would be likely to be accepted if the configuration included monovision, but rejected if it included stereovision. However, and crucially, in the case of *altcolor* there is low level of evidence of a negative association with *accept* (the probability is $1 - 0.367 = 0.633$ of a negative association with *accept*, whereas for all other variables the probability is 1).

Table 9– Summary of the posterior distributions of the parameters showing the mean and standard deviations of the distribution and the 95% credible intervals. Prob > 0 are the probabilities that the parameters are positive.

Parameter	Coefficient of	Mean	SD	2.5%	97.5%	Prob > 0
β_{00}		-1.77	0.17	-2.10	-1.46	0.000
β_{01}	monovision	2.31	0.54	1.27	3.39	1.000
β_{02}	noparallax	1.83	0.40	1.05	2.62	1.000
β_{03}	monoaudio	11.11	5.10	4.38	23.41	1.000
β_{04}	lowres	2.85	0.62	1.70	4.11	1.000
β_{05}	origcolor	1.04	0.30	0.45	1.63	0.999
β_{10}		0.02	0.16	-0.29	0.34	0.560
β_{11}	stereovision	-2.56	0.85	-4.51	-1.17	0.000
β_{12}	parallax	-1.80	0.42	-2.65	-1.01	0.000
β_{13}	spatialaudio	-3.81	0.82	-5.69	-2.46	0.000
β_{14}	highres	-4.87	1.23	-7.91	-3.08	0.000
β_{15}	altcolor	-0.10	0.31	-0.70	0.51	0.367

It is likely that when the configuration included *originalcolor* that participants may have experimented with changing it, hence the high probability of an association (0.999). However, if they had changed it, they would tend not to change it back again, hence the lower probability associated with *altcolor*. Moreover, from the posterior distributions of the β_{0j} parameters, we find that the posterior probabilities $\beta_{05} < \beta_{0j}, j = 1, 2, 3, 4$ is at least 0.966 in every case, meaning that the *originalcolor* variable has the lowest effect size.

The factor *monoaudio* also clearly has the greatest effect size since the posterior probability that $\beta_{03} > \beta_{0j}, j = 1, 2, 4, 5$ is at least 0.998 in each case (this can also be seen through the credible intervals). Participants were most likely to change from *monoaudio* to *spatialaudio*, which is not surprising given the nature of the scenario.

The factor *lowres* has the second highest effect size after *monoaudio*, meaning that it was the second most likely factor to be changed. Moreover, *highres* has the largest (negative) effect size, meaning that it was the least likely to be changed to *lowres*. This also accords with the subjective data shown in Figure 4, where resolution was considered the most important of the 5 factors.

2.8 Entropy of Eye Scanpaths

Considering the space in front of the participant, suppose that the number of states $s = 18$ (each angular segment is 10°) and that the time period is divided into 30 segments (e.g., for a 20 minute session each segment would be 40 seconds). Then we would have for each of the 29 individuals 30 entropy readings. Figure 5A shows the relationship between the mean entropy over the 29 participants for each time period by time.

It can be seen that as predicted there is a strong negative association between entropy and time ($r = -0.66$ with 95%

confidence interval -0.82 to -0.39). However, this does not show that entropy is connected with PI, since entropy could have declined just because participants got used to the environment and their scanpaths became more regular.

In order to examine whether there is a relationship with subjectively reported PI we can make use of the questionnaire variables for PI (Table 2) and consider the overall measure of PI, which here we refer to as PI_median, which is the median of the 4 variables. We use the median rather than the mean since these are ordinal variables. Now we want to see whether the PI_median scores are related to entropy in the period at which participants stopped making any further changes (i.e., after the last change that they made).

Figure 5B shows the entropy by PI_median. Contrary to expectations the relationship seems to be a curvilinear one, with low PI associated with low entropy, which then increases with PI and then curves downwards (a quadratic fit to these data is highly ‘significant’). However, the curvilinear aspect is entirely caused by the 3 observations that have low presence (2 or less). What happens if we eliminate these? Then the relationship is still strong with negative correlation ($r = -0.36$, 95% confidence interval -0.65 to 0.03).

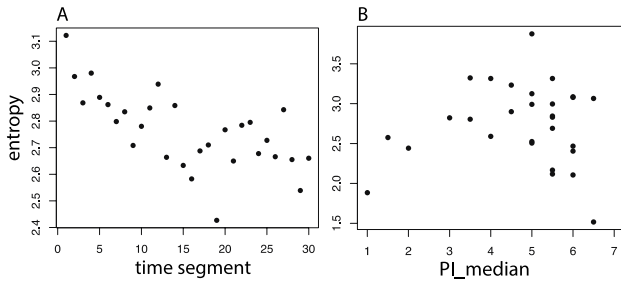


Figure 5 – Mean entropy by time and PI (A) Entropy by time segment over 30 segments (B) Mean entropy by PI_median.

A potential problem with this analysis is that we chose 18 states and 30 time periods. Perhaps the results are dependent on this choice. In order to overcome this possibility, we uniformly randomly chose between 5 and 40 time segments, and independently between 3 and 30 states, and ran the analysis for 1000 such combinations. We then computed the mean entropy over these repetitions. This results in the relationship between the entropy and subjective PI shown in Figure 6.

We see the same pattern as earlier. This is highly ‘significant’ as a quadratic fit. Moreover, if we eliminate the 3 with lowest presence (PI_median < 3), then there is a strong negative correlation between the entropy and the subjective PI: $r = -0.41$, with 95% confidence interval -0.69 to -0.02.

More formally, to analyze this we use the quadratic model:

$$\mu_i = \beta_0 + \beta_1 P_i + \beta_2 P_i^2 \quad (4)$$

$$entropy_i \sim Student_t(v, \mu_i, \sigma)$$

where P_i is the median presence for the i th individual, corresponding to $entropy_i$. We use a Student t distribution for the likelihood (the distribution of the response conditional on the parameters) since for low degrees of freedom (v) this distribution allows greater dispersion than the normal distribution, yet for high degrees of freedom ($v \geq 30$) it closely approximates the normal. The median of the distribution is μ_i (and it is the mean for $v > 1$) and $\sigma > 0$ is the scale factor.

The prior distributions are $\beta_j \sim normal(0,10)$ as before and $v, \sigma \sim Gamma(shape = 2, rate = 0.1)$. The corresponding prior 95% credible intervals for these are 2.4 to 55.7. The Stan program

was run with 3000 iterations and fully converged with results shown in Table 10. It is clear from the 95% credible intervals and the probabilities that the fit is very good. Also compare the very narrow posterior credible intervals with the prior intervals. As an example, the posterior probability $P(\beta_2 < 0) = 1 - 0.002 = 0.998$. Note that v is lower than 30 (with probability 0.79) justifying the use of the Student t distribution rather than the normal.

Figure 6 shows the scatter diagram of the observed entropy plotted against the median presence, and the quadratic (4) with the parameters replaced by their means from the first column of Table 10, illustrating the close correspondence between observed and fitted data. For example, the quadratic of model (4) has maximum at $\hat{P} = -\beta_1/2\beta_2$. Using the means of the posterior distributions of the β parameters, $\hat{P} = 3.75$ which accords well with Figure 6.

Table 10– Summary of the posterior distributions of the parameters of model (4) showing the means and standard deviations of the distribution and the 95% credible intervals. Prob > 0 are the probabilities that the parameters are positive.

Parameter	Mean	SD	2.5%	97.5%	Prob > 0
β_0	1.33	0.43	0.48	2.19	0.997
β_1	0.75	0.24	0.29	1.22	0.999
β_2	-0.10	0.03	-0.16	-0.04	0.002
σ	0.34	0.06	0.24	0.47	
v	20.84	13.70	3.83	56.48	

Using the posterior distributions of the parameters we can generate entirely new pseudo random data on model (4) and compare with the actual values, referred to as the ‘predictive posterior distributions’. For each individual i this results in a posterior distribution for predicted $entropy_i$. We can compare the means of these distributions with the original observed data. The correlation between the means of the predicted posterior distributions and the observed data is $r = 0.59$ (95% confidence interval 0.29 to 0.79).

This connection between the subjective data from the questionnaires and the entropy of the scanpath eye tracking data is fairly remarkable, since these two different scores in principle have nothing to do with one another but are nevertheless related, and in a way partially predicted by the theory. There is insufficient data to know whether the relationship actually has a quadratic shape or whether the three very low PI scores are just anomalies, but the important point is that discounting these three data points there is the negative association as predicted by the theory.

In all of the above analysis we use confidence intervals for the correlation coefficient r as an indication of the strength of the relationships, rather than as formal statistical tests.




Figure 6 – Mean entropy over 1000 iterations with uniformly randomly generated time segments and states. The quadratic fitted curve is based on the means of the posterior distributions of the parameters of the quadratic model Eq (4).

3 DISCUSSION

The theory presented in [6] postulates that the illusion of ‘being there’ (Place Illusion, PI) in the environment depicted by the virtual reality displays is based on affordances for perception in VR being via sensorimotor contingencies (SC) that match those of physical reality. It is a difficult theory to test, since there are many SC rules, and breaking some of them would be likely to result in simulator sickness. Here we presented a method whereby 4 different SC factors could be manipulated (switched) by participants, and the same for another factor that was not a sensorimotor contingency (color). We found that participants were more likely to choose to make transitions between configurations that enhanced SC factors than the change to color. This provides direct evidence for the theory, though of course, only with respect to these 4 factors. On the other hand, had participants made changes to color to the same extent that they made changes to the 4 SC factors then this would have been evidence against the theory.

PI as a direct result of the perception through everyday SC makes sense, since if we perceive in VR using our bodies in the same way as in physical reality, and each act of perception results in predicted changes to perception based on prior experience, then the simplest hypothesis for the brain to make is ‘this is where I am’. When we turn our head and eyes to look to one side, we expect to see what is to the side of us, and for sound to be modulated accordingly. If the wind was blowing in our right ear and we turn our head to the right we would be surprised if the wind still impacted our right ear rather than the front of our face (this would be an example of a breakdown in SC). When these same types of SC occur in VR this is strong evidence to the brain about where we are located (PI). This is not to say that PI cannot sometimes break [18, 19], for example, when the participant is moving around and accidentally collides with a real wall, or when the participant looks closely at an object and the pixels become very obvious. However, after a break, PI is likely to form again, since the SC remain the same [20].

The scanpath theory of visual perception [11] argues that when individuals view environments similar to ones previously seen, they tend to follow similar scanpaths. The scanpaths are sequences of saccades (rapid eye movements) and fixations (periods when the eyes are relatively stationary and focused on a specific point). Such scanpaths have low entropy because they follow a regular pattern as the eyes move between a set of fixation points. The setup of our experiment was such that through making changes, participants were attempting to transition closer to a higher level of PI. The results show that from the neutral level (4) on the subjective scale (the maximum of the quadratic occurs at 3.75 in Figure 6) greater PI was associated on the average with lower eye scanpath entropy. Jordan and Slater [13] argued on the basis of the earlier work in [11, 12] that this is because visual perception is largely top-down driven, i.e., that from minimal bottom-up cues in VR, participants rapidly infer the type of environment that they are in and then operate largely out of their own internal models based on prior perception of environments with similar features. Hence ‘being there’ in the VR is like being there in reality – minimal cues provide evidence about the type of environment, and then vision is driven by internal models.

In earlier work [6] ‘presence’ was decomposed into two orthogonal elements: PI the illusion of ‘being there’ (a perceptual illusion) and Plausibility (Psi), the illusion that the events that are happening there are really happening (a cognitive illusion), in both cases even though the participant knows for sure that these are not true. The entropy finding raises the possibility of an interesting additional dimension in the meaning of presence. There could also be an important non-conscious cognitive element that involves recognition of the type of place displayed and of having been in that

type of place before. In the case of the environment of this experiment, most people would have been before in a large hall with pillars, and also would have previously witnessed a musical band playing live in front of them. Putting these two elements together also creates a recognizable environment. A prediction that follows from this that can be experimentally tested is that if the environment were completely bizarre, a type of structure that no one could have ever experienced in reality, although SC would still lead to the illusion of being there, recognition would not be possible, and therefore scanpath entropy would not decrease over time and would be unrelated to subjective assessment of PI. To some extent this was shown in one of the control conditions of the experiment of Jordan and Slater [13] which rendered the same environment (the participant on top of the pillar) but where all the triangles that formed the scene were randomly rotated. The recognition component would be most important when VR is being used for simulations of situations that could happen in reality, for example for training purposes. In such cases we do not only want the illusion of being there, but also the sense that ‘I’ve been in this type of place before’.

The 3M method was introduced in [9] in order to examine the impact of four factors on PI or Psi. In each transition they could change the illumination model or the field-of-view or the display type (first or third person perspective) or their virtual body representation. The method was also used to show that PI and Psi relied on different configurations. Azevedo, et al. [21] examined the impact of 4 factors (vision, hearing, haptics and olfaction) on PI and Psi in the context of virtual environments that depicted outdoor scenes. They also combined the method with the use of EEG monitoring. Bergström, et al. [14] used the method to examine the impact on Psi of the same string quartet used in the current experiment of the factors gaze (whether or not the musicians sometimes looked towards the participant), sound spatialization (Mono, Stereo, Spatial), auralization (no sound reflections, incorrect or correct sound reflections in relation to the virtual room), and environment (no outside sounds, or sound that corresponded to the setting). Skarbez, et al. [22], in the context of participant interaction with virtual human characters, also examined the impact on a variation of Psi referred to as ‘coherence’ of the behaviors of virtual human characters and other objects, as well as that of their own virtual body. Debarba, et al. [23] examined the impact on Psi in relation to a virtual human and of a self-avatar of various aspects of their animation - the face, hands, upper body and lower body of the character.

The method can be used for any qualitative response to a VR experience not just presence. For example, Murcia-López, et al. [24] required participants to change factors in the configuration to optimize the quality of their experience – where they observed a virtual human giving a presentation and could manipulate the eye gaze, blinking, mouth animation, and micro expressions of the virtual speaker. Gao, et al. [25] examined the impact on the level of believability of a virtual rock climbing environment of the visual appearance of the rocks, the overall visual scene, the environmental sound and dynamic behavioral factors. Fribourg, et al. [26] used the method to discover preferences for embodiment in a self-avatar of type of its appearance, control over it and perspective position. Gonçalves, et al. [27] required participants to optimize the feeling of being in a virtual room that was a replica of a corresponding real room that they had experienced, where they could switch amongst illumination rendering methods (global illumination, ambient occlusion, screen space reflections and direct shadows) finding that global illumination was the most effective. In an Extended Reality environment Lim and Ji [28] considered presence as the illusion of being co-located with surrounding objects and the impact of four factors on this (force feedback, occlusion, lighting, and material properties).

Combining 3M with Reinforcement Learning, leading to what we have called A3M, was first introduced by Llobera, et al. [10]. An advantage of A3M is that the process of participant choice selection

amongst configurations is automated. They do not have to remember which possible changes might be made, since at each change-point the specific alternatives offered can be demonstrated and their impacts can be experienced before the choice is made. This is more important the greater the number of factors. However, the disadvantage is that a large number of trials might be needed for convergence, though this has not occurred in our applications to date, albeit with a small number of possible factor changes. Although 3M results in a Markov Chain which is a probabilistic model of participant choices, the A3M method produces a further probabilistic model which is the *policy* associated with the RL. This is the set of probabilities associated with which transitions will be offered in the context of each configuration that maximize the long-term reward to the RL. Although we have not made use of this in the current paper, it may be useful in other contexts.

The 3M method has some similarities with conjoint analysis that has been used for half a century in market research [29, 30]. It addresses the same problem as 3M – assessing how a number of factors influence an outcome such as consumer preference for a product. It is a survey method whereby respondents are presented with a set of alternatives, where each alternative is a combination (configuration) of various attributes (factors) with varying levels. Survey respondents rate these configurations, and conjoint analysis then decodes their preferences, uncovering the relative importance of each attribute. Agarwal, et al. [31] refer to it as “one of the most celebrated research tools in marketing and consumer research” and it has been applied thousands of times in different applications. Conjoint analysis has also been combined with eye tracking [32] where in addition to subjective consumer choices the amount of time that a consumer looks towards a particular item can also be exploited. Conjoint analysis has even been used in VR for analysis of design alternatives [33]. However, its more recent uses have been also for preference analysis where VR is used as a presentation mode rather than to understand issues relating to VR in itself – for example, consumers preferences for product packaging [34]. A major difference between conjoint analysis and 3M is that 3M involves real-time manipulation of the factors that then lead to transitions across the configurations, and in this sense, it is bottom-up, whereas conjoint analysis is top-down since it presents consumers with different configurations and then attempts to separate out the relative importance of the constituent factors. However, in future uses of 3M the statistical methods that have been developed over the past 50 years for conjoint analysis should be taken into account.

Presence in VR has been studied since the 1990s originating from the notion of telepresence with respect to teleoperator systems where people had the sense of being at the location of the remote robot that they were manipulating, and was originally discussed in a famous paper by Minsky [35]. Ideas from telepresence were introduced into the VR field, early influential papers being [1, 36-39]. This led to a huge amount of research covering the definition of presence, its measurement and to some extent theories about how it occurs and the idea that presence leads to realistic behavior [40]. Explanations have mainly focused on properties of the display and interactive systems, essentially a list of factors that need to be addressed, for example: field of view [41], latency [8], visual display [42-44], framerate [45], and many others. See [46] for a comprehensive and systematic review. Instead of a list our theory provides an organization by considering the extent to which each factor contributes towards the VR participant experiencing PI. It is important to note that by PI we strictly mean the illusion of ‘being there’ even though the participant knows for sure that they are not ‘there’ (and this feeling is itself part of the *quale*). PI does not include issues such as attention, interest, involvement, enjoyment, or any the other possible responses to experiencing and interacting in a virtual environment [3, 39]. For example, an individual can be in a real environment, yet be bored, not pay attention to the events around them, but in terms of SC everything that they perceive is in the environment around them, and unless they lose (or enter an

altered state of) consciousness through drugs or hypnosis, they will not have any doubt as to the environment that they are in. The theory postulates that to the extent that this is the same in VR, so the greater the likelihood that presence will be engendered.

Examination of eye scanpath entropy is not only of interest from the point of view of measurement, but its ramifications have also led to an interesting new component of presence alongside PI and Psi, which is that of Recognition: to what extent does the environment lead to (probably nonconscious) recall of other similar environments that participants have experienced? This is orthogonal to PI in the sense that based only on SC participants may experience PI without Recognition, and they may have the sense of Recognition in an environment where there is no PI. For example, viewing an environment without head-tracking may still spark Recognition, but PI will fail as soon as the participants move their head. As mentioned, we therefore put forward the hypothesis that immersion in an extremely unusual, otherworldly environment, with no familiar points of reference (for example, no floor or walls) is likely to lead to PI if SC are sufficiently supported, but are unlikely to lead to entropy changes or negative correlation of entropy with PI.

4 CONCLUSIONS

There are three main conclusions to the work described in this paper. First, we have provided empirical evidence in support of the theory that the PI component of presence is influenced by the extent to which the VR system supports sensorimotor contingencies that match those of perception in physical reality. Second, we have provided a replication of the result in [13] that eye scanpath entropy is inversely correlated with PI. However, more evidence is needed to explore our finding that the relationship between scanpath entropy and PI is curvilinear, with low levels of PI also associated with low entropy. Finally, we have introduced a potentially new component of presence. Alongside the illusions of ‘being there’ (PI) and that events are really occurring (Psi), there is the sense of Recognition (‘I have been in this type of place before’). We predict that if PI is strong but Recognition is absent then there will not a relationship between PI and scanpath entropy. Our ultimate prediction is that when PI, Psi and Recognition are all strong people will behave realistically in VR: they are there, these events are happening, I’ve been in this type of place before, therefore I respond accordingly. It would be interesting to study how behavior varies with independent variation of these three presence components.

ACKNOWLEDGMENTS

We thank the company Pico for providing the Pico Neo 3 head-mounted displays with the eye tracking capability. This research was supported by the European Research Council Advanced Grant MoTIVE (#742989), and the GuestXR project which has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 101017884. ES is partially supported by a PhD studentship from Facebook.

DATA ACCESSIBILITY

All data and the statistical analysis programs are available on the Kaggle system at these addresses:

For the analysis of transitions:

<https://www.kaggle.com/code/melslater/transitions-analysis>

For the analysis of the scanpath data:

<https://www.kaggle.com/code/melslater/scanpath>

The programs can be executed by entering edit mode.

SUPPLEMENTARY MATERIAL

Supplementary Video S1: An illustration of the scenario from the viewpoint of a participant.

REFERENCES

- [1] T. B. Sheridan, "Musings on Telepresence and Virtual Presence," *Presence: Teleoperators and Virtual Environments*, vol. 1, pp. 120-126, 1992, doi: 10.1162/pres.1992.1.1.120.
- [2] J. Lessiter, J. Freeman, E. Keogh, and J. Davidoff, "A cross-media presence questionnaire: The ITC-Sense of Presence Inventory," *Presence: Teleoperators and Virtual Environments*, vol. 10, pp. 282-298, 2001, doi: 10.1162/105474601300343612.
- [3] B. G. Witmer and M. J. Singer, "Measuring presence in virtual environments: a presence questionnaire," *Presence: Teleoperators and Virtual Environments*, vol. 7, pp. 225-240, 1998, doi: 10.1162/105474698565686.
- [4] S. Grassini and K. Laumann, "Questionnaire measures and physiological correlates of presence: A systematic review," *Frontiers in psychology*, vol. 11, p. 349, 2020, doi: 10.3389/fpsyg.2020.00349.
- [5] M. Meehan, B. Insko, M. Whitton, and F. P. Brooks Jr, "Physiological measures of presence in stressful virtual environments," *Acm transactions on graphics (tog)*, vol. 21, no. 3, pp. 645-652, 2002, doi: 10.1145/566654.566630.
- [6] M. Slater, "Place Illusion and Plausibility can lead to realistic behaviour in immersive virtual environments," *Philos Trans R Soc Lond*, vol. 364, pp. 3549-3557, 2009, doi: 10.3389/frobt.2016.00074.
- [7] J. K. O'Regan and A. Noë, "A sensorimotor account of vision and visual consciousness," *Behav Brain Sci*, vol. 24, pp. 939-1031, 2001, doi: 10.1017/S0140525X01000115.
- [8] M. Meehan, S. Razaque, M. C. Whitton, and F. P. Brooks Jr, "Effect of latency on presence in stressful virtual environments," *Virtual Reality, 2003. Proceedings. IEEE*, pp. 141-148, 2003, doi: 10.1109/VR.2003.1191132.
- [9] M. Slater, B. Spanlang, and D. Corominas, "Simulating virtual environments within virtual environments as the basis for a psychophysics of presence," *ACM Transactions on Graphics*, vol. 29, p. Paper: 92, 2010, doi: 10.1145/1833349.1778829.
- [10] J. Llobera, A. Beacco, R. Oliva, G. Şenel, D. Banakou, and M. Slater, "Evaluating participant responses to a virtual reality experience using reinforcement learning," *R. Soc. Open Sci.*, vol. 8, p. 210537, 2021, doi: 10.1098/rsos.210537.
- [11] L. Stark and S. R. Ellis, "Scanpaths revisited: Cognitive models, direct active looking," in *Eye movements: Cognition and visual perception*, D. F. Fisher, R. A. Monty, and J. W. Senders Eds.: Routledge Taylor & Francis Group, 1981, pp. 193-226.
- [12] L. W. Stark, "How Virtual Reality Works! The Illusions of Vision in Real and Virtual Environments," *Proc SPIE: Symposium on Electronic Imaging: Science and Technology*, pp. 5-10, 1995, doi: 10.1117/12.207546.
- [13] J. Jordan and M. Slater, "An Analysis of Eye Scan Path Entropy in a Progressively Forming Virtual Environment," *Presence: Teleoper. Virtual Environ.*, vol. 18, pp. 185-199, 2009, doi: 10.1162/pres.18.3.185.
- [14] I. Bergström, S. Azevedo, P. Papiotis, N. Saldanha, and M. Slater, "The plausibility of a string quartet performance in virtual reality," *IEEE transactions on visualization and computer graphics*, vol. 23, no. 4, pp. 1352-1359, 2017, doi: 10.1109/TVCG.2017.2657138.
- [15] R. S. Sutton and A. G. Barto, "Reinforcement Learning: An Introduction," 1998.
- [16] B. Carpenter *et al.*, "Stan: A probabilistic programming language," *Journal of statistical software*, vol. 76, no. 1, 2017, doi: 10.18637/jss.v076.i01.
- [17] N. P. Lemoine, "Moving beyond noninformative priors: why and how to choose weakly informative priors in Bayesian analyses," *Oikos*, vol. 128, no. 7, pp. 912-928, 2019, doi: 10.1111/oik.05985.
- [18] M. Slater and A. Steed, "A Virtual Presence Counter," *Presence: Teleoperators and Virtual Environments*, vol. 9, no. 5, pp. 413-434, 2000, doi: 10.1162/105474600566925.
- [19] B. Liebold, M. Brill, D. Pietschmann, F. Schwab, and P. Ohler, "Continuous measurement of breaks in presence: Psychophysiology and orienting responses," *Media Psychology*, vol. 20, no. 3, pp. 477-501, 2017, doi: 10.1080/15213269.2016.1206829.
- [20] M. Garau, D. Friedman, H. R. Widenfeld, A. Antley, A. Brogni, and M. Slater, "Temporal and Spatial Variations in Presence: Qualitative Analysis of Interviews from an Experiment on Breaks in Presence," *Presence: Teleoperators and Virtual Environments*, vol. 17, pp. 293-309, 2008, doi: 10.1162/pres.17.3.293.
- [21] A. S. Azevedo, J. Jorge, and P. Campos, "Combining EEG data with place and plausibility responses as an approach to measuring presence in outdoor virtual environments," *PRESENCE: Teleoperators and Virtual Environments*, vol. 23, no. 4, pp. 354-368, 2014, doi: doi.org/10.1162/PRES_a_00205.
- [22] R. Skarbez, S. Neyret, F. P. Brooks, M. Slater, and M. C. Whitton, "A psychophysical experiment regarding components of the plausibility illusion," *IEEE transactions on visualization and computer graphics*, vol. 23, no. 4, pp. 1369-1378, 2017, doi: 10.1109/TVCG.2017.2657158.
- [23] H. G. Debarba, S. Chagué, and C. Charbonnier, "On the plausibility of virtual body animation features in virtual reality," *IEEE Transactions on Visualization and Computer Graphics*, vol. 28, no. 4, pp. 1880-1893, 2020, doi: 10.1109/TVCG.2020.3025175.
- [24] M. Murcia-López, T. Collingwoode-Williams, W. Steptoe, R. Schwartz, T. J. Loving, and M. Slater, "Evaluating Virtual Reality Experiences Through Participant Choices," *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 747-755, 2020, doi: 10.1109/VR46266.2020.00098.
- [25] B. Gao, J.-I. Kim, and H. Kim, "Sensory and Perceptual Consistency for Believable Response in Action Feedback Loop," *Proceedings of Computer Graphics International 2018*, pp. 201-210, 2018, doi: 10.1145/3208159.3208171.
- [26] R. Fribourg, F. Argelaguet, A. Lécuyer, and L. Hoyet, "Avatar and Sense of Embodiment: Studying the Relative Preference Between Appearance, Control and Point of View," *IEEE Transactions on Visualization and Computer Graphics*, vol. 26, no. 5, pp. 2062-2072, 2020, doi: 10.1109/TVCG.2020.2973077.
- [27] G. Gonçalves, M. Melo, P. Monteiro, H. Coelho, and M. Bessa, "The role of different light settings on the perception of realism in virtual replicas in immersive Virtual Reality," *Computers & Graphics*, vol. 117, pp. 172-182, 2023, doi: doi.org/10.1016/j.cag.2023.10.021.
- [28] C. Lim and Y. G. Ji, "The Effects of Physical Coherence Factors on Presence in Extended Reality (XR)," *International Journal of Human-Computer Studies*, vol. 172, p. 102994, 2023, doi: 10.1016/j.ijhcs.2022.102994.
- [29] P. E. Green and V. Srinivasan, "Conjoint analysis in consumer research: issues and outlook," *Journal of consumer research*, vol. 5, no. 2, pp. 103-123, 1978, doi: 10.1086/208721.
- [30] P. E. Green and V. R. Rao, "Conjoint measurement-for quantifying judgmental data," *Journal of Marketing research*, vol. 8, no. 3, pp. 355-363, 1971, doi: 10.1177/00222437710080031.
- [31] J. Agarwal, W. S. DeSarbo, N. K. Malhotra, and V. R. Rao, "An interdisciplinary review of research in conjoint analysis: Recent developments and directions for future research," *Customer Needs and Solutions*, vol. 2, pp. 19-40, 2015, doi: 10.1007/s40547-014-0029-5.
- [32] M. Meißner and R. Decker, "Eye-tracking information processing in choice-based conjoint analysis," *International Journal of Market Research*, vol. 52, no. 5, pp. 593-612, 2010, doi: 10.2501/S147078531020151X.
- [33] J. Dijkstra and H. J. Timmermans, "Conjoint analysis and virtual reality: A review," in *4th Design & Decision Support Systems (DDSS 1998) in Architecture and Urban Planning Conference, July 26-29, 1998, Maastricht, The Netherlands*, 1998: Eindhoven University of Technology.
- [34] G. Branca, R. Resciniti, and S. M. C. Loureiro, "Virtual is so real! Consumers' evaluation of product packaging in virtual reality," *Psychology & Marketing*, vol. 40, no. 3, pp. 596-609, 2023, doi: 10.1002/mar.21743.
- [35] M. Minsky, "Telepresence," *Omni*, vol. <https://philpapers.org/rec/MINT>, no. <https://philpapers.org/rec/MINT>, pp. 45-52, 1980. [Online]. Available: <https://philpapers.org/rec/MINT>.

- [36] T. B. Sheridan, "Further musings on the psychophysics of presence," *Presence: Teleoperators and Virtual Environments*, vol. 5, pp. 241-246, 1996, doi: 10.1162/pres.1996.5.2.241.
- [37] R. M. Held and N. I. Durlach, "Telepresence," *Presence: Teleoperators and Virtual Environments*, vol. 1, pp. 109-112, 1992, doi: 10.1162/pres.1992.1.1.109.
- [38] D. Zeltzer, "Autonomy, interaction, and presence," *Presence: Teleoperators and Virtual Environments*, vol. 1, no. 1, pp. 127-132, 1992, doi: 10.1162/pres.1992.1.1.127.
- [39] T. D. Matthew Lombard, "At the Heart of It All: The Concept of Presence," *Journal of Computer-Mediated Communication*, vol. 3, no. 2, JCMC321, 1997, doi: 10.1111/j.1083-6101.1997.tb00072.x.
- [40] M. V. Sanchez-Vives and M. Slater, "From Presence to Consciousness Through Virtual Reality," *Nature Reviews Neuroscience*, vol. 6, pp. 332-339, 2005, doi: 10.1038/nrn1651.
- [41] J.-W. Lin, H. B.-L. Duh, D. E. Parker, H. Abi-Rached, and T. A. Furness, "Effects of field of view on presence, enjoyment, memory, and simulator sickness in a virtual environment," *Proceedings ieee virtual reality 2002*, pp. 164-171, 2002, doi: 10.1109/VR.2002.996519.
- [42] M. Slater, P. Khanna, J. Mortensen, and I. Yu, "Visual realism enhances realistic response in an immersive virtual environment.," *IEEE computer graphics and applications*, vol. 29, pp. 76-84, 2009, doi: 10.1109/MCG.2009.55.
- [43] I. Yu, J. Mortensen, P. Khanna, B. Spanlang, and M. Slater, "Visual realism enhances realistic response in an immersive virtual environment - Part 2," *IEEE Computer Graphics and Applications*, vol. 32, pp. 36-45, 2012, doi: 10.1109/MCG.2012.121.
- [44] C. Hendrix and W. Barfield, "Presence within Virtual Environments as a Function of Visual Display Parameters," *Presence-Teleoperators and Virtual Environments*, vol. 5, pp. 274-289, 1996, doi: 10.1162/pres.1996.5.3.274.
- [45] W. Barfield and C. Hendrix, "The Effect of Update Rate on the Sense of Presence within Virtual Environments," *Virtual Reality: The Journal of the Virtual Reality Society*, vol. 1, pp. 3-16, 1995, doi: 10.1007/BF02009709.
- [46] C. S. Oh, J. N. Bailenson, and G. F. Welch, "A systematic review of social presence: Definition, antecedents, and implications," *Frontiers in Robotics and AI*, p. 114, 2018, doi: 10.3389/frobt.2018.00114.