# Experiments in Motivating Exploratory Agents

Bobby Khaleque
Queen Mary University of London
London, UK
b.d.a.khaleque@qmul.ac.uk

Mike Cook
King's College London
London, UK
mike@possibilityspace.org

Jeremy Gow
Queen Mary University of London
London, UK
jeremy.gow@qmul.ac.uk

## ABSTRACT

Exploration is found in a variety of game genres, but there has been little research in the context of PCG. This paper investigates the potential for exploratory agents to provide feedback on how well levels support exploration, with the ultimate goal of guiding level generation. We propose several motivations which might drive exploratory behaviour and model these as metrics within an agent framework based on context steering. We present a study of how the different metrics influence exploration of six game levels. It was found that combinations of metrics lead to distinct exploratory behaviours, mostly within our expectations.

## KEYWORDS

Game AI, Agents, Exploration, Exploratory Agents, Procedural Content Generation, Level Design, Automated Level Evaluation

**ACM Reference Format:**
Bobby Khaleque, Mike Cook, and Jeremy Gow. 2024. Experiments in Motivating Exploratory Agents. In *Proceedings of the 19th International Conference on the Foundations of Digital Games (FDG 2024), May 21–24, 2024, Worcester, MA, USA.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3649921.3659850

## 1 INTRODUCTION

Exploration is a key part of the player experience in many video games e.g. Proteus, Journey, ABZU, Gone Home, or Outer Wilds. It's found in games across a wide range of genres, such as First Person Shooters, Adventure games, Open World Games, Walking simulators and Survival Games. When creating such *exploratory experiences*, designers should consider how their game environments will support exploration. There is a need to design *for* exploration.

The concept of discovery is highlighted as one of the aesthetic goals within the MDA (Mechanics, Dynamics, Aesthetics) framework for game design [7]. It describes an element of gameplay where players experience the game as uncharted territory, encouraging exploration and the unveiling of new features, areas, or story elements as they progress.

What is exploration? Meyer defines it as "behavior directed toward acquiring information about the environment." [8]. Wohlwill claims that *exploratory behaviour* involves free search of the environment in order to become familiar with its layout and features, and exploration directed at an object or feature with which one is confronted [13].

Görlitz et al [6] further go on to define three subcategories of exploratory behaviour: *affective exploration*, exploration driven by what feels best to the individual; *inspective exploration*, exploration to become more familiar/confident with the environment; *diversive exploration*, exploration in the hopes of relieving boredom. We take this distinction as a starting point for evaluating exploratory behaviour, discussed in Section 6.

Automatic evaluation of game environments in terms of how well they support exploratory behaviour is a not well studied area. One notable exception is Stahkle et al.'s PathOS framework [12] for assisting designers in level and world design. They introduced various agent profiles to motivate agent behaviour, with one profile focused on exploration. (2.4.2 describes PathOS in more detail.) This can be considered an example of an *exploratory agent*: a type of agent which traverses a level and explores it in accordance to it's features. It surveys an environment, to observe which features are available in the level, and moves in the direction towards the closest interesting target(s) or direction(s).

Cook investigated the evaluation of levels with agents using a vision-based approach. The agents had an objective function designed to achieve specific framing and player vision goals [3]. This was used in a level generation system with promising results, yet abandoned. This is another example of how exploratory agents can be used in a generative system to evaluate levels and help create more well-designed ones.

We propose several motivations which could be used to motivate exploratory agents for exploring game environments. We also introduce a framework for implementing an exploratory agent and several novel approaches to evaluating game environments. Unlike existing methodologies, which primarily focus on optimization or task-specific performance, our framework emphasizes the agent's ability to engage with virtual worlds from an exploration standpoint. This shift not only advances the field by offering a more nuanced understanding of agent behaviour for evaluating generated content, but also has practical implications for designing games that are more suitable for exploration, enhancing both player engagement and the richness of game worlds. We describe a study where we examine how our agent explores six example environments according to our aforementioned motivations (more information about these motivations is presented in Section 3).

To summarise, our research introduces a variety of motivations (operationalised via metrics) tailored to assess different aspects of exploration, offering insights into how virtual environments can support exploration. Our experimental analysis reveals distinct agent behaviours and performance generally within our expectations along with dissimilar path trajectories, aligned with our motivations given to our agent. This work not only advances our


*FDG 2024, May 21–24, 2024, Worcester, MA, USA*

ACM ISBN 979-8-4007-0955-5/24/05
https://doi.org/10.1145/3649921.3659850

understanding of exploratory agents, but also opens avenues for enhancing PCG by evaluating the potential of environments for exploration.

## 2 BACKGROUND

### 2.1 Exploration and Game Design

Spatial involvement and affective involvement from Calleja's player involvement model are quite relevant to exploration [2]. Spatial involvement is described as "players engagement with spatial qualities of virtual environment". How they might explore, navigate and internalise through an environment. This aspect of the involvement model seems to be related to Görlitz et al's definition of inspective exploration [6]. Affective involvement is described as "encompasses various forms of emotional engagement", this includes witnessing aesthetic views, which in turn may make the player motivated to see more. This would be related to the definition of affective exploration. Game design literature has provided many theories as to why a game may be considered "fun" or engaging for a player, many of which can be linked back to exploration and Görlitz et al's definitions of exploratory behaviour. Giving us an idea of what factors in exploratory experiences might affect the player experience, such as; sensational pleasure (e.g. witnessing pleasurable views), interacting with the world to learn more about it and escapism through fantastical environments.

### 2.2 Player Uncertainty

Uncertainty can be argued to come through piecing together the underlying theme(s) or narrative(s), which some exploratory experiences focus on; through investigation and observation of an unknown environment for hidden information (uncertainty from hidden game elements); or randomness (uncertainty emanating from random game elements) which are certain types of uncertainty described by Costikyan [4]. Costikyan further goes to introduce many types of uncertainty including player unpredictability, randomness, hidden information, narrative uncertainty, schedule uncertainty, and uncertainty of perception. Some types of uncertainty are relevant to exploration. For example, hidden information, which may include hiding undiscovered areas of a map within fog, or partially obscuring a view with a rock. Uncertainty of perception, "the difficulty of perceiving what's going on in the gamespace", can also be applied to exploration, partially hiding important objects in a forest where it might be barely visible or very obscured.

### 2.3 Exploration Patterns

Exploratory behavior in video games has been investigated to some degree by Si [11]. They investigated behavioural exploration patterns in games, where a study was carried out to extract behaviour patterns from 25 human participants while exploring virtual environments. The games used to investigate participant behaviour were Starcraft: Brood War and 3 game modes, a pure exploration game mode (exploring the map as fast as possible within 3 minutes), a killing game mode (hunting a certain amount of targets within 5 minutes) and a searching game mode (requiring finding an opponents' base within 4 minutes). Player think-alouds were collected and thematically analysed. Four main themes were found; strategy (what strategies people make playing the games), reasoning (how players reason about situations and options), conception (what spatial conceptions about the environments are mapped in their minds), hesitation (a reluctance to move when encountering certain situations). From these themes, 4 archetypes of behaviour patterns were identified; wanderers, seers, pathers and targeters. *Wanderers* are an archetype of players who move without a definite destination or purpose, they have no targets or awareness of map features.

*Seers* are a class of player aiming to expand their visibility span, seeking to reveal as much of the map as possible in as little time as possible. This can be considered as inspective exploration.

*Pathers* take into account terrain features to construct patterns based on prior map knowledge. Again, this may be considered inspective exploration.

*Targeters* are an objective-orientated archetype taking into account terrain features, seeking out landmarks and other identifiable objects which may serve as hints of target locations. This can also be considered a form of inspective exploration.

### 2.4 AI Agents for Exploration

The use of AI agents to assist the evaluation of generated environments is an integral part of this research project. These AI agents intend to model exploratory behaviour according to certain metrics.

*2.4.1* ***Curiosity Based Exploration.*** Pathak et al [1] presents a comprehensive study on curiosity-driven learning in artificial agents, with the main focus being on agents that operate without any external rewards. The study explores various simulated environments including games and physics simulations, examining how agents perform when driven purely by intrinsic motivation, which is operationalised through curiosity. The authors investigated the use of different feature spaces for prediction error, discovering that while random features can suffice for certain environments, learned features may offer better generalisation capabilities. The research also highlights the potential limitations of prediction-based rewards, especially in stochastic settings, and suggests further investigation into efficient handling of such environments. The techniques used by Pathak et al are different to ours – our agent is not intended to be general, in the sense it would explore many different environments using intrinsic motivation, rather we expect our agent to be given different motivations and explore in different ways, in different environments. Our agent framework is also intended to be used to evaluate environments for exploration.

*2.4.2* ***Agents to Assist Game Design.*** Stahlke et al introduce PathOS, predicting player navigation in digital games [12]. Game designers can gather data from the agents that navigate the world to improve their level and world design. The goals of the system were to reduce the burden of playtesting, accessibility for devs, ease of use for designers and generalisability. This is similar to our goals proposed with our exploratory agents.

The agents work by simulating player behaviour via motives, level geometry, agent memory, and entities are taken into account in calculating these motives as well as what kind of player profile (e.g. focused more on exploration or combat). "Each motive has a vector associated with it assigning a base favourability value to each type of game entity...".

A study was conducted with 10 participants that involved the use and evaluation of candidate applications of the system. Participants were pre-interviewed and introduced to the system, given a task to create two levels with assistance from the system.

The authors reported that in the post-task interview the impressions of the system as a design tool were quite positive, though one participant mentioned how agent behaviour was dissimilar to how a player would behave.

Nova et al introduce PathOS+ [9], an enhancement of the PathOS framework, to supplement expert evaluations with simulated player data generated by AI. This approach addresses subjectivity in expert evaluations by providing objective simulated player behaviour data, thus aiming to improve the reliability of expert assessments in games user research. The potential of PathOS+ is illustrated through its application in the analysis of gameplay, with a focus on navigation and player behaviour.

## 3 MOTIVATIONS FOR EXPLORATION

In this section, we propose different motivations, operationalised in the form of metrics (described in section 4), which can serve as reasons to explore a virtual environment. These include:

**Looking for Hidden Space**: looking for potential space behind and to the sides of an object, from the viewpoint of the player/agent or objects not seen before. For example, a large rock obscuring a building or another object. This can be considered a form of "targeter" mentioned by Si [11] and a form of inspective exploration.

**Landmarks**: looking for odd/significant objects that "stick out" in some way. For example, A house or a tent in a dense forest. This can also be considered a form of "targeter" mentioned by Si and a form of inspective exploration.

**Objects Arranged in Intentional Patterns**: such rocks in the shape of a square. This can be considered a form of inspective exploration, as objects with some designer intent in their placement could draw players to explore.

**High Points**: such as a top of a hill or mountain. This may be considered a form of inspective exploration, investigating high points for better views.

**Open Areas**: such a large open field. This can be considered a metric for a form of affective exploration where a player or agent might want to feel more "free" in open areas.

**Lighter Areas**: such as a well lit section of a dark environment. Investigating this can be considered a form of affective exploration, our agent or a player might feel drawn to lighter areas (in terms of lighting in the environment) and may even be considered a "wanderer" archetype as mentioned by Si.

## 4 AGENT METRICS

There are two different types of metrics, object-based metrics and direction-based metrics. Having this design ensures a holistic approach to guide exploration. An object-based metric alone might encourage focus solely on objects without adequately exploring the space, while a direction-based metric alone could result in emergent but shallow exploration without meaningful interaction with objects. Together, these metrics promote a balanced approach, encouraging agents to thoroughly explore their surroundings and engage with objects in the environment. In this paper, we use metric codes to simplify our figures shown in the parentheses next to each metric name.

### 4.1 Direction-Based Metrics

**Anticipation Direction** (ANTD): Checks if there is an object in a given direction. The umbra and penumbra of the object are calculated (assuming that our agent is the light source). The larger the umbra and penumbra [1], the higher the score. The umbra and penumbra, though calculations for shadows, represent the total potential space there could be behind and to the sides of the object, from the viewpoint of the agent. Having an agent know the potential space may have it react more organically without feeding it information about the space which it cannot observe with the camera. This returns a maximum value of 1 and a minimum value of 0. This metric also has an object based version, to show that direction based metrics can also be an object based metrics. [11]

**Light and Shadow** (LAS): Takes a direction and measures its light intensity. The Higher the light intensity the higher the direction is scored. The light intensity itself is the score and is between 0 and 1. The light intensity is calculated using Unity's light probes system [2] and evaluating the light's colour against black and returning how different (as a percentage) it is.

**Elevation change** (ELE): Take a direction and check if it hits a terrain vertex; if it does, then the terrain vertex's y position is taken and compared against the agent's y. If it is greater than the agent's y position then a maximum value of 1 is given, depending on how much greater the y is. The maximum is achieved when the difference is 10 units or more; every unit less than 10 is given a penalty of -0.1 until the minimum value of 0 is reached.

**Openness** (OPE): Takes a direction and measures how "open" it is by checking if there are any objects within a certain distance. If there are not, score the direction very highly (a max of 1) otherwise return a value between 0 and 1 depending on how far the object(s) are from the agent.

### 4.2 Object-Based Metrics

**Anticipation Object Detection** (ANTO): Takes an object, and checks the umbra and penumbra of the object, using our agent as a light source. It returns a maximum value of 1 and minimum value of 0. This metric was included to show that object-based metrics in this agent can also be used as direction-based metrics.

**Large Object Detection** (LOD): Takes an object and compares it against the largest object it had seen during it's run. A value between 0 and 1 is returned which represents how big (in terms of percentage) the observed object is compared to the largest one our agent had observed so far. If the object is larger than the largest object observed so far, 1 is returned and the largest object observed so far is updated to the object seen most recently. This can be considered a form of inspective exploration, as large objects can count as landmarks. This can also be considered a form of "targeter" mentioned by Si, where large objects, specifically, are searched for.

[1]https://www.astronomy.ohio-state.edu/pogge.1/Ast161/Unit2/eclipses.html
[2]https://docs.unity3d.com/Manual/LightProbes.html

**Simple Detection** (SIM): Takes an object and returns 1, regardless of anything else. This is a simple metric designed to investigate every object that our agent can observe.

**Group Detection** (GRD): Takes an object and checks if there are any other objects in a certain radius (2 units) of that object. Each object that is close adds a 0.1 to the score, to a maximum of 1.

Combinations of metrics are referred to with these codings as well e.g. ANTD and LAS would be referring to the combination of Anticipation direction and Light and Shadow metrics.

# 5 EXPLORATORY AGENT FRAMEWORK

This agent framework uses a system similar to context steering [5], in which context maps are formed for each measured direction (36 in total). A context map is a projection of the decision space of the entity onto a 1D array.

This allows for a detailed interest map capturing a thorough representation of the surroundings while not being very computationally expensive. Each direction measured exists "within view" of the agent, meaning that if the direction is within the field of view of the camera attached to the agent, it is measured.

There are a number of parameters available to set for this agent:

**Length of View**: The farthest the agent can "see" in units

**Field of View**: The angle at which the agent can "see" objects within, independent of the camera attached.

In this framework, interest maps are formed from a list of objects which are in view of the agent. It uses a camera to detect which objects are in view and only samples directions within view of its camera. The highest scoring direction is chosen to be moved in. There are three main stages to the pipeline, explained below.

**Stage 1: Selecting a subset of objects** A camera is used to survey the surrounding area of the agents. Every object which falls in the agent's camera frustum is added as an object of interest. The output from this stage is a list of objects of interest.

**Stage 2: Making Interest Judgments** The output of stage 1 is taken and an interest map, a score for each direction and an optional object associated with each direction is formed. For each direction, a direction based metric is applied to calculate the directions interest score. Also, object based metrics are applied, each object has it's direction taken and rounded to the closest direction in the direction interest map (and added to the direction interest map) before the direction score is updated.

**Stage 3: Making a Navigation Decision** Finally, the direction map of Stage 2 is used to make a navigation decision. The direction of highest interest is chosen. If there is an object which is associated with the highest scoring direction, that is chosen to be navigated towards using the navmesh and object coordinates. Objects are only associated with directions when an object based metric is being used. Otherwise, simply moving 10 steps in the highest scoring direction is chosen. If there are multiple directions that are scored as the highest, a random one is chosen. In context steering a target vector is used as the velocity or delta on the velocity, whereas our agent frameworks takes discrete steps in the "best" direction. Therefore, it cannot be considered context steering, even though it uses interest maps.

Our agent framework also contains a memory which contains all the objects it had seen and investigated during it's run. If an object had been seen before and investigated (the agent had come within 10 units of the object), it was marked as visited and is now discounted from any metric calculations. This pipeline is repeated every 2 seconds, so the agent would make a navigation decision and stick to it for 2 seconds and repeat the whole pipeline again.

## 5.1 Random Agent

The random agent was used to compare with our exploratory agent to deduce if it performed better, according to our evaluation measures detailed in the next section. The implementation of the random agent is simple. It picks a random direction, within a 135 degree angle, with a bias angle towards travelling towards the centre of the level.

The bias angle is calculated by taking the signed angle from where the agent is currently at to the centre of the level. A random direction is then chosen within a 135 degree angle and the bias angle is added to that random angle. The agent travels for 2 seconds in that direction. Every 2 seconds, the bias angle is recalculated, and a different random direction is chosen.

# 6 EVALUATION MEASURES FOR EXPLORATORY AGENTS

We would like our agents to be somewhat human-like in the sense that it would explore a level in a "good-enough" way which gives meaningful feedback to a designer or generator to improve their levels. We investigated exploratory behaviour and created two evaluation models. An inspective evaluation model, with two inspective evaluation criteria and a diversive evaluation model, with 1 novelty measure. These are discussed in the following subsections.

## 6.1 Evaluating Diversive Exploration

Diversive exploration evaluation serves as a measure of how "bored" the agent is during exploration trajectory. This is inspired by Görlitz et al's [6] definition of diversive exploration (exploring to relieve boredom). This measure can show how much novel stimuli the agent is experiencing across an episode and in turn how much novel stimuli a player, who might follow a similar exploration trajectory, may experience.

To evaluate diversive exploration, a custom novelty score was created where each type of object (e.g. trees, rocks, etc.) is given an initial novelty score of 0.1. The level is divided into a grid where each tile is 50x50 units.

When a type of object comes into view of the agent for the first time, the tile of which the agent is currently in was recorded and the novelty score associated with that tile is increased (by 0.1 when the object type is viewed for the first time, when it is considered "new"). When the object type is seen for the first time, it is marked as "seen", and a penalty is applied to the object's novelty score, taking it to 0. As time passes, the object type gains a score from 0, at a rate of 0.01 per second (to a maximum of 0.1, the initial novelty score). If the type is seen again, the object type gets the new score added to the respective grid tile in which the agent is. This way, if an type is seen in a very short amount of time after having been observed initially, the grid tile the agent is in, gains a very small amount of score compared to when the object type had been marked as "new", however, when enough time passes, the object type gets back as much score as it would have gotten when it was marked as "new". This system balances between the familiarity of seen object types and the novelty of unexplored areas. By allowing the novelty score of an type to recover over time, it ensures that the agent does not completely disregard previously seen types, which might still hold some interest after a period.

## 6.2 Evaluating Inspective Exploration

Inspective exploration evaluation serves as a measure of how much knowledge the agent has about the environment. This is inspired by Görlitz et al's definition of inspective exploration.

We derived coverage of a level for each metric from the heatmap data; we counted how many of the 50x50 regions the agent had visited with their respective metrics, and created violin plots showing the coverage across the engaging and unengaging levels. Counting the total coverage of an agent can suggest how much total knowledge the agent had about the environment, a high total coverage would suggest the agent learned a lot and was exploring in a manner suggestive of inspective exploration.

We also had an object inspection measure. This was a measure (in terms of percentage) of how many objects were seen and visited by the agent. To be counted as visited, the agent needed to come within 10 units of the object. An agent with a high total object inspection score suggests it was exploring in a manner which suggested a "want" to learn about the objects in the environment, whereas a low one implies the opposite.

# 7 EXPERIMENTS

To show our implementation of **Exploratory Agents** and how our motivations are operationalised via our metrics 4, we conducted a study with agents exploring six hand-made whitebox levels. We looked at trajectories of an agent using the above metrics as well as combinations of metrics, each with different weightings, to show possible examples of exploratory behaviour in six different environments and evaluate them with our evaluation measures.

Four of these levels were based on exploratory experiences; Journey (Level 1), Proteus (Level 2), No Man's Sky (Level 3) and Zelda: BOTW (Level 4) (particularly the starting area, The Great Plateau). Figure 1 shows top down views of these levels. The additional 2 levels were meant to be considered unengaging experiences. A dense level (level 5) consisting of many objects, all of the same type, was included. An almost empty level (level 6), consisting of few objects, all of the same type, was also included. All levels were the same size (350x350) units. Let's Plays of the first four of these levels were viewed and these were decided as "ideal" or exemplary versions of exploratory experiences. These levels were chosen to be modeled because we thought that they represented environments that supported exploratory behaviour very well. They consisted of wide open spaces which could be explored in a multitude of ways, all of which can be considered valid. The unengaging levels were designed due to the future plan to use these agents to evaluate generated content, so including levels which were considered unengaging was useful.

It is also worth noting that all colliders (except for the terrain) on all objects were turned off in all levels. This is to make sure the agent did not get stuck on any objects while exploring. We generally

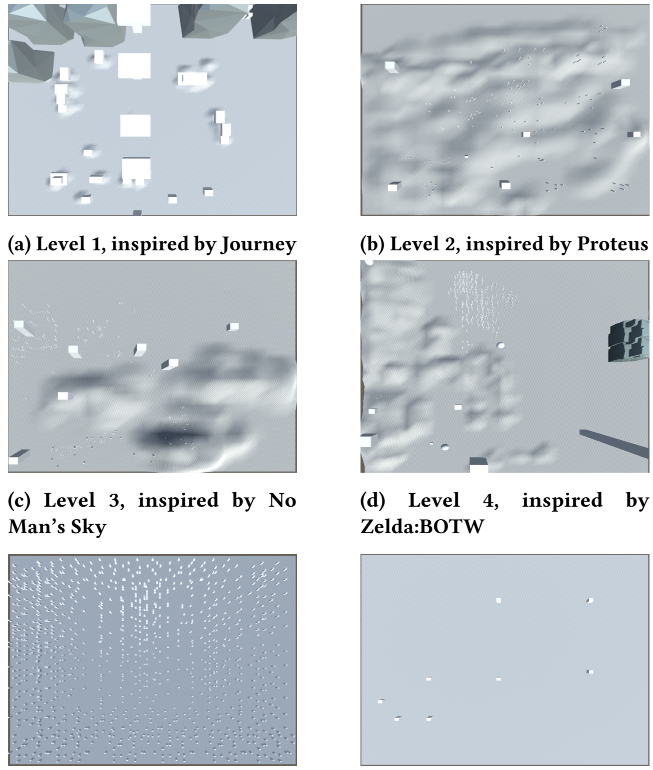

**(a) Level 1, inspired by Journey**

**(b) Level 2, inspired by Proteus**

**(c) Level 3, inspired by No Man's Sky**

**(d) Level 4, inspired by Zelda:BOTW**

**(e) Level 5**

**(f) Level 6**

**Figure 1: Top-down views of the levels tested**

expect our coverage and object inspection scores for our levels to be around 33–66% for our exemplary levels and less than 33% for our unengaging levels. This range acknowledges the complexity and unpredictability of dynamic environments, where achieving 100% inspection and coverage is unrealistic and not necessarily what a player might target. This range is good for environments where complete coverage and inspection is less critical than strategic discovery of key areas and environmental characteristics. An agent episode consisted of surveying the level for three minutes.

For all of the metrics tested, their position throughout their spawn was recorded to get path data, as well as how much time they spent in each 50x50 region of each level. A novelty score was measured, where each 50x50 region of the level was measured, in terms of a custom novelty score, as a form of diversive exploration evaluation. 2 inspective measures were also included, listed in the previous subsection. An agent episode for the random agent was simply to run it for 3 minutes, diversive and inspective evaluations for this agent were also measured. Once the agent episode is finished,the agent is respawned for another episode at the initial spawn point. The initial spawn point for each level was chosen randomly and kept the same throughout all agent episodes for each level. This process was repeated three times with different initial spawn points (chosen randomly and inspected to make sure that they were significantly different from the previous spawns). The limited length and field of view does mean the spawn point will likely greatly affect the agent paths; to obtain a broader sample, we tested 3 different spawn points on 3 levels. All individual metrics for the agent were tested. All pairs for each metric for the agent were also tested. We expect our novelty heat maps for our exemplary levels to have regions where novelty is significantly higher than the other regions of which the agent explores and some other regions aren't as high as the average. We can consider this as drama showing the agent experiences a variety of object types throughout it's trajectory. For the unengaging levels we expect the novelty heat maps to have more uniform values, as these levels consist of only one type of object, representing a non-dramatic path that may be considered unengaging.

## 7.1 Path Similarity

We also looked at how similar the behaviours were for our metrics, within the context of each level, using Earth Movers Distance (EMD) [10] to measure the similarity between the agents' paths. For each agent path, we took the time spent in each 50x50 region of the level, and compared that distribution to those of other paths for the the same level and spawn point. The EMD between two distributions $P$ and $Q$ is defined as the minimum amount of work to transform one into another, 'moving' time from one region to another. If the amount moved between two regions $i$ and $j$ is the flow $f_{i,j}$, and $d_{i,j}$ is the ground distance between those regions, the EMD for a distribution $n*m$ is defined as $EMD(P,Q) = \frac{Work(P,Q)}{Flow(P,Q)}$ and where $Work(P,Q) = \sum_{i=1}^{m}\sum_{j=1}^{n} f_{i,j}d_{i,j}$ and $Flow(P,Q) = \sum_{i=1}^{m}\sum_{j=1}^{n} f_{i,j}$. Using this measure of path similarity, we performed agglomerative hierarchical clustering along with nearest-point linkage of the metrics, for each individual map and spawn point.

## 7.2 Experiment Parameters

For the experiment conducted the following parameters were set to the values described below for the agent:

**Length of View**: 80. We thought 80 would be an appropriate distance distance the agent could observe objects as it followed its path, where it saw about 23% of the level's total length.

**Field of View**: 90. We thought this was an appropriate value for the Unity camera. A 90-degree FOV provides a good angle of vision, allowing the agent to perceive a significant portion of the environment at any given time. This mimics a natural field of view similar to human peripheral vision, making it a good balance between seeing enough of the environment without excessive distortion. Also, many first-person games use a FOV around 90 degrees as it provides a good balance for player experience.

# 8 RESULTS AND DISCUSSION

## 8.1 Inspective Exploration

To perform an evaluation on the inspectiveness of the agent, two measures were used: coverage and object inspection. The findings for both measures were quite different, particularly for each of their metric and metric combinations. However, they can be considered to perform well, from an inspective standpoint, mostly in line with our expectations as shown in the results[3].

[3]https://github.com/BKhaleque/Experiments-in-Motivating-Exploratory-Agents/tree/main

*8.1.1* ***Coverage***. The data points, two of which are shown in Figure 2, for individual spawns show variability around the average coverage. Some points fall below the lower boundary (33%), which may indicate an unfavourable spawn point. This is apparent on every level. Random is well below the lower boundary in all spawn points, as expected, and every metric on every level has much greater coverage.

For our exemplary levels, the target coverage was set between 33–66%. Across all 3 spawns, almost all metrics within this target range on average, indicating successful exploration strategies for those levels.

Metrics which come below our expected coverage value (on average) in our exemplary levels include; OPE, LAS, ANTD LAS, LAS ELE and ANTD. This may have been due to an unfavourable spawn point for these levels or due to the level features not supporting these metric's exploration behaviours, particularly in combinations. For example, in level 2 having ELE and OPE might not have been conducive to high coverage as a lot of objects are situated on high terrain points, which is what ELE is drawn to investigate, whereas avoiding objects is what OPE is drawn to investigate.

The expectation was that levels 5 and 6, being considered unengaging, would see coverage under 33%. However, the data does not entirely support this, with most metrics exceeding this threshold, indicating a higher level of engagement than anticipated. This suggests that while these levels might have been designed or considered to lack engagement, agents still explored these environments to a significant extent. High variability around the average is observed in all spawns.

However, the data from the three spawns show that most metrics fluctuate around or below 33%, with some dipping well below it. (e.g. ELE OPE, ELE, OPE and GRD ANTD in level 5 and OPE, GRD OPE, GRD LAS OPE, ELE, GRD in level 6) This outcome is in line with the expectation that the unengaging levels would generally be less conducive to high coverage, likely due to fewer interactive elements in level 6 or overly complicated navigation in level 5.

*8.1.2* ***Object Inspection***. On average, the object inspection percentages for our exemplary levels are generally higher than those for the unengaging levels, however many metrics come below our expected values for our exemplary levels, this is demonstrated in Figure 2.

For our exemplary levels, the target inspection was set between 33-66%. Across all 3 spawns, around half the metrics come within this range, on average. Object inspection is, on average, lower than coverage in our exemplary levels.

Metrics which come below our expected inspection value (on average) in our exemplary levels include; SIM GRD, LOD SIM, GRD ANTD, ANTD OPE, ANTD LAS, ANTD, ANTD ELE, LAS, LAS OPE, ELE, ELE OPE, LAS ELE and OPE. This, again, may have been due to an unfavourable spawn point for these levels or due to the level features not supporting these metric's exploration behaviours. However, it is expected at metrics like OPE and those paired with OPE will have lower inspection values, due to OPE making the agent prefer to explore areas with fewer objects in them. It is worth noting that some spawns in some levels with some spawn points do go above our expected boundary (e.g. GRD ANTD and LOD ANTD have one spawn above our expected upper value of 66%),

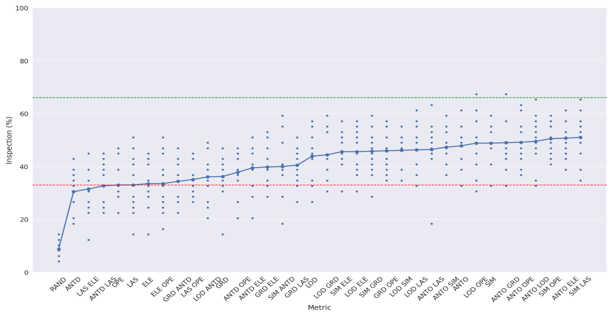


**(a) Coverage, Engaging Levels**

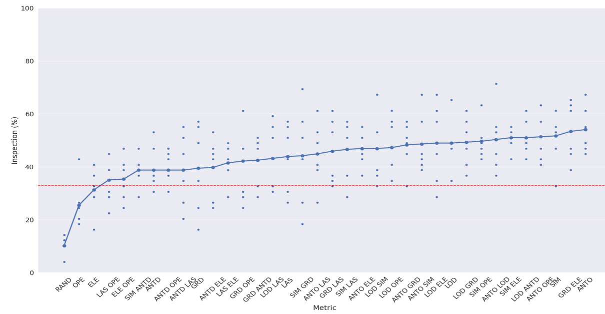


**(b) Coverage, Unengaging Levels**

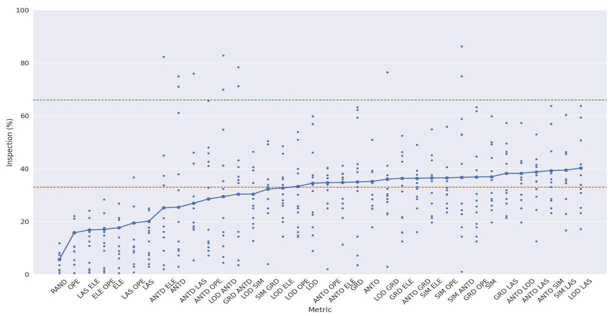


**(c) Inspection, Engaging Levels.**

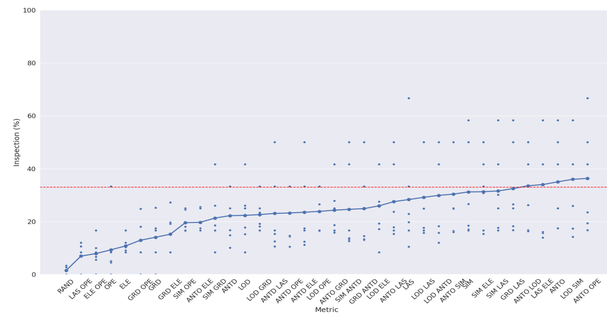


**(d) Inspection, Unengaging Levels.**

**Figure 2: Level coverage and inspection for various metric combinations, ordered by average value. Each point represents an individual trajectory from a spawn point within a level. Blue plot lines show average value across all levels/spawns. Green/red dashed lines show our expected upper and lower bounds for exploratory agents.**

indicating a very favourable spawn point for these instances. The random agent, on average and in almost every spawn, had much lower object inspection scores than any of the other metrics, which was expected

On average, object inspection scores come much higher and show much more variability in our exemplary levels than in the unengaging levels. It should be noted that some metrics exceed our expected value of below 33%. These include; GRD LAS, ANTO LOD, LAS ELE, ANTO, LOD SIM, ANTO OPE. These could indicate favourable spawn points for each of these metrics at all levels. This is further indicated by some of the data points being far below 33% in our plot.

In all our exemplary levels the random agent has much lower inspection scores than our agent, which is what we expected.

## 8.2 Diversive Exploration

Novelty values throughout their respective paths are fairly uniform throughout, with one notable hotspot (the point at which they spawn, where everything in view is considered "new"). The novelty remains fairly consistent throughout each metric path with various high and low points. This is consistent for all our exemplary levels across all spawnpoints. Not every section in each level is as novel as each other; however, there are notable novelty hotspots (apart from the spawn point). This suggests that many parts of some levels and the paths taken were more novel than others, as we expected. Regions in the unengaging levels showed a lot more uniform novelty scores, which is especially true for level 6, where the level was mostly empty. There are spikes in novelty score in certain regions (especially for level 5) though they are less pronounced than in our exemplary levels. This is consistent throughout all runs. Our agent, in all levels, compared to the random agent, show far more varied novelty scores; the main novelty hotspot for the random agent is the middle of each level, as expected.

*8.2.1* ***Similarity of paths.*** The dendrograms, shown in figure 3, formed of the mean EMDs of time spent in each region for each metric of all three spawns, show a diverse range of exploration patterns. Most metrics in our exemplary levels appear to be more similar to one another compared to our unengaging levels. This suggests that our exemplary levels were more conducive to exploration, at least for our tested metrics, than our unengaging levels, where metrics were exploring in a less directed way and dissimilarities between metrics were higher.

In terms of path variations for all tested levels, there is a large amount, all seemingly focused on their respective metrics, e.g. the group detector visits groups and pairs, anticipation detector visits objects with lots of space behind and to the sides, and large object detector visits large objects. Most singular metrics on most levels show the expected behaviour and all metrics tested show distinct behaviour.

## 8.3 Overall Findings

Overall, coverage, object inspection, novelty and path similarity measures, for all spawns, generally, come within our expectations,

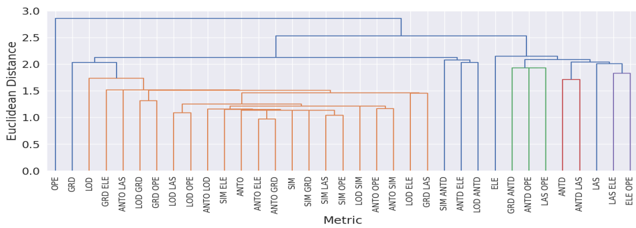


**(a) Engaging levels**

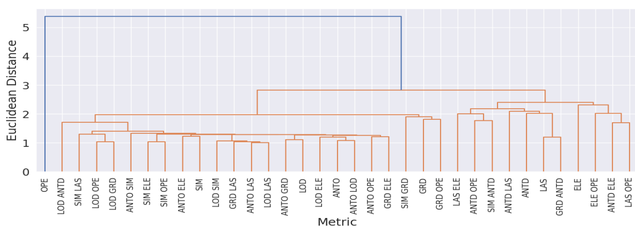


**(b) Unengaging Levels**

**Figure 3: Dendrograms Showing Clustering of Metrics Based on Path Similarity**

for our exemplary levels and our unengaging levels. There is also evidence here to suggest that our agent provides performs much more competently than our implementation of a random agent and may be used to evaluate levels in a generative system.

Compared to other implementations of exploratory agents, PathOS [12] and PathOS + [9], in particular, is that our agent framework is more focused on evaluating environments in terms of how well they support exploration in order to help developers and procedural content generators create environments that are more suited to support exploration. Motivations for exploring a 3D environment are not as thoroughly investigated compared to our work. PathOS/PathOS+ provide a more generalised agent framework for autonomous play testing more focused towards aiding developers in finding bugs and potentially predicting player behaviour.

## 9 FUTURE WORK

The agent demonstrated in this study provides competent inspective and diversive exploration according to our measures. However, there are room for improvements. Some improvements to the current metrics and more metrics shall be investigated; Perhaps different field of views and length of views for the agent might be tested to see how it performs. Investigating more/different metrics is a high priority, including:

**Landmark detection** - detecting odd/significant objects that might "stick out" in some way. This models where players who are interested in significant locations or objects might go, this may be considered inspective exploration.

**Simple pattern detection** - detection of patterns in placements of objects, such as if they were placed in a circular or square pattern, to see if there was any intention in their placement and how they could fit into the environment. This models some type of inspective exploration that players who are interested in finding out more about the story of the world or the designer intention.

**Specific object type detection** - finding objects of a certain type e.g., trees only, rocks only, or apple trees only. Some players/agents may be interested in looking for specific types of object. Perhaps this may be considered a type of affective exploration, a player might feel drawn towards specific object types.

Because the agent is going to be used for evaluating generated content, testing on a wider variety of levels, including more of which are not exemplary, may tell us more about how the agent would evaluate "bad" levels. However, our results suggest that with our current metrics and evaluation criteria, our unengaging levels were evaluated as "worse" than our exemplary levels. This suggests that the agents will be useful in evaluating generated content.

## 10 CONCLUSION

In this paper we investigated exploratory behaviour and an implementation of exploratory agents. A study is performed that introduces how exploratory behaviour is modelled via several metrics. These metrics and some of their pairs were evaluated using our own models of inspective and diversive exploration. It was found that the metrics have their own distinct behaviours and function as intended and perform mostly within our expectations. Relevant future work was also mentioned to improve the system.

## ACKNOWLEDGMENTS

This work was supported by the EPSRC Centre for Doctoral Training in Intelligent Games & Games Intelligence (iGGi) EP/S022325/1.

## REFERENCES

[1] Yuri Burda, Harri Edwards, Deepak Pathak, Amos Storkey, Trevor Darrell, and Alexei A. Efros. 2019. Large-Scale Study of Curiosity-Driven Learning. In *Seventh International Conference on Learning Representations.*

[2] Gordon Calleja. 2011. The Player Involvement Model. In *In-Game: From Immersion to Incorporation.* The MIT Press.

[3] Michael Cook. 2021. The Road Less Travelled: Trying And Failing To Generate Walking Simulators. *CoRR* abs/2104.10789 (2021). arXiv:2104.10789 https://arxiv.org/abs/2104.10789

[4] Greg Costikyan. 2013. *Uncertainty in Games.* The MIT Press.

[5] Andrew Fray. 2015. Context steering: Behavior-driven steering at the macro scale. In *Game AI Pro 2.* CRC Press, Chapter 18, 183–193.

[6] Dietmar Görlitz, Hans Joachim Harloff, Günter Mey, and Jaan Valsiner (Eds.). 1998. *Children, Cities, and Psychological Theories.* De Gruyter, Berlin, Boston. 97–98 pages. https://doi.org/doi:10.1515/9783110885194

[7] Robin Hunicke, Marc Leblanc, and Robert Zubek. 2004. MDA: A Formal Approach to Game Design and Game Research. *AAAI Workshop - Technical Report* 1 (01 2004).

[8] Jerrold S. Meyer. 1998. Chapter 22 - Behavioral Assessment in Developmental Neurotoxicology: Approaches Involving Unconditioned Behaviors Pharmacologic Challenges in Rodents. In *Handbook of Developmental Neurotoxicology*, William Slikker and Louis W. Chang (Eds.). Academic Press, San Diego, 403–426. https://doi.org/10.1016/B978-012648860-9.50029-7

[9] Atiya Nova, Stevie Sansalone, and Pejman Mirza-Babaei. 2021. PathOS+: A New Realm in Expert Evaluation. 122–127. https://doi.org/10.1145/3450337.3483495

[10] Yossi Rubner, Carlo Tomasi, and Leonidas J. Guibas. 2000. The Earth Mover's Distance as a Metric for Image Retrieval. *International Journal of Computer Vision* 40, 2 (Nov. 2000), 99–121. https://doi.org/10.1023/A:1026543900054

[11] Chen Si. 2017. *Believable Exploration: Investigating Human Exploration Behavior to Inform the Design of Believable Agents in Video Games.* Ph. D. Dissertation. University of Technology Sydney.

[12] Samantha Stahlke, Atiya Nova, and Pejman Mirza-Babaei. 2020. Artificial Players in the Design Process: Developing an Automated Testing Tool for Game Level and World Design. In *Proceedings of the Annual Symposium on Computer-Human Interaction in Play* (Virtual Event, Canada) *(CHI PLAY '20).* Association for Computing Machinery, New York, NY, USA, 267–280. https://doi.org/10.1145/3410404.3414249

[13] Joachim F. Wohlwill. 1981. A Conceptual Analysis of Exploratory Behavior. In *Advances in Intrinsic Motivation and Aesthetics*, Hy I. Day (Ed.). Springer US, Boston, MA, 341–364. https://doi.org/10.1007/978-1-4613-3195-7_14