Bournemouth University

FACULTY OF MEDIA SCIENCE AND TECHNOLOGY

MSc Human Centred Artificial Intelligence for Games Development
May 2026

Are VR and AI effective tools for creating an unsupervised learning environment for photography?

by

Josh Eyres

Faculty of Media Science and
Technology
School of Computing & Engineering
Individual Masters Project

Abstract

Traditional training and education frameworks for technical and creative disciplines such as photography rely on cost-prohibitive hardware or supervised studio environments, creating significant accessibility barriers for novice learners. Alongside this existing virtual photography within video games provide many features but function as passive afterthoughts rather than providing pedagogical feedback or tools to guide artistic and technical skill acquisition.

To resolve this limitation, this project developed an autonomous VR learning application that combines real time camera simulations with automated multimodal AI evaluation and feedback. Built within Unity 6 using the Universal Render Pipeline (URP) render graph architecture, the application simulates the core mechanics of the exposure triangle (ISO, Shutter speed and Aperture) via a physical diegetic virtual DSLR designed to reinforce real-world muscle memory support sim-to-real transfer learning. To eliminate VR sickness and maintain a smooth framerate a custom asynchronous render pass was engineered to capture and save camera colour and depth data that is subsequently passed to a local Python backend via a flask server. This backend sequentially runs two versions of Google's Neural Image Assessment (NIMA) models (technical and aesthetic image scoring) followed by mathematically calculated metrics that are then used to anchor a Llama-Vision 3.2 VLM run locally via Ollama, to generate qualitative actionable feedback for the user.

A remote visual heuristic user evaluation involving a range of photographer from novice to expert, confirmed the systems validity with 95% of respondents verifying that the application functions as an effective unsupervised learning environment. While qualitative feedback highlighted minor local model attention drift and stylized asset biases within the aesthetic scoring model. Ultimately, the study confirmed a strong consensus for successful sim-to-real cognitive transfer, demonstrating that automated, hybrid AI-driven feedback loops can reliably cultivate real-world technical confidence and self-efficacy in beginners.

Dissertation Declaration

I agree that, should the University wish to retain it for reference purposes, a copy of my dissertation may be held by Bournemouth University normally for a period of 3 academic years. I understand that once the retention period has expired my dissertation will be destroyed.

Confidentiality

I confirm that this dissertation does not contain information of a commercial or confidential nature or include personal information other than that which would normally be in the public domain unless the relevant permissions have been obtained. Any information which identifies a particular individual's religious or political beliefs, information relating to their health, ethnicity, criminal history, or sex life has been anonymised unless permission has been granted for its publication from the person to whom it relates.

Copyright

The copyright for this dissertation remains with me.

Requests for Information

I agree that this dissertation may be made available as the result of a request for information under the Freedom of Information Act.

Signed: 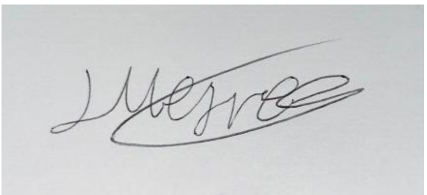

Name: Josh Eyres

Date: 22/05/2026

Programme: MSc Human Centred Artificial Intelligence for Games Development

Original Work Declaration

This dissertation and the project that it is based on are my own work, except where stated, in accordance with University regulations.

Generative AI was used for coding support during the development process and to refine readability, spelling, and grammar during the report writing process. The author reviewed and edited all output and takes full responsibility for the final content.

Signed: 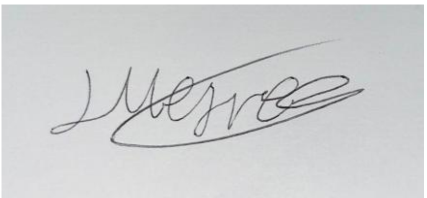

Name: Josh Eyres

Date: 22/05/2026

Acknowledgments

First of all, I would like to thank my supervisor, Fred Charles, for keeping me on track, supporting me and this project throughout the whole process and for his unwavering enthusiasm for his students as a whole. His guidance has been a massive help

I would also like to thank all of the participants who volunteered for the user study; without their help, I would not have been able to validate the technical pipeline and this project as a whole to the degree that I have.

Finally, I would like to thank my family and friends, for their support of me throughout my university career. In particular, I am deeply grateful to my father, Rob Eyres, my mother, Viki Kelley, and my close friend, George Smith, who constantly helped to motivate me, keep me on track, and acted as invaluable sounding boards for my ideas.

TABLE OF CONTENTS

Contents

LIST OF FIGURES

1 INTRODUCTION

1.1 Problem definition

Existing virtual photography and photo-mode tools rarely provide feedback that supports player creativity, helps them develops new skills, or refines prior knowledge - particularly within VR environments.

Currently virtual photography is often passive and almost an afterthought that is just put in as another feature for no real reason causing a barrier to entry for beginners who want to learn virtual or real photography techniques. Even in existing dedicated photography simulators such as “Lushfoil Photography Sim” the user gets no real feedback and while that is likely on purpose to mirror real life it does not give the player the ability to grow within the creative space that is the game. This project aims to explore exactly that by integrating AI directly to give players instant feedback and improvements.

1.2 Aims and objectives

The aim of this project is to develop an AI-powered virtual reality photography system that provides players with evaluation and feedback to help them improve the quality and composition of their photos while supporting their creativity.
Objectives:

1. Design and develop a VR environment that supports the technical requirements of the implemented photography system
2. Develop and implement an AI-driven feedback and evaluation tool within the game that is capable of assessing photos on a defined set of criteria (Composition, Lighting, etc)
3. Provide players with clear actionable feedback that supports and enhances player creativity rather than constraining it
4. Conduct user testing to evaluate the tools effectiveness, usability and impact on creative learning
5. Analyse results to determine how the feedback affects player engagement and skill development
6. Implement an effective VR photography system that mirrors a real-world camera as closely as possible.

1.3 Success criteria

Technical Success:

- Interacting with the virtual camera should mirror its real-world counterpart as accurately as possible to ensure Sim-to-Real transfer of skill
- The pipeline from photo planning – taking – to feedback should be efficient and fast
- Captured photos should be high quality and accurate to the user’s camera settings

Pedagogical Success:

- The AI should provide clear actionable feedback for the user to improve their next photo
- Do users feel more confident taking aesthetically pleasing or technically correct photos?
- The AI feedback should correctly identify technical faults (exposure, composition and focus errors) in the majority of cases

User experience Success:

- Users should be able to navigate and interact with the application without external help

- Feedback should show the AI feedback helped the user improve over time
- Feedback should show the AI scoring, and feedback motivates the user rather than demoralizing them

1.4 Research questions

1. How effective is AI feedback and evaluation for improving players creative and technical ability in photography?
2. To what extent does automated AI feedback influence user creative decision-making?
3. How translatable are virtual photography skills to the real world?

2 BACKGROUND

2.1 AI – Photo critique and Image quality assessment (IQA)

2.1.1 Image quality assessment

Computational evaluation of image quality requires a distinct separation between pixel-level quality and human perceptual preference. Traditional IQA methodologies typically rely on error sensitivity and pixel-based metrics, the most common being Mean squared error (MSE) and Peak signal-to-noise ratio (PSNR) (Wang et al., 2004). These traditional metrics fail to capture perceptual quality compared to structured approaches (Bosse et al., 2017; Hou et al., 2014).

While there is some potential for semantic image analysis with Multimodal Large Language Models (MLLMs) they have some key limitations when tasked with direct aesthetic scoring. Because MLLMs generate text autoregressively through tokenization, they inherently process numerical ratings as separate words rather than points on a continuous scale. This structural insensitivity to numeric value, combined with score quantization errors, causes MLLM-generated ratings to behave erratically or collapse into narrow, over-predictable categories (Ma et al., 2026).

To resolve the limitations both the traditional pixel-based metrics and MLLM scoring constraints this project implements a hybrid AI Pipeline for image analysis and feedback. Making use of Google’s Neural Image Assessment (NIMA) model which predicts the “distribution of human opinion scores using a convolutional neural network” (Talebi & Milanfar., 2018) to generate the “scores”. Rather than forcing a language model to guess an abstract number, the system uses NIMA to compute the numeric score, which is subsequently passed into the local Vision-Language Model's context window to drive the feedback generation loop.

2.1.2 Photo critique (LLM/VLM)

While aesthetic and technical scoring provides necessary baseline metrics, they function as “black boxes” that fail to offer actionable feedback or context to a learner. To translate these quantitative insights into pedagogical feedback current systems increasingly leverage MLLMs.

However, research emphasizes that while MLLMs excel at object recognition they struggle with subjective, abstract and human centred cognitive properties (Huang et al., 2024; Qi et al., 2025; Chan et al., 2025), such as scoring an image on its aesthetic quality. Furthermore, while there is an established demand for qualitative image assessments over rigid pixel-level metrics, current standalone vision-language methods suffer from dataset biases and fail to reliably achieve true aesthetic reasoning patterns (Liu et al., 2025).

To circumvent these operational boundaries without incurring prohibitive computational or commercial token costs, the architecture implemented an open-source Vision-Language Model, Llama 3.2 Vision, via a local pipeline. Rather than forcing the VLM to interpret composition, exposure, and aesthetics in a vacuum, the pipeline explicitly scaffolds the input context. By feeding the pre-calculated NIMA score distributions alongside deterministic, mathematically calculated metrics directly into the prompt alongside the image, the VLM is successfully anchored to a technical ground truth. This approach isolates the model's responsibility to qualitative translation and critique generation, neutralizing its scoring insensitivities to ensure accurate, stable, and pedagogy-driven feedback loops.

2.2 Game – Environments and Sim-to-real

2.2.1 Environment design and spatial dynamics

To maximize the effectiveness of a virtual training simulation, literature emphasizes a critical synergy between environment design, spatial presence and active user interaction. The level of fidelity and realism, whether defined by visual complexity, physical accuracy, or immersive depth fundamentally alter the user experience and skill transfer. High fidelity is essential to minimize perceptual errors such as depth underestimation (Jones et al., 2008), and Ragan et al. (2015) states that "Training systems for visual scanning and similar tasks should, when possible, use a level of visual complexity that is as close to the real environment as possible in order to ensure good transfer.". While Meade et al. (2019) states active navigation significantly improved spatial memory over passive observation. This requirement for active, high-fidelity exploration led to the implementation of real-time 3D environments over less interactive modalities, such as static 360-degree images or high-latency geospatial streaming APIs. To balance the demand for spatial accuracy with performance limitations, the scenes were constructed using optimized, stylized low-poly geometries.

2.2.2 Immersion, presence and embodied interaction

To facilitate meaningful skill transfer, the design of the virtual camera rejected abstract menu inputs in favour of "embodied interaction." Defined by Dourish (2001) as phenomena that occur in real time and space. Rather than relying on traditional 2d user interfaces and button presses the user must physically interact with the virtual camera to manipulate its parameters, via tangible interactions like swapping lenses, twisting a lens ring to zoom or interacting with a joystick menu. This ensures the muscle memory associated with handling a real camera is conditioned while using the virtual one. To support this the application relies on perceptual affordances, the visual appearance of an object should convey its function (Norman, 1999), for example showing a joystick on the camera body, rubber grip on the lens rings and highlights for the lens to attach to the camera.

2.3 Pedagogy

Much research has already been conducted on game-based learning as Plass et al. (2015) describe, with Piaget (1951/2013, p15) stating "the reversibility which characterises the operations of reason, while the primacy of accommodation characterises imitation and the image, and the primacy of assimilation explains play" essentially stating that people must play, explore and absorb concepts at their own pace to develop truly flexible intelligence. Alongside this VR and emerging technologies boost motivation, interest and active participation (Fu & Zhang, 2021) Subsequently improving learnability of concepts.

This architecture combined the benefits of play, exploration, and the capacity to "fail gracefully" (Plass et al., 2015) within an immersive VR environment and AI technologies, the system established an autonomous, unsupervised learning sandbox that promotes technical instruction without incurring the financial or logistical costs of physical camera equipment. Due to this setup the project satisfies a distinct technological potential explicitly highlighted by Fu and Zhang (2021). Furthermore, by integrating structured, AI-generated multimodal feedback, the system purposefully mitigates the risk of a novice learner experiencing cognitive overload when faced with open-ended experimentation. The generative pipeline acts as an evaluative guide that scaffolds the learning process, when necessary, rather than completely replacing human instruction, aligning directly with the pedagogical frameworks for human-AI creative collaboration outlined by Hutson and Lang (2023).

The ultimate validation of a virtual simulation is its capacity for “sim-to-real” skill transfer, the degree to which skills learned in the virtual space translate to a real-world scenario. While early research described that simulations needed to replicated physical likeness to provide more effective training (Hays & Singer, 1989/2012). Modern research shows “cognitive fidelity” is far more critical for technical skill acquisition (Rose et al., 2000) and that physical recreation is not necessarily a driving factor in learning. This theoretical shift allows the architecture to selectively abstract non-essential physical parameters without losing pedagogical value. For example, while the low-poly environmental assets lack photorealism, accurately replicating the core mental load of a DSLR camera provides an educationally equivalent training experience.

2.4 Virtual Production

While virtual production is a promising and important field for this type of product that fundamentally changes the creative workflow by enabling real-time visualization and interaction (Trottnow et al., 2015) within the real world, allowing for what would essentially be real-world photography with the benefits of virtuality. Research dictates that it does have significant drawbacks in its current stage: Mendez et al. (2025) state “learning curve for operating these environments nearly insurmountable for a novice user.” In their discussion on operating a professional virtual set. Secondly, both Mendez et al. (2025) and Bessa & Madureira (2025) note the high cost of entry and difficulties faced when utilising this technology in an educational context. Due to these reasons this project intentionally omitted use of virtual production within the application during the planning process, nevertheless, if a virtual production application that addressed these limitations and democratized the medium were made it serve as an invaluable tool for immersive education.

3 GAME DESIGN & DEVELOPMENT

3.1 Environments

Several environment options were considered, including photorealistic 3D, 360-degree imagery, and geospatial streaming via Cesium. Low-poly 3D environments (See figure 1) were selected to balance performance with replicating the high-fidelity spatial complexity of the real-world required for skill transfer (Ragan et al., 2015) while also enhancing immersion and overall learnability (Christopoulos et al., 2023).

Figure 1 An example Low-poly 3D environment

While 360-degree images offer performance gains and a larger range of “environments” they fundamentally restrict player agency by lacking variations such as motion parallax or dynamic depth of field essential for photo composition causing the player to essentially be taking “photos of a photo”. Furthermore, the absence of depth in a 2D projection would invalidate the planned mathematical metrics and subject detection, necessitating a vastly different approach. Alongside this, Meade et al. (2019) highlight that active navigation in 3D environments is essential for developing the spatial memory required for complex tasks.

Unlike streamed geospatial data, which lacked the functional realism required for this prototype, manual 3D world-building facilitates dynamic lighting and object manipulation. This is exemplified in the "Studio" scene (Figure 2), where players have full control over lighting and subject placement, mirroring the mechanical and creative workflows of real-world studio photography.

Figure 2 The studio environment showing plain backgrounds and many props for the user to choose from

3.2 Virtual DSLR

The core design philosophy for the virtual camera was to mimic a real-world DSLR to facilitate sim-to-real skill transfer. The system prioritizes diegetic controls, interactions existing within the 3D world, over traditional 2D overlays to build relevant muscle memory.

- **Lens interaction**: Zoom and focus are adjusted by physically rotating lens rings. This 1:1 movement provides a haptic affordance that mirrors mechanical photography. A lerped exponent is applied to the focus output, granting high-precision control at close ranges where depth-of-field is most sensitive, while allowing broader adjustments at distance.
- **Physical immersion**: Gear selection and physical movement (e.g., swapping an 18-55mm for a 70-300mm lens) forces players to engage with the technical constraints of real-world equipment, further enhancing pedagogical scaffolding.
- **Settings interaction**: Virtual touchscreens were discarded due to "ghost" presses and lack of tactile resistance. Instead, a joystick-driven physical interface was implemented on the camera body. This mimics older DSLR ergonomics, reducing cognitive load and providing "passive haptics" through direct correlation between player hand movement and in-game mechanical action.

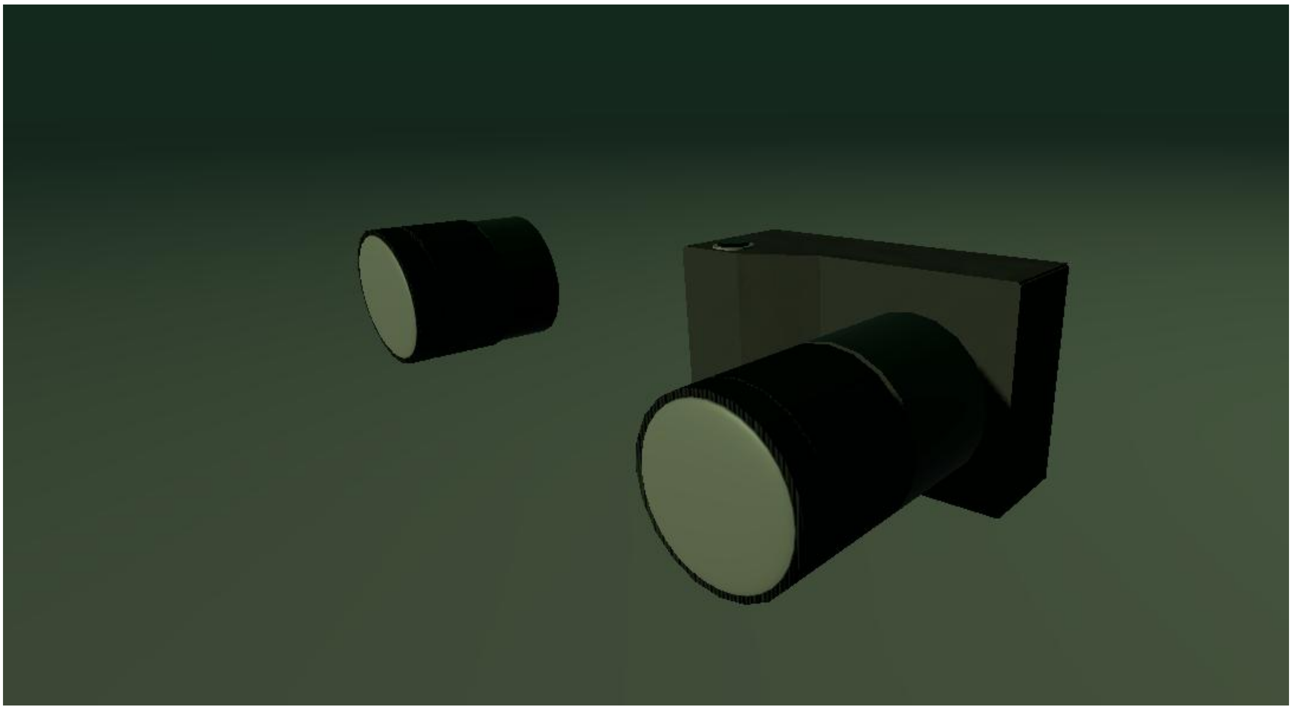

Figure 3 The in-game DSLR with different lenses made to mimic the look of a real-world DSLR

To extract high-fidelity data for AI analysis without compromising VR performance the project makes use of the unity render graph which has become the foundational architecture for the Universal render pipeline in Unity 6+. Making use of this architecture ensures the application is futureproof and leverages native engine optimizations.

Figure 4 Output of the camera saved after post processing through the render graph system

The Render Graph acts as a high-level GPU scheduler, merging passes and optimizing resource allocation. This is critical in VR to maintain consistent framerates and mitigate motion sickness during intensive capture operations. This architecture provides three primary technical advantages:

1. It ensures both colour and depth data are captured from the exact same frame, providing synchronized, consistent data for the AI analysis pipeline
2. Using render pass events, data is extracted after post processing allowing the AI to analyse the players aesthetic intent and stylized output rather than raw scene data
3. It makes use of asynchronous GPU readback allowing data to be transferred from the GPU to CPU without stalling the main thread. While it introduces a slight delay before data is available to the CPU, this prevents stuttering within the application maintaining a smooth frame rate required for user comfort.

3.3 Settings and Post processing

The prototype implements the exposure triangle (ISO, Shutter Speed, and Aperture) to provide a focused pedagogical environment. These settings are simulated via the Unity Post-Processing system, utilizing logarithmic functions to translate setting values into "stops" for accurate exposure adjustment.

- While ISO and Shutter Speed utilize the Colour Adjustments (Post-Exposure) and Film Grain overrides, motion blur is handled via a custom shader rather than the standard URP velocity blur. This allows for a more accurate simulation of motion trails based on exposure duration, better representing long-exposure photography.
- Aperture, zoom and focus all make use of the Depth of field override for creating blurriness or bokeh via direct values taken from the “physical camera”
- Post-processing effects are isolated to the camera’s viewfinder rather than the user’s overall VR vision. This prevents motion sickness and ensures the simulation mirrors real-world perception, where camera settings only alter the captured medium, not the photographer's vision.

By passing the camera settings, which shows the user’s intent, alongside the post-processed image, which shows what actually happened, the system is designed this way to allow the AI to be able to identify specifically where a player's technical understanding faltered, for example, recognizing that an overexposed image was the result of a shutter speed set too low for the current ISO.

Inputs are mapped to the Quest 3 controller's 'A' and 'B' buttons to replicate thumb-operated DSLR controls, further reinforcing sim-to-real muscle memory. The inclusion of an on-camera gallery and a dedicated Review Station align the virtual workflow with the professional photographic process of capture, review or delete in-camera, and process externally.

3.4 System Architecture

Because the AI models and metric calculations are computationally expensive, and require a Python environment, they cannot run natively on the VR headset without massive performance drops making the game unplayable. To solve this problem, we went with a 3-stage setup, we have the unity environment, a flask server that acts as a bridge, and a python backend. Allowing the heavy lifting to happen outside of the game loop, subsequently allowing the game to maintain a consistent frame rate.

The python backend needs the raw image, the depth data and the camera settings that were used to take the photo. On one hand we could simply send the raw bytes and hope the correct data is being sent, however this system wouldn't be robust or explainable. In our implementation when the image is captured in unity it saves each set of data to a file, saved with the atomic timestamp of when it was taken acting as a unique identifier to ensure the image, depth data, and metadata files are tethered to one another. The paths of these files are then sent within the flask post request, becoming a pointer-based handover, saving memory by not duplicating large data such as image buffers in the flask request. Performing the request this way allows us to check the data is sent correctly and allows for far easier debugging should issues arise, along with being able to manually audit the data if necessary.

Finally, this pipeline also allowed for much faster prototyping with modifying the AI models to run within the unity application would have taken significantly longer and is something the development timeframe did not permit.

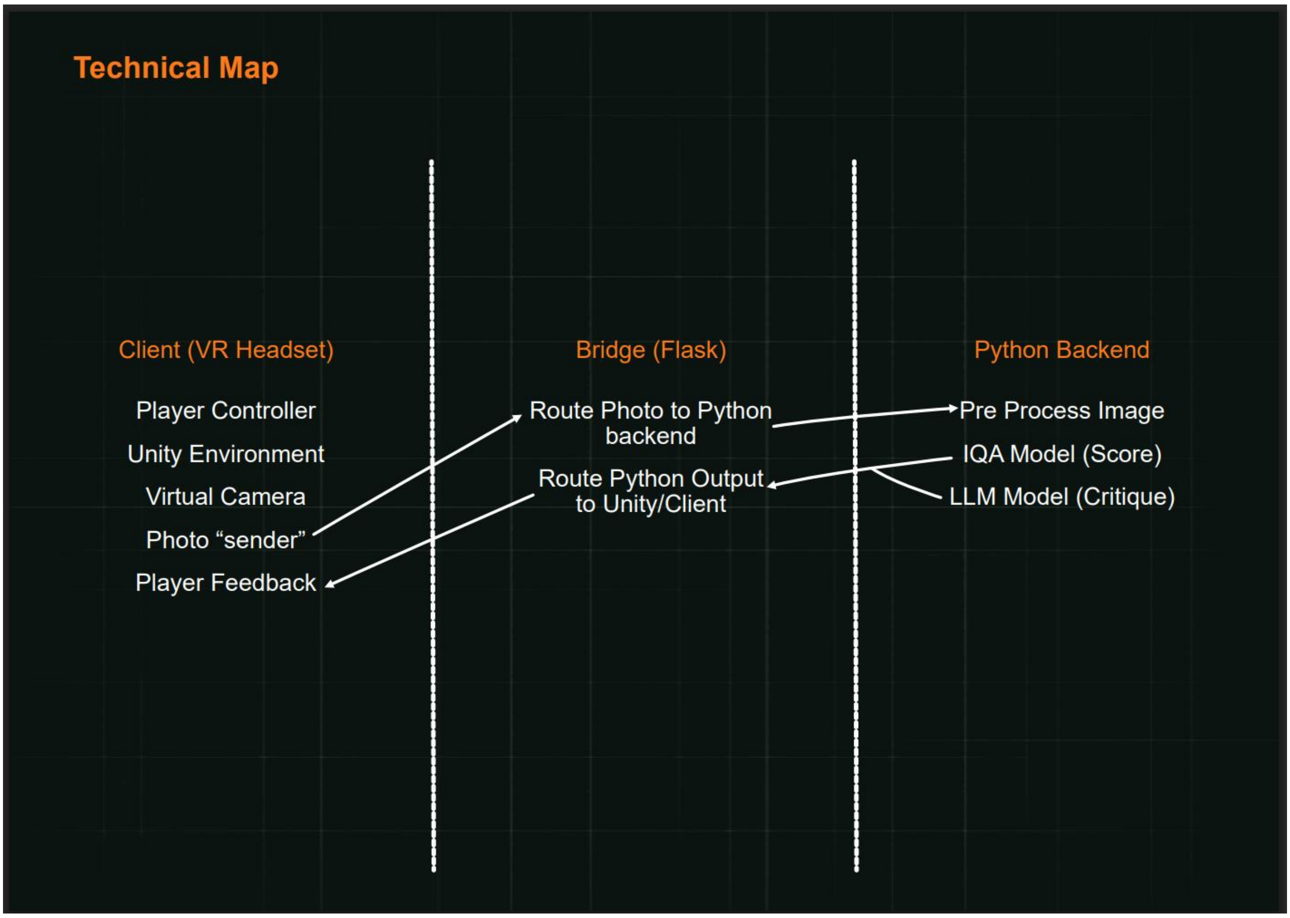

Figure 5 Technical map displaying 3-stage setup for application

3.5 Feedback UX & Diegetic UI

The display strategy for AI-generated feedback was evaluated across three methods: non-diegetic popups, on-camera UI, and dedicated in-world assets. While on-camera displays align with DSLR hardware, the limited screen space was found to increase cognitive load, forcing users to switch between a 3D photography task and a complex 2D analysis. Similarly, 2D overlays were rejected as they necessitate "context switching," breaking player immersion and presence, both of which are critical for pedagogical outcomes and skill retention.

The application utilizes a diegetic "Review Station", a tablet-like interface integrated into the 3D environment. This provides the necessary visual real estate for detailed technical analysis while mirroring the real-world transition from field capture to a dedicated workstation. By keeping the player within the 3D world, this method preserves the "flow state" and maintains a consistent level of presence during the pedagogical feedback loop.

4 AI/PYTHON BACKEND

Figure 6 Step-by-step overview of end-to-end process of taking a photo to getting feedback

The backend utilizes a multi-modal hybrid pipeline (Unity, Flask, and Python) to provide reliable, explainable pedagogical feedback. By combining Image Quality Assessment (IQA), Deterministic Metrics, and a Vision-Language Model (VLM), the system ensures that the "black box" nature of aesthetic scoring is grounded by actionable technical data. An overview of this system can be seen in Figure 6.

4.1 Image Quality Assessment (NIMA)

The first stage of the image analysis Utilizes Google’s Neural Image Assessment (NIMA) model. This model was chosen as:

- Traditional image quality assessment often focuses on pixel-level errors like blur or noise whereas NIMA is designed to predict human aesthetic preferences by training on the Aesthetic Visual Analysis data set containing about 255,000 images, rated based on aesthetic qualities by amateur photographers (Talebi & Milanfar., 2018).
- NIMA is built upon the MobileNet architecture, which minimizes computational overhead by using depth wise separable convolutions (Howard et al., 2017), a necessity for maintaining the performance of our hybrid pipeline without affecting the VR frame rate.
- Finally, a critical feature of NIMA is that it uses Earth Mover’s Distance (EMD) as a loss function.

Unlike standard models that predict a mean score, EMD allows for prediction of a probability distribution of ratings. As explained by Hou et al. (2016) “The EMD^2 loss uses the predicted probabilities of all classes and penalizes the miss-predictions according to a ground distance matrix that quantifies the dissimilarities between classes.” Essentially, an incorrect prediction is penalized for how far it is from the correct rating. By capturing the full probability distribution, the system can also potentially identify 'polarizing' images that a traditional model might miss.

This project sequentially runs two NIMA models loaded with Aesthetic and technical weights respectively. The reasoning behind this was to give more detail as to where the submitted photo may be lacking, it could get a high technical rating, but the aesthetic score be low, showing they have a grasp of the settings to take a photo, but are less aware of what makes a photo visually interesting.

Figure 7 Cropped version of the step-by-step overview showing which files are used for NIMA scores and metrics

4.2 Metrics & Subject detection

The backend makes use of three metrics to ground the AI feedback and combat hallucinations, Focus, Composition and Exposure as these are the most relevant metrics to the implemented virtual camera. Generating these gives data that helps explain NIMA scores and gives the VLM exact scoring to explain where users might have gone wrong. Finally, providing these metrics to the user along with the NIMA score and VLM feedback allows them to see specifically where their image is underperforming.

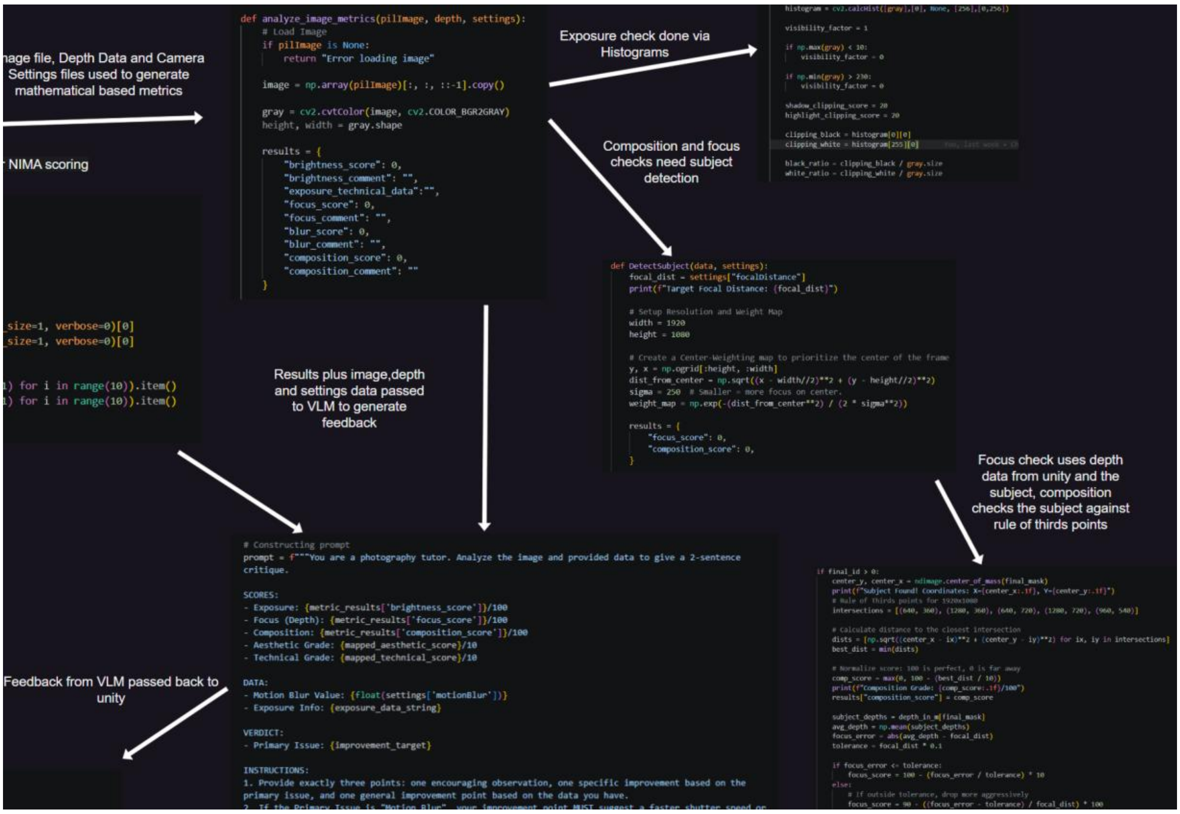

Figure 8 Cropped in version of step-by-step diagram, showing specific metric steps

4.2.1 Subject detection

For the focus and composition, we initially calculated them based on the entire image, however early in development this was proven to be unreliable and typically incorrect based on the image provided. Due to this we transitioned to determining the subject of the image and performing the focus and composition checks using that knowledge for greater accuracy overall.

While we could have made use of another AI model for this task, to reduce the computational load we calculate the subject based on the provided depth data and provided focus distance from the camera settings. Using this data allows for accurate prediction of the subject in any image provided from the application that typically wouldn't be available from a traditional photograph, because we know where the camera is focused, we know roughly where the subject should be. However, if the user gets the focus completely wrong the algorithm cannot calculate the subject for any reason a fallback is required: reverting to the previous, less accurate method of calculating the metrics based on the entire image.

To determine the subject, Sobel filters are used to find sharp depth changes between objects, isolating distinct 3D forms rather than treating the entire scene as a flat 2D plane. Then we apply a focal mask to the result of that calculation allowing for some leeway with the focus due to the depth of objects themselves, or the user not focussing the camera perfectly. We then use a binary opening function on the in-focus pixels to ensure that objects that are touching are snapped apart, finding the best subjects via labelling each "island" in the image, scoring them based on how close they are to the centre of the image to account for beginner centring bias, then using the winning id for subsequent metrics.

4.2.1.1 Composition

For the composition metric we take the winning subject and compare its central pixel position to the closest intersection point with the rule of thirds points for the given resolution, in this prototype 1920x1080, returning a score out of 100.

4.2.1.2 Focus

For the focus metric we make use of our depth data and compare the average depth of the subject against the focus distance from the camera. We apply a tolerance threshold to account for the physical depth of the subject itself, ensuring that an object isn't marked as 'out of focus' because the camera is focused on its leading edge. Generating a score out of 100 based on the "focus error".

4.2.2 Exposure

For generating the exposure score we made use of histograms allowing us to generate specific exposure metrics like clipping values, shadow density, highlight density and tonal range along with overall brightness. Applying this solution rather than an overall brightness check allowed for a far more accurate, reliable calculation and allows the system to provide the VLM with specific data to explain why the exposure was incorrect on the image.

4.3 Vision-Language Model (Llama Vision 3.2)

For the final stage of image analysis, the feedback generation we use Llama Vision 3.2 via the Ollama framework. We use the VLM rather than the standard large language model because it allows us to prompt it with the image via an adapter that integrates the pre-trained image encoder to the pre-trained language model (Meta, 2024), providing further grounding data for feedback.

By using the Ollama framework via a local API we can run the entire application locally ensuring that the image and JSON data packets are processed entirely on the host machine. Alongside this the model is open source and free making this prototype more accessible without having to pay for tokens. However, this does mean that the model performs worse than a cloud service due to limited computational resources available, overall, this setup provides a balance between speed and reasoning capability.

```
# Constructing prompt
prompt = f"""You are a photography tutor. Analyze the image and provided data to give a 2-sentence critique.

SCORES:
- Exposure: {metric_results['brightness_score']}/100
- Focus (Depth): {metric_results['focus_score']}/100
- Composition: {metric_results['composition_score']}/100
- Aesthetic Grade: {mapped_aesthetic_score}/10
- Technical Grade: {mapped_technical_score}/10

DATA:
- Motion Blur Value: {float(settings['motionBlur'])}
- Exposure Info: {exposure_data_string}

VERDICT:
- Primary Issue: {improvement_target}

INSTRUCTIONS:
1. Provide exactly three points: one encouraging observation, one specific improvement based on the primary issue, and one general improvement point based on the data you have.
2. If the Primary Issue is "Motion Blur", your improvement point MUST suggest a faster shutter speed or more stable position.
3. DO NOT mention scores, raw numbers, or describe the image.
4. DO NOT repeat any advice.
5. Max Three sentences total. Be punchy.
"""
```

Figure 9 Screenshot displaying the Llama-Vision 3.2 prompt highlighting the structure and how the data is being used to ground the model

The prompt was engineered to provide a specific information hierarchy. By giving the model, a calculated “verdict” derived from the earlier metrics stage, the VLM is ‘anchored’ to the technical ground truth. Figure 9 shows the instructions mandate a three-point structure: an encouraging observation, a specific improvement based on the primary technical “failure”, and a general tip to keep user motivated whilst improving. This structure ensures the feedback remains actionable and punchy, mirroring the critiques a student might receive from a real photography tutor.

4.4 Image assessment summary

The final stage requires packaging the multi-modal outputs into a unified data structure for transmission. Once the Llama model generates the natural language critique, it is bundled with the numerical NIMA scores and technical metrics into a final JSON object. This payload is sent back across the Flask "bridge" to the Unity client. By utilizing an asynchronous response, the system ensures that the VR environment remains responsive while the "Review Station" updates its UI components.

Figure 10 Cropped version of the step-by-step overview showing the final stage of sending the feedback back to the Unity frontend via Flask

As shown in Figure 10, the Unity client parses the JSON response to populate the specific score fields and the main critique text box on the Review Station, completing the unsupervised learning loop. Figure 11 shows an example of what the user would see after the feedback is returned to Unity.

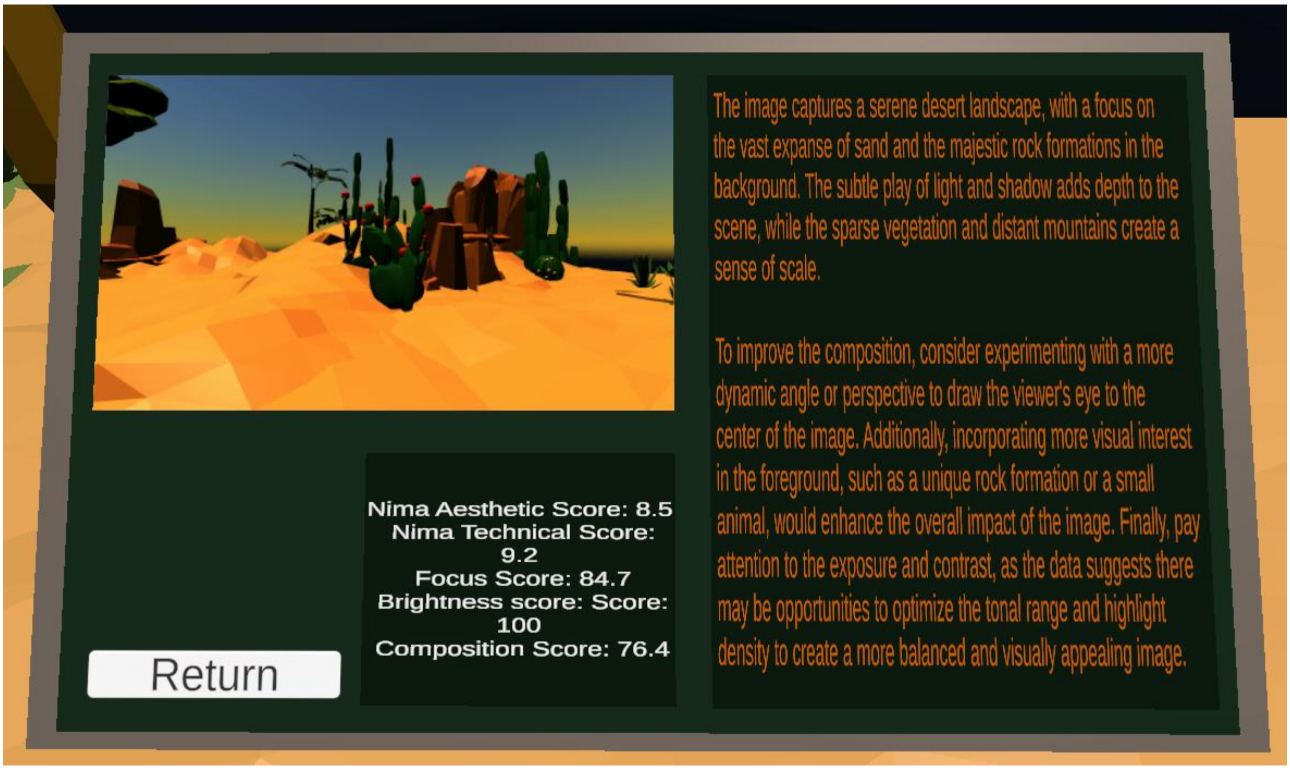

Figure 11 An example of the feedback the user sees in Unity after Python sends it back

5 STUDY

5.1 Plan

Due to hardware and logistical constraints, the study pivoted from an in-person VR trial to a remote visual heuristic evaluation. This approach utilized convenience sampling to gather a broader dataset of expert and novice feedback within the project's timeframe. The evaluation was conducted via a five-stage questionnaire designed in accordance with the BU Research Ethics Checklist. The first two stages being the consent form and demographic questions.

5.1.1 Usability and Heuristic Evaluation (RQ3)

Participants were presented with a high-fidelity video demonstration covering the end-to-end user journey, from physical lens interaction to AI feedback delivery. Usability was measured using a 5-point Likert scale mapped to Nielsen's (1994) Usability Heuristics as shown in Figure 12.

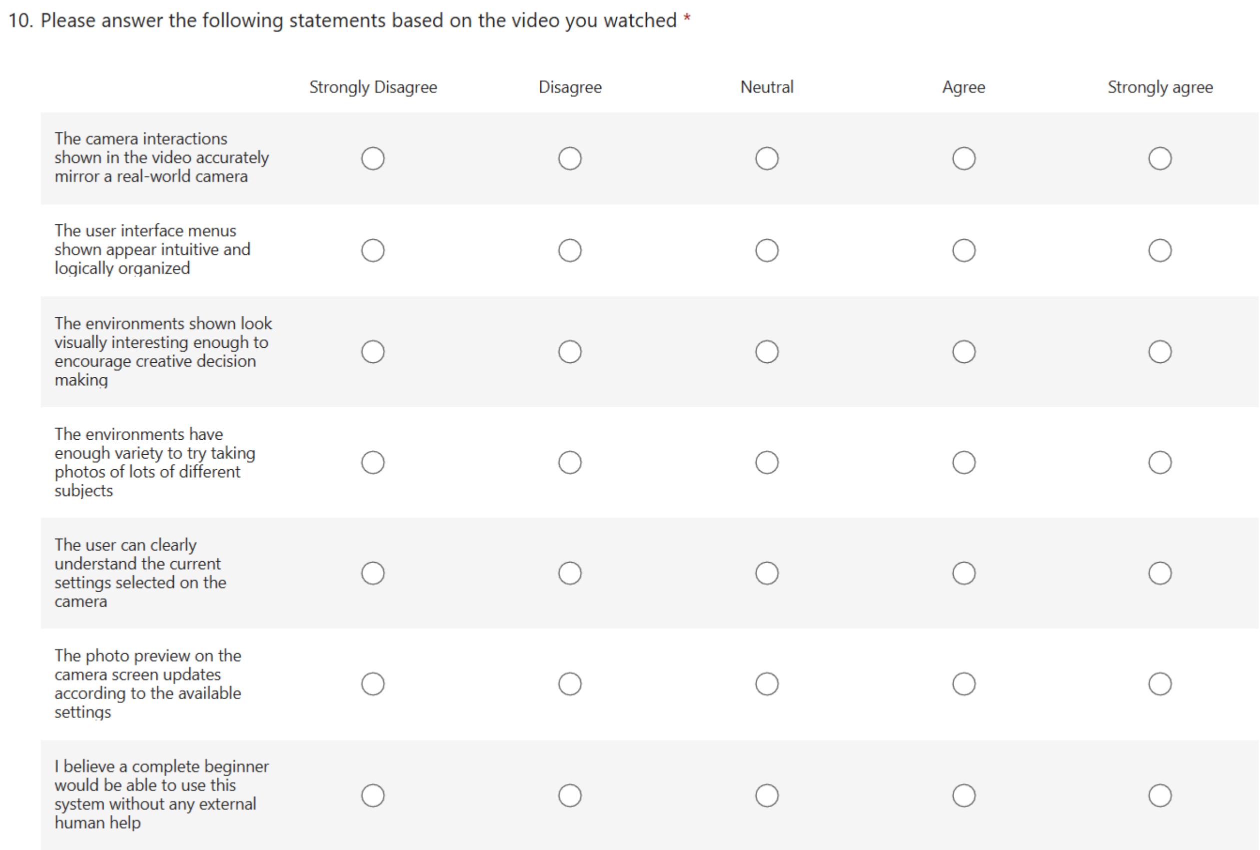

10. Please answer the following statements based on the video you watched *

	Strongly Disagree	Disagree	Neutral	Agree	Strongly agree
The camera interactions shown in the video accurately mirror a real-world camera	○	○	○	○	○
The user interface menus shown appear intuitive and logically organized	○	○	○	○	○
The environments shown look visually interesting enough to encourage creative decision making	○	○	○	○	○
The environments have enough variety to try taking photos of lots of different subjects	○	○	○	○	○
The user can clearly understand the current settings selected on the camera	○	○	○	○	○
The photo preview on the camera screen updates according to the available settings	○	○	○	○	○
I believe a complete beginner would be able to use this system without any external human help	○	○	○	○	○

Figure 12 An image showing the Likert scale statements participants answered

This allowed the system to be evaluated against industry-standard principles (see Table below) without requiring participants to have prior theoretical knowledge. In this section there is also chance for the user to provide specific feedback on usability issues that aren't covered by the Likert scale statements.

Questionnaire Statement	Nielsen's Heuristic (Nielsen, J., 1994)	Rationale
The camera interactions shown in the video accurately mirror a real-world camera	Match between system and the real world	Camera interaction should accurately reflect real-world camera function to promote

		sim-to-real skill transfer
The user interface menus shown appear intuitive and logically organized	Consistency and standards	Camera UI should follow established design convention to reduce user cognitive load
The environments shown look visually interesting enough to encourage creative decision making	Aesthetic and minimalist design	Environments should provide enough visual interest for quality photos to be taken.
The environments have enough variety to try taking photos of lots of different subjects	Aesthetic and minimalist design	Environments should have enough variety to promote creative decision making and artistic choices.
The user can clearly understand the current settings selected on the camera	Visibility of system status	Diegetic camera UI should keep the user informed of their current technical parameters
The photo preview on the camera screen updates according to the available settings	Visibility of system status	The visual changes to the camera preview should change in accordance with setting changes
I believe a complete beginner would be able to use this system without any external human help	Help and documentation	The application should be usable by any user regardless of ability levels

5.1.2 AI Reliability and Performance (RQ1 & RQ2)

The second phase focused on the perceived performance and reliability of the AI analysis pipeline. To mitigate participant fatigue while ensuring coverage of the 120 assessment images, a quasi-random distribution was employed. Participants selected a number (1–6) to be assigned a specific subset of 20 images (10 high-resolution captures and 10 corresponding AI critiques as shown in Figure 13).

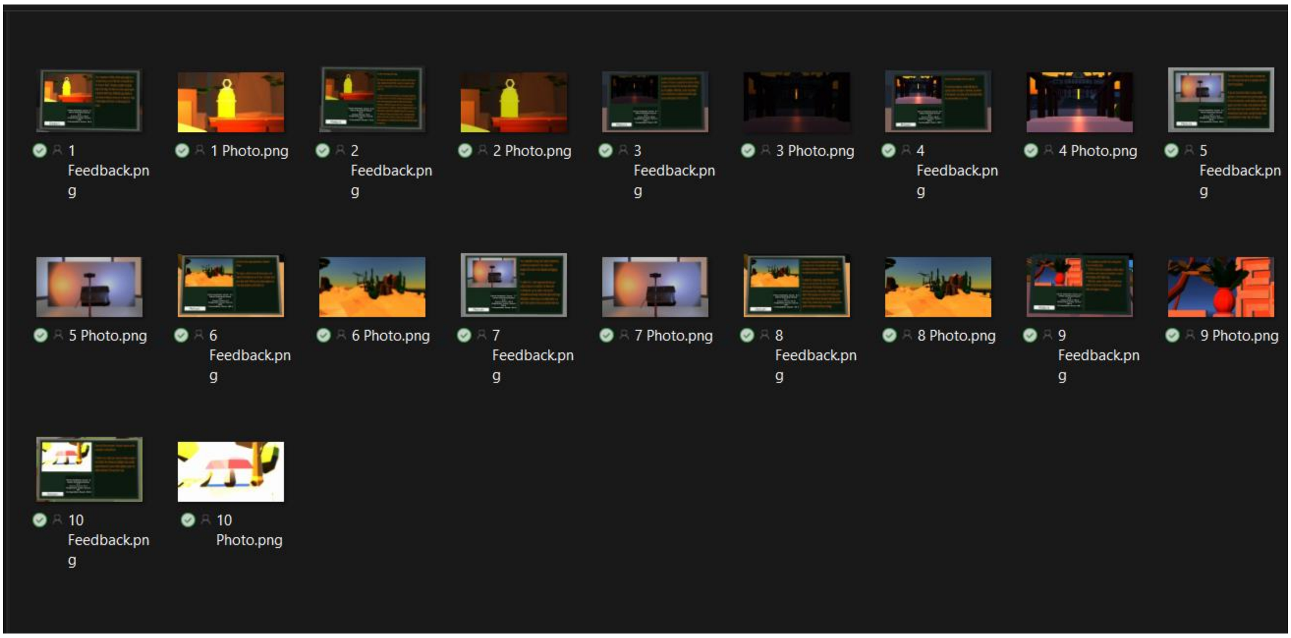

Figure 13 File structure that participants saw when reviewing image sets

This blind-choice method facilitated a balanced review of the dataset despite platform constraints. Participants evaluated the accuracy of the AI's critique relative to the visual evidence, providing primary data for RQ1 (AI effectiveness) and RQ2 (AI influence on decision-making).

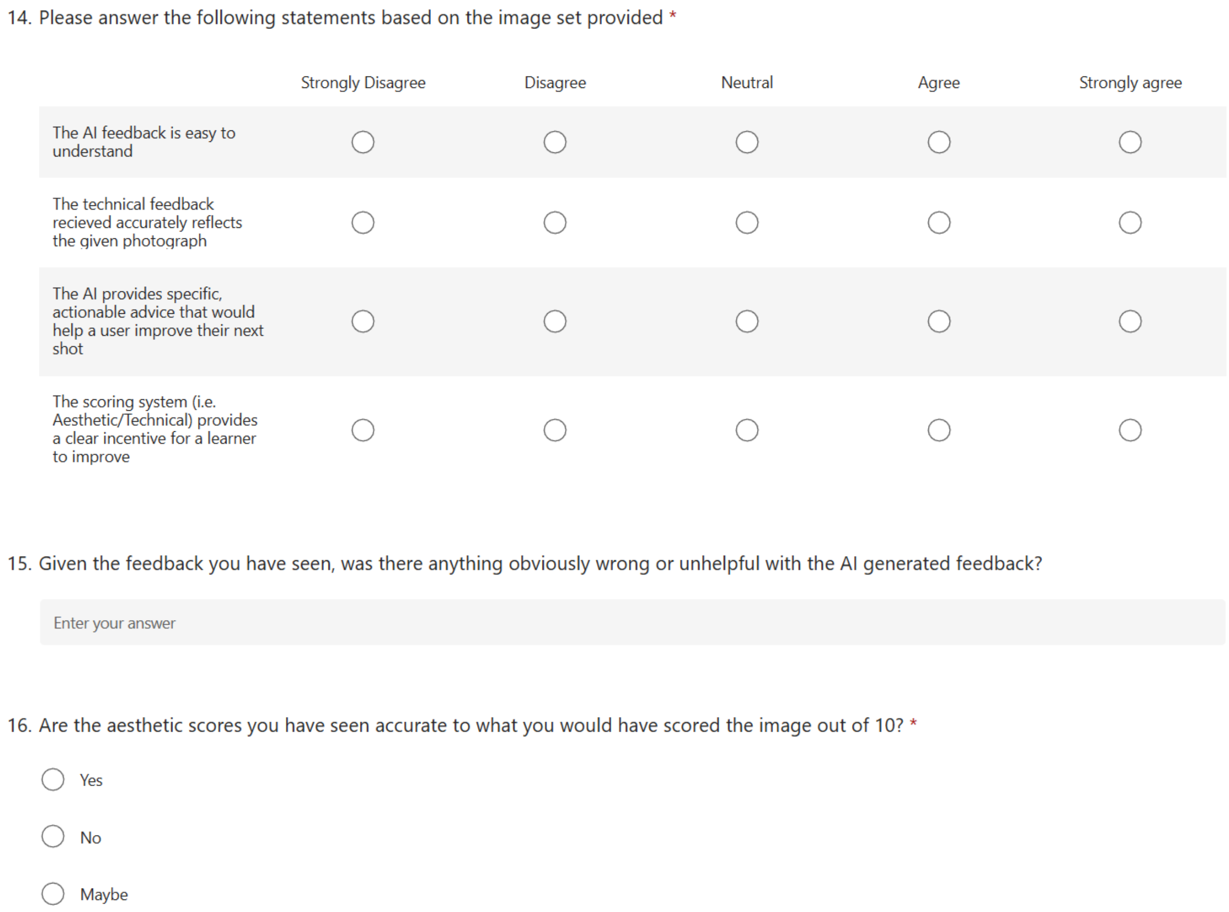

14. Please answer the following statements based on the image set provided *

	Strongly Disagree	Disagree	Neutral	Agree	Strongly agree
The AI feedback is easy to understand	◯	◯	◯	◯	◯
The technical feedback recieved accurately reflects the given photograph	◯	◯	◯	◯	◯
The AI provides specific, actionable advice that would help a user improve their next shot	◯	◯	◯	◯	◯
The scoring system (i.e. Aesthetic/Technical) provides a clear incentive for a learner to improve	◯	◯	◯	◯	◯

15. Given the feedback you have seen, was there anything obviously wrong or unhelpful with the AI generated feedback?

Enter your answer

16. Are the aesthetic scores you have seen accurate to what you would have scored the image out of 10? *

◯ Yes

◯ No

◯ Maybe

Figure 14 The first three questions participants answered on image sets, including the Likert statements

5.1.3 Reflection and Qualitative Feedback

The final section utilizes a mixed-methods approach, combining multiple-choice questions for consensus data with open-ended reflections. This allowed for the identification of usability edge cases and pedagogical gaps not captured by the previous sections, providing a comprehensive basis for the final evaluation.

5.2 Results and Evaluation

5.2.1 Demographics

Overall, the study was well balanced with an even split of men and women taking part, a range of ages with a slight bias towards the 18-24 range and a variety of educational backgrounds. Most users identified themselves as beginner photographers that were somewhat familiar with technical photography terms. With some users being of higher skill, and some having never heard of the technical terms. Most participants never use VR and completed the questionnaire on their smartphones.

5.2.2 Heuristic evaluation results

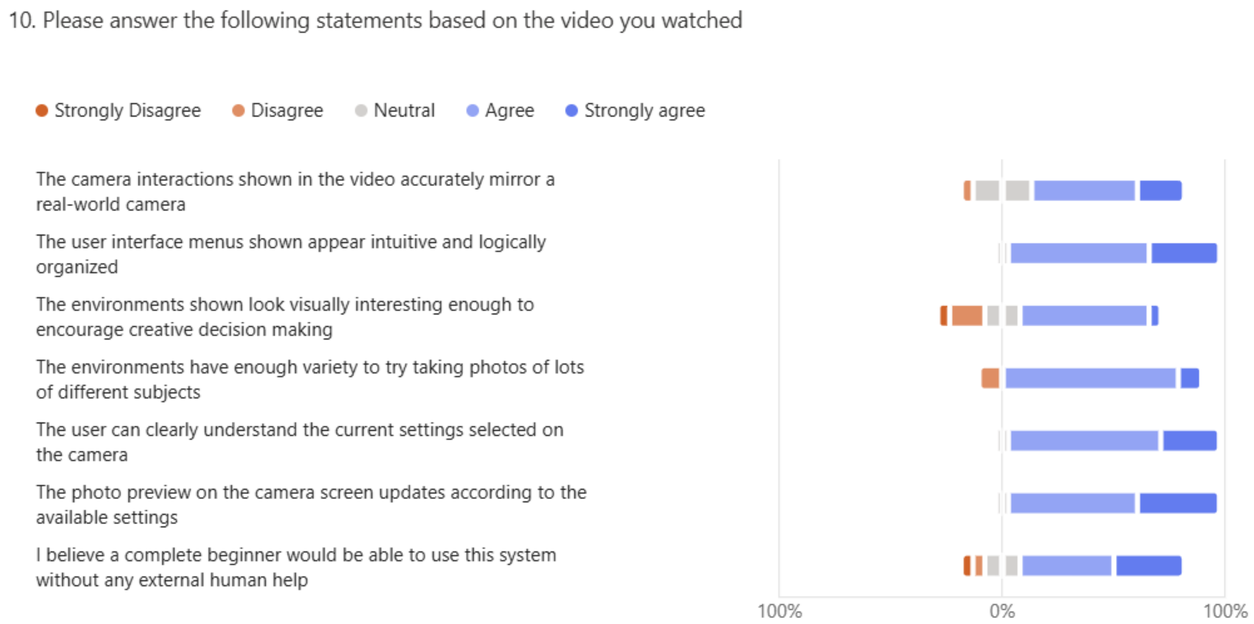

Figure 15 Likert scale responses for the video heuristic evaluation

As illustrated in Figure 15, the heuristic evaluation indicates a positive consensus across all key usability metrics. Most notably, the 'Visibility of System Status' and 'Consistency and Standards' categories received near-unanimous positive feedback, confirming that the diegetic UI and real-time camera updates are intuitive and follow established design conventions to reduce cognitive load.

While the 'Aesthetic and minimalist design' category was also viewed positively, a small cluster of neutral and negative responses suggests that some users found the environments lacked sufficient diversity. However this may be a limitation of the video demonstration format, which only showcased a subset of available levels, rather than a fundamental failure of the environment design.

Furthermore, the positive feedback regarding the 'Match between system and the real world' statement validates the decision to use physical lens rings and a joystick-driven interface. This provides primary evidence for RQ3, suggesting that the diegetic interaction model successfully mirrors a real-world DSLR to facilitate sim-to-real skill transfer.

2	anonymous	Hard to see exactly what is happening on the camera, on my smartphone
3	anonymous	The settings of the camera are clearly presented, but could use an explanation for those not familiar with cameras. Something akin to a tooltip of sorts could help?
4	anonymous	@1:16 The UI on the camera looks difficult to use and needs improved legibility. This may be down to the video and be improved when using a vr headset however. @0.52 Would be easier to provide a menu to select which lens to use instead of physical objects in the environment, especially since the focus seems to be on teaching shot compositions.

Figure 16 Specific qualitative feedback from participants on the usability issues they observed

Figure 16 shows a few specific points of usability feedback, a few mentioning the shortcomings of VR video recording, but more notably pointing out the lack of tuto rial or tooltip explanations for beginners to photography and potential streamlining of physical interaction to focus on teaching photography techniques however figure 17 also shows they didn't believe these would hinder beginners that much. Figure 15 also shows the majority of participants believe the application could be used by a complete beginner with no external help, alligning with the User Experience Success Criteria, demonstrating that the application functions effectively as an unsupervised learning environment.

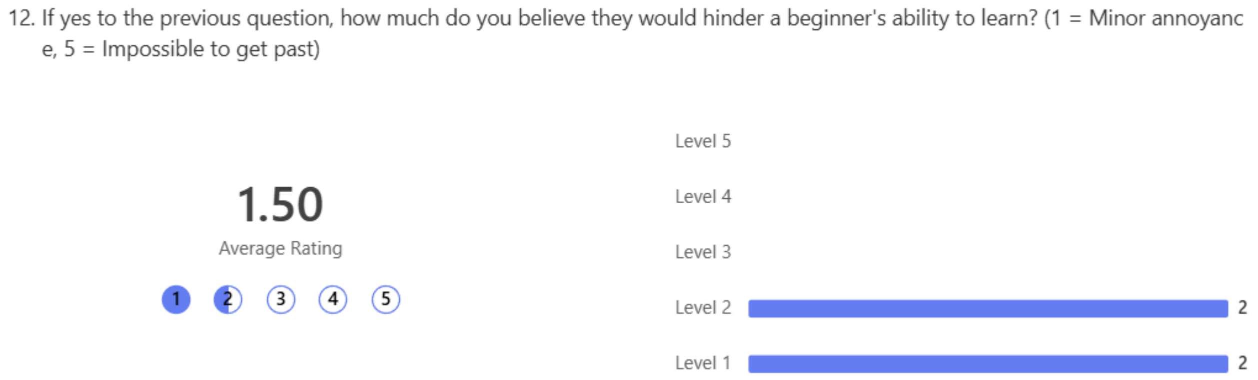

Figure 17 Participants judgements of how much the usability issues they noticed would impact a beginner

5.2.3 AI reliability and effectiveness results

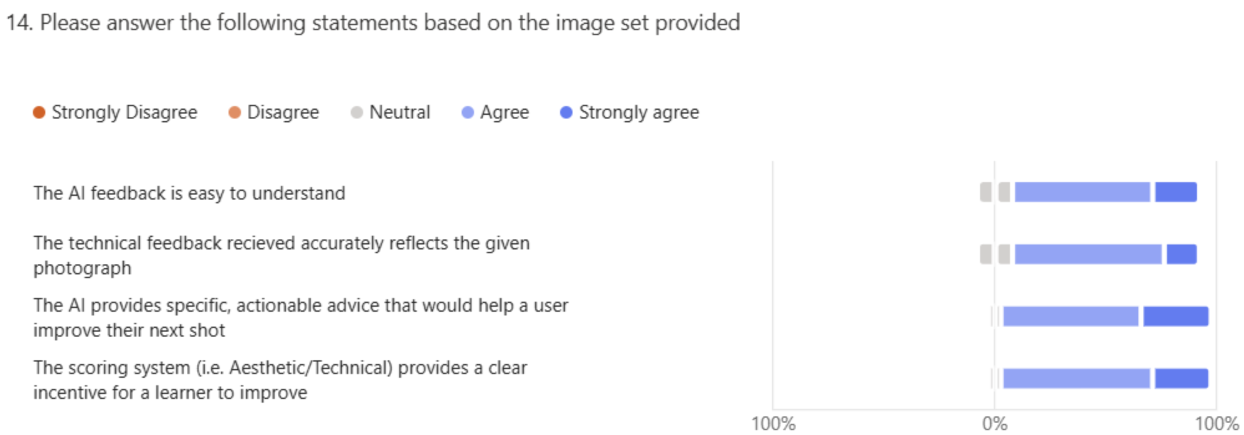

Figure 18 Participant responses to AI feedback Likert statements

The data regarding the AI pipeline indicates strong perceived reliability and pedagogical value. As shown in Figure 18, there is a near-unanimous consensus that the AI-generated feedback is understandable, accurate, and provides a clear incentive for learners to improve. This provides primary evidence for RQ1, suggesting that automated AI evaluation is a viable tool for facilitating unsupervised skill development.

2	anonymous	All-dark images should not be able to be scored - Photo #10 got scores for categories other than brightness, and the feedback did not really fit seeing as the image was completely dark and no composition is visible.
3	anonymous	If you're not a photographer, some technicalities will be unclear. Also some sentences are wordy because of the way AI writes so it could be difficult to understand. Having an option where language will be simplified and technicalities explained for beginners would be very helpful. Essentially, a skill option where AI can generate feedback appropriately to the experience of the audience (if professional photographer then ai uses more technical feedback).
4	anonymous	In the desert scene the ai suggested moving the subject closer to the background. That would be difficult in a real world scenario
5	anonymous	It felt oddly 'soft' about the entirely black image in feedback 10. Maybe make it a bit more harsh in extreme cases?
6	anonymous	The AI's use of technical language would be unhelpful to a person unfamiliar with photography terms. Perhaps an optional glossary included on the screen would resolve the issue and allow instant gratification when a beginner can apply the terms to their own photos immediately.
7	anonymous	None
8	anonymous	Doesn't tell u what the scores are out off, they can't all be out of 100

Figure 19 Specific participant feedback on generative AI feedback and suggestions

Despite the positive consensus, qualitative feedback in Figure 19 identified technical edge cases. Participant 2 noted that an entirely dark image still received scores; this is a byproduct of the Render Graph capture pipeline, which extracts depth data directly from the GPU buffer, allowing the metrics to "see" the 3D scene even when the final color render is zeroed out.

Furthermore, feedback regarding the "wordiness" of the Llama model (Participants 3 & 6) highlights a common challenge in VLM prompting known as **"**attention drift". While the 11B local model offers a balance of speed and reasoning, it occasionally prioritizes prompt-structure over conciseness. As suggested by participants, future iterations could benefit from a "difficulty slider" or a localized glossary to better tailor technical language to the user's specific experience level.

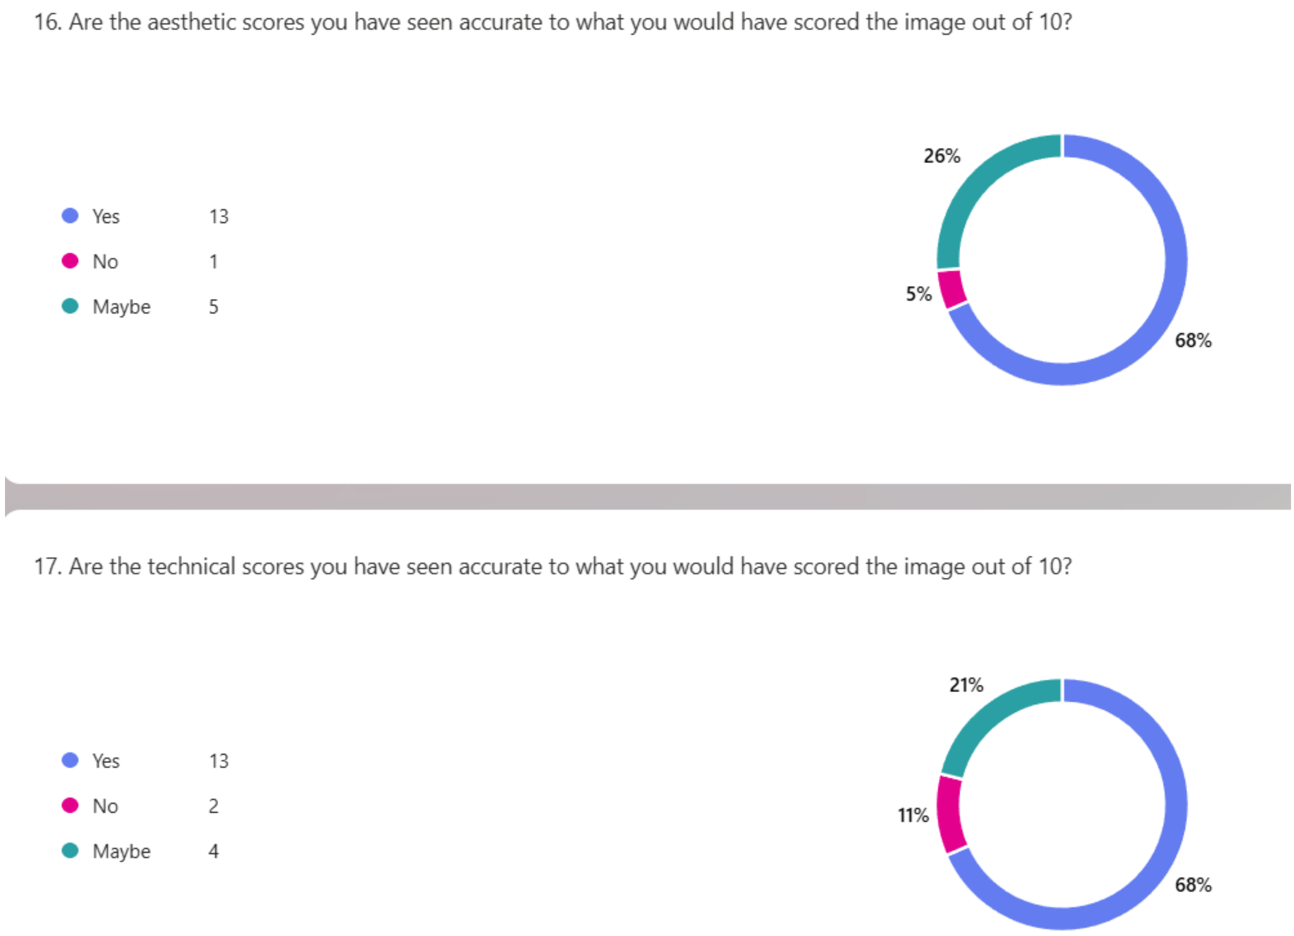

Figure 20 Participants responses to the accuracy of NIMA aesthetic and technical scores

A significant finding is the perceived utility of the deterministic metrics. While participants generally agreed with the NIMA aesthetic scores (68% 'Yes' in Figure 20), Figure 21 shows that 95% of respondents found the mathematical metrics (Exposure, Focus, Composition) to be useful. This validates the decision to include a metrics-driven stage in the pipeline. Qualitative reflections in Figure 22 reinforce this, with participants noting that while aesthetic scores are inherently subjective, technical metrics provide a "useful benchmark" and an "undeniable data point" to explain performance gaps. This finding alligns with trends noticed during the development stage that the NIMA scores were much less predictable, along with the scaling issues, making them less useful overall and possibly just unnecessary.

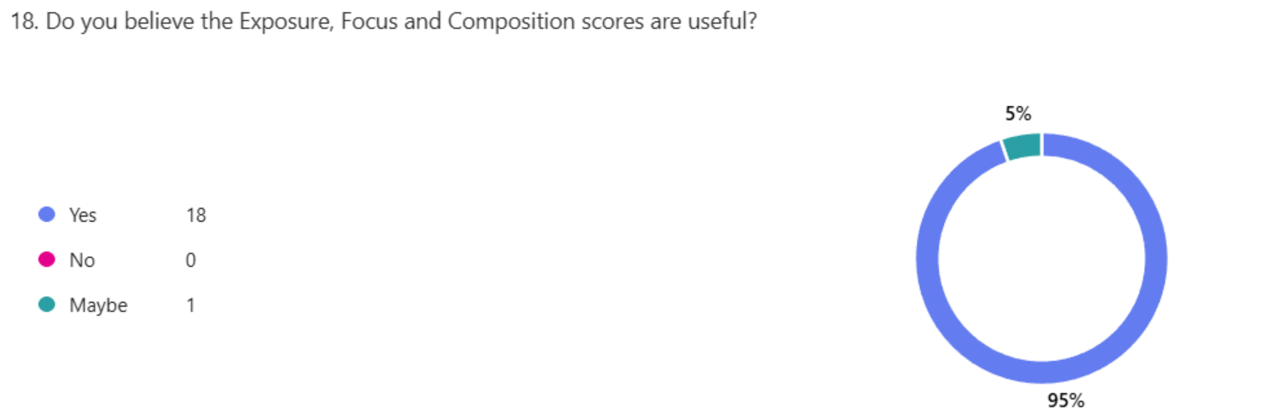

Figure 21 Participants opinion as to whether the Exposure, Focus and Composition metrics are useful for learning

ID ↑	Name	Responses
1	anonymous	Always good to have a visual score to try and better
2	anonymous	Aesthetic score is a maybe because of the before-mentioned issue.
3	anonymous	Aesthetic score is subjective that's why I placed a maybe.
4	anonymous	Any feedback is good when taking photos. In this case it allows you to focus in on areas that will make the most difference. ie a low focus score would tell me I need to refocus.
5	anonymous	The scores provide useful feedback to enable you to better your photography
6	anonymous	Focus, Exposure and Composition being separate scoring are excellent for putting into picture what exactly new/aspiring photographers can improve on. Technical scoring however felt a bit inconsistent, as seen in feedback images 1 & 2.
7	anonymous	So you can see if you require more or less of each etc
8	anonymous	The AI has accurately broken down the various aspects. The numbers are a useful benchmark for learners to see clear improvement between shots/
9	anonymous	None
10	anonymous	It shows what part of the image is lacking and shows what could be improved
11	anonymous	I find that the scores make sense very much and it reflects the state of the photo and the opinion of the AI however I find the scale weird like a white image got 3.2 and 10 (I think?) when it should have none, but overall the metric is good
12	anonymous	It give you more of a guide and reference
13	anonymous	Yes the scores are useful so you can improve the quality of your photos
14	anonymous	No

Figure 22 Participants explanations as to why they do or do not believe the metrics are useful for learning

Figure 23 shows that participants generally felt all parts of the feedback were useful with the slight exception of a few participants noting that subjective scoring such as the aesthetic score and composition scores were not as helpful as the others. Showing that the choice to include all the scoring and feedback was valid, however there is still room for improvement.

20. Are there any parts of the feedback you felt weren't as useful or entirely unneccesary?

19 Responses

ID ↑	Name	Responses
1	anonymous	No
2	anonymous	no
3	anonymous	As stated previously
4	anonymous	It is all up to the individual, I think if someone wants to practise technical side of things or aesthetics then they will look at those specific scores and ignore others. Otherwise I feel like all three are important
5	anonymous	Composition scoring is less helpful as that is, in my opinion, completely subjective and personal to the photographer
6	anonymous	No
7	anonymous	None.
8	anonymous	I'm not sure
9	anonymous	No, the AI feedback is fair and well evaluated. My only suggestion is as stated above in the first AI feedback question
10	anonymous	None
11	anonymous	If there was a new learner they might not know what the aesthetic or technical scores are, maybe would need an explanation
12	anonymous	No
13	anonymous	No, they all made sense
14	anonymous	No
15	anonymous	No
16	anonymous	No
17	anonymous	No

Figure 23 Participants explaining if they believed any of the feedback wasn't useful, only the aesthetic score was suggested to be less useful than the others

5.2.4 Reflection results

Figure 24 shows 95% of participants believe the application would be helpful for learning photography, successfully meeting the core user experience success criteria and proving the formula of this application is viable and worthy of future development and research.

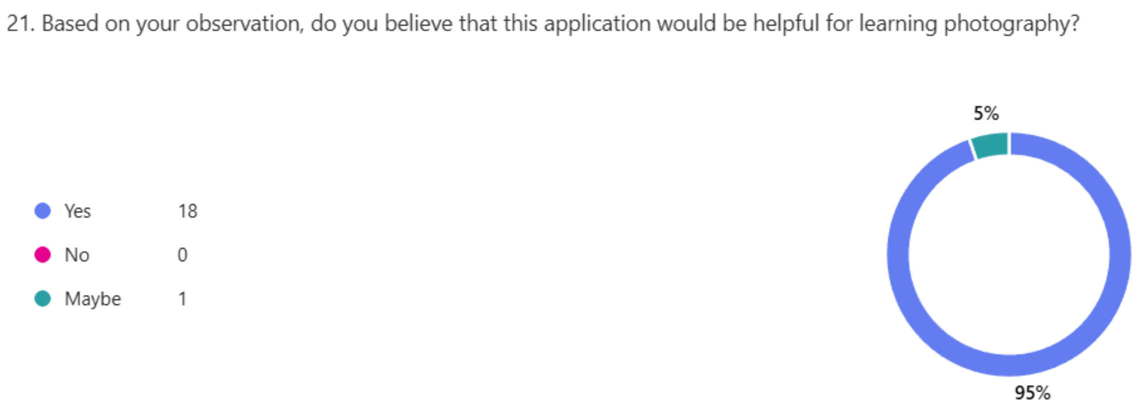

Figure 24 Participant consensus of the usefulness of the application for learning photography

The responses to question 22 (Figure 25) clearly show there is grounding for sim-to-real transfer learning with this application providing further supporting evidence for RQ3 stating that virtual photography skills are translatable to real life. This high level of confidence is notable given that the majority of the participants identified as "beginner" or "intermediate". If another study were done where these users were able to use the application we would expect the somewhat confidence to turn into extremely confident as this is likely a limitation of the study format making it more difficult to judge.

22. How confident would you feel applying the technical principles shown in the application to real life?

Figure 25 Participants confidence levels of applying the technical principles they saw within the video and image set

Finally figure 26 displays some important improvements participants would like to see in future iterations of this application, notably "skill option" settings where users have levels to select from to change how the feedback behaves or the features available to them, a dedicated tutorial/learning features and improvements to environments.

24. If this tool were made available as a standard educational program for photographers what would you like to see added or improved?

14 Responses

ID ↑	Name	Responses
1	anonymous	Unsure
2	anonymous	More diverse sceneries and settings, for example more action, moving shots, etc.
3	anonymous	As previously mentioned, have an option screen (maybe I have missed it I'm sorry) of your skill. So if you're a beginner; the ai feedback should include basic language that would also explain ALL the technical matters such as rule of thirds. Maybe it would be useful to have a demo for beginners and make them go through each technicality (aka take a photo for rule of thirds, take a photo for exposure, etc). Whereas for more professional ones kinda skip that step and go straight for it.
4	anonymous	A tripod feature and improvements to the ai responses in some areas
5	anonymous	Updated camera model to accurately reflect real cameras, which options for different models ect
6	anonymous	Some busier, larger or colorful scenarios could be nice for a wider range of possible shots.
7	anonymous	Maybe just more scenes and objects
8	anonymous	The elements such as lenses should be simplified. As while they do snap on, it feels unnecessary to have the lens as objects within the environment instead of just a UI menu. Additionally, While the AIs feedback is useful, it's terminology would be off putting to beginners who aren't familiar with the terms. As stated previously, an accessible glossary or potentially the ability to hover the cursor over terminology to get a definition would prove useful for beginners.
9	anonymous	None
10	anonymous	A tutorial about the camera and its settings and what they do.
11	anonymous	more advanced levels, alongside more basic levels, teach the person what is exposure, then teach him how to use it, then teach him more advanced techniques
12	anonymous	Moving objects
13	anonymous	Portrait pictures included

Figure 26 Participant suggestions for future development of the application

6 CONCLUSION

6.1 Limitations/Potential improvements

While the application successfully achieved its core goals, several limitations were identified during development and evaluation:

- Hardware gaps in evaluation: Due to logistical and time constraints, the study relied on a remote video and image heuristic model rather than direct in-headset testing. This restricted data regarding the application in a VR space and subsequent findings from that research.
- VLM performance: Running Llama-Vision 3.2 locally ensured zero costs and data privacy but introduced a constraint with reasoning capabilities. The model occasionally suffered from “attention drift”, resulting in disregard of given instructions, particularly about response length, where a more powerful cloud-based model could have mitigated this and produced better results.
- NIMA limitations: Due to NIMA being trained on realistic images it consistently scored the photos of the Low-Poly environments low, typically in the range of 3 - 5.5. This necessitated scaling the output for the user front end as seeing every image score in this range could be demoralizing or frustrating, making the final score very volatile. This could be solved by a longer testing period and gaining the full range of image scores, but would be better solved by a model fine-tuned, or specifically trained, for the low-poly images.

6.2 Future work

To further develop this field future would could focus on these areas:

- Model fine-tuning: Tuning a smaller local VLM model on a domain-specific data set of photography critiques would ensure conciseness and eradicate attention drift while also potentially eliminating the need for a dedicated IQA model further improving performance.
- Dynamic level design: Expanding environment selection to cover a wider variety of photography opportunities to include moving subjects, wildlife, and more complex lighting scenarios to provide opportunity for more realistic learning experiences.
- Game loop improvements: Creating dynamic options to scale both available settings and AI feedback complexity to research the effectiveness of restricting the player based on skill to improve overall learnability for their level.

6.3 Closing remarks

This project successfully demonstrates that integrating Virtual reality with a hybrid AI-pipeline creates a viable, effective and entirely unsupervised learning environment for learning creative and technical photography principles. By mirroring real-world camera functions and grounding VLM generated feedback with deterministic data, the system provides an accurate, low-risk playground for skill acquisition that also promotes sim-to-real transfer learning. Ultimately the high level of participant confidence and enthusiasm around the application recorded in the study confirms that VR and AI can reliably bridge the gap between virtual training and real-world application. The study also confirms that AI feedback does influence player decision making by providing clear, actionable advice that the user can make use of to improve their next photo.

The implementation of a native Unity Render Graph architecture fully satisfied the performance success criteria. By handling image capture asynchronously and bypassing main-thread CPU stalls, the application successfully sustained the consistent 90 FPS threshold necessary to eliminate user discomfort and simulation sickness in a virtual environment while also capturing high fidelity accurate images.

Overall, this project has been a resounding success; while having its limitations it answers the questions we set out to answer and proving that there is reason to continue developing applications like these and continue researching this subject to find improvements.

Word count (main body of the report): 6265

REFERENCES

Bessa, D., & Madureira, M. (2025, December). Teaching Virtual Production in Higher Education: Lessons Drawn from the VPSN Research Project. In Proceedings of CAGA24: The Conference of AG Animation 2024.

Bosse, S., Maniry, D., Müller, K. R., Wiegand, T., & Samek, W. (2017). Deep neural networks for no-reference and full-reference image quality assessment. IEEE Transactions on image processing, 27(1), 206-219.

Chen, Y., Han, J., Bai, T., Tong, S., Kokkinos, F., & Torr, P. (2025). From pixels to feelings: Aligning mllms with human cognitive perception of images. arXiv preprint arXiv:2511.22805.

Christopoulos, A., Pellas, N., Bin Qushem, U., & Laakso, M. J. (2023). Comparing the effectiveness of video and stereoscopic 360 virtual reality-supported instruction in high school biology courses. British Journal of Educational Technology, 54(4), 987-1005.

Dourish, P. (2001). *Where the action is: the foundations of embodied interaction*. MIT press.

Fu, Y., & Zhang, N. (2021, April). A study on the practicality of virtual reality technology in the teaching of photography in colleges and universities. In 2021 International Conference on Internet, Education and Information Technology (IEIT) (pp. 481-486). IEEE.

Hays, R. T., & Singer, M. J. (2012). *Simulation fidelity in training system design: Bridging the gap between reality and training*. Springer Science & Business Media. (Original work published 1989).

Hou, L., Yu, C. P., & Samaras, D. (2016). Squared earth mover's distance-based loss for training deep neural networks. *arXiv preprint arXiv:1611.05916*.

Hou, W., Gao, X., Tao, D., & Li, X. (2014). Blind image quality assessment via deep learning. IEEE transactions on neural networks and learning systems, 26(6), 1275-1286.

Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., ... & Adam, H. (2017). Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*.

Huang, Y., Yuan, Q., Sheng, X., Yang, Z., Wu, H., Chen, P., ... & Lin, W. (2024). Aesbench: An expert benchmark for multimodal large language models on image aesthetics perception. arXiv preprint arXiv:2401.08276.

Hutson, J., & Lang, M. (2023). Content creation or interpolation: AI generative digital art in the classroom. Metaverse, 4(1).

Jones, J. A., Swan, J. E., Singh, G., Kolstad, E., & Ellis, S. R. (2008, August). The effects of virtual reality, augmented reality, and motion parallax on egocentric depth perception. In Proceedings of the 5th symposium on Applied perception in graphics and visualization (pp. 9-14).

Liu, B., Hu, Y., Jin, S., Dou, S., Shi, G., Shao, J., ... & Huang, X. (2025). Unlocking the Essence of Beauty: Advanced Aesthetic Reasoning with Relative-Absolute Policy Optimization. arXiv preprint arXiv:2509.21871.

Ma, X., He, S., Ming, A., Zhong, H., & Ma, H. (2026, March). Regression over Classification: Assessing Image Aesthetics via Multimodal Large Language Models. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 40, No. 10, pp. 7827-7835).

Meade, M. E., Meade, J. G., Sauzeon, H., & Fernandes, M. A. (2019). Active navigation in virtual environments benefits spatial memory in older adults. Brain sciences, 9(3), 47.

Méndez-Fernández, R., Sosa-Fernández, R. D. P., Fernández-Ledo, F., & Castelló-Mayo, E. (2025, September). Democratization of Virtual Production: Usability Analysis of Three Solutions with Different Levels of Complexity: Professional, Educational and Cloud-Based. In Informatics (Vol. 12, No. 4, p. 104). MDPI.

Meta. (2024). Llama 3.2: Technical overview. Meta AI. https://ai.meta.com/blog/llama-3-2-connect-2024-vision-edge-mobile-devices/

Nielsen, J. (1994, April). Usability inspection methods. In *Conference companion on Human factors in computing systems* (pp. 413-414).

Norman, Donald A. "Affordance, conventions, and design." *interactions* 6.3 (1999): 38-43.
Piaget, J. (2013). *Play, dreams and imitation in childhood* (C. Gattegno & F. M. Hodgson, Trans.). Routledge. (Original work published 1951).

Plass, J. L., Homer, B. D., & Kinzer, C. K. (2015). Foundations of game-based learning. Educational psychologist, 50(4), 258-283.

Qi, D., Zhao, H., Shi, J., Jenni, S., Fan, Y., Dernoncourt, F., ... & Li, S. (2025). The Photographer's Eye: Teaching Multimodal Large Language Models to See, and Critique Like Photographers. In Proceedings of the Computer Vision and Pattern Recognition Conference (pp. 24807-24816).

Ragan, E. D., Bowman, D. A., Kopper, R., Stinson, C., Scerbo, S., & McMahan, R. P. (2015). Effects of field of view and visual complexity on virtual reality training effectiveness for a visual scanning task. IEEE transactions on visualization and computer graphics, 21(7), 794-807.

Rose, F. D., Attree, E. A., Brooks, B. M., Parslow, D. M., & Penn, P. R. (2000). Training in virtual environments: transfer to real world tasks and equivalence to real task training. *Ergonomics*, *43*(4), 494-511.

Talebi, H., & Milanfar, P. (2018). NIMA: Neural image assessment. *IEEE transactions on image processing*, *27*(8), 3998-4011.

Trottnow, J., Götz, K., Seibert, S., Spielmann, S., Helzle, V., Einabadi, F., ... & Grau, O. (2015, November). Intuitive virtual production tools for set and light editing. In Proceedings of the 12th European Conference on Visual Media Production (pp. 1-8).

Wang, Z., Bovik, A. C., Sheikh, H. R., & Simoncelli, E. P. (2004). Image quality assessment: from error visibility to structural similarity. IEEE transactions on image processing, 13(4), 600-612.

APPENDIX A

Project Proposal

Bournemouth University

School of Computing & Engineering

2025-26 Academic Year Individual Masters Project

Project Proposal Form

Please refer to the **Project Handbook Section 4** when completing this form. Note that your proposal should be your own original work and you must cite sources in line with university guidance on **referencing and plagiarism**[1].

Degree Title: **MSc Human-Centred Artificial Intelligence for Game Development**	**Student's Name:** Josh Eyres
	Supervisor's Name: Fred Charles
	Project Title/Area: Human-Centred AI Evaluation and Feedback for Virtual Photography in VR Games

Section 1: Project Overview

1.1 Problem definition - use one sentence to summarise the problem:

Existing virtual photography and photo-mode tools rarely provide feedback that supports player creativity, helps them develops new skills, or refines prior knowledge — particularly within VR environments.

1.2 Project description - briefly explain your project:

This project will develop an AI-powered photo evaluation and feedback system designed to support player creativity and improve photography skills within a simulated VR environment. Players will be able to enter a virtual world, capture photos, and receive tailored, AI-generated feedback on elements such as composition, lighting, and framing. The goal is to explore how intelligent feedback can enhance creative learning and skill development in immersive spaces.

1.3 Background - please provide brief background information, e.g., client, problem domain, and make reference to the literature (minimum 4-5 sources):

[1] https://libguides.bournemouth.ac.uk/study-skills-referencing-plagiarism

Bournemouth University

School of Computing & Engineering

2025-26 Academic Year Individual Masters Project

Simulation-based training is a proven method for teaching complex skills, with fields like medicine showing it's viable for " significantly enhancing healthcare professionals' learning experience and clinical competency." (Elendu et al., 2024). However, its application to creative skills remains limited, and traditional "serious games" often suffer from a lack of intrinsic motivation, failing to engage users. On the other hand, "virtual photography" has developed as a popular, starting as simple screenshots and becoming a highly engaging activity for millions of players, driven by "uncanny" graphical fidelity (Durusoy, 2018). Current photo modes are "toys, not tools"—superficial, engaging, and lacking the feedback on core principles, required for real-world "transfer of learning" (Cook et al., 2007). A solution lies in "game-based learning" (GBL), which is defined by its focus on balancing educational content with the desire to prioritise engaging gameplay (Plass et al., 2015).

This project will bridge that gap. It aims to develop and evaluate a GBL artefact that harmonises education and entertainment. By embedding a VR camera system that mimics a real-world digital camera and AI-driven feedback into an intrinsically motivating game, this project will test the hypothesis that a "fun" and engaging experience can be an effective tool for teaching real-world photography skills.

1.4 Research Questions

How effective is AI feedback and evaluation for improving players creative and technical ability in photography?
How useful is the feedback information from the AI system?
How translatable are virtual photography skills to the real world?

1.5 Aim and objectives – what are the aims and objectives of your project? should be <u>specific</u> and <u>measurable</u>:

The aim of this project is to develop an AI-powered virtual reality photography system that provides players with evaluation and feedback to help them improve the quality and composition of their photos while supporting their creativity.
Objectives:

1. Design and develop a VR environment and photography system
2. Develop and implement an AI-driven feedback and evaluation tool within the game that is capable of assessing photos on a defined set of criteria (Composition, Lighting, etc)
3. Provide players with clear actionable feedback that supports and enhances player creativity rather than constraining it
4. Conduct user testing to evaluate the tools effectiveness, usability and impact on creative learning
5. Analyse results to determine how the feedback affects player engagement and skill development

Bournemouth University

School of Computing & Engineering

2025-26 Academic Year Individual Masters Project

Section 2: Artefact

2.1 What is the artefact that you intend to produce?

A Virtual Reality game with at least one world and a photography system in place, with an AI-powered evaluation and feedback tool for the produced photos.

2.2 How is your artefact actionable (i.e., routes to implementation and exploitation in the technology domain)?

The artefact will be implemented as a functional VR prototype built in the unity game engine and applying an ai evaluation model (Likely some sort of computer vision) to analyse photos and then a second model will be used to generate the tailored feedback for it.
This artefact is actionable both within the games industry and the virtual production industry where similar systems could support player creativity and learning. The projects framework could also be extended into a full-fledged photography learning platform with things like digital tutorials as well. Finally, this project could be used as a foundation for further research into AI supporting creativity and human-AI interaction within immersive environments.

Section 3: Evaluation

3.1 How are you going to evaluate your project artefact?

The artifact will be evaluated through user studies, questionnaires, data/metric analysis and a pre-test/post-test group evaluation, where there will be two groups Group A will take some photos in the real world, read an article or watch a tutorial on photography principles then take some more photos in the real world, Group B will do the same except they will play the game instead of reading/tutorial. The results will then be analysed to see how transferrable the skills from the game are to the real world.

3.2 How does this project relate to your MSc Programme and your degree title outcomes?

This project links to my programme because it utilises skills learnt from a multitude of different units within my course, Human-Computer interaction for the user studies, Mixed reality for the VR element and Ludonarrative design for embedding the AI in some form within a game.

3.3 What are the risks in this project and how are you going to manage them?

The main risk in this project is injury or “VR sickness” while using a VR system, this will be managed by performing tests in a controlled environment with plenty of space in short periods of time, there will also be comfort features such as being able to choose between continuous movement or teleporting/snap turning to help with the motion sickness.

Bournemouth University

School of Computing & Engineering

2025-26 Academic Year Individual Masters Project

Section 4: References

4.1 Please provide references if you have used any.

Cook, D. J., Holder, L. B., & Youngblood, G. M. (2007). Graph-based analysis of human transfer learning using a game testbed. *IEEE Transactions on Knowledge and Data Engineering*, *19*(11), 1465-1478.

Elendu, C., Amaechi, D. C., Okatta, A. U., Amaechi, E. C., Elendu, T. C., Ezeh, C. P., & Elendu, I. D. (2024). The impact of simulation-based training in medical education: A review. *Medicine*, *103*(27), e38813.

Durusoy, M. (2018). In-Game Photography: Creating New Realities through Video Game Photography. *Membrana-Journal of Photography, Theory & Visual Culture*, *3*(1).

Plass, J. L., Homer, B. D., & Kinzer, C. K. (2015). Foundations of game-based learning. *Educational psychologist*, *50*(4), 258-283.

Section 5: Academic Practice and Ethics

5.1 I confirm that I have read and am familiar with and understand the University guidance on academic integrity and plagiarism.	**Yes**
5.2 I confirm that I have read and am familiar with and understand the University guidance on research ethics.	**Yes**

Note: Please complete the ethics form **only** after your supervisor has approved the proposal.

School of Computing & Engineering

2025-26 Academic Year Individual Masters Project

Section 6: Proposed Plan (please attach your Gantt chart below)

The Gantt chart below, shows the outline of the project with key dates for the project and splits the process into 7 distinct phases, additionally I am making use of Clickup, a website that this Gantt chart was made on that also has, a calendar, 'Trello-like' board, a list view and table view of all the tasks I have devised to allow me to keep on top of the project as I progress through the development of it.

School of Computing & Engineering

2025-26 Academic Year Individual Masters Project

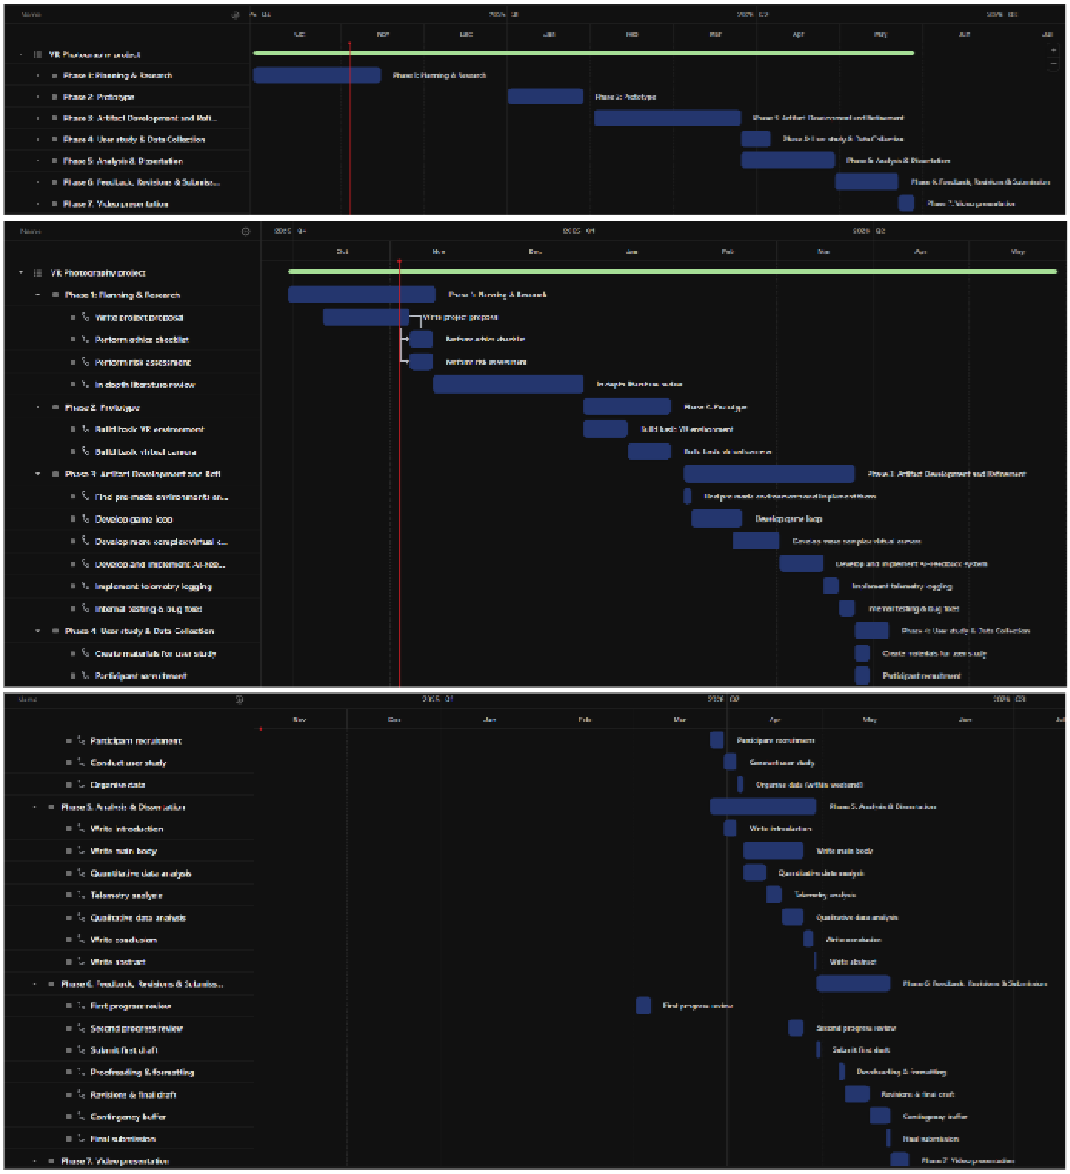

Bournemouth University

School of Computing & Engineering

2025-26 Academic Year Individual Masters Project

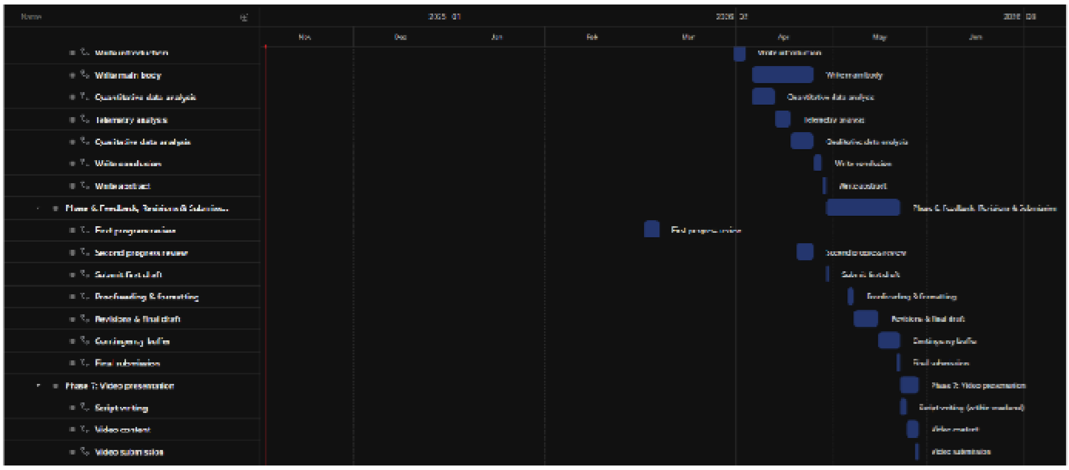

Bournemouth University

School of Computing & Engineering

2025-26 Academic Year Individual Masters Project

Name	Assignee	Due date	Priority	Status	Comments
Phase 1: Planning & Research 4	JE	11/14/25		TO DO	
Write project proposal	JE	Thu		TO DO	
Perform ethics checklist	JE	11/13/25		TO DO	
Perform risk assessment	JE	11/13/25		TO DO	
In depth literature review	JE	12/31/25		TO DO	
Phase 2: Prototype 2	JE	1/28/26		TO DO	
Build basic VR environment	JE	1/14/26		TO DO	
Build basic virtual camera	JE	1/20/26		TO DO	
Phase 3: Artifact Development and Ref... 6	JE	3/25/26		TO DO	
Find pre-made environments and impl...	JE	2/3/26		TO DO	
Develop game loop	JE	2/18/26		TO DO	
Develop more complex virtual camera	JE	3/1/26		TO DO	
Develop and Implement AI-Feedback s...	JE	3/15/26		TO DO	
Implement telemetry logging	JE	3/20/26		TO DO	
Internal testing & bug fixes	JE	3/25/26		TO DO	
Phase 4: User study & Data Collection 4	JE	4/5/26		TO DO	
Create materials for user study	JE	3/30/26		TO DO	
Participant recruitment	JE	3/30/26		TO DO	
Conduct user study	JE	4/3/26		TO DO	
Organise data	JE	4/5/26		TO DO	

Bournemouth University

School of Computing & Engineering

2025-26 Academic Year Individual Masters Project

Name	Assignee	Due date	Priority	Status	Comments
Phase 5: Analysis & Dissertation 7	JE	4/28/26		TO DO	
Write introduction	JE	4/3/26		TO DO	
Write main body	JE	4/24/26		TO DO	
Quantitative data analysis	JE	4/12/26		TO DO	
Telemetry analysis	JE	4/17/26		TO DO	
Qualitative data analysis	JE	4/24/26		TO DO	
Write conclusion	JE	4/27/26		TO DO	
Write abstract	JE	4/28/26		TO DO	
Phase 6: Feedback, Revisions & Submi... 7	JE	5/22/26		TO DO	
First progress review		3/6/26		TO DO	
Second progress review		4/24/26		TO DO	
Submit first draft	IF	4/29/26		TO DO	
Proofreading & formatting	JE	5/7/26		TO DO	
Revisions & final draft	JE	5/15/26		TO DO	
Contingency buffer	JE	5/22/26		TO DO	
Final submission	JE	5/22/26		TO DO	
Phase 7: Video presentation 3	JE	5/28/26		TO DO	
Script writing	JE	5/24/26		TO DO	
Video content	JE	5/28/26		TO DO	
Video submission	JE	5/28/26		TO DO	

APPENDIX B

Approved Ethics Checklist

Research Ethics Checklist

About Your Checklist	
Ethics ID	69994
Date Created	26/02/2026 16:36:11
Status	Approved
Date Approved	11/05/2026 14:26:07
Risk	Low

Researcher Details	
Name	Josh Eyres
Faculty	Faculty of Media, Science and Technology
Status	Postgraduate Taught (Masters, MA, MSc, MBA, LLM)
Course	MSc Human Centred Artificial Intelligence for Games Development

Project Details	
Title	Virtual Reality Photography Simulator
Start Date of Project	02/02/2026
End Date of Project	22/05/2026
Proposed Start Date of Data Collection	11/05/2026
Supervisor	Fred Charles
Approver	Fred Charles
Summary - no more than 600 words (including detail on background methodology, sample, outcomes, etc.)	
A photography simulator making use of VR and AI to create an unassisted learning environment where users have access to common settings and different environments to develop photography skills without the need to find locations or buy expensive gear, while getting personalised feedback based on the images they take through AI aesthetic/technical image asessment and LLM models.	

Filter Question: Does your study involve Human Participants?

Participants	
Describe the number of participants and specify any inclusion/exclusion criteria to be used	
Ideally the study will make use of at least 10 participants but anyone will be welcome to take part, various levels of photography knowledge will be aimed at, from complete novice to professional..	
Do your participants include minors (under 16)?	No

Are your participants considered adults who are competent to give consent but considered vulnerable?	No
Is a Disclosure and Barring Service (DBS) check required for the research activity?	No

Recruitment

Please provide details on intended recruitment methods, include copies of any advertisements.	
Due to the time constraints convenience sampling will be used for this project with some social media messages to gain any extra participants	
Do you need a Gatekeeper to access your participants?	No

Data Collection Activity

Will the research involve questionnaire/online survey? If yes, don't forget to attach a copy of the questionnaire/survey or sample of questions.	Yes
How do you intend to distribute the questionnaire?	
online	
If online, do you intend to use a survey company to host and collect responses?	Yes
If yes, please provide details of survey company.	
Microsoft Forms	
Will the research involve interviews? If Yes, don't forget to attach a copy of the interview questions or sample of questions	No
Will the research involve a focus group? If yes, don't forget to attach a copy of the focus group questions or sample of questions.	No
Will the research involve the collection of audio recordings?	No
Will your research involve the collection of photographic materials?	No
Will your research involve the collection of video materials/film?	No
Will the study involve discussions of sensitive topics (e.g. sexual activity, drug use, criminal activity)?	No
Will any drugs, placebos or other substances (e.g. food substances, vitamins) be administered to the participants?	No
Will the study involve invasive, intrusive or potential harmful procedures of any kind?	No
Could your research induce psychological stress or anxiety, cause harm or have negative consequences for the participants or researchers (beyond the risks encountered in normal life)?	No
Will your research involve prolonged or repetitive testing?	No
What are the potential adverse consequences for research participants and how will you minimise them?	
There will be minimal risks to the participants due to the study being an online questionnaire	

Consent

Describe the process that you will be using to obtain valid consent for participation in the research activities. If consent is not to be obtained explain why.	
A consent form will be provided at the start of the questionnaire.	
Do your participants include adults who lack/may lack capacity to give consent (at any point in the study)?	No
Will it be necessary for participants to take part in your study without their knowledge and consent?	No

Participant Withdrawal

At what point and how will it be possible for participants to exercise their rights to withdraw from the study?	
Participants will be allowed to withdraw from the study at any time by simply letting me, the researcher, know, verbally or via email	
If a participant withdraws from the study, what will be done with their data?	
Any data that is linked to the participant will not be saved if they choose to withdraw, however any data from the questionnaire will not be identifiable after they submit their response so if that is the case it will not be able to be destroyed	

Participant Compensation

Will participants receive financial compensation (or course credits) for their participation?	No
Will financial or other inducements (other than reasonable expenses) be offered to participants?	No

Research Data

Will identifiable personal information be collected, i.e. at an individualised level in a form that identifies or could enable identification of the participant?	No
Will research outputs include any identifiable personal information i.e. data at an individualised level in a form which identifies or could enable identification of the individual?	No

Storage, Access and Disposal of Research Data

Where will your research data be stored and who will have access during and after the study has finished.	
The research data will be stored on the cloud via microsoft forms, i am the only one who will have access to it	
Once your project completes, will your dataset be added to an appropriate research data repository such as BORDaR, BU's Data Repository?	No
Please explain why you do not intend to deposit your research data on BORDaR? E.g. do you intend to deposit your research data in another data repository (discipline or funder specific)? If so, please provide details.	
The research data will not be stored publically only the findings of the paper	

Final Review

Are there any other ethical considerations relating to your project which have not been covered above?	No

Risk Assessment

Have you undertaken an appropriate Risk Assessment?	Yes

Attached documents
VR Photography Visual Heuristics Questionnaire.pdf - attached on 11/05/2026 14:23:07

APPENDIX C

Large File Contents

/VR Photography Submission - HCAIGD Individual project - Josh Eyres s5304176

- /PythonBackend
  - Contains the Flask local web server, AI/Metric scripts, and requirements manifest for managing the Ollama VLM pipeline.
  - /models
    - Contains aesthetic and technical model weights for Google's Neural Image Assessment model
  - /src
    - Contains code required to run NIMA models
  - /requirements.txt
    - Contains python project dependencies
- /SubmissionBuild
  - Contains the compiled standalone executable, assets, data, and graphics binaries compiled via the Unity 6 Render Graph pipeline.
- /ReadME.txt
  - Provides system prerequisites, step-by-step deployment instructions, and the active public source code repository link.
- /Run.bat
  - The initialization script configured to silently boot background processes and launch the primary VR execution thread.

APPENDIX D

System Requirements and Target Environment
As a prototype developed and validated within a localized research environment, the application's performance has been verified verified using a fixed reference hardware and software configuration.

1. Reference Hardware Specification (Validated Test Environment)
The following hardware configuration was utilized to build and test the application, achieving stable rendering and execution loops:

Processor (CPU): AMD Ryzen 9 5900X (or equivalent high-performance multi-threaded processor)

Graphics Card (GPU): NVIDIA GeForce RTX 3080ti (10GB VRAM minimum recommended to sustain concurrent execution of the Unity 6 Render Graph pipeline and local vision model inference)

System Memory (RAM): 32 GB DDR4

Storage: A minimum of 12 GB of free space on a Solid-State Drive (SSD), distributed as follows:

Compiled Application Artefact: ~2 GB (for the executable game build and assets).

Local AI Infrastructure: ~10GB (allocated for the Llama 3.2 Vision model weights and Ollama environment overhead).

Source Project Files: ~17 GB (required only for full source-code deployment, development caching, and asset versioning review).)

2. Target VR Hardware Platform
Standalone / Link Device: Meta Quest 3

Connection Interface: Meta Quest Link (Oculus Link) via a high-bandwidth USB-C 3.0 cable, or a dedicated Wi-Fi 6E wireless network layer capable of handling high-bitrate spatial tracking and stereoscopic frame streaming.

3. Runtime Software Dependencies
To successfully execute the autonomous feedback loop and the underlying local artificial intelligence pipeline, the host system must have the following software environment pre-configured:

Ollama Runtime Environment: Installation of the Ollama service on the local loopback address (localhost:11434).

Model Manifests: Local deployment of the Llama 3.2 Vision weights (llama3.2-vision) via the Ollama registry to process multi-modal text and image evaluation streams.

Operating System: Windows 11 (64-bit) configured with up-to-date graphics drivers

APPENDIX E

Questionnaire full questions and results

Responses	Average Time	Duration
19	13:37	11 Days

1. Study Title: Are VR and AI effective tools for creating an unsupervised learning environment for photography?
   Principal Investigator: Josh Eyres
   Affiliation: Bournemouth University
   Contact Information: s5304176@bournemouth.ac.uk

   You are invited to participate in a study exploring the use of VR and AI as a photography tool to develop new skills and refine pri or knowledge. Your role will be to evaluate the user interface (UI), user experience (UX), functionality and realism of the applicatio n. This study is being conducted by Josh Eyres, a student at Bournemouth University. Your participation is entirely voluntary, and you have the right to withdraw at any time with no penalty.

   Purpose of the Study: The goal is to determine if a VR environment combined with AI-generated feedback can effectively facilitat e "unsupervised" learning, allowing users to develop real-world photography skills without a human instructor.

   Study Procedures: If you agree to participate in this study, you will be asked to analyse screenshots and video from the VR photo graphy application in order to critique the system against standard usability principles, assess the functionality and also realism o f the application.

   Potential Risks and Benefits:
   Risks:
   - General Fatigue: As with any computer-based task, there is a minor risk of digital eye strain from viewing the monitor.
   - Minimal Risk: There are no physical risks associated with VR hardware (motion sickness or physical injury) as no headset is used.

   Benefits:
   - Understanding how VR can be used as a Virtual learning environment in applications where it may be too expensive, impossibl e to experience scenarios etc.
   - Understanding how AI can be used as a guide as opposed to being directly supervised
   - Interface Optimization: Your analysis will contribute to designing more intuitive AI-driven learning tools.

   Confidentiality: These questionnaires will be anonymous and therefore confidential. The data from these questionnaires will be p assword protected and only accessed by the researcher.

   Right to Withdraw: Your participation in this study is entirely voluntary and you can withdraw at any time. However data from the questionnaire is anonymous so after it has has been submitted the data will no longer be able to be destroyed.

   Contact Information: Feel free to contact Josh Eyres at s5304176@bournemouth.ac.uk if you have any questions about the stud y.

   Consent: By agreeing to participate in this study you are confirming that you have read and understood this information sheet a nd you consent to participate in this study under the conditions described.

More details

- I consent to the data from this questionnaire being used for research purposes and subsequent reports... 19

2. What is your gender?

More details

Female	8
Male	10
Non-binary	0
Prefer not to say	0
Other	1

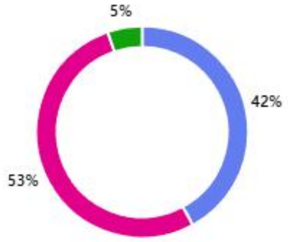

3. How old are you?

More details

Under 18	0
18-24	7
25-34	4
35-44	2
45-54	3
55+	3
Prefer not to say	0

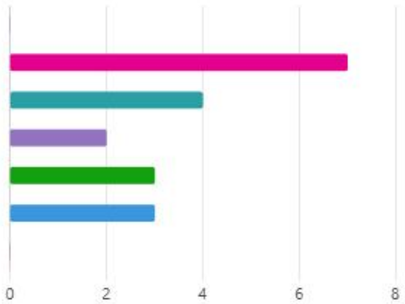

4. What is your highest level of education completed?

More details

Secondary school / High School	5
Undergraduate degree	8
Postgraduate degree (Master's, PhD, etc.)	2
Other	4

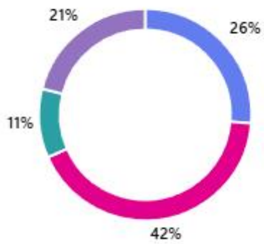

5. What is your current employment status?

More details

Student	1
Parent	1
Employed full-time	12
Employed part-time	2
Retired	1
Prefer not to say	2

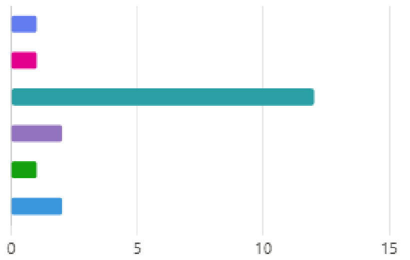

6. What level would you consider your photography ability?

More details

Beginner	11
Intermediate	6
Advanced	2
Expert	0

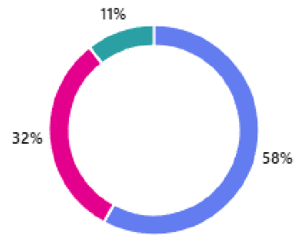

7. How familiar are you with technical photography concepts? (Rule of thirds, Exposure, etc)

More details

Very Familiar	2
Somewhat familiar	8
Heard of them	4
Never heard of them	5

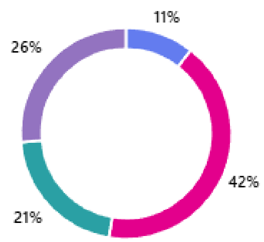

8. How frequently do you use VR

More details

- Everyday 0
- A few time a week 1
- Once a week 1
- A few times a month 4
- Once a month 2
- Never 11

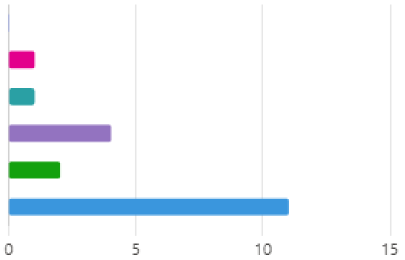

9. On what type of device are you viewing the image sets for this study?

More details

- Smartphone 11
- Tablet 2
- Laptop / Desktop Monitor 6
- Other 0

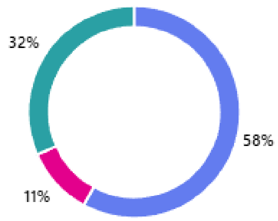

10. Please answer the following statements based on the video you watched

More details

Strongly Disagree · Disagree · Neutral · Agree · Strongly agree

- The camera interactions shown in the video accurately mirror a real-world camera
- The user interface menus shown appear intuitive and logically organized
- The environments shown look visually interesting enough to encourage creative decision making
- The environments have enough variety to try taking photos of lots of different subjects
- The user can clearly understand the current settings selected on the camera
- The photo preview on the camera screen updates according to the available settings
- I believe a complete beginner would be able to use this system without any external human help

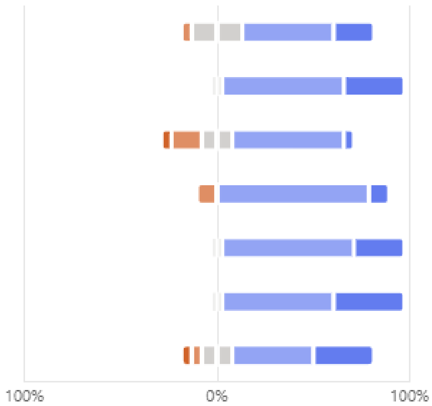

11. Did you notice any specific usability issues with the demo shown in the video? (If possible please also provide a video timestamp (e.g. @1:20) of the issue you noticed in your answer)

More details

7
Responses

Latest Responses
"No"
"No"
...

3 respondents (43%) answered camera for this question.

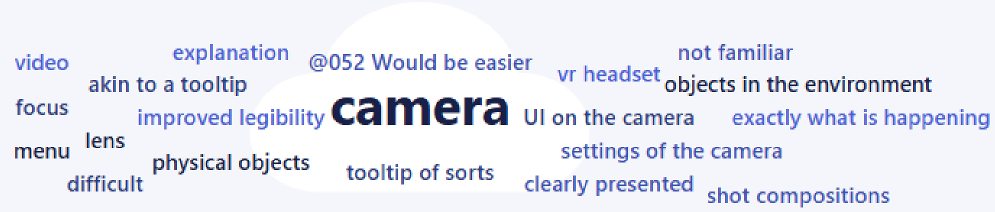

12. If yes to the previous question, how much do you believe they would hinder a beginner's ability to learn? (1 = Minor annoyance, 5 = Impossible to get past)

More details

1.50
Average Rating

1 2 3 4 5

13. Please randomly select one of the options below to choose an image set

More details

1	3
2	2
3	3
4	8
5	2
6	1

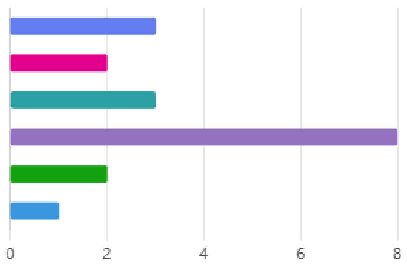

14. Please answer the following statements based on the image set provided

More details

Strongly Disagree · Disagree · Neutral · Agree · Strongly agree

The AI feedback is easy to understand

The technical feedback recieved accurately reflects the given photograph

The AI provides specific, actionable advice that would help a user improve their next shot

The scoring system (i.e. Aesthetic/Technical) provides a clear incentive for a learner to improve

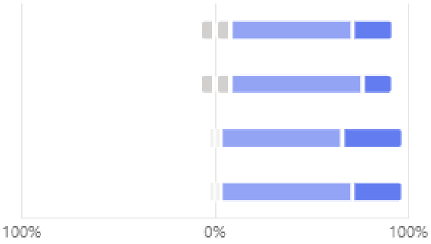

15. Given the feedback you have seen, was there anything obviously wrong or unhelpful with the AI generated feedback?

More details

10
Responses

Latest Responses
"No"
"No"
...

3 respondents (30%) answered feedback for this question.

16. Are the aesthetic scores you have seen accurate to what you would have scored the image out of 10?

More details

Yes	13
No	1
Maybe	5

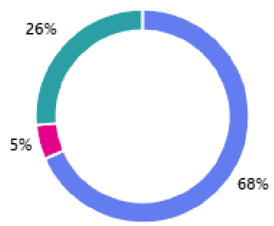

17. Are the technical scores you have seen accurate to what you would have scored the image out of 10?

More details

Yes	13
No	2
Maybe	4

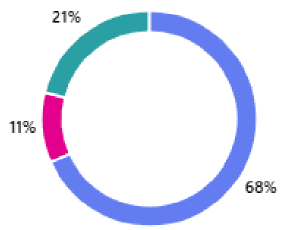

18. Do you believe the Exposure, Focus and Composition scores are useful?

More details

Yes	18
No	0
Maybe	1

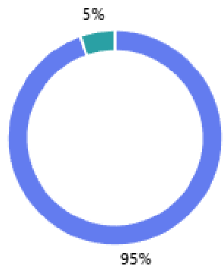

19. Please explain the answer to the previous question where possible

More details

14

Responses

Latest Responses

"No"

"Yes the scores are useful so you can improve the quality of your photos"

...

7 respondents (50%) answered score for this question.

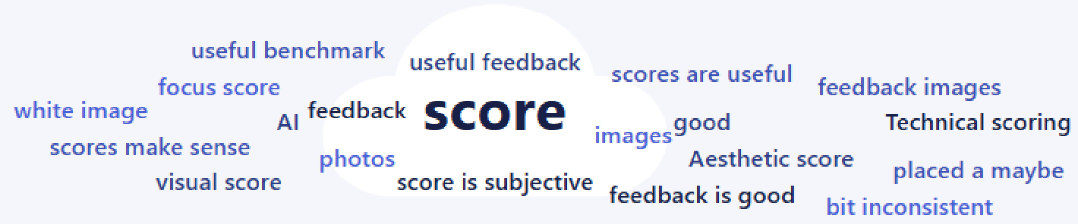

20. Are there any parts of the feedback you felt weren't as useful or entirely unneccesary?

More details

19
Responses

Latest Responses
"No"
"No"
"No"
...

2 respondents (11%) answered aesthetics for this question.

explanation
opinion
technical side of things
feedback is fair
feedback question
suggestion
stated previously
scores
aesthetics
technical
scoring is less helpful
new learner
specific scores
things or aesthetics
AI feedback
maybe would need
personal
photographer
subjective
Composition scoring

21. Based on your observation, do you believe that this application would be helpful for learning photography?

More details

Yes	18
No	0
Maybe	1

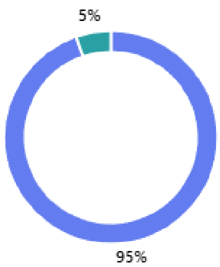

22. How confident would you feel applying the technical principles shown in the application to real life?

More details

Extremely confident	6
Somewhat confident	12
Neutral	0
Somewhat not confident	1
Extremely not confident	0

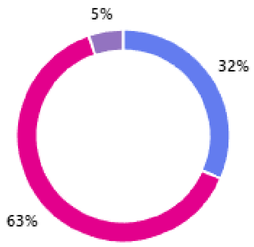

23. Do you believe that having a virtual environment to experiment with photography is a useful tool to have for aspiring photographers?

More details

Yes	18
No	0
Maybe	1

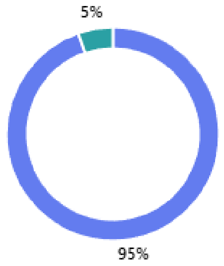

24. If this tool were made available as a standard educational program for photographers what would you like to see added or improved?

More details

14

Responses

Latest Responses

"Nothing"

"Portrait pictures included"

...

3 respondents (21%) answered objects for this question.

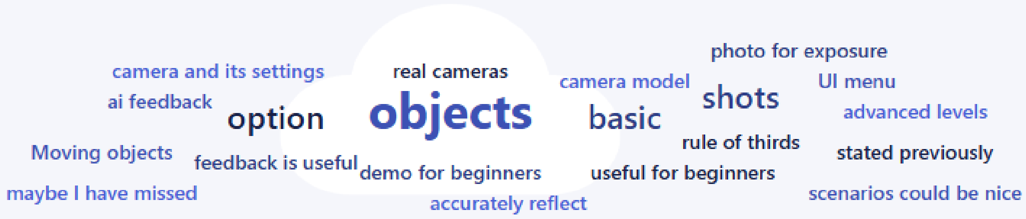

25. Are there any other visual or interaction elements you observed that did not align with your expectations of a professional photography tool?

More details

9

Responses

Latest Responses

"No"

"No"

...

5 respondents (56%) answered No for this question.

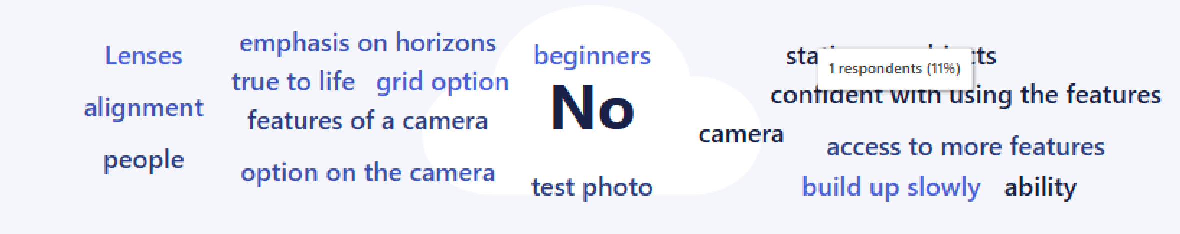

26. Do you have any extra comments about the transferrability of skills learnt not previously covered?

More details

7

Responses

Latest Responses

"No"

"No"

...

5 respondents (71%) answered No for this question.

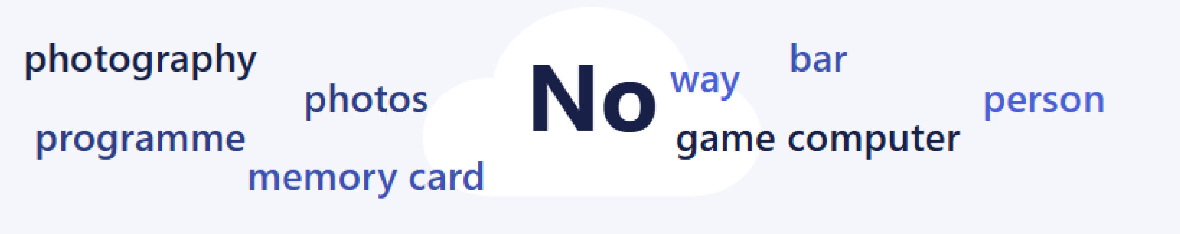

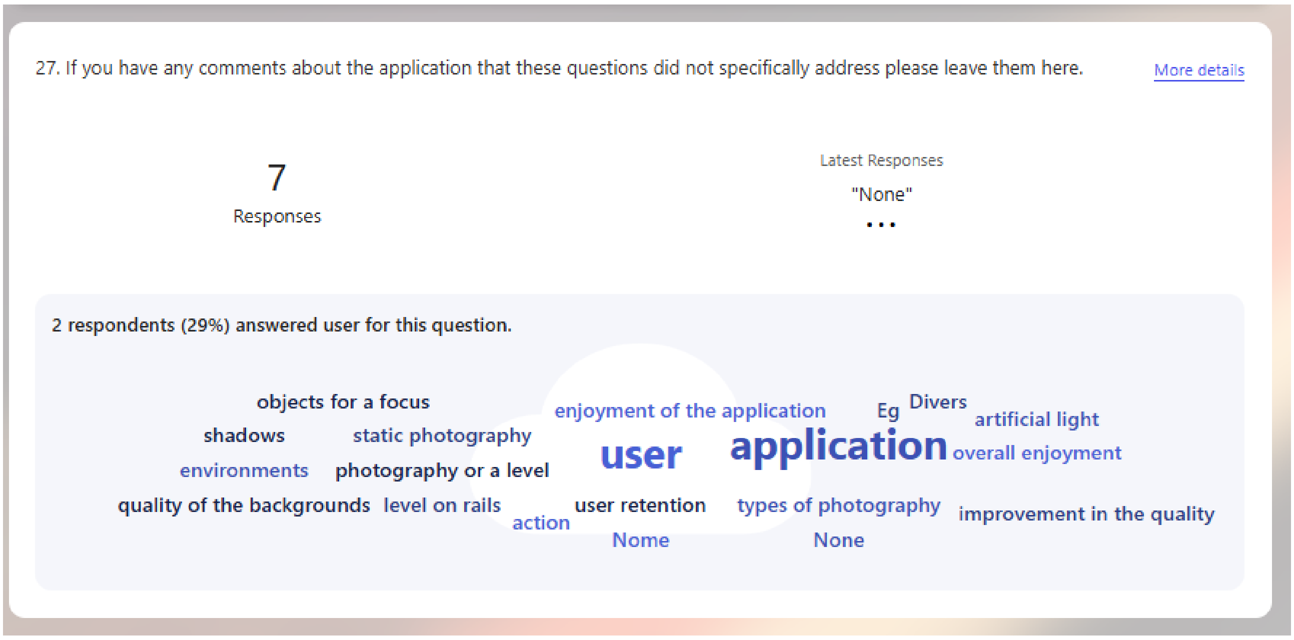
27. If you have any comments about the application that these questions did not specifically address please leave them here.
More details
7
Responses
Latest Responses
"None"
...
2 respondents (29%) answered user for this question.
objects for a focus
enjoyment of the application
Eg
Divers
artificial light
shadows
static photography
user
application
overall enjoyment
environments
photography or a level
quality of the backgrounds
level on rails
action
user retention
types of photography
improvement in the quality
Nome
None