

Automagic's Augmented Reality

Computer Graphics: Rendering, Autumn 2023

Coursework 1 of 2

Daniel Bilc, Kartic Subr

1 Overview

Automagic, a startup, has set out to develop AI for augmented reality. They have already attracted significant investment but are first looking to understand the manual workflow before committing to targeted tools for automation. They are looking to hire a creative consultant to provide an executive report. Due to the large number of applications, they have designed the following test to seek the best applicant. Your first assignment is to try to impress them!

The three skills they are looking for are: the ability to quickly learn how to use a modelling tool, documentation skills and creativity. Their test is comprised of two sections. For the first one, they want you to document the procedure for photo-realistically embedding rendered computer graphics objects within a real/natural scene. The second portion of the test involves a short demo reel that will show off your creative and technical abilities. Finally, they would like you to present your results in a short (5 minutes) presentation.

2 Specifications

2.1 Augmented reality

Goal: Augment a photograph with a photo-realistically rendered (virtual) object. You are expected to choose a unique scene and model it sufficiently to provide the illusion of augmented reality.

Constraints:

1. Lighting, reflections and shadows that are realistic (consistent with viewers' expectations);
2. Use Blender with the cycles renderer;
3. Must include a Coca-Cola can to advertise the most prominent investors of Automagic;
4. The modelled scene should be lightweight – ideally minimal for convincing augmentation; and
5. The training documentation should be concise yet comprehensive.

2.2 Animation

Goal: Model a short animation (minimum 10 seconds) to impress Automagic with your creativity and to demonstrate your ability to quickly learn a modelling tool.

Constraints:

1. The main camera must be in motion;
2. Use Blender with the cycles renderer;
3. The animation must contain non-trivial physics such as collision, wind, fire, deformation, etc.
4. Includes sophisticated materials to highlight rendering capabilities of the modelling tool.

2.3 Submission

Include three images: one without virtual object, one with virtual object and one with a real object where you placed the virtual object (even if different from the shape of the virtual object) for reference. Please see Section 4 for all deliverables.




Figure 1: Example set of deliverable images.

3 Marking scheme

This coursework is marked out of 100 and will count toward 50% of your final mark for the course.

1. Basic modelling 20
 - (a) Approximate modelling of scene geometry (5)
 - (b) Modelling scene lighting (5)
 - (c) Modelling textures (5)
 - (d) Choice of 3D model (5)
2. Visible effects 20
 - (a) Shadow of real object(s) on virtual object(s) (5)
 - (b) Texture of real object(s) reflected on virtual object(s) (5)
 - (c) Shadow of the virtual object(s) on real object(s) (5)
 - (d) Texture of virtual object(s) reflected on real object(s) (5)
3. Creative video 30
 - (a) Conception and screenplay (10)
 - (b) Modelling (5)
 - (c) Physics/animation (10)
 - (d) Rendering (5)
4. Report writing 20
 - (a) Concise presentation (5)
 - (b) Illustration (5)
 - (c) Effectiveness as a training manual (10)
5. Presentation 10
 - (a) Attendance (1)
 - (b) (coursework-relevant) Enthusiasm (2)
 - (c) Adherence to time (2)
 - (d) Props/slides (5)

This coursework is meant to be an *independent exercise* and so create your own, original scene. **Do not copy** the examples provided, and **do not collaborate** with each other to actively create your scenes. Indeed, we encourage you to discuss and brainstorm general concepts pertaining to this coursework or about the use of relevant tools.

The marks that you receive for almost every sub-item in the marking scheme will depend on the quality of implementation and not just on whether it was implemented. e.g. If the model chosen for 1.(d) is a sphere or Suzanne (default Blender monkey), you can expect 1 mark out of 5 even though you have ticked the box by choosing a model. The remaining marks are reserved for students who take extra effort. e.g. choosing a model from free online resources, making their own models, etc.

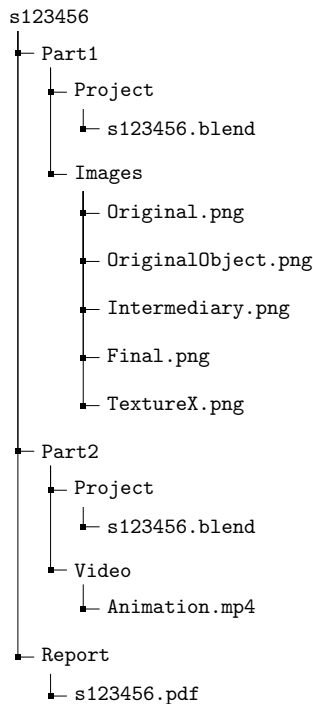
Part 2 of this assignment will be used to assess exceptionalism. That is, if you execute Part 1 of the coursework well, you can expect a mark of about 60%+. The expected time for students who have attended lectures and tutorials to attain this is about 25h. For a mark of 80%+ the video will also need to be exceptional. Please be warned that you could easily spend an inordinate amount of time tweaking each part. So, plan and prioritise!

4 Deliverables

Submit a compressed .zip with your student ID as the filename. e.g. "s123456.zip" via Learn (details to be posted soon on the course webpage). Any deviations from this submission format will incur a penalty of 10% of the final mark for all impacted sections.

4.1 Folder structure

The structure of the zip should be as follows:



Part1:

The folder **Project** will contain the project file with all the assets required to create the scene with your student ID as its name. To ensure it works on the marker's computer:

1. Open the file in Blender
2. Go to 'File' > 'External Data' > 'Pack all into .blend'
3. Save the file

The folder **Images** will contain a minimum of 5 files: *Original.png* like Figure 1a, which contains the original photo of the scene; *OriginalObject.png* like Figure 1c, which includes the original stage of the scene with a real object present where the virtual object(s) will be placed (if you are using more than one virtual object you do not need to use more than one real object to get this image); *Intermediary.png* like Figure 7a, which will contain the rendering of the image without any image editing; *Final.png* like Figure 1b, which contains the original scene with the virtual object(s) added to it; *TextureX.png* like Figure 5b, which contains the texture of the real object that you used for one of the surfaces or objects where X will start from 1. You can import as many as you want, ordering them from 1 to the last one.

Part2:

The folder **Project** will contain the project file with all the assets required to create the scene with your student ID as its name. To ensure it works on the marker's computer, follow the instructions for the blend file in part 1.

The folder **Video** will contain only one file, the video of the demo in the .mp4 format. Other popular formats, such as .mov, are allowed, but they must be viewable across platforms (Windows, MacOS or Linux machines) without special tools.

Report:

The folder **Report** will contain only one file, which will use your student ID as your name. This will be your report as a PDF. **No other format will be accepted.** Your report should have two parts: first, explain each of the seven stages of the high-level workflow at the start of Section 5 and second, explain how you made your video. Automagic might use your documentation to train incoming recruits, so please use illustrations, screenshots and example images to improve exposition. Concise and complete reports will receive more marks than complete but verbose reports. Add a limitations section where you discuss potential drawbacks of design decisions that you made to simplify the task and an acknowledgements section. There are no space constraints on the report but the lack of structure or conciseness will result in a loss of marks (under 4.(c)).

This ends the specifications for this coursework. The following sections are for guidance only, and may be ignored. You may use alternative workflows to what is suggested below. If you do so, please document them clearly.

5 Tips for augmentation

The high-level process that we recommend is:

1. plan and construct your scene carefully, taking measurements if necessary;
2. take a photograph of the scene;
3. estimate camera position and parameters relative to the scene;
4. approximately estimate and model lighting in the scene;
5. replicate crucial players virtually by modelling them (geometry and texture);
6. render virtual scene using steps 3, 4 and 5; and
7. composite the rendered image onto the photograph, with appropriate editing.

5.1 Scene

Design of the scene is crucial and dictates how easy the rest of the workflow will be. If you would like a simple scene, we recommend using objects with flat surfaces. Although cylinders, spheres or other more complex shapes will result in more spectacular effects, be warned that they will add difficulty. Try to include objects with varying gloss, colour and some with texture. Choose objects so that you maximise coverage across the marking criteria (Sec.3).

Try to simplify your lighting by the use of a spotlight. e.g. a photo taken in a dimly lit room might be easier to replicate virtually. It might be beneficial to use multiple photos in an initial study of the scene. It might be useful to grab some overhead views for easy measurement.

If possible, design a scene that you can quickly and easily assemble again if required. Alternatively, design it so it does not need to be disassembled immediately. This is not a requirement, but it might be useful if you want to tweak it later.

An example of a scene with planes, boxes and cylinders with many object properties including a mirror, is present in Figure 1a.

5.2 Camera and background

The static virtual camera will only point at one angle and from one direction to match the one you used to take the photo of the scene. You may either iteratively guess/tune camera position and parameters to align the rendered view or use programs like fSpy (see Fig 2). A quick introduction to fSpy can be seen [here](#)[1] or more in-depth [here](#)[2]. A good alignment will simplify the compositing step.

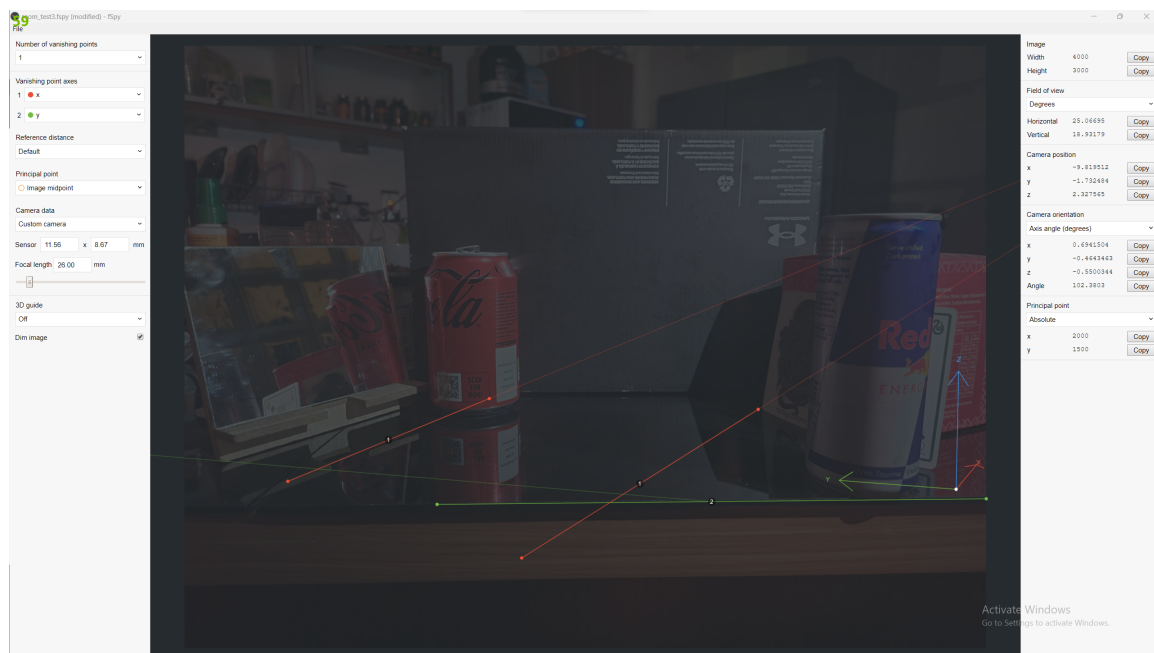


Figure 2: Background image in fSpy

5.3 Virtual object(s)

You may use one or multiple different virtual objects with different material properties to demonstrate the various effects. If using multiple virtual objects, ensure that they all interact mutually.

Keep in mind that the main goal is photo-realistic augmentation while demonstrating the visual effects in the marking scheme. The virtual objects can be either imported from an online source or created by yourself. The object's level of detail and complexity is not important, but please avoid very trivial shapes (sphere, cylinder, cubes, cuboids, torus). It is recommended to use online resources (with proper attribution) to **save time** that could instead be used to improve the final result. All assets that are used that are not created by you must be acknowledged at the end of your report in the "Acknowledgements" section.

Some examples can be seen in 3 where a Dragon[3] is used for shadows and some reflections where a Gemstone[4] is used for the reflection of real-world objects.



Figure 3: Example of virtual objects and their properties

5.4 Modeling the scene


To mimic the interaction of light between virtual objects and the environment, you will have to selectively model some real-world objects in Blender. This is especially true of the light sources. It will help to use the correct light source type as well as colour and intensity, not to mention angle and positioning, via Blender's tool for lighting. Also, place the virtual lights in the modelled scene as close to their correct location relative to the scene.

Objects that cast or receive shadows (on or from virtual objects) and objects that exchange reflections with the newly added virtual object(s) need to be modelled in Blender. The models do not have to be exact. Some examples:


1. For just shadows, you can use shadow catchers which allow you to use only the 3D model to cast shadows or receive the shadows and not have to perfectly match the object in texture and properties. This can be seen in Figure 4
2. For reflections, you will have to create/import models with real textures(example: Figure 5) and properties such that they are correctly represented in the reflections. You can still hide

them in the final image if you prefer, but you must keep their reflections. An example of reflections can be seen in Figure 6

3. For objects that have multiple sides or sub-components, you will only need a 3D representation of the part that interacts with the virtual object. For flat objects, like boxes, you can use planes to mimic just the walls that are shown in the direction of the camera.




(a) Objects used to cast and receive shadows




(b) Shadows that the objects left/picked in the scene

Figure 4: Example of using shadow catchers for creating shadow effects in the scene. red highlights objects that cast shadows and the areas where they can be seen.




(a) Virtual representation of real object with real texture




(b) Real UV map texture imported in Blender

Figure 5: Example of a real object created in Blender with imported UV mapping



(a) Objects used to cast and receive reflections



(b) Reflections that the objects left/picked in the scene

Figure 6: Example of objects and their reflection in the scene with Blue being the object having reflection and with Red the reflections themselves.

5.5 Compositing

The rendering of the scene from Blender can be a bit artificial and have quite a few effects that are not perfectly Blended in the original image. This can be seen in Figure 7a, where the mirror has a grey background, picked up from the HDRI and the floor where the Dragon is standing, which is much darker than the actual glass top. In Figure 7b, the mirror and black plane are eliminated by combining the original scene with the new render and removing the unwanted sections.

The composition was achieved using Gimp; you can learn more about it [here](#)[5]. You are not required to use it or any other image editor if the final render is already perfect.



(a) Blender rendering output



(b) Final result after image editing

Figure 7: Example of how the image can be edited to obtain a realistic image.

6 Tools

1. Blender for modelling, animation, physics and rendering -> <https://www.blender.org/>
2. fSpy for camera matching -> <https://fspy.io/>
3. Gimp for compositing -> <https://www.gimp.org/>
4. ffmpeg for making a video from rendered frames -> <https://ffmpeg.org/>

References

- [1] Drifting Walnut Productions. Quick fsptutorial- blender. <https://www.youtube.com/watch?v=iWqP1EZvUdQ>, 2023. [Video].
- [2] Bartosz Pampuch. Best way to use fspy | all options explained | full camera matching workflow with blender. <https://www.youtube.com/watch?v=daiMOYR8GS8>, 2022. [Video].
- [3] Artec Group. Dragon. <https://www.artec3d.com/3d-models/dragon>. 3D model.
- [4] Artec Group. Gemstone. <https://www.artec3d.com/3d-models/gemstone>. 3D model.
- [5] Tutorials by Zee. Gimp tutorial - how to combine / blend two pictures together. <https://www.youtube.com/watch?v=EpMhDEtnHj4>, 2020. [Video].