Automagic's Augmented Reality
Computer Graphics: Rendering, Autumn 2024
Coursework 1 of 2

Daniel Bilc, Krzysztof Grykiel, 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 (photorealistic);
2. Use Blender with the cycles renderer;
3. Must include a Duracell battery to advertise the most prominent investors of Automagic;
4. The modelled scene should be lightweight – ideally minimal for convincing augmentation; and
5. The report 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 a virtual object and one with a real object where the virtual object was placed (could be a different shape) for reference (see Sec. 4).

(a) Scene as it will be used (b) Augmented result (c) Ref. image (real object)

Figure 1: Example set of deliverable images. The real object does **not** have to resemble the virtual object, although it does in this example.(Coke was the previous sponsor, before Duracell)

3 Marking scheme

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

1. Exercises 20
   (a) Radiometry (6)
   (b) Cameras (6)
   (c) Ray Tracing 1 (4)
   (d) Monte Carlo (4)
2. AR Tool 20
   (a) Modelling (10)
   (b) Effects (10)
3. Creative video 20
   (a) Conception and screenplay (5)
   (b) Modelling (5)
   (c) Physics/animation (5)
   (d) Rendering (5)
4. Report writing 10
   (a) Conciseness (5)
   (b) Effectiveness (3)
   (c) Completeness (2)
5. Presentation 10
   (a) Attendance (1)
   (b) Keeping to allotted time (2)
   (c) Enthusiasm demonstrated (1)
   (d) Slides/props (4)
   (e) Exceptionalism (2)
6. Exceptionalism 20

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 share digital assets** with each other or actively create your scenes collaboratively. Indeed, we encourage you to discuss and brainstorm general concepts, experiences pertaining to this coursework and the use of relevant tools.

3.1 Weighted marking

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.

3.2 Showcase all effects

Furthermore bear in mind that the marking is based on **presentation** rather than simply pure effort. Make sure that all aspects of interest of your virtual image are displayed **clearly** and **prominently**. That includes: shadow of real object(s) on virtual object(s), texture of real object(s) reflected on virtual object(s), shadow of the virtual object(s) on real object(s) and texture of virtual object(s) reflected on real object(s). The given example for instance would mark poorly for the "texture of real object(s) reflected on virtual object(s)" criteria. That aspect is implemented fully via the reflections of the Swiss army knife box on the gemstones, but it is barely visible and so the Automagic management is left unimpressed.

3.3 Manage your time

The expected time to attain a 70 on this coursework is about 37 hours of directed effort. Marks beyond 80 are reserved for exceptional work. If you believe that you have such elements your submission, please include an explanation in your report. Those marks will be awarded at the discretion of the instructor and marker. Please be warned that chasing exceptional marks could consume large amounts of time and effort. So plan and prioritise!

4 Deliverables

Submit a compressed .zip with your student ID as the filename. e.g. "s123456.zip" via Learn. The structure of the zip should be as follows:

```
s123456
├── Exercises
│   ├── Radiometry.pdf
│   ├── Cameras.pdf
│   ├── RayTracing.pdf
│   └── MonteCarlo.pdf
├── Part1
│   ├── Project
│   │   └── s123456.blend
│   └── Images
│       ├── Original.png
│       ├── OriginalObject.png
│       ├── Intermediary.png
│       ├── Final.png
│       └── TextureX.png
├── Part2
│   ├── Project
│   │   └── s123456.blend
│   └── Video
│       └── Animation.mp4
└── Report
    └── s123456.pdf
```

Any deviations from this submission format will incur a penalty of 10% of the final mark for each impacted section.

Exercises: The Exercises folder will contain your solutions to the exercises in the coursework.

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 6a, 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 6 and second, explain how you made your video. 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.(a)).

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 Exercise questions

5.1 Radiometry

Design a scene where 3 points on a closed 2D shape of your choice are illuminated by 4 point lights at reasonable distances outside the shape. Show this scene by sketching an arbitrary closed shape in 2D. Mark 4 points (L_1 to L_4) anywhere at reasonable distances outside the shape and 3 points (P_1 to P_3) on the boundary of the shape. Imagine that point light sources are placed at L_i. Ensure that none of the points P_j are in complete shadow. You may choose arbitrary values for the power or flux (ϕ_i) emitted by these lights, the distance of each P_j from each of the L_i and the angle of incidence from each light. Tabulate these design choices (one 3×4 matrix for distances and another for angles). Calculate the irradiance at each P_j due to direct illumination from the lights. (hint)

5.2 Cameras

Look up the camera specifications of any mobile phone that you have access to. Report the make & model of the phone, the number of lenses it has and the focal length of the lenses. Take two pictures of an object in the foreground (at around 1m) with some objects in the background (>10m), at different zoom levels. For example, it could be a leaf on a tree with buildings in the background. Crop the zoomed-out image so that it roughly matches the zoomed-in image. Show the two original images and the crop in your solution. Use the 'image information' option to write down the camera parameters used for each image (focal length, aperture, ISO). Comment on differences in the size and shapes of foreground and background objects between the cropped image and the zoomed-in image if the focal lengths used were different (this depends on the phone).

5.3 Raytracing

Sketch a 2D scene containing the following: A virtual pin-hole camera with an image plane containing 5 pixels, one point light source, one area light source and two arbitrarily shaped objects. Mark and label the following: one primary ray, one path of length 5 from the area light source that passes through any pixel.

5.4 Monte Carlo

Recall the convergence plots shown in class while estimating the area of a circle using Monte Carlo integration. Perform the experiment for 5 different choices of radii. Include the 5 convergence plots of error vs number of samples (in log-log space), along with your chosen radii. What trends do you observe from these plots? Explain your conclusions.

6 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.

6.1 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.

6.2 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.

You can make this step much easier for yourself by ensuring that the scene you have chosen has many easily identifiable lines that you know are parallel to each other.

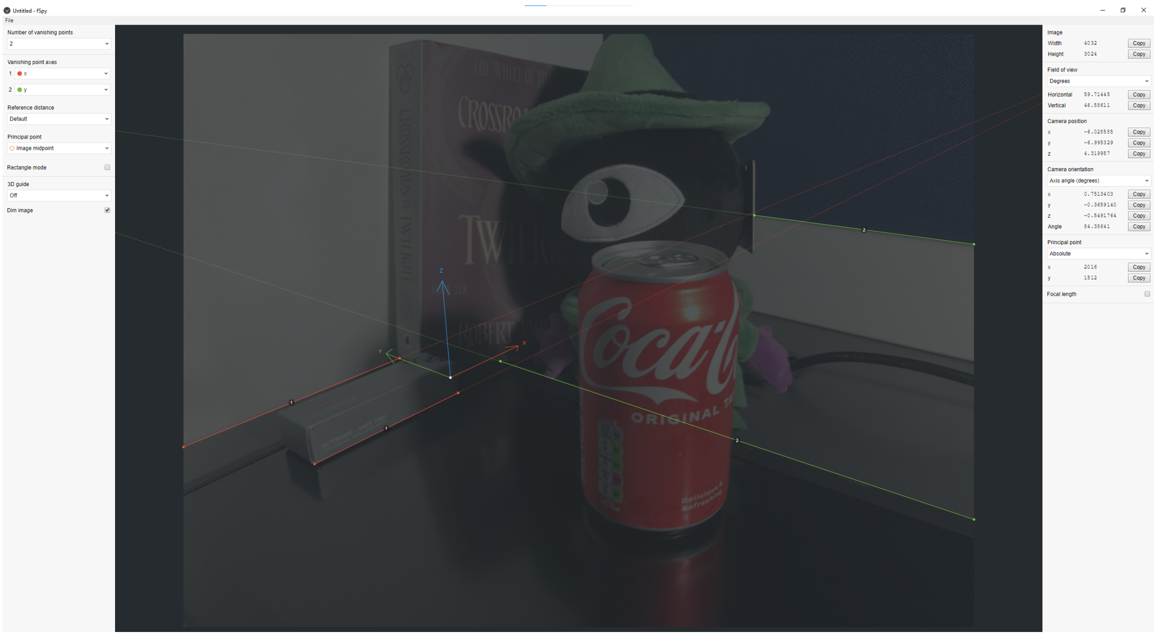

Figure 2: Background image in fSpy

6.3 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 is used for the reflection of real-world objects.

6.4 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.

Bear in mind that Blender has the capability to simulate not only hard shadows (where the light source is infinitely small) but also soft shadows where the light source can be of any size and shape. Make sure to modify your light object(s) to achieve best results.

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

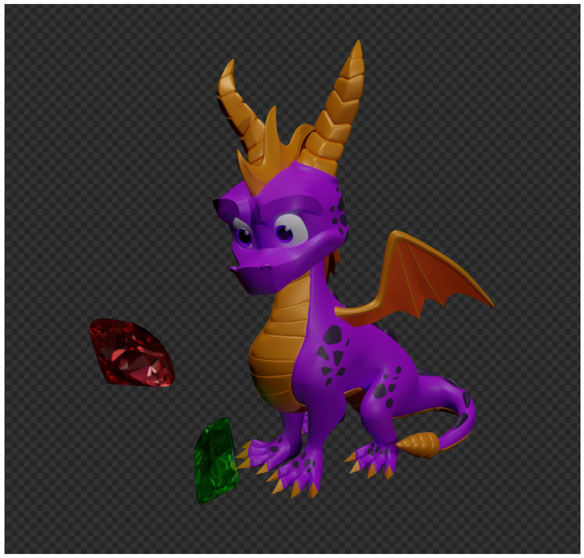

Figure 3: Example of virtual objects and their properties

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 4
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.

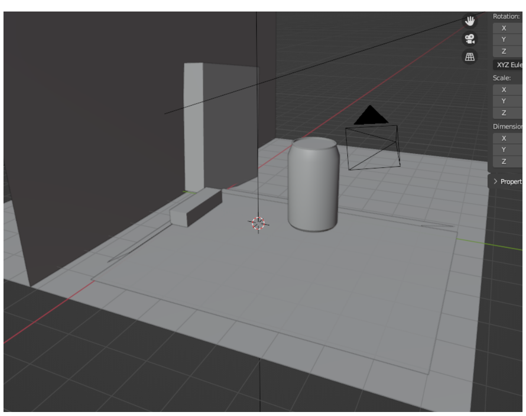

(a) Objects used to cast/receive shadows and provide reflections

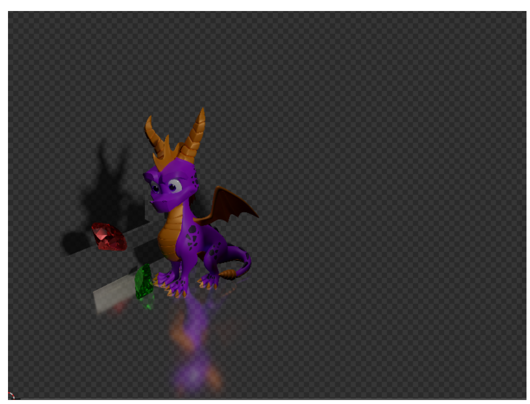

(b) Shadows and reflections that the objects left/picked in the scene

Figure 4: Example of using shadow catchers for creating shadow effects and reflective surfaces for creating reflections in the scene.

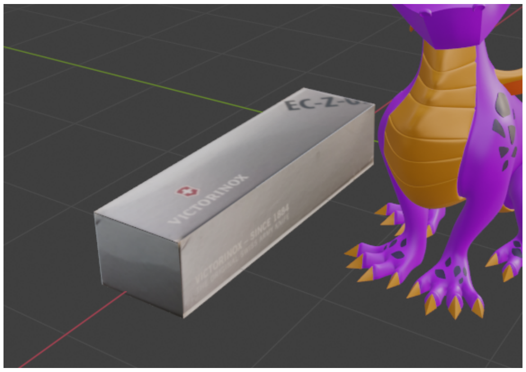

(a) Virtual representation of real object with real texture

(b) Real UV map texture imported in Blender

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

6.5 Compositing and Post-processing

Sometimes the resulting render might feel out of place for various reasons. In figure 6a for instance, the rendered parts are too crisp compared to the comparatively low quality camera image. Moreover the colours seem too vibrant. Such issues can be fixed within blender itself through the compositing tab! In the case of this example, the render from the cycles engine has been enhanced by a slight desaturation, blurring and an overlay of a noise texture in order to better fit its surroundings. The result of those actions are visible in 6b

Another option is to use image manipulation programs like gimp or photoshop to achieve similar results. If you believe the result of your render is already perfect and doesn't need any alteration you can skip this step.

(a) Blender raw render

(b) Final result after post-processing

Figure 6: Example of how the image can be processed to obtain a realistic image.

7 Tools

1. Blender for modelling, animation, physics and rendering –> Blender

2. fSpy for camera matching –> fSpy
3. Gimp for compositing –> Gimp
4. ffmpeg for making a video from rendered frames –> ffmpeg

References

[1] Drifting Walnut Productions. Quick fspy tutorial - 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] HappyMoon. Spyro statue - reignited trilogy. https://www.thingiverse.com/thing:3246882, 2018.