CS148 Homework 6 - Advanced Rendering

Grading on Monday, Nov 3 – see the pinned posts on Ed for details!

0.1 Assignment Outline

This assignment is a break from the coding last week and focuses on experimenting with the advanced rendering techniques in Blender. Specifically, you will play with depth of field, motion blur, and volume rendering.

0.2 Collaboration Policy, Office Hours, and Grading Session

These policies are the same as they were in HW2. See the HW2 document for details.

0.3 AI Policy

You may use AI tools to generate HDRI images, geometry, and materials to place in your scenes for this HW. However, you may not use them to generate anything else, especially your rendered images. All rendered images must be the final ray traced result from Blender Cycles.

Quiz Questions (1 pt)

You will be randomly asked one of these questions during the grading session:

- What about sampling causes aliasing, and what visual artifact does aliasing cause in an image? What must we do before sampling to do anti-aliasing?

- What causes motion blur to appear in images when we take photos with real life cameras? How do we modify ray tracing to account for motion blur?

- What visual effect does a depth of field produce for an image? How do we modify ray tracing to incorporate depth of field?

- What does it mean for light to be attenuated as it moves through a medium? How does the heterogeneous Beer's Law allow us to model the attenuation of light as it moves through non-homogeneous media?

- What is the difference between out-scattering vs in-scattering? Which of these two concepts causes participating media like smoke and clouds to actually have visible shape and why?
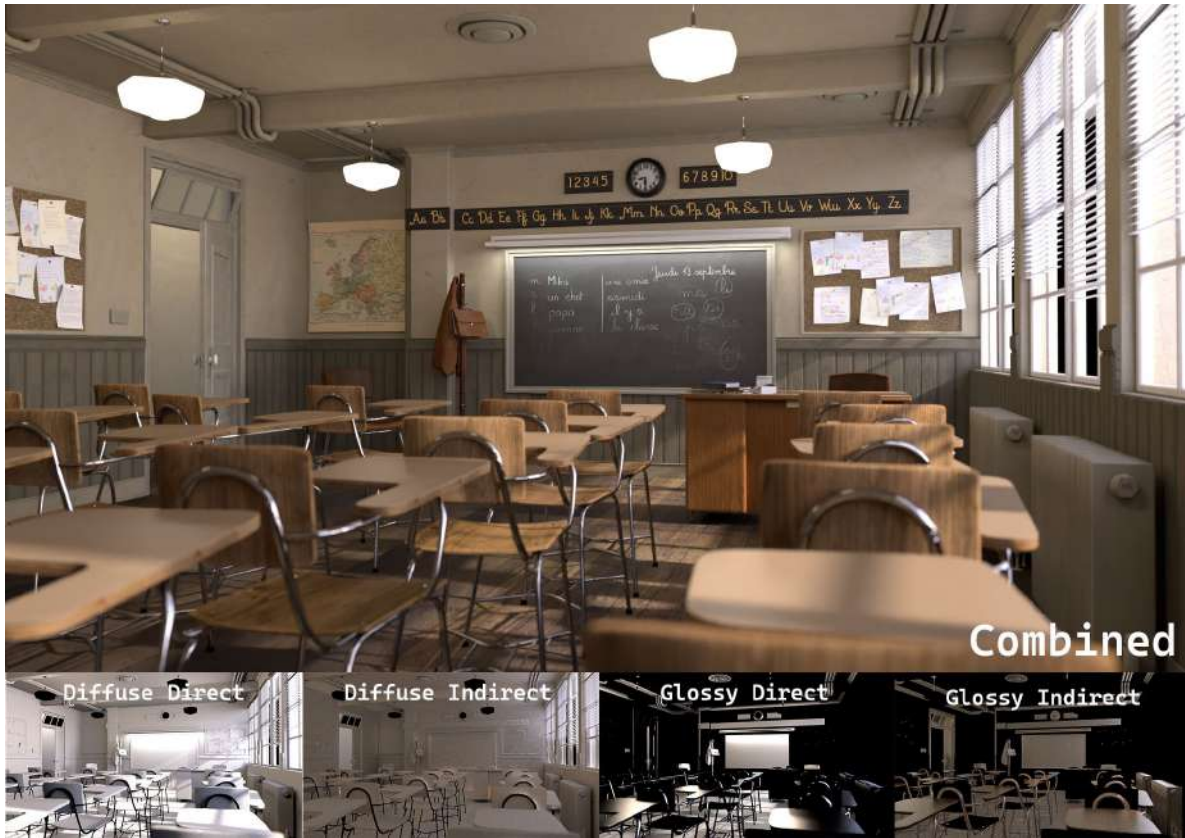
---

1 Assignment

1.1 `TODO 1`: Render Passes

Figure 1: A classroom scene with various render passes, taken from Blender demo files.

Before we get into the advanced rendering techniques, let's take a bit to review what we implemented in HWs 3 and 5 and see how it is done in Blender's Cycles render engine. The ray tracer in HW3 had the following lighting components summed up to get the final color for a ray trace pass:
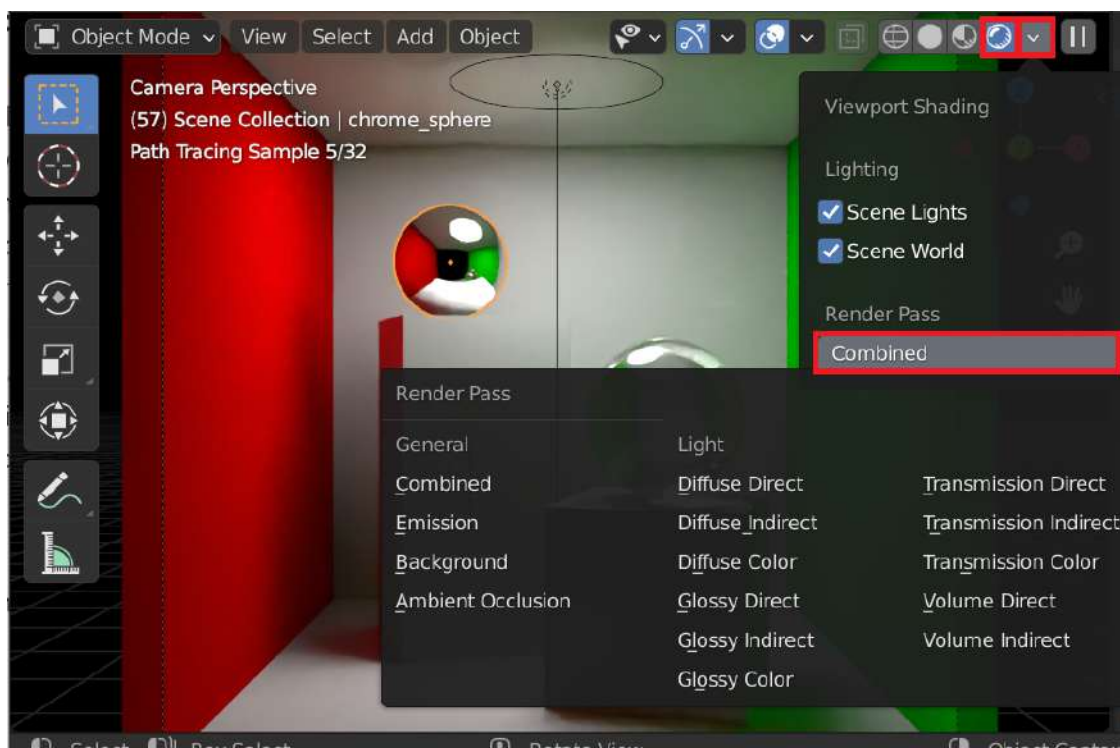
- Diffuse

- Specular

- Reflection

- Transmission

You then learned that these passes were all instances of "direct illumination." In order to achieve global illumination (i.e. color bleeding effects), we had to implement indirect lighting in HW5. While we only implemented indirect diffuse lighting (to keep the assignment simple), there also exist indirect versions of all the other above listed components.

To get a better sense of how each component contributes to a scene, we'll use Blender to view each component separately for the Cornell Box scene from HW5. Open this .blend file, which

contains the Cornell Box set up for you. The render engine has also already been set for you to Cycles. Toggle to the render preview in the viewport (If you forgot how to, then check HW4). You should see the Cornell Box scene from the last HW fully ray traced with 256 samples.

Now, the render that you see is called the "combined pass", which is the final result from combining the diffuse, specular, transmission, etc components into one image. Blender is also capable of ray tracing each component separately. To view these separate components, click the downward arrow next to the render preview toggle, then click on the drop-down menu under `Render Pass` :

Notice under the `Light` options, we have various direct and indirect options. Clicking on any of these options will show you a render of only that component, e.g. clicking on `Diffuse Direct` will give a render that shows only how direct illumination from the light interacts diffusely with the objects. **Click on each of the following components, aka render passes, and try to reason why each render pass looks the way it does based on knowing what each one is supposed to show.** If the render preview is slow, then scroll to the next page for another option.

1. `Diffuse Direct` : This render pass shows how the light(s) interacts with diffuse objects in the scene after just one direct pass (i.e. no recursive calls yet in the ray tracing algorithm). Objects that have high diffuse components in their material properties will receive a lot of direct illumination and often show up completely white (assuming the light color is white). Meanwhile, objects that are not very diffuse, or set to have no diffuse properties at all, will receive no direct illumination.

   Why do you think the spheres in the scene are completely black in this render pass, despite having color in the combined pass?

2. `Diffuse Indirect` : This render pass shows how the light(s) interacts with diffuse objects in the scene from the indirect bounces. That is, this render pass shows the indirect diffuse

3

illumination from the 2nd, 3rd, 4th, etc recursive calls of the ray trace function, omitting the illumination from the 1st call (because that would be the direct pass).

Why do you think there is red and green color in this pass while the direct pass is only black and white?

3. `Diffuse Color` : This render pass simply shows the colors of the objects in the scene based on their material properties without lighting. These are the diffuse colors that get multiplied onto every diffuse lighting calculation to get the overall, final diffuse color for the object(s). (Remember how you always multiplied the result of the ray trace function call with a "diffuse_color" array?)

   Why do you think this pass is flat shaded?

4. `Glossy Direct` : This render pass shows how the light(s) interacts with specular objects (with shiny highlights) in the scene after just one direct pass (i.e. once again, no recursive calls for direct). Objects that have high specular components in their material properties will have their specular (shiny) highlights showing in the render.

   What do you think are the white spots in this pass?

5. `Glossy Indirect` : This render pass shows the results of the recursive ray trace calls for reflection. Reflection is labeled under "Glossy Indirect", because "glossy" objects reflect light, and reflected light always takes at least two bounces (e.g. light → some object contributing its color → reflective object → camera) to reach the film plane, hence "indirect" lighting.

   Why do you think the right-most sphere looks different in this pass compared to how it looks in the combined pass? Notice how it has a black center here instead of a gray center. Where is that black center coming from?

6. `Glossy Color` : This render pass shows the reflectiveness of the object(s) in the scene using grayscale for intensity. Objects that are very reflective are near white (max intensity) while objects that are not reflective at all appear black (0 intensity).

   What do you think it means for the right-most sphere to appear mostly gray compared to the fully white sphere on the left? Can you also think of a reason to explain the white outline around the gray sphere? (Hint: it is something you coded up in HW3)

7. `Transmission Direct` : This render pass shows the results of only the first recursive ray trace calls for transmission, hence why it is a "direct" pass. This pass is often not very useful because most transmissive objects need at least 2 recursive calls to ray trace (e.g. light → enters object for transmission 1 → exists object for transmission 2 → camera). However, if there were a very thin transmissive object in the scene (e.g. a glass pane with nearly 0 depth), then the light might reach the camera on the first recursive call, and we would see something in the pass.

   **Why do you think this pass is completely black?**

8. `Transmission Indirect` : This render pass shows the results of the 2nd, 3rd, 4th, etc recursive ray trace calls for transmission.

   Why do you think the right-most sphere this time has a gray center when it had a black center in the glossy indirect pass?

9. **Transmission Color** : This render pass shows the transmissiveness of the object(s) in the scene using grayscale for intensity. Objects that are very transmissive are near white (max intensity) while objects that are not transmissive at all appear black (0 intensity).

   Why do you think only one sphere shows up in this pass?

Alternatively, besides viewing the render passes in the viewport, you can actually render each of these passes as images. To do so, go to the **View Layer Properties** in the **Properties Editor** and scroll to the **Light** field. Click the checkbox for all 9 of the above passes. Then render the scene as an image as you did in HW4. After the combined pass fully renders in the **Image Editor** , you can click on the drop-down menu in the upper-right to toggle through all the passes.
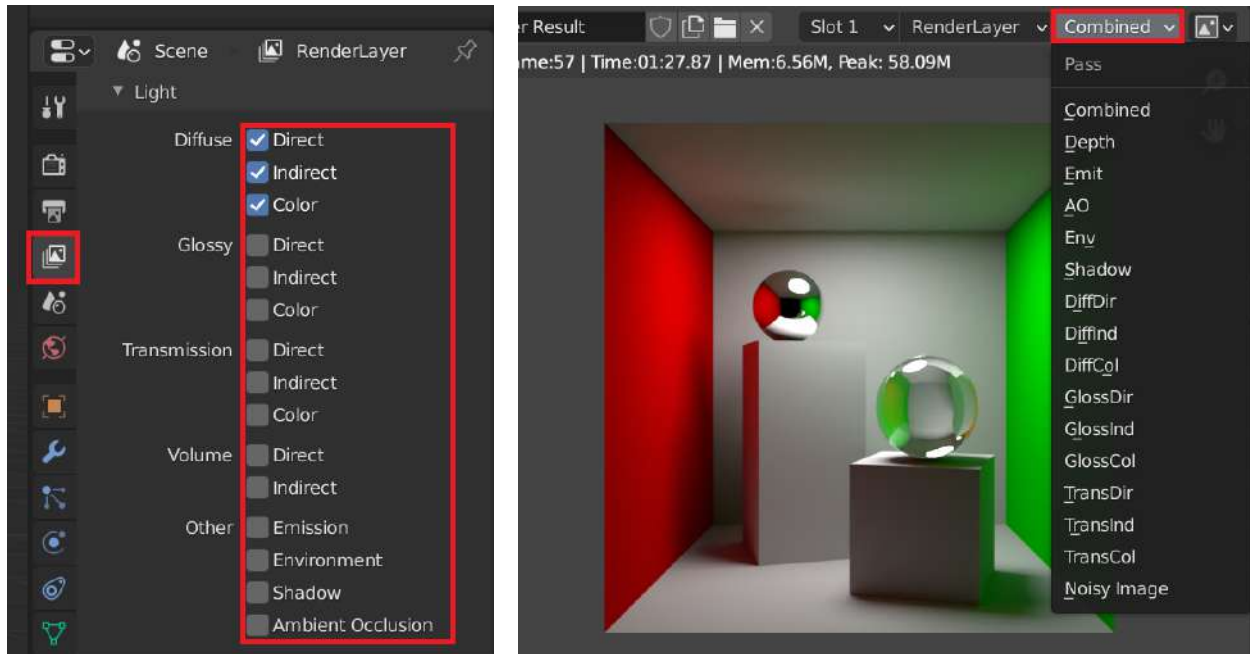
You might have noticed through all this that the ambient component of the color is missing from the render passes. This is because ambient color was color that we gave to objects when there was no influence from the light. This was to approximate the real life phenomenon where objects will emit their own light or have equal lighting from the environment to still be visible without explicit light sources. Blender does account for this phenonmenon with its **Emission** and **Background** aka **Environment** passes. You can read more in the Blender manual. Overall, the full formula that Blender uses to compute the combined pass is:

Combined = (Diffuse Direct + Diffuse Indirect) × Diffuse Color
+ (Glossy Direct + Glossy Indirect) × Glossy Color
+ (Transmission Direct + Transmission Indirect) × Transmission Color
+ Emission
+ Environment

> **Show us: (1 pt) All 9 render passes and be ready to answer 3 of the above questions in red from the set of 9, picked at random by the CA.**

1.2 TODO 2 : Motion Blur
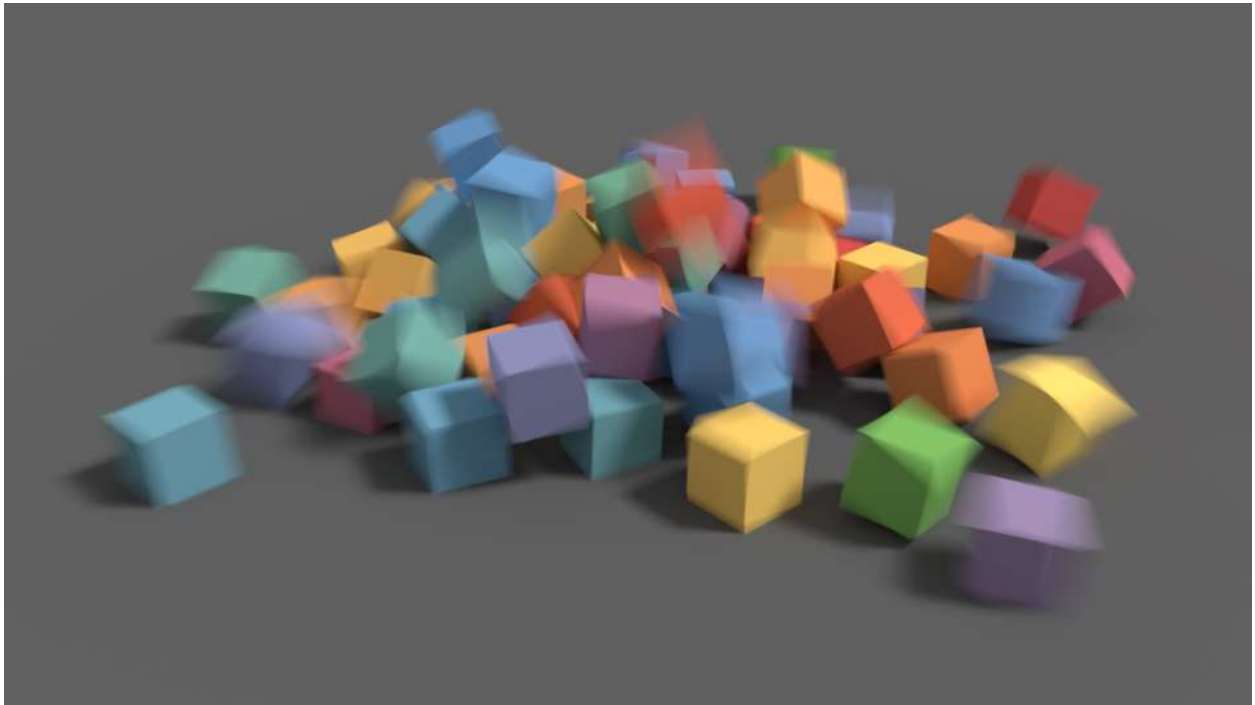
Blender provides a very simple way to set up and ray trace motion blur. Let's take a look with a simple example. First, taking the default Blender scene:

- Translate the cube to some location that you want to act as its starting point.

- Move the Timeline Editor at the bottom of the Blender interface up until you see the blue line clearly (see Figure 2).

- Right click the cube and click Insert Keyframe... (see Figure 2), then click Location (see Figure 3). This will mark the current scene as our first position in the timeline and also tell Blender that we would like to interpolate location aka position values for an animation.

- Move the blue line in the Timeline Editor to a later frame like frame 10.

- Go back and translate the cube to another location that you want to act its next position.

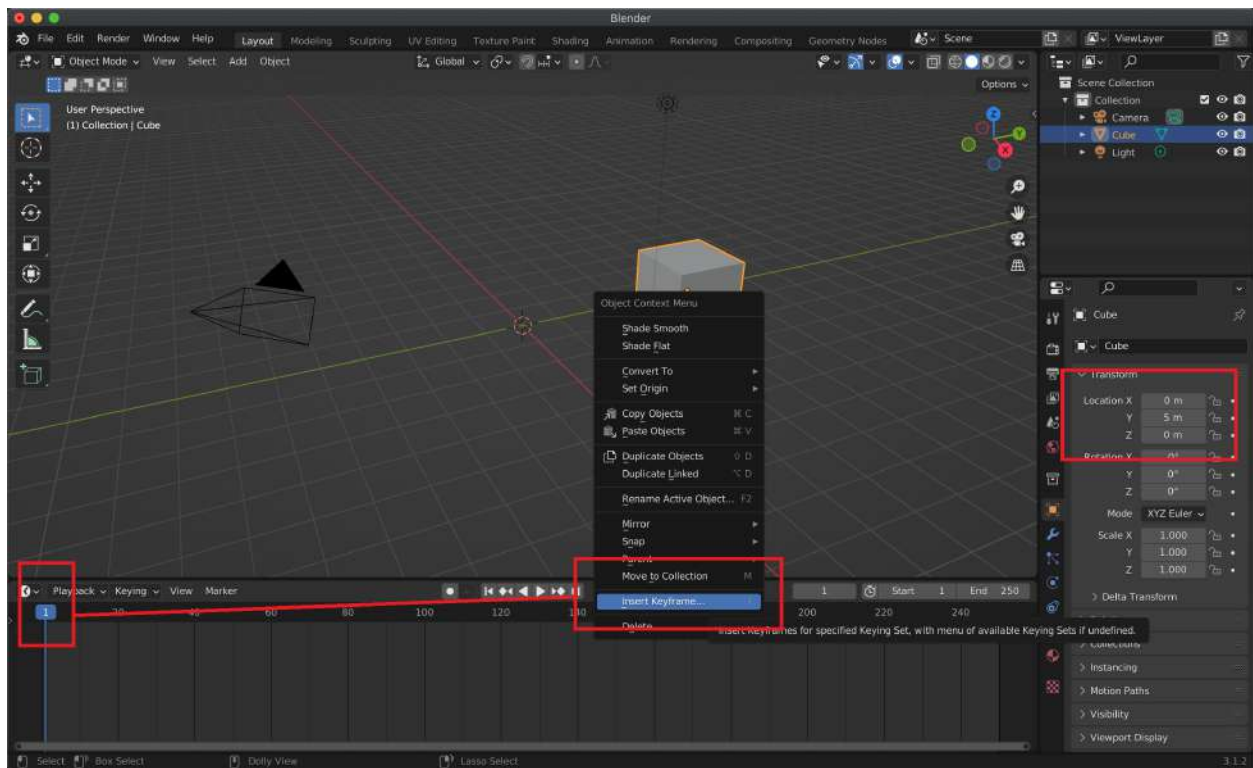- Then, using the same steps as above, insert a new location keyframe.
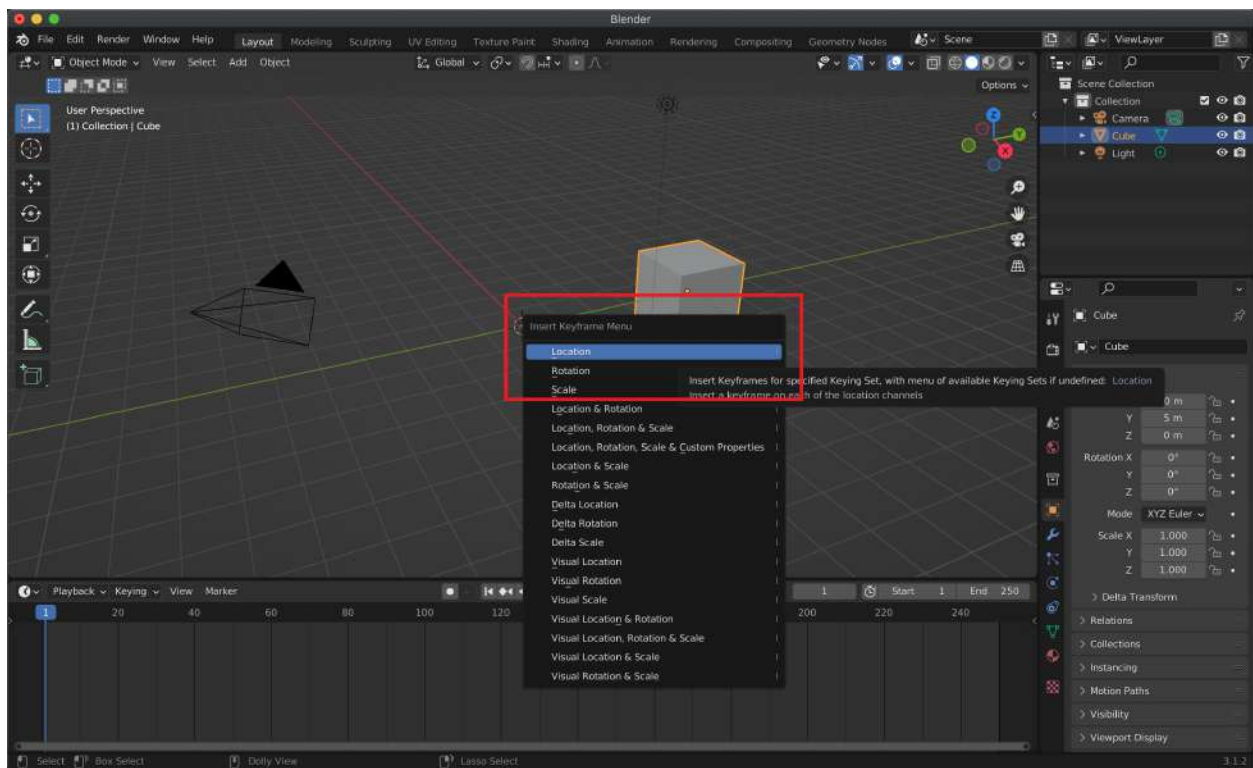
Figure 2

Figure 3

Now, if you scrub through the timeline by clicking and dragging the blue line from frame 1 to frame 10, you'll see your cube move smoothly from its initial position to its next position. Figure 4 shows our example with the cube at frame 7 of the animation and its interpolated position values.
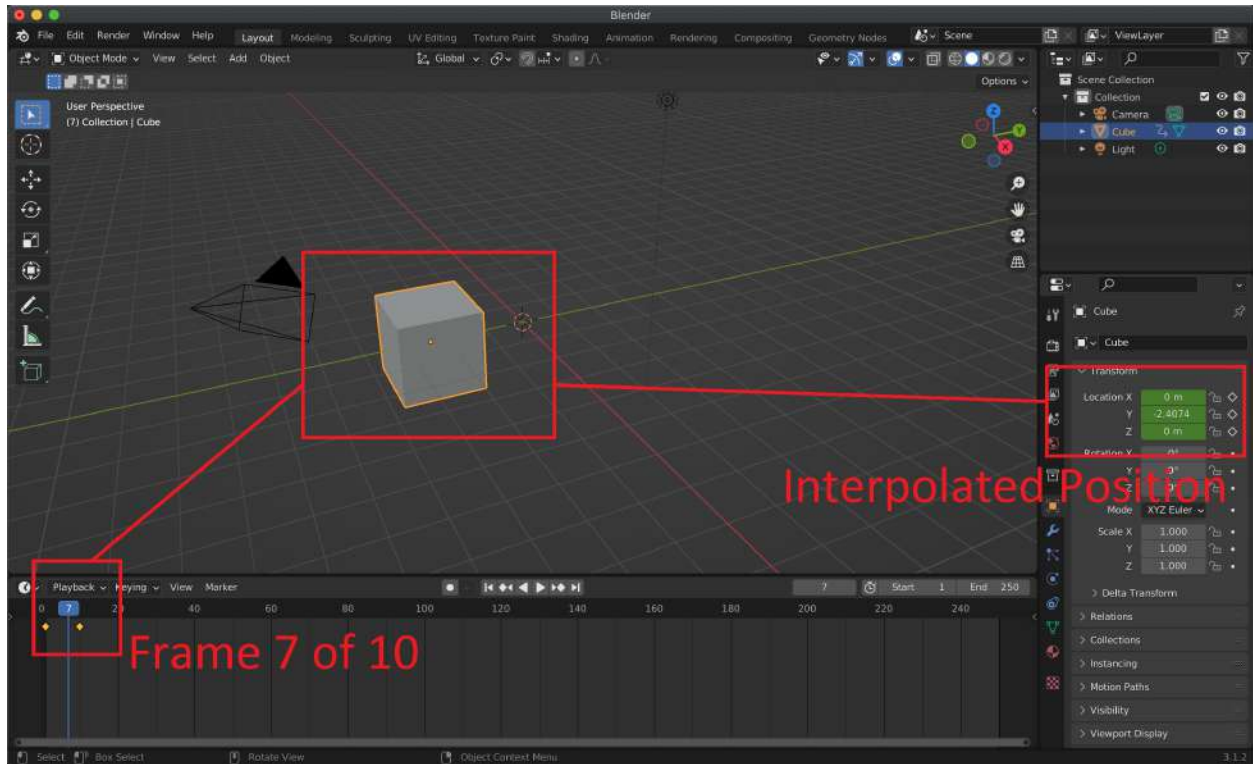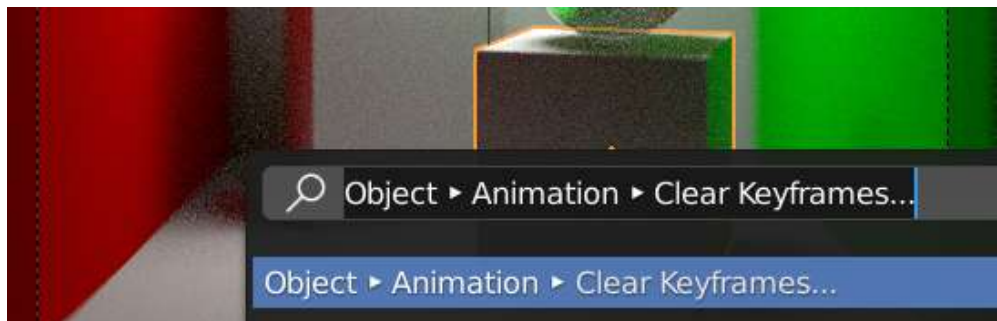

Figure 4

To have Blender ray trace your scene with motion blur, first go to Render Properties as usual to change the Render Engine to Cycles for ray tracing. Then scroll down in this Properties Editor tab to find the Motion Blur checkbox (see Figure 5). Just check the box, and Blender will do motion blur! To see, pick a frame in the middle of the animation (e.g. frame 7 of 10 in the timeline) and render your scene as an image. You should see something like that in Figure 5 pop up in your Blender render window.

If you ever want to remove all the animation data that you created for the motion blur, then you can do so by selecting all the keyframes in the Timeline Editor and deleting them (right-click → Delete Keyframes). Alternatively, you can select your motion blurred object (e.g. the cube) and press F3 to pull up the Menu Search , then type Clear Keyframes .
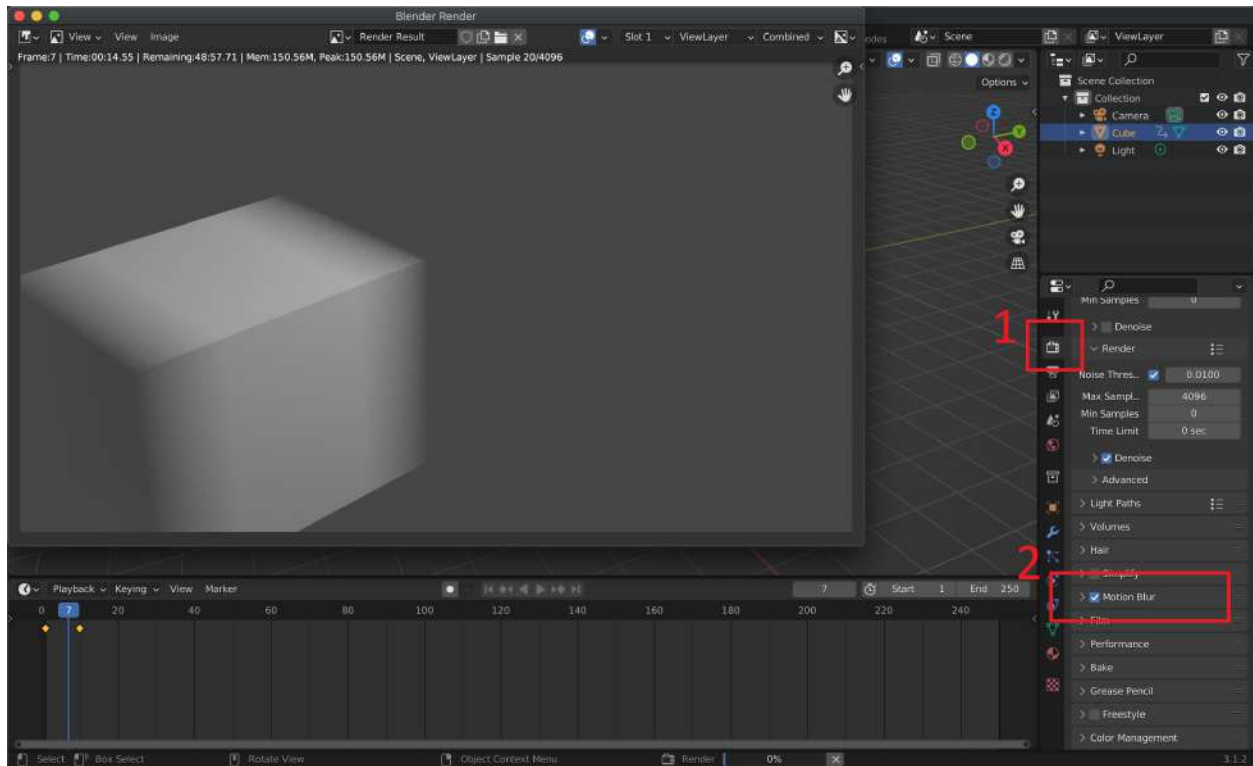

8

Figure 5

NOTE: It's VERY easy to unintentionally add TOO MUCH motion blur and make it look too exaggerated in your image. While you won't be deducted on this HW for having exaggerated motion blur, it can negatively affect your final project if you have it there. Use this as an opportunity to look up some real life photos of motion blur to get an visual intuition for what looks realistic.

1.3 TODO 3 : Depth of Field in Blender

Blender provides a very simple way to set up and ray trace depth of field. Let's take a look with a simple example. Open the .blend file containing the Cornell Box scene. Select the Camera in the Scene Collection in the upper-right of the interface. Then go to the camera's Object Data Properties in the Properties Editor . Click the checkbox for Depth of Field , and set the Focus Object to the glass_sphere . This will make our depth of field focus on the right-most sphere.

Play around with different values of F-Stop to get a feel for how it affects the strength of the blur. Lower values will increase blur, while higher values make the image more clear overall. For instance, try rendering the Cornell Box scene with depth of field focusing on the right-most, glass sphere using a F-Stop of 0.1. Then compare that sphere visually to the other sphere.
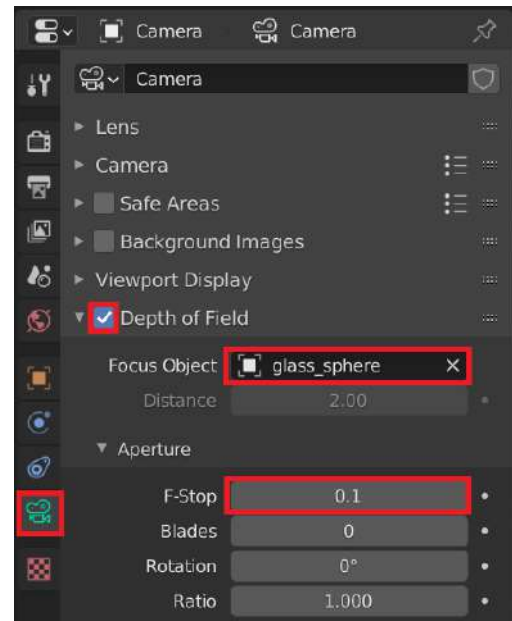
Show us: (1 pt) A reference photo or image that you took, found, or created and contains depth of field, along with your attempt to create the scene in that photo or image with the depth of field. You cannot use the Cornell Box scene. You can use placeholders for geometry, materials, and background if they are too complex for you to model at the moment. Try to replicate the camera angle and lighting as best you can though! You will need to justify to the CA how your work is a reasonable attempt at recreating the reference photo or image.

NOTE: Of all the advanced rendering techniques, depth of field is the most commonly applicable one to scenes. So most likely, you'll want to get a good intuition for how to adjust the settings above to have an easier time adding it to your final project later. It's very possible to have the depth of field options enabled, but have no actual depth of field visual effect in your image, because your settings aren't well tuned!

1.4 TODO 4 : Volume Rendering in Blender

As discussed in lecture, volume rendering is an advanced technique for achieving effects that cannot be represented by surface meshes alone. More specifically, volume rendering is designed to render volumetric effects like smoke, fire, fog, and clouds. We will try out Blender's volume rendering capabilities with a simple fog example. For more info on all the options that Blender has to offer for volume rendering, see its manual on volumetric effects (It's for an older version of Blender, but should still be relevant).

To set up our fog example, create a new default scene and enlarge our default cube. Then:

- Go to Material Properties in the Properties Editor (#1 in Figure 6).

- Click on the Surface option (#2 in Figure 6), then Remove or Disconnect (#3 in Figure 6) the default Principled BSDF material.

- Scroll down to the Volume options (#1 in Figure 7).

- Add a Principled Volume in the Volume option.

- Then, set the Density (of our fog) to 0.2 (#2 in Figure 7)

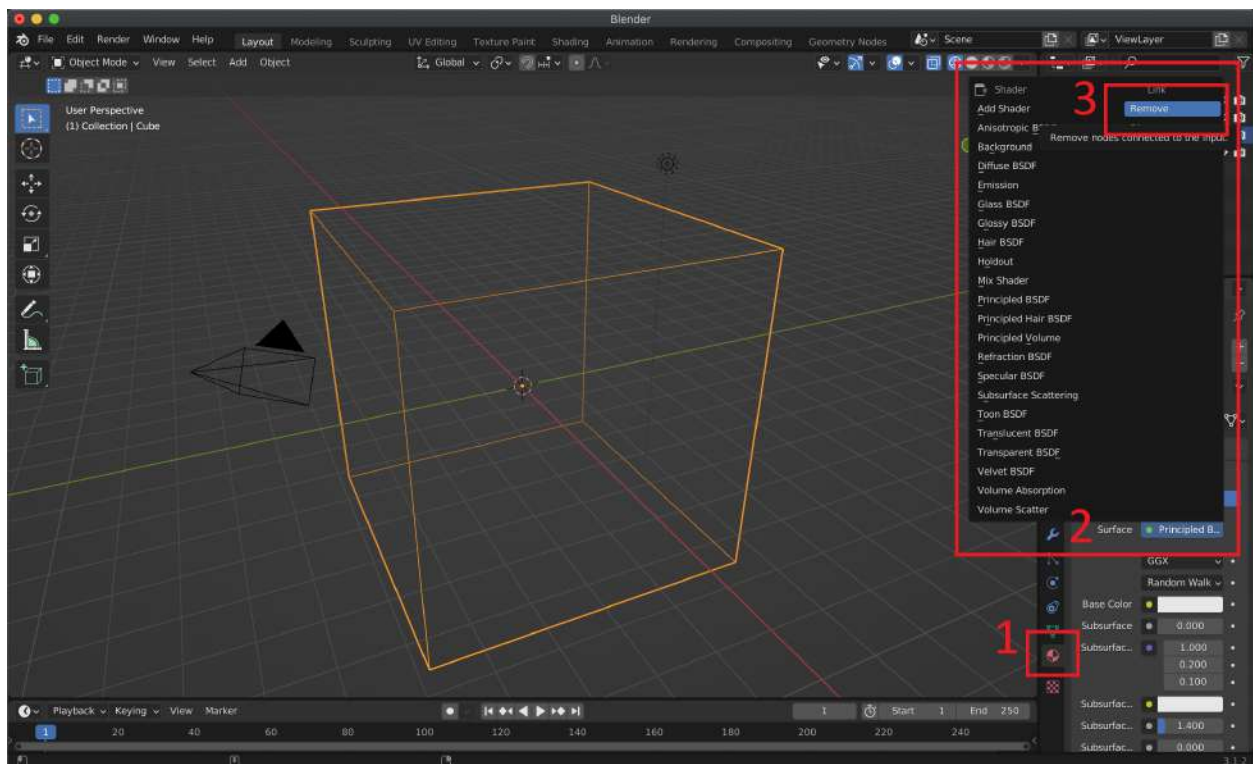- Add another cube inside our enlarged cube. You may want to switch to wireframe view when doing this.
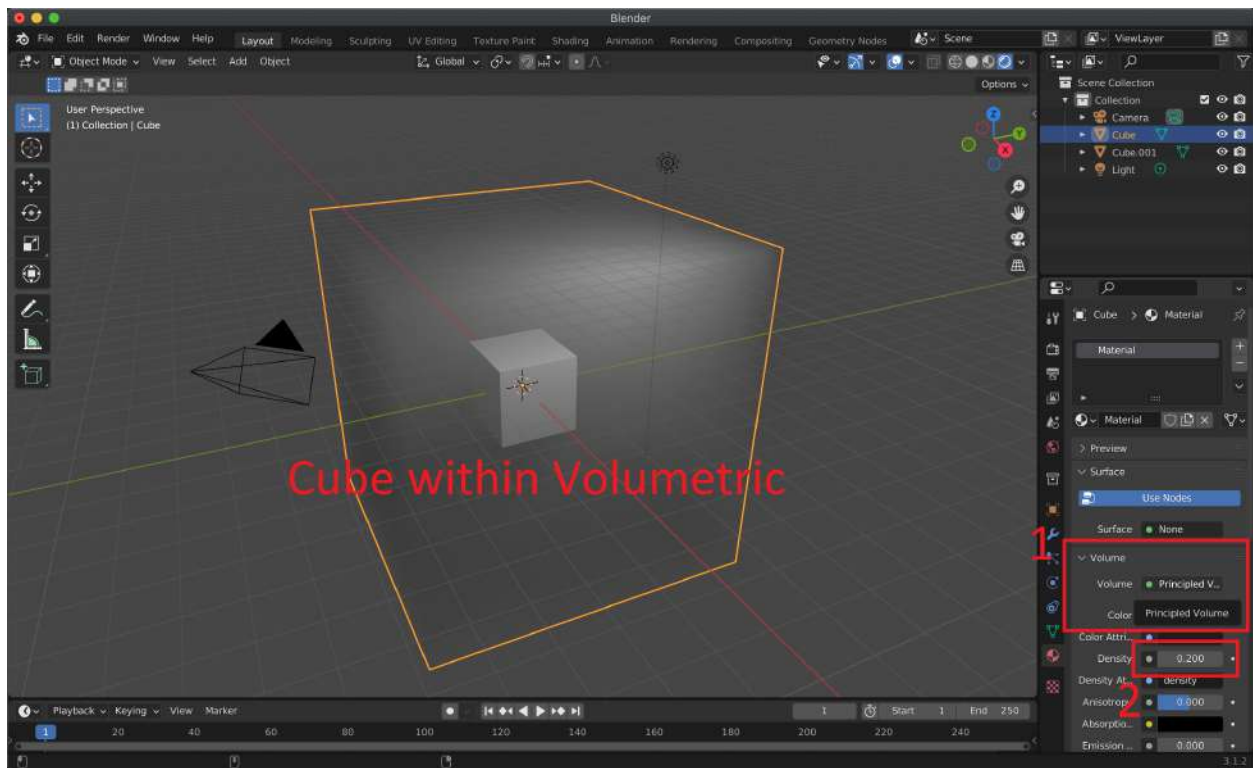

Figure 6

11

Figure 7

Set Blender to Cycles and toggle on the render preview. You should see a fog like effect inside the outer cube surrounding your inner cube. Render the image to see something similar to Figure 8.
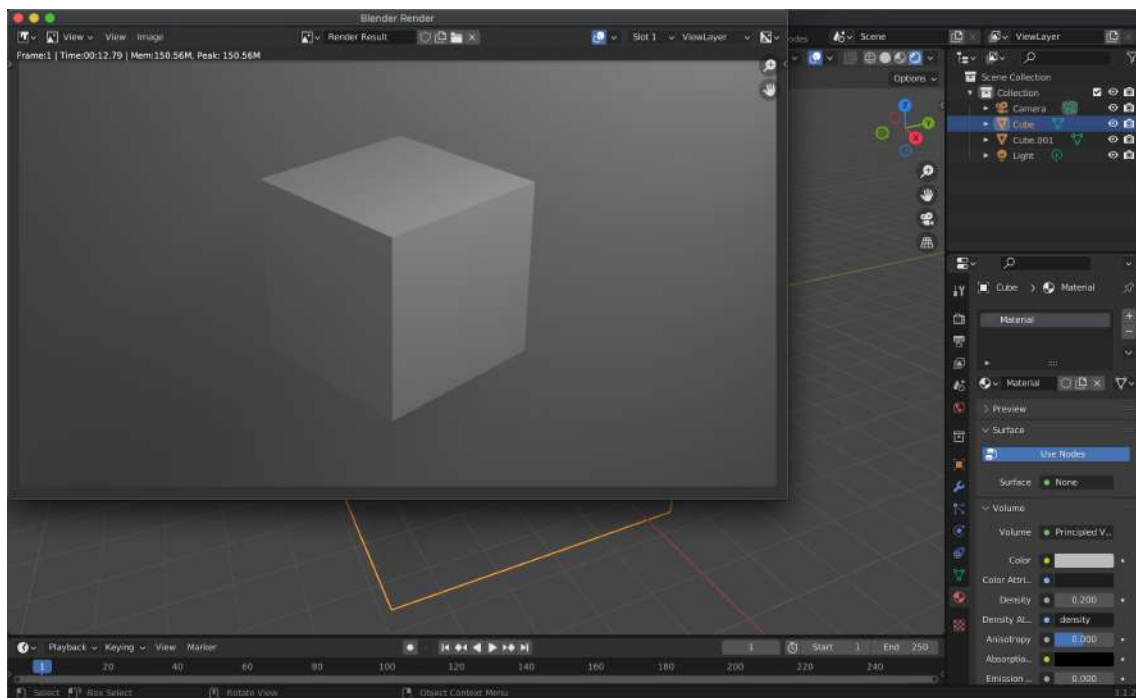

Figure 8

Try adding a fog effect to the Cornell Box scene – it results in a pretty interesting visual phenomena. First, we need to surround the entire Cornell Box scene with a volumetric cube. You can do this by adding a new cube, translating it to $(-2.75, 2.75, 2.75)$, and scaling all its dimensions by 4. When adding the material, you will have to create a `New` material first. When you render the Cornell Box scene with fog, you might notice that the shape of the light becomes visible for reasons we will discuss later in lecture! (Self-check: why does the light become visible? Ask in office hours if you're not sure!)

> **Show us: (1 pt) A reference photo or image that you took, found, or created and contains volume rendering (e.g. find one with fog, clouds, etc), along with your attempt to create the scene in that photo or image with volume rendering. You cannot use the Cornell Box scene. You can use placeholders for geometry, materials, and background if they are too complex for you to model at the moment. Try to replicate the camera angle and lighting as best you can though! You will need to justify to the CA how your work is a reasonable attempt at recreating the reference photo or image.**

> **NOTE: This is a good opportunity to look up tutorials for more complex volume rendering like clouds, fire, smoke, etc if you're interested!**