Converging verticals

Dr John Miller
Wooler, Northumberland NE71 6EL
February 2023

What is a converging vertical? Easy - it's what you get if you point your camera upwards to get everything into the frame. You know for certain that the sides of a building are vertical[1]. Your picture of it shows that the building's sides lean in towards each other. The building is real and the picture is real but the mismatch can be a trifle disturbing.

✯ ✯ ✯ ✯ ✯

The four diagrams below show ***why*** converging verticals occur.

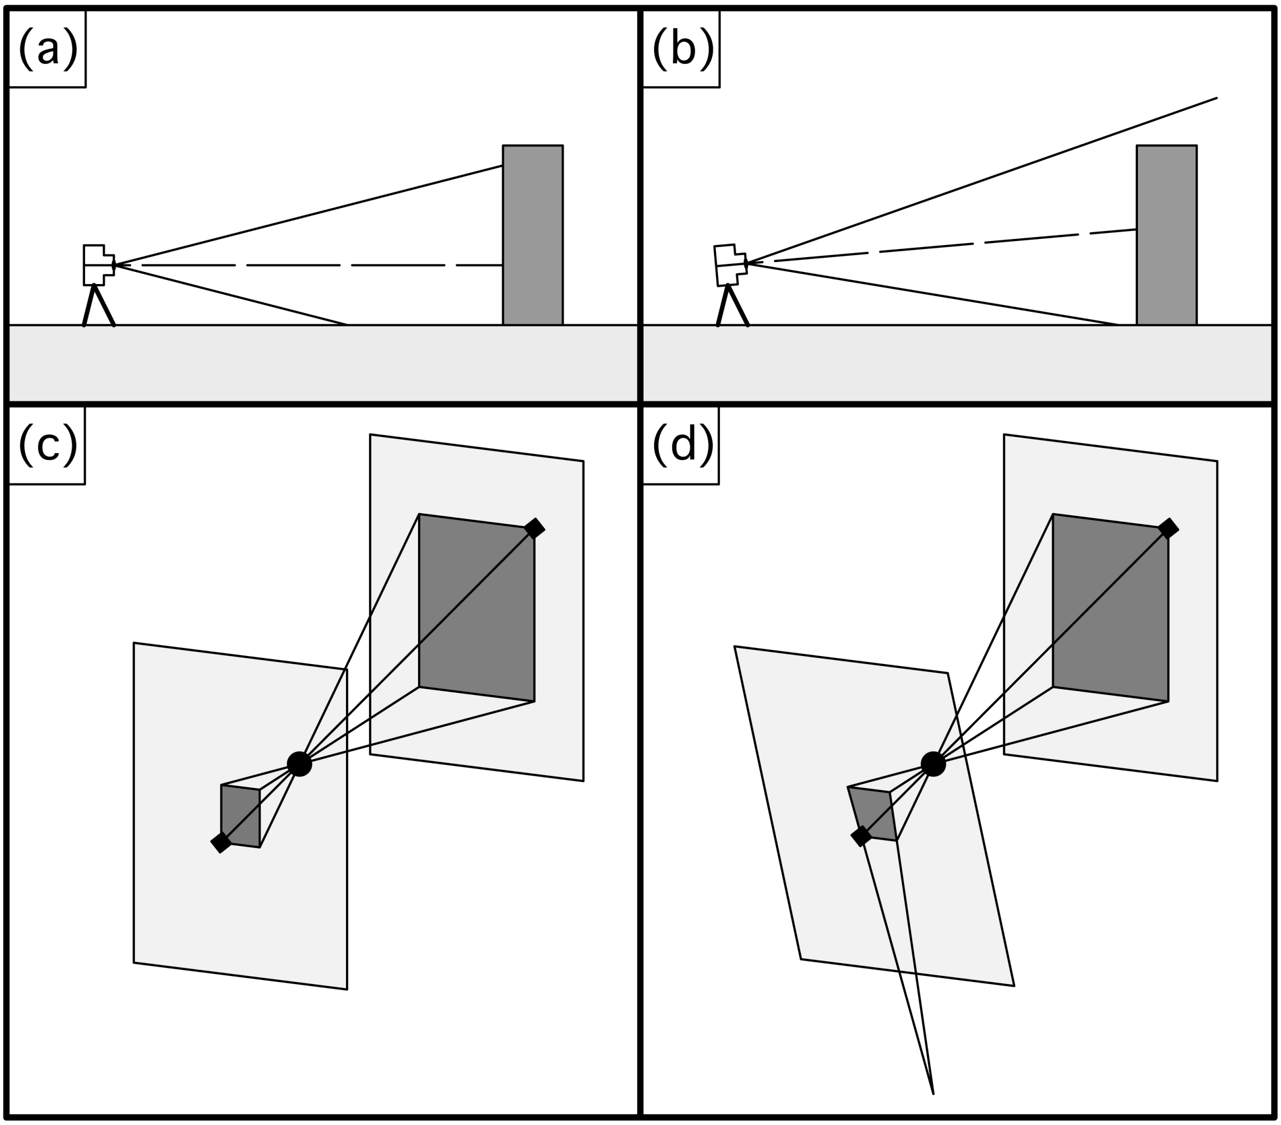

Diagrams (a) and (b) are of a camera and a building.

In Diagram (a), the camera is levelled. As a result the sensor in the camera and the front of the building are both vertical, and therefore parallel. The angle of view of the camera causes the top of the building to be outside the frame, and perhaps too much foreground to be in the frame.

In Diagram (b), the camera is pointed upward. The front of the building remains vertical but the camera's sensor is not vertical, and so they are ***not*** parallel. The camera tilt, however, brings the top of the building into the frame, and there is less foreground.

1 If they're not vertical, keep away - it might be about to fall down.

Diagrams (c) and (d) are ray diagrams showing how the subject forms an image on the camera sensor. Diagram (c) relates to Diagram (a), and Diagram (d) to Diagram (b).

In Diagrams (c) and (d), light rays from the subject travel in straight lines to the sensor. The rays pass through the camera lens and lens aperture (the aperture position is shown by a black circle) and fall on the sensor. The image on the sensor is inverted. As the subject is a rectangle only rays from its corners are needed to project the image. In Diagrams (c) and (d) subject and image are dark grey, and the subject and image planes[2] are light grey rectangles.

In Diagram (c) the subject and image planes are parallel. Compared with the subject rectangle, the image is a smaller rectangle of identical shape.

In Diagram (d) the subject and image planes are not parallel. Compared with the subject rectangle, the image is a small trapezium, not a rectangle. The longer sides of the image are not parallel: they are ***converging verticals***.

✯ ✯ ✯ ✯ ✯

The camera levelled but pointed obliquely at a vertical flat rectangular surface is effectively Diagram (d) turned onto its side. The result is an image with ***converging horizontals***.

If the camera is tilted up and pointed obliquely at a vertical rectangle, the resulting image is a quadrilateral with opposite sides converging.

From a topological viewpoint[3], all buildings are cubes. Viewed from an arbitrary direction, the image of a cube projected onto a flat surface will always have convergences. Be warned!

✯ ✯ ✯ ✯ ✯

A photographic image with converging verticals, and to a lesser extent converging horizontals, may to some degree be disagreeable. This raises the question of how to avoid convergences. There are three main ways.

First, get the image right ***before*** exposure. Level the camera: that avoids converging verticals. Pointing the camera squarely at the subject also reduces converging horizontals. Choose a lens focal length sufficiently short (i.e. wide-angle) to get all the subject into the frame. If possible, move back from the subject if your shortest lens is too long. The disadvantage is that images taken with levelled short lenses have too much unwanted foreground.

Second, use a tilt-and-shift lens[4]. After levelling the camera, the lens is shifted upward (i.e. relative to the sensor) to get the top of the subject into the frame and reduce unwanted foreground. The disadvantages are that tilt-and-shift lenses are hideously expensive, and offer little choice of lens focal length.

Third, stop worrying about getting the image right before exposure. Choose your lens focal length and point the camera so that all the subject is in the frame, then correct the convergences in the image in software. Any half-decent image editor has this correction capability, though it might be slow and fiddly to use[5]. The disadvantage is the time it takes to learn how to use complex software.

✯ ✯ ✯ ✯ ✯

There is no quick and easy way ***always*** to avoid convergences in images. The efficient way is to do the best you can before exposure, then tidy up the image in software after exposure. What you hope for, however, may not be attainable.

---

2 The front of the building is assumed be a flat surface, and the light-sensitive side of the sensor is a flat surface. The grey rectangles are extensions of those surfaces, making the angle between the two surfaces obvious.

3 "Topology is the mathematical study of the properties that are preserved through deformations, twistings, and stretchings of objects. Tearing, however, is not allowed." (Wolfram MathWorld)

4 The "shift" action removes convergences. The "tilt" action increases depth of field.

5 The easiest software to use is DxO ViewPoint, which is explicitly designed for perspective correction.