


Virilia



the whole continent, 6.4 km a pixel

A made-up continent, 7,795 by 4,570 km. The shape of the land was drawn by hand; everything else is calculated from it - where rain falls, where rivers run, what grows, where ice sits. The map you are looking at is built from that, not painted.

A river valley



48 km across, 50 m a pixel - the finest the data goes

Rivers are drawn as water, not as lines. Each one gets a real width from how much land drains into it, and a real depth from that width, so a big river is wide and dark and a small one is a thread. The dark edge along each bank is the same trick an atlas uses to make water sit in the land rather than on top of it.

Mountains



40 km across, 50 m a pixel

The shading uses four lights instead of one, which is what stops a single-sun hillshade leaving half of every range in flat shadow. Slopes are compressed rather than exaggerated: gentle ground gets pushed hard and steep ground barely at all, so a plain and a mountain can share one picture without either turning to mush or to crumpled foil.


Coast and islands



60 km across, 50 m a pixel

The sea colour comes from depth, measured off the sea bed itself rather than picked. Where the bed was never modelled - which is about half this ocean - the depth is filled in by interpolating between the parts that were, so the open water still deepens away from the shore instead of stopping at a straight line.

Ice



55 km across, 50 m a pixel

Snow, sea ice and frozen lakes all come from the climate: the warmest and coldest month at every point. The edge of the ice sheet is not a line on a map - it wanders with the ground, holding on the ridges and retreating up the valleys, the way a real ice margin does.

Same place, six zoom levels

800 m a pixel: 768 km across



400 m a pixel: 384 km across



200 m a pixel: 192 km across



100 m a pixel: 96 km across



50 m a pixel: 48 km across



25 m a pixel: 24 km across



768 km down to 24 km across

The whole map is baked into tiles, nine levels of them, 73,000 in all. The finest level is exactly 50 m a pixel - the same grid the height data is stored on - so nothing is resampled at the bottom. Zooming changes which level you are looking at, and the brightness and colour are held steady across the change so it does not flicker.

The tiling problem




left to right: alive but seams / tiles but dead / both


Each tile is drawn on its own and can never see its neighbours. Shading needs to know how steep the ground is 'normally' - measure that per tile and the tiles disagree at their edges; use one number for the whole world and flat country goes blank, as in the middle panel, where the whole river network has vanished. The fix is to work it out from a small patch around each pixel. Neighbouring tiles then get the same answer because they are asking about the same ground.

Matched against a real photo

your reference - 135 m a pixel



Virilia at 135 m a pixel



Landsat of the Croatian coast (left), Virilia (right), both 135 m a pixel

Colour was set by comparing with satellite imagery rather than by eye. Brightness, contrast and saturation were measured on both and the settings moved until they matched. The land turned out to be about four times too colourful, while the sea was already right - so the two are adjusted separately.