Guided Super-Resolution of Environmental Maps

Riccardo de Lutio, Stefano D'Aronco, Jan Dirk Wegner, Konrad Schindler

EcoVision Lab, Photogrammetry and Remote Sensing, ETH Zürich

Environmental mapping often requires extensive and expensive field work, the resolution that can be achieved for these maps is significantly lower than the ground sampling distance of modern earth observation satellites.

Guided super-resolution is a unifying framework for several computer vision tasks where the inputs are a low-resolution **source image** of some target quantity - e.g., a vegetation height map - and a high-resolution **guide image** from a different domain - e.g., a multispectral Sentinel-2 image; and the target output is a high-resolution version of the source, in our example a high-res vegetation height map (note that the sensing principle of several airborne and spaceborne imaging sensors can also be seen as guided super-resolution, where a high-resolution panchromatic channel is fused with colour channels recorded at lower resolution).

The standard way of looking at this problem is to formulate it as a super-resolution task, i.e., the source image is upsampled to the target resolution, while transferring the missing high-frequency details from the guide. Here, we propose to turn that interpretation on its head and instead see it as a pixel-to-pixel mapping of the guide image to the domain of the source image. The pixel-wise mapping is parametrised as a multi-layer perceptron, whose weights are learned by minimising the discrepancies between the source image and the downsampled target image. Importantly, our formulation makes it possible to regularise only the mapping function, while avoiding regularisation of the outputs; thus producing crisp, natural-looking images. The proposed method is unsupervised, using only the specific source and guide images to fit the mapping.

We evaluate our method by super-resolving tree height maps for Switzerland (Ginzler & Hobi 2015). We clearly outperform recent baselines in quantitative comparisons, while delivering visually much sharper outputs. We also show promising preliminary results for the task of creating dense nation-wide maps of biodiversity measures for Switzerland (BDM).

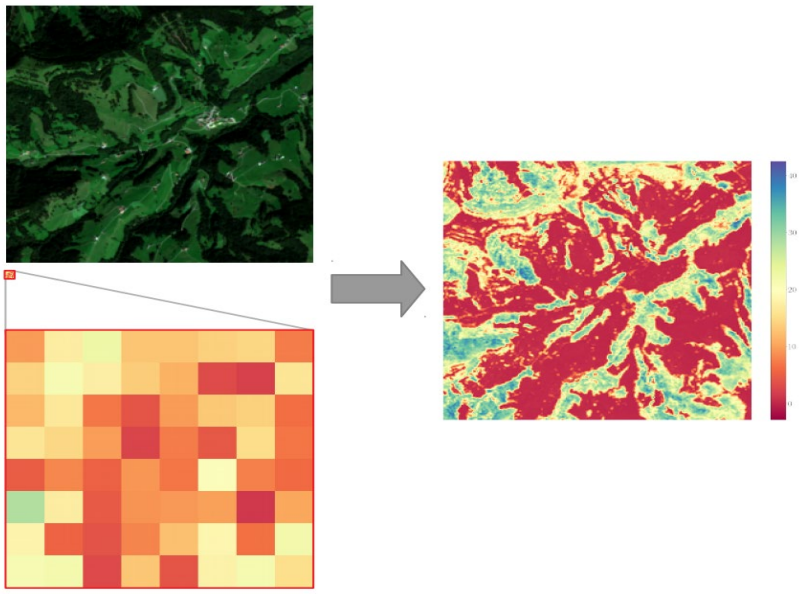

Figure 1. Guided super-resolution: given a low-resolution VHM and a high-resolution guide image, our method predicts a high-resolution VHM. The figure shows an example output of the proposed method, for an upsampling factor of 32X.

REFERENCES

de Lutio, R., D'Aronco, S., Wegner, J. D. & Schindler, K. : Guided super-resolution as pixel-to-pixel transformation, International Conference on Computer Vision (ICCV), 2019.

Ginzler, C. & Hobi, M. L. : Country-wide stereo-image matching for updating digital surface models in the framework of the Swiss National Forest Inventory, Remote Sensing, 2015.

Biodiversity Monitoring Switzerland (BDM) : data and findings up to 2018, http://www.biodiversitymonitoring.ch/en/home.html.