




Article

SEN2VEN μ S, A DATASET FOR THE TRAINING OF SENTINEL-2 SUPER-RESOLUTION ALGORITHMS

Julien Michel ^{1,†} , Juan Vinasco-Salinas ², Jordi Inglada ^{3,†} , Olivier Hagolle ^{4,†} 

- ¹ CESBIO, Université de Toulouse, CNES, CNRS, INRAE, IRD, UT3, Toulouse, France; julien.michel@cesbio.cnes.fr
- ² CESBIO, Université de Toulouse, CNES, CNRS, INRAE, IRD, UT3, Toulouse, France; juan.vinasco@cesbio.cnes.fr
- ³ CESBIO, Université de Toulouse, CNES, CNRS, INRAE, IRD, UT3, Toulouse, France; jordi.inglada@cesbio.cnes.fr
- ⁴ CESBIO, Université de Toulouse, CNES, CNRS, INRAE, IRD, UT3, Toulouse, France; olivier.hagolle@cesbio.cnes.fr
- † Current address: 18 avenue Edouard Belin BPI 2801 31401 Toulouse Cedex 9

Abstract: Boosted by the progress in deep learning, Single Image Super-Resolution (SISR) has gained a lot of interest in the Remote Sensing community, who sees it as an opportunity to compensate for satellite’s ever-limited spatial resolution with respect to end users needs. While there has been a great amount of work on network architectures in the latest years, deep learning based SISR in remote sensing is still limited by the availability of the large training sets it requires. The lack of publicly available large datasets with the required variability in terms of landscapes and seasons pushes researchers to simulate their own dataset by means of downsampling. This may impair the applicability of the trained model on real world data at the target input resolution. In this paper, we propose an open-data licenced dataset composed of 10m and 20m cloud-free surface reflectance patches from Sentinel-2, with their reference spatially-registered surface reflectance patches at 5 meter resolution acquired on the same day by the VEN μ S satellite. This dataset covers 29 locations on earth with a total of 132 955 patches of 256x256 pixels at 5 meters resolution, and can be used for the training of super-resolution algorithms to bring the spatial resolution of 8 of the Sentinel-2 bands down to 5 meters.

Keywords: Single-Image Super-Resolution, Sentinel-2, Dataset

1. Introduction

Single Image Super-Resolution is the task of obtaining a higher resolution version of a single image, using no other inputs than the image itself. In the context of remote sensing satellite imagery, higher resolution means that not only the same area on the ground is covered by a higher number of smaller pixels - which can be achieved by means of traditional spatial re-sampling techniques - but also that the super-resolved image exhibits faithful higher spatial frequency content with respect to the original image.

Super-resolution is an ill-posed inverse problem, as many higher resolution predictions can explain the same low-resolution image. Prior to the deep learning era, super-resolution was considered as a blind deconvolution problem. It has therefore been traditionally tackled by regularization constraints during optimization, in order to promote desired properties of the prediction, such as piece-wise smoothness, with Total Variation [1], or sparsity [2].

The advent of deep learning gave a fresh start to the Single Image Super-Resolution problem, for which convolutional, residual and then Generative-Adversarial architectures have been proposed with success [3,4]. Provided that enough data are available for training, we are able to estimate a non linear mapping with a few hundred thousand parameters that undoes all the high frequency damping and aliasing occurring at sensor level, and even generate plausible high resolution details past the cut-off frequency of the sensor.

In order to feed those data-hungry algorithms and lay grounds for architecture contests, several natural images dataset have been published [5–7]. Most of these datasets are put up by down-sampling the high resolution image, using the original high resolution image as a reference for training, validation and testing. In practice, this simulation of the low resolution data leads to a simplified version of the problem, by ignoring parameters of low resolution acquisition, such as noise or compression artifacts. This in turn may lead to lesser performances when applied to real world data. More importantly, to the best of our knowledge, no remote sensing Single Image Super-Resolution dataset has been made publicly available yet. There is however a public dataset dedicated to the super-resolution of Proba-V images using multiple images (Multiple Images Super-Resolution) [8].

The progress of Single Image Super-Resolution of course raised a lot of interest in the field of remote sensing optical imagery [9] and its never-ending race for better resolution. Remote sensing images are very different in nature from natural images used in the previously cited natural images dataset : the image content itself is different, consisting in Earth surface observed from an almost fixed altitude, which means object of typical sizes will always be represented by roughly the same amount of pixels. But more importantly, remote sensing images can capture earth radiance in several wavelength bands that may or may not include what is called Red, Green and Blue in natural images, with a bit depth exceeding the traditional 8-bits encountered in most natural images. Moreover, spectral bands may be acquired at different spatial resolution. There is therefore a need for a dedicated dataset that can represent all the specifics of remote sensing imagery in the development and comparison of super-resolution algorithms.

In this paper we propose SEN2VEN μ S, an open dataset for the super-resolution of Sentinel-2 images by leveraging simultaneous acquisitions with the VEN μ S satellite (see section 2.1 for a detailed description of those satellite missions). The dataset is composed of 10m and 20m cloud-free surface reflectance patches from Sentinel-2, with their reference spatially-registered surface reflectance patches at 5 meter resolution acquired on the same day by the VEN μ S satellite. This dataset covers 29 locations with a total of 132 955 patches of 256x256 pixels at 5 meter resolution, and can be used for the training of super-resolution algorithms to bring spatial resolution of 8 of the Sentinel-2 bands down to 5 meters.

2. Dataset generation

2.1. Sentinel-2 and VEN μ S missions

Sentinel-2 is the well known high revisit optical Copernicus mission operating since mid-2015 [10]. It provides a full coverage of all lands between 56° south and 84° north, every 5 days at most, in 13 spectral bands in 3 different resolution groups : 10 meters for visible bands and wide Near Infra-Red (NIR), 20 meters for red edges bands and narrow NIR and 60 meters for coastal blue and atmospheric corrections bands. Viewing angles from a given orbit are constant and lower than 12°, and orbit crosses equator at 10h30 local time. Sentinel-2 data are distributed with an open-data policy, and recurring satellites are provisioned until 2035. Sentinel-2 is therefore a major stable source for remote sensing optical imagery and widely used in many applications [11–13].

Vegetation and Environment monitoring on a New Micro-Satellite (VEN μ S) is a French-Israeli satellite providing a very high revisiting frequency of 2 days on a selection of 125 sites around the world, with constant viewing angles since 2017 [14,15]. VEN μ S provides 5 meter observations in 12 spectral bands, among which traditional visible bands, red-edge and near infra-red bands closely match Sentinel-2 spectral bands, as shown in figure 1. Note that despite its very large bandwidth that does not really match any band in VEN μ S, we chose to keep the Sentinel-2 B8 band, since it is the only Near Infra Red domain Sentinel-2 band sampled at 10 meter resolution.




Figure 1. Spectral Sensitivity Response of corresponding spectral bands between Sentinel-2 (top) and VENµS (bottom).

2.2. Product levels and processing

Remote Sensing Imagery products come in different levels depending on the applied processing. End-users usually work with products of at least level 2A, which includes, among other things:

- Ortho-rectification (geometric processing)
- Conversion of radiance to Surface Reflectance values, including estimation and compensation of aerosol content and water vapor amount
- Screening of clouds and cloud shadows

Though ESA delivers L2A products from their own processor Sen2Corr, in this work, we used L2A products generated by the MAJA open-source processing chain developed by CNES and CESBIO [16]. The rationale behind this is that both VENµS archive and a subset of the Sentinel-2 archive is produced by CNES at level 2A with MAJA, and distributed on Theia portal¹. Therefore, VENµS and Sentinel-2 images used in our dataset were produced by the exact same algorithms and code, which enforces the coherence between both products. Table 1 summarizes corresponding spectral bands between Sentinel-2 and VENµS that are sampled in the SEN2VENµS dataset. For both sensors, surface reflectance with adjacency effect compensation have been used.

Both sensors offer the same range of validity flags, namely:

- A mask of no-data pixels, which are out of the sensor swath
- A mask of clouds and clouds shadows
- A mask of saturated pixels
- A mask of geophysically invalid pixels (water, out of sight pixels due to relief, etc...)

From all those masks, we derive a single validity mask for each product, by taking their union.

¹ <https://theia.cnes.fr>

Sentinel-2 L2A products distributed by Theia are subject to the Etalab Open Licence Version 2.0². VENμS L2A products distributed by Theia are subject to the Creative Commons BY-NC 4.0³.

Sentinel-2	10m bands B2 B3 B4 B8	20m bands B5 B6 B7 B8A
VENμS	5m bands B3 B4 B7 B11	5m bands B8 B9 B10 B11

Table 1. Corresponding Sentinel-2 and VENμS bands used in the SEN2VENμS dataset. Note that B11 is used as high resolution band for both B8 and B8A Sentinel-2 bands.

2.3. Site selection

Among the available VENμS sites, we selected 29 sites covered by the L2A production of Sentinel-2 in Theia, which gave the selection presented in figure 2. Some sites with data scarcely available due to cloud occlusion or insufficient overlap were discarded. Note that since there are many more VENμS sites available, the dataset can always be extended by producing the corresponding Sentinel-2 tiles to level 2A. Each site is identified by a unique name. The swath of the VENμS instrument is 27.8 kilometers, and most sites correspond roughly to a square area.

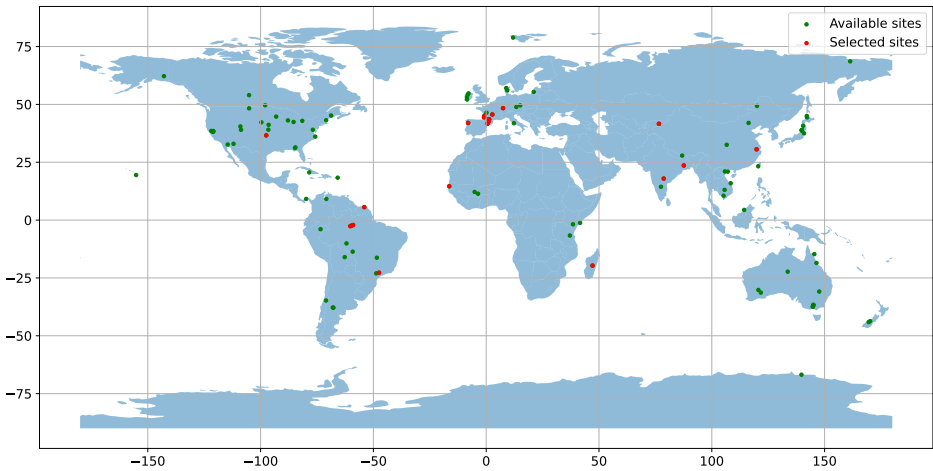


Figure 2. Map of available VENμS sites (green) and 29 selected sites (red) for the dataset.

It is important to note that if the viewing angles of VENμS are guaranteed to be constant for a given site, due to its limited swath, and to acquire as many sites as possible, most sites are acquired under a large viewing angle. As a result at least half of the selected sites have a zenith viewing angle greater than 25°, as can be seen in figure 3. On the other hand, Sentinel-2 field of view is 21°, always looking at nadir, so that the maximum zenith viewing angle is around 11°. Moreover, higher angles will also yield parallax effects, which will modify the apparent position of objects above ground. For this reason, VENμS sites with higher viewing angles might exhibit more difference with respect to the corresponding Sentinel-2 image.

2.4. Pair selection

For each selected site, we harvested the whole Theia archive in order to select pairs of VENμS and Sentinel-2 dates acquired on the same day, in order to minimize changes in image content. This leads to a variable number of pairs across sites, as shown in figure 4.

² <https://theia.cnes.fr/atdistrib/documents/Licence-Theia-CNES-Sentinel-ETALAB-v2.0-en.pdf>
³ <https://creativecommons.org/licenses/by-nc/4.0/>




Figure 3. Zenith viewing angles for the 29 selected VENµS sites

This can be explained by the fact that both the Level1 and the MAJA processors do not produce dates that are estimated as fully cloudy. We can see that sites such as FGMANAUS or ATTO, in Amazone state, Brazil, have fewer pairs than other sites. Some sites were also redesigned or split during the course of the mission, leading to artificially fewer acquisitions. A total of 579 pairs of images have been selected, with a maximum of 39 pairs for ARM site in Oklahoma state, USA, and a minimum of 4 pairs for FGMANAUS in Amazone state, Brazil.




Figure 4. Repartition of the 579 selected pairs across selected VENµS sites, sorted by increasing zenith viewing angle.

Figure 5 shows the distribution of acquisition dates for each site. We can see that a wide variety of seasons is represented, through several years, from late 2019 to September 2020.




Figure 5. Distribution of acquisition dates of selected pairs for each site. Colors are used to increase readability.

2.5. Sampling patches in pairs

Each pair of VENµS and Sentinel-2 images identified in section 2.4 underwent the same patch sampling procedure, which we further describe in this section.

2.5.1. Reprojection and common bounding box cropping144

Sentinel-2 images layout follow the MGRS tile system, where each tile correspond to145
a specific area using the local Universal Transverse Mercator (UTM) projection. VENμS146
data also follows the local UTM projection but with a grid specific to each site. A first step147
is therefore to determine the common bounding box between the Sentinel-2 and VENμS148
images.149

2.5.2. Spatial registration150

It is well known that remote sensing sensors may suffer from absolute location noise.151
According to ESA data quality reports, Sentinel-2 absolute location accuracy is 11 meters152
for 95.5% of the products. Of course this is prior to the new GRI processing that has been153
set-up in April 2021, but since the reprocessing is still undergoing, most products in Theia154
archive will have this accuracy, which corresponds to 2 pixels from VENμS. On VENμS155
side, partly because of the narrower field of view, accuracy is measured to be 3 meters for156
93% of the images [17]. Since the aim of the SEN2VENμS dataset is to serve for the training157
of super-resolution algorithms, correct alignments of images is of great importance. We158
therefore implemented the following procedure in order to register the Sentinel-2 image159
onto the VENμS image of a given pair.160

We estimate a global translation between the two images. This translation is estimated161
on the blue channel of both sensors, at 10 meters resolution, which means that VENμS162
image is first downsampled by a factor of 2 using a bicubic kernel.163

To estimate the translation, we leveraged the SIFT key-points detection and matching164
algorithms [18] and use it in a similar way as in [19]:165

1. Divide the downsampled VENμS image and the Sentinel-2 in non-overlapping corre-166
sponding patches of 366×366 pixels167
2. For each patch, compute SIFT matches168
3. Discard matches that are masked by the respective validity masks169
4. Discard matches that are further than 15 meters apart (obvious outliers)170
5. Compute the average shift in both directions from the collection of remaining matches171

The idea behind this process is that while individual SIFT matches are highly unreli-172
able, estimating a 2 parameter transform from a collection of several thousands of points is173
reliable. In all cases, because of step 4), the applied shift magnitude will not be greater than174
15 meters.175

Once the average shift is estimated, it is applied to Sentinel-2 image with bicubic176
resampling, in order to get registered Sentinel-2 images at both 10 meters and 20 meters.177

2.5.3. Patchification and invalid patch filtering178

Once data have been registered, images are divided into corresponding triplets of179
patches. All 5 meter bands from VENμS are divided into patches of 256×256 pixels.180
Consequently the corresponding patches of Sentinel-2 10 meter bands have a shape of 128
 $\times 128$ pixels, and the corresponding patches of Sentinel-2 20 meter bands have a shape of182
 64×64 pixels.183

If any pixel of any patch in the triplet (either VENμS 5 meters, Sentinel-2 10 meters184
or Sentinel-2 20 meters) is invalid according to the corresponding validity mask, due to a185
detected cloud for instance, the triplet is entirely discarded, so that after this process all186
patches are composed of valid pixels only.187

2.5.4. Radiometric adjustments188

Radiometric adjustments are required to compensate for differences in spectral sensi-189
tivity response, sensor calibration, atmospheric effects compensation and viewing angle,190
and increase the radiometric coherency of source and target images. Because those param-191
eters can vary between sites and pairs, we perform a separate adjustment for each pair. Since192
our aim is to provide a dataset for the super-resolution of Sentinel-2 images, we adjust193
VENμS radiometries to increase their coherency with Sentinel-2. In a previous work [20],194

we showed that linear regression was sufficient to compensate for discrepancies between surface reflectances from different sensors with similar spectral bands.

For each pair, we therefore perform a linear least-square radiometric fit between all source and target bands, separately for the group of 10 meter bands and the group of 20 meter bands (see table 1). This is described in equation 1 for the 10 meter bands, where W is a 5×4 containing the linear regression weights, V is the source matrix containing n rows of VENμS surface reflectances from n randomly selected pixels, and S is the target matrix containing n rows of Sentinel2 surface reflectances in the corresponding bands. The equation holds separately for Sentinel-2 20 meter bands.

VENμS patches are downsampled beforehand, so that the fit is performed at Sentinel-2 resolution. Prior to fitting, patches with a global RMSE greater than 0.2 reflectance count are discarded, in order to avoid obvious outliers.

$$W^* = \underset{W}{\operatorname{argmin}} ||VW - S||_2^2$$

$$W \in \mathbb{R}^{5 \times 4}$$

$$V = \begin{bmatrix} 1 & \rho_{1, \text{venus}, b2} & \rho_{1, \text{venus}, b4} & \rho_{1, \text{venus}, b7} & \rho_{1, \text{venus}, b11} \\ 1 & \rho_{2, \text{venus}, b2} & \rho_{2, \text{venus}, b4} & \rho_{2, \text{venus}, b7} & \rho_{2, \text{venus}, b11} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & \rho_{n, \text{venus}, b2} & \rho_{n, \text{venus}, b4} & \rho_{n, \text{venus}, b7} & \rho_{n, \text{venus}, b11} \end{bmatrix}$$

$$S = \begin{bmatrix} \rho_{1, \text{sentinel2}, b2} & \rho_{1, \text{sentinel2}, b3} & \rho_{1, \text{sentinel2}, b4} & \rho_{1, \text{sentinel2}, b8} \\ \rho_{2, \text{sentinel2}, b2} & \rho_{2, \text{sentinel2}, b3} & \rho_{2, \text{sentinel2}, b4} & \rho_{2, \text{sentinel2}, b8} \\ \vdots & \vdots & \vdots & \vdots \\ \rho_{n, \text{sentinel2}, b2} & \rho_{n, \text{sentinel2}, b3} & \rho_{n, \text{sentinel2}, b4} & \rho_{n, \text{sentinel2}, b8} \end{bmatrix} \quad (1)$$

Note that VENμS B11 band is processed twice, and fitted separately to the wide Near Infra-Red B8 channel at 10m and to the narrow Near Infra-Red B8A Sentinel-2 channel at 20m (see figure 1). This allows to build consistent 5 meter references for those two bands, as can be seen in sample patches (see figure 10).

2.5.5. Random selection and outlier removal

We perform a random selection of at most 500 patches for each pair, which in most case samples all the valid patches, since a squared VENμS site should contain at most 484 patches. This allows to limit the imbalance with respect to sites with a longer segment.

After this selection, we perform a final outlier removal. This removal has been added afterward, after finding some patches with spurious differences, such as unmasked cloud shadows or specular reflections. In order to eliminate those patches, we discard any patch for which the RMSE between Sentinel-2 10 meter band and their corresponding VENμS bands downsampled at 10 meters is greater than 0.02. This operation discards approximately 6.75% of the patches.

3. Dataset content

In this section, we further describe the characteristics of the generated dataset.

3.1. Quantitative analysis

In total, the SEN2VENμS dataset contains 132 955 patches. Figure 6 presents the number of patches sampled from each site. We can observe that only 5 sites have more than 8k patches, and that only the ARM site has more than 12k patches. On the other hand, 16 sites have less than 4k patches sampled. While this imbalance might be seen as problematic for the training of some algorithms, we must stress that this imbalance is different in nature than classical class imbalance problem in classification tasks: what really matters is variability and equity of different kinds of landscape, which can not be reduced

to sites. For instance several sites in South West of France will exhibit the same kind of landscape.




Figure 6. Total number of patches sampled from each site

Figure 7 allows to dig further into the sampling of each site. We can observe that if the maximum number of patches per pair reaches our 500 limit for most sites, the minimum is almost always very small, because of pairs suffering from heavy cloud cover. Other sites, such as ARM in Oklahoma, USA, ESGIS-B3 in West of France or LERIDA-1 in Spain show a median number of patches per pair close to the maximum, which denotes a fair sampling across pairs.




Figure 7. Statistics of number of patches per pair for each site

One important thing to know is wether VEN μ S 5 meter patches are faithful to Sentinel-2 radiometry in each band, so that we do not capture radiometric biases in the course of a learning process involving the dataset. In order to answer this question, we computed Mean Average Error and Root Mean Square Error in each band for a random selection of 200 patches involving at most 10 different pairs for each site, by first downsampling VEN μ S patches to the Sentinel-2 corresponding resolution (either 10m or 20m depending on the band). Downsampling is achieved by first convolving the 5 meter image with a gaussian spatial kernel whose standard deviation is tuned to the known values of the Modulation Transfer Function at Nyquist rate for Sentinel-2, and then decimating to the target resolution. Those metrics are presented in figures 8 and 9.

First thing to note from those two figures is that for all bands and all sites, Mean Average Error is lower than the expected absolute surface reflectance error specification of the MAJA algorithm, which is set to 0.01 for all bands. We can therefore say that our dataset has a good spectral consistency between Sentinel-2 and VEN μ S, for all bands. We can also observe that Sentinel-2 B8 band is often the band with the higher MAE and RMSE, which can be explained by its lesser match with VEN μ S band. Nevertheless, being the only 10 meter band in the Near Infra-Red, it seemed reasonable to keep it in the dataset. We can also observe a tendency of the MAE error to increase with the zenith viewing angle, which may be explained by higher BRDF and parallax effects. This trend is not confirmed for all sites however, as site BENGA, located in India, has almost the highest viewing angles and better performances than its high angle siblings. On the other hand, we can observe that the Root Mean Square Error does not show a clear trend related to the viewing angle, and that Root Mean Square Error is quite stable across sites for a given spectral band. It is worth noting that all bands have a RMSE below 0.03 for all sites.




Figure 8. Mean Average Error per band and per site computed on a random selection of 200 patches from at most 20 pairs, at Sentinel-2 resolution.




Figure 9. Root Mean Square Error per band and per site computed on a random selection of 200 patches from at most 20 pairs, at Sentinel-2 resolution.

3.2. Qualitative analysis

Figure 10 shows one sample patch per site, drawn at random. We can observe the variety of landscapes covered by the dataset, including natural, semi-natural and urban areas, as well as forest and shorelines. We can also observe the spectral consistency between Sentinel2 (odd columns) and Venus (even columns) patches. Note the high consistency of the 5 meter B8 (columns 4 and 12) and B8A (columns 8 and 16) bands with their Sentinel-2 10 meter (columns 3 and 11) and 20 meter (columns 7 and 15) corresponding bands, though being generated from the same B11 VENμS band. The wider bandwidth B8 has less distinctive features, whereas B8 accurately responds to vegetated areas, and this difference is also visible in the 5 meter reference patches.

Despite a very high consistency between VENμS and Sentinel-2 patches, local discrepancies remain between the 5 meter VENμS patch and its corresponding 10 meter or 20 meter Sentinel-2 patch. For instance, when comparing A9 with A10, we can see cars on the road on the VENμS patch that have no match in the Sentinel-2 patch. Other differences include changes in water surfaces (C1 vs. C2), and artifacts related to viewing angles (shadowed area in G1 appear larger than in G2). Nevertheless, we must stress that those discrepancies are expected when using 2 different sensors as we can not filter out or model all the differences between sensors specifications and viewing conditions. This dataset can be used as a real world case for algorithms, which should be resilient to those discrepancies.

3.3. Format and distribution

The dataset is composed of separate sub-datasets, one for each site. For each site, the sub-dataset folder contains a set of files for each date, following this naming convention as the pair id: {site_name}_{mgrs_tile}_{acquisition_date}. For each pair, 5 files are available, as shown in table 2. Patches are encoded as ready-to-use tensors as serialized by the well known Pytorch library [21]. As such they can be loaded by a simple call to the torch.load() function. Note that bands are separated into two groups (10m and 20m Sentinel2 bands), which leads to four separate tensor files (2 groups of bands × source and target resolution). Tensor shape is [n, c, w, h] where n is the number of patches, c = 4 is the number of bands, w is the patch width and h is the patch height. In order to save storage




Figure 10. Examples of patches. From left to right: columns 1-8 and 9-16 show rendering of two different patches. Columns 1 and 9: B4, B3, B7 (RGB natural) at 10m, columns 2 and 10: B4, B3, B7 (RGB natural) at 5m, columns 3 and 11: color-mapped B8 at 10m (Wide Near Infra-Red), columns 4 and 12 : color-mapped B8 at 5m (Wide Near Infra-Red), columns 5 and 13 : B7, B6, B5 color composition (Red Edge 3 to 1) at 20m, columns 6 and 14: B7, B6, B5 color composition (Red Edge 3 to 1) at 5m, columns 7 and 15: color-mapped B8A at 20m (Narrow Near Infra-Red), columns 8 and 16: color-mapped B8A at 5m (Narrow Near Infra-Red). 29 patches are displayed, one random patch for each site. Only 64×64 pixels crops of the patches are displayed in order to improve readability. High resolution and low resolution patches radiometries where scaled to 8 bits with the same scaling factors.

space, they are encoded as 16 bits signed integers and should be converted back to floating point surface reflectance by dividing each and every value by 10 000 upon reading.

File	Content
{id}_05m_b2b3b4b8.pt	5m patches (256 × 256 pix.) for S2 B2, B3, B4 and B8
{id}_10m_b2b3b4b8.pt	10m patches (128 × 128 pix.) for S2 B2, B3, B4 and B8
{id}_05m_b5b6b7b8a.pt	5m patches (256 × 256 pix.) for S2 B5, B6, B7 and B8A
{id}_20m_b5b6b7b8a.pt	20m patches (64 × 64 pix.) for S2 B5, B6, B7 and B8A
{id}_patches.gpkg	GIS file with footprint of each patch

Table 2. Naming convention for files associated to each pair. {id} is {site_name}_{mgrs_tile}_{acquisition_date}.

Each file comes with a master index.csv CSV (Comma Separated Values) file, with one row for each pair sampled in the given site, and columns as described in table 3, separated with tabs.

Column	Description
venus_product_id	ID of the sampled VENμS L2A product
sentinel2_product_id	ID of the sampled Sentinel-2 L2A product
tensor_05m_b2b3b4b8	Name of the 5m tensor file for S2 B2, B3, B4 and B8
tensor_10m_b2b3b4b8	Name of the 10m tensor file for S2 B2, B3, B4 and B8
tensor_05m_b5b6b7b8a	Name of the 5m tensor file for S2 B5, B6, B7 and B8A
tensor_20m_b5b6b7b8a	Name of the 20m tensor file for S2 B5, B6, B7 and B8A
s2_tile	Sentinel-2 MGRS tile
vns_site	Name of VENμS site
date	Acquisition date as YYYY-MM-DD
venus_zenith_angle	VENμS zenith viewing angle in degrees
patches_gpkg	Name of the GIS file with footprint for each patch
nb_patches	Number of patches for this pair

Table 3. Columns of the index.csv file indexing pairs for each site. For file naming conventions, refer to table 2.

Figure 11 shows the size in Gigabytes for each site in the dataset. Altogether the dataset weights 1116 Gb. Each site is compressed into a separate archive.




Figure 11. Uncompressed files sizes for each site in Gigabytes. The full dataset weights 116 Gb.

The SEN2VENμS dataset is distributed on Zenodo⁴ [22]. Files {id}_05m_b2b3b4b8.pt and {id}_05m_b5b6b7b8a.pt are distributed under the the original licence of the Sentinel-2 Theia L2A products, which is the Etalab Open Licence Version 2.0². Files {id}_05m_b2b3b4b8.pt and {id}_05m_b5b6b7b8a.pt are distributed under the original licence of the VENμS products, which is Creative Commons BY-NC 4.0³. Remaining files are distributed under the Creative Commons BY 4.0⁵ licence.

⁴ <https://zenodo.org/deposit/6514159>
⁵ <https://creativecommons.org/licenses/by/4.0/>

Note that even if the SEN2VEN μ S dataset is sorted by sites and by pairs, we strongly encourage users to apply the full set of machine learning best practices when using it : random keeping separate pairs (or even sites) for testing purpose, and randomization of patches accross sites and pairs in the training and validation sets.

4. Conclusion

In this paper, we propose the first dataset dedicated to the training of super-resolution algorithms for Sentinel-2, the ESA well known earth observation mission. This dataset is composed of 132 955 patches sampled from 29 sites including desertic, temperate, sub-tropical and tropical landscapes, with several acquisition dates per site spanning the 2019 - 2020 period. The dataset is composed of Sentinel-2 10m bands B2, B3, B4 and B8, as well as Sentinel-2 20m bands B5, B6, B7, B8A, with corresponding spatially registered, spectrally adjusted 5m resolution VEN μ S images acquired on same date and with same spectral bands.

As any dataset composed of real world data, the SEN2VEN μ S dataset has some limitations and discrepancies that we try to identify and analyse in this paper. First, most sites are acquired from high viewing angles at 5 meters (greater than 20°) whereas Sentinel-2 is acquired under at most 12°. This might cause differences in radiometry because of BRDF effects, and local mis-registration because of parallax effects. In addition, even if acquisitions happened the same day, there might a difference in acquisition time of at most around 1 hour due to the differences in orbits and viewing angles. Nevertheless, statistics show a very strong radiometric consistency between 5m and corresponding 10m or 20m patches. A second limitation is the imbalance of number of dates and number of patches across sites. Though this can not directly be related to an imbalance in landscape variability, extra care must be taken during training to avoid specialization on most represented sites. Third, as any real world dataset, and despite all precautions taken during its preparation, additional discrepancies may occur for some individual patches, including rapid change of landscape between the low and high resolution acquisition and spurious image quality artifacts. However we would like to stress that any super-resolution method aiming at operational use must be able to cope with such discrepancies.

It is also interesting to note that a smaller version of the SEN2VEN μ S dataset has been used in the Sentinel-HR study [23] in order to train the CARN super-resolution network [24].

Finally, it is interesting to note that the VEN μ S mission has recently entered the VM5 stage, a new phase on a different orbit, where a new set of sites are acquired with a daily revisit and a spatial resolution of 4 meter. In a few months, when the archive will be populated with those new products, this could be a promising extension of the super-resolution dataset presented in this paper.

1. Chan, T.; Wong, C.K. Total variation blind deconvolution. *IEEE Transactions on Image Processing* **1998**, *7*, 370–375. doi:10.1109/83.661187.

2. Krishnan, D.; Tay, T.; Fergus, R. Blind deconvolution using a normalized sparsity measure. In *Proceedings of the CVPR 2011*, 2011, pp. 233–240. doi:10.1109/CVPR.2011.5995521.

3. Anwar, S.; Khan, S.; Barnes, N. A deep journey into super-resolution: A survey. *ACM Computing Surveys (CSUR)* **2020**, *53*, 1–34.

4. Liu, H.; Qian, Y.; Zhong, X.; Chen, L.; Yang, G. Research on super-resolution reconstruction of remote sensing images: a comprehensive review. *Optical Engineering* **2021**, *60*, 100901.

5. Agustsson, E.; Timofte, R. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *Proceedings of the Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2017, pp. 126–135.

6. Shoeiby, M.; Robles-Kelly, A.; Wei, R.; Timofte, R. Pirm2018 challenge on spectral image super-resolution: Dataset and study. In *Proceedings of the Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, 2018, pp. 0–0.

7. Wang, Y.; Wang, L.; Yang, J.; An, W.; Guo, Y. Flickr1024: A large-scale dataset for stereo image super-resolution. In Proceedings of the Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops, 2019, pp. 0–0.

8. Märtens, M.; Izzo, D.; Krzic, A.; Cox, D. Super-resolution of PROBA-V images using convolutional neural networks. *Astrodynamics* **2019**, *3*, 387–402.

9. Rohith, G.; Kumar, L.S. Paradigm shifts in super-resolution techniques for remote sensing applications. *The Visual Computer* **2021**, *37*, 1965–2008.

10. Drusch, M.; Del Bello, U.; Carlier, S.; Colin, O.; Fernandez, V.; Gascon, F.; Hoersch, B.; Isola, C.; Laberinti, P.; Martimort, P.; et al. Sentinel-2: ESA’s optical high-resolution mission for GMES operational services. *Remote sensing of Environment* **2012**, *120*, 25–36.

11. Phiri, D.; Simwanda, M.; Salekin, S.; Nyirenda, V.R.; Murayama, Y.; Ranagalage, M. Sentinel-2 data for land cover/use mapping: A review. *Remote Sensing* **2020**, *12*, 2291.

12. Segarra, J.; Buchailot, M.L.; Araus, J.L.; Kefauver, S.C. Remote sensing for precision agriculture: Sentinel-2 improved features and applications. *Agronomy* **2020**, *10*, 641.

13. Misra, G.; Cawkwell, F.; Winkler, A. Status of phenological research using Sentinel-2 data: A review. *Remote Sensing* **2020**, *12*, 2760.

14. Ferrier, P.; Crebassol, P.; Dedieu, G.; Hagolle, O.; Meygret, A.; Tinto, F.; Yaniv, Y.; Herscovitz, J. VEN μ S (Vegetation and environment monitoring on a new micro satellite). In Proceedings of the 2010 IEEE International Geoscience and Remote Sensing Symposium. IEEE, 2010, pp. 3736–3739.

15. Dedieu, G.; Hagolle, O.; Karnieli, A.; Ferrier, P.; Crébassol, P.; Gamet, P.; Desjardins, C.; Yakov, M.; Cohen, M.; Hayun, E. VEN μ S: PERFORMANCES AND FIRST RESULTS AFTER 11 MONTHS IN ORBIT. In Proceedings of the IGARSS 2018 - 2018 IEEE International Geoscience and Remote Sensing Symposium, 2018, pp. 7756–7759. doi:10.1109/IGARSS.2018.8519207.

16. Lonjou, V.; Desjardins, C.; Hagolle, O.; Petrucci, B.; Tremas, T.; Dejus, M.; Makarau, A.; Auer, S. Maccs-atcor joint algorithm (maja). In Proceedings of the Remote Sensing of Clouds and the Atmosphere XXI. International Society for Optics and Photonics, 2016, Vol. 10001, p. 1000107.

17. Dick, A., R.J.L.R.A.P.S.C.S.D.G.H.O.B.J.P.B.R.; Moreau, A. VENS: Mission Characteristics, Final Evaluation of the First Phase and Data Production. *submitted to Remote Sensing* **2022**.

18. Lowe, G. Sift-the scale invariant feature transform. *Int. J* **2004**, *2*, 2.

19. Michel, J.; Sarrazin, E.; Youssefi, D.; Cournet, M.; Buffe, F.; Delvit, J.; Emilien, A.; Bosman, J.; Melet, O.; L’Helguen, C. A new satellite imagery stereo pipeline designed for scalability, robustness and performance. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences* **2020**, *2*, 171–178.

20. Michel, J.; Inglada, J. Learning Harmonised Pleiades and SENTINEL-2 Surface Reflectances. *The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences* **2021**, *43*, 265–272.

21. Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimeshein, N.; Antiga, L.; et al. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems* 32; Wallach, H.; Larochelle, H.; Beygelzimer, A.; d’Alché-Buc, F.; Fox, E.; Garnett, R., Eds.; Curran Associates, Inc., 2019; pp. 8024–8035.

22. Michel, J.; Vinasco-Salinas, J.; Inglada, J.; Hagolle, O. SEN2VEN μ S, a dataset for the training of Sentinel-2 super-resolution algorithms, 2022. doi:10.5281/zenodo.6514159.

23. Michel, J.; Hagolle, O.; Puissant, A.; Herrault, P.; Corpetti, T.; Nabucet, J.; Faure, J.F.; Maurel, P.; Lelong, C.; Berthier, E.; et al. Sentinel-HR Phase 0 Report. Research report, CNES - Centre national d’études spatiales ; CESBIO, 2022.

24. Ahn, N.; Kang, B.; Sohn, K.A. Fast, accurate, and lightweight super-resolution with cascading residual network. In Proceedings of the Proceedings of the European Conference on Computer Vision (ECCV), 2018, pp. 252–268.