

COP-GEN-Beta: Unified Generative Modelling of COPernicus Imagery Thumbnails

Supplementary Material

A. Appendix

In this section, we provide additional qualitative results and visualizations to further illustrate the capabilities of COP-GEN-Beta.

Figure 1 presents examples of unconditional synthesis, demonstrating the diversity and variability captured by COP-GEN-Beta across different modalities. These results highlight the model’s ability to capture the full spectrum of the training data distribution.

Figure 2 showcases an extended loop between the S2L2A and S1RTC modalities. This visualization provides insight into how the model maintains consistency across generative steps. Future work will focus on a quantitative assessment of image quality, particularly analyzing potential degradation over iterative generations.

For completeness, Figure 3 illustrates all possible conditioning combinations for COP-GEN-Beta. Each row represents an independent example, where the red boxes indicate the input modalities used to condition the generation of the remaining ones. The left side of the figure provides real images (both conditioned and unconditioned) for reference, while the right side displays the generated outputs. This visualization highlights the flexibility and robustness of COP-GEN-Beta in handling diverse cross-modal synthesis tasks, effectively leveraging many different input configurations to produce high-quality outputs.

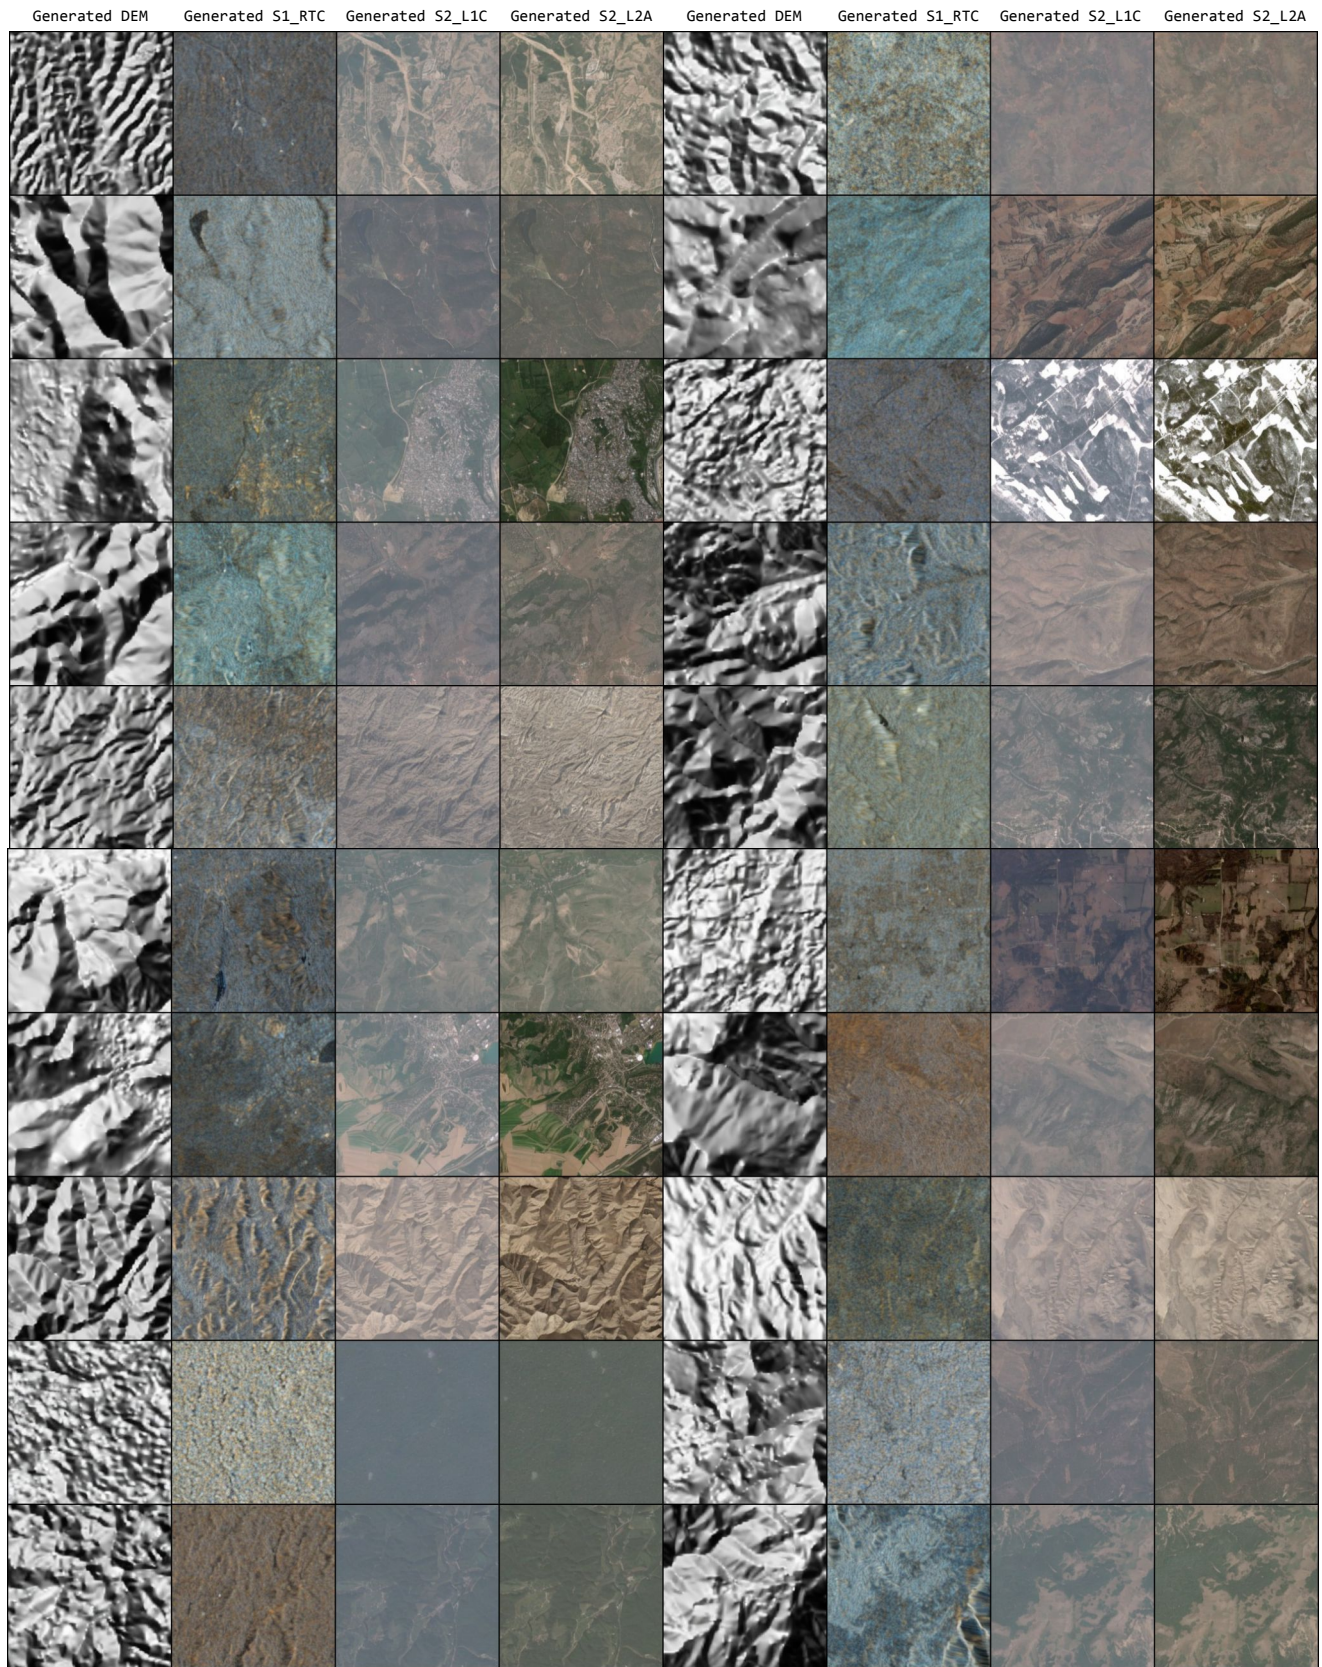


Figure 1. Unconditional joint generation of all modalities. Rows represent independent samplings (in groups of 4). The model generations exhibit high diversity, consistent with the rich and global data used for training.

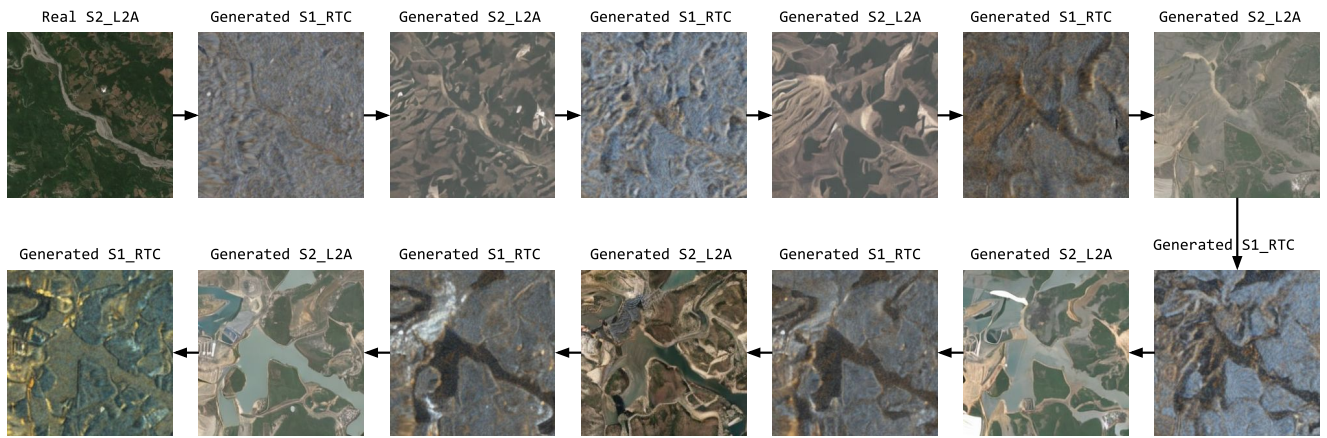


Figure 2. To analyse the model robustness to generation degradation, we perform a generation chain, starting from a real S2L2A image and iteratively conditioning the model on the previously generated image. For illustrative purposes, we only show a loop between S2L2A and S1RTC modalities repeatedly.

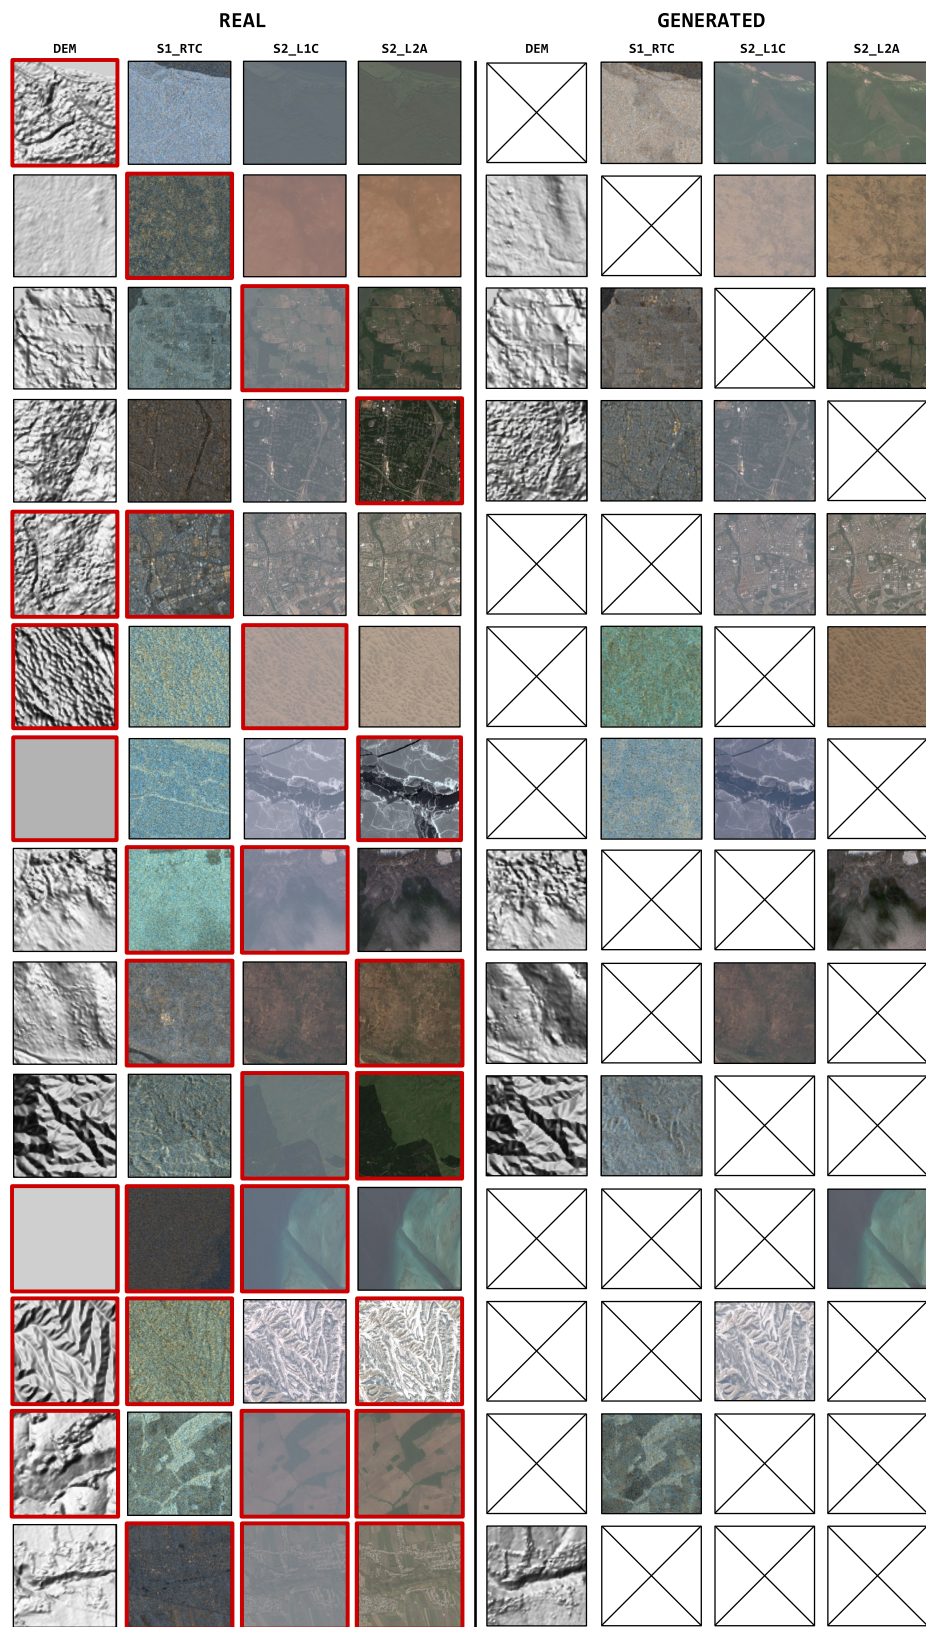


Figure 3. Illustration of all the possible combinations for conditioning COP-GEN-Beta. We show all the possible ways of conditioning our model. Each row represents an independent example. The red boxes (■) indicate the modalities used as input to condition the generation of the remaining modalities. The right side of the figure displays the generated outputs, while the left side includes the corresponding real images (both conditioned and unconditioned) for comparison.