

Accurate Georegistration of Point Clouds using Geographic Data

Supplemental Material

Chun-Po Wang Kyle Wilson Noah Snavely
Cornell University

{cpwang, wilsonkl, snavely}@cs.cornell.edu

1. Additional Alignment Results

This supplemental document shows additional alignment results for the datasets described in our main paper; these eight datasets cover a variety of places around the world, and include:

1. The Arts Quad dataset [1]
2. The Colosseum, Rome
3. Trafalgar Square, London
4. The Empire State Building, NYC
5. The Panthéon, Paris
6. Piazza del Popolo, Rome
7. Union Square, NYC
8. Top of the Rockefeller Center, NYC

For each dataset (i.e., each reconstructed structure from motion model), we compare the performance of alignment methods using different types of geographic information:

1. Using only the noisy geotags from the Flickr photos used to build the 3D SfM model
2. Using accurate geotags from registered Google Street View (GSV) panoramas, (or pre-alignment heuristics for Google Earth (GE) 3D model when nearby panoramas are not available), and
3. Refining the alignment by snapping to a Google Earth 3D model with ICP.

For each dataset, we first illustrate the overall alignment by showing our SfM model overlaid on an overhead image for each alignment method. In these visualizations, green dots represent estimated camera locations, and red dots represent a sampling of reconstructed 3D SfM points. Next, we show several photographs from each dataset, along with approximate geometric segmentations from different alignment methods as a further way of visualizing alignment accuracy. We also show a depth map, produced by rendering the Google Earth model from each photo's final viewpoint, for our final ICP-refined alignment. As in the paper, each geometric segmentation shows sky (red), vertical (yellow), and horizontal (blue) regions segmented from the Google Earth 3D mesh using the normal to each face (or the far depth buffer in the case of sky) to perform the segmentation. The segmentation results are overlaid on the photograph. These visualizations illustrate the quality of the alignment and segmentation in structures common to the photo and 3D model. Please zoom into the figures for best results.

1.1. Arts Quad

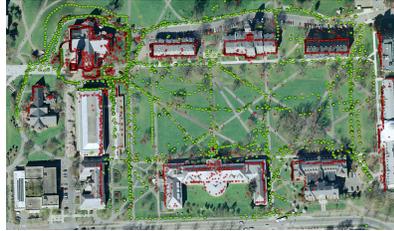
For the Arts Quad dataset, we obtained a good alignment using ICP (although the Google Street View-aligned model is also quite good). We observed that for the last example photograph below, the 3D building model in Google Earth has an incorrect altitude, and hence there is a discrepancy in the alignment to the photograph.

Overhead views of alignments (green dots: cameras, red dots: 3D SfM points)

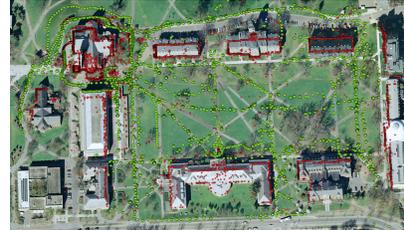
Georegistered



Georegistered with GSV data



Refined with ICP



Example photos and aligned view in Google Earth

Photo



Georegistered



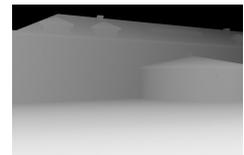
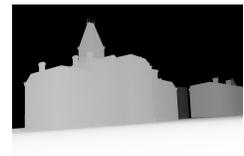
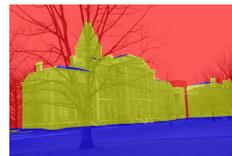
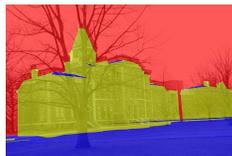
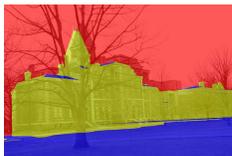
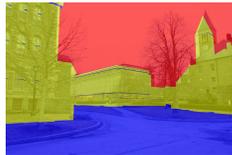
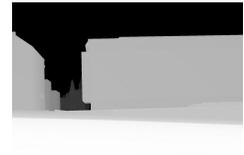
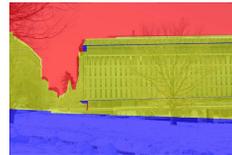
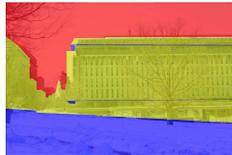
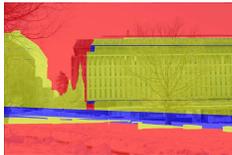
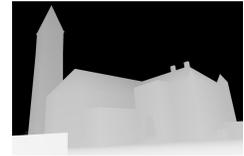
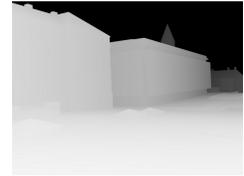
GSV data



Refined with ICP



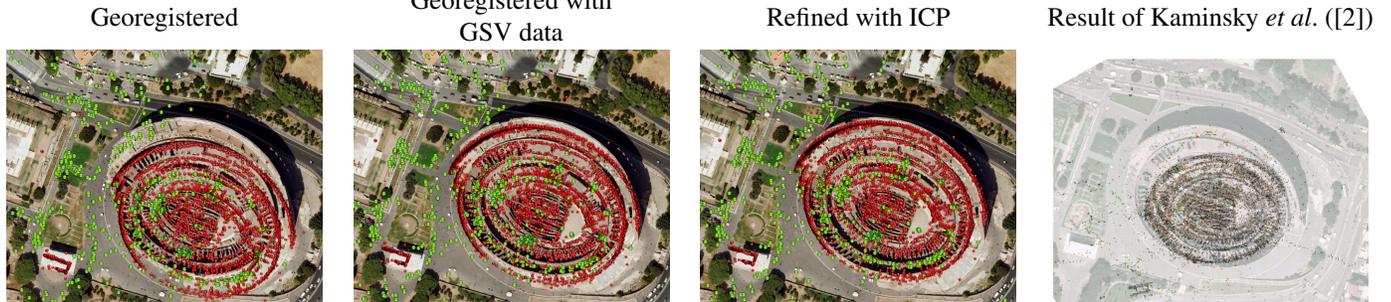
Depth map



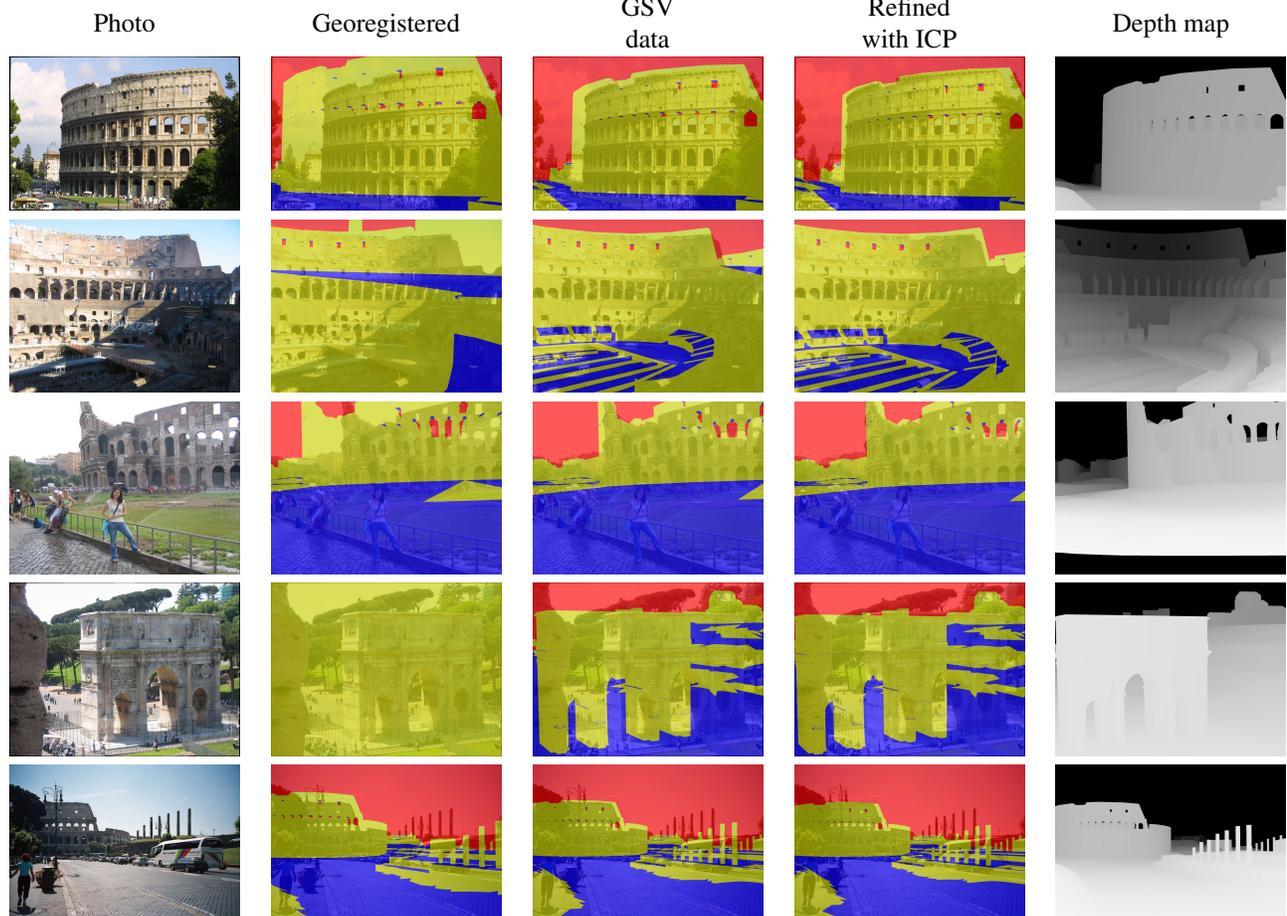
1.2. The Colosseum, Rome

For the Colosseum dataset, we also provide the overhead alignment result of Kaminsky *et al.* [2] for reference (note the misregistration). Our final ICP alignment for the Colosseum itself is visually quite accurate; however the photograph of the Arch of Constantine in the fourth example photo below does not align well with the Google Earth model. After further inspection, we believe this is an error in the SfM model, due to weak visual connections between images of the Colosseum and the Arch. The alignment in the last example photo also shows a discrepancy with the Google Earth model; we believe this is due partly to weak camera geometry, and partly to errors in the Google model. The incorrect segmentation of the ground is due to a noisy Google Earth terrain model near the area.

Overhead views of alignments



Example photos and aligned view in Google Earth



1.3. Trafalgar Square, London

This dataset contains large numbers of objects (fountains, sculptures, buildings, Big Ben) distributed at widely varying distances. Note that the overhead views are not exactly nadir; this causes some ambiguity in the building footprints in the aerial image. Big Ben (and other distant objects) are not covered by the Google Earth 3D model we downloaded, hence there is some missing geometry in segmentation and depth map for the first example image.

Overhead views of alignments

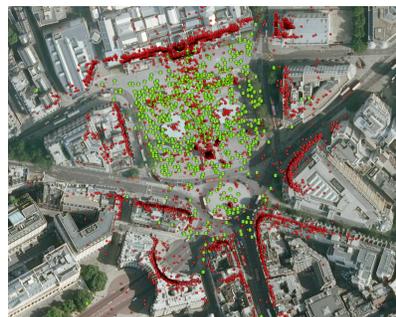
Georegistered



Georegistered with GSV data



Refined with ICP



Example photos and aligned view in Google Earth

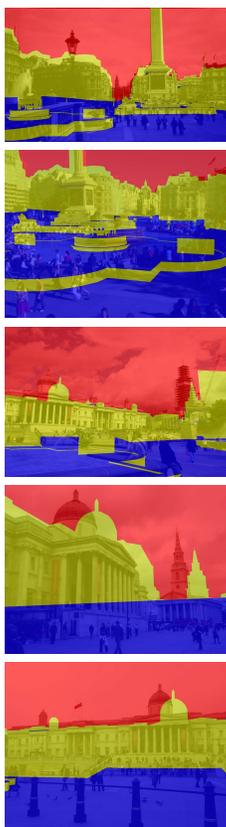
Photo



Georegistered



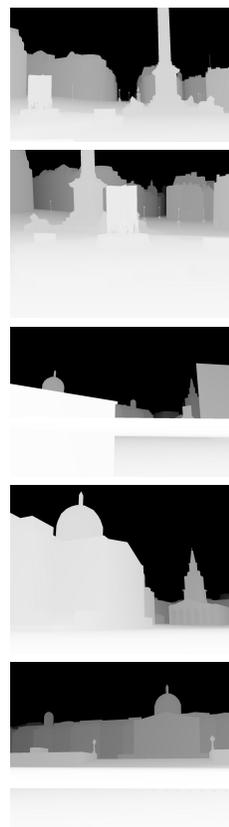
GSV data



Refined with ICP



Depth map



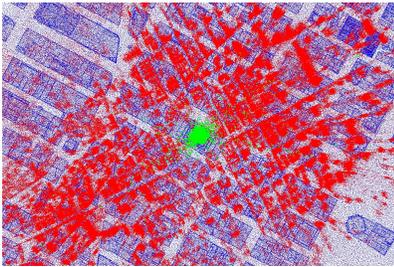
1.4. The Empire State Building, NYC

For this dataset there are no Google Street View (GSV) panoramas available, and the alignment using just georeferenced Flickr photos is incorrect in both orientation and scale, because the distances between cameras approach the scale of the error in the geotags. The alignment using our heading and scale matching heuristics with the Google Earth model already gives much better results, and ICP further improves the alignment. Note that our sky segmentations below are in general too large, because the 3D Google Earth model we downloaded is not large enough to cover the entire scope of this (many kilometer-scale) scene. For instance, the boundary of the Google Earth model can be seen in the first example.

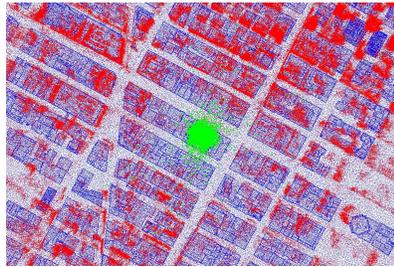
In the overhead alignment images, we show alignments to the GE model (rendered as sampled blue points) rather than aerial images because non-nadir aerial views available for this dataset make for a difficult visualization.

Overhead views of alignments

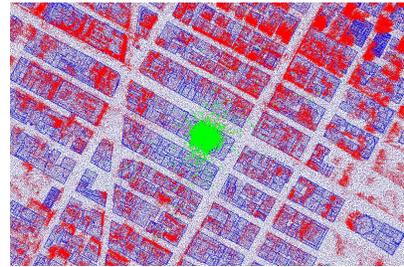
Georegistered



Georegistered with GE data



Refined with ICP



Example photos and aligned view in Google Earth

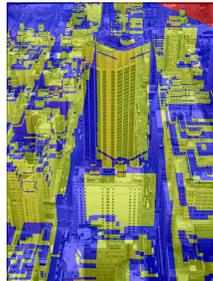
Photo



Georegistered



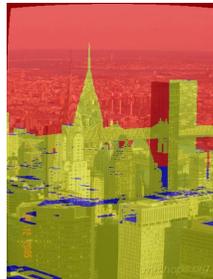
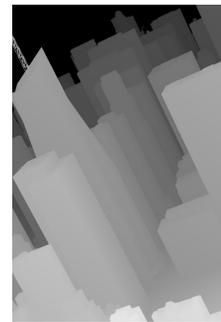
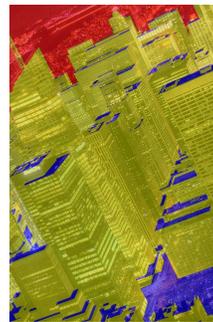
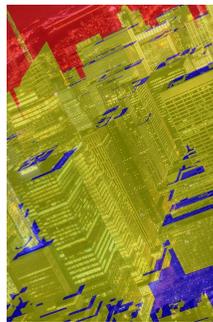
GE data



Refined with ICP



Depth map



1.5. The Panthéon, Paris

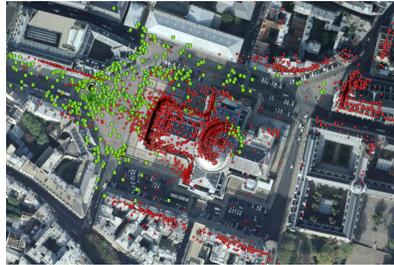
For this dataset, the use of Google Street View panoramas for alignment already achieves a good result, with ICP further refining the alignment. The last example shows a failure case, in which the estimated camera position is too far back from the building and behind another object in the GE model. In general, however, the alignments between image and model are surprisingly tight. (We observed that the Google Earth models for Paris are generally of high quality.)

Overhead views of alignments

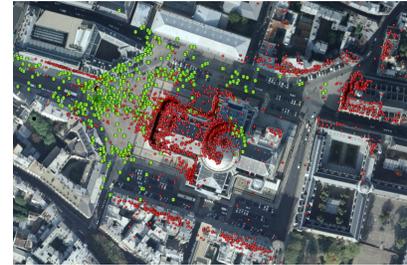
Georegistered



Georegistered with GSV data



Refined with ICP



Example photos and aligned view in Google Earth

Photo



Georegistered



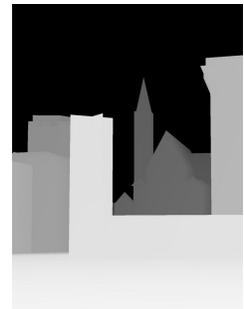
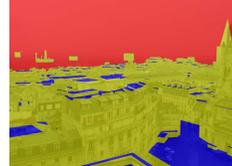
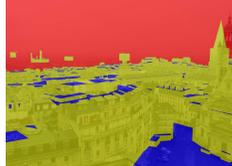
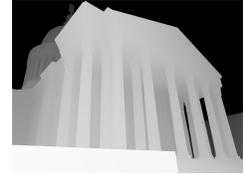
GSV data



Refined with ICP

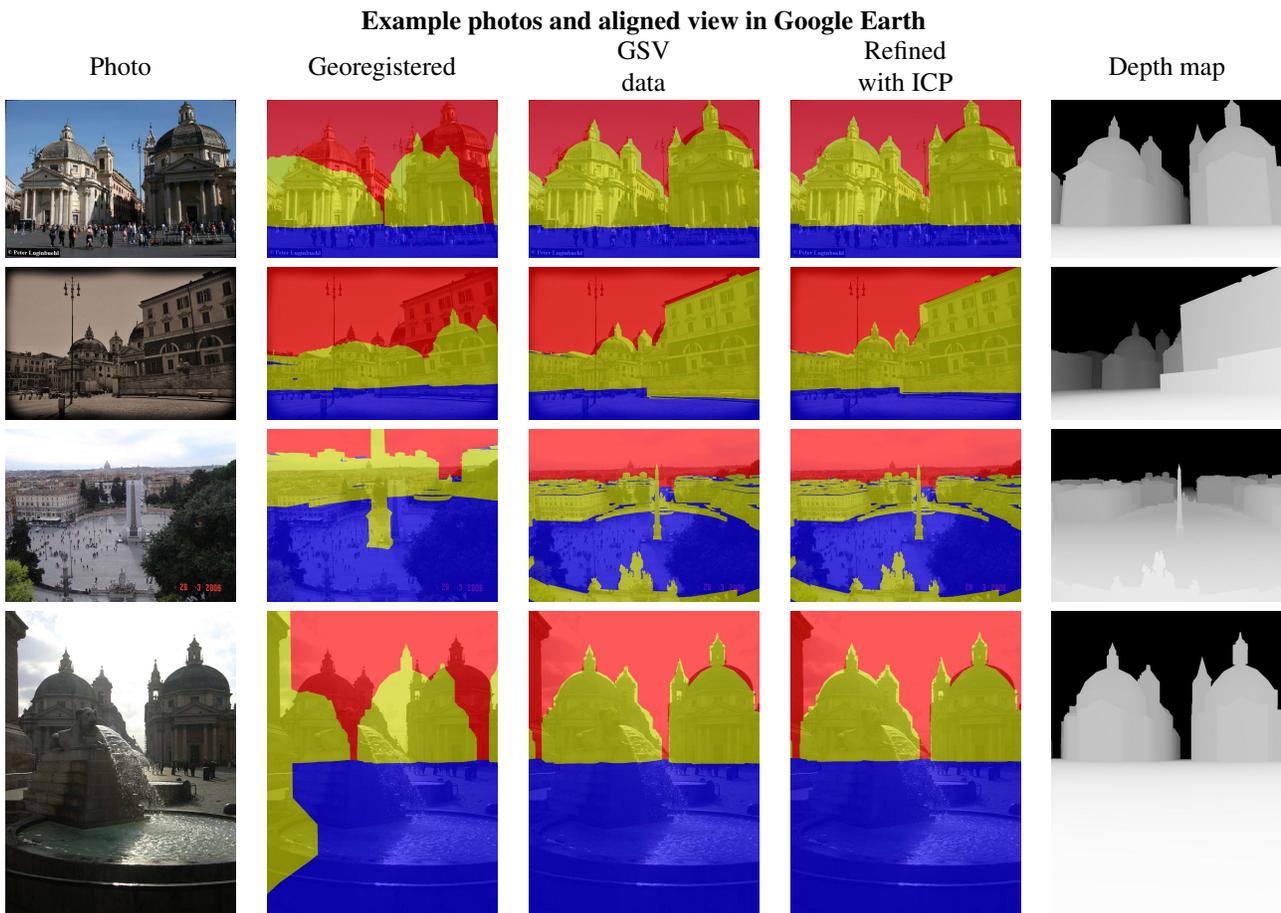
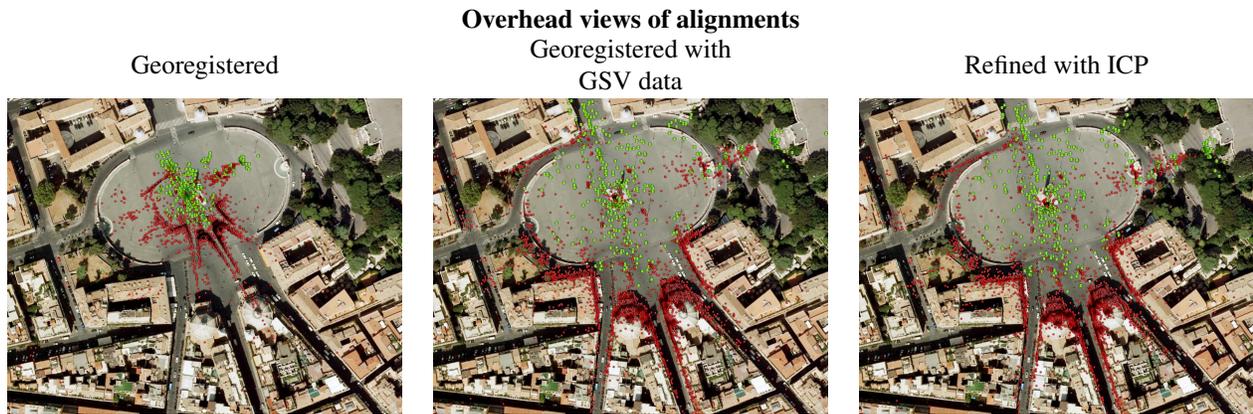


Depth map



1.6. Piazza del Popolo, Rome

We also achieved a reasonably good final alignment in this dataset, especially for the twin Santa Maria di Montesanto and Santa Maria dei Miracoli churches. Note that background city in the third example, and the fountain in the last photo, are not present in the downloaded 3D model.



1.7. Union Square, NYC

In the fourth photo below, the distant buildings are again not in the downloaded 3D model (but this could again be fixed by downloading a Google Earth model for a larger area). In the last example, the Google Street View alignment moves the camera behind another object, and thus the segmentation is erroneous.

Overhead views of alignments

Georegistered



Georegistered with GSV data



Refined with ICP



Example photos and aligned view in Google Earth

Photo



Georegistered



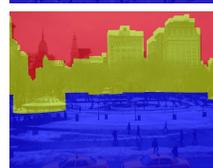
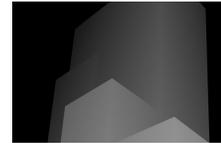
GSV data



Refined with ICP

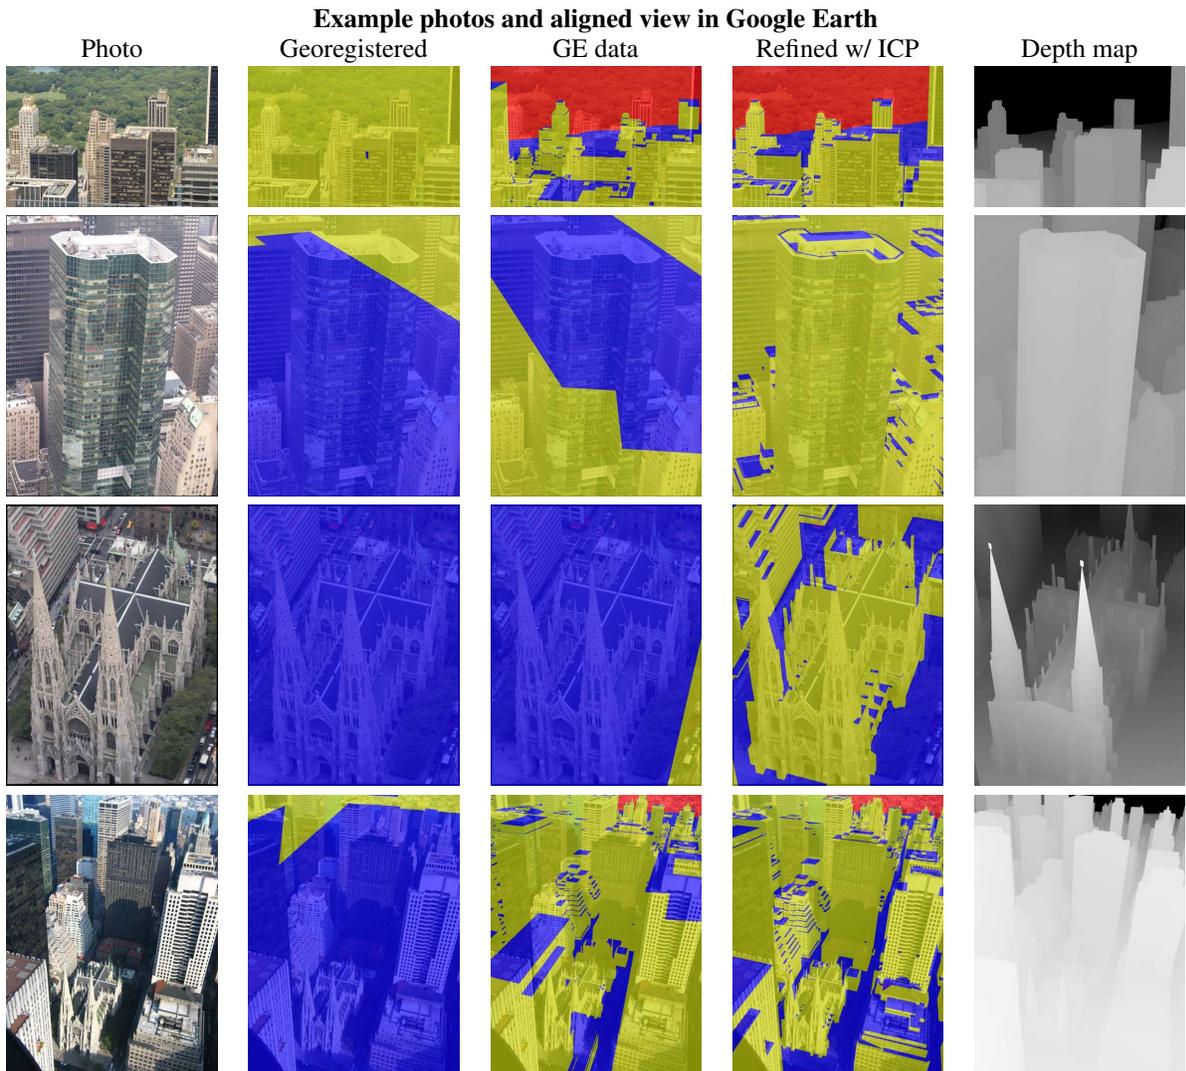
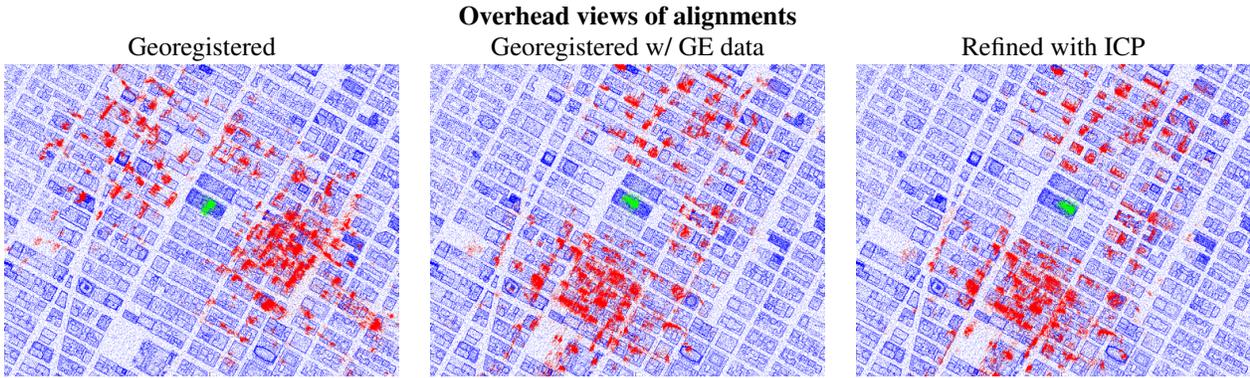


Depth map



1.8. Top of Rockefeller Center, NYC

Like the Empire State Building dataset above, there are no Google Street View panoramas available for the Top of the Rock, and our alignment using just georeferenced photos is incorrect in both orientation (by around 90 degrees) and scale. The alignment using heuristics to match the SfM model to the Google Earth model already gives a much better result, but ICP further refines the alignment and yields better segmentations and depth maps. Note the sky segmentations are again too large because our 3D model is not large enough to cover the entire scope of this scene.



References

- [1] D. Crandall, A. Owens, N. Snavely, and D. Huttenlocher. Discrete-continuous optimization for large-scale structure from motion. In *CVPR*, 2011.
- [2] R. S. Kaminsky, N. Snavely, S. M. Seitz, and R. Szeliski. Alignment of 3d point clouds to overhead images. In *Proc. CVPR Workshop on Internet Vision*, 2009.