Lifting GIS Maps into Strong Geometric Context

Raúl Díaz, Minhaeng Lee, Jochen Schubert, Charless C. Fowlkes
Computer Science Department, University of California, Irvine
{rdiazgar,minhaenl,j.schubert,fowlkes}@uci.edu

Abstract

Contextual information can have a substantial impact on the performance of visual tasks such as semantic segmentation, object detection, and geometric estimation. Data stored in Geographic Information Systems (GIS) offers a rich source of contextual information that has been largely untapped by computer vision. We propose to leverage such information for scene understanding by combining GIS resources with large sets of unorganized photographs using Structure from Motion (SfM) techniques. We present a pipeline to quickly generate strong 3D geometric priors from 2D GIS data using SfM models aligned with minimal user input. Given an image resectioned against this model, we generate robust predictions of depth, surface normals, and semantic labels. We show that the precision of the predicted geometry is substantially more accurate other single-image depth estimation methods. We then demonstrate the utility of these contextual constraints for re-scoring pedestrian detections, and use these GIS contextual features alongside object detection score maps to improve a CRF-based semantic segmentation framework, boosting accuracy over baseline models.

1. Introduction

In recent history, the problems of object detection and estimation of 3D geometry have largely been pursued independently. However, there seem to be many good arguments for why these two sub-disciplines should join forces. Accurate recognition and segmentation of objects in a scene should constrain matching of features to hypothesized surfaces, aiding reconstruction. Similarly, geometric information should provide useful features and context for object detection and recognition. The use of detailed stereo depth data has already proven to be incredibly effective in the world of object detection. More general formulations integrate structure from motion, scene and object recognition as a joint problem (e.g., [4, 5, 20]. On the other hand, work such as [19, 14, 25, 31] have focused on the role of context in a single image, operating under the assumption that the camera and scene geometry are unknown and must largely be inferred based on analysis of monocular cues. This problem is quite difficult in general although some progress has been made [18, 28], particularly on indoor scenes of buildings where the geometry is highly regular [27, 15, 16].

In this paper, we argue that a huge number of photographs taken in outdoor urban areas are really photos of "known scenes" for which rich geometric scene data exists in the form of GIS maps and other geo-spatial data resources. Robust image matching techniques make it feasible to resection a novel image against a large database of scene data and produce precise estimates of camera pose on a world-wide scale [22]. Once a test photo has been precisely localized, much of this contextual information can be easily backprojected into the image coordinates to provide much stronger priors for interpreting image contents. For the monocular scene-understanding purist, this may sound like cheating, but from a practical perspective, such strong context is already widely available or actively being assembled and should prove hugely valuable for improving the accuracy of image understanding.

The role of GIS map data in automatically interpreting images of outdoor scenes appears to have received relatively little attention in computer vision. Detailed GIS map data is used extensively in analysis of aerial images [32, 33]. More recently, a few groups have looked at using GIS data and multi-view geometry for improving object recognition. The most closely related work to ours is probably that of [24] who introduced a geographic context re-scoring scheme for car detection based on street maps. GIS data was also used in [3] as a prior for static object detection and camera localization. Finally, [7] resectioned and matched image patches in monocular test images to a database of localized images in order to automatically perform background subtraction to improve object detection.

Contributions We present a new dataset test bed consisting of over six thousand images covering a portion of a university campus for which relative camera pose and 3D coordinates of matching points has been recovered using

This work was supported by NSF IIS-1253538 and a Balsells Fellowship to RD.

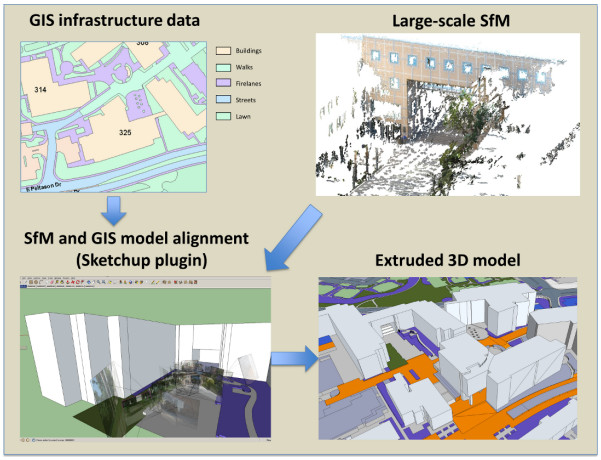

Figure 1: System overview: GIS and SfM data are merged using custom plugin inside the 3D modeling tool Sketchup to allow easy import of images and camera parameters. 2D GIS data can then be easily extruded into a full 3D model with relatively little user effort. This fused geocontext dataset provides a basis for efficiently transferring rich geometric and semantic information from the to a novel test image where it can be used to improve performance of general scene understanding in monocular images.

structure from motion. We describe a method for aligning this geometric and photometric data with 2D GIS map data derived from satellite images in order to quickly build rough 3D polygonal models of buildings, sidewalks, streets and other static structures with minimal user input. This combined geo-semantic context dataset serves as a geometric reference against which novel test images can be resectioned, providing a rich variety of geometric, geographic, and semantic features for further image analysis. We describe methods for improving the accuracy of resectioning that makes use of both triangulated 3D points and the polygonal scene model. Finally, we demonstrate the utility of the strong contextual constraints provided by this data for rescoring pedestrian detections, estimating scene depth from a single image, and improve semantic segmentation.

The rest of the paper is structured as follows: Section 2 describes the dataset construction and a 3D modeling pipeline that leverages GIS map data; Sections 3 and 4 define real world 3D geometric cues that will help reasoning about object detection and semantic segmentation respectively; Section 5 explains a validation protocol for the 3D model built and results on scene context re-scoring; Section 6 closes with a summary, discussion of limitations and future work.

2. Lifting GIS Maps

Here we describe building a test-bed dataset for leveraging strong GIS-derived context. We focus on aspects of this pipeline that emphasize bi-directional interaction where SfM can be used to enhance GIS maps (lifting to 3D) and where feedback from GIS context can improve camera resectioning and triangulation.

Image database acquisition We collected a database of 6402 images covering a large area (the engineering quad) of a university campus. Images were collected in a systematic manner using a pair of point and shoot cameras attached to a monopod. Locations were chosen so as to provide approximately uniform coverage of the area of interest. Images were generally collected during break periods when there were relatively few people present although some images still contain pedestrians and other non-rigid objects.

Running off-the-shelf incremental structure from motion (i.e., [2, 30]) on the entire dataset produces a 3D structure that is qualitatively satisfying but contains some metric inaccuracies. In particular, there is significant drift over the whole extent of the model which makes it impossible to globally align the model with GIS map data with high precision. Even the use of state of the art GPS and inertial sensors still result in inaccuracies in urban areas [21]. To guarantee the best possible alignment between camera poses estimated from bundle adjustment and map data, we developed a more scalable approach that performs bundle adjustment on local clusters of cameras and then aligns these clusters to a global model. While such local-to-global strategies are not uncommon in the SfM and SLAM literature, our contribution here is in optimizing trade-offs to assure the best possible alignment between any given image and the GIS model for later use in detection and segmentation tasks.

Local Structure from Motion To establish rough geometric layout of all cameras, we ran incremental bundle adjustment on the entire image set. We then generated overlapping spatially localized clusters of cameras with the property that each set shares many cameras with several other sets. To accomplish this we proceed in two phases.

First, we performed disjoint spatial clustering of the camera positions by running k-means on these initial camera estimates to produce disjoint clusters. In our experiments we used $K = 10$ which provide clusters with around 600 images each. Let C_i denote the set of cameras in the cluster i, m_i denote the mean of each cluster and σ_i the standard deviation around that mean. We define a pair of clusters to be neighbors based on a simple thresholding

$$\mathbf{N} = \left\{ (i,j) : \frac{d(m_i, m_j)}{\sigma_i + \sigma_j} < 3 \right\} \quad (1)$$

For each pair of clusters $(i,j) \in \mathbf{N}$ we then run bundle adjustment on the set of cameras $P_{ij} = C_i \cup C_j$.

From this large set of overlapping models we selected a small high quality set that covers all the cameras. We computed a quality measure for each camera set P_{ij} by

calculating the ratio of bad camera pairs in P_{ij} for which there were less than 50 matched points successfully triangulated during bundle adjustment. To select a subset of P_{ij} we find a minimum spanning tree of a graph whose nodes correspond to clusters C_i and whose edges correspond to pair P_{ij} with weights given by the quality of that camera set. Figure 2 depicts the clusters and the spanning tree over the clusters. The set of reconstructions corresponding to the edges of the spanning tree have the property that every camera cluster is covered, every reconstruction has significant overlap with at least two other reconstructions, and the selected reconstructions are high quality (the cameras are well connected). Since the tree is acyclic, we can easily produce an initial global alignment of all the clusters by traversing the tree in some order and using Procrustes alignment between the recovered coordinates of the overlapping cameras. For example P_{ij} and P_{jk} are aligned using their respective estimates of the cameras in C_j, using RANSAC to reject outlier cameras.

Global GIS-structure alignment We obtained a 2D GIS map of the campus maintained by the university's technical office. The map was originally constructed from aerial imagery and indicates polygonal regions corresponding to essential campus infrastructure tagged with semantic labels including building footprints, roadways, firelanes, lawns, etc. (see Figure 2).

We would like to align our reconstructions with this model. While GPS alone is not generally accurate enough, one may leverage existing sets of geo-referenced images. For example, [24] used aerial LiDAR and Google Street View images with known geo-coordinates in order to provide an absolute coordinate reference. To seed our model, we utilized hints from a user who identified corresponding 3D points in the model with 2D coordinates in the GIS map. We used vertical building corners which are easily visible and distinctive in both sets and are very close to parallel in 3D. We then applied Procrustes alignment between 3D points on the building corner and vertical lines perpendicular to the map at the corresponding building corners.

Once a particular camera set P_{ij} has been aligned to the geocoordinate system, we can easily traverse the cluster tree to automatically produce an initial alignment for all remaining camera sets. Specifying additional correspondences in other clusters and smoothing cluster alignment estimates over the tree (e.g., using dynamic programming) provides an efficient mechanism for fixing global drift seen in incremental bundle adjustment without having to re-optimize all cameras and tracks. To aid this process we built a simple tool to visualize the alignment projected back on to the GIS map to allow fast interactive refinement of the correspondences as necessary.

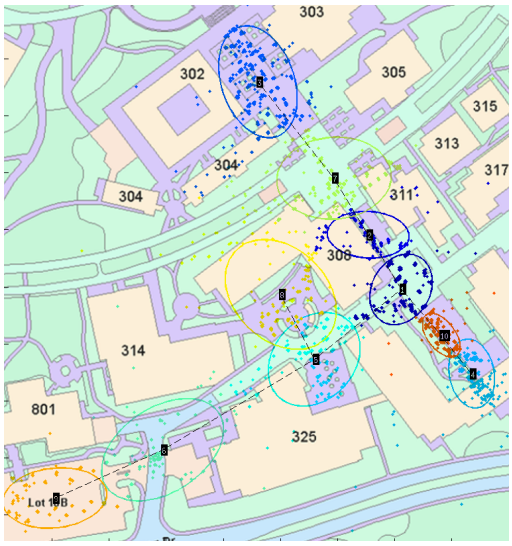

Figure 2: Spatial camera clusters and associated spanning tree overlaid on the GIS map data. Each edge in the tree defines a camera set P_{ij} for which the bundle adjustment yielded a high quality, densely connected reconstruction. Camera sets overlap and so can be automatically aligned and then further refined using user-supplied correspondences in order to correct for drift over the tree.

Image assisted 3D model construction We developed a custom plugin for the 3D modeling tool Sketchup [1] to allow efficient user-assisted "lifting" of 2D GIS map data into a full 3D model. We import the GIS 2D polygons with their corresponding semantic labels as well as selected images and their corresponding extrinsic and intrinsic camera parameters estimated previously. The user is then presented with an overlay in which the 3D model is backprojected into the selected camera viewpoint and overlayed with the image to assist in building the model. Figure 3 shows an example of this visualization.

In a few clicks, a full, 3D mesh can be easily extruded from the 2D map by lifting buildings up, carving stairs down, tilting non-fully horizontal surfaces, etc. Additional more detailed geometry can be easily created, as well as new semantic labels or corrections to the original data. With the assistance of these aligned camera views, constructing a fairly detailed model in Sketchup covering 10 buildings takes approximately 1-2 hours. This can be considered an offline task since the modelling effort is performed only once, as buildings are largely static over time.

Model-assisted Structure from Motion The previous subsections discussed how GIS information helped to achieve a better alignment of structure from motion models to real world coordinates and correct drift. On the other hand, we can also exploit the 3D model in order to detect poorly calibrated cameras and outlier 3D points in the SfM

Figure 3: We built a full 3D model from 2D GIS data by creating a toolbox that imports structure from motion reconstruction data into Sketchup. The toolbox allows to read camera information and properly display the image overlaid on the modelling interface. The user then models the 3D geometry by extruding, tilting, and carving the 2D data until it is aligned with the image.

model. This curation of the database by removing noisy elements improves test-time camera resectioning process by eliminating potential cases in which a test image 2D features might match wrongly calibrated 3D points.

Since our structure from motion models now live in real world coordinates, it is easy to define several heuristics for determining if a camera is well posed. To prune bad cameras from any model P_{ij}, we compute the camera height with respect to the ground and remove those cameras whose height is outside a reasonable range between 1 and 4 meters. Also, we measure the up vector of a camera and prune those cases where they are tilted more than 30 degrees. When pruning a camera, we also remove its point tracks from the SfM model.

Given a 3D polygonal model, many poorly recovered 3D points can be pruned by the following simple test. For each 3D point visible in a camera, we compare the depth of the point with the depth of the backprojected 3D model along the same ray. By applying a generous threshold of 25 meters depth discrepancy, the model prunes most of the wildly generated 3D points, while keeping other points of interest not accurately modeled in the GIS data (like trees, benches, trash bins, etc). A final filtering step is performed to detect (1) cameras with less than 6 track points and (2) 3D points with only 1 track. The resulting curated model then comprises all the SfM reconstruction data that is consistent with the GIS model and will serve as a fixed database against which to apply image resectioning in our tests.

Test-time camera resectioning To estimate camera pose for test images, we performed resectioned against the model using a RANSAC-based 3-point absolute pose (P3P) camera resectioning procedure, based on 2D-3D matching correspondences. Each test image was matched against every SfM cluster pair P_{ij} model included in the cluster tree. A best camera pose estimate was selected using the same match quality, height and orientation criteria used in building the original model. If a camera is geometrically reasonable in more than one model, we select the one with the highest number of matched inliers.

Resectioning of a test image produces immediate predictions of scene depth, surface orientation, and semantic labels at every image pixels. It is interesting to compare these to previous work that attempts to estimate such geometric or semantic information from a single image. By simply resectioning the image against our 3D model we are immediately able to make surprisingly accurate predictions without running a classifier on a single image patch! In the remainder of the paper we discuss how to upgrade these "blind" predictions by incorporating backprojected model information into standard detection and segmentation frameworks.

3. Strong Geometric Context for Detection

Estimating camera pose of a test image with respect to a 3D GIS model can aid in reasoning about the geometric validity of object hypotheses (e.g. pruning false positives placed in unlikely places, or boosting low-score results on suitable areas). In this section we describe methods for using these constraints to improve the performance of a pedestrian detector.

Geometric hypothesis sampling Let b be a candidate 2D bounding box in a test image I with an associated height in pixels h^{im}. If we assume that the object is resting on the ground-plane and the base of the object is visible, then we can estimate the depth z of the object with respect to the camera by intersecting the ray going from the camera center through the bounding box base with the 3D model. Importantly, unlike many previous works, the ground is not necessarily a plane (e.g., our model includes stairs and ramps). Given camera focal length f, we can estimate the height of the in world coordinates by the following expression:

$$h = \frac{z}{f} h^{im} \tag{2}$$

Unfortunately, the object's "feet" might not be visible at all times. We rely on our detector to let us know when that circumstance happens. We use a deformable part model that consists of a mixture over three different template filters: full-body, half upper body, and head. We set the image detection height h^{im} based on which mixture fires as 1, 2 or 3 times the bounding box height respectively.

When the feet are not visible, we cannot simply find the ground intersection point. Instead our strategy is to sample scene depth hypotheses z along a ray through the top center of the candidate bounding box. For each depth, we hypothesize possible world heights h and then compute the vertical elevation $e(z, h)$ of the object above (or below) the

ground surface of the 3D model. For computational efficiency we only consider a small discrete range of depth hypotheses (30cm intervals) centered around depths where an average height object is at zero elevation. We also prune those samples lying beyond walls to further reduce the pool of samples generated.

Hypothesis scoring Let $v = (s, h, e, g)$ be a hypothesized object where s is the detector score associated with the corresponding 2D bounding box, h is the hypothesized height, e is the elevation from the ground, and g is the semantic label on the ground point below the object location. To score a hypothesized detection, we interpret the detector score s as the log of a probability to which we add penalties associated with the height or elevation priors.

First, we assume a Gaussian prior on human height with mean $\mu_h = 1.7m$ similar to the work of [17]. Second, we penalize elevations e which suggest the detected pedestrian is flying above or tunnelling beneath the ground. Finally, we prune those candidate detections lying in an unlikely semantic region. Let W be the set of semantic labels in the GIS data defining a human walkable space (streets, firelanes, walks, etc.). The final penalized detection score is given by:

$$S(v) = \begin{cases} s - \alpha||\mu_h - h||^2 - \beta||e||^2 & g \in W \\ -\infty & g \notin W_s \end{cases} \quad (3)$$

where parameters α, β capture the degree of variation around the mean height as well as inaccuracies in our estimation of the height from image features. We performed exhaustive search over parameters α and β to find a global performance optimum at $\alpha = 0.4$ and $\beta = 0.3$. However, performance was robust to these parameters and setting them both to 1 provided an accuracy difference less than 0.5% of the overfitting settings.

4. Strong Geometric Context for Segmentation

The geometric and semantic information contained in the GIS data and lifted into the 3D GIS model can aid in reasoning about the geometric validity of class hypotheses (e.g., horizontal surfaces visible from the ground are not typically buildings). We describe methods for using such constraints to improve semantic segmentation performance. We follow a standard CRF-based labeling approach, built on the high quality implementation of Gould et al. [13] and an augmented set of image features from [29]. We explore simple ways to enhance this set of features using GIS data and study its influence on semantic labeling accuracy.

GIS label prior distributions The GIS-derived model provides an immediate estimation of pixel labels based on semantic labels in the original GIS map. If a camera pose is known, we can backproject the model into the image plane and transfer the polygon label in the GIS model to the projected pixel. However, camera pose estimation is not perfect and might contain minimal deviations from its ground truth pose. In order to account for slight camera pose inaccuracies, we define a feature descriptor to softly handle these cases. Given an image I, a pixel $x \in I$, and a backprojected GIS semantic label $g(x)$, we define the feature $h^r_k(x)$

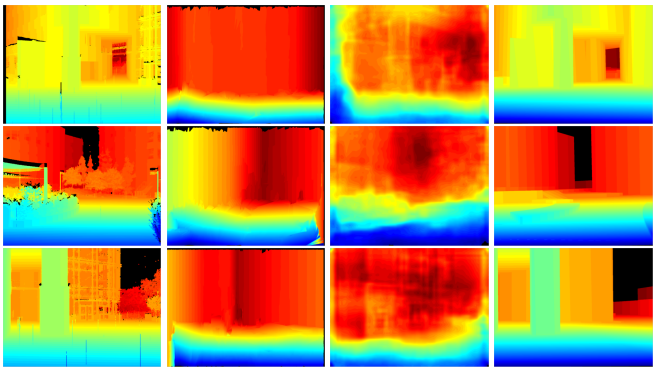

Figure 4: Qualitative depth comparison. Our GIS backprojected depth map is shown in the last column. While it lacks many details such as foliage and pedestrians which are not included in our coarse GIS-based 3D model, simply backprojecting depth provides substantially more accurate estimation than existing monocular approaches.

$$h^s_{r,k}(x) = \frac{1}{N} \sum_{y:\|y-x\|<r} [g(y) = k] \quad (4)$$

as the normalized count of class k pixels in a circular disc of radius r around x, where N is the number of pixels in the disc. In our experiments, r is defined so that the angular error of the camera pose is 0, 1, 3, and 5 degrees.

GIS surface normal distributions In a similar manner, a surface normal can be quickly estimated for any pixel by backprojecting the 3D model into the camera plane. Surface normals can be discriminative of certain classes like pavement, roads, buildings, etc. Following the same structure as in equation 5, we define the feature $h^r_n(x)$ as

$$h^n_{r,k}(x) = \frac{1}{N} \sum_{y:\|y-x\|<r} [n(y) \in N_k] \quad (5)$$

where N_k is one of 3 possible surface orientation bins: horizontal (ground), horizontal (ceiling), vertical (wall).

GIS Depth features Depth can also be efficiently estimated when from a 3D model when a camera pose is known. Following other methods like [35, 26], we extract histograms of oriented gradients [6] to encode depth variations. However, we did not find these absolute depth features to provide a substantial gain in our experiments.

Threshold	Make3D	DNN	GIS
$\delta < 1.25$	11.15%	26.03%	68.13%
$\delta < 1.25^2$	30.68%	46.24%	83.59%
$\delta < 1.25^3$	48.96%	67.15%	89.33%

Figure 5: Quantitative comparison of depth estimation methods: Make3D [28], Deep Neural Network [8], and GIS backprojection. We use the maximum allowed relative error $\max(\frac{d_{gt}}{d_{est}}, \frac{d_{est}}{d_{gt}})$, where d_{gt} is the ground truth depth and d_{est} is the estimated depth at some point. DNN model was re-scaled to overfit our ground truth data since the model provided is adapted only for indoor scenes.

DPM as a context feature Inspired by other works that try to create segmentation-aware detectors [10], we also incorporate the outputs of category-specific object detectors in our segmentation model. To do so, we collect the scores of a DPM detector for an object category c and generate a DPM feature map h_c by assigning to every pixel the maximum score of any of the candidate detection boxes intersecting the given pixel. Let Ω_c be the set of candidate detections and B_i the bounding box for the ith detection, then

$$h_c^o(x) = \max_{i \in \Omega_c}(s_i \cdot [x \in B_i]) \quad (6)$$

5. Experimental results

In our experiments we started from a test set comprising 570 pictures taken in the area covered by the 3D model. These images were taken over different days and collected several months after the initial model dataset was acquired. Of these images, 330 images were successfully resectioned (verified by visual inspection). This success rate compares favorably with typical success rates for incremental bundle adjustment (e.g., [24] register 37% of their input images). To evaluate performance of detection, we annotated these 330 test images with ground-truth bounding boxes for 1450 pedestrians visible in the images using the standard VOC procedure including tags for truncated and partially occluded cases. We also manually segmented 302 of these images using the segmentation tool provided by [23] and labeled each segment with one of 9 semantic categories.

5.1. Monocular Depth Estimation

To verify that the coarse-scale 3D GIS model provides useful geometric information, despite the lack of many detailed scene elements such as trees, we evaluated resectioning and backprojection as an approach for monocular depth estimation. While our approach is not truly monocular in the sense that it relies on a database of images to resection the camera, the test time data provided is in fact monocular. Furthermore, this type of accurate camera localization is increasingly practical for many outdoor urban environments.

To establish a gold-standard estimate of scene depth, we scanned 14 locations of the area covered by our dataset using a Trimble GX3D terrestrial laser scanner. We took the scans at in a range of resolution between 5 and 12 cm in order to keep the total scanning time fairly low, resulting in roughly a half a million 3D points per scan. We mounted a camera on top of the laser scanner and used the camera focal length to project the laser-based 3D point cloud onto the camera image plane, interpolating depth values to obtain a per-pixel depth estimate. We then resectioned the test image and synthesized the depth map predicted by our 3D GIS model.

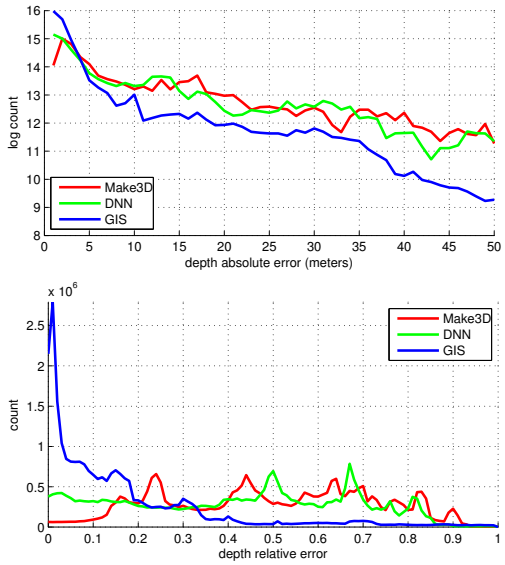

Figure 6: Accuracy of predicted depth estimates compared to gold-standard provided by a laser-scanner for 14 images. Top histogram shows the distribution of absolute depth errors between Make3D, DNN and depth computed from resectioning and backprojecting GIS derived 3D model. The bottom plot shows the distribution of relative error $|d_{est} - d_{gt}|/d_{gt}$.

Figure 5 shows quantitative results of our GIS backprojection depth estimation against other single-image depth approaches for the 14 scan dataset. We used the provided pre-trained models included in [28, 8] as baselines for comparison. Since the DNN model [8] was trained on indoor scenes, we estimated a scaling and offset for the output that minimized the relative error over the 14 images. While the reader might consider comparison with pretrained models unfair, we point out that collecting large amounts of high-quality outdoor depth data requires specialized hardware (i.e. laser scanner) while our approach makes surprisingly accurate predictions using only camera pose and a map, with no "training data" required.

5.2. Object Detection

We evaluated our geometric and semantic object rescoring scheme applied to the widely used deformable part model detector [9] implemented in [11].

Since we hypothesize full-height bounding boxes for our contextual rescoring, we apply a variant of non-max suppression we term "Full Body NMS" on these full height bounding boxes so that the suppression is applied fairly among boxes of the same aspect ratio. The output detections are benchmarked using the original (non-full height) boxes since our ground truth annotations do include truncated and occluded objects. A candidate output bounding box is considered a true positive if the standard VOC intersection over union overlap is higher than 0.5.

The standard DPM on the test set provide a baseline performance with an average precision (AP) of **0.509** using the standard NMS approach. Using the Full Body NMS improved results significantly, increasing AP to **0.564**. Even without a geometry context, we are able to prune false positives that would pass the overlap criteria when their height is simply the bounding box size rather than their full-body height. Rescoring hypotheses based on height, elevation and semantic label constraints provides a further small boost in performance, raising AP to **0.576**. Figure 7 shows the behaviour of DPM with Full Body NMS when elevation and height priors are applied. If we only consider fully visible ground-truth pedestrians and detections, the relative gain in AP is slightly higher, from **0.66** to **0.679**.

Figure 8 shows examples of DPM and DPM+GIS rescored detections at a fixed level of recall. As can be seen in the figure, geometric context can quite effectively suppress high scoring detections in improbable locations (e.g., the trees) but the same rescoring effectively boosts other geometrically plausible hypotheses near the walkable ground surfaces. As a result, we observe only small overall gains in average precision. We note that this parallels similar findings by [24] where 3D context between cars and streets allowed for improved geometric reasoning about car orientation but only small gains in detection performance (in fact their final model had lower average precision relative to a baseline DPM model).

5.3. Semantic Segmentation

We trained baseline semantic segmentation models on two different sets of training images using the code provided by [12]. We trained one model using 100 images taken from our labeled set of engineering quad images (ENGQ) and reserved 202 remaining images as held-out test data for evaluating labeling accuracy. We also trained a generic model using images collected from the online SUN dataset [34] by querying for multiple categories in the SUN label set that are semantically equivalent to the 9 categories labeled in the ENGQ images.

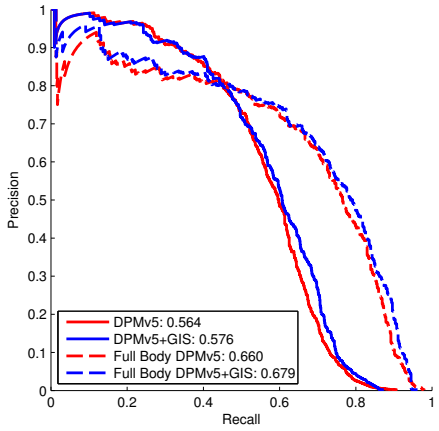

Figure 7: Precision recall curves for detector baseline using Full Body-NMS (red), and GIS-rescored (blue). Dashed lines evaluate the scenario in which only full body annotations and full body detector outputs are available.

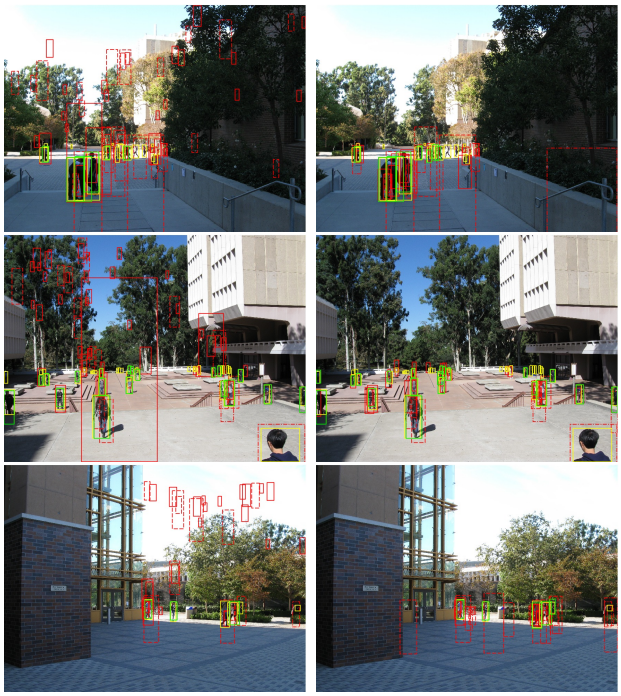

DPM DPM+GIS

Figure 8: Detection results at 0.75 recall. Yellow boxes correspond to ground truth annotations, while red and green indicate false and true positives respectively. Inferred full-body bounding boxes are drawn in dashed lines. Left column shows original DPM output; right shows GIS-rescored detections.

Detectors Improve Segmentation We collected DPM score features as described previously for two objects of interest: pedestrian and bicycle. Figure 9 shows the influence

Model	building	plants	pavement	sky	bench	wall	ped.	ped. sit	bicycle
GIS-Only	0.711	0.346	0.876	0.664	0.000	0.104	0.000	0.000	0.021
SUN	0.767	0.589	0.878	0.872	0.000	0.000	0.012	0.000	0.000
+GIS-Only	**0.787**	0.581	**0.890**	**0.880**	0.000	0.000	0.000	0.000	0.000
+DPM	0.778	**0.607**	0.876	0.871	0.000	0.000	**0.146**	0.000	**0.002**
ENGQ	0.914	0.890	0.934	0.923	0.128	0.323	0.399	0.030	0.157
+GIS	0.934	0.897	0.947	0.961	**0.170**	0.484	0.405	0.036	0.181
+DPM	0.920	0.898	0.939	0.937	0.130	0.368	**0.596**	0.024	**0.683**
+DPM+GIS	**0.937**	**0.904**	**0.949**	**0.963**	0.161	**0.494**	0.572	**0.048**	0.654

Figure 9: Quantitative segmentation results for models trained with generic (SUN) and scene specific (ENGQ) data. Accuracy is measured using average overlap score between labeled regions and ground-truth. Adding geometric (+GIS) and detection (+DPM) features outperformed the baseline models. Combining both methods gave the best overall results in the scene specific model, although some classes did not achieve their best accuracy individually.

of adding these features into the model. Besides a slight general improvement across all categories, we were able to boost pedestrian accuracy from 0.399 to 0.596 and bicycle from 0.157 to 0.683 in the scene-specific model.

GIS-aware Segmentation To evaluate the influence of GIS features, we first trained a "blind" model on the ENGQ without any image features (GIS-Only), instead using only the features extracted from the backprojected GIS model. This yielded relatively good accuracy in those categories present in the GIS data (building, plants, pavement, and sky) but poor results for many others, indicating that certain labels cannot be learned just from context features. This is quite natural since, e.g., our GIS model doesn't include detailed scene elements such as benches and provides almost no information about which particular pixels might be a bike or pedestrian on any given day.

On the other hand, combining these GIS features with standard image features gave a significant benefit, outperforming the ENGQ-trained scene-specific baseline in all categories (+GIS rows in figure 9). For example, performance on wall jumps from 0.368 to 0.494 when GIS features are added. This can be explained by the fact that buildings and retaining walls often have similar local appearance but the global context provided by the 3D model can resolve this ambiguity.

Evaluating the benefit of GIS model context is more difficult for the generic SUN-based appearance model since the GIS features were not available for SUN images. Instead, we evaluated a hybrid SUN+GIS-Only model which combined pretrained unary terms from each model. We adapted a single parameter $\alpha = 0.7$ which controlled the relative weighting of the SUN appearance unary and GIS-Only geometry unary terms to optimize performance. We found similar gains as in the scene-specific trained with ENGQ data.

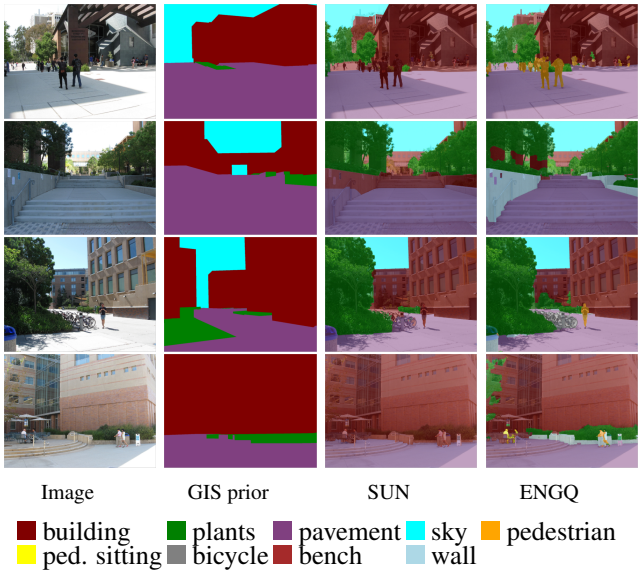

Figure 10: Qualitative segmentation results with overlaid images. Our context-exploited classifier (ENGQ) includes both GIS and DPM features, outperforming the generic baseline.

6. Discussion

The rapid growth of digital mapping data in the form of GIS databases offers a rich source of contextual information that should be exploited in practical computer vision systems. We have described a basic pipeline that allows for integration of such data to guide both traditional geometric reconstruction as well as semantic segmentation and recognition. With a small amount of user supervision, we can quickly lift 2D GIS maps into 3D models which immediately provide strong scene geometry estimates (typically less than 5-10% relative error), greatly outperforming existing approaches monocular depth estimation and providing a cheap alternative to laser range scanners. This backdrop provides a strong basis for geometric and semantic context which can be exploited to improve detection and segmentation and even make surprisingly accurate predictions without even analyzing image appearance.

References

[1] Sketchup 3d modelling tool. http://www.sketchup.com/. Accessed: 2014-11-13. 3

[2] S. Agarwal, N. Snavely, I. Simon, S. Seitz, and R. Szeliski. Building Rome in a day. *ICCV*, 2009. 2

[3] S. Ardeshir, A. R. Zamir, A. Torroella, and M. Shah. Gis-assisted object detection and geospatial localization. In *ECCV*. 2014. 1

[4] S. Y. Bao, M. Bagra, Y.-W. Chao, and S. Savarese. Semantic structure from motion with points, regions, and objects. In *CVPR*, 2012. 1

[5] N. Cornelis, B. Leibe, K. Cornelis, and L. Van Gool. 3d urban scene modeling integrating recognition and reconstruction. *IJCV*, 2008. 1

[6] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *CVPR*, 2005. 5

[7] R. Diaz, S. Hallman, and C. C. Fowlkes. Detecting dynamic objects with multi-view background subtraction. In *ICCV*, 2013. 1

[8] D. Eigen, C. Puhrsch, and R. Fergus. Depth map prediction from a single image using a multi-scale deep network. In *NIPS*, 2014. 5, 6

[9] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part-based models. *TPAMI*, 2010. 7

[10] S. Fidler, R. Mottaghi, A. Yuille, and R. Urtasun. Bottom-up segmentation for top-down detection. In *CVPR*, 2013. 6

[11] R. B. Girshick, P. F. Felzenszwalb, and D. McAllester. Discriminatively trained deformable part models, release 5. http://people.cs.uchicago.edu/ rbg/latent-release5/. 7

[12] S. Gould. Darwin: A framework for machine learning and computer vision research and development, v1.9. *The Journal of Machine Learning Research*, 2012. 7

[13] S. Gould, R. Fulton, and D. Koller. Decomposing a scene into geometric and semantically consistent regions. In *ICCV*, 2009. 5

[14] A. Gupta, S. Satkin, A. A. Efros, and M. Hebert. From 3d scene geometry to human workspace. In *CVPR*, 2011. 1

[15] V. Hedau, D. Hoiem, and D. Forsyth. Recovering the spatial layout of cluttered rooms. In *ICCV*, 2009. 1

[16] V. Hedau, D. Hoiem, and D. Forsyth. Thinking inside the box: Using appearance models and context based on room geometry. In *ECCV*. 2010. 1

[17] D. Hoiem, A. Efros, and M. Hebert. Putting Objects in Perspective. *CVPR*, 2006. 5

[18] D. Hoiem, A. A. Efros, and M. Hebert. Automatic photo pop-up. In *ACM Transactions on Graphics (TOG)*, 2005. 1

[19] D. Hoiem, A. A. Efros, and M. Hebert. Geometric context from a single image. *ICCV*, 2005. 1

[20] D. Hoiem, A. A. Efros, and M. Hebert. Closing the loop in scene interpretation. In *CVPR*, 2008. 1

[21] B. Klingner, D. Martin, and J. Roseborough. Street view motion-from-structure-from-motion. In *ICCV*, 2013. 2

[22] Y. Li, N. Snavely, D. Huttenlocher, and P. Fua. Worldwide pose estimation using 3d point clouds. In *ECCV*. 2012. 1

[23] M. Maire, S. X. Yu, and P. Perona. Hierarchical scene annotation. In *BMVC*, 2013. 6

[24] K. Matzen and N. Snavely. Nyc3dcars: A dataset of 3d vehicles in geographic context. In *CVPR*, 2013. 1, 3, 6, 7

[25] K. Murphy, A. Torralba, and W. Freeman. Using the forest to see the trees: a graphical model relating features, objects and scenes. *NIPS*, 2003. 1

[26] G. Rogez, J. S. Supancic III, M. Khademi, J. M. M. Montiel, and D. Ramanan. 3d hand pose detection in egocentric rgb-d images. *arXiv preprint arXiv:1412.0065*, 2014. 5

[27] S. Satkin and M. Hebert. 3dnn: Viewpoint invariant 3d geometry matching for scene understanding. In *ICCV*, 2013. 1

[28] A. Saxena, M. Sun, and A. Y. Ng. Learning 3-d scene structure from a single still image. In *ICCV*, 2007. 1, 5, 6

[29] J. Shotton, J. Winn, C. Rother, and A. Criminisi. Textonboost: Joint appearance, shape and context modeling for multi-class object recognition and segmentation. In *ECCV*. 2006. 5

[30] N. Snavely, S. Seitz, and R. Szeliski. Modeling the World from Internet Photo Collections. *IJCV*, 2007. 2

[31] M. Sun, S. Y. Bao, and S. Savarese. Object detection using geometrical context feedback. *IJCV*, 2012. 1

[32] H. Uchiyama, H. Saito, M. Servieres, and G. Moreau. Ar gis on a physical map based on map image retrieval using llah tracking. In *MVA*, 2009. 1

[33] L. Wang and U. Neumann. A robust approach for automatic registration of aerial images with untextured aerial lidar data. In *CVPR*, 2009. 1

[34] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *CVPR*, 2010. 7

[35] R. Yang and R. Yang. Action segmentation and recognition based on depth hog and probability distribution difference. In *Intelligent Computing Theory*. 2014. 5