Distance Bitmask: An Alternative Approach to Nearest Neighbor Distance for Image Geolocation

Eugene Polishchuk
Nathan Campbell
Jason Surratt
Ryan Smith
ryan.smith@digitalglobe.com
DigitalGlobe Intelligence Solutions

Abstract

Geolocation of terrestrial photos and videos often involves distance estimation to easily identifable objects in the image. Those distance estimates are then later used for the purpose of area reduction and ultimately identification of the actual location where the image was taken. In most cases this approach makes the assumption that the object seen in the image is the closest to the location of interest. This assumption simplifies the problem by allowing the system to precompute Nearest Neighbor Distances to well known objects and later use them to perform rapid area reduction based on the provided distance estimates. Unfortunately this simplification can be problematic in cases where there is a high likelihood that the object present in the image is actually not the closest object to the location where the image was taken. For example if you see a road in the distance in the image, there is no guarantee that there is no road right behind the location where the image was taken. Here we introduce an alternative approach to Nearest Neighbor Distance, that still allows us to perform rapid area reduction, but does not require the assumption that the object seen in the query image is the closest object to the location of interest.

I. Introduction

Successful geolocation of videos and images often relies on the ability to narrow the search space and focus computationally and analytically intensive resources on high probability locations. Our approach to search area reduction, also referred to as suitability analysis, has in the past relied on building a world model from geospatial data and using the model to encode features and objects as precomputed Nearest Neighbor Distance (NND) rasters. These features include easily identifiable manmade points of interest, terrain and land cover. They are presented to the user as a set of high quality, curated and intelligible map layers.

Nearest neighbor computation assigns each cell in a raster the Euclidean distance from the cell center to the nearest object. For our purpose this makes an assumption that the object seen in an image or video is the closest object of that type to the location of interest. The result of this approach is a precomputed NND raster for each map layer. This is trivial to compute and store, easy to understand and allows area reduction algorithms to rapidly filter search regions based on proximity to user annotated map layers in an image or video.

Acknowledgement: Supported by the Intelligence Advanced Research Projects Activity (IARPA) via Air Force Research Laboratory, contract FA8650-12-C-7214. The U.S Government is authorized to reproduce and distribute reprints for Governmental purposes not withstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, AFRL, or the U.S. Government.

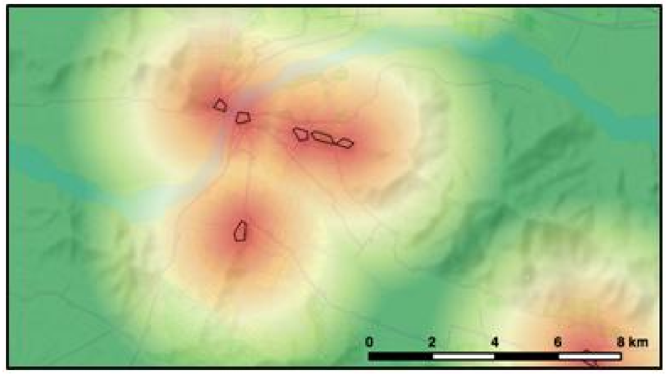

Figure 1: *Knolls - Nearest Neighbor Distance*
This visualization of a NND raster illustrates how distances are encoded. The dark outlines represent the object of interest, knolls. Warmer colors correspond to lower pixel values, or locations nearer to the object of interest.

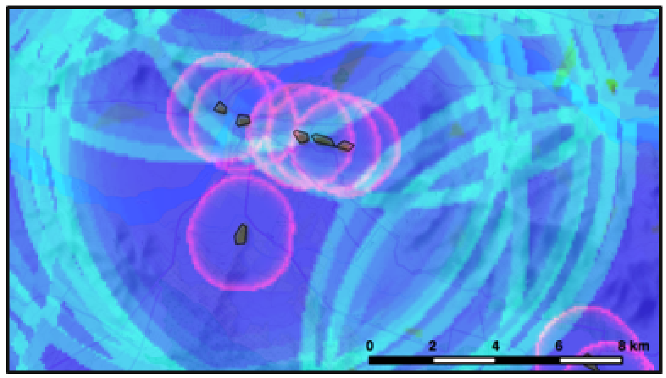

Figure 2: *Knolls - Distance Bitmask*
This visualization makes apparent the additional information captured by using a Distance Bitmask approach on the same knolls seen in Figure 1.

II. Method

Area reduction from simple proximity based suitability analysis can be improved upon if the user can estimate basic distance information to objects seen in the image or video. Utilizing distance estimates with NND rasters has inherent limitations in that only distance to the closest object is captured. For example, if you see a road in the distance in an image there is no guarantee that there is not a road right behind the location where the image was taken, beyond the field of view of the camera. This limitation can result in prioritizing incorrect regions and potentially overlooking regions containing the actual camera location. This scenario is a common occurrence in outdoor images and videos where significant terrain or large manmade objects can be visible at great distances. The NND approach potentially renders many types of features unusable in these scenarios.

In this paper we introduce an alternative approach to NND rasters called Distance Bitmask (DBM). This novel approach allows us to overcome the assumptions of NND models while maintaining a rapid and interactive suitability analysis. Figure 1 and Figure 2 show a visualization of a raster encoding NND and DBM respectively for a particular feature, which in this example is knolls.

III. Technical Overview

In an ideal world we would know the distance to every known object for every potential location in our area of interest. In most cases this is prohibitively expensive to compute and store. Instead we take all distances, up to 32 kilometers, to a particular object type and encode them in a three-band, sixteen-bit raster. Values beyond 32 kilometers are discarded as it is rarely possible to see an object at that range.

Each bit in the three band DBM raster corresponds to a distance range. A value of one indicates an object present for that distance range. For example, if band one, bit five is set there is at least one object 400-499 meters away. With this approach, we lose the number of objects a certain distance away from a particular location, but we retain the more valuable knowledge that there is an object a certain distance away from the particular location. This allows us to handle the scenario where an object seen in the image is not the closest instance of that object to the location in question. This approach is also efficient in terms of data storage space required and is also easy to use and understand.

Distance values are binned in the following fashion. Distances 0-1500 meters are stored as the first band at 100 meter resolution. The first bit in band one is reserved for containment, or where a point falls inside a polygon. The second band encodes distances 1000-9000 meters in 500 meter resolution and band three encodes 8500-32500 meters at 1500 meter resolution. Distances greater than 32500 meters are discarded. Figure 3 illustrates the complete bit definition for the Distance Bitmask (DBM) file format. DBM files can take significantly longer to create compared to NND rasters, but utilizing them to determine object's presence within a certain distance range for suitability analysis takes only a fraction of a second even for large areas.

Figure 3: *DBM Bit Definitions*

Bit	Band 1	Band 2	Band 3
0	Contained	1000-1499m	8500-9999m
1	0-99m	1500-1999m	10000-11499m
2	100-199m	2000-2499m	11500-12999m
3	200-299m	2500-2999m	13000-14499m
4	300-399m	3000-3499m	14500-15999m
5	400-499m	3500-3999m	16000-17499m
6	500-599m	4000-4499m	17500-18999m
7	600-699m	4500-4999m	19000-20499m
8	700-799m	5000-5499m	20500-21999m
9	800-899m	5500-5999m	22000-23499m
10	900-999m	6000-6499m	23500-24999m
11	1000-1099m	6500-6999m	25000-26499m
12	1100-1199m	7000-7499m	26500-27999m
13	1200-1299m	7500-7999m	28000-29499m
14	1300-1399m	8000-8499m	29500-30999m
15	1400-1499m	8500-8999m	31000-32499m

IV. Performance Impact

We applied our suitability analysis algorithms in conjunction with DBM rasters to a study area of 250,000 square kilometers in eastern Australia. The primary metric of the study was to reduce the potential search space returned by the algorithms while capturing the camera location. Out of 200 test images with known ground truth locations, 58 utilized at least one DBM layer.

For these 58 queries overall, we saw an increase of 15.6% in area reduction performance using DBM layers vs NND only layers.

At a threshold of 5% of the search area, we observed a 40% increase in the number of captured camera locations. Figure 4 illustrates that a performance increase can be seen across practically whole range of area thresholds.

Figure 4: *Performance Curve*

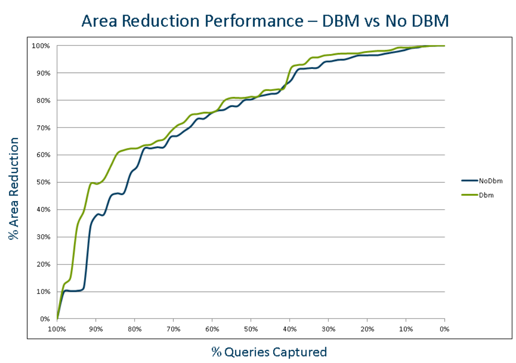

V. Future Work

The current implementation of distance bitmasks includes features whether or not they are actually visible. Modern graphics processing units (GPU) technologies enable rapid viewshed computation. We can very quickly determine visible areas from a particular location. This can be used to perform visibility filtering of objects so that only the features that are visible are encoded into the distance bitmask. Given that DBM overcomes the nearest neighbor assumption and encodes distances to all features within 32 kilometers, by filtering to only the visible features we expect to see additional improvement in area reduction performance. For features that are often prevalent in a region such as forests, a DBM range query is likely to return most of the region as the majority of the area is likely to be within range of a forest. By applying a visibility filter and including only those forests that are visible yet at a distance we can greatly improve the area reduction value of the forest map layer.

VI. Conclusion

We have demonstrated the utility of distance bitmasks as an alternative to nearest neighbor distances in the scope of area reduction for image geolocation. There is a clear path forward to further enhancing DBMs with the inclusion of visibility calculations to add significant area reduction. Additionally, we expect future improvements to DBM code to speed the initial creation time of these files. The use of the DBM file format allows us to encode more information from our world model. This enhances suitability analysis performance by allowing us to more effectively leverage distance to objects.