GeoSIM: A Geospatial Data Collection System for Participatory Urban Texture Documentation

Farnoush Banaei-Kashani Houtan Shirani-Mehr Bei Pan Nicholas Bopp
Luciano Nocera Cyrus Shahabi

Department of Computer Science
University of Southern California
Los Angeles, CA 90089-0781
[banaeika, hshirani, beipan, boppnick, nocera, shahabi]@usc.edu

Abstract

Participatory texture documentation (PTD) is a geospatial data collection process in which a group of users (dedicated individuals and/or general public) with camera-equipped mobile phones participate in collaborative collection of urban texture information. PTD enables inexpensive, scalable and high resolution data collection for urban texture mapping. In this paper, we introduce GeoSIM (Geospatial Social Image Mapping), a system we have designed and developed to enable efficient PTD. GeoSIM deploys a two-step planning process for efficient PTD. At the first step, termed "viewpoint selection", a minimum number of points in the urban environment are selected from which the texture of the entire urban environment (the part accessible to cameras) can be collected/captured. At the second step, called "viewpoint assignment", the selected viewpoints are assigned to the participating users such that given a limited number of users with various user constraints (e.g., specific participation time) users can collectively capture maximum amount of texture information within a limited time interval. Viewpoint selection and viewpoint assignment are both NP-hard problems. We present the design and implementation of GeoSIM based on our proposed heuristics for efficient viewpoint selection and viewpoint assignment that enable on-the-fly planning for PTD.

1 Introduction

The advent of earth visualization tools (e.g., Google Earth^{TM}, Microsoft Virtual Earth^{TM}) has inspired and enabled numerous applications. Some of these tools already include *texture* in their representation of the urban environment. The urban texture consists of the set of images/photos collected from the real environment, to be mapped on the façade of the 3D model of the environment (e.g., building and vegetation models) for photo-realistic 3D representation. Currently, urban texture is collected via aerial and/or ground photography (e.g., Google Street View). As a result, texture collection/documentation is 1) expensive, 2) unscalable (in terms of the required resources), and 3) with low temporal and/or spatial resolution (i.e., texture cannot be collected frequently and widely enough).

Bulletin of the IEEE Computer Society Technical Committee on Data Engineering

These limitations can be addressed by leveraging the popularity of camera-equipped mobile devices (such as cell phones and PDAs) for inexpensive and scalable urban texture documentation with high spatiotemporal resolution. With *participatory texture documentation*, termed *PTD* hereafter, a group of participants (dedicated individuals and/or general public) with camera-equipped mobile phones participate in collaborative/social collection of the urban texture information. By enabling low-cost, scalable, accurate, and real-time texture documentation, PTD empowers various applications such as eyewitness news broadcast, urban behavior analysis, real-estate monitoring, emergency-response, and disaster management (e.g., for damage assessment in case of earthquake, hurricane, and wildfire).

With this paper, we introduce *GeoSIM* (*Geospatial Social Image Mapping*), a system we have designed and developed to enable efficient PTD [1]. GeoSIM plans for efficient PTD by employing a two-step process. At the first step, called *viewpoint selection*, a set of points in the urban environment is selected from which the texture information of the entire environment (the part accessible to cameras) can be collected. We call such points as *viewpoints*. Due to the participatory nature of PTD, available resources (e.g., users' participation time) are usually limited and, therefore, it is critical to minimize the number of selected viewpoints. At the second step, termed *viewpoint assignment*, the selected viewpoints are assigned to the users for texture collection. The viewpoints must be assigned such that the texture collected during the documentation campaign (i.e., the specific time interval allocated for texture documentation) is maximized while all users' constraints are satisfied. In [7] and [6], we prove that the problems of viewpoint selection and viewpoint assignment are both NP-hard problems by reduction from the minimum set-cover problem and the team orienteering problem, respectively. Therefore, optimal implementations of viewpoint selection and viewpoint assignment are unscalable and fail to satisfy the real-time requirements of planning for efficient PTD given its participatory nature. Accordingly, we propose efficient heuristics for each of the two problems that are scalable and allow for on-the-fly planning for PTD. GeoSIM is designed and implemented based on our proposed heuristics.

The rest of this paper is organized as follows. In Section 2, we formally define our problem by describing the two-step planning process for efficient PTD. Subsequently, in Section 3 we present our corresponding two-step solution for on-the-fly PTD planning. We discuss the GeoSIM system design and implementation in Section 4. In Section 5, we discuss the related work, and finally, we conclude in Section 6.

2 Problem Definition

In this section, after explaining preliminary concepts, we formally define the viewpoint selection and viewpoint assignment problems in Sections 2.2 and 2.3, respectively.

2.1 Preliminaries

Below, first we explain how we model the 3D environment which is subject to texture documentation. Next, we define our assumed user participation model for participatory texture documentation.

2.1.1 Environment Model

Consider an urban environment which consists of various 3D elements such as buildings, trees and terrain (see Figure 1(a), for example). Suppose the environment is modeled in object-level (i.e., a 3D model exists in which the entire environment is represented by a set of objects). Here, without loss of generality, we assume the environment is modeled by the 3D TIN (Triangulated Irregular Network) model. The corresponding TIN model of the environment shown in Figure 1(a) is depicted in Figure 1(b) (shown in 2D). The *texture* of the environment is defined as the set of images mapped on the triangles of the 3D TIN model.

a. Actual Environment

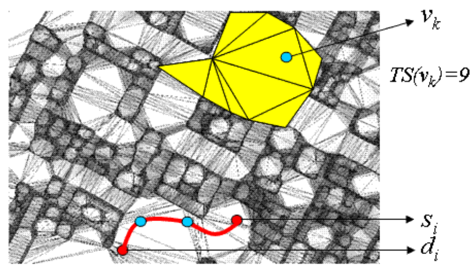

b. TIN Model

Figure 1: 3D Environment Representation for Texture Documentation

2.1.2 Participation Model

A texture documentation campaign is defined as the process of collecting and mapping the environment texture onto the corresponding TIN model of the environment during a predefined time interval T_C, termed the *campaign time* (e.g., 10:00am to 2:00pm on a particular day). We assume the urban texture remains unchanged during T_C. Suppose V is the set of points in the environment such that from each point $v \in V$ one can collect texture information by imaging the surrounding area. We call each such point v a *viewpoint*. Accordingly, we define the *texture score* $TS(v)$ of a viewpoint v as the total number of TIN triangles visible from v. For example, in Figure 1(b) the texture score of the viewpoint v_k is $TS(v_k) = 9$. Similarly, the texture score $TS(W)$ of a set of points $W \subseteq V$ is defined as the total number of TIN triangles visible from any viewpoint in W.

With participatory texture documentation, the texture collection process is implemented by a set of users U. We assume each user $u \in U$ has a set of participation constraints denoted by $c = (s, d, A)$, where s is user's starting point in the environment, d is user's desired ending point (where the user intends to leave the documentation campaign), and A is user's maximum available time for participation. Accordingly, a *participation plan* (or *participation path*) for a user u is define as a path $P_u = (s, v_1, v_2, \ldots, v_n, d)$ that starts from the starting point s and ends at the ending point d while traversing a number of viewpoints v_1 to v_n, where the user is expected to make stops for texture collection. Figure 1(b) shows a sample participation path for a user u_i (not shown in the figure) with the set of constraints $c_i = (s_i, d_i, A_i)$; in this case, the sample participation path traverses two viewpoints. Furthermore, a participation path P_u for user u is said to *satisfy* the user constrains c if and only if the total time to traverse the participation path (i.e., the actual user participation time) is less than the user available time A:

$$t_p + nt_{v_i} \leq A \tag{1}$$

where t_p is the total time to traverse the subpaths between successive viewpoints (assuming shortest path), and t_{v_i} is the time it takes to collect images at each viewpoint v_i along the path P_u. Finally, the texture score $TS(P_u)$ of the path P_u is defined to be equal to the texture score $TS(V_{P_u})$, where V_{P_u} is the set of viewpoints v_1 to v_n covered by P_u. Similarly, the texture score $TS(P_U)$ of a set of paths P_U is defined to be equal to the texture score $TS(V_{P_U})$, where V_{P_U} is the set of viewpoints covered by at least one path in P_U.

2.2 Viewpoint Selection Problem

Suppose T is the set of TIN triangles that comprise the 3D model of the target environment. Consider $T' \subseteq T$ as the subset of TIN triangles that are visible from at least one viewpoint in V (note that given a finite set of viewpoints V, there might be a non-empty set of triangles $T \backslash T'$ that cannot be texture mapped, regardless). Accordingly, we call a set of viewpoints $V' \subseteq V$ a *texture covering* set, if every triangle in T' is visible from at least one viewpoint in V'. The *viewpoint selection* problem is defined as the process of finding a texture covering set V_S with minimum size among all texture covering subsets of V.

2.3 Viewpoint Assignment Problem

Once the texture covering set V_S is identified, the viewpoints $v \in V_S$ must be assigned to the participating users $u \in U$ (by including viewpoints in their participation plan) such that the total number of the TIN triangles $t \in T'$ covered by users within the campaign time T_C is maximized. Formally, the problem of *viewpoint assignment* is defined as an optimization problem, $arg\,max_c \ TS(P_U)$, to find the set of participation plans $P_U = \{P_{u_1}, P_{u_2}, \ldots, P_{u_m}\}$ corresponding to the users u_1, u_1,…,u_m in U such that $TS(P_U)$ is maximized while each P_{u_i} satisfies the corresponding user constraints c_i.

With GeoSIM, we assume users can join the texture documentation campaign *progressively* (not necessarily at a single time instant), with a poisson arrival distribution. Accordingly, we generalize the definition of the viewpoint assignment problem by considering an iterative viewpoint assignment scheme. With this scheme, the campaign time T_C is divided into equi-length epochs, I_1 to I_l, and viewpoint assignment is repeated at each epoch to assign the remaining uncovered viewpoints (those viewpoints that are not covered at previous epochs) to the users who arrive within the current epoch I_i. With iterative viewpoint assignment the optimization problem defined above is generalized and modified to $arg\,max_c \ TS(P_{U_i})$, where $U_i \subseteq U$ is the subset of users arriving during the i-th epoch I_i.

3 On-the-Fly Planning for Participatory Texture Documentation

We briefly describe our proposed solutions that enable on-the-fly viewpoint selection and viewpoint assignment in Sections 3.1 and 3.2, respectively. The detailed description of our solutions can be found in our extended papers [7] and [6].

3.1 Viewpoint Selection Solution

With [7], we propose an efficient heuristic, termed *GVS* (short for *Greedy Viewpoint Selection*), which allows for approximate but real-time viewpoint selection with approximation guarantees. In essence, GVS is a greedy heuristic that solves a given instance of the viewpoint selection problem by *reduction to* the corresponding instance of the classical minimum set-cover problem. With our experimental results, we show that GVS finds the minimum texture covering set V_S for an area as large as Los Angeles County (covering 1183 square miles with numerous objects) in a few seconds. In this case, V_S only includes 17% of the viewpoints in V.

3.2 Viewpoint Assignment Solution

With [6], we propose two families of efficient heuristics that enable on-the-fly viewpoint assignment: individual-based heuristics and group-based heuristics. With individual-based heuristics, we generate each user participation plan exclusively, independent of those of other users. Toward that end, we reduce the viewpoint assignment problem for a single user to the classical problem of orienteering [3], and accordingly adopt and extend the most recent heuristic solutions for the orienteering problem [4, 5] to implement viewpoint assignment. Individual-based heuristics are efficient and allow for on-the-fly viewpoint assignment; however, due to their exclusive nature, the participation plans generated by these heuristics may significantly deviate from optimal plans.

Alternatively, with our group-based heuristics we consider all users as a united group of participants. This allows for optimizing the assignment of the viewpoints among all users as a group; consequently, group-based heuristics can potentially generate near-optimal plans while maintaining high efficiency. In particular, group-based heuristics implement viewpoint assignment as a two-stage process. The main idea is to break the viewpoint assignment problem into multiple disjoint and smaller subproblems (at the first stage), where each subproblem takes a limited number of viewpoints as input and, therefore, can be solved efficiently (at the second stage). Accordingly, at the first stage group-based heuristics use various measures (e.g., proximity of the users to the viewpoints) to partition the set of viewpoints in V_S into a number of subsets, one subset per each user. At the second stage, similar to the individual-based heuristics, an orienteering heuristic is adopted to assign a subset of the viewpoints in each partition to the corresponding user of the partition.

4 GeoSIM: A Participatory Texture Documentation System

Figure 2 illustrates the client-server architecture of our PTD prototype system, dubbed GeoSIM. As depicted in the figure, the GeoSIM server consists of two engines, the planning engine and the texture mapping engine, which plan users participation and successively map the collected images, respectively. In addition to the viewpoint selection and viewpoint assignment modules which correspondingly implement our viewpoint selection and viewpoint assignment solutions described in Section 3, the planning engine includes a pre-imaging module that simulates the required images by imaging the corresponding area of the 3D environment model. The pre-imaged/simulated images are used to direct users in taking the required images properly.

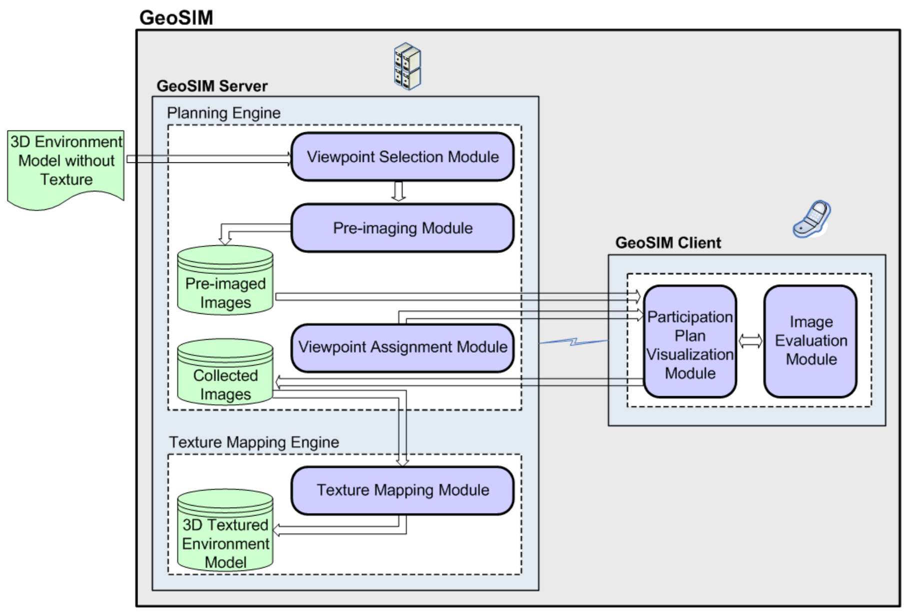

Figure 2: GeoSIM Architecture

The GeoSIM client (which is implemented as an Android application) comprises of two modules. The visualization module uses a Google Map^{TM} based interface to take user constraints, visualize the assigned user participation plan, and direct the user to take the required images. On the other hand, the image evaluation module considers various image features such as orientation, blurriness, and lighting to evaluate the quality of the images collected by the user. Accordingly, user is asked to re-take the images that are rejected by the evaluation module. Figure 3 shows the self-explanatory workflow of the participatory texture documentation process with GeoSIM. See [1] for more details about GeoSIM.

5 Related Work

Various commercial systems and research prototypes (e.g., Microsoft's Photosynth [2]) are developed that allow for texture mapping based on the images acquired by commodity cameras. In contrast, our focus is on effective planning to collect the images rather than merging the collected images for texture mapping.

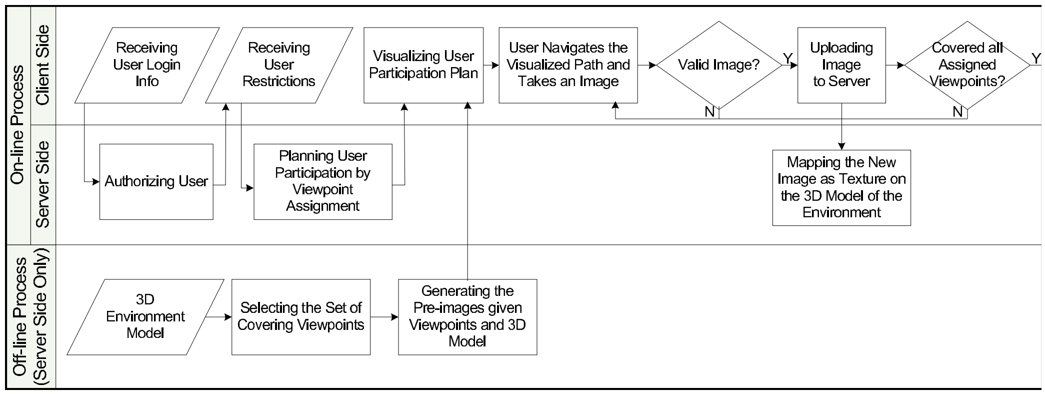

Figure 3: GeoSIM Workflow

6 Conclusion and Future Work

In this paper, we introduced the problem of geospatial planning for effective collection of texture information from urban environments. We also presented efficient heuristics that enable on-the-fly planning and described GeoSIM, the research prototype we have developed based on our solutions for participatory texture documentation. As part of our future work, we plan to extend GeoSIM to allow for documentation of data with other modalities, such as sound and temperature, and pollution.

Acknowledgement

This research has been funded in part by NSF grants IIS-0238560 (PECASE), IIS-0534761, and CNS-0831505 (CyberTrust), the NSF Center for Embedded Networked Sensing (CCR-0120778) and in part from the METRANS Transportation Center, under grants from USDOT and Caltrans. Any opinions expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

[1] Geosim. http://infolab.usc.edu/projects/GeoSIM. Last Update: April 2010.

[2] Photosynth. http://photosynth.net/. Last Update: April 2010.

[3] I. Chao, B. Golden, and E. Wasil. A fast and effective heuristic for the orienteering problem. *European Journal of Operational Research*, 88(3):475 – 489, 1996.

[4] C. Chekuri, N. Korula, and M. Pál. Improved algorithms for orienteering and related problems. In *Proceedings of SODA*, January 2008.

[5] K. Chen and S. Har-Peled. The euclidean orienteering problem revisited. *SIAM Journal of Computing*, 38(1):385–397, 2008.

[6] H. Shirani-Mehr, F. Banaei-Kashani, and C. Shahabi. Efficient viewpoint assignment for urban texture documentation. In *Proceedings of ACMGIS*, November 2009.

[7] H. Shirani-Mehr, F. Banaei-Kashani, and C. Shahabi. Efficient viewpoint selection for urban texture documentation. In *Proceedings of GSN*, July 2009.