Interactive Fusion of 360° Images for a Mirrored World

Geollery.com

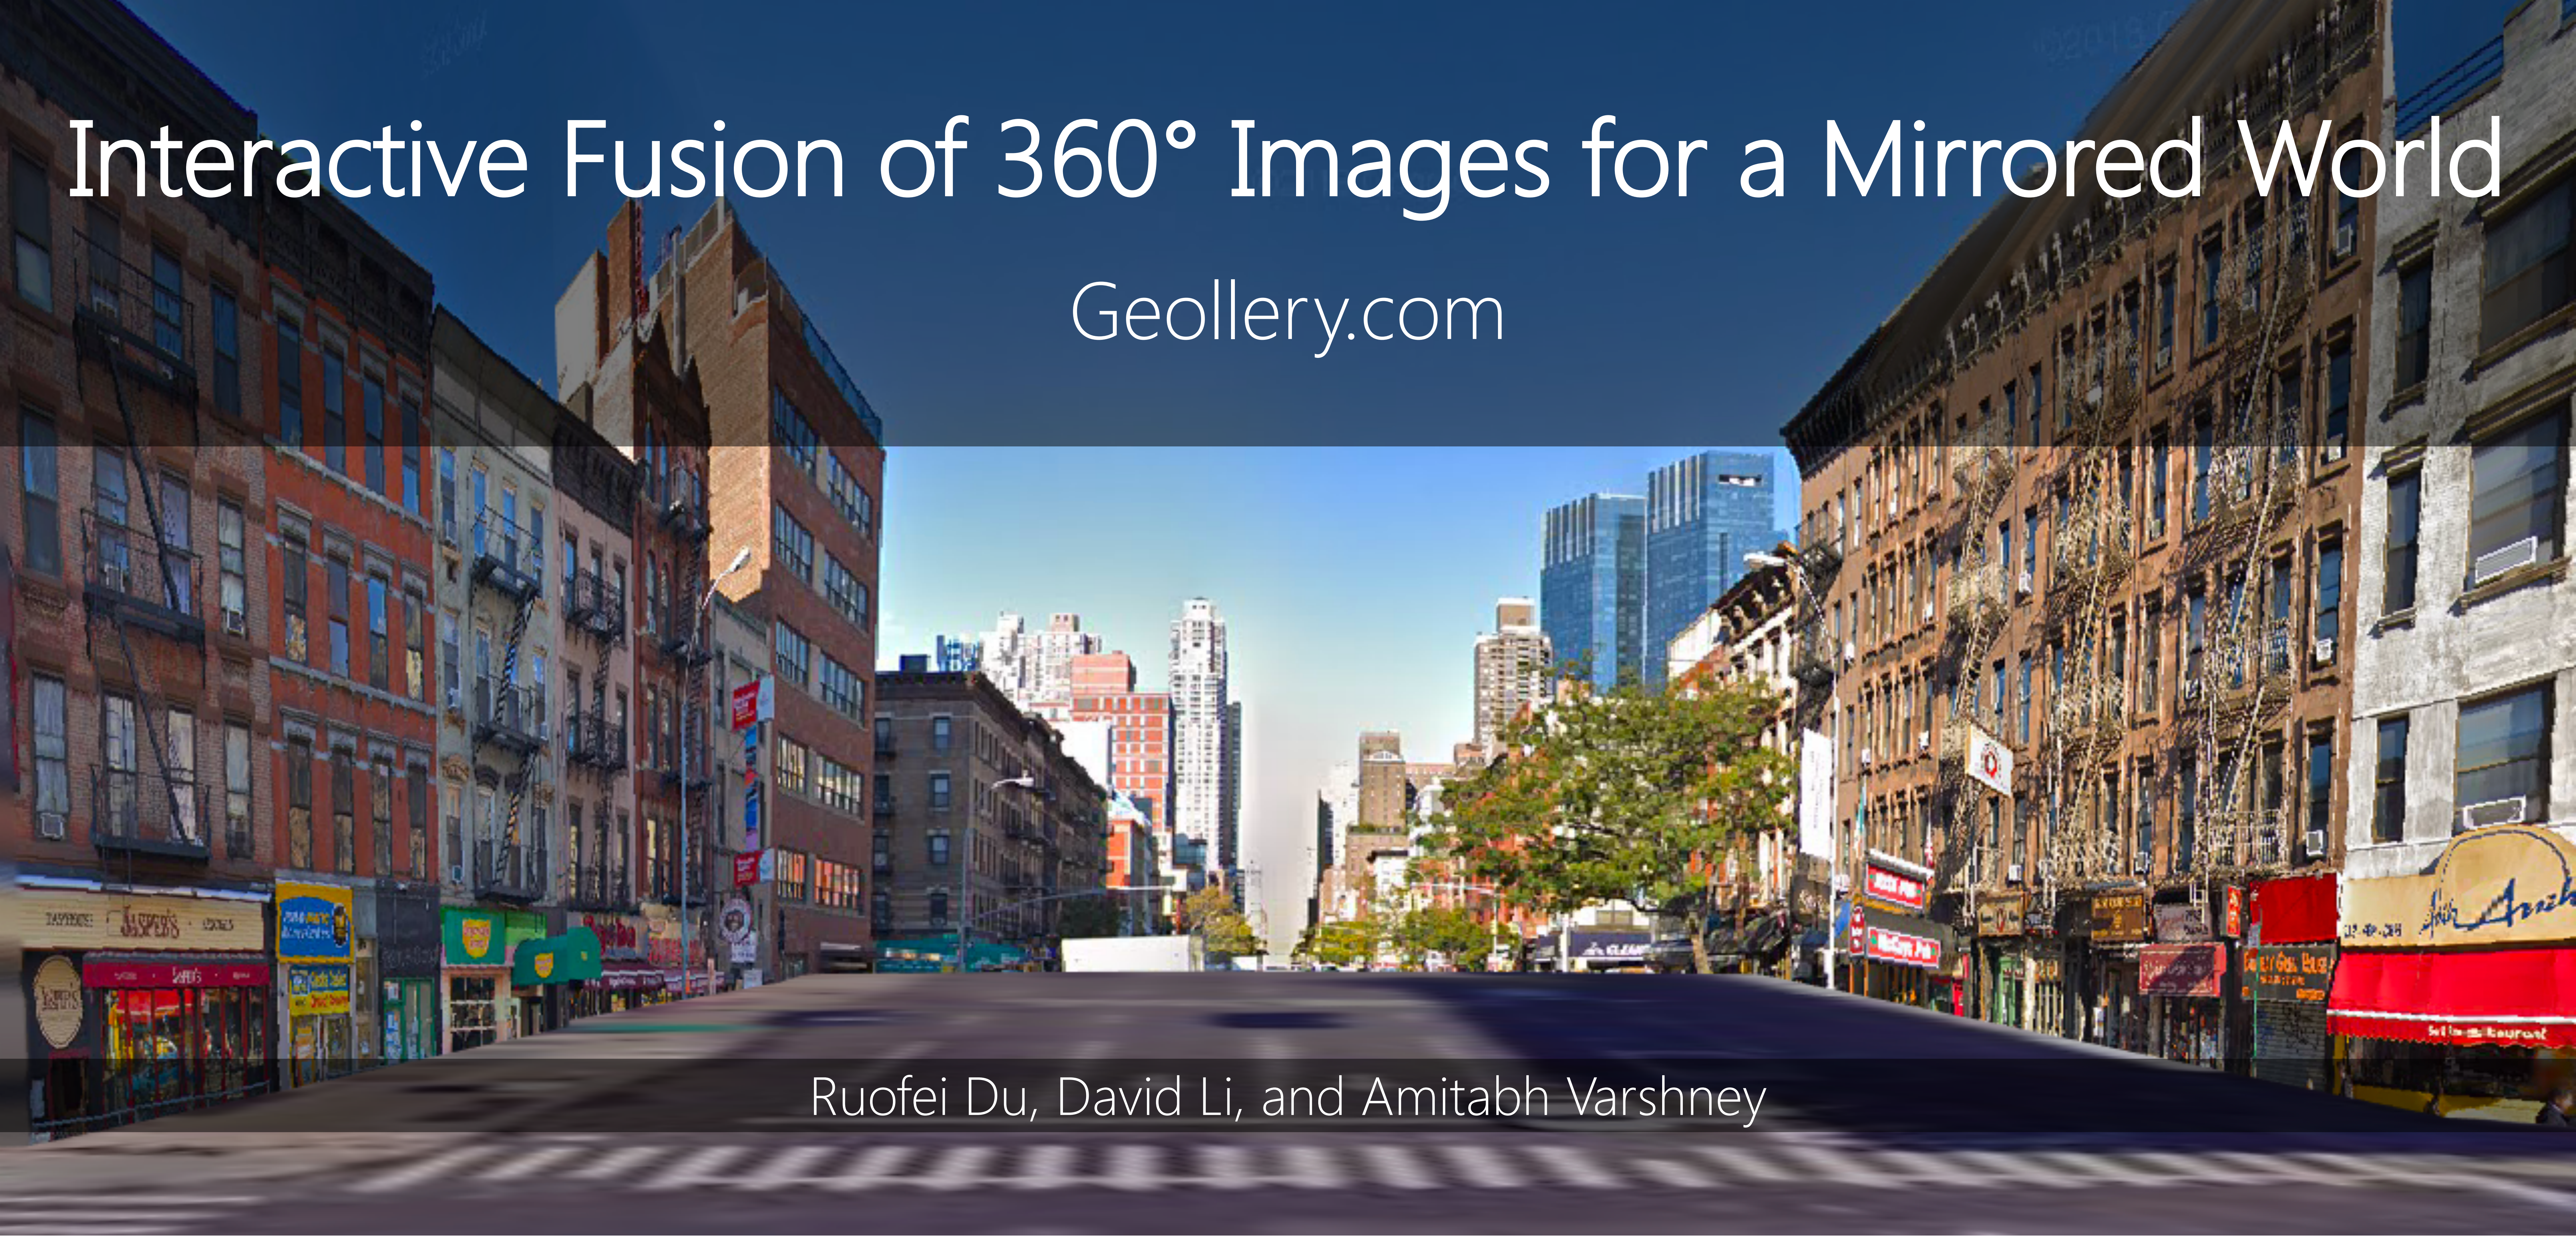

Ruofei Du, David Li, and Amitabh Varshney

Introduction

3D representations of the real world are widely used for virtual tourism, geographical education, and neighborhood auditing. Yet the availability of large-scale 3D models has been limited due to the high data, computational processing, and human labor requirements for creating such models.

In Geollery.com, we introduce an interactive pipeline to create on-demand 3D mirror world environments in real time. At a fine level of detail for close-up views, we render textured meshes using adjacent local street views and depth maps. When viewed from afar, we apply projection mappings to 3D geometries extruded from building polygons for a coarse level of detail. Our techniques are used in a mixed-reality social platform, Geollery.com, and we validate our real-time strategies on various platforms including mobile phones, workstations, and head-mounted displays.

System Overview

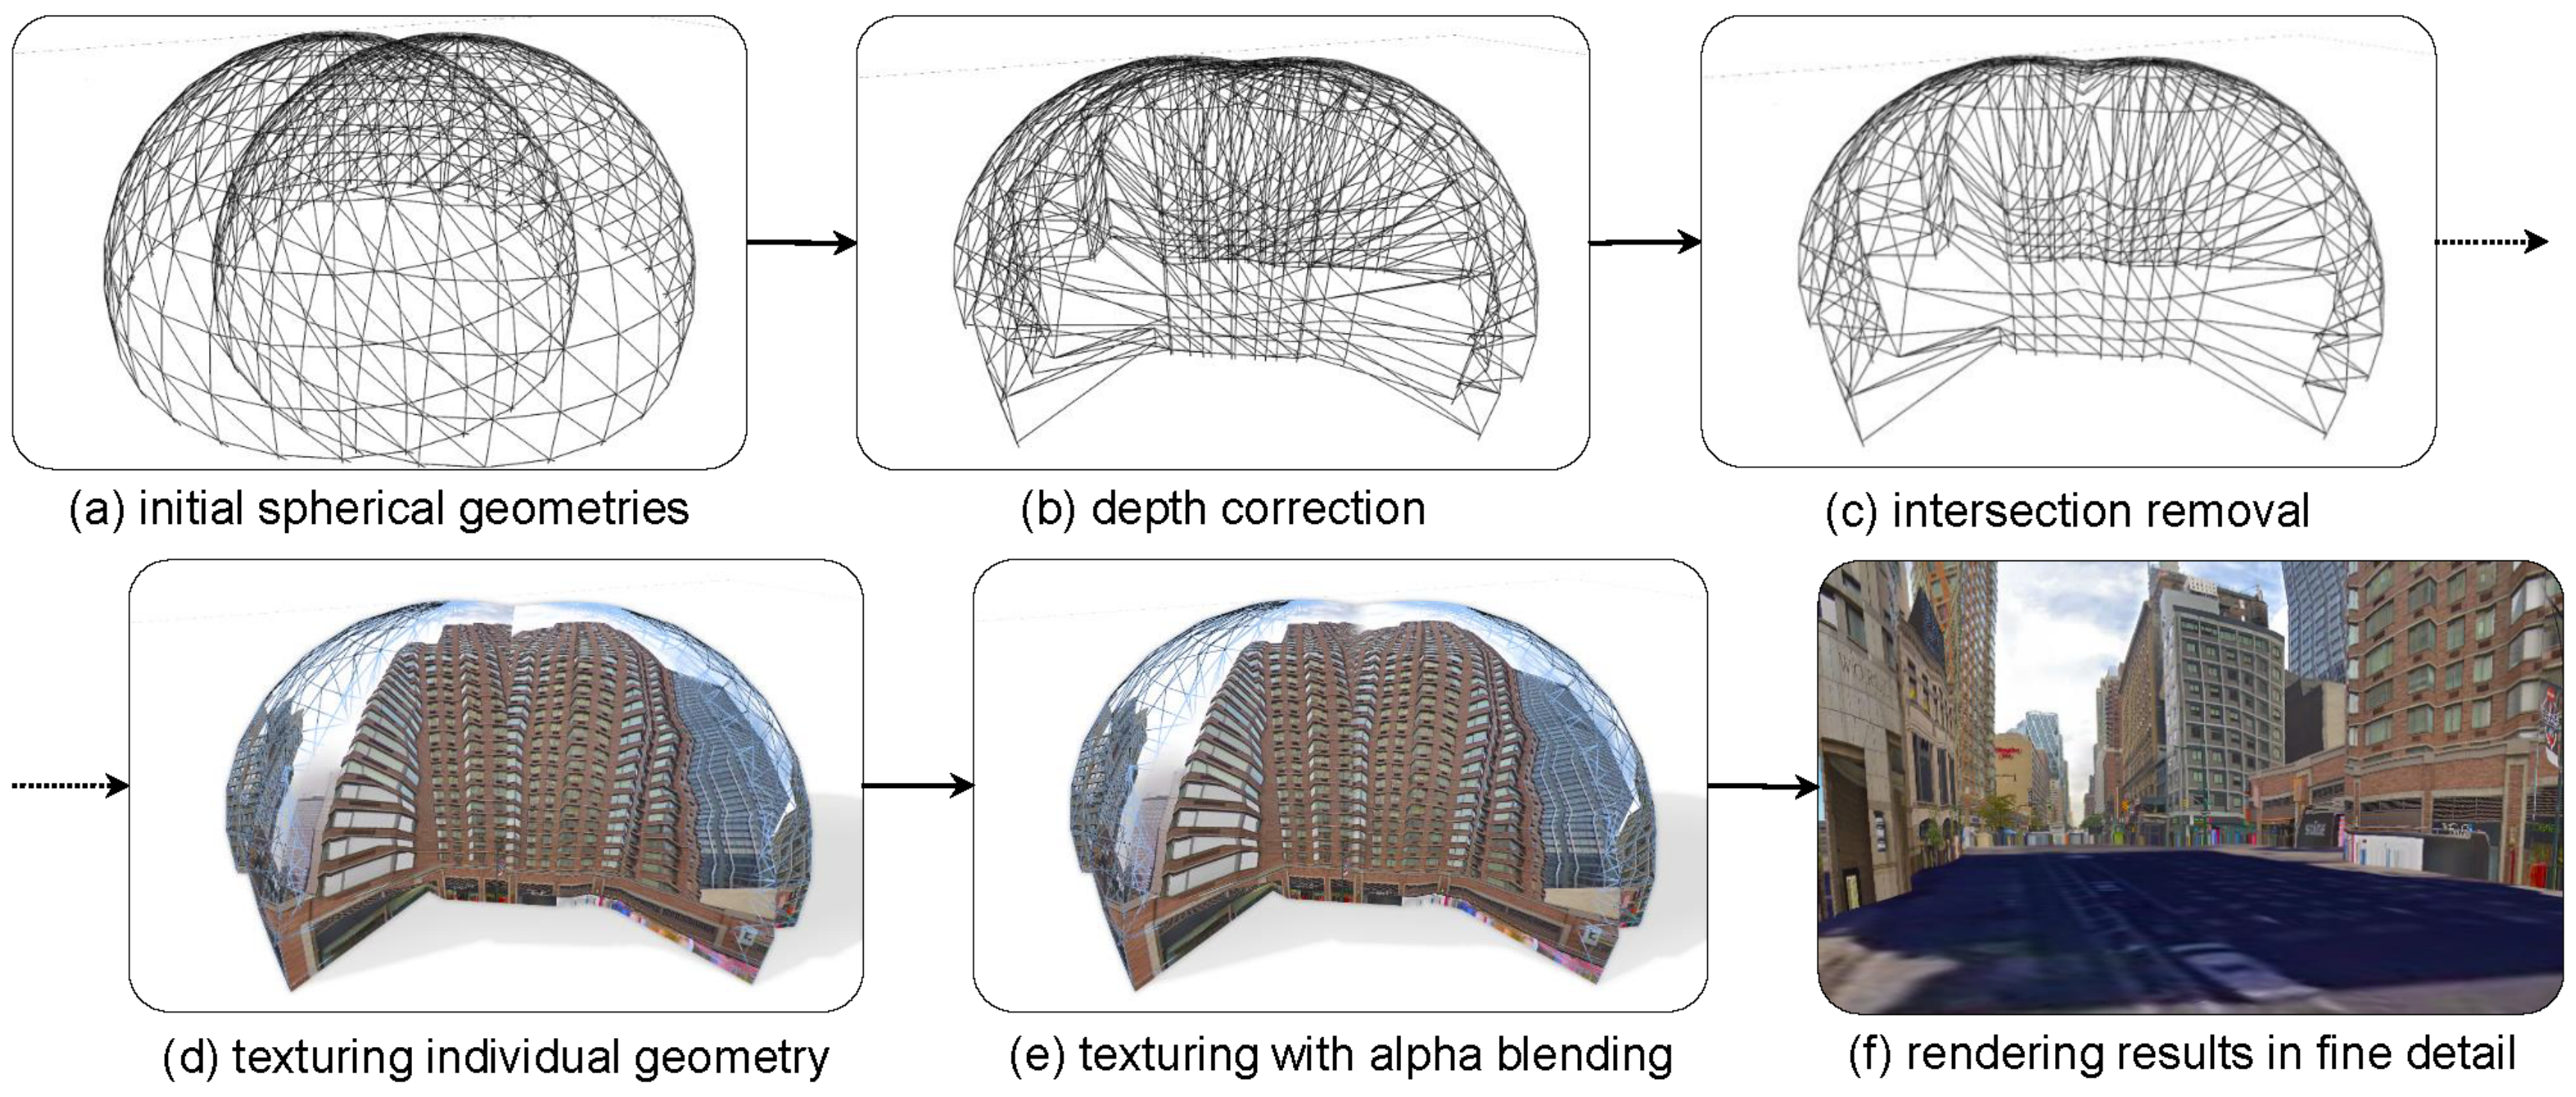
(a) initial spherical geometries
(b) depth correction
(c) intersection removal
(d) texturing individual geometry
(e) texturing with alpha blending
(f) rendering results in fine detail

In its current form, Geollery is a web-based application, which is available on desktop, mobile phones, and virtual reality headsets. The architecture of Geollery consists of a 3D world generator, a social media scraper, a distributed SQL database, a web-server powered by Apache and PHP, and optional modules such as a temporal filter, a geospatial filter, and a keyword filter. The rendering system is powered by three.js and our custom GLSL vertex and fragment shaders. Our system streams 2D map tiles, 2D building metadata, and Google Street View images based on the user's selected or real-world location. When this data is available, our system begins to construct an 3D environment. As the user moves around in the 3D environment, we dynamically stream new data and update the environment's geometries and textures.

Visualization Modes

Our system reconstructs the physical world at two levels of detail: a course level for far-away viewing and a fine level for close-up viewing. The level of detail dynamically changes depending on the position of the user's camera when rendering.

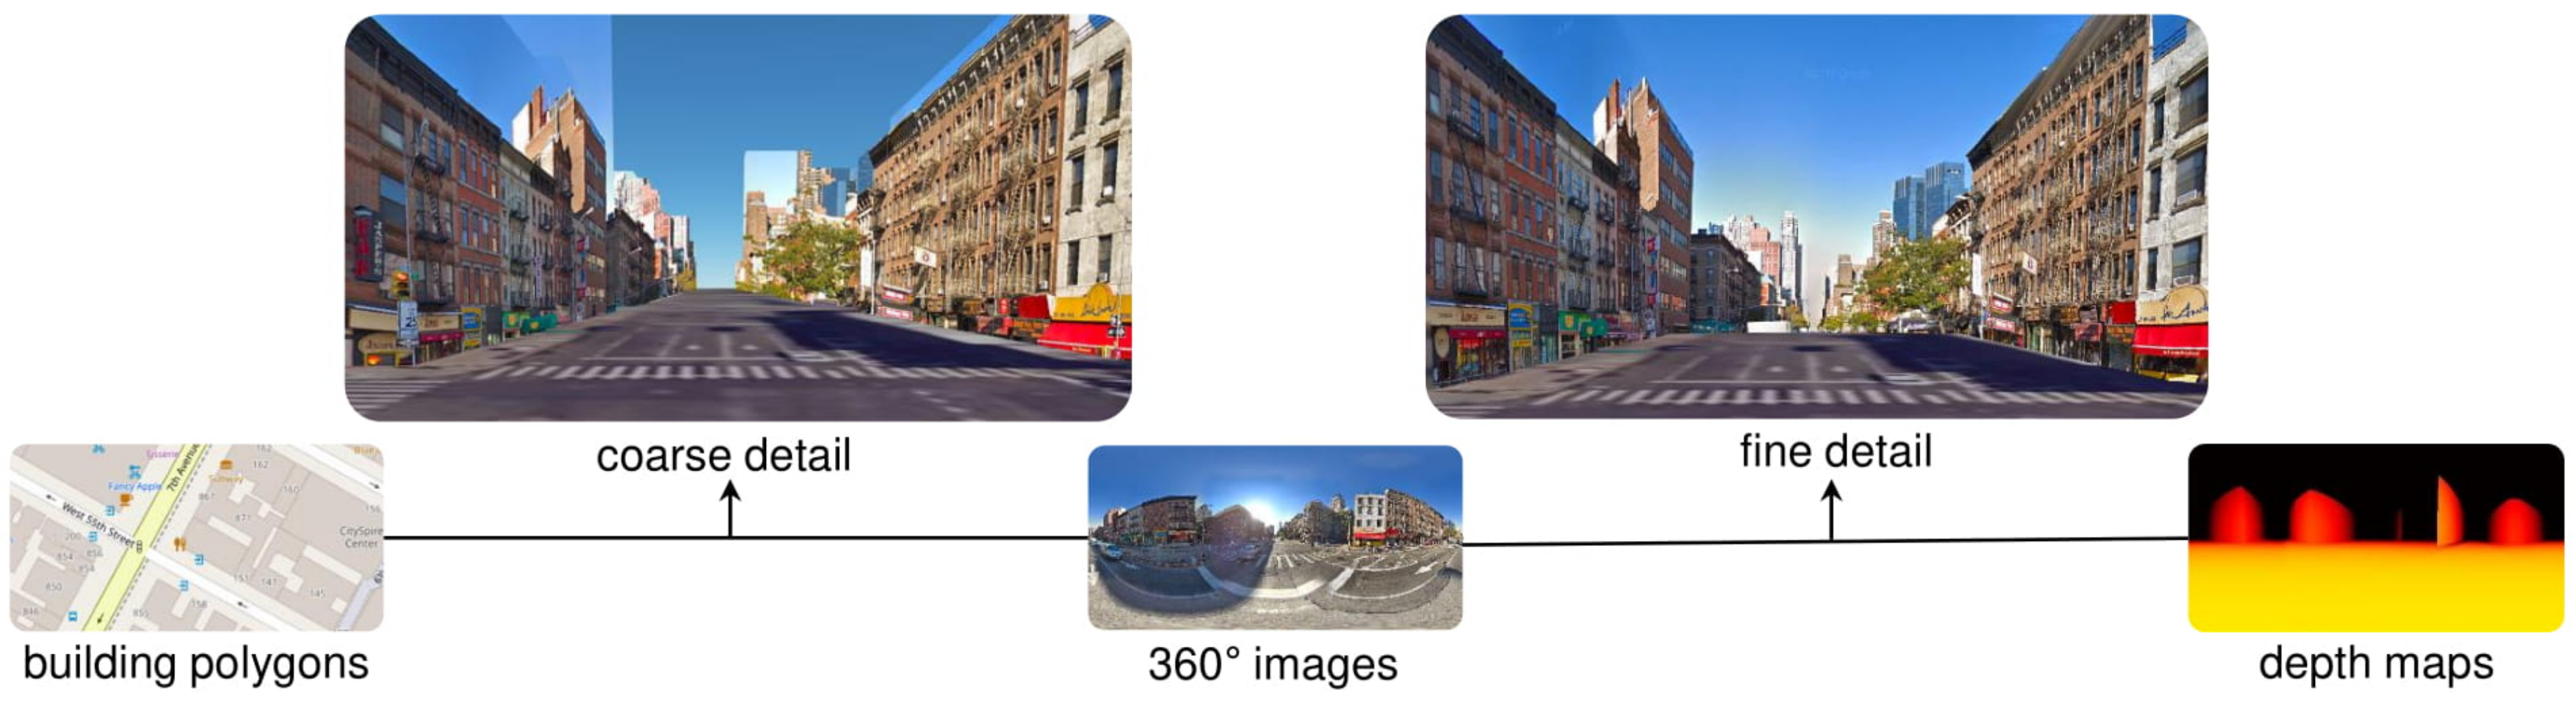

When viewed from afar, our system generates 3D buildings by extruding 2D building polygons from OpenStreetMap building metadata. Different sides of these building are then textured dynamically on the GPU with a fragment shader sampling colors from the nearest Google Street View images.

When viewed up close, our system generates a 3D skybox with a vertex shader positioning vertices at the proper depth relative to the position of each street view image. Then the fragment shader textures the interior of the geometry based on Google Street View images.

Future Work

In the future, we intend to refine our 3D reconstruction and augment it with additional data sources such as live surveillance videos. By caching street view images and building textures, we intend to progressively texture multiple sides of buildings.

UNIVERSITY OF MARYLAND

UMIACS

THE AUGMENTARIUM
VIRTUAL AND AUGMENTED REALITY LAB
AT THE UNIVERSITY OF MARYLAND

COMPUTER SCIENCE
UNIVERSITY OF MARYLAND

Ruofei Du
me@duruofei.com

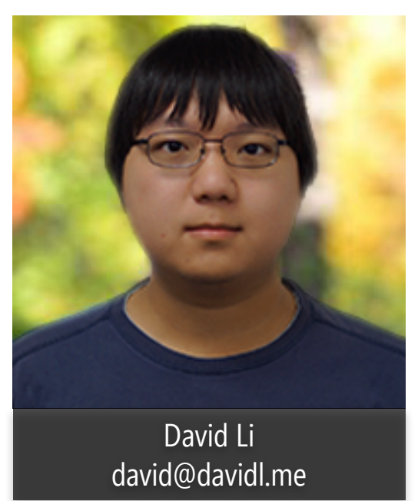

David Li
david@davidl.me

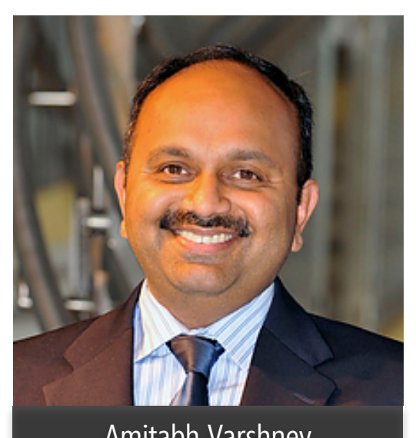

Amitabh Varshney
varshney@cs.umd.edu