

Similar Textures by Procedural Approximation

Context:

Interactive graphics applications suffer from an unprecedented increase in the amount of graphical content required to depict virtual worlds. Environments are typically represented by a mixture of polygons capturing the scene geometry and images capturing the appearance of the various surfaces. These images, called *textures*, are not only very long to create for CG artists but also occupy a large amount of space, often more than available memory. This problem worsens with urban scenery, where an entire city must be modeled. This is especially unfortunate as many applications are based on such urban environments: Navigation, urban planning, video games, emergency response training, visualization of dynamic data over a city (pollution, traffic, etc.).

Several directions have been explored to solve these issues: Image compression, data streaming from a remote server, repetition of a same image on multiple surfaces. However, these solutions typically focus on a single aspect of the problem: Creation, storage, loading or display. Often, the approach solves one aspect to the detriment of another: For instance compression algorithms are not well adapted to display at interactive frame-rates, progressive data loading introduces visible slow-downs and repetitions impair the ability to navigate as all areas become visually similar.

Proposed PhD topic:

This project proposes to study the problem of texture management as a whole, in the context of urban scenery. Our key insight is that in many cases only a small subset of the scene must be strictly identical to the reality. For instance, landmarks and historical buildings must be accurately reproduced in all details. However, for the vast majority of the other parts of the scene we only seek for a plausible reproduction, not a strictly similar one. Instead of reducing the resolution of the images on these buildings, hence introducing a very noticeable lack of details, we propose to capture the appearance of the buildings with a *procedural* model. That is, a compact data structure encoding only the important information (positions of windows, doors, type of materials) without capturing the details (the position of each brick, each cracks). The details will be generated only at the time of display, hence not requiring any storage. We can see that our approach has to consider the problem of content management as a whole: We need to infer procedural models from photographs, store them in compact data structures, load this data progressively as the user explores the scene, and finally display it at interactive frame rates. Each step corresponds to scientific and technological challenges: Capturing important information in a photograph (i.e. separating semantic from detail), studying data structures allowing compact storage and fast decoding, determining what will remain hidden or what will become visible for efficient loading, and finally generating details only in last resort to maintain small memory usage while enforcing interactive frame rates.