Reconstructing cities in the real world

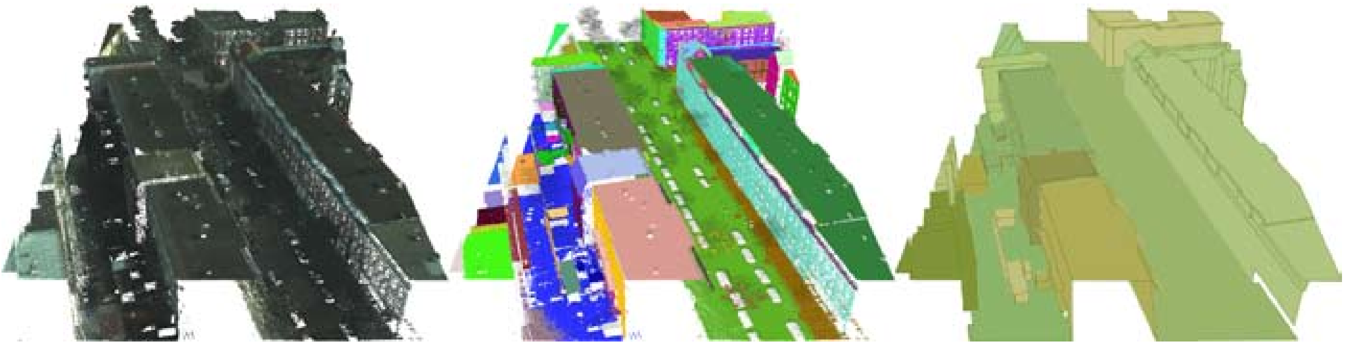

Modeling the Virtual World

Automating reconstruction to quickly create virtual copies of real cities.

Workpackage
1.1 Automatic World Generation Based on Real Data

Partners
Utrecht University
TNO Defense, Security and Safety

Budget
500.000 euro

Key Publication
Van Kreveld et al. (2009). Identifying well-covered minimal bounding rectangles in 2D point data. Proc. 25th European Workshop on Computational Geometry, pp. 277-280

Contact details
Remco Veltkamp
Utrecht University
Remco.Veltkamp@cs.uu.nl

Emergency services, municipalities, and location games require a detailed and up-to-date virtual copy of a real-world urban scene. Automating reconstruction can greatly reduce the time and cost of the process. New methods have been developed that can efficiently perform steps in recreating the world. Soon the whole world will be completely modeled in 3D.

To create a geometric 3D model of the real world, we need to reconstruct the scene from data sources such as images and laser range. The detail and complexity of the reconstructed model depend on the amount and quality of the data, while the time needed for reconstruction depends on the efficiency of the methods, especially when using larger data sets. This project is aimed at automatically reconstructing geometric models from laser range scans of large real-world urban scenes. These scans result in large clouds of points in 3D space. The methods should efficiently reconstruct a detailed and precise representation of the real scene from these points.

Bounding rectangular surfaces

The reconstruction process can be split into different parts like identifying existing infinite surfaces that pass through many points, bounding the points in these surfaces, and filling the remaining gaps where data is missing. Urban scenes contain many simple surfaces, and for the typical scene about a quarter of the surfaces are rectangular. We developed a new method for efficiently bounding the appropriate surfaces using rectangles. After different surfaces are identified in the data, each surface is checked to see if it should be bounded by a rectangle. We have developed a method to determine whether the data in a surface is well covered, such that there are no large parts in the rectangle that are void of points. The Netherlands Forensic Institute (NFI) has provided a laser range data set, and our method efficiently identifies surfaces in this data set and provides the rectangles that correctly bound each surface. Once rectangular surfaces are identified and bounded, the remaining part of the scene may be searched for increasingly more complex shapes. Alternatively, the shapes bounding the surfaces may be based on the data distribution instead of predefined forms.

Data-driven boundaries and stitching

Surfaces may be bounded using predefined shapes like rectangles, triangles, L-shapes, etc. or their boundary shape may be determined from the local data distribution. We are currently developing a method for data-driven boundary creation. This method not only takes into account the data measured in one surface, but also the shapes of neighboring surfaces. The method should result in a shape that is easy to connect to neighboring surfaces, while behaving nicely in the presence of noise and missing data.

"Soon the whole world will be modeled in 3D"

Because of measurement problems like low data resolution and occlusion, some parts of the real-world scene may not be captured in the data. Stitching together the bounded surfaces that are present in the data may reveal these holes in the data. We will develop techniques that can identify missing parts in the data and fill these parts in a realistic way.

From pictures to 3-D

MODELING THE VIRTUAL WORLD

Building virtual models of the world using pictures.

Virtual models of the real world are typically built by a specialized team of designers and developers. Consequently, this costs a lot of time, effort and money. Now, imagine being able to do away with this by just going outside, taking some pictures of the desired scene, uploading them to a computer and pressing a key.

Virtual models of real world scenes are commonly required for simulations, entertainment purposes,

“Virtual representations of existing environments can be automatically built from images”

navigation and many others. Instead of having to model these by hand, many researchers are exploring ways to do this automatically from a variety of data sources like imagery and laser measurements. Our objective is to be able to exploit the vast amount of real world imagery that is available, like the panoramic images that are acquired in large volumes, or the many image databases that are on the internet, and create complete 3D models from that data.

Humans are smart, computers are not?

Estimating the 3D structure of what we see around us is something our eyes and brain do so easily that we don’t even realize it. For a computer, working from captured images this is a lot harder. The problems start with finding corresponding points between the images. Research has turned out that even the most state-of-the-art methods are often not able to find out which parts of two images are the same when the images of a scene are taken from quite different perspectives, yet this is a key requirement in order to be able to calculate depth. Our research has analyzed the points of failure for commonly used point matching techniques. Based on these finding we try to propose improved matching algorithms that are effectively computing corresponding points between images, even when the images are taken from quite different perspectives. The matched points then are easily converted into 3D points that describe the scene.

You have a point there, but now what?

Developing new methods to calculate the 3-dimensional position of points in the scene is but the first step. A set of points is not yet a model that describes surfaces, joined together at the edges to form a closed whole. Many existing algorithms try to find these surfaces from the 3-D points only. In our research we plan to combine the points that were found with additional information contained in the images. By virtually slicing the images apart into segments along the edges of components in the scene, we identify the surfaces which are needed in the model. Exploiting the principle that ‘3 points define a plane’, by combining the segments with the position of the points, it is possible to calculate where the surfaces of the objects are in the real world. Combining these surfaces then provides the virtual model of the real world environment, exactly what is so desired.

Workpackage
1.1 Automatic World Generation Based on Real Data

Partners
TNO Defence, Security and Safety
Utrecht University

Budget
400.000 euro

Contact details
Frido Kuijper
TNO Defence, Security and Safety
frido.kuijper@tno.nl

Sensibly stuffing game worlds

Modeling the Virtual World

New techniques to orderly populate virtual worlds with smart objects.

Workpackage
1.2 Automatic creation of imaginary worlds

Partners
Delft University of Technology

Budget
500.000 euro

Key Publications

T. Tutenel et al. (2008). The role of semantics in games and simulations. ACM Computers in Entertainment 6(4):a57

T. Tutenel et al. (2009). Rule-based layout solving and its application to procedural interior generation. Proc. CASA workshop on 3D advanced media in gaming and simulation, pp. 15-24.

T. Tutenel et al. (2009). Using semantics to improve the design of game worlds. Proc. 5th conference on artificial intelligence and interactive digital entertainment, pp. 100-105.

Contact details

Rafael Bidarra
Delft University of Technology
r.bidarra@tudelft.nl

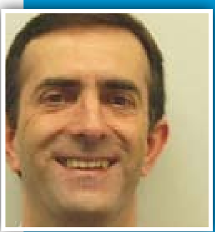

In huge urban game worlds most buildings are inaccessible since manually designing all their interiors is simply unaffordable. Enriching objects with real-world logic and mutual relationships can open up more interactive, convincing and immersive interiors. We developed semantic layout solving techniques that facilitate the automatic generation of such lively environments.

Objects in virtual worlds are mostly related to other objects in their surroundings. A desk without a chair or a computer without a keyboard appear weird, so in a working office we see these items placed together, following some logical rules. This project investigates ways to enrich virtual objects with information about their functionality and relationships to other objects, in order to automatically find suitable locations for them in a sensible context. We developed a layout solving approach that uses these semantically-rich objects to procedurally generate and populate realistic layouts for virtual environments, e.g. building interiors, gardens or city streets.

Semantic layout solving

We developed a semantic library with many classes of objects, each describing their type, functions, services, possible relationships to other objects, and also how players or AI characters can interact with them. For instance, you can define specific clearance areas around an object, e.g. to make sure there is enough space in front of a cupboard to open it.
Using this information, our semantic layout solver takes a description of all objects desired in a given scene and incrementally determines sensible locations for each of them, in order to create a complete and meaningful layout.
To specify the contents of generic scene types, e.g. a kitchen, a living room or an office, we defined an intuitive description language, mainly focussing on the object types that can or should be present in such scene, as well as on possible, scene-specific relationships and qualifiers. Based on these descriptions, every instance of these scene types can be laid out using the semantic layout solver.
Eventually, every scene instance can also be automatically refurbished or decorated using semantic filters, typically linked to some predicate from the semantic library (e.g. dirty, antique or messy), further customizing its visual details without altering its specific intent.

Rapid game prototyping

Using the above descriptions and layout solver, level designers can quickly create virtual worlds. Since all its objects include functional information, sensible object interaction is at once embedded into these virtual worlds, making them ideal as exemplar worlds for rapid prototyping. We are investigating integrated methods to specify both semantic descriptions for game worlds and gameplay-related object interaction in order to test gameplay aspects much before level designers finished the entire world.
Eventually, these prototype worlds have to be manually refined by designers. We are currently investigating how to guarantee that all sensible relationships set among objects are kept in that process. For this, in close cooperation with the other Workpackage 1.2 project (see next page), we are developing semantics-based techniques for maintaining world consistency after each manual operation. Together, such techniques can allow designers to create game worlds quickly and efficiently without losing control over the final product.

> *"Instant generation unleashes the all-access pass to an enter-anywhere world"*

From sketch to virtual world

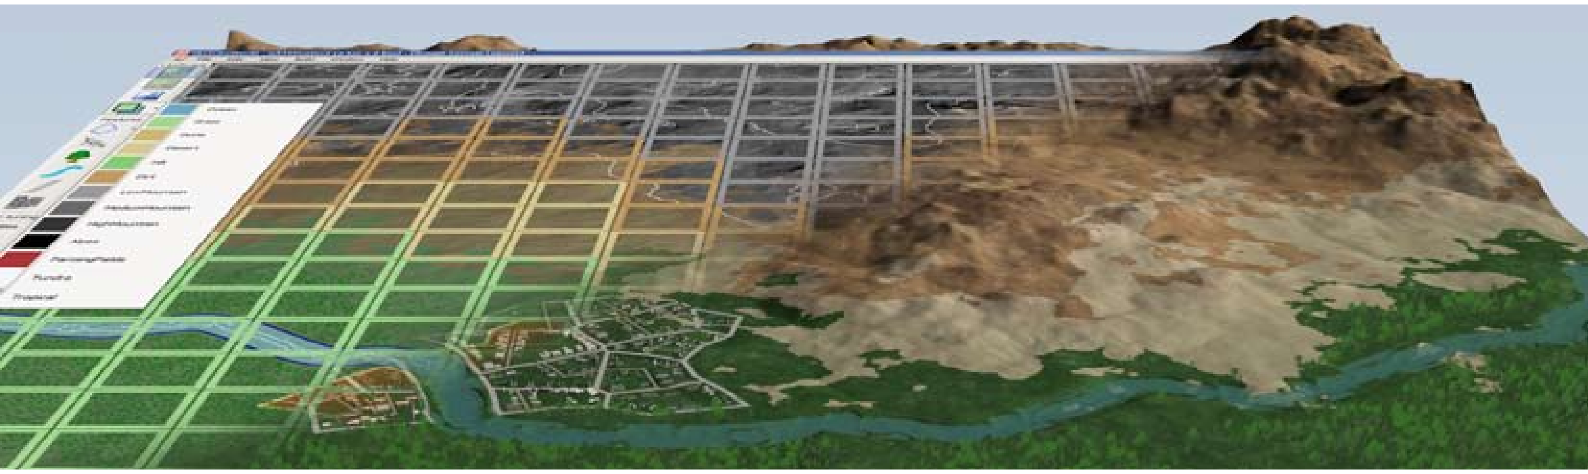

Modeling the Virtual World

Enabling non-experts to create whole virtual worlds in minutes.

Virtual world creation mostly starts with designers sketching their vision. What usually follows is months of 3D modelling and manual fine-tuning every little detail of the world. Novel modelling and interaction techniques can spare much of this routinely work. Our prototype SketchaWorld, for example, generates virtual worlds from simple 2D sketches.

Automated content creation for virtual worlds has been around for years. However, it hasn't really caught on yet, as most virtual worlds are still modelled entirely by hand! This is because current content generation methods are often limited to creating one specific type of feature. They are typically very complex and difficult to use, and hard to steer towards a desired result. Furthermore, the generated content is hard to integrate into a complete and consistent virtual world. This project develops new content generation techniques and user interaction methods for making the automatic creation of virtual worlds accessible, controllable and efficient.

> *"Procedural sketching makes creating virtual worlds fast, fun and easy"*

Interactive procedural sketching

We developed a very intuitive content generation method called procedural sketching, for designers to interactively create their virtual world. Designers paint the landscape with colours representing steep mountain ridges, green hills, barren desert, etc. On top of this, designers draw terrain features, such as rivers, forests and cities, using simple lines and shapes. Meanwhile, as you sketch each element of the virtual world, it is automatically expanded to a realistic terrain feature. Furthermore, all generated features are automatically fit with their surroundings. For instance, a road's embankment is integrated in the landscape, and when it crosses a river an appropriate bridge is inserted in place. Automated content generation techniques as these will certainly never replace manual modelling. SketchaWorld's foremost application is enabling virtual world design by non-specialists, for instance, instructors creating scenarios for serious games. In addition, it is of valuable assistance for applications as concept design, rapid prototyping, and exploration of ideas. Moreover, the generated results can always serve as a sound basis for further enrichment using traditional methods. See www.sketchaworld.com for more details. We are currently cooperating with re-lion BV in making these techniques available for Dutch military instructors, to assist them in creating new scenarios for game-based training.

Mixing manual editing with automatic generation

Using procedural sketching, designers can create complete virtual worlds in minutes. However, to create an environment that exactly matches their intent, designers need more fine-grained control on the generated content, e.g. they may wish to modify the course of a river, shift a street crossing point, or place an important building in a city. We currently investigate the integration of procedural sketching with manual editing operations. For this, in close cooperation with the other Workpackage 1.2 project (see previous page), we are developing semantics-based techniques for maintaining the world consistency after each operation. Main challenges here include preserving past manual changes whenever an area is re-generated, and balancing user control versus automatic consistency maintenance. We believe that a seamless integration of procedural sketching and manual editing can offer designers the best of two worlds, and enable them to freely experiment, eventually creating the virtual world that precisely matches their intent.

Workpackage
1.2 Automatic creation of imaginary worlds

Partners
TNO Defense, Security and Safety
Delft University of Technology

Budget
400.000 euro

Key Publications

Smelik et al. (2010). Declarative Terrain Modeling for Military Training Games. International Journal of Computer Game Technology, vol. 2010.

Smelik et al. (2010). Interactive Creation of Virtual Worlds Using Procedural Sketching. Proc. Eurographics 2010 - Short Papers.

Smelik et al. (2010). Integrating procedural generation and manual editing of virtual worlds. Proc. Fifth International Conference on Foundations of Digital Games 2010.

Contact details
Klaas Jan de Kraker, TNO
klaas_jan.dekraker@tno.nl

Affect in virtual environments

Modeling the Virtual World

Attention to affective qualities improves the validity of virtual environments.

Workpackage
1.3 CAVA: Creating Ambience by Visual and Auditory Means

Partners
Utrecht University
TNO Human Factors

Budget
500.000 euro

Key Publications

Toet, A., Welie, M. van & Houtkamp, J.M. (2009). Is a dark environment scary? CyberPsychology & Behavior, 12(4), 363-371.

Houtkamp, J.M., Schuurink, E. & Toet, A. (2008). Thunderstorms in my Computer: The Effect of Visual Dynamics and Sound in a 3D Environment. In M. Bannatyne & J. Counsell (Eds.), Visualisation, 2008 International Conference (pp. 11-17).

Los Alamitos: IEEE Computer Society. Houtkamp, J.M., Spek, E.D. van der & Toet, A. (2007). The influence of lighting on the affective qualities of a virtual theater. In Predicting the Future, 25th eCAADe Conference Proceedings (pp. 77-84). Frankfurt am Main.

Contact details
Joske Houtkamp
joske.houtkamp@cs.uu.nl

Alexander Toet
lex.toet@tno.nl

Virtual environments used for visualisation and training often do not convey the ambience of a real environment, and the urgency of a serious event. Visual and auditory effects can be used to influence the emotional response of a viewer, and enhance the validity and effectiveness of the application.

Virtual environments are used to create convincing environments for training, and visualisations for urban planning. It is important that these environments not only display spatial characteristics, but also elicit affective responses similar to those experienced by a viewer in corresponding real environments. However, on a desktop computer or projection screen much of the visual and auditory information of a real environment is lost. In what way do characteristics of the modeled environment influence the affective response of the viewer, and can we determine this response?

Ceci n'est pas une pipe, this is not a pipe

Visual effects (dynamics, illumination) and sounds in a virtual environment affect the viewers' response differently than in a real environment. For instance, adding a lot of dirt to a virtual environment will not automatically lead to a negative response; viewers may find it an interesting element in an otherwise rather sterile and dull environment.
Although game technology creates increasingly convincing virtual environments on desktop computers, they are still impoverished versions of reality. Viewers however easily imagine and add information in their minds, which is not present in the model. They decide on for instance season, ambience and materials, using simple cues such as bright colours or plain textures. Modelers can thus influence the ambience of the virtual environments intentionally or unintentionally, while the viewer is not even aware of this.
On the other hand, users have to overcome obstacles with navigation and interaction, which creates distractions and diminishes involvement with the environment displayed. As in films, visual elements and sounds that are important must be emphasized and 'hyper real' to be noticed.
To increase the validity of our studies, we used environments that were developed by E-Semble, VSTEP, and Deltares for training and examination purposes.

Engagement with the virtual world

Our research shows that some very important cues to create a desired ambience in virtual environments, such as darkness, and soundtracks containing real sounds, do not have the expected effect on the viewers' response. We have evidence that these cues are only effective when users are involved with the environment because of personal relevance for them. In our next experiments, we will further explore these issues.

> *"A virtual environment must convey a real atmosphere"*

To attain the required affective response to any virtual environment, the effects on viewers' responses must be considered during the whole design and development process, taking into account users, tasks, and context of use. At the end of this project we will draw up guide lines for design and evaluation that developers can use to create valid 'affective' virtual environments.