

Ph.D. THESIS REVIEW
(confidential, do not distribute)

Title: Procedural Modeling and Realism in Computer Graphics
Author: Jan Benes

General Comments

This is my report on the Ph.D. thesis of Mr. Jan Benes, devoted to the assessment of the realism of procedural modelling techniques in Computer Graphics. For this, he performed a detailed analysis of the meaning and implications of the concept of realism in the field, studying different approaches to quantify it, with special emphasis in its application to procedural modelling, and modelling of buildings in particular.

Since its earliest days, Computer Graphics has witnessed a great enthusiasm in research applied to the generation of realistic models in an automatic way. This is certainly not surprising considering the great needs from the film and video game industries for larger and richer synthetic models.

What is quite original in this thesis is that Mr. Jan Benes takes the more challenging route of questioning the key concept of realism in Computer Graphics in general, and in Urban Procedural Modelling in particular, which is not only a challenging task, but also one that hides a deep problem that we have faced in Computer Graphics almost since its origins: the problem of what really realism is, and how much of it we really need in the field. Both questions are not only challenging, but are almost impossible to tackle without the proper tools, which Mr. Benes develops in a quite clever and skilful manner.

I consider that the thesis introduces innovative formulations to three important problems: (1) it poses the problem of what can be considered realism in the field of Computer Graphics, which is, by definition, an ill-posed problem as there are no universal definitions or unique mathematical formulations for such a subjective attribute; (2) the specialization of the mentioned problem to the procedural modelling of buildings and urban landscapes, which is of central importance for this research line, where researchers have rarely stopped to think about these crucial issues; and (3), the proposal of a model for generating procedural road networks between, inside to and near cities, mimicking historical aspects of their evolution in order to achieve, once more, more realistic results.

The proposed solutions address relevant problems, are well thought of, and introduce a wealth of well-adapted methods. The thesis should motivate new research directions, and allow finally having realistic procedural models that are, up to a certain point, indistinguishable from their real-world counterparts. The film and video game industries should benefit from it, but I can easily see much renewed interest from other

stakeholders, like furniture builders, urban planners or even architects, to mention just a few.

The thesis covers a lot of really complex material and concepts, and its careful illustrations help understand the most intricate parts of some of the proposed ideas, as well as the problems thus solved. The literature review covered all the previous work I knew of, also providing some new references I was not aware of.

Two publications already support the novel scientific aspects of the definitions of realism in Computer Graphics and Urban Procedural Models (Chapter 4); Computer Graphics Forum is a well-established and respected international journal in computer graphics. The second publication, the method for generating road networks inside and around urban centres, as well as formulating their coalescence when their evolution makes their areas overlap (Chapter 5), was also published in Computer Graphics Forum, demonstrating again the high quality of the research presented in this thesis.

In the following section, I will provide more details about what I feel are the main contributions of this thesis.

Detailed Comments

Chapter 1 motivates very well the needs for a robust, believable and mathematical definition of realism for the field of Computer Graphics, yet the difficulties associated to work with such subjective issues. The chapter is short, but very clear, and explains well where the contributions will come from. It made the road map to follow this thesis very much appreciated.

Chapter 2 dives into the wide spectrum of the different historic definitions of realism in Computer Graphics, especially those used in rendering to compare the result of rendered simulations with other images, which could be real photographs (e.g., the Cornell Box) or synthetic reference images. The literature review is very well organized. I particularly valued that Mr. Benes revisited some of the theory under different kinds of definitions of what realism is, once some goals had been explained in more detail. In particular, Section 2.5 introduces different ways of evaluating realism, including different metrics, in a quite clear and straightforward way. This made the understanding much easier and reading much lighter. It goes to demonstrate how Mr. Benes has a clear vision of what has been done, and how everything in this field fits together. This is a quite remarkable aspect of Mr. Benes's thesis, as it deals with some of the toughest, worst posed problems in Computer Graphics in a very elegant, yet practical, way

Chapter 3 specialises into the concept of realism in procedural modelling. This chapter continues the previous work by tying together the concept of realism in different areas of procedural modelling, from game objects (e.g., the creatures in games like No Man's Sky), to vegetation and buildings. Probably, the most important part of the chapter is Section 3.2, where Mr. Benes defines different characteristics of procedural rules, like their range, desirability, plausibility, realism and desirable outputs, just to mention a few.

It is in this chapter where I would like to make a couple of very minor comments. First of all, while reading Section 3.2.4, it seemed to me that Mr. Benes is assuming that all models generated by the rules are valid models, not considering the possibility of having broken or deformed models. In my own experience, all rules have parameters that must have a validity range, outside of which the results are not guaranteed by the author. For some models, these restrictions may seem trivial (e.g., the length of a wall should be larger than a trivial minimum, such as a couple of meters, and there is no

restriction on the maximum possible length). However, in many models I have seen, including procedural models created by professional designers, and even more in the case of complex buildings, there are parameters that could not vary arbitrarily, or that could not be fiddled with without changing other parameters accordingly. Just to give an absurd example, a window width cannot usually be made larger than the width of the wall where it is inserted. Or, just to provide a more realistic example, an asset cannot be stretched in one direction more than some fraction of the real proportions without it looking deformed and “wrong”. These two extreme examples serve to illustrate my point: rules can produce broken models. In my humble opinion, this point should be mentioned, although the formulas do not need to be changed, as, again in my humble opinion, these are simply cases of generated undesirable models (in reference to Table 3.1'). So, in my opinion, the formulas and the reasoning should not be changed at all, but the possibility of having “broken” results should be (briefly) discussed in the manuscript.

My second concern is that this chapter is too heavily focused on procedural techniques that produce an object or set of objects as a result. On the other hand, there are procedural techniques that do not produce objects as buildings, trees or animals, but that produce, for instance, clouds, rivers, weather, textures or animations (e.g., the game Spore, with its Siggraph paper). All these topics have a large body of literature in themselves, so let me clarify that I am NOT asking for a thorough review of the literature here, not at all! I am just suggesting to mention that there are alternative procedural techniques that may require their own definition of realism and that the results should be evaluated using specifically tailored measures, and discharge the burden of the references to some surveys or books such as Digital design of nature, the book by Oliver Deussen; The Algorithmic Beauty of Plants, by Aristid Lindenmayer and Przemysław Prusinkiewicz; or the book Texturing and Modeling: A Procedural Approach, by Ebert et al (this one is already cited, but the citation should be updated to the 3rd edition, 2002). Also, observe that the line that divides procedural modelling from simulations is quite diffuse, so let me insist that I am only suggesting to add a few sentences to comment about these issues, and the appropriate references as further reading. In my opinion, a deeper analysis of these issues is out of scope of this thesis.

Chapter 4 continues with the topic of realism by conducting a user study that investigates the realism of procedurally generated buildings. Specifically, the user study Mr. Jan Benes proposes, is designed to determine both the contribution of fine and coarse details to the perception of realism, as well as to illuminate the more qualitative factors that make a contribution to the overall subjective perception of realism. As a result, Mr Benes and co-authors have found that this last point, i.e., the perception of realism, when applied to procedural buildings, depends largely on different scales of detail in statistically significant ways. On the other hand, they have also identified other important aspects that also contribute to the perception, such as quality of photographs (lens, composition, experiment pre-processing, etc.), texture repetition, incorrect geometry, model structure, excessive regularity, and even window detail, among others. As enumerated in the manuscript and the published paper, the main contributions of this chapter are to provide an assessment of the importance of coarser scale structure on realism, to identify a list of features that, when taken into account, either increase or diminish the realism of generated buildings, and to develop a novel methodology for procedural models using cut-out images to compare generated objects with reference photographs. The approach is useful for both procedurally and non-procedurally

generated buildings. I think it is here where Mr. Benes excels in his mastery of these really complex concepts.

Chapter 5 introduces the final major contribution of the thesis, a model for growing procedural road networks in and close to cities. Starting from an empty map, the algorithm produces an initial road network, which is grown and urban nuclei are identified. The city is initialized as a set of several smaller settlements, but, as it grows and traffic increases, new major roads are generated, resulting in the formation of a conurbation. In this algorithm, in the end even new major roads may be absorbed. This model is based on the assumption that a city cannot be meaningfully simulated without taking its neighbourhood into account. The authors consider a simple traffic simulation, which considers this neighbourhood is then used to grow new major roads and to influence the locations of minor road growth. The algorithm also introduces waterways to help position the city nuclei on the map. The resulting cities are formed by allowing several smaller settlements to grow together and to form a rich road structure, much like in real world, and require only minimal per-city input, allowing for batch generation.

At this point I would like to suggest the authors briefly mention other approaches for traffic simulation, such as micro-, meso- and macro simulators. In my opinion, the implementation in this chapter belongs to the meso simulation set, but I think a brief mention of the alternatives, together with some bibliographic references, are in order. A good starting point, in my opinion, is the corresponding Wikipedia article (https://en.wikipedia.org/wiki/Traffic_simulation), and to mention the main software in the field are Sumo or Transims (among many others), where the latter was used for the simulation of the whole Switzerland. I think a brief mention (just a few sentences) of the alternatives are in order. This would also strengthen the future work section, which explicitly mentions intra-city simulations, so adding a couple of extra references there would round this paragraph up.

Chapters 6 revisits the major contributions of the thesis, putting them in perspective against the original goals, and what they achieved. Naturally, with such high expectations with potentially unforeseen badly-defined problems like realism in Computer Graphics and Procedural Modelling, it is very complex, almost impossible, claim everything has been solved. The proposed directions of improvements are quite exciting, and I am sure that starting Ph.D. students, at Prague or anywhere, would find lots of ideas to occupy them for years to come. This shows how well Mr. Benes has mastered his own wide area of research, and his clear vision for the future.

Conclusion

The thesis of Mr. Benes addresses challenging problems dealing with realism in Computer Graphics in general, and procedural modelling of buildings in particular, and proposes strong innovative contributions to analyse these problems. Also, Mr. Benes proposes an evolution-based approach for street network growth which proves to be historically-faithful and quite practical for urban modelling.

I find the final results already quite good given the incredible dangers hidden in such subjective, ill-defined concepts. In all major contributions, I am certain that the film and video game industries would be already highly interested in these solutions, demonstrating their application to real world problems. Moreover, I can see how urban planners, for instance, could greatly benefit from these expressions, formulae and evaluation techniques to assess their models, thus helping stakeholders to better control their models and predict their overall impact.

The strength of the contributions, their originality, and their complex results and thorough analyses, demonstrate very well that Mr. Benes has become a true expert in the difficulties due to the inaccurate, complex nature of the vague definitions of realism in Computer Graphics. The associated publications and certainly those to come, also support my evaluation of the quality of this research.

For all these reasons, it is with a *very strong support* that I *accept* this thesis for its defence for a Ph.D. degree from the Faculty of Mathematics and Physics from Charles University.