EG UK Theory and Practice of Computer Graphics (2010)
John Collomosse, Ian Grimstead (Editors)

Perception of Clones in Forest Rendering

Alan Purvis[†1] and Veronica Sundstedt[‡2]

[1]Trinity College Dublin
[2]Blekinge Institute of Technology

Abstract

The application of instanced clones represents a powerful technique for reducing the time and space requirements of the storage and visualization of large populations of similar objects. This paper presents the results of several perceptual experiments into the application of cloning to plant populations, within the context of a project to explore the use of resource-acquisition based techniques to model plant distributions. The perceptive effects of clone rotation on human subjects will be explored, with the goal of stratifying clone rotations and minimizing their detection. The perceptual effects of clone number, plant species heterogeneity and appearance will also be explored.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism;

1. Introduction

While the procedural modelling of plant structures based on L-systems and other mathematical models has been subject to exhaustive studies over the last four decades [Lin68], far fewer studies have been conducted into the realistic creation of flora based ecosystems and the distribution of plants and plant species within them. What studies have been performed largely pertain to the application of the same L-system fractal mathematics to the creation of population distributions [DHL*98]. While these models do account for some real world phenomena such as the process of self thinning and clumping of plants by species [Hop54], there exist many basic natural traits which cannot be modelled within their design. The existence of heterogeneous resources across an ecosystem and the preferences of different species of plant play a large role in where different species may prosper within a given terrain. The parent project of this study seeks to explore the modelling of flora ecosystems using a resource based model of generational plant growth, considering water and light distribution, soil depth and rule based plant preferences and requirements.

The plant populations generated within our framework may contain many plants, often of the order of millions. While the visualization of these populations is not a direct part of the project, simply storing the models for each of these plants takes up a tremendous amount of computer memory and warrants a study into what memory saving techniques might be applied. Beyond simple data compression, the resulting plant populations seemed a perfect case for the application of cloning and instancing. It is however desirable that whatever usage of clones is employed to reduce the memory footprint of models have as small an impact on the human perception of the compressed model as possible.

1.1. Population Distribution Generation

Our framework for generating plant distributions takes as inputs: a height map for the terrain to be populated, a programmatic sky model describing a hemisphere of light values, information about rainfall and a series of plant species definitions, which include descriptions of plant lifespan, size, growth, aggressiveness, resource requirements and seeding mechanics. From these inputs, it can derive further information such as light and water distribution across the terrain (Figures 3, 4), and then simulate the competition of plants across the resulting terrain using a number of simple rules pertaining to self thinning and resource competition, among others (Figures 1, 2). In order to start the competition, a number of seeds are randomly "dropped" onto the terrain as if

† e-mail: purvisa@tcd.ie
‡ e-mail: veronica.sundstedt@bth.se

DOI: 10.2312/LocalChapterEvents/TPCG/TPCG10/107-114

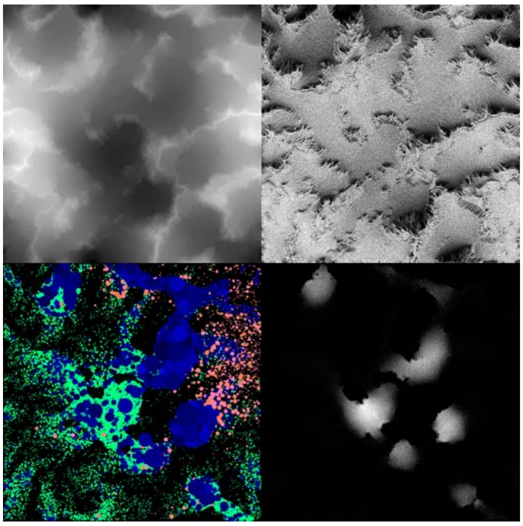

Figure 1: *Top left, clockwise: Height map, Light incident over the period of a day, Water distribution, Resulting plant distribution (coloured by species: red, green, yellow).*

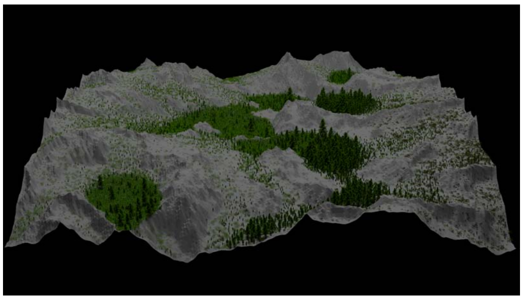

Figure 2: *3D visualisation of plant distribution shown in Figure 1.*

seeded by a bird or the wind. These seeds will then, if the environment is conducive, grow into an active ecosystem after several plant lifecycles of simulation.

2. Background and Related Work

Much prior work exists in the field of forest visualization, the most commercially visible example being SpeedTree [Int10], sold by Interactive Data Visualization and used in many high profile commercial videogames, such as Gears of War 2 [Epi08b], Unreal Tournament 3 [Epi08a] and The Elder Scrolls IV: Oblivion [Bet07], among many others. There also exists a plug in to integrate SpeedTree with the popular 3D modelling software, 3D Studio Max [Aut10] which may be used to create offline rendered images and video sequences. In academic research, work with a more ecosystem oriented basis may be seen in [DHL*98], where self thinning curves, and Hopkin's Index as a measure for plant clustering are addressed. The issues presented in rendering trees and

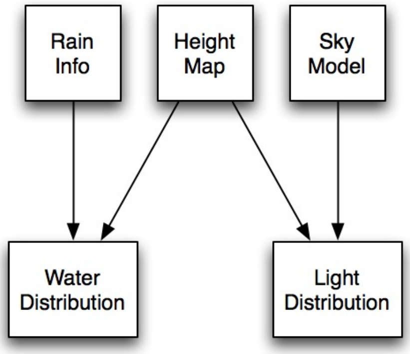

Figure 3: *Inputs and derived resources.*

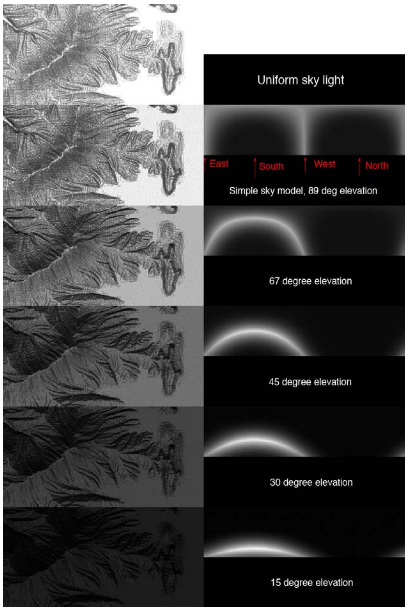

Figure 4: *Effect of varying elevation angle in sun model on a simple terrain. Left: Incident light, Right: Polar plot of sky model. North towards the top of the light maps.*

other plants with realistic shading and texturing are well addressed in [Bou08], where specific attention is given to the approximation of the effects of global illumination within the constraints of real time rendering. Several studies have been performed into the perception of clones in human models in crowds [MLD*08], and [MLH*09] and various aggregate groups [RBF08], however, the study of cloning on plant populations however has seen less attention.

3. Tree Generation Framework

For the purposes of this paper, it was decided that adoption of an L-system based tree-generation scheme would be most suitable. This decision was made largely on the prevalence of such systems, both in academic and commercial work, and on the amount of research and discussion which has been based around them, which was seen to support their position as well documented and robust method. Further, being fractal, an L-system based approach would afford us the ability to create a practically infinite supply of trees very easily for our experiments. Rather than re-implement L-system based trees from scratch, Arbaro [Ope10a] was adopted as a suitably advanced and cross platform generator, with a transparent connection to academic L-system research, as it implements directly the algorithms documented in [WP95].

3.1. Tree Models

Having selected Arbaro as our source of L-system based trees, suitable L-system descriptions were required for input before we could generate the library of models needed for our experiments. Arbaro comes bundled with a modest selection of botanically inspired description files to be used for demonstration of the software. From these, it was deemed that the "Quaking Aspen" (Species 1, Figure 5), "Black Tupelo" (Species 2, Figure 6) and "Eastern Cottonwood" (Species 3, Figure 7) reasonably represented the plant morphology of their real-world counterparts, and that they reasonably spanned the rough form of large forest trees. While generating the 54 models to be used in our experiments, it was realized however that some of the default values for random variation in these definition files were very low, often taking values prohibiting or dramatically limiting random variation between trees. This resulted in the generated trees taking very similar structural forms, especially in the lower levels of the L-system. In order to compensate for this deficiency, a value inside the definition files representing random variation in angles between subsequent branches was increased to angles of 15-45 degrees. With this concession to the original definition files, the variability of the trees generated by the descriptions was dramatically increased, and it was felt that the issue of two unique trees bearing similarities to each other was not significantly higher than the risk of the same event occurring naturally.

Figure 5: *Quaking Aspen tree model.*

Figure 6: *Black Tupelo tree model.*

Figure 7: *Eastern Cottonwood tree model.*

4. Perceptual Validation

In order to facilitate the compression of our generated flora ecosystems, it was required that we further define what features it would be important to measure in our experiments. Two experiments were designed in order to answer some questions of basic performance. Experiment one was designed to pose the questions:

1. Can participants identify two clones?
2. Does rotation do anything to mask those clones, and if so, what is the smallest angle of significant confusion?

It was decided that for experimental purity and to minimize the danger of confounding variables in the tests, the first experiment would present the tree models for comparison in individually rendered images which were tiled on screen. This would give one common and equal viewing perspective on each tree, as well as remove the potential pollutant of trees occluding each other as they might if compared in a single common scene.

Experiment two was designed as a follow on experiment from the first, and to investigate the further questions:

1. Does increasing clone number have a negative effect on identification of clones?
2. How effective is stratifying the rotation of clones by the "smallest angle of confusion" recorded in experiment one as a means of concealing clones?
3. Does species heterogeneity have an effect on clone identification rates?

As the questions posed by the second experiment relate more to the interrelations of trees, and the perception of them as a group, it was decided unlike the first experiment to present the models in single rendering with a common camera looking at an arrangement of trees, as they might be viewed in the real world.

4.1. Stimuli

Once the models had been created, it was necessary to make suitable renderings of them. For this task, the open source global illumination renderer SunFlow [Ope10b] was selected. Sunflow offered the benefits of native instancing support and a practically fast global illumination engine. It was decided that, although at the time of writing, global illumination approximations are only beginning to make it into practical real-time rendering [ML09] [WWZ*09], as technology progresses these techniques and their successors only become more widespread. Thusly it was deemed that including the effects of global illumination in this study would be more relevant to future directions the field might take.

4.1.1. Experiment One

For experiment one, 108 renderings (800x800 pixels) were made of 54 different tree models (18 models per species). Each model was rendered in an un-rotated state, and then rendered a second time under a rotation performed around its trunk, to give a total of 108 images. These images formed rotated/un-rotated pairs of each model, which the participants were later asked to pair up from a set of decoys. The rotations performed were picked from a set of nine angles (Figure 8). Accordingly, within each species, there were two unique models rotated to each of the nine angles. This system of pseudo-replication of angle was decided on as it was thought that presenting each angle with only one model could be potentially unfair, if the model assigned to an angle within a species was particularly distinctive in some way.

For the presentation of experiment one to users, a program was written which would tile images from the set of 108 generated earlier into cross shaped presentations (Figure 9). The program ran though each of the 54 rotated/un-rotated pairs in a random order for each participant, placing one image in each pair in the center of the 3x3 presentation, and the other image in the pair randomly in one of the surrounding tiles. The remaining slots were then filled with random decoy images from the same species, taking care not to accidentally introduce a second pairing to the experiment (If one

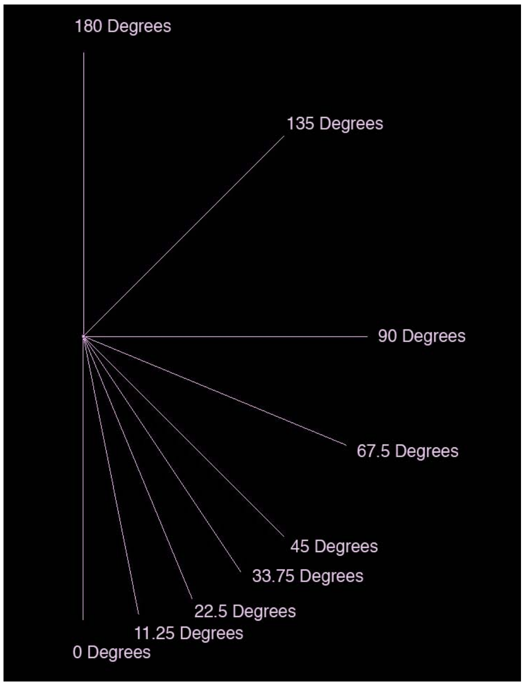

Figure 8: *Rotational permutations used for rotated pairs in experiment one.*

of a rotated/un-rotated pair was used as a decoy, the other in that pair was ruled out as a choice for future decoys within that presentation). Users were then asked to click on the surrounding image they believed to be the same model as the one in the center of the cross. The program recorded times and correct/incorrect guesses as the user progressed through the experiment.

4.1.2. Experiment Two

Experiment two called for the presentation of a larger number of trees in a single image, it was deemed that a 4 wide x 3 deep grid pattern with a downwards looking camera as demonstrated in (Figure 10) offered the best trade-off between tree number and visibility. A deeper 4x4 grid did not offer enough clarity to the back rows, and it was felt that a more shallow arrangement did not offer enough of a sense of the perception of the trees in a cluster.

The permutations decided for experiment two were to create images using each of the three tree species, with the number of clones varying from 2 to 4, with and without stratified rotation, and presented homogeneously, or with a second plant species used as a heterogeneous decoy (the participants were not told which species contained the clone). This made for a total of 3x3x2x2 = 36 presentations per participant.

The images for experiment two could not be generated on the fly for each participant as they could in experiment one,

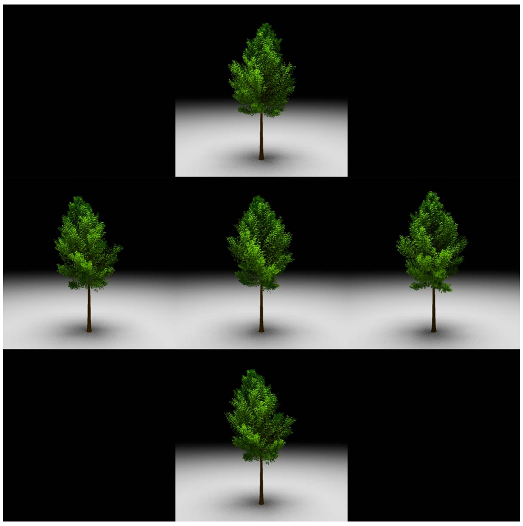

Figure 9: *Example stimuli image from experiment one (clone pair in positions 4 and 8– reading left-to-right, top-to-bottom, 15 degree rotation).*

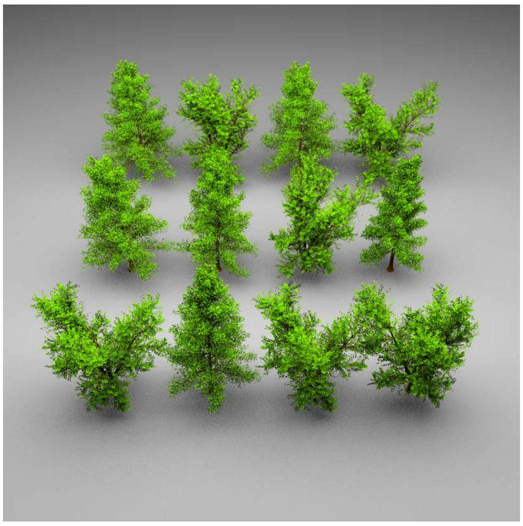

Figure 10: *Example stimuli image from experiment two, heterogeneous population, no rotation, three clones (positions 4, 9, 11– reading left-to-right, top-to-bottom).*

as it had been decided that the trees should be presented in a single rendered image. For this reason, a larger number of images was required to ensure fairness in the event of particularly distinct presentation being used to represent one of the experimental permutations. To guard against this danger, four versions of every permutation were created and one of the four variations selected at random for each participant. This required the creation of a total of 144 images.

As with experiment one, a program was written to present the permutations to the participants, enumerating all 36 permutations in a random order. Users were this time asked to click two trees that they believed to be clones. In the presence of multiple clones, they were still only asked to identify one pair. The time and correctness of their pairings were recorded.

4.2. Setup

The experiments were run as much as possible in a sound and light controlled environment and the distance between participants and the 24 inch LCD display used was loosely controlled by positioning the seat they were asked to sit on. They were offered an explanation of the task they were to perform beforehand, aided by diagrams, and asked if they had any questions or wanted anything clarified or repeated before the experiment began. Care was taken to address the level of language that a few foreign participants possessed.

At the start of the experiments and between each image, the experimental program presented users with a black screen with a white cross in the center. In order to minimize the effects of "performance pressure", participants were informed that they were not timed on this screen and that if they desired to take a break to regain composure during the experiment, that they could take any rest needed upon it. Participants were instructed on how to dismiss this screen by looking at the cross to center their vision, and clicking the mouse in their own time. All presentations to participants were made without a time limit, which the participants were made aware of, they were however asked to limit their analysis of each image to a sensible period and a suggestion of around 15-30 seconds was made.

4.3. Participants

Participants were volunteers from the university. The only discriminatory criteria applied in their selection was their own willingness to take part in our experiments. A reasonable balance of male and female participants took part. The two experiments were run with 16 and 18 participants, respectively.

5. Results

Interestingly, the time to answer recorded by the experimental program did not show much variation from question to question. The length of time spent on a given question appeared to be determined almost totally by the participant and not bear much relation to the difficulty of the question being asked. As a result of this, we decided to base our statistics off the correct/incorrect values that were simultaneously collected. One possible explanation for this is participants second guessing themselves on the easier questions in the absence of a time limit.

Experiment one shows that for a rotation of 0 degrees (identical image pair), users have a near perfect identification rate when presented with three similar decoy trees (95%, Figure 11). This detection rate remains relatively unhampered when the angle of rotation is increased to 15 degrees. Past 15 degrees however, the detection rate rapidly falls off as the rotation increases. At 45 degrees, the detection rate reaches a near-plateau around 40-45%. Considering the base-line for random guesses in this experiment is 25%, this represents a significant falloff. Given 45 degrees as the smallest angle of significant confusion, this would suggest that there are eight significantly different views of a tree. This is an interesting figure for the stratification of clone rotations, as it would suggest that if you could solve the map-colouring problem, each tree in a forest could have up to seven direct neighbours before two adjacent trees were forced into the same configuration by the pigeonhole principle.

Experiment two shows little difference between the identification rates for different plant species across all other permutations (Figure 12). Species three, the Eastern Cottonwood has a slightly higher identification rate, probably resulting from the more irregular form is possesses in comparison to the other two species, but with p = 0.5, there is no significance under a chi-squared test.

Varying clone number in experiment two has a more interesting effect on clone detection rates. Without stratified rotation, detection rate increases smoothly with clone number. With one clone (a pair), detection rates are around 37%, rising to 72.2% with three clones (a quad of identical models)(Figure 13). The statistical significance of this effect under a chi-squared test is dramatic (p = 0.00000136). In the presence of stratified rotation, however, this effect becomes more complicated, with identification rates going down as the clone number goes from 2 to 3 (Figure 13). The only possible explanation the authors can offer for this result is that participants may have been confused by the ambiguity offered by the presence of multiple correct answers to a difficult question, and this may have had some effect on their confidence on that question. Whatever the explanation is, the statistical significance of this effect is still large (p = 0.01401).

The effect of Homogeneity/Heterogeneity on detection rates within experiment two is less clear (Figure 14). Without tree rotation, there appears to be no difference in the detection rates between Homogeneous and Heterogeneous stimuli. In the presence of stratified tree rotation, the detection rates with Heterogeneous populations appears larger. this effect is not statistically significant however (p = 0.1857).

6. Conclusions and Future Work

The experiments performed, while far from encompassing all parameters and questions of interest on the application of instanced clones to trees, would suggest that there is a lot of leeway for clone repetition in forest representation. Rotation appears to be a particularly effective technique for hiding clones, especially when suitably stratified. The presence of Heterogeneous plant populations does not appear to present any obstacle to the use of instancing, at least within the balanced two-species Heterogeneity tested in our experiments.

The discovery of 45 degrees as an optimal angle for stratifying rotation would suggest that many forests could be stored with very few unique models repeated many times over. It is the authors' expectation however that as the number of trees increases and the area that each tree occupies on the screen diminishes, viewers will begin to look at a forest as a texture rather than a sequence of distinct objects, and that repeated clones may begin to appear as a recognizable pattern. To explore this possibility, it would be interesting

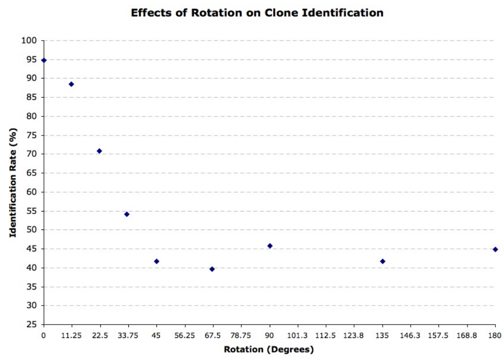

Figure 11: *Identification rate graphed against clone-pair rotation. Random guess base-line 25%.*

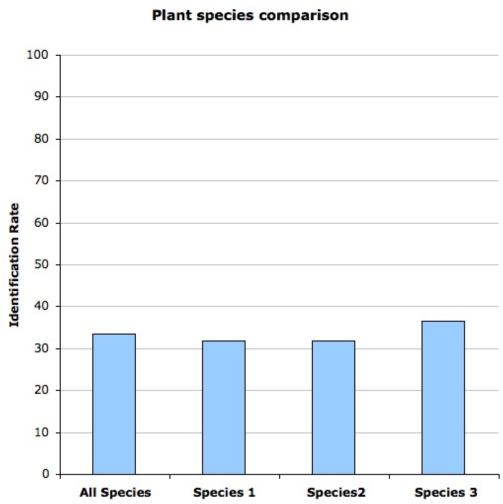

Figure 12: *Plant species comparison.*

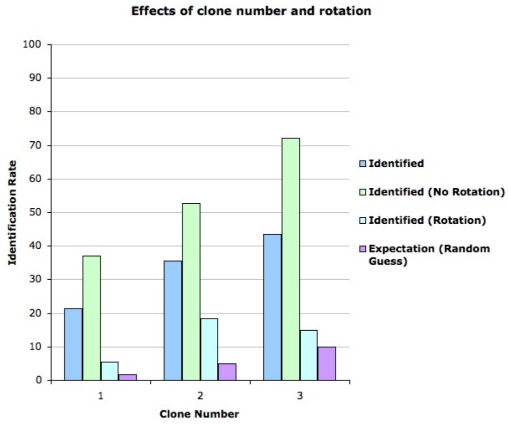

Figure 13: *Clone number comparison.*

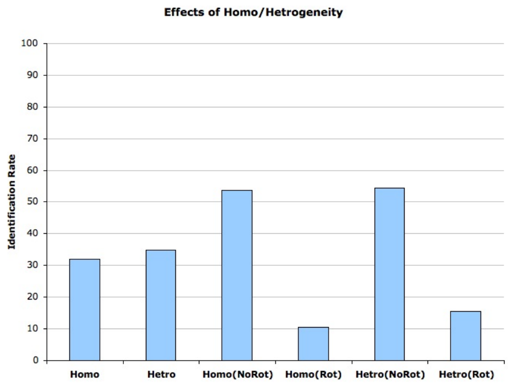

Figure 14: *Homo/Hetrogenious population comparison.*

to perform a third experiment similar to experiment two, using more trees, viewed from a greater distance and with an appropriately higher number of clones.

Another obvious future avenue would be performing further experiments to ascertain exactly why the identification rate in experiment two drops with a clone number of three and stratified rotation enabled. Performing the experiment again with a larger set of clones and seeing what happens to identification rates with more than three clones may help shed some light on this. Alternatively, exploration of human psychology and behaviour may yield new avenues if no rational explanation can be found.

Acknowledgements

We would like to thank everyone that participated in the experiments and the Author(s) of Arbaro, SunFlow and Blender [Ope10c], without which the development of this research would not have been as smooth.

References

[Aut10] AUTODESK MEDIA AND ENTERTAINMENT: Autodesk 3ds max. *www.autodesk.com* ((URL Accessed 2010)). 2

[Bet07] BETHESDA GAME STUDIOS: The elder scrolls iv: Oblivion. *Bethesda Softworks/2K Games* (2006/2007). 2

[Bou08] BOULANGER K.: Real-time realistic rendering of nature scenes with dynamic lighting. *Ph.D Thesis, University of Central Florida.* (2008). 2

[DHL*98] DEUSSEN O., HANRAHAN P., LINTERMANN B., MĚCH R., PHARR M., PRUSINKIEWICZ P.: Realistic modeling and rendering of plant ecosystems. *SIGGRAPH '98: Proceedings of the 25th annual conference on Computer graphics and interactive techniques* (1998). 1, 2

[Epi08a] EPIC GAMES: Unreal tournament 3. *Midway Games* (2007/2008). 2

[Epi08b] EPIC GAMES: Gears of war 2. *Microsoft Game Studios* (2008). 2

[Hop54] HOPKINS. B.: A new method for determining the type of distribution of plant individuals. *Annals of Botany XVIII* (1954), 213–226. 1

[Int10] INTERACTIVE DATA VISUALIZATION: Speedtree. *http://www.speedtree.com/* ((URL Accessed 2010)). 2

[Lin68] LINDENMAYER. A.: Mathematical models for cellular interactions in development, i ii. *Journal of Theoretical Biology* (1968), 280–315. 1

[ML09] MCGUIRE M., LUEBKE D.: Hardware-accelerated global illumination by image space photon mapping. *HPG '09: Proceedings of the Conference on High Performance Graphics 2009* (2009). 4

[MLD*08] MCDONNELL R., LARKIN M., DOBBYN S., COLLINS S., O'SULLIVAN C.: Clone attack! perception of crowd variety. *ACM Transactions on Graphics (SIGGRAPH 2008), 27, (3), 26:1 - 26:8, 2008* (2008). 2

[MLH*09] MCDONNELL R., LARKIN M., HERNANDEZ B., RUDOMIN I., O'SULLIVAN C.: Eye-catching crowds: Saliency based selective variation. *ACM Transactions on Graphics (SIGGRAPH 2009), 28, (3), 55:1 - 55:10, 2009* (2009). 2

[Ope10a] OPEN SOURCE: Arbaro. *http://arbaro.sourceforge.net/* ((URL Accessed 2010)). 3

[Ope10b] OPEN SOURCE: Sunflow. *http://sunflow.sourceforge.net/* ((URL Accessed 2010)). 4

[Ope10c] OPEN SOURCE / BLENDER FOUNDATION: Blender. *www.blender.org* ((URL Accessed 2010)). 7

[RBF08] RAMANARAYANAN G., BALA K., FERWERDA J.: Perception of complex aggregates. *Proceedings of SIGGRAPH 2008 (SIGGRAPH 2008, ACM Transactions on Graphics, Volume 27, Number 3)* (2008). 2

[WP95] WEBER J., PENN J.: Creation and rendering of realistic trees. *SIGGRAPH '95: Proceedings of the 22nd annual conference on Computer graphics and interactive techniques* (1995). 3

[WWZ*09] WANG R., WANG R., ZHOU K., PAN M., BAO H.: An efficient gpu-based approach for interactive global illumination. *SIGGRAPH '09: ACM SIGGRAPH 2009 papers* (2009). 4