

Haptic Granular Synthesis: Targeting, Visualisation and Texturing

Andrew Crossan,¹ John Williamson,² Roderick Murray-Smith^{1,2}

¹ Hamilton Institute, National University of Ireland, Maynooth, Co. Kildare Ireland.

² Department of Computing Science, University of Glasgow, Glasgow, Scotland.
andrew.crossan@may.ie, {jhw, rod}@dcs.gla.ac.uk

Abstract

This paper introduces the idea of haptic rendering using granular synthesis – an established technique for synthesising audio. It describes the technique along with potential application areas, and initial results from an implementation on a PHANToM force feedback device. Three main applications are considered. Firstly, rendering of probabilistic vector fields for presenting ambiguity and context information to the user. Secondly, the possibility of producing textured virtual objects using granular synthesis is discussed. Thirdly, we use the approach to display scatterplot data on haptic devices.

1. Introduction

There are many situations where non-visual presentation of data is desirable or even essential. This is particularly the case when designing systems for visually impaired users. This paper examines probabilistic methods of texturing to present information to a user through a force-feedback device. The methods outlined in this paper, and tested on a PHANToM force-feedback device [7] (from SensAble Technologies) allow the design of textures to be linked to audio files accompanied by weighting functions, which are used to blend these in space and time. We describe a haptic design environment where the system (or the user) is able to place different distribution functions around regions in the state-space, and associate these with audio files, which then render different textures. The distribution functions allow localization of effects, blends of multiple effects and could be smooth, differentiable functions, or discrete functions with compact support in the space. The distribution function could also be a corridor around a trajectory in the space. Designers could use their auditory experience to browse potential sound files, or generate or mix their own. This could be so easy to do that end-users could further customize their interfaces, or provide sample sounds that provided textures that were meaningful to themselves.

Granular synthesis is an inherently probabilistic rendering method, and can be used to display probability densities composed of a number of component densities – known as mixture models in statistics [10].

Forces synthesised in granular fashion on a force-feedback haptic device can therefore be used to attract the user to different regions of space, or to change the closed loop dynamics to make it easier to follow certain trajectories. These forces, however, as well as attracting the user to a particular point or trajectory, would carry a recognizable signal in the form of the texture or disturbance associated with its audio file. This means the user can ‘browse’ the space for the target texture. It also means that unlike alternative approaches, such as gravitational effects, the user does not merely perceive the vector field associated with the *average* of the multiple local flow systems, but perceives a *mixture* of components. With granular synthesis, the long-term time-average force at a given point in the state-space will converge towards the average of the raw flows (assuming infinite grains, and balanced power content in each audio file), but the local behaviour will contain information about multiple nearby attractors, with their signature haptic effects.

The approach is therefore ideal for the design of probabilistic model-based displays, where multiple uncertain hypotheses must be displayed, and the user is made aware of the uncertainty of the different options, and the strength of evidence for each. The likelihood of these hypotheses will typically be conditioned on the system state history and user inputs. An example might be a target acquisition task such as button-pressing, where the hypothesis H_i would be ‘user wishes to press button i ’, and would be conditioned on the position represented by the current cursor state variables x, y . Alternatively we might have a gesture recognition task, where the hypothesis might be ‘User intends to open new file’, where the user’s intention is inferred from the time-series of input actions. In this case the likelihood of the hypothesis would vary over time, as the trajectory was incrementally processed by a pattern recognition system. In both cases the density functions associated with a given sound file could represent the conditional probability of

the various hypotheses, providing the user with time and state-varying textures and attractive forces. The model and inference mechanisms map directly to display rendering. The probabilities can obviously also be context-dependent, and modelling tools such as Bayesian networks can be used to encode the appropriate models for the display.

With granular synthesis being an accepted technique for synthesising audio, development of a similar haptic rendering technique naturally lends itself to multimodal presentation of information. Using the scatterplot example described in Section 4.2, a user could simultaneously explore the data using haptic and auditory cues based on the same models. However, there are open questions as to how the audio and haptic synthesis techniques will combine. The different effective filters inherent to the different displays, and the frequencies ranges perceived by both senses will present different frequency ranges of the common audio file.

2. Introduction to Granular Synthesis

Granular synthesis is an established method for the creation of interesting and dynamic audio. The technique was originally developed in its current form by Xenakis (see [13]) in the late 1950's; see [9] for a modern overview of the topic.




Figure 1. A description of the granular synthesis process.

The fundamental idea is that a meaningful macroscopic waveform structure can be created from a large number of small waveforms. This process is

somewhat analogous to particle systems in computer graphics. All of the low-level elements (“grains”) are controlled by a small number of global parameters, usually statistical distributions over the individual element parameters. Such a formulation means that statistical models can easily be mapped to the synthesis algorithm. For example, the output probabilities from a pattern recognition process can be directly mapped to the probability of drawing a grain from a source corresponding to each of the potential goals.

In audio granular synthesis, the grains are short sections of a source waveform, e.g. a pre-recorded section of audio, although it can also be the output of other synthesis algorithms. Each grain is shaped (amplitude enveloped) to eliminate discontinuities (perceivable as clicks) and all of the active grains are mixed down into a single output stream. The particular qualities of the sound are manipulated by changing the distribution parameters. Most importantly, the distribution over waveforms from which each element is drawn can be modified to interpolate between audio textures. Normally several hundred to several thousand elements would be simultaneously present. Figure 1 shows the overall process.

The technique can be applied to any situation in which complex waveforms with consistent but rich and continuously-varying properties are required. The output of the synthesis could be transformed into time-varying force vectors in a haptic setting, to produce a rich, complex sensation, which can easily be controlled by a small number of intuitive parameters. This is particularly true when the system with which the user is interacting is based upon some statistical model, or can easily be transformed into one, e.g. in pattern recognition of context. In this case, the model can be directly translated into force output parameters, given some reasonable choice of source waveforms.

3. Force Vector Fields

Barrett *et al.* [1] demonstrate how changing the dynamics of interaction can be used to aid targeting. Using the concept of negative inertia – using additive derivative control – they demonstrate how perceived sluggishness in an isometric interface can be negated. A force vector field throughout a space can similarly be used to affect the dynamics of the system. It can be used to provide context information about the space, and can be dependent on both spatial and time-varying properties. One example of a spatial vector field would be a gravitational field. If a user were to navigate through a space with a number of fixed targets that were exerting a gravitational force on the user's cursor, the gravitational force felt would purely be as a result of the position of the cursor in the space with respect to the targets. The user

will then feel a constant gravitational pull that will provide information about the direction of and distance to nearby targets. It will provide a pull on the user that will assist in targeting tasks. As the user gets closer to a target, and it therefore becomes more certain which target is of interest to the user, it will provide a stronger pull towards the target.

However, in the event of two approximately equidistant targets that are in opposite directions from the cursor, the forces from these targets may to a large extent cancel each other out. If the user has a decision to make as to which target to aim for, the spatially distributed vector field may not always provide the information required. The user may be aware, because of the low force exerted, that there is an area of high uncertainty, but will not be given the context information required to find a particular target. In this instance, vector fields that vary in time as well as space may provide a better solution – the approach we develop is one based on haptic granular synthesis.

3.1 Granular Synthesised Methods

As in audio granular synthesis, at each time step, every target in the space will have an amplitude value associated with it. If no grain is currently being played from a source, this value will be zero. Otherwise, this value will be the amplitude of the synthesised waveform. For example, for an 8 bit sound source there are 256 possible amplitude levels. The time-varying properties of granular synthesis can be displayed to the user through a force feedback device by directly varying the force from the actuators of the device using these variations in waveform amplitude. While the frequencies in audio rendering are considerably higher than would be expected in a haptic system, the rate of playback or chosen sampling rate can be adjusted accordingly.

In a system such as the gravity example above, grains emitted by a target can be set to produce a force of magnitude determined by the current waveform amplitude in the direction of the target. For a single grain, the user would feel a short tug towards the target. With multiple grains from different sources, vector summation can be used to calculate a resultant force. In some instance, these grains will cancel each other out. However, as the grains are generated at random intervals, it is likely that the user will feel a series of short tugs towards each of the different sources (shown in Figure 2). A probability distribution function can be used to control the number of grains generated from each source. For example, if we assume that a user is more interested in a close target than a far away target, we can propose a probability density $P(\text{Interest}|\text{Distance})$ such that a higher density of grains are drawn from the closer target. The user will feel stronger tugs towards the close target, but may also be

able to perceive and locate a more distant target based on the sparse but recognisable grains generated from that target.

In an area of high uncertainty such as the saddle point (where the forces from different targets largely cancel each other out) described in the above gravity field example, the grains will be drawn in opposite directions and will cancel each other out to some extent. The user may not feel much force in any one direction, but will be able to appreciate short bursts of force in the direction of each of the targets. If they then explore the region close to this point, they can detect the change in the texture and disturbances locally, and move towards their desired goal. Texture can therefore be used as a type of quickening [5]. By structuring the texture relative to hypothesized goals, we can display to the user a prediction of the consequences of maintaining the current trajectory. We expect that most applications will use a weighted combination of textures and disturbances generated by granular synthesis.




Figure 2(a). Force vectors varying in time as the user holds the cursor between two sources. The “tugging” force vectors are visible.




Figure 2(b). Magnitude of force vectors from two sources as the user moves left to right between their associated densities (the bell curves).

4. Texturing

In real world haptic object exploration, an object's texture is often important for object identification. Texture display in a virtual environment is therefore desirable to provide the user with more information about the object and increase the fidelity of the simulation. Haptic texture rendering, which has been examined extensively in the literature, can be separated into 3 main categories: abstract models, record and playback models, and stochastic models. For example, McGee [8] examines perception of roughness using sinusoidal force feedback textures – an abstract texture model. Jansson [6] examined the perception of real world and virtual sand paper textures. The virtual models used in this study were built from recorded properties of the real world sandpapers. Fitz and Barner [3] describe a stochastic method for generating textures that can be combined with a granular synthesis approach, where the time-series would be generated by the chosen probabilistic model instead of audio files.

However, generating textures can still prove challenging for designers of virtual environments. One goal of this research is to provide a method of virtual texture generation that is easy to include in a simulation for a non-technical designer.

The changing texture of a surface is one method of conveying information to a user in a virtual environment. McGee [8] demonstrated that – despite the lack of cutaneous cues – a force feedback device can provide discriminable textures that can be ordered by their perceived roughness. Wall [15] describes the possibility of encoding information in the different haptic properties of an object – including friction, stiffness and texture – for displaying graph data to visually impaired users.

4.1 Granular Synthesis for Surface Texturing

In a granular synthesis surface texture system, there are two different methods of rendering the texture by using the waveform amplitude to vary either the height of the surface or lateral forces applied when moving across the surface. Using amplitude to vary the height of the surface seems to be the more intuitive, as with a force profile representation, the texture is invariant of related object properties such as friction and damping. Figure 3 shows the amplitudes generated when moving across a flat surface textured using granular synthesis. For one source it can be seen that the texture consists of a number of distinct bumps. The density of the bumps can be changed by increasing the number of grains generated per second. For three sources, the texture is generated from a mixture of grains from different sources, and is therefore more variable. For any point in time, the playback rate from the grain sources can be altered due the velocity of

the cursor such that a higher rate of grains is felt by the user with a higher cursor velocity.

Due to the stochastic nature of the granular synthesis process, the same area of surface will not have exactly the same texture at different points in time although the perceived qualities will be similar. The implementation described in this paper displays the haptic granular synthesis on a PHANTOM force feedback device [7]. We use the PHANTOM as it provides a high-resolution force feedback and allows intuitive interactions with a three-dimensional space. It allows a user to interact with a virtual environment through a single point of contact. The user can move freely in 6 degrees of freedom (X, Y, Z, roll, pitch, and yaw). The device can also provide 3 degrees of high-resolution force feedback to resist or assist motion in the X, Y and Z dimensions.

In the implemented system, users can select a variety of sound files and drag and drop them over a flat surface using the PHANTOM device. The width of effect of each of the files can be altered along with the number of grains generated per second.




Figure 3. Change in the height of a texture generated by granular synthesis. Top: granules from individual waveforms. Bottom: summation of the granules generated from the 3 sources.

Different feels can be generated by simply clicking and dragging the grain sources to different areas of the surface. To generate a surface that is approximately uniformly textured, a single source can be used and a uniform density can be associated with every position on the surface. This can be achieved by the normalisation of the probability density function, which ensures that the source will have the same contribution everywhere on the

surface. It is important to determine whether conditional probabilities of hypotheses are being used which sum to one at every state, or whether they integrate to one over the whole state-space.

4.2 Display of Scatterplots and Probability Density Functions

This mechanism provides an interesting approach to the display of scatterplots – an area of great interest to blind users. Existing work in haptic visualization of graphs and mathematical functions includes [14] and [16].



Figure 4(a). A cursor trace from a user exploring the granular synthesised scatterplot using a PHANToM.




Figure 4(b). A contour plot of the Gaussian distributions around the data points of the scatter plot shown in Figure 4(a).

For display of point data where all points belong to the same class, we have a common audio file for each point, and we associate the same form of density function, e.g. a Gaussian density, with each data point. This is analogous to nonparametric modelling in statistics based on Kernel density estimation, also known as Parzen

smoothing [4]. The average density then becomes a smoothed version of the original data and approximates the probability density function. In cases with multiple classes of data point, each class uses a different audio file. The changing signature files would allow users to identify clusters of data points from a common class and find outliers.

Other features of data points could potentially be represented by transformations of the associated audio files – e.g. if the data could be enhanced by a regression line, or other parametric models, such as mixture models [10], the position the granules that are taken from the audio file could be linked to the arc length of the point on the function normal to each data point.



Figure 5(a). A cursor trace from a user exploring the granular synthesised scatterplot using a PHANToM.




Figure 5(b). A contour plot of the Gaussian distributions around the data points of the scatter plot shown in Figure 5(a).

Figures 4 and 5 display two scatterplots (with 3 classes of data) built using the system described. Figure 4 shows a scatterplot with few data points where there is little overlap between the distributions. Figure 5 demonstrates a more realistic situation where more data points are included. There is more overlap between the data points and data classes. One key feature of this

