Morphing Sound Attractors

Axel Röbel
Institute for Communication Sciences, Sekr. EN-8, Technical University of Berlin
Einsteinufer 17, 10587 Berlin, Germany
email: roebel@kgw.tu-berlin.de

ABSTRACT

Dynamic modeling of sound is a new approach to model and synthesize natural sound signals and is based on a recent method of modeling dynamical systems with neural networks. The following investigation addresses the problem of controlling the characteristics of the sound signals that are obtained from a dynamical sound model. Here, we propose an algorithm for morphing between the dynamics of different sounds. The dynamics of one dimensional attractors can be used to model a large number of sound signals. In the following we will give an example of a dynamical model of a piano sound that is based on one dimensional attractors. After having motivated the restriction to one dimensional attractor topologies, we theoretically investigate into a number of properties of the proposed morphing scheme and experimentally verify our conclusions. The topics that will be discussed in the light of the special application are the selection of reconstruction parameters, necessary constraints on the network parameters and signal characteristics that can prevent a successful application of the morphing method.

Keywords: Dynamic Modeling of Sound, State Space Reconstruction, Attractor Morphing, Radial Basis Function Network.

1. INTRODUCTION

Recently it has been shown that neural networks may be used to establish a so called *dynamical model* of chaotic system dynamics given a time series of the system, only [1, 8, 9, 2]. The method is based upon a reconstruction of the systems original state space and an embedding of the systems attractor [16, 15]. Using the state space reconstruction a neural network is trained as a predictor of the system dynamics, which when used iteratively can be regarded as a model of the system dynamics at the respective attractor. Stability of the model, and even more the equivalence of the models and systems attractor are difficult to guarantee, however, it is widely accepted that by means of recurrent training algorithms proper dynamical models can be achieved [8, 2].

Application of dynamic modeling is not constrained to chaotic dynamics, and it has been shown that dynamical models of natural sound signals can be used to resynthesize sound signals with high quality [11]. Compared to standard sound synthesis methods, however, the possibilities to control the synthesized sound are somewhat limited. Due to the complicated relations between the network parameters and the attractor of the model it appears to be impossible to find any sensible algorithm that allows a user to change the model characteristics by directly altering the model parameters. Therefore, we have investigated another approach to achieve signal modification, which we call attractor morphing.

Our approach to attractor morphing is based upon a simple homotopic mixing algorithm [4] that interpolates between two progenitor sound models to obtain a new model with interpolating dynamics. The progenitor models are trained independently, and are combined linearly to achieve intermediate dynamics. It is generally accepted that harmonic sounds of musical instruments are mainly characterized by amplitude and frequency of their partials. A volume change of a natural instrument, for example, is generally expected to change mainly the amplitudes of the partials with only very minor changes in their frequency. In contrast a pitch change moves all partials synchronously such that the sound remains harmonic while minor additionally amplitude changes may occur. Consequently, most sound morphing algorithms are based upon frequency domain analysis [17]. Due to the fact that frequency and amplitude have to be controlled independently to achieve convienient control of the sound characteristics we started the investigation of the attractor morphing algorithm with the question whether the morph between two sounds which differ only with respect to pitch or amplitude can be morphed without altering any of the other fixed signal parameters. Because in this initial stage of investigation we only morph between two progenitor sounds there exists a single morphing parameter, only. In the long run, however, the method is intended to interpolate between a higher number of attractors of the same musical instrument such that a dynamical model not only of a specific sound, but, for a part of the whole instrument is obtained. For such an application the space of morphing input parameters has to be organized in such a manner that similar control inputs reflects similar dynamics in the reconstructed state space. A possible method to solve this problem in an automatic fashion is the hidden control neural network approach [3].

The results of our first experimental investigation of the proposed morphing algorithm have been promising

[12]. Morphing synthetic sounds with one half tone of pitch difference results in no perceivable alteration of amplitudes and morphing amplitude differences does not affect pitch. Small changes of the phase of the partials have been morphed independently of pitch and amplitude. Even the morph between two real world saxophone signals with a pitch difference of one half tone has been successfully achieved. Those experimental investigations, however, lack any theoretical foundation and, moreover, in some cases the morphed real world signals exhibit a disturbing level of nonlinear distortions, which was not observed for the synthetic test signals. As an additional result we have found that the morph of attractors with different topologies, for example morphing a closed line (harmonic signal) into a torus (non harmonic signal), is very difficult to achieve.

In the following article we will first explain our special interest for morphing one dimensional attractors, that is mainly due to the fact that non harmonic signals can be modeled by means of non stationary harmonic signals. As an example for such a situation we present the dynamical model of a piano sound. Then we investigate into the morphing algorithm based on some fundamental assumptions concerned with the dynamic modeling with normalized RBF networks. We restrict our focus to the problem of morphing phase or amplitude changes for signals with identical pitch and give some arguments why an independent and stable morph of phase and amplitude can be expected at least for infinitesimally small parameter changes. The case of pitch morphing is more difficulty to handle and will be considered in forthcoming investigations.
Besides the case of small parameter changes we investigate a special worst case large scale parameter change and argue that the morph of phase shifts of about π degree is impossible to achieve without altering the partials amplitude. However, it is interesting to note that in case of morphing pitch and large phase differences the impact on the amplitude is much smaller. As a last result we address the problem of nonlinear distortions and demonstrate that the problem is due to the undefined behavior of a dynamical model well apart from the trained attractor. We demonstrate that the distortion can be attenuated by means of constraining the width parameters of the Gaussians of the RBF network from below.

The article is organized as follows. In section 2 we explain some fundamental aspects of the dynamic modeling of sounds and give an example for modeling a piano signal, which motivates our interest in dynamical models of one dimensional attractors. In section 3 we introduce our morphing algorithm and give a explanation of the fundamental mechanism. In sections 4 we present morphing results for a number of artificial sound signals that are intended to verify the theoretical reasoning. Section 5 concludes with a summary and an outlook on further work.

2. ATTRACTOR MODELING OF SOUNDS

Traditional musical instruments belong to the class of dissipative, nonlinear mechanical systems. The discrete time evolution of such a system may be described in a k-dimensional state space S by means of a mapping $f(.)$

$$\vec{z}_{n+1} = f(\vec{z_n}) \quad \vec{z} \in S \quad , \tag{1}$$

which connects the system state $\vec{z}$ at time n, with the system state at the next time step. For a stable system $f(.)$ and for large $n \rightarrow \infty$ the state $\vec{z_n}$ will be confined to a bounded and closed subset $A \subset S$ of the state space, which is called an *attractor* of $f(.)$ [13]. An attractor may be as simple as a point or as complex as a fractal set if the system dynamics are chaotic. If the dynamical system generates an output (sound) signal y_n, the characteristics of this signal are closely related to the topology of the attractor [7].

For a d-dimensional attractor generating a sound signal y_n the fractal embedding theorem [16, 15] ensures under weak assumptions concerning y_n and $f(.)$, that for $D > 2d$, the set of all *delayed coordinate vectors*

$$Y_{D,T} = \{n > n_0 : \vec{y_n} = (y_n, y_{n-T}, \ldots, y_{n-(D-1)T})\}, \tag{2}$$

with an arbitrary delay time T, forms an embedding of A in the D-dimensional *reconstruction* space G. Because an embedding preserves the characteristic features of A, especially it is one to one, it may be employed for building a system model. To achieve this we use a neural network to approximate the system function $f(.)$ in the reconstructed state space G, yielding a prediction model

$$y_{n+T} = f_N(\vec{y}_n). \tag{3}$$

For stable predictor models f_N the prediction can be iterated thereby establishing a model of the system dynamics on attractor A. The building of predictor based neural network dynamical models has been successfully applied to chaotic system dynamics [8, 1] and also to non chaotic sound dynamics of saxophone, piano and speech signals[11]. Note, however, that in the case of music and speech signals the system dynamics are generally not stationary. For slowly varying dynamics this situation can be described by a system undergoing a parameter variation and, therefore, following a sequence of attractors [14, 11]. The input/output relations of a dynamical sound model for a non stationary sound is shown in Fig. 1. Note the control input on the left that is required to model different dynamics at different times. For modeling single sounds we have been able to model the attractor sequence with the control input being a fixed linearly increasing function of time.

Properties of Reconstructed Sound Attractors

It is well known that the attractors of sound signals with harmonic spectra have the topology of a closed line [6]. Consequently, many sound signals that are originated by musical instruments are well described by dynamical models with one dimensional attractors. Moreover, from

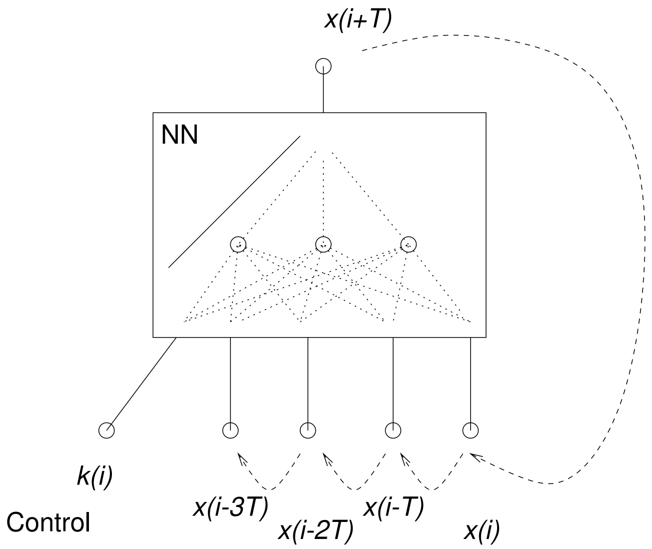

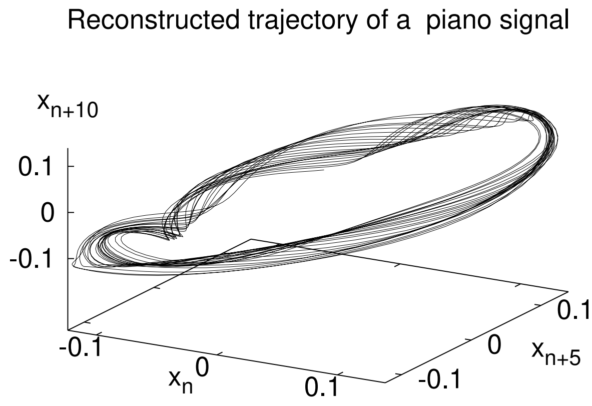

Figure 1: Input/output description of the neural network dynamical models (left) and a three dimensional reconstruction of a short segment of a piano signal (right).

the theory of linear prediction it is clear that a stationary signal that consists of a collection of k partials can be predicted without error by means of a linear predictor with $2k$ coefficients. Therefore, the prediction function f_N is linear for all points of the sound attractor if $D \geq 2k$. In contrast to the linear predictor which due to the roots on the unit circle is always unstable, the iterated prediction of the nonlinear dynamical model can be stable because the linear relation is not extended to the attractor neighborhood. While the dimension $D = 2k$ is well above the required reconstruction dimension it is often much easier to model the dynamics in this case, due to the simple linear relation that are required for prediction.

If some partials of a signal are not harmonic the attractor dimension is increased and exhibits the topology of an k-torus, where k is the number of independent base frequencies. An example is a sound of a piano, which is well known to be inharmonic. The attractor of the piano sound is a point in state space, the resting position of the string. For short times, however, we can neglect the time dependence of the dynamics and can interpret the reconstructed trajectories as part of a attractor, which due to the non harmonic partials has the topology of a torus. Because the frequency deviation of the partials from the harmonic relation is rather small a long time is needed until the trajectory describes the whole attractor. Due to energy dissipation the amplitudes of the partials change and the whole torus structure of the piano trajectory is not available. Part of the torus is shown in Fig. 1, where a three dimensional reconstruction of the decaying part of a piano sound is shown. While there does not exist a piano attractor that describes the relaxation of the string, we can describe the piano trajectory by means of a sequence of one dimensional attractors with continually changing phase and slowly decreasing amplitudes. Analysis of the piano model presented in [11] has revealed the fact that this is exactly the way the non stationary model shown in Fig. 1 is modeling the dynamics of the piano signal.

3. MORPHING DYNAMICAL MODELS

The new morphing algorithm we are going to present now, is based upon the homotopic mixing of dynamical systems [4]. In our approach we use as the progenitor systems the sound models obtained from the respective sounds. Then we construct a new sound model with an additional morphing parameter α that consists of the convex sum of the progenitor sound models, $f_1(.)$ and $f_2(.)$, following

$$f_m(\vec{y}, \alpha) = \alpha f_1(\vec{y}) + (1 - \alpha) f_2(\vec{y}). \qquad (4)$$

The morphing parameter α is confined to the interval $[0, 1]$, such that the parameterized model $f_m(., \alpha)$ establishes a linear transformation between the progenitor models. For $\alpha = 1.0$ and $\alpha = 0.0$ the morphing model $f_m(., \alpha)$ reproduces the progenitor models, and for intermediate values of α new sound dynamics are produced. Smooth changes of α result in smooth changes of the dynamical model $f_m(.)$. However, smooth changes in $f_m(.)$ does not necessarily result in smooth changes of the attractor, because for varying α all kinds of bifurcations may occur [4]. To achieve reasonable interpolating dynamics the progenitor attractors have to be geometrically and topologically similar.

To investigate into this problem let us assume that the reconstruction dimension is

$$D \geq 2k, \qquad (5)$$

with k being the number of partials in the quasi periodic signal, such that the predictors $f_1(.)$ and $f_2(.)$ are linear functions at their respective reconstructed attractors,

where the linear functions have all their roots at the unit circle. Further we assume that the nonlinear extension of the linear prediction into the neighborhood of the attractor can be described by means of a projection of the input vector $\vec{y}$ onto the attractor prior to the linear prediction. Due to the projection each of the iterated predictors is a stable oscillator. A dynamical model that is based on neural networks will hardly achieve this ideal behavior, however, if we employ a so called normalized RBF network [5], supply sufficient data, and continually increase the number of hidden units this type of prediction function is a possible limiting behavior of the network function. For finite number of units and for large distances to the attractor the projection and prediction mechanism is surely wrong. In this cases a large scale interpolation of the behavior of the predictor at different parts of the attractor will occur.
With this understanding of the prediction function we will now investigate into the results that can be expected for the above morphing algorithm and different characteristics of the changes of the signals parameters. The simplest situation occurs if both progenitor signals differ only in amplitude and phase of the partials. Therefore, we started our theoretical investigation with this case. A lot of work has to be applied before a full understanding of the proposed morphing scheme can be obtained. However, the following considerations are necessary foundations for a further development of the method.

Due to the fact that the optimum linear predictor does not depend on amplitude or phase of the predicted signal the linear functions that the predictors $f_i(.)$ calculate are identical. However, both predictors are constrained to different regions of the state space, due to the projection that is performed prior to linear prediction. In this case Eq. (4) can be expressed as

$$f_m(\vec{y}, \alpha) = \alpha \vec{a}'(\vec{y} \perp A_1) + (1 - \alpha)\vec{a}'(\vec{y} \perp A_2), \quad (6)$$

where A_i is the attractor of model i and the transposed column vector $\vec{a}'$ is the linear filter that describes both linear prediction functions. Note, that due to the linearity of the morphing equation Eq. (4) a morph between any two independent evolutions of the models $f_1(.)$ and $f_2(.)$ on their respective attractors will result in a new attractor with the same frequencies, but, with different phase andr amplitudes of the partials. If we start with an initial point $\vec{y}$ in Eq. (6) this state vector is first projected onto A_1 and A_2 to find the nearest points. By means of the linear prediction the projections are then moved one step ahead at the respective attractors and the resulting predictions are then combined by means of the linear interpolation to obtain the new state of the morphed dynamics. The next step starts again with a projection of the state vector onto the attractors. As long as this projection and the linear interpolation are inverse operations the whole process is simply an interpolation of two independent attractor motions. In such cases the morph between attractors with the same frequencies would not affect the frequency of the interpolated signals. In most cases, however, interpolation and projection are not inverse, and, therefore, frequency deviations may occur.
The closer the progenitor attractors are in the reconstructed state space the closer is the projection operation to the inverse of the linear interpolation and the smaller are frequency deviations for morphing amplitude or phase changes. If we continually change any single parameter of a progenitor signal (amplitude or phase), then the related attractor will change smoothly and will define a surface in the reconstruction space. Because all intermediate attractors are related to the same linear model, the surfaces that are related with different parameters can not intersect as long as the underlying signals are different. For small parameter changes the local change of the attractor is approximately linear and, therefore, for small changes a linear interpolation of those attractors that are related to the parameter limits will approximate the attractor for an intermediate parameter setting without considerable affect on other signal parameters. A morph by means of linear interpolation can be interpreted as a first order approximation of the dependence of the attractor on amplitude or phase changes.

For large differences of phase or amplitude, however, the linear interpolation will not be able to morph amplitude and phase changes without affecting the other parameter. This is due to the fact that the linearization of the attractor behavior around a parameter setting is only a crude approximation for larger parameter changes. A special situation is a phase shift of one partial of about π. Note, that reconstructed attractor will not change if the signals time origin is changing. Therefore, only relative phase changes of the partials are considered here. Shifting a partial phase by π yields the same result as a change in amplitude to the inverse sign. Because the signals are identical, so are the attractors and by means of morphing, a independent morph of phase or amplitude changes can not be expected. Think of two progenitor trajectories that are concentric straight spirals with identical radius where one spiral is shifted by π radiant along the axial center. Due to the ideal symmetries of this setting the morphing algorithm, driven by means of the global superposition of all the hidden units would solely change the radius of the partials. However, due to the curvatures of the attractors the ideal symmetries will never exist and, therefore, the results for morphing such signals will always show phase and amplitude changes. The amount of phase and amplitude changes that take place in practice depends on the specific signal and attractor characteristics.

The Effect of the Width of the RBF Gaussians

The above discussion is based upon a simple notion of the neural network dynamical models. Due to the fact that training data is only available at the progenitor attractors the network behavior in other regions of the state space is to some extend random. Deviations from the trained attractor may result in nonlinear signal distortions. In our first investigation of the morphing algorithm we have adapted all the network parameters without constraints

Signal abbrev.	Fundamental $w_0/2/\pi$	Amplitude a_3	Phase ϕ_3	Frequency shift w_r	Attractor dimension
BAS	$\frac{100+\pi}{6000}$	0.3	0.0	1.0	1
BASPID3	$\frac{100+\pi}{6000}$	0.3	$\frac{\pi}{3}$	1.0	1
BASPI	$\frac{100+\pi}{6000}$	0.3	π	1.0	1
HIG	$1.0595\frac{100+\pi}{6000}$	0.3	0	1.0	1

Table 1: Parameter settings of Eq. (7) for all synthetic signals that has been investigated in the following experiments.

[12]. As a result the range of the widths of the different RBF units is large. The width parameter affects the region of influence of a hidden unit, which might be rather large due to the normalization of the activation function. If two adjacent hidden units have differing width then the influence of the one with larger width is for some parameter settings global all around the hidden unit with the small width, which is only locally effective. This will not harm as long as the model is used only in a close neighborhood of the attractor, because the effect might help the model to decrease the error. However, in larger distance the effect is unpredictable and, moreover, may destroy the determination of the model behavior by the nearest point of the attractor, in effect what is the projection in the above discussion. There is a further point concerning the width parameter of the units that has to be considered. For an increasing number of RBF units the width parameter will generally decrease. As a consequence in a larger distance away from the attractor, especially if there exist regions with high curvature on the attractor, the model will exhibit regions with steep transitions. Morphing the attractors through such regions will generally result in a high amount of nonlinear distortion. This undesirable behavior of the model can be improved to some extend if the width parameters of the model is constrained from below.

Attractors with Local Loops

Local loops are typically obtained for reconstructed attractors of sound signals with an considerable amount of energy in different partials. This local loops are difficult to model, because, the points where the trajectory closely returns demands high precision of the model and are likely to cause unstable behavior. In the light of the morphing algorithm this local loops can generate further difficulties, if this local loops are not present similarly for both progenitor attractors. One goal for the attractor reconstruction is, therefore, to prevent local loops. Local loops are due to high similarity of delayed coordinate vectors over time distances other than an integer multiple of the base period of the signal. These local loops are often due to the existence of local extrema of the time signal. If the delayed coordinate vectors cover only a short time segment the edges of local and global extrema are mapped to close positions in state space such that local loop structures arise. To prevent the local loops the delayed coordinate vectors should cover a considerable amount of the base period such that effects of short time scale similarities are suppressed. Due to the fact that the delay time T is limited from above due to Nyquists sampling theorem we have to increase the reconstruction dimension until a sufficient length of the delayed coordinate vectors is achieved. The best one do is to increase the D until the product DT is equal to the base period of the signal. The increased dimension does not harm the model, despite the increased calculation demands, but is also desirable with respect to the linearity limit of the prediction function given in Eq. (5).

4. EXPERIMENTAL VERIFICATIONS

In the following section we want to experimentally verify some of the results we have obtained by means of theoretical investigation. Because we already have demonstrated that the method can be applied to real world signals [12] we restrict the following experiments to synthetic sound signals, where we have fine control to set up and study interesting situations. The mathematical description of all signals we will study in the following is

$$y_n = \sin(w_0 n) + a_3 \sin(3 w_0 n + \phi_3) + 0.2 \sin(5 w_0 n). \tag{7}$$

The specific parameter settings we choose for the following experiments are shown in Tab. (1). Note that in our previous investigation we used the same basic equation, however, with different parameters. As explained earlier, we are mainly interested in morphing harmonic signals, and, therefore all the signals we use in the following have one dimensional attractors. The signal with indicator *BAS* is the basic harmonic signal with no phase shift, for *BASPID3* and *BASPI* the third harmonic partial has been shifted by means of $\pi/3$ or π radiant respectively, *HIG* has increased pitch of one half tone.

For all the signals given in Tab. (1) we trained dynamical sound models as explained in [11]. The neural networks used are of the radial basis function type with normalized hidden units. The network implements a function

$$\vec{N}(\vec{y}_k) = \sum_j \vec{w}_j \frac{\exp\left(-(\frac{\vec{c_j}-\vec{y}_k}{\sigma_j})^2\right)}{\sum_i \exp\left(-(\frac{\vec{c_i}-\vec{y}_k}{\sigma_i})^2\right)} + \vec{b}. \tag{8}$$

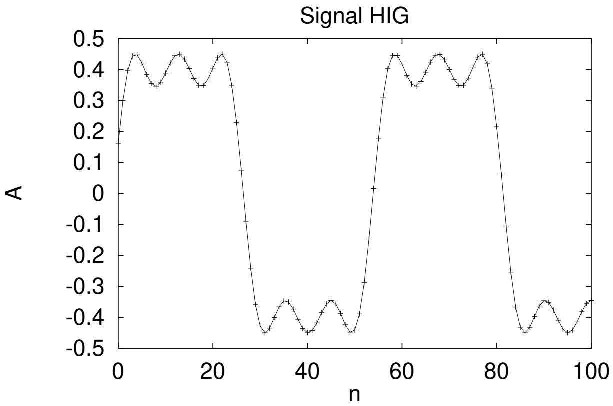

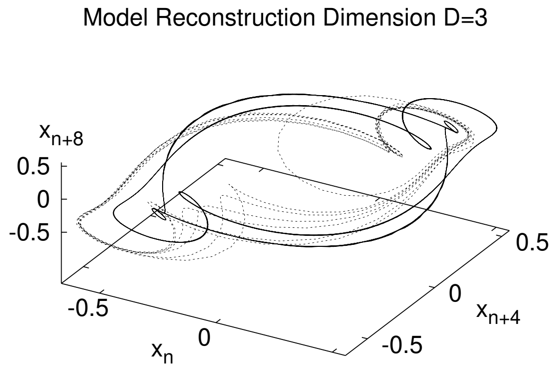

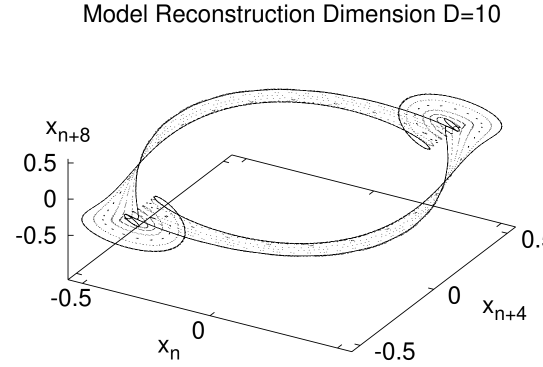

Figure 2: Results for morphing *HIG* to *BASPID3* . The time signal *HIG* (top left) and resynthesized attractors for two different reconstruction dimensions are shown. The solid lines represent the cases $\alpha = 0.0$ and 1.0 while the dotted lines represent intermediate α. The morph is unstable for $D = 3$ although a proper embedding is achieved (top right). Best results are obtained if local loops are suppressed as far as possible (bottom right).

The network parameters $\vec{w}$, $\vec{c}$ and $\vec{b}$ are adjusted by means of a standard training algorithm *RPROP* [10] to obtain optimal prediction of the following sample y_{k+T}. Iteration of the model and up sampling the output signal (to compensate for the step size T) yields a time series that, for a stable model, closely resembles the training signal. In contrast to our previous investigation we have chosen to use recurrent training of iterated prediction, because the recurrent training algorithm will generally adjust the predictor in the neighborhood of the attractor in a way to increase model stability [2]. For our morphing algorithm the design of the neighborhood of the attractor is important, and, therefore, we use it here even if the iterated prediction is generally stable without recurrent training. Because the Lyapunov exponents of our attractors are non positive, the number of recurrent iteration during training can be selected without constraints. Here we selected 6 recurrent iterations for each input vector during training.

In all following experiments we have used a delay time $T = 4$. This selection is to some extend arbitrary and is not critical as long as the subsampling of the time series by the factor Tdoes not violate the Nyquist sampling theorem. In the latter case it would be impossible to reconstruct the time series from the iterated dynamical model.

In our first experiment we want to demonstrate the impact of the reconstruction parameters on the morph. We used the signals *BASPID3* and *HIG* and trained dynamical models for different reconstructions. In the first case a reconstruction dimension $D = 3$ is used, which is just the sufficient embedding dimension for the one dimensional attractors. In the second case we increase the dimension until a linear prediction on the attractor is sufficient for prediction $D = 6$, and, in the last case we increase the dimension until all local loops are maximally suppressed. To achieve the last goal we apply a simple rule. As one can see from the display of the time series in Fig. 2 there exists two parts in one period of the time series with rather steep slopes. By selecting D such that for all times n at least one component of the delayed coordinate vectors falls into this region we increase the minimal distance for any pair delayed coordinate vectors besides the ones that are neighbors in time. Note that for all these dynamical models we achieved stable iterated prediction, such that we can interpret them as stable dynamical models. However, the models differ significantly when used for our morphing purposes. Typical results we have obtained are displayed in Fig. 2. Here we display the results not in there respective reconstructed state spaces, but, due to dimension limitations, as 3 dimensional reconstructions. For the 3 dimensional model we find that all models become unstable if α differs significantly from 0 or 1. In the case shown the model becomes unstable for $\alpha = 0.013$. In the case $D = 6$ (not shown) where linear

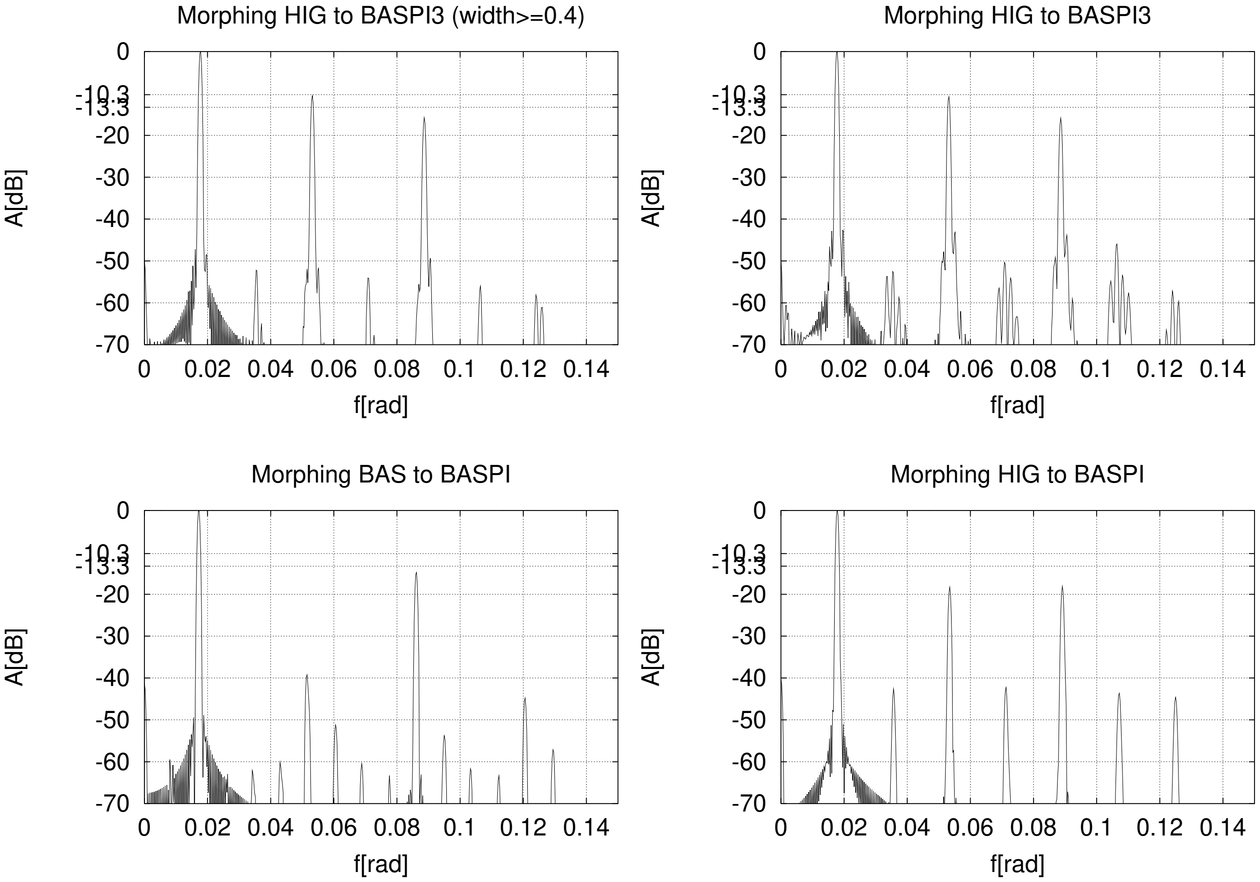

Figure 3: Spectral analysis of morphed signals with $\alpha = 0.5$. Morphing dynamical models *BAS* to *BASPID3* with unconstrained and constrained width parameter (top). The constrained models achieve less distortions. The bottom line compares two different situations with phase difference $\phi_3 = \pi$. The morph of *BAS* to *BASPI* first attenuates the second partial by approximately 25dB and then adjusts the phase. Morphing *HIG* to *BASPI* requires less attenuation of the second partial of 7dB, only. The levels of the partials of the original signals are indicated by grid lines.

relations at the attractor are achieved, the morph is stable, however, there exists nonlinear distortions. This distortions disappears for $D = 10$. In the last case the local loops are sufficiently unfolded to achieve a smooth neighborhood of the predictor around the trained attractor.

The second experiment is concerned with the impact of the width parameter of the normalized RBF networks on the distortions that are obtained during morphing. We used the same signals as before and trained two types of networks with 10 input units, 30 hidden units, 1 output unit and delay time $T = 4$. The networks differ with respect to the constraint that is applied to the width parameter. Here we compare the results obtained for unconstrained networks and for networks with the constraint $\sigma_j > 0.4$. In each case we have trained four different initializations such that all networks achieve comparable RMSE of about $4 \cdot 10^{-4}$ and are stable dynamical models. If we analyze the signal synthesized by means of the morphing model Eq. (4) with the morphing parameter set to $\alpha = 0.5$ we find that the unconstrained model obeys a significant increase of nonlinear distortion. The DFT spectra of two representive examples are shown in the top line of Fig. 3. Because only three partials are present in the signal we would like an ideal morph to be restricted to only three partials, also. From the higher number and higher amplitude of additional partials that are present when we morph the unconstrained models we conclude that the constrained models achieve a much better extrapolation of the dynamics into the neighborhood of the attractor. Similar results have been obtained for the constraints $\sigma_j > 0.5$ and $\sigma_j > 0.6$. The optimal selection of this constraint, however, which can be expected to depend on both progenitor attractors is not solved up to now.

Our last example is concerned with the morph of phase differences of about π. The dynamical models we used here have the same topology than the ones in the previous experiment. The signals we use are *BAS*, *BASPI* and *HIG*. The spectra for a morph from *BAS* into *BASPI* are shown in the bottom line of Fig. 3. We find that the third harmonic which originally had a level of -10.3dB is attenuated by about 25dB for $\alpha = 0.5$. This has been expected from the discussion in the previous section. However, if we try to morph from *BASPI* into *HIG* we find that the attenuation of the harmonic is significantly smaller, while the distortions are increased. At present we do not have a explanation for the latter observation. We conjecture, that the difference in pitch destroys the symmetries that otherwise prevent the direct morph of the phase. Note, however, that in the latter case the added partials are all harmonic and, therefore, subjectively less annoying.

5. OUTLOOK AND SUMMARY

In the present article we have investigated into a new algorithm for morphing dynamical models. We have motivated our special interest for one dimensional attractors and have analyzed some fundamental properties of reconstructed attractors of periodic or quasi periodic signals. Those properties leads us to conjecture that a reconstruction dimension that is considerably above the embedding dimension should be used to achieve best morphing results. An experimental investigation supports our conclusions. Based on a simple model of how normalized RBF networks achieve function approximation, we give theoretical support for the morphing algorithm, and argue that by means of linear interpolation of the model predictions independent morphing of small amplitude and phase changes should be possible. However, if the phase differences of the progenitor signals partials approach π an independent morph of the phase can not be achieved. We demonstrated that morphing this kind of phase differences is always accompanied by considerable attenuations of the amplitude of the respective partial.
Moreover, we argue that the extrapolation of the prediction function into the neighborhood of the trained attractor requires some constraints on the network parameters to suppress random effects and distortion during morphing.

Obviously, their is a lot of further work necessary to understand the proposed algorithm. Current investigations are dealing with the morphing of pitch changes and the stability of the intermediate attractors. As an application we investigate the morph of pitch and volume changes of our piano sound model.

REFERENCES

[1] G. Deco and B. Schürmann. Neural learning of chaotic dynamics. *Neural Processing Letters*, 2(2):23–26, 1995.

[2] S. Haykin and J. Principe. Making sense of a complex world. *IEEE Signal Processing Magazine*, 15(3):66–81, 1998.

[3] E. Levin. Hidden control neural architecture modelling of nonlinear time varying systems and its applications. *IEEE Transactions on Neural Networks*, 4(2):109–116, 1993.

[4] R. Mettin and G. Mayer-Kress. Chaotic attractors from homotopic mixing of vector fields. *International Journal of Bifurcation and Chaos*, 6(2):395–408, 1996.

[5] J. Moody and C. Darken. Fast learning in networks of locally-tuned processing units. *Neural Computation*, 1:281–294, 1989.

[6] T. Parker and L. Chua. Chaos: A tutorial for engineers. *Proc. of the IEEE*, 75(8):982–1008, 1987.

[7] T. Parker and L. Chua. *Practical Numerical Algorithms for Chaotic Systems*. Springer-Verlag, New York-Heidelberg-Berlin, 1989.

[8] J. C. Principe and J.-M. Kuo. Dynamic modelling of chaotic time series with neural networks. In G. Tesauro, D. S. Touretzky, and T. Leen, editors, *Neural Information Processing Systems 7 (NIPS 94)*, 1995.

[9] J. C. Principe, A. Rathie, and J.-M. Kuo. Prediction of chaotic time series with neural networks and the issue of dynamic modeling. *Int. Jour. of Bifurcation and Chaos*, 2(4):989–996, 1992.

[10] M. Riedmiller. Advanced supervised learning in multilayer perceptrons – From backpropagation to adaptive learning algorithms. *Computer Standards and Interfaces, Special Issue on Neural Networks*, 5, 1994.

[11] A. Röbel. Neural network modeling of speech and music signals. In *Neural Information Processing Systems 9 (NIPS 96)*, pages 779–785, 1997.

[12] A. Röbel. Morphing dynamical sound models. In *Proceedings of the 1998 IEEE Workshop on Neural Networks for Signal Processing VIII*, pages 409–418, 1998.

[13] D. Ruelle. Characteristic exponents for a viscous fluid subject to time dependent forces. *Commun. Math Phys.*, 93:285–300, 1984.

[14] D. Ruelle. Diagnosis of dynamical systems with fluctuating parameters. *Proc. of the Royal Society London*, 413:5–8, 1987.

[15] T. Sauer, J. A. Yorke, and M. Casdagli. Embedology. *Journal of Statistical Physics*, 65(3/4):579–616, 1991.

[16] F. Takens. *Detecting Strange Attractors in Turbulence*, volume 898 of *Lecture Notes in Mathematics (Dynamical Systems and Turbulence, Warwick 1980)*, pages 366–381. D.A. Rand and L.S. Young, Eds. Berlin: Springer, 1981.

[17] E. Tellman, L. Haken, and B. Holloway. Timbre morphing of sounds with unequal numbers of features. *Journal of the Audio Engineering Society*, 43(9):678–689, 1995.