
The Virtual Flute: An Advanced Fingering Guide Generated Via Machine Intelligence

Andrew Botros, John Smith and Joe Wolfe

Music Acoustics, School of Physics, University of New South Wales, Australia

Abstract

The Virtual Flute is a web service that provides many thousands of machine-predicted alternative fingerings. Alternative fingerings can offer variations in intonation and timbre, and can be easier to play in different musical contexts. Many play multiphonics. An advanced fingering guide is invaluable when exotic effects and demanding passages are required of the player. *The Virtual Flute* uses an expert system that predicts musical properties from acoustic impedance spectra. Impedance spectra for the 39,744 acoustic configurations of the flute are generated by a physical model of the instrument, some parameters of which are machine-learned. We report the construction and use of *The Virtual Flute* (<http://www.phys.unsw.edu.au/music/flute/virtual/>).

1. Introduction

When a musician is faced with a technically difficult phrase in a performance, a resource that offers an extensive guide to the instrument's available options is invaluable. One of the primary goals of music acoustics is to offer assistance to performing musicians, and *The Virtual Flute* does just that. By making appropriate physical approximations and taking advantage of computational power, *The Virtual Flute* maps the entire fingering space of the modern flute for the benefit of performers and composers.

A fingering is a given combination of pressed keys on a woodwind instrument. A pressed key closes or opens

particular tone holes along the length of the instrument; thus, a fingering defines an acoustic configuration. For the modern Boehm flute, there are 39,744 unique acoustic configurations (including both B foot and C foot models).

The range of the flute spans 44 notes in the equal tempered scale (from B3 to F \sharp 7), and for most notes, one or two fingerings are recognized as standard. Thus, given the many thousands of acoustic configurations available, alternative fingerings abound. In compositions since the 1950s, woodwind players are increasingly required to play multiphonics (chords, in which two or more notes are sounded simultaneously), microtones (notes that fall between those of standard temperaments), and notes with various timbre effects (e.g. Berio, 1958; Bartolozzi, 1967). Alternative fingerings can offer differences in pitch and pitch flexibility, timbre, stability, ease of production, or ease of fingering (particularly in rapid passages), all of which may be desirable in a given musical context.

We know of no previous attempt to present the entire list of alternative fingerings for the flute. Dick (1989) gives an extensive collection of fingerings but it is far from complete. Hence, *The Virtual Flute* completes the flute fingering landscape for many musicians worldwide.

2. An overview of *The Virtual Flute*

The musical response of the flute for a given fingering can be largely determined from the spectrum of the acoustical impedance $Z(f)$ – the ratio of acoustic pressure to volume flow of air – measured at the embouchure hole of

the flute¹. The resonances of the flute correspond closely to the minima of acoustic impedance, and flutes play notes whose frequencies are close to these resonances.

For many standard fingerings, all tone holes are closed down to a certain point and mostly open beyond that. In such a configuration, the flute acts like a tube, open at both ends, whose resonances occur approximately in a harmonic series. The effective tube length L is approximately that between the embouchure hole and the first open tone hole; the minima of $Z(f)$ correspond to standing waves with wavelengths of $2L/n$, where n is an integer. The flute can operate using one of these resonances as the fundamental, with higher harmonics contributing to the overall brightness of the played note.

Amongst all fingerings however, such simple fingerings are greatly outnumbered by cross-fingerings, where one or more tone holes are closed downstream from the first open hole. Within the instrument bore, some of the travelling wave is reflected at each open hole whilst some is transmitted downstream to the remaining open holes (Wolfe & Smith, 2003). These separate reflections can give rise to different standing waves, which the player may be able to excite simultaneously. As the lengths involved are not in general simple harmonic ratios, multiphonics can be produced. Additionally, the lack of simple harmonic resonances in cross-fingerings leads to significant differences in timbre. *The Virtual Flute* examines thousands of cross-fingerings that have not been previously documented.

Our technique for measuring the acoustic impedance of flutes provides sufficiently many accurate data for modelling purposes (Wolfe et al., 2001a, b). At the core of *The Virtual Flute* is a simple model of acoustic impedance that agrees well with our experimental measurements. Secondly, *The Virtual Flute* consists of an expert system that relates features of impedance spectra to key musical properties, such as the playability of the note associated with any given impedance minimum. Armed with these two models, a machine analyses every fingering of the flute (Figures 1 and 8 – see Section 5 – illustrate this for the fingering Th 1 - 3 | 1 2 - D_♯ C²). *The Virtual Flute* is the interface to these musical data.

The Virtual Flute was first available in 2001 (Botros et al., 2002, 2003). In 2005, a second release of *The Virtual Flute* was launched, offering a more accurate and

elegant physical model and a wider set of features for musicians. In particular, a machine learning approach is employed, resulting in an automated and woodwind-generic method of modelling acoustic impedance. This paper describes the current release of *The Virtual Flute* in its entirety.

3. A machine-learned model of acoustic impedance

3.1 One-dimensional acoustical waveguide (transmission line) models

The flute has no strong resonances above about 3 kHz, and hence cannot play in this range. Consequently, the wavelengths of the sounds of interest are much longer than the diameter, and the propagating waves within the bore are pre-dominantly planar³. Hence, a one-dimensional approximation of woodwind geometries can achieve acceptable precision with some semi-empirical corrections.

To calculate $Z(f)$ given such an approximation, a woodwind is first represented as a network of small cylindrical and conical sections. For the flute, each tone hole and each bore segment of the body is represented as a cylindrical section, whereas the tapered head joint and embouchure hole can be represented as a sequence of cylindrical or conical sections. To make a numerical calculation, one starts from the downstream end of the instrument and works back towards the embouchure. The acoustic impedance at the end of a pipe is that of the radiation field. This is used as the load impedance for the section of the bore leading to the first tone hole. The input impedance of the section is calculated using a transfer matrix (for relevant formulae, see Fletcher & Rossing, 1998). The hole is also a (short) cylindrical section whose load is either another radiation load (if open) or infinite (if closed). These two sections are combined in parallel to form the load for the next section of the bore:

$$Z_{\parallel} = \frac{Z_{\text{bore}}Z_{\text{hole}}}{Z_{\text{bore}} + Z_{\text{hole}}}. \quad (1)$$

The process continues to the input of the instrument, and the entire calculation is repeated at different frequencies to create an impedance spectrum. To compensate for higher-dimensional effects, one-dimensional models typically include a few adjustable parameters, such as end corrections at each section.

³Thus the flute is an acoustical waveguide, and many of the techniques developed for waveguide analysis, e.g. in microwaves may be applied. (The use of “waveguide” is distinct from digital waveguide modelling used in synthesis.)

¹Throughout this paper, impedance is specified on a logarithmic scale: $Z(\text{dB}) = 20 \log_{10} (|Z|/\text{Pa s m}^{-3})$. Z is a complex quantity, but examination of the phase spectrum gives little extra information (Wolfe et al., 2001a). Consequently, we show only magnitude spectra.

²Flutists write fingerings using the numbers 1, 2, 3 for the keys usually operated by the long fingers of each hand (left given first), Th for the left thumb key, and then individual names for the other keys (such as the D_♯ key). A vertical line separates the two hands.

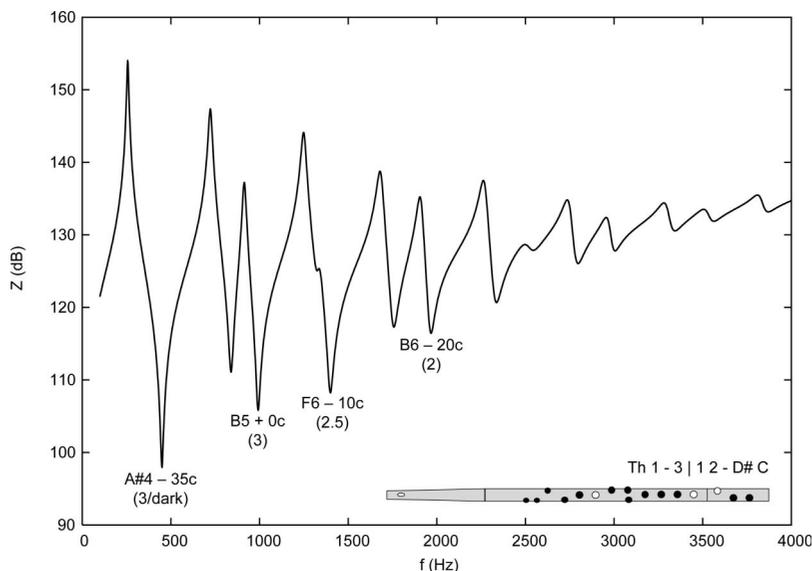


Fig. 1. *The Virtual Flute*'s impedance model and expert system predictions for a randomly chosen cross-fingering (open holes denoted white, closed holes black). Each playable minimum is labelled with its intonation (offset from equal temperament in cents) and its playability level (3 maximum). $A\sharp 4$ is predicted to be dark in timbre. Predicted multiphonics are $A\sharp 4$ & $B5$, $B5$ & $F6$, $F6$ & $B6$.

Using essentially these methods, impedance calculations have been made of a variety of wind instruments, such as the oboe (Plitnik & Strong, 1979), the flute (Coltman, 1979; Strong et al., 1985; Nederveen, 1998) and brass instruments (Caussé et al., 1984).

3.2 Adjusting parameters by machine learning

As mentioned, previous one-dimensional models have included empirical parameters to compensate for the inherent inaccuracies of their approximations. Plitnik and Strong (1979) used the measured end corrections of Benade and Murday (1967) at open tone holes. Coltman (1979) measured the end corrections of open flute tone holes, and the same end corrections were used in the model of Strong et al. (1985). Strong et al. (1985) also widened the diameter of the flute bore underneath each tone hole, and energy loss factors were increased at each hole. Coltman (1979) used increased energy loss factors throughout the entire flute.

Previous models have not had the benefit of large sets of precise $Z(f)$ measurements to validate their parameter choices. Furthermore, each parameter value is typically chosen by measurements at isolated sections of the instrument (such as a single hole), rather than by evaluating the interaction of all parameters at the instrument input (the ultimate point of interest).

Our novel approach is to use a machine to adjust multiple parameters simultaneously, iteratively improving the agreement with our database of experimental measurements. The database provides a training set for the machine learning algorithm `LEARNIMPEDANCE`. Whereas the first version of *The Virtual Flute* used an

impedance model with manually adjusted parameters (much like previous models), the current version has been optimized in a completely automated fashion. Moreover, the impedance software defines a generic schema (written in XML) for an arbitrary woodwind geometry – `LEARNIMPEDANCE` can be applied to various woodwinds without alteration.

3.3 Model parameters

We model a Pearl PF-661 B foot flute, configured in the same manner as Wolfe et al. (2001a). A total of six adjustable parameters are used – five end corrections and an energy loss factor. All of the parameters represent well-known physical effects and have been incorporated in previous models. The model geometry and adjustable parameters are shown in Figure 2 and specified in detail in Appendix A.

3.4 The `LEARNIMPEDANCE` algorithm

The goal of the `LEARNIMPEDANCE` algorithm is to minimize the error between the physical model described above and our experimental measurements. This is achieved by searching for an optimal setting of the adjustable parameters (a total of six for the flute). Given the expensive computational requirements, an exhaustive search through all parameter combinations is impossible – for the required parameter resolution, the total number of combinations is in excess of billions, and each impedance spectrum takes ~ 1 s to compute. Thus, our algorithm must perform a “greedy” search through parameter space.

The algorithm we designed is, in essence, a combination of a gradient descent search and the expectation-maximization (EM) algorithm (Dempster et al., 1977). LEARNIMPEDANCE iterates over a two-step process until little improvement is gained between iterations:

1. Determine which parameter, when adjusted in isolation, would reduce the error between the model and the experimental measurements (training set) by the largest amount.
2. Following the ordering established in step 1, adjust each parameter in turn to minimize the error.

The algorithm is specified in detail in Appendix B.

Ten impedance spectra are used to train the model – standard fingerings for B3, C4, C_♯4, F_♯4/5, B4/5, D6, D_♯6, C7, D_♯7 and E7 are included. Of these, five are cross-fingerings. The error at any point during the algorithm’s execution is the mean squared difference between the model and measured spectra (in log terms), averaged over the ten training spectra. Note that the use of log Z means that both maxima and minima contribute equally to the quality of the fit. One might argue that a flute model should favour impedance minima. We prefer, however, to use a single model for all woodwinds and thereby to constrain more tightly the performance of the acoustical model.

All end corrections are initially set to 0 and the energy loss factor is set to 1. The algorithm may only adjust parameters within a small window at each iteration (± 0.1 for the end corrections and ± 2.5 for the energy loss factor). This ensures that the model is not too heavily influenced by the first parameters that are adjusted. End corrections can be adjusted without bound, but the energy loss factor must not be less than 1.

3.5 Results and discussion

LEARNIMPEDANCE was executed over 13 iterations, a total time of 4.5 h. The final parameter values are shown in Table 1. Figure 3 compares the model with experimental measurements for three fingerings that have not been used for training.

When calculated over the frequency range of 0.2–4.0 kHz, the average root mean squared (RMS) error for

40 standard fingerings was 1.16 dB. When calculated over 0.2–3.0 kHz – the musical range of the flute – the RMS error over the 40 fingerings was 1.28 dB. This compares favourably with the first impedance model of *The Virtual Flute* (Botros et al., 2002). To 4 kHz, the original model results in an error of 1.47 dB; to 3 kHz, the error is 1.67 dB. Furthermore, whereas the current impedance model uses six adjustable parameters, the original model used nine. The advantages of the current approach are clear: it is more accurate, less complex, automated and woodwind-generic.

To demonstrate the woodwind-generic framework of LEARNIMPEDANCE, We applied the algorithm to a wooden classical flute. The instrument consisted of eight finger holes and one hole covered by a keypad. In addition to the six adjustable parameters of the modern flute, we added two end corrections (open and closed) for a fingered hole. This is necessitated because, when closed, the ball of the finger intrudes into the hole; when open, a finger lifts further from the hole than does a keypad. All other parameters matched those of the modern flute. LEARNIMPEDANCE was trained by eight impedance spectra, three of them from cross-fingerings.

LEARNIMPEDANCE was executed over 19 iterations for the classical flute (6.5 h). The final parameter values are shown in Table 1. Figure 4 compares the model with experiment for an unseen cross-fingering. Over the frequency range of 0.2–4.0 kHz, the RMS error was 0.99 dB when averaged over our full dataset of

Table 1. Final LEARNIMPEDANCE parameter values for the modern flute and classical flute models (parameter specifications given in Appendix A).

Parameter	Modern flute value	Classical flute value
m_z	14.75	14.00
$\Delta l_{\text{bore-emb}}$	-0.15	-0.02
$\Delta l_{\text{bore-hole}}$	-0.06	-0.01
$\Delta l_{\text{emb-bore}}$	0.63	0.63
$\Delta l_{\text{hole-bore}}$	0.41	0.35
$\Delta l_{\text{openhole-keypad}}$	1.07	1.90
$\Delta l_{\text{openhole-finger}}$	-	0.47
$\Delta l_{\text{closehole-finger}}$	-	-0.83



Fig. 2. The geometry and adjustable parameters of the modern flute impedance model. Bore segments and tone hole segments are modelled by cylindrical sections; head joint and embouchure hole modelled by conical sections. The adjustable parameters include five end corrections and one energy loss factor localized at the embouchure region.

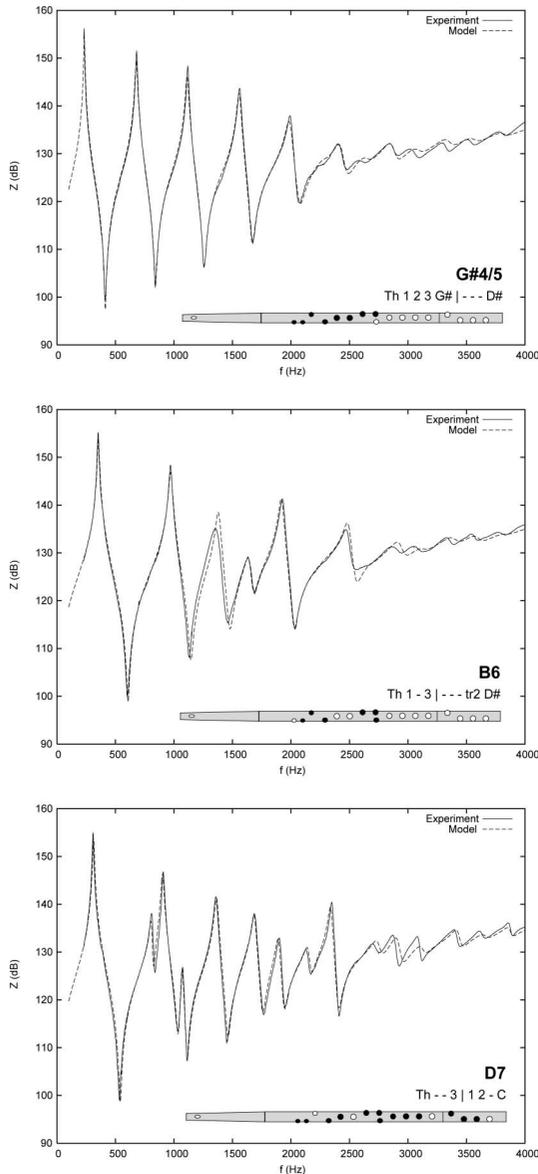


Fig. 3. Comparison of impedance model and experimental measurement for three standard fingerings of the modern flute. The measurements were not used to train LEARNIMPEDANCE. B6 is the worst fit in our dataset; $G\#4/5$ and D7 are typical fits.

19 impedance spectra. With the possible exception of $\Delta I_{\text{openhole-keypad}}$ – keypad design varies across instruments – the parameter values generally agree with those for the modern flute.

Six adjustable parameters may seem extravagant in a theoretical model: in many experimental datasets, a fit with six parameters would be doomed to success. Here, however, the datasets are very demanding. In the impedance spectrum for each measured fingering, the impedance varies by a factor exceeding a thousand, and it has typically about 20 very sharp extrema, with well-defined impedance and frequencies. Fitting the many detailed features of one such spectrum is difficult; further, there are tens of spectra to be fitted with just the one set of six

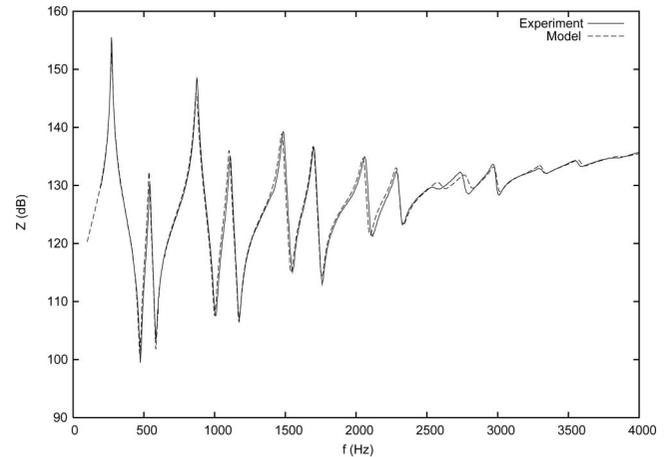


Fig. 4. Comparison of impedance model and experimental measurement for a classical flute cross-fingering. The measurement was not used to train LEARNIMPEDANCE.

parameters. Our empirical approach provides a flexible means of building accurate models. Adjustable parameters are easily added and subtracted from the model; the algorithm uses the parameters if they significantly improve the model. Whereas more complex, higher dimensional models exist for specific geometries such as single tone holes, our simple model adapts to the instrument of interest and quickly provides useful musical data.

4. Predicting musical responses from acoustic impedance spectra

4.1 Extracting features from $Z(f)$

To study the relationship between $Z(f)$ and musical response, 76 impedance spectra were selected from our experimental database for analysis (cf. Botros et al., 2002). The frequency range studied was 0.2–4.0 kHz, covering the range of all playable notes on the instrument. 957 minima are contained within the selected spectra.

For each extremum (minimum or maximum), a set of features were calculated: f_e , the frequency of the extremum; Z_e , the magnitude of $Z(f)$ at frequency f_e ; and B_e , the bandwidth of the extremum. In addition to these three parameters, each minimum was characterized by a further 11 features: the Q factor ($Q = f_m/B_m$); the absolute frequency difference and relative impedance difference of the adjacent left minimum ($\Delta f_{L\min}$, $\Delta Z_{L\min}$), right minimum ($\Delta f_{R\min}$, $\Delta Z_{R\min}$), left maximum ($\Delta f_{L\max}$, $\Delta Z_{L\max}$) and right maximum ($\Delta f_{R\max}$, $\Delta Z_{R\max}$); the number of harmonic minima (N_{harm}); and the weighted average impedance of those harmonics (Z_{harm}).

Minima at frequencies above f_m were deemed to be harmonic if their frequencies were in the range $n(1 \pm 0.05)f_m$, where n is a positive integer. N_{harm} is the total of such harmonic minima. The harmonicity function Z_{harm} is the

average impedance of the harmonic minima, where each minimum has a weighting of $1/n$.

Collectively, the 14 features represent physical parameters that we initially suspected to be musically important – for example, we expect notes to be easy to play if they are associated with strong resonances (i.e. at minima that are deep and isolated), and the harmonicity of higher minima might be important in “mode locking” of the nonlinear oscillation régime of the jet (Fletcher, 1978).

4.2 Predicting playability

The presence of a minimum in $Z(f)$ does not necessarily mean that a note can be played at that pitch. To distinguish between playable and unplayable minima, an expert system that predicts the playability of minima based on their physical features was constructed. An experienced flutist ranked the playability of notes corresponding to each of the 957 minima in our selected dataset. Minima were ranked subjectively into four levels of playability, from 3 (most playable) to 0 (unplayable). Thus, our training dataset consists of 957 instances, each consisting of the 14 minima features and the expert determination of playability.

The playability expert system of *The Virtual Flute* consists of two separate models: (i) minima are firstly predicted as playable or unplayable; (ii) thereafter, playable minima are ranked on a continuous scale of 0–3.

The first model was machine-learned using the C5.0 decision tree algorithm (Quinlan, 1993, 2005a). Cross-validation was used to test the performance of the decision tree with unseen data. In this process the expert data are randomly divided into F subsets (folds), and in each iteration a single subset is withheld from the C5.0 algorithm and used as a test set. A decision tree may be pruned to remove spurious dependence on any of the minima parameters that do not improve the error rate of the tree – C5.0 allows the analyst to specify a minimum number of training instances that at least two branches must carry at a decision node. Guided by cross-validation results, the analyst may prune a decision tree to an appropriate size. The playability tree was selected in this manner. Table 2 shows the tree size and error rate at different levels of pruning. The error rate is the 10-fold cross-validation error, averaged over 10 trials. As can be seen from the results, large trees overfit the training data with insignificant branches, performing no better than smaller trees. The selected tree, shown in Figure 5, specifies a minimum of 30 instances at each branch, resulting in a 5.2% error rate. At higher levels of pruning, the error rate rises.

The second model was also machine-learned, using the Cubist model tree algorithm (Quinlan, 1992, 2005b). The training dataset consists of only the playable minima

Table 2. C5.0 playability tree size (number of leaf nodes) and 10-fold cross-validation error rate at different levels of pruning. The pruning level is the minimum number of instances that at least two branches must carry at a decision node. Results are averaged over 10 trials. Selected model is highlighted.

Pruning level	Tree size	Error rate (%)
2	17.7	5.9
10	10.1	5.3
20	7.3	5.5
30	5.9	5.2
40	5.7	5.9
50	5.1	8.2
100	2.7	11.3
200	2.0	11.7

from the original dataset, providing 329 instances. Model trees work with continuous target functions, as required; other techniques include linear regression, instance-based locally-weighted learning and neural networks. These four techniques were compared using the Weka Experimenter (Witten & Frank, 2000); based on cross-validation results over the training data, model trees consistently performed the best. Similar to C5.0 pruning, Cubist allows the analyst to specify a minimum percentage of instances that a rule must cover. The playability model tree was selected by comparing the correlation between the training data and model tree predictions at different levels of rule coverage. Table 3 shows the cross-validation correlation at different levels of rule coverage. A level of 20% was selected based on the results, giving the following playability model (f and Δf in Hz, Z and ΔZ in dB):

RULE 1: If $Z_m > 103.2$ dB and no left-neighbouring minimum exists

$$Play = 4.4 - 0.022Z_m + 0.0002f_m + 0.0005\Delta f_{Rmax}$$

RULE 2: If $Z_m > 103.2$ dB and $\Delta Z_{Lmin} \leq 0.4$ dB

$$Play = 2.3 + 0.034\Delta Z_{Lmax} - 0.011Z_m + 0.00003f_m$$

RULE 3: If $Z_m > 103.2$ dB

$$Play = 6.4 + 0.025\Delta Z_{Lmax} - 0.041Z_m + 0.00005f_m$$

RULE 4: If $Z_m \leq 103.2$ dB

$$Play = 3$$

If more than one rule applies to a given instance, the average of the rule outputs is used as the prediction.

It is important to note that the playability models are far from definitive – indeed, there are many different

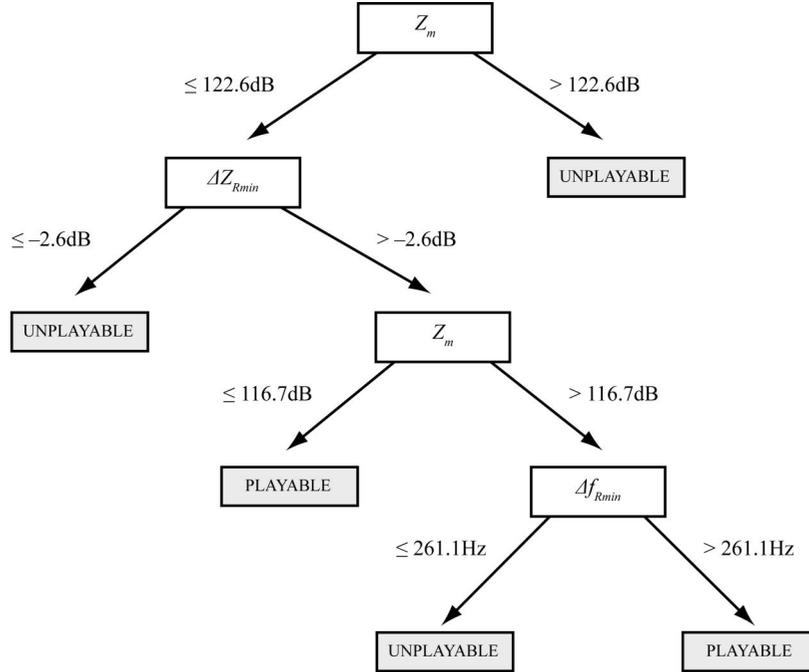


Fig. 5. The playability decision tree, constructed using C5.0.

Table 3. Cubist playability model correlation at different levels of rule coverage (10-fold cross-validation used). Selected model is highlighted.

MIN instances per rule (%)	Correlation
1	0.54
2	0.58
5	0.60
10	0.58
15	0.58
20	0.59
25	0.58
30	0.55
40	0.56

models that could have been constructed from the training data. Fundamentally, the minima parameters themselves are strongly correlated and can therefore appear interchangeably in relationships (hence, it would be misleading to regard any of the coefficients as simple weighting factors). Nevertheless, the two playability models perform reasonably well and were constructed in a robust manner. Furthermore, the data-mined rules of the models are intuitive: in general, playability is greater at lower impedance minima that are not closely neighbored by deep minima at lower or higher frequencies.

Finally, for each fingering, pairs and triplets of playable notes whose frequencies are not in a simple ratio of small integers are predicted as multiphonics.

4.3 Predicting timbre

The timbre of notes associated with minima – or, more specifically, the brightness or darkness – is predicted in a way similar to that used for playability. Our expert flutist ranked the brightness of playable minima (329 in total) into four levels, from 4 (brightest) to 1 (darkest – dark notes are characterized by a diffuse, dull quality; standard fingerings are all bright). Interestingly, the playability and brightness rankings of the 329 minima are only weakly correlated ($r = 0.12$).

As for the continuous playability scale, the brightness model was created using the Cubist model tree algorithm. Again, the Weka Experimenter demonstrated that model trees consistently outperformed the remaining three learning schemes. Table 4 shows the cross-validation correlation at different levels of rule coverage. A level of 20% was selected based on the results, giving the following brightness model (f and Δf in Hz, Z and ΔZ in dB):

RULE 1: If $\Delta f_{Lmax} > 139.2$ Hz

$$\begin{aligned}
 \text{Bright} = & 2.1 + 0.067\Delta Z_{Lmax} - 0.0153\Delta f_{Lmax} \\
 & + 0.0064\Delta f_{Rmin} - 0.054\Delta Z_{Rmax} \\
 & + 0.00026f_m - 0.0007\Delta f_{Rmax} + 0.018N_{harm}
 \end{aligned}$$

RULE 2: If $\Delta f_{Lmax} \leq 139.2$ Hz

$$\begin{aligned}
 \text{Bright} = & 0.5 + 0.00132f_m + 0.037\Delta Z_{Rmax} \\
 & + 0.023\Delta Z_{Lmax} - 0.0013\Delta f_{Rmin} \\
 & - 0.0004\Delta f_{Rmax} + 0.012N_{harm}
 \end{aligned}$$

RULE 3: If no left-neighbouring maximum exists

$$Bright = 1.3 + 0.00124f_m + 0.03\Delta Z_{R_{max}} - 0.0032\Delta f_{R_{max}} + 0.085N_{harm} + 0.0015\Delta f_{R_{min}}$$

As is intuitive, the number of harmonics is a factor in the brightness model (more harmonics give a brighter sound). We also note that a given minimum’s impedance magnitude is not directly included in the model, supporting the observation that brightness is not strongly correlated with playability (where Z_m is a strong influence). Similarly, the harmonicity of minima is not included in the model of playability.

4.4 Predicting pitch

The playing frequencies of the flute do not exactly match those of the minima in $Z(f)$, measured at room

Table 4. Cubist brightness model correlation at different levels of rule coverage (10-fold cross-validation used). Selected model is highlighted.

MIN instances per rule (%)	Correlation
1	0.65
2	0.67
5	0.68
10	0.66
15	0.67
20	0.67
25	0.65
30	0.65
40	0.56

temperature. Although the differences are a few percent or less, the effect is noticeable – up to a quarter of a semitone or so. Flutists raise the temperature and humidity of the air in the instrument, and thus raise the pitch overall. They can also vary the frequency of the resonance by varying the extent to which the lower lip covers the embouchure hole. They also vary the speed of the jet.

We choose to include all these factors (and some more subtle effects) in a single empirical function. Three experienced flutists were asked to play the same flute that was used in the experimental impedance measurements. They were asked to use their normal embouchure and to avoid correcting the pitch when and if the instrument was out of tune. Using standard fingerings, each note in the range C4 to D \sharp 7 was played four times in an ordering that decreases pitch memory. Deutsch (1978) has demonstrated that pitch memory is weakest when test tones differ by large intervals. We ordered the test notes in such a way that (i) intervals between notes were mostly large; and (ii) harmonicity between successive notes was minimized (for example, no octave intervals appear in the sequence). The sequence begins as such: D4, F \sharp 6, G \sharp 5, C4, C \sharp 5, B6, A4, D \sharp 5. . .

The pitch of each note was recorded on a tuning meter and compared to the corresponding impedance minima frequencies. The results are shown in Figure 6. The trend in pitch difference is consistent across the three flutists and agrees well with our previous study (Botros et al., 2002). We fit the results (averaged over all subjects and all four attempts per note) by a 3rd-order polynomial (f in Hz):

$$\Delta Pitch(\text{cents}) = 6.9475(\log_2 f)^3 - 197.53(\log_2 f)^2 + 1841.6(\log_2 f) - 5618.5. \tag{2}$$

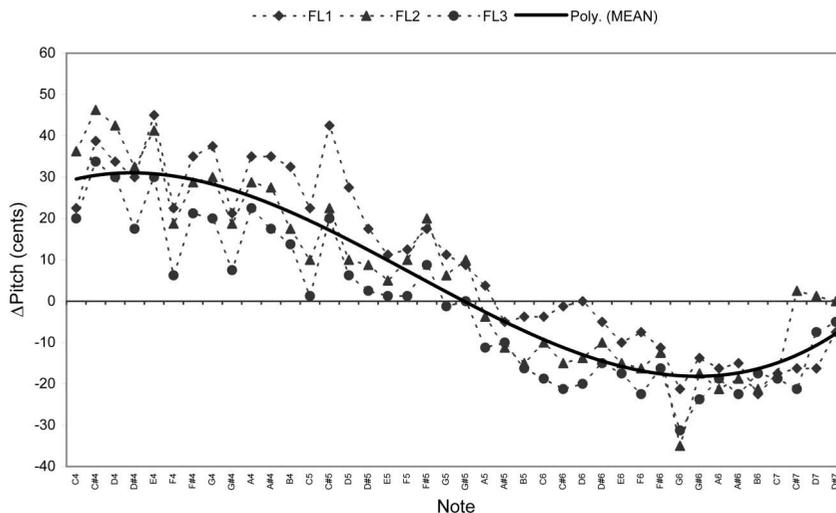


Fig. 6. Pitch difference (played frequency minus corresponding impedance minimum frequency) for three flutists. The mean difference is fit by a 3rd-order polynomial.

Different flutes and different players will give different results, and players vary frequency by more than 10 cents (a tenth of a semitone) in different circumstances and

musical contexts. Thus in general, more accurate predictions of pitch cannot be made, and so this empirical pitch correction suffices for our purposes.

© Andrew Botros 2001-2005
abotros@phys.unsw.edu.au

the virtual flute

Choose from any of the three tools below...

★ [What's new!](#)
 🗨 [Your comments](#)
 📖 [How To Guide](#)
 🏆 [Credits](#)

1. Click the keys of a fingering to search for all its predicted notes and multiphonics ... [?](#)
A note on flute models

Split E B Foot | C Foot
 Allow unconventional finger positions [What's this?](#)

2. Select a note from the list to search for an alternate fingering, trill or microtone ... [?](#)
e.g. A4, C#6. Use [standard note names](#). [F#7 and G7?](#)

SELECT NOTE

B Foot C Foot Both
 Split E No split E Both
 Allow unconventional finger positions [What's this?](#)

Fingerings MUST include the keys...

<input type="checkbox"/> BbTh or Th	<input type="checkbox"/> G	<input type="checkbox"/> D#
<input type="checkbox"/> 1	<input type="checkbox"/> F#	<input type="checkbox"/> C#
<input type="checkbox"/> A#	<input type="checkbox"/> 1	<input type="checkbox"/> C
<input type="checkbox"/> 2	<input type="checkbox"/> 2	<input type="checkbox"/> B
<input type="checkbox"/> 3	<input type="checkbox"/> 3	<input type="checkbox"/> gizmo
<input type="checkbox"/> G#	<input type="checkbox"/> tr1	
	<input type="checkbox"/> tr2	
	<input type="checkbox"/> Bblever	

OR... Include any keys

And must NOT include the keys...

<input type="checkbox"/> BbTh or Th	<input type="checkbox"/> G	<input type="checkbox"/> D#
<input type="checkbox"/> 1	<input type="checkbox"/> F#	<input type="checkbox"/> C#
<input type="checkbox"/> A#	<input type="checkbox"/> 1	<input type="checkbox"/> C
<input type="checkbox"/> 2	<input type="checkbox"/> 2	<input type="checkbox"/> B
<input type="checkbox"/> 3	<input type="checkbox"/> 3	<input type="checkbox"/> gizmo
<input type="checkbox"/> G#	<input type="checkbox"/> tr1	
	<input type="checkbox"/> tr2	
	<input type="checkbox"/> Bblever	

OR... Don't exclude any keys

3. Search for a multiphonic fingering ... [?](#)
e.g. C5&D6. Use [standard note names](#).

Select two or possibly three different notes of a multiphonic,
 or select only one note to search for multiphonics which include that note:

SELECT NOTE 1 SELECT NOTE 2 SELECT NOTE 3

B Foot C Foot Both
 Split E No split E Both
 Allow unconventional finger positions [What's this?](#)

Fig. 7. The front page of *The Virtual Flute*, showing the three available fingering search tools.

5. The Virtual Flute web service

5.1 Implementation

The *Virtual Flute* is a database-backed website that presents the machine-predicted notes and multiphonics of all flute fingerings in a manner useful to musicians. The database is preloaded with the results of a 12 h computation: for all 39,744 fingerings, (i) an impedance spectrum $Z(f)$ is calculated; (ii) minima features are extracted from $Z(f)$; (iii) playability and timbre are predicted of each minimum; (iv) pitches of playable notes are adjusted for playing conditions; and (v) multiphonics are predicted from the harmonicity of playable notes. The web service was developed following the principles of Greenspun (1999). It is located at <http://www.phys.unsw.edu.au/music/flute/virtual/>.

The *Virtual Flute* provides three tools (Figure 7). The first allows the user to enter a fingering via a graphical interface and then returns the predicted playable notes with their predicted playability, brightness and pitches. Multiphonic possibilities are also given. The second tool allows the input of a note name to search for all suitable fingerings. The third allows the user to search for multiphonics. The *Virtual Flute* caters for B foot and C foot models, with or without a “split E” mechanism.

Figure 8 shows *The Virtual Flute*’s predictions for the fingering Th 1 - 3 | 1 2 - D_♯ C. For each note listed at left, a star rating is given to indicate playability (in half-star steps, from half a star to three stars), and notes with dark timbre are indicated by the outline of a moon. Our expert flutist subjectively judged dark notes as those with $Bright < 2.75$. Multiphonic possibilities are listed at right. Multiphonics that consist of adjacent notes in the fingering’s frequency-ordered sequence are denoted with a star. We expect these multiphonics to be the most playable.

When searching for fingerings that play a given note, the user may filter the search results by (i) restricting the search to a specific flute model (e.g. C foot, split E); (ii) removing fingerings that include unconventional keys or more than one key per finger; and (iii) specifying keys that must or must not be included in the returned fingerings. The search results may be ranked by intonation (particularly useful for microtone searches), playability or darkness.

Similarly, when searching for fingerings that play a given multiphonic, the user may filter the search by items (i) and (ii) above. The user may specify a particular two-note or three-note multiphonic, or the user can provide a single note to find all multiphonics that include the given note. The latter feature is especially useful for composers. Multiphonic fingerings are ranked in a three-tiered scheme: (i) by the adjacency of the multiphonic notes in the fingering’s frequency-ordered sequence (a star denoting perfect adjacency); (ii) by the geometric mean of the individual note playabilities of the multiphonic; and (iii) by the sum of squares of the individual note pitch offsets from standard temperament (in cents).

The user can run the theoretical model for any fingering to produce $Z(f)$, calculated in real-time. The user can also post a comment on any fingering, increasing the value of *The Virtual Flute* over time.

5.2 Musical applications

Cavanagh (2005) describes a number of practical applications in performance contexts. Here we provide a few selected examples to illustrate the musical significance of *The Virtual Flute*.

Stravinsky’s *The Firebird* requires an acciaccatura from B5 to E6 to B5. The standard fingering for B5 (also used for B4) will comfortably play F_♯6, the third

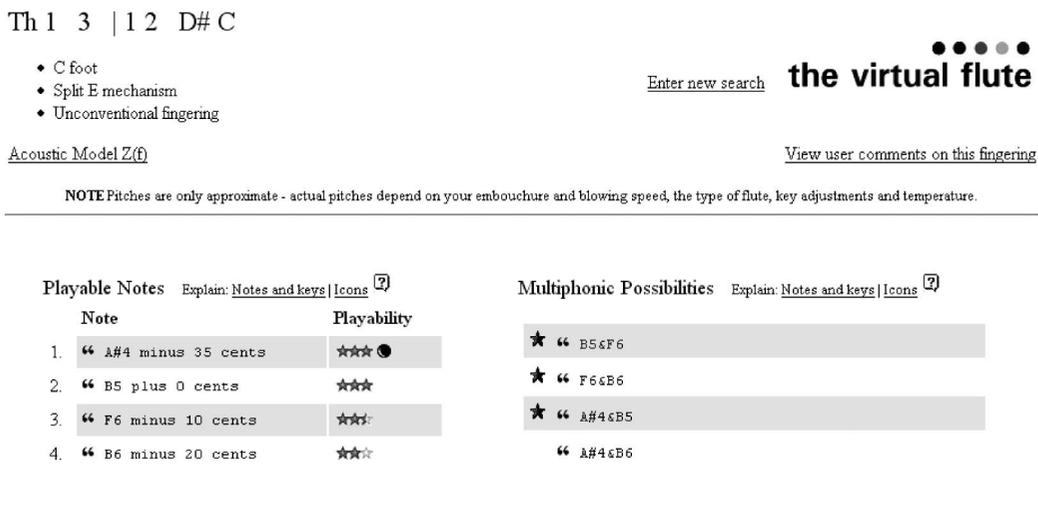


Fig. 8. Search results for the fingering Th 1 - 3 | 1 2 - D_♯ C (cf. Figure 1).

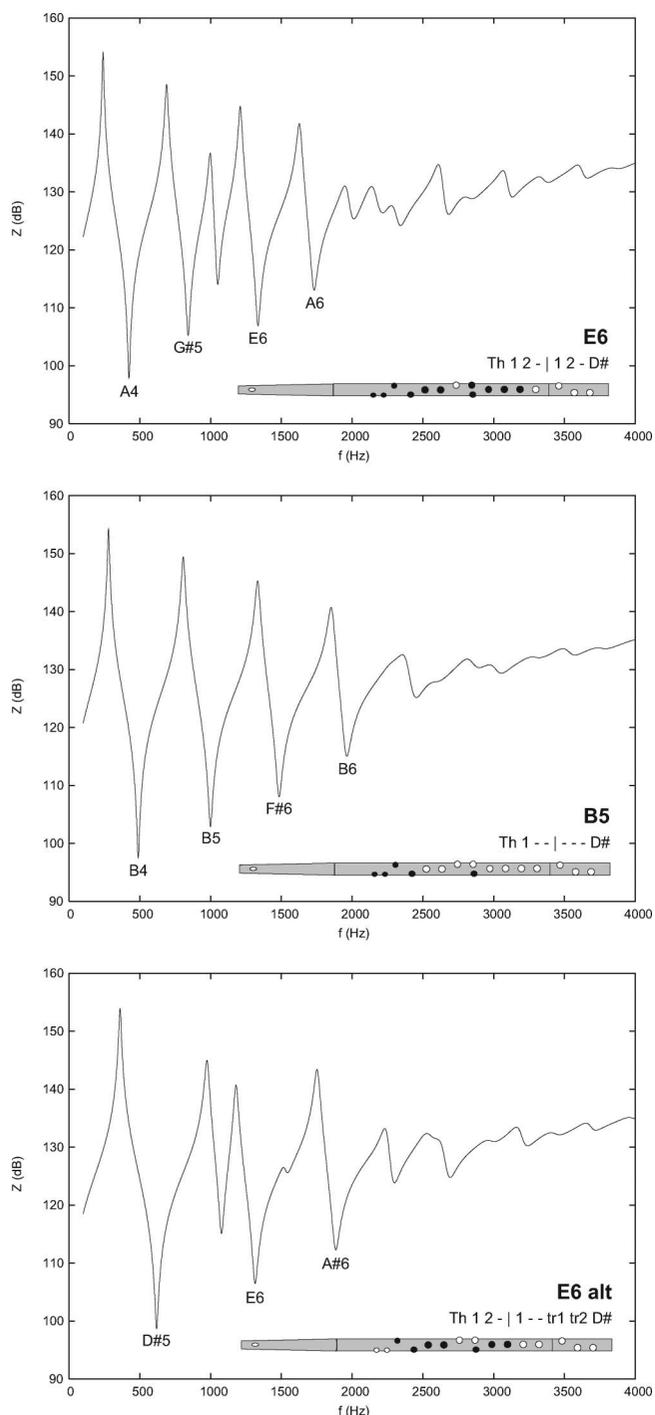


Fig. 9. Impedance model for the standard fingerings for E6 (top) and B5 (middle), and for the alternative E6 fingering (bottom). The alternative fingering has no playable minima near B5, making it an excellent B5-E6 trill.

harmonic of B4. The standard fingering for E6 will also play G \sharp 5 and A6. For a slow transition, the flutist would shorten and/or quicken the jet by just the right amount to play E6 whilst avoiding G \sharp 5 and A6. The flutist would then lengthen or slow the jet to play the B5 as before.

To play this acciaccatura (or a trill between these two notes), there is no time to adjust embouchure – one simply forms an intermediate embouchure and lets the flute fingering select the note. The danger here is that, if the embouchure compromise favours the high note, one risks playing B5-E6-F \sharp 6-B5, because the unwanted F \sharp 6 (the “split”) is only a tone above the originating E6. If the compromise is too low, one will play B5-G \sharp 5-B5, or B5-nothing-B5, because the E6 is not a particularly stable note.

Cavanagh (2005) searched for an alternative E6 fingering on *The Virtual Flute* and found Th 1 2 - | 1 - - tr1 tr2 D \sharp . (The search was structured in such a way that the returned fingerings included similar keys to the standard fingering.) The note is stable and easy to sound. Figure 9 explains how this alternative fingering proves useful. According to *The Virtual Flute* (and confirmed by players), the minimum immediately below that corresponding to E6 is unplayable, so there is little danger of the E6 “dropping down”. Consequently, the embouchure may be compromised more towards B5 and less to E6, which minimizes the chance of sounding F \sharp 6. It is a comfortable, safe solution which directly improves the quality of a performance. Although we provide an in-depth insight by analysing impedance spectra (readily provided by *The Virtual Flute*), no such expertise was required of the player. The user simply searched for E6 alternatives and experimented with the first few returned fingerings until a solution was found.

Kathleen Gallagher, who specializes in the contemporary repertoire, cites an example from Richard Barrett’s *What Remains*, composed for flute, bass clarinet and piano. The composer calls for a slurred, fast passage sharp C7, sharp B6, E7, flat D7, flat E7, D7, E \flat 7, B \flat 6. Kathleen used standard fingerings for B \flat 6, D7 and flat D7, but found the rest of the fingerings on *The Virtual Flute*, thereby creating a unique solution for a bar that has produced anxiety and performance approximations for all flutists attempting this work.

High notes are often difficult to play softly, or in tune, and flutists are keen to find improvements. Take for example F7 (the second highest note in the flute’s range): *The Virtual Flute* offers a split E fingering (1 2 - | - 2 - tr2) whose minimum is about 30% lower than that of the standard fingering. *The Virtual Flute* lists the fingering at the top of its search list when ranked by playability (Figure 10). It allows flutists of modest ability to play this very high note, and even to play it softly.

Indeed, with a database containing over 500,000 records, there’s a very strong chance that musicians will find the fingering they need in a given performance or composition. Since its launch in 2001, *The Virtual Flute* has been enthusiastically received by musicians and scientists alike. It is used hundreds of times each day, offering an easily accessible, prized service as part of our extensive Music Acoustics website.

F7 B foot search results...

· Split E mechanism.

[Enter new search](#)


NOTE Pitches are only approximate - actual pitches depend on your embouchure and blowing speed, the type of flute, key adjustments and temperature.

"Standard" Fingerings Explain: [Notes and keys](#) | [Icons](#) ^(?)

Fingering	Note	Playability	Acoustic Info	Acoustic Model
2 3 tr2 D# C#	“ 	Not predicted	Info	Z(f)

All Fingerings Explain: [Notes and keys](#) | [Icons](#) ^(?)

· Ordered by playability.

Too many fingerings? Try reducing the number by including or excluding keys in your search.

Results 1 to 50 of 54

Page ... 1 2 ▶ | ▶

Fingering	Note	Playability	Acoustic Model
1 2 2 tr2 (split E)	F7 minus 40 cents	★★★	Z(f)
1 2 2 (split E)	F7 minus 40 cents	★★★	Z(f)
1 2 2 3 tr2 (split E)	F7 minus 40 cents	★★★	Z(f)
1 2 2 tr2 C# (split E)	F7 minus 40 cents	★★★	Z(f)
1 2 2 tr2 C (split E)	F7 minus 40 cents	★★★	Z(f)
1 2 2 tr2 gizmo (split E)	F7 minus 40 cents	★★★	Z(f)
1 2 1 2 tr2 gizmo (split E)	F7 minus 40 cents	★★★	Z(f)
1 2 1 2 tr2 (split E)	F7 minus 40 cents	★★★	Z(f)
1 2 3 2 tr2	F7 minus 45 cents	★★★	Z(f)
1 2 1 2 tr2	F7 minus 45 cents	★★★	Z(f)

Fig. 10. Search results for the note F7. Fingerings are ranked by playability. The standard fingering is predicted to be unplayable.

Acknowledgements

We thank Jane Cavanagh, the flutist upon whom the expert system is based, John Tann and Paul Dickens for technical assistance, and Cécile van der Burgh and Caoimhe McMillan, our remaining experimental flutists. This research was supported by the Australian Research Council.

References

- Bartolozzi, B. (1967). *New sounds for woodwind*. London: Oxford University Press.
- Benade, A.H. (1968). On the propagation of sound waves in a cylindrical conduit. *Journal of the Acoustical Society of America*, 44(2), 616–623.
- Benade, A.H. & Murday, J.S. (1967). Measured end corrections for woodwind tone holes. *Journal of the Acoustical Society of America*, 41(6), 1609.
- Beranek, L.L. (1954). *Acoustics*. New York: McGraw-Hill, pp. 116–128.
- Berio, L. (1958). *Sequenza*. Milan: Suvini Zerboni.
- Botros, A., Smith, J. & Wolfe, J. (2002). The Virtual Boehm Flute – a web service that predicts multiphonics, microtones and alternative fingerings. *Acoustics Australia*, 30(2), 61–65.
- Botros, A., Smith, J. & Wolfe, J. (2003). The ‘Virtual Flute’: acoustic modelling at the service of players and composers. In *Proceedings of the Stockholm music acoustics conference (SMAC 03)*. Stockholm: Speech, Music and Hearing, Royal Institute of Technology, pp. 247–250.
- Caussé, R., Kergomard, J. & Lurton, X. (1984). Input impedance of brass musical instruments – comparison between experiment and numerical models. *Journal of the Acoustical Society of America*, 75(1), 241–254.
- Cavanagh, J. (2005). Find 40,000 new flute fingerings! *Pan (Journal of the British Flute Society)*, 24(2), 14–17.
- Coltman, J.W. (1979). Acoustical analysis of the Boehm flute. *Journal of the Acoustical Society of America*, 65(2), 499–506.
- Dalmont, J.-P., Nederveen, C.J. & Joly, N. (2001). Radiation impedance of tubes with different flanges: numerical and experimental investigations. *Journal of Sound and Vibration*, 244(3), 505–534.
- Dempster, A.P., Laird, N.M. & Rubin, D.B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B*, 39(1), 1–38.
- Deutsch, D. (1978). Delayed pitch comparisons and the principle of proximity. *Perception & Psychophysics*, 23(3), 227–230.

- Dick, R. (1989). *The other flute: a performance manual of contemporary techniques*, 2nd edn. New York: Multiple Breath Music.
- Fletcher, N.H. (1978). Mode locking in nonlinearly excited inharmonic musical oscillators. *Journal of the Acoustical Society of America*, 64(6), 1566–1569.
- Fletcher, N.H. & Rossing, T.D. (1998). *The physics of musical instruments*, 2nd edn. New York: Springer-Verlag, pp. 190–235.
- Greenspun, P. (1999). *Philip and Alex's guide to web publishing*. San Francisco: Morgan Kaufmann.
- Levine, H. & Schwinger, J. (1948). On the radiation of sound from an unflanged circular pipe. *Physical Review*, 73(4), 383–406.
- Nederveen, C.J. (1998). *Acoustical aspects of woodwind instruments*, 2nd edn. DeKalb: Northern Illinois University Press.
- Plitnik, G.R. & Strong, W.J. (1979). Numerical method for calculating input impedance of an oboe. *Journal of the Acoustical Society of America*, 65(3), 816–825.
- Quinlan, J.R. (1992). Learning with continuous classes. In *Proceedings of the fifth Australian joint conference on artificial intelligence (AI'92)*. Singapore: World Scientific, pp. 343–348.
- Quinlan, J.R. (1993). *C4.5: programs for machine learning*. San Mateo: Morgan Kaufmann.
- Quinlan, J.R. (2005a). *C5.0: an informal tutorial*. Rulequest Research, electronic publication available online at <http://www.rulequest.com/see5-unix.html>
- Quinlan, J.R. (2005b). *An overview of Cubist*. Rulequest Research, electronic publication available online at <http://www.rulequest.com/cubist-unix.html>
- Strong, W.J., Fletcher, N.H. & Silk, R.K. (1985). Numerical calculation of flute impedances and standing waves. *Journal of the Acoustical Society of America*, 77(6), 2166–2172.
- Witten, I.H. & Frank, E. (2000). *Data mining: practical machine learning tools and techniques with Java implementations*. San Francisco: Morgan Kaufmann.
- Wolfe, J. & Smith, J. (2003). Cut off frequencies and cross fingering in baroque, classical and modern flutes. *Journal of the Acoustical Society of America*, 114(4), 2263–2272.
- Wolfe, J., Smith, J., Tann, J. & Fletcher, N.H. (2001a). Acoustic impedance of classical and modern flutes. *Journal of Sound and Vibration*, 243(1), 127–144.
- Wolfe, J., Smith, J., Tann, J. & Fletcher, N.H. (2001b). *Acoustic impedance of classical and modern flutes: a compendium of impedance spectra, sound spectra, sounds and fingerings*. JSV+, electronic publication available online at <http://www.harcourt-international.com/journals/jsv/supplementary/suppindex.htm>.

Appendix A – Acoustic impedance model parameters

The geometries of the small sections within the model match the physical dimensions of a Pearl PF-661 B foot flute, configured in the same manner as Wolfe et al. (2001a), with the following notes:

- the length of bore sections are measured from the centre of each hole;
- the length of a tone hole is the internal height of the hole chimney;
- the tapered head joint is modelled as a single conical section with input and output radii equivalent to the head joint radii at each end;
- since the input of the instrument is at the embouchure hole, the head joint is split into two branches with lengths measured from the centre of the embouchure hole to either the cork or the flute body;
- the embouchure hole is modelled as a single conical section with input and output areas matching those of the hole on each side of the lip plate, and length matching the thickness of the plate; and
- a 5 mm long cylindrical stub is added at the input of the instrument, in accordance with the experimental setup. (The impedance of this element models the radiation impedance at the embouchure hole in a typical playing configuration.)

The adjustable parameters of the model consist of end corrections at each hole, as follows:

- $\Delta l_{\text{bore-hole}}$, an end correction to each bore section at a tone hole, expressed as a fraction of the bore radius;
- $\Delta l_{\text{bore-emb}}$, an end correction to each bore section at the embouchure hole, expressed as a fraction of the bore radius;
- $\Delta l_{\text{hole-bore}}$, an end correction to each tone hole at the bore, expressed as a fraction of the hole radius;
- $\Delta l_{\text{openhole-keypad}}$, an end correction to each open tone hole at the keypad, expressed as a fraction of the hole radius; and
- $\Delta l_{\text{emb-bore}}$, an end correction to the embouchure hole at the bore, expressed as a fraction of the embouchure radius at the bore end.

In addition to the end corrections, a multiplicative energy loss factor m_α is used for the bore within 10 mm of the embouchure hole centre (roughly a bore radius on each side). Specifically, the wave vector k can be written as the complex number:

$$k = \frac{\omega}{v} - j\alpha. \quad (3)$$

ω is the radial frequency of the wave, v is the phase velocity and α is the attenuation coefficient per unit

length of path. The present model uses Benade's (1968) formulas for v and α :

$$v \approx c \left[1 - \frac{1.65 \times 10^{-3}}{a\sqrt{f}} \right], \quad (4)$$

$$\alpha \approx \frac{3 \times 10^{-5} \sqrt{f}}{a}, \quad (5)$$

where f is the frequency of the wave within the pipe and a is the radius of the pipe. To increase the energy loss within the bore near the embouchure hole, α is multiplied by an adjustable parameter. Whilst previous models have increased attenuation factors throughout the entire flute, we choose to focus the correction at the large discontinuity between the embouchure hole and the flute bore. At a large discontinuity in an acoustical waveguide, some energy is usually lost by turbulence.

The radiation impedance of an unflanged pipe is used at the end of the flute, whereas the radiation impedance of a flanged pipe is used at open tone holes (Levine & Schwinger, 1948; Beranek, 1954). The approximation of an open tone hole by a pipe set in an infinite baffle is somewhat crude; although more involved analyses of woodwind radiation impedances have been made (Dalmont et al., 2001), we choose to absorb any differences within the selected end correction parameters.

Appendix B – The LEARNIMPEDANCE algorithm

```

SET end correction Parameter(i) := 0.0
SET  $\alpha$  correction Parameter(i) := 1.0
REPEAT
  FOREACH Parameter(i)
    SET ParameterOld(i) := Parameter(i)
  END
  FOREACH Parameter(i)
    FOR Parameter(i) := Low(Parameter(i)) TO
    HIGH(Parameter(i)) IN INCREMENT(Parameter(i))
      SET Error(Parameter(i)) = average mean
      squared error over training spectra
    END

```

```

SET ParameterError(i) := MINIMUM(Error)
SET Parameter(i) := ParameterOld(i)
END
Order parameters by increasing ParameterError(i)
FOREACH Parameter(i) in order
  FOR Parameter(i) := Low(Parameter(i)) TO
  HIGH(Parameter(i)) IN INCREMENT(Parameter(i))
    SET Error(Parameter(i)) = average mean
    squared error over training spectra
  END
  SET ParameterError(i) := MINIMUM(Error)
  SET Parameter(i) := value corresponding to
  ParameterError(i)
END
IF iteration error has not improved by more than
0.01 dB RMS
  EXIT LEARNIMPEDANCE
END
END

```

For each parameter, LEARNIMPEDANCE accepts the following input:

INC: the parameter resolution
MAXINC: the window (\pm MAXINC) in which the parameter may change in a given iteration
MIN: the absolute minimum value that a parameter may have (MIN=? if no minimum is set)
MAX: the absolute maximum value that a parameter may have (MAX=? if no maximum is set)

Thus, the FOREACH loops in LEARNIMPEDANCE are bounded as such:

Low(Parameter(i)): the higher value of (Parameter(i) – MAXINC(i)) and MIN(i)
HIGH(Parameter(i)): the lower value of (Parameter(i) + MAXINC(i)) and MAX(i)
INCREMENT(Parameter(i)): INC(i)

For the flute, end corrections have the following specification: INC = 0.01, MAXINC = 0.1, MIN = ?, MAX = ?; the α correction has the following specification: INC = 0.25, MAXINC = 2.5, MIN = 1.0, MAX = ?.