

**A Methodology for Investigation of Bowed String Performance
Through Measurement of Violin Bowing Technique**

by
Diana Young

S.M. Media Arts and Sciences
MIT, 2001

B.S. Electrical Engineering, B.A. Music
Johns Hopkins University, 1999

Performer's Certificate in Violin
The Peabody Conservatory of JHU, 1999

Submitted to the Program in Media Arts and Sciences, School of Architecture and
Planning
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy
at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2007

© Massachusetts Institute of Technology 2007. All rights reserved.

Author
Program in Media Arts and Sciences, School of Architecture and Planning
January 19, 2007

Certified by
Tod Machover
Professor of Music and Media
Thesis Supervisor

Accepted by
Andrew Lippman
Chairman, Department Committee on Graduate Students

A Methodology for Investigation of Bowed String Performance Through Measurement of Violin Bowing Technique

by

Diana Young

Submitted to the Program in Media Arts and Sciences, School of Architecture and Planning
on January 19, 2007, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

ABSTRACT

Virtuosic bowed string performance in many ways exemplifies the incredible potential of human physical performance and expression. Today, a great deal is known about the physics of the violin family and those factors responsible for its sound capabilities. However, there remains much to be discovered about the intricacies of how players control these instruments in order to achieve their characteristic range and nuance of sound. Today, technology offers the ability to study this player control under realistic, unimpeded playing conditions to lead to greater understanding of these performance skills.

Presented here is a new methodology for investigation of bowed string performance that uses a playable hardware measurement system to capture the gestures of right hand violin bowing technique. Building upon previous Hyperstring research, this measurement system was optimized to be small, lightweight, and portable and was installed on a carbon fiber violin bow and an electric violin to enable study of realistic, unencumbered violin performances. Included in the system are inertial and force sensors, and an electric field position sensor. In order to maximize the applicability of the gesture data provided by this system to related fields of interest, all of the sensors were calibrated in SI units. The gesture data captured by these sensors are recorded together with the audio data from the violin as they are produced by violinists in typical playing scenarios.

To explore the potential of the bowing measurement system created, a study of standard bowing techniques, such as *détaché*, *martelé*, and *spiccato*, was conducted with expert violinist participants. Gesture data from these trials were evaluated and input to a classifier to examine physical distinctions between bowing techniques, as well as between players. Results from this analysis, and their implications on this methodology will be presented. In addition to this examination of bowing techniques, applications of the measurement system for study of bowed string acoustics and digital music instrument performance, with focus on virtual instruments created from physical models, will be discussed.

Thesis Supervisor: Tod Machover
Title: Professor of Music and Media

A Methodology for Investigation of Bowed String Performance Through Measurement of Violin Bowing Technique

Diana Young

Submitted to the Program in Media Arts and Sciences, School of Architecture and Planning, in partial fulfillment of the requirements for the degree of Doctor of Philosophy at the Massachusetts Institute of Technology, February 2007.

Hugh Herr

Associate Professor of Media Arts and Sciences
Associate Professor of Health Sciences and Technology
MIT

Thesis Reader

Joseph Paradiso

Associate Professor of Media Arts and Sciences
MIT

Thesis Reader

Julius Orion Smith III

Professor of Music
Associate Professor of Electrical Engineering (by courtesy)
Stanford University

Thesis Reader

Acknowledgments

I would like to thank first my advisor, Tod Machover, for giving me the opportunity to study in his group and for all of his support and enthusiasm over the years.

I am also very grateful to the rest of my committee: Hugh Herr, Joe Paradiso, and Julius O. Smith III, for all of their time and guidance throughout my PhD process.

Thanks also to Barry Vercoe and Neil Gershenfeld, who have kindly shared their research facilities with me at various times, and to Knut Guettler and Sile O'Modhrain, and Stefania Serafin for their continued support, discussion, and friendship throughout the last few years.

Special thanks to all of the violinists who have made this work possible by contributing their valuable bowing data, and to those who provided great assistance and support in organizing my bowing studies: Esmé Acton and Helen Thorp of the Royal Academy of Music; and André Roy and Stephen McAdams from the Schulich School of Music of McGill University; and Knut Guettler from the Norwegian Academy of Music, and Marianne Staff.

I am very thankful to Anagha Deshmane, for her tireless work on the Bowstroke Database and numerous contributions its design; Roberto Aimi, for all of his work on the Pd recording patch used in the experiments in this thesis; George Popescu his Instron tutorial; and Todd Farrell and Marko Popovic for their assistance with the Vicon experiments.

I am very happy to be part of the Hyperbow collaboration with the Royal Academy of Music. Thanks to my colleagues there: Simon Bainbridge, Patrick Nunn, Artem Vassiliev, Philippe Kocher, Milton Mermikides, Philip Sheppard, Shu-Wei Tseng, Alex Holladay, and Peter Gregson, and to those who will join the project this year! I was also very fortunate to be a contributor to the *Toy Symphony* project. Thanks to Joshua Bell and Cora Venus Lunny, who tested the early Hyperbow with grace and patience; and to all of the members of the team.

Special thanks to Anders Askenfelt and Erwin Schoonderwaldt from KTH for their kindness and hospitality during my visit with Stefania in the summer of 2003, and for enabling early force calibration experiments with the bowing machine.

Thanks to the members of the Hyperinstruments Group (past and present), especially Roberto Aimi, Jihyon Kim, Jacqueline Karaaslanian, and Heather Childress. Also to Cornell King, Kevin Davis, John Difrancesco, Will Glesnes, and Michail Bletsas who have helped me with so many things around the lab.

Huge thanks to Linda Peterson, Pat Sokoloff, and Nikki Fortes, without whom this docu-

ment would never have been finished!

And finally, warmest thanks to my dear friends and family, especially my parents and my grandparents, and Yael, whose love and encouragement have made all the difference.

Contents

Contents	6
Figures	9
Tables	20
1 Introduction	23
1.1 Violin Performance	23
1.2 Bow-String Interaction	25
1.3 A Future Research Vision	29
1.4 Thesis Contributions	30
1.5 Thesis Outline	31
2 Toward Playable Measurement	34
2.1 Related Work	34
2.1.1 Measurement of Bowing Parameters	34
2.1.2 Bow Controllers	35
2.1.3 Performance of Bowed String Models	36
2.2 Hyperbow Performance	38
2.2.1 <i>Toy Symphony</i>	38
2.2.2 Royal Academy of Music Collaboration	40
2.3 Advancing the Hyperbow	43
3 Measurement System	45
3.1 Adding Calibration to Playability	45
3.2 Hardware Implementation	46

3.2.1	Sensing	46
3.2.2	Data Communication	54
3.2.3	Power	55
3.2.4	Ergonomics	55
3.2.5	Hardware Summary	56
3.3	Calibration and Estimation	60
3.3.1	Force Calibration	60
3.3.2	Accelerometer and Gyroscope Calibration	64
3.3.3	Position, Velocity Calibration	66
3.3.4	Kalman Filtering	69
4	Experiments	75
4.1	Preliminary Testing	76
4.2	Bowing Technique Study	77
4.2.1	Study Protocol	77
4.2.2	Experimental Setup	78
4.2.3	Participants	81
4.3	Virtual Violin Performance Study	81
4.3.1	Violin Physical Model	82
4.3.2	Study Protocol	82
4.3.3	Experimental Setup	83
5	Results	85
5.1	Technique Study Evaluation	85
5.1.1	Visual Data Evaluation	86
5.1.2	Computing the Principal Components	89
5.1.3	k-NN Classification	114
5.2	Virtual Violin Performance Evaluation	131
5.2.1	Archived Gestural Performance	131
5.2.2	Live Gestural Performance	132
6	Discussion	136
6.1	Technique Study Results	136

6.1.1	Technique Classification	137
6.1.2	Player Classification	137
6.2	Virtual Violin Study Results	139
6.2.1	Archived Gestural Performance	139
6.2.2	Live Gestural Performance	142
6.3	Sources of Error	144
6.4	Player Feedback	144
6.5	Bowstroke Database	147
6.5.1	Design Goals	149
6.5.2	Database Design	150
6.5.3	Usage to Date	154
7	Concluding Remarks	155
7.1	Contributions	155
7.2	Future Work	157
A	Violin Bowing Information	160
A.1	Bowing Terms	160
A.2	Bowing Techniques	160
B	Measurement System Schematic	162
C	Force Calibration Curves	164
C.1	Downward Force Sensor	165
C.2	Lateral Force Sensor	169
D	Bowing Technique Study Instructions	173
	Bibliography	181

List of Figures

1.1	(a) illustrates the string in Helmholtz motion; (b) describes the string velocity during the sticking and slipping intervals, from [27].	26
1.2	Schelleng diagram showing the region of playability between the minimum and maximum force required for normal violin playing with the bowing velocity kept constant.	27
1.3	Digital waveguide schematic for bowed-string instrument, from [31].	28
2.1	The Hyperbow sensing system that was developed for the <i>Toy Symphony</i> project was installed on a Jensen electric violin and a CodaBow® Conservatory™ carbon fiber violin bow. The system was optimized to be lightweight and compact and relied on minimal external electronics.	39
2.2	(a) Cora Venus Lunny performs the <i>Lullaby</i> movement of <i>Toy Symphony</i> at Media Lab Europe (Dublin, January 2002). (b) Kent Nagano conducts Deutsches Symphonie-Orchester with Cora Venus Lunny as Hyperviolin soloist (Berlin, February 2002). (c) Joshua Bell in the first complete performance of <i>Toy Symphony</i> , with the BBC Scottish Symphony Orchestra led by Gerhard Marks (Glasgow, June 2002).	39
2.3	The original Hyperbow for violin, which was used in the performances of <i>Toy Symphony</i> , was adapted for use with cello. Two of these Hyperbows for cello were built for the ongoing collaboration with the Royal Academy of Music. . .	41
2.4	(a) Royal Academy of Music masters student Shu-Wei Tseng performing Patrick Nunn’s <i>Gaia Sketches</i> on acoustic cello and Hyperbow. (b) Closeup of Hyperbow and acoustic cello. Here, the placement of the antenna for the bow position sensor has been improvised by the composers and cellists from RAM.	41

3.1 The Hyperbow sensing system that was developed for the *Toy Symphony* included electronics installed on the bow as well as an external base station. The bow electronics included two force sensors (composed of foil strain gauges mounted on the bow stick) and three dimensions (3D) of acceleration sensing. In addition, an electric field position sensor was implemented by emitting a square wave signal from the tip and the frog ends of the bow stick and detecting the amplitude of this signal from the bridge, where a receive antenna was mounted and cabled to the base station. The acceleration and force data were transmitted from the bow to the base station by means of an 900 MHz RF module, where these data were combined with the position data and connected to the serial port of a computer. 47

3.2 In typical performances, the bow and violin are both in motion, and depending on the physical styles of individual players, they may move significantly in opposition. This measurement system was designed to detect the 3D movement of both the bow and the violin, so that the bowing parameters relative to the violin could be measured. 48

3.3 The hardware for the new measurement system for violin bowing was based on the earlier Hyperbow sensing system described by Figure 3.1. Once again, two printed circuit boards (PCBs) were designed. However, several important additions were made. 3D angular velocity sensing was added to the 3D acceleration sensing on the bow by means of gyroscopes in order to create an inertial measurement unit (IMU) with six degrees of freedom (6DOF). In addition, a 6DOF IMU was also designed for use on the test violin. The electric field sensor from the Hyperbow system was expanded to include four receive antennas (one for each violin string). Finally, the wireless feature of the bow was maintained by implementing a Bluetooth® module for data communication to the violin and using a Lithium-ion battery to supply power. The violin board was powered via the USB port, which also was used for data communication (of both the bow and violin data) to the computer. 49

3.4 Each of the two force sensors used in the sensing system was comprised of four foil strain gauges, installed directly on the carbon fiber bow stick in full Wheatstone Bridge configuration. 50

3.5 The process of installing strain gauges on a test object whose surface area was small and curved, i.e. the carbon fiber violin bow, was quite difficult. Various steps had to be taken in preparation of the test surface and in the positioning of the devices. Because four gauges were used in each strain sensor and were installed within very close proximity to each other, there were further complications in applying the adhesive for a given gauge so as not to tamper with other gauges. These two uniaxial strain gauges each had nominal resistances of approximately 1000Ω and were 6.35mm long and 3.18mm wide each. The winding of the metal alloy wire that composes the gauge can be seen running back and forth in the direction parallel to the bow. The terminals of this wire are connected to solder tabs. 50

3.6 The electric field position sensor was built using a thin strip (of width a and length l) of resistive material that was adhered to the length of the bow stick, so that it would face the position antennae mounted behind the bridge of the violin during performance. A square wave signal was emitted alternatively from the tip end of the strip and the frog end of the strip. 52

3.7 For a sensor measurement of 0.3 V, the previous approximation from [57] yields a bow-bridge distance of 0.0168 m, while the current approximation from [30] yields 0.0098 m, a difference of 7 mm (0.007 m). 54

3.8 The final hardware implementation for the playable measurement system for violin bowing was installed on a Yamaha SV-200 Silent Violin and a CodaBow® Conservatory™ carbon fiber violin bow. The bow remained wireless, while the violin electronics used one USB cable for both data communication and power (as well as one audio cable not pictured). 56

3.9 One circuit board, containing one 6DOF IMU and supporting electronics for the force and position sensors, was mounted to the frog of the carbon fiber violin bow. From this view, the Bluetooth® module, microcontroller, and battery charging pads are visible. 59

3.10 The second circuit board used in the hardware implementation was mounted to the body of the Yamaha SV-200 Silent Violin, between the bridge and the tailpiece of the instrument. Similar to the circuit board on the bow, this board contains one 6DOF IMU and a Bluetooth® module. As seen here, the daughter board containing the four antennae for the electric field position sensor was placed between the two middle violin strings (A and E strings). 59

3.11 (a) In order to properly mount the bow in a vice for the Instron® force experiment, a simple device composed of LEGO® pieces was made. In addition, a device to attach to the Instron® was designed to mimic the application of force to the bow hair by a violin string. (b) This photo shows a close view of the tip of the bow just before application of force by the Instron® machine. 61

3.12 (a) This photo shows the setup for the downward force calibration experiment using the Instron® machine. (b) This photo shows the setup for the lateral force calibration experiment using the Instron® machine. 61

3.13 (a) The linear response of the downward force sensor to increasing applied force from the Instron® machine was shown at each of the thirteen test points. The plot above shows the response of the sensor at test point 9. (b) After generating calibration curves for each of the thirteen individual downward force test points along the length of the bow hair, a calibration surface was generated to illustrate the force as a function of distance from the frog. 62

3.14 (a) The linear response of the lateral force sensor to increasing applied force from the Instron® machine was shown at each of the thirteen test points. The plot above shows the response of the sensor at test point 9. (b) After generating thirteen calibration curves for each of the individual lateral force test points along the length of the bow hair, a calibration surface was generated to illustrate the force as a function of distance from the frog. 63

3.15 By placing the six inertial sensors incorporated in the bow electronics at different orientations with respect to ground, so as to maximize and minimize the effect of gravity on each, the offset for each sensor (as well as the scale of each accelerometer) was estimated. This exercise was repeated for the sensors in the violin IMU. 64

3.16	By subtracting the offsets and scaling the output as determined by the calibration exercise described above, the offsets values for each of the six inertial sensors on the bow, as well as those six on the violin were calculated. In addition, the gain of each of the accelerometers was estimated.	65
3.17	(a) Six infrared Vicon markers, visible in this photograph, were used to track the bow movement during the position calibration experiment. (b) This photograph shows a subset of the Vicon markers used to track the violin movement during the position calibration experiment. (c) A bowing guide was fashioned using LEGO® pieces in order to constrain the path of the bow during the position calibration exercise.	67
3.18	After placing the infrared markers and taking careful measurements of the location of each, models of the bow and the violin were constructed using the Vicon software.	67
3.19	Using the position data produced by the Vicon motion capture system, the data from the electric field position sensor was calibrated and then Kalman filtered. This plot shows the agreement between the two measurement systems after calibration and estimation steps.	68
4.1	This figure shows the study instructions for the <i>accented détaché</i> bowing technique, including a musical example from [10]. Participants were asked to perform each of the four lines three times (constituting one trial).	79
4.2	This figure describes the experimental setup used in the recording sessions for the bowing technique study. The top half of the figure shows the interface for the Pd recording patch, and the lower half shows the individual elements of the setup. From left to right, they are the custom violin bowing measurement system installed on a Yamaha SV-200 Silent Violin and a CodaBow® Conservatory™ violin bow; headphones; M-Audio Fast Track USB audio interface; and an Apple Macbook with a 2 GHz Intel Core Duo processor (OS X).	80

4.3 The violin physical model runs in realtime in a Pd patch. The model may be demonstrated by fixing the pitch of the output sound and manipulating, with a mouse, the sliders corresponding to the three primary bowing parameters: bow force, bow speed, and bow position (bow-bridge distance). In the virtual violin performance study conducted here, these inputs are connected to the corresponding three outputs of the measurement system’s real-time Kalman filter (running in Python) via OSC. 84

5.1 Overlay plot of the three examples of the spiccato bowing technique played on the E string by player 3. The audio waveform, as well as the data from each of the 8 bow sensors is shown. 88

5.2 Analysis of gesture data from **player 1** yielded these principal component weights for the first six components. (The weights of all of the 9152 components sum to 1.) From top to bottom: *accented détaché*, *détaché lancé*, *louré*, *martelé*, *staccato*, *spiccato*. Trials are colorized by string: G string (blue), D string (green), A string (red), E string (cyan). 91

5.3 Scatter plot of all six bowing techniques for **player 1**. *Accented détaché* (blue), *détaché lancé* (green), *louré* (red), *martelé* (cyan), *staccato* (magenta), *spiccato* (yellow). The axes correspond to the first three principal components. 92

5.4 Analysis of gesture data from **player 2** yielded these principal component weights for the first six components. (The weights of all of the 9152 components sum to 1.) From top to bottom: *accented détaché*, *détaché lancé*, *louré*, *martelé*, *staccato*, *spiccato*. Trials are colorized by string: G string (blue), D string (green), A string (red), E string (cyan). 93

5.5 Scatter plot of all six bowing techniques for **player 2**. *Accented détaché* (blue), *détaché lancé* (green), *louré* (red), *martelé* (cyan), *staccato* (magenta), *spiccato* (yellow). The axes correspond to the first three principal components. 94

5.6 Analysis of gesture data from **player 3** yielded these principal component weights for the first six components. (The weights of all of the 9152 components sum to 1.) From top to bottom: *accented détaché*, *détaché lancé*, *louré*, *martelé*, *staccato*, *spiccato*. Trials are colorized by string: G string (blue), D string (green), A string (red), E string (cyan). 95

- 5.7 Scatter plot of all six bowing techniques for **player 3**. *Accented détaché* (blue), *détaché lancé* (green), *louré* (red), *martelé* (cyan), *staccato* (magenta), *spiccato* (yellow). The axes correspond to the first three principal components. 96
- 5.8 Analysis of gesture data from **player 4** yielded these principal component weights for the first six components. (The weights of all of the 9152 components sum to 1.) From top to bottom: *accented détaché*, *détaché lancé*, *louré*, *martelé*, *staccato*, *spiccato*. Trials are colorized by string: G string (blue), D string (green), A string (red), E string (cyan). 97
- 5.9 Scatter plot of all six bowing techniques for **player 4**. *Accented détaché* (blue), *détaché lancé* (green), *louré* (red), *martelé* (cyan), *staccato* (magenta), *spiccato* (yellow). The axes correspond to the first three principal components. 98
- 5.10 Analysis of gesture data from **player 5** yielded these principal component weights for the first six components. (The weights of all of the 9152 components sum to 1.) From top to bottom: *accented détaché*, *détaché lancé*, *louré*, *martelé*, *staccato*, *spiccato*. Trials are colorized by string: G string (blue), D string (green), A string (red), E string (cyan). 99
- 5.11 Scatter plot of all six bowing techniques for **player 5**. *Accented détaché* (blue), *détaché lancé* (green), *louré* (red), *martelé* (cyan), *staccato* (magenta), *spiccato* (yellow). The axes correspond to the first three principal components. 100
- 5.12 Analysis of gesture data from **player 6** yielded these principal component weights for the first six components. (The weights of all of the 9152 components sum to 1.) From top to bottom: *accented détaché*, *détaché lancé*, *louré*, *martelé*, *staccato*, *spiccato*. Trials are colorized by string: G string (blue), D string (green), A string (red), E string (cyan). 101
- 5.13 Scatter plot of all six bowing techniques for **player 6**. *Accented détaché* (blue), *détaché lancé* (green), *louré* (red), *martelé* (cyan), *staccato* (magenta), *spiccato* (yellow). The axes correspond to the first three principal components. 102
- 5.14 Analysis of gesture data from **player 7** yielded these principal component weights for the first six components. (The weights of all of the 9152 components sum to 1.) From top to bottom: *accented détaché*, *détaché lancé*, *louré*, *martelé*, *staccato*, *spiccato*. Trials are colorized by string: G string (blue), D string (green), A string (red), E string (cyan). 103

- 5.15 Scatter plot of all six bowing techniques for **player 7**. *Accented détaché* (blue), *détaché lancé* (green), *louré* (red), *martelé* (cyan), *staccato* (magenta), *spiccato* (yellow). The axes correspond to the first three principal components. 104
- 5.16 Analysis of gesture data from **player 8** yielded these principal component weights for the first six components. (The weights of all of the 9152 components sum to 1.) From top to bottom: *accented détaché*, *détaché lancé*, *louré*, *martelé*, *staccato*, *spiccato*. Trials are colored by string: G string (blue), D string (green), A string (red), E string (cyan). 105
- 5.17 Scatter plot of all six bowing techniques for **player 8**. *Accented détaché* (blue), *détaché lancé* (green), *louré* (red), *martelé* (cyan), *staccato* (magenta), *spiccato* (yellow). The axes correspond to the first three principal components. 106
- 5.18 Top: scatter plot of *Accented détaché* for all players. Bottom: Scatter plot of *Accented détaché* for **players 2** (green), **3** (red), and **5** (magenta). The axes correspond to the first three principal components. 108
- 5.19 Top: Scatter plot of *détaché lancé* for all players. Bottom: scatter plot of *lancé détaché* for **players 2** (green), **3** (red), and **5** (magenta). The axes correspond to the first three principal components. 109
- 5.20 Top: scatter plot of *louré* for all players. Bottom: scatter plot of *louré* for **players 2** (green), **3** (red), and **5** (magenta). The axes correspond to the first three principal components. 110
- 5.21 Top: scatter plot of *martelé* for all players. Bottom: scatter plot of *martelé* for **players 2** (green), **3** (red), and **5** (magenta). The axes correspond to the first three principal components. 111
- 5.22 Top: scatter plot of *staccato* for all players. Bottom: scatter plot of *staccato* for **players 2** (green), **3** (red), and **5** (magenta). The axes correspond to the first three principal components. 112
- 5.23 Top: scatter plot of *spiccato* for all players. Bottom: scatter plot of *spiccato* for **players 2** (green), **3** (red), and **5** (magenta). The axes correspond to the first three principal components. 113

5.24 This figure shows the mean recognition rates versus the number of principal components used, produced by using the k-NN method (with one nearest neighbor) with two-thirds of each player’s gesture data as training to predict the remaining third of that player’s gesture data. 116

5.25 Mean prediction rates produced by KNN using two-thirds of the data from each of the eight players to predict the remaining one-third of all player data and increasing principal components from one to ten. 121

5.26 This bar graph shows the technique classification rates for the data contributed by each individual player. 125

5.27 Mean classification rates produced by KNN using two-thirds of the data from each of the eight players to classify the remaining one-third the data, classifying by *player*. 130

5.28 This figure shows the waveforms produced by the real electric violin and the virtual violin, and the three primary bowing parameters that produced them both. 132

5.29 This figure shows the waveforms produced by the real electric violin, and the three primary bowing parameters that were input to the model by the player when the calibration of the bow-bridge distance parameter (β) was *incorrect*, since it is calibrated in meters rather than scaled to the length of the distance between the bridge and the nut of the violin. 134

5.30 This figure shows the waveforms produced by the real electric violin, and the three primary bowing parameters that were input to the model by the player. In order to obtain a distinction between the two bowstrokes, the player made a significant pause of about one second between the downbow and the upbow. . . 135

6.1 This figure shows the audio waveforms from two different recordings provided by player 2 for the bowing technique study that were properly classified. The top plot corresponds to an example of *détaché lancé*, while the bottom plot corresponds to a *spiccato* recording. 138

6.2 This audio waveform is from a recording provided by player 6 when asked to perform the *détaché lancé* technique. The gesture corresponding to this audio was seemingly incorrectly classified as *spiccato* technique. However, as can be seen by comparing to Figure 6.1, this audio waveform bears more resemblance to a typical example of *spiccato* than to a typical example of *détaché lancé*. 138

6.3 This figure shows the effect of holding the bowing parameters constant for a duration of 1 second, at the point when the speed is zero at the moment of the change in bow direction. As can be seen in the audio waveform produced, it takes almost the entire added second for the amplitude of the sound to decrease to zero, indicating that the violin physical model does not produce realistic damping behavior. 141

6.4 This figure indicates the maximum and minimum limits for bow force, as determined by the values of bow-bridge distance and speed of the performance, given the values for μ_s , μ_d , Z_0 specific to the violin model and R calculated from data collected in an experiment by Woodhouse using a real acoustic violin [88]. 143

6.5 This figure shows the maximum and minimum limits for bow force as determined by the values of bow-bridge distance (β) and speed for the live playing case in which the β parameter was not properly calibrated. 145

6.6 This figure shows the maximum and minimum limits for bow force as determined by the values of bow-bridge distance (β) and speed for the live playing case in which the β parameter was properly calibrated. 146

6.7 Using the recently created playable measurement system developed, a large amount of audio and calibrated gesture data, as well as video, was collected in the experimental recording sessions. These data are archived in the bowstroke database. 148

6.8 This screenshot shows one of the Stroke Group pages. From this view, the audio and gesture data for this stroke may be downloaded. Also, any available related files, such as video, analysis, or synthesis files may also be downloaded or added to the My Bookshelf list. 151

6.9 This group listing is the result of a search using “sub31” (subject number 31), “detache”, “G string” as search terms. As can be seen from this screenshot, in addition to information such as user ratings and comments, the main audio file associated with each Stroke Group can be played from this view. 153

B.1 Circuit schematic for violin bowing measurement hardware. 163

C.1 Using the Instron® machine, force was applied to the bow hair at each of thirteen positions, spanning twenty-four inches, along its length. 164

C.2 These are the force calibration curves that were generated from data gathered at test points 1-4 using the Instron® machine. At each of these test points, the sensor had a good linear response. 165

C.3 These are the downward force calibration curves that were generated from data gathered at test points 5-8 using the Instron® machine. At each of these test points, the sensor had a good linear response. 166

C.4 These are the downward force calibration curves that were generated from data gathered at test points 9-12 using the Instron® machine. At each of these test points, the sensor had a good linear response. 167

C.5 These are the downward force calibration curves that were generated from data gathered at test point 13 using the Instron® machine. As shown at the other test points, the sensor had a good linear response here. 168

C.6 These are the lateral force calibration curves that were generated from data gathered at test points 1-4 using the Instron® machine. At each of these test points, the sensor had a good linear response. 169

C.7 These are the lateral force calibration curves that were generated from data gathered at test points 5-8 using the Instron® machine. At each of these test points, the sensor had a good linear response. 170

C.8 These are the lateral force calibration curves that were generated from data gathered at test points 9-12 using the Instron® machine. At each of these test points, the sensor had a good linear response. 171

C.9 These are the lateral force calibration curves that were generated from data gathered at test point 13 using the Instron® machine. As shown at the other test points, the sensor had a good linear response here. 172

List of Tables

3.1	This table indicates the dynamic range and bandwidth of the sensors used in the hardware measurement system for violin bowing.	58
3.2	At 100% duty cycle, the individual power consumptions of the key components of the bow electronics sum to approximately 160 mW.	58
3.3	Because of the specific voltage requirements of the Bluetooth® module (3.3 V \pm 0.1 V), the ADXRS300 gyroscope (5.0 V), and the importance of separating the analog and digital circuits (the other sensors were powered with 3.2 V), it was necessary to perform several power conversions in the bow electronics. The different efficiencies of these conversions are shown here.	58
5.1	Training on two-thirds of the data from each of the eight players, predicting the remaining third of each player’s data (with overall prediction of $54.9 \pm 17.1\%$) with two principal components.	117
5.2	Training on two-thirds of the data from each of the eight players, predicting the remaining third of each player’s data (with overall prediction of $69.3 \pm 11.3\%$) with three principal components.	118
5.3	Training on two-thirds of the data from each of the eight players, predicting the remain third of each player’s data (with overall prediction of $78.8 \pm 6.4\%$) with four principal components.	118
5.4	Training on two-thirds of the data from each of the eight players, predicting the remain third of each player’s data (with overall prediction of $83.3 \pm 5.2\%$) with five principal components	118

5.5	Training on two-thirds of the data from each of the eight players, predicting the remain third of each player's data (with overall prediction of $90.8 \pm 2.6\%$) with six principal components.	119
5.6	Training on two-thirds of the data from each of the eight players, predicting the remain third of each player's data (with overall prediction of $95.3 \pm 2.6\%$) with seven principal components.	119
5.7	Training on two-thirds of the data from each of the eight players, predicting the remain third of each player's data (with overall prediction of $95.8 \pm 2.0\%$) with eight principal components.	119
5.8	Training on two-thirds of the data from each of the eight players, predicting the remain third of each player's data (with overall prediction of $96.2 \pm 2.3\%$) with nine principal components.	120
5.9	Training on two-thirds of the data from each of the eight players, predicting the remain third of each player's data (with overall prediction of $95.1 \pm 3.0\%$) with ten principal components.	120
5.10	Training on players 2-8, classifying data from player 1 (overall classification rate of $91.7 \pm 20.4\%$ with ten principal components, seven nearest neighbors.	122
5.11	Training on players 1 and 3-8, classifying data from player 2 (overall classification rate of $98.6 \pm 3.4\%$) with ten principal components, seven nearest neighbors.	123
5.12	Training on players 1-2 and 4-8, classifying data from player 3 (overall classification rate of $72.2 \pm 38.6\%$) with ten principal components, seven nearest neighbors.	123
5.13	Training on players 1-3 and 5-8, classifying data from player 4 (overall classification rate of $68.1 \pm 30.0\%$) with ten principal components, seven nearest neighbors.	123
5.14	Training on players 1-4 and 6-8, classifying data from player 5 (overall prediction of $84.7 \pm 14.4\%$) with ten principal components, seven nearest neighbors.	124
5.15	Training on players 1-5 and 7-8, classifying data from player 6 (overall classification rate of $69.4 \pm 44.9\%$) with ten principal components, seven nearest neighbors.	124

5.16 Training on players 1-6 and 8, classifying data from **player 7** (overall classification rate of $75.0 \pm 35.8\%$) with ten principal components, seven nearest neighbors. 124

5.17 Training on players 1-7, classifying data from **player 8** (overall classification rate of $76.4 \pm 25.0\%$) with ten principal components, seven nearest neighbors. . 125

5.18 Individual classification rates for each technique achieved with **ten** principal components, seven nearest neighbors. An overall technique classification rate of $79.5 \pm 14.5\%$ was produced for the player-1 case. 126

5.19 Training on two-thirds of the data from each of the eight players, classifying the remaining one-third of data, classifying by player (with overall classification of $87.9 \pm 3.5\%$) with **eight** principal components. 127

5.20 Training on two-thirds of the data from each of the eight players, classifying the remaining one-third of data, classifying by player (with overall classification of $90.8 \pm 2.5\%$) with **ten** principal components. 128

5.21 Training on two-thirds of the data from each of the eight players, classifying the remain third of data, classifying by player (with overall classification of $92.0 \pm 1.8\%$) with **twelve** principal components. 128

5.22 Training on two-thirds of the data from each of the eight players, classifying the remain third of data, classifying by player (with overall classification of $92.7 \pm 2.8\%$) with **fourteen** principal components. 129

5.23 Training on two-thirds of the data from each of the eight players, classifying the remain third of data, classifying by player (with overall classification of $93.1 \pm 3.2\%$) with **sixteen** principal components. 129

5.24 Training on two-thirds of the data from each of the eight players, classifying the remain third of data, classifying by player (with overall classification of $94.1 \pm 3.1\%$) with **eighteen** principal components. 129

5.25 Training on two-thirds of the data from each of the eight players, classifying the remain third of data, classifying by player (with overall classification of $93.9 \pm 2.6\%$) with **twenty** principal components. 130

Chapter 1

Introduction

The violin is widely considered to be one of the greatest musical instruments ever developed. It is admired for its form as well as its function, and it exists as a well accepted example of the innovation in science and art. As such, the violin has great cultural meaning. The violin is capable of producing a large range of pitch and volume and is well-known for its extraordinary potential for nuanced performance. Because of these capabilities, the violin is often compared to the natural musical instrument of the human voice, which is perhaps the highest compliment possible.

1.1 Violin Performance

Just as the violin has great cultural significance, so do virtuoso violinists. Even in current times, these musicians draw a great amount of interest, as the physical facility and expressive capability they succeed in developing are truly awe-inspiring. Perhaps the most famous violin virtuoso to date was Nicolo Paganini. As described by music critic Leigh Hunt in 1831 [87],

His playing is indeed marvelous. What other players can do well, he does a hundred times better. We never heard such playing before; nor had we imagined it. His bow perfectly talks. It remonstrates, supplicates, answers, holds a dialogue. It would be the easiest thing in the world to put words to his music. We are sure that with a given subject, or even without it, Paganini's best playing could be construed into discourse by any imaginative person. [87]

The violin is an especially expressive instrument, affording its player the ability to alter the dynamic level or timbre of the notes gradually or abruptly, to play as many as four pitches at any one time, to adjust the attack and decay characteristics of each sound, vary the space and

nature of the connections between every set of notes, and a large frequency range over which to voice these notes. A violinist is granted control over every aspect of the sound produced by her instrument and throughout any playing period may recreate the character and tone continually. It is due to the possibility, the flexibility, and the versatility of the violin that the instrument is such an enigmatic and compelling vehicle of expression. As asserted by Leopold Mozart, much of the responsibility of this expression lies in bowing technique:

The present chapter will convince us entirely that the bowing gives life to the notes; that it produces now a modest, now an impertinent, now a serious or playful tone; now coaxing, or grave and sublime; now a sad or merry melody; and is therefore the medium by the reasonable use of which we are able to rouse in the hearers the aforesaid affects. [43]

It is because of this capacity for expression that it remains perhaps the most challenging of all instruments. It can be easily argued that the violin has one of the most difficult physical interfaces of any musical instrument. The violin is a small instrument and for its operation requires movements that are quite unnatural and awkward. These performance gestures range in size from extremely small to large, but it is also the number of simultaneous movements coupled with the fine precision demanded in executing them that make the violin so complicated to master. Many treatises have been written to describe the difficulties of and strategies for violin playing technique, such as [5, 22, 19].

One of the primary reasons that the violin is such a difficult instrument to master is that the playing techniques required of the left hand and the right hand are very different from each other (unlike other instruments, such as the piano, that are played with similar left and right hand techniques). The violin and the other members of the bowed string family require quite different techniques, involving control of multiple rapidly varying parameters by each hand. The left hand of a violinist not only controls pitch (by stopping the string against the surface of the violin fingerboard, effectively shortening the string length), but also can control changes in timbre by means of vibrato (small alterations in pitch). Good right hand technique requires an accomplished violinist to be proficient in a vast number of bowing methods (such as *détaché*, *martelé*, and *spiccato*) in order to obtain the range of expression demanded in standard repertoire [10, 26, 24]. The careful coordination necessary for violin playing is described by virtuoso Yehudi Menuhin in 1970 [79],

In what other art indeed does the whole body from toes to head contribute to the delivery and power, the infinite range and sensation of fingertip motion; in what

other art does the balanced body, together with the tools, become one vibrating entity? Where else indeed does one find such balanced and poised co-ordination as must needs be the violinist's? [79]

Despite the formidable obstacles imposed by the difficult playing interface of the violin, many musicians still devote themselves to its pursuit, committing years to obtain decent sound and decades to achieve mastery. As reward for this dedication, the violin offers tremendous range of tone color and expressive powers that are unsurpassed.

But violin playing technique is not just difficult to learn, it is also quite difficult to understand, even by those who have mastered it. Some violin pedagogues have done important work in documenting playing exercises, as well as postures, and physical instructions designed to improve many different techniques. However, it seems that the approach of trial and error is still the primary strategy for students of the instrument. As Leopold Auer wrote of violin bowing technique in 1921 [5],

I myself have found that there can be no exact and unalterable rule laid down indicating which one or which ones of the fingers shall in one way or another grasp and press the stick in order to secure a certain effect. Pages upon pages have been written on this question without definitely answering it. I have found it a purely individual matter, based on the physical and mental laws which it is impossible to analyze or explain mathematically. Only as the result of repeated experiment can the individual player hope to discover the best way in which to employ his fingers to obtain the desired effect. [5]

The ambiguities described by Auer, the lack of clear and precise methodologies to obtain desired results from the instrument, also likely contribute to the perception of the violin as mysterious and elusive. Of course, the difficulties of bowing technique find some explanation in the physics that govern the violin bow and strings (discussed in detail in [17]). The bow-string interaction is discussed below.

1.2 Bow-String Interaction

Why is violin playing so complicated from a physical level? The characteristic sound of a bowed string instrument is due to the phenomenon of Helmholtz motion, depicted in Figure 1.1. When Helmholtz is achieved by a player, the string forms a corner that travels in a parabolic path back and forth between the bridge and nut of the violin. When this corner is between the bridge and the bow, the string “slips” and moves in the direction opposite to that of the bow motion. When




Figure 1.1: (a) illustrates the string in Helmholtz motion; (b) describes the string velocity during the sticking and slipping intervals, from [27].

the corner is between the bridge and the nut, the string “sticks” to the bow hair and therefore moves with the same velocity as the bow. As stated by Helmholtz in [84],

During the greater part of each vibration the string here clings to the bow, and is carried on by it; then it suddenly detaches itself and rebounds, whereupon it is seized by other points in the bow and again carried forward.

This characteristic “slip-stick” behavior, in which the string slips just once in its period, occurs because of the friction component inherent in bow-string interaction. For a given bow-bridge distance, the achievement of Helmholtz motion (the goal of all bowed-string players) depends on the careful negotiation between bow speed and force. As described in [90], for steady bowing, a player must control the bow speed ν , position β (bow-bridge distance, normalized to the length between the bridge and nut of the violin), and the normal force between the bow and the string F . If F is too low, the bow will not stick to the string and will produce what is known as “surface sound”. If the force is too high, the string does not release when the Helmholtz attack occurs, and the motion becomes raucous. From these two extremes, the range of force for normal play has been analyzed as follows:

$$F_{\max} = \frac{2Z_0\nu}{\beta(\mu_s - \mu_d)} \quad (1.1)$$

$$F_{\min} = \frac{Z_0^2\nu}{2\beta^2 R(\mu_s - \mu_d)} \quad (1.2)$$

where μ_s and μ_d are the static and dynamic friction coefficients, respectively. Z is the characteristic impedance of the string and R indicates the equivalent of the rate of energy loss

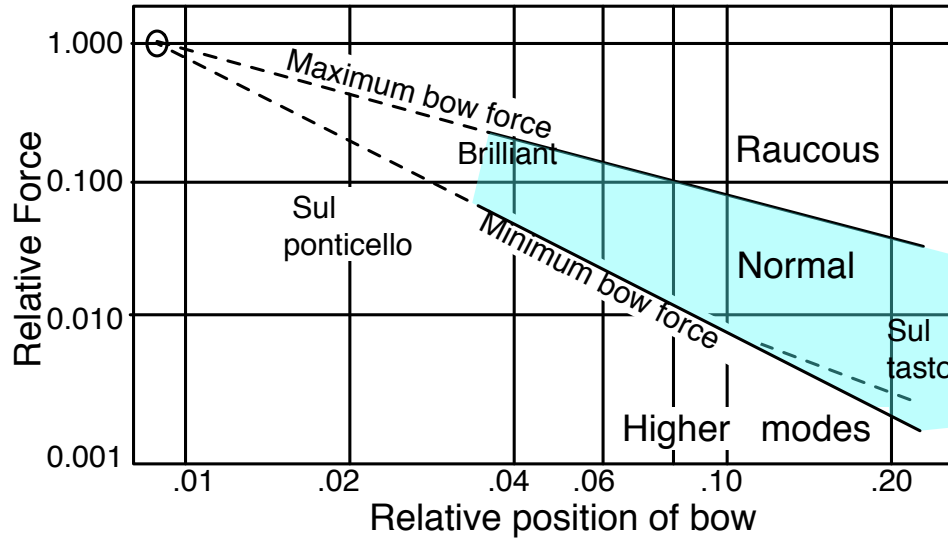


Figure 1.2: Schelleng diagram showing the region of playability between the minimum and maximum force required for normal violin playing with the bowing velocity kept constant.

into the violin body.

For constant bow speed, the dependency between the remaining two parameters, force and position (bow-bridge distance), in creating good tone may be illustrated in a two-dimensional plot. This visualization, first developed by Schelleng [63] and shown in Figure 1.2, outlines a region of “playability” wherein good tone is achieved. Note, this represents a strong idealization, as in any typical musical performance, the velocity is continually varying.

The key point of this discussion is that the relationship between bowing parameters and the sound they produce is quite complex due to the nonlinear friction mechanism described. Unlike the case of the piano playing, in which key velocity is the only physical input parameter, there is no simple mapping between inputs and output. In fact, the situation is one of a “many-to-one” mapping. Therefore, though the sound may be predicted when the bowing parameters are known, it is not possible to determine the bowing parameters used from inspection of the audio result. (In consideration of this point, audio recordings of bowed string playing can be seen as rather incomplete representations, as they do not contain all of the relevant information to reconstruct these performances.)

This point concerning the “many-to-one” mapping between the input bowing parameters and output sound can be further understood by considering a simple model of a violin (or viola, cello, or double bass). Figure 1.3 shows a digital waveguide schematic for a bowed-string instrument from [31]. In this diagram, two sets of delays are depicted, one between the bowing

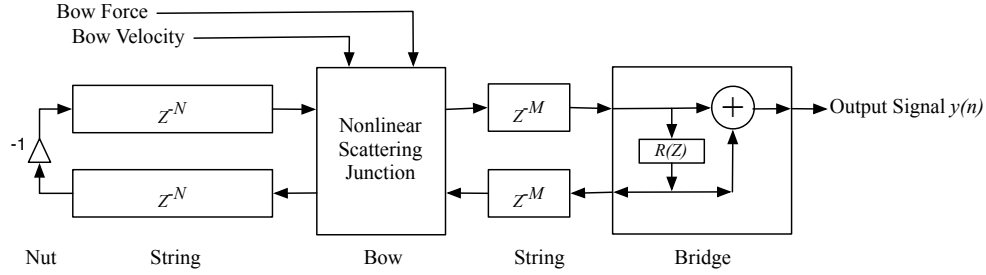


Figure 1.3: Digital waveguide schematic for bowed-string instrument, from [31].

point and the nut of the instrument, and one between the bowing point and the bridge of the instrument.

As explained in [66], when the bow applies force and velocity to the string, two traveling waves v_{on} and v_{ob} propagate between the bowing point and the nut and the bowing point and the bridge, respectively. These are represented as follows:

$$\begin{aligned} v_{on} &= v_{ib} + f/(2Z) \\ v_{ob} &= v_{in} + f/(2Z), \end{aligned} \quad (1.3)$$

where f is the frictional force and Z is the wave impedance of the string. The frictional force is represented as follows:

$$f = \mu(v - v_b)f_b, \quad (1.4)$$

where v_b is the bow velocity and f_b is the bow force, and μ is the frictional coefficient. For a given β , if $f_1 = f_2$ in two different instances of bowing, then

$$\mu_1(v - v_{b1})f_{b1} = \mu_2(v - v_{b2})f_{b2}. \quad (1.5)$$

Because there are multiple real value solutions this equation, that is, multiple values of bow force and bow velocity can be used as solutions, multiple combinations of bowing parameters can be used to achieve the same sound. The interaction is even more complicated by the additional parameter (which is not usually considered) of bow width, that is, the width of the bow hair in contact with the string. This topic was explored in depth in [59].

1.3 A Future Research Vision

Because the violin bowing parameters can not be determined from audio alone, they are the topic of keen interest in research communities. The primary interest in this research is to investigate the potential of physical measurement of violin bowing dynamics to inform understanding of bowed string performance. Central to the motivation behind this thesis are questions such as: *What do musicians do in order to achieve the sonic results they desire? How do they differ in their approaches to the same performance tasks? What are the actual limits of their physical capabilities?*

It is believed that knowledge of physical bowing gesture can enable great progress in many related fields of research. Specifically,

- **Gesture Classification.** With precise capturing of bowing parameters produced by actual players when performing different bowing techniques and musical tasks, work in the area of gesture classification of bowstrokes will be facilitated.
- **Virtual Instrument Development.** By capturing calibrated bowing parameters produced by real players in realistic playing scenarios, a new abundance of bowing gesture data will become available. These data may be used to test existing bowed string physical models and close the gap between real and virtual instrument sound quality.
- **Archiving.** The measurement of real player performance technique raises the new possibility of archiving the gesture our great performers for future generations. It is now possible to not only make audio and video recordings of the masters, but to also make complete gestural recordings as well, so that we may preserve their techniques for future study.
- **Bowed String Acoustics.** With precise and accurate measurement of real bowing parameters produced by many different players when performing different bowing techniques and musical tasks, studies of the bow-string interaction in realistic scenarios will be made possible.
- **Pedagogy.** It is also likely that new pedagogies may arise from the vast amount of player performance data, allowing more students to achieve a high level of skill.

- **Performance Practice.** Even for advanced players, more precise performance practice studies may be made possible. Practitioners will be able to precisely describe various techniques, leading to better communication of musical instructions, and historians will be better able to document practices.
- **Composition.** Information related to real player technique may also help new music composers. Currently, a common challenge faced by composers of interactive music is the availability of instrumentalists. Of course, this obstacle is common to all composers of instrumental music, as performances and rehearsals are subject to each of the player's schedules. However, in new interactive music, the challenge is greater due to the comparative lack of familiarity of the instruments involved. Often, compositions require the use of a controller or interface that has never been used before, or whose performance practice is not yet established.
- **New Instrument Mappings.** With greater understanding of the real physical mapping that exists between input (gesture) and output (sound) in a traditional acoustic instrument like the violin, researchers in the fields of new music instrument design, both acoustic and digital, may gain greater insights into new mappings. In this way, understanding how performers play well-loved instruments that have stood the test of time can help to create promising new instruments.

The primary difficulty in studying bowing parameters and their interaction in real performances is the current lack of experimental equipment capable of accurately and precisely measuring these parameters when produced by actual players. Therefore, in order to significantly advance progress in the above research areas, a new measurement system that performs this function must be created.

1.4 Thesis Contributions

For this thesis, several of the steps necessary to achieve the long-term vision described above have been undertaken. The first of these is the design and construction of an ergonomic, playable measurement system for physical violin bowing technique. A crucial component of this project is the calibration of the individual sensors included in the system as well as the inclusion of a real-time Kalman filter to improve estimation of the primary bowing parameters of most interest

in research fields related to bow-string interaction. A measurement system for violin bowing technique is designed and built to capture as precisely as possible the subtle changes in bowing gesture applied to the string by real players. In this design, two critical features are prioritized. These are *calibration* and *playability*, so that real-time bowing parameters are captured accurately and precisely, while maintaining the ergonomics and feel of the violin bow so as to allow study of natural, unencumbered bowing performances. After the construction of the measurement system, two bowing experiments are completed to determine if it is possible to sufficiently capture the salient features of violin bowing technique and examine the potential of this system as a research tool.

In the experimental stage of this thesis, the usefulness of this system for gesture classification of violin bowstrokes produced by advanced violinists is shown. In addition to this work, the applicability of the measurement system as a calibrated controller for live performance of a physical violin model is also demonstrated. The first of these experiments is a bowing technique study, in which the physical gestures used to produce various techniques are analyzed to recognize the particular techniques themselves with a standard gesture classification method. Next, an experiment is designed to explore the usefulness of this measurement system for the advancement of virtual violin research. Specifically, calibrated real player data captured by the system is used to “play” a physical violin model. Data from an archived performance, as well as data from a live real-time performance are used to drive the model, and the synthesized sound resulting from each of these two scenarios are compared.

The final main contribution of this thesis is the design and implementation of a web-accessible, searchable database containing all of the bowing data (audio, gesture, video, etc.) collected in the process of completing this thesis. This serves not only as an archive of the data, but also as a research tool to facilitate continued study of these data. Finally, all of the data collected while conducting the experiments are stored in a well-organized database in order to provide accessibility for future work by researchers interested in bowing.

1.5 Thesis Outline

1. **Introduction.** Chapter 1 introduces the motivation for this work, which is the interest in the violin as an exceptional musical instrument, as well as the common fascination with virtuosos. In addition, a future research vision based on the design of a new measurement

system for violin bowing technique is described.

2. **Toward Playable Measurement.** Chapter 2 reviews the prior art related to measurement of violin bowing parameters, as well as bow controllers for interactive music performance, with a specific focus on Hyperbow performance work.
3. **Measurement System.** Chapter 3 details the design and construction of the violin bowing measurement system, which includes force and position sensors, as well inertial sensors on both the bow and the violin. The procedures performed to calibrate the raw data from these sensors is also discussed, as is an implementation of a Kalman filter for improved estimation of the primary bowing parameters. In addition to creating a measurement system to capture the nuances of violin bowing gesture, this system is optimized to be highly playable, so that these gestures are unencumbered. The system built can be used not just as a measurement system, but also as a real-time controller. An integral part of this methodology is the manner in which all of the data are archived so as to allow further analysis and future study.
4. **Experiments.** Chapter 4 presents the experiments conducted using the new measurement system. The first is a bowing technique study, in which the audio and gesture data corresponding to performances by eight different advanced violinists of six different bowing techniques are recorded. The second experiment explores the application of the calibrated measurement system to the real-time performance of a physical violin model, in which the estimates of the three primary bowing parameters are determined in realtime (as the raw data is produced by a live player) and input to the model to achieve physical control of the synthesized sound.
5. **Experimental Results.** Chapter 5 includes the results of the bowing technique study and the virtual violin performance study. For the first study, gesture data are analyzed using standard gesture classification methods in order to discriminate the various bowing techniques. For the virtual violin study, the sound produced by the violin model in response to gesture data (force, speed, bow-bridge distance) provided by the measurement system under live control by a player is inspected.
6. **Discussion** Chapter 6 discusses the results of the previous chapter, offers further insight into the potential of the measurement system to further understanding of bowed string

performance. The bowstroke database, developed to assist this research is also described.

7. **Concluding Remarks** Chapter 7 summarizes the work described in this thesis and the contributions achieved. Future work is also discussed.

Chapter 2

Toward Playable Measurement

This thesis involves an intersection between scientific bowing parameter measurement, bow controller development, and performance of bowed-string synthesis algorithms. Related work on these topics is now reviewed.

2.1 Related Work

2.1.1 Measurement of Bowing Parameters

There is a considerable amount of bowed string acoustics research involving custom programmable bowing machines, which enable audio production from real violins generated by known bowing parameters. Such a setup allows researchers to examine the effects of different bowing parameters on violin spectrum. For instance, recent work by Paul Galluzzo includes the generation of empirical Schelleng diagrams produced using a custom programmable bowing machine [23].

Despite the obvious importance of bowing machines in studies of bowed string acoustics, measurement systems that enable studies of actual musicians' performance techniques are of great interest. Anders Askenfelt presents the first of such systems [1, 3, 2] for measuring calibrated bowing parameter data from real players. The system is installed on a traditional acoustic violin and bow. Transverse bow position is determined using a thin resistive wire inserted among the bow hairs. When the bow is in contact with the string, this wire is divided into two sections (by the string, which is resistive). These two sections are connected as one half of a Wheatstone Bridge, which is DC-driven. Velocity is derived from the bow position signal using a high-pass filter (differentiator) with a cutoff frequency of 16 Hz. The bow-bridge

measurement is performed using the same circuit as that used in the transverse bow position measurement, but it is AC-driven. To create the force sensor, the bow hair is cut from the bow and fitted with two metal strips that are attached (glued) to each end of the bow hair ribbon and the bow. Strain gauges are wired in a Wheatstone Bridge configuration and mounted on the metal strips to form the downward force sensor. Using the system described, sustained notes, scales played at different positions, and dynamic changes are studied. Askenfelt determines ranges for the three primary bowing parameters in “typical” playing conditions to be 10-50 mm, 0.2-1 m/s, and 0.5-2 N for bow-bridge distance, speed, and force, respectively. These are significant findings, though deeper investigation is certainly warranted, as these studies included only two violinist participants.

Unfortunately, the implementation of the measurement system is confined to the laboratory environment, as it is not portable. Also, because Askenfelt’s system does not measure changes in the bow angle, or tilt, which may very well impact the force sensor, this measurement method may not be complete for all playing techniques.

More recently, work has been done to examine the coordination between right hand and left hand violin technique using an infrared high-speed camera to capture finger (left hand) and bow trajectories [7]. However, this system is also confined to a laboratory environment.

2.1.2 Bow Controllers

In recent years considerable research effort has been directed toward developing new systems for electric string performance. Some of these systems augment traditional bowed string instruments to achieve new performance capabilities, such as [53, 21], and some focus on augmenting the bow itself. Several of these latter systems are of specific interest here, as they have each involved use of modern sensors to measure aspects of bowing gesture. In each, the measured parameters are extracted and used to control real-time sound synthesis or audio effects on amplified sound. Chris Chafe’s *Celletto* [14] is a custom designed electric cello supplemented with the ability to measure movements of the bowing hand. Chafe employs one of Don Buchla’s inventions, the Buchla Lightning wand, which uses IR sensing technology [62], to measure his right hand movement. This method was easily and unobtrusively implemented, as the wand housing the IR emitter is simply attached to the right wrist via a rubber band. The receiver electronics are contained within a separate box that is placed in front of the player. Because this intervention does not alter the instrument itself, the playing technique is quite unencumbered.

However, because it yields only one axis of position data and detects no force information, the system does not capture the all of the degrees of freedom of bowing motion.

The Hypercello (MIT Media Lab) [37, 38] also features a custom gestural measurement system. The component of this that detects bowing gesture includes an electric field sensor [57] to measure bow position and speed and a force sensitive resistor (FSR) to measure pressure of the right hand index finger on the bow, as a reflection of bow force.

As a component of his BoSSA (Bowed String Speaker Array) project [82], Dan Trueman created the R-bow [81]. This enhanced violin bow measures acceleration using a commercial two-axis accelerometer and, as in the Hypercello, pressure of the right hand index finger on the bow using an FSR. Another FSR, placed between the bow hair and the tip of the bow, is also used as another indicator of bow force.

In developing the eViolin, Goudeseune explored measurement of bow and violin position and speed using two commercial magnetic field sensors and a custom-designed antenna, as well as a CAVE virtual reality environment, to measure bow and violin movement [25] for applications in interactive music.

More recently, the Augmented Violin project [11, 61], modelled after the Hyperbow (discussed below) and Hypercello projects, uses a reduced sensing system that includes an accelerometer to perform gesture recognition of standard violin bowing techniques.

Though these interfaces are quite playable, they do not yield the bowing parameters of greatest interest (force, position, and speed) in calibrated S.I. units. Therefore, though useful in new music performance scenarios, they are not appropriate for use in studies requiring precise and accurate measurement of violin bowing gesture.

2.1.3 Performance of Bowed String Models

The notion of controlling a bowed string physical model with input parameters that simulate a physical gesture was first emphasized in the work of Chris Chafe [13] and Jaffe and Smith [32]. In this work, a bowed string physical model was used to produce various bowstroke techniques. Although the work illustrates the importance of physical playing techniques in the identification and characterization of musical instruments, a real-time controller was not used in these demonstrations. One of the first attempts to control a bowed string physical model using a novel controller (unlike a mouse or keyboard) was made by Claude Cadoz and colleagues (ACROE). The controller provided tactile feedback to enable the player to feel the friction of

a synthetic bowed string. However, this controller did not permit the range of violin bowing necessary for traditional technique. Recent results of this research are detailed in [20].

Experiments using a commercial Wacom tablet [85] to control a real-time waveguide bowed string model are described in [67]. Because the Wacom pen captures pressure and two axes of position information, researchers implemented a straightforward mapping of these sensor data to bow force, speed, and position. However, the dramatic difference in the ergonomics and feel of this interface in comparison to a real violin bow poses an obstacle to its use in performance scenarios.

Using the Moose, a novel haptic device built by Sile O’Modhrain and Brent Gillespie [50], experiments show that the playability of a bowed string physical model increases when tactile feedback is present [51, 49]. In order to provide tactile feedback as well as ergonomics reminiscent of a traditional violin interface, Charles Nichols presents the vBow [48]. This project incorporates encoders and motors to detect movement of a bow-like rod/string that can be played by hand and apply physical force feedback to the performer, as a violin physical model is played. Unfortunately, this interface does not permit the player to use traditional violin bowing technique, as it cannot be held as a violin due to its weight and mechanical constraints.

The Digital Stradivarius project [64] of Bernd Schoner explores the possibilities of a virtual violin not based on a physical model but on novel machine learning techniques. This project uses a bow input system similar to that featured in the earlier Hyperstring work [57] that uses an electric field sensor to measure bow position and speed, and an FSR to sense the pressure of the right hand index finger on the bow. In addition, left hand finger position is measured using a stainless steel film, installed on the fingerboard, to detect the contact point of a string when depressed by the finger. (The contact point and the two ends of the string, which is resistive, create a resistive divider used to estimate the finger position.) Using a synthesis and machine learning tool developed, “cluster-weighted synthesis”, the input parameters are mapped to an output wavetable synthesis engine. This nonlinear regression does not take into account any *a priori* models of violin physics or acoustics, but instead treats the system as a “black box”. By doing this, it cannot be easily be shown that the samples of audio data and gesture are sufficient to reproduce any violin sound, nor that the model does not introduce artifacts due to this.

2.2 Hyperbow Performance

This thesis work has been informed by several performance experiences, in which the Hyperbow controller was used by professional musicians in traditional performance scenarios. The Hyperbow was first designed and developed for use in Tod Machover's *Toy Symphony* piece. After a concert tour that included five performances of *Toy Symphony*, the Hyperbow became focus of a collaborative project with the Royal Academy of Music. The goal of this project, which is still ongoing, is to further investigate the compositional potential of the Hyperbow.

2.2.1 *Toy Symphony*

Years after the early Hyperstring research, related work was done to further develop the musical possibilities afforded by such a system. Focusing on the topic of violin performance, the Hyperbow prototype was created.

The first Hyperbow system [100], a descendent of the original MIT Hyperstring Project [38], was designed to capture elements of violin bowing gesture for use in real-time performance applications. Installed on a commercial carbon fiber violin bow and electric violin, it featured accelerometers on the frog, force sensors (composed of strain gauges) mounted on the bow stick, and an electric field position sensor that includes an antenna mounted behind the bridge of the violin. This last component of the sensor system was an adaptation of the position sensor first used in the Hypercello project [57].

The Hyperbow system, shown in Figure 2.1, was battery-powered and transmitted sensor data wirelessly via an RF communication module. An external electronics board received the data and sent it to the computer via the serial/USB port.

The Hyperbow was designed to be used in performances of Tod Machover's *Toy Symphony* piece, which included violin soloist, full orchestra, and children's singing choir, as well as a small group of children playing musical toys also developed by members of the Opera of the Future group at the Media Lab. Figure 2.2 shows the Hyperbow in performance.

This early Hyperbow work is detailed in [100, 94, 93].



Figure 2.1: The Hyperbow sensing system that was developed for the *Toy Symphony* project was installed on a Jensen electric violin and a CodaBow® Conservatory™ carbon fiber violin bow. The system was optimized to be lightweight and compact and relied on minimal external electronics.



Figure 2.2: (a) Cora Venus Lunny performs the *Lullaby* movement of *Toy Symphony* at Media Lab Europe (Dublin, January 2002). (b) Kent Nagano conducts Deutsches Symphonie-Orchester with Cora Venus Lunny as Hyperviolin soloist (Berlin, February 2002). (c) Joshua Bell in the first complete performance of *Toy Symphony*, with the BBC Scottish Symphony Orchestra led by Gerhard Marksen (Glasgow, June 2002).

2.2.2 Royal Academy of Music Collaboration

After the conclusion of the *Toy Symphony* concert tour, the Hyperbow was used in another interactive music performance in the spring of 2004. This time, the Hyperbow helped to facilitate a new interactive version of Michael Alcorn's piece, *Crossing the Threshold*, which included featured spatialized sound and graphical mappings. The work was performed by Darragh Morgan, for the opening concert of the Sonic Arts Research Centre (SARC) at Queen's University, Belfast.

Although the Hyperbow had been used in performance by several violinists by this time (who produced invaluable feedback concerning their experiences), it still could not be considered in the same class as well-established acoustic or electric instruments. These instruments not only provide wonderful new sonic possibilities to performers and composers, but also benefit from the commitment of many musicians who help to ensure their survival by creating significant repertoire to discover and showcase their capabilities. As noted by Trueman in [81], Corelli helped to establish the violin by devoting his entire compositional career to it.

In order for the Hyperbow to develop further as a performance instrument, it is essential that the Hyperbow be placed in the hands of more composers and performers. It is this belief that was at the core of the motivation for a collaborative project between the MIT Media Laboratory and the Royal Academy of Music that began in January, 2005.

This project, which is still ongoing, includes researchers from the MIT Media Laboratory, whose primary role is that of interface design, and composers and cellists from the Royal Academy of Music (RAM) [98].

Two Hyperbows for cello, one of which is shown in Figure 2.3, were built at MIT and then transferred to the RAM's permanent instrument collection. One of the goals of the project was for the artists at RAM to create, rehearse, and perform without any need for outside technology support, and to be able to develop their own individual work processes.

This collaboration began in January 2005, at which time the collaborators from MIT travelled to London for a week-long workshop. Two Hyperbows (slightly revised, as described below) were presented to the colleagues at RAM, and the week began by imparting technical knowledge of the Hyperbow system and related software such as Max/MSP.

After these introductory exercises, each composer/cellist team began the work of creating the first compositional sketches of the project. These evolved throughout the week, and on the



Figure 2.3: The original Hyperbow for violin, which was used in the performances of *Toy Symphony*, was adapted for use with cello. Two of these Hyperbows for cello were built for the ongoing collaboration with the Royal Academy of Music.

last working day we made a presentation to the RAM community.

The composers continued to independently develop their pieces through the spring of 2005, and in June 2005, the group met again, this time in Boston. After another week of concentrated work together, progress was presented in a presentation at the MIT Media Laboratory. Images of this event are included in Figure 2.4.



Figure 2.4: (a) Royal Academy of Music masters student Shu-Wei Tseng performing Patrick Nunn’s *Gaia Sketches* on acoustic cello and Hyperbow. (b) Closeup of Hyperbow and acoustic cello. Here, the placement of the antenna for the bow position sensor has been improvised by the composers and cellists from RAM.

In November 2005 at the AEC Congress 2005, the first performance of two of the works in development was given to the outside community, Patrick Nunn’s *Gaia Sketches*, and Artem Vassiliev’s *MODES*. These two pieces were performed again a month later for the third research seminar on the Hyperbow collaboration entitled “New Tools, New Uses” in December 2005.

In January 2005, the Hyperbow technology was already several years old. Although the hardware could have been significantly redesigned in preparation for this new collaboration,

this revision was postponed to a later date when the process could benefit from exposure to more composers and performers.

However, because of the different range of bowing movement required by the cello and the fact that we would be using acoustic cellos, some small but critical revisions of the existing hardware were required. Specifically, the electric field position sensor was adapted to produce a much greater amplitude of the signal emitting from the bow (so that it could be detected by the original receiver from the greater bow-bridge distance of the cello).

This project has been very fortunate to have the participation of several highly skilled, classically trained, adventurous cellists. To date, four such players, Philip Sheppard, Shu-Wei-Tseng, Alexander Holladay, and Peter Gregson, have performed the various new works with Hyperbow.

Not surprisingly, the expectations, impressions, and experiences of each individual cellist regarding the Hyperbow vary considerably. Previous experiences, such as whether or not the player has performed with an electric or amplified acoustic cello before, or the degree of familiarity with studio equipment, computers, or technology in general, are important factors. Also, the amount of time spent with the Hyperbow, in collaboration with the composers, and in rehearsal is of course critical.

Both Shu-Wei Tseng and Pater Gregson, who recently performed *Gaia Sketches* and *MODES*, respectively, took a keen interest in the technical capabilities of the Hyperbow. However, in reaction to various aspects of the experience their impressions varied. Tseng remarks, “I was immediately put into a position where I have to be fully aware of what and how I do things. This is wonderful...” Gregson observes that producing the gesture data to produce the desired sound “came down simply to feel.” Of course, their perceived ability to control the sound output of the system was also dependent on the mapping in place.

On the issue of the weight and ergonomics of the Hyperbow, feedback also differed. One player stated that the increased weight of the Hyperbow was of no concern, while another disagreed. Interestingly, criticisms regarding the carbon fiber bow itself were also expressed, as Tseng observed that it produced sound louder than usual, and Gregson noted that he found the frog to be too low. Most importantly, both cellists have expressed their interest in continued involvement with the project.

We are greatly encouraged by the progress we have achieved in the past year of our Hyperbow collaboration. The successful deployment of the Hyperbow within a new community of users, performances of new compositions, and the enthusiasm of both composers and perform-

ers for this work of art and research, are all positive results.

As we continue, we plan to further increase the number of participants involved in the project, while creating more repertoire for the Hyperbow and a greater body of knowledge regarding its related performance issues. It is our hope that through such work, we may one day be able to not only produce a more improved, refined Hyperbow, but also to help establish its performance practice.

2.3 Advancing the Hyperbow

The related research detailed above represents a great deal of experience in new bowed string performance systems and knowledge of the physical dynamics of bow-string interaction. However, there remains a noticeable lack of work in calibrated bow controllers (as well as in other controllers) for performance and measurement systems for bowing that are ergonomic and playable.

A measurement system has still not been created that: captures the relevant bowing parameters (position, speed, force normal to the string) necessary to describe the bow-string interaction; is calibrated in S.I. units; exhibits the necessary bandwidth and dynamic range to suggest it may reflect the full expressiveness of the bowing parameters; does not inhibit traditional bowing techniques; and can be used in typical performance scenarios (outside of a laboratory setting). As a result, the amount of real bowing data from actual players that can be applied in the pursuit of the research vision discussed in the previous chapter is quite limited.

Such data would greatly facilitate many areas of research related to bowed strings, as discussed in Chapter 1. In particular, the field of musical instrument synthesis, or virtual instrument development, would benefit from real player data. In this digital age, many branches of science, engineering and art are moving towards an interplay of physical and virtual. By its very nature, the virtual can transcend physical laws and the energy required to change virtual models is effortless. Physical systems are rigid, but provide tactile interfaces with which we as physical beings can interact. As with many instrument designs, people have pursued perfecting the virtual and physical representations of all instruments. Physical models of musical instruments have now achieved a level of completeness that allows them to compare favorably to their real instrument counterparts. That is, for a given set of physical input parameters, they can now produce sounds that convincingly imitate acoustic instruments. However, many of these

instruments still cannot truly substitute for traditional instruments (in compositional processes or performances), due to the lack of sophisticated physical interfaces for controlling them. In order for composers to have the full benefit of this advancing technology, they must be able to easily compose and rehearse music for virtual instruments using specially designed controllers, and players must be able to use them in live performance. In short, the physical interface, or controller, remains a crucial element of the instrument design.

There are many challenges in creating controllers to play physical models, some of which relate to aspects of new electronic instruments that are currently of strong interest, such as playability and mapping issues. When a true virtual instrument is finally created, it will be one that not only can produce the steady state sounds of the traditional instrument, but will convincingly reproduce the instrument's dynamics and transient behaviors, while still relying on traditional playing techniques.

Virtual instruments that are currently available commercially usually favor electronic keyboards communicating via MIDI as controllers for their synthesis algorithms. This paradigm poses an immediate disadvantage for non-keyboard synthetic instruments (such as strings, wind, and brass instruments), due to the inability to play the instrument using traditional techniques. In the case of strings, it is clear that such devices are not able to capture all the subtle nuances that a violin player is able to produce with right hand bowing technique, leaving much to be desired in the performance experience of these instruments.

In order to one day allow players to achieve the desired level of expressivity with a virtual bowed string instrument, it is important to match the synthesis algorithms capable of producing the sound of the instrument with a similarly sophisticated controller. Both the model and the controller must reflect the reality of physical laws. That is, they must both be carefully calibrated so that they may be well matched [75].

As part of this thesis work, a physical model incorporating state of the art research on bowed string acoustics will be tested using the bowing data collected from a previous performance by a real player. In addition, the model's behavior when driven under live control by a player will also be examined.

Chapter 3

Measurement System

3.1 Adding Calibration to Playability

In order to facilitate the experiments in this thesis, a custom measurement system was built to precisely capture the physical bow motion created by right hand violin bowing technique. To enable the measurement system to be as versatile and portable as possible, the necessary hardware for this measurement system was installed on a commercial CodaBow® Conservatory™ carbon fiber violin bow and a Yamaha Silent Violin, an electric violin, used in the experiments. So that the movement of both the bow and the violin could be tracked, individual sensor units were installed each. These sensor subsystems both include inertial sensors, as well as the necessary components for electric field position sensing of the distance between the bow and the violin. The bow sensor subsystem was also designed to include force sensors to measure the bending of the bow stick in two dimensions.

In addition to the hardware sensing system that was built, a variety of calibration exercises were also conducted in order to interpret each of the sensor outputs in S.I. units. Then, a real-time Kalman filter was implemented in order to enable good estimation of the primary violin bowing parameters, bow force (normal to the string), bow-bridge distance (β), and bow velocity, in calibrated units.

The resultant measurement system represents a large improvement on previous designs made for the Hyperbow project [100, 94, 95], because of its extended sensor system, as well as its enhancement through the calibration and estimation work accomplished. This chapter discusses the hardware implementation, calibration and estimation steps completed in the design and construction of this measurement system for violin bowing technique.

3.2 Hardware Implementation

The sensing system electronics in this system greatly resemble those in previous designs of the Hyperbow sensing system. The most recent version of the Hyperbow system is shown in Figure 2.1. The electronics used in this project are summarized in 3.1 and discussed thoroughly in [100].

In the process of designing the new measurement system, the highest priority was the goal to maintain playability of the measurement system, so that the traditional bowing technique that is of interest in this work remains unimpaired throughout the studies. Therefore, great effort was spent to ensure that the electronics created were as small and light as possible, free from unnecessary wires that would have constrained the movement of the players, and that the bow itself remains comfortable to use.

The hardware sensing system consists of four types of sensors: force sensors (composed of foil strain gauges), accelerometers, gyroscopes, and an array of four electric field position sensors. These, as well as the rest of the hardware implementation, are described below.

3.2.1 Sensing

Force Sensing

In order to sense the downward bow force, foil strain gauges with a strain limit of $\pm 1.5\%$ at room temperature from Vishay® Micro-Measurements [41] were used (as in the Hyperbow sensing system 3.1). These devices were chosen due to their large bandwidth (small capacitance and inductance) and negligible hysteresis. These features were both highly desired, as violin bowing exhibits rapid changes in force that must be accurately recorded in order to facilitate the research described in previous discussions.

Two force sensors, each composed of four foil strain gauges, in full Wheatstone Bridge configuration, were installed on the carbon fiber bow stick in order to provide “downward” (normal to the string in standard playing position) and “lateral” (orthogonal to the string in standard playing position) force measurements. The placement of these force sensors, around the middle of the bow stick, is shown in Figure 3.4.

Because foil strain gauges are designed for use primarily on flat metal structures and are generally temperature-compensated for either steel or aluminum (due to their high heat conductivity), the use of these devices on a thin, cylindrical, carbon fiber bow stick was quite

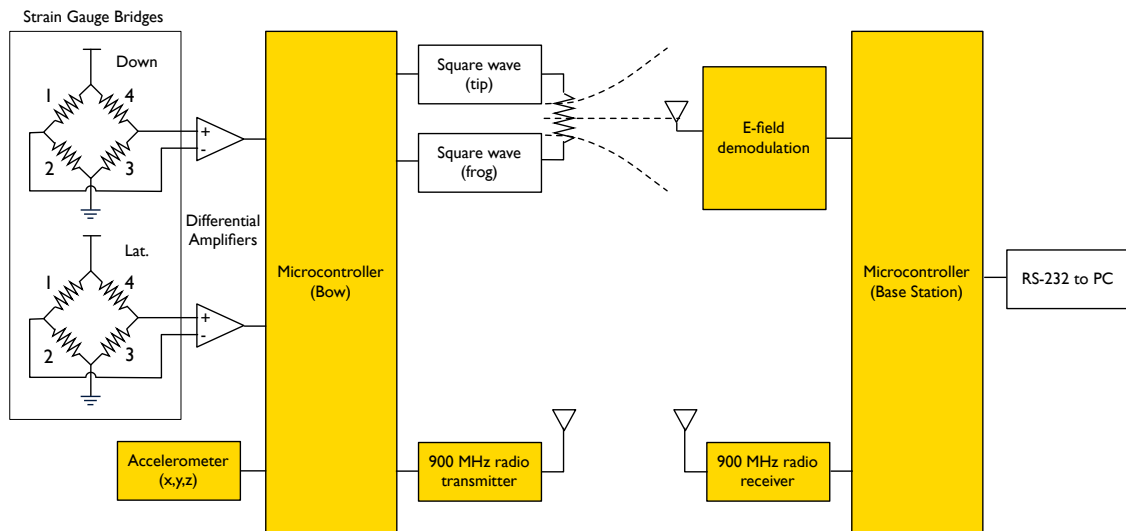


Figure 3.1: The Hyperbow sensing system that was developed for the *Toy Symphony* included electronics installed on the bow as well as an external base station. The bow electronics included two force sensors (composed of foil strain gauges mounted on the bow stick) and three dimensions (3D) of acceleration sensing. In addition, an electric field position sensor was implemented by emitting a square wave signal from the tip and the frog ends of the bow stick and detecting the amplitude of this signal from the bridge, where a receive antenna was mounted and cabled to the base station. The acceleration and force data were transmitted from the bow to the base station by means of a 900 MHz RF module, where these data were combined with the position data and connected to the serial port of a computer.




Figure 3.2: In typical performances, the bow and violin are both in motion, and depending on the physical styles of individual players, they may move significantly in opposition. This measurement system was designed to detect the 3D movement of both the bow and the violin, so that the bowing parameters relative to the violin could be measured.

nonstandard. Special care was taken in adhering the strain gauges to the round surface of the stick, both to ensure strong even bonds with the material as well as to maintain good alignment of each gauge with respect to the bow stick as well as with respect to the other gauges in the Wheatstone bridge. An example of the manner in which the gauges were installed is shown in Figure 3.5, which depicts two of the four gauges used in a lateral force sensor. After the gauges were installed and wired, they were covered with a thin layer of heat shrink tubing to protect them from damage (for instance from scraping against the violin strings during performance). The process of strain gauge installation on the bow stick and the force measurement enabled by a Wheatstone Bridge are discussed in detail in [100].

To maximize the dynamic range of the force measurement, a digital to analog converter (DAC) was used to offset any imbalance of the Wheatstone Bridge after installation.

Acceleration and Angular Velocity Sensing

In previous work, the acceleration measurements were primarily used to indicate changes in bowing direction (upbow and downbow). In this work, bow tilt was also of interest. However, three dimensions (3D) of acceleration measurement, as implemented in the Hyperbow sensing system, can only be used to estimate tilt with respect to gravity. In the design of this measure-

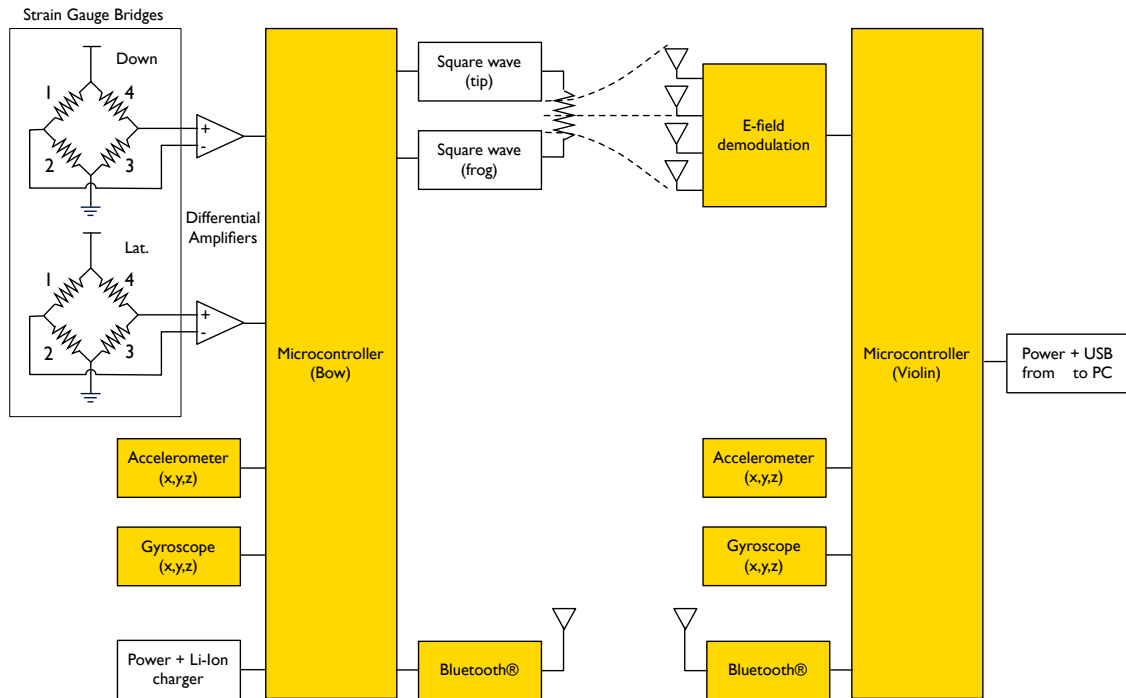


Figure 3.3: The hardware for the new measurement system for violin bowing was based on the earlier Hyperbow sensing system described by Figure 3.1. Once again, two printed circuit boards (PCBs) were designed. However, several important additions were made. 3D angular velocity sensing was added to the 3D acceleration sensing on the bow by means of gyroscopes in order to create an inertial measurement unit (IMU) with six degrees of freedom (6DOF). In addition, a 6DOF IMU was also designed for use on the test violin. The electric field sensor from the Hyperbow system was expanded to include four receive antennas (one for each violin string). Finally, the wireless feature of the bow was maintained by implementing a Bluetooth® module for data communication to the violin and using a Lithium-ion battery to supply power. The violin board was powered via the USB port, which also was used for data communication (of both the bow and violin data) to the computer.




Figure 3.4: Each of the two force sensors used in the sensing system was comprised of four foil strain gauges, installed directly on the carbon fiber bow stick in full Wheatstone Bridge configuration.




Figure 3.5: The process of installing strain gauges on a test object whose surface area was small and curved, i.e. the carbon fiber violin bow, was quite difficult. Various steps had to be taken in preparation of the test surface and in the positioning of the devices. Because four gauges were used in each strain sensor and were installed within very close proximity to each other, there were further complications in applying the adhesive for a given gauge so as not to tamper with other gauges. These two uniaxial strain gauges each had nominal resistances of approximately 1000Ω and were 6.35mm long and 3.18mm wide each. The winding of the metal alloy wire that composes the gauge can be seen running back and forth in the direction parallel to the bow. The terminals of this wire are connected to solder tabs.

ment system for violin bowing, it was important to measure bow tilt with respect to the violin (which is often related to the area of bow hair in contact with the string). This parameter could not be measured using accelerometers alone, therefore: three dimensions of angular velocity sensing were also added to the bow sensing subsystem. Together, these inertial sensors formed a six degrees of freedom (6DOF) inertial measurement unit (IMU); and another 6DOF IMU was included in the violin sensing subsystem, as seen in Figure 3.3.

The three axes of acceleration measurement (previously achieved via two ADXL202 accelerometers from Analog Devices) included in each IMU were provided by a single 3-axis

accelerometer (LIS3L06 from STMicroelectronics [76]) with a full-scale of $\pm 6g$, a bandwidth limit of 500 Hz, and a noise density of $50 \mu g/\sqrt{\text{Hz}}$. The three dimensions (3D) of angular velocity sensing on each IMU were provided using two piezoelectric vibrating gyroscopes (GYROSTAR® ENC-03M from Murata [44]) and one MEMS gyroscope (ADXRS300 from Analog Devices [18]). This configuration of gyroscopes is the same as that used in [8] and, with the choice of the LIS2L06 accelerometer, enables a completely planar implementation. Both of these types of gyroscopes are capable of sensing a maximum angular velocity of $\pm 300^\circ/s$. The former has a bandwidth limit of 1 kHz and a noise density of $0.5^\circ/s$ at 50 Hz (from [9]), while the latter has a bandwidth limit of 400 Hz and a noise density of $0.1^\circ/s/\sqrt{\text{Hz}}$.

Although the measurement system included two 6DOF IMUs to estimate the bow-bridge distance and bow velocity measurements, the errors on these two estimates increase quadratically with respect to time and linearly with respect to time, respectively. This is because these estimates are made by integrating the accelerations and angular velocities. Therefore, an additional sensor was required to provide both accurate estimates, as well as precise estimates of bow-bridge distance and bow velocity.

Position, Velocity Sensing

In order to improve the bow-bridge distance and bow velocity estimates, the original electric field bow position sensor, first designed for use in the Hypercello project and detailed in [57], and adopted in the earlier Hyperbow systems [94], was retained. The basic design of this sensor included a resistive strip extending from the frog of the bow to the tip, as shown in Figure 3.6.

The material used for this strip, which had a resistance of approximately $20k\Omega$, was a carbon-impregnated plastic from UPM [83]. From either end of this resistive strip, square wave signals were transmitted. These signals were received by an antenna mounted behind the bridge of the violin, and their corresponding magnitudes were measured to estimate bow-bridge distance and tip-frog bow position (x -axis and y -axis, respectively), as shown in Figure 3.6.

Recent improvements were made to this position sensor to increase the range (implemented for the cello Hyperbows made for the collaboration with the Royal Academy of Music described in Chapter 2). In order to decrease the attenuation of the signal from the resistive strip caused by the proximity of the player's right hand, adjustments, such as slightly decreasing the length of the strip (so that it would be further from most player's grasp) and increasing the width of the strip (to make the effective capacitive area at this end of the strip equal to that elsewhere along its length) at the frog end of the bow, were made. Also, a time domain multiple access




Figure 3.6: The electric field position sensor was built using a thin strip (of width a and length l) of resistive material that was adhered to the length of the bow stick, so that it would face the position antennae mounted behind the bridge of the violin during performance. A square wave signal was emitted alternatively from the tip end of the strip and the frog end of the strip.

(TDMA) technique was implemented to lower the power consumption of this sensor. Both of these improvements were included in this measurement system. In addition to these revisions, the bridge receive antenna was expanded to include four electrodes (instead of one as before) to enable position sensing from each of the four violin strings.

The approximation of bow-bridge distance, x , given by [57], is as follows:

$$\phi_1(x, y) + \phi_2(x, y) = B_1/x + B_2, \quad (3.1)$$

where ϕ_1 is the potential from one end of the strip, and ϕ_2 is the potential from the other end. Recently, a full analytical solution for the potential in space around a resistive strip has been provided in [30] and is presented below.

In order to make position estimates from the magnitudes of the received signals, the following analysis of the resistive strip as a transmitter was employed. According to [30], the potential (ϕ) in space around the strip is given by the following equation:

$$\phi(x, y) = (A_1 \tanh^{-1} \sqrt{\frac{x^2}{x^2 + a^2}} - A_2)(A_3 y - A_4) \quad (3.2)$$

in the $z = 0$ plane, where a is the width of the short dimension of the strip.

The scheme, similar to that used in [56], adopted for position sensor was as follows. Using the TDMA protocol, the potential of one end of the strip is raised to V , while the other end is grounded, and then the process is reversed and repeated so that the signal is emitted from both the tip and the frog ends of the bow. The potential field in space is inverted by changing the sign

on y ,

$$\phi_1(x, y) = (A_1 \tanh^{-1} \sqrt{\frac{x^2}{x^2 + a^2}} - A_2)(A_3 y - A_4) \quad (3.3)$$

$$\phi_2(x, y) = \phi_1(x, -y). \quad (3.4)$$

Therefore, the current approximation for bow-bridge distance, x , is as follows:

$$\phi_1(x, y) + \phi_2(x, y) = -A_1 A_4 \coth \sqrt{\frac{x^2}{x^2 + a^2}} + A_2 A_4. \quad (3.5)$$

As $x \rightarrow \infty$, $\phi_1(x, y) = \phi_2(x, y) \rightarrow 0$. Therefore $A_2/A_1 = \coth(1) \approx 1.3104$. If $\alpha = \frac{\phi_1(x, y) + \phi_2(x, y)}{-A_1 A_4}$, then x can be solved directly as:

$$x = \frac{i \sqrt{a \coth(\coth(1) + \alpha)}}{\sqrt{\frac{-1}{a^2} + \frac{1}{a^2} a \coth(\coth(1) + \alpha)}}. \quad (3.6)$$

Now, to solve for y :

$$\phi_1(x, y) - \phi_2(x, y) = 2A_3 y (A_1 \coth \sqrt{\frac{x^2}{x^2 + a^2}} - A_2) \Rightarrow \quad (3.7)$$

$$y = \frac{\phi_1(x, y) - \phi_2(x, y)}{2A_3 A_1 (\coth \sqrt{\frac{x^2}{x^2 + a^2}} - \coth(1))} \quad (3.8)$$

$$y = \frac{\phi_1(x, y) - \phi_2(x, y)}{-2A_3 A_4 (\phi_1(x, y) + \phi_2(x, y))}, \quad (3.9)$$

which is a linear function of $\phi_{1,2}(x, y)$. This equation is only an approximation, as it assumes that the bow is always parallel to the bridge antennae and that position $z = 0$.

The current approximation [30], given by Equation 3.5 (after substituting for A_2/A_1) is plotted in Figure 3.7, as is the previous approximation [57], given by Equation 3.1.

In practice, the potential discussed above was created by a 100 kHz square wave signal. The frequency of this signal was chosen to allow coupling between the strip and the position receiver (behind the bridge of the violin) and to facilitate easy implementation with a microcontroller timer unit. The position receive circuit contained on the violin electronics board was a gain filter stage tuned to the 100 kHz electric potential from the bow, followed by a peak detector. The low pass filter of the detector had a notch filter for 60 Hz noise, though it was later determined to be unnecessary. Therefore, no grounding cable was needed. This feature was of great benefit




Figure 3.7: For a sensor measurement of 0.3 V, the previous approximation from [57] yields a bow-bridge distance of 0.0168 m, while the current approximation from [30] yields 0.0098 m, a difference of 7 mm (0.007 m).

to the ergonomics of the system.

3.2.2 Data Communication

Two microcontrollers from Silicon Laboratories [35] were used in the measurement system, one for the bow electronics, a C8051F041, and one for the violin electronics, a C8051F320. The former was chosen because of its built-in eight-port 12-bit analog to digital converter (ADC), as well as its two-port DAC. The latter was selected because of its built-in eight-port 10-bit ADC and the ease with which it interfaces to USB.

The data communication process that was implemented is as follows. The bow microcontroller receives sensor data from each of the eight bow sensors (three accelerometers, three gyroscopes, and two force sensors) via its ADC. These data are sent to the violin board via a Bluetooth® class 2 module (the WT12, from Bluegiga Technologies [80]). In addition, the bow board also communicates the status (on/off) of the tip and frog position signals.

The violin microcontroller receives sensor data from each of the six inertial violin sensors (three accelerometers, three gyroscopes) and two position data measurements of the tip and the

frog signal strengths via its built-in 10-bit ADC. Because the two position measurements are the result of multiplexing, the received signals from each of the four position antennas, these data are de-multiplexed at this stage. The violin microcontroller combines the received bow data channels with the violin data channels and outputs all of the sensor data channels to the USB port of the computer. Therefore, the computer then receives twenty-two sensor data channels from the measurement system in one single data stream.

3.2.3 Power

The bow electronics were powered by a lithium polymer single cell rechargeable battery with a nominal voltage output of 3.7 V and a capacity of 145 mAh (the KOK145T from Kokam™ [33]). This battery was chosen because of its small package (27.5 mm x 20.5 mm x 4.5 mm) and light weight (3.5 g). An integrated switching supply and charger circuit was included. In order to save valuable space and weight on the bow PCB, a charging connector that consisted of only two charging pads and a mechanical cutout was designed to accommodate a Sony Ericsson mobile phone charging cable/connector.

The violin electronics were powered through the USB cable connected to a computer. Details of the power consumption of this system are included in the following summary.

3.2.4 Ergonomics

As emphasized earlier, a primary concern in the design of this system was to minimize the ergonomic impact of the electronics on the original violin and bow.

The most critical ergonomic issues were, of course, the weight and balance point of the violin bow. Despite the addition of the gyroscopes and the charging circuitry to the sensing system, the weight and balance point of the bow were maintained at approximately the same values as achieved with the previous system [100], in large part due to the reduction in the weight of the battery. Specifically, the weight of the bow was changed from its original weight of 61 g to 84 g, and the balance point was shifted from 21 cm (from the end of the bow stick) to 26 cm.

The test violin selected, the Yamaha SV-200 Silent Violin, provided both the convenience of an electric violin and the ergonomics of an acoustic violin. Though it does not have a full body, it does maintain the most relevant features of the acoustic violin's form. Specifically, it features a realistic outer frame that serves as a critical guide to players for left hand technique

and allows the use of a traditional chin rest and shoulder rest. In fact, these two accessories may be easily changed, as is the case with an acoustic instrument, so that a player may use his or her own personal device for the most natural and familiar ergonomics.

The final hardware implementation on a Yamaha SV-200 Silent Violin and a CodaBow® Conservatory™ is shown in Figure 3.8.



Figure 3.8: The final hardware implementation for the playable measurement system for violin bowing was installed on a Yamaha SV-200 Silent Violin and a CodaBow® Conservatory™ carbon fiber violin bow. The bow remained wireless, while the violin electronics used one USB cable for both data communication and power (as well as one audio cable not pictured).

3.2.5 Hardware Summary

The bow electronics remained battery-powered and continued to feature a wireless data channel. Also as before, the violin electronics maintained wired power and data communication. (As the violin already requires an audio cable, it is already necessarily wired.) Figure 3.3 illustrates the schematic designs of the hardware sensing system, highlighting those features that represent the new additions described.

One of the focal points of this measurement system was to make sure the dynamic range and bandwidth of the sensors and digital acquisition system are as high as possible given current devices and hopefully exceeding the Nyquist limit for a player. In order to achieve sufficient

dynamic range, all of the sensors were designed to easily accommodate the typical ranges of the bowing parameters previously determined and discussed in the previous chapter [2, 28].

On the issue of bandwidth, it was important that the measurement system developed here have the ability to capture tremolo (rapid repetition of unaccentuated notes). Recent experiments examining typical speeds of tremolo (as well as trill and vibrato) performed by professional bowed string players, including violinists, indicate that typical timings are well under 20 Hz [42]. Typical bouncing rates for a violin bow in contact with the string (as in *spiccato* or *ricochet* techniques), controlled by a machine not a player, have been measured to be between 6 and 40 Hz [4].

In addition, the ultimate bandwidth of a measurement system such as that developed here should be sufficient not only to capture typical human performances, but also the limits. Therefore, other types of human performance were considered. According to [40], on the topic of Irish tap dance, foot-tapping can reach a rate as high as 28 Hz, while hand-tapping has been measured at a frequency of 20 Hz [65]. Other work indicates that skilled drummers may produce drum-rolls at over 30 Hz (by means of passive “double-sticking”) [29, 86]. Studies of dart-throwing yield measurement of timing precision that can be faster than 1.8 ms [72]. As discussed by [55], trilling speed of concert pianists can be on the order of 16 Hz.

Though these figures are quite manageable with current electronics, it should be noted that these values do not represent limits for continuous spatial movement, which are clearly higher. Furthermore, gesture dynamics can be influenced by resonances created by mechanical contact. Such mechanical resonances can be significantly higher than human performance limits, as in the case of the bouncing bow [4], in which a 150 Hz resonance is observed in the measurement of bowing parameters.

Based on these findings, participants studied in this thesis were not expected to generate gesture data requiring a Nyquist sampling rate above 200 Hz. This digital acquisition system was designed to have 73 dB of dynamic range and a sampling rate of 2.4 kHz for each of the input channels, improving upon the figures of 60 dB and 120 Hz for the Hyperbow. Table 3.1 shows the dynamic range and bandwidth of each sensor that was included in this measurement system.

The power consumption measured for the bow electronics was between 40 mA and 50 mA. Therefore, the 3.5 g rechargeable battery with 145 mAh capacity at 3.7 V provided ≈ 3 hours of

sensor type	dynamic range	dynamic range (practice)	bandwidth
accelerometer (3 axis)	73 dB	68 dB	100 Hz
gyroscope (3 axis)	73 (x,y), 57 (z) dB	68 (x,y), 57 (z) dB	50 (x,y), 100 (z) Hz
bow force (2 axis)	N/A	68 dB	100 Hz
position (2 axis)	N/A	50 dB	25 Hz

Table 3.1: This table indicates the dynamic range and bandwidth of the sensors used in the hardware measurement system for violin bowing.

continuous use. The violin electronics were powered through a USB port, capable of supplying 100 mA at 5.0 V, which was more than sufficient for this application. The power consumption of each of the key components used in the measurement system and the power conversion stages of the bow electronics are detailed in Tables 3.2 and 3.3.

Chip	Power Consumption	Bits or dB	Range	BW
Analog Devices ADXRS300 (gyro)	6 mA @ 5.0V = 30 mW	9/57	$\pm 300^\circ/\text{sec}$	100 Hz
Murata ENC-03M (gyro) $\times 2$	4.5 mA @ 3.2V = 14 mW	11/74	$\pm 300^\circ/\text{sec}$	50 Hz
ST Micro LISL02AL (accelerometer)	0.95 mA @ 3.2V = 3.0 mW	11/74	± 6 g	100 Hz
Strain sensor $\times 2$	3.2V over $1\text{k}\Omega = 10$ mW (2/9 duty cycle)	11/74	0-10 N	100 Hz
Position sensor (TDMA)	5V over $\approx 30\text{k}\Omega = 0.83$ mW	$\approx 8/50$	0-0.7 m, 0-0.1 m	25 Hz
Bluegiga WT12 (blue-tooth radio)	22.6 mA @ 3.3V = 96.4 mW	n/a	n/a	n/a
C8051F041 (micro)	2 mA @ 3.3V = 6.6 mW	n/a	n/a	n/a

Table 3.2: At 100% duty cycle, the individual power consumptions of the key components of the bow electronics sum to approximately 160 mW.

Type	Efficiency	Current Output (mA)
3.7 Battery to 3.3V	89%	49.8
3.3 to 5.0V Analog	72%	8.56
3.3 to 3.2V Analog	97%	12.6

Table 3.3: Because of the specific voltage requirements of the Bluetooth® module ($3.3\text{ V} \pm 0.1\text{ V}$), the ADXRS300 gyroscope (5.0 V), and the importance of separating the analog and digital circuits (the other sensors were powered with 3.2 V), it was necessary to perform several power conversions in the bow electronics. The different efficiencies of these conversions are shown here.

Labelled images of the electronics mounted on the bow and the violin are shown in Figures 3.9 and 3.10.




Figure 3.9: One circuit board, containing one 6DOF IMU and supporting electronics for the force and position sensors, was mounted to the frog of the carbon fiber violin bow. From this view, the Bluetooth® module, microcontroller, and battery charging pads are visible.




Figure 3.10: The second circuit board used in the hardware implementation was mounted to the body of the Yamaha SV-200 Silent Violin, between the bridge and the tailpiece of the instrument. Similar to the circuit board on the bow, this board contains one 6DOF IMU and a Bluetooth® module. As seen here, the daughter board containing the four antennae for the electric field position sensor was placed between the two middle violin strings (A and E strings).

3.3 Calibration and Estimation

3.3.1 Force Calibration

The bow force data were calibrated (as in previous investigations) using the Instron® 1122 Universal Materials Testing Machine. Because the bow force sensor outputs are dependent on bow position, force was applied to the bow hair by the Instron® at a series of test points between the tip and the frog of the bow. Specifically, thirteen points were measured and marked out along the length of the bow stick at which to apply force from the Instron® machine.

The Instron® was controlled using Labview to apply force at a speed of (100 mm/min) up to a maximum limit and then return to its original rest position. During the calibration experiment, the bow force sensor data was recorded by the measurement system as was the Instron® load cell data. Measurements from the Instron® sensor were recorded at a sampling rate of 20Hz.

Downward Force

In order to mount the bow securely in position so that the Instron® could apply force to the bow hair, a mount was quickly built using LEGO® pieces. This assembly was designed to allow room for the electronics board and included foam to cushion the bow against the mount and provide for a more secure fit. Cable ties were used to fix the bow in place.

A second LEGO® assembly was produced to attach to the Instron® itself in order to mimic the physical contact between the bow hair and a violin string. This piece was made with a pointed edge so that it could be pressed onto the bow hair (orthogonal to the hair) in a similar manner to the way in which the bow pressed onto a violin string in normal violin playing. This, as well as the bow mount just described, are shown in Figure 3.11(a).

The bow was placed in the mount, which was then placed in a vice so that the bow extended out parallel to the ground with the bow hair pointing up, as seen in Figure 3.12(a). A closer view can be seen in Figure 3.11(b).

At each test point, a maximum force value was chosen so as to saturate the downward force sensor, though care was taken to ensure that no harm would come to the bow.

After the calibration experiment, the data collected was analyzed as follows. Using a matlab script, the data from the bow force sensor were resampled to be equal to the sampling rate of the Instron® load cell data, and then the data from each of the sensors (Instron® load cell and bow force sensor) were plotted. The script then allowed the beginning and end of the downward sloping portion of each curve to be selected by hand. Finally, the bow data from its selected




Figure 3.11: (a) In order to properly mount the bow in a vice for the Instron® force experiment, a simple device composed of LEGO® pieces was made. In addition, a device to attach to the Instron® was designed to mimic the application of force to the bow hair by a violin string. (b) This photo shows a close view of the tip of the bow just before application of force by the Instron® machine.


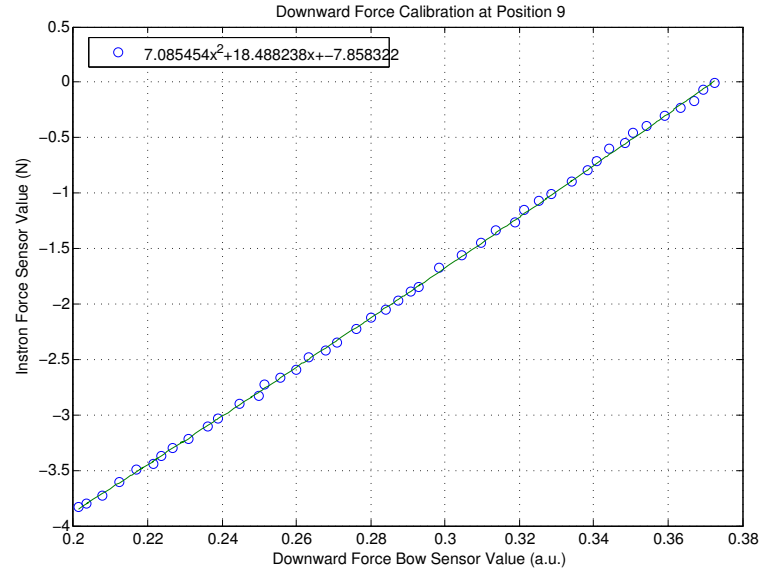



Figure 3.12: (a) This photo shows the setup for the downward force calibration experiment using the Instron® machine. (b) This photo shows the setup for the lateral force calibration experiment using the Instron® machine.

segment were plotted against the Instron® data from its selected segment and then fit using a second-order polynomial. The results of this analysis can be seen in Appendix C.

Each one of the thirteen resulting line fits was then stacked to create a 3D plot showing the dependence of the bow force sensor on force (N) and distance from the frog of the bow (m). This calibration surface can be seen in Figure 3.13(b).



(a)



(b)

Figure 3.13: (a) The linear response of the downward force sensor to increasing applied force from the Instron[®] machine was shown at each of the thirteen test points. The plot above shows the response of the sensor at test point 9. (b) After generating calibration curves for each of the thirteen individual downward force test points along the length of the bow hair, a calibration surface was generated to illustrate the force as a function of distance from the frog.

Lateral Force

The LEGO[®] assembly used to mount the bow for the downward force calibration experiment was revised to enable mounting ninety degrees from the previous orientation so that force could

be applied to the edge of the bow hair as seen in Figure 3.12(b).

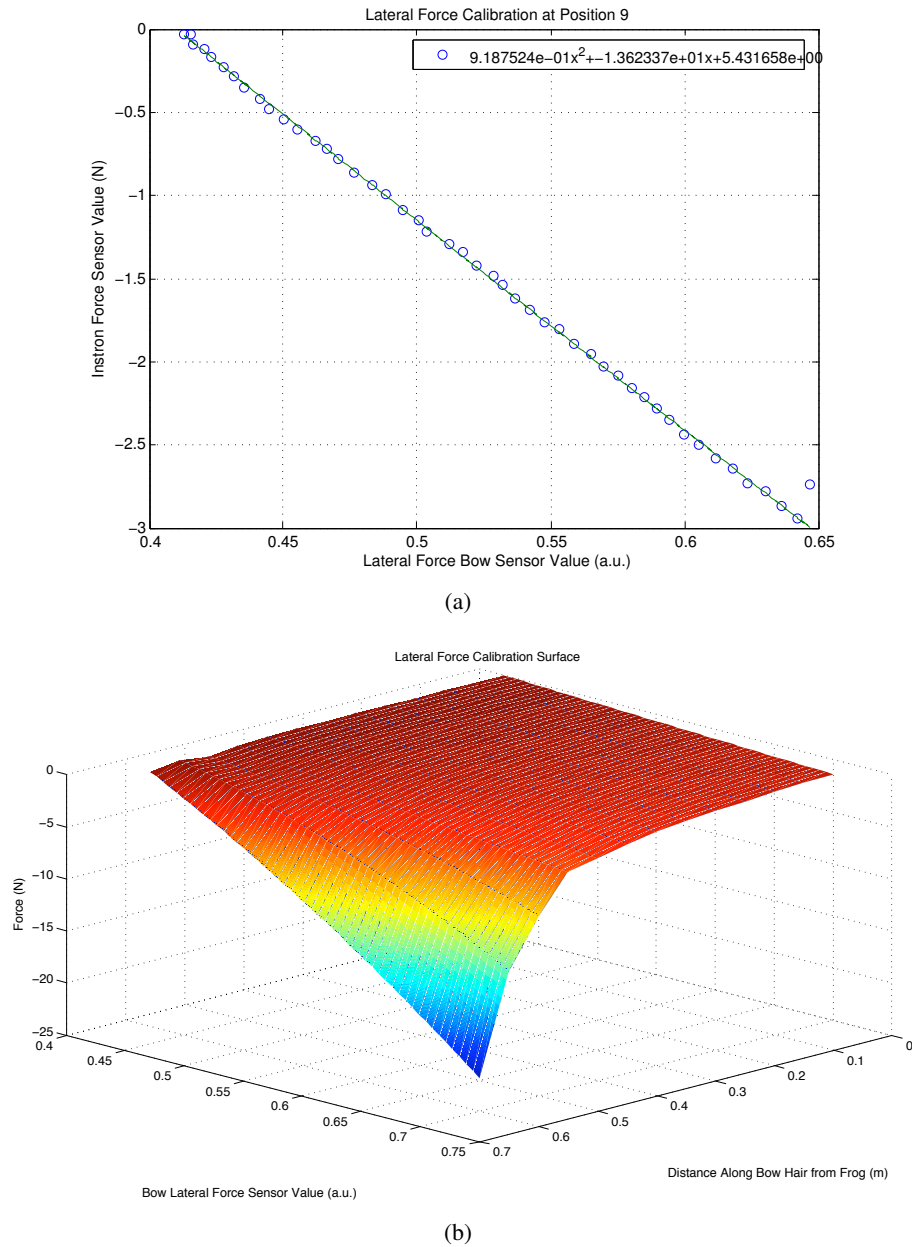


Figure 3.14: (a) The linear response of the lateral force sensor to increasing applied force from the Instron® machine was shown at each of the thirteen test points. The plot above shows the response of the sensor at test point 9. (b) After generating thirteen calibration curves for each of the individual lateral force test points along the length of the bow hair, a calibration surface was generated to illustrate the force as a function of distance from the frog.

3.3.2 Accelerometer and Gyroscope Calibration

In order to calibrate the inertial sensors on the bow and the violin, each inertial unit was fixed in six positions as depicted by Figure 3.15. In this process, each accelerometer and gyroscope pair (corresponding to each axis of measurement) was placed at rest orthogonal to ground (parallel to force of gravity) in two orientations, first maximizing and then minimizing the effect of gravity (pointed down and then pointed up). The output of each relevant sensor was recorded for at least one second to obtain a sufficient average. The output value when each accelerometer was in the downward position was taken as the value equal to $-1g$, and the output value when each accelerometer was in the upward position was taken as the value equal to $+1g$. In this way, both offset and gain were estimated for each of the accelerometers. The calibration data for the inertial sensors can be seen in Figure 3.16.




Figure 3.15: By placing the six inertial sensors incorporated in the bow electronics at different orientations with respect to ground, so as to maximize and minimize the effect of gravity on each, the offset for each sensor (as well as the scale of each accelerometer) was estimated. This exercise was repeated for the sensors in the violin IMU.

During the accelerometer calibration experiment, the value of each of the three gyroscopes was also recorded in each of the bow orientations. An average was taken of all of these measurements for each of the gyroscopes and taken as its offset value. Therefore, while offset and gain estimates were made for each of the accelerometers, only offset was estimated for the gyroscopes by this process. The gain for each of the gyroscopes was estimated by an *a priori* calculation based on the published values provided by the datasheets. Later, an optimization

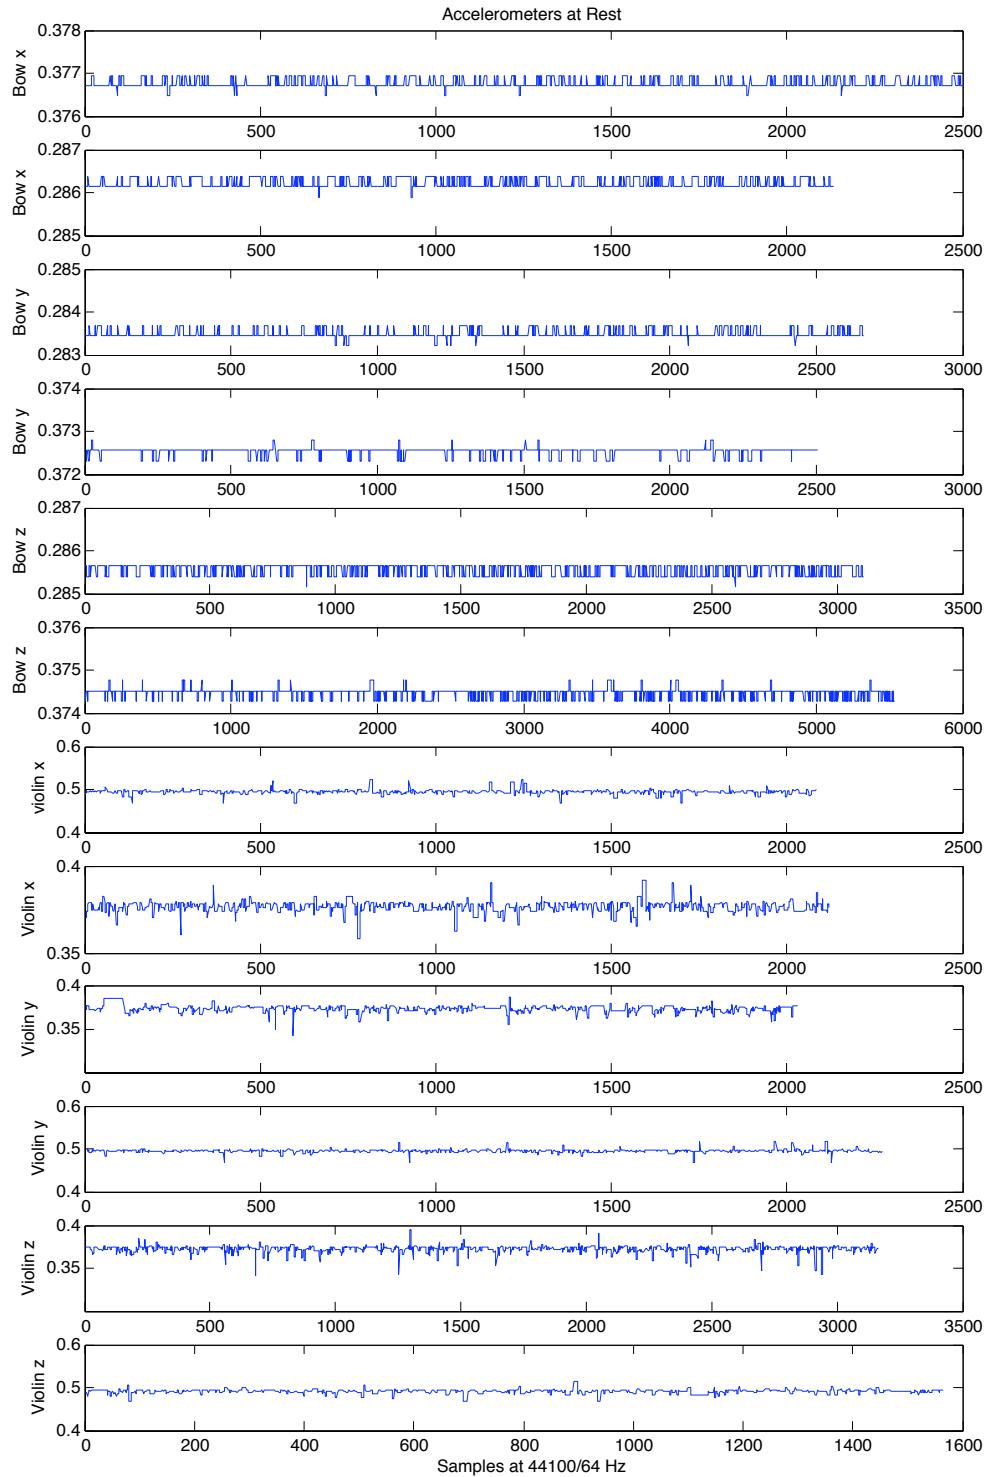


Figure 3.16: By subtracting the offsets and scaling the output as determined by the calibration exercise described above, the offsets values for each of the six inertial sensors on the bow, as well as those six on the violin were calculated. In addition, the gain of each of the accelerometers was estimated.

was performed to match the recorded angle produced by the Kalman filter with the actual angle of rotation. The details of the Kalman filter are discussed later in this chapter.

3.3.3 Position, Velocity Calibration

In order to calibrate the electric field position sensor to obtain estimates of bow-bridge distance (position along the string) and speed of the bow with respect to the violin, an experiment was conducted using a Vicon motion capture system [78] (located at MIT in the Ray and Maria Stata Center of CSAIL). This system, which provides millimeter precision at a sampling rate of 120 Hz, includes the Vicon IQ software and 16 infrared cameras in its Vicon 8I hardware.

In this experiment, both the violin and bow were fitted with visual markers and the movement of each throughout a series of bowstrokes was recorded, by both the custom measurement system designed here as well as by the Vicon motion capture system. The bow and violin were both outfitted with visual markers as seen in Figures 17(a) and 17(b). The bow used six markers, while the violin used five.

A bowing guide, as seen in Figure 17(c), was constructed so that the player would be able to easily bow with fixed bow-bridge distance. Measurements were taken at each of nine values of bow-bridge distance, while bowing at two different speeds. Recordings were also made while holding the bow fixed at designated positions along the length of the bow stick.

The position of one of the markers on the bow was plotted with respect to the violin's bridge antenna, with a local coordinate system defined by multiple markers on the violin. Then the raw bow position sensor data was input to Equations 3.6 and 3.9 and the constants were varied to minimize the error before inputting into the Kalman filter (discussed in the following section) for better estimation. Figure 3.19 shows the result of the calibration procedure and Kalman filtering.




Figure 3.17: (a) Six infrared Vicon markers, visible in this photograph, were used to track the bow movement during the position calibration experiment. (b) This photograph shows a subset of the Vicon markers used to track the violin movement during the position calibration experiment. (c) A bowing guide was fashioned using LEGO® pieces in order to constrain the path of the bow during the position calibration exercise.




Figure 3.18: After placing the infrared markers and taking careful measurements of the location of each, models of the bow and the violin were constructed using the Vicon software.




Figure 3.19: Using the position data produced by the Vicon motion capture system, the data from the electric field position sensor was calibrated and then Kalman filtered. This plot shows the agreement between the two measurement systems after calibration and estimation steps.

3.3.4 Kalman Filtering

Due to the growing errors in time from the drift in the accelerometer and gyroscope data, the linear acceleration and angular velocity estimates were refined using a Kalman filter [12]. There are two main steps in any Kalman filter algorithm: the prediction step and the update step. In the prediction step, the state of the system and its error covariance matrix are predicted using an *a priori* model. In the update step, the state and the error covariance are corrected using actual measurements and a Kalman gain matrix that minimizes the error covariance matrix.

Basic Kalman Filter

The prediction step begins with the following:

$$\hat{x}_k^- = A\hat{x}_{k-1}, \quad (3.10)$$

where \hat{x} is the vector describing the internal state of the system of interest. For the violin bowing measurement system, \hat{x} was defined as $(x \ y \ z \ u \ v \ w \ q_0 \ q_1 \ q_2 \ q_3 \ p_{\text{bias}} \ q_{\text{bias}} \ r_{\text{bias}})$, where x , y , and z are the 3D position of the bow; u , v , and w are the 3D velocity; q_0 , q_1 , q_2 , and q_3 comprise a quaternion that represents the orientation of the bow; and p_{bias} , q_{bias} , and r_{bias} are the biases on the gains of the gyroscope measurements. (The use of the quaternion instead of Euler angles θ , ϕ and γ to represent the 3D orientation of the bow will be explained later in this chapter.) The subscript k or $k - 1$ indicates a discrete time step. A is the model matrix that relates the previous state of the system (\hat{x}_{k-1}) to the current state (\hat{x}_k^-). This is the prediction step in which the algorithm predicts the current state of the system based on the model A and the previous state \hat{x}_{k-1} .

Next, the estimate of the *a priori* error covariance matrix P is predicted as

$$P_k^- = AP_{k-1}A^T + Q, \quad (3.11)$$

where Q is the process noise matrix. Q is defined assuming the measurements are independent of each other with a Gaussian probability distribution (the matrix is diagonal with variances of the process noise). P is updated in the prediction step using the model A and the process covariance matrix Q .

The update step of the Kalman filter begins by creating the Kalman gain matrix K_k .

$$K_k = P_k^- H^T (H P_k^- H^T + R)^{-1}, \quad (3.12)$$

where H is the matrix that relates the predicted state of the system to the current measurement z and R is the measurement error covariance matrix.

Using the Kalman gain matrix (K_k), the actual measurements of the system (\hat{z}_k), and the predicted state (\hat{x}_k^-) and predicted error covariance matrix (P_k^-), the internal state of the system (\hat{x}_k) and the error covariance matrix (P_k) can be updated in the current time step. That is,

$$\hat{x}_k = \hat{x}_k^- + K_k(\hat{z}_k - H\hat{x}_k^-) \quad (3.13)$$

and

$$P_k = (1 - K_k H) P_k^-. \quad (3.14)$$

Extended Kalman Filter

In order to implement a Kalman filter for this application, two modifications were made to the basic design just described. The first of these was the use of the derivatives of the prediction state necessitated by the accelerometer and gyroscope measurements, due to the fact that the outputs of each of these sensors (acceleration and angular rate) are derivatives of the states of interest (velocity and orientation). The second major alteration in this implementation was made due to the nonlinearity of the prediction and the measurement matrix. (The dynamics of a rigid body are nonlinear, as described in [12].) Because of this, an extended Kalman filter was designed as detailed below.

The first modification produced a two-stage prediction step of the internal state. The first stage predicts the derivative of the internal state,

$$\dot{\hat{x}}_k^- = A\hat{x}_{k-1}. \quad (3.15)$$

The second stage of the prediction step performs an Euler integration to obtain the prediction state:

$$\hat{x}_k^- = \hat{x}_{k-1} + \dot{\hat{x}}_k^- \cdot \Delta t. \quad (3.16)$$

The second modification required that the matrix A in Equation 3.15 be replaced by the nonlin-

ear function \hat{f} to yield:

$$\dot{\hat{x}}_k^- = \hat{f}(\hat{x}_{k-1}). \quad (3.17)$$

After making these changes to the prediction state, the prediction error covariance matrix was also revised. The derivative of the prediction error covariance matrix was computed as follows:

$$\dot{P}_k^- = A_k P_{k-1} + P_{k-1} A_k^T + Q. \quad (3.18)$$

Next, \dot{P}_k^- was integrated to obtain the prediction error covariance matrix:

$$P_k^- = P_{k-1} + \dot{P}_k^- \cdot \Delta t. \quad (3.19)$$

The matrix A_k was still required to determine the derivative of the prediction error covariance matrix and was calculated from the Jacobian of \hat{f} as

$$A_k = \begin{bmatrix} \frac{\partial f_1}{\partial x_{1,k}} & \cdots & \frac{\partial f_1}{\partial x_{n,k}} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_n}{\partial x_{1,k}} & \cdots & \frac{\partial f_n}{\partial x_{n,k}} \end{bmatrix}. \quad (3.20)$$

The update step of the extended Kalman filter follows the same form as the canonical Kalman filter, except the relationship of the measurement state to the internal state is also nonlinear, that is, $\hat{z}_{k,\text{prediction}} = \hat{g}(\hat{x}_k^-)$. This means that the H matrix now becomes H_k , which is derived from a Jacobian of the function $\hat{g}(\hat{x}_k^-)$ as

$$H_k = \begin{bmatrix} \frac{\partial g_{1,k}}{\partial x_{1,k}} & \cdots & \frac{\partial g_{1,k}}{\partial x_{n,k}} \\ \vdots & \ddots & \vdots \\ \frac{\partial g_{l,k}}{\partial x_{1,k}} & \cdots & \frac{\partial g_{l,k}}{\partial x_{n,k}} \end{bmatrix}. \quad (3.21)$$

Therefore, the Kalman gain matrix is slightly modified in order to use H_k instead of H :

$$K_k = P_k^- H_k^T (H_k P_k^- H_k^T + R)^{-1}. \quad (3.22)$$

The final update to the internal state is similar to that before and includes the new form of $\hat{z}_{k,\text{prediction}}$:

$$\hat{x}_k = \hat{x}_k^- + K_k(\hat{z}_k - \hat{z}_{k,\text{prediction}}). \quad (3.23)$$

The final error covariance matrix update now requires H_k :

$$P_k = (1 - K_k H_k) P_k^-. \quad (3.24)$$

Now that the structure of the new extended Kalman filter has been produced, the functions $\hat{f}(\hat{x})$ and $\hat{g}(\hat{x})$ must be described. The $\hat{f}(\hat{x})$ function (which predicts the derivatives of the state vector) is a set of smaller, nonlinear matrices. The position derivative update is straightforward, since the position and velocity of interest are those in the frame of the bow, not the room:

$$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{z} \end{bmatrix} = \begin{bmatrix} u \\ v \\ w \end{bmatrix}. \quad (3.25)$$

The velocity derivative update is more complex, due to the fact that the bow can be rotating.

The bow coordinate system rotated with respect to gravity is described as:

$$\begin{bmatrix} \dot{u} \\ \dot{v} \\ \dot{w} \end{bmatrix} = \begin{bmatrix} a_x \\ a_y \\ a_z \end{bmatrix} + DCM \begin{bmatrix} 0 \\ 0 \\ -g \end{bmatrix} - W_{\text{euler}} \begin{bmatrix} u \\ v \\ w \end{bmatrix} \quad (3.26)$$

where DCM is the Direction Cosine Matrix (from The MathWorks Aerospace Toolbox [39]), which transforms a vector in the room axes to the bow axes:

$$DCM = \begin{bmatrix} 1 - 2(q_2^2 + q_3^2) & 2(q_1 q_2 + q_0 q_3) & 2(q_1 q_3 - q_0 q_2) \\ 2(q_1 q_2 - q_0 q_3) & 1 - 2(q_1^2 + q_3^2) & 2(q_2 q_3 + q_0 q_1) \\ 2(q_1 q_3 + q_0 q_2) & 2(q_2 q_3 - q_0 q_1) & 1 - 2(q_1^2 + q_2^2) \end{bmatrix} \quad (3.27)$$

and finally, the W matrix uses the current rotations to apply addition acceleration moments:

$$W_{\text{euler}} = \begin{bmatrix} 0 & -(r_{\text{meas}} - r_{\text{bias}}) & q_{\text{meas}} - q_{\text{bias}} \\ r_{\text{meas}} - r_{\text{bias}} & 0 & -(p_{\text{meas}} - p_{\text{bias}}) \\ -(q_{\text{meas}} - q_{\text{bias}}) & p_{\text{meas}} - p_{\text{bias}} & 0 \end{bmatrix} \quad (3.28)$$

The quaternion (that represents orientation) derivative update is as follows:

$$\begin{bmatrix} \dot{q}_0 \\ \dot{q}_1 \\ \dot{q}_2 \\ \dot{q}_3 \end{bmatrix} = W \begin{bmatrix} q_0 \\ q_1 \\ q_2 \\ q_3 \end{bmatrix}, \quad (3.29)$$

where

$$W = \begin{bmatrix} 0 & -\frac{1}{2}(p_{\text{meas}} - p_{\text{bias}}) & -\frac{1}{2}(q_{\text{meas}} - q_{\text{bias}}) & -\frac{1}{2}(r_{\text{meas}} - r_{\text{bias}}) \\ \frac{1}{2}(p_{\text{meas}} - p_{\text{bias}}) & 0 & \frac{1}{2}(r_{\text{meas}} - r_{\text{bias}}) & -\frac{1}{2}(q_{\text{meas}} - q_{\text{bias}}) \\ \frac{1}{2}(q_{\text{meas}} - q_{\text{bias}}) & -\frac{1}{2}(r_{\text{meas}} - p_{\text{bias}}) & 0 & \frac{1}{2}(p_{\text{meas}} - p_{\text{bias}}) \\ \frac{1}{2}(r_{\text{meas}} - r_{\text{bias}}) & \frac{1}{2}(q_{\text{meas}} - q_{\text{bias}}) & -\frac{1}{2}(p_{\text{meas}} - p_{\text{bias}}) & 0 \end{bmatrix}. \quad (3.30)$$

The use of quaternions is warranted here, since this matrix update is composed of a simple polynomial. (Updating the Euler angle derivative is a much more complex set of transcendental functions.)

Finally, the derivatives of the gyroscope biases are equal to zero since we expect those values to be constant:

$$\begin{bmatrix} \dot{p}_{\text{bias}} \\ \dot{q}_{\text{bias}} \\ \dot{r}_{\text{bias}} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}. \quad (3.31)$$

The A_k matrix is determined from the Jacobian of the relationships listed in equations Equations 3.25 - 3.30, governed by Equation 3.20 (it is left out of this text due to size). The Kalman update was done twice: once for the accelerometer and gyroscope measurements ($\hat{z}_{k,\text{prediction inertial}}$) and once for the position measurements ($\hat{z}_{k,\text{prediction position}}$). These updates were performed separately since the position data was obtained at a slower rate. $\hat{z}_{k,\text{prediction inertial}}$ was determined by the actual measurement of the acceleration from the hardware and the current predicted quaternion was used to predict the acceleration from the *DCM*:

$$\hat{z}_{k,\text{prediction inertial}} = -g \begin{bmatrix} 2(q_1 q_3 - q_0 q_2) \\ 2(q_2 q_3 + q_0 q_1) \\ 1 - 2(q_1^2 + q_2^2) \end{bmatrix}. \quad (3.32)$$

This is only an approximation since this assumes the magnitude of the acceleration is always $|g|$. $\hat{z}_{k,\text{prediction position}}$ was easier:

$$\hat{z}_{k,\text{prediction position}} = \begin{bmatrix} x \\ y \\ z \end{bmatrix} \quad (3.33)$$

where the H_k matrix was obtained from Equation 3.21 in each case (it is also left out of this text due to size).

Implementation Details

In order to enable real-time operation, the Kalman filter designed was based on the open source code of the autopilot project (which provides information and instructions on how to build UAV (unmanned aerial vehicle) systems) [6] and implemented in Python. This Kalman filter program was designed to receive gesture data received through the USB port of an Apple Macbook with a 2 GHz Intel Core Duo processor (OS X) and output the filtered data to a recording program running in PureData (Pd) [60] by means of the Open Sound Control (OSC) protocol [91]. The Kalman filter was set to update the parameters at 200 Hz (same as the sampling rate of the measurement system hardware).

The performance of the Kalman filter for estimation of bow position (y -axis) is illustrated in Figure 3.19.

Chapter 4

Experiments

This chapter describes the two experimental studies conducted in this thesis work that were enabled by the creation of the measurement system for real-time violin bowing described in the previous chapter. These experiments were undertaken to explore the potential of this system and validate its use as a tool for furthering understanding of bowed string technique.

The first of these experiments was a study of common bowing techniques used in traditional Western classical violin repertoire. In this study, the measurement system was used to capture and record the bowing gestures produced by violinists when performing these techniques. The main goal of this experiment was to determine whether the measurement system is able to capture the salient features of these techniques so as to enable recognition of the individual techniques using gesture data alone (using no information related to the corresponding audio).

The second experiment undertaken in this thesis work was a brief exploration of the application of the measurement system as a real-time controller of a physical violin model, or “virtual violin”. In conducting this experiment, the quality of the measurement system’s playability and calibration were explored, as well as those features of the violin model itself. For comparison, two cases were created, one in which the model was driven using the real bowing parameters of an archived performance on the real test violin (without a live player), and one in which the model was driven in realtime while the participating violinist monitored and adapted to the synthesized sound produced.

Both of these experiments served as the pilot experiments for the violin bowing measurement system developed in this work. These experiments, as well as the preliminary testing sessions that were conducted in order to refine experimental protocol, are described in detail in

the pages that follow.

4.1 Preliminary Testing

Before beginning the final data collection for the bowing technique and virtual violin studies, a great deal of work was done to refine the experimental procedure with the help of experience gained during three sets of preliminary testing sessions. One set took place at each of the Schulich School of Music of McGill University, the Royal Academy of Music, and the Norwegian Academy of Music.

Most of the changes in experimental procedure were made to improve the communication of the kind of bowing technique desired in each part of the study. In the preliminary testing sessions, a musical excerpt for each of the techniques was provided in order to indicate a musical context to the participants. This was a helpful component of the instruction, but there still remained some ambiguity, especially given the varied backgrounds and playing experiences of the participants. (Even for those players in the same year of conservatory training, there were considerable differences in exposure to repertoire featuring these techniques.) Therefore, in the final data collection for the bowing technique study, an audio example of each bowing technique was provided to the participants, in addition to the contextual musical excerpt.

In addition to the improvements to the experimental procedure, two changes were also made to the measurement system. During the first two sets of these test sessions, those conducted at McGill and the Royal Academy of Music, only one of the four electric field position sensors (receive antennas) was in operation. In preparation for the sessions at the Norwegian Academy of Music, the remaining three position sensors were integrated into the measurement system. The second improvement to the measurement system (also made prior to the sessions at the Norwegian Academy of Music) was made when the ranges of the gyroscopes were adjusted to allow for higher values of angular velocity produced by some of the violinists. With these two hardware changes, the measurement system was now in its final state and ready for use in the following two experimental studies.

Feedback from the participants was always solicited during the preliminary testing sessions. Most of this feedback gained was in the form of differing opinions on the appropriate tempos for the various bowing techniques.

In all, ten students at the Schulich School of Music of McGill University, twelve students

from the Royal Academy of Music, one professional and one student from the greater Boston area, and eight professionals in Oslo from the community of the Norwegian Academy of Music or members of local professional orchestras participated in these preliminary test sessions.

4.2 Bowing Technique Study

Is the bowing measurement system capable of capturing the distinctions between common bowing techniques?

As already stated, the primary goal of the bowing technique study was to address this research question. In search of the answer to this question, the gesture and audio data generated by eight violinists performing six different bowing techniques on each of the four violin strings were recorded for later analysis. The details of the study protocol, experimental setup, and participants are discussed below.

4.2.1 Study Protocol

In this study each of the eight participants was asked to perform repetitions of a specific bowing technique originating from the Western “classical” music tradition. To help communicate the kind of bowstroke desired, a musical excerpt (from a work of the standard violin repertoire) featuring each bowing technique was provided from [10]. In addition, an audio example of the bowing technique for each of the four requested pitches was provided to the player. The bowing technique was notated clearly on a score, specifying the pitch and string, tempo, as well as any relevant articulation markings, for each set of the recordings.

Two different tempos were taken for each of the bowing techniques (on each pitch). First, trials were conducted using a characteristic tempo for each individual bowing technique. Immediately following these, trials were conducted using one common tempo. Though the target trials were actually those that were conducted with the same tempo across all of the bowing techniques, it was found early on that requesting performances using the characteristic tempo first enabled the players to perform at the common tempo with greater ease.

Both tempos required for each bowing technique were provided by a metronome. In some cases, a dynamics marking was written in the musical example, but the participants were instructed to perform all of the bowstrokes at a dynamic level of *mezzo forte*. Participants were instructed to take as much time as they required to either play through the musical example

and/or practice the technique before the start of the recordings to ensure that the performances would be as consistent as possible.

Three performances of each bowing technique, comprising one trial, were requested on each of the four pitches (one on each string). During the first preliminary set of recording sessions (at the Schulich School of Music of McGill University), which were conducted in order to refine the experimental procedure, participants were asked to perform these bowing techniques on the open strings. The rationale for this instruction was that the current measurement system does not capture any information concerning the left hand gestures. However, it was observed that players do not play as comfortably and naturally on open strings as when they finger pitches with the left hand. Therefore, in the subsequent recording sessions that comprise the actual technique study, the participants were asked to perform the bowing techniques on the fingered fourth interval above the open string pitch, with no vibrato.

The bowing techniques that comprised this study were *accented détaché*, *détaché lancé*, *détaché porté*, *louré*, *martelé*, *staccato*, and *spiccato*. The study instructions for the *accented détaché* technique are shown in Figure 4.1. For this technique, the characteristic tempo and the common tempo were the same (this common tempo was later used in for the subsequent bowing techniques, after trials using the characteristic tempo for each technique were recorded). The remaining instructions used in the study procedure for this experiment are included in Appendix D of this document.

4.2.2 Experimental Setup

In each trial of the bowing technique study, the physical gesture data were recorded simultaneously with the audio data produced in the performances of each technique. The experimental setup, depicted in Figure 4.2, was simple:

- custom violin bowing measurement system installed on a CodaBow® Conservatory™ violin bow [16] and the Yamaha SV-200 Silent Violin [92].
- headphones (through which the participants heard all pre-recorded test stimuli and real-time sound of the test violin)
- M-Audio Fast Track USB audio interface [36]

Detaché

Franck, *Sonata in A Major* for Violin and Piano (last movement)



Example of *accented détaché*

$\text{♩} = 84$



Figure 4.1: This figure shows the study instructions for the *accented détaché* bowing technique, including a musical example from [10]. Participants were asked to perform each of the four lines three times (constituting one trial).

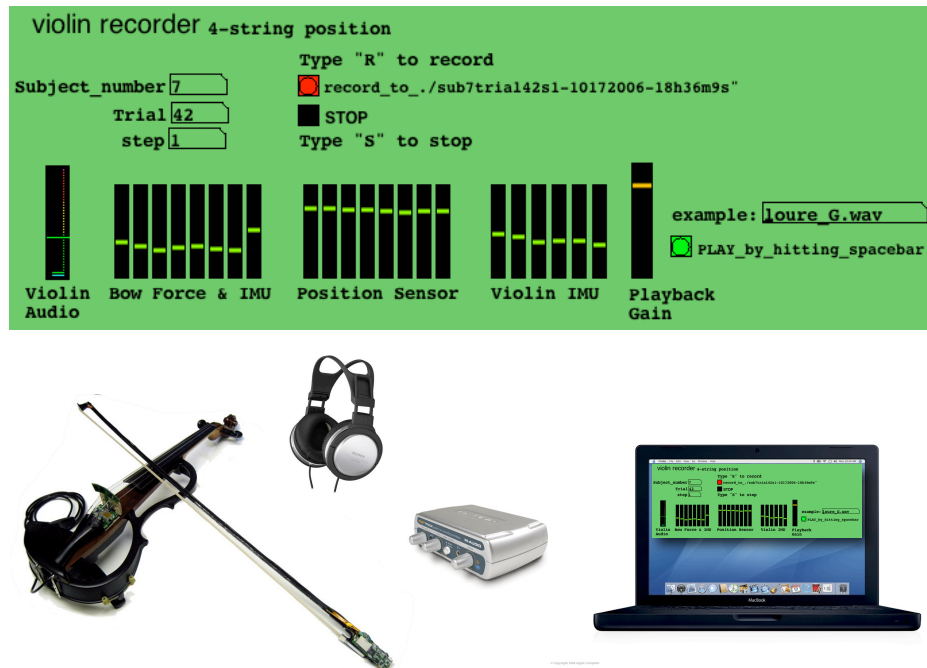


Figure 4.2: This figure describes the experimental setup used in the recording sessions for the bowing technique study. The top half of the figure shows the interface for the Pd recording patch, and the lower half shows the individual elements of the setup. From left to right, they are the custom violin bowing measurement system installed on a Yamaha SV-200 Silent Violin and a CodaBow® Conservatory™ violin bow; headphones; M-Audio Fast Track USB audio interface; and an Apple Macbook with a 2 GHz Intel Core Duo processor (OS X).

- Apple Macbook with a 2 GHz Intel Core Duo processor (OS X) running PureData (Pd) version 0.40.0-test08 [60]

The audio and the gesture data were recorded to file by means of a PD patch (shown in Figure 4.2), which encoded the gesture data as multi-channel audio in order to properly “sync” all of the data together. Each file was recorded with a trial number, repetition number, and time and date stamp. The Pure Data (Pd) patch also allowed for easy playback of recorded files used as test stimuli. (This patch was designed with the help of Roberto Aimi.)

The recordings took place in the Center for Interdisciplinary Research in Music Media and Technology (CIRMMT) of McGill University. Care was taken to create as quiet and natural a playing environment as possible.

4.2.3 Participants

The participants for the bowing technique study included eight students from the Schulich School of Music of McGill University, five of whom had taken part in the preliminary testing sessions and who therefore already had experience with the measurement system and the test recording setup. The participants were recruited by means of an email invitation and “word of mouth”, and they were each compensated 15 CAD to take part in the study.

All of the players were violin performance majors and had at least one year of conservatory level training. Most of them were also of the same approximate age. Only one player was a graduate student.

These studies received approval (application no. 0403000283) from the MIT Committee on the Use of Humans as Experimental Subjects (COUHES) [52].

4.3 Virtual Violin Performance Study

To explore the primary question of this thesis, whether instrument development may be facilitated by knowledge of player technique, two simple virtual violin validation experiments were conducted. In these, the bowing data generated by real violinists was used to drive a physical violin model to produce synthesized sound.

Currently, one of the most obvious uses for real bowing data is in the development of virtual bowed string instruments. Physical models of violins, which take as inputs physical bowing parameters, have now reached the level of sophistication at which the sound they produce compares well to that of their real instrument counterparts. However, true validation of a particular model has not yet occurred, as this requires demonstration of the same behavior as a real violin in response to realistically varying physical input parameters, and real input data has as of yet been unavailable. Also, such models are still of limited use to musicians (composers and performers), as it is not yet possible to play them using traditional physical techniques. That is, there are no realistic controllers yet available to allow players to properly explore their potential.

This thesis seeks to show that a sufficient coupling between a violin controller and physical model is now possible, as suggested in [75]. Detailed below are experiments designed to create useful comparisons between the virtual violin and the real test violin featured.

4.3.1 Violin Physical Model

The physical model used in these experiments was implemented by Stefania Serafin in Max/MSP external [68, 69, 70] and ported to the Pd environment by Julius Orion Smith III. It incorporates the computationally efficient waveguide synthesis first developed by Smith [74], with recent discoveries in the acoustics of bowed string instruments [89].

In particular, the friction interaction between the bow and the string is modeled taking into account the thermodynamic properties of rosin discovered by Jim Woodhouse and Jonathan Smith [73]. The model is calibrated from physical measurements on real violins and takes the physical parameters of bow force normal to the string, bow position (bow-bridge distance) and bow speed as inputs.

The basic operation of the model in a similar setup, as a Max/MSP external, was verified in previous experiments [99, 71].

4.3.2 Study Protocol

In the virtual violin performance study, two different cases were explored, one in which the violin physical model was driven with gesture data recorded from a previous performance by a real player, and one in which the model was played by a live musician using the measurement system as a physical controller.

Archived Performance Study

Does the virtual violin model respond realistically to calibrated data extracted from an actual performance by a real player?

The first virtual violin experiment utilized the data recorded from a previous performance on the test violin. This simple performance featured one full down-bow bowstroke, followed immediately by one full up-bow bowstroke. Both of these bowstrokes were performed on the open G string and each had a duration of approximately three seconds. The audio and gesture data were recorded using the custom measurement system designed here. In addition, the gesture was also recorded using the Vicon motion capture system, as in the position calibration experiment detailed in the previous chapter. The position data and speed data from the Vicon system were then combined with the force data from the custom measurement system.

Using these bowing gesture data as inputs, the virtual violin was “played”. That is, the three recorded primary bowing parameters were used to drive the violin physical model. The

resulting synthetic sound was recorded to a file, and then the original acoustic sound from the test violin (corresponding to the the gesture data recording) was compared to that produced by the model.

Real-Time Performance Study

Does the virtual violin model respond realistically to live control by a real player?

In the virtual violin experiment, a violinist was asked to play the violin model in realtime using the gesture measurement system as a physical controller for the model. In these trials, the player listened to the resultant real-time synthesis through headphones. The player was first asked to imitate as closely as possible the sound of the real-player performance featured in the archived performance study described above. After several trials, the task requested was modified and the player was asked instead to achieve the best quality of sound possible in two bowstrokes, one down-bow and one up-bow, of equal duration.

In this experiment, the outputs of the real-time Kalman filter (running in Python) described in the previous chapter were input to the violin model (running in Pd) using the OpenSound Control (OSC) protocol [91]. Though it is possible to produce the bowing parameters relative to the violin with the hardware sensing system that was built, for this preliminary experiment only the bow movement was used. Because of the simplicity of the performance task and the fact that the experiment was conducted while the player was seated, this was determined to be a reasonable approximation (though future studies will take full advantage of the hardware measurement system).

4.3.3 Experimental Setup

The equipment used in the experimental setup for the virtual violin performance study was the same as that in the bowing technique study. Only the software differed, as this study required the violin physical model Pd patch, shown in Figure 4.3.

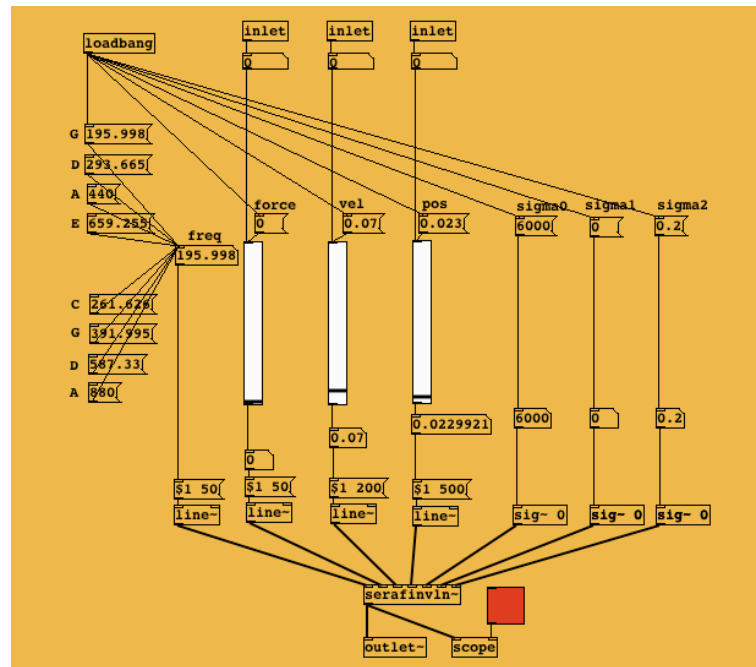


Figure 4.3: The violin physical model runs in realtime in a Pd patch. The model may be demonstrated by fixing the pitch of the output sound and manipulating, with a mouse, the sliders corresponding to the three primary bowing parameters: bow force, bow speed, and bow position (bow-bridge distance). In the virtual violin performance study conducted here, these inputs are connected to the corresponding three outputs of the measurement system’s real-time Kalman filter (running in Python) via OSC.

Chapter 5

Results

The first experimental study conducted for this thesis work was the bowing technique study, in which the gesture and audio data produced by the eight violinist participants when performing the six different bowing techniques were recorded. The evaluation of these data was crucial to determine the usefulness of the bowing measurement system that is at the core of the approach presented in this thesis. In the technique study evaluation, the data were first visually inspected in order to eliminate obvious outliers in the set, as well as to gain familiarity with the data itself.

In addition to the technique study, a study of the application of this measurement system to the real-time performance of a virtual violin model was also conducted. This chapter contains the results of each of these two studies related to bowed string performance. (A detailed description of the design of each of the studies is contained in the previous chapter.)

5.1 Technique Study Evaluation

The first goal of the technique study was to determine whether the gesture data provided by the measurement system would be sufficient to recognize the different six bowing techniques (*accented détaché*, *détaché lancé*, *louré*, *martelé*, *staccato*, *spiccato*) played by the eight violinist participants. To begin these classification explorations, only a subset of the gesture data provided by the measurement system was considered for the evaluations. Included in the analyses were data from the eight bow gesture sensors only: the downward and lateral forces; x, y, z acceleration; and angular velocity about the x, y, and z axes. This data subset was chosen to explore the progress possible with the augmented bow alone, as this method could easily be used with other instruments if shown to be sufficient.

Are the gesture data provided by this measurement system sufficient to capture the distinctions between common bowing techniques?

The second goal of the technique study was to investigate whether or not the gesture data could be used to identify the individual players themselves.

Can this measurement system be used to understand the similarities and differences in players' approaches toward the same techniques?

In order to answer these questions, a simple supervised classification algorithm was used. The k-Nearest-Neighbor (k-NN) algorithm was chosen because it is simple and robust for well-conditioned data. The k-NN algorithm requires the input data to have low dimensionality. In this study, the dimensionality of the gesture data vector used, 9152 (1144 samples in each time series x 8 gesture channels) was far too high. Therefore, the dimensionality of the gesture data set was first reduced before being input to the classifier.

5.1.1 Visual Data Evaluation

In all, there were 576 recorded files (8 players x 6 techniques x 4 strings x 3 performances of each) from the bowing technique study. Before beginning any data analysis, the bow gesture data was visually inspected to search for any obvious outlying examples. The data from each of the eight bow sensors (lateral force and downward force; x, y, z acceleration; and angular velocity about the x, y, and z axes), as well as the audio waveform produced, were plotted for each of the recorded files that comprise the bowing technique study. Each recording included 16 instances of each bowstroke. For this exercise, the entire time-series length, 1144 samples, of each gesture recording was used. (The gesture data for each of the recordings were aligned with the start of the first attack in the audio.)

For each individual player, the three performances of one technique on the first individual string were plotted, as in Figure 5.1, and visually inspected. The same was done for the player's performances of the same technique on each of the three remaining strings. Then, an overlay plot containing the data from all twelve examples (4 strings x 3 performances) of that player's technique was produced. This process was repeated once for each of the five other bowing techniques performed by that player. Then, the same was done for the data from each of the remaining seven players. In all, five overlay plots were created for each bowing technique for each player (one for the performances on each of the four strings, and one for all of the

performances of the same technique), making 240 overlay plots (8 players x 6 bowing technique x 5 overlay plots for each player) in all.

The eight channels of bow gesture data as well as the audio waveform of each trial recording were inspected for gross inconsistencies, such as timing errors made by the violinists. Throughout this visual evaluation process, 22 of the original 576 files (approximately 4%) were omitted and replaced with copies of files in the same class (technique, player, string). After performing this inspection process, the raw gesture data (now aligned with the beginning of the first audio attack of each recording and of the same length of 1144 samples), it was now possible to proceed to the dimensionality reduction phase of the analysis.




Figure 5.1: Overlay plot of the three examples of the spiccato bowing technique played on the E string by player 3. The audio waveform, as well as the data from each of the 8 bow sensors is shown.

5.1.2 Computing the Principal Components

In order to prepare the data for classification, the dimensionality of the raw data was first reduced. Principal Component Analysis (PCA), also known as the Karhunen-Loève transform, is one common technique used for this purpose [77]. PCA is a linear transform that transforms the data set into a new coordinate system such that the variance of the data vectors is maximized along the first coordinate dimension (known as the first principal component). That is, most of the variance is represented, or “explained”, by this dimension. Similarly, the second greatest variance is along the second coordinate dimension (the second principal component), the third greatest variance is along the third coordinate dimension (the third principal component), et cetera. Because the variance of the data decreases with increasing coordinate dimension, higher components may be disregarded for similar data vectors, thus resulting in decreased dimensionality of the data set.

Beginning with a p -sized set of zero-mean, normalized vectors, \mathbf{m} of dimension d , and a data matrix \mathbf{M} of dimension $d \times p$, PCA factors the covariance matrix

$$\mathbf{C} = \mathbf{M}\mathbf{M}^T \quad (5.1)$$

into

$$\mathbf{M}\mathbf{M}^T = \mathbf{W}\mathbf{\Sigma}^2\mathbf{W}^T, \quad (5.2)$$

where \mathbf{W} is the matrix of singular vectors of \mathbf{M} , or the matrix of eigenvectors of \mathbf{C} . The eigenvalues and eigenvectors of \mathbf{C} are of interest because they represent the weights and coordinate system of the principal components, respectively.

However, if \mathbf{m} is of dimension d , then the covariance matrix is of dimension d^2 , making the computation of the principal components computationally intractable for large values of d . Fortunately, when \mathbf{M} is not square and the number of samples of data used is much smaller than the dimension of the data, that is $p \ll d$, then Singular Value Decomposition (SVD) offers an efficient means of decomposing \mathbf{M} directly.

$$\mathbf{M} = \mathbf{W}\mathbf{\Sigma}\mathbf{V}^T \quad (5.3)$$

$$\mathbf{V}^T \mathbf{M} \mathbf{W} = \text{diag}(\sigma_1, \dots, \sigma_p), \quad (5.4)$$

with $\sigma_1 \geq \dots \geq \sigma_p \geq 0$, where the eigenvalues of \mathbf{C} are $\text{diag}(\sigma_1^2, \dots, \sigma_p^2)$. Therefore using this technique, a set of closely related eigenvalues of \mathbf{C} and the eigenvectors \mathbf{C} can be computed without dealing directly with \mathbf{C} , but rather with \mathbf{M} .

In the specific case of the bowing technique study, there were 576 recorded examples produced by the participants, and for each example, the 8 channels of the bow gesture data were prepared as described above. For each gesture channel, a time series that was 1144 samples long was taken. In order to ready the data for input into the SVD algorithm to compute the principal components, a vector for each of the examples was made by concatenating each of the 8 time series corresponding to the bow sensors to produce a vector of length 9152. Then, a matrix, \mathbf{M} was created, where $p = 576$ and $d = 9152$, in order to enable the following analyses.

Technique Separation

The first set of analyses were performed on the each of the data sets corresponding to each of the 8 players in order to investigate the distinctions between the techniques produced by a single individual. From the matrix \mathbf{M} , a smaller matrix composed of those 72 rows corresponding to each violinist (6 techniques x 4 strings x 3 performances of each) was taken and then decomposed using the SVD algorithm to produce the principal components.

The results of this analysis are included in the following pages. For each of the players, two figures were produced to illustrate the separability of the individual bowing techniques. The first of these was a plot showing the weights of each of the first 6 principal components for each of the 6 bowing techniques. The plots of the component weights for each player are displayed in Figures 5.2, 5.4, 5.6, 5.8, 5.10, 5.12, 5.14, and 5.16. In these plots, color was used to denote the string (G, D, A, or E) corresponding to each example.

The second figure produced for each player's data was a scatter plot of the first 3 principal components corresponding to each bowing technique. These are shown in Figures 5.3, 5.5, 5.7, 5.9, 5.11, 5.13, 5.15, and 5.17. In each of these plots, color was used to denote the 6 different bowing technique.




Figure 5.2: Analysis of gesture data from **player 1** yielded these principal component weights for the first six components. (The weights of all of the 9152 components sum to 1.) From top to bottom: *accented détaché*, *détaché lancé*, *louré*, *martelé*, *staccato*, *spiccato*. Trials are colored by string: G string (blue), D string (green), A string (red), E string (cyan).




Figure 5.3: Scatter plot of all six bowing techniques for **player 1**. *Accented détaché* (blue), *détaché lancé* (green), *louré* (red), *martelé* (cyan), *staccato* (magenta), *spiccato* (yellow). The axes correspond to the first three principal components.




Figure 5.4: Analysis of gesture data from **player 2** yielded these principal component weights for the first six components. (The weights of all of the 9152 components sum to 1.) From top to bottom: *accented détaché*, *détaché lancé*, *louré*, *martelé*, *staccato*, *spiccato*. Trials are colored by string: G string (blue), D string (green), A string (red), E string (cyan).




Figure 5.5: Scatter plot of all six bowing techniques for **player 2**. *Accented détaché* (blue), *détaché lancé* (green), *louré* (red), *martelé* (cyan), *staccato* (magenta), *spiccato* (yellow). The axes correspond to the first three principal components.




Figure 5.6: Analysis of gesture data from **player 3** yielded these principal component weights for the first six components. (The weights of all of the 9152 components sum to 1.) From top to bottom: *accented détaché*, *détaché lancé*, *louré*, *martelé*, *staccato*, *spiccato*. Trials are colorized by string: G string (blue), D string (green), A string (red), E string (cyan).




Figure 5.7: Scatter plot of all six bowing techniques for **player 3**. *Accented détaché* (blue), *détaché lancé* (green), *louré* (red), *martelé* (cyan), *staccato* (magenta), *spiccato* (yellow). The axes correspond to the first three principal components.




Figure 5.8: Analysis of gesture data from **player 4** yielded these principal component weights for the first six components. (The weights of all of the 9152 components sum to 1.) From top to bottom: *accented détaché*, *détaché lancé*, *louré*, *martelé*, *staccato*, *spiccato*. Trials are colored by string: G string (blue), D string (green), A string (red), E string (cyan).




Figure 5.9: Scatter plot of all six bowing techniques for **player 4**. *Accented détaché* (blue), *détaché lancé* (green), *louré* (red), *martelé* (cyan), *staccato* (magenta), *spiccato* (yellow). The axes correspond to the first three principal components.




Figure 5.10: Analysis of gesture data from **player 5** yielded these principal component weights for the first six components. (The weights of all of the 9152 components sum to 1.) From top to bottom: *accented détaché*, *détaché lancé*, *louré*, *martelé*, *staccato*, *spiccato*. Trials are colorized by string: G string (blue), D string (green), A string (red), E string (cyan).




Figure 5.11: Scatter plot of all six bowing techniques for **player 5**. *Accented détaché* (blue), *détaché lancé* (green), *louré* (red), *martelé* (cyan), *staccato* (magenta), *spiccato* (yellow). The axes correspond to the first three principal components.




Figure 5.12: Analysis of gesture data from **player 6** yielded these principal component weights for the first six components. (The weights of all of the 9152 components sum to 1.) From top to bottom: *accented détaché*, *détaché lancé*, *louré*, *martelé*, *staccato*, *spiccato*. Trials are colorized by string: G string (blue), D string (green), A string (red), E string (cyan).




Figure 5.13: Scatter plot of all six bowing techniques for **player 6**. *Accented détaché* (blue), *détaché lancé* (green), *louré* (red), *martelé* (cyan), *staccato* (magenta), *spiccato* (yellow). The axes correspond to the first three principal components.

Given this analysis, the gesture data from player 6 exhibited the least separability between the six different bowing techniques. As seen in the scatter plot in Figure 5.13, the first three principal components corresponding to techniques *détaché lancé* and *spiccato* were significantly similar.




Figure 5.14: Analysis of gesture data from **player 7** yielded these principal component weights for the first six components. (The weights of all of the 9152 components sum to 1.) From top to bottom: *accented détaché*, *détaché lancé*, *louré*, *martelé*, *staccato*, *spiccato*. Trials are colorized by string: G string (blue), D string (green), A string (red), E string (cyan).




Figure 5.15: Scatter plot of all six bowing techniques for **player 7**. *Accented détaché* (blue), *détaché lancé* (green), *louré* (red), *martelé* (cyan), *staccato* (magenta), *spiccato* (yellow). The axes correspond to the first three principal components.




Figure 5.16: Analysis of gesture data from **player 8** yielded these principal component weights for the first six components. (The weights of all of the 9152 components sum to 1.) From top to bottom: *accented détaché*, *détaché lancé*, *louré*, *martelé*, *staccato*, *spiccato*. Trials are colorized by string: G string (blue), D string (green), A string (red), E string (cyan).




Figure 5.17: Scatter plot of all six bowing techniques for **player 8**. *Accented détaché* (blue), *détaché lancé* (green), *louré* (red), *martelé* (cyan), *staccato* (magenta), *spiccato* (yellow). The axes correspond to the first three principal components.

Player Separation

After exploring the ability to use the principal components produced by the SVD method to separate the data for individual violinists by technique, the focus was then changed to inspect the distinctions between players approaches to the same bowing technique.

In order to explore the differences between individual approaches for the six different bowing techniques in this study, SVD was performed on all of the player data for each technique. Then the first three components produced were plotted to show what separation, if any, there was between the performances by the different players for the same bowing technique. The results of this investigation are included in the top portions of Figures 5.18 - 5.23, which show the results obtained by plotting all of the player data at once. In these plots, color was used to denote each individual player.

Though these scatter plots do not show clear separation for every one of the violinists approaches to each stroke, it was shown that some players were separable. In illustration of this point, the bottom portions of Figures 5.18 - 5.23 highlight only players 2, 3, 5 for the corresponding bowing technique.





Figure 5.18: Top: scatter plot of *Accented détaché* for all players. Bottom: Scatter plot of *Accented détaché* for **players 2** (green), **3** (red), and **5** (magenta). The axes correspond to the first three principal components.

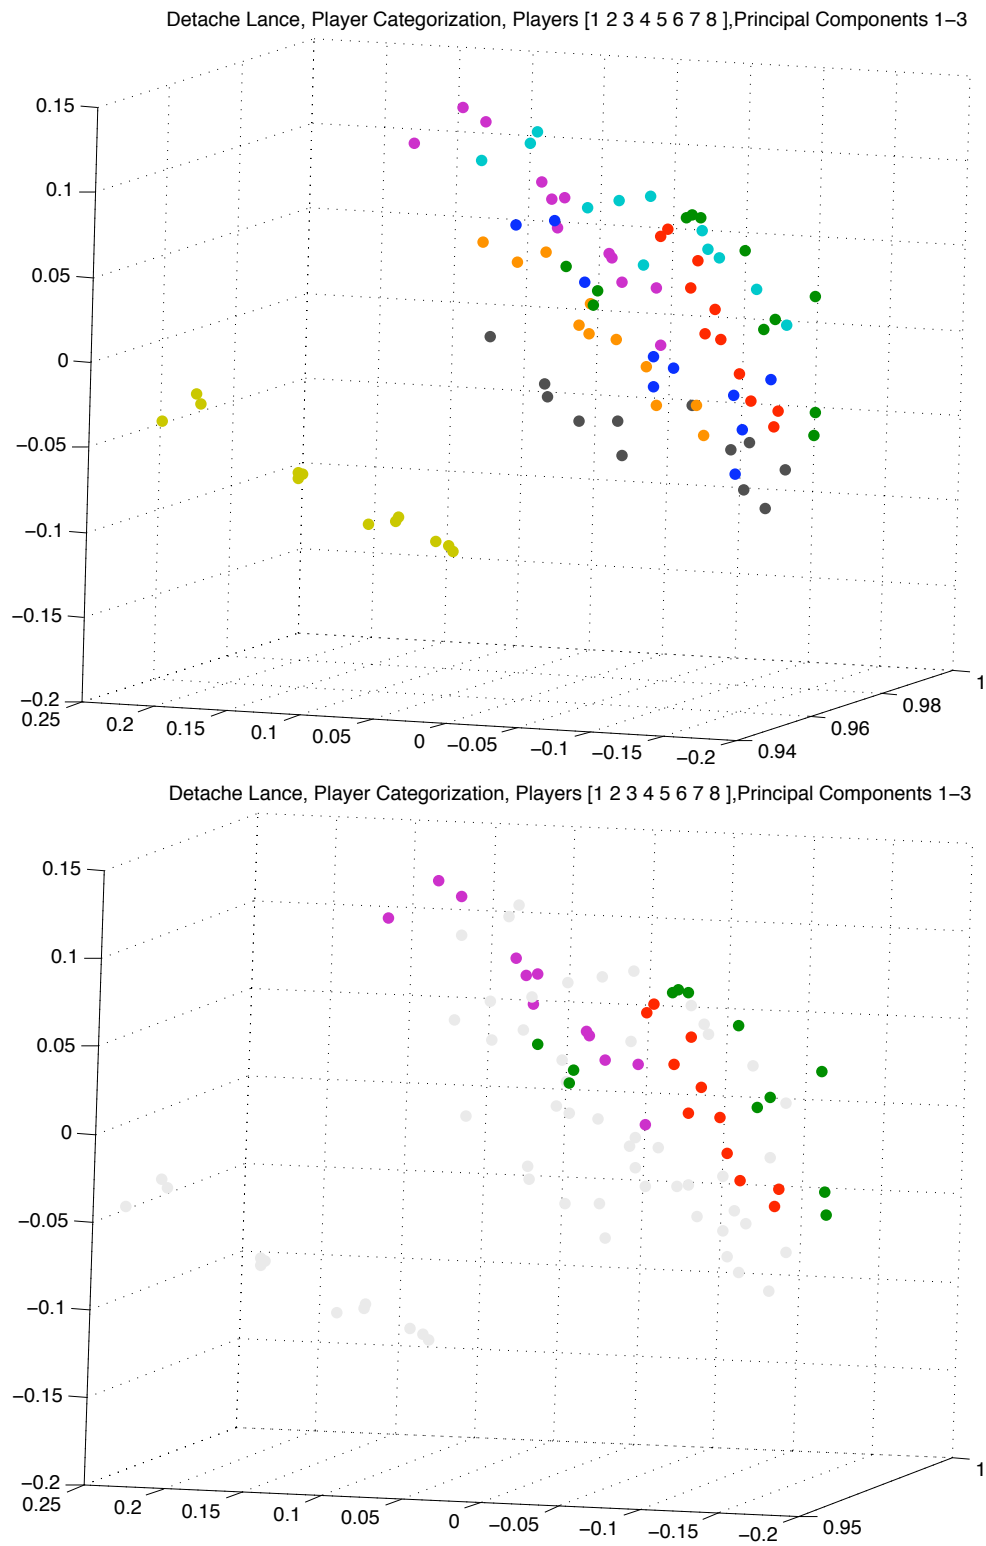


Figure 5.19: Top: Scatter plot of *détaché lancé* for all players. Bottom: scatter plot of *lancé détaché* for **players 2** (green), **3** (red), and **5** (magenta). The axes correspond to the first three principal components.

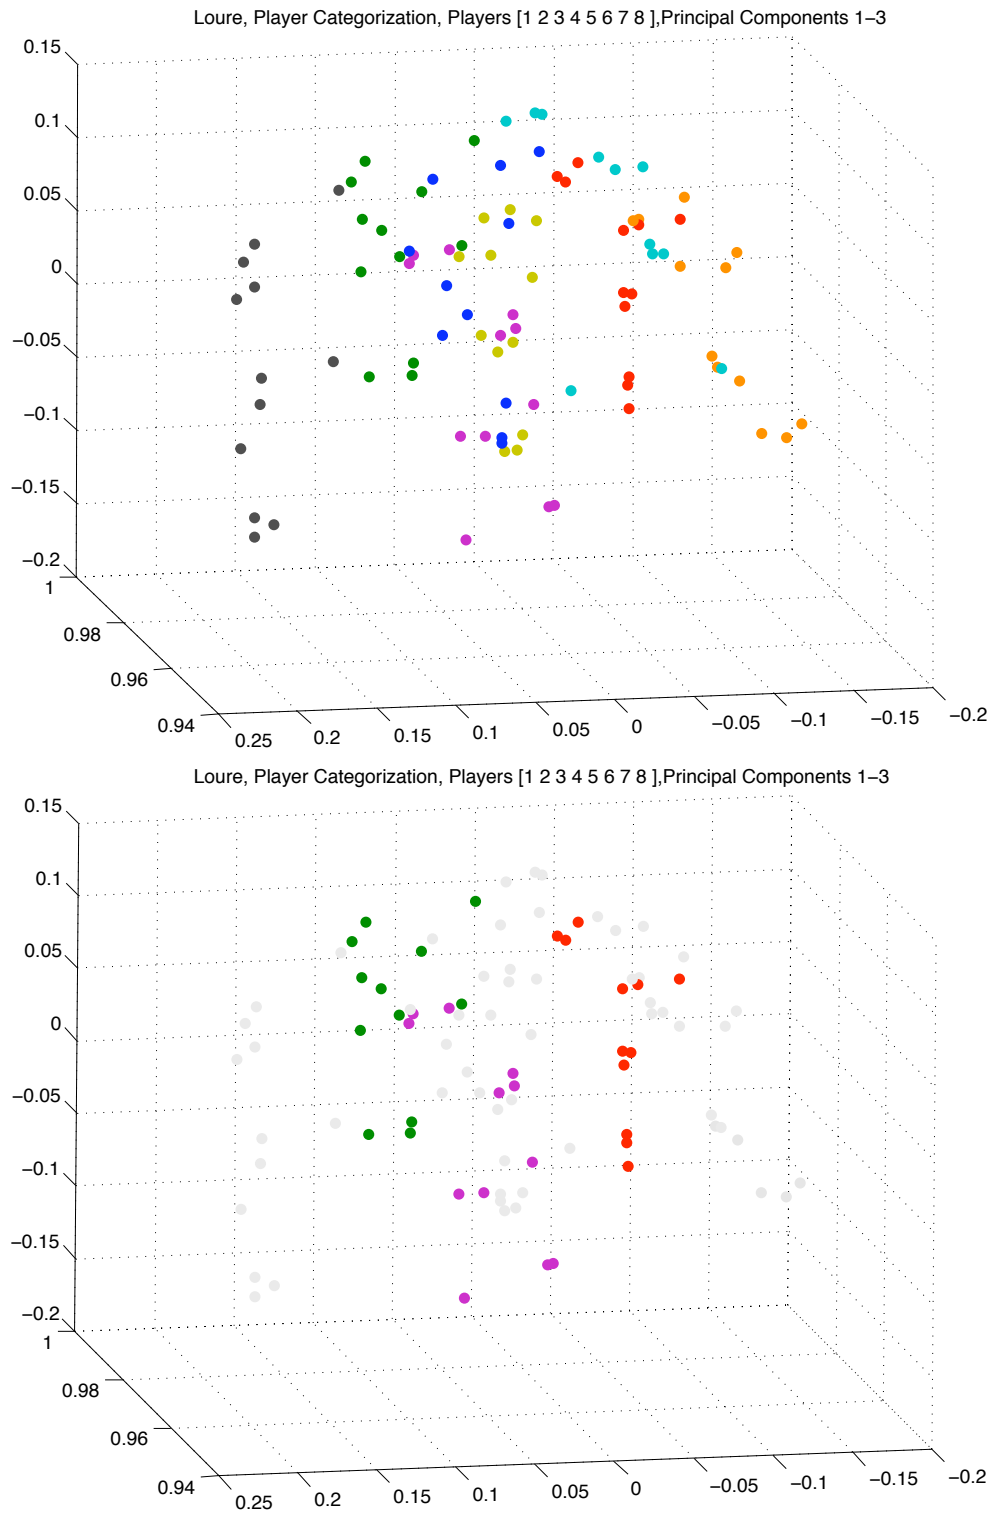


Figure 5.20: Top: scatter plot of *louré* for all players. Bottom: scatter plot of *louré* for **players 2** (green), **3** (red), and **5** (magenta). The axes correspond to the first three principal components.




Figure 5.21: Top: scatter plot of *martelé* for all players. Bottom: scatter plot of *martelé* for **players 2** (green), **3** (red), and **5** (magenta). The axes correspond to the first three principal components.

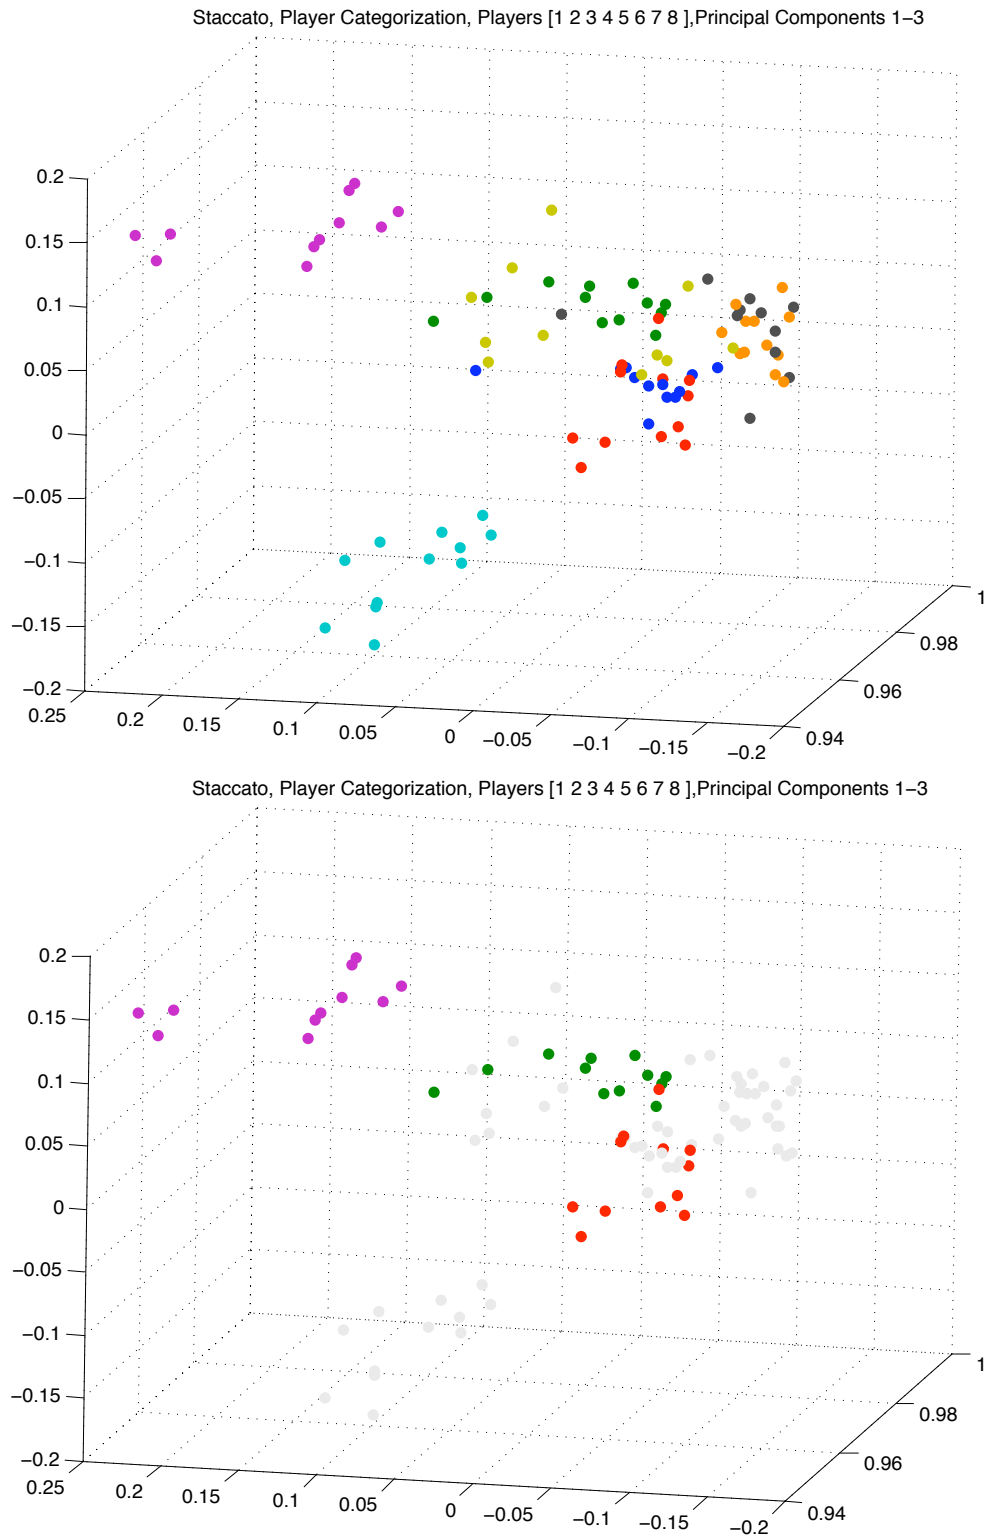


Figure 5.22: Top: scatter plot of *staccato* for all players. Bottom: scatter plot of *staccato* for **players 2** (green), **3** (red), and **5** (magenta). The axes correspond to the first three principal components.




Figure 5.23: Top: scatter plot of *spiccato* for all players. Bottom: scatter plot of *spiccato* for **players 2** (green), **3** (red), and **5** (magenta). The axes correspond to the first three principal components.

5.1.3 k-NN Classification

After conducting the explorations of the bowing technique study data using the principal components produced by the SVD method (detailed above), the challenge of classifying the data was undertaken. This classification work was done in order to investigate the ability of the gesture measurement system to capture the salient differences between the bowing techniques, that is, addressing the critical question: *Is the bowing measurement system capable of capturing the distinctions between common bowing techniques?*

Toward this goal, a k-nearest-neighbor classifier was used. Specifically, Nabney's matlab implementation, available online [46] and discussed in [47] was employed. In particular, two functions, "knn", which creates a k-NN classifier, and "knnfwd", which performs forward propagation through the k-nearest-neighbor classifier, were used.

Technique Classification

Throughout this phase of exploration, three different cases of technique classification were addressed. They are as follows:

- Intra-Player Case: In this case, a subset of the data from *each individual player* was used to train the k-NN algorithm in order to classify the remaining data from the same player by technique.
- All-Player Case: In this case, a subset of the data contributed by *all of the players* was used to train the k-NN algorithm in order to classify the remaining data from all of the players by technique.
- Player-1 Case: In this case, the data from *seven players* was used to train the k-NN algorithm in order to classify the data from the eight remaining player by technique.

In each case, the principal components of the training data set were first computed using the SVD method. The remaining data (to be classified) were then projected into the eigenspace determined by this exercise. Some number of the principal components corresponding to the training data were input to the k-NN algorithm, enabling the remaining data to be classified according to technique. For each case, a cross-validation procedure was obeyed, as this process was repeated as the training data (and the data to be classified) were rotated. The final estimates of the success of the classification cases were taken as the mean and standard deviation of the success rates of these cross-validation trials.

Intra-Player Case

How consistent are individual players throughout multiple performances of the same technique, and how well do they distinguish the various bowing techniques?

The first classification exercise applied to the technique study bow data was an inspection of the performance of each individual player.

After reducing the dimensionality of two-thirds of the data from a *single player* by computing the principal components using SVD, the k-NN algorithm was trained on the same two-thirds of the data in order to classify the remaining one-third of the player's data. The number of nearest neighbors was maintained at one, while the number of principal components was increased from one to ten. Three-fold cross validation was performed by rotating the test/predict data three times. For each number of principal components, the mean and standard deviation of the cross-validation trials were computed to determine the overall success of the technique classification.

This whole procedure was repeated for each of the other eight players. Good prediction rates for each player's gesture data were achieved for even low numbers of principal components. The results of this analysis for each individual player are shown in Figure 5.24, in which the effect of increasing the number of principal components on the overall classification success is clearly demonstrated. In fact, by increasing the number of principal components, very high prediction rates were reached for each of the eight players, and for five of the eight players rates of over 90% were achieved using only three principal components.

Of all of the eight violinists, only the gesture data provided by player 6 achieved a success rate significantly under 90% with four principal components as input to the classifier.

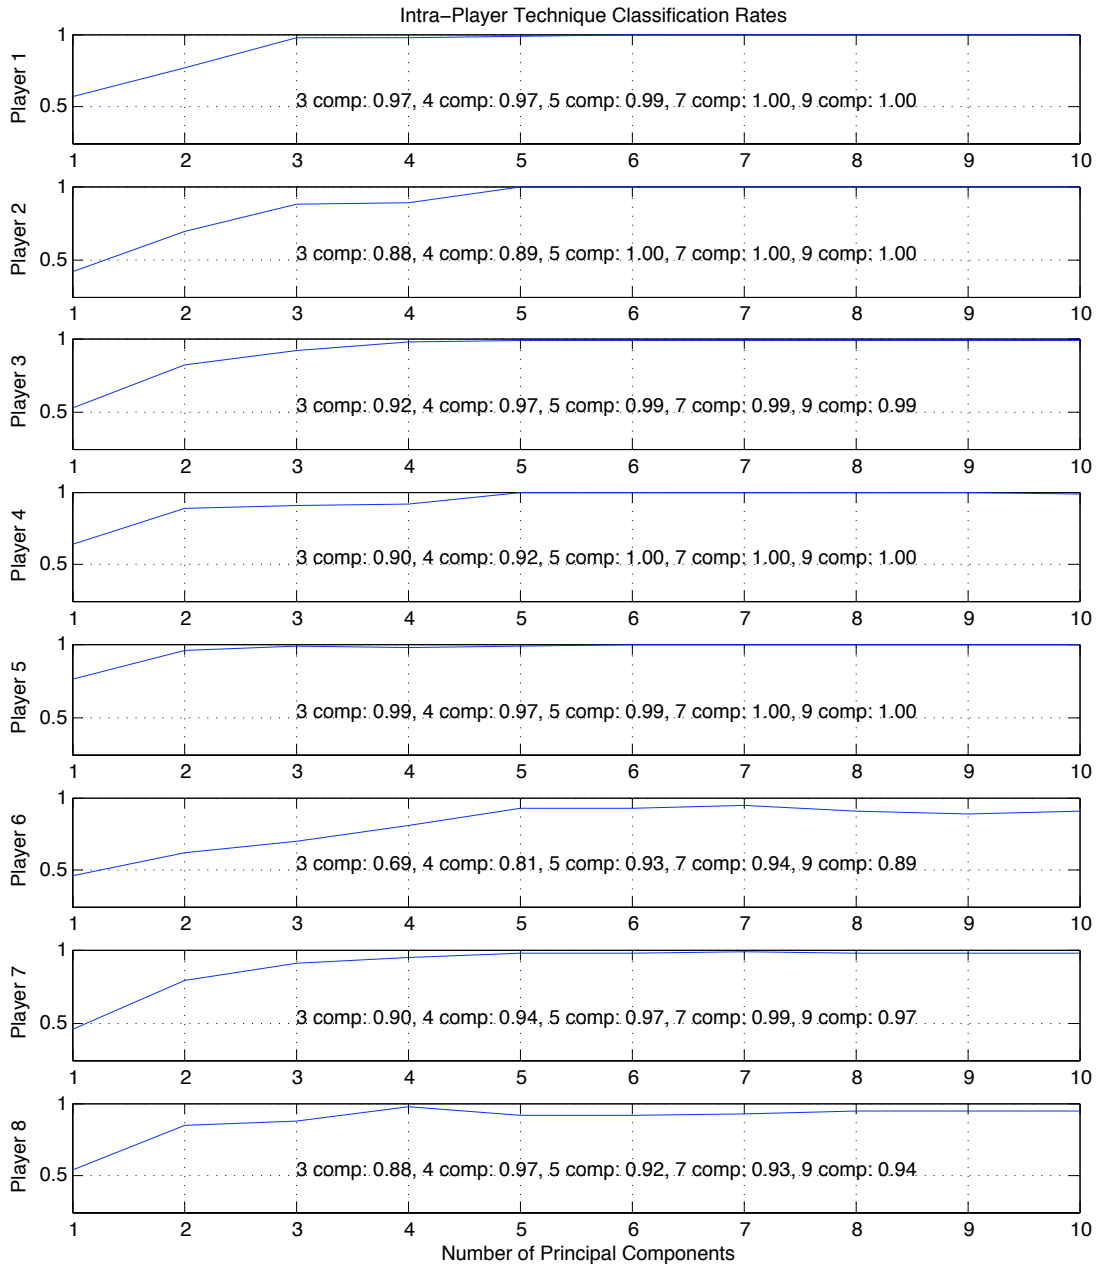


Figure 5.24: This figure shows the mean recognition rates versus the number of principal components used, produced by using the k-NN method (with one nearest neighbor) with two-thirds of each player's gesture data as training to predict the remaining third of that player's gesture data.

actual \ class.	acc. det.	det. lancé	louré	martele	staccato	spiccato
acc. det.	0.698	0.000	0.031	0.167	0.104	0.000
det. lancé	0.010	0.385	0.240	0.010	0.156	0.198
louré	0.031	0.146	0.510	0.094	0.177	0.042
martelé	0.156	0.031	0.094	0.490	0.229	0.000
staccato	0.042	0.125	0.250	0.188	0.396	0.000
spiccato	0.000	0.115	0.073	0.000	0.000	0.812

Table 5.1: Training on two-thirds of the data from each of the eight players, predicting the remaining third of each player's data (with overall prediction of $54.9 \pm 17.1\%$) with **two** principal components.

All-Player Case

Considering all of the player gesture data together, how well can a subset of the data be used to classify the remaining data?

After performing the intra-player analysis discussed above, the k-NN classifier was trained using two-thirds of the data from *all of the players* to classify the remaining one-third of *all player data* by technique.

As in the intra-player case, the number of nearest neighbors was maintained as one, while the number of principal components was increased from one to ten. Again, three-fold cross validation was performed by rotating the training data three times and the three resultant classification rates were used to determine the mean and standard deviation of the overall classification success.

The effect on the overall classification success of the number of principal components is clearly illustrated by Figure 5.25. The results of this analysis are shown in the following Tables 5.1-5.9, which show the individual confusion matrices for the k-NN classification results produced with 2-10 principal components as input. As seen in Table 5.5, an overall rate of $90.8 \pm 2.6\%$ was achieved with only 6 principal components, while using 7 principal components enables classification of 6 bowing techniques of over $95.3 \pm 2.6\%$ of the remaining data.

actual \ class.	acc. détaché	det. lancé	louré	martele	staccato	spiccato
acc. détaché	0.812	0.000	0.052	0.083	0.052	0.000
det. lancé	0.000	0.625	0.177	0.031	0.052	0.115
louré	0.021	0.177	0.635	0.000	0.083	0.083
martelé	0.073	0.042	0.010	0.656	0.219	0.000
staccato	0.062	0.104	0.062	0.198	0.573	0.000
spiccato	0.000	0.094	0.052	0.000	0.000	0.854

Table 5.2: Training on two-thirds of the data from each of the eight players, predicting the remaining third of each player's data (with overall prediction of $69.3 \pm 11.3\%$) with **three** principal components.

actual \ class.	acc. détaché	det. lancé	louré	martele	staccato	spiccato
acc. détaché	0.823	0.021	0.073	0.073	0.010	0.000
det. lancé	0.021	0.719	0.104	0.021	0.062	0.073
louré	0.031	0.115	0.792	0.000	0.042	0.021
martelé	0.094	0.031	0.010	0.719	0.146	0.000
staccato	0.010	0.031	0.073	0.094	0.792	0.000
spiccato	0.000	0.083	0.031	0.000	0.000	0.885

Table 5.3: Training on two-thirds of the data from each of the eight players, predicting the remain third of each player's data (with overall prediction of $78.8 \pm 6.4\%$) with **four** principal components.

actual \ class.	acc. détaché	det. lancé	louré	martele	staccato	spiccato
acc. détaché	0.844	0.000	0.052	0.083	0.021	0.000
det. lancé	0.000	0.781	0.073	0.031	0.052	0.062
louré	0.052	0.062	0.802	0.000	0.042	0.042
martelé	0.052	0.052	0.000	0.792	0.104	0.000
staccato	0.010	0.021	0.031	0.073	0.865	0.000
spiccato	0.000	0.052	0.031	0.000	0.000	0.917

Table 5.4: Training on two-thirds of the data from each of the eight players, predicting the remain third of each player's data (with overall prediction of $83.3 \pm 5.2\%$) with **five** principal components

actual \ class.	acc. détaché	det. lancé	louré	martele	staccato	spiccato
acc. détaché	0.875	0.000	0.031	0.073	0.021	0.000
det. lancé	0.000	0.896	0.000	0.031	0.021	0.052
louré	0.021	0.000	0.927	0.000	0.052	0.000
martelé	0.052	0.031	0.000	0.885	0.031	0.000
staccato	0.000	0.010	0.052	0.000	0.938	0.000
spiccato	0.000	0.042	0.031	0.000	0.000	0.927

Table 5.5: Training on two-thirds of the data from each of the eight players, predicting the remain third of each player's data (with overall prediction of $90.8 \pm 2.6\%$) with **six** principal components.

actual \ class.	acc. détaché	det. lancé	louré	xmartele	staccato	spiccato
acc. détaché	0.938	0.010	0.010	0.042	0.000	0.000
det. lancé	0.000	0.917	0.000	0.010	0.021	0.052
louré	0.000	0.000	0.979	0.000	0.021	0.000
martelé	0.042	0.021	0.000	0.938	0.000	0.000
staccato	0.000	0.010	0.010	0.000	0.979	0.000
spiccato	0.000	0.031	0.000	0.000	0.000	0.969

Table 5.6: Training on two-thirds of the data from each of the eight players, predicting the remain third of each player's data (with overall prediction of $95.3 \pm 2.6\%$) with **seven** principal components.

actual \ class.	acc. détaché	det. lancé	louré	martelé	staccato	spiccato
acc. détaché	0.938	0.010	0.010	0.042	0.000	0.000
det. lancé	0.000	0.938	0.010	0.010	0.021	0.021
louré	0.000	0.000	0.979	0.000	0.021	0.000
martelé	0.031	0.021	0.000	0.948	0.000	0.000
staccato	0.000	0.010	0.010	0.000	0.979	0.000
spiccato	0.000	0.031	0.000	0.000	0.000	0.969

Table 5.7: Training on two-thirds of the data from each of the eight players, predicting the remain third of each player's data (with overall prediction of $95.8 \pm 2.0\%$) with **eight** principal components.

actual \ class.	acc. détaché	det. lancé	louré	martelé	staccato	spiccato
acc. détaché	0.938	0.010	0.010	0.042	0.000	0.000
det. lancé	0.000	0.948	0.010	0.010	0.010	0.021
louré	0.000	0.000	0.990	0.000	0.010	0.000
martelé	0.031	0.021	0.000	0.948	0.000	0.000
staccato	0.000	0.000	0.010	0.000	0.990	0.000
spiccato	0.000	0.031	0.010	0.000	0.000	0.958

Table 5.8: Training on two-thirds of the data from each of the eight players, predicting the remain third of each player's data (with overall prediction of $96.2 \pm 2.3\%$) with **nine** principal components.

actual \ class.	acc. détaché	det. lancé	louré	martelé	staccato	spiccato
acc. détaché	0.927	0.010	0.010	0.052	0.000	0.000
det. lancé	0.000	0.927	0.010	0.010	0.000	0.052
louré	0.000	0.000	0.990	0.000	0.010	0.000
martelé	0.042	0.021	0.000	0.938	0.000	0.000
staccato	0.000	0.000	0.010	0.000	0.990	0.000
spiccato	0.000	0.052	0.010	0.000	0.000	0.938

Table 5.9: Training on two-thirds of the data from each of the eight players, predicting the remain third of each player's data (with overall prediction of $95.1 \pm 3.0\%$) with **ten** principal components.




Figure 5.25: Mean prediction rates produced by KNN using two-thirds of the data from each of the eight players to predict the remaining one-third of all player data and increasing principal components from one to ten.

actual \ class.	acc. détaché	det. lancé	louré	martelé	staccato	spiccato
acc. détaché	1.000	0.000	0.000	0.000	0.000	0.000
det. lancé	0.000	1.000	0.000	0.000	0.000	0.000
louré	0.000	0.000	1.000	0.000	0.000	0.000
martelé	0.000	0.250	0.000	0.500	0.250	0.000
staccato	0.000	0.000	0.000	0.000	1.000	0.000
spiccato	0.000	0.000	0.000	0.000	0.000	1.000

Table 5.10: Training on players 2-8, classifying data from **player 1** (overall classification rate of $91.7 \pm 20.4\%$ with ten principal components, seven nearest neighbors).

Player-1 Case

If the classifier is given no training data related to one player, how well can the rest of the gesture data from seven players be used to classify the first player's data?

After addressing the intra-player and the all-player technique classification cases, the third and final technique classification case, the player-1 case, was implemented. In this case, the k-NN classifier was given a training set consisting of all of the data from seven of the eight players in order to classify the eight remaining player's data by technique.

Ten principal components (corresponding to the training data set used) were taken as input, and the number of nearest neighbors was seven. This time, an 8-fold cross validation procedure was followed. The training data set was rotated 8 times to include the data from 7 of the players as input to the k-NN algorithm to enable the classification of the data from the remaining player. The confusion matrices produced by this process are displayed in following Tables 5.10 - 5.17, and a bar graph showing the overall classification rates for the data contributed by each individual player is shown in Figure 5.26.

The overall classification rates of these eight different cross validation trials were then used to compute the mean and standard deviation of the overall success of this case of classification by technique. This result was calculated to be $80.4 \pm 0.14\%$. The confusion matrix corresponding to the overall classification results for the individual techniques is shown in Table 5.18.

actual \ class.	acc. détaché	det. lancé	louré	martelé	staccato	spiccato
acc. détaché	1.000	0.000	0.000	0.000	0.000	0.000
det. lancé	0.000	0.917	0.000	0.000	0.000	0.083
louré	0.000	0.000	1.000	0.000	0.000	0.000
martelé	0.000	0.000	0.000	1.000	0.000	0.000
staccato	0.000	0.000	0.000	0.000	1.000	0.000
spiccato	0.000	0.000	0.000	0.000	0.000	1.000

Table 5.11: Training on players 1 and 3-8, classifying data from **player 2** (overall classification rate of $98.6 \pm 3.4\%$) with ten principal components, seven nearest neighbors.

actual \ class.	acc. détaché	det. lancé	louré	martelé	staccato	spiccato
acc. détaché	0.000	0.000	0.000	1.000	0.000	0.000
det. lancé	0.000	0.917	0.083	0.000	0.000	0.000
louré	0.000	0.000	1.000	0.000	0.000	0.000
martelé	0.000	0.333	0.000	0.583	0.083	0.000
staccato	0.000	0.000	0.167	0.000	0.833	0.000
spiccato	0.000	0.000	0.000	0.000	0.000	1.000

Table 5.12: Training on players 1-2 and 4-8, classifying data from **player 3** (overall classification rate of $72.2 \pm 38.6\%$) with ten principal components, seven nearest neighbors.

actual \ class.	acc. détaché	det. lancé	louré	martelé	staccato	spiccato
acc. détaché	0.333	0.000	0.000	0.667	0.000	0.000
det. lancé	0.000	0.417	0.000	0.083	0.500	0.000
louré	0.000	0.000	0.500	0.000	0.500	0.000
martelé	0.083	0.000	0.000	0.917	0.000	0.000
staccato	0.000	0.000	0.000	0.000	0.917	0.083
spiccato	0.000	0.000	0.000	0.000	0.000	1.000

Table 5.13: Training on players 1-3 and 5-8, classifying data from **player 4** (overall classification rate of $68.1 \pm 30.0\%$) with ten principal components, seven nearest neighbors.

actual \ class.	acc. détaché	det. lancé	louré	martelé	staccato	spiccato
acc. détaché	0.917	0.000	0.000	0.083	0.000	0.000
det. lancé	0.000	0.750	0.000	0.083	0.000	0.167
louré	0.000	0.000	1.000	0.000	0.000	0.000
martelé	0.333	0.000	0.000	0.667	0.000	0.000
staccato	0.000	0.000	0.000	0.250	0.750	0.000
spiccato	0.000	0.000	0.000	0.000	0.000	1.000

Table 5.14: Training on players 1-4 and 6-8, classifying data from **player 5** (overall prediction of $84.7 \pm 14.4\%$) with ten principal components, seven nearest neighbors.

actual \ class.	acc. détaché	det. lancé	louré	martelé	staccato	spiccato
acc. détaché	0.250	0.417	0.333	0.000	0.000	0.000
det. lancé	0.000	0.000	0.000	0.000	0.000	1.000
louré	0.000	0.000	1.000	0.000	0.000	0.000
martelé	0.000	0.000	0.000	1.000	0.000	0.000
staccato	0.000	0.000	0.083	0.000	0.917	0.000
spiccato	0.000	0.000	0.000	0.000	0.000	1.000

Table 5.15: Training on players 1-5 and 7-8, classifying data from **player 6** (overall classification rate of $69.4 \pm 44.9\%$) with ten principal components, seven nearest neighbors.

actual \ class.	acc. détaché	det. lancé	louré	martelé	staccato	spiccato
acc. détaché	1.000	0.000	0.000	0.000	0.000	0.000
det. lancé	0.000	0.667	0.333	0.000	0.000	0.000
louré	0.000	0.000	1.000	0.000	0.000	0.000
martelé	0.917	0.000	0.000	0.083	0.000	0.000
staccato	0.000	0.000	0.250	0.000	0.750	0.000
spiccato	0.000	0.000	0.000	0.000	0.000	1.000

Table 5.16: Training on players 1-6 and 8, classifying data from **player 7** (overall classification rate of $75.0 \pm 35.8\%$) with ten principal components, seven nearest neighbors.

actual \ class.	acc. détaché	det. lancé	louré	martelé	staccato	spiccato
acc. détaché	1.000	0.000	0.000	0.000	0.000	0.000
det. lancé	0.000	0.833	0.000	0.000	0.000	0.167
louré	0.000	0.000	0.667	0.000	0.000	0.333
martelé	0.000	0.667	0.000	0.333	0.000	0.000
staccato	0.000	0.000	0.250	0.000	0.750	0.000
spiccato	0.000	0.000	0.000	0.000	0.000	1.000

Table 5.17: Training on players 1-7, classifying data from **player 8** (overall classification rate of $76.4 \pm 25.0\%$) with ten principal components, seven nearest neighbors.




Figure 5.26: This bar graph shows the technique classification rates for the data contributed by each individual player.

actual \ class.	acc. détaché	det. lancé	louré	martelé	staccato	spiccato
acc. détaché	0.688	0.052	0.042	0.219	0.000	0.000
det. lancé	0.000	0.677	0.052	0.031	0.062	0.177
louré	0.000	0.000	0.896	0.000	0.062	0.042
martelé	0.167	0.156	0.000	0.635	0.042	0.000
staccato	0.000	0.000	0.094	0.031	0.865	0.010
spiccato	0.000	0.000	0.000	0.000	0.000	1.000

Table 5.18: Individual classification rates for each technique achieved with **ten** principal components, seven nearest neighbors. An overall technique classification rate of $79.5 \pm 14.5\%$ was produced for the player-1 case.

Player Classification

is it possible to use the bowing measurement system to recognize not only common bowing techniques, but the individuals who play them?

After investigating the potential of the gesture data collected by the bowing measurement system for use technique classification, the k-NN classification algorithm was employed to explore classification by player.

In this exploration, only one case, the all-player case, was addressed. Just as before in performing the technique classification with the data from all of the players, the k-NN classifier was implemented using two-thirds of the data from *all of the players* to classify by player the remaining one-third of data from all of the players. But this time, the classification was done by *player*. Once again, a 3-fold cross validation procedure, in which the training data set was rotated three times, was obeyed. For each training data set, the number of principal components input to the algorithm was increased from two to twenty. The overall classification results for each number of components used were determined as the mean and standard deviation of the classification rates achieved in each of the three cross-validation trials (for each number of principal components used).

Individual confusion matrices for the 8, 10, 12, 14, 16, 18, and 20 principal components used as input to the k-NN classifier are shown in Tables 5.19-5.25. With 12 principal components, a player classification rate of nearly 90%, 89.12% to be exact, is reached, and with 18 principal components, an overall classification player rate of 94.21% is obtained. Again, the effect of increasing the number of principal components was clearly demonstrated, as illustrated in the plot in Figure 5.27.

actual \ class.	pl. 1	pl. 2	pl. 3	pl. 4	pl. 5	pl. 6	pl. 7	pl. 8
player 1	0.806	0.000	0.056	0.014	0.056	0.014	0.000	0.056
player 2	0.056	0.875	0.042	0.000	0.000	0.028	0.000	0.000
player 3	0.042	0.000	0.917	0.014	0.014	0.000	0.014	0.000
player 4	0.014	0.000	0.014	0.917	0.056	0.000	0.000	0.000
player 5	0.042	0.000	0.014	0.056	0.861	0.000	0.000	0.028
player 6	0.028	0.069	0.000	0.000	0.014	0.889	0.000	0.000
player 7	0.042	0.014	0.042	0.000	0.000	0.000	0.889	0.014
player 8	0.042	0.028	0.028	0.000	0.028	0.000	0.000	0.875

Table 5.19: Training on two-thirds of the data from each of the eight players, classifying the remaining one-third of data, classifying by player (with overall classification of $87.9 \pm 3.5\%$) with **eight** principal components.

actual \ class.	pl. 1	pl. 2	pl. 3	pl. 4	pl. 5	pl. 6	pl. 7	pl. 8
player 1	0.861	0.014	0.028	0.000	0.042	0.014	0.000	0.042
player 2	0.028	0.903	0.028	0.000	0.000	0.028	0.014	0.000
player 3	0.028	0.014	0.931	0.014	0.000	0.000	0.014	0.000
player 4	0.000	0.000	0.028	0.944	0.028	0.000	0.000	0.000
player 5	0.042	0.000	0.014	0.042	0.903	0.000	0.000	0.000
player 6	0.000	0.083	0.000	0.000	0.014	0.903	0.000	0.000
player 7	0.042	0.028	0.028	0.000	0.000	0.000	0.903	0.000
player 8	0.028	0.014	0.028	0.000	0.014	0.000	0.000	0.917

Table 5.20: Training on two-thirds of the data from each of the eight players, classifying the remaining one-third of data, classifying by player (with overall classification of $90.8 \pm 2.5\%$) with **ten** principal components.

actual \ class.	pl. 1	pl. 2	pl. 3	pl. 4	pl. 5	pl. 6	pl. 7	pl. 8
player 1	0.861	0.028	0.000	0.000	0.014	0.014	0.000	0.083
player 2	0.014	0.875	0.042	0.014	0.000	0.056	0.000	0.000
player 3	0.028	0.000	0.944	0.000	0.000	0.000	0.000	0.028
player 4	0.000	0.000	0.028	0.917	0.056	0.000	0.000	0.000
player 5	0.056	0.014	0.014	0.028	0.889	0.000	0.000	0.000
player 6	0.042	0.083	0.014	0.000	0.000	0.861	0.000	0.000
player 7	0.014	0.028	0.000	0.042	0.000	0.014	0.903	0.000
player 8	0.056	0.000	0.069	0.000	0.028	0.000	0.000	0.847

Table 5.21: Training on two-thirds of the data from each of the eight players, classifying the remain third of data, classifying by player (with overall classification of $92.0 \pm 1.8\%$) with **twelve** principal components.

actual \ class.	pl. 1	pl. 2	pl. 3	pl. 4	pl. 5	pl. 6	pl. 7	pl. 8
player 1	0.917	0.014	0.000	0.000	0.042	0.014	0.000	0.014
player 2	0.014	0.889	0.028	0.014	0.000	0.056	0.000	0.000
player 3	0.028	0.000	0.931	0.028	0.000	0.000	0.014	0.000
player 4	0.000	0.000	0.014	0.958	0.028	0.000	0.000	0.000
player 5	0.028	0.014	0.014	0.028	0.917	0.000	0.000	0.000
player 6	0.000	0.028	0.000	0.000	0.000	0.972	0.000	0.000
player 7	0.042	0.028	0.028	0.000	0.000	0.000	0.903	0.000
player 8	0.014	0.014	0.028	0.000	0.014	0.000	0.000	0.931

Table 5.22: Training on two-thirds of the data from each of the eight players, classifying the remain third of data, classifying by player (with overall classification of $92.7 \pm 2.8\%$) with **fourteen** principal components.

actual \ class.	pl. 1	pl. 2	pl. 3	pl. 4	pl. 5	pl. 6	pl. 7	pl. 8
player 1	0.931	0.014	0.000	0.000	0.042	0.000	0.000	0.014
player 2	0.014	0.889	0.028	0.014	0.000	0.042	0.014	0.000
player 3	0.014	0.000	0.972	0.014	0.000	0.000	0.000	0.000
player 4	0.000	0.000	0.014	0.944	0.028	0.000	0.014	0.000
player 5	0.014	0.014	0.014	0.028	0.931	0.000	0.000	0.000
player 6	0.000	0.028	0.000	0.000	0.000	0.972	0.000	0.000
player 7	0.028	0.028	0.028	0.000	0.000	0.000	0.917	0.000
player 8	0.028	0.014	0.042	0.000	0.028	0.000	0.000	0.889

Table 5.23: Training on two-thirds of the data from each of the eight players, classifying the remain third of data, classifying by player (with overall classification of $93.1 \pm 3.2\%$) with **sixteen** principal components.

actual \ class.	pl. 1	pl. 2	pl. 3	pl. 4	pl. 5	pl. 6	pl. 7	pl. 8
player 1	0.944	0.014	0.000	0.000	0.028	0.000	0.000	0.014
player 2	0.014	0.917	0.028	0.014	0.000	0.014	0.014	0.000
player 3	0.014	0.000	0.972	0.014	0.000	0.000	0.000	0.000
player 4	0.000	0.000	0.000	0.972	0.028	0.000	0.000	0.000
player 5	0.014	0.000	0.028	0.014	0.944	0.000	0.000	0.000
player 6	0.000	0.028	0.000	0.000	0.000	0.972	0.000	0.000
player 7	0.028	0.028	0.028	0.000	0.000	0.000	0.917	0.000
player 8	0.028	0.014	0.042	0.000	0.028	0.000	0.000	0.889

Table 5.24: Training on two-thirds of the data from each of the eight players, classifying the remain third of data, classifying by player (with overall classification of $94.1 \pm 3.1\%$) with **eighteen** principal components.

actual \ class.	pl. 1	pl. 2	pl. 3	pl. 4	pl. 5	pl. 6	pl. 7	pl. 8
player 1	0.931	0.014	0.014	0.000	0.028	0.000	0.000	0.014
player 2	0.014	0.917	0.028	0.000	0.000	0.028	0.014	0.000
player 3	0.014	0.000	0.972	0.014	0.000	0.000	0.000	0.000
player 4	0.000	0.000	0.000	0.958	0.028	0.000	0.014	0.000
player 5	0.014	0.014	0.028	0.014	0.931	0.000	0.000	0.000
player 6	0.000	0.028	0.000	0.000	0.000	0.972	0.000	0.000
player 7	0.028	0.014	0.028	0.000	0.000	0.000	0.931	0.000
player 8	0.028	0.014	0.028	0.000	0.028	0.000	0.000	0.903

Table 5.25: Training on two-thirds of the data from each of the eight players, classifying the remain third of data, classifying by player (with overall classification of $93.9 \pm 2.6\%$) with **twenty** principal components.




Figure 5.27: Mean classification rates produced by KNN using two-thirds of the data from each of the eight players to classify the remaining one-third the data, classifying by *player*.

5.2 Virtual Violin Performance Evaluation

The second major study undertaken in this thesis was designed to explore the potential of the bowing measurement system as a performance tool for the control of a real-time virtual violin model.

In order to conduct this study, the calibrated gesture data from the measurement system was input to a Kalman filter (described in Chapter 3) to obtain good estimates of the three primary bowing parameters of bow force, bow speed, and bow-bridge distance. These three control parameters were then input to the physical violin model (implemented in PD), and the resultant synthesized sound was recorded and evaluated.

Two different performance scenarios were tested. The first used the bowing parameter data from an archived performance to drive the model, and the second used the data generated by a violinist in realtime. The results of both of these investigations are described below.

5.2.1 Archived Gestural Performance

Does the virtual violin model respond realistically to calibrated data extracted from an actual performance by a real player?

As described above, this virtual violin experiment relied upon the use of the Vicon motion capture system to provide input bowing parameters corresponding to an actual performance by a real player. The example performance used for this experiment was that of one simple, broad downbow stroke followed immediately by a similar upbow stroke, played on the open G string.

The position data provided by the Vicon motion capture system and the force data from the bowing measurement system were input to a Kalman filter (the force data was simply passed through the filter). This produced a data file containing estimates of the three primary bowing parameters for every time step corresponding to the gestural performance previously recorded. (The Kalman filter updates the parameters at 200 Hz, at the same rate the hardware sensor system produces the raw sensor values.) This data file was then input via Python and OSC to the physical violin model implementation in Pd, so that the model would receive a continuous data stream of the three primary bowing parameters, just as if the model was being played by the violinist in realtime. The resulting audio produced by the model was then recorded.

To the ear, one of the most striking features of the synthesized violin sound was that there was no distinct attack for the second part of the sound file corresponding to the original upbow

stroke recorded by the real player. This observation was verified by plotting the virtual violin waveform, which is displayed in Figure 5.28, along with the same visualization of the real electric violin sound for reference and the plots of the primary bowing parameters (that produced both real and virtual sound). As can easily be seen here, the violin model fails to produce a distinction between the two bowstrokes.




Figure 5.28: This figure shows the waveforms produced by the real electric violin and the virtual violin, and the three primary bowing parameters that produced them both.

5.2.2 Live Gestural Performance

Does the virtual violin model respond realistically to live control by a real player?

The second virtual violin experiment was conducted by driving the physical violin model

with gesture data produced by a violinist in realtime. The purpose of this investigation was to address the question of how well a real player is able to play the model using the measurement system as a controller.

The performance given by this violinist was quite similar to that used in the previous virtual violin performance, which consisted of one downbow and one upbow, played on the open G string. In fact, as previously discussed, the violinist participating in this study was asked to listen to the recording of the real electric violin sound from the archived performance and then asked to try to imitate it. In the event that this task was too difficult, the player was asked to produce the highest quality sound possible with the same downbow-upbow bowing as in the example.

After listening to the example, the violinist attempted to imitate both the sound quality and timing of the recording. However, after several attempts, the violinist abandoned this approach and continued according to the second instruction, trying to coax the best sound possible out of the model with bowstrokes of different lengths and speeds.

Flawed Calibration

Interestingly, in the first trials of this experiment featuring live control of the model, the bow-bridge distance was improperly calibrated. The parameter expected by the model respects the same convention as used in acoustics, in which bow-bridge distance is a unit-less number scaled to the length of the fingerboard (rather than the value in meters given by the Kalman filter). However, this adjustment was neglected at first and the value in meters was supplied directly by the Kalman filter. Figure 5.29 shows a result from this experiment. (Unlike the velocity update rate of the Kalman filter, which is equal to the sampling rate of the hardware sensing system, the Kalman filter position update is slower than the sampling rate of the hardware sensing system. Because of this, the plot of the bow-bridge distance shows a staircasing effect.)

Proper Calibration

After the error in the calibration of the bow-bridge distance was discovered, it was quickly corrected, and additional experimental trials were conducted. Figure 5.30 shows a result from a performance by the same player.




Figure 5.29: This figure shows the waveforms produced by the real electric violin, and the three primary bowing parameters that were input to the model by the player when the calibration of the bow-bridge distance parameter (β) was *incorrect*, since it is calibrated in meters rather than scaled to the length of the distance between the bridge and the nut of the violin.




Figure 5.30: This figure shows the waveforms produced by the real electric violin, and the three primary bowing parameters that were input to the model by the player. In order to obtain a distinction between the two bowstrokes, the player made a significant pause of about one second between the downbow and the upbow.

Chapter 6

Discussion

The previous chapter of this thesis presents results from two main studies related to bowed string performance using the playable, calibrated measurement system created here. The first of these is the bowing technique study, and the second is a study that explores the performance of a violin physical model, or virtual violin. This chapter discusses the results of these two studies, as well as related topics, such as possible sources of error, the violin participants, and player feedback.

In addition, the bowstroke database, which contains all of the bowing data collected in the two studies, is described.

6.1 Technique Study Results

The evaluation of the technique study data began with a simple visual inspection (to confirm the consistency of the samples among same techniques) of a subset of the gesture data (downward and lateral forces, 3D acceleration, 3D angular velocity of the bow) recorded in the data collection sessions with the violinist participants. After a small number of outlying samples were removed from the data set (4%), this data set was then prepared for more quantitative evaluation.

As described in the previous chapter, the SVD method was used to determine the principal components (thereby reducing the dimensionality of the data) of the bow gesture data, enabling new visualizations of the data. Scatter plots depicting the separability of the data by bowing technique were generated showing all twelve examples of each of the six techniques contributed by each of the eight players. Also, for each of the bowing techniques, a scatter plot depicting the

separability of the data by player was generated. These explorations were quite encouraging, and foreshadowed the successful classification rates (by technique and by player) of over 90% that were later achieved using the k-NN algorithm.

6.1.1 Technique Classification

The bowing technique study was designed to answer the question: *Is the bowing measurement system capable of capturing the distinctions between common bowing techniques?*

Three cases of technique classification were performed. In the first two, the Intra-Player and the All-Player cases, success rates of well over 90% were attained. In the third one, the Player-1 case, a success rate of only approximately 80% was reached. However, this result can to a large extent be understood by examining the individual results for data corresponding to each of the eight players.

For instance, Table 5.15 shows that 100% of the gesture data provided by player 6 when asked to perform *détaché lancé* was not classified as such, but rather was classified as *spiccato*. However, a quick inspection of data from another player, such as those of player 2, produces a simple explanation. Figure 6.1 shows two waveforms produced by player 2 that correspond to gesture data that was correctly classified. As can be seen from the plots, the differences in the amplitude and the shape are significant. Figure 6.2 shows a waveform corresponding to one of the gestures produced by player 6 when asked to perform *détaché lancé*. As is apparent from listening to this example, when comparing the waveform to the two examples in Figure 6.1, it is clear that it bears a greater similarity to a typical example of *spiccato* than to a typical example of *détaché lancé*.

6.1.2 Player Classification

While the primary goal of the bowing technique study was to investigate the potential of the measurement system to provide sufficient information to enable gesture recognition of traditional violin bowing techniques, in the course of this work it was discovered that not only is the system useful for this, but it also offers the ability to recognize players as well.

This of course was a welcome surprise, and suggests that this measurement system may be used to make quantitative comparisons between the techniques and styles of individual performers.




Figure 6.1: This figure shows the audio waveforms from two different recordings provided by player 2 for the bowing technique study that were properly classified. The top plot corresponds to an example of *détaché lancé*, while the bottom plot corresponds to a *spiccato* recording.




Figure 6.2: This audio waveform is from a recording provided by player 6 when asked to perform the *détaché lancé* technique. The gesture corresponding to this audio was seemingly incorrectly classified as *spiccato* technique. However, as can be seen by comparing to Figure 6.1, this audio waveform bears more resemblance to a typical example of *spiccato* than to a typical example of *détaché lancé*.

At first, it seemed a great disadvantage not to have “novice” players as participants to the study. However, it soon became apparent that the tasks demanded of the players, though simple, were not necessarily easy. Many participants commented on their surprise at finding the exercises challenging, and it was clear that despite their skill and experience, certain bowing techniques or bowstrokes were difficult to produce if not a part of the player’s recent regular practice. Therefore, the understanding of “novice” versus “expert” has evolved in this research to describe bowing techniques or bowstrokes, rather than to identify the players themselves, for violinists of a very high expert level may still produce bowstrokes of novice quality, and vice versa.

6.2 Virtual Violin Study Results

The virtual violin study, though limited in scope, enabled two important discoveries by considering two different cases for driving the physical model with real player gesture data. The first experiment was designed using data from an archived gestural performance and addressed the question: *Does the virtual violin model respond realistically to calibrated data extracted from an actual performance by a real player?* Conducting this experiment made it possible to test the response of the model to calibrated bowing parameters and thereby enabled a significant observation concerning the damping behavior of the violin model, discussed below.

The second experiment was designed to use the bowing parameters generated in realtime by a violinist (monitoring the virtual violin sound with earphones) to drive the violin model, in order to address the question: *Does the virtual violin model respond realistically to live control by a real player?* Though this experiment was conducted with only one violinist participant, it was very instructive, as the violinist was able to control the model with the playable measurement system to easily produce a better sonic result than was produced using the archived gesture data, even when the calibration and/or the model was flawed.

6.2.1 Archived Gestural Performance

The results from the first virtual violin performance experiment, in which the physical model was driven with the gesture data from an archived performance on a real instrument, revealed one very significant point concerning the behavior of the physical model. Though the gestural performance used to drive the model was that of one downbow stroke followed by one upbow

stroke, the synthesized violin sound produced failed to exhibit the distinction between the two strokes that was obvious to the listener in the original real audio.

This qualitative observation was confirmed quantitatively, through the analysis detailed in the previous chapter. As seen in the waveform of the virtual violin audio displayed in Figure 5.28, the absence of the second attack (of the upbow stroke) is clear. This analysis indicated that the physical model implementation had an improper response when the speed of the bow was zero (as was the case at the moment that the bow changed directions between the downbow and upbow strokes) and the bow force was nonzero.

In order to validate this assertion, a simple test was constructed. The model was once again driven with archived gesture data, but this time the data were altered so that the values of the three bowing parameters were held constant for a duration of one second from the moment of zero bow speed between the two bowstrokes. The results of this investigation are shown in Figure 6.3. As can be seen, a considerable length of time passes (well over half a second) before the synthesized audio dissipates. Therefore, it appears that this implementation of the physical violin model does not exhibit proper damping to emulate that which is typical to its real instrument counterpart (as seen in the waveform for the real electric violin in Figure 5.28).

One final investigation was conducted in order to better understand the results of this virtual violin experiment. Using the equations that govern Helmholtz motion, the bowing parameters of bow speed and bow-bridge distance were used to compute the maximum and minimum limits for bow force that would be expected to produce “good sound”. Once again, these two equations are:

$$F_{\max} = \frac{2Z_0\nu}{\beta(\mu_s - \mu_d)} \quad (6.1)$$

$$F_{\min} = \frac{Z_0^2\nu}{2\beta^2R(\mu_s - \mu_d)} \quad (6.2)$$

where μ_s and μ_d are the static and dynamic friction coefficients, respectively. Z_0 is the characteristic impedance of the string and R (often referred to as the dashpot rate of the bridge) indicates the equivalent of the rate of energy loss into the violin body.

The violin model implementation uses $\mu_s = 0.8$, $\mu_d = 0.35$ and $Z_0 = 1/0.55$. Therefore, using the values for β (bow-bridge distance) and ν (velocity) of the archived gesture performance used to drive the mode, the maximum force curve was easily generated using Equation




Figure 6.3: This figure shows the effect of holding the bowing parameters constant for a duration of 1 second, at the point when the speed is zero at the moment of the change in bow direction. As can be seen in the audio waveform produced, it takes almost the entire added second for the amplitude of the sound to decrease to zero, indicating that the violin physical model does not produce realistic damping behavior.

6.1. However, in order to calculate the corresponding minimum force curve, R was also needed. The model itself was not designed with an explicit value for R . Therefore, a realistic value for the dashpot rate for a real violin was sought elsewhere. Using data from an experiment concerning minimum bow force for violin playability conducted by Woodhouse and detailed in [88], R was estimated as follows.

From Figure 3.a of [88], which shows a plot of minimum bow force with respect to frequency, the value for minimum bow force was estimated to be 0.17 N for the fundamental frequency, 196 Hz, of the sound produced by the example bowstrokes in the recorded gesture file. Taking Woodhouse's values, $\mu_s = 0.8$, $\mu_d = 0.3$, $Z_0 = 1/(3.1\text{mN}^{-1}\text{s}^{-1})$, $\beta = 0.1$, and $\nu = 0.5\text{m s}^{-1}$, and substituting into Equation 6.2, R is then calculated as 30.6. Now, using this R and the values for μ_s , μ_d , Z_0 , β , and ν employed above to determine the maximum bow force

curve specific to this virtual violin experiment, the minimum bow force curve can be calculated.

These maximum and minimum force curves are plotted with the actual bow force (and speed and bow-bridge distance) used in this experiment in Figure 6.4. As can be seen, the bow force parameter used is not always within the playability limits for Helmholtz motion. During both the beginning and the end of the performance, the bow force value drops below the minimum value dictated by the above calculations. However, these occurrences are not surprising, as the violinist who performed the bowstrokes in question was at the time playing the real electric violin and responding only to the sound of this instrument. In Figure 6.4, closeups of the virtual violin audio waveform (included under the full audio waveform) at two different times (notated on the full waveform) can also be seen. These closeups illustrate the motion of the virtual string when the three primary bowing parameters are within the correct ranges to achieve Helmholtz motion.

It should be remembered that the specific value for R indicated by Woodhouse's data reflects the minimum bow force limit for the acoustic violin used in those experiments. It was taken here as a typical value in the absence of a value for R specific to the physical model. The absence of an explicit R (which represents all losses through through the violin body) may be one explanation for the lack of realistic damping of the violin sound observed in this virtual violin experiment.

6.2.2 Live Gestural Performance

The second virtual violin experiment, in which a violinist was given live control of the model while monitoring the synthesized sound through earphones, was quite interesting. As shown and discussed above, the violinist playing the model in realtime was able to achieve a more convincing performance of the two-bowstroke playing task, due to the production of a distinct attack for the second stroke not present in the sound produced using the archived gesture data.

The success of this live performance serves as a clear illustration of the ability of the violinist to quickly adapt to the constraints of the virtual violin (as any musician would with a new real instrument) in order to obtain better performance than created in the offline comparisons. Because of the innate ability of skilled musicians to compensate for limitations in their instruments, it is not surprising that the virtual violin performances produced in this thesis are more favorable in real-time playing demonstrations than in demonstrations relying on archived data.

The bowing parameters used in these live performances are further analyzed using the same

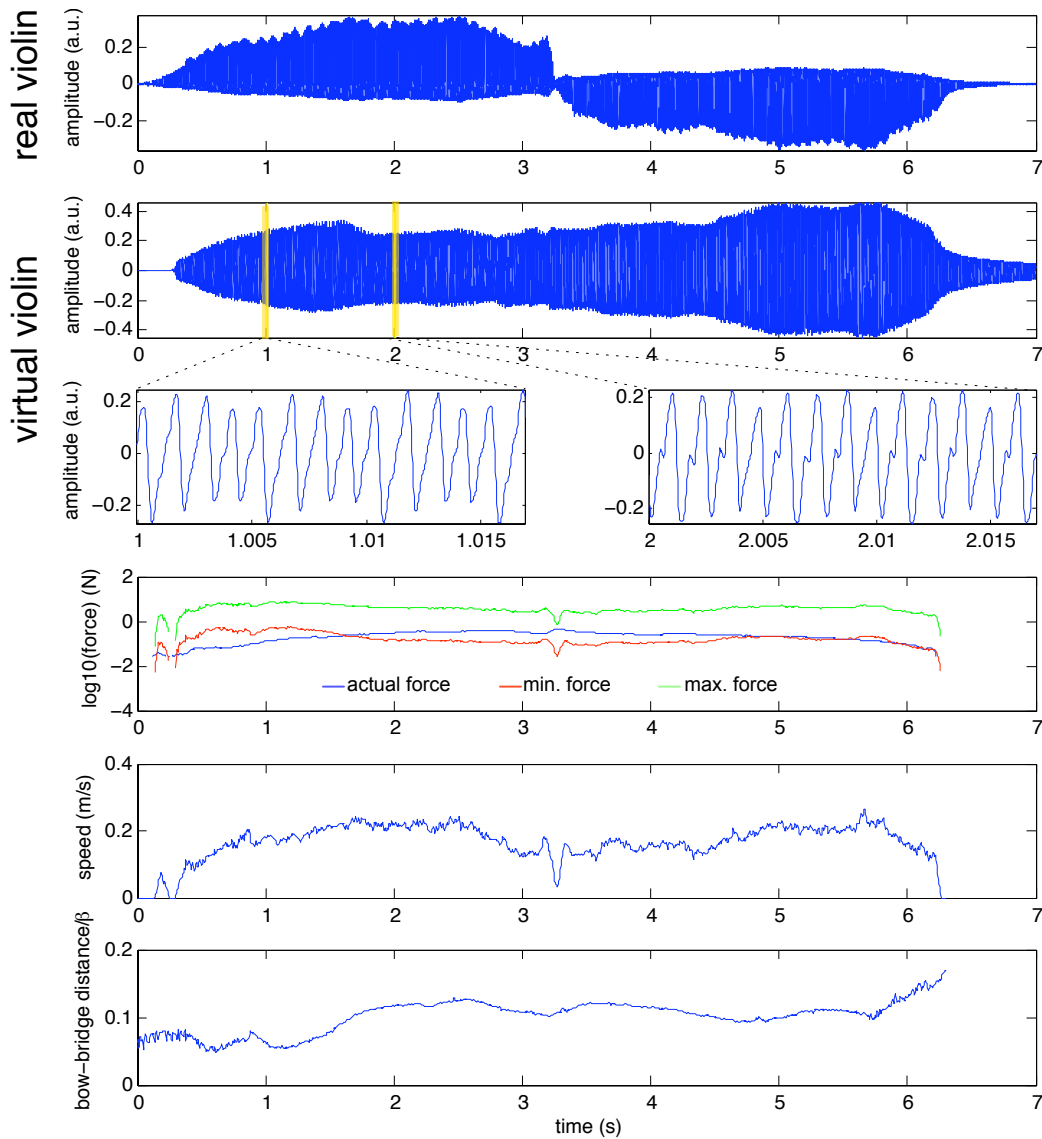


Figure 6.4: This figure indicates the maximum and minimum limits for bow force, as determined by the values of bow-bridge distance and speed of the performance, given the values for μ_s , μ_d , Z_0 specific to the violin model and R calculated from data collected in an experiment by Woodhouse using a real acoustic violin [88].

approach as above. Given Equations 6.1 and 6.2 and the value of R determined from the data in [88], the maximum and minimum force curves are calculated for the live playing case in which the bow-bridge distance (β) was improperly calibrated as well as the live playing case in which this parameter was properly calibrated. The former is shown in Figure 6.5 and the latter in 6.6.

Also included in Figures 6.5 and 6.6 are two closeups (displayed under the full waveform)

of the waveform produced by the virtual violin model in response to the live control parameters. In each of the two cases, an example was chosen at a point in the performance when the three primary bowing parameters were within the playability region described by Equations 6.1 and 6.2 (that is, the bow force was within the limits described by the maximum and minimum force curves calculated). A second example was also chosen at a point in the performance when the bowing parameters were not within the playability region.

As can be seen from these closeups of the waveforms, the model does appear to respond appropriately to the input parameters, as the shape of the “sawtooth” signal is more distinct and periodic when the bowing parameters are within the playability region.

6.3 Sources of Error

Because the bow hardware was installed in such a way to make it as light as possible and easily repairable, no enclosure was made for the electronics, nor were the electronics permanently adhered to the frog. Therefore, the bow board was fixed to the frog of the bow by means of tape and adhesive putty. It is possible that this mounting scheme enabled the board, and thus the IMU, to shift slightly throughout individual performances as well as between adjustments of the bow hair tension.

Effort was made to maintain the same audio levels of the input sound of the test violin heard by the player through headphones, and all players were strongly encouraged to imitate the sound volume as accurately as possible. However, it is possible that differences in sound level could have contributed to a distortion in the results of the bowing technique study.

6.4 Player Feedback

Due the increased weight and altered bowing point of the bow as a result of the sensor augmentation, it was expected that the violinists involved in these studies might experience some difficulty performing all of the tasks to their usual standards. However, throughout the recording sessions it was apparent that although the bow may have felt heavy to the user when first picked up and handled, once it was in motion in normal playing the weight and bowing point were still quite manageable. With very brief and limited warmup and experimentation, all of the players (even those who initially commented on the different ergonomics of the system) were able to

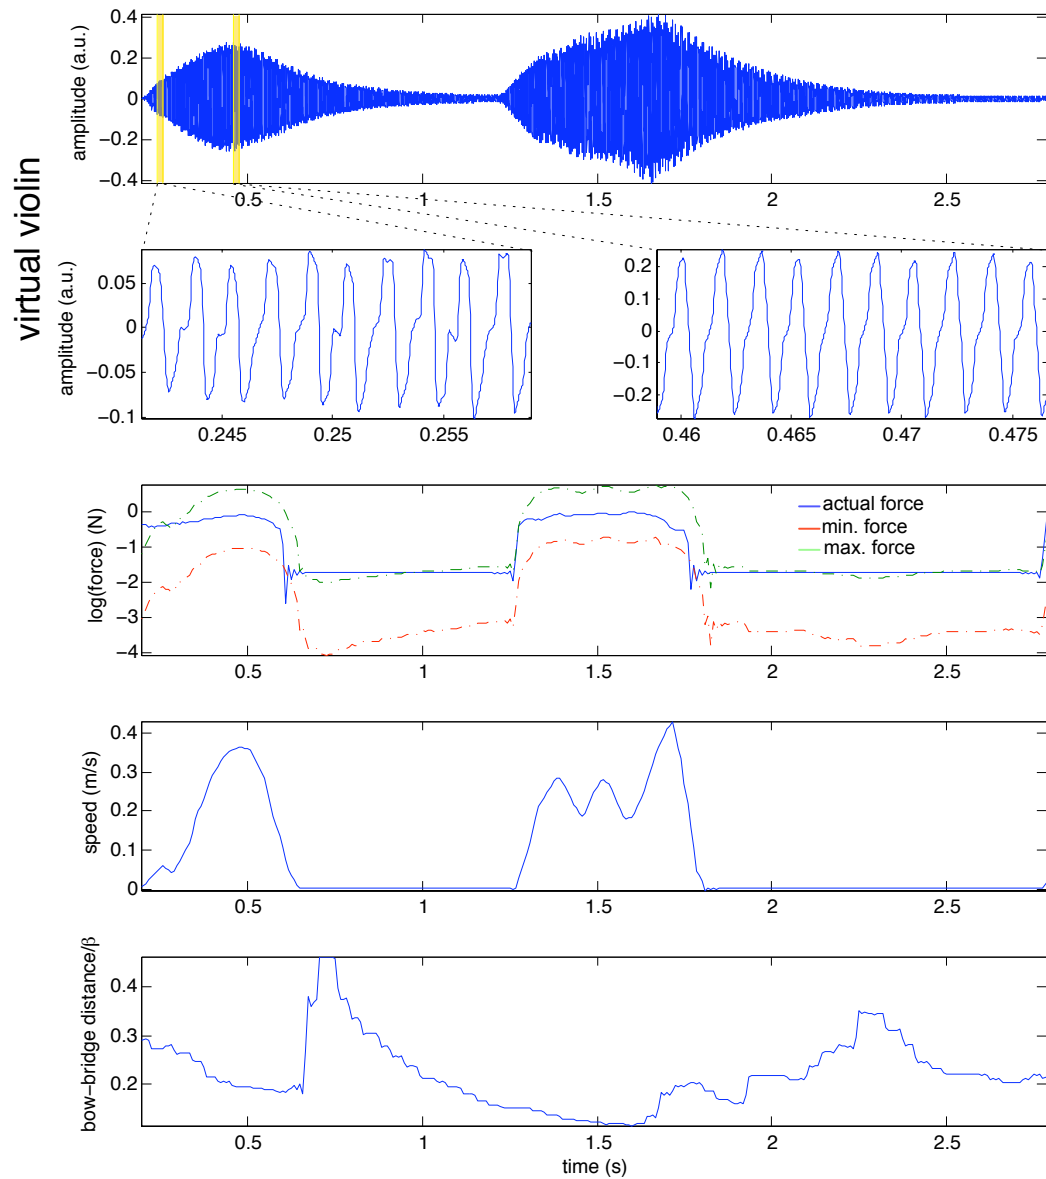


Figure 6.5: This figure shows the maximum and minimum limits for bow force as determined by the values of bow-bridge distance (β) and speed for the live playing case in which the β parameter was not properly calibrated.

achieve the bowing required without obvious difficulty. (This observation was the same as that made in previous performance work with the Hyperbow, discussed in Chapter 2 as well as in [94, 93].)

Interestingly, the study participants reported greater inconvenience overall caused by the audio recording system than by the gesture recording system. That is, the heavy audio cable (which most players draped over the left shoulder) and headphones presented a situation that

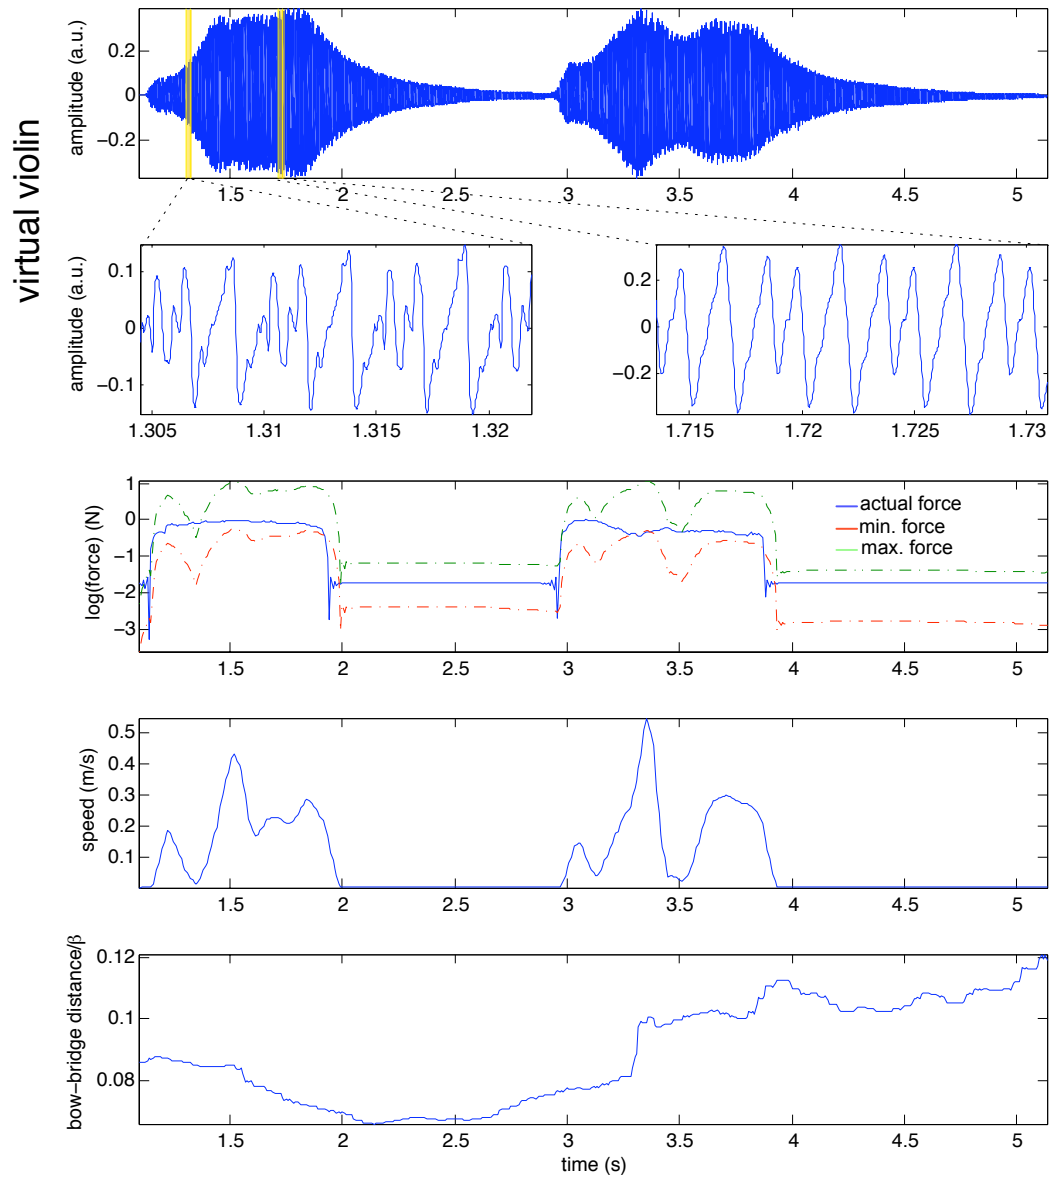


Figure 6.6: This figure shows the maximum and minimum limits for bow force as determined by the values of bow-bridge distance (β) and speed for the live playing case in which the β parameter was properly calibrated.

was quite foreign to most players included in these recording sessions, as did the use of an electric violin. The overall weight and weight distribution of the Yamaha SV-200 Silent Violin (in contrast to those features of a typical acoustic violin) were the subjects of some comment. However, there was also quite a bit of enthusiastic discussion on the novelty of playing this instrument.

Other than issues related to weight, there were two ergonomic deficiencies that produced

some slight difficulty for players. They were: the presence of the electric field position antenna array, which is directly in the line of sight of the bow-string contact point; and the height of the outer frame of the Yamaha SV-200 Silent Violin, which for a few players presented an obstacle to bowing on the higher strings near the fingerboard. The ergonomic issue presented by the height of the position antenna would not be present in future implementations for cello or double bass (due to the position of the player with respect to the bridge), but will be addressed in future violin and viola implementations.

The audio quality of the electric test violin was clearly high enough for the technique study included in this thesis work. Some of the violinists who had never played an electric violin did complain about the difference between its audio output and that of an acoustic violin. However, given the difficulties that an acoustic test violin would have presented to the recording setup, and the more dramatic physical designs of other available electric violins, the choice of the Yamaha SV-200 Silent Violin still appears to have been a good one for these studies.

6.5 Bowstroke Database

The studies conducted in the course of this thesis research have generated a large body of experimental data, more than was previously collected regarding physical violin bowing technique of real players. (As already mentioned, previous work involved much smaller groups of participants and fewer studies.) So far, the recording sessions conducted for the bowing technique study alone have generated well over 1000 experimental audio and gesture data files. In addition to these, many video recordings (and example screenshot is shown in Figure 6.7 and Vicon motion capture recordings were also made throughout this thesis work.

It was critical to establish a reliable method to store and backup these data collected in the experimental sessions already conducted as well as those to come. It was also clear that analyzing these data would be an iterative process and that any given file might need to be revisited and examined many times. Therefore, organizing these data to improve accessibility for both short and long term use was extremely important. The accessibility requirements became more complicated with the prospects of continued collaborative research on the interaction of bowing parameters in bowstrokes generated by real players and one the development of a virtual violin (complete with realistic physical controller) grew. It was obvious that progress would be faster if the exchange of data, both gesture and audio, was made easy. In considering this scenario,



Figure 6.7: Using the recently created playable measurement system developed, a large amount of audio and calibrated gesture data, as well as video, was collected in the experimental recording sessions. These data are archived in the bowstroke database.

the possibility that other independent researchers might also benefit from such a resource of bowing data was considered.

There are several recent research projects that address the issue of archiving and accessibility to experimental data. Most notably are the Rutgers MCL Interaction Capture Repository [54] and the CMU Graphics Lab Motion Capture Database [15].

The Rutgers MCL Interaction Capture Repository is a web-based resource containing force capture and motion capture data corresponding to various hand movements. Users may download these data corresponding to sixteen different manual manipulations such as “index finger surface scratching” or “index finger pad presses”. As stated by the authors, the primary purpose of this website is to further research on interaction capture and synthesis [34] by making their data available to other researchers.

The CMU Graphics Lab Motion Capture Database, another web-based resource, contains 2642 trials in 6 different motion categories: human interaction, interaction with environment, locomotion, physical activities and sports, situations and scenarios, and test motions. The database provides motion capture data files corresponding to each of the trials, as well as video files and animations. Unlike the Rutgers repository described above, the data on this website is searchable by subject number or keyword.

As part of the work in this thesis, a similar tool was created. This web-accessible, searchable

Bowstroke Database [97] was designed to serve as an archive of the bowing data collected in this research, as well as to assist in future research. (This database was completed with the help of Anagha Deshmane as part of her 2006 Undergraduate Research Opportunity (UROP) with the Opera of the Future group at the Media Lab.)

6.5.1 Design Goals

The first goal of this project was to create a *secure and reliable* archive for the experimental gesture and audio data collected using the playable measurement system for violin bowing (evolved from earlier Hyperbow designs) described above. However, due to the initial amount of collected data and the intention to increase the size of this collection, it became clear that in order to provide *fast access* to these data, the archive should be *searchable*.

As the bowing technique and physical model performance studies described in [96] progressed, other design priorities emerged, such as providing the ability to maintain not just the original bowing data but also *store related data*. For instance, when using the data to drive a bowed string model, it is advantageous to be able to store any synthesized audio produced with the original gesture and audio data, as well as any related analysis files. Further, in order to better navigate the data and keep track of any progress, the ability to annotate these files with *comments* and *ratings* became desirable.

In addition to providing an easy way for the authors to access these bowing data, it was also hoped that this database could assist other researchers. Therefore, in order to facilitate continued research by the authors and as well as by colleagues, it was important to enable *remote access* to the database. Because of this, the decision was made to make the database web-accessible, like the Rutgers and CMU projects described above. However, not only was it important to provide *downloading* functionality, but supporting easy *uploading* was also a priority. This was due both to the fact that bowing data are most often recorded in remote locations as well as to the goal of encouraging other researchers to share their discoveries made using the database. In this way, the Bowstroke Database stands in contrast to the Rutgers and CMU projects.

Because it was expected that multiple users would contribute to the database, it was important to enable each individual user to administrate his or her own files by setting *file permissions* to specify access by other users.

All of these design goals were implemented using MySQL® [45] and PHP [58] in the format described below order to create not just a data archive, but also a viable research tool for

collaboration between researchers interested in bowling.

6.5.2 Database Design

Because the body of data contained within the Bowstroke Database is quite large and is expected to grow continuously with future research, the organizational design of this resource was critical. While maintaining the design goals described above, a number of features were implemented in order to provide good accessibility to the data as well as to encourage interaction and collaboration between fellow researchers. These features are organized into the Stroke Group, Database Search, File Upload, My Files, and My Bookshelf pages of the database website.

Stroke Group

The most organizational element most central to the design of the Bowstroke Database is the Stroke Group. For each bowstroke recording (gesture and audio data) in the database there is a corresponding Stroke Group. The Stroke Group includes all files related to its original bowstroke recording, and these can be viewed on the Stroke Group page, which serves as a “homepage” for the original data.

These related files may be of several different types: original recorded data, analysis, and synthesis. Original recorded data files are divided into the categories of audio, gesture, audio and gesture (bundled together), video/Vicon motion capture. Analysis files may include audio analysis and gesture analysis produced from any of the original recorded data listed above. Finally, the synthesized audio file category is included for audio that has been synthesized by computer models using the original recorded gesture data.

To further encourage research collaboration and to improve navigation through the data, the features of user comments and ratings have been enabled on the Stroke Group page. Ratings apply to the Stroke Group as a whole, while comments allow collaborating researchers to discuss specific files in detail.

Finally, the Stroke Group page is the only place in the database where users can preview (audio and video files may be played within the browser window) and download files. Each Stroke Group page has its own BibTeX entry that may be used to reference the original recorded gesture and audio data so that researchers may document the data they use.

An excerpt of an example Stroke Group page is shown in Figure 6.8.

VIOLIN BOWSTROKES DATABASE - Stroke Group sub100trial3s3_10_09_2006You are logged in as **diana**.[Home](#) - [My Files](#) - [My Bookshelf](#) - [Upload files to Database](#) - [Search the Database](#) - [Logout](#)**Average User Rating: 0 Subject: 100 Keywords: G string, Timbre B**[Bibtex Reference](#)[Audio Details](#) | [Download This Audio File](#) | [Add This Audio File to My Bookshelf](#)**[Download Hyperbow Audio/Gesture Data](#)**[Details](#) | [Add Audio/Gesture bundle to My Bookshelf](#)

There are no individual gesture sensor files for this stroke group.

Hyperbow Video:

Preview	Comments/Description by Owner			
T3S3.mov	<i>Recorded using Vicon motion capture system (Stata Center, MIT)</i>	Details	Download	Add Video to My Bookshelf
T3S3raudio.mov	<i>This is a motion capture movie recorded in the Vicon motion capture studio, including the real audio produced by the player</i>	Details	Download	Add Video to My Bookshelf

Analysis Files:

Preview	Comments/Description by Owner			
vv-Vicon-filt_T3S3.wav	<i>This is synthesized audio (filtered) from the violin model serafinvl~, which was driven by parameters from the Vicon motion capture system and Hyperbow force.</i>	Details	Download	Add to My Bookshelf
vv-Vicon-unfilt_T3S3.wav	<i>This is synthesized audio (unfiltered) from the violin model serafinvl~, which was driven by parameters from the Vicon motion capture system and Hyperbow force.</i>	Details	Download	Add to My Bookshelf

Figure 6.8: This screenshot shows one of the Stroke Group pages. From this view, the audio and gesture data for this stroke may be downloaded. Also, any available related files, such as video, analysis, or synthesis files may also be downloaded or added to the My Bookshelf list.

Database Search

The Database Search function performs a keyword and filename search of all the files in our database. It can take multiple search terms as input in order to return a more narrow set of data. Search results are presented as a list of Stroke Groups that contain files that match the search parameters. Also displayed in this view are the initial comment made by the owner of the original bowing data for each Stroke Group, and the average user ratings. In addition, keywords for each Stroke Group are listed. These can be used to quickly refine the search results.

Because audio is a simple way to quickly preview and sift through the information contained within a Stroke Group, each Stroke Group is returned with its original audio file for quick previewing while still in the Database Search view. Users can also sort the search results by the name of the Stroke Group, the Stroke Group rating, keywords, and subject number.

Figure 6.9 shows the results of an example search.

File Upload

The File Upload feature of the database allows users to upload up to ten files at once. In order to easily identify them later, the user can specify corresponding keywords as well as a subject number for each file. The same interface can be used to upload .zip files. Each file contained in the .zip folder will take the same subject and keywords as those specified for the .zip folder. The subject and keywords can later be modified through the My Files view.

My Files

The My Files database view acts as a control panel for file information. From this mode the user can assign file categories and file permissions to his or her files in bulk.

Each file also has a File Details page, which allows the user to edit the file's keywords, Stroke Group affiliation, file owner's comments, and permissions settings (which allow researchers to specify which other users do or do not have access to the file's content).

In order to simplify navigation through the database, special attention has been given to each user's previous use of the database. In particular, the list of possible Stroke Group labels in the File Details menu (used to assign newly uploaded files to existing Stroke Groups) has been compiled from the list of Stroke Groups from which the user has previously downloaded data.

VIOLIN BOW STROKES DATABASE - Search

You are logged in as **diana**.

5 groups found related to "sub31>detache>G string"

III

Preview	Stroke Group Δ ∇	Rating Δ ∇	Keywords Δ ∇	Subject Δ ∇	Owner Δ ∇	Owner's comment
	sub31trial26s2a_09_04_2006 4 related files in Stroke Group	4	Accented Detache, G string	31	diana	<i>quarter=168, 4/4 time, 4 bars</i>
	sub31trial38s1_09_04_2006 2 related files in Stroke Group	3	Detache Lance, G string	31	diana	<i>quarter=168, 4/4 time, 4 bars</i>
	sub31trial38s2_09_04_2006 2 related files in Stroke Group	0	Detache Lance, G string	31	diana	<i>quarter=168, 4/4 time, 4 bars</i>
	sub31trial26s2b_09_04_2006 1 related files in Stroke Group	0	Accented Detache, G string	31	diana	<i>quarter=168, 4/4 time, 4 bars</i>
	sub31trial38s3_09_04_2006 1 related files in Stroke Group	0	Detache Lance, G string	31	diana	<i>quarter=168, 4/4 time, 4 bars</i>

III

Below you can search for files uploaded to the database. To search for multiple terms, separate phrases with > only.
Example: "detache>E string"

[Home](#) - [My Files](#) - [My Bookshelf](#) - [Upload files to Database](#) - [Search the Database](#) - [Logout](#)

Figure 6.9: This group listing is the result of a search using “sub31” (subject number 31), “detache”, “G string” as search terms. As can be seen from this screenshot, in addition to information such as user ratings and comments, the main audio file associated with each Stroke Group can be played from this view.

My Bookshelf

To facilitate browsing a simple bookmarking function has been included. The My Bookshelf feature is a tool to help researchers navigate the database. When a user finds a Stroke Group that contains some interesting files, he or she can add those files to his or her bookshelf.

The My Bookshelf page is easily accessible from all of the database pages and contains a list of files that have been bookmarked. Along with the name of the file, it lists the file's keywords, category, and a link back to the Stroke Group to which it belongs.

6.5.3 Usage to Date

To date, the database has served its initial purpose of providing reliable storage and organization for the data collected in this thesis work, as well as creating an easy vehicle for communication with several colleagues. It is expected that these colleagues will become more frequent users, and other researchers will be encouraged to use this resource.

The design of the database was created with the hope that it would someday be of use to other researchers interested in bowed string acoustics, gesture, and synthesis. The database may serve as a valuable research tool that may enables researchers in related fields to share results easily, by uploading audio, analysis figures, gesture data, et cetera.

In addition, it is possible that this research tool will also be of use to composers, who may use the gesture data available in the database to assist in the development of musical mappings for interactive compositions, especially when workshop time with live players is limited.

Chapter 7

Concluding Remarks

This thesis presents an approach to studying bowing technique that is based on playable measurement. To test the viability of this approach, the topic of violin bowing technique was studied in a series of bowing experiments designed to discriminate between different bowing techniques and compare the approaches of different players. Also, in collaboration with Julius Orion Smith III and Stefania Serafin, a virtual violin performance study was undertaken using gesture data captured in realtime by the measurement system.

Below is a discussion of the contributions of this thesis and future work.

7.1 Contributions

The main contributions of this thesis are summarized as follows.

Playable Calibrated Measurement System

For this thesis work, a playable measurement system for violin bowing was designed and constructed, and then calibrated using using state of the art equipment (Instron® force measurement machine, and Vicon motion capture system). This measurement system provides real-time estimation of primary bowing parameters (bow-bridge distance, bow force, and bow velocity) and 3D orientation, enabled by the implementation of a real-time Kalman filter. (The estimation of the bow-bridge distance parameter was also improved by implementing a new approximation recently published in the physics literature). Because the ergonomics of this measurement system are also optimized to enable natural, unencumbered performances by traditionally trained violinists, this system presents many new opportunities to examine the subtleties and nuances

of real violin performances. Using this system, which incorporates a Yamaha SV-200 Silent Violin, the gesture and audio data may be easily captured and recorded together using a simple Pd patch.

Demonstrated Use for Gesture Classification

The usefulness of this playable measurement system for gesture classification has been clearly demonstrated. It has been shown that the gesture data captured by this system can be used to classify different violin bowing techniques with high success rates (well over 90%). In addition to this finding, high success rates (well over 90%) were also achieved for classifying the gesture data by player.

Demonstrated Use in Performance of a Violin Model

Because the system built for this thesis is not merely a calibrated measurement system and is in fact also highly playable, it can be also used as a real-time controller for a calibrated physical model of a violin. This application was clearly demonstrated in the experiment involving live performance of the model by a real player, in which the player was able to produce articulated bowstrokes, and, in doing so, was able to compensate for a deficiency in the model, its lack of realistic damping. The discovery of this deficiency, which was made while testing the response of the model to real data archived from a previous performance (not live) on the real electric violin and was made possible by this measurement system, represents a significant finding and will facilitate future improvements to the violin model. It is also likely that additional improvements may be facilitated by the availability of bowing parameters provided by the measurement system developed in this thesis.

Bowstroke Database

While developing an archiving protocol for the immense amount of audio and gesture data collected throughout the course of this thesis work, a searchable, web-accessible database [97] was created to store all of these data. In its current form, this database serves the purpose of not only storing the original audio and gesture data files recorded in the study sessions, but also any related data file, such as video, motion capture, gesture and audio analysis, and audio synthesis files. In addition to allowing users to download data, the database interface also allows users to upload related files and add comments and ratings to assist other users and further the collective understanding of the bowing information contained.

7.2 Future Work

A great amount of future work remains to be done in order to extend that achieved in this thesis. In considering each of the contributions made in this thesis (discussed above), much can be done to improve or extend each.

The playable measurement system will be improved by revising the Kalman filter already developed to include the sensor data from the violin IMU, in order to enable relative tracking of the bow with respect to the violin. Though the electric field position sensing method has worked well in this and past projects, electric field methods have only approximate solutions. Other position measurement modalities that can measure parameters directly, such as optical methods, will be explored. A new implementation with an optical sensor for measuring bow-bridge distance might be simpler to build and calibrate, and will likely be smaller than the position antennae currently mounted behind the bridge of the test violin (which can interfere with a player's sight of the strings and bow). Also, the possibility of measuring transverse bow position with an array of force sensors (composed of strain gauges) along the length of the bow stick will also be explored. The measurement system will also be refined and slightly revised in order to pursue related study of cello, viola, and double bass bowing. Finally, in order to create a playable measurement system capable of capturing a complete gestural violin performance, left hand technique must also be captured. Future versions of the measurement system will therefore incorporate sensing for left hand finger movement and force.

The bowing technique study produced some interesting results, for both classification of bowing techniques as well as classification of individual players. Of course, many more studies must be done to explore these questions concerning differences in technique and individual approaches to the same techniques. These future studies will include more players, more techniques, and include different dynamics. Some day, it should be possible to use gesture classification to assist in future pedagogies. Since it is possible to classify different techniques from each other, with more work it should be possible to classify the quality of different performances of these techniques. For instance, in the confusion matrices included in the results of the bowing technique study, a performance of a techniques that is not included in the diagonal is one that has not been classified as expected and is likely not the examples one should choose to teach a student how to perform that particular technique.

One day, a teacher might be able to use a classification tool using gesture data only to

evaluate the quality of a given bowstroke. Because bowed string sound may be estimated with bowing parameter data, gesture measurement such as that used here may possibly be used for prediction of audio. For instance, in principle, the “real player” Schelleng diagram resulting from a playability study might provide the ability to diagnose good tone through gesture analysis alone. Such a result would make possible simple devices to monitor the quality of the tone produced by student beginners without the need to process the audio produced, or even for the player to produce sound at all. (The gesture data have a much higher compression ratio than the audio data – $2400 \text{ Hz} \times 12 \text{ bits}$ is much better than $44.1 \text{ kHz} \times 16 \text{ bits}$.)

Other bowing studies may be conducted using the playable measurement system to explore the interaction of bowing parameters that are produced by real players when a particular sound, or changes in sound, are produced. For instance, the bowing parameters at the beginnings of different strokes can be examined to try to understand how they contribute to the evolution of the sound throughout the entire stroke. Also, studies to examine the physical limits of human bowed string performance are of great interest.

In addition to the ways in which the playable measurement system may be used to study human performance dynamics particular to violin bowing, it may also be used to study instruments. A great deal of work has been done to evaluate different violins and isolate the salient differences between them by measuring their responses to various stimuli. What is proposed here is a different approach that seeks to evaluate instruments by measuring players’ responses to them. That is, by measuring the input parameters necessary to achieve the same sound on different instruments, the differences between the instruments may be defined by the different input required by them to produce the same sound.

Though the virtual violin work done in this thesis was limited, it yielded very interesting results and demonstrated that a true matching between calibrated controller and calibrated model is a tractable problem. Many more studies remain to be done, and much more development of both the model and the playable measurement system as a controller must begin. The future addition of left hand technique measurement (mentioned above) to the measurement system for this application will be a crucial step toward the goal of a truly playable virtual violin.

Progress in the development of a virtual violin might be complemented well by a future research project in which a notation for physical bowing is developed as a result of deep study of physical bowing parameters from real players. Such a notation might be used to indicate the technique or even individual style of bowing desired for a particular score of music and specific

virtual violin. This would allow a composer to write music for a virtual violin and a virtual player of his or her choosing.

As this future work proceeds, the amount of bowing data will grow. Therefore, it is important to continue development of the Bowstroke Database. In addition to encouraging other researchers to use this tool, future work on the database will most likely focus on the addition and refinement of features that facilitate collaboration and interaction between users. This will begin with Help menus and website navigation tools to assist new users while they learn how to use the database. New features may include file tagging, file sharing and implementing a process for suggestions between users, and user groups (specific, perhaps, to research fields, projects, or institutions).

In addition to researchers in bowed string synthesis, acoustics, and gesture, composers may also find the data contained within the Bowstroke Database to be of use. For instance, when composing interactive music they may use the gesture data available in the database to assist in the development of musical mappings for interactive compositions. This ability may be especially useful when workshop time with live players is limited as can often be the case.

It is also hoped that what has begun with the Bowstroke Database may be extended to another archiving project. That is, it could be used (with the playable measurement system) to archive real gestural performances by living masters of the violin (or other bowed string instruments, or other families of instruments) so that we might preserve them for future generations.

This thesis has provided significant steps to enable further study of bowed string performance and its applications, and it is hoped that this marks the beginning of many years of exciting work.

Appendix A

Violin Bowing Information

The following definitions and description are from [10].

A.1 Bowing Terms

1. **frog** - A device made of ebony, ivory, tortoise-shell, etc. at the lower end of the bow to which the hair is fastened. A screw mechanism functions to adjust the tension of the hair by sliding the frog back and forth along the stick in the mortice.
2. **tip** - The pointed end of the bow. Also refers to playing at the extreme upper part of the bow near the tip.
3. **upbow** - When the bow-hand moves toward the instrument while bowing.
4. **downbow** - When the bow-hand moves away from the instrument while bowing.

A.2 Bowing Techniques

1. **détaché** - Comprises a family of bow-strokes, played on-the-string, which share in common a change of bowing direction with the articulation of each note. *Detaché* strokes may be sharply accentuated or unaccentuated, *legato* (only in the sense that no rest occurs between strokes), or very slightly *staccato*, with small rests separating strokes.
2. **accented** (or **accentuated**) **détaché** - A percussive attack, produced by great initial bow speed and pressure, characterizes this stroke. In contrast to the *martelé*, the *accentuated*

détaché is basically a non-*staccato* articulation and can be performed at greater speeds than the *martelé* (generally used at the point). It is usually indicated only by accent marks.

3. **détaché lancé** - “Darting” *détaché*. Characteristically, a short unaccented *détaché* bow-stroke with some *staccato* separation of strokes.
4. **legato** - Bound together (literally, “tied”). Without interruption between the notes; smoothly connected, whether in one or several bows.
5. **louré** - This bowing consists of a short series of gently pulsed, slurred, *legato* notes. Varying degrees of articulation may be employed. The *legato* connection between notes may not be disrupted at all, but minimal separation may be employed. While the bow is in motion, pressure of the forefinger is applied to pulse a note, then released, and then reapplied for the following notes. This bowing is used in lyric passages to highlight expressive notes of a phrase. The range of articulation may be from a *legato* with a slight pressure on each note to gentle breaks in the sound.
6. **martelé** - Hammered; a sharply accentuated, *staccato* bowing. To produce the attack, pressure is applied an instant before bow motion begins. The string is then very taut before it is set into vibration. The bow is thrust into motion with great initial speed, and pressure is simultaneously reduced. Tension is suddenly discharged as the string snaps into maximum amplitude with an explosive accent, characteristic of the stroke. The bow is then quickly stopped or lifted (at the frog) in order that there will be a silence before the next stroke. *Martelé* differs from the *accentuated détaché* in that the latter has primarily no *staccato* separation between strokes and can be performed at faster speeds.
7. **spiccato** - Refers to a slow to moderate speed bouncing stroke. Every degree of crispness is possible in the *spiccato*, ranging from gently brushed to percussively dry.
8. **staccato** - Used as a generic term, *staccato* means a non-*legato martelé* type of short bow-stroke played with a stop. The effect is to shorten the written note value with an unwritten rest... *Staccato* may also refer to a style of bowing which may be likened to an unaccentuated *martelé*. The bow is not pressed before motion begins. Thus no sharp report, as in *martelé*, but a sort, “round” attack is produced. The bow is moved very rapidly and stopped abruptly to produce a rest before the next stroke. begins.

Appendix B

Measurement System Schematic

APPENDIX B. MEASUREMENT SYSTEM SCHEMATIC

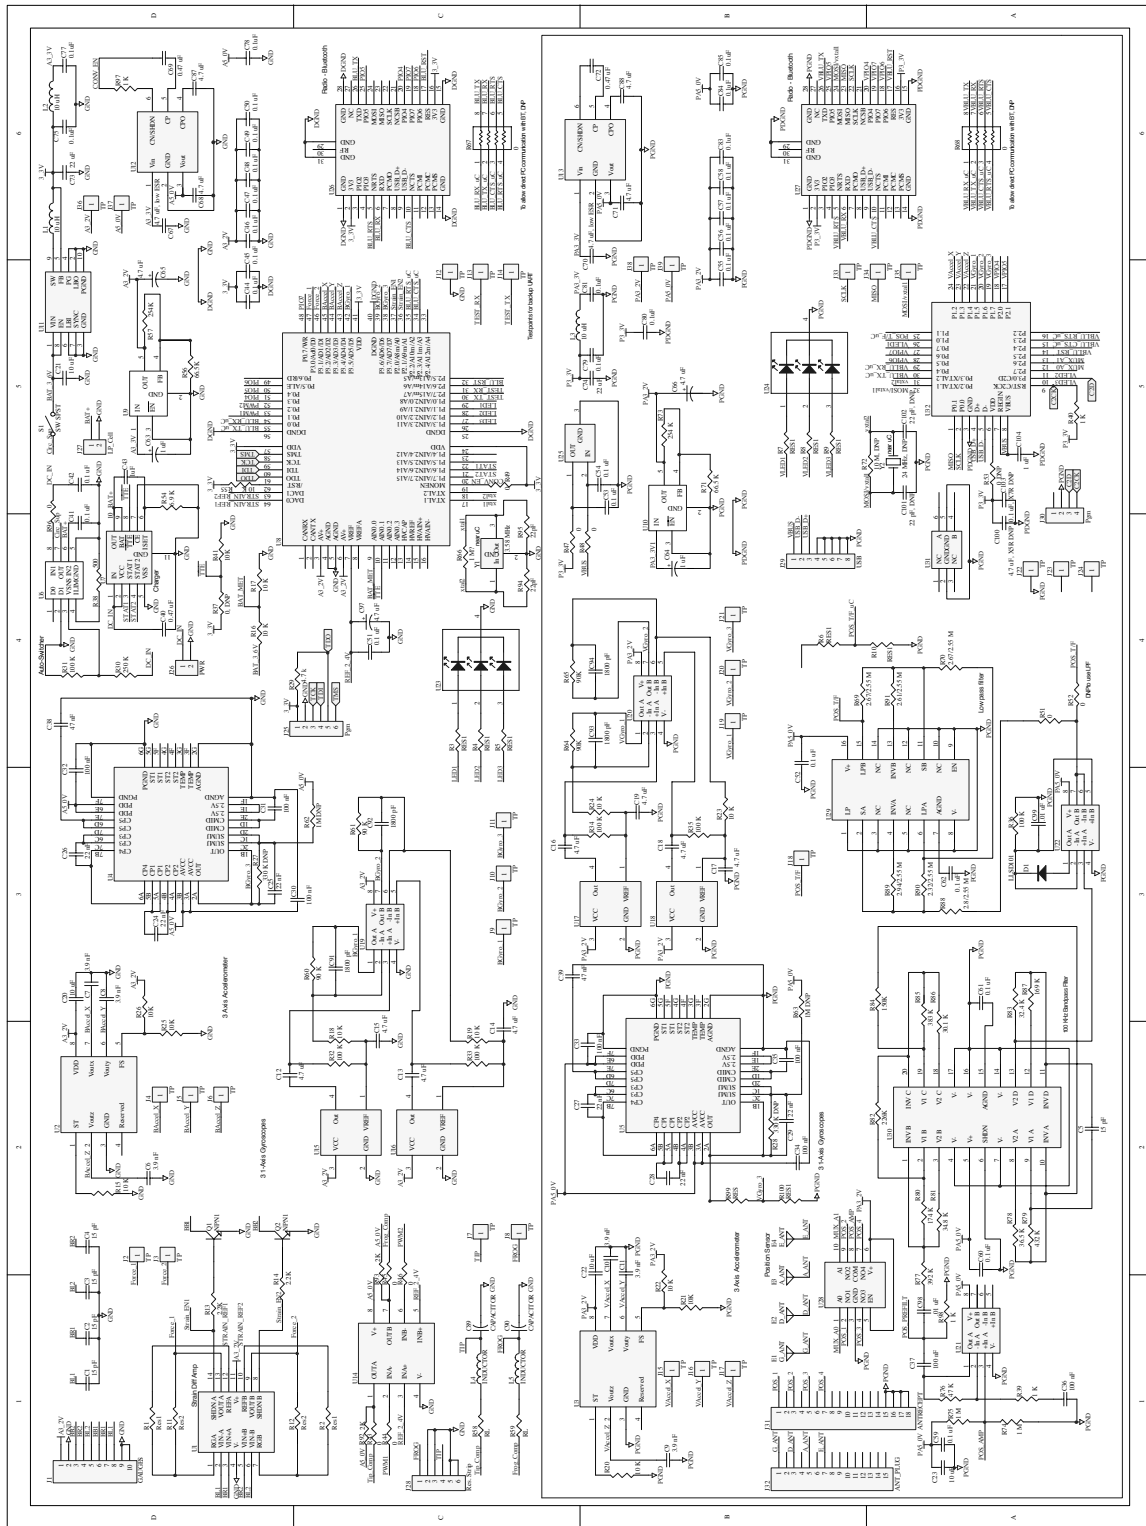


Figure B.1: Circuit schematic for violin bowing measurement hardware.

Appendix C

Force Calibration Curves

This appendix includes the individual force curves for both the downward force sensor and the lateral force sensor, obtained from the calibration exercises described in Chapter 3. Using the Instron® 1122 Universal Materials Testing Machine, increasing force was applied to the bow hair at each of thirteen positions, spanning twenty-four inches, along its length. (The thirteenth measurement position was one inch from the ferrule of the bow (frog end), and the distance between neighboring positions was two inches.)

The Instron® was controlled using Labview to apply force at a speed of (100 mm/min) up to a maximum limit and then return to its original rest position. During the calibration experiment, the bow force sensor data was recorded by the measurement system as was the Instron® load cell data. Measurements from the Instron® sensor were recorded at a sampling rate of 20Hz.

For each of these measurement positions, the force values (in Newtons) registered by the Instron® load cell are plotted against the corresponding bow sensor values. These plots show the linear performance of both force sensors.




Figure C.1: Using the Instron® machine, force was applied to the bow hair at each of thirteen positions, spanning twenty-four inches, along its length.

C.1 Downward Force Sensor

The following figures are the force curves describing the performance of the downward force sensor at thirteen positions along the length of the bow hair.

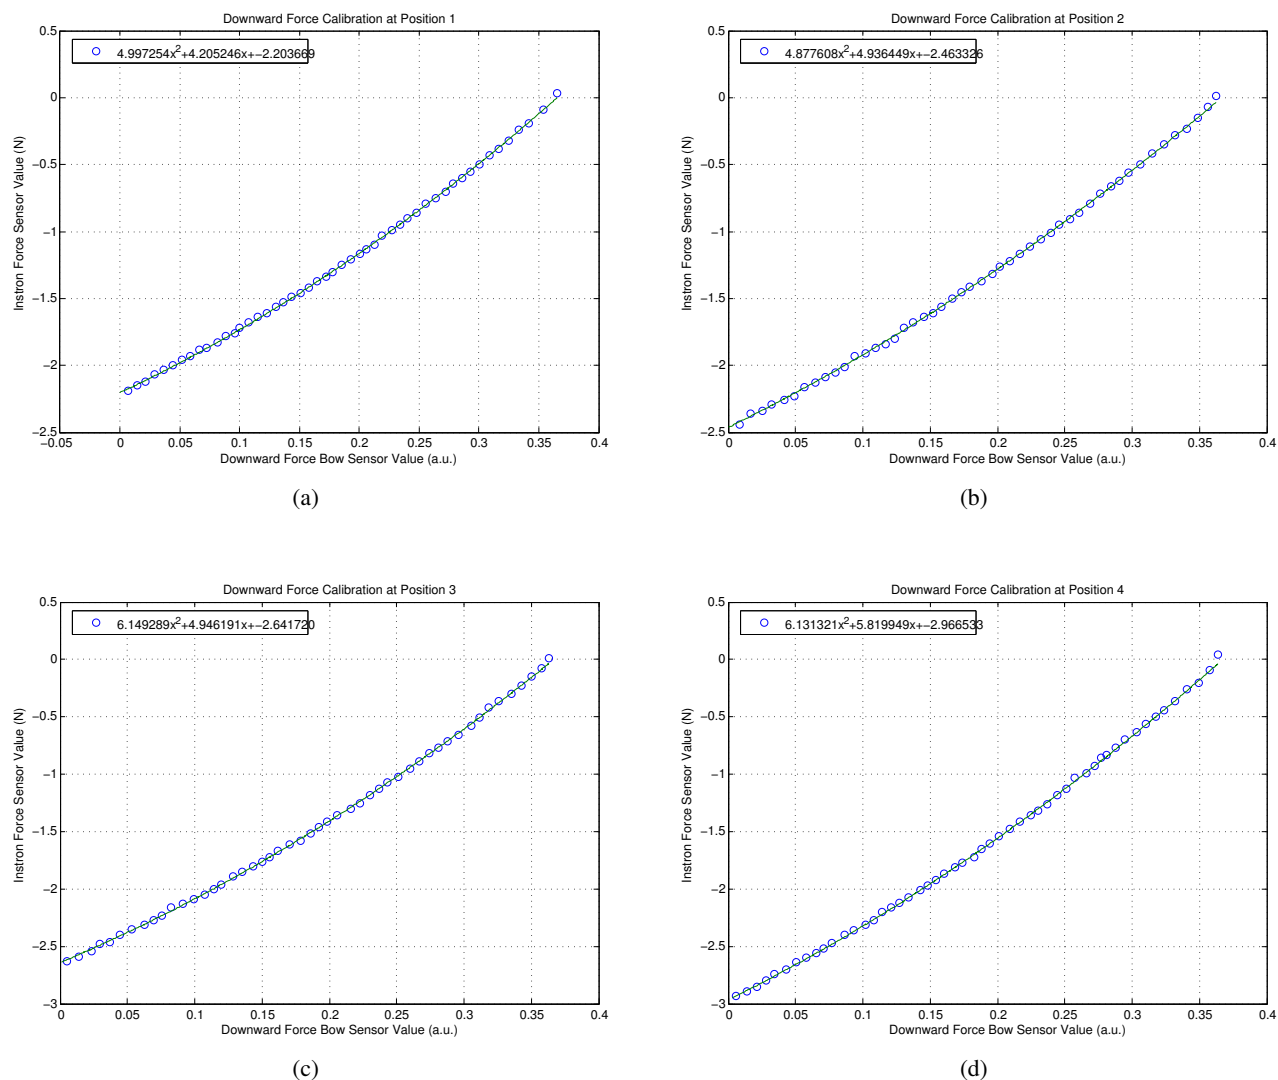

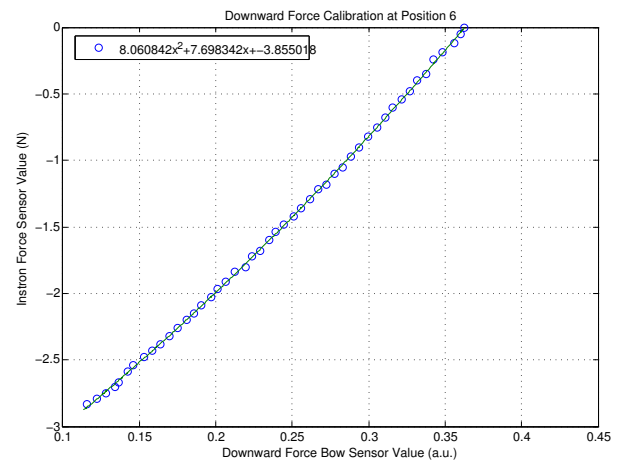



Figure C.2: These are the force calibration curves that were generated from data gathered at test points 1-4 using the Instron® machine. At each of these test points, the sensor had a good linear response.




(a)



(b)




(c)




(d)


Figure C.3: These are the downward force calibration curves that were generated from data gathered at test points 5-8 using the Instron® machine. At each of these test points, the sensor had a good linear response.




(a)



(b)



(c)



(d)

Figure C.4: These are the downward force calibration curves that were generated from data gathered at test points 9-12 using the Instron® machine. At each of these test points, the sensor had a good linear response.




Figure C.5: These are the downward force calibration curves that were generated from data gathered at test point 13 using the Instron® machine. As shown at the other test points, the sensor had a good linear response here.

C.2 Lateral Force Sensor

The following figures are the force curves describing the performance of the lateral force sensor at thirteen positions along the length of the bow hair.







Figure C.6: These are the lateral force calibration curves that were generated from data gathered at test points 1-4 using the Instron® machine. At each of these test points, the sensor had a good linear response.




(a)



(b)




(c)




(d)


Figure C.7: These are the lateral force calibration curves that were generated from data gathered at test points 5-8 using the Instron® machine. At each of these test points, the sensor had a good linear response.




(a)



(b)



(c)



(d)

Figure C.8: These are the lateral force calibration curves that were generated from data gathered at test points 9-12 using the Instron® machine. At each of these test points, the sensor had a good linear response.




Figure C.9: These are the lateral force calibration curves that were generated from data gathered at test point 13 using the Instron® machine. As shown at the other test points, the sensor had a good linear response here.

Appendix D

Bowing Technique Study Instructions

The written instructions for the *accented détaché*, *détaché lancé*, *détaché porté*, *louré*, *martelé*, *staccato*, and *spiccato* techniques recorded in the violin bowing technique study conducted in this thesis are included in the following pages. For each technique, a musical example is provided from [10].

II. Bowing Techniques

You will be asked to play a series of different bowing techniques on each of the 4 strings. For each string, you will be asked to perform the technique as notated 3 times.

Accompanying each technique is an example of that technique in the context of a piece of music. Please feel free to familiarize yourself with this musical context before beginning each set of recordings.

Detaché

Franck, *Sonata in A Major* for Violin and Piano (last movement)




Example of *accented détaché*

$\text{♩} = 84$



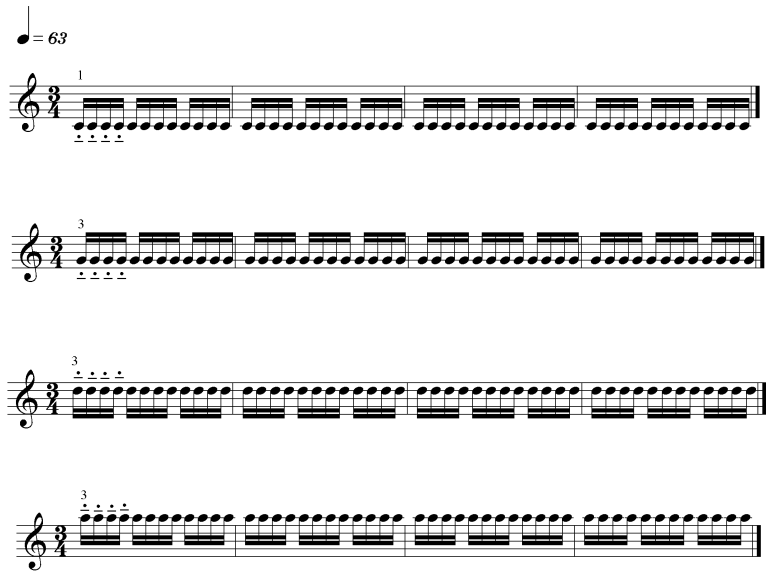
J.S. Bach, *Partita No. 2 for Violin Solo*, Ciaccona



simile

Example of *détaché lancé*

♩ = 63



Louré

Mendelssohn, *Violin Concerto in E Minor*, Op. 64
(first movement)



♩ = 132



Martelé

Vieuxtemps, *Violin Concerto No. 4 in D Minor*, Op. 31
(fourth movement)

A musical score snippet in treble clef, D minor, 4/4 time. It features a series of accented eighth notes. Above the notes, there are bowing markings: 'v' for downbow and 'w' for upbow. The first measure is marked 'energico' and 'f' (forte). The second measure has 'v w v' above it. The third measure has 'v w' above it. The fourth measure has 'v w' above it. The notes are: D4, E4, F4, G4, A4, B4, C5, B4, A4, G4, F4, E4, D4.

Example of *martelé*

♩ = 132

A musical staff in treble clef, C major, 4/4 time. It contains a sequence of 12 accented eighth notes: C4, D4, E4, F4, G4, A4, B4, C5, B4, A4, G4, F4. A first finger fingering '1' is written above the first note. Below the staff, there are 'A' markings under each note.


A musical staff in treble clef, C major, 4/4 time. It contains a sequence of 12 accented eighth notes: G4, A4, B4, C5, B4, A4, G4, F4, E4, D4, C4, B3. A third finger fingering '3' is written above the first note. Below the staff, there are 'A' markings under each note.

A musical staff in treble clef, C major, 4/4 time. It contains a sequence of 12 accented eighth notes: D4, E4, F4, G4, A4, B4, C5, B4, A4, G4, F4, E4. A third finger fingering '3' is written above the first note. Below the staff, there are 'V' markings above each note.

A musical staff in treble clef, C major, 4/4 time. It contains a sequence of 12 accented eighth notes: F4, G4, A4, B4, C5, B4, A4, G4, F4, E4, D4, C4. A third finger fingering '3' is written above the first note. Below the staff, there are 'V' markings above each note.

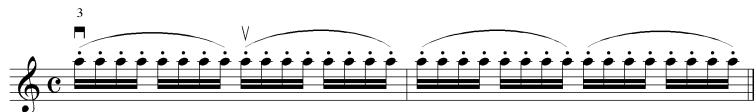
Staccato

Saint-Saens, *Introduction and Rondo Capriccioso*, Op. 28
Allegro ma non troppo



Example of the *firm staccato*

♩ = 88 (on the string)



Spiccato

(Beethoven, *String Quartet in C# Minor*, Op. 131)

Ex. 2

sul ponticello
arco
dim. *pp*

Examples of *spiccato*.

$\text{♩} = 192$

Bibliography

- [1] A. Askenfelt. Study your bowing technique! In *Proceedings of the Stockholm Music Acoustics Conference (SMAC 83)*, pages 123–391, Stockholm, 1983.
- [2] A. Askenfelt. Measurement of the bowing parameters in violin playing. II. Bow-bridge distance, dynamic range, and limits of bow force. *Journal of the Acoustical Society of America*, 86(2):503–16, Aug. 1989.
- [3] A. Askenfelt. Measurement of bow motion and bow force in violin playing. *Journal of the Acoustical Society of America*, 80(4):1007–15, Oct. 1986.
- [4] A. Askenfelt and K. Guettler. The Bouncing bow – An Experimental study. *CASJ (Series II)*, 3(6):3–8, 1998.
- [5] Leopold Auer. *Violin Playing as I Teach It*. Dover Publications, New York, reprint edition, 1980.
- [6] Autopilot Project. <http://autopilot.sourceforge.net/>.
- [7] Andres P. Baader, Oleg Kazennikov, and Mario Wiesendanger. Coordination of bowing and fingering in violin playing. *Cognitive Brain Research*, 23:436–443, 2005.
- [8] A. Y. Benbasat and J. A. Paradiso. A Compact Modular Wireless Sensor Platform. In *Proceedings of the 2005 Symposium on Information Processing in Sensor Networks (IPSN)*, 2005.
- [9] Ari Yosef Benbasat. An inertial measurement unit for user interfaces. Master’s thesis, M.I.T., 2000.
- [10] Joel Berman, Barbara G. Jackson, and Kenneth Sarch. *Dictionary of Bowing and Pizzicato Terms*. 4 edition, 1999.
- [11] Frédéric Bevilacqua, Nicolas H. Rasamimanana, Emmanuel Fléty, Serge Lemouton, and Florence Baschet. The augmented violin project: research, composition and performance report. In *Proceedings of the 2006 Conference on New Interfaces for Musical Expression (NIME-06)*, Paris, 2006.
- [12] Robert Grover Brown and Patrick Y. C. Hwang. *Introduction to Random Signals and Applied Kalman Filtering*. John Wiley & Sons, New York, 3rd edition, 1997.
- [13] C. Chafe. Simulating performance on a bowed instrument. Technical Report STAN-M-48, Center for Computer Research in Music and Acoustics (CCRMA), Department of Music, Stanford University, May 1988.

- [14] Chris Chafe. Celletto. <http://www-ccrma.stanford.edu/~serafin/NBF/img13.htm>.
- [15] CMU. CMU Graphics Lab Motion Capture Database. <http://mocap.cs.cmu.edu/>.
- [16] CodaBow. Conservatory Violin Bow. <http://www.codabow.com/>.
- [17] Lothar Cremer. *The Physics of the Violin*. MIT Press, Cambridge, MA, 1984.
- [18] Analog Devices. Angular Rate Sensor. <http://www.analog.com/en/prod/0%2C2877%2CADXRS300%2C00.html>.
- [19] Carl Flesch. *The Art of Violin Playing: Book One*. Carl Fischer, New York, reprint edition, 2000. Foreword by Anne-Sophie Mutter.
- [20] J. L. Florens and C. Henry. Bowed string synthesis with force feedback gesture interaction. In *Proceedings of the International Computer Music Conference*, pages 115–118. ICMC, 2001.
- [21] Adrian Freed, David Wessel, Michael Zbyszynski, and Frances Marie Uitti. Augmenting the Cello. In *Proceedings of the 2006 Conference on New Interfaces for Musical Expression (NIME-06)*, Paris, 2006.
- [22] Ivan Galamian. *Principles of Violin: Playing and Teaching*. Prentice-Hall, Englewood Cliffs, N.J., 1962.
- [23] Paul M. Galluzzo. *On the playability of stringed instruments*. PhD thesis, Trinity College, 2003.
- [24] Charles Gigante. *Manual of Orchestral Bowing*. Tichenor Publishing, Bloomington, Indiana, 1986.
- [25] Camille Goudeseune. *Composing with Parameters for Synthetic Instruments*. PhD thesis, University of Illinois, 2001.
- [26] Elizabeth A. H. Green. *Orchestral Bowings and Routines*. American String Teachers Association, Reston, VA, 1990.
- [27] Knut Guettler. *The bowed string: On the development of Helmholtz motion and on the creation of anomalous low frequencies*. PhD thesis, Royal Institute of Technology - Speech, Music and Hearing, 2002.
- [28] Knut Guettler. A Closer look at the string player’s bowing gestures. *CASJ (Series II)*, 4(7):12–16, 2003.
- [29] A. J. Hajian, D. S. Sanchez, and R. D. Howe. Drum roll: Increasing bandwidth through passive impedance modulation. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 2294–9, Albuquerque, April 1997.
- [30] J. A. Hernandez and A. K. T. Assiss. The potential, electric field and surface charges for a resistive long straight strip carrying a steady current. *American Journal of Physics*, 71(9):938–942, 2003.

- [31] Julius O. Smith III. A Basic Introduction to Digital Waveguide Synthesis (for the Technically Inclined). <http://ccrma.stanford.edu/~jos/swgt/swgt.html>, February 2006.
- [32] D. A. Jaffe and J. O. Smith. Performance expression in commuted waveguide synthesis of bowed strings. In *Proceedings of the International Computer Music Conference*, pages 343–346. ICMC, 1995.
- [33] Kokam. <http://www.kokamamerica.com/>.
- [34] Paul Kry and Dinesh Pai. Interaction capture and synthesis. In *ACM Transactions on Graphics (Proceedings of the SIGGRAPH conference)*, volume 25, pages 872–880, 2006.
- [35] Silicon Laboratories. <http://www.silabs.com/>.
- [36] M-Audio. Fast Track USB. http://www.m-audio.com/products/en_us/FastTrackUSB-main.html.
- [37] Tod Machover. Hyperinstruments: A Progress report 1987-1991. Technical report, MIT Media Laboratory, 1992.
- [38] Tod Machover and Joe Chung. Hyperinstruments: Musically intelligent and interactive performance and creativity systems. In *Proceedings of the 1989 International Computer Music Conference*, pages 186–190, San Francisco, CA, 1989. Computer Music Association.
- [39] The MathWorks. Aerospace Toolbox. <http://www.mathworks.com/products/aerotb/>.
- [40] L. McElligott, M. Dillon, K. Leydon, B. Richardson, M. Fernström, and J. A. Paradiso. ‘ForSe FIElds’ - Force sensors for interactive environments. In G. Borriello and L.E. Holmquist, editors, *UbiComp ’02: Proceedings of the 4th International Conference on Ubiquitous Computing*, volume 2498 of *Lecture Notes in Computer Science*, pages 168–175, London, 2002. Springer-Verlag.
- [41] Vishay Micro-Measurements. <http://www.vishay.com/company/brands/micrommeasurements/>.
- [42] Dirk Moelants. Temporal aspects of instrumentalists’ performance of tremolo, trills and vibrato. In *Proceedings of the International Symposium on Musical Acoustics (ISMA 2004)*, Nara, 2004.
- [43] Leopold Mozart. *A Treatise on the Fundamental Principles of Violin Playing*. Early Music Series. Oxford University Press, Oxford, 2nd edition, 1951.
- [44] Murata. Piezoelectric Vibrating Gyroscopes (GYROSTAR®). <http://www.murata.com/sensor/index.html>.
- [45] MySQL. <http://www.mysql.com/>.
- [46] Ian T. Nabney. Netlab neural network software. <http://www.ncrg.aston.ac.uk/netlab/index.php>.

- [47] Ian T. Nabney. *NETLAB: Algorithms for Pattern Recognition*. Advances in Pattern Recognition. Springer-Verlag, Great Britain, 2002. NETLAB toolbox available online at <http://www.ncrg.aston.ac.uk/netlab/index.php>.
- [48] Charles Nichols. The vBow: A Virtual violin bow controller for mapping gesture to synthesis with haptic feedback. *Organised Sound*, 7(2):215–220, 2002.
- [49] M. Sile O’Modhrain. *Playing by feel: Incorporating haptic feedback into computer based musical instruments*. PhD thesis, Stanford University, 2000.
- [50] M. Sile O’Modhrain and R. Brent Gillespie. The Moose: A Haptic user interface for the blind. In *STANM-95*, 1995.
- [51] S. O’Modhrain, S. Serafin, C. Chafe, and J. O. Smith. Qualitative and quantitative assessment on the playability of a virtual bowed-string instrument. In *Proceedings of the 2000 International Computer Music Conference (ICMC2000)*, pages 66–69, Berlin, August 2000.
- [52] MIT Committee on the Use of Humans as Experimental Subjects (COUHES). <http://web.mit.edu/committees/couhes/>.
- [53] Dan Overholt. The Overtone Violin. In *Proceedings of the 2005 Conference on New Interfaces for Musical Expression (NIME-05)*, Vancouver, 2005.
- [54] Dinesh Pai and Paul Kry. Rutgers MCL Interaction Capture Repository. <http://www.interactioncapture.org/>.
- [55] Caroline Palmer. Music Performance. *Annual Review of Psychology*, 48, 1997.
- [56] J. Paradiso. New Technologies for Monitoring the Precision Alignment of Large Detector Systems. *Nuclear Instruments and Methods in Physics Research*, Sec. A(386):409–420, 1997.
- [57] Joseph Paradiso and Neil Gershenfeld. Musical applications of electric field sensing. *Computer Music Journal*, 21(3):69–89, 1997.
- [58] PHP: Hypertext Preprocessor. <http://www.php.net/>.
- [59] R. Pitteroff and J. Woodhouse. Mechanics of the Contact Area Between a Violin Bow and a String. I. Reflection and Transmission Behaviour. *Acustica - Acta Acustica*, 84(3):543–62, May-June 1998.
- [60] Miller Puckette. Pure Data (Pd). <http://www.crca.ucsd.edu/~msp/software.html>.
- [61] N. H. Rasamimanana. Gesture analysis of bow strokes using an augmented violin. Master’s thesis, Université Pierre et Marie Curie, 2004.
- [62] Robert Rich. Buchla lightning midi controller. *Electronic Musician*, 7(10):102–108, 1991.
- [63] J. C. Schelleng. The Bowed string and the player. *Journal of the Acoustical Society of America*, 53:26–41, 1973.

- [64] Bernd Schoner, Chuck Cooper, and Neil Gershenfeld. Cluster-weighted sampling for synthesis and cross-synthesis of violin family instruments. In *Proceedings of the ICMC*, Berlin, 2000. ICMC.
- [65] Andras Semjen. On the Timing Basis of Bimanual Coordination in Discrete and Continuous Tasks. *Brain and Cognition*, 48:133–148, June 2002.
- [66] S. Serafin. *The sound of friction: Real-time models, playability and musical applications*. PhD thesis, Stanford University, 2004.
- [67] S. Serafin, R. Dudas, M. M. Wanderley, and X. Rodet. Gestural control of a real-time model of a bowed string instrument. In *Proceedings of the International Computer Music Conference*. ICMC, October 1999.
- [68] S. Serafin and J. O. Smith. A multirate, finite-width, bow-string interaction model. In *Proceedings of the COST G-6 Conference on Digital Audio Effects (DAFX-00)*, Verona, December 2000.
- [69] S. Serafin and J. O. Smith. Impact of string stiffness on digital waveguide models of bowed strings. *Catgut Acoustical Society Journal, Series II*, 4(4):49–52, November 2001.
- [70] S. Serafin, J. O. Smith, and J. Woodhouse. An Investigation of the impact of torsion waves and friction characteristics on the playability of virtual bowed strings. *Proceedings of 1999 Workshop on Applications of Signal Processing to Audio and Acoustics*, pages 87–90, 1999.
- [71] Stefania Serafin and Diana Young. Bowed string physical model validation through use of a bow controller and examination of bow strokes. In *Proceedings of the Stockholm Music Acoustics Conference (SMAC 03)*, Stockholm, August 2003.
- [72] Jeroen B. Smeets, Maarten A. Frens, and Eli Brenner. Throwing darts: timing is not the limiting factor. *Experimental Brain Research*, 144(2):268–274, May 2002.
- [73] J. H. Smith and J. Woodhouse. The Tribology of rosin. *Journal of the Mechanics and Physics of Solids*, 48(8):1633–81, August 2000.
- [74] J. O. Smith. Physical modeling using digital waveguides. *Computer Music Journal*, 16(4):74–91, 1992.
- [75] J. O. Smith. Virtual acoustic musical instruments: Review and update. *Journal of New Music Research*, 33(3), 2004.
- [76] STMicroelectronics. Three-Axis MEMS-based Accelerometers. <http://www.st.com/stonline/products/families/sensors/3-axis.htm>.
- [77] Gilbert Strang. *Linear Algebra and Its Applications*. Brooks Cole, Stamford, CT, 4th edition, 2005.
- [78] Vicon Motion Capture System. <http://www.vicon.com/>.
- [79] Ottó Szende and Mihály Nemessuri. *The Physiology of Violin Playing*. Collet's (Publishers) Limited, London, 1971. Foreword by Yehudi Menuhin.
- [80] Bluegiga Technologies. <http://www.bluegiga.com/>.

- [81] Dan Trueman. *Reinventing the Violin*. PhD thesis, Princeton University, 1999.
- [82] Dan Trueman and Perry R. Cook. BoSSA: The Deconstructed violin reconstructed. In *Proceedings of the ICMC*, Beijing, 1999.
- [83] UPM. <http://w3.upm-kymmene.com/>.
- [84] H. von Helmholtz. *On the Sensations of Tone*. Dover Publications, reprint edition, 1954.
- [85] Wacom. <http://www.wacom.com/>.
- [86] S.A. Wall and W.S. Harwin. Quantification of the effects of haptic feedback during motor skills tasks in a simulated environment. In *Proceedings of the 2nd PHANToM Users Research Symposium*, pages 61–69, Zurich, July 2000.
- [87] Piero Weiss and Richard Taruskin, editors. *Music in the Western World: A History in Documents*. Schirmer Books, New York, 1984.
- [88] J. Woodhouse. On the Playability of Violins. II. Minimum Bow Force and Transients. *Acustica*, 78(3):137–53, April 1993.
- [89] J. Woodhouse and P. M. Galluzzo. The bowed string as we know it today. *Acustica - Acta Acustica*, 90(4), Aug 2004.
- [90] James Woodhouse. Stringed instruments: Bowed. In Malcolm J. Crocker, editor, *Encyclopedia of Acoustics*, Wiley-Interscience, pages 1619–1626. Wiley, Cambridge, 1997.
- [91] Matthew Wright and Adrian Freed. OpenSound Control: A New Protocol for Communicating with Sound Synthesizers. In *Proceedings of the ICMC*, Thessaloniki, 1997. ICMC.
- [92] Yamaha. SV-200 Silent Violin. <http://www.global.yamaha.com/index.html>
- [93] Diana Young. The Hyperbow: A precision violin interface. In *Proceedings of the ICMC*, Goteborg, 2002. ICMC.
- [94] Diana Young. The Hyperbow controller: Real-time dynamics measurement of violin performance. In *Proceedings of the 2002 Conference on New Instruments for Musical Expression (NIME-02)*, Montreal, 2002.
- [95] Diana Young. Wireless sensor system for measurement of violin bowing parameters. In *Proceedings of the Stockholm Music Acoustics Conference (SMAC 03)*, Stockholm, August 2003.
- [96] Diana Young. *A Methodology for Investigation of Bowed String Performance Through Measurement of Violin Bowing Technique*. PhD thesis, M.I.T., 2007.
- [97] Diana Young and Anagha Deshmane. Bowstroke Database. <http://bowing.media.mit.edu/>.
- [98] Diana Young, Patrick Nunn, and Artem Vassiliev. Composing for Hyperbow: A Collaboration Between MIT and the Royal Academy of Music. In *Proceedings of the 2006 Conference on New Interfaces for Musical Expression (NIME-06)*, Paris, 2006.

- [99] Diana Young and Stefania Serafin. Playability evaluation of a virtual bowed string instrument. In *Proceedings of the 2003 Conference on New Interfaces for Musical Expression (NIME-03)*, Montreal, 2003.
- [100] Diana S. Young. New frontiers of expression through real-time dynamics measurement of violin bows. Master's thesis, M.I.T., 2001.