Engineering the violin

Artificial intelligence and simulation may well rival Stradivarius

S ANANTHANARAYANAN

Musical instruments are known for complexity and precision, and the violin ranks as one of the most complex. The present form of the violin first appeared in the early 16th century, and present violins follow designs of master violin makers from the 18th century. Extremely high quality is now attained, but it has not been possible to match 18th century masterpieces.

Sebastian Gonzalez, Davide Salvi, Daniel Baeza, Fabio Antonacci and Augusto Sarti, from the Musical Acoustics Lab at the Violin Museum of Cremona, part of the Polytechnic of Milan, and the faculty of physical and mathematical sciences, University of Chile, describe in the journal, *Scientific Reports*, how they have used Artificial Intelligence to peer into the internals of 18th century violins, to understand the factors that made for their exquisite performance.

The characteristic of musical instruments is that each kind can produce the notes of musical scales according to its own quality of sound, and do it consistently. This is to say that the notes of a piano, for instance, have a similar purity or quality, or the mix of overtones, right through the different octaves. So also, a violin, or trumpet, has its own quality, and a skilled player can bring out sounds, typically of the violin or trumpet, of different pitch and loudness, or softness, to create music.

Building instruments to do this needs accurate balancing of a number of features, so that sounds produced are consistent in quality through the range of pitch and loudness that is possible with the instrument. For instance, the thickness and material of piano strings need to be just so. And the material of the soundboard, the materials of hammers, need to match and go together, so that the piano is suitable for a range of compositions. The design of the modern piano, hence, represents elaborate and sensitive coordination of its different parts.

The violin is much smaller than a piano and has much fewer parts. Its design, however, is no less complex. With just four strings, and not longer than a child's arm, the violin can produce notes of great purity and power, over nearly four octaves. To be able to do this consistently over its range, many parts of the violin need to be accurately shaped in keeping with the properties of materials used. The design of the violin took more than two centuries to grow to its present state. And the growth took place with concerted trial and error work of violin makers, violinists and composers to make the best use of, and exacting demands on the violin's capability.

The first step in making a violin is shaping the sides, or the outline of the instrument. This is done by creating a mould, and bending selected strips of wood, called the ribs, into the desired contour. When this is made, the specially carved bottom and top are attached, to create the resonating hollow for the sound when the strings are bowed. Next, the neck, which is shaped out of a single piece of wood, is fixed to the body to complete the structure.

As shown in the picture, the strings are attached from the tailpiece, over the bridge and to the pegs. The pegs enable the tension of the strings to be adjusted and there is provision, between the bridge and the tailpiece, for fine adjustment. Once the tension of the four strings are set, the note that any string sounds can be changed by the violinist's finger pressing the string on to the fingerboard. And after everything is in place, it takes fine correction over months, or years, for the violin to attain full timbre and volume.

The current shape, or outline of the violin has come about after several iterations during the first 200 years of its development, the paper in *Scientific Reports* says. However, even geometrically exact copies cannot equal the perfection attained by old masters. This suggests that the difference may lie in the material, the wood sourced three centuries ago. Current violin makers have thus attempted "tweaking" the shape to compensate. Without, however, any notable success, as there is yet no science of how the contours affect and could improve the sound of a violin.

The authors of the paper hence used a different approach. They make use of computers and the methods of AI to create a dynamic model of the Stradivari violin and assess the role of the different aspects of its structure. Based on the evolving designs of violins down the centuries, the authors of the paper identified a list of features that affected how the violin performs, as well as units to measure the features. These measures, and the known, main sound frequencies that the violin supported, were then fed into an AI computer programme, to work out how far each of the identified features influenced the final product.

Artificial intelligence is so called because it is a computer procedure that mimics the way of learning of natural systems, like the human brain. For an infant to learn speech, for instance, the method is not of learning vocabulary and formal rules of grammar, as when an adult learns a foreign language. Instead, the infant brain collects the different sounds that it hears, and the objects, sights or sensations that accompany the sounds. Where certain sensations are repeatedly associated with particular sounds, the infant brain strengthens the connection and associates the sight or sensation with the sound alone – a first step in learning language.

When computer programmes are developed to mimic this process, we get in our hands a powerful means of "natural" learning, for computers can crunch huge data and work continuously for hours. AI programmes, known as machine learning, are hence used to analyse huge data, make weather forecasts, even train road vehicles to drive themselves, or learn language and work as a translator.

The selection of factors to be considered, the paper says, started from an original drawing by Antonio Stradivari, preserved in the Violin Museum of Cremona, Italy. We can make out from the copy shown here that the outline has been derived from nine circles of different diameters. The position of the centres and angles further specify the shape, of the "ribs" and the "top plate". The shape is thus constrained by 20 parameters, the paper says, and in addition, there are parameters of the curvature of the top plate.

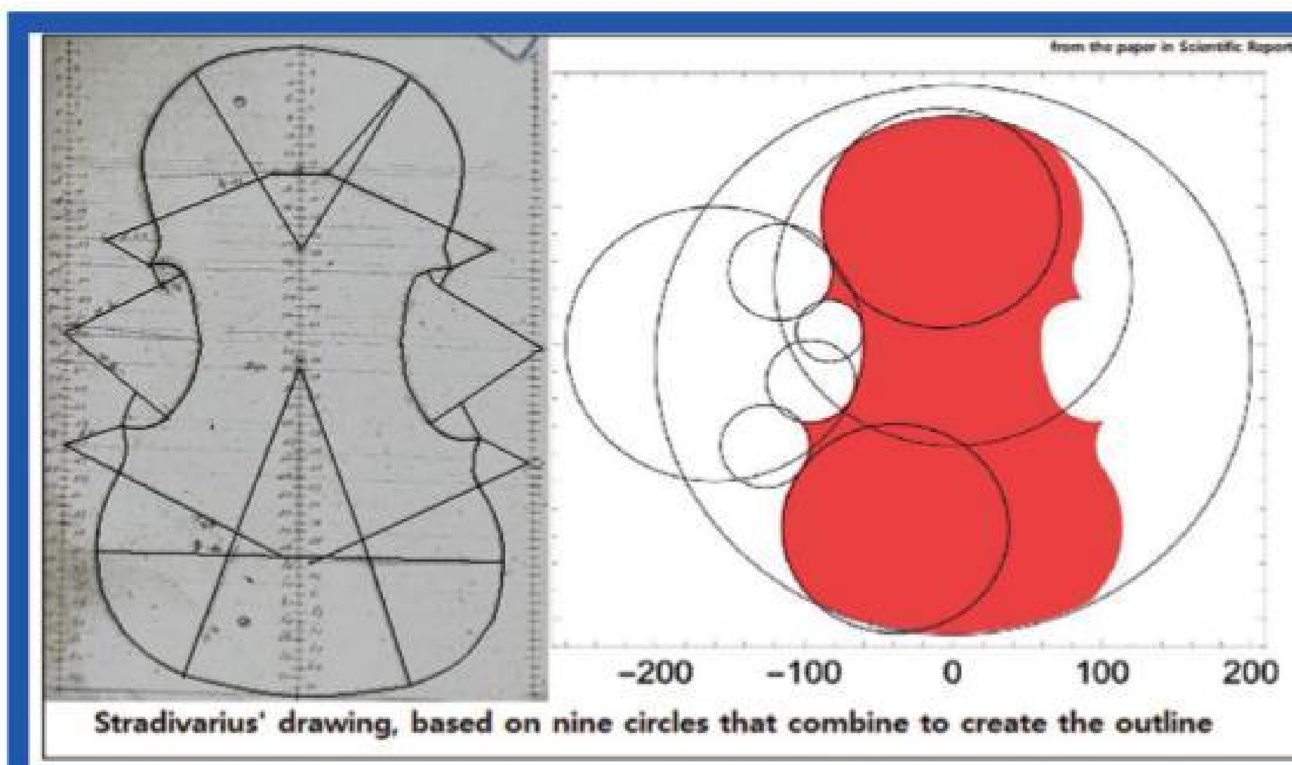

Stradivarius' drawing, based on nine circles that combine to create the outline

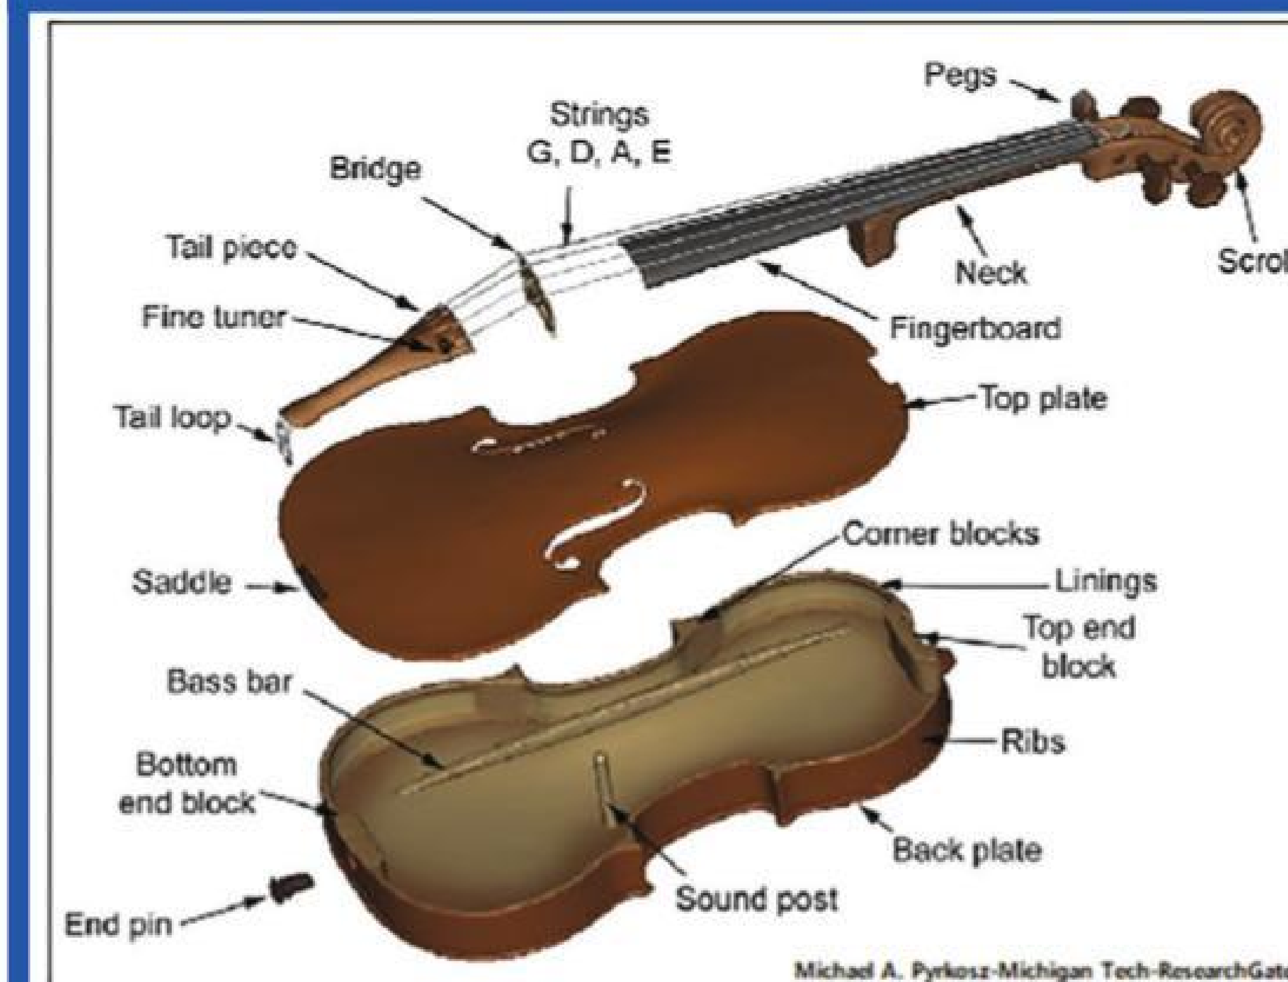

Michael A. Pyrkosz-Michigan Tech-ResearchGate

The quality of a violin is determined largely by the sound frequencies that are natural to vibrational modes of the top plate. And these frequencies depend largely on the shape of the plate. The parameters that specify the shape of the plate hence determine the frequencies the plate would support, and hence the sound quality. The relationship, however, is complex indeed, which is why studies so far have not been able to say what effect a change in one of the features would have on the sound produced.

The values of 1,750 sets of the 20 shape parameters and 10 of the corresponding known, main frequencies were then fed into a machine learning programme to "train" the programme. Next, the parameters of 250 other shapes were fed in, to see if the main frequencies predicted were close to the actual frequencies. The result was that the programme was successful – it was able to learn from the "training set" and make accurate predictions!

The work is a proof in principle, that methods of AI can be used in violin design. This trial's description was limited to natural frequencies, but the method can be extended to cover more features and overall violin performance. "The ability to predict how a violin design sounds can truly be a game changer for violin makers, as not only will it help them do better than the grand masters, but it will also help them explore the potential of new designs and materials," the paper says.

The writer can be contacted at response@simplescience.in

IT'S ALL ABOUT DENSITY

Here's how DNA replicates semi-conservatively in all organisms


Watson (left) and Crick with their DNA model in 1953

TAPAN KUMAR MAITRA

One of the most significant features of the double-helical model of deoxyribonucleic acid is that it immediately suggests a mechanism for replication. In fact, a month after James Watson and Francis Crick published their now-classic paper postulating a double helix for DNA, they followed it with an equally important paper suggesting how such a base-paired structure might duplicate itself.

Here, in their own words, is the basis of that suggestion, "Now our model for deoxyribonucleic acid is, in effect, a pair of templates, each of which is complementary to the other. We imagine that prior to duplication the hydrogen bonds are broken, and the two chains unwind and separate. Each chain then acts as a template for the formation onto itself of a new companion chain, so that eventually we shall have two pairs of chains, where we only had one before. Moreover, the sequence of the pairs of bases will have been duplicated exactly. (Watson and Crick 1953)"

The essence of the model Watson and Crick proposed for DNA is that one of the two strands of every newly formed DNA molecule is derived from the parent molecule, whereas the other strand is newly synthesised. This is called semi-conservative replication, because half of the parent molecule is retained by each daughter molecule.

Within five years of its publication, the Watson-Crick model of semi-conservative DNA replication was tested and proved correct by Matthew Meselson and Franklin Stahl. The ingenuity of their contribution lay in the method they devised, in collaboration with Jerome Vinograd, for distinguishing semi-conservative replication from other possibilities. Their studies utilised two isotopic forms of nitrogen, ^{14}N and I5N, to distinguish newly synthesised strands of DNA from old strands.

Bacterial cells were first grown for many generations in a medium containing ^{15}N-labelled ammonium chloride to incorporate this heavy (but nonradioactive) isotope of nitrogen into their DNA molecules. Cells containing ^{15}N-labelled DNA were then transferred to a growth medium containing the normal light isotope of nitrogen, ^{14}N. Any new strands of DNA synthesised after this transfer would therefore incorporate ^{14}N rather than ^{15}N.

Since ^{15}N-labelled DNA is significantly denser than ^{14}N-labelled DNA, the old and new DNA strands can be distinguished from each other by equilibrium density centrifugation. For DNA analysis, equilibrium density centrifugation often uses cesium chloride, a heavy metal salt that forms solutions of very high density. The DNA to be analysed is simply mixed with cesium chloride and the solution is centrifuged at high speed for a relatively long time (with a modern centrifuge, 80,000 rpm for eight hours, for example). As a density gradient of cesium chloride is established by the centrifugal force, the DNA molecules float "up" or sink "down" within the gradient to reach their equilibrium density positions. The difference in density between heavy (^{15}N-containing) DNA and light (^{14}N-containing) DNA causes them to come to rest at different positions in the gradient.

Using this approach, Meselson and Stahl analysed the DNA obtained from bacterial cells that were first grown for many generations in I5N and then transferred to ^{14}N for one or more additional cycles of replication What results would be predicted for a semi-conservative mechanism of DNA replication? After one replication cycle in ^{14}N, each DNA molecule should consist of one ^{15}N strand (the old strand) and one ^{14}N strand (the new strand), and so the overall density would be intermediate between heavy DNA and light DNA. The experimental re-sults clearly supported this model. After one replication cycle in the ^{14}N medium, centrifugation in cesium chloride revealed a single band of DNA whose density was exactly halfway between that of ^{15}N-DNA and ^{14}N-DNA.

Because they saw no band at the density expected for heavy DNA, Meselson and Stahl concluded that the original, double-stranded parental DNA was not preserved intact in the replication process. Similarly, the absence of a band at the density expected for light DNA indicated that no daughter DNA molecules consisted exclusively of newly synthesised nucleotides. Instead, it appeared that a part of every daughter DNA molecule was newly synthesised, while another part was derived from the parent molecule. In fact, the density halfway between that of ^{14}N-DNA and ^{15}N-DNA meant that the hybrid DNA molecules were one-half parental and one-half newly synthesised, just as pre-dicted by the semi-conservative model of replication.

The data from cells allowed to grow in the presence of ^{14}N for additional generations provided further confirmation. After the second cycle of DNA replication, for example, Meselson and Stahl saw two equal bands, one at the hybrid density of the previous cycle and one at the density of purely ^{14}N-DNA this is also consistent with a semi-conservative mode of replication.

Meselson and Stahl concluded "that the nitrogen of a DNA molecule is divided equally between two physically continuous subunits; that, following duplication, each daughter molecule receives one of these; and that the subunits are conserved through many duplications". Further experimentation was then required to prove that the "physically continuous subunits" into which DNA is partitioned are indeed separate DNA strands. Meselson and Stahl obtained this proof by heating the ^{14}N/^{15}N hybrid DNA to separate its two strands and then showing that one strand exhibited the density of a ^{15}N-containing strand and the other exhibited the density of a ^{14}N-containing strand.

Meanwhile, other researchers had used radioactive labelling and autoradiography to look at the process of DNA replication in eukaryotic chromosomes. In the end, Watson and Crick were proven right – DNA is replicated semi-conservatively in all organisms.

The writer is associate professor and head, department of botany, Ananda Mohan College, Kolkata

PLUS POINTS

Lens-less camera

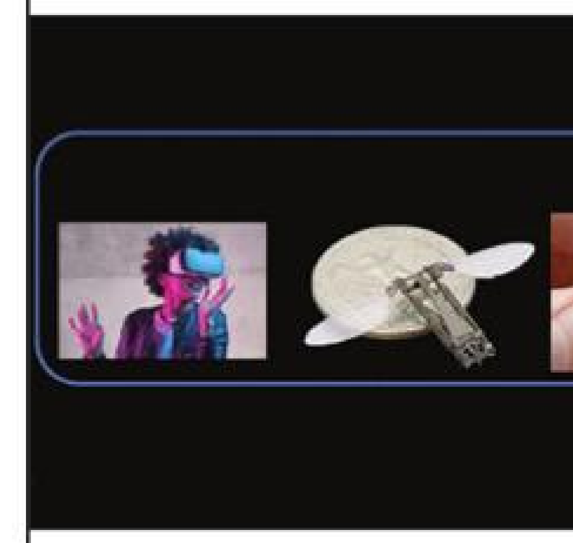

Indian Institute of Technology-Madras and Rice University, United States, have developed algorithms for lens-less, miniature cameras. Such lens-less cameras have numerous vision applications in areas such as augmented reality/ virtual reality, security, smart wearables and robotics where cost, form-factor and weight are major constraints.

In 2016, professor Ashok Veeraraghavan's lab at Rice University registered success in making a lens-less camera. They were able to develop a low-cost and low-weight ultra-thin lens-less camera where a thin optical mask was placed just in front of the sensor at a distance of approximately one mm. However, because of the absence of focusing elements, the lens-less camera captured blurred images restricting their commercial use. IIT-Madras and Rice researchers have now developed a computational solution to this problem. The team developed a deblurring algorithm, which can correct the images taken from a lens-less camera. The findings were presented as a paper at the prestigious IEEE International Conference on Computer Vision and an extended version appeared in IEEE *Transactions on Pattern Analysis and Machine Intelligence*.

The research was led at IIT-Madras by Kaushik Mitra, head of the Computational Imaging Laboratory and assistant professor, department of electrical engineering. The research team included Salman Siddique Khan, Varun Sundar and Adarsh VR from IIT-Madras. Veeraghavan led the Rice University team, which included Vivek Boominathan and Jasper Tan.

Mitra said, "Lens-less imaging is a new technology and its true potential in solving imaging/vision problems has not been exploited completely. Therefore, we are working on designing newer and better lens-less cameras using data-driven techniques, devising efficient algorithms for doing inference on lens-less captures and looking into interesting and important applications like endoscopy and smart surveillance, among other areas, where one can fully realise the benefits of lens-less imaging."

For preterm birth

A new medical device that improves the prediction of preterm birth at a fraction of the cost of current methods, will help to reduce the global number of deaths and long-term complications caused by babies born prematurely.

Pioneering research carried out on Electrical Impedance Spectroscopy at the University of Sheffield, United Kingdom, has led to the creation of the device, brought to market by EveryBaby, a UK-based company backed by South Korean investment.

Current technologies for assessing the likelihood of preterm birth, such as trans-vaginal ultrasound, are expensive and not always available – especially in poorer countries. The new device uses pioneering patented technology involving a novel method of impedance spectroscopy to pick up on changes to the composition and structure of cervical tissue as a mother nears birth. The device provides a higher degree of accuracy in an affordable, hand-held and portable device, giving huge potential for export to low-income countries.

Dilly O'Anumba, professor of obstetrics and gynaecology at the University of Sheffield and consultant at Sheffield Teaching Hospitals Foundation Trusts, said, "More than one in 10 babies are born too early and data has shown that preterm birth rates are increasing in many parts of the world.

"This pioneering technique will enable health care professionals to better prevent and manage preterm birth. It is not only more accurate than current methods but is significantly lower in cost making it more accessible, especially in low income communities where preterm birth rates are particularly high."

The World Health Organisation estimates 15 million babies are born preterm every year, with complications arising from premature births being the leading cause of death in infants under five years of age. Long term complications can be severe, and many survivors face a lifetime of disability, including learning disabilities as well as visual and hearing impairments.