Young Acousticians Network Newsletter #92 September 2020

Newsletter's Summary

Board's Highlights

News

The YAN discussed the Matlab MIRtoolbox wih Olivier Lartillot and Petri Toiviainen. Read about the interview on page 4.

Page 4

Jobs

A lot of new job offers in this month's curated list. Find your next job opportunity in this newsletter.

Page 11

Upcoming Events

September 2020

9th - 11th — ASJ Autumn 2020 — Acoustical Society of Japan: Autumn Meeting. Japan. **Virtual Conference**

16th - 17th — ICAANS 2020 — International Conference on Architectural Acoustics and Noise Control. Lisbon, Portugal. **Virtual Conference**

17th - 18th — ISMA45-ISAAC31 — ISMA45 & ISAAC31 Lectures on Advanced Techniques in Applied and Numerical Acoustics. Leuven, Belgium. **Virtual Conference**

October 2020

14th - 15th — Acusti.cat 2020 — 3rd Catalan Congress of Acoustics. Sant Cugat del Vallès, Catalonia, Spain.

21st - 21st — Quiet Drones — A Symposium on Noise from UASs/UAVs. Paris, France

November 2020

3rd - 5th — ISNVH 2020 — International Styrian Noise, Vibration and Harshness (ISNVH) Congress. Graz, Austria.

16th - 18th — Noise-Con 2020 — Noise Control Conference. New Orleans, LA, USA.

19th — The acoustics of buildings. Bari, Italy.

23rd- 25th — A&V 2020 — 1st Biennial International Conference on Acoustics and Vibration. Denpasar, Indonesia.

Upcoming Deadlines

September 2020

22nd — A&V 2020 — 1st Biennial International Conference on Acoustics and Vibration. Denpasar, Indonesia. **Abstract submission**

October 2020

31th — FA 2020 — Forum Acusticum. Lyon, France. **Paper submission**

November 2020

2nd — SAPEM 2020 — Symposium on the Acoustics of Poro-Elastic Materials. Wes Lafayette, Indiana, USA. **Abstract submission**

8th — A&V 2020 — 1st Biennial International Conference on Acoustics and Vibration. Denpasar, Indonesia. **Paper submission**

Did we miss a date ?

Behind the YAN, there're humans you can help!

The agenda listing is all gathered by hand: if you think we missed something relevant, don't hesitate to tell us!

yan@euracoustics.org

News

MIRtoolbox is a Matlab toolbox dedicated to the analysis of music and sound from audio recordings and to the extraction of musical features such as tonality, rhythm, or structures. It has also been used for non-musical applications, such as in Non Destructive Testing, and with non-audio signals. In this issue of the newsletter, the YAN discusses the MIRtoolbox with Olivier Lartillot (RITMO Centre for Interdisciplinary Studies in Rhythm, Time and Motion, University of Oslo, Norway) and Petri Toiviainen (University of Jyväskylä, Finland)

You can also check out the MIRtoolbox website at:

shorturl.at/oA038

Usability of the toolbox

In a few words, what are the strengths of the toolbox?

It offers an extensive overview of techniques for the analysis of music and sound from audio recordings. The extent of the set of features is, I would say, larger than in other toolboxes of that category. It is very easy to use, even for non-Matlab users, but is very convenient to use for expert users as well. The high modularity of the architecture makes it possible to quickly carry out analyses answering the users' needs. Technical details can be ignored in a first approach, but when needing to refine the analyses, a very large number of options and parameters are available. A large range of visualisations is also available. We have developed original features that can be useful for a large range of applications. For instance, the peak picking is very versatile, and adapts quite well to the type of signal analyzed.

Who is the typical user of the toolbox?

A little difficult to answer, based on the, grossly, 50000 of downloads and 1500 citations in papers. I would guess that the largest number of users are students, because the toolbox is used for teaching in many places, and students use it for their studies, including PhDs. Then many researchers use it also in very various contexts in different scientific domains.

> … a combination of the need to establish a software that would be state of the art in computational music analysis from audio, easy to use for both students and researchers, and that would also encourage the user to think about the whole picture in a modular way….

How often is the MIRtoolbox used by other research labs in your field?

Actually in the Music Information Retrieval field itself, it is not used very often by expert researchers, because the current trend is towards deep learning, and using as

News

input low level spectral analyses performed on other platforms than Matlab (Python in particular). But it seems that in computational musicology and in music cognition, it is used somewhat often. It is difficult to provide a statistic, but again the number of citations gives a good picture.

Other than in Music Information Retrieval, in which areas of research could people also benefit from the MIRtoolbox?

Looking at papers citing MIRtoolbox, we can see general audio analysis research, and speech in particular. For exotic audio applications, I can cite for example "Detailed modeling of cutting forces in grinding process considering variable stages of grain-workpiece micro interactions". MIRtoolbox has also been used for the analysis of non-audio signals. For instance: "Anomaly detection in earth dam and levee passive seismic data" or "Automatically detecting avalanche events in passive seismic data".

Detailed modeling of cutting forces in grinding process considering variable stages of grain-workpiece micro interactions
https://doi.org/10.1016/j.ijmecsci.2016.11.016

Anomaly detection in earth dam and levee passive seismic data using support vector machines and automatic feature selection
https://doi.org/10.1016/j.jocs.2016.11.016

Automatically Detecting Avalanche Events in Passive Seismic Data
DOI: 10.1109/ICMLA.2012.12

Could a user adapt the functions to non-musical signals? And how about outside the human audible range?

It depends on the features under consideration. General signal processing operators and general features related to dynamics and timbre should be as ready-to-use as in the musical context. Concerning pitch, the particular frequency region under consideration might need to be tuned. For rhythm, it might require particular tuning of parameters as well, and I have to say I haven't worked much on non-musical rhythm. In any case, even for music, some parameters might need to be tuned for the particular problems at hand. Especially when dealing with high-level dimensions such as structure, for instance related to speech segmentation.

As for working with signals outside the human audible range, yes. I guess the user would need to make sure that the default parameters are valid with the particular type of sound. I cannot answer precisely without doing experiments on my own.

News

History of the toolbox

How were the early days of the toolbox? Where did the idea for the toolbox come from?

I was employed as a postdoc in Petri's team for the BrainTuning project ("Tuning the Brain for Music", funded by the NEST (New and Emerging Science and Technology) program of the European Commission), which started early 2006 and whose aim was to study links between brain, music and emotion. We were working on a particular task of unveiling correspondences between audio and musical features and perceived or induced emotion, with Tuomas Eerola, at that time also in Jyväskylä, and Anders Friberg and Roberto Bresin at KTH in Stockholm. Our expertise at that time was somewhat more on the "symbolic side" of music (meaning the study of scores, or MIDI sequences) than on the "audio side". But we realized that a thorough study of emotion in music would require investigating the audio aspects of music. Luckily, I had a background in signal processing so I was in charge of developing features for audio and music analysis based on the state of the art. It looked quite daunting to me at first…

At that time, Petri already nurtured the good habit of developing and publishing Matlab toolboxes for teaching and research purposes. Tuomas developed with him the MIDItoolbox, which offers a large range of tools to analyse music from MIDI format. Released a couple of years before MIRtoolbox, it has become a reference tool in the computational musicology community. So Petri had the good idea to release another toolbox, this time for analysis from audio.

The third dynamic that enabled the conception of MIRtoolbox was that we were starting a new master program, Music Mind Technology, at the University of Jyväskylä, in which there was a course, starting in Fall 2006 about Music Information Retrieval. I was in charge of the demo part, and Petri and I decided that we would present to the students prototypes of sound and music analysis tools that would be part of MIRtoolbox, starting from basic signal processing concepts and covering progressively various music dimensions. The challenge, here, is that the students came from very different backgrounds, and many of them did not have expertise in computer science. So we had to make a toolbox very easy to use. But that constraint was actually very fruitful for the development of the toolbox. The first demo was about basic signal processing concepts such as FFT and autocorrelation. The demo went well and it looked quite promising, although I remember one student who was dubious that this tool could be used for more advanced analysis than this first demo. That was MIRtoolbox 0.1.

It was the combination of the need to establish a software that would be state of the art in computational music analysis

News

from audio, easy to use for both students and researchers, and that would also encourage the user to think about the whole picture in a modular way, that made MIRtoolbox possible.

It was very interesting to discover how the modular design enables to fuse concepts that were previously considered in separate subdomains of research, such as rhythmic vs. harmonic analysis into common modules where expertise could be exchanged between these silos.

How many developers contributed to the MIRtoolbox when it started? How about now?

Well, it was (and still is) basically me, with some code from Petri and some code available publicly as well as third party toolboxes that are included in the distribution.

How did the toolbox grow to where it is today? Were you supported by your lab to build the toolbox? How about now?

I was very lucky to have the possibility to spend a lot of my time on the development of the toolbox throughout the years, first during that initial postdoc for the BrainTuning project, then as a postdoc in the Petri's Center of Excellence in Interdisciplinary Music Research, then as an Academy of Finland research fellow, and in my subsequent research position at the Universities of Geneva, Aalborg and Oslo. The development benefitted from a large amount of interaction with students, colleagues and users.

Pasi Saari, who was a PhD student at Petri's lab, developed the "mirplayer" graphical player interface. Another student, Enrico Glerean, introduced me to version control, and helped me create a SubVersion repository that we are still using.

How do people hear about your toolbox?

We first advertised it to students in courses and in specialized summer schools, then we published articles about it and gave tutorials in various conferences. That is how it progressively got known to the research community.

About the toolbox: MIRtoolbox 1.7.2

Your toolbox is very well documented. That must have taken some time.

Thank you! It was time well invested during my postdoc years in Jyväskylä. The material was in fact side-products of lectures prepared for summer courses and conference tutorials. And it progressively grew up further after that.

News

How important was it for you to have that?

Well, we considered it as an important part of the toolbox release. It was surely a teaching support, a way to explain everything to students, users and colleagues without having to re-explain everything every time. Also it is a way for us to have a clear idea of what is in the toolbox and to remember the technical background. In fact I often (or should I say always?) go back to the documentation myself when using the toolbox. :)

The modular framework of the toolbox is quite interesting. Is this how the toolbox started, or did you later decide to reorganise it this way?

When we started thinking about building a toolbox for audio analysis, I don't think we were considering a modular framework like what the toolbox has become. Other comprehensive toolboxes have been conceived in the lab before MIRtoolbox (MIDItoolbox), but also after (MoCapToolbox). Those toolboxes are conceived as a list of functions, each performing a particular task and requiring a strict syntax concerning the list of input parameters. In contrast, as soon as I started building a first small prototype for a first demo to the students, the idea of modularity was already a central idea of my vision for the toolbox architecture. You can see below a slide that I presented to the students during that first demo

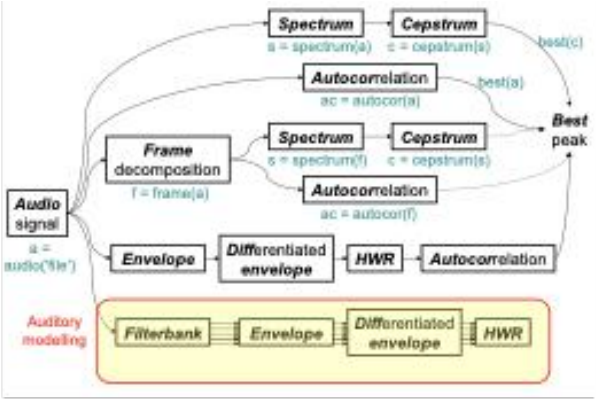

How often is the toolbox maintained now?

The latest public release (1.7.3) was in January 2019, and we should release a new update as soon as possible. These updates are not released in a very regular manner. Around the years 2008-2012 there were somewhat frequent releases (several per year), now the frequency has decreased, with sometimes one or two years between two releases.

What is built around the toolbox? (Discussion list, tweets, ...)

For MIRtoolbox 1.x, there is the discussion list, another mailing list informing about the new releases and the twitter account. For the MiningSuite, there is also a mailing list notifying about updates, as well as two discussion lists, one for users, and one for

News

developers, and the twitter account. And the usual developer network tools powered by GitHub for the project repository.

Looking at the future: MIRtoolbox 2.0

How will MIRtoolbox 2.0 be different from MIRtoolbox 1.x?

MIRtoolbox 2.0 has a completely rewritten architecture that should be more efficient. This enables also to rewrite all the code in a way that is much more readable, so that the opaque boxes become as transparent as possible, and so that all the internal processes could be understood by anyone willing to have a look. For that purpose, a new layer of syntax has been created: while MIRtoolbox 1.0 introduced this layer of syntax enabling to easily launch operations, this additional layer of syntax in MIRtoolbox 2.0/MiningSuite (cf. below) enables to write the underlying code more easily.

One important objective is to transition into a truly open-source framework, so that anybody could contribute to the further development of the project. The MiningSuite project is hosted on GitHub, and anyone can fork the repository and push changes.

Besides, the same philosophy underlying MIRtoolbox could be used outside of the purely audio MIRtoolbox context. The same design principles ruling MIRtoolbox could be applied to process symbolic representations, be it MIDI (as in MIDItoolbox) or scores more generally, better articulating audio and symbolic representations. Besides music, MIRtoolbox is used for audio analysis, but also for any kind of signal analysis. In this respect, it seems more relevant to consider the core architecture of the toolbox as independent from music or audio, and to add the audio and music components in the form of packages. That is why the new framework has the most abstract name of "MiningSuite", featuring packages called SigMinr for signal processing, AudMinr for audio analysis and MusMinr for music analysis. Other packages are considered: video analysis (VidMinr), physics and motion analysis (PhyMinr), sequence processing (SeqMinr), pattern mining (PatMinr)…

One important objective is to transition into a truly open-source framework, so that anybody could contribute to the further development of the project.

When could people expect a tested MIRtoolbox 2.0 version to roll out?

In the context of my current affiliation (RITMO, University of Oslo), I have hired a research assistant, Habibur Rahman, helping me carrying out an extensive test bed, making sure that MIRtoolbox 1.x and

News

MIRtoolbox 2.0/MiningSuite give the same results. I am so sorry this has been delayed, because I am very busy on my new research project. But well, maybe we should aim to release a public version in a few months...

Would you recommend users start with MIRtoolbox 1.72, or should they just wait for a fully tested MIRtoolbox 2.0?

They can use the current official of MIRtoolbox 1.7.2. If they are curious, they can have a look at the MiningSuite in parallel. And when it will be time to officially transition to the 2.0 version, I am preparing a documentation section explaining how to convert MIRtoolbox 1.0 syntax into MiningSuite commands.

What comes after the MIRtoolbox?

In my current research project, MIRAGE, funded for 4 years by the Norwegian Research Council, I extend further the design of the computational framework for music analysis. **One hypothesis is that pure signal processing is not sufficient to grasp the complexity of music.** So instead, we try to improve transcription techniques (using also deep learning) to represent music in the form of a symbolic representation, which is then used as the basis for musicological analyses. We try to design new music visualisation techniques as well. The rich description of music will also be used to investigate elaborate notions such as emotions, groove or mental images.

Acknowledgments:

The current development of MIRtoolbox is partially supported by the Research Council of Norway through its Centers of Excellence scheme, project number 262762, the MIRAGE project, grant number 287152 and the TIME project, grant number 249817.

Job Announcements

Student Assistant – Audiology & Signal Processing Concepts. Oticon. Smorum, Denmark.

Research Engineer – Acoustic Research. Bowers & Wilkins. Southwater, England.

Research Associate in Immersive Audio StoryTelling. Department of Electronic Engineering, University of York – Heslington Campus. Helsington, England.

Acoustic Consultant / Senior Acoustic Consultant. ACA Acoustics. London, England.

Development Engineer Acoustics. Brose Group. Belgrade, Serbia.

Research Scientist Audio/Voice. Vertex Solutions International. Brussels, Belgium.

Project Engineer – Acoustics & Vibrations Urban. TRACTEBEL. Brussels, Belgium.

PhD in Binaural 3D audio. Huawei. Munich, Germany.

Scientist position on Acoustic Virtual Reality for Environmental Noise. Empa. Dübendorf, Switzerlad.

R&D Engineer – Development Engineering. Microchip Technology. Teltow, Germany.

Signal Processing Engineer. Luko. Paris, France.

Publications

Acoustic localization of Bragg peak proton beams for hadrontherapy monitoring

INFOS

Author: Jorge E. Otero-Vega
Affiliation: Instituto de Investigación para la Gestión Integrada de Zonas Costeras (IGIC) – Universitat Politècnica de València (UPV)

Contact:
jorodve@upv.es

Hadrontherapy makes it possible to deliver high doses of energy to cancerous tumors by using the large energy deposition in the Bragg-peak. However, uncertainties in the patient positioning and/or in the anatomical parameters can cause distortions in the calculation of the dose distribution. In order to maximize the effectiveness of heavy particle treatments, an accurate monitoring system of the deposited dose depending on the energy, beam time, and spot size is necessary. The localized deposition of this energy leads to the generation of a thermoacoustic pulse that can be detected using acoustic technologies. This article presents different experimental and simulation studies of the acoustic localization of thermoacoustic pulses captured with a set of sensors around the sample. In addition, numerical simulations have been done where thermo-acoustic pulses are emitted for the specific case of a proton beam of 100 MeV.

About the author

I brought up in a little town near Bogota the capital of Colombia. I finished my sound engineer degree at the University of San Buenaventura in 2009. For five years I applied my knowledge in acoustic in the architectural and sound recording area where I develop new skills in acoustic design, building engineering, noise control inter alia. Over this time, I was working in the implementation of

new materials in the acoustic environment in order to reduce the acoustic traffic impact in the building facades. Due to the new technologies in absorbents materials, software prediction, and simulation methods I decided to start a master's engineering acoustic degree in the University Polytechnic of Valencia. In this period, I developed my final master project in the characterization and optimization of piezoelectric ceramics. As a result, I was cooperating in the study of overheated fluid applied to a black matter detector. As a result, the piezoelectric technology used in the detector gave me the possibility to work in the physics medical area. Thus, I started my thesis degree studying the localization techniques applied to the monitoring of hadrontherapy treatment. At the same time, I am collaborating with the KM3NeT project, the neutrino telescope in the Mediterranean Sea. Now, I am finishing my thesis and working in different acoustic projects.

Published in Sensors 2019, 19(9), 1971. DOI: 10.3390/s19091971