

Embedded Systems Feel the Beat in New Orleans

Highlights from the IEEE Signal Processing Cup 2017 Student Competition

Foot-tapping and moving to music is such a natural human activity, one may assume that feeling the beat in music is a simple task. Feeling the beat and then producing it, e.g., by foot tapping, is an intrinsically real-time process. As listeners, we do not wait for the beat to occur before tapping our foot; instead, we make predictions about when the next beat in the music will occur and continually revise our sense of the beat based on the accuracy of our predictions. Likewise, performing musicians have shared sense of beat, which is what allows them to play in time together.

This type of high-level music listening and understanding sits at the heart of the challenge set for this year's IEEE Signal Processing Cup (SP Cup) competition, the final stage of which concluded at the 2017 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), hosted in New Orleans, America's jazz heartland, on 5 March. The participating undergraduate cohort had to devise and construct a creative, embedded application demonstrating a real-time response to the beat of the music. Depending on the genre, composition, and rhythmic complexity of a musical piece, real-time beat tracking poses considerable challenges, which are equally present for human listeners, especially those without formal musical training. Throughout the SP Cup, the teams

confronted these challenges from both the human and computational perspectives via the choice of training and testing material, the human annotation of beat locations, the implementation and evaluation of their beat-tracking algorithms, and the response to the beat in their creative applications.

Beat tracking in music signals

The task of beat tracking of music signals has been an active area of music signal processing research for more than 25 years. While many of the earliest computational approaches sought to emulate the human process of tapping the beat in real time by making predictions of future beats [1], [2], a marked shift occurred in the early-to-mid 2000s toward offline approaches that could observe the entire musical input prior to determining beat locations.

The standard pipeline for offline beat tracking involves the explicit identification of note onset locations (or an "onset strength function," which emphasizes their location) that are subsequently passed to a tempo-estimation stage used to estimate the latent beat periodicity in the input signal, followed by the recovery of the phase (or alignment) of the beats to the music. Common techniques used to extract the beat from music signals include multiagent systems, dynamic programming, hidden Markov models, and a mixture of experts systems. Current state-of-the-art methods employ deep neural network architec-

tures to learn the relationship between labeled beat annotations in training data sets and feature representations extracted from musical audio signals, thus leveraging both advanced signal processing and machine learning.

The growth of offline approaches arose in part by the significant increase in the use of beat tracking for so-called beat-synchronous analysis as an intermediate processing step within other music signal analysis tasks, such as structural segmentation, chord detection, and music transcription. With the shift toward making multiple passes across input signals, the focus on real-time analysis was reduced. Furthermore, with a greater emphasis on the accuracy of beat tracking over computational efficiency, offline approaches also provided the opportunity for tracking the beat in music with expressive timing (i.e., changes in tempo) something that was considered impossible for real-time systems bound by the need to make predictions of future beats in the music [3].

An emerging topic related to the domain of music signal processing is creative music information retrieval, which seeks to open new possibilities for music creation, interaction and manipulation. This is facilitated by the robust analysis and interpretation of music signals [4]. For applications that target live interaction between users and/or musicians and technology, there is a compelling need to perform music signal analysis in real time. One specific motivation for the SP




FIGURE 1. An overview of real-time beat tracking. (a) An input audio signal for which the first 5 s have been acquired by a microphone. (b) A spectrogram representation of the input audio signal used to generate the onset strength function. (c) The onset strength function with overlaid beat estimates shown in black and predicted future beats in red.

Cup was, therefore, to reimagine research into beat tracking with an explicit link to real-time creative applications. From a technical perspective, real-time beat tracking, unlike offline approaches, must extract an onset strength function, estimate tempo and predict future beats based only on a continuously evolving observation of the input signal, and thus it sits firmly at the more challenging end of the spectrum. This real-time requirement also imposes strict computational limitations, a difficulty that is only increased by constraining the use of hardware in the SP Cup to embedded devices with limited computational resources. The final aspect of the competition—developing a creative application that reacts to the (predicted) beat of the music—provides an open-ended activity for the teams, but one that must be also performed in real time on the embedded device. An overview of

the process of real-time beat tracking is shown in Figure 1.

The SP Cup is an undergraduate competition organized by the IEEE Signal Processing Society (SPS) in which undergraduate students work in teams to tackle a real-life signal processing problem. Launched in 2014, the SP Cup competition has been held annually, and 2017 is the fourth edition.

To join, undergraduate students are required to form a team. Each team is composed of one faculty member to advise the team members, up to one graduate student to assist the supervisor in mentoring the team, and three to ten undergraduate students. Three top teams are selected from the initial round of competition and provided travel grants to participate in the final competition at 2017 ICASSP. The final results are described in “Winners of the SP Cup 2017.”

Tasks in the SP Cup 2017

The SP Cup challenge covered the many and multidisciplinary aspects of beat tracking, with the aim of giving students training in several areas such as music understanding and beat annotation, strategies for selecting content for training and competition, signal processing, computational optimizations for real-time performance, hardware implementations, and creative application design and development. With such a wide range of tasks and challenges to address, the SP Cup 2017 was seen as the most challenging edition so far. All of the resources related to the competition can be found at <http://sydney.edu.au/engineering/electrical/carlab/beattracking.htm>.

The open competition stage

The SP Cup started with an open competition stage from June 2016 to January 2017, consisting of two parts. The

objective for part one was to submit three 30-s musical excerpts with human-annotated beat times. The judging criteria was the quality of the beat annotations. For the first part, participants were provided with a database of 50 musical excerpts spanning a range of styles and difficulties. The database was split into two halves. One half was open, meaning that for these musical excerpts, human-annotated beat times were provided. The other half was closed so that the annotated beat times for these musical excerpts remained hidden. The purpose of the database was to assist with the development and testing of real-time beat-tracking algorithms. The task for the first part consisted of an exercise in crowdsourced beat annotation. Each participating team was required to provide human-annotated beat times for three musical excerpts of their own choosing. They also nominated one of the three musical excerpts as a challenge piece so that the beat annotations would remain hidden from the other participating teams.

To assist participants with the evaluation of their beat-tracking algorithms and give a reference for how beat-tracking accuracy would be calculated, a MATLAB evaluation script was provided. The evaluation method, extended from [5], gives an accuracy score based on a comparison of estimated beat times with annotated ground truth. It calculates the proportion of continuously correct beat estimates occurring with a perceptually specified tolerance window around the ground truth annotations. To mirror the ambiguity in human perception of the beat in music, estimated beats at perceptually related metrical levels to the ground truth annotations (e.g., twice or half the tempo of the ground truth for music in 4/4 time) were also considered correct.

The first part of the open competition was devised to serve multiple purposes. From the perspective of the teams wishing to participate, the annotation of three musical excerpts provided a relatively low barrier for entry, while also offering teams the chance to actively shape the SP Cup through their personal choice of musical content. For the organizers, the use of team submitted content led to the

creation of a totally new annotated data set for beat tracking (free from sampling bias) and, furthermore, one that could reflect the cultural diversity of the teams who participated.

For the second part of the open competition, participants had to develop and implement their beat-tracking algorithm on an embedded device (the choice was left open, but most used the Raspberry Pi for beat tracking and an Arduino for control of the output) so that it achieved real-time performance. The objectives for part two were the following:

- 1) real-time embedded software with instructions on how to run it
- 2) beat-time output for the real-time embedded device for the database and participant submitted musical excerpts
- 3) a video demonstrating real-time operation
- 4) a report in the form of an IEEE conference paper.

The judging criteria were a performance score for the real-time embedded algorithm and a creative application score. Participants then had to design and construct a creative application for their real-time beat-tracking device. In addition to submitting the beat-tracking output of their systems across all of the available musical material as well as providing source code with installation instructions, participants also had to submit a report in the form of an IEEE conference paper and post a video online demonstrating the creative application and real-time operation. This year's SP Cup is unique in that the competition included real-time constraints as well as a creative application.

The teams were evaluated on three main components submitted across both parts of the open competition. In the first part, a team of experts active in beat-tracking research assessed the subjective quality of the annotations and made corrections where necessary so as to ensure their validity as ground truth. In the second part, the submitted beat times provided by each team on the musical material without released annotations were evaluated using the publicly available MATLAB script. In addition, the creativity of the

demonstrated applications were assessed, again by a group of experts. Since each team submitted the beat-tracking software for their real-time embedded device as part of the submission for the open competition, the real-time operation and its beat-tracking output could be verified. The final score for each team was weighted across these three components with the following proportions: one-sixth for the annotations, one-half for the real-time beat-tracking accuracy, and one-third for the creative application. A breakdown of the scores as well as a written assessment by the organizing committee was provided to all teams that participated in the second part of the open competition.

Final competition

After the judging committee evaluated the submissions from the open competition, three finalist teams were chosen to advance to the final competition. Prior to attending the final event at ICASSP, each team was required to submit additional annotated challenge excerpts to be used for on-site evaluation. However, in contrast to earlier stages in this year's competition, neither the audio nor the annotations were made available to the other teams.

The final SP Cup event was held at ICASSP in New Orleans, Louisiana, on 5 March. For the first time since the inception of the SP Cup, a live demo session was included in the final event. The event started by testing the accuracy of the real-time beat-tracking embedded devices in real-world conditions with the audio of the newly submitted challenge pieces captured by microphones (Figure 2). The finalist teams were then allowed time to set up their live demos. Each team then presented its beat-tracking algorithm, its implementation, and the design and development of the creative application. This was followed by a live demonstration of the creative application and a question and answer session. The final judging committee convened and selected the first-, second-, and third-prize winners as well as present-ed honorable mentions.

Winners of the SP Cup 2017

Grand Prize: Team Beats on the Barbie

- University of New South Wales
- Undergraduate students: Angus Keatinge, Max Fisher, Jeremy Bell, and James Wagner
- Supervisor: Vidhyasaharan Sethu
- Video: <https://www.youtube.com/watch?v=VkoGZnVEsfw>
- Technical Approach: Team Beats on the Barbie (Figure S1) adapted and optimized an existing real-time beat-tracking algorithm [6] for Raspberry Pi. They controlled their creative application, a robotic drumming system (see Figures S2 and S3), using an Arduino Mega. The robotic drummer can play back a drum part encoded as an Arduino sketch, and during the final competition it accompanied team members Jeremy Bell and James Wagner in a performance of John Lennon's "Imagine." Due to the use of high-powered solenoid drivers and fast triggers, the system was able to play drum fills and was loud enough to require no additional amplification.



FIGURE S1. First place team: Beats on the Barbie.



FIGURE S2. Solenoid-based actuators for Team Beats on the Barbie.

Second Prize: Team Madmom

- Johannes Kepler University, Austria, and Télécom ParisTech, France
- Undergraduate students: Amaury Durand (Télécom ParisTech), Sebastian Pöll (Johannes Kepler University), and Raminta Balsyte (Johannes Kepler University)
- Supervisor: Sebastian Böck
- Graduate Mentor: Florian Krebs
- Video: <https://www.youtube.com/watch?v=Losv4GqsGYU>
- Technical Approach: Team Madmom (Figure S4) adapted a real-time beat-tracking system from the existing offline approach in the Madmom Python library [7] and used a recurrent neural network. To allow real-time operation, the bidirectional neural network was replaced with a unidirectional network. They controlled their creative application, a robotic drumming system (see Figures S5 and S3), using a Raspberry Pi. Instead of a preprogrammed drum pattern, the system inferred what to play based on the analysis of the rhythmic



FIGURE S3. The automated drums for Team Madmom and Team Beats on the Barbie. Both teams implement drum signals for the bass drum, the snare drum, and the hi-hat.



FIGURE S4. Second place team: Team Madmom.

structure of the input and was able to react to changes in a time signature. Team Madmom intends to make its system freely available and open source at <https://gitlab.cp.jku.at/ROBOD>.

Third Prize: Team PulseBox

- University of Maryland, United States
- Undergraduate students: William Heimsoth, Creed Gallagher, and Josh Preuss
- Supervisor: William Hawkins
- Video: <https://www.youtube.com/watch?v=KPwFnY6bJNl>
- Technical Approach: Team Pulsebox (Figure S6) developed all aspects of their system entirely from scratch.




FIGURE S5. The drum actuators for Team Madmom.



FIGURE S6. The third place team: Team PulseBox.

Their beattracking algorithm made use of a novel comb-snapping technique to maintain high temporal accuracy of the predicted beats and used machine learning to optimize multiple relevant parameters including those related to tempo adjustment, windows, and the choice of frequency bands. Their creative system, the PulseBox, (shown in Figure S7) was a light-emitting diode (LED) cube containing 245 LEDs arranged in a 7x7 grid on each of the five visible faces of the cube. The LEDs were individually configurable with 24-bit color and were programmed to react to the beat of the music with rotating shapes and patterns.

In addition to the three overall winning teams (Figure S8), the SP Cup 2017 judging committee made the following honorable mentions. Videos for these and other submissions can be found at <http://sydney.edu.au/engineering/electrical/carlab/beatracking.htm>.

Honorable Mention for Excellent Video Production and an Entertaining Concept

- Team NTHU-EECS, National Tsing Hua University, Taiwan.

Honorable Mention for Excellent Video Production and Accurate Ground Truth Annotation

- Team Impulse, Bangladesh University of Engineering and Technology, Bangladesh.

Honorable Mention for Excellence in Ground Truth Annotation and Beat-Tracking Performance

- Team Sharif University of Technology, Sharif University of Technology, Iran.





FIGURE S7. The rhythmic LED cube for Team Pulsebox.



(a)



(b)



(c)



(d)



(e)

FIGURE S8. A behind-the-scenes look at the SP Cup 2017 teams that received honorable mentions: (a) Team Sharif, University of Technology, Iran; (b) and (c) Team NTHU-EECS, Taiwan, with their metronome mechanism; (d) and (e) Team Impulse, Bangladesh, and their band of skeletons.

Highlights of technical approaches

For the real-time beat-tracking aspect of the SP Cup, many teams implemented methods inspired by or directly adapted existing approaches for beat tracking. Even in the cases where a reference implementation was publicly available, these required a very significant overhaul to make the algorithms real-time compatible and sufficiently optimized to run on the embedded devices.

From an algorithmic perspective, the great majority of submitted algorithms followed the standard approach for beat tracking by

- generating one or more onset strength signals (often across subbands) derived from time-frequency representations of the streaming input audio signal
- performing periodicity analysis on the onset strength signal(s) by means of autocorrelation or comb filtering

- estimating the phase of the beats by cross-correlation or dynamic programming, and then using phase as the reference point from which to predict future beat locations.

Many teams also included some higher-level modeling to provide a smooth output without rapid switching between metrical levels (i.e., tempo doubling or halving). Depending on the computational resources of the chosen embedded device (some of which were extremely low power), the beat-tracking approach had to be highly optimized, e.g., purely based on time-domain analysis. The most computationally expensive and ambitious approaches attempted to run state-of-the-art deep neural network architectures for beat prediction.

The technical approaches were invariably biased by the initial project description, which mentioned blinking LEDs and the Raspberry Pi and Arduino.

So, for example, the Raspberry Pi was the embedded platform used for beat tracking by the majority of teams (fourteen teams). A variety of other interesting embedded platforms were used by a single team: ARM mbed, NAO robot, STM32F4Discovery, and UDOO Quad. Many teams coupled the beat tracker with an Arduino to assist with the creative output. With regard to the programming language used for the embedded application, it was evenly distributed between C/C++ and Python. A wide variety of creative applications were demonstrated. Applications demonstrated by multiple teams were: LED displays (seven teams); screen displays (four teams); and automated drumming (two teams). The unique creative applications were a moving head, a dancing robot, a band of skeletons, a metronome follower, a vibration device for the hearing-impaired, and an encryption device.

SP Cup 2017 Statistics

In total, the teams from 20 different countries participated in SP Cup 2017. At the registration stage of the competition, 40 teams were involved with a total of 279 participants. For the first part of the open competition, 33 teams across 18 countries with more than 250 participants submitted musical excerpts (thus adding 99 new examples to the initial data set of 50 provided by the organizers). In the second part, which presented a significant increase in difficulty and submission requirements, 21 teams participated with 147 members spread across 14 countries. The countries with the most registrations were India with eight and the United States with seven.

As in previous years, the SP Cup was run as an online class on the Piazza platform, which, in addition to allowing continuous interaction with teams, also hosted the test material supporting documentation. In total, 115 students registered for the course, with approximately 220 contributions and 2,500 views of the posts. An archive of the class is available at https://piazza.com/ieee_sps/other/sp1701/home.

Since its inception, the SP Cup has received generous support from MathWorks, Inc., the maker of the popular MATLAB and Simulink platforms. MathWorks also provided funding support to the SP Cup and contributed their expertise. Each student team that registered for the SP Cup was provided complimentary software access to MATLAB and related toolboxes. After discussion with the SP Cup organizers, MathWorks provided skeleton code for real-time audio using a Raspberry Pi, which is available at <http://au.mathworks.com/matlabcentral/fileexchange/59825-real-time-beat-tracking-templates-for-ieee-signal-processing-cup-2017>. The IEEE SPS welcomes continued engagement and support from industry in future SP Cup competitions. Interested supporters may contact Dr. Patrizio Campisi, director for student services, at patrizio.campisi@uniroma3.it.

Participants' feedback

Throughout the open competition there was a great deal of interaction, not only



FIGURE 2. A Real-time beat-tracking assessment for the final competition: music was played from the Bluetooth loudspeaker and recorded by three microphones, one for each team.

through questions for the instructors posted to Piazza but also among the different student teams who often engaged in discussion over the provided responses. Indeed, these interactions were critical in expanding the flexibility of the evaluation script to correctly process music in non-4/4 meters. As organizers, we were delighted to see this collaborative spirit continue right through to the preparation for ICASSP and the final session itself. Next, we provide an overview of some feedback and perspectives received from the three winning teams.

Team Beats on the Barbie

■ “The project itself was extremely challenging. I worked on the software implementation of the algorithm, and to do this meant implementing the hardware interface on an embedded system. For me, the most challenging part of the SP Cup was setting up many different projects and libraries that often had never been tested on an embedded system to work in real time and simultaneously. This required running parts of the algorithm in different threads, modifying audio drivers, and writing low-level sound architecture code. Having these components running at the same time, and interacting with the hardware, was an amazing feeling.”

—Jeremy Bell, undergraduate

■ “I learned a lot about DSP while working on the SP Cup, and since I am undertaking more DSP courses this semester, I feel more confident in my ability to understand more complicated concepts. I think my future career will almost certainly involve signal processing, so I will take the skills I have learned in DSP beyond university as well.”

—Jeremy Bell, undergraduate

■ “I learned a lot about DSP algorithm design in general. I am also more confident in my understanding of sound architectures in Linux. I think I also learned a lot about teamwork, and what it takes to get things done under extreme time constraints.”

—Jeremy Bell, undergraduate

■ “ICASSP was my first conference as an undergraduate, and I found it incredible. The amount of state-of-the-art technology and innovative creations was overwhelming, and it was almost impossible to keep up with in lectures. I was also surprised by the number of social events that occurred at the conference. It was great to be able to interact with so many talented and like-minded people on such a casual and friendly basis throughout the conference.”

—Jeremy Bell, undergraduate

- “We have already received several offers for other events at which we will be demonstrating the system. To do this will require some refinement of the interface and additional work on the software to make it more robust. Upon the graduation of our team, we will also be creating a handover document, so that future students can continue working on the system.”

—Team Beats on the Barbie

Team Madmom

- “I am interested in all topics making the link between music and mathematics, machine learning. I was working on incorporating online and real-time processing in the Madmom library when Sebastian told me that this work would be really useful for the SP Cup.”

—Amaury Durand, undergraduate

- “[Attending ICASSP was] really rewarding, it was my first time at a conference and, even though it was difficult for me to understand the talks I went to, I found it really interesting to meet the people who work on the topics that interest me.”

—Amaury Durand, undergraduate

- “[Participating in the SP Cup] was a perfect match. I just finished my Ph.D. in (mostly) offline beat and downbeat tracking, so it was very exciting for us to see how we can transform our system to work online and on an embedded device. Of course, it was more work than expected, but definitely a very exciting and rewarding experience!”

—Florian Krebs, graduate mentor

- “The organization of everything was great, and I think there is no way to make this better. It was really great that you could organize a drum set, although this was not planned beforehand and not easy in a city that you don’t know.”

—Florian Krebs, graduate mentor

- “It was very challenging given the limited processing power of the

embedded device and extremely rewarding that it worked.”

—Sebastian Böck, supervisor.

Team Pulsebox

- “When I first heard of the topic for the 2017 SP Cup, I was very excited. As someone with a strong interest in both music theory and programming, I knew I had to get involved.”

—Creed Gallagher, undergraduate

- “One thing I learned a lot about was how to write truly speed-optimized code (Python with heavy use of NumPy). We had to push our Raspberry Pi to its limits. We also learned some lessons about the importance of effective communication and time management. We had to exercise a lot of discipline to complete such a big project on schedule.”

—Creed Gallagher, undergraduate

- “The signal processing challenge of beat tracking is incredibly complex! With so many types of songs and genres of music, there is no hard and fast rule as to what gives the best results. We ended up trying many approaches, many of which did not give as good results as we hoped. As a result, when we finally had something we felt performed well, it was incredibly satisfying.”

—Team Pulsebox

- “My senior project involves continual development of the PulseBox. I want to eventually create a 3-D holographic display that tracks both the beat and ‘mood’ of a song. Sebastian’s team convinced us that the future of musical analysis lies in the use of neural networks, which is the avenue I will be exploring.”

—Creed Gallagher, undergraduate

- “As an undergraduate, attending ICASSP was an amazing and humbling experience. I enjoyed listening in on the presentations which gave me a window into the cutting edge of SP

research. Plus, everyone was friendly and New Orleans was a fun venue.”

—Creed Gallagher, undergraduate

Forthcoming project competitions for undergraduates

The fifth edition of the SP Cup will be held at ICASSP 2018. The theme of the 2018 competition will be announced in September. Teams who are interested in the SP Cup competition may visit this link: <https://signalprocessingsociety.org/get-involved/signal-processing-cup>.

In addition to the SP Cup, the IEEE SPS recently announced the first edition of the Video and Image Processing Cup. The final competition will be held at the IEEE International Conference on Image Processing, in Beijing, China, 17–20 September. The theme of this competition is “Challenging Road Sign Detection.” For details, visit: <https://signalprocessingsociety.org/get-involved/video-image-processing-cup>.

Acknowledgments

As the SP Cup 2017 Organizing Committee, we would like to express our gratitude to all of the people who made this adventure a reality: the participating teams, the judging panel, the local organizers, the IEEE SPS Membership Board for its financial support for the drum kit rental, and MathWorks for its sponsorship. Matthew E.P. Davies is supported by Portuguese National Funds through the FCT-Foundation for Science and Technology, I.P., under the project IF/01566/2015.

Authors

Craig T. Jin (craig.jin@sydney.edu.au) is an associate professor at the University of Sydney, Australia. He is a Senior Member of the IEEE and a current member of the IEEE Audio and Acoustic Signal Processing Technical Committee (AASP TC). He initiated this edition of the SP Cup on behalf of the AASP TC, developed the competition specification and pedagogical materials, and ran the competition alongside Matthew Davies.

(continued on page 170)

Exposition, Atlanta, GA, June 2013, pp. 23.828.1–23.828.21.

[8] A Guide to the flipped classroom. The Chronicle of Higher Education (2015, Jan. 07) [Online]. Available: <http://www.chronicle.com/article/A-Guide-to-the-Flipped/151039/>

[9] J. L. Jensen, T. A. Kummer, and P. D. D. M. Godoy, “Improvements from a flipped classroom may simply be the fruits of active learning,” *CBE Life Sci. Educ.*, vol. 14, no. 1, pp. 1–12, Mar. 2015.

[10] E. F. Gehringer, “Resources for “flipping” classes,” in *Proc. ASEE Annu. Conf. Exposition*, Seattle, WA, June 2015, pp. 26.1336.1–26.1336.10.

[11] J. O’Flaherty and C. Phillips, “The use of flipped classrooms in higher education: A scoping review,” *Internet Higher Educat.*, vol. 25, pp. 85–95, Oct. 2015.

[12] Flip Learning: Research, Reports, and Studies. Flipped Learning Network. [Online]. Available: <http://flippedlearning.org/research-reports-studies/>

[13] B. J. Limbach and W. L. Waugh, “Questioning the lecture format,” *NEA Higher Educ. J. Thought Action*, vol. 20, no. 1, pp. 47–56, 2005.

[14] M. Freeman, P. Blayney, and P. Ginns, “Anonymity and in class learning: The case for electronic response systems,” *Australasian J. Educ. Tech.*, vol. 22, no. 4, pp. 568–580, 2006.

[15] J. R. Stowell and J. M. Nelson, “Benefits of electronic audience response systems on student participation, learning, and emotion,” *Teaching Psychol.*, vol. 34, no. 4, pp. 253–258, 2007.

[16] C. R. Graham, T. R. Tripp, L. Seawright, and G. Joeckel, “Empowering or compelling reluctant participants using audience response systems,” *Active Learn. Higher Educ.*, vol. 8, no. 3, pp. 233–258, 2007.

[17] Y. K. Kim and L. J. Sax, “Student–faculty interaction in research universities: Differences by student gender, race, social class, and first-generation status,” *Res. Higher Educ.*, vol. 50, no. 5, pp. 437–459, 2009.

[18] W. Griffin, S. D. Cohen, R. Berndtson, K. M. Burson, K. M. Camper, Y. Chen, and M. A. Smith, “Starting the conversation: An exploratory study of factors that influence student office hour use,” *College Teach.*, vol. 62, no. 3, pp. 94–99, 2014.

[19] P. C. Blumenfeld, E. Soloway, R. W. Marx, J. S. Krajcik, M. Guzdial, and A. Palincsar, “Motivating project-based learning: Sustaining the doing, supporting the learning,” *Educ. Psychol.*, vol. 26, no. 3–4, pp. 369–398, June 1991.

[20] H. A. Hadim and S. K. Esche, “Enhancing the engineering curriculum through project-based learning,” in *Proc. 32nd Annu. Frontiers in Education (FIE’02)*, Nov. 2002, vol. 2, pp. F3F.1–F3F.6.

[21] M. Frank, I. Lavy, and D. Elata, “Implementing the project-based learning approach in an academic engineering course,” *Int. J. Tech. Design Educat.*, vol. 13, no. 3, pp. 273–288, Oct. 2003.

[22] J. S. Krajcik and P. C. Blumenfeld, “Project-based learning,” in *The Cambridge Handbook of the Learning Sciences*, R. K. Sawyer, Ed. New York, NY: Cambridge Univ. Press, 2006, ch. 19, pp. 317–334.

[23] J. Bourne, D. Harris, and F. Mayadas, “Online engineering education: Learning anywhere, anytime,” *J. Eng. Educ.*, vol. 94, no. 1, pp. 131–146, Jan. 2005.

[24] L. Pappano. (Nov. 2, 2012). The year of the MOOC. The New York Times. [Online]. Available: <http://www.nytimes.com/2012/11/04/education/edlife/massive-open-online-courses-are-multiplying-at-a-rapid-pace.html>

[25] D. F. O. Onah, J. Sinclair, and R. Boyatt, “Dropout rates of massive open online courses: Behavioural patterns,” in *Proc. 6th Intl. Conf. Education and New Learning Technologies (EDULEARN’14)*, Barcelona, Spain, July 2014.

[26] T. A. Baran, R. G. Baraniuk, A. V. Oppenheim, P. Prandoni, and M. Vetterli, “MOOC adventures in signal processing: Bringing DSP to the era of massive open online courses,” *IEEE Signal Process. Mag.*, vol. 33, no. 4, pp. 62–83, July 2016.

[27] Khan Academy. [Online]. Available: <http://www.khanacademy.org/>

[28] R. H. Rockland, L. Hirsch, L. Burr-Alexander, J. D. Carpinelli, and H. S. Kimmel, “Learning outside the classroom—Flipping an undergraduate circuits analysis course,” in *Proc. ASEE Annu. Conf. Exposition*, Atlanta, GA, June 2013, pp. 23.854.1–23.854.8.

[29] B. Van Veen, “Flipping signal-processing instruction,” *IEEE Signal Process. Mag.*, vol. 30, no. 6, pp. 145–150, Nov. 2013.

[30] M. L. Fowler, “Flipping signals and systems—Course structure & results,” in *Proc. IEEE Intl. Conf. Acoustics, Speech, and Signal Processing (ICASSP’14)*, Florence, Italy, May 2014, pp. 2219–2223.

[31] G. J. Kim, M. E. Law, and J. G. Harris, “Lessons learned from two years of flipping Circuits I,” in *Proc. ASEE Annu. Conf. Exposition*, Seattle, WA, June 2015, pp. 26.1087.1–26.1087.12.

[32] M. G. Schrlau, R. J. Stevens, and S. Schley, “Flipping core courses in the undergraduate mechanical engineering curriculum: Heat transfer,” *Adv. Eng. Educ.*, vol. 5, no. 3, Nov. 2016.

[33] J. R. Buck, K. E. Wage, and J. K. Nelson, “Designing active learning environments,” *Acoustics Today*, vol. 12, no. 2, pp. 12–20, 2016.

[34] W. U. Bajwa. SigProcessing YouTube channel. [Online]. Available: <http://www.youtube.com/user/SigProcessing>

[35] Poll Everywhere. [Online]. Available: <http://www.pollerywhere.com/>

[36] R. A. Layton, M. L. Loughry, M. W. Ohland, and G. D. Ricco, “Design and validation of a web-based system for assigning members to teams using instructor-specified criteria,” *Adv. Eng. Educat.*, vol. 2, no. 1, pp. 1–28, 2010.

[37] CATME System. [Online]. Available: <http://info.catme.org/>



SP COMPETITIONS *(continued from page 150)*

Matthew E.P. Davies (matthew.davies@inesctec.pt) is a senior researcher at INESC TEC, Portugal. He is an IEEE Member and an associate editor for *IEEE/ACM Transactions on Audio, Speech, and Language Processing*. He is responsible for the technical components of the competition and its materials and ran the competition alongside Craig T. Jin.

Patrizio Campisi (patrizio.campisi@uniroma3.it) is a professor at Roma Tre University, Rome, Italy. He chairs the Student Service Committee of the IEEE Signal Processing Society and

served as the general chair of the 2015 Information Forensics and Security Technical Committee. He is the organizer for the SP Cup since the 2015 edition and the initiator of the Video and Image Processing Cup.

References

[1] M. Goto and Y. Muraoka, “A beat tracking system for acoustic signals of music,” in *Proc. 2nd ACM Int. Conf. Multimedia*, 1994, pp. 365–372.

[2] E. D. Scheirer, “Tempo and beat analysis of acoustic musical signals,” *J. Acoust. Soc. Am.*, vol. 103, no. 1, pp. 588–601, 1998.

[3] F. Gouyon and S. Dixon, “A review of automatic rhythm description systems,” *Comput. Music J.*, vol. 29, no. 1, pp. 34–54, 2005.

[4] X. Serra, M. Magas, E. Benetos, M. Chudy, S. Dixon, A. Flexer, E. Gómez, F. Gouyon, P. Herrera, S. Jordà, O. Paytuyvi, G. Peeters, J. Schlüter, H. Vinet, and G. Widmer. (2013). Roadmap for Music Information ReSearch. London: MIRE5 Consortium. [Online]. Available: <http://mires.eecs.qmul.ac.uk/about.html>

[5] M. E. P. Davies, N. Degara, and M. D. Plumbley, “Evaluation methods for musical audio beat tracking algorithms,” Queen Mary Univ., Centre for Digital Music, London, Tech. Rep. C4DM-TR-09-06, 2009.

[6] J. Oliveira, M. E. P. Davies, F. Gouyon, and L. P. Reis, “Beat tracking for multiple applications: A multi-agent system architecture with state recovery,” *IEEE Trans. Audio, Speech Lang. Process.*, vol. 20, no. 10, pp. 2696–2706, 2012.

[7] S. Böck, F. Korzeniowski, J. Schlüter, F. Krebs, and G. Widmer. “Madmom: A new Python audio and music signal processing library,” in *Proc. 2016 ACM Multimedia Conf.*, 2016, pp. 1174–1178.

