

Brian McFee

Assistant professor of Music Technology and Data Science

Music and Performing Arts Professions
Center for Data Science
New York University
60 5th Avenue
Room 621
New York, NY 10003

Email: brian.mcfee@nyu.edu
Homepage: <https://brianmcfee.net/>
Google Scholar: <https://goo.gl/SEdqCf>
Phone: +1 212 998 3260

Research interests

Music information retrieval, machine learning, recommender systems, multimedia signal processing

Education

UNIVERSITY OF CALIFORNIA, SAN DIEGO La Jolla, CA
Ph.D. in Computer Science, 2012.
Dissertation: *More like this: Machine learning approaches to music similarity*
Advisor: Gert Lanckriet.
C.Phil. in Computer Science, 2009.
M.S. in Computer Science, 2008.

UNIVERSITY OF CALIFORNIA, SANTA CRUZ Santa Cruz, CA
B.S. Computer Science with Highest Honors, 2003.

COLLEGE OF MARIN Kentfield, CA
A.A. Mathematics, 2001.

Employment

NEW YORK UNIVERSITY New York, NY
Assistant professor, 2018–present.

NEW YORK UNIVERSITY New York, NY
Moore-Sloan data science fellow, 2014–2018.

COLUMBIA UNIVERSITY New York, NY
Postdoctoral research scholar, 2012–2014.

UNIVERSITY OF CALIFORNIA, SAN DIEGO La Jolla, CA
Graduate research and teaching assistant, 2004–2012.

SOCIAL SOURCE FOUNDATION San Francisco, CA
Software engineer, 2005.

GROUNDSPRING.ORG San Francisco, CA
Software engineer, 2004.

TKO SOFTWARE
Quality assurance engineer, 2003.

Santa Cruz, CA

VA LINUX SYSTEMS
Artist, Linux.com, 1999–2001.

Sunnyvale, CA

Awards and honors

ISMIR Best student paper award, for **Deep salience representations for f_0 estimation in polyphonic music** (with R. Bittner as primary author), Suzhou, CN, 2017.

ISMIR Best oral presentation award, for **Analyzing song structure with spectral clustering**, Taipei, TW, 2014.

ISMIR Best poster presentation award, for **mir_eval: a transparent implementation of common MIR metrics**, Taipei, TW, 2014.

NIPS Reviewer Award, Neural Information Processing Systems, Lake Tahoe, CA, 2013.

Music Hack Day — Soundcloud prize, for **Mend-a-break**, Philadelphia, PA, 2013.

Music Hack Day — Echo Nest prize, for **Auto-chip-tune**, Boston, MA, 2012.

ISMIR Best poster presentation award, for **The natural language of playlists**, Miami, FL, 2011.

Qualcomm Innovation Fellowship, for **Location-, Demographic-, Preference- and Content-Based Music Search and Recommendation**, 2010–2011. (With L. Barrington.)

ISMIR Best presentation award, for **Heterogeneous embedding for subjective artist similarity**, Kobe, JP, 2009.

Publications

Journal articles

D. Huppenkothen, B. McFee, and L. Norén. Entropy your cohort: A transparent method for diverse cohort selection. *PLoS One*, 2020.

B. McFee, J. W. Kim, M. Cartwright, J. Salamon, R. M. Bittner, and J. P. Bello. Open-source practices for music signal processing research: Recommendations for transparent, sustainable, and reproducible audio research. *IEEE Signal Processing Magazine*, 36(1):128–137, Jan 2019. ISSN 1053-5888. doi: 10.1109/MSP.2018.2875349.

V. Lostanlen, J. Salamon, M. Cartwright, B. McFee, A. Farnsworth, S. Kelling, and J. P. Bello. Per-channel energy normalization: Why and how. *IEEE Signal Processing Letters*, 26(1):39–43, Jan 2019a. ISSN 1070-9908. doi: 10.1109/LSP.2018.2878620.

B. McFee, J. Salamon, and J. P. Bello. Adaptive pooling operators for weakly labeled sound event detection. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 26(11):2180–2193, Nov 2018. ISSN 2329-9290. doi: 10.1109/TASLP.2018.2858559

B. McFee, O. Nieto, M. Farbood, and J.P. Bello. Evaluating hierarchical structure in music annotations. *Frontiers in Psychology*, 8:1337, 2017b. ISSN 1664-1078. doi: 10.3389/fpsyg.2017.01337. URL <http://journal.frontiersin.org/article/10.3389/fpsyg.2017.01337>.

- B. McFee. resampy: efficient sample rate conversion in python. *The Journal of Open Source Software*, 2016. doi: 10.21105/joss.00125.
- Y. Vaizman, B. McFee, and G.R.G. Lanckriet. Codebook based audio feature representation for music information retrieval. *IEEE Transactions on Audio, Speech, and Language Processing*, 22(10):1483–1493, 2014.
- C. Galleguillos, B. McFee, and G.R.G. Lanckriet. Iterative category discovery via multiple kernel metric learning. *International Journal of Computer Vision*, pages 1–18, 2013.
- B. McFee, L. Barrington, and G.R.G. Lanckriet. Learning content similarity for music recommendation. *IEEE Transactions on Audio, Speech, and Language Processing*, 20(8), 2012a.
- B. McFee and G.R.G. Lanckriet. Learning multi-modal similarity. *Journal of Machine Learning Research*, 12:491–523, February 2011.
- B. McFee, C. Galleguillos, and G.R.G. Lanckriet. Contextual object localization with multiple kernel nearest neighbor. *IEEE Transactions on Image Processing*, 20(2):570–585, Feb 2011. ISSN 1057-7149.

Peer-reviewed conference publications

- V. Lostanlen, S. Sridhar, B. McFee, A. Farnsworth, and J.P. Bello. Learning the helix topology of musical pitch. In *IEEE International conference on acoustics, speech and signal processing*, ICASSP, 2019b.
- B. McFee and K. Kinnaird. Improving structure evaluation through automatic hierarchy expansion. In *20th International Society for Music Information Retrieval Conference*, ISMIR, 2019.
- A. Cohen-Hadria, M. Cartwright, B. McFee, and J.P. Bello. Voice anonymization in urban sound recordings. In *IEEE International Workshop on Machine Learning for Signal Processing*, 2019.
- C. Tralie and B. McFee. Enhanced hierarchical music structure annotations via feature level similarity fusion. In *IEEE International conference on acoustics, speech and signal processing*, ICASSP, 2019.
- M. Fuentes, B. McFee, H. Crayencour, S. Essid, and J.P. Bello. A music structure informed downbeat tracking system using skip-chain conditional random fields and deep learning. In *IEEE International conference on acoustics, speech and signal processing*, ICASSP, 2019.
- E.J. Humphrey, S. Durand, and B. McFee. OpenMIC-2018: An open dataset for multiple instrument recognition. In *19th International Society for Music Information Retrieval Conference*, ISMIR, 2018.
- M. Fuentes, B. McFee, H. Papadopoulos, S. Essid, and J.P. Bello. Analysis of common design choices in deep learning systems for downbeat tracking. In *19th International Society for Music Information Retrieval Conference*, ISMIR, 2018.
- V. Trinh, B. McFee, and M. Mandel. Bubble cooperative networks for identifying important speech cues. In *InterSpeech*, 2018.
- B. McFee and J.P. Bello. Structured training for large-vocabulary chord recognition. In *18th International Society for Music Information Retrieval Conference*, ISMIR, 2017.
- R.M. Bittner, B. McFee, J. Salamon, P. Li, and J.P. Bello. Deep salience representations for f_0 estimation in polyphonic music. In *18th International Society for Music Information Retrieval Conference*, ISMIR, 2017. Winner of the best student paper award.
- B. McFee, E.J. Humphrey, and J. Urbano. A plan for sustainable MIR evaluation. In *17th International Society for Music Information Retrieval Conference*, ISMIR, 2016.

- B. McFee, E.J. Humphrey, and J.P. Bello. A software framework for musical data augmentation. In *16th International Society for Music Information Retrieval Conference, ISMIR, 2015a*.
- B. McFee, O. Nieto, and J.P. Bello. Hierarchical evaluation of segment boundary detection. In *16th International Society for Music Information Retrieval Conference, ISMIR, 2015b*.
- B. McFee and D.P.W. Ellis. Analyzing song structure with spectral clustering. In *Proceedings of the 15th International Society for Music Information Retrieval Conference, ISMIR, 2014c*.
- C. Raffel, B. McFee, J. Salamon, E.J. Humphrey, O. Nieto, D. Liang, and D.P.W. Ellis. mir_eval: a transparent implementation of common MIR metrics. In *Proceedings of the 15th International Society for Music Information Retrieval Conference, ISMIR, 2014*.
- Z. Chen, B. McFee, and D.P.W. Ellis. Speech enhancement by low-rank and convolutive dictionary spectrogram decomposition. In *Interspeech, 2014*.
- B. McFee and D.P.W. Ellis. Learning to segment songs with ordinal linear discriminant analysis. In *IEEE International conference on acoustics, speech and signal processing, ICASSP, 2014b*.
- B. McFee and D.P.W. Ellis. Better beat tracking through robust onset aggregation. In *IEEE International conference on acoustics, speech and signal processing, ICASSP, 2014a*.
- D.K. Lim, B. McFee, and G.R.G. Lanckriet. Robust structural metric learning. In *Proceedings of the 30th annual International Conference on Machine Learning, ICML, June 2013*.
- B. McFee and G.R.G. Lanckriet. Hypergraph models of playlist dialects. In *Proceedings of the 13th International Society for Music Information Retrieval Conference, ISMIR, 2012*.
- J. Urbano, J.S. Downie, B. McFee, and M. Schedl. How significant is statistically significant? The case of audio music similarity and retrieval. In *Proceedings of the 13th International Society for Music Information Retrieval Conference, ISMIR, 2012*.
- B. McFee and G.R.G. Lanckriet. The natural language of playlists. In *Proceedings of the 12th International Society for Music Information Retrieval Conference, ISMIR, pages 537–541, 2011b*. Winner of the best poster presentation award.
- B. McFee and G.R.G. Lanckriet. Large-scale music similarity search with spatial trees. In *Proceedings of the 12th International Society for Music Information Retrieval Conference, ISMIR, pages 55–60, 2011a*.
- C. Galleguillos, B. McFee, S. Belongie, and G.R.G. Lanckriet. From region similarity to category discovery. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition, CVPR, 2011*.
- B. McFee, L. Barrington, and G.R.G. Lanckriet. Learning similarity from collaborative filters. In *Proceedings of the 11th International Society for Music Information Retrieval Conference, ISMIR, 2010*.
- N. Koenigstein, G.R.G. Lanckriet, B. McFee, and Y. Shavitt. Collaborative filtering based on P2P networks. In *Proceedings of the 11th International Society for Music Information Retrieval Conference, ISMIR, August 2010*.
- B. McFee and G.R.G. Lanckriet. Metric learning to rank. In *Proceedings of the 27th International Conference on Machine Learning, ICML, pages 775–782, Haifa, Israel, June 2010a*.
- C. Galleguillos, B. McFee, S. Belongie, and G.R.G. Lanckriet. Multi-class object localization by combining local contextual interactions. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition, CVPR, pages 113–120, 2010*.

B. McFee and G.R.G. Lanckriet. Heterogeneous embedding for subjective artist similarity. In *Proceedings of the 10th International Society for Music Information Retrieval Conference, ISMIR, 2009a*. Winner of the best presentation award.

B. McFee and G.R.G. Lanckriet. Partial order embedding with multiple kernels. In *Proceedings of the 26th International Conference on Machine Learning, ICML, pages 721–728, 2009b*.

Book chapters

B. McFee. Statistical methods for scene and event classification. In T. Virtanen, M. Plumbley, and D. Ellis, editors, *Computational Analysis of Sound Scenes and Events*, pages 103–146. Springer International Publishing, Cham, 2018. ISBN 978-3-319-63450-0. doi: 10.1007/978-3-319-63450-0_5. URL https://doi.org/10.1007/978-3-319-63450-0_5.

M. Schedl, P. Knees, B. McFee, D. Bogdanov, and M. Kaminskas. Music recommender systems. In F. Ricci, L. Rokach, and B. Shapira, editors, *Recommender Systems Handbook*, pages 453–492. Springer US, 2nd edition, 2015.

Demos, workshops, and other work

B. McFee and T. Kell. ISMIR tutorial on open source and reproducible MIR research, 2018.

T. Kell, B. McFee, and B. Lacker. Amen, amen server, amen.js: Tools for algorithmic remixing, 2017. ISMIR Late-breaking/demo.

J. Salamon, B. McFee, and P. Li. DCASE 2017 submission: Multiple instance learning for sound event detection. Technical report, DCASE2017 Challenge, September 2017.

B. McFee, E.J. Humphrey, and C. Jacoby. Pescador: a stream manager for iterative learning. In *Scientific computing with Python, SciPy, 2017a*.

B. McFee and F. Upham. Interpreting musical taste through playlists. In *Meeting of the Society for Music Perception and Cognition, SMPC, 2015*.

B. McFee, C. Raffel, D. Liang, D.P.W. Ellis, M. McVicar, E. Battenberg, and O. Nieto. librosa: Audio and music signal analysis in python. In *14th annual Scientific Computing with Python conference, SciPy, July 2015c*.

B. McFee and D.P.W. Ellis. DP₁, MP₁, MP₂ entries for MIREX₂₀₁₃ structural segmentation and beat tracking, 2013. Ninth Music Information Retrieval Evaluation eXchange.

B. McFee, T. Bertin-Mahieux, D.P.W. Ellis, and G.R.G. Lanckriet. The million song dataset challenge. In *Proceedings of the 2nd International Workshop on Advances in Music Information Retrieval, AdMIRE, 2012b*.

B. McFee and G.R.G. Lanckriet. Audio music similarity via metric learning, 2011c. Seventh Music Information Retrieval Evaluation eXchange.

B. McFee and G.R.G. Lanckriet. Integrating multi-modal data by metric learning, 2010b. 11th ACM SIGMM International Conference on Multimedia Information Retrieval.

Works in progress

R. Bittner, B. McFee, and J.P. Bello. Multi-task learning for fundamental frequency estimation in music, 2018. Submitted to IEEE-TASLP.

Invited seminars

Discovering multi-level structure in music, Data Science Tea, University of Massachusetts Amherst, 2019.

Making efficient use of musical annotations, Machine Learning for Music Discovery Workshop, ICML, Long Beach, CA, USA, 2019.

Music information retrieval, SpaRTaN/MacSeNet Summer school (SPARS conference), Lisbon, PT, 2017.

Exposing hierarchy through graph analysis, Schloss Dagstuhl seminar on Computational Music Structure Analysis, Wadern, DE, 2016.

The role of structure analysis in music discovery, Machine Learning for Music Discovery Workshop, ICML, Lille, FR, 2015.

Music information retrieval: challenges in jazz, Jazz and Technology Forum, New York, NY, 2013.

Modeling playlists by context, Fifth Recommender Stammtisch, Berlin, DE, December 2012.

Professional service

Program committee co-chair, *International Society for Music Information Retrieval (ISMIR) Conference*, 2020.

Mentor, *Women in Music Information Retrieval mentoring program*, 2017–.

Editor, *Journal of Open Source Software (JOSS)*, 2020—.

Senior program committee, *International Joint Conference on Artificial Intelligence (IJCAI)*, 2020.

Co-organizer, *OpenMIC: Open music informatics challenge*, 2016–.

Local organization, *International Society for Music Information Retrieval (ISMIR) Conference*, 2016.

Panelist, *National Science Foundation (Information & Intelligent Systems division)*, 2014.

Program committee member, *International Society for Music Information Retrieval (ISMIR) Conference*, 2014, 2017.

Committee member, *J-Disc task force*, 2014–2017.

Co-organizer of the *Million Song Dataset Challenge*, with T. Bertin-Mahieux, D.P.W. Ellis, and G.R.G. Lanckriet, 2012.

Reviewer, *Journal of Open Source Software (JOSS)*, 2018–.

Reviewer, *Journal of Open Research Software (JOSS)*, 2017.

Reviewer, *International Conference on Machine Learning (ICML)*, 2012–2018.

Reviewer, *International Conference on Learning Representations (ICLR)*, 2018–.

Reviewer, *Neural Information Processing Systems (NeurIPS)*, 2011–.

Reviewer, *Journal of Machine Learning Research*, 2011–2016.

Reviewer, International Society for Music Information Retrieval (ISMIR) Conference, 2009–.

Reviewer, IEEE Transactions on Pattern Analysis and Machine Intelligence, 2012–2017.

Reviewer, IEEE Transactions on Multimedia, 2012–2013.

Reviewer, IEEE Transactions on Audio, Speech and Language Processing, 2010–.

Program committee co-chair, Second Workshop on Music Recommendation and Discovery (WOMRAD), with A. Anglade, Ò. Celma, B. Fields and P. Lamere, 2011.

Program committee member, International Joint Conference on Artificial Intelligence, 2011.

Co-organizer of the Understanding Multiple Kernel Learning Methods Workshop, with F. Bach, G.R.G. Lanckriet, and N. Srebro, NIPS 2009.

Manager of the UC San Diego Multiple Kernel Learning Repository, <http://mkl.ucsd.edu>, 2009–2012.

University service

CDS Faculty Fellows search committee, 2019, 2020.

M.S. Admissions committee, Center for Data Science, NYU, 2015, 2017.

Doctoral candidacy examination committee, Music Technology, NYU, 2016.

Co-organizer of the Center for Data Science weekly lunch seminar, NYU, 2015.

Ph.D. committees

Guillaume Doras, UPMC, 2020.

Jong Wook Kim, Music Technology, New York University, 2019.

Brandon Morton, Electrical Engineering, Drexel University, 2016.

Colin Raffel, Electrical Engineering, Columbia University, 2016.

Teaching experience

Instructor, DS-GA 1004, Big Data, NYU, Spring 2019, 2020.

Instructor, MPATE-GE 2599, Fundamentals of Digital Signal Theory I, NYU, Fall 2018, 2019.

Co-instructor, CCRMA Workshop on Music Information Retrieval, Stanford University, Summer 2018.

Guest lecturer, ELEN6820, Speech and audio signal processing, Columbia University, 2017.

Co-instructor, DS-GA 1009, Practical training for data science, New York University, 2016.

Guest lecturer, DS-GA 1008, Deep learning, New York University, 2016.

Co-instructor, EECS E6891, Reproducing computational results, Columbia University, 2014.

Guest lecturer, CS5660, Signal and image processing, Cornell Tech, 2013.

Guest lecturer, CS5785, Modern analytics, Cornell Tech, 2013.

Teaching assistant mentor, Computer Science & Engineering, UCSD, 2010–2012.

Teaching assistant, ECE 30, Introduction to computer engineering, UCSD, 2010.

Teaching assistant, CSE 21, Mathematics for algorithms and systems, UCSD, 2009.

Teaching assistant, CSE 20, Discrete mathematics, UCSD, 2008.

Teaching assistant, CSE 100, Advanced data structures, UCSD, 2008.

Teaching assistant, CSE 101, Design and analysis of algorithms, UCSD, 2007.

Teaching assistant, CSE 151, Introduction to artificial intelligence, UCSD, 2006.

Teaching assistant, CSE 105, Theory of computation, UCSD, 2006.

Teaching assistant, CSE 3, Fluency in information technology, UCSD, 2005.

Lab tutor, CMPE 12C, Computer organization, UCSC, 2002.

EOPS tutor, Computer science and mathematics, College of Marin, 1998–2001.