Katerina Kosta

Flat 61, 1 Pilot Walk, London, SE10 0UN, UK **|** katkost@gmail.com **|** 07818958018
www.linkedin.com/in/katerinakosta **|** www.katerinakosta.com

Key Skills

Machine Learning | Generative Models |
Information Retrieval | Music/Audio Processing
Python, PyTorch, TensorFlow, scikit-learn, GitHub, TDD, CI, LaTeX

Work Experience

Senior Research Scientist
AI lab: Speech, Audio and Music Intelligence team
TikTok (BYTEDANCE)

April 2019 - Present

- Implemented end-to-end machine learning/deep learning models in the audio and symbolic music domain, specialised in generative and classification systems.
- Created pipelines for data processing and analysis.
- Designed and executed product-driven research.
- Presented results to stakeholders and external teams.
- Built containerised applications and web interfaces for demonstrations.
- Filed patents for innovative algorithms and submitted accepted conference papers accompanied with open source code.
- Delivered presentations to research labs and conferences.
- Mentored an intern which resulted in a full-time hire.
- Worked as part of a distributed team between the UK, US and China.

Machine Learning Researcher
JUKEDECK

February 2016 - April 2019

London startup for creating and productionizing automatic music generation content:
- Designed and ran experiments to improve the core deep music generation models.
- Gathered, cleaned and processed training data.
- Directly managed an intern focused on audio processing research.
- Delivered presentations to research labs and conferences (including a keynote speech at the International Conference of Computational Creativity, 2018).
- Used industry best practices throughout, including pairing, testing and code reviewing.

Research Intern
Intelligent Music System Group R&D
YAMAHA

October 2015 - December 2015

- Worked on a [project](#) aiming to re-create a piano recording on a Yamaha disklavier.
- Built a machine learning system for estimating note velocities from audio, based on non-negative matrix factorisation and source separation techniques.

Education

Ph.D. Electronic Engineering and Computer Science
Centre for Digital Music, Queen Mary University of London
2012 - 2016

- Principal's Interdisciplinary Full Scholarship.
- Thesis: "Computational Modelling and Quantitative Analysis of Dynamics in Performed Music", Supervisors: Prof. Elaine Chew, Dr. Oscar F. Bandtlow.
- Published work related to machine learning techniques for prediction and classification tasks for analysing and modelling expressiveness in music performance.
- Developed algorithms for time series analysis.
- Teaching Assistant for several undergraduate courses.

MSc. Sound and Music Computing
Music Technology Group, Universitat Pompeu Fabra, Barcelona
2010 - 2011

- Thesis: "Unsupervised Generation of Chord Sequences from a Sound Example".
- Created an unsupervised learning system that generates a sound sequence from an original audio chord sequence, leading to a top-tier conference publication.

BSc. Mathematics
National and Kapodistrian University of Athens, Greece
2004 - 2010

- Specialised in applied mathematics.

Music
Filippos Nakas Conservatory, Athens, Greece
2005 - 2010

- Diploma in classical piano.
- Music theory studies.
- Member of chamber music groups and choirs.

Hobbies

Travelling, reading books, participating in music choirs, learning new music instruments.

Publications

Kosta, K., Lu, W. C., Medeot, G., Chanquion, P. (2022). A deep learning method for melody extraction from a polyphonic symbolic music representation. In Proceedings of the 23rd International Society for Music Information Retrieval Conference (ISMIR), pp. 757-763. (link) (code)

Zhang, D., Wang, J-C., Kosta, K., Smith, J. B. L., Zhou, S. (2022). Modeling the rhythm from lyrics for melody generation of pop songs. In Proceedings of the 23rd International Society for Music Information Retrieval Conference (ISMIR). (link)

Micchi, G., Kosta, K., Medeot, G., Chanquion, P. (2021). A deep learning method for enforcing coherence in automatic chord recognition. In Proceedings of the 22nd International Society for Music Information Retrieval Conference (ISMIR), pp. 443-451. (link) (code)

Medeot, G., Cherla, S., Kosta, K., McVicar, M., Abdallah, S., Selvi, M., Newton-Rex, E., & Webster, K. (2021). A method of generating music data. US Patent App. 16/967,064. (link)

Medeot, G., Cherla, S., Kosta, K., McVicar, M., Abdallah, S., Selvi, M., Newton-Rex, E., & Webster, K. (2018). StructureNet: Inducing Structure in Generated Melodies. In Proceedings of the 19th International Society for Music Information Retrieval Conference (ISMIR), pp. 725-731, Paris, France. (link)

Kosta, K., O. F. Bandtlow, E. Chew (2018). Dynamics and relativity: Practical implications of dynamic markings in the score. Journal of New Music Research, 47(5): 438-461. (link)

Kosta, K., O. F. Bandtlow, E. Chew (2018). MazurkaBL: Score-aligned loudness, beat, and expressive markings data for 2000 Chopin Mazurka recordings. In Proceedings of the 4th International Conference on Technologies for Music Notation and Representation (TENOR), pp. 85-94, Montreal, Canada. (link)

Kosta, K. (2017). Computational Modelling and Quantitative Analysis of Dynamics in Performed Music. Ph.D. Thesis. Centre for Digital Music, Queen Mary University of London, London, UK. (link)

Kosta, K., O. F. Bandtlow, E. Chew (2017). Dynamic change points in music audio capture dynamic markings in score.18th International Society for Music Information Retrieval Conference (ISMIR), Late-Breaking and Demo Session, Suzhou, China. (link)

Kosta, K., R. Ramirez, O. F. Bandtlow, E. Chew (2016). Mapping between dynamic markings and performed loudness: A machine learning approach. Journal of Mathematics and Music, 10(2): 149-172. (link)

Kosta, K., O. F. Bandtlow, E. Chew (2016). Outliers in Performed Loudness Transitions: An Analysis of Chopin Mazurka Recordings. In Proceedings of the 14th International Conference for Music Perception and Cognition (ICMPC), pp. 601-604, July 5-9, 2016, San Francisco, California, USA. (link)

Kosta K., R. Ramirez, O. F. Bandtlow, E. Chew (2015). Predicting loudness levels and classifying dynamic markings in recorded music. In Proceedings of 8th International Workshop on Machine Learning and Music (MML2015), Machine Learning for Music Generation, Vancouver, Canada. (link)

Kosta, K., O. F. Bandtlow, E. Chew (2015). A Change-point Approach Towards Representing Musical Dynamics. In T. Collins, D. Meredith, A. Volk (eds.): Mathematics and Computation in Music: 5th International Conference, MCM 2015, London, UK, June 22-25, 2015, Proceedings, pp. 179-184, Lecture Notes in Computer Science 9110, Berlin: Springer. (link)

Kosta, K., Li, S. (2014). 2013 Performance Studies Network International Conference. Computer Music Journal, 38(2): 78-80. (link)

Kosta, K., O. F. Bandtlow and E. Chew (2014). A Study of Score Context-dependent Dynamics in Piano Performance (abstract). In Proceedings of the Performance Studies Network International Conference (PSN3), Jul 17-20, Cambridge, UK. (link)

Kosta, K., O. F. Bandtlow, E. Chew (2014). Practical Implications of Dynamic Markings in the Score: Is piano always piano? In Proceedings of the 53rd Audio Engineering Society (AES) Meeting on Semantic Audio, Jan 26-29, London, UK. (link)

Kosta, K., Y. Song, G. Fazekas, M. Sandler (2013). A Study of Cultural Dependence of Perceived Mood in Greek Music. In Proceedings of the 14th International Society for Music Information Retrieval (ISMIR), pp. 317-322, Nov 4-8, Curitiba, Brazil. (link)

Kosta, K., M. Marchini, H. Purwins (2012). Unsupervised Chord-Sequence Generation from an Audio Example. In Proceedings of the 13th International Society for Music Information Retrieval (ISMIR), pp. 481-486, Porto, Portugal. (link)