

Gus (Guangyu) Xia

1158, NYU Shanghai, Shanghai, 200122
Email: gxia@nyu.edu Tel: (412)-979-0662
Webpage: <http://www.cs.cmu.edu/~gxia/>

RESEARCH INTEREST

I design intelligent systems to *understand* and *extend* musical creativity and expression. To *understand* means to learn the musical representation conveyed through sounds, performances, and symbolic compositions. To *extend* means to use such an understanding to create artificial music partners, serving music lovers at all levels.

My team builds interactive music agents who 1) compose and arrange music via style transfer and analogy, 2) perform accompaniment and improvise expressively in concert with human musicians by learning from rehearsal experience, and 3) teach music beginners using multimodal feedbacks. Other projects include content-based music information retrieval, autonomous dancing robots, and bio-music computing using slime molds.

ACADEMIC EXPERIENCE

- Sep 2017 – Present* **NYU Shanghai**, Shanghai, China
 Assistant Professor, Computer Science
- Aug 2016 – Aug 2017* **Dartmouth College**, Hanover, New Hampshire, USA
 Neukom Postdoc Fellow, Digital Music Center
- May 2013 – Aug 2016* **Carnegie Mellon University**, Pittsburgh, Pennsylvania, USA
 Research Assistant, Machine Learning

EDUCATION

- May 2010 – Aug 2016* **Carnegie Mellon University**, Pittsburgh, Pennsylvania, USA
 Ph.D., Machine Learning
 — Thesis Topic: Expressive Collaborative Music Performance via Machine Learning
 — Advisor: Prof. Roger Dannenberg
- Aug 2006 – Jul 2010* **Peking University**, Beijing, China
 B.S., Information Management and Information System
 B.S. (Minor), Psychology
- Aug 2004 – Jul 2010* **China Conservatory of Music**, Beijing, China
 — Private DI (Chinese flute) performance study
 — Teacher: Prof. Weiliang Zhang, master DI player and the Dean of Chinese Music Department

PUBLICATIONS

- 2020 K. Chen, **G. Xia**, D. Shlomo. “Continuous Melody Generation via Disentangled Short-Term Representations and Structural Conditions,” in *Proc. 14th International Conference on Semantic Computing (ICSC)*. Expected in 2020.
- J. Jiang, **G. Xia**. “Transformer VAE: A Hierarchical Model for Structure-aware and Interpretable Music Representation Learning,” in *Proc. 45th International Conference on Acoustics, Speech and Signal Processing*. Expected in 2020.

PUBLICATIONS (continued)

- 2019
- R. Yang, D. Wang, Z. Wang, T. Chen, J. Jiang and **G. Xia**. "Deep Music Analogy Via Latent Representation Disentanglement," in *Proc. 20th International Society for Music Information Retrieval Conference*, Delft, Nov 2019.
- J. Jiang, K. Chen, W. Li, and **G. Xia**. "Large Vocabulary Chord Transcription via Chord Structure Decomposition," in *Proc. 20th International Society for Music Information Retrieval Conference*, Delft, Nov 2019.
- Y. Zhang, Y. Li, D. Chin, and **G. Xia**. "Adaptive Multimodal Music Learning via Interactive-haptic Instrument," in *Proc. 19th The International Conference on New Interfaces for Musical Expression*, Brazil, June 2019.
- R. Yang, T. Chen, Y. Zhang and **G. Xia**. "Inspecting and Interacting with Meaningful Music Representations Using VAE," in *Proc. 19th The International Conference on New Interfaces for Musical Expression*, Brazil, June 2019
- J. Jiang, **G. Xia**, R. Dannenberg. "Representing Music Structure by Variational Attention", in *International Workshop on Machine Learning for Music Discovery at ICML*, California, June 2019.
- M. Xu, Z. Wang, and G. Xia. "Transferring Piano Performance Control Across Environments," in *Proc. 44th International Conference on Acoustics, Speech and Signal Processing*, UK, May 2019.
- K. Chen, W. Zhang, S. Dubnov, **G. Xia**, W. Li. "The effect of explicit structure encoding of deep neural networks for symbolic music generation," in *2019 International Workshop on Multilayer Music Representation and Processing (MMRP)*, Milan, Jan 2019.
- 2018
- Z. Wang, and **G. Xia**. "A Framework for Automated Pop-song Melody Generation with Piano Accompaniment Arrangement." *6th National Conference on Sound and Music Technology*, Oct 2018
- G. Xia**, S.Dai. "Music Style Transfer: A Position Paper", *6th International Workshop on Musical Metacreation*, Spain, June 2018.
- G. Xia**, C. Jacobsen, Q. Chen, X-D. Yang, and R. Dannenberg. "ShIFT: A Semi-haptic Interface for Flute Tutoring." in *Proc. 18th The International Conference on New Interfaces for Musical Expression*, USA, June 2018.
- 2017
- S. Dai, **G. Xia**. "Computational Models For Common Pipa Techniques", *best student paper, the 5th National Conference on Sound and Music Technology*, Oct 2017.
- S. Dai, **G. Xia**. "Computational Models and MusicXML Definitions for Common Pipa Techniques", Late Breaking Demo, in *Proc. 18th International Society for Music Information Retrieval Conference*, Suzhou, Oct 2017.
- G. Xia**, R. Dannenberg. "Duet Interaction: Learning Improvisation Techniques for Automatic Accompaniment", in *Proc. 17th The International Conference on New Interfaces for Musical Expression*, Copenhagen, June 2017
- 2016
- G. Xia**, M. Kawai, K. Matsuki, M. Fu, S. Cosentino, G. Trovato, R. Dannenberg, S.Sessa and A. Takanishi. "Expressive humanoid robot for automatic accompaniment", *The 13th Sound and Music Computing Conference*, Hamburg, Germany, September 2016.
- G. Xia**, "Expressive and Collaborative Music Performance via Machine Learning", *CMU-MLD Doctoral Thesis*, August 2016.
- 2015
- G. Xia**, R. Dannenberg "Spectral Learning for Expressive Interactive Ensemble Music Performance", in *Proc. 16th International Conference on Music Information Retrieval*, Malaga, October 2015.
- M. Fu, **G. Xia** R. Dannenberg "A Statistical View on Piano Rolled Chords", in *Proc. 16th International Conference on Music Information Retrieval*, Malaga, October 2015.

PUBLICATIONS (continued)

- G. Xia**, R. Dannenberg “Duet Interaction: Learning Musicianship for Automatic Accompaniment”, in *Proc. 15th The International Conference on New Interfaces for Musical Expression*, Baton Rouge, June 2015.
- 2014 **G. Xia**, T. Huang, M. Yifei, R. Dannenberg, C. Faloutsos “MidiFind: Similarity Search and Popularity Mining in Large MIDI Databases”, in *Music Sound and Motion*, pp. 259 -276.
- R. Dannenberg, N. Gold, D. Liang, **G. Xia** “Methods and Prospects for Human Computer Performance of Popular Music”, in *Computer Music Journal*, 38 (2) (Summer2014).
- R. Dannenberg, N. Gold, D. Liang, **G. Xia** “Active Scores: Representation and Synchronization in Human-Computer Performance of Popular Music”, in *Computer Music Journal*, 38 (2)
- 2013 T. Huang, **G. Xia**, Y. Ma, R. Dannenberg, C. Faloutsos “MidiFind: Fast and Effective Similarity Searching in Large MIDI Databases”, in *Proc. 10th International Symposium on Computer Music and Multidisciplinary Research*, Marseille, October 2013.
- 2012 **G. Xia**, J. Tay, R. Dannenberg, M. Veloso “Autonomous Robot Dancing Driven by Beats and Emotions of Music”, in *Proc. 12th International Joint Conference on Autonomous Agents and Multi-Agent Systems*, Valencia, June 2012, pp. 205-212.
- 2011 **G. Xia**, D. Liang, R. Dannenberg, M. Harvilla “Segmentation, Clustering, and Display in a Personal Audio Database for Musicians”, in *Proc. 12th International Conference on Music Information Retrieval*, Miami, October 2011, pp.139-144.
- D. Liang, **G. Xia**, R. Dannenberg “A Framework for Coordination and Synchronization of Media”, in *Proc. 11th The International Conference on New Interfaces for Musical Expression*, Oslo, May 2011, pp. 167-172.

TEACHING EXPERIENCE

Instructor

- Spring 2020* **Class:** CSCI-SHU 360: Introduction to Machine Learning. NYU Shanghai
Co-instructor: Prof. Ross
- Spring 2018-2020* **Class:** CSCI-SHU 188: Introduction to Computer Music. NYU Shanghai
- Fall 2017-2018* **Class:** CSCI-SHU 101: Introduction to Computer Science. NYU Shanghai
- Winter 2017* **Class:** MUS102: Music, Mind, and Invention. Dartmouth College
Co-instructor: Prof. Casey

Guest Lecturer

- Spring 2017* **Computer Music and HCI**
Class: CS67/167: Introduction to Human-Computer Interaction
Instructor: Prof. Yang
- Fall 2014* **Open Sound Control and its Applications in Music Performance**
Class: 15104: Principles of Computing for Creative Practice, Carnegie Mellon University
Instructors: Prof. Dannenberg & Prof. Roberts
- Fall 2014* **The Design and Acoustic Properties of the Chinese Flute**
Class: 57337: Sound Recording, Carnegie Mellon University
Instructors: Prof. Schulz
- Fall 2013* **The History and Future of Computer-aided Music Performance**
Alumni Lecture Series of Information Science Department, Peking University
Host: Prof. Wang

Student Supervision

Current Students:

- April 2018 – present* Junyang Jiang, Master student in Machine Learning, CMU
Project: Hierarchical Music Representation Learning
- April 2018 – present* Ziyu Wang, Ph.D. student in Computer Science, NYU
Project: Interactive Music Arrangement via Deep Representation Learning
- April 2018 – present* Tianyao Chen, Ph.D. student in Computer Science, NYU
Project: A Type-oriented System for Music Structure Analysis and Music Theory Induction
- Previous Students:*
- Dec 2017 – May 2019* Maoran Xu, Undergraduate student in Applied Mathematics and Data Science, Fudan University
Project: Performance Control Transfer
Current position: Ph.D. student in Statistics, University of Florida
- May 2018 – Aug 2019* Ke Chen, Undergraduate student in Computer Science, Fudan University
Project: Conditional Deep Music Generation
Current position: Ph.D. student in Music, UCSD
- Nov 2017 – May 2018* Han Su, Undergraduate student in Computer Science & IMA, NYU Shanghai
Project: Bio-music Computing using Slime Mold
Current position: Master student in Comparative Media Study, MIT
- May 2018 – Aug 2019* Ruihan Yang, Undergraduate student in Computer Science, NYU Shanghai
Project: Representation Learning and Disentanglement for Symbolic Music
Current position: Ph.D. student in Computer Science, UCI
- Nov 2016 – Oct 2018* Shuqi Dai, Undergraduate student in Computer Science, Peking University
Project: Expressive Performance Model for Pipa
Current position: Ph.D. student in Computer Science, CMU
- Oct 2016 – Aug 2017* Qianwen Chen, Masters Student in Computer Science and Digital Arts, Dartmouth College
Project: Haptic Guidance for Flute Tutoring
Current position: Designer, Apple.
- Feb 2017 – Aug 2017* Rohan Arora, Masters Student in Computer Science, Dartmouth College
Project: Haptic Guidance for Flute Tutoring
Current position: Data & AI Products Builder, Microsoft.
- Aug 2014 – Aug 2016* Mutian Fu, Masters Student in Music Technology, CMU
Project: Piano Rolled Chord Timing Analysis for Collaborative Performance
Current position: Software Engineer, Bloomberg
- Summer 2015* Peiling Lu, Undergraduate Student in Communication Engineering, CMU
Project: Piano Pedal Timing Analysis for Collaborative Performance
Current position: Graduate student, CCRMA, Stanford University.

INDUSTRY EXPERIENCE

- Summer 2013* **Music Information Retrieval Research Intern**, The Echo Nest, Cambridge, MA
Project: Improvement on Echo Nest's Beat Tracking Algorithm
Mentor: Dr. Tristan Jehan
- Summer 2012* **Research Intern**, Gracenote Inc., Emeryville, CA
Project: Improvement on Gracenote's Music Fingerprint Algorithm
Mentor: Dr. Bob Coover

RESEARCH PROJECTS

NYU Shanghai & NYU Tandon, Shanghai & New York:

Sep 2017 – present

Project: Deep Music Representation Learning

- Designed EC²-VAE for representation disentanglement
- Designed Transformer-VAE for context-sensitive representation learning
- Designed hybrid models for deep music generation

Sep 2017 – present

Project: Multimodal Music Tutoring System

- Designed adaptive haptic feedback for flute tutoring
- Designed adaptive visual feedback for sight-playing tutoring

Dec 2017 – present

Project: Music Information Retrieval Using Deep Learning

- Built the state-of-the-art chord recognition system (first place in MIREX 2018)
- Built the state-of-the-art key recognition system (first place in MIREX 2019)

Sep 2017 – present

Project: Performance Control Transfer via Player Pianos

- Designed heuristics to build the mapping between performance control and sound
- Enabled different player pianos to “adjust” performance in different environments

Neukom Institute, Dartmouth College, USA:

Aug 2016 – Aug 2017

Project: New Instrumental Interface for Flute Tutoring

- Enabled machines to control human fingers in real time
- Designed new flute infrastructure with motors

Jul 2012 – May 2013

Project: Biocomputing for music composition

- Used live creature, slime mold, as a computational model
- Enabled slime mold to sing

Humanoid Piano Robotics Lab, TeoTronica Inc., Italy (*visiting collaboration*):

Sep 2015 – Aug 2016

Project: A Humanoid Piano Robot for Human-Computer Music Performance

- Enabled a piano robot to perform music interactively with a human musician
- Enabled a piano robot to compose an accompaniment for a given melody
- Designed algorithms generating natural facial expressions in music performance

Takanishi Lab, Humanoid Robotics Institute, Waseda University, Japan (*visiting collaboration*):

Sep 2015

Project: A Humanoid Saxophone Robot for Human-Computer Music Performance

- Enabled a saxophone robot to perform music interactively with a human musician
- Designed algorithms generating natural body gestures in music performance

Machine Learning Department, Carnegie Mellon University, USA:

Jul 2013 – Aug 2016

Project: Expressive Collaborative Music Performance via Machine Learning

- Enabled machines to learn musicianship from human rehearsals
- Enabled machines to expressively perform music with human musicians

Jul 2012 – May 2013

Project: A Fast and Effective MIDI Search Algorithm

- Aimed to search different performance MIDI files of the same composition
- The algorithm is both effective (0.95 F-score) and fast (searching 10K files in 0.1 second)
- Website: www.cmumidifind.com

Jul 2011 – Jul 2012

Project: Dancing Robots Driven by Music

- Automatic music beats and emotion analysis
- Enabled physically based robots to dance with music according to the beats and emotion

Aug 2010 – Jul 2011

Project: A Framework for Coordination and Synchronization of Media

- Designed the GUI for interactive performance based on scanned sheet music
- Automatic page turning driven by music performance

Aug 2010 – Aug 2011

Project: Rehearsal Audio Stream Segmentation and Clustering

- Designed novel and robust representation of noisy music signals
- Achieved state of the art accuracy (0.98) by ADAboost-HMM classification frameworks

National Laboratory on Machine Perception, Peking University, China:

Aug 2009 – May 2010 **Project: Automated Music Structure Analysis on a Symbolic Representation**

- Designed a hierarchical music structure analysis framework
- Detected music phrases using Conditional Random Fields
- Detected music sections using NTRPs (Nontrivial Repeating Patterns)

Oct 2009 – Jun 2010 **Project: Automated Composition Based on Music Structure Analysis**

- Designed a structural model of music styles
- Automated composition by recombining and modifying extracted phrases

MUSIC RELATED EXPERIENCE

Dizi and Xiao (Chinese Flute and Vertical Flute):

Mar 2013 – May 2016 **Soloist**, Pitt Carpathian Ensemble, Pittsburgh, PA

May 2010 **Solo Concert**, President Hall of Peking University, Beijing

Dec 2006 – Jun 2010 **Leading Soloist**, Chinese Music Institute of Peking University, Beijing

Musical and Choral:

Oct 2016 – Jun 2017 **Glee Club Member**, Dartmouth College

Apr 2014 **Role of Audrey 2**, in the musical “Little Shop of Horrors”, Carnegie Mellon University

Feb 2013 **Role of Phantom**, in the single “Music of the Night”, Carnegie Mellon University

Oct 2012 **Role of Teen Angel**, in the musical “Grease”, Carnegie Mellon University

Oct 2011 **Role of Simon Zealotes**, in the musical “Jesus Christ Superstar”, Carnegie Mellon University

Conducting:

Apr 2007 – Jun 2009 **Assistant Conductor**, Chamber Orchestra of Chinese Music Institute of Peking University, Beijing

INVITED TALKS

- Nov 2019* Better Music Representation Learning via Inductive Bias: Mind vs. Machine, *CDS, NYU*
- Nov 2019* Deep Music Analogy via Representation Disentanglement, *ISMIR*
- Oct 2019* Hierarchical Music Representation Learning Using Transformer VAE, *MARL, New York University*
- Oct 2019* Better Music Representation Learning via Inductive Bias: Mind vs. Machine, *Carnegie Mellon*
- Oct 2019* Better Music Representation Learning via Inductive Bias, *Dartmouth College*
- Sep 2019* Interpretable Music Generation via Representation Disentanglement, *CILVR, New York University*
- Aug 2019* Music Understanding and Generation via AI, *CCF Annual Conference*
- Apr 2019* Music Style Transfer via Analogy-Making, *International Forum on Statistics, SUFE*
- Apr 2019* Music Intelligence: Towards Creative AI Systems, *Zaojiu Talk*
- Jan 2019* Explicit Structure Encoding for Music Generation, *MMRP*
- Oct 2018* Towards More Expressive Artificial Music Intelligence. *Carnegie Mellon University*
- July 2018* A Framework for Automated Popular Music Arrangement, *CSMT*
- May 2018* Music Style Transfer: Some Important Concepts, *MuMe workshop, ICCC*
- Oct 2017* Computer Music as an Interdisciplinary Field, *Shanghai Science and Technology Museum*
- Jun 2017* Turing Test for the Creative Arts, *MuMe workshop, ICCC*
- May 2017* Learning Improvisation for Automatic Accompaniment, *NIME*
- Feb 2017* Expressive Human-computer Music Interaction, *Georgia Tech*
- Aug 2016* Robot Embodiment for Automatic Accompaniment, *SMC*
- Oct 2015* Spectral Learning for Expressive Interactive Ensemble Music Performance, *ISMIR*
- Sep 2015* Modeling Piano Rolled Chords and Pedal Timing, *Carnegie Mellon University*
- Jun 2015* Learning Musicianship for Automatic Accompaniment, *NIME*
- Sep 2014* Improving Automatic Accompaniment by Machine Learning, *University of California, San Diego*
- Dec 2013* A History of Automatic Accompaniment and its Future Directions, *Peking University*
- Oct 2012* A Vision of Human-Robot Music Performance, *Carnegie Mellon University*
- Jun 2012* Autonomous Dancing Robots Driven by Music, *AAMAS*
- Dec 2011* Smart Music Displays and Autonomous Dancing Robots, *Peking University*
- Jun 2011* Coordination and Synchronization of Media for Human-Computer Music Performance, *NIME*

HONORS and AWARDS

- Co-chair for MuMe workshop, 2019
- Chair for NIME Conference, *Expected 2021*
- Music Chair for ISMIR Conference, 2017
- Turing Test Competition Organizer, Dartmouth College, 2017
- 3-Minute Thesis Presentation Competition Finalist, Carnegie Mellon University, 2015
- Graduate Student Assembly Representative, Carnegie Mellon University, 2013-2014
- Team Leader for OurCS Computer Music Group, Carnegie Mellon University, 2013, 2015
- Outstanding Society Leadership Award (2 out of 260), Peking University, 2009
- President and Music Director of Chinese Music Institute, Peking University, 2007-2009
- Volunteer for Olympic Games, Beijing, 2008
- Selected Student Delegation out of 2000+ candidates for BESETOHA Forum, Tokyo, 2007.
- 1st place out of 3000+ in the National Musical Instrument Proficiency Evaluation, China, 2006

REFERENCES

Available upon request.