Harmonious Research Collaborations in Computational Musicology

Cory McKay
Marianopolis College and CIRMMT, Canada

María Elena Cuenca Rodríguez
Universidad Autónoma de Madrid, Spain

June 29, 2024. Digital Technologies Applied to Music Research Conference Lisbon, Portugal

Introduction

- Collaborations between musicologists and data scientists hold enormous potential for furthering musical knowledge
- However, there can be significant challenges with these kinds of interdisciplinary collaborations

- This talk will present some of what we have learned from working together successfully on a number of projects over the past several years as a musicologist (María Elena) and as a data scientist (Cory)

Big questions to think about

- What existing needs of music scholars can be addressed by computational approaches?
- What new, different opportunities for music scholarship do computational approaches present?
- What challenges and pitfalls can one encounter with computational approaches?
  - And how can we address them?
- How can we encourage and facilitate effective discussions and collaborations between domain experts?

Cory’s MIR to musicology origin story (1/2)

- I began my research career in music information retrieval (MIR) in 2004
- Although I eventually began to focus on multimodal approaches, my earliest work was purely on symbolic music
- I wanted to move beyond commercially-oriented or MIR-oriented priorities of the time, such trying to build systems with the highest possible classification accuracy, to using MIR techniques to actually reveal new insights about music
- I had no musicological experience at the time beyond a few courses I had taken as an undergrad, so I began to look for musicologists to collaborate with (e.g., attending meetings of the AMS)
  - Unfortunately, many (but happily not all!) of the musicologists I met at the time seemed to think that the statistical or machine learning approaches I proposed were somewhere between useless and heretical
  - Similarly, many (but happily not all!) of my MIR colleagues at the time seemed to think that musicological approaches were dated or scientifically invalid
- Rebuffed, I regretfully retreated back to the walled garden of MIR-oriented conferences

Cory’s MIR to musicology origin story (2/2)

- Happily, both fields have become more open with time, and in the mid-2010’s I was approached by Ichiro Fujinaga and Julie Cumming to join a sequence of grant applications combining MIR and musicology, in the particular context of early music
  - Delighted, I did, and delightfully we were awarded the grants
- I began attending musicology conferences again, especially MedRen, and was overjoyed to find musicologists that were open and interested in experimenting with computational approaches
  - And humbled to confirm how much I did not know and needed to learn about early music
- Since then, I have had the pleasure of carrying out and publishing work with a number of different musicologists, especially María Elena

In a reflective mood

- Why were my first attempts to find musicological collaborators unsuccessful?
- Why do MIR researchers, and data scientist more generally, still relatively rarely work together with musicologists, even now?
- What can be done to foster more interdisciplinary collaborations in the future?

The crux of the issue?

- Musicology and MIR / data science have very different:
  - Vocabularies
    - e.g., “symbolic music” or “multimodal”
  - Ways of thinking about and conceptualizing music
  - Research priorities
  - Methodologies
  - Styles of presenting and disseminating research
- These kinds of differences can make work seem inaccessible, arcane or even alienating to those not in a given field
- Demystifying these kinds of elements is an essential step towards fostering effective collaboration
  - With that in mind, I will briefly introduce the feature-based approach that is at the core of the work María Elena and I have done together

What is a “feature?”

- Information that measures a characteristic of a segment of music in a simple, consistent and precisely-defined way
- Represented using numbers
  - Can be a single value, or can be a set of related values (e.g., a vector of histogram values)
- Provides a summary description of the characteristic being measured
  - A macro rather than local view
- Usually extracted from pieces in their entirety or from large sections (e.g., mass movements)
  - But can also be extracted from smaller segments of music

Example: A simple feature

- Range: Difference in semitones between the lowest and highest pitches present

- Value of this feature for this music: 7
  - G - C = 7 semitones

Josquin's *Ave Maria . . . Virgo serena*

- Range: 34 (semitones)
- Repeated notes: 0.181 (18.1%)
- Vertical perfect 4ths: 0.070 (7.0%)
- Rhythmic variability: 0.032
- Parallel motion: 0.039 (3.9%)

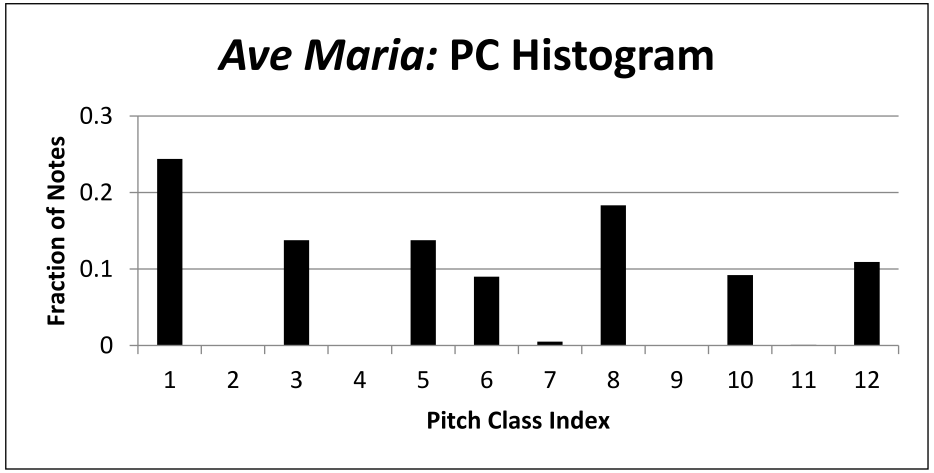

Ockeghem’s Missa *Mi-mi* (Kyrie)

- Range: 26 (semitones)
- Repeated notes: 0.084 (8.4%)
- Vertical perfect 4ths: 0.109 (10.9%)
- Rhythmic variability: 0.042
- Parallel motion: 0.076 (7.6%)

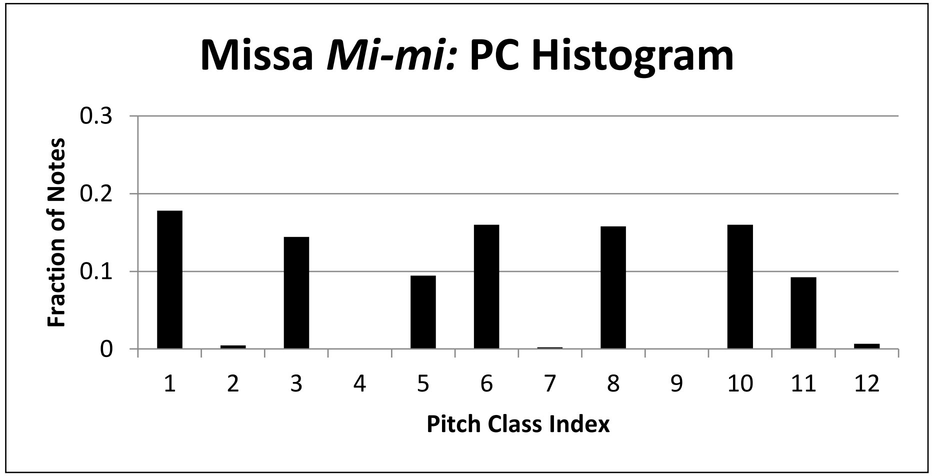

Feature value comparison

Feature	*Ave Maria*	Missa *Mi-mi*
Range	34	26
Repeated notes	0.181	0.084
Vertical perfect 4ths	0.070	0.109
Rhythmic variability	0.032	0.042
Parallel motion	0.039	0.076

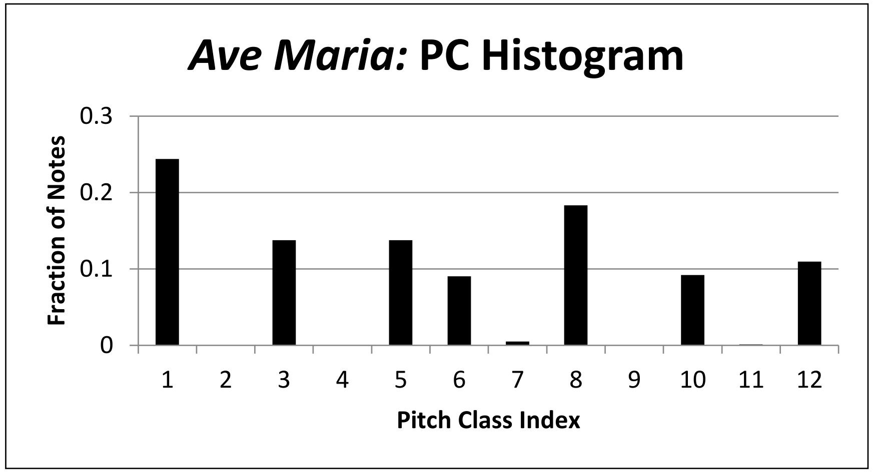

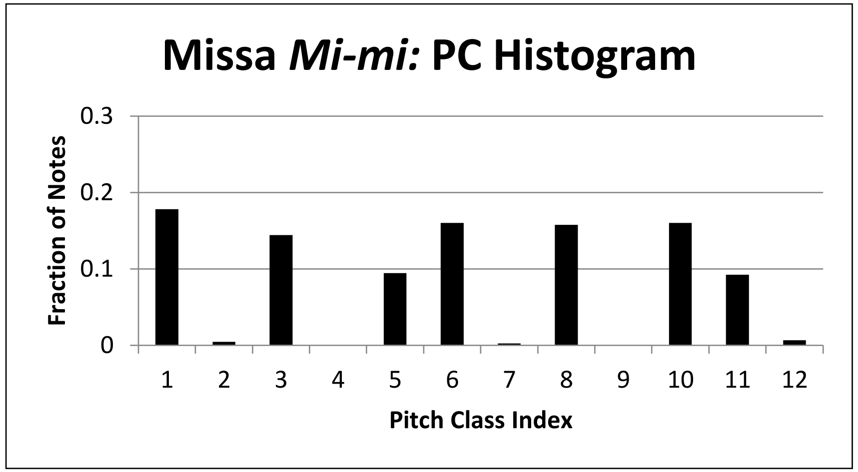

Comparing features

- Comparing pieces like this in terms of features can be revealing
  - Especially when that comparison involves hundreds or thousands of features, not just six
- Things get even more interesting when comparisons are made between hundreds or thousands of pieces, not just two
  - Especially when the music is divided into groups of interest, which can then be collectively contrasted with one another
    - e.g. comparing the styles of composers, genres, regions, time periods, etc.

How might one calculate features?

- The jSymbolic research software (McKay et al. 2018) can be used to automatically extract features from symbolic digital scores
    - Open source
    - Applicable to diverse musics
- Version 2.2 extracts 246 unique features
    - 1497 separate feature values, since many features a multi-dimensional (e.g. histogram vectors)
- The upcoming Version 3 extracts 533 unique features
    - 2040 feature values, including n-gram features

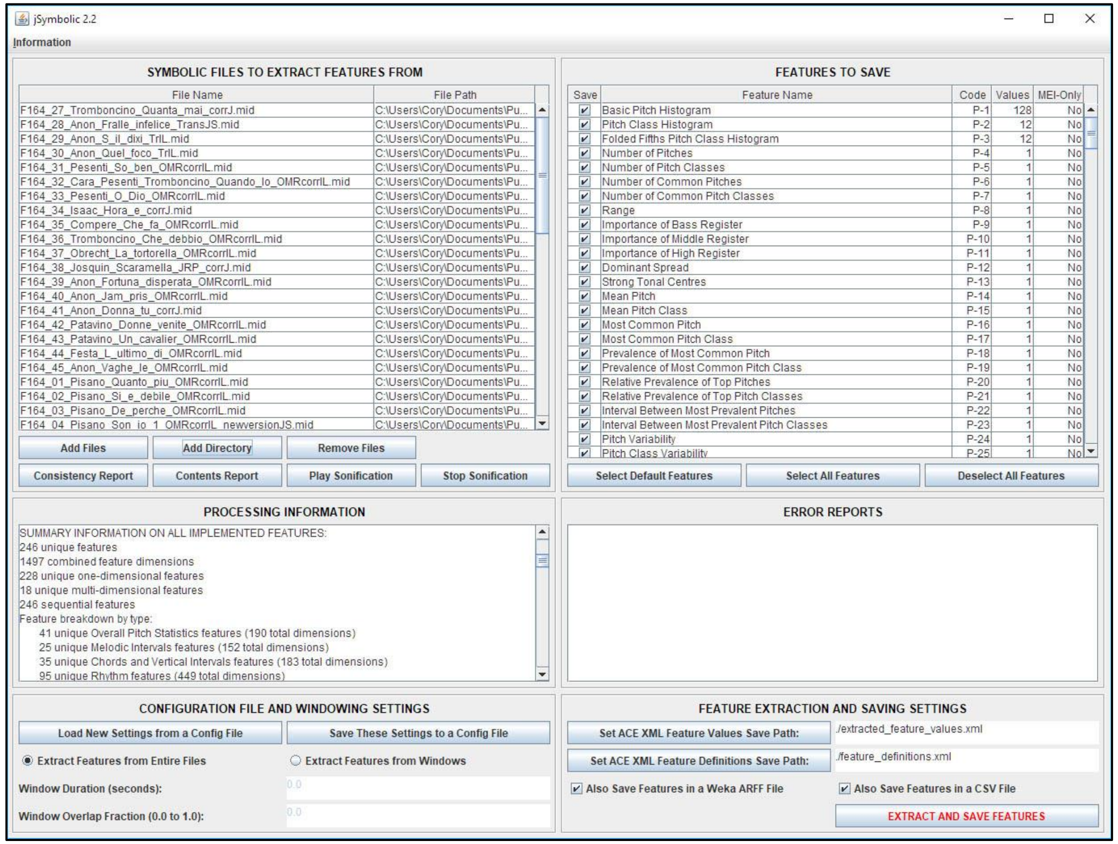

jSymbolic 2.2's feature types

- Pitch statistics
  - e.g. Range
- Melody / horizontal intervals
  - e.g. Most Common Melodic Interval
- Chords / vertical intervals
  - e.g. Vertical Minor Third Prevalence
- Texture
  - e.g. Parallel Motion
- Rhythm
  - e.g. Note Density per Quarter Note
- Instrumentation
  - e.g. Note Prevalence of Unpitched Instruments
- Dynamics
  - e.g. Variation of Dynamics

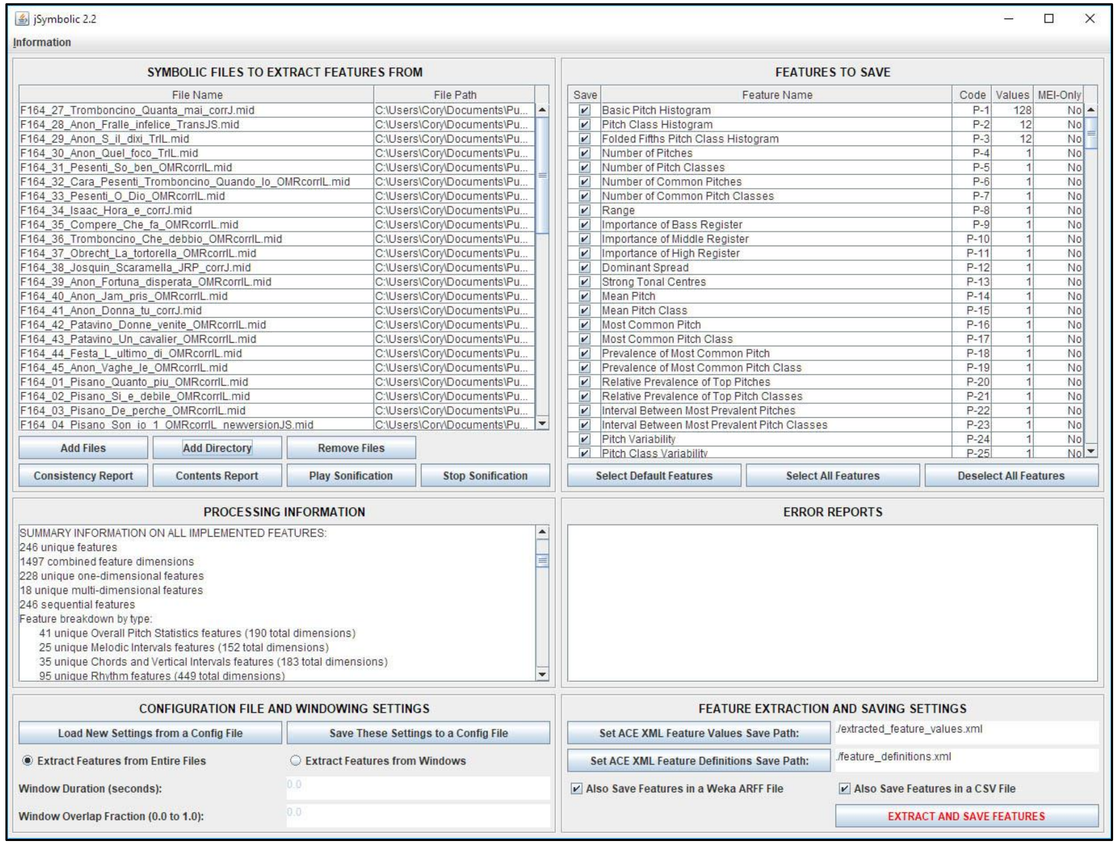

Context: Overview by María Elena

Big challenge 1

- How can we make musicological sense of data yielded by computational approaches, particularly with respect to feature data?
  - The data can sometimes be unintuitive and difficult to interpret musically
  - The sheer amount of information can be overwhelming

Ways of comparing feature values

- Manually:
  - Text editors
  - Spreadsheets
- With automatic assistance:
  - Statistical analysis software
    - e.g. SPSS, SAS, etc.
  - Machine learning and data mining software
    - e.g. Weka, Orange, etc.
    - Supervised or unsupervised
- Many of these tools can produce helpful visualizations
  - But not in music-specific ways

Future priorities

- There is an important need for new and better ways of visualizing musical data
  - Developed in consultation with musicologists!
- Fortunately, some good initial work has already been done in this direction, including:
  - CRIM heat maps (Freedman et al.)
    - As mentioned by María Elena above
  - SymPlot (Muñoz-Lago et al.)
  - Sonic Visualizer (Cannam et al.)
- Ultimately, we need more ways of presenting musical data that are flexible, easily human-interpretable and musically salient

Big challenge 2

- We must prepare and present music and musical data to computers in ways that will avoid incorrect or misleading results
  - Problems in a score or inconsistencies across scores that may be obvious and easily adjusted for in manual expert analysis can be easily missed in automatic analysis, which can deeply compromise results

Avoiding encoding bias (1/2)

- Some potential problems to be careful of in early music, for example:
  - Inconsistent encoding of accidentals corresponding to *musica ficta*
  - Choice of different rhythmic note values to denote the beat
  - Differing metrical interpretations of mensuration signs
  - Transposition to different keys
- How to avoid biased results:
  - Ideally, use digital music files that were all consistently generated using the same methodology
    - All editorial decisions (e.g. *musica ficta*) should be applied consistently and should be documented
  - If this is not possible, then exclude all features that could be sensitive to particular biases that could be present in the data

Avoiding encoding bias (2/2)

- Related publications:
  - Cumming, J., C. McKay, J. Stuchbery, and I. Fujinaga. 2018. Methodologies for creating symbolic corpora of Western music before 1600. *Proceedings of the International Society for Music Information Retrieval Conference*. 491–8.
  - Nápoles López, N., G. Vigliensoni, and I. Fujinaga. 2018. Encoding matters. Presented at the *International Conference on Digital Libraries for Musicology*.

Big challenge 3

- What constitutes an important research question?
    - Musicologists and data scientists often start with quite different goals and priorities

What constitutes an important research question?

- Efforts must be made by both types of researchers to learn about and respect one another's priorities, and to think about how they can be merged
    - e.g., data scientists must see music as intrinsically meaningful and worthy of study, not just as a toy domain to test more general technical approaches
    - e.g., musicologists must see data scientists as researchers with their own priorities, not just as technicians or tool makers whose role is limited to implementing what musicologists want done
- Fortunately, with openness, thought and communication, goals from both domains can certainly complement one another
    - Or, even better, result in important new goals and priorities that neither party might have otherwise thought to consider

A few other important challenges

- Reconciling different research cultures
  - e.g., solo papers vs. group papers
  - e.g., emphasizing personal expertise and reflection vs. “objective” data and open methodologies
  - e.g., speed of progression from research conceptualization to publication
- Usability and documentation of tools
  - GUIs make tools more accessible, but other interfaces (command line or API) can have advantages too, so ideally multiple parallel interfaces should be available
  - Good manuals and tutorials are essential in encouraging the adoption of tools
- Effective interdisciplinary research dissemination
  - e.g., Borsan et al. (2023) found that MIR research is very rarely cited in the musicological literature
  - We need more conferences like this one!

Conclusion

- Openness and commitment to good communication are essential

- It is especially important that researchers from different disciplines be open to seeing the work they are doing and the methodologies they are employing from perspectives outside their own disciplines
  - This can mean having assumptions previously taken for granted probed critically, something that should be welcomed rather than avoided
- While computational approaches can and should be used to help answer existing musicological questions, they are arguably at their most valuable when they open the door to new questions and new types of answers that nobody had previously thought to consider

Thanks for your attention

cory.mckay@mail.mcgill.ca

elenacrod@gmail.com

MARIANOPOLIS COLLEGE

CIRMMT Centre for Interdisciplinary Research in Music Media and Technology

UAM Universidad Autónoma de Madrid

Fonds de recherche Société et culture Québec

SSHRC CRSH Canada