

On Every Note a *Griff*: Looking for a Useful Representation of Basso Continuo Performance Style

Adam Štefunko¹, Carlos Eduardo Cancino-Chacón², Jan Hajič jr.¹

¹ Charles University, Faculty of Mathematics and Physics, Czech Republic

² Institute of Computational Perception, Johannes Kepler University, Austria

<stefunko@ufal.mff.cuni.cz>

Introduction

Basso continuo, the baroque improvisatory harmonic accompaniment style, requires the player to create a multiple-part improvisation based on a bass line from a musical score. This process is called *basso continuo realization*. *Continuo* practice has been resurrected in the 20th century by the Historically-Informed Performance movement ([Christensen, 2002](#)), and many harpsichord and organ players now play *basso continuo* realizations regularly. However, despite its position at the attractive intersection between a well-understood musical style and orally transmitted improvisatory practice, little work has analyzed contemporary *continuo* realizations.

The first sample of *basso continuo* practice has been recorded in *The Aligned Continuo Realization Dataset* (ACoRD),¹ and a performance-to-score alignment system was developed to match the performances in the dataset with the *basso continuo* scores they are based on ([Štefunko et al., 2025](#)). *Basso continuo* alignment, albeit imperfect in some edge cases ([Chiruthapudi et al., 2025](#)), establishes resources to start analyzing *continuo* realizations. The next missing step is finding suitable features that can be extracted from the performances and used for analysis.

We propose a feature representation called *griff*: configuration of tones played for a given score note (see Figure 1). The word *griff* stems from the *basso continuo* treatise written by [Muffat \(1699\)](#), but the idea of an idiomatic chord note arrangement is also present in other treatises ([Dandrieu, 1719](#); [Geminiani, 1756](#)).

¹ <http://hdl.handle.net/11234/1-5963>




Figure 1. *Griffs* are built using performances aligned to scores, and each score note forms its own *griff* from performance notes aligned to it (a). MIDI pitches of performance notes are grouped to vectors based on the order and proximity of their onset times (b). A *griff* is a sequence of these vectors. *Griffs* converted to chromatic intervals from the score note pitch (c) are encoded as strings (d). A note can participate in multiple consecutive griffs, e.g., the Bb is contained in both the 2nd (yellow) and 3rd (purple) *griff*, or the high E in the 3rd and 4th (red) *griff*).

Griffs become the vocabulary for analyzing *continuo* realization styles. At the same time, *griffs* are a practically lossless encoding of pitch content of a realization. We hypothesise that whereas individual MIDI pitch and interval content of performances might be too coarse to study individuality of different players' performance styles, *griffs* are features with an appropriate resolution, as well as an encoding congruent with historical ways of thinking about *continuo*. In this paper, we define the *griff* representation, describe how *griffs* are extracted from aligned *continuo* performances, and use the *griff* representation to investigate whether we can stylistically distinguish between different players' performances.

1 Building the *Griffs*

Griff extraction begins with *basso continuo* alignment mapping each score note to performance notes played to it. We propose two *griff* representations:

- *Ordered*, which take into account the order in which individual performance notes appear, and
- *Pooled*, in which all distinct pitches of performance notes aligned with the same score note form one vector.

An *ordered* griff is a matrix of performance notes aligned to a score note in which the vertical axis represents pitches encoded either as MIDI numbers or as intervals from a given baseline, and the horizontal axis portraying the order of note onsets. It is very often desirable to join notes whose offsets differ only slightly in the same vector using a customizable quantization threshold parameter. All distinct pitches whose onsets appear within a window set by this parameter form a vector, pitches played after this window form the next vector, and a *griff* is a sequence of such vectors. To distinguish between block chords and purposeful arpeggiation, we use 35ms windows, a value based on previous research ([Flossmann et al., 2009](#); [Nakamura et al., 2014](#)).

In order to work with an encoding invariant to musical transposition, we convert MIDI pitches to chromatic intervals from the score note. *Griffs* are encoded as strings by separating intervals belonging to the same vector by an underscore, and neighboring vectors by a vertical line (see Figure 1).

Griffs are designed to enable narrowing down the possible harmonic space of *basso continuo* realizations and bring out structured tonal content. To this end, *griffs* discard the absolute temporal relations between notes because rhythmic differences expressed by absolute time tend to be variations of structurally equivalent patterns, so exact timing would obscure the equivalence of these patterns. Similar representations were used by [Laaksonen and Lemström \(2025\)](#) for music pattern matching and by [Peter \(2023\)](#) in a DTW-based alignment system. Although *griffs* do not directly contain absolute timings, they can be retrieved from the original MIDI files if they are needed.

2 Analysis of the ACoRD Dataset Using *Griffs*

In order to see whether the *griff* representation is able to provide a suitable granularity to capture potential differences in different players' *basso continuo* realizations, we analyze the ACoRD dataset from the *griff* content perspective.

The ACoRD dataset consists of 175 MIDI recordings, totalling 6 hours of runtime spread across 66,967 notes. This dataset enables us to see how different players realized the same musical scores because the recordings were done by 7 different players, and each of them realized 5 *basso continuo* scores for 5 times. *Griffs* can be built thanks to performance-to-score alignments automatically generated by an alignment system developed earlier. This system leverages the state-of-the-art symbolic alignment methods, and combines them with preprocessing techniques and a time-concurrency dependent greedy algorithm ([Štefanko et al., 2025](#)). The following statistics and experiments are based on usage counts of different *griff* types. In fact, most used *griff* types are performance bass notes without harmony (724

times, 3.9%) and score notes not aligned to any performance note (534 times, 2.9%). However, since we are interested in analyzing the harmonic content of the dataset, we filter these two *griff* types out.

2.1 Griff Statistics across the Dataset

Since ground-truth alignment data are available only for a limited number of performances, we use the best performing alignment method *DualDTWNoteMatcher* (Peter, 2023) with *trimming* preprocessing (Štefunko et al., 2025). We obtain 17,152 individual *griffs* from the dataset after filtering. The basic statistics about the *griffs* is shown in Table 1. Distributions of both *griff* representations illustrated in Figure 2 follow Zipf’s law. Figure 3 shows examples of different chord patterns present in *basso continuo* playing encoded as *griffs*.

Griff Representation	Ordered	Pooled
Total Griff Types	7041	2817
Average Griff Occurrence Count	2.44	6.09

Table 1. Basic statistics of both *griff* representations. *Griff* occurrence count averaged over 17,152 *griffs*.




Figure 2. Distribution of different *griff* types in both representations.

2.2 Experiments with Performance Profiles

We investigate how individuality of players is reflected in the distribution of different *griff* types across performances. We use performance *griff* profiles, showing how often was each *griff* type played in a performance. These profiles can be further aggregated across multiple performances.




Figure 3. Three most occurring ordered *griff* types without and with vertical lines belonging to the three most frequent chord types. The referential score note and a root position of the chords are shown. The figure contains string encoding of each *griff*, frequency of occurrence and respective rank in the whole dataset.

We measure how similar are player’s realizations of a score in comparison to other player’s realizations of the score in the space of *griff* profiles. We measure cross-entropy² of normalized *griff* profiles on pairs of all 5 performances of a score by one player. We similarly measure cross-entropy on pairs of performances of two different players. Cross-entropies are computed for both *griff* representations, and for frequencies of each chromatic interval to see how adding the structure to the intervallic content can help in measuring differences in performance styles.

2.3 Experiment results

In order to study individuality of players, we need to find a level of representation that enables us to see variability across different players’ playing styles. Figure 4 presents similarity matrices of mean pairwise cross-entropy values between performances of players for both *griff* profile representations and, as a control, frequency profiles for intervals of individual realization notes from the score note (being a much less detailed representation).

²Any other similarity metric could be used, but since *griff* profiles normalized over score length are categorical distributions, the choice of cross-entropy follows from information theory. It tells us how much extra information, in terms of bits, we need if we approximate the *griff* distribution of a referential performance with the *griff* distribution of a different performance. KL-divergence would be the most preferential choice, but we can use cross-entropy directly because the distribution of the referential performance stays constant.




Figure 4. Similarity matrices of different players' mean cross-entropies in three different representations for each score.




Figure 5. Effect sizes using Cohen's d shows the magnitude of differences between mean cross-entropy of pairs of same player performances and different player performances. The figure displays effect sizes separately for each player playing each score show and all three representations.

Even though all three representations showed to be significant under the *permutation test* for most player and score pairs, effect size measured as Cohen's d for each player and each score shows that *griffs* are superior in terms of capturing consistency in different players' styles (Figure 5). This is expected: harmonic constraints of *continuo* realization heavily restrict the space of possible interval profiles for each score, so this representation is too coarse. In the *griff* profile spaces, however, performances of the same player

are more similar to each other than to performances of other players. If the *griff* space was too fine-grained to be analytically useful, we would expect each performance pair to be approximately the same distance from each other. This is far from a complete analysis, and should not be taken as confirmation that there are individual *continuo* realization styles yet, but we do take it as an indication that this “historically informed feature space” of *griffs* is worth defining and worthy of a more robust empirical validation (such as with clustering and classification experiments, which are beyond the current scope).

3 Conclusion

Griffs seem to be a historically-informed and empirically promising way of representing aligned *basso continuo* performances. They retain distinctions between individual performers’ style, and thus provide meaningful tokens for analysis of *continuo* as a living improvisatory tradition.

Acknowledgements

This research was supported by the project “**Human-centred AI for a Sustainable and Adaptive Society**” (reg. no.: CZ.02.01.01/00/23_025/0008691), co-funded by the European Union, by The Charles University Grant Agency (project no. 58126), and by the SVV project no. 260 821. The computing infrastructure was provided by the LINDAT/CLARIAH-CZ Research Infrastructure (<https://lindat.cz>) supported by the Ministry of Education, Youth and Sports of the Czech Republic (project no. LM2023062).

This work was also supported by the Austrian Science Fund (FWF) under grants PAT 8820923 (“**Rach3: A Computational Approach to Study Piano Rehearsal**”).

Statement on the use of generative AI

The code for running the experiments and creating figures in this paper was written with the assistance of GitHub Copilot.

References

- Chiruthapudi, S., Štefunko, A., Hajič jr., J., & Cancino-Chacón, C. E. (2025). Challenges in Basso Continuo Performance-to-Score Alignment. *Proceedings of the 17th International Symposium on Computer Music Multidisciplinary Research*, 924–936. <https://doi.org/10.5281/zenodo.1750251>
- Christensen, J. B. (2002). *18th Century Continuo Playing: A Historical Guide to the Basics*. Bärenreiter. <https://www.baerenreiter.com/en/product/BA08177>
- Dandrieu, J. F. (1719). *Principes de l'accompagnement du clavecin*.

- Flossmann, S., Grachten, M., & Widmer, G. (2009). Expressive performance rendering: Introducing performance context. *Proceedings of the SMC*, 155-160.
- Geminiani, F. (1756). *The art of accompaniment or A new and well digested method to learn to perform the thorough bass on the harpsichord with Propriety and Elegance*.
- Laaksonen, A., & Lemström, K. (2025). Parallelising the Polyphony: High-Performance Algorithms for Music Pattern Matching. *Proceedings of the 17th International Symposium on Computer Music Multidisciplinary Research*, 760–771. <https://doi.org/10.5281/zenodo.17496797>
- Muffat, G. (1699). *Regulae Concentuum Partiturae*.
- Nakamura, E., Nakamura, T., Saito, Y., Ono, N., & Sagayama, S. (2014). Outer-Product Hidden Markov Model and Polyphonic MIDI Score Following. *Journal of New Music Research*, 43(2), 183–201. <https://doi.org/10.1080/09298215.2014.884145>
- Silvan David Peter. (2023). Online Symbolic Music Alignment With Offline Reinforcement Learning. *Proceedings of the 24th International Society for Music Information Retrieval Conference*, 634–641. <https://doi.org/10.5281/zenodo.10265367>
- Štefunko, A., Chiruthapudi, S., Hajič jr., J., & Cancino-Chacón, C. E. (2025, September 10). Basso Continuo Goes Digital: Collecting and Aligning a Symbolic Dataset of Continuo Performance. *The AI Music Creativity Conference (AIMC)*, Brussels, Belgium. <https://doi.org/10.5281/zenodo.16946799>