Submitted exclusively to the *Journal of Mathematics and Music*
Last compiled on May 24, 2021

Some Observations on Autocorrelated Patterns within Computational Meter Identification

Christopher Wm. White [a*]

[a] *Department of Music and Dance, University of Massachusetts Amherst, USA*

()

The computational approach of autocorrelation relies on recurrent patterns within a musical signal to identify and analyse the meter of musical passages. This paper suggests that the autocorrelation process can act as a computational proxy for the act of period extraction, a crucial aspect of the cognition of musical meter, by identifying periodicities with which similar events tend to occur within a musical signal. Three analytical vignettes highlight three aspects of the identified patterns:1) that the similarities between manifestations of the same patterns are often inexact, 2) that these patterns have ambiguous boundaries, and 3) that many more patterns exist on the musical surface than contribute to the passage's notated/felt meter, each of which overlap with observations from music theory and behavioral research. An Online Supplement at chriswmwhite.com/autocorrelation contains accompanying data.

Keywords: meter; rhythm; harmony; computation; music theory

2010 Mathematics Subject Classification: 00A65; 62P15; 91E10; 94A17
2012 Computing Classification Scheme: Applied computing Sound and music computing

1. Introduction

Meter is a phenomenon of patterns. In general, theorists imagine meter as arising from a series of consistently paced accents (Lerdahl and Jackendoff 1983; Kozak 2020; Krebs 1999; Repp 1998), as involving a listener who expects that pacing to continue into the future (Hasty 1997; London 2004), and as grouping adjacent pulses to form a hierarchy of stronger and weaker pulses (Cooper and Meyer 1960; Lerdahl and Jackendoff 1983), with quicker pulses evenly dividing broader pulses by two or three, creating consistent duple or triple relationships between levels (Cohn 2001, 2020). London (2004) suggests that the process of identifying a meter involves two complementary methods: 1) period extraction, and 2) template matching. The former – and more basic – strategy identifies patterns of recurrent periodicities present in some musical signal; the latter involves an approach that categorizes and organizes events into some a priori understanding of metric hierarchies and relationships.

Computational approaches to meter can broadly be seen to fall roughly along London's divide. The first camp includes those techniques that identify periodic patterns within some musical timeline, and then equates the patterns that emerge from those patterns with metrical pulses. This general approach includes such techniques as Discrete Fourier

*Corresponding author. Email: cwmwhite@umass.edu

Transforms (Amiot 2016), wavelet analyses (Velarde, Meredith, and Weyde 2016), and resonant neural networks (Large, Herrera, and Velasco 2015). However, perhaps the most straightforward and oldest of these approaches is *autocorrelation*, a technique that identifies the periods at which similar events appear in some musical timeline (Boone 2000; Brown 1993; Eck 2006; Eerola and Toiviainen 2004; Palmer and Krumhansl 1990; Volk 2008). While the basic method has been shown to identify the notated time signature of music in various styles and genres with a reasonably high level of precision (Gouyon et al. 2006; White 2019a), researchers have supplemented this technique with such additions as Expectation Maximization (de Haas and Volk 2016) and Shannon entropy (Eck and Casagrande 2005) to increase the process's utility and precision.

Much recent computational research can be seen as modeling the processes described by a template-matching, approach (at least roughly). In many such approaches, some training session or encoded expert knowledge provides a model with expectations that allow it to categorize events into strong and weak beats, and to arrange those event into a metric hierarchy. Such approaches include Context Free Grammars (McLeod and Steedman 2017; Rohrmeier 2020), Hidden Markov Models (Khadkevich et al. 2012; Papadopoulos and Peeters 2011), and deep learning techniques such as Support Vector Machines (Durand, David, and Richard 2014), Recurrent Neural Networks (Durand et al. 2015; Böck, Krebs, and Widmer 2016), and Temporal Convolutional Networks (Böck and Davies 2020). These approaches have proven to perform quite well against a ground truth, and represent the current state of the art in meter detection within the music information retrieval literature.

However, while computational music research is often driven by the desire for an optimally efficacious model – that is, the best meter-finding model is the one whose output conforms to some expectation of desired results – such research may also be motivated by investigating the contours of the musical concept being modeled. It is with this latter purpose in mind that this paper approaches computational models of meter, and which motivates its study of a topic that, on the one hand constitutes the more elementary side of meter finding, while on the other hand represents a primary foundation of the theoretical concept of meter: period extraction. This study investigates some basic computational tactics underpinning period extraction, specifically analyzing the periodic patterns identified therein. While the specific mathematics and engineering behind the various approaches to computational period extraction substantially differ, the autocorrelation approach can be seen as representing the basic framework of many of these applications, as noted in Eck and Casagrande (2005), and Kim and Large (2017). This paper therefore adopts an autocorrelation approach to identify isochronous patterns in order to observe these patterns' attributes. Analyzing the musical characteristics of these patterns can then yield insights into the roles that period extraction and pattern finding potentially play within theoretical and cognitive concepts of meter. The broader goal of this study, then, is to provide an initial investigation into some intersections between computational, cognitive, and theoretical approaches to isochronous patterns within models of musical meter.

This paper will first offer a brief overview of the autocorrelation process, after which I will present three illustrative analytical vignettes using autocorrelation. I will highlight three particular aspects of the patterns identified in these analyses: 1) that the process often identifies patterns that are not particularly similar, but more similar than other options (or, *fuzzy patterns*), 2) that these patterns have ambiguous boundaries (or, the patterns are *unbounded*); and 3) that patterns exist on the musical surface in excess of those suggested by the notated/felt meter (or, *excessive patterns*). I then connect these aspects of computational periodic pattern identification to some broader theoretical a)

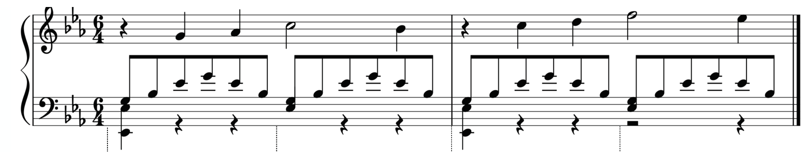

b)

Onsets Prime Vector	3	1	2	1	2	1	3	1	1	1	2	1	3	1	2	1	2	1	3	1	1	1	2	1
Onsets Rotation 1	1	3	1	2	1	2	1	3	1	1	1	2	1	3	1	2	1	2	1	3	1	1	1	2
Onsets Rotation 2	2	1	3	1	2	1	2	1	3	1	1	1	2	1	3	1	2	1	2	1	3	1	1	1
Onsets Rotation 12	3	1	2	1	2	1	3	1	1	1	2	1	3	1	2	1	2	1	3	1	1	1	2	1
Duration (in 8ths) Prime Vector	2	1	2	1	2	1	4	1	1	1	2	1	2	1	2	1	2	1	4	1	1	1	2	1
Mean Pitch (in MIDI) Prime Vector	48	58	65	67	66	58	63	58	63	67	67	58	48	58	68	67	68	58	65	58	63	67	69	58

c)

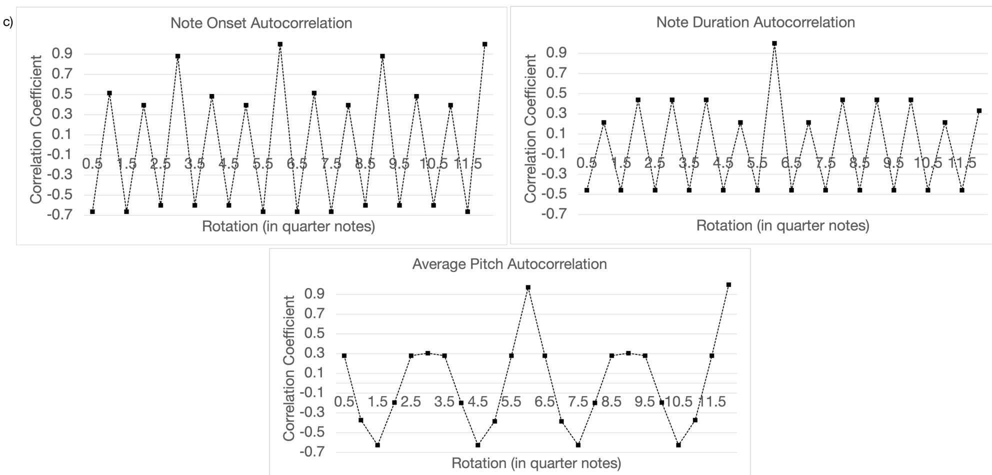

Figure 1. a) An adaptation of Fanny Mendelssohn Hensel's "Die Mainacht," Op. 9 No. 6, mm. 1-2; b) the music, represented as timelines of various parameters, such that each value corresponds to a consecutive eighth-note pulse, along with sample rotations; c) the resulting autocorrelation values for each timeline.

and behavioral aspects of musical meter. While a thorough comparison with template-matching/categorization approaches is outside the boundaries of this paper, I end by situating this paper's results within the broader scope of computational approaches to meter finding.

For optimal transparency and accessibility, the computational implementations of each of the following analyses are presented as spreadsheets in this paper's Online Supplement. In this format, readers with no computational background can examine the engineering behind my analyses, and the procedures can be replicated or implemented in any computational environment with relative ease. These materials can be found at chriswmwhite.com/autocorrelation.

2. Autocorrelation: its engineering and its usage in meter finding

At its most basic, an autocorrelation approach to meter finding frames a musical passage as a timeline of values, with each consecutive point in that timeline representing an isochronous pulse (i.e., they are separated by the same time interval). Figure 1a shows a score for the opening measures of Fanny Mendelssohn Hensel's "Die Mainacht," while FigureFigure 1b represents the events of each consecutive eighth-note pulse, and does so in three ways: as the number of note onsets at each moment, as the longest duration appearing at each moment (measured in eighth notes), and as the average pitch at each moment (measured in MIDI pitch, and counting only pitches that begin at each moment).

An autocorrelation method compares the initial – the prime – vector to each possible rotation of the timeline. Selected rotations appear under the note-onset timeline. Even visually, it is apparent that some of these rotations have more in common with the prime vector than do others. Rotating the vector by a single value – shifting each value by one eighth note – mostly aligns zeros with non-zeros: this vector is very unlike the prime. However, because the note-onset patterns of both measures are identical, rotating the vector twelve eighth-notes (or, a full measure) creates an identical correspondence. Autocorrelation quantifies these comparisons by correlating each rotation to its prime. For these calculations, a standard bivariate correlation is used (Pearson's r). Figure 1c shows the resulting correlation coefficients for all rotations for each parameter, with the x-axis metered such that each addition of 0.5 indicates a sequential rotation by eighth note. (NB: there can be some variation surrounding how to apply autocorrelation techniques; the approach outlined here rotates the whole length of the vector under consideration and consistently compares vectors of the same length to one another.)

To find the passage's meter, one can identify equally spaced peaks in the correlation coefficients, pinpointing isochronous patterns of consistent levels of similarity. In the current example, the most similarity occurs at the duration of the measure – events six quarter notes away from one another are quite similar. The note-onset and average-pitch timelines show the half-measure as the next most-similar isochronous periodicity, while the note-onset and duration timelines clearly show a quarter pulse. (The more-ambiguous quarter-note pulse of the average-pitch timeline is due to the stepwise motion of the melody and the skips to adjacent chord tones in the accompaniment: the average pitch heights of moments separated an eighth note are often quite similar in this texture!) In what follows, I use three musical excerpts to illustrate and describe the particular patterns identified by the procedure.

3. Three analytical vignettes: some representative examples of meter-finding using autocorrelation

3.1. *Procedure*

The following analyses use three parameters to construct timelines: note onsets, average pitch height, and duration. These three features are used because they have been shown to be salient to meter finding in both the cognitive and computational literature, with behavioral models showing that listeners associate a feeling of accent with relative peaks in note onsets (White 2019b), with certain durational patterns (Iversen, Patel, and Ohgushi 2008; Vos 1977; Woodrow 1951), as well as with changes in pitch contour (Huron and Royal (1996), Thomassen (1982), Acevedo, Temperley, and Pfordresher (2014), Prince and Rice (2018)). Similarly, computational models often rely on durational patterns (McLeod and Steedman (2017), Nakamura et al. (2017)) and note-onset data (Gouyon and Dixon (2005) and Gouyon et al. (2006); this parameter is generally used when analyzing symbolic score data, as onsets can serve as a proxy for loudness). Additionally, spectrogram models implicitly involve information about pitch distributions (e.g., Böck and Davies (2020); admittedly, the pitch information involved in a spectrogram will be more nuanced than the symbolic pitch data used here). Again, note that my goal here is simply to illustrate and study the patterns found within a period-extraction model: other musical features like harmonic change, melodic leaps, or phrasing – or the combination of multiple features – could also participate in the meter of these passages.

The process as implemented here identifies the shortest note duration, divides the pieces into consecutive events of that duration, and creates a timeline such that each value corresponds to each consecutive slice (this duration was the eighth note in Figure 1, for instance). For simplicity, if the quickest duration of the passage only occurs a handful of times (as will be the case for the sixteenth-note pulse in Figure 2) or has an ambiguous duration (as will be the case for the grace notes in Figure 3), the next-quickest pulse is used. The three approaches are implemented as follows: the note onset approach shows the number of notes that initiate/begin at each pulse, with zero indicating either rests or sustained events. The average-pitch approach associates MIDI numbers with each pitch onset, and averages these numbers; sustained notes do not contribute to the average pitch height. The duration approach represents the length of the longest note begun at each timepoint (with the duration generating the timeline assigned a value of 1). For simplicity and brevity, the following analyses presents a selection of these approaches, tailored to the characteristics of each example. Finally, given that this paper is interrogating the patterns used in these processes rather than advocating for a particular approach, I do not implement a rigorous meter-finding or time-signature-identification step but rather simply comment on the periodicities that would be available to such a process.

3.2. *Vignette 1: fuzzy patterns in Ludwig v. Beethoven's string quartet, op. 18, no 1*

Figure 2 shows the opening to Ludwig v. Beethoven's string quartet opus 18, number 1, with its corresponding note-onset vector. Autocorrelation identifies a triple grouping of quarter notes, and duple groupings of 16th notes and 8th notes: these are precisely the periodicities indicated by the time signature. It also identifies a sextuple and 24-tet grouping of quarter notes, thus identifying a two-measure hypermeter, and the eight-measure phrase. These groupings are noted by the brackets underneath the timeline.

Several of these identified patterns are straightforward and visually apparent. The fact that several motives explicitly return (mm. 1 and 3, mm. 9 and 11) and several rhythmic profiles return (mm. 2 and 4, mm. 6 and 8, 11 and 13, mm. 14 and 16, mm. 15 and 17) results in the a two-bar hypermetric patterns. The duple divisions of the quarter-note pulse are also straightforward to identify: more events occur on the quarter pulses than on the alternating eighth pulses. The dotted-half-note patterns, however, are less immediately apparent, but are still present. The music's recurrent motives tend to place events with longer durations at the beginnings of measures while events with shorter durations appear at the ends of measures, thereby creating patterns on onsets whose density fluctuations at the dotted-quarter pulse.

This dotted-quarter pulse exhibits what I call *fuzzy* patterns (roughly adapting the adjective's usage in Quinn, 2001). The autocorrelation procedure illustrates peaks at that pulse not because all – or even most – events at the remove of that duration are similar, but because more similarities arise at that pulse duration than at other durations (while the rotations of 3 and 9 and 15 quarter notes return a coefficient of only 0.2, the surrounding values are considerably lower).

3.3. *Vignette 2: unbounded patterns in J. Straus's An der schönen blauen Donau—*

Figure 3 illustrates a piano reduction of the opening to the Johann Straus waltz "An der schönen blauen Donau", with timelines indicating note onsets and average pitch height, now divided into each quarter pulse. The patterns of note onsets produced by the "Oom-

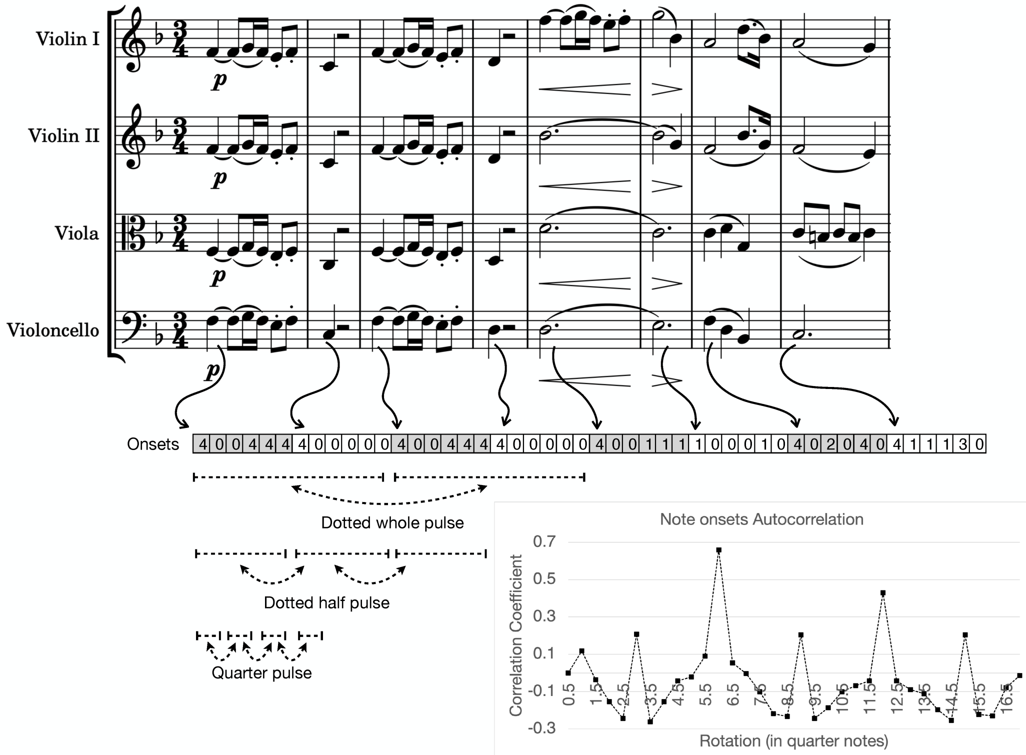

Figure 2. Ludwig v. Beethoven, String Quartet, Op 18, no 1, mm. 1-8 and accompanying autocorrelations

pa-pa" oscillations of the waltz provide a clear triple pattern; the relatively low pitches that initiate each measure create a similar three-quarter pattern within the average pitch height timeline. The melody's consistent shift between a treble register (mm. 1, 5, 9) and a soprano register (mm. 2-4, 6-8) provides a further four-measure pattern within that approach.

This example provides further instances of fuzzy patterns returned within this passage: a rotation of two-measures aligns events like the rising three-quarter-note gestures of measures 1 and 5 with the characteristic Viennese accents on the first and third beats of measures 3 and 7, resulting in a low correlation (a -.17 correlation coefficient, shown in the supplement). Again, however, this alignment produces a correlation better than the surrounding quarter notes. Such local maxima within otherwise-low coefficients can combine with spikes at the third and ninth rotations to express the dotted-quarter pulse.

More notable, however, is that the arrangement of correlations does not indicate that the first measure would be felt as an upbeat to the proceeding measures. Note that Figure 3 includes hypermetric and phrasing annotations for the passage (drawn from Rothstein (1989)) which show that – while there exists a four-bar hypermeter and a four-bar phrase – the phrase and hypermeter are out of phase, such that the first measure is felt as a hyper-upbeat to the second. The autocorrelations are agnostic to this difference, indicating only that there is a four-bar hyper-measure, not where it begins. These patterns simply rely on their periodic similarities rather than specific boundaries that denote their beginnings and endings – these patterns are unbounded. Below, I'll return to the implications of unbounded patterns as regards this procedure's identification of meter. (The "unboundedness" of the patterns in somewhat obscured in the orthogra-

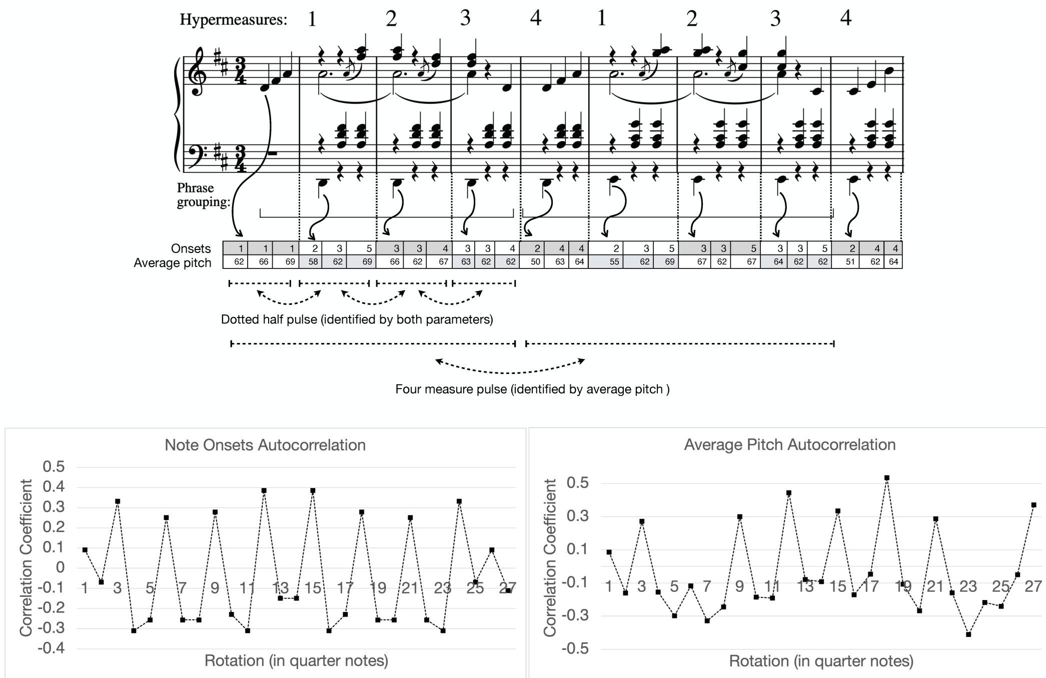

Figure 3. J. Strauss II, An der schönen blauen Donau, Op. 314, mm. 19, with annotations from the analysis of Rothstein 1989, 4-7

phy of the figures, as the brackets have beginnings and endings. It is useful to recall that, for instance, two adjacent measure-long brackets simply illustrate that events removed at the distance of a measure in that passage are similar.)

3.4. *Vignette 3: excessive patterns in Johann S. Bach's WTC 1, no. 1*

Figure 4 shows the first two measures of Prelude 1 from Johann S. Bach's *Well Tempered Clavier*, Book 1, along with duration and pitch timelines with sixteenth-note divisions. Both timelines identify periodicities at the half-note pulse, a predictable result given that Bach's figuration pattern repeats every half measure. What's more notable, however, are the peaks in correlation at 3, 5, and 7 rotations of the pitch-height timeline, indicating similar patterns of events at the remove of the dotted-eighth, five-sixteenth, and seventh-sixteenth periodicities. Similar peaks occur at the same relative points corresponding to each measure within the correlation graphs.

This unusual series of peaks is due to the repeated figuration arpeggiating the underlying chord: in each half-measure pattern, the initial two sixteenth notes are followed by two repetitions of the same three-note pitch pattern. Because of this, a rotation of three sixteenths will exactly align this trio of pitches with one another; similarly, a rotation of five sixteenths will align the second triple group in the first figure (the 6th, 7th, and 8th events in the prime timeline) with the first triple group in the second figure (the 11th, 12th, and 13th events). Additionally, exemplifying a more fuzzy pattern, the first five and last three notes of the figure produce consistently rising contours, and rotating the vector at the remove of those corresponding patterns produces higher correlations than do rotations of the surrounding values. (Notably, this analysis resonates with the 2+3+3 non-isochronous accent pattern with which Cone (1968) reads in this piece.)

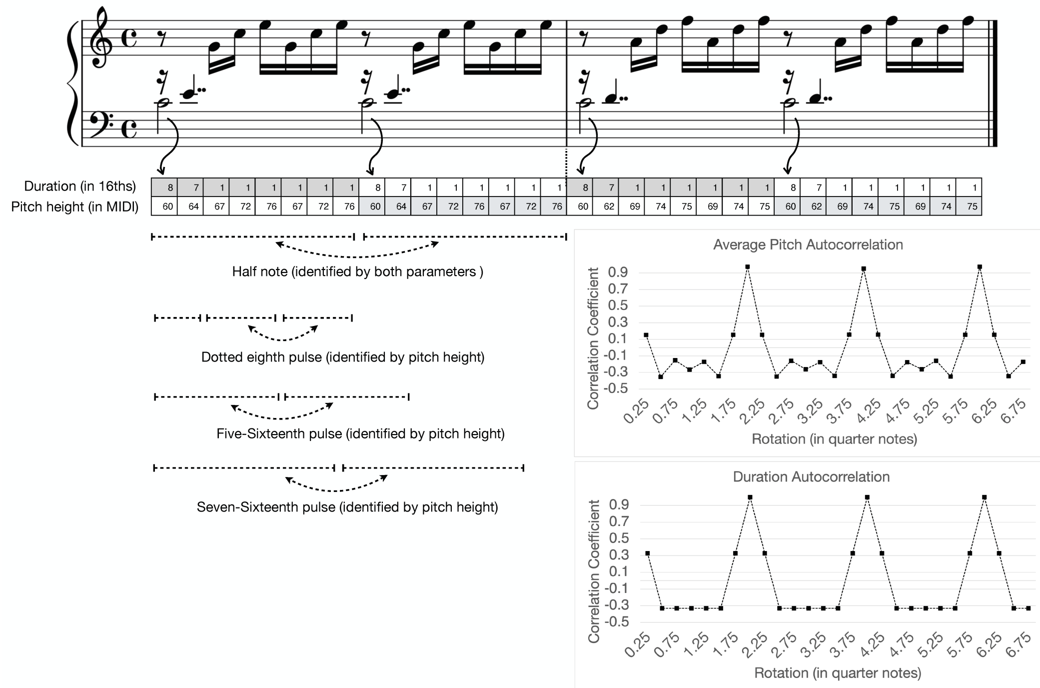

Figure 4. J.S. Bach's Well Tempered Clavier, Book 1, Prelude 1, mm. 1-2

These patterns do not align with the notated – and likely perceived – prevailing duple meter of this passage. A human analyst would would almost certainly discard these triple, quintuple, and septuple patterns as non-metrical, either because their own divisions would be non-isochronous (i.e., they mix duple and triple groups), because they do not evenly divide the clear half and whole-note pulses, or because they are insufficiently regular (they "re-begin" in each measure). However, these patterns do exist in this music, and a local, pitch-based understanding of this passage would notice such patterns. These patterns appear to exist in excess of metrical frameworks. Such *excessive* patterns are not uncommon within autocorrelated musical timelines (White 2019a), and suggest that a musical surface includes more isochronous patterns than contribute to the prevailing meter.

4. Pattern finding and musical meter

The patterns outlined offer overlaps with and insights into theoretical and cognitive conceptions of meter. In what follows, I outline some of these intersections, particularly in how fuzzy, unbounded, and excessive patterns function in these analyses.

4.1. *Fuzzy patterns*

These examples consistently feature fuzzy patterns, or patterns whose periodic similarities are identified simply because they were more similar than other periodicities. On the one hand, such patterns would seem to align with what we know about humans' capacity to project metric periodicities onto musical stimuli, even if the stimuli provide very sparse cues (Dawe, Platt, and Racine 1994; Toiviainen and Snyder 2003), conflicting/competing cues (Pfordresher 2003; Prince, Thompson, and Schmuckler 2009; Prince and Rice 2018; Repp 2007), and even no cues (Bolton 1894; Brochard et al. 2003). Furthermore, once a listener has entrained to a metric pattern, they have the capacity to continue hearing a metric pattern even in the face of conflicting or sparse information (London 2004; Krebs 1999; Lerdahl and Jackendoff 1983).

On the other hand, fuzzy patterns also have the potential to be at odds with perceived meter. There are limits to the listeners' capacities to identify and entrain to patterns in music, be the patterns too subtle to notice (Fraisse 1946; Szelag et al. 1998; White 2017; Woodrow 1909) or in conflict with a pattern from another musical domain (Ellis and Jones 2009; London, Himberg, and Cross 2009; Prince, Thompson, and Schmuckler 2009). The subtlety of some of the patterns in the Figure 4, for instance, would seem difficult – if not impossible – to hear. While a computational approach, then, might focus on identifying any pattern present within a musical timeline, some such patterns might be *too fuzzy* for a cognitive model. To align a pattern-finding model of musical meter with cognition, then, one would need to import some definition of what is "too fuzzy" to be perceived by a human listener to contribute to a felt meter.

4.2. *Unbounded patterns*

As observed in Eck and Casagrande (2005), any model based solely on recurrent patterns will not indicate the location of relatively strong beats, or even where a metric pattern starts. The pattern-finding process simply identifies that events removed by some consistent periodicity are similar. It does not say where those patterns begin. Recall that Figure 3 included two parallel annotations, that of the four-bar phrase, and that of the four-bar hypermeter, the former beginning in measure 1, the latter beginning in measure 2. A pattern-finding procedure like autocorrelation does not distinguish between these two readings: it simply indicates that events are quite similar to events that occur at the distance of four measures. In order to distinguish the two parsings of this passage, the process would need some definition of "accent" or "beginning," and such definitions are not available to a model that engages purely in period extraction. In other words, any meter-identifying process that focuses solely on isochronous patterns might be able to identify the periods with which events recur in some signal, but it will not be able to assign a downbeat within that period without some further definition of what constitutes a downbeat. I return to this issue in my final comparison with template-matching procedures. It should be noted that Eck and Casagrande (2005) do indeed devise an adaptation to autocorrelation that identifies the initiation of patterns; however, their additions result in their overall method being a hybrid between period extraction and template matching. Their adaptation uses Shannon entropy to equate more contextually unusual/less predictable events with metrical stress, something which can be seen to import a definition of accent to categorize strong and weak pulse layers.

4.3. *Excessive patterns*

In Figure 4, the triple groupings embedded within the repeated 8-note figure produced several peaks at periodicities that conflict with the notated duple meter; as noted above, an analyst or listener would find no shortage of reasons to discard these patterns in favor of a purely duple reading. But such patterns are present on the musical surface, and they do pervade the passage. As has been demonstrated elsewhere (Conklin 2010; Jehan 2005;

White 2019a), discovering patterns on a musical surface poses little problem for a computational model of meter; rather, choosing which patterns to use in the final assessment presents a challenge. The excessive patterns returned by period extraction also demonstrates a notable interconnection between meter and other types of musical patterns by highlighting the porous boundary between metric patterns and motivic/thematic groups. Indeed, while this paper highlighted the role of automated pattern finding for meter identification, the identification of recurrent strings of similar musical events has aided the study of piece-specific and style-generic event sequences (Cambouropoulos 2001; Conklin 2010) as well as musical form (Bañuelos and Orduña 2017). Indeed, researchers have studied the grey area that exists between recurrent patterns that participate in a meter and those that are simply thematic, being the focus of both music theorists (Krebs 1999; Riemann 1903) and cognitive researchers (Acevedo, Temperley, and Pfordresher 2014; Prince 2014). The output of a period-extracting model like autocorrelation explicitly shows the shared resources used by (and definitional overlap between) both metric and motivic musical patterns.

4.4. *Periodic pattern finding and meter*

In my initial discussion, I invoked London's dichotomy between period extraction and template matching as a way to heuristically divide computational approaches into those that rely on identifying surface patterns versus those that categorize pulses and events. I also noted that isochronous pattern-finding models have generally been shown to underperform models that involve some manner of expert knowledge or training session. I end this paper by considering what this apparent imbalance might tell us about the nature of meter and the cognitive act of meter finding. First, as observed in Jehan (2005), "downbeat estimation requires some fair amount of prior knowledge", and the simple act of period extraction lacks this prior knowledge. This study demonstrates exactly why this prior knowledge is required, as models such as those used in the current paper lack the ability to a) provide boundaries on the kinds of imprecise and fuzzy patterns that are considered salient, b) include some definition of accent or initiation to distinguish between levels of metric strength, and c) select from the array of possible patterns in such a way that ensures sufficient regularity and the duple and triple relationships between pulses. Periodicities on the musical surface are therefore necessary but not sufficient when determining the perceived or intended meter of a musical passage. While a full review and comparison of template-matching/categorization models is outside the purview of this paper, this insufficiency can be rectified by pre-programming a model with some definition of accent (McLeod and Steedman 2017) or including a deep-learning training session (Böck and Davies 2020).

And so: what of period extraction and pattern finding? Regardless of its relative efficacy, isolating periodic patterns within a meter finding task demonstrates the role these patterns play in the larger concept of meter. At the broadest level, this study suggests that meter is a phenomenon of interpretation and organization: while isochronous and recurrent patterns may pervade some musical signal, they remain insufficient to model musical meter on their own. More specifically: modeling period extraction allows for a closer examination of characteristics of recurrent patterns that contribute to musical meter, yielding – for instance – insights into how those patterns overlap with notions of motivic grouping. More speculatively: this study's suggestion that period extraction is insufficient to identify a piece's meter in and of itself has implications for cognitive and embodied theories of metric entrainment. And yet, a cognitive model of meter *is* a model reliant on isochronous and periodic musical patterns – from a listener's standpoint, meter allows for the organization of events into equally spaced patterns that both explain past events and predict future ones (Hasty 1997; Kozak 2020; London 2004). If a listener's experience of a musical meter involves *both* period extraction *and* template matching, this study suggests that the former is crucially reliant on the latter to produce a coherent metric assessment.

Acknowledgements

I thank Darrell Conklin for his guidance and feedback during this project's development. I would also like to thank the two anonymous reviewers for their invaluable feedback on earlier drafts.

Supplemental online material

Supplemental online material for this article can be accessed at `http://www.chriswmwhite.com/autocorrelation`. The supplement contains all data associated with the autocorrelations presented in this paper.

ORCID

Do not change this. Production will take care of it if the paper is accepted.

References

Acevedo, Stefanie, David Temperley, and Peter Q. Pfordresher. 2014. "Effects of Metrical Encoding on Melody Recognition." *Music Perception: An Interdisciplinary Journal* 31 (4): 372–386. `http://www.jstor.org/stable/10.1525/mp.2014.31.4.372`.

Amiot, E. 2016. *Music through Fourier Space: Discrete Fourier Transform in Music Theory.* Heidelberg: Springer.

Bañuelos, Cristian, and Felipe Orduña. 2017. "Dynamic Time Warping for Automatic Musical Form Identification in Symbolical Music Files." In *Mathematics and Computation in Music,* edited by Octavio A. Agustín-Aquino, Emilio Lluis-Puebla, and Mariana Montiel, Cham, 253–258. Springer International Publishing.

Böck, Sebastian, and Matthew Davies. 2020. "Deconstruct, Analyse, Reconstruct: How to improve Tempo, Beat, and Downbeat Estimation." In *Proc. of the 21st Int. Society for Music Information Retrieval Conf (virtual),* 574–582.

Böck, S., F. Krebs, and G. Widmer. 2016. "Joint Beat and Downbeat Tracking with Recurrent Neural Networks." In *Proceedings of the 17th International Society for Music Information Retrieval, New York,* .

Bolton, T. L. 1894. "Rhythm." *The American Journal of Psychology* 6 (2): 145–238.

Boone, G.M. 2000. "Marking mensural time." *Music Theory Spectrum* 22 (1): 1–43.

Brochard, R., D. Abecasis, D. Potter, R. Ragot, and C. Drake. 2003. "The "ticktock" of our internal clock: Direct brain evidence of subjective accents in isochronous sequences." *Psychological Science* 14: 362–366.

Brown, J.C. 1993. "The determination of meter of musical scores by autocorrelation." *Journal of the Acoustic Society of America* 94 (4): 1953–1957.

Cambouropoulos, Emilios. 2001. "Melodic Cue Abstraction , Similarity , and Category Formation : A Formal Model." *Music Perception* 18 (3): 347–370.

Cohn, Richard. 2001. “Complex hemiolas, ski-hill graphs and metric spaces.” *Music Analysis* 20 (3): 295–326.

Cohn, Richard. 2020. “Meter.” In *The Oxford Handbook of Critical Concepts in Music Theory,* edited by Alex Rehding and Steven Rings, 207–233. New York: Oxford University Press.

Cone, Edward T. 1968. *Music through Fourier Space: Discrete Fourier Transform in Music Theory.* New York: W.W. Norton.

Conklin, Darrell. 2010. “Discovery of distinctive patterns in music.” *Intell. Data Anal.* 14: 547–554.

Cooper, Grosvenor, and Leonard Meyer. 1960. *The Rhythmic Structure of Music.* Chicago, IL: The University of Chicago Press.

Dawe, L.A., J.R. Platt, and R.J. Racine. 1994. “Inference of Metrical Structure from Perception of Iterative Pulses within Time Spans Defined by Chord Changes.” *Music Perception* 12 (1): 57–76.

de Haas, W.B., and A Volk. 2016. “Meter Detection in Symbolic Music Using Inner Metric Analysis.” In *Proceedings of the 17th Conference of the International Society for Music Information Retrieval, New York, USA,* 574–582.

Durand, S., J. P. Bello, B. David, and G. Richard. 2015. “Downbeat tracking with multiple features and deep neural networks.” In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Brisbane,* 409–413.

Durand, S., B. David, and G. Richard. 2014. “Enhancing downbeat detection when facing different music styles.” In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Florence,* 3132–3136.

Eck, D. 2006. “Identifying metrical and temporal structure with an autocorrelation phase matrix.” *Music Perception* 24 (2): 167–176.

Eck, Douglas, and Norman Casagrande. 2005. “Finding Meter in Music Using An Autocorrelation Phase Matrix and Shannon Entropy..” In *ISMIR 2005 - 6th International Conference on Music Information Retrieval, London,* 01, 504–509.

Eerola, Tuomas, and Petri Toiviainen. 2004. “Mir in matlab: The midi toolbox.” In *Proc. 5th International Conference on Music Information Retrieval, Barcelona,* .

Ellis, R. J., and M. R. Jones. 2009. “The role of accent salience and joint accent structure in meter perception.” *Journal of Experimental Psychology: Human Perception and Performance* 35: 264–280.

Fraisse, P. 1946. “Contribution à l’étude du rhythme en tant que forme temporelle.” *Journal de Psychologie Normale et Pathologique* 39: 283–304.

Gouyon, Fabien, and Simon Dixon. 2005. “A Review of Automatic Rhythm Description Systems.” *Computer Music Journal* 29: 34–54.

Gouyon, Fabien, Gerhard Widmer, Xavier Serra, and Arthur Flexer. 2006. “Acoustic Cues to Beat Induction: A Machine Learning Perspective.” *Music Perception: An Interdisciplinary Journal* 24 (2): 177–188. http://www.jstor.org/stable/10.1525/mp.2006.24.2.177.

Hasty, Chris F. 1997. *Meter as rhythm.* New York: Oxford University Press.

Huron, David, and Matthew Royal. 1996. “What Is Melodic Accent? Converging Evidence from Musical Practice.” *Music Perception: An Interdisciplinary Journal* 13 (4): 489–516. http://www.jstor.org/stable/40285700.

Iversen, John R., Aniruddh D. Patel, and Kengo Ohgushi. 2008. “Perception of rhythmic grouping depends on auditory experience.” *The Journal of the Acoustical Society of America* 124 (4): 2263–2271. https://doi.org/10.1121/1.2973189.

Jehan, Tristan. 2005. “Downbeat prediction by listening and learning.” In *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics, New Paltz,* 11, 267 – 270.

Khadkevich, M., T. Fillon, G. Richard, and M. Omologo. 2012. “A probabilistic approach to simultaneous extraction of beats and downbeats.” In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Kyoto,* 445–448.

Kim, J. C., and Edward Large. 2017. “Entrainment of Canonical Oscillators to Complex Rhythms: Temporal Receptive Field Revisited.” *Delivered at Society for Music Perception and Cognition National Conference, San Diego.* .

Kozak, Mariusz. 2020. *Enacting Musical Time: The Bodily Experience of New Music.* New York: Oxford University Press.

Krebs, Harold. 1999. *Fantasy Pieces: Metrical Dissonance in the Music of Robert Schumann.* New York: Oxford University Press.

Large, E.W., J.A. Herrera, and M.J Velasco. 2015. “Neural Networks for Beat Perception in Musical Rhythm.” *Frontiers in Systems Neuroscience* 9: 159–173.

Lerdahl, Fred, and Ray Jackendoff. 1983. *A generative theory of tonal music.* Cambridge, MA: MIT Press.

London, Justin. 2004. *Hearing in time: Psychological aspects of musical meter.* New York: Oxford University Press.

London, Justin, Tommi Himberg, and Ian Cross. 2009. "The Effect of Structural and Performance Factors in the Perception of Anacruses." *Music Perception: An Interdisciplinary Journal* 27 (2): 103–120. http://www.jstor.org/stable/10.1525/mp.2009.27.2.103.

McLeod, Andrew, and M. Steedman. 2017. "Meter Detection in Symbolic Music Using a Lexicalized PCFG." In *Proceedings of the 14th Sound and Music Computing Conference, Espoo, Finland,* 373–379.

Nakamura, Eita, Kazuyoshi Yoshii, Shigeki Sagayama, Eita Nakamura, Kazuyoshi Yoshii, and Shigeki Sagayama. 2017. "Rhythm Transcription of Polyphonic Piano Music Based on Merged-Output HMM for Multiple Voices." *IEEE/ACM Trans. Audio, Speech and Lang. Proc.* 25 (4): 794–806. https://doi.org/10.1109/TASLP.2017.2662479.

Palmer, C., and C. L. Krumhansl. 1990. "Mental Representations for Musical Meter." *Journal of Experimental Psychology: Human Perception and Performance* 16: 728–741.

Papadopoulos, H., and G. Peeters. 2011. "Joint Estimation of Chords and Downbeats From an Audio Signal." *IEEE Transactions on Audio, Speech, and Language Processing* 19 (1): 138–152.

Pfordresher, Peter. 2003. "The role of melodic and rhythmic accents in musical structure." *Music Perception* 20: 431–464.

Prince, Jon. 2014. "Pitch Structure, But Not Selective Attention, Affects Accent Weightings in Metrical Grouping." *Journal of experimental psychology. Human perception and performance* 40: 2073–2090.

Prince, John, and T. Rice. 2018. "Effects of Metrical Encoding on Melody Recognition." *Journal of Experimental Psychology: Human Perception and Performance* 44: 1356–1367.

Prince, Jon B., William F. Thompson, and Mark A. Schmuckler. 2009. "Pitch and time, tonality and meter: how do musical dimensions combine?." *Journal of Experimental Psychology* 35 (5): 1598–1617.

Repp, Bruno H. 1998. "Variations on a Theme by Chopin: Relations between Perception and Production of Timing in Music." *Journal of Experimental Psychology: Human Perception and Performance* 24: 791–811.

Repp, Bruno H. 2007. "Hearing a melody in different ways: Multistability of metrical interpretation, reflected in rate limits of sensorimotor synchronization." *Cognition* 102 (3): 434–454. https://www.sciencedirect.com/science/article/pii/S0010027706000540.

Riemann, Hugo. 1903. *System der Musikalischen Rhythmik und Metrik.* Leipzig: Breitkopf und Haertel.

Rohrmeier, Martin. 2020. "Towards a Formalization of Musical Rhythm." In *Proceedings of the 21st Int. Society for Music Information Retrieval Conference (virtual),* 621–629.

Rothstein, William N. 1989. *Phrase Rhythm in Tonal Music.* New York: Schirmer.

Szelag, E., J. Kowalska, K. Rymarczyk, and E. Pöppel. 1998. "Temporal integration in a subjective accentuation task as a function of child cognitive development." *Neuroscience Letters* 257 (2): 69–72.

Thomassen, Joseph M. 1982. "Melodic accent: Experiments and a tentative model." *The Journal of the Acoustical Society of America* 71 (6): 1596–1605. https://doi.org/10.1121/1.387814.

Toiviainen, P., and J.S. Snyder. 2003. "Tapping to Bach: resonance-based modeling of pulse." *Music Perception* 21 (1): 43–80.

Velarde, Gissel, David Meredith, and Tillman Weyde. 2016. *A Wavelet-Based Approach to Pattern Discovery in Melodies*, 303–333. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-25931-4_12.

Volk, Anja. 2008. "The Study of Syncopation using Inner Metric Analysis: Linking Theoretical and Experimental Analysis of Metre in Music." *Journal of New Music Research* 37 (4): 259–273.

Vos, P. 1977. "Temporal duration factors in the perception of auditory rhythmic patterns." *IScientific Aesthetics/Sciences de l'Art* 1: 183–199.

White, Christopher Wm. 2017. "Relationships Between Tonal Stability and Metrical Accent in Monophonic Contexts." *Empirical Musicology Review* 12 (1): 19–27.

White, Christopher Wm. 2019a. "Autocorrelation of Pitch-Event Vectors in Meter Finding." In *Mathematics and Computation in Music, Proceedings of the 7th International Conference,* edited by Mariana Montiel, Francisco Gómez-Martín, and Octavio A. Agustín-Aquino, Lecture Notes in Computer Science, 287–296. Heidelberg: Springer.

White, Wm., Christopher. 2019b. "Influences of chord change on metric accent." *Psychomusicology: Music, Mind, and Brain* 29 (4): 209–225.

Woodrow, H. 1909. "A Quantitative Study of Rhythm." *Archives of Psychology* 18: 1–66.

Woodrow, H. 1951. "Time Perception." In *Handbook of Experimental Psychology,* edited by S.S. Stevens, 11224–1236. New York, NY: Wiley.