Annotation of Musical Scores

Requirements and considerations for digital implementation

Defining annotations

What is an annotation of a musical score?

For the purposes of this document, we define an annotation of music notation as associating some information with elements of the musical score. Thus the main participants in this annotation are '**some information**', defined very loosely, and '**elements**' of the score. The latter may be explicit graphical symbols, but they may also take more abstract forms, such as a temporal range (e.g. 'from bar 2, beat 3 to bar 4, beat 1') which need not be embodied in any explicit symbols. In many current instances, the range is graphical, with membership of the annotation implied by enclosure within a drawn symbol

What isn't an annotation of a musical score?

We are deliberately catholic in our definition. We may exclude particular annotations from consideration in particular activities (such as implementation in Verovio), but that should be explicit. The range of activities that can be considered annotations is large, but the overlap of form and function justifies this.

Our central requirement for consideration in the current project is the use of musical score in a form that could be represented in MEI (even if it currently is not represented that way). This is a purely practical consideration. This limitation of direct scope also excludes document-level annotations, which several respondents mentioned. These are important and interesting, but will not be implemented within the current project.

Why make annotations?

Annotations are used for a variety of purposes, but are very rarely carried out in a digital, interoperable, machine-readable form. Purposes include:

- **Performance and pedagogy** – often in the form of overlaid music notation augmenting or altering existing markings. Conductors and directors may want to share markings between performers (as may section leaders in orchestras), and this is accommodated in some software.

- **Analysis and criticism** – labelling or commenting on regions of music to explain aspects of it. This can take many forms, but often associates a label or some text with a musical line or region. Some analyses (e.g. voice-leading or Schenkerian analysis) annotates musical regions with fresh, often specialised, notation.

- **Edition** – qualifying editorial decisions and supporting them with justifications and citation of sources used as witnesses. Those preparing editions may also annotate drafts with corrections or layout suggestions.

- **Performance studies** – a score can be annotated with analyses of how it is performed. This may relate to other media, including sound recordings.

Making and using annotations

Probably the most common way of annotating music is by drawing on a score – often with a pencil or pen on paper, but increasingly with a stylus on a tablet or touchscreen. This approach is intuitive and immediate, and clearly very important. It often involves

circumscribing notational elements (usually with a circle) – with extra information written or drawn beside the circle – or direct notational interventions – with music notation being overwritten into the printed text.

The former of these patterns – circling notation – is common where annotation is being *created*, but less common where it is being *typeset* for readers. For example, few analytical or pedagogical texts use circling in music examples or teaching scores, though brackets and boxes may be used. This suggests that intuitive notations for creating annotations may be different from those that are judged most effective for communicating them (of course, it may also relate to limitations of printing or the imaginations of book compilers).

Representing annotations

Some forms of annotation practice are part of established methodologies and, in those cases, there may exist specialised ways of representing them. One example is the **`<harm>`** element in MEI[1], which allows the encoder to record a variety of harmonic labels or analytical marks, and comparable analytical 'spines' exist in HUMDRUM[2].

Since annotations are often the result of novel research, we also need to cater for a general case in which the information being added could not have been prescribed. Both specialised and generic forms have been approached in a range of ways in digital and computational musicology. Examples of more generalised structures being used for annotations include the Dezrann (bespoke JSON)[3], CRIM (bespoke objects with EMA URLS (see below), with some use of web annotations)[4], work by IReMUS/CEDRIC/CNAM (novel Linked Data)[5] and that built around the Music Encoding and Linked Data (MELD) framework (Linked Data web annotations with MEI)[6]. Although most of these can and do operate on music notation and audio/visual media, they are primarily symbolic in focus. In the primarily audio domain, examples include JAMS (JSON)[7], Schubert Winterreise Dataset (bespoke CSV)[8], Sonic Annotator (Linked Data)[9] and the Polifonia music annotation pattern (Linked Data based on JAMS)[10].

One widely-adopted, more abstract approach to this problem is the Web Annotation model and accompanying Linked Data ontology (Sanderson et al, 2017), which provides structures for annotating generic, web-addressable objects with generic, web-addressable annotations.

Components of an annotation

The annotation with which we are concerned in Annote is taken as associating information with elements of a musical score, so it follows that there are two main components:

---

[1] https://music-encoding.org/guidelines/v5/content/analysisharm.html#analysis
[2] See https://www.humdrum.org/guide/
[3] See https://algomus.fr/dezrann and, for example, Ma, et al. (2019).
[4] https://sites.google.com/haverford.edu/crim-project/home
[5] See Cherfi et al (2017)
[6] https://meld.web.ox.ac.uk and, for example,.
[7] https://jams.readthedocs.io/en/stable/ and Humphrey et al (2014)
[8] Weiß et al. (2024) & Zenodo. https://doi.org/10.5281/zenodo.10839767
[9] Cannam et al (2010)
[10] Berardinis et al (2023)

The two components are: the information to be associated and the musical notation to which the annotation applies. In addition to these two components, there is the annotation itself, which we will want to augment with information about its provenance – who made it and when, etc. – and perhaps about what sort of annotation it is and how it should be interpreted.

This aligns very closely with the web annotation model (see below, from https://www.w3.org/TR/annotation-model/). In this model, the 'area of interest' of the annotation is called a **Target** (which can be anything with a URI). Information that is associated with the target by the annotation is called a **Body** (the annotation itself in whatever form it takes, including arbitrary text). Annotations can include information about motivation, and record who or what created them. We shall use this as the primary terminology in the sections that follow..

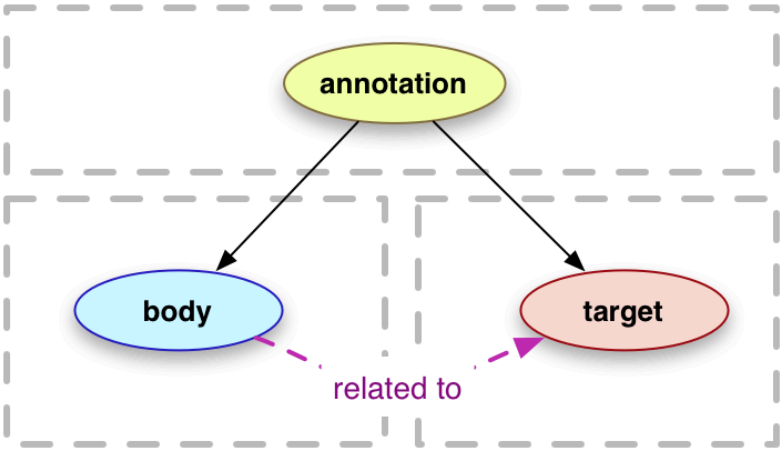

The annotation target: identifying a musical region in digital notation

Whatever form an annotation takes, it must have a target – some circumscribed area of interest within one or more documents. Broadly speaking, there are three categories of mechanism for indicating a 'target' region in music notation:

- **Graphically:** Choosing some coordinate system for the rendered notation, this approach records shapes that bound the target regions, or which describe the annotation shapes. This approach is primarily favoured by tablet-based systems for performers, and it can be challenging to preserve these annotations where scaling and re-formatting the score is possible.
- **Enumeration:** Listing notation elements to be affected by the annotation. This can be achieved directly – inline – by marking rows (in tabular formats, such as HUMDRUM) or creating parent elements (in XML-based formats, such as MEI), or

indirectly – standoff – by literal enumeration. In XML, any element with an id can be indicated using the fragment component of a URL. This approach is unambiguous and easy to compute, but it cannot indicate regions that do not contain elements (such as beats 2-4 of a bar with only a semibreve in it). The approach of using `xml:id`s is also vulnerable to changes to the source file, especially if these attributes are automatically generated.

- **Range specification**: Taking the 'axes' for 'coordinates' of a score to be metrical (horizontally) and part/staves/layers (vertical), this approach specifies a musical region in a way that is less reliant on the precise details of the format or the encoding. The model used to specify range is crucial here, and existing implementations can struggle with annotations where the target is non-metrical (such as articulation or grace notes) or with music notations for which one or other of the axes is less appropriate (for example, notations lacking measures, polyrhythms or scores with non-sequential parts). A range-based model (usually based on seconds) is near-universal for audio-oriented models (e.g. JAMS and Polifonia) and common in notational models too (using beats, as in EMA, Dezrann)

An inline targeting approach, whether graphical, enumerative or ranged, requires the annotator to have write access to the document or to copy and republish it. This has implications for keeping track of 'same' versions of the same document or edition, and a heavily annotated document would either be duplicated many times (potentially with alterations) or evolve into a single, central, very complicated file with many annotation structures within it.

It should also be noted that these approaches are generally graphical or performance-time-based, with little reference to the logical. In cases where a particular passage recurs, and where this is indicated with repeat or da segno signs, it is not clear how to distinguish between targeting all recurrences and specific ones.

Annotations can be more subtle than we have implied above. Not all annotations are really 'about' the encoding or the edition itself – they will often be directed at an abstract concept embodied by the notation, such as a theme, a phrase or a musical gesture. Annotations may also be connective, joining musical instances together, or connecting musical and non-musical elements. The Music Annotation Ontology (MAO) appears to be unique in accommodating such musical abstraction. It provides for a step of indirection, allowing annotations to target musical ideas that are then embodied in the notation files, images or recordings. This allows, for example, describing the behaviour of multiple arrangements of the 'same' musical idea (Lewis et al, 2022). The MAO is used currently used in the Beethoven in the House Annotator (Lewis et al, 2023) and in mei-friend (Goebl and Weigl, 2024; Plaksin, 2024).

Although the Web Annotation supports multiple targets, there is no explicit discussion of annotating directional relationships between targets (for example, to say 'this material is derived from this material'). An extension to the model would be required for this.

The annotation body: attaching information

The range of potential information to attach is large. Strategies for structuring annotation bodies include:

- **Textual content**: This approach allows for easy integration into visualisation applications (see, for example, the Beethoven in the House annotator, Lewis et al. 2023). It is directly supported in JAMS and within the MEI **<annot>** element (see below).
- **Notational content**: Many notational elements can be regarded or used as annotations. One common example (from our survey) is fingering indications. In MEI, these would be encoded directly in the edition, optionally using tags such as **<supplied>** to show their editorial nature.
- **Generic data structures**: If annotations are primarily quantitative, the annotations can be characterised by a range of primitive data types (labels from a limited vocabulary, numbers, vectors of numbers). This is the approach used by Sonic Visualiser
- **Bespoke data structures**: Dezrann has been extended to accommodate specialised data structures as they are needed by particular annotators. Thus, the semantics of the format evolve over time. Previous MELD applications also operated on tailored data structures.
- **Arbitrary data structures**: Use of the web annotation ontology permits the association of anything with a URI as an annotation body. Consuming software would almost certainly require further constraints to operate.

Information about the annotation itself

Associating information about the creation of an annotation – for example, attribution of authorship, creation date and motivation – is important, especially for scholarly work. This sort of ‘information provenance’ is supported to varying degrees by JAMS, web annotations, sonic annotator and others. Since this is relatively well studied, the structure of information provenance information will be considered beyond the scope of this report, though the visualisation of that information is within scope.

Implementation in MEI

Two standards for encoding annotations are particularly relevant for this research:

- The Web Annotation model (see above)
- MEI **<annot>**

We have mentioned Web Annotations above. The **<annot>** element in MEI is more complex and, currently, more ambiguous. Its primary intention, according to the guidelines, appears to be to annotate encoding processes or decisions (and sometimes ones of editorial practice): ‘In all cases, **annot** provides a comment upon a feature of the encoding’ (Dev guidelines, §9.2.13. Annotations). However, in the documentation of the element itself, the guidelines go beyond this scope: ‘The **annot** element can be used for both general comments and for annotations of the musical text.’ A new attribute may prove necessary to clarify the role of **<annot>** in a given context.

The structure of the **annot** element may be compatible with other models for annotation, and it has many pointing mechanisms associated with it, which makes it very flexible.

Several attributes are provided to identify the target (musical subject) of the annotation, including: **plist** allows reference by a list of xml:ids; **tstamp**, **tstamp2** and **dur** support time ranges; and **startid** and **endid** specify elements that bound the subject region. The text also implies that the annotation can also apply directly to its containing element, which might be confusing in practice and is implicitly indicated as unusual in the guidelines.

If the region being annotated is a combination of voices in the score, multiple parts can be indicated using the **staff** attribute. To identify a single voice within a staff, this can be combined with the **layer** attribute. Although **staff** can take a list of ids, **layer** indicates sequence number, and so would be ambiguous if, for example, an annotation applied to the upper voice of the first staff and the lower of the second. Similarly, only one pair of temporal bounds may be specified, so an irregular-edged region (such as a set of fugal entries) or a set of non-contiguous regions cannot be specified in a single annot element.

The body of the annotation – that is the information intended to attach to the music notation is, seemingly, only indicated by adding child elements. There is some variety of permitted child elements, however, presumably intended to ensure that this is expressive enough. Since the generic pointer element, **ptr**, is available, arbitrary URIs can be associated with the **annot**, although the meaning or implications of this are not explicitly spelled out in the guidelines.

It is also possible for many notational annotations to be applied directly in the music encoding. Many examples raised in our survey have their own specialised elements already, such as harmonic labels, fingerings, accidentals or editorial interventions. If there is need for a generic approach to annotations (for example, for making a set of annotations visible or hidden), then the range of forms that they take in MEI may cause difficulties. Similarly, these notational annotations are largely inline by default.

Applications and visualisations of annotations

Graphical annotations

Drawing on scores is overwhelmingly the most common form of annotation, and this is mirrored in the digital world. Software used can be generic PDF annotators, such as GoodNotes, or specialised for particular music use cases, as with Nkoda's layered, sharable annotations (Stokes, 2024), and the equivalent functionality in Enote (Pacha and Rettinghaus, 2022) and ForScore.

Although the same functionality is not available where the music is represented in a dynamic edition, investigation into how this might be implemented is clearly warranted given the demand. Such drawn annotations would need to be automatically or semi-automatically associated with a musical target area, and then would flow or scale as necessary.

Existing software is largely proprietary, but if a standard format was available for drawn annotations (comparable, for example, with those in IIIF) then transfer of annotations

between packages would be possible. Clearly this is only meaningful if the documents that provide annotation targets are also compatible.

[add screenshots]

Symbolic annotations

Dezrann

The Dezrann software and libraries (Giraud et al, 2018) are developed by the Algomus group in Lille, headed by Matthieu Giraud. The system uses the Verovio library for rendering, and so is compatible with the symbolic formats provided there. Given automatic or manual synchronisation, itself a set of annotations, Dezrann can also visualise audio and facsimile images. Annotations of different types are visualised differently, showing boxes, highlighting, vertical lines to label points in time and labelled horizontal ribbons above or below the stave to represent time regions.

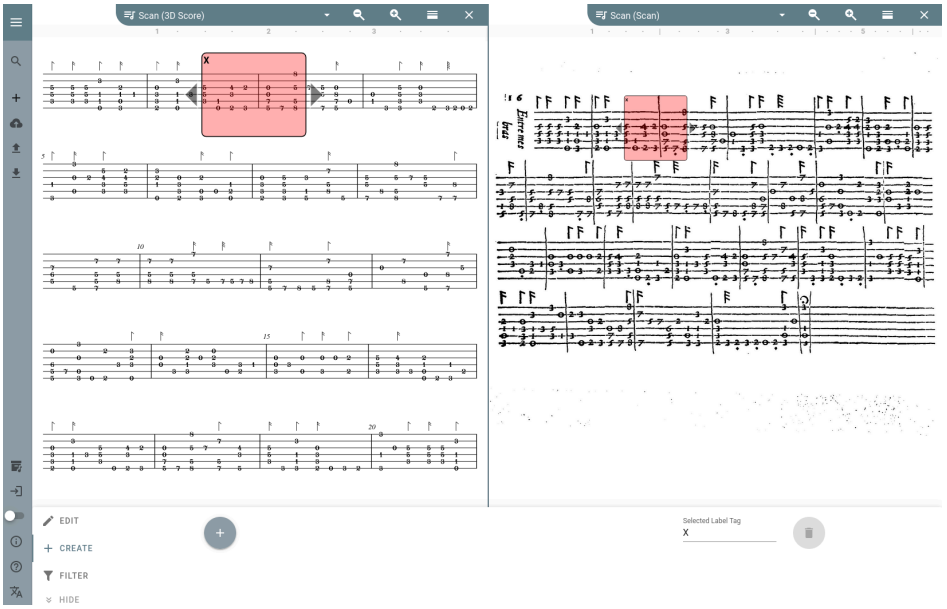

Data is stored as JSON files in a consistent, but bespoke, format, and the annotation targets are selected by staves, bars and (crotchet) beats.

Once a particular annotation type is established, data entry is usually graphical (see above illustration). Modes of visualisation generally focus on boxing (above) and bracketing (below)

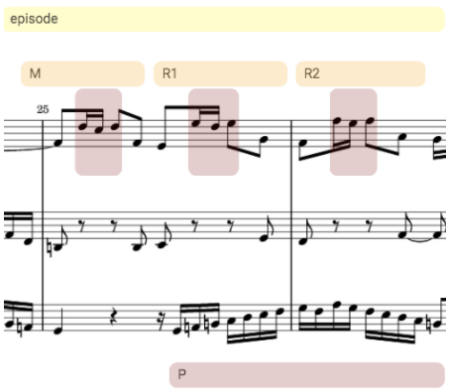

Although not mentioned by our survey respondents, some similar functionality is shown by the IREMUS lab's tools (Tonalities). This uses the Polifonia data model.

CRIM & the EMA API

The Citations in Renaissance Imitation Masses project (CRIM)[11] is the most notable user of the EMA URL schema and API. The targets of analytical assertions are specified as EMA URLs. This means that visualisations can use the EMA web service API, either directly, showing extracts from the music with relevant examples, or to provide the necessary identifiers to inform the highlighting of a full score rendered in Verovio, with the notes that are the target of observations being highlighted in the score.

MELD

Music Encoding and Linked Data (MELD)[12] is the JavaScript framework behind a selection of web applications, all of which use annotations in some form to describe relationships within multimedia resources. Verovio is used to render scores, with the resulting SVG image used to inform either interaction or visualisation, which can include bounding boxes (Weigl et al, 2021), note colouring and in-line music notation (Lewis et al, 2019).

Some applications give the ability to create annotations, and the Beethoven in the House annotator – which was developed by the University of Paderborn in collaboration with Oxford and uses MELD libraries for loading annotations – uses the Music Annotation Ontology (Lewis et al, 2022) to coordinate annotation of versions.

Verovio

Specialised elements exist in MEI for a few annotation-like structures, in these cases Verovio already provides some basic visualisation support (see figure below for visualisations of harmonic labels)

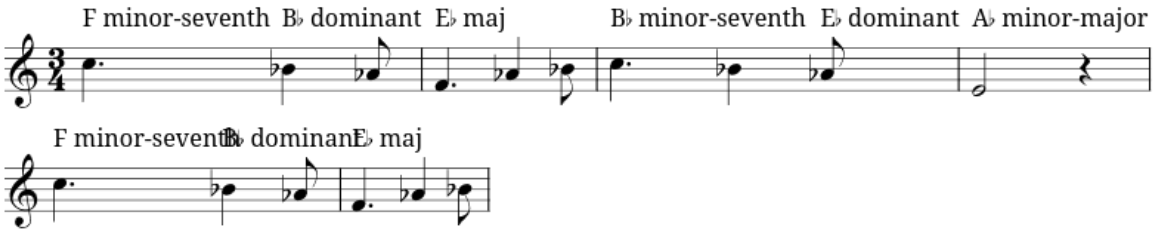

Music engraving software (Sibelius, MuseScore, etc)

Many survey respondents used engraving software for their annotations, but this was a far greater focus for dissatisfaction than more graphical applications. Frustrations included cluttered output, poor variety of symbols, poor anchoring, undesirable spacing and a lack of functionality for layering (and so hiding and showing) annotations, Whilst specialist software allows musically appropriate notation, respondents clearly found themselves restricted

[11] https://sites.google.com/haverford.edu/crim-project/home and Freedman, 2023
[12] https://meld.web.ox.ac.uk and, for example, Lewis et al, 2023

(compared with drawing) in what they could express, and with relatively little compensatory benefits.

In some cases, this meant producing pdf editions with engraving software and then annotating those.

User survey: main findings

We asked people who annotate music, whether digitally or on paper to fill in an online questionnaire. Following an initial engagement session at the IAML conference, the survey was open from 26 June-4 October 2023, and was publicised to music librarians, musicologists and digital musicologists at the IAML and DLfM conferences, through IAML, digital-musicology, musicology-all mailing lists and on Twitter/X from the @TMusicology account. 20 responses were received.

The importance of the graphical

There is a palpable frustration at the difference between the affordances of pencil and paper (or, to a slightly lesser extent, of programs that allow direct drawing on scores or images) and what is possible in existing music engravers or music editing software. This is based on a variety of factors:

- Restrictions in the range of symbols or notations available
- Perceived weaknesses in typesetting around annotations
- Shortcomings in target selection for annotations – often this is limited to single elements
- Data entry being less immediate than with a pencil or stylus

Individual needs

There is some disagreement between respondents about what is and isn't effective for visualising annotations, much of which can be attributed to different communication needs and use cases. For example, whether it is desirable for the user to interact with on-screen elements in order to see annotations will depend on how the score will be viewed by the final user – including whether it will be printed.

Also telling is that many respondents used a wide variety of visualisations themselves – embracing colours, circles, brackets, boxes, marginalia, tables and many more.

The best way to view annotations appears to be dependent on use, and here the digital has an advantage. The appearance of a drawn annotation is often related to the pencil movement required – it is based on facilitating 'data entry' – an electronic score can reconfigure its output for needs of the viewer, which may differ based on their own use, independently of the initial annotator.

In general, there seems to be dislike for annotations that overly change the spacing in a way that distorts musical logic. This may be beneficial for implementers, since the musical typesetting may be prioritised when generating a layout.

Finally, although there is heavy use of generic annotations (text and boxes) or standard notation (fingerings, phrasings, chord symbols), there are also many examples where specialised or even invented, bespoke symbols are needed and used. It seems unlikely that an attempt to codify all annotations would be successful, even if it were desirable. The ability to draw is currently the clearest (and least technical) way for such symbols to be used in a general-purpose system.

Layers of annotation

The concept of layers – groups of annotations that are associated with one another, and which can be hidden, customised, displayed or shared in one go – is popular when it occurs, and strongly desired when it doesn't. Layers are a generic user-interface model for implementing a variety of semantic situations – they can represent the annotations from a single music lesson, from a single annotator, or representing a particular analytical or text-critical unit.

This represents another case where the visualisation is calibrated to the needs of the user (removing unwanted analytical clutter).

Dynamic scores and high-quality typesetting

Drawing on a score is only guaranteed to be trivially reproducible where the layout and typesetting are static. The same is true, to a lesser extent, for marginalia. The semantics of any annotation whose targeting of music notation is solely by spatial juxtaposition tends to be dependent on layout. Although this is also true for musical symbols such as hairpins, and has proven largely solvable there, it does require appropriate treatment.

Several respondents, especially those who distribute their annotations for performers or students, particularly value a version of their edition that looks as clear and well-laid out as possible. It may be beneficial to make an explicit separation of dynamic display and final adjustment for static publication for annotations. This is comparable to the ability to position individual symbols in Verovio, something that might be used for final output, but is much less likely to be used for interactive, in-browser applications.

Provenance

Two situations described by survey contributors imply a need for tracking responsibility for annotations. Firstly, encoders may be transcribing prior annotations. Recording the circumstances of those annotations may be relevant for their work. Secondly, annotation may be collaborative, distributed or simply multi-layered, and tracking and attributing activity may be important for appropriate credit, for analysis and for layered display.

Connecting and navigating

Several contributors to our survey noted the possibilities of connecting annotations not solely to one notational region, for example:

- Connecting related formal or thematic regions

- Connecting related materials or media
- Connecting related annotations

This can be particularly useful for high-level visualisations (for thematic or formal structure), for rapid navigation and for juxtaposition of relevant material.

Requirements for annotations of musical scores

Representation requirements

We have argued elsewhere (Lewis et al 2023) that it should be possible for annotations to apply to arbitrary groupings of music notation elements. The requirements listed in this document stand, but are insufficiently precise for our current purpose. Annotation can operate through enumeration of symbols (e.g. MELD), temporal range and voice specification (e.g. EMA and Dezrann) or graphical demarcation (most tablet-based commercial apps). These three ways of selecting a target require different data models and usually imply different ways of visualisation.

Recommendations for targeting of annotations:

- **Enumeration**: this is most clearly represented by a list of all affected symbols. If this is not explicitly provided, then it must be possible to generate. Start and end symbols can only be considered equivalent if a) the order of symbols is unambiguous and b) multiple symbols ranges can be combined. Similarly, container elements (such as measures) may be used to indicate application to all child elements (symbols in that measure).
- **Temporal range:** Full, generalised expressive power for this requires that:
    - The range can encompass a set of temporal ranges for each metrically-independent musical layer.
    - The temporal boundaries can be specified in notation-relevant terms (so bars and beat units or crotchets may not be relevant for mensural notation, chant, or non-western or non-metrical scores)
- **Graphical demarcation**: This area requires more research. Clearly the shape of the demarcation should be captured, for example using a scalable graphic, but the implications of its shape and extent are likely to be ambiguous, making interpretation difficult. Thus a model that includes shape alone may not suffice for a usable representation.

In practice, we expect that implementations satisfy only some of these requirements, prioritised based on application.

Display requirements

Some requirements clearly emerge from our user survey

- **Annotation visibility should be switchable, ideally selectively (i.e. in groups)**: this is a recurring requirement in our survey

- **Music typesetting logic should not be impaired by annotations**:excess space (as in the Verovio example above) is considered disruptive to reading by several respondents
- **Annotation should be visually distinct from the main score**

Others are strongly implied

- **Annotation attribution**. It should be possible to indicate the provenance of an annotation, and that annotation may have its origins in a non-digital medium. Switchable visibility by annotation provenance is a common request.
- **Generic annotation visualisation should be functionally separated from core typesetting** - Optimal annotation visualisation is application-dependent. This suggests that for many applications, the process of drawing them should be separated from, but assisted by, the music engraving library. This complexity may later be reduced for publication by making the annotation visualisations a static part of the rendered score (for example, by serialising them as notational elements in an MEI file rather than annotations).
- **A printable/performable option with annotation material visible should be possible**

Optional requirement, about which users may disagree:

- **Interaction: 'Expand on interaction' annotation views**
    - A notable minority of survey respondents suggested interaction as a way to avoid visual clutter – that an annotation 'body' could be displayed on hover or click (as in the digital edition of Tinctoris' theoretical works)
    - Other respondents clearly have a requirement for printed materials, for which interaction would be ineffective
- **Navigation by annotation**: Depending on what the annotations are and how they relate, they may provide musically salient navigation points in a score. In other cases, a user may simply want to browse annotations. In either case, a notable minority of users would prefer that it not be assumed that the score is the primary navigation point, decorated with annotations – it may be that the annotations provide entrance points for exploring the score.

Commonly-desired visualisations:

- Target indication:
    - **Boxes and brackets** (either horizontally above to span a region, or vertically in pairs to indicate the start and end) to delimit an area
    - **Notation colouring and circles** to indicate target elements
- Body indication:
    - Insertion of **musical notation** – including notes, ornaments, fingering, phrase marks, dynamics, accidentals, etc).
    - **Text** – proximally with the target, or marginally close to it, dynamically displayed, or separately in a table

Recommendations for further user research

1. **The authoring and semantics of hand-drawn annotations in music apps**. It is a clear desire that annotation input should be easier and, preferably, should support stylus or optical input. In these cases, the meaning of the graphical entry needs to be well understood. – better than current publications suggest is currently the case. This could be carried out in partnership with software makers who have experience of purely graphical or hybrid graphical and typeset content.
2. **Effective novel visualisation methods.** It is clear that many current visualisations are limited by physical constraints that could be disregarded in some circumstances in the digital domain. For example, circling an item in a physical score will result in a drawn circle in the final visualisation, where in the digital domain, circling a note could be made to change its colour or draw a shape around it. An imaginative exploration of score visualisations to express annotations would be valuable research.

Development plan for initial integration of annotations into Verovio

Each part of the above can be considered and implemented in stages. Given the preference of users to avoid distorting typesetting logic, initial work should be based on visualisation bypassing the music layout entirely.

1. Since **annot** is comparable to other annotation models, and highly expressive, we recommend starting with this. This would still allow other mechanisms to be used and then serialised into **annot** before employing Verovio. A first step, then, is to ensure that **annot** is read and correctly interpreted.
2. A prerequisite for the target visualisations described above is knowing where the affected regions and symbols are located in the SVG generated by Verovio.
    a. Symbols can be located by providing the drawn elements with an explicit association with a given annotation (this could use @class or a custom data attribute).
    b. For regions it would be necessary for Verovio to provide some way to pass locations to the parent application this could be through a map (in the manner of tempo maps) between annotations and bounding boxes describing their regions
3. A parent annotation can then use Verovio for the information necessary to locate and visualise SVG targets. As a first step, textual annotation bodies can be added based on the location of the target. This will establish a framework for adding further annotation types and visualisations

Bibliography

Jacopo de Berardinis, Albert Meroño-Peñuela, Andrea Poltronieri, Valentina Presutti (2023). The Music Annotation Pattern. *Proceedings of the 13th Workshop on Ontology Design and Patterns (WOP 2022)*.

Chris Cannam, Michael O. Jewell, Christophe Rhodes, Mark Sandler, and Mark d'Inverno (2010) *Linked Data and You: Bringing music research software into the Semantic Web*, Journal of New Music Research, volume 39, no. 4. 313-325

Samira Si-said Cherfi, Christophe Guillotel, Fayçal Hamdi, Philippe Rigaux, and Nicolas Travers (2017) Ontology-based annotation of music scores. *Proceedings of the 9th Knowledge Capture Conference*. 1-4

Richard Freedman (2023) An Introduction to CRIM Essays and Experiments. In *CRIM Essays and Experiments* Haverford College and CRIM Project, https://crimwp.richardfreedman.sites.haverford.edu/?page_id=1126

Mathieu Giraud, Richard Groult, and Emmanuel Leguy (2018) Dezrann, a Web Framework to Share Music Analysis. *Proceedings of the International Conference on Technologies for Music Notation and Representation* (TENOR 2018). 104–11

Werner Goebl and David M Weigl (2022) Alleviating the Last Mile of Encoding: The mei-friend Package for the Atom Text Editor. *Proceedings of the Music Encoding Conference Proceedings 2021*. 31–39

Goebl and Weigl (2024) mei-friend: An Interactive Web-based Editor for Digital Music Encodings. *Journal of Open Source Software* 9(98), 6002,

Eric J. Humphrey, Justin Salamon, Oriol Nieto, Jon Forsyth, Rachel M. Bittner, & Juan Pablo Bello. (2014). JAMS: A JSON Annotated Music Specification for Reproducible MIR Research. *Proceedings of the 15th International Society for Music Information Retrieval Conference*, 591–596.

David Lewis, David Weigl, and Kevin Page (2019). Musicological Observations During Rehearsal and Performance: a Linked Data Digital Library for Annotations. *Proceedings of the 6th International Conference on Digital Libraries for Musicology* . Association for Computing Machinery, New York, NY, USA, 1–8.

David Lewis, Elisabete Shibata, Mark Saccomano, Lisa Rosendahl, Johannes Kepper, Andrew Hankinson, Christine Siegert, and Kevin Page (2022). A model for annotating musical versions and arrangements across multiple documents and media. *Proceedings of the 9th International Conference on Digital Libraries for Musicology (DLfM '22)*. Association for Computing Machinery, New York, NY, USA, 10–18.

David Lewis, Elisabete Shibata, Andrew Hankinson, Johannes Kepper, Kevin Page, Lisa Rosendahl, Mark Saccomano & Christine Siegert (2023). Supporting Musicological Investigations With Information Retrieval Tools: An Iterative Approach to Data Collection. *Proceedings of the 24th International Society for Music Information Retrieval Conference*, 795-801.

Ling Ma, Mathieu Giraud, Emmanuel Leguy (2019) Realtime collaborative annotation of music scores with Dezrann. *International Symposium on Computer Music Modeling and Retrieval (CMMR 2019)*, Marseille, France

Yevgen Mexin, Aristotelis Hadjakos, Axel Berndt, Simon Waloschek, Anastasia Wawilow and Gerd Szwillus (2017) Tools for Annotating Music Measures in Digital Music Editions. *Proceedings of the 14th Sounds and Music Computing Conference*. 279-286

Matthias Nowakowski and Aristotelis Hadjakos (2022) Online Survey on Usability and User Experience of Music Notation Editors. *Proceedings of the International Conference on Technologies for Music Notation and Representation* (TENOR 2022). 6-15

Alexander Pacha and Klaus Rettinghaus (2022) Building a Comprehensive Sheet Music Library Application. *Proceedings of the Music Encoding Conference* (MEC 2021) Alicante, Spain

Anna Plaksin. 2023. Understanding the needs of music editors in a digital world. Adding support for editorial markup to the mei-friend editor. *Proceedings of the 10th International Conference on Digital Libraries for Musicology* (DLfM '23) 40–48

Robert Sanderson, Paolo Ciccarese, and Benjamin Young (2017). *Web Annotation Data Model*. W3C Recommendation. W3C. https://www.w3.org/TR/2017/REC-annotation-model-20170223/.

Stokes, L. K. T. (2024). Nkoda. *Music Reference Services Quarterly*, *27*(2), 91–95.

Raffaele Viglianti (2016). The Music Addressability API: A Draft Specification for Addressing Portions of Music Notation on the Web. *Proceedings of the 3rd International Workshop on Digital Libraries for Musicology* (New York, USA) (DLfM 2016). Association for Computing Machinery, New York, NY, USA, 57–60.

David M. Weigl, Werner Goebl, David J. Baker, Tim Crawford, Federico Zubani, Aggelos Gkiokas, Nicolas F. Gutierrez, Alastair Porter, and Patricia Santos. (2021). Notes on the Music: A social data infrastructure for music annotation. In *Proceedings of the 8th International Conference on Digital Libraries for Musicology (DLfM '21)*. ACM, 23–31

Christof Weiß, Frank Zalkow, Vlora Arifi-Müller, Meinard Müller, Hendrik Vincent Koops, Anja Volk, & Harald G. Grohganz. (2024). Schubert Winterreise Dataset [Data set]. In *ACM Journal on Computing and Cultural Heritage* (2.1, Vol. 14, Number 2, pp. 25:1–18).

Appendices

Appendix A: List of software packages and libraries

Generalised graphical		
PDF	**PDF readers**	
	Sometimes Adobe Acrobat, sometimes unspecified (so could also be Preview, etc.). Generally used for passive reading of scores (exported from engraving packages, scanned or downloaded from sites such as IMSLP), but most support some level of annotation.	
GN	**GoodNotes**	**https://www.goodnotes.com.**
	A cross-platform (iOS, Android, Windows) note-taking application that supports annotating PDFs.	
PPT	**Powerpoint**	
	Intuitive page-based drawing with the ability to load in images and overlay shapes and texts.	
Score-optimised graphical		
FS	**ForScore**	**https://forScore.co**
NK	**Nkoda**	**https://nkoda.com**.
	Digital sheet music library and application. Supports annotation, highlighting and drawing (Stokes, 2024)	
Hen	**Henle**	**https://henle-library.com**
	Sheet music library with annotation and fingering overlay (including loading from various pre-existing fingerings)	
EN	**Enote**	**https://enote.com**
	A hybrid system, with graphical and MEI components (e.g. Rettinghaus & Pacha, 2022)	
Music engraving[13]		
Fin	**Finale**	**https://www.finalemusic.com**
Dor	**Dorico**	**https://www.steinberg.net/dorico**
Sib	**Sibelius**	**https://www.sibelius.com**
Mus	**MuseScore**	**https://musescore.org**

[13] About music engravers see, for example, Nowakowski and Hadjakos (2022)

Specialised

Edi **Edirom** **https://www.edirom.de**

Suite of tools developed by the Virtuelle Forschungsverbund Edirom (ViFE), covering a range of editorial and musicological activities, primarily centred around musical sources (see, for example, Mexin et al, 2017)

MF **mei-friend** **https://mei-friend.mdw.ac.at**

In-browser MEI code editor with live preview and many features to support the editing process (see Goebl and Weigl, 2022)

Dez **Dezrann** **https://algomus.fr/dezrann**

In-browser music exploration and annotation library, with a range of applications, and supporting video, audio and notation (Giraud et al, 2018; Ma et al, 2018)

MELD **https://meld.web.ox.ac.uk**

In-browser (and server) library for linked-data based music applications, with much use of web annotations. Supports time-based media, MEI and TEI (see Lewis et al, 2023)