

EuterPen: Unleashing Creative Expression in Music Score Writing

Vincent Cavez
LISN

Université Paris-Saclay, CNRS, Inria
Gif-sur-Yvette, France
vincent.cavez@inria.fr

Caroline Appert
LISN

Université Paris-Saclay, CNRS, Inria
Orsay, France
caroline.appert@universite-paris-saclay.fr

Catherine Letondal
Interactive Informatics

University of Toulouse - ENAC
Toulouse, France
catherine.letondal@enac.fr

Emmanuel Pietriga
LISN

Université Paris-Saclay, CNRS, Inria
Gif-sur-Yvette, France
emmanuel.pietriga@inria.fr

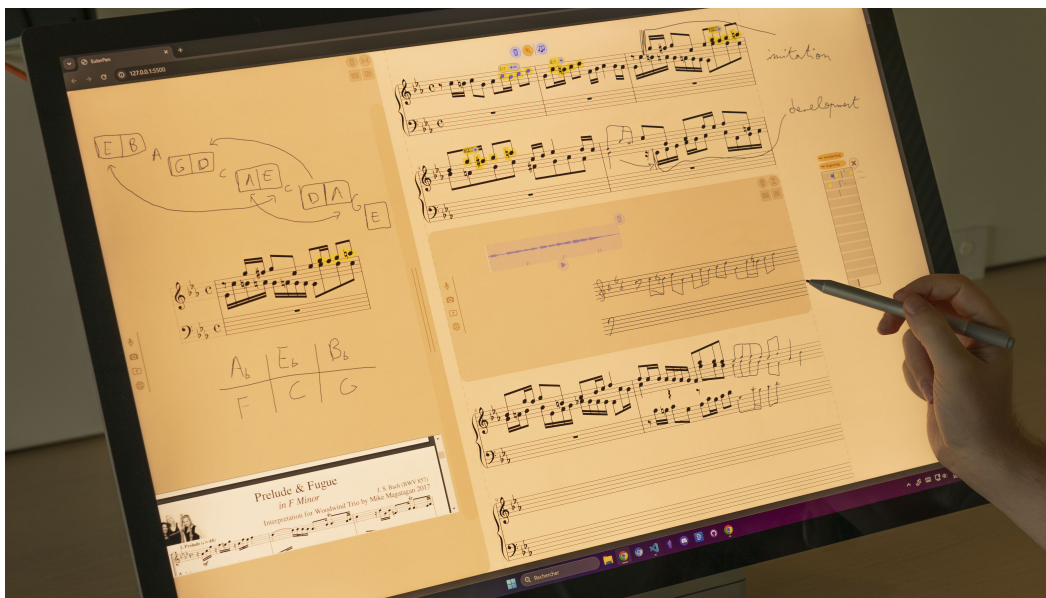


Figure 1: EuterPen running on an interactive surface. Areas around staves (called *pensieves*) are used to freely explore musical ideas while keeping the main score stable, and to store content that supports composers' creative process: documents, audio samples, diagrams. Handwritten and engraved music notations coexist and are both amenable to structured-yet-flexible interactive manipulations, including a tool to search for constrained or relaxed melodic patterns across the whole workspace.

Abstract

Music notation programs force composers to follow the many rules of the staff notation when writing music and constantly seek to optimize symbol placement, making numerous adjustments automatically. Even though this impedes their creative process, many composers still use them throughout their workflow, for lack of a better option. We introduce EuterPen, a music notation program prototype that selectively relaxes both syntactic and structural constraints while editing a score. Composers can input and manipulate

music symbols with increased flexibility, leveraging the affordances of pen and touch. They can make space on, between and around staves to insert additional content such as digital ink, pictures and audio samples. We describe the iterative design process that led to EuterPen: prototyping phases, a participatory design workshop, and a series of interviews. Feedback from the participating professional composers indicates that EuterPen offers a compelling and promising approach to music writing.

CCS Concepts

• **Human-centered computing** → **Interactive systems and tools; Graphics input devices; Touch screens.**

Keywords

digital pen, multi-touch interaction, ink, music engraving



This work is licensed under a Creative Commons Attribution 4.0 International License.
CHI '25, Yokohama, Japan

© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1394-1/25/04
<https://doi.org/10.1145/3706598.3713488>

ACM Reference Format:

Vincent Cavez, Catherine Letondal, Caroline Appert, and Emmanuel Pietriga. 2025. EuterPen: Unleashing Creative Expression in Music Score Writing. In *CHI Conference on Human Factors in Computing Systems (CHI '25), April 26–May 01, 2025, Yokohama, Japan*. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3706598.3713488>

1 Introduction

Music notation programs, or scorewriters, are software tools that assist composers in the creation of scores through typesetting music symbols on digitally rendered pages. They ensure that the syntactic rules of the modern staff notation are properly followed, and automatically optimize symbol layout to achieve good readability. As such these programs are structured editors that effectively support the process of music *engraving*, which refers to the drawing of high-quality music notation.

Composers use music notation programs not only for engraving or playing back their score, however. They use them throughout their creative process, to capture early ideas, to rework them and test alternatives, to arrange them together. These creative activities – rather exploratory and iterative in nature – call for flexibility in the interactive manipulation of the score. But as the word suggests, *engraved* music notation is not easy to modify, even in a digital format. This raises a tension between, on one hand, the need for freedom to capture fleeting thoughts, to explore different ideas, to iterate on them, to import foreign material; and, on the other hand, the need for structure and well-formedness of the notation.

In a recent study, Cavez *et al.* [12] consider this tension between flexibility and structure, and discuss what opportunities interactive surfaces – operated with pen and touch input – can bring in terms of interaction design to alleviate it. Informed by these findings, we developed EuterPen, a prototype music notation program designed specifically for interactive surfaces. EuterPen, illustrated in Figure 1, adopts a semi-structured editing approach, giving composers more freedom to manipulate the score by breaking away from some constraints of the notation and by letting composers make space within the score to explore musical ideas while preserving the main structure. Composers can interactively carve space on, around and between staves. In this extra space, composers can place and manipulate different types of content, such as drawings, audio recordings, images, foreign scores, ink annotations or empty staves to experiment with ideas. Designed for pen and touch input [22], EuterPen enables composers to intermingle both handwritten and engraved representations of the music notation. It offers a user experience that is closer to that of writing music on paper while still benefiting from the search, edit and beautification capabilities of a structured document editor. Other content types can be seamlessly integrated with staves, aiming to maintain stability of the score and preserve composers' fine adjustments – considering both global organization (bar layout) and local details (individual symbol layout).

We first motivate our work by discussing the affordances [18] of pen input and the limitations of existing approaches to music score writing. We then describe our methodology, which consisted of design and software prototyping phases, a workshop, and interview sessions with professional composers. We then introduce EuterPen, organizing the discussion according to how we break away from

the rigidity of regular music notation programs, providing composers with advanced editing capabilities based primarily on direct manipulation with a combination of pen and touch input.

2 Background and Motivation

Previous studies have shown that the creative process varies significantly from one composer to another [15, 38, 49], but that for many of them it is highly iterative [5, 6, 14]. Composers use music notation programs throughout this iterative process, including the more exploratory phases. Yet those programs tend to focus on the well-formedness of compositions, enforcing syntactic rules at every step and automatically laying out the notation to optimize legibility. This is achieved by imposing many constraints on how composers input and modify the notation. The outcome is a music score of high quality, but obtained at the cost of raising multiple usability issues [12, 31], including:

- a major loss of flexibility that adversely impacts composers' freedom to experiment and jot down ideas or even insert annotations and custom symbols;
- the mismatch between a highly non-linear creative process and tools that impose a linear input flow and suffer from much viscosity [19] when editing the notation;
- significant time spent correcting issues in the notation due to the program making wrong guesses;
- a loss of control on notation layout, both at the global level (measures and staves) and local level (notes and rests), which again causes composers to waste a lot of time adjusting elements;
- advanced notation manipulation features exposed through bloated UI panels, that require learning numerous keyboard shortcuts or navigating complex menus.

This often leads composers to use both music notation programs *and* paper to support their creative process [25, 48]. Mixing these two different media causes “*abrupt shifts*” though [16], which in turn may lead composers to keep editing on the computer even when paper could be a more appropriate medium for the task at hand.

The literature on pen + touch input (see, *e.g.*, [11, 22, 41]) suggests that interactive surfaces – both handheld such as, *e.g.*, the Apple iPad, and desktop such as, *e.g.*, the Microsoft Surface Studio – can effectively merge the advantages of both paper and digital media. They thus have the potential to resolve the tension between structure and flexibility. Some recent music notation programs such as StaffPad [44] and Symphony Pro [35] have taken a first step in this direction, but still primarily follow the same WIMP-oriented interaction model as desktop applications. They support pen and touch input, but handle those as undifferentiated generic pointers, failing to leverage their unique affordances and expressive power. The pen's potential remains largely untapped, in particular. We detail this potential, and then we give an overview of how pen input is actually used both in research projects and commercial music notation software.

2.1 Properties of Pen Input

Pen input has been studied in the HCI literature both as a generic modality and as a device adapted to specific tasks, demonstrating

interesting properties on its own and when combined with touch input. Several of these properties are of particular interest in the context of music score writing.

- *Digital pen input is both versatile and precise*: a pen can be used to select small graphical elements, to draw arbitrary shapes, and to write text. It is an excellent tool to make annotations (writing, sketching) in dense environments [41, 46, 53], to make freeform selections of sets of small items (there is much less occlusion than with touch input [37, 40]); and to perform direct manipulations in 2D canvases hosting rich content [39, 50]. In the context of a music notation program, such capabilities can enable elaborate selections and direct manipulations of the primary notation's symbols, and the input of secondary notations such as annotations [12] – all with a single input device.
- *Pen input, combined with touch, can have high expressive power*: compared to interactive systems that only feature a single type of pointer, interactive surfaces that discriminate between pen and finger touch make it possible to implement a division-of-labor strategy [22] that maps a broader set of direct manipulations to individual modalities, sometimes using one to parameterize actions performed simultaneously with the other [11, 27, 34]. Music symbols can be modified in a variety of ways, that typically require selecting different tools from a palette or invoking commands in a menu. Enabling a broader set of direct manipulations this way can streamline the editing process.
- *Digital pens are well-suited to handwritten input*: such input is not limited to raw ink but can have structure and can be made editable [43]. Beyond notes and sketches, a digital pen has the potential to let composers input music symbols by handwriting as easily as on paper, provided it is properly articulated with the music symbol recognizer that will eventually parse and beautify the handwriting.
- *Pen input can leave actionable ink traces in the workspace*: those traces can be temporary or permanent. They can be used to invoke commands [2], that can be triggered immediately or interpreted later [40, 48] – which can help avoid some cases of *premature commitment* [19]. Those traces can also be used to generate new content [20, 43, 46, 50], sometimes coupled with an ink recognizer. This can improve *closeness of mapping* [19] compared to other input methods such as, e.g., keyboard shortcuts or tool palettes, and has the potential to make music symbol input much more effective.
- *Writing with a pen favors the externalization of thoughts*: capture fleeting ideas, take notes [26, 37]. This has cognitive benefits such as lowering working memory load and facilitating idea reformulation. It is particularly important for sensemaking and data exploration tasks [28, 40, 42], but it is also key to many creative processes.

2.2 Actual Use of Pen Input

Two main strategies have been explored for music notation input with a pen: recognition of handwritten strokes; or recognition of simple predefined gestures from a controlled vocabulary. Leroy *et al.*'s system [24] was based on an early form of optical music recognition [10] that would let composers input music symbols on staves by drawing them. With Presto [1, 32], composers

would input symbols using more abstract, predefined gestures and could then perform some manipulations on them using command menus: transposing, moving elements, adding ornaments. Those first systems paved the way for future ones but were limited by the capabilities of early hardware. Current hardware has significantly more advanced capabilities, providing an opportunity to effectively use pen input and to combine it with touch to broaden music score writing beyond rudimentary symbol input, and to streamline the iterative editing process of the notation.

More recently, another line of research [16, 17, 48] has investigated a particular type of ink pen (Anoto digital pen) equipped with a digital camera to record what gets written on special sheets of so-called interactive paper [21]. Marks drawn on interactive paper are actual ink. They cannot be modified, moved or deleted, making iterations on the notation itself impossible. Targeting contemporary composers, the aim of this research was rather to enable the creation of music pieces by combining programming tools [8, 36], custom notations and score fragments using gesture-based interaction.

On the commercial side, we find music notation programs originally written for the desktop that can operate on interactive surfaces, such as Sibelius, Dorico, or Flat. These are mostly straightforward adaptations of their desktop counterparts: composers write music symbols by first selecting the proper symbol type in a palette (virtual keypad) and then clicking on staves to instantiate this symbol. They make very few adjustments to the desktop version's WIMP model [12], and treat pen and touch interchangeably for the most part. They rather promote similarity between the desktop and mobile versions of a program. They thus miss an opportunity to leverage the affordances and expressive power of pen and touch combined, and keep relying on menus, tool palettes and mode switches to give composers access to a wealth of features.

Only a few music notation programs are designed specifically for interactive surfaces such as the Apple iPad or the Microsoft Surface Pro. These include StaffPad [44], Symphony Pro [35] and Notion Mobile [33]. These truly leverage the pen, effectively supporting handwritten input of many music symbols and making the music writing experience closer to what it would be on paper. But these programs still impose significant constraints on how composers may input the notation, constraints that are related to the underlying recognition process. In order for the recognition engine to properly parse their strokes, composers must pay more attention to how they draw symbols than they typically would when writing freely on paper. This concerns not only the shape of head, stem and beams, but the order in which they are input as well. A constraint that perhaps more deeply impacts composers' creative process is related to the temporal articulation between symbol drawing by the user and recognition by the engine. Recognition is typically implemented as a greedy process: as soon as the parent measure is complete, the system automatically beautifies inked symbols and checks that the measure is well-formed: a measure must be filled correctly before proceeding to the next. And once it has been beautified, the notation is difficult to edit: as in desktop music notation programs, syntactic and structural constraints must be enforced. In a way, the notation has indeed been *engraved* (digitally speaking). Modifications can be made, but with much viscosity [19].

Taken together, these constraints impose a specific way of writing on composers that can adversely impact their creative process

as it limits their capacity to experiment, to jot ideas down [31], and to iterate on the beautified symbols. The user experience is closer to writing on paper, but the freedom inherent to it is lost. EuterPen aims to offer composers the best of both worlds: reclaim part of that freedom while keeping enough structure to enable advanced manipulations of the music notation; and reclaim another part of that freedom to let composers externalize their ideas inside the score while keeping the main structure stable.

3 Methodology

The development of EuterPen followed an iterative design process that consisted of four main phases: a first design and prototyping phase informed by previous studies; a workshop involving two professional composers; a second design and prototyping phase informed by this workshop; and a series of interviews with eight composers.

3.1 Design Journey

The first design and prototyping phase was driven by one high-level goal: to reveal the full potential of interactive surfaces for music score writing – prior work in the HCI literature suggesting that pen and touch input have the potential to resolve tensions between structure and flexibility. This phase (4 months) was primarily informed by the findings from Cavez *et al.* [12], who conducted interviews with nine composers about their creative process and how music notation programs support it.

Once developed, the first EuterPen prototype was used to illustrate a variety of possible interactions to two professional composers who accepted to participate in a half-day face-to-face workshop together with the authors. We believe it was important to make a high-fidelity prototype (see Section 3.2) available to participants even if we were still early in the design process, as composers are not necessarily very familiar with interactive surfaces. We wanted composers to be able to experience the new possibilities offered by this technology first-hand, enabling them to actually manipulate music notation with pen and touch. Discussions during the workshop were structured along different themes. First, the lead author gave an overview of EuterPen. Then, for each theme, a series of features and interaction techniques were showcased to the composers using video recordings, leading to discussions about the relevance of the feature and how it was mapped to interactions. EuterPen was running on a Microsoft Surface Studio 2 in the workshop room, and composers were encouraged to play with the system and try the features whenever they wanted. We made it clear to composers that despite the high fidelity of our prototype we were still at an early design stage, and that they should feel free to question and rethink what they were shown and propose new ideas during these thematic discussions.

As effective interaction design often lies in the careful attention to detail, the workshop then included a 1-hour brainstorming session focused on a specific operation to gather fine-grained insights. We chose copy-and-paste, an operation that is central to managing repetitions and variations in the music composition process, but often perceived by composers as tedious [12]. To avoid influencing composers during brainstorming, we did not present any solutions, but rather asked them to propose alternatives and improvements to

the conventional COPY and PASTE buttons. Workshop participants were invited to first produce around 10 ideas each, on separate post-it notes. After 15 minutes the resulting ideas were shared orally or by demonstrating how this would work in front of the EuterPen prototype.

The entire workshop was recorded (audio and video) and later transcribed by the lead author. A thematic coding [45] of those transcripts and recordings was performed by the first two authors.

Following this workshop, we entered a second design and prototyping phase (4 months) during which we developed new features and iterated upon existing ones based on the feedback from the two composers and new ideas that arose from the discussion. This led to the second version of EuterPen, which we used to validate our approach through interviews with a broader set of professional composers. Composers, as many creative professionals, are typically very busy, and we chose to conduct interviews online so that they could participate without traveling to our lab. We conducted 8 interviews, all performed using a videoconference tool. An interview lasted between 45 minutes and 2 hours 30 minutes (1 hour 40 minutes on average, more than 13 hours total). Information about the composers' profile is available in Table 1. As we did for the workshop, the key features and interactions of EuterPen were recorded and grouped thematically into coherent feature sets. Those were made available to the interviewees on the Web, embedding the videos into a Web page together with textual explanations. That Web page was shared with the interviewees several days before the actual interview took place, so that they could familiarize themselves with the different features. The Web page was actually a form, that they could use to record early comments and ratings about individual features at their convenience. During the interview, this early feedback was revisited as part of the discussion between the interviewer and the composer. The lead author conducted the interviews, and the first two authors performed the transcription and thematic coding of those recordings. Ratings are reported by thematic group in Figures 3, 5 and 8.

Finally, we entered a third, shorter design and prototyping phase (1 month) to make final improvements to the prototype. These different stages constitute EuterPen's design journey, which we discuss in Sections 4.1–4.3, emphasizing the main insights from the different stages of the process.

3.2 Software Implementation

EuterPen is developed as a Web application and can run in any browser that supports the W3C Pointer Events API [9]. It uses the VexFlow music notation API [13] to render scores as SVG elements [4], that can be imported from MusicXML documents with OpenSheetMusic [47]. The programmatic manipulation of the SVG-rendered music notation elements and all UI components is coded in JavaScript with D3 [7]. Handwritten music notation recognition is handled via a custom version of MyScript [30] configured with a specific grammar. This recognition service runs locally as a Java servlet that EuterPen queries to interpret digital-ink input.

EuterPen has been implemented to experiment with, and demonstrate, novel ways to support composers in their creative process. As such, we did not aim to develop a full-featured score editor but rather a functional prototype that can import music scores and let

Interviewee	Initials	Activities(*)	Main Style	Experience (y)	Current Software	Main Writing Tool(s)
Lucius Arkmann	LA	C, A, E, T	Modern classical	> 15	Musescore	Laptop or paper
Caroline Itier	CI	C, A, T	Jazz	> 15	Finale	Laptop or paper
Bruno Fages	BF	C, A, T	Jazz	> 30	Sibelius	Laptop or paper
<i>Anonymous</i>	AN	C, T	Modern classical	≈ 20	Sibelius	Laptop and paper
Alexandre Olech	AO	C, A, E	Modern classical	> 5	Sibelius, Dorico	Desktop
Yves Torchinsky	YT	C, A, T	Jazz, Rock	> 40	Sibelius	Laptop and paper
Olivier Sabatier	OS	C, A, T	Modern classical	> 20	Sibelius	Laptop and paper
Dominique Pifarély	DP	C	Jazz	> 45	Sibelius	Desktop, laptop and paper

Table 1: Profile of the eight composers we worked with. *Activities: C=Composition; A=Arrangement; E=Engraving for another composer; T=Teaching. All of them participated in the final interviews and two of them also participated in the design workshop (we refer to them as C₁ and C₂ for anonymity concerns).

users modify them or write music from scratch. Section 4 discusses those novel ways to support composers, focusing on novelty and design choices rather than describing every single feature.

4 Design

The design of EuterPen has been informed by two main sources: the participatory design workshop & interviews conducted as part of the present work (Section 3); and the findings from Cavez *et al.*'s recent study [12]. Those latter findings yield the following high-level design guidelines:

- the music notation program should give composers more freedom to lay elements out in the workspace, breaking away from the linear, unidimensional flow of staves: make space in and around the score to capture ideas and explore them freely, writing new content to be stored and used later, or extracting existing passages from the main flow;
- those spaces should not be limited to the primary music notation but should also let composers capture their ideas by using secondary notations and by importing foreign content inside the score such as audio samples or even pictures and videos;
- and finally, the program should let composers select and manipulate score elements at different granularities, from whole systems and measures to the individual components of a single note (head, stem and beam).

The participatory design workshop and evaluation interviews have revealed two additional, complementary principles: stability and flexibility. On one hand, composers find it important that the editing decisions they make are not overridden by the music notation program. For instance, inserting bars, adding notes or annotations can have dire consequences on the global organization of the score. But, on the other hand, they find it as important that the program allows them to explore ideas freely without disturbing what is already settled.

How EuterPen follows the above set of guidelines and principles is discussed in the next three subsections, organized according to three themes: breaking the score's linear structure (Section 4.1); breaking the score's homogeneity (Section 4.2); and breaking down musical elements (Section 4.3). In each subsection we also explain where the main ideas originated from and report on the feedback we gathered from composers.

A design scenario (see Scenario 1 below) illustrates how composers can benefit from those guidelines and principles as implemented in EuterPen. This scenario is based on the interviews conducted during the evaluation phase. It is built using composers' feedback about EuterPen interactions – feedback that is rooted in their actual work practices.

4.1 Breaking the Score's Linear Structure

A musical piece may unfold temporally, but the creative process that yields it is most often non-linear. This is a source of tension as: on one hand, the notation program should help composers comply with the rules governing the piece's temporal – inherently linear – structure; but on the other hand, it should let them freely explore and arrange different ideas by supporting a very flexible, non-linear editing workflow.

Existing music notation programs are clearly focused on enforcing temporal-structure well-formedness, providing composers with a workspace organized strictly as a linear sequence of bars forming staves or systems of staves that all abide by the declared time signature. The constraints imposed by those programs do not apply to the primary music notation only, but prevent composers from inserting other pieces of information that are key to the composition process: text annotations or foreign objects such as images that are not meant to stay but support the creative process. Composers have to cope with these constraints, when an exploratory process rather calls for those constraints to be lifted, if only temporarily. When possible, composers will cheat the program: for instance, adding dummy bars at the end of a score to test ideas that are not meant to be part of the final piece. However, many constraints cannot be circumvented, leading composers to resort to other means to capture their ideas, such as pen and paper. This has a cost though, as they will then have to juggle completely disconnected media and workspaces.

We address this global concern by designing EuterPen not as a regular *score editor*, but rather as a *music processor*.¹ While EuterPen has composers write music notation on staves as any regular score editor does, composers can interactively carve spaces between staves, that can hold a variety of contents. Two spaces coexist:

¹A term coined by LA during their interview, drawing an analogy with *word processors* that are more flexible than structured document editors and enable authors to mix different types of contents.

Scenario 1 – A Composer’s Journey – illustrates EuterPen’s features using examples of composers’ actual practices gathered from the evaluation interviews. A (*) indicates suggestions originating from composers who participated in the interviews, but not yet implemented in the EuterPen prototype.

Chris has started to think of a piece in the fugue form inspired by J.S. Bach, that he has to work on for his composition class. While he is trying some rhythmic patterns on the piano, he decides to record a few audio samples on his smartphone.

A few days later, he decides to create an EuterPen score to explore and shape his piece. He opts for a 4/4 E-b signature and writes the first two bars with the pen. To check that it sounds as he imagined, he makes the play gesture with the pen to hear those bars (see Figure 3-c and Figure 6-c) and is happy with the result. Next Chris wants to insert one of the samples recorded on his smartphone a few days ago. He opens the inter-staff *pensieve* (as in Figure 3-c) and loads them there (similar to Figure 2.c). One of those samples fits very well as a continuation of the first two handwritten bars, and he drags-and-drops it on the staff (e.g. Figure 5-b). He can then listen to the whole content, EuterPen seamlessly playing handwritten and engraved notation, as well as audio samples that have been inserted on the staves. Chris then writes a few words above one of the bars to remind him about an idea he had while exploring ideas on the piano, which was about combining background sea sounds with the pattern (he quickly sketches a diagram representing this idea using the pen, as in Figure 1). He also loads one of J.S. Bach’s fugues as a PDF in the global *pensieve* on the left of the staves (see also Figure 1).

Several days later, Chris has completed the structure of his piece in three parts, partly filled with two voices that he has already engraved. Using EuterPen, his teacher annotates the score using the pen (similar to Figure 5-c) and makes suggestions such as adding a central part aimed at developing the first theme. She also confirms that Chris could use the 1st voice as a model for the 3rd one.

Back home, Chris selects the 1st voice (as in Figure 6-a) over 3 bars and duplicates it before moving it to the 3rd voice. But he has to adjust a few notes to comply with the rules of harmony: using the pen he selects the notes, changes their height, and slightly adjusts their horizontal position so that the score remains legible (similar to Figure 8-c). Chris then proceeds to the central development that the teacher suggested. He first explores variations on the patterns using the inter-staff *pensieve*, where he can freely instantiate short staves and play them separately (Figure 3-a). He then discovers a better rhythmic form that would make sense for the whole piece: he selects the elements involved in this pattern (similar to Figure 8-a) and issues a search over the entire score – not just notes on staves but *pensieves* and annotations as well (*). This highlights all pattern occurrences throughout the workspace. Those highlights are also visible on a minimap that shows a global view of the document (see Figure 7). Chris puts all occurrences back in handwritten form and edits them all at once.

Going back to the central development, Chris wants to check Bach’s original score that he imported as a PDF, and copy a fragment of it in the *pensieve* between two staves – close to the part of his own score that echoes this passage. Copy-&-pasting this fragment from a foreign source automatically creates a hyperlink back to the original material (*). Such *provenance* information will prove useful to him when he revisits his composition in the future. Finally, Chris inserts 30 bars at once to prepare the development section, using a gesture similar to Figure 6-c. He fills those by copy-&-pasting the notes drafted in the inter-staff *pensieve*. This leaves several bars incomplete (as in Figure 3-a) but that is fine. Happy with this step, Chris feels like he now needs some time to reflect on this first draft. He closes EuterPen, confident that he will be able to quickly find the places that still need some work in his score as he has left them in handwritten form whereas the finalized bars have been engraved.



Figure 2: Carving a pensieve by (a-b) stretching space between two staves using a simple multitouch gesture, and then (c) populating that space with multiple objects using drag-&-drop. An audio sample has been dropped already, and the user is about to place an empty staff to write music on.

- the regular **temporal composition space** with a linear flow of staves, as found in all music notation programs;
- and canvases, named **pensieves**,² that composers can instantiate and populate freely with a variety of contents.

Pensieves can be used to store and retrieve any material relevant to the creative process (handwritten text annotations, drawings, pictures, audio samples and video clips). They can also be used to experiment with musical ideas, instantiating blank staves to write or copy music notation on – see Section 4.2. As such, pensieves

²Following the suggestion of a workshop participant, who drew an analogy with the thought spaces described by character Dumbledore in the *Harry Potter* books series.

exist in-between and around staves, and can be opened and closed at will.

Carving a pensieve between two staves in the temporal composition space involves holding one staff with a finger, and pulling the other staff with two fingers. This type of pensieve remains tightly integrated with the temporal composition space and scrolls along with it. A global pensieve can be shown to the side of the score as well by dragging it like a drawer. That global pensieve spans the full height of the EuterPen workspace and can be scrolled independently from the staves of the temporal composition space.

Pensieves let composers organize content *around* staves. But EuterPen also aims to give them more flexibility *inside* the staves themselves. The systematic enforcement of temporal constraints in bars represents a major source of frustration and impedes the creative process [12]. EuterPen addresses this key issue by relaxing these constraints. Staves can hold a mix of structured (engraved) spans and unstructured ones consisting of arbitrarily-long sequences of notes and rests. When writing music, composers can – but are not *required to* – draw bar lines, and then engrave or delete them at will. EuterPen also makes insertion more flexible in order to better support the creative process, which is highly non-linear, as mentioned before. A caret gesture [29] performed with the pen (see Figure 6-c) will push existing notes to make space that can then be used to write new notes. As the layout of bars on staves is something that composers consider important, the insertion strategy can be parameterized on-the-fly: adjust the amount of space depending on the (log-transformed) length of the caret gesture; after lifting the pen, choose to add space before (push content upstream) or after (push content downstream) the insertion point; and most importantly choose whether to automatically generate bar lines (which helps remain within a given duration) or not (which leaves the duration unconstrained).

EuterPen lets composers make further layout adjustments to engraved content, leveraging the structural information obtained in the process: spreading or packing notes with a two-finger pinch gesture, adding space to a bar by dragging the bar line with the pen. EuterPen does not attempt to optimize layout automatically as other score editors do, but rather gives composers the possibility to resize bars according to music notation rules on demand.

Beyond staves and pensieves, which constitute two distinct spaces with different purposes, EuterPen also lets composers make annotations anywhere in the workspace – in all areas, both temporal composition space and pensieves. Composers create those annotations seamlessly without the need to switch modes, as explained in Section 4.2.

Feedback

As illustrated in Figure 3, interviewees were very enthusiastic about EuterPen’s way of breaking the score’s linearity: canvases for free-form work and the flexibility in handling bars were unanimously judged to be useful (Q1). Interviewees did not think these exist in their current music notation programs, except for the insertion and manipulation of bars, which two composers felt they can already perform as they wish (Q2). Composers found the demonstrated interactions easy to use with the exception of one composer who

found the caret gesture to insert measures potentially difficult to perform (Q3).

The idea of providing a blank canvas for composers to freely arrange and edit measures outside the score stemmed from the opportunities identified by Cavez *et al.* [12] and eventually grew into pensieves during the workshop. As C₂ put it: “*it would be great to be able to copy, how shall I put it, copy and paste, but outside our score*”, C₁ then suggesting this could be a “*reserve of ideas, pensieve or clipboard*”. The first EuterPen prototype that C₁ and C₂ saw during the workshop only allowed carving pensieves between staves. The concept was thus expanded to include the global pensieve (on the left in Figure 1), which is not tied to a particular place in the score but rather remains fixed on screen, serving as a global storage and annotation space.


The relevance of pensieves was further confirmed during interviews: “*And yes, it leaves a little bit of.. a little bit of a laboratory in the middle of the score*” (BF). It is not just about “*having everything already on screen rather than travelling around in windows*,” but also about having space, “*a bit like having cards and moving them around*” (LA) instead of having to create extra bars, “*which is hell, because then you have to take them away*” (AO). Composers found multiple uses for pensieves. They can be simple temporary work areas, for example to “*switch the left and right hands*” (LA). Composers often saw them as a space to note comments, draw, but also to explore musical ideas, for example to “*try out fifteen different rhythms*” (LA). Staves created in pensieves (e.g. Figure 1 or Figure 2-c) are fully functional: they can be interpreted, played back, and “*re-injected into the score at a later time*” (DP), a feature that LA was enthusiastic about: “*It’s really incredible [...] adding staves on the fly underneath. Wow. Fantastic [...] You draw five lines by hand to make the poor man’s staff, and now the computer does it itself. It’s much more practical.*”

Pensieves could even be used to build a structured library of ideas: “*I can imagine this as a sort of idea box where you can have places where you have a zone with rhythmic ideas, a zone with ideas for a melodic sequence and you can take one or the other*” (C2). This notion of a library extends beyond music symbols on staves to other creative resources: musical references or sources of inspiration (as in Figure 1), as C1 described during the workshop: “*You can get the whole idea bit by bit. You don’t have time to go into all the details of the thing, but then you have to note down as many resources as possible. If, for example, I want to take inspiration from Steve Reich, then I’ll put Reich’s score on the page so I don’t have to go looking for it.*” It can also be a teaching tool, as AO points out: “*There were times when my teacher said to me: ‘Yes, but that, you have to look at Debussy, look at LA MER by Debussy. Look at all the string models’*”. Or as OS says: “*to make exercise sheets perhaps.*” He sees pensieves as a tool to “*import whatever you want. Images, text. Links to pages like that, [...] stories about music or anything else that might relate to the study of the orchestra, for example instruments in the orchestra.*”

The greater flexibility obtained from the relaxation of structural constraints within measures was also received quite positively. Composers can experiment freely with melody and rhythm without having to worry about the notation’s formal correctness, which fosters creativity: “*There’s no more alteration, there’s no more rhythmic signature in fact? You can do what you like, you’re free, aren’t you? That’s great. The example you did on the side, can you add beats,*

- Q1: Do the actions shown in this video correspond to actions you would like to do when you compose? ■ Yes ■ More or less ■ No
- Q2: Are you able to carry out this type of action with the tools currently at your disposal? ■ Yes ■ More or less ■ No
- Q3: Do you find what is shown in the video difficult to use? ■ Very easy ■ Easy ■ Neither difficult nor easy ■ Difficult ■ Very difficult

a) Work in a Free Space



Q1 ■ ■ ■ ■ ■ ■ ■ ■ ■ ■

Q2 ■ ■ ■ ■ ■ ■ ■ ■ ■ ■

Q3 ■ ■ ■ ■ ■ ■ ■ ■ ■ ■

b) Localization of Thinking Spaces




Q1 ■ ■ ■ ■ ■ ■ ■ ■ ■ ■

Q2 ■ ■ ■ ■ ■ ■ ■ ■ ■ ■

Q3 ■ ■ ■ ■ ■ ■ ■ ■ ■ ■

c) Insertion and Manipulation of Bars



Q1 ■ ■ ■ ■ ■ ■ ■ ■ ■ ■

Q2 ■ ■ ■ ■ ■ ■ ■ ■ ■ ■

Q3 ■ ■ ■ ■ ■ ■ ■ ■ ■ ■

Answers to Q2 color-coded red or yellow are to be interpreted in favor of EuterPen, as they indicate operations impossible or impractical in the score editors that interviewees currently use.

Figure 3: Breaking the score's linear structure: participants' evaluation of the features demonstrated in the videos shown during the interviews.

can you do a nine-beat bar like that, even without numbering it?" (YT), echoed by (BP) who observed that "the bar we're going to insert can be of any rhythmic signature" suggesting he would rather "define afterwards whether it's identical to the preceding bar or whether it's a completely different rhythmic signature." BF further added: "There may be an irregularity at some point in the beat, a change of time signature, for example, or a change at the end of a phrase that runs counter to the beat", which may lead to "the possibility of adding an extra beat to a bar". Similarly, composers commented positively on the management of empty space in bars. While other programs automatically insert rests to keep a bar well-formed, EuterPen leaves composers free to write down their ideas unimpeded: "So typically, when you've deleted the two beats, it hasn't put anything in their place. I like that, that there's nothing there" (LA).

Composers also appreciated not having to worry about bar layout at the start of the writing process. As AO puts it: "when I'm composing, I'm spending time doing something and I, I shouldn't be spending time doing [layout]. In fact, ideally, I should be composing." echoing observations from Bennett [5]. Layout stability – avoiding automatic optimizations – was also considered important: "Well yes, that's very practical because everything stays there. And then you do the layout again, but after you've finished the whole passage and not while you're adding notes one after the other" (LA).

4.2 Breaking the Score's Homogeneity

To better match composers' mental model and creative process, EuterPen supports two representations of the primary notation that

afford different manipulations. The first representation uses regular, beautified symbols as found in other music notation programs. The second notation is the composer's own handwriting. The two can coexist seamlessly on a staff and even be intermingled in the same bar. The former representation, because of its engraved look & feel, suggests fully-developed, finalized material; while the latter rather suggests material that is still work-in-progress and more likely to change. This is further reinforced by the differences in how composers manipulate the two representations. Handwritten symbols can easily be deleted by flipping the pen and using the eraser – as on paper. On the contrary, engraved symbols must first be selected before they can be deleted. Composers can thus easily jot ideas down by drawing symbols on the score, and they can as easily discard them with the pen's eraser, without running the risk of inadvertently modifying or deleting engraved symbols nearby.

The two representations – engraved and handwritten – can be used on all staves, in both the main composition space and pensieves. One important design decision we made was to leave handwritten music symbols seemingly uninterpreted by default. Handwritten symbols necessarily have to be interpreted when composers ask to engrave them (Figure 4-f), but this interpretation actually takes place in the background even before that, as soon as they get selected (with a lasso selection gesture). This lets composers perform operations that would typically only be possible on engraved material on handwritten notation as well – while preserving their *work-in-progress* representation. Such operations include : copy-and-pasting the selected notes (Figure 4-b & c); listening to



Figure 4: A note selection and the associated contextual menu: a) erase selection, b-c) copy-&-paste selection, d) play back selection, e) find occurrences of the selection, f/f') switch between engraved and handwritten representations.

them (Figure 4-d);³ and searching for other occurrences of the same pattern throughout the score (Figure 4-e and Figure 5-a) – both engraved and handwritten. Finally, the conversion from handwritten to engraved notation is not a one-way transformation. Composers can actually toggle between the two representations at will (provided the selected symbols were actually input with the pen in the first place), which can be useful for instance to visually revert a passage to a more *work-in-progress* look & feel.

The decision not to interpret handwritten symbols immediately was taken for two main reasons. The first reason was to give composers as much freedom as possible, preventing the parsing process from forcing a particular way of writing on composers and thus avoiding the pitfalls of greedy input interpretation commonly found in other music notation programs [12]. The second reason was to enable three very different types of ink-based input to coexist without resorting to mode switches: handwritten music symbols; arbitrary secondary notations such as text and drawings (Figure 5-c); and a variety of gesture-based commands (see Section 4.3). Composers can seamlessly perform any of these thanks to the following disambiguation strategy:

- any pen mark that 1) has been initiated in empty space, and 2) does not match one of the three predefined gestures (Figure 6), will remain uninterpreted and treated as a simple annotation (secondary notation);
- any pen mark that has been initiated on a staff will be considered a candidate music symbol for interpretation, if selected later;
- if the trace currently being inked matches one of the predefined gestures, the ink color will change, providing feedback to composers about how the gesture will be interpreted if they lift the pen then (for instance, the pin-to-play gesture in Figure 6-b);
- if this is not the interpretation they intended, composers can continue inking and lift the pen later when the gesture no longer matches any predefined one, which will be indicated by the ink reverting to its default color.

This strategy works all over the EuterPen workspace and lets composers add secondary notation not only in pensives but anywhere on staves as well, more tightly integrated with the primary notation – for instance ink marks made to circle and link different passages on the score. Such secondary notation is actually not limited to handwritten text and drawings, but can be much more heterogeneous. Composers can drop different media types directly

³Such auditory control is key to support elaborate tasks such as transcription, incrementation or modification

on the staves (replacing some bars) or very close to the staves. The list of supported media types is informed by results from Cavez *et al.*'s study [12] and by our workshop: an audio recording made on-the-spot using the tablet's microphone; a picture taken on-the-spot using its camera; a document (PDF, image, audio, video) imported from the file system or linked from a URL. The purpose and relevance of these secondary notations will evolve as the composition progresses. Composers can thus easily move, resize, show and hide, or entirely discard them at will.

Feedback

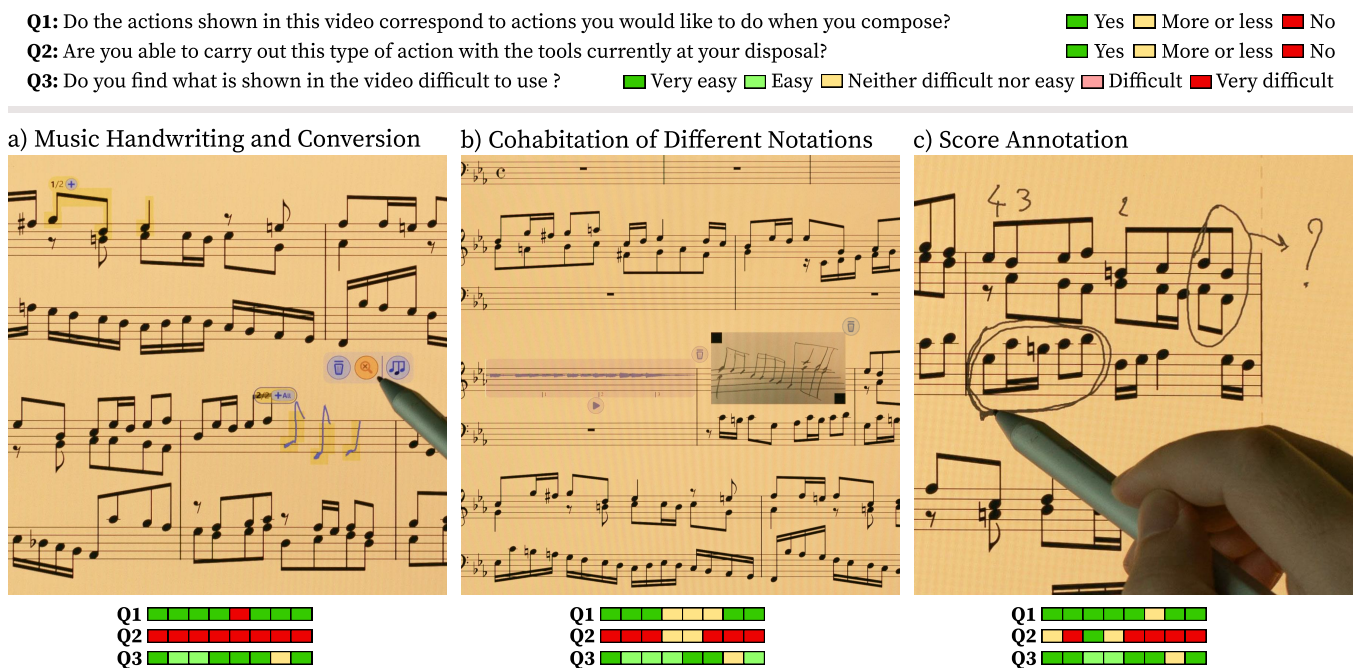
As illustrated in Figure 5, interviewees were very positive about EuterPen's way of breaking the score's homogeneity: mix handwritten and engraved notations, insert other media such as audio samples on staves, put secondary notations anywhere, were possibilities deemed useful by a majority of composers (Q1). Those features are largely absent from current music notation programs (Q2), and seemed generally easy-to-use in EuterPen (Q3).

Secondary notations are useful not only to make text annotations but to make freeform drawings as well – what LA calls “*supervised cohabitation*” – that can capture intent, help understand something, give information about musical structure (see Figure 1): “*Maderna's SERENADE POUR UN SATELLITE is a circular score*”... “*musically [...] it's illegible*” but “*in a circle you can make choices, you can go left and right ... [so] the idea of engraving it like that rather than explaining it logically, was really so that it would inspire musicians*” (LA). Drawing could also be a way to “*invent symbols that you're going to put in the score*” (AN) or to support musical ensembles as a way to “*build the score by inserting an idea of physical movement, even explaining it with a diagram, showing everyone's place, with a movement that's taken as you would on a sketch*” (YT).

The idea of integrating music passages in forms other than that of the primary notation (an audio sample, a picture of a roughly sketched passage to be transcribed - see Figure 5-b) had emerged during the workshop. Its relevance was further confirmed during the interviews: “*It would be fantastic if... [the system could play] what was engraved, written, and given that there's the MP3 afterwards, given that there's the audio afterwards, we'd really like it to follow on*” (LA). Similarly, C₁ wished they could work with audio recordings of themselves searching for ideas on the instrument and integrate those easily in the score: “*I have a lot of memos in my phone that should be developed, but they are in my phone, and I cannot put them in the score.*” Several interviewees saw potential in the combination of pensives and secondary notations as a means to notate and coordinate group work: “*It can be really useful because it transforms the score into a performance*” (DP). On the same topic YT said: “*We're in the process of doing something, creating a show with the double bass orchestra. We're making it by hand. In other words, there are lots of things. It starts with an idea. And then the work in progress happens in rehearsal, with five people, and everyone puts in their two cents.*” And since “*the sixth person wasn't there because she's away on business*”, so it's like “*an instruction manual*” for her.

The possibility to make the handwritten and engraved representations coexist had much success.⁴ Composers appreciated that

⁴One composer found this of little use (Figure 5-a), primarily because they found it difficult to draw notes on a screen.



Answers to Q2 color-coded red or yellow are to be interpreted in favor of EuterPen, as they indicate operations impossible or impractical in the score editors that interviewees currently use.

Figure 5: Breaking the score’s homogeneity: participants’ evaluation of the features demonstrated in the videos shown during the interviews.

many operations could be performed on handwritten passages as well, such as play them, copy-and-paste them, or look for other occurrences (Figure 4): “Searching for the motif to see if you’ve already written it before beautifying it, I’d use that every day” (LA). During the workshop, both C₁ and C₂ commented very positively on the possibility to immediately listen to handwritten passages using a simple gesture (Figure 6-b) without having to engrave them. But at a more fundamental level, the possibility to switch back-and-forth between engraved and handwritten representations was seen as a means for composers to keep track of where they are in their creative process: “You don’t have to ask yourself the question ‘Is this a final version, is this my draft?’” (LA), because this enables “visualiz[-ing] very, very easily a passage that needs to be reworked” (DP). As C₁ observed, it could also help composers keep track of the decisions made: “Going back in time can really help. Often you get the clean version and can’t remember how you got there.” But during the workshop composers also wondered how this would work for engraved notation not originally written by them (for instance, bars imported from a MusicXML file). C₁ emphasized that it would be “seriously awkward to see the handwriting of someone else”. In such cases, one option discussed with C₁ would be to use machine learning to train EuterPen on the composer’s own handwriting and to have it generate the handwritten representation, possibly involving *ink beautification* and *handwriting generation* mechanisms [43].

While some score editors do enable the annotation of scores, composers saw the possibility to draw and add pictures or audio recordings directly on the staves as a means to support creativity,

that could be useful for teaching as well: “We can even, while keeping what [the student] has written, add something else in manuscript [...] something that can go in another direction” (YT).

4.3 Breaking Down Musical Elements

As discussed in subsection 4.2, one of EuterPen’s core design principles is to make multiple types of ink-based input coexist without resorting to mode switches: handwritten music symbols, secondary notation, and predefined gesture-based commands. In addition to these, and consistent with current editors, EuterPen should support pen- and touch-based direct manipulation of music notation elements. This introduces further ambiguities that need to be resolved.

Predefined gesture commands. Three actions are triggered by pen-based gestures that are easy to learn and to perform (Figure 6). These gestures are recognized by a simple, robust classification algorithm that accounts for variability in gesturing. To be scale-independent, it uses the *gap* between two staff lines as a reference distance to classify input traces: *lasso-to-select* encloses an area at least $2 \cdot \text{gap}$ wide and high, ending within $2 \cdot \text{gap}$ of the start point; *pin-to-play* combines a lasso with a vertical trace at least $4 \cdot \text{gap}$ high; and finally *caret-to-insert* is a downward vertical trace of at least $4 \cdot \text{gap}$ high followed by an upward trace.

Direct Manipulation. Music notation elements that lie on staves can be moved horizontally or vertically to change the pitch or adjust the layout. For instance, when selecting a note head, a small cross will appear (see Figure 8-c), suggesting the possibility of direct manipulation in these directions. Dragging along another




Figure 6: The three predefined pen gestures available to composers, that coexist with all direct manipulations on the engraved notation as well as the input of handwritten notation – all without any mode-switch : a) lasso-to-select; b) pin-to-play; c) caret-to-insert (the amount of space to insert being controlled by the length of the upward trace, here 9 bars).

direction will initiate one of different lasso selections, as detailed next.

The digital pen is an excellent tool to write music symbols. But it is also a very precise selection tool, that can be used to delineate free-form areas. As such it is well-suited to the selection of music elements, which are composed of multiple tiny glyphs densely packed together. The modern staff notation is a multi-dimensional grammar that encodes different auditory attributes of a note (pitch, duration, *etc.*) by combining those glyphs and positioning them precisely on staves and ledges. The most frequent glyphs include the head, stem, flag, beam, dots and accidentals. Some glyphs define a note’s auditory attribute, while other glyphs are rather modifiers of that attribute. But regular score editors typically consider notes as entities that cannot be further decomposed. Selecting any constituent glyph of a note will select the whole note, that can thus only be moved (adjusting pitch), copy-&-pasted, or deleted.

EuterPen aims to break this monolithic view on notes and let composers select individual auditory attributes of a note if they want to. This is made possible by introducing novel types of selections. Composers can invoke those novel selection tools seamlessly, still without switching modes. We achieved this by adopting a strategy that disambiguates what to select based on *where* composers start their selection. A lasso enclosing multiple notes will select:

- the whole notation if initiated outside of any glyph – see Figure 6-a;
- handwritten notation only if initiated on handwritten symbols, ignoring engraved notation;
- a series of durations if initiated on a rhythm-related glyph (beam, stem, flag or rest) – see Figure 8-a;
- a series of pitches if initiated on a note head, dragging diagonally;⁵
- accidentals if initiated on an accidental;
- text if initiated on textual elements such as, *e.g.*, dynamics.

To help composers perform those different selections, EuterPen provides feedforward, dimming the glyphs that are not related to one of the selected auditory attributes until the lasso selection is completed. Figure 8-a illustrates this feedforward for a selection initiated on a rhythm-related glyph (in this case a quaver), temporarily dimming note heads and accidentals.

⁵As mentioned earlier, horizontal and vertical movements on note heads are reserved for direct manipulations.

While selecting individual auditory attributes makes little sense for coarse manipulations such as deletion, it provides composers with a whole new set of capabilities for more elaborate actions such as Find or Copy-&-Paste.

Find. Repetitions and variations are an essential part of many compositions [12]. Highlighting the occurrences of a melodic pattern is thus an important feature of music notation programs. EuterPen highlights such occurrences directly in the score, and also shows them on a minimap of the score (Figure 7-b), facilitating navigation over the result-set. Composers can then effortlessly add these occurrences to the active selection and manipulate them together – for instance to adjust their relative pitch. Combining this feature with EuterPen’s novel selection capabilities actually enables composers to make queries involving specific auditory attributes only. They are able to search for a rhythmic pattern regardless of pitch or, as illustrated in Figure 7-a, search for a pitch sequence regardless of rhythm.

Copy-&-Paste. Music has always been about repetition and variation (as AO says: “Mozart would have been so happy to be able to copy and paste notes! I think he must have copied quite a few...”), and composers copy-&-paste music notation very frequently. Available in many score editors (Figure 8-b, Q2), the possibility to duplicate music notation this way represents a major advantage of computer-based editing over paper. But the way regular score editors implement copy-&-paste is very monolithic and provides little feedforward. It makes it difficult for composers to anticipate what will happen, and sometimes yields unexpected or unwanted results. EuterPen takes a different approach to copy-&-paste. Dragging a pen-made selection with a finger will duplicate it (*drag-to-duplicate*). While dragging, notation currently at the destination will fade-out almost completely to show composers where the duplicate notation will be placed if they drop at this moment, as illustrated in Figure 8-b. Another way to control precisely where to paste is to explicitly select what to replace (*select-to-replace*). Once a selection has been copied, composers can select what symbols they want to replace with what is in the system’s clipboard.⁶ A tap on the PASTE button in the selection contextual menu (Figure 4-c) will then effectively replace the notation. Again, combining this feature with EuterPen’s novel selection capabilities opens new possibilities. By restricting

⁶Copied content is not necessarily meant to replace existing content. Drawing a small lasso in an empty spot will *insert* – rather than *replace* – the notation at this location.

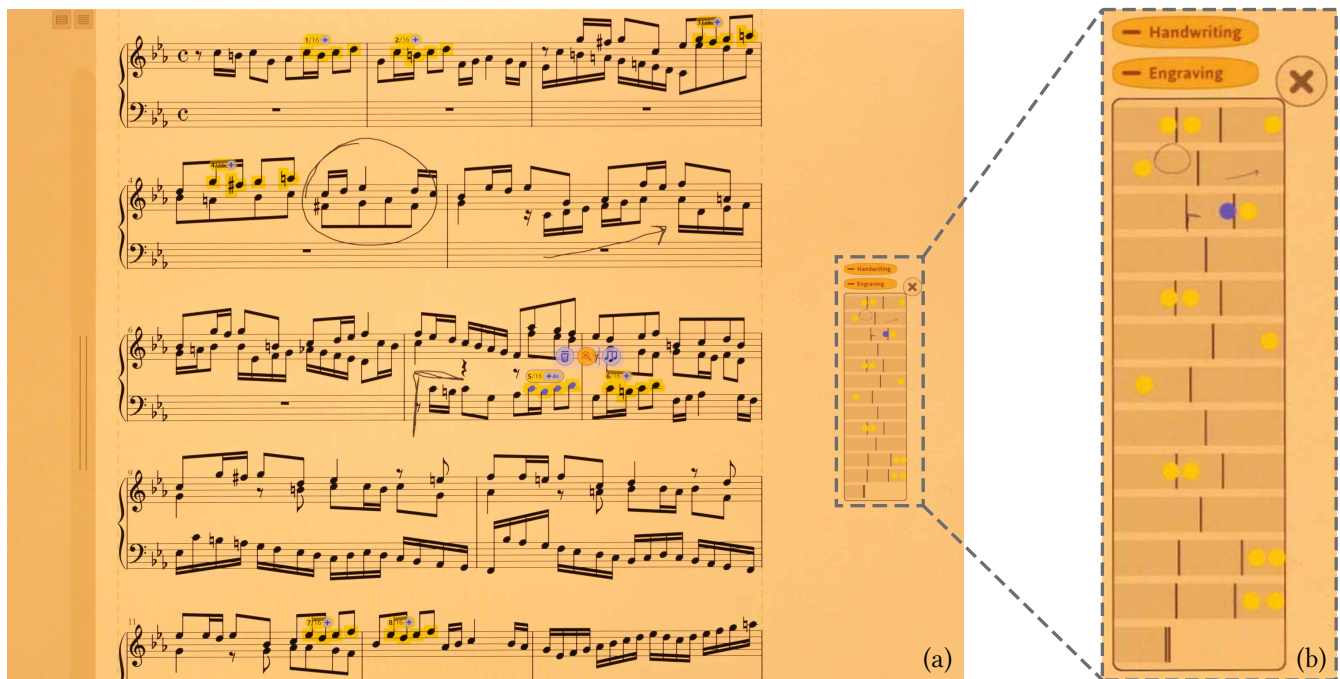


Figure 7: Searching for all occurrences of a melodic pattern based on pitch only, relaxing constraints on rhythm. a) Four note heads have been selected (colored blue, 7th staff), and all occurrences of the same pitch sequence (C-B-C-D) are highlighted yellow, regardless of their duration. Occurrences of the same sequence but transposed are also highlighted (e.g., Eb-D-Eb-F, 11th staff). b) The sixteen occurrences are also highlighted on the interactive minimap that gives an overview of the whole score. Annotations made on the score are also visible on the minimap, providing composers with navigation landmarks.

the selection to one auditory attribute only (pitch, duration), composers can for instance copy a sequence of pitches and paste them on another rhythm than the original one; or conversely they can copy a rhythm and paste it on another pitch sequence than the original one – all without having to erase and rewrite the entire sequence.

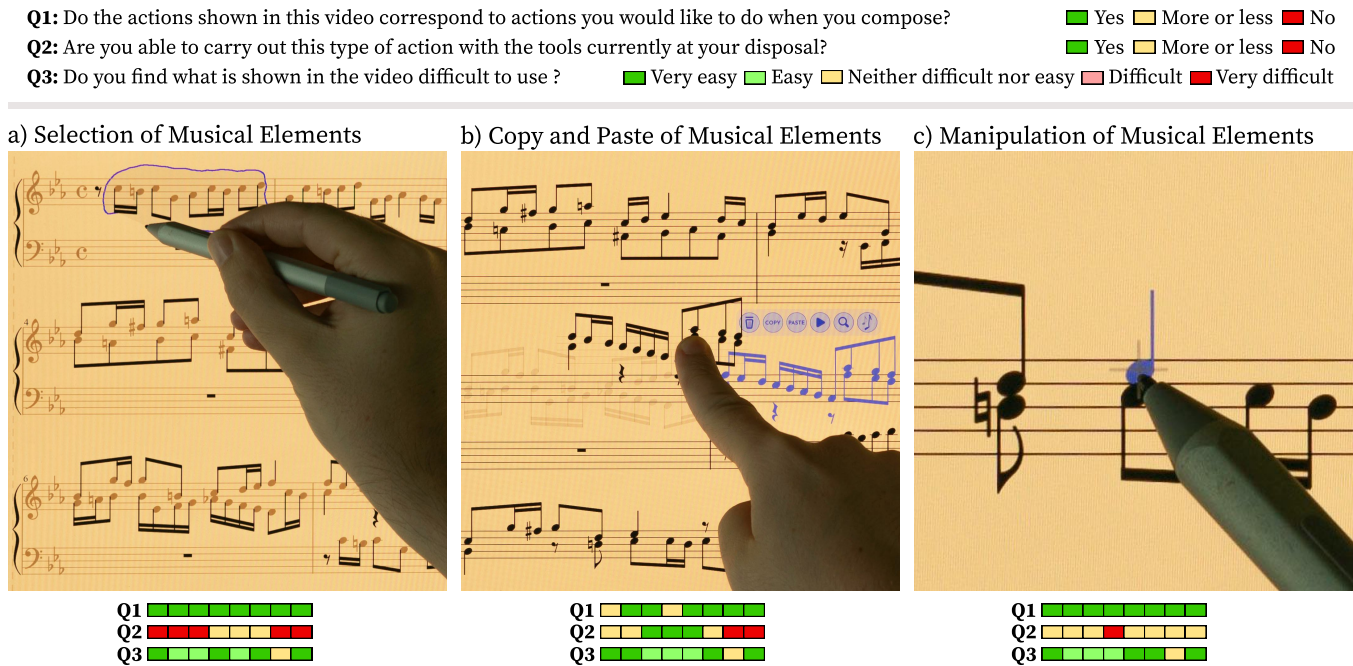
Feedback

As illustrated in Figure 8, interviewees were very enthusiastic with EuterPen’s innovative way of breaking down musical elements to select, copy-&-paste and manipulate them with precision. Composers were almost unanimous in recognising these features as useful (Q1), although half of them have already found workarounds with their current music notation programs (Q2). They found the way EuterPen supports these features easy to use (Q3).

Selection plays a key role in many operations and is of major concern to composers. Interviewees found EuterPen’s extended set of selection tools useful (Figure 8-a) and shared their thoughts about how they would use them in their work. For instance, OS would like to select melodic voices with ease “to erase and replace with something else because I made a mistake. It’s often the intermediate voices in particular that I have problems with. It’s true that you often have to erase everything and start all over again.” YT saw potential in auditory attribute selections, for instance to try different rhythms for a given pitch sequence: “If we have a melody that we like, we can

change the rhythm. I’ve sometimes started with a sequence of notes, and thought to myself ‘what can I do with this? Here, I’d like it to have this groove.’” But also the other way around, to manipulate only the melody: “We’ll change the pitches and keep a form that’s close by changing the notes [...] I’ll take this melody. So that will be my wink and reference to that melody, but I’m going to reverse it.”

Combining search with auditory attribute selections to highlight full but also partial matches for a given melodic pattern was also commented positively about. Modern classical composers seek structural clarity and our interviewees talked about how they could use relaxed queries constraining rhythm only: “When you’re doing an orchestration with lots and lots of content, it can be nice to say ‘Wait, let’s see all the people who do this rhythm’” (AO). Similarly, for melodic patterns: “[...] you search for the subject at the beginning and see it pop up in the score, and that’s it. Your analysis is done. Even hidden motifs will appear, things that can be overlooked because they’re placed between voices two and three, they’ll appear on their own” (LA). DP and BF, both jazz composers, give even more importance to “series of notes” and “series of rhythms”: “I only work with small pieces of material that I repeat and transform, and which end up in different forms, sometimes very hidden in different parts of the score” (DP). BF explains how he would use the search feature to find a formula to edit: “I know that I’ve put this formula in the middle of a development and I want to find it again. Maybe it’s quicker that way, since it’s highlighted. So it provides a focused reading of the big picture.” Composers including AN, LA and BF also expressed the



Answers to Q2 color-coded red or yellow are to be interpreted in favor of EuterPen, as they indicate operations impossible or impractical in the score editors that interviewees currently use.

Figure 8: Breaking down musical elements: participants' evaluation of the features demonstrated in the videos shown during the interviews.

need to extend the search's scope: to be able to search in pensieves and to include secondary notation as well. The minimap (Figure 7) was first suggested during the workshop, and requested again by several interviewees, but only added to EuterPen in the last prototyping phase. It shows not only melodic pattern matches during a search, but bar lines and handwritten annotations as well, as these provide strong navigation cues.

Feedback about direct manipulation of the primary notation was positive (Figure 8-c). BF emphasized the need for fine-grain manipulation to improve legibility: "when the voices are very, very close together, simply to avoid blots. And if you have accidentals in addition to this crowded writing... it gives you clarity." While most composers can already adjust symbol placement with their current editor, some complained about the rigidity of such manipulations and their lack of stability: "It moves all the notes. For example, I often have a problem with arpeggio notes that, you know, immediately change the dimensions of the score. Just by having this arpeggio, you end up with a bar that's the size of two bars. And then, typically, every time I want to move something, it's going to move the whole score" (AO). Interviewees also shared ideas about how to further improve direct manipulation of the primary notation, as briefly discussed in Section 6.

During the workshop's brainstorming session, C₁ and C₂ both imagined that sequences of notes could be dragged outside of the score, and dwelled on the concept of a rich-yet-straightforward manipulation of specific auditory attributes, either when copying or when pasting. This inspired our alternative proposals to classic

copy-&-paste, which were well received by interviewees. Commenting about *drag-to-duplicate*, AO observed: "I like the fact that you can already perform an action without having to click on a button," BF adding: "This movement, we did not have this ease with the editors. It had to be copied... Moving is complex because you often have to cancel a measure to put something back in [...]. [In the meantime you need to put] rests, because an empty measure is not tolerated by the software." Interviewees described different uses of copy-&-paste combined with auditory-attribute selections. Leveraging pensieves, AN saw potential for "a sort of idea box where we have a zone with rhythmic ideas, a zone with melodic progression ideas and we can take one or the other." LA, who insists on the importance of stability of the score, further stresses the importance of spaces where he can work on duplicated elements before integrating them into the main temporal flow: "It's a good thing that the pseudo-final material, or the engraved version in any case, is relatively fixed, so that we don't tamper with it too much, but rather copy [some of] it into draft spaces and work there. That way, there's a version that moves less, otherwise, in fact, we'd be, I think, very tempted to throw everything left and right. [Things] would be moving all the time." Other composers (BF, CI and OS) also expressed interest for auditory attribute copy-&-paste, for instance to change the rhythm of a melodic sequence, acknowledging how easy EuterPen makes such operations: "This is something we do regularly when there are repetitions of formulas [...] except that we often take only the rhythm. We take the formula and then edit the notes one by one, on the formula we've just copied. So it's a lot more laborious" (BF).

5 Discussion

Talking about EuterPen, one of our interviewees said: *“It’s not just sand. [...] You can add a lot of different things in a lot of different places, but it’s all still linked to something that’s constantly there and that you can reframe, in cases where you’ve expanded it a bit, which is the score that you’re working on. And so, around this thing that remains, you can really develop your thinking and organise your work. A supervised cohabitation”* (LA). This comment captures the balance that EuterPen seeks to strike between flexibility, structure and stability, as discussed in the previous sections.

Flexibility can be opposed to structure. It can also be opposed to stability. But in our context structure and stability are two orthogonal concerns. Regular music notation programs are heavily structured, imposing many constraints on the notation and its manipulation. But this structure does not contribute to the score’s stability, and can even play against it, structured editors sometimes moving notation automatically to enforce syntactic or layout rules.

The tension between flexibility and structure primarily exists within staves, when editing the primary notation. Structure is what enables advanced selections, copy-&-paste and direct manipulations on (groups of) notes, melodic pattern matching, well-formedness checks. But structure is also what separates writing music in a score editor from writing music on paper. By systematically enforcing the notation’s rules, regular music notation programs unnecessarily restrict composers’ means to capture their ideas, when those creative *professionals* most often know what they are doing and would prefer to fix syntactic and structural issues later. EuterPen aims to balance structure and flexibility, temporarily relaxing writing and editing constraints to foster creativity. This is for instance achieved by letting users decide when to interpret inked marks to perform advanced operations [40, 48]; or letting them turn their notation back into a “work-in-progress” state at will [51], instead of forcing a more permanent choice at the time of creation. Another way of relaxing constraints while keeping structure within reach is to trust composers with more control over the layout of symbols and measures. EuterPen allows composers to take the time and space they need to work [41, 52], providing them with the means to trigger automatic layout optimization as they see fit.

The tension between flexibility and stability takes place at a more global level. While this concern was part of the initial design ideas, the brainstorming session on copy-&-paste during the workshop emphasized its importance, which was then confirmed during the interviews. Composers want to be able to operate in the way that suits them best at a given stage, without having to change modes, without having to break what they have already built, with free spaces at hand in which to experiment. To paint a metaphor, composers we worked with ask for a tool that will let them play in sandboxes while preserving the castles they have already built. In that sense, pensieves serve two purposes: a means to store material that supports their creative process; and a means to freely explore creative ideas while preserving the stability of their work. As such, we approached the design of pensieves as canvases that can be instantiated on-demand in different places, integrated in and around the structure of the score [3] – as opposed to having a single monolithic canvas that better suits other types of tasks such as sketching [23] or expressive visualization authoring [39, 50].

One important consideration about pensieves is that participating composers did not see these flexible spaces as unstructured, but quite the contrary. Pensieves let composers drop and freely organize heterogeneous pieces of content, but they offer the same interpretative functionalities as the temporal composition space. Some composers actually went further, suggesting that pensieves could feature “drawers” to organize content, or that they could be used to organize performances of their piece. Pensieves should really be understood as spaces complementing the temporal composition space, not competing with it.

6 Conclusion and Future Work

EuterPen is a music notation program prototype that leverages the affordances of pen and touch, while selectively relaxing both syntactic and structural constraints of the music notation. This approach aims to strike a balance between the need for structure and the need for flexibility during the music composition process, in an effort to reduce constraints on the composers’ creative expression. We designed EuterPen according to an iterative process with composers, including an online questionnaire in the final stage where eight composers were invited to provide feedback on the tool’s features. Although composers reacted very positively in this first evaluation, future work should observe how composers would adopt a tool like EuterPen in their composition process with a hands-on user evaluation or, ideally, a longitudinal study performed in their work environment.

Beyond a more comprehensive evaluation, future work could also focus design efforts on leveraging structure even further. In our questionnaire, composers reacted very positively to EuterPen’s flexibility. They expressed much interest in features that benefit from structure as well – selection capabilities and the associated manipulations in particular. During their interview, one of the composers suggested pushing the direct manipulation of music symbols further, so that symbols would adapt in real-time while being moved around on the staff with the pen – for instance merging two symbols into a single (equivalent) one when they get juxtaposed. LA explained how he envisioned this: simulating the drag of an eighth rest with the pen and approaching another eighth rest, he said *“the eighth rest becomes a quarter rest [...] the software understands that you’re moving the note. It’s not just a graphical update, it’s a rhythmic update.”* Brainstorming about the same idea, DP suggested adding feedforward when dragging symbols this way: *“I think there needs to be a little barrier, a little help. Maybe it’s just when you get to the middle of the two notes, that a little gray bar appears and you think ‘Oh no, I’m going too far.’”* Such advanced direct manipulations of notes and rests would require the tight coupling of direct manipulation interactions with structural transformations to re-write the engraved notation in real-time. This would require defining an entire rule-set of syntactically correct and semantically meaningful notation transformations that could be studied in future work.

Some interviewees also discussed the possibility to automatically interpret foreign material imported in EuterPen, such as an audio sample or a photo of a few bars quickly sketched on paper (both illustrated in Figure 5-b): *“What I’d like is for the photo taken there to be directly integrated into the score. That would be great”* (AN); *“There must be a way for it to understand, ‘Oh yes, this is music, let’s*

try to play it.' [...] in fact, when you think about it, it's just music that's there, but in a different form" (LA). While this is also left as a possible avenue for future work, recent advances in Artificial Intelligence might already have made this possible.

Acknowledgments

We warmly thank the music composers who generously devoted their time to share their invaluable experience during the workshop and the interviews. We would also like to thank MyScript for providing us with handwritten music recognition software, and Ludovic David for contributing code to interactively resize measures in SVG. This project has been partially supported by grant ANR Continuum (ANR-21-ESRE-0030).

References

- [1] J. Anstice, T. Bell, A. Cockburn, and M. Setchell. 1996. The Design of a Pen-Based Musical Input System. In *Proceedings of the 6th Australian Conference on Computer-Human Interaction (OZCHI '96)* (OZCHI '96). IEEE Computer Society, USA, 260–267. doi:10.1109/OZCHI.1996.560019
- [2] Caroline Appert and Shumin Zhai. 2009. Using Strokes as Command Shortcuts: Cognitive Benefits and Toolkit Support. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston, MA, USA) (CHI '09). Association for Computing Machinery, New York, NY, USA, 2289–2298. doi:10.1145/1518701.1519052
- [3] Ian Arawjo, Anthony DeArmas, Michael Roberts, Shrutarshi Basu, and Tapan Parikh. 2022. Notational Programming for Notebook Environments: A Case Study with Quantum Circuits. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology* (Bend, OR, USA) (UIST '22). Association for Computing Machinery, New York, NY, USA, Article 62, 20 pages. doi:10.1145/3526113.3545619
- [4] Amelia Bellamy-Royds, Bogdan Brinza, Chris Lilley, Dirk Schulze, David Storey, and Eric Willigers. 2018. Scalable Vector Graphics (SVG) 2 - W3C Candidate Recommendation. <https://www.w3.org/TR/SVG2/>.
- [5] Stan Bennett. 1976. The Process of Musical Creation: Interviews with Eight Composers. *Journal of Research in Music Education* 24, 1 (1976), 3–13.
- [6] Alan F Blackwell, Thomas RG Green, and Douglas JE Nunn. 2000. Cognitive dimensions and musical notation systems. In *Workshop on Notation and Music Information Retrieval in the Computer Age at the International Computer Music Conference (ICMC)*. 6 pages.
- [7] Michael Bostock, Vadim Ogievetsky, and Jeffrey Heer. 2011. D3: Data-Driven Documents. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (2011), 2301–2309. doi:10.1109/TVCG.2011.185
- [8] Jean Bresson, Carlos Agon, and Gérard Assayag. 2011. OpenMusic: Visual Programming Environment for Music Composition, Analysis and Research. In *Proceedings of the 19th ACM International Conference on Multimedia* (Scottsdale, Arizona, USA) (MM '11). Association for Computing Machinery, New York, NY, USA, 743–746. doi:10.1145/2072298.2072434
- [9] Matt Brubeck, Rick Byers, Patrick H. Lauke, and Navid Zolghadr. 2019. Pointer Events Level 2 - W3C Recommendation. <https://www.w3.org/TR/pointerevents2/>.
- [10] Jorge Calvo-Zaragoza, Jan Hajič Jr., and Alexander Pacha. 2020. Understanding Optical Music Recognition. *ACM Comput. Surv.* 53, 4, Article 77 (jul 2020), 35 pages. doi:10.1145/3397499
- [11] Vincent Cavez, Caroline Appert, and Emmanuel Pietriga. 2024. Spreadsheets on Interactive Surfaces: Breaking through the Grid with the Pen. *ACM Trans. Comput.-Hum. Interact.* 31, 2, Article 16 (jan 2024), 33 pages. doi:10.1145/3630097
- [12] Vincent Cavez, Catherine Letondal, Emmanuel Pietriga, and Caroline Appert. 2024. Challenges of Music Score Writing and the Potentials of Interactive Surfaces. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 728, 16 pages. doi:10.1145/3613904.3642079
- [13] Mohit Muthanna Cheppudira. 2024. VexFlow - JavaScript Music Notation and Guitar Tablature. <https://www.vexflow.com>. Last accessed: 2024-08-01.
- [14] David Collins. 2005. A synthesis process model of creative thinking in music composition. *Psychology of Music* 33, 2 (2005), 193–216. doi:10.1177/0305735605050651
- [15] Göran Folkestad, David J. Hargreaves, and Berner Lindström. 1998. Compositional Strategies in Computer-Based Music-Making. *British Journal of Music Education* 15, 1 (1998), 83–97. doi:10.1017/S0265051700003788
- [16] Jérémie Garcia, Theophanis Tsandilas, Carlos Agon, and Wendy Mackay. 2012. Interactive Paper Substrates to Support Musical Creation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Austin, Texas, USA) (CHI '12). Association for Computing Machinery, New York, NY, USA, 1825–1828. doi:10.1145/2207676.2208316
- [17] Jérémie Garcia, Theophanis Tsandilas, Carlos Agon, and Wendy E. Mackay. 2011. InkSplore : Exploring Musical Ideas on Paper and Computer. In *Proceedings of the International Conference on New Interfaces for Musical Expression (NIME '11)*. Oslo, Norway, 361–366. doi:10.5281/zenodo.1178027
- [18] James J Gibson. 1979. *The Theory of Affordances. The Ecological Approach to Visual Perception*. Houghton Mifflin.
- [19] T. R. G. Green and M. Petre. 1996. Usability Analysis of Visual Programming Environments: A 'Cognitive Dimensions' Framework. *Journal of Visual Languages & Computing* 7, 2 (1996), 131–174. doi:10.1006/jvlc.1996.0009
- [20] Patrick G. T. Healey and Jean-Baptiste Thiebaud. 2007. Sketching Musical Compositions. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, Vol 29. 6 pages. <https://escholarship.org/uc/item/3kc0t66r>
- [21] D.L. Hecht. 2001. Printed embedded data graphical user interfaces. *IEEE Computer* 34, 3 (2001), 47–55. doi:10.1109/2.910893
- [22] Ken Hinckley, Koji Yatani, Michel Pahud, Nicole Coddington, Jenny Rodenhose, Andy Wilson, Hrvoje Benko, and Bill Buxton. 2010. Pen + Touch = New Tools. In *Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology* (New York, New York, USA) (UIST '10). Association for Computing Machinery, New York, NY, USA, 27–36. doi:10.1145/1866029.1866036
- [23] Jun Kato, Kenta Hara, and Nao Hirasawa. 2024. Griffith: A Storyboarding Tool Designed with Japanese Animation Professionals. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 233, 14 pages. doi:10.1145/3613904.3642121
- [24] Annick Leroy, Giovanni Müller, and Guy E. Garnett. 1994. The Design of a Pen-Based Music Notation System. In *Proceedings of the 1994 International Computer Music Conference (ICMC '94)*. 286–292.
- [25] Catherine Letondal and Wendy Mackay. 2007. The Paperoles Project: An Analysis of Paper Use by Music Composers. In *Proceedings of the 2nd International Workshop on Collaborating over Paper and Digital Documents (CoPADD)*. 27–30.
- [26] Catherine C. Marshall. 1997. Annotation: From Paper Books to the Digital Library. In *Proc. of the ACM Int. Conf. on Digital libraries*. ACM Press, 131–140. doi:10.1145/263690.263806
- [27] Fabrice Matulic and Moira C. Norrie. 2013. Pen and Touch Gestural Environment for Document Editing on Interactive Tabletops. In *Proceedings of the 2013 ACM International Conference on Interactive Tabletops and Surfaces* (St. Andrews, Scotland, United Kingdom) (ITS '13). Association for Computing Machinery, New York, NY, USA, 41–50. doi:10.1145/2512349.2512802
- [28] Hrim Mehta, Adam Bradley, Mark Hancock, and Christopher Collins. 2017. Metatation: Annotation as Implicit Interaction to Bridge Close and Distant Reading. *ACM Trans. Comput.-Hum. Interact.* 24, 5, Article 35 (nov 2017), 41 pages. doi:10.1145/3131609
- [29] Thomas P. Moran, Patrick Chiu, William van Melle, and Gordon Kurtenbach. 1995. Implicit structure for pen-based systems within a freeform interaction paradigm. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '95). ACM Press/Addison-Wesley Publishing Co., USA, 487–494. doi:10.1145/223904.223970
- [30] MyScript. 2024. Cross-platform handwriting recognition and interactive ink APIs. <https://developer.myscript.com>. Last accessed: 2024-08-01.
- [31] Chris Nash. 2015. The Cognitive Dimensions of Music Notations. In *Proceedings of the First International Conference on Technologies for Music Notation and Representation (TENOR '15)*, Marc Battier, Jean Bresson, Pierre Couprie, Cécile Davy-Rigaux, Dominique Foer, Yann Geslin, Hugues Genevois, François Picard, and Alice Tacaille (Eds.). Paris, France, 190–202.
- [32] E. Ng, T. Bell, and A. Cockburn. 1998. Improvements to a pen-based musical input system. In *Proceedings 1998 Australasian Computer Human Interaction Conference. (OzCHI)*. 178–185. doi:10.1109/OZCHI.1998.732212
- [33] Notion. 2024. Music Notation On the Go. <https://legacy.presonus.com/products/Notion-Mobile>. Last accessed: 2024-07-03.
- [34] Ken Pfeuffer, Ken Hinckley, Michel Pahud, and Bill Buxton. 2017. Thumb + Pen Interaction on Tablets. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 3254–3266. doi:10.1145/3025453.3025567
- [35] Symphony Pro. 2024. Music notation & composition app for iPad. <https://symphonypro.net>. Last accessed: 2024-06-28.
- [36] Miller Puckette. 2002. Max at Seventeen. *Computer Music Journal* 26, 4 (2002), 31–43.
- [37] Yann Riche, Nathalie Henry Riche, Ken Hinckley, Sheri Panabaker, Sarah Fuelling, and Sarah Williams. 2017. As We May Ink? Learning from Everyday Analog Pen Use to Improve Digital Ink Experiences. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 3241–3253. doi:10.1145/3025453.3025716
- [38] Hans Roels. 2016. Comparing the main compositional activities in a study of eight composers. *Musicae Scientiae* 20, 3 (2016), 413–435. doi:10.1177/1029864915624737

- [39] Hugo Romat, Caroline Appert, and Emmanuel Pietriga. 2019. Expressive authoring of node-link diagrams with graphies. *IEEE Transactions on Visualization and Computer Graphics* 27, 4 (2019), 2329–2340.
- [40] Hugo Romat, Nathalie Henry Riche, Ken Hinckley, Bongshin Lee, Caroline Appert, Emmanuel Pietriga, and Christopher Collins. 2019. ActiveInk: (Th)Inking with Data. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3290605.3300272
- [41] Hugo Romat, Emmanuel Pietriga, Nathalie Henry-Riche, Ken Hinckley, and Caroline Appert. 2019. SpaceInk: Making Space for In-Context Annotations. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology* (New Orleans, LA, USA) (UIST '19). Association for Computing Machinery, New York, NY, USA, 871–882. doi:10.1145/3332165.3347934
- [42] Abigail J. Sellen and Richard H.R. Harper. 2003. *The Myth of the Paperless Office*. MIT Press, Cambridge, MA, USA.
- [43] Lei Shi, Rachel Campbell, Peggy Chi, Maria Cirimele, Mike Cleron, Kirsten Climer, Chelsey Q Fleming, Ashwin Ganti, Philippe Gervais, Pedro Gonnet, Tayeb A Karim, Andrii Maksai, Chris Melancon, Rob Mickle, Claudiu Musat, Palash Nandy, Xiaoyu Iris Qu, David Robishaw, Angad Singh, and Mathangi Venkatesan. 2024. Inkeraction: An Interaction Modality Powered by Ink Recognition and Synthesis. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 181, 26 pages. doi:10.1145/3613904.3642498
- [44] StaffPad. 2024. Music composition app designed for pen and touch. <https://www.staffpad.net>. Last accessed: 2024-06-28.
- [45] Anselm Strauss and Juliet Corbin. 1998. *Basics of qualitative research: Techniques and procedures for developing grounded theory*, 2nd ed. Sage Publications.
- [46] Poorna Talkad Sukumar, Anqing Liu, and Ronald Metoyer. 2018. Replicating User-Defined Gestures for Text Editing. In *Proceedings of the 2018 ACM International Conference on Interactive Surfaces and Spaces* (Tokyo, Japan) (ISS '18). Association for Computing Machinery, New York, NY, USA, 97–106. doi:10.1145/3279778.3279793
- [47] OSMD Team. 2024. OpenSheetMusicDisplay. <https://opensheetmusicdisplay.org>. Last accessed: 2024-08-01.
- [48] Theophanis Tsandilas, Catherine Letondal, and Wendy E. Mackay. 2009. Musink: Composing Music through Augmented Drawing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston, MA, USA) (CHI '09). Association for Computing Machinery, New York, NY, USA, 819–828. doi:10.1145/1518701.1518827
- [49] Peter R Webster. 2002. *Creativity and music education*. Canadian Music Educators' Association, Chapter Creative thinking in music: Advancing a model, 16–33.
- [50] Haijun Xia, Nathalie Henry Riche, Fanny Chevalier, Bruno De Araujo, and Daniel Wigdor. 2018. DataInk: Direct and Creative Data-Oriented Drawing. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (CHI '18). Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3173574.3173797
- [51] Haijun Xia, Ken Hinckley, Michel Pahud, Xiao Tu, and Bill Buxton. 2017. Writ-Large: Ink Unleashed by Unified Scope, Action, & Zoom. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 3227–3240. doi:10.1145/3025453.3025664
- [52] Dongwook Yoon, Nicholas Chen, and François Guimbretière. 2013. TextTearing: opening white space for digital ink annotation. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology* (St. Andrews, Scotland, United Kingdom) (UIST '13). Association for Computing Machinery, New York, NY, USA, 107–112. doi:10.1145/2501988.2502036
- [53] Dongwook Yoon, Nicholas Chen, François Guimbretière, and Abigail Sellen. 2014. RichReview: Blending Ink, Speech, and Gesture to Support Collaborative Document Review. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology* (Honolulu, Hawaii, USA) (UIST '14). Association for Computing Machinery, New York, NY, USA, 481–490. doi:10.1145/2642918.2647390