



Contents lists available at ScienceDirect

Expert Systems With Applications

journal homepage: www.elsevier.com/locate/eswa

Late multimodal fusion for image and audio music transcription

María Alfaro-Contreras^{*}, Jose J. Valero-Mas, José M. Iñesta, Jorge Calvo-Zaragoza

University Institute for Computing Research, University of Alicante, Carretera San Vicente del Raspeig s/n, 03690, Alicante, Spain

ARTICLE INFO

Keywords:

Optical Music Recognition
Automatic Music Transcription
Multimodality
Deep learning
Connectionist Temporal Classification
Sequence labeling
Word graphs

ABSTRACT

Music transcription, which deals with the conversion of music sources into a structured digital format, is a key problem for Music Information Retrieval (MIR). When addressing this challenge in computational terms, the MIR community follows two lines of research: music documents, which is the case of Optical Music Recognition (OMR), or audio recordings, which is the case of Automatic Music Transcription (AMT). The different nature of the aforementioned input data has conditioned these fields to develop modality-specific frameworks. However, their recent definition in terms of sequence labeling tasks leads to a common output representation, which enables research on a combined paradigm. In this respect, multimodal image and audio music transcription comprises the challenge of effectively combining the information conveyed by image and audio modalities. In this work, we explore this question at a late-fusion level: we study four combination approaches in order to merge, for the first time, the hypotheses regarding end-to-end OMR and AMT systems in a lattice-based search space. The results obtained for a series of performance scenarios—in which the corresponding single-modality models yield different error rates—showed interesting benefits of these approaches. In addition, two of the four strategies considered significantly improve the corresponding unimodal standard recognition frameworks.

1. Introduction

A long-standing research problem in Music Information Retrieval (MIR) is that of attaining structured digital representations from music sources, typically known as (music) *transcription* (Serra et al., 2013). The MIR community follows two main research lines that study how to computationally solve this problem when targeting either music documents—known as Optical Music Recognition (OMR) (Calvo-Zaragoza, Hajič, & Pacha, 2020)—or acoustic music signals—namely, Automatic Music Transcription (AMT) (Benetos, Dixon, Duan, & Ewert, 2018). Despite having a similar purpose, these two fields have historically evolved in separate ways owing to differences in the nature of the data, which has resulted in specific task-oriented recognition frameworks that are most typically based on multi-stage procedures (de la Fuente, Valero-Mas, Castellanos, & Calvo-Zaragoza, 2021).

However, some recent proposals in MIR literature frame transcription problems in a sequence labeling formulation that approaches the task in a holistic or end-to-end manner (Calvo-Zaragoza & Rizo, 2018; Liu, Morfi, & Benetos, 2021): the input data—either scores or acoustic pieces—are directly decoded into a series of music-notation symbols. This makes it possible to address OMR and AMT tasks with similar recognition models that differ only as regards the input data used to train the system. A graphic illustration of these recognition approaches is provided in Fig. 1.

The common formulation makes it possible to explore possible synergies that may exist between image and audio sources. Some of the possible research avenues yet to be explored by the community are: developing common language models, devising multi-task neural architectures capable of dealing with both tasks independently in a single model, using pre-trained models with one modality and fine-tuning with the other, or multimodal image and audio transcription (Alfaro-Contreras, Valero-Mas, Iñesta, & Calvo-Zaragoza, 2022). The last case is the main focus of this work.

The use of multimodal recognition frameworks, which refer to those that contemplate multiple modalities or representations of the same datum as input, has led to advances in a variety of fields (Pitsikalis, Katsamanis, Theodorakis, & Maragos, 2017; Singh, Sangwan, & Hansen, 2012; Toselli, Vidal, & Casacuberta, 2011). The different modalities in these schemes are expected to provide complementary information to the system, eventually resulting in an overall improvement in performance. These approaches are generally classified according to the following taxonomy (Dumas, Signer, & Lalanne, 2012): (i) *feature* or *early-based* methods that directly merge the individual features of the modalities; and (ii) *decision* or *late-fusion* methods that combine the different hypothesis retrieved for each input representation. Note that the latter family of approaches allows model flexibility, since it

^{*} Corresponding author.

E-mail addresses: malfaro@dlsi.ua.es (M. Alfaro-Contreras), jvalero@dlsi.ua.es (J.J. Valero-Mas), inesta@dlsi.ua.es (J.M. Iñesta), jcalvo@dlsi.ua.es (J. Calvo-Zaragoza).

<https://doi.org/10.1016/j.eswa.2022.119491>

Received 1 September 2022; Received in revised form 19 December 2022; Accepted 28 December 2022

Available online 3 January 2023

0957-4174/© 2023 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

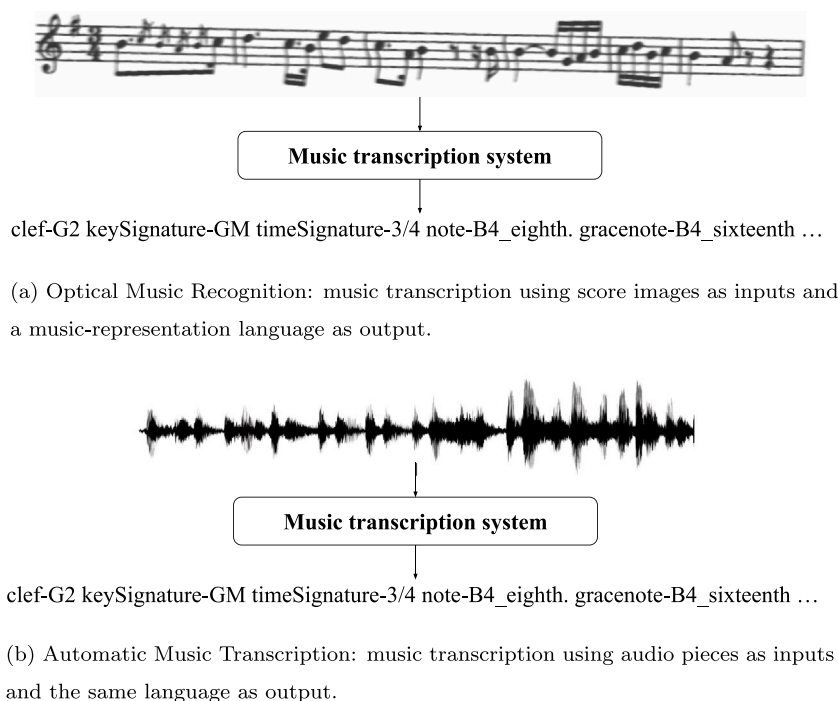


Fig. 1. End-to-end music transcription framework. OMR techniques deal with images and AMT techniques deal with audio signals; however, both tasks have to provide a result in a symbolic format that represents a piece of music.

requires only the individual hypotheses for the different modalities for the combination.

In the context of music transcription, while this might not be entirely true, both a recording and a score image can be considered as two complementary modalities of the same piece of music. On the one hand, the image provides all the information that was intended to be stored, but is in a graphic domain that must be decoded in order to interpret it as music. On the other, the audio contains information about the musical performance, but certain aspects are difficult to retrieve owing to human interpretation or certain unavoidable ambiguities (such as the clef or the meter). If the original composition were in the form of both an image and a recording, the two could consequently be combined in a synergistic manner so as to leverage the transcription advantages depicted by each modality. A promising task that has, with a few exceptions (de la Fuente et al., 2021), barely been explored in literature, is, therefore, that of developing approaches that are capable of efficiently combining the information conveyed by the image and the audio of the same piece of music for their joint transcription.

In this work, we explore multimodality in image and audio music transcription by combining each individual transcription hypothesis yielded by OMR and AMT models, respectively. More precisely, our contributions are: (i) a thorough revision and first-time application of four existing late-fusion recognition approaches to multimodal image and audio music transcription; (ii) comprehensive experimentation with nine different performance scenarios, involving cases in which OMR and AMT depict different base errors—high, medium, and low performance levels are considered—; and (iii) the demonstration of a significant improvement to the error rate with respect to existing unimodal approaches.

The remainder of the paper is organized as follows: Section 2 provides the background to this work, while Section 3 introduces the proposed approach for unimodal recognition. Section 4 thoroughly develops the multimodal image and audio transcription framework considered, and Section 5 describes the experimental setup. Section 6 presents and analyses the results, and finally, Section 7 concludes the work and discusses possible ideas for future research.

2. Related work

Traditional attempts to carry out OMR and AMT have been framed within a pipeline that divides the process into a series of independent phases (Liu & Benetos, 2021; Rebelo et al., 2012). These approaches have been highly conditioned by the nature of their input, thus resulting in incompatible individual (or unimodal) recognition frameworks. However, recent developments have formulated both tasks as *sequence labeling* (Graves, Fernández, Gomez, & Schmidhuber, 2006): the input data (either image or audio) is directly decoded into a series of music notation symbols. This is generally accomplished by employing neural end-to-end systems, and specifically Convolutional Recurrent Neural Network (CRNN) architectures owing to their competitive results. In the aforementioned approach, the convolutional stage extracts the most appropriate features for the case in hand, while the recurrent part models the temporal (or spatial) dependencies of these symbols. Given the advantage of being trained using unsegmented sequential data without any input–output alignment, the CRNN scheme introduced is usually considered together with the Connectionist Temporal Classification (CTC) loss function (Graves et al., 2006). This particular framework currently constitutes one of the state-of-the-art end-to-end approaches for both OMR and AMT tasks (Benetos et al., 2018; Calvo-Zaragoza et al., 2020). This common formulation has naturally triggered an interest in discovering intersections between these research lines, which have previously been addressed separately. Multimodality is, therefore, a suitable avenue of research by which to further study and exploit the synergies between these individual recognition schemes.

Multimodal frameworks seek to leverage the information depicted by each input modality in a synergistic manner (Simonetta, Ntalampiras, & Avanzini, 2019). This combination can be performed at two levels: at the feature or early-fusion level or at the decision or late-fusion level. The former fuses the data “as is”—the different input modalities undergo a pre-processing phase in order to extract their corresponding features, which are later merged and passed through a single processing algorithm. The latter approach, however, combines the outputs of several ad-hoc algorithms, one for each modality. In this

respect, late-fusion multimodality allows more flexibility—particularly when the data sources are significantly different from each other.

Scientific literature comprises a large number of works that consider recognition tasks within a multimodal framework. A first example is that by Singh et al. (2012), which proposes the late fusion of individual Text Recognition (TR) and Automatic Speech Recognition (ASR) systems for postal code recognition using a heuristic approach based on the edit distance. In the context of handwritten manuscripts, recent proposals rely on probabilistic frameworks to integrate individual estimations using word-graph hypothesis spaces (Granell, Martínez-Hinarejos, & Romero, 2018) or a compact version of them, namely confusion networks (CN) (Granell & Martínez-Hinarejos, 2015).

Multimodality has also benefited Gesture Recognition. For instance, Pitsikalis et al. (2017) boost the performance by re-scoring its different hypotheses using information from an ASR system. Other works have considered the use of dynamic programming techniques (Miki, Kitaoka, Miyajima, Nishino, & Takeda, 2014) or the aforementioned CN paradigm (Kristensson & Vertanen, 2011) to align different hypotheses.

With regard to the MIR field, recent studies exploit different modalities that capture complementary aspects of the same piece of music. In this respect, there are multimodal approaches for different tasks, such as artist identification, music recommendation, or instrument classification, amid others (Simonetta et al., 2019). These techniques are well known in the field of music transcription, and multimodality has been used as a means to break through certain glass-ceiling scenarios attained by single-modality approaches. For instance, research on AMT has contemplated the use of supplementary sources of information, such as harmonic information, timbre, or onset events (Benetos, Dixon, Giannoulis, Kirchoff, & Klapuri, 2013). More recently, the work by de la Fuente et al. (2021) considered that a given score image and its acoustic performance are two different modalities of the same piece to be transcribed. These authors show that transcription results can be enhanced with respect to single-modality systems when their individual performances do not greatly differ.

This work aims to further research the aforementioned avenue of multimodal OMR-AMT transcription. We assume that the score image and the recording of the same piece of music are, in some respects, two complementary sources of information. We follow a late-fusion approach to merge these two modalities, as it allows the more adjustable processing of each of them and does not require multimodal training data for the underlying models. More precisely, we study the existing solutions in related areas in order to then adapt them to the focus of the work: image and audio music transcription.

3. Unimodal neural end-to-end transcription

This section formally presents the neural end-to-end recognition framework for the stand-alone OMR and AMT processes. Considering the aforementioned *sequence labeling* formulation, our goal is to retrieve the most likely series of symbols from the input data—either a score image or an audio recording, as appropriate. Some notations will now be introduced to properly describe these design principles.

Let $\mathcal{T} = \{(x_m, \mathbf{z}_m) : x_m \in \mathcal{X}, \mathbf{z}_m \in \mathcal{Z}\}_{m=1}^{|\mathcal{T}|}$ denote a data collection in which element $x_m \in \mathcal{X}$ corresponds to the sequence $\mathbf{z}_m = (z_{m1}, z_{m2}, \dots, z_{mN}) \in \mathcal{Z}$, related by function $g : \mathcal{X} \rightarrow \mathcal{Z}$. Note that the latter space is defined as $\mathcal{Z} = \Sigma^*$, where Σ stands for the vocabulary of possible music symbols.

We have considered a CRNN scheme with the CTC training algorithm in order to approximate function g owing to its aforementioned competitive performance in the related literature. The CTC method makes it possible to train the CRNN using unsegmented sequential data, i.e., it only requires the different input signals and their associated sequences of symbols from Σ as its expected output, disregarding any input–output alignment. It is important to mention that CTC requires an additional “blank” token within the set of symbols, i.e., $\Sigma' = \Sigma \cup \{\text{blank}\}$, to enable the detection of consecutive repeated elements.

During the inference phase, CTC assumes that the architecture contains a fully-connected network of $|\Sigma'|$ neurons with a *softmax* activation. Although this inference phase can be performed in several ways (Zenkel et al., 2017), the so-called *greedy* decoding is usually considered. In this respect, assuming sequences of length K with $|\Sigma'|$ scores at the output of the recurrent layer—namely *posteriorgram*—this process retrieves the most probable symbol per step, following:

$$\hat{\pi} = \arg \max_{\pi \in \Sigma'^K} \prod_{k=1}^K y_{\pi_k}^k \tag{1}$$

where $y_{\pi_k}^k$ denotes the activation probability of symbol $\pi_k \in \Sigma'$ at time-step k and $\hat{\pi}$ is the retrieved sequence of length K .

A $B(\cdot)$ mapping function that merges consecutive repeated symbols and removes the *blank* label must eventually be applied to $\hat{\pi}$. The predicted sequence is, therefore, obtained as $\hat{z} = B(\hat{\pi})$, where $|\hat{z}| \leq K$.

3.1. Word graphs

While greedy decoding may be sufficient for unimodal recognition, a policy of this nature may be unsuitable for multimodal transcription based on late-fusion approaches. In this case, it would be appropriate to maintain as much information as possible from the original search space—i.e., the collection of possible output sequences from the stand-alone recognition models—for the appropriate integration of the different modalities. In this work, we employ lattice-based *Word Graph* (WG) representations owing to their flexibility and successful use in the related literature (Calvo-Zaragoza, Toselli, Vidal, & Sánchez, 2019; Li, Ness, Ragni, & Gales, 2019; Toselli, Vidal, Romero, & Frinken, 2016).

Definition. A WG is a weighted directed acyclic graph, represented as the 5-tuple $w = (\Sigma, V, v_i, F, E)$, where:

- Σ represents the symbol vocabulary related to the recognition task;
- V is a finite set of vertices or nodes, of which $v_i \in V$ represents the initial vertex of the graph;
- $F \subseteq (V - v_i)$ describes the possible final nodes of the structure; and,
- $E \subset (V - F) \times (V - v_i)$ denotes the finite set of edges, in which each element $e_{A,B} \in E$ —edge between vertices v_A and v_B —is labeled with the symbol $l(e_{A,B}) \in \Sigma$ and is weighted with a score $s(e_{A,B}) \in \mathbb{R}$. Note that this latter value represents the likelihood of symbol $l(e_{A,B})$ appearing between nodes v_A and v_B .

Using the above as a basis, let $\mathbf{t} = (v_1, \dots, v_{|\mathbf{t}|})$ s.t. $v_1 = v_i$ and $v_{|\mathbf{t}|} \in F$ denote a complete path within w and $T = \{\mathbf{t}_j\}_{j=1}^{|\mathcal{T}|}$ be the set of complete paths in it. Additionally, let $\mathbf{l}_\mathbf{t} = (l(e_{n,n+1})_{n=1}^{|\mathbf{t}|-1})$ be the sequence of symbols in \mathbf{t} and $s_\mathbf{t} = \prod_{n=1}^{|\mathbf{t}|-1} s(e_{n,n+1})$ be its corresponding likelihood.

When considering WG representations, the decoding process selects the most appropriate $\mathbf{t} \in T$ on the basis of certain criteria. This process is usually carried out by considering a *best path* strategy: selecting the sequence that maximizes the likelihood, i.e., $\mathbf{t}^* = \arg \max_{\mathbf{t} \in T} s_\mathbf{t}$.

In this work, WG representations are obtained by applying a decoding strategy to the posteriorgram output of the aforementioned CRNN model. This decoding is based on Weighted Finite-State Transducers (WFST) and follows the procedure explained in the Kaldi toolkit (Povey et al., 2011).

4. Late multimodal fusion

Late multimodal fusion has principally been considered in the context of recognition systems owing to its reported benefits in terms of performance improvement. In a broad sense, such approaches typically integrate a group of hypotheses into a consensus hypothesis. This work

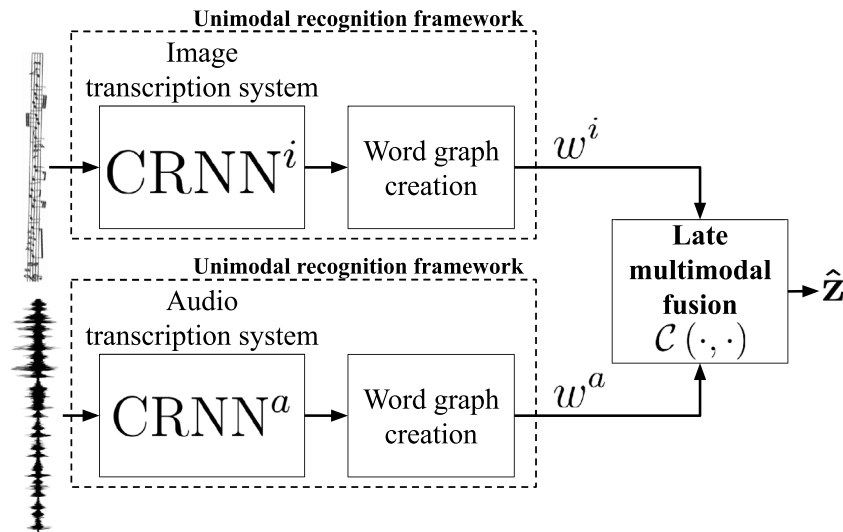


Fig. 2. Graphical example of the proposed late multimodal fusion framework: the hypotheses, w^i and w^a , depicted by the CRNNⁱ image and the CRNN^a audio transcription models, respectively, are combined according to a certain combination function $C(\cdot, \cdot)$. This combination results in the symbol sequence eventually predicted: \hat{z} .

explores the use of decision-level fusion for multimodal neural end-to-end music transcription involving image scores and audio recordings given that the objective in both domains is to retrieve the same symbolic representation. More precisely, we assess the multimodality by combining two respective WG related to an image and an audio recording of the same piece of music. We shall now introduce some notations in order to then formally define this recognition scenario.

Let superscripts i and a identify the image and audio domains, respectively. There are consequently two different representation spaces, \mathcal{X}^i and \mathcal{X}^a , which are related to the image scores and audio signals, respectively, with a single output vocabulary Σ . Furthermore, let CRNNⁱ and CRNN^a denote the two CRNN transcription models introduced in Section 3 that deal with the source domain of the image and audio data, respectively.

Given an input pair (x^i, x^a) representing a score image and an audio signal of the same piece of music, we obtain their respective w^i and w^a WG structures by employing the aforementioned CRNNⁱ and CRNN^a unimodal recognition methods. Function $C(\cdot, \cdot)$, which combines two WG and decodes the resulting structure, is then applied to (w^i, w^a) , and the predicted sequence is obtained as $\hat{z} = C(w^i, w^a)$. A graphic illustration of this process is provided in Fig. 2.

In this work, we consider and assess four different $C(\cdot, \cdot)$ combination approaches: (i) one that addresses the fusion carried out by employing the Minimum Bayes Risk criterion (Xu, Povey, Mangu, & Zhu, 2011), (ii) a second case that addresses the task from a lightly-supervised learning perspective (Fainberg, Klejch, Renals, & Bell, 2019), (iii) a third approach that follows a global alignment strategy (Granell & Martínez-Hinarejos, 2015), and (iv) a last procedure based on local alignment (de la Fuente et al., 2021). Their adaptation and first-time application to multimodal image and audio music transcription are described in the following sections.

4.1. Minimum Bayes risk approach

The work of Xu et al. (2011) presents a lattice combination method for speech recognition based on Minimum Bayes Risk (MBR) decoding. More specifically, the authors propose the combination of individual hypotheses from different systems addressing the same task to improve the stand-alone recognition rate. In this respect, we can adapt this late combination strategy to image and audio music transcription, since CRNNⁱ and CRNN^a share a common representation for their outputs.

In WG representations, the objective of MBR decoding is to find the sequence that minimizes the expected risk for a given loss function.

In our case, the loss function is the string edit distance, and MBR, therefore, seeks the set median string of the distribution provided by the WG. In late-fusion combination scenarios, MBR performs this search by considering the weighted distributions of the WG being fused. In mathematical terms, for the m th sample in set \mathcal{T} , this process is expressed as:

$$\hat{z}_m = \underset{\mathbf{l}_m^* \in \{\mathbf{l}_m^i \cup \mathbf{l}_m^a\}}{\operatorname{argmin}} \left(\alpha \sum_{\mathbf{l}_m^i \in T_m^i} P(\mathbf{l}_m^i | x_m^i) \operatorname{ED}(\mathbf{l}_m^*, \mathbf{l}_m^i) + (1 - \alpha) \sum_{\mathbf{l}_m^a \in T_m^a} P(\mathbf{l}_m^a | x_m^a) \operatorname{ED}(\mathbf{l}_m^*, \mathbf{l}_m^a) \right) \quad (2)$$

where $\alpha \in (0, 1)$ and $(1 - \alpha)$ describe the weights given to the image and audio transcription systems, respectively; sets T_m^i and T_m^a denote the sets of complete paths for the image and audio recognition methods, respectively; $\operatorname{ED}(\cdot, \cdot)$ represents the string edit distance (Levenshtein, 1966); and $P(\mathbf{l}_m^i | x_m^i)$ and $P(\mathbf{l}_m^a | x_m^a)$ denote the respective posterior probabilities of \mathbf{l}_m^i and \mathbf{l}_m^a given the corresponding score image x_m^i and audio recording x_m^a samples, which can be approximately computed by following the work of Toselli et al. (2016) as:

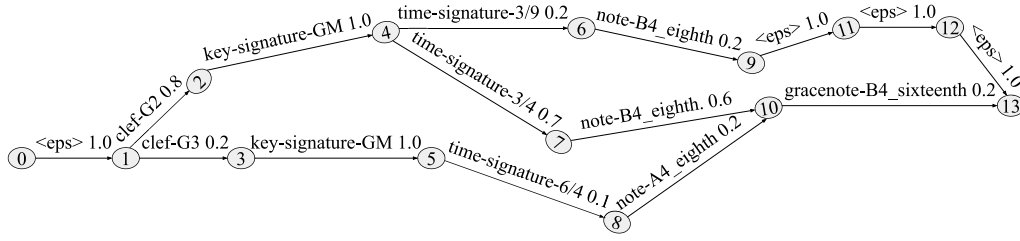
$$P(\mathbf{l}_m | x_m) = \frac{s_{\mathbf{l}_m}}{\sum_{\mathbf{t} \in T} s_{\mathbf{t}}} \quad (3)$$

In order to study the goodness of this multimodal fusion strategy, we consider different weights in order to assess the balance between image and audio domains in the range of the combination parameter $\alpha \in (0, 1)$.

4.2. Lightly-supervised approach

The second late multimodal fusion technique considered was proposed by Fainberg et al. (2019). These authors devised a method to correct inaccurate transcriptions by combining them with lattices depicted by a seed recognition model, namely *lightly-supervised* training. In this framework, a transcribed symbol is considered to be correctly estimated if, and only if, it is also present in the lattice of the recognition model.

The method fundamentally depends on a function $f_{\text{collapse}}(\cdot, \cdot)$ that collapses a WG—the compact representation of the hypothesis space—into a sequence of symbols (the transcription) when the aforementioned assumption is met. This results in a new WG with a reduced number of vertices/edges in the matching areas, and remains unaltered otherwise.



(a) Search space as word graph.



(b) Search space as confusion network.

Fig. 3. Search space representation used in the proposed late multimodal fusion framework: Fig. 3(a) represents a given w and Fig. 3(b) its corresponding c . The symbol sequence “<eps> clef-G2 keySignature-GM timeSignature-3/4 note-B4_eighth. gracenote-B4_sixteenth <eps>” is the best path (Γ^*) in both lattices.

We present the adaptation of this decision level multimodal policy to image and audio music transcription in Algorithm 1. Given an input pair, $(x^{(1)}, x^{(2)})$, describing two modality-specific representations of the same source sample, we consider the greedy-decoded hypothesis, $\mathbf{t}^{(1)}$, depicted by CRNN⁽¹⁾, as the transcription to be corrected using the hypothesis, $w^{(2)}$, which is yielded by the domain-counterpart model, CRNN⁽²⁾. In this work, the two modalities correspond to the image and audio domains. We, therefore, evaluate the technique presented when assuming model correctness in (i) the image domain, i.e., $w^{(2)} = w^I$, and (ii) the audio domain, $w^{(2)} = w^A$.

Algorithm 1: Lightly-supervised late multimodal fusion.

Input : $w^{(1)} \leftarrow$ First WG
 $w^{(2)} \leftarrow$ Second WG
Output: $\hat{z} = (\hat{z}_1, \dots, \hat{z}_{|\hat{z}|}) \leftarrow$ Retrieved sequence
1 $\mathbf{t}^{*(1)} \leftarrow l(t) \forall t \in \arg \max_{t \in T^{(1)}} s_t \quad \triangleright$ Best path of $w^{(1)}$
2 $w^{(2)} \leftarrow f_{\text{collapse}}(\mathbf{t}^{*(1)}, w^{(2)})$
3 $\hat{z} \leftarrow l(t) \forall t \in \arg \max_{t \in T^{(2)}} s_t \quad \triangleright$ Best path of $w^{(2)}$

4.3. Global alignment approach

The work of Granell and Martínez-Hinarejos (2015) presents a late multimodal fusion approach based on CN with which to combine the outputs of Handwritten Text Recognition (HTR) and ASR systems so as to transcribe handwritten documents. This method can be directly adapted to the image and audio music transcription task by making HTR correspond with OMR and ASR correspond with AMT.

CN is a different topology with which to represent lattice-type spaces. Given w , we obtain its corresponding CN, denoted as c , by following a clustering process that identifies mutually supporting and competing hypotheses in w , and establishing a total order for all of them (Mangu, Brill, & Stolcke, 1999). The resulting c is also a weighted directed acyclic graph with three particularities: (i) every path $\mathbf{t} \in T$ from v_i to v_f passes through all the nodes, i.e., $|\mathbf{t}| = |V| - 1 \forall \mathbf{t} \in T$, where v_f is the final vertex of c ; (ii) a subnetwork, denoted as $N = \{e_{A,B}\} \subseteq E$, where $|A - B| = 1$, is the set of all the edges between two consecutive nodes, and (iii) $\sum_{n \in N} s(n) = 1$, i.e., the total score of the edges contained in N adds up to 1. Meeting the first particularity implies dealing with hypotheses of different lengths from T . In this respect, one edge (at most), labeled with a special symbol <eps>, can be inserted into N . The inclusion of this symbol results in c adding paths that are not present in w . Fig. 3 provides a graphical representation of a WG together with its associated CN.

Given c^i and c^a , the CN-based combination technique seeks to create a new \hat{c} in which the error level is reduced. In a first step, c^i and c^a are aligned by means of similarity on the basis of a matching error using the Dynamic Time Warping global alignment algorithm (Müller, 2007), where the cost function is the normalized string edit distance between the given the symbols of N^i and N^a . In a second and final step, a new \hat{c} is created by merging subnetworks from the aligned c^i and c^a following:

$$s(n) = s_{\text{smoothed}}(n^i)^\alpha s_{\text{smoothed}}(n^a)^{1-\alpha} \quad (4)$$

where $\alpha \in (0, 1)$ and $(1 - \alpha)$ respectively describe the weights of the image and audio transcription systems; and $s_{\text{smoothed}}(n^i)$ and $s_{\text{smoothed}}(n^a)$ denote the Laplacian smoothed (Zhai & Lafferty, 2004) score of edges $n^i \in N^i$ and $n^a \in N^a$, respectively.

Note that if, and only if, the normalized string edit distance equals 1—the given symbols of N^i and N^a are completely different— N^i and N^a are, respectively, combined with a special subnetwork, \bar{N} , that contains only one unit-weighted edge labeled with <eps>. When this condition is not met, N^i and N^a are combined with each other. The final estimation, \hat{z} , is obtained by retrieving the best path, \mathbf{t}^* , from \hat{c} .

In order to study the balance between the relative reliability of OMR and AMT, we assess the goodness of the late multimodal technique presented when $\alpha \in (0, 1)$.

4.4. Local alignment approach

The last decision level multimodal policy considered in the present work was proposed by de la Fuente et al. (2021). These authors combined the predictions provided by end-to-end AMT and OMR systems by considering the Smith–Waterman (SW) local alignment algorithm (Smith & Waterman, 1981).

This method is presented in Algorithm 2. The inputs of the method were originally the hypotheses depicted by each single-modality model decoded by means of Eq. (1). However, in order to maintain consistency with the methodology presented in the current work, we consider the corresponding best paths \mathbf{t}^{*i} and \mathbf{t}^{*a} to be analogous instances of the former greedy-decoded hypotheses. In preliminary experiments carried out for the current work, no differences were found in the recognition results when assuming that this was the case. Given that the \mathbf{t}^{*i} and \mathbf{t}^{*a} may depict different lengths, the Smith–Waterman (SW) local alignment algorithm (Smith & Waterman, 1981) is, therefore, used to align both estimations by searching for the most similar regions between them. The final estimation is eventually obtained from these two aligned sequences, $\hat{\mathbf{t}}^{*i}$ and $\hat{\mathbf{t}}^{*a}$, by following these premises: (i)

if both sequences match a symbol, it is included in the resulting estimation; (ii) if they disagree, the symbol with the highest score is included in the estimation, and (iii) if one of the sequences lacks a symbol, that of the other sequence is included in the estimation.

Algorithm 2: Local alignment late multimodal fusion.

```

Input :  $w^i \leftarrow$  Image WG
          $w^a \leftarrow$  Audio WG
Output:  $\hat{z} \leftarrow (\hat{z}_1, \dots, \hat{z}_{|\hat{z}|})$   $\triangleright$  Retrieved sequence
1  $t^{*i} \leftarrow \arg \max_{t \in T^i} S_t$   $\triangleright$  Best path of  $w^i$ 
2  $t^{*a} \leftarrow \arg \max_{t \in T^a} S_t$   $\triangleright$  Best path of  $w^a$ 
3  $\bar{t}^{*i}, \bar{t}^{*a} \leftarrow$  SW( $t^{*i}, t^{*a}$ )  $\triangleright$  SW local alignment
4  $K \leftarrow |\bar{t}^{*i}|$   $\triangleright$  Sequence length:  $|\hat{z}| = |\bar{t}^{*i}| = |\bar{t}^{*a}|$ 
5 for  $k \in [1, \dots, K]$  do
6   if  $l(\bar{t}_k^{*i}) = "" \vee l(\bar{t}_k^{*a}) = ""$  then
7      $\hat{z}_k \leftarrow l(\bar{t}_k^{*i}) + l(\bar{t}_k^{*a})$   $\triangleright$  No symbol
8   else if  $l(\bar{t}_k^{*i}) = l(\bar{t}_k^{*a})$  then
9      $\hat{z}_k \leftarrow l(\bar{t}_k^{*i})$   $\triangleright$  Same symbol
10  else
11     $\hat{z}_k \leftarrow l\left(\arg \max_{\bar{t}_k^s = \{\bar{t}_k^{*i}, \bar{t}_k^{*a}\}} S(\bar{t}_k^s)\right)$ 
12  end if
13 end for

```

The same alignment parameters as those employed in the reference work (de la Fuente et al., 2021) have been used for the SW algorithm for comparative purposes.

5. Experimental setup

This section presents the definition of the different layers of the neural models, the evaluation protocol used, and the precise evaluation scenarios and corpus considered.

5.1. Neural network configuration

The choice of the CRNN topologies for image and audio music transcription systems is generally conditioned to the particular corpus considered, the amount of accessible data, or the computational resources available, among others. Nevertheless, these configurations usually contain a set of convolutional layers for the feature extraction process, followed by recurrent units for the dependency modeling, with a last dense unit with $|\Sigma^i|$ output units. In this work, the actual composition of each layer, depicted in Table 1, is based on that used in recent works tackling the individual OMR and AMT tasks by framing them in a sequence labeling scenario (Calvo-Zaragoza & Rizo, 2018; Liu & Benetos, 2021).

All the models were trained using the backpropagation method provided by CTC for 150 epochs using the ADAM optimizer (Kingma & Ba, 2015) with a fixed learning rate of 0.001. The batch size was fixed to 16 for the OMR system, while in the case of the AMT it was set to 4 because it is more memory-intensive.

5.2. Performance metrics

The performance of the recognition schemes presented is assessed by considering the Symbol Error Rate (SER), as occurred in previous works addressing end-to-end transcription tasks (Calvo-Zaragoza & Rizo, 2018; Román, Pertusa, & Calvo-Zaragoza, 2019). This figure of merit is computed as the average number of elementary editing operations (insertions, deletions, or substitutions) required in order to match the sequence predicted by the model with that in the ground

truth, normalized by the length of the latter. In mathematical terms, this is expressed as:

$$\text{SER} (\%) = \frac{\sum_{m=1}^{|\mathcal{S}|} \text{ED}(\hat{z}_m, z_m)}{\sum_{m=1}^{|\mathcal{S}|} |z_m|} \quad (5)$$

where $S \subset \mathcal{X} \times \mathcal{Z}$ is a set of test data—from either the image or the audio domains—, $\text{ED} : \mathcal{Z} \times \mathcal{Z} \rightarrow \mathbb{N}_0$ denotes the string edit distance, and \hat{z}_m and z_m respectively represent the estimated and target sequences.

5.3. Evaluation scenarios

We aim to provide insights into how the differences among the performances of the stand-alone transcription models affect the outcome of the combined paradigm. The definition of the evaluation scenarios is set according to different ranges of SER. We specifically consider three possible levels:

- *High*, which refers to a SER of approximately 30%. We consider this performance threshold since ineffective configurations in the related literature report error values similar to this figure. Cases above this level are, therefore, of no interest.
- *Low*, which denotes a SER of around 10%. As with the previous case, state-of-the-art transcription methods report error rates within this range, and figures lower than this threshold are, therefore, omitted from this study.
- *Medium*, which stands for an approximate SER of 20%. This intermediate error level has, therefore, been considered for reasons of the completeness of the results.

In this respect, we have considered the Camera-based Printed Images of Music Staves (Camera-PrIMuS) database (Calvo-Zaragoza & Rizo, 2018). This corpus contains 87,678 real music staves of monophonic incipits¹ extracted from the *Répertoire International des Sources Musicales* (RISM).² Different representations are provided for each incipit: an image with the rendered score (both plain and with artificial distortions), several encoding formats, and a MIDI file.

Each transcription architecture considers a particular type of data: on the one hand, the OMR model takes the artificially distorted staff image of the incipit as input, and on the other, each MIDI file in the AMT case is synthesized using the FluidSynth software³ and a piano timbre, considering a sampling rate of 22,050 Hz. A time-frequency representation based on the Constant-Q Transform was obtained, with a hop length of 512 samples, 120 bins, and 24 bins per octave, which is eventually embedded as an image that serves as the input. The height of the input considered is scaled to 64 pixels for image data, or to 256 pixels for audio data, maintaining the aspect ratio (signifying that each sample might differ in width) and converted to grayscale, with no further preprocessing.

Since this corpus was originally devised for OMR tasks, a data cleansing process was carried out to adapt it to the multimodal transcription framework presented, resulting in 22,285 incipits.⁴ We derive three non-overlapping partitions—train, validation, and test—corresponding to 60%, 20%, and 20% of the data, respectively, following a 5-fold cross-validation scheme. Note that, since the same corpus is considered for both image and audio data, both recognition tasks depict the same label space of $|\Sigma^i| = |\Sigma^a| = 1166$ tokens.

When preliminary assessing OMR and AMT models for the original partitions, the resulting error levels are framed within the *High* and

¹ Short sequences of notes, typically the first measures of the piece, used to index and identify a melody or musical work.

² <https://rism.info/>

³ <https://www.fluidsynth.org/>

⁴ Mainly, samples containing long multi-rests were removed as they barely extend the length of the score image but may span for large number of frames in the audio signal.

Table 1

Layer-wise description of the CRNN models considered. Notation: Conv($f, w \times h$) represents a convolution layer of f filters of size $w \times h$ pixels, BatchNorm performs the normalization of the batch, LeakyReLU(α) represents a Leaky Rectified Linear Unit activation with a negative slope of value α , MaxPool($w \times h$) represents the max-pooling operator of $w \times h$ dimensions and striding factors, BLSTM(n, d) denotes a bidirectional Long Short-Term Memory unit with n neurons and d dropout value parameters, Dense(n) is a fully-connected layer of n neurons, and Softmax(\cdot) represents the softmax activation. Σ' denotes the alphabet considered, including the CTC-blank symbol.

	Layer 1	Layer 2	Layer 3	Layer 4	Layer 5	Layer 6	Layer 7
OMR	Conv(64, 5 × 5) BatchNorm LeakyReLU(0.20) MaxPool(2 × 2)	Conv(64, 5 × 5) BatchNorm LeakyReLU(0.20) MaxPool(1 × 2)	Conv(128, 3 × 3) BatchNorm LeakyReLU(0.20) MaxPool(1 × 2)	Conv(128, 3 × 3) BatchNorm LeakyReLU(0.20) MaxPool(1 × 2)	BLSTM(256) Dropout(0.50)	BLSTM(256) Dropout(0.50)	Dense(Σ') Softmax(\cdot)
AMT	Conv(8, 2 × 10) BatchNorm LeakyReLU(0.20) MaxPool(2 × 2)	Conv(8, 5 × 8) BatchNorm LeakyReLU(0.20) MaxPool(1 × 2)	BLSTM(256) Dropout(0.50)	BLSTM(256) Dropout(0.50)	Dense(Σ') Softmax(\cdot)		

Low levels, respectively. To obtain the other cases, it is necessary to modify the data partitions used in each transcription model. Broadly speaking, since OMR has a base performance framed in the *Low* level, it is possible to obtain the other two cases by simply reducing the train partition; in the case of AMT, however, it is necessary to remove certain incipits from the test partition—those that greatly hinder the performance—as regards achieving the remaining cases. Please note that, since the late-fusion framework addressed in the work requires the same test set for both modalities, the test partition is constrained by AMT.

Considering the design principles presented, the approximate resulting SER rates for the different individual cases are a SER of 27% for the *High* level, a SER of 17% for the *Medium* level, and a SER of 7% for the *Low* level. The combination of those performance standards results in the nine controlled evaluation scenarios used in this work.

6. Results

This section presents the results obtained for the proposed experimental scheme. To aid in their evaluation, the following notation is introduced for the four late multimodal fusion methods and the nine scenarios considered. On the one hand, *MBR* is used for the MBR-decoding approach (Section 4.1); *Lightly_{IA}* and *Lightly_{AI}* correspond to the lightly-supervised scenario, when assuming model correctness from the audio domain, i.e., the to-be-corrected transcriptions originate from image one and vice versa, respectively (Section 4.2); *Global* is employed for the global alignment strategy based on CN (Section 4.3), and finally, *Local* corresponds to the SW local alignment procedure (Section 4.4). On the other hand, each scenario is denoted as Scenario $N_{(OMR \text{ Level Performance} - AMT \text{ Level Performance})}$, where N denotes the number of the given scenario out of the nine that are possible and *OMR Level Performance* and *AMT Level Performance* refer to the three possible levels of error rate considered—High (H), Medium (M), and Low (L)—for each unimodal model, respectively.⁵

The results obtained in terms of the SER metric for the contemplated scenarios are presented in Fig. 4 and Table 2. Since the experiments have been performed in a cross-validation scheme, the results presented for each of the cases considered correspond to the average values for the test partition in which the validation data achieved the best performance.

With regard to the aforementioned results, it is first necessary to state that the image and audio music transcription task can be successfully solved by using a late multimodal fusion approach, as some of the proposed methods employed in the present work achieved lower error rates than their corresponding unimodal baselines for all scenarios. This outcome supports the initial premise that both a recording and an image score can act as two complementary modalities of the same piece

⁵ The code developed in the work is publicly available for reproducible research at: <https://github.com/mariaalfaroc/late-fusion-music-transcription>. git.

Table 2

Best results obtained by the proposed unimodal and multimodal frameworks in terms of the Symbol Error Rate (SER) for each scenario considered. For a given scenario, underlining indicates that the corresponding late multimodal fusion approach improves both unimodal baselines and the boldface denotes that it achieves the lowest SER of all the other multimodal methods that also improve the unimodal methods.

	Scenario								
	1 _(H-H)	2 _(H-M)	3 _(H-L)	4 _(M-H)	5 _(M-M)	6 _(M-L)	7 _(L-H)	8 _(L-M)	9 _(L-L)
OMR	26.8	27.3	26.2	18.4	17.3	16.9	8.1	7.4	7.5
AMT	26.3	16.9	7.2	26.3	16.9	7.2	26.3	16.9	7.2
MBR	<u>23.3</u>	<u>16.8</u>	<u>7.2</u>	<u>17.4</u>	<u>14.4</u>	<u>7.1</u>	8.0	7.2	<u>5.8</u>
Lightly _{IA}	25.8	17.1	9.1	25.4	16.8	8.9	24.9	16.5	9.1
Lightly _{AI}	28.5	28.4	28.0	20.7	19.3	19.6	10.9	10.3	11.4
Global	<u>24.6</u>	<u>15.9</u>	<u>6.3</u>	<u>18.2</u>	<u>15.0</u>	<u>5.8</u>	<u>8.1</u>	<u>7.3</u>	<u>5.5</u>
Local	20.8	13.8	6.1	18.2	11.1	4.5	14.5	8.1	3.0

of music, since its multimodal combination yields better recognition rates than if performed in an isolated manner. We shall now examine each method in detail.

As will be noted in both Fig. 4 and Table 2, MBR is capable of retrieving improved transcription results while disregarding the actual scenario considered. In relation to the combination coefficient, α , there are three distinct situations that are consistent with the differences in the error rate of the models: (i) if both transcription systems are at the same performance level—Scenarios 1, 5, and 9—, there is an almost equal distribution of the weight, slightly shifted towards AMT; (ii) if the image system depicts considerably lower error rates than its audio counterpart—Scenarios 4 and 8—, then the inflection point of the previous weight distribution is now shifted towards the former model, especially when they are at opposite performance levels—Scenario 7—; and finally, (iii) if the audio system is that which works better—Scenarios 2, 3, and 6—, the information provided by that model has a high weight, of at least 80%, in the combination.

The Lightly method combines an inaccurate transcription—the best-path decoded hypothesis of one modality—with the WG depicted by the seed model—from the other modality. Not surprisingly, the best results provided by this approach are obtained when the corrector is the modality model with the lowest error rate: if the CRNNⁱ image model performs better than its audio counterpart, CRNN^a, Lightly_{AI} should be used rather than Lightly_{IA}, and vice versa. However, both approaches appear to be unsuitable for the image and audio music transcription task as they are generally able only to improve the recognition rate of one of the domain-specific models. On the one hand, Lightly_{IA} is able to enhance transcription results if, and only if, both unimodal baselines depict either a High or a Medium error level. When this condition is not met, it is able only to improve the recognition rate of the worst-performing single-modality transcription system. On the other hand, Lightly_{AI} boosts the audio domain results, at best, when CRNNⁱ is at a lower error rate level than CRNN^a.

The late multimodal fusion policy proposed by Global behaves in a synergistic manner, as it yields higher performance rates in all scenarios than its corresponding unimodal frameworks. Fig. 4 shows

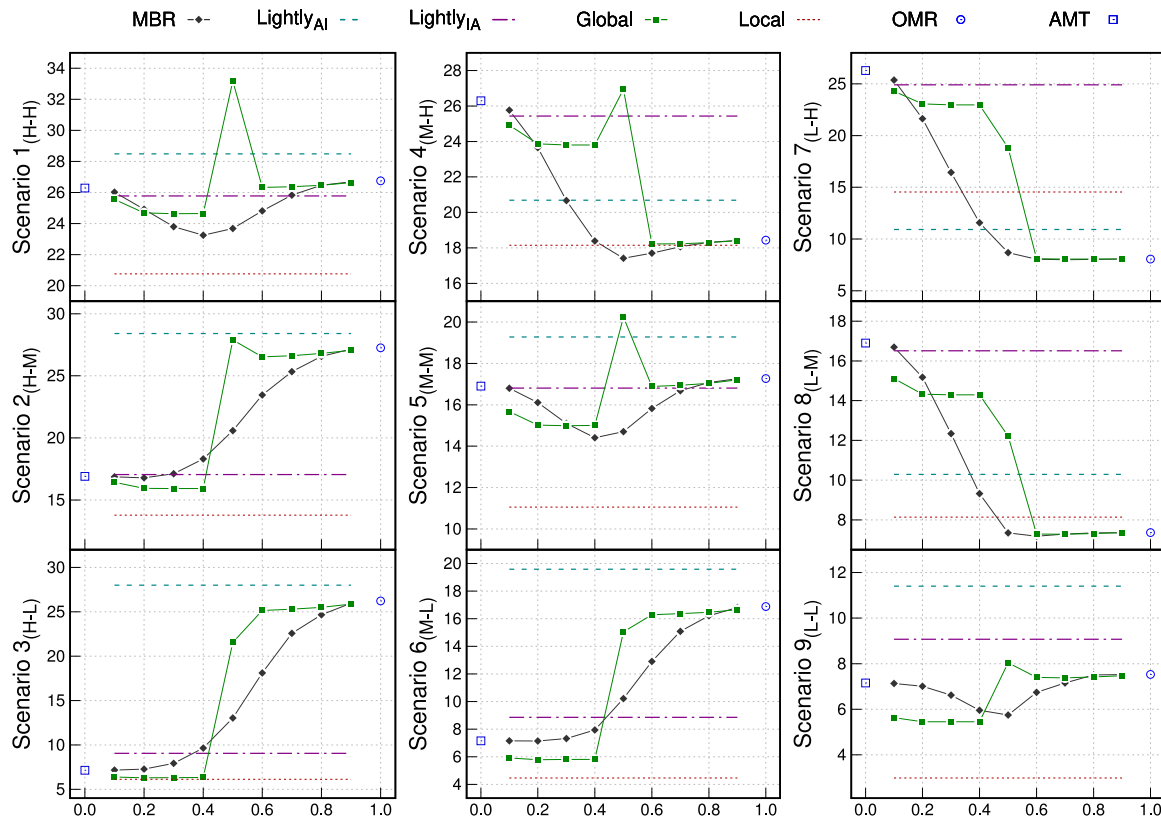


Fig. 4. Results obtained by the proposed unimodal and multimodal frameworks in terms of the Symbol Error Rate (SER) for each scenario considered when $\alpha \in (0,1)$. Please note that α is the combination coefficient that allows the weight of each of the unimodal systems involved in the late multimodal fusion to be balanced. α and $(1 - \alpha)$ specifically describe the weights given to the image and audio transcription systems. $\alpha = 1$ and $\alpha = 0$ are, therefore, considered to represent the SER figures corresponding to the image and audio domains, respectively.

two clear situations related to the balance between the reliability of the image and audio domains, controlled by the α coefficient. In this respect, when either both unimodal transcription systems are at the same performance level or the audio one has a lower error rate, 70% of the weight is given to the aforementioned modality. This distribution is reversed when the CRNNⁱ image recognition system outperforms its domain counterpart model, CRNN^a.

Finally, with regard to Local, the results indicate that the actual improvement achieved by this method depends on the scenario considered. In this respect, this method is not the best multimodal framework when the image system is at a Low error rate level and the audio system is at one of the other two levels (Scenarios 7 and 8). In these cases, it manages only to reduce the error rate of CRNN^a. However, it is necessary to state that when Local does improve the unimodal recognition rates, it is the best method 90% of the time, obtaining improvements ranging from 14% to 58% with respect to the lowest unimodal error rate. Please note that this upper bound of improvement is obtained when both music transcription systems are at the same performance level. This suggests some sort of synergistic behavior in the fusion policy as the resulting sequence takes the most accurate estimations of the OMR and AMT transcription methods.

6.1. Statistical significance analysis

The analysis shown above can be summarized in the following conclusions:

- MBR and Global provide a general improvement with respect to the standard unimodal framework. When confronting both methods, the former is a better approach than the latter: it yields higher recognition rates in more than half of the scenarios, and is the best performing late multimodal policy in a third of them. However, the differences between them are not remarkable.

Table 3

Statistical significance analysis of the base neural unimodal models of the multimodal fusion schemes for the Wilcoxon signed-rank test when considering a significance value of $p < 0.05$ for the Symbol Error Rate metric. Symbols ✓, ×, and = represent that the error of the method in the row is significantly lower than, greater than, or no different to that in the column, respectively.

	MBR	Lightly _{IA}	Lightly _{AI}	Global	Local
OMR	×	=	✓	×	=
AMT	×	=	=	×	×

- Lightly_{IA} improves both unimodal music transcription systems if, and only if, both of them are at the same error level and that level is not Low. However, Lightly_{AI} fails to decrease the error rate of both single-modality models in any scenario.
- The Local method improves both domain-specific baselines whenever CRNNⁱ is not at a Low level, because if this is the case, it requires CRNN^a to also be at this error level in order to do so. Note that, when possible, it provides the greatest improvement of all the multimodal decision level strategies 90% of the time.

In order to support the relevance of the statements above, we shall now statistically assess the results obtained. This will be done using the non-parametric Wilcoxon signed-rank test (Demšar, 2006). This analysis states that each result obtained for each scenario forms a sample of the distributions to be compared. The results obtained are shown in Tables 3 and 4.

The results obtained with a significance value of $p < 0.05$ provide a response to two main questions: (i) is the use of late multimodal fusion frameworks worthwhile?, and (ii) if so, which approach performs best? On the one hand, it is possible to respond to the first by studying Table 3, which shows that both MBR and Global significantly

Table 4

Statistical significance analysis of the late multimodal fusion schemes considering the Wilcoxon signed-rank test with a significance value of $p < 0.05$ for the Symbol Error Rate metric. Symbols ✓, ×, and = represent that the error of the method in the row is significantly lower than, greater than, or no different to that in the column, respectively.

	MBR	Lightly _{IA}	Lightly _{AI}	Global	Local
MBR	–	✓	✓	=	=
Lightly _{IA}	×	–	=	×	×
Lightly _{AI}	×	=	–	×	×
Global	=	✓	✓	–	=
Local	=	✓	✓	=	–

outperform both image and audio music transcription systems. It also shows that, while Lightly_{IA} depicts results that are statistically equal to those of both unimodal frameworks, Lightly_{AI} worsens this situation by yielding considerably higher error rates than the OMR model. Finally, Local considerably improves the audio transcription system without hindering the image system. On the other hand, the second question can be answered by studying the results depicted in Table 4, which shows that there are no meaningful differences among MBR, Global, and Local, but there are with respect to Lightly, in any of its forms, since the three former methods significantly outperform the latter.

Overall, it is now possible to state that the use of late multimodal fusion frameworks is worthwhile for image and audio music transcription, since three out of the four proposed approaches statistically improve at least one of the domain-specific recognition systems without a decline in the performance of the other.

7. Conclusions

The transcription of music sources into a structured digital format is one of the key challenges in the Music Information Retrieval (MIR) field. This problem is typically divided into two lines of research: Optical Music Recognition (OMR), when considering visual input data such as scores, and Automatic Music Transcription (AMT), when considering acoustic input data such as audio signals. While these fields have historically evolved separately, their recent definition within a sequence labeling formulation results in a common representation for their expected outputs, hence enabling being addressed within a multimodal recognition framework.

In this work, we present the first application of diverse existing late multimodal fusion approaches to related areas so as to solve the multimodal image and audio music transcription task. We specifically consider four combination functions in order to merge the hypotheses, which are formatted as word graphs, depicted by individual OMR and AMT systems. The results obtained with monophonic music data in a series of performance scenarios—in which the corresponding single-modality models yield different error rates—provided statistical evidence for the benefits of these approaches, since two of the four strategies considered significantly improved the corresponding unimodal standard recognition frameworks.

As future work, we plan to follow different research avenues. For instance, while the approaches described in this paper work at the decision level, it would be interesting to additionally explore early fusion. Experimentation may also be extended to more challenging data, such as addressing other instrument timbres, handwritten scores, or homophonic and polyphonic music.

CRedit authorship contribution statement

María Alfaro-Contreras: Conceptualization, Investigation, Methodology, Software, Validation, Writing – original draft, Writing – review & editing. **Jose J. Valero-Mas:** Conceptualization, Investigation, Methodology, Software, Validation, Writing-reviewing. **José M. Iñesta:** Investigation, Supervision, Writing – review & editing. **Jorge Calvo-Zaragoza:** Investigation, Supervision, Writing – review & editing.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: María Alfaro-Contreras reports financial support was provided by Spain Ministry of Universities. Jorge Calvo-Zaragoza reports financial support was provided by Spain Ministry of Science and Innovation. Jose J. Valero-Mas reports financial support was provided by Government of Valencia Ministry of Innovation Universities Science and Society.

Data availability

The data used is public.

Acknowledgments

This paper is part of the I+D+i PID2020-118447RA-I00 (Multi-Score) project, funded by MCIN/AEI/10.13039/501100011033. Some of the computing resources were provided by the Generalitat Valenciana and the European Union through the FEDER funding programme (IDIFEDER/2020/003). The first and second authors are respectively supported by grants FPU19/04957 from the Spanish Ministerio de Universidades and APOSTD/2020/256 from Generalitat Valenciana.

References

- Alfaro-Contreras, M., Valero-Mas, J. J., Iñesta, J. M., & Calvo-Zaragoza, J. (2022). Insights into transfer learning between image and audio music transcription. In *Proceedings of the 19th Sound and Music Computing Conference* (pp. 292–298).
- Benetos, E., Dixon, S., Duan, Z., & Ewert, S. (2018). Automatic music transcription: An overview. *IEEE Signal Processing Magazine*, 36, 20–30.
- Benetos, E., Dixon, S., Giannoulis, D., Kirchhoff, H., & Klapuri, A. (2013). Automatic music transcription: challenges and future directions. *Journal of Intelligent Information Systems*, 41, 407–434.
- Calvo-Zaragoza, J., Hajič, J., Jr., & Pacha, A. (2020). Understanding optical music recognition. *ACM Computing Surveys*, 53, 1–35.
- Calvo-Zaragoza, J., & Rizo, D. (2018). Camera-PrIMus: Neural end-to-end optical music recognition on realistic monophonic scores. In *Proceedings of the 19th International Society for Music Information Retrieval Conference* (pp. 248–255).
- Calvo-Zaragoza, J., Toselli, A. H., Vidal, E., & Sánchez, J. A. (2019). Music symbol sequence indexing in medieval plainchant manuscripts. In *International Conference on Document Analysis and Recognition* (pp. 882–887). IEEE.
- de la Fuente, C., Valero-Mas, J. J., Castellanos, F. J., & Calvo-Zaragoza, J. (2021). Multimodal image and audio music transcription. *International Journal of Multimedia Information Retrieval*, 1–8.
- Demšar, J. (2006). Statistical comparisons of classifiers over multiple data sets. *Journal of Machine Learning Research*, 7, 1–30.
- Dumas, B., Signer, B., & Lalanne, D. (2012). Fusion in multimodal interactive systems: an HMM-based algorithm for user-induced adaptation. In *Proceedings of the 4th ACM Symposium on Engineering Interactive Computing Systems* (pp. 15–24).
- Fainberg, J., Klejch, O., Renals, S., & Bell, P. (2019). Lattice-based lightly-supervised acoustic model training. In *Interspeech 20th Annual Conference of the International Speech Communication Association* (pp. 1596–1600).
- Granel, E., & Martínez-Hinarejos, C.-D. (2015). Combining handwriting and speech recognition for transcribing historical handwritten documents. In *13th International Conference on Document Analysis and Recognition* (pp. 126–130).
- Granel, E., Martínez-Hinarejos, C. D., & Romero, V. (2018). Improving transcription of manuscripts with multimodality and interaction. In *Proceedings of IberSPEECH* (pp. 92–96).
- Graves, A., Fernández, S., Gomez, F., & Schmidhuber, J. (2006). Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the 23rd International Conference on Machine Learning* (pp. 369–376). New York, NY, USA: ACM.
- Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In Y. Bengio, & Y. LeCun (Eds.), *3rd International Conference on Learning Representations*.
- Kristensson, P. O., & Vertanen, K. (2011). Asynchronous multimodal text entry using speech and gesture keyboards. In *Proceedings of the International Conference on Spoken Language Processing* (pp. 581–584).
- Levenshtein, V. I. (1966). Binary codes capable of correcting deletions, insertions, and reversals. *Soviet Physics Doklady*, 10, 707–710.
- Li, Q., Ness, P., Ragni, A., & Gales, M. J. (2019). Bi-directional lattice recurrent neural networks for confidence estimation. In *IEEE International Conference on Acoustics, Speech and Signal Processing* (pp. 6755–6759). IEEE.

- Liu, L., & Benetos, E. (2021). From audio to music notation. In *Handbook of artificial intelligence for music* (pp. 693–714). Springer.
- Liu, L., Morfi, V., & Benetos, E. (2021). Joint multi-pitch detection and score transcription for polyphonic piano music. In *IEEE international conference on acoustics, speech and signal processing* (pp. 281–285). IEEE.
- Mangu, L., Brill, E., & Stolcke, A. (1999). Finding consensus among words: lattice-based word error minimization. In *Eurospeech*. Citeseer.
- Miki, M., Kitaoka, N., Miyajima, C., Nishino, T., & Takeda, K. (2014). Improvement of multimodal gesture and speech recognition performance using time intervals between gestures and accompanying speech. *EURASIP Journal on Audio, Speech, and Music Processing*, 2014, 1–7.
- Müller, M. (2007). Dynamic time warping. *Information Retrieval for Music and Motion*, 69–84.
- Pitsikalis, V., Katsamanis, A., Theodorakis, S., & Maragos, P. (2017). Multimodal gesture recognition via multiple hypotheses rescoring. In *Gesture recognition* (pp. 467–496). Springer.
- Povey, D., Ghoshal, A., Boulianne, G., Burget, L., Glembek, O., Goel, N., et al. (2011). The kaldi speech recognition toolkit. In *IEEE Workshop on Automatic Speech Recognition and Understanding*. IEEE Signal Processing Society.
- Rebelo, A., Fujinaga, I., Paszkiewicz, F., Marçal, A., Guedes, C., & Cardoso, J. (2012). Optical music recognition: State-of-the-art and open issues. *International Journal of Multimedia Information Retrieval*, 1.
- Román, M. A., Pertusa, A., & Calvo-Zaragoza, J. (2019). A holistic approach to polyphonic music transcription with neural networks. In *Proceedings of the 20th International Society for Music Information Retrieval Conference* (pp. 731–737).
- Serra, X., Magas, M., Benetos, E., Chudy, M., Dixon, S., Flexer, A., et al. (2013). *Roadmap for music information research*. The MIReS Consortium.
- Simonetta, F., Ntalampiras, S., & Avanzini, F. (2019). Multimodal music information processing and retrieval: Survey and future challenges. In *International Workshop on Multilayer Music Representation and Processing* (pp. 10–18). IEEE.
- Singh, A., Sangwan, A., & Hansen, J. H. (2012). Improved parcel sorting by combining automatic speech and character recognition. In *IEEE International Conference on Emerging Signal Processing Applications* (pp. 52–55). IEEE.
- Smith, T. F., & Waterman, M. S. (1981). Identification of common molecular subsequences. *Journal of Molecular Biology*, 147, 195–197.
- Toselli, A. H., Vidal, E., & Casacuberta, F. (2011). *Multimodal Interactive Pattern Recognition and Applications*. Springer Science & Business Media.
- Toselli, A. H., Vidal, E., Romero, V., & Frinken, V. (2016). Hmm word graph based keyword spotting in handwritten document images. *Information Sciences*, 370, 497–518.
- Xu, H., Povey, D., Mangu, L., & Zhu, J. (2011). Minimum Bayes risk decoding and system combination based on a recursion for edit distance. *Computer Speech and Language*, 25, 802–828.
- Zenkel, T., Sanabria, R., Metze, F., Niehues, J., Sperber, M., Stüker, S., et al. (2017). Comparison of decoding strategies for CTC acoustic models. In *Proc. INTERSPEECH* (pp. 513–517).
- Zhai, C., & Lafferty, J. (2004). A study of smoothing methods for language models applied to information retrieval. *ACM Transactions on Information Systems*, 22, 179–214.