

# Musicalising sonification: Image-to-music conversion using OpenMusic

Luiz Castelões

David Cope [4] defines *sonification* as the production of sound “by translating (called normalising), in one way or another, traditionally non-audible data to the audible frequency range.” He goes on to say that “[s]ince the data typically used for data sonification has little if any inherent musical logic, the normalised output from sonification generally *has little musical use*” (p. xi, italics are mine).

Departing from Cope’s definition, as sonification becomes more and more widespread and trivial (concerning the technologies involved) the compositional challenge at hand shifts from the act of converting between different realms (image-to-sound, number-to-sound), which characterises any given sonification, to effectively obtaining consistent, original musical results through such operations.

Strategies for turning sonification into music (and finding possible intersections between musical logic and other logics in the context of sonification) have been a major field of interest for me as a composer and scholar and will constitute the main focus of this article, particularly with regard to image-to-music conversions.

## Introduction

Image-to-sound conversions (in a broader sense than image-to-music conversions)<sup>1</sup> constitute a widespread type of sonification. One could even propose that a modern/contemporary tradition of image-to-sound conversion for compositional purposes began developing even prior to computer-aided composition—a tradition that includes Scriabin, Villa-Lobos, Messiaen, and Xenakis. I will use the word “conversion” to refer to this type of sonification, instead of Cope’s “translation”, because I believe that the latter implies some shared meaning between the data to be sonified and its corresponding

---

<sup>1</sup>Even if, in a general sense, music can *make use* of any sound (Cage’s “anything goes”), that is not the same as to state that music *is* any sound for every composer and every musical work. In practical terms, the amount of effort and intent that most composers put into their compositional processes reveals that their music seeks *certain* sounds (potentially different for each composer, or each work), a much more specific territory than simply *any* sound (“whatever”, “*n’importe quoi*”). This is also valid for improvisatory, aleatoric, and “chance” approaches, which in their own ways also produce *certain* sounds as opposed to *any* sound.

sounding result (which is not necessarily the case), while the former is more neutral with regard to obtaining a shared meaning. Also, throughout this study, whenever image-to-sound sonification results in music it will be called “image-to-music conversion” so as to differentiate it from Cope’s understanding of sonification (as having “little musical use”).

Clearly there is not a unique solution to the problem of obtaining music from image-to-sound conversions, but rather multiple strategies that can serve musical solutions: not least because of the premise that music is culturally based and individually based. Therefore making music may constitute a different task (and, as such, imply different methods) for different (groups of) individuals.

Following the model of a case study, I will present some compositional strategies for making music out of sonification that have been employed in two original works: the electronic soundtrack of *VIA* (2013) and *3 Transcrições* (2011) for solo piano. These works highlight two separate aspects of the aforementioned compositional problem: contour-to-music conversion and colour-to-music conversion. Although this investigation focuses on describing a list of strategies for obtaining image-to-music conversions in these concrete compositional examples, it is also part of a broader perspective on how to musicalise (humanise, personalise) sonification.

## Contour-to-music conversion in the electronic soundtrack of *VIA*

*VIA* is an interdisciplinary project<sup>2</sup> integrating dance, architecture, design, music, video, and smartphone technology. The musical task at hand consisted of supplying 1-minute bits of music to accompany choreographies that took place in several different streets of downtown Rio de Janeiro. Each choreography was videoed *in situ* and accompanying musical segments were added later. The overall result was delivered via satellite to mobile phones as the audience walked past the streets where the dance had taken place.

Due to the project’s tight deadline, I had to make the practical decision not to compose a work of instrumental chamber music, for which the final stage of carefully transcribing the computer results to a readable score would take a longer time than available.<sup>3</sup> So the final musical segments would have to be produced for electronics, not for acoustic instruments. Furthermore, I decided to work according to Cope’s CGC (Computer-Generated Composition, i.e. generating an entire composition from one or more compositional algorithms) instead of CGA (i.e. generating musical materials that will later be worked upon to become a finished composition) [4] because the former was likely to achieve musical results within a shorter period of time—as well as presenting an interesting algorithmic challenge. I chose OPENMUSIC as the main tool for the project’s image-to-music conversions due to my previous experience with its visual programming environment and with using it to convert a variety of 2D and 3D images (see [2] and [3]).

---

<sup>2</sup>Headed by semiotician João Queiroz and dancer Daniella Aguiar and begun in 2013. For more information on *VIA* please visit the website: <http://via-locativeart.com/>.

<sup>3</sup>Following Joshua Fineberg’s invaluable advice, I always transcribe CAC (Computer-Assisted Composition) or CGA (Computer-Generated Assistance) results by hand, a slow process that however allows for an intimate control of the final results in the score.

These preliminary decisions left me with three possibilities for a sound source: a) sound synthesis (via OPENMUSIC and CSOUND, for instance); (b) exporting the generated OPENMUSIC scores to other software in order to generate sound; or (c) using a sound preview provided in OPENMUSIC to render the musical results (typically via MIDI playback).

The possibility of assembling bits of ready-made timbres (a sort of collage approach) seemed closer to the project's approach to video-dance, based on editing short images of dance taking place on the streets: with pedestrians passing by, pigeons flying around, street vendors, etc. So (a), sound synthesis, was discarded for its usually more abstract approach, whereas (b), OPENMUSIC scores exported to other software, and (c), MIDI rendering, were tested numerous times using hundreds of conversions. Eventually most of the results used in the project's final version were the product of the built-in MIDI rendering within the OPENMUSIC environment.

## Converting image contours to musical parameters

In general, the image-to-sound conversion used for this project followed the visual logic already present in traditional musical writing: that is, the  $x$ - (horizontal) axis supplied rhythm-related data, subdivided into onsets and durations, whereas the  $y$ - (vertical) axis supplied pitch-related data, within a variable pitch space that was flexibly normalisable. The process sought not only to sonify in Cope's sense, but also that this sonification generate satisfactory musical results (obviously according to the subjectivity of the composer). The patch shown in Figure 1 illustrates the overall process of conversion used in VIA. The straightforward and flexible approach to pitch is shown by an *om-scale* function tied to two *note* objects to normalise the image contour's  $y$ -axis values within any desired pitch range.

As constraints on pitch range were not demanded by any acoustic instrument, this flexibility was instead intended to allow empirical testing of the best values for the desired sounds and metaphorical "sound spectra".<sup>4</sup> For each different range of normalisation one obtains a different sonic (harmonic, intervallic, spectral) profile.

Other musical parameters, such as intensity and instrumentation/orchestration, were generated based on the same  $x$ - or  $y$ -axis input data, but with different algorithms, as shown in the following sections.

## Instrumentation and orchestration

But why use default MIDI sounds for compositional purposes in 2013?<sup>5</sup> The answer is threefold. First, the use of these sounds nowadays provides a sort of "vintage" sound profile. Second, they allow for a compositional approach based on the reuse and recombination of clichéd sounds, referencing collage and pop-art approaches. Finally, they represent an enticing compositional challenge: as they are not originally meant to

---

<sup>4</sup>Given that each MIDI note often only contributed a tiny "brick" to the overall result, the task is much like "instrumental synthesis", albeit carried out with MIDI samples instead of acoustic instruments—instrumental synthesis taken back to the studio, so to speak.

<sup>5</sup>For a reader growing up in the 1980s or 1990s, such sounds (or their earlier counterparts) already sounded old-fashioned then!

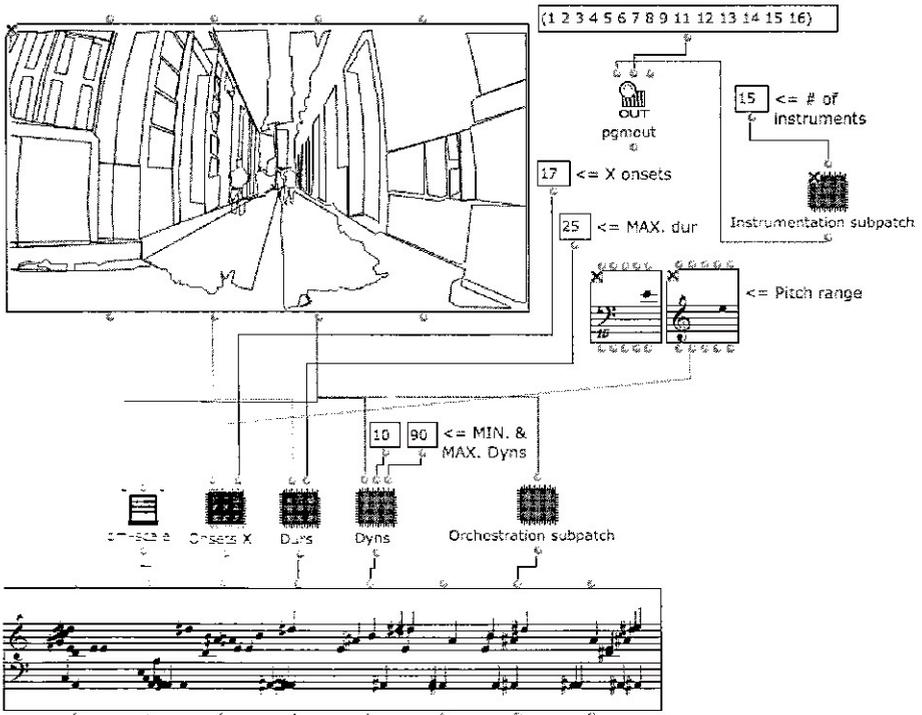


Figure 1. Overall view of one of the VIA's patches.

be presented as a final music product, but only as a sketch in the early stages of any compositional process, their use as music represents a personalised, distorted way of exposing the compositional tool at hand.

In spite of all of these arguments, to my ears most of these MIDI sounds were simply unbearable either sustained or in isolation. Therefore it was necessary at least to develop algorithms to avoid two (un)musical situations: solos and sustained sounds.

Avoidance of solos, here taken to mean timbres played alone or with significant prevalence over other timbres, was addressed and implemented by means of two simple steps analogous to traditional music approaches: *instrumentation* (the choice of sound sources, in this case particular sets of MIDI presets) and *orchestration* (the specific use one makes of these sources). I structured the OPENMUSIC *Instrumentation* patches so that any image-to-music conversion would use 15 different instruments.<sup>6</sup> In rough terms, 15 is large enough to avoid undesirable solos. This is easily feasible by assigning different MIDI programs (or instruments) to the different MIDI channels (see Figure 2).

The *Orchestration* sub-patch, in turn, divides the total pitch range in different slices ("bands", or registers), then assigns subsets of three MIDI channels (or instruments) to each range-slice, and randomly chooses one MIDI channel/instrument of each subset

<sup>6</sup>The MIDI standard provides 16 different channels, but one (channel 10) is usually reserved for percussion.

for each note. This is similar to basic acoustic orchestration (e.g. if a certain pitch is within the range of a flute, an oboe, and a clarinet, then it might be played by any of them). Several different versions of orchestration sub-patches were developed and tested. Figure 3 displays a version that segments the total pitch range into five different registers.

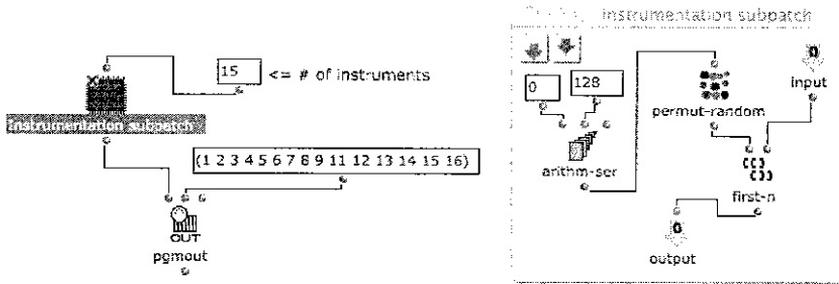


Figure 2. Left: sending out instrumentation settings to the MIDI system (one MIDI program for each channel). Right: inside the *Instrumentation* sub-patch (random selection of programs). The *Instrumentation* sub-patch is locked when an interesting combination of MIDI instruments is reached.

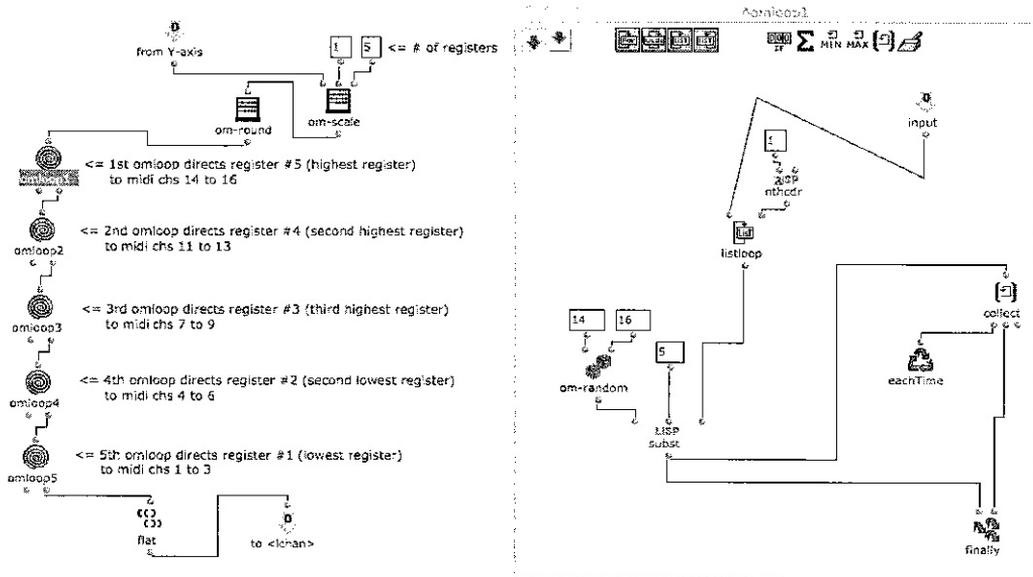


Figure 3. Left: *Orchestration* sub-patch controlling the MIDI channel assignments in Figure 1. Right: Inside the first OMLOOP.

In order to address the second point, avoiding sustained MIDI sounds, I systematically favoured small values for note durations as I normalised the image contour values. From a more global standpoint, given that resonances are not particularly attractive in these sounds, I decided to adhere to a music of attacks by limiting myself to the very beginning of each MIDI sound, its “richest” portion. The most straightforward way to implement this in OPENMUSIC is by assigning a short maximum value for *<ldur>* (the slot controlling the durations) in the *chord-seq*. Figure 4 illustrates the use of the *Durs* sub-patch, scaling down the score durations according to a small maximum value.

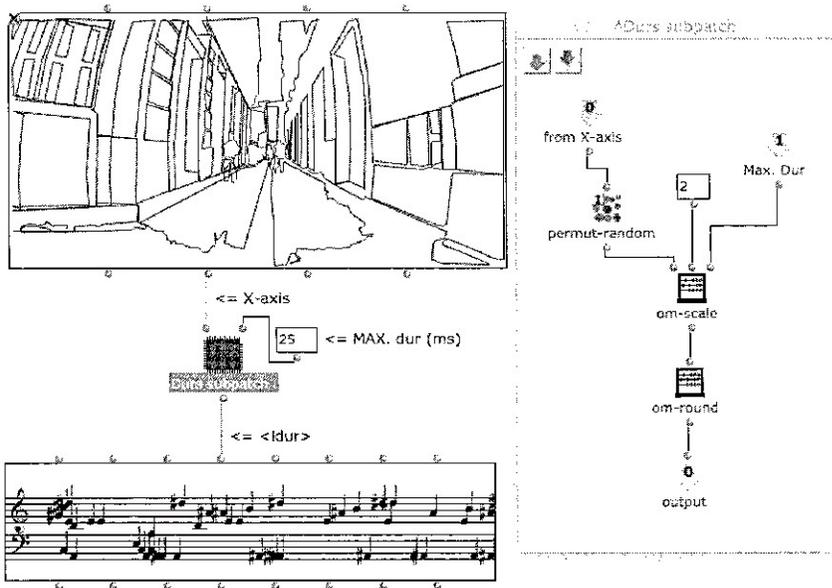


Figure 4. Setting the durations from image contour data. (N.B. this figure hides the other connections to the *chord-seq* shown in Figure 1).

By reducing the participation of each sound to a minimum, the focus of hearing is shifted from the sounds themselves to the complex networks resulting from the conversion of visual image contours to music. Each sound contributes only a tiny brick, a grain, so its individual lack of interest becomes significantly less noticeable. The overall task is more like random granular sound synthesis than writing notated music, especially taking into account that I also often compressed the inter-onset values to a minimum (see Figure 5).

In the few cases where I used recorded sound samples (instead of default MIDI sounds) to render a MIDI file generated in OPENMUSIC (by exporting it to FINALE or REAPER), I selected sounds of acoustic musical instruments used in a free, non-idiomatic way—given that the music was not intended to be written, read, or performed. The result is a sort of instrumental music produced acousmatically, one that employs tempo and gestural profiles impossible for actual musical instruments—a contrabass and a contrabassoon that play at the speed of Paganini’s frenzied violin, for instance.<sup>7</sup>

<sup>7</sup>In “Estudo Capanema” from *VIA*, audio available at: <https://soundcloud.com/lecasteloes/estudo-capanema-2013-para>.

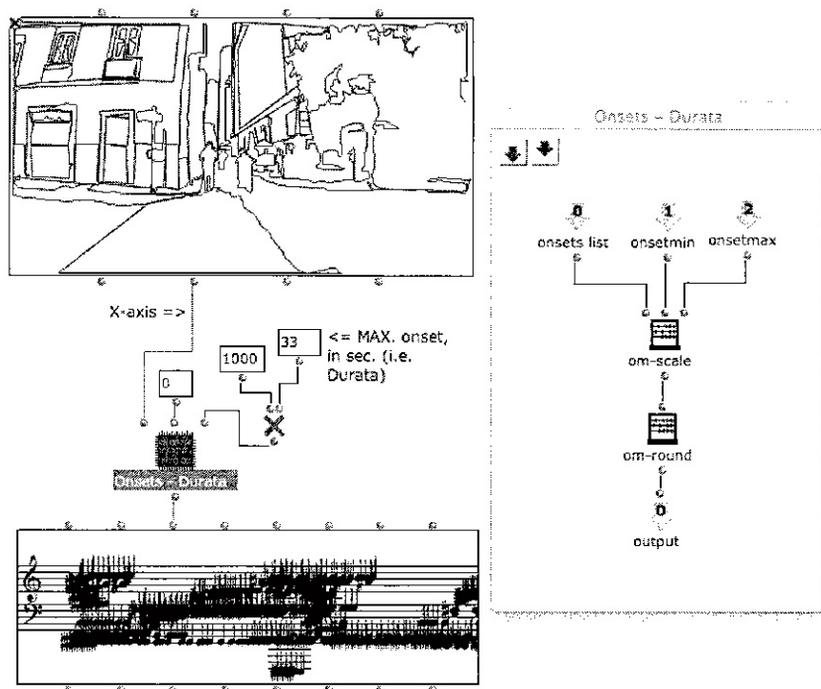


Figure 5. Scaling onset values in *VIA* (e.g. to a total duration of 33 seconds). (N.B. this figure hides the other connections to the *chord-seq* shown in Figure 1.)

## Image-to-music relationships

The implementation of these algorithms in *OPENMUSIC* was carried out so as to allow for maximum flexibility in the normalisation of the image data, like freely compressing or expanding a picture before printing it. In musical terms, this means that I was able comfortably to change normalisation values as I converted each photograph as many times as I wished, and hear the corresponding musical results, until I was fully satisfied.

These at-times-subtle changes in the normalisation values should not be underestimated because they often have a great impact on musical parameters (harmony, timbre, melody, gesture, texture, tempo, etc.). In this way, they may be the most essential factor (along with the above-mentioned instrumentation and orchestration strategies) in obtaining music from sonification—so much so that no further compositional manipulation of the input data (such as cutting, looping, superimposing, etc.) was necessary in order to obtain musical consistency in *VIA*.

In the end, hearing is the decisive factor in *VIA*'s image-to-music conversion method. It determines whether (1) there should be changes in the instrumentation and orchestration sub-patches, (2) there should be changes in the normalisation values, or (3) a musically consistent result has been reached, in which case the process comes to an end (see Figure 6).

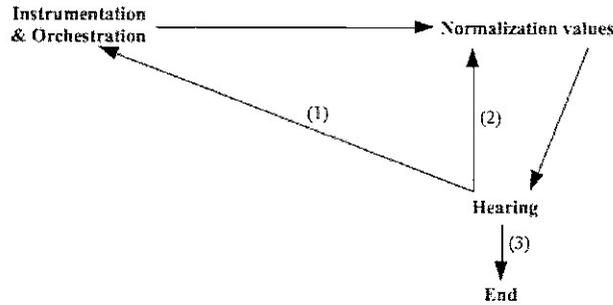


Figure 6. Image-to-music conversion method in VIA

## Music-to-video relationships

In the early stages of our project I discussed with VIA’s team the best way to apply music to video-dance for our purposes. We quickly rejected the possibility of adopting a previously composed soundtrack (totally independent from dance and images) or, on the other hand, a soundtrack that strictly followed the dancer’s bodily movements—“Mickey Mousing”. Instead, we chose to convert the visual contours of photographs from the locations where the dance had taken place. I then redrew these visual contours “by hand” (using the computer mouse) by placing each photograph as a background picture in a BPC (break-point curve) object, but always allowing a certain degree of freedom (Figures 7 and 8).

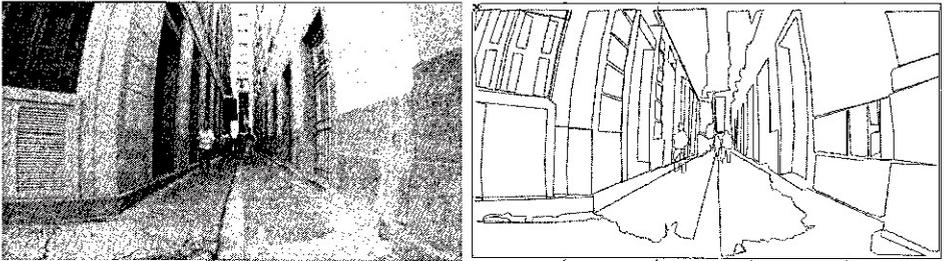


Figure 7. Redrawing visual contours in a BPC (photograph from Beco das Cancelas).

This freedom was a means not only to personalise the whole process (as another person would inevitably redraw it in a different way) but also to obtain consistent musical results: as I repeated the task over and over I gradually began to perceive general connections between the intricate visual contours and the corresponding resulting sounds, thereby allowing myself to interfere musically in the conversion process—a consequence of what Mikhail Malt calls a *solfège de modèles*, which in this case would mean *forehearing* the potential musicality beneath the images.<sup>8</sup>

<sup>8</sup>“Qu’entendons-nous par ‘solfège de modèles’? Nous ne parlons pas d’un solfège au sens d’un catalogue de modèles statiques, soit d’un solfège issu d’une typologie figée. Nous parlons plutôt du développement

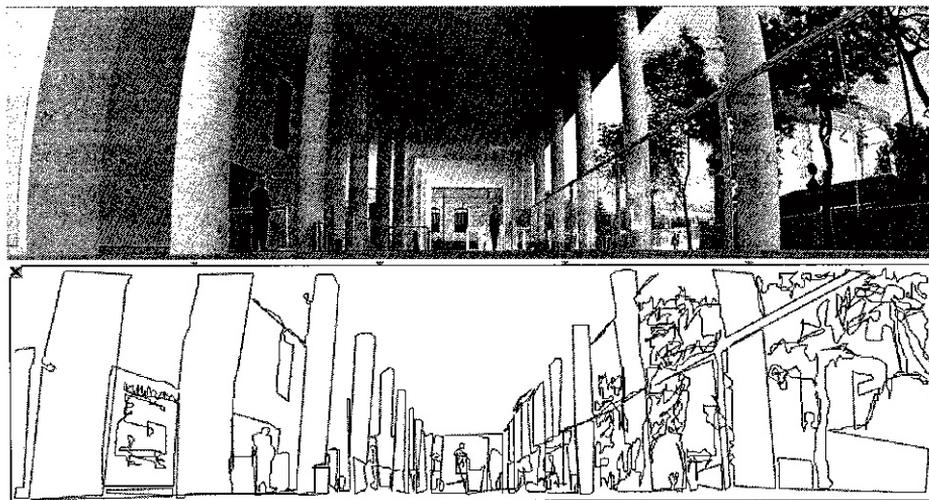


Figure 8. Redrawing visual contours in a BPC (photograph from Museu de Arte do Rio).

### Further compositional questions

Five further questions were posed and tentatively answered during the compositional process of *VIA*:

(1) What is the *space of musicality* intended by this project? Space of musicality refers here to the musical space one wishes to occupy within a certain work. Specifically, starting by imagining musical extremes (e.g., highest degree of stability opposed to the highest degree of instability), one may predetermine the musical/auditory territory—for instance characterised by a high degree of repetition, or by the alternation between repetition and contrast, or by athenaticism, etc. For the current project, a space was chosen that offered a middle ground between the high degree of communication or redundancy of pop music and the high degree of information and contrast of contemporary “art” music—that is, a space that is intelligible and enjoyable by non-musicians, young people, or specialists.

(2) What is the *magnitude* of music making? This magnitude (especially as compared to the amount of data generated by image contours) denotes the amount of data that is necessary to produce each space of musicality, or each musical work. In any given creative project that involves image-to-music conversion, one may observe a significant discrepancy between the amount of data generated from images and the required, or desired, amount of data to create a certain musical space. In the specific case of *VIA*, where the resulting musical excerpts were to be no longer than one minute, such a discrepancy did occur occasionally: the amount of data generated by image contours was of a significantly larger magnitude than the amount of data required, or desired, to create the intended space for this project. This motivated the development of criteria

---

*d'un savoir, de capacités intellectuelles et cognitives de la part du compositeur qui lui permettraient soit de contrôler et maîtriser le résultat musical issu d'un modèle génératif quelconque, soit d'établir consciemment le lien entre des représentations graphiques et/ou textuelles de certains logiciels musicaux avec un résultat musical. Solfège au sens d'aptitude à relier le comportement de deux espaces de caractéristiques différentes en tenant compte des particularités de chacun.” [5]*

intended to reduce the amount of data, including the use of only the opening musical excerpt that resulted from a given image-to-music conversion, or the simplification of the image contour done by hand in such a way that the conversion generated less data. In conclusion, creative work using image-to-sound conversion implies adjusting the magnitude of image-derived data to the magnitude of the intended musical resulting.

(3) If this magnitude is “complex”, or at least “complicated”, how does the composer relate to this complexity, and what are the best ways to obtain it and orient it toward an intended musical goal? *VIA*'s space of musicality ideally sought a level of musical complexity comparable to handmade crafts or to human activities (even though, perhaps paradoxically, the work's music was entirely realised by means of a computer). In *VIA*, the strategy for obtaining the desired complexity consisted in starting from data that was already complex and working from the macro- to the micro-level, instead of obtaining complexity by working all the way up from small building blocks. We did not seek complexity *per se*, but rather the kind of complexity that possessed musical logic. Nevertheless, it is necessary to remark that:

- The totality of the characteristics that make up such (musically logical) complexity have not been thoroughly formalised through the image-to-music algorithms developed for *VIA* and would constitute an incommensurable challenge; instead, such desired complexity is simply “encountered” and “chosen” by ear among the several musical results furnished by the converting algorithms. This complexity is partially analysed, understood, and managed, and these stages contribute to the continuous refinement of the original algorithms, an ongoing and potentially endless work in progress.
- In *VIA* the composer does not control the whole process from the start, but rather exerts gradual and partial control over the encountered, and chosen by ear, complexities generated by image-to-sound conversions.
- The production of musical results in this case is therefore significantly open to unforeseen structures.
- The final musical result is as much, or even more, “found” than “created”—in this context, music composition becomes the intellectual-auditory discovery of a pre-musical object, an abstract object with a highly musical potential.

(4) How does the magnitude of music making relate to independently developed images of choreography, as sound and image are united within an audio-visual product? There is certainly a re-signification that takes place as the finished musical component, which—let us suppose—was satisfactory enough when heard alone, joins the video component, which in turn contains the images that served as abstract data sources for the musical component. In *VIA* there was limited space for feedback between music and image; however exchanges of early mixes allowed us to learn about the particular twofold association between music and video as well as between music and dance. One of the main conclusions drawn from this dialogue was that music and choreography seemed to be more clearly related when the musical component was dense and intricate. In early mixes that had soundtracks with more empty spaces, more silences, the artificial aspect of the audio-visual collage became more apparent and it looked and sounded like an unintentional mistake, a technical flaw. This finding has further motivated the search for complexity within the musical component, as previously delineated in (3).

(5) How can one obtain *human-like* musical results while making music exclusively through a non-human device such as a computer? How can one make computer music that sounds as if humans had performed or even improvised it? Here the computer's role as a musical tool is emphasised, but without allegiance to a machine aesthetic: one in which the "watermark" [1] of the employed technology becomes too apparent, rendering the musical work that employs it obsolete even before the technology's own programmed obsolescence. Strategies for obtaining a human-like yet machine-made sonic profile include:

- ongoing refinement of the musical problems' formalisations and the implementation of algorithms that provide human-like musical results (which inevitably involve sound-to-sound intensity variation and frequent tempo fluctuation);
- emphasis on complexity (the more dense and intricate the algorithm output is, the less often one perceives the small-scale limitations of the technology employed);
- the search for asymmetrical rhythmic sequences, obtained effortlessly and at will as one copies an image by hand (or with a computer mouse);
- the search for varied instrumentation models (e.g. with 15 different musical instruments) combined with the use of aleatoric orchestration, which simulates spontaneity and surprise; and
- orchestration by registers, simulating instrumental chamber music practices.

## Colour-to-music conversion in *3 Transcrições*

*3 Transcrições* (3 Transcriptions) for solo piano<sup>9</sup> was my first attempt at making music entirely from colours. The entire piece was derived from three colour-to-music conversions—i.e. algorithms to convert colour data to quantifiable musical parameters. Again I chose OPENMUSIC to run all of the compositional algorithms; no piano was used during the compositional process. Except for a single pitch added in the second movement (the B♭6 in m. 61), the three movements stem entirely from colour-to-music conversions of a single photograph found by chance on the Web.

What the photograph portrays—a colourful van—played no role in the compositional process or in the resulting musical piece. It was chosen simply for its colour variety and because it generated satisfactory musical results. To this day, I cannot detect any meaningful (aesthetic, emotional, conceptual) connections between the source photograph and the resulting music. It is an exact example of how sonification may work *not* as a "translation".

### Colour systems: RGB, HSV, CMYK

The conversions were carried out through OPENMUSIC's *get-RGB* function, which extracts the red, green, and blue component for each pixel of a picture, as well as through two more functions that I implemented (with the invaluable assistance of mathematician Talita Sant'Ana): *RGB-to-HSV* and *RGB-to-CMYK* (see Figure 9).<sup>10</sup>

---

<sup>9</sup>Composed in 2011 and premiered in 2014 by Brazilian pianist Grazi Elis.

<sup>10</sup>Detailed explanations of equations, algorithms, and patches developed for this piece can be found in [3]. The colour-to-music functions can be downloaded at [http://www.ufjf.br/comus/cac\\_patches/](http://www.ufjf.br/comus/cac_patches/).

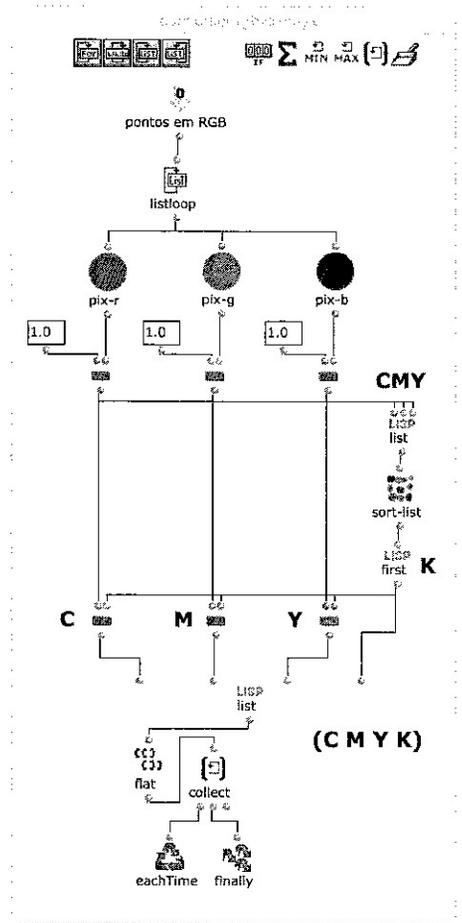
These two new functions allowed me to run and hear conversions with HSV and CMYK colour systems, enhancing my experience in colour-to-music conversion.

The different colour systems (RGB, HSV, CMYK) utilise contrasting methodologies to codify colours: RGB codifies colours by mixing numerical values of Red, Green, and Blue; HSV by mixing Hue, Saturation, and Value (Brightness); and CMYK by mixing Cyan, Magenta, Yellow, and Key (black). Therefore they produce different lists of values as they codify colours, and colour-to-music conversions based on the same image using these different systems will produce differing musical results. Figures 10 (a), (b), and (c) illustrate these differences with the results of converting the same image using the RGB, HSV, and CMYK colour systems, respectively.

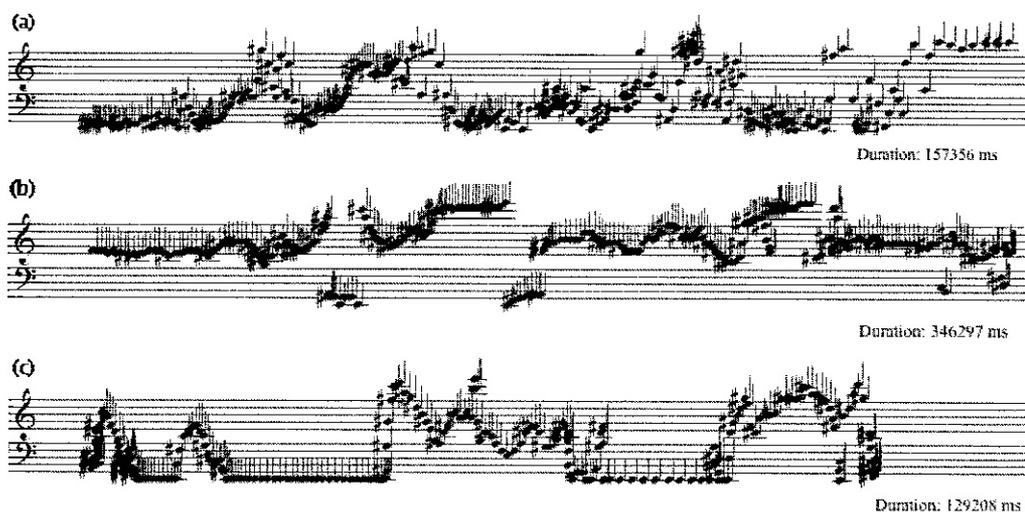
To further complicate matters, the region of the total colour space covered by each of these systems is not identical. So, even though it is possible to find intersections and convert between any two of these systems with some level of precision (using the *RGB-to-HSV* and *RGB-to-CMYK* functions), when one separates and distributes the original values of each system among various musical parameters for sonification purposes there is no hope that they will coincide.

Regardless of any evaluation of the effectiveness of each colour system (based on its total colour space, intuitiveness, etc.), I consider this methodological diversity extremely fertile for music composition because it widens the array of possible musical results derived from colour-to-music conversions. We may suppose that each of these conversions probably leaves its characteristic mark, although my research has not gone so far as to describe this mark in musical terms. This diversity also demonstrates how arbitrary the relation between colour and sound is in the context of non-synæsthetic conversion (or whenever the composer and the listener do not share a common type of synæsthesia).

In the specific case of *3 Transcrições*, only the RGB conversions have been used in the final version of the score. The choice was purely based on the quality of the resulting music.



**Figure 9.** Inside the *RGB-to-CMYK* function: an OMLOOP which converts a list of RGB points to CMYK values.



**Figure 10.** Examples of colour-to-music conversions. The same image is converted (a) using RGB: R=midicents, G=onsets, B=dynamics; (b) using HSV: H=midicents, S=onsets, V=dynamics; (c) using CMYK: C=midicents, M=onsets, Y=dynamics.

## From colour data to musical parameters

The approach to distributing colour data among musical parameters in *3 Transcrições* is in many ways similar to what I did with image contours two years later in *VIA*; that is, each list of numerical values is mapped to one or more musical parameters.

Normalisation variables are left loose enough to allow for satisfactory musical results to emerge. In the first movement R=pitch, G=rhythm (onsets+duration), and B=dynamics; in the second movement, R=dynamics, G=pitch, and B=rhythm; and in the third movement, R=rhythm, G=dynamics, and B=pitch (see Figure 11).

What distinguishes this compositional process from *VIA*'s is the particular strategies used for adapting the magnitude of image-derived data to the magnitude of music making: in *VIA*, I redrew the photographs' contours in order to simplify, musicalise, and personalise contour-to-sound conversions, whereas in *3 Transcrições* I selected only a small number of image lines (four) to be converted and then slightly filtered them, as will be shown in the next section.

## Selection of image lines and subtractive synthesis

In order to adjust the amount of colour data output by *get-RGB* to the amount of data I needed to compose a solo piano piece, I made two decisions:

First, I selected the lines of the photograph that were to be converted, given that the chosen picture had 183 lines that were not all needed to make music. This was very easy to implement in *OPENMUSIC*: I determined the position of the lines to extract by simply dividing the number of lines in the image by the number of voices (+1, so that the converted lines would never be at the borders of the photograph). This way, the extracted lines would also be equally distributed across the image.

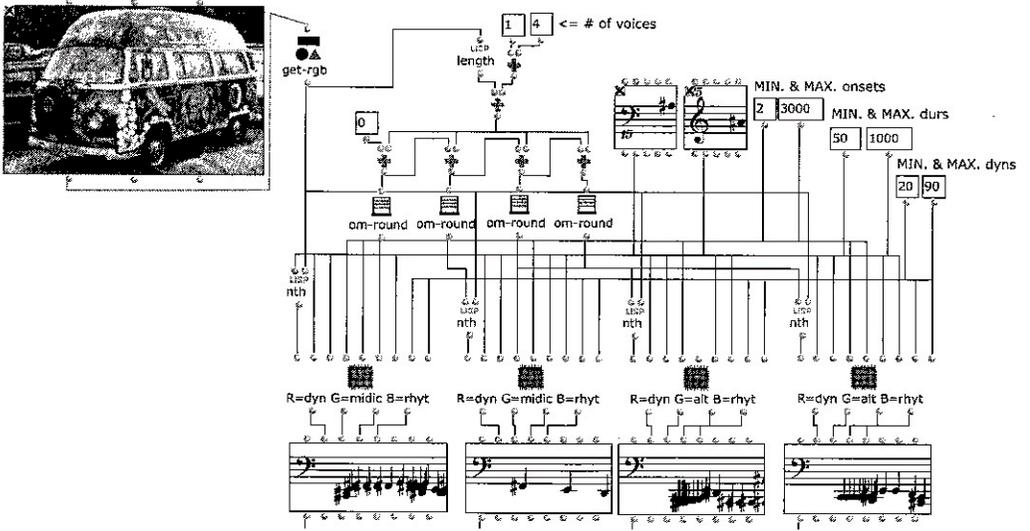


Figure 11. Overview of the patch used for the second movement of *3 Transcrições*. Selection of four horizontal lines from the picture and colour-to-music conversions.

Second, I carefully “filtered” the colour-to-music conversions by ear, very much as one does when performing subtractive synthesis. My intent with this filtering stage was twofold: (1) to further musicalise the sonification, and (2) to make the conversions playable by a human being (i.e. humanise the sonification). This filtering stage was necessary because even though the converting algorithms had been refined as closely as possible to obtain a finished composition (Cope’s CGC), the preliminary colour-to-music conversions functioned only as a sort of raw material, a block of marble that had to be sculpted further (representing a firm positioning on the side of CGA). This stage was done entirely by ear and has not been formalised in the converting algorithm. The various steps of this filtering used in the first and third movements of *3 Transcrições* are displayed in Figures 12 and 13.



Figure 12. Filtering stage in the 1<sup>st</sup> movement of *3 Transcrições*.

Figure 13. Filtering stage in the 3<sup>rd</sup> movement of *3 Transcrições*.

Once I reached satisfactory musical results by ear I went on to transcribe everything into a readable/playable musical score. This was done by hand, gesture by gesture, by carefully looking at a *chord-seq* editor grid—a slow, painstaking, yet rewarding process (see Figure 14). Transcribing music from a *chord-seq* implies not only notating pitches and durations (clearly visible on the screen), but above all hearing through what is visible in order to identify and write down what is invisible (and, therefore, does not allow itself to be transcribed by a machine). This includes the length and musical character of each musical gesture, the articulation of each sound, as well as the clearest, most musical way to fit individual durations into a larger rhythmic framework. Finding a middle ground between loyalty to the intended sound (since this piece first emerged as sound, not as a written score) and human playability was the most challenging aspect of the musical transcription (Figure 15).<sup>11</sup>

Figure 14. *chord-seq* of the beginning of *3 Transcrições* used for musical transcription.

<sup>11</sup>The audio recording of *3 Transcrições* can be heard at:  
<https://soundcloud.com/lecasteloes/3-transcricoes-2011-para-piano-solo-elis>.



Figure 15. First system of 3 *Transcrições*, after transcription.

## References

- [1] Rodolfo Caesar. *Círculos Ceifados*. Rio de Janeiro: 7 Letras, 2008.
- [2] Luiz E. Castelões. “Conversão de imagens para sons através de três classes do OpenMusic”. *Revista DAPesquisa*, 8, 2011.
- [3] Luiz E. Castelões, Talita De Oliveira, Yago Franco. “Conversores de parâmetros del color a parámetros sonoros cuantificables usando los sistemas RGB, HSV y CMYK”. *Sonic Ideas*, 7(14), 2015.
- [4] David Cope. “Preface”. In Jean Bresson, Carlos Agon, Gérard Assayag (eds.) *The OM Composer’s Book 2*. Editions Delatour France/IRCAM-Centre Pompidou, 2008.
- [5] Mikhaïl Malt. “Concepts et modèles, de l’imaginaire à l’écriture dans la composition assistée par ordinateur”. In Bruno Bossis, Anne Veitl, Marc Battier (eds.) *Musique, instruments, machines: Autour des musiques électroacoustiques. Actes actes du séminaire du MINT (vol. 2)*. Musicologie, Informatique et Nouvelles Technologies/Observatoire Musical Français/Université de Paris-Sorbonne, 2006.