

# Assisted music creation with Flow Machines: towards new categories of new

François Pachet, Pierre Roy, Benoit Carré

CTRL, Spotify [francois@spotify.com](mailto:francois@spotify.com)

**Abstract.** This chapter reflects on about 10 years of research in AI-assisted music composition, in particular during the Flow Machines project. We reflect on the motivations for such a project, its background, its main results and impact, both technological and musical, several years after its completion. We conclude with a proposal for new categories of “new”, created by the many uses of AI techniques to generate novel material.

**Keywords:** Machine-Learning · Markov chains · Applications · Music · global constraints

## 1 Background and Motivations

The dream of using machines to compose music automatically has long been a subject of investigation, by musicians and scientists.

Since the 60s, many researchers used virtually all existing artificial intelligence techniques at hand to solve music generation problems. However, little convincing music was produced with these technologies.

A landmark result in machine music generation is the Illiac suite, released to the public in 1956 [17]. The Illiac suite demonstrated that Markov chains of a rudimentary species (first order, augmented with basic generate-and-test methods) could be used to produce interesting music<sup>1</sup>.

However, the technology developed for that occasion lacked many fundamental features, to make it actually useable for concrete, professional musical projects. Notably, the experiment involved generate-and-test methods to satisfy various constraints imposed by the authors. Also, the low order of the Markov chain did not produce convincing style imitation. In spite of these many weaknesses, the Illiac suite remains today a remarkable music piece, that can still be listened to with interest.

The Flow Machines project (2012-2017)<sup>2</sup> aimed at addressing the core technical issues at stake when generating sequences in a given style. In some sense, it addressed the two main weaknesses of the Markov chains used in the Illiac suite:

---

<sup>1</sup> we invite the reader to listen to the Illiac suite, composed more than 70 years ago, to appreciate its enduring innovative character <https://www.youtube.com/watch?v=n0njBFLQSk8>

<sup>2</sup> funded by an ERC advanced grant, and conducted at Sony CSL and Université Pierre et Marie Curie, now Sorbonne Université.

the low order (and the poor style imitation quality) and the controllability, i.e. the capacity to force generated sequences to satisfy various criteria, not captured by Markov models.

### 1.1 The Continuator

More precisely, the main motivation for Flow Machines stemmed from the Continuator project. The Continuator [37] was the first interactive system to enable real-time music dialogues with a Markov model. The project was quite successful in the research community, and led to two main threads of investigation: jazz and music education. Jazz experiments were conducted notably with György Kurtág Jr and Bernard Lubat, leading to various concerts at the Uzeste festival (2000-2004) and many insights concerning the issues related to control [31].

The education experiments consisted in studying how these free-form interactions with a machine learning component could be exploited for early age music education. Promising initial experiments [39,1] led to an ambitious project<sup>3</sup> about so-called “reflexive interactions”. During this project, the Continuator system was substantially improved and extended, to handle various types of simple constraints. An interesting variant of the Continuator for music composition, called MIROR-Impro was designed, deployed and tested, with which children could generate fully-fledged music compositions, built from music generated from their own doodling [59]. It was shown also that children could clearly recognize their own style in the material generated by the system [19], a property considered as fundamental for achieving reflexive interaction.

## 2 Markov Constraints: Main Scientific Results

These promising results in the investigation of Markov sequence control led to the Flow Machines project. Technically, most of the work consisted in exploring many types of interesting constraints to be enforced on finite length sequences generated from various machine learning models, such as variable order Markov models. Other tools were developed to offer musicians a comprehensive tool palette with which they could freely explore various creative use cases of style imitation techniques.

### 2.1 The “Markov + X” roadmap

Markov models are used everywhere, from economics to Google ranking algorithms, and are good at capturing local properties of temporal sequences and abstract them into well-known mathematical concepts (*e.g.*, transition matrices or graphs). A Markov model can easily be estimated from a corpus of sequences in a given style to represent information about how musical notes follow each other. This model can then be used to produce new sequences that will sound more or

<sup>3</sup> the Miror project <http://mirorproject.eu/>

less similar to the initial sequences. Generation, or sampling, is extremely simple and consists in so-called random walks (also called drunken walks): starting from a random state, transitions are drawn randomly from the model to build a sequence step by step. However, the remarkable simplicity of random walks in Markov models meets its limitations as soon as one tries to impose specific constraints on the generated sequences. The difficulty arises when one wants to impose simple properties that cannot be expressed as local transition probabilities. For instance, imposing that the last note equals the first one, or that the notes follow some pattern, or that the total duration be fixed in advance. Even more difficult, how to impose that the generated melody is nicely “balanced”, for instance exhibiting a well-known  $1/f$  property [1], characteristic of natural phenomena?

The initial idea was to use the powerful techniques of combinatorial optimization, constraint satisfaction in particular (CP), precisely to represent Markovianity, so that other, additional properties could also be stated as constraints. Indeed, the main advantage of constraint programming is that constraints can be added at will to the problem definition, without changing the solver, at least in principle.

The idea of representing Markovianity as a global constraint was first detailed in [33]. We reformulated Markov generation as a constrained sequence generation problem, an idea which enabled us to produce remarkable examples. For instance, the *AllDifferent* constraint [57], added to a Markov constraint could produce our “Boulez Blues”: a Blues chord sequence in the style of Charlie Parker so that all chords are different! Table 2.1 shows the chord progression of the Boulez Blues<sup>4</sup>.

C7 / Fmin	Bb7 / Ebmin	Ab7 / Db7	Dbmin / Cmin
F7 / Bbmin	Eb7 / Abmin	Gmin / Gbmin	B7 / Gb7
Bmin / E7	Amin / D7	Emin / A7	Dmin / G7

However, this approach was costly, and we did not propose any boundaries on the worst-case complexity. So we started to look for efficient solutions for specific cases.

## 2.2 Positional constraints

The first substantial result we obtained was to generate efficiently Markov sequences with positional constraints (called “unary constraints” at the time). Enforcing specific values at specific positions in a finite-length sequence generated from a Markov model turned out to be solvable in polynomial time. The result described in [41] consists in first propagating all the zero probability events (filtering) and then back-propagating probabilities to ensure unbiased sampling. This result enabled us to implement an enhanced version of Continuator and triggered the development of the Flow Machines project.

<sup>4</sup> A rendering with jazz musicians can be heard at [https://www.francoispachet.fr/wp-content/uploads/2020/04/Giant\\_Steps\\_comping\\_only.mp3](https://www.francoispachet.fr/wp-content/uploads/2020/04/Giant_Steps_comping_only.mp3)

An interesting and unexpected application of this result (possibility to add positional constraints to Continuator) enabled us to introduce several types of continuations, depending on positional constraints posted on the beginning and endings of a musical phrase. Unary constraints could be used to represent various musical *intentions*, when producing a melody from a Markov model and an input melody provided in real-time. For instance, we defined the following types of melodic outputs <sup>5</sup>:

1. Continuation: input is continued to produce a sequence of the same size. A constraint is posted on the last note to ensure that it is “terminal”, i.e. occurred at the end of an input melody, to produce a coherent ending.
2. Variation: is generated by adding two unary constraints that the first and last notes should be the same, respectively, as the first and last notes of the input.
3. Answer: is like a Continuation, but the last note should be the same as the first input note. This creates a phrase that resolves to the beginning, producing a sense of closure.

Additionally, for all types of responses, unary constraints were posted on each intermediary note stating that they should not be initial nor final, to avoid false starts/ends within the melody. This use of positional constraints to bias the nature of a continuation nicely echoes the notion of conversation advocated by composer Hans Zimmer in [10]. These constraints substantially improved the musical quality of the responses generated by Continuator.

Figure 1 shows a melody from which we build a Markov model  $M$ . Figure 2 shows examples of continuations, variations and answers, built from  $M$  and the constraints corresponding to each melody type, from an input (I). It is clear that these melodies belong to their respective musical categories: Continuations end naturally, with the same 3 notes as the input (a consequence of the constraint on the last note); variations sound similar to the input; and answers sound as responses to the input. This shows how simple unary constraints can be very powerful and produce a global effect on the structure of generated melodies.

An application of positional constraints to the generation of jazz virtuoso melodies [32] was also designed, controlled by a gesture controller<sup>6</sup>.

An interesting application of positional constraints was lyric generation, satisfying rhymes and prosody constraints [5], leading to the generation of lyrics from arbitrary prosody and rhyme templates<sup>7</sup>.

In all cases, the imposition of these positional constraints, and their “back propagation” consequence on the generated sequence, produce pleasing musical

<sup>5</sup> we thank here Shai Newman for motivating us to develop these ideas

<sup>6</sup> see, e.g. a session with Virtuoso, involving Mark d’Inverno, Goldsmiths, University of London <https://www.youtube.com/watch?v=wCaC53G1EY0>, and another one at Sony headquarters <https://www.youtube.com/watch?v=iGvWVy7dglk>, and an earlier version in the hands of an enthusiastic and expressive Geber Ramalho <https://www.youtube.com/watch?v=3NWkmSKyMXY>

<sup>7</sup> A python implementation and app is available at [4]

effects, probably due to the implicit generation of meaningful "preparations" for specific, imposed events. Note that this work was extended to handle positional constraints for recurrent neural networks in [14], leading to similar effects.



Fig. 1. A training sequence.



Fig. 2. Different continuation types generated from the same training sequence: continuations (C), variations (V) and answers (A), generated from the input melody (I).

**Harmonization** An interesting application of positional constraints was the harmonization system described in [35]. In the system, positional constraints were used to enforce specific constraints on chords to be played during the onset of melody notes. Passing chords were generated in between those notes to create interesting harmonic movements (called fioritures) that would fit with the imposed melody, but not necessarily in conformant ways. Some remarkable results were obtained, such as a version of Coltrane’s Giant Steps in the style of Wagner (see figure 3).

An interesting harmonization of the title “Começar de Nuovo” composed by Ivan Lins, in the style of Take 6 was produced, and shown to Ivan Lins himself, whose reacted enthusiastically <sup>8</sup>.

These results obtained with positional constraints and Markov chains and their application paved the way for an extensive research roadmap, aimed at finding methods for controlling Markov chains with constraints of increasing complexity.

<sup>8</sup> <https://www.youtube.com/watch?v=WBvX-03qP6A>

The image shows a musical score for Coltrane's 'Giant Steps' in the style of Wagner. It consists of four staves of music. The top staff is the soprano line, which is the melody. The rest of the score is a complex harmonic accompaniment. Chord symbols are written above the notes. The chords are: B maj7, D 7, G maj7, Bb 7, Eb maj7, A min7, D 7, G maj7, Bb 7, Eb maj7, F# 7, B maj7, F min7, Bb 7, Eb maj7, A min7, D 7, G maj7, C# min7, F# 7, B maj7, F min7, Bb 7, Eb maj7, C# min7, F# 7.

**Fig. 3.** An harmonization of Coltrane’s Giant Steps in the style of Wagner. The melody is the soprano line, chord symbols are displayed, the rest is generated in the style of Wagner.

### 2.3 Meter and all that jazz

The next problem to solve was meter: notes have durations and there are many reasons to constrain the sum of these durations in music. Meter cannot be expressed easily by index positions in a sequence. As a consequence, one cannot easily generate Markov sequences of events having variable durations, while imposing a fixed total duration for instance, which is problematic for music composition. Thanks to a theorem by Khovanskii in additive number theory, we found a pseudo-polynomial solution for meter and Markov sequences. the *Meter* constraint enables the generation of Markov sequences satisfying arbitrary constraints on partial sums of durations. This important result [60] was heavily used in our subsequent systems, notably for lead sheet generation “in the style of” with arbitrary constraints [34].

Another big step was to address a recurring problems of Markov chains: increasing the order leads to solutions which contain large copies of the corpus: how to limit this effect? MaxOrder was introduced in [49] precisely to solve this problem. The solution consists in reformulating the problem as an automaton, using the framework of regular constraints [55]. We proposed a polynomial algorithm to build this automaton (max order imposes a set of forbidden substrings to the Markov sequences).

Now that we had a way to enforce basic constraints on meter and order, we investigated ways of making sequences more “human”. A beautiful result, in our opinion, was to revisit the classical result of Voss & Clarke concerning  $1/f$  distribution in music [66,67]. We looked for a constraint that biases a sequence so that its spectrum is in  $1/f$ . We showed that the stochastic, dice-based algorithm

proposed by Voss can be expressed as a tree of ternary sum constraints, leading to an efficient implementation [44]. For the first time, one could generate meter sequences in  $1/f$ . The examples in the original Voss paper did not have bars, understandably: now we could add them!

Paradoxically, the only negative result we obtained was to show that enforcing binary equalities within Markov chains was  $\#P$ -complete in the general case (as well as grammar constraints more generally) [58]. This is a counter intuitive result as this constraint seemed *a priori* simple to enforce.

## 2.4 Sampling methods

The next class of problems we addressed was sampling: how to get not only all the solutions of a constraint problem, but a distribution of typical sequences, and for more powerful graphical models than linear Markov chains, to handle polyphonic sequences. Works on sampling led to a remarkable result: that all regular constraints (as introduced by Pesant in [55]) added to Markov constraints can be sampled in polynomial time [48]. We later realised that our positional constraint algorithm was equivalent to belief propagation [52], and that meter, as well as MaxOrder [47] were regular. This led to a novel, faster and clearer implementation of metrical Markov sequences [48].

Some interesting extensions to more sophisticated constraints were studied, such as *AllenMeter* [64]. *AllenMeter* allows users to express constraints using temporal locations (instead of events) involving all the Allen relations [2]. This constraint was used to generate polyphonic music enforcing complex synchronization properties (see section 5). Palindromes were also studied (i.e. sequences that can read both forward and backward) and a beautiful graph-based solution was found to generate all palindromic sequences for 1st order Markov chains[51]. A fascinating application to the generation of cancrizan canons was experimented by Pierre Roy with promising results (see Figure 4).



**Fig. 4.** A 10-bar cancrizan canon generated by our palindrome Markov generator.

Another interesting development addressed the issue of generating meaningful variations of musical sequences. This was performed by representing musical

distance as biases of the local fields in the belief propagation algorithm [63,40]. This type of issues is now addressed typically with variational auto-encoders [7] with similar results.

These results somehow closed the chapter “Markov + X” global constraints, since most interesting constraints in music and text can be expressed as a regular constraints. This line of works generated a substantial amount of theoretical and algorithmic results [46], as well as fruitful collaborations with other CP researchers. Jean-Charles Régin and Guillaume Perez, in particular, reformulated a number of our algorithms in the framework of Multi-valued Decision Diagrams (MDD), yielding substantial gains in efficiency [54,53].

### 3 Beyond Markov Models

Markov models (and their variable order species) having been thoroughly investigated, we turned to more powerful models with the same goal in mind: finding efficient ways of controlling them. We explored the use of the *maximum entropy* principle[18] for music modelling. This model is based on the representation of binary relationships between possibly not contiguous events, thereby preventing issues related to parameter explosion inherent to higher-order Markov models. Departing from a pure filtering approach, parameter estimation is performed using high-dimension gradient search, and generation using a Metropolis algorithm. This model gave interesting results for monophonic sequences [65] and some extensions were studied to polyphonic ones as well [16]. An application to modelling expressiveness in monophonic melodies was conducted in [29] with promising results.

Other aspects of style capture and generation were considered in Flow Machines, such as the generation of audio accompaniments “in the style of”. A dynamic programming approach was developed in conjunction with a smart “audio gluing” mechanism to preserve *groove*, i.e. small deviations in the onset of events that characterize the style [56]. A convincing example can be heard in the Bossa nova orchestration of Ode to joy [38]. This result started an ongoing collaboration with Brazilian colleagues to capture Brazilian guitar styles (the Brazyle project <sup>9</sup>).

### 4 Flow Composer: the first AI assisted lead sheet composition tool

These techniques were used to develop a lead sheet generator called Flow Composer <sup>10</sup>. The lead sheet generator was trained using a unique database of lead sheets developed for this occasion, LSDB [36]. In order to use these generators, we designed the interface Flow Composer [50]. The basic idea was to let users

<sup>9</sup> [www.brazyle.com](http://www.brazyle.com)

<sup>10</sup> A retrospective analysis of the development of Flow Composer [45] describes the motivation and principles behind this project



enter arbitrary chunks of songs (melody or chords) and let the generator fill in the blanks, a process referred to now as *inpainting*<sup>11</sup>. This process took several iterations, ranging from the development of a javascript library for music web editing [23] to studies about the impact of feedback on composition with these tools [24].

Thanks to this interface, Flow Composer was intensively used by musicians, in particular by SKYGGE<sup>12</sup> to compose and produce what turned out to be the first mainstream music album composed with AI: Hello World<sup>13</sup>.

## 5 Significant Music Productions

In this section, we review some of the most significant music produced with the tools developed in the Flow Machines project. Departing from most research in computer music, we stressed from the beginning the importance of working with musicians, in order to avoid the “demo effect”, whereby a feature is demonstrated in a spectacular but artificial way, regardless of the actual possibility by a real musician to use it to make real music to a real audience. It is the opinion of the authors that some of the music described above stands out, but we invite the reader to listen to these examples to form his/her judgement.

### 1. The Boulez Blues

The Boulez Blues (see section 2.1) is a Blues chord sequence in the style of Charlie Parker (generated from a 1st order Markov model). The sequence contains only major and minor chords, in all 12 keys (i.e. 24 chords in total, 2 per bar) and additionally satisfies an *AllDifferent* constraint, i.e. all chords are different. The sequence was generated with a preliminary version of Markov constraints, that enabled the computation of the most probable sequence (using branch & bound). In that sense, we can say that the Boulez Blues is the most Parker-like Blues for which all 24 chords are different. A rendering with jazz musicians can be found at<sup>14</sup>. This remarkable Blues chord progression can be considered as an original **stylistic singularity** (of the style of Charlie Parker, using a specific constraint which clearly goes in the way of the style).

### 2. Two harmonizations with Flow Composer

Some harmonization produced with Flow Composer (see Section 2.2) stand out. We stress here the musical qualities of the Wagner harmonization of Giant Steps, and the harmonization of “Começar de Nuovo” (Ivan Lins) in the style of Take 6<sup>15</sup>. These pieces can be seen as particular instances of **style transfer** (here the style of orchestration is transferred from one piece to another).

<sup>11</sup> See video <https://www.youtube.com/watch?v=f6V8N3ZcSEs>

<sup>12</sup> SKYGGE is the artist name Benoit Carré uses for all musical productions done with AI, see <https://open.spotify.com/artist/4aGSoPUP1v4qh7RfBlvgBR/about>

<sup>13</sup> <https://www.helloworldalbum.net/>

<sup>14</sup> <https://www.francoispachet.fr/markov-contraints>

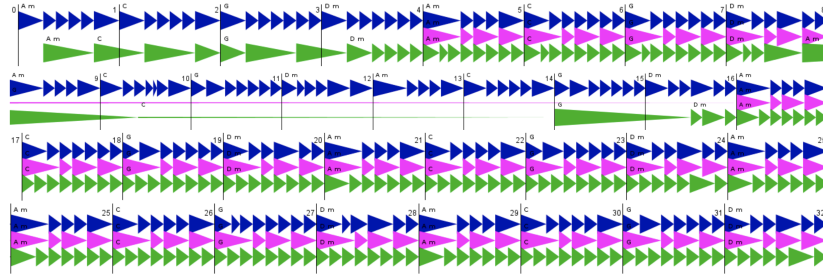
<sup>15</sup> <https://www.youtube.com/watch?v=WBvX-03qP6A>

### 3. Orchestrations of Ode to Joy

During the project, at the request of the ERC for the celebration of the 5000th grantee, we produced seven orchestrations of Beethoven’s “Ode to joy” (the anthem of the European Union). These orchestrations were produced with various techniques and in different styles [38]<sup>16</sup>.

The most notable orchestrations are:

- (a) Multi-track audio generation A particularly interesting application of the work on *AllenMeter* by Marco Marchini (see section 2.3) is the generation of multi-track audio, involving temporal synchronization constraints between the different tracks (bass, drum, chords). Figure 5 shows an excerpt of Ode to Joy in the style of Prayer in C (Lilly Wood and the Prick).



**Fig. 5.** A multi-track audio piece generated with Allen Meter: A graphical representation of the guitar (top), bass (center), and drum (bottom) tracks of Prayer in C. Each track contains 32 bars and each triangle represents a chunk. Vertical lines indicate bar separations.

- (b) Bossa Nova Orchestration with concatenative synthesis

A Bossa nova orchestration was generated for which the guitar accompaniment was produced using concatenative synthesis, applied to a corpus of guitar recordings (see Section 3). The techniques used are described in [56], and this Bossa nova example nicely emphasizes how the groove of the original guitar recording (by the first author of this paper) was preserved.

- (c) Bach like chorale using Max Entropy

An orchestration in the style of Bach was generated, using the maximum entropy principle (see Section 3), which had shown great theoretical as well as musical results on monophonic music modelling [65]. This anachronistic yet convincing orchestration paved the way for the DeepBach system, a more complete orchestration system in the style of Bach based on a LSTM architecture [15].

<sup>16</sup> the orchestrations are available on this video, which received the best video award for AAAI 2016 <https://www.youtube.com/watch?v=buXqNqBFd6E>

#### 4. Beyond The Fence

Beyond The Fence is a unique musical, commissioned by Wingspan Production. The idea was to produce a full musical show using various *computational creativity* techniques, from the pitch of the musical to the songs, including music and lyrics. The musical was produced in 2015 and was eventually staged in the Arts Theatre in London’s West End during February and March of 2016. The complex *making of* of the musical gave birth to two 1-hour documentary films, which were aired on Sky Arts under the title “Computer Says Show”.

Some of the songs were composed with a preliminary version of Flow Composer, under the control of composers Benjamin Till and Nathan Taylor. In particular, the song “Scratch That Itch” (see Figure 6) received good criticism. The song was produced by training Flow Composer on a corpus of musicals, mostly Broadway shows, in an attempt to replicate that style.

Score

**12: Scratch That Itch**  
From Beyond the Fence

Benjamin Till / Nathan Taylor /  
Computational Composer

The score is for the song "Scratch That Itch" from the musical "Beyond the Fence". It is in 8/8 time with a tempo of 102. The score includes three parts: Voice 1, Piano, and Glockenspiel. The voice part is for a character named CERIDWEN and includes the lyrics: "Now there's one thing I want to make clear, there's wo-men of all sorts by here. All". The piano part features dynamic markings of *f* and *mp*. The Glockenspiel part has a rhythmic accompaniment.

**Fig. 6.** An excerpt of “Scratch That Itch”, a song composed with Flow Composer for the musical “Beyond the Fence”.

An analysis of the production and reception of the musical was published in [8]. The overall reception was good, though difficult to evaluate, because of the large scope of the project. Concerning the songs made with Flow Composer, a critic wrote:

I particularly enjoyed Scratch That Itch, . . . , which reminded me of a Gilbert & Sullivan number whilst other songs have elements of Les Miserables. Caroline Hanks-Farmer, carns, TheatrePassion.com

This critic embodies the essence of the reception in our view: songs created did bear some analogy with the corpus, and created a feeling of *reminiscence*, at least to some listeners.

#### 5. Catchy tune snippets with Flow Composer

Before Flow Composer was used on a large scale for the Hello World album (see below), the system was used to compose fragments of songs in various styles. Some of these fragments are still today worth listening, bearing a catchiness that is absent from most of the generated music produced 4 years later. These three examples can be seen as **stylistic explorations** of various composers.

(a) in the style of Miles Davis

This song snippet was generated by training Flow Composer on a set of about 10 songs composed by Miles Davis (see Figure 7)<sup>17</sup>. These 8 bars, played in loop, produce an engaging groove though with a rather unconventional yet consistent melody.



**Fig. 7.** A song snippet composed with Flow Composer in the style of Miles Davis. Available at <https://www.francoispachet.fr/wp-content/uploads/2020/04/MilesGetLucky.mp3>

(b) From Solar to Lunar

This short song was also generated by training Flow Composer on a set of 10 songs composed by Miles Davis (see Figure 7). Additionally, the structure of the song “Solar” was reused, an instance of **templagiarism**, hence its title. A rendering is available at [https://www.francoispachet.fr/wp-content/uploads/2020/04/Miles-Davis-Mix\\_DEF.mp3](https://www.francoispachet.fr/wp-content/uploads/2020/04/Miles-Davis-Mix_DEF.mp3). This song was deemed sufficiently interesting to be performed live in London (Mark d’Inverno quartet, available at <https://www.francoispachet.fr/wp-content/uploads/2020/04/LunarPerformance.mp3>, see Figure 9) during an “AI concert”. Listening to the the tune performed by Mark d’Inverno quartet, journalist James Vincent wrote<sup>18</sup>:

The AI’s contribution was just a lead sheet — a single piece of paper with a melody line and accompanying chords — but in the hands of d’Inverno and his bandmates, it swung. They started off running through the jaunty main theme, before devolving into a series of solos that d’Inverno later informed me were *all human* (meaning, all improvised).

<sup>17</sup> A rendering is available at <https://www.francoispachet.fr/wp-content/uploads/2020/04/MilesGetLucky.mp3>, with an accompaniment in the style of the song Get Lucky by Daft Punk

<sup>18</sup> <https://www.theverge.com/2016/10/12/13247686/ai-music-composition-jazz-club-london-deep-learning>



**Fig. 8.** Lunar, a song composed with Flow Composer in the style of Miles Davis, with the structure of Solar. Available at [https://www.francoispachet.fr/wp-content/uploads/2020/04/Miles-Davis-Mix\\_DEF.mp3](https://www.francoispachet.fr/wp-content/uploads/2020/04/Miles-Davis-Mix_DEF.mp3)



**Fig. 9.** The concert of the Mark d'Inverno quartet playing Lunar at the Vortex Club in London, October 2016.

(c) in the style of Bill Evans

This short song was generated by training Flow Composer on a set of 10 songs composed by Bill Evans (see Figure 10<sup>19</sup>). The melody nicely navigates through the chord progression, sometimes in an unconventional but clever manner (bar 3). Thanks to unary constraints, the tune nicely transitions from the end to the beginning. Other songs in the style of Bill Evans were generated on the fly and performed live at the Gaité Lyrique concert (see below).

CP 2016 FC

Generated Bill Evans AB

The musical score is written in 4/4 time and consists of four staves. The chords are indicated above the notes. The first staff contains measures 1-4 with chords Bm9, C#+7, Gm9, and EbM7#11. The second staff contains measures 5-8 with chords AbM7#11, Am9, Dm9, and Gm7. The third staff contains measures 9-12 with chords Em9, Am9, D7, Gm7, C7, Am7, Abm7, Gm7, and Fm6. The fourth staff contains measures 13-16 with chords Em9, DM7, Gm7, and G#o.

**Fig. 10.** A song snippet composed with Flow Composer in the style of Bill Evans. Available at <https://www.francoispachet.fr/wp-content/uploads/2020/04/BillEvansStyle.mp3>

## 6. Daddy’s car

The song Daddy’s car was generated with Flow Composer, using a corpus of 45 Beatles song (from the latest period, considered usually as the richest and most singular in the recording history of the Beatles). Flow Composer was used to generate the lead sheet, while the lyrics and most of the orchestration were done manually, by SKYGGE. The song was released in septembre 2016 on YouTube<sup>20</sup> and received a lot of media attention<sup>21</sup>.

An interesting aspect of this song is how the system identified and reproduced a typical “pattern” found in many Beatles song. This pattern consists in a chromatic descending bass and can be found for instance in the introduction

<sup>19</sup> A rendering is available at <https://www.francoispachet.fr/wp-content/uploads/2020/04/BillEvansStyle.mp3>


<sup>20</sup> [https://www.youtube.com/watch?v=LSHZ\\_b05W7o](https://www.youtube.com/watch?v=LSHZ_b05W7o)

<sup>21</sup> about 2.5 million views as of April 2020

and in the middle of the song Michelle (which was part of the training set) (see Figure 11). This pattern was reused and instantiated in a novel context in the middle of “Daddy’s car” (see Figure 12). For that reason, the song can be considered a **pastiche**: it reuses explicitly typical patterns of the style, without any attempt at introducing novelty.



**Fig. 11.** The introduction of Michelle uses a typical harmonic pattern of the Beatles (here in D minor). Note that we show here the Real Book spelling, with a C#+ chord, which is equivalent to a DmM7.



**Fig. 12.** The song Daddy’s Car reuses the harmonic pattern of Michelle in a different context (here in A minor).

## 7. Busy P remix

Artist Busy P<sup>22</sup> launched title “Genie” in 2016<sup>23</sup>. Several remixes of the title were produced, including one made with Flow Machines<sup>24</sup>. In this remix, a variation of the chord sequence was generated with Flow Composer by SKYGGE, and then rendered using concatenative synthesis [56]. The remix features an interesting and surprising slowdown, due to an artefact in the generation algorithm, that was considered creative and kept by the remixer.

## 8. Move on

Unesco commissioned a song to be composed with Flow Machines, using material representative of 19 so-called “creative cities”. This song, called “Move On”, was composed in 2017. 20 scores of songs were gathered and used as a training set. The composition and realization was performed by SKYGGE,

<sup>22</sup> His real name is Fred Winter, French DJ, producer and ex-manager of Daft Punk

<sup>23</sup> available on Spotify at [https://open.spotify.com/track/7J02VXGmyD20uGPLTnuONu?si=QKqpmXe\\_SACtpGsI810kbw](https://open.spotify.com/track/7J02VXGmyD20uGPLTnuONu?si=QKqpmXe_SACtpGsI810kbw)

<sup>24</sup> The Flow Machine remix of Busy P’s Genie, mixed by Alf Briat: [https://www.youtube.com/watch?v=Pf\\_1\\_wDeoSY](https://www.youtube.com/watch?v=Pf_1_wDeoSY), also on Spotify at [https://open.spotify.com/track/6BoZ318dTJ7cWPENY28M2a?si=GBPppzwuQ\\_Goe79LYksFqA](https://open.spotify.com/track/6BoZ318dTJ7cWPENY28M2a?si=GBPppzwuQ_Goe79LYksFqA)

Arthur Philippot and Pierre Jouan, with the band “Catastrophe”<sup>25</sup> credited as “interpreter”. The song was released and distributed as a physical 45 RPM vinyl<sup>26</sup> in 2017. The song integrates a large number of musical elements, sounds and voices but manages to produce a coherent and addictive picture, thanks to the repeated use of smart transformations of the sources to fit the same chord progression throughout.

#### 9. Concert at Gaité Lyrique

A unique concert involving several artists who composed songs with Flow Machines was performed at Gaité Lyrique in Paris, in October 2016. This concert was probably the first ever concert showcasing pop tunes (as well as jazz) composed (in real time for some of them) with AI.

The concert<sup>27</sup> involved the following artists and songs:

(a) SKYGGGE (Daddy’s Car)

SKYGGGE performed Daddy’s car as well as several titles<sup>28</sup> which appeared later in the Hello World album (see below for descriptions).

(b) Camille Bertault

Jazz singer Camille Bertault (who has, since then, become a successful jazz interpret) sung various songs composed in real-time by Flow Composer, in the style of Bill Evans. She was accompanied by Fady Farah on the piano<sup>29</sup>. The voice of Bertault fits particularly well with the Evansian harmonies and the simple but efficient voicings played by the pianist.

(c) O (Azerty)

Artist Oliver Marguerit (aka “O”) composed a beautiful song titled “Azerty”<sup>30</sup>. Verses were generated by Flow Composer trained on a mix of “God Only Knows” by The Beach Boys, “Sea Song” by Robert Wyatt and “The Sicilian Clan” by Ennio Morricone. This song features interesting, and in some sense typical, melodic patterns created with Flow Composer that sound momentarily unconventional but eventually seem to resolve and make sense, pushing the listener to keep his attention on the generously varied melodic line. The song also features a 4 bar melody woven onto an unconventional harmony that deserves attention (see Figure 13).

(d) François Pachet (Jazz improvisation with Reflexive Looper)

Pachet performed a guitar improvisation on the jazz standard Autumn Leaves using Reflexive Looper<sup>31</sup>, a system that learns in real time how

<sup>25</sup> <http://lacatastrophe.fr/>

<sup>26</sup> it can be heard at <https://soundcloud.com/user-76901948/sets/la-playlist-des-villes-creatives-de-lunesco>

<sup>27</sup> a summary of the concert can be seen at <https://www.youtube.com/watch?v=bptKZ2ACZfQ>.

<sup>28</sup> [https://www.youtube.com/watch?v=cTPOSr\\_ehmY](https://www.youtube.com/watch?v=cTPOSr_ehmY)

<sup>29</sup> <https://www.youtube.com/watch?v=1YfKbLcjwUM>

<sup>30</sup> <https://www.youtube.com/watch?v=v69KcJBrpuo>

<sup>31</sup> <https://www.youtube.com/watch?v=8YzPaCzDDzg>





**Fig. 13.** A haunting 4 bar melody/harmony combination at the heart of the song Azerty.

to generate bass and chord accompaniments, described in [43,22] (see section 6).

(e) Kumisolo (Kageru San)

Japanese artist Kumisolo interpreted the song “Kageru San”, composed in the JPop style by SKYGGE with Flow Composer <sup>32</sup>. Flow Composer was trained with songs from the album “Rubber Soul” by The Beatles. The audio renderings were generated from excerpts of “Everybody’s Talking at Me” by Harry Wilson and “White Eagle” by Tangerine. The song, while firmly anchored in the JPop style, has a memorizable, catchy chorus.

(f) Lescop (A mon sujet)

The song “A mon sujet” <sup>33</sup> was composed by Lescop with Flow Composer trained on songs “Yaton” by Beak, “Vitamin C” by Can, “Tresor” by Flavien Berger, “Supernature” by Cerrone and “Rain” by Tim Paris. Although the resulting melodic and harmonic lines are monotonous compared to other songs, the voice of Lescop mixed with the piano and bass ostinato create a hauntingly obsessive and attaching musical atmosphere.

(g) ALB (Deedadooda)

The song Deedadooda <sup>34</sup> by emerging artist ALB was composed with a training set of 16 songs by ALB himself. It features a dialog between ALB and a virtual composer, cleverly integrated in the song itself.

(h) Housse de Racket (Futura)

Futura was composed from a training set of 6 songs by Housse de Racket themselves <sup>35</sup>. The rendering is generated from the introduction of “Love’s in Need of Love Today” by Steve Wonder, which creates an enigmatic and original Pop atmosphere on an unusually slow tempo.

10. Hello World

The most sophisticated music production of the project was the album “Hello World” <sup>36</sup>. This album features 15 songs composed with various artists who used Flow Machines at its latest stage. Composition sessions were held at Sony CSL, conducted by SKYGGE, and involved artists in several genres

<sup>32</sup> [https://www.youtube.com/watch?v=sCmp3Xp\\_PG4](https://www.youtube.com/watch?v=sCmp3Xp_PG4)

<sup>33</sup> <https://www.youtube.com/watch?v=9uTIyJ8vDRM>

<sup>34</sup> <https://www.youtube.com/watch?v=u7iG4SKbIbY&list=PLvoqwxjRRNfmn2A6e9-gdWo90RhZfLFm0&index=4>

<sup>35</sup> [https://www.youtube.com/watch?v=BU9gOX\\_H8zo](https://www.youtube.com/watch?v=BU9gOX_H8zo)

<sup>36</sup> <https://www.helloworldalbum.net>

(Pop, Jazz, Electronic Music). Hello World received considerable media attention<sup>37</sup> and reached remarkable audience stream counts (about 12 million streams). Hello World was the first mainstream music album featuring only AI-assisted compositions and orchestrations. The song list is the following:

(1) Ballad of the Shadow (2:48)

The Ballad of the Shadow is the first song written for this album. It was composed by SKYGGE with Flow Machines, inspired by jazz standards, with the goal of making a vaporwave cowboy song. The idea of shadows singing was developed into a melody with a happy mood. It resembles a cartoon-like tune when it is played at 120 bpm, but becomes melancholic when played slower, especially with ambient textures. Ash Workman and Michael Lovett detuned the drums and added some drive effects. The main theme (see Figure 14) is, again, enigmatic, unconventional yet beautifully makes its way through an unusual chord progression.




Fig. 14. An enigmatic 8 bar theme at the heart of “Ballad of the Shadow”.

(2) Sensitive (4:14)

Sensitive was composed by SKYGGE with Flow Machines, inspired by bossa novas. Lyrics and voice are from C Duncan. The song was composed from a corpus of bossa novas of the 60s. There are some patterns in the song such as major-minor progressions, that bossa nova fans will recognize and like (see Figure 15). Harmonic changes are sometimes very audacious, but the melody always stays on tracks. SKYGGE generated a voice for the melodic line with random lyrics, like a mosaic of syllables extracted from the a cappella recording of a vocalist. However, once played with a piano, the music sounded like a powerful 70s ballad. SKYGGE liked the combination of the two styles, and recorded a string arrangement written by Marie-Jeanne Serrero and live drums. C Duncan was very enthusiastic when he listened to the song, he wrote lyrics from random material inspired by key words like wind, deep and feel.

(3) One Note Samba (3:06)

<sup>37</sup> <https://www.helloworldalbum.net/press/>

The image shows a musical score for the first eight bars of the song 'Sensitive'. The music is written in a 4/4 time signature with a key signature of two flats (B-flat and E-flat). The notation is split across two staves. The first staff contains measures 1 through 4, with chords Ab (marked with a boxed 'A'), D7, Fm6, and G7b9 indicated above the notes. The second staff contains measures 5 through 8, with chords Gm6, Bbm7, F7, and Fm7 indicated above the notes. Measure 8 ends with a whole rest.

**Fig. 15.** The first bars of *Sensitive*, a song composed in the style of Bossa Novas, but rendered in a slow pop ballad.

François Pachet and SKYGGE happened to love Antonio Carlos Jobim, the great Brazilian composer, and they wanted to do a cover of his famous song *One Note Samba* with Flow Machines.

SKYGGE put together generated stems with drums, bass and pads and right away got awesome results: a singular and catchy tune with great harmonies and timbre. The resulting harmonies are a bit different from the original but do not betray the logic of the song. The chord progression brings a Jobimian touch to the melody through unexpected harmonic modulations.

A few days before, the French band “The Pirouettes” had come to the studio and had uploaded two songs from their first album “Carrément, Carrément”. Vocals were generated from these recordings for *One Note Samba*, so fans of The Pirouettes may recognize words from their original songs.

(4) *Magic Man* (2:57)

SKYGGE fed Flow Machines with French pop songs from the 80s. The machine generated a simple melody, which sounded groovy once SKYGGE rendered it with a generated choir. The title comes from a phrase that comes back frequently in the choir: “*Magic Man*”. It was a nice surprise that the machine came up with a shiny pop song title with such an electro-disco feel. Flow Machines generated guitars from an American folk stem, as well as other vocals on the verse, and SKYGGE sung over those voices to get a more complex vocal blend. He also asked the singer Mariama to sing along with the generated choir to reinforce the groove. The lyrics are a mashup from all the generated syllables. French electro band Napkey worked on the arrangement at the end of the production, and Michael Lovett added synthesizer arpeggios.

(5) *In the House of Poetry* (4:14)

SKYGGE wanted to compose a song with the enchanting charm of ancient folk melodies. He fed Flow Machines with folk ballads and jazz tunes. As a result, the machine generated melodies with a chord progressions right in that mood, and a catchy and singular melodic movement (see Figure 16). Once the verse was done, he fed Flow Machines with a jazzier style for the chorus part, in order to bring in rich harmonic modulations.

Flow Machines then proposed an unconventional and audacious harmonic modulation in the first bar of the chorus, with an ascending melody illuminating the song (see Figure 17). SKYGGGE followed up with a small variation by exploiting a 2-bar pattern generated by Flow Machines and asked the system again to generate a harmonic progression that would resolve nicely with the tonality of the verse.




Fig. 16. The catchy and singular verse of “In the house of poetry”.

He subsequently asked Kyrie Kristmanson to join, hoping she would like the song and would not be afraid of its technically challenging nature. She was indeed enthusiastic and wrote lyrics inspired by the tale The Shadow by Andersen. She focused on the part of the story where the Shadow tells the learned man what he saw in the house of Poetry. The song is divided in two parts. In the first part Kyrie sings; in the second part, Kyrie’s vocals are generated by Flow Machines from recordings of Kyrie’s voice.



Fig. 17. The chorus of “In the house of poetry” features an unconventional and audacious ascending melody.

Flow Machines generated pianos, strings and vocals from SKYGGGE’s material as well as from Kyrie’s voice. SKYGGGE played all additional instruments (drums, piano, and electric guitars).

(6) Cryyyy (3:06)

This simple melody in the style of pop tunes from the 60s was generated almost exactly in its final form. When SKYGGGE was working in the studio of Ash Workman with Michael Lovett, Ash had just received an old cabinet organ from the 70s bought on the net. They plugged it in and began to play the chords of Cryyyy. It sounded great, and they recorded

all those sounds for the song. SKYGGE wanted a melancholic but modern sound and the timbre of Mariama’s voice matched perfectly. Flutes and detuned and distorted guitars were generated by Flow Machines, and SKYGGE added some beats and deep bass.

(7) Hello Shadow (2:54)

Stromae was fascinated by the possibilities of the software, and he fed the machines with his own influences: scores and audio stems in the Cape Verdian style. He selected his favorite melodies and stems from what Flow Machines had generated. Those fragments were put together and the song was built step by step. Stromae sung a vocal line that followed the generated melody, and he improvised on the pre-chorus. The choir in the chorus was also generated. When the song was ready for final production, it was sent to singer Kiesza who loved it. She wrote lyrics inspired by the tale “The Shadow”. Kiesza envisioned a happy, shiny shadow. A most unusual and characteristic feature of this song is its first four notes of the verse, which evoke the image of a ball bouncing and rolling.

(8) Mafia Love (Oyestershell) (2:02)

This song was composed by SKYGGE with Flow Machines, inspired by pop of the 60. It is almost a direct composition by Flow Machines, with very little human edits to the melody. The song has its internal logic, like all good songs. It tells a story, though unconventionally due to its rich structure, with almost no repetition. This absence of repetition sounds seem strange at first, but the song becomes an ear worm after hearing it a couple of times. The song was rendered with a generated voice from an a cappella recording of Curtis Clarke Jr., and a generated piano track from a stem by SKYGGE.

(9) Paper Skin (3:31)

Paper Skin is an interesting, indirect use of AI. This song was built from the song “Mafia Love” (see above). JATA picked up fragments of Mafia Love for the verse and the pre-chorus. He then composed a new chorus fitting those fragments. Ash Workman added some sounds from Mafia Love in the intro and in the bridge. Paper Skin is an offshoot of Mafia Love, illustrating how a melodic line can travel ears and be transformed according to unpredictable musicians inspirations.

(10) Multi Mega Fortune (3:39)

Michael Lovett from NZCA Lines fed Flow Machines with his own audio stems, vocals, drums loops, bass and keyboards, as well as lead sheets in the style of Brit pop. A lot of material was generated, both songs and stems. Michael Lovett and SKYGGE curated the results, and Lovett wrote lyrics inspired by the tale The Shadow. The result is a synth-pop catchy tune with a distinctive gimmick generated by Flow Machines, that runs throughout the song.

(11) Valise (3:22)

When Stromae came to the studio we tried several ideas based on six lead sheets generated with Flow Machines. Between sessions, SKYGGE

explored one of those directions and generated a vocal line from a former one of Stromae’s, uploaded on Flow Machines. The lyrics produced from the generated vocals meant something different from the original song by Stromae, but they were relevant, since they addressed the theme of “luggage” (Valise, in French)’. Stromae liked the song, but at the time we focused on the song Hello Shadow and left Valise aside for a while. SKYGGE asked the French band “The Pirouettes” to sing the melody instead of using the generated voice. The Pirouettes sung in sync with the generated voice by Stromae. This song bears an uncommon yet catchy chord progression and structure. Like other songs from the album, first hearings may sound strange, but after a couple of listenings, the song becomes an ear worm. One can hear the generated choir laughing on top of this unusual chord progression.

(12) Cake Woman (4:19)

Médéric Collignon, an amazing jazz trumpet player, came to the lab full of energy, with his own audio tracks, mostly jazz progressions played on a Rhodes piano, and some bass synths. SKYGGE also brought some hip-hop grooves, and the two musicians worked on a funk pop in the style of the 80s. The generated lead sheet was simple but contained harmonic twists that Médéric digs as a jazz composer. He selected chromatic modulations that he often uses in his own scores.

When they generated audio stems for the song from all the audio material, the output was messy but sounded very exciting. In particular, the groovy Rhodes generated from Médéric’s own recordings reinforced the funkiness of the song. Pachet and SKYGGE wrote lyrics inspired by the nonsensical words of the generated voice, in the style of Surrealist poetry. They asked the young and talented jazz singer Camille Bertault to sing the song. She also performed a scat-like improvisation, echoing the trumpet solo. In Andersen’s tale, the Shadow is hiding in the coat of a “cake woman”, hence the title.

(13) Whistle Theme (2:43)

Themes from older soundtracks are often more melodic than recent ones. Today, film scores are more often based on textures than melodies. Inspired by those old soundtracks, Flow Machines generated a catchy theme for this song (see Figure 18). The whistling was backed up by airports sounds generated to match the song. Pierre Jouan from the pop band “Catastrophe” sung another song, from which the machine generated the voice heard in the song.



Fig. 18. The catchy line of “Whistle song”.

## (14) Je Vais Te Manger (3:00)

Laurent Bardainne came to the studio with his audio stems, marimba, synth bass patterns and compositions in the style of 80's pop. Flow Machines generated a few songs, and Laurent selected the good parts. Laurent and SKYGGE built the song in a few hours. They left each other without knowing what to do with their song. For over a week, SKYGGE woke up every morning with this melody stuck in his head. Looking in depth at the generated lead sheet, one can see a harmonic twist that no one would have thought of. This twist pushes the melody up and down over an audacious modulation (see Figure19). SKYGGE and Laurent asked Sarah Yu to sing. The song is about a woman who says she will eat our souls and that we are lost.



**Fig. 19.** The catchy line of “Je Vais Te Manger”.

## (15) Cold Song (3:36)

The Cold Song is a well-known part of the opera “King Arthur” by Purcell. It has been sung by many singers in many styles (Sting, Klaus Nomi). For this cover, SKYGGE was inspired by artists such as Andre Bratten, Anne Clarke and Johann Johannsson, who have used machines to produce melancholic moods. The voice is generated from an a cappella recording of singer Kyte. It turns out that the generation produced many “A I A I”, by coincidence. In this song cover, everything was produced by Flow Machines, and there was no manual production.

## 11. Hello World Remix

A few months after the release of Hello world, a remix album was released<sup>38</sup>, with remixes of 10 songs, produced by various remixers. This album did not involve the use of AI, but is an important sign that the original songs of Hello World were deemed sufficiently rich musically by several top remixers to deserve a remix.

## 6 Unfinished but promising projects

The Flow Machines project generated a rich flow of ideas and experiments. Some of these ideas did not contribute directly to music making, but are worth

<sup>38</sup> available for instance at <https://open.spotify.com/album/5TAfJI5wypZ7JW1iyki6pI?si=GXsy6uQtReCWQFgtLpETQ>

mentioning, to illustrate how an ambitious goal supported by a generous grant (the ERC) can lead to very high levels of creativity and productivity.

### 1. Continuator II

The Continuator system, which was at the root of the Flow Machines project, was substantially extended, improved, and tested, notably during the Mirror project. Extensions consisted mostly to add a constraining features, notably positional, as described in Section 2.2, thereby improving substantially the musicality of the responses. Many experiments involving children and the question of music education based on reflexive interactions were conducted and described in [59]. A composition software (Mirror-Compo) was implemented, allowing users too build fully fledged musical compositions that would reuse their doodlings (see Figure 20).



**Fig. 20.** A composition software for children, that allows the to build pieces by reusing their previous doodlings.

### 2. Film Music

Experiments in Film Music by Pierre Roy consisted in implementing the theory proposed in [30]<sup>39</sup>. The basic idea of Murphy is to consider all pairs of consecutive chords, and assign them an emotional category. It turns out

<sup>39</sup> see also the clear videos examples at <https://www.youtube.com/watch?v=YSKAt3pmYBs&t=529s>



that enforcing such typical pairs of chords can be approximated by biases of binary factors of the belief propagation algorithm. For instance, sadness is related to the pattern  $M4m$ , i.e. a major chord, followed by a minor chord 4 semitones (2 tones) above. Figure 21 shows lead sheet a generated in the style of Bill Evans (for the melody) and the Beatles (for the chords) biased to favor chord pairs related to “sadness”<sup>40</sup>.

**Fig. 21.** A lead sheet in the style of the Beatles, with a bias towards “sadness” related chord pairs: C to Em (3 times), Bb to Dm.

### 3. Reflexive Looper

Reflexive Looper is the latest instantiation of a project consisting in building an interactive, real-system dedicated to solo performance. The initial fantasy was that of a guitarist playing in solo, but in need of a bass and harmony backing that would be tightly coupled to his real time performance. The core ideas of the system are described in [43,11] and many experiments and refactorors were done during several years [28], culminating by a version implemented by Marco Marchini [22]. This system was nominated for the Guthman Instrument Competition and was demonstrated to a jury including guitarist Pat Metheny.

### 4. Interaction studies in jazz improvisation

A question that arose regularly in our work about real time assisted improvisation was how to model actual humans interacting together when they improvise. It turns out that little is known about the nature of such interactions. In particular, a key question is to which extent this interaction involves the content of the music (rhythm, harmony, melody, expressiveness)? Such a question was crucial for designing of our music interaction systems. We proposed in [42] an analytical framework to identify correlates of content-based interaction. We illustrated the approach with the analysis of interaction in a typical jazz quintet (the Mark d’Inverno quintet). We extracted audio features from the signals of the soloist and the rhythm

<sup>40</sup> audio available at [https://www.francoispachet.fr/wp-content/uploads/2020/04/Beatles\\_Sadness\\_melody\\_Bill\\_Evans.mp3](https://www.francoispachet.fr/wp-content/uploads/2020/04/Beatles_Sadness_melody_Bill_Evans.mp3)

section. We measured the dependency between those time series with correlation, cross-correlation, mutual information, and Granger causality, both when musicians played concomitantly and when they did not. We identified a significant amount of dependency, but we showed that this dependency was mostly due to the use of a common musical context, which we called the *score effect*. We finally argued that either content-based interaction in jazz is a myth or that interactions do take place but at unknown musical dimensions. In short, we did not see any real interaction!

#### 5. Feedback in lead sheet composition

We studied the impact of giving and receiving feedback during the process of lead sheet composition in [24]. To what extent can peer feedback affect the quality of a music composition? How does musical experience influence the quality of feedback during the song composition process? Participants composed short songs using an online lead sheet editor, and are given the possibility to provide feedback on other participants' songs. Feedback could either be accepted or rejected in a later step. Quantitative data were collected from this experiment to estimate the relation between the intrinsic quality of songs (estimated by peer evaluation) and the nature of feedback. Results show that peer feedback could indeed improve both the quality of a song composition and the composer's satisfaction about it. Also, it showed that composers tend to prefer compositions from other musicians with similar musical experience levels.

#### 6. Experiments in collective creativity (aka the rabbits experiment)

In order to better understand the impact of feedback and collective creation, we designed an experiment to highlight different types of creator "profiles" [12]. In this experiment, users could write captions to short comic strips, and were shown what other users in a group would do, and given the possibility "switch" to their productions. One of the conclusions is that the potential impact of implicit feedback from other participants and objectivity in self-evaluation, even if encouraged, are lessened by a bias "against change". Such a bias probably stems from a combination of self-enhancing bias and of a commitment effect


#### 7. A Comic book on style creation

A comic book was designed and produced, telling a story about a character (Max) who struggles to invent her own style [13]. The name of the character stemmed from the Max Order result, which concerns the limitation of plagiarism in Markov sequence generation [47,49]. The comic is available on the web<sup>41</sup>. This comic gave rise to the ERCcOMICS<sup>42</sup> project, which was about doing the same exercise for 16 ERC projects.

#### 8. Song composition project

<sup>41</sup> <https://www.erccomics.com/comics/max-order>

<sup>42</sup> <http://www.erccomics.com>



**Fig. 22.** The character named MaxOrder, who tries to create her own style.

In this project, the author participated actively in the composition of 2 albums: one in Pop music (Marie-Claire<sup>43</sup>) and one in jazz (Count on it<sup>44</sup>). The albums were produced and distributed. Most composition sessions were recorded on video for later analysis. The jazz album was premiered in a sold-out concert at Pizza club express, one of the main jazz venue in London.

#### 9. Automatic Accompaniment generation

Following the Flow Composer project (see above), we addressed the issue of real time accompaniments using the same techniques. We recorded jazz pianist Ray d’Inverno (see Figure 23) as well as saxophonist Gilad Atzmon, and produced various accompaniment algorithms<sup>45</sup>.

#### 10. Interview series

Several interviews composers were conducted during the project: Ennio Morricone, Ivan Lins, Benoit Carré, Hubert Mounier (L’affaire Louis Trio), Michel Cywie (composer of beautiful French pop songs in the 70s), Franco Fabbri. The initial idea was to track down what composers of famous hits thought of the song writing process, to form a book about how songwriters make hits.

#### 11. Flow-Machines radio

A prototype of an online web radio was developed. This web generated on-the-fly pieces composed by Flow Composer, rendered them with our concatenative synthesis system (ReChord), and streamed them on a dedicated web site. Users could rate the songs they listened. The goal was to obtain enough feedback on songs to use reinforcement learning to tune automatically the generators (see Figure 24).

<sup>43</sup> The album Marie-Claire by band Marie-Claire: <https://open.spotify.com/album/6fXMeoGU8KXlGM6VWQtIq6?si=rqD64rtxRcepoXGywvjtBg>

<sup>44</sup> The album Count on it, by Mark d’Inverno quintet: <https://open.spotify.com/album/1VK1b1ZCYhFA6aKzmCLfIW?si=NSorGAU4S12nCzd2AMNHSg>

<sup>45</sup> For instance, [https://www.francoispachet.fr/wp-content/uploads/2020/04/Girl\\_from\\_Ipanema\\_mix.mp3](https://www.francoispachet.fr/wp-content/uploads/2020/04/Girl_from_Ipanema_mix.mp3) is an accompaniment generated automatically from the accompaniment on “Girl from Ipanema”



**Fig. 23.** Jazz pianist Ray d'Inverno recording an accompaniment on “Girl from Ipanema”.

Radio FM<sup>2</sup> - SONY CSL

Suggestions +

Michel Legrand 4/4 Style (150 BPM)	FlowRobot	fm <sup>2</sup>
Bossa Nova Style (140 BPM)	FlowRobot	fm <sup>2</sup>
Bossa Nova Style (110 BPM)	FlowRobot	fm <sup>2</sup>
Bill Evans 3/4 Style (170 BPM)	FlowRobot	fm <sup>2</sup>
Bill Evans 3/4 Style (120 BPM)	FlowRobot	fm <sup>2</sup>

Recently played

20:47:55	Charlie Parker Style (120 BPM)	0	0	↓
20:47:22	Charlie Parker Style (130 BPM)	0	0	↓
20:46:47	Miles Davis 4/4 Style (120 BPM)	0	0	↓
20:46:11	Irving Berlin 4/4 Style (120 BPM)	0	0	↓
20:45:41	Miles Davis 4/4 Style (140 BPM)	0	0	↓

Radio FM<sup>2</sup> by FlowComposer

listener: 1

FlowComposer

Charlie Parker Style (120 BPM)

B 7 Db m Gb 7 B 7 Db m Gb 7

Blues For Alexs-1 Confirmation-1 Chery1-1 Chery1-1

B 7 Db m Gb 7 B 7 Db m Gb 7 Chery1-1

**Fig. 24.** A snapshot of a web radio broadcasting automatically generated music, with user feedback.

## 7 Impact and followup

Three years after the completion of the project, what remains the most important for us in these developments is not the technologies we built *per se*, though many strong results were obtained in the “Markov + X” roadmap, but the music that was produced with these technologies. In that respect, the utmost validation of this work lies in the following:

1. The overwhelming media reception of the album composed with it<sup>46</sup>, including some outstanding reviews<sup>47</sup>
2. The enthusiasm of all the musicians who participated in this projet. 15 songs were composed and signed by various artists (including Stromae), and this is in itself a validation, as dedicated musicians would not sign songs they are not proud of
3. The overall success of the album on streaming platforms (a total of 12 million streams as of April 2020), with Magic Man (5.8 million streams) and Hello Shadow (2.8 million) as main hits
4. The positive reception by music critics. For instance, Avdeef writes[3]: SKYGGGE’s Hello World is a product of these new forms of production and consumption, and functions as a pivot moment in the understanding and value of human-computer collaborations. The album is aptly named, as it alludes to the first words any new programmer uses when learning to code, as well as serving as an introduction to new AI-human collaborative practices. Hello, World, welcome to the new era of popular music.

Similarly, emphasising the difference between interactive AI-assisted song composition, which Flow Composer pioneered, and fully automatic composition, Miller writes [25]:

On the one hand, we have François Pachet’s Flow Machines, loaded with software to produce sumptuous original melodies, including a well-reviewed album. On the other, researchers at Google use artificial neural networks to produce music unaided. But at the moment their music tends to lose momentum after only a minute or so.

5. Sony Flow Machines  
The Flow Machines project is continuing at Sony CSL<sup>48</sup>, and research around the ideas developed in the project have been continued, for instance with new interactive editors in the same vein as Flow Composer [6]. Promising musical projects have also been setup, notably involving Jean-Michel Jarre <sup>49</sup>.
6. Beyond Flow Machines  
Since Flow Machines and Hello World, SKYGGGE has launched another album made with Artificial Intelligence, American Folk Songs<sup>50</sup>. In this album,

<sup>46</sup> <https://www.helloworldalbum.net/press/>

<sup>47</sup> e.g. “AI and humans collaborated to produce this hauntingly catchy pop music” by Melissa Avdeef, or “Is this the world’s first good robot album?” by Alex Marshall

<sup>48</sup> <https://www.sony CSL.co.jp/tokyo/2811/> and [www.flow-machines.com](http://www.flow-machines.com)

<sup>49</sup> <https://www.bbc.com/news/entertainment-arts-50335897#>

<sup>50</sup> <https://open.spotify.com/album/6NbX54o0pEZhs0jfdSYepw?si=RcTgqFqgTfC5iXkVmlW2Ug>

original a cappella songs of traditional songs (notably by Pete Seeger, with his right owner’s approval) were reorchestrated with automatic orchestration tools developed at Spotify. The techniques used are based on a belief propagation scheme [62] combined with advanced edition features [61]. Ongoing work to use large-scale listening data (skip profiles) to tune the model are ongoing with promising results [26].

## 8 Lessons learned

### 8.1 Better model does not imply better music

We experimented with various generative models: variable-order Markov models, max entropy and deep learning techniques. Paradoxically, the most interesting music generation were not obtained with the most sophisticated techniques. In spite of many attempts to use deep learning methods upfront for generation, we did not succeed in using these techniques for professional use (at the time of the project). The Bach chorales are of undisputed high quality, but they are not very *interesting* musically: they produce either “correct” material, or mistakes, but not much in between. However, we observed that the combination of various tools (*e.g.*, the lead sheet generation tool and our audio synthesis tool) produced flow states more easily than the use of single algorithms. This may have to do with the appropriation effect mentioned below.

### 8.2 New creative acts

The use of generative algorithms introduces new tasks for the user, and these tasks are actually creative acts, in the sense that they require some musical expertise or, at least intuition and intention: 1) the selection of training sets and 2) the curation of the results. Of course, training set selection and curation can be performed using random algorithms, but so far we did not succeed in automating these processes.

### 8.3 The appropriation effect

More importantly, there seems to be a tradeoff to find between the sophistication of the algorithm and the sense of appropriation of the results by the user. The IKEA effect is a cognitive bias which has been confirmed by many experiments, in which consumers place a disproportionately high value on products they partially created [20]. Many consumer studies have shown for instance that instant cake mixes were more popular when the consumer had something to do (such as adding an egg) than when the instant mix was self sufficient. It is likely that in the case of AI assisted creation, the same type of effect applies, and that users do not get a sense of appropriation if the algorithm, whatever its sophistication, does all the creative job. This is a fascinating area of study that will surely produce counter-intuitive results in the future.

## 9 Toward new Categories of *new*

The traditional conceptual landscape used to describe novelty in music is based essentially on three notions: 1) “original” songs (something new that may be inspired by preexisting work, but to an acceptable degree), 2) “covers” (new orchestrations of existing songs), 3) “plagiaristic work”, i.e. works containing segments of existing songs with little transformations <sup>51</sup>. A lot of arguments used in the debate about the role of AI in music creation touch on the notion of creativity or novelty, i.e. can AI produce really novel music material. However, we don't think the question “is AI creative” is relevant, since we see AI as a tool for creators, an opinion shared by most researchers in the field (*e.g.* [9]). Yet, we argue the idea that all the generations we described here, and the ones to come, are instances of new categories of “new”, and some of them are new in themselves. In this text, we have described several pieces generated with various techniques, and used intentionally the following terms (highlighted in red):

1. stylistic explorations: Song snippets “in the style of”, Sensitive
2. stylistic singularities: the Boulez Blues
3. reminiscences: Scratch my itch
4. pastiches and exercises: Daddy’s car and Bach chorales
5. timbre transfers: most of the renderings of the examples described here
6. templagiarism: Lunar

It is our view that these new categories may help us conceptualize the contributions of these generative technologies to music creation. Interestingly, images created with Deep Dream [27] and referred to as “hallucinations” do not have their equivalent yet in music: what would be, indeed a musical hallucination ?

## 10 Conclusion

We have described a number of research results in the domain of assisted music creation, together with what we consider are remarkable music productions. We attempted to describe why these generations are interesting, as technological artefacts, and also from a musical viewpoint. Since the Flow Machines project ended (2017), research in computer music generation has exploded, mostly due to the progress in deep learning [7]. After the launch of Daddy’s car and “Hello World”, several music titles were produced with AI <sup>52</sup>, contributing to the exploration of these technologies for music creation.

These explorations are still only scratching the surface of what is possible with AI. However, they already challenge the status and value of what is considered “new” in our digital cultures. We have sketched a draft of a vocabulary to

<sup>51</sup> Another category is “parasiting”, the action of imitating an orchestration style, without plagiarising it, in order to avoid paying royalties [21], but this category is apparently not used systematically by music professionals

<sup>52</sup> notably by artists Taryn Southern and Holly Herndon

describe and distinguish different type of “newness” in music. These categories start to have well defined meaning technically but much more is needed to give them precise definitions. Surely, more categories will arise, some will vanish. In the end, only music and words will remain, not technologies.

## 11 Acknowledgements

The Flow Machines project received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 291156. We thank the team of the musical Beyond the Fence for their insightful comments in using earlier versions of Flow Composer. We thank the researchers and engineers who participated in the Flow Machines project: Marco Marchini, Mathieu Ramona, Alexandre Papadopoulos, Jason Sakellariou, Dani Martín, Vincent Degroote, Gabriele Barbieri and Fiammetta Ghedini. We also thank Professor Mirko Degli Esposti from University Bologna, Mark d’Inverno from Goldsmiths College, University of London, Jean-François Perrot and Jean-Pierre Briot from CNRS and Sorbonne Université, Luc Steels from VUB, for their precious and friendly contributions. We thank our colleagues and friends Giordano Cabral, Geber Ramalho, Jean-Charles Régin, Guillaume Perez, Shai Newman. We also thank all the musicians who were involved in the projects, under the supervision of Benoit Carré: Stromae, Busy P, Catastrophe, ALB, Kumisolo, Olivier Marguerit, Barbara Carlotti, Housse de Racket, Camille Bertault, Lescop, Raphael Chassin, Michael Lovett, Ash Workman, C Duncan, Marie-Jeanne Serrero, Fred Decès, Mariama, Kyrie Kristmanson, The Bionix, Kiesza, JATA, The Pirouettes, Médéric Collignon, Pierre Jouan, Sarah Yu Zeebroeke, Laurent Bardainne. Finally, we thank our host institutions Sony CSL, UPMC, Spotify, the label Flow Records and the distributor Idol for hosting, supporting and valorising this research.

## References

1. Addessi, A.R., Pachet, F.: Experiments with a musical machine: musical style replication in 3 to 5 year old children. *British Journal of Music Education* **22**(1), 21–46 (2005). <https://doi.org/10.1017/S0265051704005972>
2. Allen, J.F.: Maintaining knowledge about temporal intervals. *Commun. ACM* **26**(11), 832–843 (Nov 1983). <https://doi.org/10.1145/182.358434>, <https://doi.org/10.1145/182.358434>
3. Avdeef, M.: Artificial intelligence and popular music: SKYGGGE, Flow Machines, and the audio uncanny valley. *Arts* **8**(130) (2019)
4. Barbieri, G.: Remixing Dylan lyrics with style. <https://redylan.neocities.org>, accessed: 2020-02-09
5. Barbieri, G., Pachet, F., Roy, P., Esposti, M.D.: Markov Constraints for Generating Lyrics with Style. In: et al., L.D.R. (ed.) *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI 2012)*. pp. 115–120 (2012)
6. Bazin, T., Hadjeres, G.: NONOTO: A model-agnostic web interface for interactive music composition by inpainting. *CoRR* **abs/1907.10380** (2019)



7. Briot, J.P., Hadjeres, G., Pachet, F.D.: Deep learning techniques for music generation. Springer (2019)
8. Colton, S., Llano, M.T., Hepworth, R., Charnley, J., Gale, C.V., Baron, A., Pachet, F., Roy, P., Gervas, P., Collins, N., Sturm, B., Weyde, T., Wolff, D., Lloyd, J.: The Beyond The Fence Musical and Computer Says Show Documentary. In: 7th International Conference on Computational Creativity (ICCC 2016). Paris (France) (June 2016)
9. Fiebrink, R., Caramiaux, B.: The machine learning algorithm as creative musical tool. ArXiv [abs/1611.00379](https://arxiv.org/abs/1611.00379) (2016)
10. Fisch, S.K.: Hans zimmer: Music is a conversation. <https://scottkfish.com/2017/03/19/hans-zimmer-music-is-a-conversation/>, accessed: 2020-02-09
11. Foulon, R., Roy, P., Pachet, F.: Automatic classification of guitar playing modes. In: Aramaki, M., Derrien, O., Kronland-Martinet, R., Ystad, S. (eds.) Sound, Music, and Motion - 10th International Symposium, CMMR 2013, Marseille, France, October 15-18, 2013. Revised Selected Papers. Lecture Notes in Computer Science, vol. 8905, pp. 58–71. Springer (2013). [https://doi.org/10.1007/978-3-319-12976-1\\_4](https://doi.org/10.1007/978-3-319-12976-1_4)
12. Ghedini, F., Frantz, B., Pachet, F., Roy, P.: The comic strip game: observing the impact of implicit feedback in the content creation process. In: Osman, N., Yee-King, M. (eds.) Proceedings of the First International Workshop on AI and Feedback, AInF 2015, co-located with the 24th International Joint Conference on Artificial Intelligence (IJCAI 2015), Buenos Aires, Argentina, July 26, 2015. CEUR Workshop Proceedings, vol. 1407, pp. 18–26. CEUR-WS.org (2015), <http://ceur-ws.org/Vol-1407/AInF2015paper3.pdf>
13. Ghedini, F., Pachet, F., Roy, P.: Max order: A tale of creativity. In: Yang, Q., Wooldridge, M.J. (eds.) Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015. pp. 4136–4137. AAAI Press (2015), <http://ijcai.org/Abstract/15/581>
14. Hadjeres, G., Nielsen, F.: Anticipation-RNN: enforcing unary constraints in sequence generation, with application to interactive music generation. *Neural Computing and Applications* **32**(4), 995–1005 (2020). <https://doi.org/10.1007/s00521-018-3868-4>, <https://doi.org/10.1007/s00521-018-3868-4>
15. Hadjeres, G., Pachet, F., Nielsen, F.: Deepbach: a steerable model for Bach chorales generation. In: Precup, D., Teh, Y.W. (eds.) Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017. Proceedings of Machine Learning Research, vol. 70, pp. 1362–1371. PMLR (2017), <http://proceedings.mlr.press/v70/hadjeres17a.html>
16. Hadjeres, G., Sakellariou, J., Pachet, F.: Style imitation and chord invention in polyphonic music with exponential families. *CoRR* [abs/1609.05152](https://arxiv.org/abs/1609.05152) (2016), [http://arxiv.org/abs/1609.05152](https://arxiv.org/abs/1609.05152)
17. Hiller, L.A., Isaacson, L.M.: *Experimental Music: Composition with an Electronic Computer*. McGraw-Hill (1959)
18. Jaynes, E.T.: Information theory and statistical mechanics. *Phys. Rev.* **106**(4), 620–630 (May 1957). <https://doi.org/10.1103/PhysRev.106.620>, [http://prola.aps.org/abstract/PR/v106/i4/p620\\_1](http://prola.aps.org/abstract/PR/v106/i4/p620_1)
19. Khatchatourov, A., Pachet, F., Rowe, V.: Action identity in style simulation systems: Do players consider machine-generated music as of their own style? *Frontiers in Psychology* **7** (2016)
20. Landers, R.: Unfolding the ikea effect: Why we love the things we build. *Journal of Consumer Psychology* **22**(3), 453–460 (2012)

21. Le Tavernier, S.: Private conversation (2016)
22. Marchini, M., Pachet, F., Carré, B.: Rethinking Reflexive Looper for structured pop music. In: NIME (2017)
23. Martín, D., Neullas, T., Pachet, F.: LeadsheetJS: A Javascript library for online lead sheet editing. In: Proceedings of the first TENOR Conference (05 2015)
24. Martín, D., Pachet, F., Frantz, B.: The Creative Process in Lead Sheet Composition: Perspectives from Multiple Domains, pp. 275–296. Palgrave Macmillan (08 2018)
25. Miller, A.: Creativity and AI: the next step. *Scientific American* (October 2019)
26. Montecchio, N., Roy, P., Pachet, F.: The skipping behavior of users of music streaming services and its relation to musical structure. *CoRR* **abs/1903.06008** (2019), <http://arxiv.org/abs/1903.06008>
27. Mordvintsev, A., Olah, C., Tyka, M.: Inceptionism: Going deeper into neural networks (2015), <https://research.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html>
28. Moreira, J., Roy, P., Pachet, F.: Virtualband: Interacting with stylistically consistent agents. In: de Souza Britto Jr., A., Gouyon, F., Dixon, S. (eds.) Proceedings of the 14th International Society for Music Information Retrieval Conference, ISMIR 2013, Curitiba, Brazil, November 4-8, 2013. pp. 341–346 (2013), [http://www.ppgia.pucpr.br/ismir2013/wp-content/uploads/2013/09/277\\_Paper.pdf](http://www.ppgia.pucpr.br/ismir2013/wp-content/uploads/2013/09/277_Paper.pdf)
29. Moulieras, S., Pachet, F.: Maximum entropy models for generation of expressive music. *CoRR* **abs/1610.03606** (2016), <http://arxiv.org/abs/1610.03606>
30. Murphy, S.: Transformational theory and the analysis of film music. In: *The Oxford Handbook of Film Music Studies*. Oxford University Press (2014)
31. Pachet, F.: Music interaction with style. In: in SIGGRAPH 2003 Conference Abstracts and Applications. ACM Press, San Diego, USA (August 2003)
32. Pachet, F.: Musical virtuosity and creativity. In: McCormack, J., D’Inverno, M. (eds.) *Computers and Creativity*. Springer (2012)
33. Pachet, F., Roy, P.: Markov constraints: steerable generation of Markov sequences. *Constraints* **16**(2) (2011)
34. Pachet, F., Roy, P.: Imitative Leadsheet Generation with User Constraints. In: Proceedings of the 22th European Conference on Artificial Intelligence (ECAI 2014). Prague, Czech Republic (2014)
35. Pachet, F., Roy, P.: Non-conformant harmonization: the real book in the style of Take 6. In: *International Conference on Computational Creativity*. Ljubljana (2014)
36. Pachet, F., Suzda, J., Martín, D.: A comprehensive online database of machine-readable leadsheets for Jazz standards. In: de Souza Britto Junior, A., Gouyon, F., Dixon, S. (eds.) Proceedings of the 14th International Society for Music Information Retrieval Conference (ISMIR 2013). pp. 275–280. ISMIR, Curitiba, PA, Brazil (Nov 2013)
37. Pachet, F.: The Continuator: Musical interaction with style. *Journal of New Music Research (JNMR)* **32**(3), 333–341 (2003)
38. Pachet, F.: A joyful ode to automatic orchestration. *ACM TIST* **8**(2), 18:1–18:13 (2017). <https://doi.org/10.1145/2897738>, <https://doi.org/10.1145/2897738>
39. Pachet, F., Addressi, A.R.: When children reflect on their own playing style: experiments with continuator and children. *Computers in Entertainment* **2**(1), 14 (2004). <https://doi.org/10.1145/973801.973824>, <https://doi.org/10.1145/973801.973824>

40. Pachet, F., Papadopoulos, A., Roy, P.: Sampling variations of sequences for structured music generation. In: Cunningham, S.J., Duan, Z., Hu, X., Turnbull, D. (eds.) Proceedings of the 18th International Society for Music Information Retrieval Conference, ISMIR 2017, Suzhou, China, October 23-27, 2017. pp. 167–173 (2017), [https://ismir2017.smcnus.org/wp-content/uploads/2017/10/50\\_Paper.pdf](https://ismir2017.smcnus.org/wp-content/uploads/2017/10/50_Paper.pdf)
41. Pachet, F., Roy, P., Barbieri, G.: Finite-length Markov processes with constraints. In: Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI 2011). pp. 635–642. Barcelona, Spain (Jul 2011)
42. Pachet, F., Roy, P., Foulon, R.: Do jazz improvisers really interact?: The score effect in collective jazz improvisation. In: The Routledge companion to embodied music interaction, pp. 167–176. Routledge (2017)
43. Pachet, F., Roy, P., Moreira, J., d’Inverno, M.: Reflexive loopers for solo musical improvisation. In: Mackay, W.E., Brewster, S.A., Bødker, S. (eds.) 2013 ACM SIGCHI Conference on Human Factors in Computing Systems, CHI ’13, Paris, France, April 27 - May 2, 2013. pp. 2205–2208. ACM (2013). <https://doi.org/10.1145/2470654.2481303>, <https://doi.org/10.1145/2470654.2481303>
44. Pachet, F., Roy, P., Papadopoulos, A., Sakellariou, J.: Generating 1/f noise sequences as constraint satisfaction: The Voss constraint. In: IJCAI. pp. 2482–2488. AAAI Press (2015)
45. Pachet, F., Roy, P., Papaopoulos, A.: Comments on “assisted lead sheet composition using flowcomposer”. Virtual Volume, 25th anniversary of the CP Conference (2019)
46. Pachet, F., Roy, P., Papadopoulos, A., Briot, J.P.: Constrained Markov Sequence Generation Applications to Music and Text. Springer (2019)
47. Papadopoulos, A., Pachet, F., Roy, P.: Generating non-plagiaristic markov sequences with max order sampling. In: Creativity and Universality in Language, pp. 85–103. Springer (2016)
48. Papadopoulos, A., Pachet, F., Roy, P., Sakellariou, J.: Exact sampling for regular and markov constraints with belief propagation. In: CP. Lecture Notes in Computer Science, vol. 9255, pp. 341–350. Springer (2015)
49. Papadopoulos, A., Roy, P., Pachet, F.: Avoiding Plagiarism in Markov Sequence Generation. In: AAAI. AAAI Press (2014)
50. Papadopoulos, A., Roy, P., Pachet, F.: Assisted lead sheet composition using FlowComposer. In: Rueher, M. (ed.) Principles and Practice of Constraint Programming: 22nd International Conference, CP 2016, Toulouse, France, September 5-9, 2016, Proceedings. pp. 769–785. Programming and Software Engineering, Springer International Publishing (2016)
51. Papadopoulos, A., Roy, P., Régin, J., Pachet, F.: Generating all possible palindromes from ngram corpora. In: Yang, Q., Wooldridge, M.J. (eds.) Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015. pp. 2489–2495. AAAI Press (2015), <http://ijcai.org/Abstract/15/353>
52. Pearl, J.: Reverend Bayes on inference engines: A distributed hierarchical approach. In: AAAI. pp. 133–136. AAAI Press (1982)
53. Perez, G., Rappazzo, B., Gomes, C.P.: Extending the capacity of 1 / f noise generation. In: CP. Lecture Notes in Computer Science, vol. 11008, pp. 601–610. Springer (2018)
54. Perez, G., Régin, J.: MDDs: Sampling and probability constraints. In: CP. Lecture Notes in Computer Science, vol. 10416, pp. 226–242. Springer (2017)

55. Pesant, G.: A regular language membership constraint for finite sequences of variables. In: CP. Lecture Notes in Computer Science, vol. 3258, pp. 482–495. Springer (2004)
56. Ramona, M., Cabral, G., Pachet, F.: Capturing a musician’s groove: Generation of realistic accompaniments from single song recordings. In: Yang, Q., Wooldridge, M.J. (eds.) Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25–31, 2015. pp. 4140–4142. AAAI Press (2015), <http://ijcai.org/Abstract/15/583>
57. Régin, J.C.: A Filtering Algorithm for Constraints of Difference in CSPs. In: Proceedings of the 12th National Conference on Artificial Intelligence (Vol. 1). pp. 362–367. AAAI ’94, American Association for Artificial Intelligence, Menlo Park, CA, USA (1994)
58. Rivaud, S., Pachet, F.: Sampling Markov models under constraints: Complexity results for binary equalities and grammar membership. CoRR **abs/1711.10436** (2017), <http://arxiv.org/abs/1711.10436>
59. Rowe, V., Triantafyllaki, A., Pachet, F.: Children’s Creative Music-Making with Reflexive Interactive Technology. Routledge (2016)
60. Roy, P., Pachet, F.: Enforcing Meter in Finite-Length Markov Sequences. In: de-Jardins, M., Littman, M.L. (eds.) AAAI. AAAI Press (2013)
61. Roy, P., Pachet, F.: Smart edition of MIDI files. CoRR **abs/1903.08459** (2019), <http://arxiv.org/abs/1903.08459>
62. Roy, P., Pachet, F., Carré, B.: A robust, controllable music orchestrator that’s fun to use, with commented examples. CoRR **to appear** (2020)
63. Roy, P., Papadopoulos, A., Pachet, F.: Sampling variations of lead sheets (2017)
64. Roy, P., Perez, G., Régin, J., Papadopoulos, A., Pachet, F., Marchini, M.: Enforcing structure on temporal sequences: The Allen constraint. In: CP. Lecture Notes in Computer Science, vol. 9892, pp. 786–801. Springer (2016)
65. Sakellariou, J., Tria, F., Loreto, V., Pachet, F.: Maximum entropy models capture melodic styles. Scientific reports **7**(1), 1–9 (2017)
66. Voss, R.F., Clarke, J.: ‘1/f noise’ in music and speech. Nature **258**(5533), 317–318 (1975), <https://doi.org/10.1038/258317a0>
67. Voss, R.F., Clarke, J.: “1/f noise” in music: Music from 1/f noise. The Journal of the Acoustical Society of America **63**(1), 258–263 (1978). <https://doi.org/10.1121/1.381721>, <https://asa.scitation.org/doi/abs/10.1121/1.381721>