

# Interactive Musical Partner: A Modular Human/Computer Duo Improvisation System

Jeffrey Albert

Loyola University New Orleans,  
College of Music and Fine Arts  
6363 St. Charles Ave., New Orleans, LA 70118 USA  
jvalbert@loyno.edu  
<http://www.loyno.edu/~jvalbert>

**Abstract.** The Interactive Musical Partner (IMP) is software designed for use in duo improvisations, with one human improviser and one instance of IMP, focusing on a freely improvised duo aesthetic. IMP has Musical Personality Settings (MPS) that can be set prior to performance, and these MPS guide the way IMP responds to musical input. The MPS also govern the probability of particular outcomes from IMP's creative algorithms. The IMP uses audio data feature extraction methods to listen to the human partner, and react to, or ignore, the human's musical input, based on the current MPS. This article presents the basic structure of the components of IMP: synthesis module, musical personality settings, listener system, creative algorithm, and machine learning implementation for timbral control.

**Keywords:** Improvisation, Interaction, Computer Generated Music

## 1 Introduction

The Interactive Musical Partner (IMP) is software designed for use in duo improvisations, with one human improviser and one instance of IMP, focusing on a freely improvised duo aesthetic. Two concepts are definitive for IMP. IMP is a monophonic participant in non-idiomatic improvisation (or free improvisation), and IMP deals with the ways it hears, remembers, and creates musical content with the fewest possible levels of abstraction. Whenever possible pitches are dealt with in terms of frequency and durations in terms of milliseconds. By avoiding thinking in terms of note names and note values, IMP can more easily navigate the spaces outside of tonality and tempo.<sup>1</sup> IMP strives to function in the aesthetic lineage of the freely improvised duo. David Borgo describes the music, "often dubbed 'free improvisation'" as tending to, "devalue the two dimensions that have traditionally dominated music representation - quantized

---

<sup>1</sup> There are two exceptions to this principle in which musical data for the generative algorithm is stored in an abstracted form.

pitch and metered durations - in favor of the microsubtleties of timbral and temporal modification” [4]. This is an accurate description of the musical priorities of IMP.

Of course, IMP is not the first interactive music system. There have been systems that use information from the human performer to control or influence the computer’s output, like Jean-Claude Risset’s *Duo for Piano* [11], and Richard Teitelbaum’s *Concerto Grosso for Human Concertino and Robotic Ripieno* [12]. Other systems build an output style based on the input of the human partner, or a preloaded corpus. Francois Pachet’s Continuator [9] continues the phrase of the human, and the OMax software developed at IRCAM [3] is described as co-improvising with the human.

The common thread through these previously mentioned systems is that they are all dependent upon human input. A system that interacts with human input, but does not depend upon it, is the *Voyager* system by George E. Lewis. Lewis describes *Voyager*’s structure: ”as multiple parallel streams of music generation, emanating from both the computers and the humans - a nonhierarchical, improvisational, subject-subject model of discourse, rather than a stimulus/response setup” [8]. In *Voyager*, the computer and human are equal in terms of agency; either one can make music without the other.

IMP is philosophically most similar to *Voyager*, in that IMP is not dependent upon a human partner, it simply interacts with one. IMP could play music on its own. Unlike *Voyager*, IMP has but one voice. With IMP’s single voice synthesis, a performance with IMP might superficially sound more like OMax or the Continuator, but its philosophy of interactivity and structure are much more like *Voyager*.

IMP was programmed in the Max 6 Programming environment, using externals by Tristan Jehan [7] and Tap.Tools [10], and uses The Wekinator [6] for machine learning implementation. IMP consists of: a synthesis module, a Musical Personality Settings (MPS) module, a frequency decider, a duration decider, a global variation module, a listener module, and a number of smaller decider modules. I will use the term creative algorithm to refer to the aspects of the duration decider and frequency decider that control IMP’s autonomous output.

## 2 Synthesis Module

The synthesis module is IMP’s voice. This is the section of the software that makes the sounds, and manages the messages sent from the other modules. The frequency decider, duration decider, global variation module, and listener module are all sub-patches of the synthesis module.

The synthesis module uses frequency modulation (FM) synthesis to generate its sounds [5]. IMP also employs a second order FM synthesis, meaning that there is a second modulating oscillator, which modulates the product of the first order oscillators. Variations of the combination of these two FM pairs are how IMP controls timbral variety. The primary FM pair is usually set to a simple harmonic ratio; 2 is the default. A more complex ratio on the second FM

pair allows for IMP to make its sound more complex and strident by adjusting the amount of modulation from the second modulating oscillator.<sup>2</sup> This mechanism will be discussed in greater detail in the section on Machine Learning and Timbral Interaction.

The synthesis module also contains three smaller decision modules: the `density_decider`, the `volume_decider`, and the sub-patcher called `silent_event`. The `density_decider` is the module that controls the density of events. To begin a new event (note), a decision is made as to whether the event will make sound, or be silent. This decision is weighted by the values in the density parameter of the MPS. The densest setting will have every event make sound, and the least dense setting will have no events make sound. If the next event is to have sound, the `density_decider` sends a bang on the `next_now` channel, which cues the synthesis process, and if the event is to be silent, a bang is sent to the `silent_event` sub-patcher. The `silent_event` sub-patcher receives the next duration from the `duration_decider` and lets that amount of time pass in silence before sending a bang to the `density_decider` to cue the decision process for the next event. This system allows for the MPS to control the density of the texture without necessarily changing any of the other duration parameters, so it is possible for IMP to play sparsely in a setting that still has a relatively short duration of events, or densely in a setting that has relatively long event durations.

### 3 Musical Personality Settings

One of the goals of this research was to design a system with variable sets of behavioral characteristics, or musical personalities. This is implemented in IMP through the Musical Personality Settings (MPS), which are seven separate parameters that influence various aspects of IMP's behavior. The parameters are: Density of Events, Length of Events, Rhythmic Regularity, Frequency Listeneriness, Duration Listeneriness, Melodicness, and Variation.

These parameters were selected so that various aspects of IMP's behavior could be controlled independently. The MPS are used to weight decisions made in the generative algorithms, which generate the actual data used to create the sounds (frequency, duration, etc). The MPS affect musical output by affecting the ways in which the generative algorithm makes its decisions. Each MPS parameter is controlled with a slider on the MPS interface. The interface also contains the mechanism for setting the length of an episode (or performance), and a visible timer to give the performer a reference for the amount of elapsed time since the beginning of the episode.

---

<sup>2</sup> The choice of FM synthesis as IMP's voice was as much an aesthetic decision as a technical one. While it was tempting to try to design IMP with a more organic voice, in part to try to make the human performer forget that IMP was not in fact human, I ultimately decided that giving IMP a voice that would remind the performer that IMP was not human was a better path. This decision is covered in greater detail in my thesis[2].

The Density of Events parameter controls the weighting of the `density_decider`'s decision algorithm, which decides whether an event will make sound or not. The higher this parameter is set the higher the sound to silence ratio will be. This parameter is also influenced by what is heard from the human, once an episode begins.

The Length of Events and Rhythmic Regularity parameters work together to control IMP's tempo and sense of pulse. I use these terms (tempo and sense of pulse) loosely in this context, since there is no abstraction of meter present, but there can be a sense of IMP playing faster or slower, and in more or less regular event lengths. The Rhythmic Regularity parameter controls a pool from which duration proportions are chosen in the creative algorithm, and the Length of Events parameter controls a factor that controls the speed at which these proportions are realized.

Listeneriness is a term I have coined to describe the two parameters that control IMP's responsiveness to human input. The lower the listeneriness value, the more independently IMP will behave, and the higher the value, the more IMP will derive its output from what it has heard from the human. There are two listeneriness settings; one for frequency and one for duration. Frequency Listeneriness controls the weighting of the frequency decider mechanisms and influences whether IMP's pitch output is derived from its creative algorithm or from the pool of pitches it remembers hearing from the human. Duration Listeneriness controls the weighting of the duration decider and similarly influences IMP's output in terms of duration of events.

The Melodicness parameter sets a set of pitches from which the creative algorithm chooses when IMP is generating content on its own. As the value moves from low to high the pool of available pitches moves from pentatonic sets, through major scales, melodic minor (ascending) scales, diminished scales, whole tone scales, and finally to a fully chromatic set of pitches.

The final MPS parameter is Variation. This parameter weights the decisions made by the global variation module, which controls a mechanism that causes variation of the other MPS parameters. The most often the parameters will change is once per second, and the least often they will change is every 100 seconds, with the largest possible jump on any MPS scale being 10 units, on a 128 unit scale. This keeps IMP's output from seeming static in content, but helps avoid seemingly random huge shifts in musical space as well.

## 4 Listener System

The listener module receives in the incoming audio signal from the human via the ADC, performs the audio feature extractions, and sends that extracted feature data to other IMP modules or to The Wekinator. The central component in the listener module is the `analyzer~` object, which is a Max/MSP extension programmed by Tristan Jehan. The `analyzer~` object outputs seven different audio features, and IMP uses five of those: pitch, loudness, brightness, noisiness, and attack.

Pitch is output as frequency in Hz, which is stored in the frequency decider module. Onset detection is done using a combination of pitch and amplitude analysis. A bang is sent out of the attack outlet whenever a new onset is detected. This onset bang serves two important functions. The first is that it cues the current pitch to be sent to the `heard_freq` channel. The second is that it is sent into the `onset_average` subpatcher, which is used to keep a running average of the time between the last ten onsets detected from the human input, as well as to send the elapsed time between each individual onset on the `heard_dur` channel, which goes to the duration decider list of heard durations. The loudness, brightness, and noisiness features are sent via `OpenSoundControl[1]` to the `Wekinator`, where they are used to control the timbral interaction system.

## 5 Creative Algorithm

There are two primary modules in the creative algorithm: the frequency decider, and the duration decider. Each functions very similarly, but their processes do not affect each other. The beginning of the new event process works similarly on both the frequency and duration deciders. A bang on the `next_now` channel cues each new event. That bang causes a decision to be made as to whether the next event will come from IMPs generative algorithm or IMPs pool of heard events. This is the way in which IMP can play off of what it has heard, or introduce independent new material. The frequency and duration deciders each have two sides: one side is a list of heard data, and the other side is the generative algorithm, or the part that makes IMP's original output.

On the heard data side of the duration decider, each new heard duration is entered into a list that keeps the last ten heard durations. When the `next_now` bang is sent to the heard data side, one of these ten most recently heard durations is selected randomly, and output as the next duration on the `next_dur` channel, which is received in the synthesis module and stored in the amplitude envelope until the `next_now` triggers an event and a new duration is sent. This side of the decider also keeps a running average of the list of heard durations that can be used to change the length and density MPS. This average is scaled to fit the MPS values, and every 500 ms a decision is made to change or not change the MPS based on the current average of heard durations. The average is scaled differently for the length and duration MPS, and the decision to change each MPS is made independently. This system keeps IMP in a similar density and speed area as the human improviser, but does allow for some divergence in terms of these parameters as well. It has the effect of varying how much it seems like IMP is following or ignoring the human.

On the generative algorithm side of the duration decider, IMP chooses a duration proportion from a set of lists of proportions, and that proportion is multiplied by a length factor to get the next duration. There are 15 different files of proportion values, one of which is loaded into the duration `coll` object (`dur_coll`) based on the rhythmic regularity MPS. The lower the rhythmic regularity MPS value the more varied the proportions are. The more similar the

duration proportions, the more of a sense of pulse one hears from IMP. There is a rhythmic variation decider that uses the variation MPS to change the choice of proportion `coll` file during the course of an episode. If the `next_now` bang is sent to the generative algorithm side of the duration decider, a proportion is output from the currently loaded duration `coll`.

Once a proportion is sent from the duration `coll`, it is multiplied by a length factor. This length factor is controlled by the length MPS. The shortest duration that IMP will create from its generative algorithm is 50 ms, and the longest is 2500 ms (2.5 seconds). This proportion/length factor system allows IMP to deal separately with the sense of pulse and the sense of speed. Rhythmic regularity with long lengths will feel slow but have pulse and rhythmic irregularity and short lengths will feel fast, but with little feeling of pulse. The length factor is also influenced by the input from the human, so IMP will follow the human's tempo, for the most part, although as was mentioned earlier there is a mechanism in place to keep the following from happening too closely.

The frequency decider has a very similar structure to the duration decider. On the heard data side there is a list of the last ten heard frequencies, and each new heard frequency is added to that list. If a `next_now` bang is routed to the heard data side of the frequency decider, a frequency from the list of the ten most recently heard frequencies is selected. This randomly selected frequency is output as the `next_freq`, and a loop is setup that will output the rest of the list as the next frequencies. For example if the initial `next_now` bang causes the heard frequency in index 7 on the list to be chosen, then the next three frequencies sent will be indexes 8, 9, and 10. After the end of the list is reached, the system resets to choose the next frequency from either the heard data side or the generative algorithm side. This loop system causes IMP to play not just one pitch that it has heard from the human, but a series of pitches, and in the same order that they were heard.

The generative algorithm side of the frequency decider is structured similarly to the generative side of the duration decider. There are 42 different files of sets of pitches, and one of those files is loaded into the frequency `coll` (`freq_coll`) based in the melodiousness MPS. The lower numbered sets are major pentatonic scales, and as the numbers go up they cycle through major scales, ascending melodic minor scales, diminished scales, whole tone scales, and finally a chromatic scale. There is a `melody_decider` that changes the choice of frequency `coll` file, according to the variation MPS, during the course of an episode.

Once a frequency is sent out of the current frequency `coll`, a loop is enabled that will select the next 1-5 pitches in a stepwise relationship to the original pitch within the frequency `coll`. The steps may move up or down, or any combination of up and down. This feature gives IMP's output a little more melodic coherence. While it does not eliminate large melodic leaps, it does force at least occasional stepwise motion. Each frequency is sent out on the "`next_freq`" channel, which is received in the synthesis module and stored as the frequency of the carrier oscillator until a `next_now` bang triggers an event and a new frequency is generated.

## 6 Machine Learning and Timbral Interaction

IMP uses The Wekinator[6], which is a real-time machine learning application, to analyze incoming timbral information, and to send appropriate timbral output data to the synthesis module. While IMP is playing, the Wekinator is running as a separate application. IMP's listener module sends loudness, brightness, and noisiness data to the Wekinator via OSC. The Wekinator runs these three streams of data through a neural network that outputs a single value between 0 and 127, which is sent back to IMP via OSC where it controls the timbral elements of the synthesis module.

The Wekinator must first be trained by playing tones into the feature extractor (which is part of the listener module), and assigning a value between 0 and 127 to each sound played in. This is usually done with 0 being the most pure tone, and 127 being the noisiest tone. However, if one wanted IMP to respond differently in the timbral domain, one could train the Wekinator differently. When IMP gets a 0 from the Wekinator, IMP plays its most pure tone, and a 127 gives its noisiest tone, with the varying degrees in between. With that knowledge, the Wekinator could be trained for any given input to make pure tones or noisy tones, as long as that input is associated with that value in the training stage.

For most of IMP's testing I used a set of training data comprised of solo trombone and solo saxophone recordings. This was done in hope that one universally useful set of training data could be used for all performers with IMP. That may still be possible, but a much larger sample size will be needed, so individual instrument training sets have been devised which have proved to be more accurate with smaller amounts of training data.

The value returned by the Wekinator is received in the timbral noise module. This incoming value is in constant flux, so the timbral noise module polls that value every 50 ms and keeps a running average of the ten most recent polled values, and this average is what is used to drive the timbral variations in IMP's sound. Using this running average smooths the data flow, creating a more organic, less scattered result.

The value from the Wekinator is tied to the gain on the second order modulation oscillator in the synthesis module. This means that when the human is playing pure tones, the second order modulation is turned off. As the human's sounds get noisier, the second order modulation depth is increased and IMP's tone gets more strident. After a certain threshold, the harmonicity ratio on the first order modulation begins to change to a non-harmonic ratio as well, which can get quite crunchy. This direct relationship between the timbre of the human input and the timbre of IMP is the way I prefer to play with IMP, but it is entirely dependent on how the Wekinator is trained. Different training data can produce very different results.

## 7 Conclusion

This outline of the general structure of IMP shows a system that can be expanded and varied with some ease. Each aspect of the analysis and decision making is

compartmentalized so that existing aspects may be altered without having to change the entire system, and new features can also be plugged in. At this point in its development, IMP is really just out of the proof of concept stage. IMP has been used with success in public performance, but there are still many areas of planned further development. A system for analyzing the amplitude envelopes of events heard from the human, and incorporating that information into the synthesis module's enveloping system will be the next addition, followed by an expansion of the timbral variance capabilities.

## References

1. Open sound control, <http://opensoundcontrol.org/>
2. Albert, J.V.: Interactive Musical Partner: A System for Human/Computer Duo Improvisations. Dissertation, Louisiana State University, Baton Rouge, LA (2013)
3. Assayag, G., Bloch, G., Chemillier, M.: OMAX-OFON (2006)
4. Borgo, D.: Sync or Swarm: Improvising Music in a Complex Age. Continuum International Publishing Group, New York (2005)
5. Chowning, J.: The synthesis of complex audio spectra by means of frequency modulation. In: Roads, C., Strawn, J. (eds.) Foundations of Computer Music, pp. 6 – 29. The MIT Press, Cambridge, MA (1985)
6. Fiebrink, R.: Real-Time Human Interaction with Supervised Learning Algorithms for Music Composition and Performance. Dissertation, Princeton University, Princeton, NJ (2011), [http://www.cs.princeton.edu/~fiebrink/Rebecca\\_Fiebrink/thesis.html](http://www.cs.princeton.edu/~fiebrink/Rebecca_Fiebrink/thesis.html)
7. Jehan, T., Schoner, B.: An audio-driven perceptually meaningful timbre synthesizer. Proceedings of the International Computer Music Conference (2001), [http://web.media.mit.edu/~tristan/Papers/ICMC01\\_PSE.pdf](http://web.media.mit.edu/~tristan/Papers/ICMC01_PSE.pdf)
8. Lewis, G.E.: Too many notes: Computers, complexity and culture in "Voyager". Leonardo Music Journal 10, 33–39 (2000)
9. Pachet, F.: The continuator: Musical interaction with style. Journal of New Music Research 32(3), 333–341 (Sep 2003)
10. Place, T., Allison, J.: Tap tools, <http://74objects.com/taptools/>
11. Risset, J.C., Duyne, S.V.: Real-time performance interaction with a computer-controlled acoustic piano. Computer Music Journal 20(1), 62–75 (Apr 1996), <http://www.jstor.org/stable/3681273>, ArticleType: research-article / Full publication date: Spring, 1996 / Copyright 1996 The MIT Press
12. Teitelbaum, R.: Improvisation, computers, and the unconscious mind. Contemporary Music Review 25, 497–508 (Dec 2006), 5/6

## Appendix: The Software Archive

The IMP software package is archived at <http://research.jeffalbert.com/imp/>, along with links to related publications, and available audio and video of performances with IMP.