

GenJam: An Interactive Genetic Algorithm Jazz Improviser

John A. Biles

jab@it.rit.edu

Information Technology Department
Rochester Institute of Technology
102 Lomb Memorial Drive
Rochester, NY 14623-5608

Popular version of paper 4pMU1

Presented Thursday Afternoon, December 4, 1997

134th ASA Meeting, San Diego, CA

Embargoed until December 4, 1997

Jazz improvisation is an art form requiring spontaneous creativity and subtle communication. When improvising, a jazz soloist uses a personal store of melodic ideas, adapts the ideas to meet rhythmic and harmonic constraints imposed by the rhythm section, and concurrently constructs and performs an improvised solo. Jazz soloists communicate with each other by "trading fours" or "eights," where soloists take turns improvising over a four or eight measure portion of the tune. In these "chase" choruses, one soloist often quotes or develops material played by another soloist, leading to highly interactive, musical conversations.

GenJam, the subject of this article, is a software sideman and featured soloist in the Al Biles Virtual Quintet. It is able to improvise solos in real time on arbitrary tunes, and it can trade fours or eights interactively with a human soloist. This article briefly describes how GenJam works, how it carries on conversations with human soloists, and what it is like for the only human member of the Virtual Quintet to play gigs with it.

How GenJam Works

GenJam, which stands for **Genetic Jammer**, is a computer program that learns to play jazz solos, using a software technology called **genetic algorithms**. This technology applies the principles of biological genetics to search for solutions to artificial or abstract problems. In GenJam's case, the genetic algorithm searches for pleasing melodic ideas with which it constructs jazz improvisations. Our discussion will focus first on how GenJam represents musical ideas, then on how those ideas evolve in a process analogous to what jazz musicians refer to as "woodshedding."

GenJam maintains two hierarchically related populations of melodic ideas, where each **individual** in a population is a chromosome (string of bits) that decodes to a distinct melodic idea or "lick." Individuals in the measure population (the lower-level population in the hierarchy) decode to a measure of eighth-note-length events. Individuals in the phrase population (the higher level in the hierarchy) decode to four pointers to individuals in the measure population. In other words, a phrase individual represents a four-measure phrase made up of four individuals from the measure population.

The melodic ideas represented by the individual chromosomes are abstract melodic contours that are not in any particular key. The specific pitch for a given note event in a measure chromosome is not determined until that note is about to be played during a solo. The note event is then used to look up a pitch in a scale suggested by the chord that the rhythm section is currently playing. In other words, GenJam makes all the chord changes because it plays only theoretically correct notes.

In addition to its chromosome, each individual has a numeric **fitness** value, which is an indication of how good that idea is. The fitness values come from a human **mentor**, who listens to GenJam improvise on a tune and expresses opinions in real time by typing 'g' for good or 'b' for bad whenever so moved. When the mentor types 'g', the fitness for the currently playing idea is incremented; when the mentor types 'b', the fitness for the currently playing idea is decremented. The mentor, then, provides the environment in which individuals either survive and breed, or die off. Genetic algorithms that use a human to provide fitness, rather than a programmed function to compute fitness, are called **interactive genetic algorithms**.

From time to time GenJam breeds new ideas. The better ideas in the current populations (those individuals with higher fitness values) tend to become parents, and their offspring tend to replace the worse ideas in the populations. The offspring are created using crossover and mutation, which operate on the individual chromosomes. **Crossover** involves laying two chromosomes next to each other, picking a random crossover point, and exchanging the bits on one side of the crossover point between the two chromosomes.

Mutation, in most genetic algorithms, is implemented simply by flipping an occasional random bit in the chromosome from 0 to 1 or from 1 to 0. GenJam takes a different approach by using **musically meaningful mutation**. This departure from traditional genetic algorithms causes radical rather than gradual change in the populations, which vastly speeds up the learning process and provides a clean framework for embedding musical intelligence in the search process. Compositional devices used in the mutation operators include transposition, inversion, retrograde, sorting, smoothing, and sequencing. Some mutation operators combat the tendency of genetic algorithms to converge on a few super individuals by insuring genetic diversity in the populations.

After a dozen or so generations of this genetic woodshedding, the populations tend to consist mostly of well-received ideas (high-fitness individuals), and the soloist is ready to play a real gig in public. The Virtual Quintet features about a dozen different soloists, each of which has been trained on tunes of different styles, including up-tempo bebop, bossa nova, ballad, contemporary, hard Latin, waltz, and 5/4 time.

Real-Time Interactivity

The recent addition of a pitch-to-MIDI capability allows GenJam to listen via a microphone to what the human soloist plays and to use that material in its own improvisations, specifically in chase choruses. When GenJam trades fours or eights with a human, it listens to the last four measures of what the human soloist plays, maps what it hears to its chromosome representation, mutates the chromosomes, and plays back the result as its next four or the first four of its next eight.

Pitch-tracking technology is notoriously error-prone, but GenJam is extremely fault-tolerant. Since the target of the pitch-tracking process is GenJam's chromosome representation, pitch or timing errors are not a serious problem because when the phrase is played back, the notes will always be theoretically correct, as described above. Indeed, mistakes made by the pitch tracker (or by the human soloist, for that matter) result in melodic "development" rather than error. This follows in the grand tradition of turning a bug into a feature!

Error-based development notwithstanding, GenJam's fours would be annoyingly predictable if it simply parroted back the most recent four measures it heard when playing its next four. To make it more capricious, GenJam uses its musically meaningful mutation operators to develop the measure and phrase individuals derived from the human's last four so that when it plays its next four, the material is different from but related to what the human just played. The result is a truly interactive improvisation system that is spontaneous and responsive to the human player.

Playing Gigs with GenJam

The Virtual Quintet consists of the author on trumpet and flugelhorn, GenJam on various solo instruments, and a standard rhythm trio consisting of a chord instrument, bass, and drums. Some tunes add guitar, a sustained part, and/or an extra front line horn, but the intent is to maintain the aural illusion of a quintet and not overwhelm the human soloist with elaborate arrangements.

The Quintet currently has an eclectic repertoire of over 100 tunes in a broad range of jazz, Latin and new age styles, and it has played numerous gigs in a variety of venues including concerts, lecture/demonstrations, receptions, and coffee houses. Background settings are usually very successful, but in foreground settings the lack of visual interaction becomes an issue. Audiences like to focus attention on the current soloist, but when that soloist is an inert box, its physical gestures are not exactly expressive. In particular, when GenJam takes an entire solo chorus, the human on stage is in the awkward position of being the visual focus of attention while having nothing to do.

In an attempt to deal with this problem, the quintet has experimented with **audience-mediated performance**. Each member of the audience is given a feedback paddle that is red on one side and green on the other. The audience then collectively trains a new soloist for three to five GenJam-only tunes by using the paddle to register opinions (holding up green for good, red for bad), while the human types 'g' and 'b' in proportion to the relative amount of green and red visible. After the training tunes, the human plays a collaborative tune or two with the audience's new soloist. On these tunes the audience invariably greet the human's first solo with a blizzard of red and green. Technology can be fun!

Contact Information

Al Biles (aka John A. Biles) is in San Diego from Wednesday, 12/3, through Sunday, 12/7. He is attending the ASA meeting only on Thursday, 12/4, when the jazz improvisation presentations and concert are scheduled. He is staying at the Town and Country Hotel, 619-291-7131. For more in-depth information on GenJam, including audio samples from the Virtual Quintet's CD,

visit the GenJam web site at <http://www.it.rit.edu/~jab/GenJam.html>, or email Al Biles at jab@it.rit.edu.