Widmer, head of the Machine Learning and Data Mining Group at the Austrian Research Institute for Artificial Intelligence.

The Creative Processor

With a souped-up reproducing piano and some ingenious learning machines, AI maestro Gerhard Widmer is discovering how performers unlock the art in Mozart. By Pat Blashill

A gray-blue dusk is settling over the Gothic cathedrals, palatial opera houses, and labyrinthine streets of Vienna's First District. Statues of Johann Strauss and cartoons of "Wolfie" Amadeus Mozart peek out from every other public park and store window. Here in the Austrian capital, music is an almost elemental force. You're just as likely to overhear the sound of a young diva practicing operatic vocal scales as the thump of a sleek techno-innovator like Patrick Pulsinger playing in some ancient cellar just out of sight. It's a place where very old and very new musical traditions collide and intermingle – the perfect setting for a computer scientist obsessed with examining the blips and fault lines, deviations and inventions, that transform music into something more than code and just slightly less than magic.

PHOTOGRAPH: THEODOR BARTH

It's Friday evening, and most of his fellow teachers at the University of Vienna have already gone home, but associate professor Gerhard Widmer is bounding up the stairs of a peach-colored Baroque building and into the offices of the Austrian Research Institute for Artificial Intelligence. He waves hello to his team scientists - Simon Dixon, Emilios Cambouropoulos, and Werner Goebl - and makes for a computer monitor marked EUROPA, which is jacked into an electric piano. Widmer eagerly begins to trigger several audio files. The first is a computer-generated performance of Mozart's Piano Sonata KV 331 - a stiff, note-for-note execution. Next, he cues up a version of KV 331 by Viennese pianist Roland Batik, whose playing is much more contemplative and fine. Last, Widmer plays another computer performance of the same work, this time by a "learner" machine that has taught itself a couple of rudimentary rules about Batik's technique. This rendition is delicately finessed and preternaturally beautiful. It sounds almost identical to Batik's. If anything, the computer-generated version comes off a little, well, sadder.

Widmer is rocking back and forth in his hemp-colored Nikes, beside himself with excitement. "It's using just two rules," he exclaims, "and it sounds just like Batik!"

Tonight, Widmer is a machine-learning specialist who has it all: a million-dollar research grant, a startling new set of experimental data, and a computer that can play a Mozart sonata like a Viennese master.

He's a classically trained pianist who took a left turn on the way to the conservatory and ended up in artificial intelligence. Now, backed by a prestigious grant from the Austrian government, 40-year-old Gerhard Widmer is heading up a six-year investigation into the expressive aspects of live musical performance. His goal, simply put, is to quantify the elusive, often rapturously mythologized, sound of music.

Widmer may be the most ingenious AI researcher investigating the phenomenon of creativity. Ever since John McCarthy coined the term "artificial intelligence" for a 1956 Dartmouth computer science conference, the study of the mental aerobics necessary for artistic problem solving has held an overwhelmingly seductive allure for scientists. Most researchers have concentrated on developing machines that create, and ended up inventing elaborate parlor tricks like bagpipe-tooting robots, or Ray Kurzweil's Ramona, a femme-bot metal queen who didn't exactly wow 'em at the January 2001 TED conference.

Widmer has built upon the pioneering work of David Cope - a professor of music at UC Santa Cruz who created the Emmy program, which composes original music in the style of Bach, Beethoven, and others - and UC San Diego's Harold Cohen, whose painting program, Aaron, produces canvases that suggest a mellow George Grosz. However, Widmer has flipped the focus and begun to use his computer learners not to produce whiz-bang new music, but to scrutinize the act of human musical performance in such minute detail that the creative process is illuminated as a vast universe of small steps and inventive choices. Ultimately, his work may result in a program that plays piano by approaching music as an incredibly complex system instead of as an inspired act of genius.

Widmer's "learners" examine exactly what happens between a pianist and a keyboard, seeing music as a complex system instead of an act of genius.

The traditional view of a live classical performance suggests that it's the artfulness of the musician that makes a piece come alive, transforming the score into real music, and, some would say, music into poetry. Humans are taught their craft, often verbally, sometimes by example, and even by piano teachers who explain melody simply by singing to their students. Whatever the method, warm-blooded players learn very quickly that reading, translating, and interpreting the score are all part of the game.

Yet musicologists have only recently begun looking at *exactly* what happens when performers do that thing they do. "Classic musicology has spanned hundreds of years studying compositions as if the score itself were the music," Widmer says. "There are tons of papers about the structure of Mozart sonatas or Beethoven symphonies. But the score itself is a dead piece of paper and specifies not nearly enough detail for making a performance. It leaves so many things implicit."

It's easy enough, Widmer explains, to write a computer program that translates a score into a MIDI file, the standard format for interconnecting musical instruments and computers. "But when you play it," he says, "you get a mechanical performance that doesn't sound like music at all. Something has been neglected - the essential part of the story of what the performer does to the score."

By using computers and an amazing reproducing piano called the Bösendorfer 290 SE, he intends to find out exactly what happens between the fingers of a piano player and a keyboard, and how this interaction of flesh and ivory becomes art. He subjects the chain of sounds and symbols that is music to the same glare of examination that others have turned toward DNA. Widmer's even looking for the same kind of answers: If genetic decoding is about discovering what we are physically, deconstructing and quantifying the act of making music may tell us more about what we are psychically.

He uses state-of-the-art software like the inductive logic program FOIL, and conventional machine-learning methodology, which usually involves "teaching" computers by giving them real-world information - anything from Impressionist paintings to Wall Street trends - and letting them discover the patterns and rules that may have created these things. Widmer gives a learner the written score of a Mozart piano sonata, then feeds the same machine a recording of a human performance of that sonata. The learner compares the two sets of information, scanning for deviations between how a piece is supposed to be played versus how it actually is played. Widmer asks the computer to deduce a set of rules that might explain these differences. He then feeds it the written score for an unfamiliar sonata and asks it to "play" the piece using its newly identified rules. If the result sounds

A Bösendorfer tech works on a 290 SE; the reproducing piano has sensors to send and receive data.

like music, the computer has learned something, and Widmer may have unearthed one of the thousands of previously unteachable rules that underpin live musical performance.

Widmer isn't the first AI researcher to plunge into the area of musical intelligence, but he may be the most adventurous, and his experiments could have repercussions across a spectrum of arts and sciences. Thomas Mathiesen, distinguished professor of music at Indiana University, believes that Widmer's work is part of a growing interest in performance among musicologists. "There have been lots of attempts to get at what it is that gives music life beyond the mere technical notation on the page," he explains. "[Widmer's work] is certainly an interesting approach. It seems to me that, methodologically, it makes perfect sense. It's going to reveal a great deal about the way one intelligent performer handles musical events that occur routinely in Mozart sonatas – how cadences are approached, how lines are shaped – and if you had enough of these studies, you could reverse the process and say, 'Let's look at all of these different interpretations, and see if we can find certain commonalities among them that may reveal something about the score or the underlying concept of the piece.'"

David Cope, who has continued to evolve Emmy, thinks of Widmer's work as complementary to his own. "We have no idea what music really is – we haven't a clue," he says. "We know it's important and precious, but that's it. The question is, how close can you get toward understanding music? I think Gerhard will go really far."

Bruce Buchanan, president of the American Association for Artificial Intelligence and professor of computer science at the University of Pittsburgh, agrees. "Widmer is taking a human enterprise – the performance of music, something that all of us would call creative – and he's showing us how much of that is understandable, by the same kinds of rules that we use for a lot of other activities.

"Some of the heuristics that Cope or Cohen or any of these people use are things you and I can use when we're trying to find creative solutions to other problems in other areas," Buchanan continues. "In Widmer's work, [he's found] that you're not constrained to playing quarter notes exactly as one-quarter of the measure, and being able to be liberal is *part* of what we mean by creative expression. It's true in the visual arts, too: When you put paint on canvas, if you draw a line and you slavishly keep all the paint within the line, that's one thing. But if you let it *flow* a little bit, it can take on a much nicer appearance."

Widmer's project relies on a corpus of raw musical data: MIDI files containing all 106,000 notes of the performances of 13 Mozart piano sonatas by Roland Batik, who is internationally acclaimed for his interpretations of Wolfgang Amadeus. Batik favors the pianos of Bösendorfer, a Viennese workshop whose concert instruments are nearly unrivaled; the pianomaker's name is so beloved in the Austrian capital, and so enmeshed in national history, that one of the lanes of the city's regal First District is called Bösendorferstrasse. Batik played the performances Widmer is using on a 290 SE, an instrument with a fascinating history of its own.

The 290 SE is a modified model of Bösendorfer's finest piano, the Imperial, the only concert grand still made with the nine extra sub-bass keys that enable a musician to play certain compositions by Bartok, Debussy, and Busoni. SE stands for Stahnke Electronics, which was founded by Wayne Stahnke, an American musical engineer and former NASA Jet Propulsion Laboratory contractor. In the early '80s, Stahnke – a musician who had long been fascinated with player pianos and their better-made, more musically sophisticated brethren, reproducing pianos – finally perfected his own instrument. Working for a private client who had hired him to build an exquisite reproducing piano, Stahnke souped up a Bösendorfer 290 with an array of highly sensitive electronic measuring devices, including infrared shutters, which he placed beneath the piano's keys and above its hammers. Augmented with micro-

PHOTOGRAPH: THEODOR BARTH

processors and a special program to record performance data for translation into MIDI files, the 290 SE remains unparalleled in the history of reproducing pianos. It can record and play back a human performance with such precision that, according to Bösendorfer technicians, musicians as diverse as Herbie Hancock, Dave Brubeck, and French classical master Philippe Entremont have all recorded on it out of curiosity and have come away amazed.

Stahnke licensed the SE system to Bösendorfer in 1984, and the company made 33 of the SEs commercially available before discontinuing their manufacture because of disappointing sales. Today, a 290 SE trades hands between owners to the tune of about $90,000. Wayne Stahnke no longer builds pianos - he now writes microcode for MMC, a router chip company that's a division of San Diego, California-based Applied Micro Circuits. But he remains convinced that the 290 SE offers a unique opportunity for both recording and analyzing music.

"You don't want to talk about the mechanical aspects of piano playing with some pianists, though, because they'll insist that they're putting something human into their playing," he says. "If there's anyone who believes in the mystery and power of music, it's me. But I'm an engineer as well. It's like heart surgeons - they operate while thinking of the heart as a pump. That doesn't mean they don't also think of the heart as the wellspring of our humanity, but they're just trying to make this pump work better. If you focus on the mystery in the playing of music and ignore the mechanical aspects, you do so at your own peril."

In 1998, while Widmer was teaching at the University of Vienna, he learned that Roland Batik had recorded not 1 but 13 Mozart sonatas on the 290 SE over the course of several visits to the Bösendorfer headquarters. The scientist realized such a collection of performances of one composer's work by one pianist represented a perfect set of control data for a machine-learning project. He got permission from Batik's representatives to use the recordings, and then set to work with a three-man team on the project he titled Artificial Intelligence Models of Musical Expression (AIMME). Widmer has yet to meet Batik: "He's very busy," the scientist says - which probably says more about Widmer's respect for the great musician's time than it does about the pianist's schedule.

Widmer planned to have his computers compare the sonata scores to Batik's renditions of them, but he immediately ran into a manpower problem: Batik's Bösendorfer recordings, translated into MIDI files, are computer readable. But translating the Mozart scores, which would have to be done by hand - typing code to precisely indicate the correct name, position, and properties of 106,000 notes - would have taken upward of a year. To save time, Widmer team member Emilios Cambouropoulos developed a program that worked backward through the Batik performances and extracted huge files of score data, using the record of how the pianist interpreted the music to guess how the compositions may have been originally notated. Widmer says he then "hired a bunch of students, for a lot of money and quite some time, to go through all the files - Mozart score in hand - and check every single note." The result: machine-readable Mozart.

The next step was to match up the information in the extracted score file (the record of how each sonata was notated and meant to be played) with the MIDI file translations of Batik's Bösendorfer performances (examples of how one pianist actually did play the piece) by a process Widmer calls score-performance matching. This meant another massive number-crunching task for the computers, and "eight or nine weekends" of meticulous hand-correction by the scientist himself.

"The computer has to find out which note in the notated score corresponds to which note in the performance," explains Widmer. "Technically speaking, this is not trivial, because there are lots of notes in the performance that shouldn't be in there, and the piano records it all. The pianist may have accidentally played or dropped a note, and there are also lots of ornaments in Mozart, like trills, which are just notated in the score as *tr* but add a lot of notes to the performance."

After score-performance matching each of the 13 sonatas, Widmer's machines could chart the broad patterns in Batik's playing - his deviations in tempo, volume, and articulation, all of which Widmer refers to as expression data - and graph them as performance curves. These curves, actually jagged lines across a grid, are like EKG maps of Batik's style, and they give Widmer a look at the pianist's technique, at the methods he uses to shape his performance.

These include the smallest, slightest - perhaps even unconscious or intuitive - facets of Batik's performance. Widmer's examination of these elements is the foundation of his current phase of research. He is breaking down each sonata into individual notes and short passages that may exemplify any number of properties, such as ritardando (a slowing down of the tempo) or crescendo (increased volume). Then he feeds the computer many thousands of examples of, say, ritardando passages as played by Batik, asking that it compare them to their notation in the actual score and to one another, as well as search for patterns in the pianist's playing. The learner must consider the context of each note; for example, what sounds came before and after it, and where the note falls in relation to the beginning and end of the sonata. The computer then makes a series of educated guesses that might explain why the pianist played each note as he did. After a few hours of computation, the learner delivers a long list of these guesses, which Widmer refers to as a table of rules. Put simply, these rules project that a given note with certain qualities will be played in a specific way. Each rule includes a ratio that expresses the number of examples of this kind of note being played in a particular fashion.

The immensity and complexity of this phase can't be overstated: Since any of the performance data's 106,000 notes may be

The sonatas' 106,000 notes are crunched first into performance curves, like EKG maps of piano technique, and finally a complex table of rules.

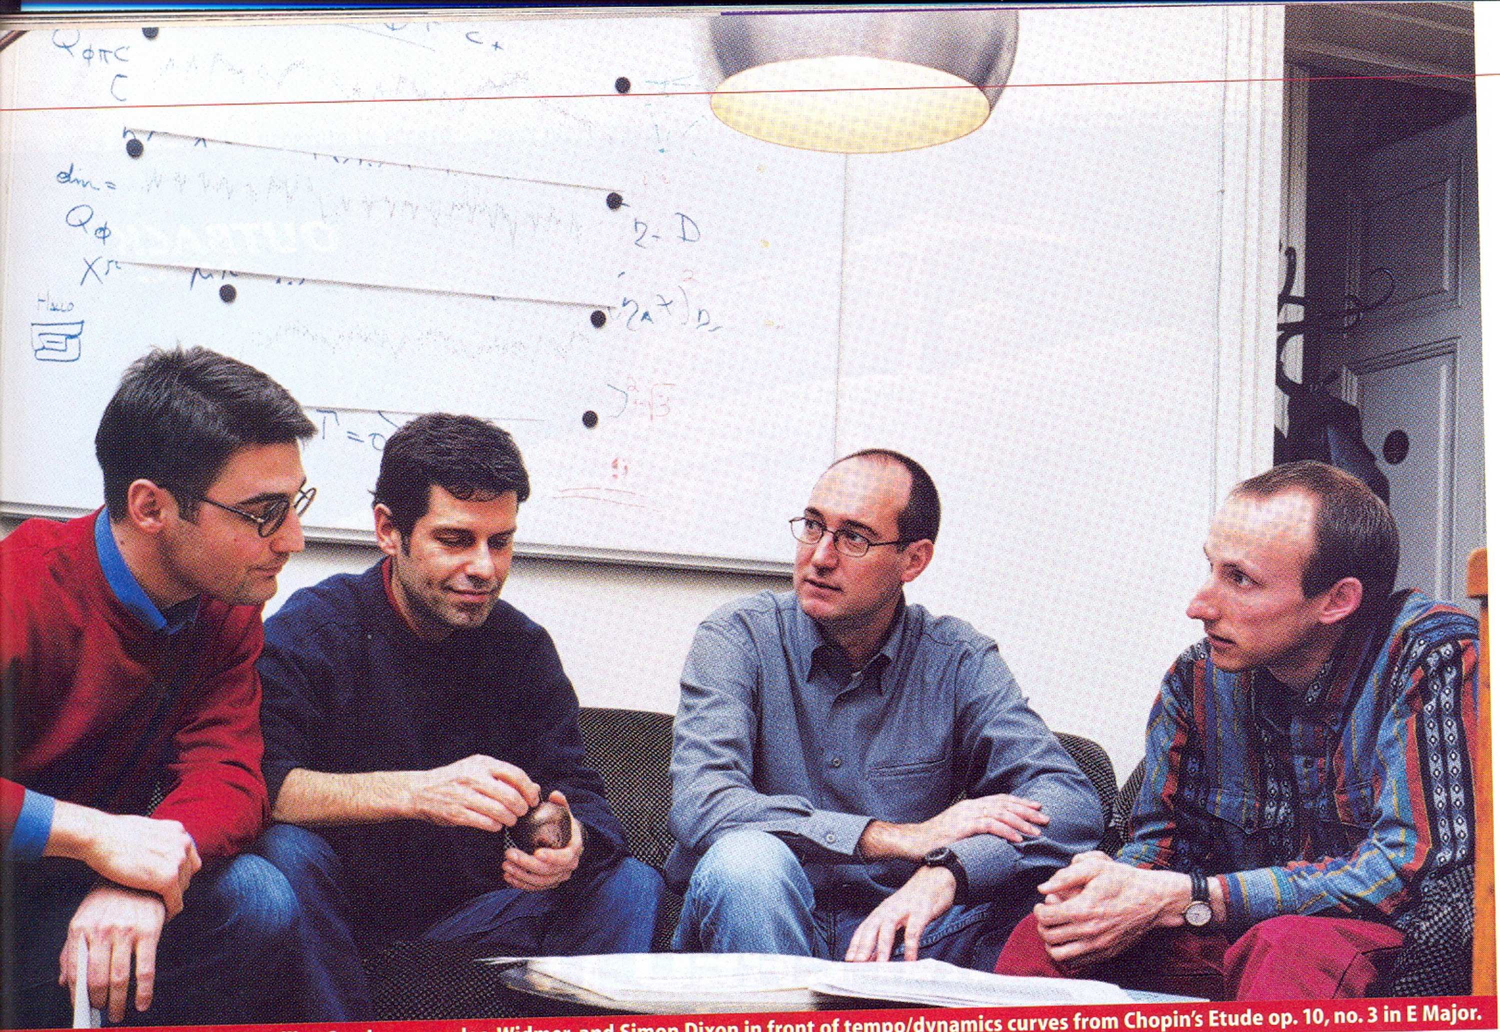

From left: Werner Goebl, Emilios Cambouropoulos, Widmer, and Simon Dixon in front of tempo/dynamics curves from Chopin's Etude op. 10, no. 3 in E Major.

examples of musical techniques that make up expressive performance, Widmer must repeatedly group and regroup thousands of single-note examples by different classifications. He might, for instance, give the computer thousands of examples of decrescendo notes played in three-quarter time. The computer returns an array of rules that may explain that aspect of the pianist's performance. Widmer repeats this process over and over again.

One night when I visit the researcher, he's examining a 10-page sheath of ritardando-note rules. It's a printout of lines that reads, in part:

```
R8: rit (8/0- 0.90/0.98/0.94/0.03) <=
[metr_strength:lte(2), dir_prev: up,
int_prev: lte (4), tempo: gt (128),
global_scaledeg: fourth}.
```

In this case, the third figure, *8/0*, is an expression of the ratio of times this particular rule predicts Batik's performance; here, one kind of ritardando note was played the same way eight different times, and there are no counterexamples of the note being played another way. The ratio is an indicator of the rule's robustness – that is, of whether it's a reliable predictor of expressive performance. In this instance, Widmer's learner seems to have discovered a robust rule about Batik's playing.

To be sure of that, Widmer spends a good deal of his time nowadays going over these rules by hand, using his own knowledge as a pianist as well as his familiarity with decades of musicology research to select those that seem interesting, surprising, or indicative of a previously known expressive performance phenomenon. The computer ends up discarding rules that are statistically weak; for example, those that merely highlight meaningless, random coincidences in the pianist's playing.

The only way to test his research and confirm he's discovered an underpinning of Batik's performance is to have one of the learners apply the rules to play a different Mozart sonata. After Widmer pushes his system equivalent of Play, the learner produces a performance curve that represents how it would execute the piece. If he wants to, Widmer can translate the new data into a MIDI file, hook it up to an audio playback system, and listen to it "live." If the computer has miscalculated and stumbled on rules that are inessential or misleading, the performance will be mechanical and awkward. But if Widmer and his system have indeed discovered valid rules about how to play Mozart, the learner's music will be nuanced, delicate, and perhaps even touching.

Gerhard Widmer didn't grow up in Vienna but in Bildstein, a tiny town in the Austrian alpine region of Vorarlberg. He was one of the first teens from that part of the country to leave home and go to gymnasium, the high school system for advanced learners. Widmer loved to play piano, but he didn't love piano classes, and his interest in taking things apart led him away from playing music for a living to studying it using computers. While doing undergraduate work at the University of Vienna and two years of computer-science study at the University of Wisconsin-Madison (where he developed a superb, Midwestern-inflected command of English), Widmer grew more interested in empirical research than in AI

PHOTOGRAPH: THEODOR BARTH

theory, and was drawn to the rigors of machine learning.

When Widmer returned to Vienna, Robert Trappl, one of his professors, invited him to pursue postgraduate work at the Austrian Research Institute for Artificial Intelligence. The institute is a think tank supported by the Austrian Federal Ministry of Education, Science, and Culture, and works closely with the University of Vienna's Department of Medical Cybernetics and Artificial Intelligence. Founded by Trappl in 1984, the institute is headquarters to nearly 30 scientists who, in addition to studying the interaction between artificial intelligence and society, are doing basic and applied research into natural-language processing, intelligent software agents, and new media.

By the mid-'90s, Widmer had chosen the perfect body of test data for his machine-learning experiments – from his first love, classical music. He began to focus on the expressive aspects of music in particular, and, in 1996, delivered a paper titled "What Is It That Makes It a Horowitz?" at the 12th European Conference on Artificial Intelligence. In 1998, Trappl urged Widmer to submit a proposal for a START grant, a research prize funded by the Austrian government, judged by an international jury of scientists from various disciplines, and awarded to a small number of researchers working across hard sciences. According to Trappl, Widmer procrastinated until the weekend before the project submission deadline, then worked almost 72 hours straight to put together a brilliant proposal and submit it just under the wire. A few months later, Widmer was named a recipient, along with a mathematician, a geophysicist, and a medical pathologist. "Using the method of machine learning, [Widmer is] looking for patterns and general principles that describe and explain artistic interpretation," announced the Austrian Science Fund, the government organization behind START, in its statement. "For musicology, this would be helpful in testing theories of musical expression, as well as in possibly discovering new, as yet unexamined principles.... This project will research the rational aspects of musical expression, especially those complex relationships between musical structure and effect which make music and its interpretation into an intellectual achievement."

In the world of academic research, especially in the Austrian corner of it, a machine-learning specialist who wins a grant for 12 million schillings (almost US$1 million) tends to stand out. Widmer has gotten all sorts of attention since beginning the AIMME project, some of it flattering, some of it a little hysterical. The Austrian magazine *Profil* touted him as a likely future recipient of the Nobel Prize. Unfortunately, the Nobel isn't awarded for computer science, or even engineering. And when Widmer was interviewed on a local radio show, one angry listener called up and accused him of wasting the government's money by examining something that's unknowable.

"One typical reaction I get from people is that they don't believe there is anything that can be studied," Widmer says with a smile. "They think music performance is something the artist just creates out of intuition or inspiration. They think expressive performance is something that cannot be measured or quantified. They say, 'Why on earth are you doing this?' I can answer this question only if you believe in the value of basic research. If you think we don't need musicology at all, or literature or history or all of that, then you will not believe this is important. But if you do believe it's important and interesting to find out about things that are characteristics of human beings, like arts and music, then you will agree this is a worthwhile field."

One pianist playing one composer is only a tiny step – and the "rules" for playing Mozart change. Now Widmer might try teaching his learners jazz.

The University of Vienna's Department of Medical Cybernetics and Artificial Intelligence occupies part of a floor in what was once the 17th-century Schottenstift monastery, and this is where Widmer spends most of every day. He shares a roomy office with one other professor, and his end of the room is a sprawl of books and papers, computer equipment, compact and Zip disks, stereo speakers, and – to one side – a Korg M1 electronic keyboard covered with a bath towel. Perhaps this is what Fatboy Slim's workstation would look like if he won a tenure track gig at Oxford.

It is here that Widmer will test and retest computer performances of the Mozart sonatas for most of the rest of this yearlong phase of the AIMME project. Afterward, he hopes to expand the project to an examination of Chopin performances by the French pianist Entremont, and to other classical performers. This could take another year. Then, well, who knows? If Widmer has any START funding left, he might begin teaching his learners jazz. Investigating the rules that underlie this highly improvisational music would present him with a fresh and awesome set of challenges, but the scientist is undaunted, and even enthralled, by the prospect. "If I had funding for 50 years," he says, "I could fill 50 years with activity, no problem."

Scrutinizing one pianist's performances of Mozart is only a tiny step. As most musicologists will point out, the "rules" for playing Mozart, if anything, have changed over the years, in response to both the kind of pianos available and the changing perceptions of what an audience wants to hear. And Widmer is quick to concede that a computer trained to play Mozart would have to learn a great deal more to be able to play the very structurally different works of Chopin, or Rachmaninoff – to say nothing of a Thelonious Monk tune or Widmer's favorite pop song, the Beatles' "Yesterday." But even if the AIMME computers learn nothing more than how to play Mozart sonatas like a 20th-century Austrian pianist, they will have done so because Widmer is certain there's a method to the magic within all music, from KV 331 to Lennon and McCartney.

"Creativity is not unrestrained," he asserts. "It's not just about making random choices or doing unpredictable things. When musicians break particular rules, they do it on purpose, and they know exactly why they do it and what effect they want to have. I've got data about Vladimir

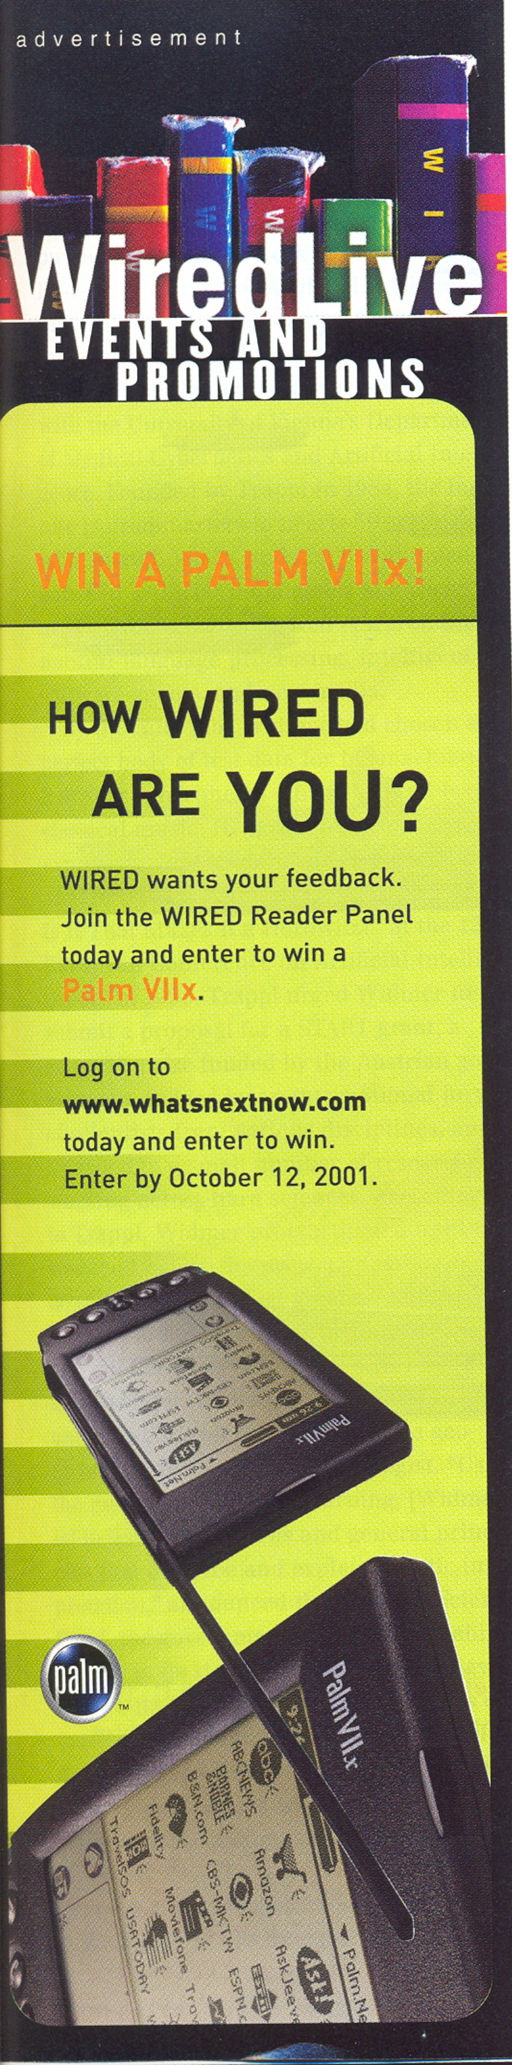

Horowitz's performances of the same piece over a period of 20 years, three different recordings, and his tempo deviations are incredibly consistent, as though he had an internal clock."

Widmer doesn't believe he'll be able to teach his learners how to play as masterfully as Horowitz anytime soon. For starters, he would need Bösendorfer SE performances by the late pianist (there are none). The AIMME team hasn't yet come up with a way to extract that sort of very precise performance data from CDs or other recordings – although they are working on the problem. But Widmer isn't interested in creating a Horowitz player program; instead, he's grappling with problems that lie at the heart of AI research.

In 1999, AI pioneer John McCarthy told an interviewer, "We understand some of the mechanisms of intelligence and not others . . . and AI research has discovered how to make computers carry out some of them and not others. If doing a task requires only mechanisms that are well understood today, computer programs can give very impressive performances on these tasks." But one of the great stumbling blocks of AI research into creativity has been finding a way to around McCarthy's dilemma: Even if you don't know how people do it, a machine can figure out how to do it. And then, in looking at how the machine did it, you gain insight into what people might actually be doing. It's a two-way arrow."

"We're using the computer for discovery," Widmer says. "We'll never get to the point of explaining the refined aspects of a Horowitz performance. But maybe we can get as far as explaining what is common between performers like Horowitz and others – the basic rules that are deeply rooted in the nature of music. My interest is in using the computer to get insight into complex phenomena."

At the end of another long day at the institute, Widmer hovers restlessly around his team as they finish up. He paces for a bit, then moves over to the piano to play a little night music. He knocks out loose, playful renditions of "Ain't Misbehavin'," "Somewhere Over the Rainbow," and the von Trapp family theme song, "My Favorite Things."

Then, out of boredom perhaps, Widmer plays the opening adagio of Mozart's KV 331. It's not a perfect performance, neither as

Widmer's breakthrough: finding a way to display expressiveness without explaining it – since we don't entirely understand the process ourselves.

explain art and music to machines, even though we don't entirely understand these creative processes ourselves.

The American Association for Artificial Intelligence's Bruce Buchanan believes that one of Widmer's most exciting breakthroughs is that he has found a way to display and describe expressiveness in music to his learners without actually *explaining* it to them. "He has created a language of musical expression before he learns the rules in that language," says Buchanan. "Within that vocabulary, he now has a structured way of talking about expressiveness in these dimensions. Given this vocabulary, he doesn't know how to tell the learner what to do; he's just telling it in this terminology: 'Figure out what to do.' That's really powerful. He's giving us a way controlled nor precise as the computer learner's rendition that Widmer was so excited about – but the scientist gives the piece something else, a sort of casual melancholy. His playing fills the room with the quietly bummed-out mood one sometimes feels in airports or laundromats. Then, without a thought, Widmer stops and the mood is gone, evanescing from the room like cigarette smoke.

If he wants to teach a computer to make such a gently wistful sound, Gerhard Widmer may indeed have a long road ahead. But he seems like someone who will enjoy the ride. ■ ■ ■

Pat Blashill (pblash@rcn.com) *writes about music for* Rolling Stone, Spin, *and other publications.*