
CALIFORNIA INSTITUTE OF TECHNOLOGY
1891
CALTECH

The Caltech Electronic Nose Project

Can we teach a computer to "smell" in the same way that we can teach it to "see"?

by Nathan S. Lewis

Although not blessed with the keenest noses in the animal kingdom, humans (in this case, the author's three-year-old son Jeffrey) can smell the difference between yum and yuck almost from birth. Photo and subject courtesy of Dr. Carol Lewis, Jet Propulsion Laboratory.

Of our five senses—sight, smell, taste, hearing, and touch—we understand three well enough to build machines that mimic them. Touch is basically a pressure sensor. There are artificial cochleas—mechanical resonators that transmute sounds into signals that our brain, or a machine, can recognize. And we can build cameras that are essentially electronic eyes. But we know very, very little about the molecular basis of taste and smell, and even less about how to model them. So my lab is trying to build something that will give a value judgment—a number—to a smell, taking design lessons from biology without necessarily mimicking the exact way that a human nose works. We can assign a visual magnitude, a brightness, to a star; can we teach a computer to "smell" in the same way that we can teach it to "see"? This project began as a crazy idea in January of 1993, but there may be something to it.

Smell is a remarkably subtle sense, because most smells are not pure substances, but complex mixtures of different molecules. There are some 700 different chemical vapors in a glass of beer, yet somehow we can take a sniff and say it's beer. The human nose is generalized enough to sense almost all possible molecules, yet discriminating enough to tell the difference between strawberries and raspberries. How can we model that?

The way that most chemists have approached this problem is epitomized by what Arnold Beckman [PhD '28] did when he invented the pH meter. He built a chemical sensor that measures the concentration of one thing (protons in water) very selectively and very sensitively. People have since extended that idea to measure other molecules, such as glucose. In almost every case, the strategy is to design a molecule that has a hole in it—a lock—such that only the right key, i.e., glucose, will fit and generate a signal. (There are, of course, more generalized sensors that measure some physical property of the molecule, but they don't really "recognize" it—they merely tell you that they've detected a molecule with, say, the same charge-to-mass ratio as the molecule you're looking for.) Nature uses the lock-and-key approach very successfully—in enzymes, for example—but it takes evolution millions of years of work to make the molecules fit just right. You can see the daunting task that a chemist would face in trying to build 700 such locks to detect the 700 odor components in a glass of beer. And we'd have to build all 700, because we don't know which components are critical for identifying the smell of beer, and determining whether it smells good or stale. And what would happen when we encountered the 701st molecule in a different odor, like in another brand of beer? We'd have to build another sensor. And we'd have to make each lock specific enough that a very slightly different molecule wouldn't also fit, because even if the other molecule fits poorly we'd still get a signal. Designing such exact locks from scratch is a very, very complex problem at the frontiers of chemistry, and hundreds of groups around the world are working on it.

We abandoned this approach in favor of a pattern-recognition strategy. We decided that the biological olfactory system must employ a set

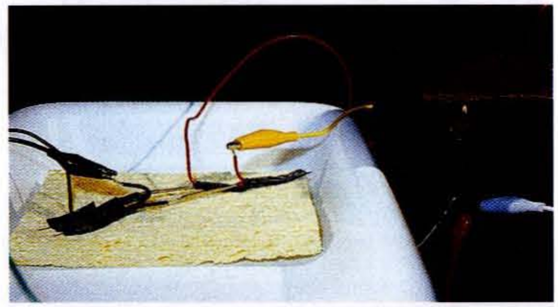

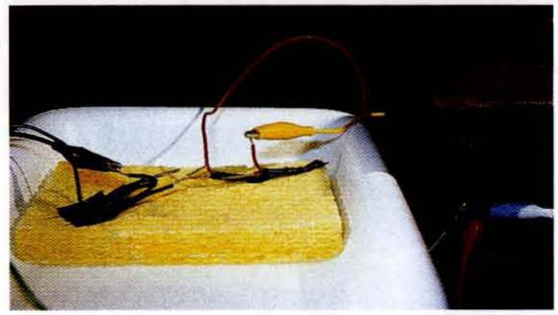

The electronic nose, high-school science-project style. Two electrodes are taped to a sponge. When the sponge is dry (top), the electrodes are in contact, completing the circuit and lighting up the bulb. As the moistened sponge swells (bottom), the electrodes move apart and break the circuit. The light goes out.

Dogs have no reason to sniff out cocaine in the wild, and yet they can be trained to do so in airports. The dogs must be learning to recognize a pattern, because one certainly couldn't train them to develop a new receptor overnight, or even in a few months.

of generalized sensors that respond to everything, but in different ways to different stimuli. Evolution might have developed specific receptors for fruits and wines, for example, but it's unlikely that dogs would have evolved receptors to smell drugs. Dogs have no reason to sniff out cocaine in the wild, and yet they can be trained to do so in airports. The dogs must be learning to recognize a pattern, because one certainly couldn't train them to develop a new receptor overnight, or even in a few months. So the task facing anyone trying to develop an artificial nose is to develop a generalized sensor whose output patterns will announce the difference between the vapors emitted by a rose and a dead fish. Then we train an electronic circuit to recognize those patterns, in the same way that signals fired to our brain get recognized as yum or yuck.

The sensor in our electronic nose must meet several basic requirements. We want it to give us an electrical signal that we can analyze on a chip. We want the signaling event to be reversible—that is, the sensor should return to its initial state when the sniff goes away, so we can use it over and over again. We want it easy to make. We want it to be stable in all sorts of environments, so we can just leave it sitting out in the air. And we want to be able to make it very small, so that we can put a million of them on a little chip.

Our solution is embarrassingly simple. In fact, I'm proud to say that a well-known physicist who wasn't familiar with this project came into my lab recently, looked at our nose and said, "This is a high-school experiment." And I said, "That's exactly right! That's what makes it so wonderful to study, because it works for anyone anywhere." Our sensor is a sponge made of insulating plastic, much like a bathtub sponge, but containing little conducting particles scattered here and there within it. When we pass a current through it, the electrons have to hop from one conductor to the next, so the sponge has a characteristic, measurable resistance. If we were to moisten the sponge, it would swell, and the conducting particles would move farther apart. It would get harder for the electrons to jump between the conductors, and the resistance would go up. Later, as the sponge dried, it would shrink, and the resistance would go back down. (If you soak a sponge, it won't shrink all the way back to its original size when it dries, but if you just add a few droplets, the swelling can be fully reversible.) The same thing happens with vapors—the sponge "sniffs" an odor by absorbing it and swelling up, causing a measurable resistance change, as you can see on the opposite page.

The linchpin of our design is to use an array of sponges with different chemical affinities. Each individual sponge will swell more (and exhibit a higher resistance) when it soaks up something it likes. For example, hydrophobic plastics don't like water at all. If you expose them to a water-like vapor, such as an alcohol, they'll repel it. The sponge won't swell much, and there's not much signal change. But hydrophobic materials *do* like oil, so an oily vapor—benzene, for example—will swell them a lot. So some of our sponges like oil better than water; some like charged molecules more than uncharged molecules, and so on. There's no lock-

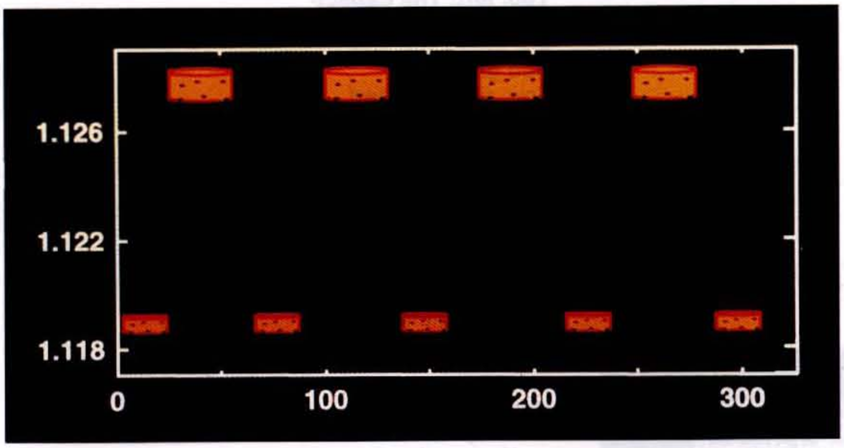

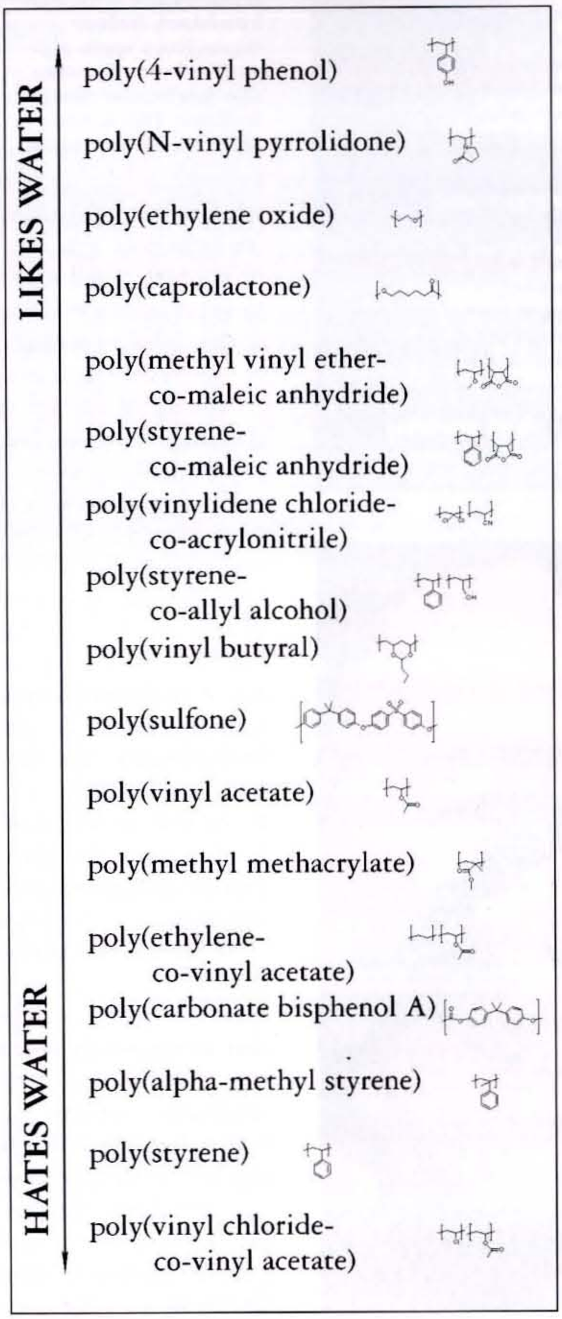

Above: When a sponge [here, poly (ethylene-co-vinyl acetate) containing carbon-black particles] sniffs something it likes (air containing 0.1 percent benzene), it swells and its resistance jumps. When the odor vanishes, the sponge shrinks, so it can be used over and over. The vertical axis is resistance in millions of ohms, and the horizontal axis is time in seconds. Right: The insulating plastics that make up the sponges can be ranked by their properties on many scales, including affinity for water. The chemical structures in brackets are the repeat units that make up the corresponding plastics.

and-key design that says that Sponge A will only respond to the molecule "methanol," and Sponge B will only respond to the molecule "benzene." We don't have to worry about the details—instead, the molecule tells *us* what its important properties are by the signals it generates in the various sponges. We don't actually know if we have enough diversity amongst our sponges yet, so the nose will evolve as we pull out sponges that don't work very well, and put in ones that we hope will work better. We're still trying to figure out how best to choose them.

We were originally going to vary the chemical affinities by modifying the conducting particles. My colleague Bob Grubbs, the Atkins Professor of Chemistry, has discovered ways of making electrically conductive plastics that you can paint on anything. [See *E&S*, Summer 1988.] But then-postdoc Mike Freund, who began this whole project (and is now an assistant professor at Lehigh University), realized that we didn't need to alter the conductor. All we really needed to do was to make one paintable conductor, and then use assorted commercial plastics with various properties, available from any supply house, as the insulators. So that, being simpler, is what we do. We dissolve the insulator, add the ingredients needed to make the conductor, and then apply the resulting solution while the reactions that make the conductor are going on. The solvent eventually evaporates, leaving us with our sniffer sponge.

In hindsight, it turns out that we didn't have to go to all the trouble of making conductive plastics. Any electrical conductor will work, as long as we can find a way to disperse it into the sponge. For example, last summer, SURF [Summer Undergraduate Research Fellowship] student Sara Beaber started using little particles of silver. And Pinocchio, our newest, most improved, nose uses carbon-black particles—the same stuff you find in asphalt and pencil lead. Postdoc Mark Lonergan did most of the work on this, aided by grad student Erik Severin, and Bob Grubbs, as usual, had the idea. Carbon black is a very stable compound, unlike the temperamental conducting polymers, and it's really easy to come by. If you break apart a Radio Shack resistor, you'll find little balls of carbon black inside.

And the way we attach wires to our sponge is incredibly inexpensive—we break apart a 10-cent capacitor. Capacitors store electric charge on thin sheets of palladium–silver foil, separated by a good insulator—a sand-like material called mica—so they don't short out. We use a belt sander to grind the top half of the capacitor down until we expose the foils, and then dip it in our

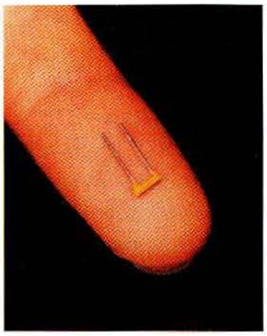

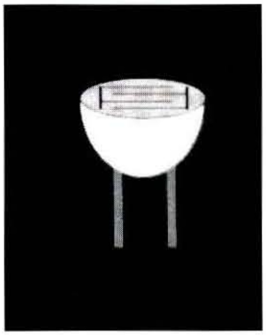

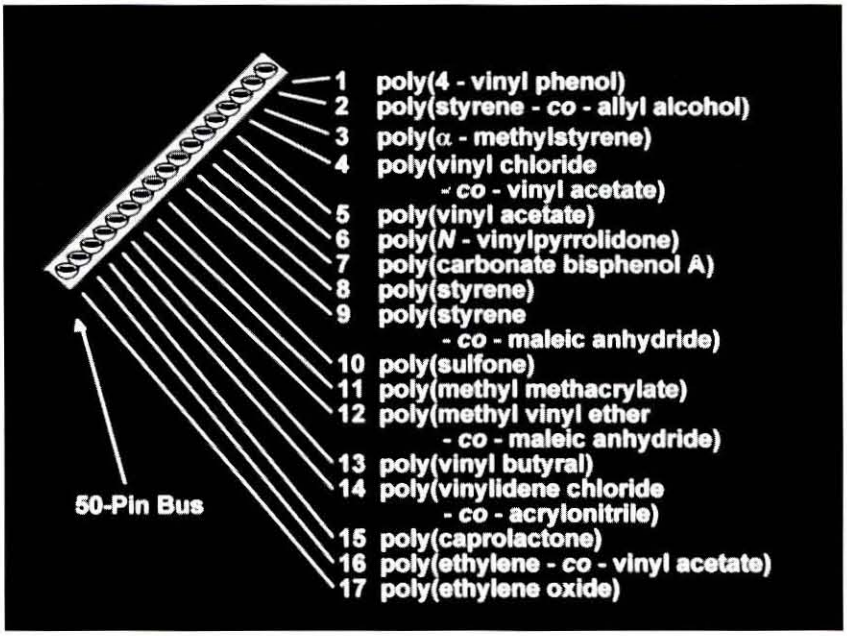

Top, left: The capacitors used in the nose are about the size of rice grains. Top, right: Sealed in epoxy within each capacitor are two sets of interleaved parallel plates, separated by an insulator. Sanding the top off the capacitor exposes a cross section through the plates. Applying a sponge coating to the exposed surface completes the circuit. Middle: Solutions of the 17 plastics listed on the previous page were doped with carbon black before capacitors were dipped in them to make this particular nose. Bottom: The author and two of his noses.

solution, bridging the insulator. Then we plug the capacitors right back into where they came from. A so-called bus chip a few centimeters long can hold a whole array of capacitors, each with a different sponge. The output signals then feed directly into the computer.

The signal's height and shape depend on both the thing being smelled, and the thing that's doing the swelling to sense the smelling. As you saw before, the resistance rises as the sponge swells, plateaus at some value characteristic of the vapor for as long as the vapor remains, and then falls off as the sponge shrinks once the vapor disappears. The swelling and shrinking rates depend on how the sniffer and the sniffed interact. A hydrophobic sponge, for example, will slurp up benzene because it's greasy, and won't let it go easily. But the same sponge won't soak up as much chloroform, and will release it faster. Right now we only look at the maximum signal-height change, but the curve's shape should provide additional, and maybe more valuable, information in the long run.

When we look at the overall pattern of all the signals from all the sponges, we get a fingerprint that—we hope!—will be different for everything that we expose the nose to. (So far, that's been true.) One sponge by itself does not identify a compound—another compound that didn't swell it as well might give a signal that's half as high, but if there were twice as much of that second compound, we might get a very similar signal. But the signals from the entire array provide a pattern that will be diagnostic of a given odor. On the facing page is an example from an array

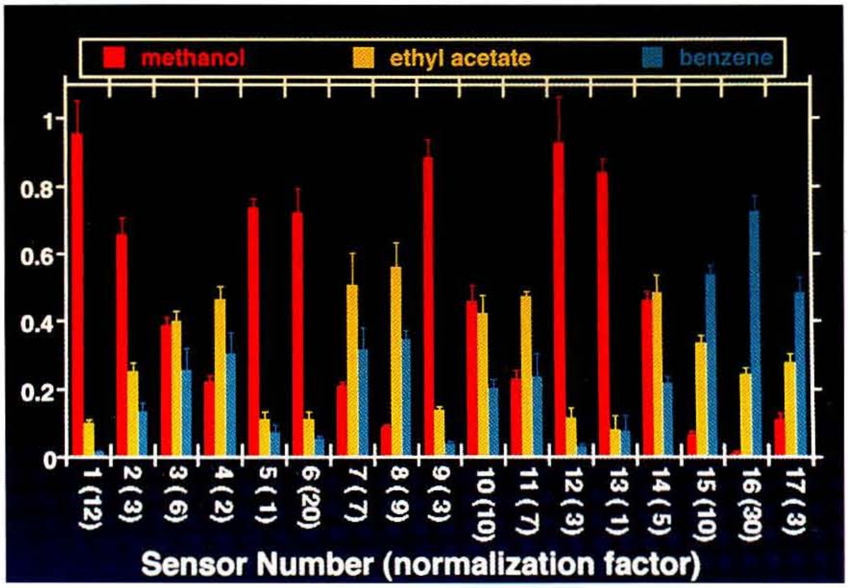

The 17-sensor carbon-black nose gave three different response patterns for three different vapors. The numbers on the vertical axis indicate the relative resistance change in each sensor. Because the response range of each indivdual sensor is different, the values were "normalized" to make them fit on a common scale by dividing them by the number shown in parenthesis below each sensor number.

When we look at the overall pattern of all the signals from all the sponges, we get a fingerprint that —we hope!— will be different for everything that we expose the nose to.

of 17 different sponges. The yellow bars show the pattern that we get for ethyl acetate, a solvent commonly found in paint thinner. The blue bars are the pattern that we get for benzene, and the red bars are methanol. You don't have to have a trained eye to see that they are different, so we can certainly distinguish them electronically.

But it's hard to quantify *how* different they are. You can't tell me if they're 10 percent different, or 20 percent. How can we teach a machine to discriminate between patterns whose differences we can't easily describe ourselves? How do we know how much leeway we can allow between two patterns and still call them a match, for example? We use a statistical method called principal component analysis (which we did not invent) to analyze the data. The method takes all the signals from the individual sensors and plots them as points in what we call odor space, in which it's easier to see the patterns. Unlike ordinary three-dimensional space, however, we have one dimension per sponge. Therefore, even though it's easier to see the patterns, the analytical process can still get quite elaborate.

Last year we did an experiment where we exposed 17 sensors to nine pure vapors—methanol, ethanol, isopropanol, acetone, ethyl acetate, chloroform, hexane, benzene, and toluene. We gave the nose sniffs of the various vapors, repeated in random order, over a period of five days. We didn't control the temperature of the room, and we didn't control the humidity in the air, so this experiment was essentially a worst-case scenario to see how well we could do. The shapes enclosing the data for each compound would have been

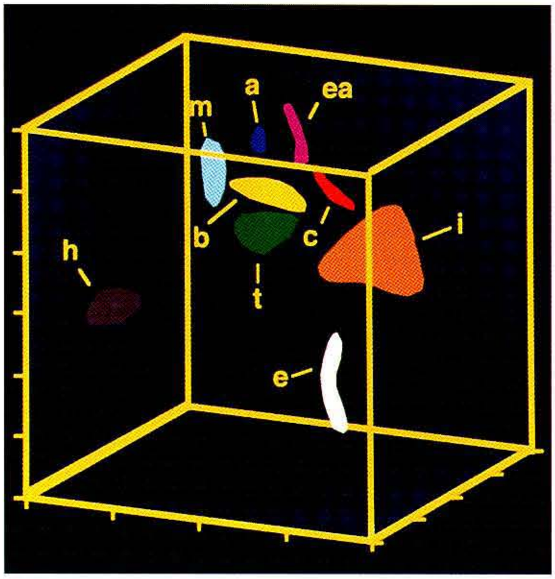

When the nose was exposed to nine different vapors, each one turned up in its own little corner of odor space once the right set of dimensions was plotted. The vapors are labeled as follows: a = acetone, b = benzene, c = chloroform, e = ethanol, ea = ethyl acetate, h = hexane, i = isopropanol, m = methanol, and t = toluene.

smaller in a controlled climate.

I can't plot a 17-dimensional space, so the plot above shows the three dimensions that contain the most differences between those nine patterns. The three alcohols (methanol, ethanol, and isopropanol—methanol has one carbon atom, ethanol two, and isopropanol three) separated very cleanly. Benzene and toluene, which are chemically only very slightly different—toluene has an extra methyl group and so is just a little bit bigger than benzene—are close together, but distinguishable. By contrast, hexane—a molecule about the same size as benzene and toluene, but with a different shape and very different properties—appears quite a distance away. And ethyl acetate and acetone (the solvent in nail-polish remover) are also chemical cousins, but they aren't as closely related to each other as benzene and toluene are, so they show up farther apart than benzene and toluene do. Chloroform, which isn't related to any of these guys, also registers separately.

A computer, of course, isn't limited to "seeing" things in three dimensions, as we are, but can look at all 17 at once.

We can also tell how much of something there is, because the responses grow larger with increasing vapor concentration. All the sponges continue to swell in approximately the same relative way as the odor gets stronger, and we retain the fingerprint.

Each coordinate axis represents some unknown property—it might be how big the molecule is, how it is shaped, how much it likes water, or, usually, some combination of properties. We've already seen how hydrophobicity works, and polarity works much the same way—we can make our sponges hospitable to positive, negative, or neutral charges. We can also discriminate between molecules of different sizes, because the plastics' pores differ in size and shape. Molecules that are too big for the pores don't fit very easily, so the sponges don't swell as much. Molecules that are smaller than the pores do fit, but not very well, and so again the sponges don't swell as much. Discovering what the coordinate axes actually correspond to is a very interesting problem. We're working very hard to try to associate the chemical and physical characteristics of the sniffed molecule with our sniffer data.

Since we have 17 dimensions to choose from, we can select the three that best discriminate between whatever specific compounds we're interested in. If I wanted, for instance, to separate chloroform and toluene, I could plot three other dimensions that would separate chloroform from toluene much better, but wouldn't separate methanol from ethanol as strongly.

A computer, of course, isn't limited to "seeing" things in three dimensions, as we are, but can look at all 17 at once. We had to learn how to analyze such data, so we're collaborating with Rod Goodman, professor of electrical engineering and director of the Center for Neuromorphic Systems Engineering, which is devoted to developing machines that mimic, on some level, the way biological brains—what's known in the trade as "wetware"—work. Rod and grad students Jeff Dickson and Alyssa Apsel are developing a model to handle our data flow based on how our brains might analyze the firing of neurons as we recognize an odor. And last summer, a SURF student of Rod's named Wei Qin set up for us a

Right: A two-dimensional plot of the nose's response to methanol and ethanol mixtures. The line of red squares indicates the response to air–methanol mixtures, and the line of green circles is for air–ethanol mixtures; in each case the deeper the color, the higher the vapor's concentration. The nose was then given whiffs of five mixtures of methanol and ethanol (methanol–ethanol ratios of 11:1, 4:1, 2:1, 1:1, and 1:2) at two different flow rates to give two sets of concentration values. The data fell neatly onto the broken lines. The arrow marked χ_M shows the direction of increasing methanol content, so where a mixture appears on the graph is directly related to its composition.
Below: If all the sensors respond linearly to individual vapors—that is, if the response increases in proportion to the vapor's concentration—the array's response to a mixture of vapors will be the sum of the responses to each vapor as if it were by itself. Here, for example, sensor 1 registers a 5 for vapor A and a 1 for vapor B, so a half-and-half mixture of the two registers as 2.5 + 0.5, or 3.

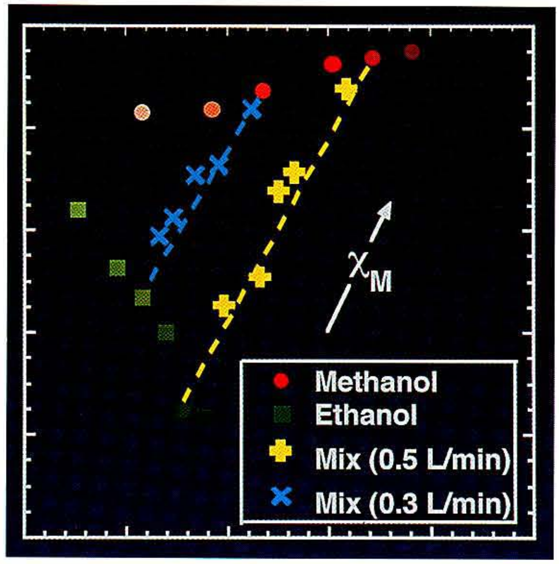

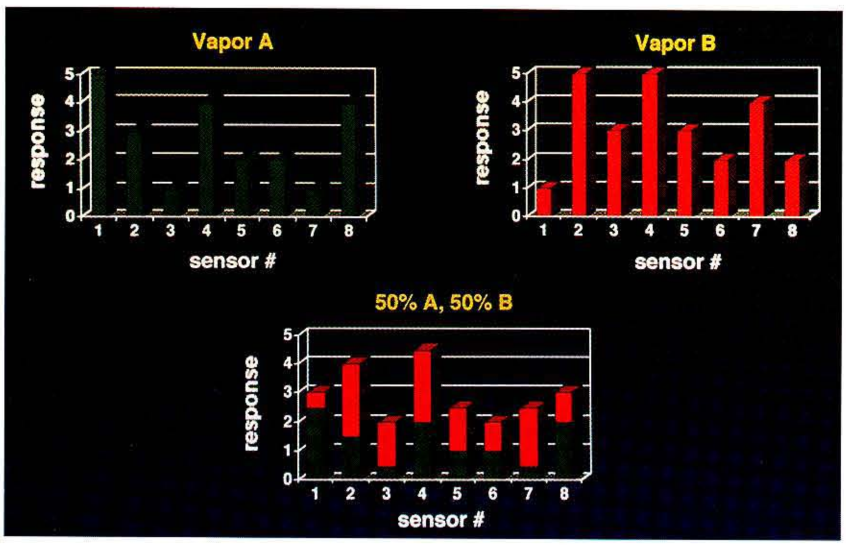

data-processing algorithm called a neural network, which basically mimics a whole bunch of interconnected nerve cells all firing messages back and forth at one another, and which can learn to recognize patterns. [See *E&S*, Summer 1990.] The network took our patterns, processed them, and identified each of our solvents by number. Such a neural network could easily be trained to recognize anything the nose can smell, as long as the sniff gives a reproducible pattern. Wei wrote the algorithm as a piece of software, but neural nets can also be built directly into chips as hardware, and Rod's working on that right now. Brett Doleman, a grad student in my group, is working with Rod's group to figure out how best to classify the different odorants.

Discriminating between pure vapors is a start, but what about mixtures? If we give the nose a mixture that's half methanol and half ethanol, the new pattern should be at the midpoint of the line segment connecting the two pure smells in odor space. Will the nose break this pattern down into the two known ones, or think it's a brand-new smell? It turns out that as long as the responses are linear, the mixture simply registers as the linear combination of the individual smells. If the responses are nonlinear, then we have to train the nose on the mixture as if it were a new compound, which is obviously a lot less useful.

Conversely, can we fool the nose by giving it a new compound? If we don't tell the nose that this is a new thing, will the nose tell us that it's smelling a linear combination of known smells? Or will the nose know that there's something new in the air? We took the data we got from

We want our nose to be able to tell the difference between a rose and a dead fish.

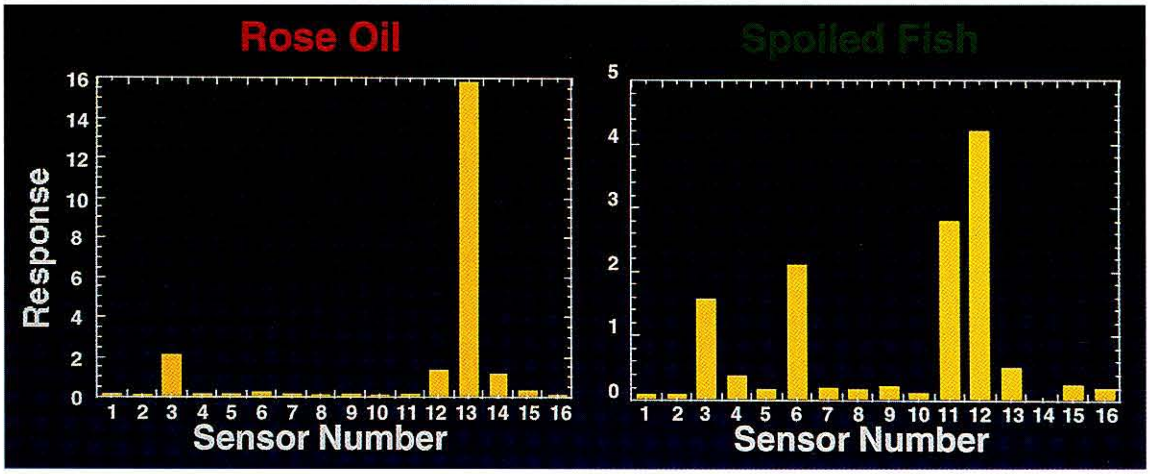

Brett Doleman (right) and Erik Severin (left) give a fish and a rose the once-over.

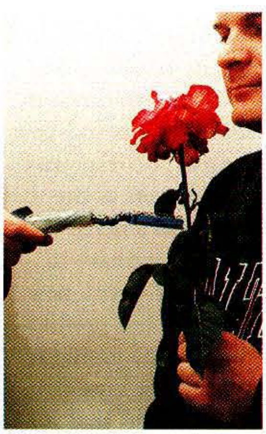

seven smells (methanol, ethanol, isopropanol, acetone, chloroform, hexane, and benzene), and tried to see if some combination of them would reproduce the pattern we got from ethyl acetate. The only stipulation was that all the components had to be positive—we didn't want a recipe that included, say, −15 percent ethanol. And with just that one constraint, we could not make the new smell out of any combination of the other seven smells. Of course, the more you know about the sample, the easier this is; the more different smells you're allowed to use, the harder it gets. There will be a happy medium somewhere, and we don't know what the trade-offs will be; but we do know that in certain instances we can't fool the nose. This is a very powerful test of the electronic nose's information content.

I said at the beginning that we want our nose to be able to tell the difference between a rose and a dead fish. So grad student Erik Severin went to the store and bought one generic fish, and put it in a flask. The human nose has evolved to smell raw meat, so the fish stank to our noses earlier than it did to the electronic nose. People were complaining by noon, but it took the nose all day to pick up the scent. Nevertheless, above right is the pattern Erik got for spoiled fish. (For unspoiled fish, the pattern is just water vapor, which we null out, so there is no pattern. We think this is what the human nose does, because people can't smell water vapor, either. It must be that our nasal sensor cells are in a constant-humidity environment, so they zero out water.) Erik also bought some rose oil, and its pattern (above, left) is quite unlike the fish's.

We can't yet tell red wine from white. We can, however, tell beer from wine from hard liquor by the alcohol content. We actually tried to tell wines apart initially—the Athenaeum is interested in sponsoring this project. In retrospect, perhaps we should have tried nulling out the water vapor with the wines, as we did with the fish.

Our electronic nose can't do what a mass spectrometer does, and say that there is one part per trillion of molecule X in the complex mixture we call "strawberries." But we don't always care about molecule X—sometimes we just want to know that it's strawberries and not raspberries. Sometimes we just want to know, does the cheese smell the same as it did yesterday, or has it rotted? The pattern-recognition approach to smelling does this very, very well.

You can imagine the quality-control applications for such a device. For example, cheese manufacturers pay people to sit on the production line and smell the cheese as it goes by. But they can only smell for two hours at a time, because their noses get saturated. And a quality-control lab can't analyze every single cheese with a gas chromatograph or mass spectrometer. You don't even know what you're looking for, necessarily—sometimes the cheese just smells bad! But a little electronic nose could just sit on the line all the time and say, "The cheese is the same as it was yesterday. The cheese is the same as it was yesterday. It's OK." The nose would beep whenever the cheese smelled different, and then you'd stop the line, and smell the cheese yourself to find out whether it really was OK or not.

Pinocchio, the new supernose, accommodates up to 20 sensors and lives in a stainless steel case on legs (right). The array of glassware in the background holds the pure liquids—the acetone, benzene, and so on—through which air is bubbled to generate the vapors that are then piped to the nose. The three black boxes are computer-controlled flow regulators. Pinocchio's case is so big because the sponges now bridge metal contacts plated onto glass slides, as seen below—an even simpler (and more reproducible) process than buffing down capacitors.

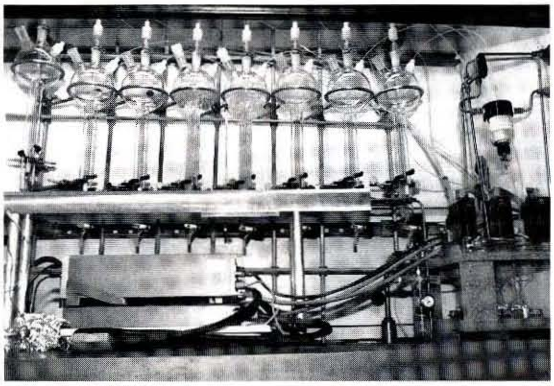

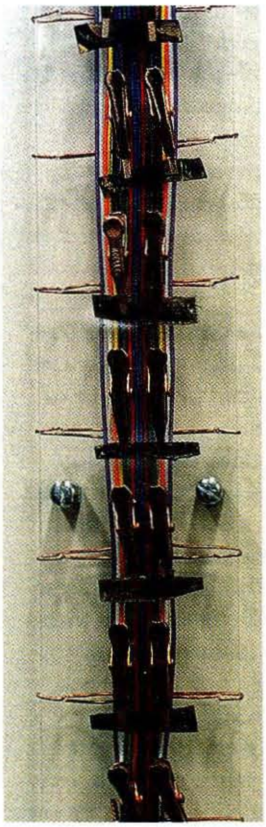

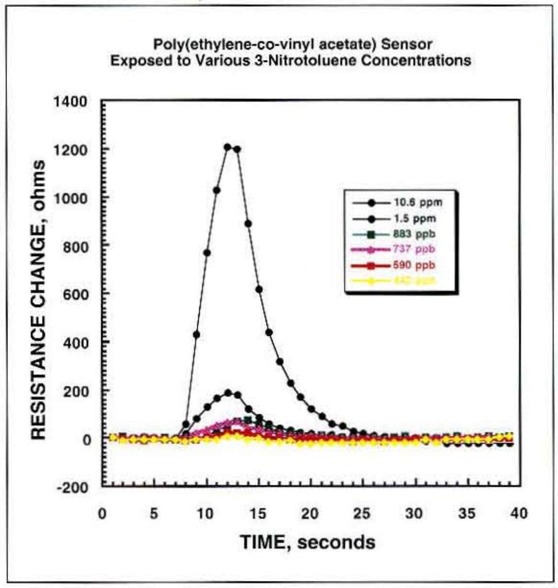

Similarly, you could program the nose to beep when a room smelled differently than normal. In a potentially hazardous situation, you might not even need to know what that difference was. You'd just leave the room (or not enter it, as the case may be), and wait until a more specific sensor had registered hydrogen sulfide from a gas leak, perhaps, and then you'd take appropriate action. NASA is interested in this for the space station, so they're helping sponsor our work. When humans will be up in confined atmospheres for years, in some cases, NASA doesn't necessarily know how to anticipate what might get into the air, and whether or not it will be safe to breathe. This way, they don't have to worry about designing a specific sensor for a substance they don't even know might be up there. The nose would just beep if something new appeared in the environment and the astronauts would reach for their oxygen masks. To this end, we are setting up a gas-handling system so that Pinocchio can try to measure toxic gases. We'd like to find out if Pinocchio can respond to gases that are odorless to us, such as carbon monoxide, but we don't know yet. There may be whole classes of gases that the nose can't smell.

I should also point out that we have no idea about the longevity of these noses. We've only been working on this project intensively for 18 months, so even our first nose isn't *that* old. It's too early to tell if this is really a durable device.

Right now, the nose's sensitivity is limited by our very primitive electronics. We use a simple voltmeter, just like the one you might have in your garage, and we can read what's known as 16-bit resolution. We can detect methanol in air down to 70 parts per million, which is roughly as good as a human nose can do. But we can detect 3-nitrotoluene, which is much less volatile, down to about 600 parts per billion, as shown above. (The less volatile a vapor is, the easier it is to detect at relatively low concentrations, because it prefers to stay liquid and is thus better held by the sponge.) We calculate that the ultimate detection limits will be about 10 parts per billion. Each sensor also needs what's called a Wheatstone bridge, which is adjusted to null out the sensor's baseline resistance. That way, we're measuring a small resistance on top of a zero. Right now, we're measuring a change of a few ohms on top of a 40,000-ohm baseline. We do care about the signal-to-noise ratio, because if there are things in very, very small concentrations that are critical for, say, distinguishing wines, we don't want to lose that information.

We can also adjust the sensor's threshold sensitivity by changing the ratio of conductor to insulator in the sponge. When the conductors are close enough to touch one another, the electrons essentially percolate from conductor to conductor through the points of contact. The electrons travel quite rapidly through the sponge (low resistance), even if they have to go through a tortuous path. On the other hand, if we swell the sponge to a little bit above that percolation threshold, they're going to have to hop across the intervening insulating regions. The resistance will jump dramatically with just that little bit of swelling. It's an on–off signal. We have

A meeting of the noses. Back row, from left: Sara Beaber, Bob Sanner, Nate Lewis, Erik Severin. Front row: Brett Doleman, Mark Lonergan.

actually shown that this works, and you can see how this could be a very sensitive alarm. The alarm wouldn't tell us what's out there, because there's no pattern of linear responses that would allow the chip to recognize what the alarming substance is. But we can set the alarm's sensitivity by adjusting the percolation threshold. The more conductor we have, the more the sponge has to swell before the last percolation pathway is broken and the resistance skyrockets.

What are we going to do next? We'd like to do what we think human and dog noses do—use a large number of incrementally different elements. We'd like to make a million sensors on a chip. We think we know how to do it; it's just a fabrication issue. The electronics aren't the problem—building and adjusting a million different Wheatstone bridges; reading out a million different resistances; making a two-dimensional grid of 1,000 by 1,000 wires; and addressing each of a million individual intersections, even if the sensors are only 10 microns big, is not stretching current chipmaking technology. Such a chip would be a modest size—one centimeter by one centimeter—much smaller than *my* nose! Overlaying the grid of wires would be a matching matrix of little wells—also easily made—to hold the sponges. The issue is, how do we make a million different plastic sponges? Bob Grubbs and I had an idea, which Bob Sanner [PhD '78], now a visiting faculty member from Lawrence Livermore National Laboratory, is trying to implement.

Mike Freund makes an olfactory value judgment.

We start with one monomer—one component of the plastic that makes up the sponge—that might like water, say, and another one that might like oil, and then spray the water-loving monomer left to right and the oil-loving one up and down while smoothly increasing the dilution of each. The wells will fill with an array of sponges gradated by water-loving-ness on one axis and oil-loving-ness on the other. It wouldn't even matter if the gradation varies slightly from chip to chip, because each chip would learn its own response. As long as the response is consistent every time the chip smells that smell, it doesn't matter what the details are.

We don't know yet how much benefit there will be in making a million different intermediate materials instead of just the two extreme cases. We do know that there's no point in doing so if all the responses are linear. Then the intermediate sponges' responses are just linear combinations of the two extremes, and there's no new information. But if the intermediate sponges behave differently, then they give us new signals to the extent that they have different swellabilities. The algorithms for this system are much like those for antenna design, it turns out, although as chemists we don't know enough about our "antennas" to decide just how many we need. That's one question we want to answer: what minimum number of elements is sufficient to distinguish very subtle differences in smells? A million channels is an awful lot of signals—can we get away with fewer? So Brett Doleman is working with Rod's group to figure out how many sensors we actually need.

And if we're building a chip with a million sensors, we could make a composite array in which some sensors beep when something appears in the environment at very low levels, and others wait a little bit and then tell you what that thing is. Or maybe you'd just get out of the room, depending on how many of the low-level sensors beep.

We'd also like to see if we can train this nose to make "human" value judgments—to say that this is a good perfume, or a bad perfume, or to set the price of a bottle of wine. Or, to restate the question more scientifically, can we assign a number to a fragrance based on these patterns? Can we assign a number to a bottle of wine or a cigar that somehow quantitatively reflects a human value judgment? This is a very interesting intellectual problem. We're working with the neural-network people to find how best to approach it. We're sniffing a fine wine versus a jug wine to see if there are any differences.

We're also interested in stereo smell. We can make a sponge so thin that it responds very quickly. It then becomes possible to use the time

One can envision a little robot equipped with stereo smell crawling along a fume-filled ventilation duct, coming to a junction and telling us that the smell is coming from the left, say, and following it back to its source.

difference between when a stimulus arrives at two separate arrays to determine where a smell is coming from. Except for cockroaches, there is no creature that has stereo smell—that can locate smells based on concentration gradients between the left and right parts of its nose. Even other insects, although they have two separate antennae, turn their heads to find out, much like we do. But one can envision a little robot equipped with stereo smell crawling along a fume-filled ventilation duct, coming to a junction and telling us that the smell is coming from the left, say, and following it back to its source. Building such a robot is, at this point, an engineering task. We know that the response is fast enough, in some systems, to allow us to build one, and we know that we can make the sensors small, but I don't know if we can make them *that* small yet. We might also want to align them along a rod, perhaps, instead of in a plane, to make insect-like antennae. It's very interesting to think about bringing the sense of smell into the same electronic regime that the sense of sight has been brought to by small TV cameras, and to use smell to guide robotic systems.

We didn't invent the idea of using conductive arrays to detect odors. The British thought of it first, although we didn't know about their work when we first started ours. In 1982, K. Persaud and G. H. Dodd built a nose that used bulk conducting polymers as the swellers, as we initially did, but there aren't really that many chemical differences between the various conducting polymers. Then, a few years later, several Japanese research groups started experimenting with tin oxide, an inorganic resistor from which you can make broadly responsive films. But in order to make the films different chemically, you have to sprinkle catalysts on the tin-oxide layer, and no one really knows how to control what those catalysts do. People have also experimented with quartz crystals, similar to what's in your watch. You launch a 100-megahertz wave, much like an ocean wave, across the surface of the crystal and look at the response. If odor molecules adsorb onto the surface, the wave's frequency will change measurably. But the electronics to launch 100-megahertz waves and then read tiny changes in their frequency are quite complex, and it's difficult to envision making an array of a million such sensors on a small chip. The crystal can also be coated with swellable plastic films, as in our work, but the signal transduction is much more difficult. The beauty of our approach is that we get all of our chemical differences from the insulating sponge, whose properties we can vary broadly and systematically in a very precisely controlled way. We only rely on the conducting phase to transduce the signal into electronic form.

In conclusion, I'd like to note that I got my first taste of research when I was an undergrad working for [Beckman Professor of Chemistry] Harry Gray. When I was ready to leave Caltech, I asked him, "What should I do? How will I know what a good project is?" Harry is a very wise person, and I remember to this day what his answer was. He said I should just follow my nose. And I did. □

Nathan S. Lewis earned his BS and MS in chemistry from Caltech in 1977, and his PhD in inorganic chemistry from MIT in 1981. He went on to Stanford and tenure before Caltech lured him back in 1988. He became a full professor in 1991.

Lewis, an electrochemist, first gained national attention in the year of cold fusion as a co-leader of the Caltech team whose meticulous experiments concluded that the phenomenon couldn't hold water, much less heat it (see E&S, *Summer 1989). But his "real" research has been in the development of liquid-based solar cells that produce electricity, chemical fuels, or both when struck by sunlight.*

Lewis, who has taught freshman chemistry for the past eight years, is also the electromotive force behind the Chemistry Animation Project (CAP) videos (see E&S, *Fall 1994).*

This article was adapted from a recent Watson lecture.