Blayney and his team conclude, possibly because of environmental exposure or household practices.

The researchers also estimate that blacks in the southeastern United States are at about the same risk of type C retrovirus and T-cell leukemia-lymphoma as are blacks in the Caribbean, yet at a considerably lower risk of such infection than persons in high-risk areas in Japan. The reason is that they compared the prevalence of antibodies to the virus in a sample of healthy blacks in Georgia to the prevalence known for persons on the Caribbean island of St. Vincent and to the prevalence known for persons on the high-risk Japanese island of Kyushu. —*J. A. Treichel*

Pre-life chemistry found in meteorite

All five of the chemical compounds that provide the genetic information for life on earth have been identified in samples of a single meteorite, a scientist this week told the American Chemical Society's annual meeting in Washington, D.C. "We have not found E.T.," says Cyril Ponnamperuma of the University of Maryland in College Park — the compounds, called bases, are merely "prebiotic" material. But confirming that the full suite of them can form on an extraterrestrial body could certainly bear on the loaded question of whether earth is one of many abodes of life or an incredibly lonely oasis in the universe. Also essential in exploring the issue is understanding the ease or difficulty with which such material can form, and Ponnamperuma reported that all five bases have also been produced in a laboratory experiment. The result, he says, "makes the creation of life chemicals appear to be simple, almost inevitable."

Four of the five bases — adenine, guanine, cytosine and thymine — comprise the "letters" of the genetic code, structured into the double helix of DNA. The fifth, uracil, along with adenine, guanine and cytosine, is part of the protein-making instructions coded in RNA. All five, according to Ponnamperuma, have now been found in samples of the Murchison meteorite, which fell in Australia in 1969. (Some of the bases had been previously reported in the Murchison meteorite and others, but the presence of the complete set was not certain.)

In the laboratory, he and colleagues exposed a model of a primitive atmosphere (methane, nitrogen and water) to an electric discharge. Past studies had produced one or more of the bases, Ponnamperuma says, "but here, for the first time in one experiment, we see that they are all there." Because the result is born of such an essentially simple process, he adds, "this I consider to be of even greater importance than the discovery of the bases in the meteorite." —*J. Eberhart*

Computing a machine's world view

When most people look at a cartoon, they generally have little trouble recognizing a set of squiggly lines drawn on paper as a familiar face, a penguin or some other figure. Somehow, the human brain can come up with the right answer even when descriptions are incomplete or partly wrong.

For years, researchers in artificial intelligence have sought ways to mimic the human brain's remarkable ability to recognize objects and ideas. Furthermore, they wanted a machine that could learn from its experiences by steadily widening its ability to identify similar objects in different settings. Last week, three scientists reported success in designing a machine that, in their view, meets these requirements and behaves more like neurons in the brain than any other model available.

Scott E. Fahlman and Geoffrey E. Hinton of Carnegie-Mellon University in Pittsburgh and Terrence J. Sejnowski of Johns Hopkins University in Baltimore described their "Boltzmann machine" at the National Conference on Artificial Intelligence, held in Washington, D.C. Their machine is designed to find and settle into a "state" that best suits or interprets a particular observation without requiring an exact match.

The machine consists of a network of simple binary units that can be either "on" or "off." These units may each be connected to, perhaps, a thousand others. A numerical weight assigned to each connection represents the strength of the link. The connection weights store knowledge within the network about the plausibility of each interpretation of an observation. In this model, a "concept" (say, a penguin or an elephant) is stored as an "on-off" pattern spread over many units. The weights act as constraints on the system.

A small number of the machine's units have connections with "the outside world" to provide for input and output of information. The rest, hidden inside, represents encoded knowledge. Fahlman says the game is to satisfy the constraints as much as possible for a given input signal. How well the constraints are satisfied can be represented by a number that behaves very much like potential energy in a physical system. The lower the "energy," the more happy the system is, says Fahlman.

In a physical system consisting of, for instance, a ball rolling through an undulating landscape of hills and valleys, the ball will try to settle into the deepest valley, where it will have its lowest potential energy. In the Boltzmann machine, this potential energy is the sum of all the unit states (1 or 0 for each) and the weights between the units. For a given input, the units make local decisions, based on constraints imposed by their neighbors, to switch on or off so that the total potential energy is a minimum. As these adjustments occur, the network eventually settles into a minimum that best satisfies the input signal and all the constraints.

A common task for such a system is to find the stored description that best matches a set of observed features, even if the match is imperfect. The most likely interpretation of, say, a diagram, will be the one that drives the network into its lowest energy state. This corresponds to the description that best fits the weighted combination of observed features and expectations. If there are several good descriptions, it is biased toward the best.

"No single weight or unit in this thing is critical," says Fahlman. "You still end up in the same general region." Many people have suggested that such a distributed representation of memory would be more reliable than a memory in which each unit represents a particular idea. It is also more consistent with what is known about the workings of the brain.

Learning occurs because input signals that do not match exactly what is already represented in memory modify the weights and unit states slightly. The researchers introduced a random noise element into their network to reduce the chances that the system will get caught in a local minimum (like a ball getting stuck in a valley that may not be the lowest point in the landscape). This probabilistic procedure, somewhat akin to "shaking" the system, also alters the weights assigned to network connections so as to reflect better the structure of the outside environment.

The Boltzmann machine, unlike conventional computers that must process information one step at a time, has a large number of processing elements working on a single task at the same time. This kind of "massively parallel" organization, Fahlman and his colleagues believe, is necessary to provide the enormous computational power that some aspects of intelligent behavior seem to require.

"This is a very new theory," says Fahlman. The researchers have done computer simulations of networks with up to 50 units, but much larger simulations are needed to study more realistic situations. Eventually, a Boltzmann machine (right now only a computer program) will be built into a silicon chip for large-scale studies.

"We need a better handle on how big a Boltzmann machine we need in order to get a given complexity of behavior," Fahlman says. "A human brain, set up this way, has probably at least 100 million units."

Fahlman concludes, "Whether what we're talking about here in the Boltzmann machine is in fact what happens in a neural network is an open question....But of the models [available] it comes closest or is least obviously wrong." —*I. Peterson*