Research News

Forecasting the Global AIDS Epidemic

Good computer models might help persuade officials of developing countries to institute anti-AIDS strategies, but modeling has proven easier said than done

Last November, President Yoweri Museveni of Uganda and about 10 of his top advisers watched as a personal computer specially programmed for the occasion drew full color graphs of AIDS prevalence in his country. The graph lines marched inexorably upward, as the computer drove the point home by emitting beeps of ever-rising pitch. Then, Ugandan AIDS officials instructed the computer model to factor in the effects of behavior changes such as condom use. Suddenly, the graph lines—and the pitch of the beeps—began to drop. The next day, Museveni reversed an earlier stand and told a Ugandan newspaper reporter that people should wear condoms to save their lives.

AIDS modelers like to cite that success story as an example of how their science can boost AIDS prevention efforts in the developing world. And such successes are desperately needed, as AIDS strikes hard at many African cities. For example, epidemiologist Joseph Konde-Lule of the University of Makerere in Kampala, Uganda, estimates that more than 20% of adults in that city are infected with the AIDS virus.

But despite the need demonstrated by numbers such as this, some prominent epidemiologists question whether current AIDS models are good enough to use to set public health policies. “I myself don't have great faith in using models for projections,” says James Chin, who is currently chief AIDS forecaster for the World Health Organization's (WHO) Global Programme on AIDS, although he plans to leave that post sometime in the next year. “I'm not sure any are ready to be used on a country-by-country basis.”

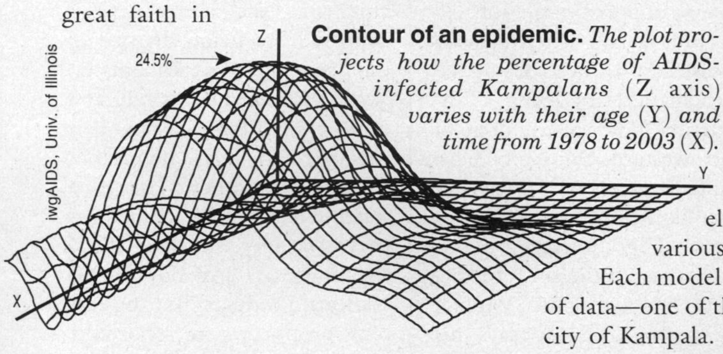

iwgAIDS, Univ. of Illinois

Contour of an epidemic. *The plot projects how the percentage of AIDS-infected Kampalans* (Z axis) *varies with their age* (Y) *and time from 1978 to 2003* (X).

There's good reason for Chin's skepticism. During the 6 to 7 year history of AIDS modeling, the researchers have found that their predictions rarely agree. That's partly because the modelers are handicapped by the poor quality of the data being fed into their computers—the old “garbage in-garbage out” problem. Especially in the developing countries, firm numbers on current AIDS prevalence rates and detailed demographic data are hard to come by, Chin says.

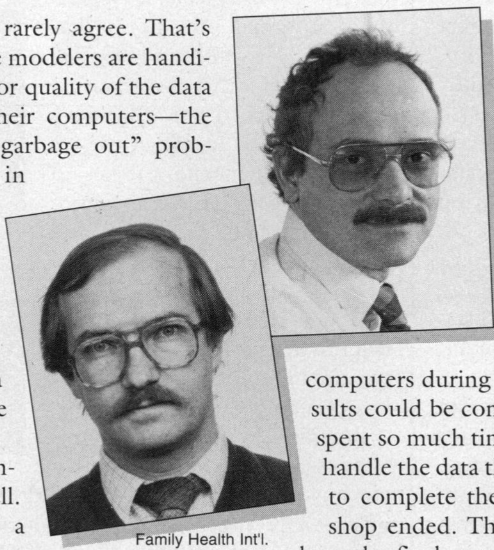

Univ. of IL

Family Health Int'l.

Tracking AIDS. *David Sokal* (left) *and Steven Seitz are trying to improve the models.*

And there's another issue as well. The models use a host of different approaches. Some take whole populations as their starting point, for example, while others look at people individually and try to project from them what will happen to the population (also see box on next page). So when the projections disagree, the researchers are hard pressed to tell where their models went wrong.

Which is why last month, in an effort to put AIDS modeling on a firmer footing, WHO and AIDSTECH, an AIDS prevention program funded by the U. S. Agency for International Development, invited three of the top modeling teams to Durham, North Carolina, for a workshop. The goal was to compare how well the models predict the outcome of various AIDS prevention plans. Each model would tackle the same set of data—one of the best available, from the city of Kampala. And if all three arrived at the same conclusions, both the models and the preventive strategies would get a welcome shot of credibility. “Having three models tuned to one real-life situation will be wonderful for future questions,” says John Stover, vice president at the Futures Group in Glastonbury, Connecticut, who helped develop the demonstration model for Uganda.

Alas, the Durham workshop illustrated all too well just how difficult it is to model AIDS. Although the original idea was to have some of the modelers run their programs on personal computers during the workshop so the results could be compared on the spot, they spent so much time arguing about how to handle the data that they didn't have time to complete the runs before the workshop ended. They now don't expect to have the final comparisons before September. But meanwhile, there is a ray of hope: Preliminary results from two of the models show about 80% agreement, says workshop organizer David Sokal of AIDSTECH. And that's already a big improvement over what happened the last time AIDS modelers got together.

That meeting was held a year and a half ago at the United Nations in New York City, as the modelers began to realize they would have to work together to improve the accuracy of their models. The first question they wanted addressed was whether there might be a simple explanation for their divergent results. Could it be just because they often tackled different data sets?

So at the New York meeting six modelers analyzed the same data set, representing a sort of generalized African AIDS situation, in the hope of reaching a consensus result. But the results, published this spring,* were disheartening. The models came up with simulated prevalence rates for AIDS virus infection that ranged from 3% to 40%.

What the researchers didn't appreciate when they set up the New York meeting, Sokal says, is the extent to which each group would have to modify the data to get them to fit the different formats used by the individual

*The AIDS Epidemic and its Demographic Consequences. Proceedings of the United Nations/World Health Organization Workshop on Modeling the Demographic Impact of the AIDS Epidemic in Pattern II Countries; UN/WHO, 1991.

models. So even though they started with the same data set, by the time the programs were up and running, the models were no longer dealing with the equivalent data.

Mindful of this hazard, the organizers of the recent Durham workshop designed their conference somewhat differently. As before, they invited a handful of researchers, each a leading proponent of a different type of model: Steve Seitz for the International Working Group AIDS model (iwgAIDS model), Bertran Auvert of INSERM, Paris, for the SimulAIDS model, and Rodolfo Bulatao of the World Bank for that institution's model. Also as before, each modeler was to tackle the same test: This time they asked what would happen to the AIDS epidemic in Kampala if WHO's Big Three AIDS prevention strategies—using condoms, treating other sexually transmitted diseases (STDs), and reducing the number of sexual partners—were put into play.

For the Durham meeting, though, the organizers allowed more time—11 days instead of the 3 that the New York participants had—so that the modelers could get together and try to work out ways of standardizing the data and the prevention strategies before feeding the information into their computers. That proved to be easier said than done, however, taking more than a week.

One typical wrangle concerned how to categorize AIDS risk groups, Sokal says. SimulAIDS has three: People with long-term partners, people with short-term partners, and people who do one-night stands. But the iwgAIDS model has four: high- and low-risk singles and high- and low-risk married people. So the researchers had to tease out the equivalent groups. For example, women who do one-night stands, from the SimulAIDS model, could be considered somewhat equivalent to single women with high-risk behaviors in the iwgAIDS model.

The World Bank's Bulatao wasn't able to stay the whole 11 days and didn't even get to start his run—he plans to do it in September—but the Seitz and Auvert groups got far enough along to do a preliminary comparison. Both their models predicted that either condom use or partner reduction alone would help blunt the AIDS epidemic. And Seitz's model further showed that a strong dose of self control—such as slashing the number of one-night stands in half—might do even better, reducing AIDS virus prevalence by nearly 60% and causing what Seitz calls "a qualitative shift" in the epidemic. He warns, however, that the projections are almost certainly too rosy because they are based on some improbable assumptions—such as that people will drastically change their sex behavior.

As for the third preventive strategy, treat-

Computing the Problem

The AIDS epidemic has spawned dozens of computer models as researchers struggle to find ways to predict the epidemic's course and to assess the effects of strategies aimed at preventing the disease. But all models intended for use in developing countries have one requirement in common: The programs have to be simple enough to run on personal computers, which are more likely to be available in those countries than the big mainframes. At a recent Durham workshop (also see story), researchers showed off a trio of such models, each a state-of-the-art representative of a different modeling style.

Steve Seitz of the University of Illinois presides over the scientific portion of the largest and most complex program, the Interagency Working Group AIDS model (iwgAIDS). Developed in concert with the U.S. Agency for International Development and the U.S. Bureau of the Census, it's a deterministic model that uses equations to simulate how different variables may interact with one another in influencing the course of the AIDS epidemic.

One of this model's chief advantages is flexibility: The user can plug in more than 100 variables—far more than most models can afford—describing a population's demographic characteristics, sexual behavior, and prevalence of AIDS virus infection. Because of the great demographic detail of iwgAIDS, it can be tailored for populations as distinct as those of rural Uganda and New York City. A presentation version of this model wowed Ugandan politicians enough to convince them, at least initially, that condom use could help cut the spread of AIDS in their country.

But some field-oriented epidemiologists say that the iwgAIDS model's biggest strength is also its biggest weakness: Most countries simply have no data on many of the variables the model uses. These include, for example, the age-specific divorce rate, how many extramarital affairs each man pursues annually, or the duration of genital inflammation in an "average" sexually transmitted disease. "I'm afraid Steve's mathematics is much, much better than any of the data we're giving him," says Robert Biggar, international AIDS coordinator of the viral epidemiology section at the National Cancer Institute.

Echoing that thought, several field epidemiologists at the Durham workshop—trench workers who actually use the models—told *Science* they prefer another model, called SimulAIDS, which was developed by Bertran Auvert, professor of statistics and epidemiology at INSERM, Paris, and his colleagues. SimulAIDS is a so-called micro model because it takes individuals, rather than populations, as its starting point. It assigns probabilities to various events in each individual's lifetime, such as marriage, visiting an AIDS-infected prostitute, and so on. The overall results for a population are then projected from the sum of what the computer calculates happened to 10,000 to 15,000 individuals.

But although it's easier to understand, SimulAIDS' conclusions are only as valid as its probability estimates. Also, it cannot consider both urban and rural populations simultaneously. And it is less adept than iwgAIDS at simulating the effects of rare events, such as drug use and homosexual contacts, which in the developing countries make very minor contributions to AIDS transmission compared to heterosexual sex.

The third model, developed by Rodolfo Bulatao of the World Bank, is slightly cheaper and more user-friendly than the other two since it doesn't require complicated data and is relatively easy to run. It simply counts deaths due to AIDS and then incorporates them into population projections for all 187 countries tracked by the Bank. Modeling cynics tend to prefer this model because it's easy to see how the data affect the results. With more detail, "you just don't have information for all those parameters—you are inventing things," says one user, modeler Alberto Palloni of the University of Wisconsin-Madison.

But other modelers insist that the World Bank program is just too simple to capture the spread of AIDS, and especially to simulate the effects of prevention programs. It steers clear, for example, of factors such as the age structure of the population that influence the spread of AIDS—but complicate the models. It also provides a less accurate simulation of the effects of sexually transmitted diseases on AIDS.

Which model approach will prove best in the end? Perhaps the answer will be that each has its uses and all will have the opportunity to help health officials in developing countries fight a deadly epidemic. ■ **E.C.**

ing the STDs, both models predicted that this would have little effect on the spread of AIDS, although that doesn't mean that STDs aren't important. Indeed, both models suggested that if the STDs could be wiped out there would essentially be no AIDS epidemic in the developing countries. But when the models take into account some gloomy realities—for example, many people who have STDs don't get severe symptoms and don't seek treatment—they predict that ordinary treatment programs may not help much against AIDS.

While the early model comparisons are encouraging, it's too soon to say whether they will help resolve the larger philosophical debate: How should models be used in setting public policy? Some, like WHO's Chin, argue vehemently that modelers ought to be very sure of their science before they take their programs abroad, because everyone's credibility will be undermined if the results change whenever a new data point comes in.

But pragmatists such as Melinda Moore, deputy director of the International Health Program Office of the Centers for Disease Control, say the modelers can't afford to wait for perfect simulations. Moore argues that the models are good enough to serve as general guides, even if they don't give precise long-term projections. "With or without data, with or without models, people are going to make decisions based on what's in their heads," she asserts. "Why not have models as an additional guide to intuition?"

Indeed, many health officials in developing countries are seeking any possible aid for convincing the powers-that-be to adopt preventive strategies. And they think that razzle-dazzle computer programs can help. "If you go before a decision maker with a personal computer and make a demonstration, that gives much credibility to what you are saying," explains Kasela Pala Mambwe of the international AIDS program in Kinshasa, Zaire.

The sad truth may be, however, that even the razzle-dazzle isn't necessarily enough. Some AIDS watchers fear that Museveni may recant his conversion. In his public remarks at the Seventh International Conference on AIDS, held in Florence in June, the Ugandan president gave only lukewarm support to condom use. Then again, Uganda's first promotional condom campaign, which was planned long before Museveni saw the computer go through its paces, hit radio, television, and pharmacies about a month ago. "Be wise," the ads urge. "Always wear Protector condoms." Perhaps the ads will give modelers some novel real-world data against which to check their results, as well as encouraging Museveni to keep up his support for condom use. ■ ELIZABETH CULOTTA

The Education of Silicon Linguists

By cramming translating programs full of facts, researchers teach them to "understand" the texts they translate

TRY TRANSLATING THE VERB "HIT" INTO Spanish: you can choose *chocar*, *golpear*, *acertar*, or about 30 other verbs. But most of the possibilities are a world of difference apart. Hit someone in the nose and the verb is *golpear*; drive a car into a tree and it's *chocar*. Likewise, the Spanish verb *comer* can be translated as "to eat," "to capture," or "to overlook." The choice you make hinges on your understanding of the word's context. Thus the paradox facing computer scientists struggling to build electronic translators that can match your average United Nations multilinguist: If accurate language translation requires understanding, how can an unthinking computer do it?

Although computers are no more thoughtful than they ever were, researchers at several universities seem to be on the homestretch in the long quest for a system that could "understand" a document and translate it accurately. At Carnegie-Mellon University, for example, a computer translator is spitting out perfect translations of television repair manuals in a matter of minutes, moving adroitly among English, Japanese, German, Spanish, and French. And at New Mexico State University and the University of Southern California (USC), computer translation systems are gearing up to tackle computer operation manuals.

No one is boasting a program that can translate James Joyce's *Ulysses*—to say nothing of *Finnegans Wake*. But these more mundane demonstrations still reflect major programming advances, which are enabling computers to store and manipulate a vast array of knowledge about the realities behind word usage. Says Yorick Wilks, head of the Computing Research Laboratory at New Mexico State: "Things are starting to happen now." Adds Charles Wayne, a program manager for the Defense Advanced Research Projects Agency (DARPA), which will begin pumping $1 million into machine translation research in September: "There's a perception that the field is poised to see some major advances."

That would represent a sea change in what had been, in recent years, a somewhat moldering field. Scientists first experimented with machine translation in the late 1940s. And there was lots of excited discussion of its promise in the 1970s, culminating in the introduction of several commercial systems. But by the early 1980s, says Eduard Hovy, who heads the machine translation effort at USC's Information Sciences Institute, the field "had died almost a complete death."

On the plus side was the track record of commercial translating systems that, according to Jaime Carbonell, director of the Center for Machine Translation at Carnegie-Mellon, have been able to trim millions of dollars from the $30 billion that Carbonell estimates governments and industries spend each year translating technical documents. But those systems are essentially dictionaries on a computer: insensitive to context, they stumble over synonyms and idiomatic expressions, picking the wrong word anywhere from 10% to 50% of the time, says Carbonell.

Know-Nothing Translation

While many machine-translation researchers are stuffing their computers with knowledge, Peter Brown, a computer scientist at IBM's Thomas J. Watson Research Center in Yorktown Heights, New York, is stuffing his computer with hundreds of millions of words worth of the Canadian Parliament's English and French proceedings. In a decidedly anti-knowledge tack, Brown has programmed the computer to compile statistics describing the relation of words in the two languages. Given any new French or English text, the computer will then be able to refer to its vast statistical tables and spit out the most probable translation—an approach Brown thinks may yield serviceable machine translations without the extensive training that knowledge-based systems require.

The underlying strategy is that of a cryptographer looking for the statistical patterns that will crack a code. "We imagine that French is a garbled version of English," he says. To