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

And that's where researchers' frustrations mounted. For decades, scientists had been hacking away at the problem of programming a crude contextual understanding into their computers, the principal hope for dramatically cutting the error rate. But the research progressed slowly, hampered by a lack of funding and computer power. Still, with hundreds of millions—perhaps billions—of dollars to be saved, universities, businesses, and the government have continued to hanker after a perfect translation system that would quickly decipher reams of technical documentation without needing a human editor to clean up mistakes.

The exponential growth in computing power—five orders of magnitude over the past 40 years—is now providing the electronic muscle for a new crack at machine translation. Added to that will be the substantial financial boost from DARPA, which will fund a knowledge-based machine translation (KBMT) collaboration between Carnegie-Mellon, New Mexico State, and USC to develop a system called Pangloss. Pangloss, its developers hope, should eventually be able to produce flawless translations of documents as complex as newspaper articles from Spanish, German, and eventually Japanese into English.

On paper the concepts underlying Pangloss and a key model, Carnegie-Mellon's current computer translating program, seem straightforward. Running on several computer workstations, Carnegie-Mellon's system parses each sentence in the starting text, then sends the various parts of speech through a "concept lexicon," a collection of hundreds of facts about the topic at hand. In the process, the computer homes in more precisely on the meaning of each word. By applying what it knows about balls to the verb in "John is hitting the ball," for example, the computer infers that the sense must be "to strike" rather than "to collide with."

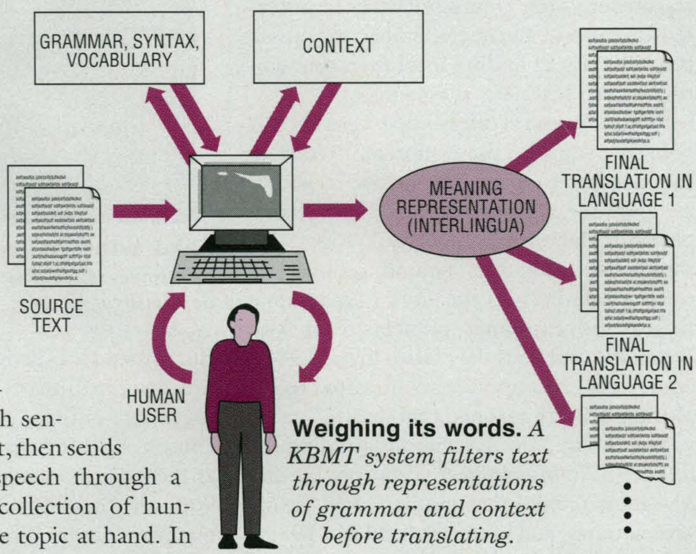

Weighing its words. *A KBMT system filters text through representations of grammar and context before translating.*

Having identified just which sense of "hit" is intended, the program finds the precise code for it in Interlingua, a generic translating language developed at Yale University and Stanford in the 1970s. Interlingua, an unambiguous code of grammar and meaning, can then readily be decoded into the target language.

One reason even Carnegie-Mellon's system hasn't gotten beyond television manuals, says Sergei Nirenburg, a Carnegie-Mellon computer scientist, is the challenge of giving the computer enough information for it to form "as realistic a model of the world as possible." Interpreting even a simple sentence accurately can take an overwhelming amount of context. To understand the sentence "John hit the first baseman," for example, the computer would need to know that players sometimes run into each other, so that here "hit" means "collided with." But if the object of the verb is "umpire," the computer needs to know enough about baseball to recognize that the more likely meaning is "strike." For each word in the translator's vocabulary a human programmer spends hours compiling contextual facts.

And Nirenburg emphasizes that it's not enough to load the computer with contextual information; the facts also have to be organized in a framework in which the computer can find the right piece of knowledge quickly enough to make the program useful. The framework serves the computer much as a net serves a fisherman: If all the facts are in place and tightly connected, no fish get away. "Otherwise there are mistakes that we catch," Nirenburg says. "For example, if the program spits out that you have a tasty car instead of a blue car, you're in trouble."

By trying out the Carnegie-Mellon system on text after text—mainly operating and repair manuals—Nirenburg and his colleagues try to detect and fill gaps in the concept lexicon. But "as the systems are fleshed out," says Carbonell, "the human component of machine translation will diminish." Other experts are more cautious: "KBMT goes in fits and starts," says Wilks, whose group at New Mexico State is developing Ultra, a translation program that asks an operator to supply the contextual knowledge whenever it stumbles during a translation. "It's impossible to predict anything about it."

And some experts in traditional, dictionary-based translating systems aren't even sure it's worthwhile trying to build a knowledgeable translator. "The problem is it's not really machine translation's goal or job to create a vast knowledge base," says Cris A. Fitch, vice president of engineering at Systran Translation Systems, Inc., the dean of machine-translation companies. Fitch says that computers already do a good job without understanding the material. "You're just not going to get a perfect translation," he says.

The DARPA project may tell who's right. The three participating universities will offer up complementary expertise: Wilks' group at New Mexico State will work on building vocabularies and parsing sentences; Carbonell's group at Carnegie-Mellon will concentrate on the concept lexicons; and Hovy's team at USC will develop routines for deciphering Interlingua into the target language.

DARPA will be keeping a close watch on the research's progress. And the technique it will use is one that is all too familiar to language students: According to Wayne, the evolving Pangloss program will sit for tests not so different from the reading comprehension section of the high school SAT test. Wayne is hoping for good performance, but he's far from Panglossian. "If we can do well with either German or Spanish," he says, "we'll be happy." ■ RICHARD STONE

ungarble the two languages, Brown has programmed his computer to examine every pair of corresponding English and French sentences in the proceedings. It scrutinizes the words in clusters of three, compiling data on repeated associations.

The cluster associations provide the context the program uses to translate individual words. The French word *prendre* usually means "take," but when the system encounters the sentence *Je vais prendre la decision*, the words flanking *prendre* enable it to home in on the correct translation: "I will make the decision."

Because the system considers words only in small clusters, it can miss the larger picture that can affect a word's meaning. But Brown plans to refine his statistical models by feeding them documents from the European Economic Community. And even some aficionados of knowledge-based translation think he may have something. "Brown's done some intriguing research," says Yorick Wilks of New Mexico State University. "Down the road, I think a blend of knowledge and statistics will be the answer." ■ R.S.