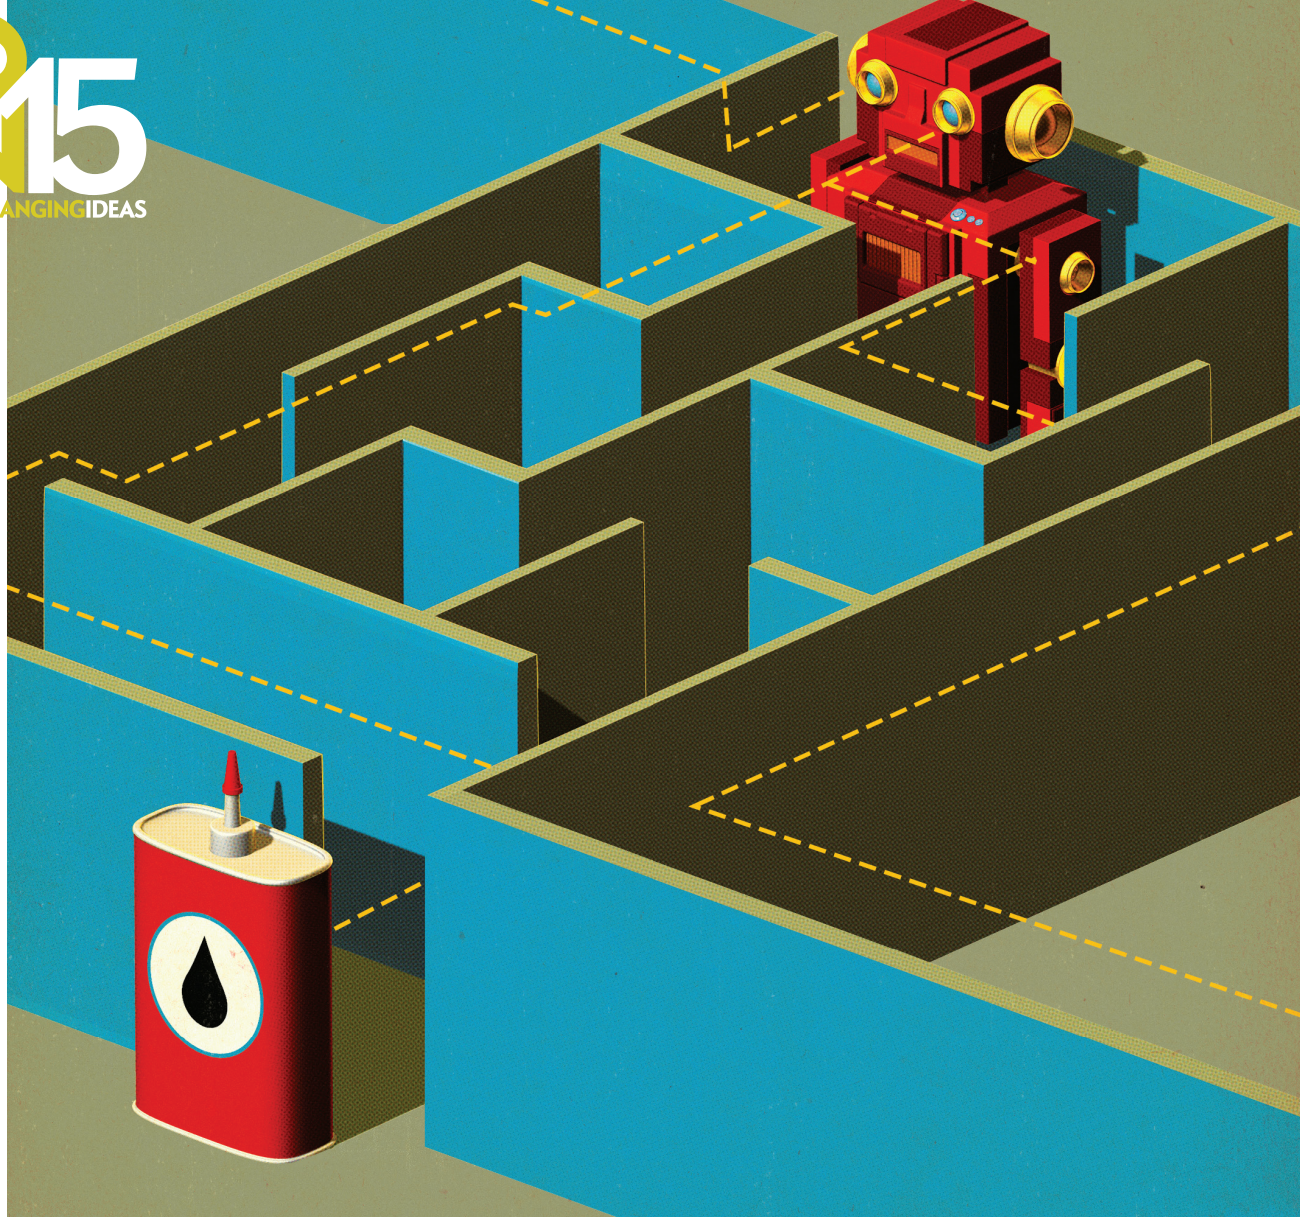

Machines That Teach Themselves

Deep-learning technology is helping A.I. fulfill its promise

Google, Facebook and other corporate giants are taking major strides in building technology that can learn on its own. Their efforts rely heavily on something known as deep learning.

Rooted in the decades-old idea that computers would be smarter if they operated more like the human brain, deep-learning networks consist of layer on layer of connected computer processing units called artificial neurons, each of which performs a different operation on the input at hand—say, an image to be classified. The difference between conventional neural networks and deep-learning ones is that the latter have many more layers. The deeper the network—the more layers—the higher the level of abstraction at which it can operate.

Deep learning gained momentum in the mid-2000s through the work of three key figures—Geoffrey Hinton of the University of Toronto, Yoshua Bengio of the University of Montreal and Yann LeCun of New York University—but it only recently began making commercial inroads. An example is the Google Photos app, which came out in May. The software can upload all the images from my iPhone, correctly identify my wife, son and grandson, and then dump their photographs in separate digital bins marked by thumbnail images. It can do this because it has learned to recognize faces through exposure to millions of images analyzed by the system. As it runs an image through each successive layer of its network, the software identifies elements within the image at an increasing level of abstraction—until it ultimately can detect the whole face within the picture.

Once it has trained on enough faces, it can spot the noses and mouths of individual people in images it has never seen before.

Deep learning can do much more than organize pictures. It may, in fact, mark a step toward artificial intelligence that exhibits intelligent behaviors virtually indistinguishable from those of its human masters. In February a team of A.I. experts from the London-based firm DeepMind (which Google bought in 2014 for $617 million) reported that it had used deep learning to build a computer that could teach itself to play dozens of Atari video games. After a lot of practice, the software beat expert human players at half of those games. A small step, but the machine age has to start somewhere. —*Gary Stix*