Case Study

Teaching robots to reproduce human handwriting and drawing

BROWN

KINOVA

BACKGROUND

Students in Professor Stefanie Tellex's graduate course "Topics in Collaborative Robotics" at Brown University had a vision: Teaching robots to infer human handwriting and immediately reproduce what they see using human-like strokes.

The thinking behind this project was that in order to collaborate with humans, a manipulator robot should be able to draw or write on a white board or even a post-it note.

Long term, the ability to write would enable a robot to put up a sign directing people that a hallway was closed, to produce art using physical mediums such as a paint brush or a pen, or to address and mail a letter. Additionally, the robot could potentially engage in teaching activities at a white board, writing a math equation or drawing a diagram. These skills rely on the ability to produce a policy to draw with a writing utensil.

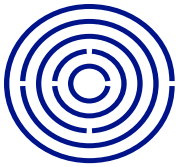

THE CHALLENGE

Up until now, work in this area required the robot to have information about the stroke order in advance, meaning that for the robot to write something, somebody would have to program the stroke orders each time.

In order for this project to succeed, students had to find a way to teach the robot to reproduce an image of just-drawn handwritten characters by inferring a plan to replicate the image with a writing utensil. To achieve this, the research team of computer scientists had to develop an algorithm using deep learning networks so the robot could analyze images of handwritten words or sketches, deduce the likely series of pen strokes that created them, and reproduce them using stroke patterns similar to human handwriting.

OBJECTIVES

The key objective was to make it possible for a manipulator robot to reproduce the target image of characters or drawings only from its pictorial information, while preserving details such as line continuity and stroke orders.

After training the model, the hope was that the robot gained general knowledge of drawing, which would allow itself to write anything at first sight.

BROWN

KINOVA

THE APPROACH

The first thing was to determine exactly what types of actions were required to reproduce human-like writing strokes. Two were identified:

- Drawing action, which carefully draws each stroke.
- Shifting action, which shifts the pen to the start of a new stroke once the current stroke is finished.

In order to incorporate this idea, the approach contained two distinct models:

- A “local” model which observed a 5x5 pixels region around the current pen-tip location and determined in which direction to move in its region as well as when to end the stroke.
- A “global” model in charge of moving the robot’s writing utensil to the next stroke of the character.

They trained two separate neural network models for the different scales with a dataset of Japanese characters, then combined them in runtime.

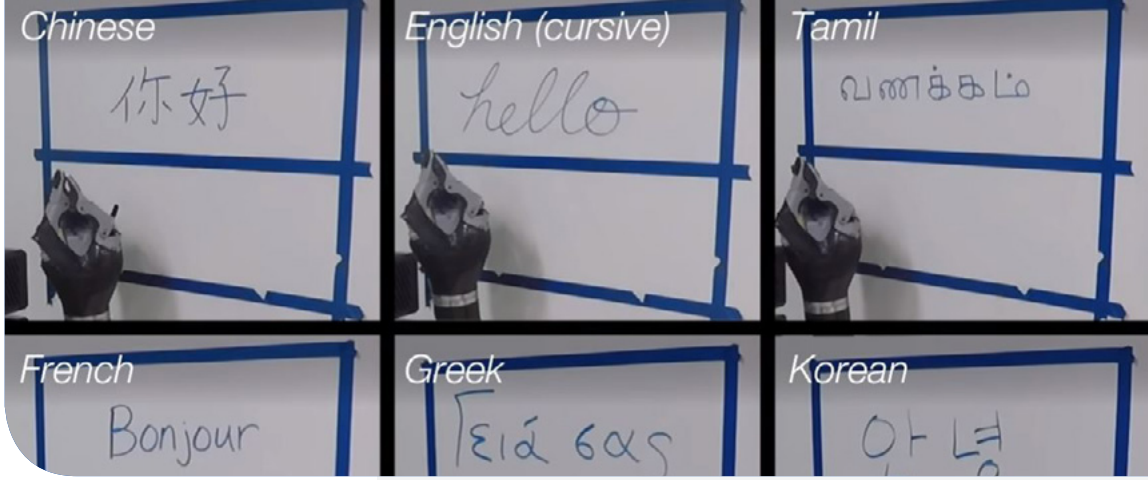

THE SOLUTION

The team of data scientists at Brown University built a system that works in real-time, enabling the robot to view an image, infer a plan to replicate it, and immediately start drawing it.

The robot draws each target stroke in one continuous drawing motion and does not rely on handcrafted rules or on predefined paths of characters. Instead, it learns to write from a dataset of demonstrations.

When the robot begins to draw, it must collect the following information before making the next movement:

- Already visited regions
- Current location
- Difference image
- Continuously connected target region

BROWN

KINOVA

THE SOLUTION

To measure performance, they introduced two metrics: pixel accuracy and stroke accuracy. Pixel accuracy measures how similar the target image and the drawn image are. They also measured the performance by stroke accuracy, which checked if the robot drew one stroke in one continuous action.

The network enabled the robot to not only reproduce characters in multiple languages, but also replicate any stroke-based drawing, immediately after seeing it for the first time.

The solution was first tested in simulation mode, then with different robots.

“I want a robot to be able to do everything a person can do,” Tellex said. “I’m particularly interested in a robot that can use language. Writing is a way that people use language, so we thought we should try this. Kinova’s robots have been reliable and easy to use for our research team and reproduced human drawn characters and images very accurately.”

- Stefanie Tellex, Assistant professor of computer science, Brown University.

BROWN

KINOVA

THE RESULTS

One of the robots used was a Kinova mobile manipulator equipped with their Jaco robotic arm and the Kinect 2 as a sensor. With its precise movement capabilities, the robot reproduced the target images very accurately.

The first tested model with Japanese characters marked 99% pixel-wise accuracy in reproduction and 96% preservation accuracy in line continuity. They then tested the system with the task of reproducing “hello” in 10 different languages as well as the sketch of the Mona Lisa. While the model had never seen those images before, it successfully reproduced them.

Overall, the approach for inferring a sequence of commands for drawing a character given a bit-mapped image of the character was a success. They demonstrated that their approach enabled the robot to gain general knowledge of handwriting and replicate the target image in both simulation and two different robotic environments.

Most significantly, the trained model accurately predicted drawing procedures for foreign characters, even those not in the training dataset, as well as arbitrary line drawings.

KINOVA

info@kinova.ca

North America: 1 855 6 KINOVA

Worldwide: 514-316-0144

kinovarobotics.com