Paper Cubes: Evolving 3D characters in Augmented Reality using Recurrent Neural Networks

Anna Fusté
MIT Medialab
afuste@media.mit.edu

Judith Amores
MIT Medialab
amores@media.mit.edu

David Ha
Google Brain
hadavid@google.com

Jonas Jongejan
Google Creative Lab
jongejan@google.com

Amit Pitaru
Google Creative Lab
pitaru@google.com

Abstract

Paper Cubes is a DIY Augmented Reality (AR) Platform that uses paper cube patterns and an AR application to teach computational concepts such as Neural Networks in a simple and engaging manner. We present an AR representation of a Recurrent Neural Network in the form of stick figures that move in the user's physical space and evolve over time. We argue that using Recurrent Neural Networks to drive agents in AR and in real-time can potentially help to generate more interactive and engaging storytelling, gaming and learning experiences.

1 Introduction

AR is an explanatory medium that takes advantage of the physical space, objects and surfaces. It gives us the opportunity to create interfaces to better understand what AI does in a visual and spatial manner. Paper Cubes [2] was created to teach basic computational skills as well as more advanced programming concepts such as Artificial Intelligence (AI) and Machine Learning (ML) [1] using AR in a more visual and engaging way.

Recurrent Neural Networks (RNN) [10] are a complex concept to visualize and understand. Some work has already been developed in generating more simple and engaging interfaces using neural networks [3,5,7] and crowd simulation [9] in virtual environments. Norton et al. [6] presented a visualization suite showing how adversarial examples fool deep learning. Teachable Machine [8] is a Google AI Experiment where the user can learn about basic AI by using their camera. There is also a good amount of research on endowing simulated characters with locomotion and other movement skills. Heess et al. [4] demonstrated how a rich environment can help to promote the learning of complex behaviors. As a result, they also presented an attractive visualization of avatars and other creatures moving in space. With Paper Cubes, we extend on this idea by using the physical environment and the digital affordances of AR to teach a group of characters how to avoid obstacles.

2 System Overview

Our system consists of 6 paper cubes that can be used to control and generate AR. We used Unity3D to create the virtual assets and animations. We used the Vuforia SDK to detect the cubes in space. Each cube has a different pattern that is recognized by the phone's camera. The Jump, Stop and Turn cubes are defined as regular computing cubes. They trigger actions in AR and establish the language between the physical environment and the digital content. The AI Cube generates an ongoing process of improvement over time. It gives intelligence to the characters making them smarter. When the AI

31st Conference on Neural Information Processing Systems (NIPS 2017), Long Beach, CA, USA.

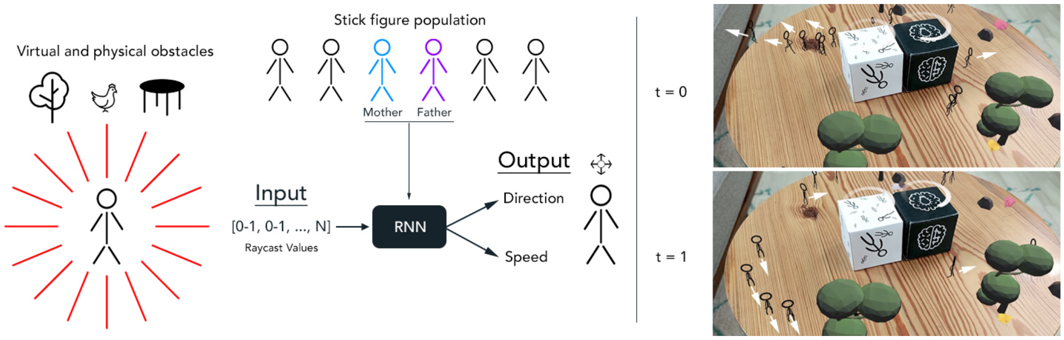

Figure 1: Paper Cubes system overview diagram and visual output in time

Paper Cube is placed next to the Start Paper Cube, a set number of stick figures will start walking in random directions.

Each stick figure is controlled by a unique RNN. The output of the network controls the actions of the stick figure, indicating direction and speed. The inputs are gathered from the stick figure's lidar system. The stick figures have a lidar-like system that detects objects close to them. This system generates a specific number of rays from the center of the stick figure outwards spread uniformly in 360 degrees. This raycast system can detect virtual obstacles such as animals or trees and physical objects such as the edges of the table. The raycast system computes the distance between the stick figures and the objects around them and creates an array of floats ranging from 0 to 1. When closer to an obstacle, the value gets closer to 1. If there is no obstacle, the value is 0. This array gets fed into the RNN as the input. The output of this RNN defines the direction and speed of each stick figure. Figure 1 shows the system overview, along with the visual output at two moments in time. Specifically, in t=0, the stick figures run at high speed towards the edges of the table and the virtual objects. In t=1, the stick figures have learned how to reduce speed and change direction when closer to obstacles.

3 Training

Unlike existing machine learning approaches that typically pre-train the weights of a neural network before deployment, we wanted to experiment with the possibility of training the neural network in real-time, while the user is playing with the cubes. For this reason, we chose to use an evolution strategy and a large population of stick figures in the environment to evolve a suitable set of weights that will guide the agents to survival. In our population, there is a maximum number of stick figures (30) that can exist in the scene at the same time. Every time a stick figure disappears, another one is created. When created, each stick figure is assigned a mother and a father from the stick figures that already exist in the scene. New stick figures learn from their assigned parents. Over time, these inherited traits will increase the chance of survival and, eventually, the stick figures learn to avoid the virtual and physical objects. The example visualization of this can be seen at: `https://experiments.withgoogle.com/ar/paper-cubes`.

4 Conclusions and Future Work

We presented Paper Cubes, an example of an engaging and creative visualization of RNNs that makes AR stick figures evolve over time. Our system is helpful for learning ML concepts in the wild, using the physical environment for an interactive and playful experience. This concept can be applied to any AR system that requires autonomous behavior agents in physical space, such as interactive education, gaming and health systems.

In future work, we plan to extend Paper Cubes into a more dynamic platform for storytelling, creativity and learning in AR. We are planning to conduct qualitative user studies to test the effectiveness of the system with children. We will test engagement and learning performance using Paper Cubes for learning ML concepts such as RNNs.

References

[1] Bishop, Christopher M. (2006). Pattern recognition and machine learning. *Springer, 2006.*

[2] Fusté, Anna, Amores, Judith & Google Creative Lab. (2017). Paper Cubes *Google AR Experiments. https://experiments.withgoogle.com/ar/paper-cubes*

[3] Ha, David & Douglas, Eck. (2017). A Neural Representation of Sketch Drawings. *arXiv preprint arXiv:1704.03477*

[4] Heess, Nicolas, et al. (2007). Emergence of locomotion behaviours in rich environments. *arXiv preprint arXiv:1707.02286*

[5] Lane, Dale. (2017) Introducing Machine Learning for kids. *http://dalelane.co.uk/blog/?p=3513*

[6] Norton, Andrew P., & Yanjun Qi. (2017). Adversarial-Playground: A visualization suite showing how adversarial examples fool deep learning. *Visualization for Cyber Security (VizSec). IEEE Symposium, 2017.*

[7] Smilkov, Daniel, et al. (2017). Direct-manipulation visualization of deep networks. *arXiv preprint arXiv:1708.03788*

[8] Støj, Use All Five, Creative Lab & PAIR teams at Google. (2017). Teachable Machine. *Google AI Experiments. https://experiments.withgoogle.com/ai/teachable-machine*

[9] Thalmann, Daniel. (2006). Populating virtual environments with crowds. In Proceedings of the 2006 ACM international conference on Virtual reality continuum and its applications (VRCIA '06). ACM, New York, NY, USA, 11-11. DOI=http://dx.doi.org/10.1145/1128923.1128925

[10] Williams, Ronald J., & Zipser, David. (1989). A learning algorithm for continually running fully recurrent neural networks. *Neural computation 1.2 (1989): 270-280.*