Demo video: https://youtu.be/sQC3y8dWLNQ

Press Release: USC Libraries Announce the Purchase, Preservation, and Investigation of Possible Unpublished Lewis Carroll Manuscript Pages

Pages seemingly handwritten and signed by Carroll depict alternate scenes from Alice tales… or are they the intended direction of the original storybook?

USC Libraries officials have announced that they were the anonymous buyers of unpublished pages apparently authored by C.L. Dodgson, famed for publishing "Alice's Adventures in Wonderland" and "Through the Looking Glass" under the pen name Lewis Carroll.

The pages were purchased at a London auction on March 19 through a third party, exciting speculation and curiosity throughout the world of Lewis Carroll scholars and enthusiasts. The university kept the purchase under tight wraps in order to surprise USC alumnus Dr. George Cassady and his wife Linda. Dr. Cassady founded the G. Edward Cassady M.D. and Margaret Elizabeth Cassady R.N., Lewis Carroll Collection in honor of his parents. Dean Catherine Quinlan will announce the acquisition at Doheny Library, during the 12th annual Wonderland Award ceremony the evening of April 21.

"These pages hold tremendous value for scholars of Carroll's work. While our experts continue to pore over the pages, we invite the public to scrutinize them as well. If they are valid, the discovery that he had actually written additional pages that were never included in his original storybook is incredible… and controversial!" said Catherine Quinlan, dean of the USC Libraries. "I'm proud that they have become part of our Cassady Carroll Collection and that we're able to make them accessible to researchers here and around the world for this crowdsourced investigation through our USC Digital Library."

"What makes these particular pages important are the similarity to Dodgson's original handwritten pages and the text of Alice's adventures," said Abby Saunders, curator of the Cassady Lewis Carroll Collection at USC. "Scanning them over, I could not remember if these scenes had actually taken place. I had to reread the original texts again and again, and I am still not quite sure… Local scholars have been flown in to examine the works, but they seem to be similarly baffled."

"These alternate pages are so important, just to get inside the head of Lewis Carroll and find out how he might have rewritten the original Alice in Wonderland books," said two-time Wonderland Award winner Andrew Woodham. "If I had known about them, my previous entries would have been totally different!"

Now that they're in the hands of the USC Libraries, the pages will be digitized and then preserved in a climate-controlled secure vault. In order for the public to examine the authenticity of the pages,

preservation specialists will carefully scan them for the USC Digital Library, and the originals will be available to the public and independent researchers – contrary to the policy of many private libraries.

Saunders led the initiative to purchase the letter, using funds donated by Wallis Annenberg that had been earmarked "for something special." And these new pages might be especially revolutionary discoveries.

Alice Through the Neural Network: Remembering a Childhood Story via Excerpts from a Computational Retelling

Artist's Statement for the 2016 USC Wonderland Award

Artwork includes:
(6) Original documents possibly handwritten and signed by Lewis Carroll
(1) Press release from USC Libraries regarding the documents' acquisition
(1) Artist's statement

Lastly, she pictured to herself how this same little sister of hers would, in the after-time, be herself a grown woman; and how she would keep, through all her riper years, the simple and loving heart of her childhood: and how she would gather about her other little children, and make THEIR eyes bright and eager with many a strange tale, perhaps even with the dream of Wonderland of long ago: and how she would feel with all their simple sorrows, and find a pleasure in all their simple joys, remembering her own child-life, and the happy summer days.
THE END

Alice in Wonderland, at its heart, is a story that we recollect from our childhoods. Coincidentally, the adventuresome tale ends with visions of a mature, adult future, recounting wistful memories of a dream of adventures past. But, as with any particular piece of literature of our childhood, do we remember the precise details?

Alice Through the Neural Network presents a "retelling" of Alice in Wonderland driven by an artificial intelligence computer system. It draws upon our nostalgia for a narrative long past, one that we cannot quite remember, but one that we imagine possessed many more fantastical adventures than we ever encountered in the story book (just as the story is retold through countless Alice-related spin-offs, parodies, and remixes). At its core, this artwork draws from a computational language model to recreate the story of *Alice in Wonderland*… just a bit differently.

This retelling of *Alice in Wonderland* emerges from a recurrent neural network, a very recent and advanced development in computational machine learning. Machine learning is a complex statistical process to model and predict data. An ordinary regression can make predictions fairly accurately, but today, most neural network approaches surpass older traditional statistical applications (especially with

noisy or complex data). In the past few years, machine learning techniques – and neural networks in particular – have allowed a computer to beat top-tier Go players,[1] have generated original artwork only imaginable in our dreams,[2] and even have helped improved predictions for medical applications (as researchers at our own University of Southern California have experimented with).[3]

The magic of a recurrent neural network is how it learns. For example, in a language-based model, it will look at the appearance of words, letters, and punctuation near each other. By looking at these relationships, and having a "memory" of what it analyzed before, a recurrent neural network can step through the basics of communication (spaces, spelling, punctuation, etc.) and then remember these first principles as it moves on to more advanced linguistic structures (the shapes of words, names, sentences, etc.). These processes result in a multi-dimensional feature space that allows the model to make statistical predictions about what words, letters, punctuation, and grammar come after each other. The recurrent neural network of *Alice Through the Neural Network* is trained as a character-level language model[4]: instead of finding and generating words, it looks at the relative likelihood of one letter or symbol appearing after another. For instance, given "w-h-a-", the model will probably output "t" to spell "what," but there is a small chance it will spell "wham" (and this depends on how the model is structured as well as how often wham appears compared to what in the original data.

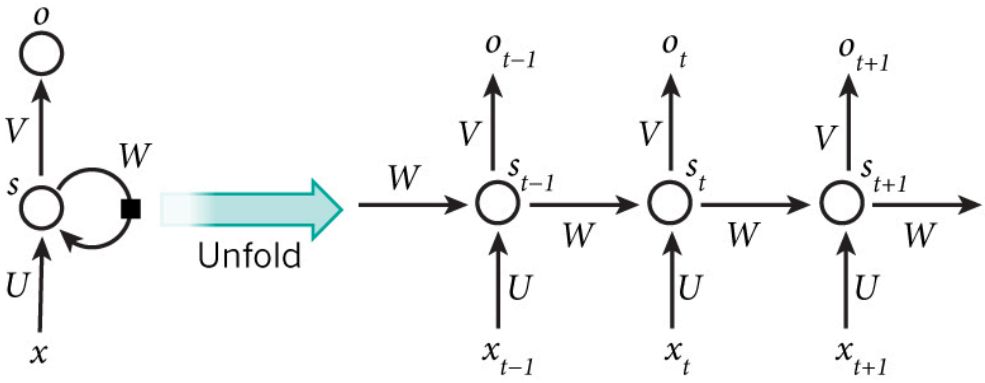

A neural network model represented over time,
with nodes having states s *at times* t*, with inputs* x *and outputs* o.[5]

Generally, what this means is that the recurrent neural network process reproduces common words and sentences, but it also introduces the ability to rearrange the text and even create new terms and lexicons! Coincidentally, the text produced by a recurrent neural network sometimes is not far off from the kind of queasy but quizzical narrative style that we already see from Lewis Carroll's most popular works, such as *Jabberwocky* or *The Hunting of the Snark*. If Carroll were alive today, given his fascination with mathematics and creative puzzles, perhaps he would be experimenting with neural networks, since they

---

[1] Yu, H. (2016). "AlphaGo's Success Shows the Human Advantage Is Eroding Fast." New York Times: Room for Debate. http://www.nytimes.com/roomfordebate/2016/03/09/does-alphago-mean-artificial-intelligence-is-the-real-deal/alphagos-success-shows-the-human-advantage-is-eroding-fast

[2] Mordvintsev, A., et al. (2015). "Inceptionism: Going Deeper into Neural Networks." Google Research Blog. http://googleresearch.blogspot.com/2015/06/inceptionism-going-deeper-into-neural.html

[3] Lipton, Z., & Elkan, C. (2016). "The Neural Network That Remembers." IEEE Spectrum. http://spectrum.ieee.org/computing/software/the-neural-network-that-remembers

[4] The neural network revolves through 59,069 iterations of 20 instances of the text of *Alice in Wonderland*. This process took around 16 hours to complete.

[5] LeCun, Y., Bengio, Y., & Hinton, G. (2015). "Deep Learning." Nature, 521. http://www.nature.com/nature/journal/v521/n7553/full/nature14539.html

seem to continue to surprise the humans developing them (even Google engineers have been surprised at what these machine learning models can figure out[6]). What is most fascinating about applying a recurrent neural network to Lewis Carroll's narrative is that it will find patterns in the original text that humans might have never foreseen, so that the model will produce a textual document that triggers familiarity and nostalgia yet fabricate something just beyond what humans might have imagined on their own.[7] In other words, the neural network not only learns how to spell words or create sentences without prior knowledge but it also embodies a type of computational creativity that can be tweaked through code.

To note, the approach that I took to recurrent neural networks employed in *Alice Through the Neutral Network* is actually a "wrong" way to do it. In an ordinary machine learning application, you are supposed to train on a large dataset of distinct textual objects. Large datasets allow better – ie., more varied and more accurate – prediction. However, Lewis Carroll's works are limited, meaning there is only a small amount of original data to work with. I turned this into an advantage though: I used multiple copies of Carroll's major works as the dataset. What this unique, "wrong" approach results in is, by repeating the same story (*Alice in Wonderland*) over and over in the base linguistic corpus document (known as the "training set"), that the neural network will reinforce particular linguistic structures that repeat (what is known as "overfitting"). In this case, it sees repetition as I copy-and-paste the narrative over and over, and it learns those repetitions.[8] From one perspective, one could argue these bits of text are actually statistically-likely versions that Carroll could have written himself. Ordinarily, the neural network would reproduce an entirely new narrative based on the calculated likelihoods. The system still does this, but now it reproduces a similar enough narrative to the original while maintaining slight differences. What this means is that the resulting new text looks familiar, but it remains slightly different, essentially resulting in a product that makes us question what we believe to be what we remember of the original stories. Coincidentally, this experience would mirror – like in the introductory quote – Alice's own experience as an adult retelling her memories of Wonderland.

The specific processes behind the creation of the neural network used for this artwork are also somewhat artistic and playful. While the computational model relies on rigid mathematics, I had to tweak a number of parameters – over the course of about 3 weeks – to settle on a model that generated the text in a way that I felt achieved my goals of creating a "very similar but slightly different" textual document. I used a text dataset of multiple copies of Carroll's two major works, comprising 8,858,426 characters in 8.9MB of raw text. The model – crafted from an open-source Python implementation of Andrej Kaparthy's char-rnn code[9] – was trained across 177,168 iterations, taking approximately 60 straight hours to complete. It should also be noted that, in general, the generated text makes sense (ie., it is readable English with rare grammatical, punctuation, or spelling errors), but the generated product is frequently not clever enough to be useful as convincing examples of near-Carrollian alternatives. After tweaking the model until it generated good-enough textual products, I then ran the final model over and over again and searched the generated text for the best artifacts that fit my goal of narrative cohesion. Theoretically, I could run the model infinite times at the click of a button, and it would generate new combinations of characters – and

---

[6] Lumb, D. (2013). "How Google's "Deep Learning" Is Outsmarting Its Human Employees," Fast Company. http://www.fastcompany.com/3022314/how-googles-deep-learning-is-outsmarting-its-human-employees

[7] For example, the brilliance behind Google's AlphaGo neural network system was that, in playing a top-ranked human player, it continually surprised professional commentators with its unique and unexpected moves. It played Go in a way that was beyond human familiarity.

[8] The reason to repeat the text of the story in the training set is to increase the amount of data that the neural network is able to learn from. However, it produces a beneficial effect that helped shape the theme of this artwork, in that it reproduces some parts of the original text while changing other parts of it slightly, creating small disjunctures in the text and therefore our memory of the original.

[9] See https://github.com/yusuketomoto/chainer-char-rnn and https://github.com/karpathy/char-rnn

therefore new text similar to the Alice narratives – every single time, but the goal was to find chunks of generated text that were triggered in just the right way that a reader might *not* question its legitimacy.

The unique nature of this artwork is of course the computational model, a media form in which no prior winners of the Wonderland Award contest have yet existed. I found it interesting that a lot of works contained in the Cassady Lewis Carroll Collection are reworkings of Carroll's *Alice in Wonderland* stories. Usually these reproductions are merely retelling the original narrative in other ways, rather than rewriting the narrative as if it were itself the original, which helped me hone in on reinterpreting Carroll's work by generating text trained on his writings. However, I also wanted to connect the computational element back to the physical element of "the archive." In visiting the Cassady Lewis Carroll Collection at the USC Libraries Special Collections in USC's Doheny Library, I was able to encounter some of the written notes and other textual documents written by Dodgson's own hand. Since the time of my undergraduate degree in English literature, I have been fascinated with how texts get made by authors beyond the text, such as editing, binding, note-taking, and interactions with audiences (and these fascinations with media put me on the trajectory of media and social science which I am currently finishing my PhD in). So, while the computational model created the seemingly-real narrative clips that generate the retold stories in this artwork, I also "retold" parts of the archive itself by fashioning documents that seem like they could have been found or placed in an archive. I printed out particular bits of text generated by the neural network in a computer font that was modeled off of Dodgson's own handwriting (used in "Alice's Adventures Under Ground," his original handwritten version). I then reworked the paper (actually soaked in tea for antique coloration!) to make them seem like lost documents that had been unearthed, now to be treasured in the USC archive. The aim was to further force those who encounter these pieces of "discovered" narrative to question their authenticity but also their own memories of Carroll's stories. I also played around with this idea of recreating the archive itself by copying an older press release from USC's libraries about other original Carroll documents, changing a few words to generate an additional fake meta-narrative (almost in the same way the neural network has generated the text for this artwork) that provokes people to take a close look at the pages.

Alice Through the Neural Network opens up our adult skepticism to childhood nostalgia by offering up a playful world of un-and-re-doing the narrative structure of *Alice in Wonderland*. It also recreates Carroll's fantastic world of both Wonderland's fantasy and his daily explorations of mathematics and complex puzzles through a novel computational approach that few have played with in prior competitions. Thank for you considering this submission.