

## Fairy Tales and Machine Learning: Retelling, Reflecting, Repeating, Recreating

Anna Ridler & Georgia Ward Dyer

*This text was included in the 2017 ARC magazine.*

In the introduction to his retelling of the Grimms' Children's and Household Tales, Philip Pullman writes 'fairy tales don't come whole and unaltered from the minds of individual writers'. The archetype emerges through countless retellings across cultures and across time; recent research analysing folk tales from Europe and Asia date the origins of some stories to thousands of years ago and the oldest – The Smith and the Devil – to the Bronze Age. These stories are told and retold, written and rewritten: Cupid and Psyche, transformed into Beauty and the Beast, transformed into Angela Carter's, The Courtship of Mr Lyon, or The Tiger's Bride. Helen Oyeyemi observes that 'when you retell a story, you're testing what in it is relevant to all times and places. Bits of it hold up, and bits of it crumble and then new perspectives come through'.

We explore this in our own versions of classic tales which are mediated through different machine learning tools; from image captioning to speech-to-text conversion. Through this mediation there often emerge striking and absurd associations between the image and the text. In Youtube & the Bass, our retelling of Beauty and the Beast, 'a person on a surfboard in a skate park' greets Beauty at the castle; Beast becomes 'a group of stuffed animals on top of a book'. The unusual and compelling chance phrases or images which surface inevitably lead us in particular directions; our process balances creative decisions with unexpected associations. We test the limits of the retelling, there are gaps in the narrative, and an important part of the meaning-making is intentionally seceded to the reader. Fairy tales after all are made up of small units of story which become building blocks that we instinctively know how to put together – the prince is always charming, the beautiful princess is always rescued - and this is possible even when the narrative is incomplete. This quality typical to fairy tales has motivated numerous scholars and literary theorists (Propp; Aarne and Thompson) to try to distill this to a fairy tale "formula" or a classification system.

Philip Pullman writes of the 'conventional, stock figures' which inhabit fairytales. While it is the recurrence of these stock figures which preserves the identity of one tale across several retellings, the surface detail of the figures in each retelling is particular to its temporal, cultural and social context. In the Nazi retelling of Little Red Riding Hood, she is saved from being eaten by a Semitic wolf by her hero in SS uniform. Our own retellings mediated through artificial intelligence tools also reflect context: that of the 'training set' that many of them were developed with. In supervised learning, a kind of machine intelligence, the program is developed by first training it on a labelled data set – for instance, in image recognition, the program is shown 'training' examples of images which have been labelled with details of what they depict. These training sets are compiled by researchers according to a variety of methodologies which

inevitably come to enshrine the cultural and social character of the content used. Take 'beauty' - when searched for on ImageNet, a popular training set used for image recognition programs - commonalities emerge across all of the results: being female, being white, being sexually provocative.

Multiples of stock figures frequently recur in fairy tales – the six sons and six daughters; the seven dwarves; the three sons set off one after the other on the same quest. The connection between repetition and remembering has been long established as part of the art of rhetoric - the earliest written reference to it dates back to 500 BC - and also for passing down stories and poems in the oral tradition. Repetition in fairy tales is often used as reinforcement or validation of a theme or element in the story. In resonance with this, image recognition programs have greater degrees of certainty in the labelling of an item if the item recognised appears multiple times in an image. For example, if an image depicts one baseball then the confidence score in the label 'baseball' would be high; however, if the image contains ten baseballs then the confidence score in the label 'baseball' would be much higher. Which is to say, that the more 'baseball-y' the image is, the more enthusiastically the program considers it a 'baseball'.

It's clear then that fairy tales are distinctly bound by certain rules – just not the ones of realism. As Pullman writes, 'realism cannot cope with the notion of multiples...[fairy tales] exist in another realm altogether, between the uncanny and the absurd'. It is perhaps this which motivated thinkers such as Freud and Jung to give psychoanalytic theories of fairy tales, likening them to dreams. Moreover, in a *mise en abyme*, the characters themselves often dream in ways that are central to the plot. In conversations we had with a research scientist at Google DeepMind, an analogy emerged between dreaming and how machine learning programs work. Our waking life experience equates to the machine learning program's 'training set'. When we dream, our brain uses this sensory data as the raw material from which to recreate a detailed and internally coherent world, just as the program takes from its training set to build up its own picture of the world and what it means. Although coherent, based on the original input, both the dream and the program are warped and imperfect as reflections of the real world; generating uncanny and absurd moments.

These moments are compelling – and it is only through their occurrence that we become aware of the program's picture of the world as an imperfect one. We do not wish to use these imperfections as criticisms of these tools – they are still developing. Those working on improving them include developmental psychologists and neuroscientists whose research on how humans learn is used to inform the design of programs that learn. These programs will undoubtedly improve their performance over time – a number of machine learning algorithms have successively outperformed predictions for their accuracy. We therefore have only a finite window of time to see their imperfections and glimpse their working. Repeating the processes that we used for our retellings now - four months on - returns different results, as the programs are consistently evaluating and adjusting their performance according to user feedback. As a program's performance becomes more streamlined, it converges on certain models of how the

world is and these models attain primacy over others. This is similar to how in the oral tradition storytellers would shape their tellings according to each particular audience. This fluidity does not exist in the printed tales which are now the most common way for us to experience them. The dominance of these printed versions has contributed to particular tellings gaining momentum as the authoritative ones.

Ultimately the raw material that both fairy tales and training sets draw from are products of people – these might have different temporal, cultural or social contexts but they are always human. Just as each retelling of a fairy tale gives us a reflection of its time, so the captions and labels given by machine learning tools give us a reflection of ourselves.