

AI Generated Chinese Kunqu Opera

錯山記



The Legend of Wrong Mountain



project

overview

The Legend of Wrong Mountain is a Chinese Kunqu Opera experience conceived, written and performed by artificial intelligence. Using machine learning technology to deconstruct and recreate the 400-year-old Chinese art form. Kunqu has developed established modes and patterns, which makes it especially suitable for neural networks to learn. By studying historical scripts, musical notations and traditional methods for creating Kunqu, we tweaked existing algorithms and devised new ones to conform to the traditional rules and norms as closely as possible. We presented this project as an actual woodblock printing script book and a video accompanied by audio.

I worked in an interdisciplinary team and took on flexible roles ranging from concept development, programming to post-production, exploring workflows of machine-generated artistic creation.

Tools

OpenCV, pix2pixHD, Detectron, OpenPose
Logic Pro, UTUA, Illustrator, Premiere Pro
Laser Cutter, Letterpress Printing



The project has been selected for paper oral presentation in **2018 NeurIPS Machine Learning for Creativity and Design Workshop**

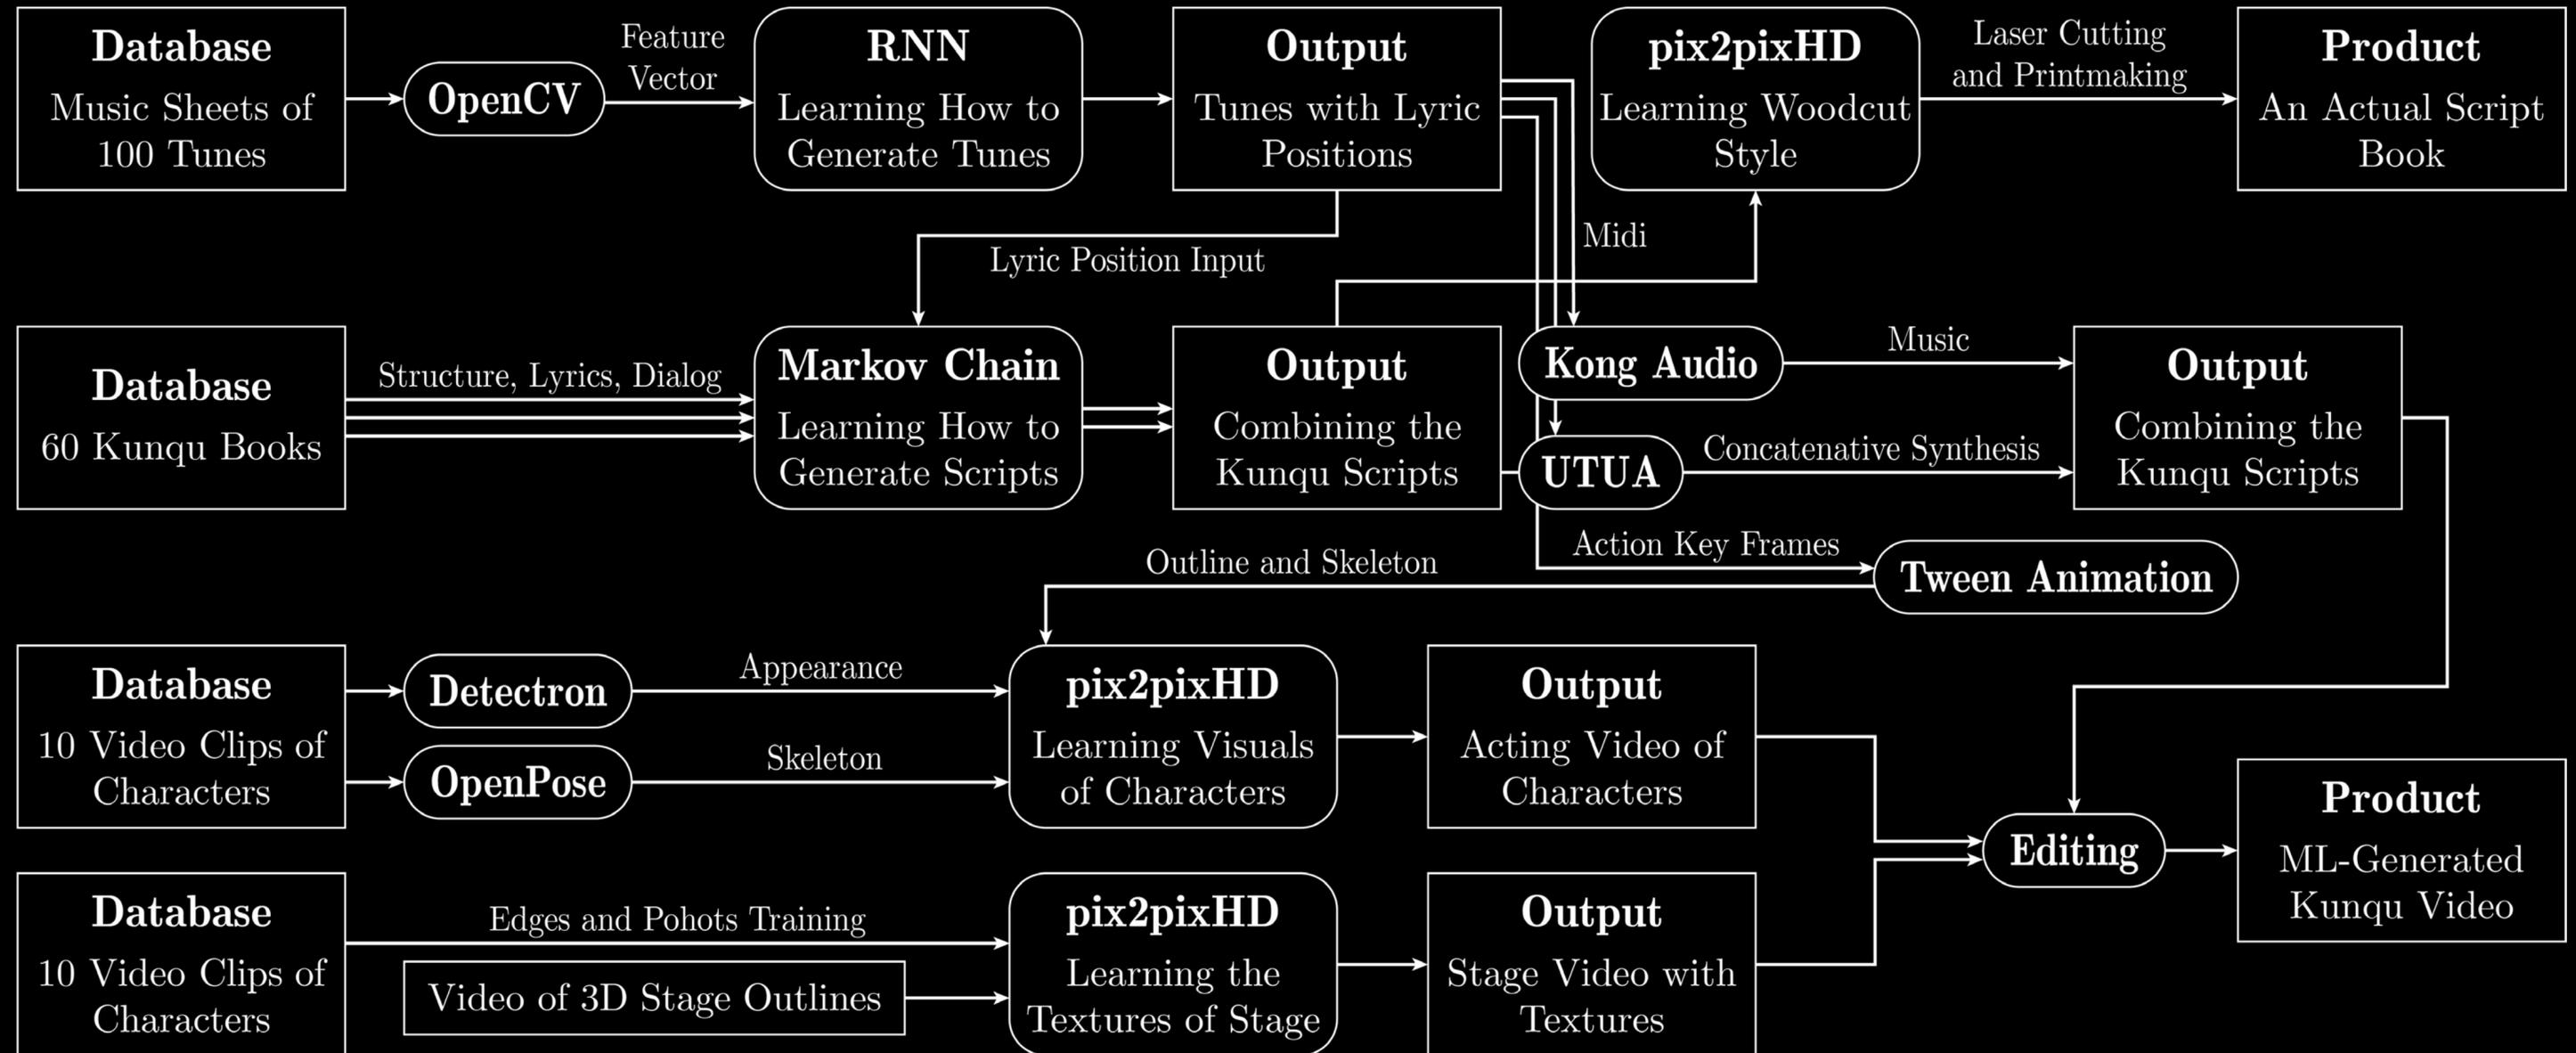


The Legend of Wrong Mountain

Trailer

<https://youtu.be/e1NICPznaD0>

workflow pipeline



Music Generation

The music of Kunqu are composed of melodies with established patterns. A hundred images containing traditional Kunqu sheet music were downloaded from Kunquwang (<http://www.52Kunqu.com>) using a scraper script. The sheets use Numbered Notation (Ziffersystem), a common format for notating Traditional Chinese music, which we wrote a new computer-vision based algorithm to parse. Then we used charRNN to generate our unique tunes. Instead of using the conventional method of encoding note-down-per-tick as a 128D vector, we concisely represented them with descriptive ASCII strings. For example, middle C is C₄. The pitch class information makes the structure easier for the learning model to grasp.

Script Generation

To learn the hierarchical structure of Kunqu scripts, we introduced a nested system with different Markov Chains at each level. The method is capable of producing outputs containing reasonable large-scale structure for chapters, dialogues, prompts and lyrics, as well as smooth sentences within the structure. For example, the sung parts of Kunqu were typically written in Qu, a form of classical Chinese poetry consisting of characters that conforms to tonal rules from melodies. Firstly, we reversed the chain for each sentence, such that the last, rhyming character is first determined, and all the previous characters are predicted backwards. Secondly, we inform the Markov Chain with tonal rules inferred from the generated melody. Thus, the generated poems not only were read smoothly but also rhyme and adhere to the traditional tonal rules.



Optical music recognition

```
IG_4:8+IA_4:8-IC_6:2+IA_5:8-IG_5:8-IE_5:8+IG_5:1-ID_5:16-ID_5:2-IC_5:2-IA_4:4-IG_4:4-IA_4:8+ID_5:4-ID_5:2-IE_5:2+IG_5:1-ID_5:2-IC_5:2-IA_4:4-IG_4:4-IA_4:8+I0:8-ID_5:4-ID_5:4-IE_5:4+IG_5:4-IA_5:1-IE_5:4-ID_5:2+IE_5:8-IE_5:4-IG_5:4-IG_5:4-IE_5:8-ID_5:8-ID_5:4+IE_5:4-IE_5:2-IE_5:2-ID_5:8-IC_5:8-IC_5:8+ID_5:2-ID_5:2-IE_5:2-IG_6:4-I0:2-IA_5:6-IG_5:1-IE_5:8+IG_5:1-ID_5:4-IC_5:8+IC_5:8+IE_5:1-IE_5:2-IE_5:8+IG_5:8-IA_5:8-IE_5:2-ID_5:2-ID_5:2-IC_5:2-IA_4:16-IA_4:16-IG_4:2-IE_4:8+IG_4:12+IC_5:8+ID_5:16-IE_5:4-IE_5:2-ID_5:2-IC_5:8+IC_8:1-ID_5:4+IE_5:4-I0:2-ID_5:2-IC_6:4-IA_5:2-IB_5:2-IG_5:4-IE_5:8-ID_5:2-IC_5:2+IE_5:8-ID_5:12-ID_5:4-IC_5:8-ID_5:8-IC_5:16-ID_5:8+IE_5:4-IA_4:8+IC_5:8+IA_5:1-IG_5:8-IA_5:16+IG_5:8-IE_5:4+IA_5:4-IG_5:2+IG_5:2-IC_5:4-I-IG_4:16-IC_5:8-IC_5:8+ID_5:8-IC_5:12+IC_5:8+IA_4:4+IA_4:4+I0:2-IG_4:2-IE_4:8-IG_4:8-IA_4:12+IG_4:4-IA_4:4+IA_4:4-IC_5:4-ID_5:8+IE_5:16-ID_5:8-ID_5:12-IG_5:2+IC_6:2-IE_5:4-IG_5:8+IG_5:
```

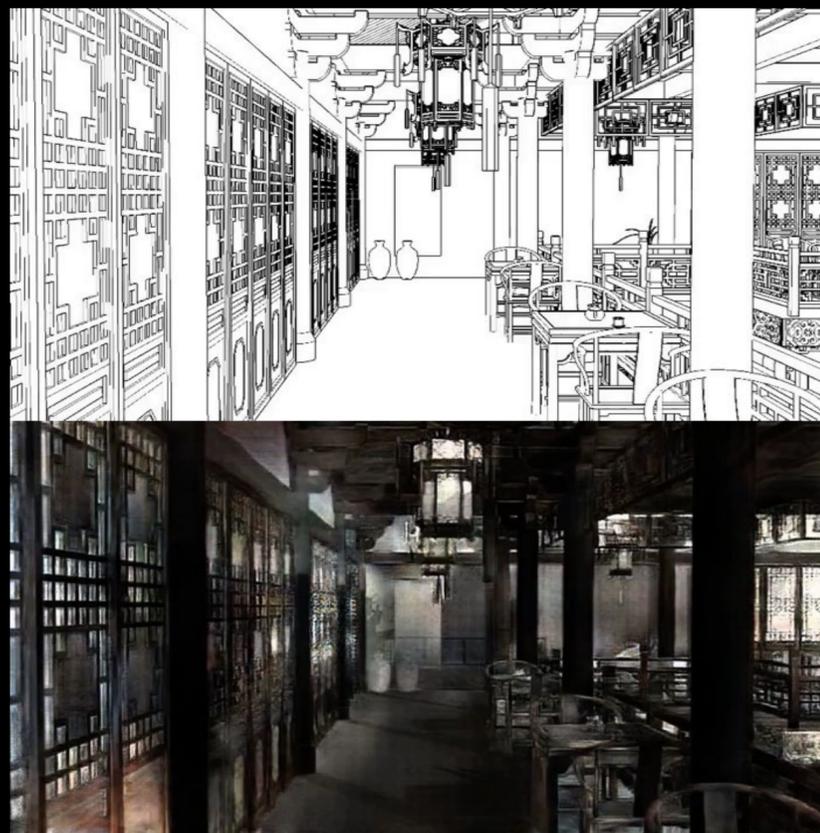
Sample melody represented by our notation

generation methods

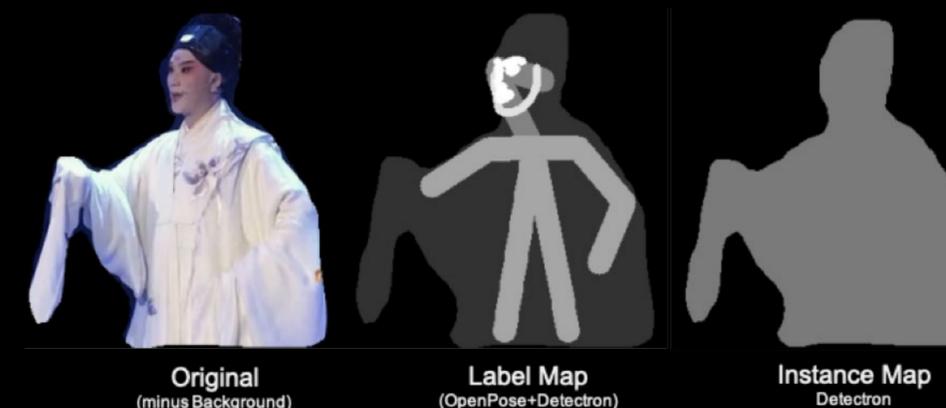
Visuals Generation

To generate graphics for the performers, we first used PIL to render new poses from scratch by implementing a simple pseudo-3D engine that contains skeleton system, physics and cloth simulation. Next, the procedurally generated poses were fed back into pix2pixHD to re-generate the photo-realistic visuals of the performers. As we had 9 pix2pixHD networks, each trained on a differently costumed character from the input videos, the network output 9 animated photographic performers that corresponds to the procedural skeletons and outlines. On the other hand, to generate the scene, We built a 3D model of the stage using SketchUp and wrote a Ruby script to control the camera. Simple outline drawings were exported as image sequences and sent to pix2PixHD. The neural network is trained using 80 photographs of historical opera houses against their extracted outlines produced by Canny edge. As the output,

pix2pixHD automatically textured our line drawings with the stylistic elements of the training set. At last, the generated images were combined into a video tour of the opera house at the start of the performance.



pix2pixHD result: scene



pix2pixHD training pipeline: character

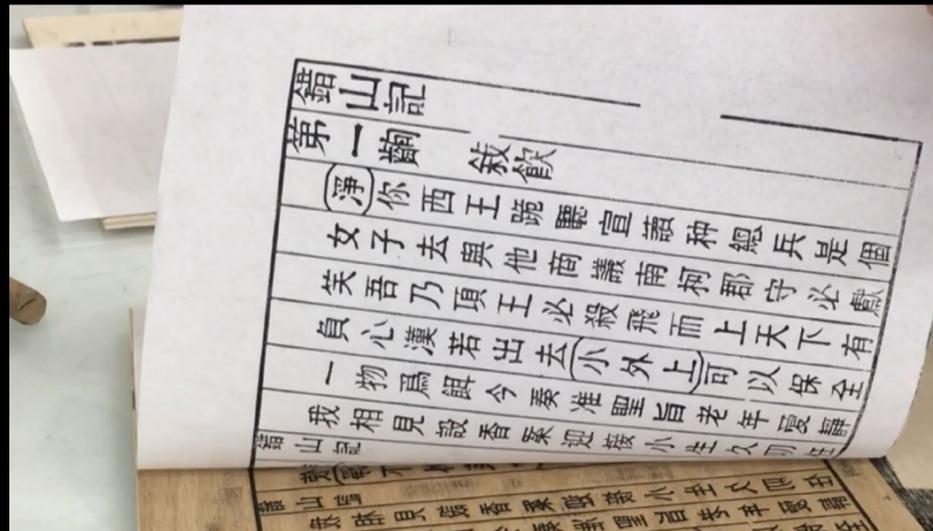
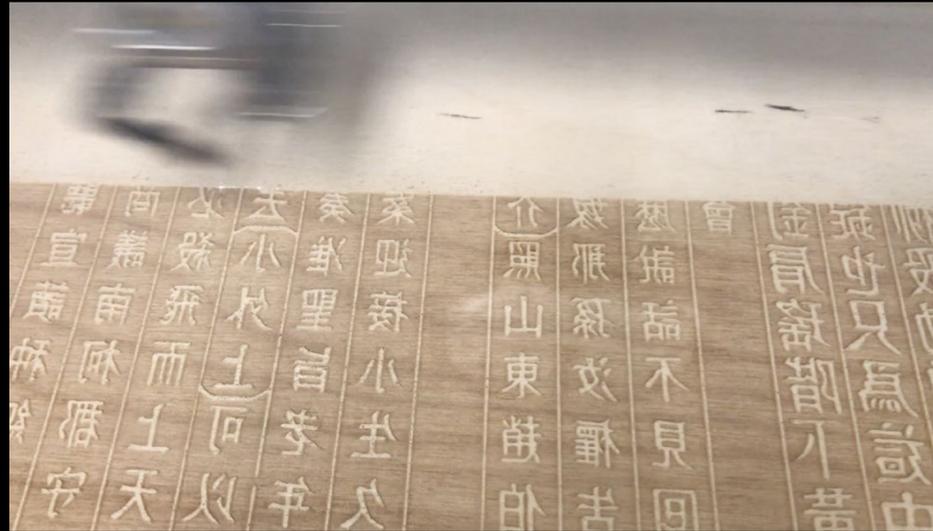


pix2pixHD result: character

generation methods

Book Generation

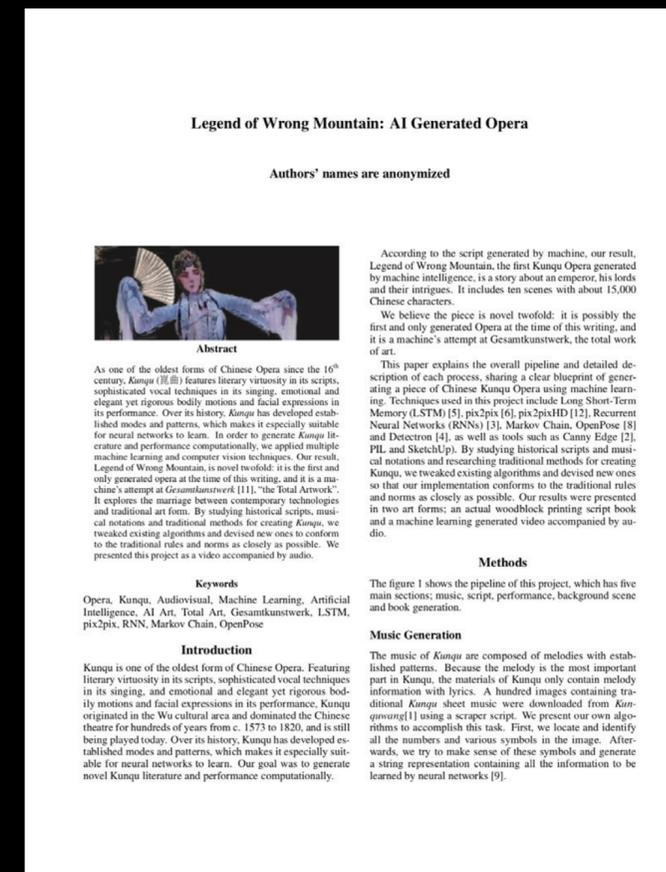
We printed and bound 8 physical books containing the script of the opera, in traditional woodcut press style. First, a pix2pix network is trained on images of Chinese characters and their woodcut counterparts, which is then used to render the characters in our script. A custom software is then written to typeset all the pages in our book, before they were sent to a laser cutter. The laser cutter prepared the printing blocks by engraving the pages onto 6mm wood pieces. Then, following traditional techniques, we applied black ink on the blocks and relief-printed onto rice paper. Finally, the books were bound with thread and needles.





Full Video with Subtitles

<https://youtu.be/QPdb28h4qjc>



Full Paper

<https://goo.gl/w9S89R>