



font-gen: Deep Models for Inferring Alternate Language Sets from Fonts

Garrick Fernandez (garrick@cs.stanford.edu)

CS230 Deep Learning, Aut 2019



Background

Fonts display a large amount of stylistic variance; as a result, the task of creating new ones is a **labor-intensive process** usually reserved for skilled artists and designers. In this paper, we examine methods of **generating characters for a language set from fonts which lack them**—methods which, with refinement, could serve as useful tools for designers looking to “internationalize” existing and in-progress fonts quickly and easily.

We present neural networks for three tasks: discrimination between in-font and not-in-font using a set of four basis characters (AHQJ), generation of all 26 Latin uppercase characters given basis characters, and generation of 46 Japanese *hiragana* given basis characters.

Data

We adapted a script [1] to produce **64x64 resolution, 8-bit grayscale images** of individual characters while preserving relative vertical alignments. The fonts were taken from the **Google Fonts** repo, which offers 2908 free and open-source fonts in various styles and weights. Of these, there are only **eight** font families that support Japanese.




Figure 1: Samples from generated datasets for each task.

Model

For the first task, we tried two “**tower**” architectures [2] one with densely connected layers in the towers, as well as one with convolutional layers. We also developed one “**shared**” architecture with convolution, to see if individual filters extracting similar features from all input characters could be useful to the network.

For the second, we took a **multitask** approach, where the most descriptive encoding from the first task is fed through shared layers generating all characters. We experimented with numbers of neurons, layers, and convolutional layers.

In the third, we attempted **transfer learning** by using the weights and network developed in the second task, and training it on the smaller set of Japanese fonts.

Experiment			MSE Loss		
#	description	parameters	train	val	test
1	1 extra conv layer	21.7M	1560.2068	2261.4195	-
2	no conv layer	21.7M	1902.4307	2645.9201	-
3	x2 neurons	43.4M	1380.6402	2182.6619	1696.2317
4	x2 decoder neurons, add layers	43.3M	1683.7724	2391.8688	-
5	x2 encoder neurons, add layers	22.1M	1598.5508	2324.1624	-
6	x2 neurons, add layers	43.6M	1432.9696	2223.8934	-

Figure 2: Architecture search in second task.

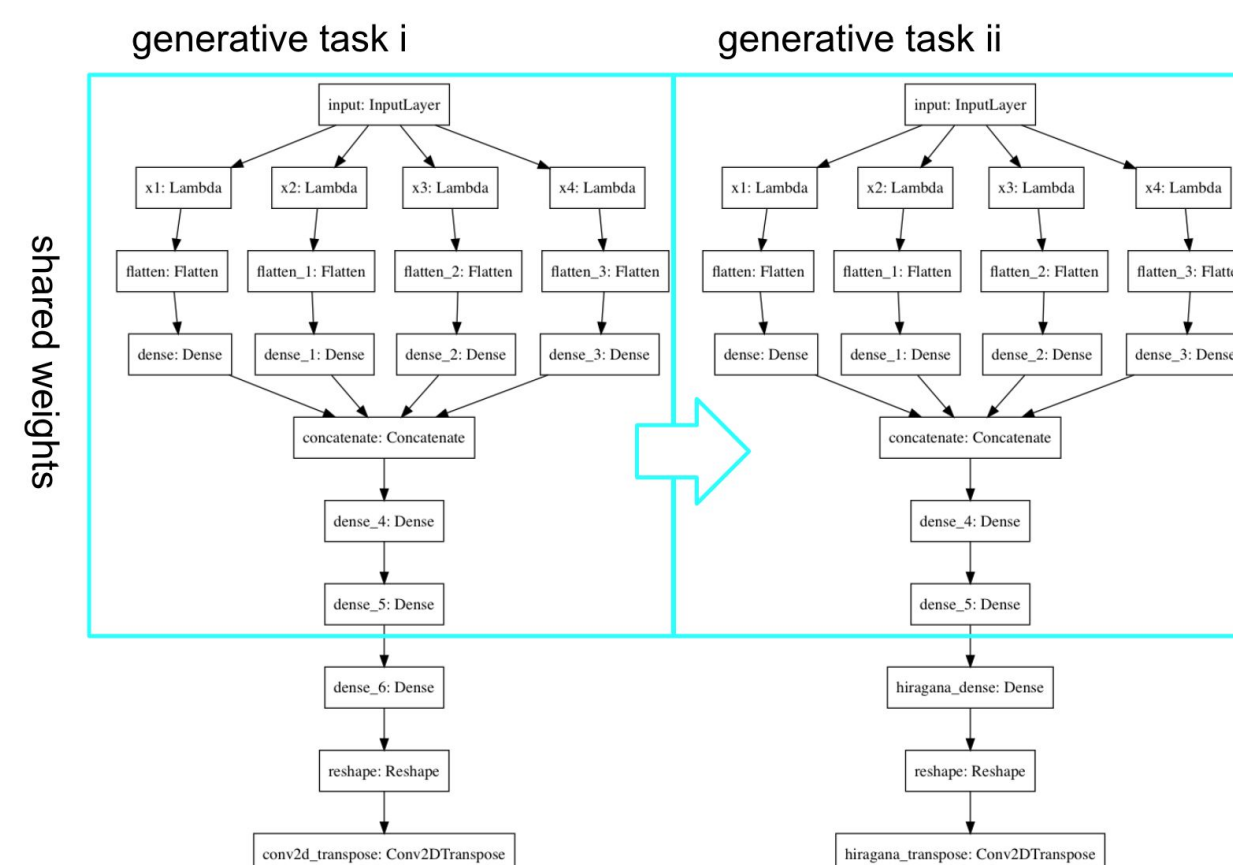


Figure 3: Architecture for third (transfer learning) task.

Results & Discussion

In the discrimination task, the tower architecture achieved **89.76% accuracy** on a test set of 332 examples, with erroneous classifications leaning towards **false negatives**.

In multitask learning, adding convolutional layers gets rid of grainy effects in generated letters. Adding **more neurons** results in an improvement in MSE loss, but adding more layers has less of a beneficial effect.

Transfer learning works well in the third task, with an improvement in training loss and validation set loss.

Predicting fonts’ Japanese sets worked well when the fonts were similar to ones we trained on. **Common features** can be seen between basis and output.

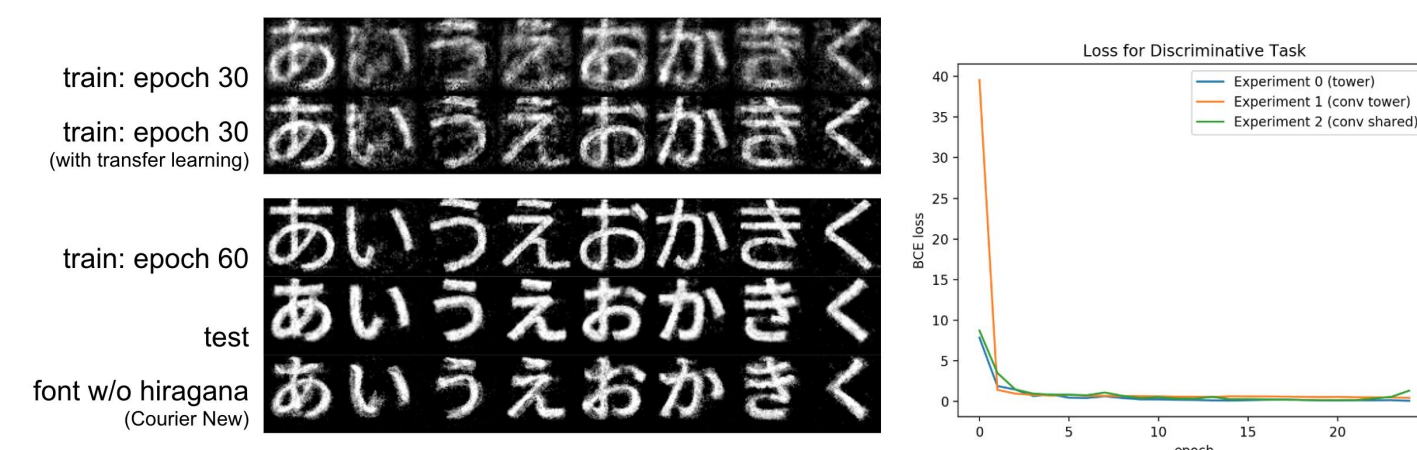



Figure 4: Training progress.



Figures 5-6: Convolution; predicting Japanese sets.

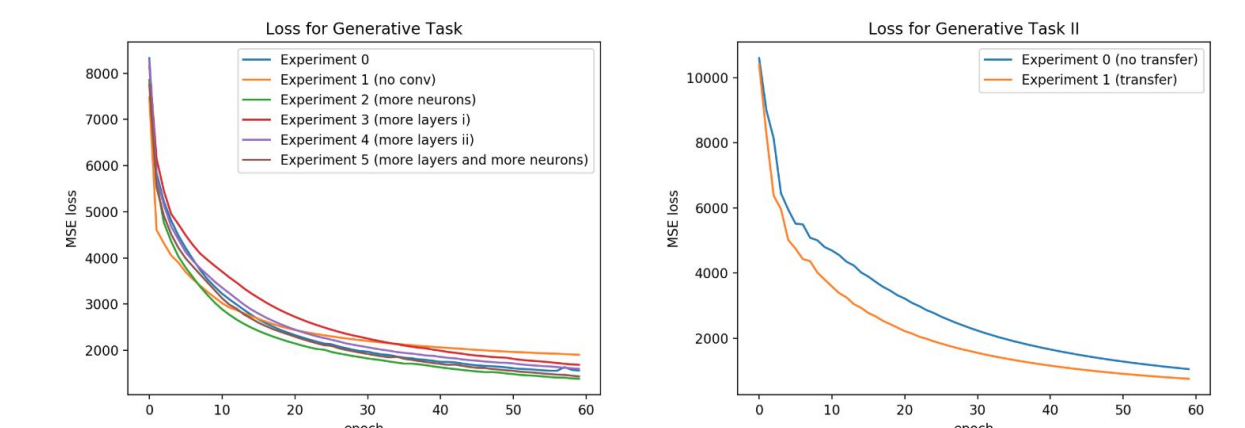


Figure 7: Loss plots for all tasks.

Conclusions

Limitations: similar fonts can have distinguishing marks, some fonts are not represented in Japanese (e.g., the monospaced-serif font Courier New or the more hand-drawn Comic Sans).

Interesting avenues for future work:

- **Generating more characters**, modified/larger networks, more fonts from more sources
- **VAEs or GANs**, **SVG decoders** for translating encoding into scale-invariant vector

Footnotes

[1] E. Bernhardsson, Analyzing 50k fonts using deep neural networks.

[2] S. Baluja, Learning typographic style: from discrimination to synthesis, arXiv preprint arXiv:1603.04000

Link to code: github.com/garrickf/font-gen

Special thanks to the CS230 course staff for inspiration and support!