Machine Knitting & Machine Learning

Lea Albaugh

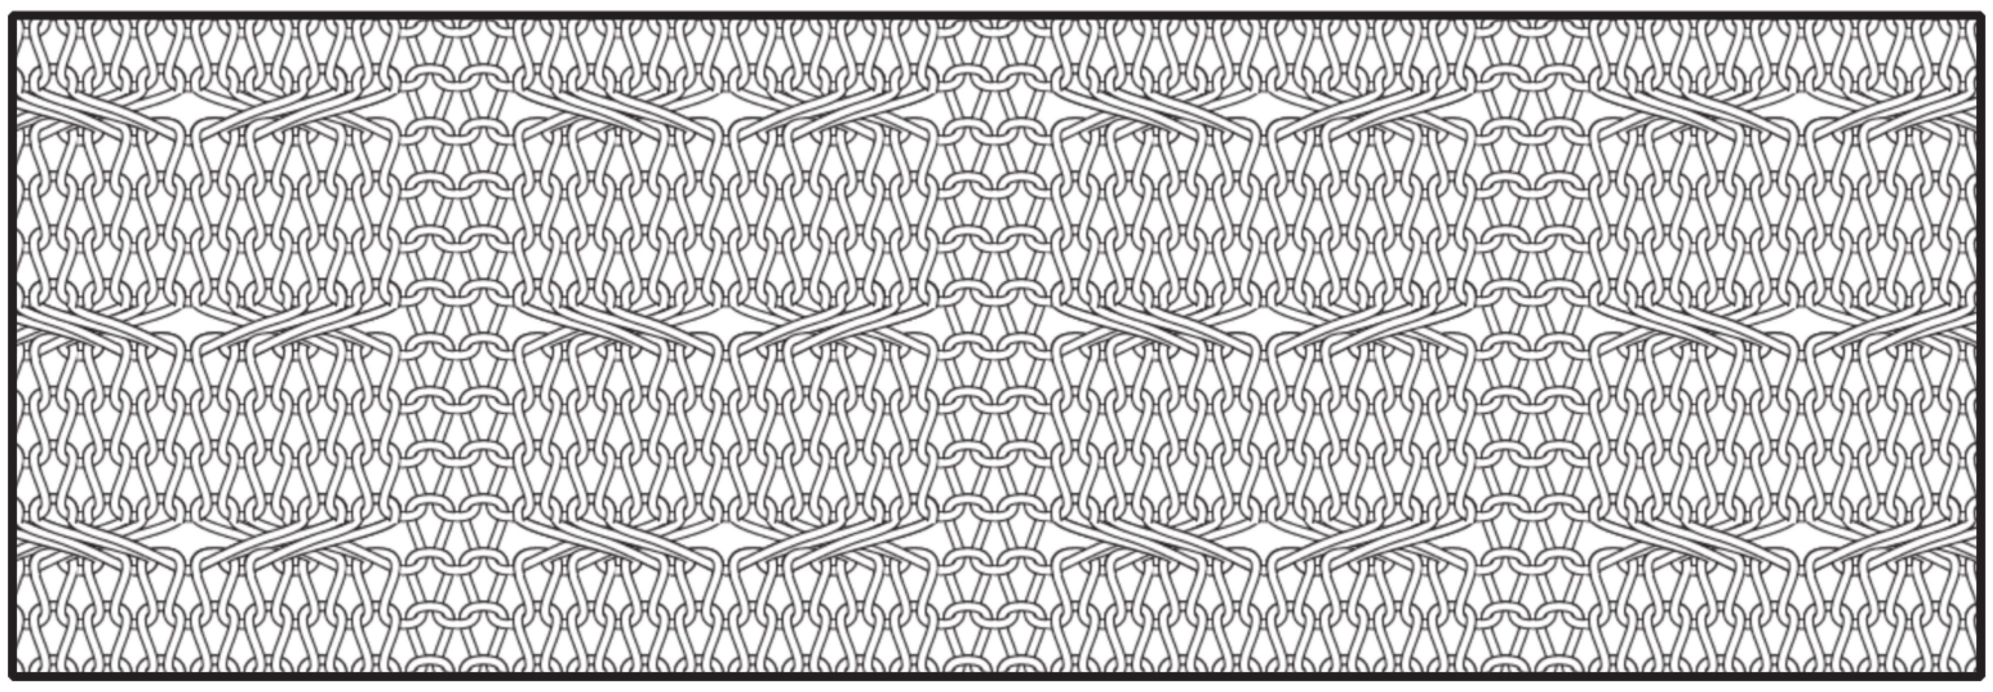

Data

As part of an ongoing research project, the CMU Textiles Lab[2] is compiling a dataset of knitting patterns that produce textured swatches: that is, flat rectangles of knitting with such stitch-level details as ribbing, lace eyelets, and cables.

The Textiles Lab's research code parses these textual patterns into a graph-based representation, with each stitch as a node that is connected to its row-wise neighbors and its column-wise parents and children. For machine knitting, each stitch is then allocated to a particular position and time of construction: in other words, the graph is laid out as a planar grid.

For texture swatches, the planar graph can be characterized with just a few pieces of information per stitch: the **direction** the stitch was pulled through its parent ("knit" vs "purl"), the **offset** between its construction location and its child's construction location, and, for stitches that overlap or get transposed, a **z-index** that describes which stitch passed on top of the other.

These three pieces of information can conveniently be packed into the three color channels of an RGB image.

Background

Knitting patterns are often expressed in "knitspeak," which is a textual representation of the operations required to produce a particular knit structure. Janelle Shane's recent machine learning art project, Skyknit[1], uses a set of knitspeak patterns as the input to a text-generating RNN. The results approximate knitspeak, though of course very few are exactly valid knitting instructions – hand knitters working with these patterns must use the knitting skill of interpretation to produce knit structures from them.

In this project, we take a different approach: we encode the input patterns into a representation that ensures the knittability of the output.

Initial GAN Results

The knit texture dataset is unfortunately fairly small as machine learning datasets go. While we augmented the data using horizontal and vertical flips and color inversion, the dataset still seemed too small to get good results from a GAN approach.

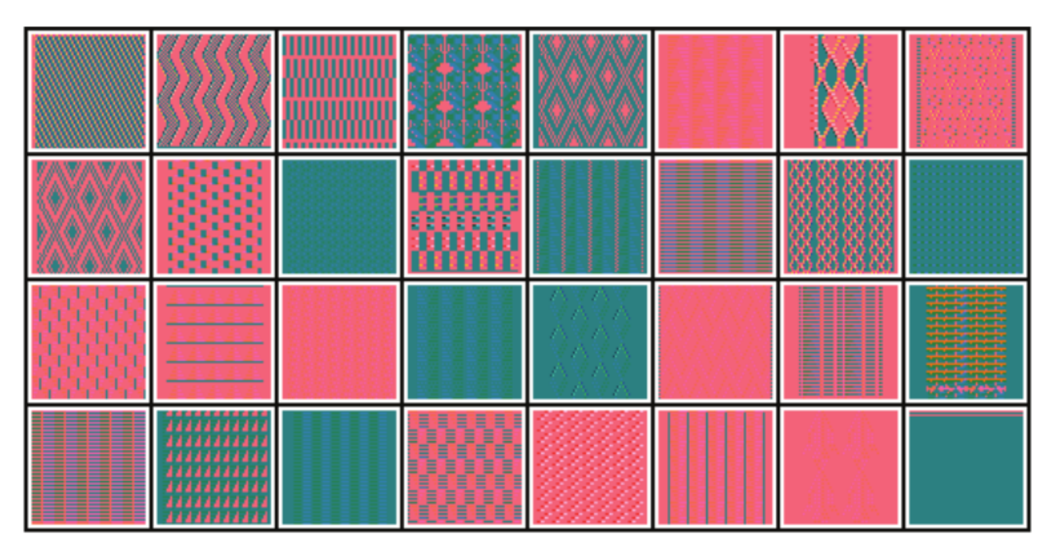

32 sample "real" input images

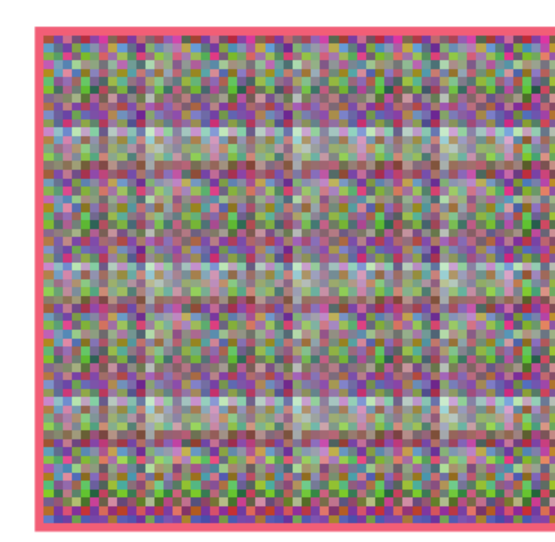

representative output

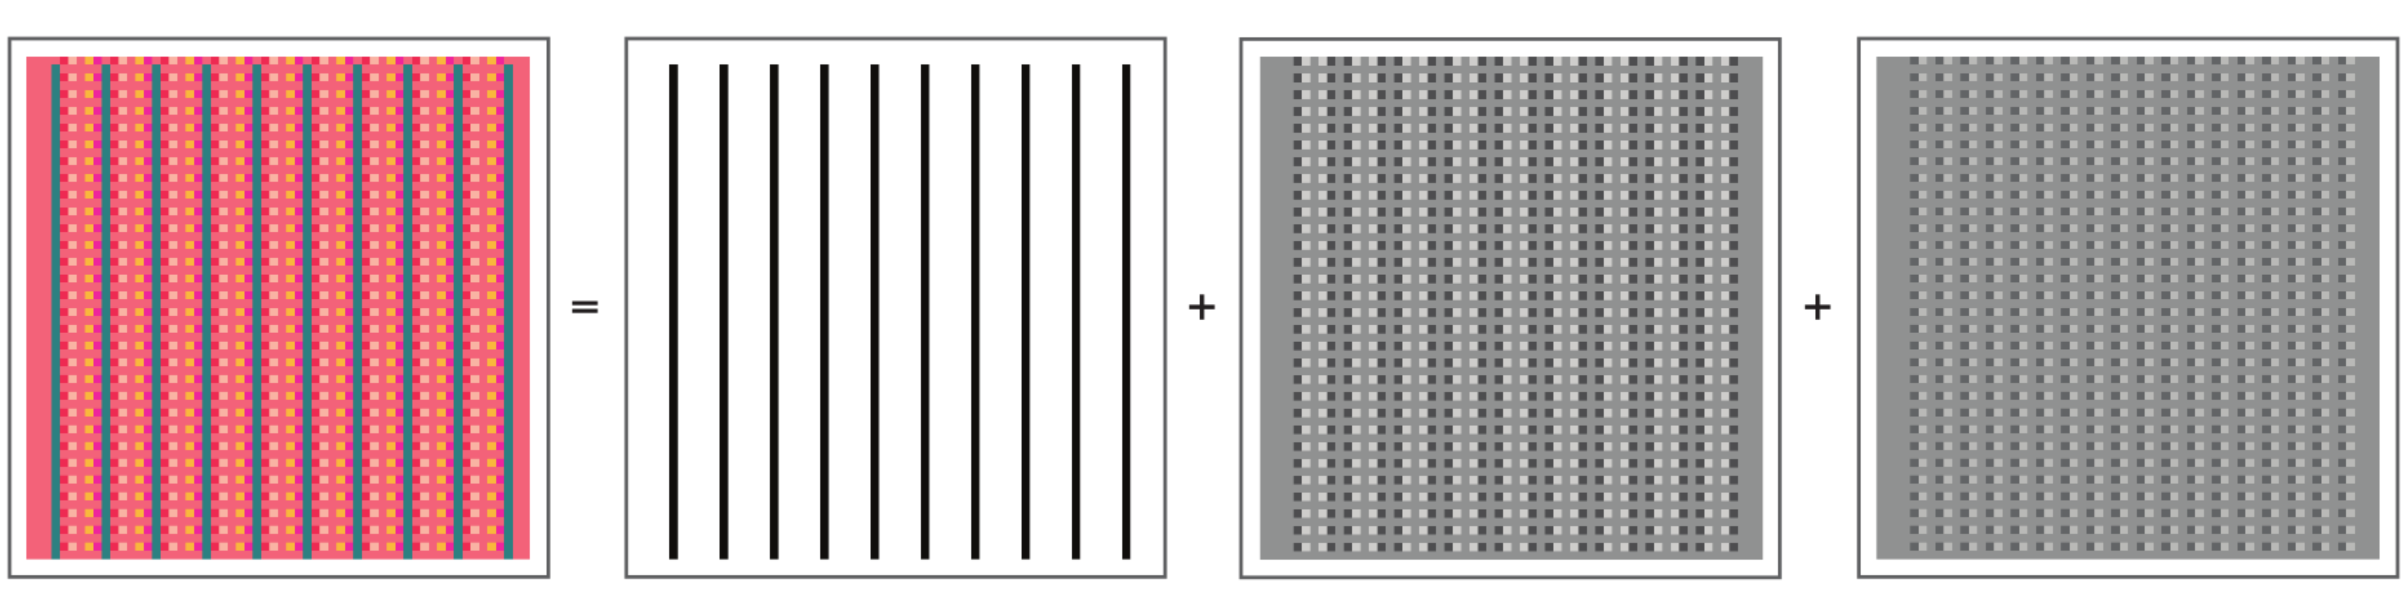

direction (knit vs purl) | z-index | offset

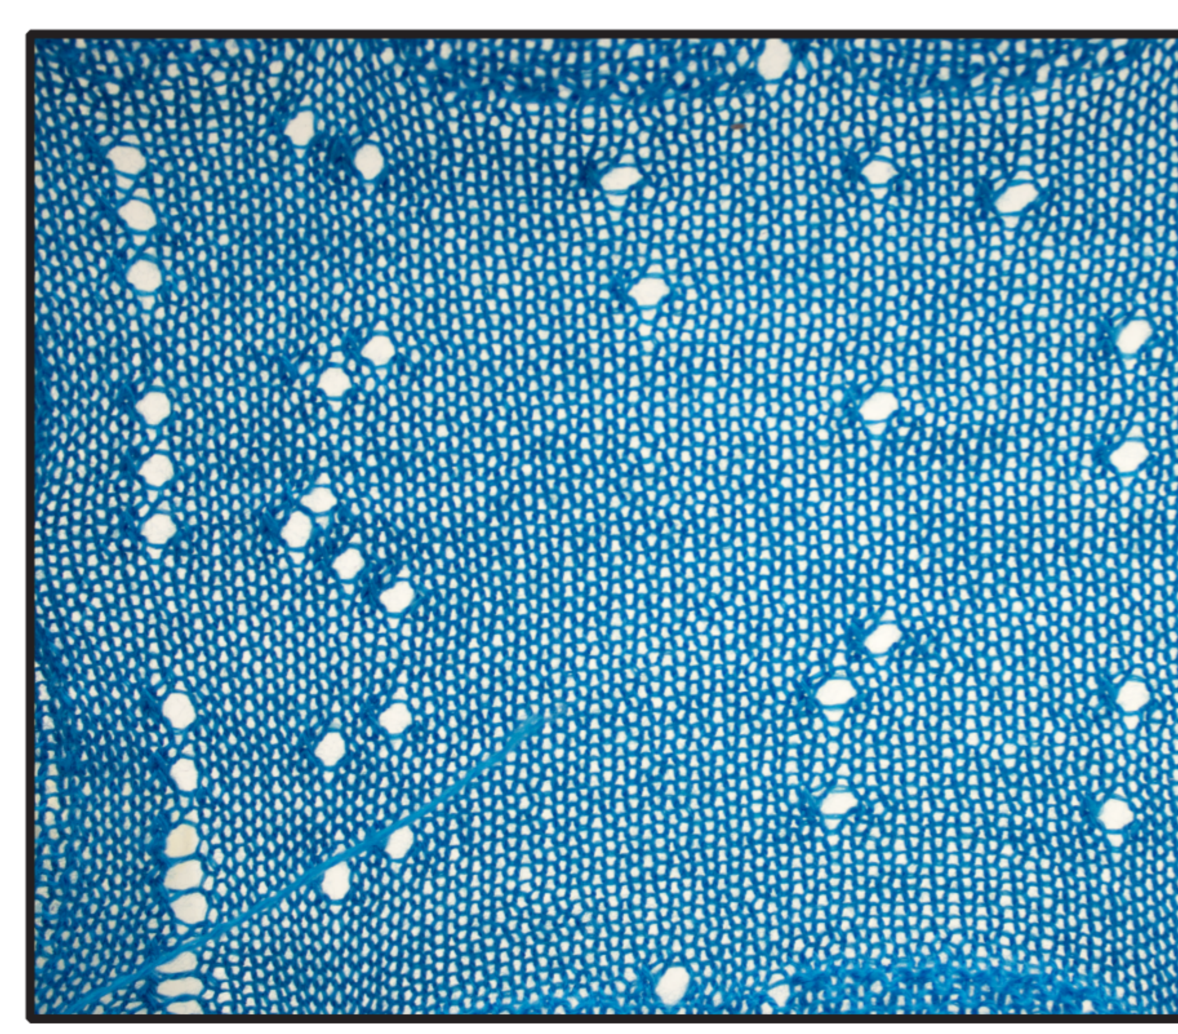

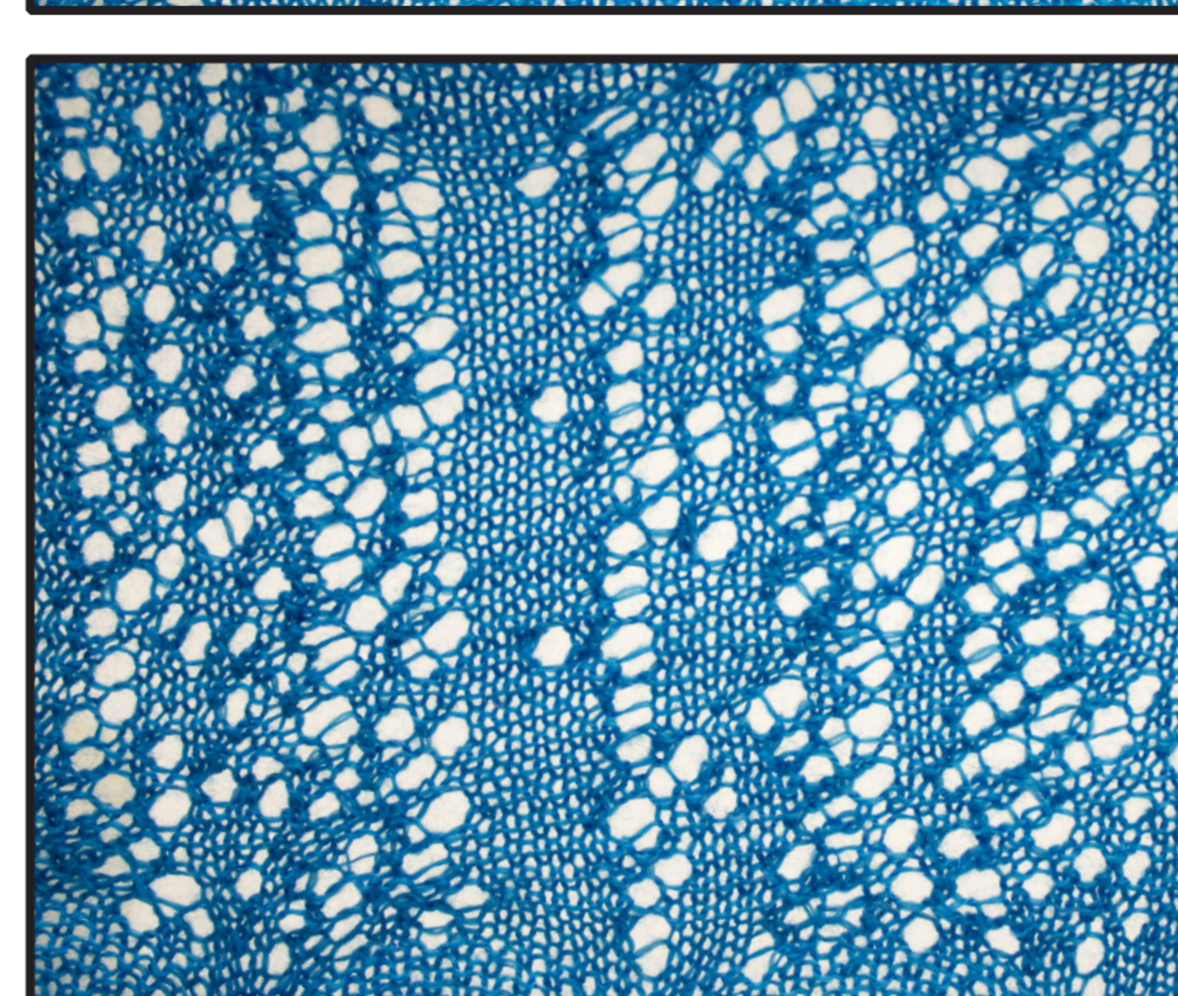

Further Results With a VAE

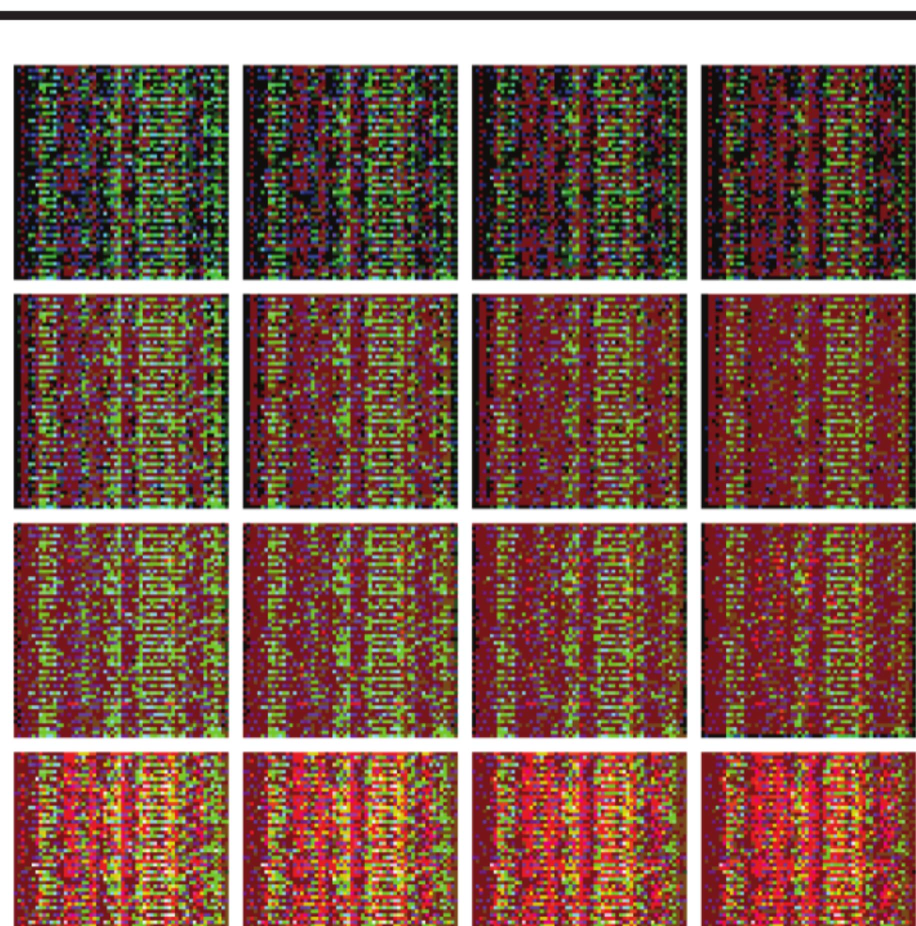

A variational autoencoder approach was tried next, with more promising results. For this iteration, the data was introduced into the network as its literal value, instead of interpolating over a 255-value color channel range.

However, while the results did capture the columnar style associated with many of the patterns in the database, the available latent space that didn't simply degenerate into all-eyelets or plain knitting was very small. Additionally, despite the simplified input data, a fair amount of deliberate thresholding was required to get good knit results.

1. http://aiweirdness.com/post/173096796277/skyknit-when-knitters-teamed-up-with-a-neural
2. https://textiles-lab.github.io/