



Estimating Yarn Length for Machine-Knitted Structures

Gabrielle Ohlson
gohlson@andrew.cmu.edu
Carnegie Mellon University Robotics
Institute
Pittsburgh, PA, USA

Angelica M. Bonilla Fominaya
abonilla@andrew.cmu.edu
Carnegie Mellon University
Pittsburgh, PA, USA

Kavya Puthuveetil*
Jenny Wang*
kavyap@andrew.cmu.edu
jennyw2@andrew.cmu.edu
Carnegie Mellon University, Robotics
Institute
Pittsburgh, PA, USA

Emily Amspoker
eamspoke@andrew.cmu.edu
Carnegie Mellon University
Pittsburgh, PA, USA

James McCann
jmccann@cs.cmu.edu
Carnegie Mellon University, Robotics
Institute
Pittsburgh, PA, USA

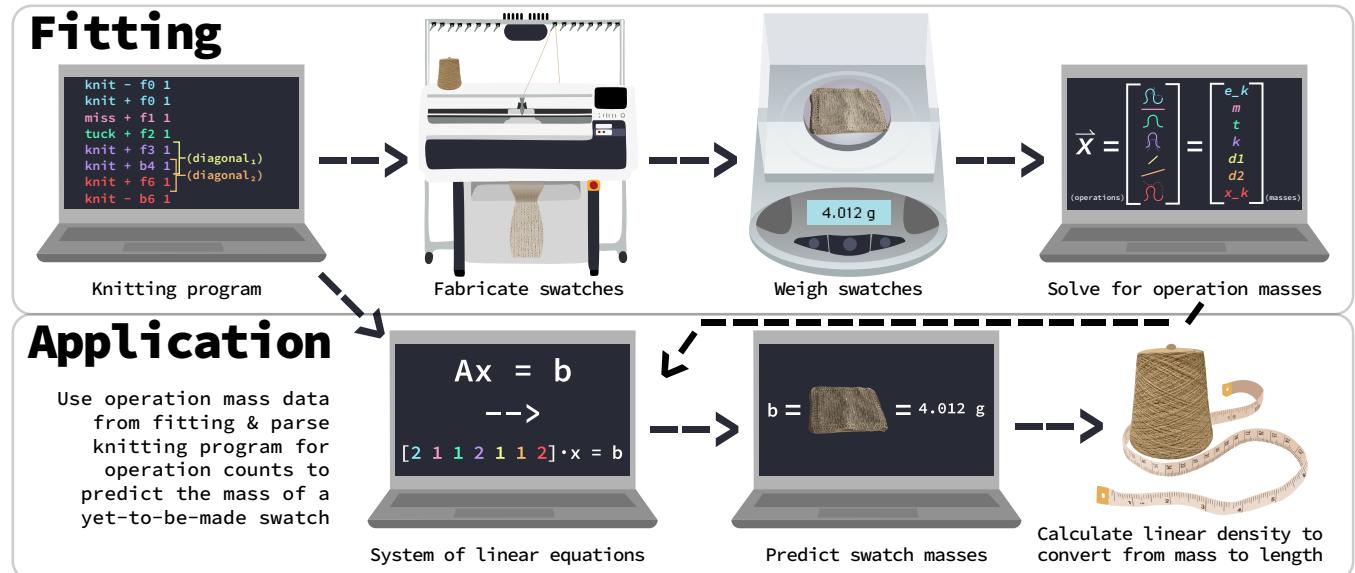


Figure 1: Workflow for estimating yarn length using our method

ABSTRACT

We show that a linear model is sufficient to accurately estimate the quantity of yarn that goes into a knitted item produced on an automated knitting machine. Knitted fabrics are complex structures, yet their diverse properties arise from the arrangement of a small number of discrete, additive operations. One can estimate the masses of each of these basic yarn additions using linear regression and, in

turn, use these masses to estimate the overall quantity (and local distribution) of yarn within any knitted fabric. Our proposed linear model achieves low error on a range of fabrics and generalizes to different yarns and stitch sizes. This paves the way for applications where having a known yarn distribution is important for accuracy (e.g., simulation) or cost estimation (e.g., design).

*Both authors contributed equally to this research.



This work is licensed under a Creative Commons Attribution International 4.0 License.

CCS CONCEPTS

- Applied computing → Computer-aided manufacturing;
- Computing methodologies → Linear algebra algorithms.

KEYWORDS

machine knitting, linear regression, yarn estimation, automated fabrication

ACM Reference Format:

Gabrielle Ohlson, Angelica M. Bonilla Fominaya, Kavya Puthuveetil, Jenny Wang, Emily Amspoker, and James McCann. 2023. Estimating Yarn Length for Machine-Knitted Structures. In *Symposium on Computational Fabrication (SCF '23), October 08–10, 2023, New York City, NY, USA*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3623263.3623355>

1 INTRODUCTION

Knitted fabrics are intricate metamaterials with structural properties that differ greatly from the bulk material in their unprocessed yarn form. Knitting can transform yarns with varying degrees of elasticity into complex, anisotropic structures with distinct material properties, determined by the fabric's "stitch pattern" — a strategic configuration of low-level yarn structures [Sanchez et al. 2023]. One can further customize knitted fabrics by altering parameters such as gauge, stitch size, and the yarn's fiber composition. All knitted fabrics are composed of a small number of basic operations, but using different combinations and arrangements of these structures can result in an infinite number of diverse fabrics.

Our model treats these complex structures by breaking them into a small number of discrete features ("basic operations") with consistent masses. These operations include the conventional yarn operations in machine knitting with a few additional versions to account for local effects between adjacent operations. The basic operation counts used by our model can be easily derived from the low-level knitout [McCann 2017] knitting-machine programming language. Generalization across varying linear densities and stitch sizes is achieved through normalization and curve fitting.

In this work, we contribute the following:

- (1) A set of basic operations (tuck, miss, knit, edge-knit, x-knit, and diagonal varieties) which are sufficient to describe the distribution of yarn in any machine-knit object (subsection 4.1).
- (2) A method for estimating the masses of these basic operations for a given knitting machine and yarn using linear regression on a small, carefully-chosen, set of knitted samples, which we refer to as the "basic set" (subsection 4.2, code on GitHub).
- (3) An experimental validation showing that these basic operations for estimation work well to predict the yarn used in a larger set of knits, referred to as the "validation set" (section 5).
- (4) Extensions to our method for efficient transfer of basic operation masses to different yarn types and stitch settings, requiring only to measure the yarn's linear density or manufacture one additional sample, respectively (section 6).

2 RELATED WORK

The complicated properties of knitted fabrics have motivated many researchers to study the physical behavior of knit objects. Knitting machines apply external forces to yarn during fabrication; accordingly, a study of machine knitting operations must consider the mechanical properties of yarn, such as bending modulus, during the formation of knit operations. Our work proposes a linear system to model these operations; however, prior research has demonstrated that there are nonlinearities to consider when deforming yarn into knitted structures.

2.1 Mechanics of Knitted Fabrics

Knitted materials deform nonlinearly under high stress. In order to model them, one needs an understanding of their material properties [Liu et al. 2017; Luo and Verpoest 2002; Mahadevan et al. 2022; Park et al. 2023; Sanchez et al. 2023; Yeoman et al. 2010]. However, under low stress, knitted fabrics exhibit linear stress-strain relationships that can be understood through their topology and stitch composition [Cachine et al. 2023; Kageyama et al. 1988; Yamada et al. 2003]. While knitted fabrics may *behave* non-linearly, we hypothesize that the *yarn distribution* of fabrics can be modeled linearly, since machine knitting is repeatable and involves only a handful of repeated, discrete operations.

2.2 Yarn Measurement and Modeling

Numerous geometric models have been proposed to predict the loop length of weft-knitted structures with parameters such as loop width, loop height, yarn thickness, and fabric thickness. However, many do not consider a loop's interactions with its neighbors [Morooka et al. 1998; Pavko-Cuden et al. 2013,?; Peirce 1947; Vékásy 1960] or are limited to a set of stitch patterns [Arbataitis et al. 2021; Ciukas and Sadauskas 2004]. It has been shown that fabric can be explicitly modeled as a tiling of optimized splines for Finite Element Analysis (FEA) simulations, but this method has not been applied to yarn measurement [Wadekar et al. 2020]. In a review of future directions for constitutive models of knitted fabrics, Gonzalez et al. highlight the importance of experimentally measuring the bending modulus, compression modulus, and length of yarn per stitch when determining how yarn elasticity and manufacturing processes individually affect fabric behavior [Gonzalez et al. 2022].

3 KNITTING BACKGROUND

Before we describe how to determine the distribution of yarn in a knit object, we review the basics of machine knitting.

Knitted fabrics have a variety of structures with distinct material properties. The way a given fabric appears, behaves, and responds to both internal and external stimuli is determined by the fabric's local geometry, which is in turn determined by the local topology of the "ribbon knots" that form the fabric [Markande and Matsumoto 2020].

Knit stitches are self-intersecting loops of yarn, created by pulling a new strand of yarn through a pre-existing loop (see Figure 2). The formation of a stitch is dependent on a loop of yarn existing below it (along the y -axis), and also the stable loops of yarn existing on either of its sides (along the x -axis). These columns of loops along the y -axis are referred to as "wales," and the rows of loops along the x -axis are known as "courses."

Without neighboring loops, a given length of yarn is not stable. The configuration of each knit stitch in relation to its neighbors also

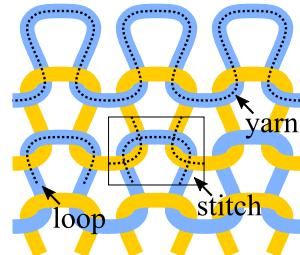


Figure 2: Knit Topology [Narayanan et al. 2019].

impacts the amount of yarn that goes into the overall fabric. For example, a loop that is deformed into a curved bell-shape requires more yarn to knit than a loop that is stretched horizontally into a nearly straight line. Similarly, the more a loop is stretched vertically, the more yarn it will require.

Knitting can be automated using either flat or circular machines (so-called because of how their needles are arranged). Our model targets industrial, flat, v-bed weft knitting machines (see subsection A.1). These are flexible machines used to manufacture knit garments, accessories, and fabrics.

4 METHOD

The core idea of our method is to model the distribution of yarn in a knitted object as a combination of seven basic operations, each of which has a consistent mass throughout the object. To obtain the total mass¹ of yarn in any knitted object, we sum the basic operation counts, multiplied by the amount of yarn in each operation type:

$$C_{\text{object}}^T \times W = w_{\text{object}} \quad (1)$$

Where $C_{\text{object}} \in \mathbb{R}^7$ is a column vector of operation counts, $W \in \mathbb{R}^7$ is a column vector of basic operation masses, and w_{object} is the final measured mass of the knit.

In this section, we first describe the basic operations used in our model, and then the set of example knits we used to estimate the masses with linear regression.

4.1 Basic Operations

We chose the basic operations used in our model to cover all the ways a knitting machine can add yarn to a knitted object, without any overlapping operations and without any operations that always co-occur. We started with the conventional knitting operations (knit, tuck, and miss). We then considered additional yarn formations that occur as a consequence of the locations of subsequent operations. For example, yarn can connect consecutive rows or cross between the needle beds, producing different lengths depending on the bed location and direction of the adjacent stitches. We attempted several categorizations of these formations to find a small set of comprehensive “basic operations” that achieves low error; for example, we explored distinguishing front knits from back knits, treating connecting lengths of yarn as separate from the surrounding stitches, and labeling operations based on their location in the fabric. Our proposed set contains the following basic operations (visualized in Figure 3):

- (1) **Knit:** a new loop is pulled through an existing loop, creating a stitch. We typically think of a *front knit* when looking at a fabric, though we can also create the inverse, a *back knit* (also known as a *purl*), by moving a loop to the other bed on the machine so that the new loop is pulled in from behind.
- (2) **Edge-knit:** knit operations along the edge of a row, including the loop. We consider edge-knits as a separate operation from knits because the carrier² turns around after these knits, which might add or remove a different amount of yarn than other knits.

¹We adopt the textiles industry convention of using mass to talk about amounts of yarn since mass is easy to measure with a scale and can be converted to length by dividing by linear density.

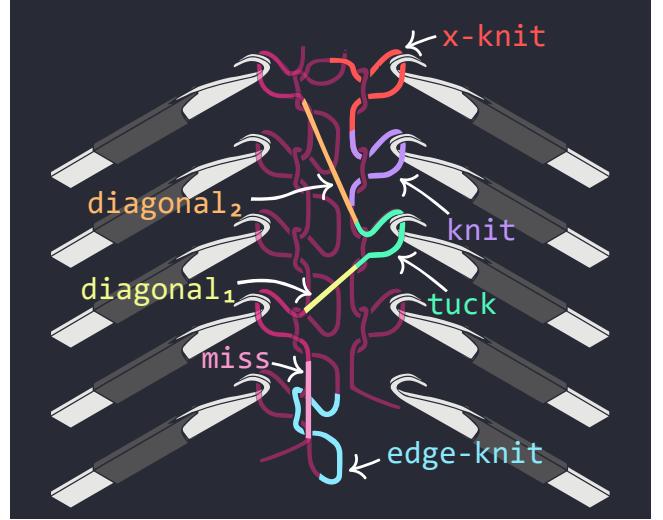


Figure 3: Basic operations sufficient to describe the yarn distribution in machine-knit objects.

- (3) **X-knit:** knits along the edge of rows when changing beds in a tubular knit. As with edge-knits, the carrier turns around, but there might also be additional mass from yarn connecting across the two beds.
- (4) **Tuck:** a new loop stacked onto an existing loop or empty needle. Unlike a knit stitch, creating a tuck is not dependent on a loop existing below, since yarn is not pulled through. Instead, the needle actuates to grab onto yarn and pulls down to deform the strand into a bell-shape.
- (5) **Miss:** the yarn carrier moves past an unactuated needle, causing the yarn to lay straight without forming a loop.
- (6) **Diagonal₁:** yarn connects two consecutive needles on *opposite* beds. Diagonal₁ occurs in patterns like rib. This operation is necessary to account for variations in the amount of yarn that goes into a set of neighboring loops, depending on their respective bed locations.
- (7) **Diagonal₂:** diagonals that skip over a needle when crossing the bed, thus adding more yarn than a plain diagonal. Diagonal₂ occurs in patterns like half-gauge rib. Additional variations on diagonal are necessary for some larger-gauge patterns, which we do not consider in this work.

Note that knitting machines also perform operations that do not add any yarn (e.g., transfers, racking, a-misses). We do not model these operations, because they do not directly contribute to the amount of yarn in a knit object.

4.2 Model Fitting

Here, we describe the computational method we use to estimate yarn quantity with linear regression. Given a set of knit fabric swatches and their masses, one can estimate the mass of each basic

²See A.1 for definitions of concepts related to automated v-bed knitting machines.

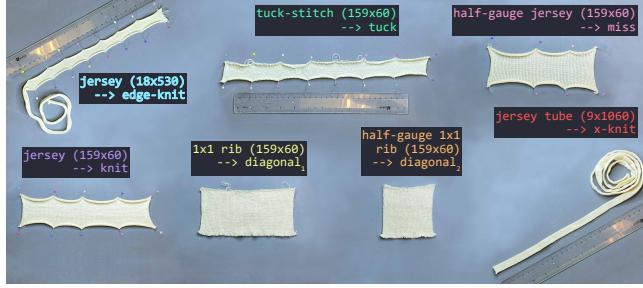


Figure 4: Photographs of the seven training swatches used in fitting our model, labeled with the operations they emphasize.

operation by solving a least-squares (i.e., linear regression) problem:

$$W \equiv \underset{x}{\operatorname{argmin}} (Ax - b)^2 \quad (2)$$

Where $A \in \mathbb{R}^{n \times 7}$ is a matrix with rows containing the basic operation counts for each of n samples; W is the column vector of per-operation masses (as per Equation 1); and $b \in \mathbb{R}^n$ is a column vector of the measured masses of the knit swatches.

Given that there are seven basic operations, it is sufficient to knit and weigh seven training fabric swatches. We designed seven linearly independent swatches that resemble a diagonal matrix to the extent that is possible based on knitability (i.e., all swatches must contain knits and edge-knits), and so that each swatch highlights a distinct basic operation.

We arrived at a set (Figure 4) consisting of a 159×60 swatch³ of 1x1 rib (alternating front and back knits to create diagonal_1 operations); a 159×60 swatch of half-gauge 1x1 rib (alternating front and back knits, skipping even needles to create diagonal_2 operations); an 18×530 swatch of jersey (all front knits; narrow to emphasize edge-knit); a 159×60 swatch of jersey (all front knits; more rectangular to emphasize knit); a 159×60 swatch of half-gauge jersey (all front knits, skipping even needles; emphasizes miss); a 159×60 swatch of tuck-stitch (a checkerboard of tucks and knits, with knits in the first and last wale for stability; emphasizes tuck); and a 9×1060 jersey tube (alternating rows of front and back knits forming a tube; narrow width to emphasize x-knit).

$$A = \begin{bmatrix} d_1 & d_2 & \text{edge-knit} & \text{knit} & \text{miss} & \text{tuck} & \text{x-knit} \\ \begin{matrix} 9480 & 0 & 118 & 9422 & 0 & 0 & 0 \\ 0 & 4740 & 118 & 4682 & 0 & 0 & 0 \\ 0 & 0 & 1058 & 8482 & 0 & 0 & 0 \\ 0 & 0 & 118 & 9422 & 0 & 0 & 0 \\ 0 & 0 & 118 & 4682 & 4740 & 0 & 0 \\ 0 & 0 & 118 & 4712 & 0 & 4710 & 0 \\ 0 & 0 & 0 & 14842 & 0 & 0 & 2119 \end{matrix} & \begin{matrix} (a) \\ (b) \\ (c) \\ (d) \\ (e) \\ (f) \\ (g) \end{matrix} \end{bmatrix}$$

Matrix of operation counts for basic swatch set:

(a) 1x1 rib (159×60), (b) half-gauge 1x1 rib (159×60), (c) jersey

³Swatch dimensions are given in needles by rows. Thus, a 159-needle-wide half-gauge swatch will only contain 80 columns of stitches, because half-gauge skips alternate needles.

(18×530), (d) jersey (159×60), (e) half-gauge jersey (159×60), (f) tuck-stitch (159×60), (g) jersey tube (9×1060)

5 RESULTS

We tested our method by producing the basic sample set described above and a large collection of validation swatches.

The code for each swatch was written in knitout [McCann 2017], which we generated in Python with some helper functions and the knitout-frontend-py [Narayanan and McCann 2019] library. Operation counts were computed with a helper program. We used a Shima Seiki SWG-091N2 15 gauge industrial knitting machine to knit the swatches. Unless otherwise mentioned, samples were knit in Tamm Petite Medium Camel (T4221) 2/30 nm acrylic yarn; and with the machine-default stitch value⁴ of 30 and leading value of 25. Yarn was supplied via one of the machine’s i-DESCS+DTC automatic yarn feeders set to “medium” tension control (7g of tension). We weighed our swatches using a shielded lab scale (Four E’s Scientific), with a maximum mass rating of 500g and claimed an error of less than 0.001g. We knitted swatches in batches, separated by waste yarn and draw threads. This was both more convenient when knitting many samples (since the machine did not need to be re-started for every swatch) and avoided any measurement errors from cast-on or bind-off patterns.

Table 1: Basic operation masses as fit from 21 swatches – three instances each of the seven basic swatches described in subsection 4.2. Note that the mass used for each swatch was the average of three weighings.

Operation	d_1	d_2	edge-knit	knit	miss	tuck	x-knit
Mass (mg)	0.10	0.17	0.49	0.39	0.14	0.41	0.84

5.1 Operation Masses

To reduce variance in our training owing to potential variability in yarn density, machine operation, and scale behavior, we chose to knit each training swatch three times and used the average of three weighings as the mass of each swatch.

Performing a linear regression resulted in the estimated basic operation masses reported in Table 1. These masses seem intuitively reasonable – knit (0.39mg) and tuck (0.41mg) use about the same amount of yarn; edge-knit (0.49mg) and x-knit (0.84mg) add more yarn because they cross between rows and between beds, respectively; and the diagonals and miss (0.10mg, 0.17mg, and 0.14mg) use less yarn because they don’t form loops. Of course, the true test of these masses is in their predictive power, which we explore next.

5.2 Validation

We produced a validation set of 55 swatches that vary in dimensions, gauge, and pattern (subsection A.2, subsection A.3). This swatch set includes samples ranging in size from 150×20 to 400×520 needles by rows, knitted in jersey, garter, rib, tuck, and seed patterns, as

⁴Stitch and leading values are in opaque, machine-dependent units that influence the shape of the needle path when making stitches. The stitch value controls how much yarn is pulled into the stitch, while the leading value controls how the already-held loop is deformed when making the stitch, and is chosen to promote stable knitting.

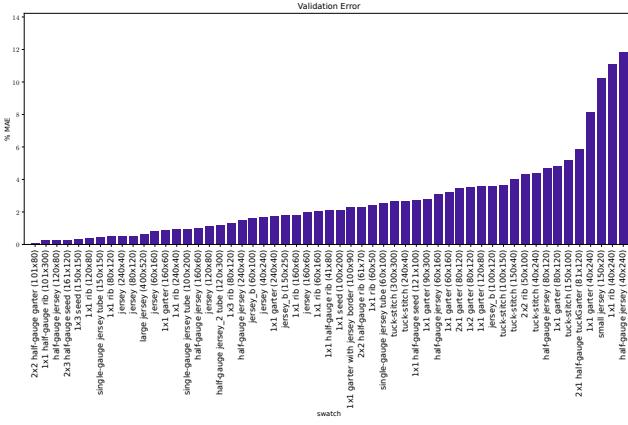


Figure 5: Difference between the predicted and measured masses of swatches in the validation set, as a percentage of measured mass.

well as some combinations of multiple patterns. We produced both single-bed and tubular versions of these samples where possible.

We used the operation masses computed from our training swatches to predict the masses of these swatches (Figure 5). The model achieves 5% or less error in all but four swatches (which have approximately 10% error).

As an additional check, we trained our model on just the validation set as well as the validation set along with our basic swatch set. The results for these sets were similar to the training swatches, with the largest deviation being about 20% difference in x-knit. (The validation set gives – in the same order as Table 1 – 0.10mg, 0.18mg, 0.14mg, 0.39mg, 0.14mg, 0.38mg, and 1.10mg for the operation masses; while training on both sets together gives 0.10mg, 0.18mg; 0.25mg; 0.39mg; 0.14mg; 0.38mg, and 0.86mg.)

6 EXTENSIONS

We were interested in understanding how our model generalized to different stitch values and yarn types and how we might expect it to perform on different machines. This section explores these extensions.

6.1 Curve Fitting for Varied Stitch Sizes

All swatches in our initial training and validation sets were knit with a stitch size of 30 (opaque machine units⁵). However, many different stitch sizes can be used within a single knit item, so it would be useful to be able to predict yarn usage for different stitch values.

Intuitively, the tuck, knit, edge-knit, and x-knit operations are all created by pulling yarn with a needle through a distance dictated by stitch size, so they should vary linearly with stitch size (at least with an ideal yarn); while the remaining operations – miss, diagonal₁, and diagonal₂, relate to the space between needles, so should not change.

To test this idea, we knitted three instances of each of the seven training swatches at stitch sizes (25, 30, 33, 34, 35, 40) and fit operation masses independently for each stitch size (Figure 6).

There are likely boundary points and nonlinearities to consider in stitch size fitting functions – the size of a loop is limited by the needle dimensions, so some stitch sizes are “un-knitable,” and the curve might behave increasingly nonlinearly as it approaches these limits. And, of course, our test yarns are not ideal, zero-bending-resistance, inextensible, zero-radius yarns. But within the regime of stitch sizes that knit without errors, it appears that response to stitch size changes is close to linear.

Table 2: Percent MAE when Generalizing to Stitch Size with Linear Fit Functions

Stitch Size	All dependent features	Only knit	Jersey (159 × 60) sample
25	3.758 ± 3.434%	4.377 ± 3.124%	3.963 ± 2.595%
33	3.154 ± 1.642%	2.561 ± 1.693%	2.058 ± 1.051%
34	6.576 ± 4.560%	6.312 ± 4.192%	5.420 ± 3.273%
35	5.088 ± 3.703%	4.262 ± 3.842%	3.895 ± 3.058%
40	5.809 ± 5.985%	5.566 ± 5.548%	5.837 ± 6.486%

6.1.1 Results. In Table 2, we compare the results when the model is trained using the full data set for a given stitch size, when the model is trained on the data for our main stitch size (30) as we interpolate the coefficients of the stitch-size-dependent features based on their respective fitting function, and when we instead curve fit to a single sample to interpolate these varying features. The curve fits are visualized in Figure 6.

We find that the estimated functions fit the knit and tuck data best, with an R^2 value of 0.928 for the linear function fitted to knit data and 0.816 for tuck. The polynomials were nearly linear and had only marginally improved R^2 values of 0.962 and 0.834, respectively. The estimations for edge-knit and x-knit contain more outliers that did not increase with stitch size, but we are able to find fitting functions for these features when including only inliers.

The linear fit function is formulated as follows: $f(x, s) = x(a \cdot s + b)$ where x is the feature mass for the main stitch size, and s is the step size from the main stitch size.

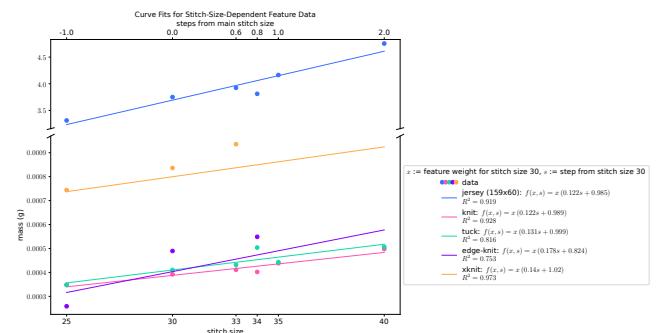


Figure 6: Stitch size fit function for jersey sample knitted in each stitch size, as well as knit, tuck, edge-knit (excluding outliers), and x-knit (excluding outliers) operations when training the model on the basic swatch set in each stitch size.

⁵Corresponding to the amount a stepper motor modifies a cam path on the carriage; and, thus, to how far a needle pulls the loop down when making a stitch.

For knit, we find that $a = 0.122$ and $b = 0.989$, with fairly similar coefficient values in the linear fit functions for the other features. This finding led us to suspect that the linear fit function for knit could accurately represent the change in the other stitch-size dependent features as well, with benefits provided from the model's tendency to fit knit better than the other features and the need for only a couple training samples when fitting one feature. Accordingly, we find that the errors are comparable when fitting only to knit as with all dependent features (with reduced error for all stitch sizes except 25), as demonstrated in Table 2.

Furthermore, we achieved even lower errors for all stitch sizes except 40 by fitting to a single sample: jersey (159×60). As this sample contains mostly knits (along with some edge-knits), it primarily represents the knit feature and can emulate the trend in knit operation mass with increasing stitch size. Thus, we can achieve satisfactory results when generalizing to new stitch sizes by knitting only one sample per stitch size, in addition to the batch of seven swatches in the main stitch size.



Figure 7: One instance of the half-gauge rib swatch from the Tamm Petit acrylic, Winning rayon, and Yeoman Elastomeric yarn training sets. The rayon swatch appears looser because of the lower bending stiffness of the thin rayon yarn. The elastomeric sample is smaller because the yarn stretches under the tension of knitting and contracts once released from the needles.

6.2 Generalizing to Different Yarns

Our model predicts the overall mass of a knit object by adding together the masses of the basic operations. We wanted to test if it could be transferred to different yarns by scaling the basic operation masses by the ratio of the linear densities of the yarns. (In effect, translating from masses of the yarn the model was fit on, to lengths, and back to masses of the new yarn.) We weigh a yarn strand of known length to calculate the linear density of a given yarn (taking the average of multiple measurements to reduce the impact of noise). Using this approach, we can represent yarn quantity at the swatch and operation levels as a mass directly correlated to the linear density of the yarn.

To see if this method of transferring the model would work, we knit a complete basic swatch set in Winning brand rayon yarn (2/30 ne: two plies of 30 yarn count filament, in 840 yds/lb units) and in Yeoman's elastomeric yarn (19% Lycra, 81% nylon, yarn count not

given). A side-by-side comparison of equivalent samples is shown in Figure 7.

We measured the linear density of the Tamm Petit acrylic yarn used in our initial sample set as 7.2×10^{-4} g/cm.⁶

Re-scaling the operation masses from our main model using the linear density of the rayon yarn (measured value: 4.1×10^{-4} g/cm; 57% of the acrylic's linear density) results in good agreement with a direct fit to the rayon sample set (Table 3), with differences around 5%. The low error we achieved on the rayon set suggests that the model could transfer well to other mechanically-similar yarns (even of different radii and material), though more data is needed to validate this hypothesis.

Table 3: Rayon operation masses. Operation masses for rayon yarn estimated by re-scaling the masses from acrylic yarn is in good agreement with fitting directly to a set of knit rayon samples.

Training method	d_1	d_2	edge-knit	knit	miss	tuck	x-knit
Re-scaled from acrylic (mg)	0.05	0.09	0.27	0.21	0.08	0.22	0.46
Fit to rayon samples (mg)	0.05	0.10	0.29	0.23	0.08	0.22	0.48

On the other hand, re-scaling using the linear density of the elastomeric yarn in its relaxed state (measured: 6.3×10^{-4} g/cm, 87% of the acrylic's linear density) deviates quite a bit from the model fit to the elastomeric sample set. This high error is unsurprising. The yarn is under tension when knitting on the machine, so the linear density we measured was likely much too large for the (stretched) yarn in the sample. These findings support the assumption that the crucial length we want to consider when estimating the quantity of yarn in a knitted structure is that of the yarn under tension.

Table 4: Elastomeric operation masses. Operation masses for elastomeric yarn estimated by re-scaling masses from acrylic yarn do not agree with those fitting directly to a set of knit elastomeric samples. Accounting for stretch improves the agreement.

Training method	d_1	d_2	edge-knit	knit	miss	tuck	x-knit
Re-scaled from acrylic (mg)	0.08	0.15	0.41	0.33	0.12	0.35	0.70
Re-scaled from acrylic with 36% LD (mg)	0.03	0.05	0.15	0.12	0.04	0.12	0.25
Fit to elastomeric samples (mg)	0.03	0.06	0.03	0.12	0.05	0.12	0.30

One interesting thing we can do, however, is to use the discrepancy between the scaled model and the model fit to the knit samples of elastomeric to approximate the axial strain in the yarn under the force of the machine when knitting. To do so, we iteratively increment the percentage of the elastomeric yarn's relaxed linear density to use for re-scaling and take the value that achieves the best agreement with the directly-fit model as our estimate for the yarn's linear density under tension. In this case, we found that 36% of the relaxed density (i.e., a 2.8x elongation of the yarn) resulted in the best fit, with an estimated linear density under tension of 2.3×10^{-4} g/cm. This is in rough agreement with an ad-hoc stretch estimate made by measuring the yarn between the yarn feeder and

⁶This differs slightly from the density implied by the "2/30 nm" yarn count – 6.7×10^{-4} g/cm – a pattern we saw repeated for all of our yarns.

the yarn carrier on the machine, then cutting the yarn off of the machine and allowing it to relax⁷. This suggests that some basic tensile testing might allow the transfer of model parameters even to elastic yarns.

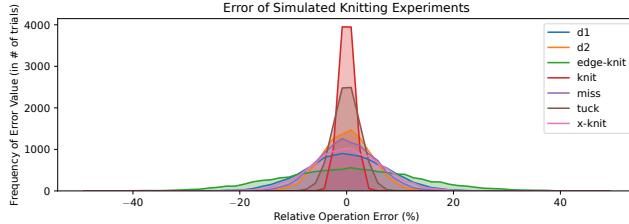


Figure 8: Histogram of the relative error in operation masses in simulated experiments following our fitting procedure. Generally, we would expect our procedure to fit the basic operation masses within 10% of the true value.

6.3 Simulated Data and Evaluating Noise

We wished to understand if our method is reproducible, or if, perhaps, we had simply been lucky in our particular measurements. To answer this question, we performed simulated experiments with machine and measurement noise levels matching those we observed and basic operations with masses within 10% of our mean measurements. We then ran 10000 iterations of this simulation. In each run, the true masses of the seven basic operations were selected independently and uniformly at random within $\pm 10\%$ of the masses reported in Table 1. Each simulated experiment followed our training procedure: knit each training swatch three times, weigh these 27 swatches three times each, and fit the operation masses to this data.

We simulated knitting and yarn density variation by summing true masses and scaling by a factor distributed normally with mean 1 and standard deviation 0.02; we simulated weighing by scaling by a factor distributed normally with mean 1 and standard deviation 0.001. (These proportional variations were estimated from the Tamm Petit sample set we knit and measured, and – owing to both the relatively small sample size and the estimate of the machine deviation not accounting for scale variation – may be a slight over-estimate, which we prefer to being overly confident.)

The histogram of relative deviations of the experimentally determined operation masses from the true masses is shown in Figure 8. We observe that in these hypothetical runs of our experiment, most of the basic operation masses are estimated within $\pm 10\%$ of their true values, with the only operation outside of that range being edge-knit. We suspect this much higher variance for edge knit is due to its relatively low representation in our chosen sample set. (By mass, the whole sample set contains 0.81g of edge knits, while it contains 22g of knit [27× more!] – the best-estimated overall.)

⁷We measured a 78 cm length of yarn under tension, which relaxed to a length of 36 cm. The strand of yarn weighed 0.015 g, corresponding to a linear density under tension of 1.9×10^{-4} g/cm and a stretch of 46%. Since we could only measure a small length of yarn when threaded on the machine, we suspect that this ad-hoc measurement is much noisier than the other linear densities we calculated – for which we took multiple measurements of long strands of yarn to reduce noise. Future work may improve measurement accuracy using a simple tension meter or tensile testing machines.

This simulation gives us some confidence that others applying our procedure, at least on knitting machines with similar repeatability, will likely succeed in estimating the operation masses.

Table 5: Noise from swatch repeatability as percent deviation for basic set in main stitch size, 30, sorted in descending order.

Swatch name (needles × rows)	Percent deviation
jersey (18 × 530)	3.394%
jersey (159 × 60)	2.021%
half-gauge jersey (159 × 60)	1.94%
tuck-stitch (159 × 60)	1.606%
1 × 1 rib (159 × 60)	1.521%
jersey tube (9 × 1060)	1.257%
half-gauge 1 × 1 rib (159 × 60)	0.436%

7 DISCUSSION

In our work, we demonstrate that linear regression can model additive yarn operations in knitted swatches. Though our proposed set of basic swatches use conventional stitch patterns⁸, our basic operations can represent the infinite number of patterns that knitting machine programmers may come up with. Our method is simple and effective, but we think some interesting further explorations remain.

Basic operations and swatches. We chose our basic operations as a simple and intuitive set, but we did not exhaustively consider alternatives. Notably, we found that some of our initial attempts at a set of basic operations led to very poor model performance (often when certain operations weighed less than 0.1 mg or were nearly linearly-dependent). A potential avenue for future work would be to stochastically or exhaustively explore possible choices of basic operations and training swatches. A series of virtual experiments, along the lines of what we explored in the extensions, might form a good basis for such an exploration.

Since the process of shaping fabrics primarily involves non-additive operations, such as loop transfers, we hypothesize that our method will transfer to non-rectangular fabrics. However, we neglected some additive yarn operations that can occur in shaped fabrics, like the ‘split’ operation, knitting through stacks of loops, and tuck-and-turns at the end of short row knitting. It would be interesting to discover if they require adding new basic operations or if they are close enough to some (combination of) basic operations to not require more training examples.

Sources of error. One possible source of systematic error in our results is variations in yarn density between different cones of the same yarn – such that basic operations that use the same length of the yarn would vary in mass. Throughout our swatch knitting, we pulled lengths of yarn (1 to 3m) from the cones we were using and computed their linear density, observing it to generally be relatively

⁸Beyond the favorable matrix properties that we get from our basic set of swatches (as demonstrated in Figure ??), using well-known stitch patterns further enables our method’s repeatability and ease of comprehension.

uniform⁹. However, this may not be the case for all yarns. It would be worthwhile to investigate if measuring the linear density of a small piece of yarn before every knit job and using it to correct for any variation would improve our model's predictions.

Future work may consider using higher-precision scales or more uniform yarn to help isolate other potential sources of error to further improve the modeling.

Alternative approaches. It is natural to ask whether a more sophisticated fitting procedure would yield better results. We explored two such directions: non-negative linear regression – which helped regularize the masses of some initial basic operations, but ended up unnecessary for the swatch set we arrived at; and convex optimization with constraints based on needle geometry and the triangle inequality – which provided very little improvement (reducing the error by less than 1%) for the complexity they added.

In Figure 9, we evaluate the three models' comparative performances on randomly selected training sets. Adding non-negativity constraints appears to improve performance with small training sets, and constrained convex optimization may help reduce the impact of outliers and overfitting. The advantages to adding constraints become less significant as the size of the training set increases, demonstrating that simply using a substantial amount of training data can achieve high performance, even with swatches that are not strategically selected.

Yarn varieties. Future work may incorporate additional yarn varieties, possibly with extensions to our model that consider the given fiber's material properties beyond linear density. In particular, interesting insights could be gained about the behavior of elastomeric yarn and a comparison of synthetic, plant-based, and animal-based fibers.

Application integrations. Given that one of our motivations for this work was to obtain better length measurements for yarn simulations (e.g., [Kaldor et al. 2008; Pizana et al. 2020; Sperl et al. 2022]), a particularly promising avenue for future work would be to use our model to initialize such simulations.

Similarly, adding yarn quantity estimation to a commercial knit design package and/or the open-source knitout toolchain would be useful.

8 CONCLUSION

In this work, we propose a simple method that can predict the amount of yarn necessary to produce an arbitrary knitted pattern with less than 5% mean absolute error. The method generalizes across arrangements of yarn operations, dimensions, swatches, stitch sizes, and yarns of various linear densities. In addition, we present a small calibration set that practitioners can use to re-fit our model for other settings. We believe that our estimates of the distribution of yarn within a knit item, when used to initialize yarn-level fabric simulation, can enable efficient computational fabrication of complex, customized structures and further the understanding of the mechanical properties embodied by these fascinating metamaterials.

⁹For our main yarn, Tamm Petite Medium Camel (T4221) 2/30 nm, linear density ranged from approximately 7.0×10^{-4} to 7.6×10^{-4} g/cm when using strands of yarn taken from different cones and at different points throughout the cone's use.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. 1955444.

REFERENCES

Edgaras Arbataitis, Daiva Mikucioniene, and Liudmyla Halavska. 2021. Flexible Theoretical Calculation of Loop Length and Area Density of Weft-Knitted Structures: Part I. *Materials* 14, 11 (2021), 3095.

Alexander Cachine, Krishna Singal, Sarah E Gonzalez, Michael S Dimitriev, and Elisabetta A Matsumoto. 2023. The effect of yarn mechanics on the non-linear elasticity of knitted textiles. *Bulletin of the American Physical Society* (2023).

R Ciukas and D Sadauskas. 2004. Theoretical determination of area density and tightness factor for weft knitted fabrics. In *Proceedings of the International Textile, Clothing and Design Conference, Dubrovnik, Croatia*. 3–6.

Sarah E Gonzalez, Michael S Dimitriev, Krishna Singal, and Elisabetta A Matsumoto. 2022. Developing a Constitutive Model for Knitted Fabrics. *APS March Meeting 2022, abstract id.Y21.001* (2022).

Mari Kageyama, S. Kawabata, and Masaka Niwa. 1988. The Validity of a “Linearizing Method” for Predicting the Biaxial-extension Properties of Fabrics. *The Journal of the Textile Institute* 79, 4 (1988), 543–567. <https://doi.org/10.1080/00405008808659163> arXiv:<https://doi.org/10.1080/00405008808659163>

Jonathan M. Kaldor, Doug L. James, and Steve Marschner. 2008. Simulating Knitted Cloth at the Yarn Level. In *ACM SIGGRAPH 2008 Papers* (Los Angeles, California) (*SIGGRAPH '08*). Association for Computing Machinery, New York, NY, USA, Article 65, 9 pages. <https://doi.org/10.1145/1399504.1360664>

Dani Liu, Daniel Christe, Bahareh Shakibajahromi, Chelsea Knittel, Nestor Castaneda, David Breen, Genevieve Dion, and Antonios Kontos. 2017. On the role of material architecture in the mechanical behavior of knitted textiles. *International Journal of Solids and Structures* 109 (2017), 101–111. <https://doi.org/10.1016/j.ijsolstr.2017.01.011>

Y Luo and I Verpoest. 2002. Biaxial tension and ultimate deformation of knitted fabric reinforcements. *Composites Part A: Applied Science and Manufacturing* 33, 2 (2002), 197–203. [https://doi.org/10.1016/S1359-835X\(01\)00106-3](https://doi.org/10.1016/S1359-835X(01)00106-3)

Kausalya Mahadevan, Vanessa Sanchez, and Katia Bertoldi. 2022. The mechanical response of periodic and aperiodic multistable knit architectures. In *APS March Meeting Abstracts (APS Meeting Abstracts, Vol. 2022)*. Article Y21.006, Y21.006 pages.

Shashank G Markande and Elisabetta A Matsumoto. 2020. Knotty knits are tangles on tori. (2020). <https://doi.org/10.48550/ARXIV.2002.01497>

James McCann. 2017. The “Knitout” (k) File Format. [Online]. Available from: <https://textiles-lab.github.io/knitout/knitout.html>.

Hideo Morooka, Yoichi Matsumoto, and Harumi Morooka. 1998. A geometric analysis of the stitch form of a circular plain knit fabric inserted over a cylinder. *Textile research journal* 68, 12 (1998), 930–936.

Vidya Narayanan and James McCann. 2019. knitout-frontend-py. <https://github.com/textiles-lab/knitout-frontend-py> [Online; accessed 30. Apr. 2023].

Vidya Narayanan, Kui Wu, Cem Yuksel, and James McCann. 2019. Visual Knitting Machine Programming. *ACM Trans. Graph.* 38, 4, Article 63 (7 2019), 13 pages. <https://doi.org/10.1145/3306346.3322995>

Jay H Park, S M Fijul Kabir, Ihsan Uluturk, Sejiro Yoshihara, Joshua Lee, and Scott Stapleton. 2023. Structure-Property Investigation of Knit Patterns on Mechanical, Thermal, and Moisture Properties. In *APS March Meeting Abstracts (APS Meeting Abstracts, Vol. 2023)*. <https://meetings.aps.org/Meeting/MAR23/Session/T09.9>

Alenka Pavko-Cuden, Ales Hladnik, and Franci Sluga. 2013. Loop length of plain single weft knitted structure with elastane. *Journal of Engineered Fibers and Fabrics* 8, 2 (2013), 155892501300800214.

FT Peirce. 1947. Geometrical principles applicable to the design of functional fabrics. *Textile Research Journal* 17, 3 (1947), 123–147.

José M. Pizana, Alejandro Rodriguez, Gabriel Cirio, and Miguel A. Otaduy. 2020. A Bending Model for Nodal Discretizations of Yarn-Level Cloth. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation (Virtual Event, Canada) (SCA '20)*. Eurographics Association, Goslar, DEU, Article 17, 9 pages. <https://doi.org/10.1111/cgf.14112>

Vanessa Sanchez, Kausalya Mahadevan, Gabrielle Ohlson, Moritz A. Graule, Michelle C. Yuen, Clark B. Teeple, James C. Weaver, James McCann, Katia Bertoldi, and Robert J. Wood. 2023. 3D Knitting for Pneumatic Soft Robotics. *Advanced Functional Materials* 33, 26 (2023), 2212541. <https://doi.org/10.1002/adfm.202212541> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1002/adfm.202212541>

Georg Sperl, Rosa M. Sánchez-Banderas, Manwen Li, Chris Wojtan, and Miguel A. Otaduy. 2022. Estimation of Yarn-Level Simulation Models for Production Fabrics. *ACM Trans. Graph.* 41, 4, Article 65 (7 2022), 15 pages. <https://doi.org/10.1145/3528223.3530167>

A Vékásy. 1960. Examination of the cover-factor and specific weight of weft-knitted or looped basis texture based on the exact value of the loop length. *Acta Technica* 31 (1960), 69–102.

Paras Wadekar, Vignesh Perumal, Genevieve Dion, Antonios Kontsos, and David Breen. 2020. An optimized yarn-level geometric model for Finite Element Analysis of weft-knitted fabrics. *Computer Aided Geometric Design* 80 (2020), 101883.

Tomoko Yamada, Noriko Ito, and Masaru Matsuo. 2003. Mechanical Properties of Knitted Fabrics Under Uniaxial and Strip Biaxial Extension as Estimated by a Linearizing Method. *Textile Research Journal* 73, 11 (2003), 985–997. <https://doi.org/10.1177/004051750307301109> arXiv:<https://doi.org/10.1177/004051750307301109>

Mark S. Yeoman, Daya Reddy, Hellmut C. Bowles, Deon Bezuidenhout, Peter Zilla, and Thomas Franz. 2010. A constitutive model for the warp-weft coupled non-linear behavior of knitted biomedical textiles. *Biomaterials* 31, 32 (2010), 8484–8493. <https://doi.org/10.1016/j.biomaterials.2010.07.033>

A APPENDIX

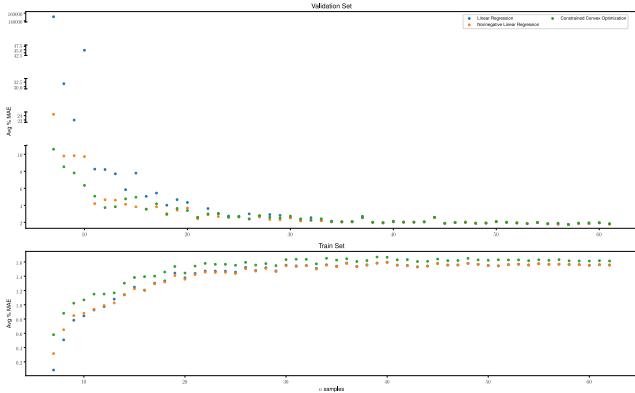


Figure 9: Comparing % MSE over 100 randomly selected training sets of size n using different modeling approaches.

A.1 Machine Knitting

To operate, automated v-bed knitting machines move a mechanism called a *carriage* across two perpendicular, flat *beds* of needles. The carriage contains cams that actuate a needle based on the operation that the machine is programmed to perform at a given time. The needles are composed of a shaft that stores loops and a hook that grabs onto new yarn. Depending on the machine model, either a sliding mechanism or a latch might be used to intersect new yarn with existing loops. During knitting, *feeders* spool yarn off of a cone, and *carriers* transport the yarn across the bed to the needle as it is actuated, laying the strand of yarn over the needle’s hook.

A.2 Stitch Patterns

A stitch pattern is a set of knitting operations that repeats throughout a given fabric. Since a stitch pattern is defined by the combination of yarn operations used in a fabric until the pattern repeats, the set of knit-able stitch patterns is theoretically infinite in size.

In this work, we chose to focus on some of the most common stitch patterns:

- (1) **Jersey:** Also known as plain knit or stockinette, jersey is fabricated using only the knit and edge-knit operations. It therefore consists of either all front knits or all back knits, the only difference being the perspective of view.
- (2) **Garter:** Defined by the alternation between rows of front and back knits. Garter tends to be stretchier in the vertical

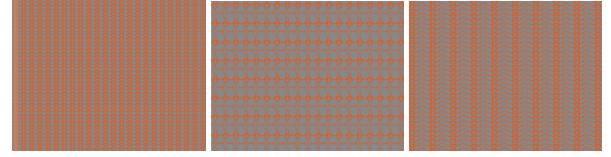


Figure 10: Jersey

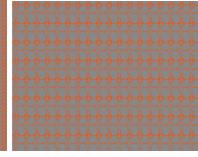


Figure 11: Garter



Figure 12: Rib



Figure 13: Seed

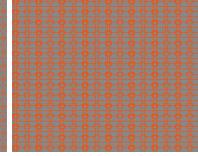


Figure 14: Tuck

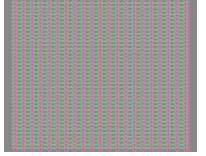


Figure 15: Tube

Knitout Visualizations

wale-direction, since the alternation between front and back knits occurs by row.

- (3) **Rib:** To fabricate rib, the alternation between front and back knits occurs from needle-to-needle in a given row. Thus, rib tends to be stretchier in the horizontal course-direction.
- (4) **Seed:** Production of this pattern involves alternating front and back knit both within and between rows, resulting in raised, seed-like bumps.
- (5) **Tuck-stitch:** The tuck-stitch pattern is created by alternating between knits and tucks on every other needle, switching which operation is performed on odd versus even needles between rows.
- (6) **Tube:** Any of the above patterns can be knitted into a tubular structure by knitting between beds on a knitting machine, as opposed to a single-bed approach.

A.3 Variations in Fabrics

A.3.1 Gauge. Though a knitting machine’s gauge is typically fixed, we can multiply the gauge by an integer factor by leaving some needles empty throughout the piece. We knitted some samples in “half-gauge,” meaning only every other needle is in use—doubling the gauge. We can also think of half-gauging as alternating between knits and misses, so this method adds more degrees of freedom to our system of equations.

A.3.2 Stitch Pattern Varieties. Within a single category of stitch patterns, there often exist numerous variations. For example, one might vary the ratio of rows composed of front (f) and back (b) knits, to make $f \times b$ garter (e.g., one row of front knits for every two rows of back knits would be 1×2 garter).