Data-driven Homogenization of Euler Beam

Conor Rowan

Spring 2024

The equations for beam bending are fourth-order and are thus quite different from the usual applications of homogenization techniques on second-order equations. It is interesting to investigate what applying the perturbation method yields on this model. The governing equation for bending of an Euler-Bernoulli beam is

$$\frac{\partial^2}{\partial x^2}\left(EI\frac{\partial^2 u}{\partial x^2}\right) = p(x)$$

where x is the axial coordinate, $E(x)$ is the Young's Modulus, $I(x)$ is the moment of inertia of the cross-section, $u(x)$ is the bending displacement, and $p(x)$ is the distributed bending force. We will carry out the usual homogenization approach: the material has multiscale behavior which depends on the fine scale coordinate $y = \frac{x}{\eta}$ where η is the width of the periodically repeating microstructure, the displacement is expanded into a coarse and fine scale contribution $u^\eta(x) = u^0(x) + u^1(x,y)$, and derivatives have a coarse and fine scale component $d/dx^\eta := \partial/\partial x + (1/\eta)\partial/\partial y$. We assume that the cross-sectional moment of inertia is the constant 1 for simplicity, as the geometry of the cross-section will not play a fundamental role in the homogenized response. Plugging the multiscale expansions into the governing equation, we have

$$\left(\frac{\partial}{\partial x} + \frac{1}{\eta}\frac{\partial}{\partial y}\right)\left(\frac{\partial}{\partial x} + \frac{1}{\eta}\frac{\partial}{\partial y}\right)\left[E(y)\left(\frac{\partial}{\partial x} + \frac{1}{\eta}\frac{\partial}{\partial y}\right)\left(\frac{\partial u^0}{\partial x} + \eta\frac{\partial u^1}{\partial x} + \frac{\partial u^1}{\partial y}\right)\right] = p(x)$$

$$= \left(\frac{\partial}{\partial x} + \frac{1}{\eta}\frac{\partial}{\partial y}\right)\left(\frac{\partial}{\partial x} + \frac{1}{\eta}\frac{\partial}{\partial y}\right)\left[E(y)\left(u^0_{xx} + \eta u^1_{xx} + 2u^1_{xy} + \eta^{-1}u^1_{yy}\right)\right]$$

where we have switched to subscript notation for derivatives. Without further expanding this expression, it can be seen that there will only be one term at the lowest order of η. Thus, the microscale equation will have no forcing, which is not what we expect from usual homogenization problems. Perhaps the error is in assuming that the first order displacement term u^0 is independent of the microscale coordinate y. Technically, this is something which needs to be proved, though it is often treated almost like a modeling choice. Removing this assumption would make the resulting expressions even more tedious. Thus, we will take a different route to explore homogenized beam bending equations. *The central insight of homogenization is that to a good approximation differential equations with high frequency periodic variations in the material tend behave like as if the material were some (unknown) constant.* The homogenization procedure provides a way of approximating this constant as a function of the periodic variations. For this reason, we refer to the homogenized response as an "effective" material property. See this plot for an example. In lieu of deriving an explicit homogenization framework that provides a means to approximate the effect properties, we can conduct a kind of "data-driven" homogenization. This is a simple process–first, we parameterize the spatial variation of the multiscale material. This is done in the following way:

$$E(x; E_0, a, b, c, \eta) = E_0 + a\sin\left(\frac{\pi x}{\eta}\right) + b\sin\left(\frac{2\pi x}{\eta}\right) + c\sin\left(\frac{3\pi x}{\eta}\right)$$

This is just one simple model of a multiscale material. There are five parameters: the "mean" E_0, the scale size η, and three weights on sine functions of different frequency. We see that the material only depends on the fine scale coordinate $y = x/\eta$. Note that the coefficients must be chosen such that the stiffness does not become negative. The next step is solve this equation exactly, which can be done by writing

$$u(x; E_0, a, b, c, \eta) = \int_0^x \int_0^y \left(\frac{1}{E(z; E_0, a, b, c, \eta)} \int_0^z \int_0^w p(\xi) d\xi dw \right) dz dy$$

and integrating numerically. When no constants of integration are included in the solution, we have that

$$u(0) = u'(0) = u''(0) = u'''(0) = 0$$

which is fine for our purposes, though this is not a typical set of boundary conditions. Ignoring the constants of integration simplifies the analysis, but does not fundamentally change the problem. *With an exact solution in hand, we can now ask: is this function well-approximated by the solution to a bending problem with the same force but a constant stiffness?* If so, we can compute this stiffness and treat it as the effective material property. In calibrating single scale solutions from the multiscale solution, let's use $p(x) = 1$ for simplicity. With constant forcing, the single scale solution with constant but unknown stiffness can be written analytically as

$$\tilde{u}(x) = \frac{1}{\tilde{E}} \frac{x^4}{24} := \tilde{C} \frac{x^4}{24}$$

which comes from repeated integration and ignoring constants of integration. We can fit the homogenization parameter $\tilde{C}$ with a least-squares procedure as follows:

$$\mathcal{E} = \int_0^L \frac{1}{2}\left(\tilde{C}\frac{x^4}{24} - u(x; E_0, a, b, c, \eta)\right)^2 dx$$

We are assuming that like typical homogenization problems, the exact solution of the multiscale Euler beam will be very well approximated by an effective stiffness. For a given set of parameters determining the material, we can determine the effective stiffness by minimizing the error:

$$\frac{\partial \mathcal{E}}{\partial \tilde{C}} = 0 = \int_0^L \left(\tilde{C}\frac{x^4}{24} - u(x; E_0, a, b, c, \eta)\right)\frac{x^4}{24}$$

$$\implies \frac{1}{\tilde{E}} = \tilde{C}(E_0, a, b, c, \eta) = \frac{\langle x^4/24, u(x; E_0, a, b, c, \eta)\rangle}{\langle x^4/24, x^4/24\rangle}$$

This is a way of computing the effective properties for a given set of parameters from the analytical solution. Repeating this process can be used to generate a data set relating the microstructure parameters to the homogenized stiffness, and this data could be fit with a neural network to act as a homogenization surrogate model. The surrogate model could be queried for new combinations of parameters within the training data set to compute the homogenized stiffness. Note that one test of whether the beam model can actually be "homogenized" is whether the effective properties computed from one force generalize to a different one. I find that the homogenized properties computed for the constant bending force work very well for different loads. This suggests that there should be a way to analytically homogenize the Euler beam bending equations. It is quite surprising the extent to which the effect of high-frequency periodic fluctuations in the material are captured by a reduced constant stiffness. See Figures 1-6 for results. See Figures 7 and 8for discussion of using a data-driven surrogate model for homogenization.

As a final demonstration of this homogenization method, we can use the surrogate model for uncertainty quantification. We assume that the parameters in the multiscale material are distributed uniformly and have no spatial correlation. To avoid extrapolating the surrogate model, these uniform distributions are defined over the same range as was used to generate the training data. At every integration point in the forward solve, the surrogate model is used to evaluate a random homogenized material property for the multiscale material. Note that it is no longer clear what would constitute an analytical solution for even a single forward solve of this problem, as we are assigned microstructures of finite size to points. See Figure 9 for results of the Monte Carlo simulation of the random material.

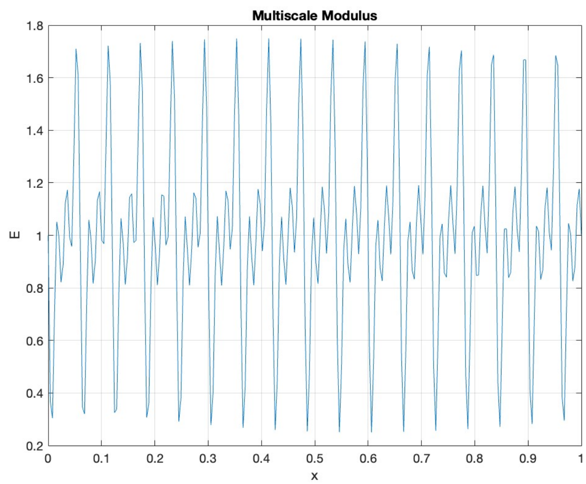

Figure 1: A particular choice of Young's Modulus displaying multiscale behavior from the given parameterization. Note that the oscillations are a large percent of the average.

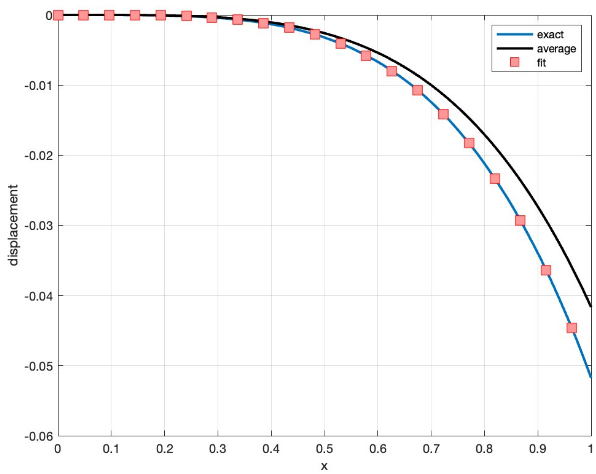

Figure 2: The exact solution computed with the multiscale modulus can be very accurately approximated by a constant modulus. Note that the multiscale beam is less stiff than the solution computed using the direct average over the microstructure E_0. This is the case for second-order homogenization problems as well. Even though the modulus displays very large oscillations, the exact solution is smooth. This departs from exact solutions in second-order homogenization problems and seems to be a function of the higher-order of differentiation.

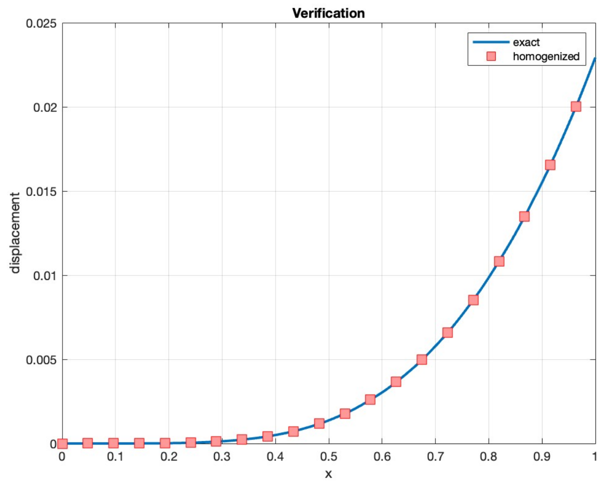

Figure 3: The homogenized material properties calibrated on a constant bending force of $p = 1$ also function to model the beam for different loading configurations. Here the force is $p(x) = (1 - x)\sin(2\pi x)$.

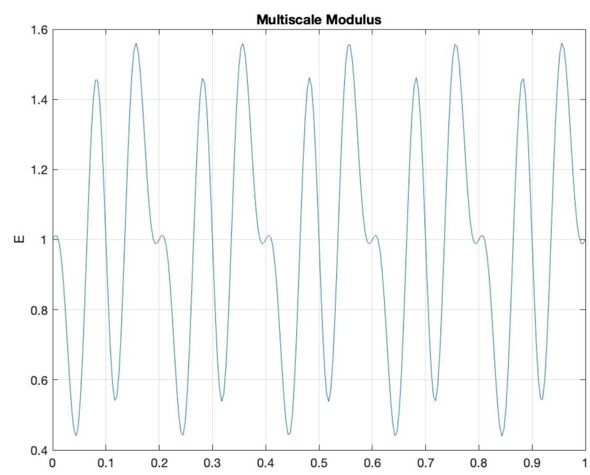

Figure 4: Another example of a multiscale microstructure generated with the Fourier-type parameterization. The scale size η is significantly larger in this case.

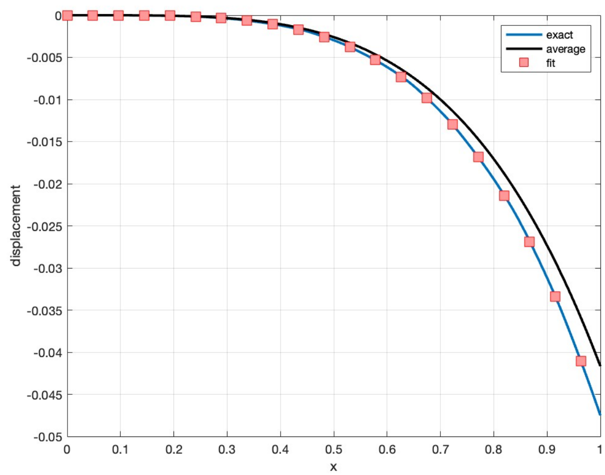

Figure 5: Exact solution fit with constant stiffness solution. Once again, the multiscale behavior manifests not in oscillations of the displacement rather as a net loss of stiffness.

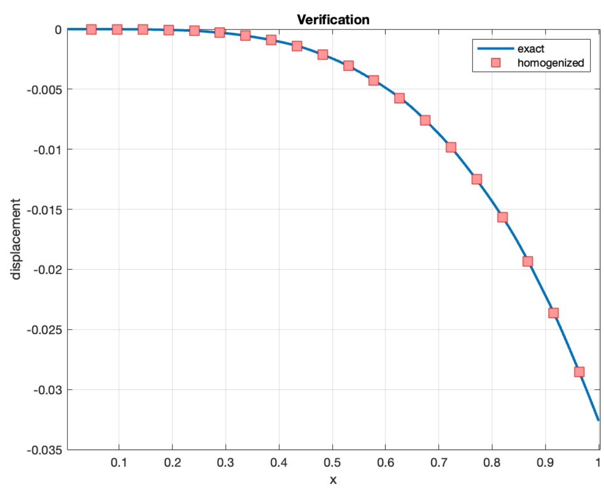

Figure 6: Verifying that the effective modulus computed using $p(x) = 1$ are accurate when used to compute the bending displacement for a distributed force of $p(x) = -\left(1 - \left(\frac{x}{2}\right)\right)^4$. The homogenized stiffness perfectly matches the exact solution.

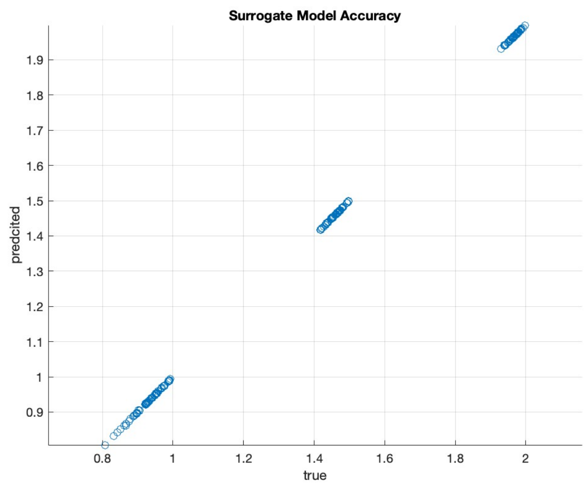

Figure 7: A two-hidden layer neural network with 12 neurons per layer is written out symbolically in MATLAB, then it is differentiated symbolically with respect to the parameters. The forward and gradient operations are written to files as numerical functions. A loss function is written as another file then fed to a quasi-Newton optimizer. The parameters of the network are tuned to map the relationship between the five parameters in the material modulus and the homogenized stiffness. The network is capable of representing this relationship nicely. Note that only 3 values of average parameter E_0 are used in training, which explains the disjoint values.

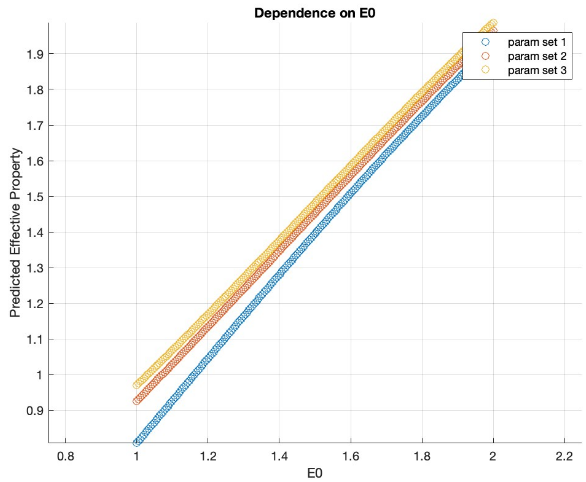

Figure 8: We can verify that the surrogate model makes "well-behaved" predictions interpolating between different values of E_0 that were not explicitly seen in the training set. Though we do not know the ground truth in this case, it is clear the interpolation behavior is quite regular. Three different sets of parameters controlling the oscillations and scale size of the microstructure are used here.

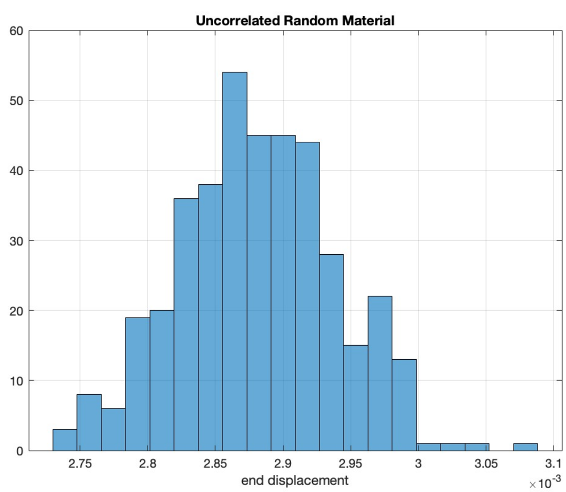

Figure 9: The parameters of the microstructure are sampled uniformly within the ranges the surrogate model was trained on, and a random effective property is assigned at each integration point. We perform 400 Monte Carlo simulations and show the distribution of end displacement values.