"Local-Global" Material Optimization

Julian Panetta

September 1, 2015

This writeup describes the material optimization method of [4] in more detail. Then, in Section 5, it attempts to explain why the method is so robust in practice.

1 Notation

We use linear elasticity to model deformations of our object, Ω. We denote the strain of displacement field $u(x)$ as $\varepsilon(u)$ and the stress as $\sigma(u) = C : \varepsilon(u)$.

2 Problem Specification

Given a loading scenario and the desired displacements of portions of the boundary, we design the material property field $C(x)$ so that the desired displacements are achieved.

The loading scenario can consist of both Dirichlet and Neumann conditions:

$$\begin{aligned} u &= u_d \quad \text{on } \Gamma_d \\ n \cdot \sigma(u) &= \tau \quad \text{on } \Gamma_n. \end{aligned}$$

Here, n is the surface normal. The boundary displacement is denoted by u_t and is defined on Γ_t. Our implementation also allows specifying loading/target conditions on a per-coordinate basis, but for notational simplicity we don't discuss this.

For the input to make sense, it should satisfy $\Gamma_n \cup \Gamma_d = \partial\Omega$, $\Gamma_n \cap \Gamma_d = \emptyset$, and $\Gamma_t \subseteq \Gamma_n$ (i.e. the full boundary's loading is specified, the Neumann and Dirichlet regions do not overlap, and the target region does not overlap the Dirichlet region).

3 Direct (Displacement) Approach

The most natural approach is to minimize the displacement deviation under the loading scenario:

$$\begin{aligned} \min_C \int_{\Gamma_t} \|u - u_t\|^2 \, \mathrm{d}A(x) & \\ \text{s.t.} -\nabla \cdot [C : \varepsilon(u)] &= 0 \quad \text{in } \Omega \\ n \cdot [C : \varepsilon(u)] &= \tau \quad \text{on } \Gamma_n \\ u &= u_d \quad \text{on } \Gamma_d \end{aligned}$$

The gradient can be computed efficiently using the adjoint method, but we've found this optimization problem difficult to solve robustly.

This approach was used in [1], and [2] uses a similar approach where the traction condition can be made a soft constraint to improve robustness. To further improve robustness and accelerate convergence, [2] searches over a reduced subspace of smooth material distributions.

4 "Local-Global" Strain-fitting Approach

Instead of fitting displacements directly, we found an iterative strain-fitting approach to work better. The idea is to run two simulations: one with the loading scenario only and the other with the target condition added as a Dirichlet constraint. These solutions are called the "Neumann" and "Dirichlet" solutions (u_N and u_D) in [4] because the former applies the Neumann condition on the target region Γ_t, and the latter enforces the target displacement using a Dirichlet condition:

$$\begin{aligned} &\text{"Neumann Problem"} && \text{"Dirichlet Problem"} \\ -\nabla \cdot [C_i : \varepsilon(u_N)] &= 0 \quad \text{in } \Omega && -\nabla \cdot [C_i : \varepsilon(u_D)] = 0 \quad \text{in } \Omega \\ n \cdot [C_i : \varepsilon(u_N)] &= \tau \quad \text{on } \Gamma_n && n \cdot [C_i : \varepsilon(u_D)] = \tau \quad \text{on } \Gamma_n \setminus \Gamma_t \\ u_N &= u_d \quad \text{on } \Gamma_d && u_D = u_d \quad \text{on } \Gamma_d \\ & && u_D = u_t \quad \text{on } \Gamma_t \end{aligned} \tag{1}$$

(recall $\Gamma_t \subseteq \Gamma_n$). Unless the current material field $C_i(x)$ already achieves the target deformation, the strain fields of the two simulations will differ. We can view $\varepsilon(u_D)$ as an estimate of the target deformation's strain field (since it integrates to the correct target boundary displacement). We can also view $\sigma(u_N)$ as an estimate of the internal stresses in the target deformation (since it satisfies the loading conditions). This suggests updating $C(x)$ to minimize the Frobenius norm distance between $\varepsilon(u_D)$ and the strain corresponding to stress estimate $\sigma(u_N)$:

$$C_{i+1} = \operatorname*{argmin}_C \|C^{-1} : \sigma(u_N) - \varepsilon(u_D)\|_F^2. \tag{2}$$

This least-squares minimization can be done over any subspace of material fields. For instance, in [4] we minimize over the space of piecewise-constant isotropic materials, parametrized by a Young's modulus and a Poisson's ratio per voxel. We add a Laplacian regularization term to encourage smooth variation, but model reduction could be used as in [2].

We can then use C_{i+1} in the next round of simulations (1). Without the Laplacian regularization, this strain-fitting can be done independently in each voxel, and our iterations alternate between global PDE solves and local fitting. This is very similar in spirit to the ARAP deformation and parametrization techniques, and we adopt the local-global name (even though the regularization term makes it more of a global-global iteration...).

5 Convergence

While it's not immediately clear that this strain-fitting local-global iteration should converge, we've found it to be extremely robust in practice. We can see why this might be the case by considering a similar but more "symmetric " fitting energy.

First, notice that instead of fitting strains as in (2), we could just as easily have fit stresses:

$$C_{i+1} = \operatorname*{argmin}_C \frac{1}{2} \int_\Omega \|\sigma(u_N) - C : \varepsilon(u_D)\|_F^2 \, \mathrm{d}x.$$

In fact, in force-feedback applications (as opposed to target deformation applications), this version is probably preferable. A natural compromise between the two is to perform the fitting in a space "halfway between" stress and strain:

$$C_{i+1} = \operatorname*{argmin}_C \frac{1}{2} \int_\Omega \|C^{-\frac{1}{2}} : \sigma(u_N) - C^{\frac{1}{2}} : \varepsilon(u_D)\|_F^2 \, \mathrm{d}x. \tag{3}$$

This symmetric fitting energy has a very nice property: we can prove local-global iterations always decrease the fitting energy (3). First, we expand the energy:

$$\frac{1}{2} \int_\Omega \|C^{-\frac{1}{2}} : \sigma(u_N) - C^{\frac{1}{2}} : \varepsilon(u_D)\|_F^2 \, \mathrm{d}x = \int_\Omega \frac{1}{2}\sigma(u_N) : C^{-1} : \sigma(u_N) + \frac{1}{2}\varepsilon(u_D) : C : \varepsilon(u_D) - \underbrace{\sigma(u_N) : \varepsilon(u_D)}_{I} \, \mathrm{d}x$$

Integrating the "I" term by parts:

$$I = -\int_\Omega \cancelto{0}{\big(\nabla \cdot \sigma(u_N)\big)} \cdot u_D \,\mathrm{d}x + \int_{\partial\Omega} \big(n \cdot \sigma(u_N)\big) \cdot u_D \,\mathrm{d}A(x).$$

5.1 Simple Case: Full-boundary Target Condition

We first consider the case where $\Gamma_n = \Gamma_t = \partial\Omega$ and $\Gamma_d = \emptyset$. In other words, tractions and target displacements are specified on the entire boundary, and no Dirichlet conditions exist in the loading scenario. Then:

$$I = \int_{\partial\Omega} \big(n \cdot \sigma(u_N)\big) \cdot u_D \,\mathrm{d}A(x) = \int_{\partial\Omega} \tau \cdot u_t \,\mathrm{d}A(x) = \text{const},$$

and the "I" term is just a constant depending on the prescribed forces and target displacements. In this simplest case, it's clear that the fitting energy (3) is really just the sum of two elastic energies (up to a constant): the elastic energies of the two simulations! Thus, we can view our "local-global" simulate-then-fit iterations for (3) as an alternating minimization,

$$\begin{aligned} C^* &= \operatorname*{argmin}_C \min_{\substack{\nabla\cdot\sigma_N = 0 \\ n\cdot\sigma_N = \tau \text{ on } \partial\Omega}} \min_{\substack{\varepsilon_D = \varepsilon(u_D) \\ u_D = u_t \text{ on } \partial\Omega}} \frac{1}{2}\int_\Omega \|C^{-\frac{1}{2}} : \sigma_N - C^{\frac{1}{2}} : \varepsilon_D\|_F^2 \,\mathrm{d}x \\ &= \operatorname*{argmin}_C \min_{\substack{\nabla\cdot\sigma_N = 0 \\ n\cdot\sigma_N = \tau \text{ on } \partial\Omega}} \min_{\substack{\varepsilon_D = \varepsilon(u_D) \\ u_D = u_t \text{ on } \partial\Omega}} \frac{1}{2}\int_\Omega \sigma_N : C^{-1} : \sigma_N + \varepsilon_D : C : \varepsilon_D \,\mathrm{d}x, \end{aligned} \tag{4}$$

proving that each step decreases the energy. In particular, the minimization of complementary potential energy over divergence-free σ_N satisfying the traction condition is equivalent to the "Neumann" simulation, and the minimization of potential energy over integrable strain fields whose corresponding boundary deformations satisfy the target conditions is equivalent to the "Dirichlet" simulation.

5.2 General Case

In general, we can decompose the "I" term into three surface integrals:

$$I = \left(\cancelto{\text{const}}{\int_{\Gamma_t}} + \int_{\Gamma_n \setminus \Gamma_t} + \int_{\Gamma_d} \right) \big(n \cdot \sigma(u_N)\big) \cdot u_D \,\mathrm{d}A(x).$$

The first integral depends only on the prescribed tractions and target displacements, as in the simple case. Thus, up to a constant, we can write

$$I = \int_{\Gamma_n \setminus \Gamma_t} \tau \cdot u_D \,\mathrm{d}A(x) + \int_{\Gamma_d} \big(n \cdot \sigma(u_N)\big) \cdot u_d \,\mathrm{d}A(x).$$

The first integral is the work done by the prescribed traction τ in the "Dirichlet problem," and the second is the work done by imposing the Dirichlet loading condition u_d in the "Neumann problem." This means we can rewrite the strain fitting energy (3) as:

$$\underbrace{\frac{1}{2}\int_\Omega \sigma(u_N) : C^{-1} : \sigma(u_N) \,\mathrm{d}x - \int_{\Gamma_d} \big(n \cdot \sigma(u_N)\big) \cdot u_d \,\mathrm{d}A(x)}_{\text{Total Complementary Potential Energy of "Neumann Simulation"}} + \underbrace{\frac{1}{2}\int_\Omega \varepsilon(u_D) : C : \varepsilon(u_D) \,\mathrm{d}x - \int_{\Gamma_n \setminus \Gamma_t} \tau \cdot u_D \,\mathrm{d}A(x)}_{\text{Total Potential Energy of "Dirichlet Simulation"}}.$$

Finally, as in (4), we can express the local-global iteration as an alternating minimization of the symmetric fitting energy:

$$\begin{aligned} C^* &= \operatorname*{argmin}_C \min_{\substack{\nabla\cdot\sigma_N=0 \\ n\cdot\sigma_N=\tau \text{ on } \Gamma_N}} \min_{\substack{\varepsilon_D=\varepsilon(u_D) \\ u_D=u_d \text{ on } \Gamma_d \\ u_D=u_t \text{ on } \Gamma_t}} \frac{1}{2}\int_\Omega \|C^{-\frac{1}{2}} : \sigma_N - C^{\frac{1}{2}} : \varepsilon_D\|_F^2 \,\mathrm{d}x \\ &= \operatorname*{argmin}_C \min_{\substack{\nabla\cdot\sigma_N=0 \\ n\cdot\sigma_N=\tau \text{ on } \Gamma_N}} \min_{\substack{\varepsilon_D=\varepsilon(u_D) \\ u_D=u_d \text{ on } \Gamma_d \\ u_D=u_t \text{ on } \Gamma_t}} \text{"Neumann" TCPE + "Dirichlet" TPE.} \end{aligned} \tag{5}$$

6 Future Work

I haven't actually tried implementing the "symmetric" local/global iteration (5) because the strain-fitting energy has been robust enough, but it would be interesting to try.

7 Acknowledgements

Considering the "symmetric" version of the strain-fitting energy (3) was suggested by Bob Kohn; he proposed the same approach in [3] to determine an object's electrical conductivity tensor distribution from boundary voltage and current measurements.

References

[1] Bernd Bickel, Moritz Bächer, Miguel A. Otaduy, Wojciech Matusik, Hanspeter Pfister, and Markus Gross. Capture and modeling of non-linear heterogeneous soft tissue. *ACM Trans. Graph.*, 28(3):89:1–89:9, July 2009.

[2] Yong Chen Hongyi Xu, Yijing Li and Jernej Barbič. Interactive material design using model reduction. *ACM Trans. Graphics*, 34(2), 2015.

[3] Robert V Kohn and Alan McKenney. Numerical implementation of a variational method for electrical impedance tomography. *Inverse Problems*, 6(3):389, 1990.

[4] Julian Panetta, Qingnan Zhou, Luigi Malomo, Nico Pietroni, Paolo Cignoni, and Denis Zorin. Elastic textures for additive fabrication. *ACM Trans. Graph.*, 34(4):135:1–135:12, July 2015.