How Data Mining Can Catalyze 3D Printing

Hod Lipson[1]

Just as we're beginning to conquer 3D printing technology, a new challenge rears its head: material complexity. Every time I give a talk about 3D printing to a group of practicing engineers, exalting the technology's benefits and promises, a skeptical engineer from the back row inevitably raise a hand and asks: How good are these 3D-printed materials? How well do we know their properties, and how do these properties depend on various settings of the printer? True, 3D printing allows us incredible freedom of form and shape. But to any practicing engineer who designs machine parts with critical function, that's not good enough.

Additive manufacturing processes for metals, plastics, and even ceramics have matured to a point where we can fabricate machine parts of incredible geometric complexity. Many of these parts and materials have been tested thoroughly and passed with flying colors. But look closely at any of these advanced processes and you'll find hundreds of parameters that can be changed by the operator. For laser sintering, for example, there are numerous temperature settings, as well as laser speed, power, pulse width, scanning patterns, and dozens more. How do these knobs change the characteristics of the printed material? Nobody has a complete answer. The problem is only exacerbated as new additive manufacturing technologies become available, each with their own parameters and side effects.

Perhaps what might break the camel's back is the expansion of new materials. If we had to deal with just a few homogenous metals and plastics, we might eventually catch up. But the rapid expansion in the range of materials, including blended alloys, graded polymer mixtures, and even biomaterials with complex rheology, is simply too much. Solid mechanics experts and material scientists have struggled for decades to model even relatively simple composite materials; there is little chance we will keep up with orders of magnitude more complexity introduced by 3D printing possibilities.

I see two solutions on the horizon: One is feedback control, and the second is data-driven material models. Feedback control printers are additive manufacturing processes that watch and learn in a closed loop. They see what they print, they measure properties, and adjust their own knobs automatically to reach a desired performance, or alert the operator if the desired properties cannot be attained. Most of today's printers operate in open loop, with perhaps a little sensor feedback for monitoring temperature and material flow. Ultimately, feedback-based printers may become the only way to keep up with the rapidly expanding range of materials. This is especially important for open markets that involve materials that are not necessarily calibrated or guaranteed, and users that are not necessarily trained.

But feedback control can only help for properties that can be measured during the printing processes, like geometric tolerances and density. For more complex properties, like fatigue life and shelf life, we'll need the help of machine learning and data mining. Imagine that every measure of mechanical property can be associated with the dozens of printer settings used to make it. That huge data source hides the material models we are looking for, and the only way to extract that knowledge will be to use machine learning methods that can distill this experimental data into useful, predictive models. At that point, not only will we know what to expect, but we can also optimize.

As the editor-in-chief, I hope that you enjoy this issue and join us as an active participant in this new journal and resource. Look out for our next issue of the Journal publishing in June 2015.

Hod Lipson
Editor-in-Chief

[1]Cornell University, Ithaca, New York.