Uncertainty in the Development and Use of Equation of State Models

V. Gregory Weirs*, Nathan Fabian*, Kristin Potter†, Laura McNamara*, and Thomas Otahal*

*Sandia National Laboratories †Scientific Computing and Imaging Institute

ABSTRACT

In this paper we present the results from a series of focus groups on the visualization of uncertainty in Equation-Of-State (EOS) models. The initial goal was to identify the most effective ways to present EOS uncertainty to analysts, code developers, and material modelers. Four prototype visualizations were developed to presented EOS surfaces in a three-dimensional, thermodynamic space. Focus group participants, primarily from Sandia National Laboratories, evaluated particular features of the various techniques for different use cases and discussed their individual workflow processes, experiences with other visualization tools, and the impact of uncertainty to their work. Related to our prototypes, we found the 3D presentations to be helpful for seeing a large amount of information at once and for a big-picture view; however, participants also desired relatively simple, two-dimensional graphics for better quantitative understanding, and because these plots are part of the existing visual language for material models. In addition to feedback on the prototypes, several themes and issues emerged that are as compelling as the original goal and will eventually serve as a starting point for further development of visualization and analysis tools. In particular, a distributed workflow centered around material models was identified. It was also found that users of material models contribute and extract information at different points in this workflow depending on their role, but encounter various institutional and technical barriers which restrict the flow of information. We expect the identification of this workflow, as well as the potential bottlenecks, to influence the development of visualizations to improve communication across this workflow and effectively express the uncertainties within the material models community.

Index Terms: H.5.1 [Information Systems]: Multimedia Information—Evaluation/methodology; J.2 [Computer Applications]: Physical Sciences and Engineering

1 INTRODUCTION

The research objective of this project is to develop effective Material Model Uncertainty Visualization (MMUV) techniques, with the eventual goal of providing a software tool to users in the in materials modeling community. In order to identify and evaluate the best methods for visualization, we decided to conduct a series of focus groups aimed at modelers, analysts and code developers. The focus group approach was chosen to better understand the needs of the community and ensure the usability, utility, and adoptability of the software produced.

Material models describe the behavior of a specific material or class of materials and are used as inputs to multiphysics numerical simulations. Because the models are based on theory, they often require empirical information to calibrate or specify free parameters and results from the models may define ranges of possible values or a collection of valid scenarios. Sources of uncertainty within simulations are abundant; this work focuses on the uncertainty arising from variability in the material models and aims at understanding how users comprehend, incorporate, and utilize this qualitative information and how to enhance understanding through visual representations.

The domain of "material modeling" is extremely complex. Material models range from the relatively simple to the highly complex, and model formulation can vary depending on the problem of interest. For example, the models selected for use in structural analysis may be quite different from those chosen for a fluid dynamics calculation. Models may also have different modes; a structural analysis may require models for plastic deformation, fracture, and crack propagation. In addition, an engineering application often requires several models for each material, and the development of each mode and type of model requires significant specialized expertise.

The material models of interest in this work fall into two classes, Equation-of-State and solid mechanics. Equation-of-state (EOS) models describe relationships between thermodynamic variables such as pressure, density, temperature, internal energy, and the speed of sound in a material. Solid mechanics models describe relationship between the stress state and the strain (or deformation) of a material in the solid state. Our research focused on EOS models because we have existing techniques for visualizing EOS surfaces, such as the pressure as a function of density and temperature. Solid mechanics models are more complicated because stress and strain are 3×3 tensor quantities, and because solid mechanical behavior is often dependent on the deformation history of the material. Visualizing solid mechanics models, or at least the quantities of greatest interest to analysts, will be the subject of future work.

Material model domain complexity makes the development of useful and usable uncertainty visualizations quite challenging. Uncertainty is an abstract concept, even when the domain of interest and its associated uncertainties are relatively uncomplicated and/or well-studied. Material science is complicated and significant knowledge gaps and aleatory uncertainties exist for most materials; yet all technology development requires at least minimal characterization of material behavior and properties.

In addition to domain complexity, organizational complexities exist. The material modeling stakeholder community comprises diverse classes of experts. A minimum set includes researchers who develop and refine models of materials; simulation code developers whose libraries incorporate material models to support computational engineering analyses; and engineering analysts developing and running simulations for which material models are an important input. Within these broad categories there is considerable diversity in research problems and work practices, ensuring that a single visualization is unlikely to support all user activities. Moreover, members of the stakeholder community are dispersed throughout Sandia and often work on a diverse range of projects that have both internal and external partners and customers. Although the work of material modelers affects the work of code developers and analysts, it is likely that neither the modelers nor the consumers (code developers, analysts) fully recognize the organizational, conceptual, or practical dependencies among the work products they generate. As a result, eliciting how material modelers and the users of material models conceptualize, represent, and analyze uncertainty is a

*e-mail:{vgweirs, ndfabia, tjotaha, lamcnam}@sandia.gov
†e-mail:kpotter@sci.utah.edu

highly challenging problem.

1.1 Equation-of-State Models

An equation of state describes relationships between thermodynamic variables for a given material. Given any two variables, all other variables can be computed through the EOS under the assumption of thermodynamic equilibrium. EOS models can cover a very wide range of conditions, and different physical phenomena dominate material behavior in different regimes [4]. The relevant physics leads to a model for the EOS in that regime. Sometimes the boundaries between regimes are sharp, such as between ice and water, but in other cases there is a gradual transition, for example between liquid and gas phases above the temperature where three phases co-exist in thermodynamic equilibrium (or triple point). For a multiphase material model, EOS models for different regimes are blended or combined to describe the behavior over many regimes.

The various theories often have parameters that must be provided to fully specify the EOS. In many cases, the particular parameter values distinguish one material from another, but some parameters may be independent of the specific material. When a theory does not provide the parameter values, they must be determined by experimental measurements or by more sophisticated theories.

Material models are input data for engineering simulation codes; the material model is not a result of the simulation. The results of the simulation are influenced by the EOS, and of course poor EOS models generally lead to inaccurate simulation results. In engineering simulation codes, EOS models are often stored in tabular form, e.g., pressure values are stored for a number of discrete density-temperature points. The points are arranged in a grid. For an arbitrary density and temperature, the pressure is interpolated from nearby points in the table. Over the range of densities and temperatures, the pressure can be shown as a surface. Most tabular EOS models contain several thermodynamic variables; the most common are density, temperature, internal energy, and pressure.

In this work we have used a Mie-Grüneisen (MG) model for aluminum to provide test data for prototype visualizations. The MG EOS describes a single phase, and has a number of parameters that engineering codes allow the user to specify. The model is defined by

$$\begin{aligned} P(\rho,E) &= P_R(\rho)+\Gamma_0\rho_0(E-E_R(\rho)) && (1)\\ E(\rho,T) &= E_R(\rho)+C_V(T-T_R(\rho)) && (2) \end{aligned}$$

where ρ is the density, P is the pressure, E is the internal energy, T is the temperature, Γ_0 is the Grüneisen parameter, ρ_0 is the reference density, and C_V is the heat capacity. The last three are constants. The MG model relates the pressure and energy to a reference Hugoniot, a special curve comprised of the locus of shock states given an initial state; the subscript R denotes the reference Hugoniot. For many metals in the solid phase, a linear Hugoniot relation expressed in the shock velocity, U_s, and the particle velocity, U_p provides an excellent fit to experimental data:

$$U_s = C_s + S_1 U_p \qquad (3)$$

where C_s is the speed of sound at reference conditions and S_1 is the slope of the linear relationship (both constants.) The Hugoniot relations can be used to derive

$$\begin{aligned} P_R(\rho) &= P_0+\rho_0 u_s u_p && (4)\\ E_R(\rho) &= E_0+(P_R(\rho)+P_0)\mu/2\rho_0 && (5)\\ T_R(\rho) &= e^{\Gamma_0\mu}[T_0+C_V^{-1}\int_0^{\mu} e^{-\Gamma_0\mu}\mu^2 u_s \frac{du_s}{d\mu}d\mu], && (6) \end{aligned}$$

where $\mu = 1-\rho_0/\rho$ and P_0, E_0, and T_0 are the initial state.

1.1.1 Sources of Uncertainty in Materials Models

Uncertainties are abundant throughout the materials models workflow. Due to the complexity of a materials model, and the expertise and time required to create an accurate model, a simple source of uncertainty comes from poor models. Unfortunately, better models can suffer from poor parameter settings which can come either from experimental measures or theoretical conditions. Changes in phase are modeled by either strict or blended boundaries and can cause uncertainties in important areas of the model. Often, there exists multiple models for a single material and it is not always clear which model to use in a particular situation, nor is it obvious the ramifications of choosing a specific model. Finally, the storage of these models as tabular data can impose interpolation errors and resolution limitations. Because these models act as inputs to simulations, traditional simulation uncertainties, such as discretization, are multiplied by uncertainties propagated from the material models. Understanding the sources and developing methods for quantifying, displaying and analyzing the uncertainty will have great impact throughout the materials modeling community.

1.2 Research Method

A focus group is a structured group interview, facilitated by a moderator, in which participants explore an issue or set of issues of research importance. Because so many disciplines use focus groups, approaches to designing, deploying, and analyzing focus groups vary tremendously. However, all focus groups begin with the same basic principle: that exchanges among participants facilitate the expression of ideas, knowledge, behaviors, and opinions that may be invisible to individualized methods (such as a questionnaire or a one-on-one interview). Groups enable researchers to access a broader range of skills and experiences than a single respondent; listening to others express ideas and opinions can spur participants to remember and share information that might not have emerged in a one-on-one setting. Not surprisingly, focus groups are an excellent way of eliciting the kinds of information that people naturally express in group settings; or for documenting how knowledge emerges in the context of group interactions.

Focus groups play an important role in computer science research and software engineering, in ways that are germane to the goals of the MMUV project. Studies of software engineering processes have used focus groups to gather data about workflow patterns in engineering teams [6]. User- and interaction-oriented design paradigms suggest the use of focus groups to gather qualitative data on user expectations and system requirements, and to evaluate prototype interfaces [11, 3, 7]. For technology developers, focus groups also afford the opportunity to demonstrate sincere interest in user concerns. In that sense, focus groups can enhance relationships between the user community and the technology developers by establishing a foundation for ongoing communication and exchange of information. Over time, the user community perceives itself as a dedicated stakeholder in the work of the technology developers, instead of recipients of a product tossed over the proverbial fence.

The focus group approach does have significant drawbacks: while participants often generate excellent contextual insights, focus group data is less useful for analyzing long-term trends or generalizing about large populations. Moreover, focus groups are prone to groupthink bias and social dominance bias (i.e., when one of the group members exerts undue influence on the interactions or content of the group, either consciously or unconsciously [8]. A structured script, pilot runs, careful moderation, and a sound qualitative sampling strategy can enhance the quality and dependability of focus group findings. Even so, software developers should treat focus groups as a starting point for technology design and evaluation, if only because because focus groups only capture information on what users "say they do – not how they actually do it" [9]. Other approaches, including observation, user participation in design teams, multiple prototyping and iterative re-designs, are necessary to develop technologies that people perceive as truly adoptable.

1.3 Focus Group Process

As noted above, one of the major goals of the MMUV project was the design and development of visualizations that would be usable and useful to experts who generate, interact with, or rely on material models in their work – what we have described above as the "stakeholder community." The domain complexity of material modeling and the organizational complexity of the stakeholder community are precisely what makes material modeling an interesting computational science, information visualization, and technology design problem; but they also make it difficult to understand what "usability" and "utility" mean in the many contexts of work where material model visualization might be useful. Accordingly, we decided to use focus groups as a way of gathering expert knowledge about material modeling challenges at Sandia, and as a way of opening a dialogue with the user community so that we could better understand the current state of practice. For specific feedback on representing uncertainty in EOS models, four visualization prototypes were developed by the MMUV project. In May, June, and July of 2011, we conducted four focus groups with participants representing the various material modeling stakeholder communities described above. Participants were technical staff at Sandia National Laboratories or in some way affiliated with Sandia.

The diversity of the material modeling stakeholder community made composition of the focus groups a bit challenging, since members of the stakeholder subcommunities tend not to interact with each other on a regular basis (more on this issue below). The success of a focus group depends on the composition of the participant pool; people need to have enough common ground that they can communicate productively about the topic of discussion, but diversity of perspective can spark insights that might not emerge in a homogeneous group. All four focus groups included representatives of each of the three primary subcommunities described above. To recruit participants, we drew on contacts from our own Sandia networks. We scheduled the focus groups and invited participants but offered no additional incentives (neither snacks nor money), and all participation was completely voluntary.

In moderating the groups, we decided to use a team facilitation approach, in which a technical leader (Weirs) and a process leader (McNamara) managed the group logistics and flow of conversation. To ensure a smooth process, we developed a script with timing notations to ensure adequate and balanced coverage of the topics of interest. In this script, we split the focus groups into four phases of discussion: an introduction, a general discussion about material modeling and uncertainty; presentation of the prototypes; and a wrap-up discussion. Rather than ask participants to dive right into assessing the MMUV prototypes, we decided to prime the discussion by asking the participants to talk about the role of material modeling in their work. In doing so, participants exchanged observations about the importance of material models for engineering analyses; identified key sources of material model uncertainties; discussed the impact of material uncertainty on their work; and described strategies for representing and managing uncertainty. These exchanges set the stage for the second half of the focus group, during which Otahal, Fabian and Potter took turns presenting and discussing their prototype designs with the group participants. As expected, putting prototype designs in front of the experts generated intense discussion about the problem of understanding and managing uncertainty in material models.

In the following sections, we describe the focus groups and summarize key themes. The prototypes are described in Section 3. The participants provided detailed comments on these prototypes and concrete suggestions for enhancing visualization utility and interactivity, as discussed in Section 4. However, the participants' discussions also illuminated the complicated organizational and technical relationships through which information about material properties and performance is exchanged and incorporated into Sandia's engineering research and development; since it provides the context for all of the participants' comments, we begin with an overview of these relationships in the next section.

2 The Distributed Workflow of Material Modeling

The focus group sessions revealed a great deal about the state of practice in material modeling and the use of material models in Sandia's research and engineering domains. Material model development and use exist in a distributed information workflow: a particular material model is developed by material modelers, incorporated into a continuum engineering simulation code by code developers, used by analysts when they run simulations for specific applications, and finally, a decision maker chooses actions that are informed by those simulations. Each of these different groups is a stakeholder in the material model, but their knowledge about the material behavior and their use of the material model vary widely. Naturally, the meaning of material model "uncertainty" also varies widely across these stakeholder roles.

Material model development begins with theories that may be incomplete, contain acknowledged gaps in applicability, or have other known deficiencies. A typical model has a number of parameters that must be determined for each particular material; these parameters are calibrated to match available experimental data, or lacking that, to data from simulations of more fundamental models such as density functional theory (DFT) or molecular dynamics (MD) simulations. Finally, while a material model may have a functional interface that accepts input values and returns output values, in many cases the model is incorporated into the engineering code in precomputed, tabular form and output values are interpolated. For a material modeler, each of these steps is a source of a different kind of error that contributes to the overall uncertainty of the model. Some material modelers in our focus groups were hesitant to even attempt to put a number on the sum uncertainty for any model, because they saw no constructive value in, essentially, guessing. With that said, serious efforts are underway for quantifying model form uncertainty, parameter uncertainty, uncertainty in experimental data, and errors in discretization, and providing this information in a form that is accessible to the engineering simulation codes.

The code developers in our focus groups are programmers, but their educational degrees were PhDs in various engineering disciplines, physics, or applied math. Development and maintenance of engineering simulation codes require a solid foundation in numerical discretization techniques as well as domain-specific expertise. While many are familiar with basic material modeling, code developers typically treat material models as black boxes and focus more on the interface between the models and discretization techniques. One reason for this separation of concerns is that the material model may have been developed at a different place, at an earlier time, or for a different code. Code developers often know some material modelers, but rarely do they know the authors of all the material models included in their code.

The essential concern of code developers is the interaction between the material models and the discretization algorithms. The discretization algorithms assume the EOS possesses certain properties, such as convexity, or a positive speed of sound. Likewise, the EOS assumes it will only be given valid input data, e.g., density and temperature points for which the EOS has been validated. In practice both sets of assumptions are sometimes violated, and occasionally the code crashes and does not produce a solution. In engineering codes, developers address these cases with a number of techniques that vary widely in their theoretical credibility.

Analysts are quite interested in how the uncertainty of the material model affects their simulation results, and less concerned about the material model uncertainty itself. Unfortunately, the effect on each simulation result is highly dependent on the details of the simulation, as well as on the specific material models used. These details determine the various regions of thermodynamic space (for EOS models) that are sampled during the simulation, and the EOS uncertainty varies significantly through the thermodynamic space. For most solid mechanics models, the situation is even more complicated – the response of the material is dependent on the path of a material element through, e.g., stress-strain space. For analysts, material model uncertainty is just one contributor to the overall uncertainty in their simulations. Analysts have minimal access to information about material models and minimal contact with material modelers. A disturbing theme, heard repeatedly, was that analysts have little guidance on choosing among the EOS models available for the same material, and do not have the resources to investigate this source of uncertainty among all the other uncertainties that affect their simulation results.

The dependence of simulation outputs on some material model uncertainties can be examined. The uncertainty of an input (e.g., a material model parameter) is described by a probability distribution, and a variety of techniques are available to propagate a number of input uncertainties to the uncertainty of a simulation output. Nonintrusive methods rely on running a number of simulations for different input values and examining the distribution of the output quantity. Intrusive techniques are also available but are very difficult to retrofit to existing engineering codes. Considering uncertain material models, this approach is effective for examining parameter uncertainties and to a lesser degree, the uncertainties that can be represented by discrete inputs. The effects of model form uncertainty and interpolation error for tabular EOSs on simulation results are not accessible in this approach.

3 PROTOTYPES

To facilitate discussion within our focus groups, we developed four visualization prototypes, each of which present uncertainty within a material model in a unique way. Participants evaluated specific features of each prototype and described scenarios in which different elements could prove helpful. We present each of the prototypes used in the focus groups, using results from a simplified equation of state simulation which produced seven realizations of a material surface.

The Mie-Grüneisen EOS was used to generate seven different realizations of the pressure surface. Parameter values were chosen to represent aluminum. The two parameters in the reference Hugoniot, C_0 and S_1 were viewed as uncertain, and were simultaneously varied within about 1% of the nominal values for aluminum. In some of the prototypes the uncertainty may be rescaled or exaggerated to improve the display. (In a production tool this would be controlled by the user.) Recall that the goal of the MMUV prototypes is to visualize the uncertainty in the material model itself, and not the uncertainty in engineering simulation results. We expect that simulation data (with or without uncertainty information) would be displayed in addition to the material model uncertainty.

The uncertainty in the EOS can be represented in a number of ways. In our example data, we must compute a measure of the uncertainty from the seven surfaces – the uncertainty is not explicitly defined. Knowing that the different surfaces were generated from particular distributions of C_s and S_1 informs the computation of the uncertainty, but this information is not always available. Alternatively, material modelers might provide an uncertainty surface for each thermodynamic variable. For developing visualization techniques, the source or definition of the uncertainty is not important; however, for analysts interpreting the visualizations it is critical to understanding what they are seeing.

3.1 Point Cloud

The first prototype implements a technique presented in [5] that renders a cloud of three dimensional points at a variable distance normal to a surface. The emphasis of this prototype is to show uncertainty in the exact location of the surface. The distance each point is away from the surface is random within a range defined by the amount of uncertainty about the surface location at a particular point. The algorithm creates a cloud of points that are further away from the surface in regions of high uncertainty, and closer to the surface in regions of lower uncertainty. Additionally, the transparency of each point can be varied with the uncertainty, thus points of higher uncertainty become more transparent. This creates a visual effect that feeds the expectation of the human visual system, where regions of low uncertainty appear crisp and solid, and regions of higher uncertainty appear hazy and indistinct. In addition, the points can be colored by another scalar value, such as temperature, and thus simultaneously convey data and uncertainty information.

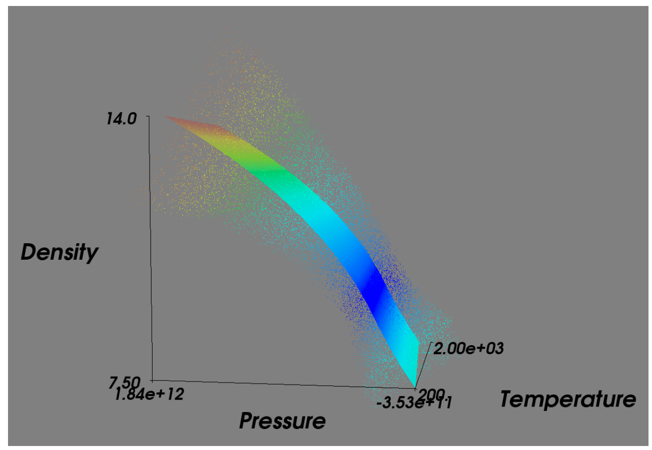

Figure 1: View of the point cloud prototype. The mean surface of the dataset can be seen in the middle of the point cloud. The point cloud is blue and opaque in regions of lower standard deviation and more red and transparent in regions of higher standard deviation, near the top of the figure.

3.2 Surface Animation

The second prototype is based on a technique described by [2] that uses animated visual vibrations of the points defining a surface to show uncertainty in the surface location. The animation draws a fixed semi-transparent surface at the mean and sweeps another solid surface through one standard deviation above and below the mean surface, with the animation transition defined by the sinusoid equation:

$$V = \frac{csin(2\pi pt + \frac{\pi}{2}) + 1}{2} + f \tag{7}$$

where V is the location of the vertex along the surface normal, c is the amplitude of the oscillation, p is the period, f is the floor of the oscillation, and t is time. The sinusoid defines a smooth transition between the floor and amplitude over time for each vertex in the surface mesh. If the floor and amplitude for each vertex corresponds to the uncertainty at that point on the surface, then the viewers eye will naturally be drawn to areas of high uncertainty as the surface animates. Other oscillation functions could be used that cause more rapid transitions between states, such as step and sawtooth functions. Figure 2 shows three frames of the animation.

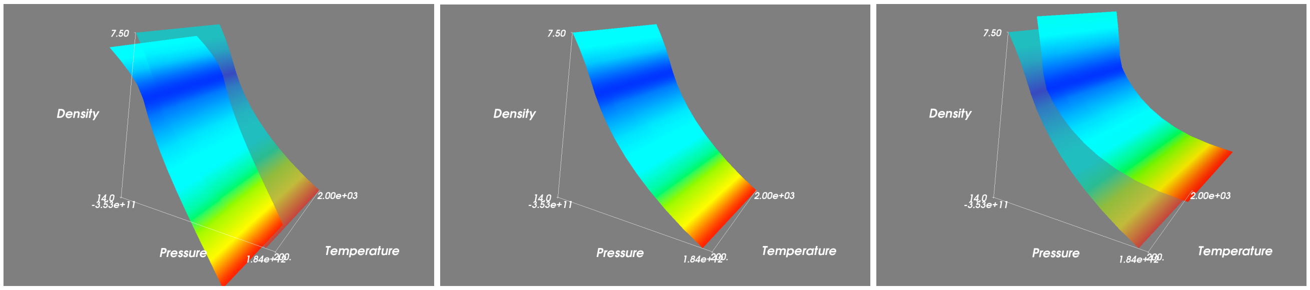

Figure 2: Three frames of the surface animation. The transparent surface represents the fixed mean surface and is visible in every frame for reference. The opaque surface sweeps through a region defined by one standard deviation distance above and below the mean surface. Here, we show the minimal, mean, and maximal position of the animated surface. Both surfaces are colored by standard deviation, with the largest point of deflection away from the mean surface existing in the red area of the bottom surface.

3.3 Bounding Statistics

The third prototype uses statistics similar to the traditional boxplot [10] to bound the valid regions of the simulation. The minimum, maximum, and mean surfaces are calculated point-wise, as well as the standard deviation between all surfaces in the simulation. The user is given control over the display of the statistical surfaces through a graphical interface which also provides options to show each of the simulation surfaces and contextual surfaces such as the mean +/- standard deviation. Data values can be colormapped onto the mean surface and the user may choose which data values are displayed. Figure 3 shows a screenshot of the prototype. The mean surface is shown centrally and flanked by the minimum and maximum surfaces which have reduced transparency to reduce visual clutter. A single surface from the simulation is shown in blue, below the mean surface. The main goal of this prototype is to show the range of possible outcomes, as well as indicate where the data is most likely to reside.

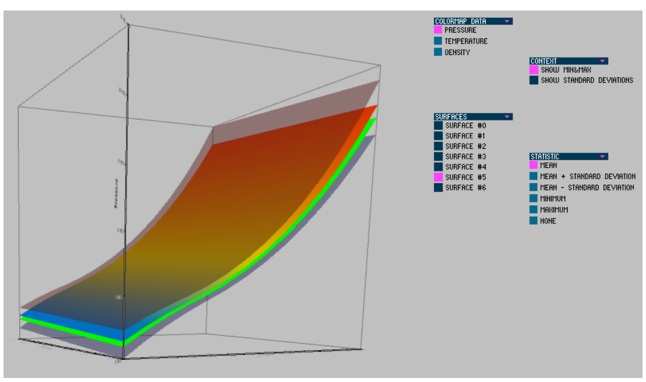

Figure 3: Prototype using bounding statistics and a graphical user interface to explore the dataset. The mean surface is shown centrally along with the minimum and maximum surfaces (with reduced transparency) and a single surface from the simulation (in green). Through a series of buttons, the user can control which statistical surface to display, contextualize surfaces within statistical bounds, and show the original simulation surfaces.

3.4 View Dependent Opacity

The final prototype uses a model similar to the Blinn lighting model [1], in that the view angle is compared with the normal of the surface at each individual point. Instead of using this to modulate the lighting, it is instead used to modulate the opacity of the surface at each point, as shown in Figure 4. At a high level this technique is used to represent a collection of surfaces, each rendered individually with a transparency associated to confidence in the surface. Thus, when the viewer is directly above the surface, the opacity is determined entirely by the surface's transparency, e.g., $\sigma_{x,y}/N$. When the viewer is at an oblique angle, the surface becomes solid, regardless of the underlying surface's transparency. In the case of this prototype, the transparency is established by the standard deviation of the surfaces as a whole at each position and is divided evenly between all surfaces, that is $\sigma_{x,y}/N$, where N is the number of surfaces rendered.

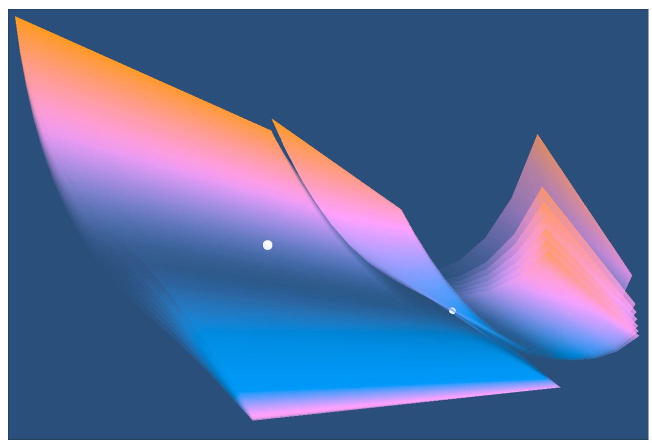

Figure 4: Screenshot of the view-dependent opacity prototype showing a spherical object embedded within two EOS datasets. The opacity of the surfaces making up the datasets depends both on the viewpoint as well as the uncertainty values. Thus, as the surface normal aligns with the view direction, the opacity is reduced with uncertainty to allow the user to easily determine the position of the object and the confidence along the surface. As the view direction diverges, the surface becomes opaque, allowing the user to clearly see the space occupied by all of the surfaces and locate the embedded objects position along the normal of each of the surfaces.

Intuitively, the motivation behind this approach can be understood by considering objects embedded within this collection of surfaces. When viewed from nearly overhead, the surfaces are transparent and objects inside are clearly visible. This allows the viewer to easily determine where the object is positioned on a 2D plane normal to the surfaces. When viewed from the side, because the surfaces become opaque, the embedded object's position along the normal of each of the surfaces becomes more apparent. This approach is not useful when the surfaces are very close together relative to the size of the surfaces.

4 Roles for Material Model Uncertainty Visualization

After reviewing and analyzing the feedback from the focus groups, several themes emerged. In this section we begin with the themes more closely tied to the prototype visualizations, then move to the broader role visualization can play in addressing the needs of the material model stakeholders.

4.1 The different prototypes were suited to different uses

Each visualization prototype has different features. Participants found that the features could be be positive or negative, depending on the use case. The point cloud prototype was effective at showing how the uncertainty varied in different areas of the domain. As intended, the points immediately conveyed a sense that the surface was not known precisely. However, because points represented the uncertainty of the material model, participants thought another mechanism would be needed to show simulation data, which would, most naturally, also be represented by points. A second concern, particularly for material modelers, was that important correlations were lost through the statistical processing of the EOS data – all the individual surfaces were averaged to compute the mean surface, and the point cloud was generated from the standard deviation with respect to the average surface.

The bounding statistics prototype also applied statistical processing, but maintained the original surfaces. While the emphasis of the prototypes was on visualization techniques, participants vocally supported the ability to display or hide the individual surfaces and statistically generated surfaces. The bounding surface technique worked well for this dataset, but participants were not sure how effective it would be for a larger number of surfaces or for multiphase surfaces, which have more geometric complexity.

The surface animation prototype was developed after the others, in response to participants' desire to see individual surfaces and variation within the set, but without overwhelming the viewer with all the information. The view-dependent opacity prototype was less intuitive to participants because initially, it wasn't clear why the opacity changed with the viewing angle. However, the representative point of simulation data allowed participants to grasp the value of seeing the EOS surface, the associated uncertainty, and simulation data at the same time – as the uncertainty increased the surfaces became more transparent, and one would see simulation data more clearly when it moved into an uncertain region, identifying a cause for concern.

4.2 Surface data is not enough

The EOS surface, whether for pressure, energy or another thermodynamic variable, provides an overview of the material behavior, but the surface in and of itself provides just the context. Participants expressed the need for various types of references to connect the surface to their understanding of the material behavior. Axes and contours labeled with numerical values would provide quantitative references. Domain-specific landmarks, such as phase boundaries, Hugoniots, isentropes and isotherms would provide a link back to traditional two-dimensional plots (see below.) For analysts, simulation results must be displayed on the surface to identify the region of interest; and the EOS uncertainty, while highly desired, is only meaningful to analysts in relation to the simulation data. Even when quantified uncertainty is not available, plotting the experimental data to which the model was calibrated would suggest a region of higher confidence to an analyst. For material modelers, simulation data is still informative, but curves with physical meaning (Hugoniots, isentropes, etc.) are more important. A key insight for visualization specialists is that making three-dimensional representations usable, useful, and adoptable across the community is likely to require careful interaction design.

4.3 Two-dimensional graphics remain the standard

The focus group participants were intrigued by the idea of having three-dimensional, interactive visualizations to enhance their understanding of material model dynamics. However, the participants also indicated that two-dimensional plots of, e.g., curves in the pressure-density plane are the existing lingua franca of material model representation; these plots are regularly shown in textbooks and the academic literature on the subject. The stakeholder community is already familiar with this visual form for material model information, so this view leverages previously developed mental models. Two-dimensional plots are relatively information impoverished, given the complexity and uncertainty that characterizes most material models; but they are easy to generate, read, and share. Three-dimensional, interactive visualizations can carry a great deal more information, and participants valued the "big-picture" view of the thermodynamic space. One participant remarked that he would start with a three-dimensional view to get a qualitative feel for the context, but wanted the ability to zoom in to a particular region of interest and see traditional two-dimensional plots because they were "more quantitative". For this application, effective visualization may depend more on the timely delivery of a key existing technique than discovering or developing a new technique.

4.4 Visualization to compare and select material models

As mentioned in the Introduction, material models are inputs to the simulation code; that is, analysts specify the material models before the simulations can be run. Analysts noted that selecting one material model from several possible choices was hindered by a lack of information about each model, and often sought out a colleague or (if available) a material modeler for guidance. Minimally, one would like to know the range of validity of each model. In addition, the intended application, any articles or documents on the theory behind the model, and the experimental data used to calibrate the model would guide the user relative to their own application. Analysts usually have some idea of the densities, temperatures, and pressures a material is likely to experience in their application. Comparative visualization of the different surfaces for those conditions, particularly if uncertainty were displayed for each candidate model, would allow analysts to choose based on quantitative information.

4.5 Visualization to analyze results

The most readily apparent role for visualization is to aid the analysis of the simulation results. According to the focus groups, there are three primary use cases. The first is to "debug" a simulation that has crashed. For some engineering codes, a majority of simulation failures can be traced to the material model; sometimes the material model is a poor model, sometimes the EOS is being sampled outside (perhaps far outside) its range of validity, and sometimes the material model catches an unrelated error committed far upstream. In all these cases, visualization of simulation data on the EOS surface can provide insight and expose patterns in failures that are difficult to identify on a case by case basis.

A second use case is to assess whether or not a simulation relies on EOS information from a region of high uncertainty. The prototypes display EOS models and their associated uncertainty; by adding the time-dependent paths of simulation data, such as element or cell values, or passive tracers used as diagnostics, an analyst could judge whether the EOS uncertainty was important for that particular simulation. For example, an analyst could conclude that while regions of high uncertainty exist in the model, none of the material in the simulation experienced conditions in those regions. Alternatively, they might observe that a critical part experienced a temperature near a phase transition, and the phase boundary is a region of high uncertainty; in this case, more simulations might be run to investigate the effect of crossing the phase boundary.

Finally, visualizing simulation data simultaneously with material model data can provide insight into simulation results. If in the previous example, the phase boundary was a region of low uncertainty, the analyst might modify their design so their part came closer to the phase transition to improve performance of the device. Of course, these are hypothetical examples, but participants had tried to answer similar questions about their own simulations and, if answers were obtainable, they often required a lot of time and effort. The material modelers who participated in the groups indicated that more detailed visualizations could support better understanding of the sources and degree of uncertainty in various regions of material behavior.

4.6 Visualization for Communication

Material modelers, code developers, and analysts all recognized that communication about material models was a weak point in their work. Analysts were concerned that they had little guidance in choosing models for their simulations. Several code developers and material modelers recalled frustration that they could not explain a particular material model issue to an analyst in a way the analyst could understand it; they could not describe a complex structure in their own mental model to someone that did not have the same mental model and domain-specific vocabulary. The visualization of material models can alleviate these issues by providing a common view for the different stakeholders. Visualization can also be used as a training and learning tool, for more experienced practitioners as well as those new to the field. Visualization enables analysts to explain the reasoning and data behind their conclusions to sponsors and decision makers. Finally, material model visualization can be used to help researchers, managers and funding agencies identify gaps in knowledge and prioritize resources to close those gaps.

4.7 Provenance

Our assumption has been that uncertainty is a mathematical object that can be quantified, propagated, plotted, and visualized. But a number of participants identified a clarity about a model's origins and history, or provenance, as highly desirable. The lack of provenance is a subjective uncertainty that cannot be quantified. While quantified uncertainties for a particular model can, in principle, be propagated across the entire distributed workflow, stakeholders would invariably be more comfortable if the provenance of the model were known. Provenance provides subjective confidence not just in the model, but in the quantified uncertainty of the model.

5 Conclusion

The visualization of uncertainty for material models can mean different things to material modelers, analysts, and the code developers that incorporate the models into the engineering simulation codes that the analysts use. The focus group approach revealed a distributed information workflow around the development and use of EOS models, and that information bottlenecks in this workflow had organizational as well as technical origins. Visualization prototypes anchored the discussions and better differentiated the perspectives of the different stakeholders. There is a clear demand for a visualization capability for EOS models and their uncertainties. This capability would improve communication across the workflow, as well as provide an analysis tool for material modelers and, if simulation data can be incorporated, for analysts and code developers. However, the complexity and diversity of material modeling and the abstract nature of uncertainty make the development of an effective tool challenging.

6 Future Work

The overall goal of this work is to develop a software tool for the material modeling community that facilitates communication and exposes the uncertainty throughout the workflow of users and developers of material models. The work presented here is the first stage of this development and has provided numerous insights into the needs of the community. While the prototypes developed were designed mainly to spur discussion within the focus groups, we have developed a list of specific "wants" of the participants, as well as a better understanding of features deemed unnecessary, ineffective, or under developed. From this feedback, we will begin development on a software tool for this community and will iterate one-on-one with select participants on the design. A notable shift in our original plan for the software tool is the disparity in needs for analysts compared to modelers. We are currently unsure as to the feasibility of satisfying the needs of both groups with a single software tool and will thus be revising our approach to reflect the results of the focus groups.

Acknowledgements

The authors deeply appreciate the time and candid discussion contributed by the focus group participants. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. This publication is based on work supported in part by Award No. KUS-C1-016-04, made by King Abdullah University of Science and Technology (KAUST) and an award from DOE NETL.

References

[1] J. F. Blinn. Models of light reflection for computer synthesized pictures. In *SIGGRAPH '77 Proceedings of the 4th annual conference on Computer graphics and interactive techniques*, New York, NY, 1977.

[2] R. Brown. Animated visual vibrations as an uncertainty visualisation technique. *Proceedings of the 2nd international conference on Computer graphics and interactive techniques in Austalasia and Southeast East Asia GRAPHITE 04*, 1(212):84, 2004.

[3] M. C. Brugnoli, J. Hamard, and E. Rukzio. User expectations for simple mobile ubiquitous computing environments. In *Proceedings of the Second IEEE International Workshop on Mobile Commerce and Services*, pages 2–10, Washington, DC, USA, 2005. IEEE Computer Society.

[4] S. Eliezer, A. Ghatak, and H. Hora. *An Introduction to Equations of State: Theory and Applications*. Cambridge University Press, 1986.

[5] G. Grigoryan and P. Rheingans. Point-based probabilistic surfaces to show surface uncertainty. *IEEE Transactions on Visualization and Computer Graphics*, 10(5):546–573, 2004.

[6] T. Hall, S. Beecham, and A. Rainer. Requirements problems in twelve software companies: an empirical analysis. *IEE Proceedings - Software*, 149(5):153–160, 2002.

[7] Z. M. Kasirun and S. S. Salim. Focus group discussion model for requirements elicitation activity. In *Proceedings of the 2008 International Conference on Computer and Electrical Engineering*, pages 101–105, Washington, DC, USA, 2008. IEEE Computer Society.

[8] A. P. Massey and W. A. Wallace. Focus groups as a knowledge elicitation technique: An exploratory study. *IEEE Trans. on Knowl. and Data Eng.*, 3:193–200, June 1991.

[9] J. Nielsen. The use and misuse of focus groups. *IEEE Softw.*, 14:94–95, January 1997.

[10] J. W. Tukey. *Exploratory Data Analysis*. Addison-Wesley, 1977.

[11] N. Van House, M. Davis, M. Ames, M. Finn, and V. Viswanathan. The uses of personal networked digital imaging: an empirical study of cameraphone photos and sharing. In *CHI '05 extended abstracts on Human factors in computing systems*, CHI EA '05, pages 1853–1856, New York, NY, USA, 2005. ACM.