

A Framework for Considering Comprehensibility in Modeling

Michael Gleicher

June, 2016

Author preprint. To appear in the journal *Big Data* with minor changes and improved formatting. Please note that the official version uses different citation numbering.

Please cite as:

Michael Gleicher. A Framework for Considering Comprehensibility in Modeling. *Big Data*. June 2016, 4(2): 75-88. doi:10.1089/big.2016.0007

Abstract

Comprehensibility in modeling is the ability of stakeholders to understand relevant aspects of the modeling process. In this paper, we provide a framework to help guide exploration of the space of comprehensibility challenges. We consider facets organized around key questions: Who is comprehending? Why are they trying to comprehend? Where in the process are they trying to comprehend? How can we help them comprehend? How do we measure their comprehension? With each facet we consider the broad range of options. We discuss why taking a broad view of comprehensibility in modeling is useful in identifying challenges and opportunities for solutions.

1 Introduction

Data-driven mathematical models—broadly, abstractions of data—have had a profound impact on a wide range of problems: they may predict unseen situations, provide compact descriptions of large data, permit inferences about populations based on samples, classify and organize data, or allow generation of new synthetic examples. Data-centric modeling is central to many “big-data” applications. Its impact has been enabled by considerable progress in the techniques for all stages of the modeling process: ever-growing datasets are used as input to sophisticated computational implementations to construct complex models that are subjected to validation and human interpretation. The combination of statistics, machine learning, distributed systems, and data management has led to impressive results. Continued development in these fields, coupled with ever more available data sets, will undoubtedly lead to even more impressive systems that apply complex methods to address important problems.

However, even as data-centric modeling becomes more capable, its potential for truly widespread adoption and impact will be limited by human-centric challenges, for the increasing sophistication of the constructed models will also make them increasingly difficult to understand. The ubiquity and potential power of these systems will raise the importance and value of **comprehensibility**, which we define as *the ability of the various stakeholders to understand relevant aspects of the modeling process*.

The goal of this paper is to provide a framework for considering comprehensibility in modeling to aid in identifying challenges and opportunities. The comprehensibility problem encompasses a wide range of scenarios such as domain scientists trying to discern meaning from their models, machine learning

What are the details of the scenario?

Attempt to abstract away from the details to find general principles of comprehensibility in modeling

Who, Why, Where? Key organizing questions

Mix-and-match stakeholders, goals, and modeling phases to identify challenges and opportunities

Who needs to comprehend?

The Stakeholders

- Developers
- Data Scientists
- Domain Experts
- Audience
- Subjects

Why want to comprehend?

The Goals

- Build Pre-Scriptiveness
- Build Trust
- Improve Performance
- Discover Causality / Build Theory
- Characterize
- Generalize
- Improve Practice

Where in the modeling process?

The Phases of the Process

- Identify Problem
- Gather Resources
- Abstract
- Design
- Build
- Evaluate
- Interpret
- Act
- Disseminate

How do we address the challenges?

Consider initial design patterns and examples

How do we measure comprehensibility?

Identify 3 main categories of approaches

Figure 1: A summary of the framework proposed in this paper. The specific lists for each question are initial organizations to show the broad range of aspects to consider.

developers trying to tune their algorithms, and the general public trying to decide whether to trust a newspaper’s prediction. A broad definition of comprehensibility admits a wide range of potential challenges to address.

A broad view of “the comprehensibility in modeling problem” is valuable because issues and opportunities can arise in so many ways. A broad view may allow for the identification of an unexpected problem, but also might suggest solutions that come from very different places. A pitfall with such breadth is that it requires some structure to organize the range of elements and to help see them in ways that expose their similarities. Therefore, this paper takes the important first step of proposing a framework for considering the range of challenges and opportunities.

The proposed framework is summarized in Figure 1. It embraces the breadth of what comprehensibility in modeling involves with a multi-faceted space with many dimensions to consider. Such a multi-faceted categorization strategy parallels some of those developed for visualization (e.g. [29]). The facets are chosen to view the modeling problem abstractly - without the specifics of a scenario. Each facet poses a question and provides a classification that can help group problems (and potential solutions) together: *Who* is comprehending? *Why* are they trying to comprehend? *Where* in the process are they trying to comprehend? *How* can we help them comprehend? and *How* do we measure their comprehension? These questions are designed to help us abstract away from another question: *What* are the specifics of the scenario?

The key message of this paper is to advocate for a broad consideration of the problem of comprehensibility in modeling. There are two main premises that are part of this. First, is that comprehensibility is an important consideration in modeling. Second, is that the comprehensibility problem is best considered in a broad and multi-faceted manner. This involves both multiple facets (or views), but also considering the breadth of options each encompasses. Such exploration purposefully encourages a holistic view of modeling as a process involving many phases, many stakeholders, many potential goals, and (therefore)

many potentials for comprehensibility problems and solutions.

A secondary aspect of this paper is to advocate more broadly for the considerations of human-centric aspects of modeling. This paper attempts to show that human-centric approaches, such as data visualization, have a wide variety of potential roles in successfully applying data to problems. This paper does not attempt to provide a survey of prior successes, but rather to provide a framework for considering comprehensibility in a broad, multi-faceted manner.

1.1 Tradeoffs in Comprehensibility

In modeling, there are many goals. For most scenarios, accuracy is a primary concern, whether it be how well the output of a predictive model matches reality, or how well a descriptive model summarizes the underlying data. However, accuracy is not the only concern in modeling: efficiency, generalizability, robustness, conciseness, verifiability, and self-consistency are some of the many other qualities that model designers and analysts must consider. These properties often form *tradeoffs* between one another: more complex algorithms may provide higher accuracy but take longer to build; models that use fewer variables may provide lower accuracy but be less susceptible to over-fitting. The proper balance in such tradeoffs depends on the context and needs.

Comprehensibility provides yet another potentially important goal. In different scenarios, its importance relative to other goals (accuracy, robustness, etc.) may vary. Tradeoffs exist. Indeed, the tradeoff between accuracy and comprehensibility is most often raised as an argument against comprehensibility. The tradeoff is often touted as inevitable (for example, [17] who “observe the inherent duality between comprehensibility and accuracy”), usually with a simple argument: more complicated models are capable of representing a richer space of functions and are therefore capable of better representing the underlying phenomena, albeit at a cost of understandability.

However, the tradeoff between comprehensibility and accuracy is not universal. In fact, comprehensibility may actually serve as a pathway to achieve other goals. In some cases (e.g. [30]), simpler (more comprehensible) models may improve accuracy. Other ways to achieve improved performance via comprehensibility are more common. For example, a tool for understanding what is happening inside a complex model can provide insights that suggest improvements (see [32] for a compelling case study).

The potential for tradeoffs with other goals provides an important motivation for understanding comprehensibility in a broad way. When tradeoffs do arise, the challenge of quantifying comprehensibility often makes sacrifices difficult: when one cannot measure how much comprehensibility is being gained for a (more easily measurable) loss in accuracy or computation speed, such a trade can be difficult to justify.

Throughout this paper, we consider making “interventions” to better meet comprehensibility goals. The term is meant to imply a choice made to improve comprehensibility. It might be an action to change an aspect of an existing process, or it could also be that the choice of a standard modeling approach in a situation already achieves good comprehensibility (either by design or by chance). An intervention may be something that trades some other property for comprehensibility, but it also might be a way to achieve comprehensibility without negatively impacting other needs.

2 A Framework for Comprehensibility

The broad definition of comprehensibility in modeling admits a wide space of scenarios. To help organize this diversity, the framework considers a set of basic questions. It suggests three questions to consider first: “*Who* is comprehending?” “*Why* are they trying to comprehend?” And, “*Where* in the modeling process are they trying to comprehend?” Considering these questions helps frame the problem so that we can better answer the two key *how* questions: “How do we help them?” and “How do we know that we have helped them?” (i.e. “How do we measure comprehensibility?”) All of these questions are intentionally abstracted from the specifics of the scenario, which are loosely grouped in the the question “What are they trying to comprehend?”

The idea of using the basic question words to create facets owes from its famous application in journalism, where it is immortalized as “the five Ws.” This was recently applied for analyzing data visualization tasks by Schulz et al. [29], and a similar set of questions is used as a framework for visualization design by Munzner [25]. As in most visualization taxonomies (see [5] for a recent survey) they are focused on building visualization solutions to data interpretation problems, whereas the questions as posed here are meant to help identify what those problems are and what the range of solutions might be (beyond just visualization).

Prior work in the areas of machine learning and modeling has taken a more narrow view on comprehensibility. The machine learning literature more often uses the term “interpretability” and almost always refers to understanding the learned model. For example, Craven [8, p.4] defines comprehensibility based on the representation of the model: “does the learning algorithm encode its model in such a way that it may be inspected and understood by humans.” This definition narrows the focus to the model itself, which is only one phase of the larger modeling process (see Section 2.4). By considering more aspects of modeling we can identify a broader range of problems and solutions. In this paper, we use the term “comprehensibility” rather than “interpretability” to avoid the latter’s connotation of focusing on interpreting the model itself.

2.1 What: Details of the Scenario

Our goal is to provide a general framework for considering comprehensibility in modeling that will apply across a wide variety of scenarios, and allowing us to develop generalizable approaches and methods to address common problems.

There is a strong tradition of trying to abstract data analysis problems. For data visualization, there has been considerable effort to develop abstract ways to consider user tasks (see [5, 29] for recent examples with extensive surveys). The provided taxonomies and categorizations can be helpful in being able to match solutions to problems, to group related solutions, to allow for comparisons, etc. However, they are too focused to provide for the broad view of the modeling process. Addressing comprehensibility in a modeling process may create a range of potential problems that could be addressed by many different types of interventions (including visualizations).

Any specific scenario of modeling involves a number of details: in some specific application or domain, with some specific type of data, for some specific model type built with some algorithm, etc. Choosing which of these details *not* to consider in building a framework is important. For example, a logical first question to ask in modeling is “*What* is being modeled?” While the thing being modeled is certainly important to modeling, many of the issues of comprehensibility apply whether we are modeling collections of literary documents, protein molecules, or financial trends.

2.2 Who: The Stakeholders

Considering comprehensibility implies that there must be someone to do the comprehending. There are many potential stakeholders in a predictive modeling application. Understanding this range of stakeholders is important. Each may have different needs for comprehension in the process and different abilities. The stakeholders in a prediction application may include:

Developers: people building tools and methods meant to work across a range of problems — for example a researcher inventing new algorithms or implementations.

Data Scientists: people who do the model building. This role involves using the general tools and methods, in a way that is somewhat detached from the domain.

Domain Experts - people who “have the problem.” They often commission the model building, and consume the results. In fields that develop general methods, such as visualization or statistics,, such people are often referred to as “domain collaborators.” The key distinction is that these stakeholders are primarily interested in the topic of the data, not the general problem of working with data.

Audience: people who ultimately get the results — for example, the audience of a scientist’s paper or journalist’s article. The definition of audience is clouded by the fact that many different stakeholders may have “audiences,” and in a sense, downstream stakeholders may be the audience of upstream ones (e.g. a scientist is the audience of a tool builder).

Subjects: people who are affected by the model, but will not work directly with it. This includes, for example, patients who may receive improved care because medical practitioners (audience) have learned from researcher’s (domain experts) predictive results, or consumers whose actions are being modeled in marketing research.

In practice there might not be clean separation between roles. One individual might have multiple roles in the process.

Different stakeholders may have different requirements. However, each category of stakeholder can potentially have multiple kinds of goals (see Section 2.3) and concerns about multiple aspects of the modeling process they are interested in (see Section 2.4).

Different stakeholders may have different abilities. One might expect statistical sophistication among data scientists, but not necessarily in a general audience such as in an article for the popular press. Within each category of stakeholder, there is the potential for a range of expertise, skills, abilities, and motivations. Novice developers or domain scientists may have different needs, and be served by different tools, than experts.

While problems (and solutions) are often specific to certain stakeholders, the issues that one group may have, and the solutions designed for them, may have implications for others. Often interventions designed with one group in mind end up helping other groups as well (see Section 3.1 for an example). In addition to considering “who is comprehending” it can be useful to think about “who else.”

Another potential distinction is between model builders and model users. This distinction is often orthogonal to the stakeholder classification listed above. The roles of builder and user may be held by many of the different stakeholder types, and often a single stakeholder may take on both builder and user roles. These two roles also provide ways to think about comprehensibility as there are comprehensibility challenges in each.

In future work, a more complete understanding of the range of stake holders in the modeling process will be valuable because it will better enable an understanding of the range of needs, and for the identification of needs and the design of approaches that address them.

2.3 Why: Goals for Comprehensibility

Comprehension of a model is rarely an end unto itself: comprehensibility is typically desired because it helps in achieving modeling goals. Being able to identify the comprehensibility goal is important in defining a comprehensibility problem, and seeing if a solution is appropriate. If there is no reason for wanting comprehensibility, there is little reason to seek it, and less reason to make tradeoffs that sacrifice other goals.

Understanding potential goals is important. It can help identify and articulate needs. It can help assess tradeoffs. It can provide the basis for assessment. It can help articulate the value of solutions. It can help create an open-mindedness towards solutions by de-coupling them from the real objectives. Craven [8] provides a list of potential goals as an argument for the value of comprehensibility. Although he considers a narrow definition of the problem (restricting it to model representations), his goals (Validation, Discovery, Explanation, Improving Accuracy, Refinement) cover a similar range to the list proposed below.

However, developing an abstract vocabulary of goals is challenging. To the extent possible, such a list should be abstracted from the other facets and aspects of scenarios, complete (in that it covers the range of goals), and non-redundant (so that specific goals fall into few categories). While this initial list of abstract goals may not achieve these aims, it provides a starting point for discussion.

Build Pre-scriptiveness: understanding how to make use of the model and results. Pre-scriptiveness is the quality of a model that allows its user to do something with a result, e.g. its ability to inform action. For example, a weather forecast can be used to choose appropriate clothing. While some “prescriptive models” explicitly consider decision making [4], decision support often relies on interpreting predictive or descriptive models.

Build Trust: understanding a prediction to appropriately trust it, or a predictive process to trust in its ability to make predictions (and to trust those predictions).

Improve Performance: understanding can lead to improvements in many of the other desired properties (e.g. accuracy and efficiency). Understanding can drive iterative refinement that is applied to improve predictive accuracy, efficiency, robustness, and even comprehensibility.

Discover Causality / Build Theory: while the modeling process typically has its main goals of modeling results (e.g. making predictions), it can often have the side effect of shedding light on the underlying process that is being modeled. While a statistical model usually uncovers correlations, not really identifying causality, it can be a useful starting point for theory building, or even an empirical approach towards testing theory. Craven [8] differentiates within the goal of theory building, separating the goals of discovery (finding unknown relationships) and refinement (refining approximately correct domain theories).

Characterize: understanding what the model can do and where it can be applied. For example, even if models do not have an explicit uncertainty characterization for their predictions or an explicit characterization of their operating regions, a deeper understanding of various aspects of the model can help create them.

Generalize: understanding can help extend modeling work to situations beyond its original goals, for example to see that methods might apply in other applications.

Improve Practice: understanding the success (or failures) in one model, modeling application, or modeling process, can help prepare the stakeholder for future applications. The pedagogical value of understanding a particular model can go beyond that specific instance, and help the learner better understand modeling in general.

Many fields, including social and management sciences, consider how people and organizations achieve similar goals in other situations.¹ This may provide a rich source of potential ideas, both for characterizing the range of tasks but also to inspire approaches to support them.

While it is tempting to associate individual goals with specific stakeholders and aspects of the modeling process (Sections 2.2 and 2.4), an important observation is that any goal might be had by a range of stakeholders, and might be addressed in a wide variety of ways.

2.4 Where: The Analysis Pipeline and Its Opportunities

Modeling is a process. The process begins before data gathering and wrangling, moves through the phases of mathematical model building and validation, and doesn’t end until at least the users have had a chance to act on the predictions made. And of course, it may not end there, as reflection, reuse, and revision make the process potentially cyclic. Comprehensibility can be a factor in any of these phases in the process. This means that comprehensibility can mean many different things.

Identifying the phases of the modeling process has been done in many ways. Often fields (such as developers of statistical techniques, or data wranglers) focus on stages they are most involved in. For example, the National Institute of Standards and Technology’s Handbook of Engineering Statistics describes modeling as a 3-phase process (model selection, fitting and validation) [1, Section 4.4.1]. The mathematics education literature sometimes provides a more comprehensive enumeration of the phases. For example, Anhalt and Cortez [3] provide a list of six steps (Analyze Problem; Develop and Formulate Model; Compute Solution; Interpret / Draw Conclusions; Validate Conclusions; Report). It is not a coincidence that this more comprehensive view comes from the community focused on how people learn

¹I thank the anonymous reviewer who suggested this connection, and provided some initial pointers into the literature.

(and understand) the modeling process: stakeholders are involved in many different activities, and all of these are likely to be important to learn (and, by extension, to understand).

The data mining community has made many attempts at trying to define a process model that identifies the common steps (see [21] and [19] for surveys). These models often include a broader range of steps, emphasizing that early aspects (e.g. planning) and late aspects (e.g. usage) are often critical to success of the overall process. While these different process models divide the steps differently, they share a common range of steps. The common CRISP-DM model breaks the process into six steps: business (domain-problem) understanding, data understanding, data preparation, modeling, evaluation, deployment.

Taking a fully broad view of data-centric modeling exposes a very broad range of activities in building and using a model, leading to a longer list of phases..

Identify the problem.

Gather Resources to address the problem. This includes data collection

Abstract the problem.

Design the model. This encompasses the “formulation” phase where the choices about the model, such as the methods to be used, are made.

Build the model. This is the computation (for example fitting) of the specific model.

Evaluate the model. This includes validation, and whatever other kinds of testing and checking might be done to assess the quality of the model.

Interpret the results.

Act on the results.

Disseminate the results.

The length of the list emphasizes the diversity and breadth of tasks. While the specifics of how the process is divided is more a matter of naming, the important lesson is that modeling involves a range of steps, from planning through dissemination, and that any stage may be the source of a comprehensibility problem, or an opportunity for an intervention (often to address an issue in a different phase).

A different way to think about the modeling pipeline is to consider the objects used by the processes, rather than the tasks (this was my approach in [12]). Each provides a tangible thing that someone might need to comprehend.

Inputs and Resources: Does the stakeholder understand the data, assumptions, and initial questions being used to build the model? For example, we might have a collection of training data or a causal theory that leads to a mathematical model.

Methods: Does the stakeholder understand the method used to build the model from the data (e.g. the algorithms)? To illustrate, in a linear support vector machine learning approach to convert training data into a classification model, the method comprises both the general method (i.e. linear SVM) and the specific implementation of that method (e.g. a particular machine learning software package). In creating a model from some underlying principles, the method might be a symbolic derivation.

Model: Does the stakeholder understand the model used to make predictions? For example, this often takes the form of the specific equations that compute predictions. For the SVM example, the model is the set of coefficients (the linear equation) created. As mentioned earlier, the notion of interpretability in machine learning tends to focus on this phase.

Outputs: Does the stakeholder understand specific predictions made by the model? The output is the prediction computed for a specific condition. It may include more than the prediction itself. For example, predictions often include confidence values, strength scores, uncertainty quantification, or sensitivity analyses. Output comprehension includes both the semantics of the result (e.g. what

does the prediction of a specific value for the predicted variable mean) and the mathematical meaning of the prediction. For a common example, a weather prediction (“20% chance of rain tomorrow”) may say little about any specific time in the day, the amount of rain that might fall, or whether it will be sunny.

Experiments: Does the stakeholder understand the results of experiments run on the model? Experiments include predictive performance evaluation (e.g. cross-fold validation or holdouts) and efficiency tests (e.g. profiling). The results of experiments can include the outputs over a set of cases (often with known ground truth), but also various statistical information.

Presentations: Does the stakeholder understand the artifacts used to present the results? This might include things such as writing that describes the results, figures, videos, spoken presentation (live or recorded), etc.

2.5 How (1): Pathways to Comprehensibility (How do we help?)

The big question for visualization research and practice is “How can I intervene to address comprehensibility in modeling?” The previous questions help us identify problems (or opportunities) and categories what solutions might address. They are intentionally abstracted from what the solutions actually do to address the problems with a goal of being open to a range of potential approaches to interventions.

Several research areas, including visualization, machine learning, and various domain sciences, provide examples of interventions that seek to address comprehensibility issues. An even broader range of work may address such problems implicitly, for example by providing better methods.

From examining a selection of representative examples, some general strategies for comprehensibility interventions begin to emerge. I provide this initial list of “design patterns for comprehensibility” to show the potential for such generalized strategies and to give some hint at the range of potential solutions that are likely to be developed. Cataloging and characterizing the range of comprehensibility interventions, in addition to creating new ones, will be important future work.

2.5.1 Aggressive Regularization

Many model building techniques explicitly encode a tradeoff between how well the model fits the training data and some measure of the representation of the answer. In optimization parlance, this is known as *regularization*. Regularization augments a primary optimization objective (e.g. how well a learned model fits its training data) with a secondary objective (e.g. the magnitude of the coefficients, the number of coefficients, or the depth of the decision tree). By penalizing answers whose representations are less desirable (by the secondary metric), better overall answers may be achieved even if the optimal value of the primary objective is not found. This secondary objective is known as a regularization term, and the amount that it is considered (relative to the primary objective) is often a tunable parameter that we might consider the degree of regularization.

Many machine learning and statistical modeling algorithms provide a parameter to tune the degree of regularization. Such parameters are often used as a tool to control over fitting. Since the metric used for regularization is often a proxy metric for comprehensibility, control of the amount of regularization can be a way to control the accuracy vs. comprehensibility tradeoff. Aggressive regularization, using large amounts to obtain simpler answers, is therefore a strategy for improving comprehensibility. A common form is to use algorithms where the regularization terms prefer sparsity (e.g. $L1$ norms in support-vector machines, LASSO, or logistic regressions).

Aggressive regularization can be straightforward to implement, as regularized modeling techniques are well established and provide a very direct control over the accuracy vs. simplicity tradeoff. However, this strategy has a number of pitfalls. Foremost, it relegates comprehensibility to a tradeoff, rather than trying to find ways to achieve comprehensibility along with other goals. Second, it assumes the regularization metric is a good measure of comprehensibility — when in practice, it is (at best) a proxy metric. Related, it is limited to proxy metrics that have been considered important by algorithm developers.

While algorithms for achieving sparsity (minimizing the number of variables) are common in machine learning and statistics, algorithms for other potential simplicity objectives are less common (see [11] and [10] for algorithms that also try to find simple coefficient values). Also, existing algorithms are often not designed for aggressive regularization, instead being tuned for smaller amounts of regularization for other purposes (like stabilizing the computation). For example, in our examples with L_1 -regularized support vector machines (SVM) [11], we found that standard SVM solvers could not achieve extreme regularized solutions (finding optimal sets of a very small number of variables).

Aggressive regularization is an attractive intervention for addressing comprehensibility concerns as it provides control over the tradeoff between accuracy and comprehensibility (or, to be more precise, properties thought to correlate with comprehensibility). Key future work includes developing more direct connections from regularization terms to comprehensibility and understanding how stakeholders can utilize control over tradeoffs to address their needs.

2.5.2 Model Rejections

Once a complex and successful model is built, that model can be used as a basis to build other models with simpler forms. An example is to first build a neural network classifier for a problem, and then to use this model to build a classifier using a different modeling strategy that is thought to be easier to comprehend, such as a rule-based or decision-tree-based classifier. While such approximation often results in poorer performance, it does not need to do so. Potentially, information gained from the process of building, using, and exploring the initial model can allow for constructing simpler models that could not have been constructed without these insights.

A canonical form of model reprojection is *rule extraction*: building decision rule classifiers from other kinds of classifiers that are considered to be more complex to interpret. Rule extraction has been explored extensively in the machine learning community (see [17] for an extensive survey). Rule extraction generally assumes that decision rules are desirable from a comprehensibility point of view, however.

Because the simpler models often have worse performance, it may make sense to use the original (complex) models for applications where accuracy matters, and reserve the simplified models for situations where interpretation is more central. [18] suggests constructing an ensemble of models at various levels of tradeoff, so that appropriate ones can be used.

Another reprojection strategy is to use aspects of a complex model as inputs to a second modeling process. For example, a first model construction may identify a set of features that are used in the training process for a more interpretable form of classifier. A concrete example is presented by Stiglic et al. [30] who use interaction analysis to identify new features that involve multiple input features and then use these features to derive simple trees. More commonly, complex neural nets can be interrogated to understand how their initial layers serve as feature finding filters — either as recognizable filters [9] or specialized detectors [22]. Indeed, some network learning approaches are specifically designed to discover structure in data [15].

Reprojection is an attractive strategy for comprehensibility interventions because it allows using a model that is otherwise successful yet addressing comprehensibility concerns. It affords the potential for designing new modeling methods specifically for their comprehensibility properties, with less concern for their ability to achieve accuracy goals. Important future work includes both how model type influences comprehensibility, as well as how to make the inherent tradeoffs transparent to stakeholders.

2.5.3 Visualizing Complexity

The previous two approaches focus on finding simpler models. An alternative is to simply accept the richness of a given model and instead find ways to better present and explore it. Such a strategy might be employed with a variety of complex objects in the modeling process. For example, specific approaches might visualize the internal representations of a deep neural network or enable visual exploration of complex validation experiments.

With a complex model, there is a dichotomy of strategies for visualization: visualizing the internal structure of the model vs. presenting the relationships encoded by the model. For example, [31] provides visualizations of neural network architectures (internals) while [7] derives visualizable representations of the relationships (sensitivities) between model inputs and outputs.

For models with very complex internal structure, direct visualization may not be possible. For example, with neural networks, visualizing the network structure may be possible for simpler networks [31], or connections between input nodes and data patterns can be seen by qualitative inspection [9]. However, as the networks grow larger and deeper, interpretation requires some degree of translation. In order to “look inside the mind” of a complex neural net (to use the terms of [16]) techniques may be needed to construct more interpretable representations. For example, [16] introduced a strategy of using sampling to infer what kinds of inputs would be required to activate different internal components. This strategy has the advantage of relating internal structures to (hypothetical) input examples, which are more likely to be familiar than activation patterns. Other methods for approximating these relationships have been developed (e.g. [20]), and the general approach of sampling to look inside a network has been empirically shown to be robust [9]. Similarly, extant methods for interpreting random forest classifiers focus on input and output relationships, rather than model internals [13].

[6] provides an example where examination of complex internal structures is successful. They allow for the exploration of collections of examples of how natural language methods represent sentences. This exploration can be used for performance improvement as it allows a developer to identify what phase of processing errors occur, and to identify opportunities to engineer new features.

A different approach to visualizing complexity is to build visual and interactive techniques that connect complex models to the users needs and workflows. An example of such an approach is the Serendip system [2] for interaction with topic models of text corpora. The system does not hide the complex internals of the models used, but does provide a connection between relevant elements and specific user tasks.

Effective communication of complexity, through visual representations and interaction, is attractive because it makes direct use of models that may be desired for other reasons (such as performance). As visualization and visual analytics methods mature to provide approaches that scale to increasing complexity, opportunities to apply this work to complex modeling applications will also improve. Important future work includes understanding how to apply general capabilities for presenting complex data to specific comprehensibility challenges.

2.5.4 Interaction around Model Construction

[24] surveys many ways that systems can enable users to interact with modeling algorithms where the algorithm itself is treated as a black box. Their work characterizes a space of opportunities for interactions *around* a model building step, such as changing inputs, altering parameters, or examining intermediate results. They do not focus on any particular task — in particular, improving comprehensibility. However, many of their strategies achieve modeling goals by a user’s extant comprehension of various aspects of the modeling process.

The space of interactions mapped out by [24] suggests a wide range of potential strategies for working with models. Many have potential applications for comprehensibility interventions. For example, they suggest “user subsetting” as an approach to achieve faster iterations in the modeling cycle. This approach might also be explored as a strategy to comprehensibility as models built on subsets may be easier to understand, and the understanding built on the reduced models may be scaled to more complex cases.

Important future work includes understanding how the range of modeling interactions might serve the range of comprehensibility needs.

2.6 How (2): Metrics (How do we Assess?)

In discussing comprehensibility, an inevitable question is “How do we measure it?” If comprehensibility is perceived to be a problem, then there is a sense of not having enough of it. If a comprehensibility intervention is proposed, we may want to assess how much more understanding is achieved. When different modeling approaches are compared, we may want to know which provides more comprehensibility. Objective measurement can be particularly important when tradeoffs are involved. While quantifying a loss in accuracy or efficiency can be easy, it is much harder to quantify the comprehensibility gained in its place. How much comprehensibility is gained by reducing a model by two variables? Is this worth a 2% decrease in accuracy?

Unfortunately, assessing comprehensibility seems quite challenging. The breadth of what comprehensibility might mean, the human-centric and often subjective nature of it, and the relative lack of attention paid to it, all complicate the development of mechanisms for comprehensibility assessment. To date, our toolbox of approaches is not very well developed. However, in the limited existing literature three main strategies seem to be emerging: direct measurement of comprehension, direct measurement of goals, and proxy metrics.

2.6.1 Direct Measurement: A direct measurement of comprehensibility tries to assess the stakeholders understanding. Such measurements are hard because they seek to measure what is happening for the stakeholder internally, and because the metrics are often not well defined. However, direct measurement is possible, for example by asking questions (either objective or subjective), or even through biometric assessment.

Assessment strategies from the learning sciences and human factors may provide a source of tools for direct measurement of comprehensibility. Already, measurements such as cognitive load and mental effort have been adapted to measure model quality in other domains [23, 33].

2.6.2 Goal Driven Metrics: A goal driven metric of comprehensibility measures the effects of comprehensibility, ideally in assessing how well its goals are being achieved. For example, if comprehensibility is sought to help improve accuracy of a model, then the accuracy gains provide a yardstick. While such measurements are attractive because they get at the real reasons for wanting comprehensibility, they cannot always attribute the results or provide ways to see other collateral costs and benefits.

Goal driven metrics are most obvious for easily quantifiable goals, such as accuracy or efficiency. However, there is the potential to turn other goals into metrics as well. For example, insight quantification approaches [26, 27] could measure how well a model supports hypothesis formulation (building theory).

2.6.3 Proxy Metrics: The common metrics for comprehensibility are measures of the underlying model itself: for example the type of model (e.g. a classifier might be a decision tree, a support vector machine, or a neural net) and measurements of the model (e.g. the depth of a decision tree, or the number of non-linear coefficients of a linear function). Such measurements do not measure comprehensibility *per se*, but rather measure model properties that are thought to correlate with comprehensibility (e.g. decision trees are thought to be easier to understand than support vector machines, and shallow trees are thought to be easier to understand than deeper ones).

Proxy metrics are attractive because they are simple, quantitative, and are often well-studied for other reasons such as to reduce overfitting. For example, the modeling communities (e.g. machine learning and statistics) have extensively explored the sparsity of models (having few non-zero coefficients). While they tout comprehensibility as one of the benefits of sparsity, they also seek sparsity for other reasons such as better generalizability and more efficient computation. This, in turn makes proxy modeling convenient *for* comprehensibility: the existence of good methods for creating and using models that optimize the proxy metrics mean that models can be created with these properties should we desire them for comprehensibility purposes (see Section 2.5.1).

There are many pitfalls to proxy metrics. Foremost, their impact on comprehensibility is not well un-

derstood or quantified. While the basic principles behind the metrics are often intuitive, a more sound understanding of them may still be helpful. For example, the oft-cited metric of sparsity is probably not incorrect: equations with fewer variables are probably easier to understand than those with more variables. However, empirical evidence, or richer cognitive rationale, for the impact on sparsity is lacking. Quantifying the benefits of 3 variables instead of 4 is still challenging. Similarly, Craven [8] suggests *syntactic complexity* as a metric for measuring the comprehensibility of rule-based models. This metric is defined based on other proxy metrics (e.g. tree depth and number of feature references), and supported based on cognitive arguments that smaller models are preferable.

Another issue with proxy metrics is that the most convenient ones are often explored to the exclusion of others. While sparsity is often discussed, others are rarely considered. However, in terms of interpreting a model, other aspects of the terms may be important. For example, having simple values for the coefficients (e.g. small integers) and choosing familiar variables (as opposed to less familiar ones) may be as important as having small numbers of variables (see [11] for a discussion).

3 Examples: The value of the broad view

The previous sections introduced a framework for considering comprehensibility in modeling. The framework attempted to be broad: considering a number of facets of problems, and taking a holistic notion of the range of modeling in each. This section examines some examples from my own work to show the benefits of such a broad view.

3.1 Validation Experiment Visualization

By considering the whole modeling process, we can identify a broader range of opportunities to address comprehensibility issues, and these interventions may help achieve a range of goals. As an example, consider the validation experiments used to test and assess models. While such experiments are an essential part of a modeling application, and are often complex and rigorous, the analysis of their results is often limited. For classification (machine learning) applications, the results are typically reported as a single number (e.g. accuracy, precision, recall, F1, etc.) or a small set of these metrics.

In [28], we introduced a visual analytics approach for examining the results of protein surface classifiers, shown in Figure 2. We sought to aid our structural bioinformatics domain collaborators (**who**: data scientists, model builders) in improving the performance of a classification system they had built (**why**: improve performance). While this **who** and **why** are standard, we chose a non-standard **where**: the validation experiments. Because our collaborators were simply using summary statistics of the experiment, this phase seemed to be under-exploited. While our assessment in this work was informal, the specific goal gave us a success criterion: if our collaborators were able to gain insights that helped them devise performance improvements. Our approach was to provide a system that allowed for overview examination and detail exploration of the entirety of the results of an experiment, which is complex because it involves dozens of hundreds of molecules, each of which is a complex 3D shape on which the specific classification examples were made (**how**: embrace and present complexity).

The surface classifier validation experiment viewer was able to achieve its initial goal of helping data scientists improve the classification performance of their system. It also served other roles. Our collaborators used the system to show their results to their audience (structural biologists interested in understanding how these molecules perform their function). By observing patterns in where classification failed, an audience in a talk was able to hypothesize about alternate (biochemical) mechanisms that might be involved in this subset of situation (**who else**: audience, **why**: Discover Causality /Build Theory).

A lesson of this example is that comprehensibility challenges, and therefore opportunities can occur in many phases of the modeling pipeline. By mixing and matching different stakeholders, goals, and pipeline phases, we can identify potential interventions. In this case, our intervention was specialized (a system for examining validation experiment results on molecular surface classification data), but it

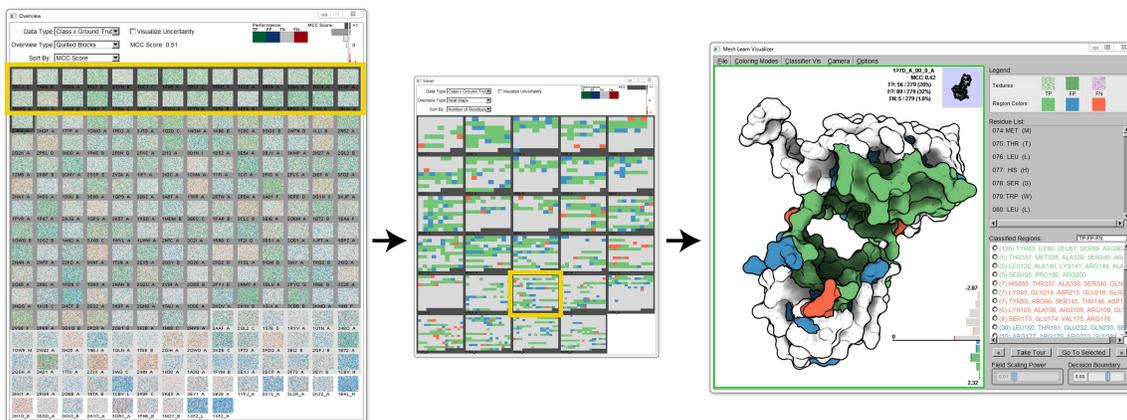


Figure 2: Visualization of a validation experiment for a DNA-binding surface classifier that allows exploration of classification results. The corpus overview (left) is configured to display each molecule in the test set as a quilted glyph and orders these glyphs by classifier performance to show how performance varies over the molecules. Those proteins that appear more green have more true positive classifications, whereas those molecular that appear more red or blue have more misclassifications (false negatives and false positives, respectively). Selected molecules (left, yellow box) are visualized as heatmaps in a subset view (middle) and ordered by molecule size to help localize the positions of errors relative to correct answers. The detail view (right) shows a selected molecule to confirm that most errors (blue, red) are close to the correctly found binding site (green).

suggests a more general class of tools.

[14] also presents a system for visualizing model errors over a test corpus aimed at identifying problematic training examples, problematic features, and opportunities for specializing models to subsets of the operating region.

3.2 Repurposing Classifiers

By considering comprehensibility as a first-class objective, we can re-use modeling technologies for different purposes. For example, classifiers are models that are typically used to classify — their standard use cases involve identifying unseen objects. However, when classifiers are constructed with comprehensibility in mind they can be used for other purposes, such as helping domain experts build theory, generate hypotheses, and organize data.

In [11], we explore such an application. Our initial goal was to help literature scholars (**who**: domain expert) generate hypotheses and identify interesting examples for detailed examination (**why**: Discover Causality /Build Theory). Specifically, they were interested in understanding the relationship between measurable features of documents (words that can be tagged as having certain rhetorical purposes) and higher-level features of the documents (such as its author, date, or genre). We sought to help them interpret the models that were being built (**where**: build model), with the idea that the internals of the model may expose connections, showing what a machine might use to classify the documents (that the scholars would classify using far different information). Unfortunately, most models were difficult to interpret, so we chose to introduce new modeling techniques that traded accuracy for comprehensibility (**how**: aggressive regularization).

We used the term *Explainer* for a classifier constructed with the purpose of being interpretable, at the expense of accuracy. An example is illustrated in Figure 3. The work showed many of the issues that must be addressed to realize such an approach. We observed that while sparsity (limited numbers of

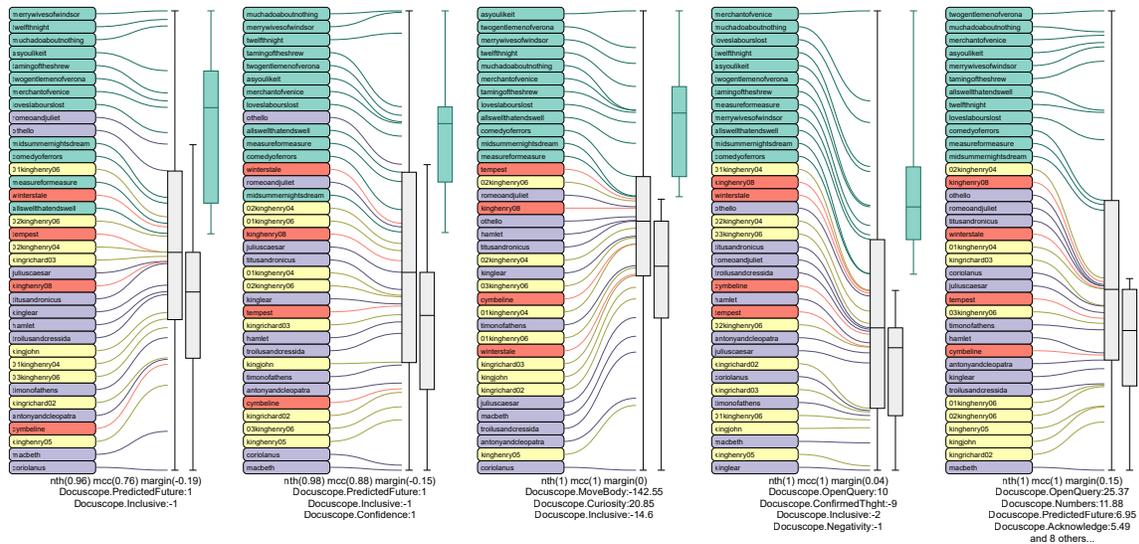


Figure 3: Visualization of example *Explainers*, classifiers constructed with tradeoffs that emphasize comprehensibility concerns [11]. In this example, Shakespeare’s 36 plays are measured with a set of 115 “Docuscope” features. Classifiers are constructed to identify the 12 comedies (green). Each column represents a linear support vector machine (SVM) classifier, with the plays sorted according to their score. The leftmost classifier uses only two features with unit coefficients. It makes several mistakes (for example, misclassifying the tragedies *Othello* and *Romeo and Juliet* as comedies), but the simplicity of the classifier makes it useful for building theory about how Shakespeare used the linguistic constructs in the different genres. In contrast, other classifiers may use more features and more complex weights to achieve better accuracy (and larger SVM margins), at the expense of how easy the functions are to comprehend.

variables) was valuable for interpretability, other properties were valuable as well, including simple coefficients (small integers) and preferring familiar variables. We observed the importance of providing the user control over the tradeoffs between the various objectives (e.g. accuracy vs. sparsity), which required developing algorithms that supported the appropriate control. We observed the importance of methods for showing the resulting models so that a user could both assess its quality (in terms of accuracy) and interpret its structure. While the initial explainers paper [11] only provided initial answers to these questions, it showed the viability of the approach by demonstrating methods for these various aspects.

Because the classifiers were simple, usually involving only two variables, they were easy to present to the (often skeptical) audience (**who else:** audience). Individual examples could be examined in detail, allowing the audience to learn about the process used (particularly, a specific choice of feature type), which helped to establish trust in the approach). However, this required development of different ways to show the results as audiences had different abilities and goals for interpreting presentations of the data.

4 Conclusion

There are many examples of work that identifies, or intervenes to address, issues of comprehensibility in modeling. These success stories are often specific: identifying a scenario with a specific domain, a specific modeling approach, a specific comprehensibility problem, and a specific intervention. A framework can help organize these examples around the various facets of the overall problem of comprehensibility in modeling. It can help identify commonalities to help generalize, as well as to help identify un-met problems and unexploited opportunities.

This paper has introduced a multi-faceted framework for considering comprehensibility in modeling. It suggests a variety of views (who, why, where, and how). Keeping a broad view of the range of answers to each question is valuable. Any stakeholder (from data source to audience) may have a comprehensibility issue. Any phase of the modeling process, from initial planning through final dissemination of results, may provide a challenge or opportunity for intervention. Many goals may be helped by comprehensibility, and a range of approaches for interventions might help address issues.

The initial framework provided in this paper may be extended by providing additional facets, or by better characterizing the range of possible answers within a facet. However, even in its current form, the framework can serve a role in the more important effort to identify comprehensibility issues, develop interventions that address them, and help modeling address even more challenging and important problems.

This paper has made an assertion that human-centric considerations of comprehensibility are applicable across a very broad range of modeling, and the very broad range of stakeholders and tasks that use modeling. Admittedly, support for this argument has been provided with limited examples across a narrow range. Moving forward, the key future work is to develop our arsenal of interventions to address the broad range of comprehensibility concerns. This includes both collecting and characterizing the existing approaches, as well as developing new ones.

The key lesson from the presented framework is that comprehensibility in modeling has numerous facets, each covering a broad range. By taking this broad view, e.g. considering the entire range of potential stakeholders or the range of steps in the modeling process, more challenges and opportunities for solutions can be identified.

Acknowledgments

This work was supported in part by NSF award 1162037, NIH award 5R01AI077376-07, and a grant from the Andrew Mellon Foundation. I would like to thank the many people who have given me feedback as I developed these ideas. In particular, I would like to thank Remco Chang for his help throughout the process, and Ross Maciejewski, Eric Alexander and Deidre Stuffer for their help in articulating the ideas in the paper.

Author Disclosure Statement

No competing financial interests exist.

Author preprint. To appear in the journal *Big Data* with minor changes and improved formatting. Please note that the official version uses different citation numbering.

Please cite as:

Michael Gleicher. A Framework for Considering Comprehensibility in Modeling. *Big Data*. June 2016, 4(2): 75-88. doi:10.1089/big.2016.0007

References

- [1] NIST/SEMATECH e-Handbook of Statistical Methods. Technical report, 2012.
- [2] E. Alexander, J. Kohlmann, R. Valenza, M. Witmore, and M. Gleicher. Serendip: Topic model-driven visual exploration of text corpora. In *2014 IEEE Conference on Visual Analytics Science and Technology (VAST)*, pages 173–182, 2014.

- [3] C. O. Anhalt and R. Cortez. Mathematical Modeling: A Structured Process. *The Mathematics Teacher*, 108(6):446–452, 2015.
- [4] J. Bertolucci. Big Data Analytics: Descriptive Vs. Predictive Vs. Prescriptive. *Information Week*, December 2013.
- [5] M. Brehmer and T. Munzner. A multi-level typology of abstract visualization tasks. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2376–85, 2013.
- [6] J. Chuang and R. Socher. Interactive Visualizations for Deep Learning. In *Visualization for Predictive Analytics Workshop*, 2014.
- [7] P. Cortez and M. J. Embrechts. Opening black box Data Mining models using Sensitivity Analysis. In *2011 IEEE Symposium on Computational Intelligence and Data Mining (CIDM)*, pages 341–348, 2011.
- [8] M. Craven. *Extracting Comprehensible Models from Trained Neural Networks*. Ph. D. Dissertation, University of Wisconsin - Madison, 1996.
- [9] D. Erhan, Y. Bengio, A. Courville, and P. Vincent. Visualizing Higher-Layer Features of a Deep Network. Technical report, Department IRO, University of Montreal, 2009.
- [10] S. Ertekin and C. Rudin. A Bayesian Approach to Learning Scoring Systems. *Big Data*, 3(4):267–276, 2015.
- [11] M. Gleicher. Explainers: expert explorations with crafted projections. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2042–51, 2013.
- [12] M. Gleicher. Position Paper: Towards Comprehensible Predictive Modeling. In *Visualization for Predictive Analytics Workshop*, 2014.
- [13] H. F. Golino and C. M. A. Gomes. Visualizing Random Forest Prediction Results. *Psychology*, 5(19):2084–2098, 2014.
- [14] D. Gotz and J. Sun. Visualizing Accuracy to Improve Predictive Model Performance. In *Visualization for Predictive Analytics Workshop*, 2014.
- [15] G. Hinton, S. Osindero, M. Welling, and Y.-W. Teh. Unsupervised discovery of nonlinear structure using contrastive backpropagation. *Cognitive science*, 30(4):725–31, 2006.
- [16] G. E. Hinton, S. Osindero, and Y.-W. Teh. A fast learning algorithm for deep belief nets. *Neural computation*, 18(7):1527–54, 2006.
- [17] J. Huysmans, B. Baesens, and J. Vanthienen. Using Rule Extraction to Improve the Comprehensibility of Predictive Models. *SSRN Electronic Journal*, 2006.
- [18] U. Johansson, R. Konig, and L. Niklasson. Automatically balancing accuracy and comprehensibility in predictive modeling. In *2005 7th International Conference on Information Fusion*, volume 2, page 7 pp., 2005.
- [19] L. A. Kurgan and P. Musilek. A survey of Knowledge Discovery and Data Mining process models. *The Knowledge Engineering Review*, 21(1):1–24, 2006.
- [20] H. Lee, R. Grosse, R. Ranganath, and A. Y. Ng. Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations. In *Proceedings of the 26th Annual International Conference on Machine Learning - ICML '09*, pages 1–8, New York, New York, USA, 2009.
- [21] O. Marban, G. Mariscal, and J. Segovia. A Data Mining & Knowledge Discovery Process Model. In J. Ponce and A. Karahoc, editors, *Data Mining and Knowledge Discovery in Real Life Applications*, chapter 1. I-Tech Education and Publishing, 2009.
- [22] A.-R. Mohamed, G. Hinton, and G. Penn. Understanding how Deep Belief Networks perform acoustic modelling. In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4273–4276, 2012.

- [23] D. Moody. Cognitive Load Effects on End User Understanding of Conceptual Models: An Experimental Analysis. In A. Benczúr, J. Demetrovics, and G. Gottlob, editors, *Advances in Databases and Information Systems: LNCS 3255*, volume 3255 of *Lecture Notes in Computer Science*, pages 129–143. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004.
- [24] T. Muhlbacher, H. Piringer, S. Gratzl, M. Sedlmair, and M. Streit. Opening the Black Box: Strategies for Increased User Involvement in Existing Algorithm Implementations. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1643–1652, 2014.
- [25] T. Munzner. *Visualization Analysis and Design*. CRC Press, 2014.
- [26] C. North. Toward measuring visualization insight. *IEEE Computer Graphics and Applications*, 26(3):6–9, 2006.
- [27] P. Saraiya, C. North, and K. Duca. An insight-based methodology for evaluating bioinformatics visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 11(4):443–56, 2005.
- [28] A. Sarikaya, D. Albers, J. Mitchell, and M. Gleicher. Visualizing Validation of Protein Surface Classifiers. *Computer Graphics Forum*, 33(3):171–180, 2014.
- [29] H.-J. Schulz, T. Nocke, M. Heitzler, and H. Schumann. A design space of visualization tasks. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2366–75, 2013.
- [30] G. Stiglic, P. Povalej Brzan, N. Fijacko, F. Wang, B. Delibasic, A. Kalousis, and Z. Obradovic. Comprehensible Predictive Modeling Using Regularized Logistic Regression and Comorbidity Based Features. *PloS one*, 10(12):e0144439, 2015.
- [31] F.-Y. Tzeng and Kwan-Liu Ma. Opening the Black Box - Data Driven Visualization of Neural Networks. In *VIS 05. IEEE Visualization, 2005.*, pages 383–390, 2005.
- [32] M. Zeiler and R. Fergus. Visualizing and Understanding Convolutional Networks. In D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, editors, *ECCV 2014*, volume 8689 of *Lecture Notes in Computer Science*, pages 818–833, Cham, 2014. Springer International Publishing.
- [33] S. Zugal, J. Pinggera, H. Reijers, M. Reichert, and B. Weber. Making the case for measuring mental effort. In *Proceedings of the Second Edition of the International Workshop on Experiences and Empirical Studies in Software Modelling - EESSMod '12*, page 1, New York, New York, USA, 2012.