Risk Analysis, Vol. 39, No. 9, 2019 DOI: 10.1111/risa.13328

Perspective

Rethinking Resilience Analytics

Daniel Eisenberg,[1] Thomas Seager,[2] and David L. Alderson[1,*]

The concept of "resilience analytics" has recently been proposed as a means to leverage the promise of big data to improve the resilience of interdependent critical infrastructure systems and the communities supported by them. Given recent advances in machine learning and other data-driven analytic techniques, as well as the prevalence of high-profile natural and man-made disasters, the temptation to pursue resilience analytics without question is almost overwhelming. Indeed, we find big data analytics capable to support resilience to rare, situational surprises captured in analytic models. Nonetheless, this article examines the efficacy of resilience analytics by answering a single motivating question: Can big data analytics help cyber–physical–social (CPS) systems adapt to surprise? This article explains the limitations of resilience analytics when critical infrastructure systems are challenged by fundamental surprises never conceived during model development. In these cases, adoption of resilience analytics may prove either useless for decision support or harmful by increasing dangers during unprecedented events. We demonstrate that these dangers are not limited to a single CPS context by highlighting the limits of analytic models during hurricanes, dam failures, blackouts, and stock market crashes. We conclude that resilience analytics alone are not able to adapt to the very events that motivate their use and may, ironically, make CPS systems more vulnerable. We present avenues for future research to address this deficiency, with emphasis on improvisation to adapt CPS systems to fundamental surprise.

KEY WORDS: Analytics; infrastructure; resilience; surprise

1. THE ALLURE OF RESILIENCE ANALYTICS

Recent advances in information technologies have reduced the marginal cost of capturing, processing, and using data by several orders of magnitude. Internet-connected devices, ranging from thermostats to toilets to traffic signals, can now function as both sensors that report on the environment and as actuators that attempt to control it. Similarly, social media applications now collect massive data from crowd-sourced observations, including earthquake and infectious disease impacts, in an effort to improve detection and response actions. This ability to connect data and decision—at large scale and in real time—creates new capabilities for society to respond and adapt to an increasingly volatile world.

This emergence of so-called big data requires associated "big data analytics" to reveal patterns, trends, and associations between physical systems and human behavior, but the technologies here are still nascent. For example, Google Flu Trends was launched in 2008 to improve the detection and prediction of influenza outbreaks in the United States by analyzing search terms. At first, it improved existing methods for rapid, reliable predictions of data later reported by the Centers for Disease Control and Prevention (Ginsberg et al., 2009), but then surprisingly stopped working (Lazer, Kennedy, King, & Vespignani, 2014) and was ultimately abandoned as a

[1]Operations Research Department, Naval Postgraduate School, Monterey, CA, USA.
[2]School of Sustainable Engineering and the Built Environment, Arizona State University, AZ, USA.
[*]Address correspondence to David L. Alderson, Operations Research Department, Naval Postgraduate School, 1411 Cunningham Road, Monterey, CA 93943, USA; dlalders@nps.edu.

 0272-4332/19/0100-1870$22.00/1 © 2019 Society for Risk Analysis

stand-alone service in 2015 (Google AI Blog, 2015). The service, which once was heralded as a "poster child" for the power of data analytics, is now regarded as a cautionary tale of "big data hubris" (Lazer et al., 2014) for modelers and users alike.

Nevertheless, data analytics for improved decision making has considerable implications for the field of risk analysis. (See, e.g., the August 2017 Special Issue of *Risk Analysis*, i.e., Choi & Lambert, 2017.) The implications for resilience are less well investigated. For the purposes of this work, we focus on analytics developed for improving the resilience of critical infrastructure systems such as those supporting hurricane prediction, flood protection, financial markets, and the electric grid. These systems are critical because they provide services that ensure the safety and function of society; are complex due to their multiple functions, interdependencies, and human–technological interactions; and they produce significant amounts of real-time data. Critical infrastructure resilience is broadly understood as the ability for these systems to adapt to adverse events that impact their functioning, where many definitions of resilience exist depending on infrastructure context (Cutter et al., 2013; Pritzker & May, 2016; The White House, 2013).

Barker et al. (2017) coined the term *resilience analytics* to mean "the systematic use of advanced data-driven methods to understand, visualize, design, and manage interdependent infrastructures to enhance their resilience and the resilience of the communities and services that rely upon them." The term emphasizes the use of analytics in support of critical infrastructure operations and management decisions. They categorize critical infrastructure as cyber–physical–social (CPS) systems composed of: (1) infrastructure networks "that enable essential 'lifeline' services for society (e.g., transportation, electric, power, communications)"; (2) service networks composed of "human systems that engage with these infrastructure systems during a disruption (e.g., emergency responders, humanitarian relief, debris removal)"; and (3) community networks of "the interconnected society that the other networks support (e.g., relationships among people and communities)." Given that the complexities of multilayered, interdependent CPS systems are beyond the comprehension of any single individual or organization, the promise of resilience analytics is that more data and advanced computational techniques will result in faster and better decision making, with fewer cascading losses, deaths, and economic impacts.

If resilience analytics can deliver on these promises and reduce the calamitous impacts of infrastructure failure and natural disaster, it is no surprise that stakeholder groups including government agencies, first responders, infrastructure providers, and others might rush to adopt the most advanced computational techniques. Government (Comptroller of the United States, 2017; U.S. Government, 2019) and the military (Eckstein, 2017) are now working quickly to try to make sense of data analytics and turn it into tactical, operational, and strategic advantage. However, without critical examination of the assumptions of existing analytics and the limitations they create, it is unclear if resilience analytics will truly result in faster and better decision making for CPS systems. Resilience analytics may go the way of Google Flu Trends and simply stop working, or worse, their use may, ironically, expose the public to grave dangers resulting from overconfidence, myopia, loss of adaptive or innovative capacity, and/or misconceptions that result in more brittle CPS systems.

In this article, we interrogate resilience analytics by answering a single question: Can big data analytics help CPS systems adapt to surprise? We find that resilience analytics may be capable to adapt to situational surprises, but will remain unable to adapt to fundamental surprises that cannot be predicted *a priori*. Unfortunately, critical infrastructure remains vulnerable to events like unprecedented hurricanes, failing dams, cascading blackouts, and stock market flash crashes. We conclude that the promise of resilience analytics is limited not by computational capability but by a lack of knowledge of how systems improvise when models are deemed useless or harmful. Future research should focus on understanding the context in which resilience analytics are used and delineating how improvisation occurs with respect to analytics. Without this knowledge, efforts to improve CPS resilience with big data analytics can lead to greater vulnerability.

2. THE ROLE OF MODELS AND MODELERS

Before we can examine resilience analytics, we must understand how they are developed and used in CPS systems. In general, analytics harness the availability of big data to inform decisions through the application of statistical and mathematical models. Because there is no single definition of "model," for simplicity we adopt the lexicon of Brown (2018): "A model is an abstraction that emphasizes certain aspects of reality to assess or understand the behavior of a system under study." This definition is broad enough to cover a range of models, including physical models (e.g., miniature vehicles or buildings), logical models (e.g., of software dependencies), mathematical models (e.g., functions), information models (e.g., infographics), and/or any combination thereof potentially scaling to represent all the working CPS infrastructure across a region. Resilience analytics can be classified in three subcategories based on their use (Barker et al., 2017; Rose, 2016):

- *Descriptive analytics* describe and help visualize the performance of a system.
- *Predictive analytics* determine complex patterns, relationships among variables, and quantify the likelihood of future events.
- *Prescriptive analytics* identify and evaluate a feasible course of action given a set of constraints, possible interventions, and objectives.

All analytics rely on embedded models that dictate their inputs, outputs, and use. Predictive and prescriptive analytics often rely on explicit models such as mathematical equations and algorithms to guide user decisions and actions. In contrast, descriptive analytics often rely on implicit models, such as the choice of what data to collect (logical models) or how to manipulate and present data (information models).

2.1. A Conceptual Look at Big Data Analytics in CPS Systems

Fig. 1 depicts a common story for how big data analytics are deployed in CPS systems to support resilience. A *Model* that captures key characteristics of the real world (e.g., here we represent the Earth as a simple circle) is embedded in *analytics* that take in high velocity, volume, and variety information assets from the real world and transform them into useful outputs for a *User* (i.e., an operator, manager, regulator, engineer, or other embedded in a CPS service and community network). The User can further query the analytics to achieve the *goals* of describing, predicting, or prescribing a particular phenomenon. The User then takes actions that change the real world, feeding back new information assets into analytic models. Together, the interaction between big data, analytics, models, users, and the real world creates a simple sociotechnical feedback loop relevant to any CPS context.

The capability for analytics to improve User decision making and actions via the feedback loop in Fig. 1 has been a major success story of the last decade and helps drive the increased use of machine learning techniques. This is emphasized in the development and deployment of machine learning models (see Huddleston & Brown, 2018, for more detailed explanation of this process). For example, supervised learning models start with an initial data set to "train" and "tune" machine learning algorithms and identify which one serves User goals. The selected model is further trained, tested, and validated against different data before being deployed for the User. A key issue in model development is the extent to which the Model can "generalize" to handle data that were not considered during training, testing, and validation. Problems can arise when the training data are not representative of the real world or when patterns in real data simply do not exist. Iterative feedback with the User as seen in Fig. 1 is helpful because the deployed analytic model can improve over time (i.e., "learn") as it interacts with more data.

However, this story of analytics is incomplete because it does not include the *Modeler* who interprets the stated goals of the User to develop the analytic model. This translation is not simple. All Modelers have their own backgrounds, beliefs, goals, needs, and other personal and contextual values comprising a *frame of reference* (also called a *preanalytic vision*; see Costanza, 2001; Schumpeter, 1954). This frame of reference may be quite different from the User's personal and contextual values that influence the User's actions (also called a *decision frame*; see Simon et al., 1987). In the case where the Modeler and the User are distinct individuals, their alignment of goals depends on communication throughout model development and use, and discrepancies among these two sets of values may limit resulting analytics. For example, it is common for a User's decision frame to change over time and render previous modeling efforts obsolete. Similarly, a Modeler's preanalytic vision may include technical decisions that reduce the accuracy, throughput, and usability of analytics with or without User knowledge. (Anecdotal "war stories" of analysts are rife with tales where misunderstanding between Modeler and User led to project failure.) Where pragmatic and technical needs conflict, User and Modeler values may differ even when they are the same individual.

Thus, it is necessary to separate the Modeler from the User and add additional dependencies into Fig. 1 that represent the influence a Modeler has

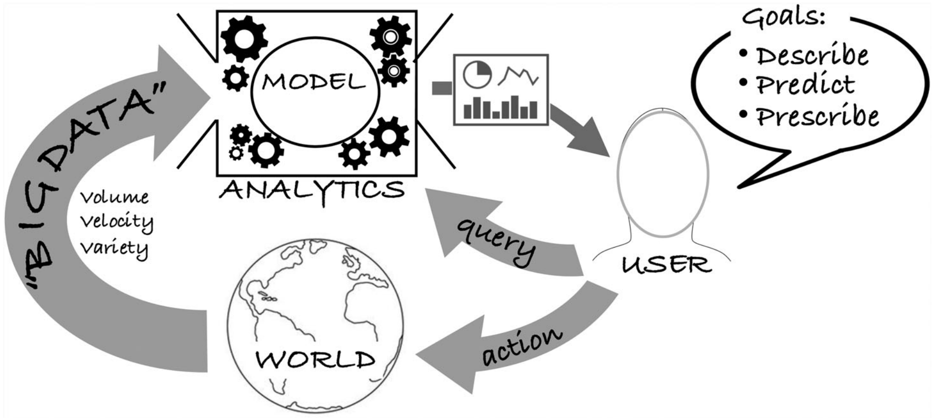

Fig. 1. A simplified representation of analytics to support user decisions. Big data from the real world are fed as input into a model where they are transformed into support for the stated goals of a User to describe, predict, or prescribe behavior in a complex cyber–physical–social (CPS) system. A User can query the model, then decide and act upon model outputs. Actions taken by a User affect the real world and subsequently feed new data inputs back into the analytic model.

on the feedback loop. Fig. 2 depicts this augmented view with Modeler and User frames of reference, the alignment of goals between them, and the actions taken by a Modeler to develop and deploy analytics. Fig. 2 is illustrated to emphasize that the Modeler's values are embedded in analytics, conceptually described by experts as "the modeler is in the model"(Woods, personal communication, 2017a).

2.2. Three Confounding Dependencies for Resilience Analytics

For resilience analytics, model development is guided by common values found across infrastructure sectors, federal agencies, and academic literature. Examples of User goals are captured in resilience definitions across the U.S. federal government, such as improving the ability for infrastructure systems "to plan and prepare for, absorb, recover from, and adapt to adverse events"(Cutter et al., 2013); the ability of service providers and organizations "to recognize threats and hazards and make adjustments that will improve future protection efforts and risk reduction measures"(The White House, 2013); and the ability of communities "to prepare for anticipated hazards, adapt to changing conditions, and withstand and recover rapidly from disruptions"(Pritzker & May, 2016). The purpose of a Modeler, then, is to produce analytics that support User actions like improved crisis planning and recovery for CPS systems.

Bringing the Modeler in the loop, however, introduces three dependencies that challenge whether resilience analytics can support these goals. The first dependency is between the User and Modeler representing the necessity for translation and interpretation of User goals and values. This dependency emphasizes that poor communication and alignment of goals may produce models that do not support User actions. George Box's famous aphorism "[a]ll models are wrong, but some are useful" reminds us that analytics are only meaningful in their decision context, which may be unavailable to a Modeler. Following Box's aphorism, *useful* resilience analytics are those that help CPS networks withstand, adapt to, and recover from disruptions. A key goal for the Modeler is to avoid deploying *useless* or *harmful* resilience analytics, that is, those that do not support these actions or make them more difficult.

The second dependency is between the Modeler and the Model, representing embedded values in data analytics. Since the Modeler's values (not the User's) are embedded in the Model, only a Modeler can fully explain how their analytics work. A growing issue with analytics is *model transparency*, that is, the extent to which the output of a model can be readily explained. Users often view analytics as a "black box" and cannot meaningfully explain how they transform inputs and/or assumptions into descriptive, predictive, and prescriptive support. In such cases, the Modeler must explain *what* the model has produced as output and *why* the result makes sense. Unfortunately, there are situations where analytic models are a black box to the Modeler as well. Automated machine learning techniques, in particular, can result in models that lack transparency to

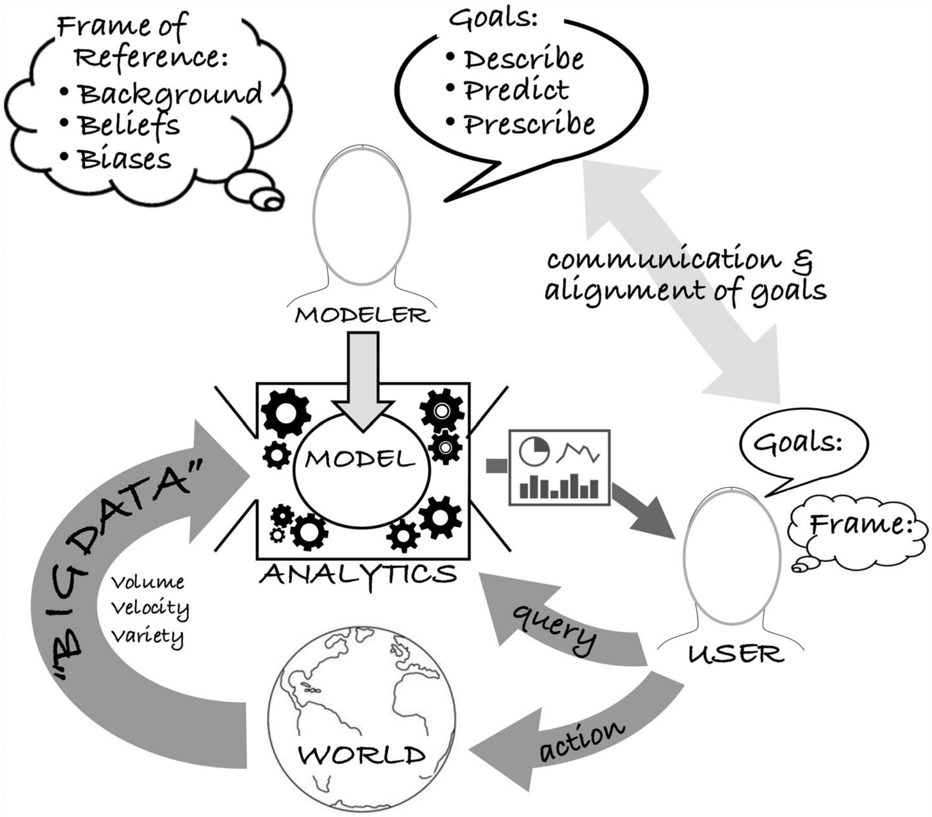

Fig. 2. The relationship between analytic models, users, modelers, and the real world. The inclusion of a "Modeler" who creates and updates big data analytics introduces additional dependencies that confound the simple model. Both the interpretation of a User's decision frame by the Modeler and the influence of real-world stimuli on Modeler may produce an incomplete preanalytic vision and lead to inappropriate analytic models.

their Modeler (Knight, 2017). For these black boxes, it may be impossible to identify how Modeler values are embedded in analytics. This makes it even more difficult to ensure resilience analytics are "useful" to a User's needs and/or reconcile differences between Modeler and User values.

The third dependency is between the real world and the Modeler representing the way in which a Modeler's values and preanalytic vision are influenced by external stimuli. Analytic models must be updated and modified regularly. The process of *sensemaking*, in particular, emphasizes that User and Modeler values will change their understanding of infrastructure over time (Weick, 1989, 2009) and will require new analytics. While some of this change is captured in the machine learning process, more dramatic shifts in values that dictate a need for a new model cannot be. For resilience, the models necessary to describe, predict, and prescribe many extreme events do not yet exist, meaning resilience analytics produced today will likely need to change in the future. Costanza (2001) elaborates this fact by emphasizing that the credibility of analytics depends on constant reconciliation between a Modeler's values and a User's values, both of which change over time: "[C]redibility proceeds from honest discussion of this underlying [preanalytic] vision and its inherently subjective elements, as well as from constant, pragmatic testing of conclusions against real-world problems, rather than by appealing to a nonexistent objectivity.... [T]he ultimate goal is therefore not truth, but quality and utility."

Together, these three dependencies create underappreciated limits on the use of resilience analytics for CPS systems.

3. SURPRISE HAPPENS

Foundational to the notion of resilience is that a system can adapt to situations or events that have never been experienced before. Specifically, for analytics to improve the resilience of CPS systems, the feedback loop in Fig. 2 needs to be responsive to extreme and/or unforeseeable natural and man-made disruptions that challenge lifeline services, infrastructure management practices, and community safety and security. We refer generally to these novel or rare situations as *surprises*, and an important question for the development of resilience analytics is: Can big data analytics help CPS systems adapt to surprise?

To answer this question, one must distinguish surprises apart from normal events (see Table I). We define *normal events* as situations where CPS systems operate according to previous beliefs, there are no contingencies to avert, and nothing new is experienced. To define surprise, we rely on the work of cognitive scientists, who typically distinguish between two types: *situational surprise* and *fundamental surprise*. Seminal work by Lanir (1986) and Wears and Webb (2014) describes four distinguishing features: "[First] fundamental surprise refutes basic beliefs about 'how things work,' while situational surprise is compatible with previous beliefs. Second, in fundamental surprise one cannot define in advance the issues for which one must be alert. Third, situational and fundamental surprise differ in the value brought by information about the future. Situational surprise can be averted or mitigated by such foresight, while advance information on fundamental surprise actually causes the surprise. … And finally, learning from situational surprise seems easy, but learning from fundamental surprise is difficult." Table I summarizes key distinctions between normal events, situational surprise, and fundamental surprise.

Situational and fundamental surprise appear within two distinct feedback loops for big data analytics, as depicted in Fig. 3. Feedback within the blue "inner loop" (denoted here as Type 1) involves data that are fed into models and can include situational surprise. Analytics like those harnessing supervised machine learning models can automatically adjust behavior to discrepancies between model outputs and observations from situational surprises.

In contrast, fundamental surprise—that challenges the decision frame of the User and/or the preanalytic vision of the Modeler—proceeds in a red "outer loop" that bypasses automated model feedback (denoted here as Type-2 feedback). Accordingly, Type-1 feedback is unresponsive to fundamental surprise because it cannot change the User or Modeler values embedded in a model. Again, this is exemplified by machine learning models that fail to "learn" when deployed in situations they were inappropriately trained for or where no underlying patterns in real data exist. There is no current way for big data analytics to adapt their internal structure to these situations on their own. A Modeler must adapt analytics when a fundamental surprise renders the current Model obsolete. In situations that change the User's decision frame, without Type-2 feedback for communication and alignment of goals, there is no way to fundamentally change the central model. Because adaptation to fundamental surprise can only occur via Type-2 feedback, even the most advanced machine learning or artificial intelligence models cannot respond to fundamental surprise, and by extension cannot adapt to it.

The key point is that the Modeler, not the Model, creates both limitations on how big data analytics support User goals and provides the solution for how to adapt analytic models to fundamental surprise. Including the Modeler "in the loop" recognizes that analytics are susceptible to having the wrong preanalytic vision during initial model development. Analytics with useless models that are predicated on misconceptions or draw our attention to the wrong things cannot be improved by better data or machine learning techniques. In fact, more and better information assets in a useless model may become harmful by reinforcing misconceptions embedded in results, leading to greater danger resulting from a false sense of security and/or poor decision making. A view of resilience analytics that only considers the inner loop in Fig. 3 is susceptible to the danger of fundamental surprise in this manner.

Similarly, CPS systems faced with fundamental surprise require a Modeler-in-the-loop to adapt big data analytics and ensure resilience to unforeseen events. When modelers identify new ways to structure analytics that were otherwise unknown during initial model development or users determine current model outputs obsolete, resilience requires the model to fundamentally change.

Thus, analytics can help CPS systems adapt to surprise *if and only if* it is a situational surprise. Big data analytics *on their own* cannot help CPS systems adapt to fundamental surprise. In situations where fundamental surprise occurs, no Model in the absence of Type-2 feedback, no matter how well informed with observation during its initial development, can achieve the goals of resilience analytics. At minimum, resilience requires Type-2 feedback to adapt models deemed useless or harmful, even if

Table I. Contrasting Features of Normal Events, Situational Surprise, and Fundamental Surprise

Normal Events	Situational Surprise	Fundamental Surprise
Events occur according to previous beliefs	Events are compatible with previous beliefs	Events refute basic beliefs about "how things work"
System and components operate as planned	Failures and system responses are well-modeled or measurable	One cannot define in advance the issues for which one must model or measure
No contingencies to avert or manage	Surprise can be averted or mitigated with information about the future	Advance information about events causes even greater surprise
Nothing new to learn	Learning seems easy	Learning is challenging
Example	*Example*	*Example*
Purchase lottery ticket, then *lose* the lottery	Purchase lottery ticket, then *win* the lottery	Do *not* purchase lottery ticket, then *win* the lottery

this form of feedback is also the reason why George Box's aphorism remains true.

4. SURPRISE IN CPS SYSTEMS

The nature of surprise suggests that the answer to our guiding question—Can resilience analytics help CPS systems adapt to surprise?—depends on the User and Modeler, not just the analytics. In general, disruptions to CPS systems can take the form of situational surprise (e.g., the failure of a known component), and here big data analytics can improve resilience (e.g., by anticipating when the component should be replaced preemptively). Searching for anomalies that cause situational surprise is a normal activity (e.g., testing and debugging) and often reveals underappreciated dependencies among physical and digital assets (Woods, 2017b). User and Modeler can automate this process with Type-1 feedback to help ameliorate situational surprises before, during, and after they occur.

However, disruptions can also take the form of fundamental surprise (e.g., failure of a previously unknown component) that changes the understanding of how the CPS system actually works. The complex nature of many CPS systems makes it difficult to know all relevant dependencies that may cause, exacerbate, or even help mitigate a disruption. These unknown dependencies are often revealed when changes in environmental context and in societal needs alter the background, beliefs, and goals of a User and/or Modeler in such a way that previous conceptions of the CPS system become irrelevant. For example, engineers are realizing that growing extremes in weather events mean that historical data are no longer representative of future events (Milly et al., 2008), radically changing the design loads for civil projects that manage lifeline services.

The ways User and Modeler adapt to these fundamentally new situations ultimately dictates the resilience of CPS systems. Specifically, the ability of a CPS system to harness Type-2 feedback and *improvise* new models and decisions when faced with fundamental surprise is critically important to the future development of resilience analytics. We consider four representative disruptions in CPS systems motivating this need.

4.1. Collective Improvisation by User and Modeler

The combined system in Fig. 3 shows that both User and Modeler can participate in Type-2 feedback when faced with fundamental surprise. The ideal situation is one where User and Modeler coordinate their efforts with sufficient time to generate new analytics that describe, predict, and prescribe previously unforeseen events.

Example: Hurricane Ophelia

The hurricane season of 2017 was marked with numerous Category-5 storms that caused catastrophic damages across the Atlantic and Caribbean. The United States alone suffered unprecedented damage to infrastructure systems in Texas, Florida, Puerto Rico, and the U.S. Virgin Islands, breaking records for inundation in Houston and infrastructure recovery time in Puerto Rico.

Notwithstanding these catastrophes, perhaps the most surprising hurricane of the season from an analytics perspective was Hurricane Ophelia. Unlike other Atlantic hurricanes that made landfall in North America, Ophelia progressed toward Europe to

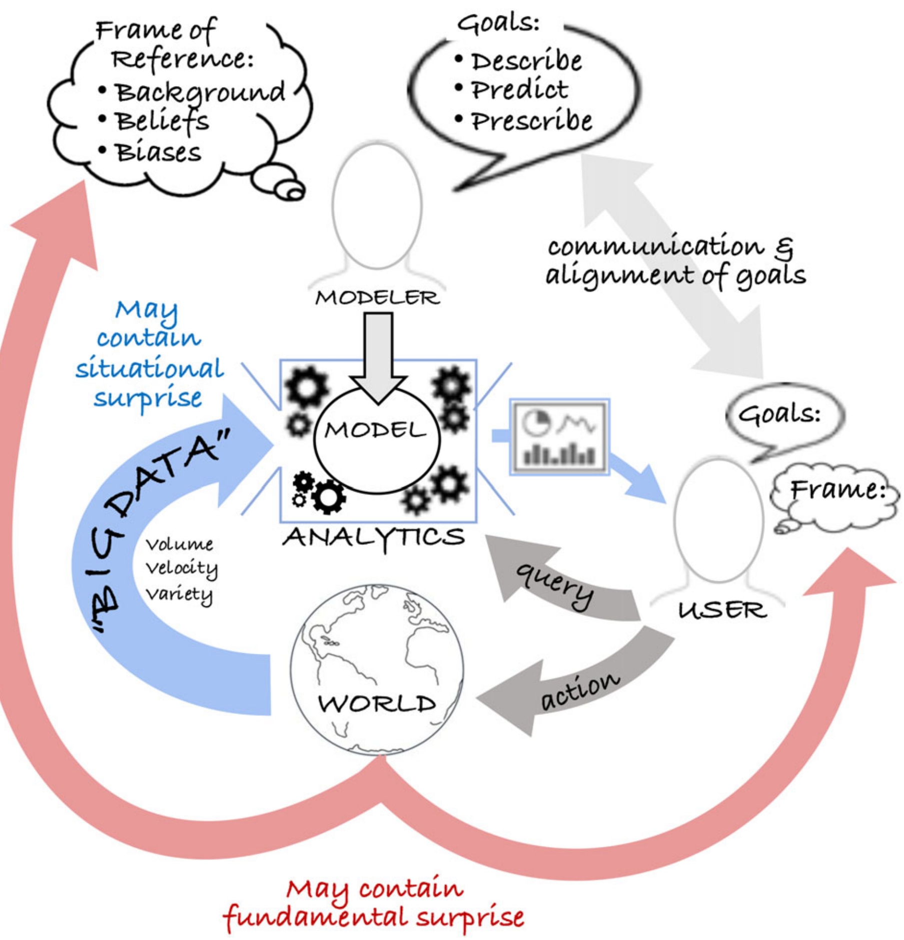

Fig. 3. Surprise from the real world creates two distinct feedback loops. Type-1 feedback (via the "inner loop" denoted in blue) that feeds big data into analytic models can contain situational surprise, in the form of data that are unexpected but still "fit" within the model structure. In the "outer loop" (Type-2 feedback, denoted in red), novel and rare experiences may characteristically change a user's decision frame or a modeler's preanalytic vision (i.e., fundamental surprise), which may upend previous model assumptions and require the development of entirely new analytics.

become "the farthest east Major Hurricane (Category-3 or higher) on record in the Atlantic Basin" (Met Éireann, Government of Ireland, 2018). Ophelia was fundamentally surprising because it followed an "extremely weird" path toward Europe and made landfall in Ireland and the United Kingdom (Mersereau, 2017). The significant deviation in path led Ophelia to reveal possibly underappreciated limitations in predictive models from the U.S. National Hurricane Center (NHC); see Fig. 4. Prior to Hurricane Ophelia, NHC models designed to predict tropical-storm-force winds were not developed to visualize past (0°, 60°N), creating a sharp corner in hurricane wind-speed predictions and visually showing fundamental surprise. The relatively slow speed of storm advance allowed User and Modeler to recognize this surprise, and an updated projection 12 hours later for the same storm corrected the issue.

4.2. Improvisation Without a Modeler

A common breakdown of Type-2 feedback is when there is no Modeler "in the loop" to update analytics (see Fig. 5). When fundamental surprise reveals to a User that the current model is flawed and/or changes a User's frame of reference, resilient crisis response may require the User to make decisions without analytic support (i.e., abandon the Model). This situation renders previously "useful" analytics "useless," and forces a User to improvise, that is, rely on own expertise and heuristics for navigating the surprise.

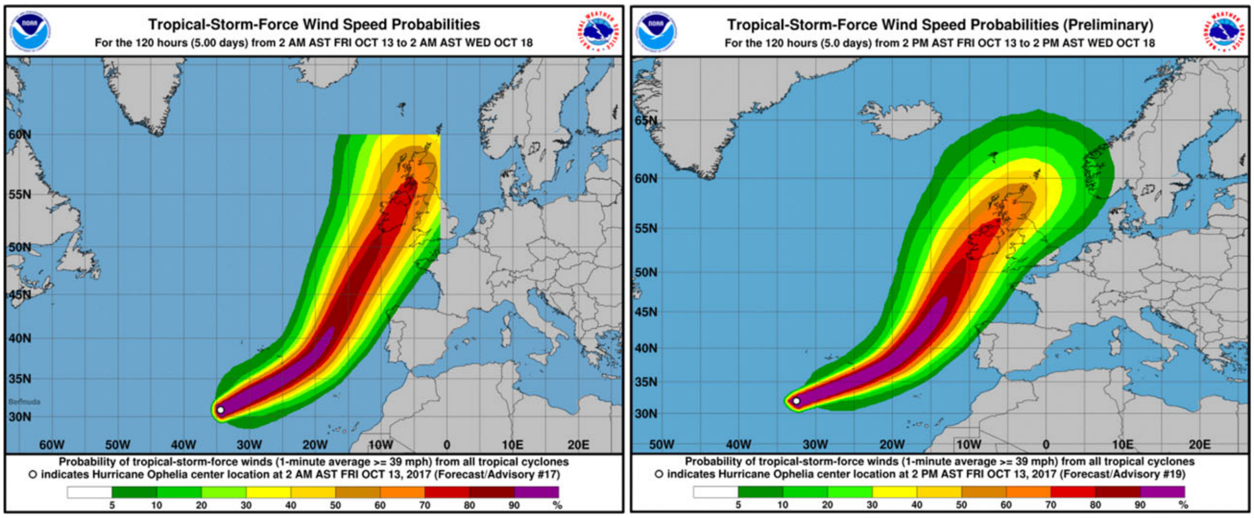

Fig. 4. U.S. National Hurricane Center (NHC) five-day probability of tropical-storm-force winds for hurricane Ophelia in October 2017. Left: Projection as of October 13, 2017 02:00 AST (U.S. National Oceanic and Atmospheric Administration [NOAA], 2019); fundamental surprise appears as missing projections beyond (0°, 60°N). Right: Projection as of October 13, 2017 14:00 AST (Met Éireann, Government of Ireland, 2018); adaptation on the part of the User and the Modeler results in updated analytics. Copyright: Public Domain, see http://www.weather.gov/disclaimer.

Example: The Near-Breaching of the Oroville Dam

The heavy rains of winter 2017 on the West Coast of the United States ended one of the most serious droughts on record in California. However, a component failure at the Oroville Dam in northern California resulted in the largest evacuation order in the history of the state, and nearly resulted in catastrophic collapse of the tallest dam in the United States (California Department of Water Resources, 2017, 2018; Hollins, Eisenberg, & Seager, 2018). Prior to the crisis, the frame of reference for the California Department of Water Resources and the Federal Energy Regulatory Commission (FERC) was to make decisions based on reservoir-wide risk analyses that identified potential dam failure modes. The most recent analysis in 2014 ignored structural problems from initial dam design, insufficient maintenance practices, and poor mountainside geological conditions. Nonetheless, given the prevailing drought conditions and the fact that the dam had successfully withstood prior flood events, none of these concerns were treated as urgent.

During February 2017, a series of rainstorms filled the reservoir, necessitating opening of the primary spillway of the Oroville Dam to reduce water levels in accordance with the operating procedures that govern dam management. However, cracks opened in the concrete surface of the primary spillway almost immediately, and dam operators were faced with the complex decision of whether to continue using the broken primary spillway, or use the ungated, emergency spillway that directed water over an unprotected, earthen section of the dam. Despite nearly five decades of dam operation, the emergency spillway had never been activated, so operators had no experiential basis from which to predict how it would perform (California Department of Water Resources, 2017, 2018; Hollins et al., 2018). Experts were divided between these two choices. Some operations personnel, management executives, and regulators favored maintaining the frame of reference that the emergency spillway would be safe to operate. However, safety engineers and emergency managers preferred continuing to use the broken primary spillway, given the lack of experience with the emergency spillway. Decisions to continue operating the broken primary spillway at reduced flow rates resulted in activation of the emergency spillway and uncontrolled flow down the hillside. Within minutes, geologists identified a fundamental surprise of downstream erosion that threatened immediate hillside collapse and reservoir containment failure. Onsite incident commanders had time only to order and implement an evacuation of local emergency personnel, leaving operators controlling the sluice gates on the primary spillway to improvise on their own. These operators quickly abandoned prior belief in a

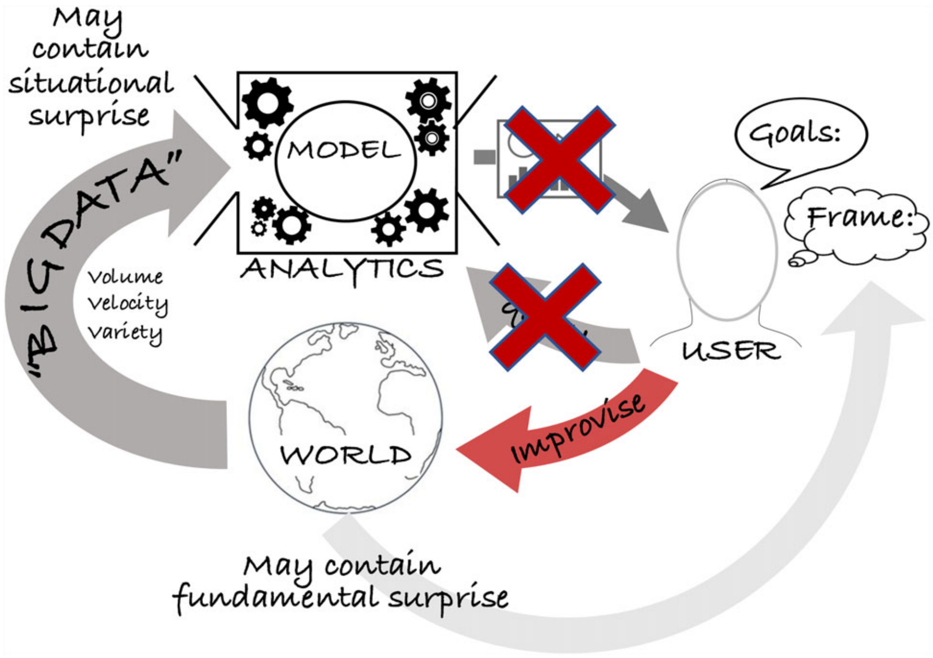

Fig. 5. Fundamental surprise experienced by a User with no Modeler-in-the-loop. When there is no "Modeler-in-the-loop" and the User deems current analytics as useless or harmful, there is no effective way to communicate new analytic needs and adapt models to fundamentally new situations. The User has little choice but to abandon the analytics (represented by the X's) and improvise actions without additional decision support.

functional emergency spillway with the improvisation action of opening the primary spillway gate to unprecedented flow rates. This flow caused even more damage to the primary spillway, but the resulting reduction in water levels redirected water from the emergency spillway prior to catastrophic collapse of the reservoir containment.

4.3. Improvisation Without a User

Another increasingly common situation that inhibits Type-2 feedback is when the role of the User is automated, and the Modeler must respond to a fundamental surprise without a User-in-the-loop (see Fig. 6). Many CPS systems and industrial controls operate on timescales faster than humans can respond. In these cases, automation of resilience analytics can create new adaptive capacity to situational surprises, but can be maladaptive in response to fundamental surprise (i.e., they preclude the possibility of a User-in-the-loop). In situations when CPS systems experience a fundamental surprise with no User, it is now up to the Modeler to improvise and try new analytics.

Example: The Arizona-Southern California Power Outage

On September 8, 2011, an 11-minute system disturbance left 2.7 million customers without power for up to 12 hours across the southwestern United States and Northwest Mexico (including the entire city of San Diego, CA; see Clark, Chester, Seager, & Eisenberg, 2018). The event was initiated by the loss of a single, extra high-voltage (500 kV) powerline near Yuma, Arizona that led to instantaneous power redistribution, voltage deviations, and overloads across the region (Veloza & Santamaria, 2016). This caused a cascade of losses as infrastructure tripped offline and initiated automatic load shedding. The Federal Energy Regulatory Commission (FERC) and North American Electric Reliability Corporation (NERC) found 27 separate issues that contributed to the cascade (Federal Energy Regulatory Commission, North American Electric Reliability Council, 2012), nearly all of which stemmed from analytics and automated controls that precluded timely operator action. Specifically, the system relied on models that made important power system information difficult for operators to access and protection schemes that tripped infrastructure

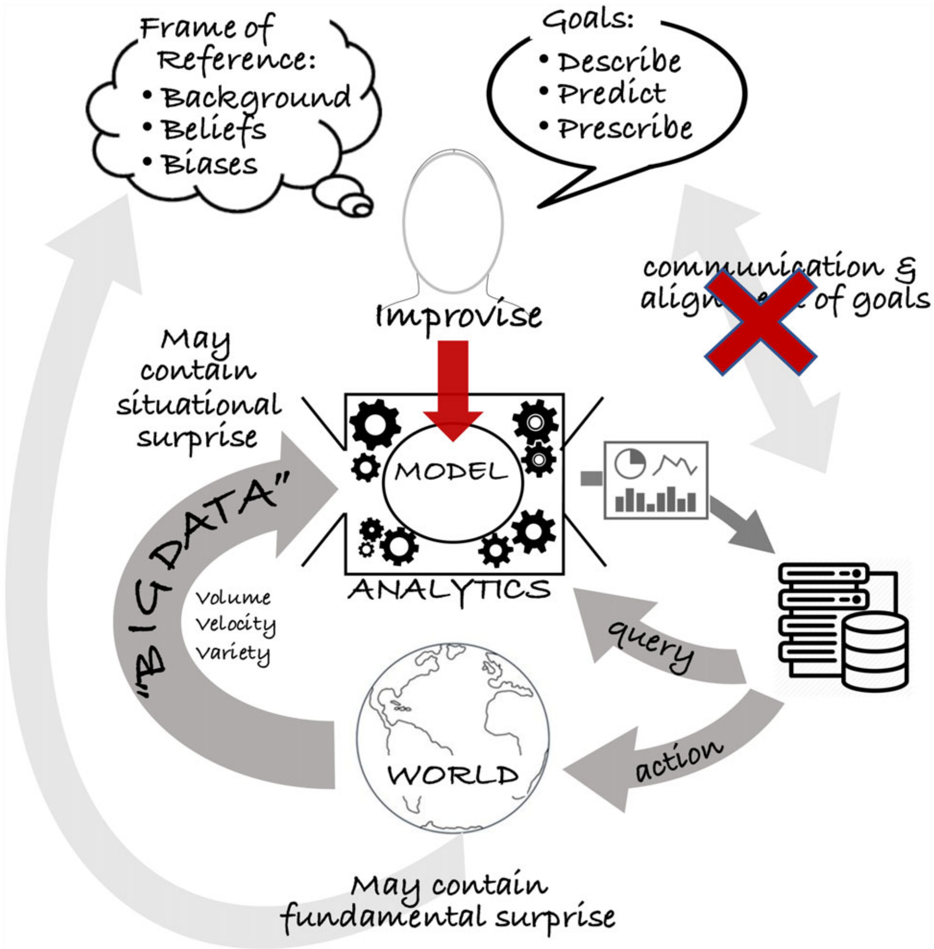

Fig. 6. Fundamental surprise experienced by modelers with no User-in-the-loop. In cases like large-scale blackouts, the User-in-the-loop is a technological system that acts at the scale and speed of infrastructure. Modelers that experience fundamental surprise in these situations have no effective way to align analytic needs to fundamentally new situations. The translation of a new frame of reference into new analytic models may be too slow to interdict maladaptive actions hard coded into existing models and technologies. Type-2 feedback is lost where red "Xs" are shown.

out of service before operators could respond. Moreover, planning models for next-day, seasonal, and long-term operations were fundamentally flawed in ways that were unknown, underappreciated, or ignored prior to the event. The combination of these issues precluded secure system operation despite the fact that the southwestern power system was designed and could have been operated to withstand the loss of the powerline. FERC and NERC concluded that the event could have been managed within the 30-minute window normally afforded to contingency operations; this resulted in six regulatory settlements between FERC and regional electric power authorities to update models and analytic capabilities above and beyond reliability standards (Federal Energy Regulatory Commission, 2015).

4.4. Improvisation Without a User or Modeler

The most challenging situation for a CPS system to respond to fundamental surprise is when the User has been replaced by automation and the Modeler is unavailable over the timescale at which the analytics need to be modified.

Example: Flash Crashes in the Stock Market

The increased use of computer-assisted trading (which itself is a form of analytics) in financial markets provides several examples where the rapidity of information flow through models and the absence of a User-in-the-loop limits the ability of the system to respond in the face of surprise. The first of these is the infamous Stock Market Crash of 1987, during which the Standard & Poor's Index fell approximately 20% in a single day. Carlson (2007) details the causes of the crash, among which was "an increase in the use of 'program trading' strategies, where computers were set up to quickly trade particular amounts of a large number of stocks, such as those in a particular stock index, when certain conditions were met." Specifically, there was a rise in the use of program trading to implement "portfolio insurance" strategies that automatically trade financial instruments such as index futures to hedge against falling markets. In the presence of an initial decline, these automated systems can trigger a cascading effect in which price declines trigger selling, which induces additional price declines, and so on. The increased use of automated quotation systems at the time further enabled this cascade to travel internationally (Roll, 1988). If human traders had been "in the loop," they may have been able to slow or arrest these cascading effects by switching trading strategies as they occurred. Ironically, Shiller (1988) argues that many investors were actually anticipating a crash, which increased the popularity of automated portfolio insurance and ultimately facilitated the crash.

Over the subsequent 30 years, human traders have been increasingly replaced by algorithms that respond at the speed of light to changes in the market. In the current environment, trading firms compete to locate their computer trading operations as close as possible to electronic exchanges so as to reap the benefits afforded by fractions of a second (Lewis & Baker, 2014). In the presence of Type-2 feedback requiring fundamental changes to analytics (due to the realization of flawed assumptions, unanticipated interactions in the market, or the discovery of flaws in the model itself), the Modeler often does not have time to revise or repair the Model. This use of high-frequency trading is largely credited with causing the Flash Crash of 2010 (Kirilenko, Kyle, Samadi, & Tuzun, 2017) during which the Dow Jones Industrial Average index dropped more than 900 points (approximately 9%) in less than five minutes, and then regained much of that over the subsequent 15 minutes. Since that time, additional mini-flash crashes have occurred (Golub, Keane, & Poon, 2012), and sometimes the only reasonable response to unexpected behavior is to shut down an exchange (McCrank, 2015).

5. A PATH FORWARD FOR RESILIENCE ANALYTICS

Recognition that CPS systems are susceptible to fundamental surprise reveals a critical need to rethink resilience analytics. Despite how useful analytics may be during situational surprises, they are insufficient for resilience. CPS systems will experience fundamental surprises that reveal "useful" analytics as outdated, incomplete, inadequate, and even harmful. The previous examples show that the answer to our guiding question is "no"—big data analytics cannot help CPS systems adapt to fundamental surprise without Modelers and Users to respond to Type-2 feedback.

A path forward for resilience analytics is to complement research on data-driven techniques with research on understanding the context-dependent implications of improvisation. When CPS systems are faced with fundamental surprise, the best (and possibly only) course of action is to improvise. A challenge for resilience analytics is that the capacity for a CPS system to improvise is dictated by the frame of reference held by individuals and groups. Combining objective analytic models with these subjective understandings of resilience remains an important area of research for CPS systems. Unfortunately, we are not yet at the point where we can deploy analytics in a resilient way because we still lack theory, methods, and tools for improvisation to adapt systems to fundamental surprise.

5.1. A Framework for Research on Improvisation with Resilience Analytics

The role of User- and Modeler-in-the-loop establishes a framework to guide future research in improvisation (see Fig. 7). Because User and Modeler serve different roles in CPS systems, they bring different capacities for improvisation. Our discussion of Type-2 feedback and examples suggest that a defining factor dictating improvisation with analytics is whether the User or Modeler is present during surprises. Thus, there are four situations, represented as quadrants in Fig. 7 and described in detail below, with different capacity to improvise: (i) user and modeler present, (ii) user alone (modeler absent), (iii) modeler alone (user absent), and (iv) user and modeler absent. Each situation requires different

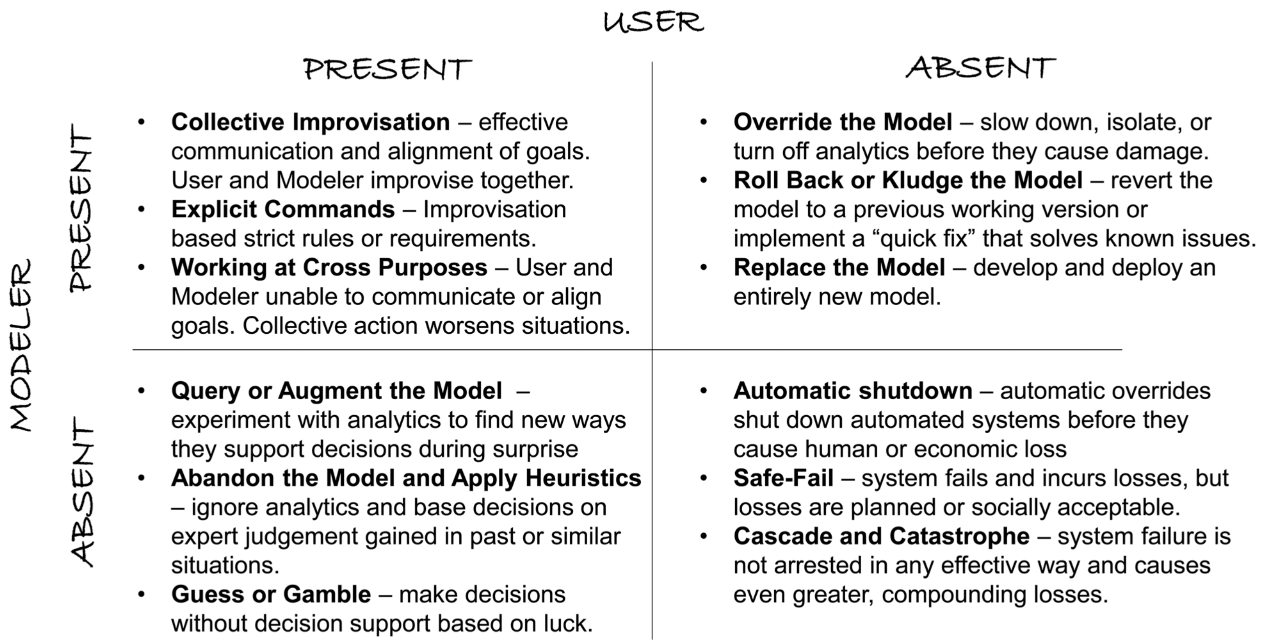

Fig. 7. Framework for research on improvisation with resilience analytics. A defining characteristic for improvisation is whether the User or Modeler is present to adapt the CPS system to fundamental surprise. This defines four situations with different capacities for improvisation: (i) user and modeler present, (ii) user alone (modeler absent), (iii) modeler alone (user absent), (iv) user and modeler absent. Future research should identify how resilience analytics support or inhibit improvisation in each situation.

kinds of improvisation to adapt CPS systems to fundamental surprise. Future development of data-driven resilience analytics should be augmented with research in improvisation for each of these situations.

5.1.1. User and Modeler Present

A CPS system with User and Modeler present has both the greatest potential for improvisation to fundamental surprise and is vulnerable to incomplete and harmful improvisation. The success of improvisation in this case is a consequence of how well Users and Modelers communicate and align analytic goals. In situations of strong alignment, tacit understandings of big data analytics enable collective improvisation, where Users and Modelers are able to improvise new models and actions without the need for explicit rules. Incomplete communication and ill-defined goals may require more explicit command and control relationships to ensure User and Modeler decisions do not interfere with each other. When communication and alignment of goals is poor, User and Modeler may actually work at cross-purposes, inhibiting crisis response and making CPS systems more vulnerable to fundamental surprise.

5.1.2. User Alone (Modeler Absent)

A CPS system with a User alone (no Modeler present) is only capable to improvise in ways that do not fundamentally change existing analytics. Normal activities like querying existing models for new results or augmenting models with new information may continue to work in the face of some surprises. A User faced with fundamental surprise will be forced to either abandon models and apply heuristics gained from past experience or make a complete guess when choosing a proper course of action.

5.1.3. Modeler Alone (User Absent)

A CPS system with the modeler alone (no user present) is only capable to improvise in ways that harness analytics for action. In some cases, the best course of action is to override the model to stop unwanted outcomes. This approach to improvisation can be improved by "rolling back" analytics to previous versions known to work or by creating a "kludge" (e.g., a patch) that solves known issues without fundamentally changing the model. In extreme situations, big data analytics can be completely replaced with new or different models that better serve CPS resilience needs.

5.1.4. User and Modeler Absent

Finally, a CPS system with neither a user nor modeler present has the least capacity for improvisation. Fully automated systems can only adapt to surprise in predetermined ways. The most common form of improvised action is when fail-safe and fail-silent systems are activated to override systems and prevent unwanted outcomes (Möller & Hansson, 2008). When these overrides are unable to contain losses, safe-fail designs may automatically absorb damages in predetermined and socially acceptable ways (Kim et al., 2017). Otherwise, data analytics on their own will lack any capacity to improvise, leading to cascading failures and catastrophic losses.

5.2. The Future of Resilience Analytics

Resilience is about adaptive capacity (Hale & Heijer, 2006; Hollnagel, Woods, & Leveson, 2006; Leveson et al., 2006; Madni & Jackson, 2009), and the resilience research community is a part of this process. Overemphasis on the benefits of data-driven methods can widen the gap between people making decisions during disruptive events (Users) and researchers developing resilience analytics to support them (Modelers). Ensuring the greatest capacity for improvisation requires a Modeler-in-the-loop who effectively communicates and aligns goals with the User.

Much of what is needed is simply the discipline to follow best practices for modeling and analysis: working to interpret a User's decision frame and translate that information into effective tools; maintaining communication between Modeler and User during both Model development and deployment; investigating events that challenge assumptions rather than discounting them. However, without more active participation by the risk analysis community, resilience analytics will remain limited, not by a lack of computational tools but by a lack of ways to adapt to fundamental surprise. A more resilient future where CPS systems are poised to adapt will require new understanding of surprise and improvisation in the context of resilience research itself.

REFERENCES

Barker, K., Lambert, J. H., Zobel, C. W., Tapia, A. H., Ramirez-Marquez, J. E., Albert, L., … Caragea, C. (2017). Defining resilience analytics for interdependent cyber–physical–social networks. *Sustainable and Resilient Infrastructure*, *2*(2), 59–67.

Brown, G. (2018). Chapter 6: Model building. In James J. Cochran (Ed.), *INFORMS analytics body of knowledge* (1st ed., pp. 231–274). Wiley Series in Operations Research and Management Science. Hoboken, NJ: Wiley.

California Department of Water Resources. (2017). *Lake Oroville spillway incident: Timeline of major events*. California Data Exchange Center. Available at: https://www.water.ca.gov/LegacyFiles/oroville-spillway/pdf/2017/Lake%20Oroville%20events%20timeline.pdf.

California Department of Water Resources. (2018). *Independent forensic team report. Oroville Dam spillway incident*. Available at: https://damsafety.org/sites/default/files/files/Independent%20Forensic%20Team%20Report%20Final%2001-05-18.pdf.

Carlson, M. A. (2007). *A brief history of the 1987 stock market crash with a discussion of the federal reserve response*. Working Paper 2007-13 in the Finance and Economics Discussion Series (FEDS), Divisions of Research & Statistics and Monetary Affairs, Federal Reserve Board, Washington, DC.

Choi, T.-M., & Lambert, J. H. (2017). Advances in risk analysis with big data. *Risk Analysis*, *37*(8), 1435–1442.

Clark, S. S., Chester, M. V., Seager, T. P., & Eisenberg, D. A. (2018). The vulnerability of interdependent urban infrastructure systems to climate change: Could Phoenix experience a Katrina of extreme heat? *Sustainable and Resilient Infrastructure*, *4*, 1–15.

Comptroller of the United States. (2017). *Data analytics to address fraud and improper payments* (Technical Report GAO-17-339SP). Washington, DC: U.S. Government Accountability Office (GAO).

Costanza, R. (2001). Visions, values, valuation, and the need for an ecological economics: All scientific analysis is based on a "pre-analytic vision," and the major source of uncertainty about current environmental policies results from differences in visions and world views. *BioScience*, *51*(6), 459–468.

Cutter, S. L., Ahearn, J. A., Amadei, B., Crawford, P., Eide, E. A., Galloway, G. E., … Zoback, M. L. (2013). Disaster resilience: A national imperative. *Environment: Science and Policy for Sustainable Development*, *55*(2), 25–29.

Eckstein, M. (2017, October 4). Navy Digital Warfare Office proving data analytics can help address nagging operational problems. *U.S. Naval Institute News*.

Federal Energy Regulatory Commission. (2015). *FERC approves final settlement in 2011 southwest blackout case*. Retrieved from https://www.ferc.gov/media/news-releases/2015/2015-2/05-26-15.asp.

Federal Energy Regulatory Commission, North American Electric Reliability Council. (2012). *Arizona-Southern California outages on 8 September 2011: Causes and recommendations*. FERC and NERC Staff. Available at: https://www.ferc.gov/legal/staff-reports/04-27-2012-ferc-nerc-report.pdf.

Ginsberg, J., Mohebbi, M. H., Patel, R. S., Brammer, L., Smolinski, M. S., & Brilliant, L. (2009). Detecting influenza epidemics using search engine query data. *Nature*, *457*(7232), 1012–1014.

Golub, A., Keane, J., & Poon, S.-H. (2012). *High frequency trading and mini flash crashes*. Retrieved from http://doi.org/10.2139/ssrn.2182097.

Google AI Blog. (2015). *The next chapter for flu trends*. Retrieved from https://ai.googleblog.com/2015/08/the-next-chapter-for-flu-trends.html.

Hale, A., & Heijer, T. (2006). Defining resilience. In E. Hollnagel, D. Woods, & N. Leveson (Eds.), *Resilience engineering: Concepts and precepts* (pp. 95–123). Aldershot, UK: Ashgate Press.

Hollins, L., Eisenberg, D., & Seager, T. (2018). Risk and resilience at the Oroville Dam. *Infrastructures*, *3*(4), 49. https://doi.org/10.3390/infrastructures3040049

Hollnagel, E., Woods, D., & Leveson, N. (Eds.) (2006). *Resilience engineering: Concepts and precepts*. Aldershot, UK: Ashgate Press.

Huddleston, S., & Brown, G. (2018). Chapter 7: Machine learning. In James J. Cochran (Ed.), *INFORMS analytics body of*

knowledge (1st ed., pp. 231–274).Wiley Series in Operations Research and Management Science. Hoboken, NJ: Wiley.
Kim, Y., Eisenberg, D. A., Bondank, E. N., Chester, M. V., Mascaro, G., & Underwood, B. S. (2017). Fail-safe and safe-to-fail adaptation: Decision-making for urban flooding under climate change. *Climatic Change*, *145*(3-4), 397–412.
Kirilenko, A., Kyle, A. S., Samadi, M., & Tuzun, T. (2017). The flash crash: High-frequency trading in an electronic market. *Journal of Finance*, *72*(3), 967–998.
Knight, W. (2017). The dark secret at the heart of AI. *MIT Technology Review*, *120*(3), 54–63.
Lanir, Z. (1986). *Fundamental surprise*. Eugene, OR: Decision Research.
Lazer, D., Kennedy, R., King, G., & Vespignani, A. (2014). The parable of Google Flu: Traps in big data analysis. *Science*, *343*(6176), 1203–1205.
Leveson, N., Dulac, N., Zipkin, D., Cutcher-Gershenfeld, J., Carroll, J., & Barrett, B. (2006). Engineering resilience into safety-critical systems. In E. Hollnagel, D. Woods, & N. Leveson (Eds.), *Resilience engineering: Concepts and precepts* (pp. 95–123). Aldershot, UK: Ashgate Press.
Lewis, M., & Baker, D. (2014). *Flash boys*. New York, NY: WW Norton.
Madni, A., & Jackson, S. (2009). Towards a conceptual framework for resilience engineering. *IEEE Systems Journal*, *3*(2), 181–191.
McCrank, J. (2015, July 8). NYSE shut down for nearly four hours by technical glitch. *Reuters*.
Mersereau, D. (2017, October 12). Hurricane Ophelia is one extremely weird storm. *Popular Science*.
Met Éireann, Government of Ireland. (2018). *An analysis of storm Ophelia which struck Ireland on the 16th October 2017*. Retrieved from https://www.met.ie/cms/assets/uploads/2018/10/Ophelia-1.pdf.
Milly, P. C., Betancourt, J., Falkenmark, M., Hirsch, R. M., Kundzewicz, Z. W., Lettenmaier, D. P., & Stouffer, R. J. (2008). Stationarity is dead: Whither water management? *Science*, *319*(5863), 573–574.
Möller, N., & Hansson, S. O. (2008). Principles of engineering safety: Risk and uncertainty reduction. *Reliability Engineering & System Safety*, *93*(6), 798–805.
Pritzker, P., & May, W. (2016). *NIST special publication 1190: Community resilience planning guide for buildings and infrastructure systems* (Vol. 1). Gaithersburg, MD: U.S. Department of Commerce, National Institute of Standards and Technology (NIST). http://doi.org/10.6028/NIST.SP.1190v1
Roll, R. (1988). The international crash of October 1987. *Financial Analysts Journal*, *44*(5), 19–35.
Rose, R. (2016). Defining analytics: A conceptual framework. *OR/MS Today*, *43*(3). Available at: https://www.informs.org/ORMS-Today/Public-Articles/June-Volume-43-Number-3/Defining-analytics-a-conceptual-framework.
Schumpeter, J. A. (1954). *History of economic analysis*. New York, NY: Oxford University Press.
Shiller, R. J. (1988). Portfolio insurance and other investor fashions as factors in the 1987 stock market crash. *NBER Macroeconomics Annual*, *3*, 287–297.
Simon, H. A., Dantzig, G. B., Hogarth, R., Plott, C. R., Raiffa, H., Schelling, T. C., … Winter, S. (1987). Decision making and problem solving. *Interfaces*, *17*(5), 11–31.
The White House. (2013). *Presidential Policy Directive—Critical infrastructure security and resilience*. Washington, DC: Author.
U.S. Government. (2019). DAP: Digital analytics program. Available at: https://digital.gov/dap/.
U.S. National Oceanic and Atmospheric Administration [NOAA]. (2019). National Hurricane Center (NHC), 5-day probability of tropical-storm-force winds for Hurricane Ophelia. Last Accessed February 2019. Available at: https://www.nhc.noaa.gov/archive/2017/OPHELIA_graphics.php?product=wind_probs_34_F120.
Veloza, O. P., & Santamaria, F. (2016). Analysis of major blackouts from 2003 to 2015: Classification of incidents and review of main causes. *Electricity Journal*, *29*(7), 42–49.
Wears, R. L., & Webb, L. K. (2014). Fundamental or situational surprise: A case study with implications for resilience. *Resilience Engineering in Practice*, *2*, 33–46.
Weick, K. E. (1989). Mental models of high reliability systems. *Industrial Crisis Quarterly*, *3*(2), 127–142.
Weick, K. E. (2009). *Making sense of the organization: The impermanent organization* (Vol. 2). Hoboken, NJ: Wiley.
Woods, D. (2017b). *STELLA report from the SNAFUcatchers workshop on coping with complexity*. Brooklyn, NY: STELLA.