

Avoiding Surprise: Augmenting Anticipatory Thinking with Scenario Explorer

Chris Argenta, PhD; Abigail Browning, PhD; Adam Amos-Binks, PhD, and Matthew Lyle

Applied Research Associates, Inc. Raleigh NC
cargenta@ara.com

Abstract

After being unpleasantly surprised in a high-value and high-risk situation, we often hear phrases like “we had all the dots, but we just didn’t connect them correctly.” This is an example of a failure in anticipatory thinking (AT), the people involved apparently knew all of the facts they needed but did not collaboratively reason over them rigorously enough. This situation differs from when possibilities are considered but determined to be too unlikely to act upon, the first is true surprise while the latter is risk management.

In this paper, we examine some of the cognitive foundations of surprise with an eye towards identifying analytic techniques that augment the support experts in systematic and collaborative anticipatory thinking. We describe Scenario Explorer, our prototype cloud-based collaborative imagination support platform. We present a set of cognitive-informed analytics techniques within Scenario Explorer augment the working memory, manage attention through systematic reasoning, and communicate prospective futures between distributed participants.

Introduction to Avoiding Surprise

Some surprises are pleasant. However, in high-value and high-risk endeavors surprises tend to be feared rather than enjoyed. Throughout history, people concerned with the downside of being surprised have employed a myriad techniques to divine the future, with wildly varying efficacies (see Pickover 2001 for an interesting survey of such techniques). In fact, organizations spend fortunes trying to stay ahead of the trends and predict future events in order to better position themselves and avoid being caught unprepared. Today, predictive analytics is a huge field dedicated to the computationally divining the future. Despite all of these efforts, we continue to be surprised and caught unprepared, even in high-value and high-risk situations where such surprises can be catastrophic.

In this research, we do not attempt to divine things about that future that cannot otherwise be known to the people involved. Instead, we focus on developing techniques to help avoid surprises that result from incomplete prospective cognition. After being surprised, how often do we hear statements like these?

“We had all the dots, but we just didn’t connect them correctly.”

“Some people expected X, others expected Y, but no one expected X and Y.”

“We did not think X was possible, but after Y happened we should have reconsidered at X.”

These are likely cases where rigorous collaborative prospective cognition may have helped participants to avoid being surprised. Our platform, Scenario Explorer, integrates techniques that attempt to address this class of surprises and shift the situation from one of complete surprise to one of managed risk through Anticipatory Thinking (AT).

Motivation: Managing Risk by Avoiding Surprise

The 2019 National Intelligence Strategy (Coats 2019) lists “Anticipatory Intelligence” as the second of seven key mission objectives. It lists foresight, forecasting, and warning essential elements of Anticipatory Intelligence. While data analytics is clearly critical for the Intelligence Community to stay ahead of world events, many of future situations of interest have not been previously observed and are therefore difficult to predict using only analytics. This places emphasis on the subject matter experts to interpret the data for situational awareness, imagine the future trajectories that may evolve, and identify key indicators that might identify which of those trajectories are occurring in the future. Anticipatory Intelligence attempts to do this well in advanced of the actual situation to provide time to understand and prepare.

Risk management is the process of identifying potential futures, estimating the likelihood and impact of these futures, and designing mitigations to manipulate the likelihood and/or impact of trajectories. We generally try to encourage futures we rate positively and avoid those we rate negatively. While we often cannot mitigate all risks, the process of risk management seeks to assess trade-offs and optimize our mitigation efforts. However, we can only perform effective risk management when we can enumerate the potential future trajectories in some form. We can accept risks, but this is a choice or trade-off. We can be surprised when our estimates of likelihood or impact prove incorrect, but these are a matter of quantifying uncertainty. However, when the realized future is significantly different from any of the potential futures that we have anticipated, we are surprised, we lose situational awareness, and we have to reassess our risks. Unfortunately, after being surprised in this way, we have less time, more pressure, and frequently fewer options for mitigation or planning.

We designed Scenario Explorer to include techniques that systematically elicit and aggregate feasible (sometimes unlikely) futures from multiple participants, automate the process of systematically estimating risk, and analyzing the resulting trajectories to extract key indicators and warnings to support situational awareness.

Some Cognitive Foundations of Surprise

The Klein model of Sensemaking (Klein 2007) proposes that the mind creates and uses data structures called frames (sometimes also called schemas) that allow us to store facts in a context. These frames represent our cognitive model of a situation based on our previous experience and knowledge. Frames embed knowledge about how to understand things we have observed, how we interpret some ambiguous knowledge, and what we expect to see in the future. Because of working memory limitations, we generally focus on a single frame. When our frame matches the real world situation, these mechanisms are effective. However, when our current frame fails to match the real world situation we tend to misunderstand, misinterpret, and mispredict. Under the Sensemaking model, when we recognize that our current frame does not fit the situation, we can decide to abandon our current frame and seek another, or modify our current frame to accommodate the new situation.

The Cognitive-Evolutionary Model of Surprise (Meyer 1997) relates to the sensemaking in that it posits an innate and unconscious cognitive mechanism that monitors the alignment of our current cognitive frame and our observations of the world. When the observations match the expectations everything works smoothly just as in the Klein model. However, when the real world observations conflict with the expectations derived from the frame, this monitor-

ing mechanism triggers the emotional response we call surprise. Surprise highlights the discrepancy, and brings ones attention to bear on the process of determining what should be done to bring reality and ones frame into harmony again.

Related research (see survey in Reizenstein 2019) attempts to qualify and quantify surprise. For example, assessing the magnitude of the surprise with respect to this difference, or creating experimental conditions to elicit surprise. While, none of this research appears to be focus directly on avoiding surprise, the concepts defined in this approach appear to lend themselves to this application.

If we accept both of these cognitive models, then we can identify two key places for which intervention might reduce the occurrence of surprise and/or minimize its magnitude:

- **Frame Curation** - If, for a situation of importance, we can improve one's frames/schemas to support a wider range of feasible situations, then we might expect there to be fewer discrepancies. Similarly, when there are discrepancies, we might want alternative frames to which we can more readily switch.
- **Discrepancy Awareness** - If we can instill into relevant frames a sensitivity to key features for discrepancy monitoring, then we might expect to catch frame-changing differences earlier. In the IC, Indicators and Warnings are methods for characterizing these features or events.

Since we natively develop frames through experience and knowledge, we can infer that people with more varied experiences and knowledge should be surprised less often. Possessing a set of related and diverse pre-existing frames might further reduce the delay and cognitive burden of formulating a new frame while in the midst of surprise. As shown in Figure 1, surprises result when our frame anticipates one thing by we are presented with another. The objective is to cultivate variations for important frames so we can easily recognize when a frame change is needed and have alternative pre-considered frames readily available.

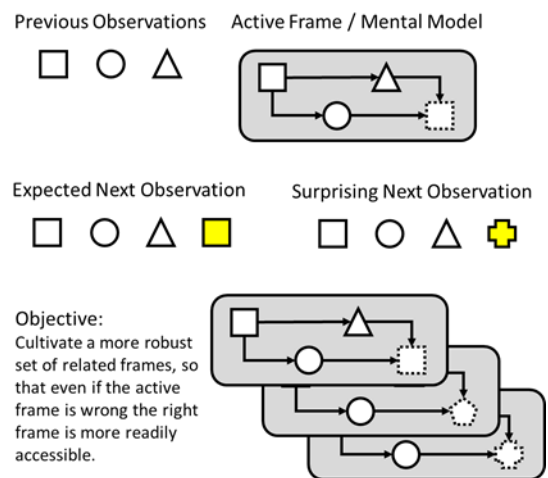


Figure 1. Observations trigger our frames, but surprises result when our frames fail to fit the observations.

One challenge is in effectively sharing those frames and composing elements of different frames. Storytelling is one transfer mechanism for situations that are not easily experienced directly. Use of a branching narrative allows us to consider many potential trajectories and how they differ.

The Role of Anticipatory Thinking in Avoiding Surprise

Anticipatory Thinking (AT) is intentionally divergent thinking that enables a person to better foresee future events (and combinations of events) and their cascade of consequences. If we could always correctly predict the sequence of events that will occur, there would be little need to consider any other scenarios. However, we often view the evolution of the situation with limited and noisy information, so despite our best predictions, we are frequently surprised. AT is similar to prediction and forecasting because they all attempt to correctly identify the scenario that is evolving. However, as shown in Figure 2, prediction and forecasting tend to prioritize Precision (being close to the correct answer and avoiding false alarms) while AT prioritizes Recall (ensuring the correct scenario in the set of answers). These different priorities lead for different analytic techniques.




Figure 2. The trade-off of Precision and Recall defines a space of future-oriented analytics. On the left, we strive to give a single close answer to achieve high precision. While on the right, we strive to ensure the single correct answer is included.

Foresight helps prevent us from finding ourselves in a situation we had not previously considered or imagined. So, if we can imagine the future, why are we still so often surprised by it? Often the answer to this question lies in the complexity of: talking through the potentially relevant things that you (or your group) know; rigorously and systematically developing a set of scenarios; and identifying which scenarios are of interest with respect to various concerns. Human working memory limitations and time constraints reduce us to thinking about a small sample from the range of feasible scenario. This sampling is often biased towards those things that we have already imagined, if not come to expect.

ARA’s Scenario Explorer platform integrates several novel structured analytic techniques that attempt to help analysts more rigorously, efficiently, and creatively explore a wide range of feasible future scenarios by:

- Enabling multiple analysts to work together to converge on a common model of the features, value systems, and timeframes that are relevant for a given project or topic.
- Eliciting the known, expected, or previously imagined future events and scenarios while expressing them consistently with respect to the common model.
- Putting the analyst in a mind-set to imagine feasible future events that have significant effects on the features in the common model.
- Combining the elicited scenarios, automatically composing the events in scenarios (potentially created by different users) to generate novel but sensible scenarios.
- Intelligently querying the combined scenarios to discover key events and potential leading indicators that can determine which scenarios are likely to be occurring.
- Exploring the effects of mitigations and manipulations to assess the sensitivity and uncertainty in trajectories, and the value of triggering potential interventions.

The primary goal of an imagination support is to automate and augment the ability of analysts to conduct these kinds of techniques interactively, reliably, and at scale. The Scenario Explorer platform is being designed and prototyped to balance the cognitive load of the analyst with the computational power of modern computers.

Imagination Support with Scenario Explorer

One thing that forecasting and AT have in common is the understanding that the future can be significantly altered by the sequence of events and states that occur. Conditional forecasts are predictions of the future assuming some given condition is true. For example, a general forecast question might be “what will the average price of an electric car be in 2030?” while a conditional forecast might be “what will the average price of an electric car be in 2030, assuming that Tesla releases the Model 3 as advertised?” The additional qualifier limits the scope of the forecast question because we no longer have to consider alternative possibilities for the Tesla Model 3 release. Conditions qualify a forecast by listing the situations under which it is believed to be correct and improving the accuracy of a forecast can sometimes be dependent on identifying the conditions under which we can expect it to be valid (usually the assumptions under which it was analyzed and computed). One can understand the impact of a conditioning by comparing the forecasted data with the qualification and without, or against forecasts with other conditions. In previous research, we showed that distinct modalities in crowd-sourced forecast responses can often be explained by differences in the conditions assumed by the forecasters.

In AT, we work to develop a divergent set of possible futures, so identifying interesting conditions that would cause us to adjust our forecasts of future states is a primary con-

cern. As shown in Figure 3, a “conditioning event” is a situation in which there are multiple possible and mutually exclusive outcomes, each outcomes resulting in a different effect on the values that would be forecast for a future state. Grouping related outcomes under a common semantic event allows us to systematically evaluate each possible outcome. By quantifying the effects that each outcome has on the features allows us to trace changes the state of the world (within the confines of this project) depending on which outcomes conditions it.

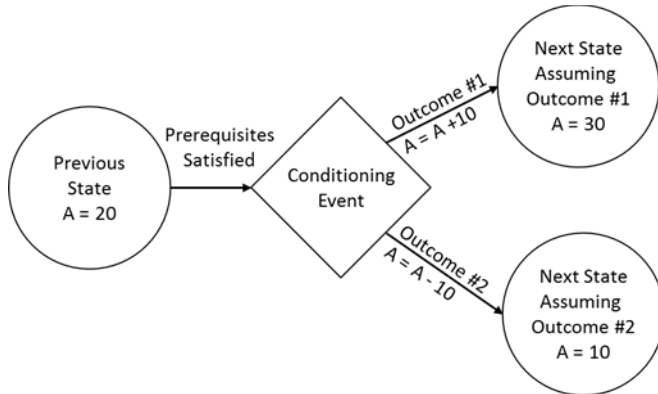


Figure 3. A conditioning event changes the context in which future events occur in a predictable way. A conditioning event can only occur when its pre-requisites are satisfied. It has (at least 1) possible outcomes, each that has a set of (at least 1) effects. The outcomes branch the futures tree, and their effects modify the values of features between the previous and next state.

For example, let a feature (A) represents the number of electric cars sold per day in some area. Today, we may sell 20 cars per day and we are interested in what future sales might look like. We identify that announcement expected next week about federal subsidies for electric cars would likely have a significant effect on that number in the future. We can imagine an announcement that increases incentives could cause the sales numbers to increase. An alternative outcome could be removing incentives which could have the effect of the sales number dropping. Consider two representations:

- **In forecasting**, we might write two conditional forecasts – one would say “what would the sales of electric cars be next week, if the announcements increased incentives?” and the other would end with “if the announcement removed incentives?” The focus is on a feature value given a specific context.
- **In anticipatory thinking**, we qualify when or how the announcement event it might occur, enumerating the possible outcomes, and estimating their effects on the features relative to their previous values. Here, the focus is on defining potential effects in a general context.

These may seem like minor differences, but the conditioning event representation allows us to take conditioning events created in one context and compose them in new

ways to create previously unconsidered scenarios. With AT we consider feasible sequences of events and interactions.

In Scenario Explorer, we represent possible futures as a tree structure. There are nodes are States that hold the values for each feature. States are anchored in both their time and their position in the tree of possible futures. The root of the tree represents our current time (i.e., “Now”) and includes the current values for each feature. Each trajectory can be viewed as a multi-variate time series covering including the changes of feature over time.

Projectors are algorithms that forecast the future value of their assigned feature based on its previous states. An example of a projector could be a function that adds compounding interest and this projector could be assigned to a feature like the amount of money in ones savings account at the bank – even with no intervening events, the amount of money in the account will increase over time. If no projector is assigned then the feature value is propagates unchanged until a conditioning event outcome changes it.

Assessing Risk requires both the likelihood and the impact or value of a State. We define Values/Impacts as a score (0.0-1.0 with an associated color scale) and we train a model to score each State in the tree. This allows display and clustering based on the applied value scale.

Conditioning Events introduce branching to the tree of futures. A scenario (or trajectory through the tree) in which Outcome #1 occurs is different from the one where Outcome #2 occurs. We assign likelihoods to each outcome (and a null outcome), resulting in each trajectory have some likelihood. Conditioning Events can be composed in multiple ways in the tree to produce many possible futures; however must fit the Conditioning Event’s pre-requisite constraints.

Mitigations are actions that can manipulate the likelihoods (e.g., make one outcome more or less likely) and effects (e.g., reduce a feature value change) of conditioning events in the tree or the value/impact of states (e.g., make some characteristics of a state more or less negative). Unlike conditioning events, each mitigation is either on or off.

Storytelling through Futures Building

The primary goal of Futures Building is to elicit interesting and feasible Conditioning Events from the analyst and display them. Futures Building accomplishes this by enabling an analyst to express a scenario that they believe to be relevant to the project. This technique starts with a tree containing only the root Now node. The analyst builds a set of scenarios by sequentially adding conditioning events that logically fit together as a form of narrative (albeit in tree form). While the analyst may be focused on expressing a single scenario (i.e., a single trajectory through a tree of possible futures) that is relevant to them, the system is automatically populates the tree with all possible compositions of the events elicited.

Cognitive Basis

Futures Building elicits branching narratives independently from multiple participants decreasing anchoring and priming effects that might occur during brainstorming.

When to Use It

Futures Building is useful when the analyst has a set of conditioning events that they wish to enter into the system that share common narrative thread. Ideally, each view instance would stand on its own as a short story of what might happen in the future. This is the default elicitation view for the Imagination Support platform.

Value Added

When analysts express their knowledge as conditioning events in the Futures Building view, they are sharing their knowledge with the team and the system using a common project model. A Futures Building session can be performed alone or as a team. When/if the entered conditioning events appear on other views they can be traced back to their origin view – a Futures Building view provide a way to group related conditioning events to more effectively explain their meaning by offering context.

The Method

Scenario Explorer treats each instance of Futures Building as a separate View. Users can collaborate on a shared Futures Building scenario or refer back to it to understand the intention of a conditioning event that might come up in another context within the system. Behind the scenes, the conditioning events being elicited are automatically being applied at the project level. So, instead of one large Futures Building View, we recommend that users create many smaller/simpler Views around specific domain-relevant narrative topics.

- Step 1. Create a New Futures Building View
- Step 2. Add New Conditioning Events
- Step 3. Integrate the Conditioning Event into a Tree

Using Extreme States to Change Your Perspective

The Extreme States technique derives from the concept of a Pre-mortem analysis (Klein 2007). Rather than starting with a blank slate and working forward, Extreme States starts with a goal State and attempts to elicit conditioning events that bring trajectories closer to this goal state. This is intended to aid the analyst in shifting their perspective to imagining how the situation might have gotten from Now to the Extreme State.

Cognitive Basis

Changing your perspective helps trigger frames that may be less accessible when in a different mindset.

When to Use It

Extreme States should be used after analysts have exhausted the conditioning events they previously expected but the tree of futures still does not contain the degree of

variation or divergence desired. If the conditioning events required to reach the goal already exist in the project, Scenario Explorer detects them and fills in the trajectory.

Value Added

It is sometimes difficult to diverge from the events we expect. Extreme States provides a method of shifting ones perspective from looking forward at a wide-open expanse of possibilities to looking back and explaining how the situation got to that point. The trajectory directly elicited is not the primary objective since it likely highly unlikely, however the individual conditioning events introduced can also be used to enrich other scenarios and introduce trajectories that are less extreme and more likely.

The Method

Scenario Explorer treats each instance of Extreme States as a separate View. An Extreme States view will likely be relatively small, since once a trajectory between Now and the Extreme State is found, the process is complete.

- Step 1. Creating an Extreme State View
- Step 2. Specifying the Extreme State
- Step 3. Checking for an Existing Trajectory
- Step 4. Adding Conditioning Events

Smart Queries Extract Warnings and Indicators

Smart Query is a technique for extracting knowledge from Scenario Explorer rather than eliciting knowledge from analysts to put into the system. Smart Queries allow one ask questions about a large number of trajectories in the tree in an intuitive manner. It provide results based on clusters of terminal states (based on similarity to query features or values/impact scores) and identifies and scores sensitivity and selectivity of conditioning event sequences that best discriminate these clusters.

For example, if we perform a Smart Query on a feature that represents the average price of electric cars, all futures could cluster into ranges of the cars being generally expensive, mid-range, or cheap. We might find that conditioning events and outcomes such as “major tax incentives given” would exist on the trajectories for mid-range and cheap, but not expensive.

Cognitive Basis

Identifying patterns that lead to futures helps analysts become sensitive to key events that indicate a class of futures.

When to Use It

A Smart Query should be used when the analyst wishes to determine which conditioning events play a critical role in differentiating possible futures with respect to a specific set of features. It allows the analyst to gain new insight either by seeing how conditioning events and outcomes (regardless of their source) influence futures.

Value Added

Smart Queries extract short stories from large data sets by finding the most salient conditioning events and outcomes for understanding how a feature might evolve. The critical conditioning events and outcomes can be used as leading indicators to determine which cluster an evolving scenario is likely to be in.

The Method

Scenario Explorer treats each instance of Smart Queries as a separate view, since they can be computationally expensive to produce; the query results are stored and updated on demand.

- Step 1. Creating a Smart Query View
- Step 2. Specifying the Query Features or Value/Impacts
- Step 3. Execute the Query

Risk Mitigation Analyses

Risk Analyses allows analysts to compare the effects of performing (or not performing) specific mitigation actions in response to triggering events within the trajectory tree. Comparing a customizable plot the trajectory risk scores side by side allows the analyst to visualize the risk reduction.

Cognitive Basis

Risk Analyses allows users to experience the effect of potential mitigation actions they imagine thus attaching additional action-oriented knowledge to relevant frame.

When to Use It

A Risk Analyses is used when the analyst wishes to understand the effect that a set of specific actions might have on the future. This can be used to determine which mitigations are most effective and what risks might be immune to them.

Value Added

A Risk Analyses automates the process of applying a set of actions in the proper context for a complex set of future trajectories. Even with simple actions, this often taxes working memory and is highly error prone when done manually.

The Method

Scenario Explorer performs Mitigation Analyses by comparing the likelihoods and impacts of potential states both with and without a selection of mitigation actions enabled.

- Step 1. Create Mitigation Analysis View in a project.
- Step 2. Define a set of Mitigations/Interventions
- Step 3. Toggle the mitigations that are enabled
- Step 4. Compare the plots and analytic summaries

What If Analyses for Understanding Sensitivities

What If Analysis is used to explore the sensitivities of the scenarios represented in the tree. A What If Analysis starts with an existing view and allows the analyst to create an

overlay on the project in which they can modify specific feature values and conditioning event outcome effects at will and without changing other views.

Cognitive Basis

Understanding sensitivities in the futures considered helps establish confidence bounds and identify assumptions that may invalidate our analyses.

When to Use It

A What If Analysis should be used when an analyst has questions about specific features and how sensitive things are to those values. For example, they can change the Now state and see the effects ripple through the tree.

Value Added

A What If Analysis allows an analyst to explore nuances of a specific feature or view of the project. These overlays do not modify or feed back into the data.

The Method

Scenario Explorer allows What If Analysis Views to modify feature values in states generated from any other view. It does this by cloning the trajectory data and only recalculating the portions of trajectories that are changed.

- Step 1. Create a What If Analysis Overlay
- Step 2. Edit the values of features in any State
- Step 3. Update the tree with the new consequences
- Step 4. Compare the original and overlay trees

Future Work

While many of the capabilities outlined in this paper exist today, many are also in a state of partial implementation, and others only mostly design. Our plan is to complete the outlined capabilities and evaluate our platform with analysts and subject matter experts.

Additionally, we have planned to continue the development of Scenario Explorer to include automated handling of historic and streaming data to enable updating of our trajectories over time and improve projector performance.

Finally, we believe that we have just opened the door to vast opportunities for additional AT analytic techniques. As we continue to develop analytics we have designed, we wish to open our platform up for other researchers in the AT community to contribute to the tools and execute experiments that leverage the quantification capabilities of platform. Please feel free to contact us to discuss collaborations.

Conclusions

Today, we are free to imagine the future in highly unconstrained ways. However, this freedom comes with some costs, which include not having systems that support, extend, and collaborate your imagination. As a result, our prefactual reasoning is limited and we are too frequently left

surprised. This is problematic in high-value and high-risk situations. This is more made more frustrating when we have the knowledge we needed to avoid the surprise but simply failed to apply enough rigor to our analyses.

By understanding cognitive models underlying sense-making and the experience of surprise, we can design tools that help us avoid surprise and better manage risks. Scenario Explorer is one such tool, and we have outlined five structured analytic techniques that we have developed to improve anticipatory thinking and to avoid being surprised.

Acknowledgement

The authors wish to acknowledge that this research has been funded by the Laboratory for Analytic Sciences (LAS). We would like to thank LAS and our government, academic, and industry partners for their support and cooperation in creating and nurturing a research community around Anticipatory Thinking and for their support of this research effort.

References

- Coats, D. 2019. National Intelligence Strategy of the United States of America. Office of the Director of National Intelligence, DC.
- Klein, G., Phillips, J. K., Rall, E., & Peluso, D. A. 2007. A Data/Frame Theory of Sensemaking
- Klein, G. 2007. Performing a project PreMortem. *Harvard Business Review*, September, 18–19.
- Meyer, W.-U., Reisenzein, R., & Schützwohl, A. 1997. Towards a process analysis of emotions: The case of surprise. *Motivation and Emotion*, 21, 251–274.
- Pickover, C. 2001. Dreaming the Future: The Fantastic Story of Prediction, *Prometheus Books*, New York
- Reisenzein, R., Horstmann, G., Schützwohl A. 2019. The Cognitive-Evolutionary Model of Surprise: A Review of the Evidence. *Topics in Cognitive Science*. 2019 Jan;11(1):50-74. doi: 10.1111/tops.12292. Epub 2017 Sep 23.