SCIENTIFIC REPORTS

OPEN

Limits of Risk Predictability in a Cascading Alternating Renewal Process Model

Xin Lin^{1,2}, Alaa Moussawi^{1,3}, Gyorgy Korniss^{1,3}, Jonathan Z. Bakdash⁴ & Boleslaw K. Szymanski^{1,2}

Most risk analysis models systematically underestimate the probability and impact of catastrophic events (e.g., economic crises, natural disasters, and terrorism) by not taking into account interconnectivity and interdependence of risks. To address this weakness, we propose the Cascading Alternating Renewal Process (CARP) to forecast interconnected global risks. However, assessments of the model's prediction precision are limited by lack of sufficient ground truth data. Here, we establish prediction precision as a function of input data size by using alternative long ground truth data generated by simulations of the CARP model with known parameters. We illustrate the approach on a model of fires in artificial cities assembled from basic city blocks with diverse housing. The results confirm that parameter recovery variance exhibits power law decay as a function of the length of available ground truth data. Using CARP, we also demonstrate estimation using a disparate dataset that also has dependencies: real-world prediction precision for the global risk model based on the World Economic Forum Global Risk Report. We conclude that the CARP model is an efficient method for predicting catastrophic cascading events with potential applications to emerging local and global interconnected risks.

A generalized Alternating Renewal Process model referred to as Cascading Alternating Renewal Process (CARP), has been recently proposed for dynamically modeling a global risk network represented as a set of Poisson processes¹. The aim was the recovery of the directly observable and hidden parameters of the model using historical data to predict future activation of global risks, and especially cascades of such spreading risk activations^{2,3}. This approach raises questions about the reliability of such recovery and the dependence of the prediction precision on the complexity of the model and the length of its historical data. Since most of the world's critical infrastructures, including the global economy, form a complex network that is prone to cascading failures⁴ this question is important but difficult to answer given the for lack of ground truth data.

Here, we discuss how the limits of prediction precision can be established for a specific system as a function of the input data size by simulating the CARP model with known parameters to generate many alternative ground truth datasets of arbitrary length. This enables us to the given system to establish the precision bound with the given length of historical data. We illustrate our approach with a model of fire risks in an artificial city that consists of modular blocks of diverse housing that are easy to assemble into an ever-growing complex system. We first simulate the fires in such cities of varying sizes, over varying time periods, in each case using different random generator seeds to create many alternative ground truth datasets. Then, using Maximum Likelihood Estimation (MLE)⁵, we recover the parameters and compare them with their values used for simulations. We also measure the recovered parameter precision as a function of (i) the complexity of the system (in our case the size of the cities) and (ii) the number of events present in the historical data. Finally, using real-world data with the developed methodology, we assess the precision of parameter recovery in the World Economic Forum model¹.

Risk Modeling. Most quantitative risk analysis models (e.g., Value at Risk and Probability-Impact models) systematically underestimate the probability and impact of worst-case scenarios (i.e., maximum loss for a given

¹Social and Cognitive Networks Academic Research Center, Rensselaer Polytechnic Institute, Troy, NY, 12180, USA. ²Dept. of Computer Science, Rensselaer Polytechnic Institute, 110 8th Street, Troy, NY, 12180, USA. ³Dept. of Physics, Applied Physics and Astronomy, RPI, 110 8th Street, Troy, NY, 12180, USA. ⁴US Army Research Laboratory, Aberdeen Proving Ground, 21005, MD, USA. Correspondence and requests for materials should be addressed to B.K.S. (email: szymab@rpi.edu)

confidence level, typically “tail” outcomes)²⁵ for catastrophic events such as economic crashes, natural disasters, and terrorist attacks⁶. Underestimation in such models is due to speciously assuming the sequences of risk materializations can be represented by normal, independent, and identically distributed (IID) random variables^{7,8} thus discounting the potential of interdependencies and interconnections between events.

Few quantitative risk models capture the non-IID properties of risk factors and their impacts⁹. The Havlin model uses branching to predict cascading failures (e.g., power grids, communication networks)¹⁰. A small fraction of initial failures could cause catastrophic damage in mutually dependent systems. The authors used the percolation theory and detected a phase transition for the robustness and functionality of the interdependent networks. The Ganin model of resilience¹¹ provides an analytical definition to criticality and a method for determining it to design more resilient technical systems¹². While both techniques go beyond simple risk models, which assume risks exhibit independent probabilities and impacts, neither aims to quantify the limits of predictability for interconnected risks. In contrast, the CARP model is a novel method in which interdependencies and interconnections are explicitly represented, and the model parameters are recovered from historical data using Maximum Likelihood Estimation. Moreover, the CARP model offers a quantitative risk assessment for interconnected conceptual models such as Reason’s Swiss cheese model of failure¹³. In Reason’s model, defenses preventing failure are layers (of Swiss cheese) and when the holes of the layers align a risk may materialize. Previously, Reason’s model has been formalized using percolation theory¹⁴, but this formalization has not been validated.

Alternating Renewal Process. The definition of Alternating Renewal Process (ARP) originated in renewal probability theory. A simple renewal process alternates between two states: the normal state, which accounts for its operational time and the abnormal (failure or repair) state, which accounts for its holding (non-operational) time. Each state is governed by a Poisson process specific to this state^{15,16}. The ARP can be used to find the best strategy for replacing worn-out machinery^{17,18}.

The CARP model expands ARP in several ways. The most important extension is to allow for defining risks as a network of risks with the given set of weighted edges representing interdependencies and probabilities of state transitions associated with each edge. Since state transitions are often observed without distinction for the inner or external, edge induced transitions, both probabilities are treated as hidden variables, since only their joined effect is observed directly. Another generalization is not restricting the number of states to two. Finally, we associate with CARP an MLE procedure for recovery of CARP parameters from historical data of the system evolution over time.

Model Structures. The city modeled is diverse. A small fraction of houses is large and placed sparsely with low fire propagation probabilities and high recovery rates. A larger fraction of houses with a medium size have a higher spatial density and fire propagation probabilities. Finally, the most densely packed small houses have the highest fire propagation probabilities, and lowest recovery rates.

Houses are grouped into blocks stitched together into a continuous city as shown in Fig. 1(a). Large and small houses sit on the West and East boundaries of the city, respectively. The North and South boundaries border all three types of houses. The houses are placed on rectangular lots whose size is commensurate with the size of the house. A large house occupies a square lot of one by one unit size. There are four such squares horizontally laid out with no space in between. Each medium house has a rectangular lot of half unit length vertically by one unit length horizontally. Each block holds 8 medium houses. Finally, repeating a similar pattern, small houses sit on a rectangular lot of one by a quarter unit size. One block contains 16 small houses. Distances between two adjacent houses vary from 1 for large houses to 1/2 for medium and 1/4 for small houses. A basic block holds 28 houses on 12 square units of land. A city of arbitrary large size can be built by repeatedly adding blocks to it. Figure 1(c) shows the degree of each house, which is defined as the number of houses within a fixed distance. As shown in Fig. 1(a), all houses within a fixed distance may catch fire from the burning center house. Hence, the small houses have a relative higher degree because of the high density of nearby houses. The surrounding houses on the boundary of the city and the border of different communities experience a slightly smaller degree compared with other same-size houses. As we explain later, we assume that, the likelihood of a house catching fire is determined by two factors: the materials the house is made of and the density of its neighbors.

Results

Unlike real-life, in simulations, we can arbitrarily vary the length of time over which we collect historical data and produce many variants of such data to measure the prediction precision of our model. We start with a mixed model with 8 large houses, 16 medium houses and 32 small houses, for a total of 56 houses. The parameter recovery is applied at seven different intervals: 100, 200, 400, 800, 1600, 3200 and 6400 time steps. The parameter recovery is run on 50 versions of historical data created by simulations running with different seeds for the random number generator to account for the randomness of the stochastic processes. To quantify the accuracy of parameter recovery, we compute the relative error between the recovered values and the target values used for creating ground truth data.

As shown in Fig. 2, the blue vertical dots represent the relative error of parameter recovery in one realization. There are 50 blue dots for each particular length of the training dataset. The scattering of blue dots represents the variation of parameter recovery accuracy. The red curve is the average value among these 50 realizations, which is very close to zero. Obviously, the variance is a better measurement for the precision since variance considers

^aResilience has a variety of definitions, but the common definition is the planning, preparation, absorption, recovery, and adaptation of systems under negative events¹¹. Predicting the limits of risk is consistent with the common definition.

Figure 1. (a) Basic block of a city. A sample city consists of four basic blocks. Three types of houses are represented as nodes. Red circles show the range of external fire spread. All houses in this circle may catch fire from the burning center node. (b) Fire propagation dynamics. This diagram shows state transitions. The nodes represent a fully operational house (green), a burning house (red) and a recovering house (blue). The lines show state transitions to a new state as solid lines, and state transitions to the same state as dotted lines. (c) House degrees in an 8 block 224 house city. The degree increases horizontally from left to right as house types change from large to medium and to small. Centrally placed small houses have the highest degrees. (d) Fraction of time each house is on fire. Results are averaged over 100 independent realizations. Each simulation runs 10^6 time units. The houses are fully operational initially. The simulations used parameter values listed in Table 2.

Parameters	α	β	γ	δ
Target value	0.00800	0.01200	0.01600	0.03200
Recovered value	0.00799	0.01199	0.01596	0.03204
Recovered value $+\sigma$	0.00853	0.01162	0.01622	0.03274
Recovered value $-\sigma$	0.00747	0.01237	0.01570	0.03134

Table 1. Parameter values for simulation of historical data with range of parameter values.

the positive and negative errors instead of canceling them out. The average of relative errors tends asymptotically to zero and the variance shrinks according to the power law with an increase in sample size which is the length of historical data series in this case. This trend is consistent with the asymptotic behavior of the MLE method¹⁹. When the sample size is very large, the relative error of realization follows a normal distribution with 0 mean value and a finite yet small variance. More data decreases relative error, and therefore it is useful to find a balance between run time and prediction accuracy.

The relative errors of $\hat{\alpha}$ and $\hat{\beta}$ are larger than that of other parameters. This is due to the combined effects of two Poisson processes causing the same transition. As defined, α and β represent the intensity of internal and external fire ignition processes respectively. In real life, it is often hard to determine the actual reason. During the parameter recovery, these two parameters influence chances of each other to start a fire, which impacts the computation of likelihood function. This nonlinear effect tangles the errors of $\hat{\alpha}$ and $\hat{\beta}$ together as larger value of $\hat{\alpha}$ can be compensated by smaller value of $\hat{\beta}$ and vice versa.

We also compute the standard deviation of relative error of recovered parameters. Figure 3(a) shows that the distribution of the standard deviation follows a power law. The standard deviation decreases very quickly and

Name	Definition of variables
$N_{1,i}$	Likelihood of house i to catch fire (internally or externally)
$N_{2,i}$	Likelihood of fire to be extinguished and reconstruction started of house i
α	Control parameter for the internal fire materialization process
β	Control parameter for the external fire materialization process
γ	Control parameter for the recovery process
δ	Control parameter for the fire extinguished process
λ_i^{int}	Intensity for the process of starting fire internally in house i : $\lambda_i^{int} = -\alpha \ln N_{1,i}$
λ_{ji}^{ext}	Intensity for the process of externally transferring fire from house j to house i : $\lambda_{ji}^{ext} = -\beta \ln N_{1,i}$
λ_i^{rec}	Intensity for the process of completing recovery of house i : $\lambda_i^{rec} = -\gamma \ln N_{2,i}$
λ_i^{etg}	Intensity for the process of extinguishing fire in house i : $\lambda_i^{etg} = -\delta \ln N_{2,i}$
p_i^{int}	Probability of internal fire ignition in house i : $p_i^{int} = 1 - e^{-\lambda_i^{int}} = 1 - N_{1,i}^\alpha$
p_{ji}^{ext}	Probability of external fire transfer from house j to house i : $p_{ji}^{ext} = 1 - e^{-\lambda_{ji}^{ext}} = 1 - N_{1,i}^\beta$
p_i^{rec}	Probability of recovery of house i : $p_i^{rec} = 1 - e^{-\lambda_i^{rec}} = 1 - N_{2,i}^\gamma$
p_i^{etg}	Probability of extinguishing fire in house i : $p_i^{etg} = 1 - e^{-\lambda_i^{etg}} = 1 - N_{2,i}^\delta$
House type	λ_i^{int} λ_{ji}^{ext} λ_i^{rec} λ_i^{etg} p_i^{int} p_{ji}^{ext} p_i^{rec} p_i^{etg}
Large	0.0073 0.0109 0.0257 0.0515 0.00730 0.01094 0.02542 0.05020
Medium	0.0096 0.0144 0.0192 0.0385 0.00959 0.01434 0.01908 0.03779
Small	0.0129 0.0193 0.0146 0.0293 0.01279 0.01913 0.01455 0.02889

Table 2. Definition of variables, intensities of Poisson processes, and the probabilities of the corresponding Bernoulli processes for the parameter values: $(N_{1,i}, N_{2,i}) = (0.4; 0.2)$ for large, $(0.3; 0.3)$ for medium, and $(0.2; 0.4)$ for small houses. $\alpha = 0.08$, $\beta = 0.012$, $\gamma = 0.016$ and $\delta = 0.032$.

Figure 2. Parameter recovery in the fire-propagation model. The x-axis includes seven time intervals: 100, 200, 400, 800, 1600, 3200 and 6400 time steps. The y axis shows the relative error for the recovered parameters. Using parameter values shown in Table 2, parameter recovery was run on 50 different historical datasets generated by simulations with different seeds for the random number generator; the results of these runs are represented by blue dots. The red dashed curves show the average values of the relative error. The visible trend is that the average of relative errors tends asymptotically to zero and the variance exhibits power law decrease as the number of time steps increases, which means more training data improves the performance of parameter recovery.

then slowly as historical data size increases. The double-logarithmic plot in Fig. 3(b) has slope close to -0.5 , which shows that the standard deviation decreases in a power-law fashion as the training data size increases.

For real-world case processes, it is impossible to get multiple historical datasets for parameter recovery, yet recovery error based on single dataset may be different from the one based on the average of such recovered values on multiple datasets. Hence, we study how sensitive our methodology is to imperfect input datasets. In

Figure 3. Comparison of standard deviation of relative error of parameter recovery. We use seven simulation time intervals: 100, 200, 400, 800, 1600, 3200 and 6400 time steps. There are 56 houses. Using parameter values shown in Table 2, parameter recovery was run on 50 different historical datasets generated by simulations with different seeds for random number generator. The plots show standard deviation of the relative error of recovered parameters averaged over 50 runs in linear scale (a) and logarithmic scale (b).

order to test this sensitivity, we compare four cases of simulations in which we record the average length of time in each state and the number of emerging fires during four simulation periods: 400, 800, 3200 and 6400. The first case is using the target values of parameters. The second case is using the averaged value of parameters recovered from 50 independent realizations. The third case employs adding one standard deviation σ to the average value of recovered parameters. The last case is subtracting one standard deviation σ from the average value of recovered parameters. The four sets of parameters employed in our simulation are listed in Table 1.

The average values of recovered parameters come from 50 independent realizations with 6400 time steps of training data. The standard deviations are: $\sigma_{\alpha} = 0.00053$, $\sigma_{\beta} = 0.000372$, $\sigma_{\gamma} = 0.00026$, $\sigma_{\delta} = 0.0007$. When adding one standard deviation σ to average value of $\hat{\alpha}$, we subtract σ from $\hat{\beta}$ and vice versa. We assume complementary recovery errors on α and β since both of them contribute to the fire triggering process so the maximum value of one is likely to be reached at the minimum of the other. This corresponds to the assumption that the unique historical dataset produces parameter values within one σ of their average values, more stringent assumption might require using broader interval around average values of parameters. In simulation, we record how long each house stays fully operational, on-fire and in recovery and record the number of new fires during the simulation. We generate over 50 different historical datasets and gather results upon reaching five simulation periods: 400, 800, 1600, 3200 and 6400. With each dataset, we run 50 independent realizations to assess the predictability of our methodology. There is a trade-off between the running time and precision. To save computational effort, we set the number of independent realizations at a moderate value of 50. Figure 4 shows that all four parameter estimations have similar precision. In Fig. 4(a), the average duration of the fully operational state is almost the same for four cases since two parameters α and β have opposite influences on fully operational houses. In Fig. 4(b,c), the gap between the cases of estimated parameters $\pm\sigma$ is very small. In Fig. 4(d), the number of emerging fires for all cases are increasing linearly as a function of time. The largest relative error is in predicting the length of fire, and even that is just below 2%. For simulation efficiency, we set the target values of parameters in such a way that the duration of staying in a state is short and of the same order for all states, while in reality they are different, for example for any house staying on fire is orders of magnitude shorter than getting rebuilt. We intentionally set large and comparable target values to generate sufficient state transitions and test the quality of parameter recovery in an arbitrary case. Table 1 shows that the estimated parameters approach the target values very well, even if the target values are of the same order as the initial settings.

In addition to the parameter recovery from different lengths of historical datasets, we study how the precision of parameter recovery varies against the complexity of the system and the standard deviation of number of fires starting each day. Figure 5(a,b) show the relative errors of recovered parameters for various city sizes (number of houses): 28, 56, 112, 224 and 448. Here, we compare two cases. In the first case, we keep three types of houses and use the parameter values shown in Table 1. The blue dots in Fig. 5(a,b) represent the results of the first case. In the second case, we initialize all houses with the same value of $N_{1,i} = 0.3$ and $N_{2,i} = 0.3$, which are equal to the values for middle houses. The red dots in Fig. 5(a,b) represent the results from this case. Therefore, we remove the influence of house types on the results. In both cases, the length of historical data is 1600 time steps and we run over 20 different datasets. Only two parameters ($\hat{\alpha}$ and $\hat{\beta}$) are shown in this figure since they are involved with the fire igniting process and another two parameters have similar trend. We find that the parameter recovery in a larger city has a smaller mean value and variance of relative errors. As the city's size increases, there are more state transitions within specific periods, leading to a more precise recovery of our parameters. To study the impact of intensity of emerging fires on recovery precision, we change the values of $N_{1,i}$ which define the intensities of fires. We compare three cases: $\sqrt{N_{1,i}}$ in red, $N_{1,i}$ in blue and $N_{1,i}^2$ in cyan. Figure 5(c,d) show the relative error of recovered

Figure 4. Impact of recovered parameter uncertainty on the prediction of model dynamics. The average length of time in each state is recorded for five time intervals: 400, 800, 1600, 3200 and 6400 time steps. The number of houses is 56. Using parameter values shown in Table 1, parameter recovery was run on 50 different historical datasets generated by simulations with different seeds for random number generator. (a) shows how long a house stays in the fully operational state which is determined by internal and external fire triggering processes. Since the effect of increasing α is reduced by the effect of decreasing β and vice versa, their combined effects is smaller than in the case of other parameters. (b,c) show the average length of time for the states of on-fire and recovery. Only one parameter determines the length of time so the gap is larger than (a) but still low. (d) shows the number of new fires during the simulation. All cases yield similar results.

Figure 5. Parameter recovery in various scenarios. (a,b) show relative errors of recovered $\hat{\alpha}$ and $\hat{\beta}$ for five city sizes: 28, 56, 112, 224 and 448. Blue color represents the city with three types of houses and different values of $N_{1,i}$ and $N_{2,i}$. Red color represents the city with identical value of $N_{1,i}$ and $N_{2,i}$. The dashed curve shows the mean values of relative error over 20 realizations and each dot represents one realization. The length of historical dataset is 1600. Parameter values are shown in Table 1. (c,d) show relative errors of recovered $\hat{\alpha}$ and $\hat{\beta}$ for three cases of $N_{1,i}$. Red dots represent the case of $\sqrt{N_{1,i}}$, blue represents the case of $N_{1,i}$ and cyan represents the case of $N_{1,i}^2$. The results come from 20 independent realizations and 6 different lengths of historical dataset: 100, 400, 800, 1600, 3200 and 6400. $N_{1,i}$ is 0.4 for large 0.3 for medium and 0.2 for small houses. $\alpha = 0.08$, $\beta = 0.012$, $\gamma = 0.016$ and $\delta = 0.032$.

Figure 6. Performance comparison of recovered parameters in the global risk network. There are 125 cases of estimated parameters. 4 different periods of time steps (120, 240, 480 and 960) are used to estimate parameters and simulate future behaviors. For each set of estimated parameter, we finish 20 realizations and average the number of materialization during the simulation. $\pm\sigma$ boundary is determined by removing 39 sets of estimated parameters with worst performance from ground truth case. (a) Histogram of number of materialization. (b) shows the boundary of $\pm\sigma$ performance for number of materialization in each period. (c) shows the number of materialization for each risk in the case of 960 time steps.

$\hat{\alpha}$ and $\hat{\beta}$. The model with more emerging fires at each day enjoys a smaller variance of relative errors. The reason is similar to that shown above: more state transitions in the historical dataset and higher precision of the recovered parameters.

Parameter Recovery Precision in Global Risk Network. Here, we show how to apply the presented approach to a disparate, real-life dataset of the global risk network, which, as with fires in cities, exhibits spreading risk activation¹. Using CARP, we estimate hidden parameters of global risks previously modeled using an Alternating Renewal Process¹. Experts from the World Economic Forum 2013 Global Risks Report²⁰ define the properties of 50 global risks grouped into five categories: economic, environmental, geopolitical, societal, and technological. These assessments include the likelihood, impact of materialization, and connections of each risk. We take the advantage of this crowd-sourcing assessment to build an interconnected network to simulate risk propagation through the system.

In the global risk network, each risk has binary states (normal and active), and the state transitions follow Poisson processes. The difference is that in the global risk network, there are three state transitions instead of four in the fire-propagation model: 1, internally igniting fire process transitioning the risk from normal to active state; 2, externally starting fire process also transitioning risk from normal to active state; and 3, starting recovery process transitioning risk from active to normal state. Hence, we have three control parameters (α , β , γ) and recover the optimal parameter values from historical events. Using data collected over a 156-month period and Maximum Likelihood Estimation, we recover the following parameter values: $\alpha = 0.003038$, $\beta = 0.00117$ and $\gamma = 3.5561$. They control internal risk materialization (α), external risk materialization (β) and recovery (γ) processes, and their detailed definitions can be found in ref.¹. Once we obtain the parameter values, we can simulate the stochastic processes driving model evolution. As in the case of the fire-propagation model presented here, we use these recovered parameters as the ground truth parameters for establishing the parameter recovery precision.

We apply the same method used in Fig. 4 to this global risk network. For the precision evaluation, we obtain multiple recovered-parameters from different lengths of alternative historical data to predict the number of risk materialization for a selected period repeatedly with different random generator seeds. First, based on the ground truth parameters recovered from the real historical data (α , β , γ), we generate 120, 240, 480 and 960 monthly time steps of alternative historical data (representing 10, 20, 40 and 80 years) for parameter recovery. Then, 125 cases of parameter recovery (α_i , β_i , γ_i , $i = 1 \dots 125$) are finished for each period of time. Next, using the recovered parameters from both real and alternative historical data, we complete 20 realizations for a prediction of 4 periods: 120, 240, 480 and 960. In each period, we use the average value of risk materialization among all 20 realizations to measure the performance of each simulated case. After that, we compute the relative error of the average number of risk materializations between the case with alternative recovered-parameters (α_i , β_i , γ_i , $i = 1 \dots 125$) and the ground truth parameters (α , β , γ). In the end, we remove the 39 simulation cases with the worst performance

(with the largest absolute relative error). The remaining 86 cases (68.8% of 125 cases) of results determine the $\pm\sigma$ boundary for the performance of the system with estimated parameters.

Figure 6(a) shows a histogram of the number of materializations in each case. Dash curve represents a Gaussian fitting over the histogram. As the length of simulation period increases, the distribution of the number of materialization gradually approaches a normal distribution. Meanwhile, the distribution shifts to the right and gets close to a steady level. In the case with longer time steps for parameter recovery and prediction simulation, the predictability is more consistent, and variance shrinks faster. Figure 6(b) shows the boundary of $\pm\sigma$ performance for the number of materializations in each period. It is evident that the distance between the boundaries is decreasing as we increase the length of the period, which implies a higher precision of prediction. Figure 6(c) shows the number of materializations for each risk in the case of 960 time steps. The difference between three cases is slight indicating a consistent prediction of estimated parameters.

Discussion

The CARP model is used to simulate and then recover parameters of heterogeneous stochastic processes. First, we created a model of fires in the cities that we use to illustrate our approach. Using assumed parameters values, we generated several historical datasets and used them to measure parameter recovery precision. The results confirmed that the accuracy of our method increases as the amount of data increases even in the presence of parameters hidden from direct observations.

Applying our approach to real-life cases, we started with the recovery of the model parameter values based on unique and limited real-life ground truth data. Then, using these values as ground truth, we finished simulations to create many alternative historical datasets. Then using these historical datasets, the parameters are recovered by applying MLE method. Next, we compared the results to the assumed ground truth values to measure the accuracy of recovery. The standard deviation of relative error of parameter recovery exhibits a power-law decay with an exponent value of -0.5 as the training data size increases. The resulting statistics enable us to verify the reliability of predictions based on originally recovered parameters. We did so by comparing original predictions to predictions based on parameters differing from their average values by the desired multiple of their standard deviations as was demonstrated for the city model. We record the duration of each state (normal, on-fire, and recovery) and the number of emerging fires within the simulated period. The largest relative error of these variables is just below 2%.

In conclusion, we showed that the CARP model is a novel approach to predict and simulate risks. It is particularly useful for modeling cascading catastrophic events and thus has potential applications for analyzing local and global risks. Local risks were demonstrated using simulated fires in cities. However, the CARP model was also successfully used to model global risks in earlier work¹. A better understanding of globally networked risks is critical to predicting and mitigating them²¹. Most of the world's critical infrastructure forms a complex, interconnected network prone to cascading failures with potentially devastating consequences to global stability²². Quantifying the limits of risk prediction, which are bounded by the amount of data, may inform earlier planning and thus potential mitigation of danger of risks spreading and its adverse consequences.

Methods

Discrete Model. Based on the structure of the fire-propagation model, we can simulate the fire cascades throughout the entire network using CARP. At time t , each house is either in state -1 (recovery), state 0 (fully operational) or state 1 (on-fire). Houses in the recovery state are under reconstruction and are immune to fire. Hence only fully operational houses are susceptible to fire. The burning (on-fire) state with certain probability switches to a state of recovery. Each house alternates between these three states.

The state transition is invoked by four types of Poisson processes. A house i transits from state 0 to 1 (on-fire) for internal reasons according to a Poisson process with the intensity λ_i^{int} . In this event, a fire starts inside of the house, caused by events such as overheating of electrical appliances, unattended stoves or other possible accidents. This house can make this transition if its neighbor j ignites the fire externally through a Poisson process with the intensity $\lambda_{ji}^{\text{ext}}$. A transition from state 1 (on-fire) to -1 (recovery) also follows a Poisson process, with the intensity λ_i^{extg} (extinguishing the fire). Finally, transition from state -1 to 0 is caused by a Poisson process with intensity λ_i^{rec} (completing the recovery process).

For all the events discussed above the exact time of occurrence is not known, hence we use a discrete time step to accommodate this uncertainty and round up the event time to the nearest integer step. Hence, the evolution of the system can be viewed as a discrete-time series of stochastic processes with three states. For convenience, we assume here that each time unit represents one day of the real world. Consequently, all fires starting on the same day are considered to be starting simultaneously. As shown in ref.¹ for real time t measured in finite time units (days here) for events generated by a Poisson process, an event happening in at most $\lceil t \rceil$ time units is identical for the corresponding Bernoulli process. The states of all houses at time t can be represented by a state vector $\vec{S}(t)$.

Here, like in ref.¹, we assume that experts provide assessments of each house's fire resistance. The value of $N_{1,i}$ represents the likelihood of house i to catch fire internally or externally, and $N_{2,i}$ represents the likelihood that the house fire is extinguished and house is rebuilt over large period of time. In the following, we define control parameters for state transitions with details listed in Table 2. The internal risk materialization is controlled by parameter α . The external risk materialization is controlled by parameter β . However, we assume that only the state transition from fully operational to on-fire state is observable, without knowing the actual reason for it. So, impact of an individual parameter α or β is hidden from direct observation. In contrast, parameters γ , controlling recovery and δ controlling fire extinguishing are independent of each other, so impact of each of the two can be observed in the changes in the evolution of corresponding underlying process.

We list events and parameters in the order compatible to the way they were ordered in the World Economic Forum model¹ in which only three events (*int*, *ext* and *rec*) and parameters (α , β , γ) exist, as defined in Table 2. We set $N_{1,i} = 0.4$ for large, 0.3 for medium and 0.2 for small houses and $N_{2,i} = 0.2$ for large, 0.3 for medium houses and 0.4 for small houses. We assign target values to our parameters $\alpha = 0.08$, $\beta = 0.012$, $\gamma = 0.016$ and $\delta = 0.032$.

To summarize, the dynamics progress in discrete steps $t = 1, \dots, T$ and the probability of transition in each step is defined by the intensity of the corresponding Poisson process as shown in Table 2. The dynamics described above and shown in Fig. 1(b) imply that the state of the system at time t depends only on its state at time $t - 1$, and therefore the evolution of the state vector $\vec{S}(t)$ is Markovian. The definitions of parameters and equations mapping them into intensities of Poisson processes are listed in Table 2.

With the assumed expert assessments and the created ground truth parameter values for the model, we simulate the evolution of house states for a particular period and record the frequency of emerging fires. In Fig. 1(d), we show the fraction of time each house is on-fire during one million time steps. The range of this on-fire fraction varies from 0.12 to 0.32. The fraction increases when the size or degree of the house increases, but areas of higher density housing suffer higher on-fire fractions than indicated by their degree.

Precision Limit of Maximum Likelihood Estimation (MLE). In our approach, we use Maximum Likelihood Estimation (MLE) to recover parameters from ground truth data. The main reason for this choice is that state transitions are governed by independent inhomogeneous probability distributions for which MLE delivers consistency and asymptotic normality with sufficient amounts of observed data²⁶. The historical data represents the combined effects of four Bernoulli processes. Our purpose is to recover the unknown parameters α , β , γ and δ mapping the expert assessments into event probabilities for Bernoulli processes of our model. We denote $\hat{\alpha}$, $\hat{\beta}$, $\hat{\gamma}$ and $\hat{\delta}$ to be the recovered (estimated) values of each parameter.

The use of MLE to find the values of hidden parameters from observed events has not been studied, yet ref.²³ indicates that it is feasible. In our approach we split one of the parameters of directly observable events into a pair of hidden (and tangled) parameters of two processes and recover these two parameters from observed events.

We denote the unknown parameters as θ . Given the n observations $x_1, x_2, \dots, x_n$, the likelihood function of this set of observations is defined as:

$$\mathcal{L}(\theta) = f(x_1, x_2, \dots, x_n | \theta) \quad (1)$$

When the distribution is discrete, f is a frequency function, defined explicitly by Eq. 7, and the likelihood function $L(\theta)$ shows the probability of observing the given data. The Maximum Likelihood method⁵ finds the values of parameters that yield the maximal probability of observing the given data. Logarithms are monotonic and therefore the likelihood and its logarithm have maximum at the same argument. Since the observed data comes from independent distributions, the logarithm of likelihood function can be written as:

$$\ln \mathcal{L}(\theta) = \sum_{i=1}^n \ln(f(x_i | \theta)) \quad (2)$$

For continuous and smooth likelihood functions, which is the case here, we can scan the parameter space in order to find the maximum point for $\ln \mathcal{L}(\theta)$. The historical data size is limited, so we need to study how this limitation affects the precision of results. In this section, we derive an estimation of the optimal parameters for large sample sizes. Under appropriate smoothness conditions, the estimate is consistent with large data sets and obeys the asymptotic normality. Detailed description of these conditions can be found in ref.¹⁹. These conditions can be summarized as: the first three derivatives of the $\frac{\partial \log(f(x_i | \theta))}{\partial \theta}$ are continuous and finite for all values of x_i and θ ; the expectation of the first two derivatives of $\frac{\partial \log(f(x_i | \theta))}{\partial \theta}$ can be obtained and the following integral is finite and positive:

$$\int_{-\infty}^{\infty} \left(\frac{\partial \log f}{\partial \theta} \right)^2 f dx \quad (3)$$

Let θ_0 denote the true value of the parameter θ and $\hat{\theta}$ to be its recovered value. Once the smoothness conditions are met, the recovered value $\hat{\theta}$ converges to the true value θ_0 as $n \rightarrow \infty$. If we normalize $\hat{\theta}$, we obtain an approximation from a normal distribution:

$$(\exists \sigma_{MLE}^2 > 0): \lim_{n \rightarrow \infty} \sqrt{n}(\hat{\theta} - \theta_0) = N(0, \sigma_{MLE}^2) \quad (4)$$

Given the definition of Fisher information²⁴ ($I(\theta)$):

$$I(\theta) = E \left[\frac{\partial}{\partial \theta} \log f(x | \theta) \right]^2 \quad (5)$$

The asymptotic normality of MLE can be written as:

$$\lim_{n \rightarrow \infty} \sqrt{n}(\hat{\theta} - \theta_0) = N \left(0, \frac{1}{I(\theta_0)} \right) \quad (6)$$

The variance of the normalized estimate decreases as $I(\theta_0)$ increases. This asymptotic variance to some extent measures the quality of MLE. Although it is hard to compute the variance analytically in our model, we know

there exists an estimation limit and the performance of MLE becomes better as the volume of given data increases. In the next section, we demonstrate that the variance indeed decreases as the size of the training dataset increases and the mean values of $\hat{\theta}$ approaches θ_0 with the error approaching 0.

Parameter Estimation in Fire-propagation Model. Given the historical data about each house transition in a finite period, we use Maximum Likelihood⁵ to find values of model parameters that make the model optimally match historical data. State transitions are independent functions with unknown parameter values. Since the historical data is generated from the discrete stochastic process, the likelihood of observing a particular sequence of risk materializations can be written as:

$$\mathcal{L}(\vec{S}(1), \vec{S}(2) \dots, \vec{S}(T)) \equiv \prod_{t=2}^T \prod_{i=1}^N P_i(t)^{S_i(t-1) \rightarrow S_i(t)} \quad (7)$$

where T is the number of time steps, N is the number of houses in the model, $S_i(t)$ is the current state of house i at time t , and $P_i(t)^{S_i(t-1) \rightarrow S_i(t)}$ is the state transition probability of house i from time $t-1$ to time t . $\vec{S}(t)$ is the vector of states for each house in the network at time t . The logarithm of this likelihood is:

$$\ln \mathcal{L}(\vec{S}(1), \vec{S}(2) \dots, \vec{S}(T)) \equiv \sum_{t=2}^T \sum_{i=1}^N \ln(P_i(t)^{S_i(t-1) \rightarrow S_i(t)}) \quad (8)$$

In the training process, we compute the probability of state transitions using p_i^{int} and p_{ji}^{ext} for transition into fire, p_i^{etg} for transition into recovery and p_i^{rec} for transition into the fully operational state. Correspondingly, there are three cases of state transitions in the historical data.

A transition from the fully operational state to the on-fire state ($0 \rightarrow 1$) happens when a house catches fire due to internal or external reasons. The probability of internal ignition is p_i^{int} . For external influence, we compute first the probability that none of the neighbors ignited this house, and then take the complement of this value. Neighbor j fails to ignite house i with probability $1 - p_{ji}^{ext}$. Taking the product over all on-fire neighbors gives us the probability that there are no externally triggered fires:

$$prod_i^{0 \rightarrow 0} = \prod_{j \in N_i} (1 - p_{ji}^{ext}) \quad (9)$$

where N_i is the set of all on-fire neighbors of house i . The complement of this product defines the probability that at least one neighbor ignites house i . Adding the probability of such external ignition to the probability of internal ignition, we obtain the total probability of a house catching fire:

$$P_i^{0 \rightarrow 1} = p_i^{int} + (1 - p_i^{int}) \left(1 - \prod_{j \in N_i} (1 - p_{ji}^{ext}) \right) \quad (10)$$

Since internal and external ignition are mutually exclusive we include a factor of $(1 - p_i^{int})$ in the external ignition probability. Accordingly, the probability of not catching on fire is:

$$P_i^{0 \rightarrow 0} = (1 - p_i^{int}) \prod_{j \in N_i} (1 - p_{ji}^{ext}) \quad (11)$$

A transition from being on fire to recovery state happens when the fire is extinguished and rebuilding process starts. This probability is defined as:

$$P_i^{1 \rightarrow -1} = p_i^{etg} \quad (12)$$

Transition from recovering to fully operational state ($-1 \rightarrow 0$) happens when a house is completely rebuilt and becomes fully operational. The corresponding probability is:

$$P_i^{-1 \rightarrow 0} = p_i^{rec} \quad (13)$$

The maximum likelihood parameters are obtained by summing the logarithms of corresponding probabilities. After scanning the potential ranges of the model parameters, we find the globally optimal values that maximize the likelihood of the historical data. The closeness of the recovered parameters to their values set in the simulations measures how precisely our model captures the dynamics of the system.

References

1. Szymanski, B. K., Lin, X., Asztalos, A. & Sreenivasan, S. Failure dynamics of the global risk network. *Scientific Reports* **5**, 10998 (2015).
2. Watts, D. J. A simple model of global cascades on random networks. *P. Natl. Acad. Sci. USA* **99**, 5766–5771 (2002).
3. Motter, A. E. & Lai, Y. C. Cascade-based attacks on complex networks. *Phys. Rev. E* **66**, 065102 (2002).
4. World Economic Forum Global Risks Report. Available at: <http://reports.weforum.org/global-risks-2015/> (Accessed: 15th January 2015).
5. Dempster, A. P., Laird, N. M. & Rubin, D. B. Maximum likelihood from incomplete data via the em algorithm. *J. R. Stat. Soc. B* **39**, NO.1 1–38 (1977).
6. Banks, E. Catastrophic risk: analysis and management. *John Wiley & Sons* (2005).

7. Taleb, N. N. The black swan: The impact of the highly improbable. *Random house*, (2007).
8. Pate-Cornell, E. On “Black Swans” and “Perfect Storms”: risk analysis and management when statistics are not enough. *Risk Analysis* **32**, 1823–1833 (2012).
9. National Research Council and others Review of the Department of Homeland Security’s approach to risk analysis. *National Academies Press* (2010).
10. Buldyrev, S. V., Parshani, R., Paul, G., Stanley, H. E. & Havlin, S. Catastrophic cascade of failures in interdependent networks. *Nature* **464**, 1025–1028 (2010).
11. Linkov, I. & Florin, M.-V. IRGC Resource Guide on Resilience. *EPFL International Risk Governance Center* (2016).
12. Ganin, A. A. *et al.* Operational resilience: concepts, design and analysis. *Scientific Reports* **6**, 19450 (2016).
13. Reason, J. Human Error. *Cambridge University Press* (1900).
14. Frosch, R. A. Notes toward a theory of the management of vulnerability. In Auerswald, P. E., Branscomb, L. N., La Porte, T. M. & Michel-Kerjan, E. O. (Editors), Seeds of disaster, roots of response: how private action can reduce public vulnerability. *Cambridge University Press*, 77–98 (2006).
15. Cinlar, E. Markov renewal processes. *Advances in Applied Probability* **1**, 123–197 (1969).
16. Cox, D. R. & Miller, H. D. The Theory Of Stochastic Processes *Methuen* (1965).
17. Dickey, J. M. The Renewal Function for an Alternating Renewal Process, Which Has A Weibull Failure Distribution and a Constant Repair Time. *Reliability Engineering and System Safety* **31**, 321–343 (1991).
18. van der Weide, J. A. M. & Pandey, M. A stochastic alternating renewal process model for unavailability analysis of standby safety equipment. *Reliability Engineering and System Safety* **139**, 97–104 (2015).
19. Cramér, H. Mathematical methods of statistics. *Princeton University Press* 500–506 (1946).
20. World Economic Forum Global Risks Report. Available at: <http://www.weforum.org/reports/global-risks-2013-eighth-edition> (Accessed: 7th January 2013).
21. Helbing, D. Globally networked risks and how to respond. *Nature* **497**, 51–59 (2013).
22. Perrow, C. Normal accidents: Living with high risk technologies *Princeton University Press* (1999).
23. DeGroot, M., Schervish, M. Probability and Statistics. *Addison-Wesley Press*, (2010).
24. Fisher, R. A. On an absolute criterion for fitting frequency curves. *Messenger of Mathematics* **41**, 155–160 (1912).
25. Linders, D. & Schoutens, W. A framework for robust measurement of implied correlation. *Journal of Computational and Applied Mathematics* **271**, 39–52 (2014).
26. Bertsekas, D., Tsitsiklis, J. Introduction to Probability, Second Edition. *Athena Scientific*, (2008).

Acknowledgements

This work was partially supported by DTRA Award No.HDTRA1-09-1-0049, by the Army Research Laboratory under Cooperative Agreement Number W911NF-09-2-0053 (the ARL Network Science Collaborative Technology Alliance). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies either expressed or implied of the U.S. Army Research Laboratory or the U.S. Government.

Author Contributions

B.K.S. conceived the research; X.L., A.M., G.K., B.K.S., and J.B. designed the research; X.L., A.M., B.K.S. implemented and performed numerical experiments and simulations; X.L., A.M., G.K., B.K.S., and J.B. analyzed data and discussed results; X.L., A.M., G.K., B.K.S., and J.B. wrote and reviewed the manuscript.

Additional Information

Competing Interests: The authors declare that they have no competing interests.

Change History: A correction to this article has been published and is linked from the HTML version of this paper. The error has not been fixed in the paper.

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2018