Hidden Markov models for estimating the restoration of interdependent infrastructure systems

M. Monsalve (1), J. C. de la Llera (2)

(1) PhD in Computer Science, mauricio.monsalve@cigiden.cl
(2) PhD in Civil Engineering, jcllera@ing.puc.cl

Abstract

If any infrastructure system is subjected to a shock, say, defined by the severe ground motions of an earthquake, it first experiences a sudden loss in functionality due to damage in its components, or due to other cascading failures. This sudden loss is followed by a recovery process that usually extends only until the service is fully restored. This curve of functionality is typically used by engineers to characterize the resilience of the affected system and different metrics exist for this purpose. In fact, the system is deemed resilient if it is functional most of the time, in other words, if it suffers little loss of functionality and recovers promptly. Therefore, to better understand the resilience of the system, it is relevant to characterize how infrastructures depend on each other to sustain their combined recovery. This work explores the use of hidden Markov models for modeling the restoration of functionality of infrastructure systems subjected to severe ground motions. Hidden Markov models are graphical models that can learn and reproduce the observed dynamics of stochastic processes and have been successfully applied in areas as diverse as finance, genomics, and speech recognition. Herein, different settings for training hidden Markov models are evaluated using real data sets that describe the restoration of functionality of a large variety of infrastructure systems from around the world, which have recently been affected by major earthquakes.

Keywords*: interdependent systems, resilience, hidden Markov models*

1 Introduction

Cities consist of interconnected networked systems, such as utilities, transportation, and healthcare, whose components depend on each other to operate, provide value, and sustain a functional society. But at the same time, these interdependencies increase the overall fragility of the combined systems by allowing disruptions to propagate through them [1][2][3]. Additionally, these interdependencies condition how these systems recover [4][5]. Thus, interdependencies play a role in both system fragility (or robustness) and recovery, i.e. in resilience.

The post-event curve of functionality, from affectation (loss of functionality) to full restoration, is typically used by engineers to measure system's resilience [6][7]. In fact, a system is deemed resilient if it suffers little loss of functionality and recovers promptly. Therefore, to better understand the resilience of the system, it is relevant to characterize how infrastructures depend on each other to sustain their combined recovery. This knowledge may help to improve seismic risk analyses, mitigate risk, and prioritize infrastructure recovery efforts.

This article explores the use of hidden Markov models (HMMs) for simulating the restoration of functionality of infrastructure systems subjected to severe ground motions. HMMs are graphical models that can learn and reproduce the observed dynamics of stochastic processes, and consist of a hidden (or latent) Markov chain and emission probabilities that associate the hidden states with observable values [8]. HMMs have been successfully applied in areas as diverse as finance [9], genomics [10], and speech recognition [8]. In seismic engineering, applications of HMMs include seismic catalog declustering [11], seismic hazard assessment [12], and even structural health monitoring [13].

2 Resilience and system restoration

2.1 *Mathematical characterization of resilience*

While there are several alternative definitions for resilience [14], the definitions usually point to the ability of a system to resist, avoid, absorb, recover, and adapt from shocks or stresses. In the case of seismic engineering, resilience is characterized by the loss of functionality a system experiences by a shock [6][7]. As Fig.1 illustrates, a disruptive event may cause a momentary loss of functionality that the system gradually resolves, until becoming (hopefully) completely functional again.

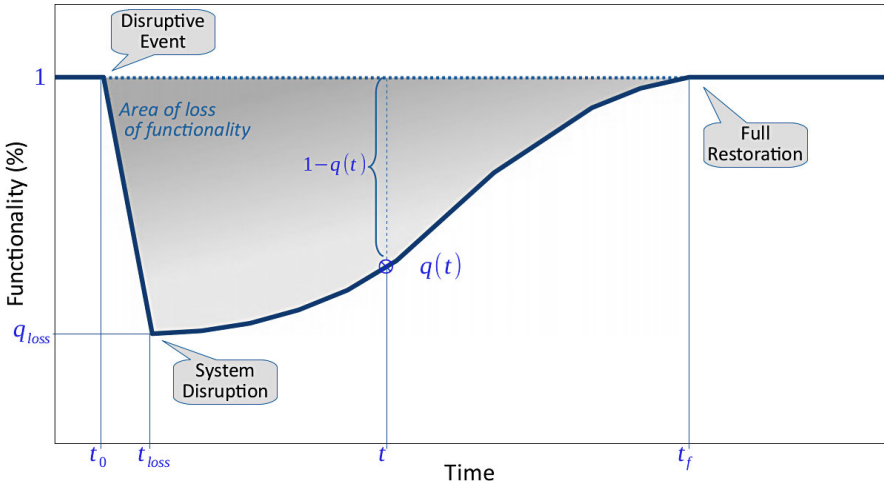

Fig. 1 – Visual schematics of the loss of functionality experienced by a system

By providing a metric for functionality, engineers are able to provide expressions for mathematically characterizing the resilience of a system after a disruptive event. These are often based on the following terms (Fig. 1):

- $q(t)$: functionality of the system at time t. It should hold that $0 \leq q(t) \leq 1, \forall\, t$.
- t_0: moment in which the disruptive event occurs.
- q_{loss}: minimum functionality reached by the system.
- t_{loss}: moment in which the functionality is minimal; note that $q_{loss} = q(t_{loss})$.
- t_f: moment in which the functionality is fully restored in the system; note that $q(t_f) = 1$.

These terms allow characterizing resilience using metrics such as

$$R = \int_{t_0}^{t_f} q(t)\,dt, \tag{1}$$

$$R_{loss} = \int_{t_0}^{t_f} (1 - q(t))\,dt, \tag{2}$$

$$R_{norm} = \frac{1}{t_f - t_0} R, \tag{3}$$

which are often called *resilience* (R, Eq. (1)), *loss of resilience* (R_{loss}, Eq. (2)), and *normalized resilience* (R_{norm}, Eq. (3)); see Chapter 3 of [6] and [7] for these and additional measures.

2.2 *Functionality restoration curves*

Following previous work [4][5], functionality restoration data after six major earthquakes were collected from the literature. The collected datasets contain time series of functionality that were compiled for three or more infrastructures recorded over the same time period, covering the same spatial areas, and whose recovery lasted more than three days. The datasets correspond to the following events:

1. 1996 Hanshin, Japan. This magnitude Mw 6.9 earthquake occurred on January 17, 1996 and was caused by an inland crustal rupture. The published curves describe the functionality of the power, gas, and water infrastructures for a period of 82 days after the event [15].

2. 2004 Mid-Niigata, Japan. This magnitude Mw 6.6 earthquake occurred on October 23, 2004. The curves describe the functionality of the power, gas, and water infrastructures for a period of 46 days after the event [16].

3. 2010 Maule, Chile. This magnitude Mw 8.8 earthquake occurred on February 27, 2010. The curves describe the functionality of the power distribution and both fixed and mobile telephony for the Maule and Bio Bio regions of Chile [17]. Both regions are treated as independent datasets herein.

4. 2011 Christchurch, New Zealand. This magnitude Mw 6.2 earthquake occurred on February 22, 2011. The curves describe the functionality of the power, gas, water, telecommunications and healthcare infrastructures for 29 days [18].

5. 2011 Tohoku, Japan. This devastating Mw 9.1 earthquake occurred on March 11, 2011 and generated a large tsunami. The curves describe the functionality of the power, gas, and water

infrastructures for 78 days [15][19]. Because of an aftershock that caused a single-day blackout, the functionalities of day 29 were replaced with the average of days 28 and 30.

6. 2016 Kumamoto, Japan. This magnitude Mw 7.0 earthquake occurred on April 16, 2016. The curves describe the functionality of the power, gas, and water infrastructures for 31 days [20].

Fig. 2 shows the collected curves. Observe that these curves are consistent with the conceptual drawing of Fig. 1.

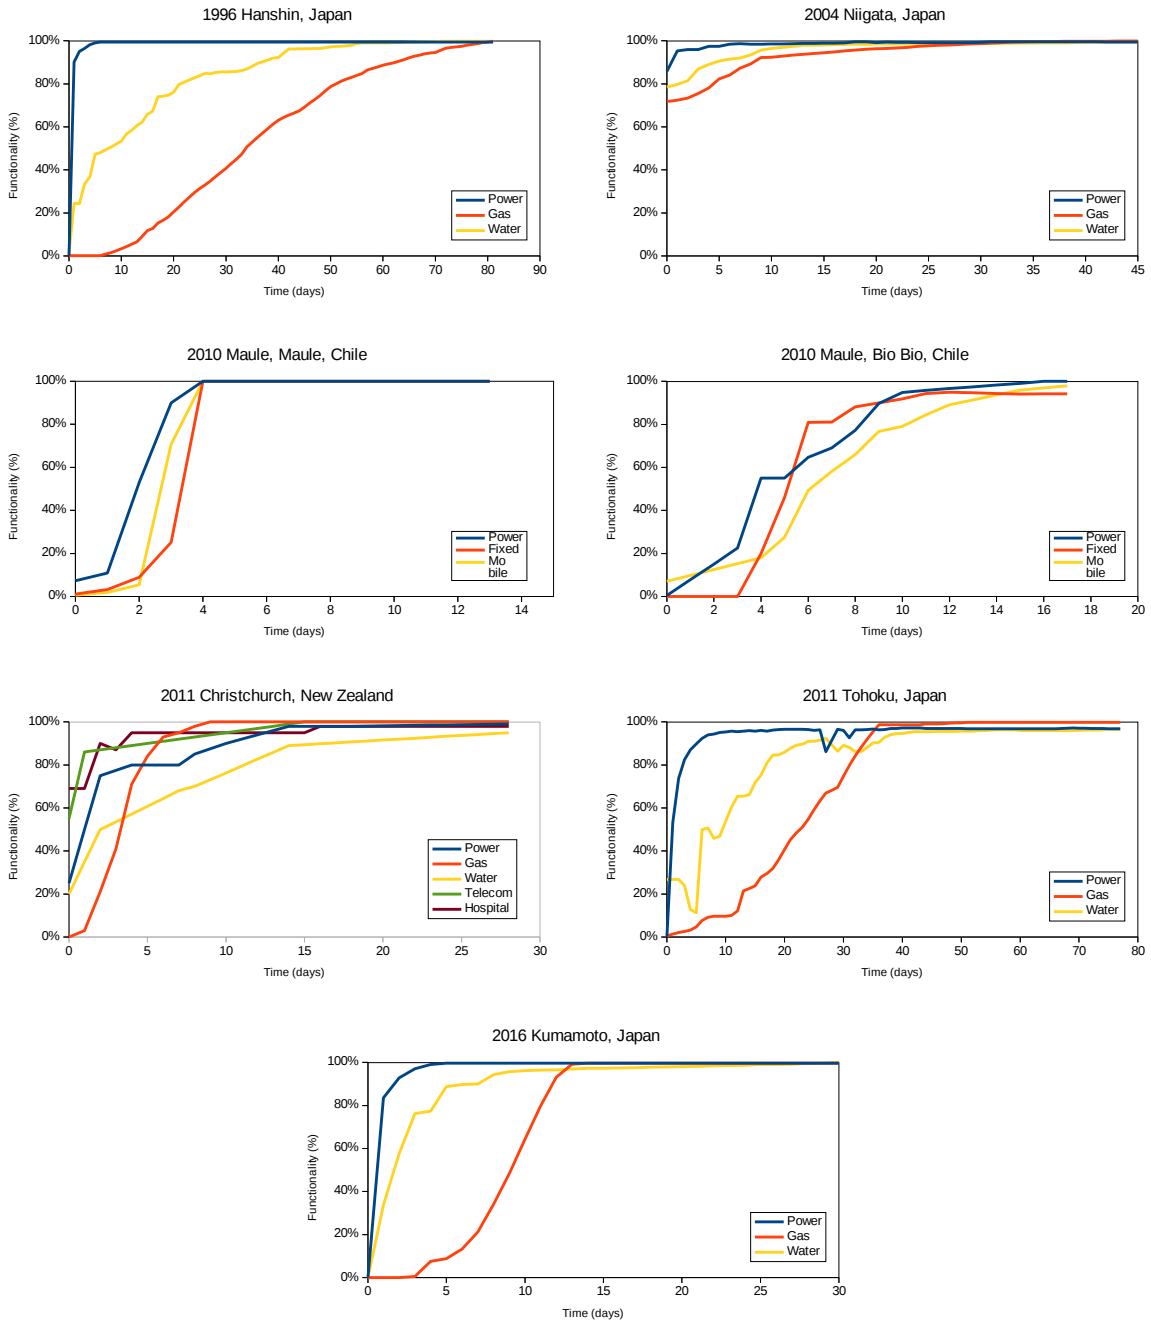

Fig. 2 – Infrastructure restoration curves compiled from the literature

3 Hidden Markov model

3.1 *Standard hidden Markov models*

HMMs are statistical abstractions that can learn and reproduce the observed dynamics of stochastic processes. Formally, a standard HMM consists of the following [8]:

- A Markov chain described as a transition matrix (M) of size $h \times h$ (the system contains h states).
- Initial state probabilities ($\pi_1, \pi_2, \ldots, \pi_h$) which indicate the starting state of the Markov chain.
- Emission probabilities ($Pr_1, Pr_2, \ldots, Pr_h$) which generate the symbols according to the current state the machine is in.

Within this structure, the Markov chain is essentially a latent or hidden variable (hence the name HMM). It is the emission probabilities that translate the hidden states into observable symbols. This gives the HMM the flexibility needed to tackle different problems: the emission probabilities may model distributions of numbers, vectors, sets, words, etc.

To illustrate how an HMM generates a sequence, consider the example shown in Fig. 3. The first state is sampled according to the start probabilities, which in this case can be either state 1 (prob. 0.6) or state 2 (prob. 0.4). Assume state 2 is chosen. Now, the first observable value can be generated by using the emission probability of state 1, which is a normal of mean 2 and standard distribution 0.1. Assume value 1.9 is sampled. After this, a transition must occur; since state 2 only leads to state 3, the latter becomes the next state visited. In state 3, the emission distribution is a normal distribution of mean 3 and standard deviation 0.1; assume 2.9 is sampled as observed value. For the next transition, state 3 may either go to state 1 or state 2; assume state 1 becomes the next state visited. And by repeating this process, the first sampled sequence of Fig. 3 is created.

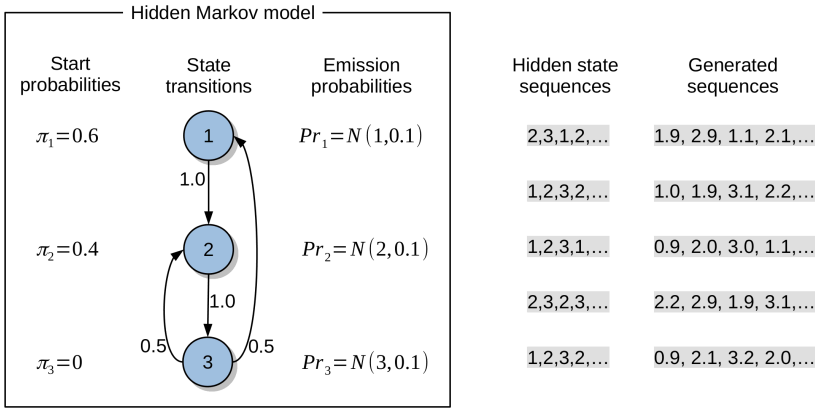

Fig. 3 – Example tiny hidden Markov model and sequences generated

Tasks associated with HMMs include sampling random observed sequences (the aim of this work), finding the probability an observed sequence is generated by the HMM, finding the most likely sequence of hidden states that produces an observed sequence (the Viterbi algorithm), and training an HMM (i.e. statistically fitting it). Training HMMs can be done in a number of ways, which include the Baum-Welch algorithm [21], Viterbi training [21], spectral inference [22], structure-inducing algorithms [23], etc. The wide variety of training algorithms stems from the fact that optimally training an HMM is a computationally difficult task.

3.2 *Custom hidden Markov model*

The standard HMM is structurally inadequate to reproduce the functionality restoration curves considered in this work. As the collected curves show, functionally is gradually restored in the system.

Therefore, a custom HMM is defined in this work, based on the autoregressive HMM (AR-HMM) [24]. In the AR-HMM, emission probabilities are replaced by conditional probabilities or, directly, transformations. These take the last observed symbol(s) and sample the next one.

In this work, states are instead assigned to two emission distributions: one for marginal distributions, i.e. $Pr_i(Y_1)$, and another for conditional distributions, i.e. $Pr_i(Y_t \mid Y_{t-1})$. The marginal distribution is used for sampling the functionality of the system right after the event, i.e. Y_1. Then, to sample Y_2, Y_3, and so on, the conditional distribution is used instead. In other words, the first symbol of the sequence is sampled like a standard HMM, while the subsequent symbols are sampled following an AR-HMM scheme. This contrast between a standard HMM and this work's implementation is shown in Fig.4.

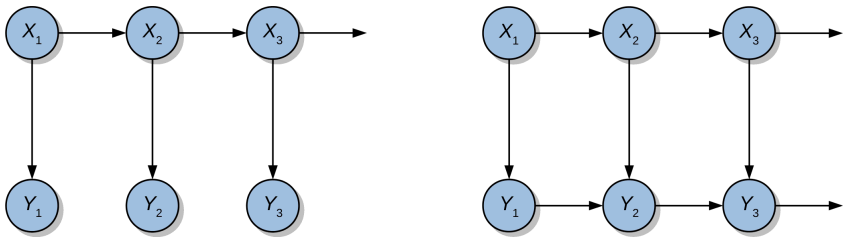

Fig. 4 – Bayesian network depiction of a standard HMM (left) and this work's HMM (right)

The emission distributions are multivariate normal distributions with diagonal covariance matrices. This also applies to the conditional distributions, where $Pr_i(Y_t \mid Y_{t-1})$ essentially model $Y_t - Y_{t-1}$ as multivariate normal. This might seem oversimplistic, but the complexity of the model is left to the structure of the HMM itself. In fact, if there are h hidden states, then there are h start probabilities, h^2 transitions, and $2h$ emission distributions in this HMM.

The custom Viterbi algorithm and Viterbi training used in this work are described in the appendix. The initial topology of the HMMs created herein is that of a complete clique, i.e. all states are reachable from each other, with equal probability. All models were built using 10 hidden states.

4 Experiments

Successful training of an HMM requires a large collection of sequences of the same phenomenon. In the case considered herein, only a single sequence is available for each system and event. This is clearly insufficient for crafting an HMM that can simulate each system's curves of functionality after an event. To address this situation, the collected curves are resampled to artificially augment their number and introduce some diversity to the training.

Let sequence $<Y_1, Y_2, \ldots, Y_n>$ describe the functionality of a system during restoration, where each $Y_t = (y_t^{(1)}, y_t^{(2)}, \ldots)$ describes the functionality of the different infrastructures at time t. The following sources of uncertainty can be used to resample the original curves, producing different ones for training:

1. Numerical precision uncertainty: functionality values are often rough estimates, a random term cand be added to values Y_t.
2. Time shift uncertainty: since functionality estimates at day t could have been ready the day before or could have been obtained at any time during day t. Thus, the value of Y_t can be replaced by that of Y_{t+1} to generate a new sequence.

3. Random starting time: if sequence $<Y_1, Y_2, \ldots, Y_n>$ describes the restoration of a system after an event, then subsequence $<Y_t, Y_{t+1}, \ldots, Y_n>$ is also a valid restoration sequence, with $1<t<n$.

In practice, the first kind of uncertainty mentioned, numerical precision, can be introduced directly into the emission probabilities of the HMM.

With the aim of evaluating if the HMMs can reproduce believable restoration curves, i.e. consistent with the data, two HMMs were crafted for each event and location: one trained with time shift uncertainty resampling, and the other trained with random starting time resampling. Sample simulated restoration curves are shown in Figs. 5 and 6. For illustrative purposes, these curves have a length of 20 days.

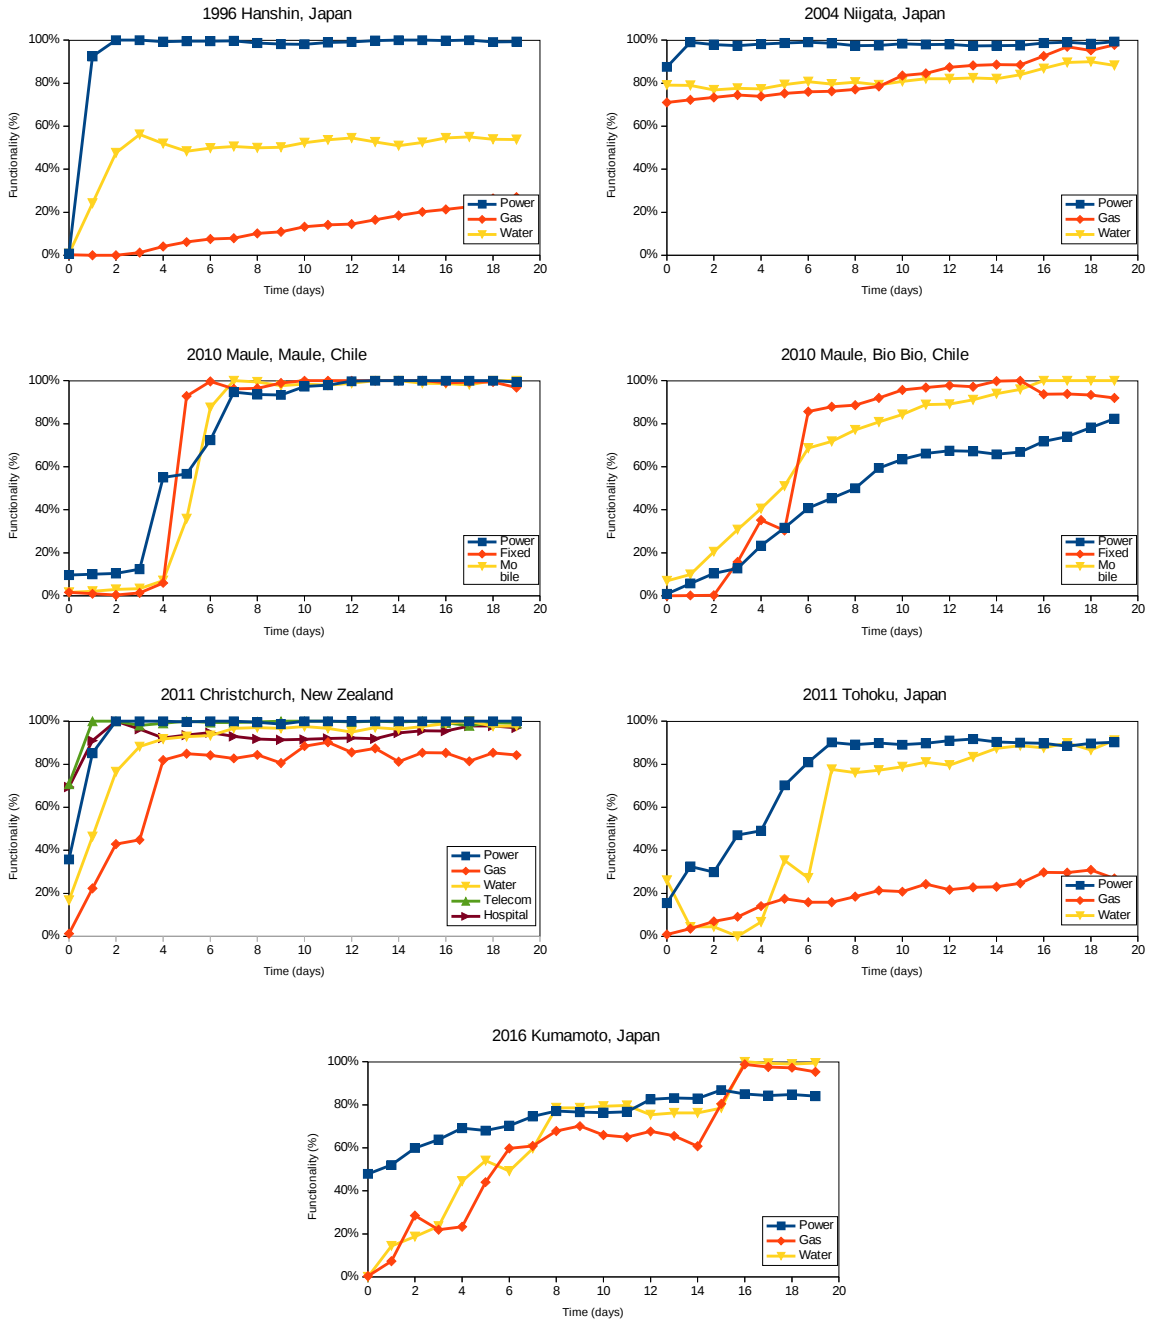

Fig. 5 – Simulated restoration curves from HMMs trained using time shift resampling.

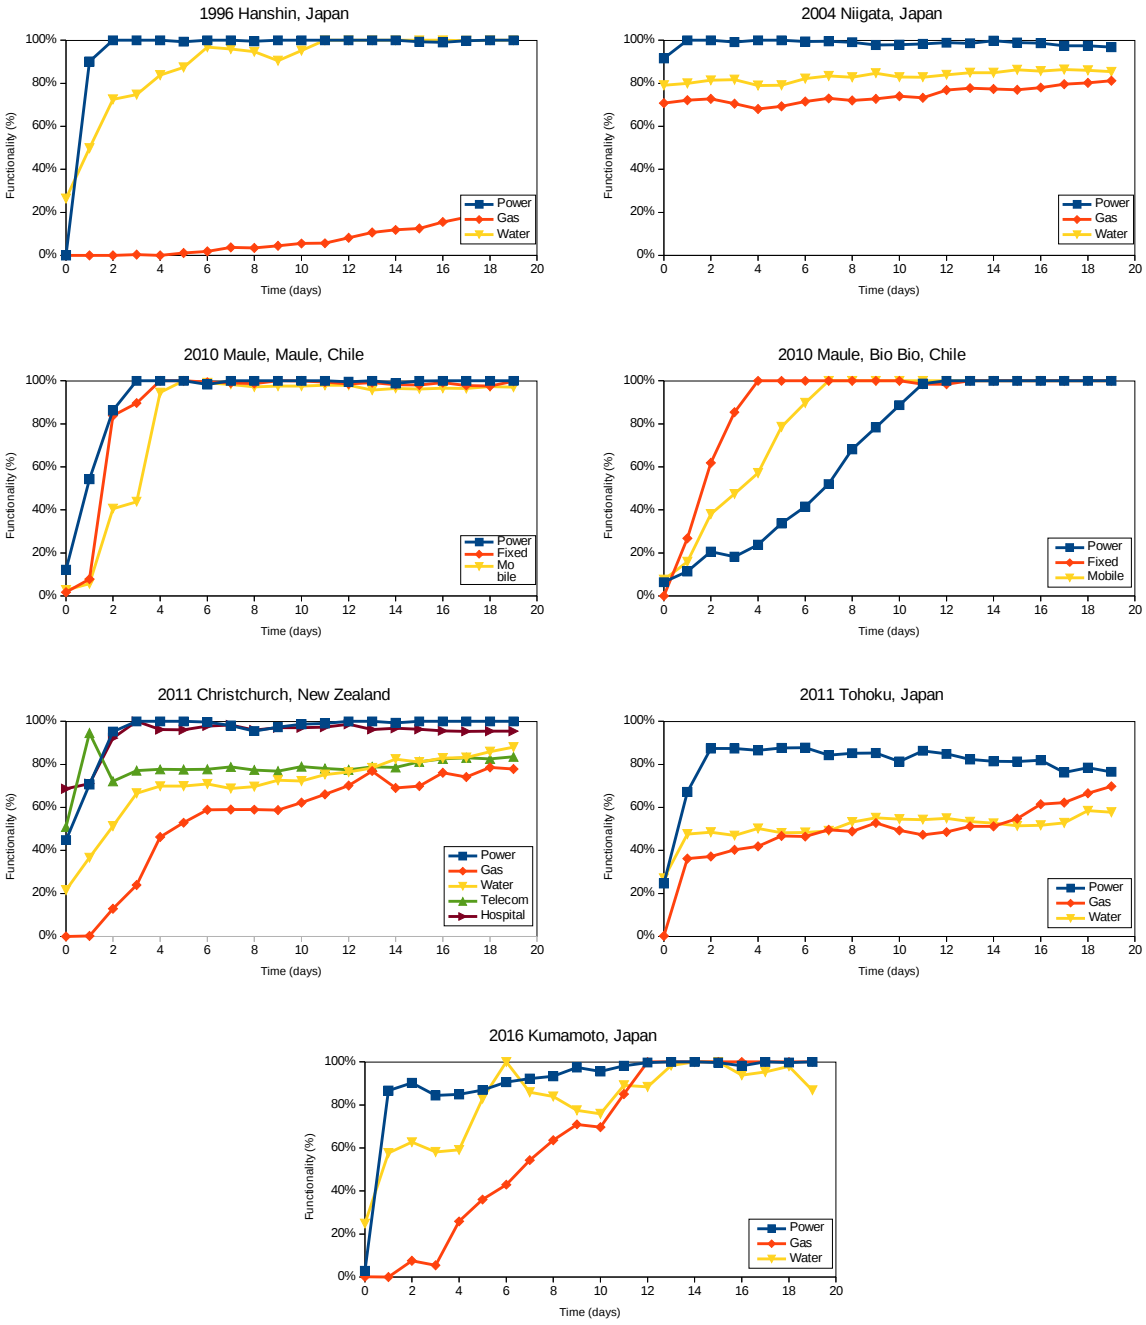

Fig. 6 – Simulated restoration curves from HMMs trained using random start resampling.

Both types of HMMs are able to more or less reproduce believable restoration curves, in that they tend to resemble the original curves, yet are still random. However, some of the curves generated did not resemble the original curves, e.g. they would stall at fragilities around 60% as if the system was fully restored. This is a consequence of the memoryless property of the Markov chain, and shows that further work is needed to improve these graphical models to obtain better performance.

Both types of training cause similar effects on the initial values generated. In general, the initial values are greater than the ones in the data, which occurs because both time shift and random start resampling supply values that may be greater than Y_0. In the case of time shift resampling, the first value supplied may be Y_1 instead of Y_0. In the case of random start resampling, Y_0 may replaced by any Y_t, which can be even greater than Y_1. Thus, the curves generated by the HMM trained on time shift

resamples tends to start on nearly the same fragilities as the original data, while the curves generated by the HMM trained on random start resamples have more varied starting fragilities.

The HMMs had issues with the 2010 Maule, Maule, Chile and 2011 Tohoku, Japan curves. In the former case, the generated curves tend to stall in the initial fragilities to then experience a sharp recovery and finally remain near 100% functionality. In the latter case, fragilities often times did not reach 100% functionality and, furthermore, experienced stages of growth and shrinkage. This was caused by the decreases that the original curves have.

5 Conclusions

Simulating the restoration of functionality of systems-of-systems permits the assessment of their resilience. Furthermore, if the simulations consider, explicitly or implicitly, their interdependencies, then it is possible to prioritize restoration of specific infrastructures to accelerate the recovery of the whole system. Following this motivation, this article presented ongoing work regarding the simulation of functionality restoration of infrastructure systems using HMMs trained on real infrastructure restoration data collected from the literature. The collected restoration curves were resampled exploiting the inherent uncertainty of the data to make it possible to train the HMMs and, subsequently, be able to simulate more rich restoration curves that still preserve characteristics of the original restoration processes.

The experiments showed that the resampling strategies can effectively be used to train HMMs that can simulate functionality restoration curves afterwards. There were, nonetheless, limitations found using this approach, both related to the HMMs themselves and the lack of data. This motivates further research on developing better customized HMMs, with more properly chosen initial configurations, and even considering training on synthetic data (restoration of systems defined in detail) together with real data.

6 Acknowledgments

This work was made possible thanks to the support received from CONICYT, Chile through: FONDECYT Postdoctorado project CONICYT/FONDECYT/3170867, FONDECYT Regular project *SIBER-RISK* CONICYT/FONDECYT/1170836, and the *National Research Center for Integrated Natural Disaster Management* (CIGIDEN) CONICYT/FONDAP/15110017.

A Appendix

A.1 *Viterbi estimation*

This appendix describes the Viterbi algorithm for the custom HMM implementation used in this work. Given an observed sequence $Y = <Y_1, Y_2, \ldots, Y_n>$, the Viterbi algorithm finds the most likely sequence of hidden states that generates it. The algorithm essentially is a recursion that constructs the sequence of states by continuously picking the subsequences with highest likelihood.

The probability a given observed sequence $<Y_1, Y_2, \ldots, Y_n>$ is generated by a sequence of states $<X_1, X_2, \ldots, X_n>$ is the following:

$$\pi_{X_1} Pr_{X_1}(Y_1) \Pi_{2 \le t \le n} m_{X_{t-1}, X_t} Pr_{X_t}(Y_t \mid Y_{t-1}). \quad (4)$$

The structure of this probability makes it amenable for optimization through recursion and dynamic programming, i.e., the Viterbi algorithm. Indeed, from some step $k > 1$, term

$$\Pi_{k \le t \le n} m_{X_{t-1}, X_t} Pr_{X_t}(Y_t \mid Y_{t-1}) \tag{5}$$

only depends on X_{k-1}. This implies that optimization of formula (X) only requires changing state variables $< X_k, X_{k+1}, \ldots, X_n >$ after having set X_{k-1}. In other words, it is an independent subproblem, if X_{k-1} is known. This is a consequence of the Markovian structure of the HMM.

Based on the above discussion, the probability of generating $< Y_1, Y_2, \ldots, Y_n >$ by the optimal sequence of states is

$$Q = \max_i [\pi_i Pr_i(Y_1) Q_{i,2}], \tag{6}$$

where $Q_{i,k}$ is defined as

$$Q_{i,k} = \max_j [m_{ij} Pr_j(Y_k \mid Y_{k-1}) \cdot Q_{j,k+1}] \quad (\forall\, 2 \le k \le n), \tag{7}$$

$$Q_{i,n+1} = 1. \tag{8}$$

The optimal sequence of states $X^* = < X_1^*, X_2^*, \ldots, X_n^* >$ is obtained by following the recursion in Eq. (6), Eq. (7), and Eq. (8).

A.2 *Viterbi training*

Traditionally, HMMs are trained using the EM algorithm. However, this algorithm is computationally expensive. The Viterbi training algorithm, instead, *translates* the set of observed sequences to a set of hidden state sequences and then uses simple summary statistics to update the structure of the HMM. This process is repeated a number of times (six) and then convergence is assumed. The following points describe how the elements of the HMM are updated using the statistics:

- Start probabilities: probability π_i is simply the proportion of times the first state of the optimal state sequences is *i*.
- Transition probabilities: probability m_{ij} is the proportion of times state *j* follows immediately after state *i*.
- State-emission samples: each state *i* is associated with two collections of observed symbols, one for inferring marginal probabilities S_i^{marg} and another for inferring conditional probabilities S_i^{cond}. Consider a sequence of observations $< Y_1, Y_2, \ldots, Y_n >$ associated with a sequence of states $< X_1, X_2, \ldots, X_n >$. Then, symbol Y_t is added to $S_{X_t}^{marg}$ and pair (Y_t, Y_{t-1}) is added to $S_{X_t}^{cond}$. After all the observations are processed (within the iteration), collections S_i^{marg} and S_i^{cond} are used to define (fit) emission probabilities $Pr_i(Y_1)$ and $Pr_i(Y_t \mid Y_{t-1})$.

References

1 Rinaldi S M, Peerenboom J P, Kelly T K. Identifying, understanding, and analyzing critical infrastructure interdependencies. *IEEE Control Systems* 2001; 21(6):11-25.

2 Buldyrev S V, Parshani R, Paul G, Stanley H E, Havlin S. Catastrophic cascade of failures in interdependent networks. *Nature* 2010; 464(7291):1025-1028.

3 Bashan A, Berezin Y, Buldyrev S V, Havlin S. The extreme vulnerability of interdependent spatially embedded networks. *Nature Physics* 2013; 9(10):667-672.

4 Monsalve M, de la Llera J C. Data-driven estimation of interdependencies and restoration of infrastructure systems. *Reliability Engineering & System Safety* 2019; 181: 167-180.

5 Monsalve M, de la Llera J C. New Models for Estimating Interdependencies and Recovery of Infrastructure Systems. *16th European Conference on Earthquake Engineering (16ECEE)*. Thessaloniki, Greece, June 18-21, 2018.

6 Tsionis G. Seismic Resilience: Concept, Metrics and Integration with Other Hazards. Joint Research Centre, Publications Office of the European Union, Luxembourg; 2014.

7 Cimellaro G P, Reinhorn A M, Bruneau M. Framework for analytical quantification of disaster resilience. *Engineering structures* 2010; 32(11):3639-3649.

8 Rabiner L R. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE* 1989; 77(2):257-286.

9 Nystrup P, Madsen H, Lindström E. Stylised facts of financial time series and hidden Markov models in continuous time. *Quantitative Finance* 2015; 15(9):1531-1541.

10 Yau C, Papaspiliopoulos O, Roberts G O, Holmes C. Bayesian non-parametric hidden Markov models with applications in genomics. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 2011; 73(1):37-57.

11 Wu Z. A hidden Markov model for earthquake declustering. *Journal of Geophysical Research: Solid Earth* 2010; 115(B3).

12 Pertsinidou CE, Tsaklidis G, Papadimitriou E, Limnios N. Application of hidden semi-Markov models for the seismic hazard assessment of the North and South Aegean Sea, Greece. *Journal of Applied Statistics* 2017; 44(6):1064-1085.

13 Mollineaux M, Rajagopal R. Structural health monitoring of progressive damage. *Earthquake Engineering & Structural Dynamics* 2015; 44(4):583-600.

14 Vugrin E D, Warren D E, Ehlen M A, Camphouse R C. A Framework for Assessing the Resilience of Infrastructure and Economic Systems. In: Gopalakrishnan K, Peeta S, editors. *Sustainable and Resilient Critical Infrastructure Systems*. Springer, Berlin, Heidelberg; 2010.

15 Nojima N. Restoration processes of utility lifelines in the great east Japan earthquake disaster, 2011. *15th World Conference on Earthquake Engineering, 15 WCEE*. Lisbon, Portugal; 2012.

16 Kajitani Y, Sagai S. Modelling the interdependencies of critical infrastructures during natural disasters: a case of supply, communication and transportation infrastructures. *International Journal of Critical Infrastructures* 2009; 5(1): 38-50.

17 Dueñas-Osorio L, Kwasinski A. Quantification of lifeline system interdependencies after the 27 February 2010 Mw 8.8 offshore Maule, Chile, earthquake. *Earthquake Spectra* 2012; 28(S1):S581–S603.

18 Zorn C R, Shamseldin A Y. Quantifying directional dependencies from infrastructure restoration data. *Earthquake Spectra* 2016; 32(3):1363-1381.

19 Cimellaro G P, Solari D, Bruneau M. Physical infrastructure interdependency and regional resilience index after the 2011 Tohoku earthquake in Japan. *Earthquake Engineering & Structural Dynamics* 2014; 43(12):1763–1784.

20 Nojima M, Maruyama Y. Comparison of functional damage and restoration processes of utility lifelines in the 2016 Kumamoto earthquake, Japan with two great earthquake disasters in 1995 and 2011. *JSCE Journal of Disaster FactSheets* 2016; FS2016-L-0005.

21 Rodríguez L J, Torres I. Comparative study of the Baum-Welch and Viterbi training algorithms applied to read and spontaneous speech recognition. *2003 Iberian Conference on Pattern Recognition and Image Analysis (IbPRIA 2003)*. Andratx, Spain: *Springer*; 2003. p. 847-857.

22 Hsu D, Kakade S M, Zhang T. A spectral algorithm for learning hidden Markov models. *Journal of Computer and System Sciences* 2012; 78(5):1460-1480.

23 Ait-Mohand K, Paquet T, Ragot N. Combining structure and parameter adaptation of HMMs for printed text recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 2014; 36(9):1716-1732.

24 Park DW, Kwon J, Lee KM. Robust visual tracking using autoregressive hidden Markov model. *2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2012)*. Providence, Rhode Island: *IEEE*; 2012. p. 1964-1971.