Research Article

Real-Time Prediction of the Trend of Ground Motion Intensity Based on Deep Learning

Tao Liu and Zhijun Dai

Institute of Geophysics, China Earthquake Administration, Beijing 100081, China

Correspondence should be addressed to Zhijun Dai; dzj@cea-igp.ac.cn

Received 9 February 2021; Revised 26 September 2021; Accepted 1 October 2021; Published 18 October 2021

Academic Editor: Luca Landi

Copyright © 2021 Tao Liu and Zhijun Dai. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

In order to predict the intensity of earthquake damage in advance and improve the effectiveness of earthquake emergency measures, this paper proposes a deep learning model for real-time prediction of the trend of ground motion intensity. The input sample is the real-time monitoring recordings of the current received ground motion acceleration. According to the different sampling frequencies, the neural network is constructed by several subnetworks, and the output of each subnetwork is combined into one. After the training and verification of the model, the results show that the model has an accuracy rate of 75% on the testing set, which is effective on real-time prediction of the ground motion intensity. Moreover, the correlation between the Arias intensity and structural damage is stronger than the correlation between peak acceleration and structural damage, so the model is useful for determining real-time response measures on earthquake disaster prevention and mitigation compared with the current more common antiseismic measures based on predictive PGA.

1. Introduction

Earthquakes are one of the natural disasters all over the world that can cause great harm to humans. With the current level of scientific research, earthquakes cannot be predicted very accurately. The main reason is that the occurrence of an earthquake is ground shaking caused by the release of a large amount of crustal energy. This is a very complicated process. Although there are certain rules to follow, there are a lot of variables involved in predicting earthquakes, many of which cannot be measured [1, 2]. Therefore, the current earthquake prediction is mainly based on the earthquake risk estimation, that is, the probability of the largest damage because of the earthquake in a certain city or building area within a certain number of years [3–5].

Although the methods such as the ETAS model have proven effective in predicting triggered earthquakes, they often fail because they often underestimate the number of future earthquakes, and the degree of underestimation is related to the time of the main shock. Joffe et al. pointed out that when the accuracy of predicting earthquakes is

insufficient to meet the demand, it is important to find a new method that can have a wider and larger source of information [2]. This almost coincides with the deep learning method.

Deep learning has attracted much attention due to its extensive success in various fields [6]. In the field of earthquake research, the core is to obtain information and knowledge by analyzing data. Deep learning technology has strong modeling, feature extraction, and data analysis capabilities and has been successfully applied to many challenging problems, such as earthquake identification and classification [7, 8], seismic phase picking [9, 10], and earthquake early warning [11, 12]. In terms of earthquake prediction, deep learning also has many applications [13], but almost all of them are long-term predictions [14, 15], short-term predictions [16], and few real-time predictions.

In the current earthquake prediction research, the output ground motion parameters mainly include peak acceleration and peak velocity, which can only reflect a specific ground motion feature and cannot fully effectively infer the damage caused by ground motion. Many scholars

have found that Arias intensity is a kind of ground motion energy parameter that includes ground motion amplitude, frequency spectrum, and duration characteristics. It has a strong correlation with disaster phenomena caused by earthquakes. It can better reflect some characteristics of ground motions and is more reliable and effective for the prediction of damage of the earthquakes [17–21].

At present, studies have shown that Arias intensity is used in the prediction of the degree of damage to short-period structures caused by ground motions [22], the probability of earthquake-induced landslides [23], the possibility of foundation failure due to earthquake-induced sand liquefaction [24], and other aspects very well. Therefore, this article uses deep learning technology to build a model and takes 30 s real-time monitoring of the horizontal acceleration time history recordings as input, which can predict the Arias intensity trend of the future 5 s acceleration time history recordings in real time.

The main contribution of our research is listed as following points:

- (1) When an earthquake occurs, the severity or development trend of the earthquake disaster can be determined in advance. The proposed model may cooperate with the earthquake early warning system to make the real-time monitoring and emergency measures of the building's seismic system be more effectively avoiding casualties and property losses.
- (2) Our model suggests that Arias intensity can be predicted in many situations.
- (3) This article is purely data-driven to build a model to explore the possibility of using current seismic acceleration recordings to predict the future.

2. Method

2.1. Fully Connected Neural Network. Commonly used neural network models include fully connected neural networks, convolutional neural networks, and recurrent neural networks. The fully connected neural network is shown in Figure 1. The neurons are laid out in layers. The leftmost layer is the input layer, and the rightmost layer is the output layer. The middle part is the hidden layer because they are invisible. All neurons in two adjacent layers are connected to each other, and each connection has a weight. Compared with the one-dimensional arrangement of neurons in each layer of a fully connected neural network, a convolutional neural network (as shown in Figure 2) is very different in structure. Each layer of neurons is arranged in a cube structure, and there is generally a pooling layer between convolutional layers to reduce the number of samples in each layer to reduce the number of parameters. The recurrent neural network (as shown in Figure 3) is somewhat similar to the full connection, but the difference is that the output of each neuron of the recurrent neural network not only depends on its input but also depends on the output of the previous neuron.

Because the recurrent neural network has problems such as gradient disappearance and training instability, the fully connected neural network is more conducive to the overall trend prediction because the neurons of each layer are connected to

FIGURE 1: Fully connected neural network.

each other, and the calculation time is shorter compared with other network structures. The shorter calculation time is more meaningful for implementing predictions, so the model created in this article is a fully connected neural network (FCNN).

The model is divided into two stages. The first stage is segmented input. This model is for real-time monitoring of the change trend of Arias intensity, that is, the change trend of acceleration. If a total of 3000 horizontal acceleration values of 30 s are used as input, the calculation amount is relatively large, and the calculation time cannot meet the requirements of rapid prediction. In addition, the input of the last 5 s is more important for the result and plays a decisive role, so we have divided different sampling frequencies according to the time position of each data point to reduce the calculation time without greatly affecting the results. Since the sampling frequency of the input data is different, the input is separated according to the sampling frequency, and according to the size of the input data, using different neural network layers and neuron numbers will have a better training effect. The second stage combines the output of the first stage to get the prediction result.

The first stage is divided into three parts. The first part is data input with a sampling frequency of 1 hz in the time range of 0~15 s, a fully connected layer, the number of inputs is 15, and the number of outputs is 5; the second part is data input with a sampling frequency of 10 hz within a time range of 15~25 s, two fully connected layers, the number of inputs is 100, and the number of outputs is 20; and the third part is data input with a sampling frequency of 100 hz within a time range of 25~30 s, three fully connected layers with the number of inputs 500, and the number of outputs is 40. The second stage combines a total of 65 outputs from the three parts of the first stage and forms a two-class model through three fully connected layers. Because the training speed of the ReLU activation function is much faster than that of the Sigmoid function and the Tanh function, the models constructed in this paper use the ReLU activation function [25]. When calculating the loss function, a Softmax layer is added to the last layer of the model. The model structure is shown in Figure 4.

FIGURE 2: Convolutional neural network.

FIGURE 3: Recurrent neural network.

2.2. Data Processing Method. Data are very critical for deep learning models. It affects the convergence speed of the model and the training effect. Therefore, it is important to preprocess the data before inputting the model to make it accurate, consistent, and applicable. The following are the data preprocessing steps (Figure 5):

- (1) *Acceleration Zoom.* The data used in this article are all downloaded from K-net and KiK-net. The acceleration of each seismic recording is not a true value but an amplified value. Since the Arias intensity needs to calculate the square of the acceleration, it will increase exponentially if it is calculated by the magnification value, and the magnification factor of each station may not be the same, so it is necessary to scale the ground motion acceleration recording according to the scale factor of each station. The unit of acceleration is gal.
- (2) *Baseline Correction.* Ground motion recordings are susceptible to low-frequency noise such as instrument noise and background noise, resulting in baseline error pollution, leading to serious drift and distortion in acceleration and the velocity and displacement waveforms obtained by integration. We can do a simple baseline correction, that is, subtract the average value from the scaled ground motion, which can greatly reduce the influence of low-frequency noise on the acceleration recording. In practical applications, the average value of the steady noise lasting 20 seconds before the earthquake can be subtracted to reduce the influence of low-frequency noise.
- (3) *Normalization.* A total of 30 seconds of data is selected, and the horizontal ground motion

acceleration is obtained according to the sampling frequencies of 1 Hz, 10 Hz, and 100 Hz in the first 15 seconds, the middle 10 seconds, and the last 5 seconds, respectively. We take the absolute average value when sampling from low frequency to high frequency. Taking 1 Hz as an example, if the ground motion recording is 100 Hz, that is, there are 100 acceleration in one second, then the absolute average value of 100 acceleration in one second is taken as one of the data in the 1 Hz sample. After obtaining a complete sample, normalization is performed; that is, the entire sample is divided by the maximum value.

2.3. Training and Testing Methods. The data are divided into a training set and a test set, which are used for model training and testing, respectively. The sample number ratio is 4:1, and the samples are not repeated between two sets. Each sample consists of an acceleration recording and a label. The acceleration recording is the 615 acceleration values with a duration of 30 seconds described above as the input to the model. The label is the trend of Arias intensity change in the future 5 seconds. If the trend increases, the label is 1, and the decrease is 0, as the actual comparison value output by the model.

After the output of the model passes through the Softmax layer, two probability values are obtained, which correspond to the probability that the model judges whether the intensity of Arias will decrease or increase in the future 5 seconds. The purpose of model training is to make the difference between the output and the label value close to 0. If it is actually increasing, then the probability of increase in the model output is close to 1, and the probability of decreasing is close to 0, and vice versa. This difference is called a loss function.

The loss function used in this paper is the sparse Softmax cross entropy, and the optimization function to reduce the cross entropy is the adaptive moment estimation optimizer (Adam), which is derived from the AdaGrad and RMSProp optimization functions. It has the following advantages: (1) it has simple implementation and efficient calculation; (2) hyperparameters hardly need to be adjusted, reducing the influence of factors to improve training efficiency; (3) the learning rate can be automatically adjusted to a certain extent; and (4) it is very suitable for large-scale data and parameter scenarios.

FIGURE 4: The model structure diagram.

FIGURE 5: Preprocessing flow chart.

2.4. Data Filtering. First of all, whether it is earthquake prediction or earthquake risk estimation, it is more realistic for larger earthquakes, so the earthquake recordings we downloaded are selected from the period 2003 to 2019, on K-net and KiK-net with a magnitude of 5.0 or higher, involving a total of 1667 stations, and divided into 5 to 5.9, 6 to 6.9, 7 to 7.9, and 8 earthquake intervals, and each of the training set and test set is guaranteed that the number of

each interval data is equal in order to get a better training effect, and the results can be compared and analyzed.

Then, when we sample the recording, each recording can be 5 s, 10 s, 15 s, and so on. Until the end of the recording as the sampling end point of each sample, take the past 30 s data according to different sampling frequencies. That is, the samples obtained in each recording include different stages, before, during, and after the earthquake. If 5 s or 10 s is not

FIGURE 6: Learning rate adjustment graph.

enough for the end point to take 30 s, the zero padding process is done in front to make the sample has the same acceleration data structure.

Finally, we count the time of earthquake occurrence in all earthquake recordings and divide the training set and the testing set by the earthquake time. The proportion of the number of the training set is 80% and the testing set is 20%, so we sort the statistical earthquake time from first to last and take four-fifths of the time as the node to divide the training set and the testing set. The training set is the earthquake that occurred before the time node, and the testing set is the earthquake that occurred after the time node.

2.5. Hyperparameter Adjustment. During model training, the adjustment of hyperparameters will have an impact on training efficiency and training results. This article involves two hyperparameters: learning rate and batch size.

- (1) *Learning Rate.* The learning rate is the degree to which the model reduces the value of the loss function each time. If the learning rate is high, the model may converge faster in the first few times, but it is likely that the model cannot reach the global optimum. If the learning rate is low, it will definitely have a great impact on training efficiency. We use the controlled variable method to measure the loss of the model while ensuring that other variables are the same. As shown in Figure 6, the learning rate is set to 0.001 in this article. The learning rate is multiplied by 0.99 in each iteration. The learning rate gradually decreases as the number of training increases, so that the loss of the model is closer to the minimum.
- (2) *Batch Size.* The batch size refers to the number of samples input to the model each time. If it is too small, it will easily bias the convergence direction of the model. If it is too large, it will easily cause the model to stay in the local optimum and fail to reach

the global optimum. According to the changes in the accuracy of the training set and the testing set in Figure 7, the model's relatively high generalization ability and the lowest over-fitting phenomenon are better, so the batch size of 50 is more appropriate.

3. Results

In this paper, a method for predicting the variation trend of ground motion Arias intensity using FCNN is proposed. After scaling, baseline correction, and normalization of the original seismic recording data, segment frequency division sampling is carried out, and FCNN is used for identification and classification. After 100 rounds of training, we selected the most accurate model parameters for analysis based on the test data set. The results show that the training is carried out by using the data set of 160000 samples, the classification is carried out according to the variation trend of Arias intensity, and the training effect of the model is tested on the test set of 40000 samples. After training, the model has the over-fitting phenomenon; that is, the accuracy of the training set increases and the accuracy of the test set decreases, as shown in Figure 8(a). Therefore, the model in this paper is taken from the highest accuracy before the continuous decline of the test data set. The average accuracy of the model is 76.5%. The model corresponds to different magnitudes of 5.0~5.9, 6.0~6.9, 7.0~7.9, and above 8.0. The average accuracy of prediction is 83.2%, 76.9%, 75.5%, and 70.6%, respectively, as shown in Figure 8(b).

Regarding the computational complexity, it represents the increasing trend of the runtime of the program with the size of the data. Usually we use big O notation to represent the computational complexity of the model. For the fully connected neural network model that has been trained and can be used in practice since the number of neural network layers and the number of neurons in each layer have been fixed, there will be no loop calculations, so its computational complexity depends entirely on the number of input

FIGURE 7: Batch size adjustment graph. (a) A graph showing the change in accuracy of the training set as the increasing number of training under different batch sizes; (b) A graph showing the change in accuracy of the testing set as the increasing number of training under different batch sizes.

FIGURE 8: The training results: (a) the corresponding relationship between the number of epochs and the accuracy rate during the training process; (b) the model accuracy rate of earthquakes with different magnitudes; (c) part of the results. The magnitudes from left to right are 5.4, 6.1, 7.3, and 9.0, respectively. The ratio of the Arias intensity in the last 5 seconds to the previous 5 seconds is 5.2, 0.61, 0.67, and 2.17. The model predicts that the probability that the intensity of the magnitude 5.4 earthquake will increase is 99.6%, that the probability of decreasing intensity of the magnitude 6.1 earthquake is 65.6%, that the probability of decreasing intensity of the magnitude 7.3 earthquake is 58.2%, and that the probability of increasing intensity of a magnitude 9.0 earthquake is 74.0%.

FIGURE 9: Model application schematic. (a) The recorded time series of acceleration with zero padding at the beginning. The red dots represent the calculated Arias intensity ratio α for each second in the next 5 seconds and the previous 5 seconds. (b) The probability that the previous 30 seconds of recordings are sampled at each red dot on the left and input the model to determine whether the Arias intensity will increase or decrease in the future. The red line represents whether the intensity will actually increase from the known recordings, that is, the label. The label with $\alpha > 1$ is 1, and the label with $\alpha < 1$ is 0.

samples. Using big O notation, the complexity is $O(n)$. Thus, the trained model is highly effective for real-time use. Our running time is less than 0.1 second for each evaluation on a single time series using a common computer.

4. Conclusion and Discussion

Through the above research and results, we can get the following preliminary conclusions:

- (1) The model predicts the future Arias intensity trend with a 76.5% overall accuracy. When an earthquake occurs, it can provide a certain reference for judging the development trend of this earthquake or the scale of damage caused by this earthquake. This model has certain feasibility for predicting the Arias intensity trend of the seismic acceleration recordings, and the sample data input to the model are only the acceleration recordings without other variables, and then a certain accuracy of the ground motion intensity trend prediction can be quickly obtained, which is meaningful. In practical applications, real-time monitoring recordings can be used as input to predict the trend of future monitoring values (Figure 9), and based on the predicted values, the damage intensity of the earthquake could be estimated more accurately in advance combining with other parameters such as propagation path and site conditions.
- (2) It can be seen from the histogram in Figure 8(b) that the accuracy of this model for predicting earthquake intensity trends from magnitude 5 to magnitude 8 earthquakes basically decreases as the magnitude increases.
- (3) In theory, any time series of images, including seismic acceleration recordings, can be represented as functions because the function may be too complicated to be implemented easily. In deep learning, as long as the number of network layers and neurons are sufficient and the training function meets the requirements, after a long period of training and debugging, the network is possible to approximate any complex function with a fairly high accuracy. Therefore, with the in-depth research on the application of deep learning, it may become a powerful tool for the analysis of ground motions.

Earthquake prediction is a world-recognized problem in seismology, and it is currently in the development and exploration stage. Due to the complexity of earthquake prediction, the current earthquake prediction is mainly based on long-term earthquake risk estimation and short-term prediction of observed earthquake precursor phenomena, almost all of which are empirical. Now, experts and scholars in the field of earthquake prediction at home and abroad have a certain understanding of medium- and long-term prediction, so the risk estimation has reference value. However, the understanding of earthquake precursor phenomena is far from reaching the level of regularity, so the success rate is not very high. Due to the rapid development of deep learning technology and the ability

to identify and judge data features that scientists cannot find at present, we have proposed a deep learning model for predicting earthquake trends and achieved a high success rate. Nevertheless, this model still has some problems to be solved urgently:

- (1) The data used in this article come from K-net and KiK-net. The earthquake recordings used are large, medium, and small earthquakes in Japan. Therefore, it is still necessary to verify whether the model is applicable to other regions. In the future, we will also add data from other countries and regions to the training data to improve the generalization ability of the model;
- (2) From Figure 8(b), it can be seen that the accuracy of the model for predicting the future 5 s Arias intensity trend of large earthquakes is not high, but it is more accurate on large earthquakes that can cause more damage. The anticipation of large earthquakes is crucial. The acceleration recordings of large earthquakes are more complex and have more influencing factors, so we will optimize the model for this point and strive to improve the accuracy of prediction of large earthquakes in the future.
- (3) The model predicts the trend of ground motion intensity in the future 5 seconds. With the optimization of the model and the increase in data in the future, it may be able to predict the trend of 10 seconds or more.

Data Availability

The earthquake recordings data used to support the findings of this study are processed from K-net and KiK-net.

Disclosure

All statements, results, and conclusions are those of the researchers and do not necessarily reflect the view of funders.

Conflicts of Interest

The authors declare that they have no conflicts of interest.

Acknowledgments

The authors are grateful for the research funding provided by the National Key Research and Development Program of China (2018YFE0109800 and 2019YFE0112500), the National Natural Science Foundation of China (51738001, U1839202, and 61472373), and the Special Fund of the Institute of Geophysics, China Earthquake Administration (DQJB21X24).

References

- [1] G. A. Sobolev, "Methodology, results, and problems of forecasting earthquakes," *Herald of the Russian Academy of Sciences*, vol. 85, no. 2, pp. 107–111, 2015.

- [2] H. Joffe, T. Rossetto, C. Bradley, and C. O'Connor, "Stigma in science: the case of earthquake prediction," *Disasters*, vol. 42, no. 1, pp. 81–100, 2018.
- [3] R. J. Budnitz, G. Apostolakis, D. M. Boore, and L. S. Cluf, *Recommendations for Probabilistic Seismic hazard Analysis: Guidance on Uncertainty and Use of Experts*, Report Prepared for the U. S. Nuclear Regulatory Commission, the U. S. Department of Energy, and the Electric Power Research Institute by the Senior Seismic Hazard Analysis Committee (SSHAC), Report No. NUREG/CR26372, Washington, DC, USA, 1997.
- [4] J. Woessner, D. Laurentiu, D. Giardini et al., "The 2013 european seismic hazard model: key components and results," *Bulletin of Earthquake Engineering*, vol. 13, no. 12, pp. 3553–3596, 2015.
- [5] M. D. Petersen, C. S. Mueller, M. P. Moschetti et al., "2018 one-year seismic hazard forecast for the central and eastern United States from induced and natural earthquakes," *Seismological Research Letters*, vol. 89, no. 3, pp. 1049–1061, 2018.
- [6] M. I. Jordan and T. M. Mitchell, "Machine learning: trends, perspectives, and prospects," *Science*, vol. 349, no. 6245, pp. 255–260, 2015.
- [7] P. S. Dysart and J. J. Pulli, "Regional seismic event classification at the NORESS array: seismological measurements and the use of trained neural networks," *Bulletin of the Seismological Society of America*, vol. 80, no. 6, pp. 1910–1933, 1990.
- [8] A. Ursino, H. Langer, L. Scarfi, G. Grazia, and S. Gresta, "Discrimination of quarry blasts from tectonic microearthquakes in the Hyblean plateau (southeastern Sicily)," *Annals of Geophysics*, vol. 44, no. 4, pp. 703–722, 2001.
- [9] T. Tiira, "Detecting teleseismic events using artificial neural networks," *Computers & Geosciences*, vol. 25, no. 8, pp. 929–938, 1999.
- [10] J. Wiszniowski, B. M. Plesiewicz, and J. Trojanowski, "Application of real time recurrent neural network for detection of small natural earthquakes in Poland," *Acta Geophysica*, vol. 62, no. 3, pp. 469–485, 2014.
- [11] A. Cuéllar, G. Suárez, and J. M. Espinosa-Aranda, "A fast earthquake early warning algorithm based on the first 3 s of the P-wave coda," *Bulletin of the Seismological Society of America*, vol. 108, no. 4, pp. 2068–2079, 2018.
- [12] L. H. Ochoa, L. F. Niño, and C. A. Vargas, "Fast magnitude determination using a single seismological station record implementing machine learning techniques," *Geodesy and Geodynamics*, vol. 9, no. 1, pp. 34–41, 2018.
- [13] K. M. Asim, A. Idris, and T. Martínez-Álvarez, "Earthquake prediction model using support vector regressor and hybrid neural networks," *Plos One*, vol. 13, no. 7, Article ID e0199004, 2018.
- [14] K. Vardaan, T. Bhandarkar, and N. Satish, "Earthquake trend prediction using long short-term memory RNN," *International Journal of Electrical and Computer Engineering*, vol. 9, no. 2, pp. 1304–1312, 2019.
- [15] Q. Wang, Y. Guo, L. Yu, and P. Li, "Earthquake prediction based on spatio-temporal data mining: an LSTM network approach," *IEEE Transactions on Emerging Topics in Computing*, vol. 8, no. 1, pp. 148–158, 2017.
- [16] J. Reyes, A. Morales-Esteban, and F. Martínez-Álvarez, "Neural networks to predict earthquakes in Chile," *Applied Soft Computing*, vol. 13, no. 2, pp. 1314–1328, 2013.
- [17] X. J. Li, J. Wu, and M. T. Gao, "Probabilistic seismic hazard analysis based on Arias intensity: a case study in Danling County, Sichuan Province and its surrounding area," *China Earthquake Engineering Journal*, vol. 40, no. 3, pp. 555–561, 2018.
- [18] X. Y. Wang, G. Z. Nie, and L. Zhang, "Relationship between landslides induced by the Wenchuan earthquake and Arias intensity," *Journal of Basic Science and Engineering*, vol. 18, no. 4, pp. 645–656, 2010.
- [19] R. W. Jibson, "Predicting earthquake induced landslide displacements using Newmark's sliding block analysis," *Transportation Research Record*, vol. 1411, pp. 9–17, 1993.
- [20] L. Cabañas, B. Benito, and M. Herráiz, "An approach to the measurement of the potential structural damage of earthquake ground motions," *Earthquake Engineering & Structural Dynamics*, vol. 26, no. 1, pp. 79–92, 1997.
- [21] R. Foulser-Piggott and P. J. Stafford, "A predictive model for Arias intensity at multiple sites and consideration of spatial correlations," *Earthquake Engineering & Structural Dynamics*, vol. 41, no. 3, pp. 431–451, 2012.
- [22] T. Travarasou, J. D. Bray, and N. A. Abrahamson, "Empirical attenuation relationship for Arias intensity," *Earthquake Engineering & Structural Dynamics*, vol. 32, no. 7, pp. 1133–1155, 2003.
- [23] R. W. Jibson, "Regression models for estimating coseismic landslide displacement," *Engineering Geology*, vol. 91, no. 2–4, pp. 209–218, 2007.
- [24] S. L. Kramer and R. A. Mitchell, "Ground motion intensity measures for liquefaction hazard evaluation," *Earthquake Spectra*, vol. 22, no. 2, pp. 413–438, 2006.
- [25] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," *Communications of the ACM*, vol. 60, no. 6, pp. 84–90, 2012.