

Papers that [reference](#) this paper.

Network compensation for missing sensors

Albert J. Ahumada, Jr. and Jeffrey B. Mulligan

NASA Ames Research Center, Human Interface Research Branch
Moffett Field, California 94035

In B. E. Rogowitz, M. H. Brill, and J. P. Allebach, eds., Human Vision, Visual Processing, & Digital Display II, Proceedings Volume 1453, SPIE, Bellingham, WA, pp. 134-146, 1991.

ABSTRACT

A network learning algorithm is presented that computes interpolation functions that can compensate for weakened, jittered, or missing elements of a sensor array. The algorithm corrects errors in translation invariance, so prior knowledge of the input images is not required.

1. INTRODUCTION

Self-adjusting networks that compensate for sensor degradation are potentially useful for autonomous vision systems and can model the self-calibration of biological vision systems. An algorithm that learns network weights to compensate for irregular sensor positioning has been previously described [1-3]. The algorithm has the error-correcting form of the delta rule or outer product rule developed by Widrow and Hoff to construct a linear transformation from a sequence of inputs and corresponding desired outputs [4]. Since correct feedback is assumed to be unavailable, feedback must come from another source, in this case, a translated view of the same scene. The rule enforces invariance of an internally reconstructed image under rigid translations of the external image. Translation invariance does not, however, uniquely specify the network weights of the linear translation. Uniqueness can be attained by fixing the transforming weights leading to one of the output units. If the input images are adequately low-pass filtered, so that the remaining sensors sample enough information to reconstruct them, the translation-invariance (TI) algorithm with one fixed receptive field can perfectly compensate for missing sensors as well as sensor position jitter and gain errors. When the images are not adequately sampled by the remaining sensors, this algorithm does not find a suitable solution. However, if we adjust only those weights leading to output units that are affected by a missing sensor, the TI algorithm can find weights similar to those found by the delta rule. In conjunction with a missing-unit-detection rule, this provides a self-correcting algorithm that can handle the loss of a moderate number of sensors.

1.1 A simplified vision system model

We will simulate a simplified visual system model with the sensors in a regular rectangular array. The sensors are connected to a corresponding array of internal units by a linear weighting network. Initially, as illustrated on the left side of Figure 1 ([GIF](#)), ([PostScript](#)), The dark parts of the squares in the figure represent the remaining the weights to corresponding units are unity and the others zero, and this identity transformation network copies the sensor outputs to internal units. An internal image is interpolated from the internal units using a discrete equivalent of the sinc interpolation function, an impulse rectangularly filtered in the spatial frequency domain. If an input image has no spatial frequency components above the Nyquist frequency of the array, the image will be correctly reconstructed. The input function illustrated is a pulse, centered over the central unit and low-pass filtered at the Nyquist frequency of the sampling array. It has zero contrast at the other sensors, so the interpolated image is the interpolation function from the central unit alone. The correctness of the interpolation illustrates that the interpolation functions are filtered pulses.

1.2 A compensating network

We simulate the loss of a sensor by removing it and its connections from the network. Connections from the remaining sensors to the unit corresponding to the missing sensor allow them to fill in an estimate of what the missing sensor would have seen, as illustrated on the right side of Figure 1 ([GIF](#)), ([PostScript](#)). This network cannot reproduce the previous input, which does not now activate any sensors, but it can reproduce images band-limited to the Nyquist frequency of a regular sampling array with the same number of sensors. The image on the right is the same pulse without its highest spatial frequency component. It now correctly reconstructed by the interpolation process from the the internal units to the internal image, which is not altered. The calculation of appropriate network transformation weights can be accomplished by iterative learning algorithms. They will be presented first for the case of partially degraded sensors.

2. PARTIALLY DEGRADED SENSORS

Imagine an array of sensor elements in a remote environment that have been functionally degraded as illustrated in Figure 2.1 ([GIF](#)), ([PostScript](#)). The dark parts of the squares in the figure represent the remaining parts of the absorbing areas of the elements. This degradation was done by randomly removing 25% of the pixel area. To a first approximation, the effects of the degradation to the i th sensor can be represented by a gain factor $g(i)$ and an offset of the sensor by the position errors, $\Delta x(i)$, $\Delta y(i)$. Figure 2.2 shows the effect on an image of a similar degradation. Each sensor was given a gain factor at random, uniformly distributed between 0.5 and 1, and position errors, independently and uniformly distributed with a range of one third the pixel width, which gives a standard deviation in the x or y direction close to that of the centroid of the degraded pixels of Figure 2.2 ([GIF](#)), ([PostScript](#)).

We will let $s'(i)$ be the value of the input image at the i th original sample point (sensor position). Let $s(i)$ be the value of the image at the jittered sensor position multiplied by $g(i)$. Let $r(j)$ be the value of the internal representation at the j th position. The internal values are a weighted linear combination of the $s(i)$,

$$r(j) = \sum_i \{s(i) w(i,j)\}. \quad (\text{Equation 2.1})$$

2.1 The delta rule

The delta rule can find appropriate weights when the desired values for the $r(j)$, the $s'(j)$, are known. This rule is diagramed in Figure 2.3 ([GIF](#)), ([PostScript](#)). This rule is also called the Widrow-Hoff rule, least mean square learning, outer product learning, and simple back-propagation [4-6]. The delta rule is error correcting. For this rule, a sequence of images are presented to the sensors. It is assumed that the correct values that the undegraded sensors would have measured are known. The error $e(j)$ for output unit j is defined to be the difference between the output of the unit $r(j)$ and the correct image value at the j th position $s'(j)$,

$$e(j) = r(j) - s'(j). \quad (\text{Equation 2.2})$$

The weight from sensor i to unit j , $w(i,j)$ is adjusted by subtracting a fraction λ of the error multiplied by the output of sensor i

$$\Delta w(i,j) = -\lambda s(i) e(j). \quad (\text{Equation 2.3})$$

2.2 The translation-invariance (TI) rule

The TI rule was developed for the case that the input image is not known. It thus allows learning to take place without a "teacher". It uses the outer product rule of Equation 2.3 but the error signal is computed differently. Figure 2.4 ([GIF](#)), ([PostScript](#)) is a diagram illustrating the error calculation for the TI rule. It is assumed that the current image was just presented with the eye or camera in a different position, and that the visual system can, effectively, internally translate its previous view of that image and sample it at the positions of the internal units. These samples, $r'(j)$, replace the $s'(j)$ in the Equation 2.2, and the errors $e(j)$ are now given by.

$$e(j) = r(j) - r'(j) . \text{ (Equation 2.4)}$$

The delta rule had no need of the interpolated internal image. The TI rule needs it for the case that the translations are not integral multiples of the the internal array step sizes. We have previously provided detailed formulas [2].

The TI criterion is blind to translation-invariant filtering of the image. In the spatial frequency domain, any gain and any phase is acceptable for every spatial frequency component. In the space domain, any impulse response (projective field) or any receptive field is acceptable as long as they are all the same. We have usually solved this problem by forcing one output position to correspond to a particular input, by keeping the weights zero from all other inputs and keeping the corresponding weight equal to one. The receptive field of this unit is then slowly copied by the learning algorithm to the receptive field of the other units. In the case of an array degraded as above, the unit with the largest average output could be selected for this position of honor.

2.3 Simulations of degradation compensation

Simulations were run of the delta and TI rules learning to compensate for the gain and jitter degradation of Figure 2.2. The images used in the simulation were noises synthesized from Fourier series components. The amplitudes of the components were made inversely proportional to their spatial frequency and their phases were chosen at random. The component spatial frequencies went up to 5 cycles per image in each dimension. The original sampling array was a rectangular 11 by 11 array. The jittered positions were quantized to one fourth the original sample spacing. Images were selected for a trial by translating the position of the noise image to a new position selected at random. These translations were also quantized to one fourth the sample spacing and wrapped around in both dimensions to eliminate border effects. After each block of 100 trials a new noise image was computed. The learning rate λ was set to one half (normalized by the sum over i of $s(i)^2$). The weight matrix was initialized to be the identity matrix. The upper left sensor was not degraded and its internal unit had a fixed receptive field. Figure 2.5 ([GIF](#)), ([PostScript](#)) shows the performance of the delta and TI rules. Both curves show a rapid initial drop in errors as the gains are corrected, followed by a much slower learning phase caused by the position jitter. The TI rule learning is slower in both phases.

3. MISSING SENSORS

3.1 Adequate sampling

What happens when sensors are lost completely? If the remaining sensors still sample enough information to reconstruct the images, (the number of remaining sensors is as large as the dimensionality of the images), then the above algorithms find compensating weights. These weights depend only on the spatial frequency range of the input images: any images spanning the same range would result in the same weights.

Figure 3.1 ([GIF](#)), ([PostScript](#)) illustrates delta and TI rules learning weights to compensate for one sensor missing from a 7 by 7 array. The noise images had one dimensional spatial frequencies up to 2 cycles per image, while the original Nyquist limit was 3 cycles per image. Other simulation details remained the same. Again, the learning is slower for the TI rule. For the delta rule, errors only occur at the output whose input is missing, so only the weights in the receptive field of that unit are altered. The TI rule, however, will also see errors at positions near the point to which a missing input has been translated. Initially it changes all the correct weights and then restores them as it finds the weights which fill in for the missing input.

3.2 Inadequate sampling

When the number of sensors is less than the number of degrees of freedom in the images, there are no weights that can perfectly fill in for a missing sensor. The delta rule with an appropriately decreasing learning rate converges to a minimum expected squared error solution [4,7]. Unfortunately, the TI algorithm does not

behave well in this situation. Figure 3.2 ([GIF](#)), ([PostScript](#)) illustrates that with only one sensor missing, the delta rule learns helpful weights, but the TI algorithm, after a good initial start, degrades performance as measured by RMS error. Initially, improved weights are found to fill in for the missing sensor, but then erosion of the originally correct weights dominates. The parameters were the same as the previous example except that the noise bandwidth extends to the Nyquist frequency of the original 7 by 7 sampling array and a small learning rate coefficient (0.01) was used so that the learning itself would not be an appreciable source of error at asymptote.

3.3 Adjustment of weights to affected outputs

One possible solution to this problem is to assume that when sensor losses occur, the system only adjusts the weights to output units affected by the losses. A drop in the average activity of an internal unit might activate the adjustment process. In biological systems the physiological consequences of losing a photoreceptor could trigger the learning process. The performance of the TI algorithm when adjustment is restricted to the weights that lead to affected outputs (TIA) is impressive in the above cases of one missing sensor. Figure 3.1 shows that the TIA rule can learn almost as fast as the delta rule when the sampling is adequate. Figure 3.2 shows that this rule can do as well as the delta rule when the sampling is inadequate.

A final example illustrates that the above conclusions hold when more than one sensor is missing. The conditions were the same as the previous case of inadequate sampling except that the size of the sampling array is increased to 11 by 11 and 30 per cent of the units were deleted at random. Figure 3.3 ([GIF](#)), ([PostScript](#)) shows that in this situation the TI rule performs worse than it does when only one sensor is missing. The larger number of missing sensors cause a faster erosion of the correct coefficients. The performance of the TIA rule is also worsened, but by a relatively small amount. The learning rates for these simulations were chosen for convenience and not intended to optimize performance. Performance closer to the range of the adequately sampled case might be achieved by using schedules with decreasing learning rates.

4. APPLICATIONS

4.1 Biological calibration

The TI theory was developed to provide a computational model to explain how an organism with a retina of irregularly arrayed sensors could effectively learn the positions of the sensors by construction appropriate interpolation functions. This application has been stressed in earlier papers [1-3,8]. A more general treatment of human visual calibration problems is given by Banks [9], who also presents the idea that "zoom" invariance could be used to calibrate the relative amplitudes of different spatial frequencies.

The question of whether TI learning could fill in for missing sensors was raised by Turano, who has been studying visual abilities of persons suffering from the degenerative disease retinitis pigmentosa [10]. During the course of the disease, patients experience progressive visual field loss accompanied by a diffuse loss of photoreceptors in remaining regions. Turano found surprisingly little correlation between performance on a bisection judgment task and extent of degeneration, and proposed that recalibration might have been partially responsible for the lack of correlation [11]. One interesting possibility suggested by the theory is that in cases of extreme degeneration, optical blur might actually aid in recalibration.

4.2 Remote sensing

Two types of remote sensing applications suggest themselves: robotic applications in which the recalibration would be performed as necessary to maintain vision capabilities, and post-processing applications for which the learning algorithm provides a simple, if not maximally efficient, search procedure for sensor calibration when overlapping images are available.

4.3 High resolution sensors

Rough [12] has developed a calibration method for x-ray lithography which is similar to this procedure in that the input image is not assumed to be perfectly known or calibrated. The input image is translated and it is assumed that corresponding points in the outputs can be found. It differs in that the distortions are assumed to be represented through a small number of global parameters, allowing more standard search procedures to be used. The present procedure would be useful for calibrating systems that have such small sensor elements that the positions and gains of the individual elements are difficult to precisely control.

5. DISCUSSION

5.1 Unknown translations

One feature of the TI models presented above is that the system is assumed to know the exact amount of the translation between the two views of the image. Last year we showed that when the position error is limited to the inter-pixel spacing, the amount of the translation can be estimated from the output images without appreciable effect on the learning [2]. We speculate that the same situation will hold when a moderate fraction of sensors are missing.

5.2 Two arrays or three views

The two views used by the TI method might come from separate arrays rather than two views from the same array at separate times. We have not tried to implement this version of the algorithm. However, in order to obtain adequate calibration, the views must be translated in more than one direction. Craik [13] deliberately damaged one retina by staring at the sun. He later observed distortions through the undamaged eye tending to compensate for distortions that remained in the damaged eye, consistent with a model in which inputs from both eyes are mutually calibrated.

If there were three arrays or three sequential views, majority rule could prevent the miscorrection which plagued the TI rule, and make unnecessary the restriction of learning to the region affected by the missing sensors.

5.3 Noise

We have ignored here the fact that sensors are generally noisy, as are input images. Sensors are not likely to be damaged in such a way as to decrease their gain and not decrease their signal to noise ratio, so that increasing gain might not be desirable. It seems unlikely that any biological destructive process causing diffuse degeneration of photoreceptors would be all or none. Cells would be expected to suffer partial damage, which could be modeled as some combination of position distortion, gain change, and noise level increase. Adding noise to the sensor model would also eliminate the artificial difference between a sensor element that is lost and one that has a very large drop in gain. We hypothesize that the effects of sensor noise would be similar to the effects of the aliasing noise that results from undersampling and that it would have the same adverse effects on the unmodified TI model.

6. CONCLUSION

When there are enough remaining sensors to adequately sample the input images, the TI algorithm described previously can construct a linear transformation compensating for gain changes, sensor position jitter, and sensor loss. However, when the images are undersampled and complete compensation is not possible, the algorithm needs to be modified. For moderate sensor losses, the algorithm works if the transformation weight adjustment is restricted to the weights to output units affected by the loss.

7. ACKNOWLEDGMENTS

This work was supported by NASA RTOP 506-71-51. We acknowledge the helpful collaborations of L. Maloney, M. Pavel, and K. Turano.

8. REFERENCES

1. L. T. Maloney and A. J. Ahumada, "Learning by Assertion: Two Methods for Calibrating a Linear Visual System," *Neural Computation*, vol. 1, pp. 392-401, 1989.
2. A. J. Ahumada, Jr. and J. B. Mulligan, "Learning receptor positions from imperfectly known motions." *Proc. SPIE Conf. on Human Vision, Visual Processing, & Digital Display*, vol. 1249, pp. 124-134, 1990.
3. A. J. Ahumada, Jr., "Learning receptor positions," In M. Landy and J. A. Movshon, eds., *Computational Models of Visual Processing*, MIT Press, Cambridge MA, in press, 1990.
4. B. Widrow and M. E. Hoff, "Adaptive switching circuits," *Institute of Radio Engineers, Western Electronic Show and Convention Record, Part 4*, pp. 96-104, 1960.
5. G. O. Stone, "An analysis of the delta rule and the learning of statistical associations," in D. E. Rumelhart and J. L. McClelland, eds., *Parallel Distributed Processing*, vol. I, MIT Press, Cambridge, Mass., chapt. 11, pp. 444-459, 1986.
6. T. Kohonen, *Self-Organization and Associative Memory*, Third Edition, Springer, New York, 1989.
7. R. R. Bitmead, "Persistence of excitation conditions and the convergence of adaptive schemes," *IEEE Transactions on Information Theory*, vol. 30, pp. 183-191, 1984.
8. [A. J. Ahumada, Jr. and J. I. Yellott, Jr.](#), "Reconstructing irregularly sampled images by neural networks." *Proceedings of the SPIE Conference on Human Vision, Visual Processing, and Digital Display*, vol. 1077, pp. 228-235, 1989.
9. M. S. Banks, "Visual recalibration and the development of contrast and optical flow perception," in A. Yonas (Ed.), *Perceptual Development in Infancy; The Minnesota Symposia on Child Psychology*, vol. 20, Erlbaum, Hillsdale, New Jersey, 1988.
10. [A. J. Ahumada, Jr. and K. Turano](#), "Calibration of a visual system with receptor drop-out." In M. Landy, L. T. Maloney, M. Pavel, eds., *Exploratory Vision: The Active Eye*, Springer-Verlag, New York, pp. 157-168, 1995.
11. K. Turano, "Bisection judgments in patients with Retinitis pigmentosa," *Clinical Vision Science*, vol. 6, pp. 119-130, 1991.
12. M. R. Raugh, "Absolute two-dimensional sub-micron metrology for electron beam lithography: A calibration theory with applications." *Precision Engineering*, pp. 3-13, 1985.
13. K. J. W. Craik, *The nature of psychology*, Cambridge University Press, Cambridge, chapt. 9-10, 1966.